The Lumen Edge Advantage

Meet the low-latency platform
powering next-gen apps
and amazing experiences.

LUMEN®
TECHNOLOGIES

Next-gen technology is powered at the edge

Businesses must successfully deploy workloads closer to where digital interactions occur—allowing them to improve user experiences, increase application performance, lower the cost of infrastructure and help reduce security and compliance risks.
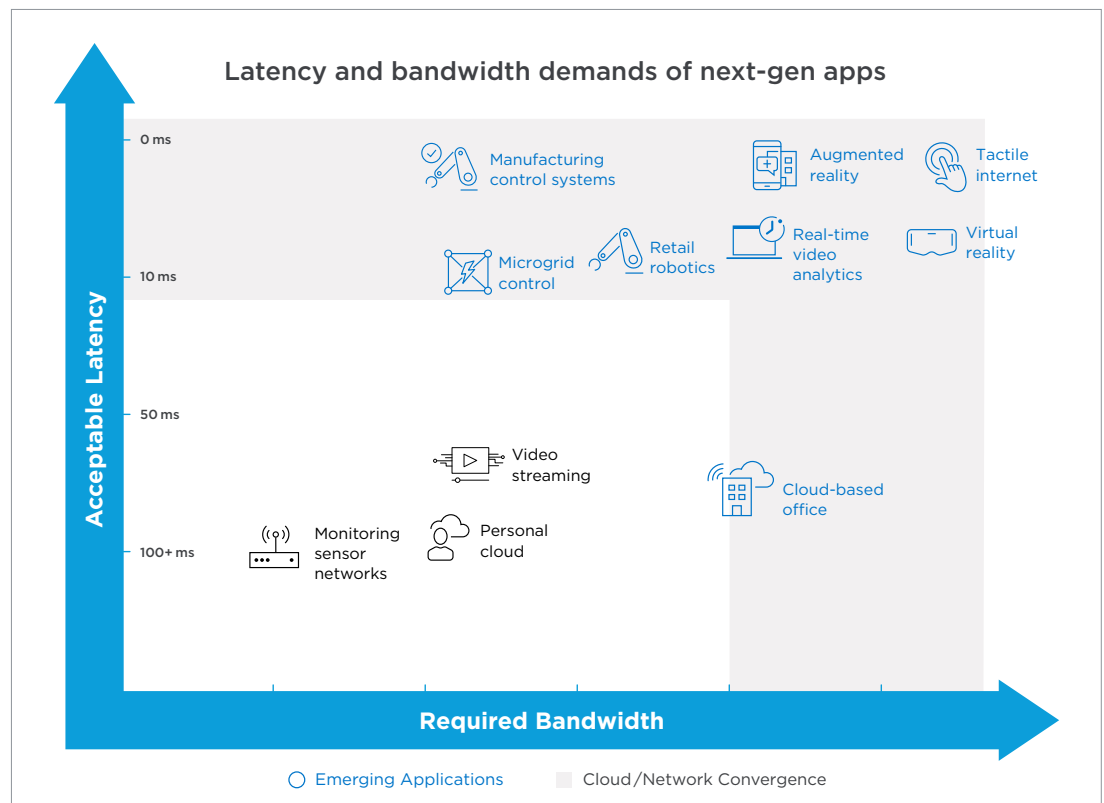
As we enter an era of autonomous vehicles, smart manufacturing, telehealth services and digital in-store kiosks, businesses in critical industries such as manufacturing, healthcare and retail compete to deliver amazing experiences—using advanced technologies such as real-time analytics, AI and augmented reality/virtual reality. Because these next-gen applications can directly improve business outcomes like improved user experience and lower costs, ensuring their optimal performance across a distributed infrastructure is vital.

LUMEN®
TECHNOLOGIES

Interaction with people and devices, dynamically, at scale

For businesses to compete on how well they can acquire, analyze and act on data, IT strategies and execution are key. Many businesses turn to central public cloud services, but a centralized cloud's long-distance latency may be a major drawback for next-gen applications. Alternatively, on-premises resources offer ultra-low latency, but the total cost of ownership (TCO) is much higher when you own the infrastructure and hire your own staff to operate it.

F1
As the degree of dynamic interaction with users or devices increases, acceptable latency decreases.

Latency and bandwidth demands of next-gen apps

- 0 ms
- 10 ms
- 50 ms
- 100+ ms

Acceptable Latency

- Manufacturing control systems
- Augmented reality
- Tactile internet
- Microgrid control
- Retail robotics
- Real-time video analytics
- Virtual reality
- Video streaming
- Cloud-based office
- Monitoring sensor networks
- Personal cloud

Required Bandwidth

○ Emerging Applications ▢ Cloud/Network Convergence

Driven mostly by TCO considerations (and to offload the IT burden), many businesses choose to migrate their applications to centralized public clouds. When deploying distributed applications, these businesses may experience disappointing performance and expected cost savings may be unrealized—especially for dynamic AI-enabled workloads. The application overhead of many "round trips" and the bandwidth consumption costs for connections between distributed locations and centralized cloud compute and storage can be very high and presents an architectural complexity that most organizations don't have the skills for.

Although the cloud offers elastic scale and the advantages of a service/subscription model, relying exclusively on public cloud services is neither efficient nor effective for applications such as real-time health monitoring, autonomous vehicle piloting and smart manufacturing. And regardless of their application strategies, businesses must comply with statutory requirements to regionally protect and govern data—a challenge in the centralized cloud.

LUMEN®
TECHNOLOGIES

Evolution from centralized to distributed computing services

The computer industry's history is one of moving from centralized to distributed computing and then uniting them in a hybrid model.

In the 1950s and 1960s, businesses and governments used centralized mainframes to process all their digital data. Users had to wait for their data processing department to work through the backlog before their job could be run. Soon, timesharing opened up access to centralized mainframes, and by the late 1970s, minicomputers had sprung up to directly serve the needs of engineering, manufacturing, finance, human resources, sales and marketing departments.
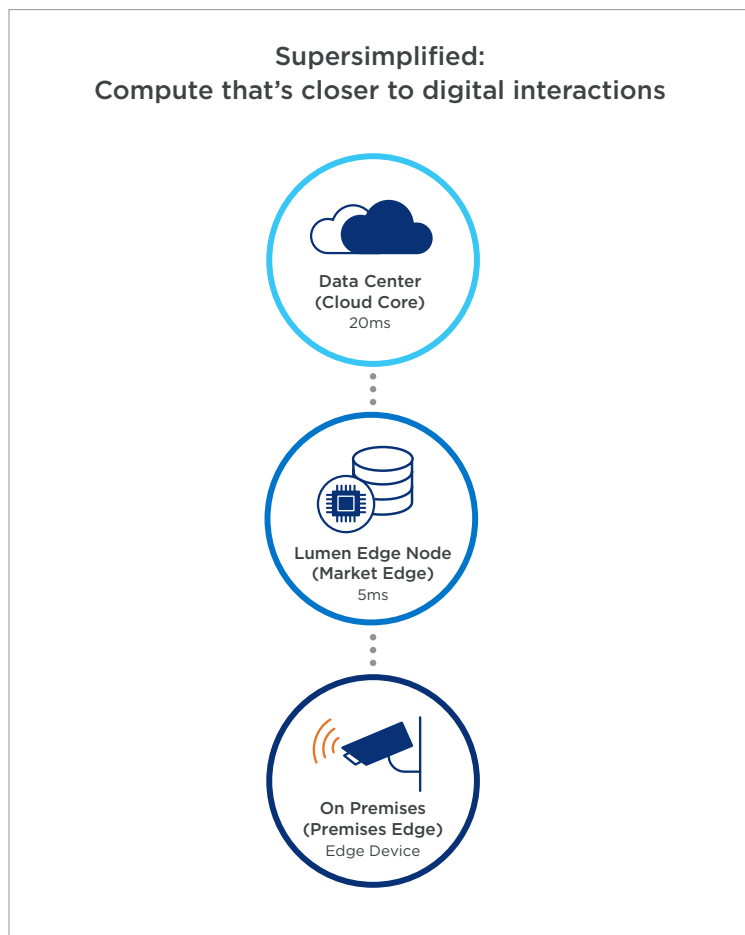
In the 1980s, personal computers empowered business users with interactive applications such as spreadsheets for "what if" analysis and word processing to edit documents on the fly. LANs emerged to connect workgroups, and in the 1990s, the World Wide Web enabled businesses and consumers to connect applications, infrastructure and operations in entirely new ways. Client/server applications such as enterprise resource planning (ERP), supply chain management (SCM) and customer relationship management (CRM) enabled corporations to redesign and globalize. Personal productivity morphed into workgroup collaboration and process automation, and more and more human activity migrated to the web.

As digital businesses such as Amazon, Microsoft and Google grew, their growing global data center footprints represented added value to them as fungible "cloud" assets that they could offer to subscribers as Infrastructure and Platform as a Service. Around the same time, Salesforce.com pioneered Software as a Service with their market-leading CRM offering, and later, the SaaS concept grew exponentially with services such as Microsoft 365.

LUMEN®
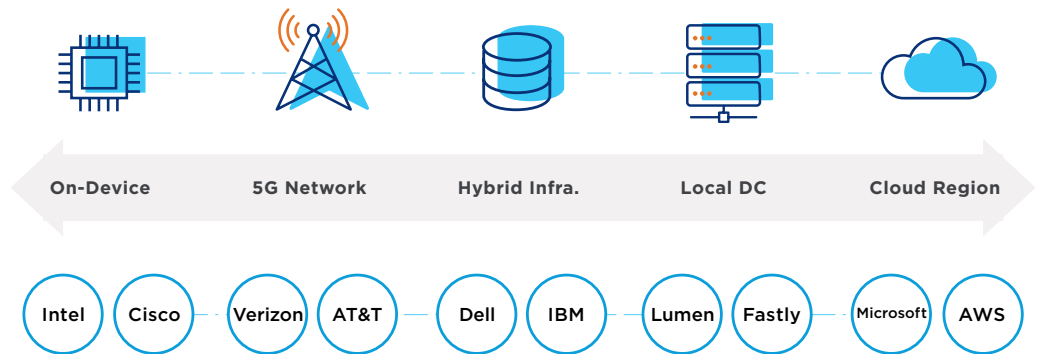TECHNOLOGIES

Edge computing completes the cloud

As necessary as the public cloud may be, centralized cloud services are not sufficient to meet the needs of all next-gen distributed applications. To dynamically interact with people and devices in near real time and at scale, business initiatives need an agile, cost-effective architecture—one that meets the rapidly growing needs of IoT workloads and data-intensive, latency-sensitive end-user experiences. That is why Lumen offers customers an open standards-based, intelligent platform built on adaptive infrastructure. It can orchestrate workloads and compute resources closer to the point of digital interaction to help reduce latency and improving application performance.

Supersimplified:
Compute that's closer to digital interactions

Data Center
(Cloud Core)
20ms

Lumen Edge Node
(Market Edge)
5ms

On Premises
(Premises Edge)
Edge Device

Computing closer to points of digital interaction is often called "edge computing." Although the term means different things to different people, we can agree it happens close to where things and people produce or consume information—that is, edge computing does not rely on a central location that can be thousands of miles away. This proximity enables applications to process and store data faster and more efficiently.

LUMEN®
TECHNOLOGIES

There is no clear consensus of where "the edge" is — ITDMs are hearing conflicting things from different tech companies

Products described as "The Edge" include...

| On-Device | 5G Network | Hybrid Infra. | Local DC | Cloud Region |

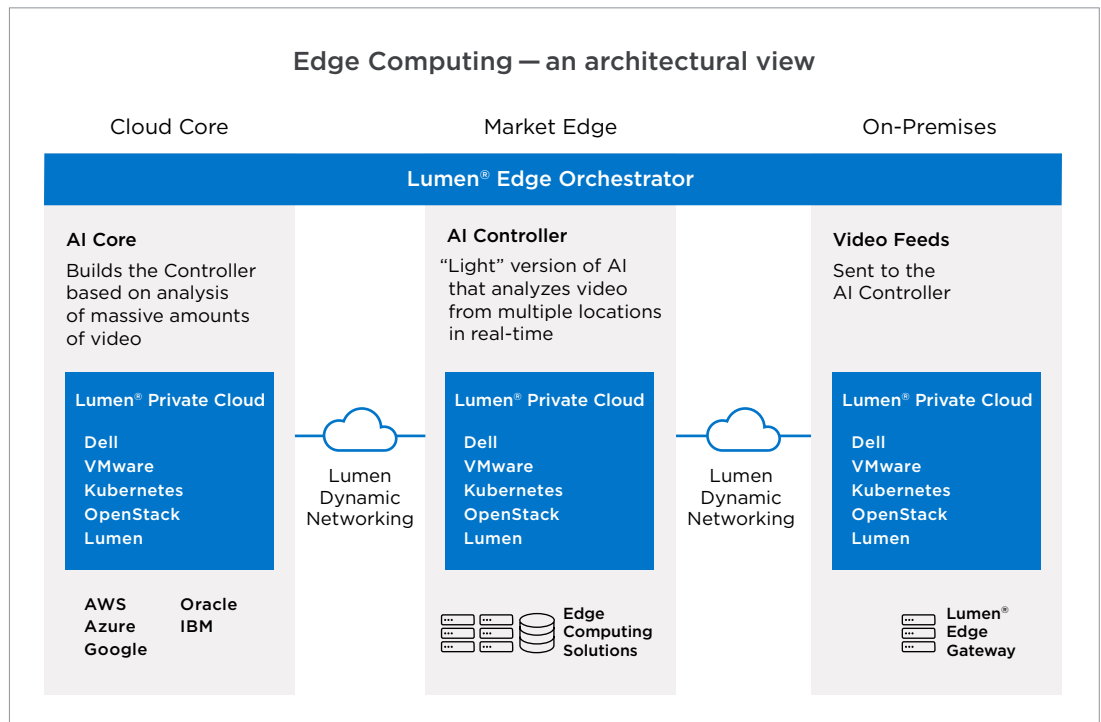| Intel | Cisco | Verizon | AT&T | Dell | IBM | Lumen | Fastly | Microsoft | AWS |

Edge computing can allude to on-premises computing where the WAN meets the LAN. But most companies would prefer to get out of the "IT business," and if they can get that same level of performance/latency in an edge cloud, that would be the best of both worlds. Edge computing systems in small data centers, closer to user and device interactions and data, open up new market opportunities for businesses across many industries.

Businesses increasingly require a range of public and private cloud environments to support their distributed applications. Combined with the intensive latency requirements of next-generation applications, edge execution venues have become a critical component of hybrid cloud strategies and the adoption of emerging technologies.

LUMEN®
TECHNOLOGIES

The power of the unified edge ecosystem

Lumen helps to improve business outcomes for customers by delivering an advanced application platform that is bolstered by one of the most interconnected global networks,[1] a broad data center footprint, connected security, and hybrid cloud services coupled with an ecosystem of technology partnerships. As a component of that platform, Lumen maintains strategic edge locations and offers Lumen® Edge Computing Solutions.

Located today in Lumen Edge facilities across many markets in the U.S., with additional edge locations planned around the world, Lumen Edge Computing offers low latency comparable to on-premises computing. In addition, we offer many of the TCO and governance benefits of the public cloud as part of a hybrid IT strategy. Lumen Edge Computing Solutions support emerging applications using AI and machine learning, robotics, real-time video analytics, manufacturing control systems, micro-grid controls and augmented reality/virtual reality among other next-gen technologies.

Edge Computing — an architectural view

Cloud Core	Market Edge	On-Premises
Lumen® Edge Orchestrator		
AI Core Builds the Controller based on analysis of massive amounts of video	**AI Controller** "Light" version of AI that analyzes video from multiple locations in real-time	**Video Feeds** Sent to the AI Controller
Lumen® Private Cloud Dell VMware Kubernetes OpenStack Lumen	**Lumen® Private Cloud** Dell VMware Kubernetes OpenStack Lumen	**Lumen® Private Cloud** Dell VMware Kubernetes OpenStack Lumen
AWS Oracle Azure IBM Google	Edge Computing Solutions	Lumen® Edge Gateway

Lumen Dynamic Networking

Lumen is ideally positioned to drive edge use cases with the ability to deliver compute on customer premises with low latency, within deep metro points of presence and central offices designed for within 5ms of latency.

LUMEN®
TECHNOLOGIES

The combination of our network strength, data center footprint, security and managed services experience bridges clouds, IT infrastructure and the edge to help workloads and applications to perform at their best, due to:

- **Exceptional assets, services and industry partnerships** around data center and cloud—customers have peace of mind knowing the Lumen Platform, along with our managed services and our distributed application and infrastructure expertise, are at their service.

- **A low-latency, dynamic network** comprising approximately 450,000 global route miles of fiber and 170,000 on-net buildings—our on-network facilities help move data fast, with low latency and reliability, even if a fault occurs.

- **Thousands of secure technical facilities and global data centers** with ITIL-based expertise in high-performance compute and multi-cloud management—customers can rest assured their applications and data are operated and managed using industry best practices.

- **Private connections to leading cloud service providers** offers customers the most direct routing of applications with the minimum number of "hops." Lumen connects customers to the public data centers of Amazon Web Services, Microsoft Azure ExpressRoute & Azure Government, Google Cloud, IBM Cloud, Oracle Cloud (along with direct peering connectivity to more than 35,000 ASNs), as well as more than 350 Lumen data centers and more than 2,200 third-party data centers via fiber in North America, Europe, the Middle East, Latin America and Asia Pacific.

- **Extensive experience in operating and integrating client requirements** in a hybrid world of cloud, network and security means Lumen customers benefit from the real-world know-how of experts facing the challenges of distributed application operations every day.

The Lumen Platform and its capabilities are grounded in extensive global infrastructure. Lumen operates one of the largest, most connected fiber networks in the world, offering dynamic connections to more than 2,200 public and private data centers, with low-latency performance via Lumen Edge Compute nodes, all delivering applications and data when and where they're needed.

With this extensive infrastructure as a foundation, the Lumen Platform combines core IT elements into a unified application delivery solution for businesses, governments and communities.

---

[1] Platform claim based on Lumen's unique combination of assets and services with the lowest network latency on the majority of the MSA routes, threat detection, data scrubbing center capacity and end-to-end protection.

[2] https://www.lumen.com/en-us/solutions/edge-computing.html

Learn more
Connect, protect and respond at the speed of business with the **Lumen Edge advantage**.[2]

877-453-8353 | lumen.com | info@lumen.com

LUMEN®
TECHNOLOGIES