Storage Challenges at Los Alamos National Lab

Aaron Torres, Gary Grider, Adam Manzanares, Alfred Torrez, Brett Kettering, Meghan McClelland - LANL

John Bent - EMC

Abstract

There yet exist no truly parallel file systems. Those that make the claim fall short when it comes to providing adequate concurrent write performance at large scale. This limitation causes large usability headaches in HPC.

Users need the following capabilities missing from current parallel file systems:

- **High bandwidth for large parallel IO using various IO patterns**
- **Minimal application tuning for IO performance**
- **Low latency for interactive work**

Problems with Modern Parallel File Systems

1. **Concurrence issues**
2. **Magic numbers**
3. **Metadata distribution**
4. **Lack of Quality of Service (QOS)**

Parallel File System Performance

GPFS

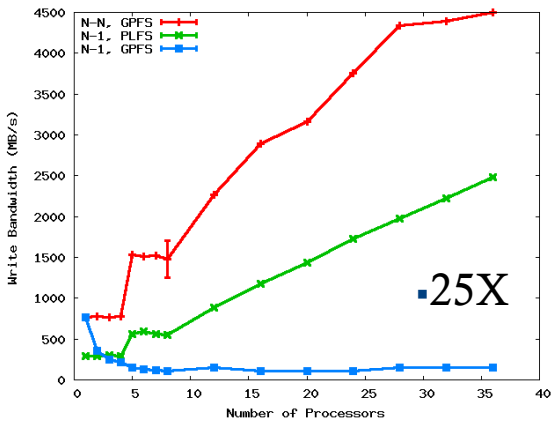

PanFS

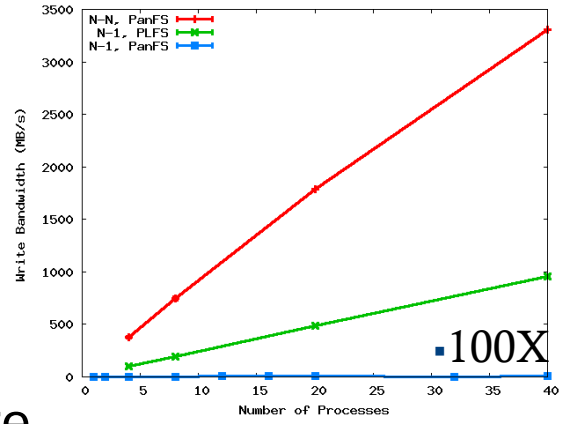

Lustre

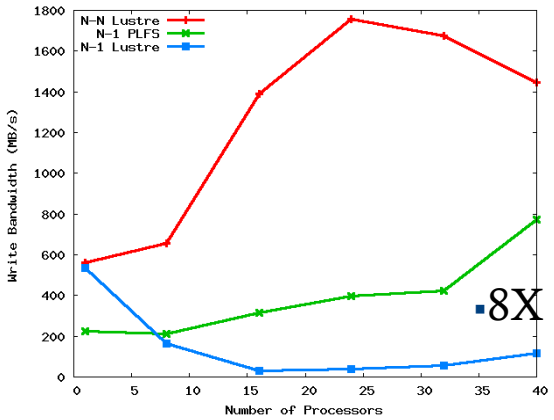

PanFS at Scale

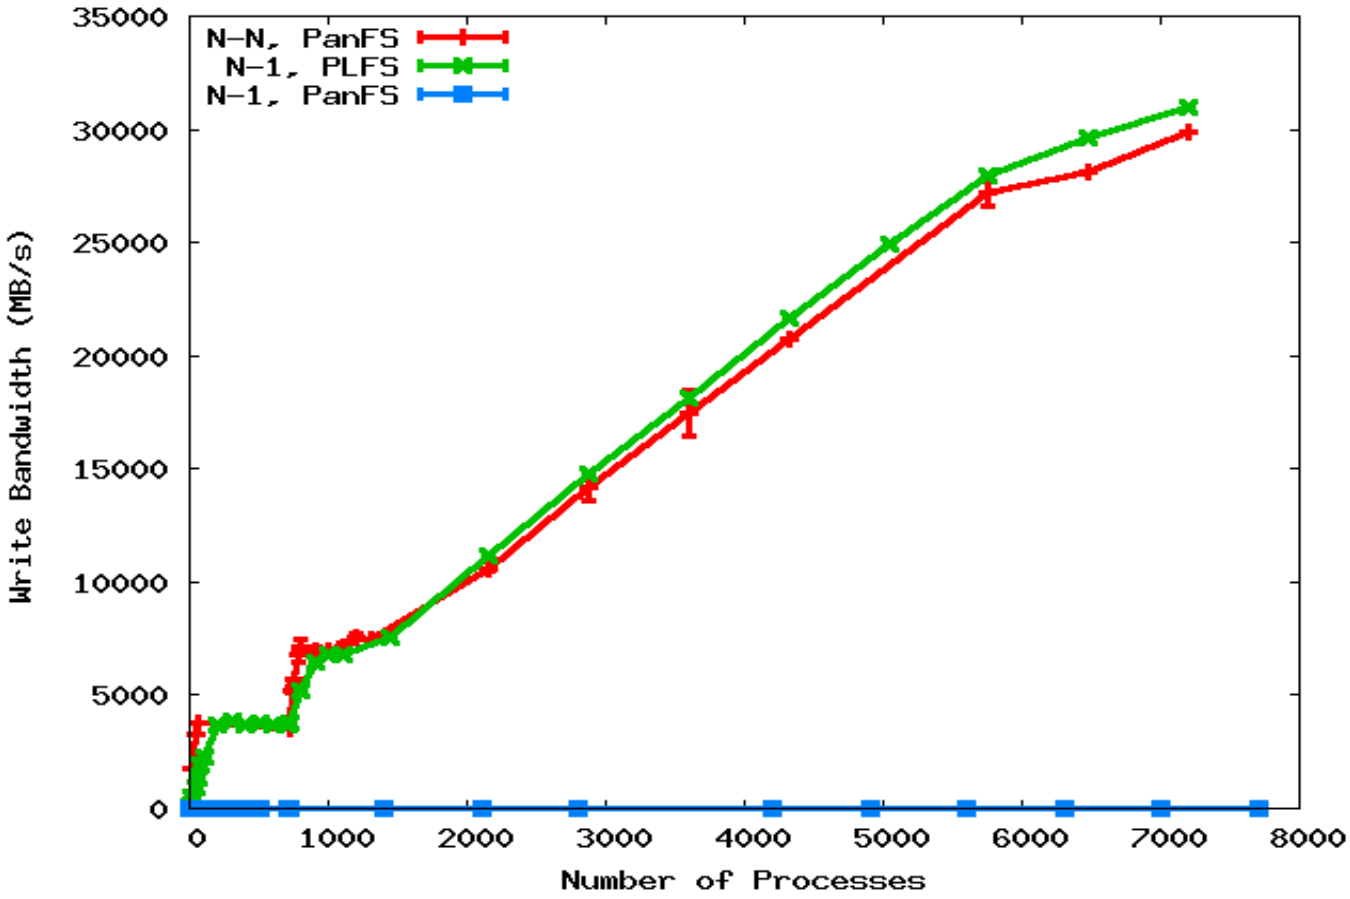

Lustre at Scale

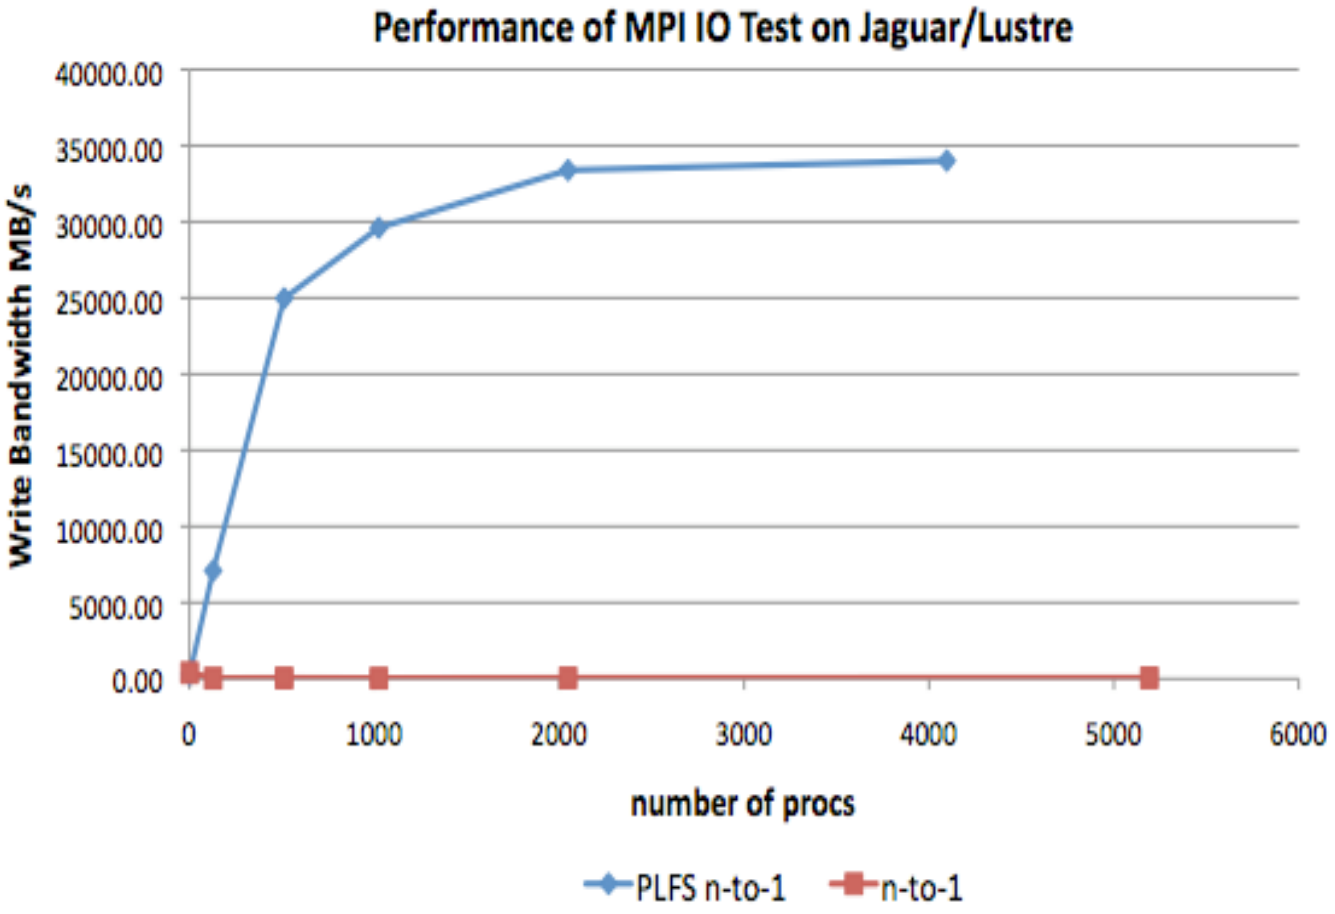

Magic Numbers

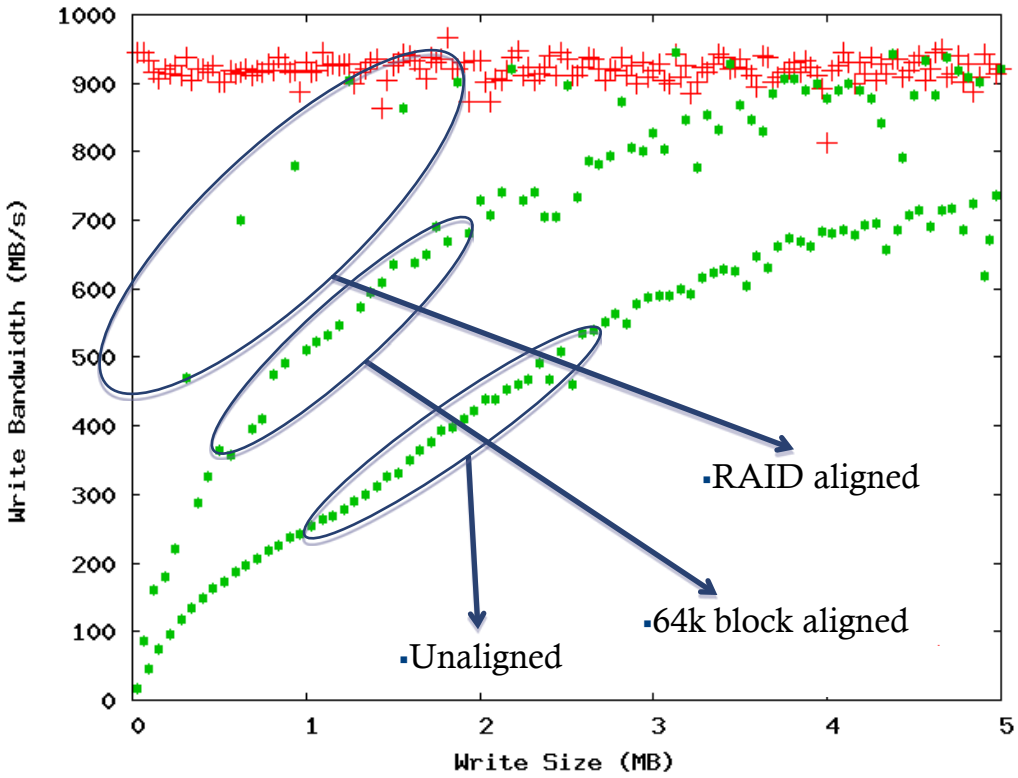

- With PLFS
- Without PLFS

Metadata Distribution

- **Hash Metadata across directories**
- **Federated File systems**

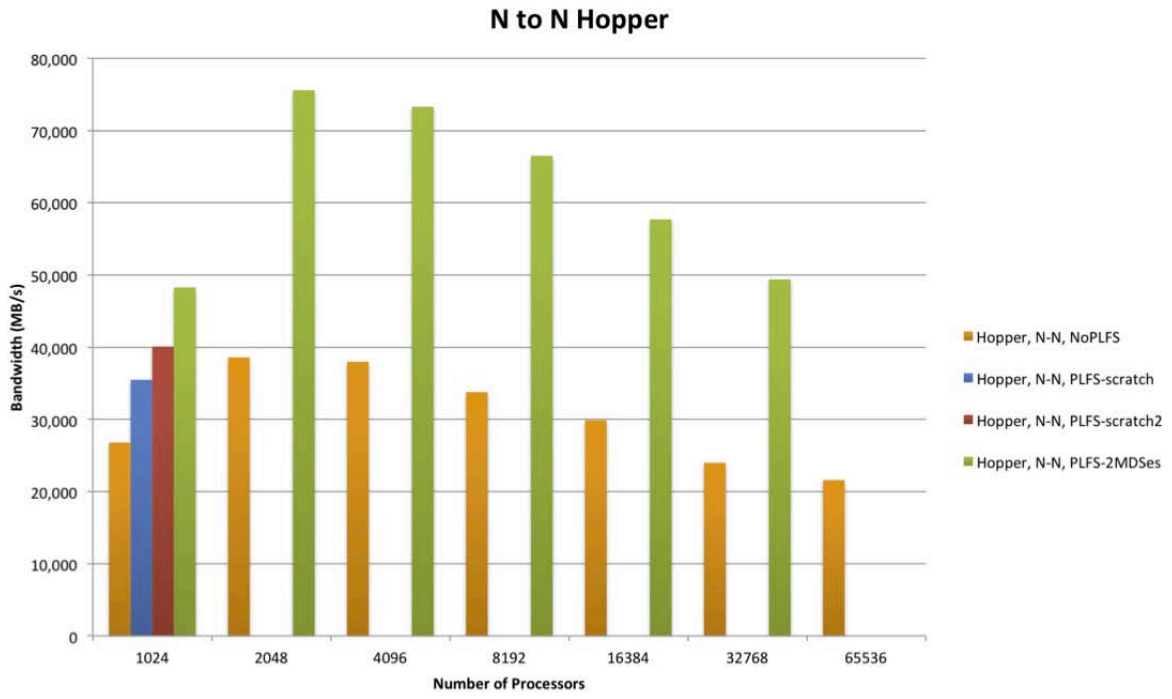

Quality of Service

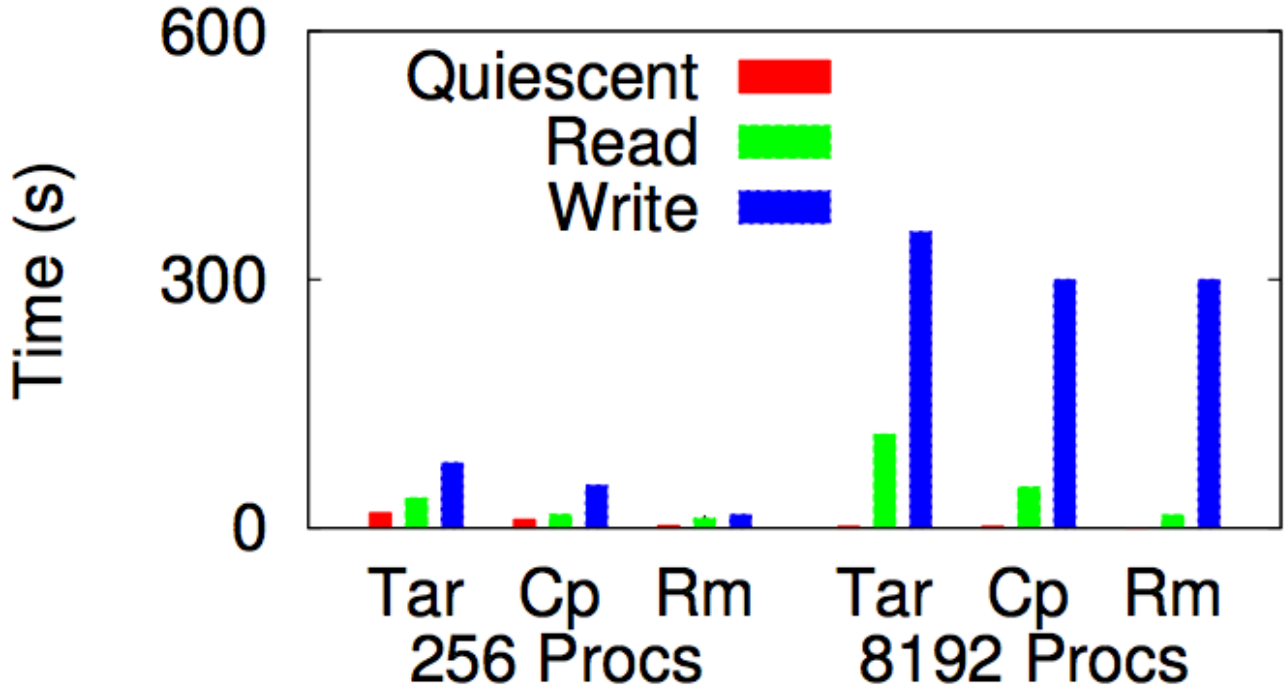

Questions?