

MIR Shared Storage Facility – ZFS Storage System Description December 08, 2017

System Overview

The Mallinckrodt Institute of Radiology has implemented a ZFS storage system that will support the data storage needs of radiology research. The ZFS solution is comprised of two separate systems – a production system and an archive system. The “live” production system is used for the storage of data that is currently being processed or with immediate and unpredictable access needs. The archive solution is for the deep-archiving of data with no immediate access needs.

The ZFS production system is comprised of both primary and disaster recovery (DR) hardware. The primary system has been engineered as a fully redundant system, no single hardware component failure will cause the system to fail. The system is currently configured to provide approximately 1PB of usable storage. Solid state technology is used as both read and write cache to enhance performance. The backend near-line SAS storage is configured as RAID7 to ensure data reliability and integrity.

Data is replicated every 30 seconds from the primary system to the physically disparate DR system. Snapshots of data that has been changed or deleted on the ZFS system are taken five nights a week. Thirty snapshots or 6 weeks of changed data are stored insuring that any accidental changes or deletions of data can be recovered.

The ZFS archive system is comprised of low-speed disks. These disk arrays will be powered down and only brought online when data retrieval is needed.

Storage is leased on a quarterly basis. Lease pricing is as follows:

Production system

FY18: \$180 per year

FY19: \$160 per year

FY20: \$160 per year

Archive system

FY18: \$60 per year

FY19: \$53 per year

FY20 \$53 per year