IT Roadmap
CONFERENCE & EXPO

IT ROADMAP:
SAN FRANCISCO

The Bro Cluster

A High-Performance NIDS Architecture for the Lawrence Berkeley National Lab

Robin Sommer
International Computer Science Institute, &
Lawrence Berkeley National Lab

Agenda

- Cyber Security at the Lawrence Berkeley National Lab
- The Bro Network Intrusion Detection System
- Addressing the 10G Challenge with the *Bro Cluster*
  - ★ Architecture
  - ★ Prototypes
  - ★ Upcoming production setup
- ★ Summary & Outlook

The Lawrence Berkeley National Laboratory (LBNL)

- National Lab managed by UC for the Department of Energy
- Main site located on a 200-acre area in the Berkeley hills

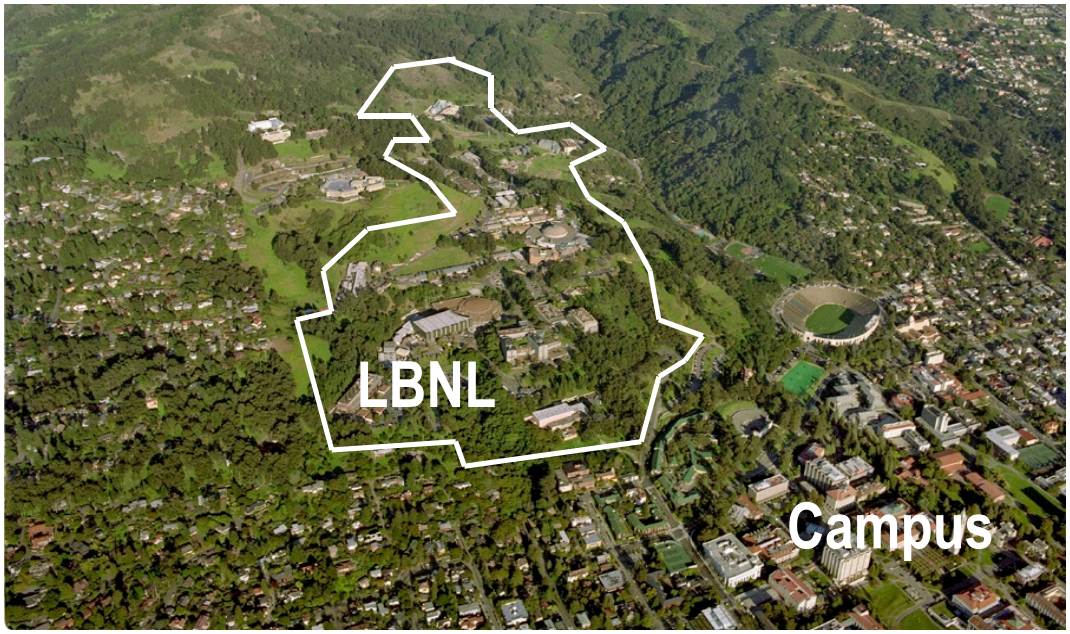

ITRoadmap
CONFERENCE & EXPO

BERKELEY LAB

Threat Landscape at LBNL

- LBNL conducts open, unclassified research
  - Research is freely shared
  - Collaborations around the world
- Wide range of research areas and very diverse user community
  - Nanotech, Energy, Physics, Biology, Chemistry, Environmental, Computing
  - Scientific facilities used by researchers around the world
  - About 3,800 employees, and 10,000 computer systems
  - Many users are transient and not employees
- Very liberal, default-allow security policy
  - Characteristic challenge for many research environments
  - What do you look for if everything is assumed to be ok?

ITRoadmap CONFERENCE & EXPO

BERKELEY LAB

Comprehensive Approach to Network Security

- **Monitoring external activity**
  - **In-depth inspection of all border traffic with the Bro NIDS**
- **Monitoring internal activity**
  - **Border taps do not see everything (internal traffic, encrypted communication)**
  - **Technology installed by the Cyber Security team throughout the Lab:**
    - Analyzing NetFlow collected from internal routers
    - Monitoring unused subnets
    - Monitoring ARP traffic
    - Central, mandatory syslog'ing for *all* Unix hosts (1.2GB/day)
    - Instrumented SSH daemons (under development)
    - Tool box for isolating hosts (external connectivity, null routing, deny boot)

Monitoring the Lab's Border with the Bro NIDS

- Bro NIDS is deployed at the LBNL upstream router
  - Sees every packet coming in or going out of the Lab
  - Can actively block attackers (and does so for about 4000 addresses a day!)

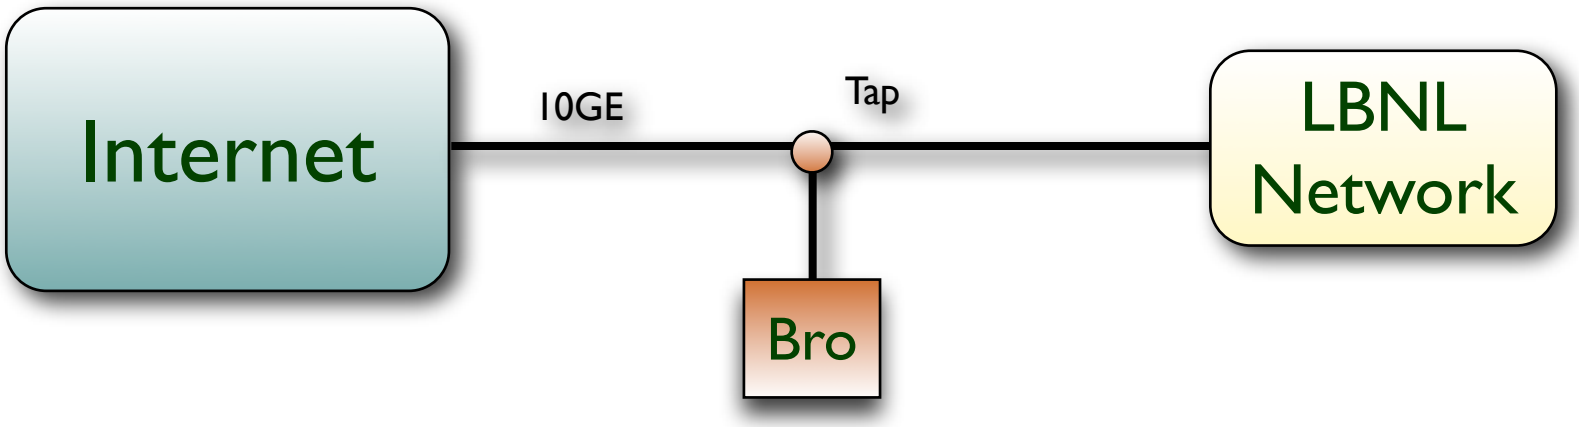

ITRoadmap
CONFERENCE & EXPO
BERKELEY LAB

The Bro Network Intrusion Detection System

- Bro is being developed at ICSI & LBNL by Vern Paxson since 1996
  - Open-source platform for in-depth monitoring on commodity hardware
- Focus is on
  - Application-level semantic analysis (rather than analyzing individual packets)
  - Tracking information over time
- Strong separation of mechanism and policy
  - The core of the system is *policy-neutral* (no notion of "good" or "bad")
- Activity-based analysis model
  - Operators program local policy using *domain-specific language*
  - Bro logs all activity comprehensively

ITRoadmap CONFERENCE & EXPO

BERKELEY LAB

The Bro Network Intrusion Detection System (2)

- Bro's analysis model differs fundamentally from other NIDS
  - Doesn't (primarily) rely on Snort-style signatures nor on anomaly detection
- Bro does require some effort to use effectively
  - Pretty complex, script-based system; no GUI, just ASCII logs
  - Requires pretty good understanding of the network
- Development is primarily driven by research
  - Lacking resources to fully polish the system
- However, our overall goal is *operational* usability
  - Goal is to bridge gap between research and operational deployment
  - Bro has been in use operationally at LBNL for more than 10 years now

ITRoadmap CONFERENCE & EXPO

BERKELEY LAB

Facing the 10G Challenge with Bro ...

- NIDSs have reached their limits on commodity hardware
  - Keep needing to do more analysis on more data at higher speeds
  - Single PC just cannot cope with >=1GE packet streams
  - Even before LBNL upgraded to 10G, it had a set of individual Bro boxes
- To address this challenge, we built the *Bro Cluster*
  - Allows us to continue operating the Bro NIDS on commodity hardware
- Cluster performs *transparent* load-balancing across PCs
  - Yields same results as a single NIDS would if it could analyze all the traffic
  - Scalable to a large number of nodes
  - Single system acts as the user interface

ITRoadmap CONFERENCE & EXPO

BERKELEY LAB

Architecture

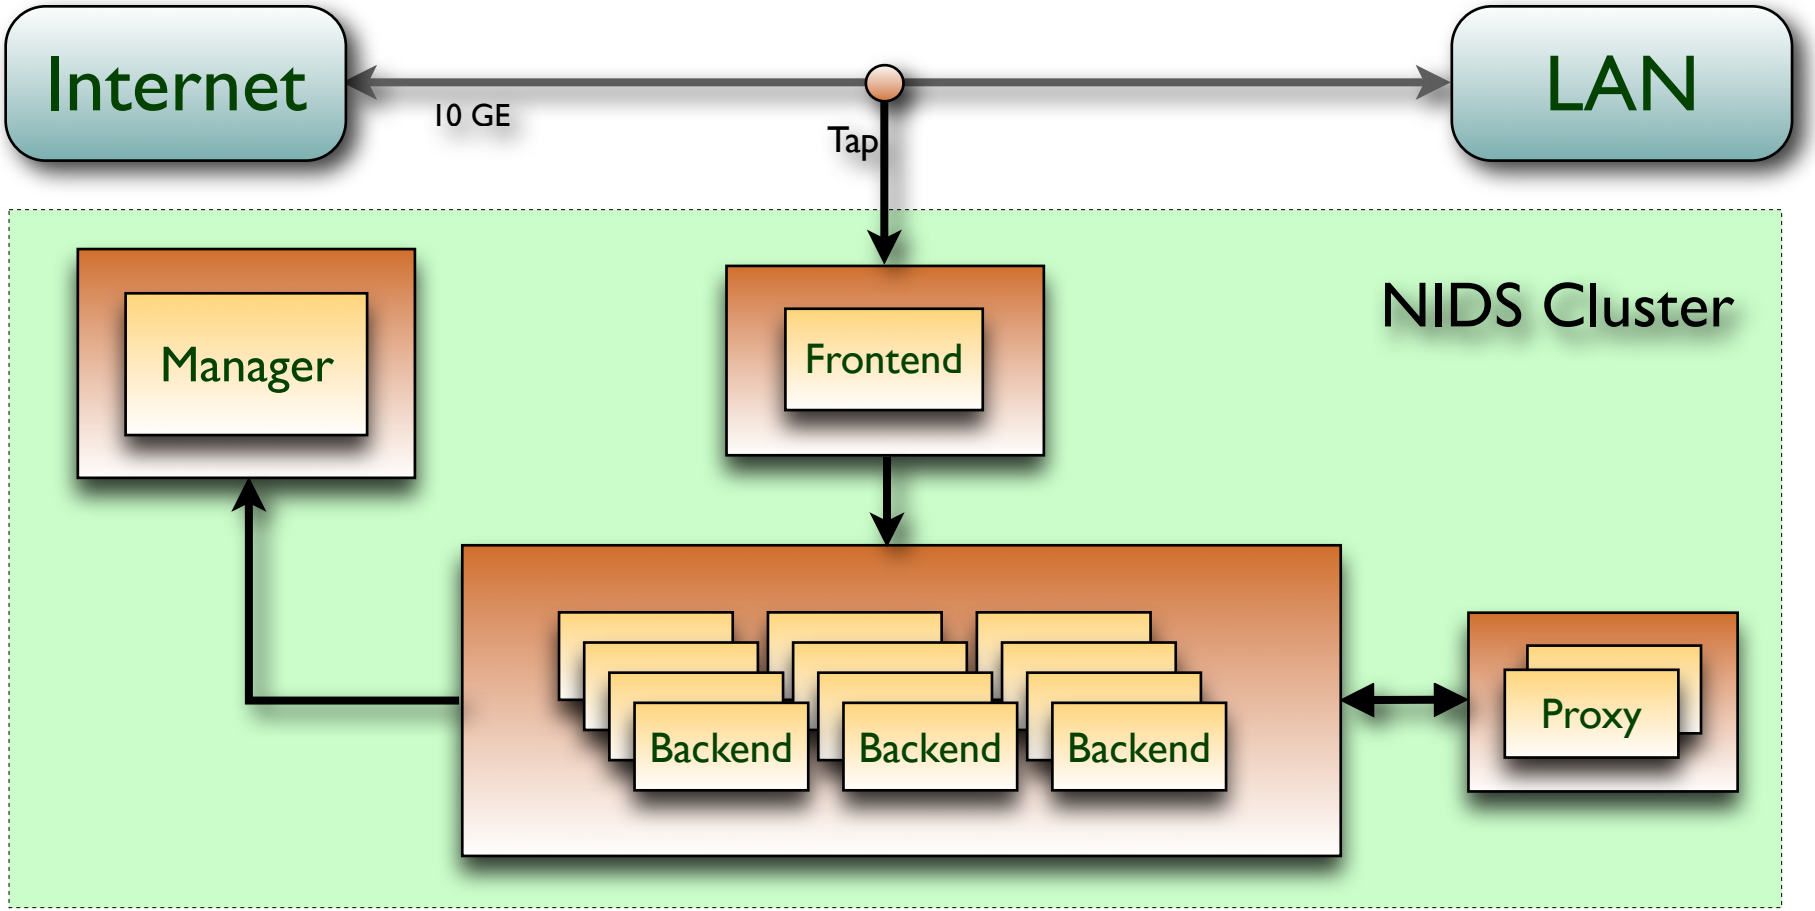

ITRoadmap
CONFERENCE & EXPO
BERKELEY LAB

Initial Prototype Cluster Setups

- Lawrence Berkeley National Laboratory
  - 10 Gbps upstream link; 1 frontend, 10 backends
- University of California, Berkeley
  - 2x1Gbps upstream links; 2 frontends, 6 backends (for only parts of the traffic)
- IEEE Supercomputing 2006
  - Conference's 1 Gbps backbone network
  - 10 Gbps High Speed Bandwidth Challenge network
- Goal: Replace current operational security monitoring at LBNL
  - We also plan to build a high-performance 10G research cluster at UCB

ITRoadmap CONFERENCE & EXPO

BERKELEY LAB

Cluster Components

- **Backends**
  - **Running Bro as their analysis engine**
  - **Using essentially the same configuration as before, just on a slice of traffic**
  - **Bro provides extensive communication facilities for sharing of low-level state**
    - Just mark an analysis variable as *synchronized* and its value will be propagated
- **Frontend**
  - **Distributes traffic across backends by rewriting MAC addresses**
    - Switch dispatches forwarded packets across the backends
  - **Initial prototype implementations**
    - Software based on open-source Click modular router platform (up to 2Gbps)
    - Customized appliance implementing rewriting in hardware

ITRoadmap CONFERENCE & EXPO

BERKELEY LAB

cFlow: A High-Performance Load-balancing Solution

- For a production front-end, LBNL worked with *cPacket Networks*
- cFlow: 10GE line-rate, stand-alone load-balancer

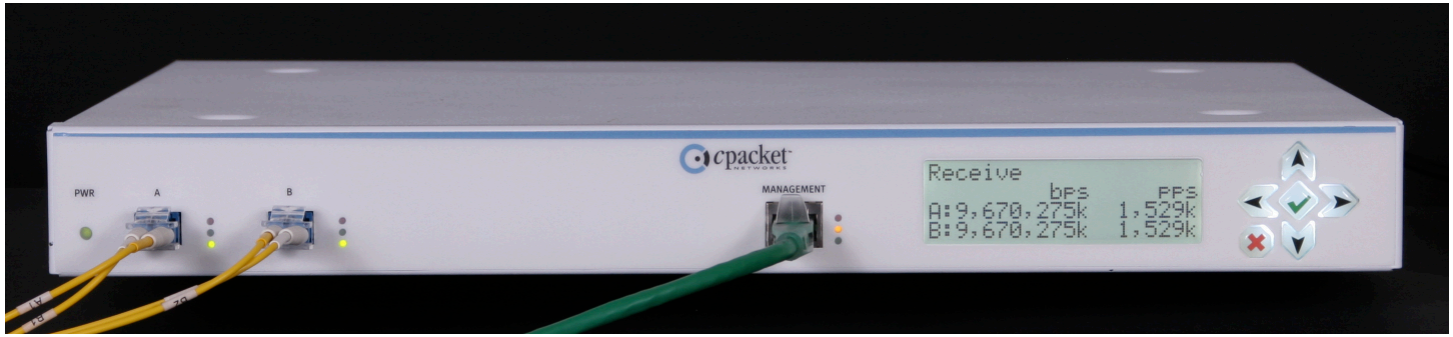

- 10GE in/out
- Web & CLI
- Filtering capabilities

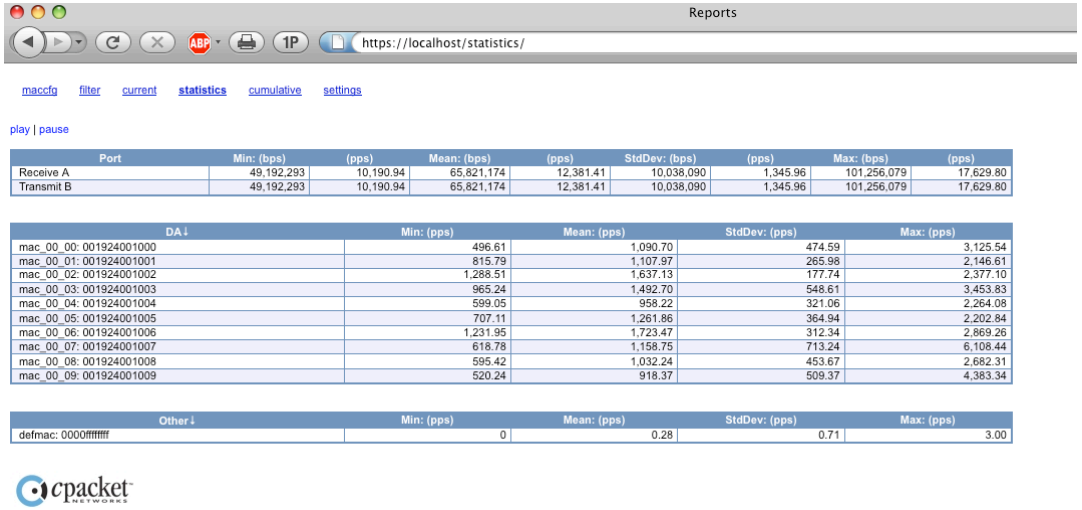

Port	Min: (bps)	(pps)	Mean: (bps)	(pps)	StdDev: (bps)	(pps)	Max: (bps)	(pps)
Receive A	49,192,293	10,190.94	65,821,174	12,381.41	10,038,090	1,345.96	101,256,079	17,629.80
Transmit B	49,192,293	10,190.94	65,821,174	12,381.41	10,038,090	1,345.96	101,256,079	17,629.80

DA↓	Min: (pps)	Mean: (pps)	StdDev: (pps)	Max: (pps)
mac_00_00: 001924001000	496.61	1,090.70	474.59	3,125.54
mac_00_01: 001924001001	815.79	1,107.97	265.98	2,146.61
mac_00_02: 001924001002	1,288.51	1,637.13	177.74	2,377.10
mac_00_03: 001924001003	965.24	1,492.70	548.61	3,453.83
mac_00_04: 001924001004	599.05	958.22	321.06	2,264.08
mac_00_05: 001924001005	707.11	1,261.86	364.94	2,202.84
mac_00_06: 001924001006	1,231.95	1,723.47	312.34	2,869.26
mac_00_07: 001924001007	618.78	1,158.75	713.24	6,108.44
mac_00_08: 001924001008	595.42	1,032.24	453.67	2,682.31
mac_00_09: 001924001009	520.24	918.37	509.37	4,383.34

Other↓	Min: (pps)	Mean: (pps)	StdDev: (pps)	Max: (pps)
defmac: 0000ffffffff	0	0.28	0.71	3.00

cpacket NETWORKS

ITRoadmap CONFERENCE & EXPO

Cluster Manager

- Interactive interface for installation, configuration, tuning, logging, ...

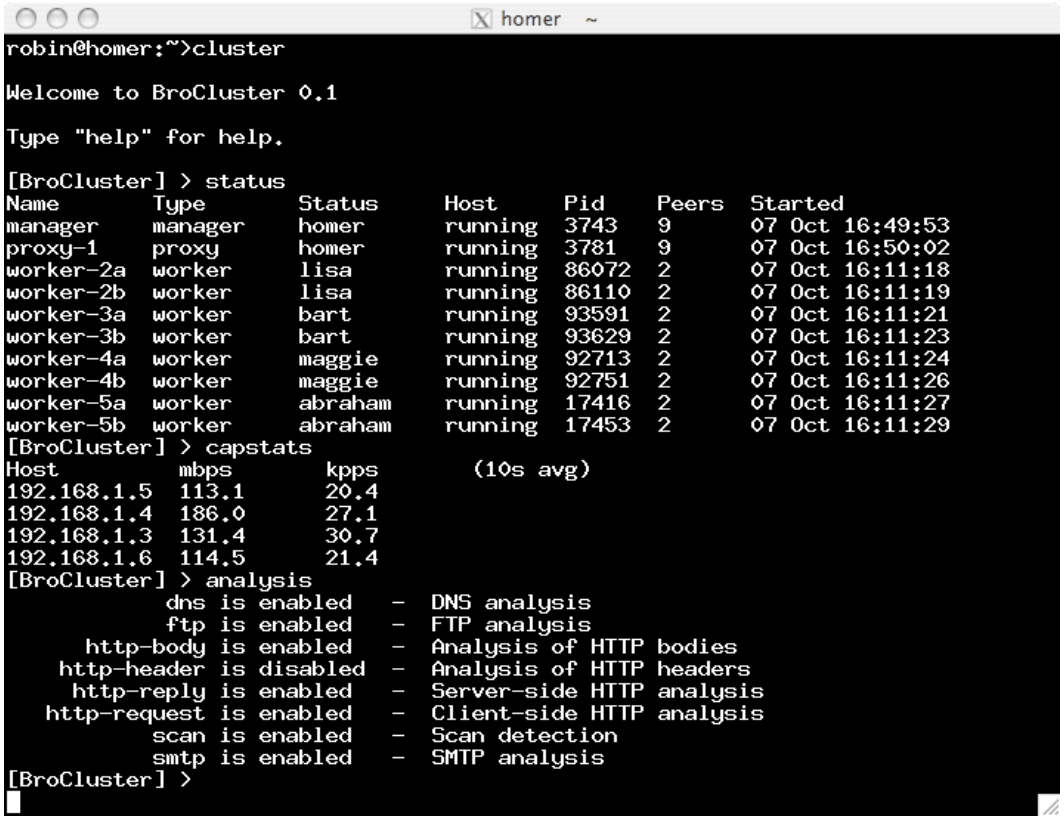

ITRoadmap CONFERENCE & EXPO

BERKELEY LAB

Summary

- LBNL faces rather unique security challenges
- Lab takes comprehensive approach to network security
  - Picking components which work best for their particular scope
- Bro NIDS has been central part of the Lab's security for many years
  - Highly flexible, *activity*-based system (rather than signatures)
  - Open-source, actively maintained, and runs on commodity hardware
- Bro Cluster provides the Lab with head-room for the years to come
- Future Plans for Bro
  - Developing a highly concurrent analysis model to exploit multi-core potential
  - In-depth application analysis

ITRoadmap CONFERENCE & EXPO

BERKELEY LAB

Thank You

Robin Sommer
International Computer Science Institute, &
Lawrence Berkeley National Lab