

Using Carrier Ethernet to Backhaul LTE

By **Michael Howard**
Co-founder and Principal Analyst
Infonetics Research, Inc.

| | |
|--|-----------|
| INTRODUCTION: IP/ETHERNET BACKHAUL AND LTE | 2 |
| LTE AND MOBILE BACKHAUL—ISSUES AND DRIVERS | 3 |
| LTE Opportunities and Challenges | 3 |
| LTE Market Growth | 3 |
| LTE Architecture for Better Economics and Performance | 5 |
| Mobile Backhaul Objectives | 6 |
| Scaling Mobile Backhaul with Lower Cost per Bit Ethernet | 6 |
| LTE Mobile Backhaul Challenges and Requirements | 7 |
| LTE Needs Packet Backhaul, but Legacy Backhaul Still Matters | 7 |
| LTE Needs Higher Backhaul Capacities | 8 |
| LTE Latency Targets Influence Backhaul Architecture | 8 |
| LTE X2 Requires Changes to Backhaul Architecture | 8 |
| LTE Internet Traffic Growth Needs Offload Designs | 8 |
| CARRIER ETHERNET SOLVES LTE BACKHAUL ISSUES | 9 |
| Many Approaches for Ethernet Backhaul for LTE | 9 |
| Ethernet Backhaul Issues of 2010 Are Being Resolved in 2011 | 10 |
| Using Ethernet in Access and Aggregation Architectures | 11 |
| Carrier Ethernet for Any Topology | 12 |
| Ethernet Transport Media: LTE Backhaul over Fiber and Microwave | 13 |
| Metro Optical Transport and P-OTS Using Ethernet | 14 |
| CONCLUSION: CARRIER ETHERNET SOLVES THE CHALLENGES OF LTE MOBILE BACKHAUL | 15 |
| APPENDIX: PACKET TRANSPORT PROTOCOLS: DEFINITIONS AND USES | 16 |
| GLOSSARY OF TERMS | 17 |
| ABOUT INFONETICS RESEARCH | 18 |

INTRODUCTION: IP/ETHERNET BACKHAUL AND LTE

Mobile backhaul today is primarily used in 2G and 3G networks, but 17 LTE networks were launched commercially in 2010, and by the end of 2012 we expect to see 64 operators worldwide with residential and business LTE subscribers. LTE will become the single mobile network for mobile phone and mobile broadband, adopted by most of the world in the next 5 to 10 years. By 2015, LTE will likely support 100 million mobile devices and will continue to rapidly gain subscribers from the 2G/3G base. Additionally, there are potentially billions of M2M (machine to machine) devices, such as smart grid, home control, and medical monitoring, for which planners are depending on ubiquitous HSPA+/LTE services.

Mobile operators and backhaul transport providers are adopting IP/Ethernet backhaul as the default technology choice to cost effectively scale instead of relying on existing TDM-based transport networks for LTE backhaul. The decision to do so is based on a number of facts:

1. In HSPA and LTE, the data plane traffic is IP (Iub interface in 3G and S1/X2 in LTE), so operators are migrating to Ethernet interfaces for base station and controller equipment
2. Ethernet has been and will be the best layer 2 transport mechanism for IP packets
3. Ethernet services and networks are less complex and less expensive than scaling TDM or operating an IP backhaul network
4. Operators trust carrier Ethernet (CE) due to the standards and testing driven by the MEF, the worldwide group comprising service providers and manufacturers (see the MEF's <http://metroethernetforum.org/InformationCenter>)

LTE AND MOBILE BACKHAUL—ISSUES AND DRIVERS


LTE Opportunities and Challenges

LTE and LTE-Advanced have better spectrum efficiency, larger spectrum bands up to 100MHz, 100Mbps downlink at high mobility and 1Gbps at low mobility (achieved through 8x8 MIMO), uplink speeds of 500Mbps, reduced latency, and backward compatibility and interworking with LTE and 3GPP legacy technologies.

LTE Market Growth

Whether measured by carrier spending or subscriber growth, LTE is coming on fast: Infonetics projects 164 million LTE subscribers by 2014.

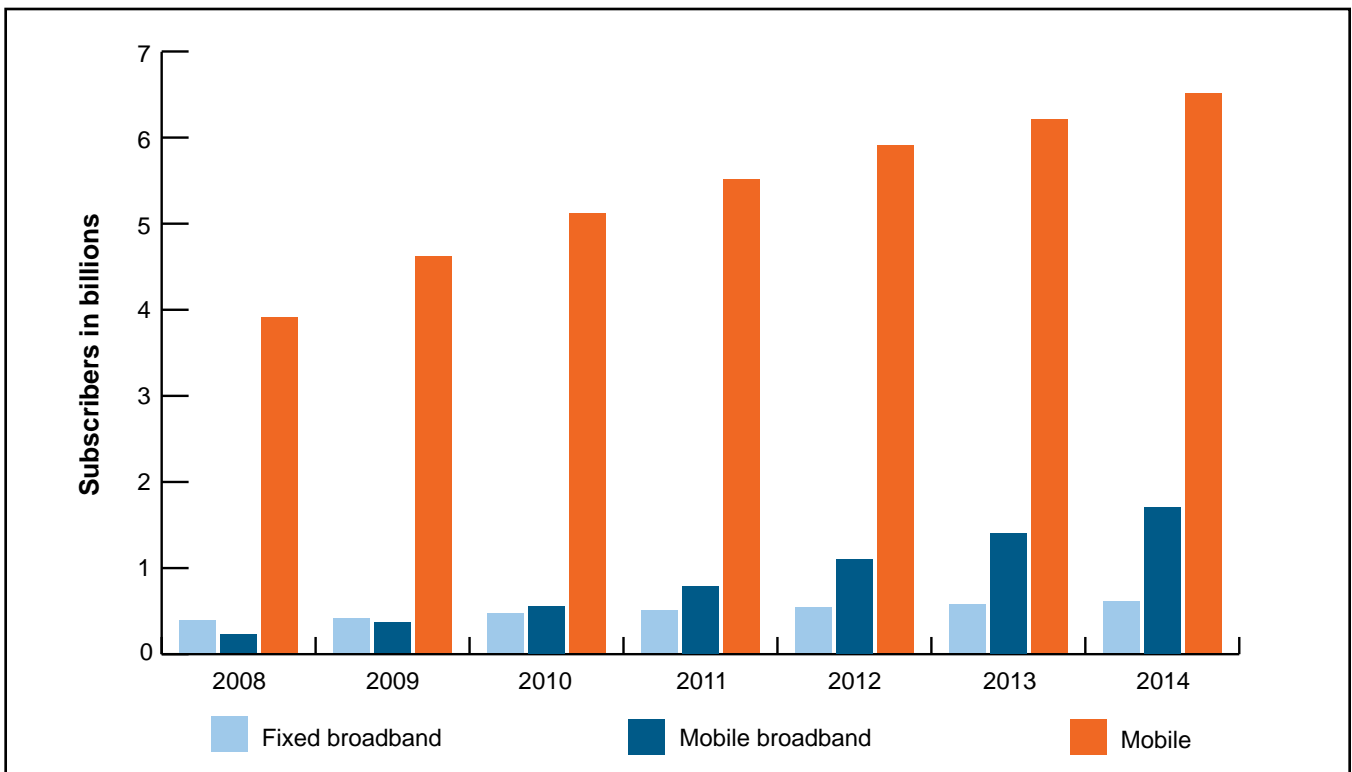
EXHIBIT 1: LTE ADOPTION GROWING RAPIDLY



Source: Infonetics Research, *LTE Infrastructure and Subscribers - Biannual Worldwide and Regional Market Size and Forecasts*, October 2010

There are well over 5 billion mobile phones in operation worldwide, but it is the smaller but rapidly growing number of data hungry mobile broadband devices that are bandwidth killers. Mobile broadband subscribers passed fixed broadband in numbers in 2010, and we expect to see 1.5 billion smartphones, netbooks, tablets and other devices running mobile-optimized OSs, such as Android, Apple Mobile, Windows 7 mobile etc., in 2014. Fixed broadband will always show slower growth as mobile broadband becomes the broadband of choice.

EXHIBIT 2: MOBILE AND MOBILE BROADBAND SUBSCRIBERS GROW FASTER THAN FIXED BROADBAND



Source: Infonetics Research, *Mobile Backhaul Equipment and Services - Biannual Worldwide and Regional Market Size and Forecasts*, September 2010

LTE Architecture for Better Economics and Performance

HSPA and LTE are designed to use IP, the basic communications protocol that moves data, voice, and video and connects users to the Internet. In most cases IP is carried in layer 2 Ethernet packets. LTE incorporates design improvements to overcome some of the limitations in 2G and 3G: 2G has a mobile core—but no packet core—and is designed for circuit switched networks only. Older 3G rollouts included a packet core option but specified the use of legacy existing TDM/ATM and therefore are operationally inefficient for large scale packet networking. HSPA and LTE offer significant advantages:

- A single cellular worldwide standard network for mobile phone and broadband allows seamless roaming
- IP enables options for the integration of fixed and mobile broadband networks
- Higher spectral efficiency and flexible frequency bands achieve lower cost per bit when supporting more users at higher speeds
- They have much lower latency than 2G and 3G (20 ms with LTE vs. 100 ms in 3G at best with HSPA), akin to fixed broadband
- A flatter IP wireless layer architecture with LTE eNodeB base stations evolving to support radio control functions, eliminating need for a base station controller and allowing a direct tunnel from cell sites to gateway site
- LTE lays a clear migration path to LTE-Advanced, expected to be deployed in 2013; LTE-Advanced has a goal of 100Mbps per user, with a future path to even higher capacities; meanwhile, HSPA+ is getting to 168Mbps, so operators plan backhaul technology and capacity upgrades for HSPA+ with LTE requirements in mind

LTE, with higher speeds, can also be a good means of access for cloud services, which many operators see as a new source of revenue growth. For example, Verizon recently announced its acquisition of cloud service specialist Terremark. Verizon's LTE rollout begins in top business markets and in airports, targeting business travelers needing premium LTE's high bandwidth. In addition, parity of bandwidth in the fixed network will allow CIOs to provide mobile broadband users with business applications designed for high speed fixed line connections, instead of developing reduced functionality versions of the applications to fit slower mobile connections.

Because today's 2G and 3G networks have wider area coverage than LTE, operators plan to deploy microcells, picocells, and femtocells (small cells) in buildings and homes to cover the gaps, mostly for LTE, and in some cases for 3G. Such small cells get their backhaul mostly over wireline networks rather than the existing macrocell backhaul. One challenge for operators is arranging for small cell traffic to be sent back to the mobile core and mobile packet core. For this, operators will need to work with building and wireline facilities including the service provider, then setup the logical pathways back to the LTE core network. Some small cells will be outdoor, using microwave to the macro cell site and adding their traffic to the existing cell site backhaul network.

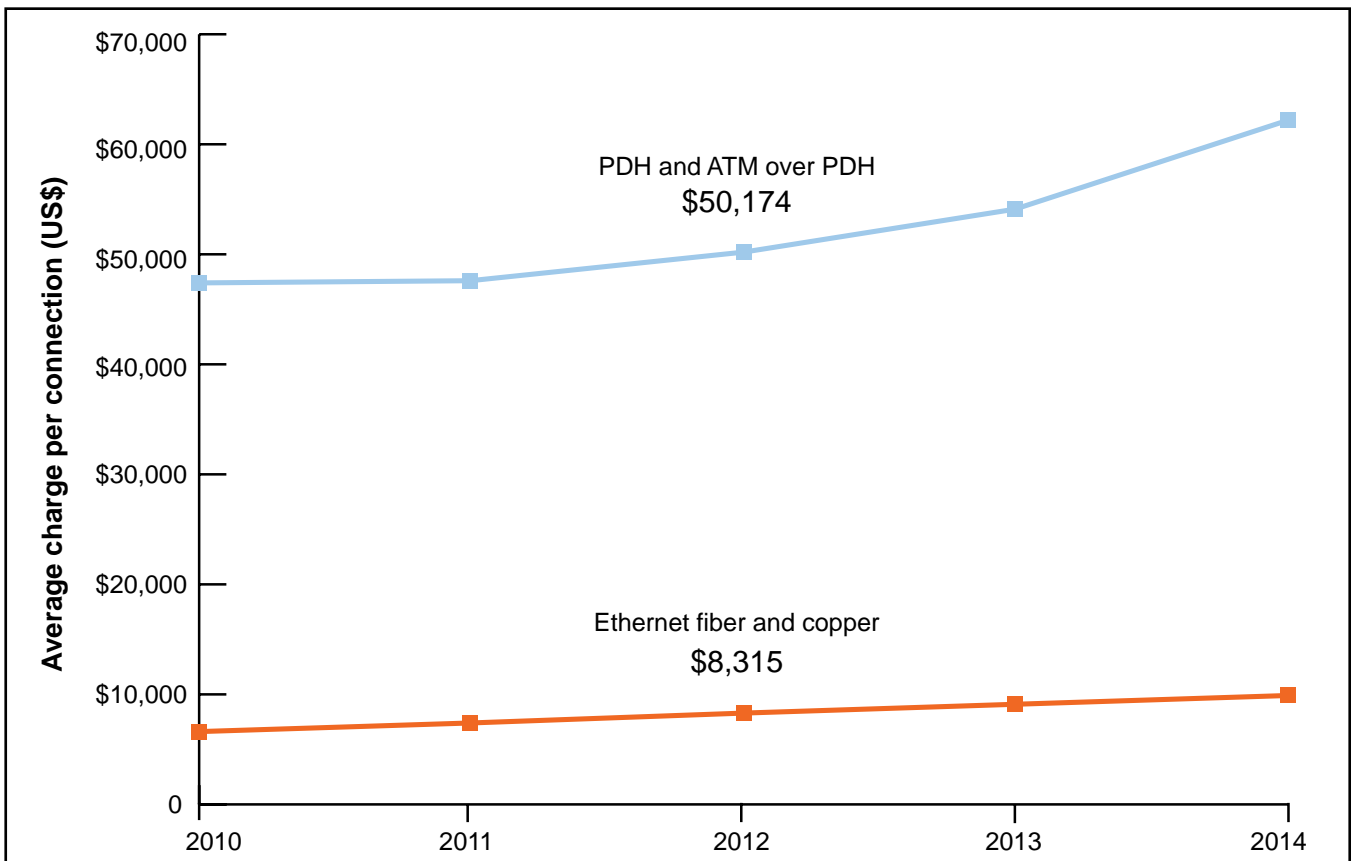
Mobile Backhaul Objectives

Scaling Mobile Backhaul with Lower Cost per Bit Ethernet

To prepare for the high bandwidth usage even for 3G's HSPA+, operators are moving from TDM to packet networks for backhaul connectivity, and that is with Ethernet services, driven by the cost savings they bring: the costs of Ethernet backhaul are well under half of TDM costs, and the higher the capacity, the more the savings. Some of this is due to the inherent statistical multiplexing efficiencies of bursty packet traffic, where capacity planning is based on average traffic volumes, compared to circuits, where it must be based on expected maximum volumes. By switching from TDM to Ethernet, the reduction in MRCs (monthly recurring charges) paid to backhaul transport providers for wireline backhaul means dramatic savings as yearly capacities increase (see Exhibit 3 below).

Ethernet is also less expensive on packet microwave compared to TDM microwave, and our research shows that 55% to 60% of cell site backhaul connections outside of North America are microwave. The biggest backhaul spending today is for upgrading TDM microwave backhaul to dual radio systems that can transmit in native TDM and native Ethernet. Then on any microwave backhaul link, the existing 2G/3G voice is put on TDM microwave, and the fast growing bulk of the traffic is put on the more efficient Ethernet microwave.

EXHIBIT 3: ETHERNET COST SAVINGS VERSUS PDH: ANNUAL MOBILE BACKHAUL SERVICE CHARGES PER CONNECTION



Source: Infonetics Research, *Mobile Backhaul Equipment and Services - Biannual Worldwide and Regional Market Share, Size and Forecasts*, November 2010

LTE Mobile Backhaul Challenges and Requirements

Backhaul networks are inherently complex, using a variety of topologies and technologies to collect data from standalone towers or from building rooftops in cities and suburbs, with a variable number of hops taken before being aggregated onto a metro network. As a result, operators highly value simplicity in deployment and operation of a technology, as they plan to scale networks to support LTE growth with a lower total cost of ownership (TCO). Operators also want the technology choice to allow network designs for deterministic and resilient services, so delay (latency) can be small and the impact of network failures on backhaul service performance minimized. Automation of network and service configuration and ongoing service performance is also critical. Operators are looking for automation and simplicity of operations to deploy or upgrade 100s to 1,000s of cell sites per city or metro region—tools that reduce human involvement in planning, provisioning, deploying, and managing the network.

We see 6 challenges for operators to meet in their move to LTE backhaul.

1. LTE uses packet switched IP—no circuit switched TDM or packet switched ATM
2. LTE supports more users with higher capacities, but with a smaller cell radius depending on spectrum used
3. LTE targets lower latency for user and signaling traffic
4. LTE has a distributed architecture with X2 interface support for eNodeB communications with neighbor eNodeBs
5. LTE usage growth can overload the mobile core
6. LTE base stations are often added to existing towers with 2G and 3G base stations

LTE Needs Packet Backhaul, but Legacy Backhaul Still Matters

LTE today supports mobile broadband data only, with no inherent voice component—it's packet switched only, with no support for circuit switching to carry 2G/3G voice. As LTE handsets become available, they will be using 2G or 3G radio support for voice, via circuit-switched fallback (CSFB). In 2012, LTE is expected to have the capability to carry voice over IP packets using the VoLTE standard.

In many parts of the world, LTE won't be ubiquitous for some time, if ever, so operators will have to support backhaul for not just LTE, but also 2G and 3G for many years to come; the backhaul network will have to be flexible to accommodate this heterogeneous traffic mix.

A key challenge is the network migration strategy to packet backhaul.

1. One option is a separate parallel packet switched network alongside the existing TDM backhaul network, since packet traffic is expected to grow significantly faster with LTE than legacy TDM voice traffic
2. A second option is a single converged packet backhaul network supporting TDM/ATM pseudowire (PWE) emulation over packet for legacy backhaul
3. In addition, there is the option to use the model of a converged packet optical transport system (P-OTS) that has native transport of legacy PDH and SDH/SONET traffic; many operators are using or planning to use P-OTS platforms in their metro aggregation network, and some are considering the use of small P-OTS access devices at the cell site; a P-OTS option can leverage an OTN and WDM wavelength layer for scaling capacity

LTE Needs Higher Backhaul Capacities

Operators planning the evolution of their backhaul architecture must be prepared to handle much greater capacities over the next few years. In the early LTE rollouts that started in 2010, carriers are deploying 50Mbps to 100Mbps cell site backhaul speeds per mobile operator. Backhaul transport providers serving 3–5 mobile operators per cell site are installing 1GE uplinks today, with many planning for 10GE uplinks at high usage cell sites now, and for most of their cell sites over the next few years. LTE-Advanced will provide 10Mbps to 100Mbps peak downstream bandwidth per user with up to 100-200 active users per cell site, so backhaul requirements will increase further. Rings must allow capacities to carry traffic of all the cell sites connected to the ring, which puts the planning requirement quickly at 10GE and multiple 10GEs for aggregation. At aggregation points in metros, even with statistical multiplexing, the long term planning requirements run quickly to 40GE and 100GE.

LTE Latency Targets Influence Backhaul Architecture

The 3G technology HSPA+ can offer bandwidth comparable to LTE. In latency, however, LTE is vastly superior: whereas HSPA+ expects 100ms latency round trip between cell site and controller, LTE expects 20ms, and Verizon Wireless' requirements for transport providers is 5ms. TeliaSonera has shown that online gamers will use LTE like a wired connection, demanding very high-speed and consistent quality connectivity to reduce latency or 'ping' in their gaming experience. LTE users accessing the Internet find that they get response times comparable to, or better than, copper-based fixed broadband, and streaming video works well. With the high quality video capabilities of the latest breed of mobile devices, such as HD-capable tablets, the demand for video to the handset could create exponential growth in the traffic overhead.

A service provider's backhaul network needs to support connectivity options that meet these very stringent latency requirements.

LTE X2 Requires Changes to Backhaul Architecture

The X2 protocol within LTE is a means for eNodeB base stations to exchange protocol messages for user handover directly with each other rather than having to send traffic up to a more centralized controller. The signaling information between the adjacent LTE eNodeB base stations is to control cell to cell handover of sessions as users move from one cell to the next. This traffic is typically between adjacent eNodeBs or among a cluster of neighbors.

Operators need to design their network with the X2 protocol in mind. Backhaul designs must allow options to switch this traffic at some nearest common aggregation point, since there is no requirement to go all the way back to the mobile core (e.g., as in 2G/3G to a GGSN or BSC/RNC). So they need to allow their backhaul architecture to facilitate eNodeB local communications, designing that part as a mesh so it can take that shortcut with lower latency rather than going back to the mobile core.

LTE Internet Traffic Growth Needs Offload Designs

With LTE being IP, it is fairly easy to shuttle Internet traffic around the mobile core; once the initial connection details—access, authorization, accounting, etc.—have been established, traffic can bypass the mobile core and go straight to the Internet. This is an issue for LTE and 3G networks; in either case, the Internet traffic doesn't have to go through mobile core, so operators are looking for ways to add more connections and flexibility to their backhaul networks to offload Internet traffic at the cell site or as near as possible before it reaches the mobile core.

Operators need packet backhaul with the ability to identify, map and prioritize high volume internet traffic toward peering points or content servers. Supporting industry defined interfaces and services at demarcation points between the mobile operator and backhaul operator is a critical step, in addition to supporting packet backhaul architecture with capabilities to handle multiple traffic classes with different performance objectives.

CARRIER ETHERNET SOLVES LTE BACKHAUL ISSUES

Many Approaches for Ethernet Backhaul for LTE

There are myriad ways operators can approach deploying LTE backhaul; a major European mobile operator calculated 158 permutations facing a service provider planning the move to LTE. In our talks with operators around the world, we find many combinations of the following being used, even on a single operator's network, and even on a single backhaul connection from cell site to mobile core.

- Packet services with point-to-point, multipoint, and point-to-multipoint service topologies
- Network topologies: hub-spoke, ring, full or partial mesh, and combinations of these
- Transport media: copper (xDSL, T1, E1, T3, E3); fiber (Ethernet, SONET/SDH/ WDM); microwave (TDM, Ethernet)
- Protocols for packet transport
 - » Ethernet PBB/PBB-TE, EoSONET/SDH/OTN, MPLS, MPLS-TP, VPLS/VPWS
 - » IP-VPNs, IP
- Methods to distribute synchronization: GPS; physical layer methods, including PDH, SONET/SDH, and SyncE (Synchronous Ethernet); and packet methods, principally IEEE 1588v2

LTE backhaul transport infrastructure can be built using different technologies with layer 2 or layer 3 capabilities. Determining the most effective and efficient mix of layer 2 and layer 3 in the backhaul network is a major issue worldwide. Some use layer 3 routers at the cell sites, and others are strongly opposed to that and want to keep backhaul as simple as possible by using Ethernet. Many operators want to keep as much of their backhaul processing as possible in layer 2, while recognizing that MPLS, MPLS-TP, pseudowires, etc., have elements of layer 2.5. Many operators believe the principal layer 2 advantages over layer 3 are simpler equipment and operations—hence lower equipment and operations cost for a lower cost-per-bit as network scales to support large capacity growth.

The backhaul provider builds a network to support the services their customer (the mobile operator) needs. Clearly mobile operators are transitioning to LTE equipment with Ethernet interfaces to support IP for S1 and X2. The key requirements for the backhaul connectivity are the capability to support different connectivity including point-to-point or multipoint, transparency to IP layer, transport network qualities including high availability, SONET/SDH-like OAM, and methods to deliver network synchronization. In any packet transport network, operators need to classify and manage traffic in a differentiated manner.

One network simplification approach is for the packet backhaul network to support Ethernet services. Ethernet is seen as the most effective method to transport IP packets. The LTE mobile operator can, for example, use MEF-compliant interfaces on the eNodeB and on the S-GW and MME. This allows the mobile operator to get MEF-compliant services at the demarcation between the mobile and backhaul networks. The mobile operator can send VLAN-tagged frames toward the backhaul network. A backhaul provider can now identify the service (e.g., VLANs), then map those frames, at the MEF-compliant interface on their equipment, to the EVCs (Ethernet virtual circuits) across the backhaul and provide SLAs on it.

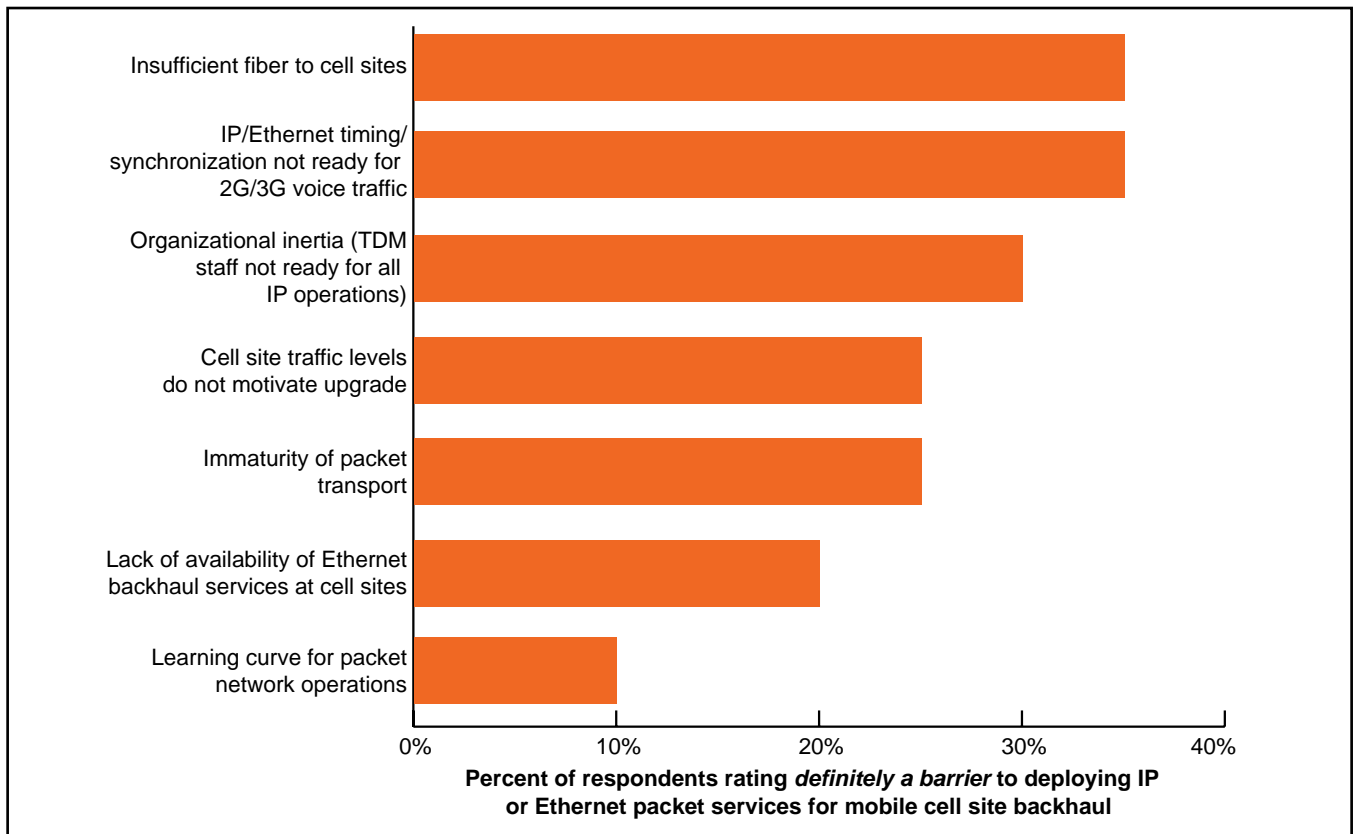
The LTE S1 interface allows an eNodeB to communicate with an S-GW or MME. The S1 interface can be supported over a point-to-point EVC (e.g., EVPL or Ethernet Virtual Private Line). The X2 interface can connect an eNodeB to, at the maximum, the 32 nearest eNodeBs, but more likely a smaller number (<10) in real deployments. A multipoint EVC can be used to support X2 among a group of eNodeBs that need to exchange protocols. When an eNodeB communicates with eNodeBs in different groups, then the X2 traffic can just be sent by eNodeBs on different VLANs to be mapped to different multipoint EVCs.

Carrier Ethernet with MEF 6 and MEF 8 services can support LTE and LTE-Advanced as well as legacy 2G and 3G, which won't disappear any time soon. Nearly all operators going to LTE are also moving to carrier class Ethernet as their backhaul transport, because of its inherent capacity improvement and opex reduction advantages.

Ethernet Backhaul Issues of 2010 Are Being Resolved in 2011

Infonetics conducted a survey of operators in March 2010, *IP/Ethernet Mobile Backhaul Strategies: Global Service Provider Survey*, in which we measured what operators indicated at the time were the barriers to deploying IP/Ethernet backhaul. Respondents were a good representation of the mobile backhaul market; respondents accounted for 42% of 2009 worldwide telecom capex, and they were a mix of incumbents, competitors, and wireless operators, from North America, Europe, and Asia. Exhibit 4 shows data from that survey. Even then, each barrier was important to only 35% or less. The barriers of March 2010 have for the most part been resolved and no longer stand in the way.

EXHIBIT 4: BARRIERS TO IP/ETHERNET BACKHAUL DEPLOYMENT IN 2010



Source: Infonetics Research, *Mobile Backhaul Equipment and Services - Biannual Worldwide and Regional Market Share, Size and Forecasts*, November 2010

With LTE rollout needing higher backhaul capacity, the problem of insufficient fiber is rapidly being fixed: North American operators especially have been urgently replacing copper with fiber in the last 2 years. Microwave and millimeter wave products are gaining in their capacities—many more of these fixed radio products have become available that can handle 400Mbps to 1Gbps.

Another third of the respondents a year ago did not trust the packet timing and synchronization capabilities of IP/Ethernet. At the end of 2009, 25 operators worldwide had sufficient confidence in IEEE 1588v2 and/or SyncE to deploy a single IP/Ethernet packet backhaul network for voice and data. Now the dam has broken: well over 100 operators were deploying a single IP/Ethernet backhaul network by the end of 2010, and more are committing each quarter.

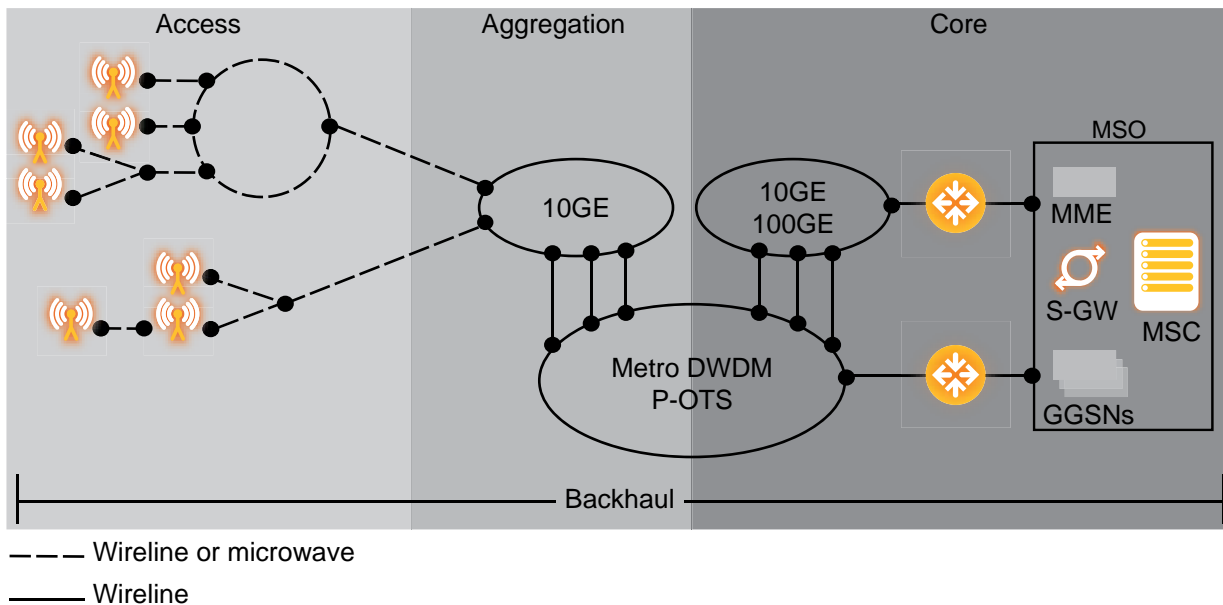
The organizational inertia that worried 30% of respondents has diminished, due to 1) the accumulation of another year’s experience with simpler connection oriented Ethernet (COE) backhaul with SONET/SDH-like operations, 2) further service standardization and interoperability specifications by the MEF and industry in general, and 3) another round of maturing backhaul product products with SLAs controls and OAM&P management tools that have appeared in the interval. Operators have increased motivation—they now add mobile backhaul as another application for Ethernet for business and broadband. Plus they are no longer pioneers; most are deploying IP/Ethernet, all are committed to IP/Ethernet, and it’s only a matter of time before a majority will be using 1588v2 and/or SyncE in a single network for voice and data.

Using Ethernet in Access and Aggregation Architectures

Mobile backhaul is normally between a cell site and an MSO site (see Exhibit 5). With LTE, the mobile backhaul service is both for cell site to/from MSO and for cell site to/from cell site. An MSO site for LTE can be in a centralized location such as a region, or it can be per metro; hence, the scope of backhaul can be different from that of a typical 2G/3G model.

Typically, the backhaul network consists of access and metro domains. The metro domain might also have one or two aggregation tiers in major cities. Operators regard their access and aggregation/metro networks as separate components: each has unique requirements. Each telco typically has one or two mobile switching offices (MSO) in a metro area, where 100s to 1,000s of sites are aggregated.

EXHIBIT 5: MOBILE BACKHAUL NETWORKS: ACCESS, AGGREGATION, CORE



Source: Infonetics Research, February 2011

The major reason many providers choose to use only layer 2 in access and aggregation networks in general—and mobile backhaul networks in particular—is to avoid adding the complexity of IP/MPLS control planes. Operators say they dictate using layer 2 only to keep operations as efficient as possible—saving staff time, minimizing elapsed time, reducing training costs, lowering the number of more expensive router engineers, and reducing the number of errors.

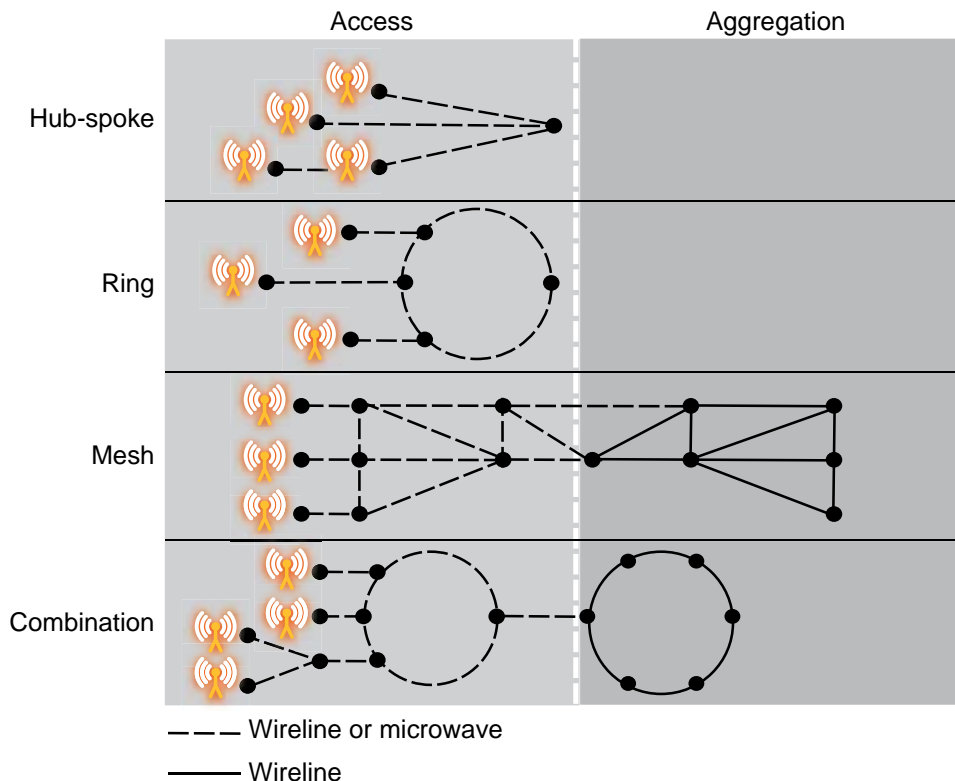
There’s an overhead cost for handling each data plane and control plane—the more there are, the more complex the network. For example, in the case of an IP-VPN, BGP is used for service discovery, and OSPF or ISIS is used for topology discovery and path computation, whereas LDP or RSVP-TE is used for signaling labels.

A carrier Ethernet access and aggregation network may not need a complicated control plane, since physical topology (fiber links) is not very complicated in most cases, and most mobile backhaul networks are simple hub-spoke. Many operators use a management-based provisioning model for most tasks—MPLS-TP or PBB-TE.

Carrier Ethernet for Any Topology

Most of the world’s legacy backhaul networks have a hub-spoke or ring access architecture, and some are taking advantage of IP/Ethernet’s flexibility to begin planning partial mesh topologies. It is common practice in hub-spoke topologies to string together a number of spokes into a serial line of cell sites, particularly when using microwave. In access networks, carrier Ethernet solutions using MEF services can run on many topologies, including parallel links or partial mesh to allow diversity, in addition to usual approaches such as an access ring. Any of the hub-spoke, ring, and mesh topologies can be used in a mobile backhaul network. A key advantage of Ethernet is that it is flexible enough to build any of them.

EXHIBIT 6: MOBILE BACKHAUL TOPOLOGIES: HUB-SPOKE, RING, MESH



Physical topologies can support all connectivity types: the service topology can sit on any combination of hub-spoke, ring, or mesh. A logical point-to-point topology, for example, can be built on a hub-spoke, ring, or mesh physical topology. In a ring, a point-to-point service is between two nodes on that ring. A multipoint service is between 2 or more nodes in the ring. In a point-to-multipoint service, one of the ring nodes is assigned to be the root and the other nodes are assigned to be leaves. So, the forwarding behavior around the ring is logically like a hub/spoke across the ring.

The capacities needed at a cell site on hub-spoke, serial, ring, or mesh configuration depend on the site's position; the last site before the hub in a series of sites must carry the accumulated traffic of the further sites and its own. Many microwave cell sites fall into such a series. Similarly, for a ring there must be enough bandwidth to accommodate all sites on the ring.

A significant design impact of topology choice is to choose the protection schemes that can be employed for high availability. Schemes such as IEEE Link Aggregation or 1+1/1:1 with parallel links or ring protection can be used in addition to path or node and link diversity. Additional mechanisms for protection around a failed section of the network instead of the entire path can also be deployed. When carriers want resiliency in their backhaul networks, they set up alternate routes from cell sites and/or from the aggregation hub point to the mobile core and to the Internet. Failure at the cell site itself often can't be helped, and some failure is tolerated due to the ability of adjacent cells to pick up sessions. Rings are an option to provide resiliency at cell sites, if required. Transport providers, who serve a number of mobile operators at a cell site, will have resilient aggregation and metro core networks—their mobile operator customers typically contract for it.

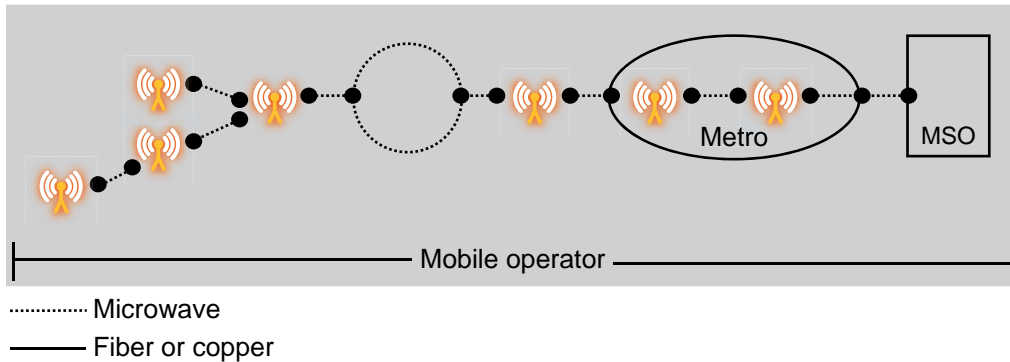
In the aggregation/metro network, topology is simpler (a few rings or in some cases mesh) and mostly fiber. In some cases microwave links are used (Exhibit 7) all the way to the MSO, but generally fiber links are used across a wireline operator's network (Exhibit 8). Resiliency is critical at aggregations sites, where traffic from hundreds of sites gets aggregated. (Unless on a ring, most access sites have no resiliency.) Carrier Ethernet allows a common data plane across access AND aggregation/metro networks.

As operators move to packet backhaul, they need to simplify and automate their tools. A metro can have 1,000s of cell sites. An operator might need to turn up 20-30 cell sites per day since operators are not looking for bandwidth upgrades or a 2G/3G-to-LTE upgrade just for a couple of sites; they want to turn on the service for as many customers as possible.

Ethernet Transport Media: LTE Backhaul over Fiber and Microwave

In most of the world (outside North America), most backhaul connections are on microwave. Those using microwave typically own the equipment, and have no monthly charges to pay a transport provider. Though many mobile operators also own wireline operations (e.g., AT&T, Verizon, DT, and Orange; even Vodafone operates DSL networks in Germany and Portugal), the wireless and wireline divisions operate separately in most cases. Thus, nearly every mobile operator uses—and pays for—a transport provider for wireline backhaul. Virtually all mobile operators operate their own microwave facilities, except specialist companies that own and operate cell sites and backhaul. A CE solution that uses MEF services also needs solid OAM mechanisms for fault management across a microwave-only or mixed microwave/fiber topology. This is important because microwave links are susceptible to weather conditions. A typical metro microwave network is shown in the Exhibit 7.


EXHIBIT 7: METRO BACKHAUL USING MOBILE OPERATOR'S MICROWAVE



Source: Infonetics Research, February 2011

When two operators are involved in the backhaul (a mobile operator and a backhaul transport provider), the interaction of those 158 choices operators face as they build their networks can cause complications: routers or Ethernet, hub-spoke or ring, layer 2 or layer 3—it's like arranging a cross cultural marriage.

EXHIBIT 8: METRO BACKHAUL USING WIRELINE TRANSPORT PROVIDER



Source: Infonetics Research, February 2011

Because of the complications of a mobile operator having self-operated microwave and working with transport provider wireline networks, plus the general desire for simplification when managing thousands of cell sites, mobile operators look for Ethernet backhaul equipment that can be used in any microwave, copper, or fiber situation.

Metro Optical Transport and P-OTS Using Ethernet

IP/Ethernet backhaul traffic around and within a metro area is carried on a metro Ethernet network. P-OTS products are made to work easily with Ethernet transport. Connection oriented Ethernet (COE) on P-OTS platforms puts Ethernet “circuits” on wavelengths; operators use it as transport for Ethernet traffic around their networks (see Exhibit 5).

Many service providers worldwide rolling out LTE have P-OTS (e.g., Verizon). P-OTS platforms can carry TDM (SONET/SDH) and Ethernet packet traffic, both natively.

CONCLUSION: CARRIER ETHERNET SOLVES THE CHALLENGES OF LTE MOBILE BACKHAUL

LTE represents a major change in mobile telephony—a single worldwide standard—and mobile broadband—widely available broadband services with capacities and latency competing with fixed broadband. LTE is designed to fit into the IP world—LTE is all IP packet with no TDM, so it is different from existing 2G/3G networks. LTE also has useful architectural changes that allow eNodeBs to talk among themselves and distribute the mobile core functions, both of which suggest a change in today's mostly hub-spoke backhaul to a partial mesh or ring design. Given that the world has gone mobile, and will do so even more in the future, much of the future of telecom rests on LTE.

Concurrently, mobile backhaul networks have been under siege for the past three years as the smartphone market, typified by the iPhone, has exploded, and with it, data usage, which stresses backhaul networks. Adding LTE will only exacerbate the backhaul problem, bringing with it the need to continue backhauling 2G and 3G traffic.

Although LTE is IP-based, many mobile operators and backhaul transport providers want to keep their access and aggregation networks simpler by avoiding layer 3 routing and avoiding the use of dynamic routing and signaling protocols across the backhaul to the cell site. They also prefer to transport not only LTE traffic, but also 2G/3G on carrier Ethernet backhaul (which they consider less complex and less expensive). Indeed, the costs of IP/Ethernet backhaul transport, whether wireline or microwave, are much lower than the costs of TDM.

In a March 2010 Infonetics survey of operators worldwide, respondents indicated many barriers to using IP backhaul; now we see well over 100 operators deploying IP Ethernet for their 2G and 3G voice and their data backhaul. Operators now trust the packet timing and synchronizing capabilities of IEEE 1588v2 and SyncE.

Carrier Ethernet interfaces and services are defined by the MEF and have been widely adopted by service providers around the world. MEF-certified CE equipment is available from all major vendors. In this paper, we have shown how CE can be used to solve the many issues and problems involved with LTE backhaul:

- Higher capacities, lower latency
- Topologies: hub-spoke, ring, mesh
- Microwave and wireline
- Mobile operator owned microwave backhaul, and transport provider wireline services
- Resiliency
- 2G, 3G, and LTE on same backhaul network

The appeal of and quest for a less complex, more efficient operation than layer 3 routing leads many mobile operators and backhaul transport providers to stipulate the use of CE for LTE backhaul.

APPENDIX: PACKET TRANSPORT PROTOCOLS: DEFINITIONS AND USES

| Protocol Name | Definition | Defining Standard | Other Names | Typical Uses |
|----------------------------------|--|-------------------|--|--|
| IEEE-defined | | | | |
| PB (a.k.a., stacked VLAN) | Stacked virtual local area network; allows VLANs within a VLAN | 802.1ad | Q-in-Q | Separate customer VLANs from provider's VLAN; used for Ethernet services, E-Line, E-LAN, wholesale to other providers, IPTV, triple play, mobile backhaul |
| PBB | Provider backbone bridging; bridges packets across provider backbone | 802.1ah | MAC-in-MAC | Scale by encapsulating customer MAC addresses and VLANs; targets same uses as stacked VLANs, but enables much greater scaling by eliminating VLAN tag and customer MAC address limitations |
| PBB-TE | Provider backbone bridging-traffic engineering; adds circuit-like deterministic characteristics, including a backup path | 802.1Qay | PBT | Scale Ethernet point-to-point connections; adds point-to-point traffic engineered connections to PBB for QoS and bandwidth management; targets same uses as stacked VLANs |
| ITU-IETF-joint definition | | | | |
| MPLS-TP | Multiprotocol label switching-transport profile; removes features not needed by connection-oriented applications, adds transport support | MPLS-TP | | Use MPLS tunnel to carry individual "circuits" like pseudowires, so can hide customer MACs and VLANs in the provider network; targets ATM, frame relay, and TDM uses and same packet uses as stacked VLANs; use for statistical multiplexing of IP, MPLS, and Ethernet in optical transport networks (much more efficient for packet transport than TDM-based transport) |
| IETF-defined | | | | |
| VPWS or Pseudowire or PWE3 | Virtual private wire service; provides a point-to-point service, eliminating a leased line | RFC 4664 | Virtual leased line (VLL) | Use to encapsulate any of ATM, frame relay, TDM, or Ethernet for packet point-to-point connection; targets ATM, frame relay, TDM uses and same packet uses as stacked VLANs |
| VPLS using LDP signaling | Virtual private LAN service using LDP signaling; a multipoint VPN service that emulates a LAN across a WAN | RFC 4762 | Multipoint to multipoint VPN, lasserre-vkompella/martini | Creates a "cloud" Ethernet service for any-to-any connections over MPLS backbone networks; targets same packet uses as stacked VLANs except for multipoint configurations |
| VPLS using BGP signaling | Virtual private LAN service using BGP signaling, a multipoint VPN service that emulates a LAN across a WAN | RFC 4761 | Multipoint to multipoint VPN, kkompella | Creates a "cloud" Ethernet service for any-to-any connections over MPLS backbone networks; targets same packet uses as stacked VLANs except in multipoint configurations |

GLOSSARY OF TERMS

| Term | Definition |
|----------------------------|---|
| APS | Automatic protection switching; a form of network resiliency involving an alternate path in case of failure of the primary path |
| E-LAN | MEF-defined E-LAN service used to create multipoint L2 VPNs and transparent LAN service; is a foundation for IPTV and multicast networks |
| E-Line | MEF-defined E-Line service used to create Ethernet private lines, virtual private lines, and Ethernet Internet access |
| eNB | see eNodeB |
| eNodeB (eNB) | Evolved NodeB, or LTE base station evolved from the 3G NodeB base station |
| E-Tree | MEF-defined Ethernet private tree (EP-Tree) and Ethernet virtual private tree (EVP-Tree) services, provide traffic separation between users with traffic from one “leaf” being allowed to arrive at one of more “roots” but never being transmitted to other “leaves” |
| EVC | MEF-defined Ethernet virtual circuit |
| G.8031 (linear APS) | G.8031/Y.1342 Ethernet linear protection defines the APS protocol and linear protection switching mechanisms for point-to-point Ethernet VLANs |
| G.8032 (ring APS) | G.8032 defines the APS protocol and protection switching mechanisms for Ethernet rings; includes bridged ring protection characteristics, architectures, and the ring APS protocol |
| GGSN | Gateway GPRS support node; responsible for the interworking between the GPRS network and external packet switched networks, like the Internet |
| ISIS-TE | Intermediate system to intermediate system--traffic engineering |
| LDP | Label distribution protocol; defined in RFC 3036; LDP is used to provide mechanisms for MPLS routers to process and route labeled traffic (LSPs) across an MPLS network |
| LSP | MPLS label switched paths; virtual circuits over an MPLS backbone |
| LTE | Long-term evolution (LTE); the 3GPP's latest standard in the mobile network technology tree, which produced the GSM/EDGE and UMTS/HSPA network technologies |
| LTE-Advanced | The next evolution after LTE, encompassing 3GPP Rel. 10, 11, and beyond |
| MEF | Metro Ethernet Forum (www.metroethernetforum.org) |
| MME | Responsible for all mobility management in LTE, which is moved into the mobile core; these functions were performed by the RNC and NodeB/BTS in 3G networks |
| MSC | The mobile switching center (MSC) is the primary service delivery node for GSM/CDMA, responsible for routing voice calls, SMS, and other services |
| MSO | Mobile switching office; a location where 2G/3G equipment such as MSCs and GGSNs reside, and where LTE mobile core equipment sits including S-GWs and MMEs |
| OSPF-TE | Open shortest path first--traffic engineering; a routing protocol that determines the best path for routing IP traffic over a TCP/IP network based on distance between nodes and several quality parameters; OSPF is an interior gateway protocol (IGP), which is designed to work within an autonomous system |
| P-OTS | Packet-optical transport system; WDM products with Ethernet switching and circuit switching (SONET/SDH crossconnect and/or OTN) across the chassis; supports connection oriented Ethernet (COE) protocols (e.g., MPLS-TP) and ROADM |
| RSVP-TE | Resource reservation protocol--traffic engineering (MPLS); used to signal point-to-point (P2P) label switched paths (LSPs) across MPLS and GMPLS networks |
| S1 | LTE S1 interface defines communication between the MME and eNodeB |
| S-GW | Serving gateway in LTE networks; manages user mobility and acts as a demarcation point between the RAN and core networks |
| SPB | Shortest path bridging; 802.1aq is the IEEE-sanctioned link state Ethernet control plane for VLANs covered in IEEE 802.1Q; SPB combines an Ethernet data path (IEEE 802.1Q or provider backbone bridges (PBB) IEEE 802.1ah) with an IS-IS link state control protocol running between shortest path bridges (NNI links) |
| T-LDP (tLDP) | Targeted LDP is used for MPLS inner label distribution (see LDP) |
| X2 | LTE X2 interface defines a direct connection between eNodeBs with stringent delay requirements |

WHITE PAPER LEAD ANALYST/AUTHOR

Michael Howard

Co-founder and Principal Analyst

michael@infonetics.com | +1 (408) 583.3351 | twitter.com/MichaelVHoward

ABOUT INFONETICS RESEARCH

Infonetics Research is an international market research and consulting analyst firm serving the communications industry since 1990. A leader in defining and tracking emerging and established technologies in all world regions, Infonetics helps clients plan, strategize, and compete more effectively.

Services

- Webinar, Conference, and Event Speaking
- Custom Brand and Demand-Side Market Research
- Custom Market Size and Forecasts
- Technology and White Papers
- Competitive Analysis and Due Diligence
- Market Share, Market Size, and Forecasts
- Enterprise/SMB and Service Provider Survey Research
- Continuous Research Services
- Service Provider Capex and Subscriber Analysis and Forecasts
- Consulting, Retainers, and Quick Consults

Coverage Areas

- Mobile and Wireless
- FMC and Femtocell
- Mobile Backhaul and Microwave
- Service Provider VoIP and IMS
- Broadband Access
- Cable, Satellite, and IPTV Video
- Enterprise Voice, Video, and UC
- Telecom and Datacom Equipment Totals
- Next Gen OSS and Policy
- Carrier Routing, Switching, and Ethernet
- Optical
- Service Provider Outsourcing
- Service Provider Capex and Subscribers
- Data Center Networks
- Security
- Enterprise Networking

SALES

Larry Howard

Vice President

larry@infonetics.com | tel: +1 408.583.3335 | fax: +1 408.583.0031

Scott Coyne

Senior Account Director - Eastern North America, Europe, and the Middle East

scott@infonetics.com | tel: +1 408.583.3395 | fax: +1 408.583.0031