White Paper

IPsec Deployment Strategies for Securing LTE Networks

Prepared by

Patrick Donegan
Senior Analyst, *Heavy Reading*

www.heavyreading.com

On behalf of

RadiSys.

www.radisys.com

May 2011

TABLE OF CONTENTS

LIST OF FIGURES

I. Network Security Risks in an Interconnected World

The mass-market penetration of low-cost and increasingly sophisticated computer software, Internet connectivity and broadband IP networks is probably the key economic and social transformation of the early 21st century. But as is often the case with any far-reaching and rapid transformation in the way people live and work, the advantages of our increasingly interconnected world are not being gained without also creating some new challenges. Ask those responsible for network planning and operations in any of the world's enterprises, public sector organizations and fixed and wireless carriers what the downside to our increasingly interconnected world is, and many of them will point to the increased vulnerability of businesses, governments and consumers to network security attacks.

To an extent, telecom carriers have always had to live with low-level criminality relating to misuse of the network, whether it be unpaid bills, illegal eavesdropping or theft of network minutes. But the ascendancy of IP as the world's predominant communications protocol has driven up security attacks. IP's inherent flexibility and mass-market accessibility enabled early generations of Internet hackers to generate spam and launch denial-of-service attacks on network infrastructure and more recently to distribute malware to PCs and other computer processor-driven machines.

Recent years have seen the threat level rise further. Early "hobbyist" hackers wanted to test themselves against network designers and disrupt services as an intellectual challenge. But the last few years have seen a marked rise in the incidents of organized crime, terrorist organizations and nation states actively using advanced networking and computing technologies to carry out large-scale theft of service; extortion; theft or distortion of highly sensitive personal, commercial or classified government information; and even disablement of a country's critical power, water, telecom or financial services infrastructures.

Some high-profile attacks of the last couple of years include:

- Attackers targeting financial service companies' networks to get customers' PINs in order to withdraw cash directly from their accounts.
- Attackers holding organizations for ransom, for example by launching or threatening to launch attacks on company websites that disable a company's ability to take payments.
- Phishing attacks driving users to disclose their credit card or bank account details.
- Large-scale theft of proprietary corporate information by means of distributed malware.
- Attacks on the Internet infrastructure of the government of Georgia, disabling it at the time of the 2008 tensions with neighboring Russia.
- The development of the extraordinarily sophisticated Stuxnet virus that infected and disabled Iranian power plants during 2010.

Richard Clarke, a partner in Good Harbor Consulting and a former Special Advisor to the U.S. President on Cybersecurity, claimed during the RSA Conference Europe in October 2010 that business and consumers are losing "billions of dollars a year" to cybercrime.

Consistent with the almost infinite possibilities created by IP, as well as the need to overcome the barriers that the networking industry has put in place to combat them, what are known as "attack vectors" in security circles – the types of attack – have grown increasingly sophisticated in recent years, with the results that we now see almost daily in media reports. To compound matters, while the level of sophistication of attacks may be increasing, in many ways hacking is actually getting easier. In its July 2010 report, Verizon Business published figures showing that 85 percent of attacks on enterprises were "considered not highly difficult." Botnet capacity can be leased from hackers via the Internet using just a credit card. Security specialist Symantec reports seeing botnets being sold online for as little as 4 cents per member bot.

II. The Mobile Network in the 2G & Early 3G Era

Until fairly recently, the mobile network has tended to feature at the margins of concern about new network vulnerabilities. The most high-profile and damaging attacks have tended to be focused on enterprise and ISP infrastructures. And there are very good reasons for that.

The mobile industry has led the rest of the telecom industry in some key aspects of security, rendering legacy mobile networks in many ways more secure than wireline networks. After all, it was GSM that gave us the world's first mass-market voice telephony device complete with end-to-end authentication and encryption. Only very recently – a remarkable 20 years after the launch of GSM – has the original GSM A5/1 encryption been shown to be vulnerable, with equipment that potentially permits eavesdropping of GSM traffic over the air. And consistent with the mobile industry's strong pedigree in anticipating future security threats, radio base station vendors are readying to deploy the 3GPP's later A5/3 algorithm to counter that exposure.

The same tight end-to-end authentication and encryption principles were carried over from 2G to 3G. Hence in 3G today, user traffic is authenticated and encrypted all the way from the handset or other end-user device all the way through to the Node B and on to the radio network controller (RNC). At the RNC, user traffic can be unencrypted because it was – and still is – considered deep enough in the operator's network to be secure.

Also, while mobile operators have invested heavily in IP networking assets in the relatively secure core of their networks, they have only very recently started deploying IP backhaul out to their cell sites. From an attacker's perspective, although large parts of an IP-oriented enterprise or ISP network have been open to attack for many years, the more remote parts of a mobile operator's 2G and 3G network have had the additional protection of TDM transport – which, having been designed from day one as a telecom-grade transmission protocol, is a very much more inherently secure protocol than IP.

III. The Mobile Network in the Mobile Broadband Era

Market and technology trends are conspiring to ensure that while many mobile operators reap the rewards of their investments in mobile broadband in the form of stabilizing ARPU declines or even ARPU growth, their transition to becoming fully-fledged ISPs in their own right is inevitably bringing them closer and closer to the global community of Internet attackers. And the evidence for that is increasing quarter on quarter. Hence when the U.K. government's "Operation White Noise" of November 2009 simulated a widespread failure of the national telecom infrastructure, the simulation included the PSTN, ISP and mobile network infrastructures.

Whether carried out for financial or political gain, or just for sport, security attacks depend on scale for their effectiveness. And the balance of scale in the data networking industry is starting to shift away from the wireline network toward the mobile network. As shown in **Figure 1**, depending on exactly how you measure it, the global number of mobile broadband connected devices either recently outstripped the number of fixed PCs or will do so at some point in the next 12 months.

Figure 1: Fixed & Mobile Broadband Subscriber Forecast

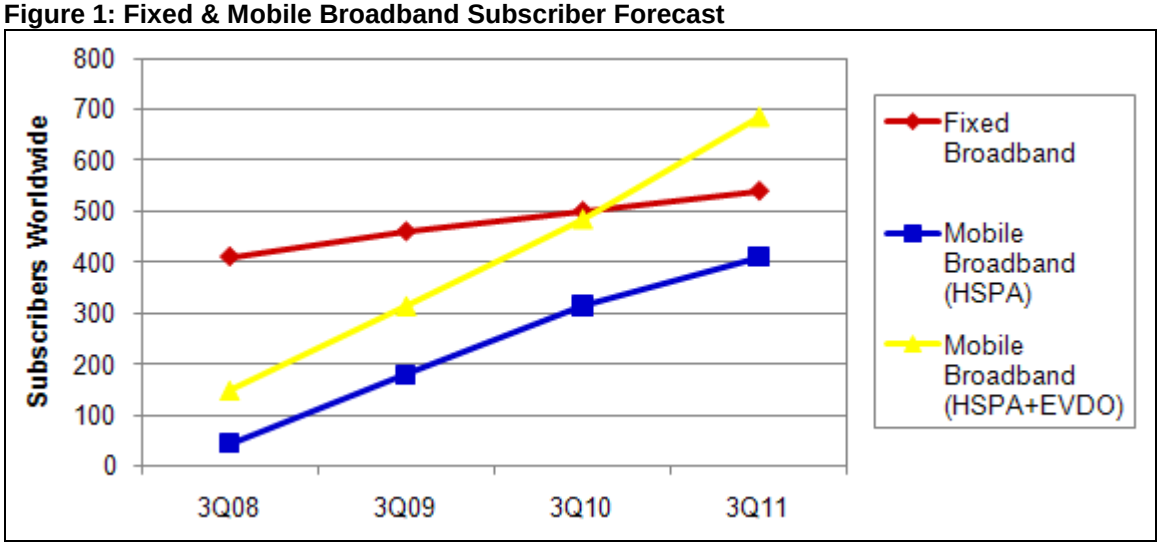

Source: Heavy Reading

An analysis of the browser and operating system (OS) markets for fixed and mobile broadband end-user devices provides a more granular perspective on why the fixed network has traditionally been more vulnerable. There is a much more fragmented market for key elements in the mobile network value chain, as compared with the more consolidated structure in the PC market. In the PC market, for example, the three variants of Microsoft Windows have more than 85 percent market share; while in smartphones, market share by OS is much more fragmented. And in the browser market, Internet Explorer is much more dominant in PCs than Opera is in smartphones.

It is important to consider this market structure from the perspective of a network attacker. Based on the volume of attacks that are directed at a particular browser or OS, hackers have traditionally preferred the consolidated market structure of the PC environment. Quite simply, it gives them a bigger target. The more fragmented market structure of the mobile environment has traditionally presented a higher barrier to entry and a higher barrier to "success." Although the mobile market will continue to remain more fragmented in this respect for the foreseeable future, in terms of absolute numbers the number of mobile devices supporting any one OS or browser will soon rival or exceed their stationery PC counterparts.

Predictably, in recent months the respective security features the Android and iPhone OSs do or don't offer has become increasingly important as successive malware attacks are reported in the media. In March 2011, for example, Google had to remove more than 50 rogue applications from

its Android Marketplace, many of which would have been capable of taking over Android smartphones and stealing data or sending expensive text messages. Google announced that it would remotely disable the malware over the air for those affected, which according to some reports amounted to more than 250,000 users. This is but one example, albeit a particularly powerful and recent one. The iPhone, Symbian and all the other smartphone OSs all have their vulnerabilities.

As shown in **Figure 2**, a large proportion of carrier respondents to the Arbor Networks Worldwide Infrastructure Security Report of 2010 cited being vulnerable to DDoS attacks across many elements of their mobile networks, with web services, data signaling gateways, mobile packet core and RAN infrastructure having proved to be the most vulnerable to date.

Figure 2: Wireless Network Infrastructure Affected by DDoS Attacks

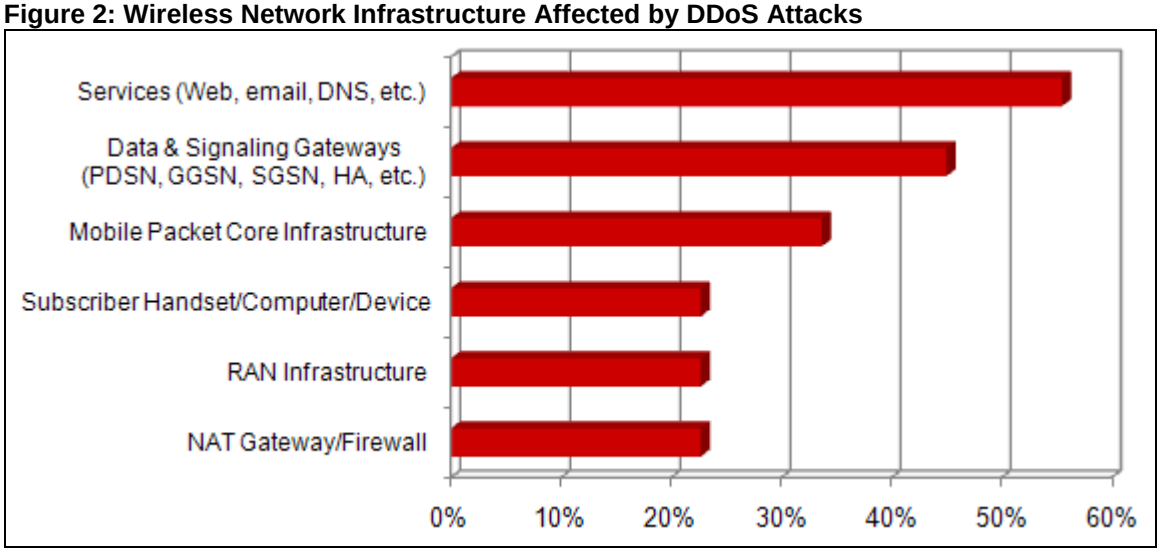

Source: Arbor Networks

If 3G were to be some kind of final destination in the evolution of the mobile network, then it is already on a collision course with the type of attack environment that the wireline and enterprise network environments have contended with for many years. But that's not the case, of course. Several operators have already launched LTE, and in that transition from 3G to LTE, operators will be exposed to that same, familiar historical pattern once again. LTE will deliver new capabilities and a better cost profile – but it will also expose the operator to new security risks.

IV. New Opportunities With LTE – But New Risks, Too

LTE will certainly reduce the cost of the network and improve the end-user experience but it will also transform the security environment. In one or two respects, security is actually enhanced in LTE as compared with 3G in that two different encryption algorithms on the LTE air interface are designed to be supported simultaneously rather than sequentially (years apart) as in 2G and 3G.

For the most part, though, the way that 3GPP fully embraces IP in LTE and the Evolved Packet System (EPS) for the first time drives the building of LTE networks, which are inherently less secure than traditional 2G and 3G networks, and which therefore require a lot of additional new security features and capabilities to close off those new exposures.

The starting point for understanding this new level of attack exposure is a high-level comparison of the 3G and LTE network architectures, shown in **Figure 3**.

Figure 3: 3G & LTE Network Architectures

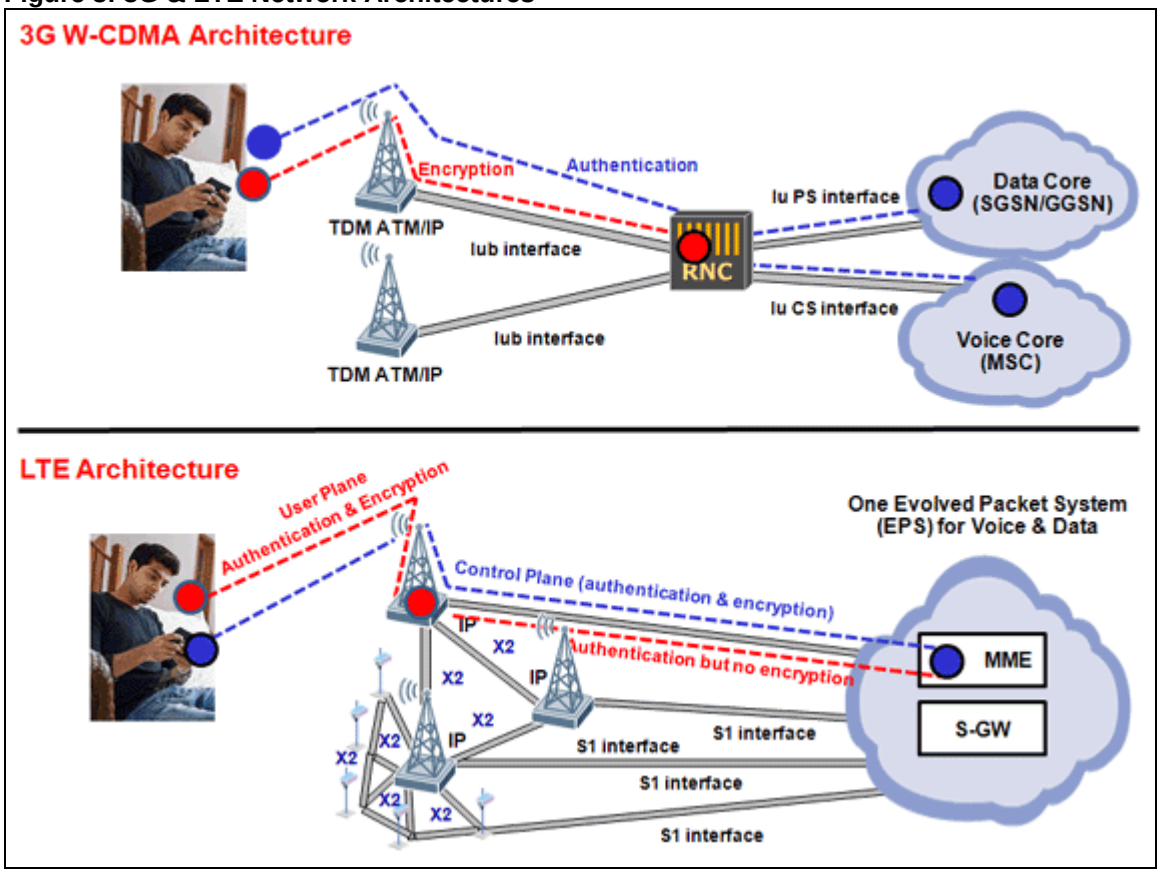

Source: Heavy Reading

As shown, the LTE architecture is much flatter and much more IP-centric. This has a number of security implications, particularly where the backhaul network is concerned.

- Over time, the LTE network will move to a much larger number of cell sites per user – or more accurately, per end point – than has been the case for 2G and 3G. From a global installed base of 2.7 million cell sites at the end of 2010, *Heavy Reading* is forecasting an increase to 3.8 million by the end of 2015. We are starting to see the adoption of small cells – femtocells and picocells – for 3G, but this trend will accelerate markedly over time, with LTE driven by the need to lower costs and increase bandwidth per subscriber. From a security perspective, the growth of small cells will pose two new challenges. More net-

work elements of any kind to manage is always more challenging. And smaller cell sites are inherently less secure, in that when they are deployed in shopping malls and on lamp posts and other locations, the operator typically cannot have the same kind of physical security that would be in place at a conventional macro cell site. The active electronics are inevitably more accessible to an attacker at a small cell site because physical barriers are much more difficult to put in place, and investment in personnel is either not feasible or not justifiable from a cost perspective.

- IP backhaul is mandatory in LTE, whereas it is optional in 3G, as previously mentioned. TDM and ATM are more esoteric and generally less well understood by attackers; IP of course is common, and commonly understood, everywhere.
- In LTE, because the Radio Resource Controller is managed by the eNode B and the Mobile Management Entity (MME) in the core, the RNC node is eliminated altogether. In the event that an attacker is able to penetrate the cell site, that attacker is therefore potentially able to have a straight shot at the core of the LTE network, the EPS, whereas in 3G his path to the core is blocked by the RNC.
- With LTE, 3GPP has created many more signaling and bearer paths between network elements than there are in 3G. LTE has the potential to support an X2 interface, which enables peer-to-peer signaling between cell sites for the first time. According to the NGMN Alliance, in LTE each eNode B may be associated with as many as 32 X2 interfaces. 3GPP also enables each eNode B to connect with several core elements – i.e., several MMEs or several security gateways – simultaneously for better load balancing or better performance. According to the NGMN Alliance each eNode B may be associated with as many as 16 S1 interfaces back toward the core. By comparison, in 3G each Node B is typically associated with just the one RNC. There are scenarios in which an attacker that is able to penetrate a cell site could therefore reach many more network elements with an attack on the LTE network than would traditionally be possible in the 3G network, by virtue of all these new logical adjacencies.
- Because in LTE the Radio Resource Control (RRC) terminates in the eNode B rather than the RNC as in 3G, the encryption of the user traffic also terminates in the eNode B. In LTE, therefore, for the first time the backhaul (S1 and X2) becomes a potential security exposure for user plane data; network control signaling (S1-MME) is equally exposed without intervention.

4.1 Other Trends in Backhaul Evolution Introduce Further Security Challenges

The previous section depicted how the LTE architecture creates a fundamentally different security environment at the cell site and across the backhaul network. But into that mix, other global trends relating to the way the backhaul is designed and delivered need to be added.

One issue is that of network sharing between mobile operators. This is a trend that has accelerated in recent years and is set to accelerate still further going forward. Whether it be sharing passive elements at the cell site, which requires more personnel from more different companies to be given access privileges, building common backhaul networks or even sharing active base station electronics, what might be considered a "walled garden" approach to building out the backhaul network is having to evolve in many cases into a more porous, shared network build model with self-evident security implications.

A related issue concerns the fact that in the transition to packet backhaul (whether for 3G, LTE or both) some operators are using E-LAN rather than E-Line wholesale Ethernet services, which has security implications, particularly when the backhaul service is leased from a third-party wholesaler. Whereas an E-Line service has the same characteristics of a private circuit that mobile operators are used to in their backhaul network, with an E-LAN service, there is always the possibility that a third-party end point will be introduced and share that E-LAN service with the operator's

backhaul service. Where an E-LAN service is managed extremely well at all times by the wholesaler, there is no security issue. But where it isn't, the mobile operator's vulnerability to having their traffic intercepted by a third party is that bit greater than in an E-Line environment.

4.2 Cell-Site Break-Ins & Attacks on the S1 & X2: Fact & Fiction

It would be nice to be able to draw a clear dividing line between attack scenarios on the mobile network that are realistic and need to be defended against, and those that can comfortably be dismissed as fanciful or fearmongering. Unfortunately, the only people really qualified to make that distinction are the high-end attackers themselves and, in any case, some of what even they consider fanciful today has a way of becoming plausible a few years down the road.

And it's important, albeit unpalatable, to consider that attackers are not always unknown external third parties. Some of the most effective attacks on telecom networks have been carried out by employees of the operator's partner companies (vendors, systems integrators, a wholesaler or other mobile operator engaged in a network sharing activity); ex-employees of the operator; or often present-day employees of the company carrying out insider attacks, leveraging their unique knowledge of the network and its weak points.

Consider first the security of the operator's sites themselves. We don't actually know how common physical intrusions into cell sites are – the operators certainly have every incentive not to talk about them – but we do know they are by no means unheard of. In February 2011, for example, hundreds of thousands of Vodafone UK's customers suffered a major outage for nearly half a day, resulting from a break-in at one of the company's central network facilities in the south of England. The motive on that occasion was theft of high-value networking equipment at the site, but the motivation and objective could easily have been different.

Anecdotal evidence would suggest that such break-ins are relatively rare. That said, the previous sections of this paper have demonstrated that the incentives to attack the mobile network are clearly increasing, while in the examples cited many of the barriers to a "successful" attack are clearly being lowered. If the rate of physical intruder attacks on cell sites is indeed as low as is widely assumed, some part of the explanation for that may be that up to now it simply hasn't been worthwhile from an attacker's perspective.

If a major central facility of Vodafone's can be taken down by intruders, you have to figure that so could one of its more remote cell sites, where the physical security is less rigorous. And then consider the minimal physical security that can be applied to small cells as they are rolled out. It then becomes increasingly difficult to consider the idea of an attacker being able to probe for weaknesses in an operator's network of thousands or tens of thousands of sites and successfully penetrate one of them and dismiss the idea of this happening as fanciful.

Assuming that an attacker gaining access to a network element isn't fanciful, since the LTE network environment is all-IP there is the potential for an attacker to plug a laptop into an eNode B or tap the connections at a cell site and gain access to its signaling to a number of other nodes in the network. Smart security processes on the operator's part (access restrictions, password protection, and the like) can erect significant barriers to that, but human error can still leave operators vulnerable even on this front, to attackers from both inside and outside the company.

V. The Option of IPsec for Enhanced LTE Security

Consistent with the mobile industry's historical leadership in network security issues, 3GPP provides a means of closing the security exposure on the S1 and X2 interfaces. Within its Network Domain Security (NDS) architecture, 3GPP recommends doing this via the use of IPSec, which is an established solution for enabling authentication and encryption of IP traffic that has been widely deployed in enterprise and wireline networking circles. 3GPP recommends a model in which IPsec tunnels are instantiated at the cell site and carry both bearer and signaling traffic across the backhaul, where they can be unencrypted in the core network by a security gateway. IPsec is already used in femtocell, I-WLAN (TTG) and UMA/GAN deployments. For LTE, most – if not all – major RAN vendors now support instantiation of IPsec tunnels in their eNode B products.

As shown in **Figure 4**, the use of IPSec is not mandatory across the S1 and X2. Rather, 3GPP provides for its usage where an operator deems that its backhaul is "untrusted." 3GPP goes no further than that, leaving it to the operator itself to determine what "untrusted" means. It could mean that the backhaul is provided by someone else; it could mean that it's shared with someone else; or it could mean that it is built out using a Layer 1 technology that is deficient according to the operator's own definition of "trusted" in some way. From an authentication perspective, 3GPP enables the mandatory IKEv2 authentication feature of IPsec to be used.

Figure 4: IPsec Options in LTE Backhaul

PLANE	REQUIREMENT FOR S1 (TRUSTED)	REQUIREMENT FOR S1 (UNTRUSTED)	REQUIREMENT FOR X2 (TRUSTED)	REQUIREMENT FOR X2 (UNTRUSTED)
User Plane	Optional	Mandatory (Tunnel or Transport mode)	Optional	Mandatory (Tunnel or Transport mode)
Control Plane	Optional	Mandatory (Tunnel or Transport mode)	Optional	Mandatory (Tunnel or Transport mode)
Management Plane	Optional	Mandatory (Tunnel mode)	Not required; X2 does not carry management traffic	

Source: Heavy Reading/3GPP

This question of whether or not IPsec is needed is one of the most contentious for network planners as they prepare for the deployment of LTE. Operators in North America have initially launched without it, but *Heavy Reading* is familiar with large European and Asian operators that intend to deploy it from launch.

As shown in **Figure 5**, 83 qualified mobile network professionals asked about this in our December 2010 annual survey on mobile backhaul gave a very wide variety of responses. The survey asked, "For the first three years following the launch of LTE, to what extent do you expect that IPsec will be needed between the LTE cell site and the LTE core?"

Figure 5: Operators Are Uncertain of the Need for IPsec

RESPONSE	ALL OPERATORS
All cell sites will need IPsec implemented	20%
At least half of all cell sites will need IPsec implemented	13%
A subset of cell sites will need IPsec implemented	19%
IPsec will probably not be needed in the backhaul	17%
IPsec will definitely not be needed in the backhaul	1%
It's still unclear at this stage	29%

Source: Heavy Reading Annual Mobile Backhaul Survey, December 2010 (n=83)

As the figure shows, about half of mobile operators reckon IPsec will be needed to some extent. About half are unclear, but inclined to think it won't be needed. Only 20 percent currently believe it will be needed at all cell sites.

5.1 Operator Uncertainty About the Case for IPsec

Figure 5 captures mobile operator opinion on the case for IPsec very well. Operators are divided on whether they are going to need to deploy it or not. Several factors explain this diversity of opinion. To begin with, the effect of the growth of mobile broadband and the implications of the LTE architecture on network security are not fully understood across the industry. Many operators are undertaking detailed studies of the implications but some are yet to develop plans for LTE at all. Put simply, most operators have not yet arrived at a detailed and conclusive risk assessment of exactly what "untrusted" means where the backhaul network is concerned and what exactly the risk is of "trusting" the backhaul and foregoing IPSec.

Opinion on IPsec also tends to vary by type of operator as well as by country. Marked variations in national security regulations have a significant effect on an operator's outlook. Its customer base does too. An incumbent operator that is the market leader in business and government subscribers tends to be more persuaded by the case for IPsec than a challenger operator whose subscriber base is largely made up of low income consumer groups like students.

The other key factor that affects an operator's outlook on IPsec is its perception of how much it is going to cost to implement and manage, how much additional complexity it's going to add in the network, whether that additional complexity is itself going to impair performance and efficiency in the network, and whether the additional cost in capex and opex is going to be worth it.

Compounding the challenge for the network planner is the fact that a decision about IPsec can't be taken in isolation. It can only be taken in the wider context of other challenging decisions relating to the LTE backhaul environment. For example, the effect of IPsec in the network is heavily dependent on whether or not the X2 interface is implemented, and if so, precisely how it is implemented. It is also heavily dependent on how the operator seeks to support QoS across the backhaul; on what packet synchronization technique is chosen and how it is implemented; and on whether the operator opts to deploy its MME and security gateway elements in the core according to a conventional, centralized model or a more IP-centric, distributed model.

Changing any one of these variables alters the model for IPsec deployment, and since many operators have yet to finalize their decisions in one or more of these areas, this explains a number of the uncertainties that many operators still have. The rest of this paper is dedicated to exploring some of the challenges associated with deploying IPsec in the LTE network and presenting a real-world outlook on how many of the real and perceived complexities can be stripped out. Doing so can enable operators to successfully close off this potentially significant security exposure in the LTE backhaul without burdening the network with unnecessary cost and complexity.

VI. Simple Models for Initial IPsec Deployment

To begin with, there's nothing to say that IPsec necessarily needs to be deployed at every single one of the operator's cell sites. As shown in **Figure 5**, a third of operators participating in *Heavy Reading*'s December 2010 mobile backhaul survey reckoned that IPsec will need to be deployed at either "a subset" or "at least half of cell sites."

One approach might be to deploy IPSec only at those sites that are considered most vulnerable from a physical security perspective – macro-cells in remote areas or small cells, for example – or those that serve particularly sensitive business or government sites. That could be seen to offer greater simplicity than a network wide deployment, although operators should also carefully consider the potential opex associated with having different kinds of security environments for encryption and authentication at different sites rather than having a single uniform approach. Having two different manual and automated authentication procedures for authenticating cell sites could get very costly and complex, for example.

A common anxiety about IPsec is that IPsec headers generate a substantial overhead – estimated by the Broadband Forum to be approximately 15 percent when implemented across both the control and user planes. There's no disputing this and it has to be addressed, but it should be relatively easy to manage in most networks, in the context of the other factors that backhaul planners are addressing.

Many operators are already used to living with a 30 percent overhead generated by the use of pseudowire encapsulation of legacy TDM and ATM traffic across their new packet backhaul deployments. Moreover, over time operators can be expected to reduce their dependency on pseudowires, particularly out at the cell site, as pure Ethernet implementations become predominant. A lot of operators are also putting very high bandwidth out at the cell site – 100 Mbit/s per cell site, or even more in some cases – so the bandwidth bottleneck is in the process of being relieved. IPSec compression may be applied to mitigate overheads as well; however, IPSec compression is not currently part of the 3GPP recommendations.

6.1 Operators Are Pursuing a One-Tunnel Model at Launch

One of the biggest concerns about using IPsec in the backhaul is the operational cost of managing huge numbers of IPsec tunnels in the network. Experience from the enterprise and wireline environment does suggest that this can be costly on opex but also imposes a substantial drain on system processing resources. In particular operators fear the complexity of having to manage multiple IPsec tunnels for multiple X2 as well as S1 (MME and U) interfaces across the network.

Consistent with this concern, *Heavy Reading*'s research has shown that those mobile operators that are moving forward with deploying it are pretty much universally looking at a single IPsec tunnel deployment model in the initial launch phase. What that consists of is one single tunnel being instantiated at the cell site – typically by the eNode B itself – and all of the user plane and control plane traffic being transported across that single tunnel to the core network, where it is unencrypted by the new 3GPP-defined security gateway. This pretty much permanent IPsec tunnel across the LTE backhaul network provides the stability that the operator requires. Critically, it avoids the highly dynamic instantiation and tearing down of tunnels that operators are wary of at the outset.

It's critical here to note that the X2 interface is not even mandatory for LTE launch. Most of the operators that have launched LTE so far do not have the X2 in service yet. For now, they are managing perfectly well with handover between LTE cells being managed across the S1. Moreover, even if the operator does want to launch LTE with the X2 in service, and even if it does want to apply IPsec to each X2 interface, X2 and S1 would be within the same network address space and domain and may be routed over the same IPSec tunnel as S1.

6.2 Implementing IPsec With QoS

QoS has been a long time coming in mobile networks but *Heavy Reading* research shows consistently that mobile operators are typically looking to implement between four and five QoS levels in across their network, especially with LTE. IPsec will typically need to be deployed with that in mind, since encrypting traffic affects the way that QoS markings can be read and adhered to by intermediate network elements and the current one-tunnel model for IPsec deployment also precludes deploying one tunnel per QoS class.

In the longer term, operators will ultimately have to look at deploying a single traffic class per IPsec tunnel. In the meantime, operators are looking at various different ways of addressing QoS in a way that better aligns with the one-tunnel model that they are assuming for the early deployment phases.

One approach is to over-provision the network. Where a backhaul wholesaler is cost-effective and flexible enough, bandwidth can be leased at a Committed Information Rate equal to the bandwidth of the tunnel. Then only the mobile operator's equipment needs to worry about classification and queuing, so the intermediate transport network doesn't need to worry about QoS.

Another option involves exploiting options in the types of encryption that IPsec allows. In addition to encrypting the whole packet, IPsec allows a model whereby the original header can be kept and only the payload is encrypted. If the latter model is implemented, then intermediate DHCP markings in the original IP header can be preserved for the intermediate transit across the S1.

6.3 Implementing IPsec With Packet Synchronization

Since it introduces additional processing and overhead in the network, network planners are wise to consider the requirements for deploying IPsec when it is implementing one of the new packet synchronization standards that are being implemented in the backhaul and have proven challenging to deploy without IPsec, let alone with it. In fact, deployed the right way, IPsec should not affect synchronization in the network at all.

Take the example of the IEEE 1588v2 standard, which many operators in Europe, Middle East, Africa and Asia are likely to deploy as they remove E1-based synchronization from their networks altogether in the transition to packet backhaul. Latency, jitter packet delay, and particularly packet delay variation in the network are what typically disrupt efficient implementations of the 1588v2 standard. The bursty nature of LTE's packet backhaul environment can render this challenging.

Several major vendors are supporting a means whereby synchronization traffic can be excluded from the IPsec tunnel and transmitted along an express path where the synchronization traffic is marked up with the highest Ethernet and DiffServ codepoint QoS level. This enables the synchronization packets to bypass all the standard queuing mechanisms in the switches and routers along the way. Leaving the synchronization traffic unencrypted would still count as a security vulnerability, but it would be a very minor one in the scheme of things and is an approach that is recommended by some of the leading LTE RAN vendors. Not only would an attack on the synchronization traffic require a very high level of sophistication, it would also typically result in a very gradual loss of synchronization that the operator would have plenty of time to respond to before it affected the user experience.

6.4 Authentication in LTE: New Applications of a Mature Technology

As previously stated, 3GPP allows operators to continue manually configuring their LTE cell sites according to the same authentication model used in 3G if they wish. Or they can obtain authentication based on IKEv2 by implementing IPsec. IKEv2 requires automated certification via a certification authority, generally assumed to be the mobile operator itself.

The prospect of rolling out PKI in the network does worry some operators, primarily because they've never used it before. But this is the wrong starting point. It assumes that the alternatives are easier when they're not. Certainly manual configuration processes are easier because they are more familiar. But it's worth recalling that they take hours, sometimes days, are prone to human error, and increasingly unmanageable as the number of network elements grows. The model is not very portable into the more dynamic era of small cell deployments, a model wherein every end point needs to have pre-shared keys or certificates manually installed into each device. A model in which the certificates expire after as little as two years is also not a model that is going to scale well in a large network, especially one where the cell density increases over time.

Leading RAN vendors have put significant development resources into automated certificate enrollment and certificate management through the use of Certificate Management Protocol (CMP) as recommended in the 3GPP specifications. These enable eNode Bs to be securely auto-configured and have their certificate managed and even replaced over the lifetime of the eNode B. Additionally, the security gateway is also mandated to support CMP for its certificate and certificate validation can be handled by OCSP. Although many operators believe they don't have the skill-sets in-house to support this adoption of PKI without investing in new people, in fact many of them do have such resources, albeit in the IT side of the company.

VII. Advanced Models for IPsec Deployment

While the vast majority of operators looking to deploy IPsec are not looking any further than the one-tunnel model in the interests of simplicity, it's clear that there are a number of other scenarios in which more than one tunnel may be needed. Among them are the following:

- **At some point as LTE traffic volumes grow, operators may want to put different types of traffic in different tunnels**. They may want separate tunnels for control and user plane traffic, or for different QoS classes. They may want to do that on the basis that there is greater risk from a resiliency perspective in having all traffic in one tunnel rather than across several on the grounds that if one tunnel can go down in the latter model and the rest of the traffic will be unaffected.
- **Some LTE operators may want to leverage the same LTE equipment for public safety services as well as mass-market services.** There are already examples of some operators looking at this kind of strategy in North America and elsewhere. In this example, an operator that wants IPsec protection for its mass-market users would inevitably require separate dedicated tunnels for its public safety customers to protect government communications from interception and/or analysis in the backhaul.
- **MVNO relationships for LTE will require second tunnels.** The primary operator will almost inevitably want to separate its traffic out from its MVNO partner in these scenarios. Indeed this scenario, as well as the previous one, may serve to drive faster adoption of IPsec by the primary operator.
- **Some Internet offload models will require separate IPsec termination points in the network.** To support non non-delay sensitive traffic being offloaded onto the Internet whole other traffic is forwarded onto the core where it will be terminated, more than one IPsec tunnel will be required again.
- **eNode B sharing models will require multiple IPsec tunnels.** Operators in Russia are deploying LTE on the basis of sharing capacity on the same eNode B. In March 2011, it was announced that in Russia, Yota will build out an LTE network that will also be shared by MTS, Megafon Vimpelcom and Rostelecom. A similar model is also being deployed in Poland. Again, one might expect these types of models to require effective isolation and protection of different traffic streams, and IPsec would appear to be the most suitable, standards-compliant approach to doing it.
- **An operator that determines it wants to move to a highly distributed model for its EPS nodes needs at least two IPsec tunnels to be instantiated by the eNode B.** In a one-tunnel model an operator is compelled to collocate its MME and security gateway in the same physical location, so that signaling and bearer traffic can be unencrypted in the same place. But this is a constraint on operators getting some of the efficiency and performance benefits of the flat IP network architecture. Operators that want to distribute their MME and security gateway elements will need the second tunnel.

VIII. Summary

With the rapid growth of mobile broadband, Internet attackers are moving their focus from the wireline to the mobile network. While this presents new challenges in itself, mobile network planners must also contend with new exposures that arise in the LTE architecture. One of the most significant exposures in LTE is the backhaul network, which no longer has the mandatory encryption of traffic in LTE that it has in 3G.

Operators must evaluate how trusted their backhaul network is in the context of growing threats from Internet attackers and the changes in the architecture. IPsec is the standard recommended by 3GPP when operators consider backhaul to be untrusted. Some common doubts and anxieties about using IPsec with LTE are well-founded, but many result from a lack of awareness, misunderstandings, or benchmarks and comparisons that are backward- rather than forward-looking.

A number of solutions pitched into the 3GPP security gateway market originate from the enterprise space. However, operators preparing for large-scale deployments of IPsec for LTE will ultimately need highly scalable platforms that are architected specifically for 3GPP Evolution and high availability. Due to the expected growth of LTE and its need for higher bandwidth, Security Gateways will need to meet or exceed high-capacity throughput performance, carrier-grade availability, and full compliance to the latest 3GPP security standards.

As this white paper has attempted to show, there is a path to deploying IPsec that provides additional – potentially critically important – protection to the LTE network, but in a manner that can be scaled according to the growth in the operator's LTE traffic and its capabilities in supporting those new requirements.

IX. Background to This Paper

9.1 About the Author

Patrick Donegan – Senior Analyst, Wireless, *Heavy Reading*
Donegan has been a telecom market journalist, analyst, and strategist for 20 years. He joined *Heavy Reading* from Nortel, having spent five years as a senior manager of strategic planning for that company's wireless business – spanning GSM, CDMA, UMTS, WiMax, and other wireless technologies. Prior to Nortel, he spent two years in business research for Motorola's Corporate Strategy Office in EMEA and two years as a wireless analyst for the Yankee Group. At *Heavy Reading*, Donegan has focused on next-generation mobile network issues. Donegan is based in the U.K. and can be reached at donegan@heavyreading.com.

9.2 Original Research

This *Heavy Reading* White Paper was commissioned by RadiSys, but is based on independent research. The research and opinions expressed in the report are those of *Heavy Reading.*

9.3 About *Heavy Reading*

Heavy Reading (www.heavyreading.com), a unit of *Light Reading* (www.lightreading.com), is an independent market research organization offering quantitative analysis of telecom technology to service providers, vendors, and investors. Its mandate is to provide the comprehensive competitive analysis needed today for the deployment of profitable networks based on next-generation hardware and software.

Heavy Reading
240 West 35th Street, 8th Floor
New York, NY 10001
United States of America
Phone: +1 212-600-3000
www.heavyreading.com