DWDM-RAM: Enabling Grid Services with Dynamic Optical Networks

S. Figueira, S. Naiksatam, H. Cohen,
D. Cutrell, P. Daspit, D. Gutierrez,
D. Hoang, T. Lavian, J. Mambretti,
S. Merrill, F. Travostino

DWDM-RAM

- DARPA-funded project
  - Santa Clara, CA
    - Nortel Networks
    - Santa Clara University
  - Chicago, IL
    - iCAIR / Northwestern University
  - Australia
    - University of Technology, Sydney

DWDM-RAM

- Goal
  - Make dynamic optical network usable by grid applications
  - Provide lightpaths as a service
  - Design and implement in prototype a new type of grid service architecture optimized to support data-intensive grid applications through advanced optical network

Why Dynamic Optical Network?

- Packet-switching technology
  - Great solution for small-burst communication, such as email, telnet, etc.
- Data-intensive grid applications
  - Involves moving massive amounts of data
  - Requires high and sustained bandwidth

Why Dynamic Optical Network?

- DWDM
  - Basically circuit switching
  - Enable QoS at the Physical Layer
  - Provide
    - High bandwidth
    - Sustained bandwidth

Why Dynamic Optical Network?

- DWDM based on dynamic wavelength switching
  - Enable dedicated optical paths to be allocated dynamically

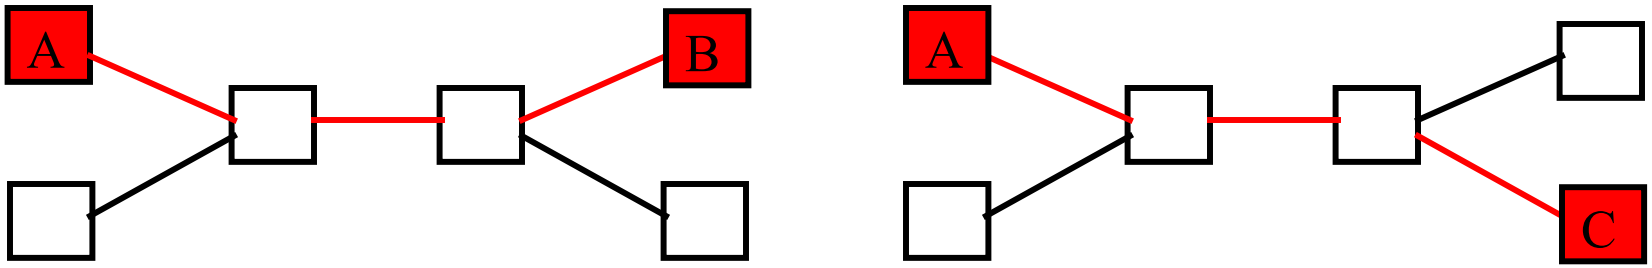

In a few seconds…

Why Dynamic Optical Network?

- Any drawbacks?
  - The overhead incurred during end-to-end path setup
- Not really a problem
  - The overhead is amortized by the long time taken to move massive amounts of data

Why Dynamic Optical Network?

When dealing with data-intensive applications, overhead is insignificant!

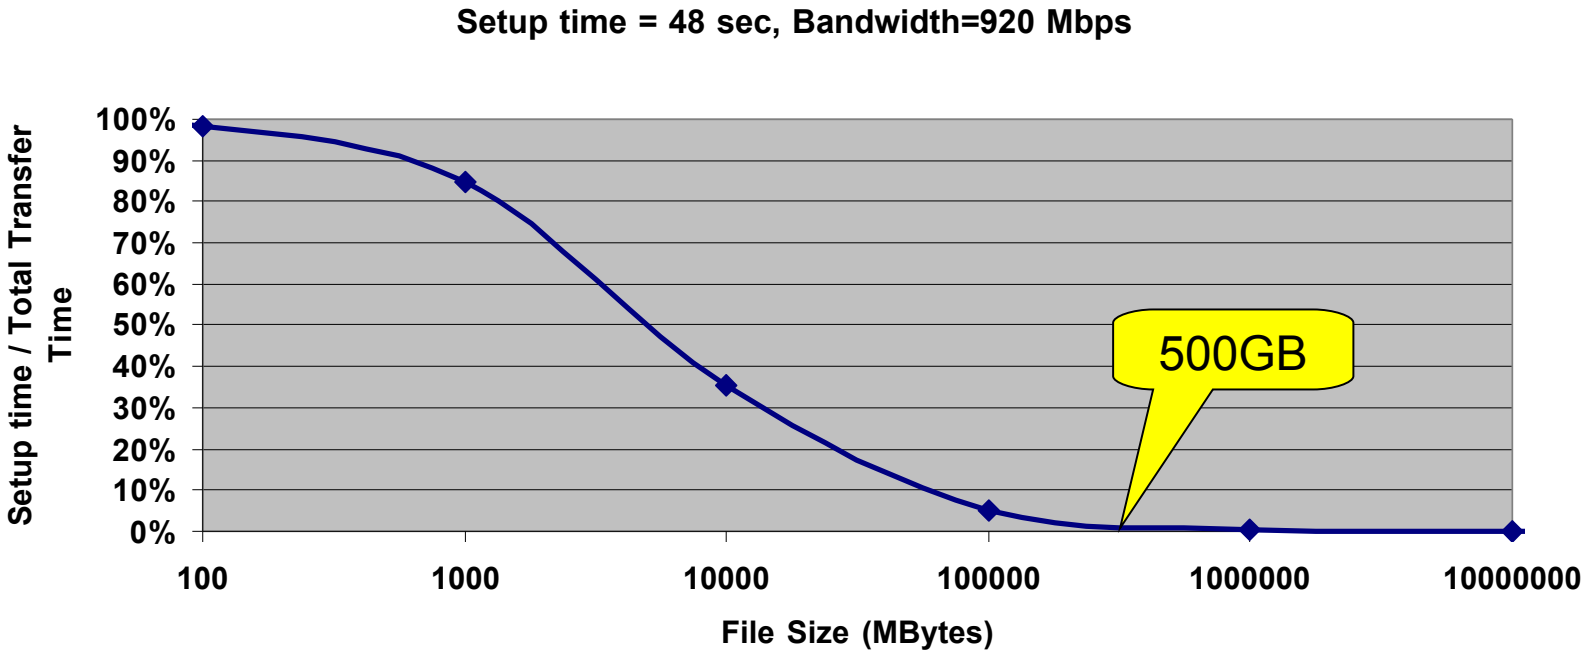

Why Grid Services?

- Applications need access to the network
  - To request and release lightpaths
- Grid services
  - Can provide an interface to allocate and release lightpaths

DWDM-RAM Architecture

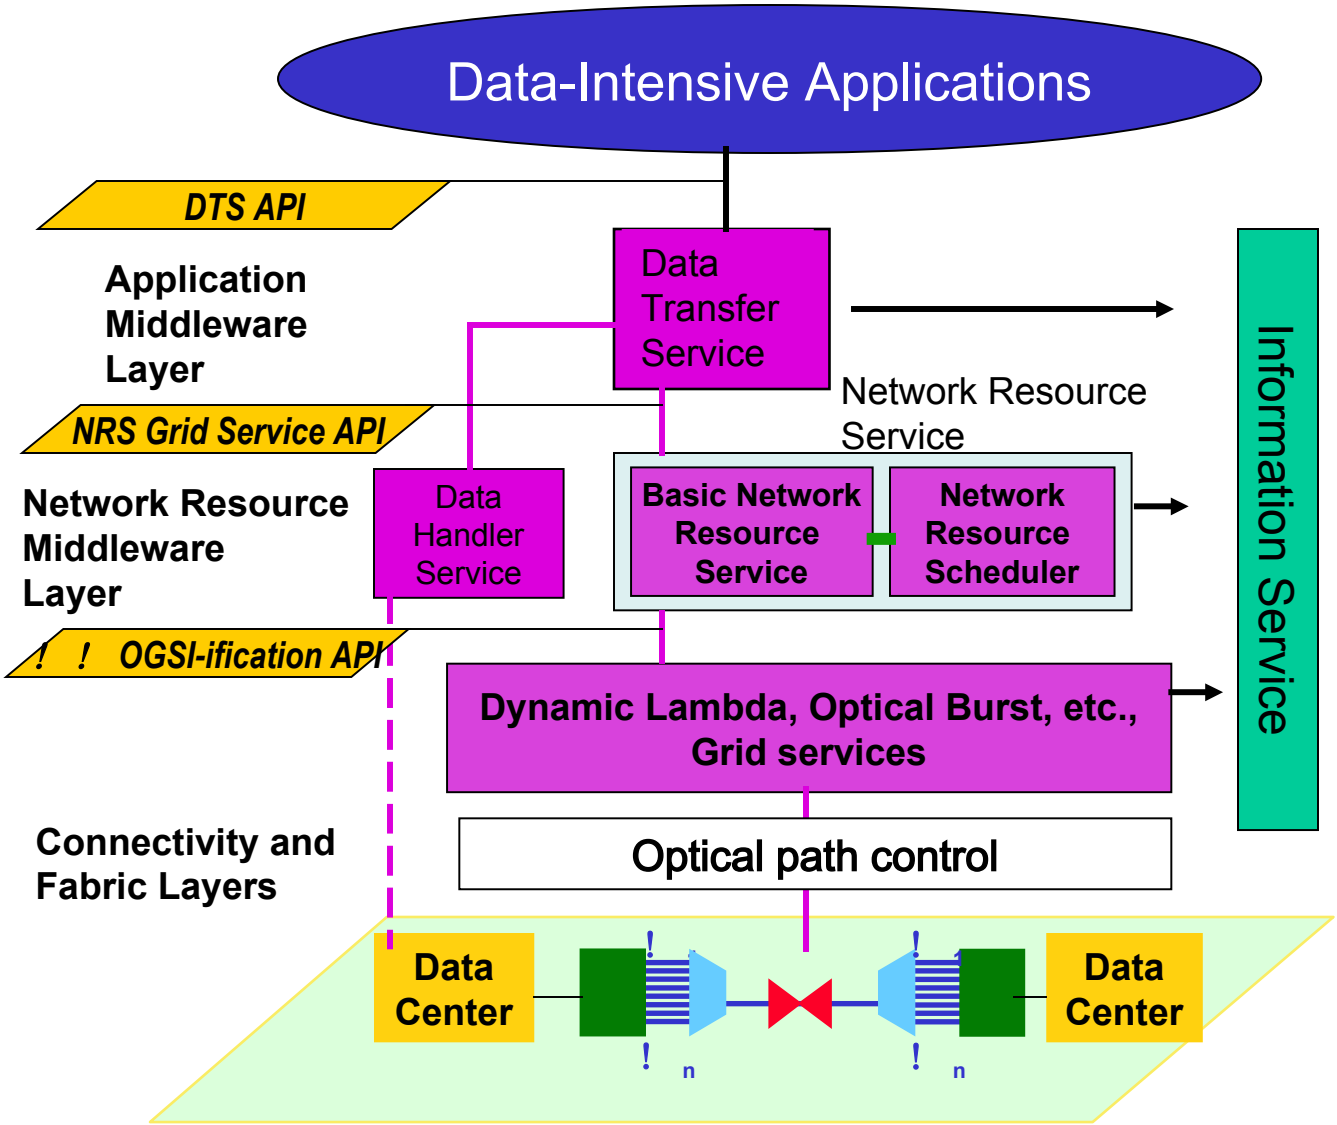

DWDM-RAM Architecture

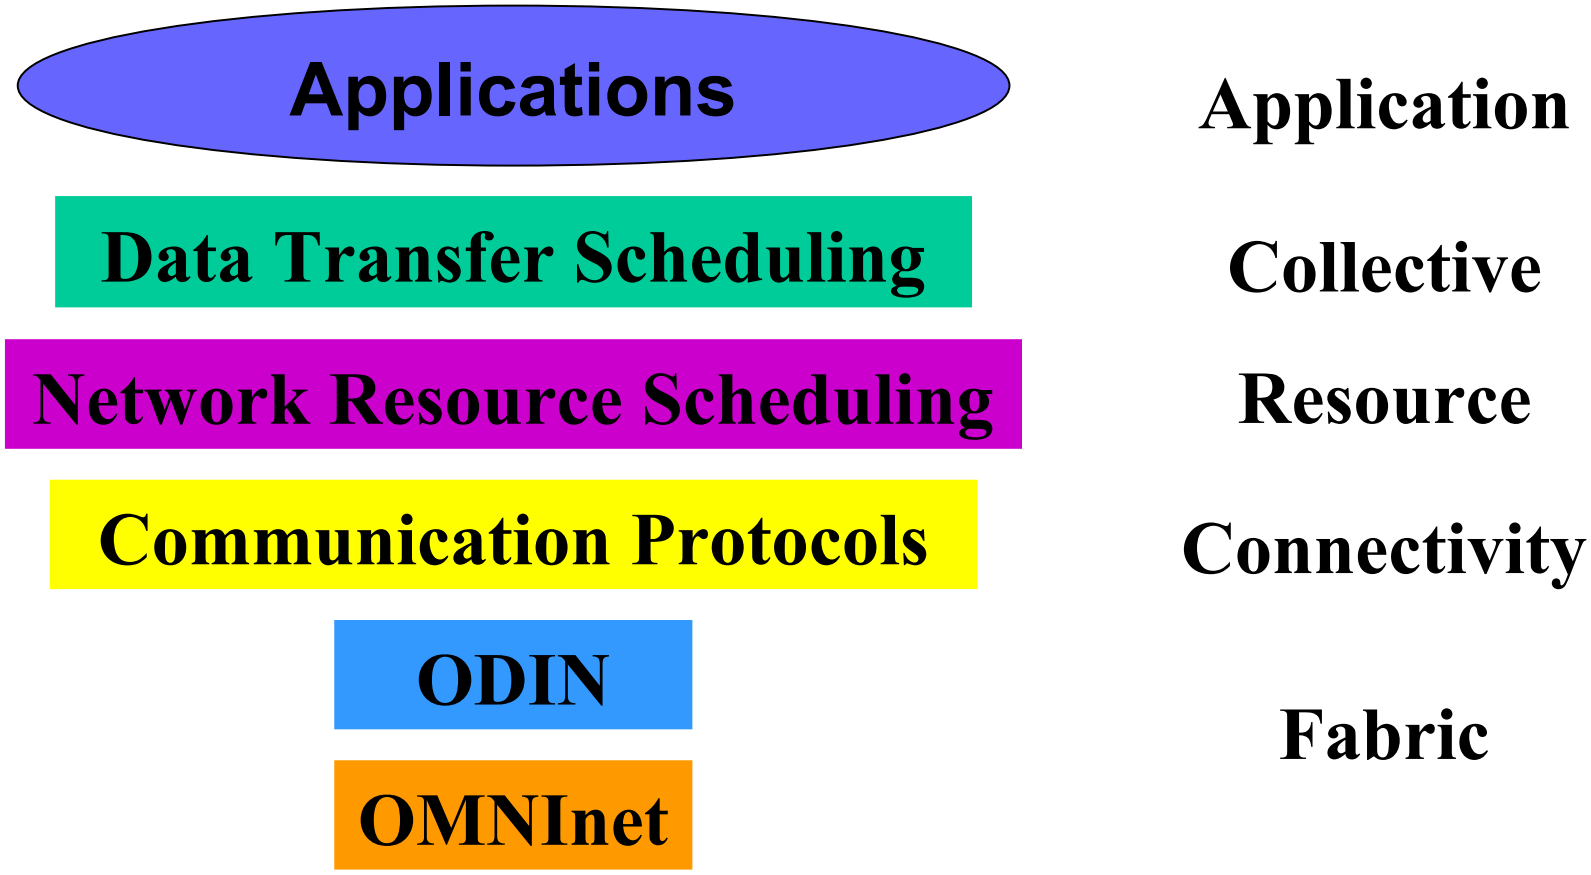

DWDM-RAM Architecture

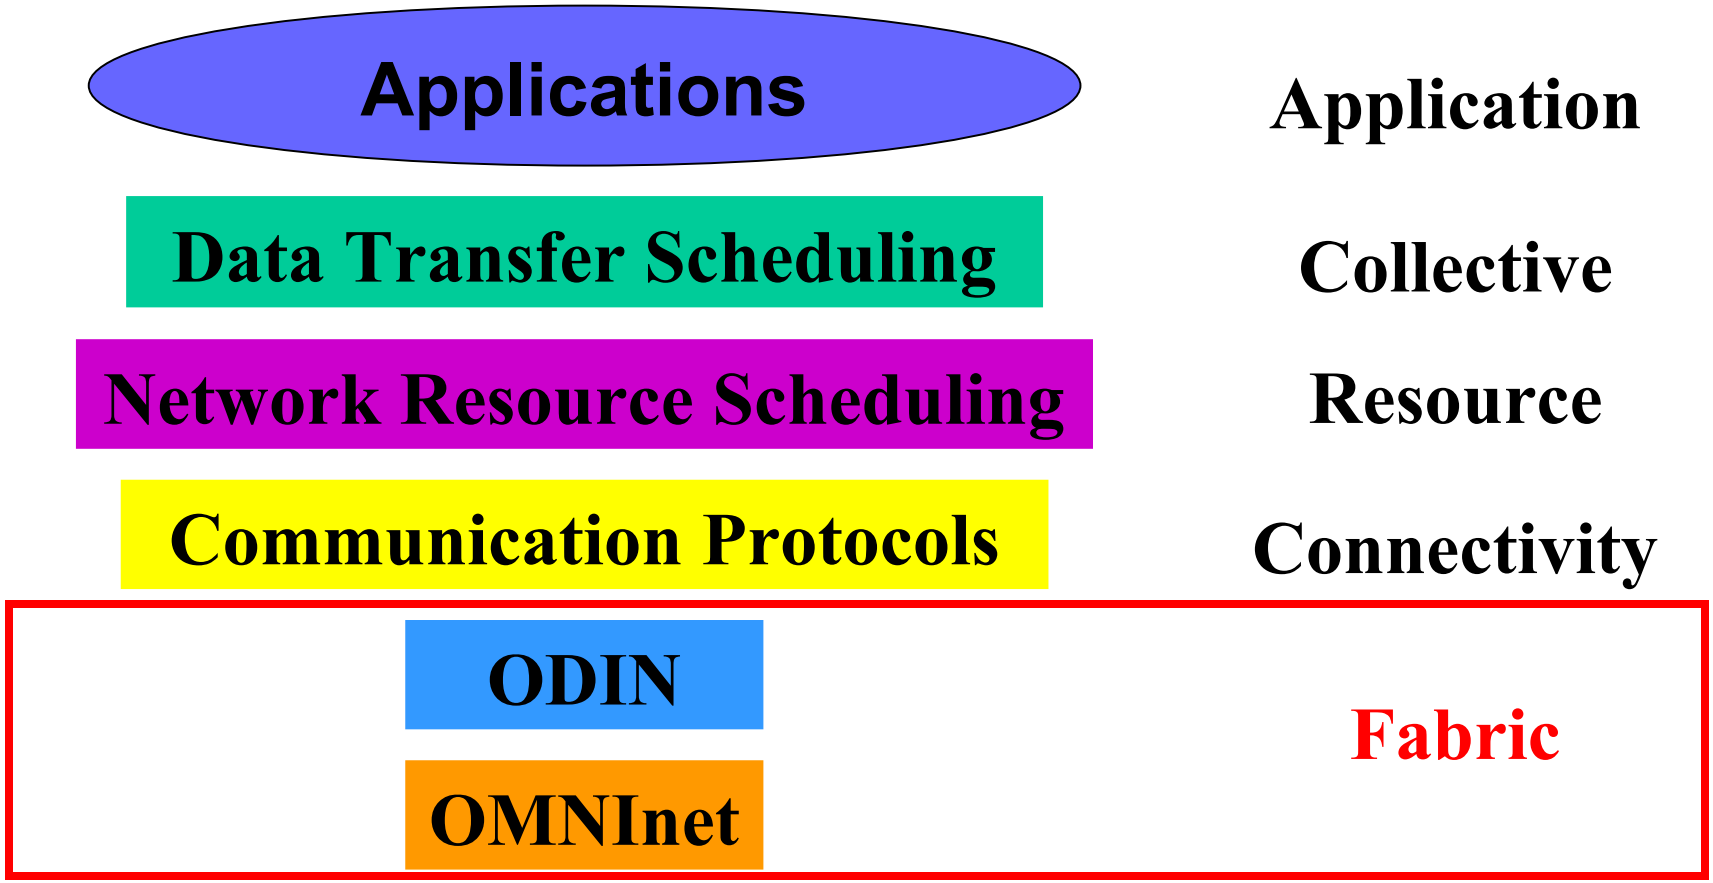

DWDM-RAM Architecture

- OMNInet - photonic testbed network
  - Four-node multi-site optical metro testbed network in Chicago -- the first 10GE service trial!
  - All-optical MEMS-based switching and advanced high-speed services
  - Partners: SBC, Nortel, iCAIR at Northwestern, EVL, CANARIE, ANL

OMNInet Core Nodes

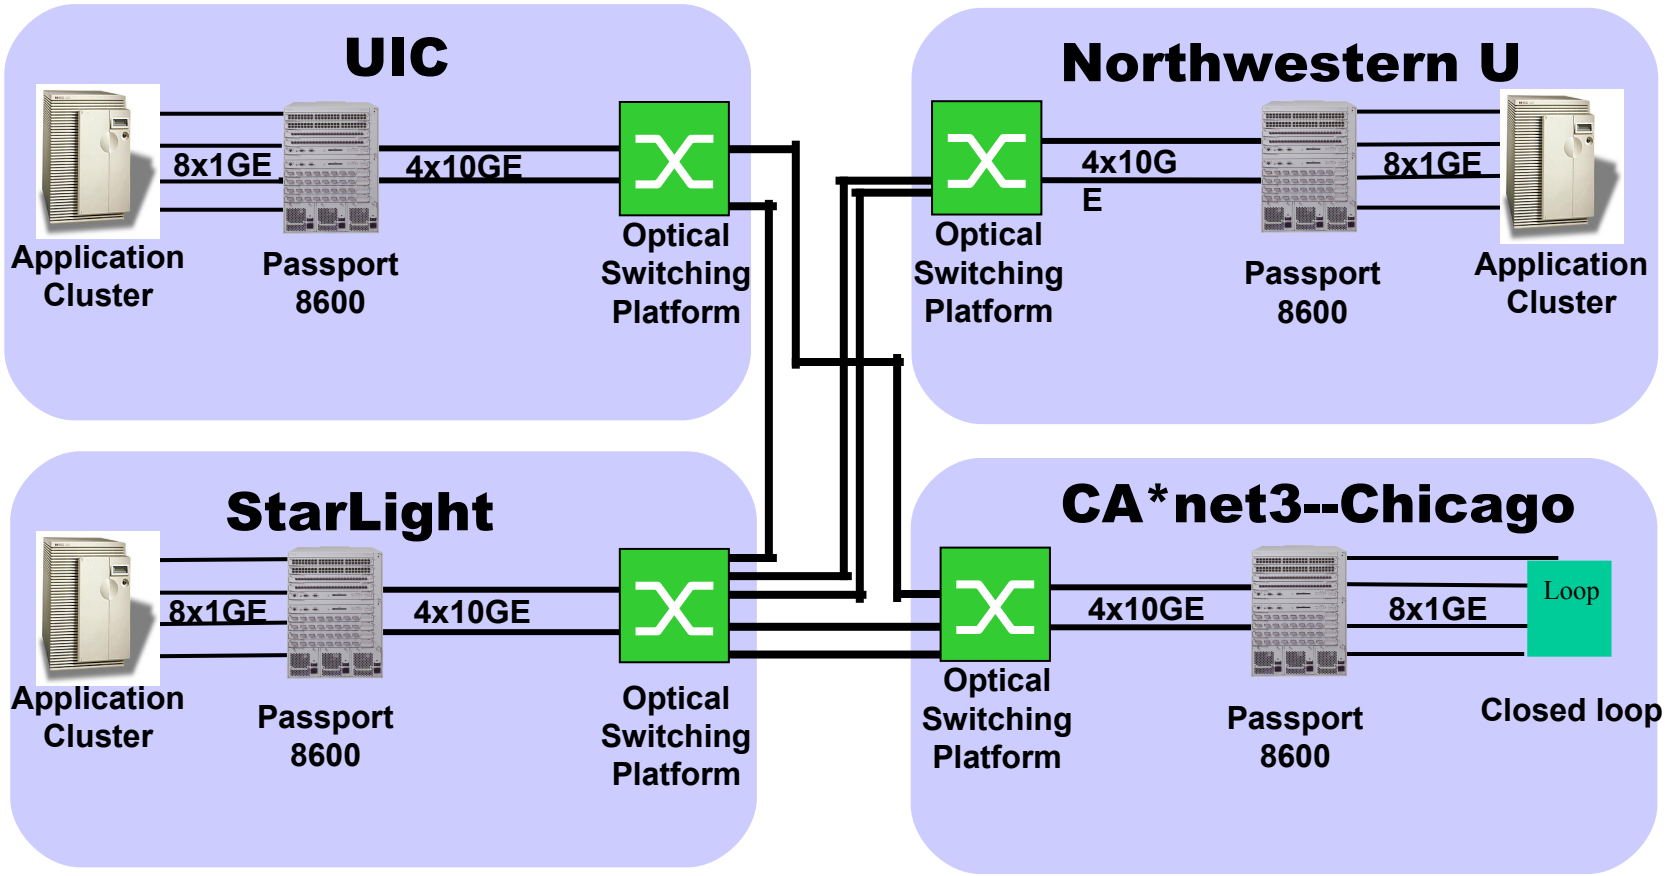

DWDM-RAM Architecture

- ODIN - Optical Dynamic Intelligent Network
  - Software suite that controls the OMNInet through lower-level API calls
  - Designed for high-performance, long-term flow with flexible and fine grained control
  - Stateless server, which includes an API to provide path provisioning and monitoring to the higher layers

DWDM-RAM Architecture

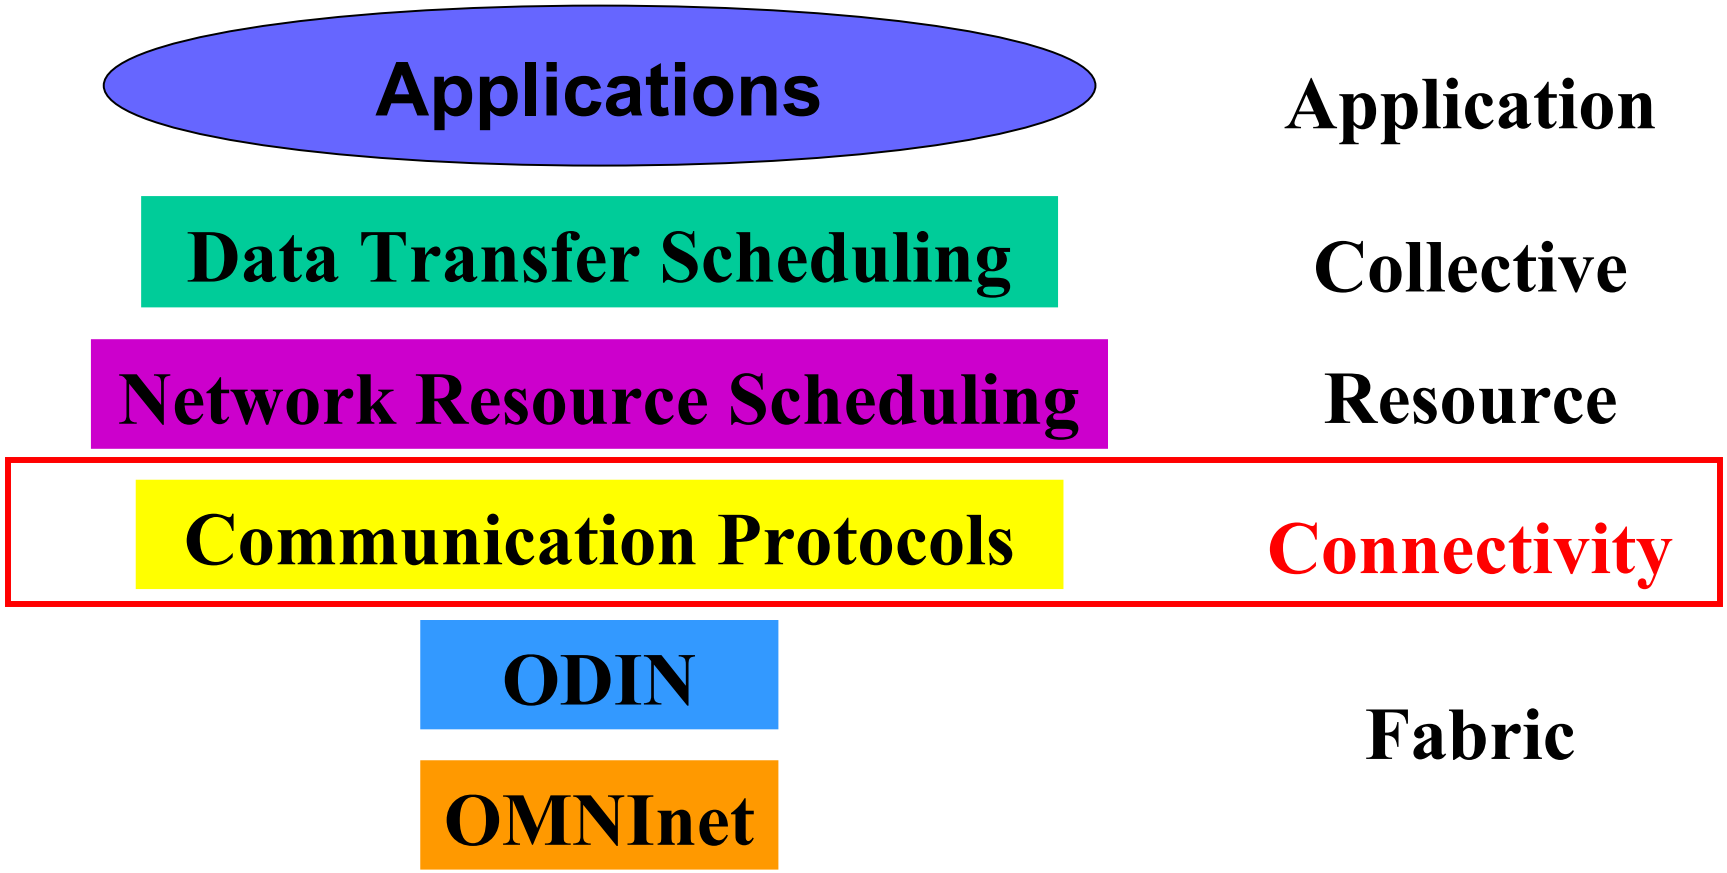

DWDM-RAM Architecture

- Communication Protocols
  - Currently, using standard off-the-shelf communication protocol suites
  - Provide communication between application clients and DWDM-RAM services and between DWDM-RAM components
  - Communication consists of mainly SOAP messages in HTTP envelopes transported over TCP/IP connections

DWDM-RAM Architecture

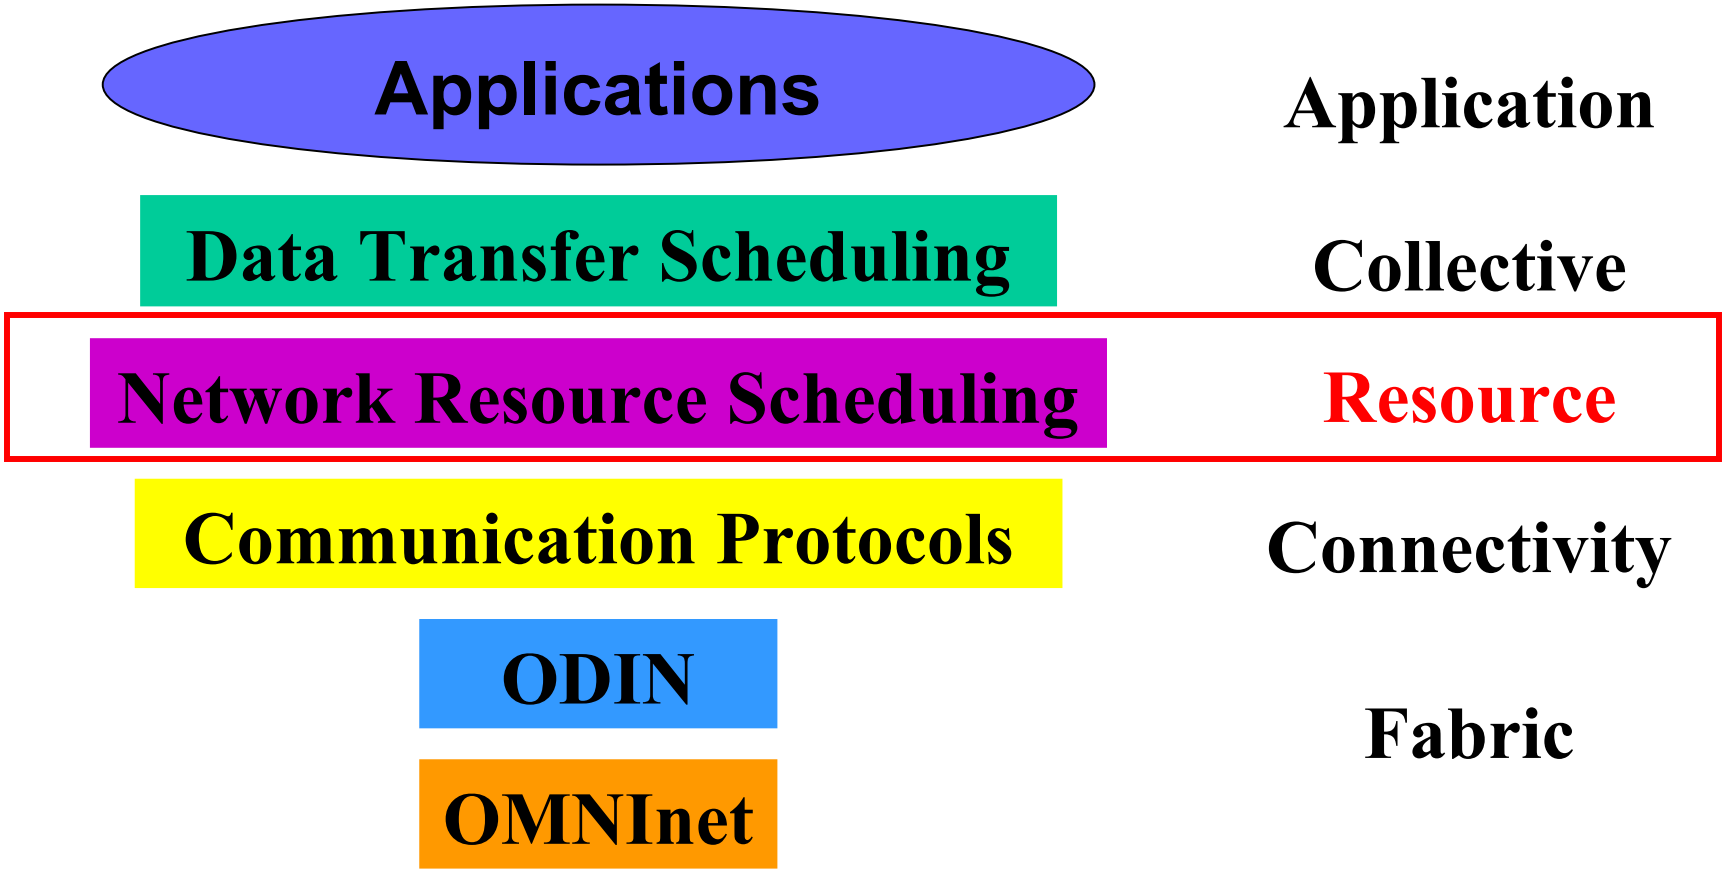

DWDM-RAM Architecture

- Network Resource Scheduling
  - Essentially a resource management service
  - Maintains schedules and provisions resources in accordance with the schedule
  - Provides an OGSI compliant interface to request the optical network resources

DWDM-RAM Architecture

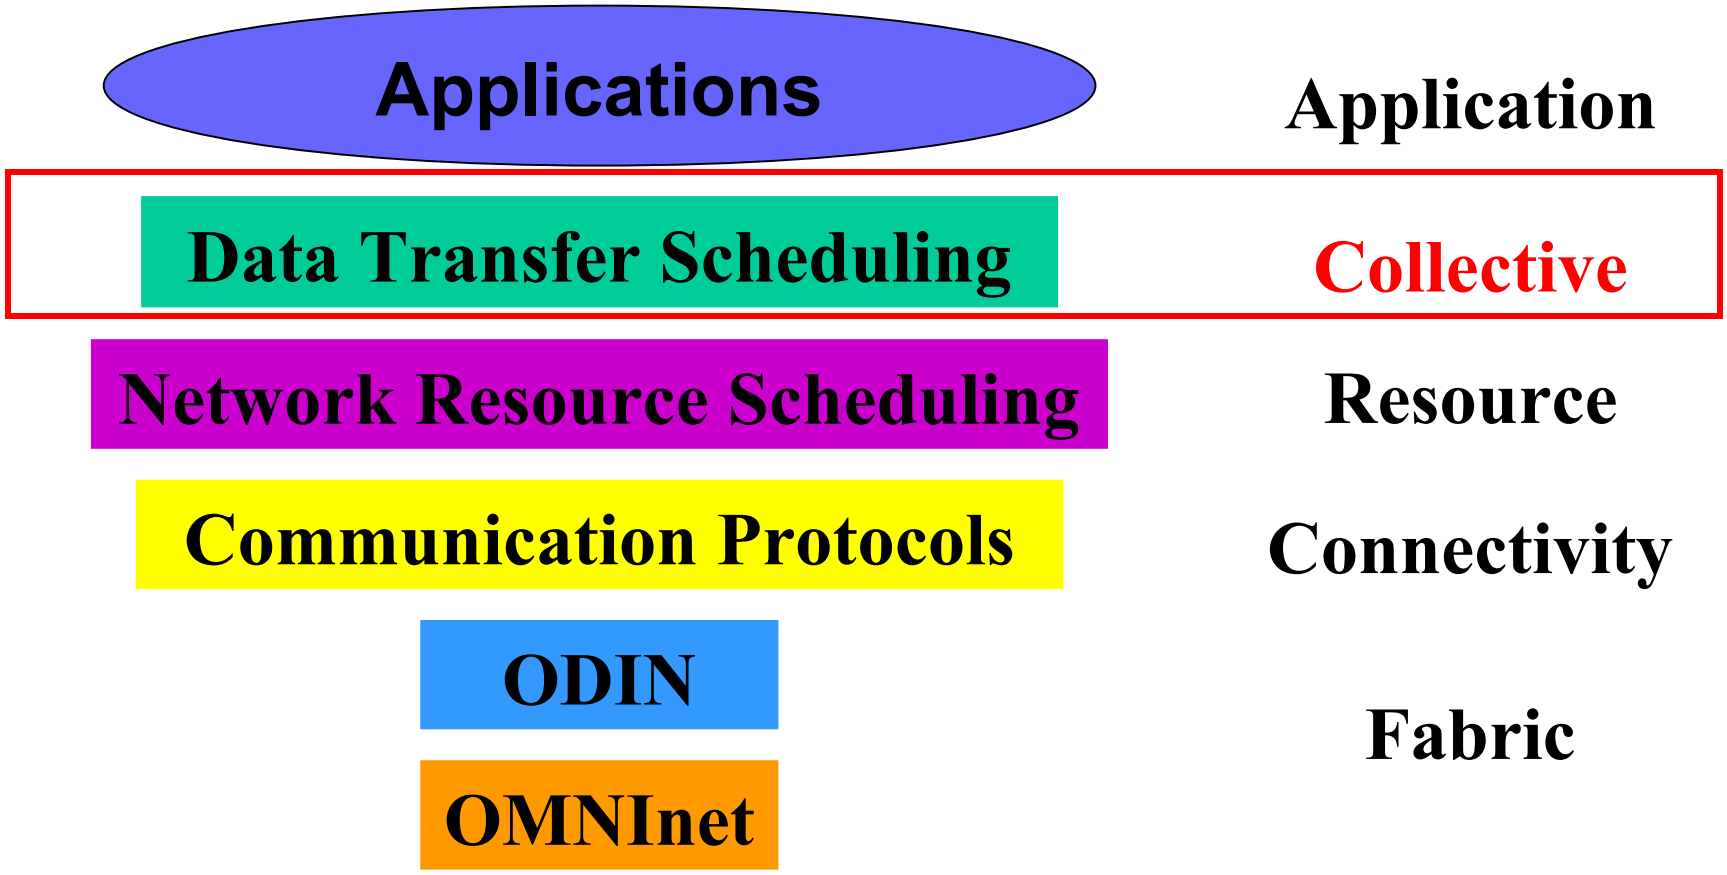

DWDM-RAM Architecture

- Data Transfer Scheduling
  - Direct extension of the NRS service, provides an OGSI interface
  - Shares the same backend scheduling engine and resides on the same host
  - Provides a high-level functionality
  - Allow applications to schedule data transfers without the need to directly reserve lightpaths
  - The service also perform the actual data transfer once the network is allocated

Data Transfer Scheduling

- Uses standard ftp
- Uses NRS to allocate lambdas
- Uses OGSI calls to request network resources

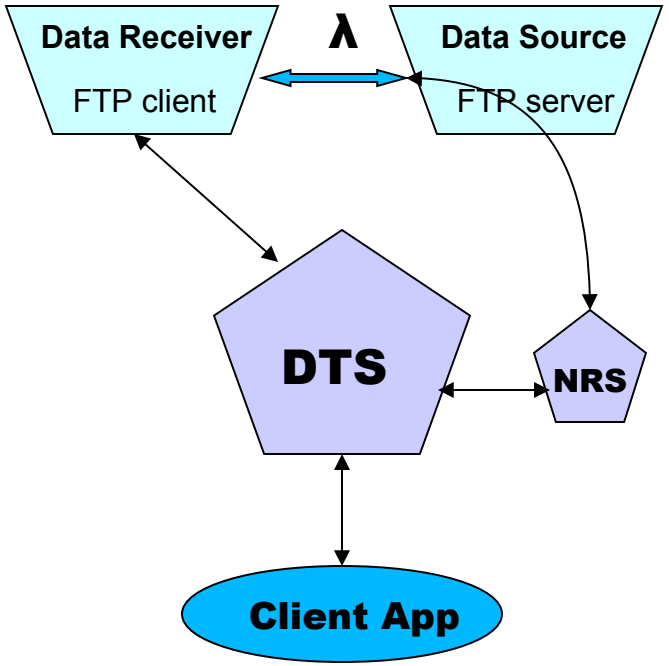

DWDM-RAM Architecture

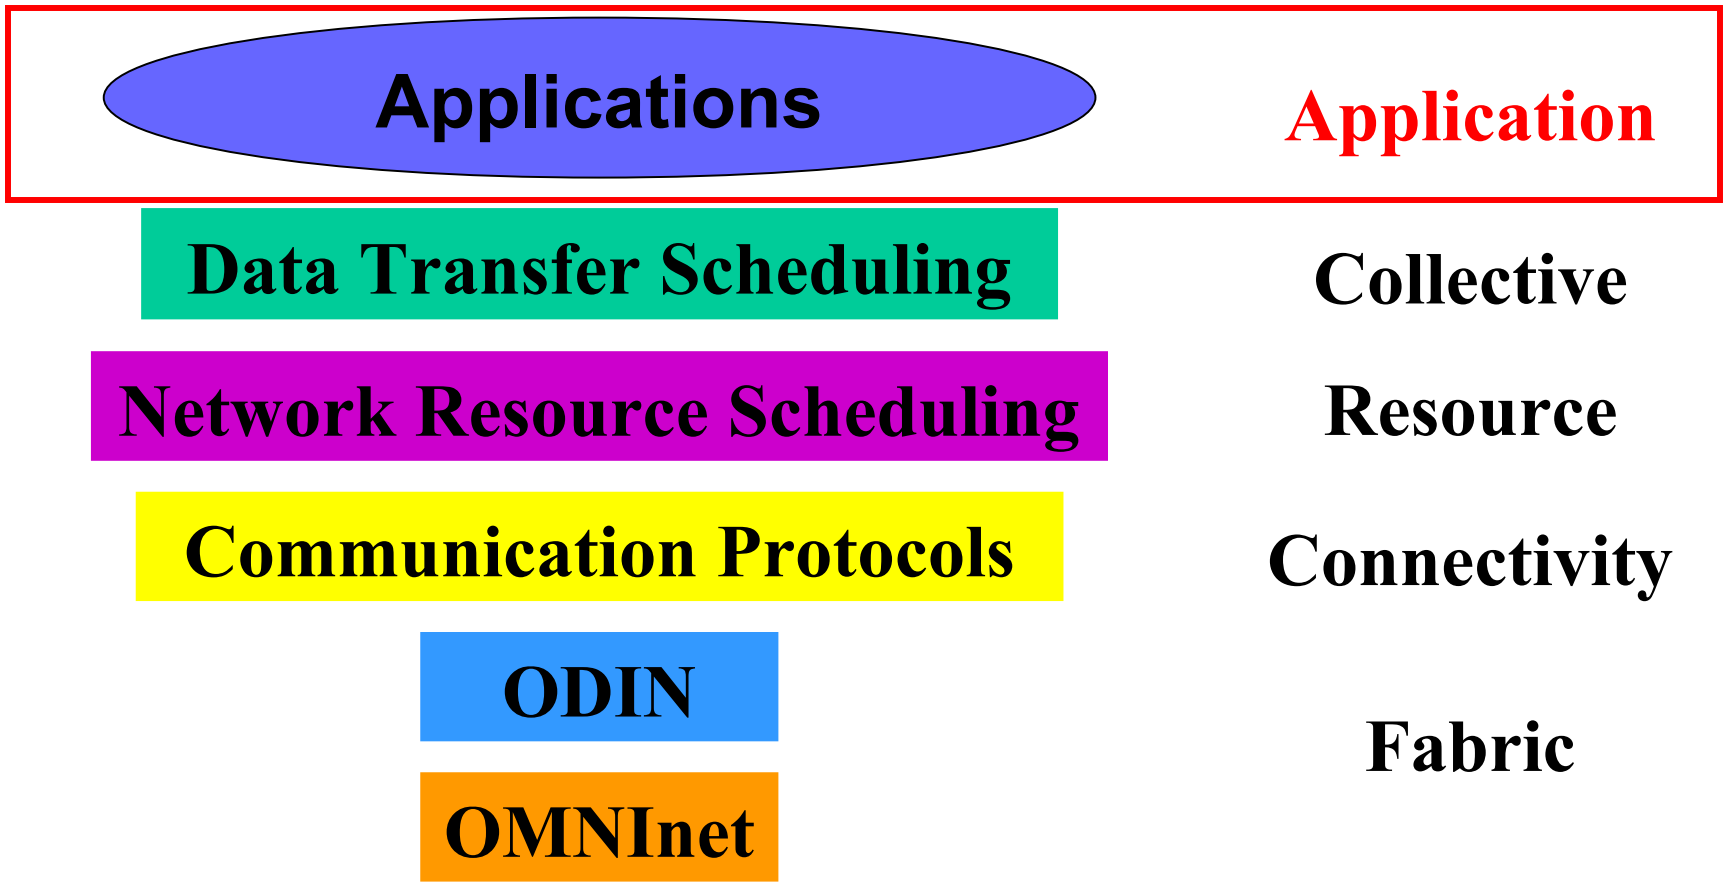

DWDM-RAM Architecture

- Applications
  - Target is data-intensive applications since their requirements make them the perfect costumer for DWDM networks

DWDM-RAM Modes

- Applications may request a data transfer

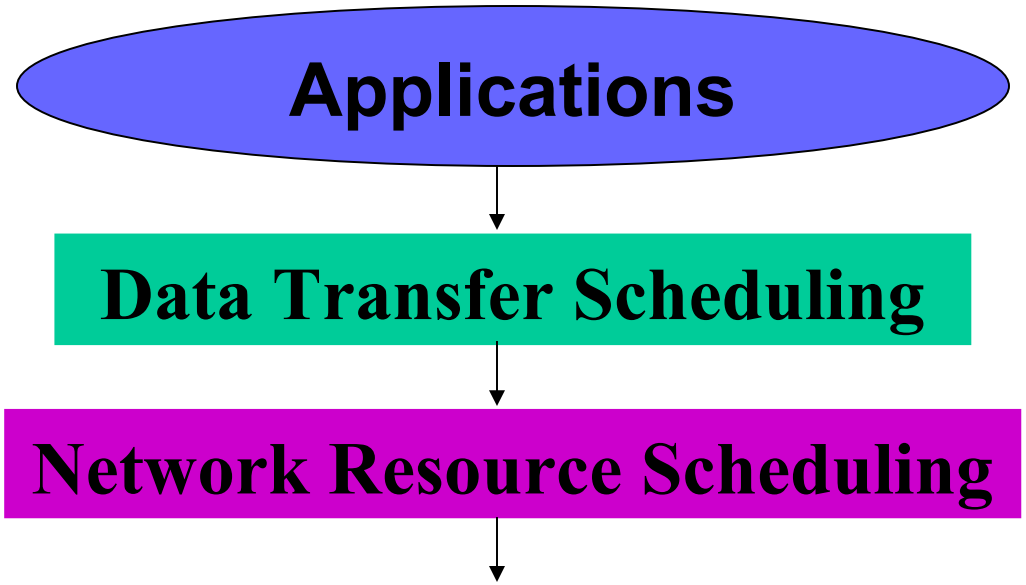

DWDM-RAM Modes

- Applications may request a network connection

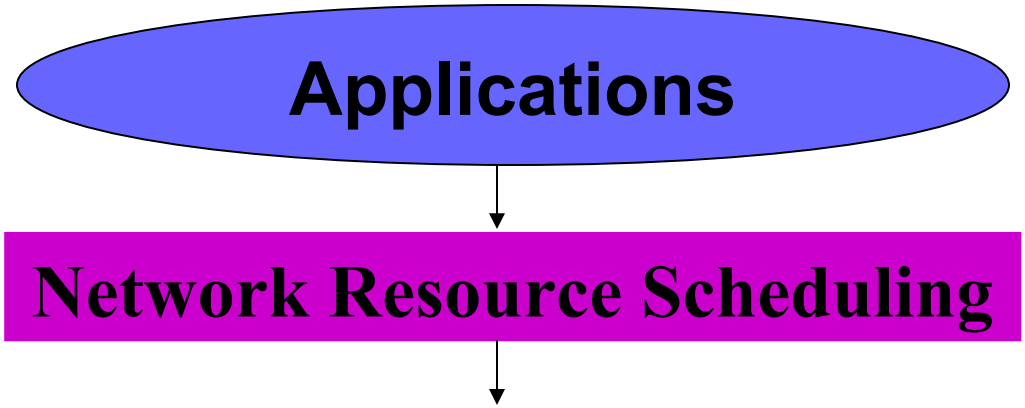

DWDM-RAM Modes

- Applications may request a set of resources through any resource allocator, which will handle the network reservation

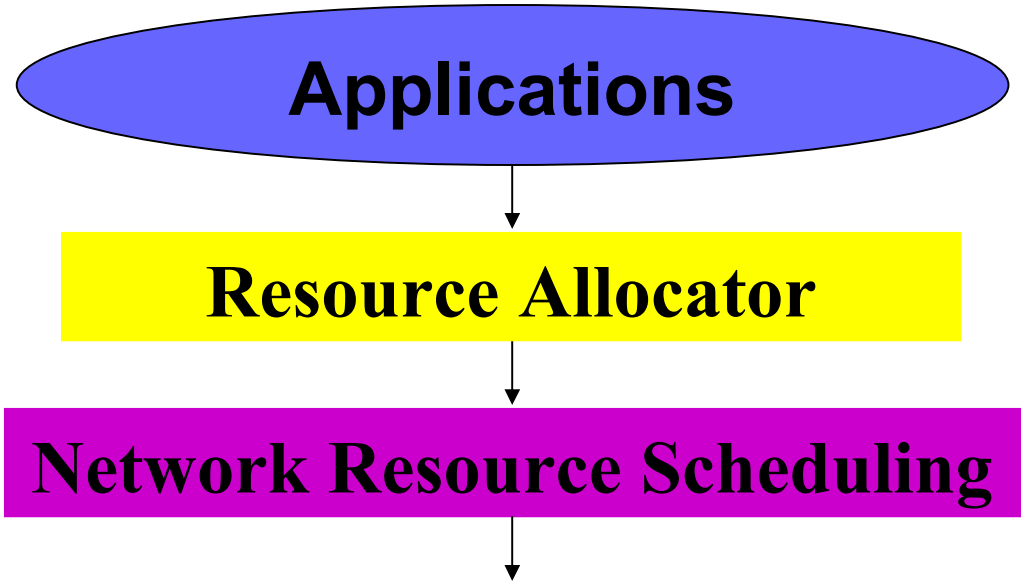

The Network Service

- The NRS is the key for providing network as a resource
  - It is a service with an application-level interface
  - Used for requesting, releasing, and managing the underlying network resources

The Network Service

- NRS
  - Understands the topology of the network
  - Maintains schedules and provisions resources in accordance with the schedule
  - Keeps one scheduling map for each lambda in each segment

The Network Service

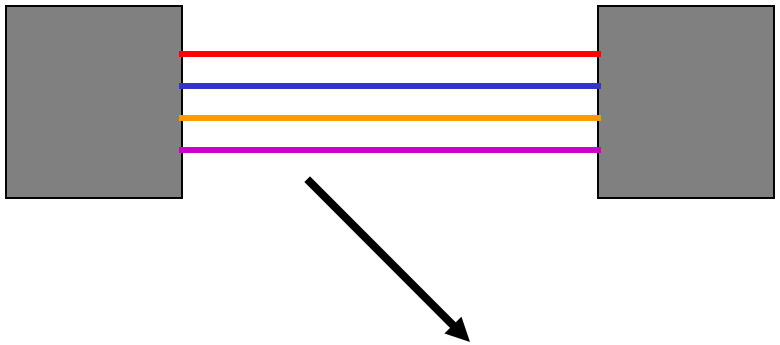

→ 4 Scheduling maps:
Each with a vector of time intervals for keeping the reservations

The Network Service

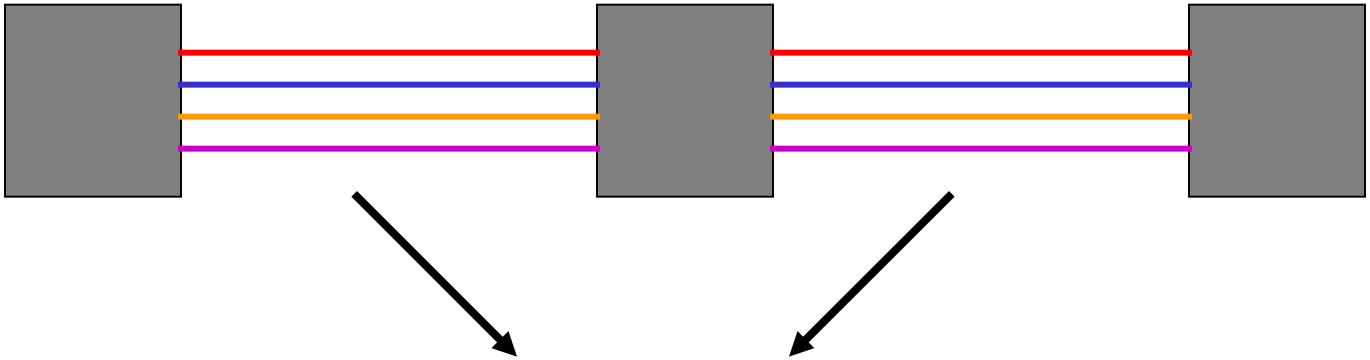

→ 4 scheduling maps for each segment

The Network Service

- NRS
  - Provides an OGSI-based interface to network resources
  - Request parameters
    - Network addresses of the hosts to be connected
    - Window of time for the allocation
    - Duration of the allocation
    - Minimum and maximum acceptable bandwidth (future)

The Network Service

- NRS
  - Provides the network resource
    - On demand
    - By advance reservation
  - Network is requested within a window
    - Constrained
    - Under-constrained

The Network Service

- On Demand
  - Constrained window: right now!
  - Under-constrained window: ASAP!
- Advance Reservation
  - Constrained window
    - Tight window, fits the transference time closely
  - Under-constrained window
    - Large window, fits the transference time loosely
    - Allows flexibility in the scheduling

The Network Service

Under-constrained window

- Request for 1/2 hour between 4:00 and 5:30 on Segment D granted to User W at 4:00
- New request from User X for same segment for 1 hour between 3:30 and 5:00
- Reschedule user W to 4:30; user X to 3:30. Everyone is happy.

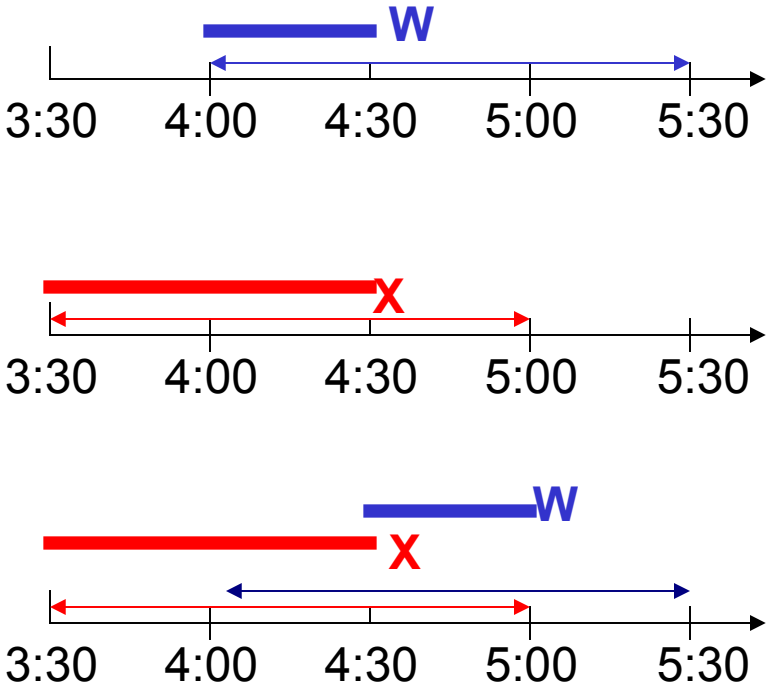

Route allocated for a time slot; new request comes in; 1st route can be rescheduled for a later slot within window to accommodate new request

Experiments

- Experiments have been performed on the OMNInet
  - End-to-end FTP transfer over a 1Gbps link
  - Goal
    - Exercise the network to show that the full bandwidth can be utilized
    - Demonstrate that the path setup time is not significant

End-to-End Transfer Time

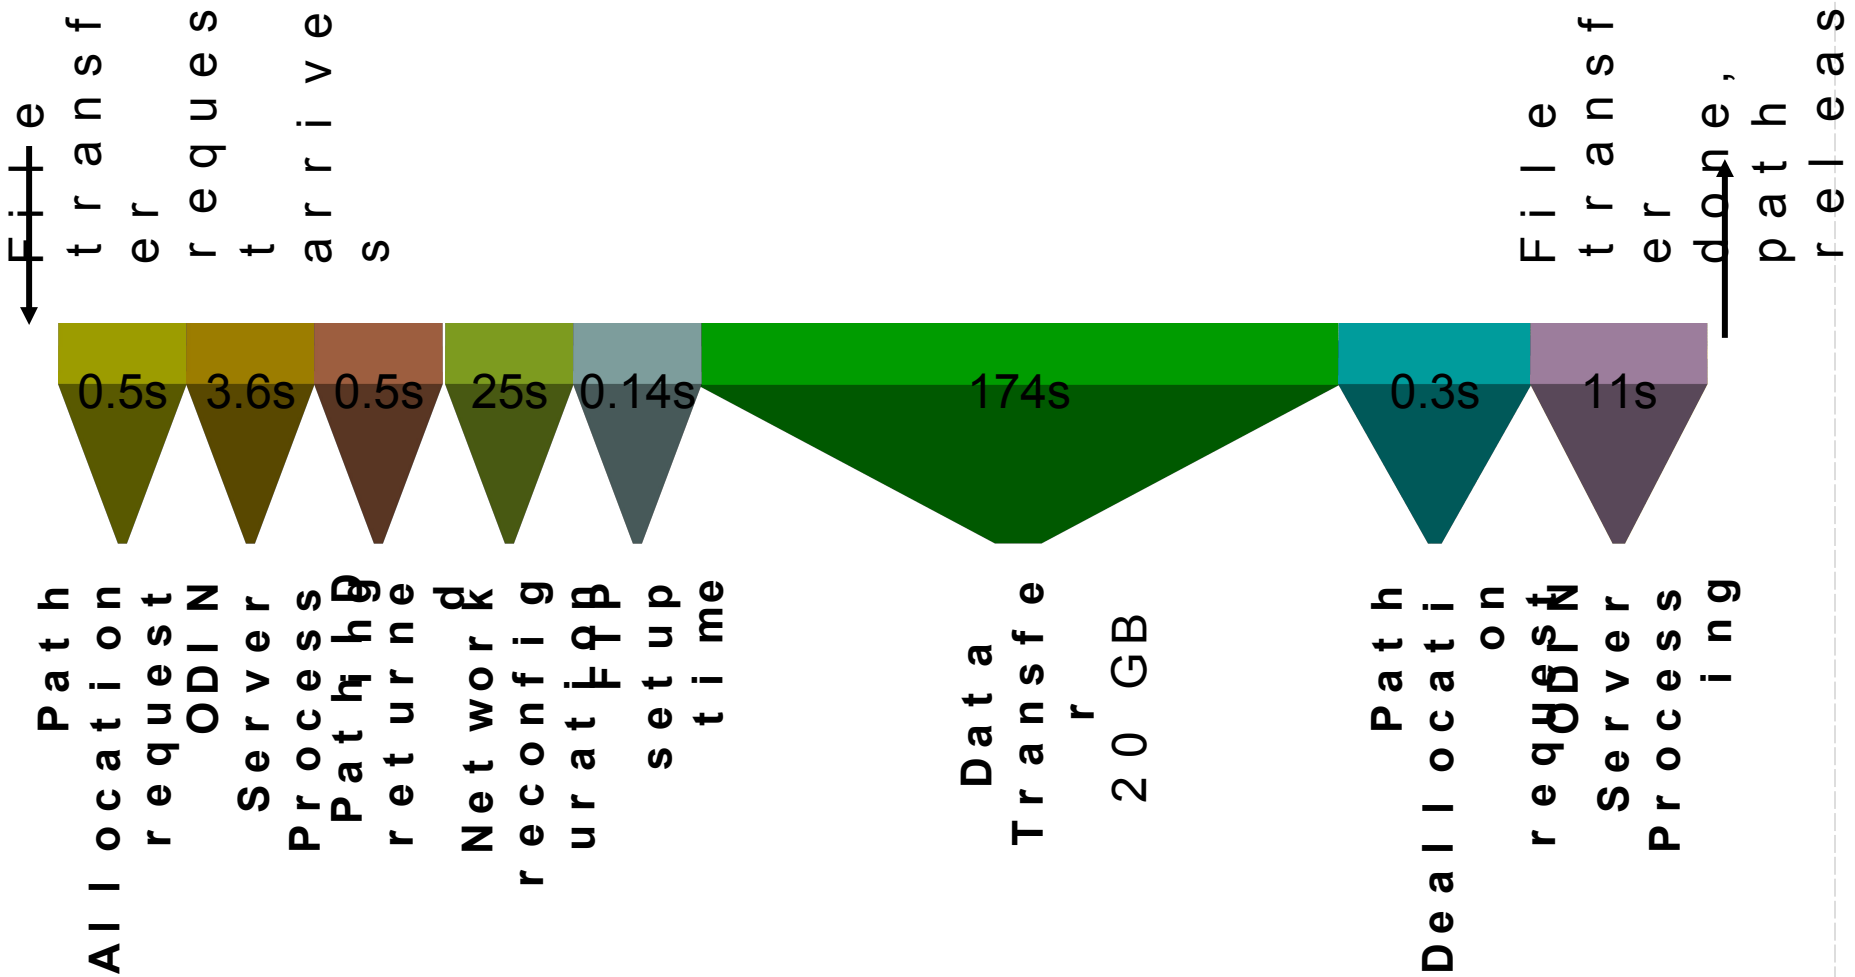

Application Level Measurements

File size:	20 GB
Path allocation:	29.7 secs
Data transfer setup time:	0.141 secs
FTP transfer time:	174 secs
Maximum transfer rate:	935 Mbits/sec
Path tear down time:	11.3 secs
Effective transfer rate:	762 Mbits/sec

20GB File Transfer

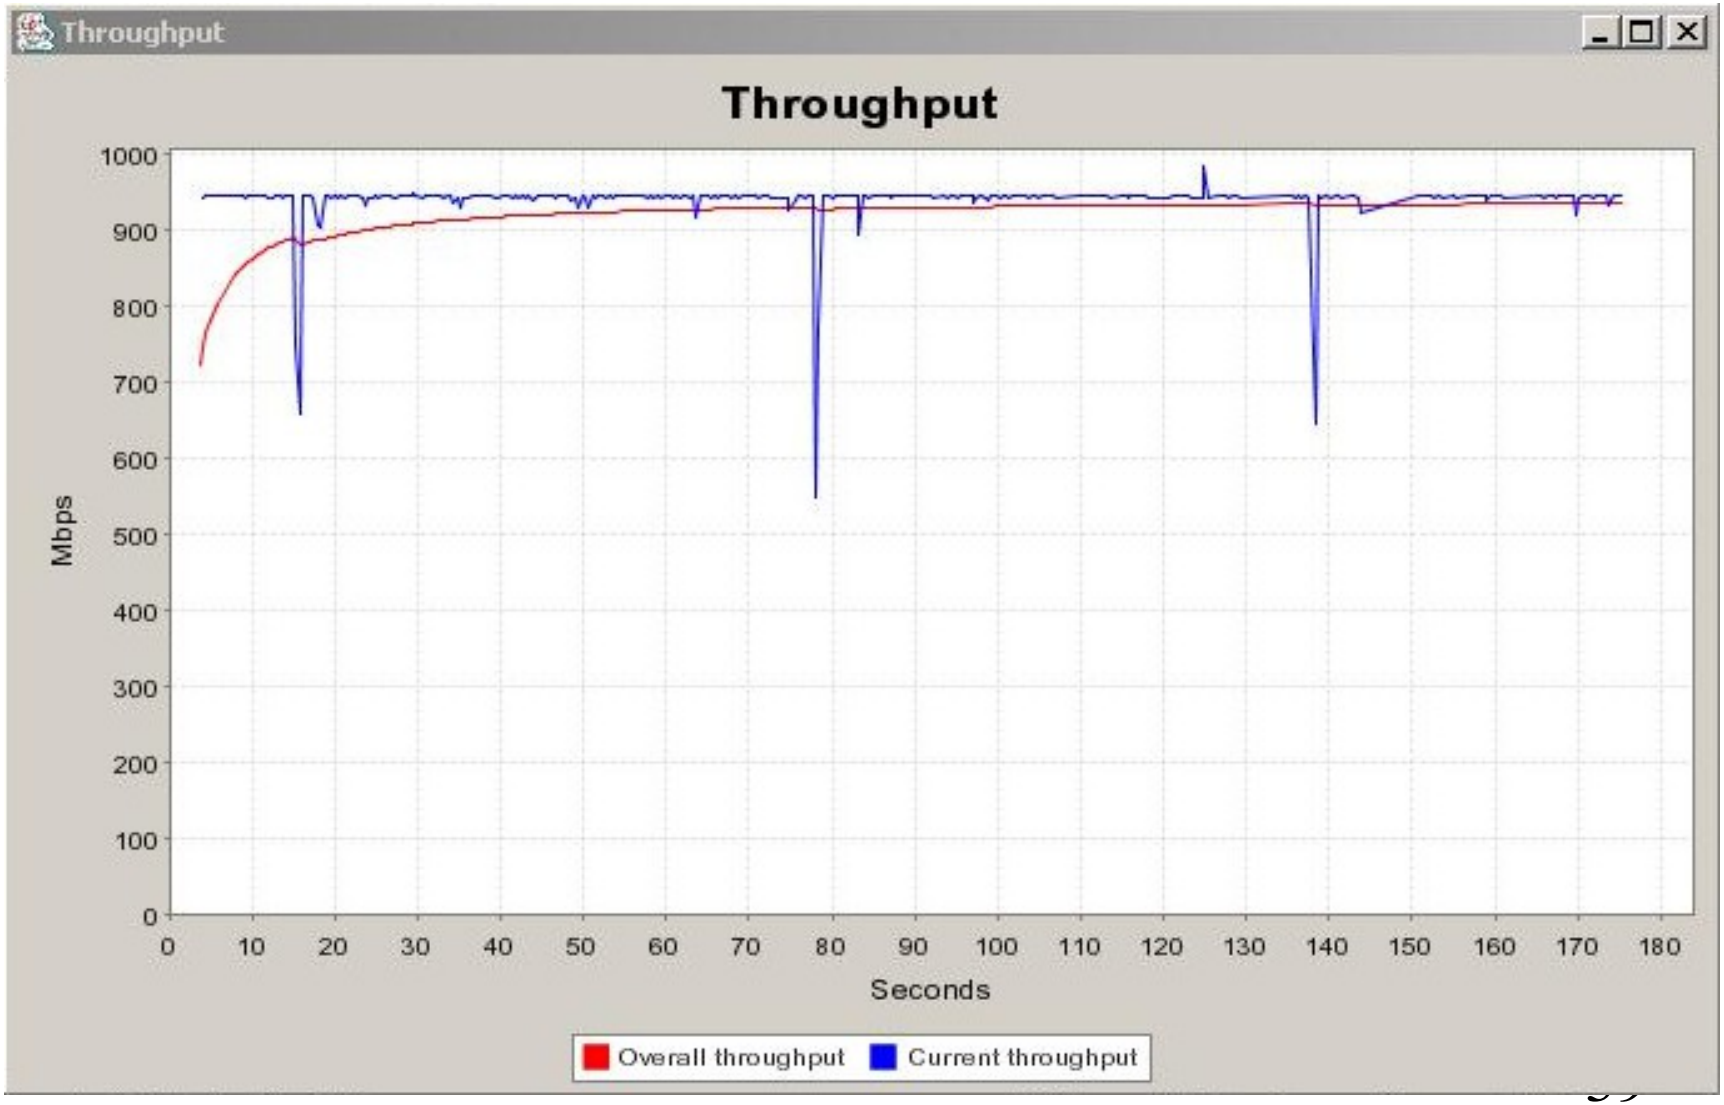

Current Status

- Allocation of one-segment lightpath
  - On demand allocation has been tested at the OMNInet
  - Advance reservation has been implemented but not tested at the OMNInet

Future Work

- Nortel Networks / SURFnet
- Lightpath allocation
  - Multiple-segment lightpaths
  - Optimized allocation when more than one path is available
- Scheduling in large-scale networks
  - Involves different administrative and/or geographic domains
  - Requires a distributed approach

Conclusion

- Dynamic optical network is a key technology for data-intensive grid computing
- DWDM-RAM's network service enables lightpaths to be provided as a primary resource