

# Metropolitan Area Network Support at Fermilab

**Phil DeMar (Fermilab), Chuck Andrews (Fermilab), Andrey Bobyshev (Fermilab), Matt Crawford (Fermilab), Orlando Colon (Fermilab), Steve Fry (Fermilab), Vyto Grigaliunas (Fermilab), Donna Lamore (Fermilab), Don Petravick (Fermilab)**

demar@fnal.gov

**Abstract:** Advances in wide area network service offerings, coupled with comparable developments in local area network technology have enabled many research sites to keep their offsite network bandwidth ahead of demand. For most sites, the more difficult and costly aspect of increasing wide area network capacity is the local loop, which connects the facility LAN to the wide area service provider(s). Fermilab, in coordination with neighboring Argonne National Laboratory, has chosen to provide its own local loop access through leasing of dark fiber to nearby network exchange points, and procuring dense wave division multiplexing (DWDM) equipment to provide data channels across those fibers. Installing and managing such optical network infrastructure has broadened the Laboratory's network support responsibilities to include operating network equipment that is located off-site, and is technically much different than classic LAN network equipment. Effectively, the Laboratory has assumed the role of a local service provider.

This paper will cover Fermilab's experiences with deploying and supporting a Metropolitan Area Network (MAN) infrastructure to satisfy its offsite networking needs. The benefits and drawbacks of providing and supporting such a service will be discussed.

## 1. Introduction

The ESnet Chicago-area metropolitan area network (MAN) is a collaborative effort between Fermilab, Argonne National Laboratory, and the US Department of Energy's Energy Sciences Network (ESnet) to provide a capacious, scalable, high bandwidth network infrastructure in support of the wide area networking needs for the two Laboratories. The MAN is built upon components of existing optical network infrastructure that were originally procured and operated independently by each Laboratory. These components have been supplemented with additional fiber cable and optical network equipment to facilitate a MAN fiber ring that provides redundant data paths for both Laboratories. Each Laboratory provides direct operational support for the underlying network infrastructure that it has contributed to the MAN. However, the MAN must function as a coherent entity, so the two Laboratories must coordinate their individual management responsibilities closely, and even integrate them where necessary. Since ESnet is the provider of wide area network services to both Laboratories, the MAN ends up defining a complex service relationship between the Laboratories and their service provider. Through the MAN, the Laboratories provide high bandwidth data channels to

ESnet. In turn, ESnet uses those channels to provide high bandwidth wide area network services to the Laboratories.

Management of the MAN reflects the complex operational relationship between the two Laboratories, as well as the service relationship between both Laboratories and ESnet. An inter-Laboratory coordination group has been convened, and meets regularly to facilitate integrated operations. The group not only deals with coordinated management and operation of the MAN, but longer term strategic planning as well.

## 2. MAN Architecture, A Layered Approach:

Architecturally, the MAN should be viewed as a three-layer optical network infrastructure.

The bottom layer is the physical layer of the MAN, consisting of fiber and associated DWDM equipment procured & operated by Fermilab & Argonne as general purpose network infrastructure for their respective facilities (Figure 1). This infrastructure provides the base MAN backbone service, but does not provide individual MAN data channels. The Fermilab infrastructure (blue) is an optical fiber connection between the Laboratory and the StarLight international optical network exchange in Chicago. That fiber terminates at both ends in Fermilab-provided Ciena OnLine Metro [1] DWDM equipment. The Argonne infrastructure (green) consists of optical fiber connecting Argonne to the Level-3 exchange point at 600 W Chicago, where a number of major Research & Education (R&E) backbone networks maintain a core presence. Argonne-managed Ciena OnLine Metro chassis are utilized at both locations to terminate the fiber. Argonne provides the fiber connecting the two Chicago points-of-presence (PoPs). A fiber connection between Fermilab (FNAL) and Argonne (ANL) closes the fiber ring, providing physical layer redundancy for the MAN. The majority of that fiber is provided & supported by Argonne, with Fermilab providing a smaller, but significant segment.

Figure 1: MAN Physical Layer

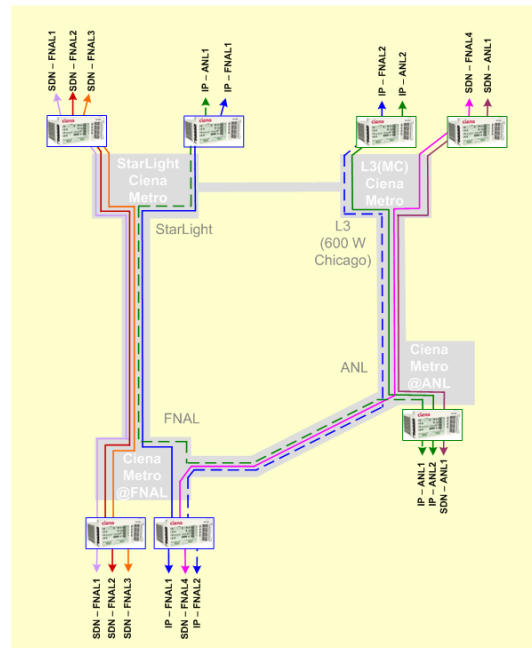
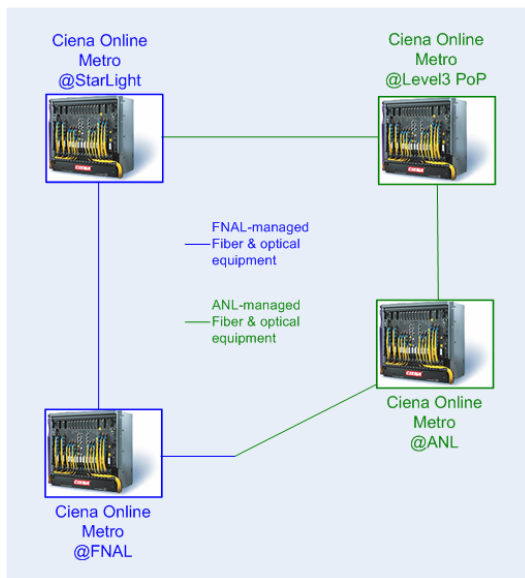


Figure 2: MAN Data Channels

The middle layer (Figure 2) of the optical network infrastructure is the DWDM equipment that provides the individual data channels for the MAN. Ciena 4200 FlexSelect [2] chassis, each capable

of supporting up to four transponder modules, are deployed at Fermilab, Argonne, StarLight, and the Level-3 PoP. The FlexSelect transponder modules support 10Gb/s LAN\_PHY ethernet transceivers. The architecture of the MAN is point-to-point 10 gigabit ethernet (10GE) channels between one of the Laboratories and one of the Chicago network exchange points. Individual 10GE channels are carried on assigned wavelengths, multiplexed across the underlying fiber infrastructure. The architecture is based on ethernet technology, not SONET ring technology. Therefore, SONET-like failover is not provided, and redundancy is left to the failover capabilities of higher level protocols. However, the ethernet-based architecture provides an infrastructure that supports multiple, dedicated 10GE channels at far lower cost than corresponding SONET technology would. In aggregate, the initial MAN FlexSelect deployment provides the 10GE channel configuration listed in the following table:

National Laboratory	StarLight	600 W Chicago	Total
ARGONNE	1	2	3
FERMILAB	4	2	6
MAN Total:			9

The top layer of the MAN is defined by the intelligent network device infrastructure into which the 10Gb/s Ethernet MAN channels terminate. Since the MAN channels are provided to ESnet so that they can provide wide-area network services back to the Laboratories, the intelligent network devices are ESnet MAN switches and routers. ESnet provides two classes of wide area service, a general-use, routed IP service, and a network bandwidth service for end-to-end data circuits. The latter is referred to as the Science Data Network (SDN) service. This service is intended to provide data movement support for large scale science experiments and projects. SDN service is not expected to have the same level of reliability as the production routed IP service.

Under the current MAN configuration, each Laboratory has one dedicated 10GE channel to the ESnet's Level-3 PoP, and a second dedicated 10GE channel to ESnet's StarLight PoP. The IP service channels terminate directly into the ESnet backbone routers at the respective PoPs. On the Laboratory end, those channels terminate in an ESnet MAN switch. The Laboratories are provided 10gigabit ethernet ports on their local ESnet MAN switch for their routed IP service. This configuration provides a very high level of redundancy to each Laboratory for their general internet access. Production IP network traffic to/from either Laboratory will follow the optimal network path from/to the ESnet backbone, effectively providing an aggregate of 20Gb/s bandwidth capacity for each site's routed IP service.

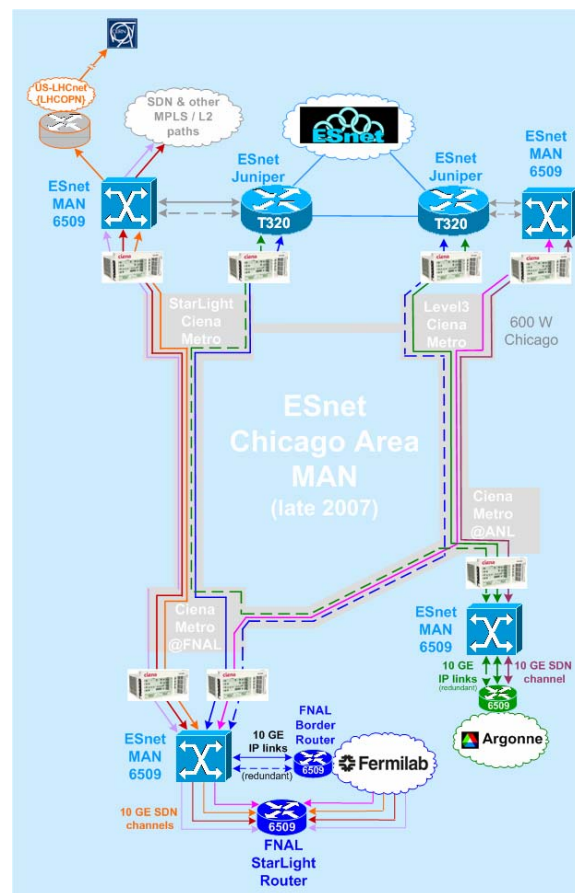


Figure 3: MAN Data Services

The initial SDN channel configuration provides four 10GE SDN channels for Fermilab, and one 10GE SDN channel for Argonne. Three of the Fermilab SDN channels terminate at StarLight, where

connections to other R&E networks for E2E circuit connections are established. The LHCOPN connection to the Tier-0 Center at CERN is one example of such an E2E circuit. All three channels are supported across the Fermilab/StarLight fiber segment. The Argonne SDN channel and the fourth Fermilab SDN channel terminate at the ESnet Level-3 PoP, where other E2E circuit connections are available. Since the fourth Fermilab SDN channel is routed via the fiber path to Argonne, it also serves a redundancy role for Fermilab's other three SDN channels. On the Laboratory end, the SDN channels terminate in the same local ESnet MAN 6509 switches.

### **3. MAN Operations and Support Issues**

Operation and management of the MAN is a cooperative effort. ANL and FNAL have formed a joint operations team responsible for defining specifications and strategies for operations, sparing, maintenance, and management of the MAN infrastructure. The team consists of network support personnel from both Laboratories, and is referred to as the FNAL ANL Network Group (FANG).

The primary focus of FANG has been on developing a common operations and maintenance support model. Key areas of operation that have been:

- Deployment of a common monitoring capability, based on Nagios [3]
- Coordination of on-call procedures & local maintenance operations
- Development of a common triage process for troubleshooting
- Definition of a common problem reporting procedure
- Establishment of common, distributed sparing plan for MAN infrastructure
- Specification & generation of common documentation on the MAN's design & configuration

The other major focus areas of FANG are strategic and tactical planning for MAN expansion and technical evolution of the MAN. Identified areas of concern are:

- Channel allocation and channel resource planning & coordination
- Topology change planning & implementation, such as a change in the local ESnet PoP
- Equipment refresh strategies & plans
- Deployment of new technology within the MAN, such as 40Gb/s channel support
- Review of provisioning agreements

### **4. MAN Technical Features & Capacity**

#### **4.1. Notable Technical Features**

The ESnet Chicago area MAN has been deployed using the latest in metropolitan area optical network equipment and technology. In particular, two features delineate the MAN's data channels from the DWDM channels that were initially deployed.

- The LAN\_PHY 10GE ethernet modules support tunable wavelength optics. This allows the modules to be used interchangeably to support a 10GE data channel across any specific wavelength desired. Sparing of 10GE transponder modules is simplified, and fast turnaround on repair of individual module failures is greatly enhanced.
- The LAN\_PHY 10GE ethernet modules also support G.709 forward error correction (FEC). Burst errors of up to 128 consecutive bytes will be detected and corrected at line rate on the receive end of each data channel.

## 4.2. MAN Capacity

The optical band in use within the MAN's Ciena OnLine Metro system is 1529.55nm – 1563.05nm, with 100GHz (0.8nm) channel spacing. The optical band is divided into eleven sub-bands, each supporting three optical channels (wavelengths). This provides 33 usable data channels between adjacent OnLine Metro chassis on the MAN ring. With the current state of the technology, each data channel is capable of supporting a 10GE connection. In the near future, 40Gb/s transponder modules, capable of supporting 4x10GE data channels, are expected to become available. Thus, aggregate capacity across any particular fiber segment within the MAN ring is currently 330Gb/s, with a potential to go higher in the future. Note that since the underlying optical technology is point-to-point data channels, not protected SONET rings, a given wavelength may actually end up being used for multiple 10GE data channels. Figure 4




Figure 4: MAN Channel Allocation

displays the current band/channel usage on the MAN. Consider channel B7C1, one wavelength within the optical band in use. Argonne (ANL) has a 10GE data channel in use on its direct fiber connection to the Level3 PoP (L3). In addition, Argonne has a separate 10GE data channel to StarLight (SL), traversing the other direction of the MAN ring. That data channel does transit through Fermilab's MAN chassis, and thus is not an available channel for Fermilab. But Argonne is able to utilize the same wavelength for two separate data channels, each traversing different segments of the MAN ring.

## 5. MAN Benefits and Drawbacks

### 5.1. MAN Benefits

The benefits of supporting a metropolitan area network in order to provide local loop data channels include:

- An economically viable means of meeting very high offsite bandwidth capacity requirements. While the cost of implementing a metropolitan optical network infrastructure is considerable, the ongoing costs associated with obtaining comparable service from a commercial provider would likely be prohibitive.
- Flexibility in configuring and changing offsite connections. There is much more agility in modifying equipment under one's direct control, rather than having to modify contractual service agreements with a provider.
- Ability to procure and implement emerging network technologies as they become available. Service providers enhance their service offerings only as market demand and internal strategic planning dictate. When 40Gb/s data channels become available and economically viable, we can implement them, without waiting for a comparable commercial service.
- Deterministic expansion capability. The incremental cost for another 10GE channel is well understood, and relatively modest. The time frame for implementation is also well defined.

## 5.2. MAN Drawbacks

There are comparable drawbacks to supporting a metropolitan area network. They include:

- A high initial deployment cost, including not only capital expenditure but personnel effort as well. In the case of the Chicago area MAN, the project was centrally-funded out of the US Department of Energy's Office of Science, after the longer term economic benefits of the project were quantified.
- Significant additional management effort. This effort is compounded by the fact that much of the equipment is remotely situated, and requires additional effort to physically access
- New technology skill sets need to be developed. Optical network technology skills are distinctly different than classical local or wide area network skills. These skill sets will most likely need to be developed in house, which will generate its own learning curve and bumps along the way.

## 5.3. Benefits versus Drawbacks

On the whole, the MAN deployment and operation experience has been a positive one. It has provided both Fermilab and Argonne with a very capacious and very flexible network infrastructure to support current requirements. The MAN also serves as a technology enabler, capable of supporting new advances in the network research environment for the Laboratories.

## References

- [1] [http://www.ciena.com/products/products\\_3369.htm](http://www.ciena.com/products/products_3369.htm)
- [2] [http://www.ciena.com/products/products\\_cn4200\\_overview.htm](http://www.ciena.com/products/products_cn4200_overview.htm)
- [3] <http://www.nagios.org/>