

# **ESnet 2013 Summer Student Project Descriptions**

## **Lawrence Berkeley National Lab, Berkeley, CA**

**<http://www.es.net/>**

Are you an exceptional software developer who likes working on truly challenging projects? Are you passionate about learning and open minded about the way that networks are built? Do you have a passion for organizing and visualizing data to aid in the understanding and development of network systems? Consider spending your summer with the research and development team for Berkeley Lab's Scientific Networking Division.

We're a dynamic organization, highly-motivated and focused on results. Our mission: to accelerate science by delivering unparalleled networking capabilities, tools, and innovations. As an organization, we are small enough to be agile, but large enough to offer rewarding challenges on a global scale. At the core of the Scientific Networking Division is ESnet, the Energy Sciences Network. ESnet interconnects the US national laboratory system, is widely-regarded as a technical pioneer, and is currently the fastest science network in the world.

We are working at the leading edge of software-defined networking, OpenFlow, dynamic network infrastructure, network visualization, network knowledge plane, multi-domain and multi-layer architectures. The successful student will be the one that brings strong and diverse coding skills and is very self motivated.

Interested students should send their resume to Brian Tierney: [bltierney@es.net](mailto:bltierney@es.net)

Possible Summer Student projects include (but are not limited to) the following:

### **Software Defined Networking with Openflow**

ESnet has designed and implemented a circuit reservation system called OSCARS (<http://www.es.net/oscars>). OSCARS is designed to make it easy to create Layer 2 circuit across multiple administrative domains. This works well in the wide area, where we can leverage MPLS. However in the end-site LANS this currently involves stitching VLANs together by hand. We are interested in exploring various ways in which OSCARS and OpenFlow (<http://www.openflow.org/>) could be used together to extend a layer-2 flow into the campus environment. In addition, doing experiments with OpenFlow and OSCARS with a potential paper submitted to a conference is highly desirable.

Skills required: Java and/or Python, basic networking knowledge, and Unix administration skills

### **Exploring Topics in Named Data Networking**

Named Data Networking (NDN) is an alternative approach to the architecture of computer networks. Its founding principle is that a communication network should allow a user to focus on

the data needed, rather than having to reference a specific, physical location where that data is to be retrieved from. We are interested in exploring the application of NDM techniques to large scientific data sets.

Skills required: Java and/or Python, basic networking knowledge

### **Distributed Lookup Service**

ESnet is designing a new distributed lookup service called sLS (simple Lookup Service). It has a flexible architecture containing one or more “core” nodes and zero or more caches, archival and replication nodes. We have a basic implementation in Java which allows clients to publish and search for records using a simple REST/JSON API with mongoDB as backend database. We are now working on providing a subscribe feature in the API. The goal of this feature is, when a client subscribes for a query, the lookup service will return the current set of records matching the query and will also automatically notify the client about any new records that come in at some later point in time. This feature is fundamental to our design as this allows nodes to subscribe to each other and thus create a distributed architecture.

We are looking for a student willing to help us with the implementation and testing the Subscribe feature for sLS. Additionally, he/she will also be deploying test instances of different lookup service nodes to evaluate the robustness of the design. This could potentially lead to a conference paper submission in this area.

Skills required: Java and/or Python, basic networking knowledge, and Unix administration skills

### **Developing Visualization Techniques for Network Monitoring Data**

Do you have a passion for organizing and visualizing data to aid in the understanding and development of network systems? ESnet collects a large amount of passive and active monitoring data, and we are looking for the best ways to present this data in interesting and useful ways. Current visualizations include <http://my.es.net> and <http://ps-dashboard.es.net>. We also want to develop new ways to visualize NetFlow data, and data in our distributed lookup service. We are looking for a talented and creative interface design student to help write new tools and improve existing tools.

Skills Required: Javascript, HTTP, familiarity with at least one of Java/Python/Perl

### **Building scalable multi-domain Topology Lookup Service for Network Circuit Setup**

ESnet has designed a new distributed lookup service called the sLS (simple Lookup Service). It has a flexible architecture containing one or more “core” nodes and zero or more caches, archival and replication nodes. We have a basic implementation in Java which allows clients to publish and search for records using a simple REST/JSON API with mongoDB as backend

database. As networks become automated, through the use of software-defined networking (SDN) and OpenFlow, automated network topology learning, its representation, and sharing a portion of topology becomes an important feature .

We are looking for a student willing to help us with the implementation and testing a topology service leveraging the sLS. This could potentially lead to a conference paper submission in this area and a demonstration at an important conference.

Skills required: Java and/or Python, basic networking knowledge

### **Optimizing OSCARS for a reduced memory footprint deployment**

ESnet has designed and implemented a circuit reservation system called OSCARS (<http://www.es.net/oscars>). The widely deployed OSCARS v0.6 comprises about a dozen distinct modules, each executing within its own Java Virtual Machine (JVM), and communicating with other modules via SOAP messages. This provides a tremendous plug-and-play flexibility, but requires a large memory footprint (about 10 GB) if all the modules are run on the same physical hosts. For many deployments in production networks, a smaller, more compact version of OSCARS is desired, where multiple modules are executed within a single JVM and communications between modules within the JVM uses a lighter weight communication scheme, such as REST.

Skills required: Java, Web Services, SOAP, REST, basic network knowledge, Unix.

### **Machine Learning Techniques applied to Network Monitoring Data**

ESnet collects a vast amount of network monitoring data. This includes passive monitoring data such as SNMP polling of routers and netflow data. It also includes active monitoring data such as bwctl and owamp data from perfSONAR (<http://www.perfsonar.net/>). From this vast amount of data we are looking to create a “knowledge plane” for the network that will lead to better overall understanding of our network. We are looking for a student experienced in machine learning techniques to work on analysis and correlation of this sort of network data. This could lead to publication of a paper if the results are significant.

Skills required: Java and/or Python, basic network knowledge, Unix.

### **Federated Login**

We are looking for a student to help develop and deploy Federated Login for various services at ESnet. These services are used internally by the staff as well as externally by Collaborators (Researchers and Scientists). The modern Identity Management System recommends having a Federated Login for granting access to the users. This eliminates the burden of managing the Collaborators(external users) Identity. We make the system trust the External Identity Providers such as Google, Facebook and Collaborators Enterprise Identity system itself. The student will be working on enabling Federated Login for one or more services.

Skills Required: Javascript, REST, Google & Facebook APIs, OpenID, OAuth & OpenID-Connect

### **ESnet API Development**

ESnet has a lot of data which is collected from a large variety of sources. In an effort to make this data easy to access we are working to create an API service to access this data. This position will work under the guidance of an experienced developer to add new data sets to the ESnet API. This will involve evaluating the existing data, designing a schema and implementing the changes. This API is used both internally and externally to access and update a wide variety of data related to ESnet.

(Note: This position is based at Argonne National Lab near Chicago)

Skills required: Python; Skills desired: Django, SQL (PostgreSQL or MySQL), REST, JavaScript

### **Network Metadata Collection and Curation**

A key issue in understanding network traffic and network problems is to having enough context to know what is going on. In this project the student will work on a new version of the Net Almanac project that provides just such information. The project centers on using the Simple Lookup Service (sLS) to store data about events (such as outages, large traffic drivers, etc) occurring on the network as well as information about network addresses and prefixes. This data can be used to better understand what is going on in the network. If time permits this information will be used to annotate visualizations of the network.

(Note: This position is based at Argonne National Lab near Chicago)

Skills required: Python; Skills desired: JavaScript, database design, visual design