WAN And LAN Speed and Service Matching – Are We Engineering It Correctly for Consumer Services Growth In The Next 5+ Years

A Technical Paper prepared for SCTE by

J.R. Flesch
Director, Advanced Technology
CommScope
3871 Lakefield Drive, Suwanee, GA 30024
Jr.flesch@commscope.com

Kurt Lumbatis, CommScope

Ian Wheelock, CommScope

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Introduction

In the monotonic drive for ever-greater bitrate consumption posed by the increased variability and sheer number of home dweller-facing digital network services (particularly immersive experiences and IoT-anchored emergency services), a review of the balancing act between WAN and LAN capacity and responsiveness begs execution. Such review seems particularly warranted at this point in the cable distribution and home Wi-Fi development epochs, given that the recalculation of transmit and receive capacities posed by the DOCSIS 4.0 WAN upgrade nearly coincides with the adoption of in-home Wi-Fi 7 LANs (themselves also potentially enhanced by higher EIRP /AFC-enabled Standard Power).

This paper will examine the implications of emerging network demands in a post-pandemic home services environment and attempt to evaluate the impact to WAN and LAN infrastructure; identifying behavioral, economic, or capacity mismatches which may alter the scope and rate of proposed network evolutions. What seems solicited is an advisory to network stakeholders on the scope of WAN and LAN futureproofing and a weighting of appropriate investment required to adequately meet anticipated mounted service connection needs in a moderately near-term adoption horizon.

Note that the exclusion of OTT WAN inroading is an intentional artifact of the constrained scope of this paper, but such competitive dismissal assumes minimal impression (consumer exploitation) in the 5-year consideration window. A brief sidebar is included for future reference and study, however, given the threat's higher global accessibility than fiber (PON), modest capability and cheaper entry ticket.

2. Bitrate and Latency, Averages and Peaks

Both the network WAN and the in-home LAN exhibit effects of mixed-flow data stochastics, though with the obvious contrast of WAN versus LAN data endpoint counts (clients and upstream sources). The WAN's high end client count and dedicated multicast video channel assets guarantee that its carry-to-capacity during peak use exhibits a higher average ratio than the bursty (though still largely video-based ABR data consumption) associated with adaptable-rate Wi-Fi traffic bound for on-premises distribution. In fact, in terms of raw bitrate capacity, Wi-Fi airtime capacity is typically leveraged to a much smaller fraction of its upper limit in the average US household (which is obviously not the case with the wireline WAN, especially during peak use periods). This begs questions around the required number of contending Wi-Fi clients (or competing in-range, interfering APs) which would render the LAN incapable of meeting its QoS aspirations – especially given next generation Wi-Fi MAC market penetration.

From the wireline WAN standpoint, DOCSIS 4.0 (in one or more of its species) promises improved duplex signaling capacity via modified exploit of the wireline's diplex split (including dynamic management of same in FDX instantiations), jettisoning video's SC-QAM in favor of full OFDM leverage and ESD (initially to 1.8 GHz) to accompany higher spectral efficiency modulation schemes. All these improvements promise higher bitrate support in both directions over the WAN.

2.1. The Wi-Fi (LAN) environment

The summary standard of on-premises network performance is Wi-Fi bitrate, and its singular comparative assessment is "the higher, the better". And while raw delivery speed makes up for many signal distribution sins, it is worthwhile to recognize that data movements between Wi-Fi endpoints lever orchestrated solicitations of clear airtime in one or more radio bands to guarantee transfer of packets, the integrated view of which amounts to flows of varying bitrate and waypoint latencies. Contention in the airtime space – especially from native-mesh-asynchronous radio sources – compounds the access issue and confounds the negotiating parties; this serves to reduce delivery determinism and inject noise into the packet goodput in the form of varying latency (jitter). In pathological cases, it can result in intermittent service interruption (if not outright failure) for any given link.

Fundamentally, Wi-Fi has evolved to include methods for framed exchange (RTS, CTS) and priority handling (WMM with its varying backoff counts) and now with the advent of MAC 7, features much broader channel BWs, higher spectral efficiency (denser QAM), cross-band channel aggregation, bi-directional MIMO and inherits OFDMA, BSS coloring and TWT from MAC 6, all of which serve to give its scheduler more tools to rapidly – and with low latency and higher dependability – transfer packets across endpoints even under challenging interference conditions. The goal, of course, is "responsive immersion" for end users in cases where content is being consumed or interacted with; and immediate cloud network connectivity for those services which operate on emergency care or home security tiers within a given residence.

Wi-Fi responsiveness has two components – bitrate and latency. The twitch gaming community (and simulation scenarios in general) have drawn attention to roundtrip command processing latency and the specific requirement of predictable, low jitter – given that now the issue of raw data pumping capacity has moved well past a 25x surfeit of capability for any mounted service in the home wireless domain. In other words, since we seem to be able to burst data to/from endpoints with low enough airtime duty cycle, it now becomes more critically important to remove queuing shims from the transfer process and more appropriately address tiers of delivery priority for the packets themselves.

On the Layer 2 front, the typical resolution for latency issues comes in the form of different priority queues (with delivery deferral of, and pre-empted delivery for, lower versus higher priorities respectively), clever multi-receiver packing of right-sized data payloads (OFDMA) and hold timeouts (buffer truncations) capping the maximum dwell time for lined-up packets. The Layer 1 mitigation involves identifying surplus airtime in underutilized Wi-Fi channels (in bands mutually supported by the attaching endpoints). Further, these types of solutions may be reactive (algorithmic responses to particular link telemetric triggers like airtime capacity or buffer lengths) or anticipatory (via managed pre-emptive adjustments given TOD and stored prior successes). Condensed, the Layer 1 consideration involves the capacity calculus of the given band and channel (so, typically, MCS x BW x SS) versus any steering options on the band/channel front for the affected client(s). If better spectrum efficiency (shorter burned airtime) is available in an alternate channel or band, the scheduler chooses that more optimal solution. And by leaning on MCS as one of the calculus components, the EIRP benefits of alternate spectrum choice will automatically receive an accounting (more power meaning better received C/N and hence higher MCS). A key leverage here is the 6 GHz spectrum, where even simple single user round-robin (non-OFDMA) scheduling can achieve 2 msec link latencies provided the served spectrum client count does not exceed a half-dozen or so participants.

The initial step in effectively handling latency expectations is to either infer requirements based upon some period of traffic observation or have the user explicitly designate a particular client/service binding which sets priority expectations for its data needs. The inferential determination can be done autonomously (and aligns with "fingerprint" methodology necessary when MAC randomization becomes a fixture in home wireless LANs – certainly likely during our study horizon). To mitigate startup jitter – and during the period when client traffic inference is still pending, the default priority handling should be set to highest (least latency and jitter). Over some initial transfers, the traffic from AP to client will establish its stochastic metrics (fingerprint) and "graduate" to the appropriate priority and latency setting. And while WMM provides a Wi-Fi mechanism which suggests leverage of 4 different priority queues, tremendous improvement for jitter comes with just the first parsing of traffic from "first-come, first served" to "warrants special handling".

Once the traffic is parsed for latency priority, the scheduler can then employ various buffer, payload packing and spectral leverages to manage the latent aspects of data connectivity for the impacted services. In the case of a low latency need, for example, a PHY optimization would consider the immense spectral width of the 6 GHz band. This large greenfield spectrum makes it straightforward to select a wide BW channel (to enhance burst speed) and locate it such that CCI is not measurably present (to guarantee a C/N which maximizes MCS and hence, bitrate); the Achilles' heel is the obvious requirement that both ends of the wireless link must support the 6 GHz band. The upside is that Standard Power and Wi-Fi 7 – both due within a year's time (or so) -- will both serve to catalyze adoption of devices which can lever that band.

2.1.1. Wi-Fi 6, Wi-Fi 7 and Standard Power

Briefly, Standard Power 6E allows the AP to paint the home expanse with up to 9 dB additional EIRP (over LPI at 160 MHz BW in the 6 GHz band) and essentially walks all 6 GHz clients up the MCS bitrate cascade (for downlink delivery), reducing consumed transmit airtime for a given packet size.

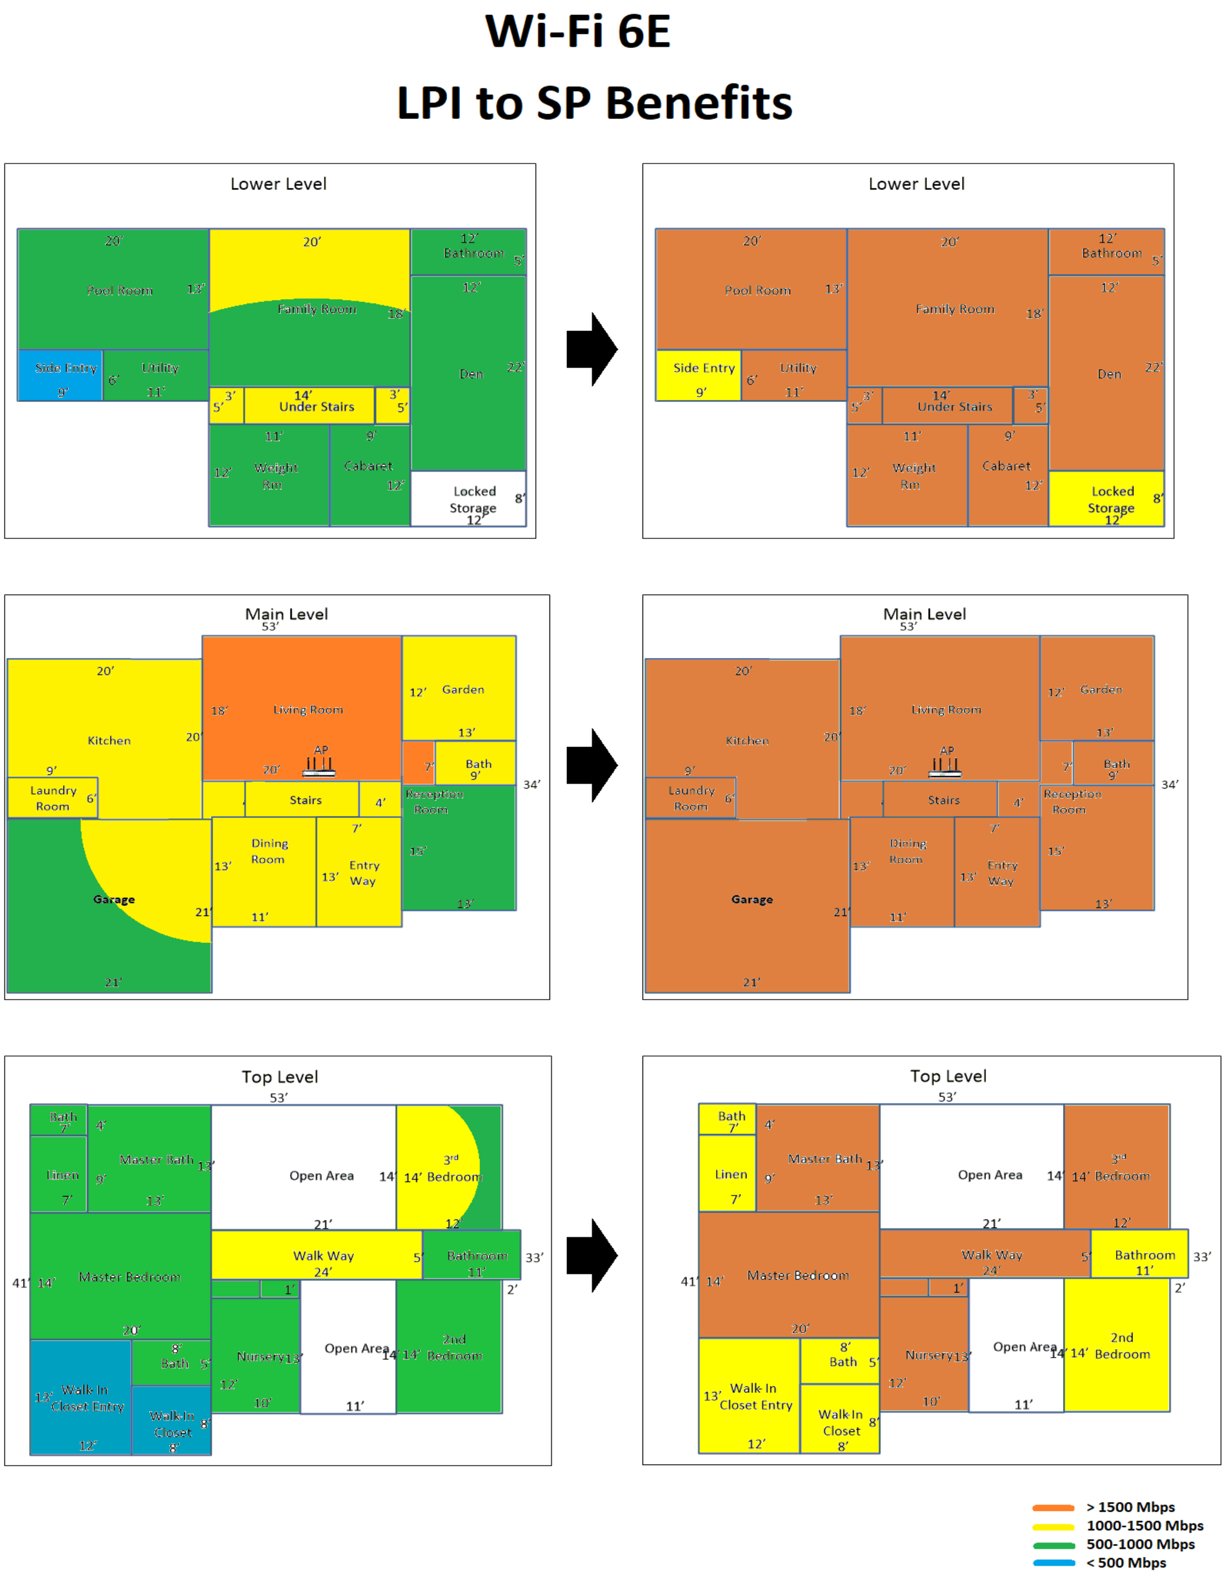

Figure 1 – Standard Power Benefit over LPI at 6 GHz in Arris Wi-Fi House

For its part, Wi-Fi 7 allows the AP to bind data delivery over multiple bands, with more granular BW leverage (think: spectrum puncturing, fractional RU use and multiband exploit), than any prior Wi-Fi MAC – and so can more cleverly leverage gaps in contested spectrum (in up to 3 bands, at that) to increase delivery rate and reliability over that of a single monolithic channel/single band at best-case C/N.

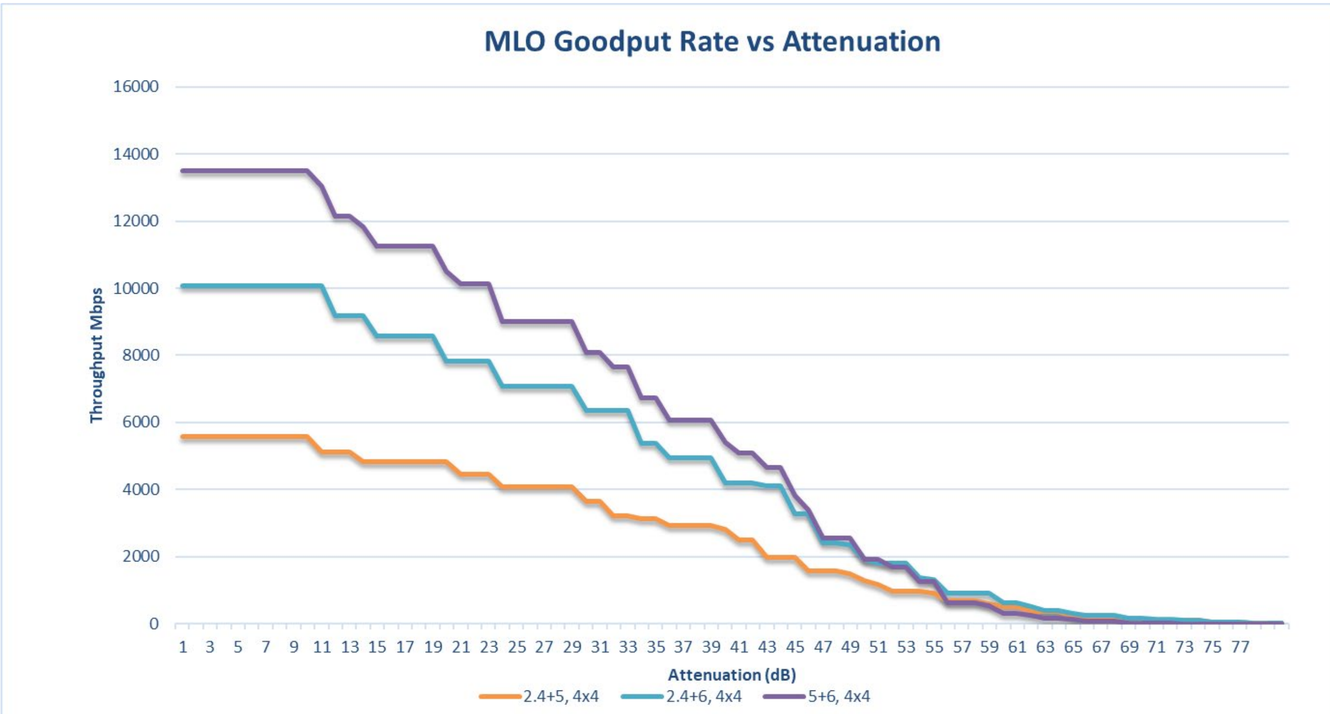

Figure 2 – Effect of MLO on Goodput at Various Bonding Options

The Wi-Fi 6E stranglehold on operations in the 6 GHz band will not even be fully two years old before Wi-Fi 7 devices will debut, beginning late this calendar year/early next. And it appears this landmark will also (roughly) coincide with the availability of Standard Power for AP devices in that same band. The significance of these coincident events manifests in outstanding bitrate reach throughout even large footprint estates – and couples this with multiband robustness and unprecedented levels of predictable, low latency signal distribution. As previously alluded, APs with Wi-Fi 7 will be able to implement schedulers which have at their disposal broader channel bandwidths than those offered by Wi-Fi 6 (320 vs 160 MHz), spectrum puncturing, MLO, OFDMA, bidirectional Mu-MIMO and TWT operations. Orchestrated data exchanges to the heterogeneous MAC limitations of respective client devices should be possible. This bonanza in Wi-Fi 7 additions to 6E – and their so-proximate release to the first 6E-capable devices – has largely rendered 6E a moot (OBE) waypoint proposition (succeeding only in fragmenting the 6 GHz Wi-Fi market in its early days, unfortunately). There exist two exceptions to this proclamation on 6E early obsolescence: “virtual wireline”, low-latency dedicated 6 GHz Wi-Fi channels between AP and gaming endpoint and AP:AP (non-backed-off power and maximum BW on a dedicated channel) trunk links for wireless extension within the premises.

Regarding the former, addition of Ethernet (or USB)-to-6E dongles to AP and client on opposing ends of a dedicated-channel 6E link provide an instantaneous means of running a (virtual) Ethernet cable between those devices and establishing near-wireline delivery latency on said link. Regarding the latter, for large-footprint homes, Wi-Fi extension via a dedicated 4SS 6E link between gateway device and mid-home can provide Wi-Fi extension with trunk capacity easily exceeding 2 Gbps – again with low single digit millisecond latency over the link. And by doubling down on this arrangement via leverage of AFC-enabled standard power, 4W EIRP at both ends of the link will be allowed; this trumps the power footprint for APs in any other Wi-Fi band. Boutique solutions, certainly – but also worthy of note in the time horizon under study.

2.2. The DOCSIS (WAN) Environment

Though (for some) glacier-paced, the onset of a DOCSIS 4.0 era has nonetheless proceeded and now finds itself well into commercial chip implementation phase within the development domains of the two most pedigreed commercial participants in the DOCSIS legacy. However, the end game – as regards deployment particulars and timing – is very much undetermined. There exist multiple permutations of DOCSIS wireline epoch which could conceivably meet network needs in the coming few years with rather steep implementation costs for some weighing on the solution calculus. Relief for the restrictive legacy return path is the only identifiable common element at this juncture, with the trunk diplex split, continued exploit of SC-QAM for video and final forward bandwidth still under evaluation – as well as the MAC subtype(s) to be exploited across the full band (such impacting the effective spectral density and hence capacity of the total spectrum).

Regarding the latter, a “budgetary” approach to upgrading the DOCSIS network might be to move to DOCSIS 3.1, slide the diplex frequency up to 204 MHz and set the forward band limit to at least 1.2 GHz. Such a conservative approach might attract some attention from (and provide temporary relief for) capex-bound systems, but for the purposes of analysis we will focus on the DOCSIS 4.0 variants matrixed with more future-proofed options on diplexer frequency and upper band edge (in anticipation of a higher adoption rate which better describes the coming WAN environment for the greatest population of NA subscribers). A brief summary of DOCSIS 4.0 options follows, in order of increasing implementation complexity.

2.2.1. DOCSIS 4.0 Enhanced Spectrum DOCSIS (ESD)

The following pictogram illustrates diplex split operations for DOCSIS 4.0 ESD in either 5x5 or 5x7 formats (referring to number of OFDM 96 MHz blocks of upstream traffic):

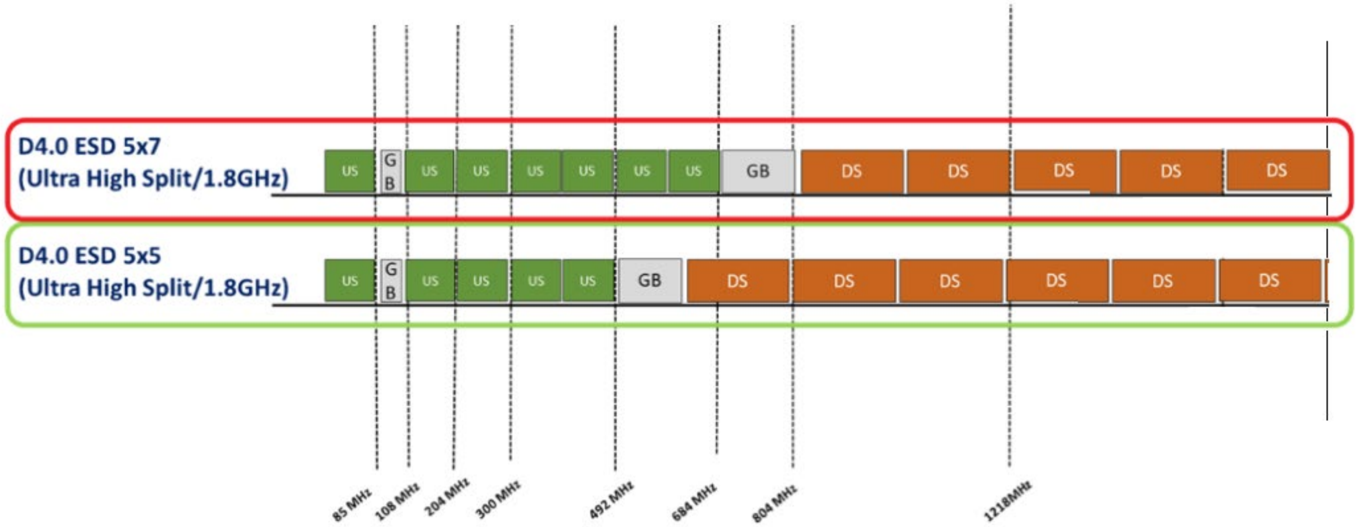

Figure 3 – DOCSIS 4.0 ESD Variants

These are but two examples; the spec lists five different diplexer options in up/downstream operation. Note the corresponding shift in the diplex corner frequencies in these two cases and the option to shift the full band forward cutoff up to 1.8 GHz if desired. If we conjure full exploit of the 1.8 GHz bandwidth with only OFDM carriage (no video SC-QAM) and maximum upstream bandwidth (so closest to most symmetric wireline operation), the EOL performance of a cascade would be capable of a downstream PHY rate exceeding 10 Gbps for all but the highest diplexer splits.

2.2.2. DOCSIS 4.0 FDX

ESD pushes the forward spectrum out and eliminates less efficient SC-QAM, but FDX goes a step further (and more complex) – shiftable duplex signaling bands (adjustable diplex frequencies). The CableLabs visualization of this looks as follows:

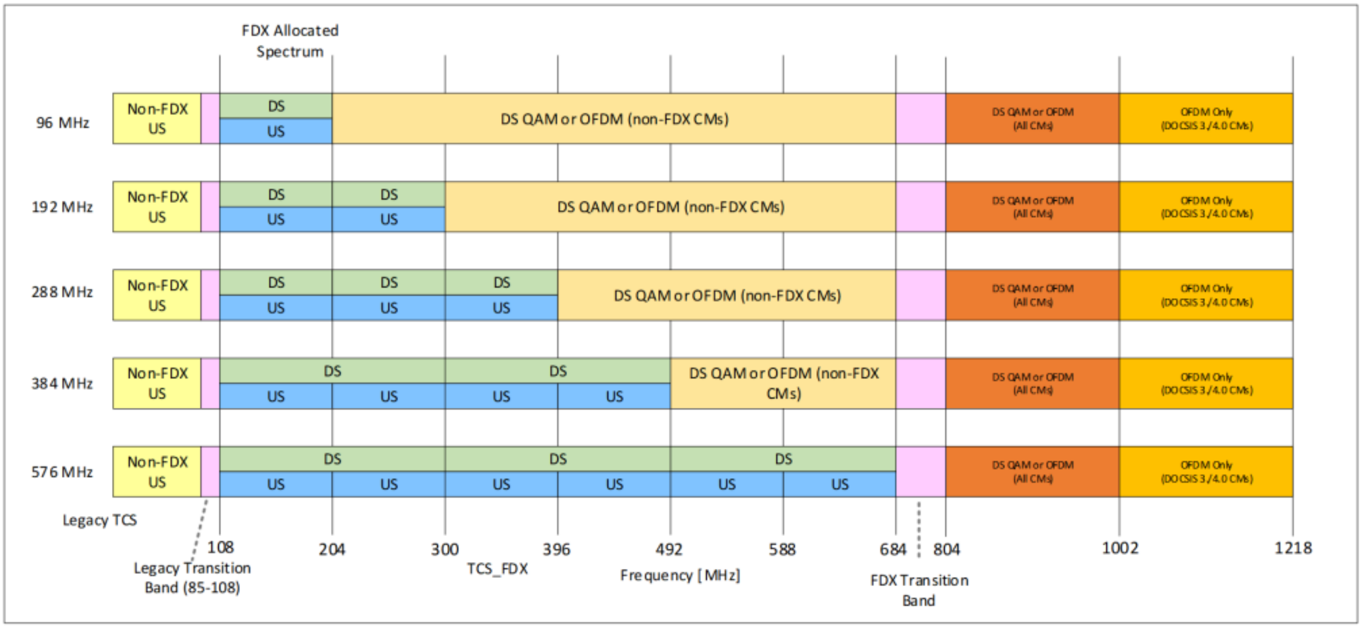

Figure 4 – DOCSIS 4.0 FDX Variants

The clear benefits on the flexibility front fall into the category of adaptable duplex ratio. Post-pandemic in particular, it is not entirely clear whether the traditional 10:1 split (in the US) for downstream:upstream capacity will continue to make sense. (The SOHO, IoT and data-sourcing aspects may see relatively higher growth rates than traditional lean-back video consumption, for example.)

3. Managing E2E Packet Latency Across WAN and LAN

Proper stewardship of packet delivery across the concatenated DOCSIS and Wi-Fi networks involves a synonymous mapping scheme for the various packet priorities within the two domains; anything less promotes thrash in the queuing arrangements on one side or the other, resulting in jitter increases as the packet queues are spasmodically serviced. Fundamentally, synonymous mapping involves marrying the wireline LLD (low latency DOCSIS) – where appropriate -- with Wi-Fi's WMM and assigning sensible, staggered priorities for the client services drawing data connectivity support from the two networks.

3.1. Low Latency DOCSIS

Pre-DOCSIS 4.0 already features an unevenly implemented technique for minimizing forwarding delays for certain priorities of packets called Low Latency DOCSIS (LLD). This is an evolutionary endpoint along a vector of buffer management trials which produced the following signatures:

(CM/CMTS RTT)

	When Idle	Under Load	99th Percentile
DOCSIS 3.0 Early Equipment	~10ms	~1000ms	~1000ms
DOCSIS 3.0 w/ Buffer Control	~10ms	~100ms	~100ms
DOCSIS 3.1 Active Queue Management	~10ms	~10ms	~100ms
Low Latency DOCSIS 3.1	~1m	~1ms	~1ms

Figure 5 – Evolution of DOCSIS pre-4.0 Roundtrip Time Latency

Fundamentally, LLD calls for a parsing of duplex service flows into LL and Classic – with a management function to apply service classifier membership to discharge data flows identified as not requiring rapid forwarding from the LL flow into the Classic mix. It also monitors queue health along both flows to make sensible, weighted exploit of the low latency path under dynamic loading conditions.

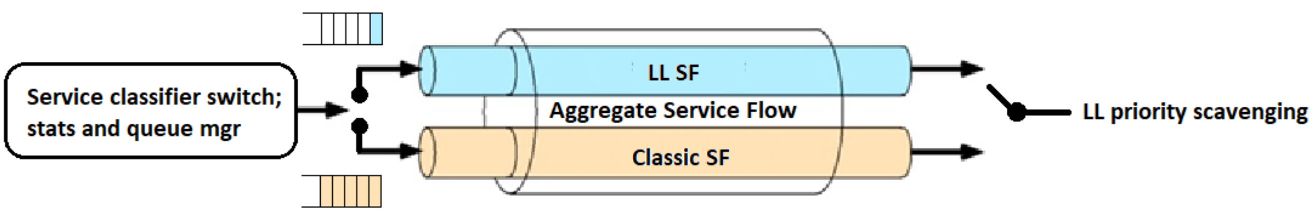

Figure 6 – 2-Component Flows in DOCSIS Pipe, Low Latency and Classic

With this simple expedient of a second, shunt buffer queue for high priority/low latency traffic, the following remarkable results are obtained in the DOCSIS domain:

Figure 7 – Differential Flow Performance Upstream for Latency and Jitter

As can easily be identified, both latency and its deviation (jitter) are remarkably constrained once LL discipline is applied in a second service flow (typically 10 msec RTT latency with 1 msec jitter, across a service group loading from 25% to 90%).

So what services require low latency assistance? The following attempts a parsing of typical home data services into “LL” and “non-LL”:

Flow Mapping for Typical Data Services

Low Latency	Bulk/Background
Web	Software Updates
Videoconferencing	Dropbox
Audio Streaming	Email
VOIP	VPN*
Gaming	Video

**** may also require LLD treatment***

Figure 8 – Candidate Service Mounts for Low Latency Grooming

Essentially, then, rate tiers provide a means of selectively monetizing progressively more aggressive user bitrate demands and this type of selectivity can now be applied to the latency domain and enhance those services requiring preferential, low latency handling.

4. Extant and Emerging Premise Services Landscape

Any well-founded attempt at aligning network data delivery to, and scavenging from, home premises necessarily must consider both present (established) service mounts and those calculated to emerge and mature – or wither and go dormant -- over the five-year consideration horizon posed in this paper to supply foundation for modification of network behaviors. A general expectation seems to be that data flows will increase (and on a duplex basis) between premise and cloud; but as always there needs to be sober justification for the timing and magnitude of capex investment to enhance data flows versus the life cycles of legacy and emerging services and the migration of opex balance sheet operating points (typically subscription revenues against licensing or other service origination/maintenance costs). This section will attempt to assign relative values to services which are expected to populate the network application service palette in the postulated time window.

4.1. Consumer Impressions of Wi-Fi Performance and Feature Value

A proper sampling of present network performance service levels and features – as perceived by the end users themselves – seems an appropriate and necessary first consideration in establishing any motivation for alteration of the current state of network data operations. This can also provide the rational basis for critique of the relative end perceived value of various modification strategies. And providing a voting mechanism on alternate value strategies would help establish success vectors (from the users' standpoints) regarding where to steer the performance of the combined cloud-to-ground (and back) environment experienced by these end users. Such a survey was constructed and the details of the canvassed topics – and the recovered results from 546 global participants – follow.

The regional breakout for the data – which sources from the CommScope associates database – exhibited the following detail:

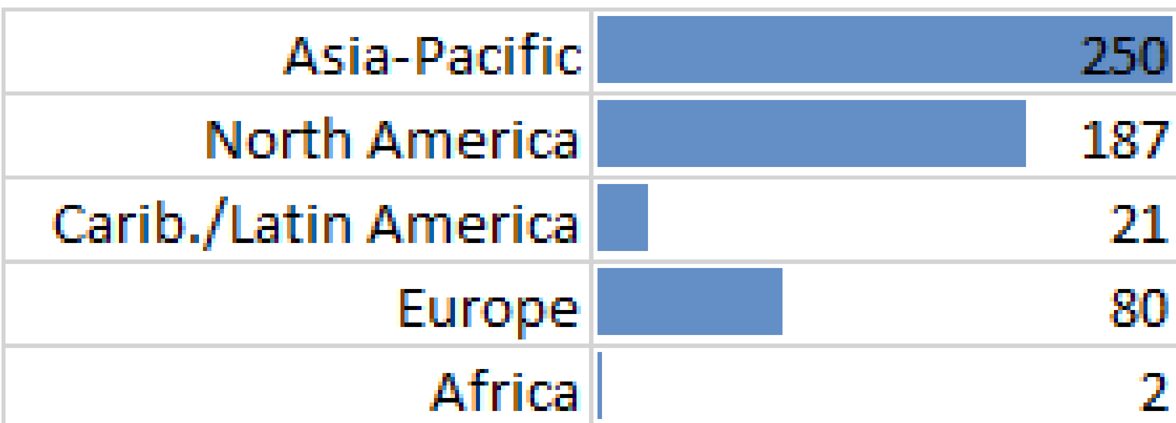

Figure 9 – Regional Basis for User Network Impressions

The survey attempted to identify users' appreciation of the common metrics used to described network data attachment and assess those parameters and values which resonated from a value standpoint. To orient some of the subscription realities, users were asked their broadband speed tier:

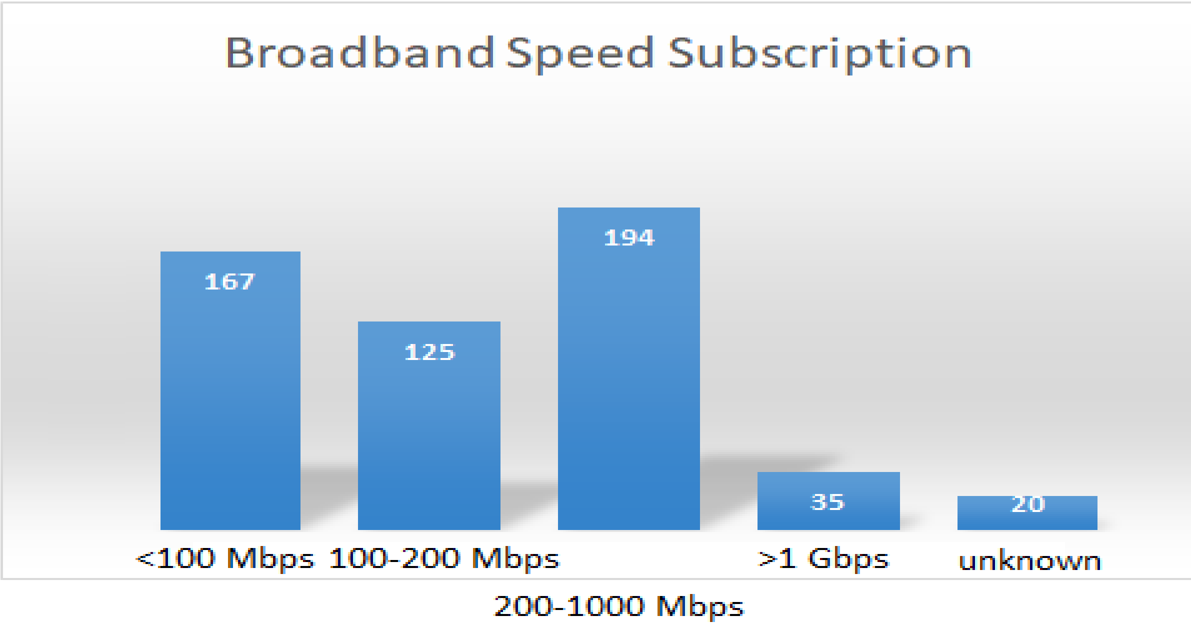

Figure 10 – Users' Broadband Speed Tiers

Note the gross preponderance of sub-Gbps subscription, with the plurality lying in the 200-1000 Mbps bucket. Cost for this privilege was solicited:

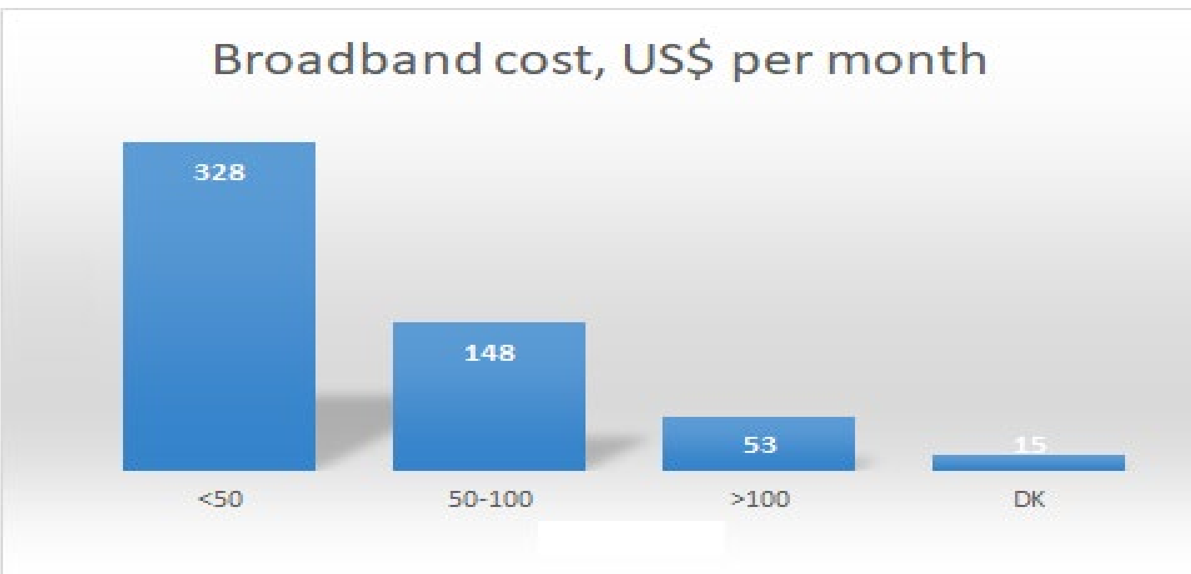

Figure 11 – Budgeted Monthly Cost for Network BB Data

Note that globally, caps appear at 1 Gbps speed and $100/month cost. In terms of satisfaction with this service exchange rate, nearly 79% of respondents indicated they were happy with the value proposition – perhaps underpinned by the fact that just over ¾ of the users had an alternate supplier handy (so, presumably, might have done a "best offer" type of alternative analysis). To add to the inertia, nearly 87% of those answering defined themselves as not actively looking for replacement broadband services.

However, in terms of total network support, just over half (56%) of respondents felt that their home Wi-Fi matched the performance of their WAN attachment – and 62% felt that, were investment in their total networking were to be made, that investment would fall to the Wi-Fi side. (There appears to be some legacy ill will or suspicion associated with Wi-Fi – perhaps created in bygone single-band days, or with a preponderance of 2.4 GHz based, single antenna clients. Certainly, Wi-Fi beyond dual-band MAC 5 – 802.11ac – has no problem with disposition of packets at the less than 10 Mbps throttling rate which represents average WAN attachment pace for dwellings!) And it is true that, wherever network issues arise, the human proximity of Wi-Fi CPE lends itself to that nearby equipment shouldering blame for connectivity issues (whether or not they originate there).

In the next section of the survey, users were tested as to specific opinions on upstream and downstream data – the latter both regarding speed and latency/jitter. First, as regards upstream connection speed:

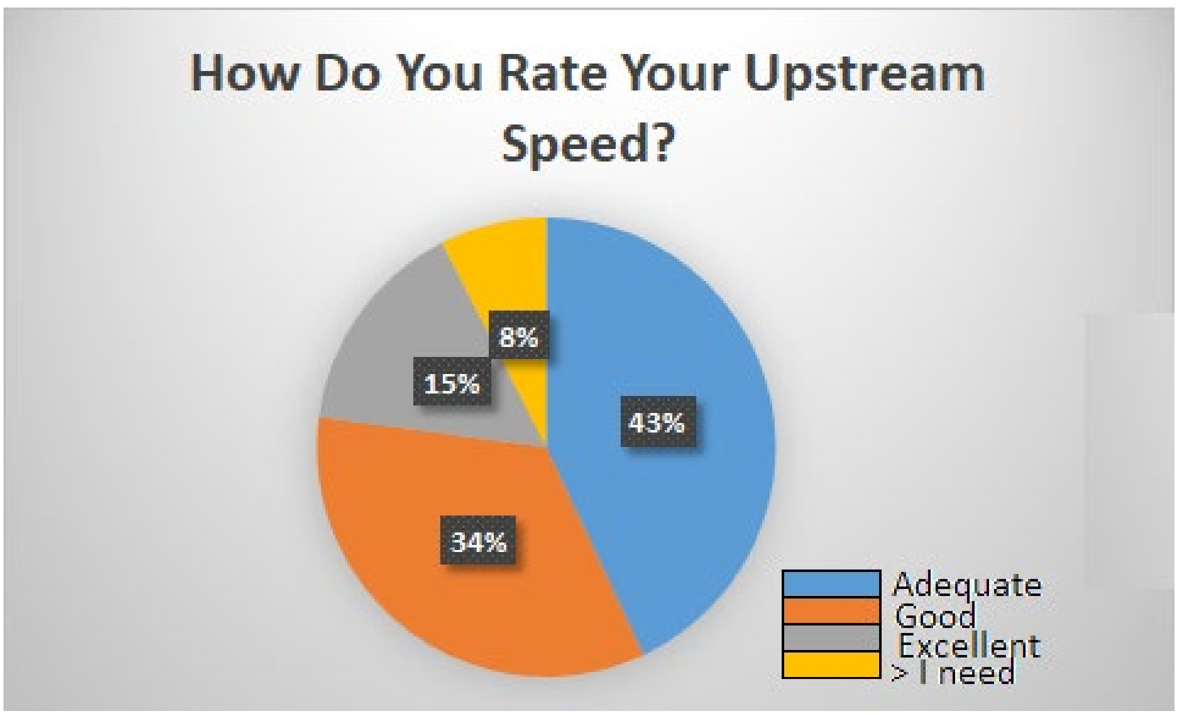

Figure 12 – Upstream Bitrate Rating

Note that nearly half the results peg the upstream as being either "good" or "excellent"; whilst an uplifting opinion, one wonders if users understand how little of their connectivity traffic is actually carried upstream – or how to quantify its "goodness". Still, their resolution on this point was tested by the alternative to either stay with their current provider at a 10% price cut or move to another provider for twice the speed at the same cost they currently experience. The result was telling:

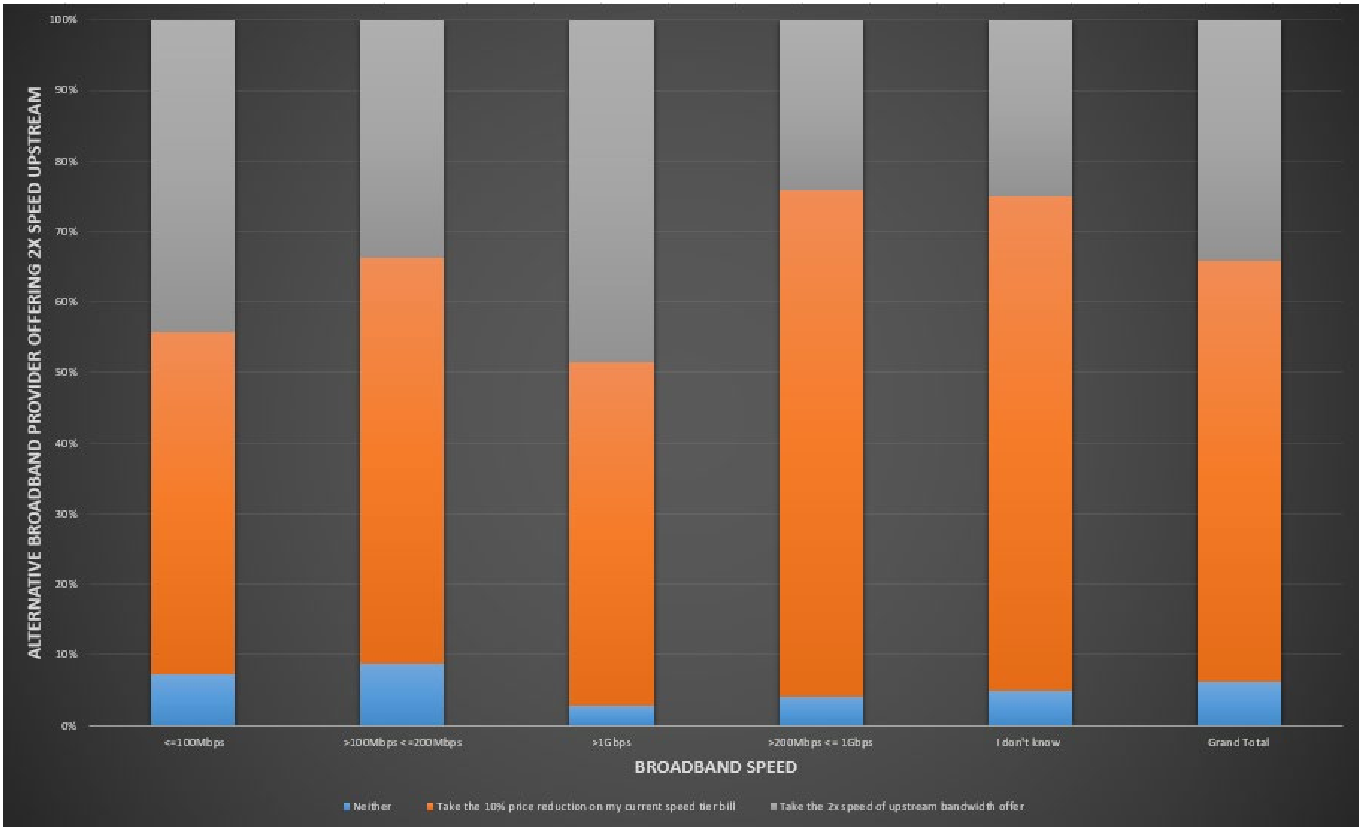

Figure 13 – The Upstream Value Challenge

Only about a third of respondents could see the worth in switching up suppliers for twice the speed at identical cost to them; still, this shows some fair sensitivity to upstream performance value – so it is noteworthy the majority push would be for less cost. (This opposite view holds in high connectivity speed regimes, so clearly even when satisfied with downlink speeds, there is opportunity to mine interest in better upstream capability). Next, the same question was posed regarding downstream performance. The response shift is interesting:

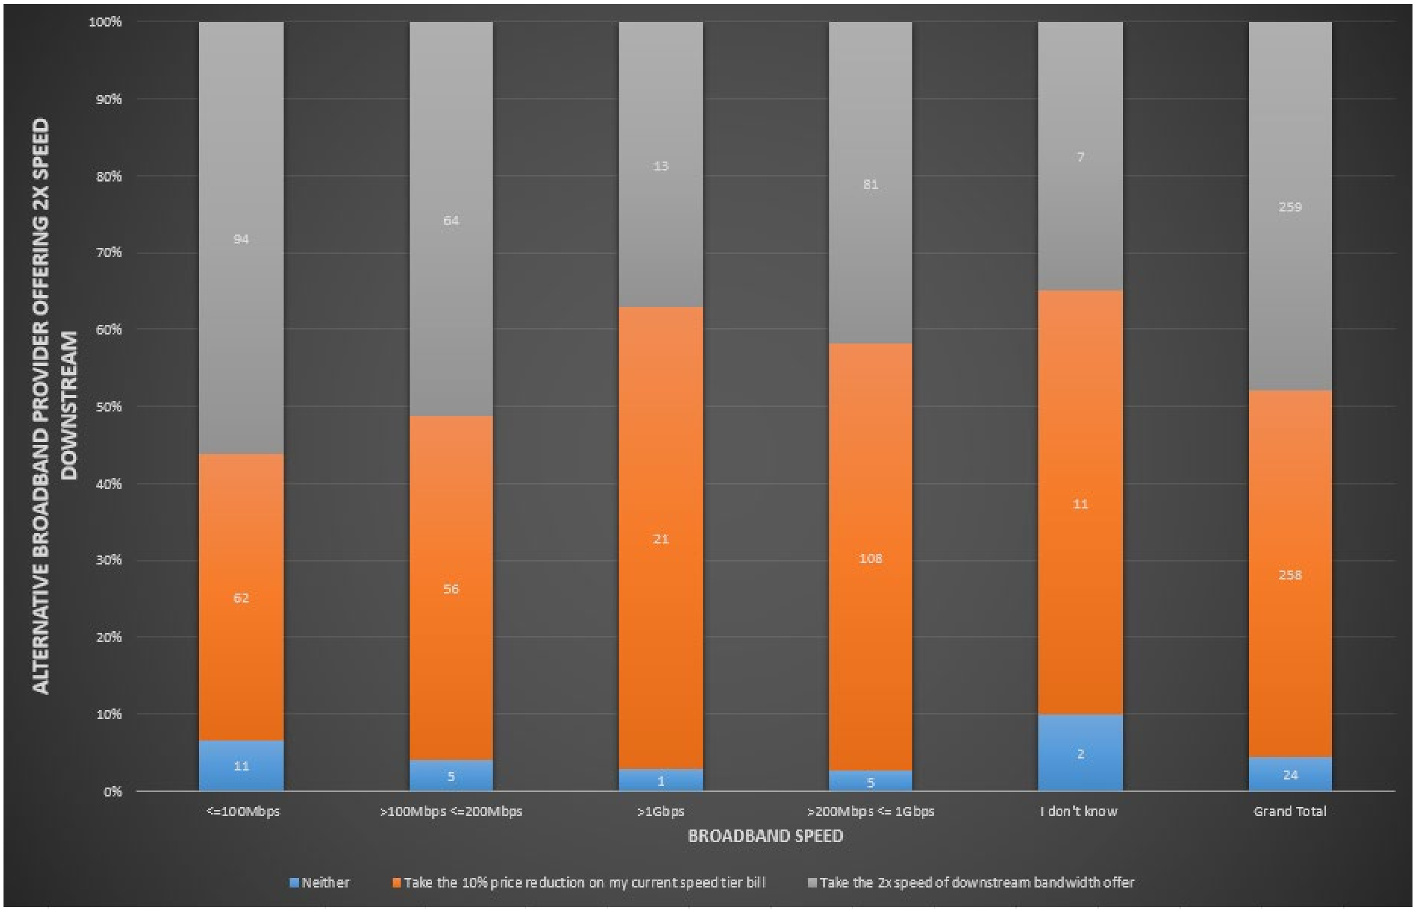

Figure 14 – The Downstream Value Challenge

Note that roughly 10% of respondents opted up for more downstream speed than were prepared to do so for upstream improvements. This suggests that they are aware of some correlation between DS bitrate and in-home services performance (while perhaps not being convinced that it is worth that much more than getting a quick 10% service cost rebate for the extant DS rate). Note that the higher the subscription speed, the more likely the user was to push for cost mitigation than more bandwidth (obviously viewing their connection as at least adequate).

The survey then shifted to probing users' feel for the value of latency and jitter in their network performance; first, as to latency:

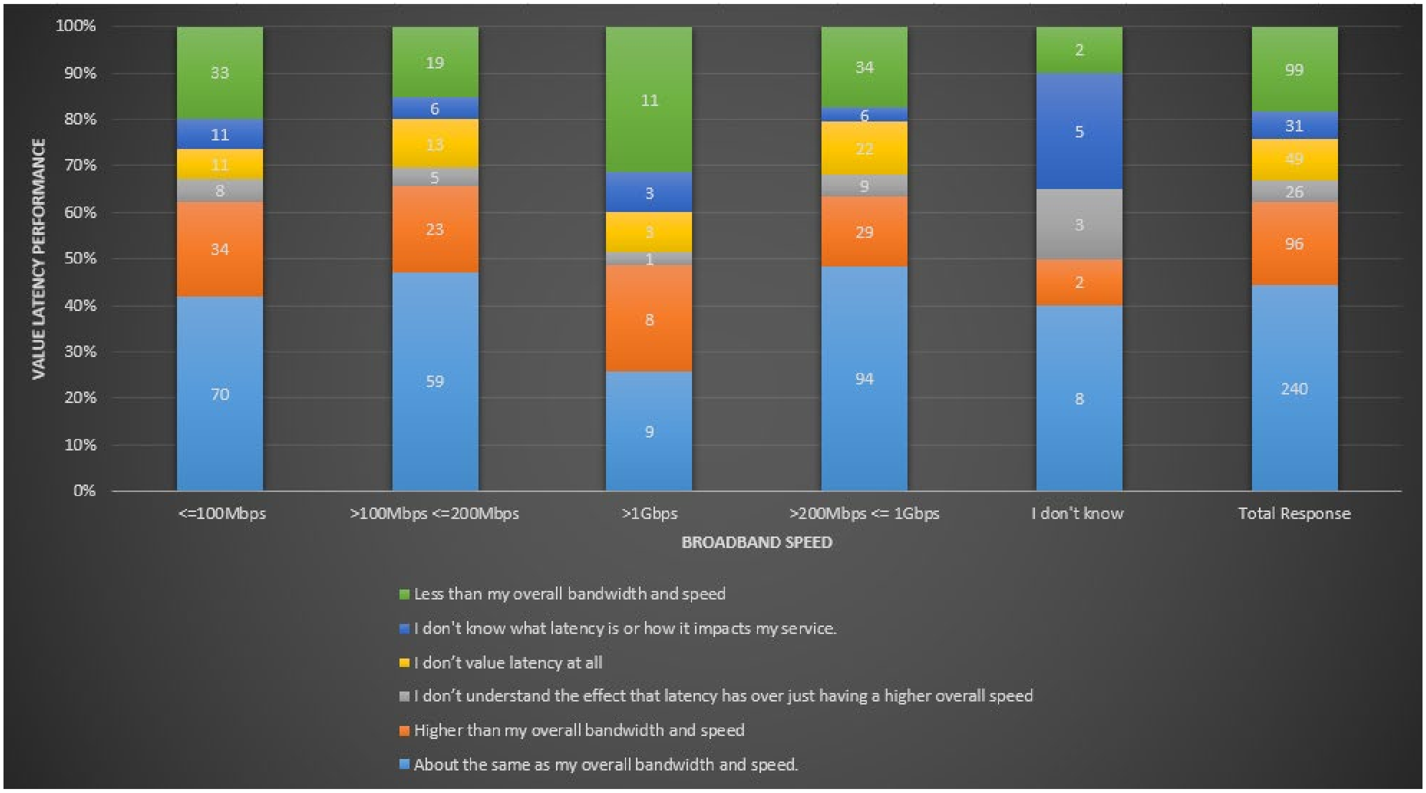

Figure 15 – Where Latency Fits for Users VS DS Bitrate

It is difficult to infer anything here: the plurality of users would have latency rank equally with bitrate as regards connection "goodness" but the weighted average of value from this chart produces no other insight – other than high bitrate users REALLY appreciate their high bitrate tier more than latency. Perhaps unsurprisingly, the jitter evaluation by users produces a near-mirror copy of the latency result:

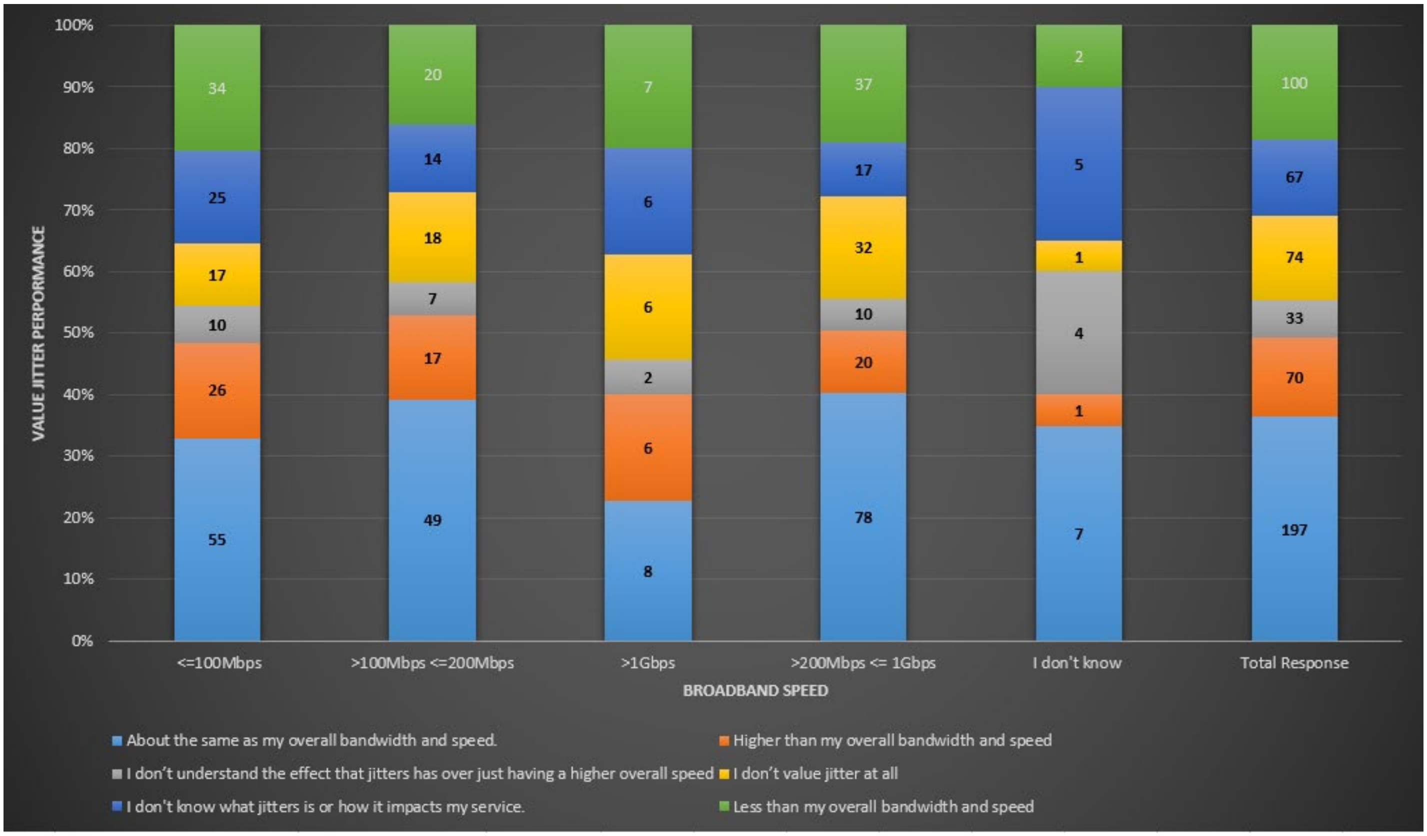

Figure 16 – Where Jitter Fits for Users VS DS Bitrate

If these results were perplexing, perhaps they appear less so when a litmus test for the magic effect of high bitrate labeling is concocted. Users were asked if, rather than receive a guarantee of bitrate, they would instead prefer a promise that all their equipment would be set to operate under the best bitrate, latency and jitter conditions possible, an astounding 2/3 replied "No". When this offer was sweetened with the inducement of a 10% rate reduction, the "yes" responses did increase – but only about 10 percentage points (to ~ 44%). So there is obviously a great deal of marketing value to network connectivity distilled to a single rate tier value. Given this reality, the respondents were quizzed as to what speed grades they looked for: the lowest to do the job, the highest offered or something in between. The responses came down as follow:

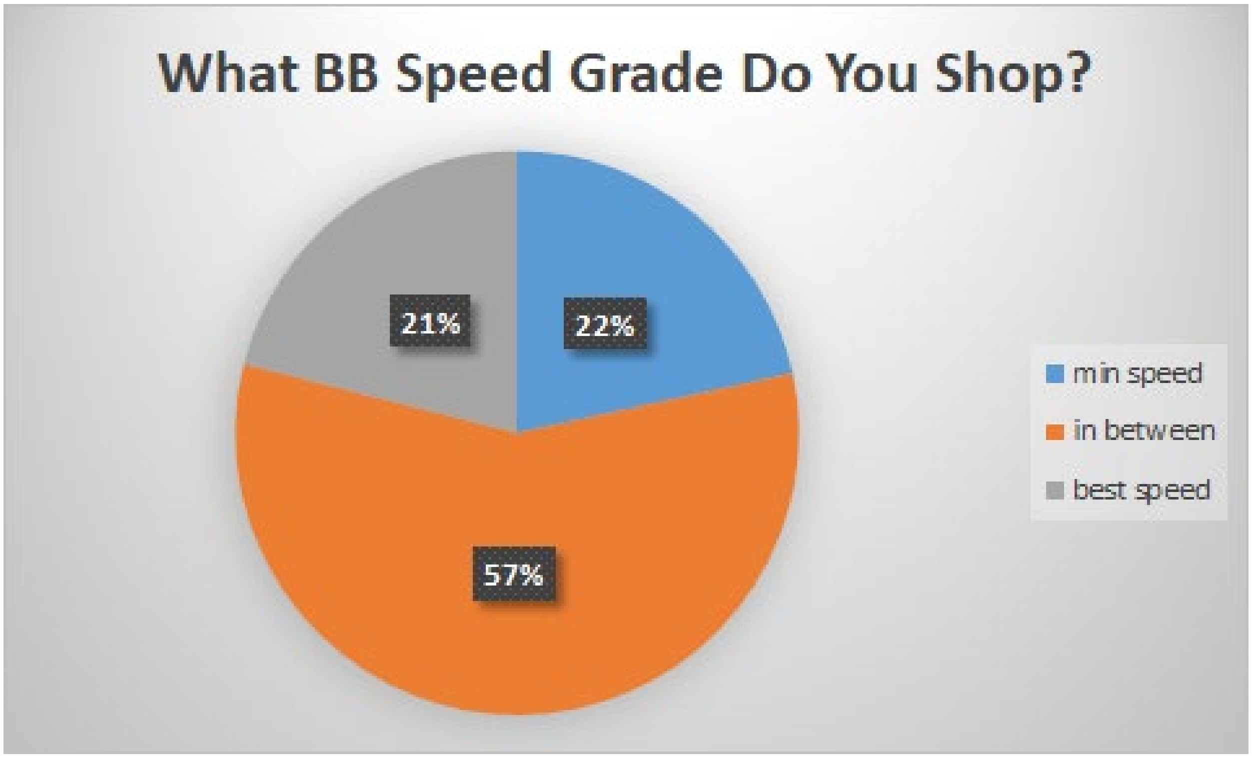

Figure 17 – BB Connection Rate Impetus to Shopping

To add a layer to this query, the users were asked if new broadband services (specifically like VR) factored into their purchase thinking for network connectivity. The responses came thus:

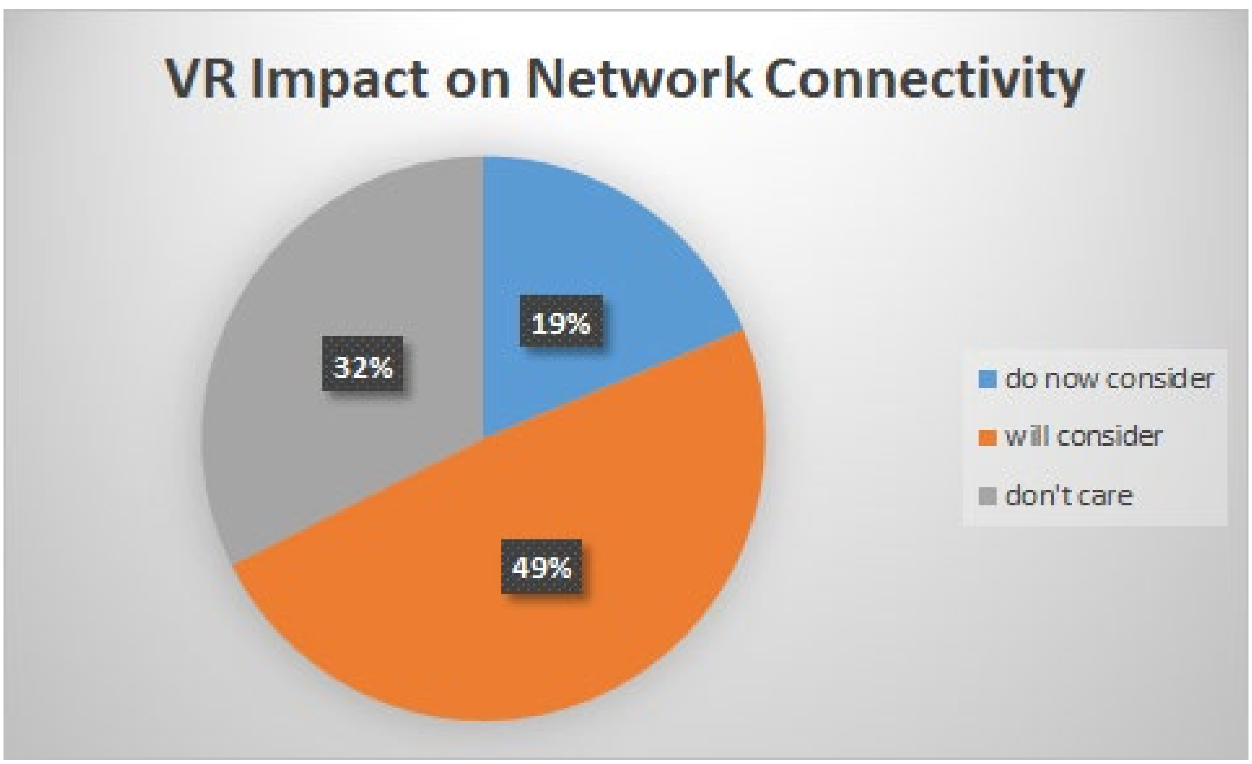

Figure 18 – VR impact to BB Connectivity Requirements

Nearly 70% of users indicated a high level of sensitivity to high bitrate demanding services emerging in their homes and appeared prepared to include consideration of these in present and future BB purchases. (This is a critical leverage point, as VR developments will very shortly push higher bitrates into the home, as we detail in following sections).

The users were then asked to rank the top attributes they examined when looking for broadband services; the rank order of the number 1 responses was distributed as follows:

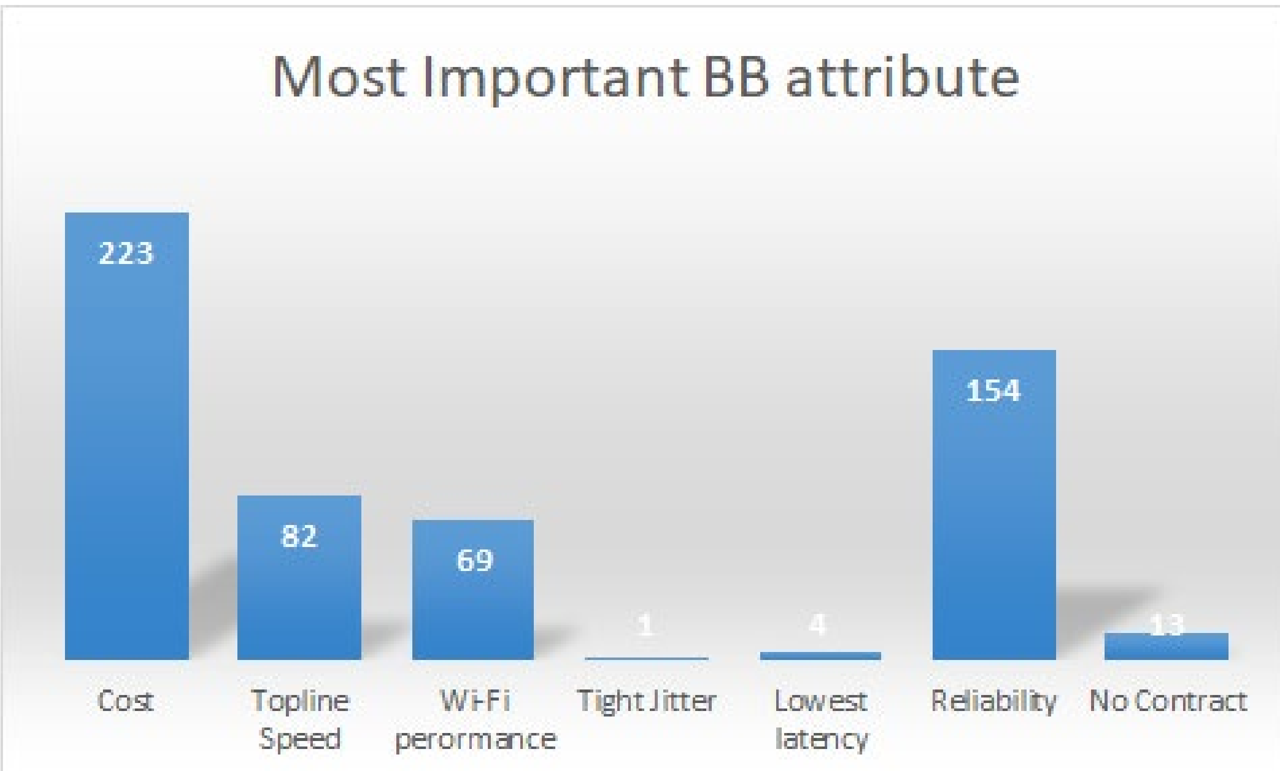

Figure 19 – Most Important Broadband Network Attribute

Cost, reliability and topline bitrate were noted – but oddly enough, Wi-Fi performance was seen to count for inclusion in broadband network performance as well. In terms of home uses for WAN services, the following were cited:

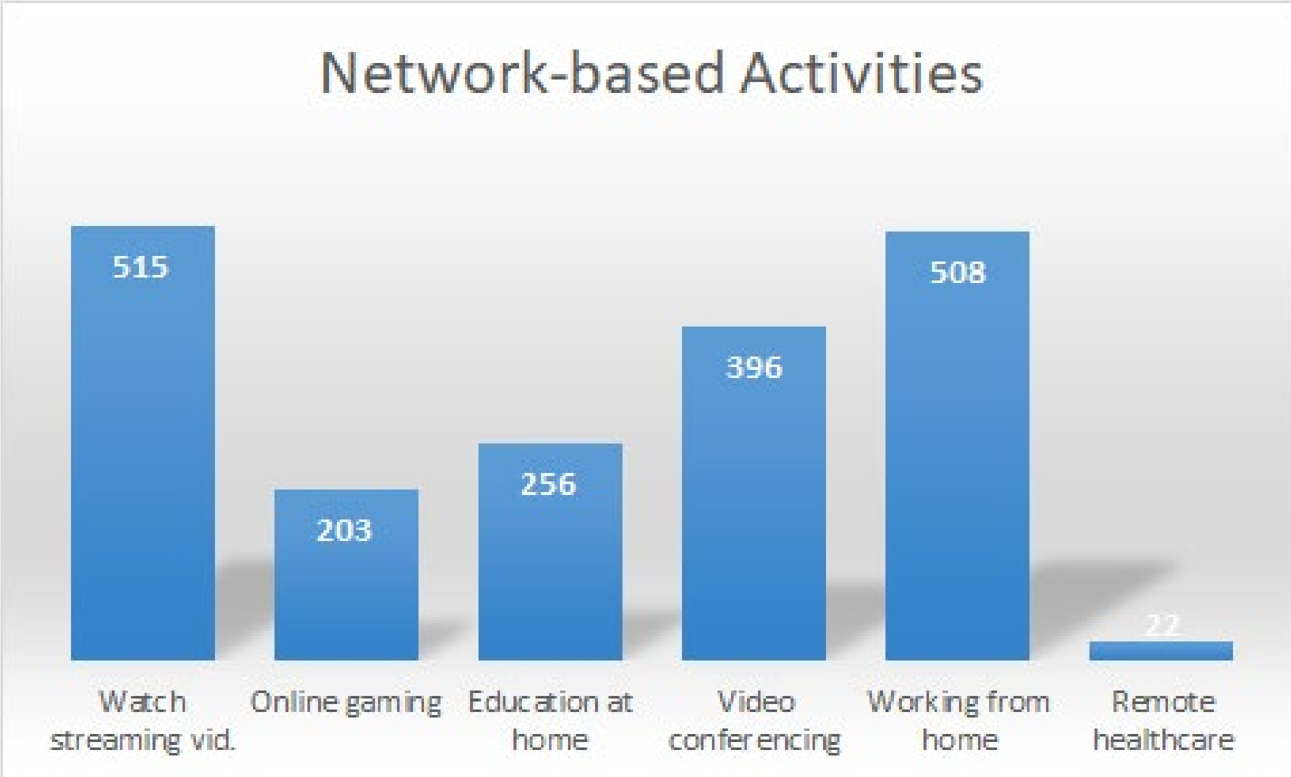

Figure 20 – Broadband Network Enabled Home Services

Unsurprisingly, watching streaming video set the bar, but in the post-pandemic world, working from home and teleconferencing received high notice. Not yet too noteworthy, but of interesting acknowledgment are the votes cast for remote healthcare; expect these to rise significantly as AIP takes hold.

A final query had to do with concurrent users in the home. Given the US household average of 2.6, perhaps these numbers are not too far-fetched:

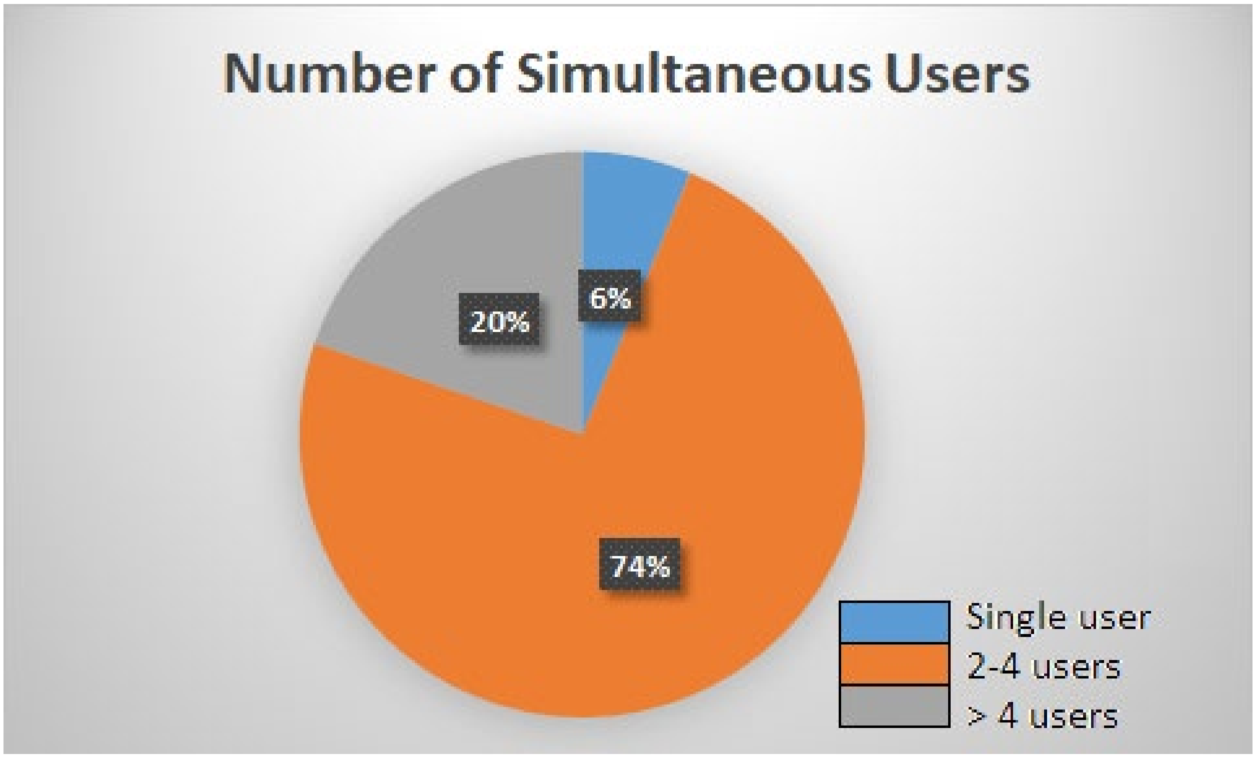

Figure 21 – Simultaneous User Leverage of WAN

The overall impression of users seems to be that they are very satisfied with their broadband investment; they value cost, speed and reliability highly (but are largely clueless on the value of latency and jitter); spend the largest chunk of time ingesting streaming video (still) and are curious about the impact higher-value video services may have on their networking environment (in which they link Wi-Fi performance directly to that of the WAN).

4.2. Premises Services Data Payloads

For the next five years (and in fairness, the conceivable future) video services will continue to represent the largest tranche of downstream service data consumption. This is not one's grandparents' video experience; broadcast channels have morphed to applications, with the appended '+' symbol an indicator of viewing options which exist underneath a brand headliner. But however the content is catalogued and presented, it is clear that lean-back video consumption will remain a lynchpin service for almost every subscriber given that current consumption metrics peg video as burning close to 80% of premise-bound downlink data payloads. And while the present HD resolution norm for premium video meters at ~ 5 Mbps given present codec efficiency and common rendering devices, growth of in-home 8K-capable TVs promises to boost the video data appetite to something in the realm of 50-100 Mbps per screen (at coming codec efficiencies and higher frame rates) in the very near future. (It is, in fact, the rate of consumer uptake of higher quality video services which most challenges the forecast of any premises data consumption budget.) In terms of perceptible wireless impact, video picture quality is also the most relatable quality metric of wireless performance accessible by the average non-technical user and provides an obvious marketing leverage point in side-by-side comparisons of network bitrate value. (Imagine the marketing penalty associated with the realization that network A can mount the latest VR headset and network B cannot.) Pegging an inflection point for adoption of super-resolution and frame rate rendering devices seems fraught with false starts, but perhaps an understanding of what is implied (in terms of network performance) might serve to frame the coming expectations:

Bandwidth Requirements of Cloud VR

- **The estimated per-user bandwidths required by strong-interaction services in the three phases of Cloud VR development are as follows:**

[1]

Phase		Fair-experience Phase	Comfortable-experience Phase	Ideal-experience Phase
Typical content resolution		2K (equivalent full-view resolution: 4K)	4K (equivalent full-view resolution: 8K)	8K/16K (equivalent full-view resolution: 12K/24K)
Typical terminal resolution		2K	4K	8K/16K
FOV		90° to 110°	120°	120° to 140°
Color depth (bits)		8	8	10~12
Coding standard		H.264/265	H.265	H.265/266
Compression ratio (I-frame/P-frame)		25/75	38/165	50/255(8K), 83/585 (16K)
Strong-interaction VR service	Typical bitrate	40 Mbit/s	90 Mbit/s	Full-view: 290 Mbit/s (12K) 1090 Mbit/s(24K) FOV: 155 Mbit/s (12K) 580 Mbit/s(24K)
	Typical bandwidth requirement	80 Mbit/s	260 Mbit/s	360 Mbit/s (8K) 1.5 Gbit/s (16K)

Note relentless increase in video bitrate

Figure 22 – Stacked Bitrate Increases Driving VR Simulation Immersion and Super-Resolution Screens

A distant, if otherwise distinct, common second service type which may be reasonably expected to continue (if not grow) its ranking among services is network gaming. (This includes immersive AR/VR equipment such as Oculus – the 80% market share leader – which is projected to exhibit a CAGR between 40 and 60% over the time period 2018 to 2025). This compound growth suggests an exit value from our 5-year window at north of $20 billion dollars of sales revenue for AR/VR. The desired network attachment profile for these types of gaming services includes a latency component which is two orders of magnitude tighter in tolerance than that required for ABR delivery of video packets for lean-back viewing. The desired jitter descriptor is a further order of magnitude more restrictive than that (so low single-digit milliseconds). Present Oculus data burn mimics that of a moderate resolution planar video screen but upcoming versions (described above) seem to intend on increasing resolution and pushing frame rates, *a la* 16K TVs. (And recall that VR/AR equipment paints not just frame content straight ahead, but also prospective adjacent frames to the sides, above and below, to reduce rendering stutter in rapidly moving head exercises). As such, it should be expected that near-term-available Oculus II and III headset gear will exhibit an even higher data appetite (from 50 Mbps to perhaps as much as 400-500 Mbps) than those 8Kp120 UHD screens which lie largely unexercised in present home viewing rooms.

AR/VR systems aside, dedicated gaming stations and PCs with legacy planar rendering surfaces will continue to see heavy use in the residences of end users during the projection period. As is the case for immersive (head-worn) rendering systems, the twitch aspects want for predictably low latency with corresponding small amounts of packet delivery jitter. In terms of raw bitrate support, however, network gaming peaks at 150 MB/hour (~ 300 kbps). As regards pipe demand, this is on the same order of magnitude as listening to streaming audio services or conducting an IP-based phone call.

An emerging video option, yet ungraduated from boutique levels of interest and deployment, is holographic imaging. Requiring multiple video planes to construct, it represents the largest potential downlink consumption in the residence (if such is so equipped). Although not yet representing notable consumer interest, it seems an inevitable (if expensive) extension along the vector of greater immersive simulation or learning. Current state-of-the-art requires a data downpipe of nearly 2 Gbps capability – several multiples of present-day “full up” home services network attachment – so when adoption becomes real, such will manifest a huge uptick in network connectivity. However, given the abbreviated window of expectations for this study, the marketing outlook for holographic support in the home may be presumed to be beyond the time scope of this analysis.

Five other, lower bitrate and less stringent latency-invoking, service profiles can easily be identified: website browsing, digital assistant interplay, live streaming, SOHO teleconferencing and IoT telemetry forwarding. Of these five, the latter three represent the largest upstream bitrate generators: camera feeds. Depending on video resolution and codec, these can vary from 500 kbps to 2 Mbps per feed and (constant rate teleconferencing aside) the common use case involves intermittent (triggered) operation, involving transfer of perhaps 7.5 MB of video/audio clip data over a 30-second capture window. From a security standpoint, it would not be unusual to expect camera coverage of all premises ingress/egress points, resulting in perhaps 4 streams worth of randomly triggered data per dwelling. In general, however, the total contribution to network connectivity (and upload at that) might only see 5 or 6 episodes dispatched per 24-hour period (so 40 MB total security payload). The teleconference load depends on meeting times (naturally), but six hours of daily meetings would contribute 5.4 GB/day worth of upstream traffic. Live streaming (Tik Tok and the like) can vary from short bursts to sustained use on the order of a teleconference hit per mounted instance.

Digital assistant interactions process audio clips of typically 48 ksps sampling rate. Assuming a multichannel far-field array capture, one could budget 400 kbps uplink and perhaps 200 kbps downlink traffic, representing around 6 Mb uplink and 1.5 Mb of downlink utilization over a 2-minute, randomly occurring interaction cycle. As with security cameras, assistants will be sprinkled throughout the home (and recall these can be apps on multiple types of CPE, including phones, remote controls, audio players and TVs). A dozen of these “conversations” among all household members might be a reasonable estimate for maximal usage during a day, which yields much less than 10 MB uplink / 2 MB downlink total payload over the course of a day.

Smart home IoT command exchanges and telemetry uploads over Wi-Fi tend to be even less of a network burden than smart assistant interchanges. If we exclude the voice-ordered aspects and consider their impact captured above, telemetry and commands might amount to no more than 30kB (up and down) worth of daily interactions; though what is critical is timely forwarding of these exchanges into, and out of, the premises, given the potential emergency nature of security or health alerts included in the payloads. (From a behavioral point of view, persistent and rapid-rate-signaled "alarm clamoring" – until such is acknowledged by the cloud receiving entity – would seem to make sense). This also presumes verbose loop-in of a cloud element, though what is expected is that routine executive control of the premises (as regards IoT) would not likely exit the home edge (and such would be hardened for loss of power and wireline WAN connectivity, at that).

4.3. Data Demands of a Model Smart Home in 2027

Combining several of our expectations, expressed above, for data interchange between home LAN and network WAN, allows us to posit the following model for a smart home in 2027. Note that we will be leveraging Arris' 5300 square foot Wi-Fi house as the representational challenge and will outfit its ~ 2x average home footprint with a full array of CPE Wi-Fi (including sub-mesh Matter support for SED – battery powered -- IoT devices, arranged so that Wi-Fi acts as aggregating long haul for the Thread edge routers commanding these distributed resources). Note that MLO triband capability will be assumed to be implemented in all link endpoints (AP and client alike). We will also liberally equip the home with A/V kiosks (2 per floor, say – though representationally, these might well be phone applications) to support voice control of IoT devices from throughout the premises.

From a human-centric viewpoint, we can suppose nine concurrent clients at peak network utilization, perhaps arranged with the following locations and application demand profiles:

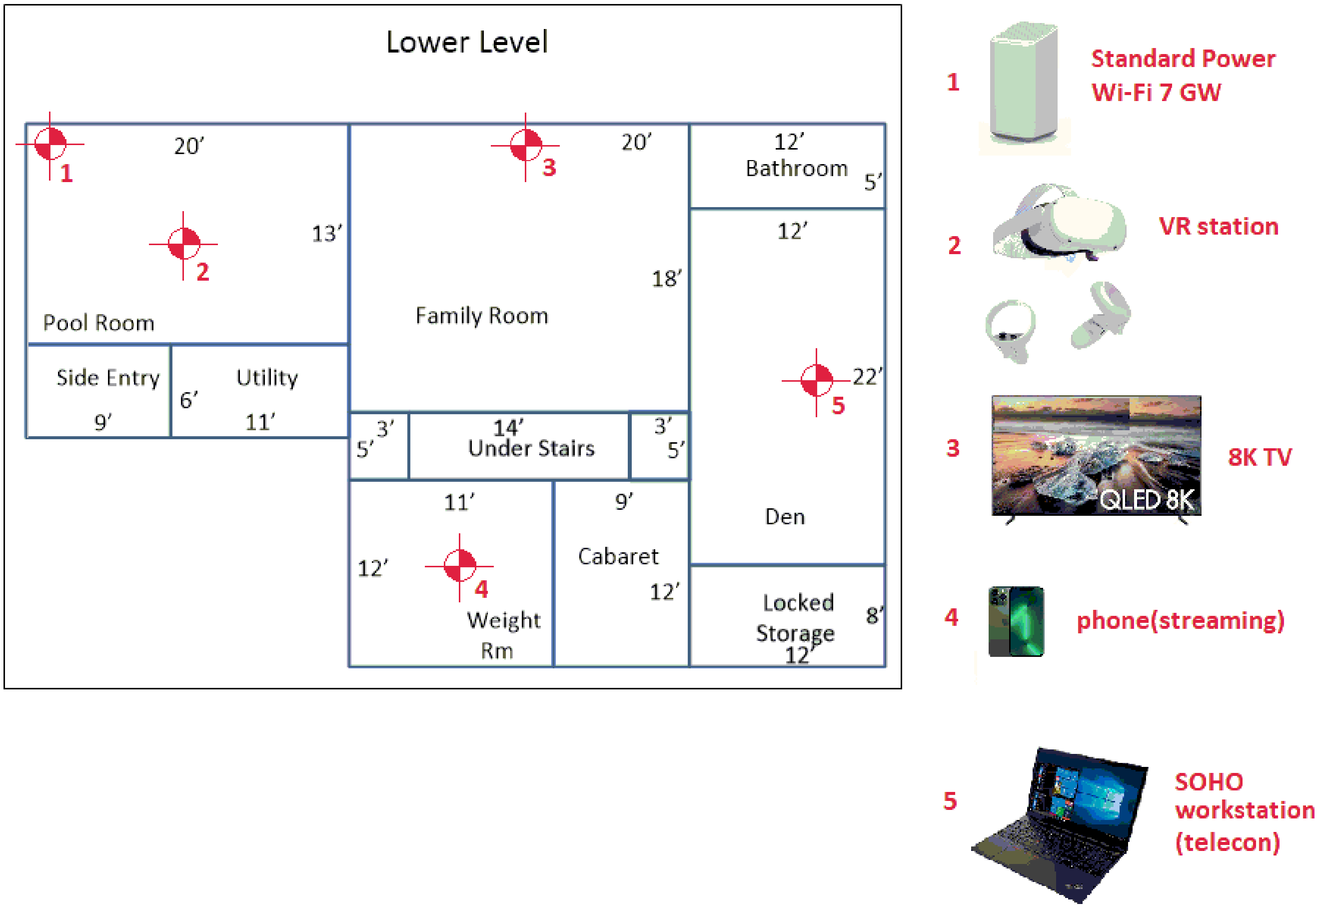

Figure 23 – Wi-Fi House Basement Device Outfitting

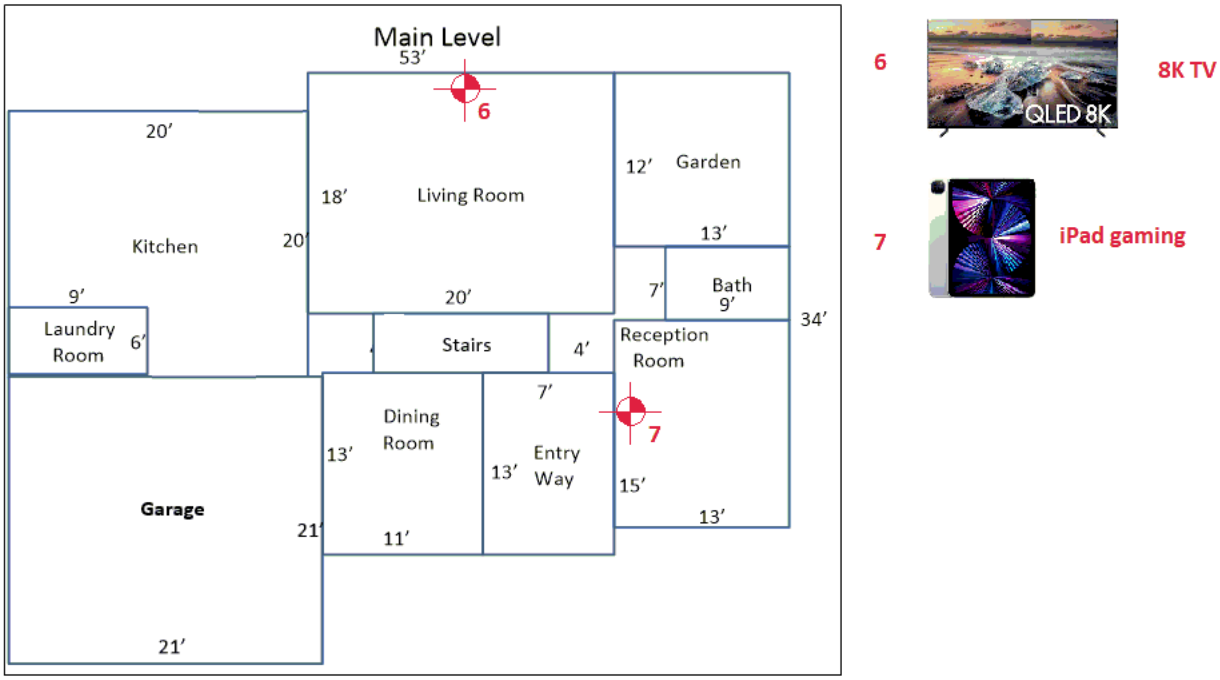

Figure 24 – Wi-Fi House Main Level Device Outfitting

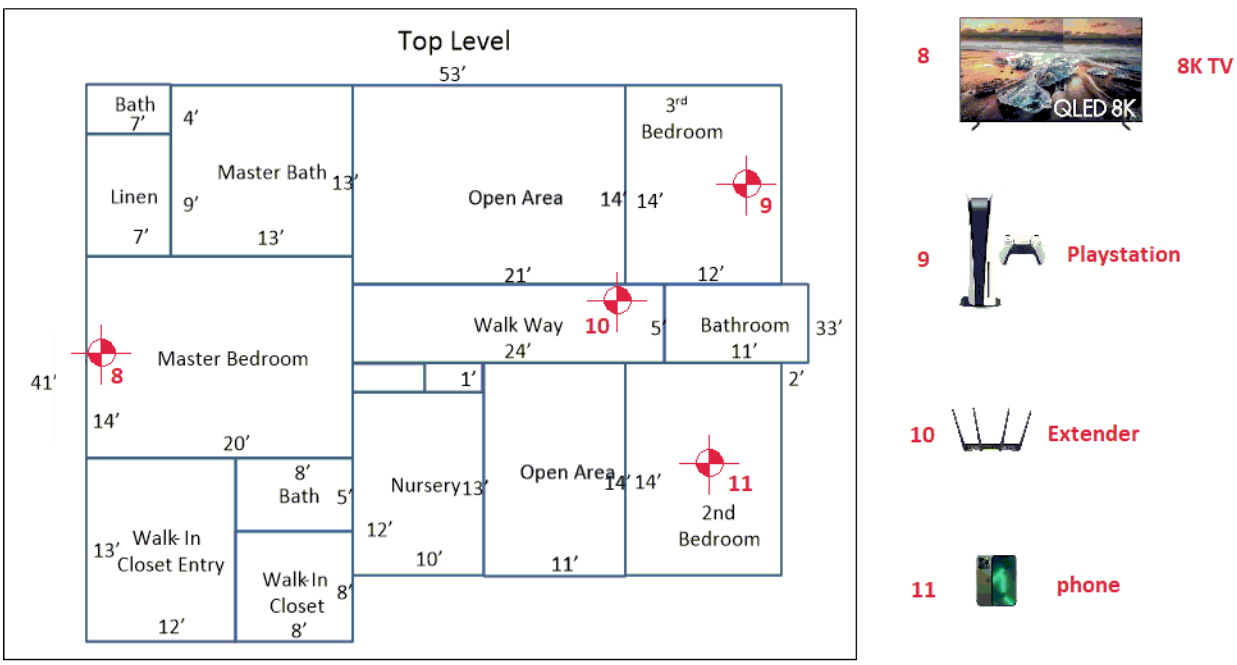

Figure 25 – Wi-Fi House Top Level Device Outfitting

So we can imagine just north of 250 DS Mbps in mounted services, roughly half of which would require the attentive grooming of low latency flow management. In addition to these dynamic client mounts, the home features external security family feeds, internal kiosks for voice command capture and a multi-drop

Thread mesh (for Matter IoT support of automated home security items such as door locks, lighting and garage door control, among others). Note that this latter, alternate-PHY mesh seeks backhaul support via Wi-Fi to the hub (edge) control functions set up in the WAN gateway.

The model home network infrastructure will adopt forward-leaning technologies and so will exhibit a main trunk extension from WAN entry point in the basement game room corner up to the third-floor juncture between the Jock and Jill bedrooms. This will feature a Standard Power and 320 MHz BW channel leverage (at 6 GHz) which will be time-split for attached client MLO/MRU support (essentially dividing client servicing duties with the WAN GW, such split determined by best sustained MCS between AP and client). As all clients will be presumed '11be capable, scheduling behavior will be expected to optimize (minimize) Wi-Fi airtime burn through exploit of appropriate BW in each of the three Wi-Fi bands. Note, however, that for MLO operations, only bonding of 5 and 6 GHz bands will be considered (conservatively capping the available Wi-Fi 7 connections).

5. Performance of the Adjoined Networks

Now we arrive at a discussion of our highly outfitted model home networked to the cloud via its DOCSIS 4.0 WAN wireline attachment. For the purposes of stressing the network attachment, the intention is to array the model with a rather large assortment of Wi-Fi dependent CPE (as described above) and amplify this data aperture by indulging multiple concurrent users with disparate services. We will examine the network behaviors in reverse order, starting with the LAN.

5.1. Stressing the Wi-Fi 7 Home Network

Given the described locations for the types of clients and their respective service mounts, the predicted airtime distribution costs for support of the services array gives us the following:

Device	Location	AP	Path Loss	Link Capacity	Service Bitrate	Service Latency	Service Jitter	Airtime %	Low Latency?
VR station	downstairs pool room	Gateway	10.8 dB	6500 Mbps	100 Mbps	5 msec	2 msec	1.54	yes
8K TV	downstairs family room	Gateway	24.2 dB	5250 Mbps	50 Mbps	250 msec	50 msec	0.95	no
phone (streaming)	downstairs weight room	Gateway	35.5 dB	1800 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.11	no
laptop (SOHO, teleconf)	downstairs den	Gateway	33.3 dB	2850 Mbps*	2 Mbps (up)	5-10 msec	2-3 msec	0.07	yes
8K TV	middle level, living room	Gateway	33.8 dB	4500 Mbps	50 Mbps	250 msec	50 msec	1.11	no
iPad, gaming	middle level, reception room	Extender	25.5 dB	4950 Mbps	350 kbps	10 msec	3 msec	0.007	yes
8K TV	upstairs, master bedroom	Extender	26.1 dB	4900 Mbps	50 Mbps	250 msec	50 msec	1.02	no
Playstation gaming	upstairs, bedroom 2	Extender	13.1 dB	6500 Mbps	350 kbps	10 msec	3 msec	0.005	yes
phone (streaming)	upstairs, bedroom 3	Extender	18.1 dB	4500 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.044	no
4SS MLO Trunk (5G+6G)	pool room to upper landing	AP/EXT	35.9 dB	8250 Mbps	52.7 Mbps	< 2 msec	< 1 msec	0.64	yes
Total, Gateway:			*client upstream EIRP limited					4.42	
Total, Extender:								1.716	
Client Bitrate Service					256.7 Mbps				

Figure 26 – In-Home Wi-Fi 7/Standard Power (with Extender) Network Performance

As can be seen, even restricting the Wi-Fi 7 MLO to bonding only 5 and 6 GHz channels (albeit, at 160 MHz and 320 MHz bandwidths, respectively), the in-home network is not at all appreciably taxed by the 256 Mbps worth of mounted client services. And it should come as no surprise that the intermittent demands of multiple security camera feeds and periodic polling of the Thread subnetwork do not impinge on operations in the measurable slightest. The leverage of standard power along with a quad-band Wi-Fi 7 extender amounts to extreme futureproofing; the WAN gateway commits to just over 4% airtime serving all its clients and the 4SS trunk to the upper far end of the home, while for its part the extender located there commits less than two percent of its time transmitting. Note that, excepting the trunk peer-to-peer 4SS link, all other links amount to 2SS connections (given client radio restrictions); and the 3 uplink feeds' bitrate capacities are all bound by client EIRP limitations (battery-driven, in the two phone cases).

These distribution costs seem extremely well controlled. What if we restrict the gateway and extender to only LPI levels at Wi-Fi 7? The following results:

Device	Location	AP	Path Loss	Link Capacity	Service Bitrate	Service Latency	Service Jitter	Airtime %	Low Latency?
VR station	downstairs pool room	Gateway	10.8 dB	6500 Mbps	100 Mbps	5 msec	2 msec	1.54	yes
8K TV	downstairs family room	Gateway	24.2 dB	4500 Mbps	50 Mbps	250 msec	50 msec	1.11	no
phone (streaming)	downstairs weight room	Gateway	35.5 dB	1750 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.065	no
laptop (SOHO, teleconf)	downstairs den	Gateway	33.3 dB	2900 Mbps*	2 Mbps (up)	5-10 msec	2-3 msec	0.055	yes
8K TV	middle level, living room	Gateway	33.8 dB	3425 Mbps	50 Mbps	250 msec	50 msec	1.46	no
iPad, gaming	middle level, reception room	Extender	25.5 dB	4500 Mbps	350 kbps	10 msec	3 msec	0.008	yes
8K TV	upstairs, master bedroom	Extender	26.1 dB	4500 Mbps	50 Mbps	250 msec	50 msec	1.11	no
Playstation gaming	upstairs, bedroom 2	Extender	13.1 dB	6025 Mbps	350 kbps	10 msec	3 msec	0.006	yes
phone (streaming)	upstairs, bedroom 3	Extender	18.1 dB	2250 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.089	no
4SS MLO Trunk (5G+6G)	pool room to upper landing	AP/EXT	35.9 dB	6000 Mbps	52.7 Mbps	< 2 msec	< 1 msec	0.88	yes
Total, Gateway:				*client EIRP limited				5.11	
Total, Extender:								2.093	
Client Bitrate Service					256.7 Mbps				

Figure 27 – In-Home Wi-Fi 7/LPI only (with Extender) Network Performance

Note that all considerations are still very well met, with only a slight ballooning in airtime for the two AP devices (certainly acceptable, given that the limits at ~ 5% and 2% are still low single digits). So it appears that, with extender, the use of Wi-Fi 7 at the two endpoints of a trunk link limited to LPI EIRP would still work very well. This begs the removal of the extender, coupled with bumping up the single gateway AP (at the WAN injection point to the home) to Standard Power to see if coverage would still be possible:

Device	Location	AP	Path Loss	Link Capacity	Service Bitrate	Service Latency	Service Jitter	Airtime %	Low Latency?
VR station	downstairs pool room	Gateway	10.8 dB	6500 Mbps	100 Mbps	5 msec	2 msec	1.54	yes
8K TV	downstairs family room	Gateway	24.2 dB	5250 Mbps	50 Mbps	250 msec	50 msec	0.95	no
phone (streaming)	downstairs weight room	Gateway	35.5 dB	1750 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.11	no
laptop (SOHO, teleconf)	downstairs den	Gateway	33.3 dB	2900 Mbps*	2 Mbps (up)	5-10 msec	2-3 msec	0.069	yes
8K TV	middle level, living room	Gateway	33.8 dB	4500 Mbps	50 Mbps	250 msec	50 msec	1.11	no
iPad, gaming	middle level, reception room	Gateway	38.0 dB	3700 Mbps	350 kbps	10 msec	3 msec	0.0095	yes
8K TV	upstairs, master bedroom	Gateway	41.5 dB	2575 Mbps	50 Mbps	250 msec	50 msec	1.94	no
Playstation gaming	upstairs, bedroom 2	Gateway	42.3 dB	2450 Mbps	350 kbps	10 msec	3 msec	0.014	yes
phone (streaming)	upstairs, bedroom 3	Gateway	47.6 dB	355 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.56	no
Total, Gateway:				*client EIRP limited				6.3025	
Client Bitrate Service					256.7 Mbps				

Figure 28 – In-Home Wi-Fi 7/Standard Power (no Extender) Network Performance

As can be seen, whole home coverage is still doable – though note the noticeable burn of airtime just to enfranchise the Tik-Tokker in the far upper bedroom. This type of airtime stress will always be the worst – trying to reach a battery-powered device a large distance from the AP and mounting a service thereon which produces upstream data (instead of consuming downstream) – hence being limited to the link MCS achievable with the relatively weak transmitter of the client device.

Still, in all these solution cases, a relatively pathological mix of concurrent services can be mounted and supported by the in-home wireless network built of Wi-Fi 7 componentry. Now the question shifts: can the DOCSIS 4 WAN adequately support these client mounts in reasonably scaled service groups – and of what particular DOCSIS 4 species must we avail ourselves?

5.2. Taxing the DOCSIS Wireline Network

From historical trends, the appearance of a 250 Mbps persistent (and pervasive) home network demand in a DOCSIS service group (SG) is not remotely projected through 2037, based upon current modeling. (Reference ULM et al 2022). Which is not to say it could not exist, merely that capacity planning using present tools has no historical basis to predict the arrival of multiple CPE with bitrate service demands at, or exceeding, 50 Mbps each. This is cataclysmic disruption, relative to present architected capacity and projected bitrate growth; at issue is whether the weighted take-up of these types of bitrate demands within multiple homes comprising a SG produces statistical impact on network service structure before historical trends suggest it might. The best way to judge the scale of this impact, perhaps, is to note that with our present Tavg value of 3.5 Mbps, UHD TV video consumption (the major single data sink per home in the network) produces SG sizes of perhaps 250 subscribers (with a forward BW of < 1 GHz). 250 Mbps represents an 80-fold increase in persistent home data consumption which would have to be defrayed by the combination of increased BW, denser spectral modulation leverage and smaller service groups. Succinctly, even the most aggressive deployment of DOCSIS 4 would find this an untenable challenge. So a bit less of an impact needs to be bitten off on first chew. (In fairness, nine concurrent users with four huge concurrent service consumers of DS data in a single HH is simply much-too-much of a service stretch example when it comes to taxing the WAN – especially if we claim prevalence in a SG, so we will be more circumspect from here on out.)

To establish a more realistic perspective (and to suppose that even 25 Mbps represented a reasonable Tavg DS service goal – note the current value is 3.5 Mbps), we can refer to the Cloonan DOCSIS capacity planning tool for the impact to SG scaling. For reference purposes, we will suppose an ESD solution to 1218 Mhz with a diplex split at 204 Mhz. This gives us roughly 1 GHz of forward BW which, if committed to OFDM parsing only (no SC-QAM) at 9.7 bps/Hz, gets us fairly near the magical 10 GBps downstream capacity.

If we further conjure a top advertised SLA of 5 Gbps, the SG scaling is then set:

Nsub (the SG size) <= (10e^9 – 1.2*5e^9) / 25e^6, or SG <= 160 (if all are at the premium SLA tier – an admittedly unusual occurrence).

This is on the smallish size (range typically varies from 100-400) but at least reads as a very workable number. The 250 Mbps Tavg, on the other hand, would produce an SG of only 16 under the same spectrum and SLA assumptions. To drive home the difficulty of this proposition, if we instead applied even a full 1.8 GHz DOCSIS network with only a modest diplex split of around mid-UHF (say 400 MHz), we would end up with forward capacity in the region of 9.7 * 800 MHz + 8.5 * 600 MHz, or just shy of 13 Gbps. At the same SLA tiering, this would put us at an SG of roughly 28 – still untenable without reworking expectations. Doing so, however – perhaps looking to support a Tavg of 50 Mbps and top SLA tier of 10 Gbps (a magical enough marketing icon) -- the SG moves to 60 (certainly a reasonable asymptote given that not all subscribers will select the premium tier, and the actual number will push into perhaps the 120+ range).

Without doubt, however, WAN network capacity becomes quickly tested when high resolution simulation environments become the norm on the client side; and the capacity "fat" at this juncture in the distribution network lies almost wholly within the LAN environment.

6. An OTT Sidebar

A brief acknowledgement of competitive (5G) forays into OTT connectivity for residential CPE in the US seems in order, if for no other reason than to inject some compounding considerations which might resolve teetering wireline network strategy decisions. As of this writing, two cellular ISPs have introduced 5G home gateway solutions, one a sub-6 GHz band version and the other a more urban-centric mmWave band variant. Given the full-USA coverage of these MNO/ISPs, it implies that both urban and suburban customers are to be granted immediate (as opposed to fiber's buildout-paced) access to premise internet coverage (as a 5G/Wi-Fi composition) which also exposes cable video consumption to co-option via streaming services.

Cable MSOs have carefully marked – and reacted to – telecom rollout of fiber. With fiber's pitch of symmetric 10Gbps connectivity (despite a lack of evidential service need for such BW) and further marketing forays into 25 and even 50 Gbps service tiering, cable has dutifully ramped up development of FDX DOCSIS 4.0 to at least place a bookmark in the derby for bragging rights to what – to-date -- may be called an unreasonably thick data pipe to the residential home. But competition for network connectivity bargains (and perhaps more concerning, in lower tiers of QoS) is already everywhere to be found, courtesy of OTT plays.

The appeal in both OTT cases lies in mining interest among these ISP budget-aware shoppers. The sub-6 GHz solution, for example, features a lowest bitrate tier which produces roughly 150 Mbps down/30 Mbps up with 25 msec latency inclusive of one Wi-Fi hop (one floor vertical) off of the modem for around $50/month (the modem being a GW device which provides dual-band Wi-Fi coverage to the premises). Via contractual agreement for a couple of years, such attachment can produce similar performance at half that subscription rate. A typical midafternoon Ookla SpeedTest sample follows:

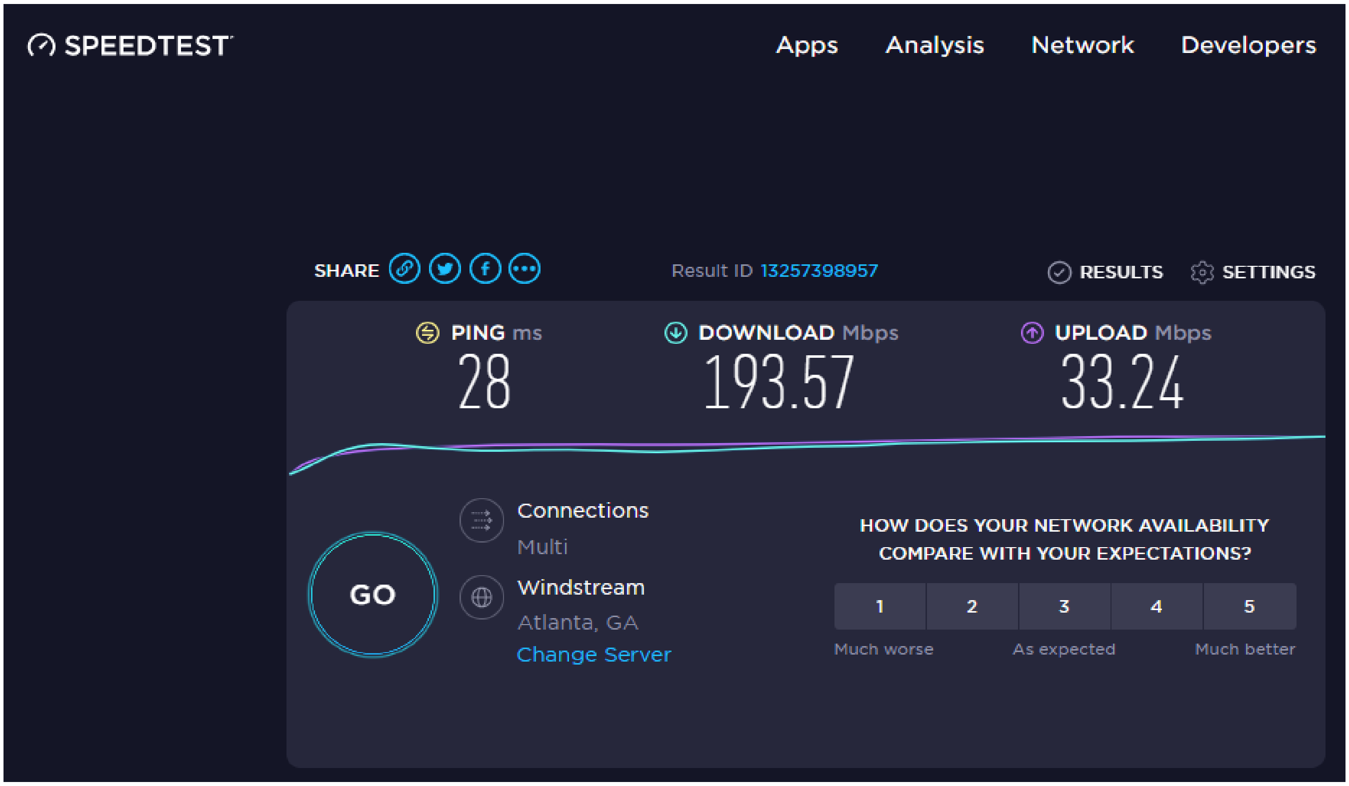

Figure 29 – Mid-day, one Wi-Fi hop OTT modem performance

In MDU and small home scenarios (with perhaps only 2 active Wi-Fi devices concurrently operating), such an arrangement can be seen to provide perfectly adequate internet connectivity, with the benefit of a single point of contact for both phone and network support and logistics. Awareness of this potential for internet fulfillment at less than eye-popping performance numbers ought to serve as a cautionary reminder to not ignore lowest QoS/bitrate markets, as opportunistic competition stands prepared to infiltrate same.

7. Conclusion

As has been shown, Wi-Fi premises bitrate budgets are much more than adequately covered for the near (and coming) term; even pathological, simultaneous adoption cases for extremely high resolution and high frame rate video displays are easily accommodated by home Wi-Fi emerging MACs during the coming five years. With proper leverage of low latency planning and scheduling tools – applied synonymously across the concatenated networks -- immersive simulation environments at modest resolutions can be created which suitably interdict artifacts in the constructed virtual domain(s) and facilitate mounting of rich gaming and instructive services. And there does not yet appear to be widespread adoption of an application or service which swamps the up- or downstream user bitrate flows over home Wi-Fi to the extent that inadequate QoS could be asserted; the home LAN, it seems, is extremely well positioned for all anticipated service mounts.

Packet latency conditioning seems acceptably met in both domains, with LLD capable of dual-piping the WAN and both WMM and buffer congestion management (by virtue of highly flexible scheduling) dissolving any congestion clots within the LAN; married, the E2E performance advertises itself to easily meet sub-20 msec type of E2E latencies just across these networks.

When it comes to raw bitrate, however, WAN prep seems perhaps not congruent with immediate market acceptance of high resolution, high scan rate immersive simulation environments becoming the sudden norm. If this type of precipitous advance in simulation entertainment occurs, then a stepped two-orders-of-magnitude increase (from a *sustained* sub-10 to multi-hundred Mbps) in on-premises average downstream bitrate demand will transpire and network capacity will be challenged as never before -- perhaps tainting even planned 10G/10G FTTP with an "inadequate" QoE estimation.

It's a gambler's paradise, perhaps; truly immersive entertainment options challenge the average home entertainment budget and so might be expected to suppress rapid adoption of the most compelling of lean-back media environments. We have been 8K-capable and pending appropriate source material for several years now, after all – granting WANs some respite in the process. And the CAPEX for system upgrades begs a lag-lag type of adoption curve, certainly. But the penalty for substandard network data support of high-end CPE involves an experiential litmus test which will be difficult to spoof if the long-awaited video nirvana does, in fact, take root. And, to be sure, just as the in-home distribution infrastructure of a very near adoption horizon seems easily capable of underpinning these viewing and gaming experiences, the wireline WAN which most commonly serves it seems, as-yet, resolutely under-capable.

Abbreviations

ABR	Adaptive Bitrate
AFC	Automated Frequency Coordination
AP	access point
AR	Augmented Reality
A/V	Audio/Video
BB	Baseband
bps	bits per second
BW	Bandwidth
CAGR	Compound Annual Growth Rate
CAPEX	Capital Expenses
CCI	Co-Channel Interference
C/N	Carrier-to-Noise
CPE	Consumer Premise Equipment
dB	Decibel
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
E2E	End-to-End
EIRP	Effective Isotropic Radiated Power
ESD	Extended Spectrum DOCSIS
FDX	Full Duplex
FEC	forward error correction
G	Giga
GB	Giga-Bytes
Gbps	Giga Bits Per Second
GW	Gateway
HD	high definition
HH	Household
Hz	hertz

IoT	Internet of Things
ISP	Internet Service Provider
Kbps	Kilo Bits Per Second
LAN	Local Area Network
LL	Low Latency
LLD	Low Latency DOCSIS
LPI	Low Power Indoor
MAC	Medium Access Control
MB	Mega-Bytes
Mbps	Mega Bits Per Second
MCS	Modulation and Coding Scheme
MLO	Multi-link Operation
MNO	Mobile Network Operator
MRU	Multiple Resource Units
Msec	Milliseconds
Mu-MIMO	Multiple User-Multiple In, Multiple Out
OBE	Overcome By Events
OFDM	Orthogonal Frequency-Division Multiplexing
OFDMA	Orthogonal Frequency-Division Multiple Access
OPEX	Operational Expenses
OTT	Over The Top
PC	Personal Computer
PHY	Physical Layer
PON	Passive Optical Network
QoE	Quality of Experience
QoS	Quality of Service
RTT	Round-Trip Time
RU	Resource Unit
SC-QAM	Single Carrier-Quadrature Amplitude Modulation
SCTE	Society of Cable Telecommunications Engineers
SG	Service Group
SOHO	Small Office / Home Office
SLA	Service Level Agreement
SS	Spatial Stream(s)
Tavg	Average Bitrate
TOD	Time of Day
TWT	Targeted Wait Time
UHD	Ultra-High Definition
US	Upstream
USB	Universal Serial Bus
VoiP	Voice over Internet Protocol
VPN	Virtual Private Network
VR	Virtual Reality
WAN	Wide Area Network
WMM	Wi-Fi Multimedia

Bibliography & References

Broadband Capacity Growth Models, John Ulm, Zoran Maricevic, Ram Ranganathan, SCTE Technical Paper 2022