

Research Storage Environment (2020-)

The following table shows bulk storage appliances purchased for JCU research.

Owner, Operator	Usable Size	Vendor Model	End-of-Life
JCU, JCU	600TiB	DELL SC4020	Q4-2021
JCU, QCIF	516TiB	DDN SFA7990E	Q1-2024

Each HPC server has internal SSD storage, as follows:

Server(s) purpose	Raw size
Login/Compute nodes	480GB (RAID-1) + 7.68TB (RAID-0)
GPU servers	240GB (RAID-1) + 960GB (RAID-0)
ESXi servers	1600GB (RAID-5)

Note: Existing GPU and ESXi servers will be decommissioned in 2021.

The following filesystems are configured specifically for JCU research use

Filesystem	Size	Shared via	Details
/gpfs01	516TiB	GPFS	Cache for ARDC/RDSI/QCIF approved allocations
/home	512TiB	NFS	Researchers' home directories
/scratch	80TiB	NFS	<i>Scratch</i> space, shared using NFS
/sw	200GiB	NFS	Software (read-only for researchers)
/fast/tmp	7TiB	-	Local (to node) SSD scratch space
/tmp	300GiB	-	Local (to node) SSD <i>scratch</i> space

For efficiency across JCU, it is best if your computational research jobs are run under /tmp, /fast/tmp or /scratch. One completed, only files that are likely to have long-term value should be moved to your home directory or ARDC/RDSI/QCIF approved allocation.

JCU's DELL storage array is very much entry-level - purchased to provide capacity, not performance.

While the DDN storage array has a much higher performance potential, its performance is limited by the number of disks (64) and SSDs (12) installed in it.

The /gpfs01 filesystem is a **medici** cache - primary copy of all data is held by QCIF (locations in Brisbane and Springfield). There is no guarantee of recall time for *offline* (not currently on JCU cache space) files. Outages at QCIF, or network issues between JCU and QCIF, will lead to **IO errors** being seen whenever you try to access/use *offline* files. Retrying your task(s) at a later time is what I would suggest. JCU HPC staff do not have access to fix issues with our medici cache that relate to QCIF(Bne) services.

Your perception of filesystem performance decreases with increasing inodes (file count). In mid-2019, HPC held over 500,000,000 nodes across 8 filesystems - having less than 10,000,000 inodes per filesystem is desirable.

As of 1-Jan-2020, the following default quotas are configured

Filesystem	Account type	Default Quotas
/home	JCU	5TiB & 250,000 inodes
/scratch	JCU	5TiB & 1,000,000 inodes
	External	100GiB & 100,000 inodes
	Delegate	100GiB & 100,000 inodes

A research data management strategy is being considered. Until this strategy is in place, individual quotas may be increased. Longer term, there may be a merit-approval and/or purchasing process for individual users obtaining an increased quota.

Click for further information about what an [inode](#) is

ARDC/RDSI/QCIF Quotas

Many JCU researchers have applied for, and been awarded, an RDSI storage allocation for their research. There is a default inode quota of 1,000,000 on all allocations. Higher inode quotas may be configured, however, requestors should consider breaking up their allocation (requesting another allocation) rather than increase inode quota. History has shown that allocations with high inode counts experience more problems than those with lower inode counts.

Allocation	Disk Quota		inode Quota
	QCIF(Bne)	JCU Cache	JCU Cache
Q0033	15T	5T	10,000,000
Q0036	20T	1T	1,000,000
Q0042	100T	55T	11,000,000
Q0043	30T	1T	1,000,000
Q0044	100T	20T	10,000,000
Q0050	12T	2T	1,000,000

Allocation	Disk Quota		inode Quota
	QCIF(Bne)	JCU Cache	JCU Cache
Q0114	10T	2T	10,000,000
Q0123	2T	2T	1,000,000
Q0124	5T	5T	1,000,000
Q0125	10T	2T	1,000,000
Q0145	35T	10T	2,000,000
Q0148	1T	1T	1,000,000
Q0149	20T	5T	1,000,000
Q0150	12T	6T	1,000,000

Allocation	Disk Quota		inode Quota
	QCIF(Bne)	JCU Cache	JCU Cache
Q0166	5T	5T	1,000,000
Q0171	1T	1T	1,000,000
Q0184	45T	8T	1,000,000
Q0188	16T	6T	1,000,000
Q0189	20T	5T	1,000,000
Q0195		512G	5,000,000
Q0199	5T	5T	1,000,000
Q0200	35T	27T	20,000,000

Allocation	Disk Quota		inode Quota
	QCIF(Bne)	JCU Cache	JCU Cache
Q0201	10T	5T	1,000,000
Q0202	10T	5T	1,000,000
Q0203	10T	5T	1,000,000
Q0208	150T	10T	1,000,000
Q0210	150T	40T	6,000,000
Q0213	200G	200G	1,000,000
Q0214	40T	40T	10,000,000
Q0217	2T	2T	5,000,000
Q0219	50T	8T	1,000,000
Q0220	60T	60T	10,000,000

Allocation	Disk Quota		inode Quota
	QCIF(Bne)	JCU Cache	JCU Cache

Q0222	8T	5T	1,000,000
Q0230	20T	5T	1,000,000
Q0252	2T	2T	1,000,000
Q0261	60T	9T	1,000,000
Q0262	512G	512G	1,000,000
Q0266	1T	1T	1,000,000
Q0269	10T	5T	4,000,000

Allocation	Disk Quota		inode Quota
	QCIF(Bne)	JCU Cache	JCU Cache
Q0308	500G	500G	1,000,000
Q0309	500G	500G	1,000,000
Q0365	1T	1T	1,000,000
Q0477	2T	2T	1,000,000
Q0478	4T	4T	1,000,000
Q0634	7T	7T	6,000,000
Q0638	15T	7T	1,000,000
Q0750	5T	5T	1,000,000
Q1116	40T	10T	1,000,000

Q0037 (Wallace2) was a 1PiB RDSI allocation that appeared to be closed (unused) at the time of migration of files to our medici environment. This combine with problems encountered during said migration, lead to a decision to breakup Q0037 into allocation that would only be housed at QCIF (Bne). To make the transfer occur in the fastest possible time content on Q0037 was also placed into many large, compressed tarballs. Given the size and workflows associated with the Wallace2 project, JCU's medici cache would struggle with reintroducing allocations below back into our cache should a researcher request it - extra capacity should be added to our DDN SFA7990E appliance should such a request come through.

	Q2024	Q2025	Q2026	Q2027	Q2028	Q2029	Q2030	Q2031	Q2536
Storage (TiB)	15	50	35	120	5	25	25	3	40
inodes	5m	10m	3m	9m	10m	10m	10m	10m	10m

Contact stefano.montanaril@jcu.edu.au if you need up increase in quota or are interested in obtaining a new allocation. ARDC/RDSI/QCIF allocations are recommended for any researcher or research group that has a storage requirement in excess of default HPC user quotas.

Other research data storage options.

HPC provides a location for the primary/trusted copy of your research data only. Other examples of locations which are responsible locations for your primary/copy of research data are:

1. Public cloud - e.g., OneDrive and AWS.
2. Education/Private Cloud - e.g., AARNet CloudStor and ARDC services.
3. Institutional/Other facilities that offer equivalent or higher level protection mechanisms than JCU HPC.

Personal computing devices (e.g., PCs, USB disks, personal NAS appliance) are definitely NOT a responsible location for the primary/trusted copy of your research data. From a technical perspective, this is because changes can occur on disks without being noticed and corrected (e.g., search the internet for "silent data corruption"). One of the bigger issues that has been brought to the attention of authorities is the lack of ability to verify/validate research if personal computing devices are involved (e.g., when they house the primary copy of data).

ESSENTIAL UNDERSTANDING (Policy)

All JCU researchers need to read and understand information within <https://www.jcu.edu.au/policy/information-and-communications-technology/information-communication-technology-acceptable-use-policy>