WHITE PAPER

Application delivery at the edge

Use cases, state of the market and expert insights.

LUMEN

What is edge and why should businesses care?

As the Fourth Industrial Revolution continues to advance, business objectives will increasingly rely on the ability to collect, analyse, and act upon data in real time. To meet these objectives, it's imperative for workloads and computing power to be placed as close as possible to where users and endpoints are located. This in turn has led to greater use of distributed architectures, including decentralised on-premise datacentres, public and private cloud, and containerised workloads.

Based on a 2020 IDC survey of over 800 C-level executives across industries,[1] **73%** view investment in edge solutions as a key strategic initiative, with two-thirds already in the process of adopting edge services into their digital workflow.

Yet where the edge actually lies and how it gets defined can vary greatly from business to business depending on where the deployment is located, what capabilities it has, or how it's connected to the rest of the architecture.

Most consumer-facing sites and applications deal with the cloud edge and/or the compute edge, which often work in a complementary fashion. The cloud edge generally refers to large-scale cloud and hosting providers. While this type of shared infrastructure usage offers large amounts of resources and high availability, its centralised nature means that latency can be a challenge, and network traffic to and from the cloud can lead to high transit costs.

The compute edge (sometimes also referred to as the fog or datacentre edge), on the other hand, utilises small datacentre locations that can be deployed on premise closer to the actual end users, thus reducing latency and leading to better overall performance. While these locations can't fully replace a cloud edge due to more limitations on capacity and resources, many services can be migrated here from the cloud to minimise latency and improve network bandwidth.

Similar to the datacentre edge are the branch edge and enterprise edge. Branch edges can be established on premise at different locations to perform specific functions that are more efficient to run on a modular basis. Examples include point of sale systems for retail companies, medical devices or health records, or local bank branches that operate independently of each other. In cases such as these, the branch edge enables latency-sensitive applications to run smoothly and makes it easier to maintain content freshness for local offers that are not shared by other branches. Meanwhile, enterprise edges are just another way of sharing resources between branch locations by grouping them together on a shared enterprise network, thus simplifying management and making it easier to scale.

1. IDC Survey

As we move toward IoT devices, we get into the device edge and the sensor edge (also sometimes referred to as mist), which must work in tandem with each other. The device edge is a nano datacentre deployed in a location like a factory floor, power plant or self-driving vehicle. The sensors located on the devices that collect data are connected to the device edge to store and process minimal amounts of data as close to real time as possible.

The nature of edge investments can vary depending on a company's sector, business model, geographical footprint and digital maturity. Those that are still setting up the foundations of their edge delivery strategies can already leverage shorter-term gains that improve the digital employee or customer experience such as optimising operational efficiency or improving application performance. After these initial edge deployments mature, the focus can shift to more long-term benefits like reducing the cost and/or complexity of their infrastructure, or opening up new revenue opportunities.

Core components of edge computing

A well-architected edge solution will address several of the key challenges that businesses face, among them security, storage, and real-time application performance. All three rate among the top priorities faced by C-level executives in the IDC survey, with security being the most prevalent; over 30% of respondents rated it as the number one challenge faced at the edge.

The truth of the matter is that the number of functions that can be served by edge computing is as vast and varied as the different types of companies that use it. Edge computing can be used for more accurate invoicing processes by connecting metered units with billing systems. It can streamline inventory management by connecting different locations with a centralised processing hub. It can be used for preemptive maintenance in manufacturing locations by detecting anomalies in the production line. It can connect autonomous vehicles in a fleet to monitor road conditions and traffic congestion.

For the purposes of this report, however, we'll be focusing primarily on how edge computing can be applied to the delivery of consumer-facing HTTP applications that are accessed by end users over the internet. Edge technology solutions are a key enabler of amazing end-user experiences. In crowded markets like retail, gaming, digital media, manufacturing and IoT, the quality of their digital experience can serve as a differentiator for customers, which explains the rapid growth in edge solutions in recent years.

When evaluating the different types of edge solutions, rapid deployment is the most common benefit that companies look for, with nearly half **(48%)** citing it as a key aspect of their decision making, just ahead of other benefits like capabilities, quality of support, cost, and computing capacity:

Question*

What do you look for when evaluating an edge solution?

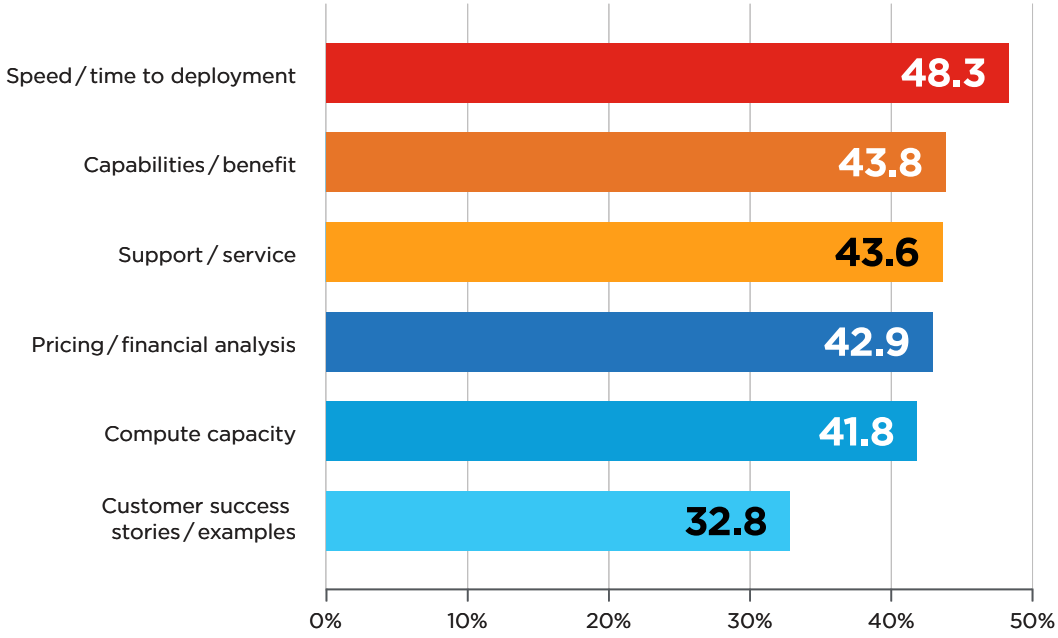

Perhaps equally important is the capability of an edge solution to minimise the amount of latency in order to maximise the performance for the end user. By deploying resources as close as possible to where data is collected and consumed, businesses can process and analyse It in real time and create a positive customer experience.

* Source: IDC Survey

LUMEN

Application delivery

Application delivery is the broadest and most complex of the core components of edge computing. It's also the one with the most direct impact on the customer experience, and therefore offers one of the best opportunities for companies to meet their business objectives.

At its heart, edge application delivery is the natural evolution of traditional content delivery. By combining the bandwidth, capacity, and throughput of the underlying network while enhancing the user experience through software-defined delivery solutions that get deployed at the edge, companies can offer the differentiated customer experiences that help them thrive in hyper-competitive markets.

These solutions can span a wide variety of functions; static or dynamic content delivery, website acceleration, data processing, image optimisation, and API caching are just a few of the services that fall under the application delivery umbrella. Regardless of the use case, they all have the same underlying goal: to improve end-user experiences and give businesses more flexibility and control over their application delivery.

The core capabilities of edge computing for application delivery improve the end-user experience by providing enterprises with more efficient processes around security, performance, scale, and engineering flexibility.

Modern DevOps and SRE organisations don't want black box solutions that leave them powerless to optimise the delivery of their applications. Modern workflows must not only be able to adapt delivery to meet the varying expectations of end users around the globe, but also offer programmability to DevOps teams that are always looking to enhance the performance and reliability of their applications.

Edge delivery providers can offer this level of programmability through bare metal services and by giving DevOps teams the ability to build directly on the underlying edge infrastructure with proprietary applications, custom code, or open-source modules.

Security

One of the unfortunate realities that comes with the rise of distributed architectures and more sophisticated applications is an increase in potential security threats. Not only are there more entry points for malicious attackers to access systems and user data, but every API and line of code within an application is a potential vulnerability that must be secured. And to make matters worse, attackers are getting more creative and sophisticated, now often deploying multi-vector attacks that combine tactics, e.g. overwhelming servers through credential stuffing plus DDoS.

However, it's not all bad news. These threats can be mitigated with a security-first edge solution in which specialised services can be rapidly deployed at the application layer to protect data and stop attacks at the origin point. And given the sensitive and valuable nature of financial data, a customer-facing site or application must be secured against the multi-faceted attacks such as DDoS, phishing, SQL injection, etc.

Network and application infrastructure providers must be able to secure content at the origin, but an advanced solution will also offer modular deployment of various edge security solutions to allow businesses to customise their security strategy around their specific use case. These can take the form of web application firewalls (WAF), bot managers, API protection, or Layer 7 DDoS mitigation. Regardless of what security solutions are deployed, it's important for them to be compatible with the underlying network architecture so that they can be implemented without a complicated installment and configuration process.

Storage

Given the emphasis placed by modern enterprises on minimising latency, the need for edge storage solutions becomes clear. Businesses must be able to access and process data close to where end users are located – and often in real time – to ensure an excellent customer experience.

According to the IDC survey, the vast majority of edge use cases **(83%)** require either real-time or frequent access to the data that gets stored at the edge. And as advances in computing power, artificial intelligence, IoT, and machine learning practices continue, the need for real-time access will only grow.

The other aspect of edge storage that companies must evaluate when shopping for solutions is the length of data retention. The need for longer data storage can vary quite a bit across industries and specific service offerings, but the ideal solution will allow flexible retention policies at competitive rates so that businesses can alter it based on their changing needs.

However, it's important to remember that storing data at the edge doesn't make much sense unless it needs to be called upon or processed in some way. Edge storage should be reserved for data that must be refreshed on a consistent basis, with the ability to aggregate, anonymise, filter, or send back to a more centralised location for long-term archiving. Given that, the compatibility with a cloud solution is an important aspect to consider here.

From content delivery to application delivery

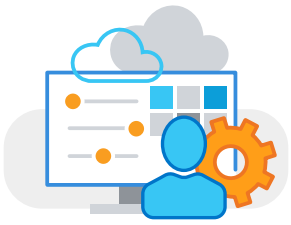

For much of the past two decades, content delivery technologies were generally viewed as being extremely relevant to media and ecommerce given the large amounts of data they had to deliver and store, but having limited use beyond that. Yet as applications have evolved to serve more dynamic content and rely on more powerful distributed computing, a modern CDN provider must be able to do far more than just deliver large amounts of data; they now must be able to transform that content to create a better, richer experience for end users.

By working in concert with hybrid cloud technologies, software-defined delivery, and edge computing capabilities, a CDN must now be part of a broader application delivery solution. And far from being the exclusive domain of media companies, application delivery is now critical to the success of businesses across industries.

Digital media

For years the media and entertainment industry has talked about delivering a viewer experience with streaming video that's 'as good as' broadcast video. More recently, however, the conversation has taken a new turn, according to Jason Thibeault, CEO of the Streaming Video Alliance. **"Now, when conditions are ideal, streaming video will always be better than broadcast."**

So how do you create those ideal conditions? High-bandwidth, low-latency CDNs are among the givens, but the industry is also looking at the difference that edge computing and edge application delivery can make. **"But the definition of the edge is changing,"** says Thibeault. **"It's moving away from the network proper to deeper within the last mile. Eventually, it will be the end device itself."**

Educating the industry about what edge is and what it means for streaming video is a task the Alliance has willingly taken on. Elevating streaming and optimising the viewer experience are among its core ambitions for the entire media and entertainment industry.

"Technology developed by Alliance members is made freely available across the industry, to help ensure high-quality delivery at scale," says Thibeault.

Although many companies are still at the stage of shifting their technology, components and workflows from on-prem hardware to the cloud, some are already much further ahead. Primarily pure-play streaming providers — driven by DevSecOps teams rather than traditional broadcast engineers — are already making advanced use of the cloud and are now investigating what the edge can deliver.

"Companies already exploring edge see the scalability it enables as a key benefit — much better than virtualisation can offer," says Thibeault. **"They're looking to get rid of hardware and instances and embed component functionality right into the fabric of the cloud to deliver serverless functionality that's scalable by nature, just like the cloud itself. So capacity issues go away, and resilience and redundancy increase."**

Looking at edge applications and capabilities that have the potential to take the streaming video experience beyond broadcast, Thibeault calls out:

Automation

Automation as an edge application could help to maintain a high-quality viewer experience by solving issues faster. AI and ML in the fog/compute edge could act on data received from the viewer's device, instead of having to transfer that data to the cloud or another central point to initiate issue resolution. So in the case of an ad failure, another ad could be served immediately; or if latency was impacting on content delivery, the stream could be switched seamlessly to an alternative CDN.

XR and immersive video

XR and immersive video could have real applicability for live sporting and other events, providing the ability to layer in features like data points or additional camera angles. **"As streaming video is based wholly on data, this will be another way to take the viewer experience beyond anything that broadcast can offer,"** says Thibeault.

DRM

One Alliance member company is already experimenting with using a distributed ledger to manage artist payments. Could a similar application at the edge replace the current process of applying DRM at the point of origin of the content?

"Perhaps we can replace the traditional DRM security wrapper with a more flexible — but equally secure — application right out at the edge on the viewer's device," suggests Thibeault. Such an application would allow viewers to easily share or move content across their devices while maintaining protection for the rights-holder.

Monitoring

Low latency is as critical for the operational experience as for the viewer experience. Getting closer to real-time monitoring will depend on moving away from on-prem hardware and virtualised instances, and pushing monitoring applications out to the edge.

"I believe that this operational demand will be as strong a driving force for moving to the edge as the desire to increase the scalability and efficiency of product features," says Thibeault.

Observability

With so much data coming from so many different components within the streaming video technology stack, finding patterns, identifying revenue opportunities, and addressing issues is much more complex than just monitoring performance data. Serverless intelligence within the edge could help post-process data, such as automatically relating multiple sources into a single data lake, and make it available to visualisation tools quicker to speed up analysis and business decisions.

Device edge

Awareness of the role of the device edge in optimising the viewer experience is growing, as companies look at deploying caching deep into the edge. This could allow a provider to pre-position content that the viewer may want to watch (based on a recommendation, for example) on or near the device, so that they don't have to backhaul it. **"That would deliver a truly great 4K viewing experience,"** says Thibeault.

At the same time, peer-to-peer networking is becoming a legitimate delivery technology, especially for live event streaming, as it improves scalability and the viewer experience with no need for more servers.

5G

Edge applications deep into the last mile only become viable when there is a low-latency network to support them. 5G will enable this. When millimetre wave technology becomes widely prevalent, streaming providers will not only be able to deliver higher-quality content, but also deliver immersive and interactive experiences which depend upon application responses times that are near real time.

> **The most innovative companies are pushing the envelope, asking: how can we make streaming more scalable, resilient, consistent and redundant, without having to stand up more boxes? Forget that, it just doesn't work anymore. Especially when you have a global audience. You've got to do something else. The edge is the answer."**
>
> — Jason Thibeault
> CEO
> Streaming Video Alliance

Definitions: from the cloud to the edge

- **Cloud** — a centralised, multi- or single tenant platform with a physical network edge.
- **Compute Edge / Datacentre Edge / Fog** — a decentralised computing infrastructure between the cloud and the devices that produce data. Fog solves latency, performance and scalability challenges in the cloud infrastructure using edge devices to perform compute, storage and networking functions locally and routed over the internet.
- **Multi-access edge computing (MEC)** — a single layer of nodes deployed at the edge to bring cloud infrastructure closer to users to improve latency and performance.
- **Device Edge / Mist** — the extreme network edge, at the consumer or viewer device, for handling requests locally.

Gaming

Video gaming is big business today, having become one of the most popular ways for people all over the world to spend their free time. But with so much competition in the market, developers and publishers must work hard to capture gamers' imagination.

"To be successful, a game obviously has to be fun to play, but it has to be properly marketed. Marketing is especially important for multiplayer games — you have to ensure there are enough people online who want to play, otherwise the game will go nowhere," says Lerika Mallayeva, Founder and CEO of DevGAMM, the biggest conference for game industry professionals in Eastern Europe.

Like the media and entertainment industry, gaming and gambling companies also face the challenge of rising consumer demands for performance and reliability while also managing the growth of mobile usage. After seeing year-over-year growth of **25.6%** in 2020, mobile gaming accounted for nearly half of the total global gaming revenue.[2]

In the face of this rising demand, the gaming industry is outgrowing the infrastructure on which it relies. With physical games giving way to online downloads, popular new releases and multi-player gameplay can put tremendous strain on the networks that deliver these applications, increasing the need for edge application delivery solutions. Compounding the problem is that end-user experience issues like lag, jitter, and latency can completely ruin the in-game experience and tarnish a provider's brand reputation within the gaming community.

2. Newzoo

Real time realism

Naturally, gamers have high expectations for the gaming experience, and developers are ready to meet them with exciting gameplay, immersive 3D landscapes, and sensory-rich feedback — all of which depend on rapid processing and delivery of vast amounts of data and content. And for competitive games — especially in the multiplayer online battle arena (MOBA) genre (that's also at the core of the esports category) — timing is everything.

"Unlike asynchronous games, where players take turns to make their moves, MOBA games are synchronous — all the players need to see everything at the same time. So if I shoot an opponent, the arrow or bullet needs to hit the target in real time. Otherwise, the experience is completely messed up, and players soon get frustrated," says Mallayeva.

Transferring and processing all of that data and content across a distributed player base — rendering landscapes, managing synchronisation, downloading large files — is hugely demanding in terms of compute power as well as network bandwidth, performance and scale. Publishers and distributors are increasingly realising that relying exclusively on centralised cloud compute can put the gaming experience at risk. In particular, there's very little tolerance for network latency. The greater the distance between a player and the computing resources delivering the game experience, the more chance there is of it being less than optimum. There's also the added burden of network transit and egress charges, which can be considerable.

To mitigate these issues and enhance the gaming experience, publishers and distributors are starting to move selected workloads and compute resources to the network edge — positioning them closer to users to leverage local content caching and data processing. With this approach, you shorten the distance to the user and reduce the network latency.

"Delivering from the edge speeds up performance and provides a more consistent and reactive gaming experience, which is really important for the big online multiplayer games," says Mallayeva.

Another way that providers can reduce latency between users is through improved network peering, which often requires a provider that offers global scale and regional connectivity to local ISP networks. By placing managed edge servers in key strategic locations, providers can bring routing and peering decisions closer to where their end users are located, thus improving the performance as those users interact with each other over the internet.

At the same time, relying on resources at the edge helps to lower network transit and egress charges, and improve your ability to scale up your games as more players come on board.

Additionally, a flexible edge application delivery solution can also help gaming and gambling providers meet the need for modular security solutions. The popularity growth of online gaming and betting platforms has increased the risk of malicious actors who try to steal sensitive user data or disrupt the unified gameplay. Due to both the growing popularity of online gaming as well as the predictability and specific locations of end-user devices, the industry saw the largest increase in the number of cyber attacks after the pandemic hit in 2020.[3]

3. Lifars.com

To counter those attacks at the point where they occur, providers must be able to deploy customisable security solutions at the edge in the form of web application firewalls (WAF), bot managers, DDoS mitigation, and API protection services. But because security threats are so widespread and varied across platforms, regions, and networks, these solutions must be flexible enough to meet the specific demands of the provider while also working seamlessly with the underlying infrastructure.

Is the future of gaming mobile?

"Although many people play on PCs and consoles, mobile devices have opened up gaming to a whole new demographic of players, especially in regions like Eastern Europe and China," says Mallayeva. **"There's a lot of interest in this growing market."**

As the popularity of mobile games increases, the games themselves will become more sophisticated — making greater use of augmented reality (AR), for example. **"We're even seeing mobile shooter games with virtual joysticks,"** says Mallayeva.

To deliver a great mobile gaming experience, publishers and developers will need to find a middle ground between delivery from the cloud and device-based processing. Relying on the former may negatively affect the experience (owing to latency), while the latter often requires gamers to have high-end devices and, even so, tends to quickly drain the device battery.

Moving portions of content, processing and application delivery to the edge will take some of the load off users' devices, and will support timely delivery of more complex and collaborative mobile games with a great player experience.

> **"For competitive and multiplayer games that rely on real-time interaction, the gameplay absolutely has to be synchronous. If the game lags or crashes, it can significantly damage the player experience. So ensuring consistent performance and minimising latency are critical to the success of these games."**
>
> **— Lerika Mallayeva**
> **Founder and CEO**
> **DevGAMM**

LUMEN

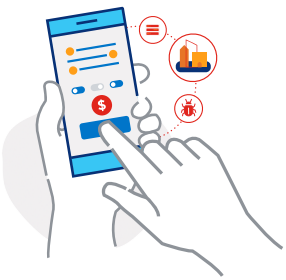

Retail

One would be hard-pressed to find a digital use case that demands lightning-fast performance and dynamic end-user experiences more than online retail and travel. Arguably the most competitive and crowded industry in the digital world, there are many different types of business that fall under the ecommerce umbrella based on the types of consumers they serve as well as the challenges they have to overcome. Therefore, this section encompasses various types of digital storefronts like luxury goods, online car dealerships, home-delivery supermarkets, travel booking sites, wholesale raw material vendors, and others.

The focus for online retailers will be on delivering a smooth, personalised experience. At a minimum, consumers aren't going to wait for a page to load. Beyond that, they want to see pages that are relevant to them — reflecting their buying history, for example, by showing new books by their preferred authors, or the latest lines from clothing brands they've bought from before.

"That means there's much more of a requirement for the data to be right there, very quickly, whatever form that personalisation takes," says Glynn Davis, Founder and Editor of Retailinsider.com, a source of retail industry insights.

In addition to being spoiled for choice with so many retailers and travel providers on the market, many ecommerce customers also have multi-channel buying habits that lead them to browse on one platform (e.g. a mobile app) but buy on another (e.g. desktop website or in store).

This is where application delivery at the edge can play a decisive role. Instead of delivering all of the data from a centralised location, such as a cloud datacentre, retailers are placing selected workloads and computing power as close as possible to where consumers are located. To meet the demand for exceptional site speed and performance, they must be able to deploy content acceleration, dynamic and static object caching, API caching, image optimisation, and instant purge solutions at the edge. Meanwhile, security imperatives must be met with flexible and customisable WAF, bot management, DDoS mitigation, SSL certificate management, and geo-fencing solutions that can be deployed on top of the network infrastructure.

Edge application delivery can make a significant impact in the retail sector when it comes to dynamic content, where latency can adversely affect the consumer experience. **"Retailers can't afford the time it takes to send all of that data back and forth — shoppers simply won't wait for a page to load,"** says Davis. **"Nor do retailers want to bear the data transit costs involved."**

Taking the concepts of dynamic content and personalisation a step further, he discusses the delivery of dynamic pricing on a per-consumer basis.

"We know that hotels and airlines use dynamic pricing, and that it works well because none of us knows what the person in the next room or the next seat was charged," he says. **"I see no reason why [more retailers] couldn't extend that to individual shoppers in other retail contexts, using computing power at the network edge to store and process what's likely to be a hefty data load — about the item, the market, the consumer's search and buying history, and so on — and display a highly personalised price."**

Compounding the importance of site and application performance is the fact that retail companies are forced to rely heavily on third-party services and APIs to meet the consumer demand for personalised, intuitive, and user-friendly shopping experiences. Doing so most often requires them to implement third-party payment processors, shopping cart and checkout workflows, pricing and inventory trackers, and a wide variety of other services on their sites.

A richer in-store experience

Davis points out that processing data at the edge also supports enrichment and personalisation of the in-store shopping journey. For many consumers, especially the younger generation, going shopping is a social experience that goes beyond buying goods. And whereas companies used to try and recreate the physical store online, now they're attempting to bring some of those digital experiences into the brick-and-mortar shop.

"There are lots of different ways for retailers to deliver those digital experiences in store, and the choices they make will depend to some degree on what they're selling," says Davis. So a store might use augmented reality (AR) to let customers virtually try on lots of pairs of sunglasses. A clothing store might team up with the coffee shop next door to send promo codes for hot drinks to its account-holders' mobile phones. A drive-through restaurant could even use automatic numberplate recognition (ANPR) to propose menus that a car's occupants have chosen in the past.

"As well as having the right technology, you also need to work with a lot of data to do that kind of thing really effectively," he adds. **"This is something that omnichannel retailers have been working on for a long time, as they bring data together from all their interactions to build a comprehensive view of each customer, and deliver a coherent experience in every channel."**

When convenience is king

Of course, there are times when consumers simply want a fast and convenient shopping experience — when buying groceries, for example. This is why some retailers are experimenting with checkout-less stores, where images and video from cameras around the store — sometimes combined with shelf-based weight sensors — ensure shoppers are charged correctly for the items they buy. **"This is another situation where you need the ability to process data close to the customer,"** says Davis. **"Otherwise, the speed and convenience benefits of checkout-less shopping simply get lost."**

Self-serve kiosks — underpinned by edge computing — provide another route to convenient shopping. They allow retailers to place their products closer to target customers (sportswear outside a stadium, say), or in locations they might otherwise struggle to reach (flip flops by the beach, for example).

And even though it's unlikely the retailer will be able to bring historical data about individual customers to the self-serve kiosk experience, there's still the opportunity to personalise it somewhat, as Davis explains:

"Cameras can be used to gauge, for example, the customer's age and gender, and show them the most appropriate products on the home screen," he says. **"Really, the closer you can bring data processing and application delivery to the customer, the greater the potential to deliver richer, more meaningful retail experiences in any channel."**

> "When retailers started building e-commerce websites, they tried to recreate the physical store online. Now it's come full circle, as they look to bring the rich, personalised, digital experience into the store. The best way to do that is to make really effective use of the data you've collected via all your interaction channels."
>
> — Glynn Davis
> Founder and Editor
> Retailinsider.com

Selecting an edge delivery solution

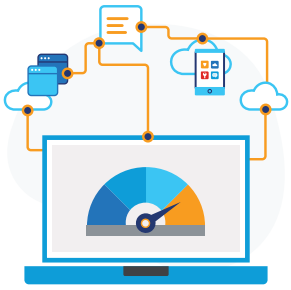

The criteria for any company searching for an appropriate edge delivery solution will obviously depend on their use case, long- and short-term needs, available budget, and digital maturity. But despite this variability, several core components separate the leading edge solutions from the rest of the market:

- Powerful and reliable underlying network and infrastructure
- Ability to quickly scale delivery and orchestrate workloads across locations
- Software-defined delivery solutions to improve the end-user experience
- Edge computing capabilities to improve response times and save bandwidth
- Modern security solutions to protect larger attack surfaces against increasingly sophisticated and varied threats
- Programmability to accommodate custom code or proprietary edge solutions
- Flexible managed service model to drive operational efficiency

Long recognised as a leader in the content delivery space thanks to the best-connected CDN on the market, Lumen has invested heavily in edge delivery technologies to create one of the most powerful and multi-faceted edge application delivery offerings available.

Physical content delivery network

The foundation for the Lumen edge application delivery solutions lies with the underlying physical network, which is interconnected by the largest fibre optic backbone in the world and offers 300 Tbps of egress capacity, including 170 Tbps of CDN server capacity. The edge data centres that power this network are strategically placed in 95+ major metropolitan areas across over 40 countries, giving content providers physical proximity to their end users wherever they might be located.

Software-defined delivery

Lumen also provides the ability to expand that network capacity with software-defined delivery capabilities at the device edge. Lumen CDN Mesh Delivery works in tandem with the CDN and uses a hybrid peer-assist approach to offer flexible scale and performance improvements for video delivery and large file downloads.

Lumen also offers mid-stream multi-CDN switching with CDN Orchestrator. Using feedback collected directly from the user devices, CDN Orchestrator can detect performance issues from the CDN and make a mid-stream switch to a different content source without the user having to refresh the page, thereby preserving the end-user experience in the face of network issues.

Edge computing

Built into its physical delivery infrastructure, Lumen offers a rich ecosystem of edge computing solutions that allow businesses to deploy application logic and data storage at the far edge of the network to improve performance and user experiences.

By partnering with best-of-breed providers that work seamlessly with the underlying architecture, Lumen's flexible approach to application delivery at the edge enables customers to eliminate the need for long deployment and configuration roadmaps, thus reducing time to market, CAPEX and OPEX while improving operational efficiency. These out-of-box solutions include security, image processing and optimisation, dynamic site or application acceleration, container orchestration, and serverless computing.

For companies that prefer to use open-source modules, their own proprietary solutions or custom code, Lumen also offers the ability to build directly onto the delivery infrastructure, thus providing a fully programmable edge for customised delivery.

Security

The network and transport layers within the Lumen CDN and IP backbone continue to provide world-class DDoS protection. In addition, Lumen CDN Edge Compute offers a variety of security solutions at the application layer including web application firewalls (WAF), bot management, Layer 7 DDoS mitigation, and API protection. This edge-based security offering provides businesses the flexibility to choose the security solutions that best suit their needs from a best-of-breed ecosystem, including ThreatX, PerimeterX, and Wallarm. Businesses also have the option of bringing their own security solutions into the Lumen ecosystem.

LUMEN

Managed services

Lumen prides itself on being flexible enough to work with enterprises of all sizes and across industries through a customisable service offering that matches an organisation's digital operational capability with the appropriate level of support. For more digitally advanced companies, Lumen edge application delivery is available on a purely self-service model. But for companies that don't yet have the ability to stand up and manage complex IT architecture on their own, the Lumen customer success teams are able to operate as an extension of their clients' operations, DevOps, and site reliability engineering (SRE) teams with fully managed or co-managed service models.

Regardless of the level of support required, Lumen customer success teams measure the outcomes that matter most to their clients: end-user experiences. With anomaly detection, automated repair services, and a customer-first approach, Lumen is committed to empowering their clients to deliver consistent, high-quality user experiences all over the globe.

Lumen's solutions are built to allow businesses of all sizes deliver fast, reliable performances and great customer experiences all over the world. To learn more, reach out to **application.delivery@lumen.com**

+44 (0) 800 496 5000 | lumen.com | uk@lumen.com