Lustre Filesystem Statistics in PBS Output File

Category: Lustre on Pleiades

For a PBS job that reads or writes to a Lustre file system, a Lustre filesystem statistics block will appear in the PBS output file, just above the job's PBS Summary block. Information provided in the statistics can be helpful in determining the I/O pattern of the job and assist in identifying possible improvements to your jobs.

The statistics block lists the job's number of Lustre operations and the volume of Lustre I/O used for each file system. The I/O volume is listed in total, and is broken out by I/O operation size.

The following Metadata Operations statistics are listed:

- Open/close of files on the Lustre file system
- Stat/statfs are query operations invoked by commands such as **`ls -l`**
- Read/write is the total volume of I/O in gigabytes

The following is an example of this listing:

```
============================================================
                LUSTRE Filesystem Statistics
------------------------------------------------------------
nbp10 Metadata Operations
       open     close      stat    statfs    read(GB)   write(GB)
       1057      1058      1394         0           2          14
Read   4KB   8KB  16KB  32KB  64KB  128KB  256KB   512KB   1024KB
         9     3     1     0     1      0      3       2      319
Write  4KB   8KB  16KB  32KB  64KB  128KB  256KB   512KB   1024KB
       138    13     1    11    36      9     21      37    12479
------------------------------------------------------------
Job Resource Usage Summary for 11111.pbspl1.nas.nasa.gov

    CPU Time Used            : 00:03:56
    Real Memory Used         : 2464kb
    Walltime Used            : 00:04:26
    Exit Status              : 0
```

The read and write operations are further broken down into buckets based on I/O block size. In the example above, the first bucket reveals that nine data reads occurred in blocks between 0 and 4 KB in size, three data reads occurred with block sizes between 4 KB and 8 KB, and so on. The I/O block size data may be affected by library and system operations and, therefore, could differ from expected values. That is, small reads or writes by the program might be aggregated into larger operations, and large reads or writes might be broken into smaller pieces. If there are high counts in the smaller buckets, you should

investigate the I/O pattern of the program for efficiency improvements.

For tips for improving Lustre I/O, see Lustre Best Practices for multiple tips to improve the Lustre I/O performance of your jobs.

Article ID: 289
Last updated: 30 Jul, 2012
Computing at NAS -> Best Practices -> Lustre on Pleiades -> Lustre Filesystem Statistics in PBS Output File
http://www.nas.nasa.gov/hecc/support/kb/entry/289/?ajax=1