Chris Saunders HPC Article
Whitney Lacy, editor
Version 3.0

SAND2018-8887R

The Hardware of Smaller Clusters

Chris Saunders and three technologists are in high demand from Sandia's deep learning teams, and they're kept busy by building new clusters of computer nodes for researchers who need the power of supercomputing on a smaller scale. Sandia researchers working on Laboratory Directed Research & Development (LDRD) projects, or innovative ideas for solutions on short timeframes, formulate new ideas on old themes and frequently rely on smaller cluster machines to help solve problems before introducing their code to larger HPC resources. These research teams need an agile hardware and software environment where nascent ideas can be tested and cultivated on a smaller scale.

Saunders and his team at Sandia's Science and Engineering Computing Environments are successfully enabling this research by creating pipelines for emerging code—from Cloud, to containers, to virtual machines—that build the right environment quickly to help teams solve their problems in a matter of days rather than months. While the larger HPC sources are available, it's these smaller clusters that can rapidly build a foundation for teams to build on for later development on larger systems.

As an example, Sandia's 6300 Capability Stewardship is one such team that uses the smaller cluster "Ray" for its Big Data for Actionable Intelligence (BDAI) project. BDAI's goal is to provide a shared understanding of all data sources collected, which includes tweets, web camera images, map quest incidents, and construction moratorium data. These wildly different data sources are normalized into a common frame of reference so that meaningful comparisons of events in space, time, and context are possible. Sandia's Ray allows a distributed streaming processing platform to efficiently consume, process, and analyze this geospatial data. The result is a quick analysis—literally in seconds—of the raw data collected.

"We rapidly build connectivity pipelines as a service to our customers," says Saunders. "We can create the infrastructure, the platform to run on, and the software it needs at an accelerated pace, without the cost of running the larger HPC supercomputers." At least as far as deep learning teams goes, smaller is sometimes better.

Chris Saunders HPC Article
Whitney Lacy, editor
Version 3.0

POTENTIAL IMAGE (s)