

# Slow Response to the "ls -l" Command

## Category: Resolved

**Problem:** Users have experienced a slow response to the Unix command `ls -l`, often waiting for minutes for the command to complete.

## Status: Resolved

### Actions

**Updated 02.03.12** - As of December 14th, 2011, all of the RAID's on which the Lustre filesystem resides were updated to a new controller, which improves the IOPS performance. For additional information, see the article [Lustre Best Practices](#).

If you are still experiencing issues with this problem, contact the NAS Control Room: (800) 331-8737, (650) 604-4444, [support@nas.nasa.gov](mailto:support@nas.nasa.gov).

**Updated 08.31.11** - NAS systems staff have deployed a second new filesystem (`/nobackupp2`). In addition, they completed the transition of users to `/nobackupp1` on August 8th.

**Updated 07.27.11** - NAS systems staff have deployed a new filesystem (`/nobackupp1`) with a new RAID controller, which improves the IOPS performance. Users on `/nobackupp10` are being transitioned to the new filesystem. Users on other `/nobackup` filesystems will be transitioned to similar, new RAID controllers.

**Updated 06.28.11** - NAS systems staff are testing new RAID controllers, which will improve Input/Output Operations Per Second (IOPS) performance, and result in an improved stats-per-second rate.

**Updated 05.26.11** - NAS systems staff began an immediate investigation, and determined the source of the problem was related to access of the metadata on the Lustre filesystem, and took the following actions:

- Implemented a change to preload metadata information into the buffer cache on all of the `/nobackup` filesystems. This improved the stats-per-second rate.
- Evaluated the solid-state drives (SSDs) to improve performance of stat calls. There are two use cases for accomplishing this: one approach is to use SSDs for filesystem journals, the other is to use them for filesystem metadata. SSDs were deployed on `/nobackupp10` for the filesystem journal.

- Reduced the default stripe count to 1 and the default stripe size to 4 MB. This change resulted in fewer stat calls and, over time, improve the `ls -l` response. (Effective January 13, 2011).
- Developed a modified `gnu tar` command, `/usr/local/bin/mtar`, which is Lustre stripe aware and will create a tar file or extract files with an appropriate sized stripe count.

## Tips

**Updated 08.02.11** - Since the default stripe count was set to 1 to improve overall performance of Lustre, users need to keep in mind that they sometimes should use larger stripe counts, particularly with large files.

- You can set the stripe count on a directory or create a file with a particular stripe count by using the command `lfs setstripe`.
- We recommend using the `mtar` command in place of `tar` whenever you create or extract from a tar file on Lustre, as `mtar` will automatically choose appropriate stripe counts.

See more information on Lustre striping in the articles [Lustre Basics](#) and [Lustre Best Practices](#).

## Background

Previously, by default, the Lustre filesystem striped data across object storage targets (OSTs) in 1-MB chunks. So, a 4-MB file would be spread across four OSTs, resulting in four stat calls to get the size of that one file.

---

Article ID: 272

Last updated: 30 Jul, 2012

Updates on Issues -> Resolved -> Slow Response to the "ls -l" Command

<http://www.nas.nasa.gov/hecc/support/kb/entry/272/?ajax=1>