CERIAS

The Center for Education and Research in Information Assurance and Security

PULSAR: Purdue Live Security Analyzer for High-Performance Research Network

Lauren Featherstun[1,4], Shivam Trivedi[1,3], Nathan DeMien[1,4], Callum Gundlach[2], Jacob Sharp[2], Brian Werts[3], Ida Ngambeki[1], Lev Gorenstein[3], Erik Gough[1,3], Preston Smith[3], Xiao Zhu[3], Sonia Fahmy[5]

1. Purdue Polytechnic Institute, 2. William Harrison High School, West Lafayette, IN, 3. Research Computing, Information Technology at Purdue (ITaP), 4. Security and Policy, ITaP, 5. Department of Computer Science, Purdue University

Abstract

As more disciplines leverage computational and data-driven modeling, the security of campus cyberinfrastructure is becoming increasingly important in order to protect intellectual property and secure a competitive advantage for researchers. Funded by the NSF Cybersecurity Innovation for Cyberinfrastructure (CICI) program, the Purdue Live Security Analyzer (PULSAR) project aims to enhance the security of Purdue's campus cyberinfrastructure by developing a cyber attack detection and response capability for the Purdue campus research network. Goals of the project include enabling domain scientists to conduct research with heightened security requirements and supporting cybersecurity education by engaging undergraduate students in the deployment and operation of advanced cyberinfrastructure. The implementation of PULSAR was led by a team of Purdue undergraduate students alongside mentors from ITaP Security and Policy and Research Computing.

Background

- The Purdue research network follows the Science DMZ model and is *designed* and *optimized* for high-performance scientific applications.
- Science DMZ Components
  - A "friction free" network path operating at high speeds with minimal packet loss
  - Use of dedicated systems for data transfer (Globus/GridFTP)
  - Regular network performance measurement and testing (perfSONAR)
  - Tailored security policies to *not* impact high-performance science environments

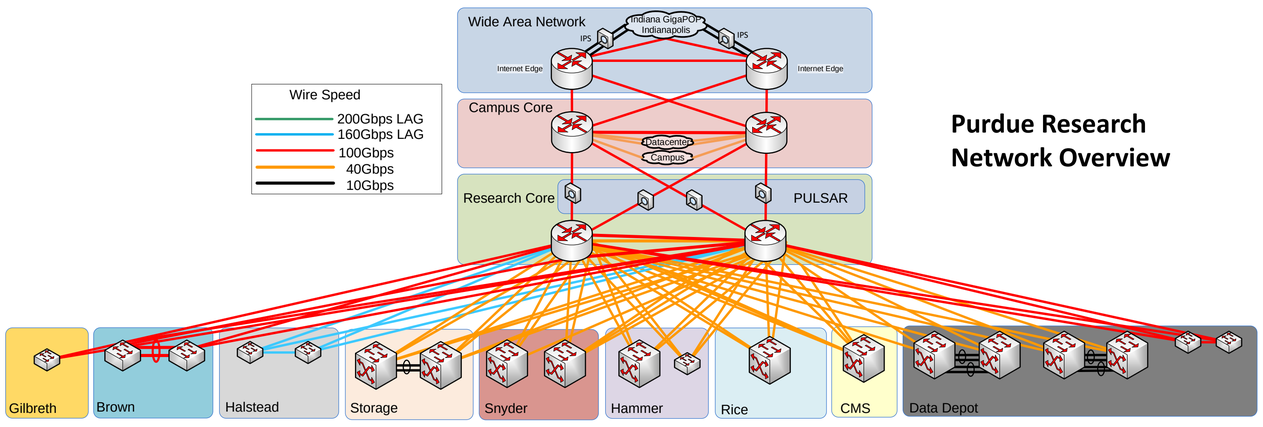

Purdue Research Network Overview

Challenges

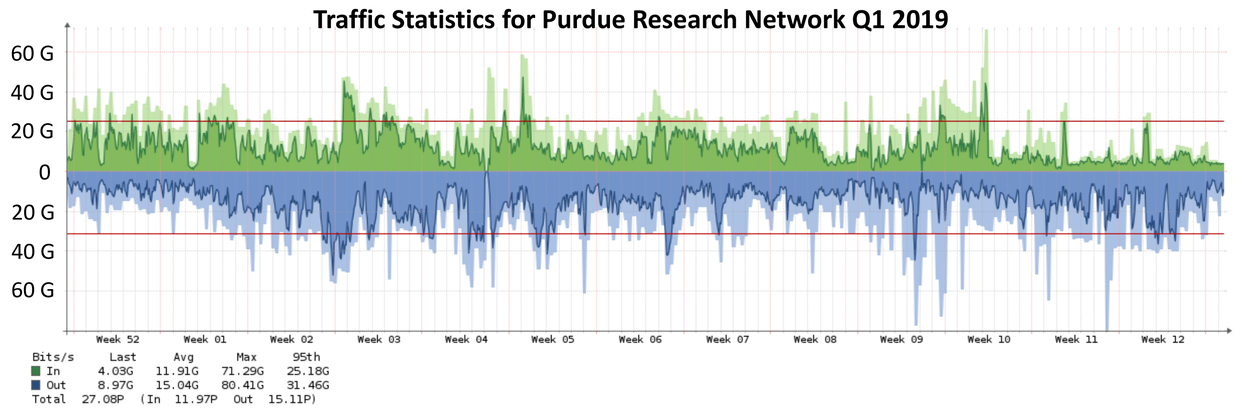

- Huge volume (steady-state traffic to campus systems between 10 and 20 Gbps, frequent network bursts of 60 Gbps through our 100 Gbps WAN link)
- Serving a large number of types of science, as well as a mixture of industry and academic collaborations, and educational work, study, and outreach
- Lots of encrypted data
- Typical inline IDS/IPS solutions are not feasible as they introduce large amounts of latency, going against the Science DMZ model

Implementation

PULSAR System Architecture and Analysis Flow

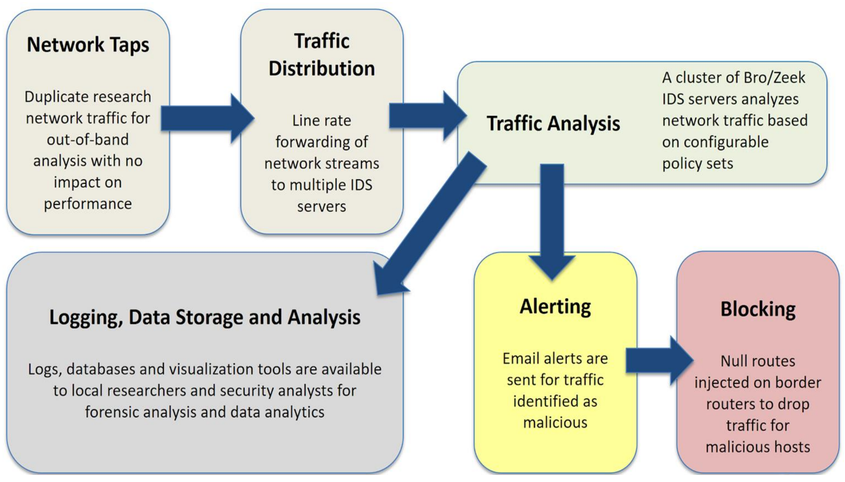

PULSAR Network Architecture, TAPs and Traffic Distribution

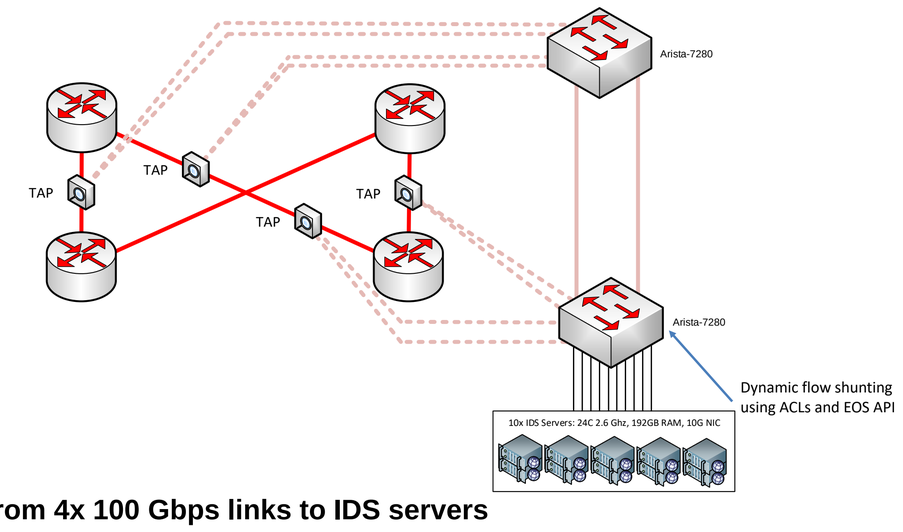

PULSAR Logging and Alerting

- PULSAR detects and separately logs any suspected security incidents
- Sends email alerts about specified incidents to security staff to be acted upon
- Local policy determines which security incidents should be separately logged, and which incidents should be alerted upon
- Example of our logging and alerting policy for SSH
  - All SSH traffic is logged
  - SSH bruteforcing attempts are logged separately
  - Only successful SSH bruteforcing logins cause notifications to be sent
- Compares all traffic against community intelligence to flag known bad actors
- Runs custom scripts including a script to do SSL fingerprinting and detect Bitcoin mining
- Blocking malicious hosts via null routes is in a testing phase

Getting data from 4x 100 Gbps links to IDS servers

- Test Access Points (TAPs) – Passively transmit full duplex TX/RX signals to 2x TX signals to Arista 7280 switches configured as Network Packet Brokers
- Traffic Distribution – Arista switches perform symmetric 5-tuple hashing (src/dst IP, src/dst port, protocol) on incoming network streams, forwarding all packets from a single stream out a single egress port
- Traffic Analysis – A cluster of ten Zeek network monitoring servers
- Flow Shunting – Zeek reaction framework, Dumbno and the Arista EOS API work together to detect bulk connections and apply dynamic ACLs to the Arista traffic distribution device, preventing Zeek workers from being overwhelmed by high speed single stream transfers

IDS Cluster Management

- Zeek IDS servers are managed via xCAT and Puppet
- Servers are stateless, boot over the network, and the OS runs in RAM

Student Mentoring

- Each student intern is paired with a mentor from the project team and mentors act as a guiding hand to lead the students
- Students document their work as they go from understanding a problem and performing research, to deciding and implementing a solution
- Technical details are discussed in bi-weekly open forums attended by all interns, mentors and project staff
- A rigorous student assessment plan is in place to determine levels of student motivation, self-directed learning, engagement, and career interest

Early Results

- On the first day of deployment, PULSAR detected a host on Purdue's internal network scanning the entire private address space for open ports associated with web and database servers
- PULSAR logged communication between an infected host and known malware servers and TOR networks

Acknowledgment

The project is supported by NSF CICI project (#1738981). We also would like to thank our former PULSAR team members Sagar Narayan and Lipu Wu.

PURDUE UNIVERSITY Discovery Park

CERIAS