PT

Transform ideas into reality with the HP Z8 Fury G5 Workstation PC

A comparison to a previous-generation workstation and a current-generation competitor

Whether you're generating three-dimensional images, shaping machine-driven interactions, or working with AI-based language technologies, the ultimate goal is the same: Get more done in less time. To test which system could help you complete your biggest projects—and save time doing so—we performed a variety of intensive tests on the HP Z8 Fury G5 Workstation PC, the HP Z8 G4 Workstation, and the Lenovo® ThinkStation® P620 Tower Workstation.

| Get back time for important projects | Less waiting for renders, more time to create | Make faster diagnoses and predictions |

| WHAT WE TESTED | CONTENT-CREATION BENCHMARKS | AI/ML WORKLOADS | RESOURCES |

We tested the following systems:

The current-gen Z8 workstation:

<u>HP Z8 Fury G5 Workstation Desktop PC</u>
1x 56-core Intel® Xeon® w9-3495X
CPU (1.9 - 4.8 GHz)
4x NVIDIA® RTX® 6000 Ada Generation GPUs

The previous-gen Z8 workstation:

<u>HP Z8 G4 Workstation Desktop PC</u>
2x 28-core Intel Xeon Gold 6258R
CPUs (2.7 - 4.0 GHz)
2x NVIDIA RTX A6000 GPUs

A current-gen competitor:

<u>Lenovo ThinkStation P620 Tower Workstation</u>
1x 64-core AMD Ryzen Threadripper PRO 5995WX
CPU (2.7 - 4.5 GHz)
2x NVIDIA RTX A6000 GPUs with NVLink

For more on our configurations and tests, read the reports:

Content-creation benchmarks

- CPU single-core performance with Cinebench R23 and Geekbench 6 Pro
- GPU performance with augmented reality and machine learning workloads with Geekbench 6 Pro
- Production rendering performance with Maxon Redshift
- Photo-editing performance with an Adobe® Creative Cloud® app with PugetBench for Lightroom Classic
- GPU rendering and multi-GPU acceleration and ray-tracing performance with Blender

What could these wins mean for you?

AI/ML testing

- Medical imaging and 3D image segmentation performance with 3D U-Net
- Speech recognition performance with RNN-T
- Natural language processing performance with BERT-99
- Image classification and detection performance with ResNet-50

What could these wins mean for you?

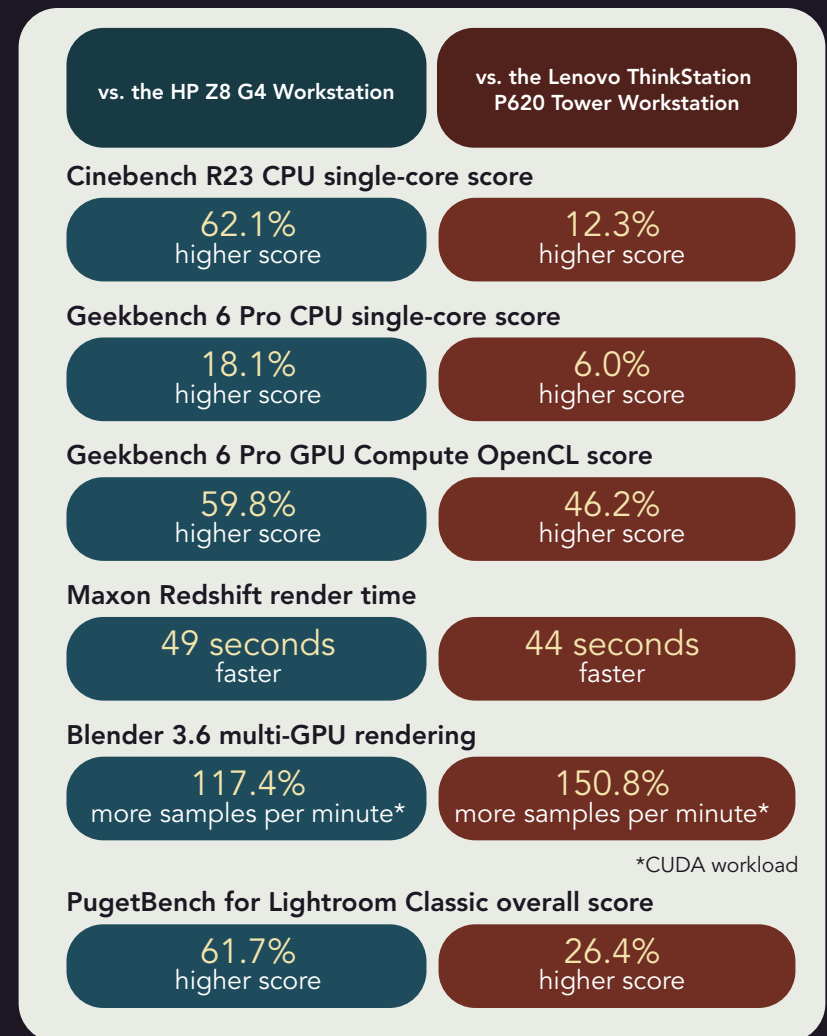| WHAT WE TESTED | CONTENT-CREATION BENCHMARKS | AI/ML WORKLOADS | RESOURCES |

Accelerate performance and expedite project cycles

What do these workloads do?

- **Cinebench R23** rates system hardware performance with Cinema 4D, a real-world 3D computer animation, modeling, simulation, and rendering software.[1]

- **Geekbench 6 Pro** quantifies how well a system performs CPU- and GPU- intensive tasks, such as image editing, image synthesis, and physics simulations.[2]

- **Maxon Redshift** and **Blender 3.6** measure 3D rendering performance, stressing a system's GPU.

- **PugetBench for Lightroom Classic** tests how well a system performs with a series of tasks in the industry-standard photo-editing application.

vs. the HP Z8 G4 Workstation	vs. the Lenovo ThinkStation P620 Tower Workstation
Cinebench R23 CPU single-core score	
62.1% higher score	**12.3%** higher score
Geekbench 6 Pro CPU single-core score	
18.1% higher score	**6.0%** higher score
Geekbench 6 Pro GPU Compute OpenCL score	
59.8% higher score	**46.2%** higher score
Maxon Redshift render time	
49 seconds faster	**44 seconds** faster
Blender 3.6 multi-GPU rendering	
117.4% more samples per minute*	**150.8%** more samples per minute*

*CUDA workload

PugetBench for Lightroom Classic overall score	
61.7% higher score	**26.4%** higher score

Accelerate performance and expedite project cycles

What could these wins mean for you?

Saving time on these workflows means creatives could spend less time waiting on technology and more time bringing their visions to life. For example, if a video team had an ambitious vision for a 3D-animated ad, an underperforming workstation might create less impressive renders or or cause them to miss the deadline. A higher performing system could help the video team quickly move their vision from the storyboard to their client's inbox without compromising quality or time.

And these test results aren't only significant for creators and editors. As technology advances, the line between creative and scientific work often blurs. The Geekbench 6 Pro tests we ran, for example, include a physics particle simulation that artists and developers might use to animate water or smoke.[3]

However, particle physics also has important applications in biomedicine, diagnostic instruments, power transmission, and more.[4] When scientific teams run particle simulations, faster workstation speeds mean less waiting on key data. The scientific process can be long and demanding, but with answers in hand sooner, physicists can move forward with important innovations—or recognize that they need to go back to the drawing board.
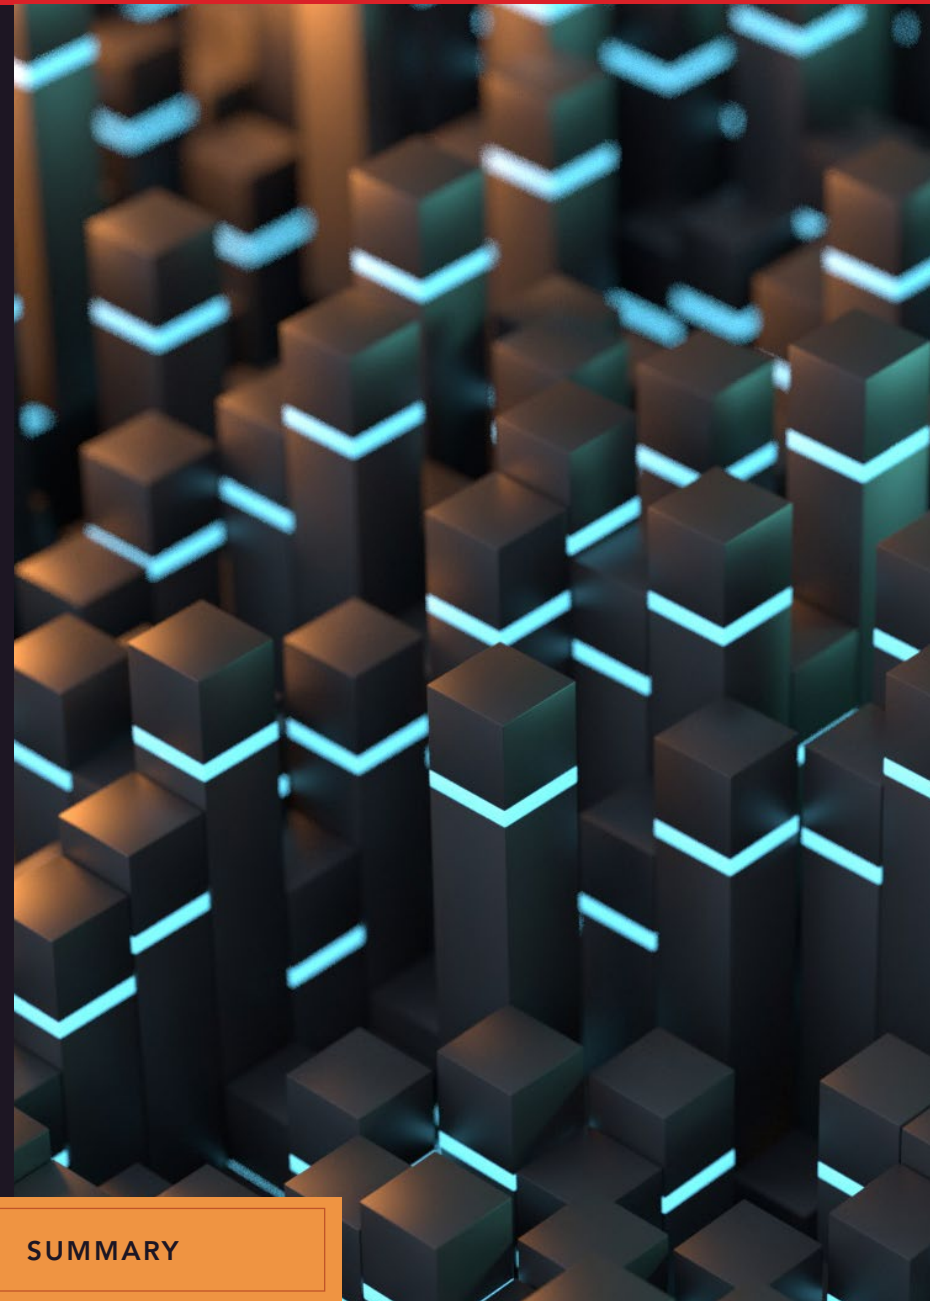
Accelerate performance and expedite project cycles

Summary

Powerful workstations are vital for companies as they prepare for the future. With better performance under intensive creative and graphics workloads, your teams could achieve more with the HP Z8 Fury G5. With less waiting on complicated processes such as renders and simulations, both content creators and technical workers can speed project completion and deliver results they can be proud of, even under tough deadlines.

To see the results of other tests, read the reports:

TESTS / RESULTS

WHY IT MATTERS

SUMMARY

WHAT WE TESTED

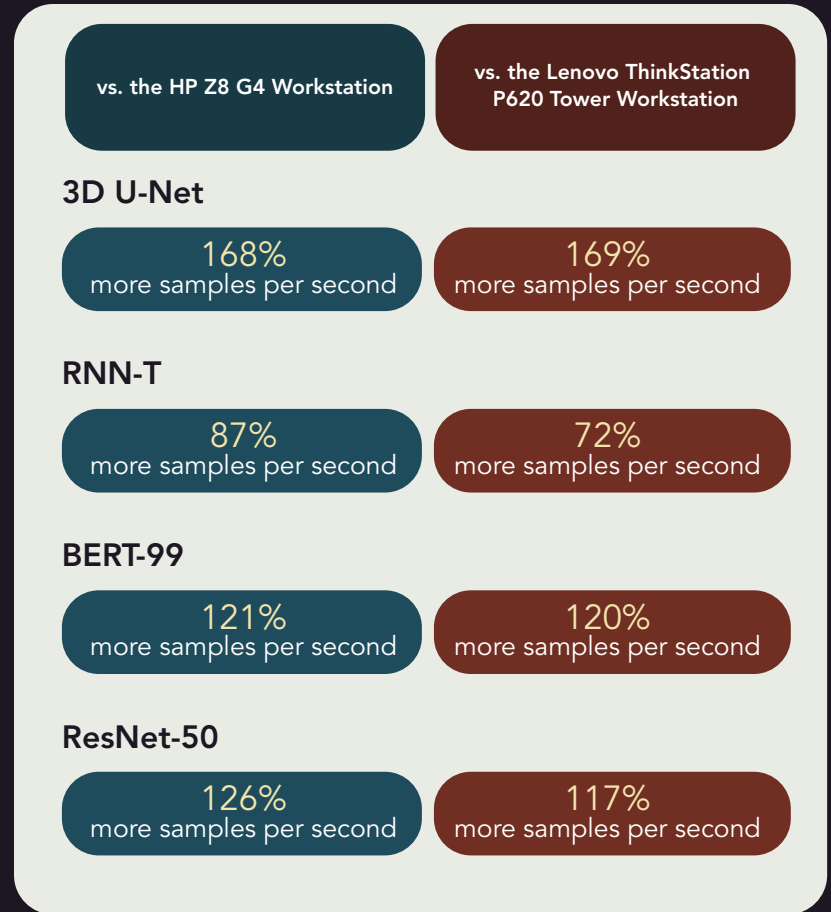CONTENT-CREATION BENCHMARKS

AI/ML WORKLOADS

RESOURCES

Tackle complex AI/ML problems

What do these workloads do?

The medical imaging, language processing, and computer vision scenarios we ran use trained models to measure how quickly each workstation processed inputs and produced results.[5]

- Medical 3D imaging: The healthcare sector uses medical imaging (e.g., X-rays, ultrasounds, MRIs, and CT scans) for medical research, disease diagnosis, and drug discovery.[6] The **3D U-Net** model we ran "performs volumetric segmentation of dense 3D images for medical use cases."[7]

- Natural language processing (NLP):

  - The recurrent neural network (RNN) aspect of NLP helps systems recognize email categories, predict stock prices, mine text, and translate different languages.[8,9] The **RNN-T** model we ran "recognizes and transcribes audio in real time."[10]

  - The **BERT** model we ran sorts and analyzes text to make accurate language predictions and respond to conversations.[11] BERT-based applications include virtual assistants (e.g., Alexa and Siri), chatbots, and video captioning.[12]

- Computer vision: Machine learning can enable computers to "see," identify, and understand objects and people in images and video. Computer vision applications include facial recognition, autonomous cars, production-line automation, and sports performance analysis.[13] The 50-layer **ResNet** model we ran "[a]ssigns a label from a fixed set of categories to an input image, i.e., applies to computer vision problems."[14]

vs. the HP Z8 G4 Workstation	vs. the Lenovo ThinkStation P620 Tower Workstation
3D U-Net	
168% more samples per second	169% more samples per second
RNN-T	
87% more samples per second	72% more samples per second
BERT-99	
121% more samples per second	120% more samples per second
ResNet-50	
126% more samples per second	117% more samples per second

TESTS / RESULTS   **WHY IT MATTERS**   **SUMMARY**

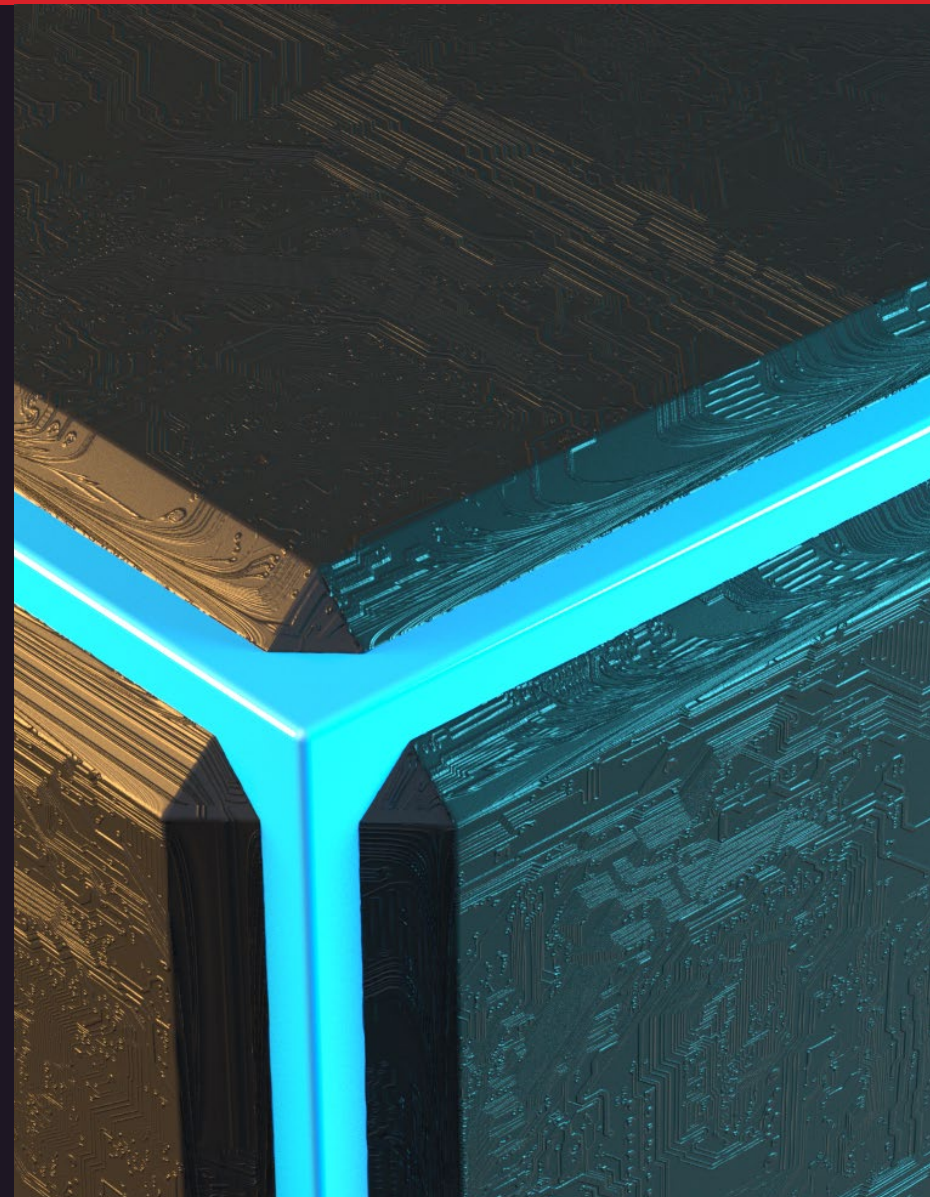**WHAT WE TESTED**   **CONTENT-CREATION BENCHMARKS**   **AI/ML WORKLOADS**   **RESOURCES**

Tackle complex AI/ML problems

What could these wins mean for you?

The final stage in the machine learning process is inference—that's the golden time when your proven ML model has all the data, training, evaluation, and tuning your experts deem necessary for that model to make informed and (hopefully) accurate predictions.[15]

Business intelligence analysts, digital transformation specialists, data analysts, and scientists need powerful workstations that can nimbly crunch through stacks of proprietary and sensitive data to get useable results as quickly as possible. For example, in the medical field, saving time can save lives when it comes to making diagnoses or planning surgeries.

With a system that can process images faster, computer vision can more quickly and effectively assist medical teams in analyzing scans, detecting anomalies, and more. So, whether your organization is analyzing images or text, the sooner you have answers, the sooner you can put that data to work.

TESTS / RESULTS

WHY IT MATTERS

SUMMARY

WHAT WE TESTED

CONTENT-CREATION BENCHMARKS

AI/ML WORKLOADS

RESOURCES
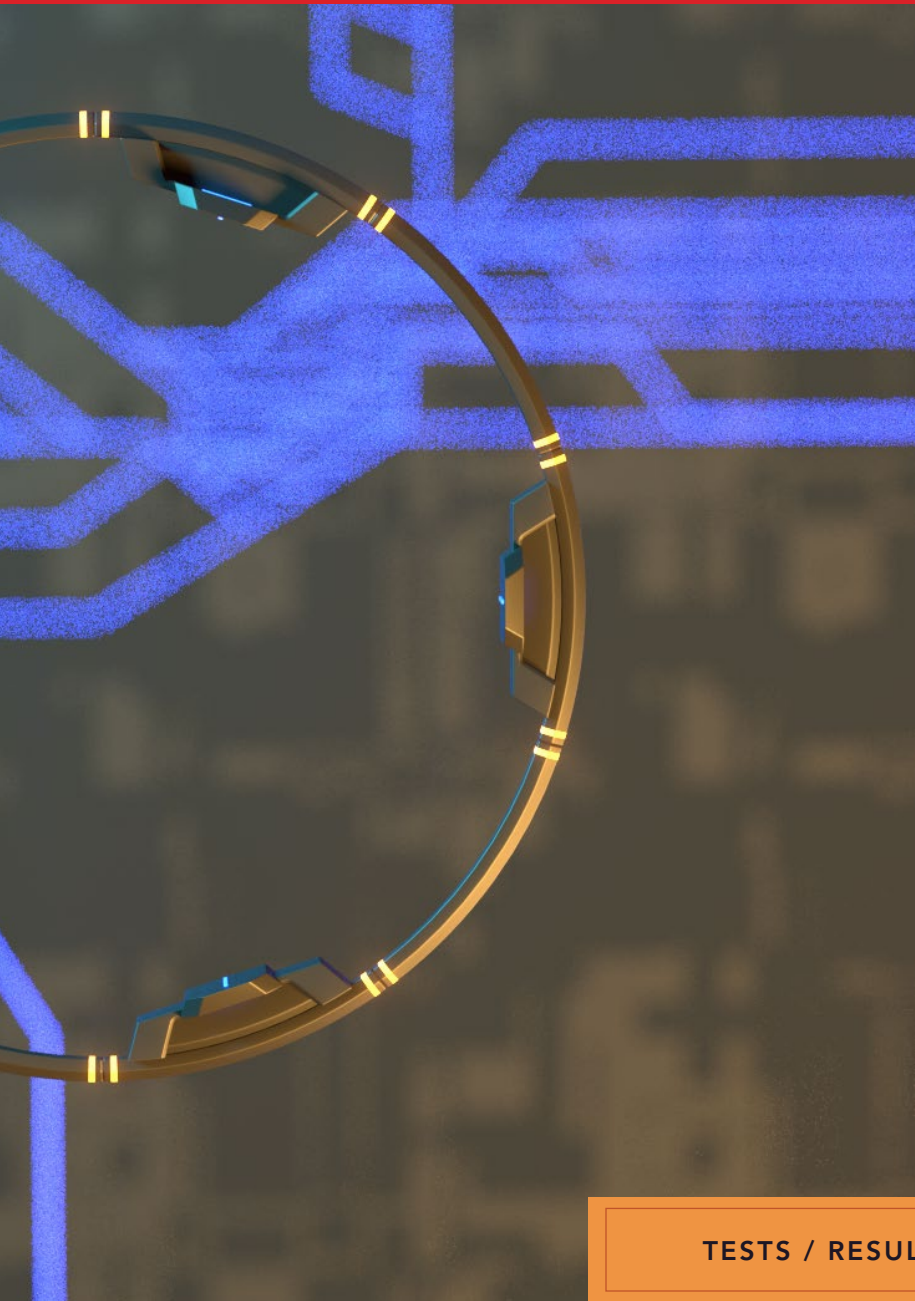
Tackle complex AI/ML problems

Summary

Saving time at any point in the machine learning process can prove invaluable. Our medical 3D imaging, NLP, and computer vision tests show that the HP Z8 Fury G5 classified more samples in less time than the other workstations we tested. With these time savings, scientists, analysts, and engineers could act on valuable insights sooner, integrating them into technologies, diagnoses, or business strategies.

For more details, including configurations, read the reports:

TESTS / RESULTS	WHY IT MATTERS	SUMMARY

WHAT WE TESTED	CONTENT-CREATION BENCHMARKS	AI/ML WORKLOADS	RESOURCES

Resources

For more details, read the reports:

See the video:

1. Maxon, "Cinema 4D," accessed August 22, 203, https://www.maxon.net/en/cinema-4d.

2. Geekbench, "Geekbench 6 GPU Compute Workloads," accessed August 23, 2023, https://www.geekbench.com/doc/geekbench6-gpu-compute-workloads.pdf.

3. Geekbench, "Geekbench 6 GPU Compute Workloads."

4. Office of High Energy Physics, "Particle Physics: Benefits to Society," accessed October 18, 2023, https://science.osti.gov/hep/-/media/hep/pdf/files/pdfs/hep_benefits_v2.pdf.

5. NVIDIA, "What is MLPerf?" accessed September 22, 2023, https://www.nvidia.com/en-us/data-center/resources/mlperf-benchmarks/.

6. Simplilearn, "Top 25 Deep Learning Application Industries," accessed September 22, 2023, https://www.simplilearn.com/tutorials/deep-learning-tutorial/deep-learning-applications.

7. NVIDIA, "What is MLPerf?" accessed September 22, 2023, https://www.nvidia.com/en-us/data-center/resources/mlperf-benchmarks/.

8. Tableau, "8 Natural Language Processing (NLP) Examples," accessed September 22, 2023, https://www.tableau.com/learn/articles/natural-language-processing-examples#:.

9. Great Learning, "What Is Recurrent Neural Network | Introduction of Recurrent Neural Network," accessed September 22, 2023, https://www.mygreatlearning.com/blog/recurrent-neural-network/.

10. NVIDIA, "What is MLPerf?" accessed September 22, 2023, https://www.nvidia.com/en-us/data-center/resources/mlperf-benchmarks/.

11. NVIDIA, "What is MLPerf?" accessed September 22, 2023, https://www.nvidia.com/en-us/data-center/resources/mlperf-benchmarks/.

12. Simplilearn, "Top 25 Deep Learning Application Industries," accessed September 22, 2023, https://www.simplilearn.com/tutorials/deep-learning-tutorial/deep-learning-applications.

13. Built In, "What Is Computer Vision?" accessed September 22, 2023, https://builtin.com/machine-learning/computer-vision.

14. NVIDIA, "What is MLPerf?" accessed September 22, 2023, https://www.nvidia.com/en-us/data-center/resources/mlperf-benchmarks/.

15. Matthew Mayo, "Frameworks for Approaching the Machine Learning Process," accessed September 20, 2023, https://www.kdnuggets.com/2018/05/general-approaches-machine-learning-process.html.

| WHAT WE TESTED | CONTENT-CREATION BENCHMARKS | AI/ML WORKLOADS | RESOURCES |

Principled Technologies®

Facts matter.®

This project was commissioned by HP.