



The Rationale for, and Business Benefits of, App QoE

INTRODUCTION

The rationale for application quality of experience (App QoE) starts with the argument that in order to improve user experience, you have to measure it. Since user experience today means the experience of using an application, like Netflix, YouTube, or Facebook, we call that “App QoE”. It’s the goal of every service provider to ensure that the App QoE of their users is as good as it can be.

Now, we argue that to improve App QoE you have to understand it. But that’s a bit of a leap. One could assume that if you measure the key performance indicators (KPIs) of each individual network node, each interface in the network, and make sure you are in the green everywhere, then you are doing everything you can to deliver a satisfactory app QoE.

This is a dangerous assumption. Network KPIs only indicate functionality, not quality. To understand what customers are actually experiencing, you need to measure KPIs against what each application requires. So, let’s dig into that metric and see what we find.

APPLICATION REQUIREMENTS, OR WHAT APPLICATIONS “FEEL” FROM THE NETWORK

The first feeling is speed. What throughput is the network delivering to the app end to end? This is not simply an access measurement, even if it’s true that most application sessions are limited by congestion in the access, but to get true application throughput, every single interface in the telco network, from access to internet edge, needs to deliver on the throughput. And not only that, the application service provider, or cache server, or whatever is the source of the data the application needs, also needs to deliver the right throughput, and be connected in the right part of the network.

And is app throughput really all about interface speed? No, QUIC and TCP sessions all “build up” speed over a ramp up period. The length of that period depends on latency and loss as well as network capacity. If the server is far away, or any of the links have interfering loss rates, then it will take a long time to build the speed up, and the application QoE will suffer.

Some applications are inherently more dependent on latency than throughput. Voice calls, and interactive games fall into this category. They typically have a fixed throughput requirement, like a voice call that requires 64 Kbps to carry the voice codec, and any more than that is unused. But for a voice call to be a good experience for the user, latency should be low and jitter (latency variance) should be low.

Other applications are very loss sensitive. Interactive web browsing is a good example. Great throughput won’t matter if the application is constantly creating new flows to get new content and every one of those flows needs to build up throughput to download the content.



So, in the model of “let’s measure this with network node KPI’s”, to produce a proxy for app QoE, you would have to accurately monitor the following in each node of the network:

- Interface utilization, to make sure they are not full
- Loss rates, both from buffer overflows and signal issues
- Latencies from propagation delay, buffer time, etc.
- Jitter

If you miss a single measurement between the user and the server, then there is a risk that, while all your measurements are green, the user is having a horrible experience.

And in the KPI model, you have no idea what the application is asking for versus what the application is really getting. Even in the case when your telco network is perfect, the application servers and the path to those servers on the internet is unknown.

Even if you do measure all the interfaces, can you tell if the application is getting what it needs if you don’t know what the application needs? If the user is watching Netflix, or on a Zoom call, or playing a first-person shooter, his application requirements from the network vary a lot.

In a mobile network for instance, the throughput, latency and loss of subscribers on different cells varies a great deal depending on signal quality and utilization of the node. Some nodes may give the users 2 Mbps, others 10 Mbps, and low utilized nodes may get several hundred Mbps. And you cannot build a mobile network that always delivers 100 Mbps to all users all the time. That would be cost prohibitive.

But unless you know what users are doing, how can you tell if you’re giving them a good experience? Maybe 2 Mbps is plenty for what users on that cell are doing right now. Or maybe it’s horrible.

INTRODUCING SCORING

This is where scoring comes into the picture. If you’re basing your App QoE model on true application per subscriber, per location, per session, metrics like throughput, latency, loss and jitter, you can now also compare these metrics against well-known thresholds for what applications require. From this baseline you can produce a score that tells how good the app QoE is for that user at that point in time, or for all users on a particular cell at any point in time, or for all users on the network using Zoom at any point in time, or simply put: For any applications, for any user, in any location, at any point in time.

To create these thresholds, Sandvine implements a method identical to the MOS (mean opinion score) scoring used for decades to measure voice call quality. We put human test subjects in a lab, ask them to use a particular application, and score their experience. During that test we randomly set their throughput, latency, loss and jitter. We repeat this experience a very large number of times until the scoring matrix has a sufficient confidence interval. And we continually repeat this test as applications change and people’s expectations change.

BUSINESS BENEFITS OF USING APPLICATION QOE

So, what are the business benefits of using application QoE and scoring over network KPIs and other methods, like active testing?

The business benefits for an operator are many, including:

- Helping Product and Marketing structure products, marketing campaigns and messaging in a differentiated way to attract more customers,



- Helping Revenue Generation enable application usage-based billing and charging methods
- Helping Revenue Assurance go after fraud in various forms (voice, video, zero rating, roaming) and simple reconciliation.

Let's focus on three pillars of business value that are consistent from telco to telco:

- Network Operations – Tactically improving the QoE by finding more faults, more quickly, with less resources, including workflow from troubleshooting customer care calls and fault tickets.
- Network Planning – Business benefits from putting CAPEX and OPEX to better use by timing new capacity more precisely (not too soon or too much, but also not too late and too little)
- Network Engineering – Using App QoE to enable traffic management schemes like Heavy User Management or Intent-Based Congestion Management

Sandvine offers all of these benefits in custom workflows for ease and speed of implementation. Let's explore each in some detail:

Network Operations

Network Operations is all about ensuring a fully functioning and high-performing network at all times. Normally this is done with a combination of active testing (drive testing, measurement tools, etc.) and KPI monitoring from OSS tools.

However, these methods fail to understand App QoE for the reasons described above. At any point in time, there can be numerous subscribers having a poor experience with an app on the network while all the KPIs are green.

Not all of these poor experiences are practically avoidable. People will be in poor coverage areas, and application servers can be overloaded or poorly engineered, giving a poor user experience.

But the majority of App QoE issues are avoidable. Identifying them is the first step.

At Sandvine, we summarize the operation benefits of app QoE like this:

“Who”?

Which subscribers are experiencing issues? What do they have in common? A particular device type? A particular App? Is it their location or the equipment there? Are they roaming or domestic? Are they all on the same usage plan?

“When”?

App QoE is an ideal leading indicator to larger networking issues. Small buildups in round trip time (RTT), loss, throughput, etc., will indicate issues sooner and more clearly than other KPIs. With small increases in loss and RTT anywhere in your network, throughput will suffer, and utilization on your interface will drop, which may not look like an issue at all and go unnoticed until calls start coming in. Getting a head start on these issues can save you from taking those calls and keep subscribers happy.

“What”?

App QoE allows you to see issues you're otherwise blind to. A particular app, like Zoom, may be scoring poorly over a long period of time before any network issue is detected. Enabling your operations centers and customer care agents to monitor and mediate issues before they become problems can avoid angry customer calls and even truck rolls.

“Where”?

App QoE comes with inherent location context. This means that if some applications are performing poorly in a particular subset of the network -- a city, a cell, a region, whatever it may be -- your operations and care teams know where to look, where to troubleshoot, and where the problem lies.



“Why”?

Troubleshooting is hard and takes experience. But having great data helps. Taking contextual App QoE information and overlaying it on the problem is a fantastic way to reduce troubleshooting time and effort. This way you can quickly isolate if the problem is with a particular location, device type, equipment vendor, link, application server, configuration setting, service tier, etc. The “why” can also tell you for a particular subscriber, whether their issues are with their device or their home, or if it’s a network problem.

All in all, it’s about solving more problems, more quickly, with less resources. It’s about a network that runs smoother and delivers a better and more reliable subscriber experience.

Network Planning

Capacity planning today looks very different in fixed versus mobile networks. In fixed networks, capacity planners pick an interface utilization metric that they are confident in, such as 70% or 50%, and when utilization hits that threshold they schedule that node for upgrade. The upgrade process itself can take several months, and the hope is that the threshold was sufficiently low that the upgrade happens before the node hits 100%. In many cases it does not and the node has to run in a congested mode for a long time before the upgrade happens.

Mobile networks typically pick a different path. They can’t afford to upgrade to 100% utilization, so instead they look at how much bandwidth the typical user is getting in a node. When that goes below 3 Mbps or whatever the threshold is, they schedule the node for upgrade. That can take at least as long as for fixed networks, and in some cases there’s really no upgrade that can be done, so they just live with a poor QoE in that location until a solution is available, such as a radio technology upgrade.

We call these methods “volumetrics based”. It’s really looking at a single variable (utilization or subscriber throughput) and trying to pick a point in time when to upgrade.

The problem is that what is fueling the growth is the applications, who are using them, how much they are being used, and how much bandwidth each application session is using. So all of that is completely missed, and the nuance of the model is simplified down to actual utilization.

So, it’s no big surprise that most upgrade decisions are wrong. They are either too soon, and you spend CAPEX and OPEX too soon and/or too much, and the resources sit unused for way longer than they need to. Or, they are too late, and subscriber QoE is diminished for a long time before an upgrade happens.

By using App QoE, operators can understand traffic growth per location much more clearly. They can see where growth is coming from, be it more subscribers using video streaming at the same time, or the streaming sessions using more bandwidth, or duration increasing. Or is it some new app like video conferencing, or the metaverse. Wherever the growth is coming from, we break it down to subscribers using an app, and applications using bandwidth. Operators can then use that data to forecast when the utilization of this interface will hit 100%, or when each subscriber’s QoE based on bandwidth available will become poor.

Having such data enables operators to plan for network upgrades and buildouts much more accurately, saving time and money, while keeping customers consistently happy.

Beyond capex forecasts, operators can do “What-If” analyses: What if Netflix video streams use 20% more bandwidth next year? What if we increase all Silver subscribers from 25 to 40Mbps? What if we get 25% more 5G users on this cell?

Bottom line, operators can spend CAPEX and OPEX dollars much more intelligently, leading to vast savings and QoE improvements.



Network Engineering

While App QoE has been established as a way to credibly and transparently assess the QoE of the user when using applications, it also serves as a great way to monitor any actions taken on the traffic to improve QoE or reduce costs.

These measures could be as simple as config changes in the RAN or Access part of the network, or involve deploying new software such as TCP Acceleration.

Without App QoE, you can't really measure the impact of these actions in how they feel to the subscriber, which also means it's hard to accurately assess the return on investment on such actions.

But if you can clearly see that for all application categories, a particular traffic action improves QoE scores from 3.5 to 3.7, and for gaming the scores went from 2.7 to 3.9, then the business benefits become clear.

The same methodology applies for Heavy User Management. If the top 0.1% are using 10% of the network resources, and you decide that this is in violation of the T&Cs of their service, or that during congestion these Heavy Users simply should be managed in favor of the other 99.9% of the users getting a better QoE. Now, regardless of how you chose to manage the traffic of these heavy users, you want to measure the impact of their traffic being reduced. Did the QoE of the other users go up? Who? Where? For what applications?

What is it that makes heavy users heavy? Should we manage all their traffic or just their file sharing, or their video? Creating the right business policy here is important and hard to do without understanding the users' traffic properly.

For the advanced user, Sandvine offers a use case called Intent Based Congestion Management. This Use Case lets the operator set "Intent" based on how they want a particular application to score. When that intent is not met, an automation engine takes various actions to try to course correct. In the end, this is much more effective and less resource intensive than manually monitoring policy and tuning variables by hand.

In summary, better classification results in better business decisions. Sandvine classifies applications with greater accuracy than anyone else and provides QoE for these applications and guided persona workflows to efficiently analyze, optimize and monetize applications.

ABOUT SANDVINE

Sandvine's cloud-based Application and Network Intelligence portfolio helps customers deliver high quality, optimized experiences to consumers and enterprises. Customers use our solutions to analyze, optimize, and monetize application experiences using contextual machine learning-based insights and real-time actions. Market-leading classification of more than 95% of traffic across mobile and fixed networks by user, application, device, and location creates uniquely rich, real-time data that significantly enhances interactions between users and applications and drives revenues. For more information visit <http://www.sandvine.com> or follow Sandvine on Twitter @Sandvine.



USA
5800 Granite Parkway
Suite 170
Plano, TX 75024
USA

EUROPE
Svärdfiskgatan 4
432 40 Varberg,
Halland
Sweden
T. +46 340.48 38 00

CANADA
410 Albert Street,
Suite 201, Waterloo,
Ontario N2L 3V3,
Canada
T. +1 519.880.2600

ASIA
RMZ Ecoworld,
Building-1, Ground Floor,
East Wing Devarabeesanahalli,
Bellandur, Outer Ring Road,
Bangalore 560103, India
T. +91 80677.43333

Copyright ©2022 Sandvine Corporation. All rights reserved. Any unauthorized reproduction prohibited. All other trademarks are the property of their respective owners.

This documentation, including all documentation incorporated by reference herein such as documentation provided or made available on the Sandvine website, are provided or made accessible "AS IS" and "AS AVAILABLE" and without condition, endorsement, guarantee, representation, or warranty of any kind by Sandvine Corporation and its affiliated companies ("Sandvine"), and Sandvine assumes no responsibility for any typographical, technical, or other inaccuracies, errors, or omissions in this documentation. In order to protect Sandvine proprietary and confidential information and/or trade secrets, this documentation may describe some aspects of Sandvine technology in generalized terms. Sandvine reserves the right to periodically change information that is contained in this documentation; however, Sandvine makes no commitment to provide any such changes, updates, enhancements, or other additions to this documentation to you in a timely manner or at all.