

Wikiprint Book

Title: File Systems

Subject: DEEP - Public/User_Guide/Filesystems

Version: 36

Date: 04.05.2024 12:55:55

Table of Contents

File Systems	3
Available file systems	3
Stripe Pattern Details	4
Additional infos	4
Notes	4

File Systems

Available file systems

On the DEEP-EST system, three different groups of file systems are available:

- the [?JSC GPFS file systems](#), provided via [?JUST](#) and mounted on all JSC systems;
- the DEEP-EST (and SDV) parallel BeeGFS file systems, available on all the nodes of the DEEP-EST system;
- the file systems local to each node.

The users home folders are placed on the shared GPFS file systems. With the advent of the new user model at JSC ([?JUMO](#)), the shared file systems are structured as follows:

- **\$HOME:** each JSC user has a folder under `/p/home/jusers/`, in which different home folders are available, one per system he/she has access to. These home folders have a low space quota and are reserved for configuration files, ssh keys, etc.
- **\$PROJECT:** In JUMO, data and computational resources are assigned to projects: users can request access to a project and use the resources associated to it. As a consequence, each user can create folders within each of the projects he/she is part of (with either personal or permissions to share with other project members). For the DEEP project, the project folder is located under `/p/project/cdeep/`. Here is where the user should place data, and where the old files generated in the home folder before the JUMO transition can be found.

The DEEP-EST system doesn't mount the \$SCRATCH and \$ARCHIVE file systems from GPFS, as it is expected to provide similar functionalities with its own parallel file systems.

The following table summarizes the characteristics of the file systems available in the DEEP-EST and DEEP-ER (SDV) systems:

Mount Point	File system	User or software	Cluster	Type	Global / Local	SW Version	Storage Features Details	Maximum Measured Performance (see benchmarks)	Description	Other
/home	zfs	zfs	zfs	GPFS replicated via NFS	Global				Home directory used only for configuration files.	
/opt/jump	zfs	zfs	zfs	GPFS replicated via NFS	Global				Project directory; GPFS does store files locally, but not suitable for performance intensive applications or benchmarking	
/work	work	work	work	DEEP 4-EST	Global	GeoFS 7.1.2			Work files	Also available in the NFS but only through 1 Gb network connection
/noarch	noarch	noarch	noarch	DEEP 4-EST	Local				Scratch files system for temporary data. Will be cleaned up after job finished	to use for long for storing temporary files
/home/scratch	home/scratch	home/scratch	home/scratch	DAM partition	Local				Scratch files system for temporary data. Will be cleaned up after job finished	"1.5 TB limit Optimize SSD Data Center (OC) PABLOCK policy PCH3 on 2.45-300 (S)Pare"
/home/scratch	home/scratch	home/scratch	home/scratch	DAM partition	Local				Scratch files system for temporary data. Will be cleaned up after job finished	"1.5 TB limit Optimize SSD Data Center (OC) PABLOCK policy PCH3 on 2.45-300 (S)Pare"
/scratch	scratch	scratch	scratch	DAM partition (no index)	Local				2.2 GB/s simple test in do-branch1	"3 TB in do-branch1 TB in do-branch2 do-branch3 do-branch4 do-branch5 do-branch6 do-branch7 do-branch8 do-branch9 do-branch10 do-branch11 do-branch12 do-branch13 do-branch14 do-branch15 do-branch16 do-branch17 do-branch18 do-branch19 do-branch20 do-branch21 do-branch22 do-branch23 do-branch24 do-branch25 do-branch26 do-branch27 do-branch28 do-branch29 do-branch30 do-branch31 do-branch32 do-branch33 do-branch34 do-branch35 do-branch36 do-branch37 do-branch38 do-branch39 do-branch40 do-branch41 do-branch42 do-branch43 do-branch44 do-branch45 do-branch46 do-branch47 do-branch48 do-branch49 do-branch50 do-branch51 do-branch52 do-branch53 do-branch54 do-branch55 do-branch56 do-branch57 do-branch58 do-branch59 do-branch60 do-branch61 do-branch62 do-branch63 do-branch64 do-branch65 do-branch66 do-branch67 do-branch68 do-branch69 do-branch70 do-branch71 do-branch72 do-branch73 do-branch74 do-branch75 do-branch76 do-branch77 do-branch78 do-branch79 do-branch80 do-branch81 do-branch82 do-branch83 do-branch84 do-branch85 do-branch86 do-branch87 do-branch88 do-branch89 do-branch90 do-branch91 do-branch92 do-branch93 do-branch94 do-branch95 do-branch96 do-branch97 do-branch98 do-branch99 do-branch100 do-branch101 do-branch102 do-branch103 do-branch104 do-branch105 do-branch106 do-branch107 do-branch108 do-branch109 do-branch110 do-branch111 do-branch112 do-branch113 do-branch114 do-branch115 do-branch116 do-branch117 do-branch118 do-branch119 do-branch120 do-branch121 do-branch122 do-branch123 do-branch124 do-branch125 do-branch126 do-branch127 do-branch128 do-branch129 do-branch130 do-branch131 do-branch132 do-branch133 do-branch134 do-branch135 do-branch136 do-branch137 do-branch138 do-branch139 do-branch140 do-branch141 do-branch142 do-branch143 do-branch144 do-branch145 do-branch146 do-branch147 do-branch148 do-branch149 do-branch150 do-branch151 do-branch152 do-branch153 do-branch154 do-branch155 do-branch156 do-branch157 do-branch158 do-branch159 do-branch160 do-branch161 do-branch162 do-branch163 do-branch164 do-branch165 do-branch166 do-branch167 do-branch168 do-branch169 do-branch170 do-branch171 do-branch172 do-branch173 do-branch174 do-branch175 do-branch176 do-branch177 do-branch178 do-branch179 do-branch180 do-branch181 do-branch182 do-branch183 do-branch184 do-branch185 do-branch186 do-branch187 do-branch188 do-branch189 do-branch190 do-branch191 do-branch192 do-branch193 do-branch194 do-branch195 do-branch196 do-branch197 do-branch198 do-branch199 do-branch200 do-branch201 do-branch202 do-branch203 do-branch204 do-branch205 do-branch206 do-branch207 do-branch208 do-branch209 do-branch210 do-branch211 do-branch212 do-branch213 do-branch214 do-branch215 do-branch216 do-branch217 do-branch218 do-branch219 do-branch220 do-branch221 do-branch222 do-branch223 do-branch224 do-branch225 do-branch226 do-branch227 do-branch228 do-branch229 do-branch230 do-branch231 do-branch232 do-branch233 do-branch234 do-branch235 do-branch236 do-branch237 do-branch238 do-branch239 do-branch240 do-branch241 do-branch242 do-branch243 do-branch244 do-branch245 do-branch246 do-branch247 do-branch248 do-branch249 do-branch250 do-branch251 do-branch252 do-branch253 do-branch254 do-branch255 do-branch256 do-branch257 do-branch258 do-branch259 do-branch260 do-branch261 do-branch262 do-branch263 do-branch264 do-branch265 do-branch266 do-branch267 do-branch268 do-branch269 do-branch270 do-branch271 do-branch272 do-branch273 do-branch274 do-branch275 do-branch276 do-branch277 do-branch278 do-branch279 do-branch280 do-branch281 do-branch282 do-branch283 do-branch284 do-branch285 do-branch286 do-branch287 do-branch288 do-branch289 do-branch290 do-branch291 do-branch292 do-branch293 do-branch294 do-branch295 do-branch296 do-branch297 do-branch298 do-branch299 do-branch300 do-branch301 do-branch302 do-branch303 do-branch304 do-branch305 do-branch306 do-branch307 do-branch308 do-branch309 do-branch310 do-branch311 do-branch312 do-branch313 do-branch314 do-branch315 do-branch316 do-branch317 do-branch318 do-branch319 do-branch320 do-branch321 do-branch322 do-branch323 do-branch324 do-branch325 do-branch326 do-branch327 do-branch328 do-branch329 do-branch330 do-branch331 do-branch332 do-branch333 do-branch334 do-branch335 do-branch336 do-branch337 do-branch338 do-branch339 do-branch340 do-branch341 do-branch342 do-branch343 do-branch344 do-branch345 do-branch346 do-branch347 do-branch348 do-branch349 do-branch350 do-branch351 do-branch352 do-branch353 do-branch354 do-branch355 do-branch356 do-branch357 do-branch358 do-branch359 do-branch360 do-branch361 do-branch362 do-branch363 do-branch364 do-branch365 do-branch366 do-branch367 do-branch368 do-branch369 do-branch370 do-branch371 do-branch372 do-branch373 do-branch374 do-branch375 do-branch376 do-branch377 do-branch378 do-branch379 do-branch380 do-branch381 do-branch382 do-branch383 do-branch384 do-branch385 do-branch386 do-branch387 do-branch388 do-branch389 do-branch390 do-branch391 do-branch392 do-branch393 do-branch394 do-branch395 do-branch396 do-branch397 do-branch398 do-branch399 do-branch400 do-branch401 do-branch402 do-branch403 do-branch404 do-branch405 do-branch406 do-branch407 do-branch408 do-branch409 do-branch410 do-branch411 do-branch412 do-branch413 do-branch414 do-branch415 do-branch416 do-branch417 do-branch418 do-branch419 do-branch420 do-branch421 do-branch422 do-branch423 do-

Stripe Pattern Details

It is possible to query this information from the deep login node, for instance:

```
manzano@deep $ fhgfs-ctl --getentryinfo /work/manzano
Path: /manzano
Mount: /work
EntryID: 1D-53BA4FF8-3BD3
Metadata node: deep-fs02 [ID: 15315]
Stripe pattern details:
+ Type: RAID0
+ Chunksize: 512K
+ Number of storage targets: desired: 4

manzano@deep $ beegfs-ctl --getentryinfo /sdv-work/manzano
Path: /manzano
Mount: /sdv-work
EntryID: 0-565C499C-1
Metadata node: deeper-fs01 [ID: 1]
Stripe pattern details:
+ Type: RAID0
+ Chunksize: 512K
+ Number of storage targets: desired: 4
```

Or like this:

```
manzano@deep $ stat -f /work/manzano
File: "/work/manzano"
ID: 0      Namelen: 255      Type: fhgfs
Block size: 524288      Fundamental block size: 524288
Blocks: Total: 120178676 Free: 65045470 Available: 65045470
Inodes: Total: 0      Free: 0

manzano@deep $ stat -f /sdv-work/manzano
File: "/sdv-work/manzano"
ID: 0      Namelen: 255      Type: fhgfs
Block size: 524288      Fundamental block size: 524288
Blocks: Total: 120154793 Free: 110378947 Available: 110378947
Inodes: Total: 0      Free: 0
```

See <http://www.beegfs.com/wiki/Striping> for more information.

Additional infos

Detailed information on the **BeeGFS Configuration** can be found [?here](#).

Detailed information on the **BeeOND Configuration** can be found [?here](#).

Detailed information on the **Storage Configuration** can be found [?here](#).

Detailed information on the **Storage Performance** can be found [?here](#).

Notes

- dd test @dp-dam01 of the DCPMM in appdirect mode:

```
[root@dp-dam01 scratch]# dd if=/dev/zero of=./delme bs=4M count=1024 conv=sync
1024+0 records in
1024+0 records out
4294967296 bytes (4.3 GB) copied, 1.94668 s, 2.2 GB/s
```

- The /work file system which is available in the DEEP-EST prototype, is as well reachable from the nodes in the SDV (including KNLs and ml-gpu nodes) but through a slower connection of 1 Gig. The file system is therefore not suitable for benchmarking or I/O task intensive jobs from those nodes
- Performance tests (IOR and mdtest) reports are available in the BSCW under DEEP-ER → Work Packages (WPs) → WP4 → T4.5 - Performance measurement and evaluation of I/O software → Jülich DEEP Cluster → Benchmarking reports:
[?https://bscw.zam.kfa-juelich.de/bscw/bscw.cgi/1382059](https://bscw.zam.kfa-juelich.de/bscw/bscw.cgi/1382059)
- Test results and parameters used are stored in JUBE:

```
user@deep $ cd /usr/local/deep-er/sdv-benchmarks/synthetic/ior
user@deep $ jube2 result benchmarks

user@deep $ cd /usr/local/deep-er/sdv-benchmarks/synthetic/mdtest
user@deep $ jube2 result benchmarks
```