



Intel® Network Builders Program Goes to the “Edge”

Introduction

The growing opportunity for “edge computing” is increasingly material and will grow more so with the advent of complementary technologies such as 5G and internet of things (IoT) devices. A number of vendors have announced or disclosed plans for edge-oriented products and services. Still there remains some confusion as customers explore options and vendors scramble to maximize their opportunity amid this growing interest.

The logic and drivers for edge computing solutions are sound. Over time customer interest and vendor marketing will match and focus on real solutions and capturing business value. But just waiting to naturally reach such equilibrium is not an option for vendors or their customers. It is critical to develop and prove business value as quickly as possible to accelerate adoption. This is the very task that the Intel® Network Builders program has long focused on and will now apply to a dedicated, focused community. This new program aims to reduce development time through verified and optimized solutions, simplify and accelerate ISV/application onboarding processes, and to amplify and promote innovation and leading network edge solutions.

This paper will provide an overview of the edge opportunity, the motivations, ecosystem, drivers, benefits, and how the Intel® Network Builders program is evolving to address these.

What Is “The Edge”? General Definition

Edge computing is the placement of data center or network-class compute and storage resources away from centralized hubs and closer to, or into, smart endpoint devices in order to optimize TCO, reduce application latency, improve service capabilities, comply with data privacy or locality requirements, and reduce application latency.

The European Telecommunications Standards Institute ([ETSI](#)) has recognized the growing interest and opportunity in enabling edge computing as a way to enable new real-time services and alleviate core network congestion. The organization has initiated the development of standards and platform definitions for multi-access edge computing ([MEC](#)).

Many point to the imminent arrival of 3rd Generation Partnership Project ([3GPP](#)) 5G technologies as synonymous with edge computing. It is undeniable that edge computing models will be dramatically enhanced and massively proliferated when complemented by the additional scale and the performance and functionality gains delivered by 5G technologies and standards. Edge computing models will indeed unlock greater business value (cost savings as well as new services opportunities) that the wide deployment of 5G enables. But edge computing is not reliant on 5G and in fact can deliver benefits in today’s network and mobile technologies.

Table of Contents

- Introduction 1
- What Is “The Edge”? General Definition 1
- Where Is the Edge? 2
- Why the Edge? 2
- What Are the Industry Drivers? 3
- Who Are Edge Solutions Providers? . 3
- What Is Intel's Role on the Edge? 4




Figure 1. Illustration of the Network Edge

Where Is the Edge?

There is no universal “edge” to all networks or compute infrastructures across the industry. In many ways, the edge is specific to the use model, workload, or customer. In general, “the edge” is the outermost layers of the end customer’s processing or network reach, before data transitions to endpoint devices or other networks. The demands on these outermost layers are growing, including the need to run an increasingly diverse set of workloads. Customers need enterprise, cloud, and communications infrastructure to adapt in response. They need these infrastructures to be capable of hosting and supporting the growing data processing, transmission, and storage needs across a much larger portion of the IT infrastructure footprint.

As shown in Figure 1, there are many physical locations and devices throughout the IT infrastructure that will host workloads meeting the definition of “edge.” And there are compelling reasons that will drive customers and service providers to utilize this infrastructure for business innovation.

Why the Edge?

Customers have readily adopted cloud from hyperscale providers and implemented on-premises private clouds. But there are classes of workloads and use models that will not be well served by either of these two general models, and there are market needs for trusted, local, distributed clouds. These clouds will provide access to compute, network, and storage resources at increasingly dispersed locations. These local clouds will support workloads and use models requiring low latency, high speed, and high scalability. They will do this via secure, controlled, and cost effective infrastructure located ever closer to where the access is desired. In short, customers/service providers will need (and have) access to

a broader set of options when it comes to deciding where to place their workloads.

Today, perhaps the primary criteria used to determine where to place workloads are:

1. **Performance** – Customers value the flexibility for a hyperscale cloud service provider (CSP) to scale and adapt flexibly to the performance and resource needs of unpredictable or dynamic requirements or for fast prototyping of new services. But other workloads, such as those with high latency sensitivity or where IT needs direct oversight of application control, are more often landed in on-premises infrastructures.
2. **Security** – Customers often value the security expertise of their CSP and the ability for that provider to be accredited by standards or regulatory entities. But there are also a growing number of instances where a customer will not/cannot place data or workloads into a public cloud due to internal risk controls, data privacy regulations, or other governance constraints and instead host these on premises.
3. **Integration** – Customers have long enjoyed the benefits of placing standalone workloads, pilots and test/development work in public cloud infrastructure and using the APIs provided by cloud providers to build cloud-ready apps. But there are still a large number of applications that have not moved to public cloud infrastructure and remain on premises. The primary reasons are that the applications and data are so complex or integrally intertwined with other applications that moving to a public infrastructure has been too daunting to undertake or that it would otherwise “break” or compromise the application.

4. **Data size** – The types of applications that have thrived on public cloud infrastructures are generally those with low-to-moderate data volumes with frequent access, or higher volumes with less frequent access. Much of this is a factor of the time and cost associated with transporting large volumes of data over distances. In many cases it is just less expensive and a better experience to keep bulk data closest to where it is needed—such as on-premises infrastructure.

In considering these key attributes, it is inevitable that the largely binary choice customers face for workload placement is challenged. Increasing demands for cost efficiency, a need for high responsiveness, and growing volumes of data from disparate, distributed sources are creating workloads and use models underserved by either a large hyperscale cloud offering or a private cloud hosted on premises.

More and more, two sizes won't fit all, and this drives the need for more local, distributed, and capable IT resources to efficiently and securely host key workloads and use models. As a result of these various requirements, many customers will ultimately adopt edge solutions to supplement what some are calling a “multi-cloud” strategy covering on-premises computing and services from multiple service providers.

What Are the Industry Drivers?

New use models and services are making it increasingly clear that more options for workloads and data placement are needed. Some of the leading use models include:

- Autonomous driving
- Data caching and storage
- Video and video analytics
- Speech recognition and voice processing
- Augmented and virtual reality and gaming
- Internet of things
- Medical

In each of these general examples, one can easily imagine ways that local access and processing provides either dramatic savings in transport costs or the low latency needed for the responsiveness of the application to meet requirements. Such a distributed model would also provide a more fine-grained ability to constrain data or applications to locales where there are corporate or governmental regulations that require such control.

Who Are Edge Solutions Providers?

Edge solutions will evolve from many technology domains as many possess valuable assets to enable the use models customers seek. Of course, the customers (led by those enterprises in the verticals above) will drive the evolution as they will ultimately pay for and benefit from these services. On the supply chain side, the ecosystem will comprise the hyperscale and other leading CSPs, communications service providers (CoSPs), independent software vendors (ISVs) and systems integrators, and hardware providers/systems vendors, as shown in Figure 2.




Figure 2. Participants in the Network Edge Ecosystem

Each of these participants in the ecosystem brings prominent assets to enable the edge components of multi-cloud solutions and expects compelling opportunities from growth in this area.

Cloud service providers are traditionally thought of in terms of their hyperscale data center cloud capabilities. They often feature several such sites to provide redundancy and geographic coverage. This semi-distributed model has been helpful for them to grow their customer base and improve availability of services, but it has limits. For example, the ability to distribute massive hyperscale data centers is hindered by regulatory and economic concerns; in many cases, the market opportunity just doesn't require such scale. But their ability to automate and integrate in hybrid cloud fashion and their established APIs make them a very powerful enabling engine for the edge opportunity.

Communications services providers (CoSP) have two primary capabilities that few if any of the other ecosystem participants have—a wide and very distributed physical footprint of real estate and established communications services providing “last mile” connectivity for fixed and mobile users. Increasingly, CoSPs have been interested in utilizing these distributed assets more effectively to grow revenue (via new services and driving new traffic) even as they modernize and update these sites' efficiency by deploying network functions virtualization (NFV) and industry-standard server platforms. With their local presence and established business relationships, CoSPs could logically be the connectivity and local hosting presence for any number of the edge-based workloads discussed previously.

Software vendors see the opportunity to provide more services to existing customers and new or enhanced services to new customers that require the low-latency, processing, or analytics or security capabilities that are best delivered on the edge. Similarly, the **hardware providers** for the cloud and communications infrastructure (systems vendors/OEMs, ODMs, network equipment manufacturers) are addressing the need for more flexible, scalable, cost-effective, and increasingly cloud-ready platforms for enabling communications and other new business services throughout the network from core to access points and beyond. The portfolio of Intel® architecture platforms provide them the opportunity to scale across this breadth of requirements efficiently. Lastly, **systems integrators** will fill their valuable role of trusted partners that help meld new technologies with business operations. They too see great opportunity in providing new services in application development and integration and partitioning of business processes across distributed cloud infrastructures and addressing the new requirements that the edge opportunity allows.

Intel is creating a network edge ecosystem within the Intel Network Builders program that will provide a focal point for the growing number of customers looking for edge network and cloud solutions. As such, this initiative can bring together all of the players involved in the development, deployment, and consumption of a variety of network edge services and solutions across businesses and verticals.

What Is Intel's Role on the Edge?

Intel has great experience in helping to bring ecosystem components together to enable major architectural changes in the IT landscape. Intel has helped enable the revolution towards cloud computing for compute, storage, and networking technologies—with common tools and platforms across those domains.

Intel platforms, products, and technologies provide the foundation for delivering the processing, storage, or connectivity capabilities to wherever they are needed. From robust Intel® Xeon® Scalable processor platforms, to low power, small form factor devices Intel offers server platforms to meet every space, power, or performance need. With high performance and reliable Ethernet, solid-state drives, accelerators, and other components, solution providers have a highly flexible set of components to build out the infrastructure to support diverse customer needs.

While Intel has value as a “trusted advisor” providing a portfolio of powerful, efficient, and open platforms for customers and the ecosystem, Intel is also uniquely positioned to help in deep optimization and debugging work in many platform technologies—processor, BIOS, memory, storage, networking, and more. Such work can be foundational to the ability to derive the performance, scalability, efficiency, and stability needed for networking and other IT/business solutions.

Intel® Network Builders has helped provide a catalyst for the burgeoning market for open, standards-based networking solutions. It has brought together the required ecosystems to develop and promote usable, scalable, and increasingly flexible and efficient solutions such that networking benefits from the same economics that compute and storage have long enjoyed thanks to cloud technologies. The mission now is to help cultivate solutions that can provide edge computing and access capabilities to deliver real customer value and to help the market absorb them through education, promotion, and collaboration from enterprises, through service providers and systems and software vendors. The emergence of edge computing is a significant component of network transformation for service providers and for enterprises alike.

Learn more about what Intel is doing to help transform networks at www.intel.com/network and about the new network edge ecosystem within Intel Network Builders at networkbuilders.intel.com/networkedgeecosystem

