High Level Architecture and APIs

DREAM's 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. "DREAM" stands for Distributed Resources for the ESGF Advanced Management. 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 **Figure**. 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.

The DREAM architecture will be designed to fulfill the following list of requirements:

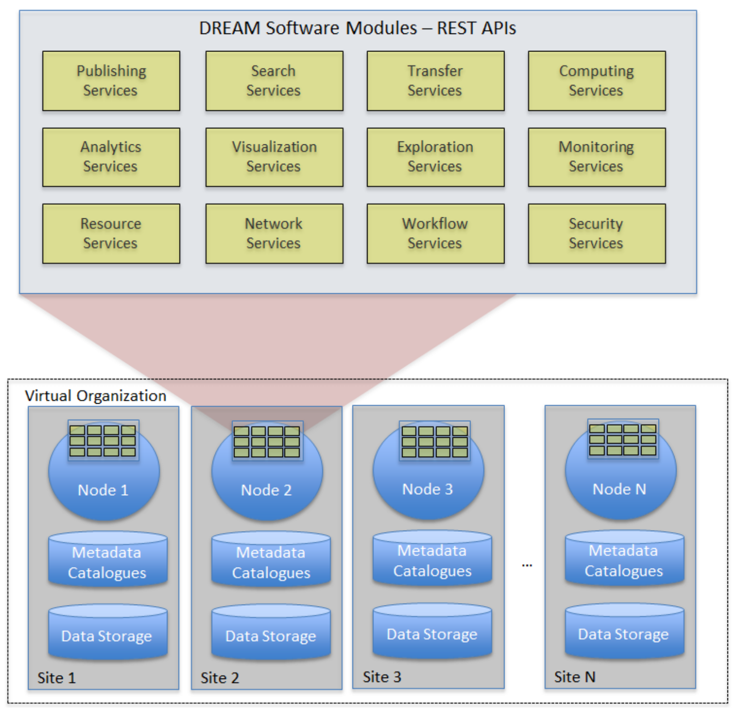

Figure 4. *High-level DREAM template architecture, highlighting its modular structure (lower box) and its distributed topology (upper box).*

- **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. Resiliency and redundancy will be built into the system, so that a client will not be affected by the planned downtime, or unexpected outage, of any single node.
- **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 (the on/off status and current computational load of the services at each node), 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 (search, data movement, computation, etc.) will be implemented through a high-performance technology that is inherently able to handle large volumes of data, in a distributed environment (see the description of each service for examples of such technologies). Second, the DREAM modular architecture will allow each node to scale "horizontally", i.e. to instantiate multiple instance of the same service to process a higher number of client requests. For example, a node could run several data servers, or split its metadata holdings across multiple catalogs that are searched simultaneously before joining the results.
- **Resiliency and Fault Tolerance**: There is a need to prioritize resiliency in the DREAM system now and as new features are developed. We know that resilience is a systems problem, not an individual component problem, and requires a systems approach. Therefore, we will ensure that the DREAM system produces a meaningful error response. In many cases, new features should be introduced when previous features are sufficiently robust. This priority shift will set a good precedent in subsequent DREAM development as well as provide a superior experience for the scientific end users.
- **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. Implementation of the security services will be based on industry standards and the latest cryptography libraries, so to guarantee maximum protection versus unauthorized access.