

Storage Infrastructure on ASC

- ▶ IBM Spectrum Scale File System
- ▶ Used for software and job scripts
- ▶ All NVMe storage
- ▶ Quotas
 - ▶ 100 GB (strict)
 - ▶ 10^6 inodes
- ▶ Accessible with **\$HOME** environment variable
 - ▶ /home/fs70XXX/username

```
1 VSC > mmlsquota --block-size auto -j home_fs7XXXX home
2                                     Block Limits
3 Filesystem type      blocks      quota      limit  in_doubt  grace |   File Limits
4 home      FILESET    32.99G      100G      100G      0         none | 289088 1000000 1000000 0         none

1 VSC > df -h $HOME
2 Filesystem      Size  Used Avail Use% Mounted on
3 home            100G   33G   68G   33% /home
4 VSC > df -i $HOME
5 Filesystem      Inodes  IUsed  IFree IUse% Mounted on
6 home            1000000 289092 710908   29% /home
```

- ▶ IBM Spectrum Scale File System
- ▶ Can be used for most kind of I/O
- ▶ Tiered storage
- ▶ Quotas
 - ▶ 100 GB (can be extended on request)
 - ▶ 10^6 inodes
- ▶ Accessible with **\$DATA** environment variable
 - ▶ `/gpfs/data/fs70XXX/username`


```
1 VSC > mmlsquota --block-size auto -j data_fs7XXXX data
2                                     Block Limits
3 Filesystem type      blocks      quota      limit  in_doubt  grace |   File Limits
4 data      FILESET    3.105T    9.766T    9.766T   3.137G   none | files quota      limit in_doubt  grace
                                     1798839 2000000 2000000    311    none

1 VSC > df -h $DATA
2 Filesystem      Size  Used Avail Use% Mounted on
3 data            9.8T  3.2T  6.7T  32% /gpfs/data
4 VSC > df -i $DATA
5 Filesystem      Inodes   IUsed  IFree IUse% Mounted on
6 data            2000000 1798835 201165   90% /gpfs/data
```

► Spectrum Scale policy



nvme pool



```
RULE ['initPlacement']  
  SET POOL 'nvme pool'
```

```
RULE 'MigrateData' MIGRATE FROM POOL 'nvme pool'  
  THRESHOLD(80,60)  
  WEIGHT(access_age_in_days*KB_ALLOCATED)  
  TO POOL 'sas pool'
```

- ▶ Should be used for local I/O
- ▶ 480 GB local SSD (VSC-4)
- ▶ 2 TB local NVMe (VSC-5)
- ▶ Accessible on compute nodes
 - ▶ /local
- ▶ Data gets deleted after the job
 - ▶ Write results to **\$DATA** to make data persistent

- ▶ Can be used for heavy local I/O
- ▶ Up to half of the nodes memory
 - ▶ ** Data resides in the shared memory (RAM) of the node **
- ▶ Accessible on compute nodes
 - ▶ /tmp
- ▶ Data gets deleted after the job
 - ▶ Write results to **\$DATA** to make data persistent

- ▶ Backups of user files is solely managed by the user
- ▶ Backups are used for disaster recovery only
- ▶ Backups are performed on the best effort basis
- ▶ Backed up filesystems:
 - ▶ \$HOME
 - ▶ \$DATA
- ▶ Project manager can exclude \$DATA filesystem from backups
 - ▶ service.asc.ac.at
- ▶ Setup continuous backups of user data to remote Machines
 - ▶ documentation: <https://docs.vsc.ac.at/storage/backup.html>

```
1 scp <source> <destination>
2 scp <me>@vsc3.vsc.ac.at:~/<myfile.c> .
3 scp myfile.c <me>@vsc4.vsc.ac.at:~/
```

```
1 rsync [OPTIONS] <source> <destination>
2 maybe useful: -av, --exclude=/daten/temp/...
```

Copy - alternative via FileZilla

