Simplifying Data Management Middleware Through Content-Based Network Services

Big data has emerged as a driving force for scientific discoveries. Large scientific instruments (e.g., colliders, light sources, and telescopes) generate exponentially increasing volumes of data. Primary examples include the Large Hadron Collider (LHC) experiments, which generate hundreds of petabytes of data per year now, and are projected to reach the exabyte/yr level in 3-5 years. The largest "big data" science experiments consist of global collaborations with hundreds of sites and thousands of scientists. To enable scientific discovery within such large collaborations, the experiment data needs to be collected, processed, indexed, archived, and analyzed as a collective effort across the collaboration. The sheer volume of data within a large scale science experiment normally taxes the collective computing and storage resources available to the collaboration. In order to make the most efficient use of those resources, highly distributed computing systems have been developed to service the computing work flows. Within such systems, remotely-located available computing resources can be put to use if local resources are not available. The high speed networks connecting the data centers that support these distributed computing resources have become a critical component in the overall distributed computing system. However, those high speed networks merely move data from one location to another. Layers of middleware are required to provide the functions of data management & discovery, data security, and data transfer. While the current generation of middleware is highly effective in managing these distributed computing systems, it is highly complex and adds performance overhead to that data management and movement. Put another way, while today's middleware may be highly effective, it is not highly efficient.

Looking out over the long term (10yr) time frame, we anticipate the following scenario for large-scale science "big data" experiments:

- Data volumes will increase 1-2 orders of magnitude. This is based on projects by existing large-scale science experiments, such as LHC and Climate Science
- Large-scale science distributing computing environments will be more complex, particularly with regard to CPU resources. Cloud services are already emerging as cost-competitive alternatives for CPU resources. High Performance Computing (HPC) resources may also emerge as a viable option for distributed computing systems with high throughput computing characteristics.

The combination of sharply increased data volumes and more complex distributed computing environments will present extremely high hurdles for existing methods of middleware-driven data management and movement. We believe a major challenge facing big data science in the exascale era will be the development of more efficient models for data management and movement. The current location-centric architecture for IP network services makes a middleware-driven approach necessary. However, new network architectures have been proposed that could significantly reduce and substantially simplify middleware for data management and movement in large distributed computing systems. Of particular interest is Named Data Networking (NDN), which proposes a network service based on a data-centric (content) architecture, instead of a location-centric (network address) one. Under NDN, many current middleware services involving data location and transfer services could be provided as basic network services. NDN caching and repository functions could similarly optimize data movement, simplifying the overall data management load on the distributed computing middleware. In the coming exascale computing / terabit networking environment, such middleware simplifications and performance optimizations would be highly beneficial, if not essential, for the highly distributed computing systems that large-scale science experiments are expected to utilize. Content-driven data retrieval techniques are already successfully being used within a number of large-scale science projects, including all four of the LHC experiments. These content-driven techniques are currently implemented in middleware. We encourage the Office of Science to investigate content-driven network services for large-scale science, with particular emphasis on the development of synergies between such services and data management methods employed by large-scale science.