preseem

FIXED WIRELESS NETWORK REPORT

2020 Q1 EDITION

Executive Summary

The Preseem Fixed Wireless Network Report leverages Preseem's data set to provide a view into fixed wireless networks across providers and different kinds of equipment.

Key insights from this edition include:

- Fixed wireless networks show little throughput degradation during peak which indicates that they are not heavily oversubscribed
- The average fixed wireless subscriber uses just over 4.4 Mbps when active
- The average fixed wireless subscriber uses 8.2 GB of data per day for a total of 244 GB per month
- Cambium and Ubiquiti access point equipment dominate fixed wireless deployments
- Over 60% of access points are less than 3x oversubscribed

CONTENTS

Overview

Welcome to the 2020 Q1 edition of the Preseem Fixed Wireless Network report. Each fall we release a new Fixed Wireless report with new topics and investigations. In the spring, we release an updated version of the fall report which has the most recent numbers from Preseem's rapidly growing customer base. This edition represents the Spring 2020 report.

In order to measure, analyze, and optimize QoE, Preseem collects detailed metrics on subscriber, network equipment and overall network performance from our customer base of fixed wireless Internet providers, often referred to simply as wireless ISPs (WISPs). Preseem ingests billions of metrics per-day from WISPs across the U.S., with a smaller number coming from Canadian and international markets. This report leverages this huge data pool to present a view of the fixed wireless industry across service providers and vendors.

The goal of this report is to show the real-world experience of fixed wireless subscribers, networks and equipment. As such, all the data in this report comes from fixed wireless networks.

It is our hope that this information is useful to WISPs as a way to benchmark their businesses against the wider broadband ecosystem, and that it also helps others understand fixed wireless networks.

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

Peak vs. Off-Peak

Most networks exhibit great variation in their load over the course of the day. This cycle, combined with the fact that network performance typically only degrades when the network is busy, means that simple numbers like the average rate over the day, hour or even minute are basically useless as tools to understand the subscriber experience or network performance.

Many of the metrics presented in this report are taken at "peak time." There are many simple and unsatisfactory methods to determine the peak time, such as approximating "prime time," but these methods fail to capture the 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 characterize network performance. As such, the numbers presented throughout this report aim to reflect the typical customer experience when the highest number of subscribers are trying to use the service and the performance is at its worst.

Percentiles

Throughout this report, we use several statistics to describe the data sets. These include statistics like the 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 it can be very misleading. 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 57 billion dollars. Sounds good, but this is pretty misleading.

Similarly, it's easy to see that using the maximum value as a way to summarize a data set could paint a misleading picture. The maximum value of your net worth and Jeff Bezos' is 114 billion dollars, for example.

Percentiles

Like average and max, percentiles are another tool to summarize a data set. Percentiles are 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). To calculate the 50th percentile, we first order all the elements of data set from smallest to largest to get: 1,2,2,2,3,4,4,5,5,6,100

The 50th percentile is the value at which 50% of the values 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 6th element and get the value 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

Throughput

Throughput refers to the network capacity (in bits/sec) received or sent by each active subscriber in the indicated time period.

Throughout this 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 (AP).

The figure below compares the download throughput achieved by WISP subscribers during the busiest (peak time) against 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 the busier times. Pat yourselves on the back WISPs!

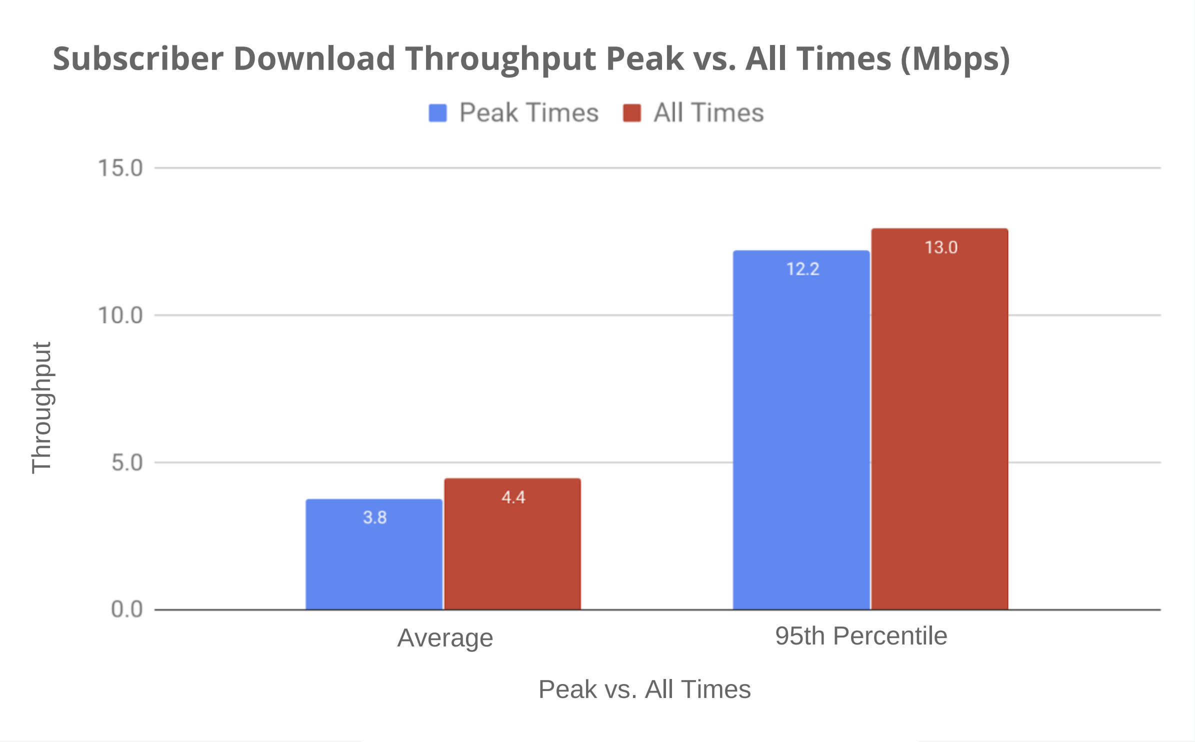

The upload throughput numbers are more confusing because most subscribers do not stress the upload direction of their connection. So while these numbers are "real" and do indicate the actual customer experience, they do not necessarily indicate the upload performance that a subscriber could achieve if desired. The 95th percentile value 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 between peak and off-peak. This indicates that WISP networks may be more congested in the upstream during peak than is commonly assumed. While subscribers care most about download performance, it is important to note that a congested upload path can cause download throughput problems because of TCP acknowledgement starvation. Similarly, a congested upload path can cause packet loss which is also bad for the customer experience.

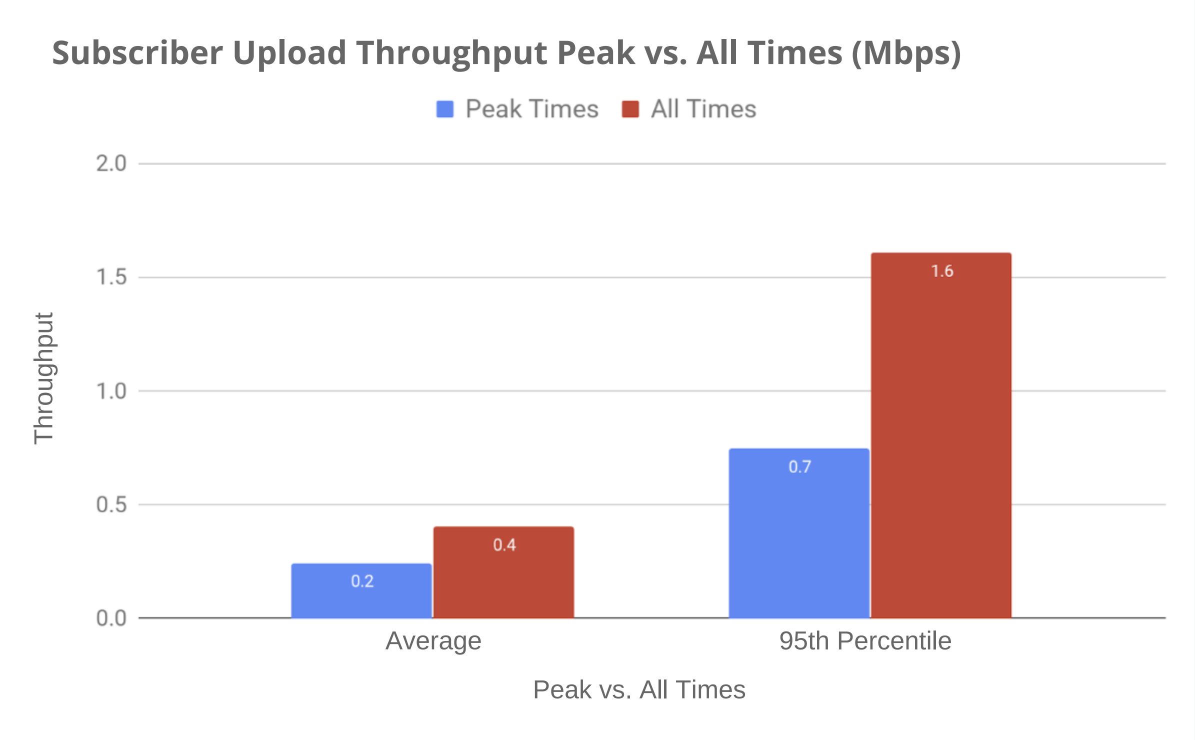

Latency

Preseem measures latency by tracking the round trip time for individual TCP segments to obtain a detailed view of the latency in the access network. This approach results in thousands of latency samples per second, per subscriber. This is fundamentally different than 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. Therefore, it is quite likely that the latency in networks without such optimization is significantly higher.

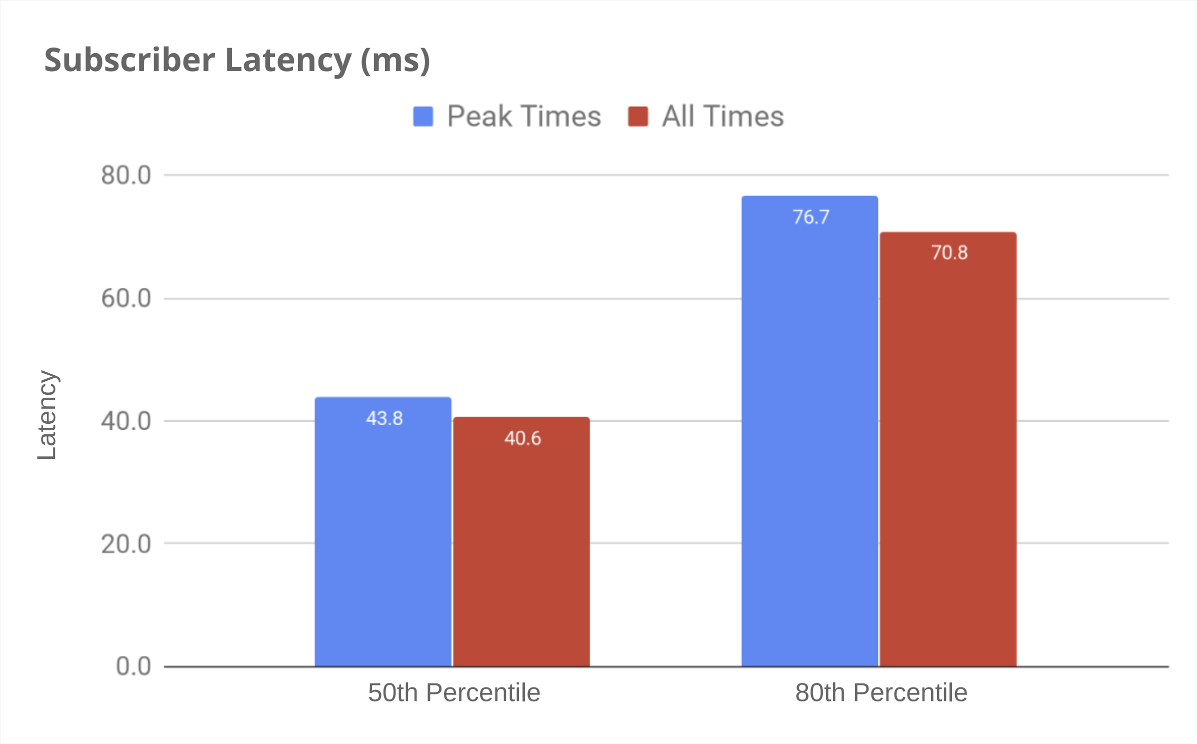

Where Does Latency Come From?

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 simply 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 crazy stuff like buy shorter fiber cables).

Queuing delay: Queuing 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 has 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 is a buffer that is 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 queuing delay through active queue management (AQM) techniques that greatly improve the subscriber quality of experience (QoE), even when the network or the subscriber's connection has reached its capacity.

Frame aggregation: In a sense, this is a type of queuing. However, due to its prevalence in wireless networks, it's worthwhile to discuss separately. In order 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 WISP subscriber over the day or month. From the perspective of the subscriber experience, the total usage isn't very instructive since a large amount of usage consumed during off-peak has less of an impact on the perceived network quality compared to a smaller amount of usage during peak.

The average subscriber download usage for WISP subscribers is 8.8 GB/day, or 244 GB/month. As expected, the average hides the significant variation that occurs between subscribers.

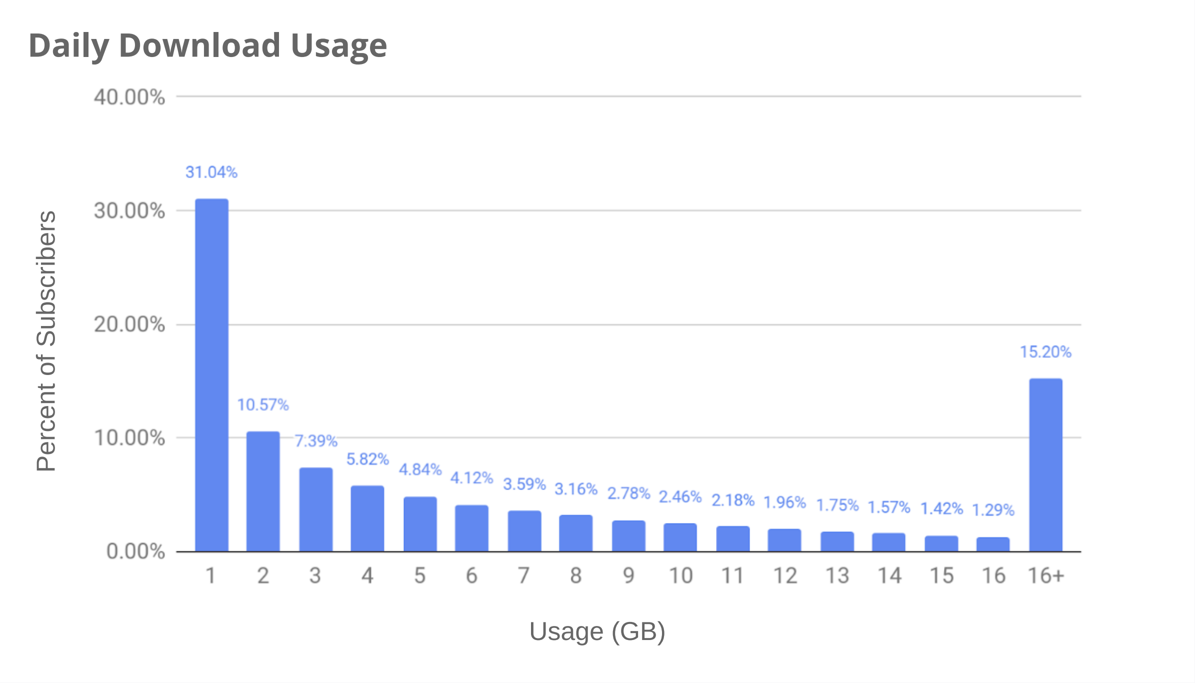

Here we see that just over 31% of subscribers use less than 1 GB of download usage per day, while over 15% use more than 16 GB per day.

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.

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 access point 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 access point models.

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 and 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 access points.

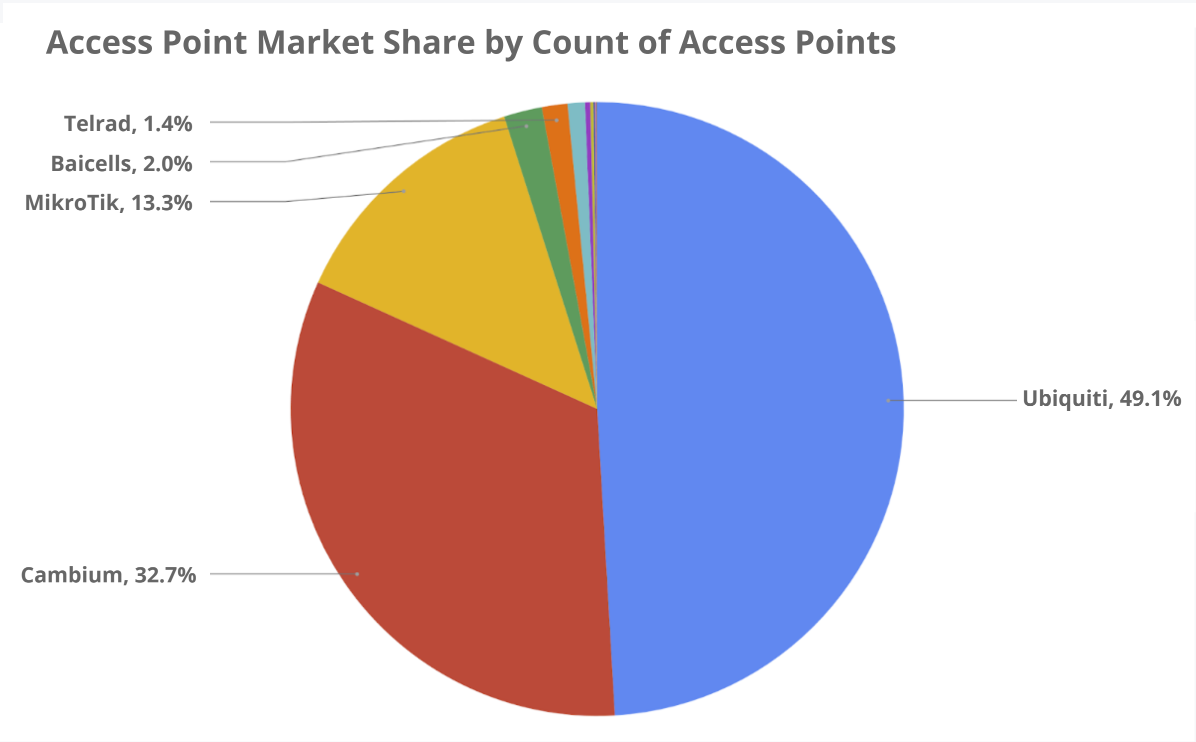

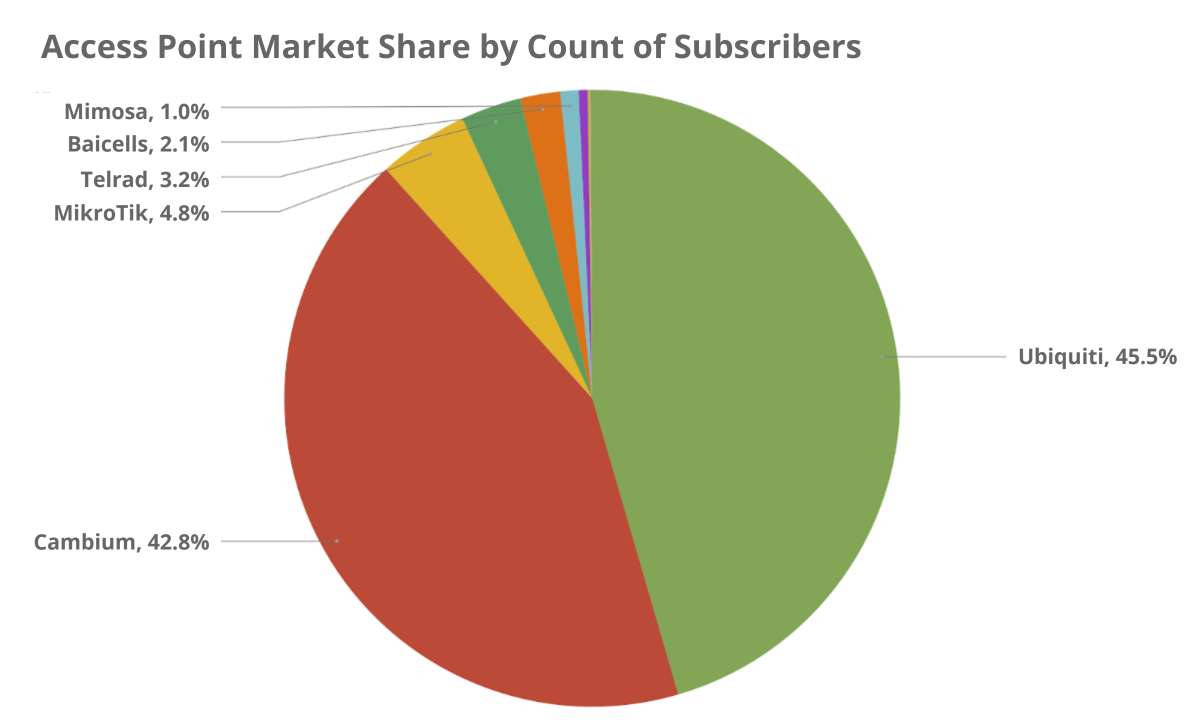

The market share gets a lot more complicated when the access point model is added to the picture. Comparing access point market share by device count and subscribers clearly shows the particular models that support a large number of subscribers.

The PMP 450m has 1.7% of the access point market share by element count, but 9.79% by subscriber count.

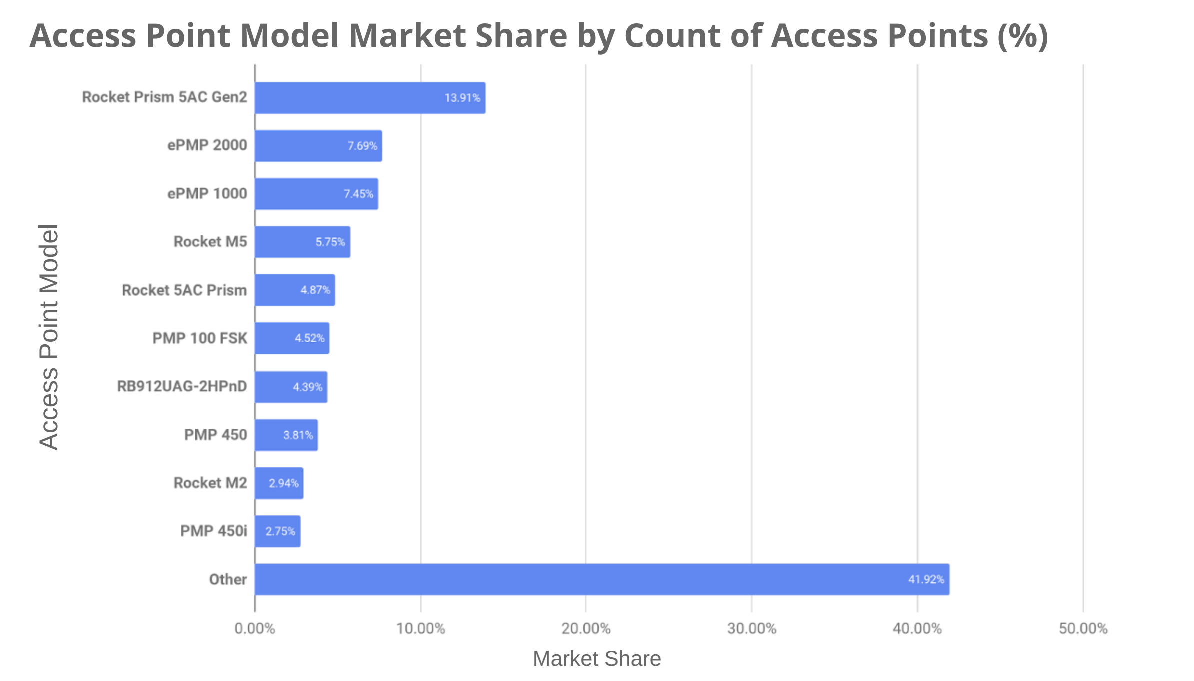

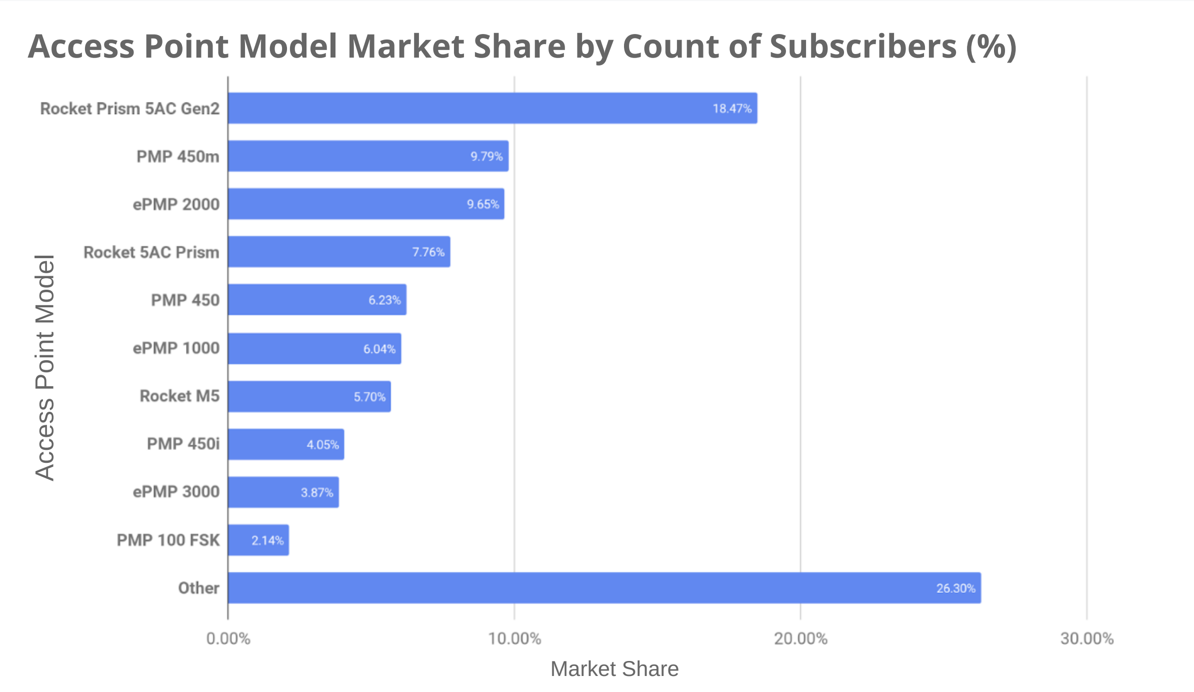

Connected Subscriber Count

As one might expect from the difference in access point market share by element and by subscriber count, there is significant variation in how many subscribers WISPs connect to individual access points. The chart below shows the number of access points bucketed by the number of subscriber radios connected to the access point.

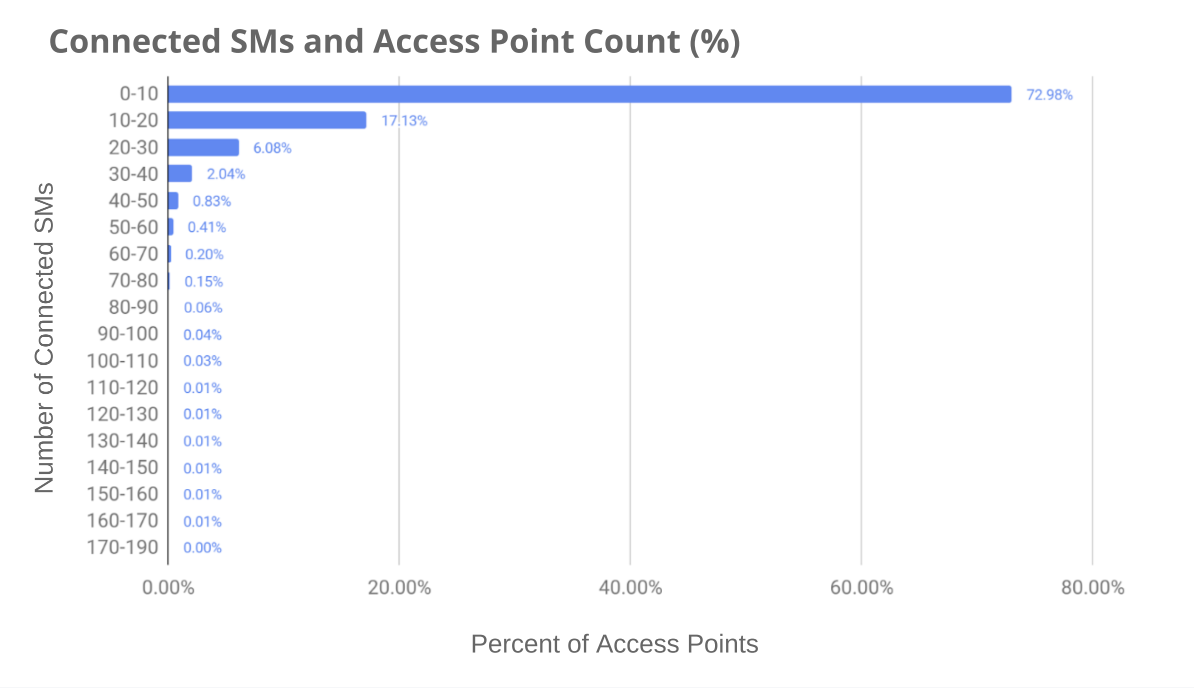

Surprisingly, the data shows that over 70% of access points have 0-10 subscribers attached. Breaking this down by a few of the top access point models paints a different picture.

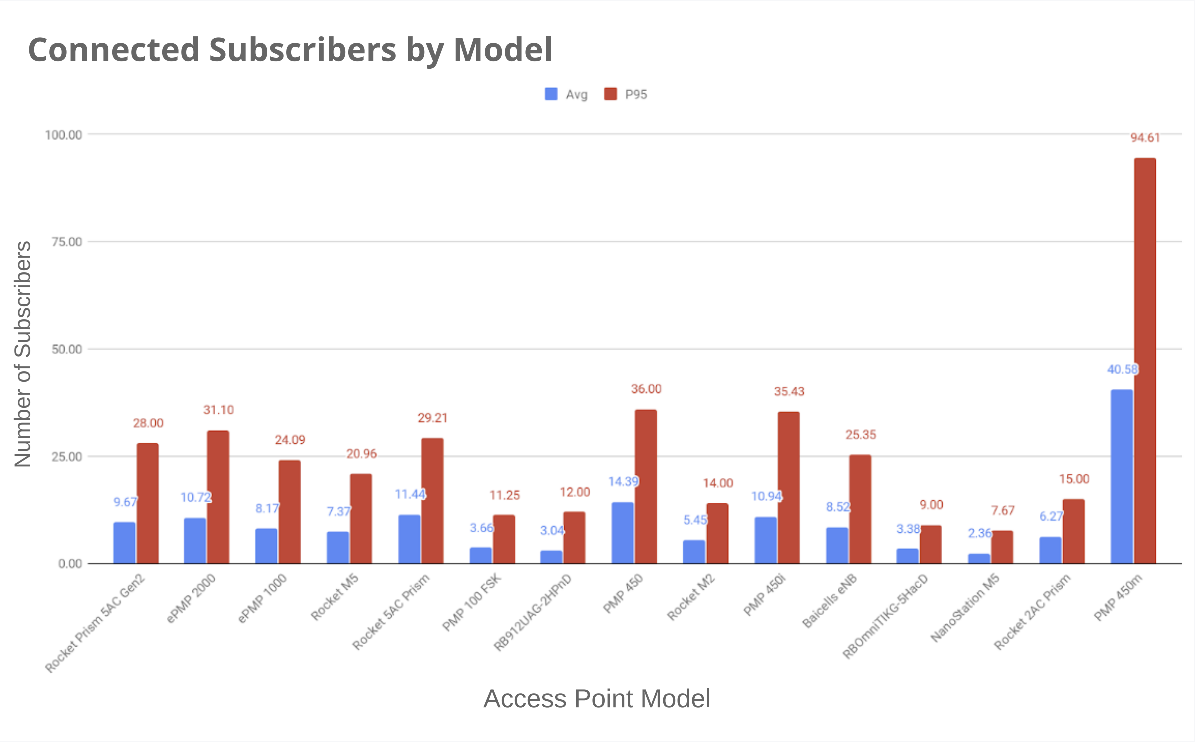

Here we see that some access point models are typically deployed with many more subscribers than the overall average. For example, PMP 450m deployments average almost 41 connected subscribers.

Throughput

Since measuring throughput outside of the peak times provides little insight into the subscriber experience—because the network is not loaded—the throughput numbers below were taken from the busiest times of the day for each individual access point.

Download Throughput

The data shows that, in terms of the number of access points, almost 39% of deployed access points deliver less than 10 Mbps of real-world throughput. This is a somewhat surprising result, but again, looking at the more modern equipment shows a very different pattern.

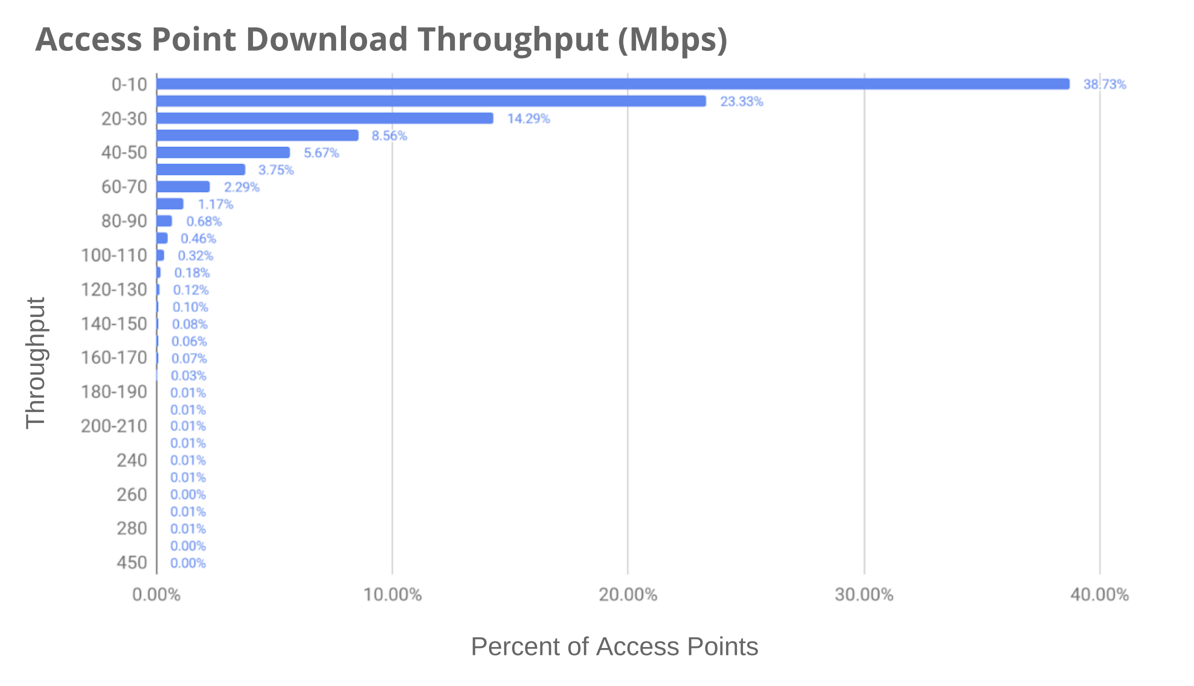

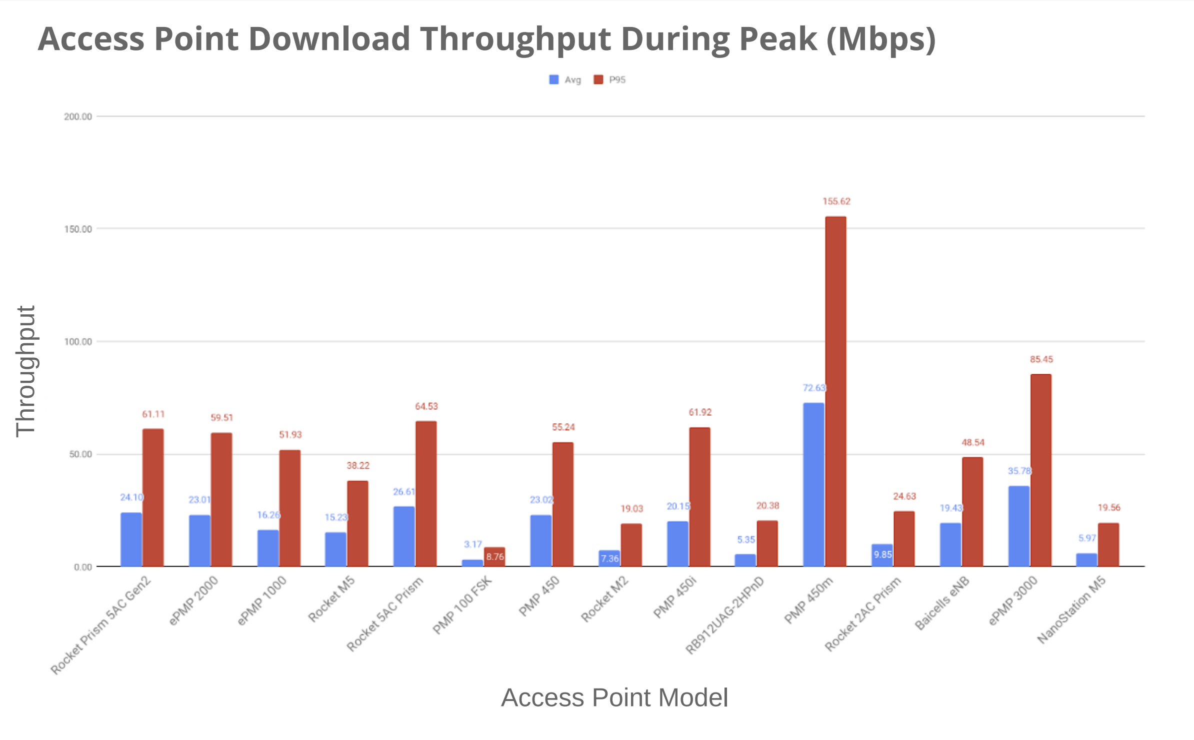

Upload Throughput

Characterizing upload performance is more difficult due to the fact that 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 WISP access points deliver less than 10 Mbps of upload throughput during the times of the day with the highest demand.

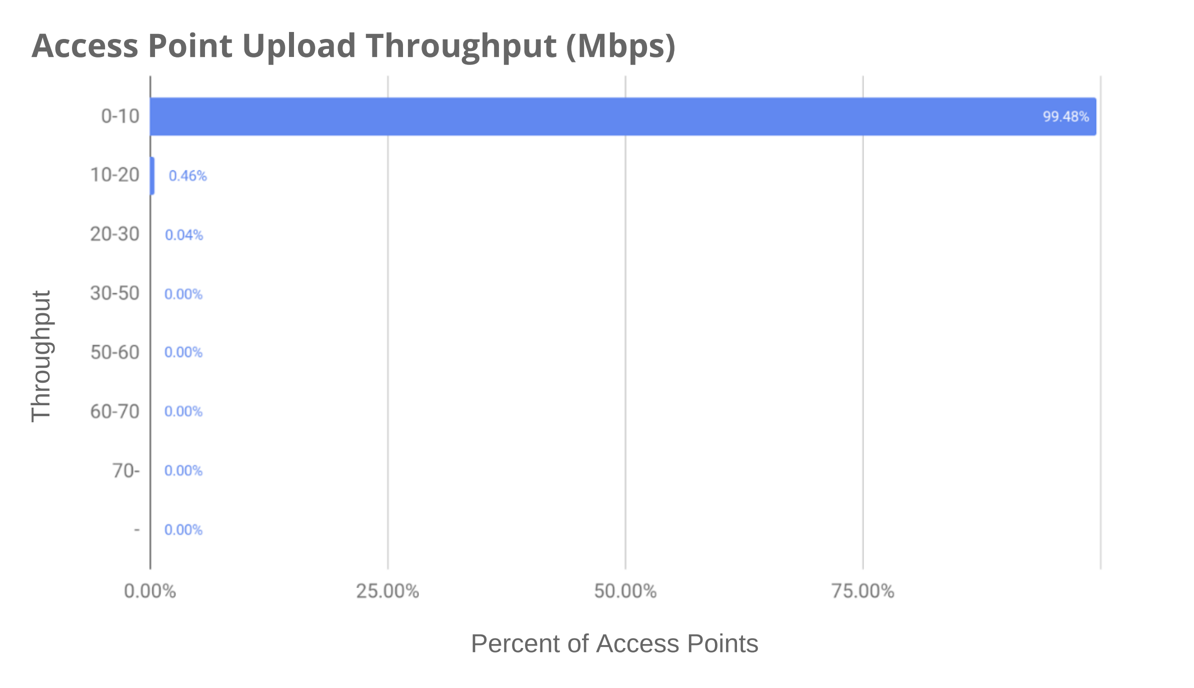
Access Point Upload Throughput (Mbps)
Throughput
0-10
10-20
20-30
30-50
50-60
60-70
70-
-
99.48%
0.46%
0.04%
0.00%
0.00%
0.00%
0.00%
0.00%
0.00%
25.00%
50.00%
75.00%
Percent of Access Points

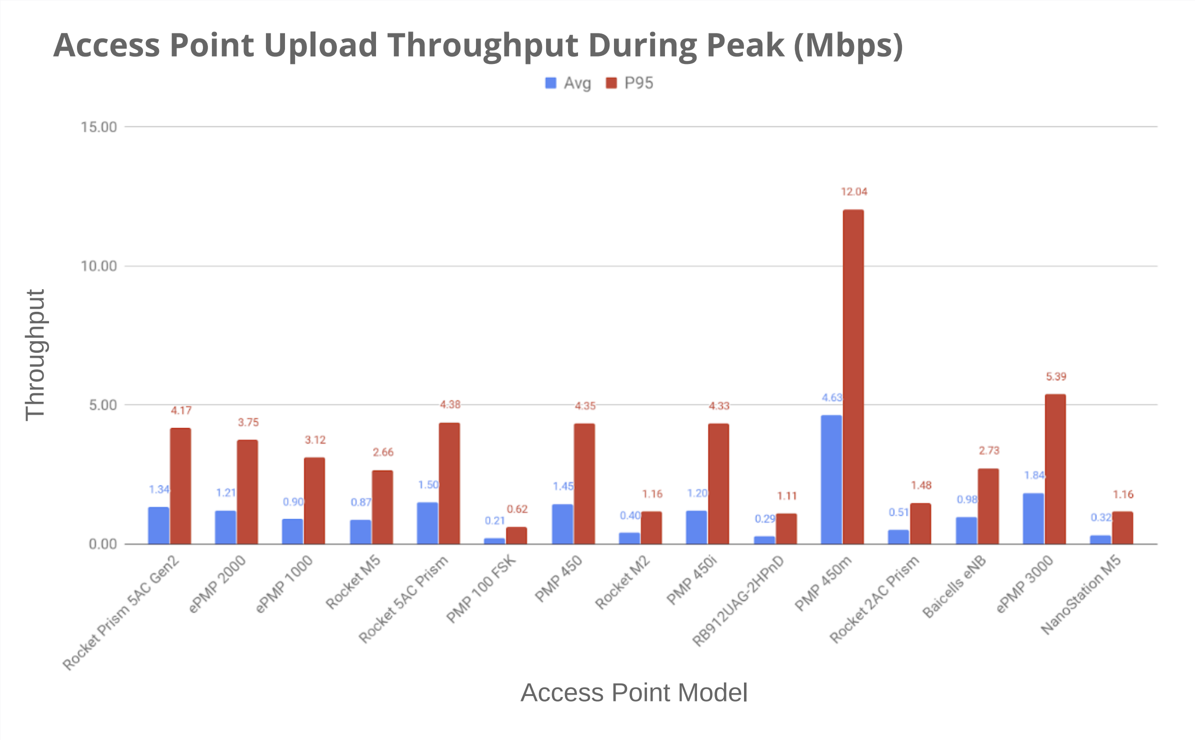
Access Point Upload Throughput During Peak (Mbps)
Avg
P95
Throughput
15.00
10.00
5.00
0.00
Rocket Prism 5AC Gen2
ePMP 2000
ePMP 1000
Rocket M5
Rocket 5AC Prism
PMP 100 FSK
PMP 450
Rocket M2
PMP 450i
RB912UAG-2HPnD
PMP 450m
Rocket 2AC Prism
Baicells eNB
ePMP 3000
NanoStation M5
Access Point Model

Latency

Across all access points, the distribution of latency follows an interesting pattern. Most access points deliver service with less than 100 ms of latency during peak times, but a significant number are over that benchmark.

What Is A Good Latency Value?

Latency requirements differ a lot by application. High latency has very little effect on Netflix but has a large impact on gaming, for example. 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.

Note that the values shown in this report reflect 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.

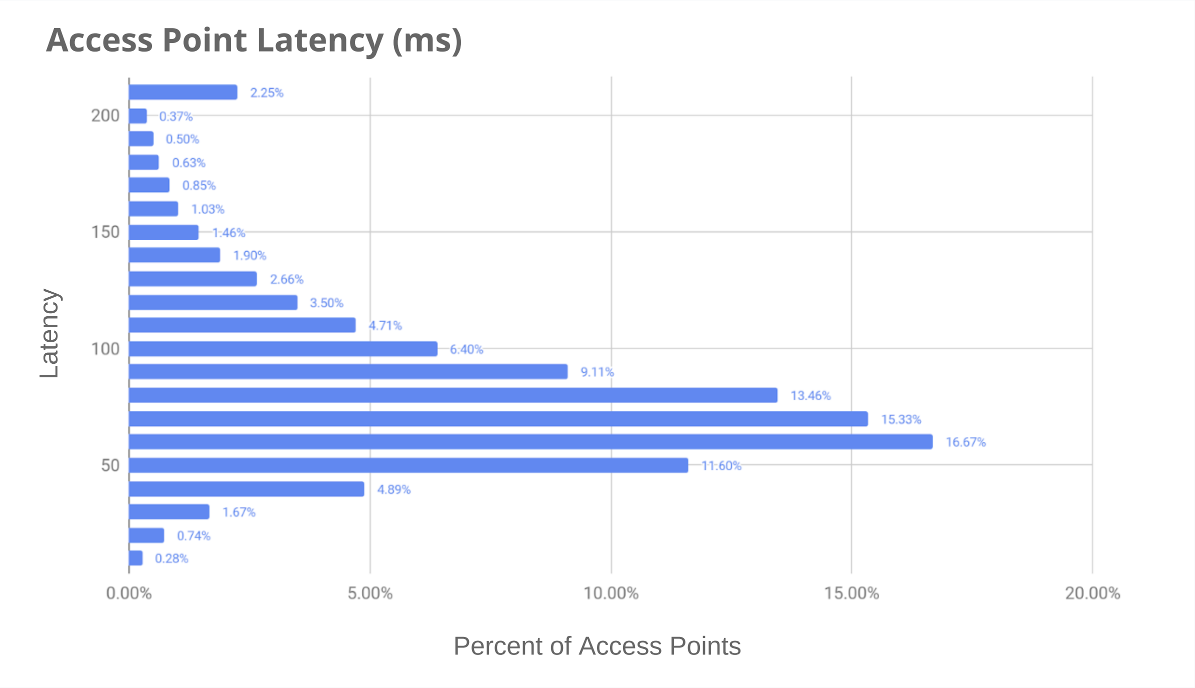

Comparing latency across access point models shows significant variation and a general trend towards newer AP models having better latency characteristics.

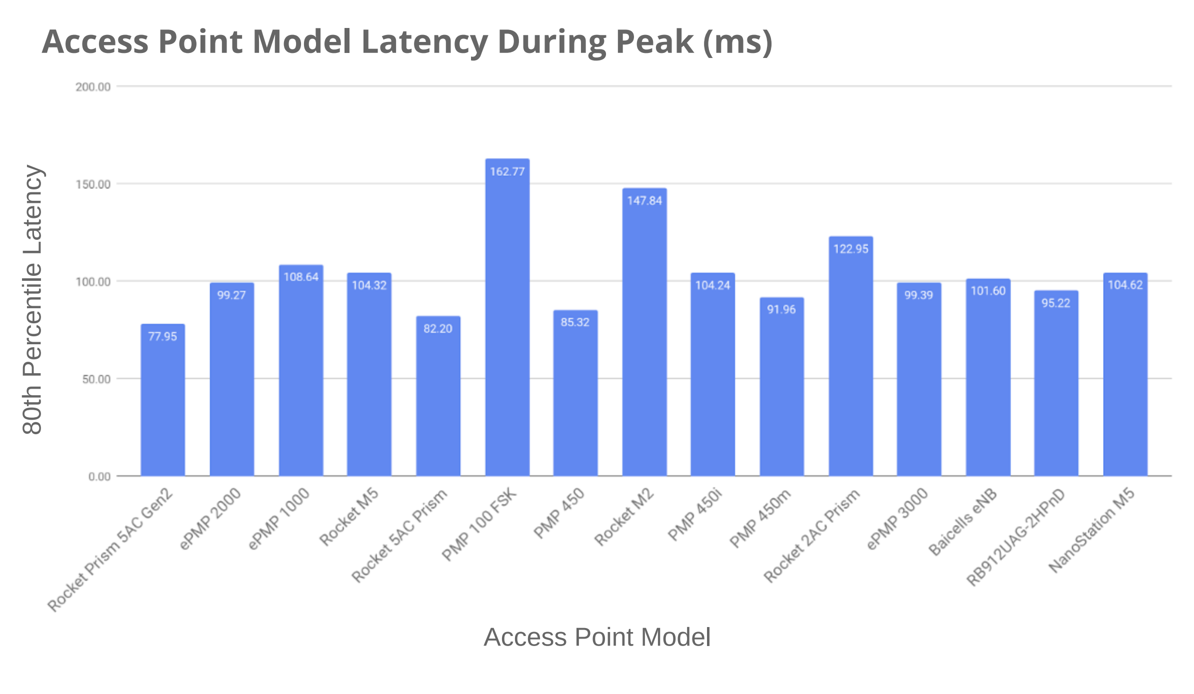

RF Channel Width

Besides the 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 RF Channel Width

Across all of the access points that Preseem monitors, about 67% use 20 MHz channels, with 10 MHz and 40 MHz channels being the next most common. A future edition of this report may compare channel width over time to get a sense of where the industry is going.

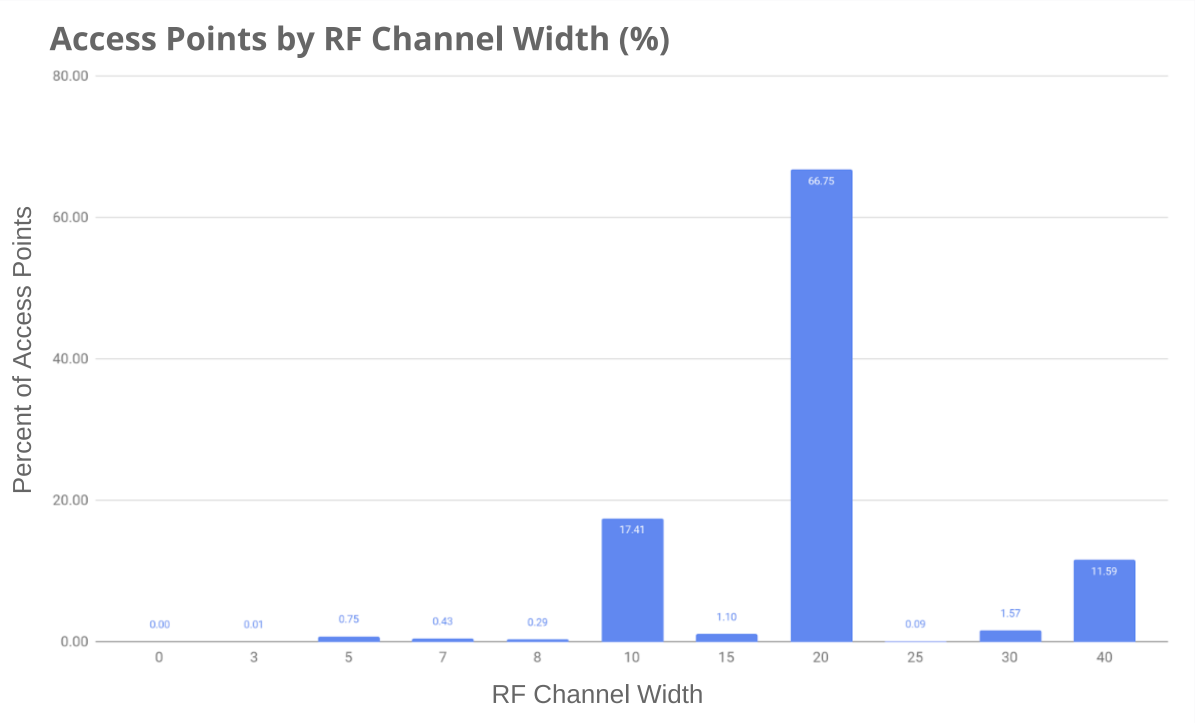

Download Rate by RF Channel Width

By looking at all APs aggregated by channel width, we can see that, while a larger channel width does increase throughput, the effect is not always equal to the increase in the channel size. For example, going from a 10 MHz channel to a 20 MHz channel increases throughput from 20.22 Mbps to 53.67 Mbps, or 165%. However, the increase from a 20 MHz channel to a 40 MHz channel results in a throughput increase from 53.67 to 89.53, or 66%.

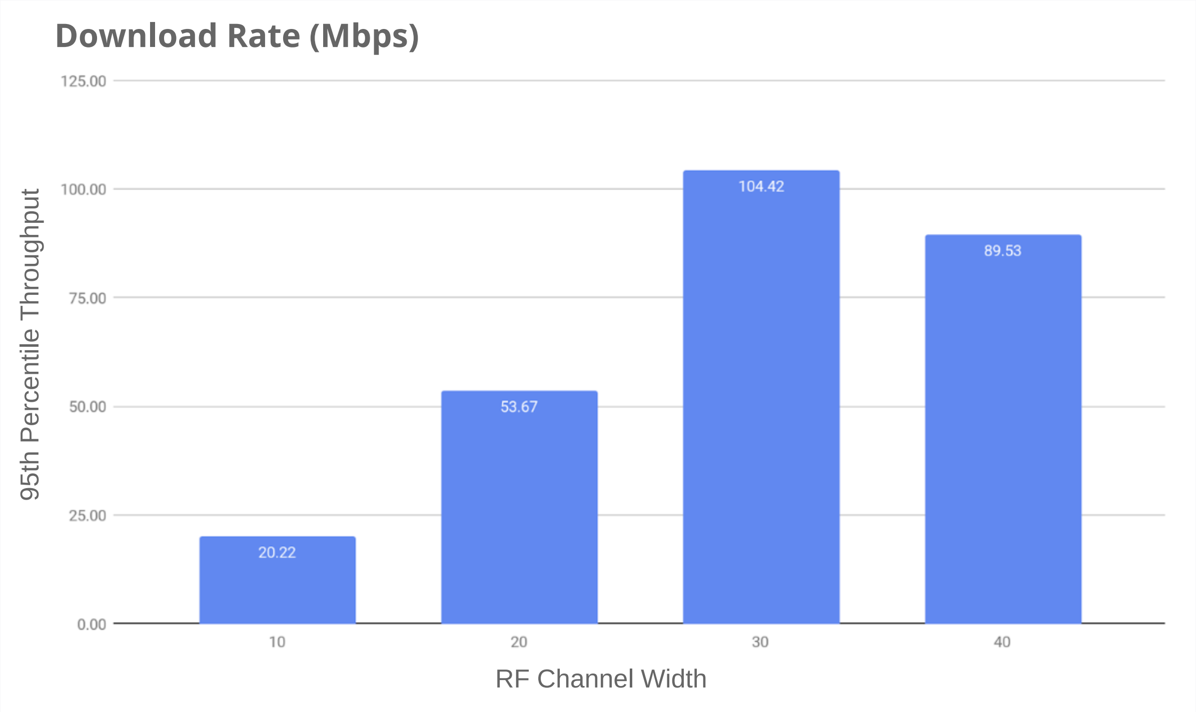

As usual, the results are more nuanced when looking at specific AP models as compared to the overall aggregation. However, the trend still holds that doubling channel width does not double throughput.

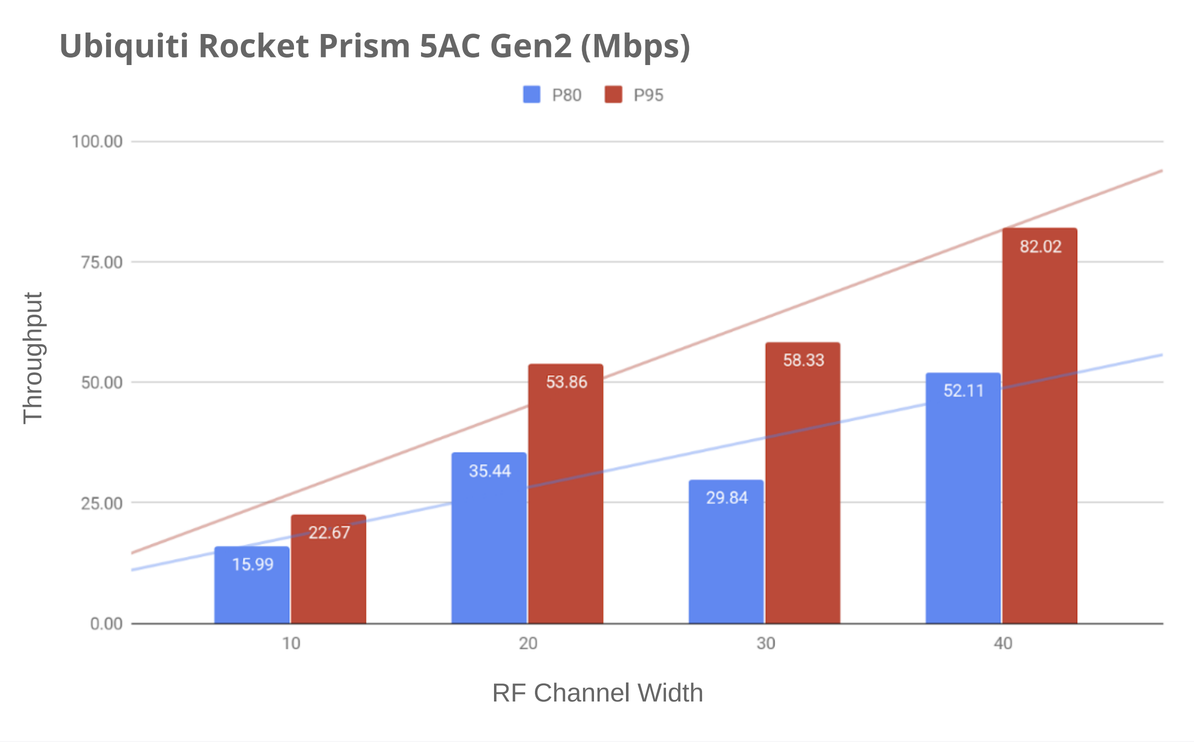

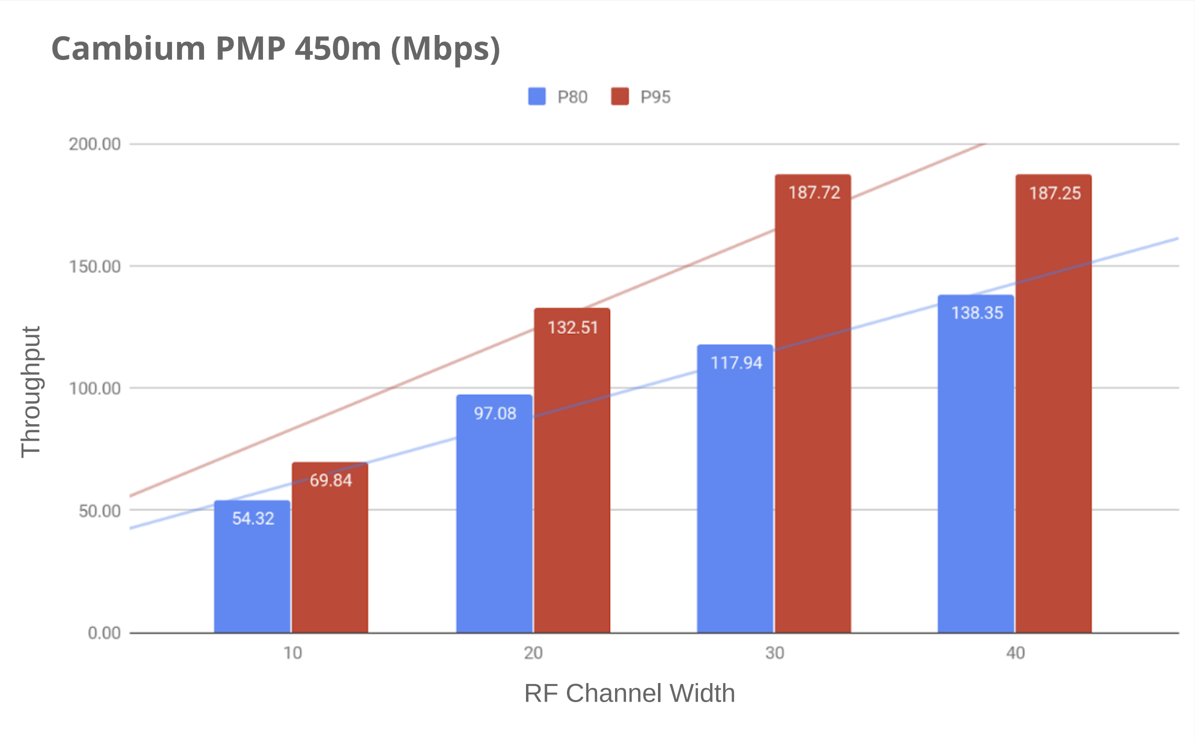
Cambium PMP 450m (Mbps)
P80
P95
200.00
150.00
100.00
50.00
0.00
Throughput
54.32
69.84
97.08
132.51
117.94
187.72
138.35
187.25
10
20
30
40
RF Channel Width

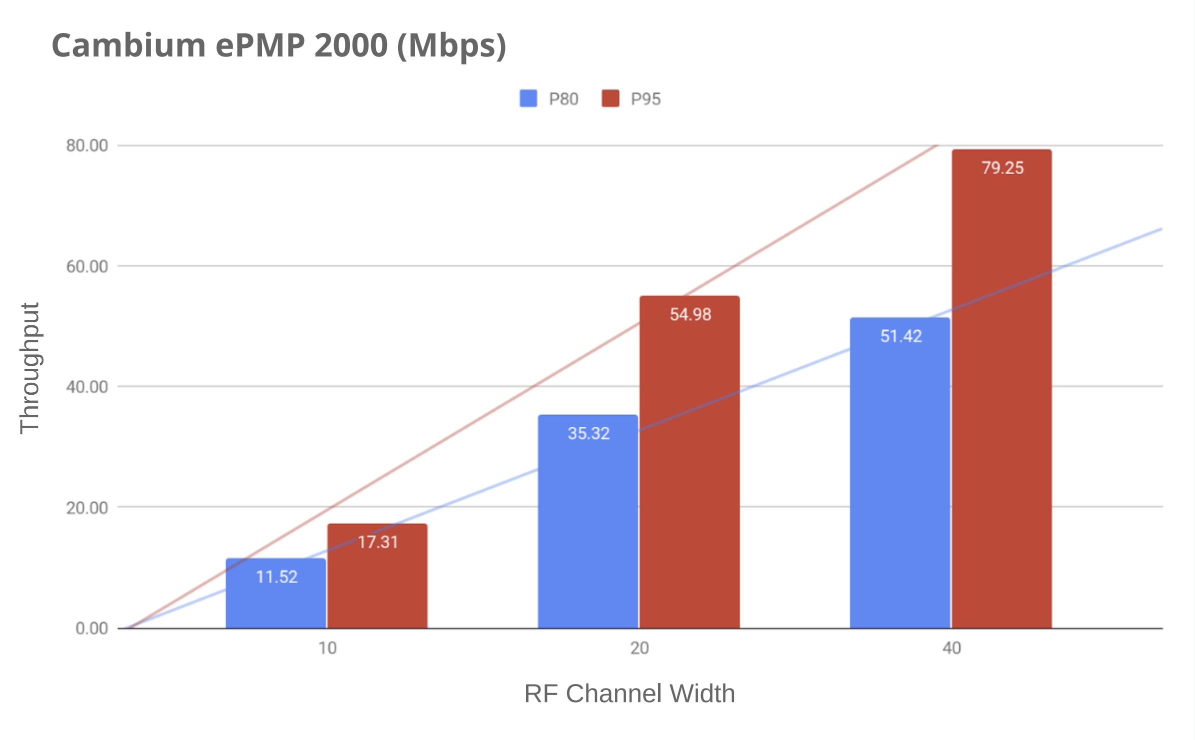
Cambium ePMP 2000 (Mbps)
P80
P95
80.00
60.00
40.00
20.00
0.00
Throughput
11.52
17.31
35.32
54.98
51.42
79.25
10
20
40
RF Channel Width

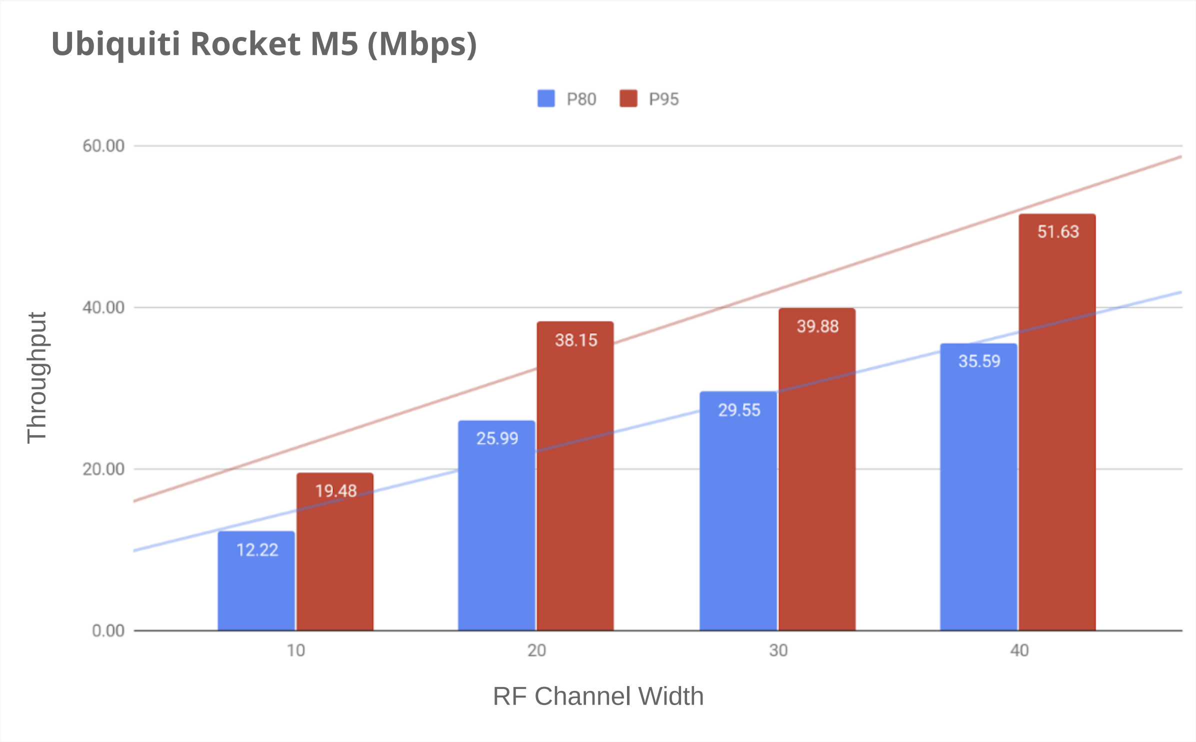

Connected SMs by RF Channel Width

The following chart shows the number of connected subscriber modules (SM) by channel width across all access point models.

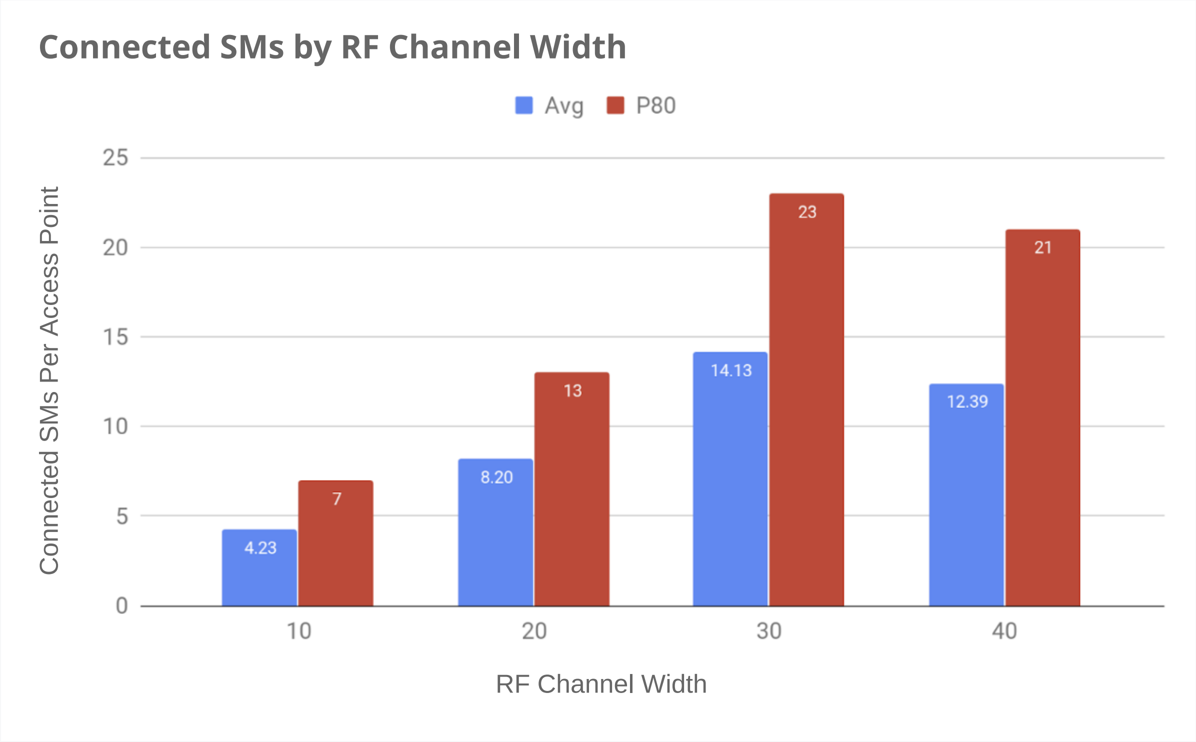

For both the 40 MHz and 30 MHz channel cases, the number of connected SMs is greater than the number of SMs connected to access points using 20 MHz channels.

Subscriber Throughput by RF Channel Width

Of course, the number of connected subscribers is only part of the picture as it doesn't relate directly 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 WISPs use wider channels to handle more subscribers instead of delivering higher per-subscriber throughput.

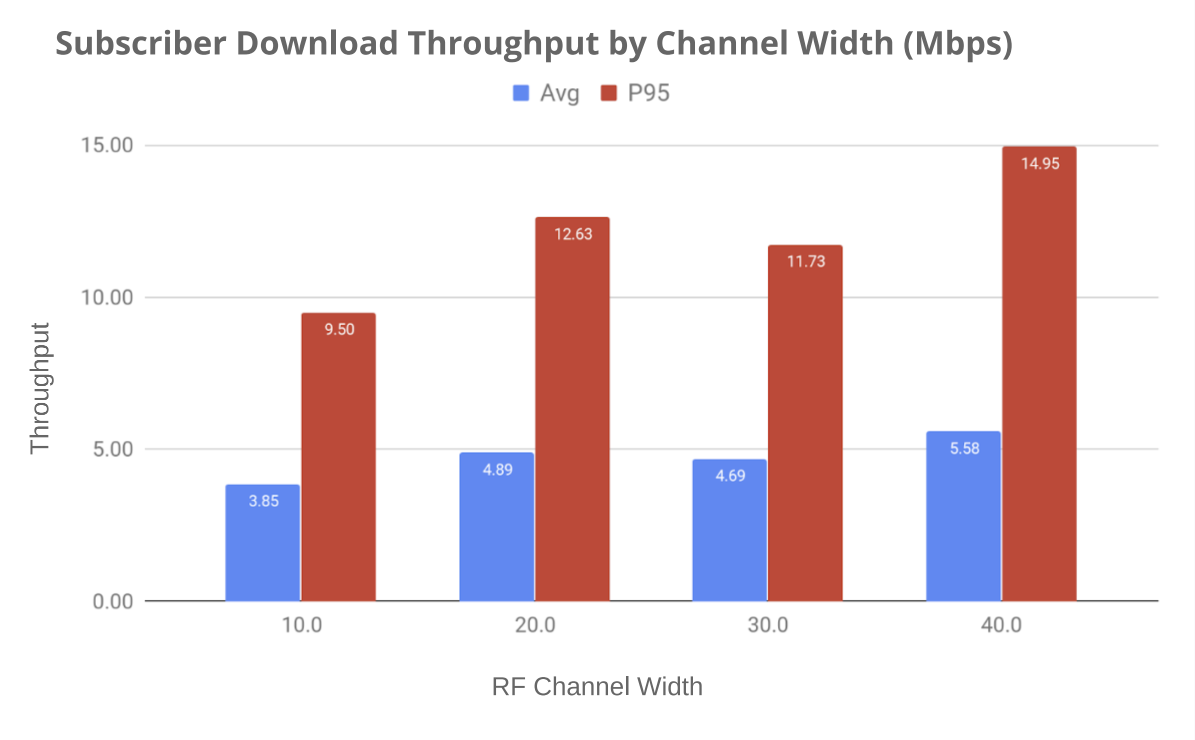

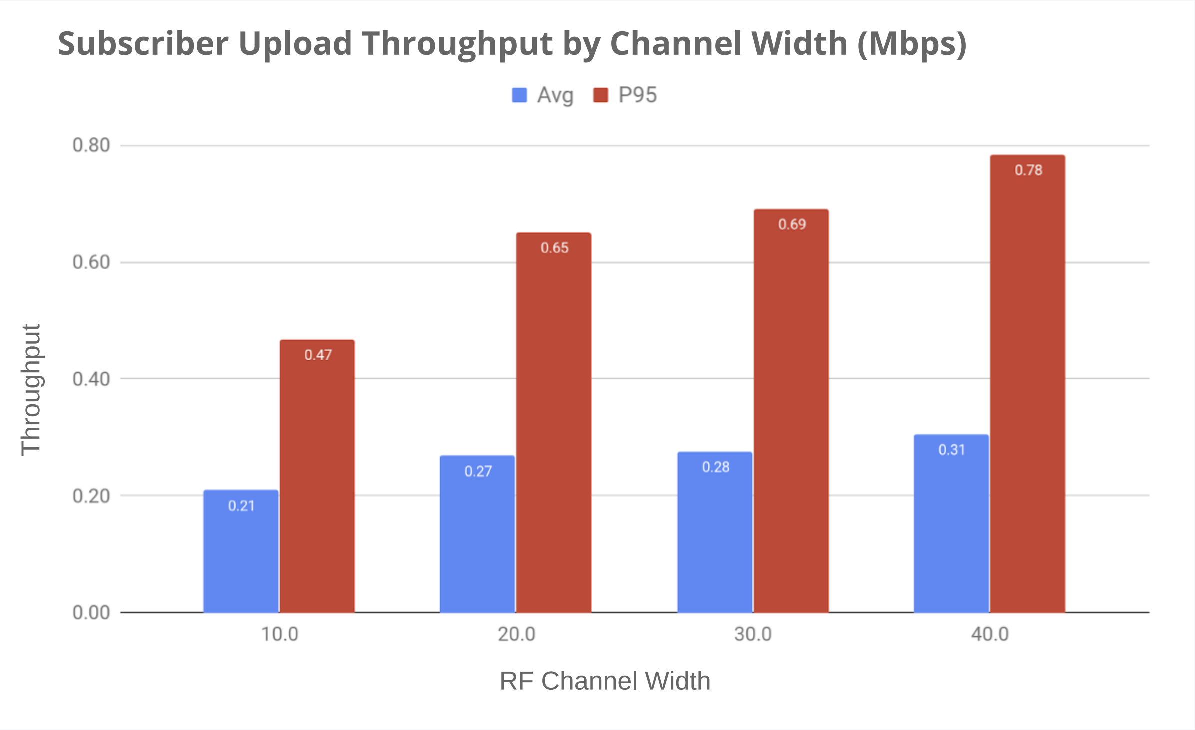

Oversubscription Ratio

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 just doesn't work.

The next chart provides some insight into how much WISPs 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 the WISP has sold twenty, 10 Mbps plans on an AP that typically achieves 50 Mbps, then the oversubscription ratio is:

(20 x 10) / 50 = 4

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.

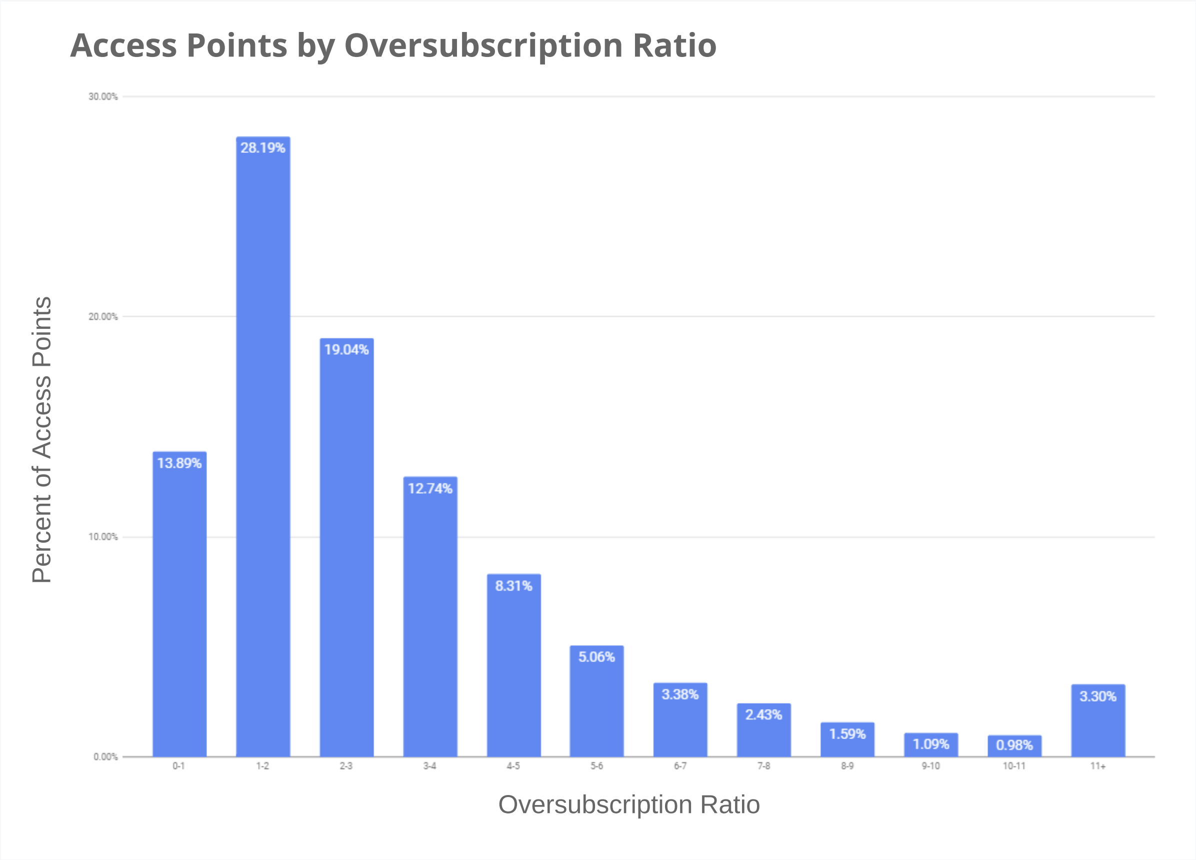

This chart shows that over 28% of APs are between one and two times oversubscribed. As can be seen in the model-specific charts below, the oversubscription distribution is relatively stable across many models, with the PMP 450m being an outlier. This suggests that the 450m supports a much higher level of oversubscription than these other models.

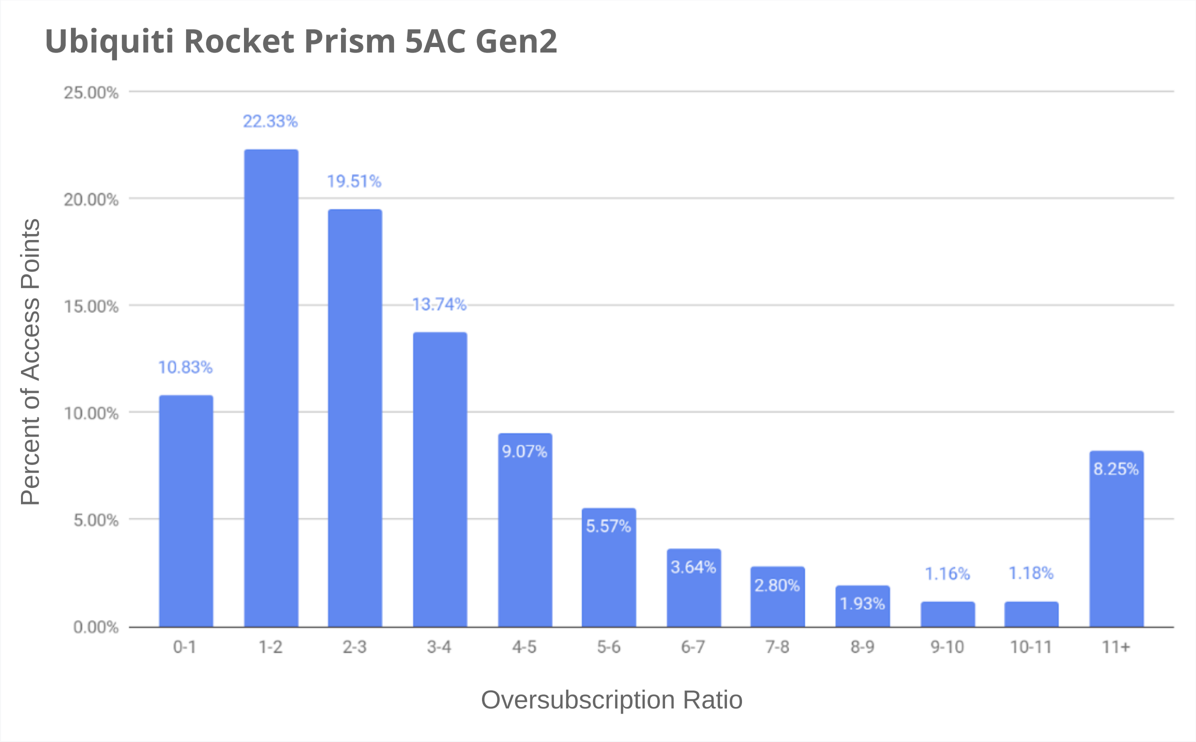

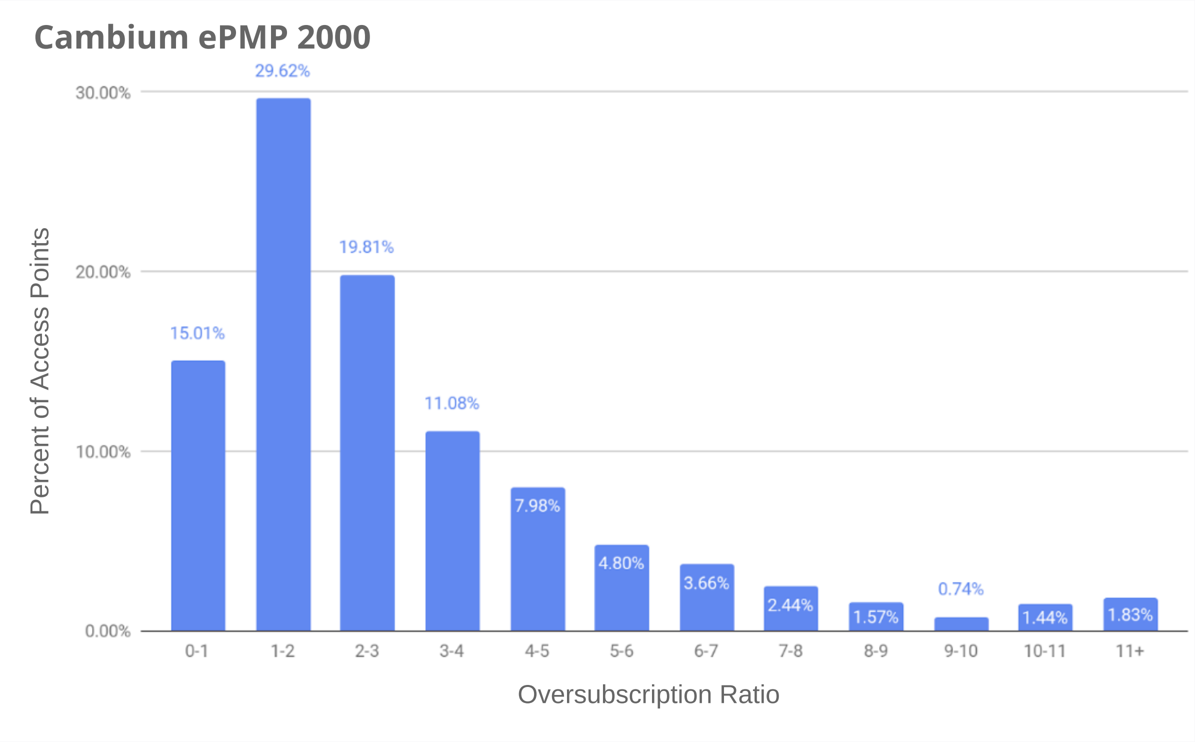
Cambium ePMP 2000
Percent of Access Points
30.00%
20.00%
10.00%
0.00%
15.01%
29.62%
19.81%
11.08%
7.98%
4.80%
3.66%
2.44%
1.57%
0.74%
1.44%
1.83%
0-1
1-2
2-3
3-4
4-5
5-6
6-7
7-8
8-9
9-10
10-11
11+
Oversubscription Ratio

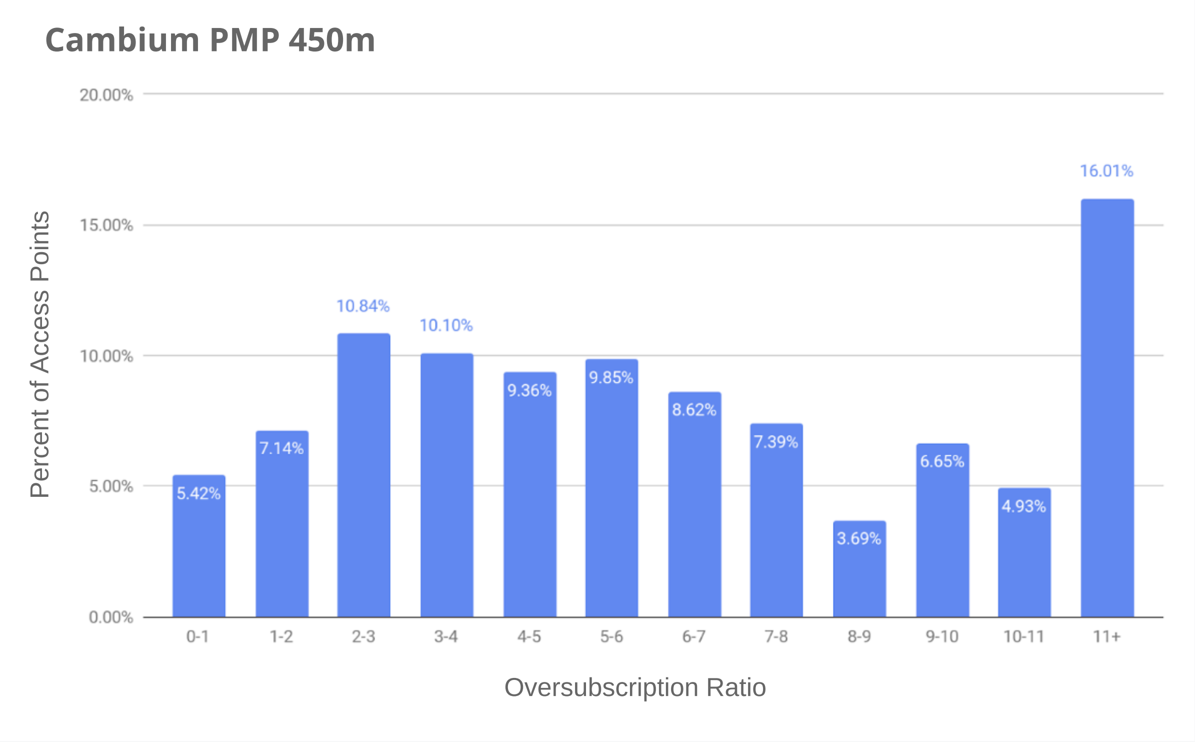
Cambium PMP 450m
Percent of Access Points
20.00%
15.00%
10.00%
5.00%
0.00%
5.42%
7.14%
10.84%
10.10%
9.36%
9.85%
8.62%
7.39%
3.69%
6.65%
4.93%
16.01%
0-1
1-2
2-3
3-4
4-5
5-6
6-7
7-8
8-9
9-10
10-11
11+
Oversubscription Ratio

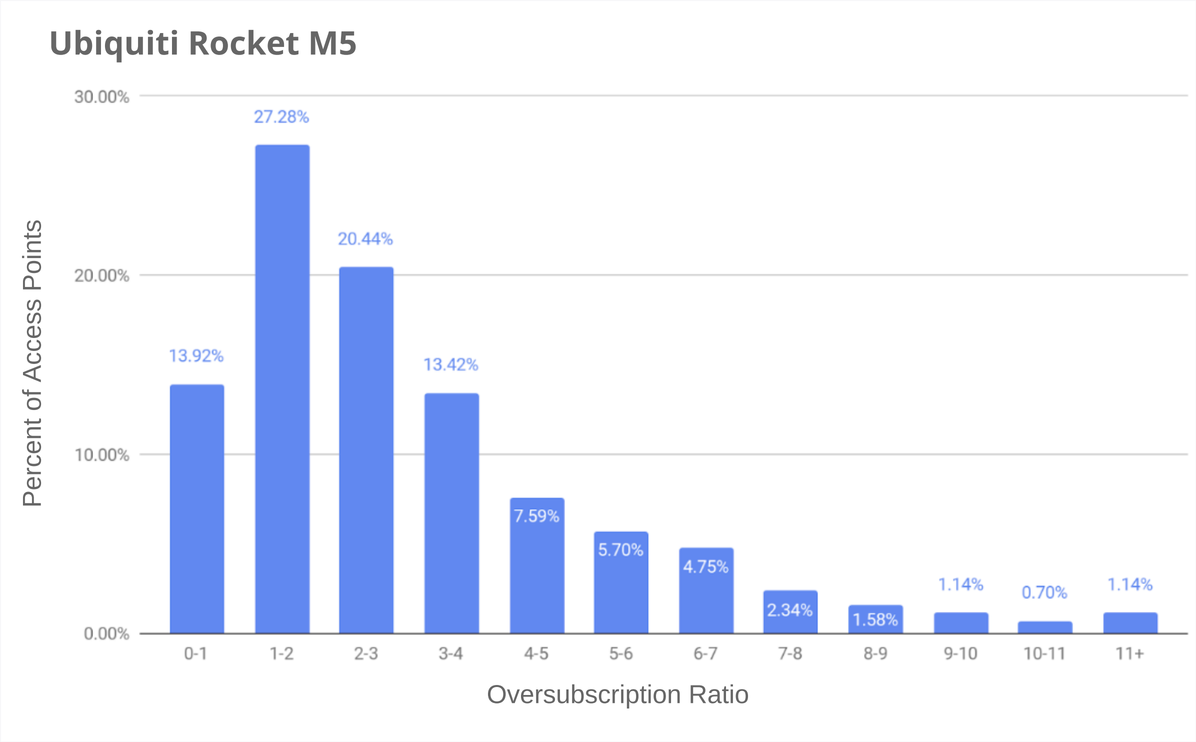

Oversubscription in and of itself is not bad. However, higher oversubscription ratios generally lead to lower subscriber throughput during peak, and therefore a poorer subscriber experience.

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 WISP access points have less than 10 attached subscribers, and that the majority of access points are less than three times oversubscribed. Other interesting results include:

- Fixed wireless networks show little throughput degradation during peak which indicates that they are not heavily oversubscribed
- The average fixed wireless subscriber uses just over 4.4 Mbps when active
- The average fixed wireless subscriber uses 8.2 GB of data per day for a total of 244 GB per month
- Cambium and Ubiquiti access point equipment dominate fixed wireless deployments
- Fixed wireless networks leverage wider channel width to deliver service to a higher number of subscribers instead of delivering higher throughput to the same number of subscribers per access point
- Over 60% of access points are less than 3x oversubscribed

preseem

QoE Monitoring & Optimization for
Fixed Wireless Networks

Connect with us

/preseem /preseem @preseem preseem.com 1-833-PRESEEM