ClearML Announces Broadest Support for AI Infrastructure Orchestration and Compute Management

ClearML is the first and only AI platform to have Kubernetes, Slurm, PBS, and bare metal compatibility for seamless orchestration of AI and ML workloads

SAN FRANCISCO, CA, US, May 16, 2024 /EINPresswire.com/ -- ClearML, the leading open source, end-to-end foundational platform for unleashing AI in the enterprise, today announced the release of new AI orchestration and compute management capabilities, making it the first AI platform to support Kubernetes, Slurm, PBS, and bare metal for seamless orchestration of AI and machine learning workloads. ClearML now offers the broadest support for AI and HPC workloads in the marketplace.

The company's newly released functionality enables AI practitioners to automate manual or repetitive tasks as well as offers broadened AI infrastructure management capabilities to include computing clusters using Simple Linux Utility for Resource Management (SLURM) or Altair PBS. The inclusion of Slurm and PBS, popular open-source workload managers commonly utilized in high-performance computing (HPC) environments, further enhances ClearML's offerings that supports all popular Kubernetes variants as well as bare metal.

"Customer AI deployments are entering a new era of complexity, spanning across diverse environments such as cloud, edge, and on-premises data centers," said Moses Guttmann, Co-founder and CEO of ClearML. "To navigate this complexity and ensure optimal performance, sophisticated scheduling and orchestration is paramount. ClearML's new capabilities reduce the overhead of managing and controlling AI infrastructure, empowering AI Builders to scale their AI and machine-learning workflows with unprecedented flexibility, ease and efficiency – giving organizations ultimate control over their AI infrastructure at any scale while seamlessly integrating ClearML."

Guttmann noted that ClearML also expanded its scheduling and triggering capabilities to further boost an AI team's efficiency and productivity with the ability to run tasks automatically based on predetermined times or events, an extension of ClearML's "set-it-and-forget-it" approach to eliminating manual tasks.

Now, DevOps can focus on what matters in getting AI to production, rather than spend time on laborious tasks such as storage maintenance, babysitting AI workflows, provisioning machines, or doling out credentials. ClearML continues to streamline manual and mundane tasks while decreasing friction and overhead for AI team admins, allowing them to spend less time on setup and more time on innovation so they can deliver faster time to value and drive costs down.

This announcement follows the company's most recent release (announced at NVIDIA GTC 2024), which enables granular management and visibility of compute resource allocations and includes open source fractional GPU capabilities. The company's expanded capabilities in orchestration, compute management, and AI infrastructure control establishes ClearML as the most comprehensive platform available for AI Builders and DevOps professionals, who use ClearML to build, train, and deploy models at any scale on any AI infrastructure. AI teams can work on shared data and build models seamlessly from anywhere in the world on any AI workload, compute type, or infrastructure – regardless of whether they're on-prem, cloud, or hybrid; with Kubernetes, Slurm, PBS, or bare metal; and with any type of GPU.

With ClearML, AI Builders and DevOps teams gain ultimate control and granular visibility over what resources and fractions of resources each team or group can access and easily self-serve resources without changing their existing AI/ML workflows. According to the company's recent survey report, The State of AI Infrastructure 2024, 25% of the 1,000 IT leaders surveyed stated that their company uses Slurm or another open source tool for scheduling and job management. Since Slurm is Linux-native and designed for handling AI/HPC workloads, it is also widely used by many of the world's most advanced supercomputers.

ClearML's Slurm/PBS integration enables AI teams to get more out of their Slurm/PBS computing clusters with a single line of code. AI/HPC jobs can be launched from anywhere (code, command line interface, Git, or web UI) and organizations can monitor their Slurm/PBS queues on the platform's orchestration dashboards. In this way, ClearML helps integrate HPC workloads into an organization's CI/CD infrastructure, so that customers can securely launch jobs on their cluster from an external endpoint. ClearML's Slurm/PBS support creates transparency for teams and enables them to focus expensive resources on delivering innovation rather than them standing by, idling and unutilized.

As well, ClearML's extended capabilities enable AI builders to use the ClearML platform for building scheduling logic and seamlessly passing it through to Slurm/PBS for execution. Now, organizations can leverage the best parts of Slurm/PBS without the extensive coding typically required.

To learn more, visit the company's most recent blog posts:

- ClearML Supports Seamless Orchestration and Infrastructure Management for Kubernetes, Slurm, PBS, and Bare Metal: (https://clear.ml/blog/clearml-supports-seamless-orchestration-and-infrastructure-management-for-kubernetes-slurm-pbs-and-bare-metal)

- How ClearML Helps Teams Get More out of Slurm: (https://clear.ml/blog/how-clearml-helps-teams-get-more-out-of-slurm)

Get started with ClearML by using our free tier servers (https://app.clear.ml) or by hosting your own (https://clear.ml/docs/latest/docs/deploying_clearml/clearml_server). If you need to scale your ML pipelines and data abstraction or need unmatched performance and control, please request a demo (https://clear.ml/demo).

About ClearML
As the leading open source, end-to-end solution for unleashing AI in the enterprise, ClearML is used by more than 1,600 enterprise customers to develop a highly repeatable process for their end-to-end AI model lifecycle, from product feature exploration to model deployment and monitoring in production. Use all of our modules for a complete ecosystem or plug in and play with the tools you have. ClearML is an NVIDIA DGX-ready Software Partner and is trusted by more than 250,000 forward-thinking Data Scientists, Data Engineers, ML Engineers, DevOps, Product Managers and business unit decision makers at leading Fortune 500 companies, enterprises, academia, and innovative start-ups worldwide. To learn more, visit the company's website at https://clear.ml.

Adam Brett
Crenshaw Communications
+1 212-367-9700
email us here
Visit us on social media:
Twitter
LinkedIn

This press release can be viewed online at: https://www.einpresswire.com/article/711802805