

The Earth System Grid: Turning Climate Model Datasets into Community Resources

David E. Bernholdt*, Kasidit Chanchio, Mei Li Chen, Line Pouchard, Oak Ridge National Laboratory; Ian T. Foster, Veronika Nefedova, Argonne National Laboratory; Alex Sim, Arie Shoshani, Lawrence Berkeley National Laboratory; Bob Drach, Dean N. Williams, Lawrence Livermore National Laboratory; David Brown, Luca Cinquini, Peter Fox, Jose Garcia, Don E. Middleton, Gary Strand, National Center for Atmospheric Research; Shishir S. Bharathi, Ann Chervenak, Carl Kesselman; University of Southern California;

Summary

Global climate simulations generate large volumes of data (1-100 TB per run) that is of interest to a variety of researchers in different locations, often distant from the supercomputer on which the simulation is actually run. The goal of the Earth System Grid effort is to develop an environment that enables management, discovery, and access to these widely distributed data resources by a widely distributed user community. As part of the ESG team, ORNL researchers are focusing on issues of metadata and ontologies, to support the management of climate data from diverse sources, and on monitoring the distributed system for faults and resource utilization.

The goal of the Earth System Grid (ESG, <http://www.earthsystemgrid.org>) project is to develop an environment that enables the management, discovery, distributed access, processing, and analysis of distributed terascale climate research data.

Global climate simulations, carried out at a number of supercomputer sites around the country, produce datasets, typically composed of thousands of files totaling roughly 1-100 TB or more. These datasets are of interest to numerous climate researchers for different purposes, and must also be archived. Tracking these datasets and moving them to where they are needed for archiving or analysis are, at present, tedious and time-consuming tasks.

The Earth System Grid project brings together remote access protocols from the environmental science community, grid-based technologies for authentication, data discovery, and resource access, tools for storage management, and new tools, to create a dis-

tributed environment allowing researchers easy discovery and use of new and existing climate datasets. The result is a large, distributed software system that links most of the major US climate computational and archive sites with tools for high throughput data movement, metadata cataloging and searching.

Recent work at ORNL has focused on two areas: the importance of metadata and ontologies, and the development and deployment of monitoring capabilities for the complex software system that comprises the ESG.

Metadata and Ontologies

Descriptive information about the datasets registered with the ESG (their *metadata*) is of crucial importance in allowing researchers to locate datasets of interest to them. However the specific details of what is contained in the metadata will depend on the climate modeling software that produced it and other factors. In an environment where

* (865) 574-3147, bernholdtde@ornl.gov

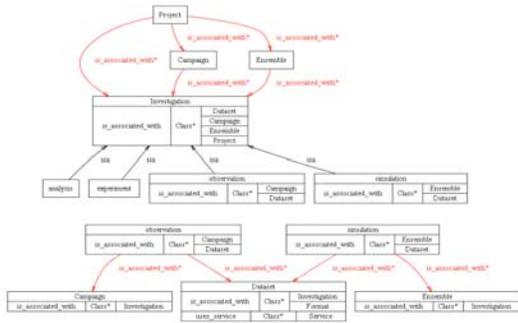


Figure 1. Two ontological entity-relationship diagrams classifying scientific products from ESG-acquired data.

different types of datasets are used, or in the federation of multiple data archives, *ontologies* become extremely important to allow interconversion of metadata and searches across different types of metadata. Whereas a traditional approach would require N^2 distinct one-to-one mappings between different metadata *schemas*, it is possible to define a mapping between each metadata schema and the unifying ontology, reducing the effort to just N mappings.

Ontological tools have been used at every stage of development of the metadata schemas currently in use in the ESG to check their formal consistency and correctness. We plan to extend this work to allow the federation of ESG data with that held by related efforts, such as the National Environment Research Council (NERC) data grid project in the UK.

Monitoring the Earth System Grid

In a distributed computing environment such as the ESG, failures, scheduled down time of components, and resource saturation, may significantly affect the ability of the ESG system to respond to a given request. Clearly it is important to provide both the users, and the system software with status and load information about the various components, but with a software system that is so large and complex, it is also important that the information be focused, to avoid overwhelming the user.

ORNL researchers are deploying prototype tools for fault monitoring and resource utilization tracking to ESG nodes in order to better frame the particular requirements posed by a distributed data management system like the ESG. The fault monitoring software is based on a prototype developed by the DOE Science Grid SciDAC project, but with the capability to define “composite” services, which are defined hierarchically, to facilitate focusing the information for the user. It is also important to be able to monitor utilization of resources which may be the subject of contention by multiple users, such as wide-area network bandwidth, and access to storage archives (with a fixed number of tape drives) and potentially CPU and disk resources on front-end systems for these storage archives. For this reason, we are in the initial stages of deployment of the Resource Usage Data and Accounting (RUDA) software, also developed by the DOE Science Grid project, to monitor and track resource utilization in the ESG.

SERVICE	NAME	STATUS	DESCRIPTION
All ESG Services	esg.ncar.*		All components at NCAR
Data Transport	esg.ncar.hrm	Up	HRM's components, gridFTP
Replica Location Service	esg.ncar.rls	Down	rls server, mysql
openDAPg Server	esg.ncar.openDAPg	Up	dods server, mysql

Machine Name	OS	#CPUs	MemTot (MB)	freeSwap (KB)	freeMem (KB)	IdleCPU
CAS-enabl	Sleepy	Solaris	2	512	542984	246240 99

PROJECT	USR	CPU Time (sec)	Wall Time (sec)	Mem/HWM (KB)	chrgUnit
science grid		9935	43634558	181896	120627028
	chertni	9047	9679886	109552	120415818
	kaandit	888	33954672	72344	211210
system		3	6982331	14168	315694
	postfix	3	3521716	12360	313886
	daemon	0	3460615	1808	1808
ESG		100079	52843299	464512	10770336
	astrg	4846	47217370	233408	415246
	grm	95233	5625929	231104	10355090

Figure 2. Screen shots of the ESG resource monitoring (foreground) and fault monitoring (background) tools.

For further information on this subject contact:

Mary Ann Scott, Program Manager
 Mathematical, Information, and Computational
 Sciences Division
 Office of Advanced Scientific Computing Research
 Phone: 301-903-6368
 scott@er.doe.gov