CEDPS: The Center for Enabling Distributed Petascale Science

A. Baranovski, S. Bharathi, J. Bresnahan, A. Chervenak, I. Foster (PI)*, D. Fraser, T. Freeman, D. Gunter, K. Jackson, K. Keahey, C. Kesselman, D. E. Konerding, N. Leroy, M. Link, M. Livny, N. Miller, R. Miller, G. Oleynik, L. Pearlman, J. M. Schopf, R. Schuler, B. Tierney

Summary

Petascale science is an end-to-end endeavor, involving not only the creation of massive datasets at supercomputers or experimental facilities, but the subsequent analysis of that data by a user community that may be distributed across many laboratories and universities. The SciDAC Center for Enabling Distributed Petascale Science (CEDPS) is developing tools to support this end-to-end process. These tools include services for the reliable, high-performance, secure, and policy-driven placement of data within a distributed science environment; tools and techniques for the construction, operation, and provisioning of scalable science services; and tools for the detection and diagnosis of failures in end-to-end distributed applications. In each area, we build on a strong base of existing technology and have already made useful progress. We are working closely with several groups already, and look forward to working with additional SciDAC application and technology projects to realize the promise of petascale science.

Introduction

Many SciDAC applications require interactions with large data sources, either to transfer data from the creation site to somewhere accessible to end users and collaborators or to move analysis codes to be closer to large data sets themselves. It is critical to the success of these applications to address the issues of smart data transport, the ability to do analysis at the data source site instead of remotely, and to address both these concerns in a reliable manner.

CEDPS has defined a pragmatic program of interaction with end users, prototyping, deployment, refinement, and expanded use within the SciDAC community to addresses these concerns. Initial work with the Earth System Grid (ESG) and Open Science Grid (OSG) shows that CEPDS is addressing critical concerns for these SciDAC projects, and is enabling additional functionality and reliability needed to produce new scientific results.

Technology Impact

In *data placement*, we are developing tools and techniques for reliable, high performance, secure, and policy driven placement of data within a distributed science environment. We are constructing a managed object placement service (MOPS) that represents a significant enhancement to today's GridFTP and provides management of space, bandwidth, memory usage, number of connections, and other resources needed to transfer data more efficiently among storage systems. Building on this base, we are developing end-to-end data placement

*foster@mcs.anl.gov

services that implement different data distribution and replication behaviors.

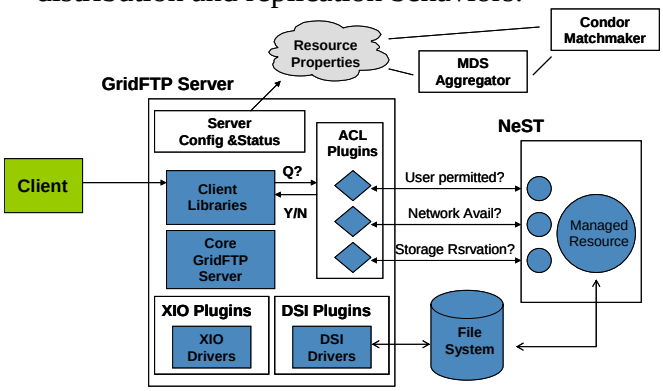

Figure 1: The MOPS architecture.

In *scalable services*, we are developing tools to simplify the process of Grid-enabling applications and dynamically provisioning resources for their execution. We are creating tools that automatically enable applications to act as remotely accessible services by adding appropriate interfaces, authorization, persistence, and other capabilities. Services constructed in this way allow the end-user to take advantage of sophisticated and complex Grid resource management machinery that can provision resources according to time-varying application needs, without having to be exposed to the details by leveraging virtualization technology.

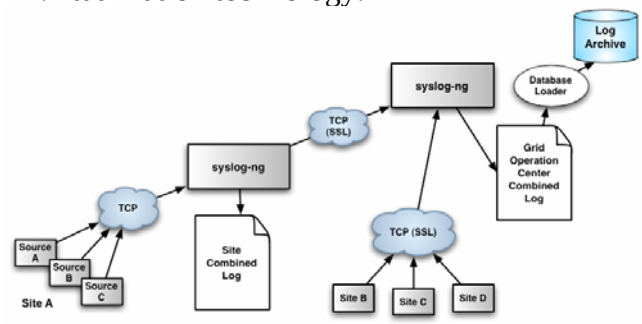

Figure 2: Syslog-ng deployment on OSG.

Finally, in *troubleshooting*, we are developing tools for the detection and diagnosis of failures in end-to-end data placement and distributed application hosting configurations. We are constructing an end-to-end monitoring architecture that uses instrumented services to provide detailed data for both background collection and run-time event-driven collection. We are also constructing new monitoring analysis tools able to detect failures and performance anomalies and predict system behaviors using archived data and event logs.

Impact in the SciDAC community

For over 10 years, members of the DOE SciDAC community have been making use of some of the basic software components that form the basis for the CEDPS tools, including GridFTP, the Replica Location Service (RLS), the Monitoring and Discovery Service (MDS), and the Workspace Service, developed at Argonne and USC/ISI, and NetLogger and pyGlobus, developed at LBNL, all part of Globus; dCache, developed at FNAL; and NeST, part of Condor.

Our ongoing work with application teams is allowing us to extend the functionality and usability of our tools to meet specific end-to-end needs of SciDAC applications. For example, over 12,000 GridFTP servers were used during calendar year 2006, establishing a large user base who will then consider deploying the expanded functionality available through MOPS. Likewise, OSG is currently experiencing a greater than 15% failure rate for job submissions and file transfers, a problem we are directly addressing through our troubleshooting work. The STAR community, among others, is using our services software to ensure consistency of the provisioned resources in their production environment.

For further information on this subject contact:

Ian Foster, Argonne National Laboratory, foster@mcs.anl.gov, (630) 252-4619
Jennifer M. Schopf, Argonne National Laboratory, jms@mcs.anl.gov, (630)252-3313