

Measuring Matters: The Plume Quality of Experience Metric

Introducing a more accurate way for communications
service providers to measure customer experience.




Table of contents

- 3 Executive summary
- 4 Introduction
- 5 The evolving connected-home environment
- 7 A problem with tradition: why Quality of Service measures aren't enough
- 9 The Plume QoE: a better approach centered on customer experience
- 14 Conclusion

Executive summary

Traditional assessments of home-internet connection quality depend on metrics that don't accurately capture customer experience. They may reliably determine whether a network is fast or slow, but they rarely consider the specific devices found in someone's home, or what those devices need to work as intended. They also usually ignore factors that exist beyond the router, like WiFi coverage quality, device location, interference, and congestion. The end result is that consumers are often left dissatisfied, and looking for alternative communication service providers.

With this in mind, Plume has developed a new approach to measuring connection and service quality called "Quality of Experience," or QoE. Like conventional quality of service measures, QoE considers network throughput, but it also looks at factors like device type, the applications a device runs, and each device's level of activity. The purpose of the QoE is to determine what a device requires for optimal performance and then compare that to what it's actually getting. CSPs can then use this information to automatically make adjustments to their service—and optimize connectivity in real-time.



Introduction

In any competitive industry, long-term business success depends on a company's ability to meet and exceed customer expectations. For customers of Communications Service Providers (CSPs), those expectations almost always center on their experience, or their perceptions about the quality of their digital connections.

Most customers simply expect the WiFi in their home to be reliably fast and trouble-free. They're typically happy as long as their WiFi matches the performance requirements of the devices they own.

The challenge for CSPs is that standard approaches to assessing connection quality rely on metrics that don't reflect the diversity of devices and applications their customers use. Measurements of signal strength and retry counts, for example, aren't accurate indicators of quality because they rely on assumptions that may not be correct depending on the network and its operating environment. While they can offer information on *quality of service*—whether a network is fast or slow—they can't determine what a specific customer's devices need to achieve optimal performance.

To fill this gap in the CSP business model, Plume has developed a new approach to measuring connection quality that captures the customer's actual experience in their home. We call this measurement "QoE," or *Quality of Experience*.

What follows is a closer look at how we calculate the QoE metric—and insight into the ways CSPs can use it to provide exactly the level of service their customers demand.



The evolving connected-home environment

To understand the challenges CSPs face in monitoring and addressing WiFi connection quality, it's helpful to first paint a clear picture of the evolving digital environment in the typical modern home.

As recently as the early 2000s, most homes with internet service only had connections through desktops or laptop computers. Since then, however, with the growth of the Internet of Things (IoT), the production and distribution of connected consumer electronics have barely kept up with demand.

Researchers have estimated that, across all industries, the IoT expanded from a total of 7 billion devices in 2018 to more than 30 billion devices just two years later. By 2025, according to analysts, the worldwide IoT will include at least 75 billion connected devices.¹

The IoT explosion has been similarly impressive at the consumer level over the past few years. Worldwide, the number of connected devices managed by consumers rose from an estimated 5.24 billion in 2017 to nearly 13 billion in 2020.² By 2022, analysts predict, the market for home-based IoT devices—including “smart” connected devices controlling things like lighting and entertainment systems—will top \$53 billion globally.³

Because the majority of connected consumer devices access the internet via WiFi, their proliferation in people's homes has led to dramatic increases in network load. While most digital devices on their own tend to require relatively little bandwidth, the data they consume collectively can push a home network to its limit.

¹ Gilad David Maayan, “The IoT Rundown for 2020: Stats, Risks, and Solutions,” Security Today, January 13, 2020


² “The Internet of Things (IoT)* units installed base by category from 2014 to 2020,” Statista Research Department, January 14, 2021

³ “Forecast market size of the global smart home market from 2016 to 2022,” Statista Research Department, January 22, 2021



Bandwidth demand has become particularly problematic in the wake of the pandemic. With more people working and learning from home, and turning to home-based entertainment because of the need for social distancing, CSPs have seen huge shifts in network usage patterns as devices connect to the internet at all times of day. Plume's own analysis of 14 major US metropolitan areas found the number of households that were online during the workday increased by 105 percent between January and April 2020, just a month after COVID-19-related lockdowns began.⁴ Similar increases occurred in Canada and Europe, and it's widely expected the work-from-home trend will continue post-pandemic as organizations recognize the advantages of operating virtually.⁵

ACTIVE HOUSEHOLDS ONLINE GROWTH



⁴ Joan Engebretson, "Report: Work-Day Home Wi-Fi Usage Doubles with COVID-19 but Hits Plateau," Telecompetitor, April 8, 2020
⁵ Bill McFarland, "Connectivity and Network Performance During a Crisis," Plume technical paper, May 2021
⁶ "American Customer Satisfaction Index Telecommunications Report 2019-2020," American Customer Satisfaction Index, LLC

WHY CUSTOMERS HAVE TROUBLE CONNECTING IN THEIR HOMES

- 1. Poor coverage:** Certain areas of the home may lack sufficient coverage because of inadequate signal strength.
- 2. Location issues:** In homes with multi-access point (AP) systems, the client may not connect to the optimal AP as it is moved to different locations.
- 3. Interference:** Overlapping transmissions from neighbors can lead to increased interference.
- 4. Congestion:** Competing devices within the home can lead to self-interference.

The big question for CSPs is what this means for their ability to provide services that deliver on customer expectations. One recent report from the American Customer Satisfaction Index, covering a range of telecommunication industries, noted that the services CSPs offer have become "more critical than ever before."⁶ The report found that while customer-satisfaction rates have risen slightly since the pandemic began, these gains could be short-lived if CSPs don't find solutions to connection issues that typically arise in the modern connected home.



A problem with tradition: why Quality of Service measures aren't enough

CSPs have traditionally used so-called “Quality of Service” (QoS) measures to assess performance of the connections they provide to customer devices. These measures were perfectly sufficient when raw broadband data transmission was the main factor determining connection quality, but that’s no longer the case in the modern home.


Simply put, QoS is a reasonable indicator of whether a network is running fast or slow, but it doesn’t offer adequate information on what’s happening beyond the connection to the customer’s router. Back when people accessed broadband internet primarily through a physical cable—and were mostly connecting computers—QoS was all you needed. Today, however, most connections are wireless and the range of device types being connected is enormous, which means typical QoS metrics often come up short.

DEFINITION: QUALITY OF SERVICE (QOS)

The measurable end-to-end performance properties of a network service, which can be guaranteed in advance by a Service Level Agreement between a user and a service provider, so as to satisfy specific customer requirements. Note: These properties may include throughput (bandwidth), transit delay (latency), error rates, priority, security, packet loss, packet jitter, etc.

Source: National Institute of Standards and Technology⁷





To assess network operation quality, most WiFi QoS metrics consider factors like signal strength, data rate, and congestion and interference. But RSSI (received signal strength indicator) measurements, PHY (physical layer) rates, and WiFi retry counts all rely on certain assumptions that usually don't apply to today's complex network-operating environments. In the past, CSPs could safely guess that the typical home would have low external interference and minimal channel contention, and that backhaul speeds wouldn't matter. But they can't do that now—and most consumers aren't interested in network performance, anyway.

In fact, consumers typically only care about the performance of their WiFi-connected devices and the applications they run—in other words, their experience. Using QoS alone to determine customer satisfaction means that CSPs may misdiagnose a customer's connection problems in their home.

A customer with a “smart” WiFi thermostat who receives reliable service at 1 megabyte per second (Mb/s) to that thermostat will probably be satisfied with that connection. But a customer with a 4K resolution set-top box won't be happy even with 20 Mb/s service to the set top box. With traditional QoS metrics, IoT devices with lower networking requirements—like a WiFi thermostat—might be identified as needing attention, while streaming devices that require higher data-transfer rates may not be noticed at all.

QOS LIMITATIONS


Traditional Quality of Service metrics often lead CSPs to misidentify connection problem areas.

- **43%** of video-streaming devices identified as problematic using QoS are actually functioning perfectly fine, while **66%** of video-streaming devices that do have real connection issues aren't identified at all with these metrics
- Only **0.3%** of IoT devices identified by QoS as operating at a suboptimal level are truly unhappy, while **77%** of IoT devices with legitimate connection problems are missed using traditional QoS.

Source: Plume Cloud data⁸

The bottom line for CSPs is that determining customer satisfaction with traditional QoS metrics comes with substantial risk. If a customer consistently has to deal with poor video quality while they're trying to stream shows or play their favorite games, there's a good chance they'll think their service isn't good enough and take their business to another provider.

⁸ Bill McFarland, “Yesterday's metrics fall short for today's smart homes,” Plume, January 30, 2020



The Plume QoE: a better approach centered on customer experience

Knowing that limiting customer churn is critical to any CSP's survival, Plume has developed a way to assess connection quality that goes well beyond the traditional measures provided by QoS. The Plume Quality of Experience (QoE) metric considers not only how a network is running, but also the various devices and applications being used and what they require in order to be “happy.”

Unlike QoS, which typically considers just the last wireless hop to a client, QoE looks at every hop across the network. QoE can tell a CSP whether each and every connected device in a home is getting the throughput it needs. It’s a clear indicator for whether each device is connected in a way that could make its users frustrated, or if that connection is exactly what it should be—and therefore able to satisfy that customer.

QOS VS QOE

Like QoS, QoE factors throughput, but it does so using more information. It also considers device “need,” and whether a device is active on the network. The result is a much more accurate understanding of the experience users have on a particular device.

	Quality of Experience (QoE)	Quality of Service (QoS)
Portion of network	All hops to gateway	Last hop to device only
Throughput	Physical layer (PHY) rate, packet retry rate (PRR), interference, congestion (the true net throughput)	PHY rate, interference
“Need”	Device type, recent usage, broadband speed	Not factored
Activity	Active vs. inactive devices marked and scored differently	Not factored

Consider, for example, a set-top box that isn't providing the level of resolution the customer expects. With traditional QoS, the CSP would look at the last hop to the client (without considering what the device is or what its throughput needs are) to determine whether it's connected to that access point well. If it is connected as would be expected, they would then incorrectly conclude the device seems to be "happy," or functioning as it should.

With QoE, on the other hand, the CSP would know what the device is and have a clear understanding of its unique throughput needs, and they would examine every wireless hop from that device to the point where the connection enters the home. In addition to the physical layer data rate (how fast the data is traveling), QoE also factors packet retry rate (an indication of packet errors), interference from neighbors, and congestion (interference from within the home). And finally, it factors activity to understand the relationship between network problems and the devices that are using the network when those problems are observed.

With all of this information at their fingertips, a CSP might then determine that one of the various hops was bad, or that different hops on the same frequency channel were interfering with each other, for example. Whatever the case, they would find the cause of the customer's issues and could then take steps to resolve them.



Measuring QoE: determining device "happiness"

Because the Plume QoE is fundamentally about deciding whether a device is "happy" or not, the measurements used to calculate QoE depend on device type and the applications they can potentially run. The first step in the calculation involves determining the *real-time estimated throughput* for the device. This value is then compared to the device's *desired throughput* to arrive at its QoE metric.

Measuring real-time estimated throughput

Real-time estimated throughput (potential throughput) is the estimated WiFi throughput capacity that can be achieved over any number of hops from a device to the WiFi gateway. This value varies from device to device, depends on the device's location within the home, and is calculated in real-time using measurements of the following network attributes:

- Radio-frequency (RF) hops from end device to the WiFi gateway.
- Airtime available on each channel used for each RF hop/link.
- Physical transmission rates used for each RF hop/link.
- Packet retry rate for each RF hop/link.
- Contention resulting from channel reuse within the WiFi network.
- External neighbor interference.
- Current data usage on the end device.

Real-time estimated throughput can also be thought of as the end-to-end bandwidth available for a device or AP. In mathematical terms, it's the sum of the calculated additional throughput that a device could receive (which factors number of hops and their frequency channels, free airtime available to transmit, and rate at which they can operate) and the current traffic usage at the device.

Real-time estimated throughput = Additional throughput to device + Estimated current traffic usage at device

Measuring desired throughput

A device’s desired throughput depends in large part on its class and category, but also on whether it is active or idle. By considering two other factors that influence desired throughput—current traffic usage and ISP connection speed—it’s possible to calculate “happy throughput” for a device, or the bandwidth it requires to run its applications at any given time.

Happy throughput = min (historically observed peak traffic usage, device throughput need, ISP throughput)

THROUGHPUT NEEDS, BY DEVICE

Device Class	Device Category	Active Need	Idle Need
IoT devices	Thermostat, WiFi speaker, watch, etc.	3 Mbps	1 Mbps
Cameras	Doorbell, VOIP gateway, etc.	15 Mbps	5 Mbps
Mid-range devices	Gaming system, set-top box, home theater, etc.	22.5 Mbps	7.5 Mbps
High-end devices	Laptop, desktop computer, smartphone, etc.	45 Mbps	15 Mbps
2×2 dual-band AP	Pod, smart mesh router, etc.	150 Mbps	50 Mbps
4×4 dual-band AP	VF extender, PowerPod, etc.	450 Mbps	150 Mbps
4×4 tri-band AP	SuperPod, etc.	600 Mbps	200 Mbps

When a device is running a high-capacity application, the network load is increased. A device using a low-capacity application, on the other hand, may not require the ISP throughput it has available. A laptop computer, for example, could potentially require 45 Mbps or more when it’s active. But if that computer is being used solely for reading email, its “happy throughput” value will be much less.

Calculating the QoE metric

Once real-time, estimated throughput and happy throughput are determined, it’s possible to compare these two numbers to see where a device falls on a “Happiness Index”. Defined as the overall state of the device based on its usage and the capacity it has available, this value is then mapped to a QoE score.

Happiness index value = real-time estimated throughput/happy throughput

Happiness Index Value	QoE Score
0 to 0.75	1
0.75 to 1.25	2
1.25 to 2	3
High-end devices	4
2×2 dual-band AP	5

The higher a device’s QoE score, the better off it is (at that moment) on the network to which it’s connected. For example, consider the following scenario:

A 2×2 802.11ac laptop computer is connected to a 4×4 802.11ac dual-band WiFi extender, which in turn is connected to a gateway over WiFi link. Each link is operating on channel 44, and the ISP broadband speed is 250 Mbps. The customer is streaming a movie, which consumes 5 Mbps of traffic, on average.

Real-time estimated throughput = min (0.8, 0.6) × 1/(1/(500×0.75) + 1/(800×0.9)) + 5 = 153 Mbps

Happy throughput = min (5 Mbps, 45 Mbps, 250 Mbps) = 5

Happiness index value = 153 Mbps/5 Mbps = 30


QoE score = 5

In this case, the device is in good shape—it has everything it needs to be happy. If the network was highly congested, on the other hand, and the device was further away from the WiFi extender, its QoE would be much lower (and its status, poor).

From data to action: putting the QoE score to work

Calculating QoE scores is one thing, but how can a CSP then take those scores and turn them into information it can use?


The answer can be found in Plume's Haystack software suite, which gives providers access to current and historical QoE scores as well as the metrics used to calculate them. Navigating their Haystack dashboard, a CSP analyst can see how long customers spend in good or poor QoE conditions and identify the reasons for that performance.

An abstract graphic on the left side of the page, consisting of a series of concentric, overlapping blue wireframe loops that resemble a stylized signal or network pattern.

They may see that a low score on a particular device resulted from an insufficient physical layer data rate, or that the score was due to interference from a nearby home. They could also use analytics and other tools to look at QoE values across a specific group of households, or across all of the households in their network. And finally, they could leverage that data to make changes to their service or recommendations to their customers—often before those customers even knew they had a problem.

They may find, for example, that a certain customer would benefit from having a repeater in their home. Or they might see that another customer's poor WiFi signal—and resulting inability to stream videos online—was due to a baby monitor located near their router. A CSP with that knowledge in hand could pass its observations on to the customer so they could address the problem on their own. Or in cases where a network reconfiguration might be handled remotely, the CSP could make the changes themselves.

Because QoE provides an accurate view of what's actually happening with the networks in people's homes, it allows CSPs to make proactive decisions and stay a step ahead of any potential problems. It's not just a snapshot of broadband data transmission—the metric you get with QoS. It's a clear picture of each customer's entire networking environment and their experience with their connected devices.



Conclusion

As the IoT continues to expand and the number of connected home devices expands with it, shifting from QoS to QoE is a way for CSPs to keep their customers happy. Plume's own data, in fact, shows that CSPs who have already adopted QoE metrics were able to maintain their customers' Quality of Experience even as load and congestion hit record levels early on in the pandemic.⁹

It won't be long before the IoT grows to the point where such network stress is the norm rather than the exception. When that day comes, the CSPs that excel will be the ones who understand their customers and their unique needs. With QoE, these CSPs will be able to focus their resources on the right places—and deliver the right solutions where they're needed the most.

To learn more about Plume, visit our [website](#) or [contact us](#) today.



⁹ Bill McFarland, "Managing the Home Network's Connectivity Conundrum," Multichannel News, June 1, 2020

