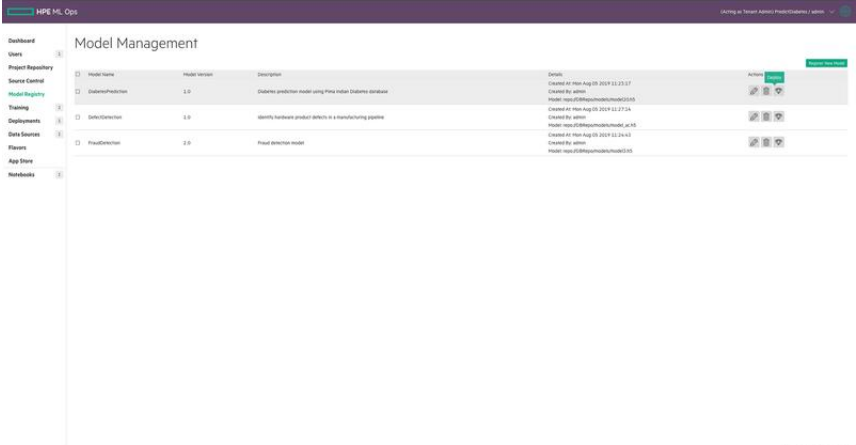


HPE EZMERAL MACHINE LEARNING OPS 5-YEAR SELECT SUBSCRIPTION 9X5 SUPPORT E-LTU (R4T73AAE)

Artificial Intelligence and Analytics



The screenshot shows the 'Model Management' interface in the HPE ML Ops console. It features a table with columns for Model Name, Model Version, Description, and Date. The table lists three models: 'DefaultPrediction', 'DefaultClassification', and 'DefaultDetection'. Each model entry includes a 'Created At' timestamp, a 'Created By' user, and a 'Model Name' field. The interface also includes a sidebar with navigation options like Dashboard, Users, Project Repository, Source Control, Model Registry, Training, Deployments, Data Sources, Flavors, App Store, and Hardware.

Model Name	Model Version	Description	Date
DefaultPrediction	1.0	Default prediction model using Prime Intellect feature extraction	Created At: Mon Aug 09 2021 11:23:17 Created By: admin Model Name: DefaultPredictionModel1000
DefaultClassification	1.0	Default classification model using Prime Intellect feature extraction	Created At: Mon Aug 09 2021 11:27:18 Created By: admin Model Name: DefaultClassificationModel1000
DefaultDetection	1.0	Default detection model using Prime Intellect feature extraction	Created At: Mon Aug 09 2021 11:28:43 Created By: admin Model Name: DefaultDetectionModel1000

WHAT'S NEW

- KubeFlow 1.3 (security hardened)?
and Model Monitoring

OVERVIEW

Much like pre-DevOps software development, data science organizations still spend a significant amount of time and effort when moving projects from development to production. Model

- Spark Operator Add-on, Spark History Server, Spark Thrift Server, Apache Livy and Hive Metastore
- Delta Lake support
- Spark versions supported: Apache Spark 2.4.7 and Apache Spark 3.1.2
- HPE Ezmeral Runtime Analytics for Apache Spark
- Introducing HPE Ezmeral Runtime Enterprise New UI with Enhanced UX: improved UX for data scientists; Packaged libraries for simplified coding experience and Notebook magics for KubeFlow

version control and code sharing is manual, and there is a lack of standardization on tools and frameworks, making it tedious and time-consuming to productize machine learning models.

HPE Ezmeral Machine Learning Ops (HPE Ezmeral ML Ops) extends the capabilities of the HPE Ezmeral Runtime Enterprise and brings DevOps-like agility to enterprise machine learning. With the HPE Ezmeral ML Ops, enterprises can implement DevOps processes to standardize their ML workflows.

HPE Ezmeral ML Ops provides data science teams with a platform for their end-to-end data science needs with the flexibility to run their machine learning or deep learning (DL) workloads on-premises, in multiple public clouds, or a hybrid model and respond to dynamic business requirements in a variety of use cases.

FEATURES

Faster Time to Value

Manage and provision infrastructure through an intuitive graphical user interface.

Provision development, test, or production environments in minutes as opposed to days.

Onboard new data scientists rapidly with their choice of tools and languages without creating siloed development environments.

Improved Productivity

Data scientists spend their time building models and analyzing results rather than waiting for training jobs to complete.

HPE Ezmeral Runtime Enterprise helps ensure no loss of accuracy or performance degradation in multi-tenant environments.

Increase collaboration and reproducibility with shared code, project and model repositories.

Reduced Risk

Enterprise-grade security and access controls on compute and data.

Lineage tracking provides model governance and auditability for regulatory compliance.

Integrations with third-party software provides interpretability.



High-availability deployments help ensure critical applications do not fail.

Flexible and Elastic

Deploy on-premises, cloud, or in a hybrid model to suit your business requirements.

Autoscaling of clusters to meet the requirements of dynamic workloads.



For additional technical information, available models and options, please reference the [QuickSpecs](#)

HPE POINTNEXT SERVICES

[HPE Pointnext Services](#) brings together technology and expertise to help you drive your business forward and prepare for whatever is next.

Operational Services from HPE Pointnext Services

[HPE Pointnext Tech Care](#) provides fast access to product-specific experts, an AI-driven digital experience, and general technical guidance to help enable constant innovation. We have reimagined IT support from the ground up to deliver faster answers and greater value. By continuously searching for better ways to do things—as opposed to just fixing things that break—HPE Pointnext Tech Care helps you focus on achieving your business goals.

[HPE Pointnext Complete Care](#) is a modular, edge-to-cloud IT environment service that provides a holistic approach to optimizing your entire IT environment, and achieving agreed upon IT outcomes and business goals through a personalized and customer-centric experience. All delivered by an assigned team of HPE Pointnext Services experts.

HPE Integration and Performance Services help you customize your experience at any stage of your product lifecycle with a menu of services based on individual needs, workloads, and technologies.

- Advise, design, and transform
- Deploy
- Integrate and migrate
- Operate and improve
- Financial Services
- Greenlake Management Services
- Retire and sanitize
- IT Training and personal development

Other related services

[HPE Education Services](#) delivers a comprehensive range of services to support your people as they expand their skills required for a digital transformation. Consult your HPE Sales Representative or Authorized Channel Partner of choice for any additional questions and support options.

Defective Media Retention is optional and allows you to retain Disk or eligible SSD/Flash Drives replaced by HPE due to malfunction.




HPE GREENLAKE

[HPE Greenlake](#) is HPE's market-leading IT as-a-Service offering that brings the cloud experience to apps and data everywhere – data centers, multi-clouds, and edges – with one unified operating model. HPE GreenLake delivers public cloud services and infrastructure for workloads on premises, fully managed in a pay per use model.

If you are looking for more services, like **IT financing solutions**, please [explore them here](#).

Make the right purchase decision.
Contact our presales specialists.

[Find a partner](#)



© Copyright 2022 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.

Parts and Materials: HPE will provide HPE-supported replacement parts and materials required to maintain the covered hardware. Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

Kubernetes® is a registered trademark of the Linux Foundation in the United States and other countries, and is used pursuant to a license from the Linux Foundation. LINUX FOUNDATION and YOCTO PROJECT are registered trademarks of the Linux Foundation.

Image may differ from the actual product
[PSN1012683988USEN](#), May, 2022.