Pipeline
Technology for Service Providers.
www.pipelinepub.com Volume 13, Issue 4

QoE: A new era of test and measurement to determine the sweet spot of network provisioning

By: Chris Chapman

Monumental changes in how companies test and measure their networks and network devices do happen; but they tend to happen slowly, only once every ten years or so. Today, we're at a tipping point. Companies need to move beyond legacy Quality of Service (QoS) tests and begin focusing on Quality of Experience (QoE).

What QoE uncovers is the actual experience of the people who use the networks and the application services riding on top. It calculates metrics based on historical trends and measurements against baselines that define minimum acceptable quality.

Initially, in the decade spanning 1995-2005, the focus was on packets, which provided a single metric analysis of bandwidth and packet loss. The idea back then was that if the network slowed or was dropping packets, the solution was to add more bandwidth. During that decade, network faults persisted even as network equipment advanced and performance improved. But those faults were mitigated by the fact that application performance—in relation to the network—wasn't as critical. If the network went down, there were work-arounds. Many business processes could still be done manually, with paper and pen, and voice and data networks were still siloed.

In mid-2000, the network's importance began to change. Voice over IP (VoIP) started making headway as a business tool, as did video. Because of this, the network pipes became increasingly crowded. QoS became the buzzword, but testing and analysis still focused on packets and loss. The prevailing mindset continued to be around bandwidth: the more bandwidth and less packet loss, the happier the end user. Over time, however, that correlation fell apart.

To compensate, the industry began testing networks in labs before deploying them. But as companies added services on top of their pre-tested networks, they still had problems. The reason? The services were simply outstripping the abilities of existing testing techniques. The techniques provided very little insight into the relationship between the lab-tested networks and real-world scenarios.

This all brings us to QoE and today's paradigm shift in the evolution of network test and measurement. QoE goes above and beyond the traditional metrics such as bandwidth, connections per second, or open connections, which are narrowly-focused metrics that measure specific engineering attributes of a device. For example, open connections and connections per second provide insights into table scaling and bandwidth provides information about forwarding efficiency—but neither of these metrics put it all together and directly measure how users perceive quality over time. Instead, metrics need to be analyzed as a system that includes multiple layers of the network architecture and multiple protocols used in the delivery of an application. Finally, the metric should be economically meaningful. This is how test and measurement techniques can catch failure impact, which in turn leads to lost revenue.

It is important to note pre-Internet quality expectations have set the bar for QoE—the long-established five-9s uptime that network providers have often quoted as the reliability of their networks.

Businesses today cannot afford network failures or lags. In fact, it can cost a company thousands of dollars for every hour that a customer relationship management (CRM) is down. Even if the CRM system just slows, a company can still lose money. Consumers are equally wedded to foolproof network performance, whether there's an outage or the network simply slows. Just a few years ago, when the Internet slowed or went down, consumers felt the inconvenience if they couldn't play a video game or check email. Today, however, consumers are paying money for online services like Netflix, Hulu or HBO Now, and they expect that service to work on demand.QoE is a top-of-stack metric which by definition measures the impact of all lower layer events and interactions; that is, how an application network service uses orchestrated protocols such as HTTP, SIP, and video, which in turn drive transport layer protocol such as TCP or UDP, which then load the network with bandwidth. It also helps to mitigate negative economic impact. If there's poor quality, a user will call customer service or issue a help ticket. Enough times, and an unhappy Netflix customer might cancel the service. When we talk about measuring Quality of Experience, what we are really saying is "what is my experience with the service now, and how does previous history with my experience impact my perception?"

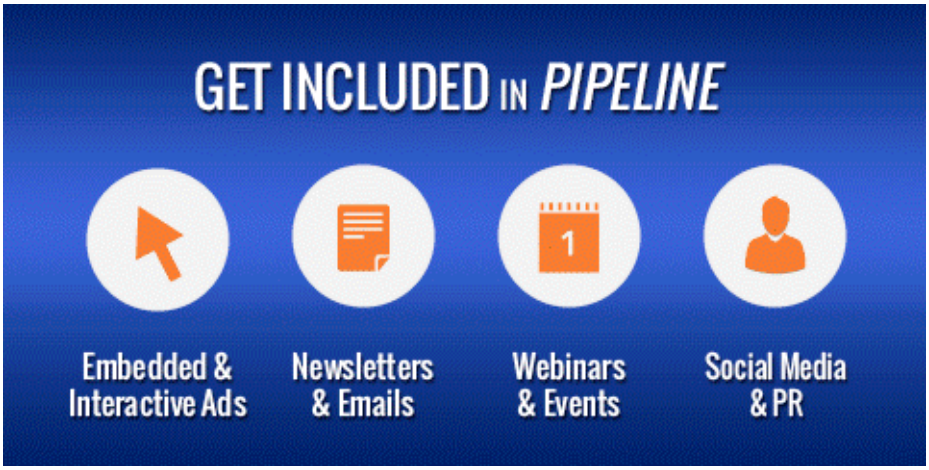

It is important to note pre-Internet quality expectations have set the bar for QoE—the long-established 99.999% uptime that network providers have often quoted as the reliability of their networks. All new services are converging to those expectations. The web is a utility that should be always-on, and if a web page loads more slowly than expected, a user's experience is diminished. Just a matter of seconds can impact perception.

As such, modern networks must be built to deliver a high QoE and robust predictability. In practical terms, this means that prior to testing, traffic flows across the network need to be well understood and for each one of the services, a minimum definition for acceptable quality must be established. For example, a modern webpage will have around 200 URLs forming the single page. In order for the instantaneous quality experience of this page to be considered high, all 200 URLs need to render in one to two seconds. Further, the variance from loading the

same page overtime needs to be very small. Lastly, any hard failure, such as a broken link or page, immediately reduces the quality of experience to unacceptable. With this definition of service QoE, we now have a measurement stick to determine the maximum number of concurrent users and user rate that we can add to a device under test.

In the lab, test and measurement systems need to generate stateful services, such as stateful CRM or stateful firewalls in order to more accurately and reliably measure provision scale and provision rate. Instead of determining that a firewall needs 100 gigabits of bandwidth, you'll be able to determine that the firewall can support 50,000 users concurrently across a device and still get a minimum QoE.

In addition tests need to be conducted to determine the sweet spot between two opposing parameters: the most expensive, foolproof service versus the cheapest service that also supports as many users as possible. This helps us to understand the economics of deployment of the device under test.

We all know that, in the fog of measurement uncertainty, we tend to throw additional resources at the device in order for users to have a good, quality experience. But the costs per user will be high; not just initially, but also longer-term. That's because we'll overpay for performance we don't need with a service contract that can represent 15 percent to 20 percent per year of the cost of the device.

That extra cost is often passed onto users, which in turn could drive the users away. Of course, if we rely too much on traditional engineering metrics like bandwidth we are likely to get an overly optimistic view of what we can provision in the network, which in turn leads to unintentional under-provisioning of a device. The net effect: user experience becomes unpredictable and

That extra cost is often passed onto users, which in turn could drive the users away.

customers leave because of their perception that service quality is poor. Losing customers is not only directly related to lost revenue, it also adds costs. The cost of recapturing customers can be quite expensive. In fact, it can cost a company about the percent of the money it could have made from a customer just to get that customer back. With QoE as the yardstick, we can better navigate the tension between under and over provisioning a network.

No less important is the testing of network services with security as a function of QoE. Network services should be tested during attacks using emulators to determine whether a device can defend against an attack while still delivering a service and transporting valid traffic. Not only is the security tested, the network performance is measured.

Test and measurement techniques have evolved as network requirements have changed and network equipment has advanced. We are now entering a new era of testing with QoE, and it's an era that's positioned well to deliver the most impactful insights ever. With QoE, businesses can better understand the quality of an application and the underlying network and devices—all from the user's perspective. In turn, they'll find that sweet spot of provisioning that is economically advantageous.