Scientific Data Management Integrated Software Infrastructure Center

R. D. Burris*, D. L. Million (ORNL)
M. K. Gleicher (Gleicher Enterprises)

Summary

This document describes work done by listed ORNL staff to support and enhance the activities of the SciDAC Scientific Data Management (SDM) Integrated Software Infrastructure Center (ISIC). The focus of this work has been to acquire, implement, enhance and support equipment and software which together present a "place to be" for other SDM researchers – a place in which they can develop their science, test their software and verify the satisfactory operation of that software in the infrastructure of the installation. In addition to these activities, development activities to improve access to storage researchers – primarily improvements to the "hsi" application, have also been performed.

"A Place To Be"

The SDM ISIC involves research by staff at eight installations, often collaborating on individual pieces of the overall ISIC product. Rather than require each site to establish a testbed, and each participant to do development at one or more such testbeds, the ISIC decided to use the resources of the ORNL Probe installation as its "place to be" and to designate it as the ISIC's "facility".

One major function of the ISIC facility is support for data reduction and analysis research. That activity makes heavy use of a four-processor IBM RS/6000, a four-processor eight-gigabyte Linux machine and a four-node Linux cluster. Oracle, R, Ggobi and other software products have been provided to support their activities.

Another major activity of the ISIC concentrates on the Parallel Virtual File System (PVFS). To augment the value of that product, PVFS has been installed on the cluster mentioned above and on a second four-node cluster being used to develop a mechanism connecting PVFS to the High Performance Storage System (HPSS), work being carried out by ORNL staff.

The SDM ISIC identified integration with the Grid as one of its initial goals. Two nodes in the facility are being used to support this work – an IBM RS/6000 and a Sun E250.

Improving Access to Massive Storage

Another major activity of the SDM ISIC is improving access to mass storage systems. This goal is being addressed by improvements to two applications – Hierarchical Storage Interface (HSI) and Hierarchical Resource Management (HRM).

The HSI application provides access to HPSS; it is in use at most HPSS installations and is both the most flexible and fastest data transfer mechanism available to HPSS users. In its typical use, an entire file is processed – either being stored into or retrieved from HPSS. However, in the case of very large files, it is frequently true that only portions of the file are needed. HSI has consequently been enhanced to make it possible to retrieve one or many portions of an

* 865-574-0972, rdb@ornl.gov

individual file, thereby improving responsiveness and reducing waste. It has also been ported to the Cray X1, thereby providing access to HPSS to users of that machine. Reliability has been improved by better retry capability and aggregate performance has been improved by supporting multiple concurrent transfers (making better use of network bandwidth). Finally, HSI now allows users to be authenticated by Grid certificates.

Work on the HRM application is being done by NERSC staff. Resources in the facility, including one of the Sun nodes and associated disk and tape capacity, host HRM. The application has recently been upgraded and is being used in production transfers of Climate T170 results to the National Center for Atmospheric Research.

Two additional projects, each associated with HRM, will also be sited in the SDM facility. The first will transfer roughly a terabyte of data from the STAR application at Brookhaven to further validate ISIC-supported bit-mapped index software. The second will use the facility to test and validate cache replacement algorithms under development in the ISIC.

Finally, HRM will be implemented on a node that supports the network tuning software being developed in the Net100 and Web100 projects.

Support for another SciDAC Application, the Terascale Supernova Initiative (TSI), has been a focus of various efforts at ORNL. Early steps toward providing the type of network capability needed by TSI have been undertaken in the SDM facility, namely the establishment of a Logistical Networking presence. Logistical Networking, a research activity of the University of Tennessee's Logistical Computing and Internetworking Laboratory, will provide to TSI the ability to store raw or analyzed data in network-attached resources and to be able to provide those data to its member institutions more rapidly and efficiently than current file-transfer mechanisms. Within the facility, one node is being used for development and test and a second node, provided by the University, will participate in a dedicated Logistical Network with nodes at several other TSI sites around the country. Tighter integration of Logistical Networking, Grid applications and HRM is being considered; if pursued, the work will take place in the SDM facility.

Integration

One of the key elements of the SDM ISIC will be integration of its products into an easy-to-use whole. Pursuant to that goal, the SDM facility will be used to prototype that capability. Specific actions will include supporting two software stacks – one which will be used by applications on a Linux cluster and one on IBM RS/6000 systems. Each stack will support HDF5 and Parallel NetCDF data formats and libraries and provide access to HPSS.

The various projects described herein all support the SciDAC activities of MICS. Two SciDAC applications – Climate and TSI – provide focus to research into data analysis, data reduction, improved access to mass-storage systems and effective network utilization. In addition, other MICS-funded activities including Logistical Networking, Net100 and HPSS development are utilized and enhanced by the work being done in the facility.

For further information on this subject contact:
John R. van Rosendale
Program Manager
Phone: (301)903-3127
E-mail: johnvr@er.doe.gov