

Grids and HPC: not as different as you might think?

I/O Forwarding in Grids

Jason Cope^{*1}, Kamil Iskra^{†1}, Dries Kimpe^{‡1,2}, and Robert Ross^{§1}

¹*Mathematics and Computer Science Division, Argonne National Laboratory*

²*Computation Institute, University of Chicago*

Abstract Traditionally there has been little interaction between the Grid and High-Performance Computing (HPC) storage research communities. Grid research often focused on optimizing data accesses for high-latency, wide-area networks while HPC research focused on optimizing data accesses for local, high-performance storage systems. Recent software and hardware trends are blurring the distinction between Grids and HPC. In this paper, we investigate the use of I/O forwarding – a well established technique in leadership-class HPC machines – in a Grid context. We show that the problems that triggered the introduction of I/O forwarding for HPC systems also apply to contemporary Grid computing environments. We present the design of our I/O forwarding infrastructure for Grid computing environments. Finally, we discuss the advantages our infrastructure provides for Grids, such as simplified application data management in heterogeneous computing environments and support for multiple application I/O interfaces.

Keywords ioforwarding, hpc, grid, io

1 Introduction

Grid computing environments, such as the National Science Foundation (NSF) funded TeraGrid project, have recently begun deploying massively-parallel computing platforms similar to those in traditional HPC centers. While these systems do not support distributed or multi-resource MPI applications [8, 3], they do support a variety of HPC applications well suited for tightly-coupled resources, including high-throughput workloads [19] and massively-parallel workloads [6]. To efficiently connect these resources, TeraGrid has focused on enhancing Grid data services. This trend is evident in the goals for the emerging third phase of TeraGrid operations, known as TeraGrid “eXtreme Digital” (XD).

This shift in resource usage and deployments aligns Grids more closely with traditional HPC data-centers, such as the DOE leadership computing facilities at Argonne National Laboratory and Oak Ridge National Laboratory. This realignment poses several data access challenges. One such challenge is enabling efficient, remote data

access by Grid applications using large numbers of processing elements. Massively-parallel applications can overwhelm file systems with large numbers of concurrent I/O requests. Leadership-class computing platforms face a similar data access problem for local data access to high-performance storage systems. Grid computing platforms experience similar problems for both local and remote data accesses. Existing Grid data management tools do not address the impact of increased concurrency on remote data access performance or account for the limited capacity of network and storage resources as application data continues to increase.

In this paper, we describe how I/O forwarding can improve the performance of Grid application data accesses to both local and remote storage systems. In the following sections, we present our I/O forwarding infrastructure for Grid computing environments and how this infrastructure optimizes application remote data accesses in Grids. Section 2 presents I/O forwarding and its use in HPC. Section 3 describes typical I/O mechanisms used by Grid applications and how I/O forwarding integrates into Grids. In Section 4, we describe related work and conclude this paper.

2 HPC I/O

In this section, we introduce the concept of I/O forwarding, followed by a description of our portable, open source implementation.

2.1 A Revised I/O Software Stack

The current generation of leadership-class HPC machines such as the IBM Blue Gene/P supercomputer at Argonne National Laboratory or the *Roadrunner* machine at the Los Alamos National Laboratory consist of a few hundred thousand processing elements. Future generations of supercomputers will incorporate millions of processing elements. This significant increase in scale is brought about by an addition in the number of nodes along with new multi-core architectures that can accommodate an increasing number of processing cores on a single chip.

While the computation power of supercomputers keeps increasing with every generation, the same is not true for

*Email: copej@mcs.anl.gov

†Email: iskra@mcs.anl.gov

‡Email: dries@uchicago.edu

§Email: rross@mcs.anl.gov

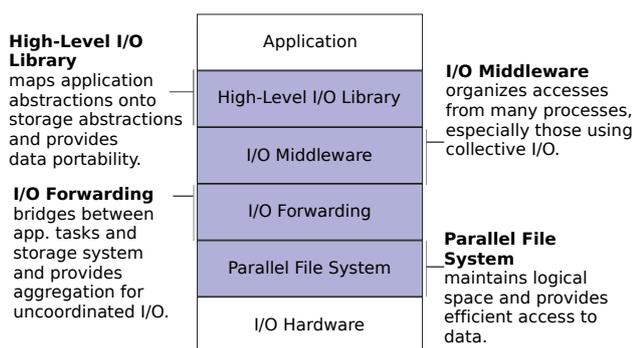


Figure 1: I/O Forwarding in HPC systems

their I/O subsystems. The data access rates of storage devices has not kept pace with the exponential growth in processing performance. In addition to the growing bandwidth gap, the increase in compute node concurrency has revealed another problem: the parallel file systems available on current leadership-class machines, such as PVFS2 [5], GPFS [14], Lustre [1] and PanFS [12] were designed for smaller systems with fewer file system clients. While some of these file systems incorporate features for enhanced scalability, they are often not prepared to deal with the enormous increase in clients brought on by the increasing trend towards more concurrency.

MPI-IO, distributed as part of the MPI library, is the standard parallel I/O API for HPC systems. In certain cases, by using *collective I/O*, the MPI-IO implementation is able to reduce the number of requests made to the filesystem. However, not all applications use the MPI-IO interface or are able to use collective I/O, so improvements made at the MPI-IO layer may not be available to the entire spectrum of scientific applications. Parallel high-level libraries such as Parallel-NetCDF [11] use MPI-IO and as such face many of the same limitations outlined above. POSIX implementations and serial high-level libraries are an artifact from an earlier generation and are only available on current HPC systems to support legacy applications.

To address this I/O bottleneck, another layer needed to be introduced into the I/O software stack. Clients, instead of directly making requests to the parallel filesystem, forward their I/O operations to an I/O forwarder node, which performs the I/O operation on their behalf. One I/O node is typically responsible for 32 to 128 compute clients. Due to its position in the I/O path, the I/O forwarder is able to perform a wide range of optimizations that were not previously possible. For example, it can aggregate requests of unrelated software running on multiple compute nodes. As such, it reduces both the number of requests and the number of clients visible to the parallel filesystem. Since the I/O forwarding software – running on the I/O node – does not share any resources (CPU or memory) with the compute clients, it is free to dedicate memory and compute power to optimizing I/O traffic without slowing down computation.

Another benefit of moving the actual I/O calls to the forwarder is that the compute client can be simplified. Instead of requiring a full I/O stack, it only needs to be able to send and receive requests to the I/O forwarder. The I/O forwarder then takes care of using the correct protocol to access the remote filesystem. Likewise, authentication (to the remote filesystem) can be handled by the I/O forwarder. This enables compute clients to use a simpler, local authentication scheme to authenticate to the I/O forwarder. Figure 1 shows the resulting I/O software stack.

2.2 I/O Forwarding Scalability Layer (IOFSL)

In view of the importance of I/O forwarding in HPC systems, it is desirable to have a high quality implementation capable of supporting multiple architectures, file systems and high-speed interconnects. While a few I/O forwarding solutions are available for the IBM Blue Gene and other leadership class platforms, such as the Cray XT, they are each tightly coupled to one architecture [20, 7]. The lack of an open-source, high-quality implementation capable of supporting multiple architectures, file systems and high-speed interconnects has hampered research and makes the deployment of novel I/O optimizations difficult.

To address this issue, we created a scalable, unified I/O forwarding framework for high-performance computing systems called the I/O Forwarding Scalability Layer (IOFSL) [2]. IOFSL includes features such as the coalescing of I/O calls on the I/O node, reducing the number of requests to the file system, and full MPI-IO integration, which translates into improved performance for MPI applications. Ongoing work includes the integration of some other techniques for improving HPC I/O performance, such as [21] and [13].

3 Grid Data Access

Two approaches to application data accesses in Grids have emerged. They are described in section 3.1. Section 3.2 describes how IOFSL can be used to improve the performance and enhance the useability of these approaches.

3.1 Traditional Grid I/O

The first approach stages data at the resource where the application executes or offloads data locally generated by an application to a remote storage system. This approach often uses GridFTP to perform bulk data transfers between the high-performance storage systems attached to Grid resources. While this approach offers good performance, as remote I/O is only used for staging files in and out the local storage, it has a number of drawbacks. For one, it is hard to maintain consistency between the local and remote copy. The second issue is related to the access granularity. Typically, the whole file needs to be transferred, reducing efficiency if the application only requires a subset of the file data.

The second approach is to host data on wide-area file systems. These file-systems construct a distributed,

shared storage space, which is mounted locally on each Grid resource to provide local application access. Examples of Grid specific filesystems include Gfarm [16] and Chirp [17]. These filesystems typically do not provide traditional I/O semantics and are currently not well supported by parallel applications. For example, in Gfarm, files are basically write-once and parallel read-write I/O has to be emulated through versioning and creating new files[15].

In addition to these Grid specific filesystems, traditional HPC filesystems such as Lustre and GPFS, have been adapted for Grid environments. While these do offer familiar parallel I/O facilities, the high latencies and large number of filesystem clients severely limits their performance and stability.

3.2 I/O Forwarding in a Grid environment

When designing IOFSL, portability and modularity was an important goal. IOFSL does not make any assumptions about operating system kernels, interconnects, filesystems or machine architectures. A such, it can be easily retargeted to other environments, such as computational Grids.

In large HPC systems, I/O forwarding isolates local compute clients, connected by a high bandwidth, low latency interconnect from the more distant, higher latency parallel filesystem. At the same time, it protects the filesystem from being crippled by a storm of requests, by aggregating and merging requests before sending them to the filesystem. From the viewpoint of the remote filesystem, this reduces the number of visible clients and requests, hence increasing performance.

In a Grid environment, these optimizations are also applicable, albeit on a different scale. While latencies might be much higher, the same discontinuity exists when an application, running on a local Grid resource needs to fetch data from a remote data store. As is the case in large HPC systems, a large number of simultaneous requests to a remote site might adversely affect the stability and throughput of the remote file server. This observation is valid both for data staging and wide area Grid filesystems.

Figure 2 shows the location of I/O forwarding in a Grid environment. Being located at the gateway between the local compute resources and the remote data, IOFSL acts as both a connection and request aggregator: local applications can share the same set of outgoing connections, increasing efficiency and reducing the load on the remote filesystem. For example, if GridFTP is used as a data transport between the site where data is stored and the site where data is consumed or generated, when using IOFSL, the number GridFTP connections will not depend on the number of clients. Instead, each I/O forwarder can be configured to use an optimal number of GridFTP connections to obtain the data. Clients interacting with the I/O forwarder will transparently share existing connections when possible.

Another important advantage of deploying I/O forwarding in a Grid environment is that, to the client software,

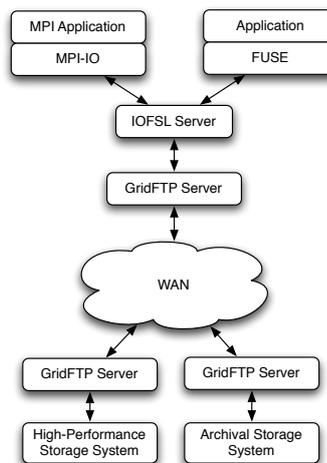


Figure 2: I/O Forwarding in a Grid environment

IOFSL can offer a more familiar access API. Currently, IOFSL implements 2 client side APIs: POSIX and MPI-IO. For POSIX, there is a FUSE and SYSIO implementation. The former enables redirecting I/O accesses of unmodified binary applications. While the latter requires relinking applications with the SYSIO library, it provides support on platforms that do not support FUSE (for example, minimal OS kernels such as Cray's Catamount kernel[9] or IBM BG/P's compute node kernel).

By directly supporting MPI-IO, the defacto I/O API for parallel MPI programs, IOFSL enables unmodified MPI applications (such as parallel analysis or visualization software) to transparently access remote data using GridFTP or other protocols not normally supported by HPC software. In this case, IOFSL effectively acts as a bridge between a local HPC programs and remote Grid-style storage.

Dedicating some nodes as I/O forwarders also helps with high latency network links, a typical problem when spanning multiple remote sites using a POSIX-like filesystem such as Lustre or GPFS. By using local system memory of the I/O forwarders for buffering read and write data, IOFSL is able to transform synchronous client accesses into asynchronous remote accesses, reducing the detrimental effects high latency links. Often requested data can be buffered locally, where it can be accessed over a low-latency high bandwidth network.

As IOFSL transparently captures all accesses, all the optimizations discussed above apply to unrelated programs; i.e. if two or more independent applications request the same data only one request needs to be made to the remote site.

4 Related Work and Conclusions

4.1 Related Work

In [4], a method is described to allow MPI-IO access to GridFTP stores. It differs from our work in that the MPI application itself makes the GridFTP connection, as opposed to the I/O forwarder node when IOFSL is used. This precludes optimizations such as request merging or link aggregation.

Stork [10] tries to improve I/O access time by explicitly scheduling data staging. While IOFSL will also buffer data using local temporary storage, it does this transparently – without explicit data staging – and on a sub-file granularity.

Condor [18] enables remote I/O by shipping I/O operations back to the submission site. It requires application to relink with the condor library. While our approach also uses function call forwarding, the calls are not shipped to the remote site but to local aggregators.

4.2 Conclusions

An overview was given of IOFSL, the I/O forwarding Scalability Layer. We describe the concept of I/O forwarding in HPC systems and show how the same technique can be applied in a Grid environment. We discuss its advantages and disadvantages, and show how it enables connecting existing HPC and posix applications with Grid data stores.

References

- [1] Lustre: A scalable, high-performance file system. Cluster File Systems Inc. white paper, Nov. 2002.
- [2] N. Ali, P. Carns, K. Iskra, D. Kimpe, S. Lang, R. Latham, R. Ross, L. Ward, and P. Sadayappan. Scalable I/O Forwarding Framework for High-Performance Computing Systems. In *IEEE Int'l Conference on Cluster Computing (Cluster 2009)*, September 2009.
- [3] G. Allen, T. Damlitsch, I. Foster, N. Karonis, M. Ripeanu, E. Seidel, and B. Toonen. Supporting efficient execution in heterogeneous distributed computing environments with cactus and globus. In *Proceedings of SC 2001*, November 10-16 2001.
- [4] T. Baer and P. Wyckoff. A parallel I/O mechanism for distributed systems. In *cluster*, pages 63–69. IEEE, 2004.
- [5] P. H. Carns, W. B. Ligon III, R. B. Ross, and R. Thakur. PVFS: A parallel file system for Linux clusters. In *Proceedings of the 4th Annual Linux Showcase and Conference*, pages 317–327, 2000.
- [6] L. Grinberg and G. Karniadakis. A scalable domain decomposition method for ultra-parallel arterial flow simulations. *Communications in Computational Physics*, 4(5):1151–1169, 2008.
- [7] K. Iskra, J. W. Romein, K. Yoshii, and P. Beckman. ZOID: I/O-forwarding infrastructure for petascale architectures. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 153–162, Salt Lake City, UT, 2008.
- [8] N. Karonis, B. Toonen, and I. Foster. Mpich-g2: A grid-enabled implementation of the message passing interface. *Journal of Parallel and Distributed Computing*, 2003.
- [9] S. Kelly and R. Brightwell. Software architecture of the light weight kernel, Catamount. In *Proceedings of the 2005 Cray User Group Annual Technical Conference*, 2005.
- [10] T. Kosar and M. Balman. A new paradigm: Data-aware scheduling in grid computing. *Future Generation Computer Systems*, 25(4):406–413, 2009.
- [11] J. Li, W. Liao, A. Choudhary, R. Ross, R. Thakur, W. Gropp, R. Latham, A. Siegel, B. Gallagher, and M. Zingale. Parallel netCDF: A high-performance scientific I/O interface. In *ACM/IEEE Conference on Supercomputing*, Phoenix, AZ, Nov. 2003.
- [12] D. Nagle, D. Serenyi, and A. Matthews. The Panasas ActiveScale storage cluster—delivering scalable high bandwidth storage. In *ACM/IEEE Conference on Supercomputing*, Pittsburgh, PA, Nov. 2004.
- [13] P. Nowoczynski, N. Stone, J. Yanovich, and J. Sommerfield. Zest: Checkpoint storage system for large supercomputers. In *3rd Petascale Data Storage Workshop Supercomputing*, pages 1–5, November 2008.
- [14] F. Schmuck and R. Haskin. GPFS: A shared-disk file system for large computing clusters. In *USENIX Conference on File and Storage Technologies*, Monterey, CA, 2002.
- [15] O. Tatebe, Y. Morita, S. Matsuoka, N. Soda, and S. Sekiguchi. Grid datafarm architecture for petascale data intensive computing. In *ccgrid*, page 102. Published by the IEEE Computer Society, 2002.
- [16] O. Tatebe, N. Soda, Y. Morita, S. Matsuoka, and S. Sekiguchi. Gfarm v2: A Grid file system that supports high-performance distributed and parallel data computing. In *Proceedings of the Computing in High Energy and Nuclear Physics Conference (CHEP'04)*, 2004.
- [17] D. Thain, C. Moretti, and J. Hemmes. Chirp: a practical global filesystem for cluster and Grid computing. *Journal of Grid Computing*, 7(1):51–72, 2009.
- [18] D. Thain, T. Tannenbaum, and M. Livny. Distributed computing in practice: The Condor experience. *Concurrency and Computation Practice and Experience*, 17(2-4):323–356, 2005.
- [19] M. Wilde, I. Ioan Raicu, A. Espinosa, Z. Zhang, B. Clifford, M. Hategan, S. Kenny, K. Iskra, P. Beckman, and I. Foster. Extreme-scale scripting: Opportunities for large task-parallel applications on petascale computers. *Journal of Physics: Conference Series*, 180(1), 2009.
- [20] H. Yu, R. K. Sahoo, C. Howson, G. Almasi, J. G. Castanos, M. Gupta, J. E. Moreira, J. J. Parker, T. E. Engelsiepen, R. Ross, R. Thakur, R. Latham, and W. D. Gropp. High performance file I/O for the Blue Gene/L supercomputer. In *International Symposium on High-Performance Computer Architecture*, Feb. 2006.
- [21] F. Zheng, M. Abbasi, C. Docan, J. Lofstead, Q. Liu, S. Klasky, M. Prashar, N. Podhorszki, K. Schwan, and M. Wolf. Predata - preparatory data analytics on petascale machines. In *Proceedings of 24th IEEE International Parallel and Distributed Processing Symposium*, 2010.