NThF3.pdf

AT&Ts Photonic Backbone Design Options

Kathleen A Tse
AT&T Labs, Room C1-3D11, 200 Laurel Ave South, Middletown, NJ 07748, USA,
Telephone: 732-420-2612 katse@.att.com

Abstract:. As AT&T plans for the future, there are a number of critical decisions required to define a new 100G backbone. These decisions are driven by specific carrier characteristics, which are described in this paper.

OCIS codes: (060.2330) Fiber optics communications

AT&T's Coast to Coast Photonic network provides high bandwidth connectivity to a large variety of customers, both internal and external. Today's 40G backbone is designed for lowest cost, through a combination of high capacity express routes between large cities and regional feeder routes that connect in the many smaller locations where service is needed. With both of those layers fully built out and growing today, AT&T is focused on the next generation of capabilities, including even higher bandwidth connections (100G and above), service connectivity and restoration, and continuously lowering costs by improving fill and leveraging technology. There are a number of choices that need to be made by any Service Provider selecting their next generation backbone, and those choices can vary based on the types of services and customers that are served, as well as the rate of growth and scale that is needed.

Overbuild or Leverage Deployed Base

The decision on when to light up a new fiber is one that is very dependent on the carrier and its needs. If fiber is a scarce commodity, it can be more economical to take the performance hit and co-propagate the higher bandwidth waves with the current traffic. If traffic is growing quickly and fiber is available, it is most likely better to optimize performance and overbuild with a system that is optimized for the new rate and has a longer life at low cost.

This decision is particularly important for coherent 100G without DCM. The penalties of co-propagating with 10G on-off keyed signals, and the added penalty of the DCM to the coherent transmission can both cause a large hit to engineering rules [3]. Adding guard bands to avoid adjacency adds complexity to a mesh network, where now it is necessary to keep aware of restrictions when setting up a new connection. In most cases when fiber is available, and the fiber can support the new technology, then it makes economic sense to start with a new 100G system rather than limp along with a mix of old and new technology, even if it can be made to work.

Integrating Long Reach Optics

Many router and switch suppliers show great advantages that can be had by integrating the long reach optics directly in the router/switch. In order to do this integration, the DWDM system needs to be from the same supplier as the switch or else you need a DWDM that is capable of accepting foreign wavelengths from another source.

The main advantage of this integrated optics solution is to avoid the cost of the two short reach interfaces that are used cross office. By avoiding those two short-reach clients, it is possible to save capital, especially at the beginning-of-life for new bit rate clients that have not yet been produced at volume. Over time as the cost comes down there is still a savings, but that typically is quite small compared to the cost of the circuit. Another potential advantage of integration is the control of the wavelength by the router or switch. The tuning of the wavelength and the detection of errors can now be done directly by the router/switch, and in future applications could be used for innovative routing and signaling without the need for a multi-vendor control plane [1].

There are a number of disadvantages of this solution, however. If the same vendor is used for both the switch and the DWDM system (or even just the costliest part of the DWDM, the transponder), then you are reliant on a single company to be a leader in both technologies in order to remain competitive with costs and features. When new capabilities or innovative technologies emerge, other suppliers may be first to market and integration may be difficult. Pricing leverage becomes difficult, because there is no easy way to separate the functions and let multiple suppliers compete on a level playing field. And for foreign optics, most DWDM suppliers have less generous engineering rules when the wavelength is coming from a competitor than from their own system. Other challenges for integration of the optics include the need for regeneration along the route, the maintenance information from different systems and often different monitoring groups in a carrier, the responsibility for troubles at a high speed optical meet and the challenges of testing on the high-speed side of a signal with complex modulation if troubles occur.

Some carriers today are successfully using integrated optics. AT&T has used integrated optics on some of our very first DWDM platforms. In most cases, integrated optics are chosen because the supplier serves only one type of customer, for example all of the wavelengths originate in a single type of router or switch. When multiple types of client equipment exist, then the integrated optics can only be used on some of the wavelengths, where others need to be served by a traditional transponder with short reach clients. Once it is necessary to mix and match, the complexity of the system becomes higher and the driver for the integration is less interesting.

Another challenge of integrated optics comes at the time of network retirement. In AT&T's network, the lifetime of the terminal equipment such as switches and routers has proven shorter than that of the wavelengths and DWDMs. In this case, it is not possible to easily reuse the wavelengths on a DWDM for the new router/switch without purchasing additional costly transponders.

For all of these reasons, a carrier needs to carefully consider the value equation for integrated optics at every layer of the network. Integrated optics makes the most sense for places where all of the wavelengths come from a single source such as a router, where the distances are relatively short not requiring regeneration or suffering from reduced engineering rules and where the lifetime of the terminating equipment is expected to be extremely long. Also for short wavelengths, the cost of the client plays a larger role of the cost of the service than in the long distance services. The choice of supplier is also very important, they need to have the expertise to be a technology and cost leader both in the switch/router and on the DWDM space.

On AT&T's backbone, given the multiple types of clients ranging from routers to important private line customers and traffic over our mesh-restorable CNI (Competitive Network Infrastructure) and the long distances required between large cities, saving short reach interfaces on some of the wavelengths may not be worth the complexity of optics integration. This is a question that we continually reevaluate for every new system and every piece of our network, and we look to see if in the future there are new router capabilities or cost points that change the equation.

Inclusion of Packet or Grooming Functionality

Especially as wavelength rates grow to 100G and service rates remain at the lower rates, there must be a grooming functionality at large nodes on the backbone to properly pack the wavelengths and provide efficient capacity [2]. Some suppliers are developing products that integrate the grooming function including the potential for packet fabrics directly onto the DWDM system itself. This provides a great deal of flexibility and control, similar to integrating the high speed photonics in the switch/router from the earlier discussion. An added benefit is now control of the ROADM cross-connect, so that both the grooming switch and the ROADM could be used to set up and direct traffic between clients [4].

This integration of packet or grooming works very well in areas of the network that are single vendor, have relatively small growth of groomed traffic (the size of these switches are still rather small), and where the life of the packet fabric is long compared to the cost of the system. For large multi-vendor/multi-service networks, it can make more sense to have the grooming function distinct from the DWDM system, similar to the arguments made for not integrating optics into the router/switch.

Conclusion

As large, multi-service carriers such as AT&T make decisions on the design of the next photonic backbone, they can be quite different than those of carriers that are smaller in service volume or network reach or that carry exclusively a single type of service on their backbone. Each of these decisions is based on the cost of doing business, the turnover rate of technology and the feature advantages of integration. As AT&T plans the next generation backbone, it builds in the capabilities needed to support large scale, large footprint and multi-service functionality to keep the costs at the lowest possible point and serve the largest customer base reliably.

References

[1] O. Gerstel, "IP and Optical Integration in Dynamic Networks," OFC/NFOEC 2009, paper OML7, March 22-26 2009, San Diego, CA

[2] L. Ostar, "Sub-Wavelength Traffic on an Optical Network," OFC/NFOEC 2009, paper NThE5, March 22-26 2009, San Diego, CA

[3] M. Taylor "Coherent Detection Method using DSP for Demodulation of Signal and Subsequent Equalization of Propagation Impairments", p 674-676, IEEE Photonics Technology Letters, vol. 16(2), February 2004

[4] S. Elby, "Transport Evolution," OFC/NFOEC 2009, Service Provider Summit keynote, March 22-26 2009, San Diego, CA