



The Edge Infrastructure Handbook

2021 Edition

Comprehensive Guide to Designing and
Deploying Scalable Edge Networks



1.0

Introduction

03-06

[Edge Computing](#)

[Your Data Stranded in the Wild](#)

[Extend Beyond Traditional Data Centers](#)

[Success at the Edge](#)

2.0

Real-world Examples

08-40

[Audi](#)

[Philips](#)

[ViewSonic](#)

[ExxonMobil](#)

[HOOBOX](#)

[Sensormatic](#)

[Accenture](#)

[Rakuten Mobile](#)

[Lumen](#)

[Verizon](#)

3.0

Key Considerations for Edge Deployments

42-53

[Hardware for Varied Data Types](#)

[Capitalize on Easy-to-use Software Tools](#)

[Ready-to-Deploy Solutions](#)

[Implement AI at the Edge](#)

[Hardware-based Security](#)

4.0

Discovering Solutions

55-58

[Expert Support From a Trusted Partner Ecosystem](#)

[Scalable, Repeatable Solutions](#)

[Your Fast Path to Edge Development](#)

Enable better business outcomes “born at the edge.”

In the age of IoT, Industry 4.0, and 5G Networks, even the smallest technological advancements can result in financial returns or a competitive advantage. And with the boom of new edge devices, sensors, and controllers generating an exponential growth in data, companies are realizing even more value in harnessing the power of edge computing.


While connectivity, bandwidth, and power may be subject to disruption, migrating computational power from the cloud closer to where data is generated can turn milliseconds of latency gain into millions in revenue. All via greater efficiency, automation, and instantaneous on-the-spot analysis.

The real value is in recapturing precious latency and monetizing untapped data by putting AI right at the data source, enabling better business outcomes to be born at the edge.

It's time to change your approach if your data is stranded "in the wild."

Businesses across nearly every category are accruing new value by acting on unstructured data at its source, even in the most remote or demanding environments, with every advance in computing power and edge-based intelligence.

But deploying and maintaining edge resources comes with its own set of challenges because, by nature, they are decentralized. Companies realize new capabilities where direct, near real-time, multidirectional communication among machines, sensors, and other devices are paramount. These deployments require a shift in the computing and outcome design process that relies on nearly instantaneous decision-making powered by artificial intelligence (AI) and machine learning (ML). And with connectivity changes factored in, physical security and data encryption concerns increase exponentially.



Gartner predicts that by 2025, 75% of enterprise-generated data will be created and processed at the edge — outside a traditional centralized data center or cloud.¹

Implement time-sensitive solutions that extend beyond traditional data centers.

The cloud-centric model that dominates much of the enterprise technology today requires a full and detailed reappraisal when building out edge systems. While “round trip to the cloud” decision support works for some applications, for others, it may limit or even prohibit near real-time decision-making—the core value of edge-focused intelligence.

Some traditional cloud-centric architectures don’t embrace a more distributed model and rely on all intelligence occurring in the cloud far away from the data sources. This consumes unnecessary bandwidth and, consequently, is likely to result in potentially crippling latency issues.

Conversely, the ability to process and analyze data in near real-time at the edge can reduce the cost and complexity associated with sending vast amounts of data to the cloud. That’s because the data generated at the edge can be enormous, making it uneconomical and inefficient to send, store, and process it all in remote, centralized, cloud-based data centers.

Take the industrial sector as an example, where factories and power plant sensors, cameras, and robots could generate petabytes of data. Adding edge computing to a company’s cloud computing infrastructure could lead to financial rewards and an even lower total cost of ownership (TCO), and help achieve the Industry 4.0 objectives of driving manufacturing forward with augmented decision-making and advanced automation.



“Highly time-sensitive IoT applications in manufacturing and other domains require immediate analysis of, or response to, the sensed data. For example, they could warn of roof collapse, air quality, temperature, and humidity in underground mines and indoor manufacturing plants; identify workers who don’t have the appropriate gear or adequate training for a particular machine; or prevent unauthorized workers from entering restricted areas.”²

Success at the edge requires expertise in solving the complexities of IT and operational divides.

Each edge deployment introduces unique sets of challenges that can further complicate the implementation of a new, affordable, secure, and resilient edge solution. At Intel, we've helped companies across a broad spectrum of industries implement tens of thousands of edge deployments, generating real business value. With that experience comes our ability to address many concerns companies face, such as:

- Building new solutions that complement or replace existing legacy systems
- Maximizing interoperability with future-proofing solutions
- Uniting operations technology (OT) and information technology (IT) requirements
- Deploying reliable AI solutions at the edge
- Maintaining data and network security as edge devices introduce new vulnerabilities

Every edge deployment may face unique challenges, but that doesn't make the issues insurmountable. Read on to learn more about the ways Intel has helped companies find solutions that worked for them.

Up to 70% of enterprises will run varying levels of data processing at the IoT edge by 2023.³

¹Rob van der Meulen, Gartner, [Edge computing promises near real-time insights and facilitates localized actions](#), Oct. 3, 2018.
²Georgakopoulos et al. (2016). [Internet of Things and Edge Cloud Computing Roadmap for Manufacturing](#). IEEE Cloud Computing.
³IDC FutureScape, [Worldwide IoT 2020 Predictions](#), October 2019.



1.0

Introduction

03-06

[Edge Computing](#)

[Your Data Stranded in the Wild](#)

[Extend Beyond Traditional Data Centers](#)

[Success at the Edge](#)

2.0

Real-world Examples

08-40

[Audi](#)

[Philips](#)

[ViewSonic](#)

[ExxonMobil](#)

[HOOBOX](#)

[Sensormatic](#)

[Accenture](#)

[Rakuten Mobile](#)

[Lumen](#)

[Verizon](#)

3.0

Key Considerations for Edge Deployments

42-53

[Hardware for Varied Data Types](#)

[Capitalize on Easy-to-use Software Tools](#)

[Ready-to-Deploy Solutions](#)

[Implement AI at the Edge](#)

[Hardware-based Security](#)

4.0

Discovering Solutions

55-58

[Expert Support From a Trusted Partner Ecosystem](#)


[Scalable, Repeatable Solutions](#)

[Your Fast Path to Edge Development](#)

Tens of thousands of repeatable and scalable deployments across the globe. *Yours is next.*


From repeatable factory automation platforms to resilient restaurants and retail stores to cloud-native 5G networks, industries across the board are increasingly striving to generate business value at the edge. Lowering costs; speeding time to market; and improving safety, productivity, and efficiency make successfully incorporating edge computing vital in an increasingly competitive marketplace.

With tens of thousands of deployments helping companies of every conceivable type, Intel brings years of expertise in aligning use cases, bridging technical gaps, solving common integration challenges, and delivering hundreds of production-ready packages—all designed to simplify edge-to-cloud development. From helping automate an Audi factory to enabling Sensormatic's retail social distancing solutions, Intel partners with businesses to help simplify technical complexities and alleviate cost burdens. When you're ready to build and evolve your own intelligent edge, we'll be ready to help.



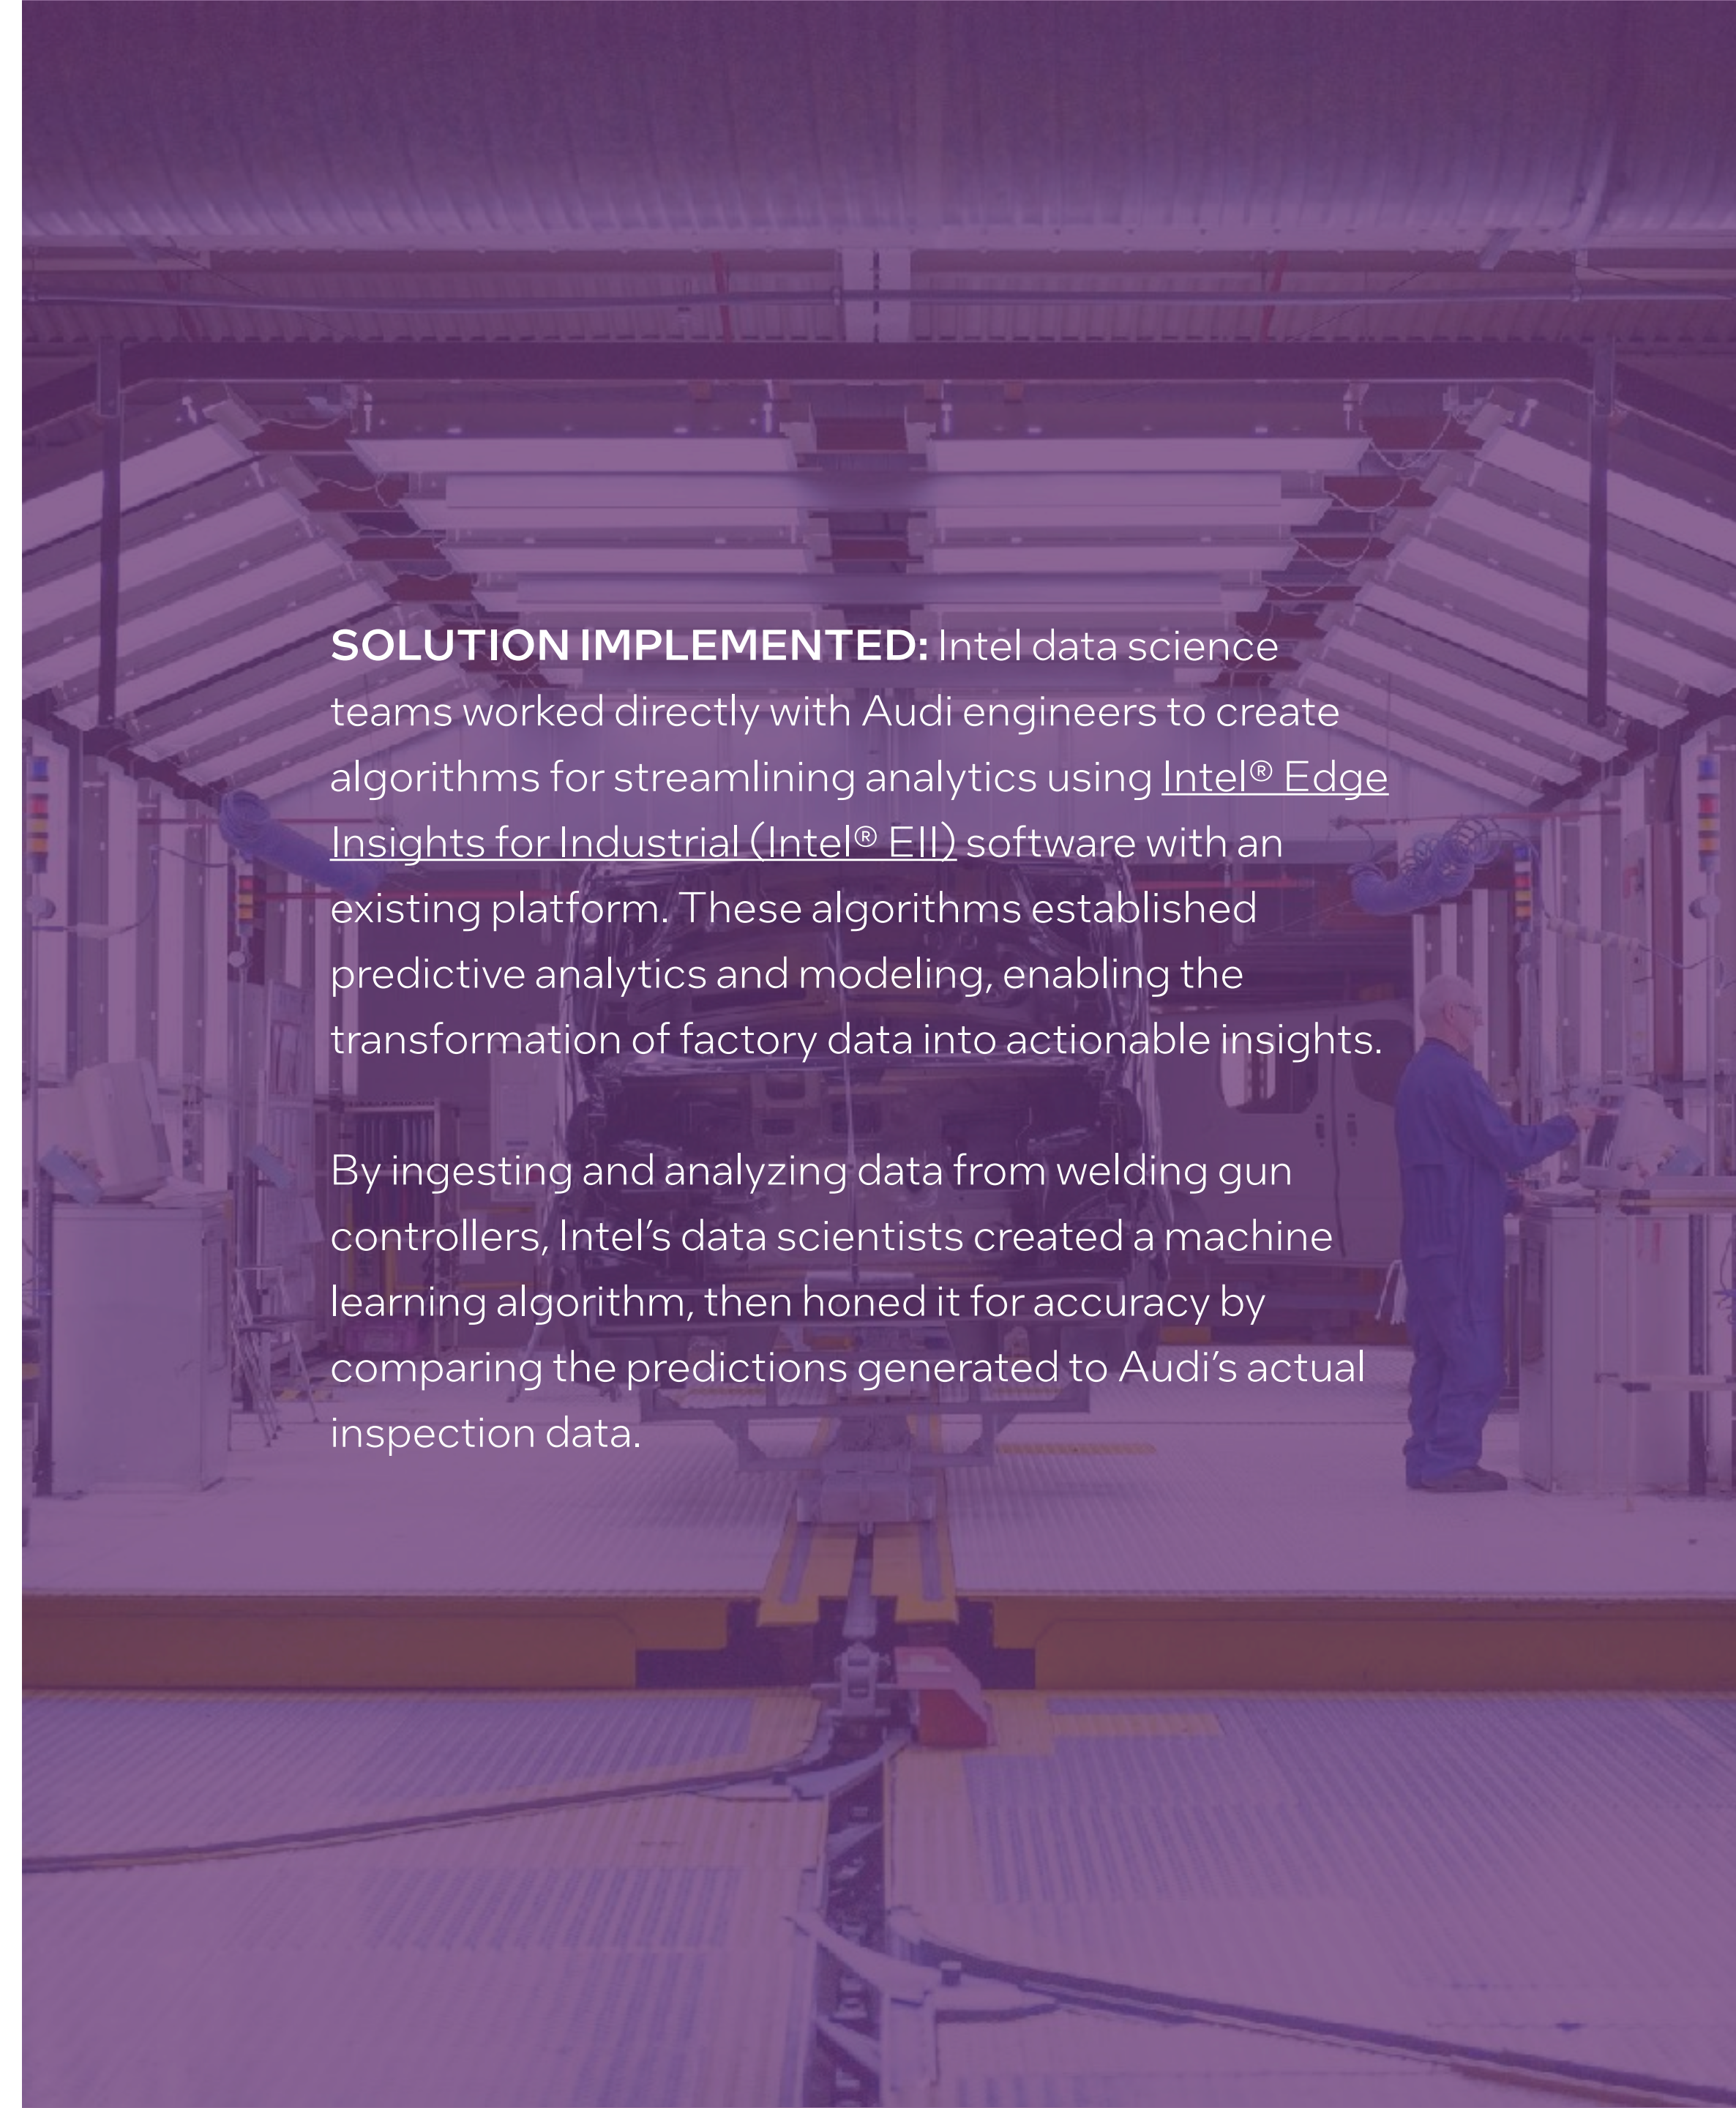
“Intel taught us how to understand the data, how to use the algorithms to analyze data at the edge, and how we can work with data in the future to improve our operations on the factory floor.”⁴

– Mathias Mayer, leader of automation technology planning, Audi



PROBLEM FACED: At the Audi factory in Neckarsulm, a city in southern Germany, they assemble 1,000 vehicles daily, each of which has 5,000 individual welds for a total of 5 million welds in a single day.

As the standard and historical industry protocol dictates, quality control inspections were done manually by pulling a single vehicle of the 1,000 made each day off the production line for 18 engineers to inspect using ultrasound probes. This procedure was costly and labor-intensive, and only the problems with that single vehicle would be correctable, leaving unanswered questions about the quality of the other 999 cars produced each day. As a result, Audi's challenge was to develop a feasible, cost-effective way to inspect 5 million welds on 1,000 vehicles each day.



SOLUTION IMPLEMENTED: Intel data science teams worked directly with Audi engineers to create algorithms for streamlining analytics using [Intel® Edge Insights for Industrial \(Intel® EII\)](#) software with an existing platform. These algorithms established predictive analytics and modeling, enabling the transformation of factory data into actionable insights.

By ingesting and analyzing data from welding gun controllers, Intel's data scientists created a machine learning algorithm, then honed it for accuracy by comparing the predictions generated to Audi's actual inspection data.

RESULTS: By creating a data-driven platform solution powered by Intel and Nebbiolo Technologies, Audi implemented a solution to meet their goal to inspect 100% of the 5 million welds they make each day with a very high degree of accuracy.⁵

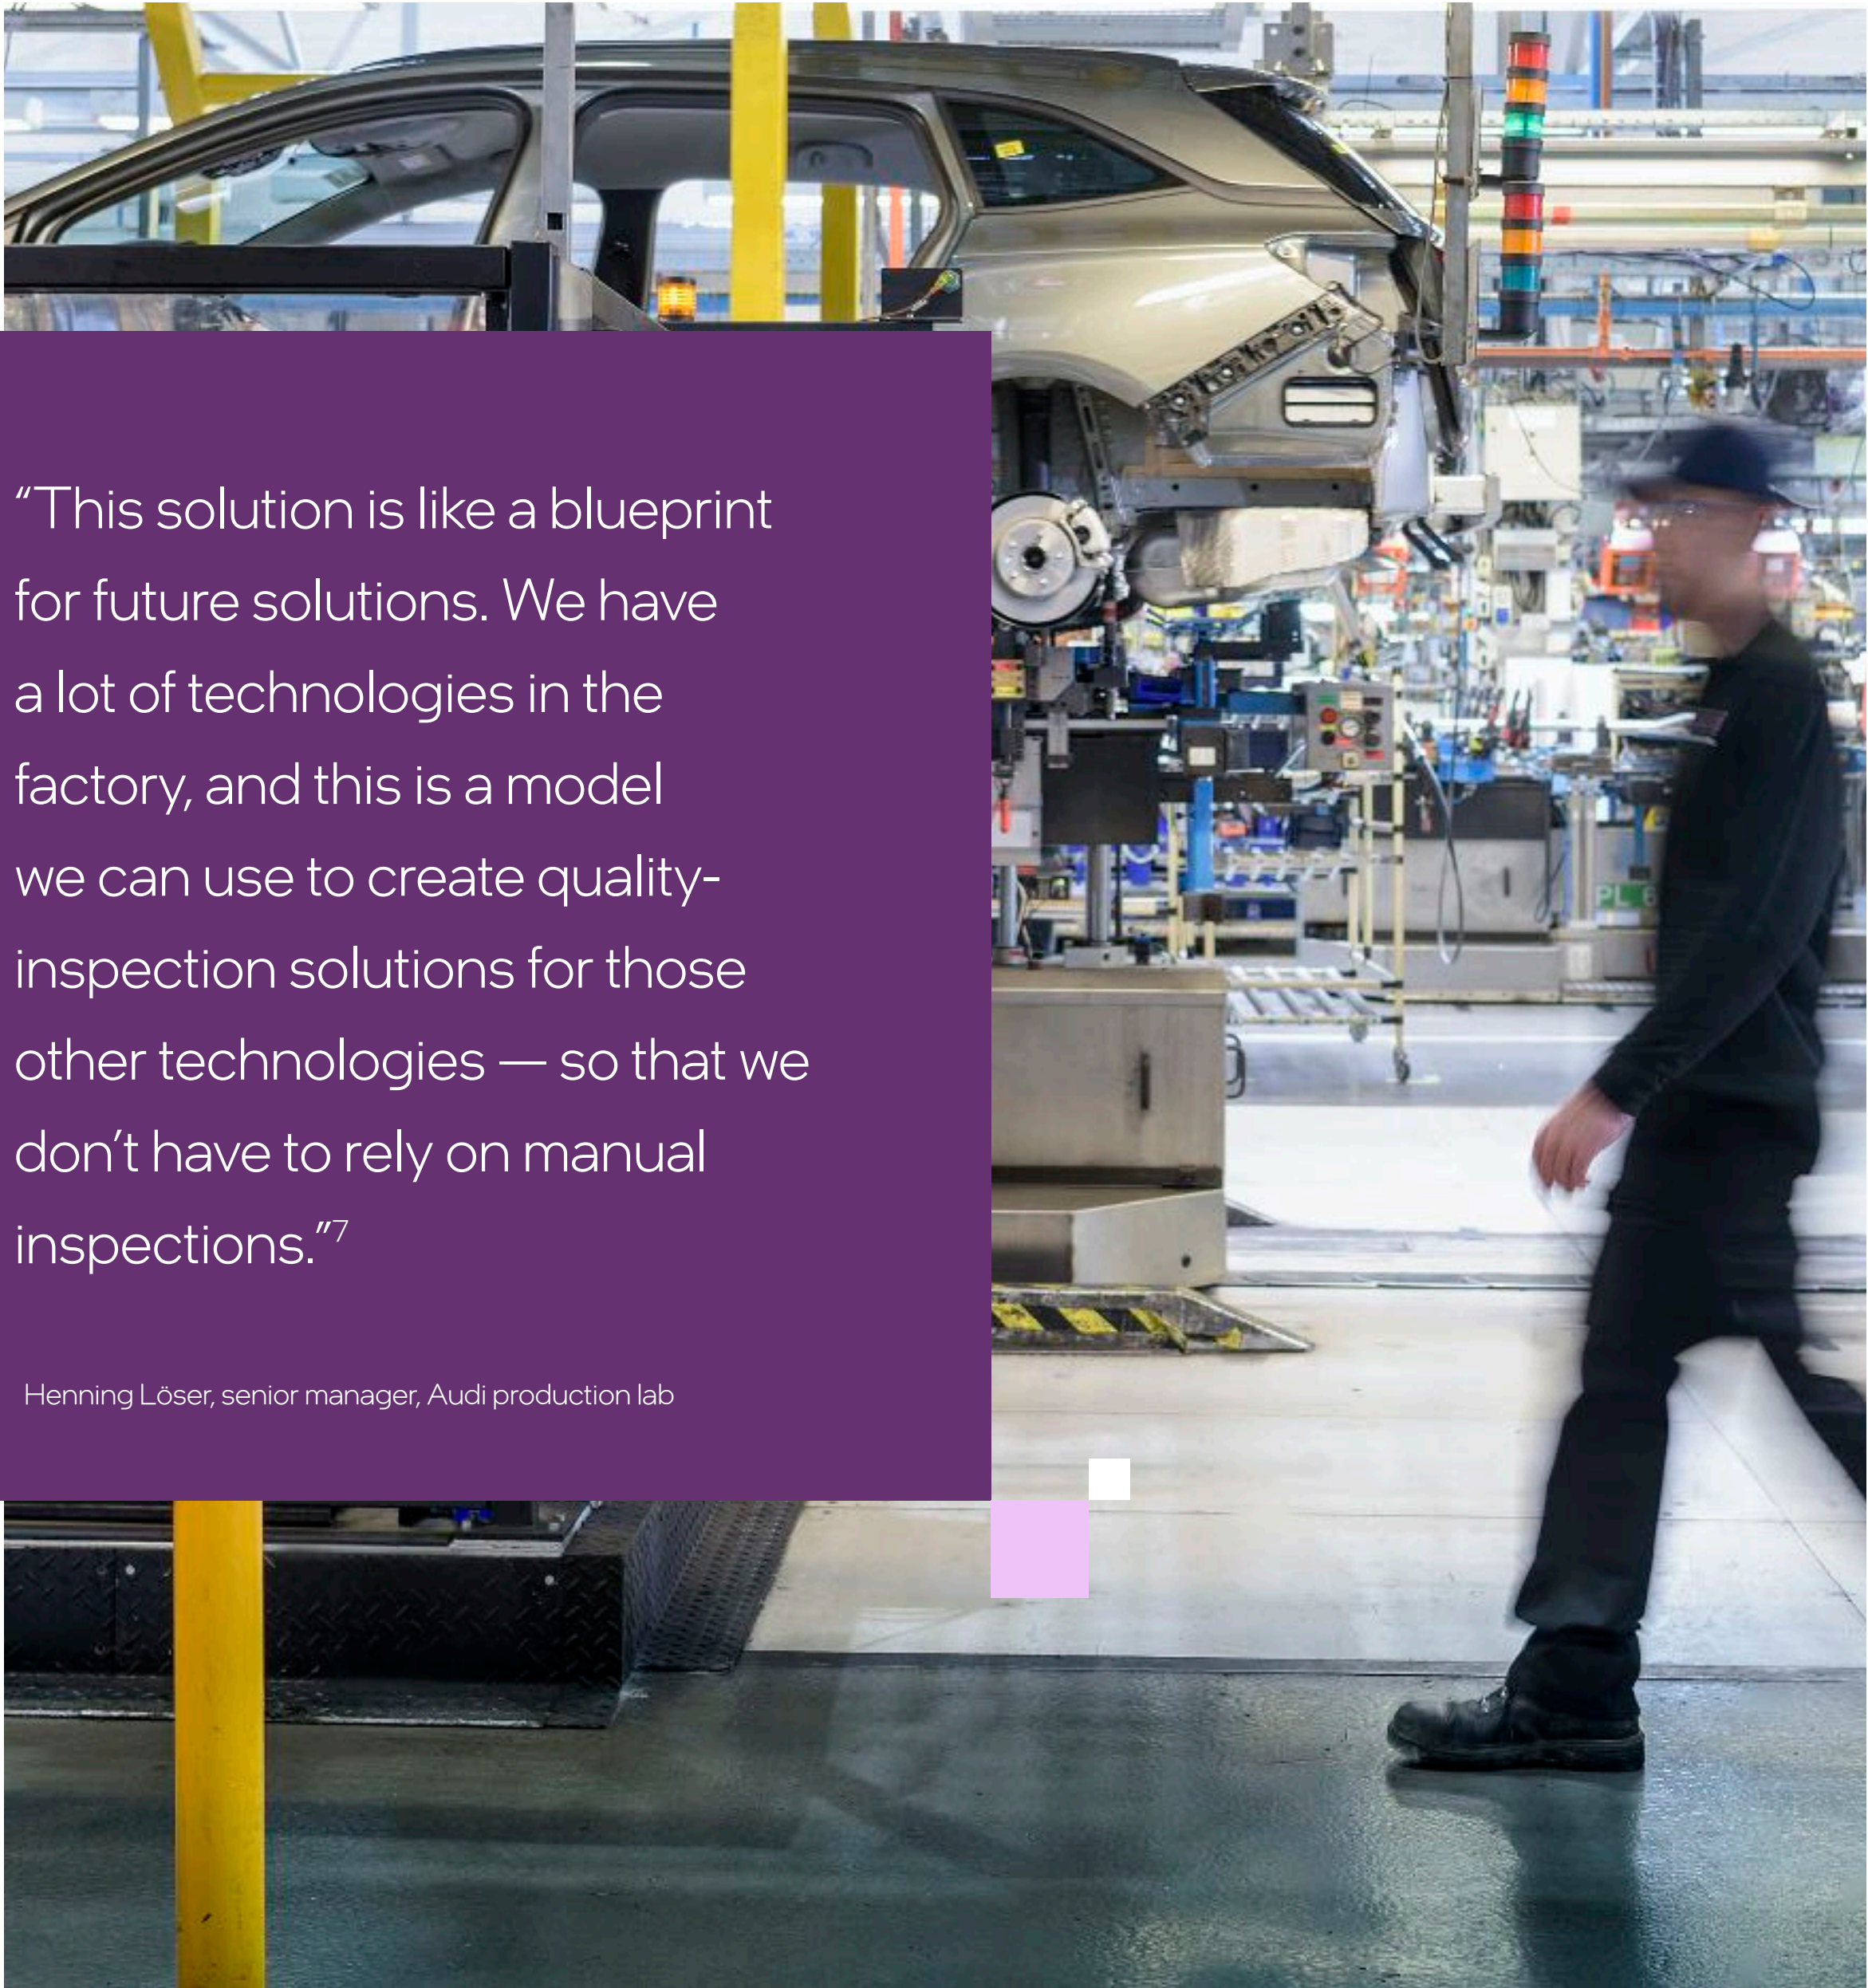
A dashboard helps Audi employees visualize data, while the system alerts technicians upon detecting a faulty weld, as well as potential configuration changes that could minimize or even eliminate faults entirely.

According to Michael Häffner, head of production planning, automation and digitization at Audi, “At the Neckarsulm factory, we are already seeing a 30%-50% reduction in labor costs.”⁶

Häffner emphasized that increasing automation and efficiency is not about replacing workers, but about giving them new knowledge and skills and creating new opportunities for them. It’s also a necessity, since many skilled factory workers are retiring and taking valuable knowledge with them. Automating some of those jobs and channeling younger employees in new directions is good for both the workers and business.

Intel helped create a scalable, flexible platform solution that not only improved quality control for spot welding, but also set the foundation for other use cases involving robots and controllers like riveting, gluing, and painting.

Find out more >



“This solution is like a blueprint for future solutions. We have a lot of technologies in the factory, and this is a model we can use to create quality-inspection solutions for those other technologies — so that we don’t have to rely on manual inspections.”⁷

Henning Löser, senior manager, Audi production lab

Other industrial use cases

Explore more >

Condition-based monitoring and predictive maintenance

Machines produce tremendous volumes of raw data that are cost and time prohibitive to send to central locations. Hyperconverged edge data centers perform analytics and filter data locally for faster time-to-asset condition insight. This opens the possibility of maintenance contracts based on actual machine health.

Defect detection using machine vision

Deliver the highest-quality product possible using machine vision to capture and analyze images to validate features and check for defects. With near real-time image processing, QA is streamlined, reducing downtime and defects while increasing production.

AI-based robotic process automation (RPA)

RPA systems perform repetitive and potentially hazardous tasks with greater speed and accuracy than humans. RPA systems read constant streams of machine data to flag anomalies, perform raw materials checks, and track production hours and materials used, requiring tremendous amounts of low-latency compute.

Precision monitoring and control

A key Industry 4.0 goal is to use data from multiple machines, processes, and systems to adapt the manufacturing process in real time; machine learning can help determine the best action. Edge computing collects, aggregates, and filters the data, then executes inference.



Augmented reality

Wearable and hand-held augmented reality (AR) and virtual reality (VR) devices can speed up on-the-job safety training, provide hands-free assembly instructions, display equipment stats like temperature and vibration to assist maintenance, and detect product faults during quality inspections. Edge computing enables near real-time, up-to-date information and alleviates on-device processing, making headsets or eyewear lighter and more user-friendly.

^{4,6,7} <https://www.intel.com/content/www/us/en/customer-spotlight/stories/audi-automated-factory.html>
⁵ <https://www.intel.com/content/dam/www/public/us/en/documents/case-studies/audis-automated-factory-closer-to-industry-case-study.pdf>

Sensormatic

by Johnson Controls

“As retailers face their new reality as a result of COVID-19, we believe that the industry needs to begin to restore consumer confidence, and that starts with helping to create a safer store environment.”⁸

– Subramanian Kunchithapatham, VP of engineering, Sensormatic Solutions

Sensormatic

by Johnson Controls

PROBLEM FACED: More than 20 years ago, ShopperTrak installed sensors throughout thousands of retail stores nationwide, quickly becoming the leader in shopper traffic monitoring and conversion rate analysis. The insights gleaned from such newly detailed and previously unobtainable data led to permanent changes in in-store design and staffing protocols to maximize efficiency.

Decades later, the emergence of the COVID-19 pandemic and the rapidly changing regulations affecting retail operations led ShopperTrak — now Sensormatic — to seek a more efficient system to monitor social distancing and occupancy rules. An automated system would reduce shopper anxiety while avoiding the cost and risk of using staff to implement the necessary changes.

SOLUTION IMPLEMENTED: Sensormatic had previously brought in Intel data scientists to create a video-based AI solution and were well into that project when the pandemic created a vital need to accelerate deployment time.

Intel leveraged existing sensor and camera platforms, combining them through custom-configured Intel® Vision products and Intel® processors designed to aggregate, analyze, and run AI models in individual stores. A smart hub at each location aggregates data, pairing them with software. In turn, each localized hub powered by [Intel® Movidius™ VPU](#)s and [Intel® Atom®](#), [Intel® Core™](#), and [Intel® Xeon® Scalable processors](#) is optimized for vision workloads with the [Intel® Distribution of OpenVINO™ Toolkit](#). This allows stores to use individualized dashboards that perform more complex analysis, including monitoring near real-time store insights and customer behavior predictions.

RESULTS: Once fully deployed, the new system created a long list of improvements and capabilities, including:

- Tracking social distance by recognizing differences between family groups in proximity versus close interactions between strangers*
- Near real-time notifications when a particular store reaches maximum capacity as defined by applicable local laws
- Providing near real-time staffing recommendations based on occupancy trends—like when to station employees at entrances to enforce crowd control
- Adding digital signage at entrances that show current capacity and provide the ability to deploy a “traffic light” to indicate when to enter safely
- In the near future, face mask compliance technology will notify staff when store rules or local laws are being violated

Counting people and social distancing is just the beginning. With an Intel processor-based, AI-ready smart hub in each location, businesses over time can take advantage of emerging applications for shopper engagement, shelf replenishment, quick service restaurants, exterior curbside wait time, and more. As the retail landscape continues to evolve and add even more AI capabilities, stores will continue to develop safer and smarter ways to shop.

Learn more >

“With Intel® Vision Portfolio, from Intel® Movidius™ Vision Processing Units to Intel® Core™ processors and Intel® Xeon® Scalable processors, and the Intel® Distribution of OpenVINO™ toolkit development tools, we have been able to tailor solutions that help equip retailers with the right technology along with near real-time insights to make business-critical decisions and ensure customers’ confidence in safety.”⁹

Subramanian Kunchithapatham, VP of engineering,
Sensormatic Solutions



*Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel's [Global Human Rights Principles](#). Intel's products and software are intended only to be used in applications that do not cause or contribute to a violation of any internationally recognized human right.

Other retail edge use cases

[Explore more >](#)

[Predictive inventory and supply chain control](#)

With its ability to read codes, text, and numbers, machine vision is widely used for identification across a range of industries, including automotive, consumer goods, electronics, packaging, and pharmaceuticals. Machine vision can help manage, track, and analyze inventory levels, ensuring critical materials are in the hands of those who need them most.

[Interactive displays for custom experiences and intelligence](#)

Today's digital signs and interactive kiosks are connected, intelligent, and responsive, making experiences more convenient and personalized. Smart sensors and cameras make it possible for kiosks to recognize products, respond to

touchless gestures, or even address loss prevention. Vision capabilities for AI-powered analytics let retailers know when a message is truly effective.

[Autonomous stores for contactless convenience and safety](#)

Data flowing from supermarket shelves carries important context such as inventory, traffic flow, purchase frequency, and linger time. This data, combined with sensors, computer vision, and robotics, is driving increasingly automated checkouts and "just-walk-out" purchase experiences.



[In-store loss prevention with video analytics and AI](#)

Footage of shoppers is streamed to on-premises edge gateways, where machine learning can quickly identify potentially criminal behavior. This can help store security immediately stop shoplifting, label switching, and other criminal activity.

*Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel's [Global Human Rights Principles](#). Intel's products and software are intended only to be used in applications that do not cause or contribute to a violation of any internationally recognized human right.

^{8,9}<https://www.intel.com/content/www/us/en/corporate-responsibility/covid-19-response-sensormatic-article.html>

PHILIPS

“Intel® Xeon® Scalable processors appear to be the right solution for this type of AI workload. Our customers can use their existing hardware to its maximum potential, while still aiming to achieve quality output resolution at exceptional speeds.”¹⁰

– Vijayananda J., data and AI, chief architect and fellow, Philips

PHILIPS

PROBLEM FACED: Modern medicine relies on medical imaging to accurately diagnose and treat ailments throughout the human body. The only prominent hardware solution to accelerate the deep learning required to advance medical discovery has long been graphics processing units (GPUs). GPUs work well with images by design, but they also have inherent memory constraints that data scientists have had to work around when building some models.

SOLUTION IMPLEMENTED: Intel and Philips partnered to accelerate deep learning inference on central processing units (CPUs) in key medical imaging use cases. CPUs, in this case [Intel® Xeon® Scalable processors](#) don't have the same memory constraints as GPUs and can accelerate complex hybrid workloads, including larger, memory-intensive models typically found in medical imaging. For many AI workloads, Intel Xeon processors offer an affordable, flexible platform for AI models and can better meet data scientists' needs than GPU-based systems, utilizing [Intel® Xeon® Scalable Processor](#), [Intel® Distribution of OpenVINO™ Toolkit](#), and [Intel® DevCloud](#).¹¹

The use of CPUs allows for rapid and deeper medical imaging, enabling AI techniques such as object detection and segmentation to help radiologists identify issues faster and more accurately using images as detailed as 1GB or even higher—none of which requires expanding capability via expensive new infrastructure. This enables Philips to offer AI solutions at a lower cost to its customers and can help with more effective prioritization of cases, better outcomes for more patients, and reduced costs for hospitals.¹²

RESULTS: Using Intel® Xeon® Scalable processors and the Intel® Distribution of OpenVINO™ toolkit, Philips worked with Intel to test two healthcare use cases for deep learning inference models: one on X-rays of bones for bone age prediction modeling, the other on CT scans of lungs for lung segmentation.

In these tests, the bone age prediction model went from an initial baseline test result of 1.42 images per second to a final tested rate of 267.1 images per second after optimizations—a remarkable 188 times improvement. In addition, there was an equally notable 38 times increase in speed for the lung segmentation model over the baseline measurements.^{13*} Such dramatic increases in diagnostic ability can be expected to result in a direct effect on overall positive patient outcomes.

Running healthcare deep learning workloads on CPU-based devices offers direct benefits. Companies like Philips process large images like those commonly found in radiology without slowing down workflow or impacting model accuracy, allowing them to offer AI-based services while maintaining costs for their end customers.



Learn more >



¹⁰ <https://newsroom.intel.com/news/intel-philips-accelerate-deep-learning-inference-cpus-key-medical-imaging-uses/>
^{11, 12, 13} <https://techdecoded.intel.io/resources/performance-ai-driven-medical-imaging-efficiently-and-cost-effectively-on-intel-cpu-based-systems/>
* Workload and system configuration appears in Appendix A in link. Results may vary.



“For Project: CORail, AI is empowering our engineers to achieve more and learn faster when it comes to growing the coral reef. It empowers the solution to gather data in a non-intrusive manner, allowing the scientists and data engineers to gather data from the reef with minimal disruption to this fragile ecology.”¹⁴

– Jason Mitchell, managing director/high tech industry, Accenture



PROBLEM FACED: Natural coral reefs are among the world’s most diverse ecosystems, and protecting them is vital to maintain the planet’s overall health. Per the United Nations Environment Program, not only do coral reefs provide income and food for over a billion people, protect coastlines from tropical storms, and generate nearly \$10 billion in tourism and recreation, but coral reefs are also critical to survival for thousands of animal and plant species.¹⁵

Unfortunately, due to global climate change and the planet’s resultant warming, overfishing, bottom trawling, and poorly planned coastal development, many reefs have been damaged, and some destroyed. Because of the damage being done, careful monitoring of reefs throughout the world to measure changing pH in seawater, as well as other human-made effects, is now vital for climate change research.

In the past, coral reef monitoring required human divers to capture photos, videos, and data that had to be analyzed on the surface. These traditional methods carry a unique risk that can affect overall accuracy. For example, divers’ presence can influence wildlife behavior, which can skew results, not to mention that a typical diver can only do 30-minute individual sessions, several times a day, without putting themselves at physical risk.



SOLUTION IMPLEMENTED: Accenture, Intel, and the Sulubaaï Environmental Foundation announced Project CORaiL, an artificial intelligence (AI)-powered solution to monitor, characterize and analyze coral reef resilience. The Accenture Applied Intelligence Video Analytics Services Platform (VASP) uses smart cameras and video analytics powered by Intel® Xeon® Scalable processors, Intel® FPGA Programmable Acceleration Cards, Intel® Movidius™ VPU, and the Intel® Distribution of OpenVINO™ toolkit.

Launched in 2019, this pilot project monitors the reef surrounding Pangatalan Island in the Philippines and has allowed researchers to use the 40,000+ images collected to measure the effects of climate change on the reef in near real-time without affecting wildlife in any significant way.

Patrick Dorsey, VP of product marketing, programmable solutions group, Intel, reports that “VASP uses AI to count and classify the marine life, with the data then sent to a surface dashboard, where it provides analytics and trends to researchers in near real-time, enabling them to make data-driven decisions to protect the coral reef.”¹⁶

RESULTS: Utilizing AI to solve the data collection problem and analyze the information gathered, engineers and oceanographers have begun restoring and supplementing the damaged reefs in the Philippines. They built a [Sulu-Reef Prosthesis](#), a concrete underwater platform to support coral fragments on the edge of collapse, allowing the coral time to recover and regenerate before suffering further damage. Once the overall positive preliminary results can be established as accurate in the longer term, participants are hopeful that the results can be replicated to repair damaged reefs worldwide.

Learn more >

“Project: CORaiL is a great example of how edge computing can augment researchers by having an always-on solution to capture data about the coral reef and marine life. Edge analytics allow researchers, doctors, farmers, retailers, and others to understand and act upon large amounts of visual and other unstructured data quickly instead of having to manually interpret and identify trends.”¹⁷

Patrick Dorsey, VP of product marketing, programmable solutions group, Intel

ViewSonic®

“Engagement is the premise of hybrid learning. How do we engage students, and how do we keep them engaged? So, we’re introducing ways of exploring things that weren’t possible before. You can grab a video, you can grab a 3D object, you can get kids participating remotely in the lesson.”¹⁸

– Craig Scott, CTO, ViewSonic



PROBLEM FACED: As the COVID-19 pandemic increasingly changed all aspects of life in spring 2020, thousands of schools across the U.S. had no choice but to begin the unprecedented experiment of distance learning—full time.

With so few tools available to allow for an adaptable learning experience, educators were forced to use whatever solutions were available, no matter how inefficient or cumbersome. And as the summer began and preparations for the upcoming school year started, no one could truly predict what would happen — whether schools would resume in-person education or be forced to continue online.

SOLUTION IMPLEMENTED: ViewSonic decided to create a distance learning initiative, a solution dubbed myViewBoard® classroom, and asked Intel for help with accelerating computational capability. myViewBoard™ classroom needed to feature a unique ability: allowing local educators to shift their entire programs back and forth among in-person, virtual, or a hybrid model, depending on local conditions and regulations.

They focused on building a flexible, equal parts teacher- and student-centric platform that put learning via engagement first. By working with Intel to create the ViewSonic ViewBoard, powered by Intel Unite® and myViewBoard™ classroom with Intel® Distribution of OpenVINO™ toolkit, ViewSonic enabled adaptable solutions and delivered in-person whiteboarding capabilities with the ability to prepare and present from anywhere.

RESULTS: ViewSonic teamed with Intel to implement AI and deep learning solutions as part of the myViewBoard™ software ecosystem, utilizing the [Intel® Distribution of OpenVINO™ toolkit](#) to enable facial expression recognition and mood indexing.* This innovative system enables teachers to gauge a student's mood, interest, and understanding of lessons at a glance, then adjust methodology to help alleviate confusion, frustration, and boredom.

Simultaneously, ViewBoard interactive displays deliver dynamic interactive capabilities. Teachers can engage students by imparting lessons to the whiteboard, using live annotation, and adding interesting videos and pictures directly on the device.

myViewBoard then connects the whiteboard experience to students engaged in distance learning, allowing teachers the flexibility to teach in-person and online simultaneously. Students, regardless of physical location, can collaborate online by working on a shared canvas, gaining the ability to learn from anywhere via lessons that can be recorded and shared on-demand.

Learn more >



"We can make the engagement experience much better for the teacher and students. myViewBoard collects data anonymously in a classroom and then offers real-time insights to teachers so they can make targeted adjustments — ultimately making the learning experience even better. A lot more things can happen as long as we keep having a conversation between the education community and technology community to talk about those problems and let us show them what's possible."¹⁹

Wei Oania, education vertical GM, Intel Internet of Things group

*Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel's [Global Human Rights Principles](#). Intel's products and software are intended only to be used in applications that do not cause or contribute to a violation of any internationally recognized human right.

^{18, 19} <https://www.intel.com/content/www/us/en/customer-spotlight/stories/brentwood-academy-customer-spotlight.html>

Rakuten Mobile

“The innovations developed by Intel and our other technology partners have enabled us to deploy a full, end-to-end virtualized mobile network that dramatically reduces our reliance on dedicated hardware and legacy infrastructure. As a result, we are able to reach new cloud-like operational efficiency and can pass our cost-savings onto our customers.”²⁰

– Tareq Amin, CTO, Rakuten Mobile

Rakuten Mobile

PROBLEM FACED: Rakuten Mobile wanted to build a new cloud with cellular connectivity rather than retrofit a cloud to existing cellular connectivity. Its CTO knew precisely what he needed — a better way to upgrade the network to 5G. Any solution would have to: use the existing spectrum; work for current and future customers; and integrate with cellular content, ecommerce financial services, gaming, shopping, and other cloud-based applications, all in a service-oriented way.

SOLUTION IMPLEMENTED: Rakuten Mobile is building the world's first fully virtualized, end-to-end, cloud-native mobile network.²¹ By harnessing proven cloud technologies and underpinning its network with Intel technologies like [FlexRAN](#), [Intel® Xeon® Scalable processors](#), [Field Programmable Gate Arrays \(FPGAs\)](#), and the [OpenNESS toolkit](#), Rakuten has a powerful foundation to deliver immersive experiences and support a variety of edge computing applications. Placing servers at the edge means very rapid response times and almost no latency in applications and rich media content. These edge servers running on Intel technology are powerful enough to provide Rakuten Mobile's customers with immersive experiences, such as HD video streams of live sporting events and other entertainment while supporting developers and companies creating new, previously unimagined edge applications.

RESULTS: Rakuten Mobile was able to achieve cloud-based automation and cloud-like operational efficiency, eliminating reliance on previously utilized hardware and legacy infrastructure. The new mobile network was launched entirely on the cloud with end-to-end automation and programmable infrastructure that helped Rakuten minimize operating expenditure and the labor required to deploy and operate the network. It also removes issues caused by human error and allows faster response to problems. All of this enables Rakuten to ensure stability, scalability, and agility while achieving economies of scale that keep the capital and operating costs of the network efficient so they can pass their cost savings onto its customers.

Learn more >

“People talk about a flexible, virtualized network, but Rakuten actually did it. When you are a pioneer, you are going to run into challenges. The learnings and experiences that Rakuten Mobile has had will benefit the entire industry as we go into the 5G world.”²²

Caroline Chan, VP and general manager, Intel Network Business Incubator Division



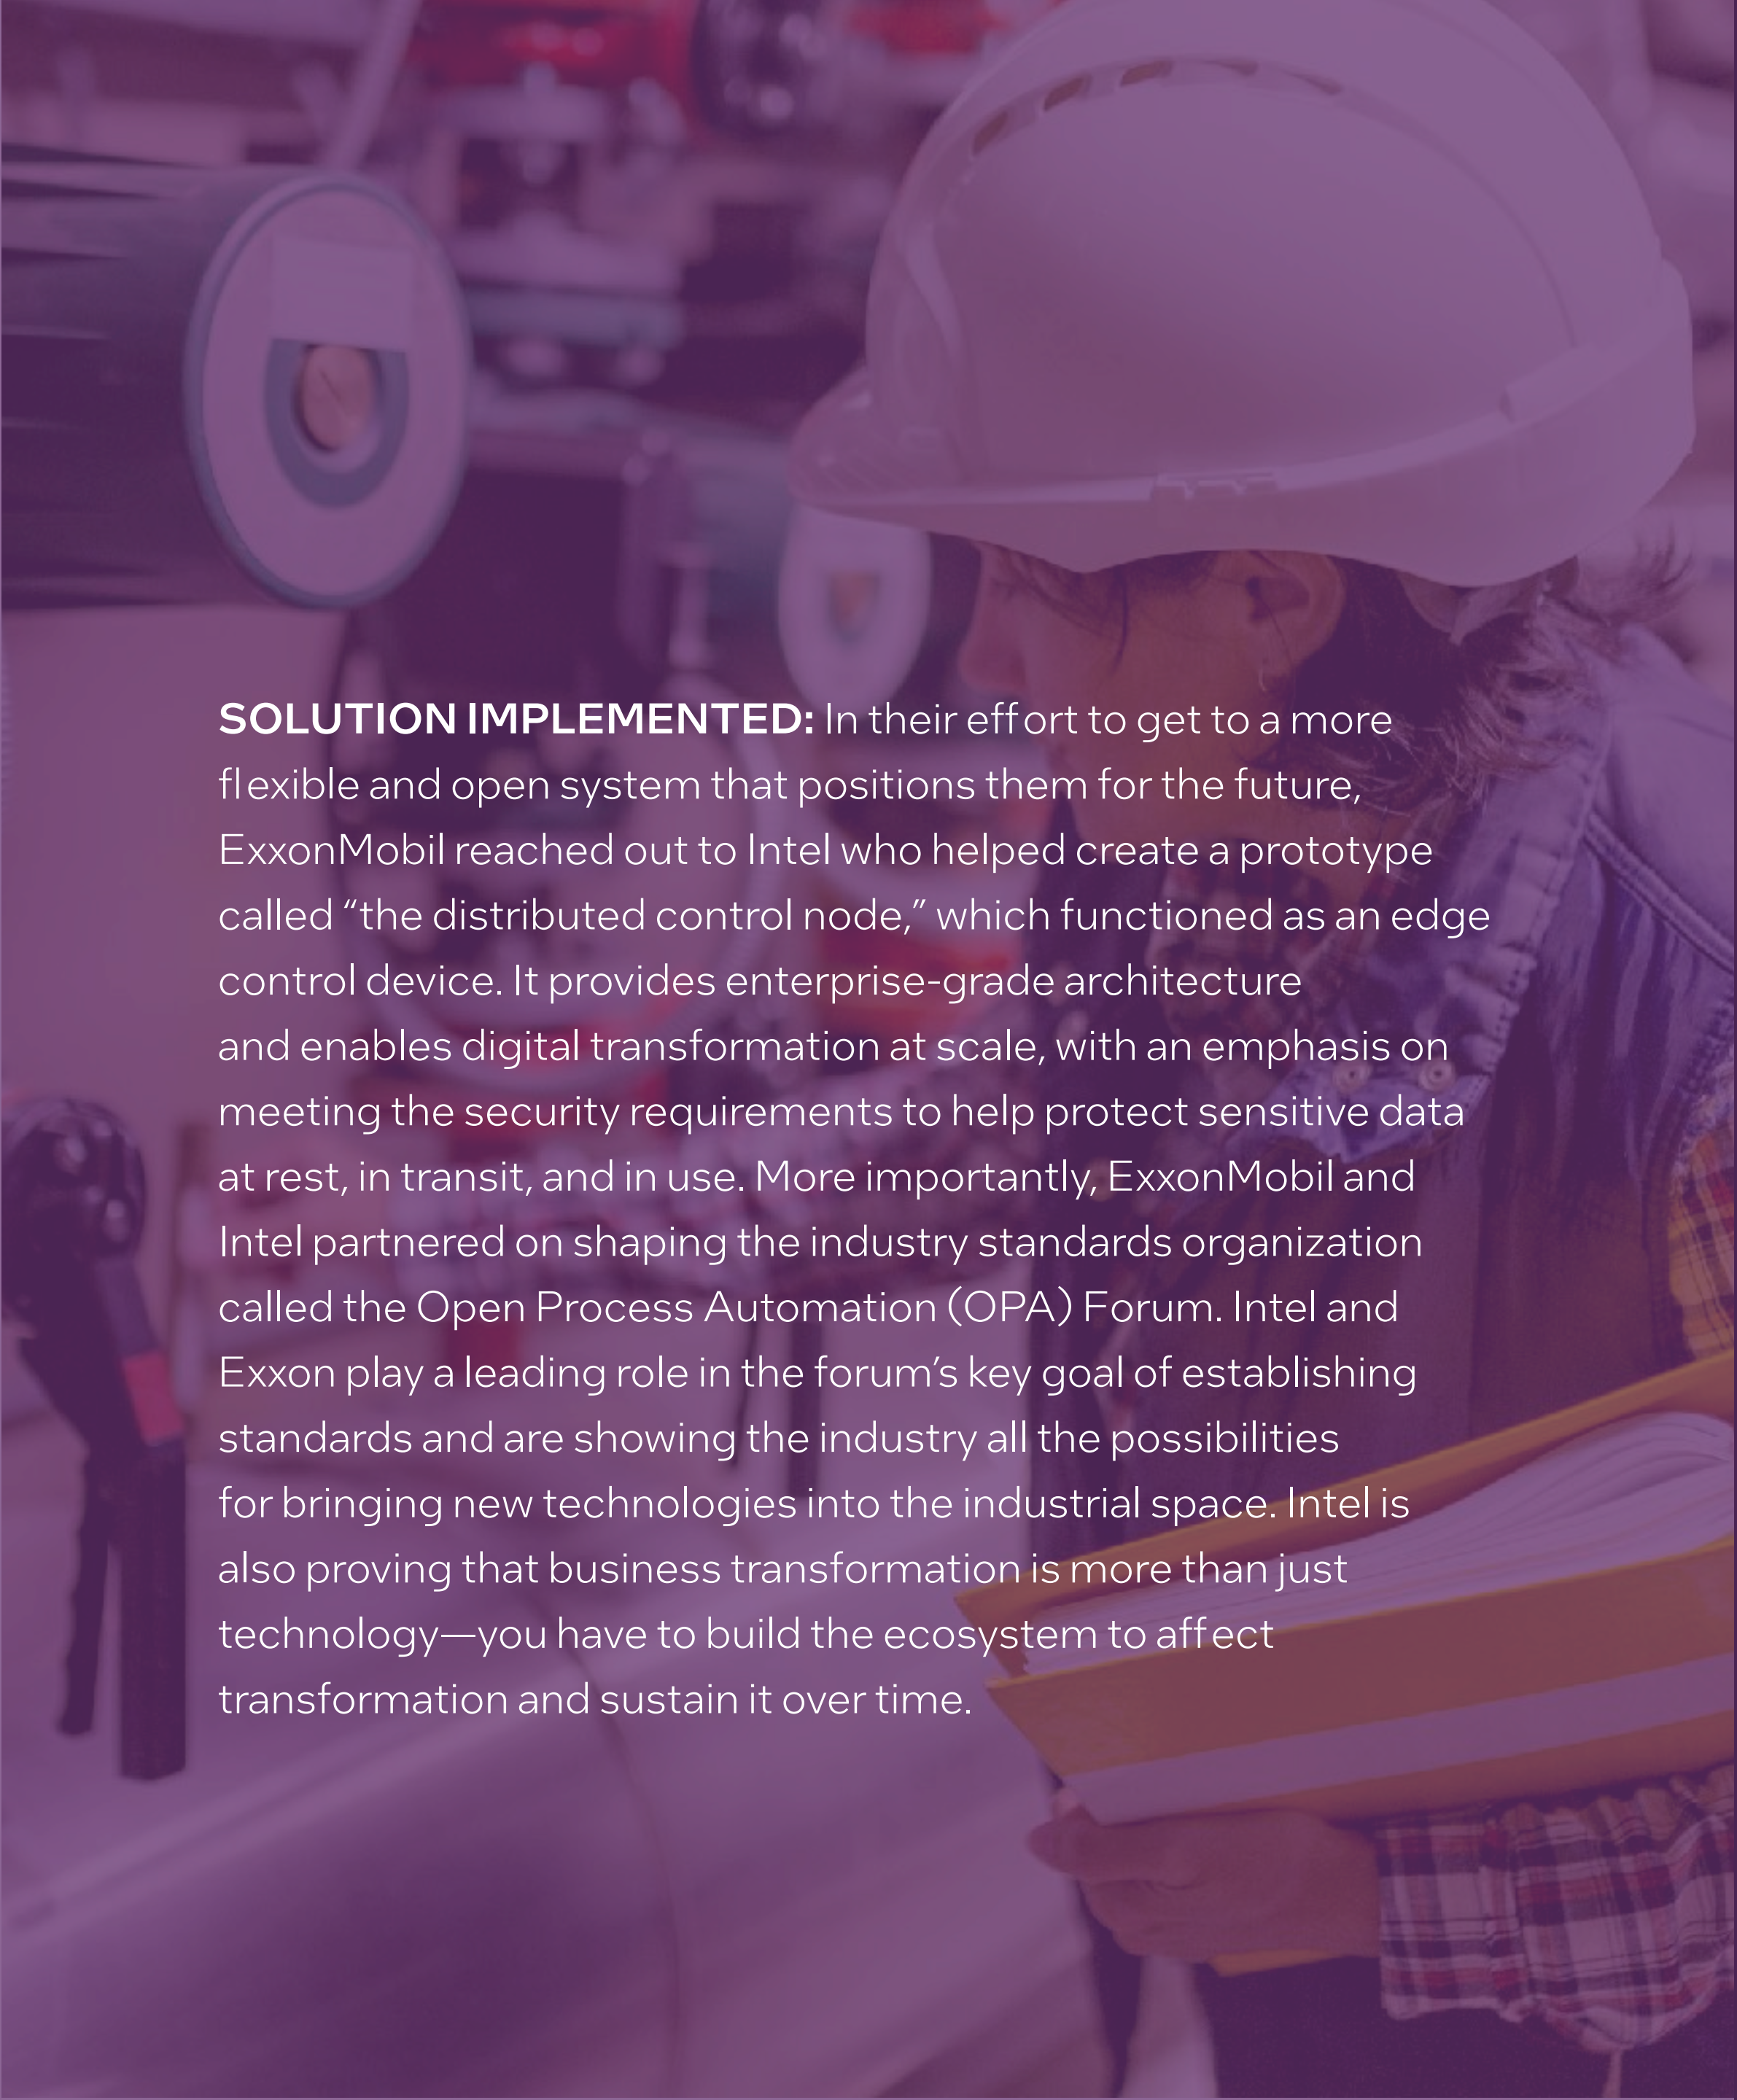
"In an industrial control context, cybersecurity is a very real concern for us. We are talking about risk to life, risk to our neighbors who live near our manufacturing facilities, and we take that responsibility very seriously. So, as we pursue these new technologies, it's always done with the constraints and the realities that we have to do it in a secure way."²³

– Don Bartusiak, chief engineer, process control, ExxonMobil



PROBLEM FACED: ExxonMobil, like many in the industry, has a large base of proprietary control systems at their manufacturing facilities, refineries and chemical plants worldwide, which are decades old. They knew that they needed to update these industrial control systems to be more cost-efficient and prepare for future needs, but the high cost of replacement and limited capacity to access the current technology was a challenge.

Exxon recognized that they would need a more flexible, open, and interoperable system to allow for innovation and value capture. Equally important was the need to address modern cybersecurity threats and align the industry on standards that support interoperability and portability so that it's easier and more cost-effective to move from closed proprietary systems to open standards based systems. Since innovation is something that organizations are trying to capture and harness as part of their transformation, ExxonMobil and others in the industry had to find ways to overcome these barriers.




SOLUTION IMPLEMENTED: In their effort to get to a more flexible and open system that positions them for the future, ExxonMobil reached out to Intel who helped create a prototype called “the distributed control node,” which functioned as an edge control device. It provides enterprise-grade architecture and enables digital transformation at scale, with an emphasis on meeting the security requirements to help protect sensitive data at rest, in transit, and in use. More importantly, ExxonMobil and Intel partnered on shaping the industry standards organization called the Open Process Automation (OPA) Forum. Intel and Exxon play a leading role in the forum’s key goal of establishing standards and are showing the industry all the possibilities for bringing new technologies into the industrial space. Intel is also proving that business transformation is more than just technology—you have to build the ecosystem to affect transformation and sustain it over time.

RESULTS: ExxonMobil knew that the industry at large shared similar pain points, so they reached out across different industry sectors that use industrial control systems. Today, the OPA Forum has become a breakthrough program for a multitude of industries that utilize industrial control systems as it lays the groundwork for future collaboration and innovation based on rapidly changing digital capabilities.

Learn more >



ExxonMobil, with their partnership with Intel, are committed to creating a competitive advantage for manufacturing facilities through the use of automation and process control technologies. Having access to open process automation systems is foundational to future automation activities and is a major investment in the future of your company.



²³ <https://www.intel.com/content/www/us/en/customer-spotlight/stories/exxonmobil-customer-spotlight.html>



“We’re seeing an increasing demand for computing capabilities at the edge of the network, where workloads move closer to digital interactions between people, things, and business models, and where the latency, bandwidth, and geography around those workloads are all critical considerations.”²⁴

– David Shachochis, VP of enterprise technology and field CTO, Lumen



previously CenturyLink

PROBLEM FACED: Lumen saw increasing demand for high-performance applications on its network at the edge, where latency, bandwidth and geography all require new computing capabilities. Emerging business models for retail, robotics, adaptive manufacturing, and AI-powered video analytics require critical considerations for the convergence of IT and OT.


SOLUTION IMPLEMENTED: Intel helped Lumen push its private cloud to the edge with network-optimized Intel® Xeon® Scalable processors and Intel® Optane™ SSDs. This meant deploying computing at several locations, including on-premises (targeting zero latency), central office edge (targeting one millisecond of latency), content delivery network and regional edge (targeting five milliseconds of latency), as well as the traditional private cloud. The combination helped achieve the necessary **performance to run emerging applications for retail robotics, adaptive manufacturing, and AI-powered video analytics, while keeping within its power footprint requirements.**

RESULTS: This collaboration with Intel allowed Lumen to enable customers to tap into the power of the software-defined data center and robust edge networks. In the future, edge computing is expected to work in tandem with a range of wireless and wireline networks to deliver solutions that build enterprise business value around various workloads. These include, but aren't restricted to, retail robotics, adaptive manufacturing, AI-fueled video analytics, low-latency transactions, and more.

Learn more >

“By working closely with Intel on the rollout of our edge services, we were able to leverage network-optimized Intel® Xeon® Scalable processors to provide virtual network functions within the same socket as the processors we used to run standard enterprise workloads. This allows us to be more flexible with our edge customers’ needs by customizing performance while keeping within our power footprint requirements. It also supports our model of deploying a wide array of enterprise workloads on a single standard implementation, which helps us manage costs while maximizing performance.”²⁵

Kevin McBride, principal architect for edge cloud, Lumen




^{24, 25} Quote is from the [video](https://newsroom.intel.com/video-archive/video-centurylink-and-the-rise-of-edge-services/) featured at <https://newsroom.intel.com/video-archive/video-centurylink-and-the-rise-of-edge-services/>



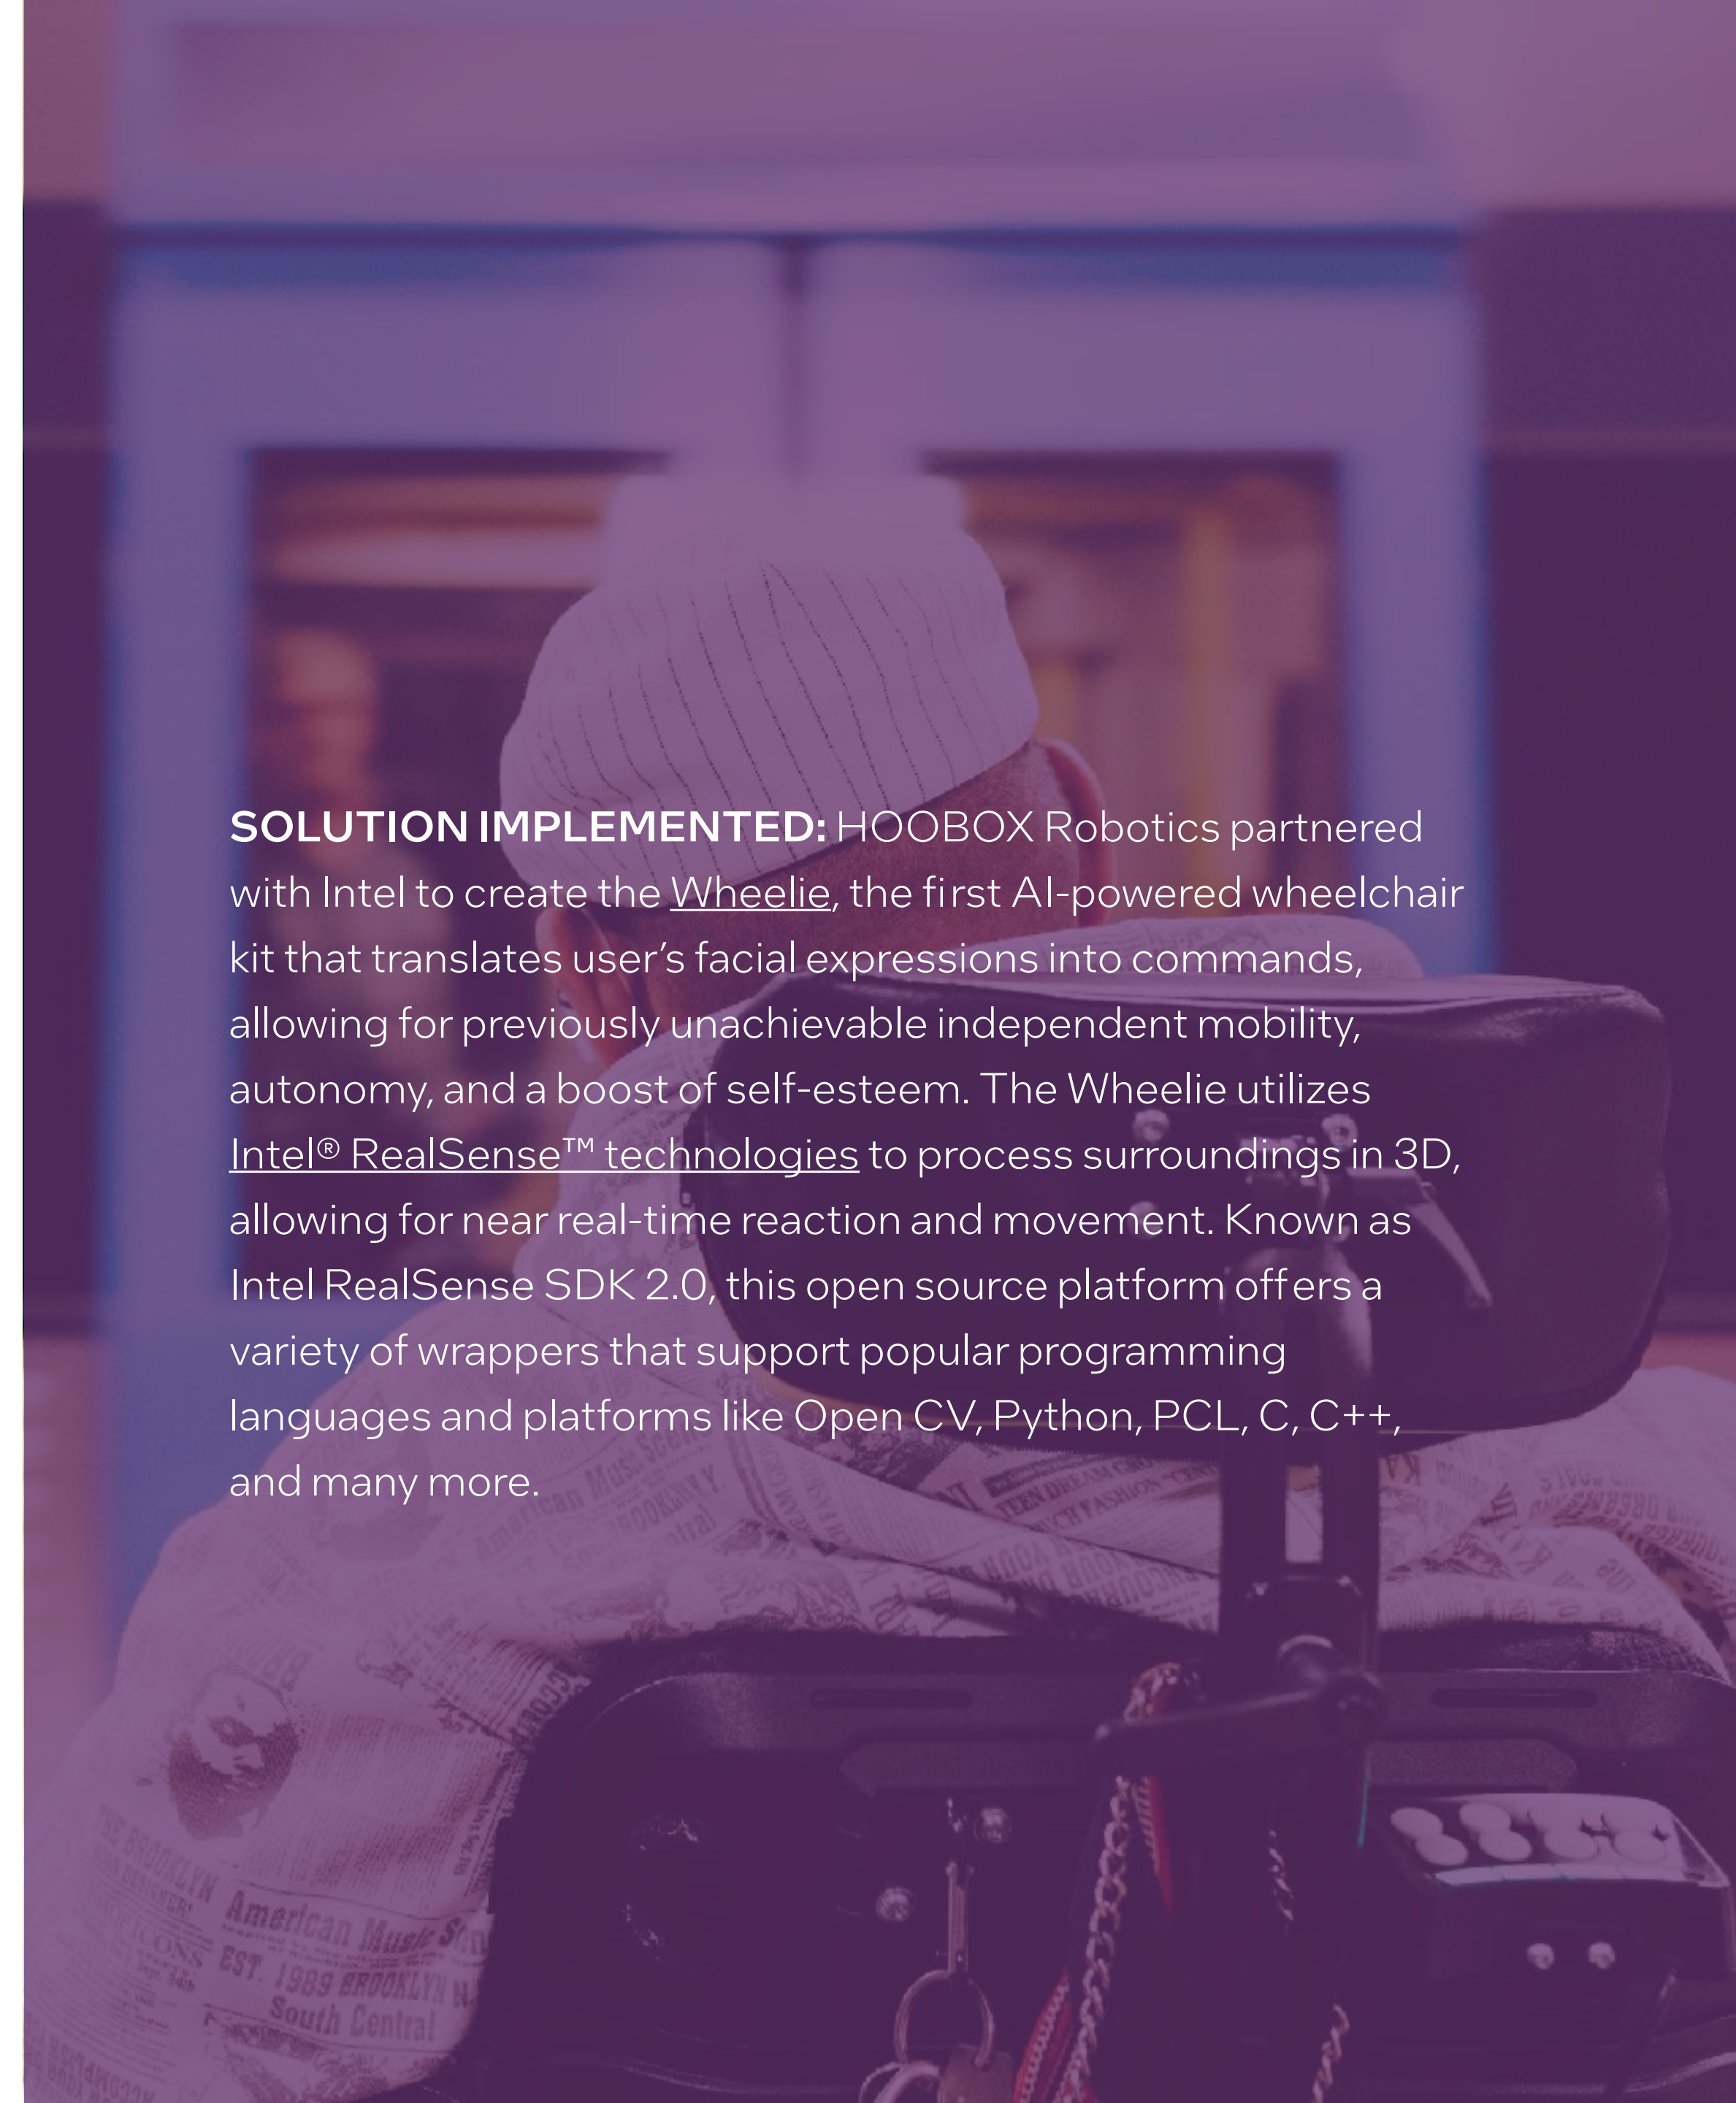
“Without the help of Intel® AI, we could not build this kind of high-precision technology.”²⁶

– Dr. Paulo Gurgel Pinheiro, co-founder & CEO, HOOBOX Robotics



PROBLEM FACED: Up to 500,000 people suffer spinal cord injuries annually, hindering not only their mobility, but also their autonomy in general.²⁷ When an injury restricts overall movement, day-to-day life becomes exponentially more difficult and, inevitably, quality of life suffers.

Existing wheelchair design runs the gamut from very basic to technologically advanced, but all suffer from the inability to allow quadriplegic users to control mobility in near real-time.



SOLUTION IMPLEMENTED: HOOBOX Robotics partnered with Intel to create the Wheelie, the first AI-powered wheelchair kit that translates user's facial expressions into commands, allowing for previously unachievable independent mobility, autonomy, and a boost of self-esteem. The Wheelie utilizes Intel® RealSense™ technologies to process surroundings in 3D, allowing for near real-time reaction and movement. Known as Intel RealSense SDK 2.0, this open source platform offers a variety of wrappers that support popular programming languages and platforms like Open CV, Python, PCL, C, C++, and many more.

RESULTS: Utilizing Intel technology, HOOBOX created a wheelchair guidance system to meet the unique needs of quadriplegics, giving them a complete autonomous mobility system. [Intel® RealSense™](#) cameras capture and map 11 facial expressions with 99.9% accuracy to generate chair commands.²⁸ Running on the Intel NUC, the [Intel® Distribution of OpenVINO™ toolkit's](#) facial recognition* generates 18.3 times faster performance than previously available solutions, allowing users to generate near real-time wheelchair response.²⁹ The kit can be installed on any motorized wheelchair in just seven minutes, giving users mobility almost immediately.³⁰

Learn more >

“Together, Intel AI technology and the Intel AI Builders program transform prototype algorithms into production-ready optimized solutions that solve large-scale, real-world problems. This is a game-changing opportunity for AI startups and the AI community at large.”³¹

Paulo Pinheiro, HOOBOX CEO

*Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel's [Global Human Rights Principles](#). Intel's products and software are intended only to be used in applications that do not cause or contribute to a violation of any internationally recognized human right.

^{26, 28, 30, 31} <https://www.intel.com/content/www/us/en/artificial-intelligence/hooobox.html>. Workload and system configuration appears in link. Results may vary.

²⁷ <https://www.who.int/news-room/fact-sheets/detail/spinal-cord-injury>

²⁹ <https://www.intel.com/content/www/us/en/artificial-intelligence/solutions/computer-vision-solutions-from-hooobox-enable-independent-mobility-for-quadruplegics.html>



“Virtualizing the entire network from the core to the edge has been a massive, multi-year redesign effort of our network architecture that simplifies and modernizes our entire network.”³²

– Adam Koeppe, senior vice president of technology and planning, Verizon



PROBLEM FACED: Remaining a leader in the ultra-competitive telecommunications market means quickly adapting to 5G standards. This requires pushing compute to the edge, closer to where data is created and consumed, for both new monetization opportunities and operational savings. Service providers like Verizon look to partners like Intel to help architect their RAN infrastructure for a diverse range of 5G and edge services. Intel helped Verizon as it deployed the world's first successful end-to-end virtualized 5G data session in a commercial network, emphasizing the readiness of virtualized radio access network (vRAN) technology.



SOLUTION IMPLEMENTED: Verizon worked with partners like Intel as they deployed the world's first successful end-to-end fully virtualized 5G data session in a commercial network, emphasizing the readiness of virtualized radio access network (vRAN) technology.³³ They were able to virtualize the RAN, thereby decoupling software functionality and enabling the network on general-purpose hardware, leading to greater flexibility and agility during the introduction of new products and services.

Moreover, the move to a cloud-native, container-based virtualized architecture with standardized interfaces can improve service delivery speeds, scalability and cost-efficiency in networks.

Verizon utilized [Intel® Xeon® Scalable processor](#) for Layer 1 and 2 processing, an [Intel® FPGA Programmable Acceleration Card \(Intel® FPGA PAC\) N3000](#) for workload acceleration and fronthaul processing, and an [Intel® Ethernet Network Adapter XXV710](#) for network I/O connectivity. In addition, Intel's FlexRAN software reference architecture helped accelerate the virtualization development.

RESULTS: To maintain its place as a leader in [multi-access edge computing \(MEC\)](#), Verizon’s RAN virtualization is poised to lower the barrier to entry for new vendors in the ecosystem, allowing them to accelerate innovation, reduce operating costs, and lay the groundwork for flexible network and cloud infrastructure closer to the customer, eventually leading to single-digit latency. Key 5G use cases focused on providing the best, most efficient network for customers that heavily rely on the programmability of virtualized networks.

These kinds of advancements in virtualization technology will enable Verizon to improve its deployments of massive-scale IoT solutions. Verizon will maintain its industry leadership into the future with more robust consumer devices, solutions, AR/VR, remote healthcare, autonomous robotics in manufacturing, and smart city solutions.

Learn more >



“Fully virtualized 5G networks can only be delivered through strong industry collaboration, and this recent achievement is the culmination of years of work between Verizon, Intel, Samsung and Wind River. As our customers implement virtualized radio access network architectures, Intel is bringing together the products, ecosystem and expertise to help accelerate commercial solutions into the marketplace.”³⁴

Dan Rodriguez, corporate VP and general manager of the Network Platforms Group, Intel

³² <http://www.globenewswire.com/news-release/2020/08/25/2083452/0/en/Verizon-paves-way-for-wide-scale-5G-mobile-edge-compute-with-first-successful-end-to-end-fully-virtualized-5G-data-session-in-the-world.html>

³³ <https://www.verizon.com/about/news/verizon-fully-virtualized-5g-data-session>
³⁴ <https://newsroom.intel.com/articles/intel-key-role-worlds-first-successful-end-to-end-virtualized-data-session/#gs.o2dj5v>



1.0

Introduction

03-06

[Edge Computing](#)

[Your Data Stranded in the Wild](#)

[Extend Beyond Traditional Data Centers](#)

[Success at the Edge](#)

2.0

Real-world Examples

08-40

[Audi](#)

[Philips](#)

[ViewSonic](#)

[ExxonMobil](#)

[HOOBOX](#)

[Sensormatic](#)

[Accenture](#)

[Rakuten Mobile](#)

[Lumen](#)

[Verizon](#)

3.0

Key Considerations for Edge Deployments

42-53

[Hardware for Varied Data Types](#)

[Capitalize on Easy-to-use Software Tools](#)

[Ready-to-Deploy Solutions](#)

[Implement AI at the Edge](#)

[Hardware-based Security](#)

4.0

Discovering Solutions

55-58

[Expert Support From a Trusted Partner Ecosystem](#)

[Scalable, Repeatable Solutions](#)

[Your Fast Path to Edge Development](#)

Key Considerations for Scalable and Secure Edge Deployments

The complexities of network design, deployment management, performance trade-offs, and total cost of ownership (TCO) issues present challenges for both network and infrastructure architects. Teams must balance existing investment and architectural strategies with a plan toward modernization, taking into consideration all aspects of edge computing, networks, and the cloud, which may often stretch years into the future.

Coordinating business operations with the needs of information technologists (IT), operations technologists (OT), and application developers is vital to ensuring the best business outcomes.


So where do you start?

Start with hardware for networking, storing, and processing varied data types.

The cloud-centric model that dominates much of enterprise technology today requires companies to reevaluate their infrastructure because of the increased volume of data and new data processing requirements. While some companies need to upgrade their existing infrastructures, others need to add new capabilities as they move to the edge so they can capture and store more data, analyze it faster, and act on it sooner.

Given the diverse nature of edge workloads, companies should consider various data types and processing requirements as well as performance, latency, and efficiency when designing their infrastructure. Intel gives architects the flexibility to choose from a range of CPUs, ASICs, VPU, and FPGAs to enable complete edge systems and ensure an ideal fit for virtually any project.

With many options based on flexible Intel Xeon processors featuring built-in AI, **Intel offers complete, full-stack, edge-ready technology** to meet broad application performance, latency, and efficiency needs.



Intel technology gives you the ability to build—and evolve—a resilient edge.

Choosing the most efficient and economical solution depends on your unique business challenges. The right solution must meet requirements for workload performance, size or power constraints, and budget. Intel offers a range of computing options to enable IoT solutions from the edge to the cloud that help your enterprise realize the best return on your investments.

Intel CPUs are the foundation of the edge

[Intel® Xeon® Scalable processors](#) deliver the high performance and reliability needed to accelerate compute, storage, memory, and networking in remote and on-premises data centers. [Intel® Core™ processors](#) feature a range of enhancements, including precision time controls, integrated remote management, functional safety, and long life to support critical near real-time edge applications like manufacturing robotics, remote healthcare, and retail transformation.*

Computer vision and deep learning

[The Intel® Movidius™ Vision Processing Unit \(VPU\)](#) is a dedicated hardware accelerator for deep neural network inference, enabling demanding computer vision and edge AI workloads in areas such as visual retail, security and safety, and industrial automation.

Customize devices to accelerate key workloads

[Intel® FPGAs \(Field Programmable Gate Arrays\)](#) are powerful integrated circuits that can be reconfigured in the field time and time again. Our FPGAs offer particular

value in high-throughput, low-latency applications and support a number of usage models such as security, I/O, and networking.

Rigorous in-network edge acceleration

Blend off-the-shelf network-optimized CPUs with a wide range of in-network accelerators for specific cost, performance, and time-to-market needs. Accelerators range from custom-built ASICs to [structured Intel® eASIC™](#) devices to fully programmable [Intel® FPGA SmartNIC platforms](#).

Agility within the data center to deliver services efficiently and cost-effectively

With [Intel® Ethernet Network Adapters](#), innovations like Hardware-Enhanced Precision Protocol (PTP) and Application Device Queues (ADQ) enhance application response time predictability, helping replace other dedicated hardware and save on overall solution costs. Workloads requiring customizable solutions can benefit from networks built with [Intel® Programmable Ethernet Switch Products](#).

Store more and access data faster

[Intel® Optane™ technology](#) is a new class of memory and storage designed to break through data bottlenecks. This revolutionary innovation bridges critical gaps in the storage and memory hierarchy, delivering persistent memory, large memory pools, fast caching, and fast storage.

[Intel® Optane™ persistent memory \(Intel® Optane™ PMem\)](#) provides dense, extremely responsive, highly predictable data access for new business models that require low-latency AI inference at the edge, so data can be accessed, processed, and analyzed in near real-time to deliver deep insights and improve operations.


[Intel® Optane™ Solid State Drives \(SSDs\)](#) fill the gap between high-performing volatile memory and lower-performing NAND storage and hard disk drives (HDDs) to help remove data bottlenecks, accelerate transactions, and improve time to insights.

* Not all features are available on all SKUs.

Capitalize on open, containerized, easy-to-use software and tools.

Distributed edge solutions can be complex to build and manage, so developers look for ways to streamline their workflows and speed deployment with a focus on open standards that support containerized and cloud-native development.

Intel's software packages and offerings help simplify optimizations for edge solutions, including computer vision and deep learning applications. Whether you are building from scratch or customizing a ready-made reference implementation, we offer helpful resources to accelerate your development.



Intel's software packages, tools, and resources help solve real development challenges at the edge.

Bring down your barriers to edge intelligence

With the [Intel® Edge Software Hub](#), software developers can customize, validate, and deploy use case-specific solutions faster and with greater confidence. It gives you quick access to a growing list of containerized, vertical-specific software packages and easy-to-use developer tools that help you overcome the complex requirements for edge use cases and enable you to rapidly develop deployable edge solutions.

Get software reference designs for top use cases

Edge Insights are free software reference designs and tools from our Software Hub that simplify edge AI solutions and accelerate top edge use cases for sensing, analytics, machine vision, automation, and more. The [Intel® Edge Insights for Industrial](#) overview can help you learn about our pre-validated, ready-to-deploy software reference design for video and time series data ingestion, and the [Intel® Edge Controls for Industrial](#) article describes a range of IT-style technologies like containerization, virtualization, and orchestration.

Once you're ready to take the next step towards deployment, be sure to check out our collection of recipes: [Edge Insights for Industrial](#), [Edge Controls for](#)

[Industrial](#), [Edge Insights for Vision](#), and [Converged Edge Insights](#) for 5G and next-gen networks.

Get fast AI application performance feedback


[Intel® DevCloud for the Edge](#) provides you with speedy feedback via a virtual AI prototyping tool, helping you reduce time and cost as you compare hardware options for optimal performance.

Speed private edge network deployment

Built on open standards, APIs, and containers, [Intel® Smart Edge](#) is a commercial out-of-the-box multi-access edge computing (MEC) platform for private LTE and 5G networks. It allows enterprises to run their private mobility (4G/5G) applications and third-party virtual network functions (VNFs) on the premises' edge in a highly secure and scalable way.

Target and remove data transport bottlenecks

See what every packet is doing, down to the nanosecond with [Intel® Deep Insight Network Analytics Software](#). This software harnesses in-band network telemetry data from switches, vSwitches, and NICs to monitor, identify, and resolve network bottlenecks for better flow.



Tap into open source software tools and libraries:

Deploy high-performance deep learning (DL) Inference

The [Intel® Distribution of OpenVINO™ toolkit](#) extends deployment of high-performance deep learning inference across multiple configurations of Intel® hardware solutions (including accelerators) used to solve a variety of tasks, including computer vision, speech recognition, and natural language processing.

Deliver enhanced visual experiences faster

[Open Visual Cloud](#) is a collection of open source stacks and pipelines built with optimized ingredients to encode, decode, inference and render. This creates a reusable developer environment that eases testing, evaluation, and deployment of services, including video on demand (VOD) and livestreaming with SVT-AV1.

Deploy cloud-native network edges faster


[OpenNESS \(Open Network Edge Services Software\)](#) is open source software that simplifies the complex orchestration and management of edge services across diverse network platforms and access technologies.



Get to market quickly and affordably, with ready-to-deploy solutions and support tailored to your industry.

Streamlining integration and deployment complexity saves both labor and overall investment, though it is dramatically easier said than done. Intel has been working closely with partners in the edge ecosystem for years. The company has become a trusted expert in delivering scalable and repeatable solutions to help reduce the time, cost, and risk of deployments.

With tens of thousands of edge deployments and counting, we’ve aligned common use cases, helped fix technical complexities, and vetted hundreds of repeatable production-ready packages, enabling any company to deploy edge applications quickly and easily.



Minimize integration headaches and bridge development gaps with an experienced ecosystem of providers and ready-to-deploy solutions.

With more than **1,200 solution providers** delivering more than **300 test- and deployment-ready packages**, we're clearing a simpler path to your business edge.

Ready-to-deploy AI and IoT solutions to help you deploy edge applications with speed

- [Intel® IoT Market Ready Solutions](#) – Deploy IoT today with commercially ready Intel market-ready solutions.
- [Intel® IoT RFP Ready Kits](#) – Accelerate time to market and reduce complexity with integrated market-specific solutions that have been deployed in the field and are both scalable and repeatable.
- [Intel® Select Solutions](#) – Optimize analytics, AI, network, and media deployments with rigorously benchmark-tested and verified solutions optimized on Intel® Xeon® Scalable processors.
- [Intel's Converged Edge Reference Architecture \(CERA\)](#) – CERA unifies and converges IoT with wireless infrastructure network technology to simplify workload convergence at the edge.

Ecosystem programs to help you get to market faster

Deep technical collaboration, expertise, training, tools, and ready-to-deploy solutions await you in our multiple AI, network, and edge-specific ecosystem programs.

- [Intel® AI Builders](#) – Harness AI quickly with over 300 leading global AI software, hardware, and service providers with more than 150 solutions across diverse use cases and markets.
- [Intel® AI: In Production](#) – Access partner solutions focused on AI at the edge and help reduce development and collaboration challenges.
- [Intel® Network Builders](#) – Get access to content and resources to help you leverage the capabilities of Intel products.
- [Intel® Partner Alliance](#) – Accelerate your growth to give you the edge in building and selling the most innovative solutions and services.


Implement artificial intelligence (AI) at the edge and unlock new business value.

Businesses across all industries are finding new ways to harness data by bringing AI closer to where the data is created — in network edge devices, on-premise servers and gateways, and smart and connected endpoints. In almost all cases, these solutions have similar goals: increased efficiency, improved decision-making, added value for customers, and employee safety.

But deploying the infrastructure needed to support AI at scale remains a challenge, given the complexity of achieving high performance and accuracy while managing costs and utilization.

With inference increasingly happening outside of data centers, recapturing latency and monetizing untapped data by implementing AI closer to the data source is becoming more critical. That's why Intel has designed a range of products for AI computing, from workload-focused CPU and GPU products to programmable components and deep learning accelerators that are optimized for inference, enabling powerful vision, speech recognition, audio processing, and more at the edge.

With extensive experience and proven use cases in nearly every industry — including defect detection, worker and public safety, robotics, supply chain management, imaging diagnostics, and enhanced service delivery — Intel offers edge technology solutions that enable businesses to seamlessly build and deploy AI applications at scale.



Whether you are building a smart camera or managing thousands of autonomous robots, Intel helps bring intelligence to your edge*.

Gain new insights to your business

[Intel® Edge Insights for Industrial](#) brings the ease of AI deployment in industrial environments through capabilities such as data ingestion, processing, and transmission. These optimized edge analytics deployments can help lead to improved industrial product quality, operational performance, prediction of downtime, and automated operational flow.

Take advantage of deep learning edge compute

Start with [Intel® Xeon® Scalable processors](#), the foundation for deep learning inferencing with AI-enhanced capabilities integrated in the silicon. Increase the power of your AI workloads with [Intel® GPUs](#). Improve performance in machine and deep learning applications where low latency is critical with [Intel® FPGAs](#). Enable computer

vision and edge AI workloads on intelligent cameras and edge servers with [Intel® Movidius™ Vision Processing Units \(VPUs\)](#).

Achieve faster deep-learning inference from edge to cloud

The [Intel® Distribution of OpenVINO™ toolkit](#) optimizes model performance across all our edge to cloud hardware to help you accelerate AI workloads, including high-performance computer vision, deep learning inference applications, audio, speech, language, and recommendation systems.

Harness AI faster with an ecosystem of industry partners

Join [Intel® AI Builders](#) to get access to technical enablement resources, co-marketing and matchmaking opportunities, and optimized market-ready AI partner solutions.

Accelerate your AI deployments with Intel® Select Solutions for AI

Accelerate AI inferencing and deployment on an optimized, verified infrastructure based on industry-standard, general-purpose Intel® hardware and technology.

Reduce development challenges with Intel® AI: In Production

Access Intel® technologies, software tools, development kits, code samples, and solutions to help our ecosystem to help accelerate time to production.

[Use AI to improve network performance and network automation](#)

Intel is working with partners and customers across a number of industry ecosystems to solve network optimization and automation problems utilizing the power of AI.

A few examples where AI is applied in the network for improving operations include seeing this as closed-loop automation, Encrypted Traffic Analytics, and Radio Optimization for 5G beamforming.

Create end-to-end solutions for federated learning in medical imaging

Federated learning, a distributed machine learning approach that enables organizations to collaborate on deep learning projects without sharing sensitive data like patient records, is becoming increasingly important in healthcare and medical research. For example, [Intel worked with the University of Pennsylvania](#) to co-develop technology to enable a federation of 29 international healthcare and research institutions to train AI models that identify brain tumors using a privacy-preserving technique.

Help protect your data from edge to cloud with hardware-based security.

Cybersecurity has risen to become a critical element of the connected world, and security must be part of a cohesive defense strategy that extends beyond a single computing device, networking environment, software application, or individual. Some industries like healthcare, data security, safety, and privacy have redefined how companies implement solutions to keep proprietary, sensitive, and personal information safe across the computing spectrum.



Edge-to-cloud security, delivered.

At Intel, [security remains an ongoing priority](#) that guides everything we do to architect, build, and support our products. Our security design principles have led us to develop new capabilities like hardware-based acceleration of cryptography applications to help keep data private, trusted execution to help protect applications running on the top layers of software, and tamper-resistant storage designed to help prevent sensitive information from being deciphered and cloned. Keep reading to learn more about Intel’s ongoing efforts to support our customers with advanced security technologies.

[Intel® Software Guard Extensions \(Intel® SGX\)](#)

Intel® SGX helps protect your most sensitive workloads by using secure enclaves—isolated, nonaddressable, encrypted regions of memory created by the processor. To help keep your most important information private, confidential, and unmodified while it’s in use in memory, SGX authenticates that the data has not been compromised even if the OS, hypervisor, or other software is breached.

[Intel® Secure Device Onboard \(Intel® SDO\)](#)

“Hard-coded” manual provisioning methods to onboard IoT devices can be costly and unsecure. Intel® SDO is a neutral onboarding service with simplified ecosystem enablement to put devices into service faster.

[Intel® QuickAssist Technology \(Intel® QAT\)](#)

Intel® QAT accelerates and compresses cryptographic workloads by offloading the data to hardware, working in concert with standard OpenSSL protocols. This makes it easier for developers to integrate built-in cryptographic accelerators in network and security applications.

[Intel® Trusted Execution Technology \(Intel® TXT\)](#)

Intel® TXT hardens platforms to address the increasing and evolving security threats across physical and virtual infrastructure by complementing antivirus software.

1.0

Introduction

03-06

[Edge Computing](#)

[Your Data Stranded in the Wild](#)

[Extend Beyond Traditional Data Centers](#)

[Success at the Edge](#)

2.0

Real-world Examples

08-40

[Audi](#)

[Philips](#)

[ViewSonic](#)

[ExxonMobil](#)

[HOOBOX](#)

[Sensormatic](#)

[Accenture](#)

[Rakuten Mobile](#)

[Lumen](#)

[Verizon](#)

3.0

Key Considerations for Edge Deployments

42-53

[Hardware for Varied Data Types](#)

[Capitalize on Easy-to-use Software Tools](#)

[Ready-to-Deploy Solutions](#)

[Implement AI at the Edge](#)

[Hardware-based Security](#)

4.0

Discovering Solutions

55-58

[Expert Support From a Trusted Partner Ecosystem](#)

[Scalable, Repeatable Solutions](#)

[Your Fast Path to Edge Development](#)

4.0

Get up to speed with proven solutions and expert support **you** **can trust**

Edge computing solutions based on Intel have already powered tens of thousands of deployments, generating real value for businesses across all industries.

Are you ready to be next?

intel®




Expert support from a **trusted partner ecosystem.**

Thanks to Intel® technologies and a mature partner ecosystem, more secure, scalable, and repeatable solutions are available to help you reduce time, cost, and risks so you can get to market faster.

The Intel ecosystem is built on open standards-based solutions, allowing you to customize everything for your specific needs. Combined with Intel® IoT Market Ready Solutions, collaboration among machines, humans, and enterprise systems has never been easier.

- [Intel® IoT Market Ready Solutions \(Intel® IMRS\)](#) – The Intel IoT partner ecosystem offers solutions that are scalable, adaptable, vetted, and ready for immediate implementation.
- [Intel® Solutions Marketplace](#) – Connect with Intel® partners and discover products and solutions to solve your business needs.
- [Intel Customer Spotlight](#) – Explore the latest customer stories, case studies, and testimonials.
- [Select Solutions](#) – Simplify and accelerate infrastructure deployment with rigorously benchmark-tested and verified solutions optimized on Intel® Xeon® Scalable processors.



Scalable, repeatable solutions ready for deployment in nearly **every industry.**

Every industry has its challenges, but they all share the same goal: [Do more with data](#) in more places, do it faster to reduce operational costs and enhance productivity, and use the data to uncover new ways to compete.

We've spent years solving the complexities of IT and operational divides, validating edge software, and collaborating with our partners to offer data-intensive workload solutions that are optimized, adaptable, vetted, and ready for implementation across multiple industries.

Industrial

Put your data to work harnessing the power of industrial IoT, enabling digital transformation and Industry 4.0. Perform analytics in near real-time, consolidate workloads, optimize production, implement predictive maintenance, and enable automation. [Learn more >](#)

Retail

Intel's retail solutions utilize analytics and AI to help retailers of every kind deliver an experience that is more engaging, curated, and convenient for all parties. [See how >](#)

Healthcare

Intel uses technology to help improve patient experiences and outcomes while streamlining workloads and reducing costs for providers and researchers. [Explore more >](#)

Smart cities

Intel's IoT-based vision for the future of city planning uses AI, vision, and connectivity to put data to work, enabling environmental monitoring, smart lighting, mobility, public kiosks, and more. [Find out how >](#)

Education

Intel is utilizing new tools and techniques to create immersive learning environments that nurture the mindsets and skill sets students will need for jobs in the Fourth Industrial Revolution. [Discover more >](#)


Go to the Intel® IoT Market Ready Solutions (Intel® IMRS) to find offerings for your industry. [Find industry solutions >](#)

Your fast path to edge development.

Intel helps make experimenting, testing, and creating even the most complex edge solutions easier, with less work beforehand.

With a focus on open standards and support for containerized and cloud-native development, our software packages, hardware, and tools are designed to streamline workflows and speed deployment.

- [Intel® Developer Zone](#) – The official source for developing on Intel hardware and software.
- [Intel® Edge Software Hub](#) – Go faster at the edge with industry-specific packages, reference architectures, and Intel resources.
- [Intel Resource & Design Center](#) – Technical resources, documentation, tools, and support for designers, engineers, and developers.



Notices and Disclaimers

Performance varies by use, configuration, and other factors. Learn more at [www.Intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Intel technologies may require enabled hardware, software, or service activation.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Intel contributes to the development of benchmarks by participating in, sponsoring, and/or contributing technical support to various benchmarking groups, including the BenchmarkXPRT Development Community administered by Principled Technologies.

Statements in this document that refer to future plans or expectations are forward-looking statements. These statements are based on current expectations and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in such statements. For more information on the factors that could cause actual results to differ materially, see our most recent earnings release and SEC filings at www.intc.com.

Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel's [Global Human Rights Principles](#). Intel's products and software are intended only to be used in applications that do not cause or contribute to a violation of any internationally recognized human right.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

The image features a dense, low-angle view of server racks in a data center. The racks are filled with various electronic components, including circuit boards and connectors, many of which are highlighted in a bright yellow color. The perspective creates a sense of depth and complexity. Overlaid on this background is the Intel logo, which consists of a small blue square followed by the word "intel" in a white, lowercase, sans-serif font. A registered trademark symbol (®) is positioned to the upper right of the word "intel".

intel®