White Paper

New Backhaul Requirements for LTE, LTE-Advanced & Beyond

Prepared by

Patrick Donegan
Chief Analyst, Heavy Reading
www.heavyreading.com

on behalf of

www.juniper.net

October 2015

The Market Landscape for Mobile Operators

For licensed mobile network operators, the demand side of the next five years is hugely exciting, even mouth-watering. What wouldn't businesses in other sectors – such as the oil and gas, financial services, mining, agriculture and automotive industries – give for a mobile network operator's demand outlook promising a very steep, pretty much inexorable, year-on-year increase in demand for their products and services over the coming five years and beyond?

Specifically, licensed mobile operators can look forward to strong growth in demand for all of the following:

- The number of connections (driven increasingly by the Internet of Things)
- The number of sessions per user
- Higher download and upload speeds
- More real-time services
- Higher security
- Higher reliability

Crafting strong business performance out of this abundant demand – and the huge increase in bandwidth required to cope with it – certainly won't be straightforward. Although some leading operators in the U.S. and elsewhere have already shown how to leverage the potential of 4G for superior "monetization," other 3G and 4G operators are still languishing with flat or even deteriorating financial performance.

Five Pillars of Growth

Mobile operators can think in terms of five pillars of growth that will drive their businesses over the next few years:

1. **Network capacity needs to be substantially expanded year-on-year,** by means of adding raw spectrum, cell site densification and cutting-edge spectral efficiency techniques.
2. **New revenue-generating applications and services, such as voice over LTE (VoLTE) and LTE Broadcast, need to be deployed.** VoLTE is now being rolled out aggressively worldwide. Led by operators in South Korea, the U.S. and Europe, momentum behind LTE Broadcast is also starting to build.
3. **A step change is needed in network performance, as measured by critical metrics and requirements such as latency and synchronization, to deliver these new capabilities.** The network must be designed in such a way, and new performance targets must be attained, so that efficiency and revenue-generation goals cited above can be delivered competitively.
4. **Business relationships with partners and competitors need to be optimized.** Where over-the-top (OTT) providers such as Google and Facebook are concerned, operators need to regularly review the extent to which these companies are partners, competitors, or a little of both.
5. **Capitalizing on SDN.** Software-defined networking (SDN) promises a slew of opportunities to reduce cost, improve network performance and grow revenues. Keeping at least in step with these opportunities is a key enabler for supporting the other four pillars of growth, as well as the operator's broader business objectives.

The Evolution of the RAN

The roadmap for the RAN in delivering on the operator's objectives is well-understood and well-documented. The primary requirements are:

- **More spectrum:** Operators will inevitably continue to bring more and more spectrum online. In addition to growing capacity, LTE carrier aggregation will continue to increase the speeds available to individual users.
- **Cell site densification via more macro cells, micro cells and particularly small cells:** Just a few short years ago, the assumption was that small cells might serve as a supplement to the macro and microcell network in high-density urban areas. Nowadays, network architects increasingly believe that capacity requirements will be such that small cells will carry the bulk of the urban traffic in a few years, and that they will be supplemented by a macro and micro layer. As part of this roadmap, many operators are considering Centralized or Cloud RAN (C-RAN) or fronthaul architectures. Here the baseband is pooled and serves separate RF heads (or small cells) via a standard interface over distances that may range from a few meters to 25 kilometers. From a backhaul perspective, operators clearly prefer fiber for small cells, whether they require conventional backhaul or fronthaul according to a C-RAN architecture. That said, operators recognize that they will also require wireless backhaul where fiber is either not available or cost-prohibitive.
- **Leveraging LTE-Advanced (LTE-A) features for better spectral efficiency and better performance:** *Carrier Aggregation* is already extensively deployed worldwide. New features such as *Enhanced Inter Cell Interference Coordination (eICIC)* and *Coordinated MultiPoint (CoMP)* transmission/scheduling will also be rolled out, often in a small cell or HetNet context. CoMP combines signals from multiple eNodeBs to improve cell edge throughput, often in a small cell or HetNet context. Some leading RAN vendors reckon that UpLink CoMP can deliver an average throughput gain of 150 percent at the cell edge. Market momentum appears to be behind Intra-site UpLink CoMP at this time. eICIC enables better interference management between layers in HetNets, using and reusing some of the same frequencies across the macro and small cell layers. SK Telecom, for example, expects to reduce inter-cell interference by 15 percent when it rolls out eICIC in 2016.

The Evolution of the Backhaul Network

RAN requirements tend to get the most attention because the RAN is where the operator spends most money. But it is equally true that investment in the backhaul needs to stay in lockstep with investment in the RAN – otherwise, the expected performance gains from the RAN simply won't be realized.

The Looming New Backhaul Bottleneck

Having outlined the "Five Pillars" of success, as well as some of the specific upcoming changes in the RAN as operators evolve through releases of LTE, LTE-A and with an eye toward future 5G requirements, this paper focuses on the role of the backhaul network in delivering on the operator's objectives over the next five years. Many LTE-A features, such as CoMP and eICIC, can be implemented in the RAN with a relatively simple software upgrade. But as the following sections of this paper show, re-equipping the backhaul to support some of those new RAN features is going to be a lot more involved.

There is a new bottleneck emerging in the backhaul. The original bottleneck was a pure capacity bottleneck, identified when operators determined that the model of scaling E1s and DS1s for 3G data traffic wasn't financially viable. This was overcome by the replacement of TDM with IP backhaul.

The new bottleneck that is looming with further rollouts of LTE and LTE-A does in part relate to capacity: Some operators are finding that their ability to deliver the required backhaul capacity for 4G is challenged. As shown in **Figure 1**, Vodafone is an example of a leading operator that is prioritizing the buildout of high-capacity backhaul, which it defines as capable of supporting up to 1 Gbit/s per site. More than 90 percent of Vodafone's European footprint supports high-capacity backhaul, the majority of which is served by microwave in the last mile.

Figure 1: Vodafone Group's Buildout of High-Capacity Backhaul

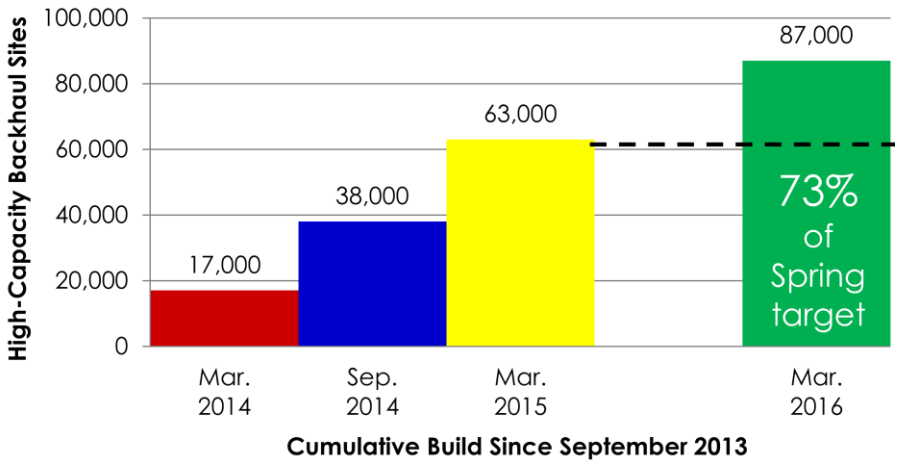

Source: Vodafone

What's different with the new bottleneck up ahead is that in addition to a capacity element, there's also a performance as well as a security element to it. New backhaul challenges that involve performance, synchronization and security include:

- Real-time applications such as VoLTE and some OTT services introduce new latency requirements into the backhaul.
- New end-user services such as LTE Broadcast and new network capacity features such as CoMP and eICIC have much more stringent synchronization requirements than anything mobile operators have tried to achieve before.
- The long-term implication of the growth in small cells is that the option of using small cells as a "slow lane" or for "offload" offering inferior quality and reliability will only remain viable in the near term. Once small cells become the primary workhorse of the urban network, the small cell layer and its supporting backhaul must be built for the highest performance requirements.
- In the LTE era, security is both a new revenue opportunity and a potential bottleneck for the first time. Whereas security was something that came bundled in with 2G and 3G systems, it is now something mobile operators need to design, deploy and manage as a set of requirements that are unique to its own business model and network assets, including in the backhaul.

New Latency Requirements

Latency requirements are changing across the mobile network, including in the backhaul. For VoLTE, network planners need to ensure that there is enough capacity in the network, and that the prioritization accorded to each audio packet is enforced end-to-end and across each of the RAN, backhaul and core domains. This enables a stream of VoLTE packets to be transmitted with sufficiently low latency, jitter and delay variation to ensure a high-quality user experience of a VoLTE call.

The backhaul is unique among all domains in that packet prioritization – hence application performance – has to be enforced across an often highly heterogeneous network environment in terms of the coexistence of different layers and different protocols, as well as different physical layers such as fiber, XDSL and microwave, each of which introduce different characteristics from a latency perspective.

Engineering Latency for VoLTE

The new challenge where new low-latency services such as VoLTE are concerned lies in the way the different services now have to be mapped across multiple Differentiated Services Code Point (DSCP) bits for classifying traffic and providing QoS at both L3 and L2. And whereas with email and Web browsing the user's experience isn't impacted if three quarters of the latency is experienced in one direction of a link, and a quarter in the return path, in the case of VoLTE that kind of sub-optimal engineering will deliver a sub-standard user experience. Nor is it just the operator's own services that require lower latency. If operators are to persuade OTT providers to pay them a premium for delivering some services, lower latency will be a prerequisite in some cases.

Engineering Latency for CoMP

In current LTE networks, where the X2 interface is typically turned around in the core, 10 ms of latency is perfectly tolerable on the X2. With CoMP, however, the coordination between the eNode Bs at handover requires a much lower X2 latency. The amount of latency on the X2 materially affects the performance gains that CoMP can deliver, with the result that X2 latency needs to be lower than 5 ms to drive truly excellent performance.

Synchronization: A New Differentiator

The synchronization requirements for CoMP, eICIC and LTE Broadcast are a lot more challenging for mobile operators than any synchronization challenge they have faced previously. Generally speaking, this is still not well enough understood.

In most synchronization approaches deployed to date, backhaul network elements have tended to be transparent to the implementation. The frequency synchronization required between adjacent base stations – the only type of synchronization required for the most widely deployed 2G, 3G and 4G networks – is either embedded in an E1/T1 or transported transparently via the Global Navigational Satellite System (GNSS) or IEEE 1588v2 Point-to-Point (PTP) protocol over a packet backhaul network or the ITU's Synchronous Ethernet (SyncE) protocol. SyncE also provides frequency synchronization but requires on-path support in all network elements across the backhaul.

Frequency Synchronization Is Relatively Easy to Achieve

Frequency synchronization is relatively straightforward because recovering frequencies takes place entirely independently of variations in network traffic. If the transmission of frequency signals is impacted by jitter or delay in the network, there is no impact. The system counts the signals, puts in a buffer, and averages out in order to meet the 3GPP target of ±50 parts per billion, which is the standard for GSM, W-CDMA and LTE FDD.

Figure 2: 3GPP's Synchronization Requirements

APPLICATION	FREQUENCY	PHASE	NOTES
3G	50 ppb	Not required	None
LTE-FDD	50 ppb	Not required	
LTE Broadcast	50 ppb	±10 µs	Inter-cell time difference
eICIC (LTE-A)	50 ppb	±1.5-5 µs	
CoMP (LTE-A)	50 ppb	±5 µs	Moderate to tight
CoMP (LTE-A)	50 ppb	±1.5 µs	Very tight

Source: Heavy Reading

The Brand-New Challenge of Phase Synchronization

Going forward, the success or otherwise of the new synchronization model required for CoMP, eICIC and LTE Broadcast will be heavily dependent on the capabilities supported in the network equipment, network design and engineering across the backhaul domain. The key to this upcoming set of new requirements in the backhaul landscape is that CoMP, eICIC and LTE Broadcast all require phase synchronization which is a very different proposition to the frequency synchronization that most mobile operators are used to.

Whereas frequency synchronization restricts the deviation in the rate at which a clock at a cell site can tick relative to the rate of the Primary Reference Clock (PRC) – together with buffer for averaging it out to achieve target – the requirements for phase synchronization are much more stringent.

Phase synchronization restricts the deviation in the exact moment when a cell site clock can tick relative to the PRC. As shown in **Figure 2**, 3GPP restricts the deviation down to between 1 µs and 5 µs, depending on the use case for phase synchronization. And because there can't be any buffering or averaging, attaining these stringent targets requires raising the bar significantly in terms of backhaul performance to eliminate or reduce the risk of jitter and delay impacts.

Exceeding the phase synchronization target for LTE Broadcast will cause video broadcasts to be interrupted – and we know that users (especially those willing to pay a premium) have a notoriously low tolerance threshold for poor video quality. Depending on exactly how wide of the mark it is, missing the phase synchronization target will also cause CoMP implementations to fall short – potentially well short – of the 150 percent cell edge throughput gains that are potentially achievable.

Supporting Phase Synchronization in the Backhaul Network

From a synchronization perspective, the backhaul ceases to be a "dumb pipe" when the operator prepares to roll out LTE Broadcast and many of the LTE-A network features. Instead, capabilities built into the backhaul become critical to achieving the operator's synchronization objectives.

In order to get down to the required levels, operators need to think in terms of shaving tens of nanoseconds – even individual nanoseconds – out of the delay budget to get to the required level of accuracy. Just as – or more – importantly, packet delay variation (PDV) needs to be very tightly controlled for phase synchronization. This can be a bigger challenge than absolute latency, particularly in multi-hop microwave backhaul networks.

It can be tempting to target the least tight requirement, particularly if the first use case only requires that. It's far better to take the long-term view and design the backhaul with the tightest possible requirements in mind, however. The incremental spending up front to engineer the network optimally from the outset will inevitably be less than the cost of engineering for looser requirements at the outset – and then having to reengineer for tighter requirements further down the road for subsequent use cases.

Among the specific things operators need to be thinking about in the way that they procure and design for evolving the backhaul are the following:

- **The use of PTP Mini Grand Master clocks for clock distribution at the edge rather than relying on a Grand Master deeper in the core.** This reduces the number of hops in the network across which the phase synchronization needs to be supported.
- **The support of on-path PTP in all backhaul network elements.** Included in this are the IEEE-defined boundary clock (BC) for PTP, which acts as the master clock *vis-à-vis* subsequent downstream clocks for phase synchronization. Also included is the transparent clock (TC), which supports a correction field within the PTP packet that is dynamically updated with the real-time delay through each element. The ITU has defined two classes of error for phase synchronization. At minimum, operators should select from vendors that comply with the ITU's Class A, which allows for the introduction of no more than 50 ns of time error in a given network element. Better still, they should choose from vendors that support Class B, which mandates no more than 20 ns. At this point in time, support for phase synchronization is by no means a simple "tick box" exercise for vendors. Some legacy equipment won't support phase synchronization at all, for example. Operators continue to report considerable variation in the quality of different vendor implementations.
- In addition to the very stringent new phase synchronization requirements for CoMP, eICIC and LTE Broadcast, **phase synchronization will also be required for some other small-cell use cases**, including potentially some in-building scenarios. In the longer term, although the launch of the first 5G networks is still some years away, it's a safe bet that with cell site densification and spectrum efficiency requirements set to intensify, synchronization requirements will become still more exacting with 5G. Hence equipment deployed in the next couple of years will need to also be 5G-ready as far as possible.

The Case for PTP & an "Any Two Will Do" Synchronization Strategy

The above section focuses on the new challenges of phase synchronization and the role of the PTP protocol in supporting it. In concluding this section, it's worth noting that

phase synchronization isn't required *instead of* frequency synchronization for LTE Broadcast and LTE-A features; it's required *in addition to* frequency synchronization.

It's also worth noting that PTP is not the only technology or standard capable of supporting phase synchronization: Global Positioning System (GPS) can also support it (although SyncE can't – it can only support frequency synchronization). It's our view that all three primary global synchronization standards – PTP, GNSS and SyncE – have a long-term role to play in providing synchronization for 4G networks.

Figure 3: Mobile Operator Adoption Forecast for Synchronization Standards

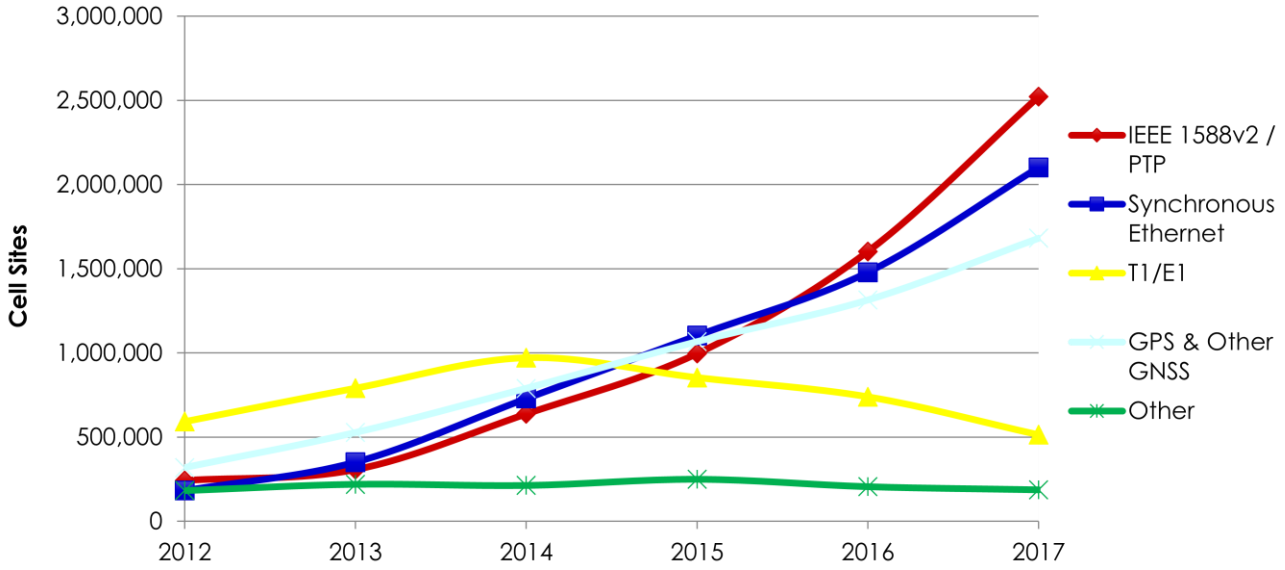

Source: Heavy Reading

When making a choice of synchronization standard for phase synchronization, PTP has some significant advantages over GNSS:

- GNSS is vulnerable to compromise by jamming.
- It is unsuitable for indoor, below-ground, and some street-level environments, because of the need for line-of-sight to the satellite constellation.
- There are also political concerns in many markets relating to the U.S. government's ownership of GPS. If anything, these concerns have been heightened in the wake of the so-called Snowden revelations.

Although PTP and GNSS can each support both frequency and phase synchronization, the trend among operators, which we expect to see consolidate further going forward, is to use two or even all three standards in the same network according to an "any two will do" approach.

Using more than one standard offers the benefit of providing backup in the event that one solution fails. In the case of phase synchronization, many operators also want to dedicate a technology exclusively to it, rather than have its resources tied down – and its performance potentially comprised – by having to simultaneously support frequency synchronization. Hence we expect combinations to become the norm.

Evolution of the 4G Backhaul Security Model

Security needs to feature prominently in mobile operators' monetization strategies in the coming years. Throughout much of the 2G and 3G eras, investment in security has been considered as a cost – something to protect against eavesdropping. And most of the security requirements (such as encryption) were bundled in with standard 3GPP-compliant products. As the LTE network is rolled out, security must increasingly be thought of as a revenue enabler. Specific mobile broadband revenue opportunities can be targeted, but on condition that specific security is also factored in.

The basic network security model changes significantly with the initial launch of LTE and is continuing to change as the market and technology evolve. The LTE architecture introduces a change in the security model at launch.

As shown in **Figure 4**, whereas 3GPP provides embedded encryption from the handset to the radio network controller (RNC) in 3G, in 4G there is no RNC equivalent. As a result, in 4G the embedded 3GPP encryption terminates at the eNodeB, which leaves the S1 and X2 interfaces unencrypted unless the operator chooses to add its own encryption to these links.

Figure 4: The LTE Architecture

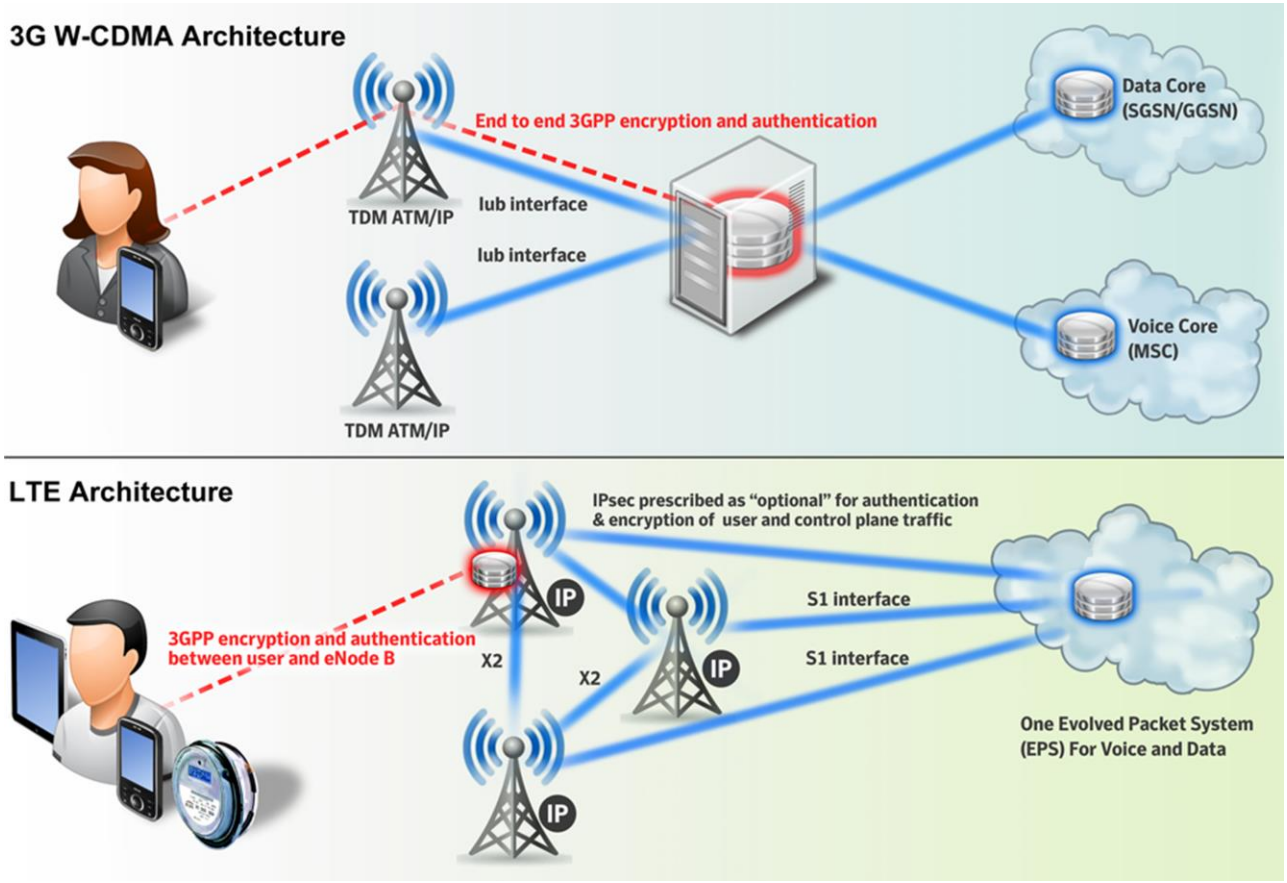

Source: Heavy Reading

Four other aspects of the current and evolving market environment fundamentally change the security model as LTE continues to roll out.

- Increasing volume, variety and sophistication of network security attacks of all kinds, including against the mobile network.
- A rapid acceleration in the number of network end points – whether they are small cells, "Things" deployed in the Internet of Things, or smartphones and other devices.
- An increasing number of the new network endpoints will be in physically accessible locations, hence vulnerable to tampering.
- In the last couple of years, mobile operator report an increasing trend whereby they are seeing the beginnings of new attacks coming from laptops and smartphones in the RAN. Where these are LTE devices, the attack traffic is coming in via the S1 and X2. In Heavy Reading's November 2014 survey on mobile network security 46 percent of a sample of 64 qualified mobile operator respondents stated that they are seeing compromised subscriber devices participating in DDoS attacks on their mobile network.

Using 3GPP's SEG for S1 & X2 Encryption

3GPP provides for encryption of S1 and X2 traffic using IPsec where the operator considers the backhaul to be "untrusted." This is a highly subjective term, so unsurprisingly the pattern of adoption around the world thus far is variable. From a regional perspective, there has been strong adoption in Europe and the Middle East.

In terms of the use cases, there has also been a lot of variation:

- Some operators, such as T-Mobile, have mandated a policy of encrypting all LTE traffic across all of their cell sites, across all of their properties.
- Some, such as Orange, are only using encryption at certain types of sites which they can consider untrusted – such as those where they use leased backhaul, or where they have potentially vulnerable small cells.
- Some are encrypting some but not all S1/X2 traffic, such as one Middle Eastern operator that is only encrypting its management traffic, but not its user or control plane traffic.

Additional Use Cases for the 3GPP SEG

3GPP's security model for LTE provides for instantiation of IPsec tunnels at the eNodeB and the termination of those tunnels in the Security Gateway (SEG) in the core of the network. There is ample opportunity to support other security capabilities with the SEG, as well. For example:

- **Authentication of eNodeBs using PKI.** At the point when eNodeBs are first installed, operators can adhere to the manual, so-called "shared secret," authentication model, whereby a field engineer manually enters a key at the cell site during the initial setup process; or, they can automate that process by leveraging the 3GPP-recommended Public Key Infrastructure (PKI) authentication architecture. This leverages a Certification Authority (CA) to authenticate a signed certificate that is pre-installed in the eNodeB at the factory (as supported by most RAN vendors). The trouble with the traditional manual inputting of shared secret keys is that it is prone to human

error and expensive from an opex perspective. That risk will only increase with the introduction of so many more end points in the network. Preferring the manual shared secret model over the more secure PKI authentication model will leave the network more exposed to so-called "man in the middle" attacks, whereby a rogue eNodeB is used to eavesdrop or carry out other attacks on the network.

- **Firewalling within the SEG.** In light of the growing risk of security threats originating in the RAN, at every point where S1 interfaces are terminated in the core, operators need to consider comprehensive LTE security with much the same suite of firewall, IDS/IPS and enhanced threat detection capabilities that they have always needed on the Gi interface between the core and Internet.

Distributing the SEG

As already discussed, one of the implications of the upcoming phase synchronization requirements for features such as CoMP is that operators need to consider distributing the Grand Master PTP function out to the edge to reduce the number of hops that synchronization packets need to pass through.

Evolving requirements for the X2 interface can lead to a similar conclusion in favor of distribution. Today, the X2 interface is typically hairpinned in the core. In the case of an LTE-A feature such as CoMP, however, the X2 latency can't be any more than a handful of milliseconds. Operators therefore need to consider pushing the X2 processing out from the core to the edge at an aggregation or hub site. And where the operator determines that it wants to encrypt the X2 interface, it follows that that operator needs to distribute the SEG functionality out to that same aggregation or hub site in the form of a SEG blade or other device.

Additional Backhaul Security Requirements

As the LTE network is densified, operators also need to look at additional security requirements. As discussed, PKI provides a means for the authentication of eNodeBs throughout a HetNet environment, but that doesn't embrace the authentication of other network elements that need to be deployed in exposed, potentially accessible areas. To ensure authentication of these other network elements, such as smaller form factor cell site routers, other authentication standards such as IEEE 802.1x need to be considered.

And while IPsec provides robust encryption of S1 and X2 traffic, it has a security vulnerability in the service provider environment, in that the protocol is designed to automatically reestablish an encrypted connection in the event that the connection is momentarily broken, including by an attacker seeking to introduce a network monitoring or other device. To ensure that any introduction of new devices onto the network path following a break in the connection is immediately flagged up, operators also need to consider using a protocol such as the IEEE's 802.1 AE or MACsec.

SDN in the Backhaul

Backhaul networks tend not to be front and center of the telecom industry's focus when it comes to transforming the network with NFV and SDN. Both nevertheless have a role to play in the evolution of these networks going forward.

The way the backhaul network operates today, network nodes typically have little or no visibility into conditions on other links in the network that could serve as an alternative path for the packets they are forwarding. Sometimes, the routing protocols that determine the paths of network flows have limited visibility into link conditions even on the links that they themselves are serving. This situation results in what is fundamentally inefficient routing of traffic, hence an inefficient use of backhaul network resources.

The introduction of an SDN paradigm into the backhaul – whereby the controller has link visibility and control of all network elements and links in software – promises a step change in the speed with which network inefficiencies can be identified and corrected, in a manner consistent with real-time variations in the mobile operator's service requirements. SDN can therefore be used toward the end of routing optimization in the backhaul network, including as the network densifies with small cells. Optimal paths can be calculated according to a number of different criteria, whether it be network performance data; utilization of other network elements such as a cell site or EPC element; or poor weather-driven fluctuations in the availability of microwave links.

Routing Requirements for Backhaul Evolution

Some mobile operators still prefer L2 switching to L3 IP/MPLS routing in the backhaul. The clear trend of the last few years, however, has been for operators to increase their reliance on L3. IP/MPLS has grown it share to become the preferred technology of mobile operators in the aggregation or hub layer of the backhaul. Recent years have also seen growing adoption of L3 in the access layer.

Key factors driving many mobile operators to require L3 functionality in the backhaul are as follows:

- **Assumptions about network layers, topologies and technologies.** Faced with a highly layered network with multiple points of concentration – perhaps leveraging different physical fiber and microwave transport pipes, multiple vendors and even a combination of self-provided and leased backhaul – many operators prefer L3 as a means of providing a single unifying protocol running across the top of all this complexity.
- **Assumptions about bandwidth requirements, growth in network end points and path diversity, as well as assumptions about emerging performance requirements such as latency.** Making aggressive assumptions about traffic growth, cell site growth and the need for more intelligence and more dynamic path diversity for LTE and LTE-A is driving many operators toward L3 in the backhaul.

Emerging Backhaul Router Requirements

The emerging demands outlined in the previous sections point to some redesigning of the backhaul network from a routing perspective, to a requirement for new features across vendors routing portfolios, as well as a rebalancing of demand for different router types.

High-Capacity Core Routers

The characteristics and requirements for high-capacity core routers deployed where the backhaul network meets the core is arguably least affected. The capacity requirements in those devices will continue to scale up, albeit those requirements

tend to be driven as much, if no more, by the demands of the wireline network rather than the mobile network itself. The roadmaps for these core routers certainly need to align with the new LTE feature requirements such as support for PTP phase synchronization and the 3GPP SEG but for the most part their role in the network isn't likely to undergo a great deal of change.

1U Cell Site Routers

The changes in requirements for 1U cell site routers are likely to be somewhat greater. Small cells drive requirements for miniaturization as well as outdoor hardening requirements such for the IP65 standard. With the growing adoption of Single RAN base stations which can support 2G, 3G or 4G in a single platform, some mobile operators believe that the L3 capability built into these platforms means that a dedicated cell site router at each and every site may no longer be required at every site.

Use cases some operators may be considering restricting the use of 1U devices to include delivering L3 VPNs to enterprises; where the device serves the indoor unit (IDU) for a microwave outdoor unit (ODU), and is therefore powering the radio head via Power over Ethernet (PoE); and where it is serving as a hub for connecting to small cells surrounding the macro cell and is driving the small cell backhaul radios via PoE.

Aggregation or Hub Site Routers

The greatest change in requirements is likely to come in the case of aggregation or hub site routers. As this paper has already pointed out, the demands of security and phase synchronization already make the case for PTP Grand Master and 3GPP SEG functionality currently hosted in the core to be pushed out to the edge, which means pushed out onto hub site routers.

Figure 5: Distribution of Functionality From the Core to the Edge for LTE-A

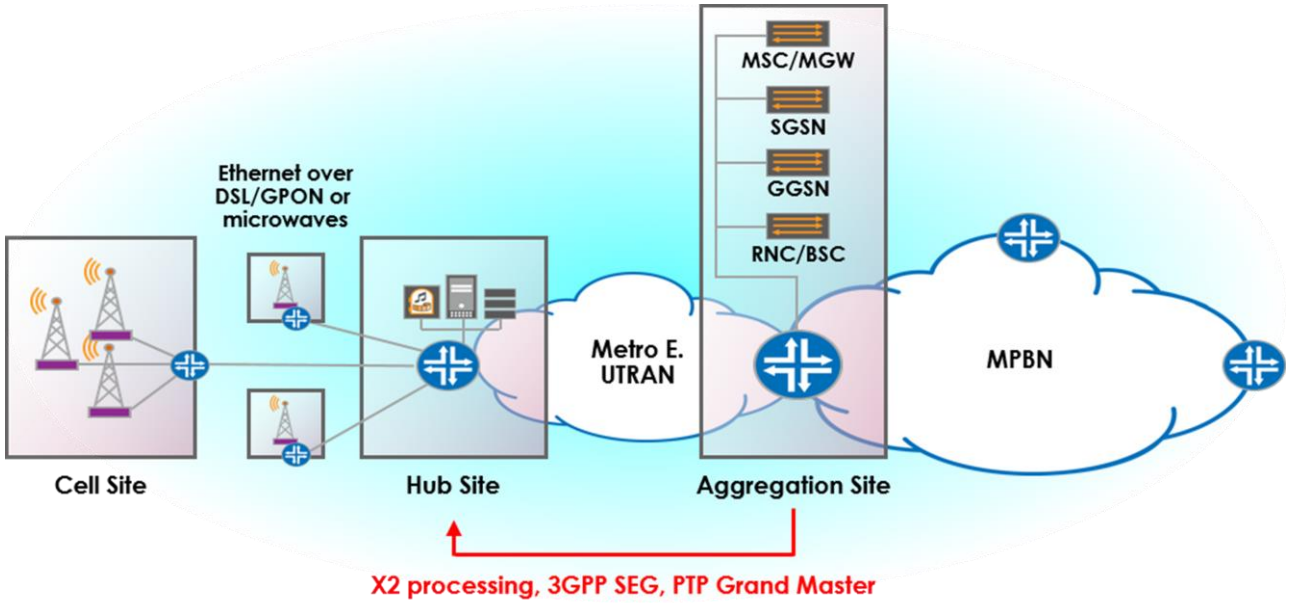

Source: Heavy Reading

There is more to it than that, though. As already discussed, in the case of CoMP, the latency requirements for X2 are of the order of a few milliseconds. That also drives a logic of network feature distribution where handling of the X2 is concerned. So whereas the X2 interface is currently turned around in the core in the LTE network, with LTE-A there will be a strong case for supporting the X2 closer to the neighboring cell sites in order to reduce latency. As with the case for distributing the Grand Master and the 3GPP SEG, this requirement points to a case for handling the X2 interface in an aggregation or hub site, rather than at the core.

As the network densifies with smaller cells, more hub sites are going to be required than previously. In addition, some existing hub sites that have supported the first phases of macro-cellular LTE traffic will not be well suited to the emerging hub site requirements for synchronization, security and X2 support with LTE-A and LTE Broadcast. These requirements for distribution of functionality out from the core also align well with ETSI's Mobile Edge Computing (MEC) standardization effort, which is seeking to drive the creation of an IT service environment for application developers and content providers at the edge of the mobile network.

Summary

Emerging capacity and performance requirements with new LTE services and business models, and with new LTE-A network features, are placing brand new demands on the backhaul network. The topology, engineering and features supported in the backhaul to date have generally served the operator well, but these new requirements – revolving around latency, synchronization and security – mandate a root-and-branch review of network design and network equipment functionality.

Operators need to focus not just on selecting the right vendor partners for evolving the backhaul; they will also need to undertake a review of their existing backhaul architecture and topology, with a view to refreshing the design for an optimal balance of core, aggregation and hub sites to support these new requirements.

About Juniper

Juniper Networks (NYSE: JNPR) is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.juniper.net or connect with Juniper on Twitter and Facebook.