# DREAM: Distributed Resources for the ESGF Advanced Management

Dean N. Williams[1], Luca Cinquini[3], V. Balaji[3], Dan Duffy[4], Valerio Pascucci[5]

[1]Lawrence Livermore National Laboratory
[2]Jet Propulsion Laboratory
[3]Princeton University
[4]Goddard Space Flight Center
[5]University of Utah

ESGF Earth System Grid Federation

## Motivation

Distributed Resources for the Earth System Grid Federation (ESGF) Advanced Management (DREAM) provides a new way to access large data sets across multiple DOE, NASA, and NOAA compute facilities, which will immediately improve climate research efforts as well as numerous other data-intensive applications. With its customizable user interface that communities of scientists can use to interact with each other, DREAM will provide a host of underlying services that can be adopted in part or as a whole, including services for publishing, discovering, moving, and analyzing data. The approach is applicable to field with large or small global data used for analytic data processing needs, including climate, biology, hydrology, and other science domains, resource management, security, and client applications.

DREAM is focused on providing services for dynamic management of resources held at diverse and scattered U.S. sites. Some of the present issues that we seek to address include how to:

- **Publish** dynamic resources into a virtual organization;
- **Discover** distributed resources and services through an intuitive, yet powerful query syntax;
- **Link** together various phases of a scientific execution;
- **Monitor and expose** the dynamic storage and computational resources, throughout the federation;
- **Execute** parallel distributed analytics across DOE, NASA, and NOAA data centers; and
- **Explore, analyze, and visualize** large-scale data remotely using commodity hardware.

## Approach

Our approach generalizes the current operational infrastructure used by the ESGF into a template architecture that hides each implementation layer behind a well-defined Application Programming Interface (API), so that different communities may decide to adopt or swap any single part. The deployable software stack will include the following modules:

- **Publishing services** (reference implementation based on ESGF publishing software);
- **Search services** (reference implementation based on Solr-Cloud and ESGF Search API);
- **Transfer services** (reference implementation based on Globus/GridFTP);
- **Computation services** (reference implementation based on Ultrascale Visualization Climate Data Analysis Tools (UV-CDAT) and ESGF Computing API);
- **Resource monitoring and allocation** (to be developed from scratch);
- **Security services** (reference implementation based on open standards and ESGF security infrastructure);
- **User interface** (reference implementation based on CoG web knowledge environment); and
- **Exploration services** (remote analysis and visualization based on streaming, multi-resolution data).

All APIs will be defined through a dynamic peer-review process and published in prominent web-accessible locations, and all implementations will be open-source, so they can be reviewed and improved by the community at large.

## Architecture

Our goal is to design a template architecture and build a reference implementation for the management, access and analysis of scientific resources in a distributed environment. We expect that DREAM, when applied to a generic scientific field, will provide new ways for scientists to leverage large-data resources to conduct research, thus greatly increasing the scientific throughput of that discipline.

The DREAM system architecture is schematically represented in the **Figure** below. The two basic principles behind the DREAM design are "**modularity**" and "**abstraction**". Modularity means that DREAM will be structured not as a single monolithic system, but rather as a composition of interacting software services, which are packaged and can be installed individually and independently. The functionality of each service will be abstracted in a well-defined API, so that each service can be easily invoked by other services and clients without worrying about the underlying implementation details. All service APIs will be defined to conform to the Representational State Transfer (REST) web service paradigm, which will allow simple invocation by standard web-enabled clients. In general, software systems that are both modular and abstract have an intrinsic longer longevity, because each service can be evolved or replaced individually, without affecting the backward compatibility with other parts of the systems, or its clients.

ESGF Peer-to-Peer Architecture

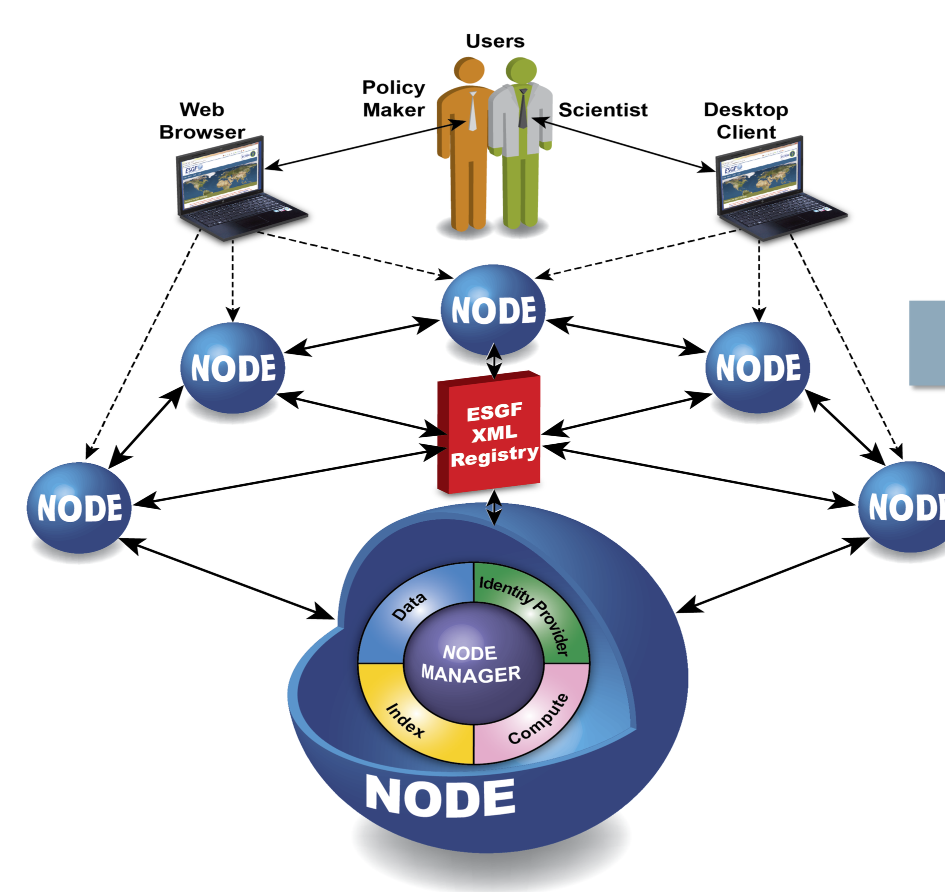


DREAM Template Architecture

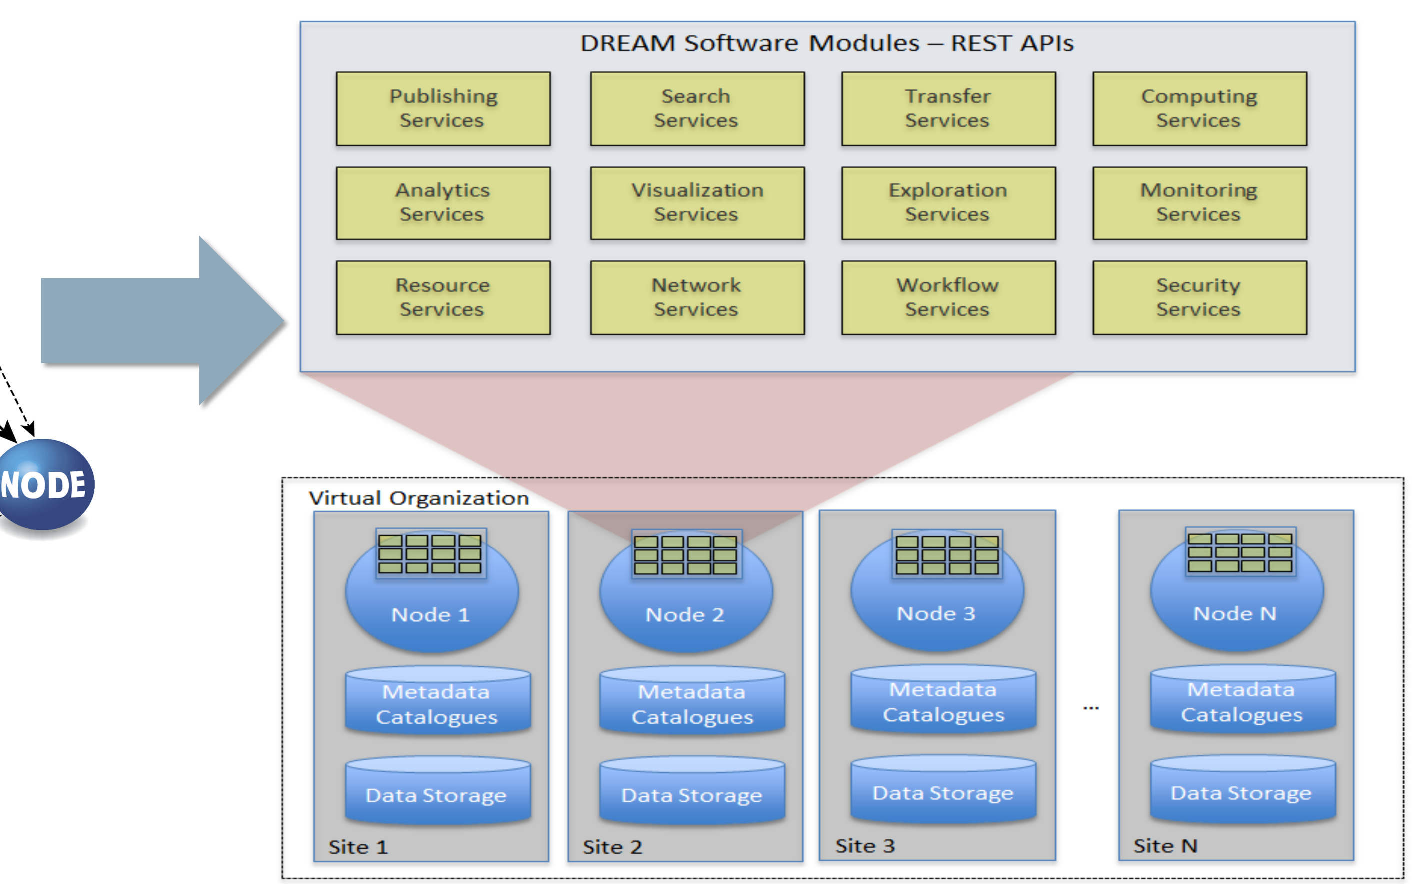


## Requirements

**Distributed**: DREAM will support a system composed of geographically distributed nodes, each with its own set of data resources, metadata catalogs, and software services. The DREAM protocols and services will unite all nodes in a single federation, so that a client will be able to discover, search, download and execute resources independent of their physical location, as if they were held and run on a single server.

**Dynamic**: the DREAM architecture will be designed to support a highly dynamic system, which will expose with minimum delay all data and services that are available at that time throughout the federation. For example, data collections produced by a model run, data streams originating from the real-time processing of a field instrument, or new derived products produced by a scientist running their processing code, could all be published into the system and immediately returned as results of federation-wide searches. Additionally, the system will continually provide an up-to-date report on the state of its components and will be able to automatically direct client requests where resources are available.

**Scalable**: the DREAM architecture will be able to scale to the "Big Data" volumes that are expected in several scientific fields (climate, astronomy, genomics, etc.) in the next 5-10 years. Scalability will be achieved through a two-fold approach. First, each service (shown in **Figure**) will be implemented through a high-performance technology that is inherently able to handle large volumes of data, in a distributed environment. Second, the DREAM modular architecture will allow each service to be instantiate multiple times (for example, for searching, downloading, etc.) so that the load can be spread across multiple servers".

**Resiliency and Fault Tolerance**: The DREAM architecture will be designed to include redundant components for all critical services (such as search, authentication, data download and visualization), and to execute automatic failover in case any component becomes unavailable. When backup services are not available, we will ensure that DREAM produces meaningful error responses to human and machine clients.

**Secure**: DREAM will support a distributed and federated security model, whereby each node will maintain complete control over the policies for accessing its local data and computational resources, while federation-wide authentication and authorization services will be responsible for enforcing these policies homogenously through the system. Single-sign-on and federated access control will allow users to register and authenticate only once, and then propagate their identity and attributes as they access resources through the system, or request the system to access resources on their behalf.

## Motivating Use Case

One of the first applications of the DREAM architecture will be the Accelerated Climate Modeling for Energy (ACME) project, which is sponsored by the Earth System Modeling program within the U.S. Department of Energy's Office of Biological and Environmental Research. Its primary goals is to advanced model development by executing and end-to-end workflow infrastructure that automates labor intensive tasks, providing intelligent support for complex tasks and reducing duplication of effort through collaboration support.

ACME End-to-End Workflow Process

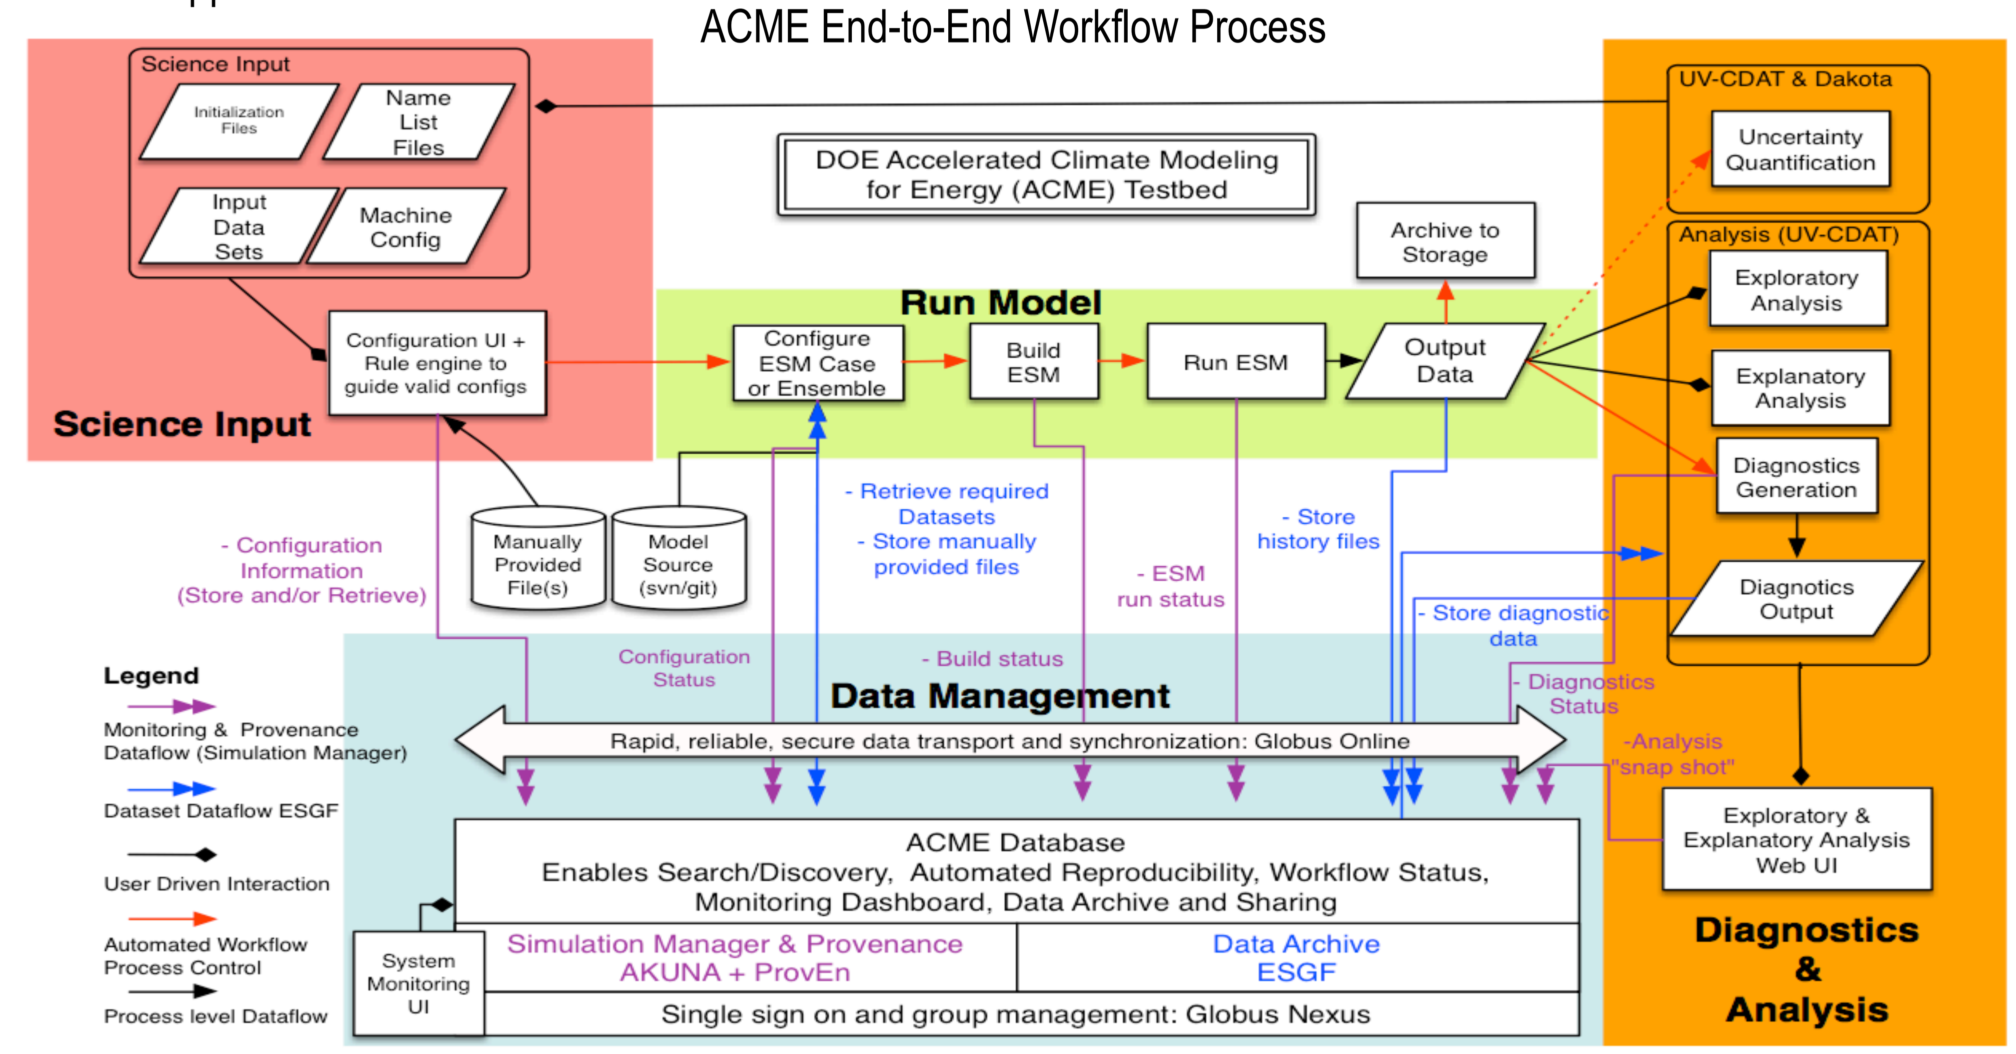


This work is supported by the ESGF consortium: DOE, NASA, NOAA, IS-ENES, NCI.