

Bright Cluster & The Cluster I/O Challenge

“Storage has become a dominant component in many clusters. Initially, NFS was used to move data through a cluster. As clusters grew to hundreds and thousands of nodes, however, the demand for increased data I/O rates grew and NFS began to fall down. NFS was not designed to satisfy the appetite of HPC clusters” (“Welcome to the World of HPC Storage”, by Jeffrey B. Layton, Cluster World Magazine, Volume 2 Number 10).

A large number of companies and projects (including IBRIX, Polyserve, Panasas, Terrascale, Data Direct Networks, Lustre, PVFS, GFS, iSCSI and others) attempt to solve the cluster I/O challenge through a multitude of methods. The problem with these methods is that they are variously intrusive, including replacement of native file systems and/or extremely expensive - attributes cluster admins & architects very much want to avoid!

Third-Party File Systems Not the Answer

Bright Cluster solves the Cluster I/O challenge unobtrusively. Rather than replacing native file systems, Bright Cluster Pulse uses standard network protocols over Ethernet, operates solely in user space and is easily integrated with existing scheduler environments for automated workflow and just-in-time data delivery.

Pulse replaces NFS and data staging by transferring data directly from primary storage to the local node hard drives. Applications can then access the data as needed with certainty and speed from the local disc.

Making only a single unicast connection (rather than potentially thousands) to any file server in the network, Bright Cluster™ Pulse™ delivers simultaneous data transfer to thousands of nodes at the highest possible speed – the Data Absorption Rate of the cluster.

The Data Absorption Rate of a cluster is the speed that a node can accept data. Small files can be absorbed at full network speed because they fit in a node's memory buffer. Files or directories of files larger than the memory buffer though must begin writing to the local hard drive before all the data can be absorbed. As the write speed to a hard drive is much slower than the network transfer speed the result is a data jam that causes thrashing and packet loss in the data stream. Pulse regulates the data flow at the absorption threshold resulting in a smooth data stream that delivers the maximum possible throughput.

Expensive Storage Not The Answer

Solving the Cluster I/O problem through the use of high bandwidth storage and networks is yet another widely used solution to the problem. The obvious downside is that this approach becomes expensive fast!

Bright Cluster Pulse delivers optimum cluster I/O performance using existing storage and/or low cost file servers. There is no need to invest in high-end storage solutions because Pulse makes only a single connection to the data source and delivers data directly to the destination cluster. This eliminates server thrashing, replaces the tedious process of data staging, and results in faster, streamlined data flow using low cost storage filers.

For more information on Bright Cluster and Bright Cluster Pulse, visit www.brightcluster.com.

