

Intel AI Solutions: AI Builders – A Thriving Enterprise Ecosystem

VERTICAL PARTNERS

RETAIL	HEALTHCARE	BFSI	TRANSPORTATION	MEDIA, TRAVEL & ENTERTAINMENT	SEC & GOVT	SW TOOLS & SERVICES	PROF. SERVICES	AGRI	TELECOM

HORIZONTAL PARTNERS

OEM				SYSTEM INTEGRATORS		
BUSINESS INTELLIGENCE & ANALYTICS	VISION	CONVERSATIONAL BOTS	NATURAL LANGUAGE	AI TOOLS & CONSULTING	AI PAAS	BIG DATA

For a complete list, visit builders.intel.com/ai/solutionscatalog



 Click on a vertical to jump to that section



**WIN
1-PAGERS**

HEALTH & LIFE SCIENCES

CLOUD SERVICE PROVIDERS

HIGH-PERFORMANCE COMPUTING (HPC)

FINANCIAL SERVICES INDUSTRY (FSI)

INDUSTRIAL AND MANUFACTURING

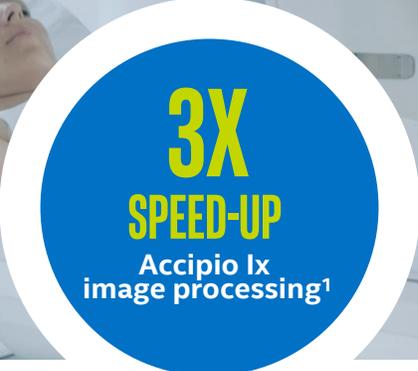
RETAIL

ENERGY

TRANSPORTATION AND LOGISTICS

OTHER





ACCIPIO doesn't replace radiologists, because we're much more than just pixel detectors. ACCIPIO is an adjunct to us as radiologists. Even for experts who are highly sensitive to spotting anomalies, it augments our skills. That makes life better for patients."

Ajay Choudhri, MD, System Chair of Radiology, Capital Health

Improving Stroke Care and More With Artificial Intelligence-Aided Medical Imaging

Capital Health is using MaxQ AI's ACCIPIO medical diagnostics solution, running on Artificial Intelligence (AI) technologies from Intel, to help patients with stroke, traumatic brain injury, or other causes of a brain hemorrhage get the care they need when they need it. The Intel® Distribution of OpenVINO™ toolkit and other Intel® software libraries were used to optimize the solution, which runs on Intel® Core™ i7 processors, for maximum performance and flexibility at the network edge. With Accipio's AI solution providing a triage of CT scans, treatment teams can proceed with greater confidence in their diagnoses. The solution can flag exams with subtle signs of intracranial hemorrhage and provide near real-time access to Accipio results, helping Capital Health to keep pace with their stroke center's commitment to read and report on CT scans within seven minutes.

Products and Solutions

[Intel® Core™ i7 processors](#)

[Intel® Distribution of OpenVINO™ toolkit](#)

Industry

Health and Life Sciences

Organization Size

1,001-5,000

Country

Italy

Partners

[MaxQ AI](#)

Learn more

[Case Study](#)

¹ For more complete information about performance and benchmark results, visit <https://www.intel.ai/case-studies/capitalhealth-maxqai/>





CorporateHealth International Improved Video Capsule Endoscopy with AI

CorporateHealth is a healthcare services company with the mission to bring medical device innovation to scale. Video capsule endoscopy is a technology that is used to take pictures of the small and large bowel, avoiding diagnostic colonoscopies – which are invasive and time consuming. CHI is pioneering the use of AI in the procedure, to make it even more efficient. Like many organizations experimenting, trialing, and iterating with AI, CHI has adopted a ‘toe in the water’ approach. This exploratory strategy is aimed at proving the value step-by-step, rather than automating every stage of the process immediately. Over the next few years CHI is hoping to optimize the diagnostic process further as time not spent analyzing video can be spent on patient care instead.

Products and Solutions

[Intel® Xeon® Scalable processors](#)

Industry

Health and Life
Sciences

Organization Size

11-50

Country

Denmark

Learn more

[Video](#)



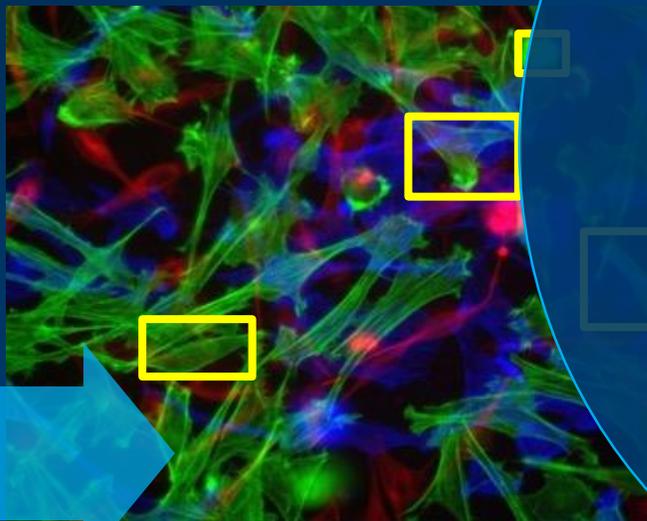
Solution: Drug Discovery Novartis*

IMAGENET*



224 X 224 X 3

26X
Larger



1024 X 1280 X 3

RESULT

6.6X FASTER TIME-TO-TRAIN

For processing a 10k image dataset, reducing the training time from 3.4 hours to 31 minutes with over 99% accuracy¹



Collaborator:

Novartis International AG*, based in Switzerland, and one of the largest pharmaceutical companies in the world

Challenge:

High content screening of cellular phenotypes is a fundamental tool supporting early stage drug discovery. While analyzing whole microscopy images would be desirable, these images are more than 26x larger than images found typically in datasets such as ImageNet*. As a result, the high computational workload would be prohibitive in terms of deep neural network model training time.

Solution: Intel and Novartis teams were able to scale to more than 120 (3.9Megapixel) images per second with 32 TensorFlow* workers.

Configuration: A cluster consisting of eight Intel® Xeon® processor servers using an Intel® Omni-Path Fabric interconnect and TensorFlow* optimized for Intel architecture.

*Other names and brands may be claimed as the property of others.

¹ Based on speedup for 8 nodes relative to a single node. Node configuration: CPU: Xeon Gold 6148 @ 2.4GHz, 192GB memory, Hyper-threading: Enabled. NIC: Intel® Omni-Path Host Fabric Interface, TensorFlow: v1.7.0, Horovod: 0.12.1, OpenMPI: 3.0.0. OS: CentOS 7.3, OpenMPU 23.0.0, Python 2.7.5. Patch Disclaimer: Performance results are based on testing as of May 25th 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. product can be absolutely secure. Time to Train to converge to 99% accuracy in model. Source: <https://newsroom.intel.com/news/using-deep-neural-network-acceleration>. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit <http://www.intel.com/performance>.

CLOUD SERVICE PROVIDERS





G-CORE LABS

High-Performance Infrastructure Underpins G-Core Labs' Public Cloud

G-Core Labs is an international cloud and edge leader in content delivery and broadcasting, hosting, security solutions and public cloud services. When online gaming company Wargaming needed to provide a seamless gaming experience for its millions of users, G-Core Labs built a world-class content delivery network (CDN) infrastructure, and then offered it to a wide range of other companies. G-Core Labs has now developed a public cloud service based on Intel® technologies, including the 2nd gen Intel® Xeon® Scalable processor. The new IaaS service has already won new customers across workloads including online gaming, high-definition video streaming and high-definition radio streaming. G-Core Labs is undertaking a proof of concept with Intel® Optane™ persistent memory to accelerate AI workloads and in-memory databases, which are essential for online gaming.

“The best thing about working with Intel is the expertise the team brings. They help us to understand how we can optimize our infrastructure and workloads to take advantage of the hardware and software available. We have an extensive research and development team, and they are in touch with Intel’s software and engineering teams who can help them with additional insights.”

Vsevolod Vayner, Cloud Platform Department Head, G-Core Labs

Products and Solutions

[2nd Gen Intel® Xeon® Scalable processors](#)

[Intel® Optane™ persistent memory](#)

[Intel® SSD D3-S4610 Series](#)

Industry

Cloud Services

Organization Size

51-200

Country

Luxembourg

Learn more

[Case Study](#)



HIGH-PERFORMANCE COMPUTING (HPC)





HLRN Brings Advanced Performance to HPC

HLRN operates one of the most powerful computers in Germany and the world, serving over 100 universities and over 120 research institutions. HLRN needs to account for the steadily increasing demand for capacities in HPC and data analytics for a myriad of scientific disciplines. HLRN selected new Intel® Xeon® Platinum 9200 processors for advanced performance in its next-generation supercomputer to enable significant computation gains and improved efficiency. HLRN-IV will be approximately six times as fast as the prior systems offering 16 PetaFLOPs of performance¹.

“We are not looking for peak theoretical performance, but for real system performance, and the Intel® Xeon® Scalable processor advanced performance is giving us this kind of performance. In addition, the CPU is quite good for artificial intelligence and machine learning.”

Prof. Ramin Yahyapour,
Professor of Computer Science at University Göttingen, Managing Director of GWDG

Products and Solutions

[Intel® Xeon® Platinum 9000 processors](#)
[Intel® Omni-Path Architecture](#)

Industry

Public Sector

Organization Size

10,001+

Country

Germany

Partners

[Atos](#)

Learn more

[Case study](#)
[Video](#)



Customer Success Story



MAX PLANCK INSTITUTE
FOR HUMAN COGNITIVE AND BRAIN SCIENCES



MAX PLANCK COMPUTING & DATA FACILITY
RECHENZENTRUM GARCHING DER MAX-PLANCK-GESELLSCHAFT

15X
REDUCTION
in memory
requirement
for image
processing¹

“The 3D data volume is at least a 1000 times larger than the previous 2D data volume, making the analysis and evaluation of individual layers by human experts impossible. By contrast, with Intel® OpenVINO™ processing times of one 3D image are now under an hour.”

Andreas Marek, Senior HPC Expert and Lead of the Data Analytics Group, Max Planck Computing and Data Facility

Researchers Accelerating the Analysis of 3D Brain Scans for Deeper Understanding

Researchers at the Max Planck Institute for Human Cognitive and Brain Sciences are seeking a deeper understanding of the human brain. Trained neural networks must process the large 3D images of cell structures produced by high-resolution microscopy producing 3D image data larger than the previously used 2D images by a factor of 1,000. Even high-performance computing (HPC) systems struggle to process the data volumes of this size. By adopting the Intel® Distribution of OpenVINO™ Toolkit, researchers unleashed the power of neural networks, enabling faster insights from 3D image data while minimizing hardware requirements. Because the project-specific workload was a perfect match for the performance characteristics of Intel® Optane™ persistent memory, the team achieved excellent results in a cost-efficient manner.

Products and Solutions

[2nd Gen Intel® Xeon® Scalable processors](#)

[Intel® Optane™ persistent memory](#)

[Intel® Distribution of OpenVINO™ Toolkit](#)

Industry

Research

Organization Size

10,001+

Country

Germany

Learn more

[Case Study](#)



FINANCIAL SERVICES INDUSTRY (FSI)



Case Study: Sentiment Analysis

Allgoo

RESULTS

Sentiment Analysis model

UP TO **1.47X** INCREASE

In inference performance over baseline for a Intel® Xeon® Scalable processor



Intel MKL-DNN

NEW

Partner: Allgoo, a software provider to the financial market, major banks, brokers, trading desks and investment funds in Brazil.

Challenge: Identify the risk profile of the major banks' customers to offer them products according to their risk appetite and tolerance. Allgoo's challenge was to have a powerful hardware to build their model of sentiment analysis and increase inference performance to be able to offer on demand investor profiles to their customers.

Solution: Allgoo's Sentiment Analysis model was developed using the powerful Intel Xeon Scalable processors. The solution also took advantage of Intel MKL-DNN optimized tools to have their expected performance. This improves the banks capabilities to offer personalized investment options to their customers.

Solution Catalog: [Link](#)

*Other names and brands may be claimed as the property of others.

Configuration: NEW: Tested by Intel as of 06/05/2019. 2 socket Intel® Xeon® Platinum 8153 Processor @2.00GHz, 16 cores per socket, 32 threads/socket, HT On, Turbo ON, Total Memory 394GB, BIOS: SE5C620.86B.00.01.0015.110720180833, Ubuntu 18.04.2 LTS, Kernel 4.15.0-46-generic, BIOS SE5C620.86B.00.01.0015.110720180833, Deep Learning Framework: Intel Optimized TensorFlow 1.13.1, MKL version: 2019.4, custom test data, for a custom CNN topology, tested using batches of 50, and 50 epochs.

BASELINE: Tested by Intel as of 06/05/2019. 2 socket Intel® Xeon® Platinum 8153 Processor @2.00GHz, 16 cores per socket, 32 threads/socket, HT On, Turbo ON, Total Memory 394GB, BIOS: SE5C620.86B.00.01.0015.110720180833, Ubuntu 18.04.2 LTS, Kernel 4.15.0-46-generic, BIOS SE5C620.86B.00.01.0015.110720180833, Deep Learning Framework: TensorFlow 1.13.1 from pip official channel (pypi), custom test data, for a custom CNN topology, tested using batches of 50, and 50 epochs.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary.

You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more

complete information visit <http://www.intel.com/performance>. Performance results are based on testing as of August 2018 and may not reflect all publicly available security updates.

See configuration disclosure for details. No product can be absolutely secure.

INDUSTRIAL AND MANUFACTURING



Smart Factories: Highly Connected and Instrumented



Customer: Intel Technology & Manufacturing

Challenge: Intel factories are continually challenged to improve efficiency and speed without sacrificing quality.

Solution: Increased factory automation and advanced analytics have helped reduce human error and have improved Intel's factory yield. We also use machine learning to evaluate vast amounts of yield data to determine types of defects. These insights help us to further improve factory processes.

Intel's smart factories are among those that now use edge computing and the Internet of Things (IoT) to enable automated control systems with real-time data. This data is categorized and prioritized in off-line systems as big data for ongoing analysis and decision making.

WHITE PAPER
MAY 2016

IT@INTEL
Driving Improvement in Manufacturing through Advanced Data Analytics



WHITE PAPER
APRIL 2015

IT@INTEL
Using Big Data in Manufacturing at Intel's Smart Factories



WHITE PAPER
JUNE 2016

IT@INTEL
Developing a Scalable Predictive-Maintenance Architecture



Embracing a world that connects and fundamentally changes how we use information technology in business.

Dr. Steve Chadwick
Senior Principal Engineer

Steve Chadwick
Manufacturing IT Principal Engineer, Intel IT

Duncan Lee
Manufacturing IT Principal Engineer, Intel IT

Steve J. Meyer
Manufacturing IT Principal Engineer, Intel IT

Joe Sartori
Industry Engagement Manager, Intel IT

Executive Overview
In the rapidly changing business and technology environment, manufacturers are striving to stay competitive, now and into the future, by increasing productivity and lowering costs.

Like many companies, Intel has evolved over the decades to provide competitive advantages through sophisticated automation—smart manufacturing—which has helped its factories increase product yields and quality, reduced costs, and improve safety. Intel's smart factories are among those that now use edge computing and the Internet of Things (IoT) to enable automated control systems with real-time data. This data is categorized and prioritized in off-line systems as big data for ongoing analysis and decision making.

Today, Intel manufacturing consistently reaps the benefits of improvement in the following areas:

- **Reduced costs.** Accurate and timely information in the hands of engineers improves product cycle time, process equipment maintenance, and other factors that save money.
- **Accelerated velocity.** Automated product flow enables easy access of products to available tools for processing, reducing bottlenecks and wait times.
- **Improved quality.** Statistical process control, advanced process control, and decision-support systems produce quality products, allowing engineers to focus on opportunities for improvement.

Through real-time capabilities, automation based on Intel provides a competitive advantage by using IoT and edge computing in manufacturing. Intel's investment in real-time data and provides significant value to the products we create.

Executive Overview

To gain and maintain a competitive advantage in today's world, manufacturers from all industries are turning to Industrial Internet of Things (IIoT) technology. IIoT can help identify excursions in standard processes, prevent unscheduled downtime due to tool failure, decrease overall maintenance costs, and improve productivity. But implementing IIoT solutions can be challenging.

Intel has used IIoT in our factories for years, and our solutions are often built on existing tools and processes. As we continue to learn, we now focus on connecting the unconnected and interconnecting intelligent devices to reduce strain on our network. In addition, we develop autonomous functionality that moves capabilities to the edge. Our goal is to design solutions that achieve interoperability and scalability that meet our needs today and in the future.

In early pilots, we reduced unscheduled downtime due to fan filter unit (FFU) failure by 300 percent

to monitor the health of our fan filter maintenance, our goals included minimizing maintenance, defining baselines for FFUs, alerting in real time, and more accurately predicting failures with GE Digital to develop a solution based on IIoT Gateways and GE's Predix platform.

Intel IT Peer Network

INTEL IT'S CONTINUED INVESTMENT IN FACTORY AUTOMATION PAYS BIG DIVIDENDS

By Joe Sartori, Intel IT

Increasing Product Quality and Yield Using Machine Learning



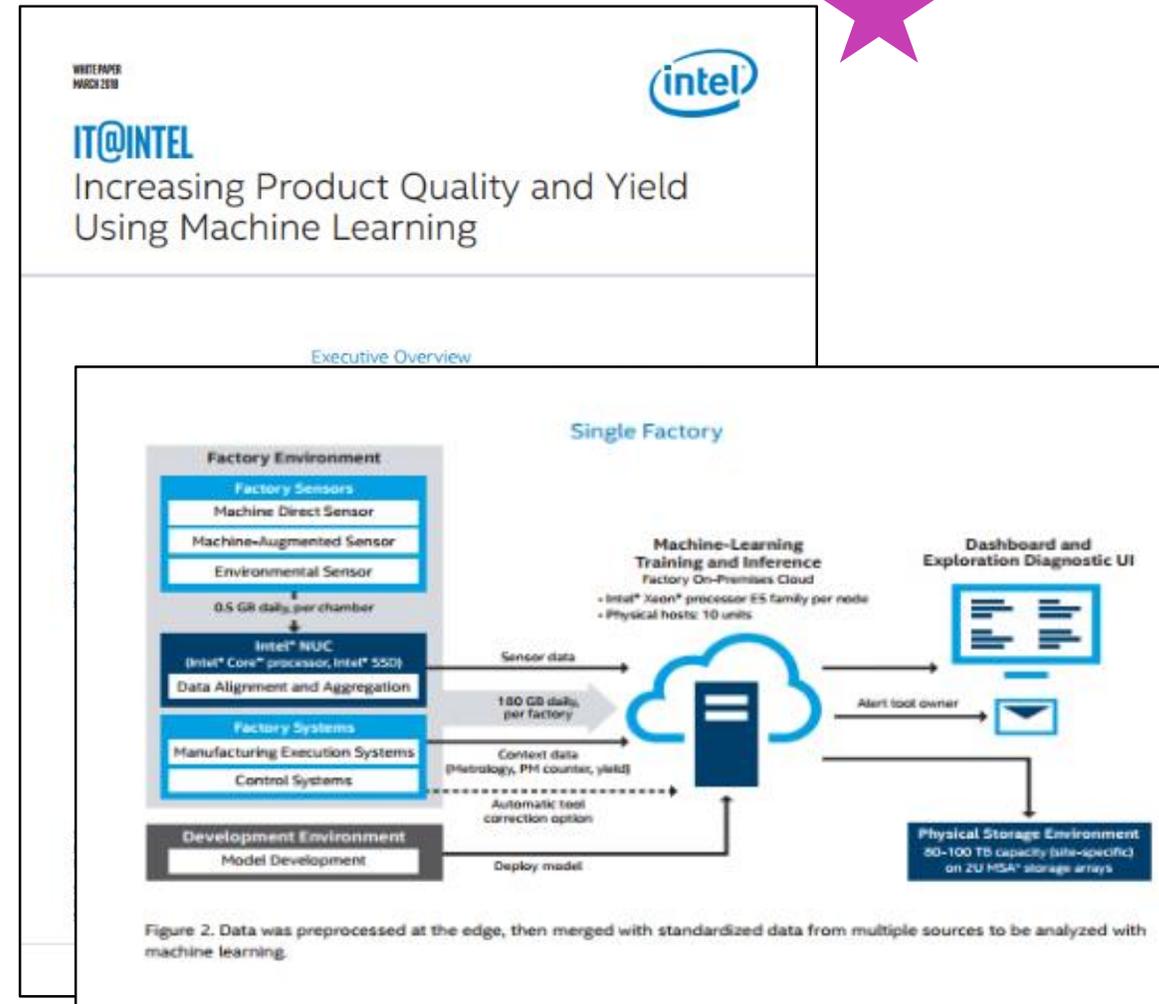
Customer: Intel Technology & Manufacturing Group

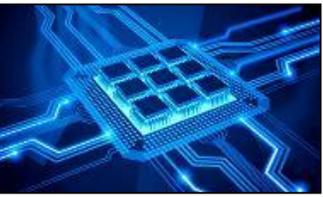
Challenge: Manufacturers across industries strive to improve throughput, yield, and product quality for better forecasting, cost reduction, and a competitive advantage. But, inconsistency in equipment performance and difficulty in predicting maintenance requirements often lead to quality issues and longer time-to-market.

Solution: We are using machine learning to predict issues with tool and relay forecasts in an intuitive, visual format, using customized frontend applications for large-scale activities, and web-based solutions for operational activities. Our predictive metrology solution connects process tools and collects rich data to provide:

- Better process equipment control and performance
- A high degree of certainty of product quality and conformity to specifications
- Verifiable engineering lead improvements with process diagnostics

The environment is supported by a variety of off-the-shelf products, such as third-party cloud solutions for computing and storage. We used standard Intel NUC hardware at the edge with an Intel Core i7 processor in virtual machines (VMs). We used an Intel® Xeon® processor E5-2643 for data acquisition and storage, which provided excellent compute power. The most important aspect is that the factory must be networked to collect and transmit the data, especially to achieve near real-time insights.





AI Reduces Costs & Accelerates Time-to-Market

Public

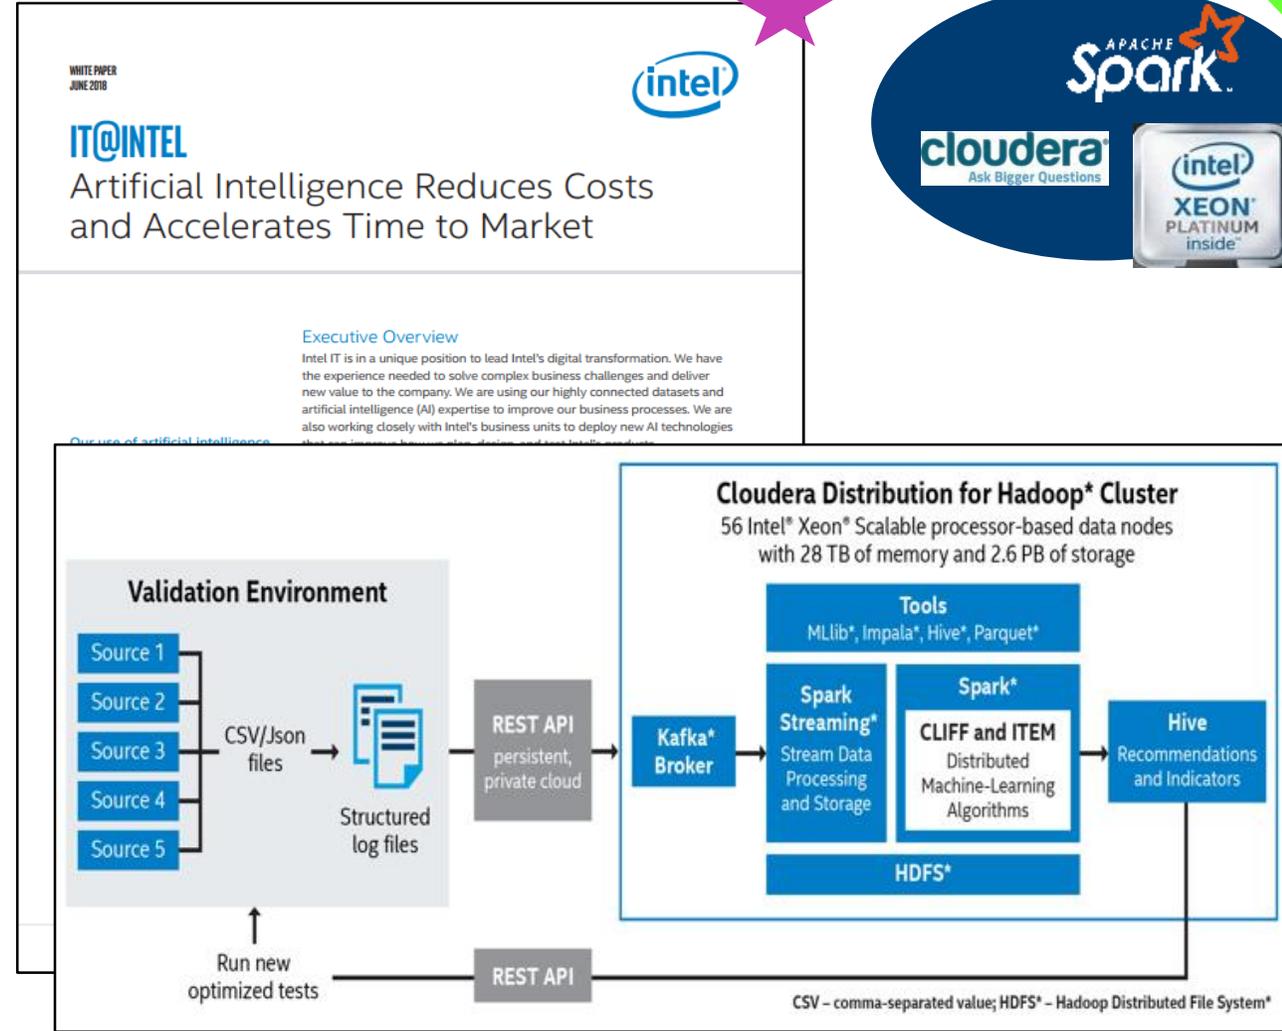
Customer: Intel Platform Engineering Group

Challenge: Semiconductor design-phase product validation is the one of the most expensive and time-consuming product-development processes, consuming up to 50 percent of the development cycle and requiring several iterations of design and prototyping.

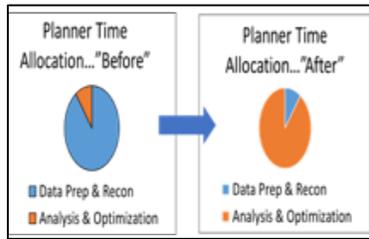
Solution: Intel IT collaborated with our product development teams to optimize the pre-silicon validation stage by automating and augmenting human validation capabilities. We developed two AI-based solutions that speed time-to-market and reduce costs:

- CLIFF (Coverage LIFt Framework) speeds validation by creating new tests for hard-to-validate functionalities, which then discovers hidden bugs as early as possible. On average, CLIFF improves the targeted functionalities coverage by 230x, compared to standard regression tests.
- ITEM (Intelligent Test Execution Management) creates the best testing suite each week. ITEM ensures that, from a bug-finding and functionality-coverage perspective, the teams run the most cost-effective tests. ITEM reduced the number of required tests by 70%.

CLIFF and ITEM run in our Cloudera Distribution of Hadoop with high-performance servers based on Intel® Xeon® Scalable processors. Multiple input datasets are converted into many thousands of structured log files each week. The system processes many dozens of GBs per week, and can store many TBs across multiple product development teams.



Optimizing Inventory via Advanced Analytics



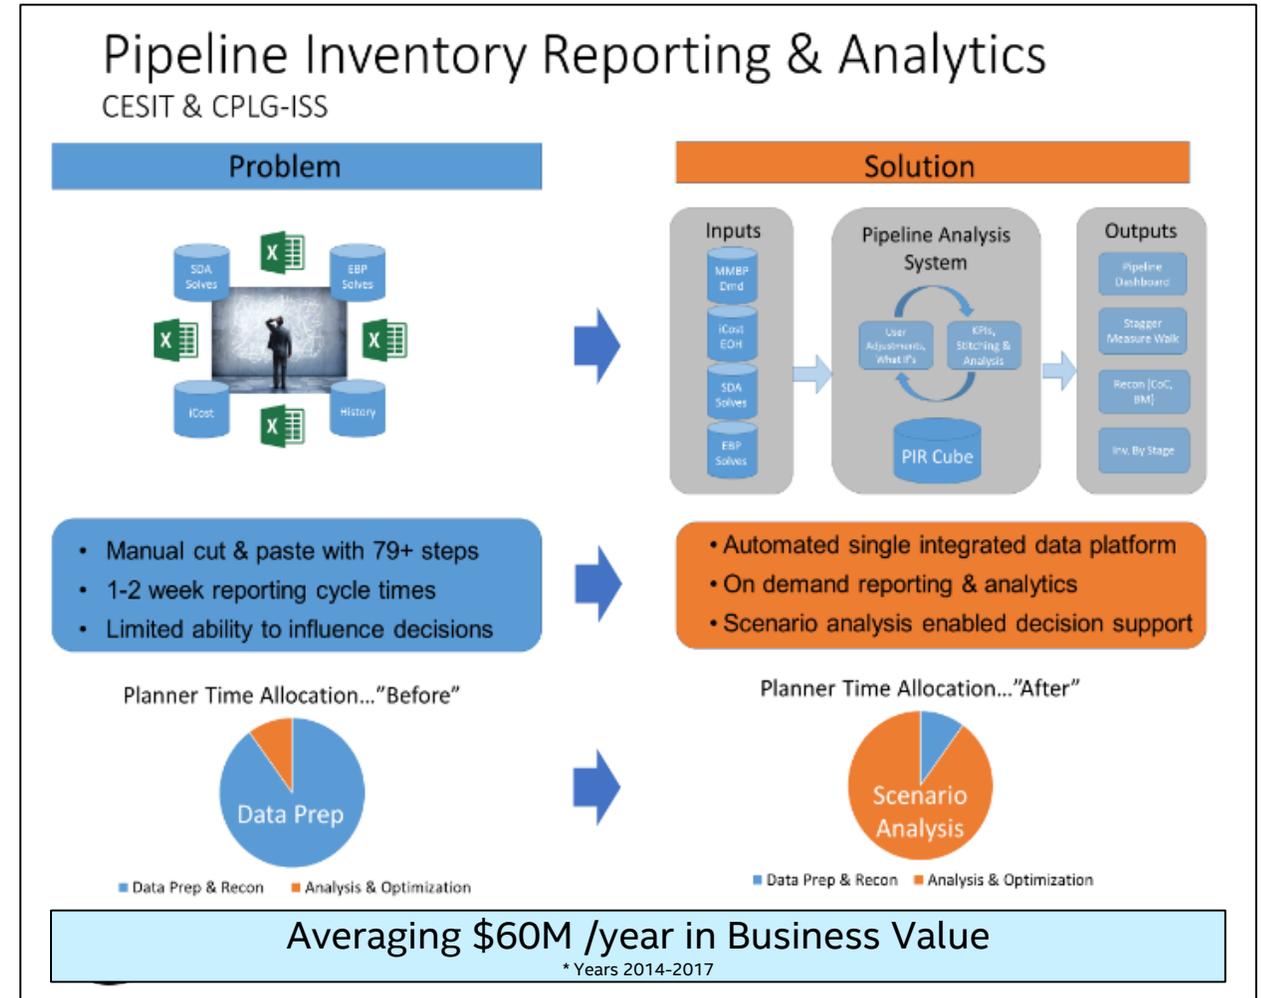
Customer: Intel Planning, Order Fulfillment and Finance

Challenge: Intel's Planning organization needed to significantly shorten the time to generate inventory reports and optimize inventory to meet demand.

Solution: Intel IT and Planning built a new planning system that shortens time to gather data and generate reports. The resulting system delivers an average of **\$60M in business value per year** by reducing pipeline inventory. Lower inventory levels mean money that was tied up in finished goods can now be used for higher-value activities.

We are investigating additional use cases where the automated inventory management system can improve supply chain management. These use cases include producing integrated business planning reports that can help inform outsourced manufacturing decisions, lower supply chain costs, and enable flexible procurement vendor choices.

Today, this workload runs in an Intel Data Center. In 2018, the inventory build data will be moved to our Integrated Data Platform (IDP). IDP combines SAP HANA and Cloudera Distribution of Hadoop, running on Intel Xeon Gold 6132 processors.



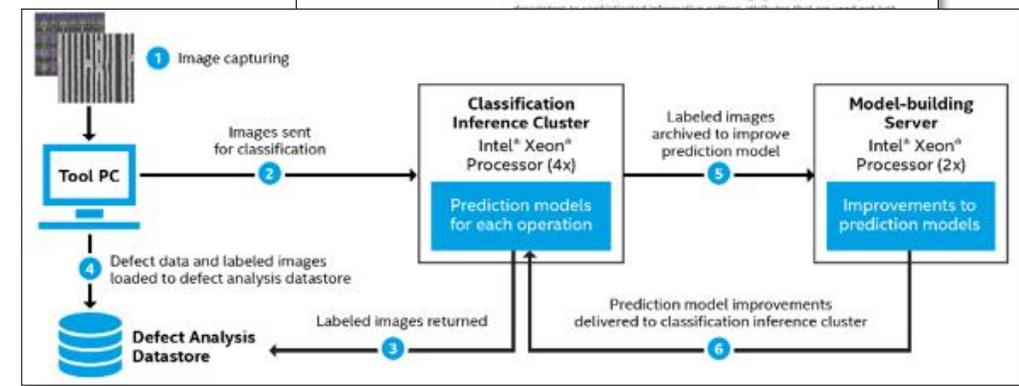
Faster, More Accurate Defect Classification Using Machine Vision

Customer: Intel Technology & Manufacturing Group

Challenge: Highly sensitive automated defect classification (ADC) using machine vision (MV) and machine learning (ML) can increase early defect detection and improve classification accuracy and consistency. This allows factories to rapidly identify and correct defects to maximize machine capacity and process high-quality products.

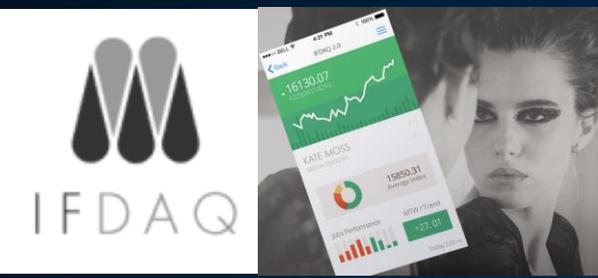
Solution: At Intel, we use MV and ML in a variety of manufacturing use cases, some of which are quite complex. During fabrication, we collect images from multiple channels and compute thousands of features ranging from standard shape descriptors to sophisticated informative pattern attributes that are used not just for defect classification, but also to provide insights into mechanisms generating specific defect modes. The high-volume ADC solution we developed for wafer fabrication uses Intel® Xeon® processors and has resulted in greater accuracy with a larger number of classification types. We also use MV and ML to repurpose existing images and identify defects in our assembly and test factories. Our ADC solution has delivered the benefits at a scale and level of accuracy that we could not have achieved even with the most experienced technicians.

We are now able to measure and classify more than 80 percent of the wafers that Intel's factories produce with a higher degree of accuracy, without increasing our total cost of ownership (TCO).



RETAIL





Case Study: Predicting the market development and performance of fashion and luxury brands

RESULTS

UP TO 3.5X increase

In **Training** Performance on 2nd Gen Intel Xeon Scalable using Intel optimized Tensorflow



Public

NEW

Partner:

IFDAQ commenced as a scientific research project in 2008, when a consortium of AI pioneers and Big Data leaders mutually explored new technologies to refine the world's largest fashion data asset.

Challenge:

The fashion and luxury industry is highly complex and insight into fundamental KPI's gives companies valuable information which they can use to improve their competitive edge. IFDAQ is looking for compute performance which can deliver these results to their customers in a timely and cost efficient manner.

Solution:

The solution helps in visualizing, predicting and monitoring the market's dynamics and impacts in real-time and with meaningful transparency. They were able to predict careers and performances of professional fashion models by calculating a final benchmark value from decisive factors under the most complex conditions. This guarantees an accurate and precise rating of the entity and shows 3.5x improvement in training performance using Intel optimized Tensorflow on 2nd Gen Intel Xeon Scalable processor

*Other names and brands may be claimed as the property of others.
Configuration: IFDAQ Configuration: NEW: Tested by Intel as of 08/06/2019. 2 socket Intel® Xeon® Gold 6248 Processor, 20 cores per socket, OS Ubuntu 18.04.2 LTS, Deep Learning Framework: Intel Optimized TensorFlow 1.12.0, custom test data, for Feed Forward (single layer), Batch size : 200, python 3.6
BASELINE: Tested by Intel as of 08/06/2019. 2 socket Intel® Xeon® Gold 6248 Processor, 20 cores per socket, OS Ubuntu 18.04.2 LTS, Deep Learning Framework: TensorFlow 1.12.0, custom test data, for Feed Forward (single layer),Batch size : 200
Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.
Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary.
You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit <http://www.intel.com/performance>. Performance results are based on testing as of August 2018 and may not reflect all publicly available security updates.
See configuration disclosure for details. No product can be absolutely secure.

Case Study: Predicting the market development and performance of fashion and luxury brands



RESULTS

UP TO 2.4X

Improvement in Inference on 2nd Gen Intel Xeon Scalable processor using Intel optimized Tensorflow



Public

NEW

Customer:

IFDAQ commenced as a scientific research project in 2008, when a consortium of AI pioneers and Big Data leaders mutually explored new technologies to refine the world's largest fashion data asset.

Challenge:

The fashion and luxury industry is highly complex and insight into fundamental KPI's gives companies valuable information which they can use to improve their competitive edge. IFDAQ is looking for compute performance which can deliver these results to their customers in a timely and cost efficient manner.

Solution:

The solution helps in visualizing, predicting and monitoring the market's dynamics and impacts in real-time and with meaningful transparency. The solution is able to predict careers and performances of professional fashion models by calculating a final benchmark value from decisive factors under the most complex conditions. This guarantees an accurate and precise rating of the entity and shows 2.4x improvement in Inference performance using Intel optimized Tensorflow on 2nd Gen Intel Xeon Scalable processor

*Other names and brands may be claimed as the property of others.

Configuration: IFDAQ Configuration: NEW: Tested by Intel as of 08/06/2019. 2 socket Intel® Xeon® Gold 6248 Processor, 20 cores per socket, OS Ubuntu 18.04.2 LTS, Deep Learning Framework: Intel Optimized TensorFlow 1.12.0, custom test data, for Feed Forward (single layer), Batch size : 200, python 3.6

BASELINE: Tested by Intel as of 08/06/2019. 2 socket Intel® Xeon® Gold 6248 Processor, 20 cores per socket, OS Ubuntu 18.04.2 LTS, Deep Learning Framework: TensorFlow 1.12.0, custom test data, for Feed Forward (single layer), Batch size : 200

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary.

You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more

complete information visit <http://www.intel.com/performance>. Performance results are based on testing as of August 2018 and may not reflect all publicly available security updates.

See configuration disclosure for details. No product can be absolutely secure.

ENERGY





**80%
REDUCTION**
in inference time¹

“We are using Pipe Sleuth to inspect our sewer network. It is a very innovative solution that dramatically increases inspection productivity and significantly reduces costs while at the same time improving the overall defect detection rate. Having the option to run Pipe Sleuth on our existing Intel-based platforms was an added benefit.”
Tom Kuczynski, Vice President, Information Technology, DC Water

Automating Pipeline Inspection with Computer Vision and Deep Learning

DC Water distributes drinking water and collects and treats wastewater for more than 672,000 residents and 17.8 million annual visitors in the District of Columbia. The manual review of and classification of video scans of sewers and utilities pipeline infrastructure can be time consuming. DC Water was looking for a better way of conducting analysis and reporting of pipelines to optimize service and repair and infrastructure replacement and lower costs. DC Water, in collaboration with Wipro, developed Pipe Sleuth to automate the process of identification, annotation, scoring/grading of pipeline health and reporting of pipeline defects. Optimized Pipe Sleuth with Intel® Xeon® Scalable processors and Intel® Distribution of OpenVINO™, resulted in faster time-to-market, cost savings on analysis and allows for more spending on capital improvements.

Products and Solutions

[Intel® Xeon® Scalable processors](#)
[Intel® Distribution of OpenVINO™ toolkit](#)

Industry

Energy & Utilities

Organization Size

1,001-5,000

Country

United States

Partners

[Wipro](#)

Learn more

[Case Study](#)
[Podcast](#)

¹ For more complete information about performance and benchmark results, visit <https://www.intel.ai/solutions/dcwater-wipro/>



TRANSPORTATION AND LOGISTICS





Yunda Express Improves the Efficiency of their Delivery Logistics System

Yunda Express has developed rapidly with a comprehensive service network covering the entire delivery chain. With a skyrocketing volume of business, its traditional mode of manual operation is now time-consuming, labor-intensive and unable to guarantee the desired quality. Using automation and artificial intelligence (AI) to save cost and enhance efficiency of the delivery logistics system has become an important means for Yunda to tackle these challenges. Yunda deployed Analytics Zoo, a unified big data analytics and AI platform, as well as Intel® Xeon® Scalable processors, and carried out an all-level technical collaboration with Intel. Intel helped Yunda build high-efficiency AI applications for several key links including “package size measurement”, “data center anomaly detection” and “shipment quantity prediction”.

“The introduction of advanced products and technologies such as the Analytics Zoo platform and Intel® Xeon® Scalable processors can help us greatly improve the efficiency of the entire express delivery chain, optimize resource utilization, significantly reduce operational costs, and effectively tackle the challenges for future development.”

Peiji Li, Chief Architect of Yunda Express, Yunda Co., Ltd.

Products and Solutions

[Intel® Xeon® Scalable processors](#)
[Analytics Zoo](#)

Industry

Logistics & Supply Chain

Organization Size

5,000-10,000

Country

People's Republic of China

Learn more

[Case Study](#)

OTHER





kakao



“Using Intel solutions helps Kakao improve performance, simplify maintenance and save time while providing a better quality of AI service.”

Biho Kim, Manager of Voice Processing Part, Kakao

Improved AI Performance for Speech Recognition Workloads Through Optimization

Kakao Corp is the #1 mobile life platform company in Korea, creating a more comfortable and exciting world by innovating the way we connect. Kakao hosts more than tens of thousands of physical servers, virtual machines and containers for various workload such as communication, portal, contents, commerce, game and many more. Kakao was most interested in finding a way to improve speech recognition inference performance and velocity for the company’s artificial intelligence (AI) speaker service running on Intel® Xeon® processors. Kakao’s proof of concept (PoC) system with 2nd Gen Intel® Xeon® Scalable processors, delivered up to 41% performance improvement with a combination of Intel® Advanced Vector Extensions 512 (Intel® AVX-512), Intel® Deep Learning Boost (Intel® DL Boost) and Intel® Math Kernel Library (Intel® MKL).

- Products and Solutions**
- [2nd Gen Intel® Xeon® Scalable processors](#)
- [Intel® Math Kernel Library](#)
- [Intel® Advanced Vector Extensions 512](#)

Industry	Organization Size	Country	Learn more
Information and Communication Technology	1,001-5,000	South Korea	About Kakao

¹ The results were provided by Kakao and were based on its internal tests. For more complete information, about performance and benchmark results, please contact Kakao.

Solution: Call Center Routing Gigaspaces*

RESULT

“BigDL and AI portfolio provide an infrastructure-optimized solution for deep learning workloads leveraging Intel® Xeon® Scalable processors [...] with low risk and TCO.



Client: GigaSpaces*, provider of in-memory computing platform for fast data analytics and extreme transaction processing.

Challenge: Lowering the call center routing latency to improve customer experience through customer/data-360 integration. Goal is to automatically trigger transactional workflows to route call to correct agent based on prediction criteria and scoring.

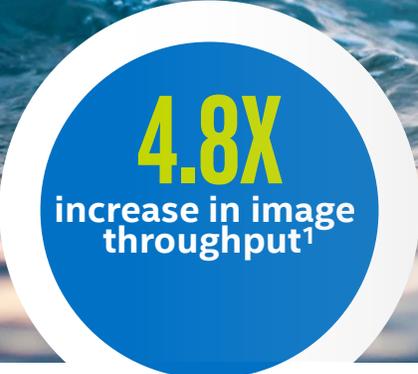
Solution: GigaSpaces* InsightEdge* and BigDL, running on Intel® Xeon® & 5th Gen Intel® Core™ processors, enable real-time analysis & business logic workflows. Continuous training models are built & deployed with no downtime.

Source: Performance numbers provided by GigaSpaces
<https://blog.gigaspaces.com/gigaspaces-to-demo-with-intel-at-strata-data-conference-and-microsoft-ignite/>

*Other names and brands may be claimed as the property of others.
 Intel does not control or audit third-party benchmark data or the web sites referenced in this document.
 You should visit the referenced web site and confirm whether referenced data are accurate.



KONGSBERG



“We have been working with TensorFlow a lot, but the resources usually assume that you will be using GPUs. Working with Intel has enabled us to optimize our solution for CPUs, so we can benefit from using a more standardized server platform.”

Jaakko Saarela, project Manager, Kongsberg Maritime

Accelerating Image Recognition for Kongsberg Maritime’s Marine Navigation Solution

Kongsberg Maritime has a vision to improve safety and increase the efficiency of shipping and is pioneering autonomous ships and using artificial intelligence (AI) to support crews with navigation at sea. The company’s AI solutions have traditionally been based on GPUs, but Kongsberg Maritime wanted to use CPUs so it can simplify and consolidate its servers more easily. Working with Intel, Kongsberg Maritime was able to increase its solution’s performance on standard server hardware by 4.8X¹ on one of the company’s demonstrator projects. The Intel® Distribution of OpenVINO toolkit™ was used to accelerate the performance of Kongsberg Maritime’s TensorFlow* model running on Intel® Xeon® Scalable processors.

Products and Solutions

[Intel® Xeon® Scalable processors](#)
[Intel® Distribution of OpenVINO™ toolkit](#)

Industry

Maritime

Organization Size

5,001-10,001

Country

Norway

Learn more

[Case Study](#)

¹ For more complete information about performance and benchmark results, visit <https://www.intel.ai/case-studies/kongsberg-maritime/>

