

# QCIF-JCU Additional Node

QCIF-JCU opted to follow the QCIF-UQ lead (choosing SGI) in the first round of RDSI node funding spend. A significant delay in order submission was due in part to changing personnel/positions and a typo (Email didn't go to the correct person). For ~\$235K, we received the following:

Hardware	Qty	Details (per unit)	Configuration
SGI IS5500	1	2 active-active storage controllers 16 x (8Gb/s FC ports) 180 x (3TB, 7200RPM NL-SAS disks)	Dynamic Disk Pools in use 17 x (20TB physical volumes) Approx. 320TB of <i>cooked</i> storage
Brocade 6510 FC Switch	2	48 x (16Gb/s FC ports) 24 x (8Gb/s SFPs licensed)	Two fabrics (for redundancy)
SGI C2108 (NAS services)	2	2 x (Intel E5-2670 CPU – 8-cores, 2.6GHz), 256GB DDR3 memory (1600MHz), 2 x (200GB SSD), 8 x (8Gb/s FC ports), 2 x (QDR IB ports), 2 x (10GbE IP ports), 4 x (1GbE IP ports)	RAID-1 system volume All FC paths are active Striped slicing of volumes (XVM) Filesystems: 100TB, 3x50TB, 38TB, 25TB, 13TB 13TB filesystem used for Quorum <i>disk</i> Bonding (active-failover) used for public networks
Edge-corE GbE Switch	1	48 x (1GbE IP ports)	Used for BMC connectivity
Redundant Power Unit	1	4 x (DC power outlets)	Provides redundant power for switches

A limited amount (300TB) of vault licensing has also been purchased - tape media will be purchased at a later time. IB connected storage arrays could not be purchased due to insufficient capacity on current IB switches (1U each) infrastructure. Growth of available IB port count was expected to be 2x the cost of an FC switched solution.

An existing VMware cluster is currently being used to provision many services for JCU researchers and their collaborators - e.g., websites and databases. The servers behind this cluster are fully populated (in terms of PCIe slot and FC/IP port consumption) and designed for reliability/resiliency rather than performance. As a result of their production status and being fully populated already, these servers cannot be connected to RDSI storage. In Sep-2013, a single server was built from spare parts to provide urgently needed RDSI services (using RDSI storage). This unmaintained server is only expected to be required until March-2014 (at the latest).

Note that JCU is not involved in NeCTAR and there is no plan to downgrade from VMware to OpenStack. Researchers with demands that cannot be met by JCU infrastructure will be pushed to NeCTAR facilities (e.g., QCloud).

## Relationship to QCIF-UQ

A plan exists to asynchronously replicate some/most QCIF-JCU data to the QCIF-UQ node. An identified group of researchers will be able to work on data locally (on HPC cluster) and have data presentation services visible to the outside world from QCloud (QCIF-UQ). Where possible, a redundant front-end service will be available on JCU infrastructure in order to minimize outage impacts (whether scheduled or unscheduled).

## RDSI Services'

DashNet infrastructure is not scheduled to be online at JCU until Mar-2014. Until then, we have no 10GbE connectivity and only have two GbE (Cu) ports available for use outside the JCU firewall. The ESX server providing virtual resources was built by JCU from spare parts as a way of offering services outside the JCU firewall now.

Service Description	Built on	Status	Connectivity	Service IP Address
Block storage (FC)	Physical	Operational	Internal	
Block storage (iSCSI)	Physical	Considering	Internal	
Fileshares (CIFS)	Virtual	Planning	Outside FW	202.6.115.222
Fileshares (NFS)	Physical	Operational	Inside FW*	137.219.15.64
File transfer (FTP)	Virtual	Operational	Inside FW*	202.6.115.211
File transfer (GridFTP)	Virtual	Configuring	Outside FW	202.6.115.212
File transfer (SCP, SFTP)	Physical	Operational	Inside FW*	137.219.15.64
File transfer (Tsunami-UDP)	Virtual	Operational	Outside FW	202.6.115.211
Easy access (XtreemFS, OwnCloud)	Virtual	Considering	Outside FW	202.6.115.223
Easy access (WebDAV)	Virtual	Planning	Outside FW	202.6.115.221
Authentication (LDAP)	Virtual	Considering	Outside FW	Software decisions to be made

\* Temporary firewall (FW) exceptions are already in place for services running inside the firewall at present. Services available outside the JCU firewall have an IP address in the 202.6.115.192/26 network.

## Ingest Summary (Recent)

Much research data was moved onto a HSM managed filesystem (JCU HPC) prior to receipt of RDSI infrastructure. No accurate record was made of times spent on ingest tasks - durations below have been calculated from relevant Email records. The dates provided below correspond to when the last upload took place. There is a large number of tarballs and gzipped files in the initially ingested data (due to lack of space at the source). Most of these will probably be expanded/uncompressed in future.

Date	Project	Allocation (GB)	Ingest source	Ingested (GB)	Ingested (inodes)
14-Oct	CC Impacts	100,000	HPC - HSM managed filesystem	40,320	3,142,365
14-Oct	CTBCC	50,000	HPC - HSM managed filesystem	43,284	925,705
14-Oct	Wallace	50,000	HPC - HSM managed filesystem	15,775	44,067,322
10-Sep	Weather	25,000	HPC - HSM managed filesystem	15,047	8,733,742
09-Sep	MangroveWatch	25,000	8 USB disks (so far)	3,811	884,375

### Issues

- MangroveWatch - Thousand of files on the USB disks were found to be corrupted (couldn't be read).
- MangroveWatch - Some of the files on USB disks were clearly not related to research.
- CTBCC/Weather/CC Impacts/Wallace - Many millions of files (source) were on tape and only 8 tape drive mechanisms exist.
- CTBCC/Weather/CC Impacts/Wallace - Lack of free disk space at the source slowing ingest, some recall commands needed to be repeated up to 7 times.