

Overview

HPE Machine Learning Development Environment

Built upon the open-source **Determined** model development platform, **HPE Machine Learning Development Environment** allows model developers and researchers to focus on building better models faster, by reducing complexity, and removing the need to write boilerplate code associated with managing ML infrastructure.

It easily integrates with popular ML **frameworks** & **tools**, and supports cloud or on-prem **infrastructure** environments, with a consistent user experience (UX).

Our platform also makes it easy for IT and MLOps teams to setup and share AI infrastructure to improve collaboration & productivity for ML teams, and reduce costs.

With HPE Machine Learning Development Environment, you can easily:

- Train models faster
 - Build better models
 - Manage and share your AI infrastructure
 - Track and reproduce your work
 - Integrate into your enterprise
-

What's New

- Run ML and HPC jobs alongside each other on the same cluster, with support for workload managers like **Slurm®** or **PBS®** and secure container runtimes like Singularity/Apptainer, Podman, or NVIDIA® Enroot.
 - Use Role-Based Access Controls (**RBAC**) to authorize development and MLOps teams to securely collaborate and share ML resources and artifacts.
 - Train models on NVIDIA or AMD® GPUs without any code changes, with foundational support for accelerator heterogeneity.
 - Use **DeepSpeed®** for 3D-Parallel (Data-, Model-, and Pipeline-parallel) distributed training, to speed up training of large models like **GPT-NeoX**.
 - Use PyTorch® Distributed Data Parallel (DDP) for flexibility and choice of distributed training strategies.
 - Version, annotate, and organize trained models in the Model Registry so that MLOps teams can effectively collaborate with model developers to manage your models' lifecycle.
 - Define your own logic to coordinate across multiple trials within an experiment.
 - Implement your own **custom hyperparameter search** algorithms, ensembling, active learning, neural architecture search, and reinforcement learning.
 - New SKUs
 - R9T40AAE - HPE ML Dev Env SW 8 GPU 3yr Sub E-RTU
 - R9H29AAE - HPE ML Dev Env SW 3yr Sub E-RTU
 - R9Y51AAE - HPE ML Dev Env SW 4yr Sub E-RTU
 - R9Y52AAE - HPE ML Dev Env SW 5yr Sub E-RTU
-

Standard Features

Train Models Faster

Train models faster, and at-scale, without changing your model code, or writing cumbersome boilerplate to manage your ML infrastructure.

We take care of ML infrastructure provisioning, networking, data loading, checkpointing, fault-tolerance and high-availability - i.e., we make distributed model training easy, fast, and cost-efficient.

HPE Machine Learning Development Environment ships cutting-edge **distributed training** strategies and techniques:

- **DeepSpeed**: for 3D-Parallel (Data-, Model-, and Pipeline-parallel) distributed training, to speedup training of large models like **GPT-NeoX**, and beyond!
 - **Horovod**: for easy-to-use data-parallel distributed training.
 - PyTorch Distributed Data Parallel (**DDP**): for flexibility and choice of distributed training strategies.
-

Build Better Models

Find better model configurations efficiently with a production-grade implementation from the creators of the cutting-edge Asynchronous Successive Halving (**ASHA**) Hyperband for **hyperparameter tuning**.

Define your own logic to coordinate across multiple trials within an experiment.

Implement your own custom hyperparameter search algorithms, ensembling, active learning, neural architecture search, and reinforcement learning methods.

Efficiently Manage and Share Your AI Infrastructure

Efficiently manage and share your on-prem or cloud GPUs & accelerators with Machine Learning workflow-aware smart scheduling and resource management to improve productivity and collaboration for your ML development and operations teams.

Consistent User Experience for deployments ranging from laptop to a supercomputer scale, and everything in-between:

- Baremetal
- Virtual Machine (incl. cloud-native and on-prem Infrastructure-as-a-Service solutions)
- Kubernetes®
- Slurm
- PBS

Run ML and HPC jobs alongside each other on the same cluster, with support for workload managers like Slurm or PBS, and secure container runtimes like Singularity/Apptainer, Podman, or NVIDIA® Enroot.

Seamlessly use spot or preemptible instances to manage cloud costs.

Train models on NVIDIA or AMD® GPUs without any code changes, with foundational support for accelerator heterogeneity.

Track and Reproduce Your Work

Easily track and reproduce your work with experiment tracking that works out-of-the-box: covering model code, configuration, hyperparameters, metrics, and checkpoints.

Version, annotate, and organize trained models with our built-in Model Registry, enabling MLOps teams can effectively collaborate with model developers to manage your models' lifecycle.



Standard Features

Integrate Into Your Enterprise

Authenticate users using enterprise Single Sign-On (SSO) services provided by Active Directory®, Okta®, PingID®, etc., with support for OpenID Connect (OIDC) and SAML.

Integrate with user provisioning systems to automate the onboarding and offboarding of your teams, with support for SCIM.

Use Role-Based Access Controls (RBAC) to authorize development and MLOps teams to securely collaborate and share ML resources and artifacts.

Supported Hardware

HPE Machine Learning Development Environment can be **deployed** on hardware equipped with NVIDIA or AMD GPUs, on a variety of on-prem or cloud environments.

Hardware

- The master node should be configured with at least four Intel Broadwell or later CPU cores, 8 GB of RAM, and 200 GB of free disk space. The master node does not need GPUs.
- Each GPU-equipped compute node should be configured with at least two Intel Broadwell or later CPU cores, 4 GB of RAM, and 50 GB of free disk space.
 - If you are using NVIDIA GPUs, compute capability 3.7 or greater are required. These include K80, P100, V100, and A100.
 - AMD GPUs that are supported include MI210, MI250, and MI250X

Notes: Most of the disk space required by the master is for the experiment metadata database. If PostgreSQL is set up on a different machine, the disk space requirements for the master are minimal.

Supported Operating Systems and Platforms

- Red Hat® Enterprise Linux (RHEL®); SUSE® Linux Enterprise Server (SLES®); Ubuntu®

HPE Machine Learning Development Environment Products

HPE Machine Learning Development Environment is licensed on a per-GPU basis for the duration of the stated subscription term.

Description	SKU
HPE ML Dev Env SW 1-19 GPU 1yr Sub E-RTU	R8W23AAE
HPE ML Dev Env SW 20-99 GPU 1yr Sub E-RTU	R8W26AAE
HPE ML Dev Env SW 100+ GPU 1yr Sub E-RTU	R8W27AAE
HPE ML Dev Env SW 8 GPU 3yr Sub E-RTU	R9T40AAE
HPE ML Dev Env SW 3yr Sub E-RTU	R9H29AAE
HPE ML Dev Env SW 4yr Sub E-RTU	R9Y51AAE
HPE ML Dev Env SW 5yr Sub E-RTU	R9Y52AAE



Service and Support

HPE Pointnext - Service and Support

Get the most from your HPE Products. Get the expertise you need at every step of your IT journey with **HPE Pointnext Services**. We help you lower your risks and overall costs using automation and methodologies that have been tested and refined by HPE experts through thousands of deployments globally. HPE GreenLake Cloud Services, **Advisory and Professional Services**, focus on your business outcomes and goals, partnering with you to design your transformation and build a roadmap tuned to your unique challenges to help deploy AI. Our **Professional** and **Operational Services** can be leveraged to speed up time-to-production, boost performance and accelerate your business. HPE Pointnext specializes in flawless and on-time implementation, on-budget execution, and creative configurations that get the most out of software and hardware alike.

HPE Technology Services Support Credits

Offer flexible services and technical skills to meet your changing IT demands. With a menu of services that are tailored to suit your needs, you get additional resources and specialist skills to help you maintain the peak performance of your IT. Offered as annual credits, you can plan your budgets while proactively responding to your dynamic business.

HPE Support Center

The HPE Support Center is a personalized online support portal with access to information, tools and experts to support HPE business products. Submit support cases online, chat with Hewlett Packard Enterprise experts, access support resources or collaborate with peers.

Learn more at <http://www.hpe.com/support/hpesc>



Summary of Changes

Date	Version History	Action	Description of Change
06-Mar-2023	Version 1	New	New QuickSpecs



Copyright

Make the right purchase decision.
Contact our presales specialists.



© Copyright 2023 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

For hard drives, 1GB = 1 billion bytes. Actual formatted capacity is less.

a50006978enw - 17088 - WorldWide - V1 - 06-March-2023