A Tale Of Two Technologies: The Stories Of IPv6 And 5G

Yaakov Stein
Forbes Councils Member
Forbes Technology Council
Sep 20, 2022,07:45am EDT

Yaakov (J) Stein is CTO at Allot.

In recent times, the two technologies that jointly brought ubiquitous online connectivity to the world have confronted existential difficulties brought on by their own success. IPv4 (Internet Protocol version 4) faced address exhaustion, and 4G (the 4th generation cellular technology) encountered bandwidth exhaustion.

In both cases, optimal addressing of the resource exhaustion challenge entailed a new technology, IPv6 (Internet Protocol version 6) and 5G (the 5th generation cellular technology), respectively. Both new technologies were proposed long before the relevant resource was completely exhausted, both were universally accepted as the appropriate long-term solution and, in both cases, short-term stop-gap measures were available to extend the life of the existing technology until the new one could be deployed.

However, the two new technologies experienced quite different adoption rates. RFC 2460 defining IPv6 was originally published in December of 1998, but IPv6 today is still less than half of the internet volume. So, acceleration from 0 to 50% took over 25 years.

On the other hand, while the early release of 5G (called R15) was completed at the end of 2018, despite global setbacks due to Covid-19, 5G population coverage reached 25% by the end of 2021 and is expected to reach 50% before 2025. Acceleration from 0 to 50% is predicted to take about six years.

Why the huge difference in ramp-up times?

Short-Term Solutions

One explanation could be the applicability of the aforementioned short-term measures.

The two limitations of 4G are limited base-station system bandwidth of 20 MHz and woefully inadequate spectrum allocations. 5G sports both larger system bandwidths and new band allocations.

However, 4G learned to cope with its limitations. To overcome the 20 MHz limitation, 4G invented “carrier aggregation” and “dual connectivity” that allow one cell site to combine several system bandwidths. The stop-gap solution for insufficient spectrum allocation is “frequency reuse,” although were it not for 5G, new spectral assets could have been allocated to 4G.

IPv4 addresses (of which there are about 4.3 billion) are a scarce resource, and their exhaustion was understood two decades ago. IPv6 remedies this by allowing for 340 undecillion (a trillion trillion trillion) addresses.

But IPv4 accommodates other addressing schemes, and “dual-stack” is similar in function to “dual connectivity.” The analog to frequency reuse is “NAT,” which allows using the same IPv4 addresses over and over again.

So, the stop-gap measures in both cases are quite similar in functionality and efficacy and cannot explain the huge difference in ramp-up times.

Standardization And Regulation

Were there major differences in the height of standardization and regulatory hurdles?

The main cellular standardization body is the 3GPP, but actual allocation of frequency bands is the job of the ITU (an agency of the UN), which holds world radiocommunication conferences (WRCs) held every three to four years. After bands have been globally allocated, every mobile network operator desiring to exploit one needs to participate in national auctions and pay dearly for every MHz.

The main organization mandating internet protocol standards is the Internet Engineering Task Force (IETF), but the actual allocation of addresses is globally allocated by the Internet Assigned Numbers Authority via five Regional Internet Registries (RIRs).

While the hierarchies are roughly similar, IP address allocations do not have to wait for periodic conferences, and addresses are much less expensive than bandwidth. So, all other things being equal, one would have expected IPv6 penetration to have been faster than 5 G's.

Generation After Generation

There is, however, another subtler difference between the internet and cellular standardization processes.

Internet protocol improvements are traditionally incremental. Everything that works today must continue to work tomorrow, with end-users free to upgrade at their own pace. For example, the latest protocol upgrade (QUIC) is built over one of the oldest layers (UDP) precisely to avoid breaking anything.

In contrast, cellular technologies are divided into distinct generations, each generation occupying a decade. The analog first generation was the only cellular option during the 1980s, being replaced by the digital second generation (2G) in the 1990s. 3G reigned during the 2000s, and our current internet-focused 4G in the 2010s. Each generation represents a paradigm shift and, while required to coexist with the previous ones, is never interoperable with them. Mobile networks undergo hugely expensive forklift upgrades every decade.

The Economic Impact

While the internet protocol suite slowly evolves, cellular discards interoperability every 10 years. And this same mindset follows over to vendors, too. We have all seen people queue for hours to purchase the latest iPhone, although their current one works perfectly well and fulfills all their needs. On the other hand, no one spontaneously upgrades their WiFi access points. And let's be honest—without auto-upgrades, you wouldn't voluntarily upgrade your operating system even to guard against really wicked malware.

Why are people willing to accept this abnormal behavior when it comes to smartphones and cellular generations? It all comes down to hype and advertising. 5G proponents have been telling us that 5G will drive the global economy, creating millions of new jobs and trillions of additional dollars of global GDP. Such hype often becomes self-fulfilling, but it is notoriously difficult to check its veracity.

Takeaways

This same dichotomy arises for all long-lived technologies—infuriatingly slow evolution versus breakneck revolution. For example, one popular family of computer processors slowly evolves, maintaining interoperability at the price of increased size and power consumption, while another stays lean but forces programmers to adapt their code over and over. Similarly, traditional car manufacturers make a plethora of incremental improvements while mavericks went "hybrid," then "all-electric" and soon "autonomous."

Think about the next generation of a technology you depend on or develop. Should you keep the historical baggage and add only critically needed new features, or would you prefer to start afresh, knowing that you have to convince the world to throw out the old and pay for the new?

It is a minor miracle that IPv6 arrived at all, although it would have been impossible to continue without it. And 5G would have arrived at exactly the same time even were 4G to suffer from no disadvantages.

And whether or not 5G succeeds, expect 6G to arrive around 2030!