Real-time backhaul assurance to enhance QoE

The evolution in backhaul monitoring in LTE networks

By Monica Paolini, Senza Fili

SENZA FILI

Sponsored by

VIAVI

1. Introduction

Backhaul assurance's role expands in response to higher traffic complexity and use of carrier Ethernet backhaul

LTE brings the much needed performance and capacity improvement over 3G that enables operators to provide better service and QoE to their subscribers. But to leverage the new network capabilities, operators need to manage traffic more actively – indeed, proactively – to prevent service issues from manifesting. They need to be able to monitor, troubleshoot and optimize each element in their network, while at the same time keeping track of QoE and end-to-end network performance in real time. Legacy networks, less complex and more homogeneous than LTE networks, do not require – or allow – this intense level of management of network resources. They are easier for mobile operators to monitor and operate.

Today, mobile operators are modifying their backhaul to support these complex heterogeneous networks with latency-sensitive applications, which require real-time, QoE-based optimization if operators are to make more efficient and profitable use of their network resources. The introduction of LTE and the overall network evolution affect backhaul and, specifically, backhaul assurance primarily along two dimensions: the overall changes in traffic dynamics and traffic management, and specific changes in backhaul technology and provisioning.

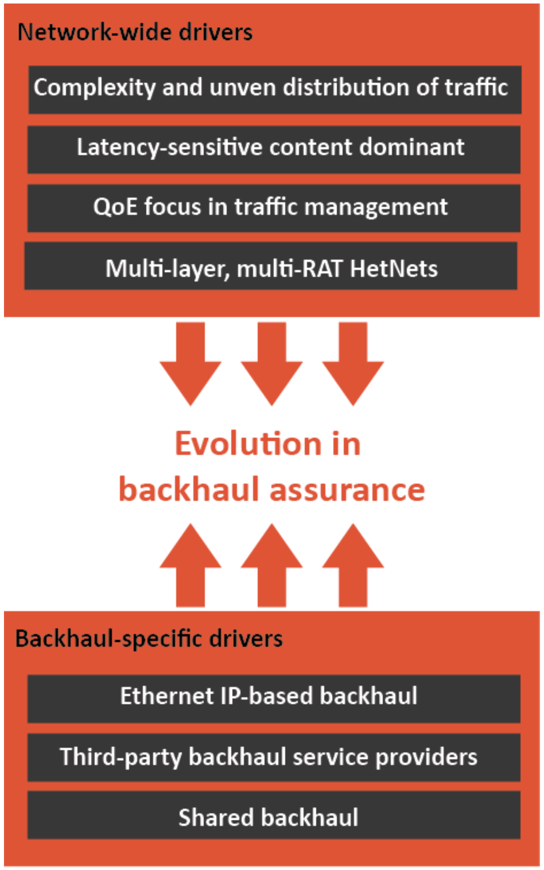

Source: Senza Fili

Not only are we seeing a staggering increase in traffic volumes; the complexity of traffic is increasing, and mobile operators have to manage traffic flows tied to different applications that have different requirements, are extremely variable in spatial and temporal distribution, and are subject to complex, real-time policy enforcement. And they have to achieve this management in networks with multiple layers and multiple RATs. Operators want to use their network resources as efficiently as they can, and to keep their subscribers happy even under the most demanding application requirements: to do so, they have to explicitly monitor and optimize QoE. Backhaul still has to provide the required capacity, as it has in legacy networks, but it also has to address the traffic complexity and latency sensitivity appropriately to avoid becoming the performance bottleneck in mobile networks.

Management of mobile backhaul is made even more complex by the expanding adoption of true IP-based Ethernet technologies, and by the fact that backhaul provisioning is increasingly shared and managed by third-party service providers. Operators have less direct control over the backhaul, and find it more difficult to gain visibility into it, at a time when the relevance of control and visibility have grown along with the need to manage traffic more actively. Backhaul assurance is essential to giving mobile operators the tools they need to monitor and troubleshoot their networks end-to-end and address appropriately any performance issue that may arise within the backhaul portion of their networks.

2. Traffic growth and focus on QoE require a new approach to traffic management

Video and voice lead to the dominance of real-time data

Mobile traffic continues to grow relentlessly – from 3.7 to 30.6 TB/month over the 2015–2020 period, with a 53% CAGR, according to Cisco VNI. The traditional response to increased demand has been to add cell sites or sectors to increase capacity. This is no longer sufficient – and it is a financially challenging proposition for mobile operators when used alone: it requires large investments that are not backed by a corresponding increase in revenue, because ARPUs are stable or even declining in most markets.

Mobile operators are discovering that they need to manage traffic more actively to drive resource utilization up, because this allows them to extract more value from the deployed infrastructure, and contain or postpone the need for expensive network expansion. With a more proactive traffic management approach, mobile operators can purposely allocate network resources to maximize QoE – giving their subscribers the best experience their networks can support.

Most of the attention in the wireless industry today focuses on the increase in data traffic, but equally important is the change in traffic characteristics, especially distribution and complexity. Initially all traffic was voice. Texting added some amount of data, but the volumes were always limited and the requirements easy to meet.

Today most mobile traffic (more than 90% in developed markets) is data, and with VoLTE, voice too becomes an instance of data traffic. Mobile video will increase from 2 to 23 TB/month, an 11-fold increase, between 2015 and 2020, and will account for 75% of total mobile data traffic by 2020. The requirements operators must meet to provide a good subscriber experience become more stringent with the increasing prominence of real-time traffic such as voice and video. Conversational video traffic, such as Apple’s FaceTime and Microsoft’s Skype and Lync services, requires voice and video clarity with no perceptible delay or packet loss, and is more sensitive to latency issues than streaming services.

The evolution of traffic type

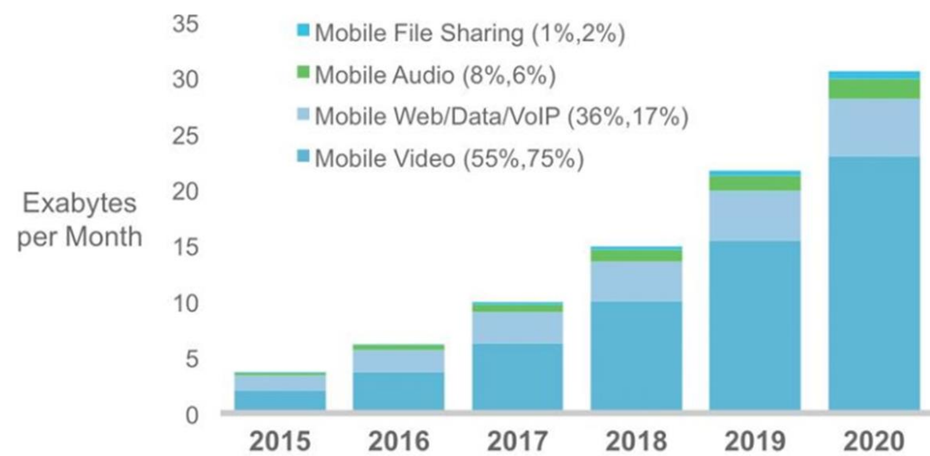

Figures in parentheses refer to 2015 and 2020 traffic share.
Source: Cisco VNI Mobile, 2016

3. Traffic complexity and uneven distribution grow in LTE networks

QoE becomes the target of network traffic optimization

The shift to IP data does not make it easier to manage traffic priorities. Furthermore, the way we use data and the requirements for different data flows have added complexity in managing traffic end-to-end in mobile networks. The table on the right lists different drivers responsible for the increase in data complexity that affect the way operators manage their overall networks and, specifically, their backhaul. The increased use of video and the introduction of VoLTE are the changes that have had the largest impact to date. We expect the other drivers, such as IoT, to take on a large role in shaping traffic management in the future.

Voice and video provide a good illustration of the impact of traffic complexity on network management. Real-time traffic types such as video and voice have similar requirements in terms of latency, jitter and packet loss that sets them apart from other data streams. Yet operators typically treat video and voice differently. Because of the importance of voice quality for subscriber retention, operators may want to give VoLTE priority over all other data services, including streaming video. Because of the high bandwidth requirements of video, they may want to limit the bandwidth allocated to video traffic in networks that are at capacity or congested. In addition, because of the special requirements of VoLTE, operators have to treat VoLTE traffic differently from OTT voice services. Similarly, they may set higher performance targets for conversational video than for streaming video, because subscribers are likely to be more sensitive to the quality of conversational video.

As a result, mobile operators need to manage traffic more carefully to drive resource utilization. This translates into the need to manage and monitor traffic not as a homogeneous flow of packets, but as a concurrent set of flows.

As traffic flows through the network end-to-end, operators need to know what the performance level is, both at different locations within the network and from the subscriber perspective in terms of QoE. They need to know this from multiple

Traffic characteristics that affect network management
Traffic type. Requirements for different types of traffic (e.g., voice, video, or best-effort data) vary greatly in terms of bandwidth, latency, jitter, packet loss, and mobility. Voice remains a special case, with subscribers strongly sensitive to degraded quality.
Application or service type. The same traffic type may be transmitted as a different service or within a different application. For instance, subscribers can get streaming video within OTT applications such as Netflix, or as a conversational video for an OTT application such as WebEx or Zoom, or an operator-managed ViLTE service. Video traffic may also be encrypted or not, and optimized by the content provider or the operator.
Spatial distribution. Usage is extremely concentrated geographically in a small part of the network – specific venues, central metropolitan areas – leading to congestion in specific areas.
Temporal distribution. The network traffic load changes throughout the day and week as subscribers travel to and from work, and go out at night and on weekends.
Microbursts. Data traffic is inherently spiky at the millisecond level and below. This may cause congestion in the network even though, when looking at transported traffic averaged over time, the traffic load on the network appears to be operating within capacity.
Policy, traffic prioritization. The mobile operator may use policy to prioritize traffic or allocate it to specific RATs, channels or infrastructure elements (e.g., macro or small cells).

dimensions – by application, type of traffic, and location – and they need to know it in real time. For both monitoring and troubleshooting the network, operators also need a precise understanding of what is required to ensure good QoE, and to prevent or solve performance issues. For that, they need visibility into the network at different levels of granularity to see how, for instance, application, traffic and location interact with each other, without succumbing to unmanageable complexity.

Traditionally, operators have relied on historical network KPIs that provide an averaged view of the performance of network elements. Although this data is still valuable, and undoubtedly operators will continue to use them to assess network performance, historical averaged KPIs do not have the granularity needed to assess network performance in real time, how it relates to QoE, and what the bottlenecks in the network are.

For instance, an operator may decide to give priority to voice and selected video services, and ensure that the latency is low for this type of service. However, this may drive up latency for applications like web access, messaging or downloads, and this is acceptable because increased latency there is likely to go unnoticed by subscribers. As a result, the averaged network latency may be higher than if all traffic were treated equally, but the latency for the selected voice and video services may be low, and hence in line with the operator's performance targets.

How should operators leverage the increase in traffic complexity to their advantage? What targets should operators pursue to get the best QoE? In the voice-dominated networks of the past, the answer was straightforward: operators' main goal was to maximize voice capacity, measured in erlangs. In 3G networks, increasing data capacity and lowering latency became essential targets. In 4G networks, with the emphasis shifting toward QoE, the targets of optimization have become more complex to define.

Metrics like capacity and latency are still crucial, but they have to be optimized for specific traffic flows or, as they are increasingly called, specific network slices, rather than for the overall traffic to and from the RAN. Network slices are logically separated traffic streams that may be defined by traffic type, application, target device, service, or other parameters.

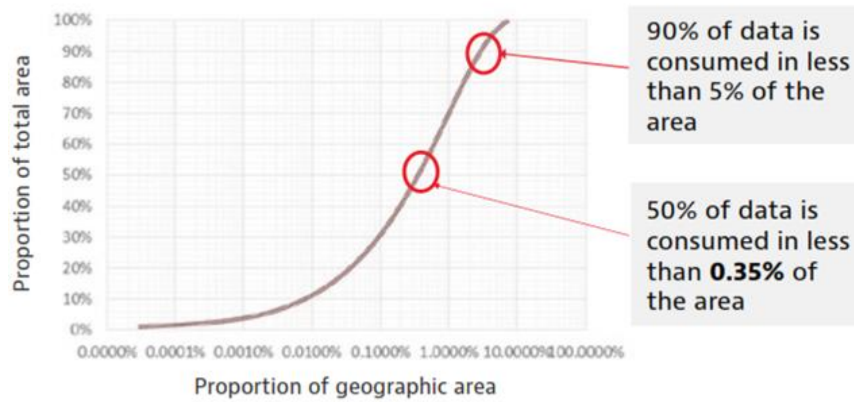

Non-uniformity in demand by location *Source: Viavi*

The goal for operators is no longer to have the lowest latency and highest capacity at the network level, but to have the lowest latency, highest capacity, or both for the network slices that matter most to the operator, or that need it most. This approach may require – as a side effect – that network slices deemed to have a lower priority or less stringent requirements end up having a degraded performance in terms of KPIs, but still retain a good QoE.

While this approach increases the complexity of traffic management, it opens new opportunities for mobile operators to allocate network resources in a more efficient way, which if implemented properly should raise the QoE within the existing network – thus removing or postponing the need for capex for capacity expansion. It also enables mobile operators to define a traffic management strategy as a differentiator from other operators, and use it as a competitive tool.

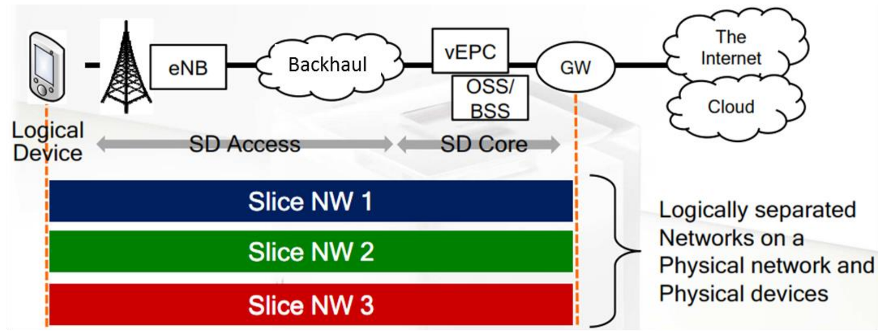

Source: KDDI

Sliced networks for connected cars

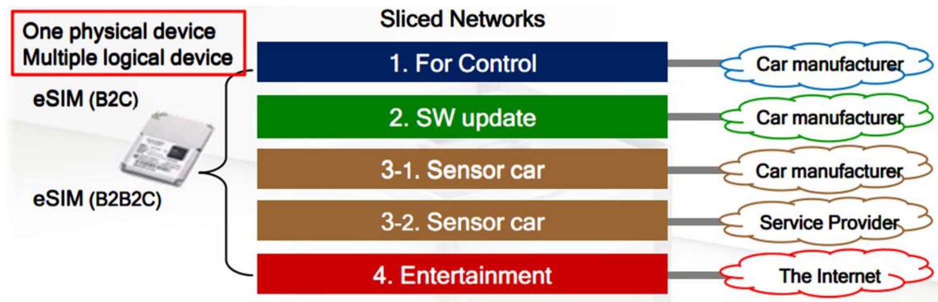

Sliced networks

Source: KDDI

4. Backhaul has to support application-based, real-time traffic management

QoE metrics take center stage in backhaul assurance

As operators learn to deal with more complex and uneven traffic distribution in real time, mobile backhaul has to work within this new framework for performance assurance and traffic management, and avoid becoming the bottleneck that degrades QoE. To do so, backhaul has to be more than a high-capacity pipe. It has to accommodate different sources of traffic and meet the different requirements set by factors such as application type, location, RAN conditions and policy. This has to happen in real time to be effective.

While QoE metrics gain prominence in assessing backhaul performance, they do not directly drive the assessment of backhaul performance. Operators have to relate QoE measurements to KPIs and to the performance of different elements in the network. QoE metrics, though, are difficult to quantify because they are inherently more subjective than KPIs, and there is no industry-wide definition of QoE measurements for data traffic. Even more challenging is the need to relate QoE to network performance – including backhaul performance. Low QoE for video, for instance, may be due to problems with the handset, RAN congestion, backhaul limitations, policy enforcement, or a bottleneck in the interface with the internet if the video is not cached.

Backhaul assurance is crucial to ensuring that backhaul supports the new mobile operator requirements. Along with other types of performance and service assurance, it has to move beyond averaged historical KPIs in order to identify and resolve performance issues in real time, at the granularity level that is required. To succeed in this task, backhaul assurance has to work within the wider context of end-to-end network assurance. When the operator spots an issue that degrades network performance or QoE at the end-to-end level, it has to identify the source within the network. Backhaul assurance is one of the tools operators can use to go deeper in their assessment of network performance, and either exonerate backhaul or establish its role in the problem.

RAN evolution expands backhaul requirements

Multiple RAT interfaces. LTE networks coexist side by side with 2G and 3G networks, with Wi-Fi for both residential and workplace offload, and with carrier Wi-Fi. LTE unlicensed is the latest addition to the mix, and although it is a version of LTE that works in the 5 GHz unlicensed band, it introduces significant differences from LTE in licensed bands, partially due to the support of LAA for listen-before-talk, or LBT, to manage interference with Wi-Fi.

More spectrum bands. Operators need and use more spectrum to meet the increase in data traffic. Carrier aggregation enables operators to use licensed spectrum they own, or can acquire, to transmit efficiently within multiple bands.

Operators are more eager to use unlicensed spectrum with carrier Wi-Fi, LWA or LTE unlicensed on an opportunistic basis, because unlicensed spectrum provides a valuable increase in capacity where those bands are not congested.

Regulators are trying to allocate additional spectrum for mobile traffic – e.g., the 3.5 GHz band in the USA. With 5G, mobile operators hope to use spectrum above 6 GHz, which can support very high capacity in dense environments.

Small cells and other sublayer elements. Densification is necessary to increase network capacity to meet increasing traffic demand. In addition to outdoor small-cell deployments, it will include indoor femto-cell and small-cell deployments, DAS, and carrier Wi-Fi networks.

SON. To manage the coexistence of multiple elements with overlapping coverage areas, automation is necessary to fine-tune the RAN in near-real time. SON treats the network elements and capacity as dynamically changing, and modifies RAN settings to optimize the use of network resources.

5. Multiple RATs, bands and layers coexist in HetNets

Backhaul assurance to operate across RAN elements and backhaul solutions

As traffic and traffic management solutions evolve, so do the RAN infrastructure and its operations. In the RAN, the transition is toward less homogeneous networks in which multiple elements coexist and are increasingly integrated.

Deeper integration across networks – e.g., LTE and Wi-Fi – allows mobile operators to allocate traffic to specific RAN resources, depending on the capabilities of RAN elements, real-time RAN conditions, subscriber location within the footprint, demand, and policy. The flexibility in managing traffic flows within the RAN makes the effective RAN capacity dynamic and affects backhaul requirements, which change correspondingly in time.

Operators have to ensure the backhaul meets the RAN requirements during network deployment, but as RAN elements change, they have to check that RAN requirements continue to be met. This is especially true in small-cell deployments with multi-hop backhaul, in which cells can be added to a local topology (e.g., hub-and-spoke or mesh topologies) more frequently than in a macro-only scenario.

The heterogeneous mix of RAN elements creates a more complex environment for backhaul assurance, because backhaul requirements vary for each element. Monitoring and troubleshooting HetNets, especially when they include a small-cell layer, have to take into account factors such as load sharing, aggregation, visibility and infrastructure sharing, which are less relevant or do not apply in a macro-only environment.

Small cells' impact on backhaul requirements
The higher number of RAN endpoints increases the need for scalable, low-complexity, cost-effective solutions, which nevertheless provide full functionality, resiliency and high capacity.
Infrastructure installed on non-telecom assets, closer to the ground but close to an aggregation point, imposes limits on the choice of backhaul solutions. At many locations, fiber is not available or cost effective, and LOS or NLOS wireless backhaul has to be used instead. Multiple backhaul solutions with varying performance characteristics are often deployed within the same footprint, increasing the complexity of monitoring and troubleshooting backhaul.
Multi-hop backhaul in hub-and-spoke or mesh topologies further increases the complexity of backhaul requirements and management. Requirements vary, and visibility may be lost or limited at different locations within the local network.
Small-cell networks are designed to grow organically as demand grows, with the addition of small cells to the existing footprint as the need arises. In a hub-and-spoke or mesh topology, such additions often change the backhaul requirements of multiple links within the network.
The introduction of the X2 interface in LTE networks to coordinate transmission among overlapping or adjacent network elements allows mobile operators to improve RAN resource utilization, but generates higher levels of signaling and imposes additional requirements – especially for latency – in the backhaul. X2-based signaling remains in the RAN – it is not sent to the core – making it difficult for mobile operators to monitor it and troubleshoot any problems that may originate from it.
Neutral-host models are emerging to make small-cell deployments cost effective, scalable, and easier to deploy and manage. They typically require a shared backhaul link managed by a third-party service provider, which may or may not be the neutral-host provider. While this arrangement gives operators flexibility and cost reduction, it limits their visibility into the backhaul up to the aggregation point, and possibly further if transport from the aggregation point to their core network is shared.

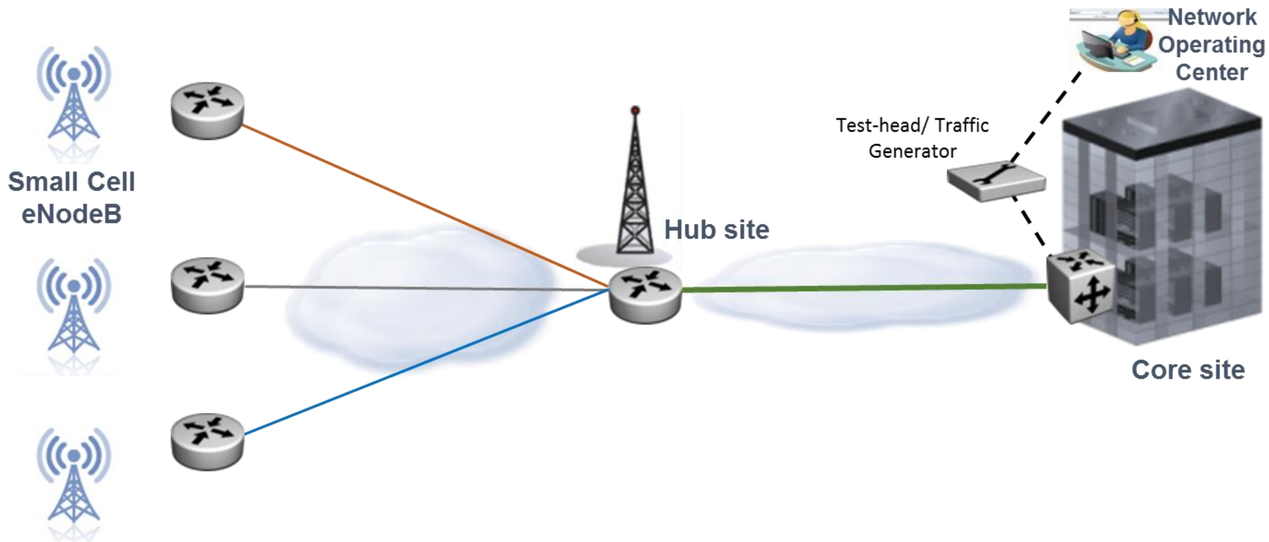

Small-cell hub-and-spoke connectivity and Ethernet backhaul

Source: Viavi

6. Ethernet backhaul is cost effective, but OAM can be challenging

New backhaul assurance solutions are needed to meet new requirements

Operators no longer use TDM-based private circuits for backhaul. Ethernet MPLS-based backhaul can now deliver scalable, resilient, carrier-grade performance in a cost-effective way and support legacy technologies such as TDM, making it possible to support 2G, 3G and 4G concurrently over the same link.

While the standards include the functionality mobile operators require, they may not provide the network-fault and performance monitoring data that operators need, especially in multivendor environments, or where backhaul is shared or provided by third parties (see next two sections).

In some cases, operators resort to using NIDs that give them more visibility into backhaul performance and better troubleshooting capabilities, but NIDs also introduce additional cost and complexity in the management of backhaul. NIDs' limited scalability and cost can be an issue in macro-only networks, but become a more severe liability in multilayer HetNets, in which the number of endpoints -- small cells or other sublayer RAN elements – and the variety of backhaul solutions sharply increase. Mobile operators have started to deploy smart SFP transceivers as an alternative. They are more cost effective, have a smaller footprint requirement, and allow operators to achieve the monitoring accuracy and resolution required to manage complex backhaul networks.

When deploying small cells, operators face a bigger challenge, because they have to keep costs lower than in the macro network, but their OAM requirements are unchanged. Backhaul assurance becomes all the more important, to ensure that operators benefit from the cost savings of carrier Ethernet backhaul. The increased complexity in traffic composition and distribution, and the need to monitor and troubleshoot performance on the basis of real-time QoE and RAN-condition data, expand the relevance and required functionality of backhaul assurance.

Smart SFP *Source: Viavi*

Drivers for carrier Ethernet and IP/MPLS backhaul

Lower costs

Shared IP backhaul is less expensive than TDM private lines, and provides more flexibility for bandwidth pricing.

Legacy support

MPLS-enabled backhaul supports multiple technologies, including legacy ones such as TDM.

Support for guaranteed SLAs

SLAs may include committed information rates, committed burst rates, excess information rates, and random early discards.

Improved support for QoS

Class-of-service options are supported.

Ethernet OAM standards

These have introduced OAM capabilities to Ethernet to support network-fault and performance management. Key Ethernet OAM standards are:

- IEEE 802.3ah for the access link (Ethernet first mile)
- IEEE 802.1ag for the connectivity layer (connectivity fault management)
- IEEE 802.1aj for managing customer demarcation devices
- ITU-T-Y.1731 (network and service layer OAM)
- RFC-2544 and ITU-T-Y.1564 (service level validation)
- RFC-5357 (Two-Way Active Measurement Protocol, or TWAMP)
- MEF E-LMI to manage the UNI and to auto-configure the CE

7. Backhaul provided by third-party service providers reduces operators' control

Visibility into third-party backhaul is crucial for OAM

Over recent years, mobile operators have faced strong pressure to lower per-bit costs, because they have to carry a much heavier traffic load but do not see a corresponding increase in service revenues. As a result, they have started to accept infrastructure-sharing and neutral-host arrangements that provide cost savings, but also limit the control and visibility they have into their networks.

Backhaul is an example of this. One of the cost advantages of carrier Ethernet comes from the fact that backhaul can be provisioned and managed by a third party, and that it can be shared with other mobile operators or service providers. This cost advantage has been a major driver for the adoption of carrier Ethernet and, as a result, operators are increasingly sharing backhaul and leasing it from third parties.

To preserve performance of their networks, however, it is imperative that they deploy solutions that give them the necessary visibility into the backhaul – initially during the activation testing phase, and subsequently for monitoring and troubleshooting. Frequently, mobile operators want to conduct the testing, monitoring and troubleshooting of backhaul links independently from the service providers, both to confirm the performance data they receive from them, and to collect data that the backhaul provider may not collect.

The data that operators collect from these solutions may enable them to relate backhaul performance to RAN performance and QoE more accurately. In turn, this enables them to manage the RAN more effectively as well, including identifying microbursts. When backhaul assurance is part of an end-to-end network assurance solution, having more-granular information on the backhaul may enable them to identify the source of QoE issues more precisely.

Assessing leased or shared backhaul networks

Ensure that SLA terms are met, during initial deployment and subsequent network upgrades and expansion (e.g., addition of small cells, carriers, or sectors), as well as during regular operations (i.e., monitoring and troubleshooting)

Monitor performance at the application and service levels and in real time, to ensure that backhaul does not become a bottleneck in RAN performance or have an adverse impact on QoE

In shared deployments, ensure that the operator gets access to a fair share of the backhaul resources

8. Implications

With LTE and carrier Ethernet, visibility into backhaul becomes more crucial to preserving end-to-end network performance and QoE

With LTE, backhaul requirements have not just expanded, they have changed qualitatively. Operators need to monitor and troubleshoot their networks in real time, taking into account QoE. They need to do this for the end-to-end network and for the backhaul as well.

Greater traffic complexity and extreme variability through time and space compel operators to monitor and troubleshoot backhaul performance more actively, increasingly in real time or near real time to address performance issues as they arise, or to prevent them.

Operators need to assess backhaul performance with a view of its impact on QoE. Performance metrics averaged over time and across the network are still useful, but operators are transitioning to network monitoring and assurance platforms that can operate at the traffic flow or network slice level, based on factors such as traffic type, service and application.

The RAN has become more complex, dynamic and dense because of the introduction of HetNets; densification with small cells, DAS and other sublayer RAN elements; and LTE Advanced functionality such as CA. These changes affect backhaul requirements and increase the complexity of monitoring and troubleshooting backhaul performance.

Carrier Ethernet delivers cost savings, but also bring challenges for managing backhaul performance. Operators need to actively control OAM to ensure that backhaul links do not become the bottleneck and, when that does happen, operators must reliably identify the causes.

Backhaul has become a service that operators often lease from third parties and share with other service providers. Operators need to make certain they have the backhaul assurance tools to verify that the backhaul's performance meets the agreed SLAs, and that they get their fair share of backhaul resources.

How can backhaul assurance meet the challenges of real-time traffic monitoring and optimization?

A conversation with Assaji Aluwihare, Director of Strategic Marketing, Viavi Solutions

We further explored the role and relevance of backhaul assurance in this conversation with Assaji Aluwihare, Director of Strategic Marketing at Viavi Solutions.

You can watch the interview here.

Monica: Assaji, can you tell us what Viavi Solutions is doing to help operators run their mobile networks, and with backhaul assurance specifically?

Assaji: Viavi Solutions provides end-to-end visibility solutions for various applications on the network, from the handset all the way to the core of the network. One piece of that is our backhaul solutions.

Our backhaul solutions are designed to ensure the quality of the backhaul network throughout the life cycle: from the time you build that backhaul and you turn it up, to when you monitor and optimize and troubleshoot it.

Monica: Most of the time operators are worried about having enough capacity in their RAN, but the backhaul can also be the bottleneck. You do need to monitor it. How has monitoring the backhaul changed over, say, the last ten years, as mobile networks have evolved?

Assaji: Initially backhaul was TDM based, when we had 2G and 3G networks and the primary application was voice. With the emergence of 4G LTE, which is IP based, Ethernet became the preferred backhaul technology. What happened is that the Ethernet network is not only carrying 4G now, but it's also carrying the 2G and 3G traffic as a pseudowire network.

Today, operators also provide networks for business services or residential broadband, and the Ethernet network is being shared with those applications as well.

The Ethernet network itself is evolving. It is getting more complicated. The use cases that the operators are seeing on the network – and certainly those that mobile operators are seeing – are getting more complicated.

For example, mobile operators now are seeing many, many, many more devices, many different types of devices. When the iPhone was introduced, that created a completely different kind of behavior on the network.

You're seeing different new applications, hundreds of new applications being released every day that drive different signaling schemes, different activities on the network. You're seeing small cells coming into play, driving much more complex backhaul and aggregation schemes.

What this means is that the network is becoming very complicated. It's not only a question of capacity, it's a question of complexity. Managing and prioritizing the different types of traffic on that backhaul networks becomes very difficult.

That's what the operator's new challenge is. We're seeing this in the marketplace; all the mobile operators in the world are either evolving to have a backhaul management or monitoring strategy, or growing what they already had, adding more functionality.

The evolution of VoLTE is another driver. VoLTE is now bringing in conversational voice traffic onto a network that was designed to carry data, and that is a challenge as well.

Monica: The need to monitor backhaul performance has become even more important where the backhaul is shared, not just among different applications, but also among different providers, or leased from third parties.

Assaji: Yes, absolutely. As operators are building out their network to increase their coverage and capacity, they're putting cells in more different places – in places where they didn't traditionally have backhaul access. They are now leasing different types of backhaul from different third-party operators, and in some cases they're actually sharing backhaul: two or more operators are sharing the backhaul.

This drives additional requirements. We have a customer who is using our solution to ensure that in a shared backhaul scenario, they are getting the fair share of bandwidth and performance that they're paying for.

In the case of third-party operators, there are wireline providers that are selling backhaul connectivity to these mobile operators. In those cases these mobile operators want to deliver and sell SLAs that include not only basic uptime, but also more and more complex SLAs that guarantee latency and other performance metrics.

In those cases, the wireline operators are using these assurance solutions to be able to characterize that backhaul and provide that data to their customers, the mobile operators.

We're also finding that the mobile operators themselves, in many cases, are driven by the voice quality or the service quality that they want to deliver. In some cases, these operators are being ranked, in the press and other venues, based on the quality of their networks.

Some of those operators are coming to us to help them ensure the quality of their backhaul, even though they do not own that backhaul. They're deploying end-to-end monitoring solutions to ensure the backhaul from the cell site to the core through a third-party network.

Mobile operators are using the assurance solution to ensure performance, even though they are not running that backhaul network. In that scenario, they'll monitor their backhaul, and if they see a problem, they will push the issue to their backhaul provider.

Monica: How have your solutions for backhaul assurance evolved through the years to address all these new requirements?

Assaji: There are many standards-driven processes for managing a backhaul network. The Metro Ethernet Forum drives some of these processes.

Typically a backhaul network has three phases of its life cycle. There is the activation phase of that network. Then there's a second operational phase, which is when the network is running and you do performance monitoring on that network. The third phase is the fault-finding and troubleshooting phase of managing a network.

Each of these is a standards-driven activity. In the activation phase, standards such as RFC 2544 and Y.1564 are prominent. In the performance-monitoring phase, the common Layer 3 standard is known as TWAMP. These are very important methodologies for testing backhaul. At Viavi Solutions, we have end-to-end solutions for all these phases.

As the networks evolve, we augment these solutions to meet new needs. Latency is becoming very critical for conversational voice with VoLTE. The latency from end to end is important, but it's also important to measure the latency in both the uplink and the downlink directions to make sure that the customer experience is managed.

Another problem that we're seeing in these networks is microbursts. In TCP networks, you tend to have the clumping of packets that cause bursts for short periods in the network. Typical performance-monitoring systems use synthetic traffic, so you engineer the testing and the performance based on what you expect the network to be doing.

As the network becomes more complicated, you can't engineer the network to manage this anymore. What you find is that you have things like microbursts that you can't detect with traditional monitoring schemes.

We have developed a solution in which we've augmented these traditional methods with our PacketPortal Intelligence Visibility device, where we can use real, live traffic to characterize the performance of the network. We can find microbursts in real, live traffic.

With the TWAMP-standard mechanism, we can also test from end to end, from the cell site to, say, the MSC. Our solution can segment the test and pinpoint where the exact problem is.

If end-to-end TWAMP finds an issue, you have to go to each of the individual spots in the network, and do tests between intermediate points. With our testing scheme you can immediately pinpoint where the problem is. That's enabled by our PacketPortal Intelligence Visibility devices. We have a solution that has really evolved past where the standards are.

Monica: You mentioned the spikiness of the data. It's clearly a concern in the RAN, but it's also a concern for the backhaul as well, and you need to measure it. But in order to measure it, you have to get away from the average capacity, the average performance. You have to start looking at this in real time. But what is the right time resolution?

Assaji: That's a great question. If you're measuring throughput in a network, typical systems today use network-element data. Network-element data measures throughput on 15-minute or 5-minute cycles at best.

If you want to find a microburst, you really need to be measuring at a millisecond level. Our system is able to measure at the tens or hundreds of a millisecond level, and find microbursts in the network at a resolution far more granular than where those would be averaged out in a typical network-element–based system.

Now, performance monitoring – and how real-time and granular the data needs to be – really does depend on the application. In financial applications, say if you're delivering backhaul to a financial customer, you need to have very high-resolution performance-monitoring data, which we can provide.

It depends on the application. Finding microbursts has to be in the millisecond range. But performance data has been historically gathered on the order of minutes.

Monica: As we are move to real-time, you have to deal with much more frequent measurements of network performance, and you tend to have much more data. There's complexity, and there is sheer volume. What do you hear operators say, and can you help them leverage the information that is valuable to them without them being overwhelmed by it?

Assaji: Viavi Solutions has a wide range of solutions. We have everything from testing the RAN all the way through the core, for signaling testing or mobile assurance through the core to the EPC itself. We have a wide range of solutions.

Our vision is that, as networks become more complex and virtualized, operators are going to need near-real-time intelligence, not only to provide QoE to the customer, but also to be able to understand the customer's experience and react to it quickly. And also to feed the systems that are managing this software-driven network.

We are building systems at the backhaul level and across the entire network. Our solutions gather data from across the network at a high resolution in terms of time, and feed them into a mediation platform that normalizes and correlates the data, and that in turn feeds this data into other systems that can generate QoE-driven KPIs, or feed third-party systems for policy, orchestration or similar functions. That's the vision we're driving to across all our product lines at Viavi.

Monica: The ability to have end-to-end visibility is crucial, because even if you see that the backhaul is not performing as it should, the problem might still be somewhere else. You need to first be able to identify the source of the problem. You have to do this on an end-to-end basis.

Assaji: That's exactly right. You have to be gathering data from points throughout the network, feeding it into a system that can then correlate the data in real time. If you see a bad user experience on a certain application, you can then work through the network and the stack to find exactly where that problem is. Those are the types of solutions that we're bringing to market.

Monica: What are the most compelling needs that you hear about from operators? And do they vary across the globe?

Assaji: They do vary across the globe. Some operators are just evolving to 4G, and others are evolving to VoLTE. Most operators are now concerned about 4G and VoLTE, the complexity of the data network, and performance monitoring. You see that over and over again.

Virtualization is a big challenge for many operators. A few are ahead, but everybody is thinking about it. How will assurance solutions become virtual, so that they can be deployed on x86 devices in a data center? But, also, how do they evolve to become a real-time intelligence system that feeds these virtualized and software-driven networks?

Those are primary issues that we see our operators dealing with right now, pretty much across the globe. Everyone sees that coming. There are different states of the life cycle. But it's evolving to the same problem.

Monica: Clearly latency is very important for voice, and this has always been the case. Except that before, it was switched, and now it's packetized. How does that change the way you do backhaul assurance on voice traffic? Are there different requirements there in order to make sure that voice works properly?

Assaji: With packetized voice, latency becomes less predictable. That's why jitter, latency and packet loss on a stream-specific basis is very important. Quality of service is implemented at different levels for different types of services.

Because of this complexity, the backhaul assurance solution has to be able to handle and monitor the performance of multiple streams and characterize these streams differently, so that the different services get the right amount of priority.

It is a complex situation. You engineer the network based on how you expect the network to behave. But when there's so much other complexity in the network, you have to continuously monitor to ensure that that's happening.

Monica: HetNets add another dimension for change. In HetNets you have different types of RAN elements. You might have a macro, you might have a small cell or a micro cell, in indoor or outdoor locations, or a DAS.

Mobile operators often use multiple backhaul technologies. Does your solution work across all of them in the same way? Are there differences in how you monitor backhaul across technologies – for instance between wireline and wireless?

Assaji: Our solutions do take into account the complexity of the backhaul due to HetNets and overlay networks, and due to the complexity of the small cells. All of them drive more complex aggregation schemes on the backhaul network. You see hub-and-spoke architectures, and similar ones.

We take all of that into account. If you deploy our PacketPortal IV devices at each of these aggregation points, we can build a view of the network's performance in a unique way that is different from typical systems that are based on sending packets end to end, to an end point. It is the complex aggregation schemes, the complex backhaul, where we provide highly differentiated visibility into performance.

Monica: Is virtualization going to have a major impact on backhaul assurance?

Assaji: Absolutely. Backhaul monitoring is based on some active testing where you send packets over the network. All of that has to change. We have to evolve that, to virtualize how we do that.

Not only that, but again, feeding systems that drive real-time management of a software-driven network, such as a policy system or an orchestration system, becomes critical. A backhaul monitoring system will have to get information in real time, back to your initial question. You have to get that data and feed those virtualized systems to drive those software networks.

Monica: This is going to be, as we move forward, there is more virtualization. Operators need to work more in real time than they currently do.

Assaji: Right. More meshy clouds interacting with each other. They're going to have to know how to measure performance – not only end-to-end performance through those clouds, but at various interaction points inside those clouds, as well.

Monica: What's going to happen over the next few years? What will Viavi Solutions be focusing on in terms of innovation and new solutions?

Assaji: Our big focus at Viavi Solutions will continue to be on end-to-end visibility. At every point in the network, throughout the network, from the handset all the way to the core, and on the backhaul side, we focus on bringing in more innovative performance-monitoring schemes that can deal with real-time data and network virtualization. Virtualization is a key focus for the evolution of our solution – not only to virtualize, but to deal with the constructs of virtualized networks as well.

Monica: Because a virtualized network is one where equipment changes to support the same function, you have a network that is much more inherently dynamic.

Assaji: That's exactly right. You can't instrument a virtualized network the way you could a traditional network, where you could monitor or test at two different points. You have to evolve your solutions now to be able to dynamically monitor across the mesh and understand path changes, things like that, very quickly.

Monica: I guess that really changes the traditional split between tests and monitoring, because you're basically testing and monitoring at the same time continuously.

Assaji: This is a perfect point. It is a completely different process change. Let's be realistic. There is still a physical instantiation of the network. There is a test phase that you have to do when that physical network is built out.

But after that, the network is truly dynamically changing, so there is no delineation between a testing phase and a monitoring phase. As the network changes dynamically, we have to dynamically ensure that the paths are verified as they are built. And then monitor them as you go on.

As the orchestrator changes the path, we have to test that everything is done appropriately, and then monitor that path – potentially for a very short time before it switches again.

Glossary

2G	Second generation
3G	Third generation
4G	Fourth generation
ARPU	Average revenue per user
B2B2C	Business to business to consumer
B2C	Business to consumer
BSS	Business support systems
CA	Carrier aggregation
CAGR	Compound annual growth rate
CE	Customer Edge
DAS	Distributed antenna system
E-LMI	Ethernet Local Management Interface
eNB	Evolved NodeB
EPC	Evolved Packet Core
eSIM	Embedded subscriber identity module
HetNet	Heterogeneous network
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of things
IP	Internet Protocol
ITU	International Telecommunication Union
IV	Intelligent Visibility
KPI	Key performance indicator
LAA	Licensed-assisted access [LTE]
LBT	Listen-before-talk
LOS	Line of sight
LSP	Label switch path
LTE	Long Term Evolution
LWA	LTE Wi-Fi aggregation
MEF	Metro Ethernet Forum
MPLS	Multiprotocol Label Switching
MSC	Mobile switching center
NID	Network interface devices
NLOS	Non line of sight
OAM	Operations, administration and maintenance
OSS	Operations support system
OTT	Over the top
QoE	Quality of experience
RAN	Radio access network
RAT	Radio access technology
RFC	Request for Comments
SFP	Small form-factor pluggable [transceiver]
SLA	Service level agreement
SON	Self-organizing network
SW	Software
TCP	Transmission Control Protocol
TDM	Time division multiplexing
TWAMP	A Two-Way Active Measurement Protocol
UNI	User Network Interface
vEPC	Virtual EPC
ViLTE	Video over LTE
VoIP	Voice over IP
VoLTE	Voice over LTE

About Viavi Solutions

Viavi Solutions (NASDAQ: VIAV) software and hardware platforms and instruments deliver end-to-end visibility across physical, virtual and hybrid networks. Precise intelligence and actionable insight from across the network ecosystem optimizes the service experience for increased customer loyalty, greater profitability and quicker transitions to next-generation technologies. Viavi is also a leader in anti-counterfeiting solutions for currency authentication and high-value optical components and instruments for diverse government and commercial applications. Learn more at www.viavisolutions.com and follow us on Viavi Perspectives, LinkedIn, Twitter, YouTube and Facebook.

About Senza Fili

Senza Fili provides advisory support on wireless data technologies and services. At Senza Fili we have in-depth expertise in financial modelling, market forecasts and research, white paper preparation, business plan support, RFP preparation and management, due diligence, and training. Our client base is international and spans the entire value chain: clients include wireline, fixed wireless, and mobile operators, enterprises and other vertical players, vendors, system integrators, investors, regulators, and industry associations. We provide a bridge between technologies and services, helping our clients assess established and emerging technologies, leverage these technologies to support new or existing services, and build solid, profitable business models. Independent advice, a strong quantitative orientation, and an international perspective are the hallmarks of our work. For additional information, visit www.senzafiliconsulting.com, or contact us at info@senzafiliconsulting.com or +1 425 657 4991.

About the author

Monica Paolini, PhD, is the founder and president of Senza Fili. She is an expert in wireless technologies and has helped clients worldwide to understand technology and customer requirements, evaluate business plan opportunities, market their services and products, and estimate the market size and revenue opportunity of new and established wireless technologies. She has frequently been invited to give presentations at conferences and has written several reports and articles on wireless broadband technologies. She has a PhD in cognitive science from the University of California, San Diego (US), an MBA from the University of Oxford (UK), and a BA/MA in philosophy from the University of Bologna (Italy). You can contact Monica at monica.paolini@senzafiliconsulting.com.