



Experiment, Prototype and Build the Future of AI on Workstation Solutions

Solutions Guide



Meeting the Demands for AI Computing Resources

Generative AI is bringing profound change across industries, accelerating the adoption of AI-infused technologies at an incredible scale. These new AI-powered workflows offer the promise of new levels of creativity and productivity, improving efficiency across industries.

But they also require significantly more computing power than before. The models used for generative AI are very large, taking weeks or months to train and test on clusters of servers. The results are highly complex models, capable of understanding language, voice, and audio—or creating content such as articles, images, and music, and much more.

95% of executives believe generative AI will compel their organization to modernize its technology architecture¹. The rush to rapidly add AI computing power to data centers and increase the availability of accelerated cloud instances is straining the availability of hardware. This makes it difficult to meet the continually growing demand.

Key Challenges

- **Hardware:** Current desktop systems often lack the processing power needed for AI workflows, leading to greater reliance on already strained cloud and data center resources.
- **Training and fine-tuning:** AI model size continues to grow, taking months to train and taxing already oversubscribed data center and cloud instance resources. Off-the-shelf models trained on public data require fine-tuning with domain-specific data to provide content relevant to business purposes.
- **Workflow complexity:** Modern professional workflows require running multiple applications simultaneously to maximize productivity. Adding AI-augmented tools and applications puts additional requirements on current computing solutions.

*Please note that several of the products (HP AI Studio GA 2025 waitlist in US/UK, HP Boost GA 2025 in US/UK), and as such, they may not be generally available to all customers at this time. Availability is subject to change, and specific eligibility criteria apply.

We encourage you to reach out to your representative for further details on the products mentioned, including information on how to apply for access to those in the early access phase.



HP Z AI Workstations Benefits

- Provide additional AI computing resources to augment data center and cloud instances for development and R&D tasks.
- HP AI Studio, HP Z Boost, and a pre-loaded Data Science Software Stack streamline and simplify workflows.
- Large GPU memory configurations enable AI-augmented, multi-application workflows that maximize productivity.
- Enterprise-grade solutions designed to work 24/7/365 and are widely available worldwide.

Recommended HP Z Desktop and Mobile Workstations for Using and Developing AI

	HP ZBook 8	HP ZBook X	HP ZBook Ultra	HP ZBook Fury	HP Z1 Tower	HP Z2 Mini/SFF/Tower	HP Z4/Z4 Rack	HP Z6/Z6 A	HP Z8/Z8 Fury
Using AI	• Gen AI/content creation, editing and optimization • Project management and monitoring			• 3D modeling and animation • Data visualization		• Real-time visualizations and simulations • Virtual production, visual effects and compositing/finishing • Reality and motion capture			
Model Fine Tuning & Retrieval Augmented Generation (RAG)	• Exploratory data analysis • Data cleaning and pre-processing • Model development			• End-to-end data processing • Machine learning development • Concept validation		• Advance visualization • Deep learning model training • Computer vision			
Ultimate mobility and battery life					Ultimate performance				

HP Z AI Workstations Accelerated by NVIDIA RTX PRO™

Our comprehensive AI portfolio is two-fold:

- **Differentiated Hardware:** HP Z AI workstations accelerated by NVIDIA RTX PRO™¹ offer incredible desktop computing performance—perfectly suited for AI training, inference, and data science workflows.
- **Solutions Innovations:** Tools like HP AI Studio and a pre-loaded Data Science Software Stack streamline workflows and enhance efficiency and productivity.

Together, these robust workstations and solutions are optimized for processing smaller AI models locally with exceptional efficiency, while also providing seamless integration with data center and cloud resources for the development of larger, more complex models. Our most powerful HP Z desktop AI workstation with up to four RTX PRO 6000 Blackwell Max-Q Workstation Edition Series GPUs per workstation for an incredible 14.2 petaFLOPS of combined compute performance and 384 gigabytes (GB) of total system GPU memory delivers up to 14,044 AI TOPS for advanced AI.

As generative AI tools become integral to professional workflows, the need for powerful GPUs grows. Running large language models (LLMs) like chatbots and code copilots locally on workstations further amplifies these demands. HP Z AI workstations powered by NVIDIA RTX PRO² are designed to handle these intensive workloads.

Key Workloads

HP Z AI Workstations for Data Science


HP Z AI workstations, equipped with up to four NVIDIA RTX PRO GPUs and substantial GPU memory, are ideal for data science tasks. These powerful workstations reduce latency, facilitating real-time data preprocessing, exploration, visualization, and evaluation of features and models. This saves time and valuable

data center and dedicated cloud compute resources. Additionally, our Data Science Stack Manager³ preloads and automatically updates your workstation with popular open-source software and development tools, allowing your teams to get up and running quickly. It integrates seamlessly with the HP Z Data Science Software Stack, available on Ubuntu or Windows, and includes open-source software like TensorFlow and PyTorch, as well as developer tools like Git and Visual Studio Code. Furthermore, it supports NVIDIA CUDA-X libraries and RAPIDS data processing and machine learning libraries—these workstations deliver accelerated workflows for data preparation, model training, and visualization. On the HP Z AI Workstation, NVIDIA's cuDF accelerates pandas operations by up to 56×, while NVIDIA's cuML achieves up to 65× faster Scikit-Learn workloads, 76× faster HDBSCAN clustering, and 220× faster UMAP for dimensionality reduction—compared to CPU-only systems. This enables rapid experimentation and iteration at enterprise scale without modifying existing workflows.

HP Z AI Workstations for Using and Developing AI

For businesses or individuals starting with AI or using smaller models, HP Z AI workstations accelerated by NVIDIA RTX PRO² offer a robust and cost-effective solution for AI R&D. A multi-GPU system can augment data center or cloud resources by efficiently handling the developing needs of small teams or workgroups, allowing for local AI model deployment and faster iteration.

In an environment where many competitors are cloud-centric, focusing on on-prem solutions addresses the needs of organizations that prefer or require data processing and AI model development to occur within their own infrastructure for reasons like enhanced security, regulatory compliance, or data sensitivity. Developer teams can fine-tune models as needed and experiment with various datasets and data sizes to optimize results, all while conserving data center or cloud computing resources.





HP Z AI Workstations for Inference

Applications with AI-enabled features—Adobe® Photoshop’s® Neural Filters, DaVinci Resolve’s face tracking, NVIDIA Broadcast’s noise and room echo removal, and image denoising in major rendering software—have been available for several years. Generative AI is bringing new levels of capabilities and efficiency to professionals, demanding more computing power and GPU memory. Professionals often work with high-resolution content and use multiple applications simultaneously, which increases these requirements.

As generative AI tools become integral to professional workflows, the need for powerful GPUs grows. Running large language models (LLMs) like chatbots and code copilots locally on workstations further amplifies these demands. HP Z AI workstations accelerated by NVIDIA RTX PRO² are designed to handle these intensive workloads. The NVIDIA RTX PRO 6000 Blackwell Workstation Edition GPU², with 96GB of GPU memory, provides the raw AI computing power and memory necessary to manage high-resolution generative AI content, iterate designs, and integrate with other applications seamlessly, without compromising performance or content fidelity.

Recommended HP Z Desktop and Mobile Workstations for Using AI-enabled Applications



HP Z AI Studio

US and UK only

HP Z AI workstations with HP AI Studio⁴ and NVIDIA® NGC™ provide a full-stack solution for AI development. They offer GPU-accelerated models and SDKs to build and deploy solutions across various AI workflows, addressing challenges around trust, security, and scalability.

- Streamline model search and access
- Test and develop directly within HP AI Studio
- Collaborate remotely with HP AI Studio

HP AI Studio's powerful analytics tools provide an overview of available compute resources, helping teams plan and perform operations efficiently. HP AI Studio ensures consistency in AI development workflows by allowing teams to easily identify the right options, avoiding confusion from different versions of data sets or developer utilities. Its microservices offer optimized model performance with enterprise-grade security, support, and stability.


HP Z Boost

US and UK only

Optimize access to existing and future GPU compute and deliver a performance boost to your Data Scientists and AI developers when they need it. HP Z Boost facilitates client device access to up to 4 GPUs on a host workstation (must be a HP Z workstation) via standard network protocol and VPN/encryption. One or more users can access the shared GPUs in a single host through GPU pools. Boost offers many benefits.

- **On-demand performance boost** – choose shared GPUs using a simple drop-down menu when users need them.
- **Accelerate AI innovation** – access GPU power when you need it for faster model development.
- **Efficient cost management** – reduce variable costs associated with cloud compute⁶ and make the most of idle GPU compute you own now and in the future.
- **Anytime productivity** – tackle complex AI workflows from nearly anywhere⁷ with GPU access through standard network protocols.
- **Securely train AI models** - locally or via VPN/encryption in an individual user session.

On-Demand GPU Compute, Turning Your Workstation into a Sharable Resource



Enterprise-Ready Solutions

HP Z AI workstations accelerated by NVIDIA RTX PRO² are designed and built for demanding enterprise deployments. Powered by the latest generation of workstation CPUs, HP Z AI workstations are ready to tackle demanding AI development workflows. The latest generation of HP Z desktop and mobile AI workstations are available now and ready to ship.

With a full stack of enterprise-level deployment, support, and optimization tools, HP Z AI workstations easily fit into existing IT infrastructure, providing drop-in solutions for AI training, developing, deploying, and inferencing on the desktop through HP AI Studio⁴. The NVIDIA[®] GPU² architecture scales with the use of HP Z Boost⁵ from cloud to data center, desktop, mobile, and embedded device, supporting the same software stack across devices, which enables users to move workloads seamlessly between them.

Ready to get started?

To learn more about the HP Z AI workstations, visit:
hp.com/z

To learn more about NVIDIA RTX PRO, visit:
nvidia.com/ai-workstations

To contact sales, visit:
hp.com/us-en/workstations/z-support.html



1. Source: <https://www.accenture.com/us-en/blogs/blogs-careers/the-future-making-it-human>

2. Products sold separately and must be configured at purchase.

3. HP Z Data Science Stack Manager requires Windows 10 version 21H2 (Build 19044) and higher or 64-bit Ubuntu 20.04 and is available on select Z workstations.

4. HP AI Studio requires Windows 10 (build 19041 and higher) or Windows 11 OS, Intel Core™ i5 12th gen or higher, AMD Ryzen™ 9 or higher processors, minimum 50 GB of available storage, 16GB of RAM, and Internet access. To enable GPU compute, HP AI Studio requires any NVIDIA® GPU compatible with driver version 528.89 or newer. A minimum of 8GB of VRAM is recommended. Join waitlist for access to HP Z AI Studio early access program.

5. HP Z Boost requires client device with Windows 10 or Windows 11 OS and Vulkan Runtime for devices without a GPU, or Ubuntu 22.04 LT or greater, Debian 11 or greater with libatomic1 or libnuma1 packages installed. Host device for sharing GPU(s) must be a HP Z workstation with NVIDIA professional grade workstation GPU(s) of Pascal architecture of later with an installed driver version 535 or greater, and a CUDA driver version 11.2 through 12.3, Windows 10 or Windows 11 OS, or Ubuntu 22.04 LT or greater or Debian 11 or greater with libatomic1, libnuma1, libvulkan1, libgl1, libglb2.0-0 packages installed. Network connection required. Join waitlist to be considered for selection for early access.

6. Cost reduction of HP Z Boost relative to variable costs with cloud compute.

7. Network connection required.

© Copyright 2025 HP Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. Microsoft and Windows are U.S. registered trademarks of the Microsoft group of companies. Apple, Mac, macOS and MacBook are registered trademarks of Apple Inc.