

## **Wikiprint Book**

**Title: File Systems**

**Subject: DEEP - Public/User\_Guide/Filesystems**

**Version: 36**

**Date: 19.04.2025 17:15:53**

## Table of Contents

<b>File Systems</b>	<b>3</b>
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	User or Software	Cluster	Type	Global / Local	SW Version	Storage Feature Details	Maximum Measured Performance (see benchmarks)	Description	Other
/home	/home/username	SDV, DEEP-EST	GPFS exported via NFS	Global				Home directory, used only for configuration files.	
/opt/gpfs	/opt/gpfs/username	SDV, DEEP-EST	GPFS exported via NFS	Global				Project directory; GPFS does store large files, but not suitable for performance intensive applications or benchmarking	
/work	/work/username	DEEP-EST	BeaGFS	Global	BeaGFS 7.1.2			Work file system	Also available in the NFS but only through 1 Gb network connection
/scratch	/scratch	DEEP-EST	4fs local partition	Local*				Scratch file system for temporary data. Will be cleaned up after job finished	to use for long for storing temporary files
/home/username	/home/username	DAM partition	local SSD (NV)	Local*				Scratch file system for temporary data. Will be cleaned up after job finished	"1.5 TB local Optane SSD Data Center (DC) P4800X PCIe4, PCIe3 at 2.5T, 3D (SPhare)"
/home/username	/home/username	DAM partition	local SSD (NV)	Local*				Scratch file system for temporary data. Will be cleaned up after job finished	"1.5 TB local Optane SSD Data Center (DC) P4800X PCIe4, PCIe3 at 2.5T, 3D (SPhare)"
/scratch	/scratch	DAM partition (no node)	DCPM in replicated mode	Local*		2.2 GB/s single disk test in dc-dmcc1		Scratch file system for temporary data. Will be cleaned up after job finished	"3 TB in dc-dmcc1, 2 TB in dc-dmcc2, 1027 GB in dc-dmcc3, 1027 GB in dc-dmcc4, 1027 GB in dc-dmcc5, 1027 GB in dc-dmcc6, 1027 GB in dc-dmcc7, 1027 GB in dc-dmcc8, 1027 GB in dc-dmcc9, 1027 GB in dc-dmcc10, 1027 GB in dc-dmcc11, 1027 GB in dc-dmcc12, 1027 GB in dc-dmcc13, 1027 GB in dc-dmcc14, 1027 GB in dc-dmcc15, 1027 GB in dc-dmcc16, 1027 GB in dc-dmcc17, 1027 GB in dc-dmcc18, 1027 GB in dc-dmcc19, 1027 GB in dc-dmcc20, 1027 GB in dc-dmcc21, 1027 GB in dc-dmcc22, 1027 GB in dc-dmcc23, 1027 GB in dc-dmcc24, 1027 GB in dc-dmcc25, 1027 GB in dc-dmcc26, 1027 GB in dc-dmcc27, 1027 GB in dc-dmcc28, 1027 GB in dc-dmcc29, 1027 GB in dc-dmcc30, 1027 GB in dc-dmcc31, 1027 GB in dc-dmcc32, 1027 GB in dc-dmcc33, 1027 GB in dc-dmcc34, 1027 GB in dc-dmcc35, 1027 GB in dc-dmcc36, 1027 GB in dc-dmcc37, 1027 GB in dc-dmcc38, 1027 GB in dc-dmcc39, 1027 GB in dc-dmcc40, 1027 GB in dc-dmcc41, 1027 GB in dc-dmcc42, 1027 GB in dc-dmcc43, 1027 GB in dc-dmcc44, 1027 GB in dc-dmcc45, 1027 GB in dc-dmcc46, 1027 GB in dc-dmcc47, 1027 GB in dc-dmcc48, 1027 GB in dc-dmcc49, 1027 GB in dc-dmcc50, 1027 GB in dc-dmcc51, 1027 GB in dc-dmcc52, 1027 GB in dc-dmcc53, 1027 GB in dc-dmcc54, 1027 GB in dc-dmcc55, 1027 GB in dc-dmcc56, 1027 GB in dc-dmcc57, 1027 GB in dc-dmcc58, 1027 GB in dc-dmcc59, 1027 GB in dc-dmcc60, 1027 GB in dc-dmcc61, 1027 GB in dc-dmcc62, 1027 GB in dc-dmcc63, 1027 GB in dc-dmcc64, 1027 GB in dc-dmcc65, 1027 GB in dc-dmcc66, 1027 GB in dc-dmcc67, 1027 GB in dc-dmcc68, 1027 GB in dc-dmcc69, 1027 GB in dc-dmcc70, 1027 GB in dc-dmcc71, 1027 GB in dc-dmcc72, 1027 GB in dc-dmcc73, 1027 GB in dc-dmcc74, 1027 GB in dc-dmcc75, 1027 GB in dc-dmcc76, 1027 GB in dc-dmcc77, 1027 GB in dc-dmcc78, 1027 GB in dc-dmcc79, 1027 GB in dc-dmcc80, 1027 GB in dc-dmcc81, 1027 GB in dc-dmcc82, 1027 GB in dc-dmcc83, 1027 GB in dc-dmcc84, 1027 GB in dc-dmcc85, 1027 GB in dc-dmcc86, 1027 GB in dc-dmcc87, 1027 GB in dc-dmcc88, 1027 GB in dc-dmcc89, 1027 GB in dc-dmcc90, 1027 GB in dc-dmcc91, 1027 GB in dc-dmcc92, 1027 GB in dc-dmcc93, 1027 GB in dc-dmcc94, 1027 GB in dc-dmcc95, 1027 GB in dc-dmcc96, 1027 GB in dc-dmcc97, 1027 GB in dc-dmcc98, 1027 GB in dc-dmcc99, 1027 GB in dc-dmcc100, 1027 GB in dc-dmcc101, 1027 GB in dc-dmcc102, 1027 GB in dc-dmcc103, 1027 GB in dc-dmcc104, 1027 GB in dc-dmcc105, 1027 GB in dc-dmcc106, 1027 GB in dc-dmcc107, 1027 GB in dc-dmcc108, 1027 GB in dc-dmcc109, 1027 GB in dc-dmcc110, 1027 GB in dc-dmcc111, 1027 GB in dc-dmcc112, 1027 GB in dc-dmcc113, 1027 GB in dc-dmcc114, 1027 GB in dc-dmcc115, 1027 GB in dc-dmcc116, 1027 GB in dc-dmcc117, 1027 GB in dc-dmcc118, 1027 GB in dc-dmcc119, 1027 GB in dc-dmcc120, 1027 GB in dc-dmcc121, 1027 GB in dc-dmcc122, 1027 GB in dc-dmcc123, 1027 GB in dc-dmcc124, 1027 GB in dc-dmcc125, 1027 GB in dc-dmcc126, 1027 GB in dc-dmcc127, 1027 GB in dc-dmcc128, 1027 GB in dc-dmcc129, 1027 GB in dc-dmcc130, 1027 GB in dc-dmcc131, 1027 GB in dc-dmcc132, 1027 GB in dc-dmcc133, 1027 GB in dc-dmcc134, 1027 GB in dc-dmcc135, 1027 GB in dc-dmcc136, 1027 GB in dc-dmcc137, 1027 GB in dc-dmcc138, 1027 GB in dc-dmcc139, 1027 GB in dc-dmcc140, 1027 GB in dc-dmcc141, 1027 GB in dc-dmcc142, 1027 GB in dc-dmcc143, 1027 GB in dc-dmcc144, 1027 GB in dc-dmcc145, 1027 GB in dc-dmcc146, 1027 GB in dc-dmcc147, 1027 GB in dc-dmcc148, 1027 GB in dc-dmcc149, 1027 GB in dc-dmcc150, 1027 GB in dc-dmcc151, 1027 GB in dc-dmcc152, 1027 GB in dc-dmcc153, 1027 GB in dc-dmcc154, 1027 GB in dc-dmcc155, 1027 GB in dc-dmcc156, 1027 GB in dc-dmcc157, 1027 GB in dc-dmcc158, 1027 GB in dc-dmcc159, 1027 GB in dc-dmcc160, 1027 GB in dc-dmcc161, 1027 GB in dc-dmcc162, 1027 GB in dc-dmcc163, 1027 GB in dc-dmcc164, 1027 GB in dc-dmcc165, 1027 GB in dc-dmcc166, 1027 GB in dc-dmcc167, 1027 GB in dc-dmcc168, 1027 GB in dc-dmcc169, 1027 GB in dc-dmcc170, 1027 GB in dc-dmcc171, 1027 GB in dc-dmcc172, 1027 GB in dc-dmcc173, 1027 GB in dc-dmcc174, 1027 GB in dc-dmcc175, 1027 GB in dc-dmcc176, 1027 GB in dc-dmcc177, 1027 GB in dc-dmcc178, 1027 GB in dc-dmcc179, 1027 GB in dc-dmcc180, 1027 GB in dc-dmcc181, 1027 GB in dc-dmcc182, 1027 GB in dc-dmcc183, 1027 GB in dc-dmcc184, 1027 GB in dc-dmcc185, 1027 GB in dc-dmcc186, 1027 GB in dc-dmcc187, 1027 GB in dc-dmcc188, 1027 GB in dc-dmcc189, 1027 GB in dc-dmcc190, 1027 GB in dc-dmcc191, 1027 GB in dc-dmcc192, 1027 GB in dc-dmcc193, 1027 GB in dc-dmcc194, 1027 GB in dc-dmcc195, 1027 GB in dc-dmcc196, 1027 GB in dc-dmcc197, 1027 GB in dc-dmcc198, 1027 GB in dc-dmcc199, 1027 GB in dc-dmcc200, 1027 GB in dc-dmcc201, 1027 GB in dc-dmcc202, 1027 GB in dc-dmcc203, 1027 GB in dc-dmcc204, 1027 GB in dc-dmcc205, 1027 GB in dc-dmcc206, 1027 GB in dc-dmcc207, 1027 GB in dc-dmcc208, 1027 GB in dc-dmcc209, 1027 GB in dc-dmcc210, 1027 GB in dc-dmcc211, 1027 GB in dc-dmcc212, 1027 GB in dc-dmcc213, 1027 GB in dc-dmcc214, 1027 GB in dc-dmcc215, 1027 GB in dc-dmcc216, 1027 GB in dc-dmcc217, 1027 GB in dc-dmcc218, 1027 GB in dc-dmcc219, 1027 GB in dc-dmcc220, 1027 GB in dc-dmcc221, 1027 GB in dc-dmcc222, 1027 GB in dc-dmcc223, 1027 GB in dc-dmcc224, 1027 GB in dc-dmcc225, 1027 GB in dc-dmcc226, 1027 GB in dc-dmcc227, 1027 GB in dc-dmcc228, 1027 GB in dc-dmcc229,

## 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
```