Enabling Science Through Cyber Security At 100G

Submitted by:

Rosio Alvarez, Ph.D.
Chief Information Officer, Berkeley Lab
RAlvarez@lbl.gov

Project team:

IT Division, Cyber Security Team

- Aashish Sharma
- Craig Leres
- James Welcher
- Jay Krous - Group Lead
- Miguel Salazar
- Partha Banerjee
- Vincent Stoffer

Project Description

Summary:

Berkeley Lab's scientific needs require a high speed network with non-disruptive cyber security protections. To meet this need, Berkeley Lab's Cyber Security Team designed and implemented an innovative system capable of monitoring a 100Gbps network. This system is a reference implementation for similar institutions moving to 100Gbps networking. The system has been running in production at Berkeley Lab since January 2015 and in April 2015 the Cyber Security Team won an Innovative Technical Achievement award from the DOE CIO for their work on this project.

Background:

Lawrence Berkeley National Lab (Berkeley Lab) conducts unclassified research across a wide variety of scientific disciplines. Massive amounts of data generated by scientific instruments and high performance computers are a fundamental part of the research and analysis that takes place at the Lab. High speed networks are critical to moving this data between instruments, computational resources and collaborators within the UC, DOE, and around the world. To handle these ever-expanding network volumes, 100Gbps (100G) speeds have become the standard for the border connections of

many large research institutions, including Berkeley Lab.

While 100G links are becoming more prevalent, the ability of security operations to maintain network monitoring at this traffic volume has not kept pace. Comprehensive network monitoring is a significant challenge and can become a barrier to implementation of 100G, or worse, security monitoring requirements can be weakened to expedite an implementation. The Berkeley Lab scientific mission and fundamental approach to cyber security required that we overcome these challenges and in June 2014 the Cyber Security Team began a project to design and implement a system capable of monitoring a 100G network.

Challenges/Requirements:

Scientists, researchers and students need fast and unfettered access to their data across a vast array of devices and instruments. Cyber Security's primary requirement was to put a network monitoring system into production that met or exceeded the existing security monitoring capabilities while supporting 100G network speeds.

In order to perform comprehensive network monitoring at high speed, three components are necessary:

- a mechanism to distribute 100G network traffic to analysis hosts
- a mechanism to further divide the traffic on the host into smaller pieces
- a network intrusion detection system (IDS) to perform analysis on the traffic

When we began to search for monitoring solutions capable of operating at 100G speeds, we were unable to find any comprehensive solutions available in the commercial marketplace. Certainly pieces of a solution are available, however, no single product exists that is able to fulfill the requirements and provide all of the necessary components. Through collaboration with our peers in the R&E community, and by leveraging the team's deep experience in network monitoring, intrusion detection and systems administration we designed and implemented a system capable of meeting our requirements.

Solution:

The basic methodology of our solution is to break down the 100G network connection into smaller pieces of traffic, while preserving the affinity of individual sessions to a single analysis process. We then distribute that analysis across many dozens or hundreds of worker processes, allowing the system to scale up to speeds of 100G. We also implemented a unique traffic reduction mechanism called shunting to further shed

load from the analysis pipeline.

We designed the system to be modular, adaptable and flexible. Components of the system can be adjusted to fit the budget, capabilities and resources of other institutions. Our specific approach to 100G network monitoring uses the following technological components:

- Arista switches to distribute 100G network traffic
- Myricom network interface cards to further divide the traffic on the host
- Bro Network Monitor for IDS

Arista:
Arista, a commodity network hardware provider, offers equipment which supports 100G network interfaces. The Arista equipment also supports a feature called Data Analyzer (DANZ) which allows the device to operate solely as a traffic distribution device for network monitoring. We feed our optical network taps to the Arista device and perform a specific type of traffic distribution (called symmetric hashing) to evenly distribute the traffic to multiple analysis nodes at 10G speeds. This is the first step of breaking 100G traffic down into pieces small enough to distribute for IDS analysis.

Myricom:
As the traffic arrives from the Arista device at 10G, this is still too much for a single IDS process to handle. We use the Myricom network interface cards (and Sniffer 10G software) to receive the traffic and further divide it into smaller pieces which can be fed directly to a Bro worker process. This replicates the previous distribution step but at the host level. To ensure distribution across CPU's on the host, individual CPU cores are bound to a fraction of the traffic from the Myricom card. This finally reaches a volume that can be effectively analyzed by a single Bro IDS worker process.

Bro:
A powerful network monitoring platform and IDS, Bro was developed at Berkeley Lab in the 1990s. Bro is now maintained by the International Computer Science Institute (ICSI) mainly through NSF grants and in collaboration with partners including NCSA and Berkeley Lab. Bro natively supports a cluster configuration which allows many Bro “workers” to operate on smaller segments of a larger volume of traffic. As described above, these small segments reach the Bro software and distributed analysis can be performed across multiple CPU cores and/or distinct server nodes. A central “manager” correlates and collects the logs, and performs notices and actions for the Bro cluster. Bro’s powerful and flexible capabilities make it especially well suited for the security challenges of open networks like Berkeley Lab. During the traffic analysis, Bro can also

categorize traffic it believes is uninteresting and remove it from the monitoring pipeline through a process called shunting.

Shunting:

We implemented a novel approach to reduce analysis load on the IDS. Bro classifies well known but uninteresting flows and dynamically removes them from analysis. Using Bro's reaction framework, Arista's JSON API and custom scripting, the shunting system can surgically remove specific unwanted large data flows while still maintaining security awareness, context, and metadata for the flow. As a result of this near real-time feedback loop, a vast amount of uninteresting traffic bypasses analysis by the IDS, reducing the requirements for IDS hosts and processing resources. Berkeley Lab is able to reduce the amount of traffic processed by the IDS from 20-25Gbps to 2-5Gbps, a savings of almost 10x.

Solution diagram:

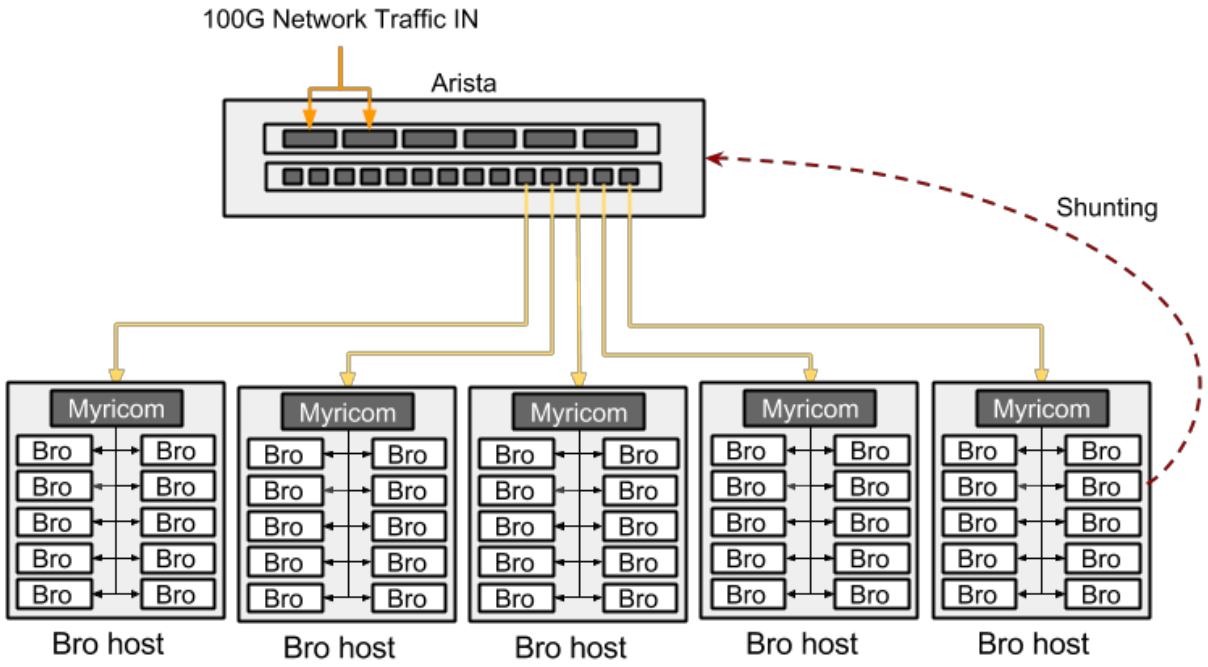

Status/Timeline:

The project commenced in June 2014 and we put the system into production in January 2015. We continue to make improvements and conduct measurements and testing for shunting policies, hardware settings, and IDS configurations. We have been assisting other institutions, including other UCs, in their own 100G deployments and are finishing our reference documentation to be released in Summer 2015.

Why does the project deserve a Sautter Award?

- The project enables research needs of the institution, scientists, and students
- Novel approach to a complex technical challenge with no commercial solution
- A reference implementation to be broadly shared
- Significant cost savings by using commodity hardware and open-source software
- Collaboration between UC, DOE and other research institution colleagues to define the architecture
- On time completion meant no disruption for other groups or services as the 100G network was put into production - the key success criteria for the project

Customer satisfaction data:

Our periodic survey of Lab scientists and staff demonstrates a high degree of satisfaction (97% positive) with the Cyber Security program in general. This reflects overwhelming support for our non-disruptive approach to security monitoring which includes the capabilities provided by the 100G monitoring project.

"Berkeley Lab's solution for 100G monitoring is elegant, affordable, and highly effective. Hundreds of US campuses, including most of the UCs, are now deploying secure enclaves for data transfer ('Science DMZs') and upgrading to 100G technology. These campuses need a solution for real-time traffic analysis that does not impair scientific productivity, and Berkeley Lab's open design is an excellent model for others to adopt."

Gregory Bell, Ph.D.
Director, Energy Sciences Network (ESnet)
Director, Scientific Networking Division
Lawrence Berkeley National Laboratory

“The Berkeley Lab Cyber Security Team's 100G monitoring solution is a novel and exciting design that allows LBLnet to offer high speed scientific networking with full cyber security protections at a reasonable cost. Their use of cost-effective off-the-shelf components and open source software and their plan to publish implementation details has generated significant interest from research and education institutions from around the nation and the world.”

Rune Stromsness
Group Lead, IT Networking Department (LBLnet)
Lawrence Berkeley National Laboratory