df and du Disk space usage

CONCEPTS

df asks the filesystem how full it is; du walks the tree and adds up files. Their numbers can differ.

These are GNU flags. BSD/macOS df and du differ: no --inodes, --apparent-size or --output.

df disk is full?
↓ `which mount`
du -x which directory?
↓ `drill down`
du -d 1 which subdir / file?

Full disk but df -i shows 100% IUse? You ran out of inodes (many tiny files), not bytes.

DF: FREE SPACE

`df -h`
Human-readable sizes (powers of 1024)

`df -h` *`<path>`*
Only the filesystem containing <path>

`df -hT`
Show filesystem type column

`df -i`
Inodes instead of blocks

`df -H`
Sizes in powers of 1000 (like disk vendors)

`df -t` *`<type>`*
Only filesystems of this type

`df -x tmpfs -x devtmpfs`
Hide filesystem types

`df -l`
Local filesystems only (skip NFS)

`df --total`
Add a grand total row

`df --output=source,avail,target`
Pick the columns (see detail for fields)

DU: DIRECTORY SIZE

`du -sh` *`<dir>`*
Total for one directory

`du -h -d 1` *`<dir>`*
Size of each direct subdirectory

`du -sh *`
Size of each item in the cwd

`du -ah` *`<dir>`*
Every file, not just directories

`du -h -d 1` *`<dir>`* **`| sort -h`**
Sort smallest to biggest

`du -ah` *`<dir>`* **`| sort -rh | head`**
Top 10 space users

`du -xh -d 1` *`<dir>`*
Don't cross filesystem boundaries

`du -ch` *`<file>`***`...`**
Per item plus grand total

`du -t` *`<size>`* **`-ah`** *`<dir>`*
Skip entries under <size> (negative: over)

`du --exclude='`*`<pattern>`***`'`** *`<dir>`*
Skip matching names

`du --apparent-size -sh` *`<dir>`*
Content size, not disk blocks

`du -S` *`<dir>`*
Exclude subdirectories from each dir's size

`du --time -h -d 1` *`<dir>`*
Add newest modification time

`du --inodes -d 1` *`<dir>`* **`| sort -n`**
File count per directory

RECIPES

Find what fills the disk

```
df -h
# pick the full mount, then descend
du -xh -d 1 / | sort -h | tail
du -xh -d 1 /var | sort -h | tail
```