



adaptive.hpc/ai/ml-as-a-service

an All-inclusive AI Training Solution for Enterprises of Any Size

Pulling together an AI/ML initiative is no small feat. Businesses can face several challenges when implementing AI/ML programs. Many enterprises do not have high-performance computing capabilities on-premises. Cost-prohibitive commercial tools and infrastructure, long-term commitments to Cloud Service Providers, setting up the necessary infrastructure for data storage, processing, model training, and choosing the appropriate AI/ML tools and frameworks are just some of the pitfalls.

adaptive.hpc/ai/ml-as-a-service is an end-to-end high-performance hybrid cloud or cloud-native AI/ML solution that is cost-effective for SMBs or enterprises of any size. In addition to our engineering staff, Adaptive Computing partners with scientific researchers at major universities, and award winning HPC/AI/ML experts to set up and guide customers based on the needs of their organizations.

adaptive.hpc/ai/ml-as-a-service is a complete AI/ML development platform, hosted in Adaptive Computing's on-premises data center, hosted in your data center, or hosted in the Cloud. adaptive.hpc/ai/ml-as-a-service is one of several service and software offerings from Adaptive Computing and includes 120+ open-source applications with additional frameworks and tools, the adaptive.hpc/ai/ml-as-a-service management software stack, and cloud-based or on-premises supercomputing infrastructure, all for a very competitive monthly price.

AI Training as-a-Service for Generative AI and Large Language Models

How it works

Select your preferred Cloud Service Provider and use the credentials associated with your account, or launch your workflows on CloudMart: the multi-cloud compute exchange where you can compare prices and infrastructure for specific requirements.

- A100 GPUs, H100 GPUs, AMD, Dell, HPE, NVIDIA servers, fast interconnects, and storage on-premises in our data center or in the Cloud
- Choose from several deployment models according to your organization's needs

[Deployment Models](#)

adaptive.hpc/ai/ml-as-a-service Technology Stack

The adaptive.hpc/ai/ml-as-a-service software stack includes Moab for dynamic scheduling and the On-Demand Data Center (ODDC) for multi-cloud access and automated infrastructure deployment.

- Uniform multi-cloud provider access through the ODDC web interface (GUI) to launch your AI/ML workloads
- Secure access to cloud resources - ODDC requires you to utilize keys that are 4096-bit class security
- Scale compute resources to multiple nodes to address the demands of your AI/ML requirements

[See Training Video](#)

E4S - Extreme-scale Scientific Software Stack

Install AI/ML and HPC tools and libraries correctly in an integrated and interoperable software stack to deploy on the 4 major cloud platforms! **E4S** was assembled under a DOE (U.S. Department of Energy) grant and is included in the Adaptive Computing Technology Stack.

E4S provides a broad collection of HPC and AI/ML software packages targeting GPUs from NVIDIA, AMD, and Intel. E4S includes TensorFlow, PyTorch, Keras, Horovod, and LBANN for AI/ML development.

Launch 120+ packages from the ODDC web interface into any of the 4 major CSPs.

- Specialized hardware (GPUs, network adapters) supported by E4S on AWS, Azure, GCP, and OCI
- Moab/Torque scheduler integration with MPI environments
- Seamless integration between on-premises infrastructure and the commercial cloud infrastructure

X-ScaleAI Stack

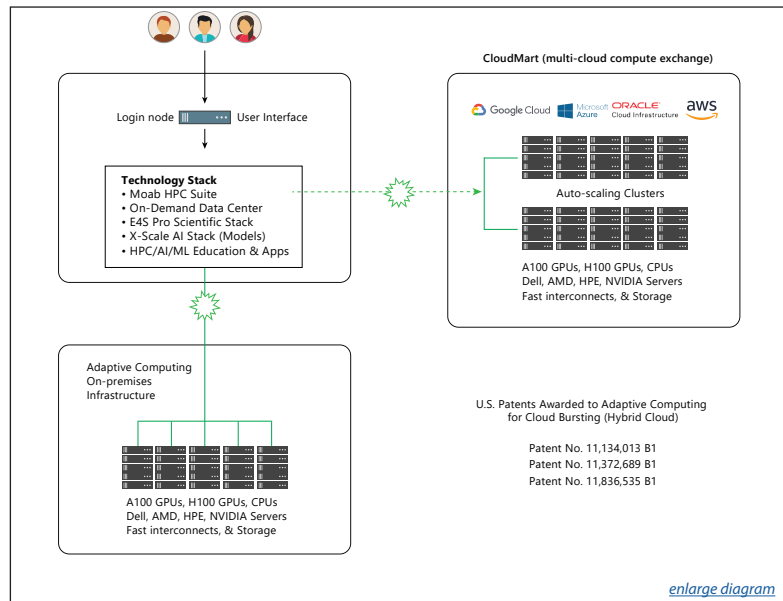
The **X-Scale AI Stack** is a high performance solution for distributed training for your complex AI problems. High-Performance for Distributed Pre-Training, Fine-tuning, and Inference.

- Supports a range of models.
 - Large Language Models (LLMs):
 - Including Llama-3, OLMo, Pythia, and BERT
 - Vision Models
 - ResNet, U-Net, ViT
- Scalable model checkpoint and restart support for long-running training and fine-tuning applications



Benefits

- Simplicity of implementation
- Flexible accelerated infrastructure
 - Compute (A100 and H100 GPUs)
 - Data Storage
 - Fast Networking
 - Free Data Egress on OCI
- Interchangeable AI/ML tool sets
- Easy and clear user experience
- No lock-in to any single Cloud Service Provider
- Lower monthly cost
- AI supercomputer in the Cloud
- Accessible from a web browser
- Generative AI model training
- LLM training and inference
- Enterprise support
- Scale to multiple nodes
- Deliver hybrid cloud from on-premises
- Only pay for what you use



adaptive.hpc/ai/ml-as-a-service hosted by Adaptive Computing, hosted in your data center, or hosted in the Cloud

Applicable to Any Industry, Adaptive Computing Offers Advanced Computing to Companies Requiring Significant Compute Resources for AI Training and Deployment

Several industries heavily rely on significant compute resources for AI and deep learning training and inference. The demand for computational power is particularly high in applications that involve large datasets, complex models, and real-time decision-making.

- **Healthcare:** AI is used in medical imaging, drug discovery, and genomics. Deep learning models are trained on large datasets of medical images and patient records, requiring substantial computational power.
- **Finance:** Financial institutions leverage AI for algorithmic trading, fraud detection, risk management, and customer service. These applications often involve processing vast amounts of financial data in real-time.
- **Automotive:** The automotive industry employs AI for autonomous driving, computer vision, and predictive maintenance. Training deep learning models for image recognition and decision-making in autonomous vehicles demands significant computational power.
- **Retail:** Retailers use AI for demand forecasting, personalized marketing, and supply chain optimization. Deep learning models are trained on large datasets to understand customer behavior and optimize business operations.
- **Manufacturing:** AI is applied in manufacturing for quality control, predictive maintenance, and process optimization. Training deep learning models for image analysis and anomaly detection in manufacturing processes requires considerable computational resources.
- **Energy:** The energy sector uses AI for predictive maintenance of equipment, energy grid optimization, and exploration. Training models to analyze sensor data and optimize energy systems involves substantial computing power.
- **Telecommunications:** Telecommunication companies use AI for network optimization, predictive maintenance, and customer service. Analyzing large datasets to improve network performance and enhance customer experience requires powerful computational infrastructure.
- **Entertainment and Media:** The entertainment industry uses AI for content recommendation, video analysis, and virtual production. Training deep learning models for video processing and content generation can be computationally intensive.
- **Aerospace:** Aerospace engineers use AI for design optimization, simulation, and predictive maintenance. Training deep learning models for aerodynamic analysis and structural optimization involves significant computational resources.
- **Government and Defense:** Government agencies and defense organizations utilize AI for various purposes, including surveillance, threat detection, and decision support. Training models for image recognition and complex decision-making tasks requires substantial compute power.



Headquarters: 1100 5th Ave South, Suite 201, Naples, FL, 34102
 Phone: +1 239-330-6093 | [Contact Us](#) | info@adaptivecomputing.com

© 2024 Adaptive Computing Enterprises, Inc. All rights reserved.

About Adaptive Computing

Adaptive Computing is a global software company headquartered in Naples, Florida, USA and has provided advanced applications and tools to the High-Performance Computing industry for over two decades with hundreds of deployments on the world's largest computing installations. Adaptive Computing products and services are used by organizations of all sizes across a broad range of industries such as High-Tech Manufacturing, Aerospace Engineering, Defense, Universities and Research Labs, Life Sciences, Oil and Gas Exploration, Financial Services, and Data Analytics. Some of the world's largest clusters, grids, and data centers use Adaptive's Moab HPC Suite and Cloud Solutions to maximize performance and value, simplify management, and create a competitive advantage.