Administration Manual for Infscape UrBackup Appliance 1.x

Infscape UG (haftungsbeschränkt)

February 21, 2020

Contents

1 Introduction

Infscape UrBackup Appliance is an advanced client/server backup system with clients for Windows, Linux and macOS, server to server replication and hybrid backup to cloud (including AWS S3 compatible). Backups are deduplicated and compressed before being stored to local RAID data storage or the cloud storage, which leads to lower costs due to the decreased storage usage especially for long-term backup archival. The appliance does both image and file backups, allowing the appropriate selection of the preferred backup method depending on organization requirements and procedures. A wide range of applications such as databases are directly supported by the clients, making setting up backups easy.

2 Appliance setup

2.1 Networking

The first setup step is to setup networking, as otherwise network connectivity cannot be guaranteed. The network can be setup manually or via DHCP.

2.2 Account and appliance name

As a second step the UrBackup appliance needs account details (username and password). If you do not have a account yet, you can create one. The account is used for password reset, optionally for encrypted settings storage and updates.

2.3 Storage

Please see section 3 for a description of the available storage options. For the cloud storage or wrie-back cached auto-layout RAID first add the write-back cache disk (e.g. NVMe) and select the appropriate use. Afterwards the RAID hard disks can be added or, optionally, the RAID metadata SSDs.

3 Backup/archival storage

UrBackup is optimized for fast incremental backups. Each of those backups can be deleted at any point, that is, there is no special relationship between backups such that some backups have to be deleted with other backups like in traditional backup systems, obliviating the need for regular full backups. Additionally UrBackup does file level deduplication, such that most peaces of data are stored only once on backup storage. This causes a larger number of non-sequential writes compared to traditional (optimized for tapes) backup systems.

Taking a step back UrBackup replicates the file or volume changes of each client at certain time intervals on the backup storage. Often the client will modify a single block multiple times or add new data, but the worst case is that the backup server has to "catch up" with a lot of random writes during a backup. So even a single client with an SSD can keep a backup storage comprised of spinning disks busy, such that backups take a long time. A consumer SSD can do over 10000 random write operations per second, while a consumer hard disk can do only slightly over 100. In this case a backup of 1 hour of worst-case changes would take 100 hours. Disks of workstations are mostly idle and most changes sequential or new data. Database servers can be experience this worst-case scenario in practice, however.

The UrBackup appliance offers three different storage modes. The first is suitable for a backup server with only few clients and a maximum of two disks in a RAID1 or RAID0 configuration, the second uses a bucket via an Amazon S3 compatible API as backup storage and the third is a local RAID with automatic layouting and up to 32 parity disks. The cloud storage and the local raid require a fast local disk (NVME) as write-back cache. Everything is written to this cache first,

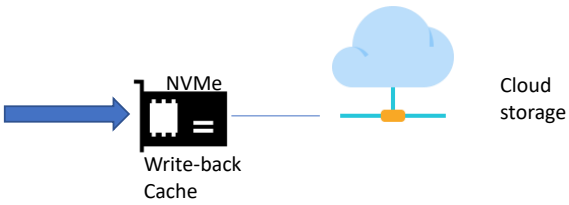

Figure 1: Example cloud storage configuration

then compressed and encrypted, then written to the cloud storage/RAID. With the local cache large enough, most random writes or reads will be satisfied by the cache and the storage will have a high number of Input/Output Operations Per Second (IOPS), while the cloud storage/RAID can store the data space efficiently.

3.1 Simple RAID1/RAID0 storage

Then simple RAID1/RAID0 storage should be used if the backup appliance has only few clients and the backup storage consists only of one or two disks. This storage cannot be converted to the write-back cached auto-layout RAID later.

If this option is selected the appliance will use a plain btrfs RAID1/RAID0 and use LUKS for full disk encryption.

3.2 Cloud storage

The UrBackup Appliance can use a bucket in a Amazon S3 compatible cloud storage as backup storage. A local disk as write-back cache is mandatory. Cloud storage is setup by selecting the local disk, and then selecting it as cache for the cloud storage. The connection details of the cloud storage can be setup in the settings. All backups will be stored into the cloud storage. The "archive to cloud" feature will be disabled in turn. The appliance will create a large amount of object in the configured bucket. If the operating system disk or the cloud storage cache disk is lost the objects in this bucket will be enumerated during recovery. Please evaluate carefully which cloud storage satisfies e.g. the recovery time. In most cases it doesn't make sense to setup a cloud object storage only for use with UrBackup appliance. In this case you should have a look at the write-back cached auto-layout RAID.

See figure 1 for an illustation of the cloud storage setup.

3.3 Write-back cached auto-layout RAID

A RAID implementation with write-back cache and automatic layouting optimized for use as backup/archival storage. The write-back cache is mandatory. The RAID shares some features with the cloud storage implementation. All writes to the RAID are checksummed and problems such as the write-hole in traditional RAID56 is avoided.

3.3.1 Automatic layout

The RAID will automatically configure itself into groups with appropriate RAID levels. Disks with same sizes will automatically be arranged into the same groups. There are two parameters which influence how the disks are arranged. The maximum failure probability and the target overhead. The first specifies the maximum failure probability per year that the whole RAID is allowed to have. If you select a failure probability of <= 5% the probability that the RAID will fail in a year is max. 5 in 100. The overhead influences the performance of the RAID (the more overhead the more performance). If you specify an overhead of 50% it will create a RAID10 layout. But the backup storage will only have 50% auf the cumulative disk capacity. The

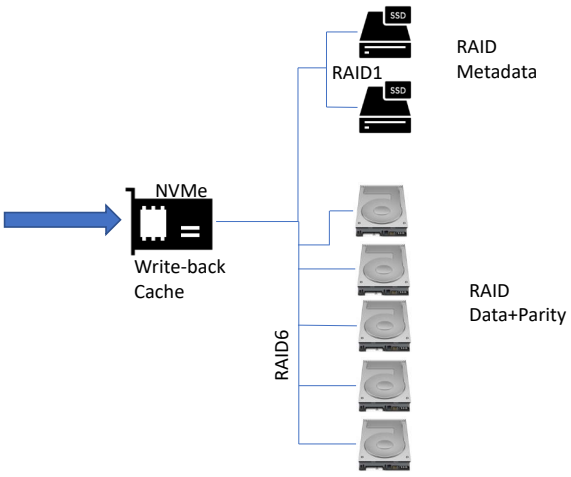

Figure 2: Example RAID configuration with five hard disks and two solid state disks as RAID metadata storage

overhead is a target, which the layout sometimes does not satisfy, while the failure probabilty is always satisfied. You can expirement with different disk setups and target failure and overhead percentages at `https://raidsim.urbackup.org/`.

The layouts are calculated by running simulations, taking into account for example unrecoverable read errors (URE) during recovery from a disk failure. The simulations currently do not use disk specific failure information.

3.3.2 Separate RAID metadata disks

Every RAID disks contains a small database with data allocation information (RAID metadata). This database is read and written to mostly non-sequentially. This database can be moved to a mirrored set of SSDs to further improve RAID performance. Its size is approximately 1GB per 1TB of RAID disk.

See figure 2 for an example layout with five hard disks and two seperate solid state disks for RAID metadata.

3.3.3 Sharing RAID disks with other appliances

RAID disks can be shared with other appliances. Enable "Share RAID disks over the network (tcp port 3541)" in the advanced RAID storage settings on the appliance with the RAID disks to be shared, then reboot the appliance. Afterwards go to the storage settings in the other appliance and select "+Add remote RAID disk(s)", select "RAID disks on remote UrBackup appliance (tcp port 3541)". After clicking on "Discover disks" it shows the disks on the remote appliance. Then they can be added to the local RAID as shared RAID disks.

3.3.4 Mirroring to S3

The complete RAID can be mirrored to S3 (compatible) storage. Select "Mirror RAID to S3" in the advanced RAID storage settings, afterwards configure the S3 access key, secret access key bucket name, storage class and S3 endpoint URL. The RAID will be mirrored during the "RAID cloud mirror window" (e.g. 1-7/1-5 to mirror between 1am and 5am every day). To recover from a mirror, setup a new appliance, then select the S3 bucket as "cloud cached" backup storage. Afterwards you'll have a working appliance that stores to S3 having a local cache (see 3.2), this setup can then be switched back to a RAID setup in the advanced cloud drive settings.

3.3.5 Defragmentation

The RAID storage can become fragmented decreasing performance (more random IO). Enabling "Nightly defrag of data disks" in the advanced RAID settings combats fragmentation. The RAID defrag will run nightly after cleanup has finished, for a maximum of "Max time to spend on nightly data disk defragmentation" minutes.

3.4 Shared cloud storage and auto-layout RAID features

3.4.1 Write-back cache

All writes go first to the write-back cache and are then written after significant delay to the RAID disks/cloud storage. This is coordinated with UrBackup, such that write-back happens after a backup is finished. This causes all the writes to the RAID to be sequential, which is optimal for spinning disks, and all writes to the cloud storage to be new objects, which is necessary for some object storage consistency models. If the cache is large enough reads can also be mostly satisfied by it. The cache is crash-persistent, excluding the currently active transaction. The cache should be scaled such that changing data remains in cache between two incremental backups. For example, if a client is backed up every 5 hours and during those 5 hours 500GB is backed up from other clients, the cache should be larger than 500GB such that the data from the last backup remains in cache and isn't evicted.

3.4.2 Compression and encryption

All data is compressed and encrypted before writing it to the RAID/cloud storage. Currently zlib is used for compression. Autheticated encryption is used (AES-GCM).

3.4.3 Background compression

If the RAID/cloud storage is idle it performs background compression. The much better, but also more CPU intensive, LZMA compression is used for this. If the background compression yields better results than the first zlib compression data will be rewritten to RAID/cloud storage with the better compression result.

3.4.4 Dynamic variable object/stripe size

The RAID/cloud storage dynamically sizes the stripe/object sizes. Stripes/objects which have partial writes automatically get transformed into smaller stripes/objects. This improves long-term performance, as the compressability and write performance of large stripes/objects is better, but decreases short-term performance because some large stripes/objects have to be rewritten as smaller stripes/objects in-order to decrease the large stripe/object write overhead.

3.5 Advanced cloud storage/auto-layout RAID settings

There are some advanced cloud storage/auto-layout RAID settings. Usually you should not change them. Please read the following sections and make sure you understand them before changing them.

3.5.1 Maximum number of active transactions

The cloud storage/auto-layout RAID works with transaction in order to ensure that the cache device breaking does not corrupt the storage. The more frequent the transactions, the less data is lost (last few backups) if the cache device fails. Increasing the maximum number of active transactions decreases the number of backups which are lost after a cache device failure (e.g. from last hour with five transactions to last 30 minutes with 20 transactions), but it also causes the

cache device to be used less efficiently and causes more IO to the cache device, which may cause decreased overall performance. By default this is "20" for cloud storage and "5" for RAID.

3.5.2 Size of in-memory write-back cache

To decrease the amount of data written to the flash cache, it can be cached only in-memory instead. This causes more data to be read and written to from the cloud storage/RAID and delays backup commits, i.e. more last backups are lost after power loss/reboot. This is because data is written directly from memory to RAID and data in RAM is lost during power loss/reboot. Per default this is disabled (by setting to "0"). To protect agains out-of-memory errors, this is limited to half the available RAM if the machine has less than 20GiB RAM and to available RAM minus 10GiB if the machine has more than 20GiB RAM.

3.5.3 Increased number of crash-persisted cache items

In case the appliance crashes or has a hardware or power failure some cache items are lost and have to be reloaded from RAID or cloud storage. If this option is enabled this number is decreased and more items are persisted across crashes. This comes with at a cost of a small increase in IO to the cache device during operation. By default this is enabled for cloud storage and disabled for RAID.

3.5.4 Protect cache data with checksums against bit-flips

Bit flips in data on the cache device can irrecoverably corrupt the whole storage. If this option is enabled the data on the cache device is protected against bit flips. This comes at the cost of a significant increase of IO to the cache device. If your cache device protects data sufficiently against bit flips you can disable the software protection to increase cache device performance. By default this is enabled.

3.5.5 Compression method

Method and level used to compress data before writing it to storage (S3 or RAID). LZMA5 is the best compression while being very slow and CPU intensive even for decompression. ZSTD19 has only slighly worse compression ratio while being faster when decompressing. ZSTD7 is as fast as zlib compression while having a better compression ratio then zlib. ZSTD3 is for if you want the compression to be fast/not CPU intensive.

3.5.6 Hard idle background compression

When the backup storage is idle, it compresses some data on the cache device with a better, but CPU intensive, compression algorithm (LZMA/ZSTD19) in the background. You should disable this if you do not value the compression ratio much, want to save energy or if the CPU is shared with other VMs. By default this is enabled, except for CPU throttled Amazon EC2 instances.

3.5.7 Background compression re-write percentage

Compressed items will be rewritten to cloud storage/RAID if the background compression causes the items size to be the specified percentage or less (e.g. 80%) of the normally compressed item. This causes a bit of addional IO during operation because compressed sizes have to be tracked and it causes additional IO to the cloud storage/RAID because objects/stripes have to be rewritten with the better compressed versions. By default this is set to "95"

3.5.8 Background compression method

Method with which data gets compressed when backup storage is idle. LZMA5 gives slightly better compression ratio while being a lot slower to decompress (when accessing the compressed data from storage). ZSTD19 is recommended and the default.

3.5.9 Allow cache evictions

Allow data that is only present in cloud storage/RAID and not in cache. If this is disabled, all data can be read even if the cloud storage/RAID is lost. When disabling, make sure that the cache device size is larger than the allocated storage amount, as it will get stuck if the cache device is too small to hold all data. This is enabled by default.

3.5.10 Minimum commit delay

The appliance only commits to storage infrequently. All backups that have not been commited are lost if the appliance is suddenly rebooted e.g. because of power loss. This setting specifies the minimum amount of time in minutes that the appliance waits after a backup has been finished before it commits to backup storage. A higher value (e.g. to "60 min") increases throughput and with RAID storage decreases the amount of random IO. A lower value (e.g. "5 min") reduces the amount of last backups lost after a sudden reboot. Default: "60 min"

3.5.11 Maximum commit delay

See previous section. Before the appliance commits to storage it replicates backup metadata to backup storage. If that replication does not finish within the specified amount of time (e.g. "60 min") the appliance will commit without the replication having been finished, reducing the amount of last backups lost after a sudden reboot. Default: "60 min"

3.5.12 Compress cache

Compress the RAID/cloud cache with this compression method. This increases the size of the RAID/cloud cache in exchange for increased CPU usage and slightly worse cache performance (increased cache latency depending on unused CPU). LZO gives slightly faster compression/decompression then ZSTD while having a significantly worse compression ration. ZSTD is recommended. Default: "None"

3.6 Client installation

3.6.1 Windows/Mac OS X client installation

- Add a new Internet client on the status page.
- Download the client installer for the Internet client and send it to the new client. Alternatively, create a user for the new client (in the settings) and send the user the username/password. The user can then download the client installer from the server on the status page and install it.
- Select the backup paths you want to backup on the client or configure appropriate default directories to backup on the server (see section 7.3.4).
- The server will automatically start backups once the client is connected.

3.6.2 Automatic rollout to multiple Windows computers

First, if you want to deviate from the default backup path selection, configure the general default backup paths so that the correct folders get backed for each client (see section 7.3.4). Then install the clients using one of the following methods.

For internet clients:

You can use `https://installercreator.urbackup.org` to create a custom installer for your server that creates a client on the server, then downloads the client installer from the server and finally installs it. There are multiple options to allow e.g. silent installation, installation without tray icon, etc.

Alternatively adapt the script at `https://urbackup.atlassian.net/wiki/display/US/Download+custom+client+installer+via+Python` to your server URL and settings and create a python executable from the modified script via e.g. cx_Freeze (`http://cx-freeze.sourceforge.net/`) or py2exe. Executing the python executable on the client automatically creates a new internet client on the server, downloads a custom client and runs the installer. You could also e.g. add the silent install switch ("/S") when starting the downloaded installer such that it needs no user intervention.

3.6.3 Client installation on Linux

- Add a new Internet client on the status page (selecting the client behind NAT option).
- Follow the instructs for Linux on the page after the client was added. Alternatively download the client installer for the Internet client and send it to the new client. Another alternative is to create a user for the new client (in the settings) and send the user the username/password. The user can then download the client installer from the server on the status page and install it.
- Select the backup paths you want to backup on the client via command line ("urbackup-clientctl add-backupdir –path /" or configure appropriate default directories to backup on the server (see section 7.3.4).
- The server will automatically start backups once the client is connected.

4 Security

4.1 Server webinterface rights management

The server web interface is protected by a pretty standard user system. You can create, manage and delete accounts. Those accounts are only linked loosely to clients by rights management. An admin account can do everything including browsing file backups of all clients. The web interface allows one to create a limited account that can only browse backups and view statistics from one client. The more sophisticated rights editor can be used to allow an account to access several clients or to limit some aspects. For example you could setup an account which can do everything except browse backups. Following domains, with which you can limit or expand an account's rights, are currently available:

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

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

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

4.2 Client security

The client only processes commands if the server or the interface process supplies it with credentials. A random identity and private/public key pair.

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

pw.txt:

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

pw_change.txt

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

- Accept a new server

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

If you want to manually add a server to 'server_idents.txt' you need to remove the preceding '#I' and '#' at the end of the contents of 'server_ident.key'. After installation the 'server_idents.txt' does not exist and the client core process accepts(and adds) the first server it sees (with the public key of the server). After that no other servers with different credentials are accepted and you need to add their credentials either manually, or via clicking on the popup box, once the client has detected the new server. This prevents others from accessing files you want to be backed up in public places.
If you want to have several servers to be able to do backups of a client you have two options. Either you manually supply the server credentials to the client (by copying them into 'server_idents.txt') or you give all servers the same credentials by copying the same 'server_ident.key', 'server_ident_ecdsa409k1.priv' and 'server_ident_ecdsa409k1.pub' to all servers.

4.3 Internet mode security

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

5 Backup process

5.1 Pre and post backup scripts on client and server

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

5.1.1 Client pre and post backup scripts

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

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

5.1.2 Server post backup scripts

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

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

6 Server to server replication/remote control

Infscape Urbackup Appliance can replicate file and image backups to other Infscape UrBackup Appliances. Replication is supported passively and actively, that is you can replicate to an appliance even if it is byhind NAT or firewall and no ports are forwarded.

To connect two appliances you need to first add a replication port on one appliance entering the name of the other appliance. Then add a replication destination on the other appliance entering the authentication key shown for the new port. Replication uses the same port as configured in the settings as in the "Internet" tab (default: 55415). The appliance where you add the replication

destination will connect to the appliance where the replication port is configured.

As next step either of the appliances can be configured to replicate backups to the other appliance, either all backups or only backups of a certain client. You can also configure a window during which this replication should happen and limit the replication speed. The replication window sytax is identical to the usual window sytax (see section 7.3.1). Different speeds can be configured at different windows (see section 7.1.8).

When "Remote control" is enabled for a port/destination admins on the other appliance will be shown a merged web interface showing clients, backups and settings of both appliances.

7 Server settings

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

7.1 Global Server Settings

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

7.1.1 Server URL

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

7.1.2 Do not do image backups

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

7.1.3 Do not do file backups

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

7.1.4 Autoupdate clients

If this is checked the server will send new versions automatically to its clients. The UrBackup client interface will ask the user to install the new client version. If you check silent autoupdate (see Section 7.1.4) it will update in the background. The installer is protected by a digital signature. Default: Checked.

7.1.5 Max number of simultaneous backups

This option limits the number of file and image backups the server will start simultaneously. You can decrease or increase this number to balance server load. A large number of simultaneous backups may increase the time needed for backups. The number of possible simultaneous backups is virtually unlimited. Default: 100.

7.1.6 Max number of recently active clients

This option limits the number of clients the server accepts. An active client is a client the server has seen in the last two month. If you have multiple servers in a network you can use this option to balance their load and storage usage. Default: 10000.

7.1.7 Cleanup time window

UrBackup will do its clean up during this time. This is when old backups and clients are deleted. You can specify the weekday and the hour as intervals. The syntax is the same as for the backup window. Thus please see section 7.3.1 for details on how to specify such time windows. The default value is 1-7/3-4 which means that the cleanup will be started on each day (1-Monday - 7-Sunday) between 3 am and 4 am.

7.1.8 Total max backup speed for local network

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

All speed settings can have different values for different windows. See first how to specify a window at section 7.3.1.

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

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

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

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

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

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

7.2 Mail settings

7.2.1 Mail server settings

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

Settings	Description	Example
Mail server name	Domain name or IP address of mail server	mail.example.com
Mail server port	Port of SMTP service. Most of the time 25 or 587	587
Mail server username	Username if SMTP server requires one	test@example.com
Mail server password	Password for user name if SMTP server requires credentials	password1
Sender E-Mail Address	E-Mail address UrBackup's mail reports will come from	urbackup@example.com
Send mails only with SSL/TLS	Only send mails if a secure connection to the mail server can be established (protects password)	

Check SSL/TLS certificate	Check if the server certificate is valid and only send mail if it is	
Server admin mail address	Address for fatal errors (such as if an emergency cleanup fails or other fatal errors occur)	

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

7.2.2 Configure reports

To specify which activities with which errors should be sent via mail you have to go to the 'Logs' page. There on the bottom is a section called 'Reports'. There you can say to which email addresses reports should be sent(e.g. user1@example.com;user2@example.com) and if UrBackup should only send reports about backups that failed/succeeded and contained a log message of a certain level.
If you select the log level 'Info' and 'All' a report about every backup will be sent, because every backup causes at least one info level log message. If you select 'Warning' or 'Error' backups without incidents will not get sent to you.

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

7.3 Client specific settings

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

[1] See section 7.3.2 for time specific intervals.

Minimal number of full file backups	Minimal number of full file backups for this client. If the server ran out of backup storage space the server can delete full file backups until this minimal number is reached. If deleting a backup would cause the number of full file backups to be lower than this number it aborts with an error message.	2
Maximal number of incremental image backups	Maximal number of incremental image backups for this client. If the number of incremental image backups exceeds this number the server will start deleting old incremental image backups.	30
Minimal number of incremental image backups	Minimal number of incremental image backups for this client. If the server ran out of backup storage space the server can delete incremental image backups until this minimal number is reached. If deleting a backup would cause the number of incremental image backups to be lower than this number it aborts with an error message.	4
Maximal number of full image backups	Maximal number of full image backups for this client. If the number of full image backups exceeds this number the server will start deleting old full image backups.	5
Minimal number of full image backups	Minimal number of full image backups for this client. If the server ran out of backup storage space the server can delete full image backups until this minimal number is reached. If deleting a backup would cause the number of full image backups to be lower than this number it aborts with an error message.	2
Delay after system start up	The server will wait for this number of minutes after discovering a new client before starting any backup	0 min
Backup window	The server will only start backing up clients within this window. See section 7.3.1 for details.	1-7/0-24
Max backup speed for local network	The server will throttle the connections to the client to remain within this speed (see 7.1.8 for setting speed with window).	-
Perform auto-updates silently	If this is selected automatic updates will be performed on the client without asking the user	Checked
Soft client quota	During the nightly cleanup UrBackup will remove backups of this client if there are more backups than the minimal number of file/image backups until this quota is met. The quota can be in percent (e.g. 20%) or absolute (e.g. 1500G, 2000M).	""
Excluded files	Allows you to define which files should be excluded from backups. See section 7.3.3 for details	""
Default directories to backup	Default directories which are backed up. See section 7.3.4 for details	""
Volumes to backup	Specifies of which volumes an image backup is done. Separate different drive letters by a semicolon or comma. E.g. 'C;D'. Use the special setting "ALL" to backup all volumes and "ALL_NONUSB" to backup all volumes except those attached via USB.	C
Allow client-side changing of the directories to backup	Allow client(s) to change the directories of which a file backup is done	Checked

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

7.3.1 Backup window

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

1-7/0-24: Allow backups on every day of the week on every hour.
Mon-Sun/0-24: An equivalent notation of the above
Mon-Fri/8:00-9:00, 19:30-20:30;Sat,Sun/0-24: On weekdays backup between 8 and 9 and between 19:30 and 20:30. On Saturday and Sunday the whole time.

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

7.3.2 Advanced backup interval

Similar to the backup speed limit (see section 7.1.8) the backup intervals can be specified for different time intervals by combining them with a backup window (see previous section 7.3.1) separated by "@". The most specific backup interval will then be used.

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

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

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

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

7.3.3 Excluded files

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

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

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

```
*/Temp/*
```

You can also give the full local name

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

or the name you gave the location e.g.

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

Rules are separated by a semicolon (";")

7.3.4 Default directories to backup

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

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

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

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

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

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

Directory flags Each directory to backup has a set of flags. If you do not specify any flags the default flags will be used. Otherwise only the flags you specify are used.
Flags are specified by appending them after the backup location name (separated by "/"). Flags themselves are sparated by ",".

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

If you want to set the optional flag:

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

(The first three are default flags)

7.3.5 Virtual sub client names

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

Separate the virtual sub client names via "|". E.g.

```
system-files|user-files
```

7.4 Internet settings

Settings	Description	Default value
Internet server name/IP	The IP or name the clients can reach the server at over the internet	""
Internet server port	The port the server will listen for new clients on	55415
Connect via HTTP(S) proxy	Clients will use the configured HTTPS proxy to connect to the server. The appliance is pre-configured as HTTPS proxy allowing connections only to 127.0.0.1:55415. If you enter "127.0.0.1" as Internet server name and the appliance https address as proxy, clients will connect via https to your appliance.	-
Do image backups over internet	If checked the server will allow image backups for this client/the clients	Not checked
Do full file backups over internet	If checked the server will allow full file backups for this client/the clients	Not checked
Max backup speed for internet connection	The maximal backup speed for the Internet client. Setting this correctly can help avoid saturating the Internet connection of a client (see 7.1.8 for setting speed with window)	-
Total max backup speed for internet connection	The total accumulative backup speed for all Internet clients. This can help avoid saturating the server's Internet connection (see 7.1.8 for setting speed with window)	-
Encrypted transfer	If checked all data between server and clients is encrypted	Checked
Compressed transfer	If checked all data between server and clients is compressed	Checked
Calculate file-hashes on the client	If checked the client calculates hashes for each file before the backups (only hashes of changed files are calculated). The file then does not have to be transferred if another client already transferred the same file	Not checked
Connect to Internet backup server if connected to local backup server	If checked the client will connect to the configured Internet server, even if it is connect to a backup server on the local network.	Not checked
Do not start file backups if current estimated data usage limit per month is smaller than	The UrBackup server will not schedule file backups if the estimated amount of data the client can transfer with its current Internet connection is smaller than the specified value. See 7.4.1 on how the data usage limit is estimated.	5000 MB
Do not start image backups if current estimated data usage limit per month is smaller than	The UrBackup server will not schedule image backups if the estimated amount of data the client can transfer with its current Internet connection is smaller than the specified value. See 7.4.1 on how the data usage limit is estimated.	20000 MB

7.4.1 Data usage limit estimation

Currently there are to machisms to estimate how much data the client can transfer per month with its current Internet connection. If the client is running on Windows 10, the client is connected via Wifi and the Wifi connection is set to be metered, the data usage limit is estimated at 1GB per month. Otherwise the dataplan database at `https://github.com/uroni/dataplan_db` is used to estimate the data usage limit via the clients hostname. You can see the hostname of your Internet connection e.g. on `http://ipinfo.io/`. If it is missing in the database, please contribute to the dataplan database. The dataplan database contains both hostnames where there is estimated to

be a usage limit (e.g. mobile phone connection, plane, ...) and where they are estimated to be unlimited. If the the hostname is estimated to be unlimited, UrBackup will always start backups disregarding the setting. If there is no information about the hostname in the database, the data usage limit per month is estimated at 1TB. So if you set one of the settings above 1TB, backups will only start if the hostname is specified as unlimited in the dataplan database. If the hostname has a limit in the database, backups will only start if this limit is greater than the respective setting for the client.

7.5 Advanced settings

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

7.6 Time to wait for file system syncs to complete before throttling UrBackup

If there are a lot of simultaneous backups, backups may not complete their final step (syncing to backup storage) for a long time. The appliance will throttle all other backups to allow finished backups to complete if the final step (syncing) takes longer than the configured amount of time. Default: "60 min"

7.7 List of server IPs (proxys) from which to expect endpoint information (forwarded for) when connecting to Internet service

If clients connect via HTTPS proxy, UrBackup will expect endpoint information from the proxy if the proxy is in this list of server IPs (see `https://blog.urbackup.org/299` for details). The integrated pre-configured web server forwards this information, therefore the default is "127.0.0.1".

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

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

7.9 Debugging: Verify file backups using client side hashes

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

7.10 Maximum number of simultaneous jobs per client

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

7.11 Volumes to snapshot in groups during image backups

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

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

7.12 Volumes to snapshot in groups during file backups

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

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

7.13 Windows components backup configuration

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

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

7.14 File backup search database

The appliance creates databases for fast file searches in file backups. If the setting is "Full" it'll create and update the list of files plus an index that accelerates searches. The index uses up some relatively small amount of space on the system volume. With "Partial" it won't create or update the index, making searches slower, but also decreasing space usage significantly. With "Disabled" no file list will be created and updates, disabling the searches in file backups features for this client/group/globally.

8 Restoring backups

UrBackup protects whole machines from disaster by creating image backups and a users or servers files by creating file backups. Because the file backups size can usually be reduced by focusing on the most important data on a machine they can usually be run more often than the image backups. It makes sense to use image and file backups in tandem, backing up the whole machine less regularily than the important files via file backups.

8.1 Restoring image backups

Image backups can be restored with a Debian GNU/Linux based bootable CD/USB-stick. During image restore the machine to be restored must be reachable without network address tranlation from the server (or you forward the client ports in sections 9.1 to the restore client). While Linux supports many mainboards, disk controllers etc. you should always verify that the restore CD works on your specific hardware especially if you use exotic or new hardware. Drivers and firmware for some wireless devices and a program to configure is included but restoring via a wired network connection will be less trouble and faster and should be preferred. The restore itself is easy to use. After startup it will look for a backup server. If it does not find one, you can enter the backup server's IP/hostname and change your networking settings. After a backup server is found it will ask for a username and password. Use for example your admin account to access all clients and their image backups. Then you can select one image backup, select the disk you want to restore to and then it will restore. The target disk must be at least as large as the disk which was image backupped. Some hardware changes may cause Windows to bluescreen on startup after restore. If the startup repair fails, you may have to do a repair install using a Windows disk. You should test the different hardware combinations beforehand if you plan on restoring Windows to different hardware.

8.2 Restoring file backups

When performing file backups Infscape UrBackup Appliance creates a file system snapshot identical to the client's file system at that point in time. Those backups can be accessed via Windows file sharing (samba). The appliance automatically creates Windows file sharing users with identical passwords and access rights as the web interface users, so you can login as *admin* and access all backups and copy files in order to access or restore them.

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

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

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

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

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

9 Miscellaneous

9.1 Used network ports

The Server binds to following default ports:

Port	**Usage**	**Incoming/Outgoing**	**Protocol**
80	HTTP web interface	Incoming	TCP
55415	Internet clients	Incoming	TCP
35623	UDP broadcasts for discovery	Outgoing	UDP

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

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

9.2 Nightly backup deletion

Infscape UrBackup Appliance automatically deletes old file and image backups during the cleanup time window. Backups are deleted when a client has more incremental/full file/image backups then the configured maximum number of incremental/full file/image backups. Backups are deleted until the number of backups is within these limits again.
If the administrator has turned automatic shut-down on, this clean up process is started on server start up instead (as the server is most likely off during the night). Deleting backups and the succeeding updating of statistics can have a huge impact on system performance.

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

9.3 Emergency cleanup

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

9.4 Archiving

UrBackup has the ability to automatically archive file backups. Archived file backups cannot be deleted by the nightly or emergency clean up – only when they are not archived any more. You can setup archival under Settings->Archival for all or specific clients. When an archival is due and the the server is currently in a archival window (See 9.4.1) the last file backup of the selected type will be archived for the selected amount of time. After that time it will be automatically not archived any more. You can see the archived backups in the "Backups" section. If a backup is archived for only a limited amount of time there will be a time symbol next to the check mark. Hovering over that time symbol will tell you how long that file backup will remain archived.

9.4.1 Archival window

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

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

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

```
*;*;*;5
```

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

```
*;1;*;*
```

and "Archive every" something like 2-27 days.
One can add several values for every field by separating them via a comma such that

```
*;*;*;3,5
```

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