

Preseem Fixed Wireless Network Report

2022 Edition

Contents

Introduction	3	Throughput	21
Executive Summary	3	Download Throughput	22
Overview	4	Upload Throughput	23
Sidebar: Peak vs. Off-Peak	5	Latency	24
Sidebar: Percentiles	7	Sidebar: What is a Good Latency Value?	25
Subscriber Metrics	9	Access Point Subscriber Capacity	28
Throughput	9	RF Channel Width	29
Active Throughput by Plan	11	Percent of APs by Channel Width	29
Latency	12	Download Rate by Channel Width	30
Sidebar: Where Does Latency Come From?	13	Connected CPE Radios by RF Channel Width	31
Subscriber Usage	15	Subscriber Throughput by RF Channel Width	32
Subscriber Usage by Speed Plan	16	Oversubscription Ratio	33
Speeds Attained by Speed Plan	17	Summary	36
Access Point Metrics	18	Appendix	37
AP Market Share	19		
Connected Subscriber Count	20		

Executive Summary

Preseem's annual Fixed Wireless Network Report uses our extensive data set to provide a unique view into fixed wireless networks across providers and different kinds of equipment.

Key insights from this year's edition include:



6 Mbps When Active

The average fixed wireless subscriber uses around 6 Mbps when active (1Y 17%↑)



Speed Plans

Speed plans over 100 Mbps often sit idle



Data Usage

The average fixed wireless subscriber uses 10.6 GB of data per day, for a total of 329 GB per month (1Y 20%↑)



Oversubscription Rate

Over 59% of access points are less than 3x oversubscribed



Consumer Habits

Many users don't consume more data as their speed plan increases



Improved Latency

Latency has improved year over year, indicating that operators are keeping up with consumer demand

Overview

Welcome to the 2022 edition of the Preseem Fixed Wireless Network Report!

By downloading this report, you've demonstrated your commitment to staying up to date with the latest trends and data in the industry. You'll also be able to compare your network's performance against ISPs worldwide, giving you unique insights that we hope will benefit both your business and your subscribers.

At Preseem, we help ISPs thrive by improving how they measure and optimize the subscriber quality of experience (QoE). As part of this work, we collect detailed metrics on subscribers, equipment, and overall network performance from our growing customer base of hundreds of fixed wireless ISPs around the world.



Like all big data sets, there are possible biases in this data. We've done our best to be agnostic, but this is not a scientific paper that controls for all confounding effects or uses other scientifically rigorous methods. Nevertheless, we believe this presents a solid, real-world view of the ISP industry.

Preseem collects billions of metrics from ISPs every day. This report uses our huge data pool to present an exclusive and in-depth analysis of the fixed wireless industry across service providers and vendors. New to this year's report are insights on subscriber usage by speed plans, speed attained by plans, active throughput by plan, and AP capacity.

The goal of this report is to show the real-world experience of fixed wireless subscribers, networks, and equipment. We hope this report is useful as a way to benchmark your ISP against the wider broadband ecosystem, and that it also helps others understand fixed wireless networks.

What's New

If you want to jump right to the new reports, check out:

- [Active Throughput by Plan](#)
- [Subscriber Usage by Speed Plan](#)
- [Speeds Attained by Speed Plan](#)
- [Access Point Subscriber Capacity](#)

SIDEBAR

Peak Vs. Off-Peak

Most networks exhibit great variation in their load during the day.

Given that performance typically only degrades when the network is busy, simple numbers like the average rate for a day, hour, or minute are essentially useless when trying to measure the overall subscriber experience or network performance.



INTRODUCTION



Many of the metrics presented in this report are taken at “peak time.” There are many simple and unsatisfactory methods to determine peak time (such as approximating “prime time”), but these methods fail to capture variation within those periods.

Preseem’s approach is to calculate the minutes in the day with the highest demand (not just throughput) and use the metrics at these times to show true network performance.

As such, the numbers presented in this document aim to represent the typical subscriber experience when the highest number of subscribers are trying to use the service and performance is at its worst.

SIDEBAR

Percentiles



Throughout this report, we use several statistics to describe the data sets. These include stats like average, maximum, and percentiles. Average and maximum are straightforward, but what's with this percentile stuff?



Average and Max

Average is a simple statistic that we all use every day but which can actually be very deceptive. For example, if you and Jeff Bezos are the only people in the room, then the average person in the room has a net worth of over 60 billion dollars. Sounds good, but you can see how this is misleading.

Similarly, using the maximum value as a way to summarize a data set can also paint a misleading picture. For example, the maximum value of your net worth and Jeff Bezos' is 122 billion dollars.



Percentiles

Percentiles are another tool to summarize a data set, and are particularly useful when simpler statistics are misleading.

Imagine you have the following 11-item data set:

5,100,1,2,2,4,5,6,3,4,2

The average of this data set is 12.18 and the maximum is 100. Neither of those statistics are very useful. As an alternative, consider the 50th percentile (aka the median or P50). To calculate the 50th percentile, we first order all the numbers from smallest to largest to get:

1,2,2,2,3, 4, 4,5,5,6,100
50%

The 50th percentile is the value at which 50% of the numbers in the data set are below and 50% of the values are above. In this simple 11-item data set, we can jump to the sixth element and get 4, which is the 50th percentile, or median value. Similarly, the 80th percentile is the value at which 80% of the data set is below and 20% is above, and so on.

Subscriber Metrics

The metrics in this section present a high-level overview of the fixed wireless subscriber experience across all types of networks and equipment.

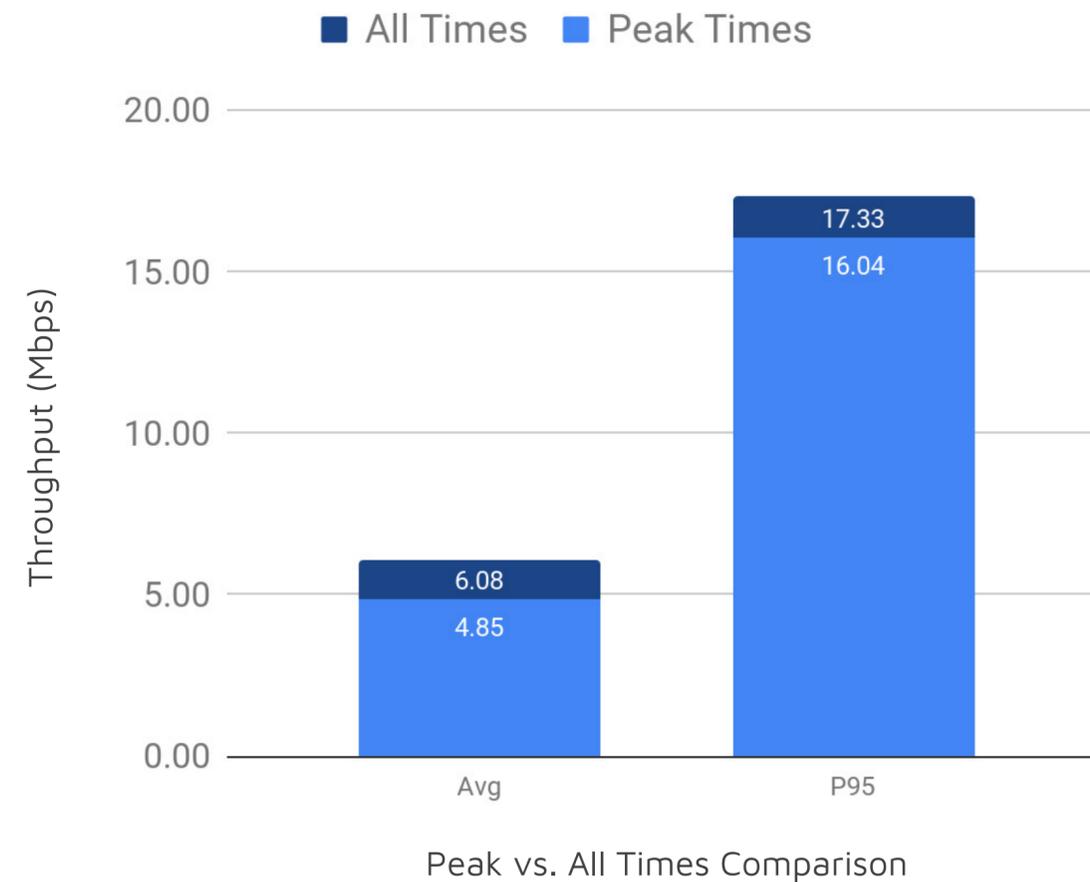
Throughput

This chart compares the download throughput achieved by ISP subscribers during the busiest (peak time) and other times of the day. The difference between peak and off-peak is surprisingly small. This indicates that, on the whole, subscriber throughput does not degrade significantly during busier times. Pat yourselves on the back, ISPs! Subscriber download throughput for all times has risen 17% year over year.



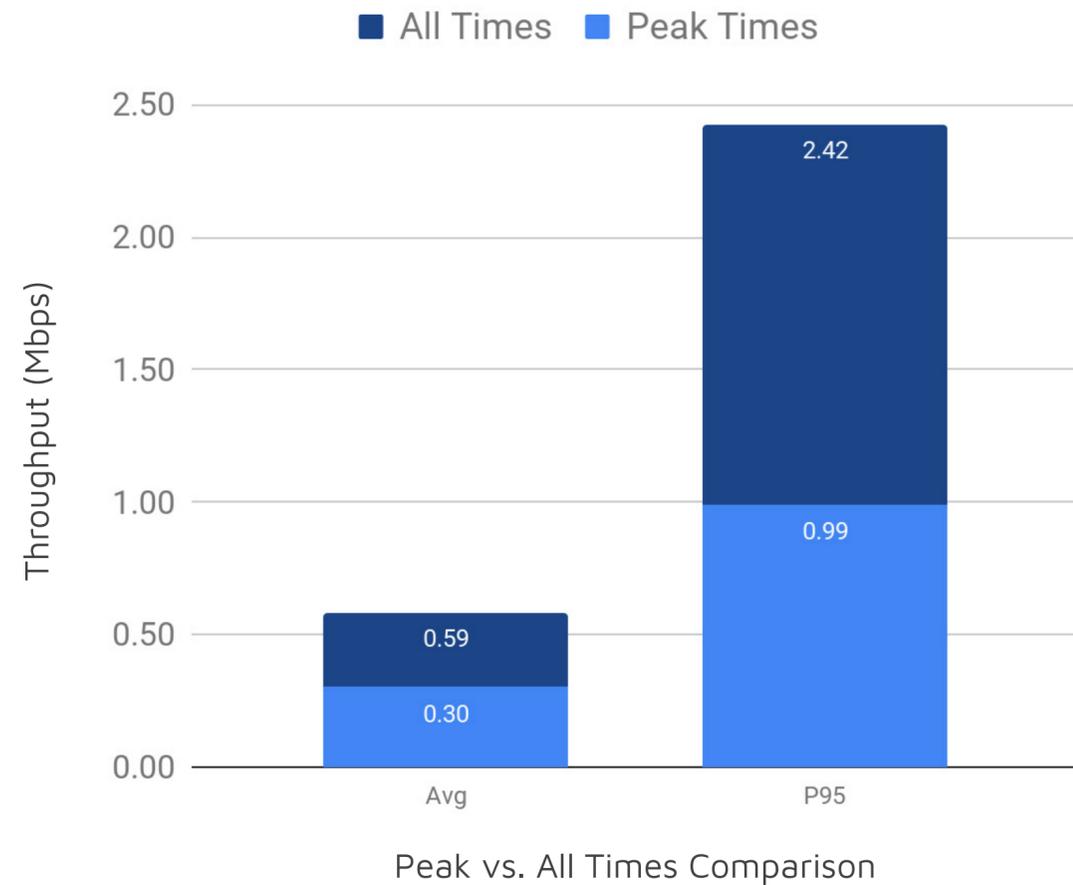
In the Fixed Wireless Network Report, we distinguish between active subscribers (those actively using the Internet) from connected subscribers, which refers to the number of radios or subscriber modules attached to an access point.

Subscriber Download Throughput Peak vs. All Times



The upload throughput numbers are a little more confusing because most subscribers don't stress the upload direction of their connection.

Subscriber Upload Throughput Peak vs. All Times



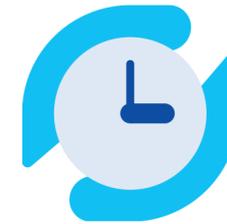
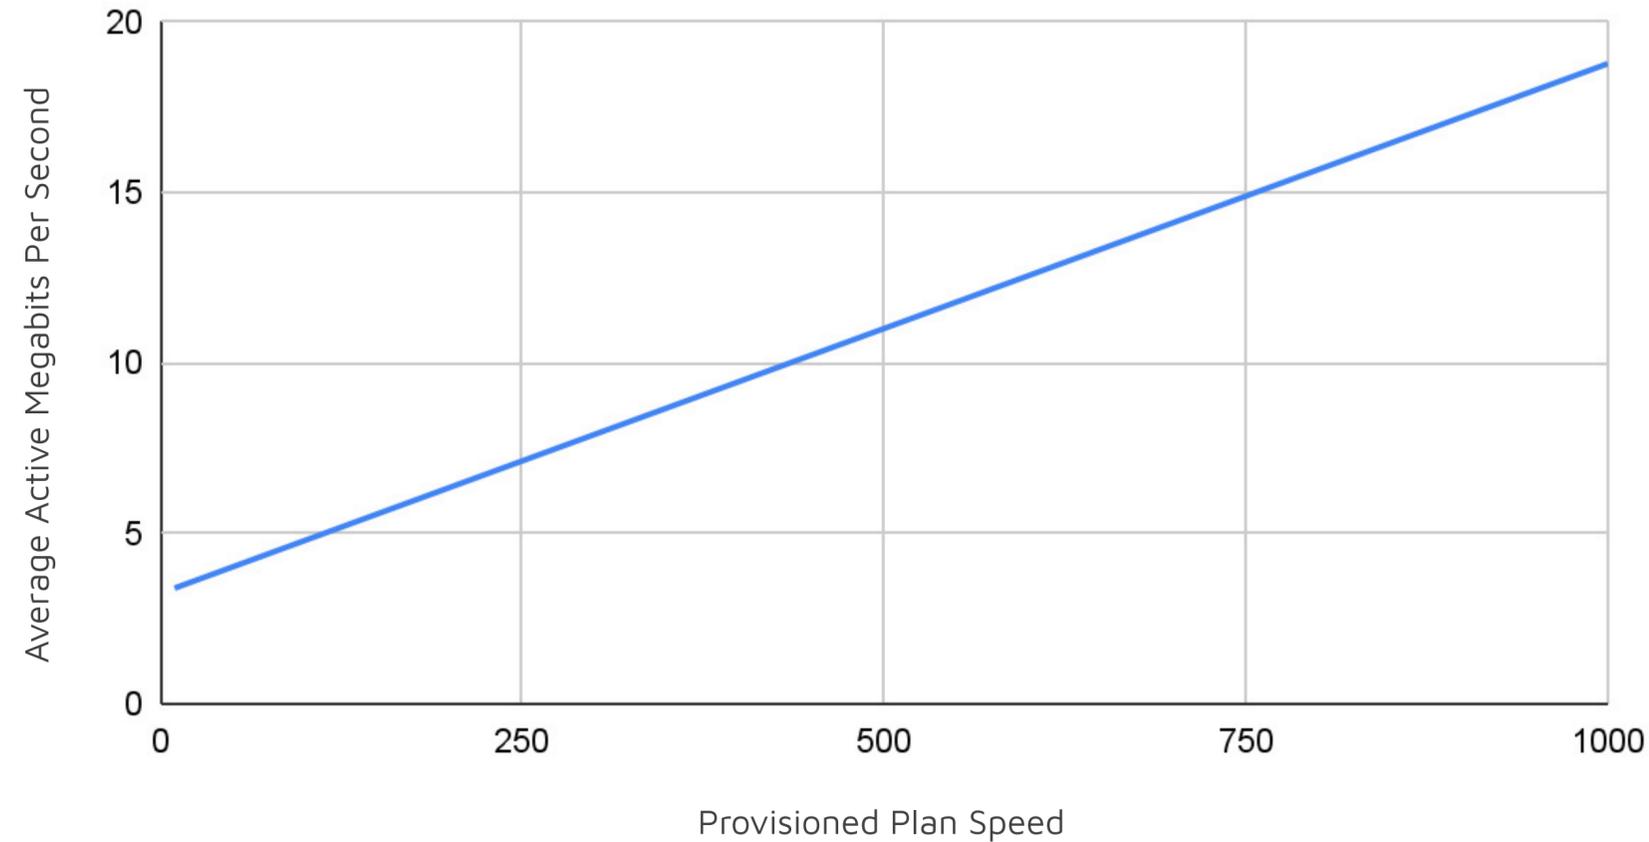
So while these numbers are real and indicate the actual subscriber experience, they don't necessarily indicate the upload performance that a subscriber could achieve if desired. The 95th percentile value (P95) may be a better indication of the rates that are possible in this case.

In comparison to download throughput, the per-subscriber upload rates show a larger difference in the ratio between peak and off-peak. This indicates that networks may be more congested in the upstream during peak than normally assumed. While subscribers care most about download performance, it's important to note that a congested upload path can cause download throughput problems because of TCP acknowledgment starvation. Similarly, a congested upload path can cause packet loss, which is also bad for subscriber QoE.

Active Throughput by Plan

Per-subscriber download throughput can also be compared across the spectrum of speed plans.

Median Used Download Rate by Plan Speed (Predicted)

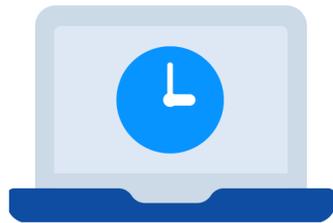


In general, even when including other access technologies than fixed wireless (such as fiber optic, cable, and DSL), we see a clear trend.

Using our high-resolution data set of user behavior, we take the average throughput sample of the average user when active and compare it across speed plans. This insight can help operators plan network capacity, whether adding new network capacity with new hardware or simply introducing faster plans to consume more of the existing capacity. The resulting slope is approximately a 3:2 (1.5:1) ratio. That is, you have to triple the provisioned plan speed to double the required demand.

Latency

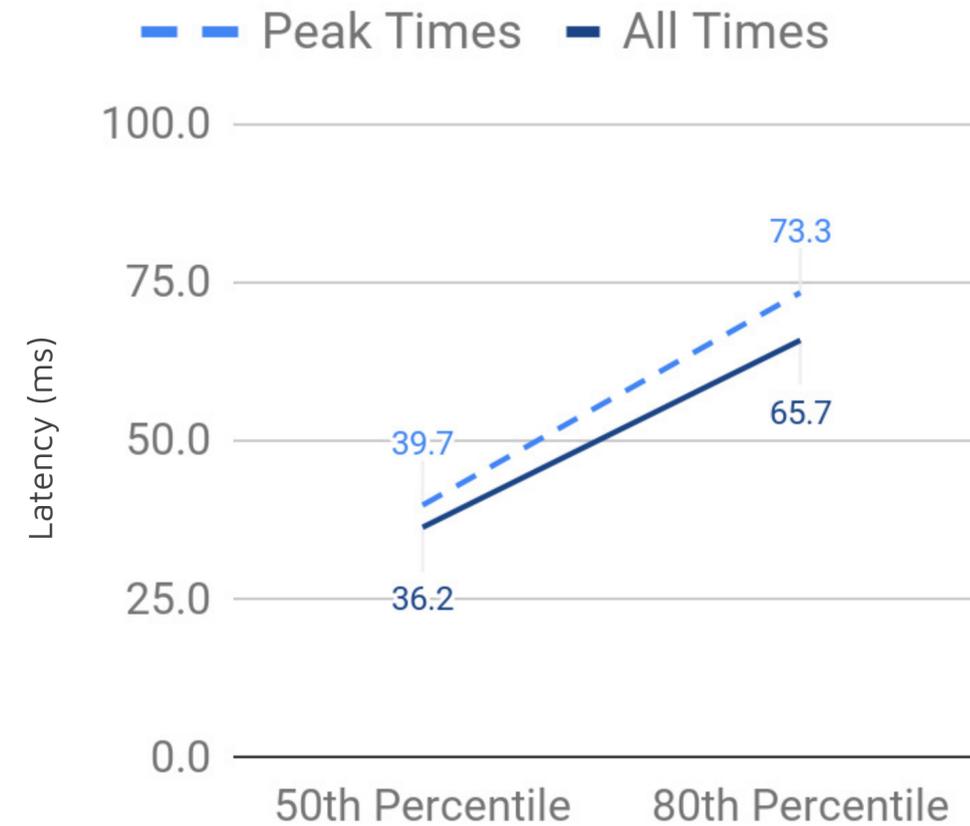
Preseem measures latency by tracking the round trip time for individual TCP segments.



This gives us a detailed view of the latency in the access network, and results in thousands of latency samples per second, per subscriber. This is fundamentally different from an ICMP ping-based latency measurement because it measures true end-to-end latency, including the latency in the subscriber's home.

Somewhat surprisingly, the latency difference between peak and off-peak times is relatively small. Note that these metrics are collected from networks where Preseem is deployed to optimize latency and the subscriber experience. It's likely that the latency in networks without such optimization is significantly higher. Peak latency has dropped nearly 4% year over year.

Subscriber Latency



SIDEBAR

Where Does Latency Come From?



Latency, or delay, is the time it takes for data to move through the network. There are many different sources of latency.



Propagation latency

Propagation latency is the time it takes for the electromagnetic or optical transmission to move from point A to point B. Unless you discover new physics, you can't do much about this (although high speed traders do things like buy shorter fiber cables).



Queueing Delay

Queueing delay is the largest source of latency in a network. When a device starts receiving a packet, it must hold onto that packet until it's been completely received and then begin transmission on the output port. For example, receiving or transmitting a 1500-byte Ethernet packet at 1 Mbps takes 12 ms—that's the best case. Typically, there's a buffer used to absorb bursts and enable prioritization. The size and techniques used to manage this buffer drastically affect the latency it introduces. Bad buffer management results in the dreaded Bufferbloat problem. Preseem reduces queueing delay through active queue management (AQM) techniques that greatly improve subscriber QoE, even when the network or the subscriber's connection has reached its capacity.



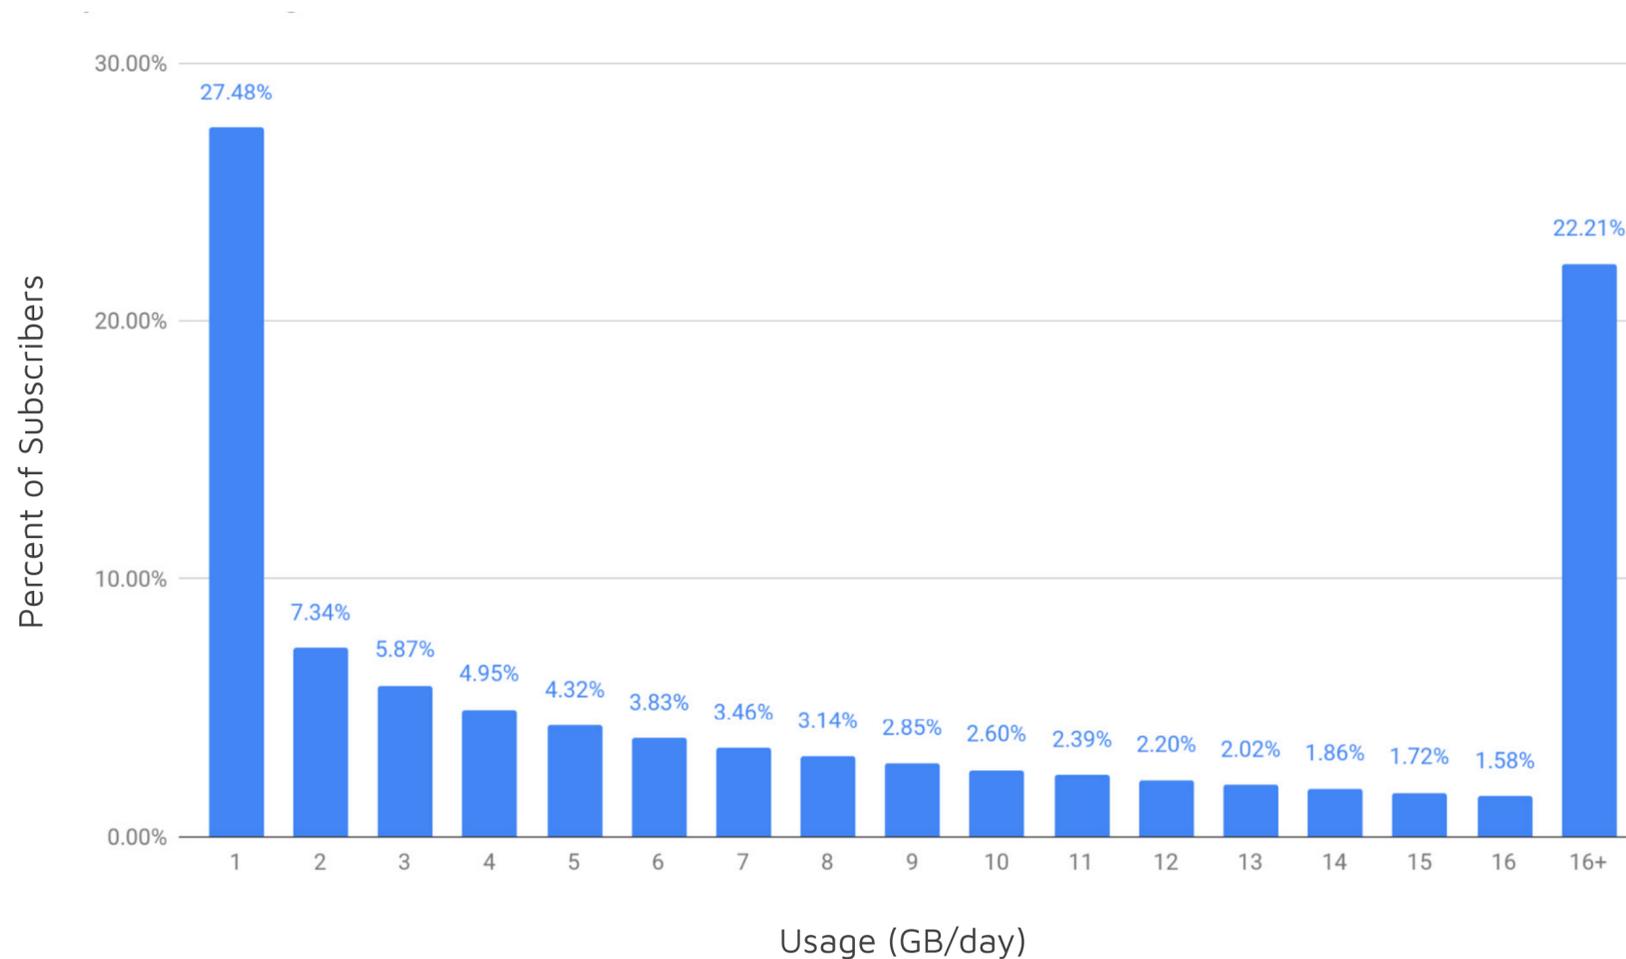
Frame Aggregation

In a sense, this is a type of queueing but because it's so prevalent in wireless networks it's worthwhile to discuss separately. To achieve higher throughput, many wireless technologies aggregate several Ethernet/IP frames into one radio frame. This optimizes for throughput at the expense of latency, as the access point waits some predefined amount of time to construct the aggregate before transmission.

Subscriber Usage

Subscriber usage refers to the total number of bytes transferred by a subscriber over a day or month.

Daily Download Usage



From the perspective of the subscriber experience, the total usage isn't very instructive. That's because a large amount of usage consumed during off-peak time has less of an impact on perceived network quality than a smaller amount of usage during peak.

The average download usage for fixed wireless subscribers is 10.6 GB per day or 329 GB per month. As expected, the average hides the significant variation that occurs between subscribers. The balance is trending toward heavier users.

Here we see that 27% of subscribers (1Y 20%↓) use less than 1 GB of download usage per day and just over 22% (1Y 30%↑) use more than 16 GB per day (half a TB per month).

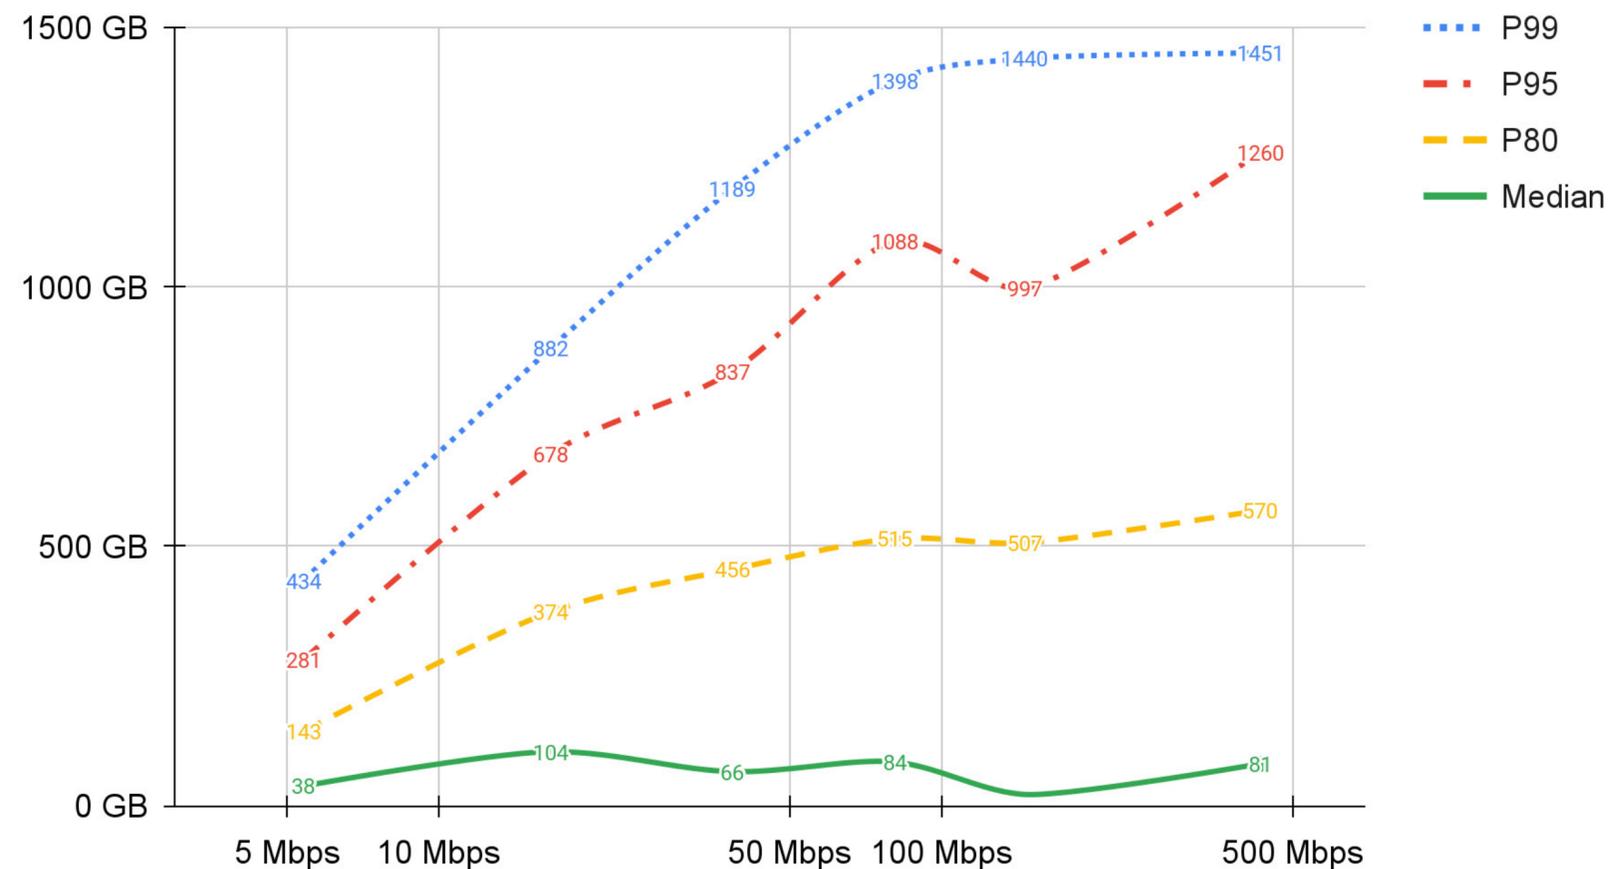
Subscriber Usage by Speed Plan

In the previous chart, we refer to the “average” user. This doesn’t present the whole picture of user behavior, however. By breaking down subscriber usage by speed plan, we can see the relationship between plan and total data consumed. In all cases, there’s a large gap between the median user and users in higher percentiles.

In fact, the median user doesn’t really use more data as the speed plan increases. At the 80th percentile and above, we do see steady increase but even in the 99th percentile we see a rapid falloff above 75 Mbps. This likely suggests that gigabit plans are not required for the vast majority of subscribers.

The underlying data for this chart includes all technologies for which Preseem collects user data: fixed wireless, cable, DSL, and fiber. Given that the majority of plan speeds greater than 250 Mbps are delivered over fiber, we can see that capacity constraints are a factor in smaller plans but only up to a point. From a consumption perspective, raising plan speeds doesn’t proportionately break the bank in terms of resources consumed.

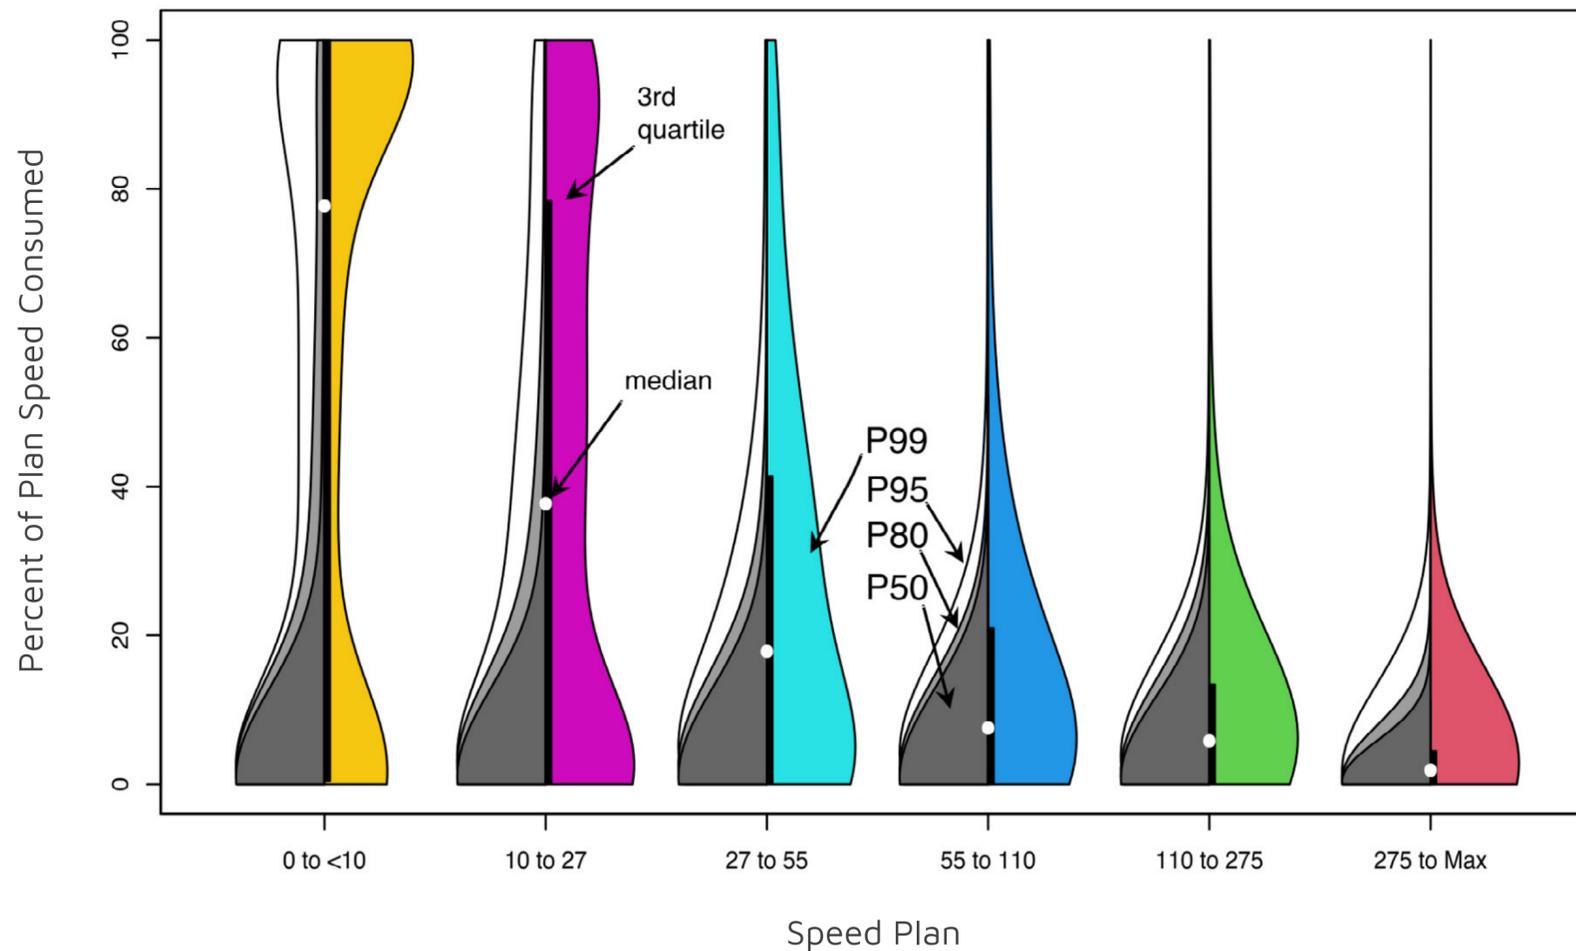
Plan Speed Vs. Monthly Download Usage



Speeds Attained by Speed Plan

Given that an individual subscriber's demand is idle much of the time, there are some interesting results when we compare different percentiles of utilization at different plan speeds.

Density Plot of Speeds Achieved per Download Plan



Here we analyze user consumption habits. To look at a single percentile of speeds attained across all users could be misleading, so we compare four different percentile curves, just as we did in the previous graph: P50 (the median), P80, P95, and P99.

The median speed consumed doesn't vary as much as the extremes. Looking at the 99th and 95th percentiles, we can see that users consume their full plan more of the time in the smaller than in the larger plans. This may seem obvious, but it also shows that user needs are being well met in plans below 100 Mbps. In all cases, most users do not max out their bandwidth.

On the vertical axis, we represent speed attained as a percentage of plan speed, which allows clear comparison of performance across speed tiers. In all cases, 100% of plan speeds is attainable. This implies that, while oversubscription (and underprovisioning) can and do occur, in the vast majority of cases it's not a deciding factor in customer service delivery and not common network practice.

Access Point Metrics

Preseem collects and utilizes many access point metrics when measuring and optimizing subscriber QoE. This section presents an analysis of Preseem metrics grouped by access point model and vendor information.

You might say, “This doesn’t look anything like the access point spec sheets! These numbers look too low!”

Remember that these are real-world throughput numbers observed by Preseem, not the highest attainable figures. For example, if a model T access point is capable of 100 Mbps, but every model T Preseem sees only has one subscriber, then the reported rates for model T access points will be low. However, this extreme scenario is unlikely for any but the most rare of AP models.



AP Market Share

In order to understand the access point market share, we look at two metrics:



The percentage of the fixed wireless market by the number of access points

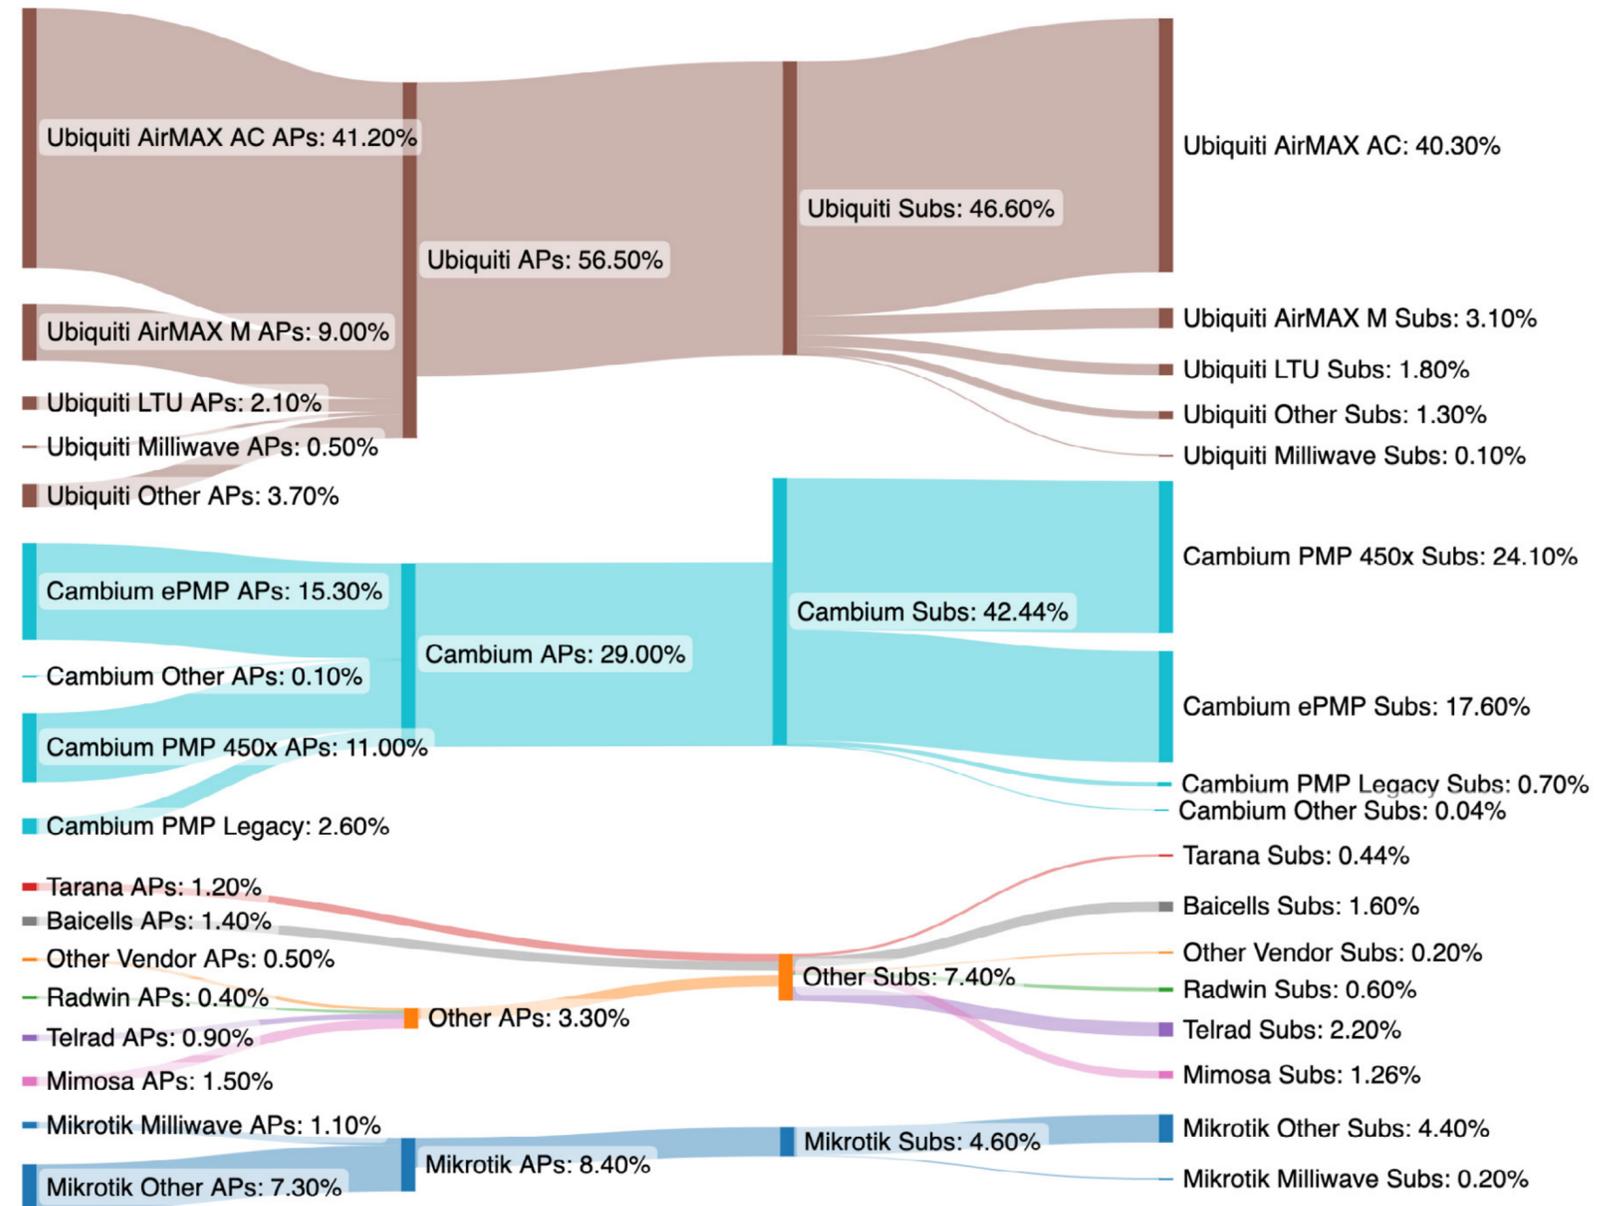


The percentage of the market by subscriber count

In both cases, the clear leaders are Cambium and Ubiquiti, with Cambium supporting a higher density on a smaller number of APs.

Reading from the left, the first column shows percentage of the Preseem market by Access Point. The two central columns break down by vendor: first APs, then Subscribers. The far right column breaks down subscriber totals by individual models or series.

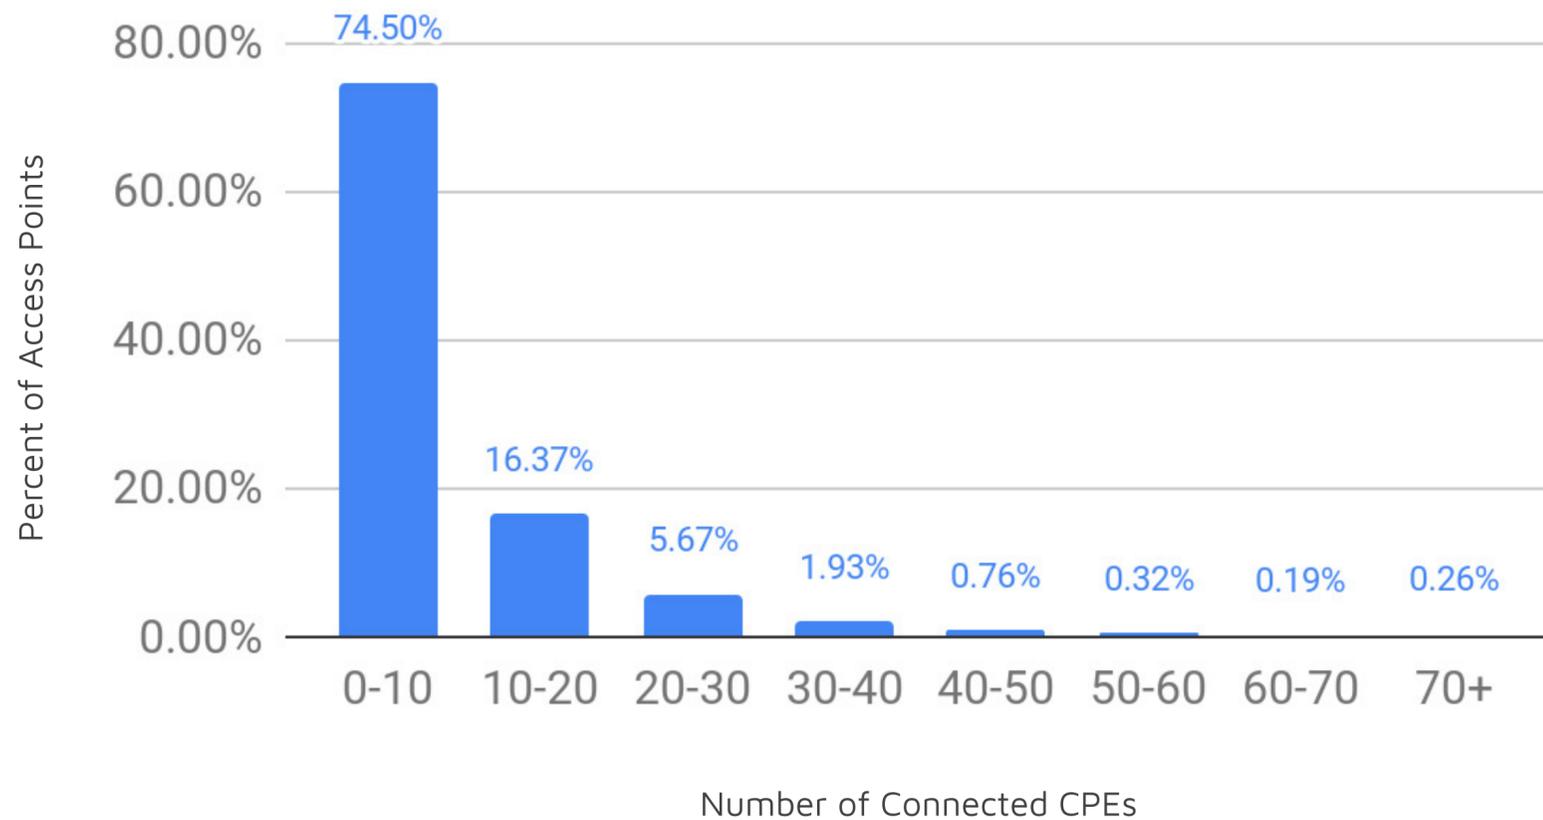
Access Point Market Share Analysis In Percent



Connected Subscriber Count

As you might expect from the difference in AP market share by element and subscriber count, there's significant variation in how many subscribers ISPs connect to individual APs.

Connected CPEs and Access Point Count (%)

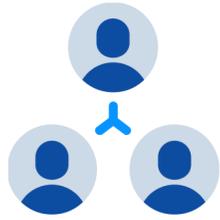


The chart on the left shows the number of access points bucketed by the number of subscriber radios connected to the AP.

Surprisingly, the data shows that over 74% (1Y 3.4%↑) of access points have 0-10 subscribers attached.

Breaking this down by a few of the top AP models paints a different picture.

Here we see that some AP models are typically deployed with many more subscribers than the overall average.

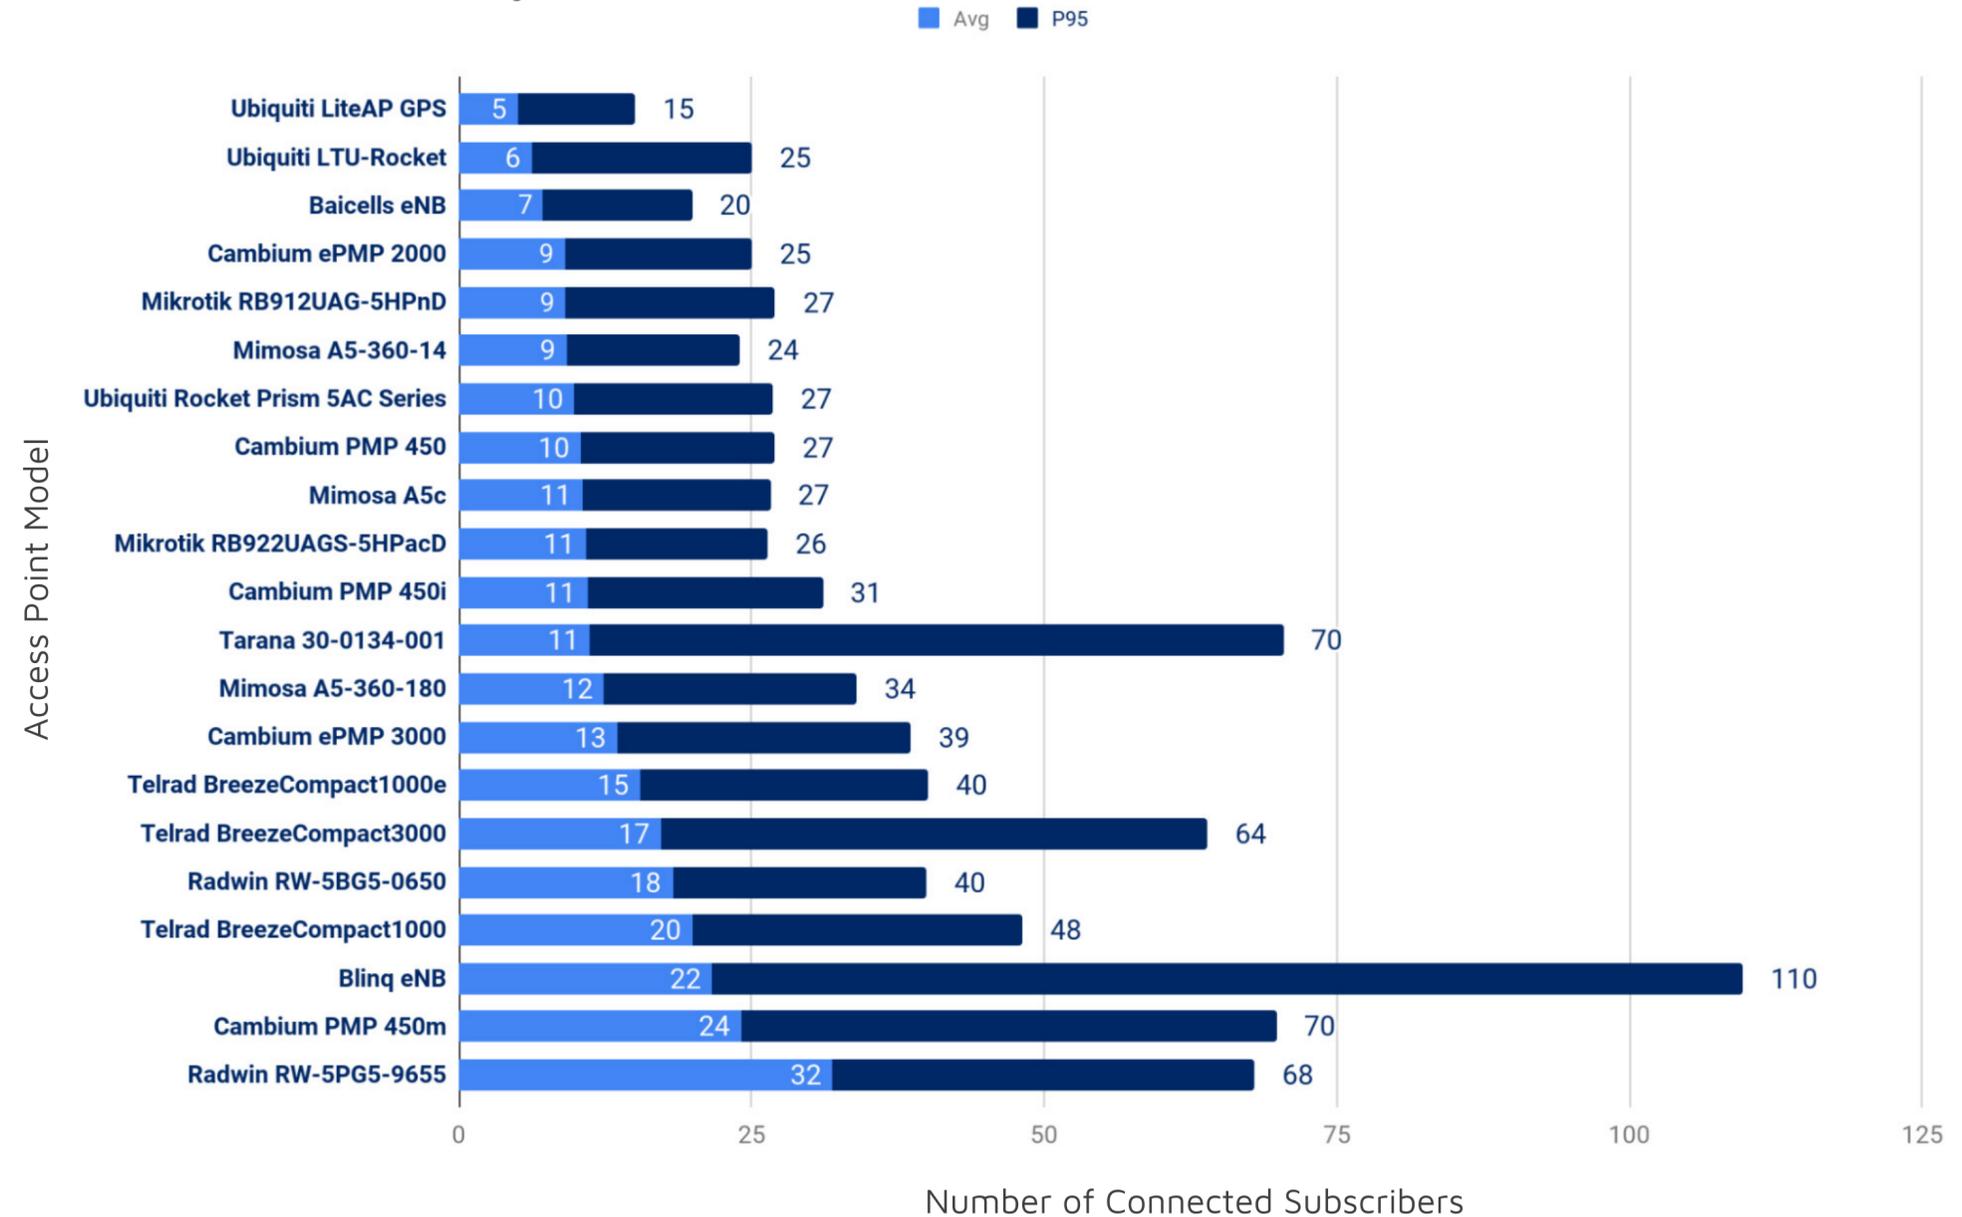


In some cases, deploying with many subscribers has a negative impact on QoE. Compare this with model latency on page 27.

Throughput

Since measuring throughput outside of peak times provides little insight into the subscriber experience—because the network isn’t loaded—the throughput numbers that follow were taken from the busiest times of the day for each individual access point.

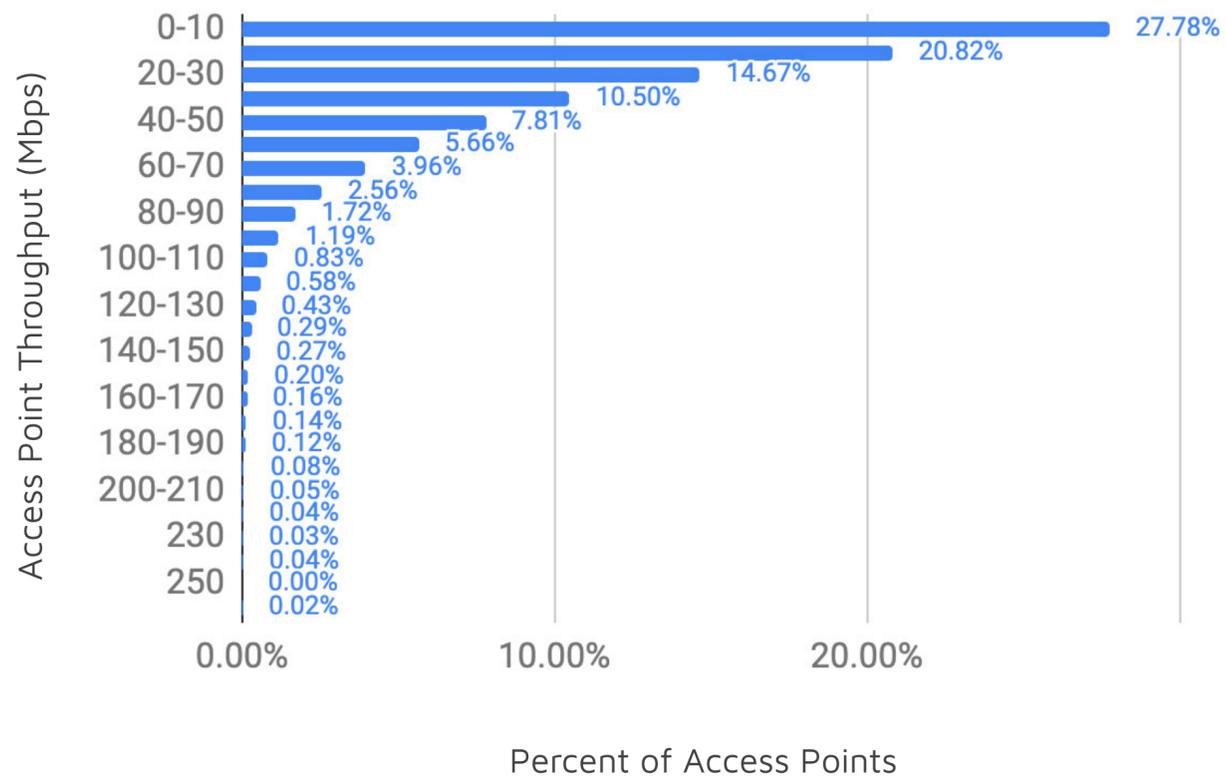
Connected Subscribers by AP Model



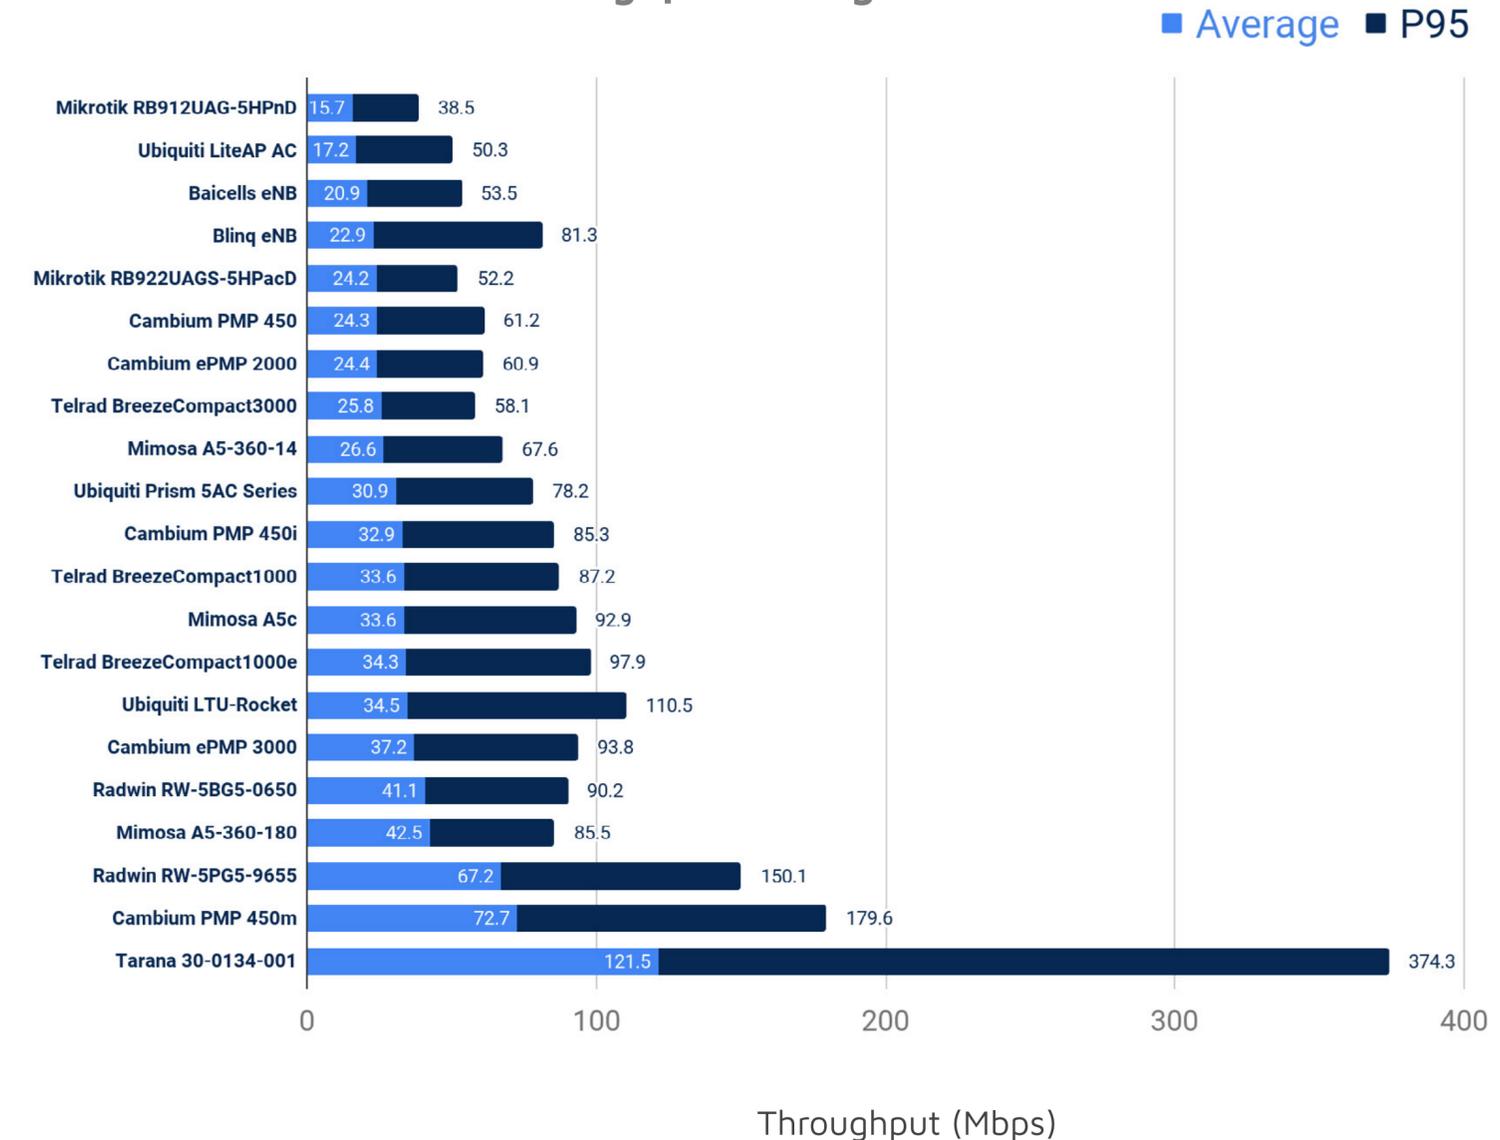
Download Throughput

The data shows that, in terms of the number of access points, almost 28% (1Y 11.7%↓) of deployed APs deliver less than 10 Mbps of real-world throughput. This is somewhat surprising but looking at the more modern equipment shows a very different pattern.

Access Point Download Throughput (Mbps)



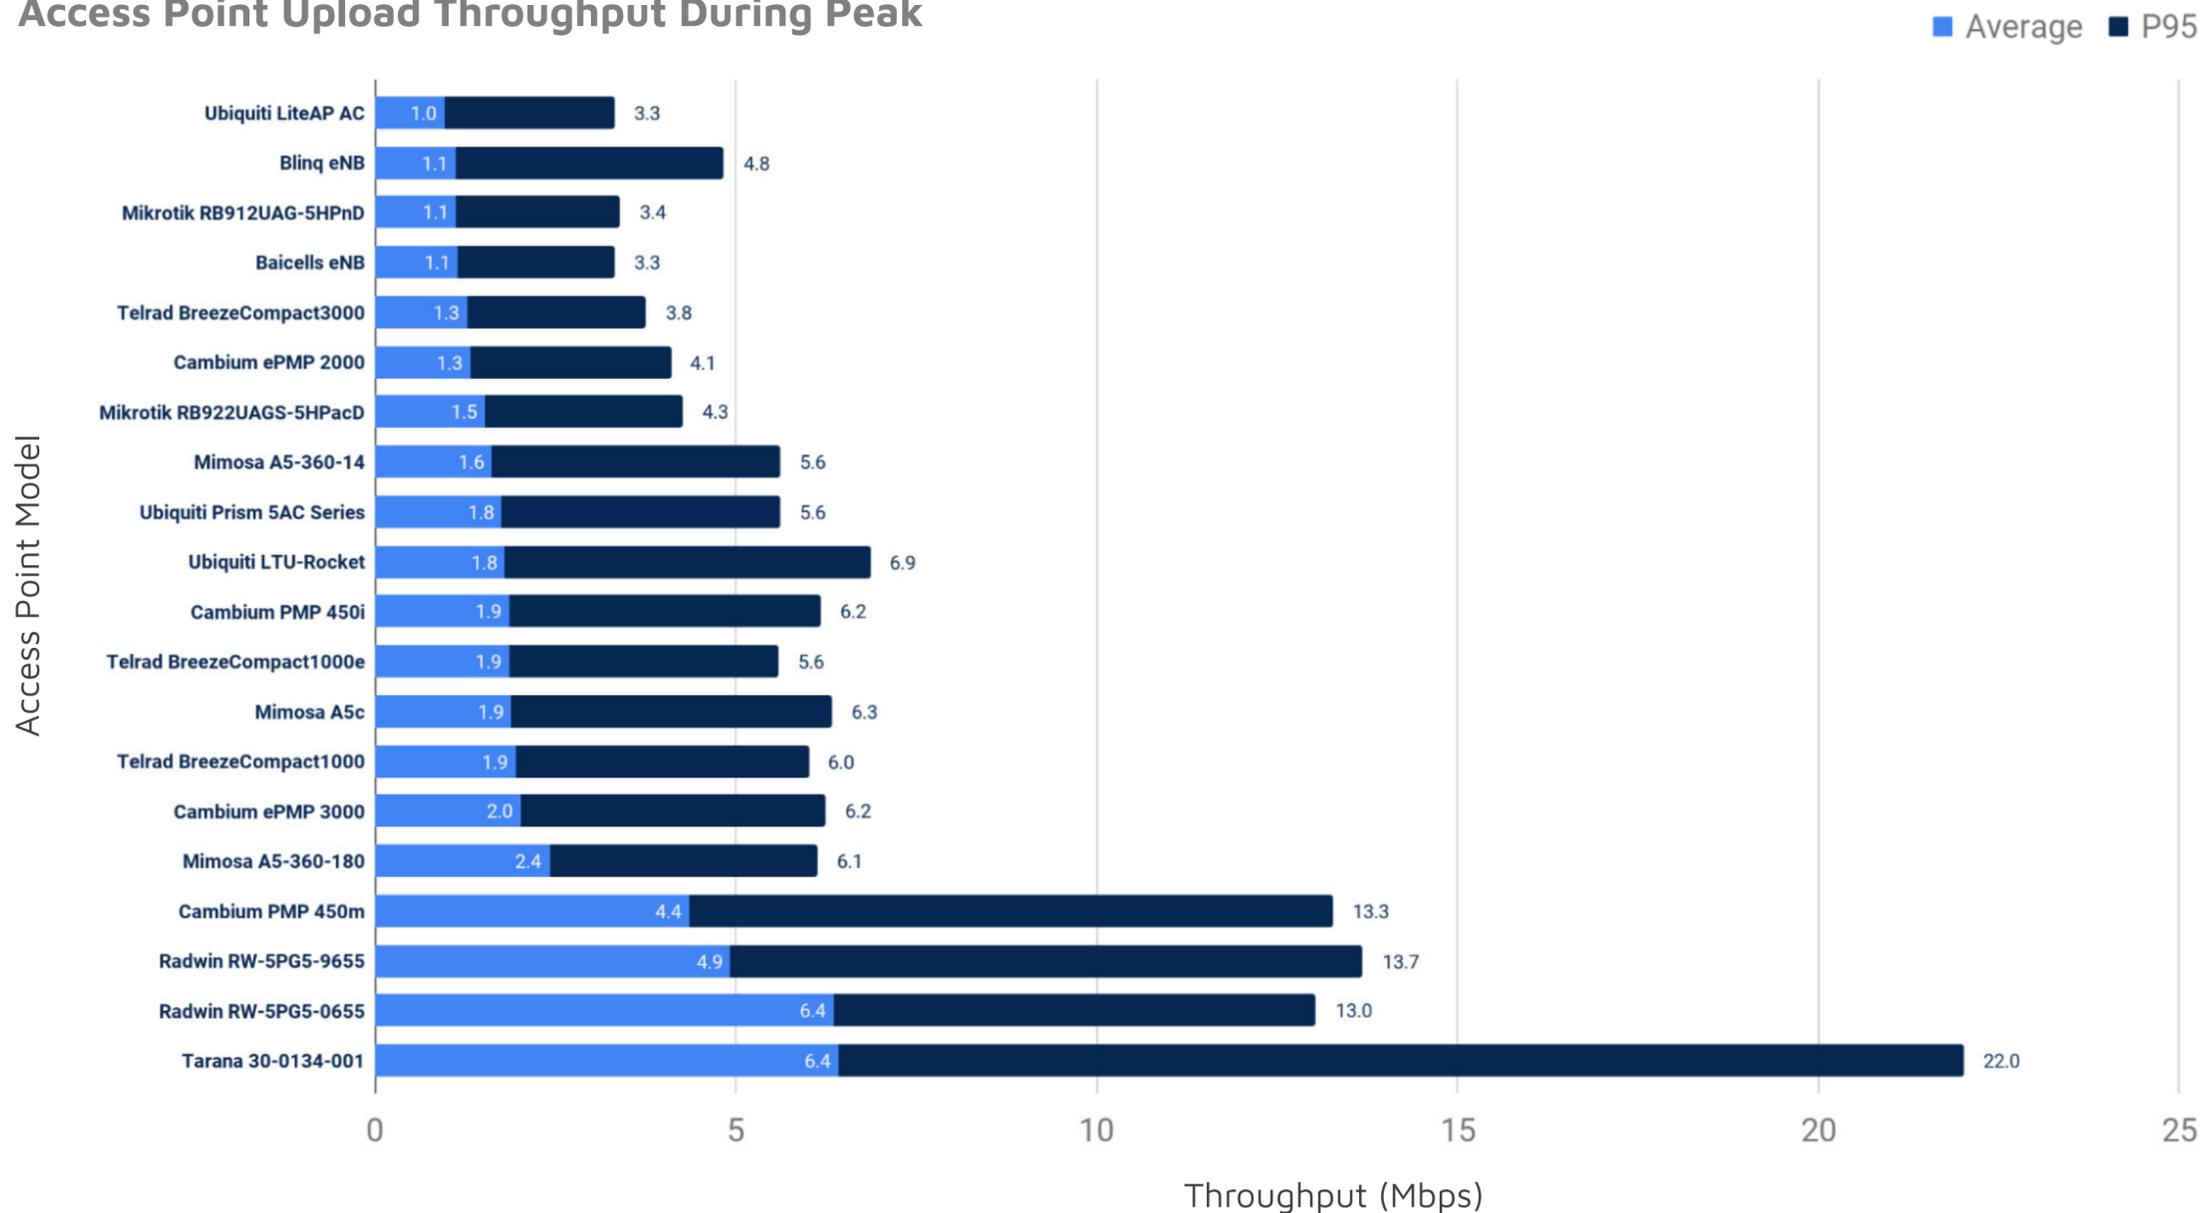
Access Point Download Throughput During Peak



Upload Throughput

Characterizing upload performance is more difficult because demand is often lower than what the network is capable of (see earlier discussion). However, there are still some interesting insights to be gained. In particular, almost all APs deliver less than 10 Mbps of upload throughput during the times of the day with the highest demand.

Access Point Upload Throughput During Peak



Latency

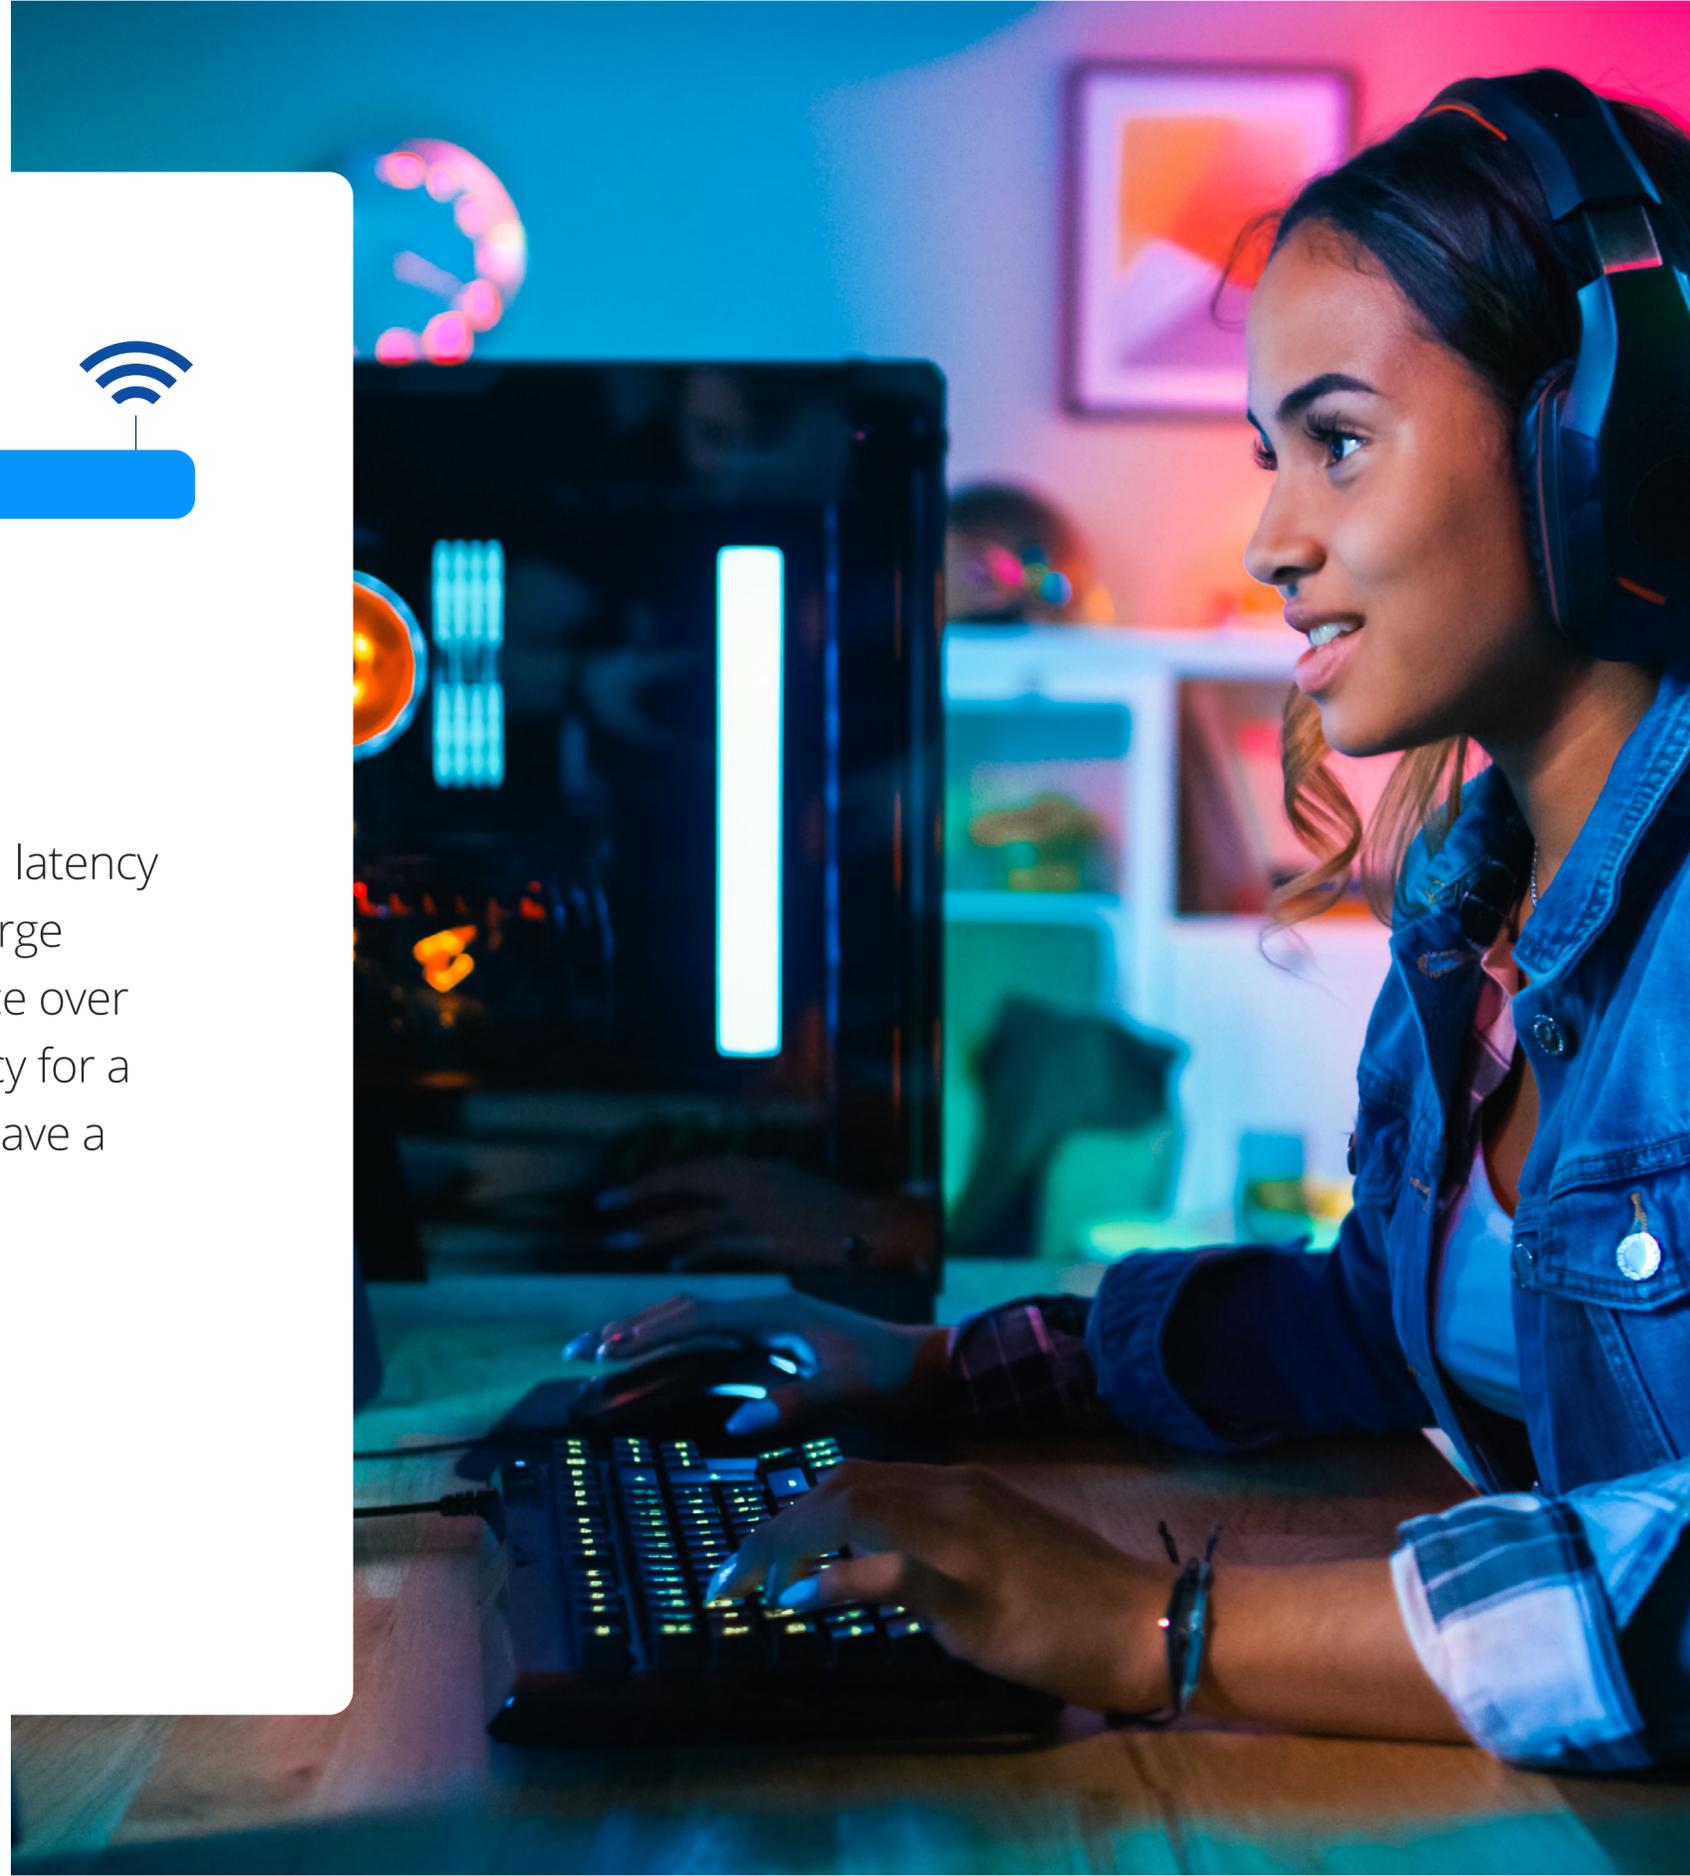
Across all access points, the distribution of latency follows an interesting pattern. Most APs deliver service with less than 100 ms of latency during peak times, but a significant number are over that benchmark. Latencies have improved since our last report, indicating that operators are upgrading their networks appropriately to keep up with rising demand.



SIDEBAR

What is a Good Latency Value?

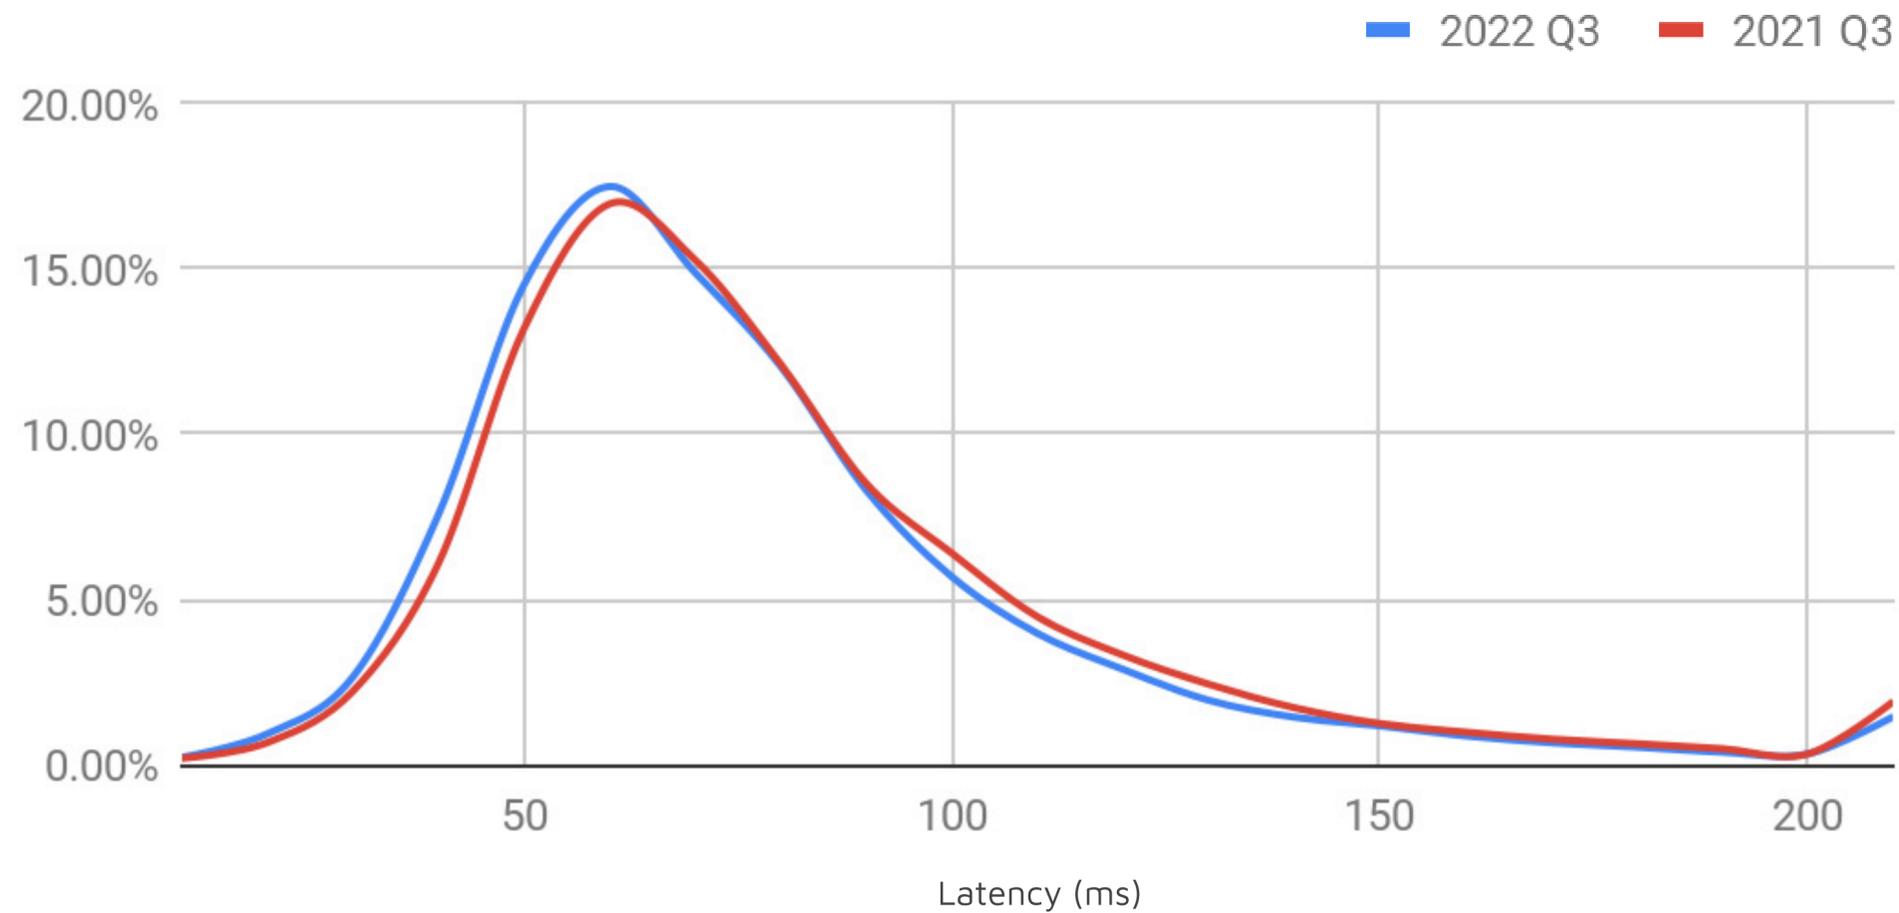
Latency requirements differ greatly by application. High latency has very little effect on Netflix, for example, but has a large impact on gaming. A simple point of comparison is Voice over Internet Protocol (VoIP). Typically, the end-to-end latency for a VoIP call needs to be less than 150 ms for the user to have a good experience.



ACCESS POINT METRICS

Note that the values shown in this report represent the latency from Preseem to the subscriber and back, and as a result do not include the rest of the path. Therefore, the values here need to be lower than 150 ms to achieve a good VoIP experience.

Latency: 2022 Q3 Vs. 2021 Q3

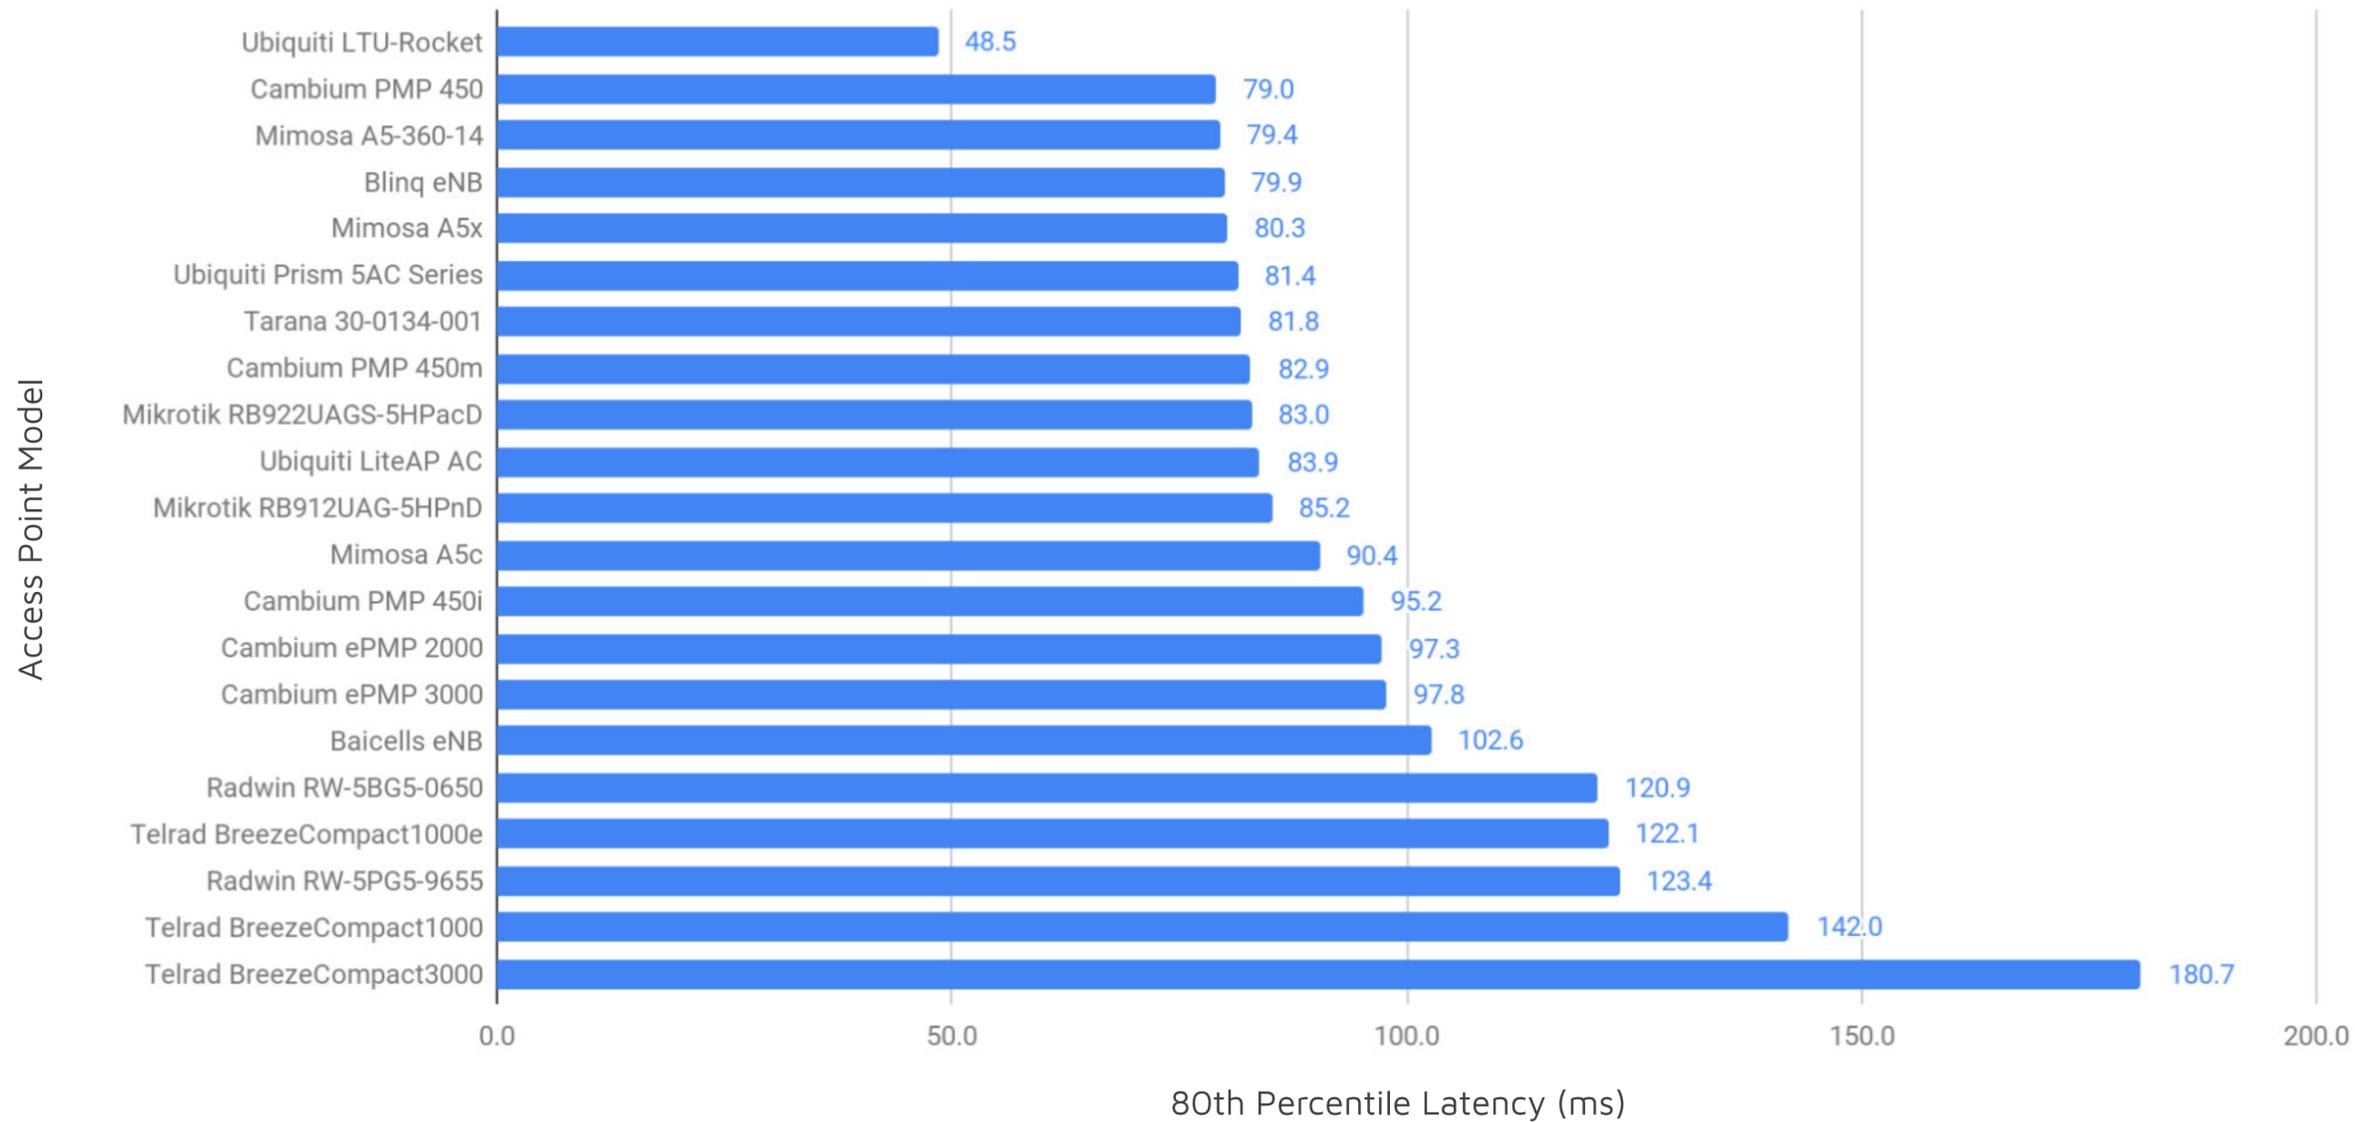


ACCESS POINT METRICS

Comparing latency across access point models shows significant variation and a general trend towards newer AP models having better latency characteristics. Subscriber QoE is strongly affected by latency, so this is an encouraging trend to see.



P80 Latency by Access Point Model During Peak



Access Point Subscriber Capacity

Airtime, bandwidth, and latency are all important metrics to understand, but what really matters is the number of subscribers that an AP can handle while still delivering a good experience.

Determining the capacity of an access point in terms of subscribers is complex, requires deep knowledge, and often comes down to intuition built up through experience.

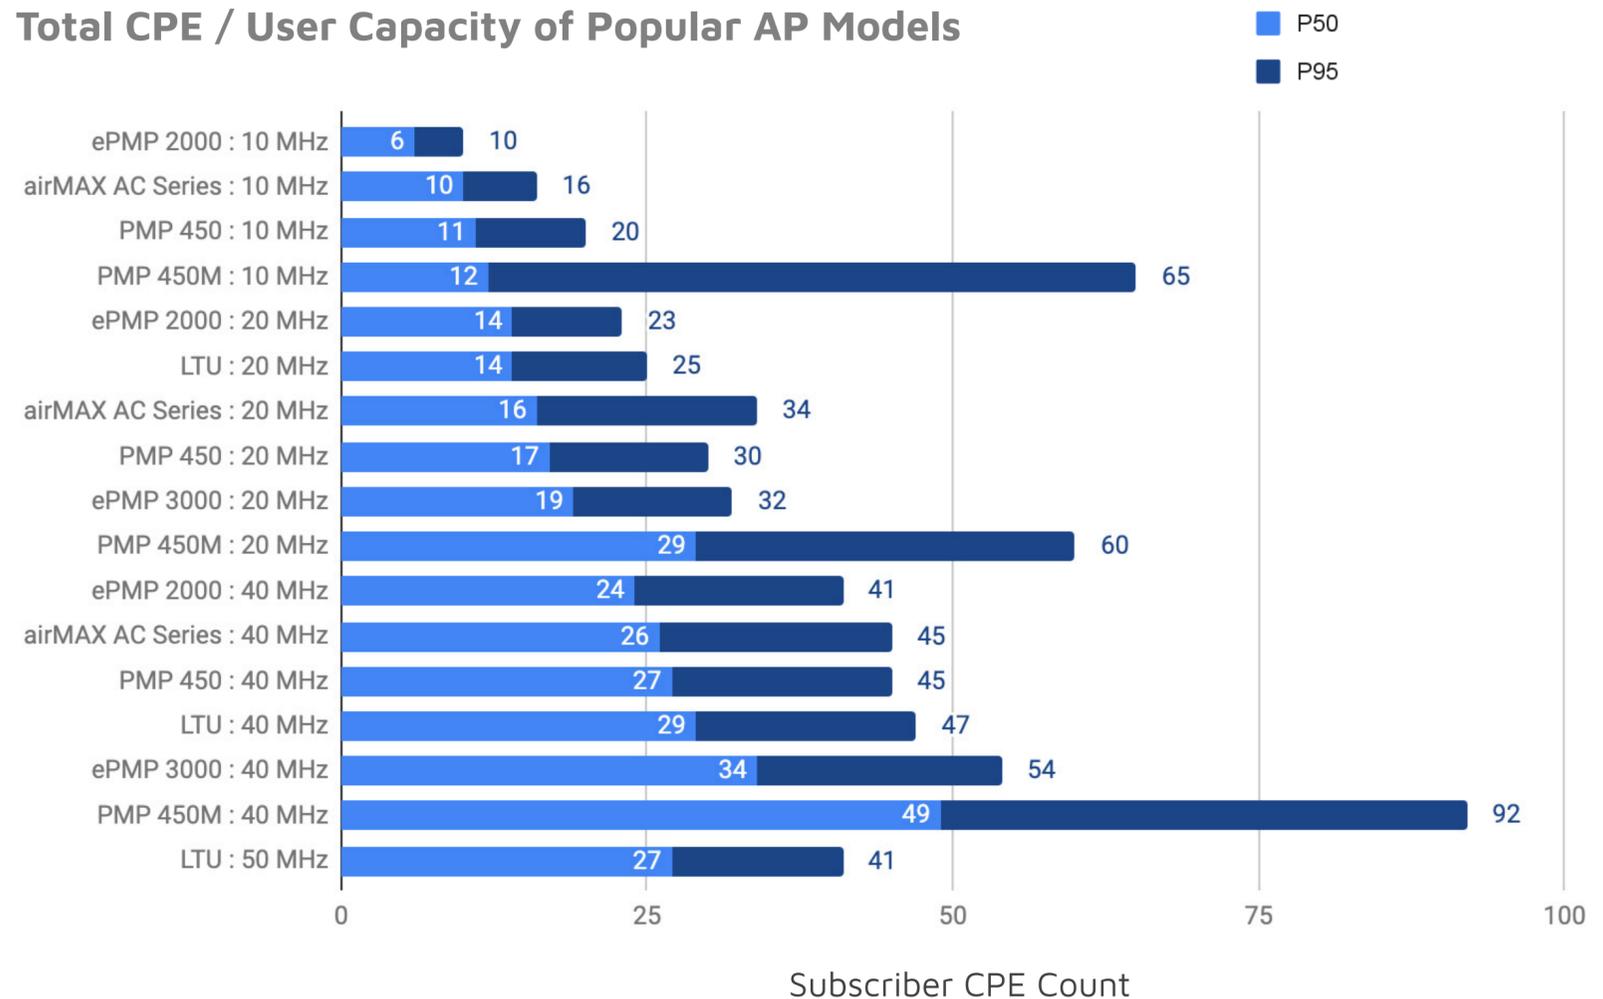
Preseem solves this problem by building a global performance model for each access point and configuration, and combining it with a model of subscriber behavior unique to each service provider. This answers the question of AP capacity in the simplest form possible—the available capacity in terms of subscribers.

This chart summarizes the subscriber capacity across all Preseem customers for some common access point models and channel widths. Many models commonly operate at between 50-60% of the achievable (P95) efficiency. In the case of the PMP 450M, adding more users actually increases the efficiency up to a point, due to increase in opportunities for MIMO transmission.



Access point capacity, represented in terms of subscribers, can be used to drive marketing decisions and provides a common reference point between sales and network engineering teams.

Total CPE / User Capacity of Popular AP Models



RF Channel Width

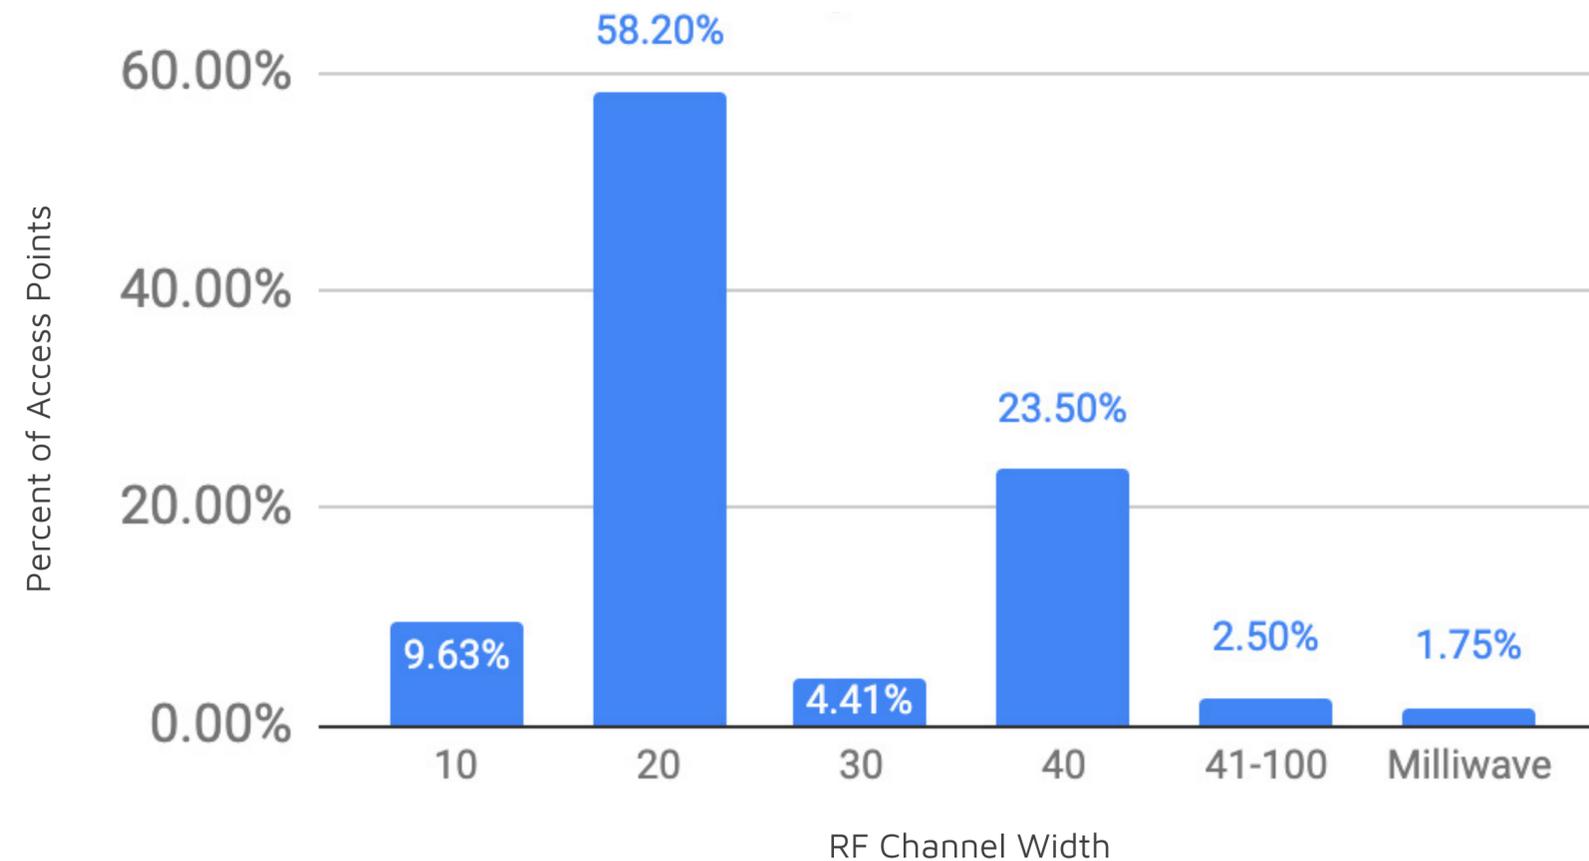
Besides obvious items like the location and AP model, the choice of channel width is one of the more important decisions that needs to be made for every site. In this section, we look at what channel widths are used and how that impacts the network and subscriber experience.



Percent of APs by Channel Width

Across all the access points that Preseem monitors, 58% (1Y 9.4%↓) use 20 MHz channels, with 40 MHz and 10 MHz channels being the next most common. Wider channels are on the rise, so we begin reporting them for the first time. A small percentage of channels are narrower than 10 MHz.

Percent of Access Points by RF Channel Width

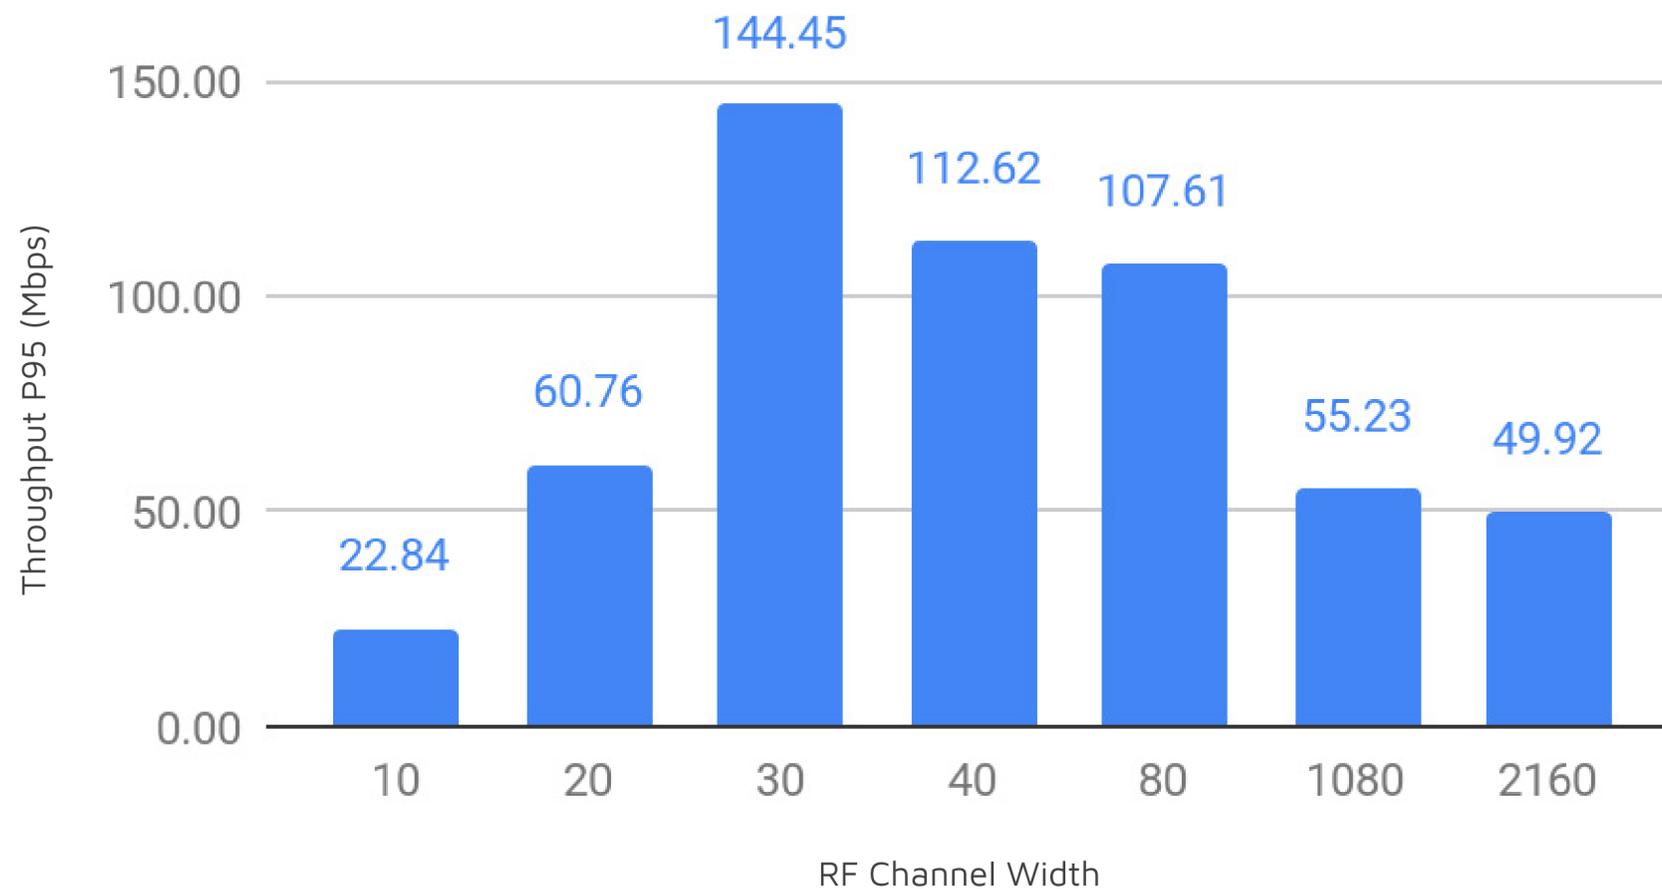


Download Rate by Channel Width

By looking at all APs aggregated by channel width we can see that while a larger channel width increases throughput, the effect is not always equal to the increase in the channel size. Larger channel sizes are growing in popularity but are not as fully utilized, and will show more potential as airtime consumption increases.



Download Rate



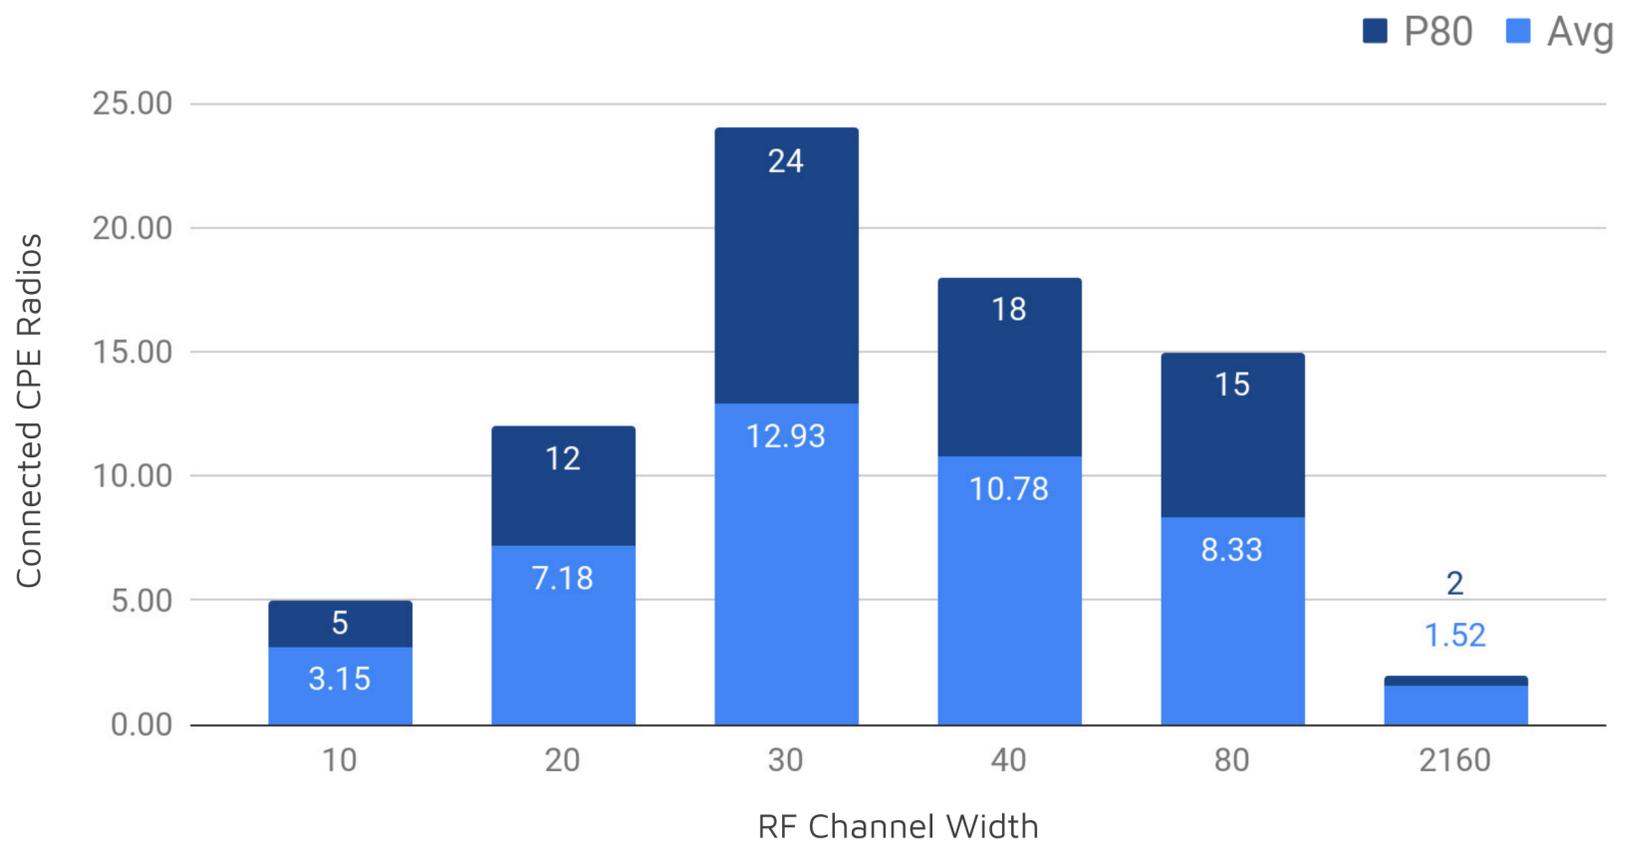
Connected CPE Radios by RF Channel Width

The following chart shows the number of connected customer premise equipment radios (CPEs) by channel width across all access point models.

Milliwave radios are on the rise, using very large channel widths, but average very few connections per radio due to the high density required. Milliwave has an effective range of just hundreds of feet, but delivers very high capacity.

30 MHz and 80+ MHz channels are still relatively unpopular, so while we've included them in the chart, less reliable conclusions can be drawn.

Connected CPE Radios by RF Channel Width

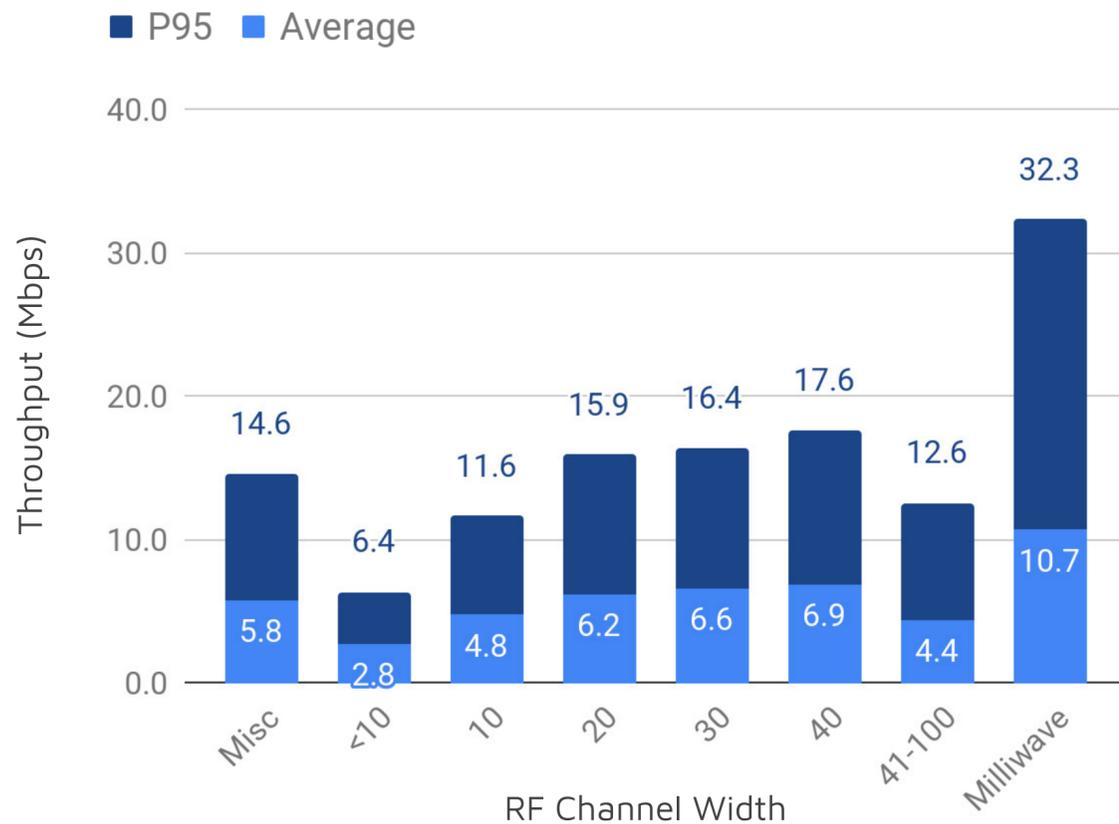


Subscriber Throughput by RF Channel Width

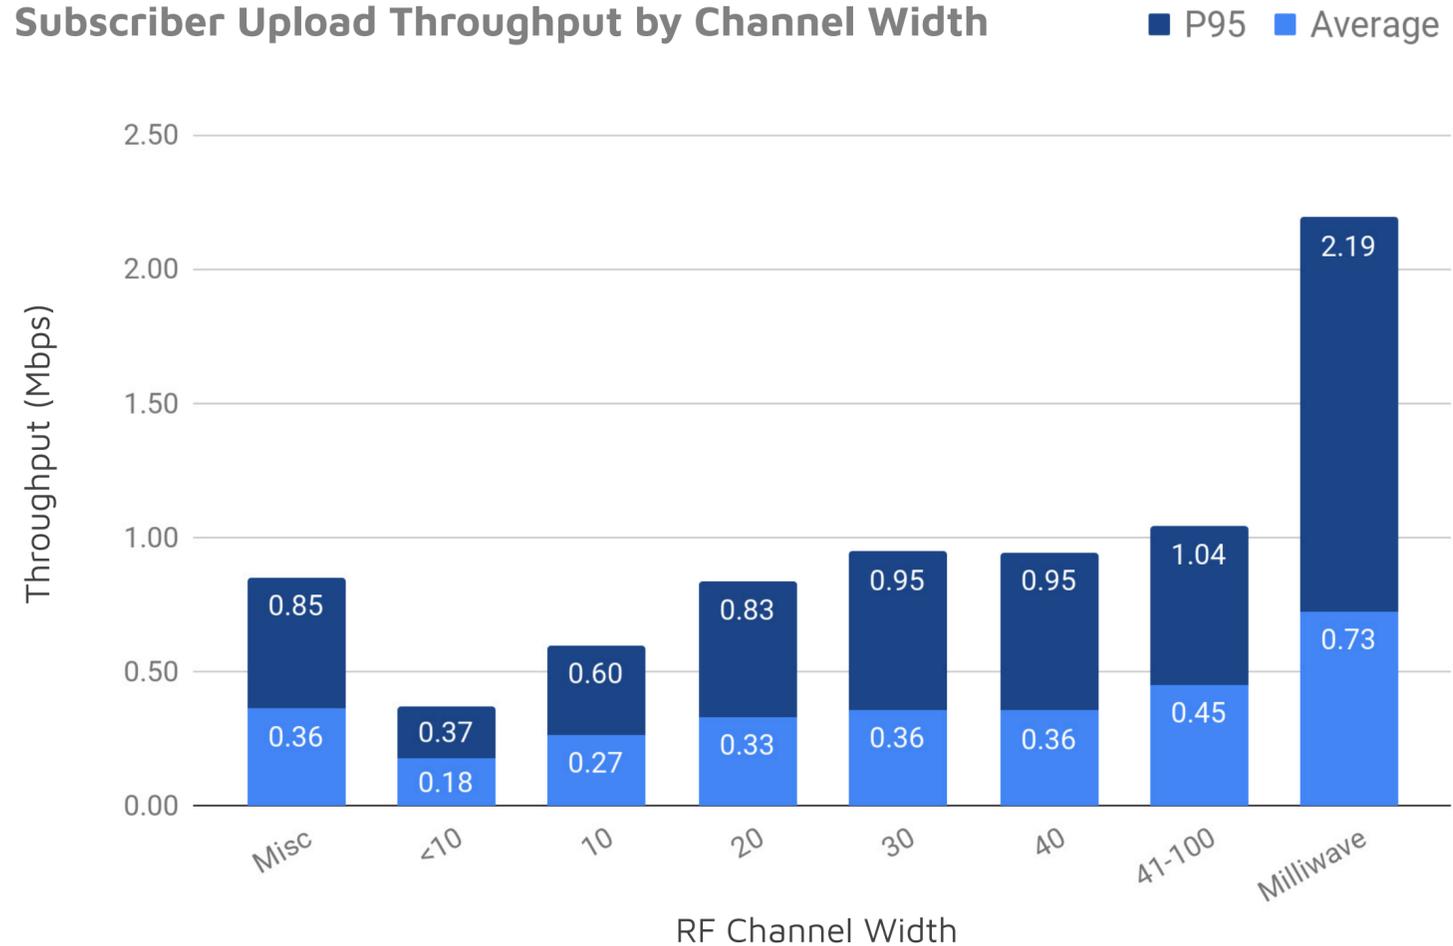
Of course, the number of connected subscribers is only part of the picture, as it doesn't relate perfectly to the subscriber experience.

The next two charts show the typical per-subscriber download and upload throughput by channel width. The results show little increase in either download or upload throughput across channel widths. This suggests that ISPs use wider channels to handle more subscribers instead of delivering higher per-subscriber throughput.

Subscriber Download Throughput by Channel Width



Subscriber Upload Throughput by Channel Width



Oversubscription Ratio

In most networks, some amount of oversubscription is normal. For example, a wiring closet switch may have twenty 1G ports with a single 10G port to the core network. This results in a 2:1 (sum of port rates/uplink port rate) oversubscription ratio.

Internet provider networks are no different in this regard. No ISP can afford to provision enough bandwidth from the edge to the transit point for every subscriber to use their entire plan rate at the same time. The business model simply doesn't work.



Like all the data in this report, the typical throughput is based on APs of a given model as seen across all Preseem customers, not the marketing spec sheets.

The next chart provides some insight into how much ISPs oversubscribe their wireless access networks. The formula used to calculate oversubscription is simply:

Sum of subscriber plan rates on AP /
typical throughput for that AP model
and channel width

For example, if an ISP has sold twenty 10 Mbps plans on an AP that typically achieves 50 Mbps, then the oversubscription ratio is:

$$20 * 10 / 50 = 4$$

OVERSUBSCRIPTION RATIO

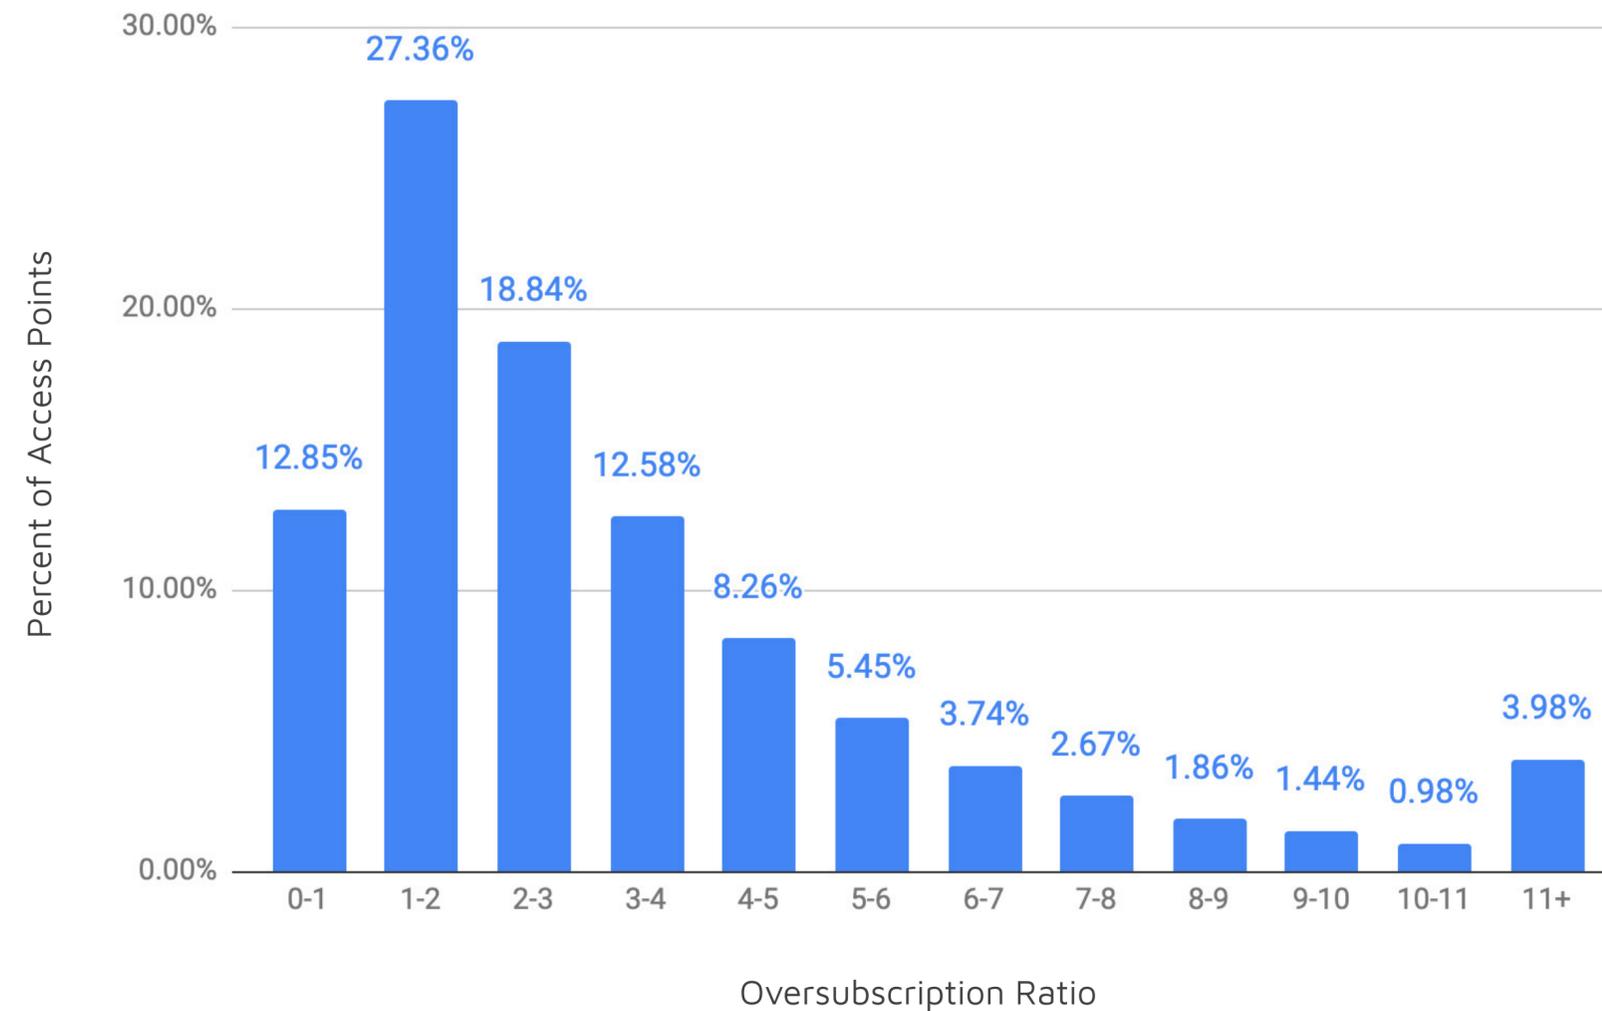


This chart shows that 40% of APs are under two times oversubscribed.

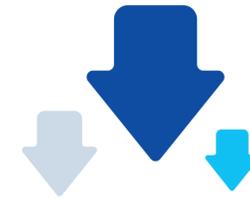
Oversubscription in and of itself is not bad. However, higher oversubscription ratios generally lead to lower subscriber throughput during peak, and therefore a poorer subscriber experience.

There are some outliers that deliver more throughput and are capable of sustaining a higher oversubscription ratio without damage to the subscriber experience.

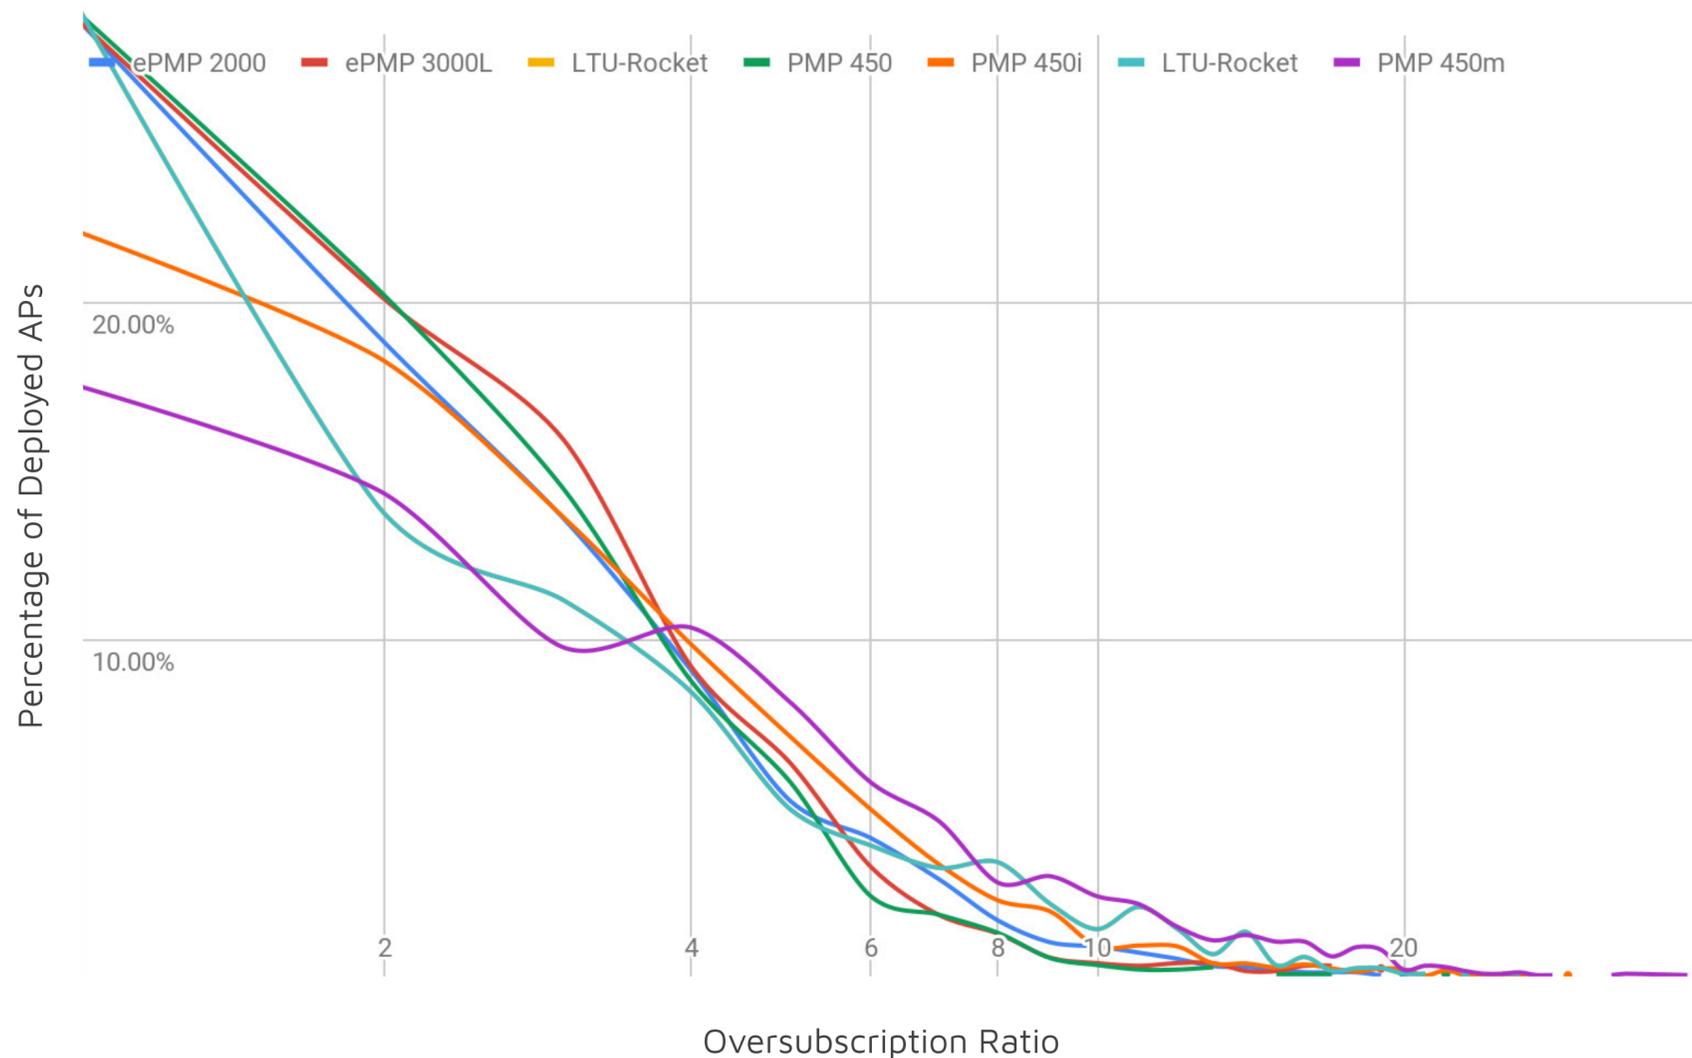
Access Points by Oversubscription Ratio



A Comparison of Some Popular Access Point Models by Oversubscription Ratio



Oversubscribed Ratio by Percentage of Deployed APs



As can be seen in the chart, the oversubscription distribution is relatively stable across many models, with the PMP 450m being an outlier. This suggests that the 450m supports a much higher level of oversubscription than these other models.

Summary

This report leverages Preseem's data to provide a view into the fixed wireless industry that we hope is both interesting and enlightening. Among the most surprising results is that the majority of ISP access points have less than 10 attached subscribers, and that the majority of APs are less than three times oversubscribed.

Highlights include:



6 Mbps When Active

The average fixed wireless subscriber uses around 6 Mbps when active (1Y 17%↑)



Speed Plans

Speed plans over 100 Mbps often sit idle



Data Usage

The average fixed wireless subscriber uses 10.6 GB of data per day, for a total of 329 GB per month (1Y 20%↑)



Oversubscription Rate

Over 59% of access points are less than 3x oversubscribed



Consumer Habits

Many users don't consume more data as their speed plan increases



Improved Latency

Latency has improved year over year, indicating that operators are keeping up with consumer demand

Appendix

Glossary of acronyms and terms (in alphabetical order):

Acronyms

AP: Access Point

AQM: Active Queue Management

CPE: Customer Premises Equipment, network-related equipment owned or leased by an internet subscriber and located within their home

ICMP: Internet Control Message Protocol, a way for network devices to exchange messages and communicate problems with data transmission

ISP: Internet Service Provider

QoE: Quality of Experience (of each internet subscriber)

RF: Radio Frequency

TCP: Transmission Control Protocol, a communications standard that permits computers to communicate and exchange data and messages with each other

Terms

Access Point: A hardware device that allows Wi-Fi devices to connect to a network

Bufferbloat: Experiencing slow internet due to other bandwidth-intensive traffic on your network, e.g. trying to stream a video while someone else in the home is gaming online

Channel Width: The size of a Wi-Fi channel, dictating how much data can pass through and at what speed

Latency: The time it takes for data to be transferred from its original source to its destination, measured in milliseconds

Subscriber: An individual or household that has a current internet plan with an ISP

Throughput: The rate of speed at which data travels on a network

Wiring Closet: A room in a building dedicated to housing electrical wiring and networking hardware



QoE Monitoring & Optimization for Fixed Wireless Networks

STAY CONNECTED

1-833-773-7336
hello@preseem.com

295 Hagey Boulevard
1st Floor, West Entrance
Waterloo, ON N2L 6R5



ANOTHER GREAT READ:



Learn more about how you can improve subscriber Quality of Experience (QoE) at:

preseem.com