JR RICHARDSON
Engineering for the Masses
Jr.richardson@cox.com

Leveraging Embedded Fiber Optic Infrastructures for Scalable Broadband Services

Preface

When telecommunications companies sell broadband access, existing methods and infrastructure are leveraged whenever possible. A component becoming outdated often requires a complete overhaul of the system. Sometimes, the racks are the only thing salvageable. This paper outlines a scalable architecture capable of delivering broadband services while maximizing revenue from embedded fiber optic infrastructures.

Many approaches to building fiber optic networks accommodate logical and physical protected network architectures. Architecture decisions are driven by capital budgets, management processes, customer base, service delivery, product set, operational expediency and risk analysis. Sparsely populated rural communities can interconnect by long, aerial fibers run down a poll line. These aerial fibers are rarely cut down by an off-course motorist. In some instances, trenching along a railway can be a great investment. Leasing space on poles or conduits owned by regional utility companies presents unique challenges. Densely populated urban environments can be the most costly per mile where underground boring is required.

As you can see, environmental, financial and logistical issues are all factors in building fiber optic networks. In building the physical fiber plant upfront and leveraging a scalable architecture, a business endeavors to gain maximum revenue now and in the future.

The scalable network model in this paper is designed to accommodate requirements for a large customer network build. This network model could very easily be adaptable to many different networks.

Global System requirements

The customer's network system resides within a 40 mile metropolitan area. The wide area connectivity network will be referred to as a Metropolitan Area Network (MAN). The magnitude of this endeavor presents itself through pure volume of locations, being 140 sites converging back to one data center or central location. The geographic density of the communication company's embedded fiber optic network enabled the company to build fiber into every site location with low construction cost.

The customer's network is deliberately overbuilt in terms of network capacity. Maintenance challenges are eased by centralized network services, hosted in the datacenter. Rather than to overbuild standard 0-3 year capacity, a 0-10 year capacity model is selected. The network services include Voice, Data and Video. The desired access media of the network at each site is Gigabit Ethernet. Each site is allotted a committed information rate (CIR) back to the data center of 100 Megabits. Any two sites, anywhere in the MAN, require the ability to burst and fully utilize 1 Gbps throughput to the data center. This forces the network distribution and core to provide additional capacity beyond the cumulative CIR of the edge sites.

Network Capacity Requirements

To accommodate the dedicated traffic of 100 Meg CIR per site, the MAN-to-LAN interface at the data center must be a minimum of 14 Gigabit. Enabling the ability for two locations to burst to full Gig speeds added another 1.8 Gig to the MAN-to-LAN interface. 15.8 Gig is the current requirement to provide non-blocking capacity to and from the data center for all 140 sites.

Data Center Function

The data center will host all services, Storage Area Network (SAN), Voice over IP (VoIP) Call Center, Internet Gateway, Video Gateway, Network Domain services, Email, Distance Learning Academy, Video Conferencing and a host of new applications.

Future Network Capacity Requirements

The network must have the ability to scale to a 1gbps CIR for each site. This equates to 140 gig backbone non-blocking into the data center. Ultimate scalability is needed in the original network design to accommodate this capacity without the need to build more fiber or change out any CPE devices.

Core Network Backbone Technology

The core backbone is a 2 fiber pair ring (4 fibers east and west) with hub sites residing on the main ring. The Access rings extend out from the hubs to pick up each site. See Figure 1.

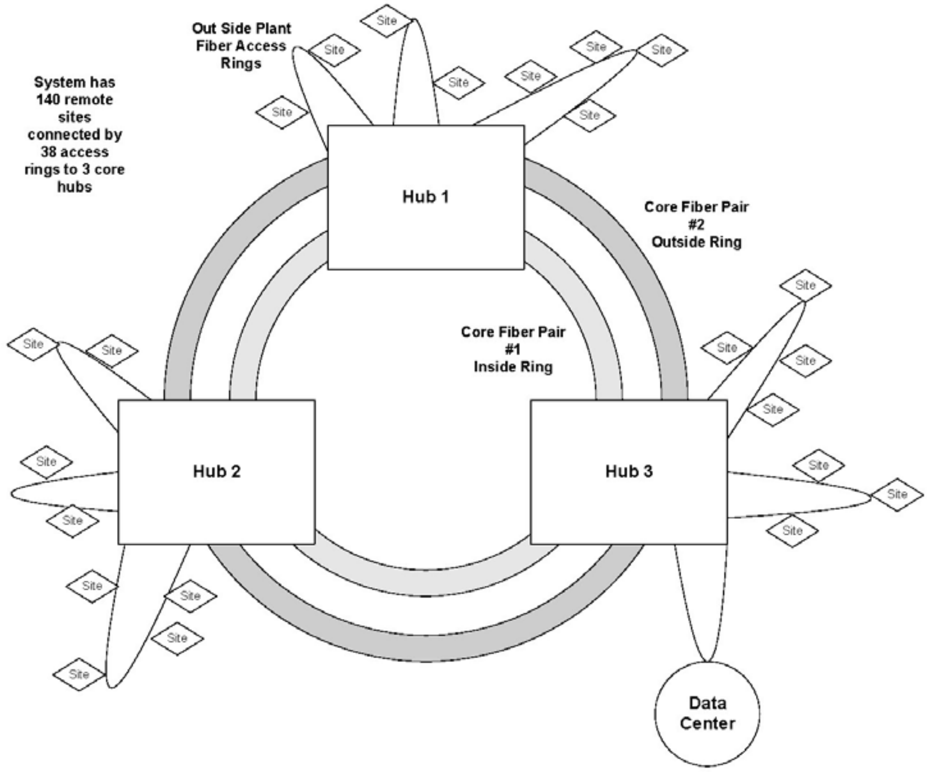

Figure 1

Ultimate scalability on the core ring is accomplished using Dense Wave Division Multiplexing (DWDM) technology.

Currently these types of systems are allowing for 40-10 Gig lambda availability at 100 gig spacing between lambdas. 50 gig spacing is also available but at much higher cost. Depending on the level of service needed, the DWDM system can be active, passive, protected, transparent, integrated, dynamic, incorporate optical switching, routing, ring, mesh, managed, unmanaged and a host of other features. Greater capital investment will inevitably bring greater flexibility but also increase operational complexity. All Venders considered for this particular project offer scalable options so it's not necessary to light-up all lambdas at once. This helps keep capital cost focused on current requirements but the core model will ultimately scale to accommodate 400 gig of bandwidth on a single fiber pair.

Access Ring Topology

The embedded fiber optic access ring network is laid out in a ring fashion extending out from a hub facility. Each access ring is a fiber pair going east and west from the hub. The sites are picked up in a daisy-chain fashion and the ring is terminated into a parent node at the hub. The physical out side plant picked up 2 to 7 sites per ring. 38 access rings were needed to bring in all 140 sites.

For architectural explanation, I will focus on an access ring with 7 schools. See Figure 2.

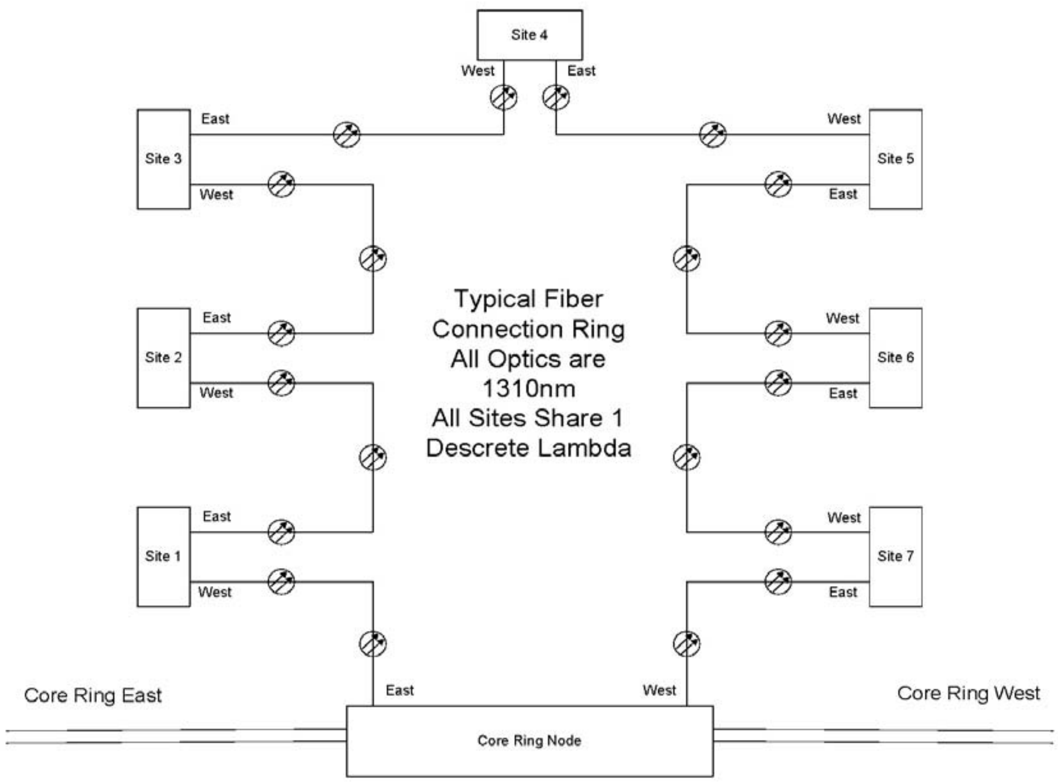

Figure 2

Starting at the hub location, a pair of fibers head east to the first site where it then terminates into the west end of a protected optical add/drop service unit. Out of the east end of the unit at site 1, the fiber continues on to site two and terminates in the same fashion. The ring continues on to site three, four, five, six and seven. The ring comes back into the west end of the hub completing a seven node ring protected add/drop sub-network.

All optics in the access ring will initially use the same wavelength 1310nm. All sites have relatively short interconnected distances but intermediate range optics should be used initially. Light levels should be padded when necessary to prevent optics from being damaged.

Initial Access Ring Service Level

The initial sub network or access ring supplies the contractual service levels within a 2.5 Gig or OC-48 level capacity, given a 100 Meg CIR to each location. This adds up to 700 Meg. If 2 sites burst up to 1 Gig at the same instance, this adds a further 1.8 Gig to the network load. This consumes the remainder of the access ring bandwidth while maintaining the dedicated 100 Meg CIR for the remaining 5 sites. See Figure 3.

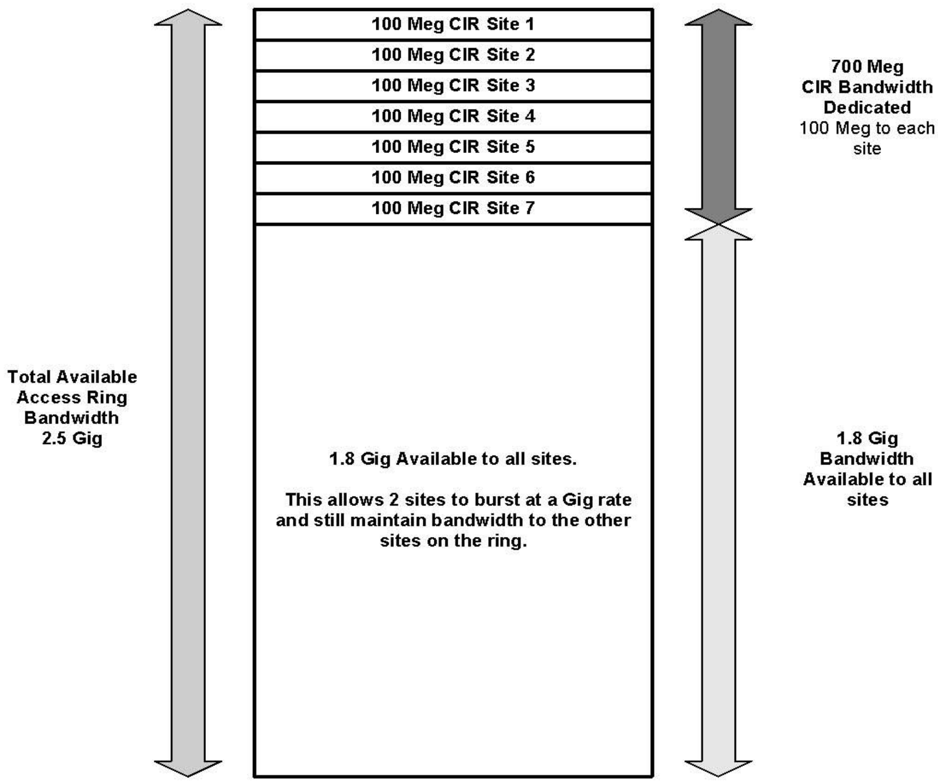

Figure 3

Access Capacity Augmentation

Traffic analysis provides network utilization statistics and predicates the cost of expansion. At what utilization level and upgrade is needed should be defined within the service level agreement to the customer. When that time comes, if the access ring requires a 10 gig or OC-192 level capacity then all optics on the access ring must be upgraded, provided the CPE can operate at those speeds.

Considerable capital expense may initially prevent embedding the network with 10Gig optics. Using lambda manipulation within a scalable design allows the use of the same optics and still upgrade where needed. The costs of optical networking components drop about 20% a year. The longer the wait is to upgrade, less expense will be incurred.

Add and Drop a Lambda

A lambda is a discrete wavelength or channel within an envelope of multiple wavelengths or bands, operating on a single fiber. Each lambda is capable of carrying up to 10 Gigabits per second. It is possible to add and drop lambdas as you need to on a specific fiber optic span. This can be accomplished actively with an optical transponder or passively with optical separating and combining filters already tuned to a discrete wavelength. See Figure 4.

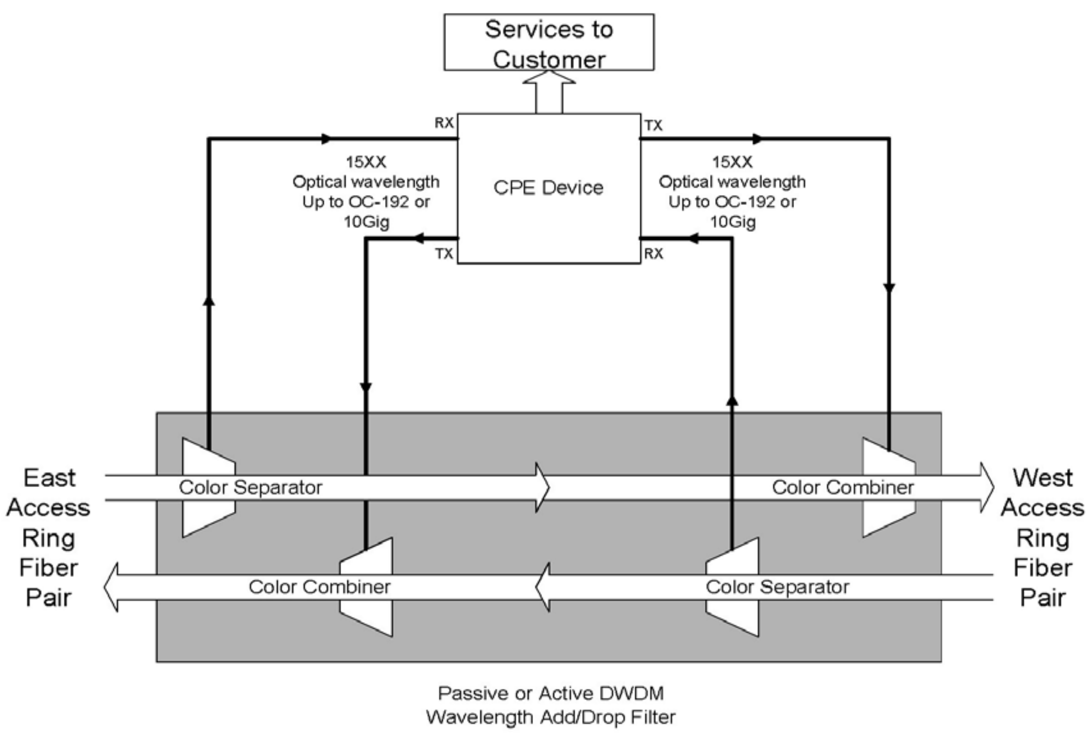

Figure 4

End-to-End Scale and Compatibility

Considerations of future growth must include end-to-end architecture. The critical compatibility is between the core and the edge sub-networks. The core network in this model is a DWDM system operating in the 15XX ITU grid spectrum. It makes sense to operate the sub-network access rings, initially, on a 1310 wavelength to eliminate interference down the road.

CPE Consideration

The CPE capabilities have a lot to do with the architectural decisions; can the CPE change its optics? If not, an active transponder will be needed at the end sites. Cost is driven upward. Modular optics on the CPE is a big plus for flexibility to grow and expand the initial design. This usually has some higher deployment cost but measured against active transponder units at all CPE sites, it's a good investment. Modular optics was chosen for this particular project, allowing the use of passive filters.

Increasing Access Ring Capacity

A ring upgrade is performed without service interruption. The service equipment in the network has built-in protection with ring traffic capable of switching east and west. This capability is required during a ring upgrade.

All sites on the access ring, including the hub equipment terminating the ring, require a passive filter set. Pre-consideration of DB loss in the optical budgets must be calculated to determine if the ring can handle the added filters. The CPE on the access rings were initially deployed with intermediate reach optics and light levels padded as needed. Once all the filters on an access ring are installed, it will be necessary to remove some padding as light loss will incur from connections on the filter sets.

The ring upgrade process begins at a hub site, and proceeds to the next site on the ring and so on. Proceeding east or west is irrelevant. At the hub, active traffic must flow in the opposite direction of the filter set installation path at the sites. Once the filter is in place at a site, traffic for that particular site must be switched to the opposite direction. While the filter sets on the access ring are upgraded, the ring will operate in a failed (ring-cut) mode. An actual fiber cut in the out side plant will eliminate the ability to upgrade without service interruption. At the individual or multiple sites that will be using a new 15XX lambda, a discrete filter set for that particular wavelength must be utilized. The hub site will also need to filter in and out the original 1310 wavelengths plus the new 15XX lambda. See Figure 5.

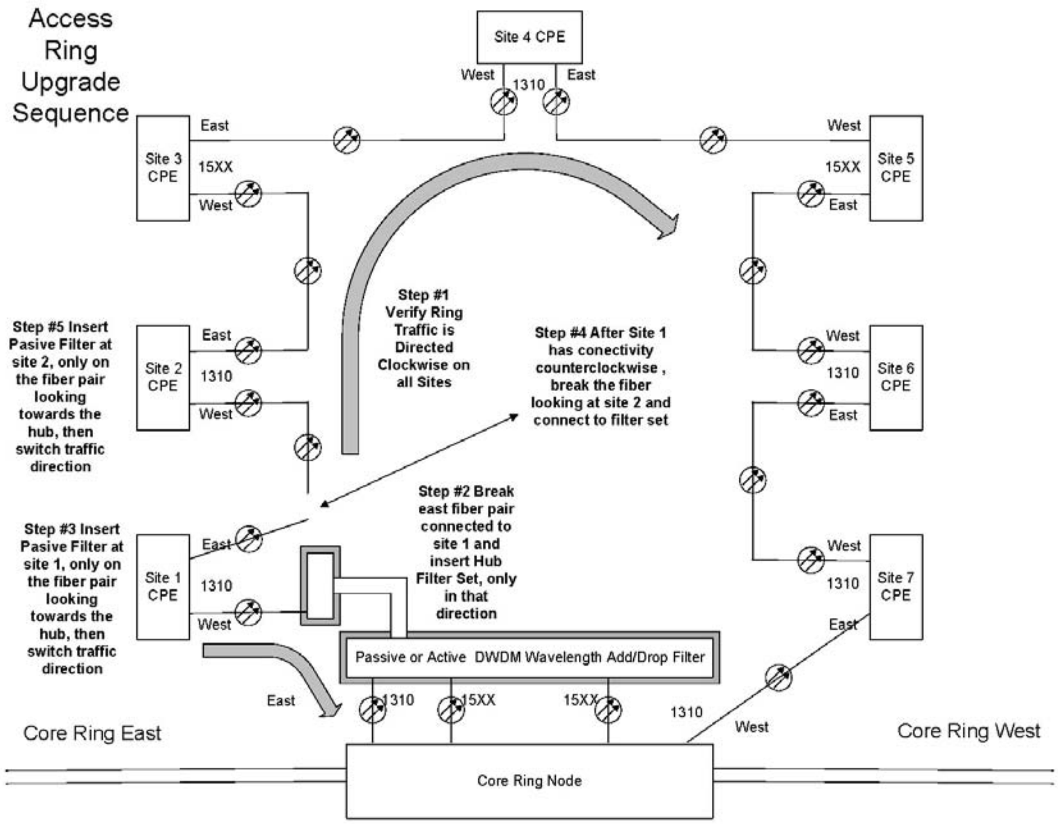

Figure 5

The CPE's ability to modify its own optics, determines the need for an active transponder at the upgrade sites and the hub. Calculations on the optical budgets must be performed to ensure the correct strength optics are in place for the new 15XX wavelengths and also for the 1310 wavelength sites that will be jumping spans (Optical path between two active components, transmitter and receivers). See Figure 6.

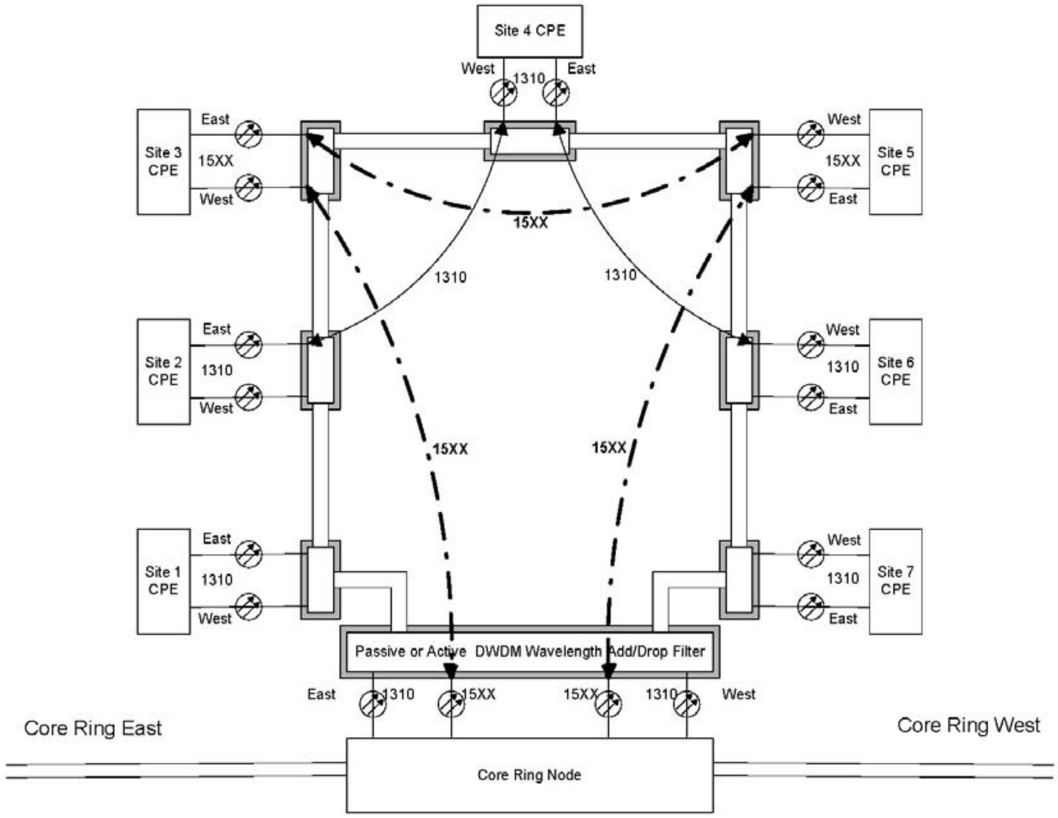

Figure 6

Conclusion

The initial architecture allows the access ring sub-network to grow and incur upgrade cost incrementally, predicated on capacity needs. The upgrade path uses filter sets and transponder equipment only where necessary, keeping expansion cost minimal. Hence, the original CPE and network equipment remain the same for contractual terms and beyond. In this manner an architectural or equipment forklift is prevented.