

A Meaningful Internship: Come join the best and brightest minds in the world at one of the most innovative and creative multidisciplinary research institutions engaged in strategic science on behalf of U.S. national security. The work that we do at Los Alamos National Laboratory (LANL) matters to our country and the world.

HPC Data Movement and Storage Team: Upcoming Intern Project Opportunities

Project: Lustre Networking Multi-Rail Functionality and Performance Evaluation

Mentors: (Dominic Manno, Brad Settlemyer, David Bonnie)

A relatively new feature of Lustre is the ability for Lustre's networking layer to use multiple network interfaces on a single host. This adds complexity around configuration and failure modes, but enables higher single server performance and potentially better storage system efficiency. This work will include testing end-to-end functionality (interface configuration, tunable compatibility, version compatibility, failure modes, etc) and performance characterization using standard tests and tools.

Preferred skills

- Interest in HPC and storage systems
- Understanding of networking and concepts
- Strong analytical skills
- Strong linux fundamentals
- Scripting experience (bash, csh, python, etc)
- Exposure to storage systems (ZFS, ext2/3/4, XFS, lustre)

Project: Programming for Emerging Storage Systems and System Software

Mentors (Dominic Manno, Brad Settlemyer, David Bonnie)

Storage systems are evolving as technology, such as flash, becomes economically viable. Vendors implementing cutting edge hardware solutions often approach LANL to help gain insight into how these systems could move into the real world (HPC applications). Work in this area includes potential modifications to filesystems, enhancing toolsets, fixing bugs, developing new solutions, etc.

Preferred skills

- Knowledge of and interest in filesystem algorithms and data structures
- Strong experience with C programming
- Experience with Linux and Command Line Interface
- Experience with code build systems and software
- Interest in HPC and storage systems at scale



***Make a Difference** - At LANL, we're determined to harness science and imagination to make the world a better and safer place by solving complex problems others can't. In LANL's High Performance Computing Division, we are pushing the computing pendulum, designing the future, and empowering scientists across the national laboratory to make an impact. When you're in charge of making a difference, there's no limit to what you can do.*

HPC Data Movement and Storage Team: Upcoming Intern Project Opportunities

Project: Software Defined Storage Evaluation

Mentors (Dominic Manno, Brad Settlemyer, David Bonnie)

While hardware block-based storage systems are a well-known quantity; software defined block storage systems are far less so. As new technologies come to bare from vendors LANL often evaluates and offers feedback to these companies to better their product for HPC workflows. Work in this area includes filesystem configuration/tuning, testing hardware, bug fixes, finding bottlenecks anywhere in the stack in order to increase efficiency and make the storage system faster.

Preferred skills

- Interest in HPC and storage systems
- Comfortable with computer hardware
- Strong analytical skills
- Benchmarking experience
- Experience with linux and scripting (bash, csh, Python, etc)

HPC Networking Team: Upcoming Intern Project Opportunities

Project High Speed Network Monitoring (Jesse Martinez/Brett Holman, Lead Mentors)

The HPC Networking Team is responsible for many types of high-speed networks that comprise a critical part of LANL's HPC Infrastructure, including InfiniBand (IB) and Intel's Omni-Path Architecture (OPA). Over the past years, there have been many improvements to these high-speed networks that have been developed for optimizing speed and performance. This project will focus on investigating recently-released, vendor-proposed improvements, including prototype testing on our high-speed networks to determine if there are gains/benefits associated with fine tuning these fabrics and making use of all the enhancements new software stacks provide. Some of these improvements include enhanced features for dynamic routing, more robust error monitoring, routing improvements to reduce congestion/loops, etc. Additionally, the project may support the network monitoring infrastructure as it is expanding to be more modular and include more automation with existing infrastructure in HPC to make administration streamlined.

Preferred Skills

- Interest in HPC and Network Administration
- Experience with Linux Operating Systems
- Experience in scripting/compiled languages: Python, Bash, C

A Meaningful Internship: Come join the best and brightest minds in the world at one of the most innovative and creative multidisciplinary research institutions engaged in strategic science on behalf of U.S. national security. The work that we do at Los Alamos National Laboratory (LANL) matters to our country and the world.

HPC Monitoring, Data Analytics, and Security Development Team Environments Group: Upcoming Intern Project Opportunities

Project High Throughput Message Brokers (Mike Mason, Lead Mentor)

HPC systems have many of high volume data sources, specifically at LANL - LDMS, SEDC and our high speed network monitoring script, HSNMon, can each generate millions of messages and close to 1TB of data per day. These sources along with, less frequent but just as important, syslog and application data must be collected, routed through our network and delivered to several different data consumers for analysis and visualization. We wish to test several message brokers like RabbitMQ, Kafka, Prometheus and others for throughput and usefulness in our environment. Several data stores are used in our environment including Elasticsearch, OpenTSDB and Splunk Metric Indexes, determining the optimal connections between these data stores and message brokers is important for a fault tolerant and reliable HPC monitoring solutions.

Project General Infrastructure Tools Projects (Lead Mentors: Edward Rose)

The Monitoring Infrastructure Tools team is part of the user facing HPC-Environments group and is always looking for qualified students for monitoring tools related projects. Our team assists with web development, website maintenance, and server administration. Our web application projects a continuation of operations efforts often involve managing role-based access control (RBAC) to web and database content as well as writing Perl, Python and Ruby programs to interact with LDAP and MySQL servers. All of these projects involve understanding the interconnection between our Systems, Networks and Business Processes, and require strong sysadmin, scripting, and systems analysis skills. Students would become familiar with the HPC IT Request/Change Tracker system, work queue management, communications and alert routing (email, chat, Splunk) as well as learn sysadmin skills.

Preferred skills

- Interest in HPC systems, web development, systems administration
- Strong analytical skills
- IT Ticket System Experience / ITIL / CMM
- Experience with Linux (CENTOS/RHEL, SUSE, Ubuntu, etc.) and scripting (bash, Perl, Python, etc.), Object-Oriented programming languages (C, C++, C#, Ruby, etc.)
- Sendmail/List Serve type system experience
- js, Bootstrap, HTML & CSS, LAMP stack, Git, Splunk
- LDAP / OpenLDAP, SAML, Shibboleth
- Web Graphics Design / Aesthetics / Dreamweaver



***Make a Difference** - At LANL, we're determined to harness science and imagination to make the world a better and safer place by solving complex problems others can't. In LANL's High Performance Computing Division, we are pushing the computing pendulum, designing the future, and empowering scientists across the national laboratory to make an impact. When you're in charge of making a difference, there's no limit to what you can do.*

HPC Application Readiness Team: Upcoming Intern Project Opportunities

Project **Application Test Harness Integration Project**
(Howard Pritchard, Lead Mentor)

The HPC ENV AR team has recently developed a python-based application test harness system. This test harness is intended to facilitate testing of LANL proxy and production applications on new computer architectures, as well as provide for regular testing of these applications as new compilers, system software, etc. are installed on our HPC production systems.

This summer project will involve adding one or more of LANL's proxy applications to the list of applications being tested using the framework. Elements of the project include learning how to build and run the applications on an HPC cluster, writing scripts to build and run the application within the application test harness, and archiving data generated by the application in existing application performance databases.

Required skills for this project are

- advanced Python programming skills
- experience using MySQL or similar databases
- basic C/C++ programming skills

Desired skills include

- experience with a time series database such as InfluxDB
- experience with data analysis frameworks such as Grafana
- experience with Jupyter notebooks
- experience with Jenkins or similar technologies.
- Experience with HPC parallel program models such as MPI and OpenMP would also be a plus.

A Meaningful Internship: Come join the best and brightest minds in the world at one of the most innovative and creative multidisciplinary research institutions engaged in strategic science on behalf of U.S. national security. The work that we do at Los Alamos National Laboratory (LANL) matters to our country and the world.

HPC Machine Learning Team, Design Group: Upcoming Intern Project Opportunities

Project Descriptive Similarity Measures for Computer-Generated Text (Lissa Moore & Nathan DeBardeleben, Lead Mentors)

Troubleshooting HPC failures and other problems can sometimes rely on being able to quickly compare large numbers of text log files and determine the relevant similarities between them. This project involves the research and development efforts to help automate the pieces of this process that can be automated, and to generally make the human troubleshooters more efficient. The ideal end product of this work is a tool for text log similarity that can be deployed to HPC human operators and researchers, although, depending on the length of appointment, the end product may be research results that lead to tool development in the future.

Preferred qualifications

- Pursuing bachelor's degree or higher in machine learning (natural language processing preferred), statistics, applied mathematics, computer science, or a related field
- Advanced knowledge of and experience with python and related machine learning packages (scikit learn, keras, tensorflow, gensim)
- Knowledge of and experience with natural language processing techniques and human language technology, including techniques for very short documents
- Knowledge of and experience with explainable machine learning techniques (TCAV, LIME)

Project Modernizing Monitoring via Machine Learning (Lissa Moore, Lead Mentor)

The target and scope of this project is flexible and can be tailored to a student based on skills and interest. LANL has an ongoing effort focused on the use and development of machine learning techniques specifically for HPC-related problems that include relevant data gathered from the monitoring system. In the past, this has included sensor telemetry data, computer-generated text logs, and slurm job logs. Depending on the length of appointment and student level, the desired project outcome can range from an academic publication through a deployed tool based on research results.

Preferred qualifications

- Pursuing bachelor's degree or higher (PhD preferred) in machine learning, statistics, applied mathematics, computer science, or a related field
- Advanced knowledge of and experience with machine learning methods and statistics
- Advanced knowledge of and experience with python and related machine learning packages (scikit learn, keras, tensorflow)
- Knowledge of and experience with explainable machine learning techniques (TCAV, LIME)



***Make a Difference** - At LANL, we're determined to harness science and imagination to make the world a better and safer place by solving complex problems others can't. In LANL's High Performance Computing Division, we are pushing the computing pendulum, designing the future, and empowering scientists across the national laboratory to make an impact. When you're in charge of making a difference, there's no limit to what you can do.*

Questions? How to Apply? → HPCRecruits@lanl.gov

HPC Platform Design Team: Upcoming Intern Project Opportunities

Project Containerized System Services for HPC Clusters (Cory Lueninghoener, Lowell Wofford, Lead Mentors)

High Performance Computing clusters are frequently designed with a single system that hosts system services for the rest of the cluster. On LANL systems, these services include Slurm for scheduling, HTTP for system package repositories, DHCP and TFTP for system booting, Ansible for system configuration, and NTP for time synchronization. This design is simple to set up, but having all of these services run on a single system is not scalable and introduces a single point of failure to the cluster. A more flexible design containerizes these applications and runs them across several systems with high availability, potentially serving multiple clusters at once. This project will explore this idea in a concrete way, looking at which container runtime engine works best in an HPC environment and which services are easy or hard to migrate to this model.

Preferred skills

- Linux container experience with Docker, Kubernetes, or similar tools
- Experience with Linux command line tools
- Experience compiling software from source
- Experience programming with Python, C, or both
- Interest in Linux at the operating system kernel level

A Meaningful Internship: Come join the best and brightest minds in the world at one of the most innovative and creative multidisciplinary research institutions engaged in strategic science on behalf of U.S. national security. The work that we do at Los Alamos National Laboratory (LANL) matters to our country and the world.

HPC Computer System Professional (CSP) Team: Upcoming Intern Project Opportunities

Project: High Speed Network Monitoring

Mentors: (Chase Harrison/Jesse Martinez, Lead Mentors)

CSP members participate in day-to-day network activities, including supporting the ethernet and high-speed networks for the super computers and storage systems. Projects will include helping to develop more efficient and secure ways of supporting and monitoring the networks, including developing tools and procedures that could help with efficiency of the administration and of the fabrics. Projects will be focused on developing network skills from the TCP/IP stack to RDMA transports that network administrators need to support these high bandwidth/low latency technologies. Overall, you will work with the HPC Network Team to design and deploy new infrastructure and work with cutting edge technologies!

Preferred Skills

- Interest in HPC and Network administration
- Experience with scripting languages (Python, Perl, bash, etc.)
- Experience with Linux Operating Systems and Command Line Interface (CLI)

Project: HPC Storage Systems - Deployment, Integration and Administration

Mentors (Alex Parga, Jarrett Crews)

The CSP team is responsible for assisting with the integration and deployment of multi-petabyte HPC storage systems, both independently and collaboratively with other members of the team or group, after receiving initial direction and requirements from technical project leads. Specific Tasks/scenarios that the selected candidate will engage in: deploy and test new hardware; troubleshoot and diagnose of system failures; modify existing systems, software and methods; participation in knowledge sharing across teams. Storage specific work in this area includes potential modifications to filesystems, filesystem configuration/tuning, testing hardware, bug fixes, finding bottlenecks anywhere in the stack in order to increase efficiency and make the storage system faster.

Preferred skills

- Interest in HPC and storage systems
- Comfortable with computer hardware
- Strong analytical skills
- Benchmarking experience
- Experience with linux and scripting (bash, csh, Python, etc.)
- Comfortable with C programming



***Make a Difference** - At LANL, we're determined to harness science and imagination to make the world a better and safer place by solving complex problems others can't. In LANL's High Performance Computing Division, we are pushing the computing pendulum, designing the future, and empowering scientists across the national laboratory to make an impact. When you're in charge of making a difference, there's no limit to what you can do.*

Questions? How to Apply? → HPCRecruits@lanl.gov

HPC Computer System Professional (CSP) Team: Upcoming Intern Project Opportunities (Handout IX)

Project: System Administration of HPC Clusters

Mentors (Sam Sanchez, Jennifer Gonzales, Conor Robinson)

Work under the supervision of HPC System administrators to provide technical assistance in building, managing and upgrading HPC clusters ranging from 400 to 1100+ nodes. Using configuration management and version control participate in monthly upgrades. Build a production like cluster in a virtual environment for deploying and testing upgrades. Assist in the further development of Interstate, a LANL written administration tool for transitioning clusters states from production to maintenance/upgrading mode.

Preferred skills

- Interest in HPC and system administration
- Knowledge of configuration management and version control.
- Experience with Linux and Command Line Interface
- Programming experience in python or relative programming language.

A Meaningful Internship: Come join the best and brightest minds in the world at one of the most innovative and creative multidisciplinary research institutions engaged in strategic science on behalf of U.S. national security. The work that we do at Los Alamos National Laboratory (LANL) matters to our country and the world.

About the HPC Computer System Professional (CSP) Team

The High Performance Computing (HPC) Computer System Professional (CSP) within the HPC Systems Group (HPC-SYS) team provides vanguard monitoring, support, testing, and maintenance for existing supercomputers, as well as deployment support for the future systems that comprise LANL's HPC production capability. Team members receive broad technical exposure to high performance computing and develop the technical skills in system, network, and file system administration to support, maintain, and continually evolve world-class supercomputers that rely upon data movement and storage systems supporting 100 Petabytes of capacity, as well as high speed networks (including Ethernet, InfiniBand, and Omni-Path networks) operating at up to 100 Gigabit per second of throughput. The CSP Team seeks highly motivated, productive, inquisitive, and multi-talented candidates who are equally comfortable working independently as well as part of a team. Join us!

About the HPC Data Movement and Storage Team

The High Performance Computing (HPC) Data Storage Team provides vanguard production support, research, and development for existing and future systems that feed and unleash the power of the supercomputer. The Data Storage Team designs, builds and maintains some of the largest, fastest and most complex data movement and storage systems in the world, including systems supporting 100 Petabytes of capacity. We provide storage systems spanning the full range of tiers from the most resilient archival systems to the pinnacle of high-speed storage, including all-flash file systems and systems supplying bandwidth that exceeds a terabyte per second to some of the largest and fastest supercomputers in the world. Innovators and builders at heart, the Data Storage team seeks highly motivated, productive, inquisitive, and multi-talented candidates who are equally comfortable working independently as well as part of a team. Team member duties include: designing, building, and maintaining world-class data movement and storage systems; evaluating and testing new technology and solutions; system administration of HPC storage infrastructure in support of compute clusters; diagnosing, solving, and implementing solutions for various system operational problems; tuning file systems to increase performance and reliability of services; process automation.

About the HPC Networking Team

The High Performance Computing (HPC) Networking Team provides vanguard production support, research, and development for existing and future high-speed networks and interconnects that feed and unleash the power of the supercomputer. The Networking Team designs, builds and maintains some of the largest, fastest, and secure networks for both data movement and system capability in the world, including systems supporting up to 100 Gigabit per second of throughput and continuing to grow. We provide network technology spanning the full range of tiers from campus networks to highest speed cluster interconnects for some of the largest and fastest supercomputers in the world. The HPC Networking Team is responsible for all aspects of networking within the HPC environment across three separate networks. This includes Ethernet, InfiniBand, and Omni-Path networks used within more than a dozen high performance computing resources. This also includes significantly complex Ethernet networks that control and manage access between the clusters and external resources. The successful candidate will support design, deployment, and maintenance efforts for networks within the HPC environment. Builders and problem-solvers at heart, the HPC Networking Team seeks highly motivated, productive, inquisitive, and multi-talented candidates who are equally comfortable working independently as well as part of a team.



Make a Difference - At LANL, we're determined to harness science and imagination to make the world a better and safer place by solving complex problems others can't. In LANL's High Performance Computing Division, we are pushing the computing pendulum, designing the future, and empowering scientists across the national laboratory to make an impact. When you're in charge of making a difference, there's no limit to what you can do.

About the HPC Environments Group

The High Performance Computing Environments group provides direct support for usage and support of HPC systems within HPC Division, the Laboratory, and the NNSA complex.

HPC Environments designs and manages the working environment for customer computing, data management, and data analysis, and provides direct customer support and interaction at any technical depth.

The group executes scientific software engineering and research for internal and external customers and designs and executes low-level monitoring to provide real-time information for usage and support of the full HPC environment.

About the HPC Design Group

The High Performance Computing Design Group focuses on future technologies and systems related to HPC while providing technical resources when needed to the more production focused HPC Groups. Areas of focus include I/O and storage, future HPC architectures, system management, hardware accelerators, and reliability and resiliency. Production timescales of projects vary from weeks in the future for production deployments to 10 years or more for some of the reliability and future architecture work.

Where You Will Work

Our diverse workforce enjoys a collegial work environment focused on creative problem solving, where everyone's opinions and ideas are valued. We are committed to work-life balance, as well as both personal and professional growth. We consider our creative and dedicated scientific professionals to be our greatest assets, and we take pride in cultivating their talents, supporting their efforts, and enabling their successes. We provide mentoring to help new staff build a solid technical and professional foundation, and to smoothly integrate into the culture of LANL.

Los Alamos, New Mexico enjoys excellent weather, clean air, and outstanding public schools. This is a safe, low-crime, family-oriented community with frequent concerts and events as well as quick travel to many top ski resorts, scenic hiking & biking trails, and mountain climbing. The short drive to work includes stunning views of rugged canyons and mesas as well as the Sangre de Cristo mountains. Many employees choose to live in the nearby state capital, Santa Fe, which is known for world-class restaurants, art galleries, and opera.

About LANL

Located in northern New Mexico, Los Alamos National Laboratory (LANL) is a multidisciplinary research institution engaged in strategic science on behalf of national security. LANL enhances national security by ensuring the safety and reliability of the U.S. nuclear stockpile, developing technologies to reduce threats from weapons of mass destruction, and solving problems related to energy, environment, infrastructure, health, and global security concerns.

The High Performance Computing (HPC) Division provides production high performance computing systems services to the Laboratory. HPC Division serves all Laboratory programs requiring a world-class high-performance computing capability to enable solutions to complex problems of strategic national interest. Our work starts with the early phases of acquisition, development, and production readiness of HPC platforms,

***A Meaningful Internship:** Come join the best and brightest minds in the world at one of the most innovative and creative multidisciplinary research institutions engaged in strategic science on behalf of U.S. national security. The work that we do at Los Alamos National Laboratory (LANL) matters to our country and the world.*

and continues through the maintenance and operation of these systems and the facilities in which they are housed. HPC Division also manages the network, parallel file systems, storage, and visualization infrastructure associated with the HPC platforms. The Division directly supports the Laboratory's HPC user base and aids, at multiple levels, in the effective use of HPC resources to generate science. Additionally, we engage in research activities that we deem important to our mission.



***Make a Difference** - At LANL, we're determined to harness science and imagination to make the world a better and safer place by solving complex problems others can't. In LANL's High Performance Computing Division, we are pushing the computing pendulum, designing the future, and empowering scientists across the national laboratory to make an impact. When you're in charge of making a difference, there's no limit to what you can do.*