

Administration Manual for UrBackup Server 2.1.x

Martin Raiber

January 17, 2017

Contents

1	Introduction	4
2	Installation	4
2.1	Server installation	4
2.1.1	Server installation on Windows	4
2.1.2	Server installation on Ubuntu	5
2.1.3	Server installation on Debian	5
2.1.4	Server installation on other GNU/Linux distributions or FreeBSD	5
2.1.5	GNU/Linux server installation hints	6
2.1.6	Operating system independent server installation steps	6
2.2	Client installation	6
2.2.1	Windows/Mac OS X client installation	6
2.2.2	Automatic rollout to multiple Windows computers	7
2.2.3	Client installation on Linux	7
3	Architecture	8
3.1	Server architecture	8
3.2	Client architecture	8
4	Security	9
4.1	Server webinterface rights management	9
4.2	Make webinterface accessible via SSL	9
4.2.1	Apache configuration	9
4.2.2	Lighttpd configuration	10
4.3	Client security	10
4.4	Transfer security	11
4.5	Internet mode security	11
5	Client discovery in local area networks	11
6	Backup process	12
6.1	File backup	12
6.2	Image backup	13
6.3	Collision probabilities	13
6.3.1	File backup collision probability	13
6.3.2	Image backup collision probability	13
6.4	Pre and post backup scripts on client and server	14
6.4.1	Client pre and post backup scripts	14
6.4.2	Server post backup scripts	14

7	Internet clients	15
7.1	Automatically push server configuration to clients	15
7.2	Download a preconfigured client installer	15
7.3	Manually add and configure clients	15
7.4	File transfer over Internet	15
8	Server settings	16
8.1	Global Server Settings	16
8.1.1	Backup storage path	16
8.1.2	Server URL	16
8.1.3	Do not do image backups	16
8.1.4	Do not do file backups	16
8.1.5	Automatically shut down server	16
8.1.6	Download client from update server	16
8.1.7	Show when a new server version is available	17
8.1.8	Autoupdate clients	17
8.1.9	Max number of simultaneous backups	17
8.1.10	Max number of recently active clients	17
8.1.11	Cleanup time window	17
8.1.12	Automatically backup UrBackup database	17
8.1.13	Total max backup speed for local network	17
8.1.14	Global soft file system quota	18
8.2	Mail settings	18
8.2.1	Mail server settings	18
8.2.2	Configure reports	19
8.3	Client specific settings	19
8.3.1	Backup window	21
8.3.2	Advanced backup interval	21
8.3.3	Excluded files	22
8.3.4	Default directories to backup	22
8.3.5	Virtual sub client names	23
8.4	Internet settings	23
8.4.1	Data usage limit estimation	24
8.5	Advanced settings	24
8.5.1	Enabling temporary file buffers	24
8.5.2	Transfer modes	24
8.5.3	Incremental image backup styles	25
8.5.4	Full image backup styles	25
8.5.5	Database cache size during batch processing	25
8.6	Use symlinks during incremental file backups	25
8.7	Debugging: End-to-end verification of all file backups	25
8.8	Debugging: Verify file backups using client side hashes	26
8.9	Periodically readd file entries of internet clients to database	26
8.10	Create symbolically linked views for each user on the clients after file backups	26
8.11	Maximum number of simultaneous jobs per client	26
8.12	Volumes to snapshot in groups during image backups	26
8.13	Volumes to snapshot in groups during file backups	26
8.14	Windows components backup configuration	26
9	Restoring backups	26
9.1	Restoring image backups	27
9.2	Restoring file backups	27

10	Miscellaneous	28
10.1	Manually update UrBackup clients	28
10.2	Logging	28
10.3	Used network ports	28
10.4	Mounting (compressed) VHD files on GNU/Linux	28
10.5	Mounting VHDs as a volume on Windows	29
10.6	Decompress VHD files	29
10.7	Assemble multiple volume VHD images into one disk VHD image	29
10.8	Migrate non-btrfs backup storage	29
11	Storage	30
11.1	Nightly backup deletion	30
11.2	Emergency cleanup	30
11.3	Cleanup for servers with file backups with lots of files	30
11.4	Cleaning the storage folder of files not known by UrBackup	31
11.5	Archiving	31
11.5.1	Archival window	31
11.6	Suitable file systems	32
11.6.1	Ext4/XFS	32
11.6.2	NTFS	32
11.6.3	btrfs	32
11.6.4	ZFS	32
11.7	Storage setup proposals	33
11.7.1	ZFS	33
11.7.2	Btrfs	34

1 Introduction

UrBackup is a client/server backup system. This means there is a server which backs up clients. Accordingly UrBackup is divided into a client and server software. The client software currently runs on Windows, Linux and Mac OS X with only the Windows client being able to perform image backups. The server part of UrBackup runs on Windows, GNU/Linux and FreeBSD.

A lot of effort in UrBackup was made to make setup as easy as possible. If you are okay with the default settings (see section 8) the only thing you need to define on the server side is where backups should be stored. On the clients you only need to say which directories should be backed up. If server and clients are in the same subnet the server will automatically discover the clients and then start backing them up (for details see section 5). This also makes building a decentralized backup strategy very easy, as e.g. one backup server per subnet is responsible for backing up all clients in this subnet. If a computer is moved from one subnet to another this new client is discovered and the server in the new subnet automatically takes over backing it up. If you want to implement something like this, you should also read the section on security (see 4) for details on when a client accepts a server.

Installation instructions are in section 2. The interested administrator should also read up on the general UrBackup architecture (section 3), how the backups are stored and performed (section 6), proposals on which file systems are suited (section 11.6) and take a look at some sample a setup with the next generation file systems ZFS and btrfs at section 11.7.

Since version 1.0 UrBackup can also backup clients over the Internet without VPN (see section 7), and archive backups (see section 11.5).

2 Installation

This section will explain how to install the UrBackup server on various operating systems and how to distribute and install the UrBackup client.

2.1 Server installation

2.1.1 Server installation on Windows

- Download the NSIS (.exe) or MSI installer. You can only use the MSI installer, if you have a 64-bit operating system and at least Windows Vista/2008.
- Install the UrBackup Server.
- Go to the web interface (<http://localhost:55414>) and then go to the settings and configure the folder where UrBackup should store the backup. This folder should have following properties:
 - It should be on a NTFS formatted volume (not ReFS or FAT).
 - There should be enough free space to accommodate the backups
 - Preferably the volume should be dedicated to UrBackup backups
 - The volume should be persistently online while the UrBackup Server instance is running. UrBackup does not support different backup volumes/drives
 - While migration is possible it will be lengthy and difficult. So best plan ahead.
 - You can easily increase the size of the backup storage volume, if you use Windows dynamic volumes or a hardware raid. If you are using a plain volume change it to a dynamic volume before the first backup.
 - Turn on compression for the urbackup folder (in Explorer: Right click and properties). If you are not using a really old computer it should pay off without decreasing the backup speed. Possible exception: If you plan to backup files with more than 50GB or

turn off the image compression and plan to backup volumes with more than 50GB you should not turn on compression. NTFS cannot compress files larger than about 50GB.

- Alternative to the compression you can use the offline dedup and compression build into Windows Server 2012
- Disable 8.3 name generation on the volume. See <https://support.microsoft.com/en-us/kb/121007> on how to do this. 8.3 name generation causes errors in rare cases, lowers performance and the 8.3 names are only used in rare cases.
- If you are using Windows Server 2008(R2) (or the equivalent Desktop versions) you should consider applying hotfix <https://support.microsoft.com/en-us/kb/967351> and then formatting the backup storage volume with

```
Format <Drive:> /FS:NTFS /L
```

This prevents a potential issue with large files on those operating systems.

Continue with the operating system independent installation steps in section 2.1.6.

2.1.2 Server installation on Ubuntu

Install the UrBackup Server via running following commands:

```
sudo add-apt-repository ppa:uroni/urbackup
sudo apt-get update
sudo apt-get install urbackup-server
```

See section 2.1.5 for further installation hints and 2.1.6 for operating system independent installation steps.

2.1.3 Server installation on Debian

Follow the download link for Debian on <http://urbackup.org/download.html>. Packages are available for Debian stable, testing and unstable for the CPU architectures i686 and AMD64.

The Package can be installed via:

```
sudo apt-get update
sudo dpkg -i urbackup-server*.deb
sudo apt-get -f install
```

See section 2.1.5 for further installation hints and 2.1.6 for operating system independent installation steps.

2.1.4 Server installation on other GNU/Linux distributions or FreeBSD

Baring details on <http://urbackup.org/download.html> you need to compile the server.

- Download the urbackup server source tarball and extract it.
- Install the dependencies. Those are gcc, g++, make, libcrypto++ and libcurl (as development versions).
- Compile and install the server via *./configure*, *make* and *make install*.
- Run the server with *urbackupsrv run*.
- Add */usr/bin/urbackupsrv run -daemon* to your */etc/rc.local* to start the UrBackup server on server start-up.

See section 2.1.5 for further installation hints for GNU/Linux systems and 2.1.6 for operating system independent installation steps.

2.1.5 GNU/Linux server installation hints

Go to the webinterface (<http://localhost:55414>) and configure the backup storage path in the settings. A few hints for the backup storage:

- It should be easily extendable, which can be done by using a hardware raid, the volume manager LVM or the next generation file systems btrfs and ZFS.
- You should compress the file backups. This can be done by using ZFS (<http://zfsonlinux.org/>) or btrfs.
- Prefer btrfs, because UrBackup can put each file backup into a separate sub-volume and is able to do a cheap block based deduplication in incremental file backups. See section 11.7.2 on how to setup a btrfs backup storage. You should set a generously low soft file system quota (see section 8.1.14) if using btrfs, because btrfs currently still has issues in out-of-space situations and may require manual intervention.

2.1.6 Operating system independent server installation steps

After you have installed the UrBackup server you should perform following steps:

- Go to the user settings and add an admin account. If you do not do this everybody who can access the server will be able to see all backups!
- Setup the mail server by entering the appropriate mail server settings (See section 8.2.1).
- If you want the clients to be able to backup via Internet and not only via local network, configure the public server name or IP of the server in the Internet settings (See section 8.4).
- If you want the clients to be able to access their backups via browser and “right click -> Restore/access backups“ enter a server URL. E.g. <http://backups.company.com:55414/>. Make sure your DNS is configured such that *backups.company.com* points to the internal IP of the backup server if accessed from the internal network and points to the external IP otherwise. You should put a real web server in front of UrBackup and setup SSL (See section 4.2).
- If you want to get logs of failed backups go the “Logs” screen and configure the reports for you email address.
- Change any other setting according to your usage scenario. See section 8 for descriptions of all available settings.
- Install the clients (see section 2.2).

2.2 Client installation

2.2.1 Windows/Mac OS X client installation

If you plan on using the client in the same local network as the server, or the client is in your local network during setup time:

- Download the client from <http://www.urbackup.org>.
- Run the installer.
- Leave the backed up items at the default, manually select paths to backup or configure the client from the server (see section 8.3.4).
- The server will automatically find the client and start backups.

If the client is only reachable via Internet/through NAT:

- Add a new Internet client on the status page.
- Download the client installer for the Internet client and send it to the new client. Alternatively, create a user for the new client (in the settings) and send the user the username/password. The user can then download the client installer from the server on the status page and install it.
- Select the backup paths you want to backup on the client or configure appropriate default directories to backup on the server (see section 8.3.4).
- The server will automatically start backups once the client is connected.

This is the easiest method to add new internet clients. Other methods to add internet clients are described in section 7.

2.2.2 Automatic rollout to multiple Windows computers

First, if you want to deviate from the default backup path selection, configure the general default backup paths so that the correct folders get backed for each client (see section 8.3.4). Then install the clients using one of the following methods.

On local network:

Add the MSI client installer as group policy to the domain controller. Alternatively you can use the NSIS (.exe) installer with the switch “/S” to do a silent install and use something like “psexec”. The server will automatically find and backup the new clients.

For internet clients:

Adapt the script at <https://urbackup.atlassian.net/wiki/display/US/Download+custom+client+installer+via+Python> to your server URL and create a python executable from the modified script via cx_Freeze (<http://cx-freeze.sourceforge.net/>). Executing the python executable on the client automatically creates a new internet client on the server, downloads a custom client and runs the installer. You could also add the silent install switch (“/S”) when starting the downloaded installer such that it needs no user intervention.

2.2.3 Client installation on Linux

If you plan on using the client in the same local network as the server, or the client is in your local network during setup time:

- Download the portable binary Linux client from <http://www.urbackup.org>.
- Run the installer.
- Select one of the available snapshot mechanisms. If none is available consider installing your Linux on LVM or btrfs. Otherwise you will have to stop all applications during backups which are modifying files via pre/post-backup scripts.
- The server will automatically find the client and start backups.

If the client is only reachable via Internet/through NAT:

- Add a new Internet client on the status page.

- Download the client installer for the Internet client and send it to the new client. Alternatively, create a user for the new client (in the settings) and send the user the username/password. The user can then download the client installer from the server on the status page and install it.
- Select the backup paths you want to backup on the client via command line (“urbackup-clientctl add-backupdir -path /” or configure appropriate default directories to backup on the server (see section 8.3.4).
- The server will automatically start backups once the client is connected.

3 Architecture

UrBackup is divided into a server and a client software part. The server is responsible for discovering clients, backing them up, deleting backups if the storage is depleted or too many backups are present, generating statistics and managing client settings. The client listens to server commands which tell it e.g. that a file list should be build or which file the server wants to download. The server also starts a channel on which the clients can request the server to start a backup or to update the client specific settings.

3.1 Server architecture

The server is organized into a core part and an interface. Currently only a web interface is available. The web interface is accessible via FastCGI (on port 55413) and HTTP (on port 55414). You can use the FastCGI port to make the webinterface accessible via SSL (using e.g. apache web server). For details on that see section 4.2. The server core part consists of several threads with different tasks. One thread discovers new clients, another checks if a client needs to be backed up, while others send pings to clients to see if they are still alive or send them the current backup status. One updates file statistics or deletes old backups. Because there are so many threads UrBackup server profits from modern multi core CPUs (the more cores the better!).

3.2 Client architecture

The client is divided into a core process and an interface process. The interface process displays the tray icon and the dialogues and sends settings and commands to the core client process. The core client process listens on port 35622 UDP for UDP broadcast messages from the server and on receiving one sends a message with its name back to the server. As name the Windows computer name is used. It listens on port 35623 TCP for commands from the client interface process and the server and on port 35621 TCP for file requests from the server. The server establishes a permanent connection to each client on its command port with which the clients can request backups or change their settings. The core client process is responsible for building a list of all files in the directories to be backed up. This list is created in the UrBackup client directory as 'urbackup/ data/ filelist.ub'. To speed up the directory list creation directories to be backed up are constantly watched via the Windows Change Journal. The Windows Change Journal can only be used for whole partitions. Thus the first time a directory on a volume is added the UrBackup core client process reads all the directory entries on the new volume into the client database file in 'urbackup/backup_client.db'. After a volume is successfully indexed the database is constantly updated to be in sync with the file system. Thus if large changes in the volume occur the database gets updated more often. This does not have a big performance penalty as only directories are saved in the database. The updating is done every 10 seconds or if a file list is requested. The server downloads the file list from the client and starts the backup by downloading changed or new files from the build in client file server. The image backup is done using only the command port.

4 Security

4.1 Server webinterface rights management

The server web interface is protected by a pretty standard user system. You can create, manage and delete accounts. Those accounts are only linked loosely to clients by rights management. Be aware that after first installing UrBackup there is no administrator password set and everybody can see all backed up files! If you want to limit access you should immediately go to the account management in the settings and create an administrator account and set its password.

An admin account can do everything including browsing file backups of all clients. The web interface allows one to create a 'limited' account that can only browse backups and view statistics from one client. The more sophisticated rights editor can be used to allow an account to access several clients or to limit some aspects. For example you could setup an account which can do everything except browse backups. Following domains, with which you can limit or expand an account's rights, are currently available:

Domain	Description
browse_backups	Browse and download files from file backups
lastacts	View the last actions (file or image backups) the server did (including backup size and duration)
progress	View the progress of currently running file or image backups
settings	Allows settings to be changed
client_settings	Allows client specific settings to be changed
status	Allows the current status to be viewed (last seen, last file backup and last image backup)
logs	View the logs which were creating during backups
manual_archive	Manually archive file backups
stop_backup	Stop backups for client on the server
piegraph*	View statistics
users*	Get client names
general_settings*	Change general settings (like backup storage path)
mail_settings	Change the mail server settings
usermod*	Create, change and delete users
remove_client*	Remove clients and delete all their backups
start_backup*	Start backups for a client on the server
download_image	Download images of volumes from the server via restore CD

You can set the domains not marked with stars(*) either to one or several client ids (separated by ',') or to 'all' - meaning the account can access all clients. The entries with stars(*) have to be set to 'all' or 'none' and don't allow client ids. In order to be able to view statistics you need to set both 'piegraph' and 'users' to 'all'. There is a special domain 'all' which is a wild card for all domains (this means if you set 'all' to 'all' the account has the right to do everything).

Currently a user needs the "status" right for at least one client, in order for the user to be able to log in.

4.2 Make webinterface accessible via SSL

The server web interface is accessible via FastCGI (on port 55413 TCP). With this you can connect UrBackup with pretty much every modern web server and thus make the web interface accessible via SSL. This section will describe how to do this with apache and lighttp.

4.2.1 Apache configuration

Either add a symlink to the 'www' UrBackup directory or define it as an alias. For the symlink method you need to go to your SSL webroot and then do e.g.:

```
ln -s /usr/share/urbackup/www urbackup
```

Be sure you have set 'Option +FollowSymLinks' in the web server configuration on the directory you link into. From now on it is assumed that urbackup should be accessible via `https://hostname/urbackup`. Download and install 'libapache2-mod-fastcgi' (this may have another name on other distributions and maybe in a "non-free" section). Add following line to the 'fastcgi.conf':

```
FastCgiExternalServer /var/www/urbackup/x -host 127.0.0.1:55413
```

The path depends of cause on where your web root is and where you want the web interface to be. UrBackup should now be accessible via apache.

4.2.2 Lighttpd configuration

Link the urbackup/www directory into the webroot as described in the apache configuration. Add

```
include "conf.d/fastcgi.conf"
```

to your 'lighttpd.conf' file. Then add

```
fastcgi.server = (  
    "/urbackup/x" =>  
    (( "host" => "127.0.0.1",  
        "port" => 55413  
    ))  
)
```

to the 'fastcgi.conf' file.

4.3 Client security

UrBackup Client only answers commands if the server or the interface process supply it with credentials. The server credential is saved in '/var/lib/urbackup/server_ident.key'. If it does not exist the server will randomly generate it the first time it runs. The server identity is also confirmed by private/public key authentication. If not present the server will generate a private and public ECDSA key in 'server_ident_ecdsa409k1.priv' and 'server_ident_ecdsa409k1.pub'.

The client interface credential is generated in the same way and resides in 'pw.txt' and 'pw_change.txt' in the UrBackup directory on the client. To give the client core process interface commands you need the contents of 'pw.txt' or 'pw_change.txt' depending on what the command is:

pw.txt:

- Getting the current status
- Get the paths which are backed up during file backups
- Get the incremental file backup interval
- Start backups
- Pause backups

pw_change.txt

- Change the paths which are backed up during file backups
- Get all settings
- Change all settings
- Get log entries/logs

- Accept a new server

Per default only privileged users can access 'pw_change.txt'. On Windows this leads to a elevation prompt on selecting a menu item which requires the contents of 'pw_change.txt'. If you want to allow the commands without elevation prompt, either disable UAC or change the permissions on 'pw_change.txt' to allow non-privileged users read access. The client core process saves the server credentials from which it accepts commands and which it allows to download files in 'server_idents.txt' - one credential per line. The server's public key is also saved in 'server_idents.txt'.

If you want to manually add a server to 'server_idents.txt' you need to remove the preceding '#I' and '#' at the end of the contents of 'server_ident.key'. After installation the 'server_idents.txt' does not exist and the client core process accepts (and adds) the first server it sees (with the public key of the server). After that no other servers with different credentials are accepted and you need to add their credentials either manually, or via clicking on the popup box, once the client has detected the new server. This prevents others from accessing files you want to be backed up in public places.

If you want to have several servers to be able to do backups of a client you have two options. Either you manually supply the server credentials to the client (by copying them into 'server_idents.txt') or you give all servers the same credentials by copying the same 'server_ident_key', 'server_ident_ecdsa409k1.priv' and 'server_ident_ecdsa409k1.pub' to all servers.

4.4 Transfer security

The transfer of data from client to server is unencrypted on the local network allowing eavesdropping attacks to recover contents of the data that is backed up. With this in mind you should use UrBackup only in trusted local networks. Be aware that it is really easy to redirect traffic in local networks if no action has been taken to prevent that.

4.5 Internet mode security

The Internet mode uses strong authentication and encryption. The three way handshake is done using a shared key and PBKDF2-HMAC using SHA512 with 20000 iterations. The data is encrypted and authenticated using AES-GCM. Additionally the local network server authentication via server identity key and ECDSA private/public key authentication is done.

5 Client discovery in local area networks

UrBackup clients should be discovered automatically given that server and client reside in the same sub-network. The client discovery works as follows:

The UrBackup server broadcasts a UDP message every 50 seconds on all adapters into the local subnet of this adapter. (On Linux you can configure which network adapters UrBackup should broadcast on.) On receiving such a broadcast message the client answers back with its fully qualified domain name. Thus it may take up to 50 seconds until a client is recognized as online.

If the client you want to backup is not in the same subnet as the server and broadcast packages therefore do not reach the client you can add its IP or host name manually by clicking "add new client" on the status page and then selecting "Discover new client via IP/hostname hint". The server will then additionally send an UDP message directly to that entered IP or resolved host name allowing switches to forward the message across subnet boundaries. Be aware though that all connections are from server to client. If you have NAT between server and client, you should use "Internet clients" (see section 7). Using "Internet clients" all connections are from client to server.

6 Backup process

This section will show in detail how a backup is performed.

6.1 File backup

- The server detects that the time to the last incremental backup is larger than the interval for incremental backups or the last time to the last full backup is larger than the interval for full backups. Backups can be started on client requests as well.
- The server creates a new directory where it will save the backup. The schema for this directory is YYMMDD-HHMM with YY the year in a format with two decimals. MM the current month. DD the current day. And HHMM the current hour and minute. The directory is created in the backup storage location in a directory which name equals the client name.
- The server requests a file list construction from the client. The client constructs the file list and reports back that it is done.
- The server downloads 'urbackup/data/filelist.ub' from the client. If it is an incremental backup the server compares the new 'filelist.ub' with the last one from the client and calculates the differences.
- The server starts downloading files. If the backup is incremental only new and changed files are downloaded. If the backup is a full one all files are downloaded from the client.
- The server downloads the file into a temporary file. This temporary file is either in the urbackup_tmp_files folder in the backup storage dir, or, if you enabled it in the advanced settings, in the temporary folder. On successfully downloading a file the server calculates its hash and looks if there is another file with the same hash value. If such a file exists they are assumed to be the same and a hard link to the other file is saved and the temporary file deleted. If no such file exists the file is moved to the new backup location. File path and hash value are saved into the server database.
- If the backup is incremental and a file has not changed a hard link to the file in the previous backup is created.
- If the backup is incremental, "Use symlinks during incremental file backups" is enabled and a directory with more than 10 files or folders is unchanged, it is symbolically linked to the same folder in the last backup. Because the last backup will probably be deleted before the current backup, the folder is first moved to a pool directory (".directory_pool" in the client folder) and then linked from both places. The reference count of the directory is increased/decreased every time another symbolic link is created/removed to that directory.
- If the client goes offline during the backup and the backup is incremental the server continues creating hard links to files in the previous backup but does not try to download files again. The files that could not be downloaded are then not saved into the server side file list. If the backup is a full one and the client goes offline the backup process is interrupted and the partial file list is saved, which includes all files downloaded up to this point.
- If all files were transferred the server updates the 'current' symbolic link in the client backup storage location to point to the new backup. This only happens if the client did not go offline during the backup.

6.2 Image backup

The server detects that the time to the last full image backup is larger than the interval for full image backups, the time to the last incremental backup is larger than the interval for incremental image backups or the client requested an image backup. The server then opens up a connection to the client command service requesting the image of a volume. The client answers by sending an error code or by sending the image. The image is sent sector for sector with each sector preceded by its position on the hard disk. The client only sends sectors used by the file system. If the backup is incremental the client calculates a hash of 512 kbyte chunks and compares it to the previous image backup. If the hash of the chunk has not changed it does not transfer this chunk to the server, otherwise it does. Per default the server writes the image data directly into a VHD file. If enabled in the advanced configuration the server writes the image data to a temporary file first. The temporary files have a maximum size of 1GB. After this size is exceeded the server continues with a new temporary file. The image data is written to a VHD file in parallel and is located in the client directory in the backup storage location. The VHD file's name is 'Image_<Volume>_<YYMMDD_HHMM>.vhd'. <Volume> being the drive letter of the backed up partition and YY the current year, MM the current month, DD the current day in the month and HHMM the hour and minute the image backup was started.

Since UrBackup Server 1.4 the VHD-files are compressed by default. This can be disabled in the image backup settings section. On Windows there are no tools to directly mount compressed VHD files. For mounting them on Linux see section 10.4. For decompressing the image files such that they can be mounted on Windows see 10.7. The compressed VHD files have the extension '.vhdz'. The VHD files are compressed in 2MB blocks using GZIP compression with normal compression level.

6.3 Collision probabilities

In this section we will look at the probability that the UrBackup backup system considers data the same, even though it is different. This can be caused by a hash collision (data has the same hash, even though the data is different). If happening, a collision can lead to files being incorrectly linked or blocks in image backups not transferred.

6.3.1 File backup collision probability

UrBackup uses SHA512 to hash the files before file deduplication. In comparison ZFS uses SHA256 for block deduplication. The choice of SHA512 is safer. The Wikipedia page for "Birthday attack" has a probability table for SHA512. According to it one needs $1.6 * 10^{68}$ different files (of same size) to reach a probability of 10^{-18} of a collision. It also states that 10^{-18} is the best case uncorrectable bit error rate of a typical hard disk. To have $1.6 * 10^{68}$ different files of 1KB you need $1.4551915 * 10^{56}$ EB of hard disk space. So it is ridiculously more likely that the hard disk returns bad data or the data gets corrupted in RAM, rather than UrBackup linking the wrong files to each other.

6.3.2 Image backup collision probability

For the blocks in an image backup SHA256 is used. They are 512kbyte in size. The chance of a hash collision with SHA256 is $1:(2^{256})$ ($p = \frac{1}{2^{256}}$) for two hashes. In the worst case you have $2TB/512kbyte = 4194304$ different blocks in an incremental image backup. The chance of having a collision in any of the 4194304 blocks (the worst case) is then $1 - (1 - \frac{1}{2^{256}})^{4194304} \approx 3.6 * 10^{-71}$. This is again ridiculously low compared to e.g. the probability of 10^{-18} of a typical hard disk having a uncorrectable bit error.

6.4 Pre and post backup scripts on client and server

UrBackup calls scripts previous and after backups on both the server and the client. This section will list the called scripts and the script parameters.

6.4.1 Client pre and post backup scripts

On Linux the clients pre and post backups scripts are searched for `/etc/urbackup/` or `/usr/local/etc/urbackup/` (depending on where urbackup is installed). On Windows they are searched for per default in `C:\Program Files\UrBackup` with a `.bat` file extension. All scripts except `prefilebackup.bat` on Windows have to be created first.

Script	Description	Parameters	On failure (return code not zero)
prefilebackup	Called before a file backup (before snapshot/shadowcopy creation).	1: "0" for full backup "1" for incremental file backup. 2: Server token. 3: File backup group	Indexing fails and backup is not started
postfilebackup	Called if a file backup successfully finished	No parameters	Ignored
preimagebackup	Called before a image backup (before snapshot/shadowcopy creation).	1: "0" for full backup "1" for incremental file backup. 2: Server token.	Image backup fails
postimagebackup	Called if a image backup successfully finished	No parameters	Ignored

6.4.2 Server post backup scripts

On Linux the post backup scripts are searched for in `/var/urbackup` or `/usr/local/var/urbackup` (depending on where urbackup is installed). On Windows they are searched for per default in `C:\Program Files\UrBackupServer\urbackup` with a `.bat` file extension. All scripts have to be created.

Script	Description	Parameters	On failure (return code not zero)
post_full_filebackup	Executed after a full file backup finished	1: Path to file backup. 2: "1" if successful, "0" otherwise. 3: File backup group	Backup fails
post_incr_filebackup	Executed after a incremental file backup finished	1: Path to file backup. 2: "1" if successful, "0" otherwise. 3: File backup group	Backup fails
post_full_imagebackup	Executed after a full image backup finished	1: Path to image backup file. 2: Image letter. 3: "1" if successful, "0" otherwise	Backup fails
post_incr_imagebackup	Executed after a incremental image backup finished	1: Path to image backup file. 2: Image letter. 3: "1" if successful, "0" otherwise	Backup fails

7 Internet clients

UrBackup is able to backup clients over the internet, enabling mixed LAN and Internet backups. This can be useful e.g. for mobile devices which are not used in the LAN all the time, but are connected to the Internet. As it causes additional strain on the backup file system this feature is disabled by default. You need to enable and configure it in the settings and restart your server to use it. The minimum you have to configure is the server name or IP on which the backup server will be available on the Internet. As you probably have a Firewall or Router in between backup server and Internet you also need to forward the configured port (default: 55415) to the backup server.

There are three ways to configure the clients illustrated in the three following sections.

7.1 Automatically push server configuration to clients

If the client is a mobile device it is easiest to let the server push its name and settings to the client. This will happen automatically. The server will also automatically generate a key for each client and push that one to the client as well. This assumes that the local area network is a secure channel. If a client has been compromised you can still change the key on the server and on the client.

7.2 Download a preconfigured client installer

If you enabled “Download client from update server” in the server settings, or manually downloaded a client update and put it in the appropriate folder (see section 10.1), users can download a preconfigured client installer from the server on the “Status” page. The installer includes the server IP/domain name, the port it is listening on, the clientname and the authentication key. This allows the installed client to automatically connect to the server the installer was downloaded from.

7.3 Manually add and configure clients

UrBackup also allows manually adding clients and manually configuring the shared key. To add such a client following steps are necessary:

1. Go to the “Status” screen as administrator
2. Click on “Add new client” and under “Name of new Internet client/client behind NAT” enter the name of the Laptop/PC you want to add.
3. UrBackup will show a generated authentication key for that client.
4. On the client go to the settings and enter the same key there in the internet settings and the name you gave the client in the “client” tab. Also enter the public IP or name of your backup server and the port it is reachable at.
5. The server and client should now connect to each other. If it does not work the client shows what went wrong in the “Status” window.

7.4 File transfer over Internet

If a client is connected via Internet UrBackup automatically uses a bandwidth saving file transfer mode. This mode only transfers changed blocks of files and should therefore conserve bandwidth on files which are not changed completely, such as database files, virtual hard disks etc.. This comes at a cost: UrBackup has to save hashes of every file. If the hashes of a file are not present e.g. because Internet mode was just enabled, they are created from the files during the backup and may therefore slow down the backup process.

8 Server settings

The UrBackup Server allows the administrator to change several settings. There are some global settings which only affect the server and some settings which affect the client and server. For those settings the administrator can set defaults or override the client's settings.

8.1 Global Server Settings

The global server settings affect only the server and can be changed by any user with "general_settings" rights.

8.1.1 Backup storage path

The backup storage path is where all backup data is saved. To function properly all of this directories' content must lie on the same file system (otherwise hard links cannot be created). How much space is available on this file system for UrBackup determines partly how many backups can be made and when UrBackup starts deleting old backups. Default: "".

8.1.2 Server URL

URL to which the client will browse if a user selects "Access/restore backups". For example "http://backups.company.com:55414/". Default: "" (If empty "Access/restore backups" will not be available on the clients.)

8.1.3 Do not do image backups

If checked the server will not do image backups at all. Default: Not checked.

8.1.4 Do not do file backups

If checked the server does no file backups. Default: Not checked.

8.1.5 Automatically shut down server

If you check this UrBackup will try to shut down the server if it has been idle for some time. This also causes too old backups to be deleted when UrBackup is started up instead of in a nightly job. In the Windows server version this works without additional work as the UrBackup server process runs as a SYSTEM user which can shut down the machine. On Linux the UrBackup server runs as a limited user which normally does not have the right to shut down the machine. UrBackup instead creates the file '/var/urbackup/shutdown_now', which you can check for existence in a cron script e.g.:

```
if test -e /var/urbackup/shutdown_now
then
shutdown -h +10
fi
```

Default: Not checked.

8.1.6 Download client from update server

If this is checked the server will automatically look for new UrBackup client versions. If there is a new version it will download it from the Internet. The download is protected by a digital signature. Default: Checked.

8.1.7 Show when a new server version is available

If this is checked the server will show on the status page (to admins only) if there is a new server version available. Default: Checked.

8.1.8 Autoupdate clients

If this is checked the server will send new versions automatically to its clients. The UrBackup client interface will ask the user to install the new client version. If you check silent autoupdate (see Section 8.1.8) it will update in the background. The installer is protected by a digital signature. Default: Checked.

8.1.9 Max number of simultaneous backups

This option limits the number of file and image backups the server will start simultaneously. You can decrease or increase this number to balance server load. A large number of simultaneous backups may increase the time needed for backups. The number of possible simultaneous backups is virtually unlimited. Default: 10.

8.1.10 Max number of recently active clients

This option limits the number of clients the server accepts. An active client is a client the server has seen in the last two month. If you have multiple servers in a network you can use this option to balance their load and storage usage. Default: 100.

8.1.11 Cleanup time window

UrBackup will do its clean up during this time. This is when old backups and clients are deleted. You can specify the weekday and the hour as intervals. The syntax is the same as for the backup window. Thus please see section 8.3.1 for details on how to specify such time windows. The default value is 1-7/3-4 which means that the cleanup will be started on each day (1-Monday - 7-Sunday) between 3 am and 4 am.

8.1.12 Automatically backup UrBackup database

If checked UrBackup will save a backup of its internal database in a subdirectory called 'urbackup' in the backup storage path. This backup is done daily within the clean up time window. Default: Checked. If you backup a lot of files and the database gets large consider disabling this and performing the database backup via another method, e.g. by installing the UrBackup client on the server. You should backup "/var/urbackup" or "C:

Program files
UrBackupServer
urbackup".

8.1.13 Total max backup speed for local network

You can limit the total bandwidth usage of the server in the local network with this setting. All connections between server and client are then throttled to remain under the configured speed limit. This is useful if you do not want the backup server to saturate your local network.

All speed settings can have different values for different windows. See first how to specify a window at section 8.3.1.

You can set different speeds at different times by combining the speed setting with a window, separated by "@".

If you want a default speed limit of 60 MBit/s and 10 MBit/s during working hours (Mon-Fri, 8am to 6pm):

```
60;10@Mon-Fri/8-18
```

The most specific speed limit will be used, so adding an extra rule for 80 MBit/s for 12am to 1pm works as expected regardless of order:

```
60;10@Mon-Fri/8-18;80@1-7/12-13
```

You can also specify speed limits as a percentage of the maximum speed. So e.g. 50%. This can again be combined with windows separated by “@”. If a percentage value is specified, UrBackup will run at full speed at certain intervals until the speed stabilizes to its maximum value. When it is not testing for the maximum speed it will throttle down the the percentage of maximum speed you specified.

8.1.14 Global soft file system quota

During cleanups UrBackup will look at the used space of the file system the backup folder is on. If the used space is higher than the global soft file system quota UrBackup will delete old backups if possible, till the used space is below the quota. Be aware that not only UrBackup’s files count against the quota, but other files as well. You can specify the quota via a percentage of total space, or by a size. For example let the size of the Backup device be 1 Tera-byte: If you set the global file system quota to "90%", UrBackup will delete old backups as soon as more than about 900 Giga-bytes of the available space is used. You could also directly set the quota to 900 Giga-bytes by setting it to "900G". Other units are possible, e.g. "900000M" or "1T".

8.2 Mail settings

8.2.1 Mail server settings

If you want the UrBackup server to send mail reports a mail server should be configured in the 'Mail' settings page. The specific settings and their description are:

Settings	Description	Example
Mail server name	Domain name or IP address of mail server	mail.example.com
Mail server port	Port of SMTP service. Most of the time 25 or 587	587
Mail server user-name	Username if SMTP server requires one	test@example.com
Mail server password	Password for user name if SMTP server requires credentials	password1
Sender E-Mail Address	E-Mail address UrBackup’s mail reports will come from	urbackup@example.com
Send mails only with SSL/TLS	Only send mails if a secure connection to the mail server can be established (protects password)	
Check SSL/TLS certificate	Check if the server certificate is valid and only send mail if it is	
Server admin mail address	Address for fatal errors (such as if an emergency cleanup fails or other fatal errors occur)	

To test whether the entered settings work one can specify an email address to which UrBackup will then send a test mail.

8.2.2 Configure reports

To specify which activities with which errors should be sent via mail you have to go to the 'Logs' page. There on the bottom is a section called 'Reports'. There you can say to which email addresses reports should be sent (e.g. user1@example.com;user2@example.com) and if UrBackup should only send reports about backups that failed/succeeded and contained a log message of a certain level.

If you select the log level 'Info' and 'All' a report about every backup will be sent, because every backup causes at least one info level log message. If you select 'Warning' or 'Error' backups without incidents will not get sent to you.

Every web interface user can configure these values differently. UrBackup only sends reports of client backups to the user supplied address if the user has the 'logs' permission for the specific client. Thus if you want to send reports about one client to a specific email address you have to create a user for this client, login as that user and configure the reporting for that user. The user 'admin' can access logs of all clients and thus also gets reports about all clients.

8.3 Client specific settings

Settings	Description	Default value
Interval for incremental file backups	The server will start incremental file backups in such intervals. ¹	5h
Interval for full file backups	The server will start full file backups in such intervals. ¹	30 days
Interval for incremental image backups	The server will start incremental image backups in such intervals. ¹	7 days
Interval for full image backups	The server will start full image backups in such intervals. ¹	30 days
Maximal number of incremental file backups	Maximal number of incremental file backups for this client. If the number of incremental file backups exceeds this number the server will start deleting old incremental file backups.	100
Minimal number of incremental file backups	Minimal number of incremental file backups for this client. If the server ran out of backup storage space the server can delete incremental file backups until this minimal number is reached. If deleting a backup would cause the number of incremental file backups to be lower than this number it aborts with an error message.	40
Maximal number of full file backups	Maximal number of full file backups for this client. If the number of full file backups exceeds this number the server will start deleting old full file backups.	10
Minimal number of full file backups	Minimal number of full file backups for this client. If the server ran out of backup storage space the server can delete full file backups until this minimal number is reached. If deleting a backup would cause the number of full file backups to be lower than this number it aborts with an error message.	2

¹See section 8.3.2 for time specific intervals.

Maximal number of incremental image backups	Maximal number of incremental image backups for this client. If the number of incremental image backups exceeds this number the server will start deleting old incremental image backups.	30
Minimal number of incremental image backups	Minimal number of incremental image backups for this client. If the server ran out of backup storage space the server can delete incremental image backups until this minimal number is reached. If deleting a backup would cause the number of incremental image backups to be lower than this number it aborts with an error message.	4
Maximal number of full image backups	Maximal number of full image backups for this client. If the number of full image backups exceeds this number the server will start deleting old full image backups.	5
Minimal number of full image backups	Minimal number of full image backups for this client. If the server ran out of backup storage space the server can delete full image backups until this minimal number is reached. If deleting a backup would cause the number of full image backups to be lower than this number it aborts with an error message.	2
Delay after system start up	The server will wait for this number of minutes after discovering a new client before starting any backup	0 min
Backup window	The server will only start backing up clients within this window. See section 8.3.1 for details.	1-7/0-24
Max backup speed for local network	The server will throttle the connections to the client to remain within this speed (see 8.1.13 for setting speed with window).	-
Perform auto-updates silently	If this is selected automatic updates will be performed on the client without asking the user	Unchecked
Soft client quota	During the nightly cleanup UrBackup will remove backups of this client if there are more backups than the minimal number of file/image backups until this quota is met. The quota can be in percent (e.g. 20%) or absolute (e.g. 1500G, 2000M).	""
Excluded files	Allows you to define which files should be excluded from backups. See section 8.3.3 for details	""
Default directories to backup	Default directories which are backed up. See section 8.3.4 for details	""
Volumes to backup	Specifies of which volumes an image backup is done. Separate different drive letters by a semicolon or comma. E.g. 'C;D'. Use the special setting "ALL" to backup all volumes and "ALL_NONUSB" to backup all volumes except those attached via USB.	C
Image backup file format	Either standard VHD (VirtualHard Disk), compressed VHD (VHDZ) or if backup storage is on the btrfs file system "Raw copy-on-write file".	VHD btrfs: Raw cow file
Allow client-side changing of the directories to backup	Allow client(s) to change the directories of which a file backup is done	Checked

Allow client-side starting of incremental/full file backups	Allow the client(s) to start a file backup	Checked
Allow client-side starting of incremental/full image backups	Allow the client(s) to start an image backup	Checked
Allow client-side viewing of backup logs	Allow the client(s) to view the logs	Checked
Allow client-side pausing of backups	Allow the client(s) to pause backups	Checked
Allow client-side changing of settings	If this option is checked the clients can change their client specific settings via the client interface. If you do not check this the server settings always override the clients' settings.	Checked
Allow clients to quit the tray icon	Allow the client(s) to quit the tray icon. If the tray icon is quit current and future backups are paused.	Checked

8.3.1 Backup window

The server will only start backing up clients within the backup windows. The clients can always start backups on their own, even outside the backup windows. If a backup is started it runs till it is finished and does not stop if the backup process does not complete within the backup window. A few examples for the backup window:

1-7/0-24: Allow backups on every day of the week on every hour.

Mon-Sun/0-24: An equivalent notation of the above

Mon-Fri/8:00-9:00, 19:30-20:30;Sat,Sun/0-24: On weekdays backup between 8 and 9 and between 19:30 and 20:30. On Saturday and Sunday the whole time.

As one can see a number can denote a day of the week (1-Monday, 2-Tuesday, 3-Wednesday, 4-Thursday, 5-Friday, 6-Saturday, 7-Sunday). You can also use the abbreviations of the days (Mon, Tues, Wed, Thurs, Fri, Sat, Sun). The times can either consist of only full hours or of hours with minutes. The hours are on the 24 hour clock. You can set multiple days and times per window definition, separated per ", ". If specifying intervals with '-' the starting and ending day are included in the interval. You can also set multiple window definitions. Separate them with ";".

8.3.2 Advanced backup interval

Similar to the backup speed limit (see section 8.1.13) the backup intervals can be specified for different time intervals by combining them with a backup window (see previous section 8.3.1) separated by "@". The most specific backup interval will then be used.

For example, the default backup interval should be one hour and at night (8pm to 6am) it should be 4 hours:

```
1;4@1-7/20-6
```

If additionally the backup interval should be 6 hours during the week-end:

```
1;4@1-5/18-6;6@6,7/0-24
```

8.3.3 Excluded files

You can exclude files with wild card matching. For example if you want to exclude all MP3s and movie files enter something like this:

```
*.mp3;*.avi;*.mkv;*.mp4;*.mpg;*.mpeg
```

If you want to exclude a directory e.g. Temp you can do it like this:

```
*/Temp/*
```

You can also give the full local name

```
C:\Users\User\AppData\Local\Temp\*
```

or the name you gave the location e.g.

```
C_\Users\User\AppData\Local\Temp
```

Rules are separated by a semicolon (";")

8.3.4 Default directories to backup

Enter the different locations separated by a semicolon (";") e.g.

```
C:\Users;C:\Program Files
```

If you want to give the backup locations a different name you can add one with the pipe symbol ("|") e.g.:

```
C:\Users|User files;C:\Program Files|Programs
```

gives the "Users" directory the name "User files" and the "Program files" directory the name "Programs".

Those locations are only the default locations. Even if you check "Separate settings for this client" and disable "Allow client to change settings", once the client modified the paths, changes in this field are not used by the client any more.

Directory flags Each directory to backup has a set of flags. If you do not specify any flags the default flags will be used. Otherwise only the flags you specify are used.

Flags are specified by appending them after the backup location name (separated by "/"). Flags themselves are separated by ",".

Flag	Description	Default
optional	Backup will not fail if the directory is unavailable	Not default
follow_symlinks	Symbolic links which point outside of the specified directory will be followed	Default
symlinks_optional	Backup will not fail if a symbolic link cannot be followed	Default
one_filesystem	Files outside of the first encountered file system will be ignored and not backed up	Not default
require_snapshot	Fail backup if no snapshot/shadow copy can be created of the location	Not default
share_hashes	Share file hashes between different virtual clients	Default
keep	Keep deleted files and directories during incremental backups	Not default

If you want to set the optional flag:

```
C:\Users|User files/follow_symlinks,symlinks_optional,share_hashes,optional
```

(The first three are default flags)

8.3.5 Virtual sub client names

Virtual sub clients allow you to have different file backup sets with one client. Once you specify virtual sub clients, multiple clients will appear with the name “clientname[subclientname]”. You can change all file backup specific options for that client, such as default directories to backup, incremental file backup interval, max number of incremental file backups, . . . The virtual sub client will always be online while the main client (“clientname”) is online.

Separate the virtual sub client names via “|”. E.g.

`system-files|user-files`

8.4 Internet settings

Settings	Description	Default value
Internet server name/IP	The IP or name the clients can reach the server at over the internet	""
Internet server port	The port the server will listen for new clients on	55415
Do image backups over internet	If checked the server will allow image backups for this client/the clients	Not checked
Do full file backups over internet	If checked the server will allow full file backups for this client/the clients	Not checked
Max backup speed for internet connection	The maximal backup speed for the Internet client. Setting this correctly can help avoid saturating the Internet connection of a client (see 8.1.13 for setting speed with window)	-
Total max backup speed for internet connection	The total accumulative backup speed for all Internet clients. This can help avoid saturating the server’s Internet connection (see 8.1.13 for setting speed with window)	-
Encrypted transfer	If checked all data between server and clients is encrypted	Checked
Compressed transfer	If checked all data between server and clients is compressed	Checked
Calculate file-hashes on the client	If checked the client calculates hashes for each file before the backups (only hashes of changed files are calculated). The file then does not have to be transferred if another client already transferred the same file	Not checked
Connect to Internet backup server if connected to local backup server	If checked the client will connect to the configured Internet server, even if it is connect to a backup server on the local network.	Not checked
Do not start file backups if current estimated data usage limit per month is smaller than	The UrBackup server will not schedule file backups if the estimated amount of data the client can transfer with its current Internet connection is smaller than the specified value. See 8.4.1 on how the data usage limit is estimated.	5000 MB
Do not start image backups if current estimated data usage limit per month is smaller than	The UrBackup server will not schedule image backups if the estimated amount of data the client can transfer with its current Internet connection is smaller than the specified value. See 8.4.1 on how the data usage limit is estimated.	20000 MB

8.4.1 Data usage limit estimation

Currently there are two mechanisms to estimate how much data the client can transfer per month with its current Internet connection. If the client is running on Windows 10, the client is connected via Wifi and the Wifi connection is set to be metered, the data usage limit is estimated at 1GB per month. Otherwise the dataplan database at https://github.com/uroni/dataplan_db is used to estimate the data usage limit via the client's hostname. You can see the hostname of your Internet connection e.g. on <http://ipinfo.io/>. If it is missing in the database, please contribute to the dataplan database. The dataplan database contains both hostnames where there is estimated to be a usage limit (e.g. mobile phone connection, plane, ...) and where they are estimated to be unlimited. If the hostname is estimated to be unlimited, UrBackup will always start backups disregarding the setting. If there is no information about the hostname in the database, the data usage limit per month is estimated at 1TB. So if you set one of the settings above 1TB, backups will only start if the hostname is specified as unlimited in the dataplan database. If the hostname has a limit in the database, backups will only start if this limit is greater than the respective setting for the client.

8.5 Advanced settings

In this section you will find global server settings which you only have to change for heavy or custom workloads. Most settings will need a server restart to come into effect.

8.5.1 Enabling temporary file buffers

Earlier versions of UrBackup always saved incoming data from clients first to temporary files and then copied it to the final destination (if the data is new) – the rationale being, that the final destination may be slow and you want to get the data from the client as fast as possible.

With UrBackup 1.1 this default behaviour was changed to directly copy the data to the final backup storage. The two settings allow you to re-enable the old behaviour, e.g., because your backup storage is slow because it is deduplicated. If you re-enable it make sure you have at least 1GB of space for each client, and at least as much space as the biggest file you are going to backup times the number of clients or maximum amount of simultaneous backups (whichever is lower), on your temporary storage. You can change the temporary storage directory via the environment variable *TMPDIR* on GNU/Linux and in the server settings on Windows.

8.5.2 Transfer modes

UrBackup has different transfer modes for files and images. Those are:

- *raw*. Transfer the data as 'raw' as possible. This is the fastest transfer mode and uses the least amount of CPU cycles on server and client.
- *hashed*. Protects the transferred data from bit errors by hashing the data during the transfer. This uses CPU cycles on the client and the server.
UrBackup uses TCP/IP to transfer the images and files. TCP/IP implements its own bit error detection mechanism (CRC32). If the network induces a lot of bit errors and if a lot of data is transferred (>2TB), however, the bit error detection mechanism of TCP/IP is not enough to detect all occurring errors. The 'hashed' transfer mode adds an additional layer of protection to make bit errors less probable. You do not need to use the hashed transfer mode if you backup via a Internet mode connection with enabled encryption, as the encryption layer already protects the integrity of the transmitted data.
- *Block differences - hashed*. Only available for file backups (as it is automatically done for images). Blocks of the transferred files are compared using CRC32 and MD5 hash functions. Only blocks which have changed are sent over the network. In cases where only some blocks of a file change, this reduces the amount of transferred data. It also causes more messages

to be sent between server and client and uses CPU cycles, which is why it is only enabled for Internet clients per default.

8.5.3 Incremental image backup styles

You can select on which backup incremental image backups should be based. Either on the last full or incremental image backup or on the last full backup (sometimes called differential). Keep in mind that image backups can only be deleted if all the image backups they are based on are also deleted.

This option has no effect with the btrfs file system and the raw image file format. There, image backups are always based on the last full or incremental image backup.

8.5.4 Full image backup styles

You can select if full images should transfer all data or only data which changed from the last incremental or full image backup (synthetic full image backup). A synthetic full backup will transfer only changed blocks and store all blocks in the VHD/VHDZ file.

This option has no effect with the btrfs file system and the raw image file format because full images will be disabled automatically there.

8.5.5 Database cache size during batch processing

Amount of memory used for the database cache during batch processing.

8.6 Use symlinks during incremental file backups

If enabled UrBackup will use symbolic links to link unchanged directories with more than 10 directory/files. This will greatly improve the incremental file backup speed, if only few directories are changed, as less hard links have to be created and hard linking operations are expensive on some file systems such as e.g. NTFS and spinning disks.

The disadvantages are:

- UrBackup needs to keep track of the number of times a directory is used. For hard links the operating system does this. UrBackup may be slower and less reliable doing this.
- If the backup storage folder is moved, all symbolic links are invalidated. To restore the symbolic links please run on Windows

```
C:\Program Files\UrBackupServer\remove_unkown.bat
```

and on Linux while the server is not running:

```
urbackupsrv remove-unknown
```

Per default symbolic links are used. If the btrfs mode is used this option has no effect.

8.7 Debugging: End-to-end verification of all file backups

This is a setting for debugging purposes or for the paranoid. If end-to-end verification is enabled UrBackup clients will create file hashes for every file for every file backup reading every file that is to be backed up. At the end of the backup process the hashes of the files stored on the server are compared to the hashes calculated on the client. If hashes differ the backup fails and an email is sent to the server admin.

8.8 Debugging: Verify file backups using client side hashes

At the end of file backups the server will go over all files in the backup and compare the file hashes with the client-side hashes.

8.9 Periodically readd file entries of internet clients to database

If checked UrBackup will periodically refresh hashes in its database such that backup deletions do not cause UrBackup to re-download files that are already present in the backup storage.

8.10 Create symbolically linked views for each user on the clients after file backups

After a successful file backup UrBackup will create symbolically linked views for each used on the client machine on the backup server. Those views can then be made accessible e.g. via samba file sharing.

8.11 Maximum number of simultaneous jobs per client

Maximum number of simultaneous jobs per client and all its virtual sub-clients. Increase this if you want it to e.g. simultaneously perform image and file backups.

8.12 Volumes to snapshot in groups during image backups

Specifies which volumes UrBackup should snapshot together when doing image backups. When snapshotted together volumes are consistent with each other. If for example one volume contains the database and the other the database log file, volumes should be snapshotted together to get a valid backup.

Can either be “all” to snapshot all volumes simultaneously, a list of volumes separated by comma (e.g. “C,D”) or a list of lists separated by the pipe symbol (e.g. “C,D|E,F”).

8.13 Volumes to snapshot in groups during file backups

Specifies which volumes UrBackup should snapshot together when doing file backups. Currently only the supported via the Windows client. When snapshotted together volumes are consistent with each other. If for example one volume contains the database and the other the database log file, volumes should be snapshotted together to get a valid backup.

Can either be “all” to snapshot all volumes simultaneously, a list of volumes separated by comma (e.g. “C,D”) or a list of lists separated by the pipe symbol (e.g. “C,D|E,F”).

8.14 Windows components backup configuration

Configuration of the Windows component backup via Windows backup API. Currently, this is best configured on the client via GUI (accessible from the tray icon or by running *UrBackupClient.exe select WindowsComponents*) or copy and pasting from a client configured via GUI.

A setting of “*default = 1*” (the default) will automatically backup Microsoft Exchange, Microsoft SQL Server and Microsoft Hyper-V. A setting of “*default = 0*” will disable Windows component backup.

9 Restoring backups

UrBackup protects whole machines from disaster by creating image backups and a users or servers files by creating file backups. Because the file backups size can usually be reduced by focusing on the most important data on a machine they can usually be run more often than the image

backups. It makes sense to use image and file backups in tandem, backing up the whole machine less regularly than the important files via file backups.

9.1 Restoring image backups

Image backups can be restored with a Debian GNU/Linux based bootable CD/USB-stick. During image restore the machine to be restored must be reachable without network address translation from the server (or you forward the client ports in sections 10.3 to the restore client). While Linux supports many mainboards, disk controllers etc. you should always verify that the restore CD works on your specific hardware especially if you use exotic or new hardware. Drivers and firmware for some wireless devices and a program to configure is included but restoring via a wired network connection will be less trouble and faster and should be preferred. The restore itself is easy to use. After startup it will look for a backup server. If it does not find one, you can enter the backup server's IP/hostname and change your networking settings. After a backup server is found it will ask for a username and password. Use for example your admin account to access all clients and their image backups. Then you can select one image backup, select the disk you want to restore to and then it will restore. The target disk must be at least as large as the disk which was image backed up. Some hardware changes may cause Windows to bluescreen on startup after restore. If the startup repair fails, you may have to do a repair install using a Windows disk. You should test the different hardware combinations beforehand if you plan on restoring Windows to different hardware.

9.2 Restoring file backups

When performing file backups UrBackup creates a file system snapshot identical to the client's file system at that point in time. As such, you can simply make those file system snapshots available to the clients via any file sharing protocol or application, such as Windows file sharing (or samba), FTP/SFTP, WebDAV...

You can also create a user for client(s), which allows the user to browse all backups of the client(s) via the UrBackup web interface and download individual files or whole directories as ZIP (limited to max. 4GB).

Since UrBackup 2.0.x users can directly access the web interface from the client if a server URL is configured. Either they right-click on the UrBackup tray icon and then click "Access/restore backups" which opens the browser, or they can right click a file/directory in a backup path and then click on "Access/restore backups" to access all backups of a file/directory.

When browsing backups the web interface will show a restore button if the client is online. The restore will ask for user confirmation. If the client includes a GUI component (tray icon), the user confirmation will popup for all active users on the client to be restored. If not acknowledged in time (timeout) or if declined the restore will fail. You can change this behaviour in *C:\Program files \UrBackup \args.txt* by changing "default" to "server-confirms" on Windows, or by changing the restore setting in */etc/default/urbackupclient* or */etc/sysconfig/urbackupclient* on Linux.

UrBackup is setup this way because a theoretical data loss scenario is an attacker taking control of your backup server, deleting all backups and then deleting all files on the clients via restores.

On Linux (and the other operating systems) you can also restore via command line from the client using *urbackupclientctl browse* and *urbackupclientctl restore-start*.

10 Miscellaneous

10.1 Manually update UrBackup clients

You should test UrBackup clients before using them on the clients. This means UrBackup should not automatically download the newest client version from the Internet and install it. This means disabling the autoupdate described in Section 8.1.8. You can still centrally update the client from the server if you disabled autoupdate. Go to <https://hndl.urbackup.org/Client/> and download all files in the *update* folder of the current client to */var/urbackup* on Linux and *C:\Program Files\UrBackupServer\urbackup* per default on Windows. UrBackup will then push the new version to the clients once they reconnect. If you checked silent autoupdates, the new version will be installed silently on the clients, otherwise there will be a popup asking the user to install the new version.

10.2 Logging

UrBackup generally logs all backup related things into several log facilities. Each log message has a certain severity, namely *error*, *warning*, *info* or *debug*. Each log output can be filtered by this severity, such that e.g. only errors are shown. Both server and client have separate logs. During a backup process the UrBackup server tries to log everything which belongs to a certain backup in a client specific logs and at the end sends this log to the client. Those are the logs you see on the client interface. The same logs can also be viewed via the web interface in the “Logs” section. One can also send them per mail as described in subsection 8.2.2.

Everything which cannot be accredited to a certain client or which would cause too much log traffic is logged in a general log file. On Linux this is */var/log/urbackup.log* on Windows *C:\Program files\UrBackupServer\urbackup.log* for the server per default. The client has as defaults */var/log/urbackup_client.log* and *C:\Program files\UrBackup\debug.log*. Per default those files only contain log messages with severity *warning* or higher. In Windows there is a *args.txt* in the same directory as the log file. Change *warn* here to *debug*, *info* or *error* to get a different set of log messages. You need to restart the server for this change to come into effect. On Linux this depends on the distribution. On Debian one changes the setting in */etc/default/urbackupsrv*.

10.3 Used network ports

The Server binds to following default ports:

Port	Usage	Incoming/Outgoing	Protocol
55413	Fast CGI for web interface	Incoming	TCP
55414	HTTP web interface	Incoming	TCP
55415	Internet clients	Incoming	TCP
35623	UDP broadcasts for discovery	Outgoing	UDP

The Client binds to following default ports (all incoming):

Port	Usage	Protocol
35621	Sending files during file backups (file server)	TCP
35622	UDP broadcasts for discovery	UDP
35623	Commands and image backups	TCP

10.4 Mounting (compressed) VHD files on GNU/Linux

If you compiled UrBackup with fuse (file system in user space) support or installed the Debian/Ubuntu packages the UrBackup server can mount VHD(Z) files directly. You compile UrBackup with fuse support by configuring:

```
./configure --with-mountvhd
```

You will be able to mount a VHD(Z) file e.g. with

```
urbackupsrv mount-vhd --file /media/backup/urbackup/testclient/\
Image_C_140420-1956.vhdx --mountpoint /media/testclient_C
```

All files in the backed up “C” volume will then be available read-only in `/media/testclient_C`. Unmount the mounts created by UrBackup (see output of `mount`), to stop the background process.

10.5 Mounting VHDs as a volume on Windows

Starting from Vista Windows can mount VHD files directly. If the VHD files are compressed, they need to be decompressed first (see next section 10.6). Then go to system settings->Manage->Computer management->Disk management->Other Actions->Add Virtual Hard Disk. Mount the VHD file read only. The image will appear as another drive in the explorer. It may not work if the VHD file is on a network drive.

10.6 Decompress VHD files

Consider using the method described in the next section 10.7 to decompress VHD files.

If you want to mount the VHD files on Windows and they are compressed (file extension is VHDZ), you need to decompress them first. Use

```
C:\Program Files\UrBackupServer\uncompress_image.bat
```

for that. Calling the batch file without parameters will open a file selection screen where you can select the VHDZ file to be decompressed. A temporary inflated copy is created and renamed in-place once the decompression is done. If the image is incremental the parent-VHD is automatically decompressed as well. If you want to prevent this please use the method described in section 10.7 to build a separate uncompressed image. All the image files will still have the VHDZ extension, as otherwise it would have to change database entries, but the files will not be compressed anymore. On Linux the same thing can be done with `urbackupsrv decompress-file -f [filename]`.

10.7 Assemble multiple volume VHD images into one disk VHD image

UrBackup stores each volume of an image backup separately. If you want to boot an image backup, without using the restore CD, as a virtual machine you have to re-assemble multiple volumes into one disk VHD image. On Windows this can be done by running `C:\Program Files\UrBackupServer\assemble_disk_image.bat`. In a first step it will ask for the VHD images to assemble. Select e.g. `Image_C_XXXXX.vhd` and `Image_SYSVOL_XXXXX.vhd`. The source images can also be incremental or compressed. Then it will ask where the output VHD disk image should be saved. After that it will write the master boot record from `Image_C_XXXXX.vhd.mbr` and the contents of the volumes into the output disk image at the appropriate offsets.

On Linux the same thing can be done with

```
urbackupsrv -a /full/path/Image_C_XXXXX.vhdx -a /full/path/Image_SYSVOL_XXXXX.vhdx \
-o full_disk.vhd
```

This tool can also be used to decompress images without decompressing their parents by selecting a single VHD file as input.

10.8 Migrate non-btrfs backup storage

UrBackup has a build-in migration function which allows you to migrate the backup storage to different devices with minimal downtime. This only works with plain backup storage, that is, it does not work with the special btrfs mode described in section 11.7.2 or copy-on-write images described in section 11.7.1. The migration functionality is only suited for migration and not for e.g. mirroring the backups. Compared to using e.g. `rsync`, the build-in migration has some advantages:

- It can resume the migration at backup level
- It automatically disables all clean-ups to minimize the amount of necessary changes during migration
- UrBackup continues to backup clients which are not currently being migrated and migrates new backups automatically
- UrBackup handles the hard-links correctly

Create a file named “*migrate_storage_to*” in `/var/urbackup` or `C:\Program Files\UrBackupServer\urbackup` with the only content of the file being the path where you want to migrate the backups to. Afterwards restart UrBackup server and the migration should begin. You can see the migration progress on the activities page. After the migration is finished change the backup storage path to the new location in the settings or mount the new backup storage at the old location. Afterwards run “remove unknown” described in section 11.4.

11 Storage

The UrBackup server storage system is designed in a way that it is able to save as much backups as possible and thus uses up as much space on the storage partition as possible. With that in mind it is best practice to use a separate file system for the backup storage or to set a quota for the ‘urbackup’ user. Some file systems behave badly if they are next to fully occupied (fragmentation and bad performance). With such file systems you should always limit the quota UrBackup can use up to say 95% of all the available space. You can also setup a soft quota within UrBackup (see section 8.1.14) which causes UrBackup to delete backups to stay within this quota, if possible.

11.1 Nightly backup deletion

UrBackup automatically deletes old file and image backups between 3am and 5am. Backups are deleted when a client has more incremental/full file/image backups then the configured maximum number of incremental/full file/image backups. Backups are deleted until the number of backups is within these limits again.

If the administrator has turned automatic shut-down on, this clean up process is started on server start up instead (as the server is most likely off during the night). Deleting backups and the succeeding updating of statistics can have a huge impact on system performance.

During nightly backup deletion UrBackup also tries to enforce the global and client specific soft quotas. It is only able to delete backups if a client has already more backups than the configured minimal number of incremental/full file/image backups.

11.2 Emergency cleanup

If the server runs out of storage space during a backup it deletes backups until enough space is available again. Images are favoured over file backups and the oldest backups are deleted first. Backups are only deleted if there are at least the configured minimal number of incremental/full file/image backups other file/image backups in storage for the client owning the backup. If no such backup is found UrBackup cancels the current backup with a fatal error. Administrators should monitor storage space and add storage or configure the minimal number of incremental/full file/image backups to be lower if such an error occurs.

11.3 Cleanup for servers with file backups with lots of files

UrBackup’s database is in a mode which enables high concurrency. Since the cleanup procedure can sometimes be bottlenecked by the database it may be advisable to switch the database into a mode which allows less concurrency but is fast for some operations for the cleanup procedure.

This is not possible while UrBackup is running, so you should tweak the backup window such that you can be sure there are no backups running at some point. Then you can stop the server run the cleanup separately by calling

```
urbackupsrv cleanup --amount x
```

on GNU/Linux or on Windows:

```
cleanup.bat x
```

Where x is the percent of space to free on the backup storage or the number of Bytes/ Megabytes/ Gigabytes e.g. “20G” or “10%”. If it should only delete old backups use “0%”.

11.4 Cleaning the storage folder of files not known by UrBackup

Sometimes, e.g., by using a database backup, there are backups in the storage directory which UrBackup does not know about, i.e., there are no entries for those backups in the database. Or there are entries in the database which are not in the storage directory (any more). In those cases the command

```
urbackupsrv remove-unknown
```

on GNU/Linux or on Windows:

```
remove_unknown.bat
```

removes files and folders in the urbackup storage directory which are not present in the UrBackup database.

This also goes over all symlinks and corrects them if necessary.

11.5 Archiving

UrBackup has the ability to automatically archive file backups. Archived file backups cannot be deleted by the nightly or emergency clean up – only when they are not archived any more. You can setup archival under Settings->Archival for all or specific clients. When an archival is due and the the server is currently in a archival window (See 11.5.1) the last file backup of the selected type will be archived for the selected amount of time. After that time it will be automatically not archived any more. You can see the archived backups in the “Backups” section. If a backup is archived for only a limited amount of time there will be a time symbol next to the check mark. Hovering over that time symbol will tell you how long that file backup will remain archived.

11.5.1 Archival window

The archival window allows you to archive backups at very specific times. The format is very similar to *crontab*. The fields are the same except that there are no minutes:

Field	Allowed values	Remark
Hour	0-23	
Day of month	1-31	
Month	1-12	No names allowed
Day of week	0-7	0 and 7 are Sunday

To archive a file backup on the first Friday of every month we would then set “Archive every” to something like 27 days. After entering the time we want the backups archived for we would then add

;;*;5

as window (hour;day of month;month;day of week). To archive a backup every Friday we would set “Archive every” to a value greater than one day but less than 7 days. This works because both conditions have to apply: The time since the last backup archival must be greater than “Archive every” and the server must be currently in the archive window.

Other examples are easier. To archive a backup on the first of every month the window would be

;1;;*

and “Archive every” something like 2-27 days.

One can add several values for every field by separating them via a comma such that

;;*;3,5

and “Archive every” one day would archive a backup on Wednesday and Friday. Other advanced features found in *crontab* are not present.

11.6 Suitable file systems

Because UrBackup has the option to save all incoming data to temporary files first (see Section 8.5.1) and then copies them to the final location in parallel backup performance will still be good even if the backup storage space is slow. This means you can use a fully featured file system with compression and de-duplication without that much performance penalty. At the worst the server writes away an image backup over the night (having already saved the image’s contents into temporary files during the day). This section will show which file systems are suited for UrBackup.

11.6.1 Ext4/XFS

Ext4 and XFS, are both available in Linux and can handle big files, which is needed for storing image backups. They do not have compression or de duplication though. Compression can be achieved by using a fuse file system on top of them such as fusecompress. There are some block-level de-duplication fuse layers as well, but I would advise against them as they do not seem very stable. You will have to use the kernel user/group level quota support to limit the UrBackup storage usage.

11.6.2 NTFS

NTFS is pretty much the only option you have if you run the UrBackup server under Windows. It supports large files and compression as well as hard links and as such is even more suited for UrBackup than the standard Linux file systems XFS and Ext4.

11.6.3 btrfs

Btrfs is a next generation Linux file system comparable to ZFS. It supports compression and offline block-level deduplication. UrBackup has a special snapshotting backup mode which makes incremental backups and deleting file backups much faster with btrfs. With btrfs UrBackup also does a cheap (in terms of CPU und memory requirements) block-level deduplication on incremental file backups. See 11.7.2 for details. UrBackup also has a special copy-on-write raw image backup format which allows “incremental forever” style image backups.

11.6.4 ZFS

ZFS is a file system originating from Solaris. It is available as a fuse module for Linux (*zfs-fuse*) and as a kernel module (*ZFSOnLinux*). There are licensing issues which are preventing ZFS from directly integrating with Linux. If you want the most performance and stability an option would be using FreeBSD (e.g. as FreeNAS). ZFS has some neat features like compression, block-level

de-duplication, snapshots and build in raid support that make it well suited for backup storage. How to build a UrBackup server with ZFS is described in detail in section 11.7.1. UrBackup also has a special copy-on-write raw image backup format which allows “incremental forever” style image backups which works on ZFS. Section 11.7.1 describes how to set it up.

11.7 Storage setup proposals

In this section a sample storage setup with ZFS is shown which allows off-site backups via Internet or via tape like manual off-site storage and a storage setup using the Linux file system btrfs using the btrfs snapshot mechanism to speed up file backup creation and destruction and to save the file backups more efficiently.

11.7.1 ZFS

Note: It is assumed that UrBackup runs on a UNIX like system such as Linux or BSD that supports ZFS. We will use all ZFS features such as compression, de-duplication and snapshots. It is assumed that the server has two hard drives (sdb,sdc) dedicated to backups and a hot swappable hard drive slot (sdd). It is assumed there is a caching device to speed up de-duplication as well in /dev/sde. Even a fast USB stick can speed up de-duplication because it has better random access performance than normal hard disks. Use SSDs for best performance.

First setup the server such that the temporary directory (/tmp) is on a sufficiently large performant file system. If you have a raid setup you could set /tmp to be on a striped device. We will now create a backup storage file system in /media/BACKUP.

Create a ZFS-pool 'backup' from the two hard drives. The two are mirrored. Put a hard drive of the same size into the hot swappable hard drive slot. We will mirror it as well:

```
zpool create backup /dev/sdb /dev/sdc /dev/sdd cache /dev/sde -m /media/BACKUP
```

Enable de-duplication and compression. You do not need to set a quota as de-duplication fragments everything anyway (that's why we need the caching device).

```
zfs set dedup=on backup
zfs set compression=on backup
```

Now we want to implement a grandfather, father, son or similar backup scheme where we can put hard disks in a fireproof safe. So each time we want to have an off-site backup we remove the hot swappable device and plug in a new one. Then we either run

```
zpool replace backup /dev/sdd /dev/sdd
```

or

```
zpool scrub
```

You can see the progress of the re-silvering/scrub with 'zpool status'. Once it is done you are ready to take another hard disk somewhere.

Now we want to save the backups on a server on another location. First we create the ZFS backup pool on this other location.

Then we transfer the full file system (otherserver is the host name of the other server):

```
zfs snapshot backup@last
zfs send backup@last | ssh -l root otherserver zfs recv backup@last
```

Once this is done we can sync the two file systems incrementally:

```
zfs snapshot backup@now
ssh -l root otherserver zfs rollback -r backup@last
zfs send -i backup@last backup@now | ssh -l root otherserver zfs recv backup@now
```

```
zfs destroy backup@last
zfs rename backup@last backup@now
ssh -l root otherserver zfs destroy backup@last
ssh -l root otherserver zfs rename backup@last backup@now
```

You can also save these full and incremental zfs streams into files on the other server and not directly into a ZFS file system.

Copy-on-write raw image backups with ZFS Since UrBackup 2.1.x ZFS storage also allows the copy-on-write raw image backup format to be used. This format has no size limitation and allows for an “incremental forever” style image backup. UrBackup puts each image backup into a separate dataset and stores the image as a single big file. Compression and unused area management are done by ZFS.

In order to create and remove ZFS snapshots UrBackup installs a setuid executable *urbackup_snapshot_helper*. However, currently ZFSOnLinux calls *mount* in turn which will fail when *urbackup_snapshot_helper* is run as non-privileged user. So you must run UrBackup server as *root* user (*urbackupsrv run -u root*). You should create a separate ZFS dataset in which the image backups will be stored e.g. *tank/images*. UrBackup server will read this dataset from */etc/urbackup/dataset* and the backup storage path from */etc/urbackup/backupfolder*. Setup by e.g.

```
mkdir -p /etc/urbackup
echo "tank/images" > /etc/urbackup/dataset
echo "/mnt/BACKUP/urbackup" > /etc/urbackup/backupfolder
```

Afterwards test if everything works correctly by running as root:

```
urbackup_snapshot_helper test
```

You should then be able to enjoy much faster incremental file backups which use less storage space and “incremental forever” style image backups.

Other caveats: UrBackup uses “file hole punching” to remove unused areas from the image file after an incremental backup. FreeBSD currently does not support this, so it will not be able to remove those unused areas after incremental image backups on FreeBSD. It sets them to zero instead, though, so if you enabled ZFS compression unused areas will not take up much space.

11.7.2 Btrfs

Btrfs is an advanced file system for Linux capable of creating copy on write snapshots of sub-volumes. For UrBackup to be able to use the snapshotting mechanism the Linux kernel must be at least 3.6.

If UrBackup detects a btrfs file system it uses a special snapshotting file backup mode. It saves every file backup of every client in a separate btrfs sub-volume. When creating an incremental file backup UrBackup then creates a snapshot of the last file backup and removes, adds and changes only the files required to update the snapshot. This is much faster than the normal method, where UrBackup links (hard link) every file in the new incremental file backups to the file in the last one. It also uses less metadata (information about files, i.e., directory entries). If a new/changed file is detected as the same as a file of another client or the same as in another backup, UrBackup uses cross device reflinks to save the data in this file only once on the file system. Using btrfs also allows UrBackup to backup files changed between incremental backups in a way that only changed data in the file is stored. This greatly decreases the storage amount needed for backups, especially for large database files (such as e.g. the Outlook archive file). The ZFS deduplication in the previous section (11.7.1) saves even more storage, but comes at a much greater cost in form of a massive decrease of read and write performance and high CPU and memory requirements.

With btrfs UrBackup can also use a special raw image file format. This format has no size limitation and allows for an “incremental forever” style image backup. UrBackup puts each image backup into a separate subvolume and stores the image as a single big file. Compression and unused area management are done by btrfs.

In order to create and remove btrfs snapshots UrBackup installs a setuid executable *urbackup_snapshot_helper*. UrBackup also uses this tool to test if cross-device reflinks are possible. Only if UrBackup can create cross-device reflinks and is able to create and destroy btrfs snapshots, is the btrfs mode enabled. *urbackup_snapshot_helper* needs to be told separately where the UrBackup backup folder is. This path is read from */etc/urbackup/backupfolder*. Thus, if */media/backup/urbackup* is the folder where UrBackup is saving the paths, following commands would properly create this file:

```
mkdir /etc/urbackup
echo "/media/backup/urbackup" > /etc/urbackup/backupfolder
```

You can then test if UrBackup will use the btrfs features via running

```
urbackup_snapshot_helper test
```

If the test fails, you need to check if the kernel is new enough and that the backup folder is on a btrfs volume.

You should then be able to enjoy much faster incremental file backups which use less storage space and “incremental forever” style image backups.