

Using 'mtar' to Create or Extract Tar Files on Lustre

Category: Lustre on Pleiades

NAS's in-house developed **mtar** program is a modification of GNU **tar** version 1.25. It is exactly equivalent to **tar** except that, if it detects a Lustre filesystem, then it restripes files as they are "tarred" and/or "untarred" for better performance. Specifically:

- The stripe count of files extracted on a Lustre filesystem will be dynamically selected based on the original file size, so, small files will be extracted with small stripe counts and large files will be extracted with large stripe counts
- Tar files created on a Lustre file system will have a stripe count based on the sum of the sizes of all component files

TIP: We recommend using **mtar** in place of **tar** when creating or extracting from a tar file on Lustre.

Currently, the number of stripes set by **mtar** is essentially the number of gigabytes of that file (for disk storage, 1 GB = 10⁹ bytes), limited by the number of object storage targets in that Lustre filesystem.

Tar files created with **gzip** (**-z**), **bzip2** (**-j**), and arbitrary compression (**--use-compress-program**) will preserve the striping of the uncompressed tar file.

Using mtar

mtar is available in `/usr/local/bin` on the Pleiades front-ends (pfe[20-27], bridge[1-4]). Usage of **mtar** is exactly the same as **tar** and you don't have to know anything extra, as it all happens automatically.

The following example demonstrates its usage and the comparison between **mtar** and **tar**. Note that some output has been removed for clarity.

```
%ls -l *_file
-rw----- 1 zsmith s0101 16800000112 Aug  3 14:58 17g_file
-rw----- 1 zsmith s0101  1200000008 Aug  3 14:51 2g_file
-rw----- 1 zsmith s0101           1215 Aug  3 15:04 2k_file
```

```
%lfs getstripe *_file
17g_file
lmm_stripe_count:  1
2g_file
lmm_stripe_count:  1
2k_file
```

lmm_stripe_count: 1

Notice that the default stripe count is 1 on all Pleiades Lustre filesystems.

Comparison of tar and mtar	
tar	mtar
<pre>%tar cvf tar.tar 17g_file 2g_file 2k_file %lfs getstripe tar.tar tar.tar lmm_stripe_count: 1</pre>	<pre>%mtar cvf mtar.tar 17g_file 2g_file 2k_file %lfs getstripe mtar.tar mtar.tar lmm_stripe_count: 19</pre>
<pre>%tar xvf tar.tar %lfs getstripe *_file 17g_file lmm_stripe_count: 1 2g_file lmm_stripe_count: 1 2k_file lmm_stripe_count: 1</pre>	<pre>%mtar xvf mtar.tar %lfs getstripe *_file 17g_file lmm_stripe_count: 17 2g_file lmm_stripe_count: 2 2k_file lmm_stripe_count: 1</pre>
<pre>%tar xvf mtar.tar %lfs getstripe *_file 17g_file lmm_stripe_count: 1 2g_file lmm_stripe_count: 1 2k_file lmm_stripe_count: 1</pre>	<pre>%mtar xvf tar.tar %lfs getstripe *_file 17g_file lmm_stripe_count: 17 2g_file lmm_stripe_count: 2 2k_file lmm_stripe_count: 1</pre>
<pre>%tar zcvf tar.tgz tar.tar %lfs getstripe tar.tgz tar.tgz lmm_stripe_count: 1</pre>	<pre>%mtar zcvf mtar.tgz mtar.tar %lfs getstripe mtar.tgz mtar.tgz lmm_stripe_count: 19</pre>

Notice that the **tar**-created archive has a default stripe count, while the **mtar**-created archive has a stripe count based on the sizes of component files. In addition, **tar**-extracted files all have a default stripe count, while **mtar**-extracted files have a variable stripe count depending on size. Also notice that using **mtar** with compression preserves striping of the uncompressed tar file.

The **mtar** script was created by NAS staff member Paul Kolano.

Article ID: 291

Last updated: 17 Dec, 2012

Computing at NAS -> Best Practices -> Lustre on Pleiades -> Using 'mtar' to Create or Extract Tar Files on Lustre

<http://www.nas.nasa.gov/hecc/support/kb/entry/291/?ajax=1>