

CRAY URIKA-XC AI AND ANALYTICS SOFTWARE SUITE

Machine learning, deep learning, and large-scale analytics workloads put heavy demands on compute resources. Meet them with supercomputer power and scale. The Cray® Urika®-XC AI and analytics software suite brings data-intensive analytics, AI, and large-scale graph processing to the massively scalable Cray® XC™ series supercomputer. Combine your simulation, analytics, and AI on a single platform and transform your datasets into rich insights.

COMPLETE SOLUTION FOR SIMULATION, ANALYTICS AND AI

Data-intensive approaches demand data-intensive systems. Commodity clusters can neither support the magnitude of available data nor the machine learning and deep learning, simulation, and analytics workloads dependent on it. You need supercomputer power and scale.

The Cray Urika-XC suite of applications and tools extends the XC system's core supercomputing capabilities, enabling converged simulation and analysis workloads that fully leverage the supercomputer's massive scale.

TOOLS THAT DELIVER RICHER INSIGHTS

From an operational perspective, technology must yield results and maximize your investment. From a scientific perspective, technology must deliver transformative insights quickly.

The Urika-XC software suite meets these needs. It brings together modern tools and languages that power richer insights at scale and reduce the need for already scarce talent.

It consists of two core components:

- **Open Source Analytics.** Popular frameworks, languages, and tools enables data-intensive analytics, machine learning, and deep learning.

- **Cray Graph Engine.** A unique in-memory semantic graph database handles the largest and most demanding use cases requiring graph discovery and pattern matching.

DRIVE DEEP LEARNING TO SCALE

You can capitalize on the XC series supercomputer's scale and performance for your deep learning-powered research with Cray-developed libraries for distributed training (CrayPE ML plugin) and hyperparameter optimization (HPO).

The CrayPE ML plugin is delivered in conjunction with popular machine intelligence frameworks (TensorFlow™, PyTorch, and Keras) and leverages the Aries™ interconnect to give you scaling to more than 500 CPU and GPU nodes. Cray's HPO libraries eliminate much of the guesswork required by a data scientist to set up a complex model and reduces the time required to achieve model accuracy.

BETTER TCO FOR MORE DISCOVERY

The Urika-XC suite on the XC supercomputer gives your users a complete portfolio of simulation and data-intensive processing capabilities. When you can run analytics, AI, and simulation workloads on the same system you eliminate the cost of buying and maintaining a purpose-built analytics machine.



OPEN SOURCE ANALYTICS

The Cray open source analytics package gives you frameworks, programming environments, and distributed deep learning — all complemented by Cray expertise and support.

We've integrated popular open-source packages into ready-to-run containers designed to take advantage of the performance and scale of the XC series systems. And because we provide direct support for upgrades and bug fixes, it eliminates the need to continually integrate updated open source packages.

With the open source analytics package comes frameworks that can be used to dynamically allocate distributed clusters of nodes for large-scale work:

Apache Spark™ fast and general-purpose engine for large-scale data processing with an extensive set of libraries and packages for analytics and machine learning uses

- Java, Scala, Python and R
- MLlib, Graph-X, Spark SQL and Spark Streaming
- Maven project management
- sbt (Scala build tool)
- BigDL

Distributed Dask and Python for open data science analytics

- Anaconda Python-based open source data science tools and package management for out-of-the-box data science
- Distributed Dask for dynamically allocating a distributed memory cluster for Python-based analytic processing

Jupyter Notebook for interactive data science

- Live, shareable documents for data preparation, code development, modeling simulation, and machine learning

Deep Learning at Scale for AI

- Google TensorFlow, PyTorch, and Keras frameworks for rapid deployment
- CrayPE Machine Learning plugin and Horovod for CPU and GPU scaling and simplified ease of use
- Cray hyperparameter optimization for data scientist productivity and faster time to accuracy

Frameworks are optimized to run in conjunction with existing system workload managers and you can easily add updates and new features.

Programming with Big Data in R (pbdR) for statistical computing at scale

- Packages built to exploit the scale of the XC Series for large-scale analytics and data science

CRAY GRAPH ENGINE

The Cray Graph Engine (CGE) is a graph analytics platform optimized for large-scale pattern matching and graph discovery applications.

Leveraging the Aries network and software infrastructure, this in-memory semantic graph database enables real-time analytics on the largest and most complex datasets, using graph data structures and problems.

It features highly optimized support for inference, deep graph analysis and pattern-based queries. CGE is based on W3C industry standards and implemented using HPC technologies. It has been tested against a 400 billion collection of tuples and shows a nearly 100x performance speed-up over Spark-based GraphX.

