

MARKET PERSPECTIVE

Intel Positions Itself as a Key Enabler for Edge Solutions

Max Pepper

Dave McCarthy

Jennifer Cooke

EXECUTIVE SNAPSHOT

FIGURE 1

Executive Snapshot: Intel Positions Itself as a Key Enabler for Edge Solutions

Intel is tapping into its experience across multiple industries to power solutions and partners that deliver edge solutions. When mapping out an edge strategy, industries need digital solutions that are secure, rapidly deployed, and affordable. As a trusted partner with exposure to challenges across many vertical industries, Intel is producing chips optimized for the challenges of edge compute. In addition, it is investing in software, APIs, and partners to accelerate the shift to new edge locations.

Key Takeaways

- Intel has a long history of solving IT and operational challenges with hardware innovation. As the compute landscape extends to new locations, the company has made significant investments in not only chip technology but also software tools for developers and partners to streamline the move to the edge.
- Intel is innovating to provide technology and solutions that simplify ongoing management. Autonomous operations, hardware-based security, and an ecosystem approach to edge compute will enable the rapid deployment and resilient operation of new solutions.
- As the competitive landscape continues to evolve, Intel is doubling down on its central mission: to be an enabler of innovation.

Recommended Actions

- The edge compute ecosystem is evolving rapidly, leaving little time to build from the ground up. As the race to innovate and develop solutions becomes more intense, leaning on proven hardware and software platforms, as well as trusted partners with industry-specific experience, will improve chances for success.
- Ensuring security is table stakes for edge compute deployments. Consider platforms infused with silicon-level telemetry to reduce risk and defend against threats.
- AI and high-performance compute workloads will strain edge hardware and critical infrastructure. Seek providers and solutions that can tackle the increased demands on compute resources with purpose-built and optimized hardware.

Source: IDC, 2021

NEW MARKET DEVELOPMENTS AND DYNAMICS

As edge computing continues to radically expand with growth of IoT and 5G technology, Intel has been formulating and executing its edge strategy. Over the past few years, Intel has shifted its efforts into high gear. The company is fast becoming a serious competitor with unique advantages, including a diverse product portfolio increasingly optimized for edge-specific uses, a robust set of developer tools to drive edge use, and a large-scale partner ecosystem to support edge infrastructure.

Industries are looking for secure, resilient, and proven solutions as they extend compute resources to new locations. 77% of U.S. organizations say that edge is a strategic investment for their business, and they expect that edge resources will allow faster decision making and increased productivity and efficiency. The need for edge compute is great, and the landscape of providers is broad.

Rapidly deploying solutions allows little time for building from the ground up. Focused on outcomes and driving value, these industries seek to put less attention on assembling solutions and more on ways to gain value from them. They need trusted partners to build a scalable foundation for edge compute. This is where Intel sees an opportunity to provide value and accelerate edge deployments. The company's strategy extends beyond providing chips optimized for edge deployments. With its deep understanding of the importance of security and speed in successful edge deployments, Intel is building a foundation upon which developers can quickly bring solutions to market.

Intel is tapping into its experience across multiple industries to power solutions and partners that deliver edge solutions. When mapping out an edge strategy, industries need digital solutions that are secure, rapidly deployed, and affordable. They need proven solutions that understand the importance of securing data and physical infrastructure. As a trusted partner with many years of experience, Intel is working across the ecosystem in new ways to drive solutions that simplify ongoing management.

The company has a long history of solving IT and operational challenges. Intel understands and appreciates the value of speed and security. Key to rapid deployment and ongoing management of edge solutions is shifting to more autonomous operations. Intel has designed its technology to reduce the strain of hands-on management by enabling autonomous controls. The need for greater autonomy and self-regulation becomes much greater when logistics prevent human intervention.

The technical and organization challenges of shifting from centralized, core compute to a distributed strategy are well understood by Intel. The company defined itself as the technology to disrupt an era dominated by centralized, mainframe computing. Intel semiconductors proliferated servers and PCs that lived outside of a datacenter. As a trusted, experienced partner that was a central player in the shift from mainframe to distributed computing, Intel is uniquely qualified to guide the industry into the next phase of edge computing.

Intel's Edge Portfolio

Hardware

Intel's product portfolio, with a consistent release schedule, has been key to building up Intel's Edge strategy. Often typecast as a company with a single architecture, the company is expanding its reach in four main computing chip types – CPUs, GPUs, application-specific integrated circuit (ASICs), and field-programmable gate arrays (FPGAs). Intel has made immense strides in most of these areas related to edge via the evolution of its new multifaceted XPU strategy. This includes a new proprietary chip type, the vision processing unit (VPU).

CPU

In April 2021, Intel announced new Ice Lake 10nm processors, optimized for high-performance workloads. Ice Lake processors can be configured with up to 40 cores and have 8 memory channels, driving performance up 46% according to Intel. With built-in acceleration and new instructions, Ice Lake provides a significant performance boost for demanding new AI, networking, and cloud workloads. On the networking side, Intel launched the Atom P5900, a 10nm system on chip (SoC) built for the high-bandwidth and low-latency demands required of 5G wireless base stations. It also recently launched its Atom x6000E, or "Elkhart Lake", series, which is designed specifically for IoT usage.

In addition to these larger-scale processors, Intel also installed the latest technology on its desktop processors that can be critical to enhancing edge performance. Its recently unveiled 11th Generation Intel vPro Platform includes Intel's SuperFin technology and Hardware Shield, providing high levels of AI performance and productivity potential along with top-level security.

As Intel continues to steadily release cutting-edge technology with features that help improve efficiency and lower latency, it continues to build itself up as a heavyweight hardware provider in the edge field. In 2021, Intel released its "Rocket Lake" 11th Generation Intel Core S-Series Desktop Processors, with 19% gen-over-gen instructions per cycle improved, and is expected to release "Alder Lake" processors that will be Intel's most power-scalable system-on-chip processor with a more enhanced version of Intel's 10nm SuperFin technology.

Intel's Tiger Lake announcement targets the industrial IoT segment with a processor optimized for AI and security applications.. Tiger Lake is the 11th Gen Core processor launched for ultrathin laptops as well as IoT applications that require high-speed processing, computer vision, and low-latency functionality.

VPU

Intel's Movidius VPU, or vision processing unit, came from Intel's acquisition of Movidius in September 2016. The VPU is a low-powered processor specifically designed for edge computing processes focused on image processing with purpose-built "SHAVE" processor cores to maximize performance for space. Intel disclosed the Gen 3 version of the processor under the code name of "Keem Bay" in late 2019, which features 64-bit memory bandwidth and projects improved performance for power and more inferences per mm² compared with the previous Myriad X chip.

Due to the low power needs and versatility of the VPU, these processors have played a major role in recent edge collaborations. Most notably, in late 2020, these chips were included in the PhiSat-1 satellite, the first time a satellite with local AI processing onboard has been launched, which included a new type of hyperspectral-thermal camera. The processor is being used to remove images taken with a new type of hyperspectral-thermal camera onboard that includes cloud cover to save bandwidth. The VPU is also part of the SoM for Microsoft's new Azure Percept Edge AI development platform, meant to combine hardware and services to make implementation of AI at the edge easier through Azure.

FPGA and eASIC

Intel has played a major role in boosting the use of field-programmable gate arrays, or FPGAs, to better support the company and its partners' edge capabilities. These semiconductor-integrated circuits, unlike a standard CPU that is fully baked in, allow a programmer to use a chip for several complex functions up to the point of acting as a comprehensive multicore processor. This allows for better optimization for workloads, particularly "smart" workloads that allow for adaptability within a

process. It also allows for dynamic reprogramming to match specific workload needs with less power requirements. While useful for many edge-related workloads, including data analytics and media streaming, the biggest impact FPGAs have is in significantly lowering the cost and improving performance for the deep neural networks used for AI.

Intel has made major investments into this technology, releasing several series ranging from the entry-level Intel MAX and Cyclone series to the more performance-intensive Intel Stratix and Agilex series. The latter two are particularly important for Intel's edge support efforts due to their high-performance capabilities, particularly with the Agilex, which is Intel's first FPGA built on 10nm process technology and 2nd Generation Intel Hyperflex FPGA Architecture. Toward the end of 2020, Intel also announced the Intel Open FPGA Stack, which gives developers a scalable and open infrastructure to build systems with FPGA-based acceleration systems. This makes FPGAs much more accessible due to simplifying the development of FPGA-based accelerators as well as taking a step toward open source usage of FPGA tools versus the closed source nature of competitors.

Intel also made a major announcement about its eASIC line. ASICs, or application-specific integrated circuits, are critical in wireless and cloud environments to allow programmable logic, and FPGAs are closely tied to these chips since their flexibility is useful in lowering the cost and risk of prototyping. However, FPGAs can become a liability beyond the prototyping stage due to their lack of efficiency compared with cost and performance. To streamline and improve ASIC development and use while further integrating itself in edge and IoT development, Intel acquired eASIC in 2018. This company designs "structured ASIC" chips, which, in simple terms, are between an ASIC and an FPGA. They have fewer efficiency issues than using an FPGA and are allowed for faster ASIC development.

In late 2020, Intel announced the first new version of the structured ASIC technology to come from eASIC since its acquisition. The new eASIC N5X is the first structured ASIC to be designed specifically with Intel's FPGA system. This makes mapping complete FPGA system-on-chip designs to an ASIC much easier, since many designs are taking advantage of the built-in ARM-based embedded processing systems on Intel's FPGAs. It is also important to note that the eASIC N5X can map designs from other FPGA technologies, making it much easier for Intel's technology to be integrated into chipsets previously using other ASICs. This gives Intel a leg up when volume production is increasing as it can better meet these needs compared with other competitors, helping make the company a foremost producer for these chips in edge products.

Software

As it builds out its product portfolio, Intel also has robust software offerings to have a fully realized edge ecosystem infrastructure to support developers using its edge hardware, no matter the vertical usage.

OpenVINO and Open Visual Cloud

Intel's Open Vision Inference and Neural Network Optimization toolkit, or OpenVINO, is made of two components: the Model Optimizer and the Inference Engine. The Model Optimizer converts a trained neural network from the source framework to an open source intermediate representation for inference operations. The Inference Engine performs several functions, including running inference operations on input data, model loading, and compiling management, and execution support for Intel's hardware types including Intel CPUs, GPUs, VPUs, and even GNAs through various software plug-in architectures. Along with the toolkit, developers can also use the Open Model Zoo repository on GitHub to help develop deep learning modules.

This toolkit has been central to improving the AI-based workloads capabilities of Intel's hardware products, particularly for image processing and deep learning inference tasks. The ability to easily switch hardware can be critical as models are moved into production and away from ideal deployment environments. This can save time and ensure optimal implementation of equipment for workloads in a variety of environments. It also allows for easier scaling of applications among different end users, making the practicality of edge uses much higher.

In addition to OpenVINO, Intel also improved its visual-oriented support for edge applications via its Open Visual Cloud, its software stack for cloud-native development optimized for Intel CPUs and accelerators. It provides open source resources via reference pipelines as Dockerfiles that can help speed development. With Open Visual Cloud, developers gain access to rapidly deployable solutions to encode, decode, infer, and render visual cloud services for media processing, media analytics, immersive media, cloud graphics, and cloud gaming. Open Visual Cloud will also enable lower-latency media processing and analytic solutions to improve performance for network edge deployments.

Multi-Access Edge Computing

Intel has made significant investments into multi-access edge computing (MEC) recently which, while low key, are important in better enabling organizations introducing Intel-based edge applications into their office infrastructure. The most important investment Intel has made in this area was its 2019 acquisition of Pivot Technology Solution's Smart Edge platform, which made Pivot Technology one of Intel's preferred partner resellers. This acquisition let Intel take center stage in not only providing the hardware to move businesses on to the cloud and at the edge via 5G but now providing a major software offering for its customers on which to platform their products. Intel also opened new revenue opportunities for both itself and its customers while lowering TCO for Intel's edge products.

Intel has also made major inroads in terms of MEC development with its Smart Edge Open toolkit (formerly Open Network Edge Service Software). This MEC software toolkit enables edge platforms to optimize onboard and manage applications and network functions across all networks based on consistent and standardized APIs, and its open source distribution makes creating edge computing applications for cloud and IoT developers much simpler. The main specific capabilities it offers include simplifying cloud-to-edge migration by abstracting network complexity, enabling secure onboarding and application management via a web-based GUI, and providing a modular, cloud-native microservices-based architecture for building functionalities. It currently has five releases, with each including improvements to core MEC and network elements, platform optimization, and application support for other toolkits and software. This toolkit is a certified Kubernetes distribution and has become a popular platform for 5G deployment, most recently with Google, Red Hat, and several others.

oneAPI

One other critical toolkit Intel released as it started expanding its edge footprint is its oneAPI toolkit. With Intel having discrete CPU, GPU, ASIC, and FPGA offerings, oneAPI provides a unified programming model that allows for a simplified application development across all these different computing architectures while offering full native code performance. This is mainly done by offering a new open, direct programming language called Data Parallel C++, or DPC++. Based on C++, it also includes a compiler supporting OpenMP supporting continuity with existing codes in other commonly used languages in this field. DPC++ also incorporates SYCL from The Khronos Group for data-parallel and heterogenous programming support across CPUs and accelerators and language extensions from its open community. As an API-based tool, it also offers libraries that work across several workload domains, with functions that are custom coded based on each target architecture.

Cross-architecture computing solutions will be critical to the growth of edge with the expanding number of architectural solutions available across platforms. oneAPI is particularly important in combination with Intel's XPU group to create an industry benchmark ultra-heterogenous computing concept. Intel has continued to push this concept of "no transistor left behind" with the release of oneAPI Gold, the next step beyond oneAPI beta, in late 2020. Not only does it include new hardware capabilities and instructions – further helping set up software and hardware packages to complement each other – but it is available for free on Intel's DevCloud platform, making it more powerful in allowing for more hardware support than would be possible through local integration. This toolkit is one of the more ambitious undertakings for Intel on expanding edge potential and should make it much easier to both develop and integrate new edge applications through Intel devices.

Partnerships and Ecosystem

Along with developing this software-based support for its edge ecosystem, Intel has built out several partners for its edge ecosystem. For network edge, Intel has developed its Network Builders Edge Ecosystems, currently led by ADVA, Advantech, Cloudfy, Lanner, Nearby Computing, Nokia, Parallel Wireless, QCT, Red Hat, Wind River, Wipro, World Wide Technologies, and ZTE, along with 84 other partners. On the on-premises side, Intel has its Edge AI Ecosystem, which includes Foxconn, QNAP, DycodeX, Sulubaai, and Hitachi.

Intel and Ericsson continue to collaborate on joint solutions for radio access network (RAN) deployments. With incubation projects for 5G use case solutions, Intel's compute and processing capabilities are meshed with Ericsson's RAN algorithms and software to create virtualized and cloud-native solutions.

On the software side, Intel has its Edge Software Hub, with nearly 50 partners helping optimize Intel's edge software offerings. Intel has also been using its Intel Retail Edge program to set up hundreds of ecosystem partnerships. For developers, DevCloud for the Edge allows users to build prototypes and experiment with AI workloads for computer vision. Developers can run AI applications remotely on Intel hardware, with the ability to view telemetry on intensity and use conditions to determine the best hardware for their solution.

ADVICE FOR THE TECHNOLOGY SUPPLIER

Enterprises and service providers are increasingly understanding and appreciating the complexity of assembling all the pieces to deploy resilient and secure IT service in new edge locations. As their edge innovations mature, they are realizing the importance of partnering with technology providers that can drive efficiencies and best practices in operations. Technology suppliers that are able to bring together hardware and software solutions will be better prepared to compete in an increasingly solutions-focused ecosystem of partners.

Over the past few years, Intel's customers have become competitors and branched out to drive innovation in different directions. The complex ecosystem of customers, partners, and suppliers continues to evolve. With this evolution, Intel's strategy also needs to evolve to underscore the company's position as an enabler of innovation. Autonomous operations are critical to a solid edge compute foundation. Edge deployments need technology that can speed deployment, ensure security of data and physical infrastructure, and enable remote monitoring. Intel's investments in secure hardware and open APIs are key to affirming the company's position as a key driver of edge innovation.

With new CEO Pat Gelsinger at the helm, significant strategic changes could be on the horizon. Although Intel has a solid record of innovating, having a fresh new approach and strategy as the ecosystem matures may be just the boost the company needs to elevate its position.

LEARN MORE

Related Research

- *Intel Provides New Tools to the Cybersecurity Task: The Results Could Be Game Changing* (IDC #US47495021, March 2021)
- *IDC Connections: Delivering Innovation at the Edge* (IDC #US47563821, March 2021)
- *The Most Critical Features for Edge-Deployed Servers* (IDC #US47538221, March 2021)
- *IDC Market Glance: Edge Infrastructure, 1Q21* (IDC #US47329921, January 2021)

Synopsis

This IDC Market Perspective analyzes Intel's approach to powering edge solutions. We provide an overview of Intel technology optimized for edge deployments, tools to drive developer innovation, and its strategy for leveraging the partner ecosystem to deliver secure and resilient edge solutions.

"As edge computing is increasingly becoming a critical part of any organization's IT infrastructure, Intel is in a strong position as an industry leader trusted in leading the way under changing technological dynamics," said Max Pepper, senior research analyst, Infrastructure Systems, Platforms and Technologies at IDC. "It has made key investments in technology underpinning edge, and with its new leadership, it has an opportunity to make even greater gains under a fresh new direction."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or web rights.

Copyright 2021 IDC. Reproduction is forbidden unless authorized. All rights reserved.

