Nokia Siemens Networks
LTE-capable transport: A quality user experience demands an end-to-end approach

White Paper

1. Executive summary: The need for end-to-end transport network planning

The rising tide of data traffic is putting transport networks under more pressure than ever and the arrival of LTE will only accelerate that process. This paper focuses on what is needed for an LTE-capable transport network to deliver an optimized end-user experience.

The first requirement is for cost-effective capacity. Should transport networks be dimensioned to meet average demand or peak demand, for example? Communications service providers (CSPs) could approach the dimensioning of their networks in a number of ways, but the optimum solution will essentially strike a balance between providing the maximum capacity for users and keeping the transport network economically and technically feasible.

Latency also has a considerable impact on user satisfaction, especially in delay-sensitive applications such as online gaming. If the reaction time of the network is too long, a high-speed connection won't do much to improve the experience. Network providers can control some aspects of latency, but it also depends on external factors, such as the distance between the user and the content. An approach that delivers the best possible latency won't go wrong, provided it is economically viable.

Quality of Service (QoS) differentiation enables CSPs to manage the performance of different streams of traffic. QoS can be a powerful tool for managing the user experience, but it must be managed end-to-end. LTE radio QoS has to be aligned with the QoS in the transport network, for instance.

Synchronization is essential in all telecommunications networks. A number of different strategies can be adopted in transport networks and they all have their advantages and limitations. Hybrid systems, where a mix of synchronization technologies is used, are likely to be commonplace.

Service assurance and network security are the other factors that play key roles in determining the user experience in LTE-capable transport networks.

Ultimately, sound network planning and an end-to-end approach to network operations will determine how well these emerging transport networks perform.

2. The impact of LTE on transport networks

LTE promises a whole new mobile broadband experience for everyone, with throughput rates beyond 100 Mbit/s and short latency of around 20 ms or better. It's an experience formerly available only from fixed connections. One thing is clear, however. All the progress made in the radio and core subsystems won't count for much unless the underlying transport architecture is ready to deliver the key performance indicators (KPIs) required to support such a lofty value proposition.

There is a general consensus in the industry that only a packet-based transport network will be able to meet the challenge. However, there are still unresolved issues around transport and they tend to revolve around three topics:
- How should we provide the user experience? What throughput and latency values are required and how can we achieve them?
- How can we do it all cost-efficiently and bring the price per bit down?
- What is the right transformation strategy? What is the optimum target architecture and how can we get there?

In this white paper we will focus mainly on the first question and discuss the requirements for an LTE-capable transport network. We'll look at how to provide the best – or more precisely optimized - user experience. This naturally includes a look at how CSPs design and implement the most reliable transport networks.

3. Capacity and dimensioning

HSPA, HSPA+, LTE and LTE-A each promise to deliver progressively higher data rates, so how should the underlying transport network be dimensioned? What is the capacity requirement for each base station (BTS)? For the latter there are two basic approaches: The bottom-up approach is based on actual traffic model predictions, while the top-down alternative is based on the bitrates possible with different air interface technologies.

3.1 BTS transport capacity: bottom-up

Network dimensioning has traditionally used the bottom-up approach. A traffic model is calculated for a time period based on certain assumptions. The model then produces estimates that can be used to dimension the transport network.

The obvious advantage of this approach is the scientific basis for the estimates, which are based on experience. It also is independent of the actual radio technology used and could be used

to plan radio and transport network capacity development over time.

3.2 BTS capacity based on air interface bit rate limitations

On the other hand, many CSPs do not have previous experience with the uptake of data services. Flat rates and large data bundles typically make predictions difficult. In short, if the sort of educated guess used previously is not practical, the other option is to do a top-down calculation based on the air-interface bitrates of different radio technologies to achieve an estimate of the user plane traffic. The following figure shows the theoretical maximum bitrates available for certain configurations. Note that those peak values are only for a single sector, so a three-sector site would have to serve three times these peak rates.

Dimensioning a network based on peak rates is looking very much at the "worst case" and will result in over-dimensioning. It's therefore useful to

consider the realistic peak bitrates, which can normally be achieved within the cell. The above figure shows average cell throughput rates based on simulations, which were carried out by 3GPP considering a certain user distribution in the cell, terminal mobility, interference etc.

When calculating the total transport capacity needed per BTS, full peak bitrate dimensioning might result in values that are too high. Dimensioning based only on the average might result in values that are too low and cause regular congestion. A good compromise might therefore be to use a so-called "single-peak, all-average" model, as shown in the next figure.

In this model the user traffic requirement of the BTS is presumed to be either the aggregated average capacity of all cells or the peak capacity of one cell. Planners use whichever value is higher, so that the advertised user service peak rates can be momentarily supported in any given cell, although the advertised user service rate will be only a fraction of the cell peak rate.
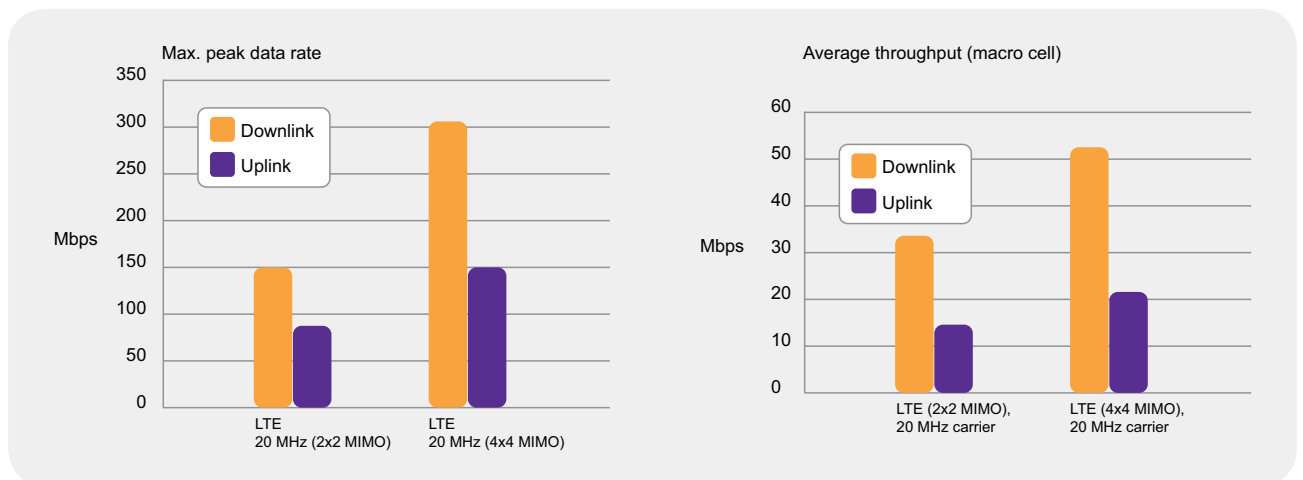

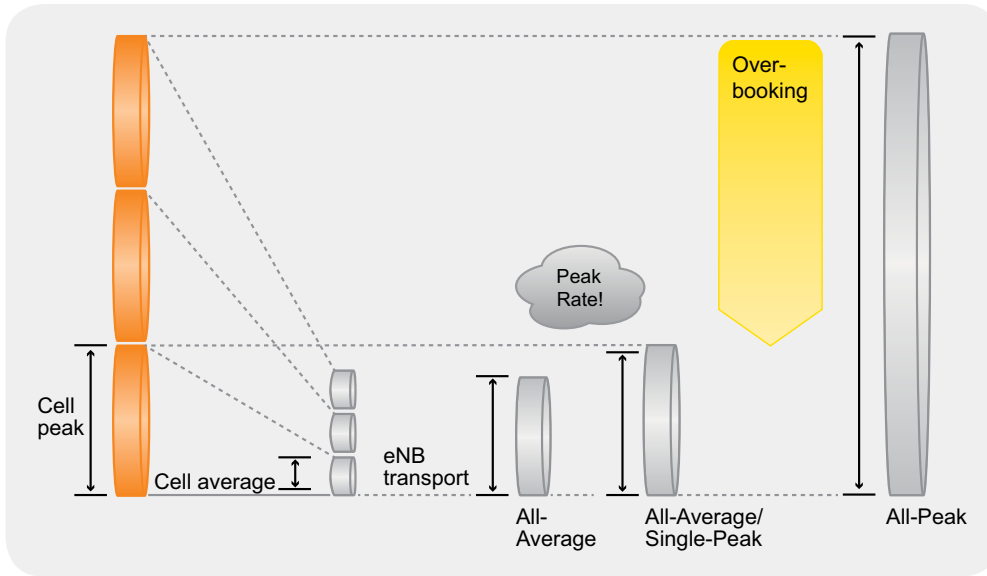Figure 1: Maximum peak vs. average data rates

Figure 2: Calculating the single-peak, all-average data rate

The final step needed to obtain the actual bitrate required for the BTS's S1 interface includes some overhead calculations. The air interface overhead is stripped out and the transport and possible IPSec overheads are added in. Of course, signaling and management traffic should also to be taken into account.

3.3 Traffic aggregation

Moving beyond the first link (or "last mile") connecting the actual BTS deeper into the network, aggregation

and overbooking become even more essential to ensure transport efficiency. In fact, aggregation should be carried out close to the BTSs (for example, in MWR hub sites) to really leverage this advantage. The example Figure 3 shows how the multiplexing gain depends on the number of aggregated BTSs, based on certain assumptions.

It is obvious that relatively large overbooking advantages can be gained by adopting hub-aggregation, whereas continued higher regions of the network does not have as large

an impact and the gain ultimately levels out.

3.4 LTE capacity requirements and the X2 Interface

One other peculiarity of LTE networks as compared to traditional 3G networks is the X2 interface, which plays an important role in the handover of connections between neighboring BTSs. During the handover procedure the radio link to the terminal is interrupted for a short time, typically between 60 and 70ms. Downlink packets arriving at the BTS formerly hosting the terminal will be forwarded to the new BTS, connecting the terminal via the X2 interface until the EPC has switched the S1 path to the new BTS.

In that sense, the X2 interface creates another set of traffic flows directly between neighboring BTSs. This traffic is extremely bursty, since it occurs mostly during brief handover phases. Studies show that it will normally be in the range of only 3% of the total S1 traffic.

It is therefore a relatively minor factor in planning, and the capacity certainly does not require the installation of dedicated physical transmission links between neighboring BTSs.
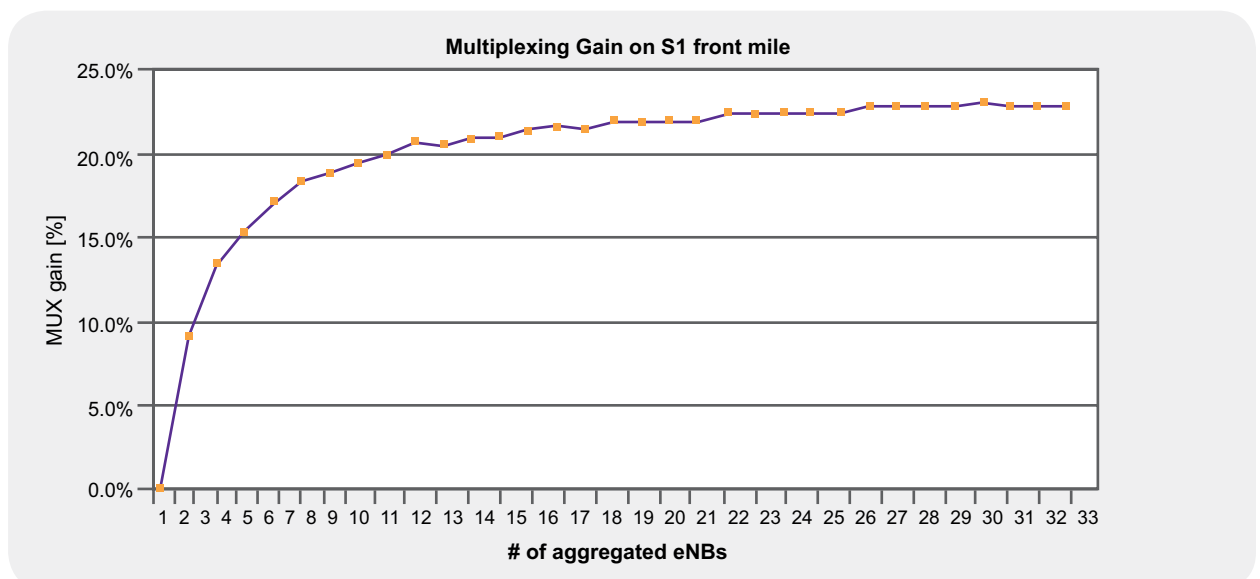

Figure 3: Multiplexing gain dependent on the number of aggregated BTSs

3.5 Striking a balance between capacity and economy

A proper bottom-up planning process for transport network capacities would probably provide the most accurate (and differentiated) results when dimensioning an LTE transport network. However, in many cases a top-down plan based on the air interface peak and average rates will be a more feasible way forward. In this context we suggest the use of the "single-peak all-average" model for the individual cell sites. This is essentially a compromise between providing the maximum capacity for users and keeping the transport economically and technically feasible. At the same time aggregation in the network is essential in order to leverage multiplexing gains, preferably close to the BTS sites.

4. Latency

Latency (or delay) is another factor that affects subscribers' service experience. From the user perspective, latency is essentially the time it takes for a data packet to travel from the terminal via the mobile network to the content server on the internet and vice versa.

There are several components affecting the final latency experienced by the subscriber. There's the system's inherent latency that depends on the radio technology used (BTSs and their controllers, gateways and so on). Then there are additional delays arising from the transport network, from the connectivity between the CSP's network and the internet and the time needed to reach the actual server running the requested service. On top of this there may also be a queuing delay within any of the various nodes if there's any congestion.

From the CSP perspective there are two elements to this latency. One is the delay introduced by the operator's network, which is the round trip time between the user's handset and the operator's internet gateway. The CSP can influence and optimize this delay. However, the other component is the time it takes for the data to travel from this gateway to the actual content server and back, and the CSP has no direct influence over it.

Latency is considered by many to be as important as the actual capacity supported, since it governs parameters such as the time it takes for a requested internet page to display. If the reaction time is too long, a high-speed connection can't do much to improve the experience. This is just one example where latency plays a role. Different services have different latency requirements.

4.1 Radio technology inherent latency

LTE offers hugely improved inherent latency values compared with other radio technologies, such as 3G or even HSPA. From 60 ms in HSPA, latency is reduced to about 20 ms in LTE (all roundtrip times).

Note that these values only take into account the radio and core components of latency. They ignore
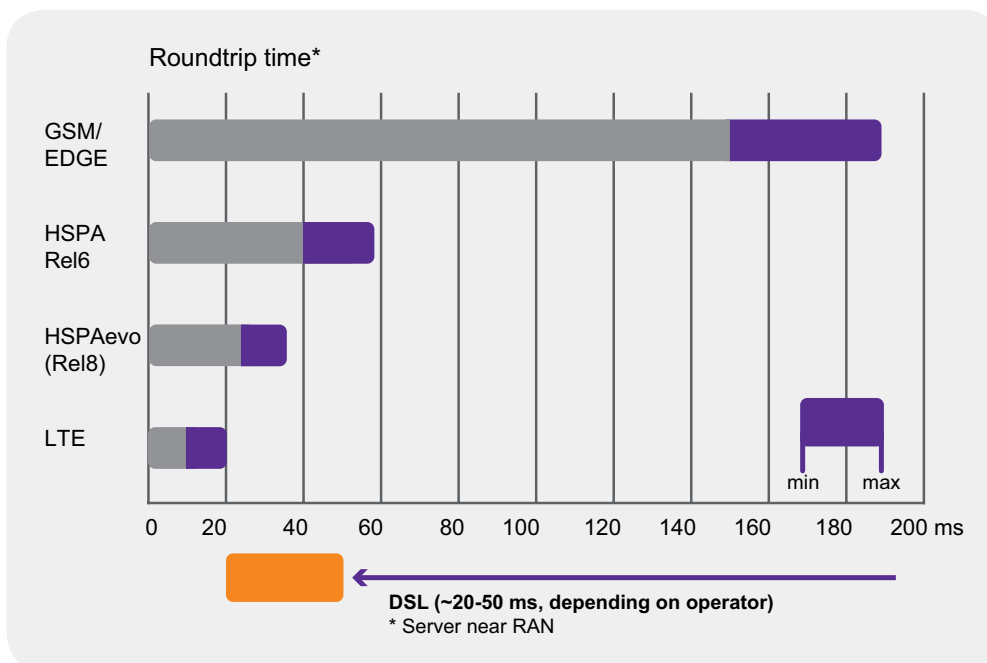

Figure 4: Roundtrip delays

that the fact that the physical transport can (and will) contribute significantly to the overall latency. In other words, the low latency promised by LTE will only be experienced by the user if the underlying transport also supports low latency.

4.2 Latency in transport and its origin

Propagation delay: The speed of light is finite, which leads to a round trip time of about 1ms per 100km. This shows the impact of topology on overall latency.

Buffering and queuing delay: Packet-based transport systems use a number of buffering and queuing mechanisms, each of which adds delay. Proper link planning will minimize this effect.

Transmission delay: A data packet takes a certain amount of time to be transmitted based on its length and the bandwidth of the connection. For large packets this can lead to delays of several ms over small bandwidth connections.

Signal processing delay: The more signal processing takes place within a signal path, the longer the delays. Therefore the sheer number of processing nodes plays a role, as does the difference between simply connecting on the optical level and processing actual routing operations.

4.3 Latency recommendations

In contrast with some of the earlier radio technologies, it is normally not the radio network systems (such as the control plane) that impose practical limits on the delay budget in LTE. Instead it's the way users experience different services. The experience typically becomes unacceptable long before network systems run into trouble.

The acceptable latency depends on the service type. 3GPP has indicated certain one-way delay goals for specific services in TS 23.203 (Table 1).

It is mainly online gaming, video conferencing and machine-to-machine (M2M) applications that drive latency requirements. For example, they have proposed a 50ms delay budget for online gaming. However it should be noted that the focus of the 3GPP document is on functionality in the core (including service prioritization) and not the transport network itself. Therefore a fixed (one-way) delay of 20 ms was assumed for the transport network.

The experience of many services depends more on actual latency than on the available bandwidth. TCP is used to shift a large part of the non-real time internet traffic and uses a handshake to secure transmission. This means that the achievable bitrates

and hence the download time for most applications (web pages, music, video, software and so on) is depending on the total round-trip time (including LTE radio, mobile backhaul, LTE core and the internet domain).

This can be improved by activating specific options in the protocol stack (such as TCP Window Scaling) or using multiple concurrent TCP sessions. The applicability of these improvement options depends on the particular operating systems used on terminals and servers, as well as on the actual application.

Finally, some industry bodies have issued recommendations for the permissible delay in mobile backhaul. These are based on the considerations already mentioned. For example, the NGMN defined a limit of 10 ms for two-way delay, and 5 ms if the CSP requires it ("NGMN-optimized backhaul requirements", released August 2008).

However, such recommendations have to be seen in context. It will be difficult to stick to such low delays over large geographies, since a transmission distance of only 1,000 km completely exhausts this delay budget. In this case, a delay budget of 40 ms from the BTS to the EPC could be seen as a good compromise that still allows providers to offer the most demanding real-time gaming services.

4.4 X2 latency requirements

The X2 interface also has its own latency requirements. It might seem at first glance as if these latency requirements would very stringent. However, it has to be considered that during the handover phase the radio link to the user terminal will be interrupted for a short time anyway. Any forwarding of packets faster than between 60 and 70 ms therefore serves no real purpose.

Given that an LTE transmission network should be designed with stringent delay targets, the X2 interface does not significantly change things. In particular, it does not mandate the use of direct inter-BTS connectivity.

Guaranteed Bit Rate	Delay Budget	Loss Rate	Application Example
GBR	100 ms	10^{-2}	Conversational Voice
	150 ms	10^{-3}	Conversational Voice, Live Streaming
	50 ms	10^{-3}	Real Time Gaming
	300 ms	10^{-6}	Buffered Streaming
Non - GBR	100 ms	10^{-6}	IMS signaling, Control plane
	300 ms	10^{-6}	Buffered Streaming, TCP applications (specific service)
	100 ms	10^{-3}	Interactive Gaming, Live Streaming
	300 ms	10^{-6}	TCP applications (premium bearer)
	300 ms	10^{-6}	Default Bearer

Table 1: Some applications are more sensitive to latency issues than others

4.5 Minimize latency to optimize the user experience

Any latency requirements are driven primarily by the targeted user experience. In that sense, an approach that delivers the best possible latency won't go wrong, provided it is economically viable.

Truly end-to-end optimization for latency has to take into account a number of factors, including topology (distances and the number of processing nodes) and the distance between the EPC and the internet peering point, as well as proper link planning and dimensioning.

The delay outside of the operator's network also deserves some attention. If content is stored literally at the other end of the world, delay values will be very high in any case. Content buffering and similar methods are therefore set to become increasingly important.

5. Quality of Service

5.1 What is Quality of Service?

Quality of Service (QoS) differentiation enables CSPs to segregate the flow of traffic and this allows them to monitor and manage the performance of different streams individually. Such a differentiation could be made on the basis of applications or services (with real-time services typically being more critical), subscribers (for example, with gold, silver or bronze profiles) or operators (especially in situations such as transport network sharing).

From the transport point of view traffic flows can be assigned to QoS classes based on a number of parameters, such as the required packet delay, delay variation and packet loss. Such parameters are typically universally good in lightly loaded networks. However, as discussed before, the economic transport of LTE data rates will lead to a certain overbooking and congestion. In this environment, QoS is the tool that guarantees that, say, voice traffic packets preferential treatment compared with peer-to-peer traffic.

In that respect it is useful to differentiate between Quality of Experience (QoE) and QoS. The former describes the quality of the end-user experience, while the latter is the method used to manage this experience.

A complete QoS differentiation solution spans the whole network. The core is responsible mainly for QoS management, such as the definition and dissemination of respective QoS policies, including the use of technologies such as DPI. Both the radio access and core networks take care of QoS control, tagging respective traffic packets using VLAN priority bits or DSCP (DiffServ Code Points) values in the IP header, for example. QoS enforcement is carried out by the radio (for the air interface) and transport (for the transport network) systems.

5.2 QoS enforcement and transport QoS

With all this in mind, the QoS requirements for a transport network to support LTE are thus about guaranteeing appropriate service levels for each service in terms of packet delay, delay variation and packet loss.

The basic functions implemented in transport network elements are prioritization and capacity reservations. In IETF standards they are referred to as Differentiated Services (DiffServ) and Integrated Services (IntServ) respectively. Element implementation is often a combination of these principles. There will be some resource reservations for parts of the traffic and prioritization will be used in scheduling.

Prioritization ("soft" QoS)
Queuing systems in various elements enforce prioritization. In the air interface there is a packet or frame scheduler that prioritizes the data. In addition, transport resource management algorithms or multiplexing algorithms can be used.

Queuing mechanisms will typically include strict priority (for high-priority traffic) and weighted fair queuing for the lower priority classes. The number of queues that can be used to differentiate the traffic are important, since this determines the level of granularity, or the number of different traffic classes that can be distinguished. The maximum of classes that can be differentiated with VLAN p-bits on the Ethernet layer is eight, so this might be considered a useful number of queues.

Resource reservation ("hard" QoS)
Admission control estimates whether there will be sufficient resources for each new connection or traffic flow. This functionality is mandatory when implementing guaranteed bit rate connections. Static resource reservation – perhaps via a Network Management System (NMS) – is a good option in the mobile backhaul

sector, for example, by using MEF-type services with a Committed Information Rate (CIR) and Peak Information Rate (PIR).

CSPs can deliberately limit the throughput of a connection by buffering the data so as not to exceed the pre-defined maximum bit rate (shaping), or by dropping packets that would exceed the maximum bit rate (policing). Traffic shaping and policing can also be used within queuing systems as congestion control mechanisms.

QoS implementation

In a real case, as mentioned before, the classification and tagging of traffic is carried out both by the BTSs and by the gateways, based on the information collected in the gateways and the policy server.

For the most typical case of service-based differentiation, this classification relies on the QCI (QoS class identifier) as defined by 3GPP (Table 2). The QCI value references a certain application type and is used within the access network as a reference for controlling packet forwarding treatment.

QCI values are translated into a packet priority marking (DSCP value and/or VLAN p-bits) applied by the BTSs and gateways. Similarly, the control, management and synchronization plane traffic is marked to ensure it receives the right priority treatment on the outgoing interfaces.

Based on these QoS markings, the transport network elements in the packet's path can then ensure that each packet is handled according to its required forwarding behavior, for example, by assigning it to the correct queues. This can be combined with connection admission control for the transmission network elements, adding a component of hard QoS.

5.3 QoS must be managed end-to-end

QoS can be a powerful tool to achieve a QoE for the LTE end user. It can take into account the requirements of different services, as well as the SLA purchased by the customer. It helps manage resources in congested environments, especially where there's pressure on radio access and the mobile backhaul domain. QoS is an enabling technology for a viable business case.

However, QoS must be managed consistently end-to-end. LTE radio QoS has to be aligned with the implementation in the transport network, but the transport network also has to cater for the QoS needs of 3G or – even more stringently – 2G packet traffic.

Note, however, that in the same way as latency, any operator can only control QoS within his own network. As soon as traffic leaves for the internet, the treatment is essentially best effort.

LTE Radio domain				LTE Transport domain
LTE Traffic Class	QCI	Resource Type		DSCP (Ethernet p bits)
Conversational Voice	1	GBR		46 (5)
Conversational Video	2			26 (3)
Real Time Gaming	3			46 (5)
Non-conversational Video	4			28 (3)
IMS signaling	5	non-GBR		34 (4)
Voice, video, interactive gaming	6			18 (2)
Video (buffered streaming)	7			20 (2)
TCP-based (e.g. www, email, ftp, p2p file sharing etc)	8			10 (1)
	9			0 (0)
C-plane				46 (5)
M-plane				34 (4)
S-plane				46 (5)
ICMP				10 (1)

Table 2: An example for mapping radio QoS onto transport QoS

6. Synchronization

6.1 Synchronization definition and requirements

Synchronization has always been of vital importance in telecommunication networks. In mobile networks it is needed for the air interface to enable smooth handovers and for aligning coding procedures.

There are two main flavors in synchronization. The most common is frequency synchronization. A standard 3GPP requirement for all radio technologies is to deliver an accuracy of the modulated carrier frequency of better than 50ppb for macro cells.

In addition, TDD technologies such as TD-LTE or WiMAX and some features such as Multimedia Broadcast Multicast Service (MBMS) also require highly accurate time (or more precisely, phase) synchronization. In the case of TD-LTE, the maximum timing error at the air interface must not exceed 1.5μs. For more detail, refer to table 3.

6.2 Synchronization options with packet transport

There are a number of ways of providing high-accuracy synchronization information. The most obvious is to use the synchronization clock output from co-located TDM-based equipment and thus effectively relieve the packet transport of synchronization duties. This is only possible in fully hybrid transport networks and does not cover the need for phase synchronization.

Another obvious method would be to use GPS receivers at each cell site, effectively covering all the synchronization needs that could possibly arise. However, the cost might be prohibitive, and with cells getting smaller and indoor-hotspot coverage more of a requirement, it might not be physically possible to use GPS receivers everywhere.

There are two main methods for packet-based synchronization:

6.2.1 IEEE 1588-2008
The IEEE1588 solution is standardized and is by far the most common implementation of packet synchronization. It consists of a Grandmaster (server) at a core site and Timing Slaves (clients) implemented in either the BTS or a transport network element, such as a cell site device. The master and slaves communicate through a bidirectional IP protocol called PTP (Precision Time Protocol) containing time stamps. The slaves apply intelligent algorithms to recover from the received packet stream the original clock information at the Grandmaster.

IEEE1588 can provide both frequency and phase synchronization. However, if used for phase synchronization, all the nodes in the transport network between the master and slave have to provide on-path support with so-called boundary or transparent clock functions. IEEE1588-2008 can run over any kind of IP and/or Ethernet network.

6.2.2 Synchronous Ethernet
Synchronous Ethernet (SyncE) is defined in G.8261/8262/8264 as an SDH-like enhancement for transporting frequency information on the physical layer of an Ethernet link. In contrast with IEEE1588, which is essentially a layer 3 technology, frequency synchronization will be extracted directly from the Ethernet interface at the BTS. Unlike packet-based synchronization (IEEE1588-2008, NTP), the stability of the recovered frequency does not depend on the network load and impairments.

SyncE has to be implemented at all intermediate nodes on the synchronization traffic path. In addition, it does not provide phase information, so it cannot be the only synchronization mechanism in the case of TD-LTE, for example.

6.2.3 A choice of three
In fully packet-based networks, only three mechanisms really can be used: GPS, IEEE1588 and SyncE. Any of these mechanisms will be useful for LTE FDD. For TD LTE either IEEE1588 with on-path support or GPS is the tool of choice. Of course it is also possible to combine different mechanisms.

Standard	max. frequency error at air interface	max. timing error at air interface
WCDMA FDD	± 50 ppb (Wide Area BTS) ± 100 ppb (Medium Range BTS) ± 100 ppb (Local Area BTS)	No requirements.
LTE FDD	± 50 ppb	No requirements.
LTE TDD	± 50 ppb	± 1.5 μs
GSM	± 50 ppb ± 100 ppb Pico Class BTS	no requirements.

Table 3: Synchronization is critical

7. Service assurance

The transformation towards packet-based traffic raises service assurance challenges. Where TDM-based technologies had a wealth of operation, administration, and maintenance (OAM) tools available, packet technologies mostly did not support them in the past. However, this issue is now being addressed by the relevant industry bodies and standards. Service assurance capabilities are increasingly being implemented in packet transport equipment.

There are different standards and concepts available depending on the networking layer being used. They typically include a set of functions that enable detection of network faults and the measurement of network performance, as well as the distribution of fault-related information.

7.1 IP layer OAM

There are multiple options and protocols (standardized and proprietary) available on the IP layer to provide OAM functionality. Some of them allow CSPs to deploy test traffic to measure the throughput, delay, delay variation and packet loss between two points in the network. Different DSCP values can be assigned to the test traffic, allowing engineers to measure the network behavior of various service classes.

7.2 Ethernet OAM

There are several Ethernet OAM-related standards available, which address either a single link or multiple links

Link Layer OAM (IEEE 802.3ah)
This looks at a single link and includes functions for discovering and monitoring the link, as well as indicating remote node failures

Ethernet Service OAM (IEEE 802.1ag / ITU-T Y.1731), a.k.a. "Service Layer OAM"
These functions allow monitoring of end-to-end connectivity and performance between the nodes in an Ethernet domain. Additional functions are included to support resilience.

Similar measurements are possible on the Ethernet layer per VLAN / p-bit for measuring throughput, delay, delay variation and packet loss per service and priority class.

7.3 Practical implementations

Practical implementations can differ significantly. Some or all of the OAM functions may be implemented and the granularity of performance counters will vary (including real-time vs. history counters). The number of active counters may also be limited.

Those functions can be implemented in the BTS and core elements (particularly important for true end-to-end OAM), or in the transport elements. It can be useful to have dedicated probes in the network to monitor specific points in the network or to compare the service level of leased lines with contracted SLAs by placing the probes directly at the endpoints.

8. Network security

8.1 The need for transport security in LTE

Packet traffic is vulnerable to hacker attacks. Methods have evolved rapidly, with cheap hardware providing hackers with high processing power and better tools. In addition, BTSs that were traditionally located in secure, locked sites are increasingly set up in public places.

Furthermore, there are two major differences that make security different in LTE transport networks, compared with WCDMA:

1. The air interface encryption of the user-plane traffic is terminated at the BTSs, so user-plane traffic in the LTE mobile backhaul network is

not secured by radio network layer protocols.

2. Since the LTE network architecture is flat, other BTSs, the EPC nodes (MME, S-GW) and other nodes in the core network become directly IP-reachable from BTS sites. If physical access to the site cannot be prohibited, a hacker could connect his device to the network port, attack these network elements and cause significant network outages.

Transport security features are mandatory unless both the mobile backhaul network and the BTS site are secure. IPSec provides a comprehensive set of security features (traffic authentication, encryption, integrity protection), solving both the problems mentioned above. The 3GPP security architecture is based on IPSec and Public Key Infrastructure (PKI). IPSec is applied between Security Gateways (SEG), which are typically located at the cell site and at the border between the trusted and untrusted network.

Probably the most efficient solution for the realization of the SEG function at the cell site is the integration of the SEG in the BTS itself. This minimizes physical accessibility, which is especially important in easily accessible hot-spot cells. It also reduces the need for additional equipment and reduces the site footprint. However, care should be taken that any such integrated solution has the necessary throughput to support LTE data rates, for instance, in a highly-loaded three-sector cell base station, and does not add significantly to the overall delay.
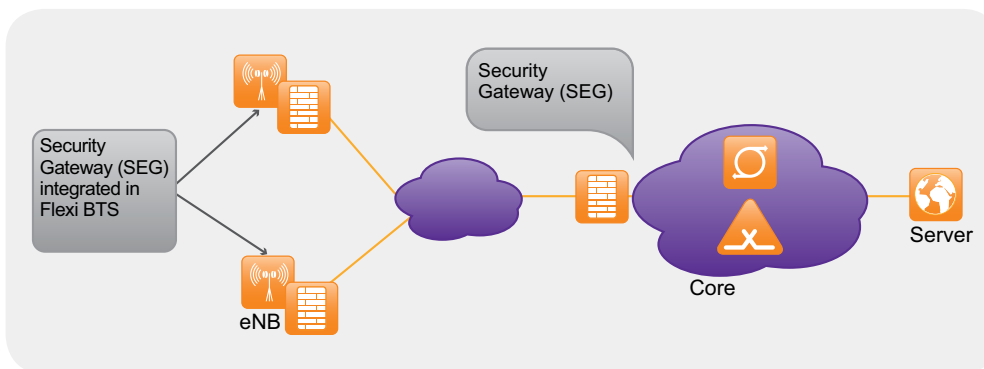
Note that such IPSec implementations and the architecture decisions necessary for the X2 interface are closely connected. After all, it is not only the user plane (S1) and management plane traffic that should be encrypted, but also the hand-over user traffic via the X2 interface. Considering that each connection requires an IPSec tunnel, security architectures can get quite complicated in a fully meshed architecture.

Figure 6: SEGs guard the CSP's Security Domain

9. Conclusion

This paper has discussed the main requirements for an effective and efficient transport network to support LTE and legacy mobile technologies.

LTE transport cannot be separated from the LTE radio and core systems because these systems have to interact at too many points, whether for effective synchronization, QoS implementation, security or service assurance. Parameters such as capacity and latency have to be planned in an end-to-end manner, because the weakest link will be the breaking point.

Close cooperation is therefore required between the relevant technical teams. A partner that is knowledgeable and experienced in all aspects of network optimization can prove valuable for CSPs looking to combine maximum subscriber satisfaction with a viable business case.

10. Glossary

- CSP: Communication Service Provider
- DSCP: Diffserv Code Point – QoS tag on IP layer
- EPC: Evolved Packet Core
- HSPA: high speed packet access
- IPSec: IP encryption methodology
- LTE-A: Long Term Evolution – Advanced
- MME: Mobility management entity – part of the EPC
- MWR: Microwave radio
- p-bit: Priority bit – QoS tag on Ethernet layer
- PKI: Public key infrastructure – key sharing concept for IPSec
- PTP: Precision Time Protocol
- QoE: Quality of Experience
- QoS: Quality of Service
- S1 interface: the logical interface between BTS and S-GW and MME gateways / evolved packet core (EPC)
- SEG: Security gateway
- S-GW: Service Gateway – part of the EPC
- X2: the logical interface between neighboring BTSs, used e.g. during hand-over

www.nokiasiemensnetworks.com