

Gadi /scratch File Management

Version: PUBLIC 10 Jun 2022

Overview

To meet the increasing demand for more space on /scratch from our researchers and to reduce the number of 'forgotten' temporary files, NCI is introducing a new file management policy for the Gadi /scratch file system. This new policy will automatically clean up files not accessed for more than 100 days and so create more space available for research projects.

This policy change will facilitate greater fairness in the use of temporary scratch storage for all NCI projects.

The removal of old /scratch files is a three stage process.

Step 1 - Files not accessed for more than **100 days** are moved from project directories on /scratch into a quarantine space. *Once a file has been moved to quarantine it will no longer be accessible to its owner, the project, or to any HPC jobs run by the project or collaborating projects with read access, regardless of file permission settings.*

Step 2 - Files remain in quarantine for **14 days**. *During this quarantine period files maybe recovered by the **file owner** and restored to active use if needed.*

Step 3 - Any files remaining in quarantine at the end of the 14-day quarantine period will be **deleted**. Deletion from the quarantine space is automatic, and final. *After a file is deleted, it cannot be recovered. All users are reminded that the /scratch file system is intended to store working files only. Data that researchers or projects wish to keep for an extended period of time must be copied from the /scratch filesystem to the project's /g/data space, archived to massdata (tape) or downloaded to local storage.*

Implementation Schedule

This new /scratch file management procedure is a significant shift in the way the /scratch file system is managed, so it will be progressively introduced in May-June 2022, with full implementation from 1 July 2022. NCI must implement this /scratch file system change before the 2022 Q3 (July) downtime as it will support essential tuning and reconfiguration of the /scratch file system in a full production, peak performance capacity.

The introduction of this /scratch file system policy may quarantine a large number of files for any projects which have accrued substantial /scratch usage on Gadi. To make this process more manageable for users the policy will be implemented in stages according to the following schedule:

1. 17 May 2022 : Files within /scratch project directories which have not been accessed for **365 days** will be quarantined. Any quarantined files that are not recovered by project users within the initial 14-day quarantine period 17-31 May will be automatically deleted at the end of the quarantine period, 31 May.
2. 14 Jun 2022 : Files on /scratch which have not been accessed for **100 days** will be quarantined. Any files remaining in quarantine at the end of the 14-day quarantine period will be automatically deleted.
3. From 21 Jun 2022 /scratch files not accessed for more than 100 days will be quarantined on a continuous, rolling basis. The automated quarantine-expiry process will run each day, adding any files with atime greater than 100 days to the quarantine space, and deleting any files that have resided in quarantine for 14 days.

File Management Utility: *nci-file-expiry*

A new utility, *nci-file-expiry*, can be used to identify and restore files from quarantine space. The document below contains more information and usage examples. The command option "--help" can also provide usage and syntax information.

Note that the accumulation of many files on /scratch for some users can cause *nci-file-expiry* to run slowly, or to run out of working memory before the utility prints a report. A user having ~10**5 or more files in the "Warning" state may need to run *nci-file-expiry* via a PBS job script to allow sufficient cpu time and memory to generate an inventory of scratch files. This potential memory issue will only affect projects that have an extremely large number of scratch files. It is expected to ease after the initial quarantine-expiry cycles in May-June.

It is important to note that users must manage any files they own. *A file in quarantine space can be identified and restored to active status only by its owner (userid), regardless of the files permission settings.* It is not possible to implement a role-based (e.g. Lead CI) scratch file management utility at this time. Lead CI's should ask all project members to manage their own files. Use the "ls -al" command to list file ownership and permission information.



GadiSystem-Gad...22-1629-37.pdf

(Right click to open this PDF document in a new browser tab.)

Key Considerations

Important points to remember about the new /scratch file management process:

- All users and project teams are encouraged to be proactive in managing their temporary storage on /scratch. The /scratch file system is intended for temporary, working storage. If you need persistent storage, use the /g/data or massdata systems, or download to your local filesystem.
- If you have files on /scratch that you do not need please delete them as soon as possible. You do not need to wait for the automated quarantine-expiry process script to run.
- Users must manage files their own files. It is not possible to implement role-based, project access for /scratch file management at this time.
- All projects with active NCMAS allocations now have /g/data directories. Default allocation is 2.5 GB/KSU.
- Stakeholder projects will get /g/data allocations per entitlements and demand. For /g/data allocations, please contact your scheme allocation manager. NCI (help@nci.org.au) can help put you in touch with the appropriate scheme managers if needed.
- Project default scratch quotas will be raised at the time of July quarterly maintenance. (Note that default scratch quotas are still necessary to protect file system stability.)
- Projects expecting to use large amounts of scratch capacity will still need to request appropriate quotas. Consultation with NCI HPC and Storage groups may be needed for projects intending to use peak-scale scratch capacity.
- Large scratch requests (e.g. ≥ 10 TB) from projects with compute allocations of less than 1 MSU/year, or projects without demonstrated track records, will be accommodated in phases with advice from NCI Storage and HPC groups.
- In general, exceptions to the scratch file expiry policy will not be permitted. If you need advice or assistance to prepare for the full implementation of scratch file quarantine-expiry contact NCI user support: help@nci.org.au.
- NCI may make adjustments to quarantine-expiry parameters to ensure operational stability of Gadi and the /scratch file system. If any changes become necessary, they will be communicated in advance to the user community via the Gadi MOTD, the NCI newsletter, and directed email information campaigns.