

Operationalize AI/ML for the enterprise with cnvrg.io and Red Hat OpenShift MLOps solution

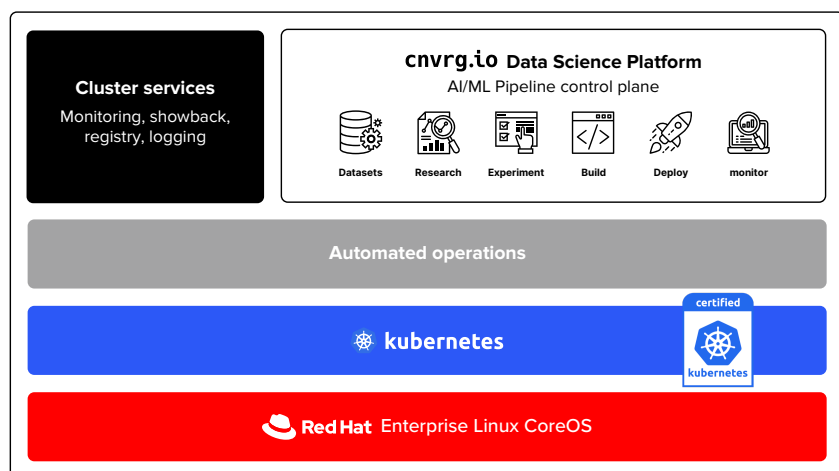


INDUSTRY CHALLENGES

Machine Learning (ML) is rapidly becoming essential to all businesses and organizations around the world. However, this means that IT and DevOps teams are now facing the challenge of standardizing ML workloads and provisioning cloud, on-premise and hybrid compute resources that support the dynamic and intensive workflows that ML jobs and pipelines require. With all the many tools, scripts, plug-ins and disconnected stacks, developers and data scientists spend over 65% of their time on DevOps and managing infrastructure resources requests and hybrid cloud compute. This manual and labor intensive work pulls them from doing what they were hired to do - deliver high impact ML models.

SOLUTION: ACCELERATE ML WORKFLOWS WITH END-TO-END MLOPS

cnvrg.io on Red Hat OpenShift delivers an out of the box solution that empowers enterprise data science and DevOps teams to better manage infrastructure in the hybrid cloud and accelerate the ML workflow in one automated and unified platform. On top of Openshift, cnvrg.io has also collaborated with NVIDIA DGX systems and NGC (Nvidia GPU Cloud) to provide a GPU optimized, distributed AI OS solution with NVIDIA GPUs and OpenShift foundations.



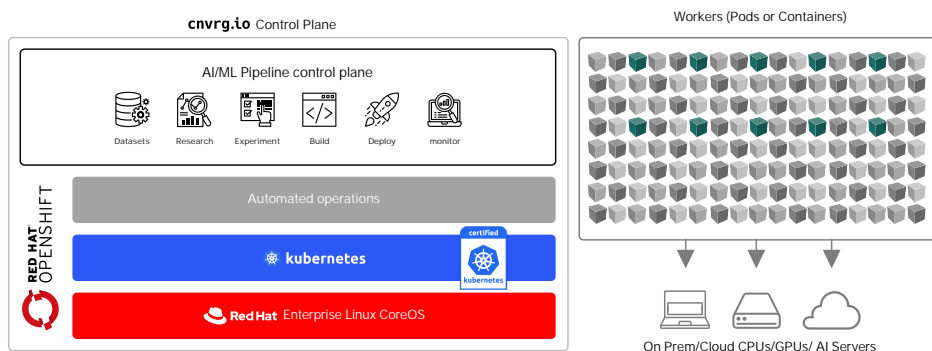
CNVRG.IO & OPENSIFT OUT-OF-THE-BOX SOLUTION

- Managed Kubernetes deployment on any cloud or on-premises environment
- Fully automated installation and life cycle management
- All tools data scientists need for ML/AI development: from research to deployment
- Open & flexible, code-first data science platform, which integrates open source tools



CNVRG.IO AND RED HAT OPENSIFT - BETTER TOGETHER

The cnvrg.io data science platform is certified for Red Hat OpenShift Container Platform and included in the OpenShift Service Catalog. The combined Red Hat and cnvrg.io solution delivers AI lifecycle management, and simplified MLOps to enterprise DevOps and data science teams across industries. Red Hat OpenShift, the leading enterprise Kubernetes application platform, enables data scientists to launch flexible, container-based jobs and pipelines, as well as providing infrastructure teams with the capabilities to manage and monitor ML workloads in a single managed and cloud-native environment. cnvrg.io then enables data scientists to rapidly launch ML workloads on remote clusters without tinkering with infrastructure or complicated configuration. For infrastructure teams, cnvrg.io provides the ability to manage all ML compute resources in a unified and secure environment with advanced monitoring and administration capabilities built in.



ABOUT CNVRG.IO

cnvrg.io is an AI OS that helps enterprises and data scientists manage, build and automate machine learning pipelines. cnvrg.io's ML platform helps data scientists spend less time on DevOps, with advanced MLOps solutions, data version control and management, experiment tracking, model management, model monitoring and more. cnvrg.io is a code-first platform and container-based making it agile and flexible to operate on any infrastructure whether in the cloud or on-prem. Its collaborative end-to-end solution enables teams to streamline machine learning and deploy high impact models in half the time. Expert and novice data scientists enjoy the flexibility to use any tool, language or environment while maintaining full control of their models in training and in production.

ABOUT RED HAT

Red Hat is the world's leading provider of open source software solutions, using a community-powered approach to provide reliable and high performing cloud, Linux, middleware, storage, and virtualization technologies. Red Hat also offers award winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT.



The AI infrastructure ecosystem is growing rapidly. We're collaborating with cnvrg.io as part of our OperatorHub to help provide an end-to-end MLOps solution that data scientists, IT and DevOps engineers need to effectively manage, build and deploy machine learning at the enterprise level. cnvrg.io is a great choice for OpenShift and GPU's to run ML workloads, and enables enhanced collaboration between data scientists and engineers for accelerated ML deployment across hybrid cloud.



Tushar Katarki
Product Manager,
Red Hat OpenShift

FIND OUT MORE

Website: www.openshift.com

Blog: www.openshift.com/blog

Website: www.cnvrg.io/solutions/red-hat-openshift/

Blog: www.cnvrg.io/blog/