



LIGHTWAVE®

100G metro transformation

By MIKE SABELHAUS

Metro networks have
unique requirements.
Now 100G technology has
begun to catch up
with them.

FOR MANY YEARS, most metro and backbone optical transport networks have been based on WDM architectures, since the ability to carry multiple wavelengths over a single fiber provides significant savings versus single-channel networks or overlaying multiple networks for each service offering. However, WDM systems operating at 10 Gbps per wavelength have quickly reached capacity on heavily used routes.

Carriers therefore need a cost-effective way to expand capacity, while at the same time reducing the cost per bit transported. To provide immediate relief, vendors have developed 40-Gbps transponders and muxponders that can be used with existing WDM networks. Yet 40 Gbps isn't enough; ever-increasing data growth requires deployment of even higher-capacity WDM systems. That means a need for more channels operating at higher data rates – in other words, 100 Gbps per wavelength, not only for the long haul, but for metro networks as well.

The emerging 100G metro

Carriers have deployed 100G in long-haul applications for the last few years because of the technology's ubiquity, modulation format commonality, and downward price migration. So much so, in fact, that the per-10G pricing of a 10×10G muxponder is competitive with the standard 10G being deployed in long-haul networks today.

As the potential for metro/regional 100G applications has surfaced, the industry has been tasked to further reduce costs. But is the timing right for 100G metro now, given that we have not completely traveled down the metro cost reduction curve? The answer is yes. While 10G is still number one in volume shipments, 40G appears to have peaked; 100G is where the provider community is now concentrating.

One could argue that current 100G systems (i.e., those based on coherent transmission technology) are still too expensive for metro applications, or that the economics of long haul do not translate to some metro, data-center, or other

MIKE SABELHAUS is principal optical solutions architect for [Fujitsu Network Communications Inc.](#) He focuses on next generation optical strategies and markets, guiding the company's technology development to meet the needs of communications service providers.






FIGURE 1. The growth of 100G deployments in metro networks is expected to ramp quickly in the next few years.

SOURCE: OVUM

nontraditional carrier networks. But standardization of future 100G optics and new component types promise to reduce costs on a more aggressive curve than we have seen previously.

Growth in 100G deployments

By surveying the network segments, it's apparent that backbone deployment has taken the traditional first lead for 100G optical transport technologies, which is not surprising. But over the two-plus years that 100G has been available, metro demand has surged and will match long haul in 2014 (see Figure 1).

A good reason for that is the multiplicity of components with a common heritage. The Optical Internetworking Forum's 100G specification work was well-supported and adopted across multiple vendors.

In addition, components as well as multisource agreements (MSAs) came online in record time. Second generation MSAs with lower power consumption, as well as C Form-factor Pluggable (CFP) optical transceivers, will also spread into the metro market and possibly back into long haul.

The lesson here is that vendor interworking, standardized specifications, and common technology will offer interoperability, lower cost, and ultimately, new applications – such as metro networks.

Overview of 100G metro equipment drivers

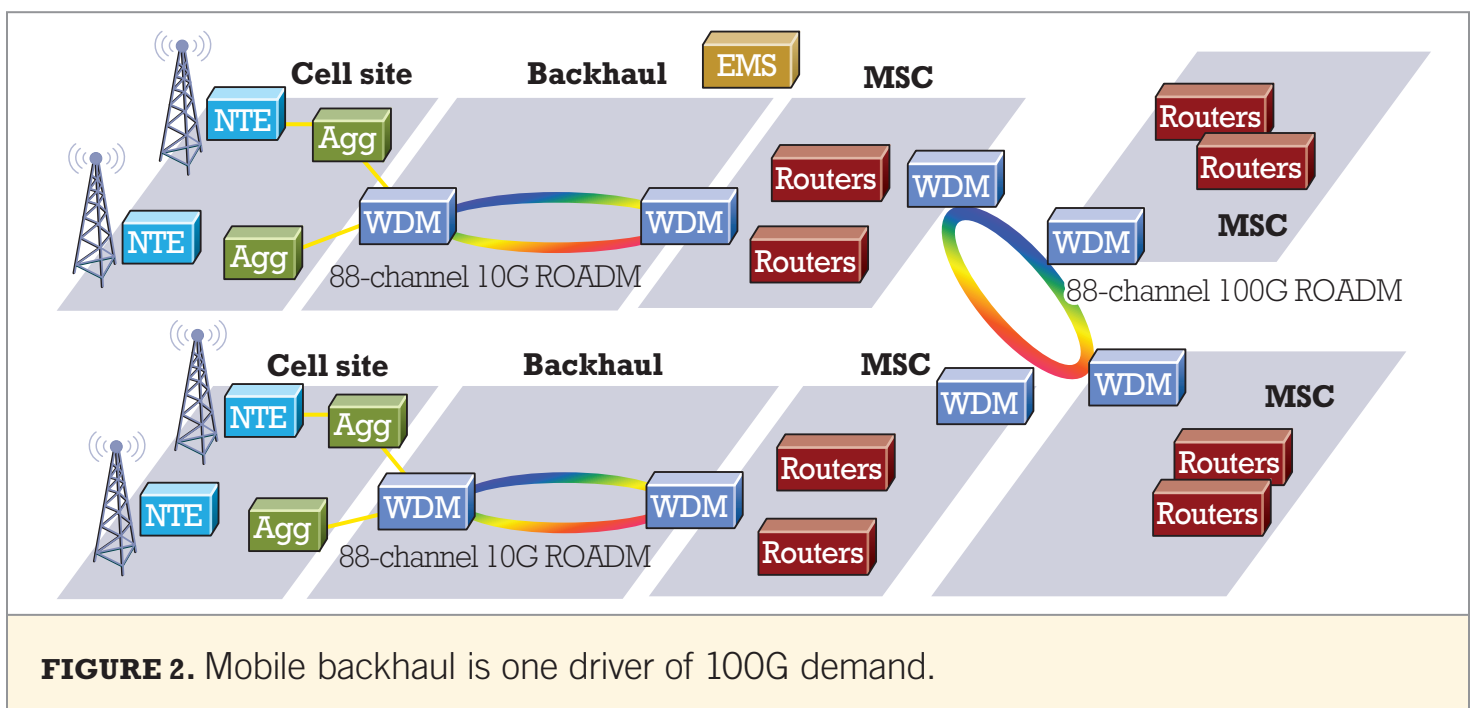
Few would dispute that demand for Internet access and streaming video are fueled by the likes of Amazon and Netflix. YouTube and Netflix combined

now account for over 50% of downstream traffic, according to one fixed-access report. The data center, with its servers, storage arrays, and telecom equipment, has become the hub for these operations. We see content providers, traditional telecommunications carriers, IT hardware/software players, and hosting-based companies all owning and operating data centers. As the number of data centers in suburban areas grows, the need to interconnect them is pushing “lateral/metro” networks and other high-speed infrastructure extensions.

Long-term evolution (LTE) cell sites are also driving 100G metro equipment sales. The growth of mobile devices and the introduction of LTE, with its increased bandwidth demands, have certainly influenced backhaul requirements. The

many new towers coming online, as well as architectural changes, are pushing the need for fiber rings and extensions that will see backhaul rates migrating from 1G to 10G. As a result, the interconnection of the head-end routers at the mobile switching center (MSC) needs to migrate to 100G (see Figure 2).

Fiber relief in congested metro areas is probably the most easily understood 100G application. While some carriers may have an abundance of fiber between central offices (COs), not all are in that position. The number of 10G circuits over the CO-to-CO WDM network has chewed up the available wavelengths in many networks; these 10G circuits are now being swapped out for 100G circuits to offer better resource use and significant cost reduction.



Other equipment drivers include network architecture changes where Carrier Ethernet, IP virtual private network (VPN), and Optical Transport Network (OTN) capabilities have created new momentum to migrate from smaller separate networks to new unified optical backbones. This trend means larger pipes and 100G connectivity. For example, consider 10×10G muxponders. Unless all of the 10G traffic in the resulting 100G wavelengths will go to the same location, enabling OTN switching at natural locations for aggregation nodes makes the most efficient use of resources (see Figure 3). A study with at least one carrier has found savings of more than 20% using OTN switching.

Likewise, carriers may tend to put different services on different wavelengths. Instead of having, say, three partially filled wavelengths, OTN enables packing these three services onto a single wavelength. This drive for further efficiency is another application that is driving 100G transport.

Challenges for 100G metro networking

The future therefore

seems bright for metro 100G deployments, but challenges remain. For instance, as optical speeds increase, it becomes more difficult to overcome optical transmission impairments and still achieve acceptable performance. Obstacles such as chromatic dispersion, polarization mode dispersion, optical passband questions, and nonlinearities increase with 100G.

Fortunately, the modulation schemes, coherent receivers, and digital signal processing (DSP) technology in the latest 100G optical systems will minimize the effects of many of these transmission problems just as they do in long haul networks. But the metro will present some new challenges as 100G migrates from the network core to the edge. We'll note here that there are predominantly two different 100G transport methods for the metro: multicarrier direct detection and single carrier coherent. However,

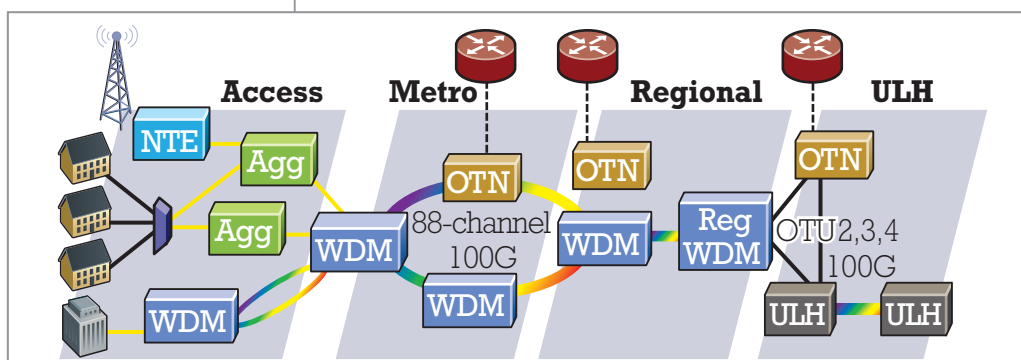


FIGURE 3. Network architecture changes to adopt OTN switching for aggregation is driving 100G demand as well.

FEATURE

100G metro transformation




FIGURE 4. As 100G technology evolves, we're likely to see different approaches pursued for different applications.

we'll limit this discussion to the latter.

Concerns about migration strategies to 100G in the metro are warranted and extend beyond capacity/cost questions. Space and power consumption are also critical challenges in metro networks. The following items also are likely to be key considerations:

- Unlike long haul, metro networks may contain a large number of reconfigurable optical add/drop multiplexers (ROADMs). The 100G performance needs to handle the increase in loss that the number of ROADMs will create.
- Older fiber may mean inconsistent transmission quality, so metro 100G needs to handle high-loss fiber.
- Legacy 10G and the new 100G wavelengths need to coexist without wasting bandwidth on guardbands.
- Low power consumption and

high port density will be needed for the metro.

- Low latency is critical; financial applications, among many other metro-centric use cases, demand low latency.
- Client/line side universal slots and pluggability offer flexibility and lower spares management and logistics costs and therefore should be supported.
- Standardized data-rate and modulation formats are required for interoperability; that's a key lesson from the early days of 40G.

Finally, the simple fact is that 10/100G metro "hybrid" networks will be a reality for some time. While the newer 100G wavelengths based on DP-QPSK 100G modulation and coherent detection can coexist with 10G wavelengths and their on/off modulation and direct-

detection schemes, the truth is that hybrid networks likely will have been designed using dispersion compensation modules to extend 10G reach. Therefore, 100G performance will be suboptimal (due to regeneration and amplification needs as well as the distances involved) on legacy links in comparison to greenfield deployments.

Optical implementations

While metro and long haul applications have different requirements, equipment vendors have increasingly sought to design hardware that will work in both parts of the network. That means they must develop lower-cost 100G technology for the metro in identical form factors to those needed for the long haul. To achieve this feat, equipment vendors will likely consider

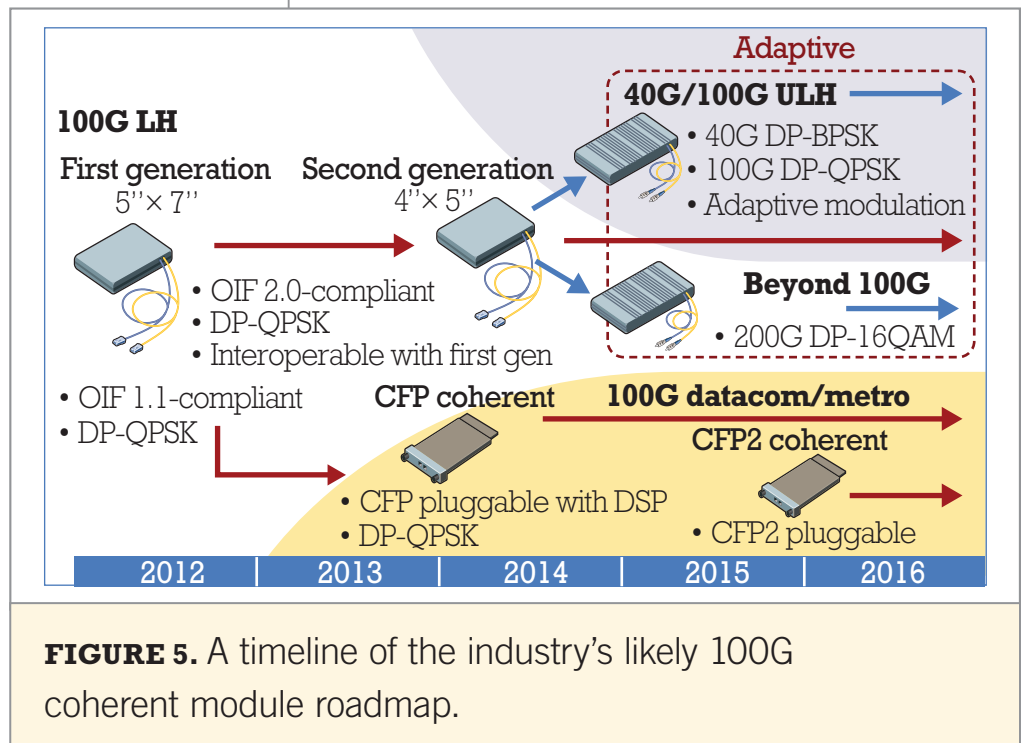
CFP or CFP2-based coherent modules as well as the ability to support multiple modulation schemes.

As the industry matures, however, it's likely we will see a shift to multiple 100G implementations optimized for individual applications (see Figure 4). Some

module examples are as follows:

- 100G MSA (DP-QPSK coherent) for high-performance, regional, and long haul routes.
- 100G CFP (DP-QPSK coherent) for metro performance, smaller size, lower cost, and lower power consumption.
- Several options are under consideration for 100G access applications, such as 100G discrete multi-tone (DMT), 100GBase-ER4 CFP, and 4×28G WDM CFP.

Considering the optical-component options that will be available, a general timeline becomes apparent for the rollout of next generation coherent MSAs as well as CFP and CFP2 modules (see Figure 5). Second generation MSAs are attractive for their power savings



(40%) and improved density. These units can help achieve a 75% reduction in required transponder space.

As coherent components roll out, the dilemma of whether to design-in a CFP coherent module (with onboard DSP) or CFP2 (with off-board DSP) will become a question for manufacturers. While not shown in Figure 5, CFP2 and CFP4 packages will become the default client side optics. Recent industry events have shown [a surge of CFP4 units](#), and some component manufacturers have considered skipping CFP2.

There are some questions surrounding the use of full-band tunable optics, among them: “Is there a reason for it to be pluggable?”

A “pluggable” slot is generally designed as a “universal slot” that could be used for client- or network-side optics. That’s typically possible only with a CFP implementation (at least for now). Since the client side is moving to CFP2/CFP4, the result may be the loss of the ability to have such universal slots unless coherent modules keep pace.

Without the benefit of universal slots, one could question the cost and performance of pluggable interfaces relative to fixed modules or discrete optics. So questions about which design approach will dominate network side

optics, even for the metro, will remain.

Regardless of the outcome, the new client side pluggable units and the new ADC/DSP technology with support for multiple modulation schemes will drive 100G performance and density into metro WDM networks. Standardization will drive competition, increase multisource availability, and decrease cost. Industry alignment behind CFP vs. CFP2 vs. CFP4 optics will take center stage.

The time for metro 100G is at hand

The 100G market has had a healthy start within the metro segment. New and otherwise familiar metro applications are using 100G; 10×10G is favored at this time, while 100G as a service will take time to develop. We will continue to see 10G/100G mixed deployments. In terms of pricing, we have reached 100G card normalized price parity with 10×10G.

There are challenges in the network and optical areas, but new optical-component and equipment manufacturers will continue to drive down cost, increase density, and drive metro rollout. ♦

Note: Amazon, YouTube, Netflix, and Hulu brands are the property of their respective owners.