# Gaming Latency vs. Engagement and Potential Impacts of Lower-Latency DOCSIS & Wi-Fi

A Technical Paper prepared for SCTE by CableLabs

**Stephen Froehlich**
Sr. Strategist
CableLabs
858 Coal Creek Cr. / Louisville, CO 80027
+1 303 661 3708
S.Froehlich@cablelabs.com

**Jacob Malone**
Dir & Principal Strategist
CableLabs
858 Coal Creek Cr. / Louisville, CO 80027
+1 303 661 3775
J.Malone@cablelabs.com

 

# Table of Contents

**Title** **Page Number**

1. Introduction 4
2. Data Source & Methodology 4
   2.1. Data Source 4
   2.2. Data Limitations 4
   2.3. Analysis Methodology 5
3. Interaction of Latency and Game Engagement 7
   3.1. Latency Overview 7
   3.2. Relationship Between Active Unique Days & Latency 7
4. Wired Latency 9
   4.1. Overall Wired Latency 10
   4.2. Routing Delay 11
   4.3. Queueing Delay 12
5. Wi-Fi Queueing Delays 16
   5.1. Wi-Fi Share 16
   5.2. Wi-Fi Delay 17
   5.3. Performance by Operator & Gaming Platform + WMM Packet Marking 18
6. Conclusion 22

Acknowledgements 22

Abbreviations 22

Bibliography & References 23

Appendices 23
Appendix 1 – Regression of Unique Active Days vs. Latency & Operator 23

## List of Figures

**Title** **Page Number**

Figure 1 – Per-User Latency by Measurement Quantile 7
Figure 2 – Latency vs. Unique Active Days by Measurement Quantile 8
Figure 3 – Correlation of Latency by Quantile to Unique Active Days 9
Figure 4 – 99%ile Wired Latency by Operator 10
Figure 5 – 99%ile Wired Latency U.S. Map 11
Figure 6 – Routing Delay by Operator 12
Figure 7 - % of Latency due to Queueing 13
Figure 8 – Queuing Delay Dist. by Operator 14
Figure 9 – Queuing Delay Dist – LLD Counterfactual 15
Figure 10 – Est. 99%ile Latency w/ Full LLD Deployment 16
Figure 11 – Wi-Fi Share 17
Figure 12 – Wi-Fi Delay by Operator 18
Figure 13 – Wi-Fi Performance by Operator and Gaming Platform 19

 

# List of Tables


**Title** **Page Number**

Table 1 - Distributional Comparison of $latency_{NN}$ + $jitter_{NN}$ to Raw Latency, Duration: 700 Seconds ........... 6

Table 2 - Regression of Wi-Fi log(latency) vs. ISP & Gaming Platform ........................................ 21

Table 3 - Regression Output of Unique Active Days per user vs. Z-score of Latency & Operator by Measurement Quantile ........................................ 24

# 1. Introduction

As the industry gives monitoring and optimizing latency more attention, there is an increasing interest in understanding the value created by that network investment. Gamers are often mentioned as a key user segment that would value such investment because of the improvements to their user experience when playing games online.

Using a dataset from Network Next that includes 10-second latency data for every session of a multiplayer, competitive game during the first 2 months after its launch, we estimate that a one standard deviation improvement in 99$^{th}$ percentile (abbreviated as "99%ile" for the remainder of the paper) latency results in an 8% increase in the number of unique days of game play on average (approximately 0.4 extra unique days).

In addition, when a 7ms ceiling is applied to users' queuing delay, to estimate the impact of Low Latency DOCSIS® technology, we find that some cable operators could see up to a 50ms improvement in 99%ile latency, while the improvement for other operators is smaller due to variation in queuing delay across cable operators. However, the user experience of gamers on all cable broadband networks are likely to benefit from the greater consistency in observed latency.

Lastly, the Network Next data allows analysis of Wi-Fi performance by gaming platform (PS4, Xbox One, Switch, Windows) on a per operator basis. Wi-Fi connections are used by 68% of the players of this game. Given that Wi-Fi is one of the largest sources of latency on the network, any complete solution to improving latency must address home Wi-Fi networks.

One potential technology to improve Wi-Fi latency is Wi-Fi Multimedia (WMM) tagging. For this game, the Nintendo Switch uses WMM. We find that for the Nintendo Switch, Wi-Fi latency is 46% lower than that of a PlayStation 4 (PS4), a console that does not use WMM. We also find there is substantial variation in Wi-Fi performance across operators.

# 2. Data Source & Methodology

## 2.1. Data Source

For the first two months following the launch of a game in the summer of 2021, Network Next provided CableLabs latency observations for each 10-seconds of activity. Each observation included the minimum latency observed during that 10-second window, jitter (defined as the standard deviation from the minimum latency observed during that 10-second window), and packet loss. Each 10-second observation is a summary of 100 pings.

The metadata include the user ID, a session ID, a postal code for the user and for the server, a timestamp, the gaming platform, whether the session was over the wired or wireless interface on the device, and the ISP name. The global data set includes 6.4 million players of the game, playing 260 million sessions, and 9.8 billion valid 10-second observations. When filtered down to the top 14 operators in the lower 48 states of the U.S., the data set includes 1.6 million players of this game who played for at least an hour.

## 2.2. Data Limitations

The Network Next data set comes with two notable limitations. First, the relationship between engagement and latency that we observe are not externally valid for other games or applications. This is because the sample only contains data on a single game, which limits how generalizable our results are, and it has a unique relationship between latency and engagement, which is discussed in Section 3.2. In short, maximum engagement for this game peaks where there are more players (i.e., the latency experienced by *most* players), not at the *minimum* latency observed; this relationship creates an upside down "U" shape that is shown below. Network Next says they have verified this result but say it is unique to this game based on their experience studying other games. Other games tend to have flat engagement out to approximately 100ms of latency, and then it decreases. However, more generally, we hypothesize that the relationship between lower latency and engagement are related to a game's characteristics. For example, players of real-time puzzle games may value latency improvements differently than those who play fighting games.

Second, some believe (Broadband Internet Technical Advisory Group, 2022) that online gaming is sensitive to the 99th percentile packet latency. This sample only provides minimum latency and standard deviation every 10 seconds. By comparison, we try to estimate the relationship between the measured minimum latency in the Network Next data to packet latency percentile, using empirical data collected at CableLabs, but this relationship is imprecise. Also, as shown in Table 1, the relationship between Network Next and actual latency measurement quantiles includes a couple of step functions that could easily move around based on network conditions, adding uncertainty to the meaning of the lower quantiles.

### 2.3. Analysis Methodology

We considered four different variables to proxy for game engagement:

1. Total game time played
2. Unique active days
3. Time between first and last session
4. Average time per Unique Active Day

Each of these four variables correlate well with latency, but unique active days was chosen because it had the strongest correlations, consistent region-to-region, and closely relates to how Network Next tracks retention. Using unique active days also mutes the effect of differences in free time between gamers. Finally, we propose, without empirical evidence or market research to confirm, that choosing to play a game again on a different day likely speaks better to a gamer's enjoyment from playing more so than a longer playing session on a single day.

Based on modelling against lab data and input from Network Next, we find *latency + jitter* for each 10-second observation (summary of 100 pings) is the best stand-in for actual latency. For each user, the primary input variable used is their quantiles of 10-second *(latency + jitter)* observations.

As discussed above, the Network Next measurement system is based on 10-second windows, where 100 "ping" packets are sent and summarized in terms of minimum measured round-trip time and the root mean square deviation from that minimum. For the round-trip time *(rtt)* of the 100 pings in each 10-second observation:

$$latency_{NN} = min(rtt)$$

$$jitter_{NN} = \sqrt{\sum \frac{\left(rtt - min(rtt)\right)^2}{100}}$$

 

To estimate the relationship between a Network Next's $latency_{NN} + jitter_{NN}$ measurements (i.e., individual summaries of a 10-second window ping) and the raw latency distribution over a period of observation, we measured raw per-packet latency for a cable modem connection carrying a mix of traffic for 700 seconds and calculated the $latency_{NN} + jitter_{NN}$ on 10-second windows. We then compare percentiles of the $latency_{NN} + jitter_{NN}$ distribution to percentiles of the raw latency distribution. The results are summarized in Table 1 below. For example, we find that the 99$^{th}$ percentile $latency_{NN} + jitter_{NN}$ value is roughly equivalent to the 96$^{th}$ percentile raw latency value.

**Table 1 - Distributional Comparison of $latency_{NN} + jitter_{NN}$ to Raw Latency, Duration: 700 Seconds**

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Network Next Quantile | 4.0% | 14.0% | 30.0% | 45.0% | 56.0% | 75.0% | 90.0% | 95.0% | 99.0% |
| Est. CL Latency Measurement WG Quantile | 25.0% | 52.3% | 57.1% | 78.1% | 90.1% | 94.0% | 95.4% | 95.7% | 96.1% |

Regression analysis is used in Section 3 to estimate correlations between latency and game engagement. As described earlier in the first part of Section 2.2 and more completely below in Section 3.2, this game has a unique upside down "U" shape in the relationship between latency and engagement. The upward slope approaching the peak (i.e., the portion of the curve that describes how the lowest latencies in the sample relate to engagement) is filtered out, so only data from the peak and after is used in the regressions. More exactly, mean engagement is calculated for gamers (minimum 10 gamers) by operator and state and the maximum value is found ($L_{pe}$). User data between $L_{pe}$ and $3*L_{pe}$ are used in the regressions. All regressions include operator fixed effects to account for any time invariant operator-specific variance in the data.

To determine wired queuing delay, the mean routing delay *(r)* was determined for each user connection to a given server from a given location over a given ISP. R is estimated by the mean of the 10-second minimum latency for each combination of user, game server, and ISP. Then queuing delay for each 10-second observation was calculated as $latency + jitter - r$. Reported quantiles in Section 4 are of this distribution.

Cleanly isolating Wi-Fi delay was difficult. For a single user, differentiating jitter due to wired queuing from jitter due to Wi-Fi queuing is impossible. Therefore, we matched pairs of wired and wireless users from the same postal code connecting to the same server from the same ISP at the same time using the same gaming platform. As there were only two wired Switch players in the entire data set, Switch wireless gamers were matched with Windows wired (Dowle & Srinivasan, 2021) gamers. This resulted in a set of 42,252 user pairs with at least an hour of overlapping activity. The latency (*latency + jitter)* of the wired user was subtracted from the latency of the wireless user to get an estimate of latency from Wi-Fi for the wireless user. After that, quantiles of Wi-Fi latency (latency + jitter) were calculated on a per-wireless-user basis.

An exponential regression of ISP and gaming platform dummy variables was performed to see the effects of different ISPs, gaming platforms, and the interactions between the two on Wi-Fi latency.

 

# 3. Interaction of Latency and Game Engagement

In this section, the focus is on latency and game engagement. In particular, how different latency percentiles compare to each other and whether the relationships with engagement are significant. Overviews of latency and its relationship with game engagement are first, followed by the regression analysis.

## 3.1. Latency Overview

Figure 1 shows the cumulative distribution of latency per user measured at different quantiles (note the x-axis is log scale). While quantiles 75% or below are generally under 70ms, latency grows quickly at higher quantiles, especially the 99%ile.

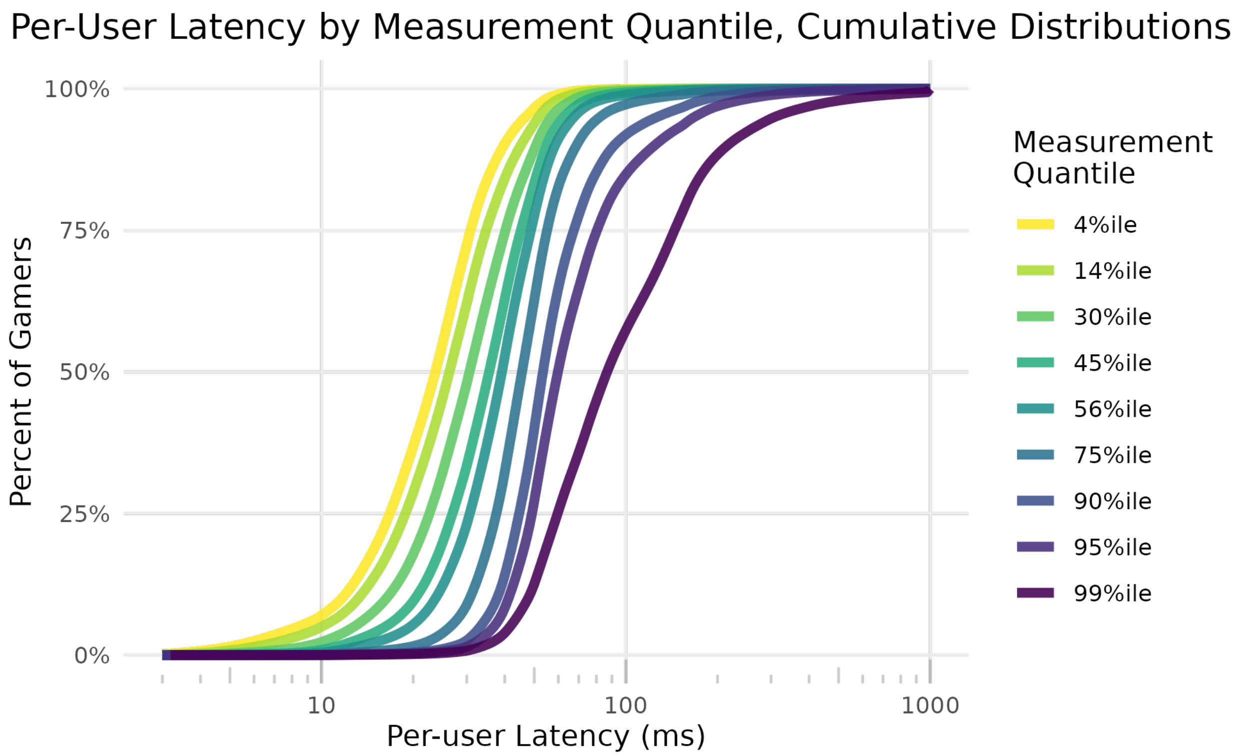


**Figure 1 – Per-User Latency by Measurement Quantile**

## 3.2. Relationship Between Active Unique Days & Latency

As shown in Figure 2, the game used in this study has a unique response curve between latency and unique active days playing the game. The peak unique number of active days is wherever most players are and is near the average latency for each operator, causing occasional double peaks or at least visible “shoulders”. The shoulders in the 56%ile, 75%ile, 95%ile, and 99%ile curves are caused by the presence of an operator whose mass of latency is at that point. The plot uses a logarithmic x-axis to emphasize the slope of each curve.

According to Network Next, other games have flat engagement to about 100ms and then it falls off. However, in this game, those with better-than-average 95%ile latency play significantly fewer days than those with average latency. While the cause remains unknown, Network Next is working with the game developer to test several hypotheses.

As you can see from the plot, the steepest falloff is in the 56%ile and the 75%ile, and this will be reflected in the regressions. However, given the step-function nature of the correlations found in Table 1 this could be an artifact of the measurement system instead of actual engagement. Nevertheless, all measurement quantiles seen here have a significant correlation to gamer engagement both above and below the peak engagement.

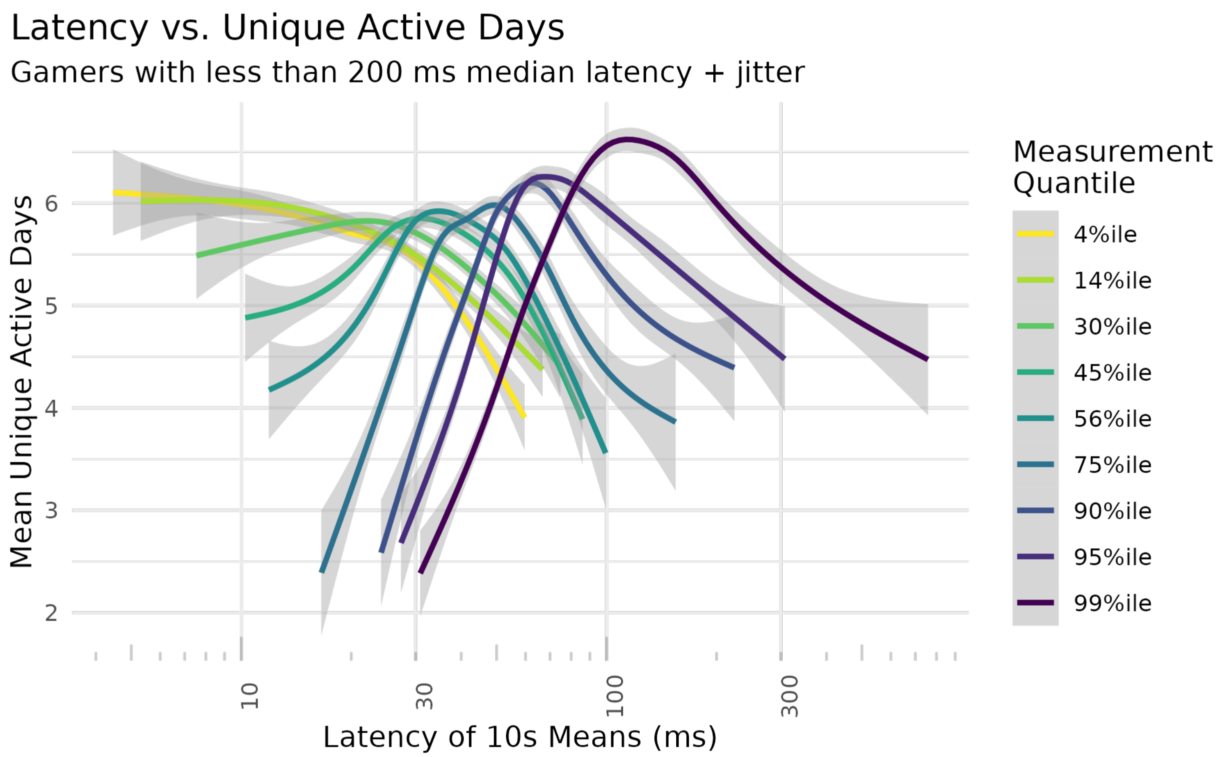


**Figure 2 – Latency vs. Unique Active Days by Measurement Quantile**

Following the regression analysis methodology outlined above, we estimated three different regressions to better understand different parts of the latency-engagement relationship. In all of the regressions, unique number of active days is the y-variable. Our x-variable includes three different latency calcuations. The first is latency in 10ms increments, the second is median latency, and the third is user z-scores. The results of these three regressions are summarized below in Figure 3, where each panel is a different regression. Each dot represents a estimate coefficient in the regression on the latency variable. Note that the negative correlations in Figure 3 are inverselty related to engagement: the more negative a coefficient is, the more engagement there is. The primary findings of these regressions is the following.

1. The positive slope in the top panel is misleading. While we do find that improvements in lower latency quantiles is larger, the next two regressions add perspective. That is, 99%tile latency does matter. However, the key result is intuitive. Median latency, for example, is experienced much more frequently by a gamer than 99%tile latency.
2. When engagement is measured relatively for each user, where their median value for each percentile is used, we find the reduction in engagement from worse latency measurements across

 

percentiles flat until the 99%tile. This suggests the importance of higher percentiles is more important than the first panel suggests.

3. Finally, a z-score is a useful way to summarize how extreme a value is relative to a mean using the distributions standard deviation. In the third panel, the results suggest that extremely large latencies at all percentiles similarly impact engagement, even out to the 99%tile.

Based on these results, we will focus on 99%ile measurements for the remainder of the paper.

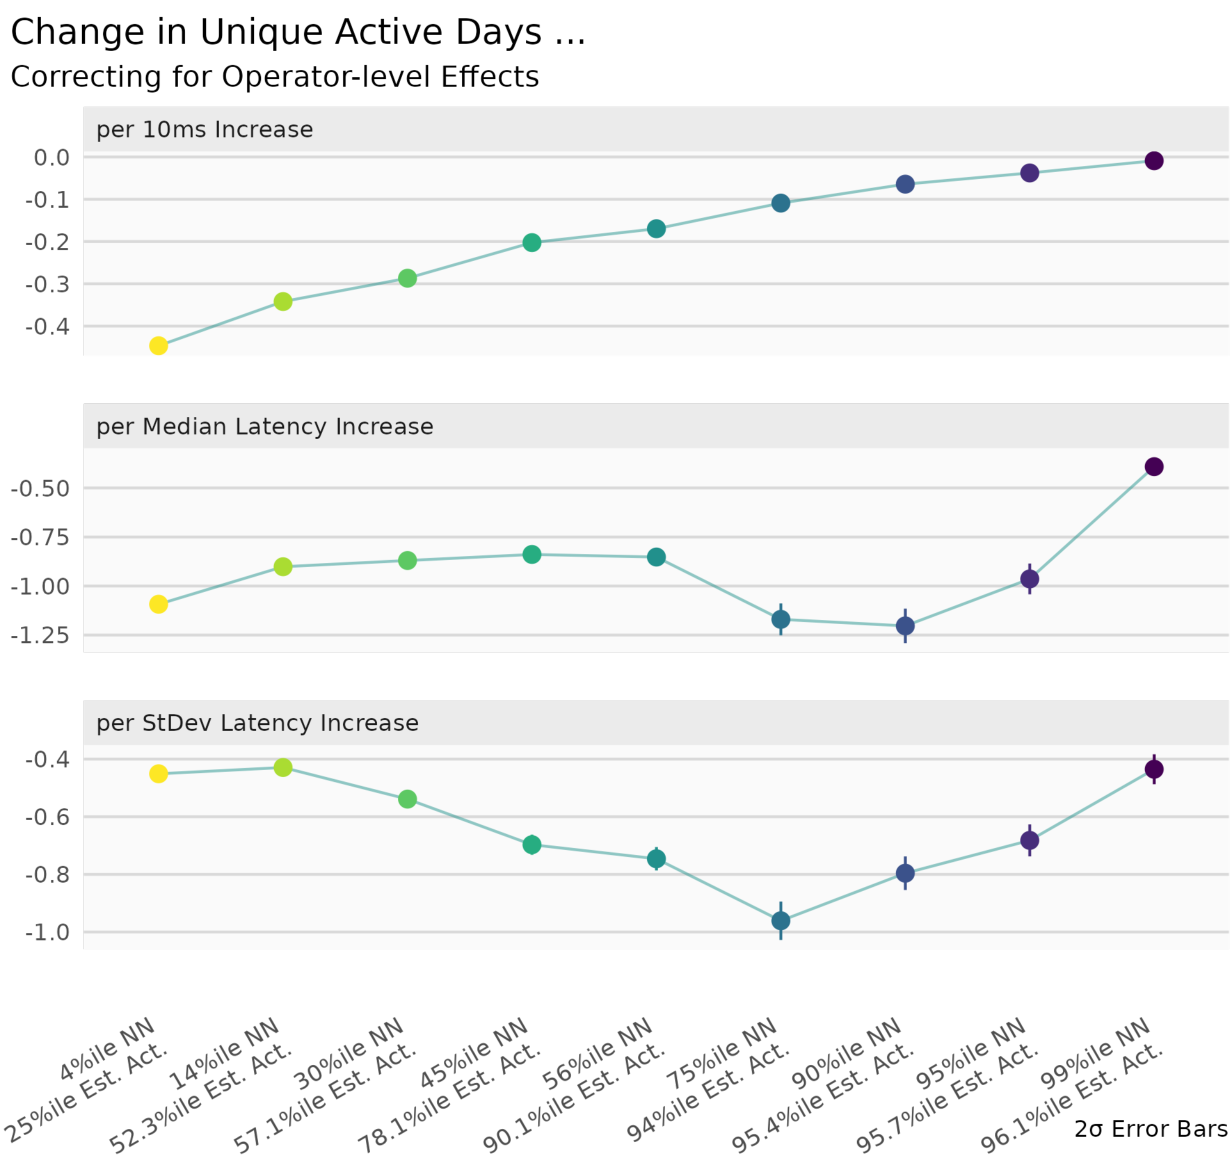


**Figure 3 – Correlation of Latency by Quantile to Unique Active Days**

## 4. Wired Latency

This section begins by quantifying the latency distribution of wired gamers and separating that latency into routing latency and 99%ile queuing latency. Next, improvements to 99%ile queuing latency are examined and with estimates of what improvements to wired latency might be expected from a full-scale deployment of Low Latency DOCSIS technologies.

### 4.1. Overall Wired Latency

This section uses PS4, Xbox, and Windows players with a wired connection to their home gateway.

Total 99%ile latency L is separated into a routing delay (R) and a queuing delay (Q).

$$R_{99\%ile} + Q_{99\%ile} = L_{99\%ile}$$

The 99%ile queuing delay (Q) is the 99%ile latency (L) minus an estimate of routing delay (R).

$$Q_{99\%ile} = Quantile_{99\%}(L - R)$$

As depicted in Figure 4, with one exception, cable and telco latencies were generally comparable during May & June of 2021. 99%ile latency generally ranges from 30 to 200ms for cable operators with a somewhat wider range for most telcos.

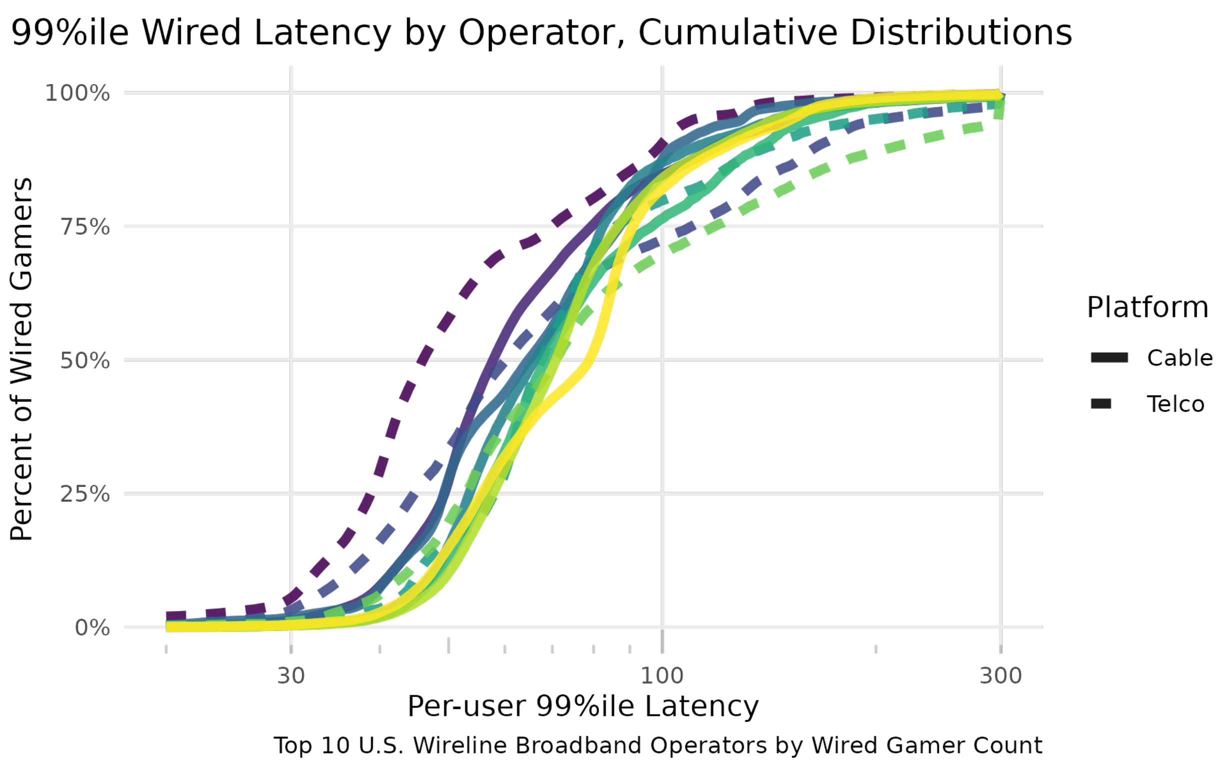


**Figure 4 – 99%ile Wired Latency by Operator**

Figure 5 is a map of latency within the lower 48 states of the U.S. It shows 99%ile wired latency by location. The sample is biased away from dense urban centers to avoid overplotting, so it is not representative. The line from each user to a game server represents the game server most often used by that gamer while each gamer's icon is colored according to his or her latency.

Two things stand out on this map:

 

1. Some gamers primarily use a server other than their nearest one. In fact, outside of a few states such as Illinois, Michigan, Georgia, and Florida this seems to be the norm rather than the exception. This result suggests that game matching algorithms are using more criteria than just the closest server location (at least, sometimes as shown by this game). *(Note: Lines that appear to go from one server location to another such as from Chicago to Omaha are an illusion and are virtually passing through "Chicagoland" from users in Northern New England.)*
2. The worst latency performers are generally telco customers in rural locations (likely DSL) such as the Mississippi Delta, east Texas, eastern Arizona, and central Nevada.

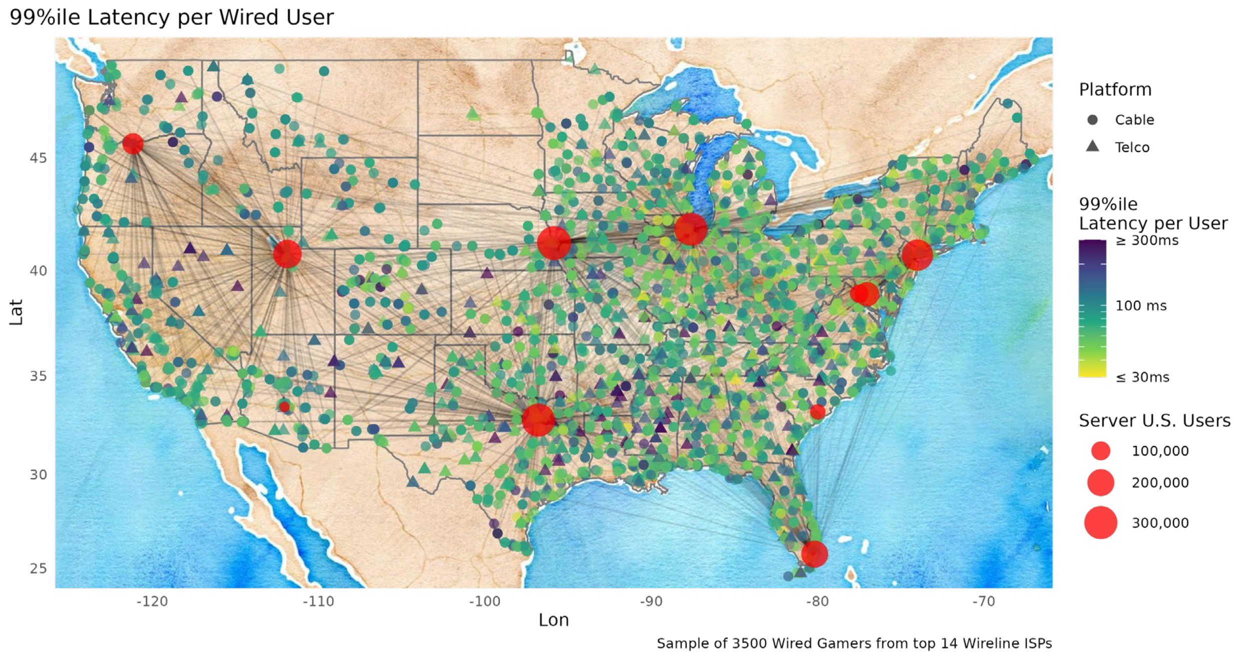


**Figure 5 – 99%ile Wired Latency U.S. Map**

## 4.2. Routing Delay

This section is presented to help readers compare the size and shape of routing delay distributions to the wired queuing delay distribution in Figure 8, the Low Latency DOCSIS counterfactual in Figure 10, and the Wi-Fi queuing delay in Figure 12 & Figure 13.

In general, routing delay ranges from 10 to 100ms in the lower 48 states of the U.S., and as the map in Figure 5 implies, is primarily determined by geographic proximity to the relevant game server for that operator's user population.

 

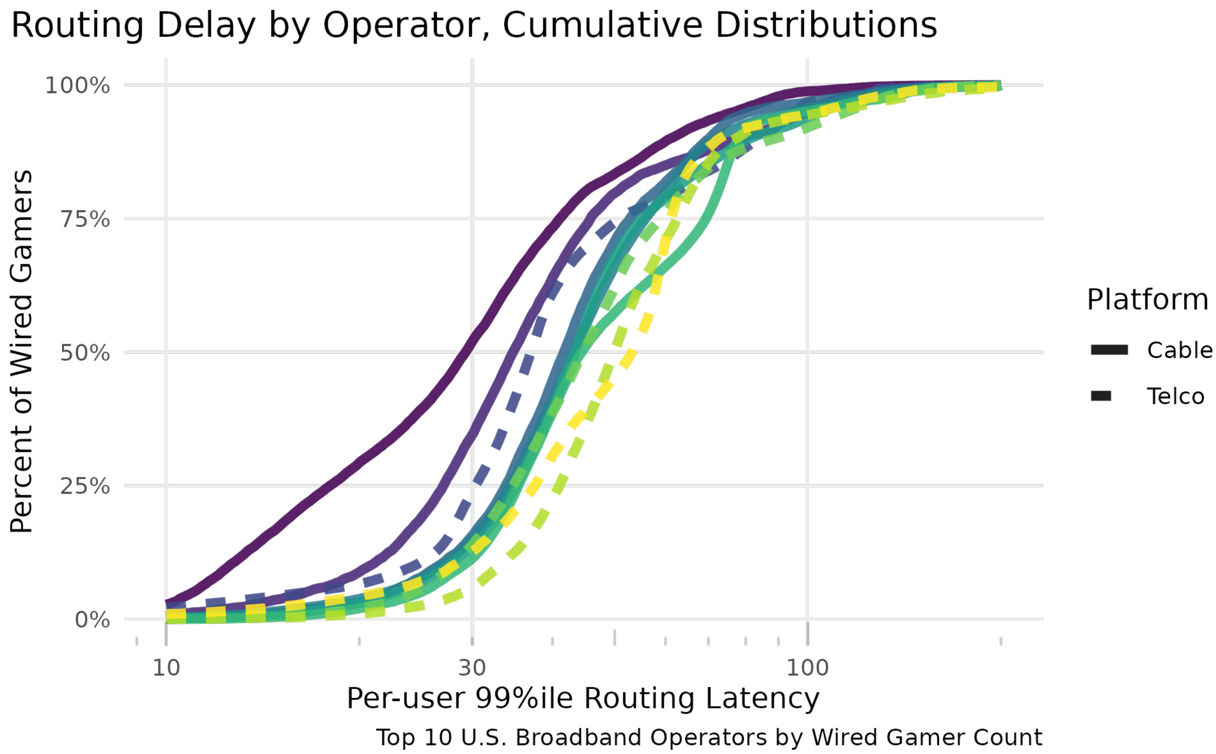


**Figure 6 – Routing Delay by Operator**

### 4.3. Queueing Delay

While routing structures are unique to each cable operator, the access network and the queuing delays that happen there are more a function of DOCSIS technology. Therefore, queuing delay and the potential impact of Low Latency DOCSIS technology as a way to improve queuing delays is of high interest. This section will examine how much of an impact wired network queuing delays (assumed to be primarily access network queuing delays) typically have and then the impact Low Latency DOCSIS could have on them.

To give an idea of the relative scale of routing and queuing delay for different operators: For cable operators, 25% of gamers in the study had at least 45% of their total 99%ile latency come from transient, assumably queuing-based delays. For telcos, the top 25% starts at 34% of 99%ile latency coming from queuing. That is, for these gamers, queuing delay is a larger share of overall latency on cable networks compared to telco. The results are summarized on a per operator basis in Figure 7.

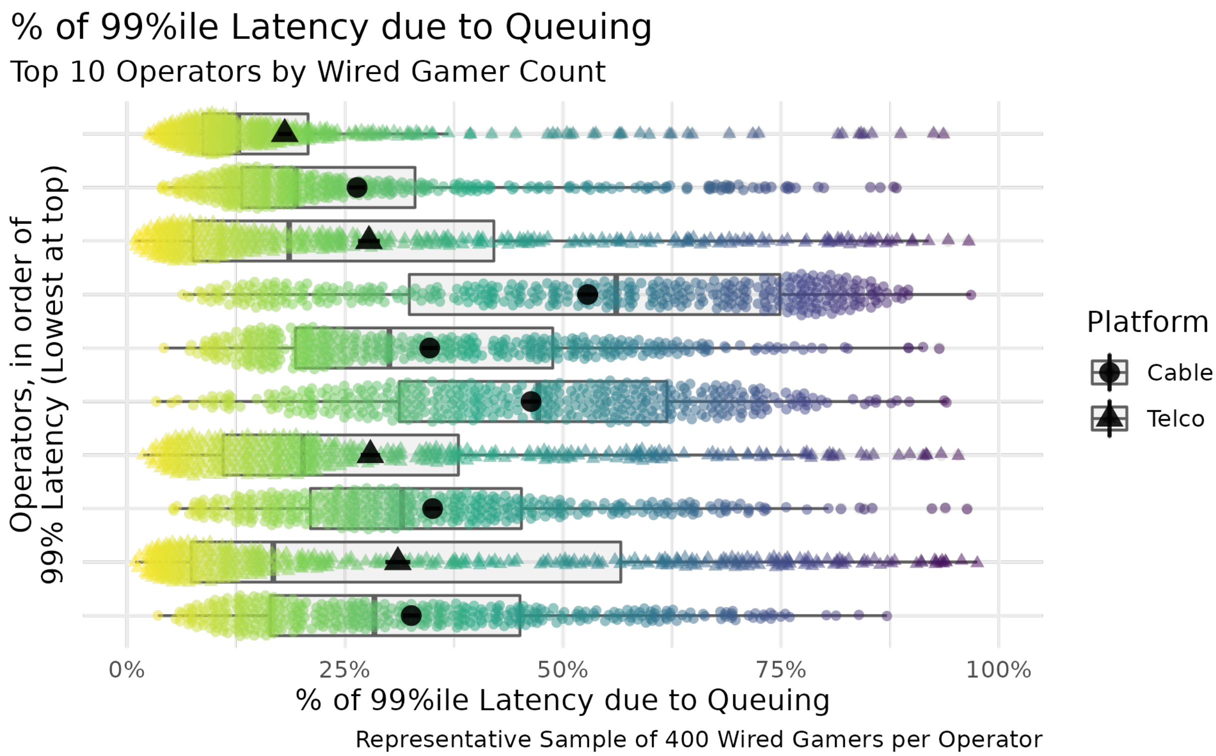


**Figure 7 - % of Latency due to Queueing**

Figure 8 shows the cumulative distributions of queuing delay for each of the top 10 U.S. wireline Internet service providers and provides some guidance on what Low Latency DOCSIS® (LLD) should be able to provide. The queuing delay experienced by users typically ranges from 2ms to more than 200ms. In general, the best telco customers experience less than 3ms queuing delay while the best cable customers experience less than 11ms. However, with one exception, all the telcos fall behind the cable curves in the top quartile. The 6 largest cable broadband operators fall into three distinct groups with a singular queuing delay leader, a close pack of three in the middle, and then two others.

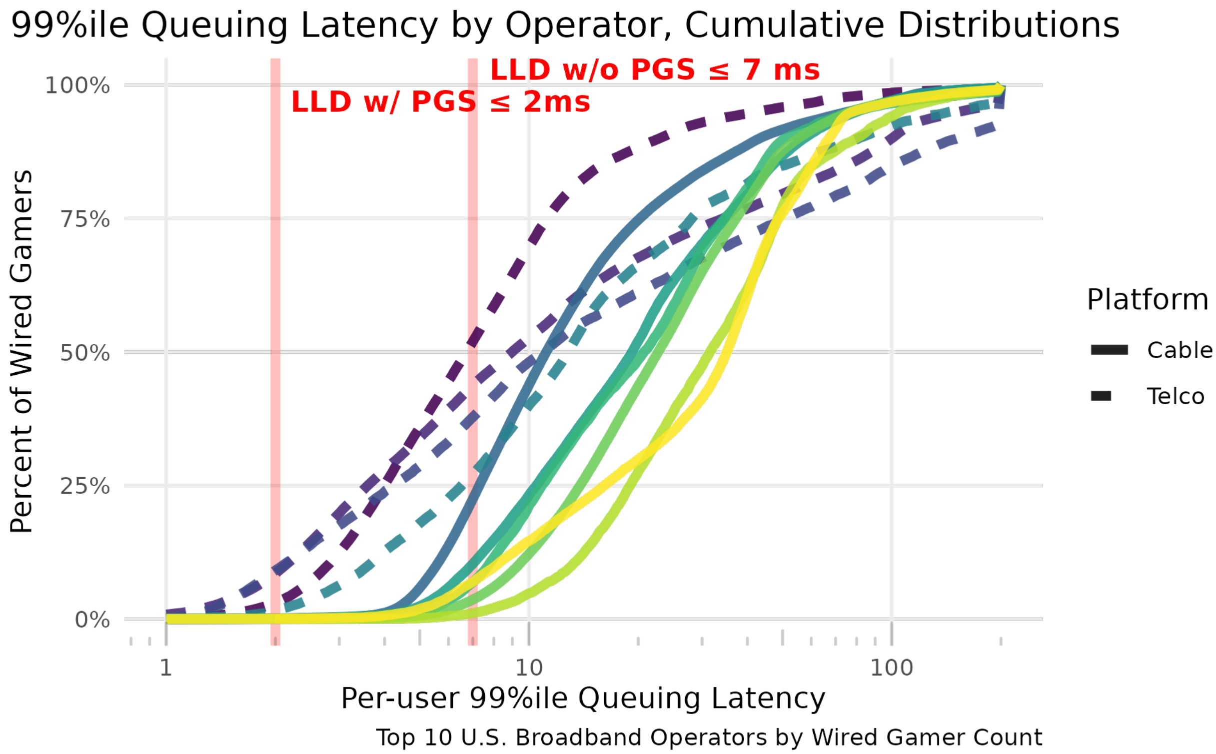


**Figure 8 – Queuing Delay Dist. by Operator**

The expectation of Low Latency DOCSIS is that it should make cable operators queuing delay much more competitive to telcos. Without the Proactive Grant Service (PGS), LLD should be able to deliver queuing delays reliably 7ms or below. With PGS, that falls to 2ms but at the cost of allocating some upstream bandwidth to the proactive grants. This is reflected in the annotations in Figure 8.

Taking that a step further, Figure 9 is a *simplistic* imagining of what might happen if the top 6 cable broadband providers in the U.S. enabled LLD across their entire installed base. Suddenly the queuing delay cumulative distributions would appear as they do in Figure 9, resulting in an average 38% improvement in 99%ile latency for cable broadband customers.

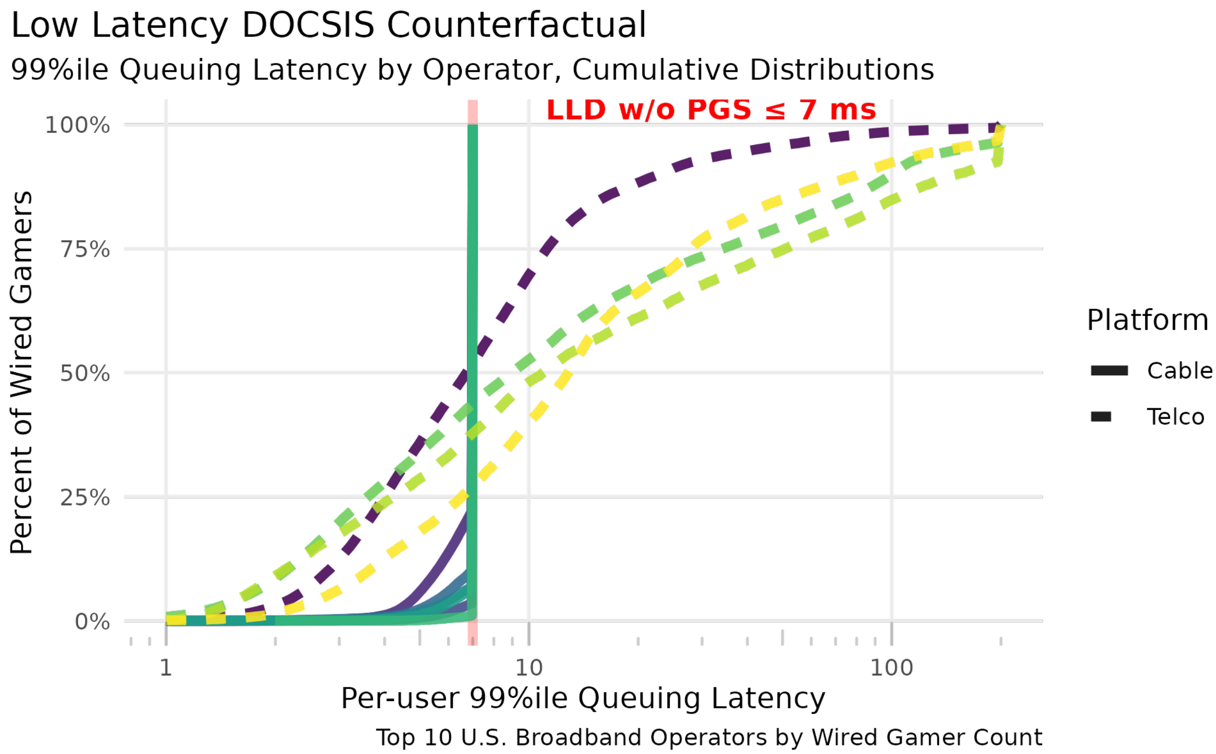


**Figure 9 – Queuing Delay Dist – LLD Counterfactual**

Finally, Figure 10 estimates what the overall 99%ile latency curves would look like for the top 6 cable operators if LLD were fully enabled on each of their networks. As you can see, the results vary wildly with the current queuing delay leader (top left) making relatively modest gains but the 99%ile queuing delay of other operators improving by 50ms or more for a significant fraction of their customers. However, the improved consistency in latency from lower jitter is understated by this figure. On today's network, latency + jitter can exceed 100ms with too much regularity to support the requirements of future real-time applications and user experiences.

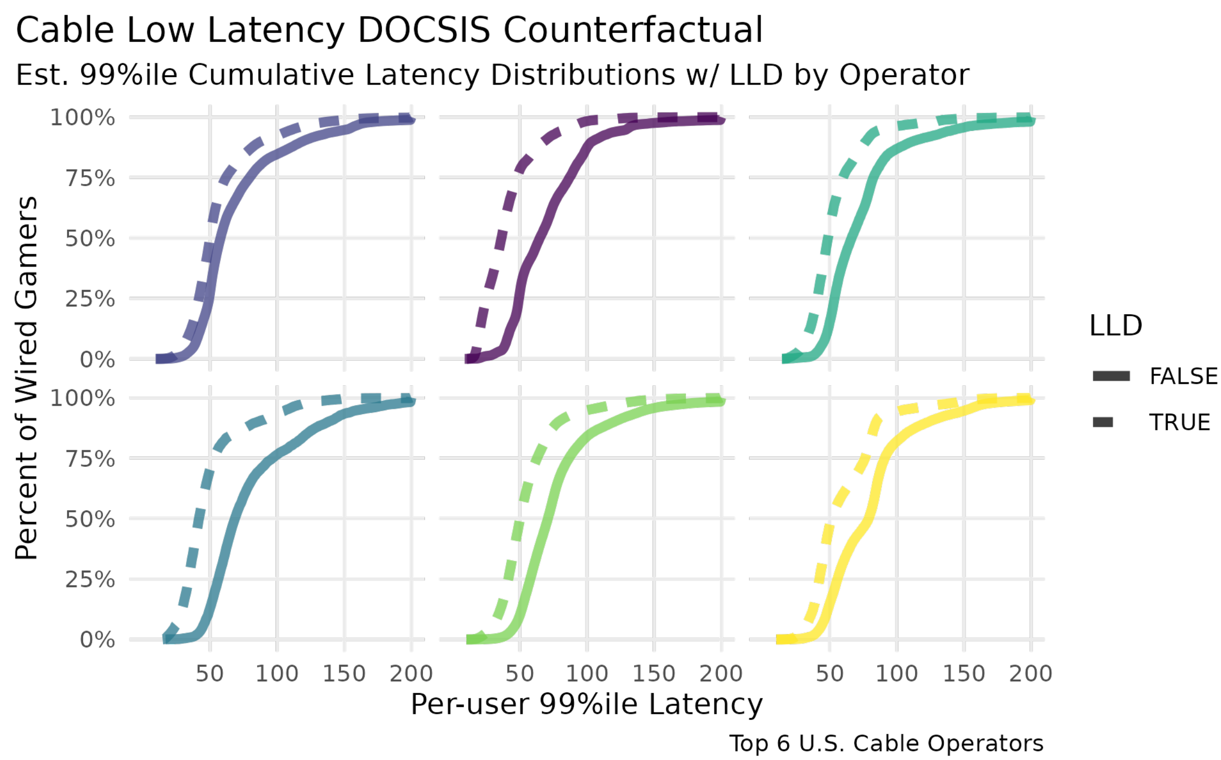


**Figure 10 – Est. 99%ile Latency w/ Full LLD Deployment**

## 5. Wi-Fi Queueing Delays

### 5.1. Wi-Fi Share

For this game, 68% of players connect via Wi-Fi and 32% on wired. A full breakout of gaming platforms can be found below in Figure 11. Xbox One is the most popular Wi-Fi platform, followed by the PS4, then the Switch, and last is Windows. Wi-Fi delays can be larger than all wired delays combined. Therefore, improving Wi-Fi must be considered in any end-to-end network latency plans.

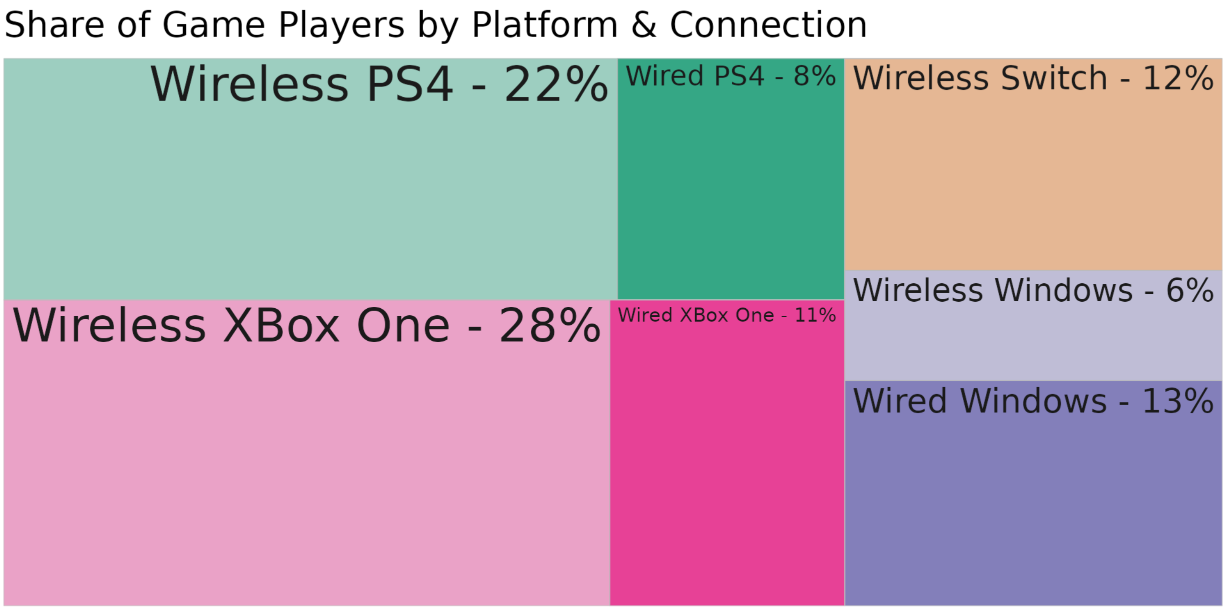


**Figure 11 – Wi-Fi Share**

## 5.2. Wi-Fi Delay

For the 42,252 Wi-Fi gamers where we could find a wired match with at least one hour of overlap, their 99%ile Wi-Fi delay ranges all the way from 4ms to near 300ms. Most operators are grouped together, but two telco operators stand out for the worst 45% of their matched gamers. Each operator in this plot has at least 400 matched gamers.

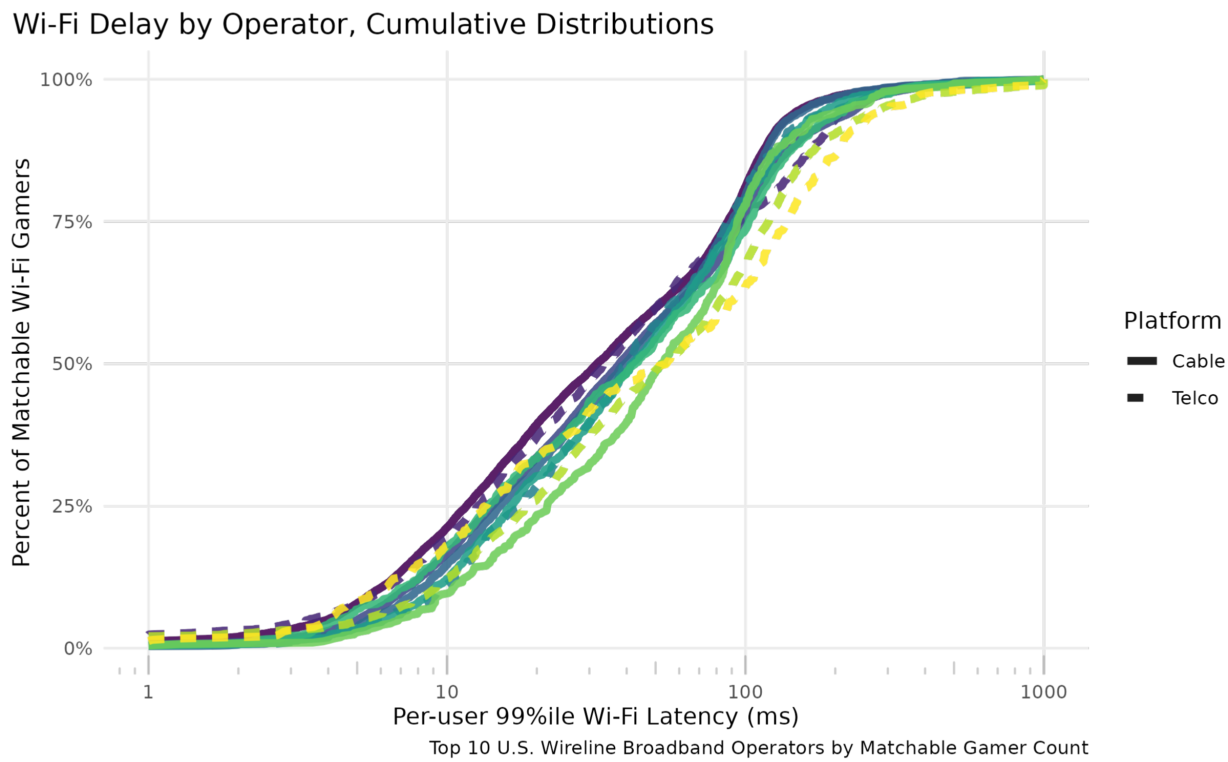


**Figure 12 – Wi-Fi Delay by Operator**

### 5.3. Performance by Operator & Gaming Platform + WMM Packet Marking

One unique feature of this data set is that Network Next's Wi-Fi stack included WMM packet marking on the Nintendo Switch and Windows Wi-Fi clients while the network stacks on the Xbox One and the PS4 did not. Figure 13 provides insights into the impact of this and other differences in the Wi-Fi implementation on different devices on 99%ile Wi-Fi latency performance. This gives us a chance to compare the performance of each of these network stacks as they interact with the home gateways used by customers of different Internet service providers.

In general, the Nintendo Switch has the best Wi-Fi performance, followed by MS Windows clients, followed by PS4 clients, with Xbox having the worst performance.

Other things to note:

- For T1, T2, and C4, it appears that there are a few gateways that have noticeably worse performance with the Switch than other gateways used by customers of the same operator. (Step function in the middle of the figure.)
- In general cable operators are outperforming the telcos in Wi-Fi delay, but not by much.

 

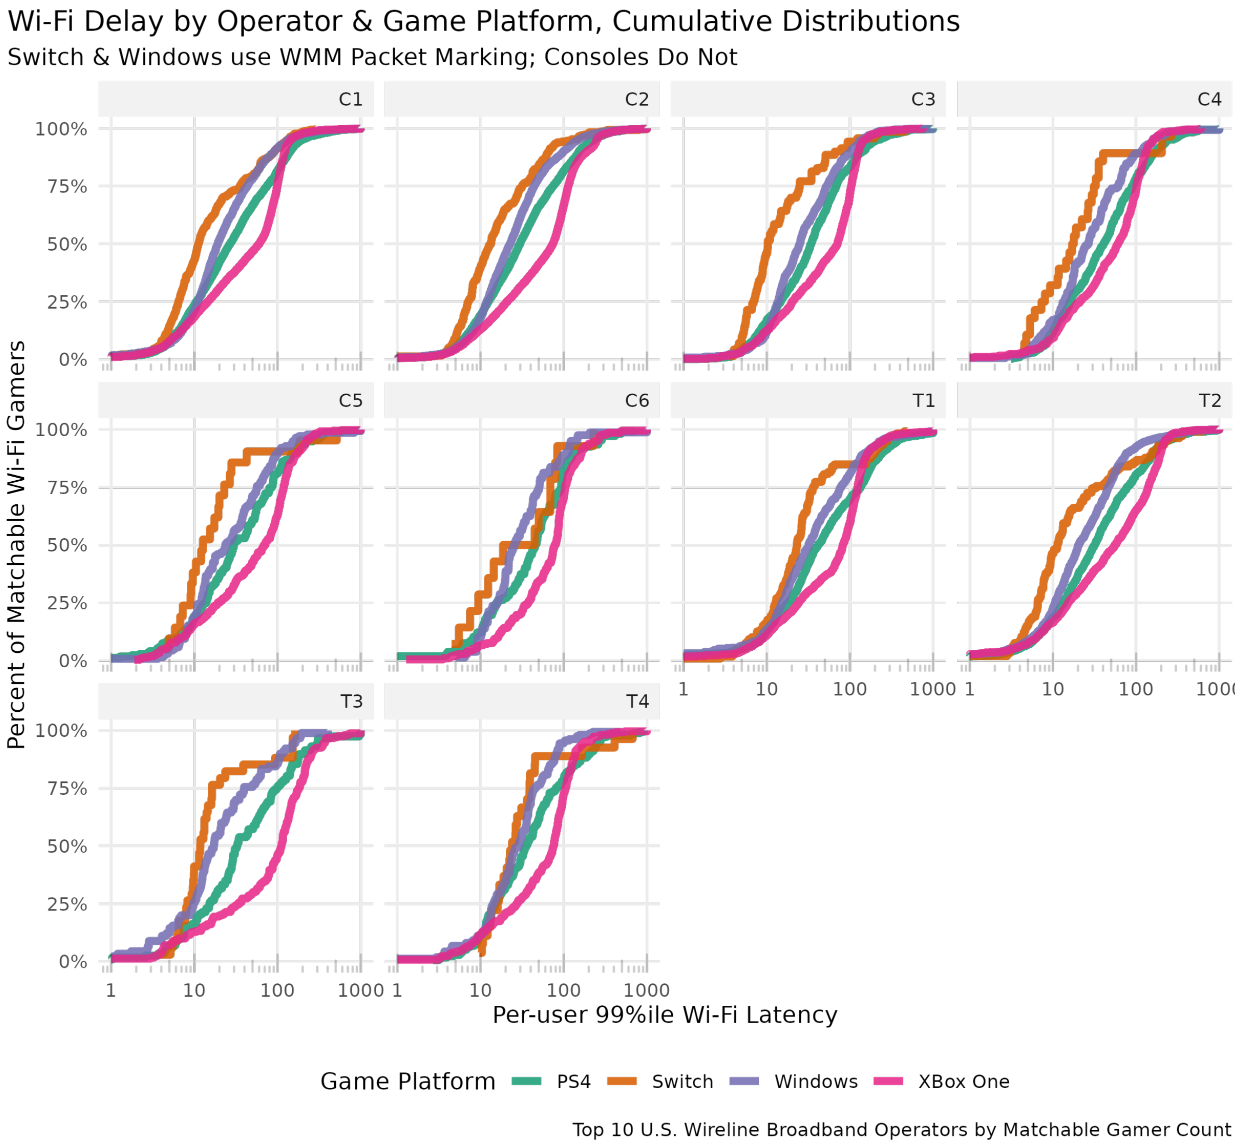


**Figure 13 – Wi-Fi Performance by Operator and Gaming Platform**

Finally, referring to Table 2, we run a dummy variable regression of the top 10 ISPs interacting with the four gaming platforms against log of latency. The base case in this dummy variable regression is ISP C1 + PS4.

While the other measurement quantiles are presented in Table 2, the 99%ile alone will be discussed:

- The geometric mean intercept is 28ms, and each ISP except C5 results in a statistically significant increase of latency over C1 (who is the Wi-Fi latency leader) ranging from a small 11% increase for C2 to a 61% increase for T1.
- The Nintendo Switch has 46% lower latency than the PS4 and Windows Wi-Fi has a 23% improvement, meaning that WMM packet marking has a significant and large effect on Wi-Fi latency compared to the unmarked PS4 packets (and even more so compared to the Xbox One). We note, however, that WMM works well on packets like those studied in this sample (i.e.,

 

gaming), but may not be sufficient for other types of traffic (e.g., real-time video). Therefore, other technologies like L4S AQM implemented over Wi-Fi are things to be considered.

- The Xbox One Wi-Fi implementation results in a 34% worse latency over the PS4
- ISP C2 has a unique issue with the Xbox One (it does quite well otherwise), seeing a 21% increase in latency compared to other operators serving the Xbox One. ISP T3 seems to have a similar problem.

**Table 2 - Regression of Wi-Fi log(latency) vs. ISP & Gaming Platform**

| | 4%ile | 14%ile | 30%ile | 45%ile | 56%ile | 75%ile | 90%ile | 95%ile | 99%ile | Max |
|---|---|---|---|---|---|---|---|---|---|---|
| Geometric Mean (Intercept - ISP C1 + PS4) | 1.43 ms *** | 1.83 ms *** | 2.41 ms *** | 3.03 ms *** | 3.6 ms *** | 5.16 ms *** | 8.5 ms *** | 12.4 ms *** | 27.71 ms *** | 106.71 ms *** |
| ISP C2 | -17% *** | -21% *** | -20% *** | -17% *** | -13% *** | -6% ** | 3% | 5% | 11% *** | 6% * |
| ISP C3 | 1% | 7% * | 11% ** | 14% *** | 16% *** | 19% *** | 22% *** | 25% *** | 20% *** | 0% |
| ISP C4 | 16% *** | 24% *** | 26% *** | 30% *** | 33% *** | 40% *** | 47% *** | 47% *** | 43% *** | 21% ** |
| ISP C5 | -13% ** | -14% ** | -11% | -8% | -3% | 2% | 10% | 10% | 14% | 1% |
| ISP C6 | -1% | 7% | 9% | 7% | 8% | 14% | 25% * | 31% * | 40% ** | 32% * |
| ISP T1 | 26% *** | 33% *** | 35% *** | 36% *** | 37% *** | 38% *** | 45% *** | 49% *** | 61% *** | 51% *** |
| ISP T2 | 13% *** | 9% *** | 6% * | 5% | 5% | 5% | 5% | 5% | 13% ** | 16% ** |
| ISP T3 | 1% | 10% | 13% | 9% | 8% | 5% | 10% | 14% | 37% ** | 73% *** |
| ISP T4 | 1% | 6% | 18% ** | 20% ** | 20% ** | 23% ** | 38% *** | 36% *** | 36% ** | 31% ** |
| Switch | -3% | -3% | -6% | -8% | -10% * | -15% ** | -25% *** | -35% *** | -46% *** | -48% *** |
| Windows | -11% *** | -16% *** | -19% *** | -21% *** | -23% *** | -24% *** | -27% *** | -29% *** | -23% *** | -31% *** |
| XBox One | -9% *** | -13% *** | -16% *** | -18% *** | -19% *** | -21% *** | -22% *** | -20% *** | 34% *** | 0% |
| ISP C2 + Switch | -1% | -3% | 5% | 4% | 2% | -4% | -7% | -3% | -4% | -6% |
| ISP C3 + Switch | -16% * | -21% * | -22% * | -25% * | -25% * | -29% * | -27% * | -26% | -17% | -16% |
| ISP C4 + Switch | -16% | -27% * | -27% | -30% * | -32% * | -34% * | -32% | -21% | -13% | -13% |
| ISP C5 + Switch | 31% * | 29% | 43% | 40% | 30% | 23% | 9% | 4% | -1% | -17% |
| ISP C6 + Switch | 16% | 18% | 32% | 35% | 32% | 28% | 18% | 23% | 28% | -26% |
| ISP T1 + Switch | -12% * | -10% | -8% | -11% | -12% | -11% | -1% | 9% | 9% | -9% |
| ISP T2 + Switch | -22% *** | -27% *** | -24% ** | -21% * | -19% * | -16% | -11% | -3% | -4% | 11% |
| ISP T3 + Switch | 3% | -4% | -3% | 1% | 1% | -3% | -17% | -18% | -23% | -51% ** |
| ISP T4 + Switch | 22% | 18% | 4% | 10% | 14% | 12% | 12% | 19% | 46% | 48% |
| ISP C2 + Windows | 7% ** | 10% *** | 9% ** | 8% * | 7% * | 5% | 4% | 5% | 1% | 8% |
| ISP C3 + Windows | -5% | -8% | -8% | -6% | -4% | 0% | 4% | 6% | 5% | 15% |
| ISP C4 + Windows | -13% * | -19% ** | -24% *** | -29% *** | -30% *** | -30% *** | -27% ** | -24% * | -10% | 28% |
| ISP C5 + Windows | 6% | 7% | 5% | 2% | 0% | 1% | 0% | 8% | 9% | 8% |
| ISP C6 + Windows | 0% | -6% | -5% | 1% | 0% | 2% | 4% | -3% | -1% | -13% |
| ISP T1 + Windows | -3% | -1% | 4% | 7% | 8% | 8% | 11% | 13% | -4% | -4% |
| ISP T2 + Windows | -4% | -4% | -8% | -10% * | -11% * | -10% | 3% | 3% | -11% | -11% |
| ISP T3 + Windows | 1% | -11% | -18% | -20% | -23% | -27% * | -25% | -16% | -34% * | -32% * |
| ISP T4 + Windows | 7% | 11% | -1% | -3% | -1% | -2% | -10% | -8% | -9% | -13% |
| ISP C2 + XBox One | 8% *** | 10% *** | 12% *** | 13% *** | 13% *** | 15% *** | 20% *** | 25% *** | 21% *** | 32% *** |
| ISP C3 + XBox One | -3% | -5% | -5% | -6% | -5% | -3% | -1% | 0% | 2% | 14% |
| ISP C4 + XBox One | -11% ** | -14% ** | -15% ** | -15% ** | -17% ** | -20% ** | -18% ** | -11% | -16% * | -11% |
| ISP C5 + XBox One | 9% | 11% | 11% | 12% | 10% | 18% | 19% | 20% | 12% | 35% * |
| ISP C6 + XBox One | -1% | 0% | 2% | 9% | 11% | 16% | 27% | 32% * | 16% | 3% |
| ISP T1 + XBox One | -1% | 1% | 1% | -1% | -2% | -2% | 1% | 2% | -8% | -2% |
| ISP T2 + XBox One | 1% | 3% | 3% | 4% | 4% | 8% | 13% * | 12% * | 0% | 12% |
| ISP T3 + XBox One | 0% | -10% | -12% | -5% | 1% | 18% | 39% ** | 69% *** | 40% * | 9% |
| ISP T4 + XBox One | 21% *** | 24% ** | 17% * | 18% | 19% | 21% | 16% | 18% | 0% | 10% |
| **Adj. $r^2$** | **4%** | **4.10%** | **3.50%** | **2.90%** | **2.60%** | **2.20%** | **2.30%** | **2.60%** | **6%** | **3.40%** |

As is typical * = p < 5%, ** = p < 1%, *** = p < 0.1%.

## 6. Conclusion

There are several opportunities for cable operators to make more enjoyable gaming experiences for their customers by reducing latency.

The first priority is to be able to understand how to measure latency in a way that correlates well to a gamer's enjoyment of this game, and the takeaway is that all of the measurement quantiles up to and including the 99%ile are well correlated to users' engagement in this game with an across-the-board one standard deviation improvement in latency likely resulting in between 0.4 and 1 (7% - 18%) more unique active days played of this game.

The second priority and biggest opportunity for differentiation compared to telcos is Wi-Fi. While there is hope of substantial improvements with Wi-Fi 6E, which will feature new (greenfield) 6 GHz channels where all the clients and base stations will include major latency-improving features such as OFDMA. However, the state of the game console cycle combined with the fact that the honeymoon period where the 6 GHz channels are unpopulated will come to an end leads us to emphasize that improvements here will require full ecosystem collaboration. For most cable operators, broader WMM adoption could result in halving the Wi-Fi lag for the Xbox One and the PS4.

Low Latency DOCSIS will have a significant effect on wired queuing delays measured at the 99%ile with some operators able to save many of their customers up to 50ms 99%ile round trip time.

Finally, routing latency should not be neglected. While it is beyond the scope of this research, there are indications in the data that one telco has been able to achieve results in routing latency vs. distance-to-server over and above the other ISPs, and there is therefore a real opportunity to bring this number down.

# Acknowledgements

We would like to thank Greg White of CableLabs for his immense assistance modelling the 10-second summarization technique used by Network Next on lab latency data he has and using that model to help us adjust the quantile range we were looking at. This paper would be far harder to understand without his help.

We would also like to thank Glenn Fielder, Tapan Desai, and Andrew Baumbach of Network Next for sharing several of their current hypotheses and their excellent support of this data set.

# Abbreviations

| | |
|---|---|
| DOCSIS® | Data Over Cable Service Interface |
| DSL | digital subscriber line |
| LLD | low latency DOCSIS |
| ms | millisecond |
| NN | Network Next |
| %ile | percentile |
| PGS | proactive grant service |
| PS4 | PlayStation 4 |
| rms | root mean square |

 

| | |
|---|---|
| StDev | standard deviation |
| WMM | Wi-Fi multimedia |

# Bibliography & References

Broadband Internet Technical Advisory Group. (2022, 01 10). *Latency Explained*. Retrieved from BITAG: https://bitag.org/latency-explained.php

Dowle, M., & Srinivasan, A. (2021). data.table: Extension of `data.frame`. Retrieved from https://CRAN.R-project.org/package=data.table

Network Next. (2020, 10 22). *Introducing the Network Next Analytics Portal*. Retrieved from Network Next: https://www.networknext.com/post/introducing-the-network-next-analytics-portal

R Core Team. (2022). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from https://www.R-project.org/

Subspace. (2021, 05 03). *How to Ensure Your Matchmaking Experience for Multiplayer Games is Fast and Full*. Retrieved from Subspace: https://subspace.com/resources/multiplayer-matchmaking

# Appendices

## Appendix 1 – Regression of Unique Active Days vs. Latency & Operator

Table 3 includes a summary of the full regression coefficients and significance levels for each measurement quantile. (The regression is run separately at each measurement quantile.) The highest $r^2$ occurs at the 75%ile, but the strongest effect for improvement (such as by deployment of Low Latency DOCSIS®) occurs in the 99%ile measurements, so that is what is discussed here.

The inclusion of operator effects (that some operators have more engaged gamers than others) was necessary to get a clean correlation for the main variable.

Notably:

- The average player on operator C1 plays 5.1 unique active days of this game.
- For every standard deviation of increase in 99%ile latency, a gamer plays 0.4 unique days less.
- ISPs C7 and T6 both have significantly higher engagement than the base case operator (C1), while T3, T5, and possibly C2 and T4 have lower base case engagement.

 

**Table 3 - Regression Output of Unique Active Days per user vs. Z-score of Latency & Operator by Measurement Quantile**

| | 4%ile | 14%ile | 30%ile | 45%ile | 56%ile | 75%ile | 90%ile | 95%ile | 99%ile | Max |
|---|---|---|---|---|---|---|---|---|---|---|
| Mean Unique Active Days (Intercept) | 5.6 *** | 5.6 *** | 5.7 *** | 5.1 *** | 5.1 *** | 4.6 *** | 4.4 *** | 4.6 *** | 5.1 *** | 8.5 *** |
| **Per StDev Increase in Latency** | **-0.5** *** | **-0.4** *** | **-0.5** *** | **-0.7** *** | **-0.7** *** | **-1** *** | **-0.8** *** | **-0.7** *** | **-0.4** *** | **0.4** *** |
| ISP C2 | 0.3 *** | 0.1 *** | 0.1 *** | 0.4 *** | 0.3 *** | -0.4 *** | -0.1 | -0.2 * | -0.2 * | -0.3 *** |
| ISP C3 | 0 | 0.1 *** | -0.1 * | 0.1 | 0.1 | -0.1 | -0.2 | -0.3 ** | -0.1 | -0.4 ** |
| ISP C4 | -0.5 *** | 0.3 *** | -0.1 * | 0.9 ** | 0.8 * | 0.4 * | 0.5 * | 0.3 | -0.1 | -0.4 |
| ISP C5 | -0.3 *** | -0.2 ** | -0.6 *** | -0.2 | -0.3 ** | -0.4 ** | -0.1 | -0.1 | -0.1 | -1.2 *** |
| ISP C6 | -0.4 *** | -0.3 *** | -0.1 | 0.2 | -0.6 *** | -0.6 *** | 0.1 | -0.1 | -0.1 | -1.6 *** |
| ISP C7 | -0.2 *** | -0.3 *** | -0.1 | 0.1 | 0 | -0.1 | 0.1 | -0.1 | 0.7 *** | 1.1 * |
| ISP C8 | 0 | -0.1 | -0.3 ** | 0 | 0.3 ** | -0.2 | 0.1 | 0.1 | 0.1 | -1 *** |
| ISP T1 | -0.1 * | -0.1 *** | -0.2 *** | 0 | 0.1 | -0.2 * | 0.1 | 0.1 | 0.1 | -1.4 *** |
| ISP T2 | -0.4 *** | -0.4 *** | -0.9 *** | -0.5 * | -0.4 | -0.7 ** | -0.3 | -0.4 * | -0.1 | 0.1 |
| ISP T3 | -0.3 *** | -0.2 *** | -0.4 *** | -0.3 *** | -0.3 ** | -0.5 *** | -0.1 | -0.1 | -0.3 *** | -1.7 *** |
| ISP T4 | -0.6 *** | -0.2 * | -0.5 *** | 0.3 *** | -0.3 * | -0.5 *** | 0 | -0.3 ** | -0.3 * | -1.5 *** |
| ISP T5 | -0.5 *** | -0.7 *** | -0.7 *** | -0.3 * | -0.7 *** | -0.7 *** | -0.2 | -0.3 * | -0.5 *** | -2.5 *** |
| ISP T6 | -0.4 *** | -0.2 | -0.3 * | -0.1 | 0.2 | -0.1 | 0.3 | 0.3 | 0.5 ** | 0 |
| **Adj. $r^2$** | **0.8%** | **0.7%** | **1.2%** | **1.9%** | **2.5%** | **4.9%** | **3.4%** | **2.6%** | **1.1%** | **1.4%** |

As is typical * = p < 5%, ** = p < 1%, *** = p < 0.1%.

 