

Jobs and Services with napp-it 4ai (client server)

Contents

1. Jobs and Services with napp-it 4ai (client server)

1.1 Replication jobs

planned:

1.2 Sync jobs ...

2. Services

2.1 ...

1. Jobs and Services

Jobs are created in the menu **Jobs** and can be listed per job type. Jobs can be run manually, time-based via the auto.pl service, or at boot. Click the job id to change its parameters. The Jobs menu shows an overview of the latest job actions. Click a job type, e.g. *replicate*, to see the latest logs of these jobs, and click "last" to see the details of the last job run.

Backup jobs usually follow the **3-2-1 rule** as a minimum requirement: three copies of the data (ideally with snap versioning hourly, daily, ...), on two different media, one of them offsite.

napp-it 4ai client-server supports backup/restore jobs on mixed servers with any OS from FreeBSD, Linux, Illumos and OSX to Solaris and Windows. OS-independent procedures are normally difficult or impossible, as every OS provides different tools, variants and options. This is solved with the **cs-tools** for every OS, which not only provide a uniform interface but also capabilities that are normally not available.

napp-it backup jobs work with the cs-tools (single-file binaries for every OS, built with Go). The cs-tools are part of napp-it 4ai/cs but can also be used standalone, e.g. in your own scripts or batch files: <https://www.napp-it.org/cs-tools.html>

1.1 Replication jobs

Part 1: ZFS replication, the "traditional" way of transferring data with ZFS (menu Jobs > Replication).

Two ZFS datasets (filesystems or zvols) are kept in sync, locally or over the network. Replication is based on ZFS snapshots and transfers only new or changed ZFS data blocks as a data stream, which makes it extremely fast. Even petabyte high-load servers with millions of files can be kept in sync within minutes, including open files in their latest state. Over the network, mbuffer or netcat (unencrypted) or ssh (encrypted) is often used for transport. In principle you can replicate between OpenZFS, Illumos ZFS and Qnap ZFS as independent ZFS developments, but not with the original Oracle Solaris ZFS.

The **cs-stream** tool replaces, for every OS, simple netcat or mbuffer (faster buffered transfer), ssh (encryption) or pv (reduced data rate so that a slow internet connection is not blocked). Even unencrypted ZFS filesystems are transferred securely through a tunnel encrypted with a one-time password, see <https://github.com/guenther-alka/cs-stream>

When creating a replication job (any server to any server), besides times and other transport settings, **keep** and **hold** can be set per job. Replication snaps are numbered consecutively per job. The last two snaps are always kept. To continue an incremental replication, an identical snap is required on both sides (same job number, same sequence number).

hold Number of the newest replication snaps to keep: n = days, e.g. `8` , or ns, e.g. `8s` (last 8 snaps).

keep Retention rules based on hours, days, months or years, e.g. `hours:24,days:32,months:12,years:2`

Rule	Meaning
hours:24	for the last 24 hours, one snap per hour is blocked from deletion
days:32	for the last 32 days, one snap per day is blocked from deletion
months:12	for the last 12 months, one snap per month is blocked from deletion
years:2	for the last 2 years, one snap per year is blocked from deletion

No duplicate snaps are kept if several rules apply to one snap. Older replication snaps that are blocked neither by a hold nor by a keep rule are deleted.

The blocking keep rule is part of the snap name, e.g. `daten1/backup_iso/nvme_nfs@1716449451_repli_zfs-kd0830_omnios_nr_439` is a replication of the job with id 1716449451 and sequence number 439 for the day of Aug 30 of the current year.

2. Services

Services are background daemons that are started either on demand or automatically together with napp-it.

(more to come)