

## ZFS in ESGF publication

Institutions store their datasets in different formats according to their own needs. Publication to projects, such as CORDEX, require from common formats that datasets must follow. Here we present a use case of ZFS to prepare data for publication.

### Background

Suppose that we have a zfs like this:

```
someUser@someHost# zfs list -r tank/test
NAME                                USED  AVAIL  REFER  MOUNTPOINT
tank/test                           104M  66.2G   23K    /tank/test
tank/test/datasetA                   104M  66.2G  104M    /tank/test/datasetA
```

Imagine that /tank/test/datasetA contains various .nc files that for legacy reasons differ in their metadata from CORDEX required metadata and they must be modified in order to be published in ESGF. How can we effectively maintain two versions of the datasets?

### ZFS snapshots and clones

In first place, we would create a snapshot of the filesystem. This would not have any additional cost, since zfs snapshots only require disk space if the files are modified.

```
# zfs snapshot tank/test/datasetA@today
# zfs list -r tank/test
NAME                                USED  AVAIL  REFER  MOUNTPOINT
tank/test                           104M  66.2G   23K    /tank/test
tank/test/datasetA                   104M  66.2G  104M    /tank/test/productA
tank/test/datasetA@today              0      -  104M    -
```

Now, we can change dataset attributes, for example, via ncatted and we would have two datasets: the modified one "tank/test/datasetA" and the legacy one "tank/test/datasetA@today" having required only the disk space for the original dataset.

We also can make clones of the tank/test/datasetA@today in order to modify the legacy dataset, since zfs snapshots are read-only filesystems.

```
# zfs clone tank/test/datasetA@today tank/test/datasetA
# zfs list -r tank/test
NAME                                USED  AVAIL  REFER  MOUNTPOINT
tank/test                           104M  66.2G   23K    /tank/test
tank/test/datasetA                   104M  66.2G  104M    /tank/test/datasetA
tank/test/datasetA@today              0      -  104M    -
tank/test/datasetAClone               0  66.2G  104M    /tank/test/datasetAClone
```

This clone can be promoted in case we need to use the legacy dataset again.

```
# zfs promote tank/test/datasetAClone
# zfs list -r tank/test
NAME                                USED  AVAIL  REFER  MOUNTPOINT
tank/test                           104M  66.2G   23K    /tank/test
tank/test/datasetA                    0  66.2G  104M    /tank/test/datasetA
tank/test/datasetA@today              0      -  104M    -
tank/test/datasetAClone               104M  66.2G  104M    /tank/test/datasetAClone
```

For more information see <http://docs.oracle.com/cd/E19253-01/819-5461/gbcxz/index.html>.