Riservatezza per i nostri backup

Campagna per la tutela di backup abbandonati, indifesi e maltrattati

Andrea fRANz Francesconi
andrea.francesconi@gmail.com
https://www.linkedin.com/in/andreafrancesconi/

IMOLUG
Imola & Faenza Linux User Group

Backup

In information technology, a backup, or the process of backing up, refers to the **copying and archiving of computer data** so it may be used to restore the original after a data loss event.

…

Backups have two distinct purposes. The primary purpose is **to recover data after its loss**, be it by data deletion or corruption. Data loss can be a common experience of computer users; a 2008 survey found that 66% of respondents had lost files on their home PC. The secondary purpose of backups is **to recover data from an earlier time**, according to a user-defined data retention policy, typically configured within a backup application for how long copies of data are required.

…

(https://en.wikipedia.org/wiki/Backup)

Backup

"Il backup è quella cosa che andava fatta prima"
cit. anonimo

(Good) Reasons

- YOUR files, YOUR data
- Hard drive failures
- Notebook/tablet/smartphones/* get lost or stolen
- Human error
- Ransomware

Approach

Local

VPS

Disconnected/
offline media

Cloud

Multiple
bck media

Media at
another
place

Cloud

- Constraints
- Attention points

Borg

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.

(https://github.com/borgbackup/borg/)

Borg

- Borg is a fork of Attic
- Borg was created in May 2015
- Borg is written in Python (with a little bit of Cython and C for the performance critical parts)
- Borg 1.1 is the current stable series of Borg
- Active community
- Additional resources:
  - BorgWeb
  - borgmatic

Features

- Space efficient storage
- Speed
- Data encryption
- Compression
- Off-site backups
- Backups mountable as filesystems
- Easy installation on multiple platforms
- Free and Open Source Software

Big picture

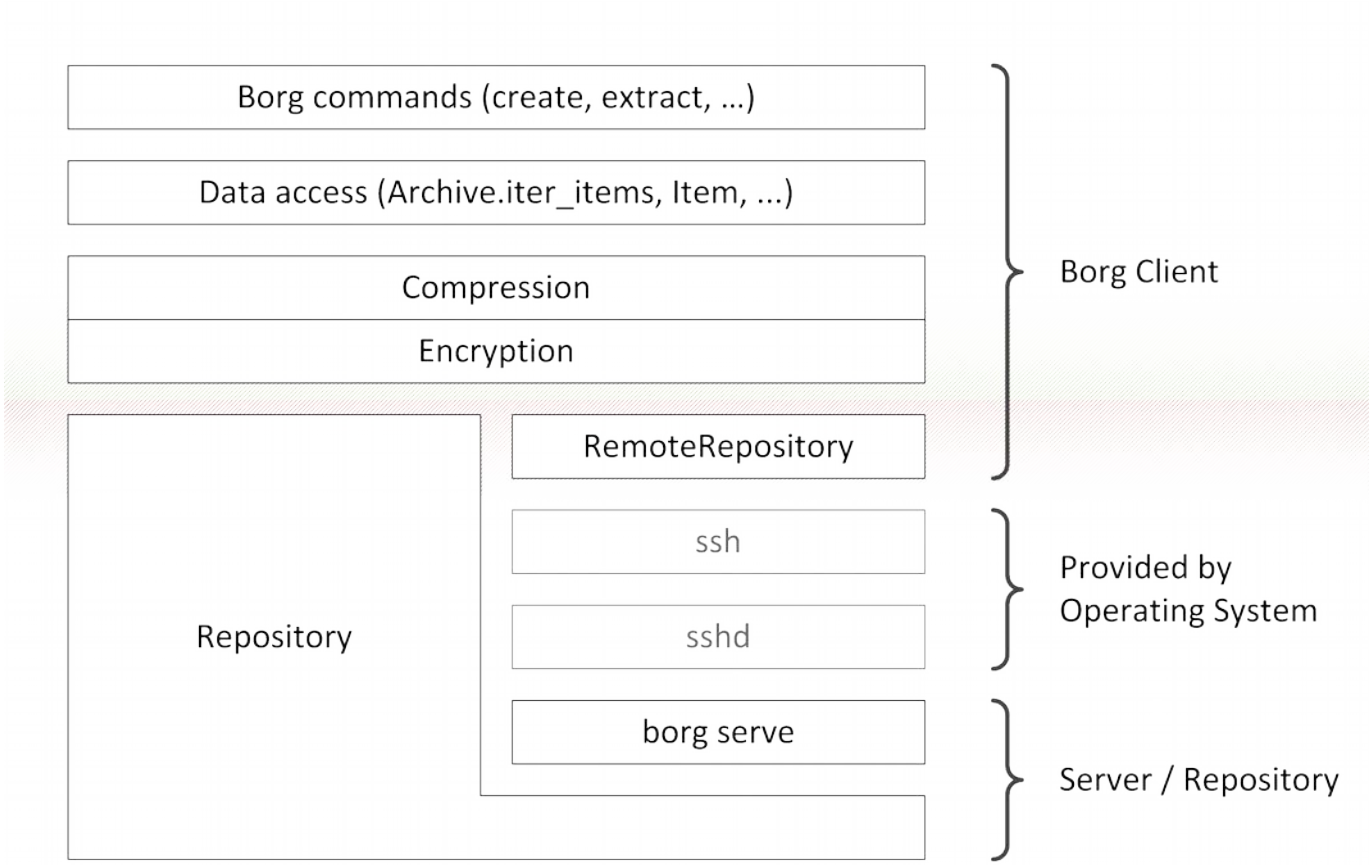

(https://borgbackup.readthedocs.io/en/stable/internals.html)

Scenario

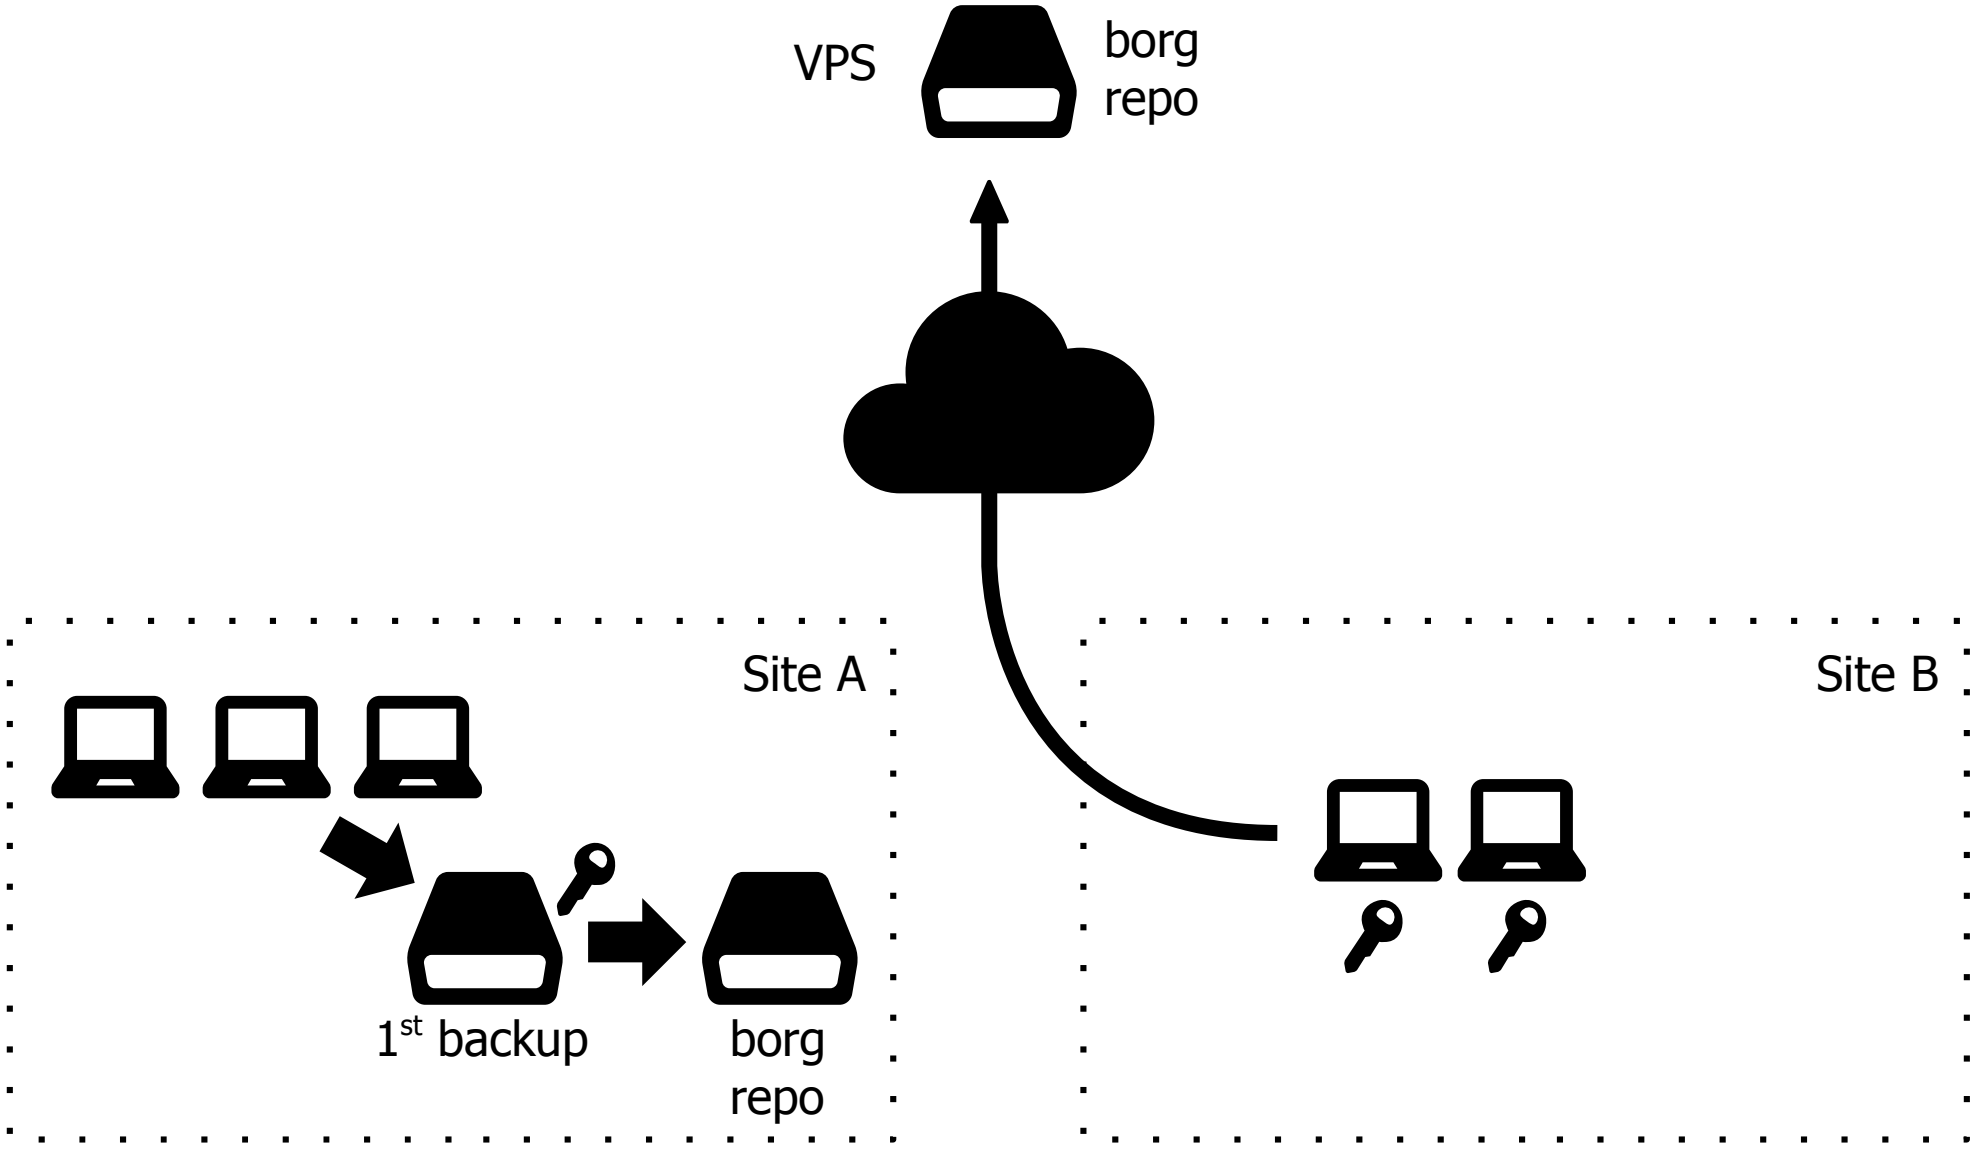

Internals - Deduplication

- 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

Internals - Encryption

- Encryption can be enabled or disabled at repository creation time
- When repository encryption is enabled all data is encrypted using 256-bit AES encryption and the integrity and authenticity is verified using HMAC-SHA256
- All data is encrypted on the client before being written to the repository
- 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

(https://borgbackup.readthedocs.io/en/stable/quickstart.html#repository-encryption)

Attack surface

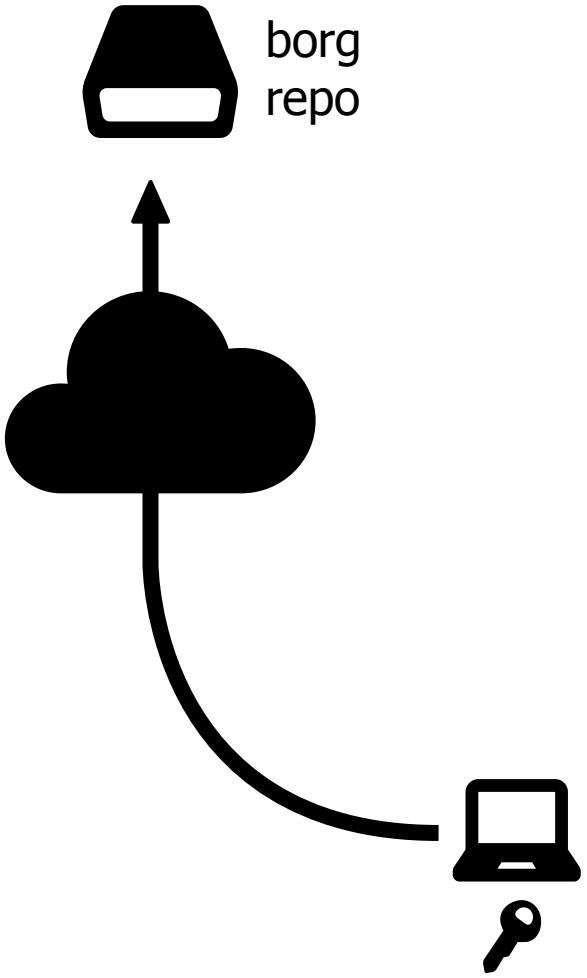

- The attack model of Borg is that the environment of the client process (e.g. borg create) is trusted and the repository (server) is not. The attacker has any and all access to the repository, including interactive manipulation (man-in-the-middle) for remote repositories.
- Furthermore the client environment is assumed to be persistent across attacks (practically this means that the security database cannot be deleted between attacks).
- Under these circumstances Borg guarantees that the attacker cannot
  - modify the data of any archive without the client detecting the change
  - rename, remove or add an archive without the client detecting the change
  - recover plain-text data
  - recover definite (heuristics based on access patterns are possible) structural information such as the object graph (which archives refer to what chunks)
- The attacker can always impose a denial of service per definition (he could forbid connections to the repository, or delete it entirely).

Stats

```
# cat 20171022020001.borg
...
A /mnt/500/bck_daily/daily.0/bck_.../.../Scanner/20171020082131494.pdf
A /mnt/500/bck_.../daily.0/.../.../c_.../05_Venerdi.bck
A /mnt/500/bck_.../daily.0/.../.../c_.../LOG Backup AUTOMATICO di
05_Venerdi.log
{
    "archive": {
        "command_line": [
            "/usr/local/bin/borg",
            "create",
            "--verbose",
            "--filter", "AME",
            "--list",
            "--stats",
            "--json",
            "--show-rc",
            "--compression", "lz4",
            "--files-cache=mtime,size",
            "::20171022020001",
            "<dir list>"
        ],
        "duration": 202.273784,
        "end": "2017-10-22T02:03:33.000000",
        "id":
"ad5abea7aa41057b4212285f2d1c4518d21530b8f047641988b8c8aea5be7187",
        "limits": {
            "max_archive_size": 0.0001528742917941076
        },
        "name": "20171022020001",
        "start": "2017-10-22T02:00:11.000000",
        "stats": {
            "compressed_size": 18195213921,
            "deduplicated_size": 100626197,
            "nfiles": 36625,
            "original_size": 21280134658
        }
    },
    "cache": {
        "path":
"/.../.cache/borg/4ea7cb905897ffbd4b7bded644494403ef50
2f15e15591c73cba0223db7f6761",
        "stats": {
            "total_chunks": 519876,
            "total_csize": 219455208356,
            "total_size": 256499577884,
            "total_unique_chunks": 28989,
            "unique_csize": 17563927155,
            "unique_size": 20938709885
        }
    },
    "encryption": {
        "keyfile": "/.../.config/borg/keys/....2",
        "mode": "keyfile"
    },
    "repository": {
        "id":
"4ea7cb905897ffbd4b7bded644494403ef502f15e15591c73cba0
223db7f6761",
        "last_modified": "2017-10-22T02:03:33.000000",
        "location":
"ssh://bck_usr_...@.../mnt/bck_.../borg"
    }
}
terminating with success status, rc 0
```

Stats

```
# cat 20171022020001.borg
...
A /mnt/500/bck_daily/daily.0/bck_.../.../Scanner/20171020082131494.pdf
A /mnt/500/bck_.../daily.0/.../.../c_.../05_Venerdi.bck
A /mnt/500/bck_.../daily.0/.../.../c_.../LOG Backup AUTOMATICO di
05_Venerdi.log
{
    "archive": {
        "command_line": [
            "/usr/local/bin/borg",
            "create",
            "--verbose",
            "--filter", "AME",
            "--list",
            "--stats",
            "--json",
            "--show-rc",
            "--compression", "lz4",
            "--files-cache=mtime,size",
            "::20171022020001",
            "<dir list>"
        ],
        "duration": 202.273784,
        "end": "2017-10-22T02:03:33.000000",
        "id":
"ad5abea7aa41057b4212285f2d1c4518d21530b8f047641988b8c8aea5be7187",
        "limits": {
            "max_archive_size": 0.0001528742917941076
        },
        "name": "20171022020001",
        "start": "2017-10-22T02:00:11.000000",
        "stats": {
            "compressed_size": 18195213921,
            "deduplicated_size": 100626197,
            "nfiles": 36625,
            "original_size": 21280134658
        }
    },
```

Job parameters

```
    "cache": {
        "path":
"/.../.cache/borg/4ea7cb905897ffbd4b7bded644494403ef50
2f15e15591c73cba0223db7f6761",
        "stats": {
            "total_chunks": 519876,
            "total_csize": 219455208356,
            "total_size": 256499577884,
            "total_unique_chunks": 28989,
            "unique_csize": 17563927155,
            "unique_size": 20938709885
        }
    },
    "encryption": {
        "keyfile": "/.../.config/borg/keys/....2",
        "mode": "keyfile"
    },
    "repository": {
        "id":
"4ea7cb905897ffbd4b7bded644494403ef502f15e15591c73cba0
223db7f6761",
        "last_modified": "2017-10-22T02:03:33.000000",
        "location":
"ssh://bck_usr_...@.../mnt/bck_.../borg"
    }
}
terminating with success status, rc 0
```

Stats

```
# cat 20171022020001.borg

A /mnt/500/bck_daily/daily.0/bck_.../.../Scanner/20171020082131494.pdf
A /mnt/500/bck_.../daily.0/.../.../c_.../05_Venerdi.bck
A /mnt/500/bck_.../daily.0/.../.../c_.../LOG Backup AUTOMATICO di
05_Venerdi.log
{
    "archive": {
        "command_line": [
            "/usr/local/bin/borg",
            "create",
            "--verbose",
            "--filter", "AME",
            "--list",
            "--stats",
            "--json",
            "--show-rc",
            "--compression", "lz4",
            "--files-cache=mtime,size",
            "::20171022020001",
            "<dir list>"
        ],
        "duration": 202.273784,
        "end": "2017-10-22T02:03:33.000000",
        "id":
"ad5abea7aa41057b4212285f2d1c4518d21530b8f[illegible]",
        "limits": {
            "max_archive_size": 0.0[illegible]
        },
        "name": "20171022020001",
        "start": "2017-10-22T02:00:11.000000",
        "stats": {
            "compressed_size": 18195213921,
            "deduplicated_size": 100626197,
            "nfiles": 36625,
            "original_size": 21280134658
        }
    },
```

```
    "cache": {
        "path":
"/.../.cache/borg/4ea7cb905897ffbd4b7bded644494403ef50
2f15e15591c73cba0223db7f6761",
        "stats": {
            "total_chunks": 519876,
            "total_csize": 219455208356,
            "total_size": 256499577884,
            "total_unique_chunks": 28989,
            "unique_csize": 17563927155,
            "unique_size": 20938709885
        }
    },
    "encryption": {
        "keyfile": "/.../.config/borg/keys/....2",
        "mode": "keyfile"
    },
    "repository": {
        "id":
"4ea7cb905897ffbd4b7bded644494403ef502f15e15591c73cba0
[illegible]6761",
        "last_modified": "2017-10-22T02:03:33.000000",
        "location":
"ssh://bck_usr_...@.../mnt/bck_.../borg"
    }
}
terminating with success status, rc 0
```

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

'A' = regular file, added
'M' = regular file, modified
'U' = regular file, unchanged
'E' = regular file, an error happened while accessing/reading this file

Stats

```
# cat 20171022020001.borg
...
A /mnt/500/bck_daily/daily.0/bck_.../.../Scanner/20171020082131494.pdf
A /mnt/500/bck_.../daily.0/.../.../c_.../05_Venerdi.bck
A /mnt/500/bck_.../daily.0/.../.../c_.../LOG Backup AUTOMATICO di
05_Venerdi.log
{
    "archive": {
        "command_line": [
            "/usr/local/bin/borg",
            "create",
            "--verbose",
            "--filter", "AME",
            "--list",
            "--stats",
            "--json",
            "--show-rc",
            "--compression", "lz4",
            "--files-cache=mtime,size",
            "::20171022020001",
            "<dir list>"
        ],
        "duration": 202.273784,
        "end": "2017-10-22T02:03:33.000000",
        "id":
"ad5abea7aa41057b4212285f2d1c4518d21530b8f047641988b8c8aea5be7487",
        "limits": {
            "max_archive_size": 0.0001528742917941076
        },
        "name": "20171022020001",
        "start": "2017-10-22T02:00:11.000000",
        "stats": {
            "compressed_size": 18195213921,
            "deduplicated_size": 100626197,
            "nfiles": 36625,
            "original_size": 21280134658
        }
    },
    "cache": {
        "path":
"/.../.cache/borg/4ea7cb905897ffbd4b7bded644494403ef50
2f15e15591c73cba0223db7f6761",
        "stats": {
            "total_chunks": 519876,
            "total_csize": 219455208356,
            "total_size": 256499577884,
            "total_unique_chunks": 28989,
            "unique_csize": 17563927155,
            "unique_size": 20938709885
        }
    },
    "encryption": {
        "keyfile": "/.../.config/borg/keys/....2",
        "mode": "keyfile"
    },
    "repository": {
        "id":
"4ea7cb905897ffbd4b7bded644494403ef502f15e15591c73cba0
223db7f6761",
        "last_modified": "2017-10-22T02:03:33.000000",
        "location":
"ssh://bck_usr_...@.../mnt/bck_.../borg"
    }
}
terminating with success status, rc 0
```

Archive name
Start, duration and stop time
Return code of the job

Stats

```
# cat 20171022020001.borg
...
A /mnt/500/bck_daily/daily.0/bck_.../.../Scanner/20171020082131494.pdf
A /mnt/500/bck_.../daily.0/.../.../c_.../05_Venerdi.bck
A /mnt/500/bck_.../daily.0/.../.../c_.../LOG Backup AUTOMATICO di
05_Venerdi.log
{
    "archive": {
        "command_line": [
            "/usr/local/bin/borg",
            "create",
            "--verbose",
            "--filter", "AME",
            "--list",
            "--stats",
            "--json",
            "--show-rc",
            "--compression", "lz4",
            "--files-cache=mtime,size",
            "::20171022020001",
            "<dir list>"
        ],
        "duration": 202.273784,
        "end": "2017-10-22T02:03:33.000000",
        "id":
"ad5abea7aa41057b4212285f2d1c4518d21530b8f047641988b8c8aea5be7187",
        "limits": {
            "max_archive_size": 0.0001528742917941076
        },
        "name": "20171022020001",
        "start": "2017-10-22T02:00:11.000000",
        "stats": {
            "compressed_size": 18195213921,
            "deduplicated_size": 100626197,
            "nfiles": 36625,
            "original_size": 21280134658
        }
    },
    "cache": {
        "path":
"/.../.cache/borg/4ea7cb905897ffbd4b7bded644494403ef50
2f15e15591c73cba0223db7f6761",
        "stats": {
            "total_chunks": 519876,
            "total_csize": 219455208356,
            "total_size": 256499577884,
            "total_unique_chunks": 28989,
            "unique_csize": 17563927155,
            "unique_size": 20938709885
        }
    },
    "encryption": {
        "keyfile": "/.../.config/borg/keys/....2",
        "mode": "keyfile"
    },
    "repository": {
        "id":
"4ea7cb905897ffbd4b7bded644494403ef502f15e15591c73cba0
223db7f6761",
        "last_modified": "2017-10-22T02:03:33.000000",
        "location":
"ssh://bck_usr_..._@.../mnt/bck_.../borg"
    }
}
terminating with success status, rc 0
```

compressed_size: size after compression ~18GB
deduplicated_size: dedup size against current repo ~100MB
nfiles: nr. of regular files in the repo
original_size: size of files+metadata before compression ~21GB

Stats

```
# cat 20171022020001.borg
...
A /mnt/500/bck_daily/daily.0/bck_.../.../Scanner/20171020082131494.pdf
A /mnt/500/bck_.../daily.0/.../.../c_.../05_Venerdi.bck
A /mnt/500/bck_.../daily.0/.../.../c_.../LOG Backup AUTOMATICO di
05_Venerdi.log
{
    "archive": {
        "command_line": [
            "/usr/local/bin/borg",
            "create",
            "--verbose",
            "--filter", "AME",
            "--list",
            "--stats",
            "--json",
            "--show-rc",
            "--compression", "lz4",
            "--files-cache=mtime,size",
            "::20171022020001",
            "<dir list>"
        ],
        "duration": 202.273784,
        "end": "2017-10-22T02:03:33.000000",
        "id":
"ad5abea7aa41057b4212285f2d1c4518d21530b8f0
        "limits": {
            "max_archive_size": 0.0001528742
        },
        "name": "20171022020001",
        "start": "2017-10-22T02:00:11.000000",
        "stats": {
            "compressed_size": 18195213921,
            "deduplicated_size": 100626197,
            "nfiles": 36625,
            "original_size": 21280134658
        }
    },
```

```
    "cache": {
        "path":
"/.../.cache/borg/4ea7cb905897ffbd4b7bded644494403ef50
2f15e15591c73cba0223db7f6761",
        "stats": {
            "total_chunks": 519876,
            "total_csize": 219455208356,
            "total_size": 256499577884,
            "total_unique_chunks": 28989,
            "unique_csize": 17563927155,
            "unique_size": 20938709885
        }
    },
    "encryption": {
        "keyfile": "/.../.config/borg/keys/....2",
        "mode": "keyfile"
    },
    "repository": {
        "id":
"4ea7cb905897ffbd4b7bded644494403ef502f15e15591c73cba0
223db7f6761",
        "last_modified": "2017-10-22T02:03:33.000000",
        "location":
"ssh://bck_usr_...@.../mnt/bck_.../borg"
    }
}
terminating with success status, rc 0
```

total_chunks: number of chunks ~520k
unique_chunks: number of unique chunks ~29k
total_size: total uncompr. size of all chunks multiplied with their reference counts ~219GB
unique_csize: compressed and encrypted size of all chunks ~17GB

Stats

```
# borg info /mnt/borgxxx::test03
Archive name: test03
Archive fingerprint: d68d24f7ee999849995b9a77d36e6ea2723f25d6286ed35fda45b5c4a5b03fd5
...
Time (start): Tue, 2017-10-24 07:20:57
Time (end): Tue, 2017-10-24 10:43:49
Duration: 3 hours 22 minutes 51.76 seconds
...
Command line: /usr/local/bin/borg create -s --list --compression none
/mnt/borgxxx::test03 /mnt/bck_xxx/xxx
Utilization of maximum supported archive size: 0%
------------------------------------------------------------------------------

                       Original size      Compressed size    Deduplicated size
This archive:              840.50 GB             840.50 GB              4.59 MB
All archives:                1.46 TB               1.46 TB            188.64 GB

                       Unique chunks         Total chunks
Chunk index:                   56132               514617
#
```

Stats

```
# borg info /mnt/borgxxx::test03
Archive name: test03
Archive fingerprint: d68d24f7ee999849995b9a77d36e6ea2723f25d6286ed35fda45b5c4a5b03fd5
...
Time (start): Tue, 2017-10-24 07:20:57
Time (end): Tue, 2017-10-24 10:43:49
Duration: 3 hours 22 minutes 51.76 seconds
...
Command line: /usr/local/bin/borg create -s --list --compression none
/mnt/borgxxx::test03 /mnt/bck_xxx/xxx
Utilization of maximum supported archive size: 0%
------------------------------------------------------------------------------
                       Original size      Compressed size    Deduplicated size
This archive:              840.50 GB             840.50 GB              4.59 MB
All archives:                1.46 TB               1.46 TB            188.64 GB

                       Unique chunks         Total chunks
Chunk index:                   56132               514617
#
```

```
# df -k /mnt/borgxxx
Filesystem                  1K-blocks      Used Available Use% Mounted on
/dev/mapper/wd_2tb_sb-borgxxx 515010816 184303004 304477028  38%
/mnt/borgxxx
#
```

Remote repo

- Friend's SSH
- VPS
- Cloud services (e.g. rsync.net)

Tips

- Backup side 
  - analysis
  - schedule
  - monitoring

- Restore side 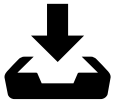
  - Test, test and test
  - Make a backup copy of key file/repo config file
  - Keep your passphrase at a safe place

Q&A

- Q & A ❓
- Thanks! 👍