HEAVY READING

White Paper

Adaptive QoE Scorecards for the IP Realm

Prepared by

Jim Hodges
Senior Analyst, Heavy Reading
www.heavyreading.com

on behalf of

PROCERA

www.proceranetworks.com

October 2015

Introduction

On some level, network operators have been attempting to quantify end-user subscriber service experience for decades since the introduction of time-division multiplexing (TDM)-based digital switching. Given the voice-centric nature of the services landscape, measurements tended to be heavily based on sampling voice quality and creation of mean opinion score (MOS) methodology. As the industry started the IP migration in the fixed network domain almost 15 years ago, it became clear that new approaches to measuring or scoring user experience were necessary given the data-centric nature of these early voice over IP (VoIP) networks.

As a result, we witnessed the introduction of the IP scorecard, which was designed to estimate user quality of experience (QoE) by focusing on assessing underlying network performance by *sampling* metrics, such as packet losses, latency and throughput. While this approach has merit from a benchmark perspective, like the voice MOS approach, it is now facing methodology and overall relevance questions for several reasons.

At the top of the list, we believe, is the fact that existing QoE scorecards are heavily network focused and make broad brush assumptions on how network performance translates into end-user QoE, rather than actually measuring end-user QoE. Therefore, the industry is once again moving to adopt new QoE measurement and scorecard approaches that reflect this fundamental shift.

Consequently, the focus of this white paper is to examine the value proposition associated with leveraging powerful real-time measurement capabilities to provide a holistic view of subscriber QoE in the IP realm.

The Rise of Adaptive QoE

As touched upon above, even though the concept of IP network-based QoE is relatively well understood, given the speed of innovation in IP services, QoE is evolving into a more powerful and adaptive construct in response to several IP evolutionary factors. Before we examine the impact of these factors in detail, it will be helpful to first document the differences between current IP QoE methodology and this new approach. While these are numerous, one fundamental difference stands out.

In the current IP QoE approach, shown in **Figure 1**, QoE calculations are based on sampling transport network performance and then applying assumptions on application performance and ultimately user experience. For example, in this scenario, since there are two sample points, there are only two potential QoE outcomes that can be assigned to subscribers. The outcome is the creation of a top-down QoE model that delivers a common and theoretical view of user experience for groups of users vs. individual users.

Figure 1: Current QoE Reference Architecture

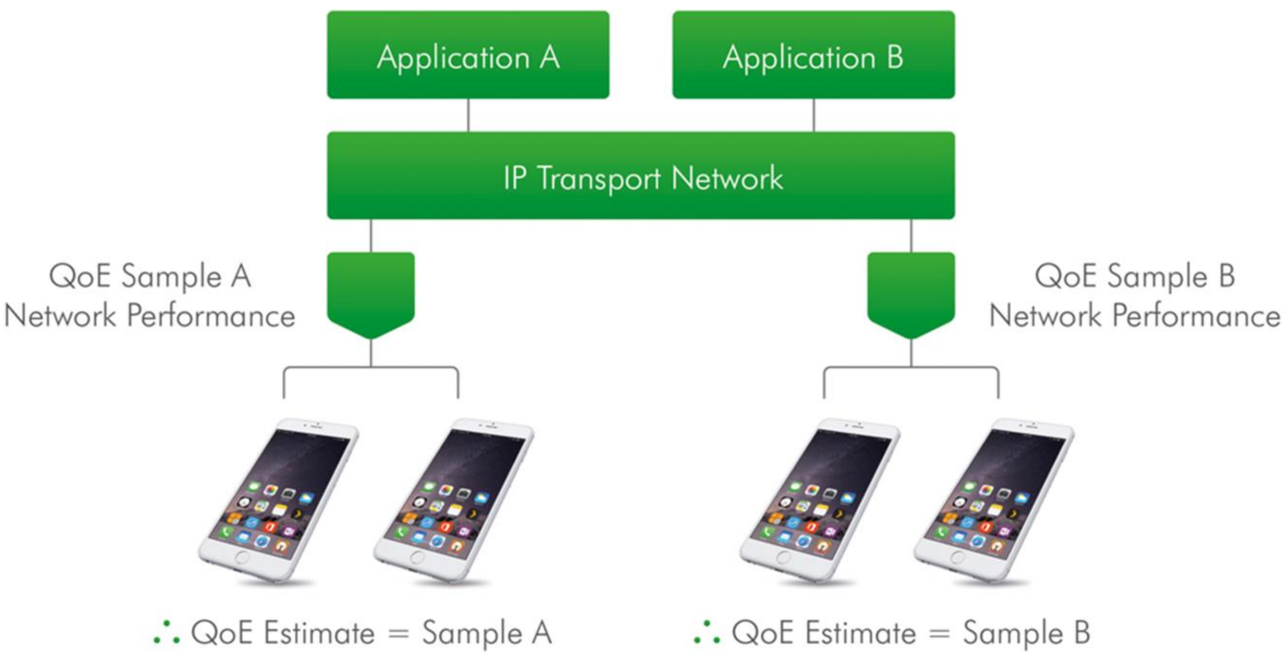

Source: Heavy Reading

Given these shortcomings, a new approach, which we refer to as Adaptive QoE (A-QoE) is gaining market momentum. This adaptive approach reflects the fact that QoE is no longer based on applying network sample data; rather, the QoE model adapts to subscriber usage patterns, and changes to create individualized QoE subscriber metrics based on both network performance and the applications that they consume.

A key benefit of this approach is that QoE can be regularly updated based on which applications are being utilized and their performance. This, in turn, enables the network operator to implement policy rules to optimize application and network performance in real time to adapt to changes in traffic patterns. As a result, this approach delivers a more complete and realistic network level view of performance since the assessment is made by aggregating all the individual subscriber QoE metrics vs. sampling and estimating.

Therefore, network and application performance assessment is much more complete, given that it utilizes a bottom-up vs. the top-down design methodology of the current approach (see **Figure 2**). By that, we mean that each subscriber QoE measurement reflects their individual mix of applications and network performance based on location.

Figure 2: Adaptive QoE Reference Architecture

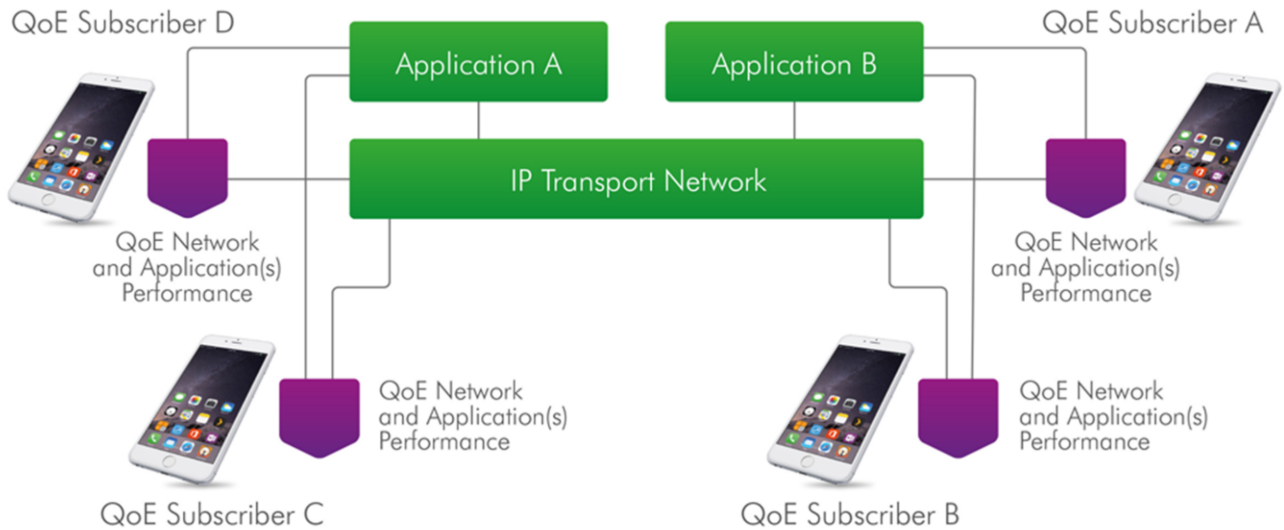

Source: Heavy Reading

Although A-QoE can be applied to enhance existing IP networks, there are several evolutionary factors that we believe will influence IP network design and make the implementation of A-QoE even more compelling in the near future. These factors, which we address below, are:

- Cloud NFV-enabled data centers
- Mobile IP connectivity
- Application evolution
- Self-provisioning services

Cloud NFV-Enabled Data Centers

There is little if any disagreement within the telecom industry that network functions virtualization (NFV) will indelibly change how applications are both deployed and developed. And given that we are now entering the NFV commercialization phase, the future is now. NFV, therefore, will also change the very nature of how QoE is assessed, since the use of distributed application VNFs, spun up based on demand, is not readily applicable to the existing model, which assumes that application software is centralized and deployed in a static location.

However, with NFV, applications are not bound by network boundaries. Moreover, this distributed software application model could have considerable deviation in software performance, so measuring actual application performance of individual users is crucial for ensuring customer service is not compromised on any level. As a result of these changes, we believe that NFV will need to rely on A-QoE techniques to adequately capture the impact of *all* the dynamic factors that shape user experience.

Mobile IP Connectivity

A second factor that will shape the future of QoE measurement is the definition and ultimately deployment of 5G networks. While it is far too early to speculate on what a finalized reference architecture will even look like, there is little doubt that 5G will take mobile broadband to the next level in terms of increased connectivity speeds and lower-latency application performance. And certainly, as we have seen with 3G and 4G, users will not only come to expect lower latency and higher-speed access over previous generational deployments, they will also expect that network performance will translate into a noticeably improved personalized user experience.

Factor in the NFV considerations noted above and it becomes clear that the network of tomorrow is no longer a pipe dream that is decades away. Instead it is rapidly coming into focus and will heavily rely of real-time individual subscriber-level QoE measurement techniques.

Application Evolution

The advancement of mobile IP connectivity will not only enhance the performance of existing services, it will also fuel major changes in the applications themselves. And while it is extremely difficult to predict which applications will succeed or fail in the future, Heavy Reading believes that technologies such as virtual reality (VR) will become more prevalent beyond simply gaming, tying together connected applications such as telemedicine and enhanced HD video content delivery. Since VR applications are extremely sensitive to QoE impairments, network sampling approaches will be of little value in assessing the overall immersive user experience for these critical applications.

Self-Provisioning Services

The manner in which subscribers will gain access to applications is also changing. More and more, network operators and subscribers are embracing a self-provisioning automated model for new services. This not only represents a lower-cost provisioning model, it also aligns with the NFV service agility model to shorten time to market.

The other advantage for network operators is extending field of sight for customer care, enabling front-line customer-care staff to more proactively manage network issues, offer loyalty promotions and minimize subscriber churn by giving them a complete profile of the applications the subscriber has recently used, as well as the real performance level and QoE metrics of these applications. Again, it's worth noting that this is only possible with A-QoE vs. the traditional network QoE approach.

Introducing the Adaptive QoE Scorecard

A direct outcome of the transition from QoE to A-QoE is the introduction of a new type of QoE score, the A-QoE scorecard that captures and summarizes user experience on an individual level. The A-QoE scorecard is incrementally more powerful than traditional scorecards since it leverages additional software intelligence and analytics capabilities to assist network operators to optimize both application and network resources in real time. It also leverages this application intelligence to rationalize network performance and user experience. What we mean by that is with the current traditional QoE scorecards, it's possible to have high levels of network performance and poorer levels of application performance due to multiple factors, including device performance and even policy control network settings.

Since an A-QoE scorecard is created for every user, it's possible to create a universal frame of reference and identify which users and which of their applications are suffering from these unique impairments, so that they can be resolved accordingly. In addition, since every subscriber receives a scorecard, it is possible to compare the actual levels of performance on a subscriber level to achieve an "apples to apples" comparison vs. the more subjective, network-based estimate of subscriber user experience. The inverse is also true: When network performance is poor, it's possible to determine which particular user(s) and which applications are most affected based on A-QoE scorecard metrics.

Another important consideration, as shown in **Figure 3**, is that as network operators improve the level of network performance, the A-QoE scorecard metrics of each subscriber should also generally improve and provide valuable and granular insight into the overall effectiveness of additional network investment and design on a bottom-up basis. The business value of this data is that it provides real-world input capturing the potential to upsell premium threshold services (e.g., 150Mbit/s download service) on either a temporary or subscription basis.

Figure 3: The Adaptive QoE Scorecard

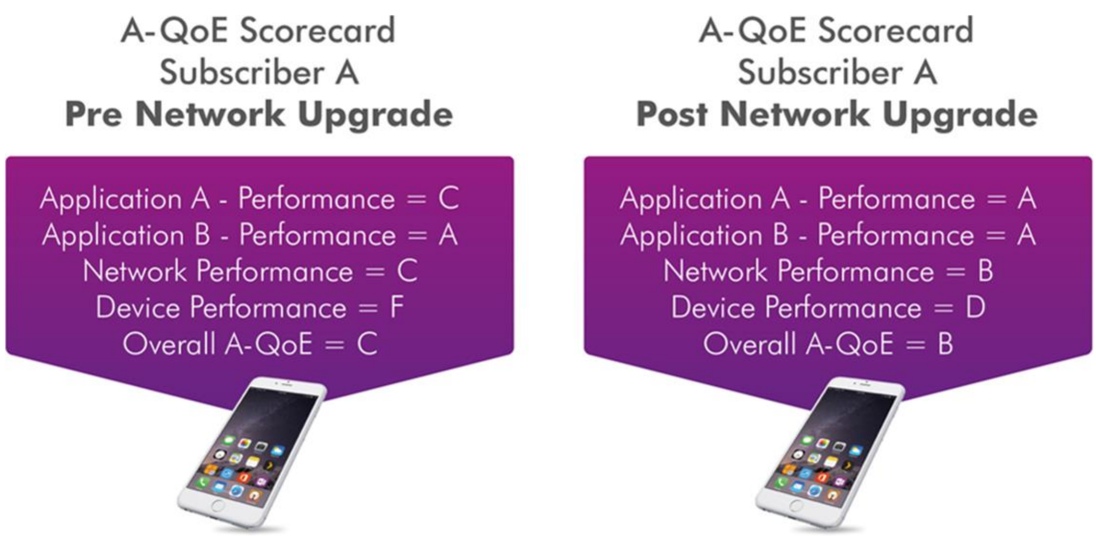

Source: Heavy Reading

A-QoE Scorecard Use Cases

In this final section of the white paper we document in greater detail the technical and business value that the implementation of A-QoE scorecards can have on networks. To accomplish this, we analyze two key use cases below.

Use Case 1: RAN/WiFi Harmonization

When mobile operators first started to roll out WiFi access, the original intent was to offload capacity from 3G networks and to provide alternative access options for roamers. And as WiFi access improved over time, subscribers have become more receptive to using WiFi for primary application support. The dilemma for network operators is that, while subscribers perceive the value, WiFi has become difficult to monetize, given original pricing strategies. Moreover, given that WiFi networks were built as overlay networks, with little monitoring or performance measurements, it is still very difficult to truly measure QoE, even though subscribers increasingly expect their mobile operators to deliver consistent QoE levels on their mobile services.

Figure 4: RAN/WiFi Harmonization

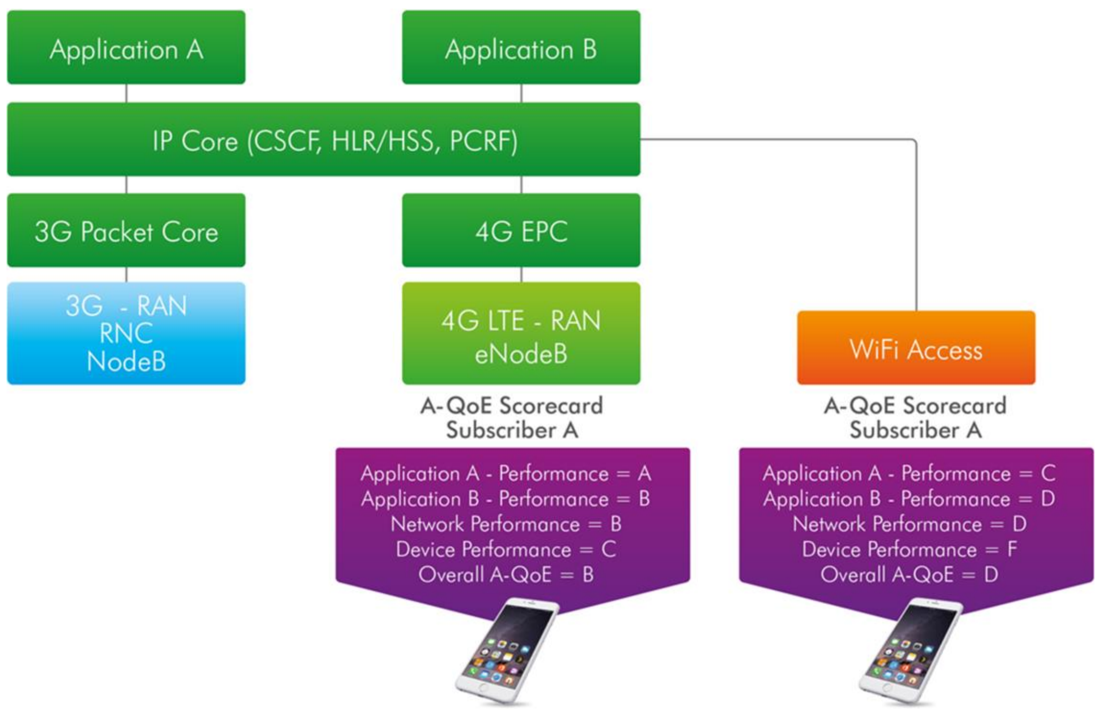

Source: Heavy Reading

Therefore, as an industry, we are now moving to a tighter integration between WiFi and 3G/4G technologies to support handover and common QoE metrics to assist in the monetization process. This is an important step, since many mobile operators are currently running 3G, 4G and WiFi networks that are stitched together based on a "patchwork" network model that hinders monetization and even uniform service delivery. As a result, A-QoE scorecards can deliver a strong value proposition, since they measure QoE for each subscriber for each radio technology.

This not only provides granular data on network and application performance, it opens the door to monetize on-demand access upsell opportunities. In this case, as shown in **Figure 4**, since the WiFi network access is underperforming, the network operator can offer a premium performance upsell opportunity on the 4G network to Subscriber A. Since A-QoE scorecards are created for the same subscriber, the boost in performance and related increase in user satisfaction level can be measured and quantified.

Use Case 2: Mobile Core Virtualization

We have already touched on the impact on virtualization, which will fundamentally change how both core network and application resources will scale, and even where they are delivered relative to the location of the end user. Given that we have very little commercial data to confirm the capacity and service agility gains upon which NFV is predicated, we believe that A-QoE scorecards will become a crucial validation component. For example, **Figure 5** captures Subscriber A, a 4G subscriber who is active in his home network, mobile and changing locations.

Figure 5: Mobile Core Virtualization

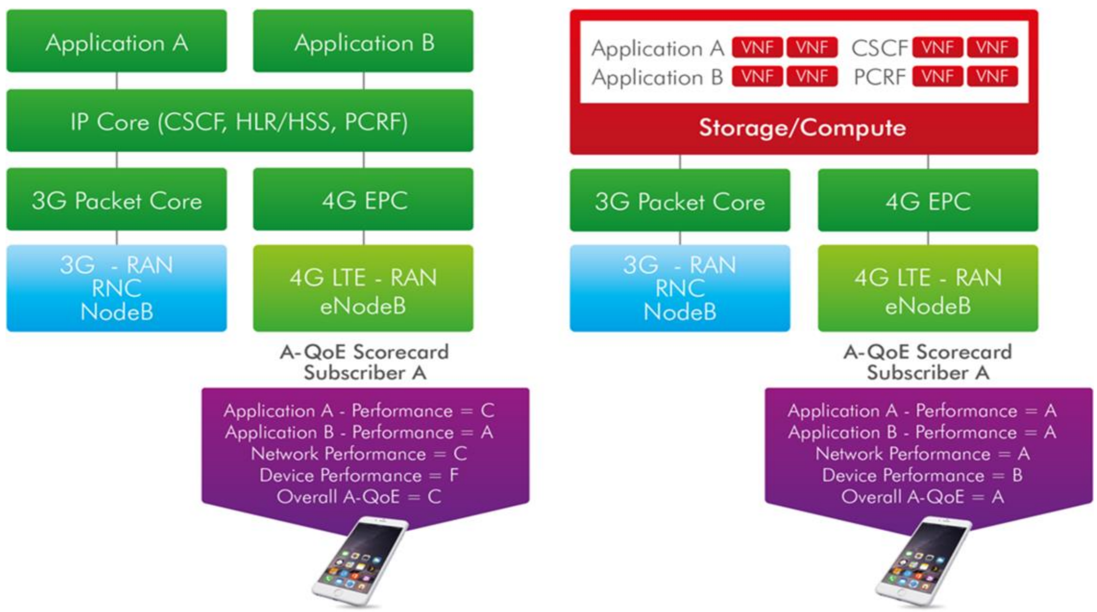

Source: Heavy Reading

At left, the figure shows a subscriber served by facilities-based IMS and application servers; at right, the same subscriber is served by IMS, PCRF and application server VNFs. Since A-QoE scorecards are created for the same subscriber, it provides a definitive view of network and application performance. In this scenario, the virtualized configuration delivers superior performance and, hence, a higher A-QoE score. However, given the number of VNFs required, there is a strong possibility that not all will meet performance thresholds, so it is also crucial to understand which subscribers running which applications fell below A-QoE thresholds. This is especially true as network operators gradually increase load on the virtualized network and applications.

Conclusion

Network operators have never before faced so many new opportunities and so many new challenges. This is reality, and the risk/reward model of the IP realm. Realistically, it will not change for many years until virtualized network penetration reaches a commercial "tipping point."

Therefore, having a deep understanding of subscriber-level QoE embodied in scorecards is no longer simply a "nice to have," but is fundamental to ensure network operators have the correct data to enable them to make both the correct business and technical decisions to drive new revenue streams and grow subscriber base.

Given this dynamic environment, the very definition and scope of QoE scorecards has made a fundamental evolutionary turn. They are no longer generic, network-based benchmarks, but granular and rich sources of actionable data capable of adapting to changes in network, application and device performance.

About the Author

Jim Hodges
Senior Analyst – Core & Edge Network Evolution

Jim leads Heavy Reading's research on the impact of SDN, NFV and D-NFV on the control plane and application layers at the core and edge. This includes the evolution path of SIP applications, unified communications (UC), IP Multimedia Subsystem (IMS), session border controllers (SBCs), Diameter signaling controllers (DSCs), IP exchange (IPX) and WebRTC.

Jim is also focused on the impact of NFV and SDN on data center application delivery controllers (ADCs). He authors Heavy Reading's NFV and SDN Market Trackers. Other areas of research coverage include subscriber data management (SDM), policy control, fixed-line TDM replacement and managed services evolution.

Hodges joined Heavy Reading from Nortel Networks, where he tracked the VoIP and application server market landscape, as well as working on the development of ANSI-41 Wireless Intelligent Network (WIN) standards. Additional industry experience was gained with Bell Canada, where he performed IN and SS7 network planning, numbering administration and definition of regulatory-based interconnection models.

Hodges is based in Ottawa and can be reached at hodges@heavyreading.com.