BERKELEY LAB

LAWRENCE BERKELEY NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY

100G Monitoring at LBNL

Vince Stoffer

Cyber Security Engineer at Berkeley Lab

security@lbl.gov

Agenda

- Background
- 100G monitoring challenges
- Berkeley Lab solution
- Questions

80 Years of World-Leading Team Science at Lawrence Berkeley National Laboratory

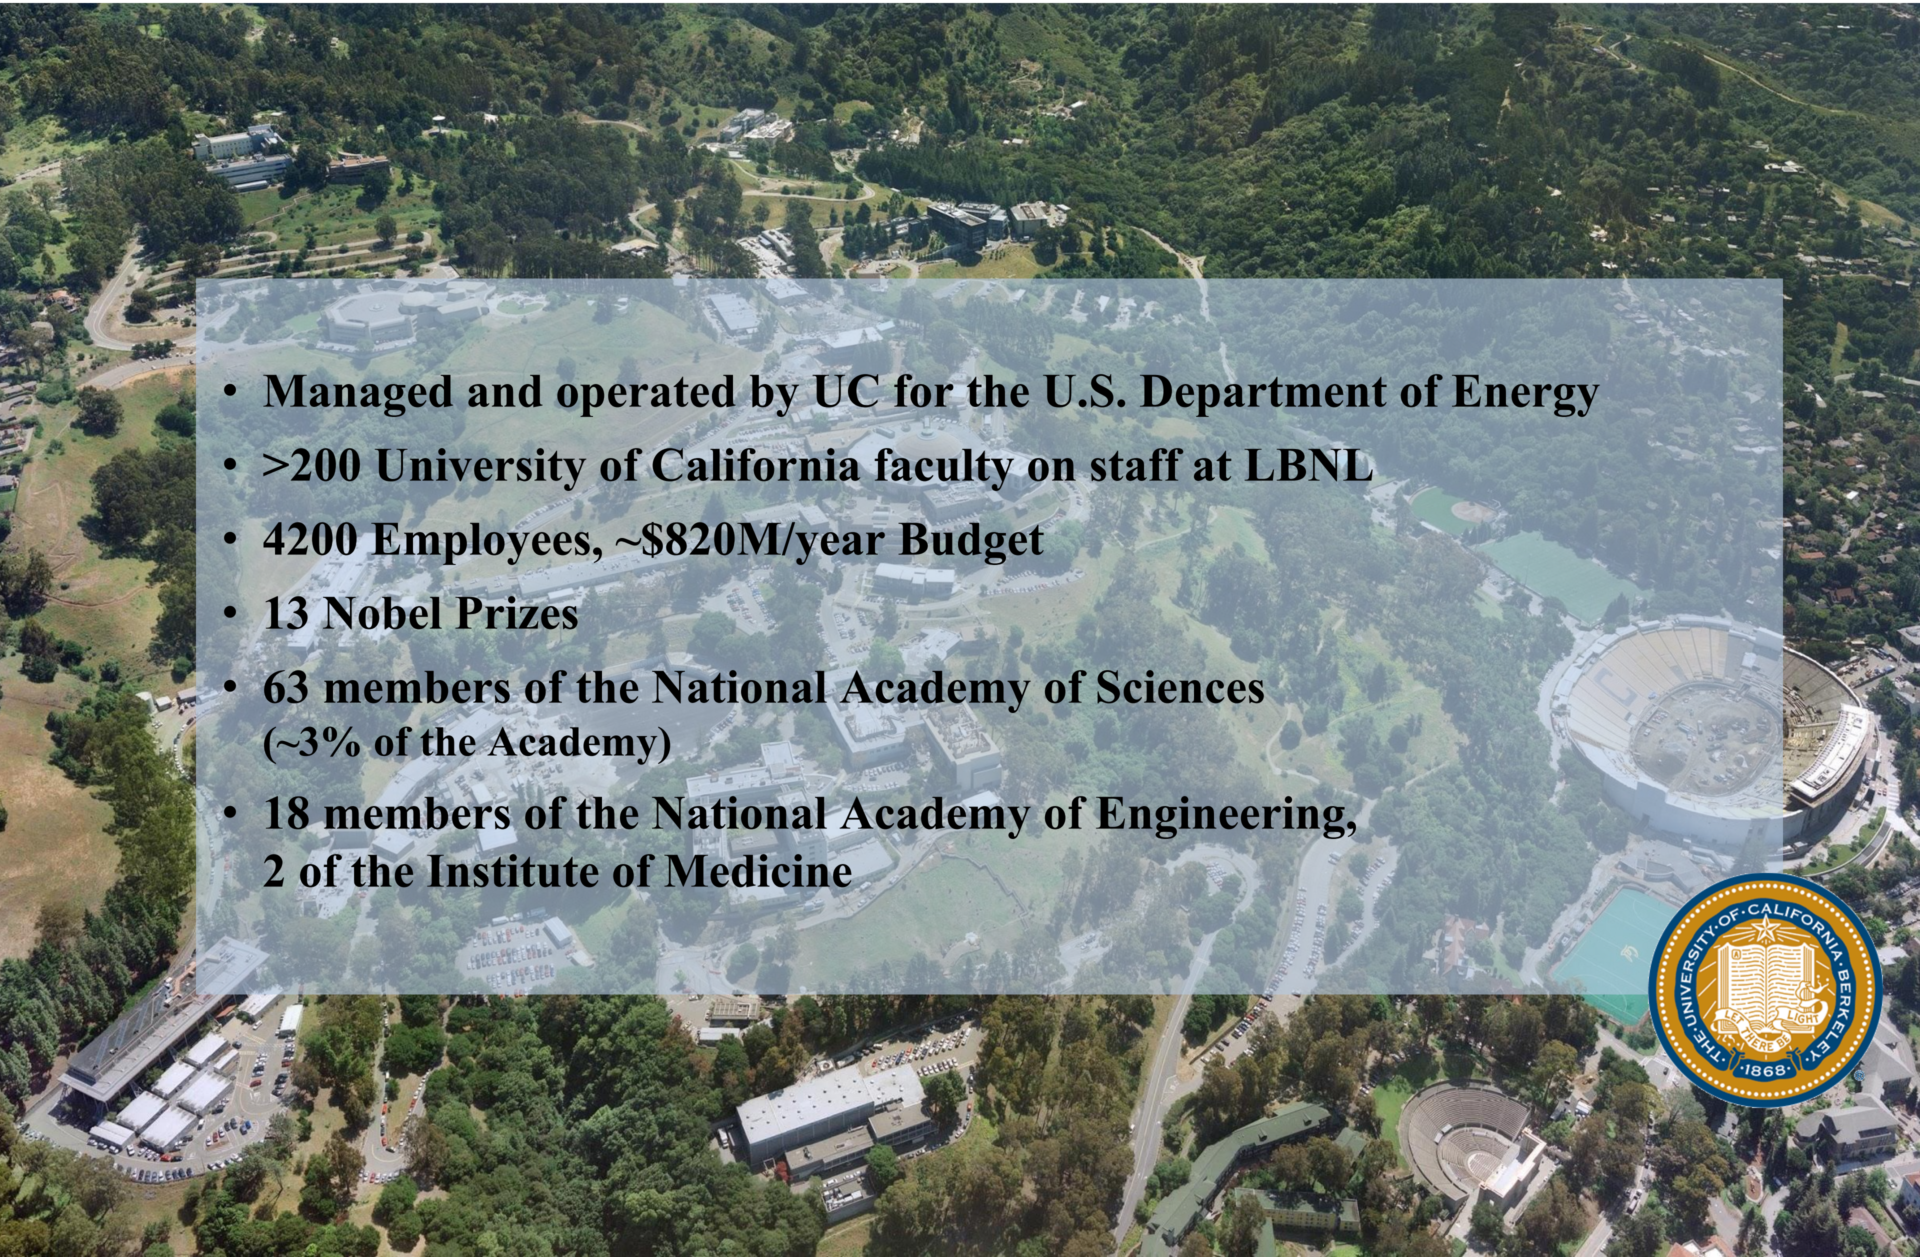

- **Managed and operated by UC for the U.S. Department of Energy**
- **>200 University of California faculty on staff at LBNL**
- **4200 Employees, ~$820M/year Budget**
- **13 Nobel Prizes**
- **63 members of the National Academy of Sciences (~3% of the Academy)**
- **18 members of the National Academy of Engineering, 2 of the Institute of Medicine**

World-Class User Facilities
Serving the Nation and the World

Advanced Light Source

Joint Genome Institute

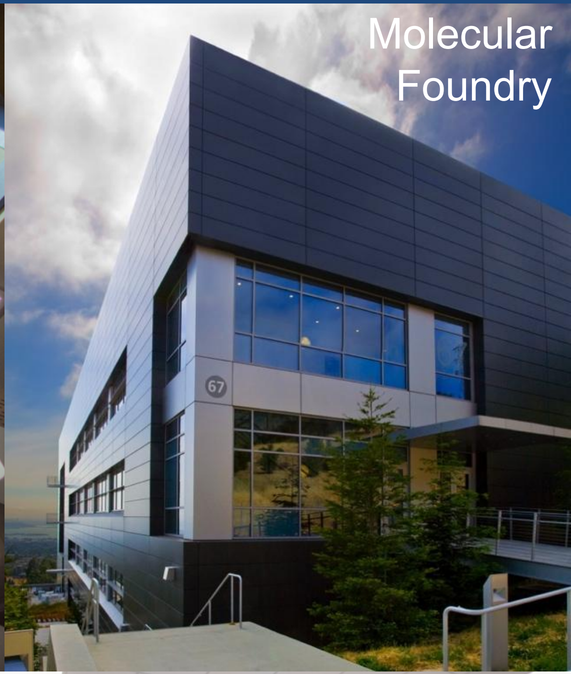

Molecular Foundry

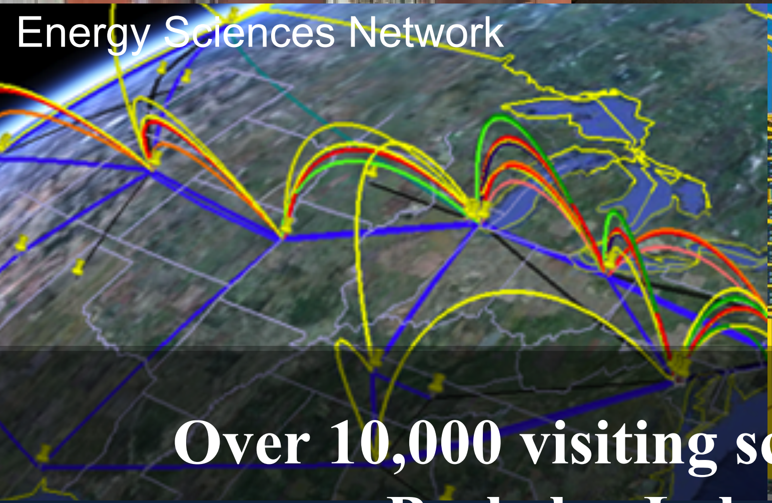

Energy Sciences Network

FLEXlab

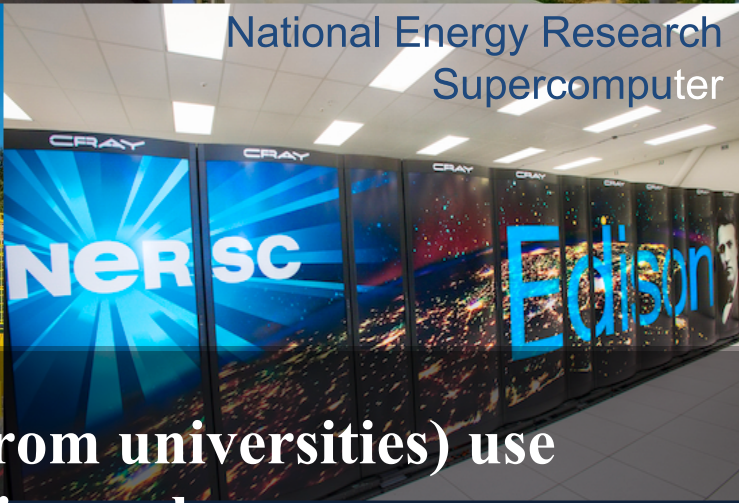

National Energy Research Supercomputer

Over 10,000 visiting scientists (~2/3 from universities) use Berkeley Lab research facilities each year

"Scientific progress will be completely unconstrained by the physical location of instruments, people, computational resources, or data"

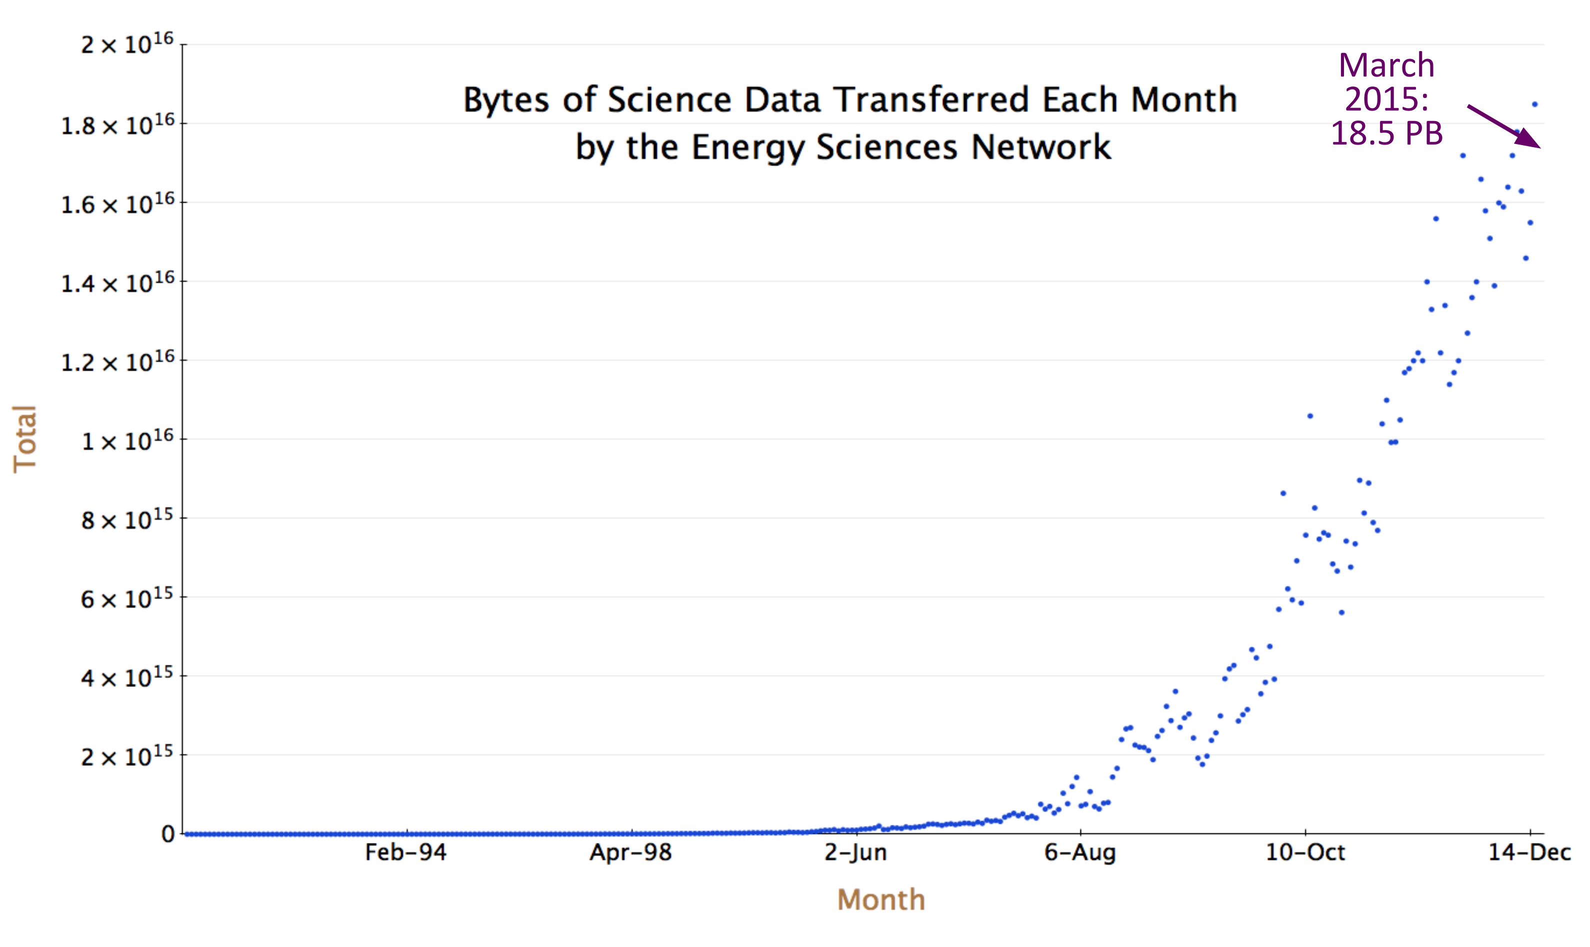

100G monitoring challenges

- No commodity solution
- 100G interfaces expensive
- Ability to scale up

Solution Overview

- Scale up a Bro cluster
- New components
  - Traffic distribution
  - Host distribution
  - Shunting

What is Bro?

www.bro.org

Bro platform

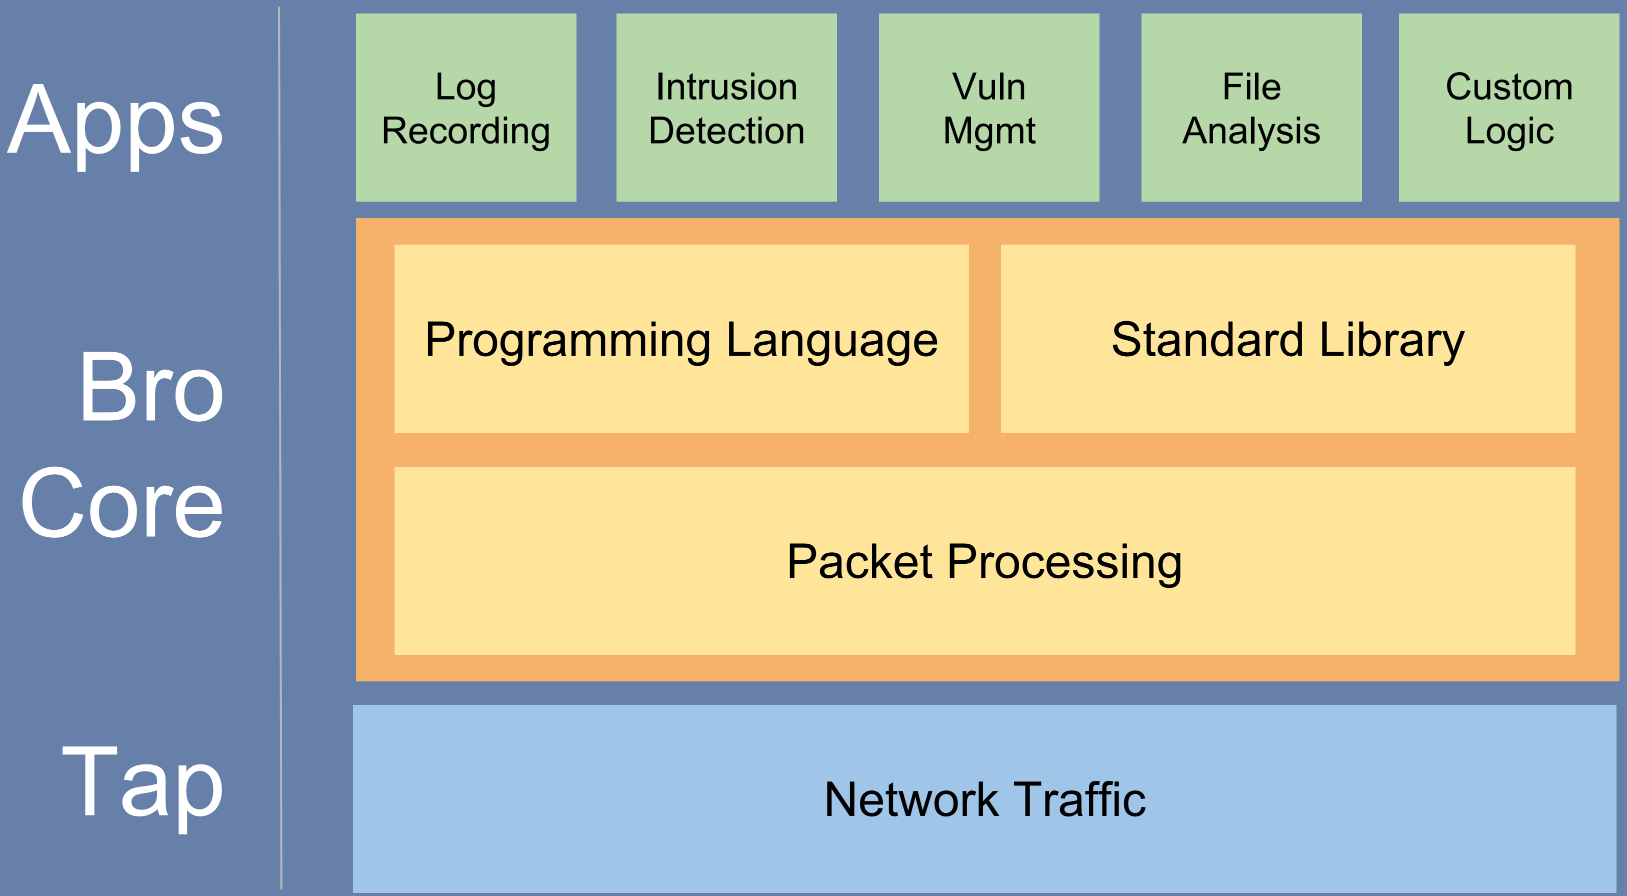

Apps

- Log Recording
- Intrusion Detection
- Vuln Mgmt
- File Analysis
- Custom Logic

Bro Core

- Programming Language
- Standard Library
- Packet Processing

Tap

- Network Traffic

Bro Clustering

- Native in Bro
- Scales horizontally
  - Across nodes and local CPUs
- Manager for all configs and logs

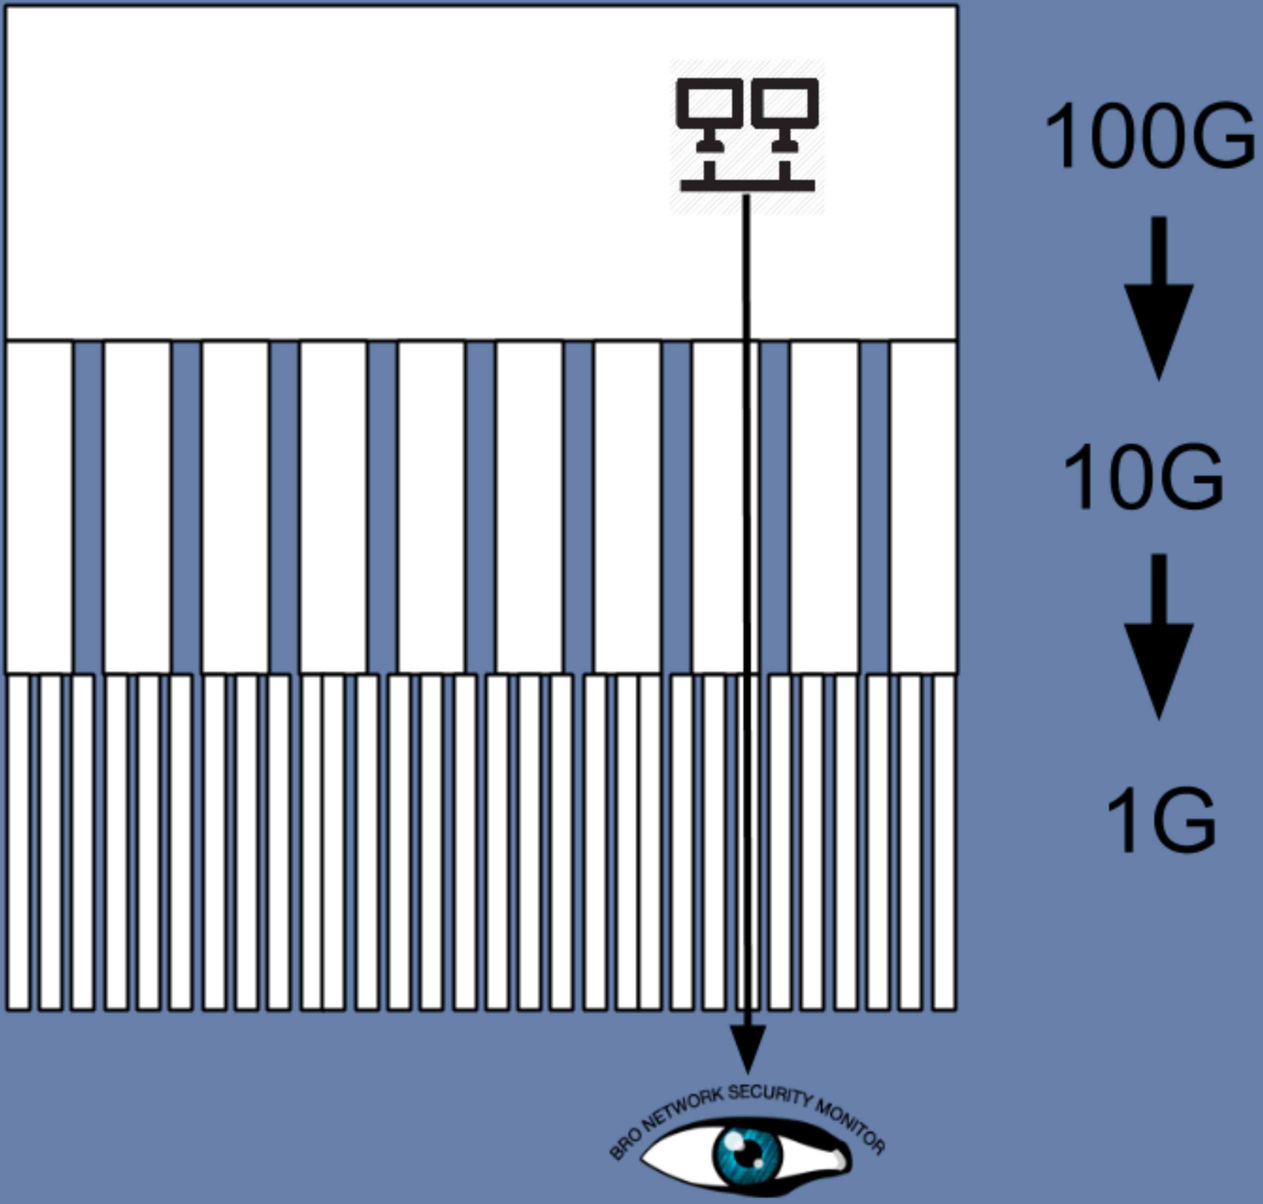

100G

10G

1G

Traffic Distribution - Arista

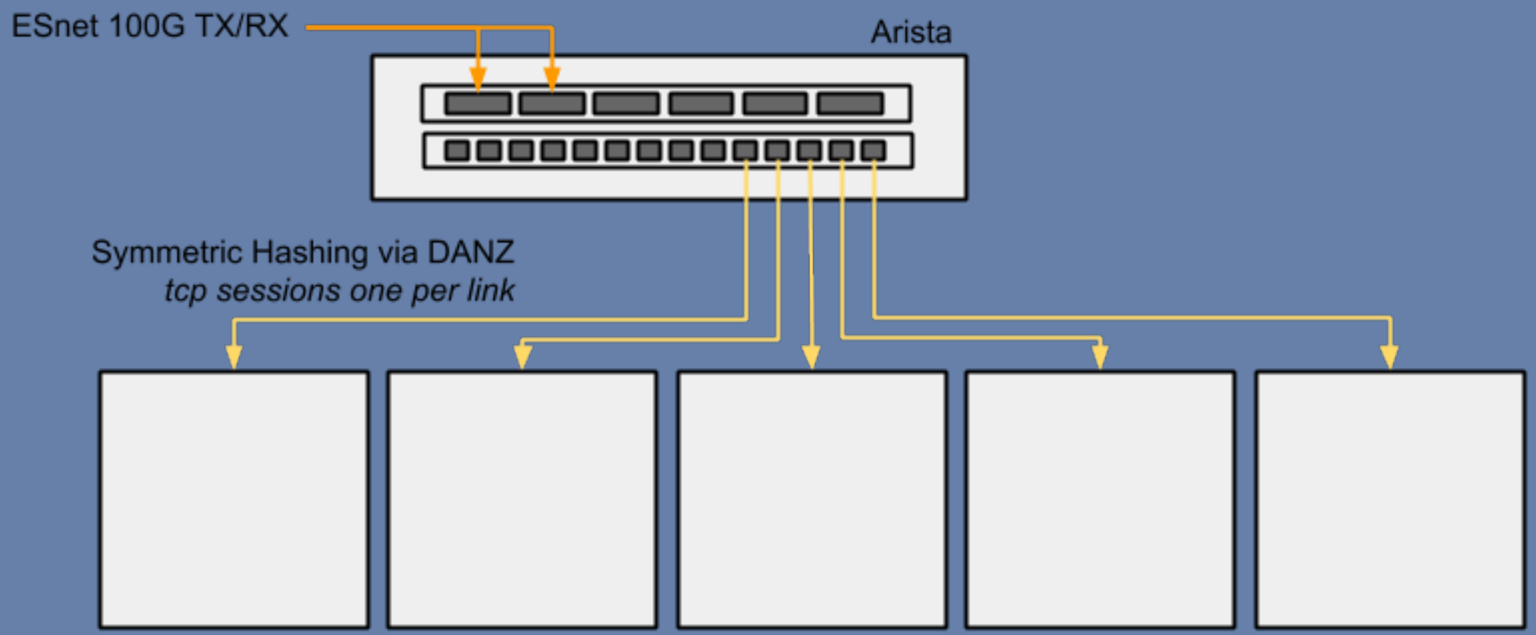

Why we chose Arista

- DANZ
- Easy to use API
  - dynamic shunting!
- Relatively low cost
- Lots of peers using
- Flexible interface including GUI

Host Distribution - Myricom

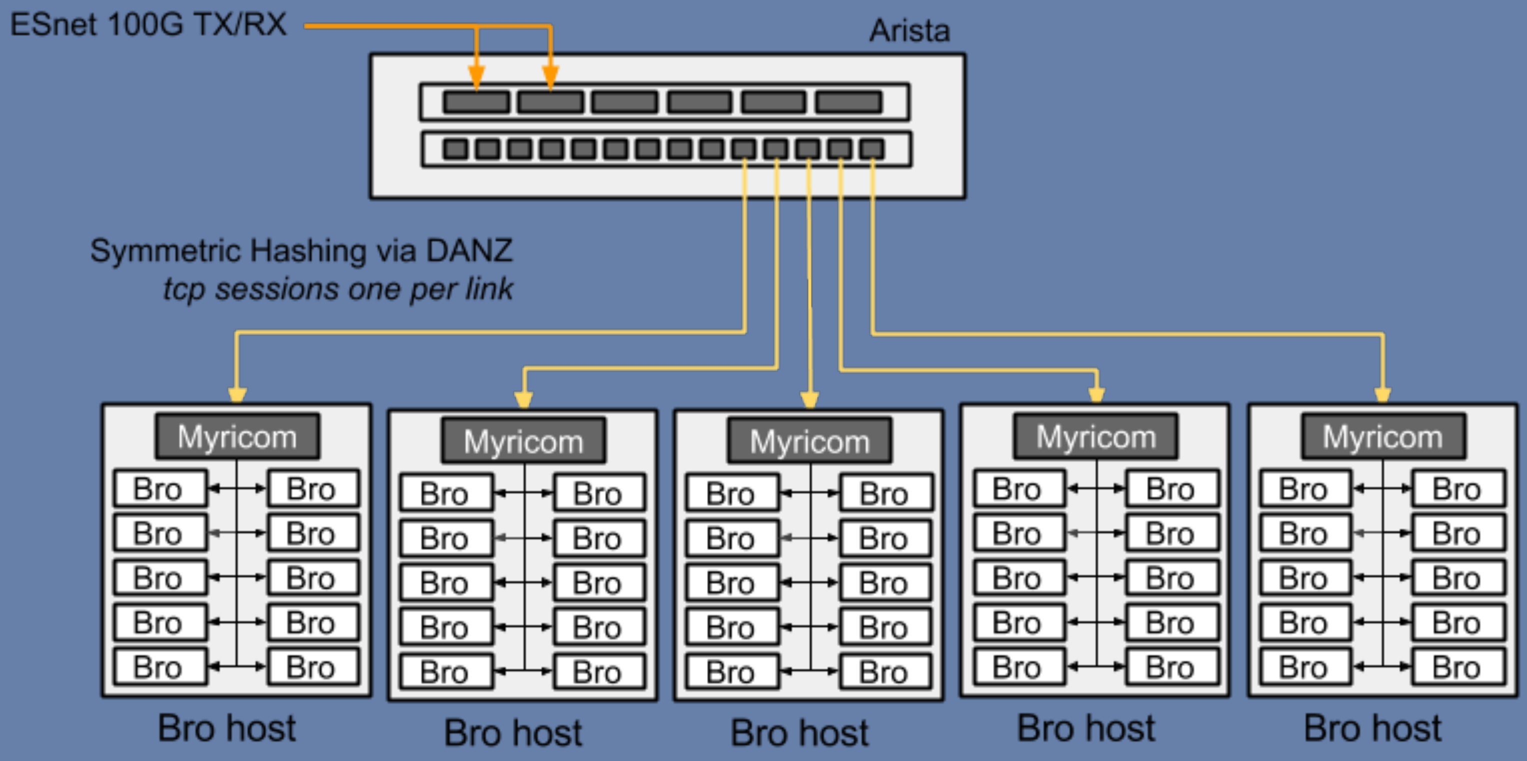

Why we chose Myricon

- Sniffer10G
  - Support for Linux, FreeBSD
  - Myricom 10G cards only
  - Supports multiple tools in 3.0

Myricom feeds to Bro workers

- Each server
  - One myricom card
  - 10 Bro processes
    - each getting 1/10 traffic
    - each pinned to a CPU
- Add servers to scale

Shunting

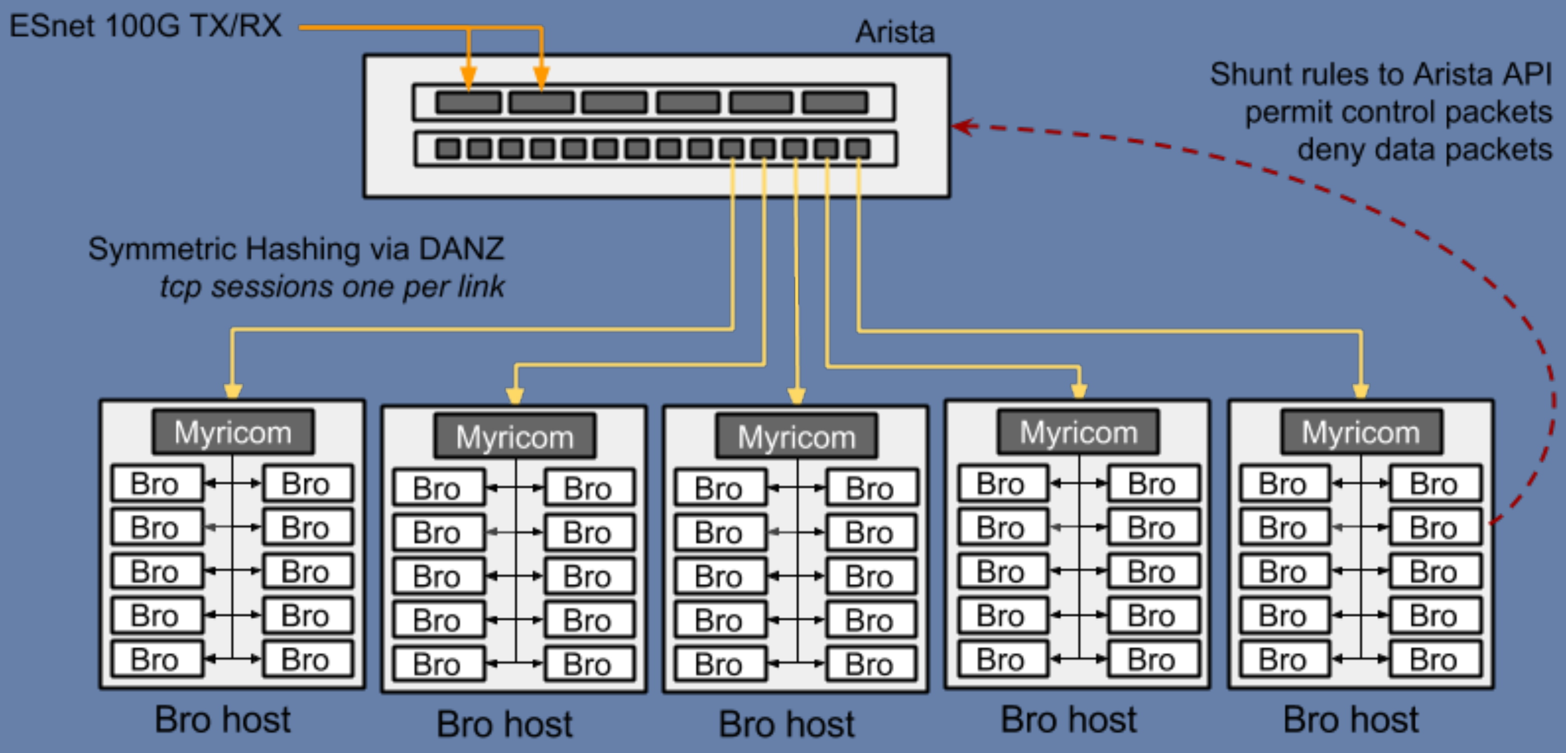
ESnet 100G TX/RX
Arista
Shunt rules to Arista API
permit control packets
deny data packets
Symmetric Hashing via DANZ
tcp sessions one per link
Myricom
Bro
Bro host

Shunting philosophy

- “Heavy Tail Effect*” a small number of flows will dominate the overall volume of data
- Detect and remove the data component of “heavy tail” flows, analysis load is reduced

*Scott Campbell NERSC

Filtering large data flows

Past:

- Nothing
- Static filtering of IPs
- Rigid
- Difficult

Shunting:

- Dynamic
- Allow control traffic
- Near real time
- Targeted
- Adaptable

Shunting script

- Python program for shunting
  - by Justin Azoff - NCSA
- Uses Arista JSON API to add ACLs which allow only control packets
- Bro’s reaction framework feeds data real-time

Deny rules example

- Connection details are preserved
  - Allow control packets
  - Deny data packets
  - Bro conn logs maintained

Shunting examples

- Bro dynamically determines protocol
- HTTP and SSH
  - shunt after 128Mb
- GridFTP (Globus)
  - shunt after 2Mb
  - harder due to:
    - multiple streams
    - changing ports

Shunting in action, April 16

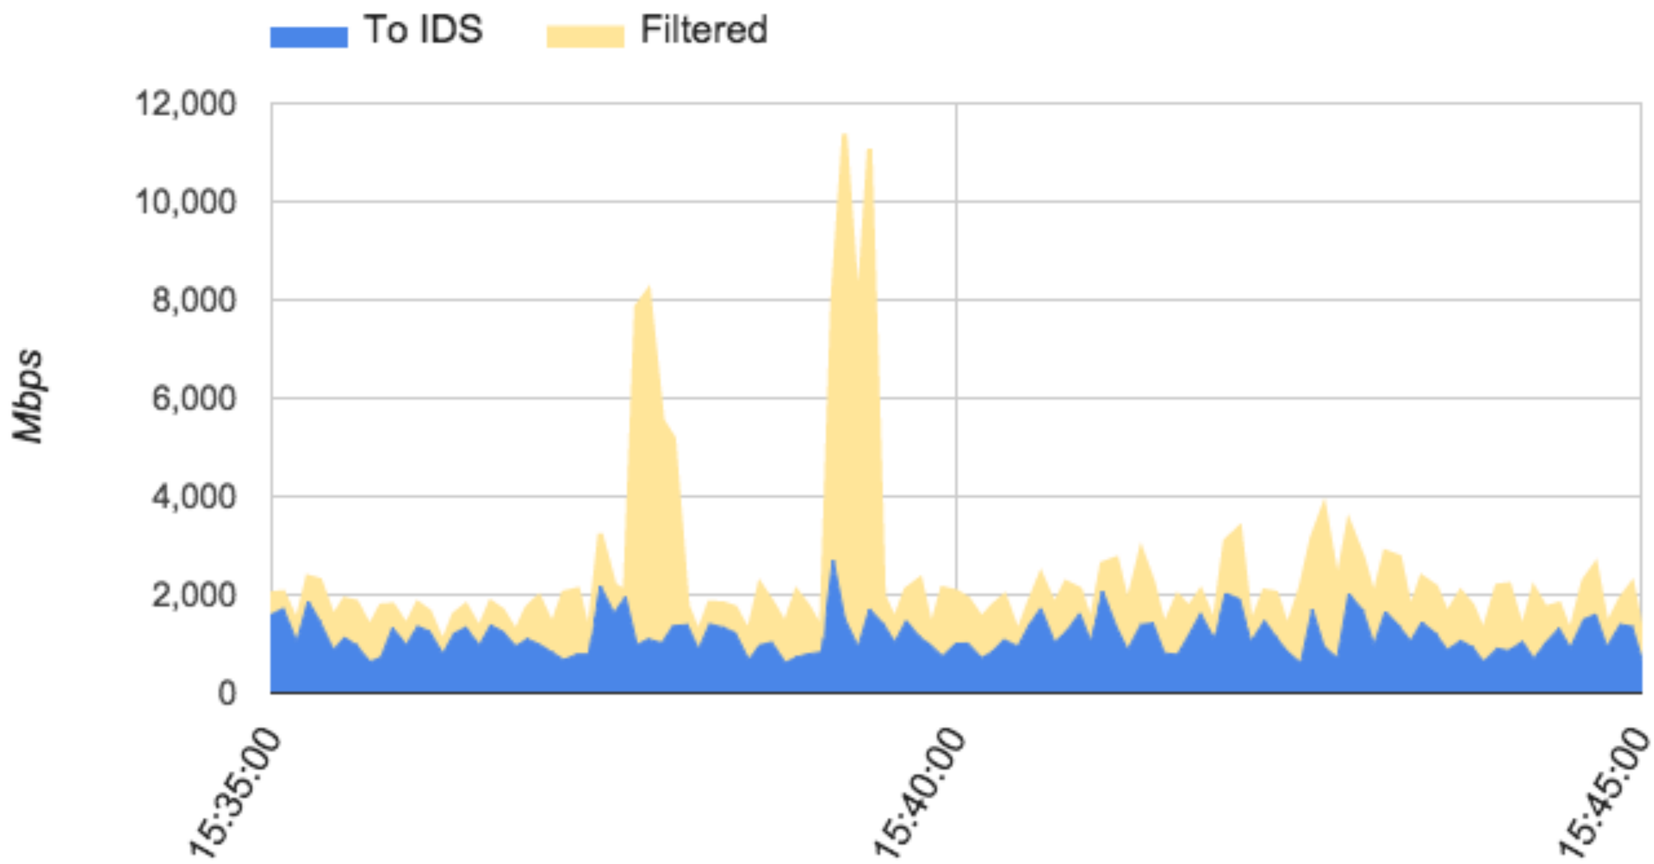

Status

- In production since Jan 2015
- Seeing average traffic 3-5 Gbps with spikes to 20-25 Gbps
- Shunting reduces this to 1-10Gbps
- Can handle to 50Gbs - add more hardware to scale up further

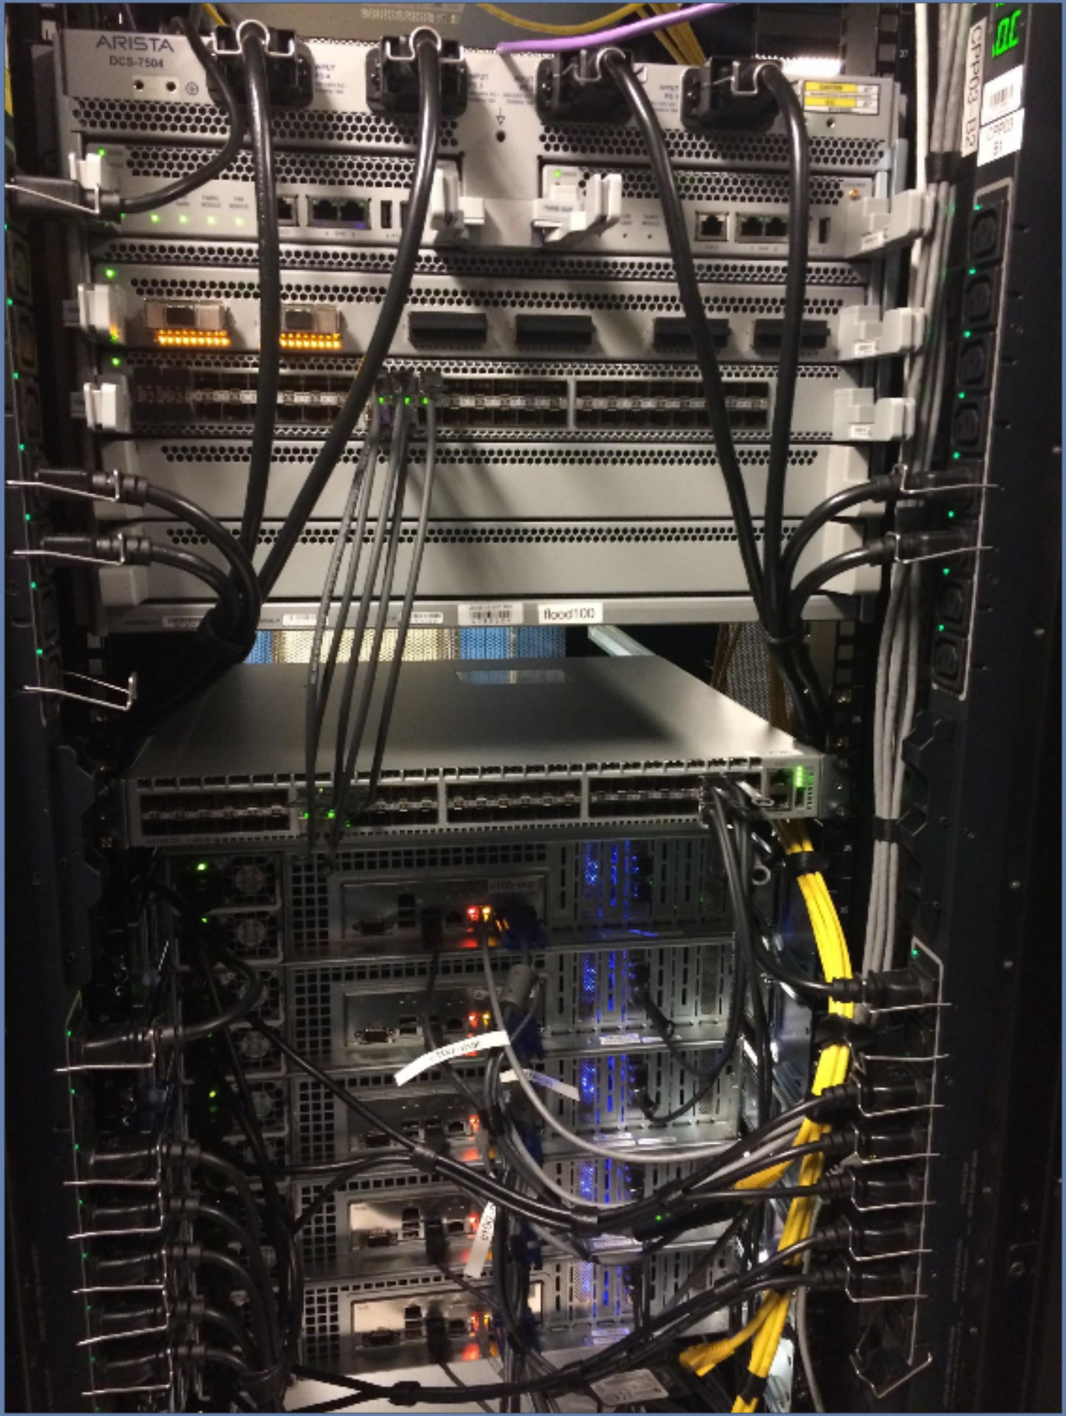
ARISTA
DCS-7504
flood100
CPP03-B2

Alternative architectures

Traffic Distribution	Host Distribution	IDS	OS
● Arista	● Myricom + sniffer drivers	● Bro	● FreeBSD
● Brocade ● Endace ● Gigamon ● OpenFlow	● PF_RING ● Packet Bricks + netmap ● Endace DAG	● Snort ● Suricata	● Linux

Next steps

- Berkeley Lab 100G technical doc
- Multiple 100G links!
- Experiment with shunting thresholds and other protocols

Discussion / Questions?

security@lbl.gov