

1. What is adaptive.hpc/ai/ml-as-a-service?

adaptive.hpc/ai/ml-as-a-service is a complete AI/ML development offering, hosted by us, hosted in your data center, or hosted in the Cloud. The adaptive.hpc/ai/ml-as-a-service offering is one (1) in a family of as-a-service offerings from Adaptive Computing.

2. What differentiates adaptive.hpc/ai/ml-as-a-service from other AI/ML development offerings?

- Lower price: 10-20 times less per month than other AI/ML development offerings
- Adaptive Management Software: Moab HPC Suite and On-Demand Data Center solutions
- E4S Extreme-scale Scientific Stack: over 120 open-source software applications and optimization tools assembled under a U.S. Department of Energy grant

3. What applications are included with adaptive.hpc/ai/ml-as-a-service?

adaptive.hpc/ai/ml-as-a-service includes over 120 open-source software applications and optimization tools assembled under a U.S. Department of Energy grant including TensorFlow, PyTorch, Keras, Horovod, and LBANN for AI/ML development. The official name for the application stack is [E4S](#) – Extreme-scale Scientific Software Stack. ([applications list](#))

4. What industries are supported by adaptive.hpc/ai/ml-as-a-service?

adaptive.hpc/ai/ml-as-a-service supports all industries and the Performance Intensive Computing as a Service (PICaaS) market as defined by IDC ([link](#)), a leading industry analyst group, and includes support for: High-Performance Computing (HPCaaS), Artificial Intelligence (AI), Machine Learning (ML), GPUs, IoT, Big Data Analytics, Bare Metal as a Service (BMaaS), Application Portability for Disaster Recovery, Cloud Migration, and Accelerated Enterprise Computing.

5. What is the cost for adaptive.hpc/ai/ml-as-a-service?

adaptive.hpc/ai/ml-as-a-service is available for less than the cost of two Senior Software Engineers per month for unlimited users. There is a one-time set up fee based on the selected deployment model. Deployment hosted by Adaptive Computing is \$10,000, hosted in your data center is \$50,000, and if hosted in the cloud is \$25,000*. There are no additional costs for adaptive.hpc/ai/ml-as-a-service.

Number of Users	Adaptive Computing Technology Stack	E4S Application Software Support**	Cloud-Usage Monthly Amount	Total Price per Month***
1-10 Users	\$2,500.00	\$2,500.00	\$10,000.00	\$15,000.00
11-20 Users	\$5,000.00	\$5,000.00	\$10,000.00	\$20,000.00
21-50 Users	\$7,500.00	\$7,500.00	\$10,000.00	\$25,000.00
Unlimited Users	\$10,000.00	\$10,000.00	\$10,000.00	\$30,000.00

*For customers that on-board to adaptive.hpc/ai/ml-as-a-service by March 31st, the one-time set up fee is waived.

**This cost is eliminated if you bring your own AI/ML development toolkit.

***The price per month includes the full Adaptive Computing Technology Stack, 1-HPC Head node and 2-HPC Compute nodes. Additional/different compute (such as GPUs) is priced separately.

6. Still need more information? Click on one of the links below:

- adaptive.hpc/ai/ml-as-a-service [datasheet](#)
- adaptive.hpc/ai/ml-as-a-service [application list](#)
- adaptive.hpc/ai/ml-as-a-service Video (coming soon)
- adaptive.hpc/ai/ml-as-a-service ODDC Training Demo – [Launching E4S from ODDC on AWS](#)
- adaptive.hpc/ai/ml-as-a-service Training Demos – [AWS](#), [GCP](#), [OCI](#), [Azure](#)

7. Is training available for adaptive.hpc/ai/ml-as-a-service?

Group classes are available at a cost of \$500 per day, and private classes are \$1,000 per day.

8. Are there any additional costs for adaptive.hpc/ai/ml-as-a-service?

No, there are no additional costs.

9. How do I get started with adaptive.hpc/ai/ml-as-a-service?

- Request an adaptive.hpc/ai/ml-as-a-service formal quotation [here](#)
- Complete an adaptive.hpc/ai/ml-as-a-service [on-boarding form](#)
- Email Arthur Allen at aallen@adaptivecomputing.com for more information

Sign up for a 30-day [free trial here](#)