NetForecast®

2022 Internet Latency Benchmark Report

Peter Sevcik, Alan Jones, Andy Lacy, and Rebecca Wetzel

NetForecast®

NetForecast Report NFR5149
September 2023

EXECUTIVE SUMMARY

A truly fast internet experience requires low latency. Internet service providers (ISPs) continue to ramp up bandwidth, calling more bandwidth higher "speed"— but above 30Mbps the user experience levels out, and latency becomes a more noticeable performance impediment than bandwidth [1]. The rate at which data flows into and from your home over your local connection is governed by bandwidth, whereas the speed at which data flows between any server on the internet and your home is governed by latency. Because latency matters to the user experience and because internet latency varies significantly by ISP and city, NetForecast set out to measure continuously and report periodically on internet latency within and between major US cities starting in 2020. This report details the latency results for all of 2022.

NetForecast collects end-to-end test data to produce a latency score using Application Performance Index (Apdex) methodology [2] that is sensitive to deviations from normal. The methodology is performance-comparative, hence is a benchmark approach. Whereas the widespread practice of averaging results over many samples conceals instances in which performance deviations are significant, NetForecast's methodology flags critical deviations from normal baseline performance, allowing realistic performance assessment and meaningful comparisons across cities and service providers over time.

This report documents internet latency performance from January through December 2022. Figure 1 shows that Verizon delivered the best overall ISP performance among the six service providers measured, and Washington experienced the best overall performance among the 11 cities measured.

The data also show that the internet does not deliver the same user experience everywhere across the US, even in these "high tech cities"—and **latency in the last mile is the major contributor to user frustration**.

Q1 City	Q1 Score	Q2 City	Q2 Score	Q3 City	Q3 Score	Q4 City	Q4 Score
Washington	0.94	Washington	0.96	Washington	0.96	Washington	0.96
Los Angeles	0.92	New York	0.94	Miami	0.94	New York	0.94
New York	0.92	Los Angeles	0.92	New York	0.94	Miami	0.93
Seattle	0.90	Phoenix	0.92	Los Angeles	0.93	Seattle	0.92
Phoenix	0.89	Seattle	0.92	Phoenix	0.91	Los Angeles	0.90
Dallas	0.88	Miami	0.91	Seattle	0.91	San Jose	0.89
Atlanta	0.87	Chicago	0.89	Chicago	0.88	Phoenix	0.89
Chicago	0.87	Dallas	0.88	San Jose	0.88	Dallas	0.88
Miami	0.86	Atlanta	0.86	Dallas	0.86	Atlanta	0.87
San Jose	0.85	San Jose	0.84	Atlanta	0.85	Chicago	0.86
Denver	0.82	Denver	0.81	Denver	0.77	Denver	0.77

Q1 ISP	Q1 Score	Q2 ISP	Q2 Score	Q3 ISP	Q3 Score	Q4 ISP	Q4 Score
Verizon	0.99	Verizon	0.99	Verizon	0.99	Verizon	1.00
Comcast	0.91	Cox	0.93	Comcast	0.92	Comcast	0.91
Cox	0.91	Comcast	0.92	Cox	0.92	Cox	0.90
AT&T	0.90	AT&T	0.91	AT&T	0.92	AT&T	0.89
CenturyLink	0.82	Charter	0.81	Charter	0.82	Charter	0.85
Charter	0.80	CenturyLink	0.81	CenturyLink	0.78	CenturyLink	0.80

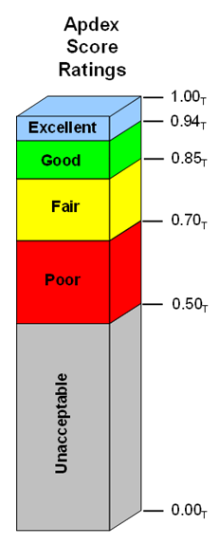

Figure 1 – Performance Benchmark Rankings

NetForecast measurements are made using the RIPE atlas network. As a RIPE Atlas Ambassador, we also operate Atlas anchor servers in the US.

NetForecast's QMap™ Internet Latency Benchmark Service builds on over two decades of internet performance testing and analysis. This report is based on tests run every hour of every day to nearby and distant servers, generating more than 40 million latency tests per month. Other entities exclusively test to nearby servers, and they do so either only when requested by users who may be experiencing perceived poor performance, or only during select periods during a few weeks each year. NetForecast uses the Apdex methodology to find and highlight times when performance is impaired and ***to give voice to the user's abnormal experience***. We determine baseline conditions, and then evaluate how current measurements deviate from that baseline. The consumer subscriber lines from which we measure are located within 11 major metropolitan areas, encompassing over 25 percent of US households.

Improvements Over Our Past Reports

This report is the continuation of our previous reports beginning with the 2020 report [3]. In 2022 we significantly expanded the measurement footprint and increased the number of testing probe-server pairs. We learned from previous reports to add enhanced analytics to remove questionable measurements. Many readers wanted to better understand the meaning of a good vs poor Apdex score. To that end, we added a new measure of poor performance magnitude in both RTT values in milliseconds and multiples of poor vs good latency RTTs. And as with all big data studies, one should not just look at top-level averages. We show interesting insights when looking at the data in detail.

FRAMING LATENCY REPORTING OBJECTIVE

Network latency is a fundamental performance property of the internet. The internet is a datagram packet transport system comprised of many component subnetworks. All packets must traverse circuits, switches, and routers. Furthermore, packets associated with a use case often traverse multiple networks operated by different service providers. Given the highly variable paths taken across this complex system, packets traveling between any two physical locations will have variable end-to-end delays known as latency.

An overwhelming majority of internet use cases rely on low or consistent latency; some require both. Consistency is more important because applications and protocols have been designed to align their behavior to the prevalent latency observed. Very low latency with very little variation is an emerging requirement for advanced use cases such as virtual reality (VR) and augmented reality (AR). Data in this study can be viewed as foundational information for VR and AR but they are out of this report's scope.

This study reports on latency measured across the broad continental United States internet. We are reporting on the quality of service (QoS), not quality of experience (QoE). Good QoS is essential for good QoE. But given a long history of advances in applications' ability to operate over a wide range of latency scenarios, we are documenting QoS as the fundamental pillar of QoE.

Latency is the elapsed time between when a data packet leaves a user's device, arrives at a destination server, and a response packet returns from that server to the user's device. This elapsed time is referred to as round trip time (RTT), and it is measured in milliseconds (msec). Latency is a complex topic which was recently studied by a Broadband Internet Technical Advisory Group (BITAG), Technical Working Group [4].

Not only is application responsiveness a significant issue for human users of interactive applications, but it is also increasingly important for autonomous systems, smart homes, logistics systems, health monitors, etc.

ISPs often advertise bandwidth as "speed," thus promoting a narrative that subscribers should buy higher bandwidth services to improve application responsiveness. This was true when bandwidth delivered was measured in single digits, however above 30 Mbps the benefits of higher bandwidth are marginal at best. Given recent technology and application initiatives, the performance focus is shifting to improving latency.

ISPs are deploying new low-latency access technologies like low-latency-DOCSIS (LLD) for cable and 5G for fixed wireless networks, which reduce latency across the last mile (modem-to-cable headend, or mobile-to-tower). These new "low latency" solutions were not widely deployed to be included in this 2022 study.

Applications and protocols have sophisticated mechanisms that adapt to high latency with practical limits across the US. However, if latency increases significantly from the acclimated norm, then applications can become erratic, freeze, or even stop operating. These "poor performance" incidents are clearly visible to a human user or can impact the reliability of automated systems operating on the internet. The actual impact on people and systems is not within the scope of this study.

MEASUREMENT AND ANALYSIS METHODOLOGY

NetForecast's Internet Latency Benchmark uses a rigorous methodology to measure and report end-to-end QoS representing typical use cases. This report, which is a part of an ongoing series, provides detailed insights into performance variation across locations and service providers over time. The data, which is summarized by major metropolitan area and by ISP, can be used to assess how well a metropolitan area or an ISP is performing relative to others.

Near vs Far Latency

When accessing content, users' devices are generally directed to a server within a nearby metro area, or to a server that is significantly distant. This structure is essentially a binary user-to-content assignment, where the content is either near or far from the user. In the case of the US, we define the binary assignment as:

- **Near** path is to a location within the user's physical metro area (e.g., many content providers pay to store information on content distribution networks, where it is served locally).
- **Far** path is to the origin content location across a substantial portion of the US internet.

Users are unaware of which path their content is traversing, and the path may change during a single session. A typical use case operates over both near and far paths simultaneously. Content providers that can afford to place their content in many metro locations attempt to deliver from local servers to optimize application responsiveness.

What We Test

As Figure 2 shows, NetForecast conducts separate tests within and between cities to measure the performance a user experiences accessing content located within the local metropolitan area (near) and beyond (far). The inner blue circle showing near testing covers the ISP's last mile, local peering, metro area networks and metropolitan data center access. Near tests (solid black arrows) simulate consumers accessing content hosted in edge service provider data centers or delivered via Content Delivery Networks (CDNs). The outer blue ring showing far testing covers the local (near) ISP's last mile, distant peering, middle-mile (transit) ISPs, transit-to-transit peering, and distant data center access. Far tests (dashed red arrows) simulate users accessing content hosted at origination data centers and from sources that do not use edge or CDN services.

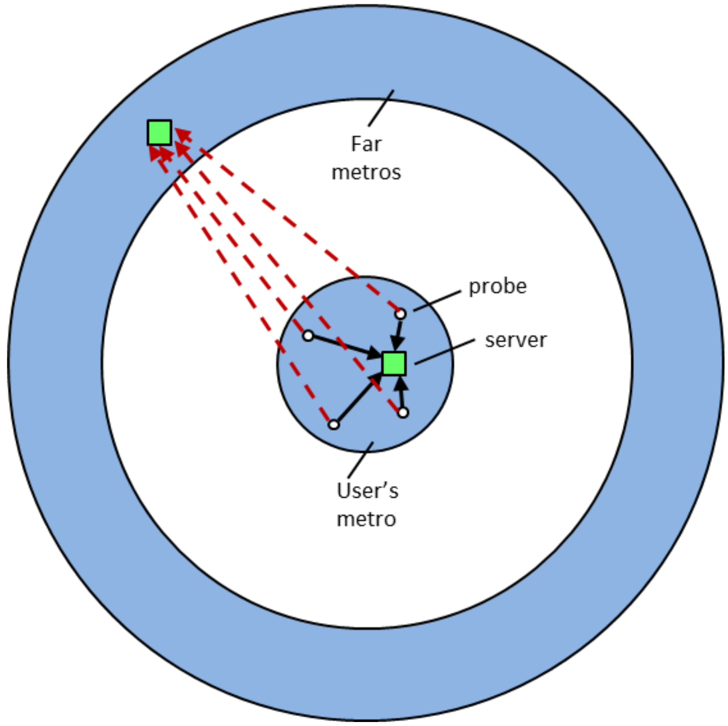

Figure 2 – NetForecast Near and Far Test Architecture

Comprehensive USA Coverage

Figure 3 shows the transit paths for NetForecast's near and far tests. Near tests are conducted within the metropolitan areas shown by the blue circles, and far tests follow the transit paths shown by red lines. The metropolitan areas are defined by a circle with a 450-mile straight line radius from the city center.

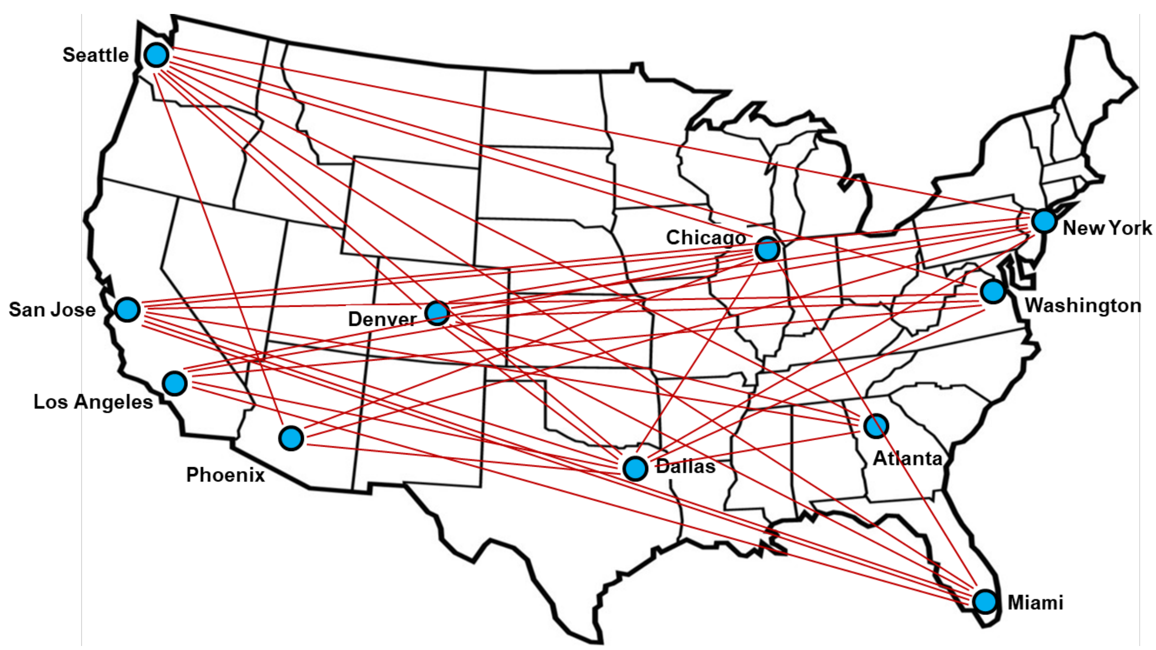

Figure 3 – NetForecast Near and Far Test Paths

The Test Plan

NetForecast actively measures latency among 11 major US metropolitan areas. Ten of the cities are used by the FCC Measuring Broadband America program which is focused on bandwidth tests. We added Phoenix in 2022 which brings the total to 11. Tests are performed by RIPE Atlas probes which consumers have installed on their home networks. The Atlas system is an open measurement platform which users can access by installing a small probe and institutions can access to study test results. The volunteers who installed the probes are not recruited by NetForecast.

Identifying Probes

At the start of 2022, we studied the Atlas database of then-current data to select probes for this study by using the following guidelines:

- Select a center point of each metro area for the 11 cities.
- Determine which ISPs are providing service to the probes in the metro.
- Identify probes accessing the internet from each ISP.
- For each ISP find probes within an expanding radial distance from the metro center until:
    - at least 4 but not more than 10 probes are identified and
    - the limit of 450 miles from the metro center has not been exceeded.
- Place the selected probes on the city's local probes list.

Identifying Target Servers

Servers used for the tests are identified within each metro city. These are Atlas Anchors or other high-availability and high-performance servers. In all cases, they are hosted at major datacenters near a local internet exchange point. There are at least 2 target servers within each metro area. Place the selected servers on the city's local servers list.

Defining Near and Far Tests

Create the following probe to target server pairings for each city.

- Near tests: All probes in a metro test to every server in the same metro
- Far tests: All probes in a metro test to two servers in each of four distant cities

The Test Matrix

Probe-target pairs are carefully selected to generate a comprehensive US internet performance view as shown in the table below. The matrix shows the 55 total unique city test pairs (11 near and 44 far).

City Testing Matrix				
Near City	**Far City A**	**Far City B**	**Far City C**	**Far City D**
New York	Seattle	San Jose	Denver	Dallas
Washington	San Jose	Seattle	Denver	Dallas
Chicago	Denver	Dallas	San Jose	Miami
Dallas	Chicago	Denver	Seattle	New York
San Jose	Miami	New York	Chicago	Dallas
Los Angeles	Washington	New York	Dallas	Chicago
Seattle	Washington	Miami	Denver	Chicago
Denver	Dallas	Chicago	Miami	Seattle
Atlanta	San Jose	Seattle	Dallas	Denver
Miami	Los Angeles	Seattle	Chicago	Denver
Phoenix	New York	Seattle	Dallas	Chicago

The tests associated with each of the 11 key cities are shown as rows in the matrix. For example, New York occupies the top row with the following test arrangements:

Near: New York probes to New York servers
Far A: New York probes to Seattle servers
Far B: New York probes to San Jose servers
Far C: New York probes to Denver servers
Far D: New York probes to Dallas servers

Programing Tests

NetForecast tests using probes in consumer homes directly connected to major broadband service providers' routers (no Wi-Fi is involved). We perform standard ICMP [5] "ping" tests between the probes and reference servers (targets) located near the largest US Internet Exchange Points (IXPs). Each probe-server pair executes a 3-ping test round every 15 minutes.

What is Measured

The latency measurements operate every hour of the day from consumer homes where users and devices are also generating internet traffic. Our latency value is equivalent to the latency observed by the home users during our test. *The latency values recorded are therefore observed latency*.

Recording Results

All tests are recorded by the RIPE Atlas system. NetForecast downloads the results daily into a database on a cloud service.

Data Analysis Using Apdex Methodology

We use the Apdex methodology to evaluate the measured RTT data. The methodology has two key aspects. First is the need to define the T and F thresholds as shown in Figure 4. These thresholds define the QoS states when applications using the internet are:

- **Satisfied** – the application has easily acclimated its behavior to the predominant latency range between the two locations the application is using to communicate.
- **Tolerated** – the application notices that latency has materially increased, but it knows how to adapt its behavior to still maintain a tolerable level of service.
- **Frustrated** – Latency has increased to a level where the application can no longer provide tolerable behavior, so the service is negatively impacted.

The above descriptions apply to any end-to-end application. If the application is serving humans in real time, then there are commensurate QoE stages as shown in Figure 4. However, the service may be completely unattended such as smart home monitoring, software updates, business processes, funds transfers, industrial factory controls, logistics data, shipping information, etc. The point is that if internet latency moves to the tolerated or frustrated stages, then any use case is at risk. This report describes conditions that are impacting QoS. Readers may want to know, "Well how will that impact gaming or browsing or streaming movies?" Answering those questions requires knowing the actual use case and specific application.

QoS applies to internet infrastructure. It is the essential foundation upon which consumer devices, value-added services and applications deliver consumer services. QoE of the final product which users see is based on the infrastructure's QoS. It is incumbent upon the infrastructure providers to maintain consistent performance. Abnormal deviations at the foundation level cause abnormal deviations at the user level. ISPs should not assume that specific adjustments to QoS will improve QoE. Doing so opens the door to assumptions about QoS-to-QoE relationships, often leading to troubling unintended consequences. The simplest and most advantageous strategy for service providers is to assure that QoS generally keeps improving or at least stays consistent.

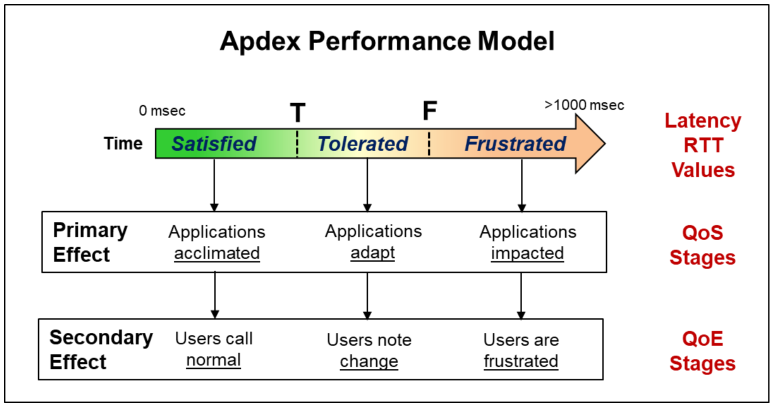

Figure 4 – The Apdex Performance Model

The Apdex process has three phases: 1) determine the T and F thresholds which define the three performance ranges; 2) count the frequency of samples within each range; 3) use the Apdex formula to calculate an Apdex value for the dataset under evaluation.

Selecting T and F thresholds is critical to generating meaningful Apdex values. In this case, to understand QoS of the internet in general and local ISPs in particular, we rely on the fact that the internet is generally operating properly for most subscribers under a broad mix of use cases. Therefore, we expect the data to present "normal" latency under "good" conditions. We found a few key days in the beginning of 2022 that met this criterion using pro forma calculations showing low variance along with minimal difference between near and far tests. Using the key good dates, we calculated means and standard deviations for each of the 55 probe-server pairs. We used those values to determine a custom T and F for each probe-server pair which remained fixed for the duration of this study. This is known as the *baseline model*.

The main analysis proceeds by grouping all RTT values within the satisfied, tolerated, and frustrated statistical bins for each of the 55 test pairs. Once so grouped, we analyze the:

- magnitude of RTT shifts across the bins, and
- frequency of samples within each bin.

The frequency of samples is used to calculate the Apdex value using the following formula (see Figure 5). Results of the formula are fractional with values ranging from zero to one. This simple zero-to-one scale is a convenient way to compare latency across a wide range of scenarios.

$$\text{Apdex Score} = \frac{\text{Satisfied count} + \frac{\text{Tolerated count}}{2}}{\text{Total sample count}} = \text{Score (0 to 1)}$$

Frustrated count have zero value

Figure 5 – Apdex Formula

Measurement, Integrity, and Analysis Assurance Steps

During the analysis process, several checks are performed to ensure a correct and complete report. The following is a list of critical automated measurement assurance algorithms operating before the Apdex summary analysis is performed.

Packet Loss Treatment
One or more of the three pings in a test round may be lost (no reply received). If this occurs for one or two of the ping tests in the round, Atlas supplies the mean with the number of successful tests. That value is recorded as the RTT. However, if all three tests in a round have no reply, then Atlas records a -1 ms value (clearly incorrect value). This is treated as complete packet loss. If the packet loss is isolated, then the RTT value is converted to a significant RTT value commensurate with the need for TCP to retransmit data.

Missing Server Detection and Mitigation
If packet loss repeats on a probe-server pair, then an automated algorithm is used to detect server loss rather than packet loss. Once a server is determined to be offline, all tests involved with that server are removed from the dataset until the server reappears.

Measurement Confidence
Results of each test pair are analyzed to assure that they are within the 95th Confidence Interval. The result of that analysis is compared to the baseline model that generated the T and F values for the test pair. If the 95th percentile point is outside the bounds of the T and F model, then the probe is removed from the study.

Minimum ISP Test Locations
Some of the above assurance algorithms may reduce the number of probes within a city metro such that the minimum number of probes testing on a specific ISP in that city falls below the desired threshold. If that occurs, all test results for that ISP in that city are removed from the study for the entire study period.

How Apdex Thresholds Normalize Near and Far Results

The T and F threshold process normalizes the effect of distance (speed-of-light issue) such that the near and far test pairs show nearly the same Apdex scores, as seen in Figure 9. The daily near and far Apdex scores plotted over the 12-month period in Figure 6 attest to the fact that this approach did normalize for distance. The difference between near and far Apdex scores was, on average, 0.014 and falling such that by October it was below 0.010 and remained so through December.

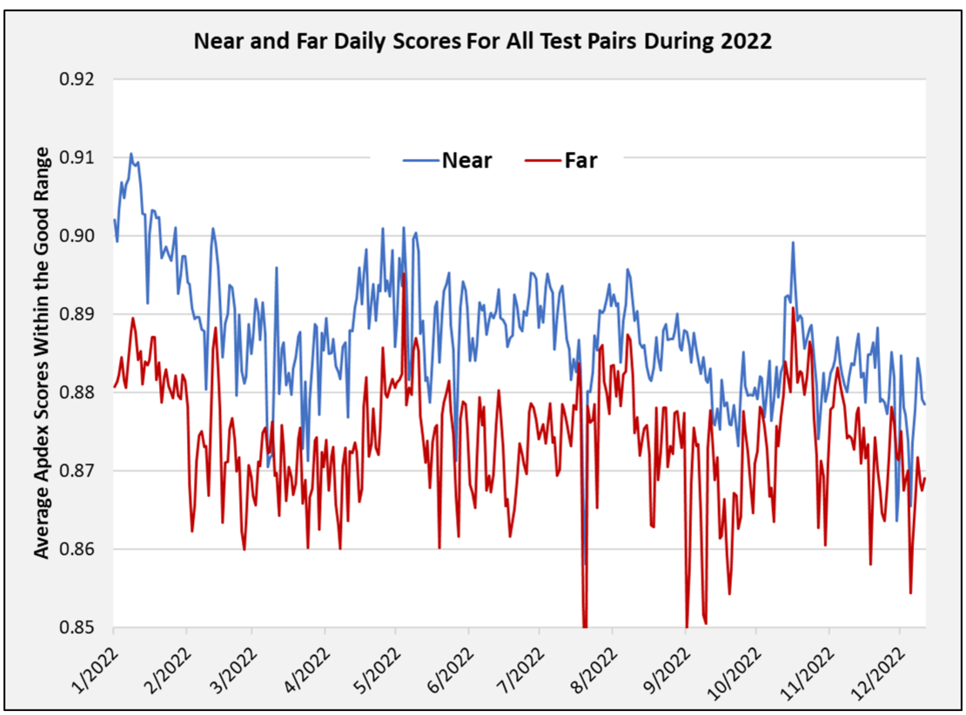

Figure 6 – Daily Near and Far Apdex Scores For 2022

How NetForecast's Benchmarking Differs from Other Testing

Unlike other testing services such as Ookla's Speedtest.net and SamKnows' testing for the FCC's Managing Broadband America Program, NetForecast's test results incorporate middle-mile networks and the contribution of both latency and loss into a single score that more accurately reflects the actual system performance. Also, NetForecast tests every hour of every day, not just when a user thinks the internet is slow or during a designated short testing period. NetForecast has many probes testing every hour, which total about 40 million tests per month.

While other testers focus on speed, NetForecast focuses on measuring latency and loss because they are critical parameters affecting the user experience. The responsiveness of interactive applications such as web browsing does not improve above 30Mbps. Since most US broadband customers experience speeds of 30Mbps or higher, speed-test results are not particularly informative. Other testers commonly present the results as averages, which can hide a long tail of serious performance degradation, thus masking critical instances when subscribers experience inferior performance.

KEY LATENCY FINDINGS

The internet is a complex system which is constantly changing and often experiencing unusually high latency events. The following discussion follows the Apdex three-range model (satisfied, tolerated, frustrated) to explain how latency increases due to distance and packet queueing delays.

Using Apdex to Identify Network Noise

Figure 7 shows the mean value of the satisfied range RTTs in each of the 55 test pairs during 2022 Q1. The values follow a linear pattern with an R^2 value of 0.97. The gray zone shows the RTT limit driven by the speed of light over distance. The 55 green dots do not conform exactly to the linear projection because each probe-to-server pair has a unique internet path. Packets do not follow the shortest great circle route.

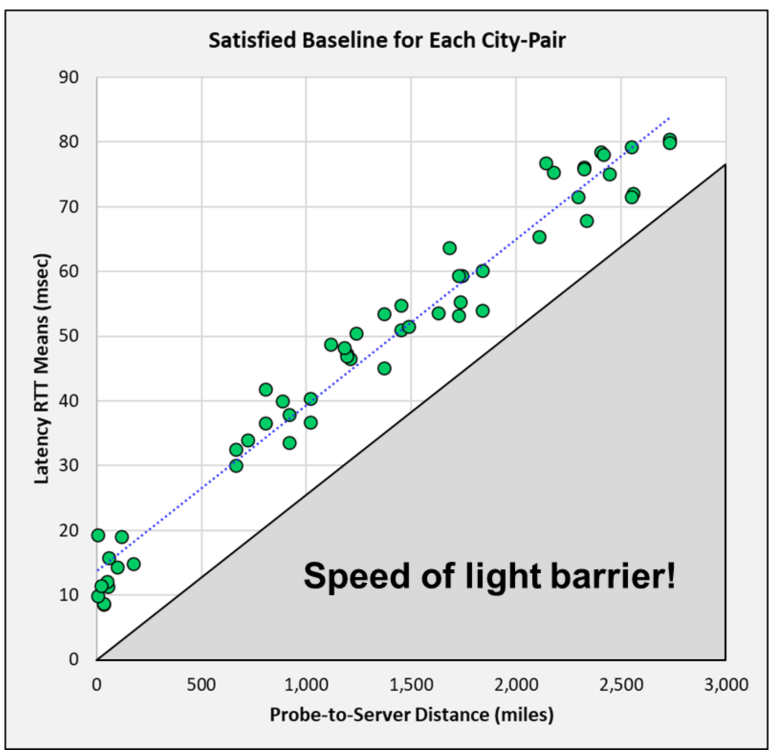

Figure 7 – RTT means for the Satisfied Range

Figure 8 provides a similar 2022 Q1 view showing all three Apdex performance ranges. It also shows the T and F thresholds used for this 2022 report. Furthermore, it identifies the clear near and far groups of test pairs. This figure provides a wealth of interesting findings.

The T and F thresholds shown as solid and dashed black lines are not perfectly linear with distance. This is because we used true RTT values on the baseline dates to determine custom T and F for each of the 55 test pairs. Therefore, the T and F values are also not following the great circle route due to general internet routing during the baseline period.

The geographic Near (within a metro) and Far (across the US) test groups are clearly delineated. By design, there are no test pairs in the approximately 200-to-650 mile distance span. (Note: measured 650 miles is greater than the design 450 mile limit because internet paths do not travel in a straight line.)

Here we see that tolerated is also very predictable because the tolerated range is between T and F values. However, frustrated has no "range" upper limit (any values above F are actual values measured). We can see that the mean of frustrated samples can go to extreme values. Notice that the red frustrated means have the same random and extreme patterns in both the near and far geographic zones. Frustrated values behave as noise irrespective of distance!

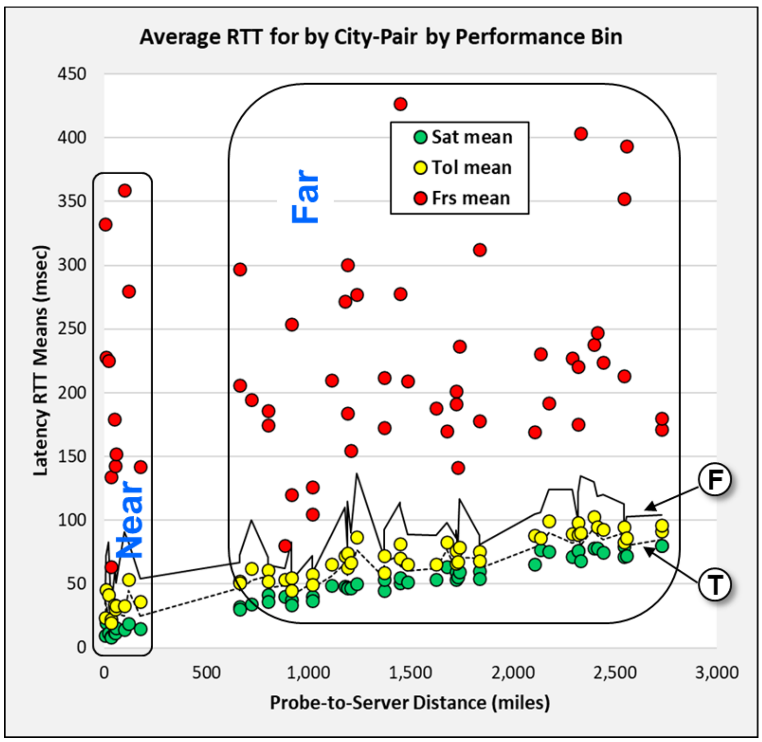

Figure 8 – RTT means for All Apdex Performance Ranges

In order to better understand how distance impacted frustrated latency, we group the RTT measurements into three general distance categories.

Near – all test within the probe's metro area (0-to-200 miles)
Far1 – tests past the metro area but generally within 1 or 2 time zones (650-to-1700 miles)
Far2 – tests past the metro area but generally within 3 or 4 time zones away (1700-to-2750 miles)

We then calculate the cumulative distribution function of all frustrated mean daily RTTs in the three distance groups for the entire 2022 measurement period. The results are shown in Figure 9.

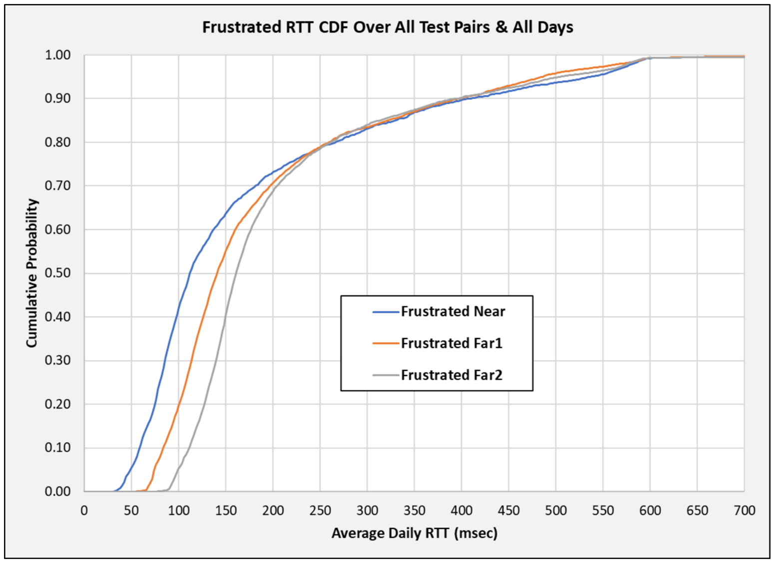

Figure 9 – Frustrated RTT Cumulative Distribution Functions

Figure 9 clearly shows that frustrated performance starts with different RTTs for each distance group but quickly converges to deliver essentially the same performance. The near tests, which are solely dependent upon the last mile ISP, quickly have the same performance as both far test groups. In fact, after about 200 msec, they all converge into a single performance profile! If frustrated latency were only caused by transit ISPs, then the near curve would be dramatically different. Under such a scenario, the near curve (blue line) would climb to 1.00 and flatten at much lower RTTs thus separating from both far curves. However, frustrated performance is equally evident across all network distances.

Figure 10 further validates this conclusion. Here we show the probability of frustrated events across not just the 55 city pairs but also the ISPs in the test program that serve each of the metro cities. There are 145 such unique ISP-near-far test pairs.

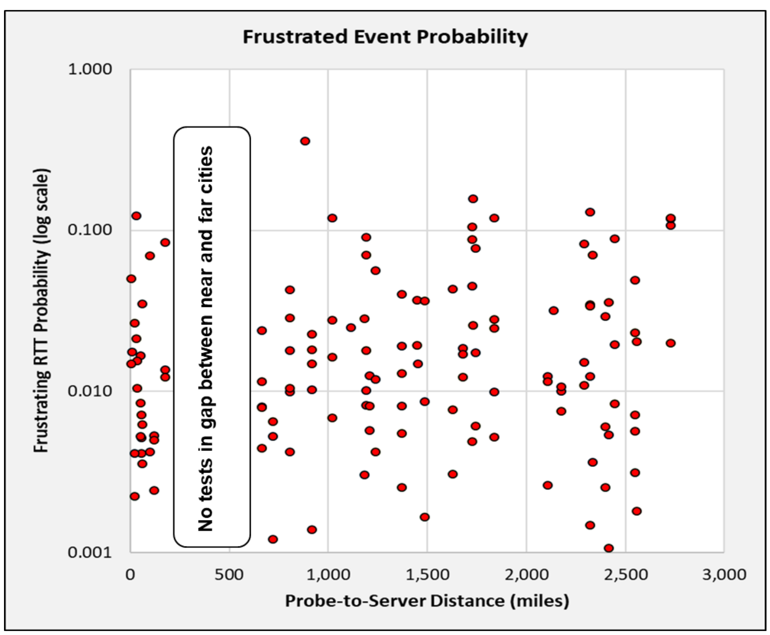

Figure 10 – Frustrated Events are Random

The red dots in Figure 10 show frustrated events occurring with over three orders-of-magnitude probability. There is no clustering or pattern of near vs far events. There are no increasing slope patterns which would indicate a distance relationship. This is internet performance noise which impacts any internet application or user. Frustrated events are caused within the last mile delivered by local ISPs.

Source of Frustrated Latency Events

Given,

that probes are hard-wired via Ethernet cables to the consumer's home router which eliminates the home Wi-Fi network from the probe-server path,

And,

Figures 8, 9, 10 clearly show that frustrated events are not driven by distance, transit networks, or datacenter conditions,

Then,

the only common place left generating the random frustrated events is the **consumers' last-mile local ISP**.

Investigating Network Noise by Last-Mile Technology

Are the frustrated events less frequent when grouped by last-mile technology? Figure 11 shows the overall probability of frustrated events for copper, cable, and fiber last mile solutions by quarter in 2022. Copper has significantly more frustrated events and their frequency climbed each quarter. Cable comes in second in frequency, while fiber has a significantly lower frequency of frustrated events.

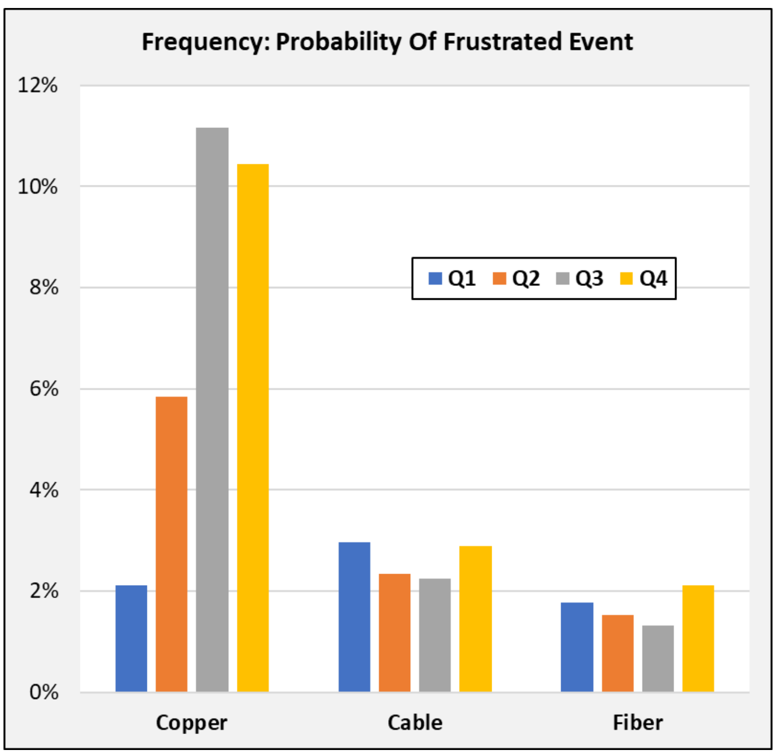

Figure 11 – Noise Frequency by Last Mile Technology

The actual mean latency RTT for these technologies is shown in Figure 12. It is interesting that the magnitude of frustrated events paints a picture that is opposite of the frequency shown in Figure 11. Copper had the highest frequency but lowest magnitude, while fiber had the lowest frequency with the highest magnitude. Cable has a middle-ground profile. However, the cable profile of frustration about 2.5% of the time, with an average RTT of 200 msec, is certainly a subscriber concern.

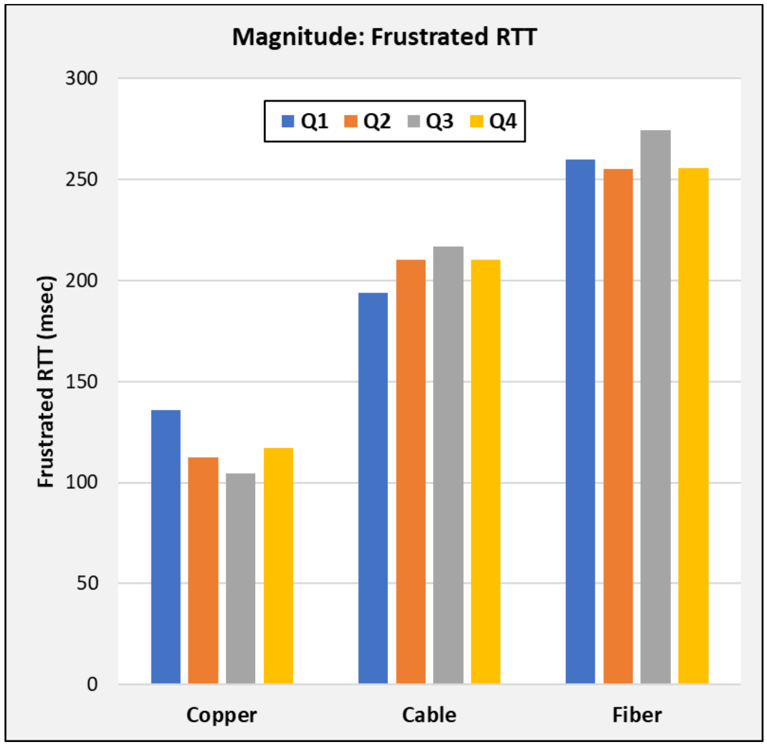

Figure 12 – Noise Magnitude by Last-Mile Technology

The Apdex methodology has three distinct phases. First, one needs to identify boundaries between satisfied, tolerated, and frustrated performance known as the T and F values. This section of our report makes a strong case for understanding the frustrated events. Most importantly, it shows that frustrated events A) are real and B) significant enough to matter to the application and user. The second phase of the Apdex methodology is to count how often satisfied, tolerated, or frustrated events occur. The third phase uses the Apdex formula to calculate an Apdex score for the specific subset of measurements of interest. The Apdex formula is based on the frequency of occurrence, not the magnitude of good or poor performance. Magnitude is covered in the selection of T and F boundaries.

If a sample set has all satisfied values, then the Apdex score is 1.00 or the best achievable. However, any number of tolerated or frustrated samples reduce the score. The relative occurrence of frustrated samples has the largest negative impact on the score (and, of course, the user). Frustrated counts are the leading cause of poor Apdex scores on page 1 and the following pages of this report.

DETAILED RESULTS FOR EACH QUARTER

Figure 13 shows the detailed Apdex scores for ISPs servicing each of the cities studied. These near and far values were averaged and summarized by city or ISP to generate the benchmark results shown in Figure 1.

The Apdex charts confirm that the normalization of distance worked properly. The average near and far scores across all of 2022 are just 0.014 apart. However, when viewed at the individual ISP level below, there are clearly larger distinctions in near and far performance. The QoS provided by local ISPs is important. Consumers, application developers, and content providers should take these differences into account. The Internet *is not the same everywhere* across the US.

Apdex Scores by Quarter

Figure 13 – Apdex Scores by Quarter in 2022

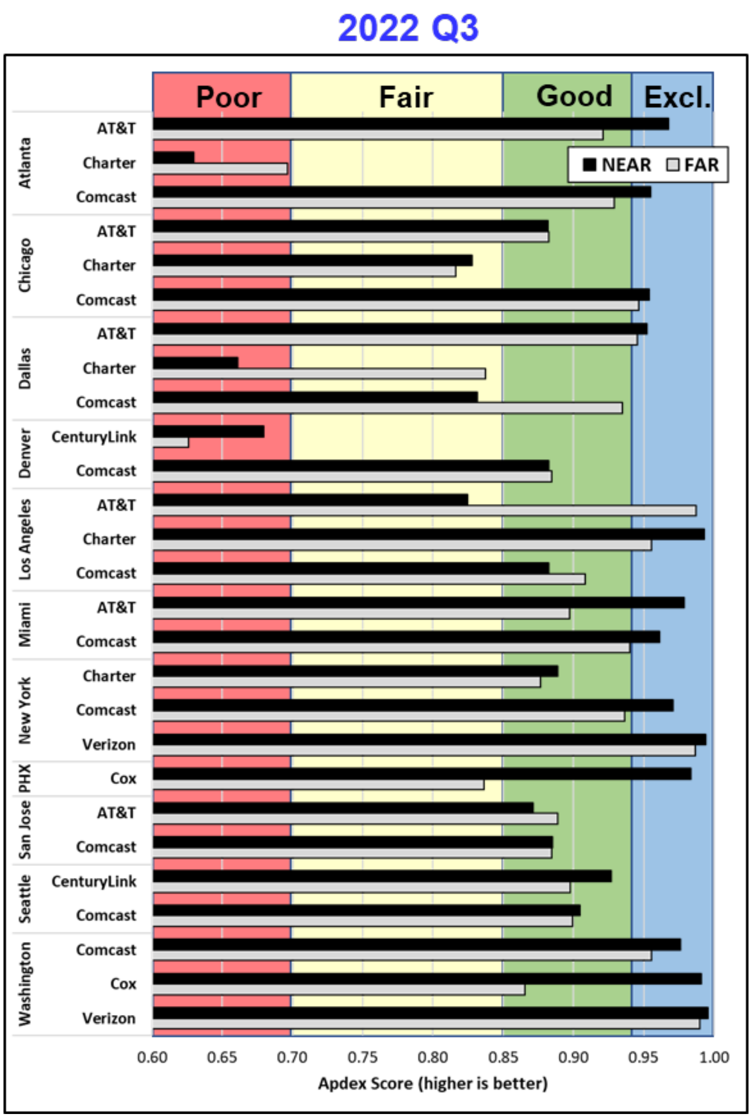

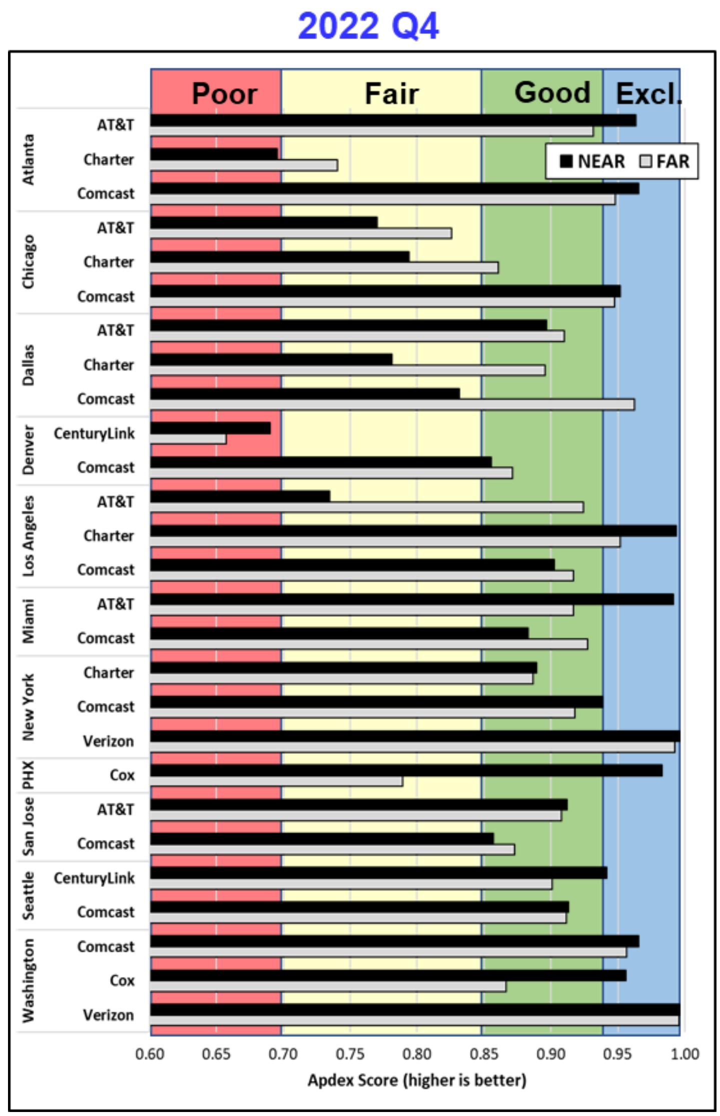

Figure 13 – Apdex Scores by Quarter in 2022 Continued

Frustrated Magnitude by Quarter

Figure 14 shows the detailed results for magnitude as the ratio of frustrated/satisfied RTTs. Note that the horizontal axis is a log scale. Also note that a near frustrated RTT has a more significant impact than a far frustrated.

Near magnitude change shift from satisfied to frustrated is always larger than the equivalent shift at far distances. That is due to near satisfied RTTs being very short, but the non-distance-governed frustrated RTTs in the *near geographic zone* are just as large as far frustrated. This is a byproduct of frustrated events' large magnitude and random occurrence characteristics shown in Figures 8, 9, and 10. Services that host content "close to the user" should take this into consideration. A frustrated event in the near zone can have a greater negative impact on applications than if the content came over a longer distance.

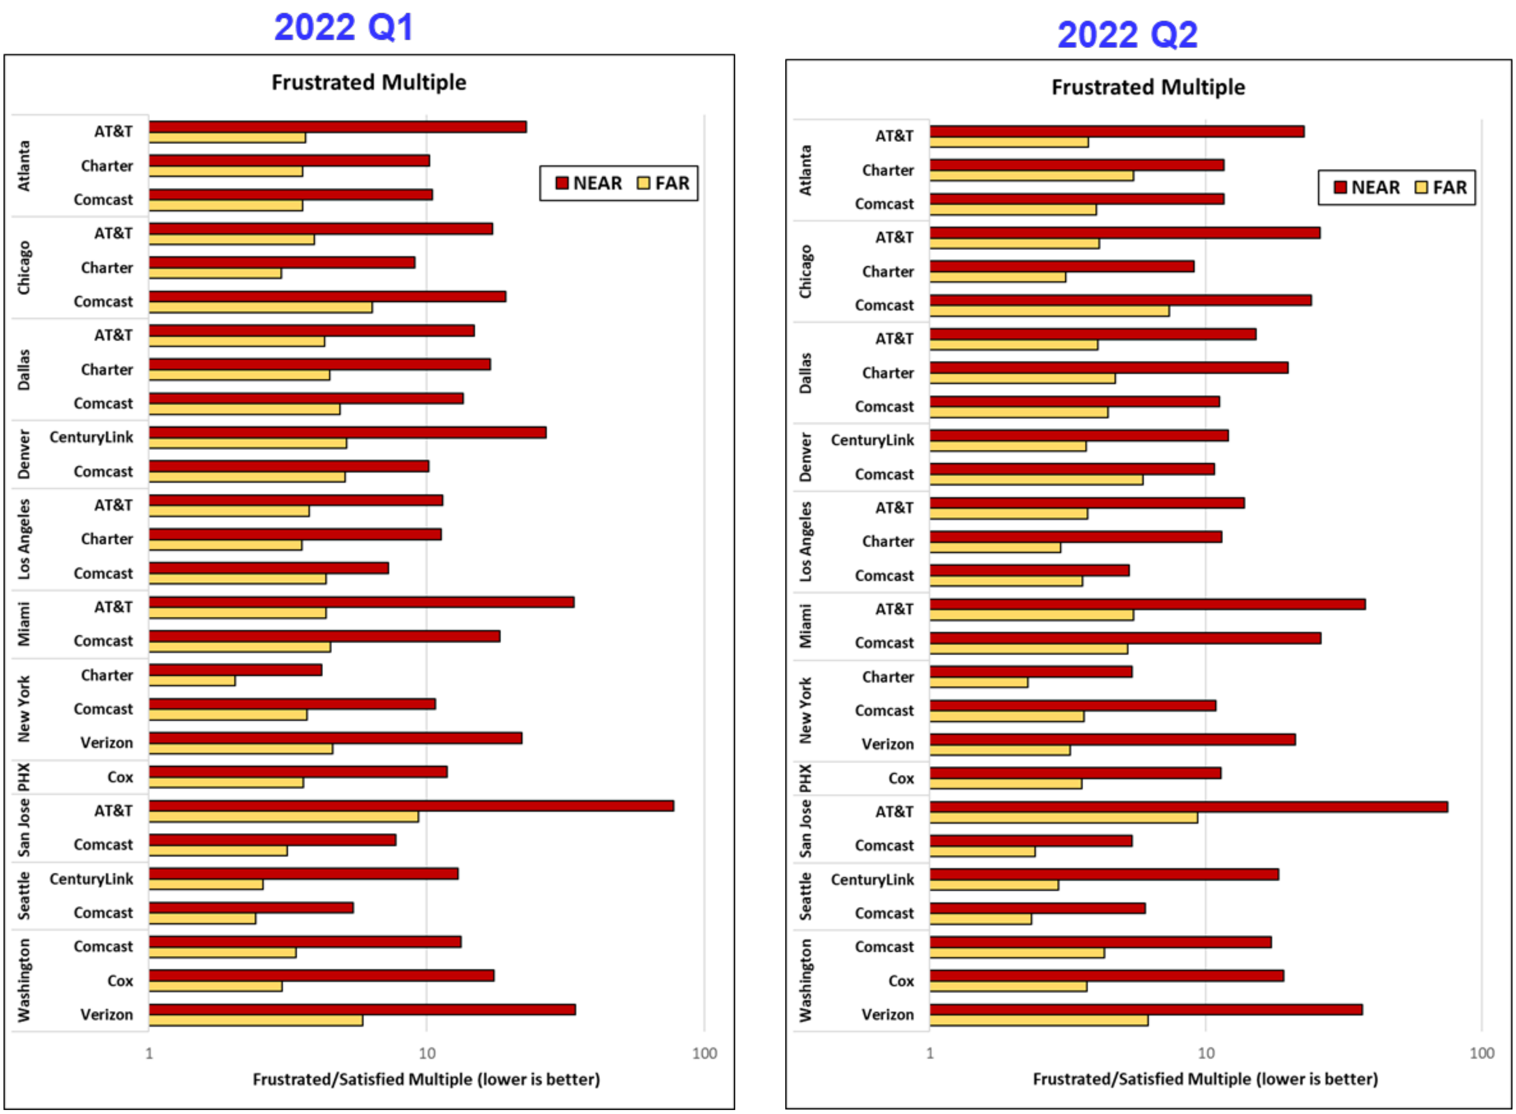

Figure 14 – Frustrated vs. Satisfied Magnitude by Quarter in 2022

2022 Q3

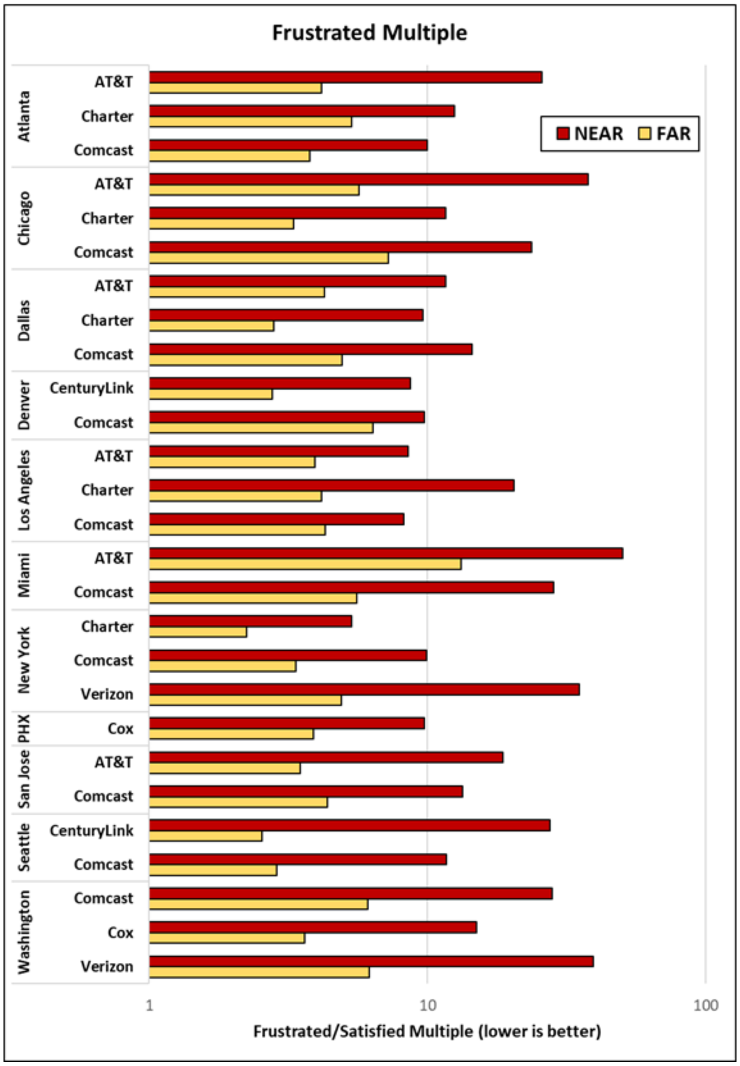

2022 Q4

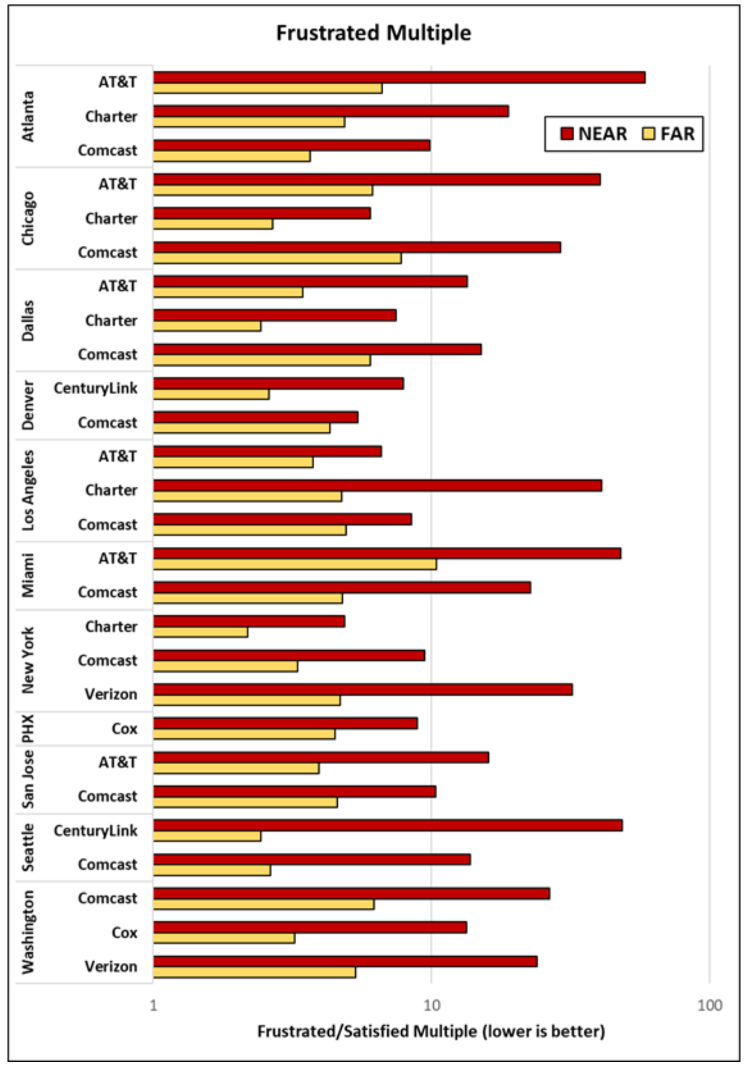

Figure 14 – Frustrated vs. Satisfied Magnitude by Quarter in 2022 Continued

ABOUT THE AUTHORS

Peter Sevcik is the Founder of NetForecast and is a leading network performance expert. An internet pioneer, Peter was among the first to measure and develop internet performance improvement techniques. He helped design more than 100 government, corporate, and commercial networks. In addition, Peter invented the Apdex performance reporting methodology, and has co-patented application response-time prediction and network congestion management algorithms.

Alan Jones is NetForecast's Chief Technologist. He has led teams in developing products and internal infrastructure for some of the largest telecom companies in the world. After eight years in cellular handset design and testing, he spent over a decade working on test systems for mobile networks. He currently works with mobile and cloud-based product development.

Andy Lacy is NetForecast's Chief Data Scientist responsible for developing deep analytics that uncover trends and anomalies in large datasets gathered by NetForecast tools or by client infrastructure. He developed key network metrics that account for subtle application and protocol behavior. He spent much of his career designing and developing real-time embedded systems and communications protocols where he was awarded several patents. He is a certified Scrum Master with a commitment to continuous improvement, test driven development.

Rebecca Wetzel is the President of NetForecast, and an internet industry veteran. She helped realize the commercialization of the internet in its early days and worked to design and market some of the internet's first value-added services such as IP-based VPNs, web hosting, and managed firewall services, as well as internet protocol testing services. She also spent many years as an internet industry analyst and consultant.

REFERENCES

1. FCC, Tenth Measuring Broadband America Fixed Broadband Report,
   Chart 9, Average webpage download time, by advertised download speed
2. Apdex Users Group, https://www.apdex.org
3. 2020 Internet Latency Benchmark Report, Peter Sevcik, Alan Jones, Rebecca Wetzel, https://www.netforecast.com/wp-content/uploads/NFR51382020NetForecastLatencyBenchmarkReport.pdf
4. Latency Explained, Broadband Internet Technical Advisory Group (BITAG), Technical Working Group Report, Jan 10, 2022., https://www.bitag.org/documents/BITAG_latency_explained.pdf
5. RFC792 Internet Control Message Protocol