it starts with

intel

Building a Secure and Resilient Education Infrastructure in the Age of AI

As artificial intelligence (AI) gains momentum, universities and school communities are discovering myriad ways to benefit from AI-powered tools. Yet, for all its upside, integrating AI into an educational environment can bring many challenges. This is particularly true when it comes to AI infrastructure planning. Developing a successful strategy calls for collaboration and clear communication about needs and goals among educators, administrators, and IT. It also requires careful evaluation of many factors and foresight about scenarios that may lie ahead.

For IT, the task requires balancing multiple priorities. Maximizing investment value is crucial, but IT must also consider how best to protect data against ever-more sophisticated cybersecurity threats. Beyond infrastructure choices, there's a need to create avenues for implementers, students, and educators alike to learn AI skills and keep those skills up to date.

In this eBook, we draw on Intel's long experience with AI and technology innovation to identify key considerations in building a robust, scalable infrastructure. We suggest a phased approach that starts with discussing where to place AI workloads. It continues with assessing hardware needs for specific AI solutions to find the right balance between cost and performance, followed by identifying software resources that simplify deployment. In each phase, protecting systems and data underpins planning—it is always top of mind and possible only with technology you trust.

We encourage you to engage with educators and tech experts as you strategize. Their perspectives enrich your planning with valuable context and clarity about AI's opportunities, risks, and challenges in education.

Please view this eBook as a guide. My team and I hope you find it helps you identify the right choices to meet your organization's unique needs.

Cigdem Ertem
Global Director
Public Sector and Education Sales

Intel

Where Education Organizations Are in Their AI Journey

Artificial Intelligence (AI) is introducing new ways to teach, manage, and learn in educational institutions around the world. Educators have been among the first to explore AI's potential. Yet, most schools and school districts are still in the early stages of their AI journey. As the number of AI-based apps and use cases grows, interest is strong, and IT decision-makers are considering the path ahead. It's a challenging task because while adopting AI in education creates infinite opportunities, it prompts almost as many choices. One thing is certain—implementing AI in education responsibly, efficiently, and cost-effectively will require careful infrastructure planning.

In our previous eBook, *Getting Real About Artificial Intelligence in Education*, we offered examples of AI-based apps already in use. These include AI assistants that help gauge student progress and provide real-time feedback or tutor students and answer their questions anytime. They also include apps that automate routine tasks such as taking attendance and scoring exams and essays. Such innovative tools free up teachers to focus on what matters most: developing engaging lessons, building meaningful relationships with students, and giving constructive feedback.

Apps also are available to assist with administrative tasks. These tools automate student registration and attendance, assist with scheduling, optimize physical resource allocation, and more—again, allowing increased time for higher-value activities.

Making the most of such AI-enabled opportunities calls for setting up environments in which to teach AI skills. These AI labs will need to adapt and scale over time to keep up with the rapid growth and evolution of AI usage.

For all the potential benefits, integrating AI into an educational environment brings a host of complex requirements. From an infrastructure and operational standpoint, there are many factors to evaluate and scenarios to anticipate. Determining how to maximize the value of your investment is a top priority—but far from the only one. For example, IT must also consider how best to safeguard student and staff data in the face of sophisticated AI-powered cybersecurity threats and how to address the IT skill gap.

As with the adoption of any new technology, there are concerns about AI that need to be addressed. A recent survey of universities found that ethical issues and the complexity of integrating AI into existing systems present the greatest challenges.[1]

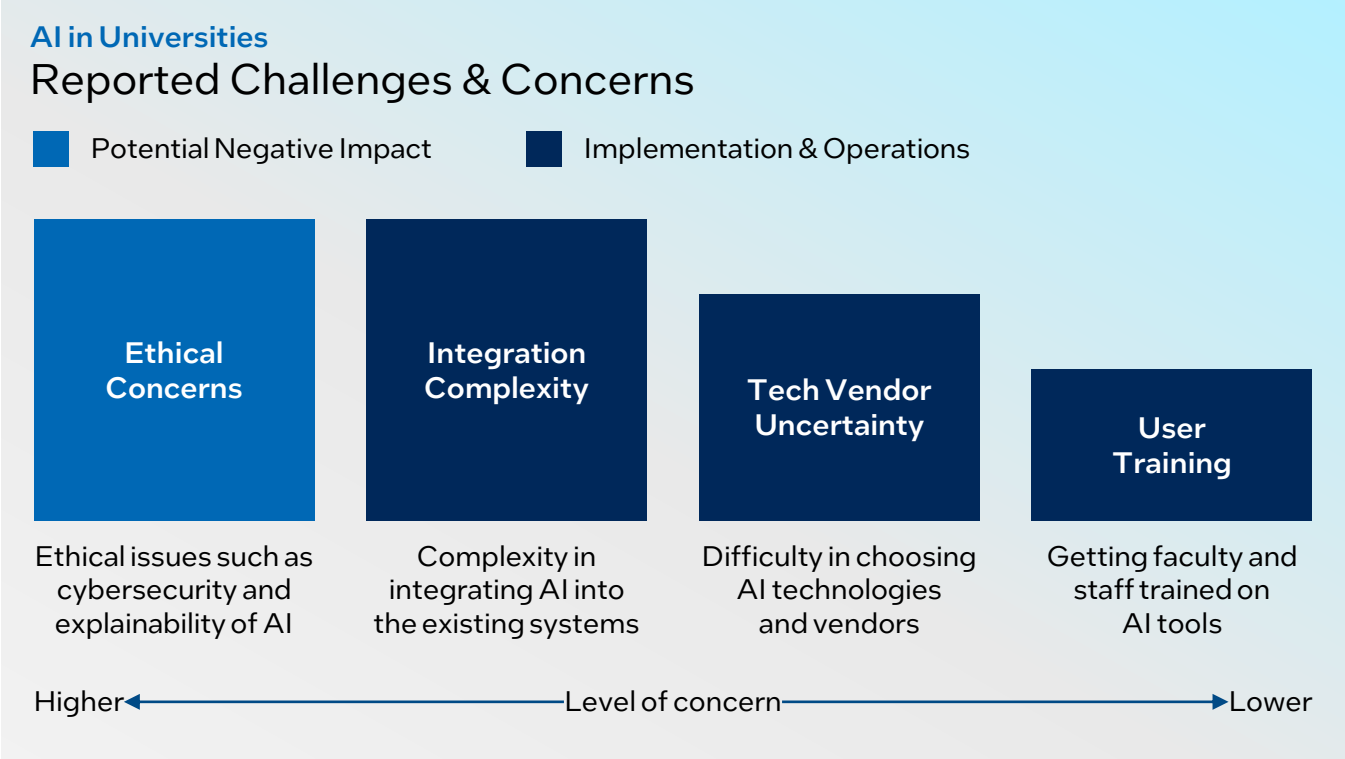

With a legacy of AI and technology innovation and strong partnerships, Intel is uniquely positioned to help educational institutions meet these challenges. This eBook represents one step in our overall effort. In it, we review key IT infrastructure considerations that will assist education IT decision-makers in making optimal decisions to fit their specific circumstances.

Key Considerations

The right infrastructure support is essential to successful AI adoption. AI workloads can put considerable strain on technology infrastructure, and education environments have unique considerations.

Strategizing to create a scalable and robust system begins with a discussion of where to place AI workloads—in the cloud, on-premises, locally on devices, or in a hybrid environment. Next comes an assessment of hardware needs for various AI solutions with the goal of finding the right balance between cost and performance. Then, identifying and leveraging software resources can ease and accelerate the path to deployment. Finally, in every phase of planning, security is foundational. Protecting critical systems and data is paramount—and possible only with technology you trust. To assist with the planning process, the discussion that follows explores these considerations in greater detail.

Infrastructure and Scalability

Different environments for AI workloads—in the cloud versus on-prem versus local, on-device—offer different advantages and disadvantages. A careful analysis, in the context of an educational institution's specific needs, informs the decision about workload placement.

Running AI workloads in cloud environments helps schools, districts, and universities access AI capabilities more quickly and cost-efficiently. These companies have invested heavily in developing their cloud AI platforms—giving customers the ability to deploy AI solutions without spending large amounts on infrastructure and personnel.

From Cloud to Edge – Optimizing AI Workloads

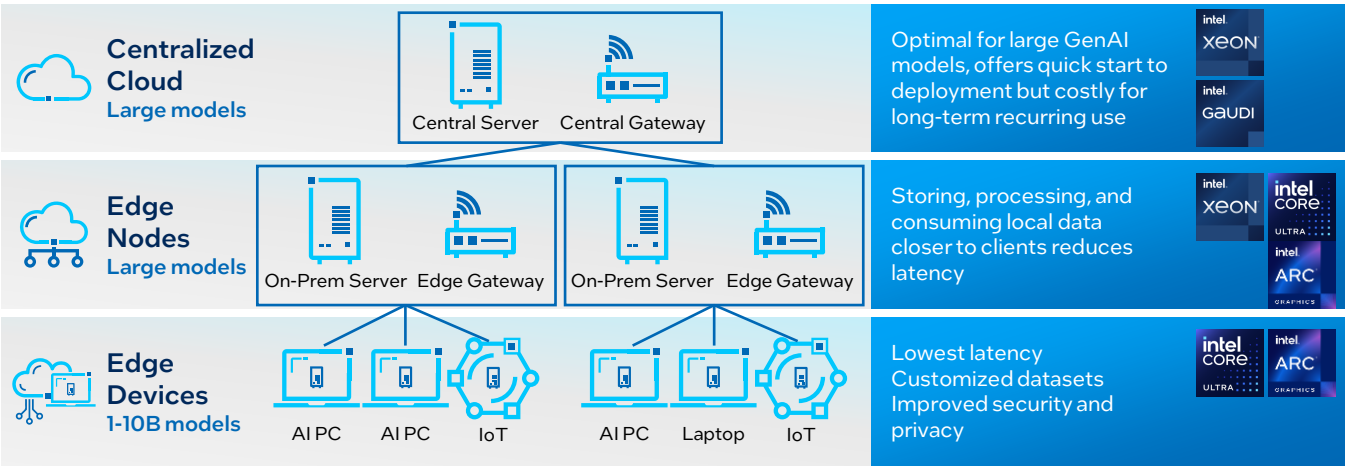

As education IT teams explore their cloud AI possibilities, they can use solutions from Intel such as Intel® Xeon® Scalable processors with built-in AI accelerators. They can also consider solutions such as those with Intel® Gaudi® 3 AI accelerators in IBM Cloud and Denver Dataworks.

A helpful new option is the ability to place AI workloads directly on innovative AI PCs powered by Intel® technology. Leveraging the PC as a local device to support AI workloads offers several distinct advantages that are not available in the cloud. For example, lower latency and swift response times facilitate more real-time AI experiences and allow for AI applications to run even when the PC is offline or not connected to the internet. Processing data locally also enables an additional layer of privacy and security. Similarly, customizing models using data specific to the user creates a possibility for increased personalization and customization. AI PCs allow for easier scaling and ease of learning with many pedagogies giving more options to learners and educators. All this comes with the potential for cost savings, in the form of the lower SaaS costs achievable as AI and machine learning become ubiquitous in client apps.

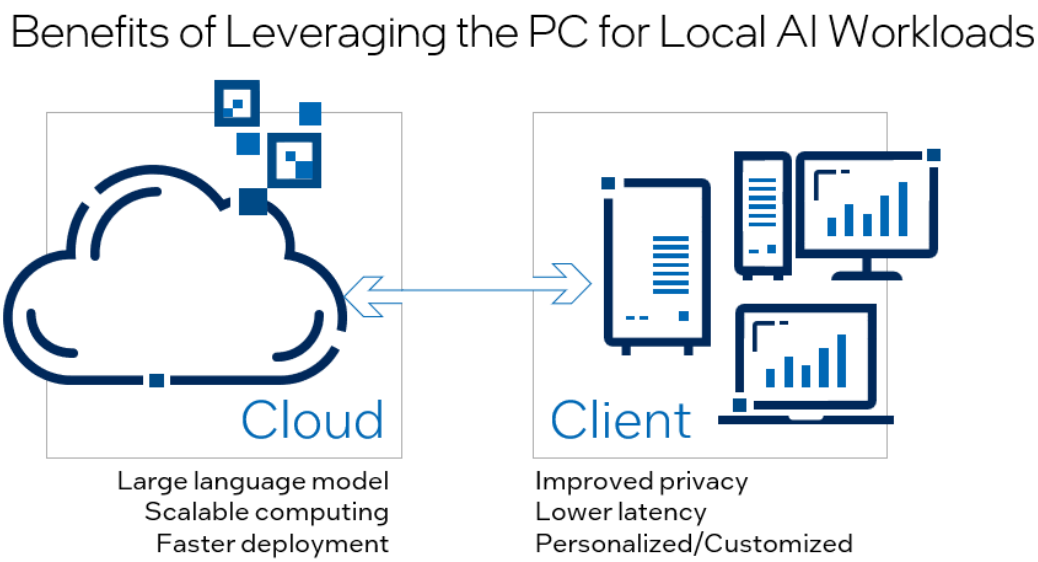

Extending AI to the PC opens other new opportunities as well. Applications that use any combination of models—cloud, client, or **hybrid AI**—can better respond to user needs and increase the efficiency of the virtuous operations cycle to enable more predictive, proactive, and self-healing solutions.

Assessing End-to-End Hardware Needs for AI Solutions

When it comes to AI hardware, balancing cost and performance can be a challenge. The hardware requirements for any given AI solution vary based on factors including workload complexity, performance expectation, and deployment location. Various phases of the AI pipeline also introduce different computing needs. To help determine the right fit for a specific deployment or hardware offering, here are compiled guidelines that span today's high-priority AI use cases.

Computer vision systems strengthen school security, inside and outside, by scanning and analyzing video footage to detect suspicious activity or abnormal events. They can also be used for parking management and real-time campus mapping, including people counting, for more efficient use of services such as building heating and cooling.

Learn More

Machine learning (ML) can identify skill gaps and provide personalized learning to meet the needs of individual students. By automating administrative and routine tasks, ML also boosts efficiency and accelerates data-driven decision-making.

Learn More

Generative AI (GenAI) makes it easier for educators to create engaging content and materials. By providing instant access to vast amounts of information, it can help students find further clarification of topics covered in class.

Learn More

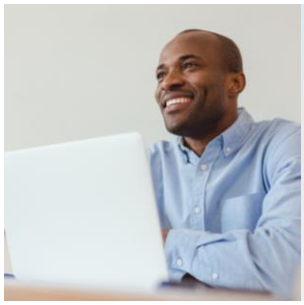

Recommendation systems can be trained to consider a student's behaviors and preferences when recommending learning materials. They have the potential to provide a personalized learning path in which the system adapts to the student's learning requirements and abilities.

Learn More

The power of AI is fully realized when hardware and software solutions are paired. Computer vision combined with GenAI can document team project videos, help with presentations generated from videos, and inspire further content creation with video, painting, and other media brought together.

Take advantage of integrated accelerator engines

To help optimize cost, performance, and energy consumption, Intel® Xeon® Scalable processors come equipped with integrated acceleration capabilities. These accelerators make them well-suited for many AI-related tasks across the pipeline, from training to deployment. Intel collaborates with AI learning software vendors, enabling them to take advantage of these capabilities right out of the box to boost performance and accuracy for educators and students.

Find out more about Intel® accelerator engines for AI:

- Intel® Advanced Vector Extensions 512: Accelerate classical machine learning
- Intel® Advanced Matrix Extensions: Accelerate deep learning training and inferencing workloads, including natural language processing (NLP)
- Intel® Deep Learning Boost (Intel® DL Boost): Expedite answers with smarter math for neural networks

Intel® hardware for the AI pipeline

Using Intel® AI hardware, IT can speed and run AI workloads on their preferred platform for optimal cost efficiency.

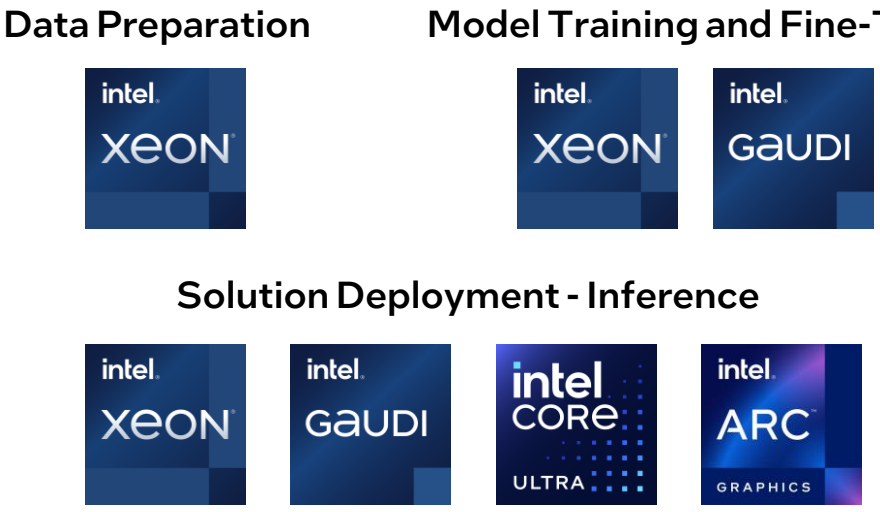

Intel® Gaudi® AI accelerators

Intel® Gaudi® AI accelerators and Intel® Gaudi® software are designed to bring a new level of computing advantages and choice to training and inference, whether in the cloud or on-premises.

We aim to make the benefits of AI deep learning more accessible to more academic institutions, removing barriers to bring deep learning advantages to students, educators, and administrators everywhere.

AI PC powered by Intel®: Why it is important

While cloud processing is optimal in certain circumstances, many have concerns about the fact that, currently, AI applications for education reside primarily in the cloud. In addition to raising data security concerns, this solution also limits access for students and educators who need AI-computing capabilities when offline.

Cognizant of these limitations, and in the interest of furthering both equity and access, Intel brought more than ten years of R&D to bear and developed a new architecture that marries high performance with the flexibility to run AI workloads in the cloud or locally on the PC. The AI PC powered by Intel® is any PC running on the latest Intel® Core™ Ultra processors (Series 2) with built-in AI capabilities. This new generation of PCs brings computing to a new era of supercharged speed, efficiency, creativity, and productivity while maintaining the highest standards for protecting privacy and security.

The AI PC shifts some workloads traditionally run in the cloud to your PC. Moving these workloads to the AI PC means it's more responsive to the needs of education IT teams and users. The shift reduces processing times, maximizes flexibility, and provides additional user data protection because the data is hosted on the user device, which reduces the need to send personal information through the network and helps in supporting greater data geo-compliance requirements.

Not only can the AI PC offload workloads from the cloud, but it can also decide which of its three engines should process each workload to give the user the best performance. The central processor (CPU) provides a fast response and works well for smaller workloads that don't require much computing power. The graphics processor (GPU) specializes in processing large workloads that require high throughput, and it can process multiple tasks on parallel paths. The neural processing unit (NPU) is ideally suited for parallel tasks and power-sensitive workloads and can help save energy.

intel CORE ULTRA

The Right Balance of Power and Performance for Modern Computing

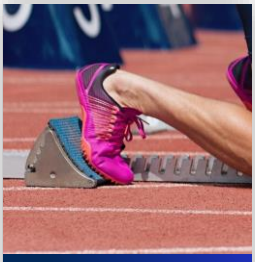

CPU
Fast Response

Ideal for lightweight, single-inference, low-latency AI tasks

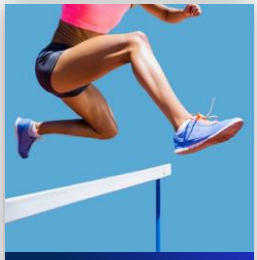

GPU
Performance Parallelism & Throughput

Ideal for AI infused in Media/3D/Render pipeline

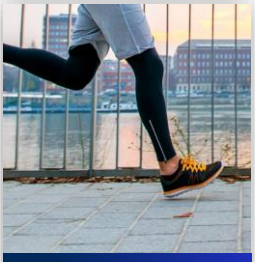

NPU
Dedicated Low Power AI Engine

Ideal for sustained AI and AI offload

Jump-start AI Development

Embrace the open-source opportunity

As with hardware, software changes can help educational institutions maximize the benefits of AI,. Identifying what changes to make—and how—may initially seem challenging. The good news is that open-source models and tools can be used to underpin analytics software like Open Platform for Enterprise AI (OPEA), where modules such as document summarization and chatbot are already developed and student can learn by pointing it to their data making it easier to understand and adopt. Intel has invested heavily in the open-source community to help facilitate AI success and provide optimizations for popular open-source machine and deep learning frameworks.

Maximize results at the edge

As organizations continue implementing AI and deploying it throughout their operations, AI at the edge will become even more critical. Intel offers a range of software resources that can help deliver outstanding results at the edge, including the Intel® Distribution of OpenVINO™ toolkit and ready-made oneAPI toolkits. OpenVINO's sophisticated optimization framework maximizes performance through Intel GPU XMX acceleration and NPU capabilities, allowing maximum flexibility for developers through intelligent workload distribution across architectures.

Making Generative AI More Accessible for Real-World Scenarios

OpenVINO™ toolkit is an open-source toolkit that accelerates AI inference with lower latency and higher throughput while maintaining accuracy, reducing model footprint, and optimizing hardware use. By allowing developers to write once and deploy anywhere, it streamlines AI development and integration of deep learning in domains like computer vision, large language models (LLM), and generative AI.

Remember, you don’t have to start from scratch

Similar to how open-source components help accelerate development, pre-built toolkits and reference deployments can shorten time to value. Examples of these ready-to-use solutions include the AI tools and reference kits available from Intel.

Optimize your path to production

With Intel® Tiber™ AI Cloud, speed up AI development using Intel-optimized software on the latest Intel® Core™ Ultra processor, Intel® Xeon® processor, Intel® Gaudi® AI accelerator, and GPU compute. Intel’s cloud allows developers to develop and optimize AI models and applications, run small- and large-scale training and inference workloads, and deploy with the best price-performance.

Security

Beyond hardware and software choices, security remains a key concern when deploying AI applications. Protecting critical systems and data is possible only with technology you trust. Intel products are designed, manufactured, and maintained using the industry's best security practices. Our innovative security capabilities and robust processes to find and address vulnerabilities are designed to help you meet today's security challenges.

Intel leads the **silicon industry in product security assurance**.[2] Not only do we implement industry best practices, but we also help develop and define them through our research investments and involvement with industry committees. We continually raise the bar on platform security by delivering a more secure platform for the entire device lifecycle through design, vulnerability monitoring, transparent communication, and proactive patching.

Secure by design for modern computing and AI

The Intel vPro® platform meets IT goals by focusing on these three device security priorities to help keep your data secure throughout its journey from edge devices such as PCs to the data center and the cloud:

Protect: Help protect data and applications against modern-day advanced threats.

Detect: Improve real-time threat detection and prevention efficacy using innovative methods, including memory scanning or behavioral tracking.

Recover: Use modern remote manageability tools to return systems to a known good state when the unexpected occurs.

As security threats become more sophisticated, the need for advanced detection and protection technologies also grows. The Intel vPro® platform takes a holistic approach to security, delivering protection at all layers.

Security for AI: Extend peace of mind for AI PC applications by tracking their use, setting policies, and avoiding data poisoning or model theft. The Intel vPro® platform provides a more secure computing foundation for AI apps and data with a robust set of security features that reduce the attack surface by 70 percent compared to older devices.[3] Its multi-layer hardware-based security strategy includes advanced threat protections to detect file-less malware attacks to memory, ransomware, and crypto-jacking threats.

Specifically, Intel® Threat Detection Technology (Intel® TDT) is the only AI-based silicon security supported across a billion PCs that leverages silicon telemetry and provides AI-enhanced defense against sophisticated attacks. Results are proven. SE Labs reported that devices with Intel® Threat Detection Technology (Intel® TDT) helped Endpoint Detection & Response (EDR) software detect attacks 24 percent more effectively than devices without Intel TDT.[4]

In addition, Intel vPro® application and data protections help prevent attacks from traditional malware against applications, login credentials, and data, while below-the-OS security technologies help prevent attacks against firmware and help protect apps from attacks that run through the BIOS.

AI for Security: Intel® Core™ Ultra processors extend the NPU capabilities to third-party security software vendors, enabling applications to become more effective and efficient at detecting and responding to threats. Leveraging the NPU also enables or enhances security use cases like deepfake detection, data loss prevention, anti-phishing, and fileless attack detection—all possible thanks to Intel's XPU strategy.

35% reduction in time spent on security breaches[5]

At Intel, our vigilance toward preventing the exposure of sensitive, proprietary, or confidential data extends well beyond our multi-layer security strategy. We are also working on Confidential AI, an effort to protect data and meet security, privacy, and compliance standards.

What is Confidential AI?

Confidential AI exists at the intersection of artificial intelligence (AI) and confidential computing, bridging the gap between Zero Trust policies designed to secure private data and generative AI, which often relies on cloud compute power to be trained and process complex tasks and requests. For organizations to trust in AI tools, technology must exist to protect from exposure inputs, trained data, generative models, and proprietary algorithms.

Confidential AI utilizes confidential computing principles and technologies to help protect data used to train LLMs, the output generated by these models, and the proprietary models themselves while in use. As processing moves more from the cloud to the edge with hybrid AI, confidential computing allows applications to access data stored on local devices without having the data travel to a cloud or edge server. Through vigorous isolation, encryption, and attestation, confidential AI prevents malicious actors from accessing and exposing data, both inside and outside the chain of execution.

Intel's approach to Confidential AI

AI will only be genuinely accessible to everyone when it is built ethically and responsibly. Intel collaborates with technology leaders across the industry to deliver innovative ecosystem tools and solutions that will make using AI more secure while helping organizations address critical privacy and regulatory concerns at scale.

To help you address the new security realities of education AI, Intel offers the most comprehensive confidential computing portfolio in the industry today. Key technologies include:

- Intel® Software Guard Extensions (Intel® SGX), which enables you to create data enclaves to ensure data is not mixed and remains secure when used for new reporting or learning initiatives.
- Intel® Trust Domain Extensions (Intel® TDX), which increases confidentiality at the VM level, enhances privacy, and allows you to gain control over your data.
- Intel® Trust Authority (now Intel® Tiber™ Trust Services), which makes it possible to take confidential computing to the next level with a zero-trust attestation SaaS that verifies the trustworthiness of compute assets at the network, edge, and in the cloud.

Intel's multi-layer approach to security helps defend against modern threats at each layer: hardware, BIOS/firmware, hypervisor, VMs, OS, and applications. Our platforms and technologies drive the convergence of AI and confidential computing, enabling educational institutions to secure diverse AI workloads across the entire ecosystem. That's why you can count on Intel as a trusted partner to help protect students, educators, administrators, and their data.

Developing AI Literacy

The success of Intel's effort to bring AI everywhere in education depends, in part, on building educators' and students' technology skills. To that end, we have established programs such as the oneAPI Academic Program for university professors and the Intel® Edge AI Certification Program for college-level computer science students. Both programs focus on fostering development skills.

To further encourage development, we make available pre-built toolkits such as Intel® Distribution of OpenVINO™ toolkit and Intel® Gaudi® AI accelerator. These tools help speed the development process while ensuring that the end result meets users' specific needs.

oneAPI Academic Program

Intel's oneAPI Academic Program provides university professors around the world industry-shaping concepts that can quickly and efficiently be integrated into STEM courses. The program offers a wide range of educational resources—including teaching kits, access to the Intel® Tiber™ AI Cloud, and support from Intel for workshops and hackathons—for AI, high-performance computing (HPC), and heterogeneous parallel programming. Students benefit from early access to the latest technological developments, invitations to exclusive events and trainings, and a real-life compute environment to conduct class assignments and research. More than 75 universities globally currently participate.

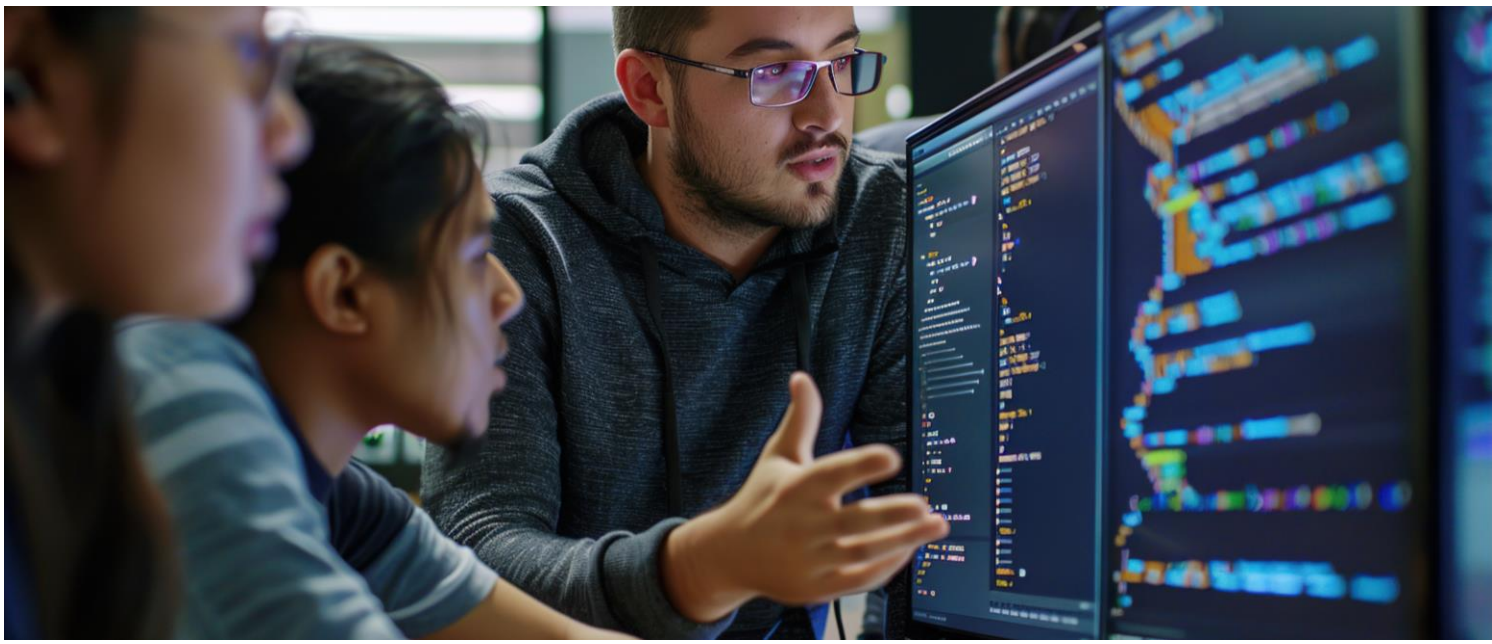

Intel® Edge AI Certification

This self-paced, online training course includes virtual classroom instruction and hands-on projects for computer science students to learn how to use the latest Intel developer tools and platforms to create their own portfolio of edge AI solutions.

These programs are building equity on a global scale and helping local communities develop high-demand skills that prepare their students to succeed in the jobs of tomorrow.

Of course, skill training is not just for developers and coders. At the K-12 level and beyond, we have collaborated closely with educators and education leaders around the world to develop programs that improve digital skills for a wide range of students and teachers.

AI for Youth

The Intel® AI for Youth (AI4Y) curriculum equips K-12 students with technical knowledge and social skills required for a digital world. Through more than 200 hours of hands-on content including computer vision, natural language processing, and data analytics, the program enhances AI awareness while demystifying and democratizing the technology. Teachers can design activities based on their own pedagogy and preferred method of delivery, using AI to enrich the learning experience for students with different styles of learning. Currently, over half a million students in more than 25 countries are using AI4Y. A sister program, Intel® AI for Future Workforce, is helping students in 113 community colleges and 500 universities and vocational schools in more than 10 countries prepare to join the digital workforce.

Intel® Skills for Innovation

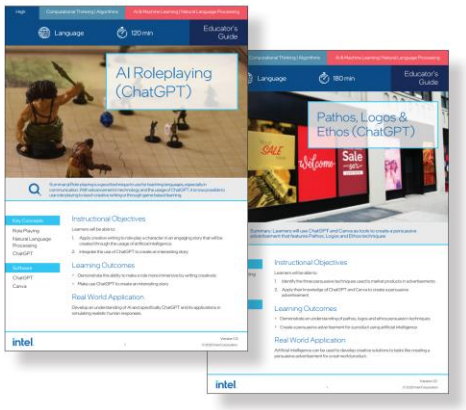

Intel® Skills for Innovation (Intel® SFI) Professional Development empowers educators as they assume their role as leaders of learning experiences and facilitators of future-ready skill-building. Intel® SFI modules and courses help educators understand new skill requirements, create action plans, test the viability of learning environments, develop competencies, and adopt technology-supported, skills-based learning models across the entire education ecosystem.

24 Starter Packs with AI-infused content engage educators and make it easy for them to learn as they tackle course design, multidisciplinary collaboration, and lesson planning for elementary, middle, and high school. Intel® SFI is being used in 150 countries by 105 partner organizations. More than 140,000 educators are engaged with Intel® SFI, and more than 750,000 students have used the program to date.

What's Next

Our collective work toward ethical, responsible AI implementation and skills development continues. While we are just at the beginning of the journey, it's clear that AI will play an important role in the path ahead. The goal for education IT teams is to provide the robust, scalable, cost-effective IT infrastructure needed to harness AI's power to advance learning and academic research.

AI workloads place significant demands on IT infrastructure in terms of performance, security, and data governance. That's why careful planning is essential. Identifying in advance the components of a well-designed, resilient, open infrastructure that can scale flexibly helps ensure educational institutions optimize their investments, now and in the future.

Intel stands ready to help you meet these challenges. With our vast ecosystem of ISVs, OSVs, and software vendors, we have the scale, expertise, and resources to meet your unique needs. We encourage you to review the recommendations on the next page and then contact your local partner to consult on the right infrastructure strategy for your organization.

Delivering responsible AI

Intel is committed to evolving best methods, principles, and tools to ensure responsible practices in AI product development and deployment. Intel collaborates with academia and industry partners to advance research in this area while also evolving platforms to make responsible AI solutions computationally tractable and efficient. Intel's mission is to help equip your organization with AI that delivers valuable results while being equitable and ethical. Learn more about our approach to responsible AI by visiting our website and downloading our eBook, *Getting Real About Artificial Intelligence in Education*.

7 Steps to Help You Implement AI Responsibly

At Intel, we understand that technology enables progress, but it's up to educators and education IT decision-makers to adopt processes, controls, and workflows that maximize results and minimize risks. We have developed the following seven essential elements to help you implement AI responsibly wherever you are on your AI journey.

Recommendations for Getting Started

Establish a Responsible AI workgroup or committee

Establish a multidisciplinary Responsible AI workgroup or committee to evaluate technology tools/options.

Encourage pilots

Validate tools, technologies, and programs within a pilot group to eliminate kinks before rolling out on a larger scale.

Provide professional development for educators and staff

Provide professional development for educators and staff to unlock the full potential of AI and other digital tools.

Celebrate success

Celebrate success stories to encourage adoption of methods that lead to superior outcomes.

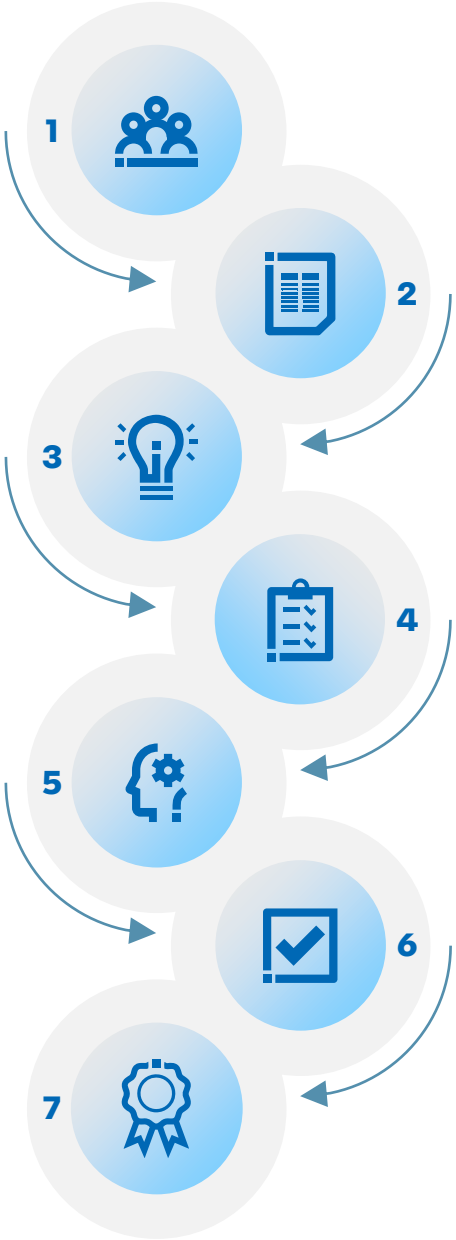

Publish AI policies and guidance

Publish AI policies, standards, ethical guardrails, processes, and guidance to keep everyone on track and collaborating effectively.

Audit your IT infrastructure

Audit your IT infrastructure with future AI-based use cases in mind.

Review policies regularly

Review and update policies regularly to address issues and essential changes.

Learn More Intel's Solutions for AI in Education:

- Intel.com/education
- Intel.com/ai
- Intel.com/aipc
- Intel® Tiber™ AI Cloud
- Intel® Distribution of OpenVINO™ Toolkit Coursework
- Intel® Tiber™ AI Cloud Training
- Intel® Digital Readiness
- Skillsforinnovation.intel.com
- Intel Responsible AI Program

Endnotes

[1] Digital Education Council. "Embracing AI: What Are University Leaders Concerned About?" March 25, 2024. https://www.digitaleducationcouncil.com/post/embracing-ai-what-are-university-leaders-concerned-about

[2] Ranked No. 1 on Forbes' 2023 "America's Most Cybersecure Companies" list and recognized by ABI Research for Product Security Assurance. https://www.intel.com/content/www/us/en/security/security-as-a-component-of-tech.html

[3] Based on IOActive's "Intel vPro 13th Gen Attack Surface Study" published March 2023 (commissioned by Intel), which evaluates Intel vPro devices powered by 13th Gen Intel Core processors against 4-year-old Intel-based PCs. Additional details at www.intel.com/performance-vpro. No product or component can be absolutely secure.

[4] Based on SE Labs "Enterprise Advanced Security (Ransomware)" – Intel Threat Detection Technology study published March 2023 (commissioned by Intel) analyzing ransomware detection rates with Intel TDT versus non-hardware-based solutions. Details at www.intel.com/Performance-vPro. Results may vary. https://selabs.uk/reports/enterprise-advanced-security-ransomware-intel-threat-detection-technology-2023-02/

[5] Based on "The Total Economic Impact™ of the Intel vPro Platform," an Intel-commissioned study by Forrester Consulting, January 2024, which surveyed 500 ITDMs at enterprises across the world using Intel vPro®, including US, Canada, France, Germany, UK, Australia, China, India, and Japan. For the study's findings, Forrester aggregated the data and experiences from the interviewees into a composite organization with an assumed revenue of $1 billion per year and 10,000 employees.