



**Borg Documentation**  
*Release 1.0.13.dev25+g8297828*

**The Borg Collective**

2018-11-05



---

# Contents

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	What is BorgBackup? . . . . .	2
1.2	Helping, Donations and Bounties . . . . .	3
1.3	Links . . . . .	4
1.4	Compatibility notes . . . . .	4
<b>2</b>	<b>Installation</b>	<b>5</b>
2.1	Pre-Installation Considerations . . . . .	5
2.2	(G)LIBC requirements . . . . .	5
2.3	Distribution Package . . . . .	6
2.4	Standalone Binary . . . . .	6
2.5	Features & platforms . . . . .	7
2.6	From Source . . . . .	8
<b>3</b>	<b>Quick Start</b>	<b>13</b>
3.1	Important note about free space . . . . .	13
3.2	A step by step example . . . . .	13
3.3	Automating backups . . . . .	15
3.4	Pitfalls with shell variables and environment variables . . . . .	15
3.5	Backup compression . . . . .	16
3.6	Repository encryption . . . . .	16
3.7	Remote repositories . . . . .	17
<b>4</b>	<b>Usage</b>	<b>19</b>
4.1	General . . . . .	19
4.2	borg init . . . . .	24
4.3	borg create . . . . .	25
4.4	borg extract . . . . .	28
4.5	borg check . . . . .	30
4.6	borg rename . . . . .	31
4.7	borg list . . . . .	32
4.8	borg delete . . . . .	33
4.9	borg prune . . . . .	35
4.10	borg info . . . . .	36
4.11	borg mount . . . . .	38
4.12	borg umount . . . . .	39
4.13	borg key export . . . . .	39

4.14	borg key import	40
4.15	borg change-passphrase	41
4.16	borg serve	42
4.17	borg upgrade	43
4.18	borg break-lock	46
4.19	Miscellaneous Help	46
4.20	Debug Commands	48
4.21	Additional Notes	49
<b>5</b>	<b>Deployment</b>	<b>55</b>
5.1	Machines	55
5.2	User and group	55
5.3	Folders	56
5.4	Restrictions	56
5.5	Client	57
5.6	Ansible	57
5.7	Enhancements	58
5.8	See also	58
<b>6</b>	<b>Frequently asked questions</b>	<b>59</b>
6.1	Can I backup VM disk images?	59
6.2	Can I backup from multiple servers into a single repository?	60
6.3	Can I copy or synchronize my repo to another location?	60
6.4	Which file types, attributes, etc. are <i>not</i> preserved?	60
6.5	Are there other known limitations?	60
6.6	Why is my backup bigger than with attic? Why doesn't Borg do compression by default?	61
6.7	How can I specify the encryption passphrase programmatically?	61
6.8	When backing up to remote encrypted repos, is encryption done locally?	61
6.9	When backing up to remote servers, do I have to trust the remote server?	61
6.10	How can I protect against a hacked backup client?	61
6.11	How can I protect against a hacked backup server?	62
6.12	How can I protect against theft, sabotage, lightning, fire, ...?	62
6.13	How do I report security issue with Borg?	62
6.14	Why do I get "connection closed by remote" after a while?	62
6.15	Why am I seeing idle borg serve processes on the repo server?	63
6.16	The borg cache eats way too much disk space, what can I do?	63
6.17	If a backup stops mid-way, does the already-backed-up data stay there?	63
6.18	How can I backup huge file(s) over a unstable connection?	64
6.19	Can I backup my root partition (/) with Borg?	64
6.20	If it crashes with a UnicodeError, what can I do?	64
6.21	I can't extract non-ascii filenames by giving them on the commandline!?	64
6.22	Can Borg add redundancy to the backup data to deal with hardware malfunction?	65
6.23	Can Borg verify data integrity of a backup archive?	65
6.24	I am seeing 'A' (added) status for a unchanged file!?	65
6.25	It always chunks all my files, even unchanged ones!	65
6.26	Is there a way to limit bandwidth with Borg?	66
6.27	I am having troubles with some network/FUSE/special filesystem, why?	66
6.28	Why does running 'borg check --repair' warn about data loss?	67
6.29	Why isn't there more progress / ETA information displayed?	67
6.30	Why am I getting 'Operation not permitted' errors when backing up on sshfs?	67
6.31	Requirements for the borg single-file binary, esp. (g)libc?	67
6.32	Why was Borg forked from Attic?	68
6.33	What are the differences between Attic and Borg?	68

<b>7</b>	<b>Support</b>	<b>71</b>
7.1	Issue Tracker . . . . .	71
7.2	Chat (IRC) . . . . .	71
7.3	Mailing list . . . . .	71
7.4	Twitter . . . . .	72
7.5	Bounties and Fundraisers . . . . .	72
7.6	Security . . . . .	72
7.7	Verifying signed releases . . . . .	72
<b>8</b>	<b>Resources</b>	<b>75</b>
8.1	Videos, Talks, Presentations . . . . .	75
8.2	Software . . . . .	75
<b>9</b>	<b>Important notes</b>	<b>77</b>
9.1	Pre-1.0.9 manifest spoofing vulnerability (CVE-2016-10099) . . . . .	77
9.2	Pre-1.0.9 potential data loss . . . . .	78
9.3	Pre-1.0.4 potential repo corruption . . . . .	78
<b>10</b>	<b>Changelog</b>	<b>81</b>
10.1	Version 1.0.12 (2018-04-08) . . . . .	81
10.2	Version 1.0.11 (2017-07-21) . . . . .	83
10.3	Version 1.0.11rc1 (2017-06-27) . . . . .	83
10.4	Version 1.0.10 (2017-02-13) . . . . .	85
10.5	Version 1.0.10rc1 (2017-01-29) . . . . .	86
10.6	Version 1.0.9 (2016-12-20) . . . . .	87
10.7	Version 1.0.9rc1 (2016-11-27) . . . . .	88
10.8	Version 1.0.8 (2016-10-29) . . . . .	89
10.9	Version 1.0.8rc1 (2016-10-17) . . . . .	90
10.10	Version 1.0.7 (2016-08-19) . . . . .	92
10.11	Version 1.0.7rc2 (2016-08-13) . . . . .	92
10.12	Version 1.0.7rc1 (2016-08-05) . . . . .	93
10.13	Version 1.0.6 (2016-07-12) . . . . .	94
10.14	Version 1.0.6rc1 (2016-07-10) . . . . .	95
10.15	Version 1.0.5 (2016-07-07) . . . . .	95
10.16	Version 1.0.4 (2016-07-07) . . . . .	96
10.17	Version 1.0.3 (2016-05-20) . . . . .	98
10.18	Version 1.0.2 (2016-04-16) . . . . .	98
10.19	Version 1.0.1 (2016-04-08) . . . . .	99
10.20	Version 1.0.0 (2016-03-05) . . . . .	100
10.21	Version 1.0.0rc2 (2016-02-28) . . . . .	101
10.22	Version 1.0.0rc1 (2016-02-07) . . . . .	102
10.23	Version 0.30.0 (2016-01-23) . . . . .	103
10.24	Version 0.29.0 (2015-12-13) . . . . .	105
10.25	Version 0.28.2 (2015-11-15) . . . . .	106
10.26	Version 0.28.1 (2015-11-08) . . . . .	107
10.27	Version 0.28.0 (2015-11-08) . . . . .	107
10.28	Version 0.27.0 (2015-10-07) . . . . .	109
10.29	Version 0.26.1 (2015-09-28) . . . . .	110
10.30	Version 0.26.0 (2015-09-19) . . . . .	110
10.31	Version 0.25.0 (2015-08-29) . . . . .	111
10.32	Version 0.24.0 (2015-08-09) . . . . .	113
10.33	Version 0.23.0 (2015-06-11) . . . . .	115
10.34	Attic Changelog . . . . .	116
<b>11</b>	<b>Internals</b>	<b>121</b>

11.1	Repository and Archives	121
11.2	Lock files	121
11.3	Config file	122
11.4	Keys	122
11.5	Segments and archives	122
11.6	The manifest	123
11.7	The Archive	123
11.8	The Item	124
11.9	Chunks	124
11.10	Indexes / Caches	125
11.11	Indexes / Caches memory usage	126
11.12	Encryption	127
11.13	Key files	128
11.14	Compression	129
<b>12</b>	<b>Development</b>	<b>131</b>
12.1	Contributions	131
12.2	Branching model	132
12.3	Code and issues	133
12.4	Style guide	133
12.5	Continuous Integration	133
12.6	Output and Logging	133
12.7	Building a development environment	133
12.8	Running the tests	133
12.9	Regenerate usage files	134
12.10	Building the docs with Sphinx	134
12.11	Using Vagrant	134
12.12	Creating standalone binaries	135
12.13	Creating a new release	135
<b>13</b>	<b>Authors</b>	<b>137</b>
13.1	Borg authors (“The Borg Collective”)	137
13.2	Attic authors	137
<b>14</b>	<b>License</b>	<b>139</b>

# CHAPTER 1

## Introduction

```
tw@tux:~/borg/demo$ # change some data slightly:
tw@tux:~/borg/demo$ echo "some more data" > data/one_file_more
tw@tux:~/borg/demo$ # now create a second backup:
tw@tux:~/borg/demo$ borg create --stats --progress repo::backup2 data
-----
Archive name: backup2
Archive fingerprint: 5737afe8ad2cda7667973b7f2e1d83f097ef3117b5753a38ba7664b616f
bdc5a
Start time: Sat Oct 24 22:28:24 2015
End time: Sat Oct 24 22:28:27 2015
Duration: 3.41 seconds
Number of files: 31557

Original size      Compressed size    Deduplicated size
This archive:      257.06 MB          114.47 MB          45.19 kB
All archives:      514.12 MB          228.92 MB          114.31 MB

Unique chunks      Total chunks
Chunk index:       33733              68060
-----
tw@tux:~/borg/demo$ # wow, that was a lot faster!
tw@tux:~/borg/demo$ # notice that "Deduplicated size" for "This archive"!
tw@tux:~/borg/demo$ # borg recognized that most files did not change and
tw@tux:~/borg/demo$ # deduplicated them
```

<sup>1</sup> <https://asciinema.org/a/28691?autoplay=1&speed=2>

## 1.1 What is BorgBackup?

BorgBackup (short: Borg) is a deduplicating backup program. Optionally, it supports compression and authenticated encryption.

The main goal of Borg is to provide an efficient and secure way to backup data. The data deduplication technique used makes Borg suitable for daily backups since only changes are stored. The authenticated encryption technique makes it suitable for backups to not fully trusted targets.

See the [installation manual](#)<sup>5</sup> or, if you have already downloaded Borg, `docs/installation.rst` to get started with Borg.

### 1.1.1 Main features

**Space efficient storage** Deduplication based on content-defined chunking is used to reduce the number of bytes stored: each file is split into a number of variable length chunks and only chunks that have never been seen before are added to the repository.

A chunk is considered duplicate if its `id_hash` value is identical. A cryptographically strong hash or MAC function is used as `id_hash`, e.g. `(hmac-)sha256`.

To deduplicate, all the chunks in the same repository are considered, no matter whether they come from different machines, from previous backups, from the same backup or even from the same single file.

Compared to other deduplication approaches, this method does NOT depend on:

- file/directory names staying the same: So you can move your stuff around without killing the deduplication, even between machines sharing a repo.
- complete files or time stamps staying the same: If a big file changes a little, only a few new chunks need to be stored - this is great for VMs or raw disks.
- The absolute position of a data chunk inside a file: Stuff may get shifted and will still be found by the deduplication algorithm.

#### Speed

- performance critical code (chunking, compression, encryption) is implemented in C/Cython
- local caching of files/chunks index data
- quick detection of unmodified files

**Data encryption** All data can be protected using 256-bit AES encryption, data integrity and authenticity is verified using HMAC-SHA256. Data is encrypted clientside.

**Compression** All data can be compressed by lz4 (super fast, low compression), zlib (medium speed and compression) or lzma (low speed, high compression).

**Off-site backups** Borg can store data on any remote host accessible over SSH. If Borg is installed on the remote host, big performance gains can be achieved compared to using a network filesystem (`sshfs`, `nfs`, ...).

**Backups mountable as filesystems** Backup archives are mountable as userspace filesystems for easy interactive backup examination and restores (e.g. by using a regular file manager).

**Easy installation on multiple platforms** We offer single-file binaries that do not require installing anything - you can just run them on these platforms:

- Linux

---

<sup>5</sup> <https://borgbackup.readthedocs.org/en/stable/installation.html>



- Mac OS X
- FreeBSD
- OpenBSD and NetBSD (no xattrs/ACLs support or binaries yet)
- Cygwin (not supported, no binaries yet)
- Linux Subsystem of Windows 10 (not supported)

### Free and Open Source Software

- security and functionality can be audited independently
- licensed under the BSD (3-clause) license, see [License](#)<sup>6</sup> for the complete license

## 1.1.2 Easy to use

Initialize a new backup repository and create a backup archive:

```
$ borg init /path/to/repo
$ borg create /path/to/repo::Saturday1 ~/Documents
```

Now doing another backup, just to show off the great deduplication:

```
$ borg create -v --stats /path/to/repo::Saturday2 ~/Documents
-----
Archive name: Saturday2
Archive fingerprint: 622b7c53c...
Time (start): Sat, 2016-02-27 14:48:13
Time (end):   Sat, 2016-02-27 14:48:14
Duration: 0.88 seconds
Number of files: 163
-----

```

	Original size	Compressed size	Deduplicated size	
This archive:	6.85 MB	6.85 MB	30.79 kB	<-- !
All archives:	13.69 MB	13.71 MB	6.88 MB	

```

-----

```

	Unique chunks	Total chunks
Chunk index:	167	330

```

-----
```

For a graphical frontend refer to our complementary project [BorgWeb](#)<sup>7</sup>.

## 1.2 Helping, Donations and Bounties

Your help is always welcome! Spread the word, give feedback, help with documentation, testing or development.

You can also give monetary support to the project, see there for details:

<https://borgbackup.readthedocs.io/en/stable/support.html#bounties-and-fundraisers>

<sup>6</sup> <https://borgbackup.readthedocs.org/en/stable/authors.html#license>

<sup>7</sup> <https://borgweb.readthedocs.io/>

## 1.3 Links

- [Main Web Site](#)<sup>8</sup>
- [Releases](#)<sup>9</sup>, [PyPI packages](#)<sup>10</sup> and [ChangeLog](#)<sup>11</sup>
- [GitHub](#)<sup>12</sup> and [Issue Tracker](#)<sup>13</sup>.
- [Web-Chat \(IRC\)](#)<sup>14</sup> and [Mailing List](#)<sup>15</sup>
- [License](#)<sup>16</sup>
- [Security contact](#)<sup>17</sup>

## 1.4 Compatibility notes

EXPECT THAT WE WILL BREAK COMPATIBILITY REPEATEDLY WHEN MAJOR RELEASE NUMBER CHANGES (like when going from 0.x.y to 1.0.0 or from 1.x.y to 2.0.0).

NOT RELEASED DEVELOPMENT VERSIONS HAVE UNKNOWN COMPATIBILITY PROPERTIES.

THIS IS SOFTWARE IN DEVELOPMENT, DECIDE YOURSELF WHETHER IT FITS YOUR NEEDS.

Security issues should be reported to the [Security contact](#)<sup>18</sup> (or see `docs/support.rst` in the source distribution).

19 20 21 22 23

---

<sup>8</sup> <https://borgbackup.readthedocs.org/>

<sup>9</sup> <https://github.com/borgbackup/borg/releases>

<sup>10</sup> <https://pypi.python.org/pypi/borgbackup>

<sup>11</sup> <https://github.com/borgbackup/borg/blob/master/docs/changes.rst>

<sup>12</sup> <https://github.com/borgbackup/borg>

<sup>13</sup> <https://github.com/borgbackup/borg/issues>

<sup>14</sup> <http://webchat.freenode.net/?randomnick=1&channels=%23borgbackup&uio=MTY9dHJlZSY5PXRydWUa8>

<sup>15</sup> <https://mail.python.org/mailman/listinfo/borgbackup>

<sup>16</sup> <https://borgbackup.readthedocs.org/en/stable/authors.html#license>

<sup>17</sup> <https://borgbackup.readthedocs.org/en/stable/support.html#security-contact>

<sup>18</sup> <https://borgbackup.readthedocs.org/en/stable/support.html#security-contact>

<sup>19</sup> <https://borgbackup.readthedocs.org/en/stable/>

<sup>20</sup> <https://travis-ci.org/borgbackup/borg>

<sup>21</sup> <https://codecov.io/github/borgbackup/borg?branch=master>

<sup>22</sup> <https://bestpractices.coreinfrastructure.org/projects/271>

<sup>23</sup> <https://www.bountysource.com/teams/borgbackup>

There are different ways to install Borg:

- *Distribution Package* - easy and fast if a package is available from your distribution.
- *Standalone Binary* - easy and fast, we provide a ready-to-use binary file that comes bundled with all dependencies.
- *From Source*, either:
  - *Using pip* - installing a source package with pip needs more installation steps and requires all dependencies with development headers and a compiler.
  - *Using git* - for developers and power users who want to have the latest code or use revision control (each release is tagged).

## 2.1 Pre-Installation Considerations

## 2.2 (G)LIBC requirements

Borg uses some filesystem functions from Python's *os* standard library module with *follow\_symlinks=False*. These are implemented since quite a while with the non-symlink-following (g)libc functions like e.g. *lstat* or *lutimes* (not: *stat* or *utimes*).

Some stoneage systems (like RHEL/CentOS 5) and also Python interpreter binaries compiled to be able to run on such systems (like Python installed via Anaconda) might miss these functions and Borg won't be able to work correctly. This issue will be detected early and Borg will abort with a fatal error.

For the Borg binaries, there are additional (g)libc requirements, see below.

## 2.3 Distribution Package

Some distributions might offer a ready-to-use `borgbackup` package which can be installed with the package manager. As Borg is still a young project, such a package might be not available for your system yet.

Distribution	Source	Command
Arch Linux	[community] <sup>24</sup>	<code>pacman -S borg</code>
Debian	Debian packages <sup>25</sup>	<code>apt install borgbackup</code>
Gentoo	ebuild <sup>26</sup>	<code>emerge borgbackup</code>
GNU Guix	GNU Guix <sup>27</sup>	<code>guix package --install borg</code>
Fedora/RHEL	Fedora official repository <sup>28</sup>	<code>dnf install borgbackup</code>
FreeBSD	FreeBSD ports <sup>29</sup>	<code>cd /usr/ports/archivers/py-borgbackup &amp;&amp; make install clean</code>
macOS	Brew cask <sup>30</sup>	<code>brew cask install borgbackup</code>
Mageia	cauldron <sup>31</sup>	<code>urpmi borgbackup</code>
NetBSD	pkgsrc <sup>32</sup>	<code>pkg_add py-borgbackup</code>
NixOS	.nix file <sup>33</sup>	N/A
OpenBSD	<b>‘OpenBSD ports’_</b>	<code>pkg_add borgbackup</code>
OpenIndiana	<b>‘OpenIndiana hipster repository’_</b>	<code>pkg install borg</code>
openSUSE	<b>‘openSUSE official repository’_</b>	<code>zypper in borgbackup</code>
Raspbian	Raspbian testing <sup>34</sup>	<code>apt install borgbackup</code>
Ubuntu	Ubuntu packages <sup>35</sup> , Ubuntu PPA <sup>36</sup>	<code>apt install borgbackup</code>

Please ask package maintainers to build a package or, if you can package / submit it yourself, please help us with that! See #105<sup>37</sup> on github to followup on packaging efforts.

If a package is available, it might be interesting to check its version and compare that to our latest release and review the *Important notes*.

## 2.4 Standalone Binary

<sup>24</sup> <https://www.archlinux.org/packages/?name=borg>

<sup>25</sup> <https://packages.debian.org/search?keywords=borgbackup&searchon=names&exact=1&suite=all&section=all>

<sup>26</sup> <https://packages.gentoo.org/packages/app-backup/borgbackup>

<sup>27</sup> <https://www.gnu.org/software/guix/package-list.html#borg>

<sup>28</sup> <https://apps.fedoraproject.org/packages/borgbackup>

<sup>29</sup> <http://www.freshports.org/archivers/py-borgbackup/>

<sup>30</sup> <https://caskroom.github.io/>

<sup>31</sup> <http://madb.mageia.org/package/show/application/0/release/cauldron/name/borgbackup>

<sup>32</sup> <http://pkgsrc.se/sysutils/py-borgbackup>

<sup>33</sup> <https://github.com/NixOS/nixpkgs/blob/master/pkgs/tools/backup/borg/default.nix>

<sup>34</sup> <http://archive.raspbian.org/raspbian/pool/main/b/borgbackup/>

<sup>35</sup> <http://packages.ubuntu.com/xenial/borgbackup>

<sup>36</sup> <https://launchpad.net/~costamagnagianfranco/+archive/ubuntu/borgbackup>

<sup>37</sup> <https://github.com/borgbackup/borg/issues/105>

**Note:** Releases are signed with an OpenPGP key, see [Security](#) for more instructions.

---

Borg binaries (generated with [pyinstaller](#)<sup>38</sup>) are available on the [releases](#)<sup>39</sup> page for the following platforms:

- **Linux:** glibc >= 2.13 (ok for most supported Linux releases). Maybe older glibc versions also work, if they are compatible to 2.13.
- **Mac OS X:** 10.10 (does not work with older OS X releases)
- **FreeBSD:** 10.2 (unknown whether it works for older releases)

To install such a binary, just drop it into a directory in your `PATH`, make borg readable and executable for its users and then you can run `borg`:

```
sudo cp borg-linux64 /usr/local/bin/borg
sudo chown root:root /usr/local/bin/borg
sudo chmod 755 /usr/local/bin/borg
```

Note that the binary uses `/tmp` to unpack Borg with all dependencies. It will fail if `/tmp` has not enough free space or is mounted with the `noexec` option. You can change the temporary directory by setting the `TEMP` environment variable before running Borg.

If a new version is released, you will have to manually download it and replace the old version using the same steps as shown above.

## 2.5 Features & platforms

Besides regular file and directory structures, Borg can preserve

- Hardlinks (considering all files in the same archive)
- Symlinks (stored as symlink, the symlink is not followed)
- Special files:
  - Character and block device files (restored via `mknod`)
  - FIFOs (“named pipes”)
  - Special file `contents` can be backed up in `--read-special` mode. By default the metadata to create them with `mknod(2)`, `mkfifo(2)` etc. is stored.
- Timestamps in nanosecond precision: `mtime`, `atime`, `ctime`
- Permissions:
  - IDs of owning user and owning group
  - Names of owning user and owning group (if the IDs can be resolved)
  - Unix Mode/Permissions (`u/g/o` permissions, `suid`, `sgid`, `sticky`)

On some platforms additional features are supported:

---

<sup>38</sup> <http://www.pyinstaller.org>

<sup>39</sup> <https://github.com/borgbackup/borg/releases>

Platform	ACLs <sup>42</sup>	xattr <sup>43</sup>	Flags <sup>44</sup>
Linux x86	Yes	Yes	No
Linux PowerPC			
Linux ARM			
Mac OS X	Yes	Yes	Yes (all)
FreeBSD	Yes	Yes	
OpenBSD	n/a	n/a	
NetBSD	n/a	No <sup>2</sup>	
Solaris 11	No <sup>3</sup>		n/a
OpenIndiana			
Windows (cygwin)	No <sup>4</sup>	No	No

Some Distributions (e.g. Debian) run additional tests after each release, these are not reflected here.

Other Unix-like operating systems may work as well, but have not been tested at all.

Note that most of the platform-dependent features also depend on the file system. For example, `ntfs-3g` on Linux isn't able to convey NTFS ACLs.

## 2.6 From Source

### 2.6.1 Dependencies

To install Borg from a source package (including pip), you have to install the following dependencies first:

- Python 3<sup>45</sup> >= 3.4.0, plus development headers. Even though Python 3 is not the default Python version on most systems, it is usually available as an optional install.
- OpenSSL<sup>46</sup> >= 1.0.0, plus development headers.
- `libacl`<sup>47</sup> (that pulls in `libattr`<sup>48</sup> also), both plus development headers.
- `liblz4`<sup>49</sup>, plus development headers.
- some Python dependencies, pip will automatically install them for you
- optionally, the `llfuse`<sup>50</sup> Python package is required if you wish to mount an archive as a FUSE filesystem. See `setup.py` about the version requirements.

If you have troubles finding the right package names, have a look at the distribution specific sections below and also at the Vagrantfile in our repo.

In the following, the steps needed to install the dependencies are listed for a selection of platforms. If your distribution is not covered by these instructions, try to use your package manager to install the dependencies. On FreeBSD, you may need to get a recent enough OpenSSL version from FreeBSD ports.

---

<sup>42</sup> The native access control list mechanism of the OS. This normally limits access to non-native ACLs. For example, NTFS ACLs aren't completely accessible on Linux with `ntfs-3g`.

<sup>43</sup> extended attributes; key-value pairs attached to a file, mainly used by the OS. This includes resource forks on Mac OS X.

<sup>44</sup> aka *BSD flags*.

<sup>2</sup> Feature request #1332

<sup>3</sup> Feature request #1337

<sup>4</sup> Cygwin tries to map NTFS ACLs to permissions with varying degrees of success.

<sup>45</sup> <https://www.python.org/>

<sup>46</sup> <https://www.openssl.org/>

<sup>47</sup> <https://savannah.nongnu.org/projects/acl/>

<sup>48</sup> <https://savannah.nongnu.org/projects/attr/>

<sup>49</sup> <https://github.com/Cyan4973/lz4>

<sup>50</sup> <https://pypi.python.org/pypi/llfuse/>

After you have installed the dependencies, you can proceed with steps outlined under *Using pip*.

## Debian / Ubuntu

Install the dependencies with development headers:

```
sudo apt-get install python3 python3-dev python3-pip python-virtualenv \
libssl-dev openssl \
libacl1-dev libacl1 \
liblz4-dev liblz4-1 \
build-essential
sudo apt-get install libfuse-dev fuse pkg-config      # optional, for FUSE support
```

In case you get complaints about permission denied on `/etc/fuse.conf`: on Ubuntu this means your user is not in the `fuse` group. Add yourself to that group, log out and log in again.

## Fedora / Korora

Install the dependencies with development headers:

```
sudo dnf install python3 python3-devel python3-pip python3-virtualenv
sudo dnf install openssl-devel openssl
sudo dnf install libacl-devel libacl
sudo dnf install lz4-devel
sudo dnf install gcc gcc-c++
sudo dnf install fuse-devel fuse pkgconfig          # optional, for FUSE support
```

## openSUSE Tumbleweed / Leap

Install the dependencies automatically using zypper:

```
sudo zypper source-install --build-deps-only borgbackup
```

Alternatively, you can enumerate all build dependencies in the command line:

```
sudo zypper install python3 python3-devel \
libacl-devel liblz4-devel openssl-devel \
python3-Cython python3-Sphinx python3-msgpack-python \
python3-pytest python3-setuptools python3-setuptools_scm \
python3-sphinx_rtd_theme python3-llfuse gcc gcc-c++
```

## Mac OS X

Assuming you have installed [homebrew](http://brew.sh/)<sup>51</sup>, the following steps will install all the dependencies:

```
brew install python3 lz4 openssl
brew install pkg-config          # optional, for FUSE support
pip3 install virtualenv
```

For FUSE support to mount the backup archives, you need at least version 3.0 of FUSE for OS X, which is available via [github](https://github.com), or via homebrew:

<sup>51</sup> <http://brew.sh/>

```
brew cask install osxfuse
```

## FreeBSD

Listed below are packages you will need to install Borg, its dependencies, and commands to make fuse work for using the mount command.

```
pkg install -y python3 openssl liblz4 fusefs-libs pkgconf
pkg install -y git
python3.4 -m ensurepip # to install pip for Python3
To use the mount command:
echo 'fuse_load="YES"' >> /boot/loader.conf
echo 'vfs.usermount=1' >> /etc/sysctl.conf
kldload fuse
sysctl vfs.usermount=1
```

## Windows 10's Linux Subsystem

---

**Note:** Running under Windows 10's Linux Subsystem is experimental and has not been tested much yet.

---

Just follow the Ubuntu Linux installation steps. You can omit the FUSE stuff, it won't work anyway.

## Cygwin

---

**Note:** Running under Cygwin is experimental and has only been tested with Cygwin (x86-64) v2.5.2. Remote repositories are known broken, local repositories should work.

---

Use the Cygwin installer to install the dependencies:

```
python3 python3-devel python3-setuptools
binutils gcc-g++
libopenssl openssl-devel
liblz4_1 liblz4-devel
git make openssh
```

You can then install pip and virtualenv:

```
easy_install-3.4 pip
pip install virtualenv
```

## 2.6.2 Using pip

[Virtualenv](https://pypi.python.org/pypi/virtualenv/)<sup>52</sup> can be used to build and install Borg without affecting the system Python or requiring root access. Using a virtual environment is optional, but recommended except for the most simple use cases.

---

<sup>52</sup> <https://pypi.python.org/pypi/virtualenv/>



---

**Note:** If you install into a virtual environment, you need to **activate** it first (source borg-env/bin/activate), before running borg. Alternatively, symlink borg-env/bin/borg into some directory that is in your PATH so you can just run borg.

---

This will use pip to install the latest release from PyPi:

```
virtualenv --python=python3 borg-env
source borg-env/bin/activate

# install Borg + Python dependencies into virtualenv
pip install borgbackup
# or alternatively (if you want FUSE support):
pip install borgbackup[fuse]
```

To upgrade Borg to a new version later, run the following after activating your virtual environment:

```
pip install -U borgbackup # or ... borgbackup[fuse]
```

### 2.6.3 Using git

This uses latest, unreleased development code from git. While we try not to break master, there are no guarantees on anything.

```
# get borg from github
git clone https://github.com/borgbackup/borg.git

virtualenv --python=python3 borg-env
source borg-env/bin/activate # always before using!

# install borg + dependencies into virtualenv
pip install sphinx # optional, to build the docs
cd borg
pip install -r requirements.d/development.txt
pip install -r requirements.d/fuse.txt # optional, for FUSE support
pip install -e . # in-place editable mode

# optional: run all the tests, on all supported Python versions
# requires fakeroot, available through your package manager
fakeroot -u tox
```

---

**Note:** As a developer or power user, you always want to use a virtual environment.

---



This chapter will get you started with Borg. The first section presents a simple step by step example that uses Borg to backup data. The next section continues by showing how backups can be automated.

### 3.1 Important note about free space

Before you start creating backups, please make sure that there is *always* a good amount of free space on the filesystem that has your backup repository (and also on `~/cache`). A few GB should suffice for most hard-drive sized repositories. See also *Indexes / Caches memory usage*.

If Borg runs out of disk space, it tries to free as much space as it can while aborting the current operation safely, which allows to free more space by deleting/pruning archives. This mechanism is not bullet-proof though. If you *really* run out of disk space, it can be hard or impossible to free space, because Borg needs free space to operate - even to delete backup archives. There is a `--save-space` option for some commands, but even with that Borg will need free space to operate.

You can use some monitoring process or just include the free space information in your backup log files (you check them regularly anyway, right?).

Also helpful:

- create a big file as a “space reserve”, that you can delete to free space
- if you use LVM: use a LV + a filesystem that you can resize later and have some unallocated PEs you can add to the LV.
- consider using quotas
- use *prune* regularly

### 3.2 A step by step example

1. Before a backup can be made a repository has to be initialized:

```
$ borg init /path/to/repo
```

2. Backup the `~/src` and `~/Documents` directories into an archive called *Monday*:

```
$ borg create /path/to/repo::Monday ~/src ~/Documents
```

3. The next day create a new archive called *Tuesday*:

```
$ borg create -v --stats /path/to/repo::Tuesday ~/src ~/Documents
```

This backup will be a lot quicker and a lot smaller since only new never before seen data is stored. The `--stats` option causes Borg to output statistics about the newly created archive such as the amount of unique data (not shared with other archives):

```
-----
Archive name: Tuesday
Archive fingerprint: ↵
↳bd31004d58f51ea06ff735d2e5ac49376901b21d58035f8fb05dbf866566e3c2
Time (start): Tue, 2016-02-16 18:15:11
Time (end):   Tue, 2016-02-16 18:15:11

Duration: 0.19 seconds
Number of files: 127
-----

```

	Original size	Compressed size	Deduplicated size
This archive:	4.16 MB	4.17 MB	26.78 kB
All archives:	8.33 MB	8.34 MB	4.19 MB

  

```
-----

```

	Unique chunks	Total chunks
Chunk index:	132	261

```
-----
```

4. List all archives in the repository:

```
$ borg list /path/to/repo
Monday                Mon, 2016-02-15 19:14:44
Tuesday              Tue, 2016-02-16 19:15:11
```

5. List the contents of the *Monday* archive:

```
$ borg list /path/to/repo::Monday
drwxr-xr-x user  group          0 Mon, 2016-02-15 18:22:30 home/user/Documents
-rw-r--r-- user  group       7961 Mon, 2016-02-15 18:22:30 home/user/Documents/
↳Important.doc
...
```

6. Restore the *Monday* archive by extracting the files relative to the current directory:

```
$ borg extract /path/to/repo::Monday
```

7. Recover disk space by manually deleting the *Monday* archive:

```
$ borg delete /path/to/repo::Monday
```

---

**Note:** Borg is quiet by default (it works on WARNING log level). Add the `-v` (or `--verbose` or `--info`) option to adjust the log level to INFO and also use options like `--progress` or `--list` to get progress reporting during

command execution.

### 3.3 Automating backups

The following example script backs up `/home` and `/var/www` to a remote server. The script also uses the `borg prune` subcommand to maintain a certain number of old archives:

```
#!/bin/sh
REPOSITORY=username@remoteserver.com:backup

# Setting this, so you won't be asked for your repository passphrase:
export Borg_PASSPHRASE='XYZ10ngandsecurepa_55_phrasea&&123'
# or this to ask an external program to supply the passphrase:
export Borg_PASSCOMMAND='pass show backup'

# Backup all of /home and /var/www except a few
# excluded directories
borg create -v --stats \
  $REPOSITORY::'{hostname}-{now:%Y-%m-%d}' \
  /home \
  /var/www \
  --exclude '/home/*/.cache' \
  --exclude /home/Ben/Music/Justin\ Bieber \
  --exclude '*.pyc'

# Use the `prune` subcommand to maintain 7 daily, 4 weekly and 6 monthly
# archives of THIS machine. The '{hostname}-' prefix is very important to
# limit prune's operation to this machine's archives and not apply to
# other machine's archives also.
borg prune -v --list $REPOSITORY --prefix '{hostname}-' \
  --keep-daily=7 --keep-weekly=4 --keep-monthly=6
```

### 3.4 Pitfalls with shell variables and environment variables

This applies to all environment variables you want borg to see, not just `BORG_PASSPHRASE`. The short explanation is: always `export` your variable, and use single quotes if you're unsure of the details of your shell's expansion behavior. E.g.:

```
export Borg_PASSPHRASE='complicated & long'
```

This is because `export` exposes variables to subprocesses, which borg may be one of. More on `export` can be found in the “ENVIRONMENT” section of the `bash(1)` man page.

Beware of how `sudo` interacts with environment variables. For example, you may be surprised that the following `export` has no effect on your command:

```
export Borg_PASSPHRASE='complicated & long'
sudo ./yourborgwrapper.sh # still prompts for password
```

For more information, refer to the `sudo(8)` man page and `env_keep` in the `sudoers(5)` man page.

**Tip:** To debug what your borg process is actually seeing, find its PID (`ps aux|grep borg`) and then look into `/proc/<PID>/environ`.

---

## 3.5 Backup compression

Default is no compression, but we support different methods with high speed or high compression:

If you have a fast repo storage and you want some compression:

```
$ borg create --compression lz4 /path/to/repo::arch ~
```

If you have a less fast repo storage and you want a bit more compression (N=0..9, 0 means no compression, 9 means high compression):

```
$ borg create --compression zlib,N /path/to/repo::arch ~
```

If you have a very slow repo storage and you want high compression (N=0..9, 0 means low compression, 9 means high compression):

```
$ borg create --compression lzma,N /path/to/repo::arch ~
```

You'll need to experiment a bit to find the best compression for your use case. Keep an eye on CPU load and throughput.

## 3.6 Repository encryption

Repository encryption can be enabled or disabled at repository creation time (the default is enabled, with *repokey* method):

```
$ borg init --encryption=none|repokey|keyfile PATH
```

When repository encryption is enabled all data is encrypted using 256-bit [AES](https://en.wikipedia.org/wiki/Advanced_Encryption_Standard)<sup>53</sup> encryption and the integrity and authenticity is verified using [HMAC-SHA256](https://en.wikipedia.org/wiki/HMAC)<sup>54</sup>.

All data is encrypted on the client before being written to the repository. This means that an attacker who manages to compromise the host containing an encrypted archive will not be able to access any of the data, even while the backup is being made.

Borg supports different methods to store the AES and HMAC keys.

**repokey mode** The key is stored inside the repository (in its “config” file). Use this mode if you trust in your good passphrase giving you enough protection. The repository server never sees the plaintext key.

**keyfile mode** The key is stored on your local disk (in `~/.config/borg/keys/`). Use this mode if you want “passphrase and having-the-key” security.

In both modes, the key is stored in encrypted form and can be only decrypted by providing the correct passphrase.

For automated backups the passphrase can be specified using the `BORG_PASSPHRASE` environment variable.

---

<sup>53</sup> [https://en.wikipedia.org/wiki/Advanced\\_Encryption\\_Standard](https://en.wikipedia.org/wiki/Advanced_Encryption_Standard)

<sup>54</sup> <https://en.wikipedia.org/wiki/HMAC>

---

**Note:** Be careful about how you set that environment, see [this note about password environments](#) for more information.

---

**Warning:** The repository data is totally inaccessible without the key and the key passphrase.

Make a backup copy of the key file (`keyfile` mode) or repo config file (`repokey` mode) and keep it at a safe place, so you still have the key in case it gets corrupted or lost. Also keep your passphrase at a safe place.

You can make backups using `borg key export` subcommand.

If you want to print a backup of your key to paper use the `--paper` option of this command and print the result, or this print [template](#) if you need a version with QR-Code.

A backup inside of the backup that is encrypted with that key/passphrase won't help you with that, of course.

## 3.7 Remote repositories

Borg can initialize and access repositories on remote hosts if the host is accessible using SSH. This is fastest and easiest when Borg is installed on the remote host, in which case the following syntax is used:

```
$ borg init user@hostname:/path/to/repo
```

Note: please see the usage chapter for a full documentation of repo URLs.

Remote operations over SSH can be automated with SSH keys. You can restrict the use of the SSH keypair by prepending a forced command to the SSH public key in the remote server's `authorized_keys` file. This example will start Borg in server mode and limit it to a specific filesystem path:

```
command="borg serve --restrict-to-path /path/to/repo",no-pty,no-agent-forwarding,no-  
↪port-forwarding,no-X11-forwarding,no-user-rc ssh-rsa AAAAB3[...]
```

If it is not possible to install Borg on the remote host, it is still possible to use the remote host to store a repository by mounting the remote filesystem, for example, using `sshfs`:

```
$ sshfs user@hostname:/path/to /path/to  
$ borg init /path/to/repo  
$ fusermount -u /path/to
```

You can also use other remote filesystems in a similar way. Just be careful, not all filesystems out there are really stable and working good enough to be acceptable for backup usage.





Borg consists of a number of commands. Each command accepts a number of arguments and options. The following sections will describe each command in detail.

## 4.1 General

### 4.1.1 Positional Arguments and Options: Order matters

Borg only supports taking options (`-v` and `--list` in the example) to the left or right of all positional arguments (`repo::archive` and `path` in the example), but not in between them:

```
borg extract -v --list repo::archive path # good and preferred
borg extract repo::archive path -v --list # also works
borg extract -v repo::archive path --list # works, but ugly
borg extract repo::archive -v --list path # BAD
```

This is due to a problem in the `argparse` module: <http://bugs.python.org/issue15112>

### 4.1.2 Repository URLs

**Local filesystem** (or locally mounted network filesystem):

`/path/to/repo` - filesystem path to repo directory, absolute path

`path/to/repo` - filesystem path to repo directory, relative path

Also, stuff like `~/path/to/repo` or `~other/path/to/repo` works (this is expanded by your shell).

Note: you may also prepend a `file://` to a filesystem path to get URL style.

**Remote repositories** accessed via ssh `user@host`:

`user@host:/path/to/repo` - remote repo, absolute path

`ssh://user@host:port/path/to/repo` - same, alternative syntax, port can be given

**Remote repositories with relative pathes** can be given using this syntax:

`user@host:path/to/repo` - path relative to current directory

`user@host:~/path/to/repo` - path relative to user's home directory

`user@host:~other/path/to/repo` - path relative to other's home directory

**Note:** giving `user@host:./path/to/repo` or `user@host:~/path/to/repo` or `user@host:~other/path/to/repo` is also supported, but not required here.

**Remote repositories with relative pathes, alternative syntax with port:**

`ssh://user@host:port/./path/to/repo` - path relative to current directory

`ssh://user@host:port~/path/to/repo` - path relative to user's home directory

`ssh://user@host:port~other/path/to/repo` - path relative to other's home directory

If you frequently need the same repo URL, it is a good idea to set the `BORG_REPO` environment variable to set a default for the repo URL:

```
export BORG_REPO='ssh://user@host:port/path/to/repo'
```

Then just leave away the repo URL if only a repo URL is needed and you want to use the default - it will be read from `BORG_REPO` then.

Use `::` syntax to give the repo URL when syntax requires giving a positional argument for the repo (e.g. `borg mount :: /mnt`).

### 4.1.3 Repository / Archive Locations

Many commands want either a repository (just give the repo URL, see above) or an archive location, which is a repo URL followed by `::archive_name`.

Archive names must not contain the `/` (slash) character. For simplicity, maybe also avoid blanks or other characters that have special meaning on the shell or in a filesystem (borg mount will use the archive name as directory name).

If you have set `BORG_REPO` (see above) and an archive location is needed, use `::archive_name` - the repo URL part is then read from `BORG_REPO`.

### 4.1.4 Type of log output

The log level of the builtin logging configuration defaults to `WARNING`. This is because we want Borg to be mostly silent and only output warnings, errors and critical messages.

Log levels: `DEBUG < INFO < WARNING < ERROR < CRITICAL`

Use `--debug` to set `DEBUG` log level - to get debug, info, warning, error and critical level output.

Use `--info` (or `-v` or `--verbose`) to set `INFO` log level - to get info, warning, error and critical level output.

Use `--warning` (default) to set `WARNING` log level - to get warning, error and critical level output.

Use `--error` to set `ERROR` log level - to get error and critical level output.

Use `--critical` to set `CRITICAL` log level - to get critical level output.

While you can set misc. log levels, do not expect that every command will give different output on different log levels - it's just a possibility.

**Warning:** Options `--critical` and `--error` are provided for completeness, their usage is not recommended as you might miss important information.

**Warning:** While some options (like `--stats` or `--list`) will emit more informational messages, you have to use INFO (or lower) log level to make them show up in log output. Use `-v` or a logging configuration.

### 4.1.5 Return codes

Borg can exit with the following return codes (rc):

```
0 = success (logged as INFO)
1 = warning (operation reached its normal end, but there were warnings -
  you should check the log, logged as WARNING)
2 = error (like a fatal error, a local or remote exception, the operation
  did not reach its normal end, logged as ERROR)
128+N = killed by signal N (e.g. 137 == kill -9)
```

If you use `--show-rc`, the return code is also logged at the indicated level as the last log entry.

### 4.1.6 Environment Variables

Borg uses some environment variables for automation:

#### General:

**BORG\_REPO** When set, use the value to give the default repository location. If a command needs an archive parameter, you can abbreviate as `::archive`. If a command needs a repository parameter, you can either leave it away or abbreviate as `::`, if a positional parameter is required.

**BORG\_PASSPHRASE** When set, use the value to answer the passphrase question for encrypted repositories. It is used when a passphrase is needed to access an encrypted repo as well as when a new passphrase should be initially set when initializing an encrypted repo.

**BORG\_PASSCOMMAND** When set, use the standard output of the command (trailing newlines are stripped) to answer the passphrase question for encrypted repositories. It is used when a passphrase is needed to access an encrypted repo as well as when a new passphrase should be initially set when initializing an encrypted repo. If **BORG\_PASSPHRASE** is also set, it takes precedence.

**BORG\_PASSPHRASE\_FD** When set, specifies a file descriptor to read a passphrase from. Programs starting borg may choose to open an anonymous pipe and use it to pass a passphrase. This is safer than passing via **BORG\_PASSPHRASE**, because on some systems (e.g. Linux) environment can be examined by other processes. If **BORG\_PASSPHRASE** or **BORG\_PASSCOMMAND** are also set, they take precedence.

**BORG\_DISPLAY\_PASSPHRASE** When set, use the value to answer the “display the passphrase for verification” question when defining a new passphrase for encrypted repositories.

**BORG\_LOGGING\_CONF** When set, use the given filename as INI<sup>55</sup>-style logging configuration.

**BORG\_RSH** When set, use this command instead of `ssh`. This can be used to specify `ssh` options, such as a custom identity file `ssh -i /path/to/private/key`. See `man ssh` for other options.

**BORG\_REMOTE\_PATH** When set, use the given path as borg executable on the remote (defaults to “borg” if unset). Using `--remote-path PATH` commandline option overrides the environment variable.

<sup>55</sup> <https://docs.python.org/3.4/library/logging.config.html#configuration-file-format>

**BORG\_FILES\_CACHE\_TTL** When set to a numeric value, this determines the maximum “time to live” for the files cache entries (default: 20). The files cache is used to quickly determine whether a file is unchanged. The FAQ explains this more detailed in: *It always chunks all my files, even unchanged ones!*

**TMPDIR** where temporary files are stored (might need a lot of temporary space for some operations)

**Some automatic “answers” (if set, they automatically answer confirmation questions):**

**BORG\_UNKNOWN\_UNENCRYPTED\_REPO\_ACCESS\_IS\_OK=no (or =yes)** For “Warning: Attempting to access a previously unknown unencrypted repository”

**BORG\_RELOCATED\_REPO\_ACCESS\_IS\_OK=no (or =yes)** For “Warning: The repository at location ... was previously located at ...”

**BORG\_CHECK\_I\_KNOW\_WHAT\_I\_AM\_DOING=NO (or =YES)** For “Warning: ‘check –repair’ is an experimental feature that might result in data loss.”

**BORG\_DELETE\_I\_KNOW\_WHAT\_I\_AM\_DOING=NO (or =YES)** For “You requested to completely DELETE the repository *including* all archives it contains:”

Note: answers are case sensitive. setting an invalid answer value might either give the default answer or ask you interactively, depending on whether retries are allowed (they by default are allowed). So please test your scripts interactively before making them a non-interactive script.

**Directories:**

**BORG\_KEYS\_DIR** Default to ‘~/config/borg/keys’. This directory contains keys for encrypted repositories.

**BORG\_SECURITY\_DIR** Default to ‘~/config/borg/security’. This directory is used by Borg to track various pieces of security-related data.

**BORG\_CACHE\_DIR** Default to ‘~/cache/borg’. This directory contains the local cache and might need a lot of space for dealing with big repositories).

**Building:**

**BORG\_OPENSSL\_PREFIX** Adds given OpenSSL header file directory to the default locations (setup.py).

**BORG\_LZ4\_PREFIX** Adds given LZ4 header file directory to the default locations (setup.py).

Please note:

- be very careful when using the “yes” sayers, the warnings with prompt exist for your / your data’s security/safety
- also be very careful when putting your passphrase into a script, make sure it has appropriate file permissions (e.g. mode 600, root:root).

## 4.1.7 Resource Usage

Borg might use a lot of resources depending on the size of the data set it is dealing with.

If one uses Borg in a client/server way (with a ssh: repository), the resource usage occurs in part on the client and in another part on the server.

If one uses Borg as a single process (with a filesystem repo), all the resource usage occurs in that one process, so just add up client + server to get the approximate resource usage.

**CPU client:** borg create: does chunking, hashing, compression, crypto (high CPU usage) chunks cache sync: quite heavy on CPU, doing lots of hashtable operations. borg extract: crypto, decompression (medium to high CPU usage) borg check: similar to extract, but depends on options given. borg prune / borg delete archive: low to medium CPU usage borg delete repo: done on the server It won’t go beyond 100% of 1 core as the code is currently single-threaded. Especially higher zlib and lzma compression levels use significant amounts of CPU cycles. Crypto might be cheap on the CPU (if hardware accelerated) or expensive (if not).

**CPU server:** It usually doesn't need much CPU, it just deals with the key/value store (repository) and uses the repository index for that.

borg check: the repository check computes the checksums of all chunks (medium CPU usage) borg delete repo: low CPU usage

**CPU (only for client/server operation):** When using borg in a client/server way with a `ssh:-type` repo, the ssh processes used for the transport layer will need some CPU on the client and on the server due to the crypto they are doing - esp. if you are pumping big amounts of data.

**Memory (RAM) client:** The chunks index and the files index are read into memory for performance reasons. Might need big amounts of memory (see below). Compression, esp. lzma compression with high levels might need substantial amounts of memory.

**Memory (RAM) server:** The server process will load the repository index into memory. Might need considerable amounts of memory, but less than on the client (see below).

**Chunks index (client only):** Proportional to the amount of data chunks in your repo. Lots of chunks in your repo imply a big chunks index. It is possible to tweak the chunker params (see create options).

**Files index (client only):** Proportional to the amount of files in your last backups. Can be switched off (see create options), but next backup might be much slower if you do. The speed benefit of using the files cache is proportional to file size.

**Repository index (server only):** Proportional to the amount of data chunks in your repo. Lots of chunks in your repo imply a big repository index. It is possible to tweak the chunker params (see create options) to influence the amount of chunks being created.

**Temporary files (client):** Reading data and metadata from a FUSE mounted repository will consume up to the size of all deduplicated, small chunks in the repository. Big chunks won't be locally cached.

**Temporary files (server):** None.

**Cache files (client only):** Contains the chunks index and files index (plus a collection of single- archive chunk indexes which might need huge amounts of disk space, depending on archive count and size - see FAQ about how to reduce).

**Network (only for client/server operation):** If your repository is remote, all deduplicated (and optionally compressed/ encrypted) data of course has to go over the connection (ssh: repo url). If you use a locally mounted network filesystem, additionally some copy operations used for transaction support also go over the connection. If you backup multiple sources to one target repository, additional traffic happens for cache resynchronization.

In case you are interested in more details (like formulas), please see *Internals*.

## 4.1.8 Units

To display quantities, Borg takes care of respecting the usual conventions of scale. Disk sizes are displayed in *decimal*<sup>56</sup>, using powers of ten (so kB means 1000 bytes). For memory usage, *binary prefixes*<sup>57</sup> are used, and are indicated using the *IEC binary prefixes*<sup>58</sup>, using powers of two (so KiB means 1024 bytes).

## 4.1.9 Date and Time

We format date and time conforming to ISO-8601, that is: YYYY-MM-DD and HH:MM:SS (24h clock).

For more information about that, see: <https://xkcd.com/1179/>

---

<sup>56</sup> <https://en.wikipedia.org/wiki/Decimal>

<sup>57</sup> [https://en.wikipedia.org/wiki/Binary\\_prefix](https://en.wikipedia.org/wiki/Binary_prefix)

<sup>58</sup> [https://en.wikipedia.org/wiki/IEC\\_80000-13#Prefixes\\_for\\_binary\\_multiples](https://en.wikipedia.org/wiki/IEC_80000-13#Prefixes_for_binary_multiples)

Unless otherwise noted, we display local date and time. Internally, we store and process date and time as UTC.

## 4.2 borg init

```
usage: borg init [-h] [--critical] [--error] [--warning] [--info] [--debug]
                [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                [--remote-path PATH] [-e {none,keyfile,repokey}] [-a]
                [REPOSITORY]

Initialize an empty repository

positional arguments:
  REPOSITORY            repository to create

optional arguments:
  -h, --help            show this help message and exit
  --critical            work on log level CRITICAL
  --error              work on log level ERROR
  --warning            work on log level WARNING (default)
  --info, -v, --verbose work on log level INFO
  --debug              work on log level DEBUG
  --lock-wait N        wait for the lock, but max. N seconds (default: 1).
  --show-rc            show/log the return code (rc)
  --no-files-cache     do not load/update the file metadata cache used to
                        detect unchanged files
  --umask M            set umask to M (local and remote, default: 0077)
  --remote-path PATH  use PATH as borg executable on the remote (default:
                        "borg")
  -e {none,keyfile,repokey}, --encryption {none,keyfile,repokey}
                        select encryption key mode (default: "repokey")
  -a, --append-only    create an append-only mode repository
```

### 4.2.1 Description

This command initializes an empty repository. A repository is a filesystem directory containing the deduplicated data from zero or more archives. Encryption can be enabled at repository init time.

### 4.2.2 Examples

```
# Local repository (default is to use encryption in repokey mode)
$ borg init /path/to/repo

# Local repository (no encryption)
$ borg init --encryption=none /path/to/repo

# Remote repository (accesses a remote borg via ssh)
$ borg init user@hostname:backup

# Remote repository (store the key your home dir)
$ borg init --encryption=keyfile user@hostname:backup
```

Important notes about encryption:

It is not recommended to disable encryption. Repository encryption protects you e.g. against the case that an attacker has access to your backup repository.

But be careful with the key / the passphrase:

If you want “passphrase-only” security, use the `repokey` mode. The key will be stored inside the repository (in its “config” file). In above mentioned attack scenario, the attacker will have the key (but not the passphrase).

If you want “passphrase and having-the-key” security, use the `keyfile` mode. The key will be stored in your home directory (in `.config/borg/keys`). In the attack scenario, the attacker who has just access to your repo won’t have the key (and also not the passphrase).

Make a backup copy of the key file (`keyfile` mode) or repo config file (`repokey` mode) and keep it at a safe place, so you still have the key in case it gets corrupted or lost. Also keep the passphrase at a safe place. The backup that is encrypted with that key won’t help you with that, of course.

Make sure you use a good passphrase. Not too short, not too simple. The real encryption / decryption key is encrypted with / locked by your passphrase. If an attacker gets your key, he can’t unlock and use it without knowing the passphrase.

Be careful with special or non-ascii characters in your passphrase:

- Borg processes the passphrase as unicode (and encodes it as utf-8), so it does not have problems dealing with even the strangest characters.
- BUT: that does not necessarily apply to your OS / VM / keyboard configuration.

So better use a long passphrase made from simple ascii chars than one that includes non-ascii stuff or characters that are hard/impossible to enter on a different keyboard layout.

You can change your passphrase for existing repos at any time, it won’t affect the encryption/decryption key or other secrets.

## 4.3 borg create

```
usage: borg create [-h] [--critical] [--error] [--warning] [--info] [--debug]
                 [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                 [--remote-path PATH] [-s] [-p] [--list]
                 [--filter STATUSCHARS] [-e PATTERN]
                 [--exclude-from EXCLUDEFILE] [--exclude-caches]
                 [--exclude-if-present FILENAME] [--keep-tag-files]
                 [-c SECONDS] [-x] [--numeric-owner] [--noatime] [--noctime]
                 [--timestamp yyyy-mm-ddThh:mm:ss]
                 [--chunker-params CHUNK_MIN_EXP,CHUNK_MAX_EXP,HASH_MASK_BITS,HASH_
↪WINDOW_SIZE]
                 [--ignore-inode] [-C COMPRESSION] [--read-special] [-n]
                 ARCHIVE PATH [PATH ...]
```

Create new archive

positional arguments:

ARCHIVE	name of archive to create (must be also a valid directory name)
PATH	paths to archive

optional arguments:

-h, --help	show this help message and exit
------------	---------------------------------

(continues on next page)

(continued from previous page)

```

--critical          work on log level CRITICAL
--error            work on log level ERROR
--warning          work on log level WARNING (default)
--info, -v, --verbose
                  work on log level INFO
--debug           work on log level DEBUG
--lock-wait N      wait for the lock, but max. N seconds (default: 1).
--show-rc          show/log the return code (rc)
--no-files-cache   do not load/update the file metadata cache used to
                  detect unchanged files
--umask M          set umask to M (local and remote, default: 0077)
--remote-path PATH use PATH as borg executable on the remote (default:
                  "borg")
-s, --stats        print statistics for the created archive. Requires
                  -v/--verbose.
-p, --progress     show progress display while creating the archive,
                  showing Original, Compressed and Deduplicated sizes,
                  followed by the Number of files seen and the path
                  being processed, default: False
--list            output verbose list of items (files, dirs, ...).
                  Requires -v/--verbose.
--filter STATUSCHARS only display items with the given status characters
-e PATTERN, --exclude PATTERN
                  exclude paths matching PATTERN
--exclude-from EXCLUDEFILE
                  read exclude patterns from EXCLUDEFILE, one per line
--exclude-caches   exclude directories that contain a CACHEDIR.TAG file
                  (http://www.brynosaurus.com/cachedir/spec.html)
--exclude-if-present FILENAME
                  exclude directories that contain the specified file
--keep-tag-files   keep tag files of excluded caches/directories
-c SECONDS, --checkpoint-interval SECONDS
                  write checkpoint every SECONDS seconds (Default: 1800)
-x, --one-file-system
                  stay in same file system
--numeric-owner    only store numeric user and group identifiers
--noatime          do not store atime into archive
--noctime          do not store ctime into archive
--timestamp yyyy-mm-ddThh:mm:ss
                  manually specify the archive creation date/time (UTC).
                  alternatively, give a reference file/directory.
--chunker-params CHUNK_MIN_EXP,CHUNK_MAX_EXP,HASH_MASK_BITS,HASH_WINDOW_SIZE
                  specify the chunker parameters. default: 19,23,21,4095
--ignore-inode     ignore inode data in the file metadata cache used to
                  detect unchanged files.
-C COMPRESSION, --compression COMPRESSION
                  select compression algorithm (and level): none == no
                  compression (default), lz4 == lz4, zlib == zlib
                  (default level 6), zlib,0 .. zlib,9 == zlib (with
                  level 0..9), lzma == lzma (default level 6), lzma,0 ..
                  lzma,9 == lzma (with level 0..9).
--read-special     open and read block and char device files as well as
                  FIFOs as if they were regular files. Also follows
                  symlinks pointing to these kinds of files.
-n, --dry-run      do not create a backup archive

```



### 4.3.1 Description

This command creates a backup archive containing all files found while recursively traversing all paths specified. Paths are added to the archive as they are given, that means if relative paths are desired, the command has to be run from the correct directory.

When giving ‘-’ as path, borg will read data from standard input and create a file ‘stdin’ in the created archive from that data.

The archive will consume almost no disk space for files or parts of files that have already been stored in other archives.

The archive name needs to be unique. It must not end in ‘.checkpoint’ or ‘.checkpoint.N’ (with N being a number), because these names are used for checkpoints and treated in special ways.

In the archive name, you may use the following placeholders: {now}, {utcnow}, {fqdn}, {hostname}, {user} and some others.

To speed up pulling backups over sshfs and similar network file systems which do not provide correct inode information the `--ignore-inode` flag can be used. This potentially decreases reliability of change detection, while avoiding always reading all files on these file systems.

The mount points of filesystems or filesystem snapshots should be the same for every creation of a new archive to ensure fast operation. This is because the file cache that is used to determine changed files quickly uses absolute filenames. If this is not possible, consider creating a bind mount to a stable location.

See the output of the “borg help patterns” command for more help on exclude patterns. See the output of the “borg help placeholders” command for more help on placeholders.

Borg respects the `nodump` flag. Files flagged `nodump` will be marked as excluded (x) in `--list` output.

### 4.3.2 Examples

```
# Backup ~/Documents into an archive named "my-documents"
$ borg create /path/to/repo::my-documents ~/Documents

# same, but verbosely list all files as we process them
$ borg create -v --list /path/to/repo::my-documents ~/Documents

# Backup ~/Documents and ~/src but exclude pyc files
$ borg create /path/to/repo::my-files \
  ~/Documents \
  ~/src \
  --exclude '*.pyc'

# Backup home directories excluding image thumbnails (i.e. only
# /home/*/.thumbnails is excluded, not /home/*/*/.thumbnails)
$ borg create /path/to/repo::my-files /home \
  --exclude 're:^(/home/[^/]+/\.thumbnails/'

# Do the same using a shell-style pattern
$ borg create /path/to/repo::my-files /home \
  --exclude 'sh:/home/*/.thumbnails'

# Backup the root filesystem into an archive named "root-YYYY-MM-DD"
# use zlib compression (good, but slow) - default is no compression
$ borg create -C zlib,6 --one-file-system /path/to/repo::root-{now:%Y-%m-%d} /

# Make a big effort in fine granular deduplication (big chunk management
```

(continues on next page)

(continued from previous page)

```

# overhead, needs a lot of RAM and disk space, see formula in internals
# docs - same parameters as borg < 1.0 or attic):
$ borg create --chunker-params 10,23,16,4095 /path/to/repo::small /smallstuff

# Backup a raw device (must not be active/in use/mounted at that time)
$ dd if=/dev/sdx bs=10M | borg create /path/to/repo::my-sdx -

# No compression (default)
$ borg create /path/to/repo::arch ~

# Super fast, low compression
$ borg create --compression lz4 /path/to/repo::arch ~

# Less fast, higher compression (N = 0..9)
$ borg create --compression zlib,N /path/to/repo::arch ~

# Even slower, even higher compression (N = 0..9)
$ borg create --compression lzma,N /path/to/repo::arch ~

# Use short hostname, user name and current time in archive name
$ borg create /path/to/repo::{hostname}-{user}-{now} ~
# Similar, use the same datetime format as borg 1.1 will have as default
$ borg create /path/to/repo::{hostname}-{user}-{now:%Y-%m-%dT%H:%M:%S} ~
# As above, but add nanoseconds
$ borg create /path/to/repo::{hostname}-{user}-{now:%Y-%m-%dT%H:%M:%S.%f} ~

# Backing up relative paths by moving into the correct directory first
$ cd /home/user/Documents
# The root directory of the archive will be "projectA"
$ borg create /path/to/repo::daily-projectA-{now:%Y-%m-%d} projectA

```

### 4.3.3 Notes

- the `--exclude` patterns are not like tar. In tar `--exclude .bundler/gems` will exclude `foo/.bundler/gems`. In borg it will not, you need to use `--exclude '*/.bundler/gems'` to get the same effect. See `borg help patterns` for more information.

## 4.4 borg extract

```

usage: borg extract [-h] [--critical] [--error] [--warning] [--info] [--debug]
                  [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                  [--remote-path PATH] [--list] [-n] [-e PATTERN]
                  [--exclude-from EXCLUDEFILE] [--numeric-owner]
                  [--strip-components NUMBER] [--stdout] [--sparse]
                  ARCHIVE [PATH [PATH ...]]

```

Extract archive contents

positional arguments:

```

  ARCHIVE      archive to extract
  PATH         paths to extract; patterns are supported

```

(continues on next page)

(continued from previous page)

```

optional arguments:
  -h, --help                show this help message and exit
  --critical                work on log level CRITICAL
  --error                   work on log level ERROR
  --warning                 work on log level WARNING (default)
  --info, -v, --verbose    work on log level INFO
  --debug                  work on log level DEBUG
  --lock-wait N            wait for the lock, but max. N seconds (default: 1).
  --show-rc                show/log the return code (rc)
  --no-files-cache         do not load/update the file metadata cache used to
                           detect unchanged files
  --umask M                set umask to M (local and remote, default: 0077)
  --remote-path PATH       use PATH as borg executable on the remote (default:
                           "borg")
  --list                    output verbose list of items (files, dirs, ...).
                           Requires -v/--verbose.
  -n, --dry-run            do not actually change any files
  -e PATTERN, --exclude PATTERN
                           exclude paths matching PATTERN
  --exclude-from EXCLUDEFILE
                           read exclude patterns from EXCLUDEFILE, one per line
                           only obey numeric user and group identifiers
  --strip-components NUMBER
                           Remove the specified number of leading path elements.
                           Pathnames with fewer elements will be silently
                           skipped.
  --stdout                 write all extracted data to stdout
  --sparse                  create holes in output sparse file from all-zero
                           chunks

```

#### 4.4.1 Description

This command extracts the contents of an archive. By default the entire archive is extracted but a subset of files and directories can be selected by passing a list of PATHs as arguments. The file selection can further be restricted by using the `--exclude` option.

See the output of the “`borg help patterns`” command for more help on exclude patterns.

#### 4.4.2 Examples

```

# Extract entire archive
$ borg extract /path/to/repo::my-files

# Extract entire archive and list files while processing
$ borg extract -v --list /path/to/repo::my-files

# Extract the "src" directory
$ borg extract /path/to/repo::my-files home/USERNAME/src

# Extract the "src" directory but exclude object files
$ borg extract /path/to/repo::my-files home/USERNAME/src --exclude '*.o'

```

(continues on next page)

(continued from previous page)

```
# Restore a raw device (must not be active/in use/mounted at that time)
$ borg extract --stdout /path/to/repo::my-sdx | dd of=/dev/sdx bs=10M
```

**Note:** currently, `extract` always writes into the current working directory (“.”), so make sure you `cd` to the right place before calling `borg extract`.

## 4.5 borg check

```
usage: borg check [-h] [--critical] [--error] [--warning] [--info] [--debug]
                 [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                 [--remote-path PATH] [--repository-only] [--archives-only]
                 [--repair] [--save-space] [--last N] [-P PREFIX]
                 [REPOSITORY_OR_ARCHIVE]

Check repository consistency

positional arguments:
  REPOSITORY_OR_ARCHIVE
                        repository or archive to check consistency of

optional arguments:
  -h, --help            show this help message and exit
  --critical            work on log level CRITICAL
  --error              work on log level ERROR
  --warning            work on log level WARNING (default)
  --info, -v, --verbose
                        work on log level INFO
  --debug              work on log level DEBUG
  --lock-wait N        wait for the lock, but max. N seconds (default: 1).
  --show-rc            show/log the return code (rc)
  --no-files-cache    do not load/update the file metadata cache used to
                        detect unchanged files
  --umask M            set umask to M (local and remote, default: 0077)
  --remote-path PATH  use PATH as borg executable on the remote (default:
                        "borg")
  --repository-only   only perform repository checks
  --archives-only     only perform archives checks
  --repair            attempt to repair any inconsistencies found
  --save-space        work slower, but using less space
  --last N            only check last N archives (Default: all)
  -P PREFIX, --prefix PREFIX
                        only consider archive names starting with this prefix
```

### 4.5.1 Description

The check command verifies the consistency of a repository and the corresponding archives.

First, the underlying repository data files are checked:

- For all segments the segment magic (header) is checked
- For all objects stored in the segments, all metadata (e.g. crc and size) and all data is read. The read data is checked by size and CRC. Bit rot and other types of accidental damage can be detected this way.

- If we are in repair mode and a integrity error is detected for a segment, we try to recover as many objects from the segment as possible.
- In repair mode, it makes sure that the index is consistent with the data stored in the segments.
- If you use a remote repo server via ssh:, the repo check is executed on the repo server without causing significant network traffic.
- The repository check can be skipped using the `--archives-only` option.

Second, the consistency and correctness of the archive metadata is verified:

- Is the repo manifest present? If not, it is rebuilt from archive metadata chunks (this requires reading and decrypting of all metadata and data).
- Check if archive metadata chunk is present. if not, remove archive from manifest.
- For all files (items) in the archive, for all chunks referenced by these files, check if chunk is present. If a chunk is not present and we are in repair mode, replace it with a same-size replacement chunk of zeros. If a previously lost chunk reappears (e.g. via a later backup) and we are in repair mode, the all-zero replacement chunk will be replaced by the correct chunk. This requires reading of archive and file metadata, but not data.
- If we are in repair mode and we checked all the archives: delete orphaned chunks from the repo.
- if you use a remote repo server via ssh:, the archive check is executed on the client machine (because if encryption is enabled, the checks will require decryption and this is always done client-side, because key access will be required).
- The archive checks can be time consuming, they can be skipped using the `--repository-only` option.

## 4.6 borg rename

```
usage: borg rename [-h] [--critical] [--error] [--warning] [--info] [--debug]
                  [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                  [--remote-path PATH]
                  ARCHIVE NEWNAME
```

Rename an existing archive

positional arguments:

```
  ARCHIVE      archive to rename
  NEWNAME      the new archive name to use
```

optional arguments:

```
-h, --help            show this help message and exit
--critical            work on log level CRITICAL
--error              work on log level ERROR
--warning            work on log level WARNING (default)
--info, -v, --verbose
                    work on log level INFO
--debug              work on log level DEBUG
--lock-wait N        wait for the lock, but max. N seconds (default: 1).
--show-rc            show/log the return code (rc)
--no-files-cache     do not load/update the file metadata cache used to
                    detect unchanged files
--umask M            set umask to M (local and remote, default: 0077)
--remote-path PATH  use PATH as borg executable on the remote (default:
                    "borg")
```

## 4.6.1 Description

This command renames an archive in the repository.

## 4.6.2 Examples

```
$ borg create /path/to/repo::archivename ~
$ borg list /path/to/repo
archivename                               Mon, 2016-02-15 19:50:19

$ borg rename /path/to/repo::archivename newname
$ borg list /path/to/repo
newname                                    Mon, 2016-02-15 19:50:19
```

## 4.7 borg list

```
usage: borg list [-h] [--critical] [--error] [--warning] [--info] [--debug]
                [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                [--remote-path PATH] [--short] [--list-format LISTFORMAT]
                [-P PREFIX]
                [REPOSITORY_OR_ARCHIVE]

List archive or repository contents

positional arguments:
  REPOSITORY_OR_ARCHIVE
                        repository/archive to list contents of

optional arguments:
  -h, --help            show this help message and exit
  --critical            work on log level CRITICAL
  --error              work on log level ERROR
  --warning            work on log level WARNING (default)
  --info, -v, --verbose
                        work on log level INFO
  --debug             work on log level DEBUG
  --lock-wait N        wait for the lock, but max. N seconds (default: 1).
  --show-rc           show/log the return code (rc)
  --no-files-cache    do not load/update the file metadata cache used to
                        detect unchanged files
  --umask M           set umask to M (local and remote, default: 0077)
  --remote-path PATH  use PATH as borg executable on the remote (default:
                        "borg")
  --short             only print file/directory names, nothing else
  --list-format LISTFORMAT
                        specify format for archive file listing (default:
                        "{mode} {user:6} {group:6} {size:8d} {isotime}
                        {path}{extra}{NEWLINE}") Special "{formatkeys}" exists
                        to list available keys
  -P PREFIX, --prefix PREFIX
                        only consider archive names starting with this prefix
```

## 4.7.1 Description

This command lists the contents of a repository or an archive.

## 4.7.2 Examples

```
$ borg list /path/to/repo
Monday                               Mon, 2016-02-15 19:15:11
repo                                  Mon, 2016-02-15 19:26:54
root-2016-02-15                       Mon, 2016-02-15 19:36:29
newname                               Mon, 2016-02-15 19:50:19
...

$ borg list /path/to/repo::root-2016-02-15
drwxr-xr-x root root          0 Mon, 2016-02-15 17:44:27 .
drwxrwxr-x root root          0 Mon, 2016-02-15 19:04:49 bin
-rwxr-xr-x root root 1029624 Thu, 2014-11-13 00:08:51 bin/bash
lrwxrwxrwx root root          0 Fri, 2015-03-27 20:24:26 bin/bzcmp -> bzdiff
-rwxr-xr-x root root    2140 Fri, 2015-03-27 20:24:22 bin/bzdiff
...

$ borg list /path/to/repo::archiveA --list-format="{mode} {user:6} {group:6} {size:8d}
↳ {isomtime} {path}{extra}{NEWLINE}"
drwxrwxr-x user user          0 Sun, 2015-02-01 11:00:00 .
drwxrwxr-x user user          0 Sun, 2015-02-01 11:00:00 code
drwxrwxr-x user user          0 Sun, 2015-02-01 11:00:00 code/myproject
-rw-rw-r-- user user 1416192 Sun, 2015-02-01 11:00:00 code/myproject/file.ext
...

# see what is changed between archives, based on file modification time, size and
↳ file path
$ borg list /path/to/repo::archiveA --list-format="{mtime:%s}{TAB}{size}{TAB}{path}
↳ {LF}" |sort -n > /tmp/list.archiveA
$ borg list /path/to/repo::archiveB --list-format="{mtime:%s}{TAB}{size}{TAB}{path}
↳ {LF}" |sort -n > /tmp/list.archiveB
$ diff -y /tmp/list.archiveA /tmp/list.archiveB
1422781200      0      .                1422781200      0      ↳
↳ .
1422781200      0      code                1422781200      0      ↳
↳ code
1422781200      0      code/myproject      1422781200      0      ↳
↳ code/myproject
1422781200      1416192 code/myproject/file.ext | 1454664653      ↳
↳1416192 code/myproject/file.ext
...
```

## 4.8 borg delete

```
usage: borg delete [-h] [--critical] [--error] [--warning] [--info] [--debug]
                  [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                  [--remote-path PATH] [-p] [-s] [-c] [--force]
                  [--save-space]
                  [TARGET]
```

(continues on next page)

(continued from previous page)

```

Delete an existing repository or archive

positional arguments:
  TARGET                archive or repository to delete

optional arguments:
  -h, --help            show this help message and exit
  --critical            work on log level CRITICAL
  --error              work on log level ERROR
  --warning            work on log level WARNING (default)
  --info, -v, --verbose
                        work on log level INFO
  --debug              work on log level DEBUG
  --lock-wait N        wait for the lock, but max. N seconds (default: 1).
  --show-rc            show/log the return code (rc)
  --no-files-cache     do not load/update the file metadata cache used to
                        detect unchanged files
  --umask M            set umask to M (local and remote, default: 0077)
  --remote-path PATH   use PATH as borg executable on the remote (default:
                        "borg")
  -p, --progress       show progress display while deleting a single archive
  -s, --stats          print statistics for the deleted archive. Requires
                        -v/--verbose.
  -c, --cache-only     delete only the local cache for the given repository
  --force              force deletion of corrupted archives, use --force
                        --force in case --force does not work.
  --save-space         work slower, but using less space

```

### 4.8.1 Description

This command deletes an archive from the repository or the complete repository. Disk space is reclaimed accordingly. If you delete the complete repository, the local cache for it (if any) is also deleted.

### 4.8.2 Examples

```

# delete a single backup archive:
$ borg delete /path/to/repo::Monday

# delete the whole repository and the related local cache:
$ borg delete /path/to/repo
You requested to completely DELETE the repository *including* all archives it
↳contains:
repo                Mon, 2016-02-15 19:26:54
root-2016-02-15    Mon, 2016-02-15 19:36:29
newname            Mon, 2016-02-15 19:50:19
Type 'YES' if you understand this and want to continue: YES

```



## 4.9 borg prune

```
usage: borg prune [-h] [--critical] [--error] [--warning] [--info] [--debug]
                 [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                 [--remote-path PATH] [-n] [--force] [-s] [--list]
                 [--keep-within INTERVAL] [-H HOURLY] [-d DAILY] [-w WEEKLY]
                 [-m MONTHLY] [-y YEARLY] [-P PREFIX] [--save-space]
                 [REPOSITORY]
```

Prune repository archives according to specified rules

positional arguments:

REPOSITORY repository to prune

optional arguments:

```
-h, --help show this help message and exit
--critical work on log level CRITICAL
--error work on log level ERROR
--warning work on log level WARNING (default)
--info, -v, --verbose work on log level INFO
--debug work on log level DEBUG
--lock-wait N wait for the lock, but max. N seconds (default: 1).
--show-rc show/log the return code (rc)
--no-files-cache do not load/update the file metadata cache used to
detect unchanged files
--umask M set umask to M (local and remote, default: 0077)
--remote-path PATH use PATH as borg executable on the remote (default:
"borg")
-n, --dry-run do not change repository
--force force pruning of corrupted archives
-s, --stats print statistics for the deleted archive. Requires
-v/--verbose.
--list output verbose list of archives it keeps/prunes.
Requires -v/--verbose.
--keep-within INTERVAL keep all archives within this time interval
-H HOURLY, --keep-hourly HOURLY number of hourly archives to keep
-d DAILY, --keep-daily DAILY number of daily archives to keep
-w WEEKLY, --keep-weekly WEEKLY number of weekly archives to keep
-m MONTHLY, --keep-monthly MONTHLY number of monthly archives to keep
-y YEARLY, --keep-yearly YEARLY number of yearly archives to keep
-P PREFIX, --prefix PREFIX only consider archive names starting with this prefix
--save-space work slower, but using less space
```

### 4.9.1 Description

The prune command prunes a repository by deleting all archives not matching any of the specified retention options. This command is normally used by automated backup scripts wanting to keep a certain number of historic backups.

As an example, “-d 7” means to keep the latest backup on each day, up to 7 most recent days with backups (days without backups do not count). The rules are applied from hourly to yearly, and backups selected by previous rules do not count towards those of later rules. The time that each backup starts is used for pruning purposes. Dates and times are interpreted in the local timezone, and weeks go from Monday to Sunday. Specifying a negative number of archives to keep means that there is no limit.

The “-keep-within” option takes an argument of the form “<int><char>”, where char is “H”, “d”, “w”, “m”, “y”. For example, “-keep-within 2d” means to keep all archives that were created within the past 48 hours. “1m” is taken to mean “31d”. The archives kept with this option do not count towards the totals specified by any other options.

If a prefix is set with -P, then only archives that start with the prefix are considered for deletion and only those archives count towards the totals specified by the rules. Otherwise, *all* archives in the repository are candidates for deletion! There is no automatic distinction between archives representing different contents. These need to be distinguished by specifying matching prefixes.

## 4.9.2 Examples

Be careful, prune is a potentially dangerous command, it will remove backup archives.

The default of prune is to apply to **all archives in the repository** unless you restrict its operation to a subset of the archives using --prefix. When using --prefix, be careful to choose a good prefix - e.g. do not use a prefix “foo” if you do not also want to match “foobar”.

It is strongly recommended to always run `prune -v --list --dry-run ...` first so you will see what it would do without it actually doing anything.

There is also a visualized prune example in `docs/misc/prune-example.txt`.

```
# Keep 7 end of day and 4 additional end of week archives.
# Do a dry-run without actually deleting anything.
$ borg prune -v --list --dry-run --keep-daily=7 --keep-weekly=4 /path/to/repo

# Same as above but only apply to archive names starting with the hostname
# of the machine followed by a "-" character:
$ borg prune -v --list --keep-daily=7 --keep-weekly=4 --prefix='{hostname}-' /path/to/
↳repo

# Keep 7 end of day, 4 additional end of week archives,
# and an end of month archive for every month:
$ borg prune -v --list --keep-daily=7 --keep-weekly=4 --keep-monthly=-1 /path/to/repo

# Keep all backups in the last 10 days, 4 additional end of week archives,
# and an end of month archive for every month:
$ borg prune -v --list --keep-within=10d --keep-weekly=4 --keep-monthly=-1 /path/to/
↳repo
```

## 4.10 borg info

```
usage: borg info [-h] [--critical] [--error] [--warning] [--info] [--debug]
                [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                [--remote-path PATH]
                ARCHIVE
```

Show archive details such as disk space used

(continues on next page)

(continued from previous page)

```
positional arguments:
  ARCHIVE                archive to display information about

optional arguments:
  -h, --help            show this help message and exit
  --critical            work on log level CRITICAL
  --error              work on log level ERROR
  --warning            work on log level WARNING (default)
  --info, -v, --verbose
                       work on log level INFO
  --debug              work on log level DEBUG
  --lock-wait N        wait for the lock, but max. N seconds (default: 1).
  --show-rc            show/log the return code (rc)
  --no-files-cache     do not load/update the file metadata cache used to
                       detect unchanged files
  --umask M            set umask to M (local and remote, default: 0077)
  --remote-path PATH   use PATH as borg executable on the remote (default:
                       "borg")
```

### 4.10.1 Description

This command displays some detailed information about the specified archive.

Please note that the deduplicated sizes of the individual archives do not add up to the deduplicated size of the repository (“all archives”), because the two are meaning different things:

**This archive / deduplicated size = amount of data stored ONLY for this archive = unique chunks of this archive.**

**All archives / deduplicated size = amount of data stored in the repo = all chunks in the repository.**

### 4.10.2 Examples

```
$ borg info /path/to/repo::root-2016-02-15
Name: root-2016-02-15
Fingerprint: 57c827621f21b000a8d363c1e163cc55983822b3afff3a96df595077a660be50
Hostname: myhostname
Username: root
Time (start): Mon, 2016-02-15 19:36:29
Time (end):   Mon, 2016-02-15 19:39:26
Command line: /usr/local/bin/borg create -v --list -C zlib,6 /path/to/repo::root-2016-
↳ 02-15 / --one-file-system
Number of files: 38100

                Original size      Compressed size      Deduplicated size
This archive:   1.33 GB              613.25 MB            571.64 MB
All archives:   1.63 GB              853.66 MB            584.12 MB

                Unique chunks        Total chunks
Chunk index:    36858                48844
```

## 4.11 borg mount

```
usage: borg mount [-h] [--critical] [--error] [--warning] [--info] [--debug]
                 [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                 [--remote-path PATH] [-f] [-o OPTIONS]
                 REPOSITORY_OR_ARCHIVE MOUNTPOINT

Mount archive or an entire repository as a FUSE filesystem

positional arguments:
  REPOSITORY_OR_ARCHIVE  repository/archive to mount
  MOUNTPOINT             where to mount filesystem

optional arguments:
  -h, --help            show this help message and exit
  --critical            work on log level CRITICAL
  --error              work on log level ERROR
  --warning            work on log level WARNING (default)
  --info, -v, --verbose
                       work on log level INFO
  --debug              work on log level DEBUG
  --lock-wait N        wait for the lock, but max. N seconds (default: 1).
  --show-rc            show/log the return code (rc)
  --no-files-cache     do not load/update the file metadata cache used to
                       detect unchanged files
  --umask M            set umask to M (local and remote, default: 0077)
  --remote-path PATH   use PATH as borg executable on the remote (default:
                       "borg")
  -f, --foreground     stay in foreground, do not daemonize
  -o OPTIONS           Extra mount options
```

### 4.11.1 Description

This command mounts an archive as a FUSE filesystem. This can be useful for browsing an archive or restoring individual files. Unless the `--foreground` option is given the command will run in the background until the filesystem is unmounted.

The `BORG_MOUNT_DATA_CACHE_ENTRIES` environment variable is meant for advanced users to tweak the performance. It sets the number of cached data chunks; additional memory usage can be up to ~8 MiB times this number. The default is the number of CPU cores.

For mount options, see the `fuse(8)` manual page. Additional mount options supported by borg:

- `allow_damaged_files`: by default damaged files (where missing chunks were replaced with runs of zeros by borg check `-repair`) are not readable and return EIO (I/O error). Set this option to read such files.

When the daemonized process receives a signal or crashes, it does not unmount. Unmounting in these cases could cause an active `rsync` or similar process to unintentionally delete data.

When running in the foreground `^C/SIGINT` unmounts cleanly, but other signals or crashes do not.

## 4.12 borg umount

```
usage: borg umount [-h] [--critical] [--error] [--warning] [--info] [--debug]
                  [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                  [--remote-path PATH]
                  MOUNTPOINT

un-mount the FUSE filesystem

positional arguments:
  MOUNTPOINT          mountpoint of the filesystem to umount

optional arguments:
  -h, --help          show this help message and exit
  --critical          work on log level CRITICAL
  --error            work on log level ERROR
  --warning          work on log level WARNING (default)
  --info, -v, --verbose
                    work on log level INFO
  --debug            work on log level DEBUG
  --lock-wait N      wait for the lock, but max. N seconds (default: 1).
  --show-rc          show/log the return code (rc)
  --no-files-cache  do not load/update the file metadata cache used to
                    detect unchanged files
  --umask M          set umask to M (local and remote, default: 0077)
  --remote-path PATH use PATH as borg executable on the remote (default:
                    "borg")
```

### 4.12.1 Description

This command un-mounts a FUSE filesystem that was mounted with `borg mount`.

This is a convenience wrapper that just calls the platform-specific shell command - usually this is either `umount` or `fusermount -u`.

### 4.12.2 Examples

```
$ borg mount /path/to/repo::root-2016-02-15 /tmp/mymountpoint
$ ls /tmp/mymountpoint
bin boot etc      home lib lib64 lost+found media mnt  opt  root  sbin  srv  _
↳tmp  usr  var
$ borg umount /tmp/mymountpoint
```

## 4.13 borg key export

```
usage: borg key export [-h] [--critical] [--error] [--warning] [--info]
                      [--debug] [--lock-wait N] [--show-rc]
                      [--no-files-cache] [--umask M] [--remote-path PATH]
                      [--paper] [--qr-html]
                      [REPOSITORY] [PATH]
```

(continues on next page)

(continued from previous page)

```

Export the repository key for backup

positional arguments:
  REPOSITORY
  PATH                  where to store the backup

optional arguments:
  -h, --help            show this help message and exit
  --critical            work on log level CRITICAL
  --error              work on log level ERROR
  --warning            work on log level WARNING (default)
  --info, -v, --verbose
                       work on log level INFO
  --debug              work on log level DEBUG
  --lock-wait N        wait for the lock, but max. N seconds (default: 1).
  --show-rc            show/log the return code (rc)
  --no-files-cache    do not load/update the file metadata cache used to
                       detect unchanged files
  --umask M            set umask to M (local and remote, default: 0077)
  --remote-path PATH  use PATH as borg executable on the remote (default:
                       "borg")
  --paper              Create an export suitable for printing and later type-
                       in
  --qr-html            Create an html file suitable for printing and later
                       type-in or qr scan

```

### 4.13.1 Description

If repository encryption is used, the repository is inaccessible without the key. This command allows to backup this essential key. Note that the backup produced does not include the passphrase itself (i.e. the exported key stays encrypted). In order to regain access to a repository, one needs both the exported key and the original passphrase.

There are two backup formats. The normal backup format is suitable for digital storage as a file. The `--paper` backup format is optimized for printing and typing in while importing, with per line checks to reduce problems with manual input.

For repositories using keyfile encryption the key is saved locally on the system that is capable of doing backups. To guard against loss of this key, the key needs to be backed up independently of the main data backup.

For repositories using the repokey encryption the key is saved in the repository in the config file. A backup is thus not strictly needed, but guards against the repository becoming inaccessible if the file is damaged for some reason.

## 4.14 borg key import

```

usage: borg key import [-h] [--critical] [--error] [--warning] [--info]
                       [--debug] [--lock-wait N] [--show-rc]
                       [--no-files-cache] [--umask M] [--remote-path PATH]
                       [--paper]
                       [REPOSITORY] [PATH]

```

Import the repository key from backup

positional arguments:

(continues on next page)

(continued from previous page)

```

REPOSITORY
PATH          path to the backup

optional arguments:
-h, --help          show this help message and exit
--critical          work on log level CRITICAL
--error            work on log level ERROR
--warning          work on log level WARNING (default)
--info, -v, --verbose
                  work on log level INFO
--debug            work on log level DEBUG
--lock-wait N      wait for the lock, but max. N seconds (default: 1).
--show-rc          show/log the return code (rc)
--no-files-cache  do not load/update the file metadata cache used to
                  detect unchanged files
--umask M          set umask to M (local and remote, default: 0077)
--remote-path PATH use PATH as borg executable on the remote (default:
                  "borg")
--paper            interactively import from a backup done with --paper

```

#### 4.14.1 Description

This command allows to restore a key previously backed up with the export command.

If the `--paper` option is given, the import will be an interactive process in which each line is checked for plausibility before proceeding to the next line. For this format `PATH` must not be given.

### 4.15 borg change-passphrase

```

usage: borg change-passphrase [-h] [--critical] [--error] [--warning] [--info]
                               [--debug] [--lock-wait N] [--show-rc]
                               [--no-files-cache] [--umask M]
                               [--remote-path PATH]
                               [REPOSITORY]

Change repository key file passphrase

positional arguments:
  REPOSITORY

optional arguments:
-h, --help          show this help message and exit
--critical          work on log level CRITICAL
--error            work on log level ERROR
--warning          work on log level WARNING (default)
--info, -v, --verbose
                  work on log level INFO
--debug            work on log level DEBUG
--lock-wait N      wait for the lock, but max. N seconds (default: 1).
--show-rc          show/log the return code (rc)
--no-files-cache  do not load/update the file metadata cache used to
                  detect unchanged files
--umask M          set umask to M (local and remote, default: 0077)

```

(continues on next page)

(continued from previous page)

```
--remote-path PATH    use PATH as borg executable on the remote (default:
                      "borg")
```

### 4.15.1 Description

The key files used for repository encryption are optionally passphrase protected. This command can be used to change this passphrase.

Please note that this command only changes the passphrase, but not any secret protected by it (like e.g. encryption/MAC keys or chunker seed). Thus, changing the passphrase after passphrase and borg key got compromised does not protect future (nor past) backups to the same repository.

### 4.15.2 Examples

```
# Create a key file protected repository
$ borg init --encryption=keyfile -v /path/to/repo
Initializing repository at "/path/to/repo"
Enter new passphrase:
Enter same passphrase again:
Remember your passphrase. Your data will be inaccessible without it.
Key in "/root/.config/borg/keys/mnt_backup" created.
Keep this key safe. Your data will be inaccessible without it.
Synchronizing chunks cache...
Archives: 0, w/ cached Idx: 0, w/ outdated Idx: 0, w/o cached Idx: 0.
Done.

# Change key file passphrase
$ borg change-passphrase -v /path/to/repo
Enter passphrase for key /root/.config/borg/keys/mnt_backup:
Enter new passphrase:
Enter same passphrase again:
Remember your passphrase. Your data will be inaccessible without it.
Key updated
```

## 4.16 borg serve

```
usage: borg serve [-h] [--critical] [--error] [--warning] [--info] [--debug]
                 [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                 [--remote-path PATH] [--restrict-to-path PATH]
                 [--append-only]
```

Start in server mode. This command is usually not used manually.

optional arguments:

```
-h, --help            show this help message and exit
--critical            work on log level CRITICAL
--error              work on log level ERROR
--warning            work on log level WARNING (default)
--info, -v, --verbose
```

(continues on next page)



(continued from previous page)

```

work on log level INFO
--debug                work on log level DEBUG
--lock-wait N          wait for the lock, but max. N seconds (default: 1).
--show-rc              show/log the return code (rc)
--no-files-cache      do not load/update the file metadata cache used to
                      detect unchanged files
--umask M              set umask to M (local and remote, default: 0077)
--remote-path PATH    use PATH as borg executable on the remote (default:
                      "borg")
--restrict-to-path PATH
                      restrict repository access to PATH. Can be specified
                      multiple times to allow the client access to several
                      directories. Access to all sub-directories is granted
                      implicitly; PATH doesn't need to directly point to a
                      repository.
--append-only          only allow appending to repository segment files

```

### 4.16.1 Description

This command starts a repository server process. This command is usually not used manually.

### 4.16.2 Examples

`borg serve` has special support for ssh forced commands (see `authorized_keys` example below): it will detect that you use such a forced command and extract the value of the `--restrict-to-path` option(s). It will then parse the original command that came from the client, makes sure that it is also `borg serve` and enforce path restriction(s) as given by the forced command. That way, other options given by the client (like `--info` or `--umask`) are preserved (and are not fixed by the forced command).

```

# Allow an SSH keypair to only run borg, and only have access to /path/to/repo.
# Use key options to disable unneeded and potentially dangerous SSH functionality.
# This will help to secure an automated remote backup system.
$ cat ~/.ssh/authorized_keys
command="borg serve --restrict-to-path /path/to/repo",no-pty,no-agent-forwarding,no-
↪port-forwarding,no-X11-forwarding,no-user-rc ssh-rsa AAAAB3[...]

```

## 4.17 borg upgrade

```

usage: borg upgrade [-h] [--critical] [--error] [--warning] [--info] [--debug]
                  [--lock-wait N] [--show-rc] [--no-files-cache] [--umask M]
                  [--remote-path PATH] [-p] [-n] [-i] [--force] [--tam]
                  [--disable-tam]
                  [REPOSITORY]

```

upgrade a repository from a previous version

positional arguments:

REPOSITORY path to the repository to be upgraded

optional arguments:

(continues on next page)

(continued from previous page)

```

-h, --help                show this help message and exit
--critical                work on log level CRITICAL
--error                  work on log level ERROR
--warning                 work on log level WARNING (default)
--info, -v, --verbose    work on log level INFO
--debug                  work on log level DEBUG
--lock-wait N             wait for the lock, but max. N seconds (default: 1).
--show-rc                show/log the return code (rc)
--no-files-cache         do not load/update the file metadata cache used to
                        detect unchanged files
--umask M                set umask to M (local and remote, default: 0077)
--remote-path PATH       use PATH as borg executable on the remote (default:
                        "borg")
-p, --progress           show progress display while upgrading the repository
-n, --dry-run            do not change repository
-i, --inplace            rewrite repository in place, with no chance of going
                        back to older versions of the repository.
--force                  Force upgrade
--tam                    Enable manifest authentication (in key and cache)
                        (Borg 1.0.9 and later)
--disable-tam           Disable manifest authentication (in key and cache)

```

### 4.17.1 Description

Upgrade an existing, local Borg repository.

#### When you do not need borg upgrade

Not every change requires that you run `borg upgrade`.

You do **not** need to run it when:

- moving your repository to a different place
- upgrading to another point release (like 1.0.x to 1.0.y), except when noted otherwise in the changelog
- upgrading from 1.0.x to 1.1.x, except when noted otherwise in the changelog

#### Borg 1.x.y upgrades

Use `borg upgrade --tam REPO` to require manifest authentication introduced with Borg 1.0.9 to address security issues. This means that modifying the repository after doing this with a version prior to 1.0.9 will raise a validation error, so only perform this upgrade after updating all clients using the repository to 1.0.9 or newer.

This upgrade should be done on each client for safety reasons.

If a repository is accidentally modified with a pre-1.0.9 client after this upgrade, use `borg upgrade --tam --force REPO` to remedy it.

If you routinely do this you might not want to enable this upgrade (which will leave you exposed to the security issue). You can reverse the upgrade by issuing `borg upgrade --disable-tam REPO`.

See <https://borgbackup.readthedocs.io/en/stable/changes.html#pre-1-0-9-manifest-spoofing-vulnerability> for details.

## Attic and Borg 0.xx to Borg 1.x

This currently supports converting an Attic repository to Borg and also helps with converting Borg 0.xx to 1.0.

Currently, only LOCAL repositories can be upgraded (issue #465).

Please note that `borg create` (since 1.0.0) uses bigger chunks by default than old `borg` or `attic` did, so the new chunks won't deduplicate with the old chunks in the upgraded repository. See `--chunker-params` option of `borg create` and `borg recreate`.

`borg upgrade` will change the magic strings in the repository's segments to match the new Borg magic strings. The keyfiles found in `$ATTIC_KEYS_DIR` or `~/.attic/keys/` will also be converted and copied to `$BORG_KEYS_DIR` or `~/.config/borg/keys`.

The cache files are converted, from `$ATTIC_CACHE_DIR` or `~/.cache/attic` to `$BORG_CACHE_DIR` or `~/.cache/borg`, but the cache layout between Borg and Attic changed, so it is possible the first backup after the conversion takes longer than expected due to the cache resync.

Upgrade should be able to resume if interrupted, although it will still iterate over all segments. If you want to start from scratch, use `borg delete` over the copied repository to make sure the cache files are also removed:

```
borg delete borg
```

Unless `--inplace` is specified, the upgrade process first creates a backup copy of the repository, in `REPOSITORY.upgrade-DATETIME`, using hardlinks. This takes longer than in place upgrades, but is much safer and gives progress information (as opposed to `cp -al`). Once you are satisfied with the conversion, you can safely destroy the backup copy.

**WARNING:** Running the upgrade in place will make the current copy unusable with older version, with no way of going back to previous versions. This can **PERMANENTLY DAMAGE YOUR REPOSITORY!** Attic **CAN NOT READ BORG REPOSITORIES**, as the magic strings have changed. You have been warned.

### 4.17.2 Examples

```
# Upgrade the borg repository to the most recent version.
$ borg upgrade -v /path/to/repo
making a hardlink copy in /path/to/repo.upgrade-2016-02-15-20:51:55
opening attic repository with borg and converting
no key file found for repository
converting repo index /path/to/repo/index.0
converting 1 segments...
converting borg 0.xx to borg current
no key file found for repository
```

### 4.17.3 Upgrading a passphrase encrypted attic repo

`attic` offered a “passphrase” encryption mode, but this was removed in `borg 1.0` and replaced by the “repokey” mode (which stores the passphrase-protected encryption key into the repository config).

Thus, to upgrade a “passphrase” attic repo to a “repokey” borg repo, 2 steps are needed, in this order:

- `borg upgrade repo`
- `borg migrate-to-repokey repo`

## 4.18 borg break-lock

```
usage: borg break-lock [-h] [--critical] [--error] [--warning] [--info]
                    [--debug] [--lock-wait N] [--show-rc]
                    [--no-files-cache] [--umask M] [--remote-path PATH]
                    [REPOSITORY]

Break the repository lock (e.g. in case it was left by a dead borg.)

positional arguments:
  REPOSITORY            repository for which to break the locks

optional arguments:
  -h, --help            show this help message and exit
  --critical            work on log level CRITICAL
  --error              work on log level ERROR
  --warning            work on log level WARNING (default)
  --info, -v, --verbose work on log level INFO
  --debug              work on log level DEBUG
  --lock-wait N        wait for the lock, but max. N seconds (default: 1).
  --show-rc            show/log the return code (rc)
  --no-files-cache     do not load/update the file metadata cache used to
                        detect unchanged files
  --umask M            set umask to M (local and remote, default: 0077)
  --remote-path PATH   use PATH as borg executable on the remote (default:
                        "borg")
```

### 4.18.1 Description

This command breaks the repository and cache locks. Please use carefully and only while no borg process (on any machine) is trying to access the Cache or the Repository.

## 4.19 Miscellaneous Help

### 4.19.1 borg help patterns

Exclusion patterns support four separate styles, `fnmatch`, `shell`, `regular expressions` and `path prefixes`. By default, `fnmatch` is used. If followed by a colon (`:`) the first two characters of a pattern are used as a style selector. Explicit style selection is necessary when a non-default style is desired or when the desired pattern starts with two alphanumeric characters followed by a colon (i.e. `aa:something/*`).

`Fnmatch`<sup>59</sup>, selector `fm`:

This is the default style. These patterns use a variant of shell pattern syntax, with `*` matching any number of characters, `?` matching any single character, `[...]` matching any single character specified, including ranges, and `[!...]` matching any character not specified. For the purpose of these patterns, the path separator (`\` for Windows and `/` on other systems) is not treated specially. Wrap meta-characters in brackets for a literal match (i.e. `[?]` to match the literal character `?`). For a path to match a pattern, it must completely match from start to end, or must match from the start to just before a path separator. Except for the root path, paths will never end in the path separator when matching is attempted. Thus, if a given pattern ends in a path separator, a `*` is appended before matching is attempted.

---

<sup>59</sup> <https://docs.python.org/3/library/fnmatch.html>

Shell-style patterns, selector *sh*:

Like `fnmatch` patterns these are similar to shell patterns. The difference is that the pattern may include `**/` for matching zero or more directory levels, `*` for matching zero or more arbitrary characters with the exception of any path separator.

Regular expressions, selector *re*:

Regular expressions similar to those found in Perl are supported. Unlike shell patterns regular expressions are not required to match the complete path and any substring match is sufficient. It is strongly recommended to anchor patterns to the start (`^`), to the end (`$`) or both. Path separators (`'` for Windows and `'` on other systems) in paths are always normalized to a forward slash (`'`) before applying a pattern. The regular expression syntax is described in the [Python documentation for the `re` module](#)<sup>60</sup>.

Prefix path, selector *pp*:

This pattern style is useful to match whole sub-directories. The pattern `pp:/data/bar` matches `/data/bar` and everything therein.

---

**Note:** *re:*, *sh:* and *fm:* patterns are all implemented on top of the Python SRE engine. It is very easy to formulate patterns for each of these types which requires an inordinate amount of time to match paths. If untrusted users are able to supply patterns, ensure they cannot supply *re:* patterns. Further, ensure that *sh:* and *fm:* patterns only contain a handful of wildcards at most.

---

Exclusions can be passed via the command line option `-exclude`. When used from within a shell the patterns should be quoted to protect them from expansion.

The `-exclude-from` option permits loading exclusion patterns from a text file with one pattern per line. Lines empty or starting with the number sign (`#`) after removing whitespace on both ends are ignored. The optional style selector prefix is also supported for patterns loaded from a file. Due to whitespace removal paths with whitespace at the beginning or end can only be excluded using regular expressions.

Examples:

```
# Exclude '/home/user/file.o' but not '/home/user/file.odt':
$ borg create -e '*.o' backup /

# Exclude '/home/user/junk' and '/home/user/subdir/junk' but
# not '/home/user/importantjunk' or '/etc/junk':
$ borg create -e '/home/*/junk' backup /

# Exclude the contents of '/home/user/cache' but not the directory itself:
$ borg create -e /home/user/cache/ backup /

# The file '/home/user/cache/important' is *not* backed up:
$ borg create -e /home/user/cache/ backup / /home/user/cache/important

# The contents of directories in '/home' are not backed up when their name
# ends in '.tmp'
$ borg create --exclude 're:^(/home/[^/]+\.\tmp/' backup /

# Load exclusions from file
$ cat >exclude.txt <<EOF
# Comment line
/home/*/junk
*.tmp
```

(continues on next page)

---

<sup>60</sup> <https://docs.python.org/3/library/re.html>

(continued from previous page)

```
fm:aa:something/*
re:~/home/[^/]\.tmp/
sh:/home/*/\.thumbnails
EOF
$ borg create --exclude-from exclude.txt backup /
```

## 4.19.2 borg help placeholders

Repository (or Archive) URLs, `--prefix` and `--remote-path` values support these placeholders:

`{hostname}`

The (short) hostname of the machine.

`{fqdn}`

The full name of the machine.

`{now}`

The current local date and time.

`{utcnow}`

The current UTC date and time.

`{user}`

The user name (or UID, if no name is available) of the user running borg.

`{pid}`

The current process ID.

`{borgversion}`

The version of borg, e.g.: 1.0.8rc1

`{borgmajor}`

The version of borg, only the major version, e.g.: 1

`{borgminor}`

The version of borg, only major and minor version, e.g.: 1.0

`{borgpatch}`

The version of borg, only major, minor and patch version, e.g.: 1.0.8

Examples:

```
borg create /path/to/repo::{hostname}-{user}-{utcnow} ...
borg create /path/to/repo::{hostname}-{now:%Y-%m-%d_%H:%M:%S} ...
borg prune --prefix '{hostname}-' ...
```

## 4.20 Debug Commands

There are some more commands (all starting with “debug-“) which are all **not intended for normal use** and **potentially very dangerous** if used incorrectly.

For example, `borg debug-put-obj` and `borg debug-delete-obj` will only do what their name suggests: put objects into repo / delete objects from repo.

Please note:

- they will not update the chunks cache (chunks index) about the object
- they will not update the manifest (so no automatic chunks index resync is triggered)
- they will not check whether the object is in use (e.g. before `delete-obj`)
- they will not update any metadata which may point to the object

They exist to improve debugging capabilities without direct system access, e.g. in case you ever run into some severe malfunction. Use them only if you know what you are doing or if a trusted Borg developer tells you what to do.

## 4.21 Additional Notes

Here are misc. notes about topics that are maybe not covered in enough detail in the usage section.

### 4.21.1 Item flags

`borg create -v --list` outputs a verbose list of all files, directories and other file system items it considered (no matter whether they had content changes or not). For each item, it prefixes a single-letter flag that indicates type and/or status of the item.

If you are interested only in a subset of that output, you can give e.g. `--filter=AME` and it will only show regular files with A, M or E status (see below).

A uppercase character represents the status of a regular file relative to the “files” cache (not relative to the repo – this is an issue if the files cache is not used). Metadata is stored in any case and for ‘A’ and ‘M’ also new data chunks are stored. For ‘U’ all data chunks refer to already existing chunks.

- ‘A’ = regular file, added (see also *I am seeing ‘A’ (added) status for a unchanged file!?* in the FAQ)
- ‘M’ = regular file, modified
- ‘U’ = regular file, unchanged
- ‘E’ = regular file, an error happened while accessing/reading *this* file

A lowercase character means a file type other than a regular file, borg usually just stores their metadata:

- ‘d’ = directory
- ‘b’ = block device
- ‘c’ = char device
- ‘h’ = regular file, hardlink (to already seen inodes)
- ‘s’ = symlink
- ‘f’ = fifo

Other flags used include:

- ‘i’ = backup data was read from standard input (stdin)
- ‘-’ = dry run, item was *not* backed up
- ‘?’ = missing status code (if you see this, please file a bug report!)

### 4.21.2 `--chunker-params`

The chunker params influence how input files are cut into pieces (chunks) which are then considered for deduplication. They also have a big impact on resource usage (RAM and disk space) as the amount of resources needed is (also) determined by the total amount of chunks in the repository (see *Indexes / Caches memory usage* for details).

`--chunker-params=10,23,16,4095` results in a fine-grained deduplication and creates a big amount of chunks and thus uses a lot of resources to manage them. This is good for relatively small data volumes and if the machine has a good amount of free RAM and disk space.

`--chunker-params=19,23,21,4095` (default) results in a coarse-grained deduplication and creates a much smaller amount of chunks and thus uses less resources. This is good for relatively big data volumes and if the machine has a relatively low amount of free RAM and disk space.

If you already have made some archives in a repository and you then change chunker params, this of course impacts deduplication as the chunks will be cut differently.

In the worst case (all files are big and were touched in between backups), this will store all content into the repository again.

Usually, it is not that bad though:

- usually most files are not touched, so it will just re-use the old chunks it already has in the repo
- files smaller than the (both old and new) minimum chunksize result in only one chunk anyway, so the resulting chunks are same and deduplication will apply

If you switch chunker params to save resources for an existing repo that already has some backup archives, you will see an increasing effect over time, when more and more files have been touched and stored again using the bigger chunksize **and** all references to the smaller older chunks have been removed (by deleting / pruning archives).

If you want to see an immediate big effect on resource usage, you better start a new repository when changing chunker params.

For more details, see *Chunks*.

### 4.21.3 `--noatime` / `--noctime`

You can use these `borg create` options to not store the respective timestamp into the archive, in case you do not really need it.

Besides saving a little space for the not archived timestamp, it might also affect metadata stream deduplication: if only this timestamp changes between backups and is stored into the metadata stream, the metadata stream chunks won't deduplicate just because of that.

### 4.21.4 `--umask`

If you use `--umask`, make sure that all repository-modifying borg commands (create, delete, prune) that access the repository in question use the same `--umask` value.

If multiple machines access the same repository, this should hold true for all of them.

### 4.21.5 `--read-special`

The `--read-special` option is special - you do not want to use it for normal full-filesystem backups, but rather after carefully picking some targets for it.



The option `--read-special` triggers special treatment for block and char device files as well as FIFOs. Instead of storing them as such a device (or FIFO), they will get opened, their content will be read and in the backup archive they will show up like a regular file.

Symlinks will also get special treatment if (and only if) they point to such a special file: instead of storing them as a symlink, the target special file will get processed as described above.

One intended use case of this is backing up the contents of one or multiple block devices, like e.g. LVM snapshots or inactive LVs or disk partitions.

You need to be careful about what you include when using `--read-special`, e.g. if you include `/dev/zero`, your backup will never terminate.

Restoring such files' content is currently only supported one at a time via `--stdout` option (and you have to redirect stdout to where ever it shall go, maybe directly into an existing device file of your choice or indirectly via `dd`).

To some extent, mounting a backup archive with the backups of special files via `borg mount` and then loop-mounting the image files from inside the mount point will work. If you plan to access a lot of data in there, it likely will scale and perform better if you do not work via the FUSE mount.

## Example

Imagine you have made some snapshots of logical volumes (LVs) you want to backup.

---

**Note:** For some scenarios, this is a good method to get “crash-like” consistency (I call it crash-like because it is the same as you would get if you just hit the reset button or your machine would abruptly and completely crash). This is better than no consistency at all and a good method for some use cases, but likely not good enough if you have databases running.

---

Then you create a backup archive of all these snapshots. The backup process will see a “frozen” state of the logical volumes, while the processes working in the original volumes continue changing the data stored there.

You also add the output of `lvdisplay` to your backup, so you can see the LV sizes in case you ever need to recreate and restore them.

After the backup has completed, you remove the snapshots again.

```
$ # create snapshots here
$ lvdisplay > lvdisplay.txt
$ borg create --read-special /path/to/repo::arch lvdisplay.txt /dev/vg0/*-snapshot
$ # remove snapshots here
```

Now, let's see how to restore some LVs from such a backup.

```
$ borg extract /path/to/repo::arch lvdisplay.txt
$ # create empty LVs with correct sizes here (look into lvdisplay.txt).
$ # we assume that you created an empty root and home LV and overwrite it now:
$ borg extract --stdout /path/to/repo::arch dev/vg0/root-snapshot > /dev/vg0/root
$ borg extract --stdout /path/to/repo::arch dev/vg0/home-snapshot > /dev/vg0/home
```

## 4.21.6 Append-only mode

A repository can be made “append-only”, which means that Borg will never overwrite or delete committed data. This is useful for scenarios where multiple machines back up to a central backup server using `borg serve`, since a hacked machine cannot delete backups permanently.

To activate append-only mode, edit the repository config file and add a line `append_only=1` to the `[repository]` section (or edit the line if it exists).

In append-only mode Borg will create a transaction log in the `transactions` file, where each line is a transaction and a UTC timestamp.

In addition, `borg serve` can act as if a repository is in append-only mode with its option `--append-only`. This can be very useful for fine-tuning access control in `.ssh/authorized_keys`

```
command="borg serve --append-only ..." ssh-rsa <key used for not-always-trustable_
↳backup clients>
command="borg serve ..." ssh-rsa <key used for backup management>
```

Please note that if you run `borg init` via a `borg serve --append-only` server, the repository config will be created with a `append_only=1` entry. This behaviour is subject to change in a later borg version. So, be aware of it for now, but do not rely on it.

### Example

Suppose an attacker remotely deleted all backups, but your repository was in append-only mode. A transaction log in this situation might look like this:

```
transaction 1, UTC time 2016-03-31T15:53:27.383532
transaction 5, UTC time 2016-03-31T15:53:52.588922
transaction 11, UTC time 2016-03-31T15:54:23.887256
transaction 12, UTC time 2016-03-31T15:55:54.022540
transaction 13, UTC time 2016-03-31T15:55:55.472564
```

From your security logs you conclude the attacker gained access at 15:54:00 and all the backups were deleted or replaced by compromised backups. From the log you know that transactions 11 and later are compromised. Note that the transaction ID is the name of the *last* file in the transaction. For example, transaction 11 spans files 6 to 11.

In a real attack you'll likely want to keep the compromised repository intact to analyze what the attacker tried to achieve. It's also a good idea to make this copy just in case something goes wrong during the recovery. Since recovery is done by deleting some files, a hard link copy (`cp -al`) is sufficient.

The first step to reset the repository to transaction 5, the last uncompromised transaction, is to remove the `hints.N` and `index.N` files in the repository (these two files are always expendable). In this example N is 13.

Then remove or move all segment files from the segment directories in `data/` starting with file 6:

```
rm data/**/{6..13}
```

That's all to it.

### Drawbacks

As data is only appended, and nothing removed, commands like `prune` or `delete` won't free disk space, they merely tag data as deleted in a new transaction.

Be aware that as soon as you write to the repo in non-append-only mode (e.g. `prune`, `delete` or create archives from an admin machine), it will remove the deleted objects permanently (including the ones that were already marked as deleted, but not removed, in append-only mode).

Note that you can go back-and-forth between normal and append-only operation by editing the configuration file, it's not a "one way trip".

## Further considerations

Append-only mode is not respected by tools other than Borg. `rm` still works on the repository. Make sure that backup client machines only get to access the repository via `borg serve`.

Ensure that no remote access is possible if the repository is temporarily set to normal mode for e.g. regular pruning.

Further protections can be implemented, but are outside of Borg's scope. For example, file system snapshots or wrapping `borg serve` to set special permissions or ACLs on new data files.

### 4.21.7 SSH batch mode

When running Borg using an automated script, `ssh` might still ask for a password, even if there is an SSH key for the target server. Use this to make scripts more robust:

```
export BORG_RSH='ssh -oBatchMode=yes'
```



This chapter will give an example how to setup a borg repository server for multiple clients.

### 5.1 Machines

There are multiple machines used in this chapter and will further be named by their respective fully qualified domain name (fqdn).

- The backup server: *backup01.srv.local*
- The clients:
  - John Doe’s desktop:  *johndoe.cln.local*
  - Webservers 01:  *web01.srv.local*
  - Application server 01:  *app01.srv.local*

### 5.2 User and group

The repository server needs to have only one UNIX user for all the clients. Recommended user and group with additional settings:

- User:  *backup*
- Group:  *backup*
- Shell:  */bin/bash* (or other capable to run the  *borg serve* command)
- Home:  */home/backup*

Most clients shall initiate a backup from the root user to catch all users, groups and permissions (e.g. when backing up  */home*).

## 5.3 Folders

The following folder tree layout is suggested on the repository server:

- User home directory, `/home/backup`
- Repositories path (storage pool): `/home/backup/repos`
- Clients restricted paths (`/home/backup/repos/<client fqdn>`):
  - johndoe.clnt.local: `/home/backup/repos/johndoe.clnt.local`
  - web01.srv.local: `/home/backup/repos/web01.srv.local`
  - app01.srv.local: `/home/backup/repos/app01.srv.local`

## 5.4 Restrictions

Borg is instructed to restrict clients into their own paths: `borg serve --restrict-to-path /home/backup/repos/<client fqdn>`

The client will be able to access any file or subdirectory inside of `/home/backup/repos/<client fqdn>` but no other directories. You can allow a client to access several separate directories by passing multiple `--restrict-to-path` flags, for instance: `borg serve --restrict-to-path /home/backup/repos/<client fqdn> --restrict-to-path /home/backup/repos/<other client fqdn>`, which could make sense if multiple machines belong to one person which should then have access to all the backups of their machines.

There is only one ssh key per client allowed. Keys are added for `johndoe.clnt.local`, `web01.srv.local` and `app01.srv.local`. But they will access the backup under only one UNIX user account as: `backup@backup01.srv.local`. Every key in `$HOME/.ssh/authorized_keys` has a forced command and restrictions applied as shown below:

```
command="cd /home/backup/repos/<client fqdn>;
    borg serve --restrict-to-path /home/backup/repos/<client fqdn>",
    no-port-forwarding,no-X11-forwarding,no-pty,
    no-agent-forwarding,no-user-rc <keytype> <key> <host>
```

---

**Note:** The text shown above needs to be written on a single line!

---

The options which are added to the key will perform the following:

1. Change working directory
2. Run `borg serve` restricted to the client base path
3. Restrict ssh and do not allow stuff which imposes a security risk

Due to the `cd` command we use, the server automatically changes the current working directory. Then client doesn't need to have knowledge of the absolute or relative remote repository path and can directly access the repositories at `<user>@<host>:<repo>`.

---

**Note:** The setup above ignores all client given commandline parameters which are normally appended to the `borg serve` command.

---

## 5.5 Client

The client needs to initialize the *pictures* repository like this:

```
borg init backup@backup01.srv.local:pictures
```

Or with the full path (should actually never be used, as only for demonstrational purposes). The server should automatically change the current working directory to the *<client fqdn>* folder.

```
borg init backup@backup01.srv.local:/home/backup/repos/johndoe.clnt.local/pictures
```

When *johndoe.clnt.local* tries to access a not restricted path the following error is raised. John Doe tries to backup into the Web 01 path:

```
borg init backup@backup01.srv.local:/home/backup/repos/web01.srv.local/pictures
```

```

~~~ SNIP ~~~
Remote: borg.remote.PathNotAllowed: /home/backup/repos/web01.srv.local/pictures
~~~ SNIP ~~~
Repository path not allowed

```

## 5.6 Ansible

Ansible takes care of all the system-specific commands to add the user, create the folder. Even when the configuration is changed the repository server configuration is satisfied and reproducible.

Automate setting up an repository server with the user, group, folders and permissions a Ansible playbook could be used. Keep in mind the playbook uses the Arch Linux [pacman](https://www.archlinux.org/pacman/pacman.8.html)<sup>61</sup> package manager to install and keep borg up-to-date.

```

- hosts: backup01.srv.local
  vars:
    user: backup
    group: backup
    home: /home/backup
    pool: "{{ home }}/repos"
    auth_users:
      - host: johndoe.clnt.local
        key: "{{ lookup('file', '/path/to/keys/johndoe.clnt.local.pub') }}"
      - host: web01.clnt.local
        key: "{{ lookup('file', '/path/to/keys/web01.clnt.local.pub') }}"
      - host: app01.clnt.local
        key: "{{ lookup('file', '/path/to/keys/app01.clnt.local.pub') }}"
  tasks:
    - pacman: name=borg state=latest update_cache=yes
    - group: name="{{ group }}" state=present
    - user: name="{{ user }}" shell=/bin/bash home="{{ home }}" createhome=yes group="{
↪{ group }}" groups= state=present
      - file: path="{{ home }}" owner="{{ user }}" group="{{ group }}" mode=0700_
↪state=directory
      - file: path="{{ home }}/.ssh" owner="{{ user }}" group="{{ group }}" mode=0700_
↪state=directory
      - file: path="{{ pool }}" owner="{{ user }}" group="{{ group }}" mode=0700_
↪state=directory
      - authorized_key: user="{{ user }}"
                          key="{{ item.key }}"

```

(continues on next page)

<sup>61</sup> <https://www.archlinux.org/pacman/pacman.8.html>

(continued from previous page)

```
        key_options='command="cd {{ pool }}/{{ item.host }};borg serve --
↳restrict-to-path {{ pool }}/{{ item.host }}",no-port-forwarding,no-X11-forwarding,
↳no-pty,no-agent-forwarding,no-user-rc'
        with_items: "{{ auth_users }}"
        - file: path="{{ home }}/.ssh/authorized_keys" owner="{{ user }}" group="{{ group }}"
↳" mode=0600 state=file
        - file: path="{{ pool }}/{{ item.host }}" owner="{{ user }}" group="{{ group }}"
↳mode=0700 state=directory
        with_items: "{{ auth_users }}"
```

## 5.7 Enhancements

As this chapter only describes a simple and effective setup it could be further enhanced when supporting (a limited set) of client supplied commands. A wrapper for starting *borg serve* could be written. Or borg itself could be enhanced to autodetect it runs under SSH by checking the *SSH\_ORIGINAL\_COMMAND* environment variable. This is left open for future improvements.

When extending ssh autodetection in borg no external wrapper script is necessary and no other interpreter or application has to be deployed.

## 5.8 See also

- [SSH Daemon manpage](#)<sup>62</sup>
- [Ansible](#)<sup>63</sup>

---

<sup>62</sup> <http://www.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/man8/sshd.8>

<sup>63</sup> <https://docs.ansible.com>



---

## Frequently asked questions

---

### 6.1 Can I backup VM disk images?

Yes, the [deduplication](https://en.wikipedia.org/wiki/Data_deduplication)<sup>64</sup> technique used by Borg makes sure only the modified parts of the file are stored. Also, we have optional simple sparse file support for extract.

If you use non-snapshotting backup tools like Borg to back up virtual machines, then the VMs should be turned off for the duration of the backup. Backing up live VMs can (and will) result in corrupted or inconsistent backup contents: a VM image is just a regular file to Borg with the same issues as regular files when it comes to concurrent reading and writing from the same file.

For backing up live VMs use file system snapshots on the VM host, which establishes crash-consistency for the VM images. This means that with most file systems (that are journaling) the FS will always be fine in the backup (but may need a journal replay to become accessible).

Usually this does not mean that file *contents* on the VM are consistent, since file contents are normally not journaled. Notable exceptions are ext4 in data=journal mode, ZFS and btrfs (unless nodatacow is used).

Applications designed with crash-consistency in mind (most relational databases like PostgreSQL, SQLite etc. but also for example Borg repositories) should always be able to recover to a consistent state from a backup created with crash-consistent snapshots (even on ext4 with data=writeback or XFS).

Hypervisor snapshots capturing most of the VM's state can also be used for backups and can be a better alternative to pure file system based snapshots of the VM's disk, since no state is lost. Depending on the application this can be the easiest and most reliable way to create application-consistent backups.

Other applications may require a lot of work to reach application-consistency: It's a broad and complex issue that cannot be explained in entirety here.

Borg doesn't intend to address these issues due to their huge complexity and platform/software dependency. Combining Borg with the mechanisms provided by the platform (snapshots, hypervisor features) will be the best approach to start tackling them.

---

<sup>64</sup> [https://en.wikipedia.org/wiki/Data\\_deduplication](https://en.wikipedia.org/wiki/Data_deduplication)

## 6.2 Can I backup from multiple servers into a single repository?

Yes, but in order for the deduplication used by Borg to work, it needs to keep a local cache containing checksums of all file chunks already stored in the repository. This cache is stored in `~/ .cache/borg/`. If Borg detects that a repository has been modified since the local cache was updated it will need to rebuild the cache. This rebuild can be quite time consuming.

So, yes it's possible. But it will be most efficient if a single repository is only modified from one place. Also keep in mind that Borg will keep an exclusive lock on the repository while creating or deleting archives, which may make *simultaneous* backups fail.

## 6.3 Can I copy or synchronize my repo to another location?

Yes, you could just copy all the files. Make sure you do that while no backup is running. So what you get here is this:

- client machine —borg create—> repo1
- repo1 —copy—> repo2

There is no special borg command to do the copying, just use `cp` or `rsync` if you want to do that.

But think about whether that is really what you want. If something goes wrong in `repo1`, you will have the same issue in `repo2` after the copy.

If you want to have 2 independent backups, it is better to do it like this:

- client machine —borg create—> repo1
- client machine —borg create—> repo2

## 6.4 Which file types, attributes, etc. are *not* preserved?

- UNIX domain sockets (because it does not make sense - they are meaningless without the running process that created them and the process needs to recreate them in any case). So, don't panic if your backup misses a UDS!
- The precise on-disk (or rather: not-on-disk) representation of the holes in a sparse file. Archive creation has no special support for sparse files, holes are backed up as (deduplicated and compressed) runs of zero bytes. Archive extraction has optional support to extract all-zero chunks as holes in a sparse file.
- filesystem specific attributes, like ext4 immutable bit, see #618<sup>65</sup>.
- For hardlinked symlinks, the hardlinking can not be archived (and thus, the hardlinking will not be done at extraction time). The symlinks will be archived and extracted as non-hardlinked symlinks, see #2379<sup>66</sup>.

## 6.5 Are there other known limitations?

- A single archive can only reference a limited volume of file/dir metadata, usually corresponding to tens or hundreds of millions of files/dirs. When trying to go beyond that limit, you will get a fatal `IntegrityError` exception telling that the (archive) object is too big. An easy workaround is to create multiple archives with less items each. See also the *Note about archive limitations* and #1452<sup>67</sup>.

---

<sup>65</sup> <https://github.com/borgbackup/borg/issues/618>

<sup>66</sup> <https://github.com/borgbackup/borg/issues/2379>

<sup>67</sup> <https://github.com/borgbackup/borg/issues/1452>

## 6.6 Why is my backup bigger than with attic? Why doesn't Borg do compression by default?

Attic was rather unflexible when it comes to compression, it always compressed using zlib level 6 (no way to switch compression off or adjust the level or algorithm).

Borg offers a lot of different compression algorithms and levels. Which of them is the best for you pretty much depends on your use case, your data, your hardware – so you need to do an informed decision about whether you want to use compression, which algorithm and which level you want to use. This is why compression defaults to none.

## 6.7 How can I specify the encryption passphrase programmatically?

The encryption passphrase or a command to retrieve the passphrase can be specified programmatically using the `BORG_PASSPHRASE` or `BORG_PASSCOMMAND` environment variables. This is convenient when setting up automated encrypted backups. Another option is to use key file based encryption with a blank passphrase. See *Repository encryption* for more details.

---

**Note:** Be careful how you set the environment; using the `env` command, a `system()` call or using inline shell scripts (e.g. `BORG_PASSPHRASE=hunter12 borg . . .`) might expose the credentials in the process list directly and they will be readable to all users on a system. Using `export` in a shell script file should be safe, however, as the environment of a process is [accessible only to that user](#)<sup>68</sup>.

---

## 6.8 When backing up to remote encrypted repos, is encryption done locally?

Yes, file and directory metadata and data is locally encrypted, before leaving the local machine. We do not mean the transport layer encryption by that, but the data/metadata itself. Transport layer encryption (e.g. when ssh is used as a transport) applies additionally.

## 6.9 When backing up to remote servers, do I have to trust the remote server?

Yes and No.

No, as far as data confidentiality is concerned - if you use encryption, all your files/dirs data and metadata are stored in their encrypted form into the repository.

Yes, as an attacker with access to the remote server could delete (or otherwise make unavailable) all your backups.

## 6.10 How can I protect against a hacked backup client?

Assume you backup your backup client machine C to the backup server S and C gets hacked. In a simple push setup, the attacker could then use borg on C to delete all backups residing on S.

---

<sup>68</sup> <https://security.stackexchange.com/questions/14000/environment-variable-accessibility-in-linux/14009#14009>

These are your options to protect against that:

- Do not allow to permanently delete data from the repo, see *Append-only mode*.
- Use a pull-mode setup using `ssh -R`, see #900<sup>69</sup>.
- Mount C's filesystem on another machine and then create a backup of it.
- Do not give C filesystem-level access to S.

## 6.11 How can I protect against a hacked backup server?

Just in case you got the impression that pull-mode backups are way more safe than push-mode, you also need to consider the case that your backup server S gets hacked. In case S has access to a lot of clients C, that might bring you into even bigger trouble than a hacked backup client in the previous FAQ entry.

These are your options to protect against that:

- Use the standard push-mode setup (see also previous FAQ entry).
- Mount (the repo part of) S's filesystem on C.
- Do not give S file-system level access to C.
- Have your backup server at a well protected place (maybe not reachable from the internet), configure it safely, apply security updates, monitor it, ...

## 6.12 How can I protect against theft, sabotage, lightning, fire, ... ?

In general: if your only backup medium is nearby the backedup machine and always connected, you can easily get into trouble: they likely share the same fate if something goes really wrong.

Thus:

- have multiple backup media
- have media disconnected from network, power, computer
- have media at another place
- have a relatively recent backup on your media

## 6.13 How do I report security issue with Borg?

Send a private email to the *Security* if you think you have discovered a security issue. Please disclose security issues responsibly.

## 6.14 Why do I get “connection closed by remote” after a while?

When doing a backup to a remote server (using a `ssh: repo URL`), it sometimes stops after a while (some minutes, hours, ... - not immediately) with “connection closed by remote” error message. Why?

That's a good question and we are trying to find a good answer in #636<sup>70</sup>.

---

<sup>69</sup> <https://github.com/borgbackup/borg/issues/900>

<sup>70</sup> <https://github.com/borgbackup/borg/issues/636>

## 6.15 Why am I seeing idle borg serve processes on the repo server?

Maybe the ssh connection between client and server broke down and that was not yet noticed on the server. Try these settings:

```
# /etc/ssh/sshd_config on borg repo server - kill connection to client
# after ClientAliveCountMax * ClientAliveInterval seconds with no response
ClientAliveInterval 20
ClientAliveCountMax 3
```

If you have multiple borg create ... ; borg create ... commands in a already serialized way in a single script, you need to give them `-lock-wait N` (with N being a bit more than the time the server needs to terminate broken down connections and release the lock).

## 6.16 The borg cache eats way too much disk space, what can I do?

There is a temporary (but maybe long lived) hack to avoid using lots of disk space for `chunks.archive.d` (see #235<sup>71</sup> for details):

```
# this assumes you are working with the same user as the backup.
# you can get the REPOID from the "config" file inside the repository.
cd ~/.cache/borg/<REPOID>
rm -rf chunks.archive.d ; touch chunks.archive.d
```

This deletes all the cached archive chunk indexes and replaces the directory that kept them with a file, so borg won't be able to store anything "in" there in future.

This has some pros and cons, though:

- much less disk space needs for `~/.cache/borg`.
- chunk cache resyncs will be slower as it will have to transfer chunk usage metadata for all archives from the repository (which might be slow if your repo connection is slow) and it will also have to build the hashtables from that data. chunk cache resyncs happen e.g. if your repo was written to by another machine (if you share same backup repo between multiple machines) or if your local chunks cache was lost somehow.

The long term plan to improve this is called "borgception", see #474<sup>72</sup>.

## 6.17 If a backup stops mid-way, does the already-backed-up data stay there?

Yes, Borg supports resuming backups.

During a backup a special checkpoint archive named `<archive-name>.checkpoint` is saved every checkpoint interval (the default value for this is 5 minutes) containing all the data backed-up until that point.

Checkpoints only happen between files (so they don't help for interruptions happening while a very large file is being processed).

This checkpoint archive is a valid archive (all files in it are valid and complete), but it is only a partial backup (not all files that you wanted to backup are contained in it). Having it in the repo until a successful, full backup is completed is

<sup>71</sup> <https://github.com/borgbackup/borg/issues/235>

<sup>72</sup> <https://github.com/borgbackup/borg/issues/474>

useful because it references all the transmitted chunks up to the checkpoint. This means that in case of an interruption, you only need to retransfer the data since the last checkpoint.

If a backup was interrupted, you do not need to do any special considerations, just invoke `borg create` as you always do. You may use the same archive name as in previous attempt or a different one (e.g. if you always include the current datetime), it does not matter.

Borg always does full single-pass backups, so it will start again from the beginning - but it will be much faster, because some of the data was already stored into the repo (and is still referenced by the checkpoint archive), so it does not need to get transmitted and stored again.

Once your backup has finished successfully, you can delete all `<archive-name>.checkpoint` archives.

## 6.18 How can I backup huge file(s) over a unstable connection?

You can use this “split trick” as a workaround for the in-between-files-only checkpoints (see above), huge files and a instable connection to the repository:

Split the huge file(s) into parts of manageable size (e.g. 100MB) and create a temporary archive of them. Borg will create checkpoints now more frequently than if you try to backup the files in their original form (e.g. 100GB).

After that, you can remove the parts again and backup the huge file(s) in their original form. This will now work a lot faster as a lot of content chunks are already in the repository.

After you have successfully backed up the huge original file(s), you can remove the temporary archive you made from the parts.

We realize that this is just a better-than-nothing workaround, see [#1198<sup>73</sup>](#) for a potential solution.

Please note that this workaround only helps you for backup, not for restore.

## 6.19 Can I backup my root partition (/) with Borg?

Backing up your entire root partition works just fine, but remember to exclude directories that make no sense to backup, such as `/dev`, `/proc`, `/sys`, `/tmp` and `/run`.

## 6.20 If it crashes with a UnicodeError, what can I do?

Check if your encoding is set correctly. For most POSIX-like systems, try:

```
export LANG=en_US.UTF-8 # or similar, important is correct charset
```

## 6.21 I can't extract non-ascii filenames by giving them on the commandline!?

This might be due to different ways to represent some characters in unicode or due to other non-ascii encoding issues.

If you run into that, try this:

- avoid the non-ascii characters on the commandline by e.g. extracting the parent directory (or even everything)

---

<sup>73</sup> <https://github.com/borgbackup/borg/issues/1198>

- mount the repo using FUSE and use some file manager

## 6.22 Can Borg add redundancy to the backup data to deal with hardware malfunction?

No, it can't. While that at first sounds like a good idea to defend against some defect HDD sectors or SSD flash blocks, dealing with this in a reliable way needs a lot of low-level storage layout information and control which we do not have (and also can't get, even if we wanted).

So, if you need that, consider RAID or a filesystem that offers redundant storage or just make backups to different locations / different hardware.

See also #225<sup>74</sup>.

## 6.23 Can Borg verify data integrity of a backup archive?

Yes, if you want to detect accidental data damage (like bit rot), use the `check` operation. It will notice corruption using CRCs and hashes. If you want to be able to detect malicious tampering also, use an encrypted repo. It will then be able to check using CRCs and HMACs.

## 6.24 I am seeing 'A' (added) status for a unchanged file!?

The files cache is used to determine whether Borg already "knows" / has backed up a file and if so, to skip the file from chunking. It does intentionally *not* contain files that have a modification time (mtime) same as the newest mtime in the created archive.

So, if you see an 'A' status for unchanged file(s), they are likely the files with the most recent mtime in that archive.

This is expected: it is to avoid data loss with files that are backed up from a snapshot and that are immediately changed after the snapshot (but within mtime granularity time, so the mtime would not change). Without the code that removes these files from the files cache, the change that happened right after the snapshot would not be contained in the next backup as Borg would think the file is unchanged.

This does not affect deduplication, the file will be chunked, but as the chunks will often be the same and already stored in the repo (except in the above mentioned rare condition), it will just re-use them as usual and not store new data chunks.

If you want to avoid unnecessary chunking, just create or touch a small or empty file in your backup source file set (so that one has the latest mtime, not your 50GB VM disk image) and, if you do snapshots, do the snapshot after that.

Since only the files cache is used in the display of files status, those files are reported as being added when, really, chunks are already used.

## 6.25 It always chunks all my files, even unchanged ones!

Borg maintains a files cache where it remembers the mtime, size and inode of files. When Borg does a new backup and starts processing a file, it first looks whether the file has changed (compared to the values stored in the files cache). If the values are the same, the file is assumed unchanged and thus its contents won't get chunked (again).

---

<sup>74</sup> <https://github.com/borgbackup/borg/issues/225>

Borg can't keep an infinite history of files of course, thus entries in the files cache have a “maximum time to live” which is set via the environment variable `BORG_FILES_CACHE_TTL` (and defaults to 20). Every time you do a backup (on the same machine, using the same user), the cache entries' ttl values of files that were not “seen” are incremented by 1 and if they reach `BORG_FILES_CACHE_TTL`, the entry is removed from the cache.

So, for example, if you do daily backups of 26 different data sets A, B, C, ..., Z on one machine (using the default TTL), the files from A will be already forgotten when you repeat the same backups on the next day and it will be slow because it would chunk all the files each time. If you set `BORG_FILES_CACHE_TTL` to at least 26 (or maybe even a small multiple of that), it would be much faster.

Another possible reason is that files don't always have the same path, for example if you mount a filesystem without stable mount points for each backup or if you are running the backup from a filesystem snapshot whose name is not stable. If the directory where you mount a filesystem is different every time, Borg assume they are different files.

## 6.26 Is there a way to limit bandwidth with Borg?

There is no command line option to limit bandwidth with Borg, but bandwidth limiting can be accomplished with [pipeviewer](#)<sup>75</sup>:

Create a wrapper script: `/usr/local/bin/pv-wrapper`

```
#!/bin/sh
## -q, --quiet           do not output any transfer information at all
## -L, --rate-limit RATE limit transfer to RATE bytes per second
RATE=307200
pv -q -L $RATE | "$@"
```

Add `BORG_RSH` environment variable to use pipeviewer wrapper script with ssh.

```
export BORG_RSH='/usr/local/bin/pv-wrapper ssh'
```

Now Borg will be bandwidth limited. Nice thing about pv is that you can change rate-limit on the fly:

```
pv -R $(pidof pv) -L 102400
```

## 6.27 I am having troubles with some network/FUSE/special filesystem, why?

Borg is doing nothing special in the filesystem, it only uses very common and compatible operations (even the locking is just “mkdir”).

So, if you are encountering issues like slowness, corruption or malfunction when using a specific filesystem, please try if you can reproduce the issues with a local (non-network) and proven filesystem (like ext4 on Linux).

If you can't reproduce the issue then, you maybe have found an issue within the filesystem code you used (not with Borg). For this case, it is recommended that you talk to the developers / support of the network fs and maybe open an issue in their issue tracker. Do not file an issue in the Borg issue tracker.

If you can reproduce the issue with the proven filesystem, please file an issue in the Borg issue tracker about that.

---

<sup>75</sup> <http://www.ivarch.com/programs/pv.shtml>



## 6.28 Why does running ‘borg check –repair’ warn about data loss?

Repair usually works for recovering data in a corrupted archive. However, it’s impossible to predict all modes of corruption. In some very rare instances, such as malfunctioning storage hardware, additional repo corruption may occur. If you can’t afford to lose the repo, it’s strongly recommended that you perform repair on a copy of the repo.

**In other words, the warning is there to emphasize that Borg:**

- Will perform automated routines that modify your backup repository
- Might not actually fix the problem you are experiencing
- Might, in very rare cases, further corrupt your repository

In the case of malfunctioning hardware, such as a drive or USB hub corrupting data when read or written, it’s best to diagnose and fix the cause of the initial corruption before attempting to repair the repo. If the corruption is caused by a one time event such as a power outage, running *borg check –repair* will fix most problems.

## 6.29 Why isn’t there more progress / ETA information displayed?

Some borg runs take quite a bit, so it would be nice to see a progress display, maybe even including a ETA (expected time of “arrival” [here rather “completion”]).

For some functionality, this can be done: if the total amount of work is more or less known, we can display progress. So check if there is a `--progress` option.

But sometimes, the total amount is unknown (e.g. for `borg create` we just do a single pass over the filesystem, so we do not know the total file count or data volume before reaching the end). Adding another pass just to determine that would take additional time and could be incorrect, if the filesystem is changing.

Even if the fs does not change and we knew count and size of all files, we still could not compute the `borg create` ETA as we do not know the amount of changed chunks, how the bandwidth of source and destination or system performance might fluctuate.

You see, trying to display ETA would be futile. The borg developers prefer to rather not implement progress / ETA display than doing futile attempts.

See also: <https://xkcd.com/612/>

## 6.30 Why am I getting ‘Operation not permitted’ errors when backing up on sshfs?

By default, `sshfs` is not entirely POSIX-compliant when renaming files due to a technicality in the SFTP protocol. Fortunately, it also provides a [workaround](#)<sup>76</sup> to make it behave correctly:

```
sshfs -o workaround=rename user@host:dir /mnt/dir
```

## 6.31 Requirements for the borg single-file binary, esp. (g)libc?

We try to build the binary on old, but still supported systems - to keep the minimum requirement for the (g)libc low. The (g)libc can’t be bundled into the binary as it needs to fit your kernel and OS, but Python and all other required

<sup>76</sup> <https://unix.stackexchange.com/a/123236>

libraries will be bundled into the binary.

If your system fulfills the minimum (g)libc requirement (see the README that is released with the binary), there should be no problem. If you are slightly below the required version, maybe just try. Due to the dynamic loading (or not loading) of some shared libraries, it might still work depending on what libraries are actually loaded and used.

In the borg git repository, there is `scripts/glibc_check.py` that can determine (based on the symbols' versions they want to link to) whether a set of given (Linux) binaries works with a given glibc version.

## 6.32 Why was Borg forked from Attic?

Borg was created in May 2015 in response to the difficulty of getting new code or larger changes incorporated into Attic and establishing a bigger developer community / more open development.

More details can be found in [ticket 217<sup>77</sup>](#) that led to the fork.

Borg intends to be:

- simple:
  - as simple as possible, but no simpler
  - do the right thing by default, but offer options
- open:
  - welcome feature requests
  - accept pull requests of good quality and coding style
  - give feedback on PRs that can't be accepted “as is”
  - discuss openly, don't work in the dark
- changing:
  - Borg is not compatible with Attic
  - do not break compatibility accidentally, without a good reason or without warning. allow compatibility breaking for other cases.
  - if major version number changes, it may have incompatible changes

## 6.33 What are the differences between Attic and Borg?

Borg is a fork of [Attic<sup>78</sup>](#) and maintained by “[The Borg collective<sup>79</sup>”](#).

Here's a (incomplete) list of some major changes:

- more open, faster paced development (see [issue #1<sup>80</sup>](#))
- lots of attic issues fixed (see [issue #5<sup>81</sup>](#))
- less chunk management overhead (less memory and disk usage for chunks index)
- faster remote cache resync (useful when backing up multiple machines into same repo)

---

<sup>77</sup> <https://github.com/jborg/attic/issues/217>

<sup>78</sup> <https://github.com/jborg/attic>

<sup>79</sup> <https://borgbackup.readthedocs.org/en/latest/authors.html>

<sup>80</sup> <https://github.com/borgbackup/borg/issues/1>

<sup>81</sup> <https://github.com/borgbackup/borg/issues/5>

- compression: no, lz4, zlib or lzma compression, adjustable compression levels
- repokey replaces problematic passphrase mode (you can't change the passphrase nor the pbkdf2 iteration count in "passphrase" mode)
- simple sparse file support, great for virtual machine disk files
- can read special files (e.g. block devices) or from stdin, write to stdout
- mkdir-based locking is more compatible than attic's posix locking
- uses fadvise to not spoil / blow up the fs cache
- better error messages / exception handling
- better logging, screen output, progress indication
- tested on misc. Linux systems, 32 and 64bit, FreeBSD, OpenBSD, NetBSD, Mac OS X

Please read the *Changelog* (or `docs/changes.rst` in the source distribution) for more information.

Borg is not compatible with original attic (but there is a one-way conversion).



Please first read the docs, the existing issue tracker issues and mailing list posts – a lot of stuff is already documented / explained / discussed / filed there.

### 7.1 Issue Tracker

If you've found a bug or have a concrete feature request, please create a new ticket on the project's [issue tracker](#)<sup>82</sup>.

For more general questions or discussions, IRC or mailing list are preferred.

### 7.2 Chat (IRC)

Join us on channel #borgbackup on chat.freenode.net.

As usual on IRC, just ask or tell directly and then patiently wait for replies. Stay connected.

You could use the following link (after connecting, you can change the random nickname you get by typing “/nick mydesirednickname”):

<http://webchat.freenode.net/?randomnick=1&channels=%23borgbackup&uio=MTY9dHJlZSY5PXRydWUa8>

### 7.3 Mailing list

To find out about the mailing list, its topic, how to subscribe, how to unsubscribe and where you can find the archives of the list, see the [mailing list homepage](#)<sup>83</sup>.

---

<sup>82</sup> <https://github.com/borgbackup/borg/issues>

<sup>83</sup> <https://mail.python.org/mailman/listinfo/borgbackup>

## 7.4 Twitter

Follow @borgbackup for announcements. You can also add @borgbackup if you would like to get retweeted for a borg related tweet.

Please understand that Twitter is not suitable for longer / more complex discussions - use one of the other channels for that.

## 7.5 Bounties and Fundraisers

We use [BountySource](#)<sup>84</sup> to allow monetary contributions to the project and the developers, who push it forward.

There, you can give general funds to the borgbackup members (the developers will then spend the funds as they deem fit). If you do not have some specific bounty (see below), you can use this as a general way to say “Thank You!” and support the software / project you like.

If you want to encourage developers to fix some specific issue or implement some specific feature suggestion, you can post a new bounty or back an existing one (they always refer to an issue in our [issue tracker](#)<sup>85</sup>).

As a developer, you can become a Bounty Hunter and win bounties (earn money) by contributing to Borg, a free and open source software project.

We might also use BountySource to fund raise for some bigger goals.

## 7.6 Security

In case you discover a security issue, please use this contact for reporting it privately and please, if possible, use encrypted E-Mail:

Thomas Waldmann <[tw@waldmann-edv.de](mailto:tw@waldmann-edv.de)>

GPG Key Fingerprint: 6D5B EF9A DD20 7580 5747 B70F 9F88 FB52 FAF7 B393

The public key can be fetched from any GPG keyserver, but be careful: you must use the **full fingerprint** to check that you got the correct key.

## 7.7 Verifying signed releases

[Releases](#)<sup>86</sup> are signed with the same GPG key and a .asc file is provided for each binary.

To verify a signature, the public key needs to be known to GPG. It can be imported into the local keystore from a keyserver with the fingerprint:

```
gpg --recv-keys "6D5B EF9A DD20 7580 5747 B70F 9F88 FB52 FAF7 B393"
```

If GPG successfully imported the key, the output should be (among other things): ‘Total number processed: 1’.

To verify for example the signature of the borg-linux64 binary:

---

<sup>84</sup> <https://www.bountysource.com/teams/borgbackup>

<sup>85</sup> <https://github.com/borgbackup/borg/issues>

<sup>86</sup> <https://github.com/borgbackup/borg/releases>

```
gpg --verify borg-linux64.asc
```

GPG outputs if it finds a good signature. The output should look similar to this:

```
gpg: Signature made Sat 30 Dec 2017 01:07:36 PM CET using RSA key ID 51F78E01
gpg: Good signature from "Thomas Waldmann <email>"
gpg: aka "Thomas Waldmann <email>"
gpg: aka "Thomas Waldmann <email>"
gpg: aka "Thomas Waldmann <email>"
gpg: WARNING: This key is not certified with a trusted signature!
gpg: There is no indication that the signature belongs to the owner.
Primary key fingerprint: 6D5B EF9A DD20 7580 5747 B70F 9F88 FB52 FAF7 B393
Subkey fingerprint: 2F81 AFFB AB04 E11F E8EE 65D4 243A CFA9 51F7 8E01
```

If you want to make absolutely sure that you have the right key, you need to verify it via another channel and assign a trust-level to it.





This is a collection of additional resources that are somehow related to borgbackup.

### 8.1 Videos, Talks, Presentations

Some of them refer to attic, but you can do the same stuff (and more) with borgbackup.

- [BorgBackup Installation and Basic Usage](#)<sup>87</sup> (english screencast)
- [TW's slides for borgbackup talks / lightning talks](#)<sup>88</sup> (just grab the latest ones)
- [Attic / Borg Backup talk from GPN 2015 \(media.ccc.de\)](#)<sup>89</sup>
- [Attic / Borg Backup talk from GPN 2015 \(youtube\)](#)<sup>90</sup>
- [Attic talk from Easterhegg 2015 \(media.ccc.de\)](#)<sup>91</sup>
- [Attic talk from Easterhegg 2015 \(youtube\)](#)<sup>92</sup>
- [Attic Backup: Mount your encrypted backups over ssh \(youtube\)](#)<sup>93</sup>
- [Evolution of Borg \(youtube\)](#)<sup>94</sup>

### 8.2 Software

- [BorgWeb - a very simple web UI for BorgBackup](#)<sup>95</sup>

---

<sup>87</sup> <https://asciinema.org/a/28691?autoplay=1&speed=2>

<sup>88</sup> <https://slides.com/thomaswaldmann>

<sup>89</sup> [https://media.ccc.de/browse/conferences/gpn/gpn15/gpn15-6942-attic\\_borg\\_backup.html#video](https://media.ccc.de/browse/conferences/gpn/gpn15/gpn15-6942-attic_borg_backup.html#video)

<sup>90</sup> <https://www.youtube.com/watch?v=Nb5nXEKSN-k>

<sup>91</sup> [https://media.ccc.de/v/eh15\\_-\\_49\\_-\\_saal\\_-\\_201504042130\\_-\\_attic\\_-\\_the\\_holy\\_grail\\_of\\_backups\\_-\\_thomas#video](https://media.ccc.de/v/eh15_-_49_-_saal_-_201504042130_-_attic_-_the_holy_grail_of_backups_-_thomas#video)

<sup>92</sup> <https://www.youtube.com/watch?v=96VEAAFJw>

<sup>93</sup> <https://www.youtube.com/watch?v=BVXDFv9Ymp8>

<sup>94</sup> [https://www.youtube.com/watch?v=K4k\\_4wDkG6Q](https://www.youtube.com/watch?v=K4k_4wDkG6Q)

<sup>95</sup> <https://borgweb.readthedocs.io/>

- some other stuff found at the [BorgBackup Github organisation](#)<sup>96</sup>
- [borgmatic](#)<sup>97</sup> - simple wrapper script for BorgBackup that creates and prunes backups

---

<sup>96</sup> <https://github.com/borgbackup/>

<sup>97</sup> <https://torsion.org/borgmatic/>

This section provides information about security and corruption issues.

### 9.1 Pre-1.0.9 manifest spoofing vulnerability (CVE-2016-10099)

A flaw in the cryptographic authentication scheme in Borg allowed an attacker to spoof the manifest. The attack requires an attacker to be able to

1. insert files (with no additional headers) into backups
2. gain write access to the repository

This vulnerability does not disclose plaintext to the attacker, nor does it affect the authenticity of existing archives.

The vulnerability allows an attacker to create a spoofed manifest (the list of archives). Creating plausible fake archives may be feasible for small archives, but is unlikely for large archives.

The fix adds a separate authentication tag to the manifest. For compatibility with prior versions this authentication tag is *not* required by default for existing repositories. Repositories created with 1.0.9 and later require it.

Steps you should take:

1. Upgrade all clients to 1.0.9 or later.
2. Run `borg upgrade --tam <repository> on every client` for *each* repository.
3. This will list all archives, including archive IDs, for easy comparison with your logs.
4. Done.

Prior versions can access and modify repositories with this measure enabled, however, to 1.0.9 or later their modifications are indiscernible from an attack and will raise an error until the below procedure is followed. We are aware that this can be annoying in some circumstances, but don't see a way to fix the vulnerability otherwise.

In case a version prior to 1.0.9 is used to modify a repository where above procedure was completed, and now you get an error message from other clients:

1. `borg upgrade --tam --force <repository> once with any client` suffices.

This attack is mitigated by:

- Noting/logging `borg list`, `borg info`, or `borg create --stats`, which contain the archive IDs.

We are not aware of others having discovered, disclosed or exploited this vulnerability.

Vulnerability time line:

- 2016-11-14: Vulnerability and fix discovered during review of cryptography by Marian Beermann (@enkore)
- 2016-11-20: First patch
- 2016-12-20: Released fixed version 1.0.9
- 2017-01-02: CVE was assigned
- 2017-01-15: Released fixed version 1.1.0b3 (fix was previously only available from source)

## 9.2 Pre-1.0.9 potential data loss

If you have archives in your repository that were made with `attic <= 0.13` (and later migrated to `borg`), running `borg check` would report errors in these archives. See issue #1837.

The reason for this is a invalid (and useless) metadata key that was always added due to a bug in these old `attic` versions.

If you run `borg check --repair`, things escalate quickly: all archive items with invalid metadata will be killed. Due to that `attic` bug, that means all items in all archives made with these old `attic` versions.

## 9.3 Pre-1.0.4 potential repo corruption

Some external errors (like network or disk I/O errors) could lead to corruption of the backup repository due to issue #1138.

A sign that this happened is if “E” status was reported for a file that can not be explained by problems with the source file. If you still have logs from “`borg create -v --list`”, you can check for “E” status.

Here is what could cause corruption and what you can do now:

1. I/O errors (e.g. repo disk errors) while writing data to repo.

This could lead to corrupted segment files.

Fix:

```
# check for corrupt chunks / segments:
borg check -v --repository-only REPO

# repair the repo:
borg check -v --repository-only --repair REPO

# make sure everything is fixed:
borg check -v --repository-only REPO
```

2. Unreliable network / unreliable connection to the repo.

This could lead to archive metadata corruption.

Fix:

```
# check for corrupt archives:
borg check -v --archives-only REPO

# delete the corrupt archives:
borg delete --force REPO::CORRUPT_ARCHIVE

# make sure everything is fixed:
borg check -v --archives-only REPO
```

3. In case you want to do more intensive checking.

The best check that everything is ok is to run a dry-run extraction:

```
borg extract -v --dry-run REPO::ARCHIVE
```



### 10.1 Version 1.0.12 (2018-04-08)

#### Bug fixes:

- repository: cleanup/write: invalidate cached FDs, tests
- serve: fix exitcode, #2910
- extract: set bsdfargs last (include immutable flag), #3263
- create -timestamp: set start time, #2957
- create: show excluded dir with “x” for tagged dirs / caches, #3189
- migrate locks to child PID when daemonize is used
- Buffer: fix wrong thread-local storage use, #2951
- fix detection of non-local path, #3108
- fix LDLP restoration for subprocesses, #3077
- fix subprocess environments (xattr module’s fakeroot version check, borg umount, BORG\_PASSCOMMAND), #3050
- remote: deal with partial lines, #2637
- get rid of datetime.isoformat, use safe parse\_timestamp to parse timestamps, #2994
- build: do .h file content checks in binary mode, fixes build issue for non-ascii header files on pure-ascii locale platforms, #3544 #3639
- remove platform.uname() call which caused library mismatch issues, #3732
- add exception handler around deprecated platform.linux\_distribution() call

#### Other changes:

- require msgpack-python >= 0.4.6 and < 0.5.0, see #3753

- add parens for C preprocessor macro argument usages (did not cause malfunction)
- ignore corrupt files cache, #2939
- replace “modulo” with “if” to check for wraparound in hashmap
- keymanager: don’t depend on optional readline module, #2980
- exclude broken pytest 3.3.0 release
- exclude broken Cython 0.27(.0) release, #3066
- flake8: add some ignores
- docs:
  - create: document exclusion through nodump
  - document good and problematic option placements, fix examples, #3356
  - update docs about hardlinked symlinks limitation
  - faq: we do not implement futile attempts of ETA / progress displays
  - simplified rate limiting wrapper in FAQ
  - twitter account @borgbackup, #2948
  - add note about metadata dedup and `-no[ac]time`, #2518
  - change-passphrase only changes the passphrase, #2990
  - clarify encrypted key format for borg key export, #3296
  - document sshfs rename workaround, #3315
  - update release checklist about security fixes
  - docs about how to verify a signed release, #3634
  - chunk seed is generated per `/repository/`
- vagrant:
  - use FUSE for macOS 3.7.1 to build the macOS binary
  - use python 3.5.5 to build the binaries
  - add exe location to PATH when we build an exe
  - use https pypi url for wheezy
  - netbsd: bash is already installed
  - netbsd: fix netbsd version in PKG\_PATH
  - use self-made FreeBSD 10.3 box, #3022
  - backport fs\_init (including related updates) from 1.1
  - the boxcutter wheezy boxes are 404, use local ones
- travis:
  - don’t perform full Travis build on docs-only changes, #2531
  - only short-circuit docs-only changes for pull requests



## 10.2 Version 1.0.11 (2017-07-21)

Bug fixes:

- use limited unpacker for outer key (security precaution), #2174
- fix paperkey import bug

Other changes:

- change `--checkpoint-interval` default from 600s to 1800s, #2841. this improves efficiency for big repositories a lot.
- docs: fix OpenSUSE command and add OpenSUSE section
- tests: add tests for `split_lstring` and `paperkey`
- vagrant:
  - fix `openbsd` shell
  - backport `cpu/ram` setup from master
  - add `stretch64` VM

## 10.3 Version 1.0.11rc1 (2017-06-27)

Bug fixes:

- performance: rebuild hashtable if we have too few empty buckets, #2246. this fixes some sporadic, but severe performance breakdowns.
- Archive: allocate zeros when needed, #2308 fixes huge memory usage of mount (8 MiB × number of archives)
- IPv6 address support also: Location: more informative exception when parsing fails
- borg single-file binary: use `pyinstaller v3.2.1`, #2396 this fixes that the prelink cronjob on some distros kills the borg binary by stripping away parts of it.
- extract:
  - warning for unextracted big extended attributes, #2258
  - also create parent dir for device files, if needed.
  - don't write to disk with `--stdout`, #2645
- archive check: detect and fix missing all-zero replacement chunks, #2180
- fix (de)compression exceptions, #2224 #2221
- files cache: update inode number, #2226
- borg rpc: use limited `msgpack.Unpacker` (security precaution), #2139
- Manifest: use limited `msgpack.Unpacker` (security precaution), #2175
- Location: accept `//servername/share/path`
- fix `ChunkIndex.__contains__` assertion for big-endian archs (harmless)
- create: handle `BackupOSError` on a per-path level in one spot
- fix error msg, there is no `--keep-last` in borg 1.0.x, #2282
- clamp (nano)second values to unproblematic range, #2304

- fuse / borg mount:
  - fix st\_blocks to be an integer (not float) value
  - fix negative uid/gid crash (they could come into archives e.g. when backing up external drives under cygwin), #2674
  - fix crash if empty (None) xattr is read
  - do pre-mount checks before opening repository
  - check llfuse is installed before asking for passphrase
- borg rename: expand placeholders, #2386
- borg serve: fix forced command lines containing **BORG\_** env vars
- fix error msg, it is `-keep-within`, not `-within`
- fix borg key/debug/benchmark crashing without subcommand, #2240
- chunker: fix invalid use of types, don't do `uint32_t >> 32`
- document `follow_symlinks` requirements, check `libc`, #2507

### New features:

- added `BORG_PASSCOMMAND` environment variable, #2573
  - add minimal version of in repository mandatory feature flags, #2134
- This should allow us to make sure older borg versions can be cleanly prevented from doing operations that are no longer safe because of repository format evolution. This allows more fine grained control than just incrementing the manifest version. So for example a change that still allows new archives to be created but would corrupt the repository when an old version tries to delete an archive or check the repository would add the new feature to the check and delete set but leave it out of the write set.
- borg delete `-force -force` to delete severely corrupted archives, #1975

### Other changes:

- embrace y2038 issue to support 32bit platforms
- be more clear that this is a “beyond repair” case, #2427
- key file names: limit to 100 characters and remove colons from host name
- upgrade FUSE for macOS to 3.5.8, #2346
- split up parsing and filtering for `-keep-within`, better error message, #2610
- docs:
  - fix caskroom link, #2299
  - address SSH batch mode, #2202 #2270
  - improve remote-path description
  - document snapshot usage, #2178
  - document relative path usage, #1868
  - one link per distro in the installation page
  - development: new branching model in git repository
  - kill api page
  - added FAQ section about backing up root partition

- add bountysource badge, #2558
- create empty docs.txt requirements, #2694
- README: how to help the project
- note -v/--verbose requirement on affected options, #2542
- document borg init behaviour via append-only borg serve, #2440
- be clear about what buzhash is used for (chunking) and what it is not used for (deduplication)- also say already in the readme that we use a cryptohash for dedupe, so people don't worry, #2390
- add hint about chunker params to borg upgrade docs, #2421
- clarify borg upgrade docs, #2436
- quickstart: delete problematic BORG\_PASSPRHASE use, #2623
- faq: specify "using inline shell scripts"
- document pattern denial of service, #2624
- tests:
  - remove attic dependency of the tests, #2505
  - travis:
    - \* enhance travis setuptools\_scm situation
    - \* install fakeroot for Linux
  - add test for borg delete --force
  - enable remote tests on cygwin (the cygwin issue that caused these tests to break was fixed in cygwin at least since cygwin 2.8, maybe even since 2.7.0).
  - remove skipping the noatime tests on GNU/Hurd, #2710
  - fix borg import issue, add comment, #2718
  - include attic.tar.gz when installing the package also: add include\_package\_data=True

## 10.4 Version 1.0.10 (2017-02-13)

### Bug fixes:

- Manifest timestamps are now monotonically increasing, this fixes issues when the system clock jumps backwards or is set inconsistently across computers accessing the same repository, #2115
- Fixed testing regression in 1.0.10rc1 that lead to a hard dependency on py.test >= 3.0, #2112

### New features:

- "key export" can now generate a printable HTML page with both a QR code and a human-readable "paperkey" representation (and custom text) through the --qr-html option.

The same functionality is also available through [paperkey.html](#), which is the same HTML page generated by --qr-html. It works with existing "key export" files and key files.

### Other changes:

- docs:

- language clarification - “borg create –one-file-system” option does not respect mount points, but considers different file systems instead, #2141
- setup.py: build\_api: sort file list for determinism

## 10.5 Version 1.0.10rc1 (2017-01-29)

### Bug fixes:

- borg serve: fix transmission data loss of pipe writes, #1268 This affects only the cygwin platform (not Linux, BSD, OS X).
- Avoid triggering an ObjectiveFS bug in xattr retrieval, #1992
- When running out of buffer memory when reading xattrs, only skip the current file, #1993
- Fixed “borg upgrade –tam” crashing with unencrypted repositories. Since *the issue* is not relevant for unencrypted repositories, it now does nothing and prints an error, #1981.
- Fixed change-passphrase crashing with unencrypted repositories, #1978
- Fixed “borg check repo::archive” indicating success if “archive” does not exist, #1997
- borg check: print non-exit-code warning if –last or –prefix aren’t fulfilled
- fix bad parsing of wrong repo location syntax
- create: don’t create hard link refs to failed files, mount: handle invalid hard link refs, #2092
- detect mingw byte order, #2073
- creating a new segment: use “xb” mode, #2099
- mount: umount on SIGINT/^C when in foreground, #2082

### Other changes:

- binary: use fixed AND freshly compiled pyinstaller bootloader, #2002
- xattr: ignore empty names returned by llistxattr(2) et al
- Enable the fault handler: install handlers for the SIGSEGV, SIGFPE, SIGABRT, SIGBUS and SIGILL signals to dump the Python traceback.
- Also print a traceback on SIGUSR2.
- borg change-passphrase: print key location (simplify making a backup of it)
- officially support Python 3.6 (setup.py: add Python 3.6 qualifier)
- tests:
  - vagrant / travis / tox: add Python 3.6 based testing
  - vagrant: fix openbsd repo, #2042
  - vagrant: fix the freebsd64 machine, #2037 #2067
  - vagrant: use python 3.5.3 to build binaries, #2078
  - vagrant: use osxfuse 3.5.4 for tests / to build binaries vagrant: improve darwin64 VM settings
  - travis: fix osxfuse install (fixes OS X testing on Travis CI)
  - travis: require succeeding OS X tests, #2028
  - travis: use latest pythons for OS X based testing

- use pytest-xdist to parallelize testing
- fix xattr test race condition, #2047
- setup.cfg: fix pytest deprecation warning, #2050
- docs:
  - language clarification - VM backup FAQ
  - borg create: document how to backup stdin, #2013
  - borg upgrade: fix incorrect title levels
  - add CVE numbers for issues fixed in 1.0.9, #2106
- fix typos (taken from Debian package patch)
- remote: include data hexdump in “unexpected RPC data” error message
- remote: log SSH command line at debug level
- API\_VERSION: use numberspaces, #2023
- remove .github from pypi package, #2051
- add pip and setuptools to requirements file, #2030
- SyncFile: fix use of fd object after close (cosmetic)
- Manifest.in: simplify, exclude \*.{so,dll,orig}, #2066
- ignore posix\_fadvise errors in repository.py, #2095 (works around issues with docker on ARM)
- make LoggedIO.close\_segment reentrant, avoid reentrance

## 10.6 Version 1.0.9 (2016-12-20)

### Security fixes:

- A flaw in the cryptographic authentication scheme in Borg allowed an attacker to spoof the manifest. See [Pre-1.0.9 manifest spoofing vulnerability \(CVE-2016-10099\)](#) above for the steps you should take.  
CVE-2016-10099 was assigned to this vulnerability.
- borg check: When rebuilding the manifest (which should only be needed very rarely) duplicate archive names would be handled on a “first come first serve” basis, allowing an attacker to apparently replace archives.  
CVE-2016-10100 was assigned to this vulnerability.

### Bug fixes:

- borg check:
  - rebuild manifest if it’s corrupted
  - skip corrupted chunks during manifest rebuild
- fix TypeError in integrity error handler, #1903, #1894
- fix location parser for archives with @ char (regression introduced in 1.0.8), #1930
- fix wrong duration/timestamps if system clock jumped during a create
- fix progress display not updating if system clock jumps backwards
- fix checkpoint interval being incorrect if system clock jumps

Other changes:

- docs:
  - add python3-devel as a dependency for cygwin-based installation
  - clarify extract is relative to current directory
  - FAQ: fix link to changelog
  - markup fixes
- tests:
  - test\_get\_(cachelkeys)\_dir: clean env state, #1897
  - get back pytest's pretty assertion failures, #1938
- setup.py build\_usage:
  - fixed build\_usage not processing all commands
  - fixed build\_usage not generating includes for debug commands

## 10.7 Version 1.0.9rc1 (2016-11-27)

Bug fixes:

- files cache: fix determination of newest mtime in backup set (which is used in cache cleanup and led to wrong “A” [added] status for unchanged files in next backup), #1860.
- borg check:
  - fix incorrectly reporting attic 0.13 and earlier archives as corrupt
  - handle repo w/o objects gracefully and also bail out early if repo is *completely* empty, #1815.
- fix tox/pybuild in 1.0-maint
- at xattr module import time, loggers are not initialized yet

New features:

- borg umount <mountpoint> exposed already existing umount code via the CLI api, so users can use it, which is more consistent than using borg to mount and fusermount -u (or umount) to un-mount, #1855.
- implement borg create --noatime --noctime, fixes #1853

Other changes:

- docs:
  - display README correctly on PyPI
  - improve cache / index docs, esp. files cache docs, fixes #1825
  - different pattern matching for --exclude, #1779
  - datetime formatting examples for {now} placeholder, #1822
  - clarify passphrase mode attic repo upgrade, #1854
  - clarify --umask usage, #1859
  - clarify how to choose PR target branch
  - clarify prune behavior for different archive contents, #1824

- fix PDF issues, add logo, fix authors, headings, TOC
- move security verification to support section
- fix links in standalone README (:ref: tags)
- add link to security contact in README
- add FAQ about security
- move fork differences to FAQ
- add more details about resource usage
- tests: skip remote tests on cygwin, #1268
- travis:
  - allow OS X failures until the brew cask osxfuse issue is fixed
  - caskroom osxfuse-beta gone, it's osxfuse now (3.5.3)
- vagrant:
  - upgrade OSXFuse / FUSE for macOS to 3.5.3
  - remove llfuse from tox.ini at a central place
  - do not try to install llfuse on centos6
  - fix fuse test for darwin, #1546
  - add windows virtual machine with cygwin
  - Vagrantfile cleanup / code deduplication

## 10.8 Version 1.0.8 (2016-10-29)

### Bug fixes:

- RemoteRepository: Fix busy wait in call\_many, #940

### New features:

- implement borgmajor/borgminor/borgpatch placeholders, #1694 {borgversion} was already there (full version string). With the new placeholders you can now also get e.g. 1 or 1.0 or 1.0.8.

### Other changes:

- avoid previous\_location mismatch, #1741

due to the changed canonicalization for relative pathes in PR #1711 / #1655 (implement ../ relpath hack), there would be a changed repo location warning and the user would be asked if this is ok. this would break automation and require manual intervention, which is unwanted.

thus, we automatically fix the previous\_location config entry, if it only changed in the expected way, but still means the same location.

- docs:
  - deployment.rst: do not use bare variables in ansible snippet
  - add clarification about append-only mode, #1689
  - setup.py: add comment about requiring llfuse, #1726
  - update usage.rst / api.rst

- repo url / archive location docs + typo fix
- quickstart: add a comment about other (remote) filesystems
- vagrant / tests:
  - no chown when rsyncing (fixes boxes w/o vagrant group)
  - fix fuse permission issues on linux/freebsd, #1544
  - skip fuse test for borg binary + fakeroot
  - ignore security.selinux xattrs, fixes tests on centos, #1735

## 10.9 Version 1.0.8rc1 (2016-10-17)

### Bug fixes:

- fix signal handling (SIGINT, SIGTERM, SIGHUP), #1620 #1593 Fixes e.g. leftover lock files for quickly repeated signals (e.g. Ctrl-C Ctrl-C) or lost connections or systemd sending SIGHUP.
- progress display: adapt formatting to narrow screens, do not crash, #1628
- borg create --read-special - fix crash on broken symlink, #1584. also correctly processes broken symlinks. before this regressed to a crash (5b45385) a broken symlink would've been skipped.
- process\_symlink: fix missing backup\_io() Fixes a chmod/chown/chgrp/unlink/rename/... crash race between getting dirents and dispatching to process\_symlink.
- yes(): abort on wrong answers, saying so, #1622
- fixed exception borg serve raised when connection was closed before repository was opened. add an error message for this.
- fix read-from-closed-FD issue, #1551 (this seems not to get triggered in 1.0.x, but was discovered in master)
- hashindex: fix iterators (always raise StopIteration when exhausted) (this seems not to get triggered in 1.0.x, but was discovered in master)
- enable relative pathes in ssh:// repo URLs, via ./relpath hack, #1655
- allow repo pathes with colons, #1705
- update changed repo location immediately after acceptance, #1524
- fix debug get-obj / delete-obj crash if object not found and remote repo, #1684
- pyinstaller: use a spec file to build borg.exe binary, exclude osxfuse dylib on Mac OS X (avoids mismatch lib <-> driver), #1619

### New features:

- add “borg key export” / “borg key import” commands, #1555, so users are able to backup / restore their encryption keys more easily.

Supported formats are the keyfile format used by borg internally and a special “paper” format with by line checksums for printed backups. For the paper format, the import is an interactive process which checks each line as soon as it is input.

- add “borg debug-refcount-obj” to determine a repo objects’ referrer counts, #1352

### Other changes:

- add “borg debug ...” subcommands (borg debug-\* still works, but will be removed in borg 1.1)



- setup.py: Add subcommand support to build\_usage.
- remote: change exception message for unexpected RPC data format to indicate dataflow direction.
- improved messages / error reporting:
  - IntegrityError: add placeholder for message, so that the message we give appears not only in the traceback, but also in the (short) error message, #1572
  - borg.key: include chunk id in exception msgs, #1571
  - better messages for cache newer than repo, #1700
- vagrant (testing/build VMs):
  - upgrade OSXfuse / FUSE for macOS to 3.5.2
  - update Debian Wheezy boxes, #1686
  - openbsd / netbsd: use own boxes, fixes misc rsync installation and fuse/llfuse related testing issues, #1695 #1696 #1670 #1671 #1728
- docs:
  - add docs for “key export” and “key import” commands, #1641
  - fix inconsistency in FAQ (pv-wrapper).
  - fix second block in “Easy to use” section not showing on GitHub, #1576
  - add bestpractices badge
  - link reference docs and faq about BORG\_FILES\_CACHE\_TTL, #1561
  - improve borg info –help, explain size infos, #1532
  - add release signing key / security contact to README, #1560
  - add contribution guidelines for developers
  - development.rst: add sphinx\_rtd\_theme to the sphinx install command
  - adjust border color in borg.css
  - add debug-info usage help file
  - internals.rst: fix typos
  - setup.py: fix build\_usage to always process all commands
  - added docs explaining multiple –restrict-to-path flags, #1602
  - add more specific warning about write-access debug commands, #1587
  - clarify FAQ regarding backup of virtual machines, #1672
- tests:
  - work around fuse xattr test issue with recent fakeroot
  - simplify repo/hashindex tests
  - travis: test fuse-enabled borg, use trusty to have a recent FUSE
  - re-enable fuse tests for RemoteArchiver (no deadlocks any more)
  - clean env for pytest based tests, #1714
  - fuse\_mount contextmanager: accept any options

## 10.10 Version 1.0.7 (2016-08-19)

### Security fixes:

- borg serve: fix security issue with remote repository access, #1428 If you used e.g. `--restrict-to-path /path/client1/` (with or without trailing slash does not make a difference), it acted like a path prefix match using `/path/client1` (note the missing trailing slash) - the code then also allowed working in e.g. `/path/client13` or `/path/client1000`.

As this could accidentally lead to major security/privacy issues depending on the pathes you use, the behaviour was changed to be a strict directory match. That means `--restrict-to-path /path/client1` (with or without trailing slash does not make a difference) now uses `/path/client1/` internally (note the trailing slash here!) for matching and allows precisely that path AND any path below it. So, `/path/client1` is allowed, `/path/client1/repo1` is allowed, but not `/path/client13` or `/path/client1000`.

If you willingly used the undocumented (dangerous) previous behaviour, you may need to rearrange your `--restrict-to-path` pathes now. We are sorry if that causes work for you, but we did not want a potentially dangerous behaviour in the software (not even using a for-backwards-compat option).

### Bug fixes:

- fixed repeated `LockTimeout` exceptions when borg serve tried to write into a already write-locked repo (e.g. by a borg mount), #502 part b) This was solved by the fix for #1220 in 1.0.7rc1 already.
- fix cosmetics + file leftover for “not a valid borg repository”, #1490
- Cache: release lock if cache is invalid, #1501
- borg extract `--strip-components`: fix leak of preloaded chunk contents
- Repository, when a `InvalidRepository` exception happens:
  - fix spurious, empty `lock.roster`
  - fix repo not closed cleanly

### New features:

- implement borg debug-info, fixes #1122 (just calls already existing code via cli, same output as below tracebacks)

### Other changes:

- skip the `O_NOATIME` test on GNU Hurd, fixes #1315 (this is a very minor issue and the GNU Hurd project knows the bug)
- document using a clean repo to test / build the release

## 10.11 Version 1.0.7rc2 (2016-08-13)

### Bug fixes:

- do not write objects to repository that are bigger than the allowed size, borg will reject reading them, #1451.  
Important: if you created archives with many millions of files or directories, please verify if you can open them successfully, e.g. try a “borg list REPO::ARCHIVE”.
- lz4 compression: dynamically enlarge the (de)compression buffer, the static buffer was not big enough for archives with extremely many items, #1453
- larger item metadata stream chunks, raise archive item limit by 8x, #1452

- fix untracked segments made by moved DELETES, #1442

Impact: Previously (metadata) segments could become untracked when deleting data, these would never be cleaned up.

- extended attributes (xattrs) related fixes:
  - fixed a race condition in xattrs querying that led to the entire file not being backed up (while logging the error, exit code = 1), #1469
  - fixed a race condition in xattrs querying that led to a crash, #1462
  - raise OSError including the error message derived from errno, deal with path being a integer FD

Other changes:

- print active env var override by default, #1467
- xattr module: refactor code, deduplicate, clean up
- repository: split object size check into too small and too big
- add a transaction\_id assertion, so borg init on a broken (inconsistent) filesystem does not look like a coding error in borg, but points to the real problem.
- explain confusing TypeError caused by compat support for old servers, #1456
- add forgotten usage help file from build\_usage
- refactor/unify buffer code into helpers.Buffer class, add tests
- docs:
  - document archive limitation, #1452
  - improve prune examples

## 10.12 Version 1.0.7rc1 (2016-08-05)

Bug fixes:

- fix repo lock deadlocks (related to lock upgrade), #1220
- catch unpacker exceptions, resync, #1351
- fix borg break-lock ignoring BORG\_REPO env var, #1324
- files cache performance fixes (fixes unnecessary re-reading/chunking/ hashing of unmodified files for some use cases):
  - fix unintended file cache eviction, #1430
  - implement BORG\_FILES\_CACHE\_TTL, update FAQ, raise default TTL from 10 to 20, #1338
- FUSE:
  - cache partially read data chunks (performance), #965, #966
  - always create a root dir, #1125
- use an OrderedDict for helptext, making the build reproducible, #1346
- RemoteRepository init: always call close on exceptions, #1370 (cosmetic)
- ignore stdout/stderr broken pipe errors (cosmetic), #1116

New features:

- better borg versions management support (useful esp. for borg servers wanting to offer multiple borg versions and for clients wanting to choose a specific server borg version), #1392:
  - add BORG\_VERSION environment variable before executing “borg serve” via ssh
  - add new placeholder {borgversion}
  - substitute placeholders in –remote-path
- borg init –append-only option (makes using the more secure append-only mode more convenient. when used remotely, this requires 1.0.7+ also on the borg server), #1291.

Other changes:

- Vagrantfile:
  - darwin64: upgrade to FUSE for macOS 3.4.1 (aka osxfuse), #1378
  - xenial64: use user “ubuntu”, not “vagrant” (as usual), #1331
- tests:
  - fix fuse tests on OS X, #1433
- docs:
  - FAQ: add backup using stable filesystem names recommendation
  - FAQ about glibc compatibility added, #491, glibc-check improved
  - FAQ: ‘A’ unchanged file; remove ambiguous entry age sentence.
  - OS X: install pkg-config to build with FUSE support, fixes #1400
  - add notes about shell/sudo pitfalls with env. vars, #1380
  - added platform feature matrix
- implement borg debug-dump-repo-objs

## 10.13 Version 1.0.6 (2016-07-12)

Bug fixes:

- Linux: handle multiple LD\_PRELOAD entries correctly, #1314, #1111
- Fix crash with unclear message if the libc is not found, #1314, #1111

Other changes:

- tests:
  - Fixed O\_NOATIME tests for Solaris and GNU Hurd, #1315
  - Fixed sparse file tests for (file) systems not supporting it, #1310
- docs:
  - Fixed syntax highlighting, #1313
  - misc docs: added data processing overview picture

## 10.14 Version 1.0.6rc1 (2016-07-10)

New features:

- borg check --repair: heal damaged files if missing chunks re-appear (e.g. if the previously missing chunk was added again in a later backup archive), #148. (\*) Also improved logging.

Bug fixes:

- sync\_dir: silence fsync() failing with EINVAL, #1287 Some network filesystems (like smbfs) don't support this and we use this in repository code.
- borg mount (FUSE):
  - fix directories being shadowed when contained paths were also specified, #1295
  - raise I/O Error (EIO) on damaged files (unless -o allow\_damaged\_files is used), #1302. (\*)
- borg extract: warn if a damaged file is extracted, #1299. (\*)
- Added some missing return code checks (ChunkIndex.\_add, hashindex\_resize).
- borg check: fix/optimize initial hash table size, avoids resize of the table.

Other changes:

- tests:
  - add more FUSE tests, #1284
  - deduplicate fuse (u)mount code
  - fix borg binary test issues, #862
- docs:
  - changelog: added release dates to older borg releases
  - fix some sphinx (docs generator) warnings, #881

Notes:

(\*) Some features depend on information (chunks\_healthy list) added to item metadata when a file with missing chunks was “repaired” using all-zero replacement chunks. The chunks\_healthy list is generated since borg 1.0.4, thus borg can't recognize such “repaired” (but content-damaged) files if the repair was done with an older borg version.

## 10.15 Version 1.0.5 (2016-07-07)

Bug fixes:

- borg mount: fix FUSE crash in xattr code on Linux introduced in 1.0.4, #1282

Other changes:

- backport some FAQ entries from master branch
- add release helper scripts
- Vagrantfile:
  - centos6: no FUSE, don't build binary
  - add xz for redhat-like dists

## 10.16 Version 1.0.4 (2016-07-07)

New features:

- `borg serve --append-only`, #1168 This was included because it was a simple change (append-only functionality was already present via repository config file) and makes better security now practically usable.
- `BORG_REMOTE_PATH` environment variable, #1258 This was included because it was a simple change (`--remote-path` cli option was already present) and makes borg much easier to use if you need it.
- Repository: cleanup incomplete transaction on “no space left” condition. In many cases, this can avoid a 100% full repo filesystem (which is very problematic as borg always needs free space - even to delete archives).

Bug fixes:

- Fix wrong handling and reporting of `OSError`s in `borg create`, #1138. This was a serious issue: in the context of “borg create”, errors like repository I/O errors (e.g. disk I/O errors, ssh repo connection errors) were handled badly and did not lead to a crash (which would be good for this case, because the repo transaction would be incomplete and trigger a transaction rollback to clean up). Now, error handling for source files is cleanly separated from every other error handling, so only problematic input files are logged and skipped.
- Implement fail-safe error handling for `borg extract`. Note that this isn’t nearly as critical as the `borg create` error handling bug, since nothing is written to the repo. So this was “merely” misleading error reporting.
- Add missing error handler in directory attr restore loop.
- `repo`: make sure write data hits disk before the commit tag (#1236) and also sync the containing directory.
- `FUSE`: `getxattr` fail must use `errno.ENOATTR`, #1126 (fixes Mac OS X Finder malfunction: “zero bytes” file length, access denied)
- `borg check --repair`: do not lose information about the good/original chunks. If we do not lose the original chunk IDs list when “repairing” a file (replacing missing chunks with all-zero chunks), we have a chance to “heal” the file back into its original state later, in case the chunks re-appear (e.g. in a fresh backup). Healing is not implemented yet, see #148.
- fixes for `--read-special` mode:
  - ignore known files cache, #1241
  - fake regular file mode, #1214
  - improve symlinks handling, #1215
- remove passphrase from subprocess environment, #1105
- Ignore empty index file (will trigger index rebuild), #1195
- add missing placeholder support for `--prefix`, #1027
- improve exception handling for placeholder replacement
- catch and format exceptions in arg parsing
- helpers: fix “undefined name ‘e’” in exception handler
- better error handling for missing repo manifest, #1043
- `borg delete`:
  - make it possible to delete a repo without manifest
  - `borg delete --forced` allows to delete corrupted archives, #1139
- `borg check`:

- make borg check work for empty repo
- fix resync and msgpacked item qualifier, #1135
- rebuild\_manifest: fix crash if 'name' or 'time' key were missing.
- better validation of item metadata dicts, #1130
- better validation of archive metadata dicts
- close the repo on exit - even if rollback did not work, #1197. This is rather cosmetic, it avoids repo closing in the destructor.
- tests:
  - fix sparse file test, #1170
  - flake8: ignore new F405, #1185
  - catch "invalid argument" on cygwin, #257
  - fix sparseness assertion in test prep, #1264

Other changes:

- make borg build/work on OpenSSL 1.0 and 1.1, #1187
- docs / help:
  - fix / clarify prune help, #1143
  - fix "patterns" help formatting
  - add missing docs / help about placeholders
  - resources: rename atticmatic to borgmatic
  - document sshd settings, #545
  - more details about checkpoints, add split trick, #1171
  - support docs: add freenode web chat link, #1175
  - add prune visualization / example, #723
  - add note that Fnmatch is default, #1247
  - make clear that lzma levels > 6 are a waste of cpu cycles
  - add a "do not edit" note to auto-generated files, #1250
  - update cygwin installation docs
- repository interoperability with borg master (1.1dev) branch:
  - borg check: read item metadata keys from manifest, #1147
  - read v2 hints files, #1235
  - fix hints file "unknown version" error handling bug
- tests: add tests for format\_line
- llfuse: update version requirement for freebsd
- Vagrantfile:
  - use openbsd 5.9, #716
  - do not install llfuse on netbsd (broken)

- update OSXfuse to version 3.3.3
- use Python 3.5.2 to build the binaries
- glibc compatibility checker: scripts/glibc\_check.py
- add .eggs to .gitignore

## 10.17 Version 1.0.3 (2016-05-20)

Bug fixes:

- prune: avoid that checkpoints are kept and completed archives are deleted in a prune run), #997
- prune: fix commandline argument validation - some valid command lines were considered invalid (annoying, but harmless), #942
- fix capabilities extraction on Linux (set xattrs last, after chown()), #1069
- repository: fix commit tags being seen in data
- when probing key files, do binary reads. avoids crash when non-borg binary files are located in borg's key files directory.
- handle SIGTERM and make a clean exit - avoids orphan lock files.
- repository cache: don't cache large objects (avoid using lots of temp. disk space), #1063

Other changes:

- Vagrantfile: OS X: update osxfuse / install lzma package, #933
- setup.py: add check for platform\_darwin.c
- setup.py: on freebsd, use a llfuse release that builds ok
- docs / help:
  - update readthedocs URLs, #991
  - add missing docs for “borg break-lock”, #992
  - borg create help: add some words to about the archive name
  - borg create help: document format tags, #894

## 10.18 Version 1.0.2 (2016-04-16)

Bug fixes:

- fix malfunction and potential corruption on (nowadays rather rare) big-endian architectures or bi-endian archs in (rare) BE mode. #886, #889
- cache resync / index merge was malfunctioning due to this, potentially leading to data loss. borg info had cosmetic issues (displayed wrong values).
- note: all (widespread) little-endian archs (like x86/x64) or bi-endian archs in (widespread) LE mode (like ARMEL, MIPSEL, ...) were NOT affected.
- add overflow and range checks for 1st (special) uint32 of the hashindex values, switch from int32 to uint32.
  - fix so that refcount will never overflow, but just stick to max. value after a overflow would have occurred.



- borg delete: fix `-cache-only` for broken caches, #874  
Makes `-cache-only` idempotent: it won't fail if the cache is already deleted.
- fixed borg create `-one-file-system` erroneously traversing into other filesystems (if starting fs device number was 0), #873
- workround a bug in Linux fadvise `FADV_DONTNEED`, #907

Other changes:

- better test coverage for hashindex, incl. overflow testing, checking correct computations so endianness issues would be discovered.
- reproducible doc for `ProgressIndicator*`, make the build reproducible.
- use latest llfuse for vagrant machines
- docs:
  - use `/path/to/repo` in examples, fixes #901
  - fix confusing usage of “repo” as archive name (use “arch”)

## 10.19 Version 1.0.1 (2016-04-08)

New features:

Usually there are no new features in a bugfix release, but these were added due to their high impact on security/safety/speed or because they are fixes also:

- append-only mode for repositories, #809, #36 (see docs)
- borg create: add `-ignore-inode` option to make borg detect unmodified files even if your filesystem does not have stable inode numbers (like sshfs and possibly CIFS).
- add options `-warning`, `-error`, `-critical` for missing log levels, #826. it's not recommended to suppress warnings or errors, but the user may decide this on his own. note: `-warning` is not given to borg serve so a `<= 1.0.0` borg will still work as server (it is not needed as it is the default). do not use `-error` or `-critical` when using a `<= 1.0.0` borg server.

Bug fixes:

- fix silently skipping EIO, #748
- add context manager for Repository (avoid orphan repository locks), #285
- do not sleep for >60s while waiting for lock, #773
- unpack file stats before passing to FUSE
- fix build on illumos
- don't try to backup doors or event ports (Solaris and derivatives)
- remove useless/misleading libc version display, #738
- test suite: reset exit code of persistent archiver, #844
- RemoteRepository: clean up pipe if remote open() fails
- Remote: don't print tracebacks for Error exceptions handled downstream, #792
- if `BORG_PASSPHRASE` is present but wrong, don't prompt for password, but fail instead, #791
- ArchiveChecker: move “orphaned objects check skipped” to INFO log level, #826

- fix capitalization, add ellipses, change log level to debug for 2 messages, #798

Other changes:

- update llfuse requirement, llfuse 1.0 works
- update OS / dist packages on build machines, #717
- prefer showing `-info` over `-v` in usage help, #859
- docs:
  - fix cygwin requirements (gcc-g++)
  - document how to debug / file filesystem issues, #664
  - fix reproducible build of api docs
  - RTD theme: CSS !important overwrite, #727
  - Document logo font. Recreate logo png. Remove GIMP logo file.

## 10.20 Version 1.0.0 (2016-03-05)

The major release number change (0.x -> 1.x) indicates bigger incompatible changes, please read the compatibility notes, adapt / test your scripts and check your backup logs.

Compatibility notes:

- drop support for python 3.2 and 3.3, require 3.4 or 3.5, #221 #65 #490 note: we provide binaries that include python 3.5.1 and everything else needed. they are an option in case you are stuck with < 3.4 otherwise.
- change encryption to be on by default (using “repokey” mode)
- moved keyfile keys from `~/.borg/keys` to `~/.config/borg/keys`, you can either move them manually or run “borg upgrade <REPO>”
- remove support for `-encryption=passphrase`, use `borg migrate-to-repokey` to switch to repokey mode, #97
- remove deprecated `-compression <number>`, use `-compression zlib,<number>` instead in case of 0, you could also use `-compression none`
- remove deprecated `-hourly/daily/weekly/monthly/yearly` use `-keep-hourly/daily/weekly/monthly/yearly` instead
- remove deprecated `-do-not-cross-mountpoints`, use `-one-file-system` instead
- disambiguate `-p` option, #563:
  - `-p` now is same as `-progress`
  - `-P` now is same as `-prefix`
- remove deprecated “borg verify”, use “borg extract `-dry-run`” instead
- cleanup environment variable semantics, #355 the environment variables used to be “yes sayers” when set, this was conceptually generalized to “automatic answerers” and they just give their value as answer (as if you typed in that value when being asked). See the “usage” / “Environment Variables” section of the docs for details.
- change the builtin default for `-chunker-params`, create 2MiB chunks, #343 `-chunker-params` new default: 19,23,21,4095 - old default: 10,23,16,4095

one of the biggest issues with borg < 1.0 (and also attic) was that it had a default target chunk size of 64kiB, thus it created a lot of chunks and thus also a huge chunk management overhead (high RAM and disk usage).

please note that the new default won't change the chunks that you already have in your repository. the new big chunks do not deduplicate with the old small chunks, so expect your repo to grow at least by the size of every changed file and in the worst case (e.g. if your files cache was lost / is not used) by the size of every file (minus any compression you might use).

in case you want to immediately see a much lower resource usage (RAM / disk) for chunks management, it might be better to start with a new repo than continuing in the existing repo (with an existing repo, you'd have to wait until all archives with small chunks got pruned to see a lower resource usage).

if you used the old `--chunker-params` default value (or if you did not use `--chunker-params` option at all) and you'd like to continue using small chunks (and you accept the huge resource usage that comes with that), just explicitly use `borg create --chunker-params=10,23,16,4095`.

- archive timestamps: the 'time' timestamp now refers to archive creation start time (was: end time), the new 'time\_end' timestamp refers to archive creation end time. This might affect prune if your backups take rather long. if you give a timestamp via cli this is stored into 'time', therefore it now needs to mean archive creation start time.

New features:

- implement password roundtrip, #695

Bug fixes:

- remote end does not need cache nor keys directories, do not create them, #701
- added retry counter for passwords, #703

Other changes:

- fix compiler warnings, #697
- docs:
  - update README.rst to new changelog location in docs/changes.rst
  - add Teemu to AUTHORS
  - changes.rst: fix old chunker params, #698
  - FAQ: how to limit bandwidth

## 10.21 Version 1.0.0rc2 (2016-02-28)

New features:

- format options for location: user, pid, fqdn, hostname, now, utcnow, user
- borg list `--list-format`
- borg prune `-v --list` enables the keep/prune list output, #658

Bug fixes:

- fix `_open_rb` noatime handling, #657
- add a simple archivename validator, #680
- borg create `--stats`: show timestamps in localtime, use same labels/formatting as borg info, #651
- llfuse compatibility fixes (now compatible with: 0.40, 0.41, 0.42)

Other changes:

- it is now possible to use “pip install borgbackup[fuse]” to automatically install the llfuse dependency using the correct version requirement for it. you still need to care about having installed the FUSE / build related OS package first, though, so that building llfuse can succeed.
- Vagrant: drop Ubuntu Precise (12.04) - does not have Python >= 3.4
- Vagrant: use pyinstaller v3.1.1 to build binaries
- docs:
  - borg upgrade: add to docs that only LOCAL repos are supported
  - borg upgrade also handles borg 0.xx -> 1.0
  - use pip extras or requirements file to install llfuse
  - fix order in release process
  - updated usage docs and other minor / cosmetic fixes
  - verified borg examples in docs, #644
  - freebsd dependency installation and fuse configuration, #649
  - add example how to restore a raw device, #671
  - add a hint about the dev headers needed when installing from source
  - add examples for delete (and handle delete after list, before prune), #656
  - update example for borg create -v --stats (use iso datetime format), #663
  - added example to BORG\_RSH docs
  - “connection closed by remote”: add FAQ entry and point to issue #636

## 10.22 Version 1.0.0rc1 (2016-02-07)

### New features:

- borg migrate-to-repokey (“passphrase” -> “repokey” encryption key mode)
- implement --short for borg list REPO, #611
- implement --list for borg extract (consistency with borg create)
- borg serve: overwrite client’s --restrict-to-path with ssh forced command’s option value (but keep everything else from the client commandline), #544
- use \$XDG\_CONFIG\_HOME/keys for keyfile keys (~/.config/borg/keys), #515
- “borg upgrade” moves the keyfile keys to the new location
- display both archive creation start and end time in “borg info”, #627

### Bug fixes:

- normalize trailing slashes for the repository path, #606
- Cache: fix exception handling in \_\_init\_\_, release lock, #610

### Other changes:

- suppress unneeded exception context (PEP 409), simpler tracebacks
- removed special code needed to deal with imperfections / incompatibilities / missing stuff in py 3.2/3.3, simplify code that can be done simpler in 3.4

- removed some version requirements that were kept on old versions because newer did not support py 3.2 any more
- use some py 3.4+ stdlib code instead of own/openssl/pypi code:
  - use os.urandom instead of own cython openssl RAND\_bytes wrapper, #493
  - use hashlib.pbkdf2\_hmac from py stdlib instead of own openssl wrapper
  - use hmac.compare\_digest instead of == operator (constant time comparison)
  - use stat.filemode instead of homegrown code
  - use “mock” library from stdlib, #145
  - remove borg.support (with non-broken argparse copy), it is ok in 3.4+, #358
- Vagrant: copy CHANGES.rst as symlink, #592
- cosmetic code cleanups, add flake8 to tox/travis, #4
- docs / help:
  - make “borg -h” output prettier, #591
  - slightly rephrase prune help
  - add missing example for –list option of borg create
  - quote exclude line that includes an asterisk to prevent shell expansion
  - fix dead link to license
  - delete Ubuntu Vivid, it is not supported anymore (EOL)
  - OS X binary does not work for older OS X releases, #629
  - borg serve’s special support for forced/original ssh commands, #544
  - misc. updates and fixes

## 10.23 Version 0.30.0 (2016-01-23)

### Compatibility notes:

- you may need to use -v (or –info) more often to actually see output emitted at INFO log level (because it is suppressed at the default WARNING log level). See the “general” section in the usage docs.
- for borg create, you need –list (additionally to -v) to see the long file list (was needed so you can have e.g. –stats alone without the long list)
- see below about BORG\_DELETE\_I\_KNOW\_WHAT\_I\_AM\_DOING (was: BORG\_CHECK\_I\_KNOW\_WHAT\_I\_AM\_DOING)

### Bug fixes:

- fix crash when using borg create –dry-run –keep-tag-files, #570
- make sure teardown with cleanup happens for Cache and RepositoryCache, avoiding leftover locks and TEMP dir contents, #285 (partially), #548
- fix locking KeyError, partial fix for #502
- log stats consistently, #526
- add abbreviated weekday to timestamp format, fixes #496

- strip whitespace when loading exclusions from file
- unset LD\_LIBRARY\_PATH before invoking ssh, fixes strange OpenSSL library version warning when using the borg binary, #514
- add some error handling/fallback for C library loading, #494
- added BORG\_DELETE\_I\_KNOW\_WHAT\_I\_AM\_DOING for check in “borg delete”, #503
- remove unused “repair” rpc method name

### New features:

- borg create: implement exclusions using regular expression patterns.
- borg create: implement inclusions using patterns.
- borg extract: support patterns, #361
- support different styles for patterns:
  - fnmatch (*fm*: prefix, default when omitted), like borg <= 0.29.
  - shell (*sh*: prefix) with \* not matching directory separators and \*\*/ matching 0..n directories
  - path prefix (*pp*: prefix, for unifying borg create pp1 pp2 into the patterns system), semantics like in borg <= 0.29
  - regular expression (*re*:), new!
- –progress option for borg upgrade (#291) and borg delete <archive>
- update progress indication more often (e.g. for borg create within big files or for borg check repo), #500
- finer chunker granularity for items metadata stream, #547, #487
- borg create –list now used (additionally to -v) to enable the verbose file list output
- display borg version below tracebacks, #532

### Other changes:

- hashtable size (and thus: RAM and disk consumption) follows a growth policy: grows fast while small, grows slower when getting bigger, #527
- Vagrantfile: use pyinstaller 3.1 to build binaries, freebsd sqlite3 fix, fixes #569
- no separate binaries for centos6 any more because the generic linux binaries also work on centos6 (or in general: on systems with a slightly older glibc than debian7)
- dev environment: require virtualenv<14.0 so we get a py32 compatible pip
- docs:
  - add space-saving chunks.archive.d trick to FAQ
  - important: clarify -v and log levels in usage -> general, please read!
  - sphinx configuration: create a simple man page from usage docs
  - add a repo server setup example
  - disable unneeded SSH features in authorized\_keys examples for security.
  - borg prune only knows “–keep-within” and not “–within”
  - add gource video to resources docs, #507
  - add netbsd install instructions

- authors: make it more clear what refers to borg and what to attic
- document standalone binary requirements, #499
- rephrase the mailing list section
- development docs: run build\_api and build\_usage before tagging release
- internals docs: hash table max. load factor is 0.75 now
- markup, typo, grammar, phrasing, clarifications and other fixes.
- add gcc gcc-c++ to redhat/fedora/corora install docs, fixes #583

## 10.24 Version 0.29.0 (2015-12-13)

### Compatibility notes:

- when upgrading to 0.29.0 you need to upgrade client as well as server installations due to the locking and commandline interface changes otherwise you'll get an error msg about a RPC protocol mismatch or a wrong commandline option. if you run a server that needs to support both old and new clients, it is suggested that you have a "borg-0.28.2" and a "borg-0.29.0" command. clients then can choose via e.g. "borg --remote-path=borg-0.29.0 ...".
- the default waiting time for a lock changed from infinity to 1 second for a better interactive user experience. if the repo you want to access is currently locked, borg will now terminate after 1s with an error message. if you have scripts that shall wait for the lock for a longer time, use --lock-wait N (with N being the maximum wait time in seconds).

### Bug fixes:

- hash table tuning (better chosen hashtable load factor 0.75 and prime initial size of 1031 gave ~1000x speedup in some scenarios)
- avoid creation of an orphan lock for one case, #285
- --keep-tag-files: fix file mode and multiple tag files in one directory, #432
- fixes for "borg upgrade" (attic repo converter), #466
- remove --progress isatty magic (and also --no-progress option) again, #476
- borg init: display proper repo URL
- fix format of umask in help pages, #463

### New features:

- implement --lock-wait, support timeout for UpgradableLock, #210
- implement borg break-lock command, #157
- include system info below traceback, #324
- sane remote logging, remote stderr, #461:
  - remote log output: intercept it and log it via local logging system, with "Remote: " prefixed to message. log remote tracebacks.
  - remote stderr: output it to local stderr with "Remote: " prefixed.
- add --debug and --info (same as --verbose) to set the log level of the builtin logging configuration (which otherwise defaults to warning), #426 note: there are few messages emitted at DEBUG level currently.
- optionally configure logging via env var BORG\_LOGGING\_CONF

- add `-filter` option for status characters: e.g. to show only the added or modified files (and also errors), use “`borg create -v -filter=AME ...`”.
- more progress indicators, #394
- use ISO-8601 date and time format, #375
- “`borg check -prefix`” to restrict archive checking to that name prefix, #206

Other changes:

- `hashindex_add` C implementation (speed up cache re-sync for new archives)
- increase FUSE `read_size` to 1024 (speed up metadata operations)
- `check/delete/prune -save-space`: free unused segments quickly, #239
- increase rpc protocol version to 2 (see also Compatibility notes), #458
- silence borg by default (via default log level WARNING)
- get rid of C compiler warnings, #391
- upgrade OS X FUSE to 3.0.9 on the OS X binary build system
- use python 3.5.1 to build binaries
- docs:
  - new mailing list [borgbackup@python.org](mailto:borgbackup@python.org), #468
  - `readthedocs`: color and logo improvements
  - load coverage icons over SSL (avoids mixed content)
  - more precise binary installation steps
  - update release procedure docs about OS X FUSE
  - FAQ entry about unexpected ‘A’ status for unchanged file(s), #403
  - add docs about ‘E’ file status
  - add “borg upgrade” docs, #464
  - add developer docs about output and logging
  - clarify encryption, add note about client-side encryption
  - add resources section, with videos, talks, presentations, #149
  - Borg moved to Arch Linux [community]
  - fix wrong installation instructions for archlinux

## 10.25 Version 0.28.2 (2015-11-15)

New features:

- `borg create -exclude-if-present TAGFILE` - exclude directories that have the given file from the backup. You can additionally give `-keep-tag-files` to preserve just the directory roots and the tag-files (but not backup other directory contents), #395, attic #128, attic #142

Other changes:

- do not create docs sources at build time (just have them in the repo), completely remove `have_cython()` hack, do not use the “mock” library at build time, #384



- avoid hidden import, make it easier for PyInstaller, easier fix for #218
- docs:
  - add description of item flags / status output, fixes #402
  - explain how to regenerate usage and API files (build\_api or build\_usage) and when to commit usage files directly into git, #384
  - minor install docs improvements

## 10.26 Version 0.28.1 (2015-11-08)

Bug fixes:

- do not try to build api / usage docs for production install, fixes unexpected “mock” build dependency, #384

Other changes:

- avoid using msgpack.packb at import time
- fix formatting issue in changes.rst
- fix build on readthedocs

## 10.27 Version 0.28.0 (2015-11-08)

Compatibility notes:

- changed return codes (exit codes), see docs. in short: old: 0 = ok, 1 = error. now: 0 = ok, 1 = warning, 2 = error

New features:

- refactor return codes (exit codes), fixes #61
- add `--show-rc` option enable “terminating with X status, rc N” output, fixes 58, #351
- borg create backups atime and ctime additionally to mtime, fixes #317 - extract: support atime additionally to mtime - FUSE: support ctime and atime additionally to mtime
- support borg `--version`
- emit a warning if we have a slow msgpack installed
- borg list `--prefix=thishostname- REPO`, fixes #205
- Debug commands (do not use except if you know what you do: debug-get-obj, debug-put-obj, debug-delete-obj, debug-dump-archive-items).

Bug fixes:

- setup.py: fix bug related to BORG\_LZ4\_PREFIX processing
- fix “check” for repos that have incomplete chunks, fixes #364
- borg mount: fix unlocking of repository at umount time, fixes #331
- fix reading files without touching their atime, #334
- non-ascii ACL fixes for Linux, FreeBSD and OS X, #277
- fix `acl_use_local_uid_gid()` and add a test for it, attic #359
- borg upgrade: do not upgrade repositories in place by default, #299

- fix cascading failure with the index conversion code, #269
- borg check: implement ‘cmdline’ archive metadata value decoding, #311
- fix RobustUnpacker, it missed some metadata keys (new atime and ctime keys were missing, but also bsdfags). add check for unknown metadata keys.
- create from stdin: also save atime, ctime (cosmetic)
- use default\_notty=False for confirmations, fixes #345
- vagrant: fix msgpack installation on centos, fixes #342
- deal with unicode errors for symlinks in same way as for regular files and have a helpful warning message about how to fix wrong locale setup, fixes #382
- add ACL keys the RobustUnpacker must know about

Other changes:

- improve file size displays, more flexible size formatters
- explicitly commit to the units standard, #289
- archiver: add E status (means that an error occurred when processing this (single) item)
- do binary releases via “github releases”, closes #214
- create: use -x and --one-file-system (was: --do-not-cross-mountpoints), #296
- a lot of changes related to using “logging” module and screen output, #233
- show progress display if on a tty, output more progress information, #303
- factor out status output so it is consistent, fix surrogates removal, maybe fixes #309
- move away from RawConfigParser to ConfigParser
- archive checker: better error logging, give chunk\_id and sequence numbers (can be used together with borg debug-dump-archive-items).
- do not mention the deprecated passphrase mode
- emit a deprecation warning for --compression N (giving a just a number)
- misc .coveragerc fixes (and coverage measurement improvements), fixes #319
- refactor confirmation code, reduce code duplication, add tests
- prettier error messages, fixes #307, #57
- tests:
  - add a test to find disk-full issues, #327
  - travis: also run tests on Python 3.5
  - travis: use tox -r so it rebuilds the tox environments
  - test the generated pyinstaller-based binary by archiver unit tests, #215
  - vagrant: tests: announce whether fakeroot is used or not
  - vagrant: add vagrant user to fuse group for debianoid systems also
  - vagrant: llfuse install on darwin needs pkgconfig installed
  - vagrant: use pyinstaller from develop branch, fixes #336
  - benchmarks: test create, extract, list, delete, info, check, help, fixes #146

- benchmarks: test with both the binary and the python code
- archiver tests: test with both the binary and the python code, fixes #215
- make basic test more robust
- docs:
  - moved docs to borgbackup.readthedocs.org, #155
  - a lot of fixes and improvements, use mobile-friendly RTD standard theme
  - use zlib,6 compression in some examples, fixes #275
  - add missing rename usage to docs, closes #279
  - include the help offered by borg help <topic> in the usage docs, fixes #293
  - include a list of major changes compared to attic into README, fixes #224
  - add OS X install instructions, #197
  - more details about the release process, #260
  - fix linux glibc requirement (binaries built on debian7 now)
  - build: move usage and API generation to setup.py
  - update docs about return codes, #61
  - remove api docs (too much breakage on rtd)
  - borgbackup install + basics presentation (asciinema)
  - describe the current style guide in documentation
  - add section about debug commands
  - warn about not running out of space
  - add example for rename
  - improve chunker params docs, fixes #362
  - minor development docs update

## 10.28 Version 0.27.0 (2015-10-07)

New features:

- “borg upgrade” command - attic -> borg one time converter / migration, #21
- temporary hack to avoid using lots of disk space for chunks.archive.d, #235: To use it: `rm -rf chunks.archive.d ; touch chunks.archive.d`
- respect XDG\_CACHE\_HOME, attic #181
- add support for arbitrary SSH commands, attic #99
- borg delete -cache-only REPO (only delete cache, not REPO), attic #123

Bug fixes:

- use Debian 7 (wheezy) to build pyinstaller borgbackup binaries, fixes slow down observed when running the Centos6-built binary on Ubuntu, #222
- do not crash on empty lock.roster, fixes #232

- fix multiple issues with the cache config version check, #234
- fix segment entry header size check, attic #352 plus other error handling improvements / code deduplication there.
- always give segment and offset in repo IntegrityErrors

Other changes:

- stop producing binary wheels, remove docs about it, #147
- docs: - add warning about prune - generate usage include files only as needed - development docs: add Vagrant section - update / improve / reformat FAQ - hint to single-file pyinstaller binaries from README

## 10.29 Version 0.26.1 (2015-09-28)

This is a minor update, just docs and new pyinstaller binaries.

- docs update about python and binary requirements
- better docs for `--read-special`, fix #220
- re-built the binaries, fix #218 and #213 (glibc version issue)
- update web site about single-file pyinstaller binaries

Note: if you did a python-based installation, there is no need to upgrade.

## 10.30 Version 0.26.0 (2015-09-19)

New features:

- Faster cache sync (do all in one pass, remove tar/compression stuff), #163
- `BORG_REPO` env var to specify the default repo, #168
- read special files as if they were regular files, #79
- implement borg create `--dry-run`, attic issue #267
- Normalize paths before pattern matching on OS X, #143
- support OpenBSD and NetBSD (except xattrs/ACLs)
- support / run tests on Python 3.5

Bug fixes:

- borg mount repo: use absolute path, attic #200, attic #137
- chunker: use `off_t` to get 64bit on 32bit platform, #178
- initialize chunker fd to -1, so it's not equal to `STDIN_FILENO` (0)
- fix reaction to "no" answer at delete repo prompt, #182
- setup.py: detect lz4.h header file location
- to support python < 3.2.4, add less buggy argparse lib from 3.2.6 (#194)
- fix for obtaining `char *` from temporary Python value (old code causes a compile error on Mint 17.2)
- llfuse 0.41 install troubles on some platforms, require < 0.41 (UnicodeDecodeError exception due to non-ascii llfuse setup.py)

- cython code: add some int types to get rid of unspecific python add / subtract operations (avoid undefined symbol FPE... error on some platforms)
- fix verbose mode display of stdin backup
- extract: warn if a include pattern never matched, fixes #209, implement counters for Include/ExcludePatterns
- archive names with slashes are invalid, attic issue #180
- chunker: add a check whether the POSIX\_FADV\_DONTNEED constant is defined - fixes building on OpenBSD.

Other changes:

- detect inconsistency / corruption / hash collision, #170
- replace versioneer with setuptools\_scm, #106
- docs:
  - pkg-config is needed for llfuse installation
  - be more clear about pruning, attic issue #132
- unit tests:
  - xattr: ignore security.selinux attribute showing up
  - ext3 seems to need a bit more space for a sparse file
  - do not test lzma level 9 compression (avoid MemoryError)
  - work around strange mtime granularity issue on netbsd, fixes #204
  - ignore st\_rdev if file is not a block/char device, fixes #203
  - stay away from the setgid and sticky mode bits
- use Vagrant to do easy cross-platform testing (#196), currently:
  - Debian 7 “wheezy” 32bit, Debian 8 “jessie” 64bit
  - Ubuntu 12.04 32bit, Ubuntu 14.04 64bit
  - Centos 7 64bit
  - FreeBSD 10.2 64bit
  - OpenBSD 5.7 64bit
  - NetBSD 6.1.5 64bit
  - Darwin (OS X Yosemite)

## 10.31 Version 0.25.0 (2015-08-29)

Compatibility notes:

- lz4 compression library (liblz4) is a new requirement (#156)
- the new compression code is very compatible: as long as you stay with zlib compression, older borg releases will still be able to read data from a repo/archive made with the new code (note: this is not the case for the default “none” compression, use “zlib,0” if you want a “no compression” mode that can be read by older borg). Also the new code is able to read repos and archives made with older borg versions (for all zlib levels 0..9).

Deprecations:

- `--compression N` (with `N` being a number, as in 0.24) is deprecated. We keep the `--compression 0..9` for now to not break scripts, but it is deprecated and will be removed later, so better fix your scripts now: `--compression 0` (as in 0.24) is the same as `--compression zlib,0` (now). BUT: if you do not want compression, you rather want `--compression none` (which is the default). `--compression 1` (in 0.24) is the same as `--compression zlib,1` (now) `--compression 9` (in 0.24) is the same as `--compression zlib,9` (now)

### New features:

- `create --compression none` (default, means: do not compress, just pass through data “as is”. this is more efficient than `zlib level 0` as used in borg 0.24)
- `create --compression lz4` (super-fast, but not very high compression)
- `create --compression zlib,N` (slower, higher compression, default for `N` is 6)
- `create --compression lzma,N` (slowest, highest compression, default `N` is 6)
- honor the `nodump` flag (`UF_NODUMP`) and do not backup such items
- `list --short` just outputs a simple list of the files/directories in an archive

### Bug fixes:

- fixed `--chunker-params` parameter order confusion / malfunction, fixes #154
- close fds of segments we delete (during compaction)
- close files which fell out the lru-cache
- `fsync DONTNEED` now is only called for the byte range actually read, not for the whole file, fixes #158.
- fix issue with negative “all archives” size, fixes #165
- `restore_xattrs`: ignore if `setxattr` fails with `EACCES`, fixes #162

### Other changes:

- remove fakeroot requirement for tests, tests run faster without fakeroot (test setup does not fail any more without fakeroot, so you can run with or without fakeroot), fixes #151 and #91.
- more tests for archiver
- `recover_segment()`: don’t assume we have an fd for segment
- lru-cache refactoring / cleanup, add `dispose` function, `py.test` tests
- generalize `hashindex` code for any key length (less hardcoding)
- lock roster: catch file not found in `remove()` method and ignore it
- travis CI: use requirements file
- improved docs:
  - replace hack for `libfuse` with proper solution (install `libfuse-dev`)
  - update docs about compression
  - update development docs about fakeroot
  - internals: add some words about lock files / locking system
  - support: mention `BountySource` and for what it can be used
  - theme: use a lighter green
  - add `pypi`, `wheel`, `dist` package based install docs
  - split install docs into system-specific preparations and generic instructions

## 10.32 Version 0.24.0 (2015-08-09)

Incompatible changes (compared to 0.23):

- borg now always issues `-umask NNN` option when invoking another borg via ssh on the repository server. By that, it's making sure it uses the same umask for remote repos as for local ones. Because of this, you must upgrade both server and client(s) to 0.24.
- the default umask is `077` now (if you do not specify via `-umask`) which might be a different one as you used previously. The default umask avoids that you accidentally give access permissions for group and/or others to files created by borg (e.g. the repository).

Deprecations:

- “`-encryption passphrase`” mode is deprecated, see #85 and #97. See the new “`-encryption repokey`” mode for a replacement.

New features:

- borg create `-chunker-params ...` to configure the chunker, fixes #16 (attic #302, attic #300, and somehow also #41). This can be used to reduce memory usage caused by chunk management overhead, so borg does not create a huge chunks index/repo index and eats all your RAM if you back up lots of data in huge files (like VM disk images). See docs/misc/create\_chunker-params.txt for more information.
- borg info now reports chunk counts in the chunk index.
- borg create `-compression 0..9` to select zlib compression level, fixes #66 (attic #295).
- borg init `-encryption repokey` (to store the encryption key into the repo), fixes #85
- improve at-end error logging, always log exceptions and set `exit_code=1`
- LoggedIO: better error checks / exceptions / exception handling
- implement `-remote-path` to allow non-default-path borg locations, #125
- implement `-umask M` and use `077` as default umask for better security, #117
- borg check: give a named single archive to it, fixes #139
- cache sync: show progress indication
- cache sync: reimplement the chunk index merging in C

Bug fixes:

- fix segfault that happened for unreadable files (chunker: `n` needs to be a signed `size_t`), #116
- fix the repair mode, #144
- repo delete: add `destroy` to allowed rpc methods, fixes issue #114
- more compatible repository locking code (based on `mkdir`), maybe fixes #92 (attic #317, attic #201).
- better Exception msg if no Borg is installed on the remote repo server, #56
- create a `RepositoryCache` implementation that can cope with `>2GiB`, fixes attic #326.
- fix `Traceback` when running `check -repair`, attic #232
- clarify help text, fixes #73.
- add help string for `-no-files-cache`, fixes #140

Other changes:

- improved docs:

- added docs/misc directory for misc. writeups that won't be included "as is" into the html docs.
  - document environment variables and return codes (attic #324, attic #52)
  - web site: add related projects, fix web site url, IRC #borgbackup
  - Fedora/Fedora-based install instructions added to docs
  - Cygwin-based install instructions added to docs
  - updated AUTHORS
  - add FAQ entries about redundancy / integrity
  - clarify that borg extract uses the cwd as extraction target
  - update internals doc about chunker params, memory usage and compression
  - added docs about development
  - add some words about resource usage in general
  - document how to backup a raw disk
  - add note about how to run borg from virtual env
  - add solutions for (ll)fuse installation problems
  - document what borg check does, fixes #138
  - reorganize borgbackup.github.io sidebar, prev/next at top
  - deduplicate and refactor the docs / README.rst
- use borg-tmp as prefix for temporary files / directories
  - short prune options without "keep-" are deprecated, do not suggest them
  - improved tox configuration
  - remove usage of unittest.mock, always use mock from pypi
  - use entrypoints instead of scripts, for better use of the wheel format and modern installs
  - add requirements.d/development.txt and modify tox.ini
  - use travis-ci for testing based on Linux and (new) OS X
  - use coverage.py, pytest-cov and codecov.io for test coverage support

I forgot to list some stuff already implemented in 0.23.0, here they are:

New features:

- efficient archive list from manifest, meaning a big speedup for slow repo connections and "list <repo>", "delete <repo>", "prune" (attic #242, attic #167)
- big speedup for chunks cache sync (esp. for slow repo connections), fixes #18
- hashindex: improve error messages

Other changes:

- explicitly specify binary mode to open binary files
- some easy micro optimizations



## 10.33 Version 0.23.0 (2015-06-11)

Incompatible changes (compared to attic, fork related):

- changed sw name and cli command to “borg”, updated docs
- package name (and name in urls) uses “borgbackup” to have less collisions
- changed repo / cache internal magic strings from ATTIC\* to BORG\*, changed cache location to .cache/borg/ - this means that it currently won’t accept attic repos (see issue #21 about improving that)

Bug fixes:

- avoid defect python-msgpack releases, fixes attic #171, fixes attic #185
- fix traceback when trying to do unsupported passphrase change, fixes attic #189
- datetime does not like the year 10.000, fixes attic #139
- fix “info” all archives stats, fixes attic #183
- fix parsing with missing microseconds, fixes attic #282
- fix misleading hint the fuse ImportError handler gave, fixes attic #237
- check unpacked data from RPC for tuple type and correct length, fixes attic #127
- fix Repository.\_active\_txn state when lock upgrade fails
- give specific path to xattr.is\_enabled(), disable symlink setattr call that always fails
- fix test setup for 32bit platforms, partial fix for attic #196
- upgraded versioneer, PEP440 compliance, fixes attic #257

New features:

- less memory usage: add global option `--no-cache-files`
- check `--last N` (only check the last N archives)
- check: sort archives in reverse time order
- rename `repo::oldname newname` (rename repository)
- create `-v` output more informative
- create `--progress` (backup progress indicator)
- create `--timestamp` (utc string or reference file/dir)
- create: if “-” is given as path, read binary from stdin
- extract: if `--stdout` is given, write all extracted binary data to stdout
- extract `--sparse` (simple sparse file support)
- extra debug information for ‘fread failed’
- delete `<repo>` (deletes whole repo + local cache)
- FUSE: reflect deduplication in allocated blocks
- only allow whitelisted RPC calls in server mode
- normalize source/exclude paths before matching
- use `posix_fadvise` to not spoil the OS cache, fixes attic #252
- toplevel error handler: show tracebacks for better error analysis

- sigusr1 / sigint handler to print current file infos - attic PR #286
- RPCError: include the exception args we get from remote

Other changes:

- source: misc. cleanups, pep8, style
- docs and faq improvements, fixes, updates
- cleanup crypto.pyx, make it easier to adapt to other AES modes
- do os.fsync like recommended in the python docs
- source: Let chunker optionally work with os-level file descriptor.
- source: Linux: remove duplicate os.fsencode calls
- source: refactor `_open_rb` code a bit, so it is more consistent / regular
- source: refactor indicator (status) and item processing
- source: use `py.test` for better testing, `flake8` for code style checks
- source: fix `tox >=2.0` compatibility (test runner)
- pypi package: add python version classifiers, add FreeBSD to platforms

## 10.34 Attic Changelog

Here you can see the full list of changes between each Attic release until Borg forked from Attic:

### 10.34.1 Version 0.17

(bugfix release, released on X)

- Fix hashindex ARM memory alignment issue (#309)
- Improve hashindex error messages (#298)

### 10.34.2 Version 0.16

(bugfix release, released on May 16, 2015)

- Fix typo preventing the security confirmation prompt from working (#303)
- Improve handling of systems with improperly configured file system encoding (#289)
- Fix “All archives” output for attic info. (#183)
- More user friendly error message when repository key file is not found (#236)
- Fix parsing of iso 8601 timestamps with zero microseconds (#282)

### 10.34.3 Version 0.15

(bugfix release, released on Apr 15, 2015)

- `xattr`: Be less strict about unknown/unsupported platforms (#239)
- Reduce repository listing memory usage (#163).

- Fix BrokenPipeError for remote repositories (#233)
- Fix incorrect behavior with two character directory names (#265, #268)
- Require approval before accessing relocated/moved repository (#271)
- Require approval before accessing previously unknown unencrypted repositories (#271)
- Fix issue with hash index files larger than 2GB.
- Fix Python 3.2 compatibility issue with noatime open() (#164)
- Include missing pyx files in dist files (#168)

### 10.34.4 Version 0.14

(feature release, released on Dec 17, 2014)

- Added support for stripping leading path segments (#95) “attic extract –strip-segments X”
- Add workaround for old Linux systems without acl\_extended\_file\_no\_follow (#96)
- Add MacPorts’ path to the default openssl search path (#101)
- HashIndex improvements, eliminates unnecessary IO on low memory systems.
- Fix “Number of files” output for attic info. (#124)
- limit create file permissions so files aren’t read while restoring
- Fix issue with empty xattr values (#106)

### 10.34.5 Version 0.13

(feature release, released on Jun 29, 2014)

- Fix sporadic “Resource temporarily unavailable” when using remote repositories
- Reduce file cache memory usage (#90)
- Faster AES encryption (utilizing AES-NI when available)
- Experimental Linux, OS X and FreeBSD ACL support (#66)
- Added support for backup and restore of BSDFlags (OSX, FreeBSD) (#56)
- Fix bug where xattrs on symlinks were not correctly restored
- Added cachedir support. CACHEDIR.TAG compatible cache directories can now be excluded using --exclude-caches (#74)
- Fix crash on extreme mtime timestamps (year 2400+) (#81)
- Fix Python 3.2 specific lockf issue (EDEADLK)

### 10.34.6 Version 0.12

(feature release, released on April 7, 2014)

- Python 3.4 support (#62)
- Various documentation improvements a new style
- attic mount now supports mounting an entire repository not only individual archives (#59)

- Added option to restrict remote repository access to specific path(s): `attic serve --restrict-to-path X` (#51)
- Include “all archives” size information in “-stats” output. (#54)
- Added `--stats` option to `attic delete` and `attic prune`
- Fixed bug where `attic prune` used UTC instead of the local time zone when determining which archives to keep.
- Switch to SI units (Power of 1000 instead 1024) when printing file sizes

### 10.34.7 Version 0.11

(feature release, released on March 7, 2014)

- New “check” command for repository consistency checking (#24)
- Documentation improvements
- Fix exception during “attic create” with repeated files (#39)
- New “-exclude-from” option for `attic create/extract/verify`.
- Improved archive metadata deduplication.
- “attic verify” has been deprecated. Use “attic extract -dry-run” instead.
- “attic prune -hourly/dailyl...” has been deprecated. Use “attic prune -keep-hourly/dailyl...” instead.
- Ignore `xattr` errors during “extract” if not supported by the filesystem. (#46)

### 10.34.8 Version 0.10

(bugfix release, released on Jan 30, 2014)

- Fix deadlock when extracting 0 sized files from remote repositories
- “-exclude” wildcard patterns are now properly applied to the full path not just the file name part (#5).
- Make source code endianness agnostic (#1)

### 10.34.9 Version 0.9

(feature release, released on Jan 23, 2014)

- Remote repository speed and reliability improvements.
- Fix sorting of segment names to ignore NFS left over files. (#17)
- Fix incorrect display of time (#13)
- Improved error handling / reporting. (#12)
- Use `fcntl()` instead of `flock()` when locking repository/cache. (#15)
- Let `ssh` figure out port/user if not specified so we don’t override `.ssh/config` (#9)
- Improved `libcrypto` path detection (#23).

### 10.34.10 Version 0.8.1

(bugfix release, released on Oct 4, 2013)

- Fix segmentation fault issue.

### 10.34.11 Version 0.8

(feature release, released on Oct 3, 2013)

- Fix xattr issue when backing up sshfs filesystems (#4)
- Fix issue with excessive index file size (#6)
- Support access of read only repositories.
- **New syntax to enable repository encryption:** `attic init --encryption="none|passphrase|keyfile"`.
- Detect and abort if repository is older than the cache.

### 10.34.12 Version 0.7

(feature release, released on Aug 5, 2013)

- Ported to FreeBSD
- Improved documentation
- Experimental: Archives mountable as fuse filesystems.
- The “user.” prefix is no longer stripped from xattrs on Linux

### 10.34.13 Version 0.6.1

(bugfix release, released on July 19, 2013)

- Fixed an issue where mtime was not always correctly restored.

### 10.34.14 Version 0.6

First public release on July 9, 2013



This page documents the internal data structures and storage mechanisms of Borg. It is partly based on [mailing list discussion about internals](#)<sup>98</sup> and also on static code analysis.

## 11.1 Repository and Archives

Borg stores its data in a *Repository*. Each repository can hold multiple *Archives*, which represent individual backups that contain a full archive of the files specified when the backup was performed. Deduplication is performed across multiple backups, both on data and metadata, using *Chunks* created by the chunker using the [Buzhash](#)<sup>99</sup> algorithm.

Each repository has the following file structure:

**README** simple text file telling that this is a Borg repository

**config** repository configuration

**data/** directory where the actual data is stored

**hints.%d** hints for repository compaction

**index.%d** repository index

**lock.roster and lock.exclusive/\*** used by the locking system to manage shared and exclusive locks

## 11.2 Lock files

Borg uses locks to get (exclusive or shared) access to the cache and the repository.

The locking system is based on creating a directory *lock.exclusive* (for exclusive locks). Inside the lock directory, there is a file indicating hostname, process id and thread id of the lock holder.

There is also a json file *lock.roster* that keeps a directory of all shared and exclusive lockers.

---

<sup>98</sup> <http://librelist.com/browser/attic/2014/5/6/questions-and-suggestions-about-inner-working-of-attic/T1/textgreater{}>

<sup>99</sup> <https://en.wikipedia.org/wiki/Buzhash>

If the process can create the *lock.exclusive* directory for a resource, it has the lock for it. If creation fails (because the directory has already been created by some other process), lock acquisition fails.

The cache lock is usually in *~/.cache/borg/REPOID/lock.\**. The repository lock is in *repository/lock.\**.

In case you run into troubles with the locks, you can use the `borg break-lock` command after you first have made sure that no Borg process is running on any machine that accesses this resource. Be very careful, the cache or repository might get damaged if multiple processes use it at the same time.

## 11.3 Config file

Each repository has a `config` file which is a INI-style file and looks like this:

```
[repository]
version = 1
segments_per_dir = 10000
max_segment_size = 5242880
id = 57d6c1d52ce76a836b532b0e42e677dec6af9fca3673db511279358828a21ed6
```

This is where the `repository.id` is stored. It is a unique identifier for repositories. It will not change if you move the repository around so you can make a local transfer then decide to move the repository to another (even remote) location at a later time.

## 11.4 Keys

The key to address the key/value store is usually computed like this:

`key = id = id_hash(unencrypted_data)`

The `id_hash` function is:

- sha256 (no encryption keys available)
- hmac-sha256 (encryption keys available)

As the `id` / `key` is used for deduplication, `id_hash` must be a cryptographically strong hash or MAC.

## 11.5 Segments and archives

A Borg repository is a filesystem based transactional key/value store. It makes extensive use of `msgpack`<sup>100</sup> to store data and, unless otherwise noted, data is stored in `msgpack`<sup>101</sup> encoded files.

Objects referenced by a key are stored inline in files (*segments*) of approx. 5MB size in numbered subdirectories of `repo/data`.

They contain:

- header size
- crc
- size
- tag

---

<sup>100</sup> <https://msgpack.org/>

<sup>101</sup> <https://msgpack.org/>



- key
- data

Segments are built locally, and then uploaded. Those files are strictly append-only and modified only once.

Tag is either `PUT`, `DELETE`, or `COMMIT`. A segment file is basically a transaction log where each repository operation is appended to the file. So if an object is written to the repository a `PUT` tag is written to the file followed by the object id and data. If an object is deleted a `DELETE` tag is appended followed by the object id. A `COMMIT` tag is written when a repository transaction is committed. When a repository is opened any `PUT` or `DELETE` operations not followed by a `COMMIT` tag are discarded since they are part of a partial/uncommitted transaction.

## 11.6 The manifest

The manifest is an object with an all-zero key that references all the archives. It contains:

- version
- list of archive infos
- timestamp
- config

Each archive info contains:

- name
- id
- time

It is the last object stored, in the last segment, and is replaced each time.

## 11.7 The Archive

The archive metadata does not contain the file items directly. Only references to other objects that contain that data. An archive is an object that contains:

- version
- name
- list of chunks containing item metadata (size: count \* ~40B)
- cmdline
- hostname
- username
- time

### 11.7.1 Note about archive limitations

The archive is currently stored as a single object in the repository and thus limited in size to `MAX_OBJECT_SIZE` (20MiB).

As one chunk list entry is ~40B, that means we can reference ~500.000 item metadata stream chunks per archive.

Each item metadata stream chunk is ~128kiB (see hardcoded `ITEMS_CHUNKER_PARAMS`).

So that means the whole item metadata stream is limited to ~64GiB chunks. If compression is used, the amount of storable metadata is bigger - by the compression factor.

If the medium size of an item entry is 100B (small size file, no ACLs/xattrs), that means a limit of ~640 million files/directories per archive.

If the medium size of an item entry is 2kB (~100MB size files or more ACLs/xattrs), the limit will be ~32 million files/directories per archive.

If one tries to create an archive object bigger than `MAX_OBJECT_SIZE`, a fatal `IntegrityError` will be raised.

A workaround is to create multiple archives with less items each, see also [#1452](#)<sup>102</sup>.

## 11.8 The Item

Each item represents a file, directory or other fs item and is stored as an `item` dictionary that contains:

- path
- list of data chunks (size: count \* ~40B)
- user
- group
- uid
- gid
- mode (item type + permissions)
- source (for links)
- rdev (for devices)
- mtime, atime, ctime in nanoseconds
- xattrs
- acl
- bsdfiles

All items are serialized using `msgpack` and the resulting byte stream is fed into the same chunker algorithm as used for regular file data and turned into deduplicated chunks. The reference to these chunks is then added to the archive metadata. To achieve a finer granularity on this metadata stream, we use different chunker params for this chunker, which result in smaller chunks.

A chunk is stored as an object as well, of course.

## 11.9 Chunks

The Borg chunker uses a rolling hash computed by the [Buzhash](#)<sup>103</sup> algorithm. It triggers (chunks) when the last `HASH_MASK_BITS` bits of the hash are zero, producing chunks of  $2^{\text{HASH\_MASK\_BITS}}$  Bytes on average.

Buzhash is **only** used for cutting the chunks at places defined by the content, the buzhash value is **not** used as the deduplication criteria (we use a cryptographically strong hash/MAC over the chunk contents for this, the `id_hash`).

---

<sup>102</sup> <https://github.com/borgbackup/borg/issues/1452>

<sup>103</sup> <https://en.wikipedia.org/wiki/Buzhash>

`borg create --chunker-params CHUNK_MIN_EXP,CHUNK_MAX_EXP,HASH_MASK_BITS,HASH_WINDOW_SIZE` can be used to tune the chunker parameters, the default is:

- `CHUNK_MIN_EXP = 19` (minimum chunk size =  $2^{19}$  B = 512 kiB)
- `CHUNK_MAX_EXP = 23` (maximum chunk size =  $2^{23}$  B = 8 MiB)
- `HASH_MASK_BITS = 21` (statistical medium chunk size  $\sim 2^{21}$  B = 2 MiB)
- `HASH_WINDOW_SIZE = 4095 [B]` (`0xFFF`)

The buzhash table is altered by XORing it with a seed randomly generated once for the repository, and stored encrypted in the keyfile. This is to prevent chunk size based fingerprinting attacks on your encrypted repo contents (to guess what files you have based on a specific set of chunk sizes).

For some more general usage hints see also `--chunker-params`.

## 11.10 Indexes / Caches

The **files cache** is stored in `cache/files` and is used at backup time to quickly determine whether a given file is unchanged and we have all its chunks.

The files cache is a key -> value mapping and contains:

- key:
  - full, absolute file path `id_hash`
- value:
  - file inode number
  - file size
  - file `mtime_ns`
  - list of file content chunk id hashes
  - age (0 [newest], 1, 2, 3, ..., `BORG_FILES_CACHE_TTL - 1`)

To determine whether a file has not changed, cached values are looked up via the key in the mapping and compared to the current file attribute values.

If the file's size, `mtime_ns` and inode number is still the same, it is considered to not have changed. In that case, we check that all file content chunks are (still) present in the repository (we check that via the chunks cache).

If everything is matching and all chunks are present, the file is not read / chunked / hashed again (but still a file metadata item is written to the archive, made from fresh file metadata read from the filesystem). This is what makes borg so fast when processing unchanged files.

If there is a mismatch or a chunk is missing, the file is read / chunked / hashed. Chunks already present in repo won't be transferred to repo again.

The inode number is stored and compared to make sure we distinguish between different files, as a single path may not be unique across different archives in different setups.

Not all filesystems have stable inode numbers. If that is the case, borg can be told to ignore the inode number in the check via `-ignore-inode`.

The age value is used for cache management. If a file is "seen" in a backup run, its age is reset to 0, otherwise its age is incremented by one. If a file was not seen in `BORG_FILES_CACHE_TTL` backups, its cache entry is removed. See also: *It always chunks all my files, even unchanged ones!* and *I am seeing 'A' (added) status for a unchanged file!?*

The files cache is a python dictionary, storing python objects, which generates a lot of overhead.

Borg can also work without using the files cache (saves memory if you have a lot of files or not much RAM free), then all files are assumed to have changed. This is usually much slower than with files cache.

The **chunks cache** is stored in `cache/chunks` and is used to determine whether we already have a specific chunk, to count references to it and also for statistics.

The chunks cache is a key -> value mapping and contains:

- key:
  - chunk id\_hash
- value:
  - reference count
  - size
  - encrypted/compressed size

The chunks cache is a hashindex, a hash table implemented in C and tuned for memory efficiency.

The **repository index** is stored in `repo/index.%d` and is used to determine a chunk's location in the repository.

The repo index is a key -> value mapping and contains:

- key:
  - chunk id\_hash
- value:
  - segment (that contains the chunk)
  - offset (where the chunk is located in the segment)

The repo index is a hashindex, a hash table implemented in C and tuned for memory efficiency.

Hints are stored in a file (`repo/hints.%d`).

It contains:

- version
- list of segments
- compact

hints and index can be recreated if damaged or lost using `check --repair`.

The chunks cache and the repository index are stored as hash tables, with only one slot per bucket, but that spreads the collisions to the following buckets. As a consequence the hash is just a start position for a linear search, and if the element is not in the table the index is linearly crossed until an empty bucket is found.

When the hash table is filled to 75%, its size is grown. When it's emptied to 25%, its size is shrunk. So operations on it have a variable complexity between constant and linear with low factor, and memory overhead varies between 33% and 300%.

## 11.11 Indexes / Caches memory usage

Here is the estimated memory usage of Borg - it's complicated:

```
chunk_count ~= total_file_size / 2 ^ HASH_MASK_BITS
repo_index_usage = chunk_count * 40
```

```
chunks_cache_usage = chunk_count * 44
files_cache_usage = total_file_count * 240 + chunk_count * 80
mem_usage ~= repo_index_usage + chunks_cache_usage + files_cache_usage = chunk_count * 164
    + total_file_count * 240
```

Due to the hashtables, the best/usual/worst cases for memory allocation can be estimated like that:

```
mem_allocation = mem_usage / load_factor # l_f = 0.25 .. 0.75
mem_allocation_peak = mem_allocation * (1 + growth_factor) # g_f = 1.1 .. 2
```

All units are Bytes.

It is assuming every chunk is referenced exactly once (if you have a lot of duplicate chunks, you will have less chunks than estimated above).

It is also assuming that typical chunk size is  $2^{\text{HASH\_MASK\_BITS}}$  (if you have a lot of files smaller than this statistical medium chunk size, you will have more chunks than estimated above, because 1 file is at least 1 chunk).

If a remote repository is used the repo index will be allocated on the remote side.

The chunks cache, files cache and the repo index are all implemented as hash tables. A hash table must have a significant amount of unused entries to be fast - the so-called load factor gives the used/unused elements ratio.

When a hash table gets full (load factor getting too high), it needs to be grown (allocate new, bigger hash table, copy all elements over to it, free old hash table) - this will lead to short-time peaks in memory usage each time this happens. Usually does not happen for all hashtables at the same time, though. For small hash tables, we start with a growth factor of 2, which comes down to  $\sim 1.1x$  for big hash tables.

E.g. backing up a total count of 1 Mi (IEC binary prefix i.e.  $2^{20}$ ) files with a total size of 1TiB.

1. with `create --chunker-params 10,23,16,4095` (custom, like borg < 1.0 or attic):  
mem\_usage = 2.8GiB
2. with `create --chunker-params 19,23,21,4095` (default):  
mem\_usage = 0.31GiB

---

**Note:** There is also the `--no-files-cache` option to switch off the files cache. You'll save some memory, but it will need to read / chunk all the files as it can not skip unmodified files then.

---

## 11.12 Encryption

AES<sup>104</sup>-256 is used in CTR mode (so no need for padding). A 64bit initialization vector is used, a HMAC-SHA256<sup>105</sup> is computed on the encrypted chunk with a random 64bit nonce and both are stored in the chunk. The header of each chunk is: TYPE (1) + HMAC (32) + NONCE (8) + CIPHERTEXT. Encryption and HMAC use two different keys.

In AES CTR mode you can think of the IV as the start value for the counter. The counter itself is incremented by one after each 16 byte block. The IV/counter is not required to be random but it must NEVER be reused. So to accomplish this Borg initializes the encryption counter to be higher than any previously used counter value before encrypting new data.

To reduce payload size, only 8 bytes of the 16 bytes nonce is saved in the payload, the first 8 bytes are always zeros. This does not affect security but limits the maximum repository capacity to only 295 exabytes ( $2^{64} * 16$  bytes).

---

<sup>104</sup> [https://en.wikipedia.org/wiki/Advanced\\_Encryption\\_Standard](https://en.wikipedia.org/wiki/Advanced_Encryption_Standard)

<sup>105</sup> <https://en.wikipedia.org/wiki/HMAC>

Encryption keys (and other secrets) are kept either in a key file on the client ('keyfile' mode) or in the repository config on the server ('repokey' mode). In both cases, the secrets are generated from random and then encrypted by a key derived from your passphrase (this happens on the client before the key is stored into the keyfile or as repokey).

The passphrase is passed through the `BORG_PASSPHRASE` environment variable or prompted for interactive usage.

## 11.13 Key files

When initialized with the `init -e keyfile` command, Borg needs an associated file in `$HOME/.config/borg/keys` to read and write the repository. The format is based on `msgpack`<sup>106</sup>, base64 encoding and `PBKDF2`<sup>107</sup> SHA256 hashing, which is then encoded again in a `msgpack`<sup>108</sup>.

The internal data structure is as follows:

**version** currently always an integer, 1

**repository\_id** the `id` field in the `config` INI file of the repository.

**enc\_key** the key used to encrypt data with AES (256 bits)

**enc\_hmac\_key** the key used to HMAC the encrypted data (256 bits)

**id\_key** the key used to HMAC the plaintext chunk data to compute the chunk's id

**chunk\_seed** the seed for the buzhash chunking table (signed 32 bit integer)

Those fields are processed using `msgpack`<sup>109</sup>. The utf-8 encoded passphrase is processed with `PBKDF2`<sup>110</sup> (`SHA256`<sup>111</sup>, 100000 iterations, random 256 bit salt) to give us a derived key. The derived key is 256 bits long. A `HMAC-SHA256`<sup>112</sup> checksum of the above fields is generated with the derived key, then the derived key is also used to encrypt the above pack of fields. Then the result is stored in a another `msgpack`<sup>113</sup> formatted as follows:

**version** currently always an integer, 1

**salt** random 256 bits salt used to process the passphrase

**iterations** number of iterations used to process the passphrase (currently 100000)

**algorithm** the hashing algorithm used to process the passphrase and do the HMAC checksum (currently the string `sha256`)

**hash** the HMAC of the encrypted derived key

**data** the derived key, encrypted with AES over a `PBKDF2`<sup>114</sup> SHA256 key described above

The resulting `msgpack`<sup>115</sup> is then encoded using base64 and written to the key file, wrapped using the standard `textwrap` module with a header. The header is a single line with a MAGIC string, a space and a hexadecimal representation of the repository id.

---

<sup>106</sup> <https://msgpack.org/>

<sup>107</sup> <https://en.wikipedia.org/wiki/PBKDF2>

<sup>108</sup> <https://msgpack.org/>

<sup>109</sup> <https://msgpack.org/>

<sup>110</sup> <https://en.wikipedia.org/wiki/PBKDF2>

<sup>111</sup> <https://en.wikipedia.org/wiki/SHA-256>

<sup>112</sup> <https://en.wikipedia.org/wiki/HMAC>

<sup>113</sup> <https://msgpack.org/>

<sup>114</sup> <https://en.wikipedia.org/wiki/PBKDF2>

<sup>115</sup> <https://msgpack.org/>

## 11.14 Compression

Borg supports the following compression methods:

- none (no compression, pass through data 1:1)
- lz4 (low compression, but super fast)
- zlib (level 0-9, level 0 is no compression [but still adding zlib overhead], level 1 is low, level 9 is high compression)
- lzma (level 0-9, level 0 is low, level 9 is high compression).

Speed: none > lz4 > zlib > lzma Compression: lzma > zlib > lz4 > none

Be careful, higher zlib and especially lzma compression levels might take a lot of resources (CPU and memory).

The overall speed of course also depends on the speed of your target storage. If that is slow, using a higher compression level might yield better overall performance. You need to experiment a bit. Maybe just watch your CPU load, if that is relatively low, increase compression until 1 core is 70-100% loaded.

Even if your target storage is rather fast, you might see interesting effects: while doing no compression at all (none) is a operation that takes no time, it likely will need to store more data to the storage compared to using lz4. The time needed to transfer and store the additional data might be much more than if you had used lz4 (which is super fast, but still might compress your data about 2:1). This is assuming your data is compressible (if you backup already compressed data, trying to compress them at backup time is usually pointless).

Compression is applied after deduplication, thus using different compression methods in one repo does not influence deduplication.

See `borg create --help` about how to specify the compression level and its default.





This chapter will get you started with Borg development.

Borg is written in Python (with a little bit of Cython and C for the performance critical parts).

## 12.1 Contributions

... are welcome!

Some guidance for contributors:

- discuss about changes on github issue tracker, IRC or mailing list
- make your PRs on the `master` branch (see *Branching Model* for details)
- do clean changesets:
  - focus on some topic, resist changing anything else.
  - do not do style changes mixed with functional changes.
  - try to avoid refactorings mixed with functional changes.
  - if you need to fix something after commit/push:
    - \* if there are ongoing reviews: do a fixup commit you can merge into the bad commit later.
    - \* if there are no ongoing reviews or you did not push the bad commit yet: edit the commit to include your fix or merge the fixup commit before pushing.
  - have a nice, clear, typo-free commit comment
  - if you fixed an issue, refer to it in your commit comment
  - follow the style guide (see below)
- if you write new code, please add tests and docs for it
- run the tests, fix anything that comes up

- make a pull request on github
- wait for review by other developers

## 12.2 Branching model

Borg development happens on the `master` branch and uses GitHub pull requests (if you don't have GitHub or don't want to use it you can send smaller patches via the borgbackup *Mailing list* to the maintainers).

Stable releases are maintained on maintenance branches named `x.y-maint`, eg. the maintenance branch of the 1.0.x series is `1.0-maint`.

Most PRs should be made against the `master` branch. Only if an issue affects **only** a particular maintenance branch a PR should be made against it directly.

While discussing / reviewing a PR it will be decided whether the change should be applied to maintenance branch(es). Each maintenance branch has a corresponding *backport/x.y-maint* label, which will then be applied.

Changes that are typically considered for backporting:

- Data loss, corruption and inaccessibility fixes
- Security fixes
- Forward-compatibility improvements
- Documentation corrections

### Maintainer part

From time to time a maintainer will backport the changes for a maintenance branch, typically before a release or if enough changes were collected:

1. Notify others that you're doing this to avoid duplicate work.
2. Branch a backporting branch off the maintenance branch.
3. Cherry pick and backport the changes from each labelled PR, remove the label for each PR you've backported.
4. Make a PR of the backporting branch against the maintenance branch for backport review. Mention the backported PRs in this PR, eg:

Includes changes from #2055 #2057 #2381

This way GitHub will automatically show in these PRs where they were backported.

### Historic model

Previously (until release 1.0.10) Borg used a “merge upwards”<sup>116</sup> model where most minor changes and fixes were committed to a maintenance branch (eg. `1.0-maint`), and the maintenance branch(es) were regularly merged back into the main development branch. This became more and more troublesome due to merges growing more conflict-heavy and error-prone.

---

<sup>116</sup> [https://git-scm.com/docs/gitworkflows#\\_merging\\_upwards](https://git-scm.com/docs/gitworkflows#_merging_upwards)

## 12.3 Code and issues

Code is stored on Github, in the [Borgbackup organization](#)<sup>117</sup>. [Issues](#)<sup>118</sup> and [pull requests](#)<sup>119</sup> should be sent there as well. See also the *Support* section for more details.

## 12.4 Style guide

We generally follow [pep8](#)<sup>120</sup>, with 120 columns instead of 79. We do *not* use form-feed ( $\text{^L}$ ) characters to separate sections either. Compliance is tested automatically when you run the tests.

## 12.5 Continuous Integration

All pull requests go through [Travis-CI](#)<sup>121</sup>, which runs the tests on Linux and Mac OS X as well as the flake8 style checker. Additional Unix-like platforms are tested on [Golem](#)<sup>122</sup>.

## 12.6 Output and Logging

When writing logger calls, always use correct log level (debug only for debugging, info for informative messages, warning for warnings, error for errors, critical for critical errors/states).

When directly talking to the user (e.g. Y/N questions), do not use logging, but directly output to stderr (not: stdout, it could be connected to a pipe).

To control the amount and kinds of messages output to stderr or emitted at info level, use flags like `--stats` or `--list`.

## 12.7 Building a development environment

First, just install borg into a virtual env as described before.

To install some additional packages needed for running the tests, activate your virtual env and run:

```
pip install -r requirements.d/development.txt
```

## 12.8 Running the tests

The tests are in the `borg/testsuite` package.

To run all the tests, you need to have `fakeroot` installed. If you do not have `fakeroot`, you still will be able to run most tests, just leave away the `fakeroot -u` from the given command lines.

To run the test suite use the following command:

---

<sup>117</sup> <https://github.com/borgbackup/borg/>  
<sup>118</sup> <https://github.com/borgbackup/borg/issues>  
<sup>119</sup> <https://github.com/borgbackup/borg/pulls>  
<sup>120</sup> <https://www.python.org/dev/peps/pep-0008/>  
<sup>121</sup> <https://travis-ci.org/borgbackup/borg>  
<sup>122</sup> <https://golem.enkore.de/view/Borg/>

```
fakeroot -u tox # run all tests
```

Some more advanced examples:

```
# verify a changed tox.ini (run this after any change to tox.ini):
fakeroot -u tox --recreate

fakeroot -u tox -e py34 # run all tests, but only on python 3.4

fakeroot -u tox borg.testsuite.locking # only run 1 test module

fakeroot -u tox borg.testsuite.locking -- -k "not Timer" # exclude some tests

fakeroot -u tox borg.testsuite -- -v # verbose py.test
```

Important notes:

- When using `--` to give options to `py.test`, you **MUST** also give `borg.testsuite[.module]`.

## 12.9 Regenerate usage files

Usage documentation is currently committed directly to git, although those files are generated automatically from the source tree.

When a command is added, a commandline flag changed, added or removed, the usage docs need to be rebuilt as well:

```
./setup.py build_usage
```

However, we prefer to do this as part of our releasing preparations, so it is generally not necessary to update these when submitting patches that change something about the command line.

## 12.10 Building the docs with Sphinx

The documentation (in reStructuredText format, `.rst`) is in `docs/`.

To build the html version of it, you need to have sphinx installed:

```
pip3 install sphinx sphinx_rtd_theme # important: this will install sphinx with
↳ Python 3
```

Now run:

```
cd docs/
make html
```

Then point a web browser at `docs/_build/html/index.html`.

The website is updated automatically through Github web hooks on the main repository.

## 12.11 Using Vagrant

We use Vagrant for the automated creation of testing environments and borgbackup standalone binaries for various platforms.

For better security, there is no automatic sync in the VM to host direction. The plugin *vagrant-scp* is useful to copy stuff from the VMs to the host.

Usage:

```
# To create and provision the VM:
vagrant up OS
# To create an ssh session to the VM:
vagrant ssh OS command
# To shut down the VM:
vagrant halt OS
# To shut down and destroy the VM:
vagrant destroy OS
# To copy files from the VM (in this case, the generated binary):
vagrant scp OS:/vagrant/borg/borg.exe .
```

## 12.12 Creating standalone binaries

Make sure you have everything built and installed (including llfuse and fuse). When using the Vagrant VMs, pyinstaller will already be installed.

With virtual env activated:

```
pip install pyinstaller # or git checkout master
pyinstaller -F -n borg-PLATFORM borg/__main__.py
for file in dist/borg-*; do gpg --armor --detach-sign $file; done
```

If you encounter issues, see also our *Vagrantfile* for details.

---

**Note:** Standalone binaries built with pyinstaller are supposed to work on same OS, same architecture (x86 32bit, amd64 64bit) without external dependencies.

---

## 12.13 Creating a new release

Checklist:

- make sure all issues for this milestone are closed or moved to the next milestone
- check if there are any pending fixes for security issues
- find and fix any low hanging fruit left on the issue tracker
- check that Travis CI is happy
- update `CHANGES.rst`, based on `git log $PREVIOUS_RELEASE..`
- check version number of upcoming release in `CHANGES.rst`
- verify that `MANIFEST.in` and `setup.py` are complete
- `python setup.py build_usage` and commit
- tag the release:

```
git tag -s -m "tagged/signed release X.Y.Z" X.Y.Z
```

- create a clean repo and use it for the following steps:

```
git clone borg borg-clean
```

This makes sure no uncommitted files get into the release archive. It also will find if you forgot to commit something that is needed. It also makes sure the vagrant machines only get committed files and do a fresh start based on that.

- run tox and/or binary builds on all supported platforms via vagrant, check for test failures
- create a release on PyPi:

```
python setup.py register sdist upload --identity="Thomas Waldmann" --sign
```

- close release milestone on Github
- announce on:
  - Mailing list
  - Twitter
  - IRC channel (change `/topic`)
- create a Github release, include:
  - standalone binaries (see above for how to create them)
    - \* for OS X, document the OS X Fuse version in the README of the binaries. OS X FUSE uses a kernel extension that needs to be compatible with the code contained in the binary.
  - a link to `CHANGES.rst`

### 13.1 Borg authors (“The Borg Collective”)

- Thomas Waldmann <tw@waldmann-edv.de>
- Antoine Beaupré <anarcat@debian.org>
- Radek Podgorny <radek@podgorny.cz>
- Yuri D’Elia
- Michael Hanselmann <public@hansmi.ch>
- Teemu Toivanen <public@profnetti.fi>
- Marian Beermann <public@enkore.de>
- Martin Hostettler <textshell@uchuujin.de>

Borg is a fork of Attic.

### 13.2 Attic authors

Attic is written and maintained by Jonas Borgström and various contributors:

#### 13.2.1 Attic Development Lead

- Jonas Borgström <jonas@borgstrom.se>

#### 13.2.2 Attic Patches and Suggestions

- Brian Johnson

- Cyril Roussillon
- Dan Christensen
- Jeremy Maitin-Shepard
- Johann Klähn
- Petros Moisiadis
- Thomas Waldmann



# CHAPTER 14

---

## License

---

Copyright (C) 2015-2018 The Borg Collective (see AUTHORS file)  
Copyright (C) 2010-2014 Jonas Borgström <jonas@borgstrom.se>  
All rights reserved.

Redistribution and use in `source` and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of `source` code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.