

UB News Highlights

The University at Buffalo's Center for Computational Research (CCR) is a leading academic supercomputing facility that maintains a high-performance computing (HPC) environment, high-end visualization laboratories, and support staff with expertise in advanced computing, modeling and simulation, and data analytics. Together these resources support CCR's core mission of: enabling research in Western New York, providing hi-tech workforce training, and fostering economic development and job creation within Western New York (WNY) and New York State (NYS).

Open XDMoD

Open XDMoD is an open source tool to facilitate the optimization and management of high performance computing systems that is widely utilized by academic, industrial and governmental HPC centers worldwide. Open XDMoD's capabilities include monitoring metrics such as utilization, providing quality of service metrics designed to proactively identify underperforming system hardware and software, and reporting job level performance data for every job running on the HPC system without the need to recompile applications.

The Open XDMoD portal provides a rich set of features accessible through an intuitive graphical interface. Metrics provided include: number of jobs, CPUs hours consumed, wait time, job sizes, and wall time, in addition to many others. Open XDMoD also supports extensive analysis of job level performance including CPU user fraction, cache and memory utilization, and I/O rates on a job by job basis. Open XDMoD provides reporting capabilities that include charting, data export, and custom report generation. Open XDMoD is available at open.xdmod.org

A version of Open XDMoD, namely XDMoD (<https://xdmod.ccr.buffalo.edu>), was developed to monitor the NSF supported portfolio of supercomputers that fall under the XSEDE program. This work is sponsored by NSF under grants ACI 1025159 and ACI 1445806.

CCR Partners with Red Hat through XDMoD and Aristotle Projects

CCR uses the RedHat supported Performance Co-Pilot (PCP) software as one of the main data providers for our Open XDMoD software. PCP runs on each compute node in our HPC cluster, collecting cpu and memory utilization, hardware performance counters, and access to GPUs and other accelerators. This data is combined with accounting data that allows job level analysis. CCR has written several data collectors for PCP and has had these accepted into the upstream PCP code. Work is ongoing to extend this data collection workflow to analyze cloud resources.

Aristotle is a \$6-million NSF Data Infrastructure Building Blocks (DIBBs) project led by Cornell, the University at Buffalo, and UC Santa Barbara. The project goal is to (1) serve as a model for providing cloud resources locally while providing elasticity and access to resources such as software, datasets, or hardware not available locally, (2) provide relevant metrics that facilitate the exchange of resources among a federation of institutions via the incorporation and further development of Open XDMoD, (3) predict the durability of AWS spot pricing via a new tool called DrAFTS which may help researchers lower their costs of computing.

