

Bringing AI Everywhere

Highlights

- ▣ GenAI and the enterprise
- ▣ Planning your enterprise AI strategy
- ▣ Getting into production
- ▣ Managing cost
- ▣ The Intel AI stack in the enterprise
- ▣ Bringing AI everywhere
- ▣ AI in action
- ▣ Building a better future

Embracing the AI revolution

In today's rapidly evolving landscape, artificial intelligence (AI) is not just a buzzword but a transformative force reshaping industries, economies, and lives. The speed and extent of adoption are expected to be faster than the internet era, the mobile era, and the cloud era.

An enterprise-centric perspective

When investing in technology to bring AI to fruition, two pivotal needs arise: maximizing value and accelerating innovation. It's about realizing visionary projects at a lower cost and breaking free from the "AI failure" trap to achieve market success with unprecedented speed.

Maximize value:

Realizing your vision economically

Embarking on the AI journey promises immense potential but can come with high costs. The key to unlocking this potential lies not in extravagant spending but in leveraging the platforms and tools you already have and know to scale AI across your enterprise. It's about finding the most cost-effective way to harness the power of AI to streamline operations, enhance decision-making, and, ultimately, drive down costs. This approach doesn't just aim for efficiency; it seeks to redefine value creation in an AI-driven era, ensuring that every investment propels you closer to your vision, economically and effectively.

Accelerate innovation:

Navigating the path to success

The path to innovation is fraught with the risk of "AI failure" - projects that falter, not from a lack of ambition but from a disconnect between technology and real-world application. Accelerating innovation means moving beyond this trap. It's about identifying the right use cases to quickly bring value to your business, and enterprise-ready solutions that get you into production fast. This journey is not just about speed but about smart acceleration - where agility meets precision, ensuring that each step forward is as impactful as it is swift.

The journey into AI is not just a technological venture but a strategic one, requiring a nuanced understanding of both the opportunities and challenges. It's a journey of exploration, requiring curiosity, vigilance, and an unwavering commitment to not just envisioning a better future but actively building it, one responsible, sustainable, and secure step at a time.

At Intel, we recognize this balance of needs and aspirations, ensuring our AI initiatives align closely with the real-world needs of our customers and society.

Start with a strategic intent

AI integration has the potential to touch every corner of the enterprise, demanding a review of strategies and operations. From enhancing efficiency across functions to redefining the nature of work, AI's influence is undeniable. Although large enterprises have the most to gain by successfully scaling AI across their organizations, companies of all sizes can benefit from new levels of innovation and efficiency.

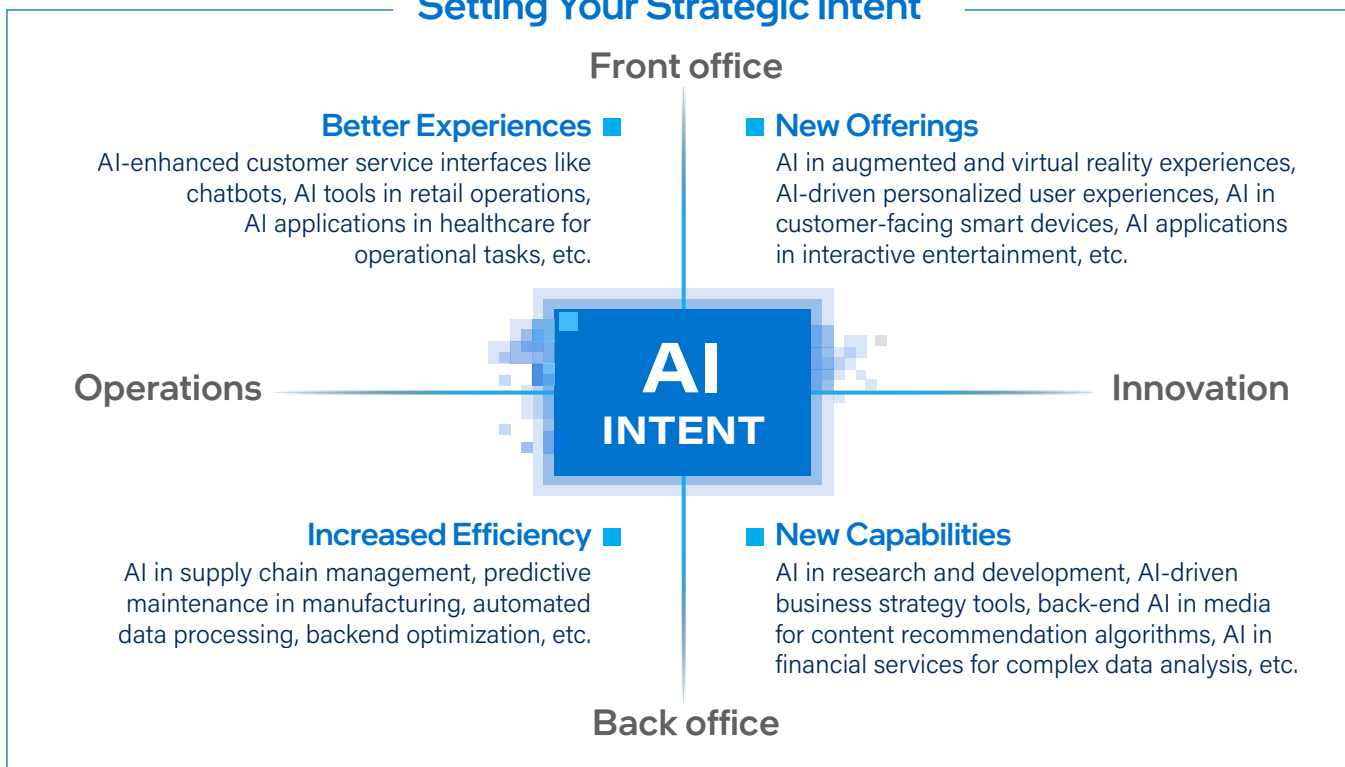
AI Transformation

Just like the digital transformation wave that swept across industries, companies now face the imperative of AI transformation. This transition involves leveraging AI to enhance efficiency and effectiveness across all functions. Businesses need to act swiftly and purposefully to integrate AI technologies, ensuring they harness the full potential of AI to stay ahead in their respective fields.

AI Intent

Defining a clear AI intent is crucial for businesses embarking on this journey. The strategy could be innovation-led, focusing on creating new value and opportunities, or it might be centered on operational efficiency, streamlining existing processes. Whether AI is applied in the front end, enhancing customer experiences, or in the backend, optimizing behind-the-scenes operations, understanding the specific goals and needs is essential for success.

Setting Your Strategic Intent



Enterprise-ready solutions

In a hypercompetitive world, speed to market is just as critical as innovation. Production-ready solutions from a solution integrator, software vendor, or cloud service provider can help enterprises scale out AI while making the desired business impact.

Considering Intel

- With tens of thousands of real-world deployments, Intel and our partners have become experts in market- and domain-specific enterprise AI use cases.
- Intel's global ecosystem helps you quickly deploy the AI that's right for your business—cloud and data center, in the network, on the PC, and at the edge.

Healthcare

AI is transforming healthcare by streamlining administrative tasks and enhancing care. It automates paperwork and patient notes, allowing healthcare workers to focus on care. AI can speed up medical imaging and enable remote patient care so that more patients can be served.

Samsung Medison achieved more accessible and cost-effective AI medical imaging by speeding up throughput for real-time ultrasound imaging with Intel® Core™ Ultra processors, compared to previous generation processors with a discrete GPU.¹

Finance

Across the financial services industry, AI is being used to capture real-time insights from massive amounts of user and financial data. Financial services organizations are putting these insights to use to stop fraud, optimize profits, predict outcomes, and dramatically streamline customer experiences.

Data teams in finance, insurance, and banking are able to improve fraud detection and risk management. In this domain, we've collaborated with our software partner, Fortanix to create the Intel® Security Solution for Fortanix Confidential AI.²

Retail

AI is reinventing the retail landscape. Machine learning is improving the guest experience and streamlining inventory management practices. Retailers that harness AI to connect with customers and operate more efficiently will be better positioned to thrive.

Hellometer's AI timer, optimized with Intel OpenVINO™ software, enables restaurant operators to improve drive-thru service speed and increase sales by 47 seconds on average, translating into about \$130k in added revenue per location.³

Manufacturing

In manufacturing, AI can help improve product quality and yield, reduce downtime, automate processes, and inform new solutions. Machine vision and industrial robotics can revolutionize asset tracking, predictive maintenance, and factory automation.

The Siemens SIMATIC Industrial PC is empowering better product quality and yield in the automated factory, delivering up to 4.25x higher inference performance for object detection and up to 3x higher inference performance for image classification.⁴

Training a model vs. running one

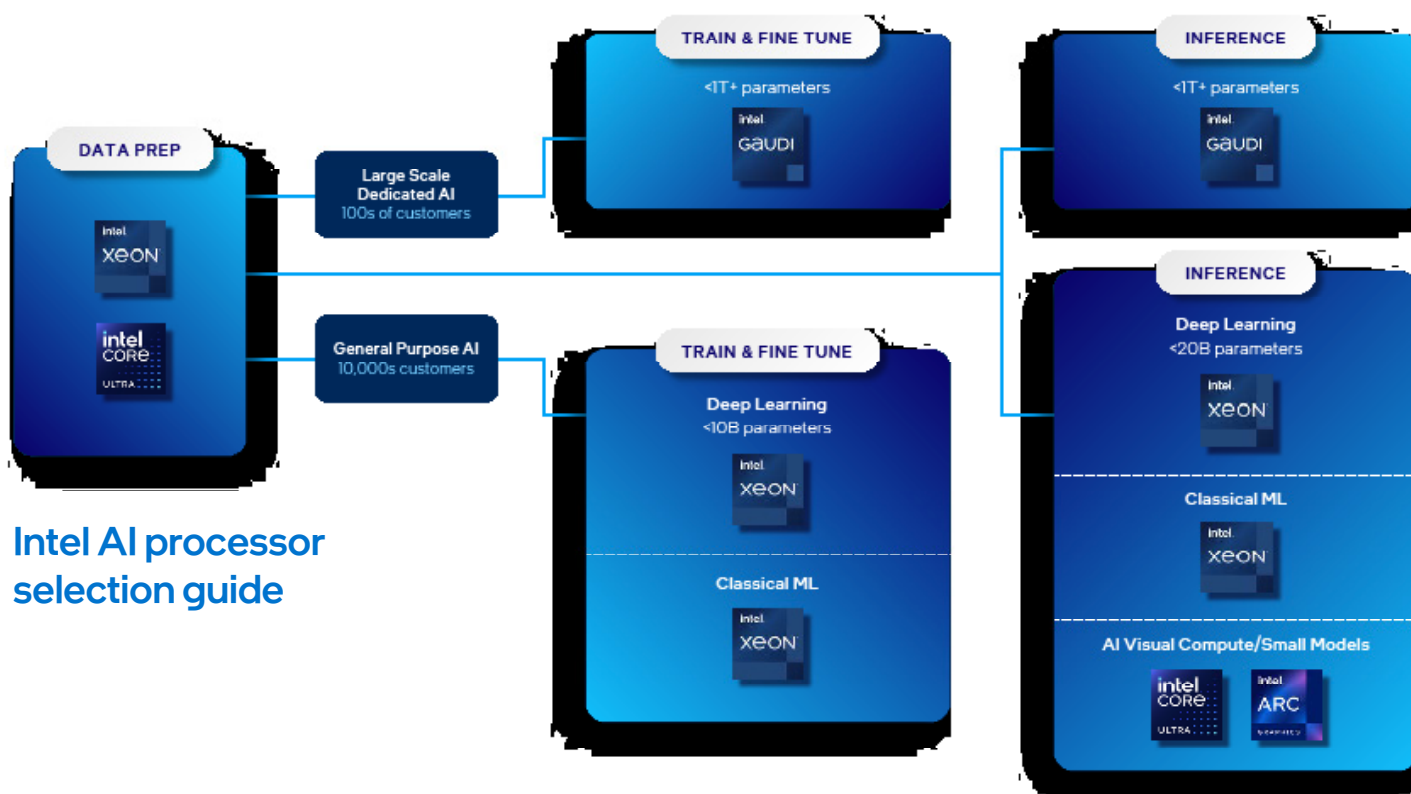
Using a model in the real world (inference) is quite different from creating it (training). Training needs high-powered data center infrastructure. This is where GPUs have become popular. But inference happens everywhere—from the cloud and network to much lighter devices, like PCs, point-of-sale terminals, robots, and CT scan machines, running alongside many other applications—and the solutions will differ based on the scale of data retrieval and the types of genAI models used.

The center of gravity for enterprise AI is inference

How and where an enterprise runs inference has become the top driver for AI infrastructure considerations. Where inference runs impacts the cost, efficiency, and privacy of your AI infrastructure. The most successful strategies combine lightweight, real-time insight from PCs and edge devices with deeper context from the cloud, allowing constant interaction across various model types, latency needs, and regulatory restrictions.

Considering Intel

- The Intel infrastructure you have and know is increasingly engineered for enterprise AI—our goal is to make every enterprise device as productive for GenAI as possible. By running AI on the platforms you already use, you can maximize the value of your investments and lower your total cost of ownership.
- The same Intel® Xeon® CPUs that run your everyday workloads—whether on prem or in the cloud—are also great for inference, especially if you choose the latest processors with built-in AI acceleration.
- If you need higher performance, maximize the value of your investments with Intel® Gaudi® AI accelerators or Intel® GPUs.
- Intel® Core™ Ultra processors offer substantial inference capabilities and performance in edge devices and AI PCs.
- This common, edge-to-cloud Intel foundation makes it easier to run AI everywhere with minimal software changes.



Intel AI processor selection guide

Understanding the AI data pipeline

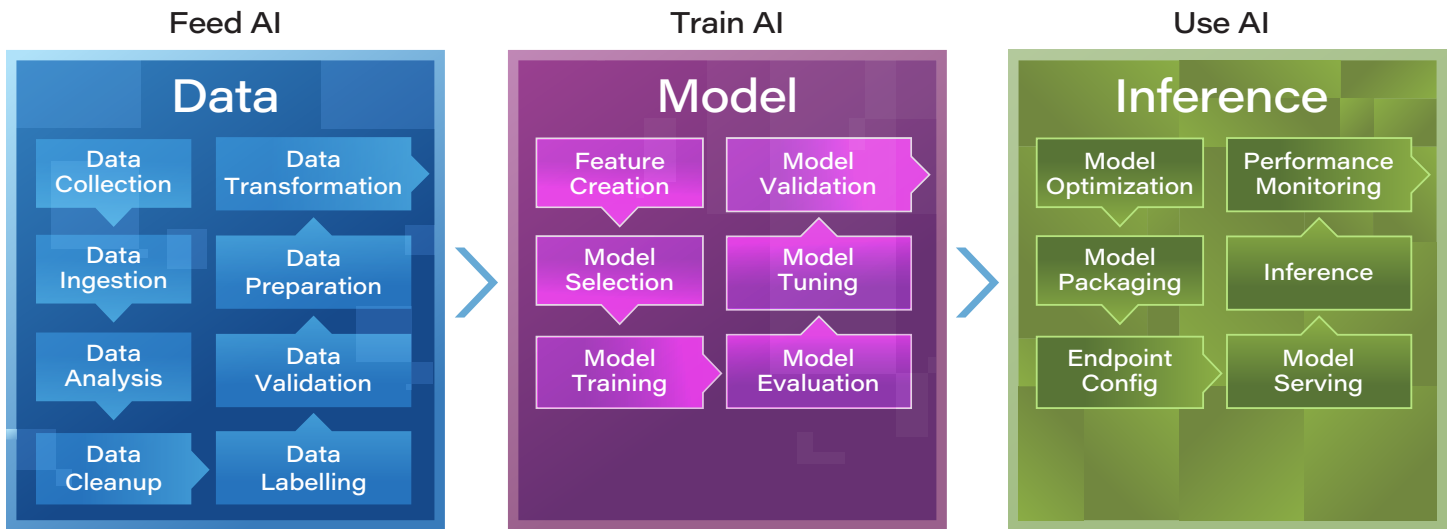
Integrating AI into your enterprise is about more than running a model. Enterprises generate an incredible amount of data, and it is estimated that ~90% of it is unstructured. Your holistic data pipeline includes data ingestion, database and storage, training or fine-tuning, and inference across multiple clouds or hardware types.

The good news is that most enterprises won't need to train a model from scratch. Fine-tuning a model or using retrieval-augmented generation (RAG) can help you secure a competitive advantage by using your data to customize a model or its output.

- **Feed AI – the Data Phase:** involves carefully collecting, augmenting, and preprocessing data.
- **Train AI – the Model Phase:** selection, development, and training of accurate and reliable models.
- **Use AI – the Inference Phase:** the effective application of these models to make informed decisions or automate processes.

Considering Intel

- Intel platforms excel throughout the entire AI data pipeline, from data collection to inference, offering exceptional value.
- All major AI frameworks have been optimized by using oneAPI libraries, helping ensure high performance on Intel® hardware.



RAG: GenAI deployment model in the enterprise

GenAI inaccuracies and hallucinations are unacceptable in enterprise environments, posing critical reliability risks. To realize AI's full value, enterprises will need solutions capable of securely producing business-specific results using their proprietary data and business models. Retrieval-augmented generation (RAG) is a technique for enhancing the accuracy and reliability of GenAI models with data pulled from an external trusted source. The RAG workflow includes data retrieval and generation. Data retrieval will happen across every layer of the enterprise (PC, edge, data center) and include standard text

as well as multimodal data (video, audio, and increasingly other sensors).

A significant amount of compute will be spent extracting context from this data using traditional data engineering methods. These workloads are well suited for our CPUs (Xeon and Core). Once data retrieval is complete, the contextual information is encoded into a larger prompt, called embedding, which is then fed into a GenAI model to generate a response. This part of the workflow is well suited for accelerators (Gaudi and GPU).



GenAI and the enterprise

The advent of Generative AI (GenAI) is transforming the business landscape, offering unprecedented opportunities for innovation, efficiency, and customer engagement. As enterprises seek to leverage these technologies, they encounter a unique set of challenges and considerations. From ensuring data privacy and integrity to integrating AI seamlessly into existing systems, the journey to harnessing the power of GenAI is complex yet rewarding.

GenAI innovations are expected to be adopted quickly by enterprises of all sizes. Here are six key points that encapsulate the essence and strategic considerations for integrating GenAI into the enterprise.

- 1. Enhancing enterprise efficiency through GenAI:** Exploring how GenAI technologies not only streamline operations and enhance productivity but also drive innovative use cases within enterprises. Setting the right AI intent to ensure that applications of GenAI align with business objectives will deliver tangible benefits.
- 2. Privacy and data concerns in building GenAI systems:** Addressing the critical importance of data privacy and security when enterprises embark on building their GenAI-led systems. Pre-trained nimble models and the need for robust data handling and protection mechanisms will be on the rise.
- 3. Ensuring reliability and accuracy with GenAI:** Focusing on the challenges of maintaining high levels of reliability and

accuracy in GenAI outputs is crucial for enterprise decision-making. Retrieval-Augmented Generation (RAG) methods will enhance the quality of GenAI responses, mitigating the risk of inaccuracies and hallucinations.

- 4. The center of gravity for enterprise AI is inference:** As "off-the-shelf" models proliferate, enterprises are thinking less about creating new models (training), and more about ensuring their data can be used to create custom output (inference). How and where an enterprise runs inference has become the top driver for AI infrastructure considerations.
- 5. Overcoming integration challenges for GenAI in enterprise systems:** Discussing the technical and operational hurdles of integrating GenAI into existing enterprise ecosystems. Highlighting strategies for seamless integration, the advantages of employing compatible systems for both operational and AI platforms to ensure efficiency and scalability.



Bringing AI everywhere

In the unfolding age of AI, Intel is helping enterprises build an AI-driven future that transcends the boundaries of mere technological advancements. “Bringing AI Everywhere” is not just a message; it’s our commitment to enterprises as they navigate how to realize value in this quickly evolving landscape.

Pioneering a Ubiquitous AI Future: The Intel Advantage

Bringing AI Everywhere requires a diverse offering of hardware and software. Intel stands unique in this endeavor as a world leader in offering technologies spanning this entire spectrum with an open, secure, and responsible approach.

Our strategy is bold, yet simple—infuse AI computing capabilities into all our technologies. From the AI PC and the dynamic edge to the robust data center and the expansive cloud, our heterogeneous architectures enable AI at every touchpoint in your enterprise. With Intel® Xeon® and Intel®

Core™ Ultra processors, Intel® Data Center GPUs (Max Series and Flex Series), Intel® Gaudi® deep learning accelerators, and Intel® Arc™ discrete graphics, we deliver competitive, performant, open-standard solutions. Our offerings empower enterprises to deploy AI at scale, catering to a full spectrum of workloads and usage models with agility and precision.

A New Horizon of Possibilities

Welcome to the new era of AI, where business as usual evolves into the realm of the unprecedented. With Intel, you can embark on a journey where AI is not a distant dream but an immediate reality where the future is not just bright—it’s brilliant.

Open software environment

Your success in AI will be defined by software—and by your developer teams. Whether they are working in data prep, training, or deployment, your developers will likely have a strong preference in the libraries, frameworks, and models they use.

Flexibility will be key, and an open software environment will enable your developers to use the tools they want on their hardware of choice. OpenVINO can maximize the impact of their efforts by allowing them to develop AI applications once, and then run them across multiple hardware types without extensive changes to their code.

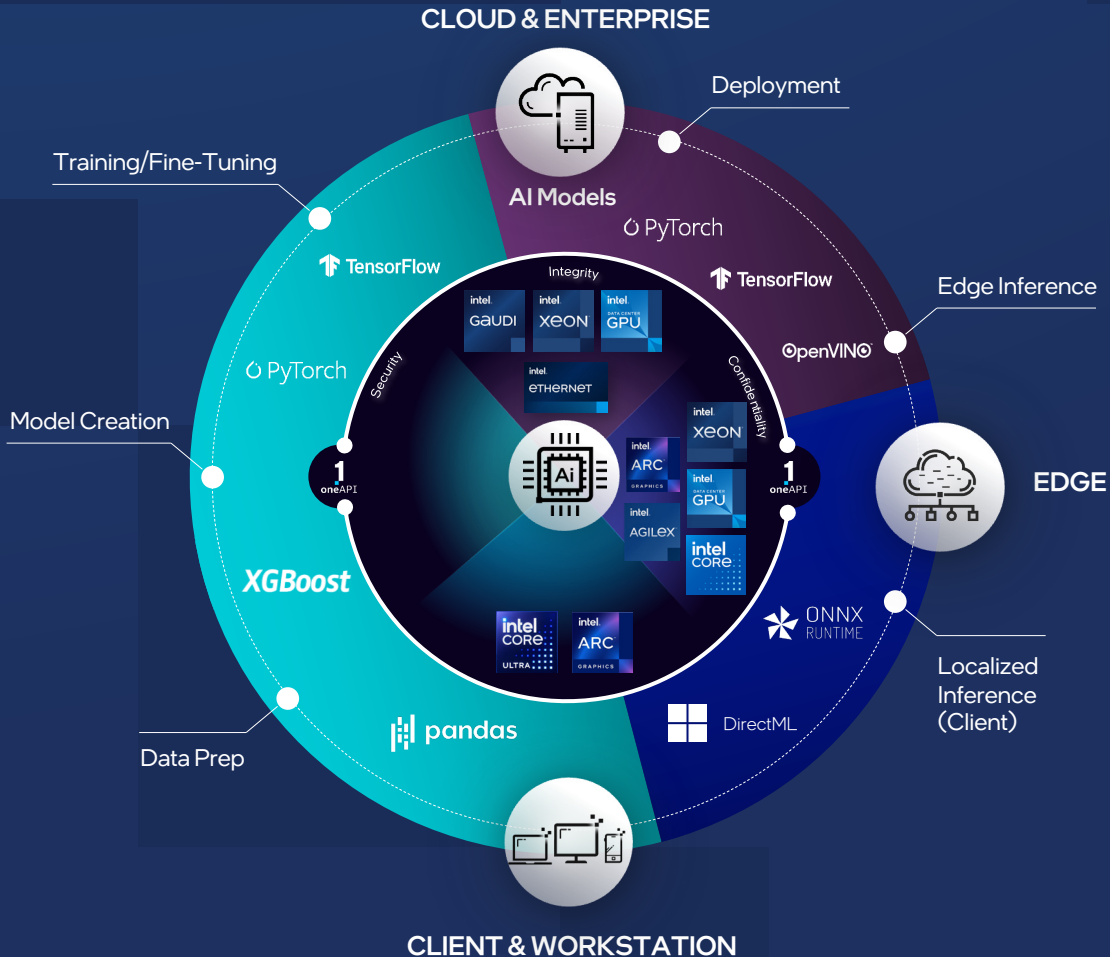
- Intel provides a comprehensive portfolio of AI development software for all your needs, including data preparation, training, inference, deployment, and scaling.
- All Intel AI development software is built on the foundation of a standards-based, unified oneAPI programming model with interoperability, openness, and extensibility as core tenets.
- Developers can implement “write once, deploy anywhere” hybrid AI vision today leveraging Intel’s unmatched installed base and OpenVINO, a model deployment tool built leveraging years of edge AI experience that has seen over 100% year-over-year growth.

Our Approach

- Enterprises should be utilizing an open software stack that performs optimally across all processors and accelerator options.
- The right software helps enterprises easily deploy AI solutions, including GenAI and RAG, across cloud and data center, network, PC, and edge.
- PyTorch is a preferred framework for building LLMs and Generative AI. Intel AI software portfolio is hardware-agnostic, includes enterprise-grade security, networking, and system management features, supports inference based on PyTorch and OpenVINO, supports fine-tuning based on PyTorch, and includes vertical-specific libraries for data retrieval.
- Choosing a pre-trained model means faster time to get insights for your enterprise by quickly fine-tuning it with your specific data. This enables you to quickly finalize a model

for deployment. Intel partners with Hugging Face, which has over 120,000 models to choose from for your needs. Models can be best developed on standard frameworks like TensorFlow and PyTorch.

- Deployment of the inference model requires not just deployment, but optimization and OpenVINO can help with this. Intel OpenVINO Toolkit can be a great asset in deploying vision, GenAI, and NLP models across a variety of deployment targets, on any Intel product from CPU to GPU.
- Our oneAPI programming model serves as the software foundation of our AI strategy. The standard programming model allows developers to get the value across different hardware architectures from one source code. The true democratization of AI allows developers to build and deploy AI everywhere.



Bringing AI everywhere

In a world rapidly reshaped by artificial intelligence, Intel's AI approach aligns cutting-edge technology with the essential pillars of Business, People, and Society. This vision is not just about technological excellence—it's about creating a synergy that propels businesses forward, enhances human potential, and fosters a more equitable society.



Business: Driving Innovation and Value

In the first half of this report, we focused on the enterprise perspective. For businesses, it is all about transforming potential into performance. Our technologies are designed to empower companies to innovate at speed, unlock efficiencies, and create new value in a competitive global marketplace. From start-ups to multinational corporations, Intel provides the tools to build AI solutions that are not just robust and scalable, but also streamlined for ease of use. We're dedicated to helping businesses of all sizes harness the power of AI, turning today's ambitions into tomorrow's breakthroughs.

People: Enhancing Lives and Work

At the heart of every technological leap are the people it serves. Intel's AI is people-centered, aiming to enhance individual capabilities, automate mundane tasks, and offer new ways for personal and professional growth. We

believe in making AI accessible, in creating interfaces that are intuitive, and in ensuring that AI serves as a partner to amplify human potential. Intel's AI PC vision to power over 100 million AI PCs with AI accelerators through 2025 has a key role here. By prioritizing people, Intel's AI initiatives strive to unlock creativity, foster learning, and support more meaningful and rewarding experiences in everyday life.

Society: Building a Better Future Together

The implications of AI on society are profound. Intel is committed to responsible AI that advances the public good, addresses societal challenges, and promotes sustainability and equity. Intel envisions a society where AI acts as a force multiplier for good, driving social progress, enhancing public services, and contributing to a fairer and more inclusive world. Through responsible and ethical AI development, we aim to ensure that the benefits of AI are shared by all.

AI in action

AI is changing everything—from the way we work and run our businesses to how we live and interact in society. The idea of bringing AI everywhere in the enterprise is about giving us tools that understand and adapt to our needs. For individuals, this means AI helps us by giving “augmented intelligence”—it’s like having a superpower that makes us more efficient and our work easier. As we explore the potential of AI through the following personas, we see a future where technology enhances human capabilities at every level.



Dave – The Executive Sponsor

Dave, a visionary C-suite executive, dreams big when it comes to AI. When he wanted to initiate his dream project, he was taken aback by the quotes for the required infrastructure—both the cost and the months-long wait were beyond reason. He realized that their critical project wasn’t a priority for some vendors.

The turning point came with the discovery of Intel’s Gaudi platform, which not only enabled Dave’s team to swiftly commence their project but also offered unparalleled support for both model training and inference. This dual capability of Gaudi significantly accelerated the development and deployment phases, ensuring the project’s timely launch. Dave now looks back with pride and satisfaction, aware that without Gaudi’s advanced solutions, the transformative impact on the company might still be on hold, awaiting the initially ordered infrastructure.



Dr. Amelia - The Scientist

Dr. Amelia, a renowned geneticist, is on the verge of a breakthrough in DNA sequencing that could revolutionize personalized medicine. Working with enormous datasets, the complexity of her simulations had always been a bottleneck. Her AI PC, however, has changed the game. This technological marvel enables her to decode the complexities of human DNA with unparalleled speed and efficiency. The processor’s built-in AI acceleration is instrumental in processing vast datasets, dramatically accelerating her genetic simulations. This capability transforms her dream into a tangible reality, where treatments are uniquely tailored to individual genetic blueprints, and the future of medicine is personalized and precise.

Dr. Amelia’s AI PC is powered with a revolutionary 3D performance hybrid architecture integrating CPU, GPU, and NPU in a single package. Specifically, the NPU’s efficient handling of sustained AI workloads dramatically enhances her genetic data analysis, with performance no longer a constraint but a catalyst.



Sophia - The Business Professional

Sophia, a mid-level manager, navigates the complex world of business with an eye on leveraging AI to enhance her team's productivity while upholding her company's strict security standards. Her AI PC, equipped with advanced security features and capable of parallel processing, becomes a central hub for driving efficiency. Sophia's vision of a dynamic, AI-enhanced workplace that's secure and efficient is now a daily reality, balancing progress with protection.

Sophia's AI PC is powered by the Intel® Core™ Ultra processor, integrating advanced AI acceleration directly into every task. This processor also features cutting-edge security capabilities to detect and neutralize cyber threats. Her AI PC delivers enhanced experiences by automating transcriptions, managing calendars, and taking meeting notes while maintaining the highest data security standards. Its advanced AI acceleration ensures that Sophia's workplace benefits from the latest in AI innovation, keeping her team productive and her company's data secure.



Maya - The Developer

Maya, with her developer's intuition for innovation, finds her sandbox in her AI PC. Here, she experiments, develops, and tests AI applications, pushing the boundaries of what's possible. The system's capability to perform offline AI processing empowers her to build and refine models without constant internet connectivity, allowing her to innovate anywhere.

Maya leverages Intel's oneAPI platform for application development, appreciating the cohesive and unified programming model it provides. With oneAPI x2, she has access to a broad Intel portfolio, including powerful tools and libraries that streamline the coding process across various architectures. This approach has significantly reduced complexity in her development workflow, enabling her to write code once and deploy it across different compute nodes without worrying about the underlying hardware. The ease of cross-architecture application development has made Maya more productive, turning her focus towards innovation rather than compatibility issues.



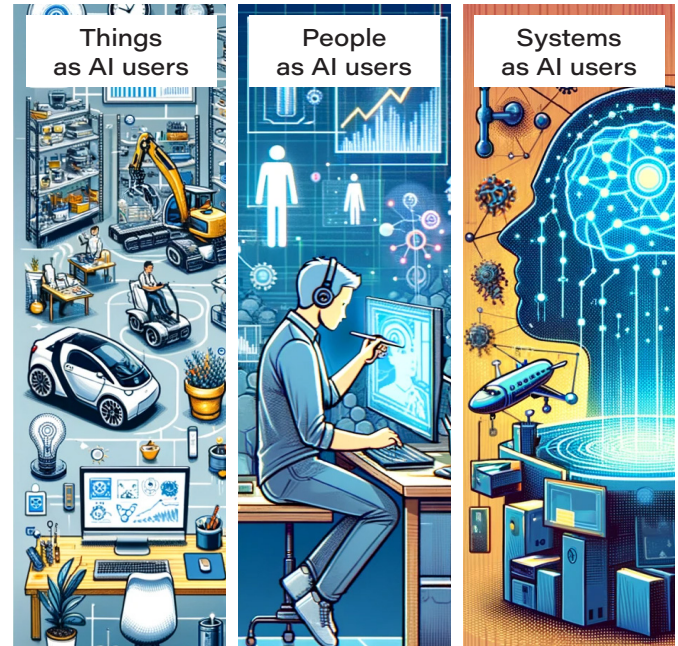
Ali - Data Scientist

Ali, a data scientist, specializes in crafting cutting-edge AI algorithms for natural language processing. His work demands formidable computing power to train intricate models, highlighting the critical role of advanced technology in his research and development efforts.

Intel® Xeon® processors have become indispensable for Ali, furnishing him with a high-performance computing environment tailored for his needs. Thanks to the Xeon CPU's integrated AI acceleration, Ali has witnessed a drastic reduction in model training time—from weeks to mere days—thus accelerating his research cycle. The processors' superior memory and processing capabilities enable him to handle large datasets with greater efficiency, paving the way for quicker scientific advancements.

The future of work

AI is revolutionizing the workplace, heralding a new era where humans, objects, and systems seamlessly interact. As we navigate this transformative landscape, it's clear that AI will have three main users: people, things, and systems. This means your next collaborator could be an algorithm or a smart device, with whom you'll communicate and collaborate. Together, you'll share tasks, leveraging each other's strengths to achieve common goals. This evolution in workforce dynamics promises to introduce innovative job roles and collaborative models, offering both exciting opportunities and challenges in managing and organizing the workforce of the future.



Spotlight: AI PC

Intel powers the AI PC. Now, AI is for everyone in the enterprise.

Productivity is at the heart of every business. With the advent of the AI PC, Intel puts intelligence in the hands of employees everywhere to revolutionize efficiency and creativity. To help create AI PCs that meet the needs of various and unique use cases, Intel has collaborated with more than 100 software vendors to optimize over 300 AI apps and features for Intel® Core™ Ultra processors. Our new 3D performance hybrid architecture integrates CPU, GPU, and NPU into a single package, each designed to handle specific AI tasks.

GPU

High Throughput

Ideal for AI-accelerated digital content creation

NPU

Low Power

Ideal for sustained AI workloads and AI offload for battery life

CPU

Fast Response

Ideal for low-latency AI workloads



Building a Better Future

Society is on the brink of a major shift, with AI at the heart of this change. Intel's vision of Bringing AI Everywhere is a commitment to infusing intelligent, responsive technology into every facet of daily life. This vision extends beyond the personal and professional realms into the broader societal landscape, promising to redefine how we approach education, healthcare, sustainability, and science.

Education

The integration of AI into educational systems is poised to fulfill the long-standing potential of transforming teaching and learning processes. AI PCs with AI-driven software, tuned to individual learning styles and interests, will deliver personalized content, maintain engagement, and provide immediate feedback. This technology will support teachers by assessing students' grasp of subjects and advising on career planning, enhancing the educational experience without replacing the invaluable human connection between teachers and students.


Science

A pivotal moment in scientific discovery has been highlighted in a recent study in Nature magazine, showcasing the monumental impact AI has had on material science. Traditionally, the search for novel crystal structures involved time-consuming and costly trial-and-error methods. However, computational approaches over the last decade have led to the discovery of 28,000 new materials. The recent integration of AI-guided methods has shattered previous limitations, with the discovery of 2.2 million materials, equivalent to about 800 years of knowledge. These advancements underscore the unparalleled scale and precision AI brings to scientific exploration, opening new horizons for innovation and application.

Can protein engineering be moved to the computer?

Proteins are the engines of all organisms, from bacteria and plants to humans. They make all the amazing phenomena of what we call life. But the problem is the proteins are natural. They are not optimized for industrial use. A single protein can take more than a year for companies to engineer. It costs millions of dollars. It's a very complex process. The number of combinations is extraordinary, 20 to the power of 300, and the most painful part is it just very often fails. So now, imagine that you can make a few tweaks to these proteins and then make them do exactly what you need. Scala Biodesign, one of the outstanding participants in Intel's Ignite program, wants to engineer the proteins at a fraction of the time and cost. They are moving protein engineering from the lab to the computer to make this protein engineering project succeed. Imagine a world where Therapeutics and vaccine development become much easier and simpler. Imagine a world where producing food will be much more sustainable and economical.

<https://www.scala-bio.com/>



Healthcare

Through advanced data analysis, AI is enabling faster discovery of treatment pathways, drug design, and side effect prediction, offering hope for cures to many incurable diseases. An AI-led technological leap forward will also accelerate medical breakthroughs, including personalized medicine tailored to an individual's genetic makeup. This fusion of AI with healthcare enables a future extending quality of life where longevity is not just a possibility but a reality.

Smart cities

Smart cities, powered by AI, envision urban environments that are more livable, efficient, and sustainable. AI-driven traffic management systems reduce congestion and pollution, while smart grids optimize energy consumption to lower carbon footprints. In public services, AI enhances safety through predictive policing and emergency response systems. These intelligent infrastructures not only improve the quality of urban life but also pave the way for cities that grow in harmony with the planet.

Can AI also shape the future of internet usage?

By harnessing the power of AI, even the future of internet usage can be reshaped. One of the standout stars of Intel's Ignite program, DeepRender, is using AI to tackle one of the digital world's growing pains: internet bandwidth issues. As we stream, browse, and download more than ever, the bandwidth strain becomes increasingly noticeable. DeepRender has developed an innovative method that uses artificial intelligence to significantly compress data without losing quality, providing a smoother online experience. With this technology, for example, a high-resolution movie we stream can be compressed for transmission and then restored to its original quality on an AI PC, eliminating the need for heavy data traffic over the internet.


<https://deeprender.ai/>

Spotlight: Intel's Edge Solutions

Empowering real-time decisions

Intel's Edge solutions represent a suite of technologies designed to bring computing closer to the source of data generation, optimizing response times and reducing bandwidth usage. These solutions empower businesses to process and analyze data in real-time, directly where it's collected—be it in retail environments, manufacturing facilities, or urban infrastructure. By leveraging powerful Intel processors and AI capabilities, Edge solutions enable smarter decision-making, enhanced operational efficiency, and improved customer experiences. Intel provides the hardware, software, and frameworks necessary to deploy applications across various industries, ensuring seamless integration and the ability to run sophisticated AI models at the edge. This approach not only accelerates innovation but also offers the flexibility to address complex challenges in an increasingly connected world.

[Get started with Intel AI Edge Reference Kits](#)



To start your AI journey with Intel visit, www.Intel.com/AI

1. Based on internal test results conducted by Samsung Medison. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Results may vary. For details, see: intel.com/content/www/us/en/products/details/embedded-processors/core-ultra.html
2. For details on Intel® Security Solution for Fortanix Confidential AI, see: intel.com/content/dam/www/central-libraries/us/en/documents/2022-05/dpg-sgx-fortanix-snapshot.pdf
3. Based on internal test results conducted by Hellometer. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Results may vary. For details, see: restaurantdive.com/press-release/20230126-hellometer-partners-with-intel-to-revolutionize-the-fast-food-experience-1/
4. Based on internal test results conducted by Siemens. See N21 and N22 at intel.com/processorclaims: 4th Gen Intel Xeon Scalable processors. Results may vary.

Availability of accelerators varies depending on SKU. Visit the [Intel Product Specifications page](#) for additional product details.

Performance varies by use, configuration, and other factors. Learn more at 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.

Your costs and results may vary.

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

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

© 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.

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.

