BERKELEY LAB

Distributed Systems Department

Bringing the Grid to Life

Computational Research Division

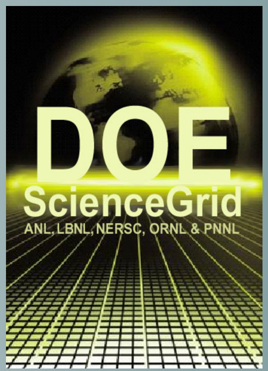

www.doesciencegrid.org

DOE Science Grid

The vision for "Grids" is to revolutionize the use of computing in science by making the construction and use of large-scale systems of diverse resources as easy as using today's desktop environments.

The DOE Science Grid is contributing to this vision by providing advanced distributed computing infrastructure on a scale that supports the level of scientific computing necessary for DOE to accomplish its scientific missions. The Science Grid is a collaboration between Lawrence Berkeley National Laboratory, Argonne National Laboratory, Oak Ridge National Laboratory, Pacific Northwest National Laboratory and DOE's National Energy Research Scientific Computing Center.

To accomplish this goal, the Science Grid is enabling access to a variety of computing and storage resources through standard Grid middleware tools. It is also providing a testbed for developing and deploying new middleware and Grid services where they can be tested and utilized by application scientists.

SciDAC
Scientific Discovery through Advanced Computing

Developing and effectively using application programs in distributed environments such as collaboratories or computational Grids presents a number of challenges. The Distributed Systems Department at the U.S. Department of Energy's Lawrence Berkeley National Laboratory has significant experience enabling applications to function efficiently in such environments, and is applying that experience to Computing and Data Grids.

We are working on several projects and have developed a variety of tools to help support scientists addressing complex and large-scale computing and data analysis problems in distributed environments, such as the DOE's Science Grid and Scientific Discovery through Advanced Computing (SciDAC) collaboratories.

Pervasive Collaborative Computing Environment

The Pervasive Collaborative Computing Environment (PCCE) provides a persistent and seamless environment that supports the entire continuum of interactions between collaborators. The goal of the PCCE is to leverage the Grid middleware environment and develop tools that support the day-to-day connectivity and underlying needs of a group of collaborators. It will allow participants to securely locate each other; use asynchronous and synchronous messaging; share documents, progress and applications; and hold videoconferences. Workflow tools will enable coordination of Grid computing processes and human tasks. These PCCE tools are helping to provide lightweight, non-intrusive, and flexible ways for collaborators to stay in touch and work together.

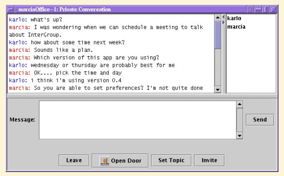

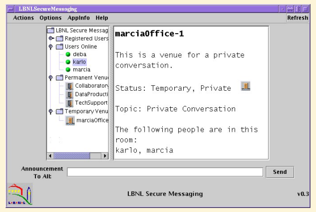

Reliable and Secure Group Communication

Many distributed applications require reliable, ordered delivery of messages and membership services for a group of processes. However, IP multicast provides only unreliable, unordered delivery of messages, and TCP only provides reliable delivery between two processes.

The goal of this project is to develop the components needed for a peer-to-peer group communication infrastructure that provides reliability, security, and fault-tolerance while enabling scalability on the Internet scale. In this project we are developing protocols called InterGroup to provide reliable multicast capabilities and Secure Group Layer to provide security. This group communication infrastructure will complement and enhance the existing communication capabilities of the Grid. Internet scalability is achieved in the protocol by keeping track of the currently active senders in a group and by providing flexible message delivery options.

Scalable and Secure Peer-to-Peer Information Sharing Tool

Groups collaborating on scientific experiments have a need to share information and data among the collaborators. In a typical scientific collaboration, there are many different locations where data would naturally be stored. These storage locations are usually determined by factors such as where the data is generated, who generated the data, where it is maintained, where it is analyzed, and/or the amount of storage space available at a location. These factors are different for each piece of data and often result in the data files being stored in a wide variety of locations, making it difficult for collaborators to find what they need.

This project is developing a lightweight, scalable and secure file-sharing system that will enable scientists to store and manage their files on local storage facilities while sharing them with remote participants. A core component of this system will be the authorization mechanisms. The underlying infrastructure that will enable this is the secure and reliable group communication system.

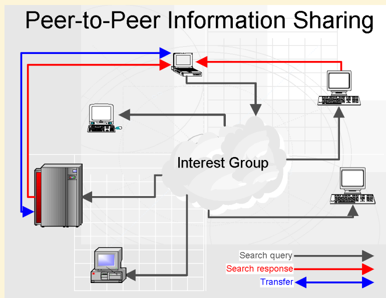

Distributed Monitoring Framework

The goal of the Distributed Monitoring Framework (DMF) is to improve end-to-end data throughput of data-intensive applications in high-speed wide-area network environments, and to provide the tools for performance analysis and fault detection in a Grid computing environment. This monitoring framework will provide accurate, detailed and adaptive monitoring of all distributed computing components, including the network. Analysis tools will be able to use this monitoring data for real-time analysis, anomaly identification and response.

The main components of the DMF are instrumentation, sensors, sensor management, event publication and event archiving. DMF includes tools to make it easy to non-intrusively add instrumentation to Grid middleware and to publish this event data in a standard manner. Network and host sensors, combined with instrumented applications, allow end-to-end performance analysis. A Sensor Management System allows secure control of the distribution and execution of monitoring sensors. The event archive will be extremely high performance and scalable to ensure that it does not become a bottleneck.

Net100

High-performance research networks have reached speeds of tens of gigabits per second, but real-world transfers on these networks achieve barely a fraction of the network's potential. Net100 is developing the tools to allow scientists and researchers to achieve the highest possible data transfer rates without the assistance of networking experts. Net100 is a collaboration between Pittsburgh Supercomputing Center, Oak Ridge National Laboratory and Lawrence Berkeley National Laboratory. Built on Web100 (PSC, NCAR, NCSA) and NetLogger (LBNL), Net100 modifies operating systems to respond dynamically to network conditions and make adjustments in network transfers, sending data as fast as the network will allow. The LBNL component is called the Network Tools Analysis Framework (NTAF). NTAF is a network testing environment used to define network measurement tools, network experiments, how often to run the experiments and where to send the results. NTAF runs the tests on a mesh of test hosts and sends the results to an archive service for historical analysis.

Self-Configuring Network Monitor

Comprehensive end-to-end and top-to-bottom monitoring is critical for developing and debugging high performance Grid applications. Without information about a stream from intermediate hops within the network, the end-to-end system is often unable to identify and diagnose problems within the network. This service is largely unavailable to the application developer except in testbed environments. Increasingly the approach of these applications is to rely on automatic tuning of transport parameters such as TCP window size, parallel streams, etc.

In this project we are designing and implementing a self-configuring, passive network monitoring system that uses special request packets to automatically activate monitoring along the network path between communicating endpoints. The ultimate goal of this monitor is to provide on-demand, application-to-application monitoring capabilities throughout the interior of the interconnecting network domains. A prototype of the monitor has been developed and deployed and is accepting requests for monitoring on the DMZs at several sites. Applications wishing to make use of the monitors run a GUI to request monitoring of their traffic.

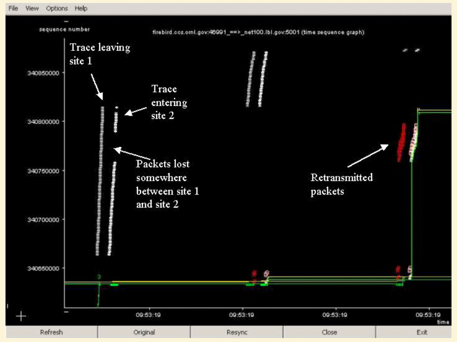

Distributed Security Architectures

Strong but easy-to-use access control of Grid resources and services is essential to scientific collaboration on the Grid. Akenti is a certificate-based authorization service designed to control access in a distributed environment where the resources, users and stakeholders are dispersed geographically and administratively. Users are identified by X.509 identity certificates or Globus proxies passed through SSL/TLS connections. Access policy is specified in a number of certificates, which may be signed by multiple stakeholders and may be stored remotely from the resource.

Akenti is called by a resource gateway via a simple interface. Given the user's identity credential and the name of the resource for which access is requested, the Akenti server locates all the relevant policy statements, verifies them and returns the allowed actions to the server. The Akenti system includes support for stakeholders to create and sign policy statements and to verify user access to their resources. The Akenti server also provides auditing of all accesses to resources.

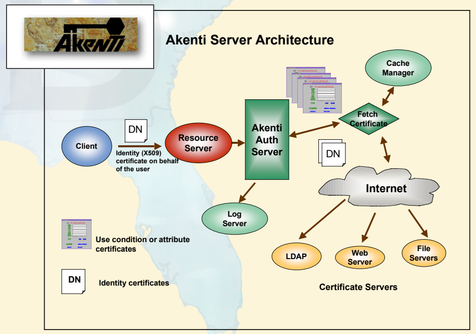

Python CoG Kit

The Python CoG Kit provides a mapping between Python and the Globus Toolkit. It extends the use of Globus by enabling access to advanced Python features such as events and objects for Grid programming. It hides much of the complexity of Grid programming behind simple object-oriented interfaces and supports rapid development of high-level Grid services, while providing access to the underlying performance of the Globus Toolkit.

The goal of this project is to develop a common set of reusable components in Python for accessing Grid services, with a focus on supporting the development of science portals, problem solving environments, and science applications that access Grid resources. This toolkit is currently being widely used in the Grid environment.

Grid Services

This project is developing a unifying standards-based framework for Grid services. It represents the next paradigm shift and will lead to the development of next-generation Grid capabilities. This framework will provide a set of easy to use Grid-oriented services and will enable the Grid and application communities to develop their own services with clearly defined portability and deployment constraints.

This project will use Web services in Grid computing as a unifying, standards-based framework for service description, discovery, and invocation. We will also use technologies developed within the Grid community (particularly within the Globus Toolkit) to provide added value to Web services, by addressing federation, security, reliable invocation, and lifetime/state management issues not supported by current Web services technologies. Our goal in this project is to design, develop, and apply the Grid services framework in support of DOE science applications.

For more information visit http://www.itg.lbl.gov