

Storage Policy

This document defines the CRC's storage policy (SP) for ND Faculty regarding storage acquisition, installation, and support. Under this program the CRC provides storage specification, acquisition expertise, installation, deployment, performance tuning, security via access control, and hardware/software maintenance. The CRC will review the SP annually and make revisions in line with CRC Faculty Advisory Council (FAC) members.

Service Definitions

- **High Performance Storage (CRC /scratchX)** – Distributed storage space associated with an active analysis or computation. The storage is for **temporary** use; files older than published thresholds are deleted. Disk drives are redundant but data is **not** backed up.
-- *X in scratchX represents the max file lifetime (in days) after which it is auto deleted.*
- **User Space Storage (CRC AFS Cell)** – User space storage that has persistence as long as the user intends to reference the data for ongoing research projects. Accessibility, sharing and capacity are higher priority than performance. The space is redundant, backed up, and can be configured for sharing via the web upon request.
- **Backup Storage** – In addition to the backup of data in CRC provided user space, the CRC can provide backup for critical servers (not workstations) used in research.
- **Archival Storage** – Data stored indefinitely for historical reference as dictated by government or commercial contract obligations. The CRC does not provide archival storage but instead works with the ND Library through their CurateND service.

Technical Definitions

- **Redundant/Resilient Storage** – Data has some degree of local redundancy to demonstrate resiliency to hardware failure. This may be accomplished via hardware or software RAID (Redundant Array of Independent Disks).
- **Backup Storage** – Duplicate data copies (snapshots in time) are stored in a separate location from the primary copy at set time intervals. Backups are kept for up to 90 days and are not archival. This is often accomplished via offsite tape or disk array.

Storage Requests

The CRC maintains a list of high performance, user space, and archival storage resources on our wiki <http://crc.nd.edu/wiki>. Continued support for existing resources is revalidated annually. The CRC reserves the right to discontinue support for underutilized, underperforming, or vendor end of life storage resources with a commitment to help transfer persistent data to new resources. Storage resources proposed for discontinued support will be mailed to the CRC user's list at least 3 months prior to support termination. Faculty can submit to the CRC a justification for continued storage support. All individual storage requests above Tier 1 (see below) must be submitted to crsupport@nd.edu and include the specific research, quota, and time requirements (an indefinite storage timeline is not viable).

Storage Security and Data Classification Policy

The CRC recommends that only data classified as public or internal should be processed in shared computing resources. If sensitive or highly sensitive data need to be processed then special arrangements must be made and dedicated resources may be required.

CRC users are required to responsibly comply with all ND information security policies http://policy.nd.edu/policy_files/InformationSecurityPolicy.pdf.

CRC Allocations

Tier 1 (no cost)

Each new CRC user is granted:

- 100GB quota of user space (crc.nd.edu AFS cell) - *automatically*
- 500GB quota of high performance space (/scratch*) - *by request*

Tier 2 (no cost, written justification required, research group quota governs)

Individual users can request up to:

- 1TB quota of high performance scratch space

Research Groups (Individual Faculty/PI) can request up to^(1,2):

- 4TB of persistent user or group space (in AFS)
- Temporary HPC scratch* storage based on scientific need and availability

Note 1: For clarification we calculate allocations based on the PI/Research Group sponsor; this applies directly to all T&R faculty. In the case of RAP appointments the member is required to demonstrate that they are an official PI on externally funded grants.

Note 2: Unfortunately we cannot accommodate "pooling" of storage resources as too many groups have tried to add faculty members to the "pool" who do not perform computational research.

Tier 3 (faculty pays for hardware and software, written justification required)

Research groups requiring greater than 4TB of persistent storage may purchase additional storage in increments of 1TB based on recent CRC purchase prices. Based on the availability of funds, the CRC covers the personnel costs of operations. The faculty member pays for base hardware and software to meet the storage requirements.

For FY 2016 -

High Performance Storage: \$565 per usable³ TB (1 time cost for 5yr HW warranty)

User Space Storage: \$200 per usable³ TB (1 time cost for 5yr HW warranty)

Backup Space Storage: \$50 per usable³ TB (tape cost)

³usable accounts for capacity loss due to RAID or incremental backup

CRC Storage Acquisition and Operations

Specialty storage acquisition/hardware/software/operation requests specific to a research requirement will be negotiated on a case by case basis.