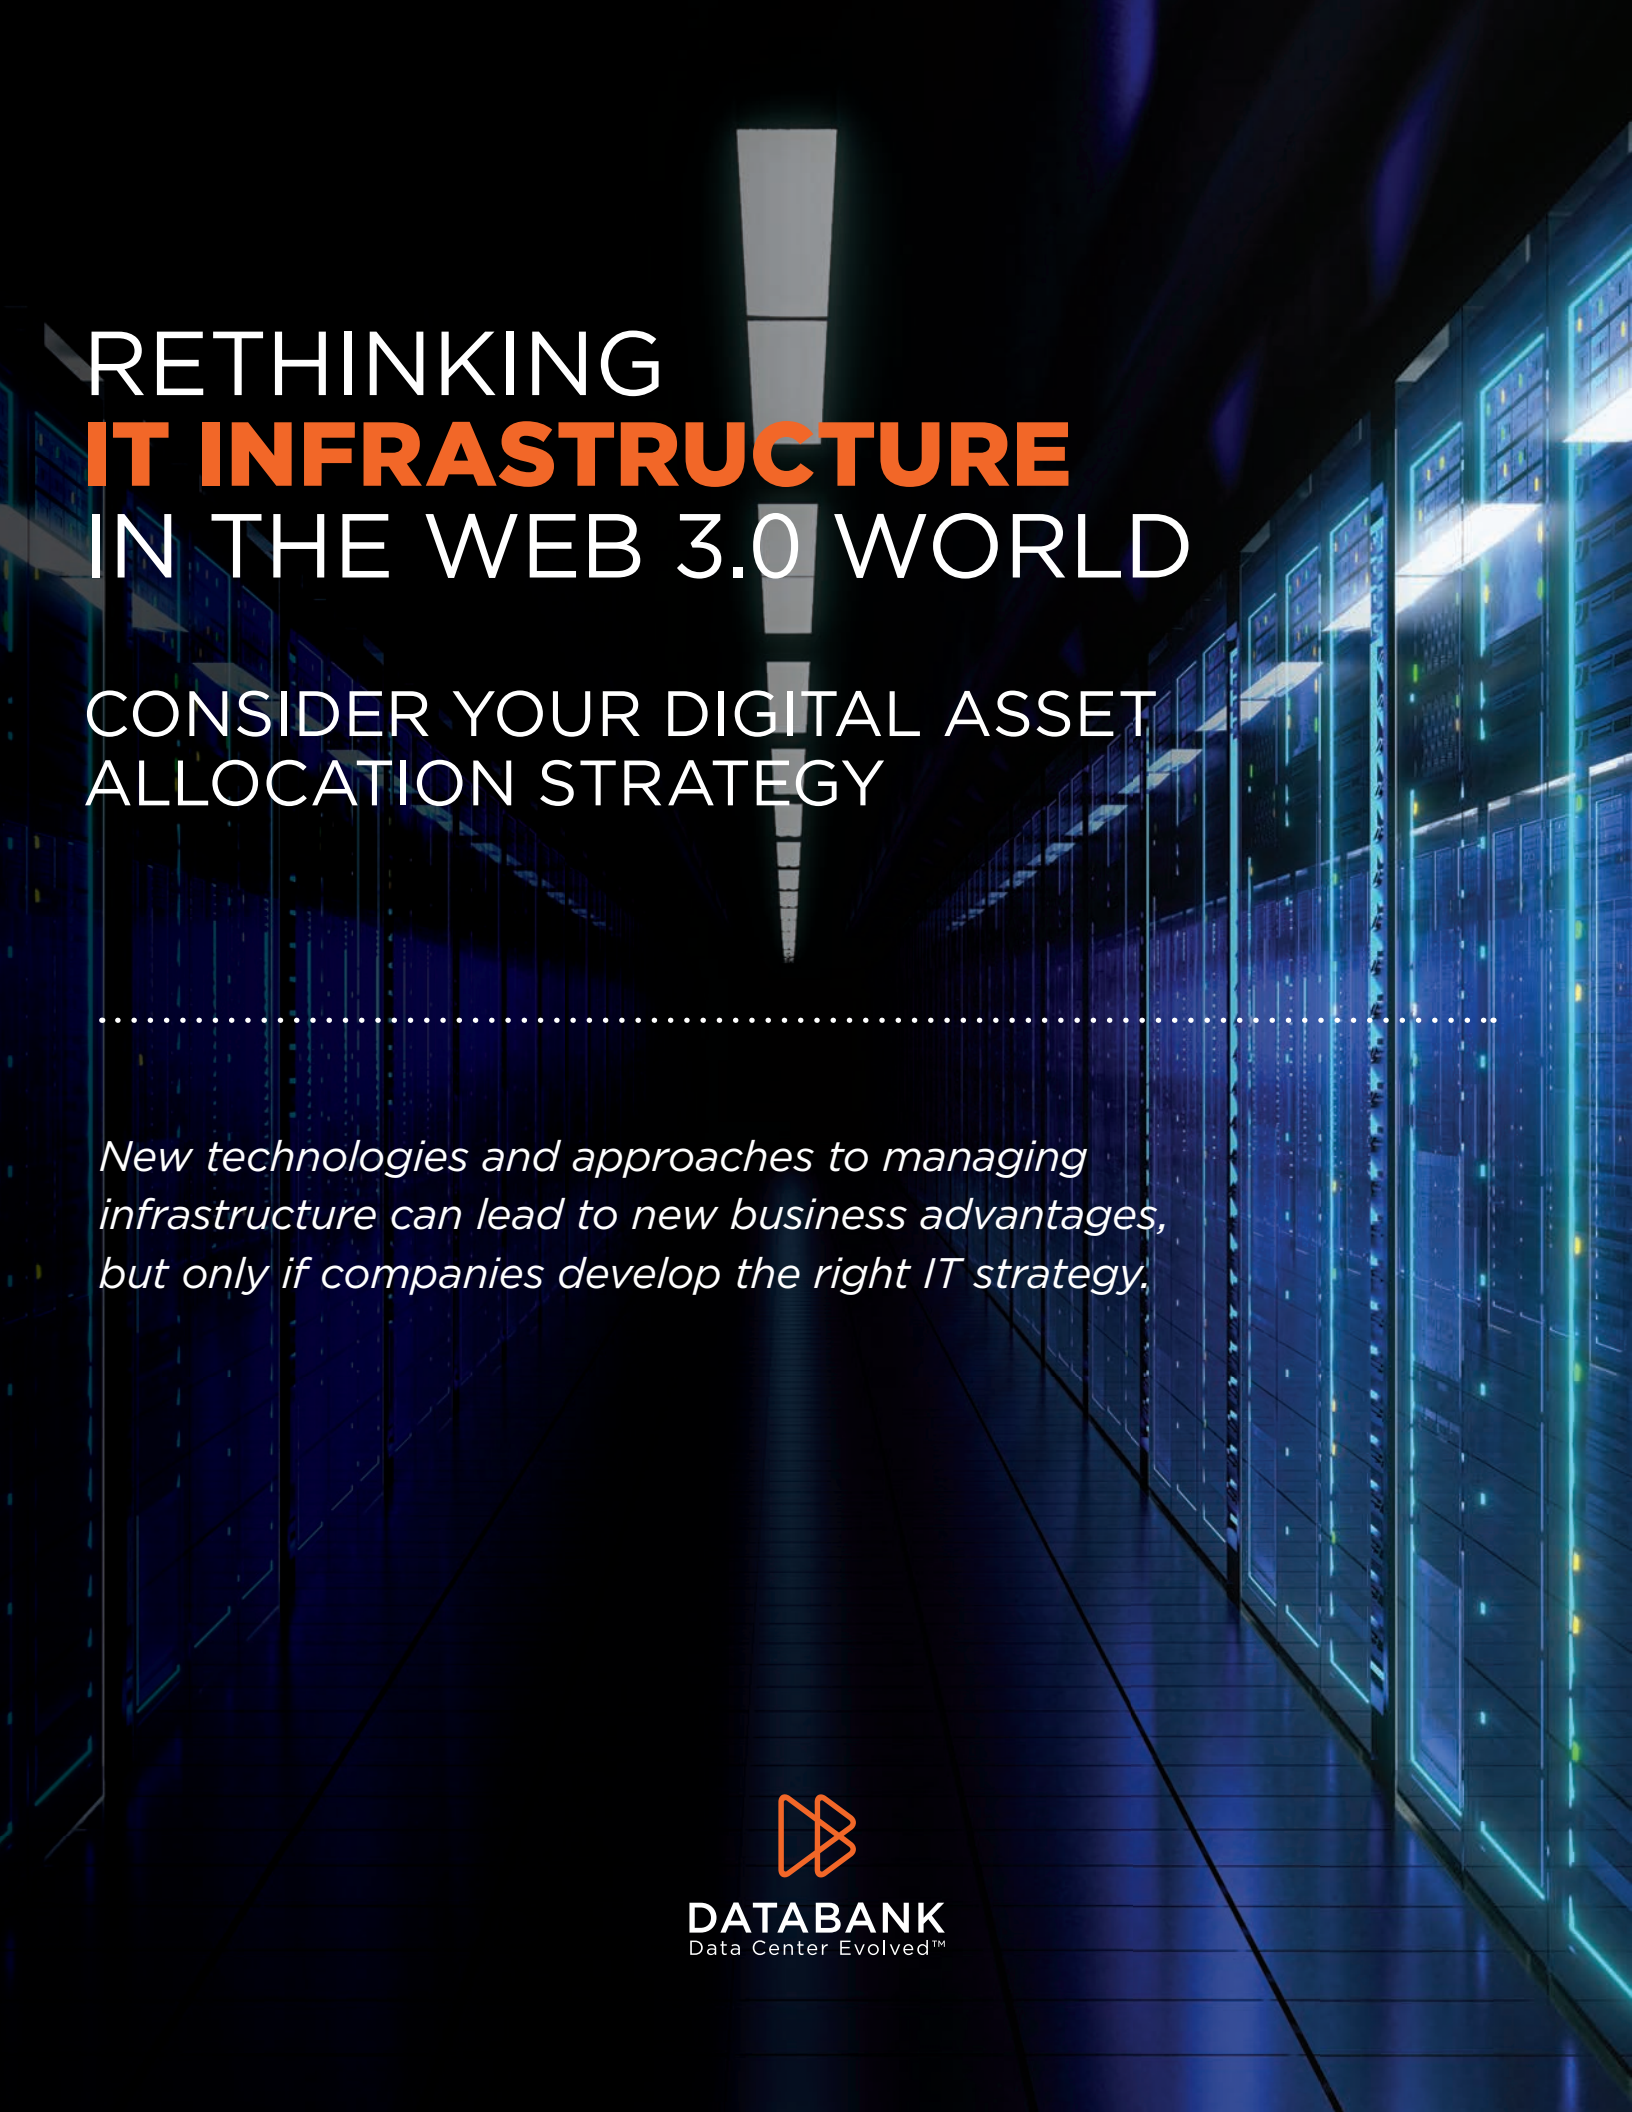

RETHINKING **IT INFRASTRUCTURE** IN THE WEB 3.0 WORLD

CONSIDER YOUR DIGITAL ASSET ALLOCATION STRATEGY

New technologies and approaches to managing infrastructure can lead to new business advantages, but only if companies develop the right IT strategy.

DATABANK
Data Center Evolved™

Table of Contents

Executive **Summary**

Web 3.0 and other innovations present real business opportunities for those organizations that understand how IT infrastructure will become more agile and edge-oriented as a result. In order to capitalize on all that Web 3.0 has to offer—and get a jump on the competition—companies of all sizes must examine their existing infrastructure, related interconnections, and their overall strategy related to their entire architecture.

This white paper introduces the concept of a "Digital Asset Allocation Strategy" and how it can be used to adapt to the Web 3.0 era. The paper also discusses software-driven infrastructure and the ubiquitous edge—additional topics today's leaders must understand to shape successful IT strategies.

Welcome to **Web 3.0**

Today, we have officially entered a third phase of the internet—Web 3.0.

As background, Web 1.0 was the first stage of the World Wide Web revolution, complete with static web pages and a model where most users consumed content but didn't actually create it. Web 2.0 marked a shift to user-generated content and a new focus on community-based collaboration and information sharing.

Web 3.0 builds on these important evolutionary steps but is now driven by a number of important technological trends:

- Thanks to IoT, everyone and every device can now produce data (and content).
- The highly anticipated 5G network standard promises to significantly increase mobile bandwidth and give consumers faster, easier access to applications and data.
- The emergence of hyper-local, latency-sensitive applications requires new strategies to provide the best services to customers.

In order to capitalize on all that Web 3.0 has to offer—and get a jump on the competition—companies of all sizes must examine their existing infrastructure, related interconnections, and their overall strategy related to their entire architecture.

Web 1.0

Web 2.0

Web 3.0

The Need for an **Effective Digital Asset Allocation Strategy**

As leaders plan their IT strategy, there are two major variables to consider.

1. First, infrastructure today is much more agile and software-driven. We're moving away from traditional telecommunications-based models, where assets and connections were once physical entities, provisioned slowly, and managed by the provider.

   Now, thanks to digitally-driven models, these same assets are increasingly virtual, provisioned instantly and on-demand, and managed by the users themselves.

2. In addition, we're seeing a shift to a geographic, localized edge model that provides faster, better services to customers, consumers, and end users. However, it's important to note that the edge is not simply a matter of geography or distance; instead, it's a ubiquitous edge that can exist in different modes.

   Sometimes the edge will be geographic, extending from Tier 1 cities to more rural markets. It may also be application-specific, depending on a company's core business and application latency-sensitivity, which can vary depending on whether it provides content, networking, or computing.

All of this points to the need for a well-defined digital asset allocation strategy, which is a topic we'll cover in this white paper. First, it's important to take a step back and look at how we got here.

Digital Asset Allocation

Your infrastructure, apps and data on the right platform, in the right place, at the right time

Digital Asset Allocation

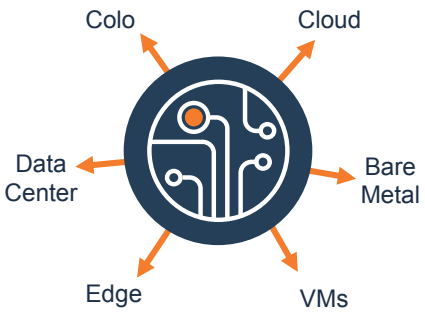

Financial Asset Allocation

Infrastructure Becomes **Code**

Software-defined networking (SDN) first emerged in the early 2010s as a way to make IP networks more flexible—by centralizing the intelligence and control of an inherently decentralized system. Initial SDN platforms focused on wide-area networks (WANs) and provided a mechanism to abstract or separate the forwarding process of network packets (the data plane) from the routing process (the control plane).

At the same time, and perhaps not unsurprisingly, a similar process was taking place in the world of computers. Memory, processing, and storage all became abstracted in early virtual machine platforms such as VMware. In both cases, *software* was the driving factor.

More recently, this virtualization has begun to take place in the data center itself. As the use of public cloud platforms has grown, IT infrastructure managers have grown accustomed to automation and the control the cloud enables, especially for the various workload connections. Cloud on-ramps have thus emerged, allowing managers to provision connections to and from their infrastructure in data centers and multiple cloud platforms.

As multi-tenant data centers like DataBank have increasingly become neutral network interconnection hubs, data center operators have upgraded their portals to allow customers to order and provision virtual cross-connects between customers, workloads, clouds, data sources, and applications. Again, software is at the heart of it all, effectively transforming infrastructure into code.

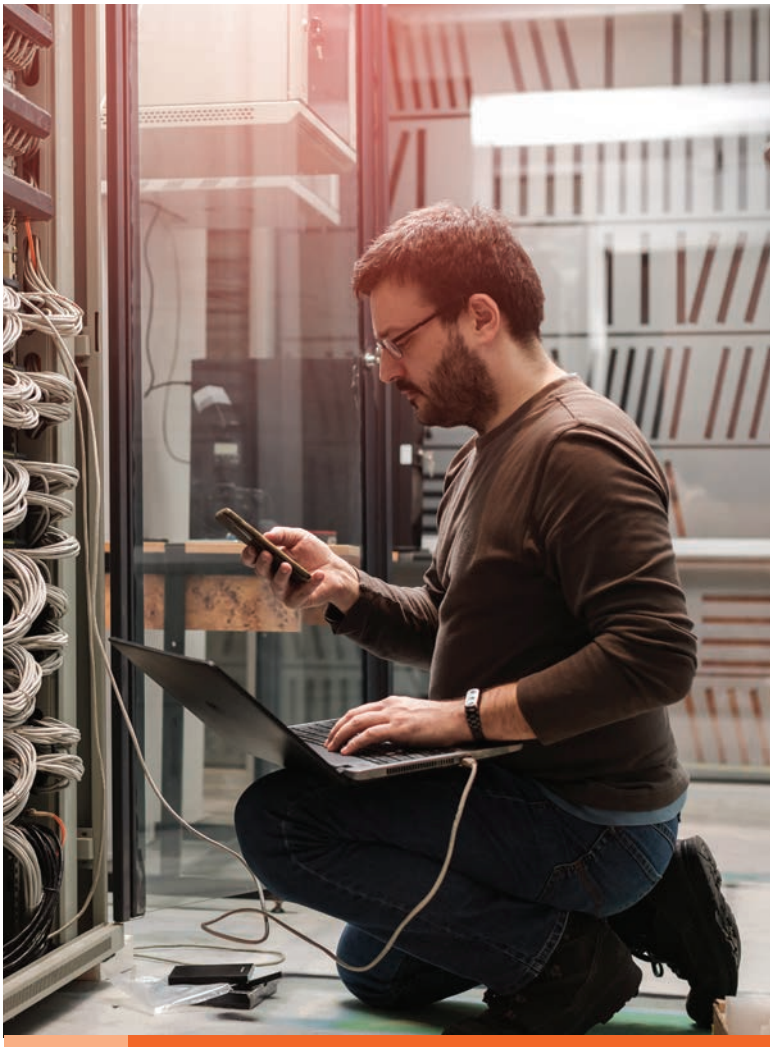

Software-Defined **Data Platforms**

We're just beginning to understand the full implications of the emergence of software-defined data platforms. What we do know is they're clearly going to result in an inherently different user experience for IT managers. In particular, our awareness and perception of individual infrastructure components will certainly change.

Rather than having to think about the many discrete, finite infrastructure elements involved from multiple providers—servers, storage arrays, data centers, network ports, and devices—infrastructure managers will begin to look at them as one giant pool of programmable resources. They can plug into (or out of) these resources on-demand and manipulate them from a single pane of glass, or even the palm of their hand.

This has already happened in the public cloud where applications may be running on many different virtual machines. Infrastructure managers, however, do not care exactly which servers or where they are located. As far as they are concerned, a virtual server has been spun up to run the workload, and if the capacity of a virtual server is exceeded, it can burst out to another.

Similarly, in the data center, a software-defined platform lets network operators and other users create connections without awareness of the specific switches or routers involved. One network port can be quickly reconfigured to connect to a different network port, and the bandwidth for those connections can be increased or decreased as needed, limited only by the physical port size.

Infrastructure as LEGOs™ vs. Infrastructure as Minecraft™

For anyone old enough to have teenage kids, software-defined infrastructure is not unlike the evolution in how we used to build with LEGOs and how our kids now build with Minecraft. With LEGOs, our imaginations were limited by the discrete, physical blocks scattered across the floor. With Minecraft, our kids have a virtually unlimited ability to create anything and even program new building blocks as they see fit (With no sharp blocks for parents to step on).

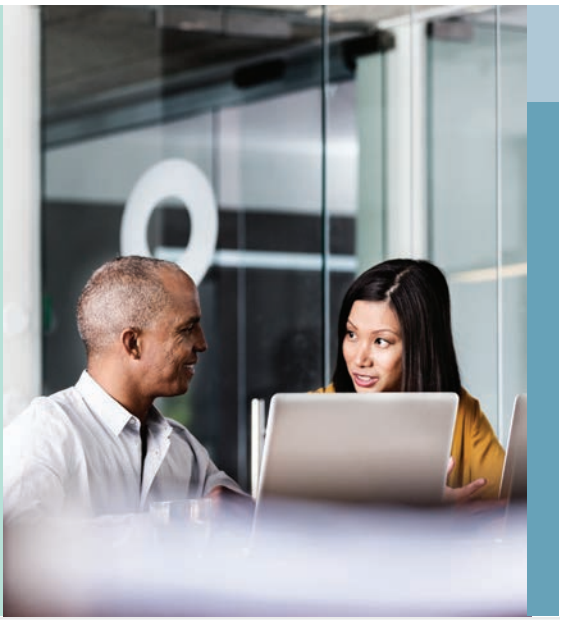

Paving the Way to a **Better Experience**

Not only will the software-defined experience be different, for IT managers, it will be better.

A number of things happen during the shift from a telco-centric model of infrastructure to a software-centric model of infrastructure:

- Formerly fixed and centralized facilities like data centers and interconnects become more distributed and flexible.
- Physical connections and cross-connects become more virtual.
- Slow, manual provisioning and changes to networks become instant and programmable.
- Instead of service providers managing everything, IT managers can self-administer their infrastructure.

Looking further into the future, it's possible we might see human intervention removed entirely. For example, applications may be self-aware enough to diagnose the need for additional connectivity, engage the network through APIs, and establish additional bandwidth or connections on their own.

Such an experience will lead to real business benefits:

1. Greater agility: IT leaders will be able to better adjust and align their infrastructure with the demands of their applications in real-time, eliminating the need to overbuild capacity.
2. More efficiency: Organizations will be able to automate provisioning and management efforts, and reduce the burden on existing IT staff.
3. Improved scalability: IT infrastructure will be better able to scale to meet the demands of new applications at the edge and to improve overall reach to partners, service providers, and systems.

Where managing infrastructure used to be difficult, time-consuming, and expensive, a software-centric model virtually eliminates all of these challenges and delivers an experience that can be managed from a single app.

Looking further into the future, it's possible we might see human intervention removed entirely.

The Next Implication of Web 3.0: **The Ubiquitous Edge**

The combination of software-centric assets and architecture is already changing the way IT leaders deploy and manage their infrastructure. Yet, it will change even more as business and technology strategies move to the edge.

Over the last few years, the edge has become the most talked-about aspect of IT infrastructure, largely driven by the arrival of the first 5G networks. As a result, the vision of the edge that has predominantly taken hold is one of small, modular data centers—sitting at the base of cell towers to speed the delivery of content and application data by being "one hop away" from end user devices (i.e., cell phones). This is a misleading over-simplification.

The reality is more complex, akin to "The Blind Men and the Elephant" fable. It is now clear that the edge means something very different, depending on one's perspective.

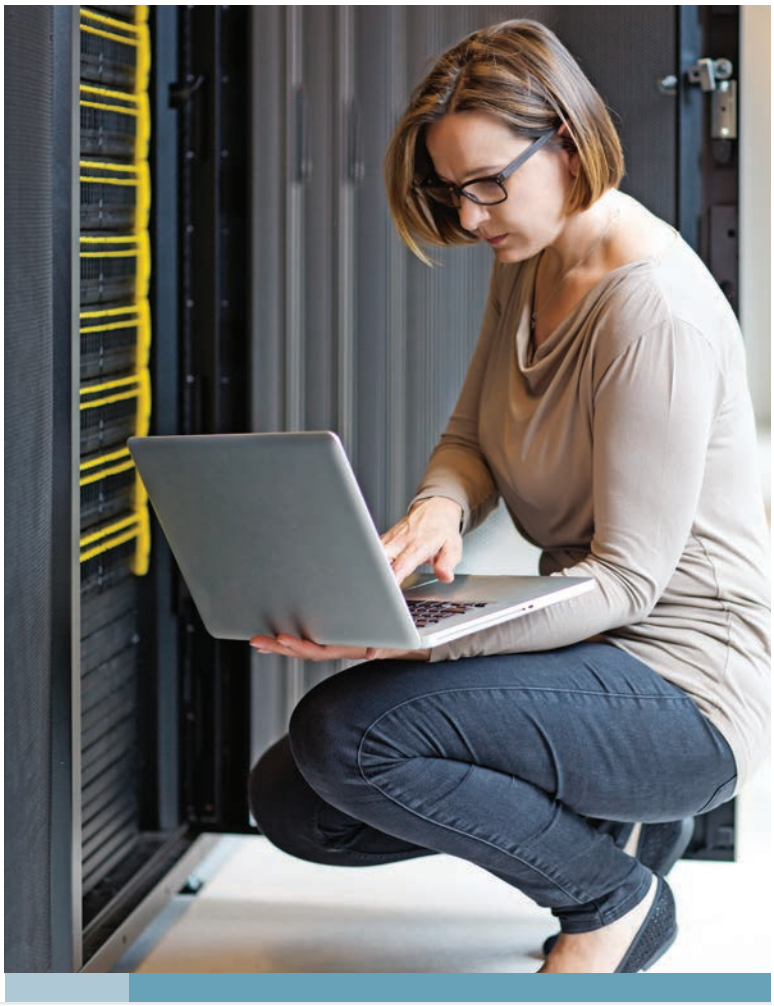

Over the last few years, the edge has become the most talked-about aspect of IT infrastructure, largely driven by the arrival of the first 5G networks.

The Edge: Driven by **Geographic, Application,** and **Latency Requirements**

In one sense, the edge is geography-specific. Today, much of the internet's content is still centralized in data centers located in Tier 1 and 2 metros. Emerging IoT applications and sensors may someday create the need for an edge that can support compute and storage in remote locations, but it's just as likely the edge will continue to be built out in large cities as well.

This is the edge that is actually getting the most notice today, as hyperscalers like AWS establish their first forays into the edge in large metros. This makes perfect sense and follows the simple logic of the famous bank robber Willie Sutton, who when asked why he robbed banks, replied, *"Because that's where the money is."* Large metros are still where the majority of users and data are located.

In another sense, the edge is application-specific, and the configuration of its infrastructure varies accordingly. The edge for a hyperscaler building a new availability zone in a Tier 2 metro requires several thousand square feet of space and several megawatts of power, but this doesn't warrant building a dedicated facility.

Instead, an existing, multi-tenant data center (MTDC) provides the ideal solution. For a network provider looking to deploy new 5G towers—or extend an undersea cable landing further inland—the edge may require far less space and power. It also might be in a unique location where no data center is available. In that case, a modular, micro-data center—with half-cabinets and kWs of power is the perfect option— such as the EdgePresence platform in which DataBank has invested.

There is also an element of the edge that is performance/latency-specific. For instance, video and e-commerce applications may continue to be adequately served by data centers, cloud platforms, and CDNs in the Tier 1 and 2 metros where they exist today. Those applications perform fine even with the 15-100ms of latency that occurs between a user and a server hundreds of miles away.

However, augmented reality or gaming applications that require sub 15ms response times may need to reside far closer to end users. Here again, that could be a modular solution in one geography or a large MTDC in another.

The edge is application-specific, and the configuration of its infrastructure varies accordingly.

A Ubiquitous, Multi-Modal Edge Across **Three Platforms**

This is why, at DataBank, we think of the edge as ubiquitous—everywhere and multi-modal—existing in different forms, depending on the business and application it's supporting.

To support this strategy, we developed three specific edge platforms:

- Our **Hyperscale Edge** provides a geographic footprint ideal for webscale cloud, content, and compute resource providers seeking to build or expand their availability zones in Tier 1 and 2 metros. This will establish their edge for the next 10-15 years.

  DataBank's 60+ data centers in the U.S. are spread across 29 Tier 1 and 2 metros, putting our entire U.S. portfolio within 50 miles of half the U.S. population, and providing best-in-class reach and scalability. DataBank typically deploys multiple facilities within a metro, and many of them are integrated with Zayo Networks fiber. This means hyperscale customers can build multi-node availability zones in a metro with uncongested access to long-haul networks that connect to other availability zones.

What Kind of Partner Do You Need?

Bringing together three key elements

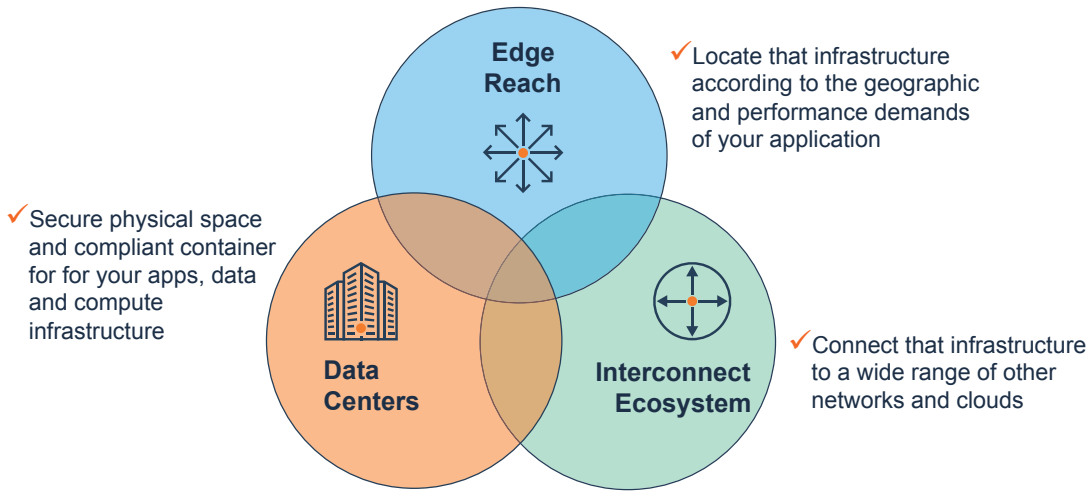

- Our **Interconnect Edge** gives network providers, SaaS developers, and enterprises seeking a hybrid IT strategy the open and neutral-access interconnections their applications demand. DataBank maintains 20 major interconnect locations in 13 markets that support over 30,000 cross-connects. These include some of the most notable carrier hotels such as 1 Wilshire in Los Angles, 111 8th Avenue, and 60 Hudson in New York City, 1950 North Stemmons in Dallas, and 1500 Champa in Denver.

  These interconnect hubs and our multi-node cloud on-ramps provide unfettered access to dozens of cloud platforms, hundreds of carriers, and thousands of miles of fiber. Also, with new advancements in our portal, customers have software-defined control of this ecosystem, allowing them to tap into any cloud platform provider, carrier, or enterprise application partner quickly.

- Our **Low-Latency Edge** meets the demands of customers who need to support extremely latency-sensitive or geography-specific applications. DataBank utilizes the EdgePresence modular, micro-data center platform, which provides fully secure, SOC-2 certified containers with redundant 48kW power sources, 300+ square feet of raised floor capacity, 10 racks (or 20 half racks), and four 5-ton cooling units.

  These EdgePods can be deployed virtually anywhere in a matter of days or weeks, rather than months, and they tether to existing DataBank facilities. This allows applications and interconnections to be established beneath 5G cell towers, at cable head-ends, or anywhere an application needs to be put, and as close to user populations as necessary for minimal latency.

EdgePods (See diagram) can meet the demands of the Low-Latency Edge. Three Edge Platform options can help IT leaders develop an effective digital asset allocation strategy to balance IT requirements and position their organization for maximum success.

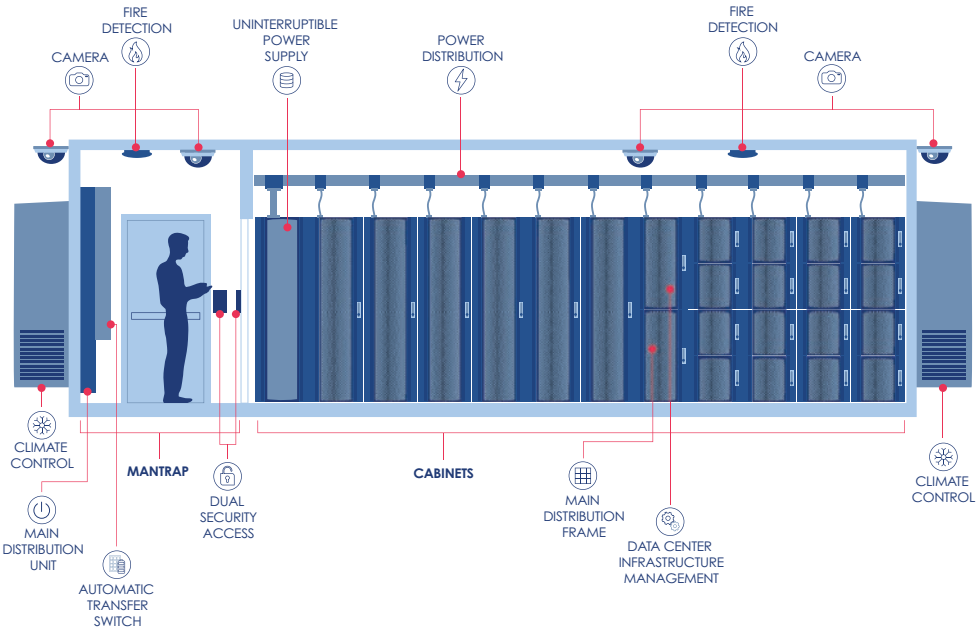

DataBank Corners **the Edge**

With these three edge platforms, DataBank is poised to help any company meet the edge demands they have today and tomorrow, regardless of application or location. We don't stop there, though. Our Data Center Evolved™ model provides all the options you need to make sure your infrastructure at the edge is managed, secure, and compliant. In addition, our Customer Portal gives you all the visibility and control you want to make sure your applications are on the right infrastructure, at the right time, in the right place, and for the right audiences.

DataBank is poised to help any company meet the edge demands they have today and tomorrow, regardless of application or location.

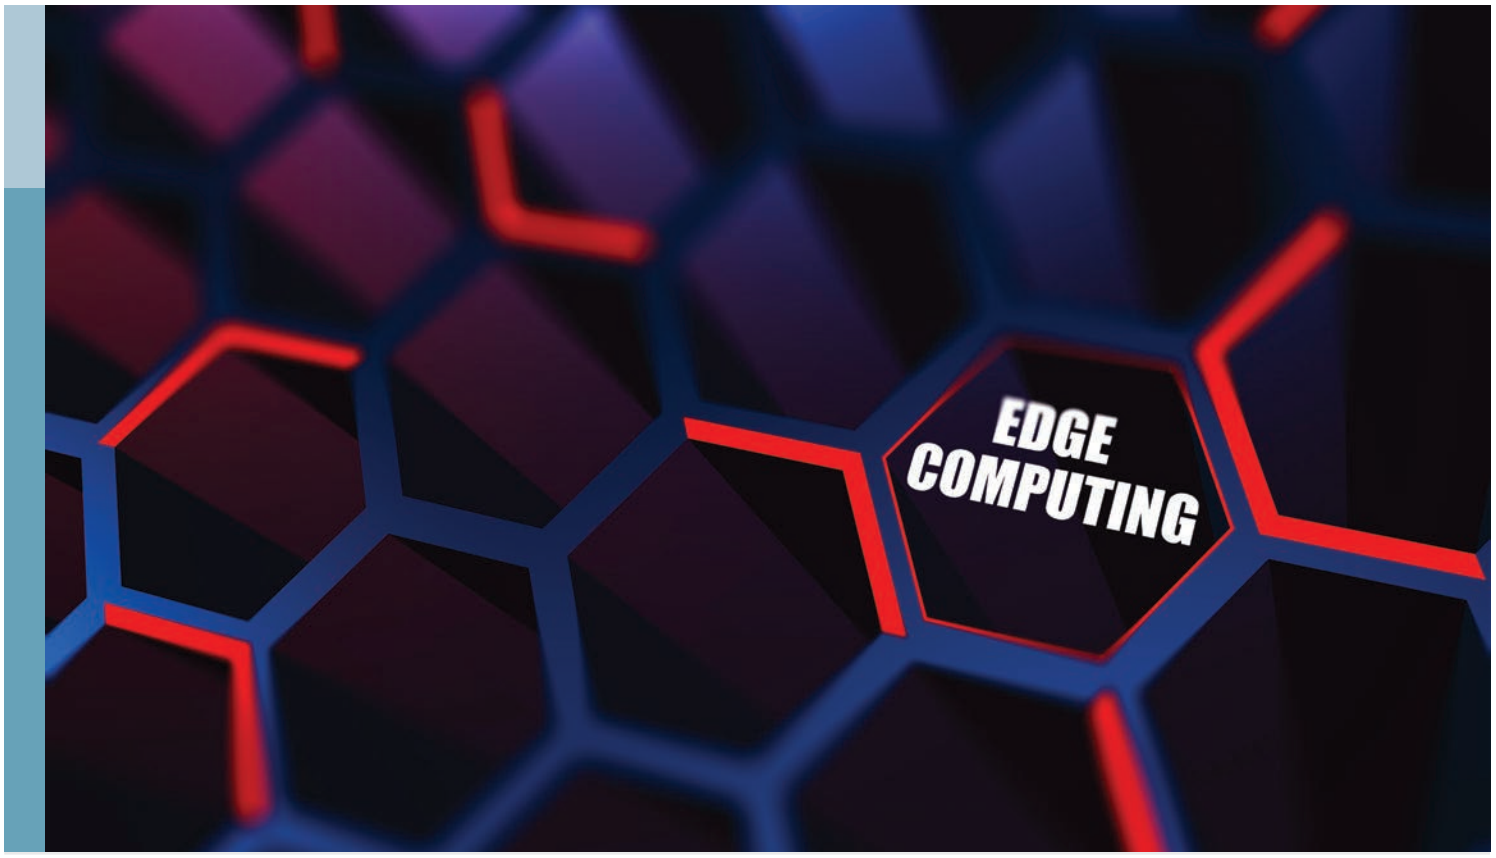

About **DataBank**

DataBank enables the world's largest enterprises, technology, and content providers to consistently deploy and manage their infrastructure, applications, and data on the right platform, at the right time, in the right place. Our colocation and edge infrastructure footprint consists of 60+ data centers and 20 interconnection hubs in 29 markets, on-ramps to an ecosystem of cloud providers, and a modular edge data center platform with virtually unlimited reach. We combine these platforms with contract portability, managed security, compliance enablement, hands-on support, and a guarantee of 100% uptime availability, to give our customers absolute confidence in their IT infrastructure and the power to create a limitless digital future for their business. DataBank is headquartered in the historic former Federal Reserve Bank Building, in downtown Dallas, TX. For additional information on DataBank locations and services, please visit **www.databank.com.**