

A background image showing a blurred perspective of a road or tunnel with blue and white light trails, suggesting speed and technology.

Time to embrace the Any-to-Any shift

By Brian Doricko

CyrusOne vice president
of Strategic Sales

When the COVID-19 pandemic forced the world into lockdown mode in March 2020, it put unexpected stresses on many of our systems. Businesses of all kinds had to figure out how to continue to serve their customers, employers had to figure out how to keep their people working and consumers had to figure out how to get the products they need for everyday living. Over a year later, we've proven once again how resourceful we are as a species.

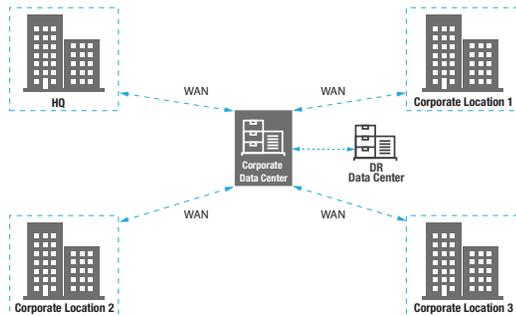
But this massive disruption also gave us an opportunity to test in real-time and in an escalated way how the IT network would hold up with unprecedented demand as we stayed home for work and school and ordered our consumer goods online.

Traffic patterns for our external and internal clients changed in a dramatic, if not violent, shift. Prior, mostly offices and internal resources accessed systems in predictable ways from predictable locations. During the pandemic, it turned into an

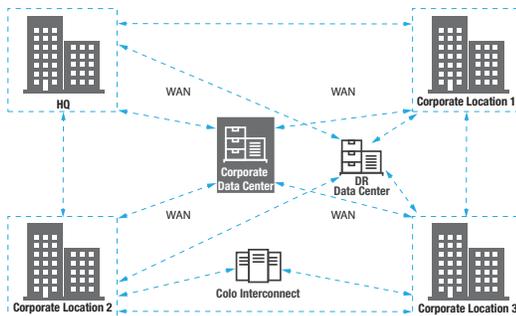
Any-to-Any flow for internal systems, cloud and hybrid cloud. This included all communications types, including internal users, the external channels like supply chain partners, SaaS, and digital services, and most certainly customers, consumers, work/school-from-home users.

While it has allowed and supported access, the current hub-and-spoke system often forces traffic to take longer, slower paths than the direct routes Any-to-Any boasts. Today, traffic patterns usually go from a remote worker to a corporate network location and then out to access another service elsewhere. We in the global data center industry remain resolute and resilient while playing a crucial role in ensuring access to critical information and services, as well as supporting key parts of the economy. But we also need to help clients deploy infrastructure that minimizes latency and distance between routes. This may mean hub-and-spoke must evolve to Any-to-Any.

1980 - 2000 Hub & Spoke



2000 - 2015 Hub & Spoke + Interconnection & Colocation



The architecture going forward must be even more efficient, and data center providers must redouble their efforts to build more capacity across key markets as cloud adoption accelerates. To enable more services and meet rapidly growing demand, the investment into a longer term Any-to-Any architecture becomes critical.

We simply need to build smarter going forward. And we need to ask ourselves four important questions:

1) How do we optimize connectivity?

Every successful project begins with well-defined goals, acknowledgement of existing limitations and a road map containing a clear start and finish. So, if the objective is to evolve IT in a way that enables the business to best meet its corporate and shareholder goals, we must optimize IT from technical, economic, security and compliance, and flexibility standpoints.

The current state of our connectivity infrastructure might be best described by “Heisenberg’s Uncertainty” Principle: “The more certain we are about our location, the less certain we can be of where we are going. Conversely, the more certain we are about where we are going, the less certain we can be about where we are.”

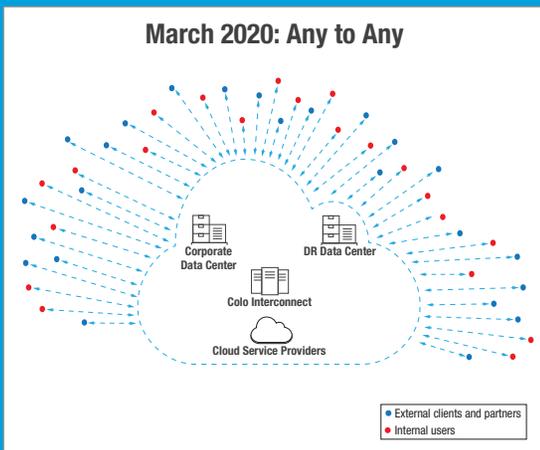
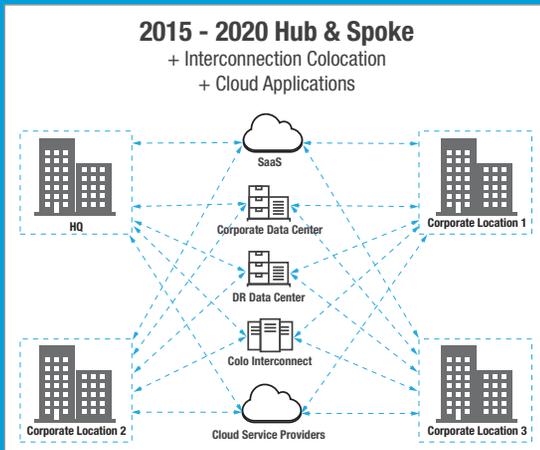
He was describing particle physics, not packet flow. Yet, his wisdom seems to hold true in this new normal.

Many organizations recognized this crossroads in 2020. I’m not sure anyone can claim they predicted what was going to happen in 2020. No one planned for a lockdown, school closures, working from home, the elimination of most types of travel in total or in part, and all the other challenges the pandemic forced upon the world.

Existing infrastructure is primarily designed around a hub-and-spoke legacy architecture. This worked well in a world where communication was predictable, particularly about where traffic flows were. It resembles the modern airline transportation system or old Public Switched Telephone Network (PSTN) infrastructure. But we no longer build hub-and-spoke phone networks. So why would we continue to build IT networks this way?

Today, organizations of all sizes have started to develop cloud, hybrid-cloud and SaaS strategies with a clearer understanding of which applications best fit into which model. The lockdown means data and applications are now Any-To-Any, not hub-and-spoke. This applies to clients who access enterprise applications and internal users, and it applies to the-system to-system cloud, SaaS, grid, compute farms and more.

The idea of “following the packet” (the information that’s travelling) seems the best way to begin to steer future architecture investment toward Any-to-Any, which allows lowest latency and cost available. And perhaps not adding more, bigger, faster to existing infrastructure is also a good starting point.



Organizations, vendors, cloud and SaaS providers, major compute players, and security and compliance organizations also seem to align to the idea that we should listen to the client because it's in the best interest of meeting enterprise goals.

Organizations of all types should try to define the packet flow within their infrastructures. Increasingly, Any-to-Any seems to be the rule, which makes the central traffic flow of the past less than optimal from a latency standpoint, less flexible and more expensive.

Imagine if cellphones routed all communications through central switching offices the way the PSTN did? Instead, we have cell towers, cable-enabled WIFI, fiber meshes and technologies like 5G ramping up. Further, security and encryption have evolved rapidly with a “Zero Trust” policy governance that defines who accesses what, and when, based upon that user’s profile, wherever they may be.

2) How will customers use the hybrid cloud?

The world is hybrid. Your users talk to each other anywhere and everywhere all the time.

Some businesses store some of their data in the cloud, some in their own data centers and some in data centers built by third-parties like CyrusOne. But all must serve a sophisticated pattern of needs for businesses on their multi-year transformation across the “Cloud Continuum.”

Enterprise businesses are in varying stages of their journeys to find the right solution for their cloud or IT infrastructure. Some may have just started

migrating to the cloud. They discover quickly that changes in technology, applications, business, and regulatory and compliance standards can make accurately forecasting future data center requirements challenging. Others are years-in experts who know their needs for today and tomorrow. They are ready to scale and adjust to meet evolving data center requirements.

[What can we say in just a graf or two that connects hybrid cloud usage/how customers will use hybrid to our notion we need to move to Any-to-Any?]

3) Where does edge computing fit in?

Edge computing is a system that brings computing and data storage closer to a location where it's in demand to create better response times and conserve bandwidth.

For the world of data centers and IT, a more client-focused definition around service enablement is more appropriate. The edge delivers digital services to any user in the most technically and economically efficient way, without compromise for security.

Consider which is faster, a Ferrari or a dump truck? Most would say a Ferrari. We might differ. Networks are not speed limits or measurements on speed. They are a measure of how much payload (data) can be delivered in a period of time.

So, return to the question of which is faster. What If we had to deliver a 2-ton payload over 10 miles? Now, which is faster? Obviously, the dump truck.

In the same way, if the speed of light is 186,000 miles per second, reduced in fiber to about 125,000 miles per second, this means light (data) can be delivered 125 miles in 1 millisecond, or 750 miles in 5 milliseconds. This obviously assumes clean fiber pathways. It also makes sense when seeing the massive investments in fiber the biggest tech, cloud, SaaS and social media organizations make. And the massive facilities they run in states, such as Iowa, Oklahoma, Oregon, and Minnesota – not exactly the places we think of to best serve users. And why does that make sense to the tune of billions of dollars? Because, Einstein was (mostly, setting aside new quantum paradigms) correct.

Physics proves the massive economic advantages of going to major lower-cost, central data centers is often the better choice for building future architecture. We might be prudent to carve out applications that have a “first-or-nothing” job to do. This means things like trading – if your competitor gets there first, you lose on the revenue from that trade. Or if you serve adds online, it means if someone jumps in front of you to serve and add, you miss that opportunity.

The above is especially true when considering all the costs of small deployments compared with large ones. This shift has already begun. The heavy investments in fiber by the biggest cloud and SaaS companies to connect back to the massive data centers in places like Iowa, Minnesota, Oklahoma, Oregon or Nevada, as we mentioned, bear this out.

4) Finally, what should flexibility look like?

Every customer and enterprise are on some different level of needs and usage, so being flexible is the key to servicing clients today and in the future. No two companies will have same needs and no two companies will get the same solutions. Flexibility should be in the design architecture, software-designed networks, which is more economical in the long term. The more flexible, any-to-any architecture works best whether it's an Global 1000 or a hyperscaler.

In summary, flexibility is optimizing the ability for an organization to deliver online applications and services anywhere, anytime. This demands an architecture that drifts away from the legacy hub-and-spoke model that served us well for a couple decades and moves into an Any-to-Any world.

Have you followed your packets yet?