NetForecast®

2020 Internet Latency Benchmark Report

Peter Sevcik, Alan Jones, and Rebecca Wetzel

April 2021

EXECUTIVE SUMMARY

How quickly an online application responds to a user's needs directly affects user productivity and satisfaction. The combined impact of network latency and packet loss on the user Quality of Experience (QoE) is at least as important—and in many cases more important—than speed. High latency and loss slow application response times for end users. Thus, NetForecast's independent Internet benchmarking service collects end-to-end test data to produce a score using Application Performance Index (Apdex) methodology [1] that uniquely incorporates the contribution of latency and loss into an informative single latency and loss score that is sensitive to deviations from normal. The benchmark incorporates three key factors: latency, consistency, and destination bias. This is a performance-comparative report, hence the term benchmark.

Whereas the common 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 covers latency and loss performance by city and by service provider during the period from March through December 2020. Figure 1 shows that Comcast and Verizon delivered the best overall ISP performance among the five ISPs measured and Atlanta and Washington experienced the best overall performance among the ten cities measured.

CITY BENCHMARK		
Rank	City	Score
1	Atlanta	0.94
2	Washington	0.92
3	San Francisco	0.85
4	New York	0.83
5	Denver	0.82
6	Los Angeles	0.80
7	Chicago	0.77
8	Miami	0.75
9	Seattle	0.74
10	Dallas	0.57

ISP BENCHMARK		
Rank	ISP	Score
1	Comcast	0.95
1	Verizon	0.95
3	CenturyLink	0.86
4	AT&T	0.85
5	Charter	0.79

Figure 1 – Performance Benchmark Rankings

NetForecast uses the publicly available RIPE Atlas measurement platform for our reporting. We selected approximately 500 Atlas probes from the available population of about 2,500 probes operating in the United States during the measurement period. As a RIPE Atlas Ambassador, we also operate Atlas anchor servers in the US. We appreciate RIPE for making the Atlas network available for internet performance research, and hope that our data analysis will benefit the entire internet community.

RIPE NCC
RIPE NETWORK COORDINATION CENTRE

NetForecast®

Click here to see daily US latency performance.

NetForecast Report NFR5138

Since 1999, NetForecast has been conducting in-depth analyses of internet performance to uncover where degradations occur and their impact on the end-user experience. NetForecast's QMap™ Internet Latency Benchmark Service is an outgrowth of over 20 years of network performance testing and analysis.

This report documents the results from a large-scale latency and loss measurement project funded and operated by NetForecast with no outside influence.

The report is based on 60 million latency tests over 10 months in 2020 using a rigorous and consistent methodology by which we can identify and document large and small performance shifts. The consumer subscriber lines from which the measurements were made are located within 10 major metropolitan areas, which encompass some 25 percent of US households.

WHY IS LATENCY IMPORTANT?

Latency is a fundamental, and often underrated, network performance parameter that significantly affects users' experience. It 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 (ms).

Interactions with internet servers require many round trips (aka turns) driven by applications as well as network protocols. Application responsiveness (or lag time) reflects the cumulative effect of RTT multiplied by the number of turns. Not only is application responsiveness a significant issue for human users of interactive applications it is increasingly important for autonomous systems. Latency can be improved by moving servers closer to users to reduce RTT, and by ensuring optimal network paths between users and the content they are accessing.

ISPs often advertise bandwidth as "speed," thus promoting a narrative that subscribers should buy higher speed services to improve application responsiveness. While this was once true when bandwidth delivered was measured in single digits, above 30 Mbps the benefits of higher bandwidth are marginal at best [2]. Given new technology and application initiatives the performance focus will shift to improving latency.

NEAR VERSUS FAR LATENCY

When accessing content, users 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 content in content distribution networks, where it is served locally).
- **Far** path is to the origin 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 interactive session. A typical use case operates over both near and far paths simultaneously. Content providers that can afford to place their content in many distant locations attempt to deliver from local servers to optimize application responsiveness.

ISPs are deploying new low-latency access technologies like low-latency-DOCSIS (LLD) for cable and 5G for wireless networks, which reduce latency across the last mile (modem-to-cable headend, or mobile-to-tower). However, these technologies cannot reduce the majority of latency in our near or far measurements.

NEAR/FAR LATENCY BIAS

An access ISP delivers packets to/from either near or far destination servers. Once the network distance to those servers is normalized, the near and far latency Apdex scores should be similar. A similar near/far bias score is a value within 0.05 Apdex points. Near and far Apdex scores from the same city hour (or ISP hour) with greater than 0.05 difference reflect a near or far bias. When near latency has a significantly better score than far latency, a latency bias favors near servers or services, and when far latency scores significantly better than near, a latency bias favors far servers or services. We define no bias hours as the hours when the average difference between near and far Apdex score is less than 0.05. NetForecast calculates how often paths to near and far destinations have similar Apdex scores—a value we call near/far bias. Near/far bias is shown as a 0-to-1 score where 1 is best (no bias hours).

KEY REPORT FINDINGS

Finding #1: Overall Internet Performance Improved During the Pandemic

The beginning of this report period coincided with the onset of the worldwide COVID-19 pandemic in March 2020. At the time, many reports and articles claimed that the US internet was struggling to support workers and students shifting from offices and schools to working and studying from home. The pandemic did change internet usage, as traffic that had traversed corporate networks moved to consumer networks. Fortunately, much of that traffic occurred during the day when consumer networks had previously been underutilized.

As Figure 2 shows, although broadband networks experienced unforeseen high demand that needed to be provisioned quickly, the overall latency picture was remarkably steady. The daily Apdex score averaged across all 10 cities in the report showed a 0.002 per month improvement rate as highlighted by the dashed trend line.

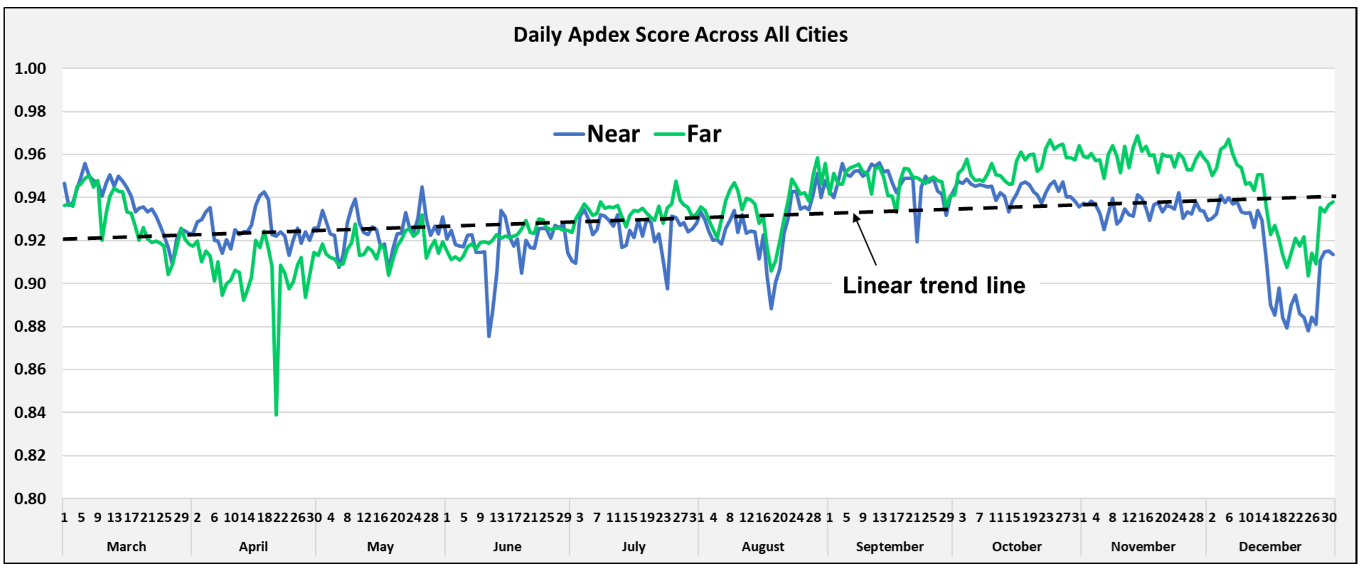

Figure 2 – National US Internet Latency Trend in 2020

Near and far scores matched well. Interestingly, near latency had slightly better Apdex scores than far from March through May, when the scores flipped to slightly better far over near scores. The distinction of near and far essentially disappeared in September, but then the scores again diverged with far scoring better than near through December.

The single day drop of the far score on April 21, 2020 was caused by a fiber cut to a major transit network in the Midwest US.

Finding #2: Near/Far Bias Existed 40% of the Time

Since the Apdex methodology was applied to provide an equivalent Apdex score for near and far destinations, ideally all near and far test result pairs should have the same score. Figure 3 shows how often near/far bias occurred each month. During the 10-month measurement period, near/far bias occurred in 40 percent of the site days.

Near bias was more prevalent in the early months, diminishing by November. Far bias was generally consistent at about 20 percent of the site days without much variation. The conclusion is that near bias occurred more often.

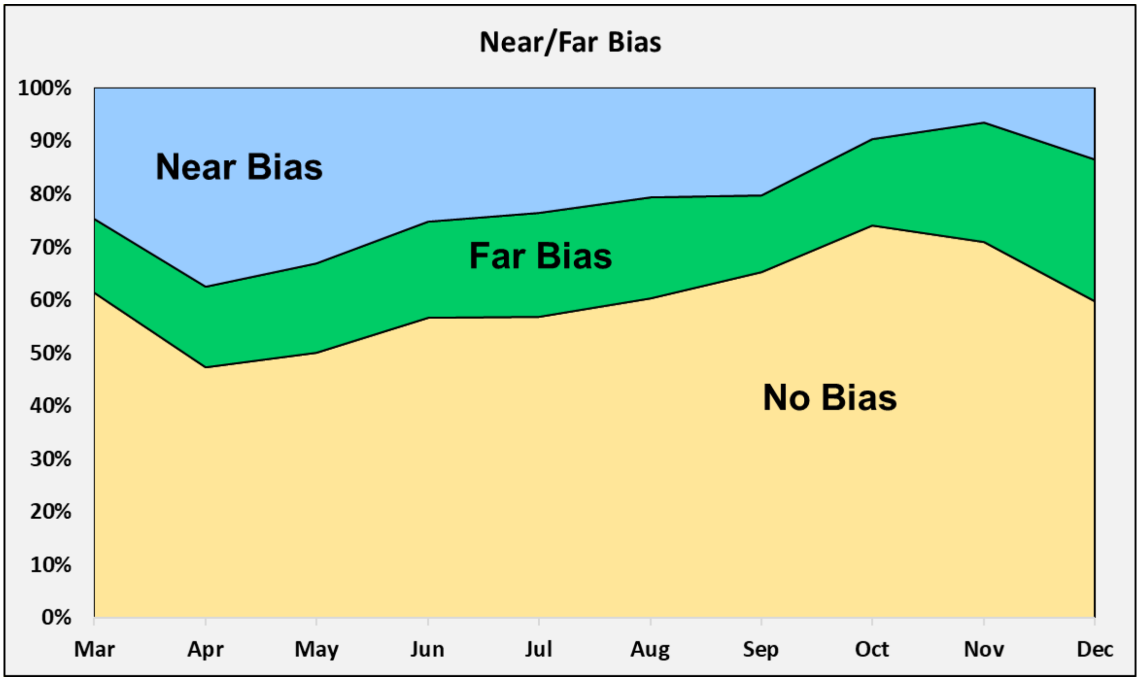

Figure 3 – Near/Far Bias by Month

The effect is that subscribers were sent to suboptimal content locations. Since users are accustomed to interacting with content in either near or far locations, switching users from one to the other will change the latency users experience, and they will notice. Predictable application performance is critical to user satisfaction. Large swings in near/far latency performance negatively impact the user QoE.

Finding #3: Latency Constantly Changed

Despite the steady latency in the aggregate view shown in Figure 2, when looked at in detail we see a very dynamic situation. Figure 4 shows the city and ISP benchmarks recalculated in the same manner as in Figure 1 (average of latency, consistency and near/far bias), but on a monthly basis.

The take-away is that when examined in detail, internet performance is highly variable. In fact, half of the cities and half of the ISPs changed their relative rankings each month!

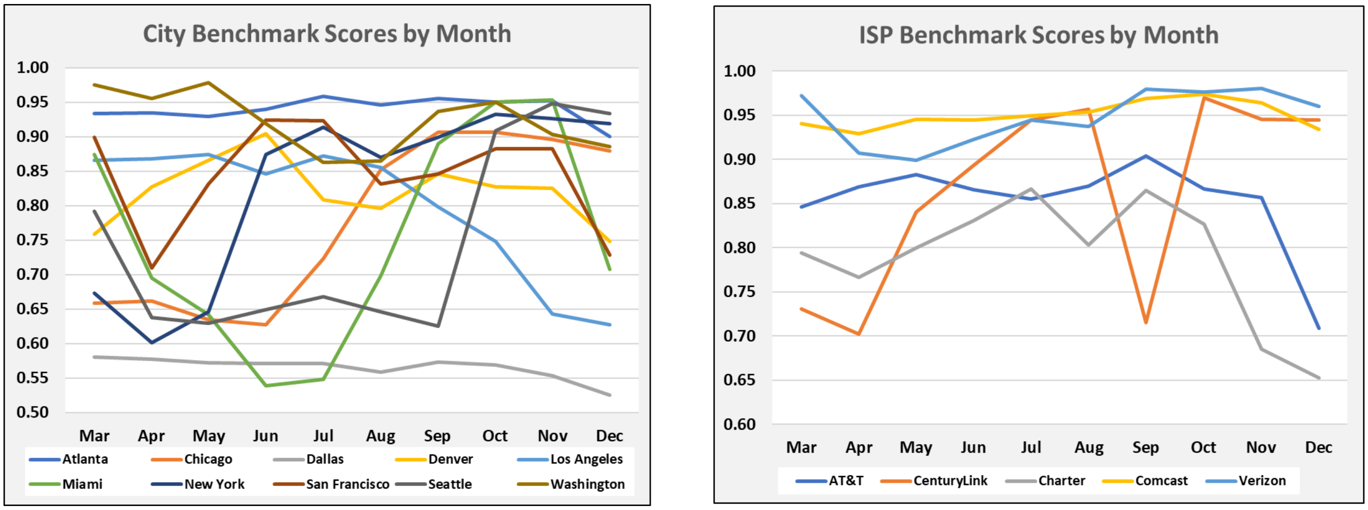

Figure 4 – Benchmark Results by Month

Although averages across broad time spans over many locations are informative, vital information is revealed in detailed and focused views of the data. The latter part of this report describes insights from detailed views, which can help service providers to improve service quality, content providers to save money, and consumers to make informed choices.

MEASUREMENT AND ANALYSIS METHODOLOGY

NetForecast's Internet Latency Benchmark uses a rigorous methodology to measure and report on the typical user experience. This report, which is the first in 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.

NetForecast tests using probes in consumer homes directly connected to major broadband service providers' routers (no Wi-Fi is involved). We perform standard ICMP [3] "ping" tests between the probes and reference servers (targets) located near the largest US Internet Exchange Points (IXPs). Probe-target pairs are carefully selected to generate a comprehensive US internet performance view.

What We Test

As Figure 5 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). Near testing covers the ISP's last mile, local peering, metro area networks and metropolitan data center access. Near tests simulate consumers accessing content hosted in edge service provider data centers or delivered via Content Delivery Networks (CDNs). Far testing covers the ISP's last mile, distant peering, middle-mile (transit) ISPs, transit-to-transit peering, and distant data center access. Far tests simulate users accessing content hosted at origination data centers and from sources that do not use edge or CDN services.

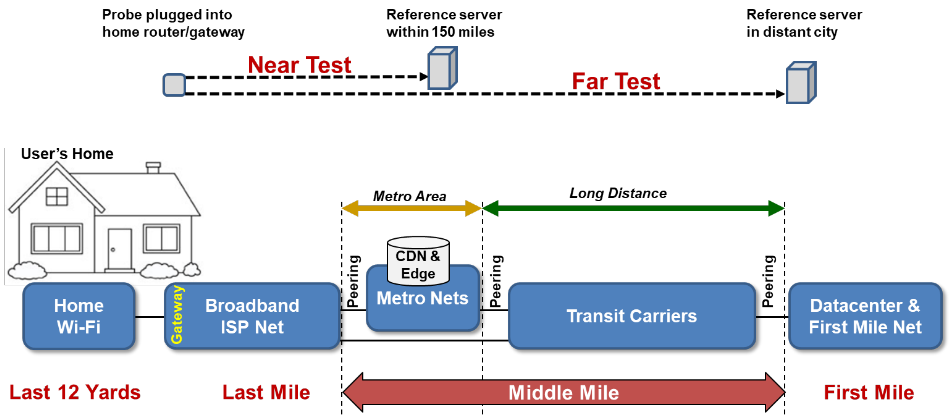

Figure 5 – NetForecast Near and Far Test Architecture

Figure 6 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 green lines. The metropolitan areas are defined by a circle with a 150-mile radius from the city center. The solid coast-to-coast long-distance paths are roughly equidistant, (great circle route). The dashed paths within the central US, although shorter, are also generally equidistant.

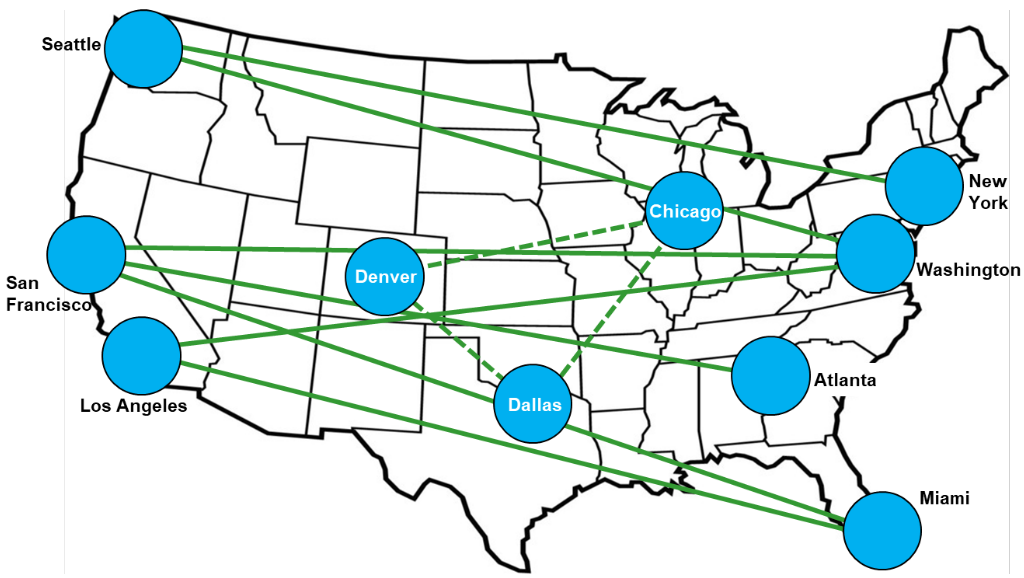

Figure 6 – NetForecast Near and Far Transit Paths

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 user experience. 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 2.7 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 [2]. 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 unhappy users, thus masking critical instances when users experience poor performance.

How NetForecast Uses Apdex

Apdex is an open standard for measuring the degree to which measured performance meets user expectations. NetForecast counts the number of RTT values and the number of test results that experienced packet loss, and places the counts into three user perception bins: **satisfied**, **tolerating**, and **frustrated**. The zones are bounded by a value "**T**" that delineates the satisfied and tolerating bins, and a value "**F**" that delineates the tolerating and frustrated bins (see Figure 7). The range of latency values is zero to infinity, where low values are better than high values.

> Users tolerate an experience that is four times slower than the fundamental response time, but they become frustrated when it is 16 times slower.

The satisfied zone encompasses the latency users have learned to expect—i.e., performance that is perceived as "normal." In the tolerating zone users notice the delay between a click and a reply, and in the frustrated zone the delay is so long that users are likely to give up. The typical application of Apdex relies on human judgment to set T and F. Often the choice reflects the biases of the people performing the analysis. Software developers tend to set the values too low because their perspective is how quickly their code responds running on a nearby dedicated server, while executives tend to set the values too high to paint a favorable picture with high Apdex scores.

Note that packet loss results are counted in the frustrated zone bin. Packet loss causes TCP to retransmit or video to pause as buffer gaps must be filled. Consequently, loss causes a large latency increase that leads to user frustration. Integrating loss and latency into a single Apdex score is valuable because it provides a single score that covers the two largest contributors to response time degradation.

Each ping test was generated by a sequence of three pings with return times. If one, two or all three pings had a return value then those values were averaged. However, if none of the three pings had a return packet received, then it was counted as a loss. The percent of loss is the total number of losses divided by total attempts in the hour for each city or ISP hour.

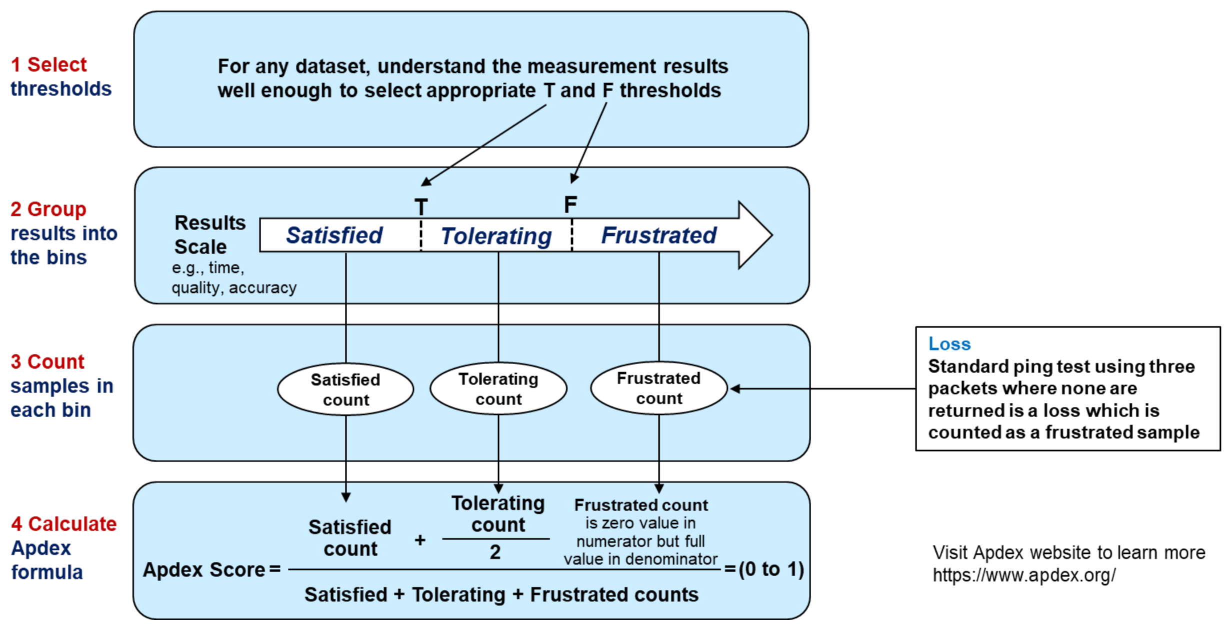

Figure 7 – The Apdex Process

NetForecast uses an unbiased automated statistical algorithm to select T and F values. This permits dynamically selecting T and F as needed and across distinct datasets. We applied the selection algorithm to baseline best-case test data across the total population. The T and F values remained fixed throughout the measurement period.

It is important to ensure that T and F latency values truly represent what users consider tolerating and frustrating. The baseline analysis shows that mean latency in the tolerating zone is four times that of the satisfied zone, and the mean latency in the frustrated zone is 16 times that of the satisfied zone. Four and 16 times are sufficiently significant increases in the fundamental response time of any application such that a user will consider the experience tolerable and frustrating.

Application redesign, protocol acceleration, or other techniques can improve application responsiveness. But soon the "new quicker" performance becomes the new normal. The user's perception of responsiveness is reset to the faster normal, which is represented by the satisfied zone of underlying network latency. Even a "quick" service will have the same four times or 16 times increases in application level responsiveness. *For this reason, the tolerating and frustrated performance zones apply to all interactive internet services*.

Once the samples are collected, NetForecast calculates the hourly Apdex score using the Apdex formula in step 4 above. The Apdex formula converts many measurements into a single score on a uniform scale of 0 to 1 (0 = no users satisfied or tolerating, 1 = all users satisfied).

Reporting Near and Far Latency with Apdex

NetForecast measures and assesses whether subscribers accessing near and far content experience the same relative latency impacts. Given that users cannot differentiate between a near or far path, but have traversed these paths before during "best case" periods, they are "trained" to view each "best case" as satisfactory quality performance. We assume the following:

- Users spend most of their online time interacting with familiar services/sites.
- Familiarity with performance causes users to accept that performance as normal.
- By repeatedly using the same online service, normal performance is perceived as satisfactory.
- Thus, when a service traverses a near or far path, even though latency is different, users are accustomed to that experience and perceive it to be satisfactory.

For this report, we converted hourly RTT results into an Apdex score showing when latency deviated from expected best performance. Apdex scores range from 0 (poorest latency) to 1 (best latency). Apdex thresholds were selected during low-traffic periods to determine baseline Apdex thresholds. Thresholds were selected for each metro area and tailored to normalize near and far scores to the same Apdex value.

The implication is that users within a metropolitan area where near and far Apdex scores are similar over a long period are receiving a satisfactory experience. If Apdex scores are the same but fall below the long-term baseline, then the user experience is degrading uniformly across the entire network infrastructure. If Apdex scores for near and far were similar, but begin to diverge, the path with the lower relative Apdex score has network characteristics that should be investigated.

How Apdex Thresholds Normalize Near and Far Results

After tuning T and F thresholds to properly reflect latency shifts that represent significant QoE changes, we adjusted the thresholds to account for latency differences due to probe-target distances. This normalizes near and far tests so they produce the same Apdex score for the same normalized QoE. Using the baseline analysis approach described previously, we finalized T and F for each metro area. Once validated, T and F remained fixed throughout the 10-month measurement period.

The daily near and far Apdex scores plotted over the 10-month period in Figure 2 attest to the fact that this approach did normalize for distance. First, the overall score differences are small despite vastly different actual near and far probe-server distances shown in Figure 6. Second, near/far advantages flip during the period. During September, the scores are nearly identical. No part of our testing and analysis approach changed over the 10-month period. Changes in Figure 2 are clear evidence that internet performance was changing.

Apdex Scoring Summary

The number of probes per metro area included in the report varied based on the number of available probes per ISP and the number of ISPs serving the area. All ping tests were initiated from the probes to target servers assigned to each metro area. Each probe was assigned to a set of near targets and far targets. We continually monitored the availability of probes and targets, occasionally shifting probe-target pairs to ensure availability, while carefully maintaining test integrity. An Apdex score was generated for each city and ISP hour, for a total of 216,000 Apdex scores which form the basis for the benchmark results.

Latency Factor

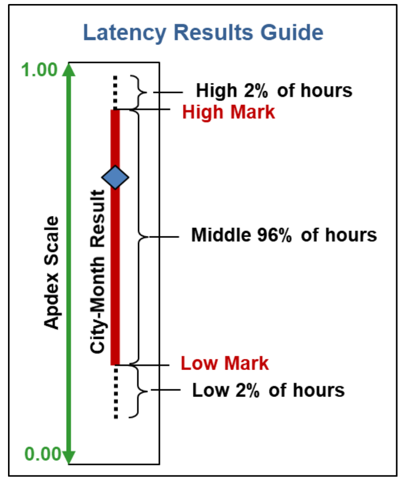

The aggregate monthly city latency view is the mean of approximately 720 near and 720 far hourly Apdex scores per city. We also calculated the 96th percentile of the respective 720 Apdex scores, which trims the top and bottom 2 percent of Apdex values. The top (highest score) and bottom (lowest score) of the 96th percentile define the range of values around the mean. The larger the range, the more variable the Apdex scores were for a given month.

Figures 8 and 9 show the monthly near and far latency results by city. As shown in the Latency Results Guide, a blue diamond marks the Apdex score mean and a red bar shows the span of the 96th percentile range.

The overall benchmark latency factor by city shown in Figure 10 is the mean of both near and far means across all 10 months for each city.

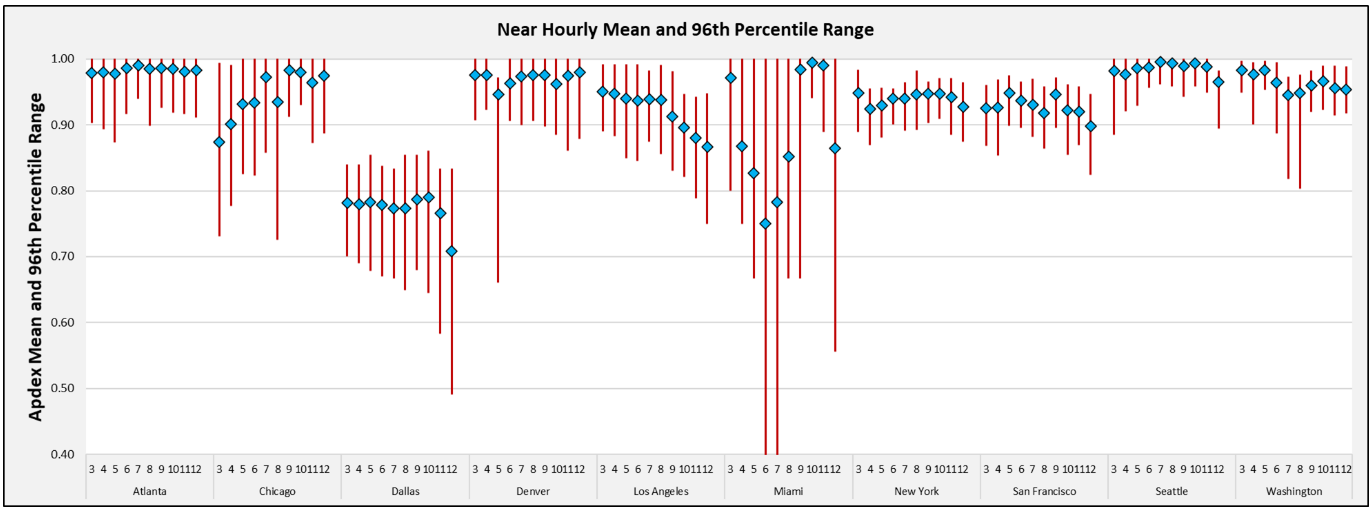

Figure 8 – City Near Latency Apdex Scores by Month

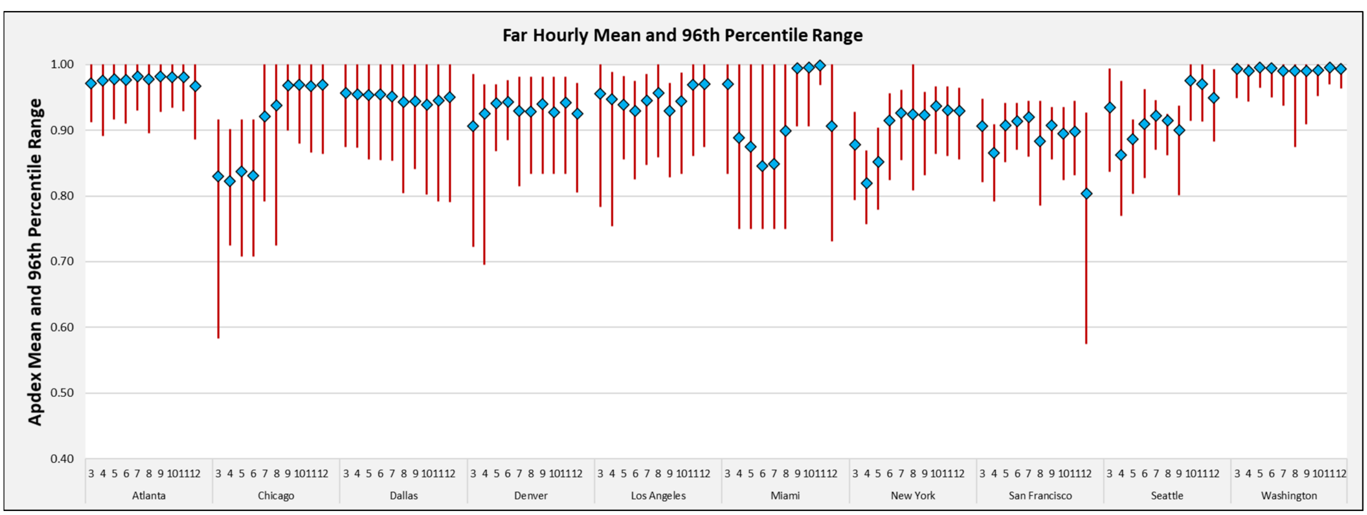

Figure 9 – City Far Latency Apdex Scores by Month

CITY LATENCY	
City	**Apdex**
Atlanta	0.98
Washington	0.98
Seattle	0.95
Denver	0.95
Los Angeles	0.93
Chicago	0.93
New York	0.92
San Francisco	0.91
Miami	0.91
Dallas	0.86

The overall latency Apdex scores when near and far are averaged are generally high as seen in Figure 10. The spread from Atlanta to Miami is just 7 Apdex points. Dallas is an outlier at 12 points below the highest-scoring city. The Dallas latency factor score was hurt by consistently poor near Apdex scores in the Dallas metro area as seen in Figure 8.

Figure 10 – City Latency Benchmark Factor

Consistency Factor

Consistency of Apdex latency scores is represented by the red bars in Figures 8 and 9. The larger the range, the more variable the Apdex scores for a city month. We calculate consistency as 1 minus the 96th percentile range as shown in the Consistency Factor Guide.

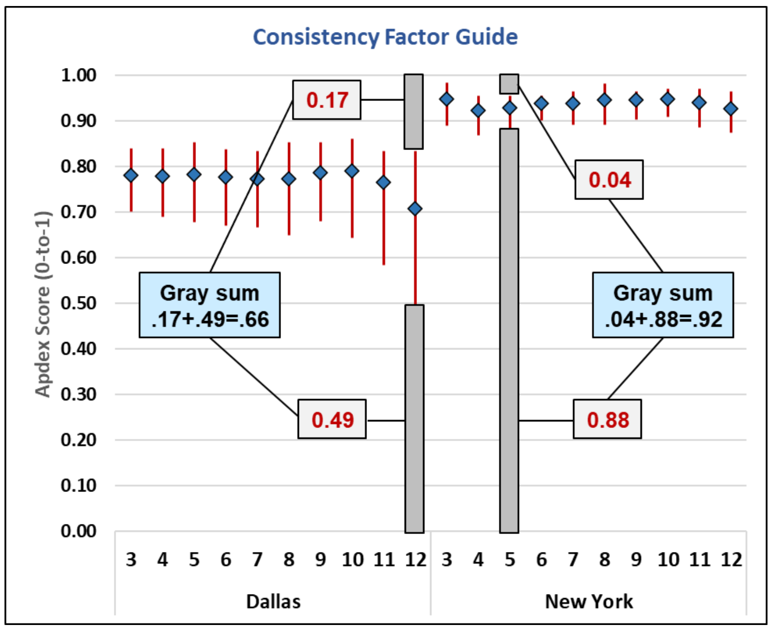

The gray bands above and below the two red bars (96th percentiles) cover the distances from the top of the red bar to Apdex 1.00 and from the bottom to Apdex 0.00. The shorter the red bar, the larger the gray band. Summing the gray bar Apdex values of top and bottom distances, yields a consistency score of 0.66 for Dallas in December and 0.92 for New York in May. New York had notably more consistent latency.

The overall benchmark consistency factor by city is shown in Figure 11.

CITY CONSISTENCY	
City	**Range**
Washington	0.93
Atlanta	0.91
Seattle	0.91
New York	0.90
San Francisco	0.89
Denver	0.85
Los Angeles	0.85
Chicago	0.82
Dallas	0.82
Miami	0.74

Consistency scores exhibit a 19 Apdex point spread between Washington, which was most consistent, and Miami, which was least consistent. Although Dallas had the worst latency, it fared better in consistency. The outlier was Miami, which experiences very inconsistent latency.

Figure 11 – City Consistency Benchmark

Near/Far Bias Factor

As mentioned previously, once normalized for distance the user experience should yield the same Apdex score for both near and far tests. If near latency has a significantly higher score than far latency, there is a bias towards a better experience accessing near servers or services. The converse is also true if far scores are better than near. In that case there is a bias toward a better experience accessing far servers or services.

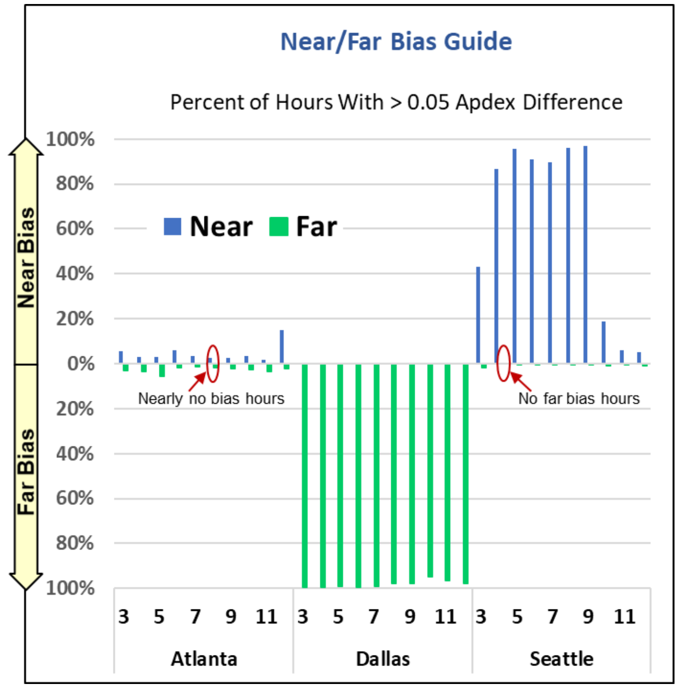

We define the Apdex score value of 0.05 as a significant level of difference. We define a destination bias hour as:

<u>Near Bias</u> = Near_score minus Far_score > 0.05
<u>Far Bias</u> = Far_score minus Near_score > 0.05
<u>No Bias</u> = neither of the above are true

We then determine for each city hour how often near or far bias hours occurred as a percentage of all hours in that month. The Near/Far Bias Guide shows examples of the three possible outcomes. Blue bars show near bias while green bars show far bias in each month. No bias over 720 hours in a month is extremely rare. Atlanta has one month with almost no bias.

Dallas experienced significant near bias for all measured months (horizontal axis is months 3 to 12). Seattle had significant far bias, which dissipated in October through December. In contrast, Atlanta experienced essentially no bias.

Figure 12 shows the percent of hours with bias for each city month. It clearly shows that there was significant bias in many cities and months. Atlanta and Washington were the only cities with small bias levels. Los Angeles shows a steady transition from near to far bias. Miami started in March with little bias but then amplified with both near and far bias hours than are a by-product of the erratic latency scores in Figures 8 and 9. It then settle down to less bias in September through November. December shows far bias reappearing.

The bias benchmark factor is defined as the lack of any destination bias. Both near and far are treated equally with equal Apdex scores. Therefore, the bias benchmark factor is 1 minus the sum of all near and far bias hours divided by all the hours across the 10-month measurement period. The equation is below.

$$\boldsymbol{Bias\ Factor = 1 - \frac{\sum Near\_Bias_{hours} + \sum Far\_Bias_{hours}}{10months}}$$

The results of the calculation are the bias benchmark factors shown in Figure 13. This factor like the previous two operates in the same unitless 0-to-1 range where 0 is worst and 1 is best.

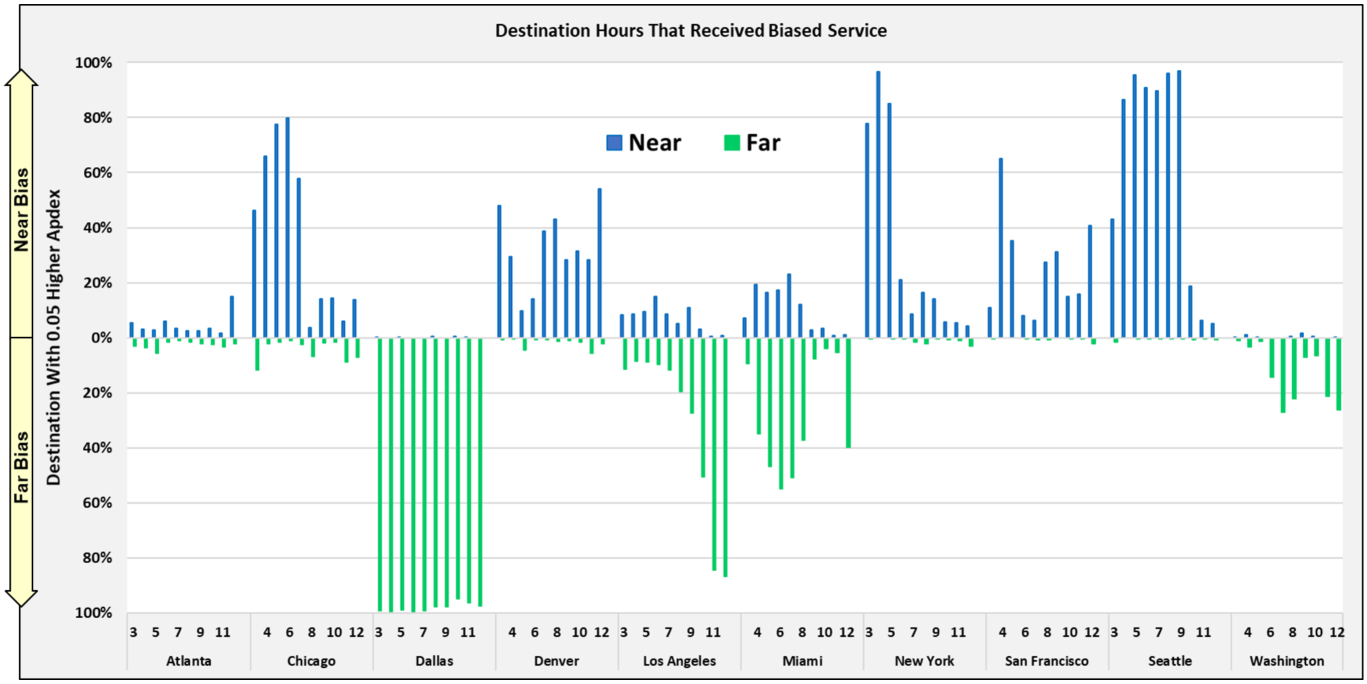

Figure 12 – City Percent of Near/Far Bias Hours by City-Months

CITY NEAR/FAR BIAS	
City	**Percent**
Atlanta	0.93
Washington	0.87
San Francisco	0.74
Denver	0.66
New York	0.66
Los Angeles	0.61
Miami	0.61
Chicago	0.58
Seattle	0.37
Dallas	0.02

The near/far bias benchmark factors have the largest high-to-low range of the factors in the study. Atlanta and Washington are clearly best with very little near/far bias while Seattle and Dallas are clearly worst with very high new/far bias. In fact, most of the cities had poor bias results relative to the top two cities.

Figure 13 – City Near/Far Bias Benchmark Factor

Latency Factor

Figures 14 and 15 show the monthly ISP near and far latency results using the same approach as in the city charts. The benchmark latency factor by ISP is the mean of both near and far across all 10 months for each ISP as shown in Figure 16.

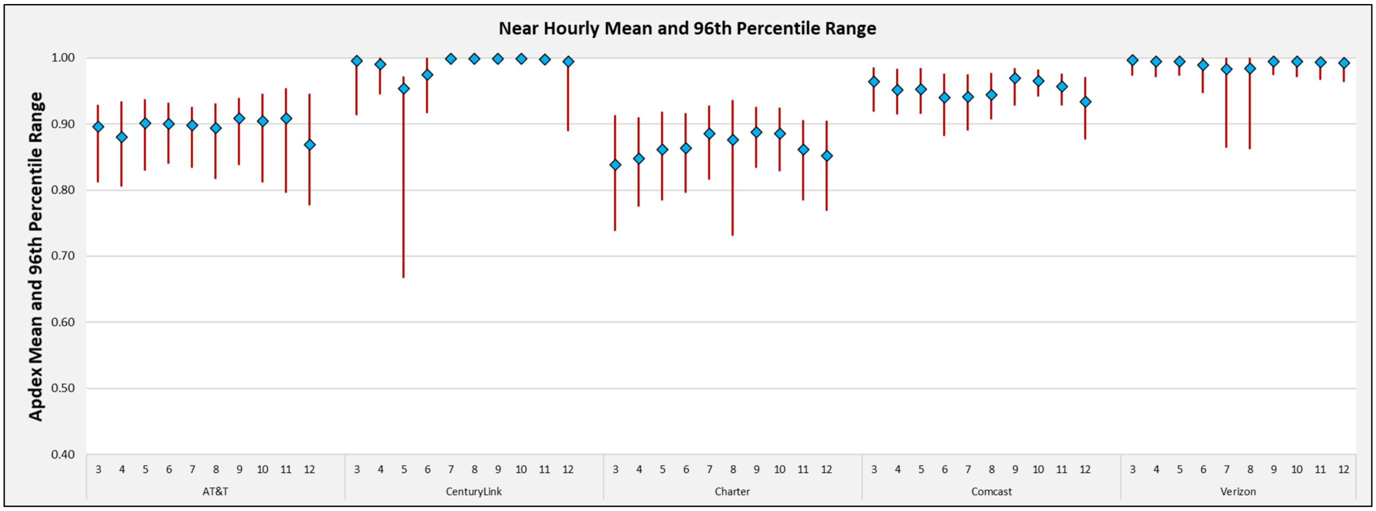

Figure 14 – ISP Near Latency Apdex Scores by ISP-Month

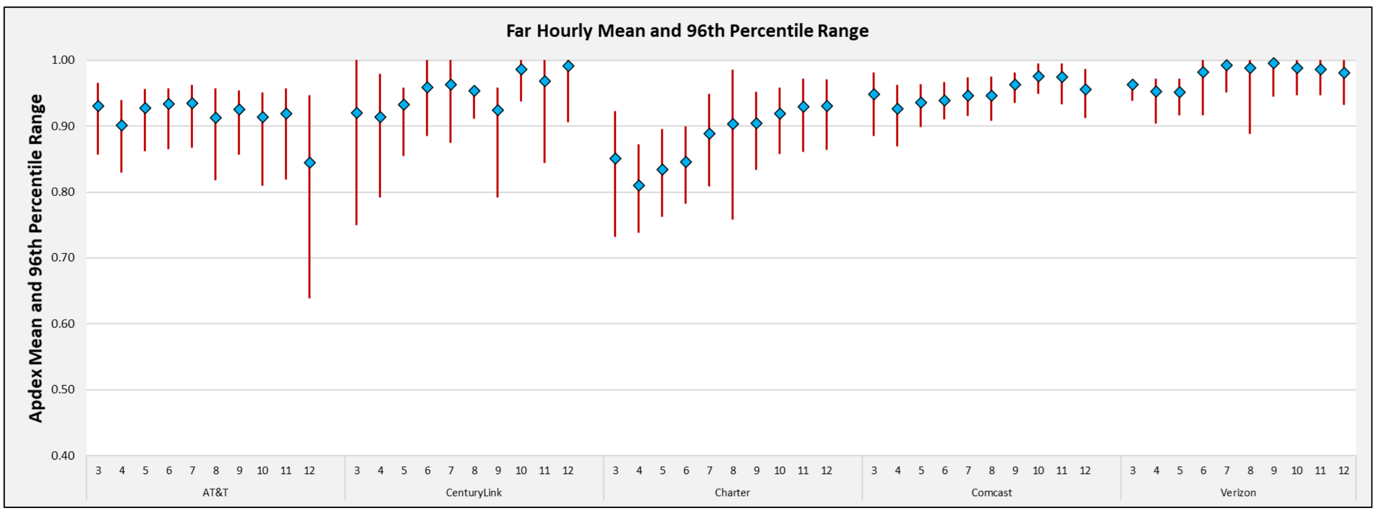

Figure 15 – ISP Far Latency Apdex Scores by ISP-Month

ISP LATENCY	
City	**Apdex**
Verizon	0.94
Comcast	0.93
CenturyLink	0.87
AT&T	0.87
Charter	0.86

Overall latency Apdex scores (the average of near and far means) are generally high. The spread from Verizon to Charter is just 8 Apdex points. When compared on an aggregate basis, near and far performance is nearly identical for city-month and ISP-month Apdex scores. Overall means and medians for cities and ISPs are within less than 0.01 Apdex points. This is expected since the city and ISP results are based on the same data set.

Figure 16 – ISP Latency Benchmark Factor

Consistency Factor

The overall benchmark consistency factor by ISP is the mean of both near and far for all 10 months for each ISP. It is shown in Figure 17.

ISP CONSISTENCY	
City	**Range**
Verizon	0.94
Comcast	0.93
CenturyLink	0.90
AT&T	0.87
Charter	0.86

ISP Apdex consistency scores match the latency Apdex results. At this level of aggregation, where the data is grouped into only five ISP buckets, variation is minimal.

Figure 17 – ISP Consistency Benchmark Factor

Near/Far Bias Factor

The near/far bias benchmark factor is defined as the lack of bias where near and far are treated equally. The methodology and bias factor equation is the same as used in the city analysis. The percent of near/far bias hours is shown in Figure 18.

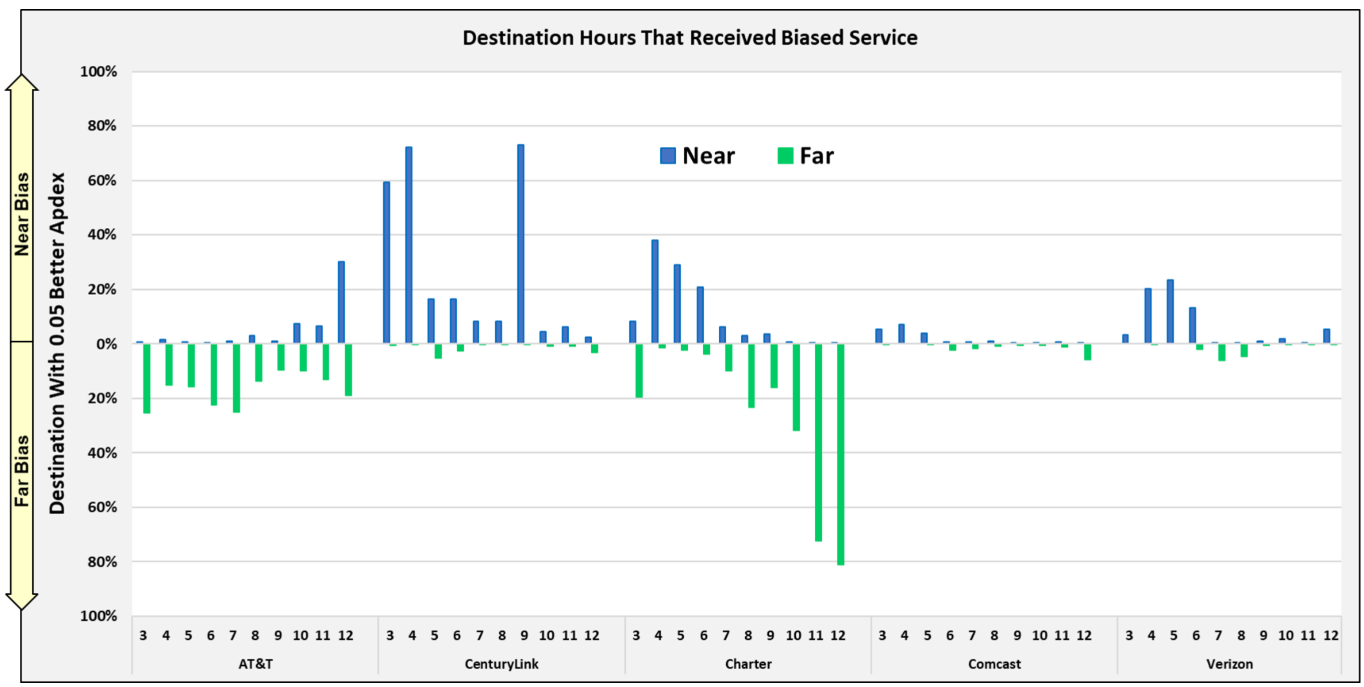

Figure 18 – Percent of Near/Far Bias Hours by ISP-Months

Charter has an interesting trend pattern moving from high near bias to very little bias and then high far bias over the 10-month measurement period.

Near/far bias uses the same formula as in city bias across all the hours across the 10 month measurement period. The results are shown in Figure 19.

The destination bias scores are clearly grouped with scores in the 0.90-to-1.00 range, no scores between 0.80-to-0.90, scores in 0.70-to-0.80 range and finally a score in 0.60-to-0.70 range. The bottom three ISPs clearly have a lot of near/far bias relative to the top two ISPs.

ISP NEAR/FAR BIAS	
City	**Percent**
Comcast	0.97
Verizon	0.92
AT&T	0.78
CenturyLink	0.72
Charter	0.63

Figure 19 – ISP Near/Far Bias Benchmark Factor

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 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 Director of Software Development. He has lead 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.

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. Apdex Users Group website
2. FCC, Tenth Measuring Broadband America Fixed Broadband Report,
   Chart 9, Average webpage download time, by advertised download speed
3. RFC792 Internet Control Message Protocol