www.pipelinepub.com
Volume 19, Issue 2

Preparing for the Metaverse

By: Mark Cummings, Ph.D.

The metaverse will create a set of new demands on our communications systems that are different from what previous changes in technology and living conditions created. The advent of the metaverse will bring stringent demands for reduced latency that will accelerate a move away from Transmission Control Protocol/Internet Protocol (TCP/IP), which may result in a return to circuit switching or something similar. There is likely to be a stream of new standards and products that have a foundation drawing heavily on circuit-switched concepts.

These new standards and associated products, though, will take some time to get to large installed bases. As metaverse applications take off, communication service providers (CSPs) and hyperscalers are likely to face demands for low-latency infrastructure. One way to fill the gap before new technology solutions appear is to dust off the expertise in and software support for virtual packet switching. Making the existing installed base perform with minimum latencies while onboarding new technologies will require fine granularity orchestration. CSPs, cloud providers, and the vendors supporting them should be preparing for this now.

Previous network challenges

Two big changes our networks faced in our recent past were:

- Moving from email and web browsing to watching movies over the Internet.
- Explosion of Zoom (and similar services) triggered by the pandemic.

The first was handled by the deployment of local points of presence where the content was cached closer to the consumer. The second generated an overall increase in traffic, but there were generally relaxed quality of service (QoS) constraints. We came to accept that there would be:

- Delays between when someone spoke and another party heard what was said.
- Difference in the timing of seeing lips moving and words heard.
- Static images of people's heads and shoulders.

- Frozen screens
- People dropping off and coming back on.
- Some countries and network operators blocking services.
- Problems with privacy and security.

Metaverses present network challenges that are fundamentally different.

The new network challenges of the metaverse

In a metaverse, people from all over the world want to interact with each other and virtual environments in a fashion similar to how they interact with each other in real-world environments. Users will connect via interfaces such as haptic interfaces, gloves, glasses, body-tracking sensors, static full-motion video, moving full-motion video, and more. Users may be abstracted as avatars, represented by real-time images of themselves, or combinations of these two. There will be an expectation of full-body images that move and interact in real time.

Similarly, environments may be:

- Fully designed and created virtually.
- Naturally occurring static spaces.
- Naturally occurring dynamic spaces captured in real time.
- Combinations of all of these.

Whether virtual, natural, or combinations, environments may have inputs coming from many sources around the globe. Not only must users interact with each other, but they also must interact with their environments. Some of these interactions may be obvious, such as picking up an object. But others may not be so obvious. For example, the wind must cause a virtual character's (avatar, real, or combo) clothes to flutter. Light and shadow must interact. Waves must break around things in the water. Why the use of the plural "metaverses"? Because users are likely to interact with portfolios of metaverses. For example, in a metaverse version of a conference, there may be a lobby environment, several auditoriums for presentations, some small breakout rooms, a coffee shop, a restaurant, and more. As virtual conferences develop, they may be placed in virtual cities with all the environments that implies. The cities may be actual cities that change in concert with local weather, news, and more. Or they may be imaginary cities on other planets. There may be different platforms hosting different metaverses. There may also be technology differences. Whatever evolves, there are likely to be tightly or loosely linked groups of metaverses.

Where will users be located when they are participating in metaverses? In scenarios currently being discussed, actual users may be inside physical buildings, specialized facilities, outdoors, vehicles, and more.

There may be some entirely local metaverses. But even for local events, there will be a pull to go global. Let's take "Hardly Strictly Blue Grass" (HSBG) as an example. Today, HSBG is a three-day series of free concerts held annually in San Francisco's Golden Gate Park. During the pandemic, part of it was held online. Musicians still traveled to a stage in San Francisco (at a secret location), but their performances were available on the web and watched by people all over the world. The pandemic produced this as an emergency response without an existing metaverse infrastructure to build on.

Let's imagine that HSBG decides to go to a metaverse. Now, the global audience can interact with each other. The musicians don't have to travel—in fact, the musicians don't even have to be co-located. There may be vendor booths that users can enter to purchase music, swag, and more. At some music festivals, there are food vendors. You could imagine a food delivery service partnering with the festival to deliver food to audience members. HSBG is free and held during daylight hours. In the evening, many of the musicians in town for HSBG play in clubs and halls—these venues could also be part of the metaverse. This is just an illustration, but when creative people fully engage with metaverses, the results will probably be well beyond what can be imagined now.

Many special events like HSBG, plus ongoing virtual worlds, will dramatically increase communication volumes. But that is not too hard to handle with today's technology. The greatest challenge will be to meet the latency requirements. For example, such a metaverse HSBG could be hybrid—that is, both in person in Golden Gate Park in San Francisco and virtual with attendees and performers from around the world. Delivering the performances from the physical stages could be handled similarly to streaming movies or Zoom. However, virtual participation, particularly concerning performances that combine both physically present and virtually distanced performers, will create demands for very low latency. These demands for very low latency will be difficult for today's packet-switched networks to meet.

Packet switching and latency

Packet switching played a key role in our transition from analog to digital communications and is a fundamental part of our digital infrastructure today. The most common form of packet switching in use today is built on TCP/IP. There are many benefits enabled by packet switching. But there are costs as well—and one of the cost factors involves latency.

To understand latency, it is helpful to use an oversimplified description of how a packet network works. This description is designed to make sense to people not familiar with the details of packet networking and to be consistent with the basic understanding that those with detailed technical knowledge have. In a packet-switched system, every node in the network has to process the packet. This means, at a minimum, the node has to read the top-level packet address and make a decision about what next node to send it to. Nodes may also read some, or all, of the progression of lower-level

addresses (packets are like a letter sent to a friend who then puts the received envelope in another envelope and sends it to you). The node may also do error detection and or correction and other processing. The node may also have to generate a log message for every packet it handles. It may have to wait till the log buffer has loaded a previous log entry to the log database before generating the next log message. The semiconductors that support this processing may be very fast. Still, it takes time—and that time compounds with the number of nodes that have to be traversed. Modern networks have a very large number of nodes.

Circuit switching and latency

In a circuit-switched network, intermediate nodes do not have to do this kind of processing. Instead, at the beginning of the communication session, there is call setup processing. Call setup involves a set of messages and responses to those messages that tell intermediate nodes how to arrange themselves to allow data to pass to its intended destination. Call setup takes time—but, once call setup is done, there is little or no processing delay in the network. Thus, circuit switching is transmission efficient when communications sessions are long in duration and volume-heavy.

Virtual circuit-switched networks

Virtual circuit-switched networks were developed and widely used in the early days of packet switching to ease the transition and in the middle years to improve efficiency. Those mechanisms developed in the middle years are contained as legacy capabilities in many nodes today. When activated, they examine flows of packets. When they find that a high level of traffic entering one node is always bound for the same second node, the system stops looking at addresses and forwards all packets to that second node.

The metaverse and latency

As we can see in the HSBG example, metaverse interactions are likely to have the key characteristics that make it hard for packet switching and easy for circuit switching. The added latency created by all the internode processing in a packet network adds significant latency to the nonreducible latency from propagation delay. Because once started, metaverse interactions are likely to continue for some time, they are not so concerned by call setup latency. Also, once set up, they are likely to have significant volume—thus, reducing concerns about fill efficiency.

Likely outcomes

We are already seeing the beginnings of a search for alternative networking approaches. Many people working on 6G cellular are working on alternatives to TCP/IP. We are likely to see similar efforts in the fiber and other terrestrial network technologies. These are likely to have a foundation drawing heavily on circuit-switched concepts. But these new standards and associated products will take some time to get to large installed bases.

As metaverse applications take off, CSPs and cloud providers are likely to face demands for low latency from existing infrastructure. One way to fill the gap before new technology solutions appear is to dust off the expertise in and software support for virtual packet switching. Making the existing installed base perform with lower latencies while onboarding new technologies will require very good intersystem cooperation between all the components in large networks. Large networks are going to become more complex and volatile. The only way to achieve that kind of intersystem cooperation is with fine-grained orchestration. Quickly developing and fielding this fine-grained orchestration is critical. CSPs, cloud providers, and the vendors supporting them should be preparing for this now.

Note: the author was a very early and strong proponent of packet-switching. He contributed to the first packet-switching standard (X.25) and designed the first international packet-switching network. He continues to work on standards for, and designs of, semiconductors and software for LANs, WLANs, WiFi, and cellular using packet switching as appropriate. He sees the evolution of technology as having pushed many portions of our information technology space to and beyond the edges of what packet-switching technology can deliver.