


**TeraGrid™**
[About](#)
[News](#)
[Outreach](#)
[Science Gateways](#)
[User Support](#)
[User Support](#)
[Resources](#)
[Access](#)
[Compute](#)
[Data & Visualization](#)
[Portal](#)

## User Support & Documentation

### Data Storage

[Home](#) > [User Support](#) > [Data](#) > Data Storage

#### File Management Overview

Users control where their data resides; selection of appropriate storage is important for efficient management of job output. Data can be stored in a user's home directory, in a temporary location or in archival (mass) storage.

- Home directories are permanent, but have relatively small quotas. This space is visible by all the nodes in the cluster, including the login nodes. Data in this area is never purged.
- Temporary storage in scratch is shared and has more space than home directories, but depends on concurrent use by other users. Scratch space is accessible to all nodes in a TeraGrid cluster, including the login nodes. The storage is shared with other users that may running on the cluster at the same time, and may be physically co-located with other logical storage areas.
- Parallel file systems are fast, large, but temporary; this space is accessible to all nodes in a TeraGrid cluster/machine, including the login nodes. This storage may be physically co-located with other logical storage areas (e.g. TG\_CLUSTER\_SCRATCH).
- Archival (mass) storage offers the advantage of long-term retention of large amounts of data, is accessible from all sites, but offers slower access time; those with gridFTP server front-ends offer faster data transfer time. SRB is a special system for movement and mass storage of data and for organization of collections. It is available from the sites specified in the table below.

It is helpful to familiarize yourself with the characteristics of file systems at locations where you choose to store data in order to make the best choices:

- Speed
- Visibility
- Quotas
- Backup & Purge policies

#### Speed

In general, if you plan on moving data across sites, use a fast file system, that is, a parallel file system for temporary storage of intermediate to large quantities of data.

#### Visibility

Differences exist in the underlying storage locations that are mounted at each site. Environment variables that are in common use across all sites provide a common syntax for referring to the location of storage that is available at each site, thereby hiding the underlying differences in their paths. See the [Environment](#) page of this guide for more information on environment variables. However, to move data from one resource to another, explicit paths must be specified rather than using environment variables.

#### Quotas

Home directories at each TeraGrid site have enforced quotas. Scratch and parallel file systems share the total space. In the table below, space available on scratch and parallel file systems is dependent upon concurrent use. Use the `df` command to display available space before sending large data outputs to these file systems.

#### Backup & Purge Policies

Home directories at each site are backed up; scratch, parallel file systems, and archival storage are not. Regardless of backup and purge policies at individual sites, users are advised to back up valuable data frequently. Backups may be made to the user's local system or to other archival storage on the TeraGrid; in the case of HPSS, requesting two copies of data will result in the creation of copies on two separate tapes. Please note that the tapes are physically stored in the same facility.

See the Data Storage Policies and Specs table below for detailed backup and purge policies at each site. For detailed information on each site's storage policies, type `tg-policy -data` from the command line to see detailed data policy information of the site you are logged into. Some examples of running `tg-policy` at various sites can be found [here](#).

#### On this page

- [Data Storage File Systems & Policies](#)
- [Archival Storage](#)

- [SRB](#)
- [HPSS](#)
- [golem](#) - PSC
- [DiskXtender](#) - NCSA
- [Data Migration Facility](#) - TACC

#### Related links:

- [Data Resources](#)
- [Current Data Collections](#)
- [to-policy Usage and Examples](#)

#### Need Help?

- Phone Toll-free  
1.866.907.2383
- [Submit a Ticket \(online form -- fastest\)](#)
- [Submit a Ticket via email](#)
- [TeraGrid Knowledge Base](#)



### Data Storage File Systems & Policies

Review the table below as well as the [Resource page](#) and details at specific sites

Site	File System	Type	Quota	Purge	Backup
IU	<b>Big Red</b>				
	HOME (/N/u/<username>/BigRed)	NFS	10 TB	none	nightly
	TG_CLUSTER_PFS	GPFS			
	TG_CLUSTER_GPFS	GPFS		60 days*	None; user can get account on IU HPSS for backup
	TG_CLUSTER_SCRATCH (/N/gpfsbr/<username>)	GPFS			
	TG_NODE_SCRATCH (/scratch)	Local scratch		60 days*	None; user can get account on IU HPSS for backup
	TG_COMMUNITY (/N/comm)	NFS	As requested		
IA-64 Linux Cluster	<b>IA-64 Linux Cluster</b>				
	TG_CLUSTER_HOME		5 GB	never	nightly
	TG_CLUSTER_PFS (scratch)	GPFS	none	Files not accessed in 5 days	none

NCSA	TG_NODE_SCRATCH (Node local scratch)		none	files are automatically deleted after batch job is completed	none
	TG_CLUSTER_SCRATCH (NFS Scratch)	NFS	none	Time since last access: File Size Removed after ≤ 25 MB 14 days > 25 GB 5 days	none
	<b>SGI Altix</b>				
	TG_CLUSTER_HOME		50 GB	never	nightly
	TG_CLUSTER_SCRATCH (scratch)	CXFS	none	Time since last access: File Size Removed after ≥ 10 MB 4 days < 10 GB 14 days	none
	<b>Xeon Linux Cluster</b>				
	TG_CLUSTER_HOME		5 GB	never	nightly
	TG_CLUSTER_SCRATCH (scratch)	Lustre	none	Time since last access: File Size Removed after ≥ 1 MB 1 day < 1 GB 14 days	none
	TG_NODE_SCRATCH		none	Files are automatically deleted after batch job is completed	none
	<b>IBM p690</b>				
	TG_CLUSTER_HOME		2 GB	never	nightly
	TG_CLUSTER_SCRATCH (scratch)	GPFS	none	Time since last access: File Size Removed after ≥ 1 MB 1 day < 1 GB 14 days	none
	<b>UniTree</b>				
		archival	none		
ORNL	TG_CLUSTER_HOME	NFS	none	none	daily
	TG_CLUSTER_SCRATCH	PVFS2	none	none	none
	\$TG_NODE_SCRATCH	local	none	14 days	none
File System		Type	Quota	Purge	Backup
PSC	<b>HP Marvel (Rachel)</b>				
	HOME		1 GB	never	nightly
	LOCAL		6 TB / machine	each compute node volatile, use in job or move to repository ASAP	
	<b>Cray XT3 (BigBen)</b>				
	HOME		2 GB	never	nightly
	TG_CLUSTER_SCRATCH	Lustre	none	as needed	never
Purdue	<b>golem</b>				
	SGI Origin 300	archival	contact admin for files ≥2TB	never	
	TG_CLUSTER_HOME			never	nightly
	scratch (TG_CLUSTER_SCRATCH)	NFS			user
SDSC	TG_CLUSTER_HOME			never	nightly
	scratch (TG_CLUSTER_SCRATCH)	NFS		4 days	user
	GPFS (DataStar) (TG_CLUSTER_PFS) (TG_CLUSTER_GPFS)	parallel (fast)	130 TB (shared)	4 days	user
	GPFS (IA64 Cluster) (TG_CLUSTER_PFS) (TG_CLUSTER_GPFS)	parallel (fast)	50 TB (shared)	4 days	user
	SAMQFS (in development)	archival			
	HPSS	archival			user
	SRB	archival/ collections			
File System		Type	Quota	Purge	Backup
TACC	<b>Lonestar</b>				
	TG_CLUSTER_HOME	NFS	200 MB	none	daily
	TG_CLUSTER_SCRATCH	Lustre parallel filesystem	750 GB	10 days	none
	TG_NODE_SCRATCH	Local /tmp	None FS=56 GB	At End of Job	none
	TG_COMMUNITY	NFS	None FS=538 GB	none	none
	TG_EXAMPLES (subset of TG_COMMUNITY)	NFS	None FS=538 GB	none	none
	<b>Maverick</b>				
	TG_CLUSTER_HOME	Local UFS /home	100 MB	none	daily
	TG_CLUSTER_PFS	Local UFS /home	100 MB	none	daily
	TG_CLUSTER_SCRATCH	QFS /work	None FS=8.7 TB	none	none
	TG_NODE_SCRATCH	Local UFS /tmp	27 GB	none	none
	TG_COMMUNITY	NFS /san/hpc	256 GB FS=6.0 TB	none	none

	TG_EXAMPLES (subset of TG_COMMUNITY)	NFS /san/hpc/ tg-examples	256 GB FS=6.0 TB	none	none
UC/ ANL	TG_CLUSTER_HOME	NFS		never	nightly
	TG_CLUSTER_SCRATCH	GPFS parallel		21 days, as needed	user
	TG_NODE_SCRATCH	local scsi		21 days, as needed	user
	TG_CLUSTER_PFS	GPFS parallel		21 days, as needed	user
	TG_CLUSTER_PVFS	PVFS parallel		21 days, as needed	user
	TG_CLUSTER_GPFS	GPFS parallel		21 days, as needed	user
	TG_GLOBAL_PFS	GPFS parallel	see <a href="#">GPFS documentation</a> for quota, purge, and backup policies		
	TG_GLOBAL_GPFS	GPFS parallel			

Temporary storage locations on the TeraGrid and [permanent \(archival or mass\) storage](#) in a UNIX-like file system (UniTree at NCSA or HPSS at SDSC) or in a collection that is organized for searching and sharing (SRB).

This guide will point out where interfaces to data and transfer commands vary and will provide links to more detailed explanations at individual sites. Please call or send e-mail to the [Help Desk](#) for individualized assistance.



## Archival Storage

A wide range of data resources are available for researchers; computation allocations have access to limited archival storage. However, for larger collections, long-term preservation of data, and staging of data collections, independent data resource collections are available. See the [Data Resources page](#) for details and links to allocations information.

## Storage Resource Broker (SRB)

The Storage Resource Broker, a data management tool, may be used for storage, replication, archiving, third-party copying and movement of large TeraGrid data sets across distributed, heterogeneous storage systems. It uses its own set of commands. SRB is available to users through either a compute allocation, through which users are automatically entitled to an SRB account, or through an independent data allocation for storing data and data collections on tape or disk that does not require a compute allocation. Any TeraGrid user can use an SRB client to download data that is available in public collections.

## SRB servers and clients on TeraGrid

- The command line client is available from all TeraGrid compute resources
- A web-based client is available from any web browser
- Other client interfaces, written in a wide range of languages are available for use from within Windows for incorporation into applications
- SRB servers are available for data resources at SDSC

## Obtaining an SRB account (SDSC)

To activate or request an SRB account, go to the [SDSC SRB Account Activation](#) page. SDSC users who have received a compute allocation automatically receive SRB accounts; you must activate your account before using it. To see instructions specific to computing allocations, go to the [SDSC SRB Account Activation](#) page and select "I have a compute allocation, and I will be using gsi-enabled authentication (certificate proxies)." Then click on "View Instructions". Users with only data/SRB allocations should select "I have an SRB data allocation."

## Using our SRB account

1. See the [SDSC SRB User Guide](#) or the [SRB software site](#) for information about using SRB.
2. Using SRB from the command line requires the use of a configuration file called **.MdasEnv**. This file is created in a .srb directory under your home directory. Copy and paste the following sample into your ~/srb/.MdasEnv file, replacing the parameters in brackets with those you received with your SRB account information:

```
mdasCollectionHome '/home/<user-name>.teragrid'
mdasDomainName 'teragrid'
srbUser '<user-name>'
#AUTH_SCHEME 'ENCRVPT1'
AUTH_SCHEME 'GSI_AUTH'
srbHost 'srb.sdsc.edu'
srbPort '7321'
defaultResource '<resource>'
SERVER_DN '/C=US/O=NACPI/OU=SDSC/CN=Storage Resource Broker/USERID=srb'
```



## High Performance Storage System (HPSS)

HPSS is available at SDSC and IU. Each site has an HPSS user guide with additional information.

### SDSC

Archival storage is available at SDSC on the [HPSS archival storage system](#). HPSS uses HSI, an FTP-like interface. HSI supports wildcards for local and HPSS pathname pattern matching, and provides recursion for many commands, including the ability to store, retrieve, and list entire directory trees, or change permissions on entire trees. It may be used interactively or in batch mode, and may be included in UNIX pipes.

### Indiana (IU)

IU's [High Performance Storage System \(HPSS\)](#) holds 2.4 Petabytes of disk-cached, web-publishable tape storage. This storage space can be accessed via GridFTP to [gridftp.archive.iu.teragrid.org](#) from any site.

## golem - PSC

Golem is a SGI Origin 300 that runs a combination disk and tape archival system. Initially, files moved to golem reside on disk. Factors such as file size and time of last access determine when a file gets migrated to tape. When you access a migrated file, it is automatically read in from tape.

Golem supports GridFTP transfers using the address [tg-gridftp.psc.teragrid.org](#). For more information, including common commands used for file transfer, visit the [golem user guide](#).

## DiskXtender Mass Storage System (MSS) - NCSA

TeraGrid users may also use the [DiskXtender mass storage system](#), housed at NCSA, for permanent storage of large file sets. A [proxy certificate](#) is required for access. Globus-url-copy (see [Grid:FTP](#) page or [globus-url-copy at NCSA](#)), and [UberFTP](#) transfer methods are supported. Note that gsiscp is not supported.

## Data Migration Facility - TACC

To provide long-term, reliable data storage, TACC operates a four processor SGI Origin 2000 with four gigabytes of fast, dynamic RAM and 1.3 terabytes of high performance, high availability fiber channel RAID-3 disks. This archive system is configured for dedicated file service using SGI's Data Migration Facility (DMF) to migrate files to a tape archival system. The disk farm on the Origin 2000 acts as a cache for recently accessed files. Files are permanently stored in two StorageTek PowderHorn 9310 automated cartridge systems. TACC's data archive is exported as a filesystem named /archive mounted on each supercomputer at the TACC via Network File System (NFS). Fast access from the high performance computing (HPC) machines at TACC is provided via the 800 megabit per second HIPPI (High Performance Parallel Interface) local area network. See the [TACC archival user guide](#) for more information.



The TeraGrid project is funded by the [National Science Foundation](#) and includes 11 partners:  
[Indiana](#), [LQNI](#), [NCAR](#), [NCSA](#), [NICS](#), [ORNL](#), [PSC](#), [Purdue](#), [SDSC](#), [TACC](#) and [UC/ANL](#).

Please email [help@teragrid.org](mailto:help@teragrid.org) with questions or comments.

The TeraGrid 1.0 Transitional CSS compliant.

