Mirroring Earth System Grid Data Sets

Erin Brady
University of Richmond

Ann Chervenak
USC Information Sciences Institute

The Earth System Grid (ESG) [1-3] provides secure access to climate simulation data located at storage sites across the U. S. ESG provides a collaborative environment that allows researchers to search for, download, and publish climate model datasets. ESG data sets currently consist of over 230 Terabytes stored in millions of files. Until recently, replication of ESG data sets was considered impractical because of their large size. With the growing importance of the Intergovernmental Panel on Climate Change (IPCC) data sets hosted at the ESG site at Lawrence Livermore National Laboratory, several sites around the world have expressed interest in hosting a replica or mirror of a subset of IPCC data. The goal of these mirror sites is to provide reliable access to these datasets for local scientists and to reduce wide area data access latencies. Our project provides the design and implementation of this data mirroring functionality, which integrates several components of the federated ESG architecture currently under development.

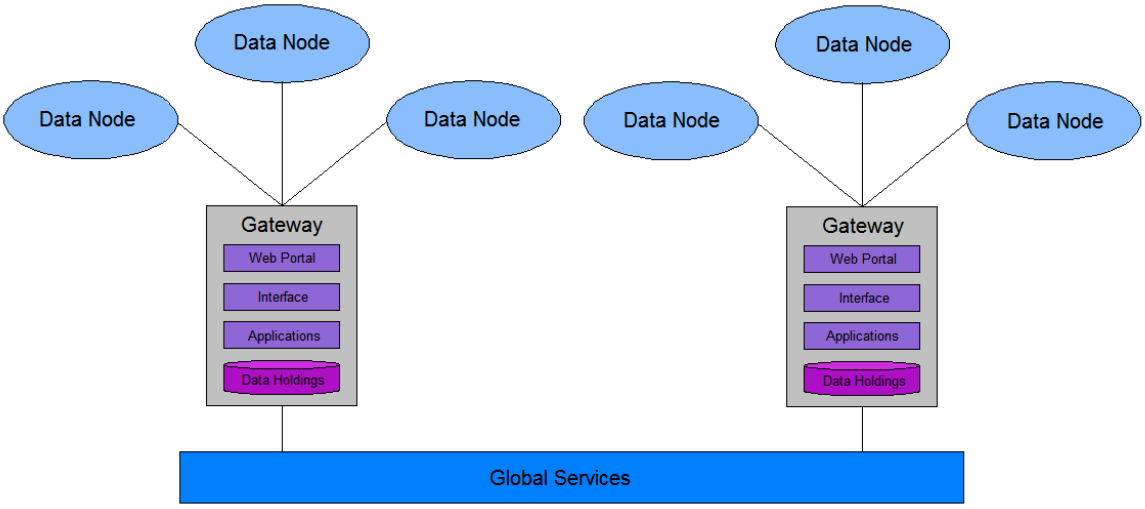

Figure 1. Federated ESG Architecture

Figure 1 illustrates the federated ESG architecture, which includes Data Nodes and Gateways. Data Nodes host data archives, where datasets are physically stored either on disk or a mass storage system. Nodes provide publication and retrieval services for data, allowing data to be served to users at different sites. Each Node has a parent Gateway, which provides the user with metadata search and data access functionality. A Gateway can be connected to several Nodes and contains detailed metadata about the datasets published by those Nodes. In addition, each Gateway provides access to search level metadata about every dataset in the distributed system, so that users can locate data from any Gateway in ESG.

The design of our data mirror tool is illustrated in Figure 2. To mirror an ESG dataset, our tool must locate the files to be replicated, transfer the dataset's files from the original Data Node to the mirror site, and publish the mirrored dataset and its associated metadata at the mirror site's Gateway. These three tasks are performed using services or clients being developed for the current ESG implementation. To locate a dataset that will be replicated, our mirror tool uses an ESG metadata client API developed at LLNL to query the metadata catalog at the Gateway where the dataset was originally published. This query returns an XML file containing metadata information for the dataset, including the location and size of the files in the dataset and the name of the climate model that produced the dataset. The mirror tool parses the physical addresses of the files from this metadata and creates a transfer request for a data movement client, with the parsed addresses as the source of the transfers and the mirror site location as the destination.

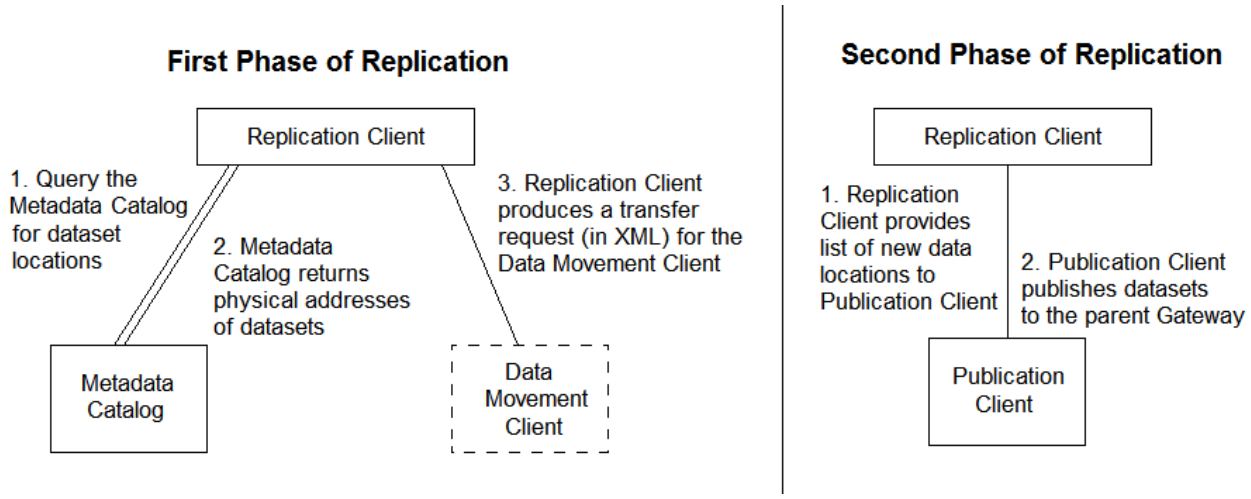

Figure 2. Data Mirror Tool Design and Interface

We issue transfer requests to a bulk data movement client developed at Lawrence Berkeley National Laboratory. Initially, we are using the Data Mover Lite client; eventually, we will use the Bulk Data Movement Client being developed for ESG. Each transfer request contains the physical locations and sizes of the files to be transferred and may also contain checksums that can be used by the transfer client to ensure that the files have been transferred accurately.

Once the dataset has been successfully transferred to the mirror site, it must be published at the mirror site's Gateway to make the dataset discoverable and available for access. The mirror tool calls a publication client developed at LLNL. This client scans the dataset to extract its metadata and publishes it in the metadata catalog at the mirror site's Gateway.

In summary, the goal of the ESG data mirror tool is to integrate existing ESG functionality for metadata query, data transfer and publication to create mirror sites for fast and reliable access to ESG data sets by a growing number of users around the world.

Currently, the mirror tool can complete the first stage of replication successfully. For each dataset the user wishes to replicate, our tool queries the Gateway for the physical addresses of all the dataset's files. The tool produces transfer requests correctly formatted for the DML client. We plan to implement the second part of the replication client, which calls the publishing client at the mirror site, in the coming months. In addition, instead of using DML to perform data transfer requests, we will use the Bulk Data Movement client once its development is complete.

[1] ESG Project, "The Earth System Grid," www.earthsystemgrid.org, 2005.

[2] D. Bernholdt, S. Bharathi, D. Brown, K. Chancio, M. Chen, A. Chervenak, L. Cinquini, B. Drach, I. Foster, P. Fox, J. Garcia, C. Kesselman, R. Markel, D. Middleton, V. Nefedova, L. Pouchard, A. Shoshani, A. Sim, G. Strand, D. Williams, "The Earth System Grid: Supporting the Next Generation of Climate Modeling Research," *Proceedings of the IEEE,* vol. 93, pp. 485- 495, March 2005 2005.

[3] D. N. Williams, R. Ananthakrishnan, D. E. Bernholdt, S. Bharathi, D. Brown, M. Chen, A. L. Chervenak, L. Cinquini, R. Drach, I. T. Foster, P. Fox, D. Fraser, J. Garcia, S. Hankin, P. Jones, D. E. Middleton, J. Schwidder, R. Schweitzer, R. Schuler, A. Shoshani, F. Siebenlist, A. Sim, W. G. Strand, M. Su, and N. Wilhelmi, "The Earth System Grid: Enabling Access to Multi-Model Climate Simulation Data," *Bulletin of the American Meteorological Society (BAMS),* vol. 90, pp. 195–205, February 2009 2009.