

Application QoE

Monitoring end user experience of enterprise applications

Introduction

The performance of Enterprise WiFi networks has traditionally been measured in terms of metrics such as radio link quality (e.g. RSSI, SNR, retry rate), packet latency and user throughput. However, with more and more enterprises leveraging the IT infrastructure for business critical tasks and increasing reliance on cloud applications, what really matters is the end user experience. For instance, while using a conferencing application such as Zoom, it's the end user's perception of call quality that is the true indicator of network performance. Similarly, the responsiveness of the Google Docs application while collaboratively editing a document online is the true measure of network performance. This calls for measuring the Quality of Experience (QoE) of enterprise applications to determine if the underlying network is performing as per expectations or not.

Quality of Experience

According to ETSI TR 102 643¹:

“A measure of user performance based on both objective and subjective psychological measures of using an ICT service or product.”

It further clarifies that *“the appropriate psychological measures will be dependent on the communication context. Objective psychological measures do not rely on the opinion of the user (e.g. task completion time measured in seconds, task accuracy measured in number of errors). Subjective psychological measures are based on the opinion of the user (e.g. perceived quality of medium, satisfaction with a service).”*

In the enterprise scenario, where users access a multitude of applications on a daily basis, it is near impossible for a network administrator to obtain subjective measures of user satisfaction for each application. Hence, objective measures need to be determined so that a non-intrusive estimation of QoE can be performed.

As per ITU-TG.1011²:

In principle, assessment of quality of experience (QoE) must be performed using subjective tests with metrics, such as the mean opinion score (MOS). However, it is also possible and sometimes more convenient to estimate QoE based on objective testing and associated quality estimation models. Through different quality estimation models, it is possible to measure or calculate the objective parameters affecting QoE, in order to evaluate QoE. Subjective testing needs more resources and effort, because it requires human subjects. On the other hand, objective measurement and automatic calculation using appropriate quality estimation models is generally much faster and cheaper, but the accuracy of the final evaluation depends on the accuracy of these models.

In light of the above, there are two challenges:

1. Identify an objective estimation model to measure QoE using a minimal set of parameters
2. Measure the QoE using the model in a non-intrusive manner without involving the end user herself

Arista’s approach to application QoE measurement is rooted in the ‘Cognitive WiFi’ vision, where the system learns on its own. This is exemplified by the ‘VoIP QoE’ estimation method which is used to determine the QoE of popular audio-video communication (AVC) applications such as Google Hangout/Meet, Microsoft Teams/Skype, Zoom etc. As shown in Figure 1, during an ongoing VoIP session, the WiFi AP extracts the model parameters from the traffic flow and this data is fed to the QoE engine which has been trained for this purpose.




Figure 1.

¹ ETSI TR 102 643 V1.0.2 (2010-01): Human Factors (HF);Quality of Experience (QoE) requirements for real-time communication services, Jan 2010

² ITU-T SERIES G: TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS - Multimedia Quality of Service and performance –Generic and user-related aspects: Reference guide to quality of experience assessment methodologies, June 2015

Measuring QoE of TCP Applications

AVC applications typically use UDP as the transport protocol and their QoE can be characterized in terms of measurable parameters such as bitrate, latency, jitter. However, the vast majority of enterprise applications run on TCP. Unlike the high degree of 2-way interactivity in the AVC applications, non-AVC applications such as Web browsing and video streaming are more transactional in nature. Hence, a different QoE estimation approach is required since the nature and context of communication is very different.

The QoE model developed by Arista for TCP applications is based on the following:

1. Analysis of TCP Sequence and ACK numbers present in TCP header to infer packet losses in the network.
2. Estimation of the time lost in a TCP connection due to packet drops in the network.
3. For a given TCP flow, the model is used to extract the following information:
4. Active time - Duration when actual packet exchange takes place between the two endpoints
5. Recovery Time per direction - The amount of time TCP was in a wait state due to potential packet losses in the network.

The Figure 2 below shows an example of this:

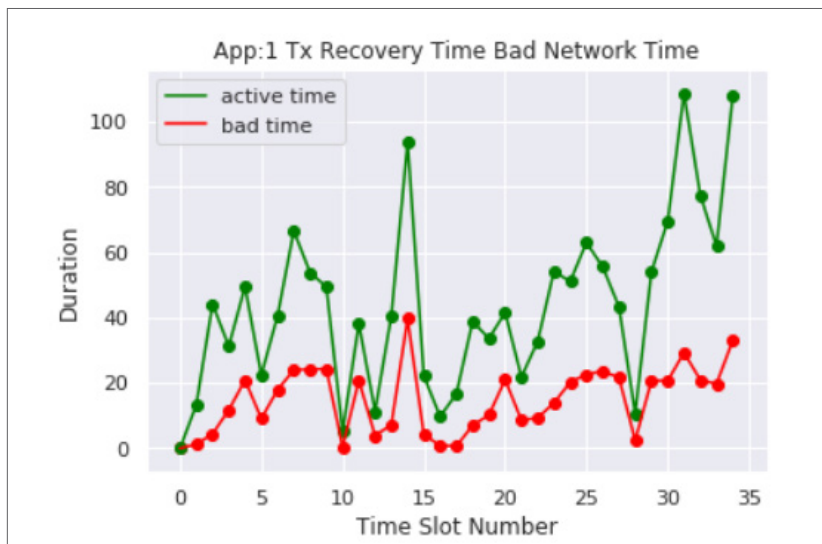


Figure 2.

The x-axis represents Time Slots, which are the basic units of time over which Active Time and Recovery Time values are computed. In a given slot, Recovery Time ('bad time') represents the amount of time for which the TCP stream was stalled. If we consider slot 10 in the figure above, Recovery Time was 0 seconds while Active time was 5 seconds. Hence, the QoE is assumed to be excellent. However, in slot 5, Active Time was 20 seconds while Recovery Time was 10 seconds. In other words, for 50% of the time, the TCP stream was stalled. Hence, the QoE can be said to be poor.




Figure 3.

We see another sample chart in Figure 3, where the bad time value is always zero, indicating excellent QoE.

For any real web application, there are typically multiple parallel TCP sessions per application session. The method described above considers all the sub-sessions while computing Active and Recovery times, as shown in Figure 4.

In Figure 4, the application session consists of 5 TCP flows, with different packet arrival processes.




Figure 4.

Implementation

The QoE estimation method described above requires each AP to monitor the TCP flows, extract the relevant information and send it to the cloud, where CV-CUE processes the data and publishes it on the Application Health Dashboard. Each card pertains to a specific application and the red portion of horizontal bar shows what percentage of time the QoE was deemed to be poor for that application. For example, we can see that Flipkart and Teams users had a very good QoE while users who accessed Facebook and Google services saw a poor QoE.

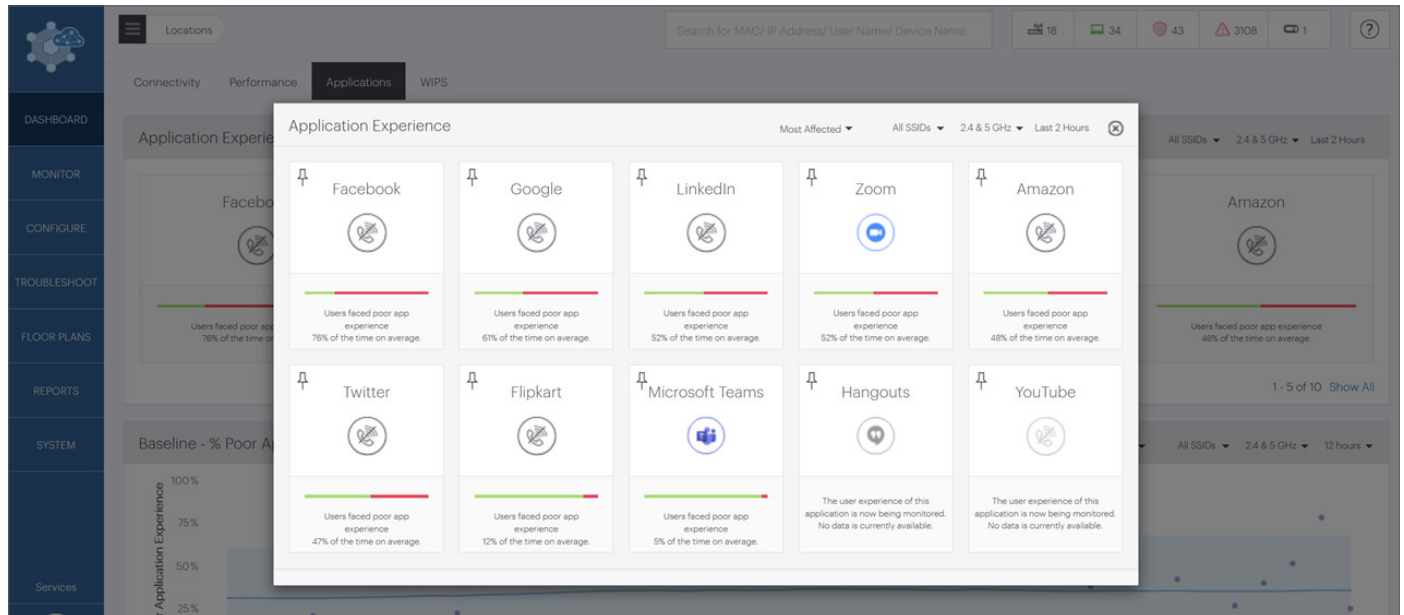


Figure 5.

In addition to the summary view shown in Figure 5, which pertains to the last 2 hours only, the dashboard also provides a historical view of QoE, stretching upto a month.



Figure 6.

Here in Figure 6, the baseline represents the average behavior and each dot reflects the QoE as measured during a given 15-min interval. This data can be viewed as an aggregate over all selected applications or for a specific application, thus giving a more granular view of the QoE. The red dots are intervals when the QoE was much higher than the baseline, indicating a network anomaly. A rapidly varying baseline itself is an anomaly as it indicates a possible instability in the network. The chart on bottom-right shows the number of clients in each QoE bucket, with QoE worsening as we move from left to right. This gives an overall sense of the application QoE at a given location.

Each customer (or CV-CUE instance) can be configured to track up to 25 applications. This includes both AVC and non-AVC applications. A network administrator can select this set from an extensive list comprising hundreds of applications.

The feature shown in Figure 7 is supported by Arista’s Wave 2 and WiFi 6 Access Points.




Figure 7.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390
Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062

