


Lenovo and Intel: Building a responsible future for AI

Ever since the concept was first introduced, artificial intelligence (AI) has been hailed as a transformative technology that will reshape humanity's future for the better in ways we can't even imagine. And along with those aspirational promises came apocalyptic warnings of how AI could end up taking over our entire existence once it was unleashed.

We're still far from seeing the total potential AI has to offer. But as AI continues to grow and evolve, we've gotten a look at both good and bad examples of what this technology is capable of.



A delicate balancing act

Alongside all the amazing process improvements and added value that AI has delivered in business, healthcare, manufacturing, and many other applications, there have been examples of discriminatory screening processes for hiring and uses of facial recognition that violate privacy laws.

What is certain is that AI will continue to become a larger and larger part of our everyday lives.

The question then becomes: How do we ensure we fulfil the promises and avoid the negatives?



92%

of executives say that GenAI is key to reinventing at scale and at speed.¹

As a leading developer of AI technology, Intel is helping organizations build a responsible future for AI together.

“Intel and Lenovo share a common vision of bringing AI everywhere. With Xeon® 6 processors, we are helping businesses use the power of AI to capture the value of their proprietary data and generate real-world business outcomes.”

JUSTIN HOTARD

EVP and General Manager of the Data Center and AI Group, Intel



Lenovo ThinkSystem SC750 V4 Neptune®, powered by the Intel® Xeon® 6 processor — performance and efficiency without compromise.

**Smarter
technology
for all**

Lenovo

The rise of responsible AI initiatives

As the use of AI grows across multiple industries, so does the concern about its responsible use.



73%

of business leaders say AI guidelines are indispensable ...



But only **6%**

have established them in their companies.²

Many businesses like **Google** and **Microsoft**, along with international organizations including **UNESCO** and the **European Union**, have put together guidelines for the responsible development of AI initiatives. The US President also issued an **executive order** on the topic. All these guidelines recognize that meaningful implementation of the principles would have to be built from the bottom up. That is, consideration for the responsible and ethical use of AI starts with how the AI is built and trained.

While the categories may differ slightly from one list to another, they tend to cover the same areas of concern. At Lenovo, we've organized responsible AI principles into six pillars you'll find on the following pages, along with some questions you should ask both yourself and your development partners as you embark on any AI initiatives.



Lenovo ThinkSystem SC750 V4 Neptune[®], powered by the Intel[®] Xeon[®] 6 processor — performance and efficiency without compromise.

**Smarter
technology
for all**


Lenovo

1. Diversity and inclusion

Responsible AI must work for everyone, equally, everywhere. Protecting against bias, prejudice, or discrimination needs to take place during the data collection and training process to include diverse and representative samples of the population. It must also happen while in use during inferencing activities to monitor the learnings and output.

Questions to ask:

- What processes are in place to test and monitor for potential biases during the entire lifecycle of the AI system (e.g., biases due to possible limitations in the composition of the data sets used)?
- Have your data scientists been trained to avoid introducing unintentional bias?
- Did you consider the diversity and representativeness of end users and/or subjects in the data?



Women make up **only 24%** of computer science PhD graduates in North America.³



Lenovo ThinkSystem SC750 V4 Neptune[®], powered by the Intel[®] Xeon[®] 6 processor — performance and efficiency without compromise.

Smarter
technology
for all

Lenovo






2. Privacy and security

AI should protect the privacy of both individuals and groups in its inputs and outputs. That means excluding data that was gathered in ways that violate privacy. Nor should AI provide any results or outcomes that violate the privacy of the subjects. AI applications must also be secure against cybersecurity threats and AI-specific threats such as data poisoning.

Questions to ask:

- Was your data collected in ways that exclude private or individually identifiable elements?
- Have you considered the impact of the AI system on the right to privacy for data subjects and users?
- Have you thoroughly audited the AI system for cybersecurity risks?



85% of recent cybersecurity attacks were powered by generative AI.⁴



Lenovo ThinkSystem SC750 V4 Neptune[®], powered by the Intel[®] Xeon[®] 6 processor — performance and efficiency without compromise.

Smarter
technology
for all

Lenovo



3. Environmental and social impact

All AI projects should be evaluated for their potential impact on both the environment and society at large. From an environmental perspective, the energy requirements for training and running a generative AI model can have a significant and measurable carbon footprint. Running trained AI inference workloads requires less energy.

Potential social impacts can include disruption of democratic decision-making processes, targeting marginalized populations, or encouraging addictive behaviors through media manipulation.

Questions to ask:

- What tools and technologies are you deploying to minimize the environmental impact of the AI system's development, deployment, and ongoing use?
- Have you assessed the societal impact of the AI system's use beyond the end user?
- Have you considered what negative effects could come from the AI system?



Training a single AI model can produce **626,000 pounds of CO₂ equivalent.**⁵



Lenovo ThinkSystem SC750 V4 Neptune[®], powered by the Intel[®] Xeon[®] 6 processor — performance and efficiency without compromise.

Smarter
technology
for all

Lenovo

4. Accountability and reliability

Ultimately, the responsibility for decisions made by an AI application must be held to humans. Failing to identify the people responsible ahead of time can result in finger-pointing when negative outcomes occur, which hinders solving the problem and moving forward. In addition, AI applications need to be reliable. The results should be predictable and repeatable without introducing erroneous results. Users should be able to recognize and react to failures in the AI outcomes.

Questions to ask:

- What human oversight exists throughout the building, training, and testing of the AI model? Have you assigned responsibility in case of AI failure?
- How traceable is the development process, the sourcing of training data, and the logging of the AI system's process and outcomes (both positive and negative)?
- Have you stress-tested the AI system for reproducibility, minimum points of failure, adversarial attacks, etc.?



57%

of organizations are concerned about managing data security for GenAI.⁶



Lenovo ThinkSystem SC750 V4 Neptune[®], powered by the Intel[®] Xeon[®] 6 processor — performance and efficiency without compromise.

Smarter
technology
for all

Lenovo

5. Explainability

To the casual observer, AI systems can appear to be “black boxes” where the input and output are known but the decision-making process is a mystery. But for critical decisions that affect things like business outcomes, healthcare decisions, or financial transactions, the reasons for the decisions need to be understandable and explainable to all parties involved.

Questions to ask:

- Can the AI’s decisions be explained to the users in easy-to-understand terms?
- Do these explanations rely on additional algorithms for enhancement? If so, are those also explainable?
- Are the descriptions of the data sets accurate, complete, and standardized?



Businesses that attribute **at least 20% of EBIT** to their use of AI are more likely than others to follow best practices that enable explainability.⁷



Lenovo ThinkSystem SC750 V4 Neptune®, powered by the Intel® Xeon® 6 processor — performance and efficiency without compromise.

Smarter
technology
for all

Lenovo




6. Transparency

Users should know who or what is making decisions behind the screen — whether they're interacting with an AI agent or a live person. In addition, it's important that both developers and users know what data was used to make a decision and what version of the AI model made the decision.

Questions to ask:

- In situations where it occurs, do you clearly communicate to users that they're interacting with an AI application instead of a human?
- Do you have measures in place to continuously monitor the quality of the input data?
- If a newer version of the AI model is released and doesn't perform as well, do you have traceability to go back and correct the model?



A recent study of 10 key foundational AI models showed that **only two of the models** scored more than 50% on 100 transparency indicators.⁸



Lenovo ThinkSystem SC750 V4 Neptune®, powered by the Intel® Xeon® 6 processor — performance and efficiency without compromise.

Smarter
technology
for all

Lenovo

Building a responsible AI future together

Designing and implementing responsible AI models that deliver reliable, actionable insights takes a very specific set of skills and extreme attention to detail. Lenovo and Intel are ready to help you achieve success with your responsible AI goals.

Working under the principles from the Lenovo Responsible AI Committee, the Lenovo AI Discover Center of Excellence brings together Lenovo and Intel AI experts to help your developers create and accelerate the delivery of responsible AI applications and AI inferencing models.

Technical engineers, partners, and data scientists optimize your AI codes using Intel technology and open-source frameworks on ThinkSystem servers like the Lenovo ThinkSystem SC750 V4 Neptune[®], powered by Intel[®] Xeon[®] 6 processors. Intel[®] Xeon[®] offers power-efficient performance and sustainability features to help lower total cost of ownership.

6th Generation Lenovo Neptune Liquid Cooling enables the most demanding AI and HPC workloads while delivering 100% heat removal and an up to 40% reduction in power consumption.^{9,10}

Get started on the road to responsible AI today. Talk to your Lenovo representative and visit the [Intel AI Alliance page](#) to learn more.

Sources


- 1 Accenture, "Reinventing Enterprise Operations: How reinvention-ready companies are driving growth and relevance with gen AI," October 2024
- 2 ZDNet, "Everyone wants responsible AI, but few people are doing anything about it," August 2023
- 3 Computing Research Association, "2023 Taulbee Survey," May 2024
- 4 CFO, "85% of Cybersecurity Leaders Say Recent Attacks Powered by AI: Weekly Stat," August 2023
- 5 University of Massachusetts, "Energy and Policy Considerations for Deep Learning in NLP," June 2019
- 6 Deloitte, "Now decides next: Moving from potential to performance," August 2024
- 7 McKinsey, "Why businesses need explainable AI — and how to deliver it," September 2022
- 8 Axios, "AI developers are failing on transparency, new index shows," October 2023
- 9 Lenovo data based on internal Lenovo ISG research
- 10 Based on Lenovo internal testing against similar air-cooled systems in a typical data center



Lenovo ThinkSystem SC750 V4 Neptune[®], powered by the Intel[®] Xeon[®] 6 processor — performance and efficiency without compromise.

Intel, the Intel logo, OpenVINO, and the OpenVINO logo are trademarks of Intel Corporation or its subsidiaries.

© Lenovo 2025. All rights reserved. v1.00 January 2025.



**Smarter
technology
for all**

Lenovo