



**Performance Testing Results**  
**As of March, 2008**



Terascale, Inc.  
145 Bodwell St  
Avon, MA  
[www.terascala.com](http://www.terascala.com)

## Lustre Performance on Terascale Hardware

### Summary:

A series of performance tests were run on a customer installation of the Terascale RTS 1000 storage appliance and up to 64 client nodes. The overall solution consisted of a single Metadata server with 4 Object Store Targets. A series of tests consisting of both striped and non-striped arrays were run.

Peak performance was over 420 MB / sec per OST (Object Store Target) while average performance across a range of clients exceeded 360 MB /sec per OST

Note that for these particular tests the non-striping approach showed better performance. Striping benefits in situations where a large number of clients are accessing large files simultaneously.

### Testing Environment:

Terascale RTS1000:

1 Metadata server - dual core 2.4 GHz AMD Processor, 4GB RAM, SDR InfiniBand HCA, 6 500 GB disk drives

4 Object Store Targets - dual core 2.4 GHz AMD Processor, 4GB RAM, SDR InfiniBand HCA, 8 500 GB disk drives

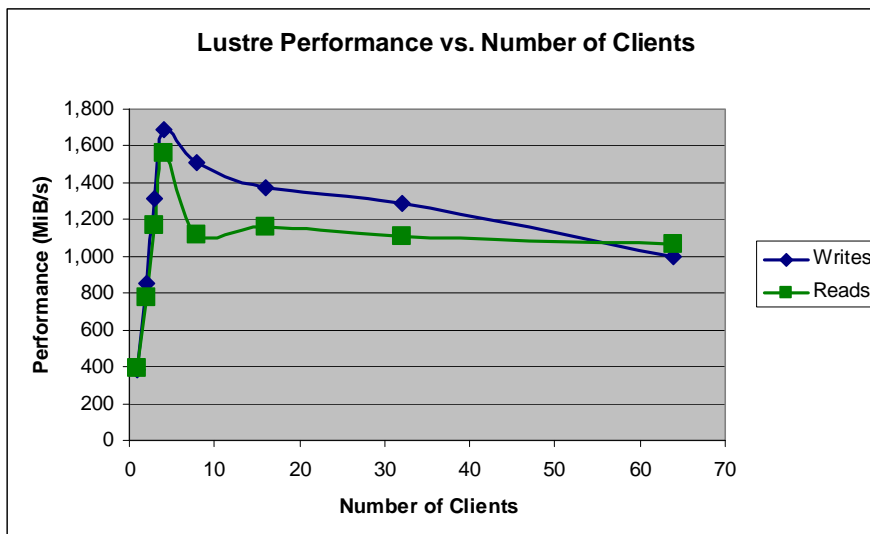
Client nodes:

Dual core 3.6 GHz AMD Processor, with 6 GB of RAM, and a SDR InfiniBand HCA

File size for testing: 10GB

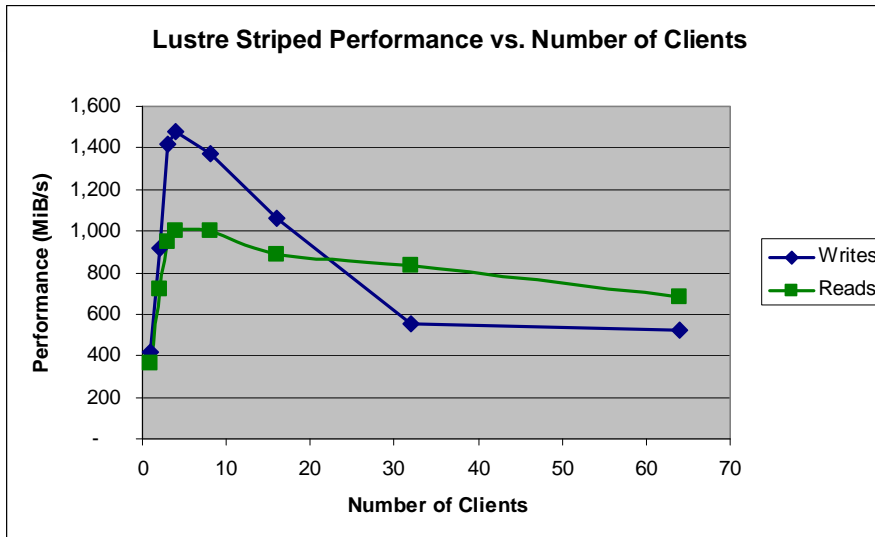
### Non-Striping Tests:

Clients were divided up equally across all 4 OSTs. Tests were run 5 times for 1-4 clients, and 3 times for more than 4 clients.



### Striping Tests:

A variety of stripe sizes were tested (1M, 2M, 4M, 8M, 32M, and 256M). Once the number of clients was high enough, the stripe size does not appear to significantly impact the performance. The numbers below represent the averages of the different stripe sizes.



### A Note about Performance Variability:

Performance varies depending on the location of the file on the disk and on disk fragmentation. It is possible for a client to write a large file to a particular Lustre OST (Object Storage Target) at the rate of 400 MB/s, and then write another large file to the same OST at the rate of 300 MB/s.

We have run tests that write to the whole OST. Based on these tests, the long-term average Write performance is 367 MiB/s per OST. The performance generally grows linearly with the number of OSTs.

### About Terascale

*Terascale develops high throughput, high capacity and cost-effective storage solutions. Its unique storage appliance approach is changing the dynamics of the performance driven computing market, enabling existing users to do more for less while enabling new users to maximize the capabilities of their processing infrastructure. Founded in 2005, Terascale is based in Avon, MA.*



Terascale, Inc.  
145 Bodwell St.  
Avon, MA 02322  
Tel: 508-588-1501  
Email: [sales@terascal.com](mailto:sales@terascal.com)  
[www.terascal.com](http://www.terascal.com)