



Drive Revenue Growth and Improve Customer Experience with Faster, More Effective AI

Leadership performance with the world's best CPU for AI

4th Generation Intel® Xeon® Scalable Processors with Intel® Advanced Matrix Extensions (Intel® AMX) Outperform AMD EPYC¹



Do More for Your Customers and Business with Better Generative AI and Large Language Models

4th Gen Intel Xeon delivers up to

↑ 5.6X
Higher BERT Large performance

AND

4.7X
Better performance/watt for running a BERT Large workload than the 4th Gen AMD EPYC²

BERT Large is a widely used machine learning model for natural language processing (NLP) and language based AI applications.

Reduce Total Cost of Ownership (TCO) Across Your Server Fleet for AI

↓ As much as 79%
Lower TCO than the 4th Gen AMD EPYC while running a BERT Large workload³




Solve common problems

- Better inform business decisions to drive revenue growth
- Reduce repetitive tasks, costs, and time for your business
- Create new content quickly and efficiently
- Improve customer retention and acquisition with better sentiment analysis
- Analyze large amounts of data faster
- Better understand, examine, and respond to customer queries
- More personalized customer suggestions
- Enable more responsive smart assistants and chatbots
- Improve text prediction speed and accuracy

Lower the Carbon Footprint of Your Data Center with Leadership Performance in AI⁴

Fewer servers to manage



50

servers with 4th Gen AMD EPYC processors

vs.

9

servers with 4th Gen Intel Xeon Scalable processors

Lower costs

Saves



\$1.38M
USD

Power demand reduced

Saves



424,000
kWH/year

More sustainable data center operations

Reduces



719,546
kg CO₂ emissions over 4 years

¹ 4th Gen Intel® Xeon® processor leadership over AMD EPYC Genoa as of June 28, 2023. Above data reflects Intel Xeon 8462Y+ vs. AMD EPYC 9354.

² <https://www.intel.com/content/www/us/en/content-details/781683/4th-gen-xeon-outperforms-competition-on-real-world-workloads.html>, slide 4, accessed July 6, 2023.

³ <https://www.intel.com/content/www/us/en/content-details/781683/4th-gen-xeon-outperforms-competition-on-real-world-workloads.html>, slide 5, accessed July 6, 2023.

⁴ Results are based on BERT Large. <https://www.intel.com/content/www/us/en/content-details/781683/4th-gen-xeon-outperforms-competition-on-real-world-workloads.html>, slide 5, 17, accessed July 6, 2023.

Performance varies by use, configuration, and other factors. Learn more at www.Intel.com/PerformanceIndex. Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure. Your costs and results may vary. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

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

© 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. 0723/KO/HBD/PDF

Intel Xeon is architected for AI with Intel AMX. This leading AI acceleration is built into each Intel Xeon processor core to help you do more for your business. Find out more about Intel Xeon Scalable processors with built-in accelerators at intel.com/acceleratorengines.

intel
xeon®