DWDM-RAM: A Data Intensive Grid Service Architecture Enabled by Dynamic Optical Networks

T. Lavian[1], J. Mambretti[2] ,D. Cutrell[1] , H. Cohen[1], S. Merrill[1], R. Durairaj[1], P. Daspit[1] I. Monga[1], S. Naiksatam[3], S. Figueira[3]. D. Gutierrez[4], D. Hoang[5], F. Travostino[1]
[1]Nortel Networks Labs, [2]iCAIR Northwestern University, [3]SCU, [4]Stanford University, [5]UT Sydney.
{tlavian, hcohen, dcutrell, smerrill, radurai, , pdaspit, imonga, travos}@nortelnetworks.com, j-mambretti@northwestern.edu, {sfigueira, snaiksatam}@scu.edu, degm@stanford.edu, dhoang@it.uts.edu.au

Abstract

Next generation applications and architectures (for example, Grids) are driving radical changes in the nature of traffic, service models, technology, and cost, creating opportunities for an advanced communications infrastructure to tackle next generation data services. To take advantage of these trends and opportunities, research communities are creating new architectures, such as the Open Grid Service Architecture (OGSA), which are being implemented in new prototype advanced infrastructures.
The DWDM-RAM project, funded by DARPA, is actively addressing the challenges of next generation applications. DWDM-RAM is an architecture for data-intensive services enabled by next generation dynamic optical networks. It develops and demonstrates a novel architecture for new data communication services, within the OGSA context, that allows for managing extremely large sets of distributed data. Novel features move network services beyond notions of the network as a managed resource, for example, by including capabilities for dynamic on-demand provisioning and advance scheduling. DWDM-RAM encapsulates optical network resources (Lambdas, lightpaths) into a Grid Service and integrates their management within the Open Grid Service Architecture. Migration to emerging standards such as WS-Resource Framework (WS-RF) should be staright forward.

In initial applications, DWDM-RAM targets specific data-intensive services such as rapid, massive data transfers used by large scale eScience applications, including: high-energy physics, geophysics, life science, bioinformatics, genomics, medical morphometry, tomography, microscopy imaging, astronomical and astrophysical imaging, complex modeling, and visualization.

Keywords: *DWDM, 10Gbs Lambda-switched networks, Data-intensive applications, eScience, Optical network resources, Scheduled Grid services, Dynamic wavelength provisioning, ODIN, OGSA, OGSI.*

Extended abstract

Common architectures that underlie traditional data networks do not incorporate capabilities required by Grids. They are generally designed to optimize the relatively small data flow requirements of consumer services and the managed services of enterprises on a common core infrastructure oriented to the requirements of general, common communication services. Many Grid applications are data-intensive, requiring specialized services and infrastructure to manage efficiently multiple, large-scale data flows of multiple terabytes and even petabytes. Such capabilities are not effectively possible in traditional routed packet data networks. Typical Grid applications require the management of highly distributed resources within dynamic environments. Basic problems related to these requirements are common to almost all Grid environments, e.g., matching multiple, and potentially conflicting, application requirements to diverse, distributed resources within a dynamic environment. Others are more specific to addressing the complexity of utilizing methods for data provisioning for large-scale data flows.

Although the existing communications model of packet-switching has served well for transporting burst transmission of short data packets, e.g., for consumer oriented email and web applications, it has not been sufficiently adaptable to meet the challenge of large scale data. In order to provide inexpensive access to advanced computation capabilities and extremely large data sets, it is necessary to create a new type of infrastructure which manages codependent storage and network resources. The DWDM-RAM project contributes to these efforts by encapsulating "optical network resources" into the Grid services framework to support dynamically provisioned data-intensive transport services. One of its novel features is the explicit representation of future scheduling in the data and network resource management model. Another consists of a set of protocols for managing dynamically provisioned wavelengths.

The proposed architecture consists of two layers between an application and the underlying optical network: the application middleware layer and the resource middleware layer. The application middleware layer speaks for the requirements of the application and presents an interface to users. This layer also shields the application from all aspects of sharing and managing the required resources. The resource middleware layer provides services that satisfy the resource requirements of the application, as specified or interpreted by the application middleware layer services. This layer contains interfaces and services that initiate and control sharing of the underlying resources.

At the application middleware layer, the Data Transfer Service (DTS) presents an interface between the system and an application. It receives high-level client requests, policy-and-access filtered, to transfer specific named blocks of data with specific advance scheduling constraints. It employs an intelligent strategy to schedule an acceptable action plan that balances user demands and resource availabilities. The action plan involves advance co-reservation of network and storage resources. The application expresses its needs only in terms of high level tasks without knowing how they are processed at the layers below. It is this middleware layer that shields the application from low level details by translating application-level requests to its own tasks of coordinating and controlling the sharing of a collective set of resources.

The network resource middleware layer consists of three services: the Data Handler Service (DHS), the Network Resource Service (NRS) and the Dynamic Lambda Grid Service (DLGS). Services of this layer initiate and control sharing of resources. The DHS deals with the mechanism for sending and receiving data and effectuates the actual data transfer when needed by the DTS. Central to our intended architecture is the Network Resource Service. It is this service that makes use of the Dynamic Lambda Grid Service encapsulating the underlying optical network resources into an accessible, schedulable Grid service. The NRS receives requests from the DTS, as well as requests from other services such as Grid services (both scheduled and on-demand). It maintains a job queue and allocates proper network resources according to its schedule. To allow for extensibility and reuse, the Network Resource Service can be decomposed into two closely coupled services: a Basic Network Resource Service and a Network Resource Scheduler. The Basic Network Resource Service presents an interface to the Data Transfer Service for making network service requests and handles multiple low level services offered by different types of underlying networks and switching technologies. The Network Resource Scheduler is responsible for implementing an effective schedule for network resources sharing. The Network Resource Scheduler can be developed independently of the Basic Network Resource Service. This provides the NRS the flexibility to deal with other scheduling schemes as well as other types of dynamic underlying networks.

DWDM-RAM is experimental as well as conceptual. It is being implemented in prototype on OMNInet. OMNInet is an operational metro-area photonic testbed located in Chicago. It was developed under partnership between Nortel, SBC and iCAIR/Northwestern University to explore photonic networking issues and demonstrate the value of providing 10GE services on wavelengths, supported by advanced photonic networking technology and new types of dynamic intelligent control signaling and related protocols, such as ODIN.

The DWDM-RAM architecture, integrated with OMNInet's dynamically switched wavelength services supports: 1) both on-demand and scheduled data retrieval, 2) a meshed DWDM switched network capable of establishing an end-to-end lightpath in seconds, 3) bulk data-transfer facilities using lambda-switched networks, and 4) out-of-band tools for adaptive placement of data replicas. In summary, the DWDM-RAM architecture illustrates the value of scheduling high-bandwidth network resources through dynamic lightpath provisioning. It makes data services easily accessible via a Grid service interface, by encapsulating the the optical network resources (e.g., Lambdas, lightpaths) and integrating their management to OGSA framework.

[The DWDM-RAM project is sponsored by DARPA under TIA F30602-98-2-0194]