Network Security Today:
Finding Complex Attacks at 100Gb/s

Robin Sommer

International Computer Science Institute, &
Lawrence Berkeley National Laboratory

robin@icsi.berkeley.edu
http://www.icir.org/robin

TU München, September 2012

Outline

Outline

Today's Threats.

Deep Packet Inspection at High Speed.

Collective Intelligence.

The Old Days ...

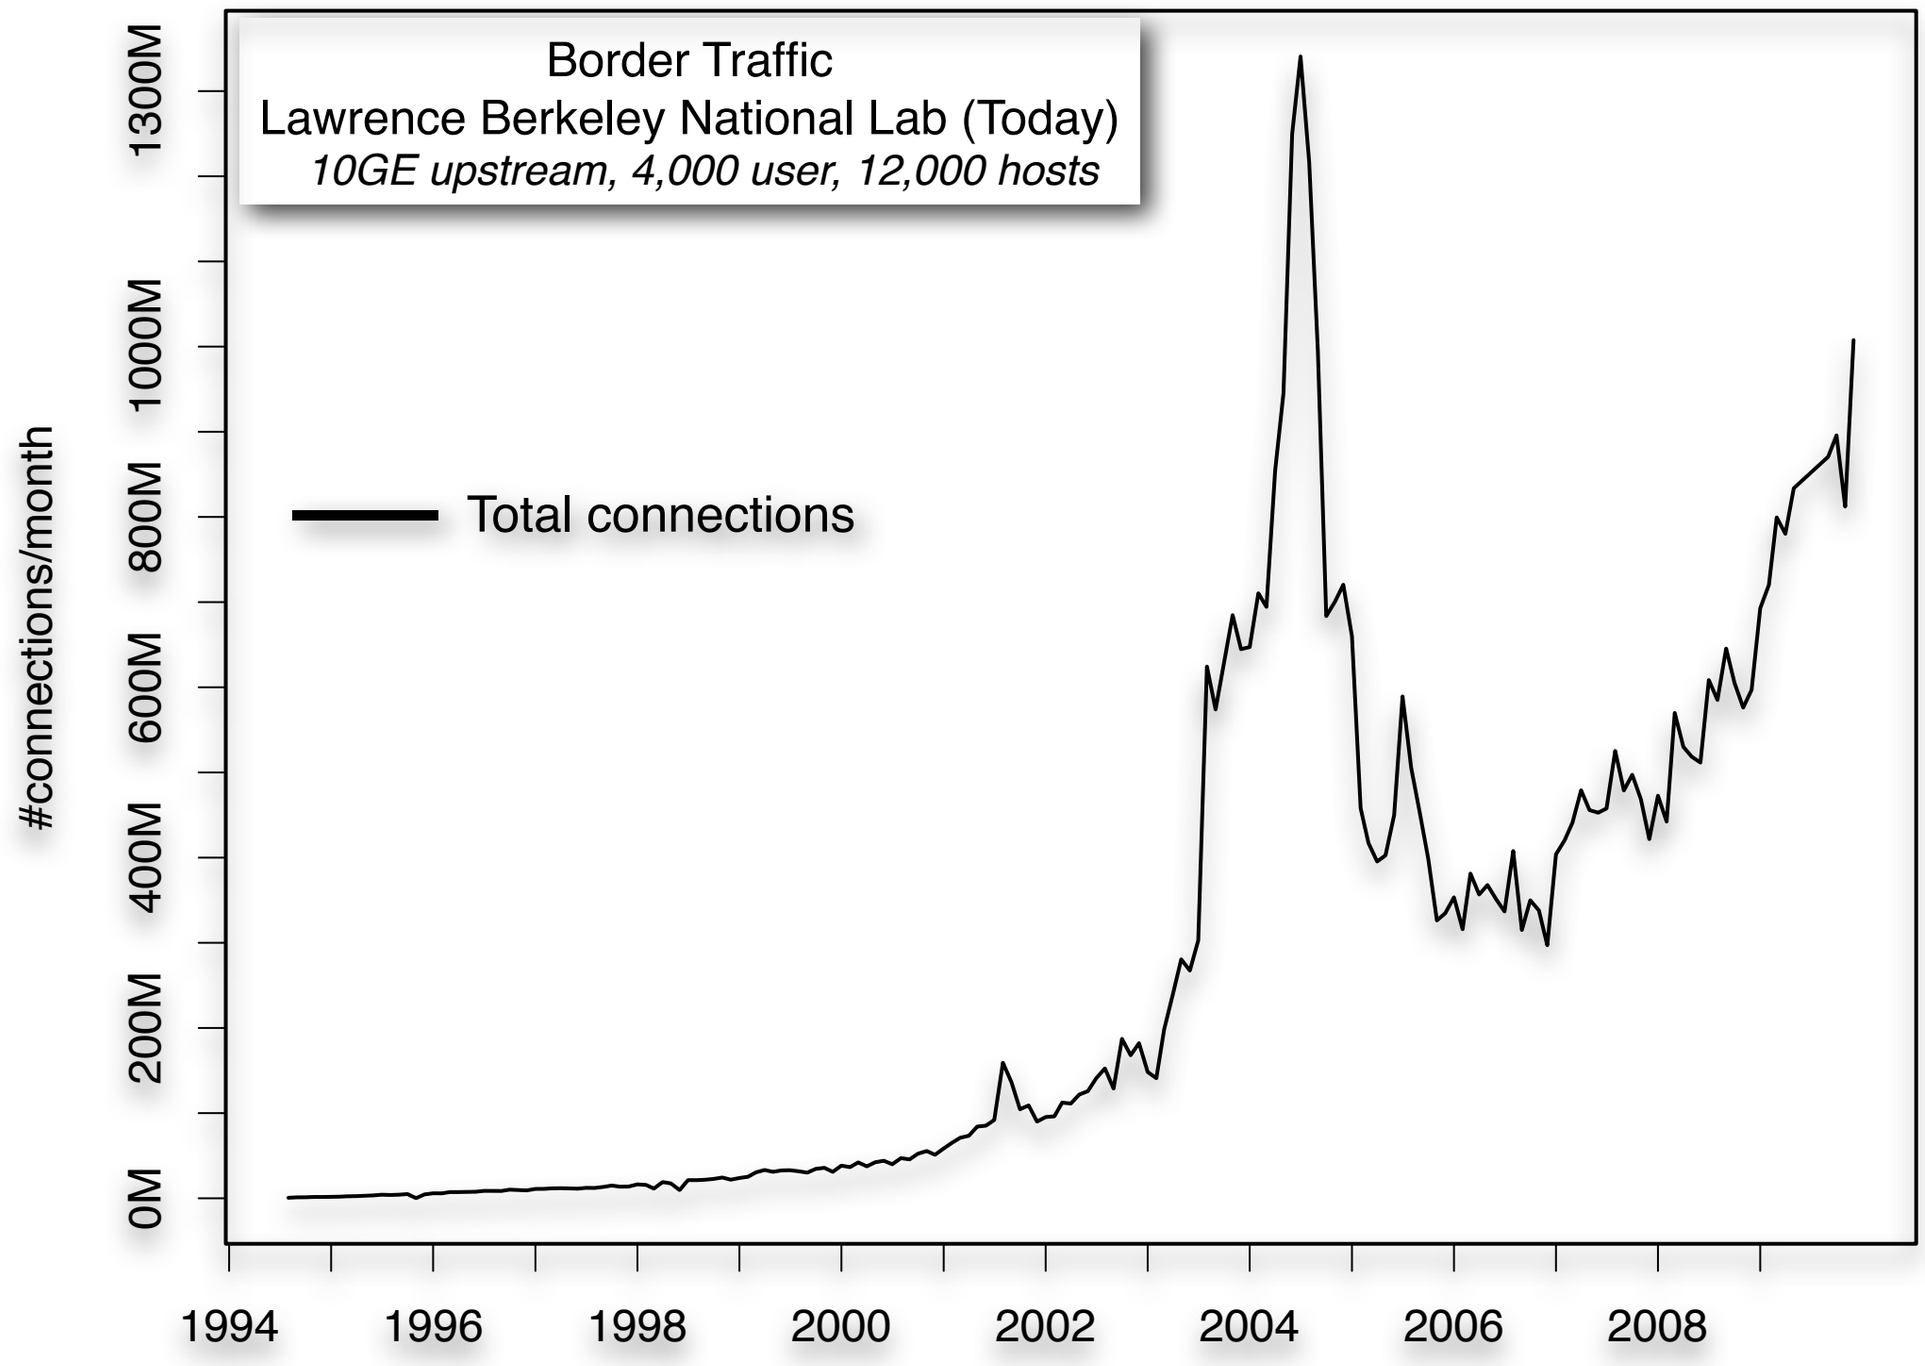

Data: Lawrence Berkeley National Lab

The Old Days ...

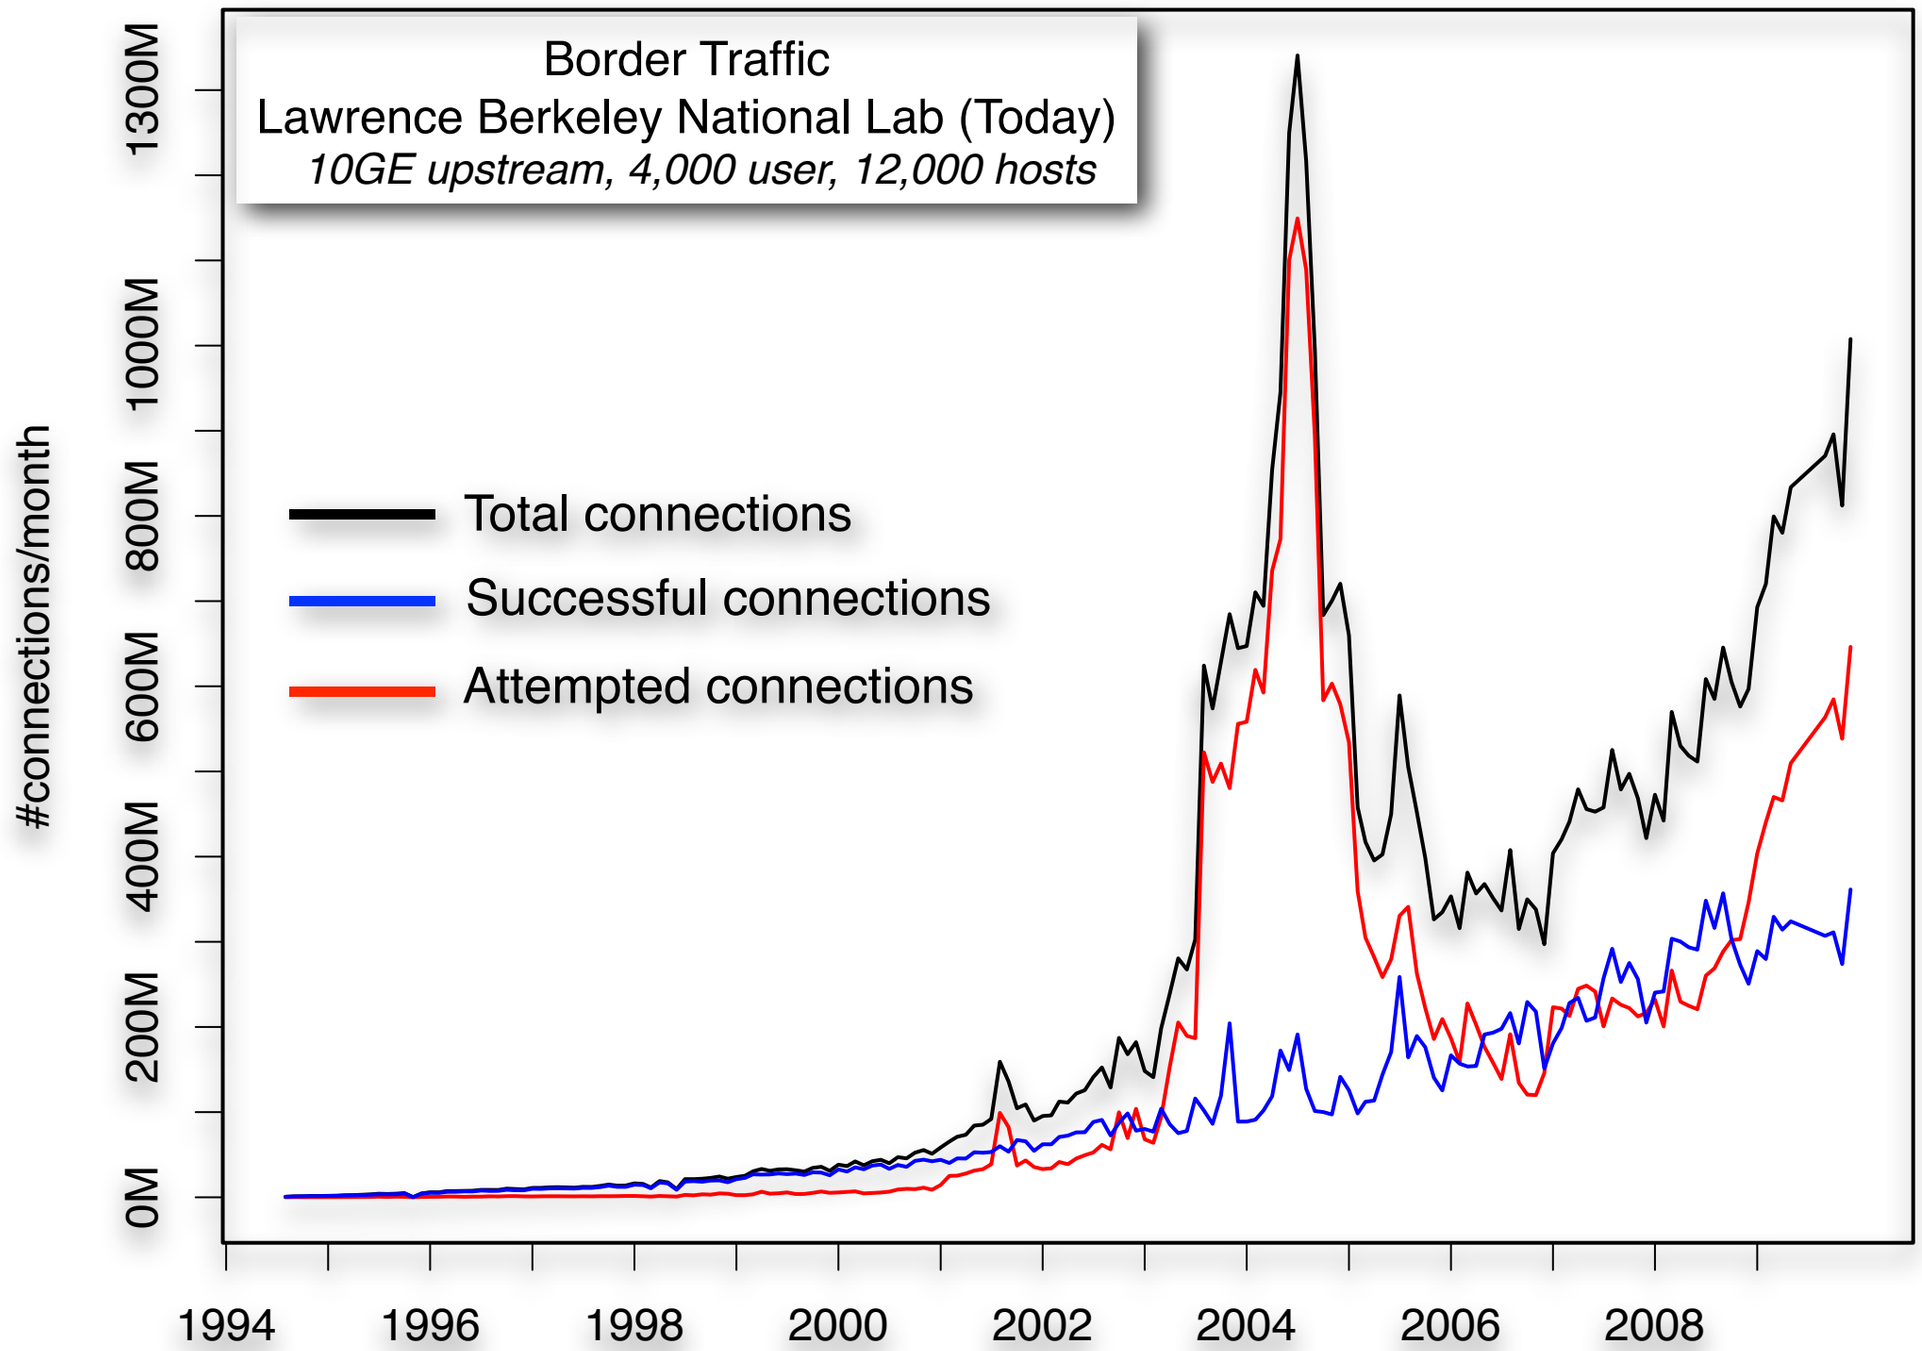

Data: Lawrence Berkeley National Lab

The Old Days ...

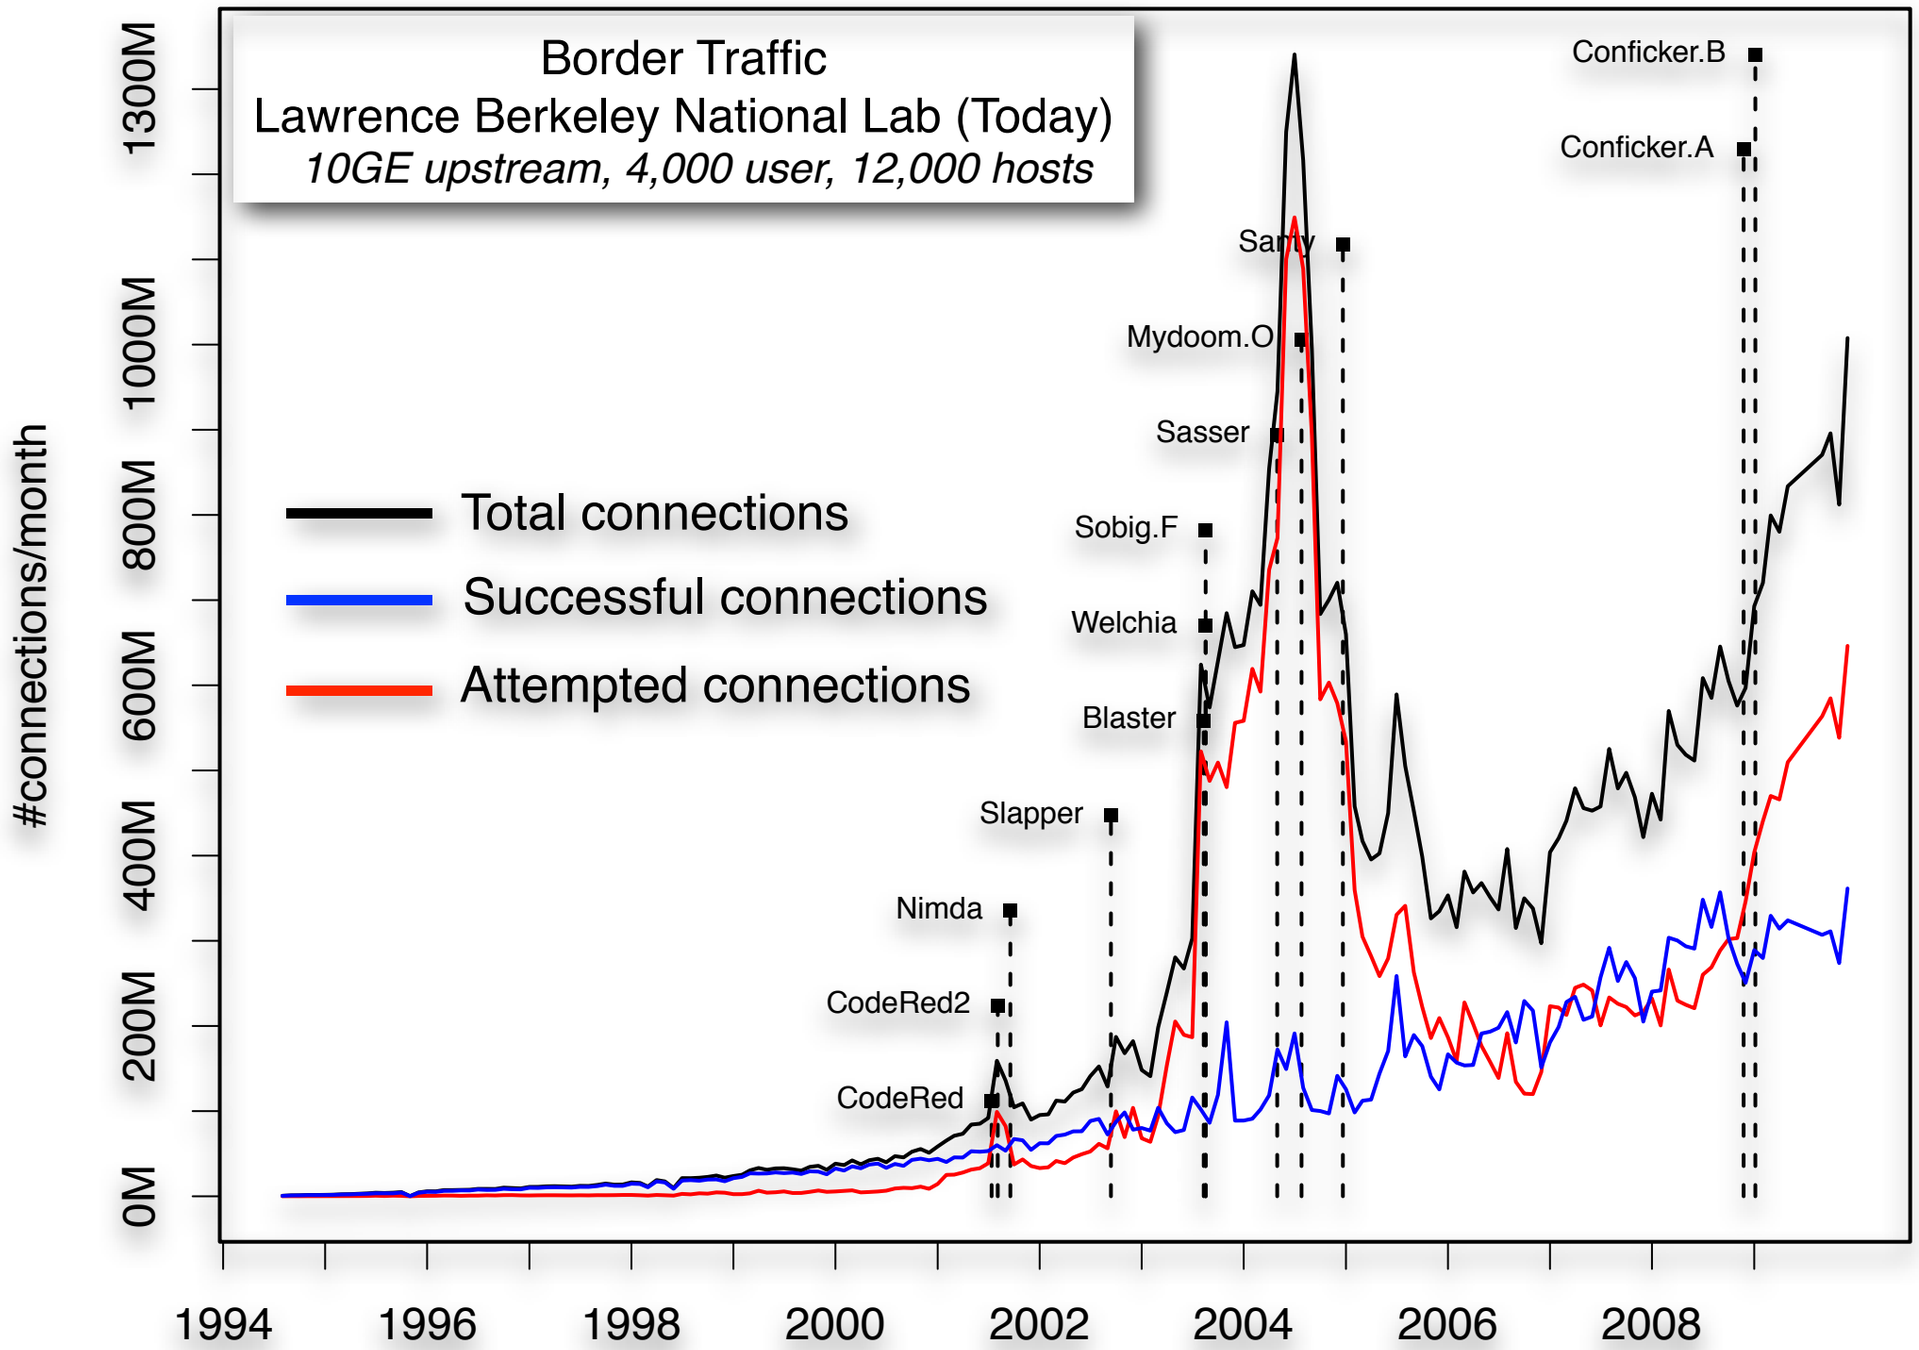

Data: Lawrence Berkeley National Lab

Trend 1: Commercialization of Attacks

Trend 1: Commercialization of Attacks

Attacks aimed at making a profit.

- Selling (illegal) goods and services.
- Exfiltrating information.

Thriving underground economy.

- Empowered by virtually endless supply of “bots”.
- Everything is on sale (“crime-as-a-service”).

“Pay Per Install” Services

“Pay Per Install” Services

Crime Economics

Crime Economics

Accelerated arms race.

Innovative, fast moving attackers.

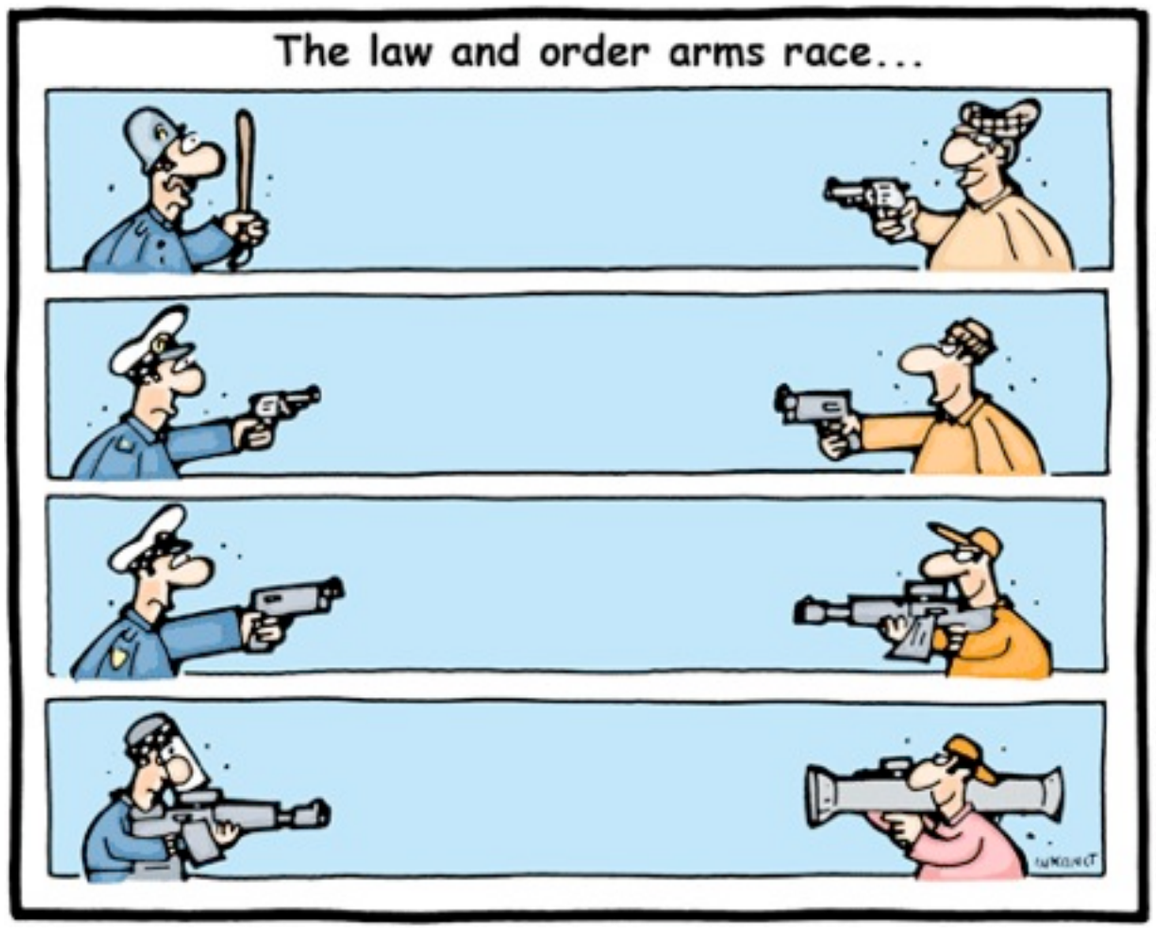

Crime Economics

Accelerated arms race.
Innovative, fast moving attackers.

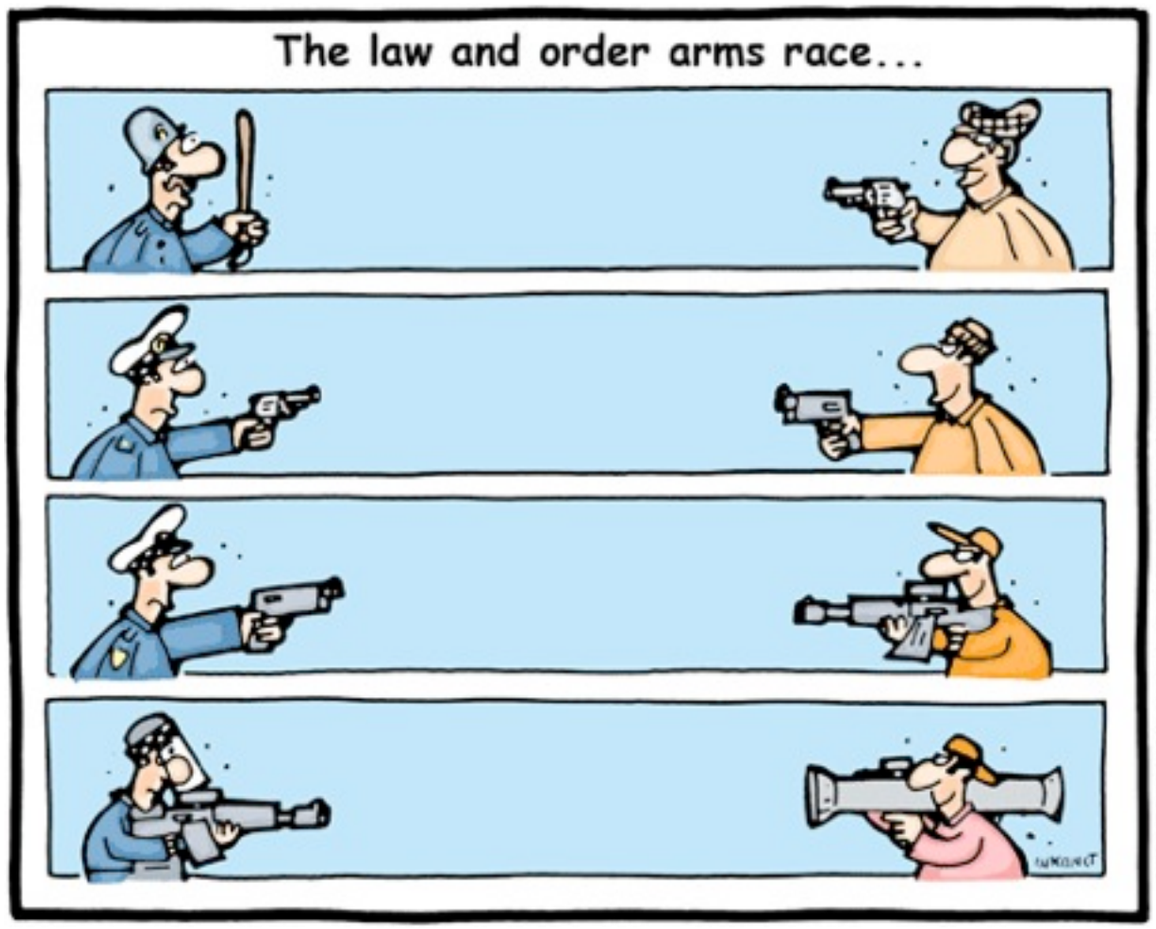

Bear race.
If attack pays, it's good enough.

Trend 2: Highly Targeted Attacks

Trend 2: Highly Targeted Attacks

High-skill / high-resource attacks.

Targeting **you.**

Extremely hard to defend against.

Trend 2: Highly Targeted Attacks

High-skill / high-resource attacks.

- Targeting **you.**
- Extremely hard to defend against.

Typical Instances

- Activist hacking.
- “Advanced Persistent Threats”.

Trend 2: Highly Targeted Attacks

High-skill / high-resource attacks.

Targeting **you.**

Extremely hard to defend against.

Typical Instances

Advanced Persistent Threat (APT). MANDIANT defines the APT as a group of sophisticated, determined and coordinated attackers that have been systematically compromising U.S. government and commercial computer networks for years. The vast majority of

Source: MANDIANT

Targeted Attacks: APTs

EXPLOITATION LIFE CYCLE

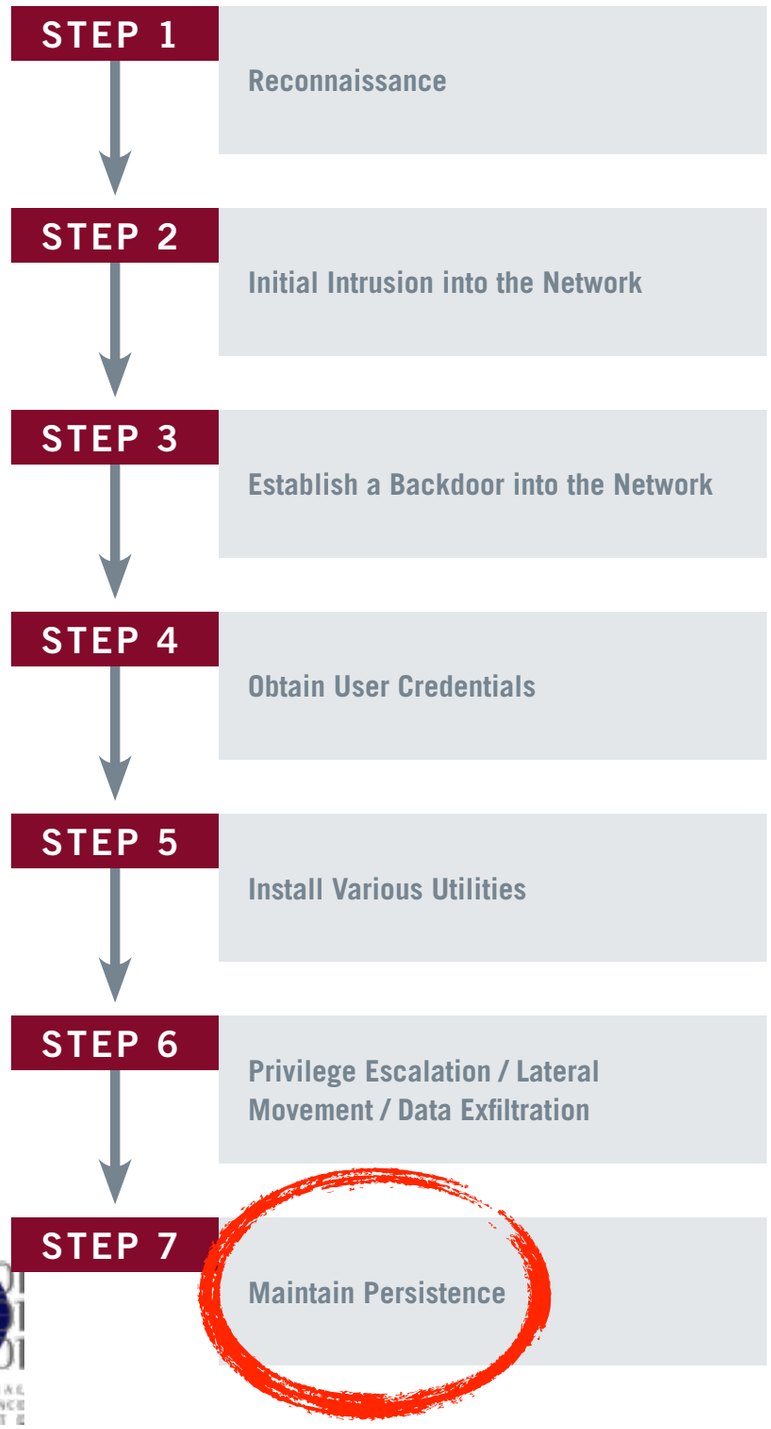

Source: MANDIANT

Targeted Attacks: APTs

EXPLOITATION LIFE CYCLE

- **STEP 1** — Reconnaissance
- **STEP 2** — Initial Intrusion into the Network
- **STEP 3** — Establish a Backdoor into the Network
- **STEP 4** — Obtain User Credentials
- **STEP 5** — Install Various Utilities
- **STEP 6** — Privilege Escalation / Lateral Movement / Data Exfiltration
- **STEP 7** — Maintain Persistence

APT MALWARE COMMUNICATION

100% of APT backdoors made only outbound connections

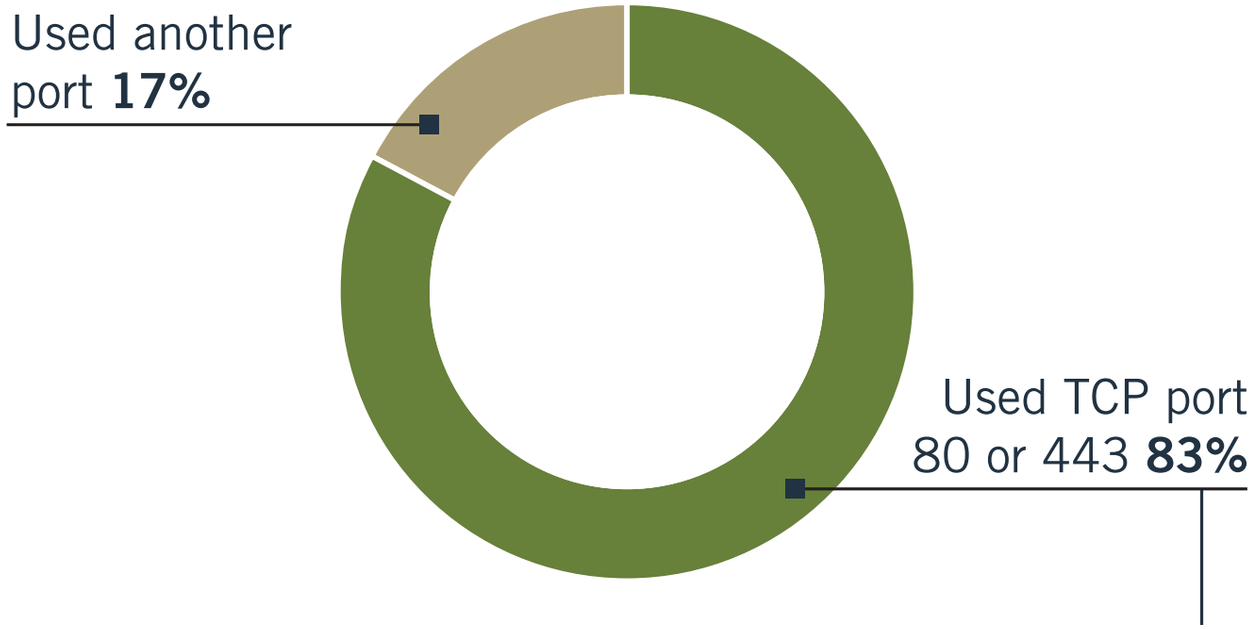

In no instance was any APT malware written or configured to listen for inbound connections.

Source: MANDIANT

Challenges for Defenders

Challenges for Defenders

Varying threat models.

No ring rules them all.

Challenges for Defenders

Varying threat models.

No ring rules them all.

Volume and variability.

Network traffic is an enormous haystack.

Challenges for Defenders

Varying threat models.

No ring rules them all.

Volume and variability.

Network traffic is an enormous haystack.

Semantic complexity.

The action is really at the application-layer.

Analyzing Semantics

Analyzing Semantics

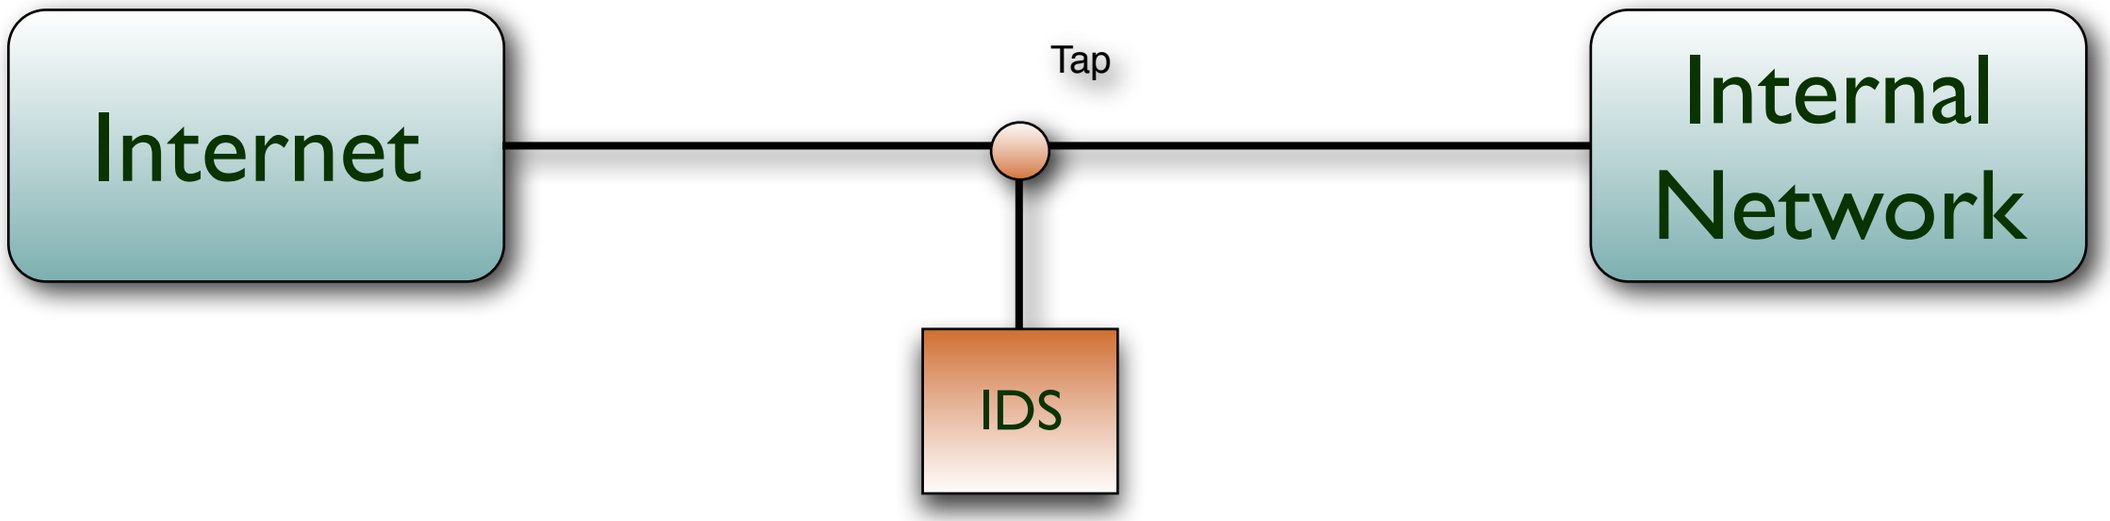

Analyzing Semantics

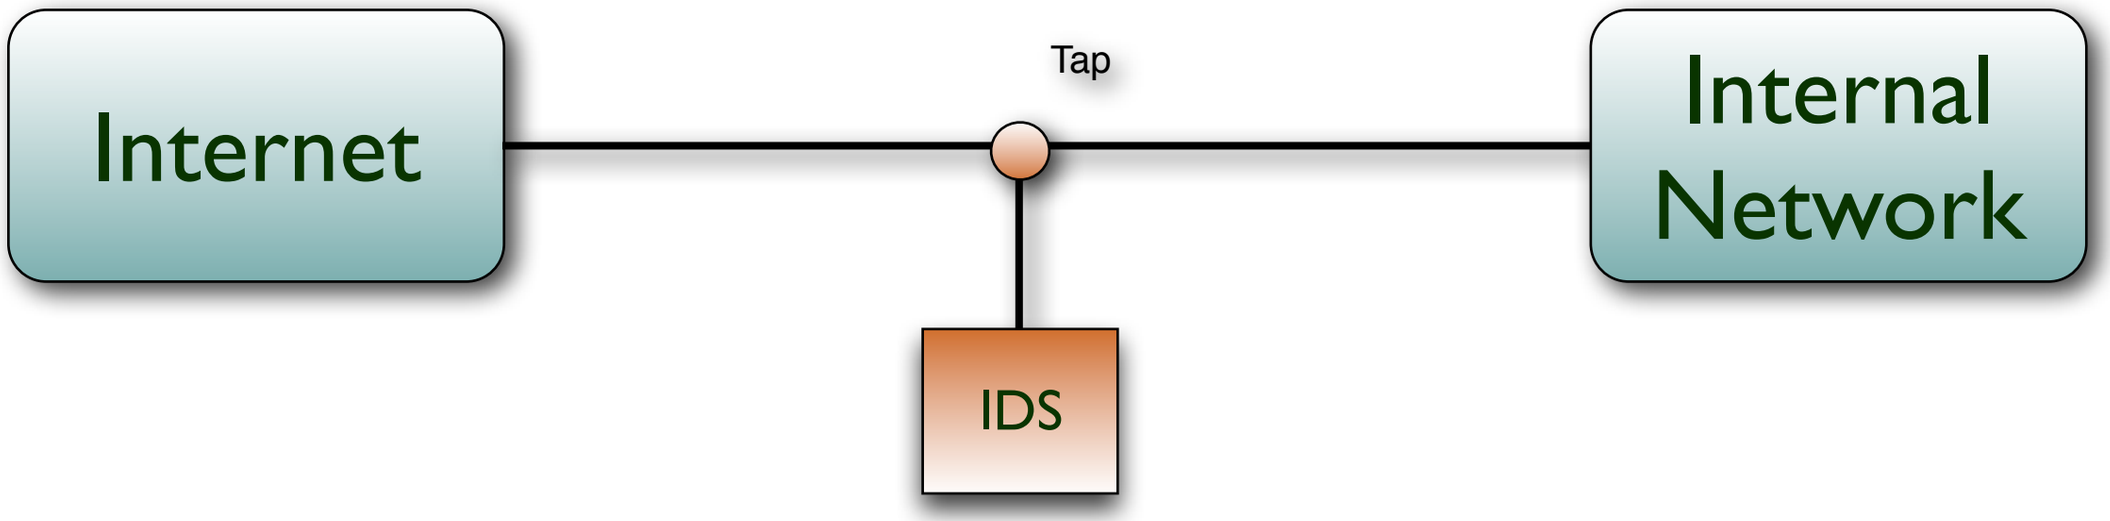

Example: Finding downloads of known malware.

1. Find and parse all Web traffic.
2. Find and extract binaries.
3. Compute hash and compare with database.
4. Report, and potentially kill, if found.

Deep Packet Inspection at High Speed

Back in 2005 ...

Back in 2005 ...

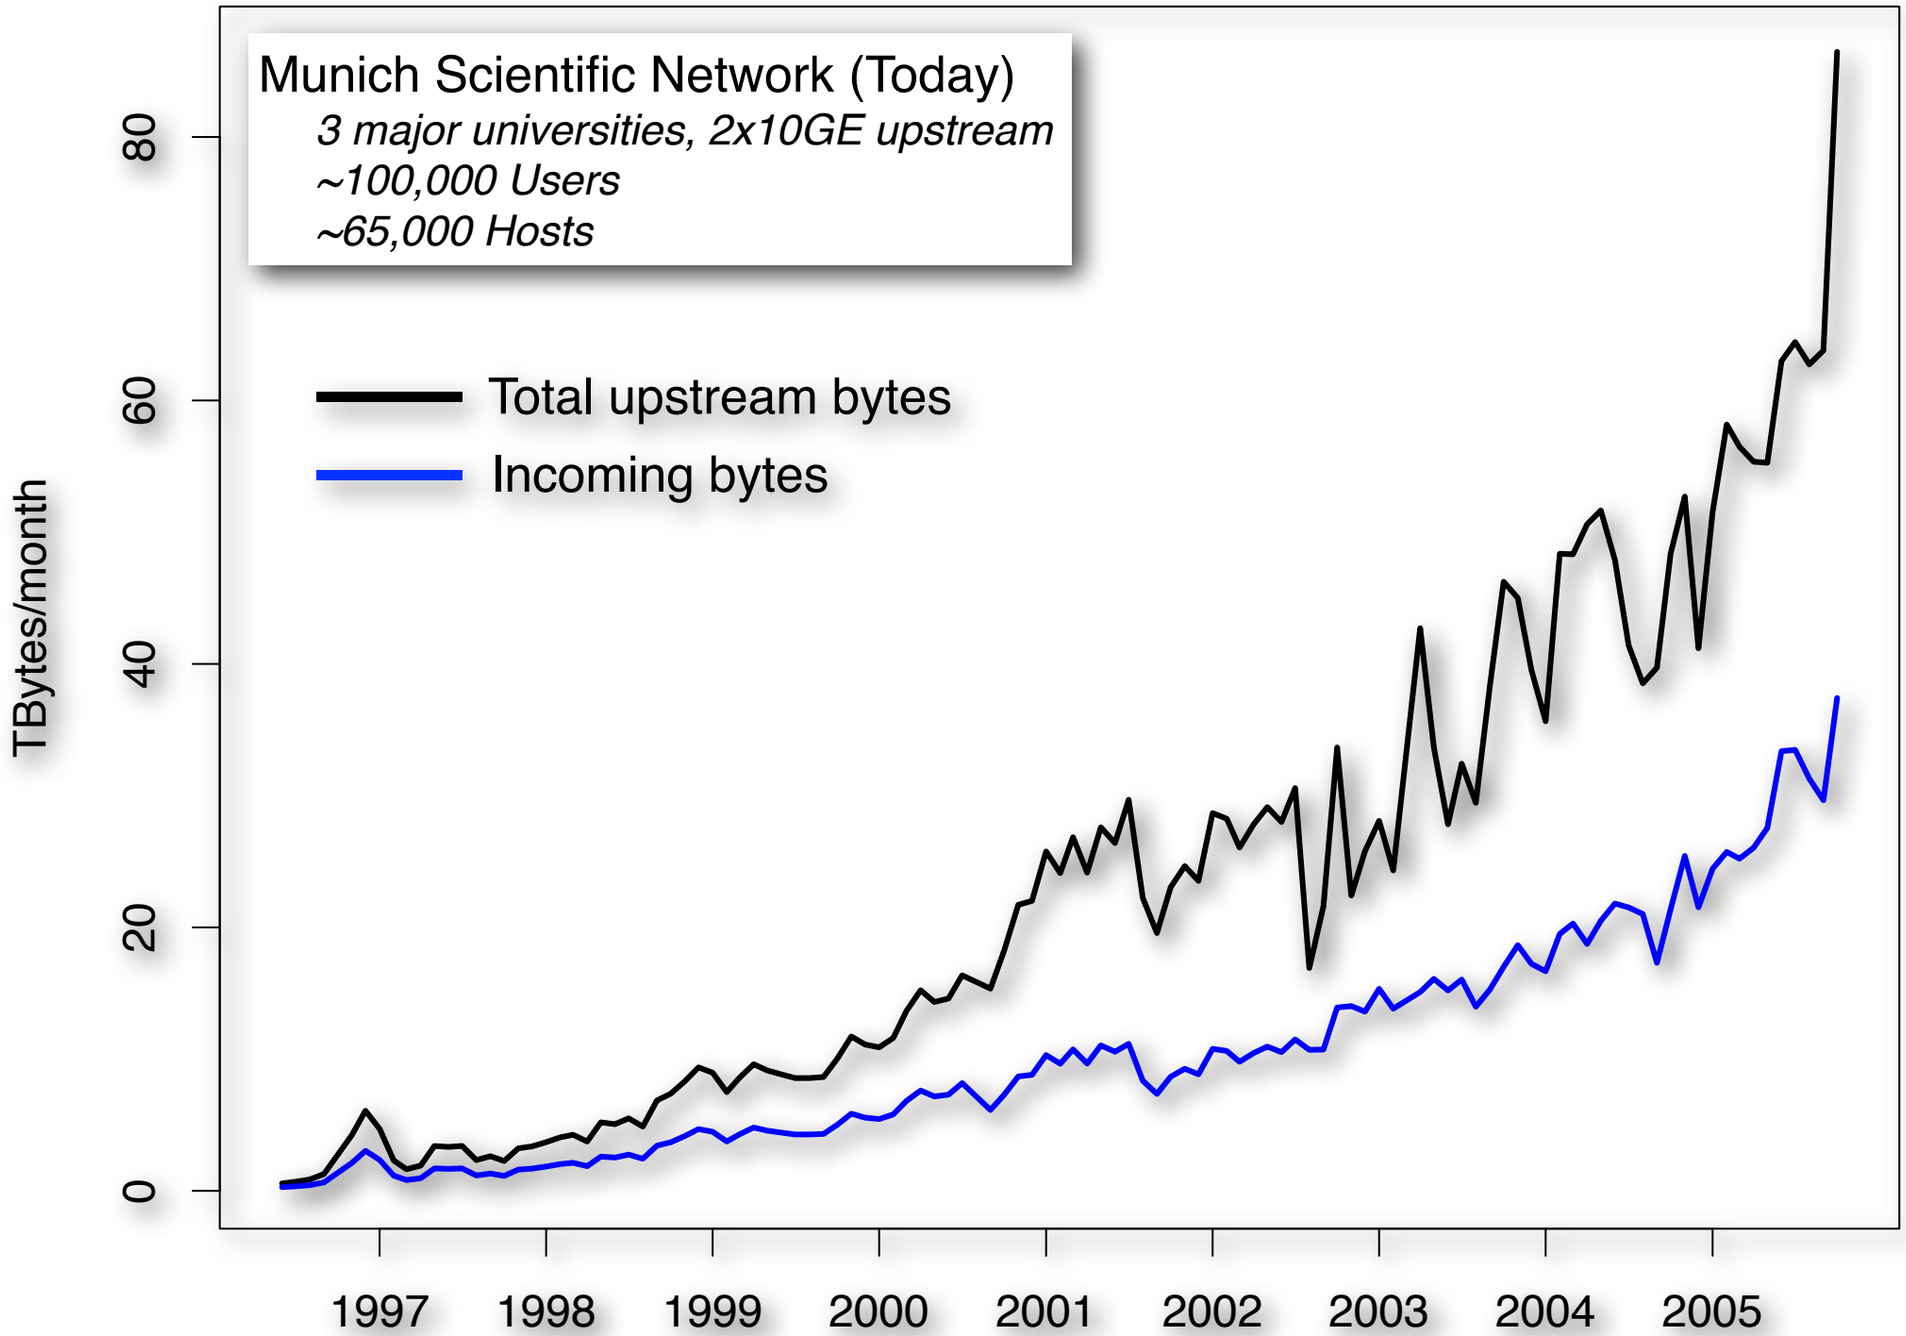

Data: Leibniz-Rechenzentrum, München

Today ...

Today ...

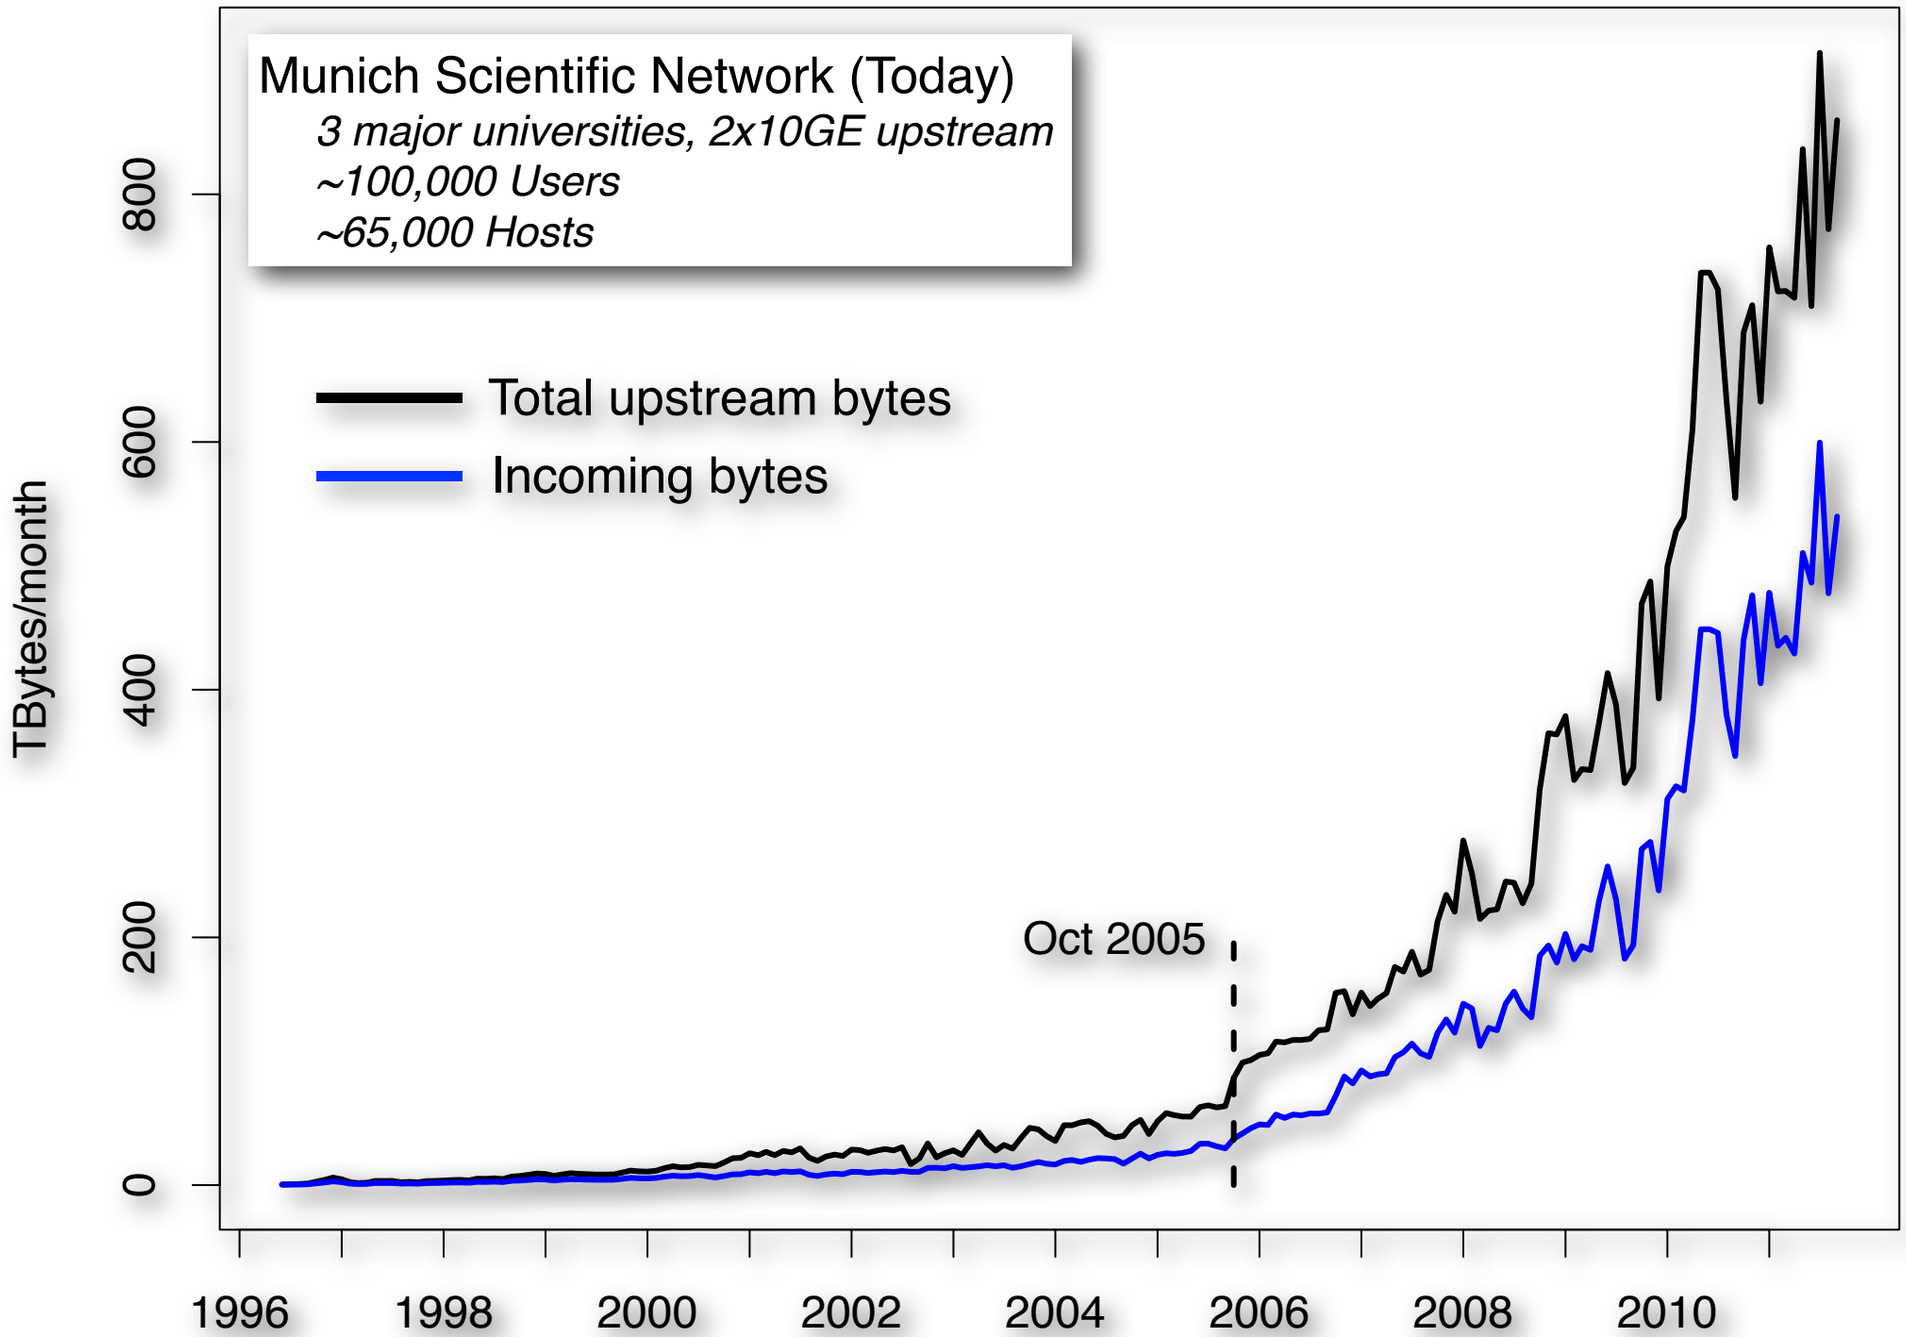

Data: Leibniz-Rechenzentrum, München

Traditional Gap: Research vs. Operations

Traditional Gap: Research vs. Operations

Conceptually simple tasks can be hard in practice.

- Academic research often neglects operational constraints.
- Operations cannot leverage academic results.

Traditional Gap: Research vs. Operations

Conceptually simple tasks can be hard in practice.

- Academic research often neglects operational constraints.
- Operations cannot leverage academic results.

We focus on working *with* operations.

- Close collaborations with several large sites.
- Extremely fruitful for both sides.

Research Platform: Bro

Research Platform: Bro

Originally developed by Vern Paxson in 1996.

Open-source, BSD-license, maintained at ICSI.

In operational use since the beginning.

Conceptually very different from other IDS.

http://www.bro-ids.org

Bro Script Example: Matching URLs

Task: Report all Web requests for files called "`passwd`".

Bro Script Example: Matching URLs

Task: Report all Web requests for files called "`passwd`".

```
event http_request(c: connection,           # Connection.
                   method: string,          # HTTP method.
                   original_URI: string,    # Requested URL.
                   unescaped_URI: string,   # Decoded URL.
                   version: string)         # HTTP version.
{
    if ( method == "GET" && unescaped_URI == /.*passwd/ )
        NOTICE(...); # Alarm.
}
```

"Who's Using It?"

Installations across the US

Universities
Research Labs
Supercomputer Centers
Industry

Examples

Lawrence Berkeley National Lab
Indiana University
National Center for Supercomputing Applications
National Center for Atmospheric Research

... and many more sites

Fully integrated into *Security Onion*

Popular security-oriented Linux distribution

Recent User Meetings

Bro Workshop 2011 at NCSA
Bro Exchange 2012 at NCAR

Each attended by about 50 operators from from 30-35 organizations

Bro History

1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012

Vern writes 1st line of code

Bro History

1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012

Vern writes 1st line of code

LBNL starts using Bro operationally

v0.2
1st CHANGES entry

v0.4
HTTP analysis
Scan detector
IP fragments
Linux support

v0.6
RegExps
Login analysis

v0.7a48
Consistent CHANGES

v0.7a90
Profiling
State Mgmt

v0.7a175/0.8aX
Signatures
SMTP
IPv6 support
User manual

0.8a37
Communication
Persistence
Namespaces
Log Rotation

v0.8aX/0.9aX
SSL/SMB
STABLE releases
BroLite

v1.0
BinPAC
IRC/RPC analyzers
64-bit support
Sane version numbers

v1.1/v1.2
when Stmt
Resource tuning
Broccoli
DPD

v1.3
Ctor expressions
GeoIP
Conn Compressor

v1.4
DHCP/BitTorrent
HTTP entities
NetFlow
Bro Lite Deprecated

v1.5
BroControl

Bro SDCI

Bro 2.0
New Scripts

Bro 2.1
IPv6
Input Framework

Bro History

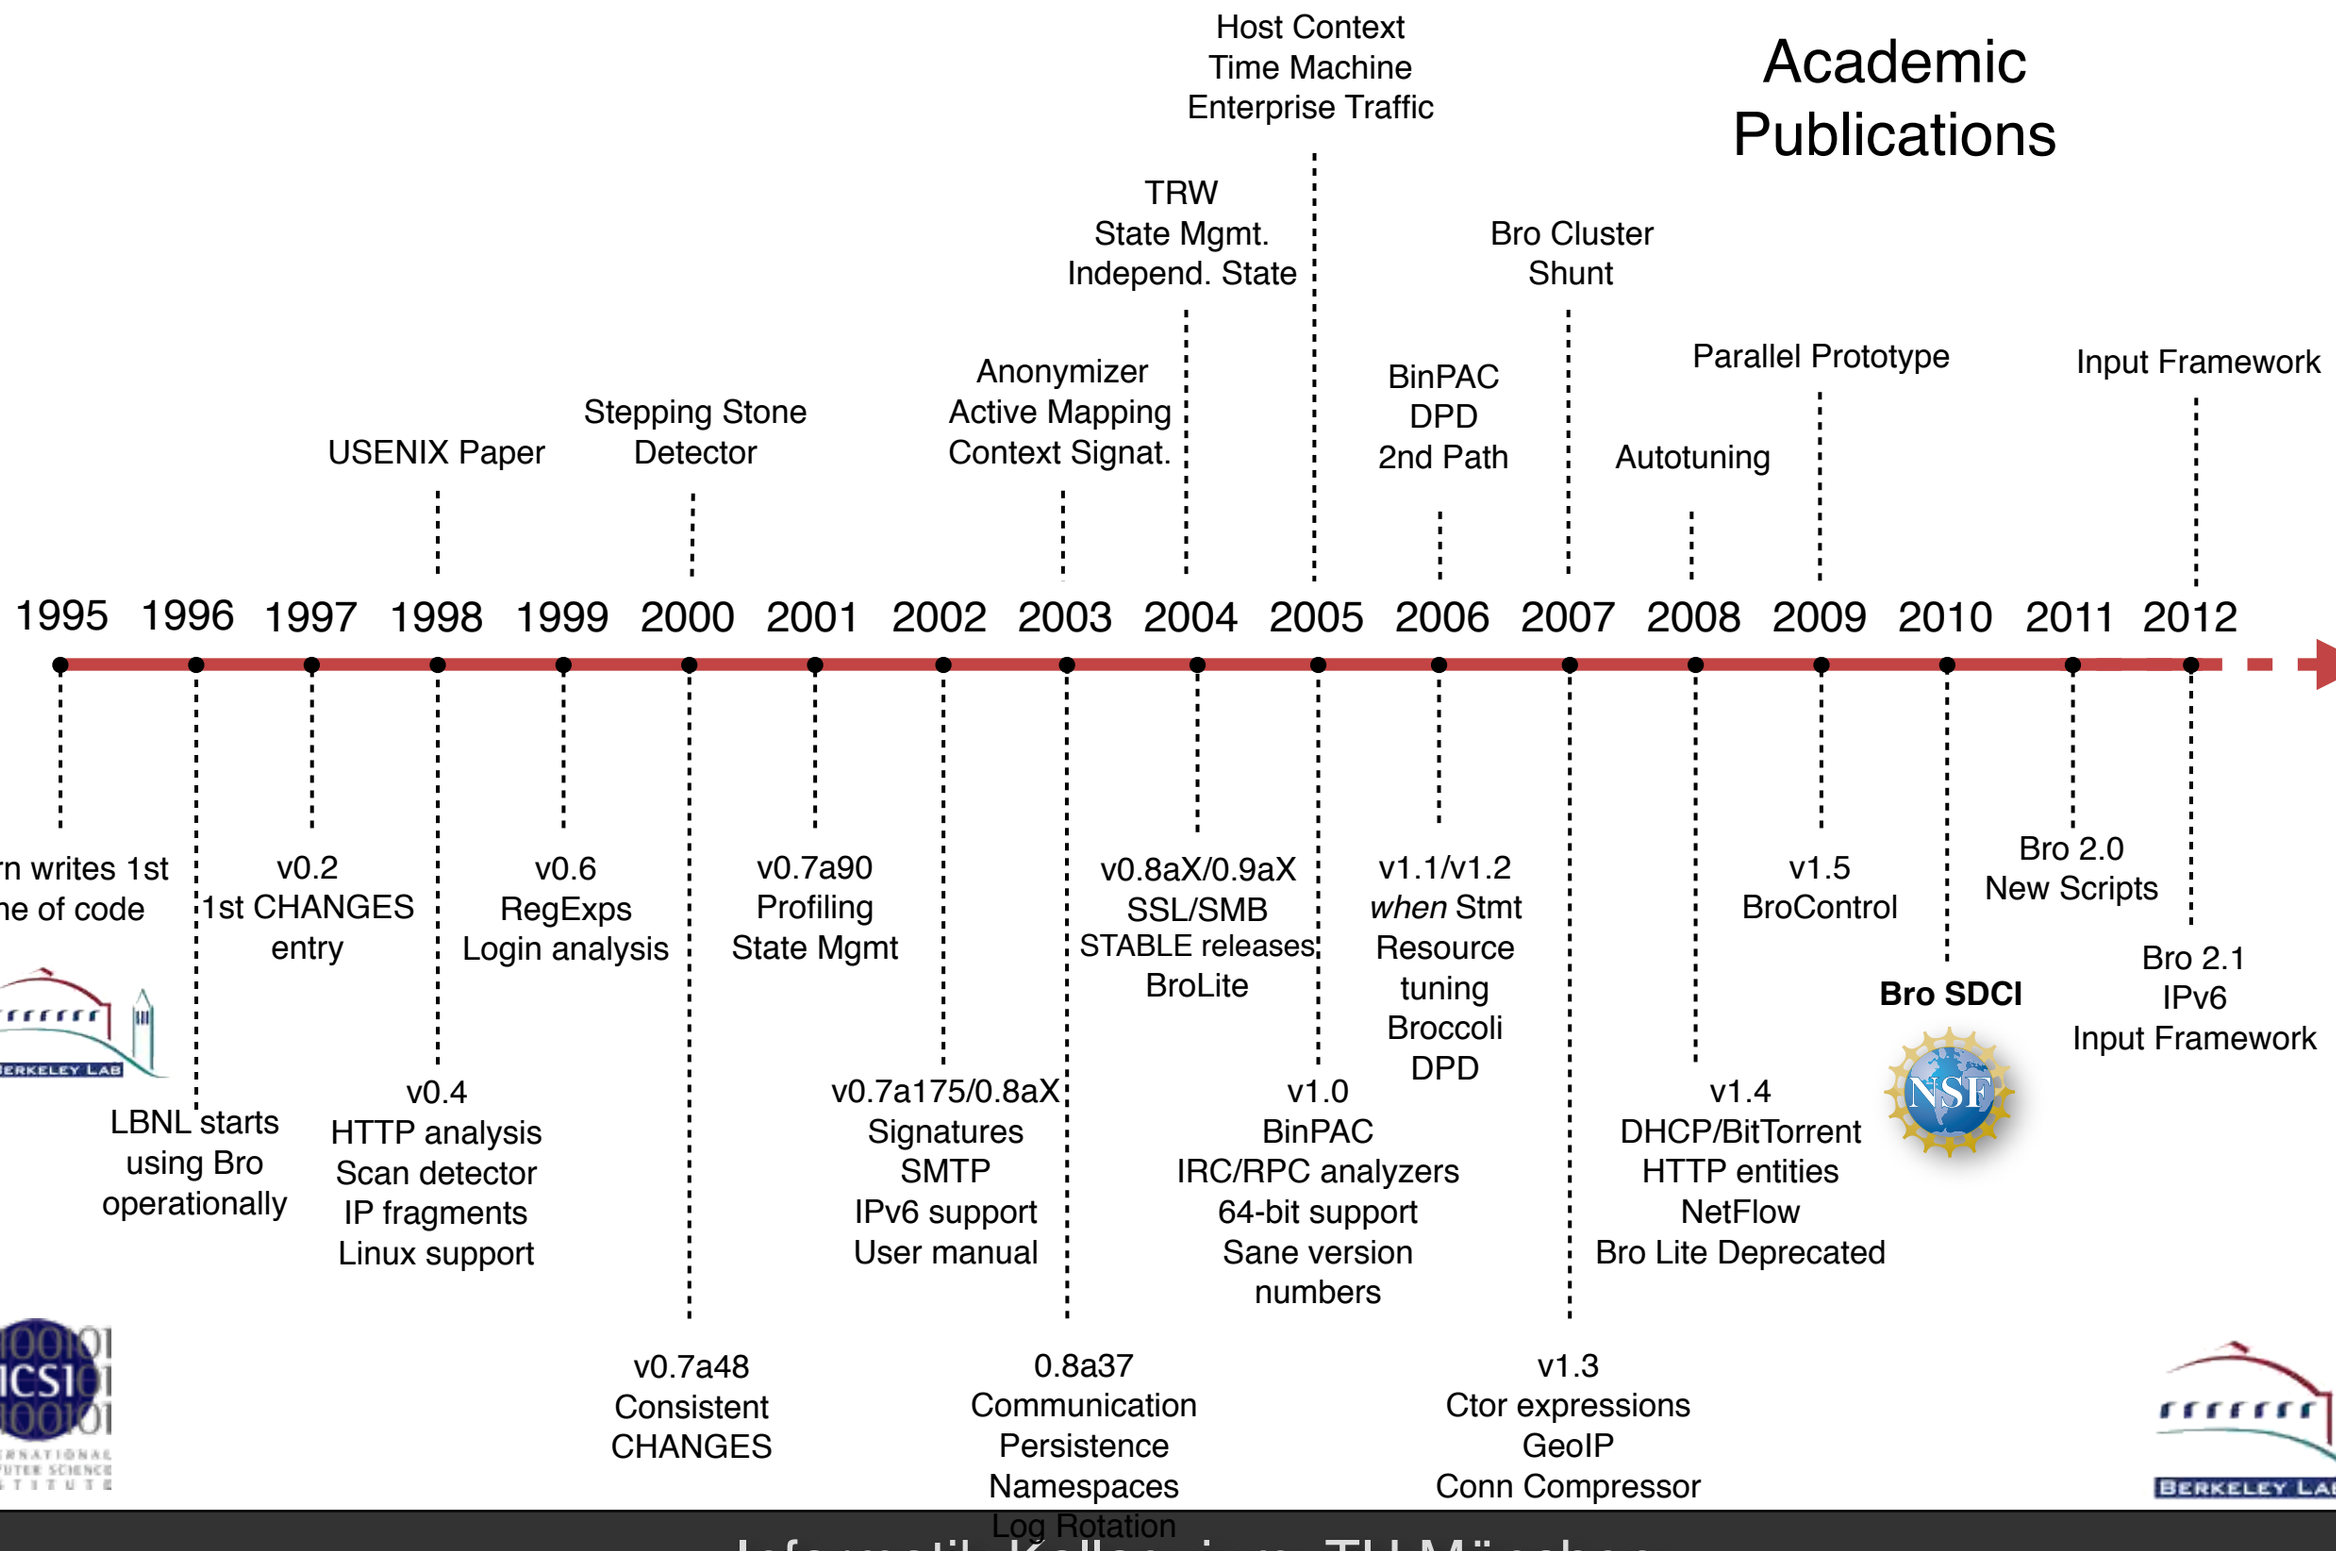
Academic Publications
Host Context
Time Machine
Enterprise Traffic
TRW
State Mgmt.
Independ. State
Bro Cluster
Shunt
Anonymizer
Active Mapping
Context Signat.
BinPAC
DPD
2nd Path
Parallel Prototype
Autotuning
Input Framework
USENIX Paper
Stepping Stone
Detector
1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012
Vern writes 1st line of code
v0.2
1st CHANGES entry
v0.6
RegExps
Login analysis
v0.7a90
Profiling
State Mgmt
v0.8aX/0.9aX
SSL/SMB
STABLE releases
BroLite
v1.1/v1.2
when Stmt
Resource tuning
Broccoli
DPD
v1.5
BroControl
Bro 2.0
New Scripts
Bro SDCI
Bro 2.1
IPv6
Input Framework
LBNL starts using Bro operationally
v0.4
HTTP analysis
Scan detector
IP fragments
Linux support
v0.7a175/0.8aX
Signatures
SMTP
IPv6 support
User manual
v1.0
BinPAC
IRC/RPC analyzers
64-bit support
Sane version numbers
v1.4
DHCP/BitTorrent
HTTP entities
NetFlow
Bro Lite Deprecated
v0.7a48
Consistent CHANGES
0.8a37
Communication
Persistence
Namespaces
Log Rotation
v1.3
Ctor expressions
GeoIP
Conn Compressor

Bro History

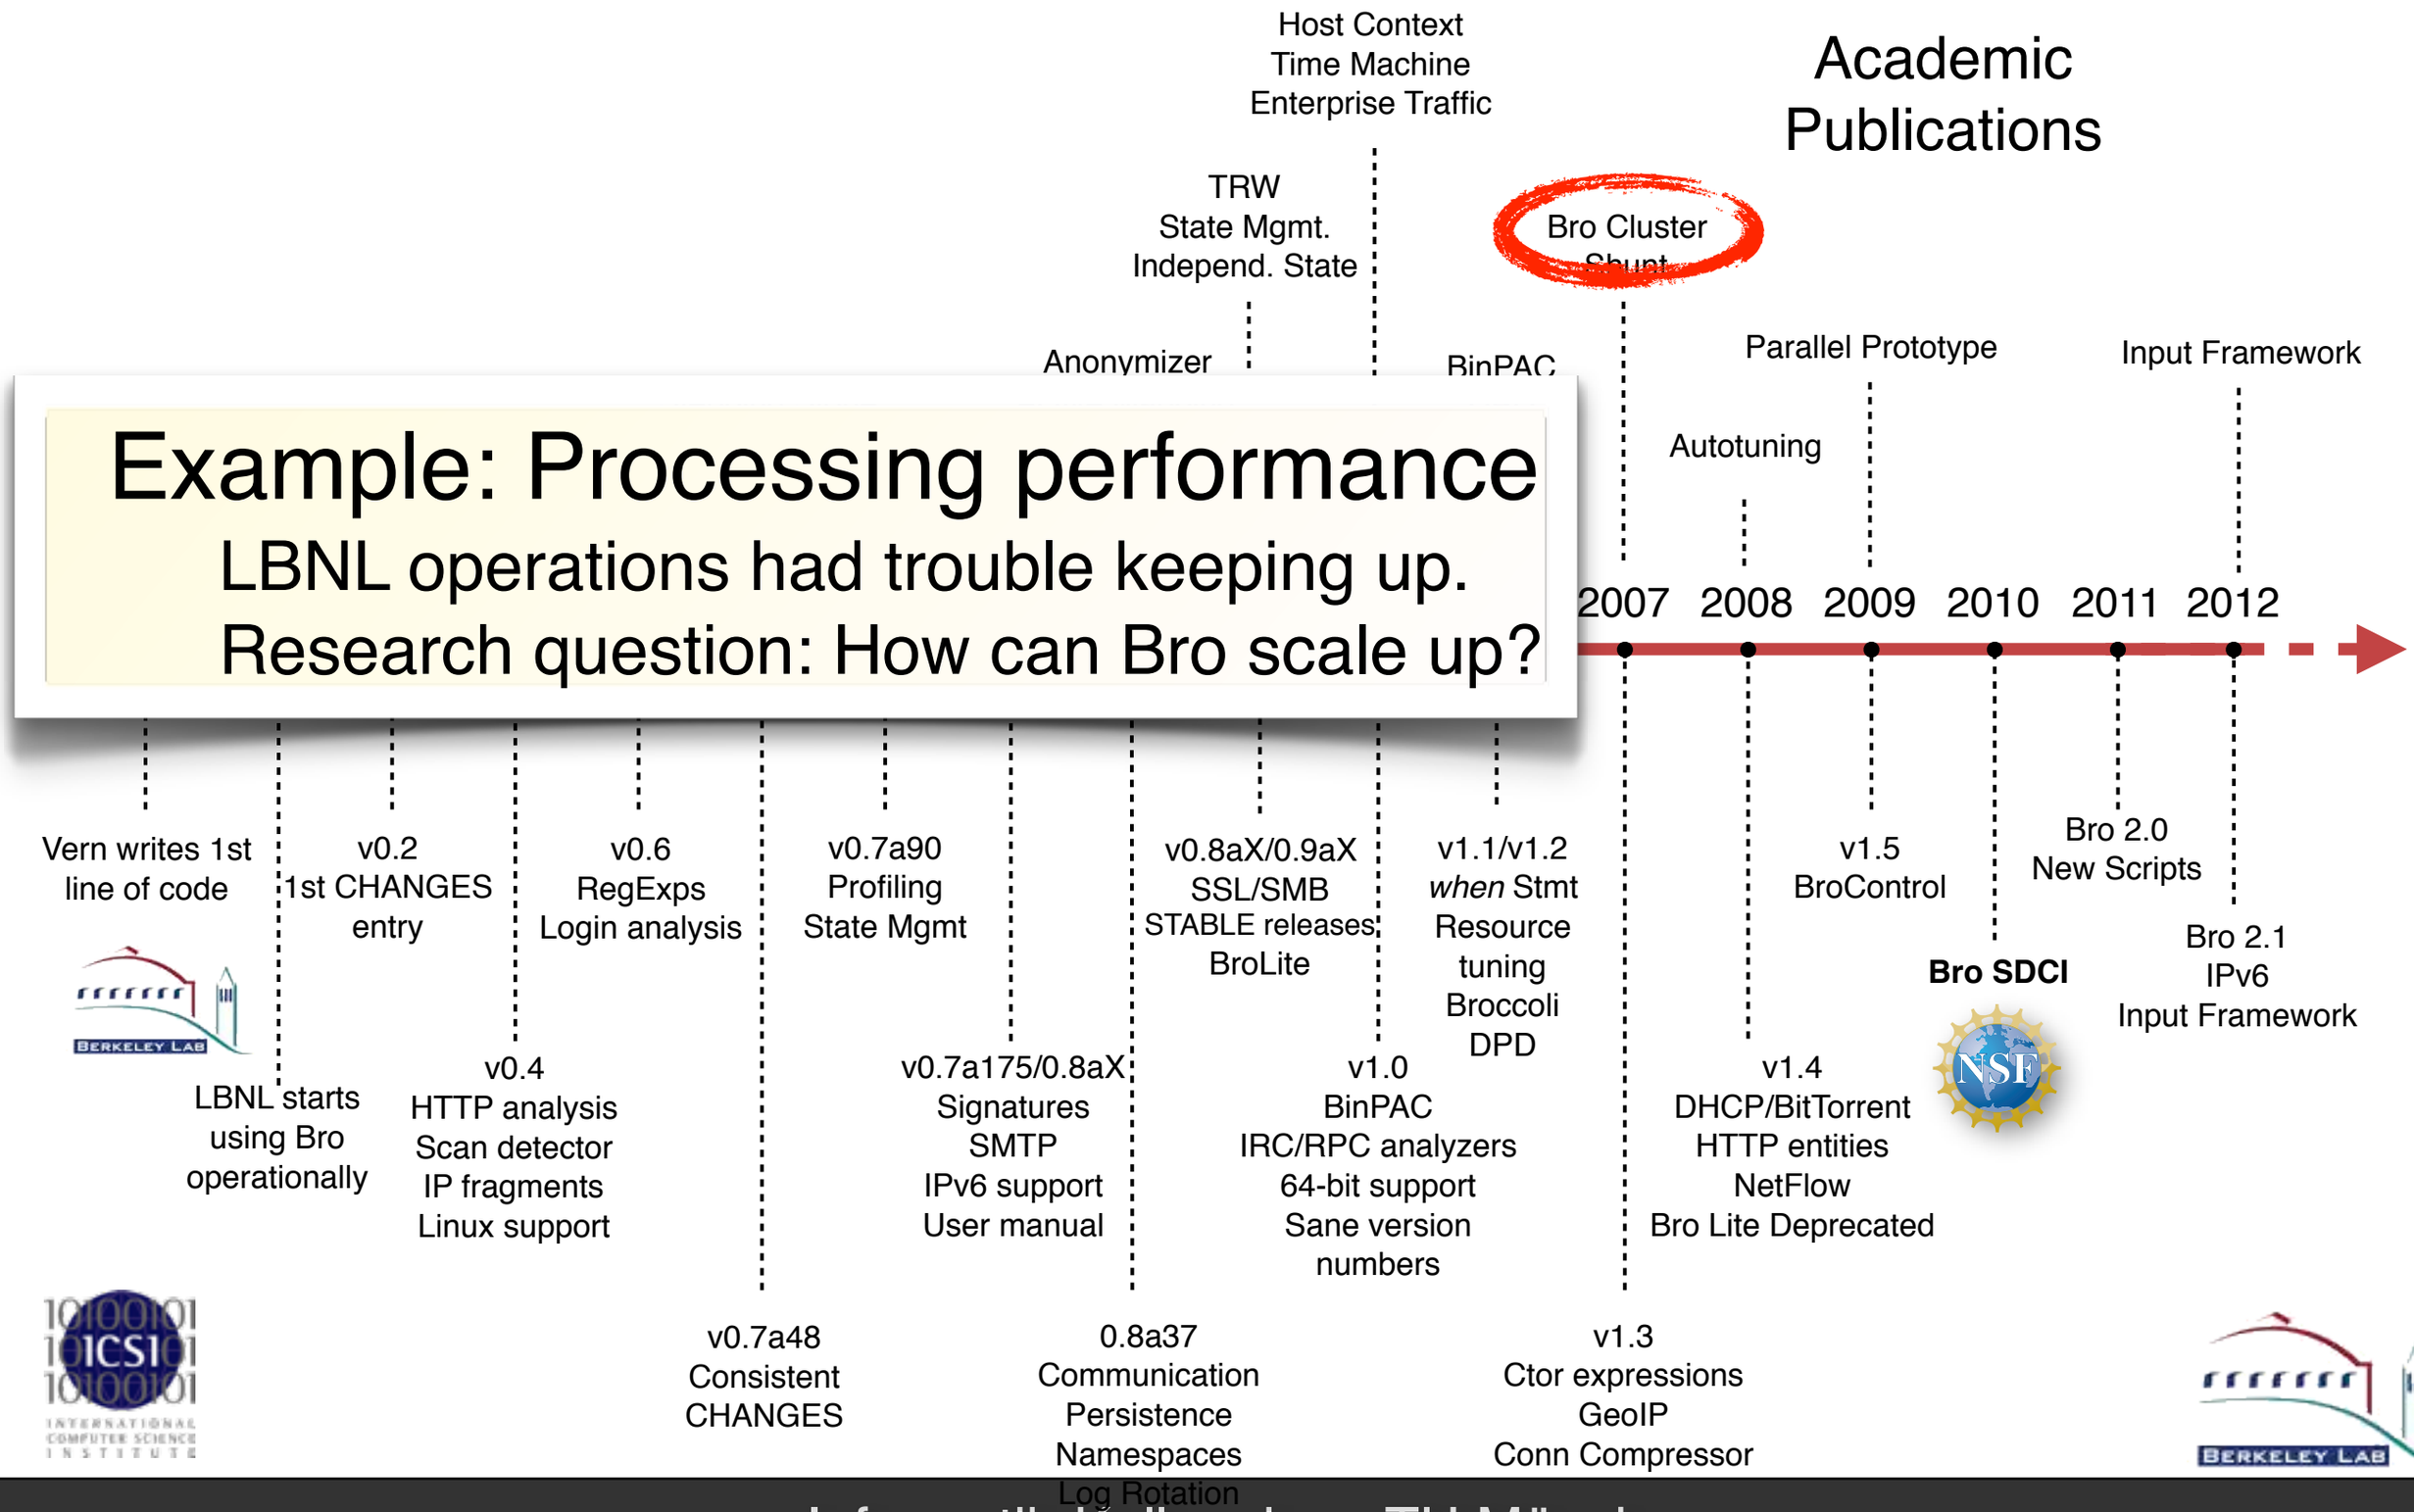

Example: Processing performance

LBNL operations had trouble keeping up.

Research question: How can Bro scale up?

Load-balancing Architecture

Load-balancing Architecture

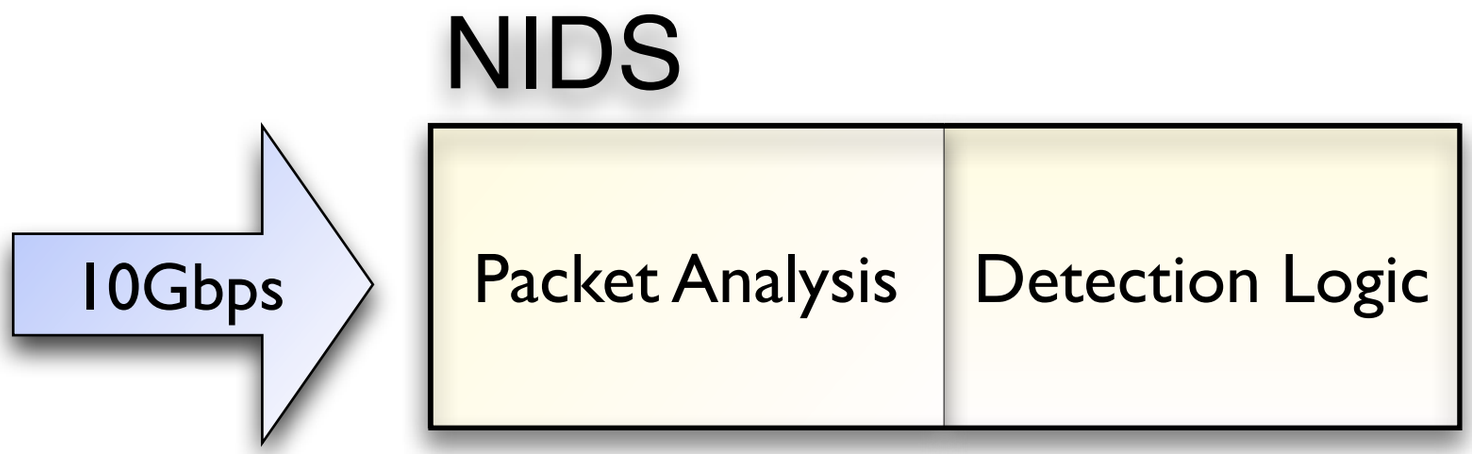

Load-balancing Architecture

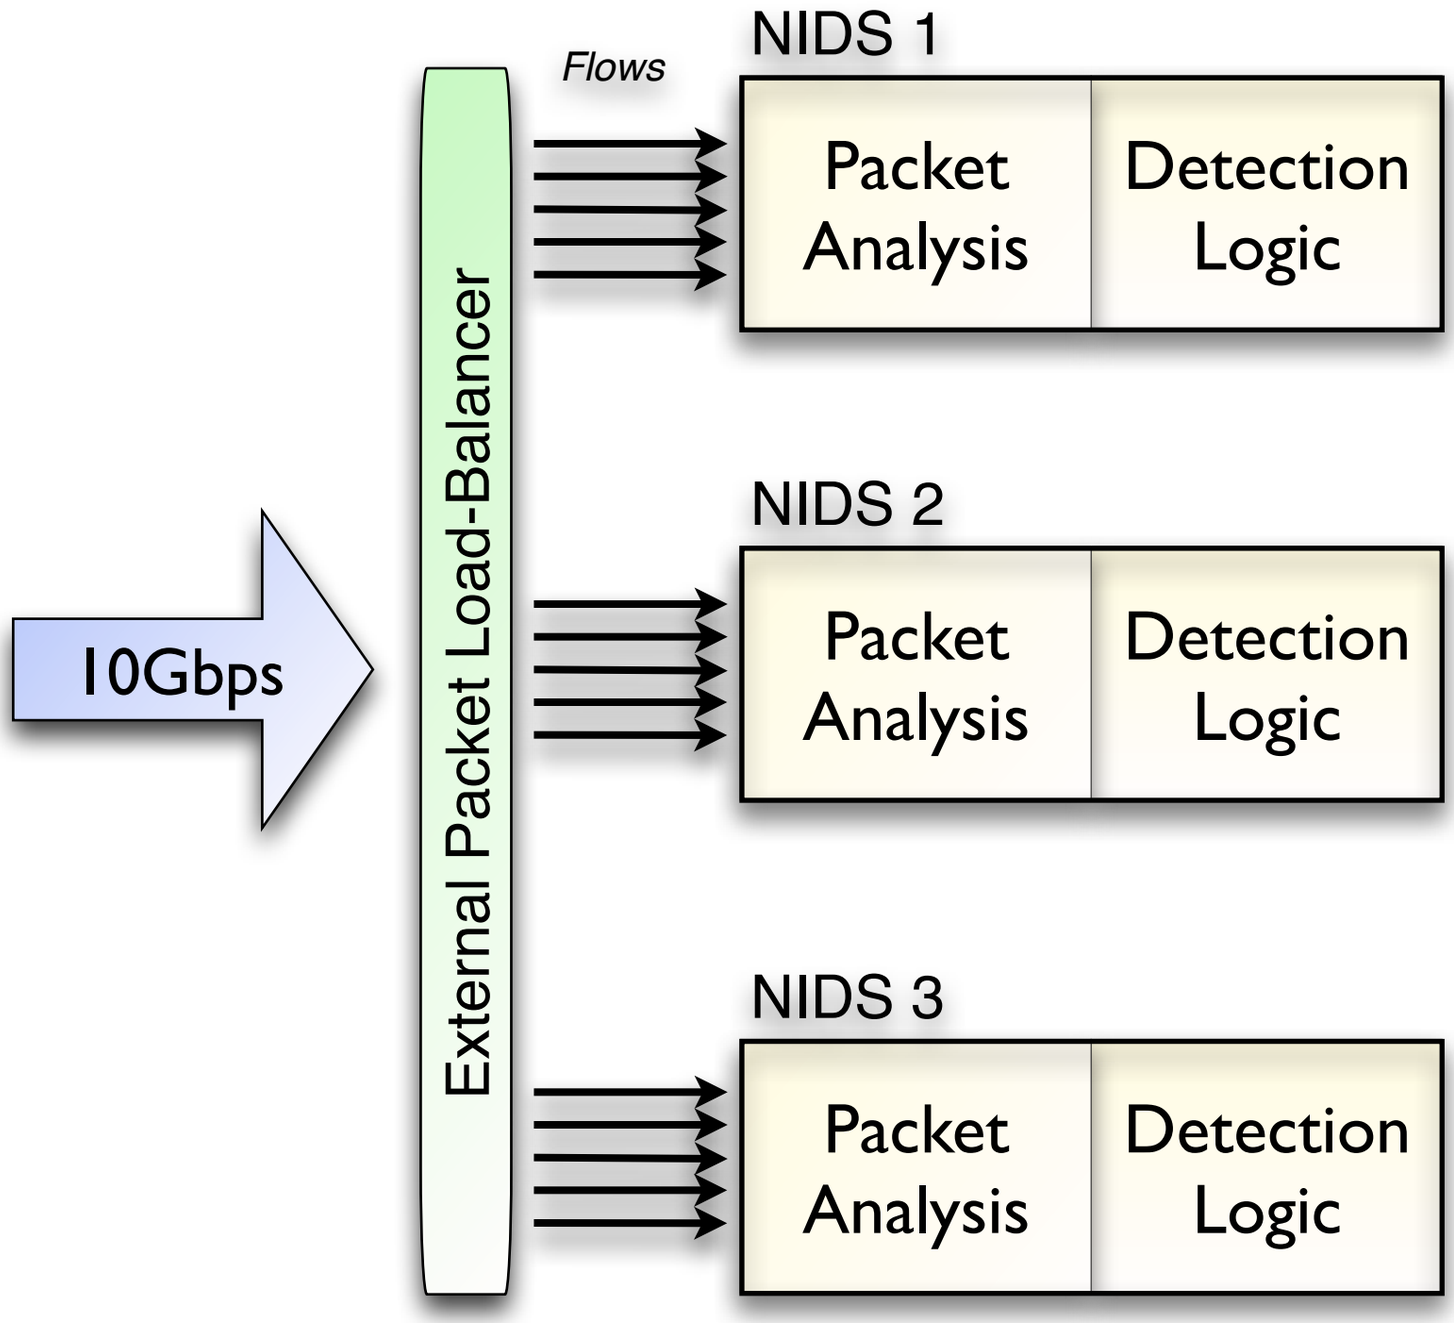

Load-balancing Architecture

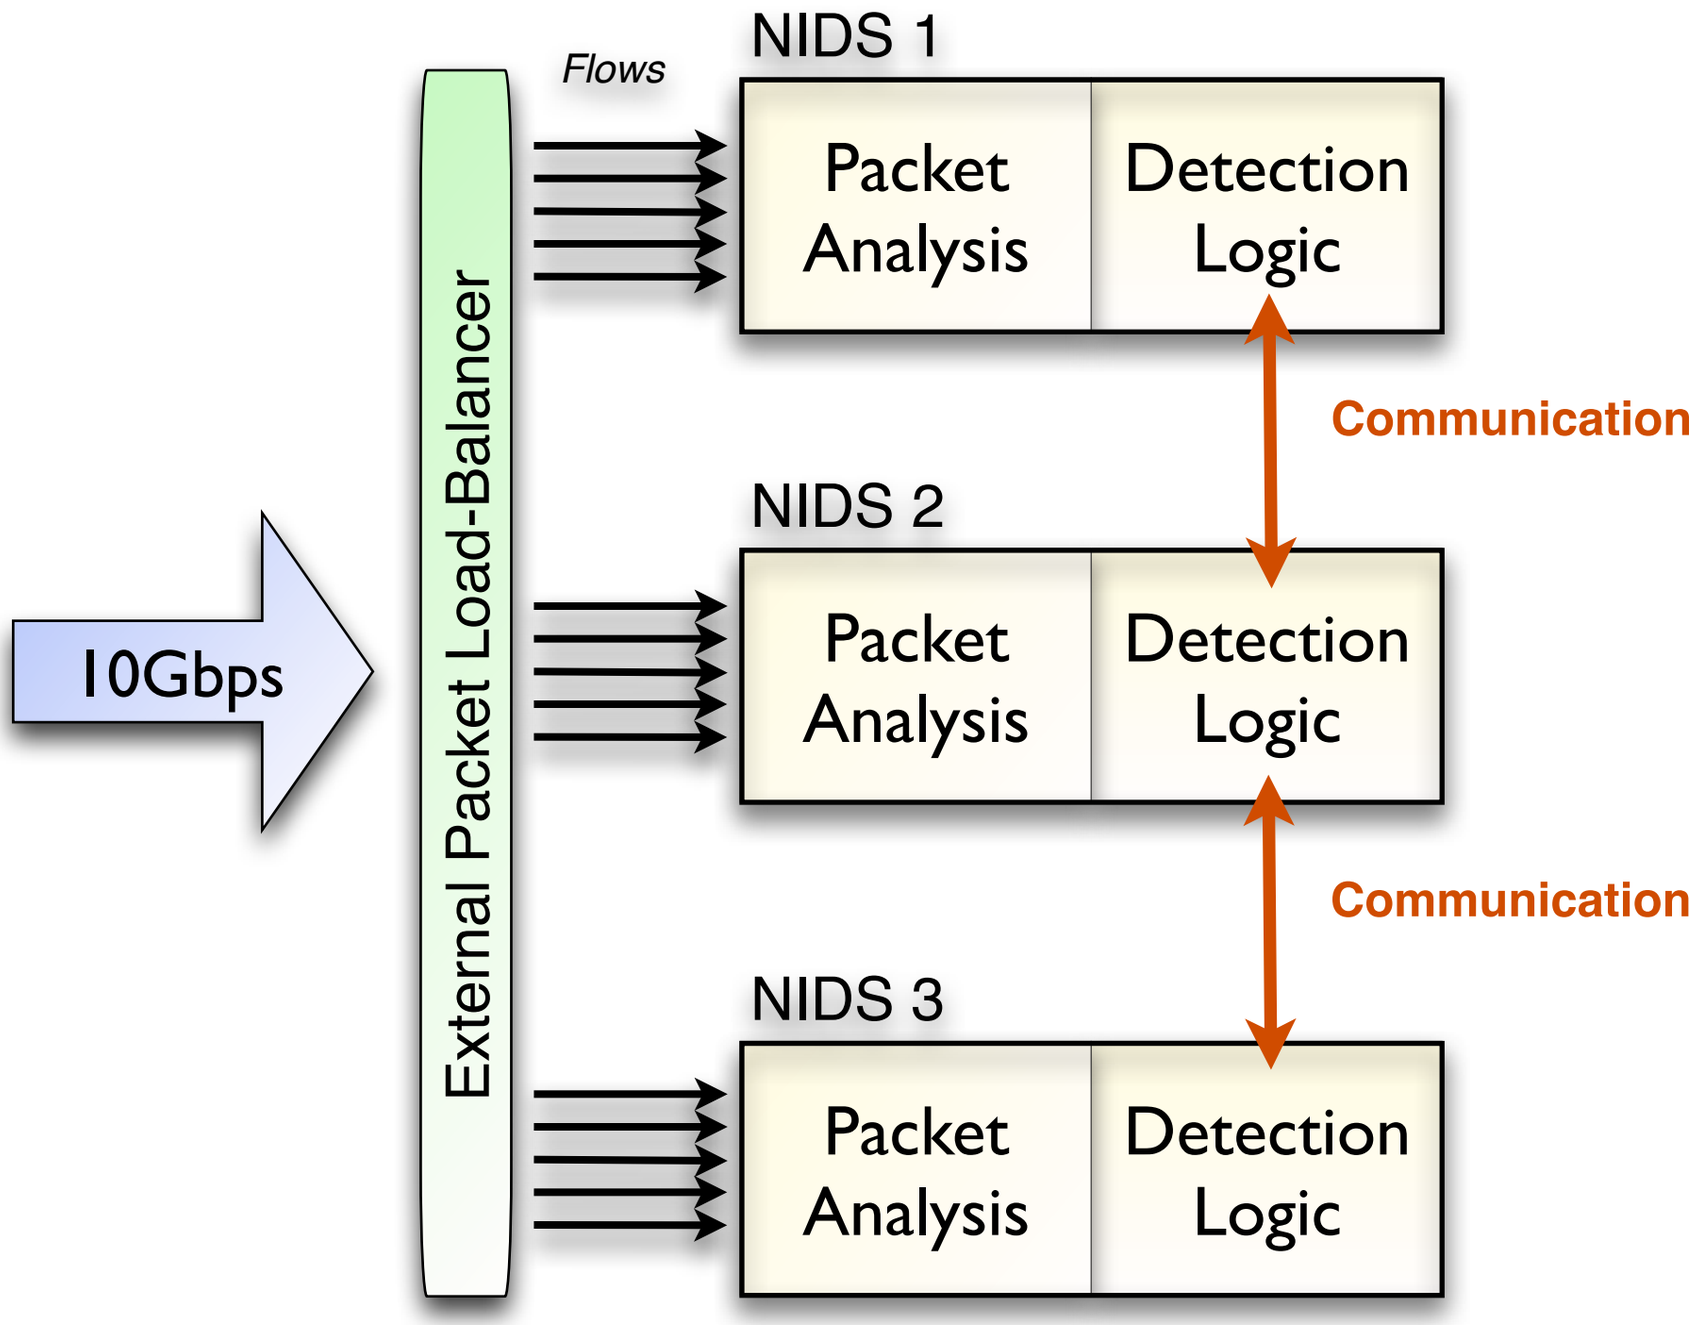

Load-balancing Architecture

"Bro Cluster"

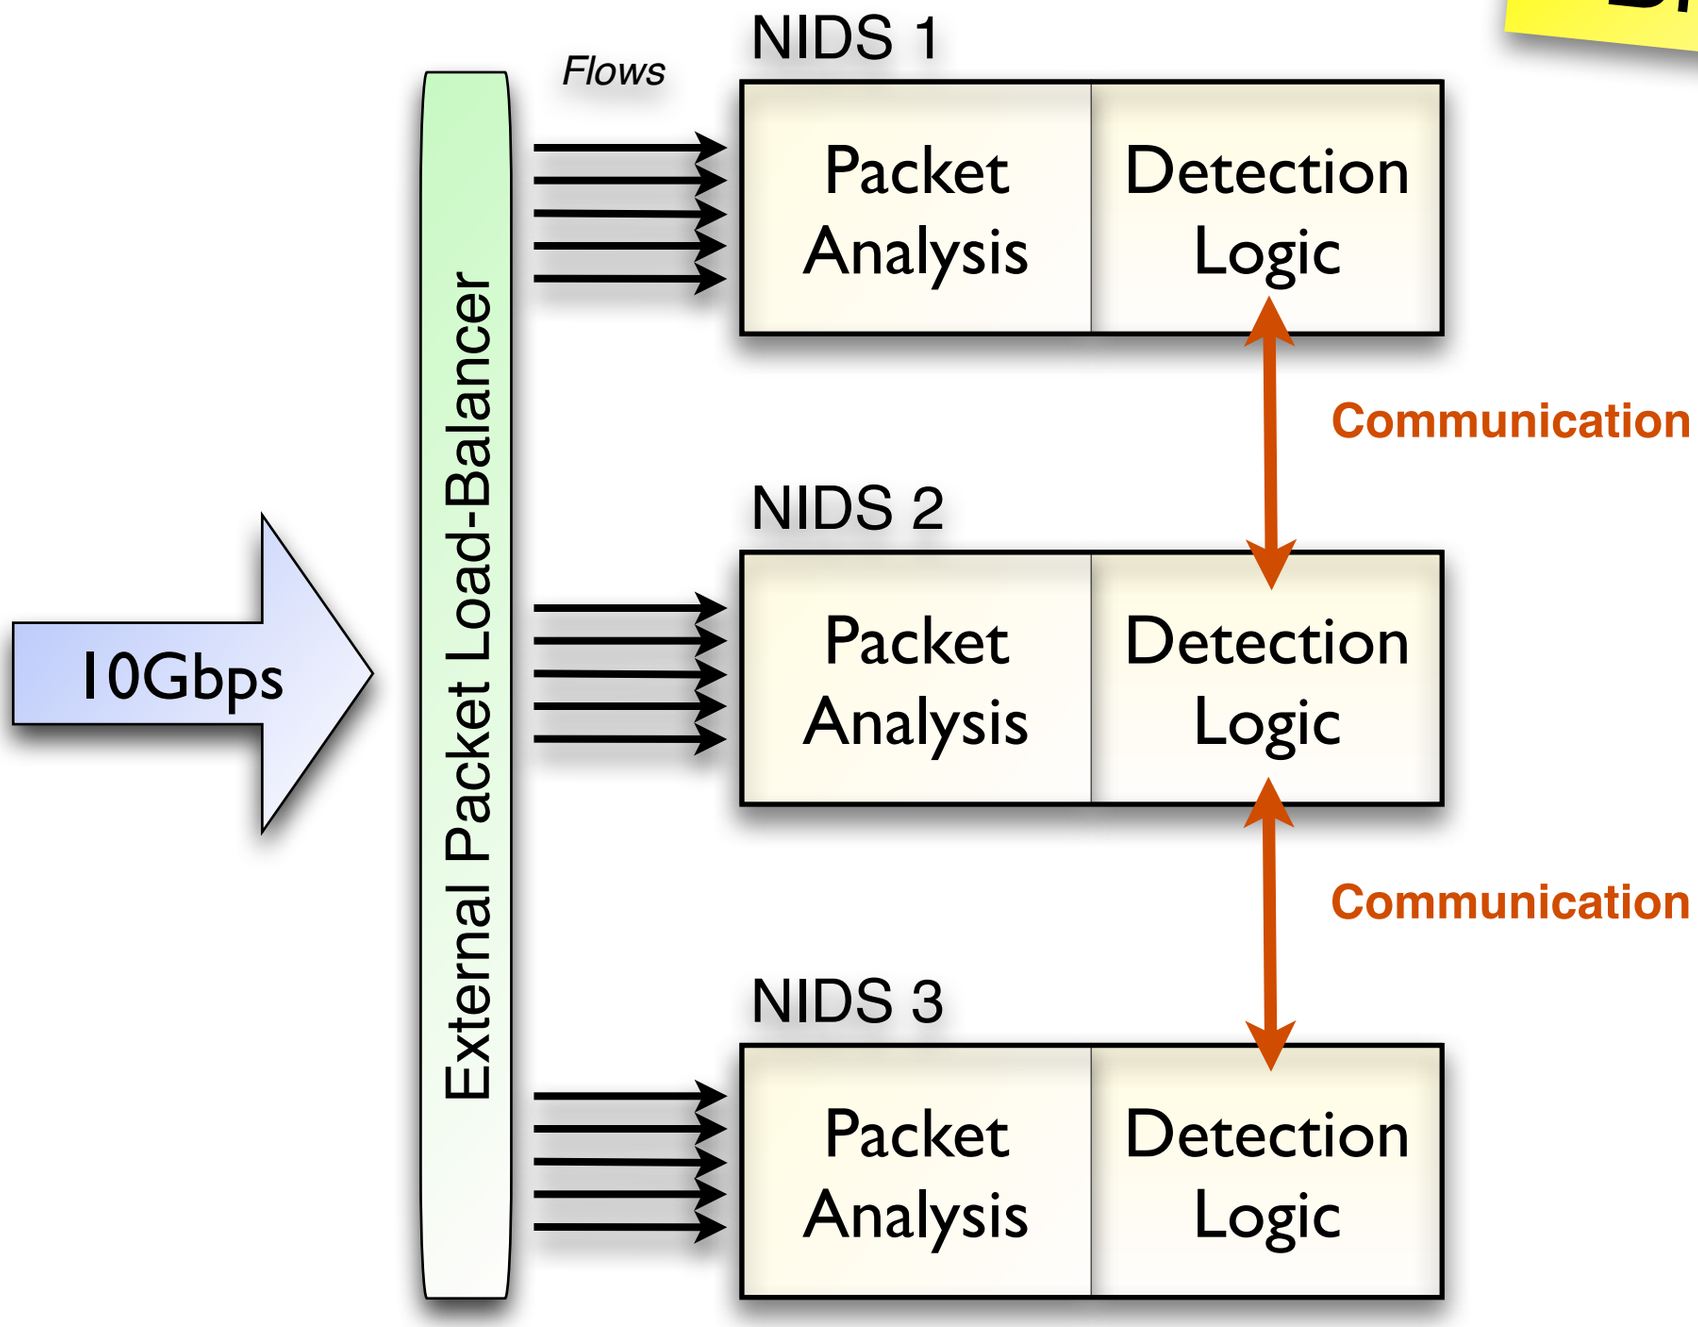

Cluster goes Operation

Cluster goes Operation

Load-balancer operates a line-rate.

1. Receive packet.
2. Calculate hash.
3. Rewrite MAC address.
4. Send packet out.

Cluster goes Operation

Load-balancer operates a line-rate.

1. Receive packet.
2. Calculate hash.
3. Rewrite MAC address.
4. Send packet out.

Research prototype limited to 2 Gb/s.

Linux box using kernel-level *Click*.

Cluster goes Operation

Load-balancer operates a line-rate.

1. Receive packet.
2. Calculate hash.
3. Rewrite MAC address.
4. Send packet out.

Research prototype limited to 2 Gb/s.

Linux box using kernel-level *Click*.

LBNL wanted reliable 10 Gb/s device.

No robust line-rate solution available in 2007.

Eventually contracted vendor to build device.

A Production Load-Balancer

cFlow: 10GE line-rate, stand-alone load-balancer

10 Gb/s in/out
Web & CLI
Filtering capabilities

Available from cPacket

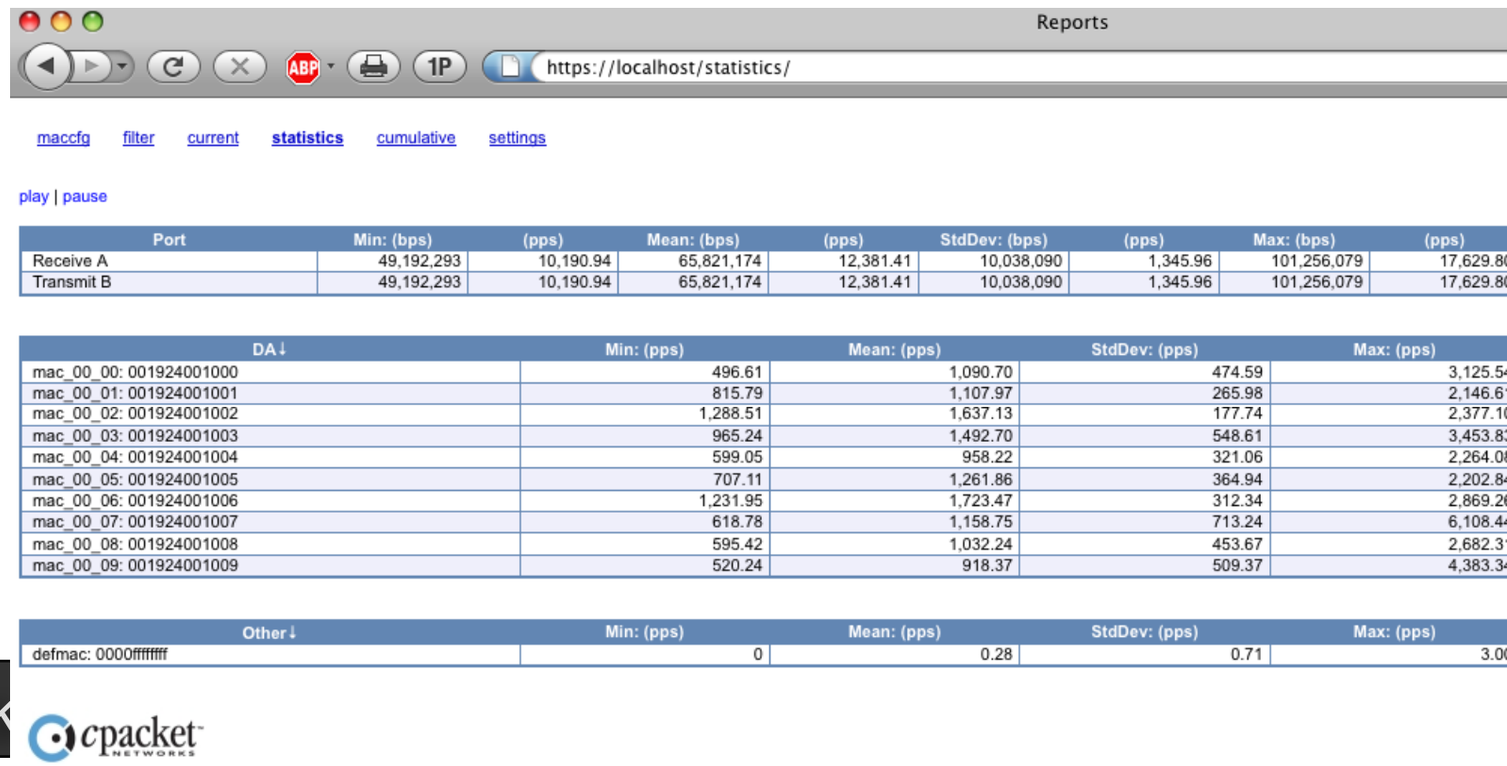

Port	Min: (bps)	(pps)	Mean: (bps)	(pps)	StdDev: (bps)	(pps)	Max: (bps)	(pps)
Receive A	49,192,293	10,190.94	65,821,174	12,381.41	10,038,090	1,345.96	101,256,079	17,629.8
Transmit B	49,192,293	10,190.94	65,821,174	12,381.41	10,038,090	1,345.96	101,256,079	17,629.8

DA↓	Min: (pps)	Mean: (pps)	StdDev: (pps)	Max: (pps)
mac_00_00: 001924001000	496.61	1,090.70	474.59	3,125.5
mac_00_01: 001924001001	815.79	1,107.97	265.98	2,146.6
mac_00_02: 001924001002	1,288.51	1,637.13	177.74	2,377.1
mac_00_03: 001924001003	965.24	1,492.70	548.61	3,453.8
mac_00_04: 001924001004	599.05	958.22	321.06	2,264.0
mac_00_05: 001924001005	707.11	1,261.86	364.94	2,202.8
mac_00_06: 001924001006	1,231.95	1,723.47	312.34	2,869.2
mac_00_07: 001924001007	618.78	1,158.75	713.24	6,108.4
mac_00_08: 001924001008	595.42	1,032.24	453.67	2,682.3
mac_00_09: 001924001009	520.24	918.37	509.37	4,383.3

Other↓	Min: (pps)	Mean: (pps)	StdDev: (pps)	Max: (pps)
defmac: 0000ffffffff	0	0.28	0.71	3.0

A Production Load-Balancer

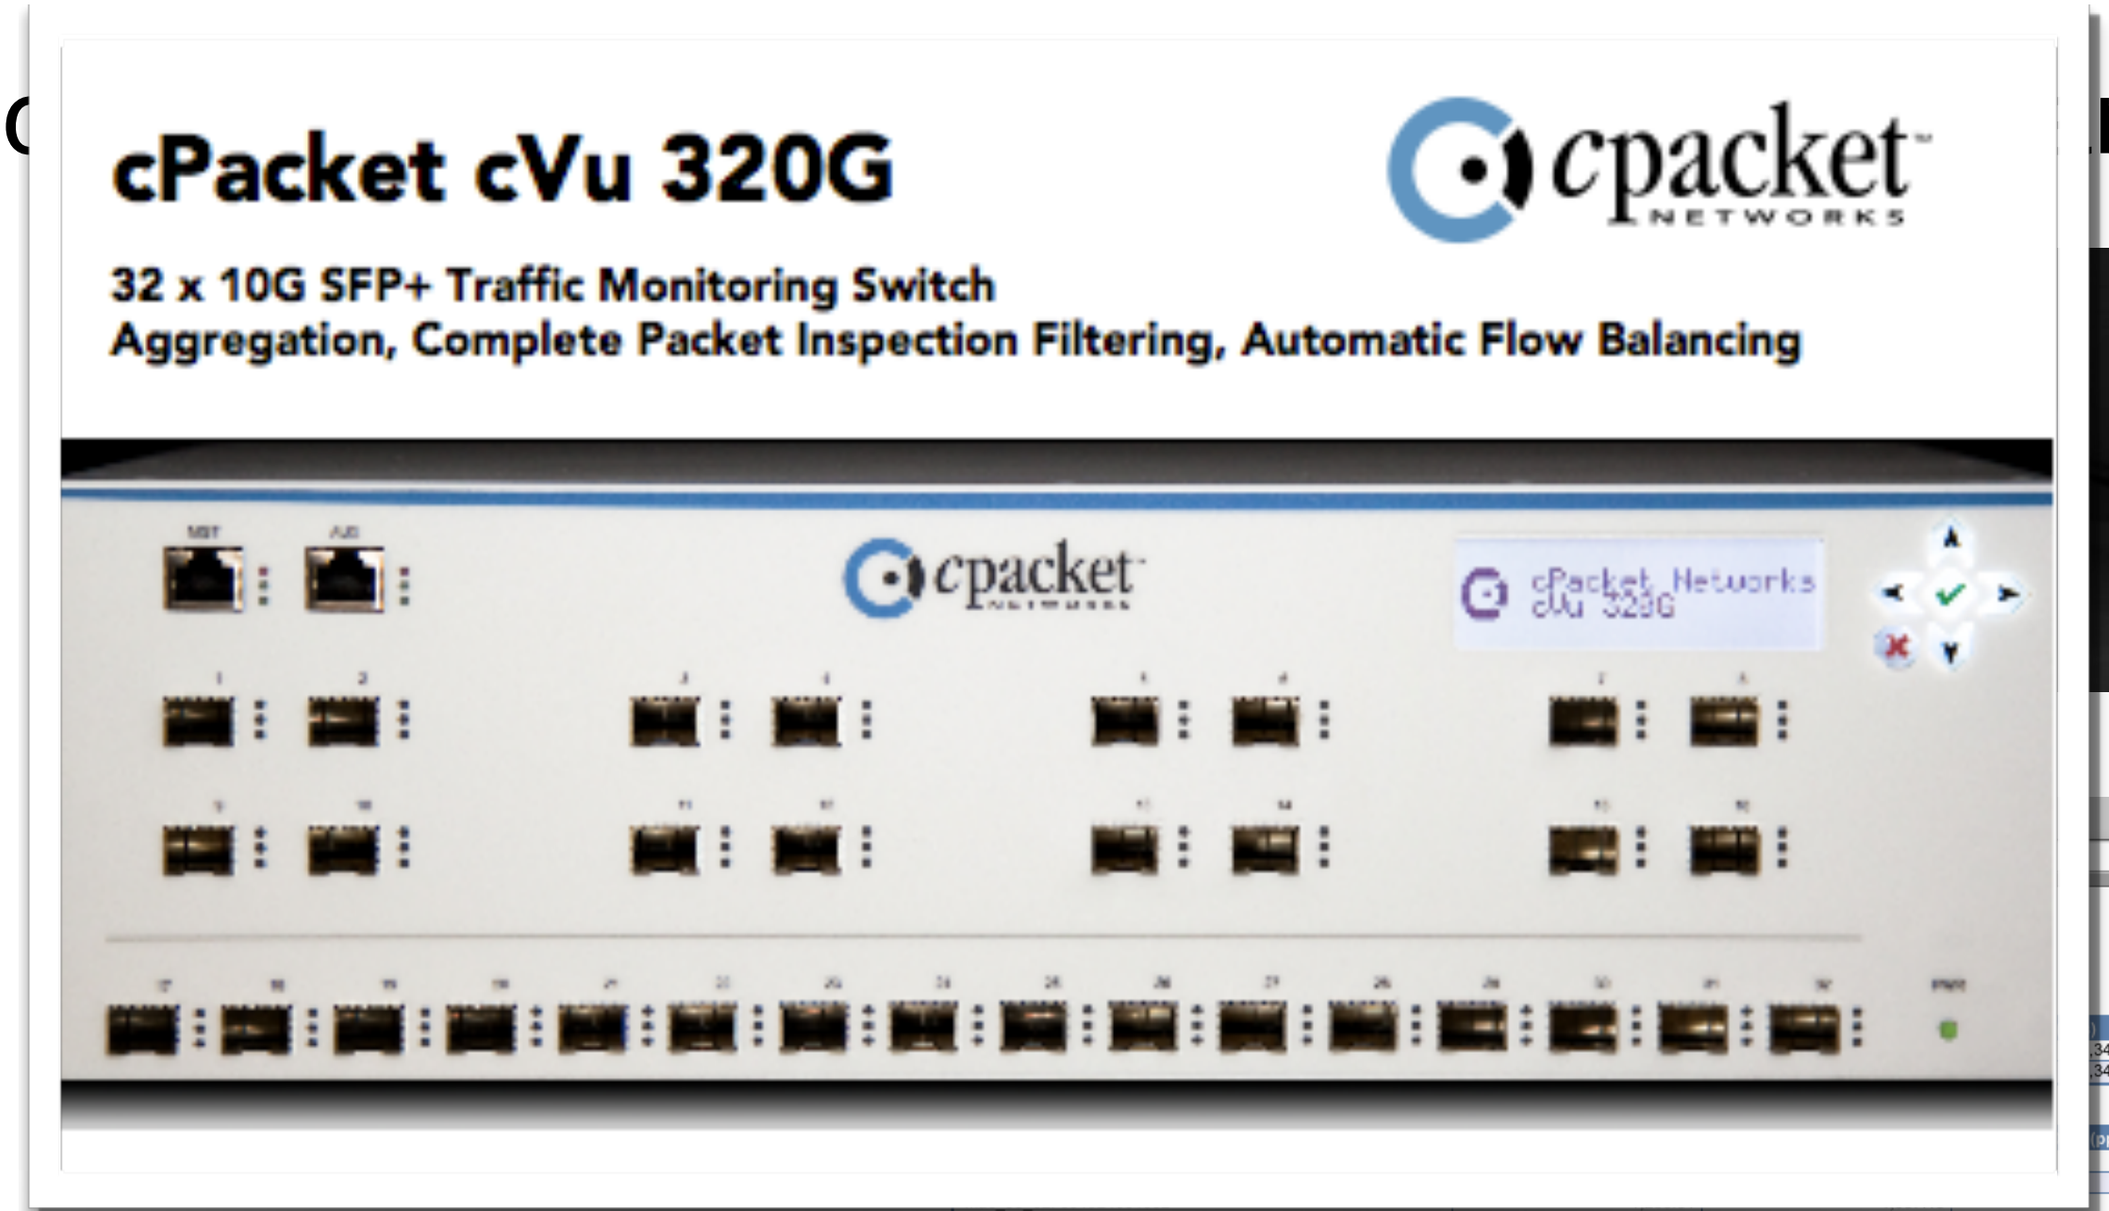

Available from cPacket

	Max: (bps)	(pps)
	101,256,079	17,629.80
	101,256,079	17,629.80

	(pps)	Max: (pps)
	474.59	3,125.54
	265.98	2,146.61
	177.74	2,377.10
mac_00_03: 001924001003	548.61	3,453.83
mac_00_04: 001924001004	321.06	2,264.08
mac_00_05: 001924001005	364.94	2,202.84
mac_00_06: 001924001006	312.34	2,869.26
mac_00_07: 001924001007	713.24	6,108.44
mac_00_08: 001924001008	453.67	2,682.31
mac_00_09: 001924001009	509.37	4,383.34

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

Indiana University

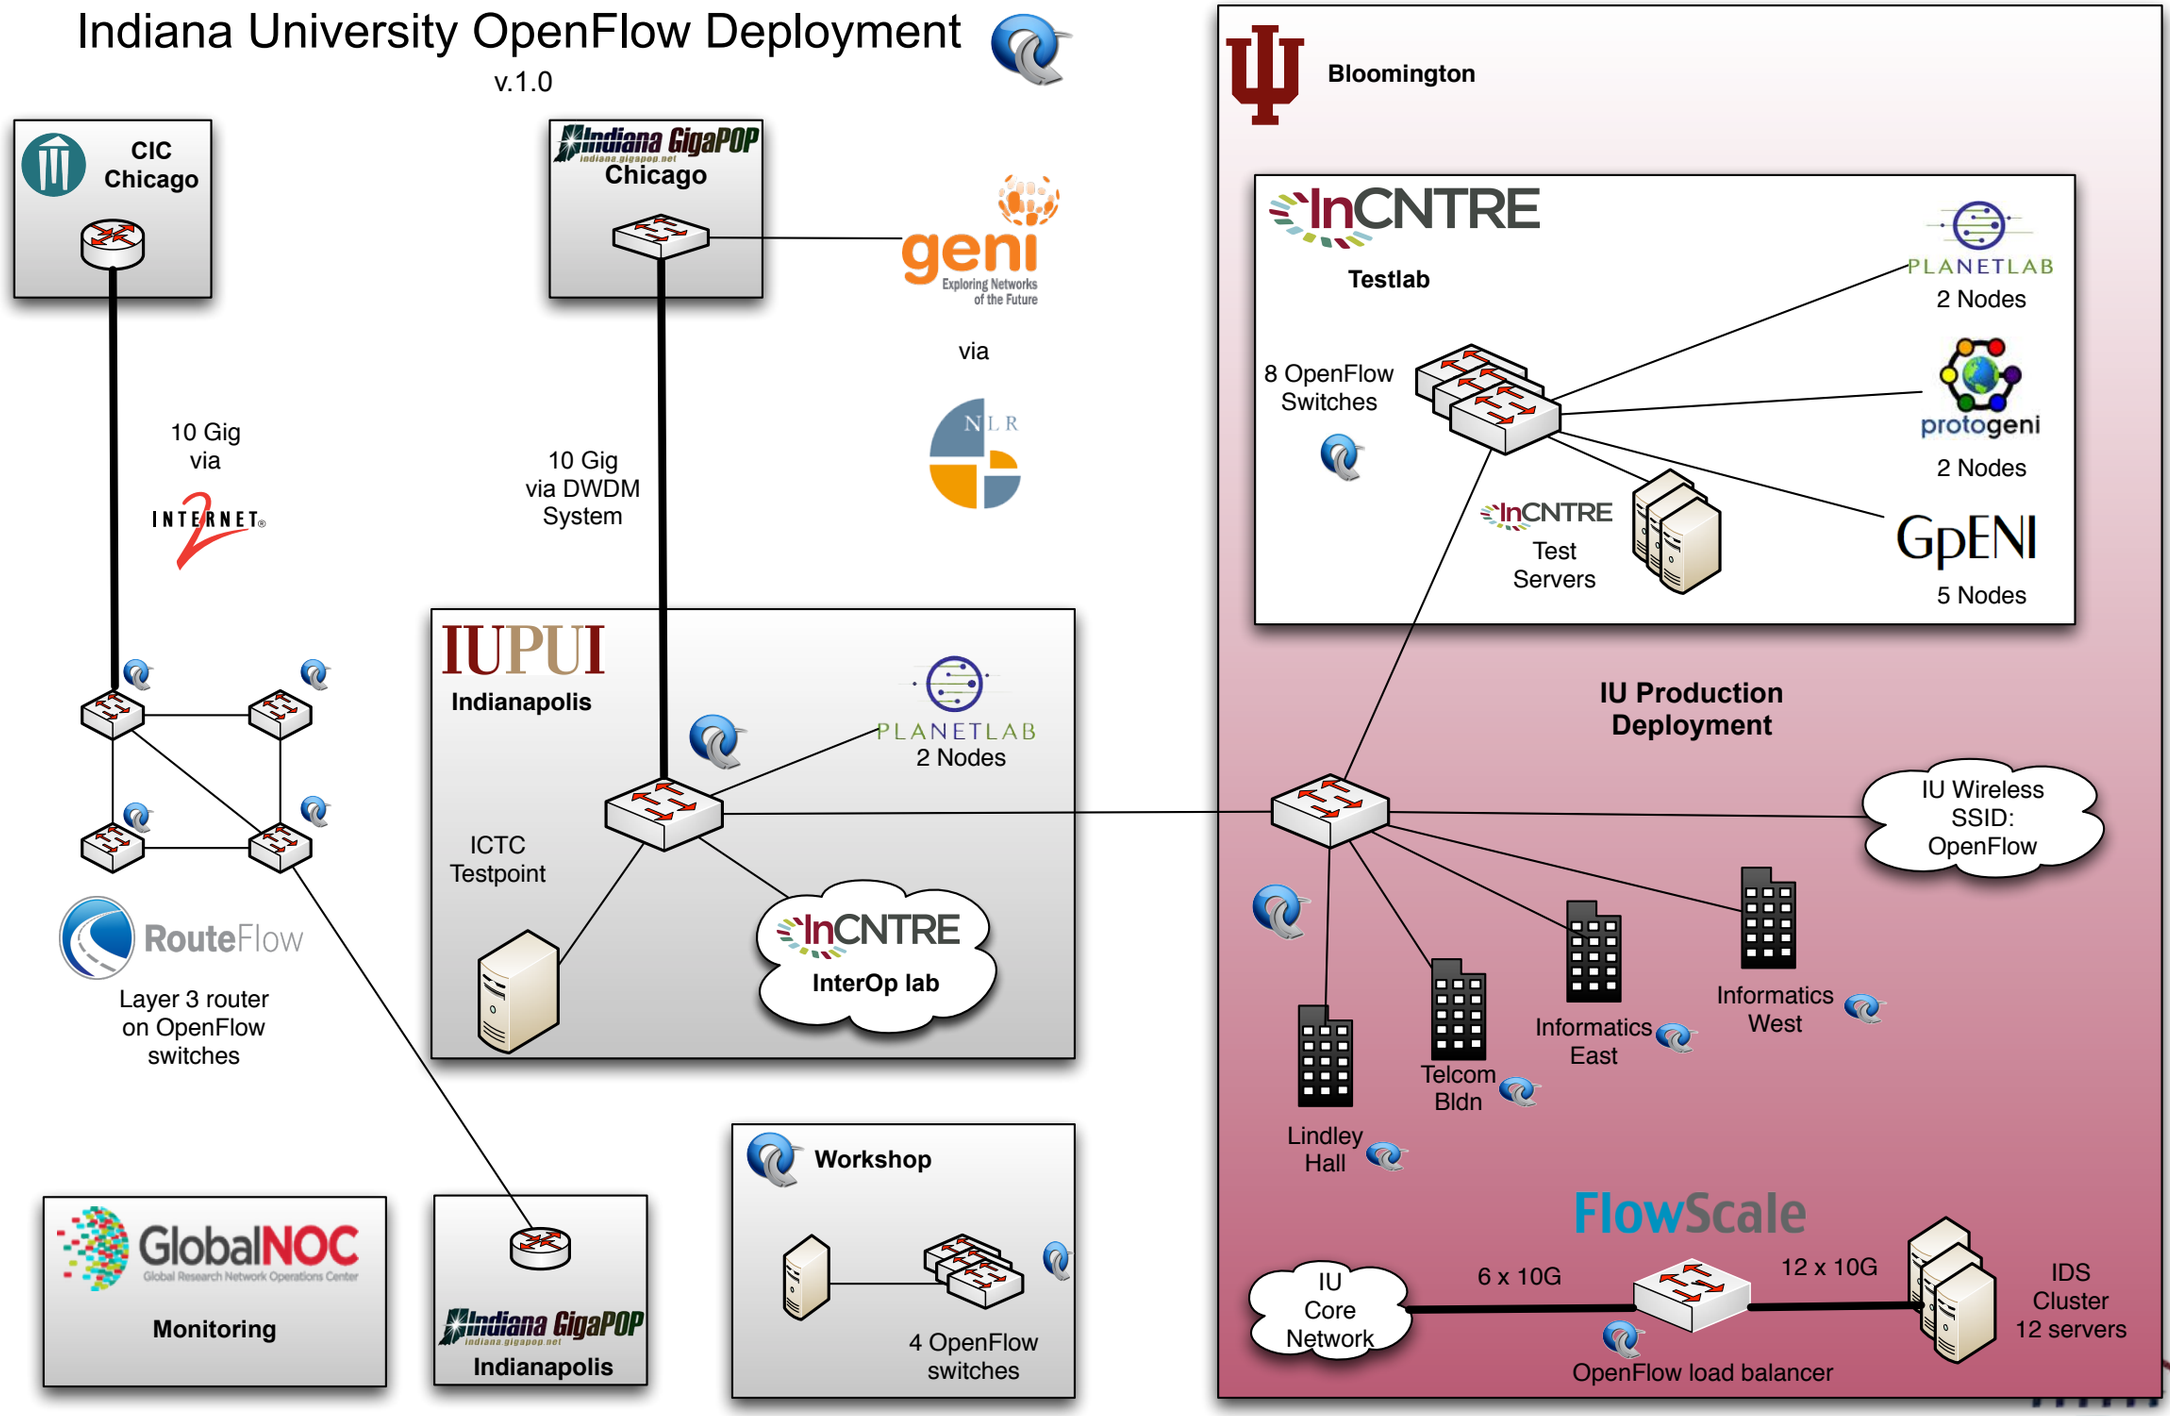

Source: Indiana University

Indiana University

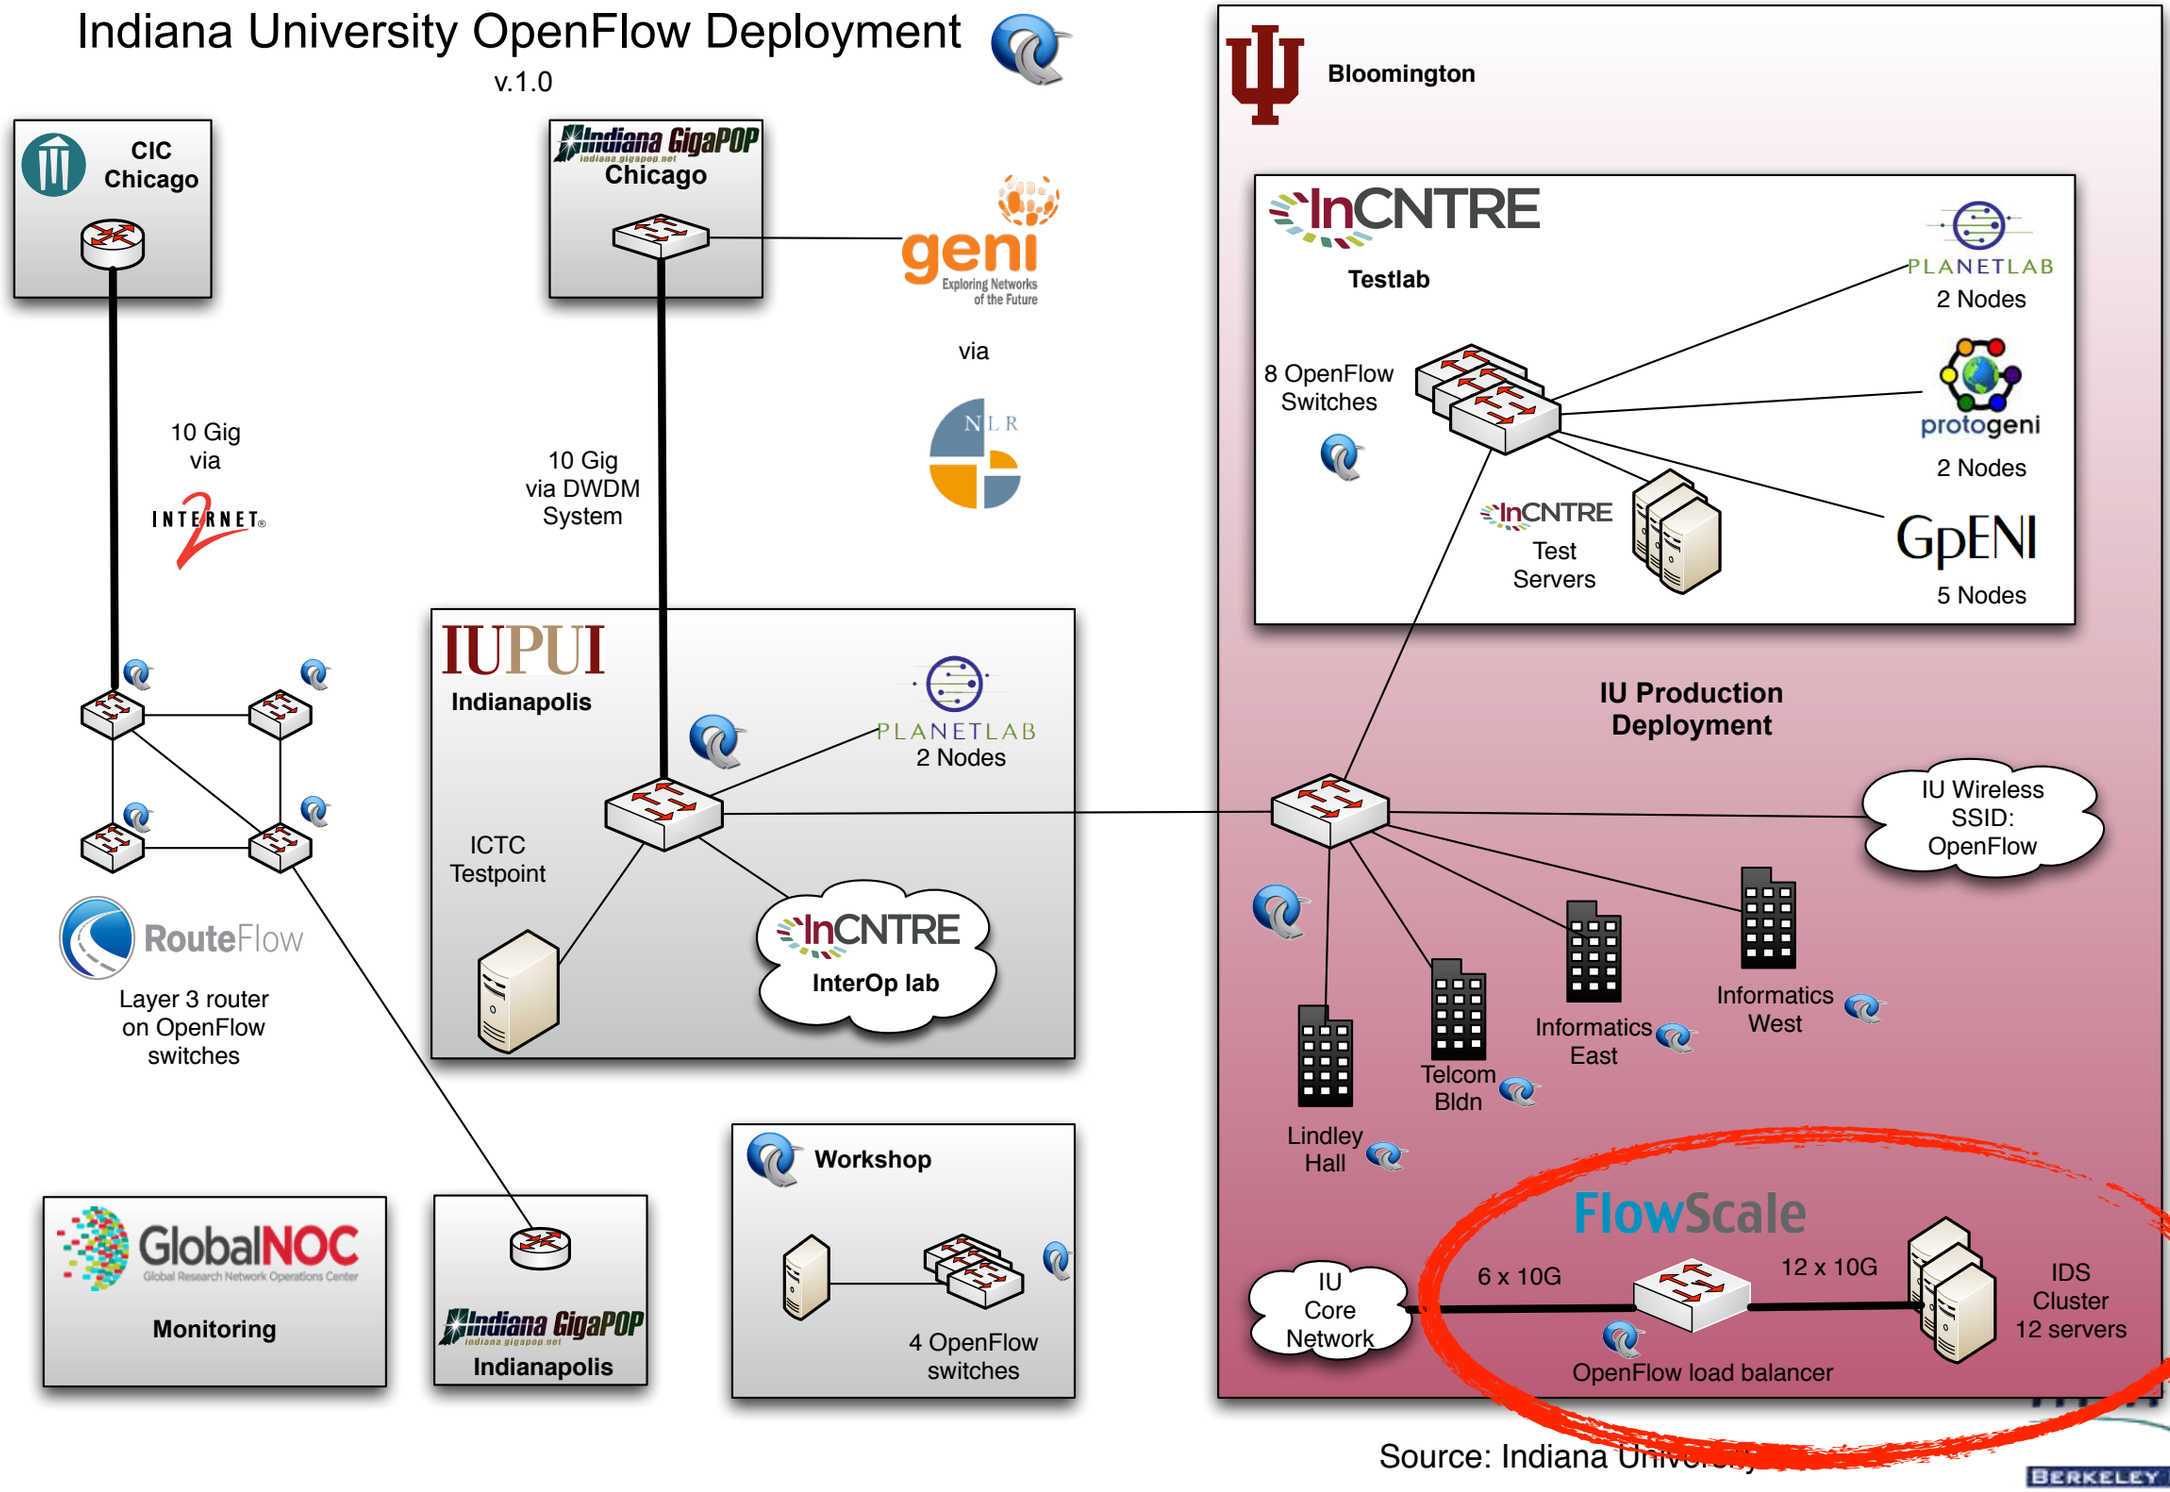

Source: Indiana University

Open-Flow Cluster at Indiana University

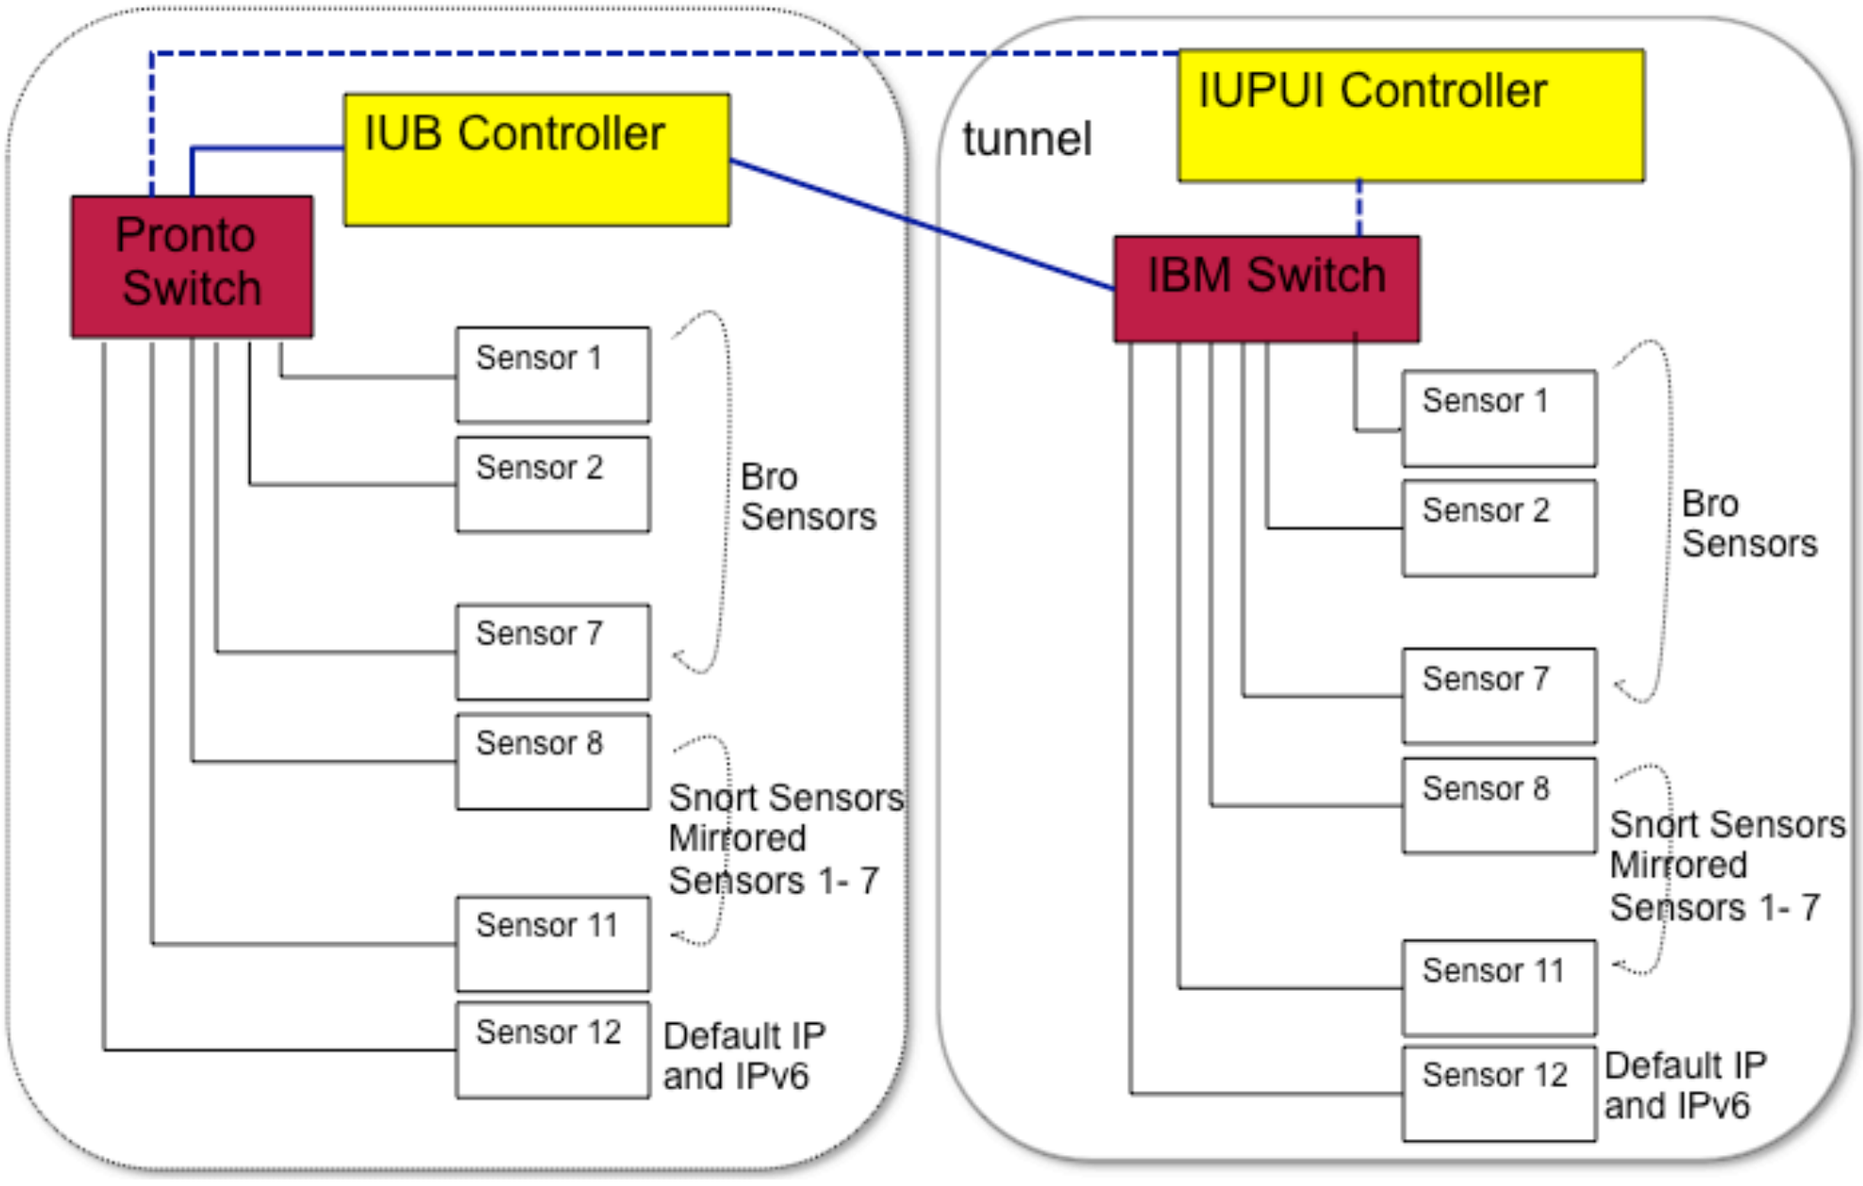

Source: Indiana University

Next Stop: 100 Gb/s

DOE/ESNet
100G Advanced Networking Initiative

Source: ESNet

Production Backbone in Planing

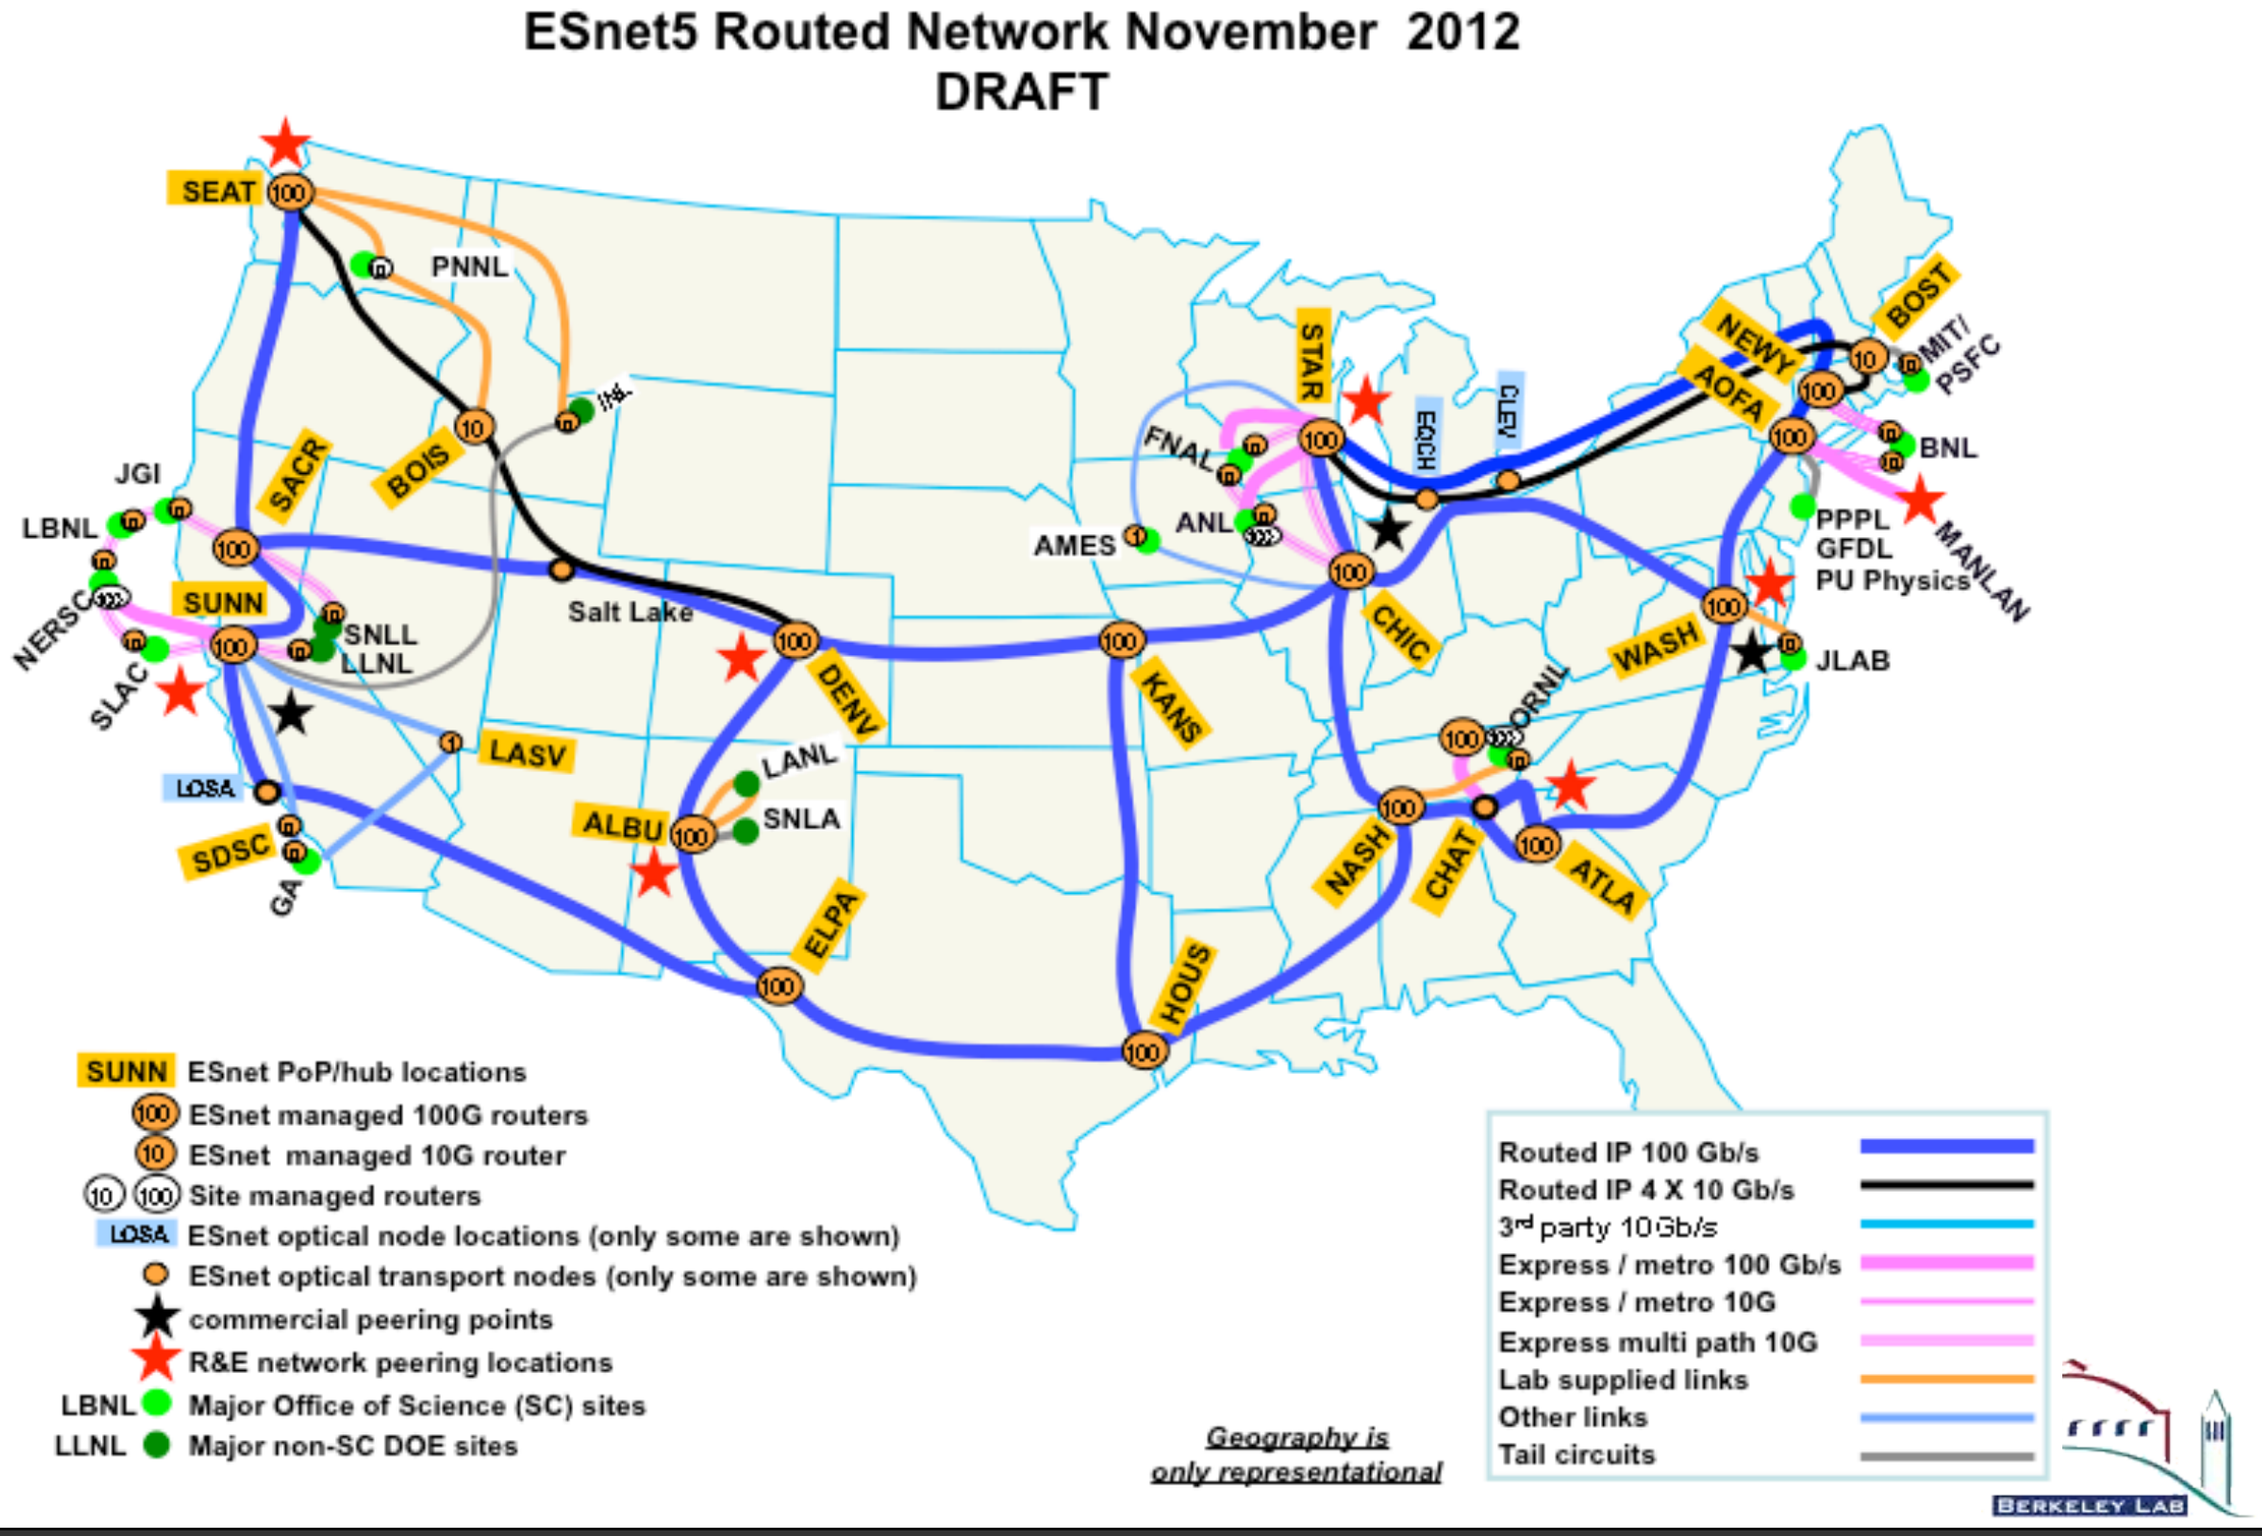

100 Gb/s Load-balancer

100 Gb/s Load-balancer

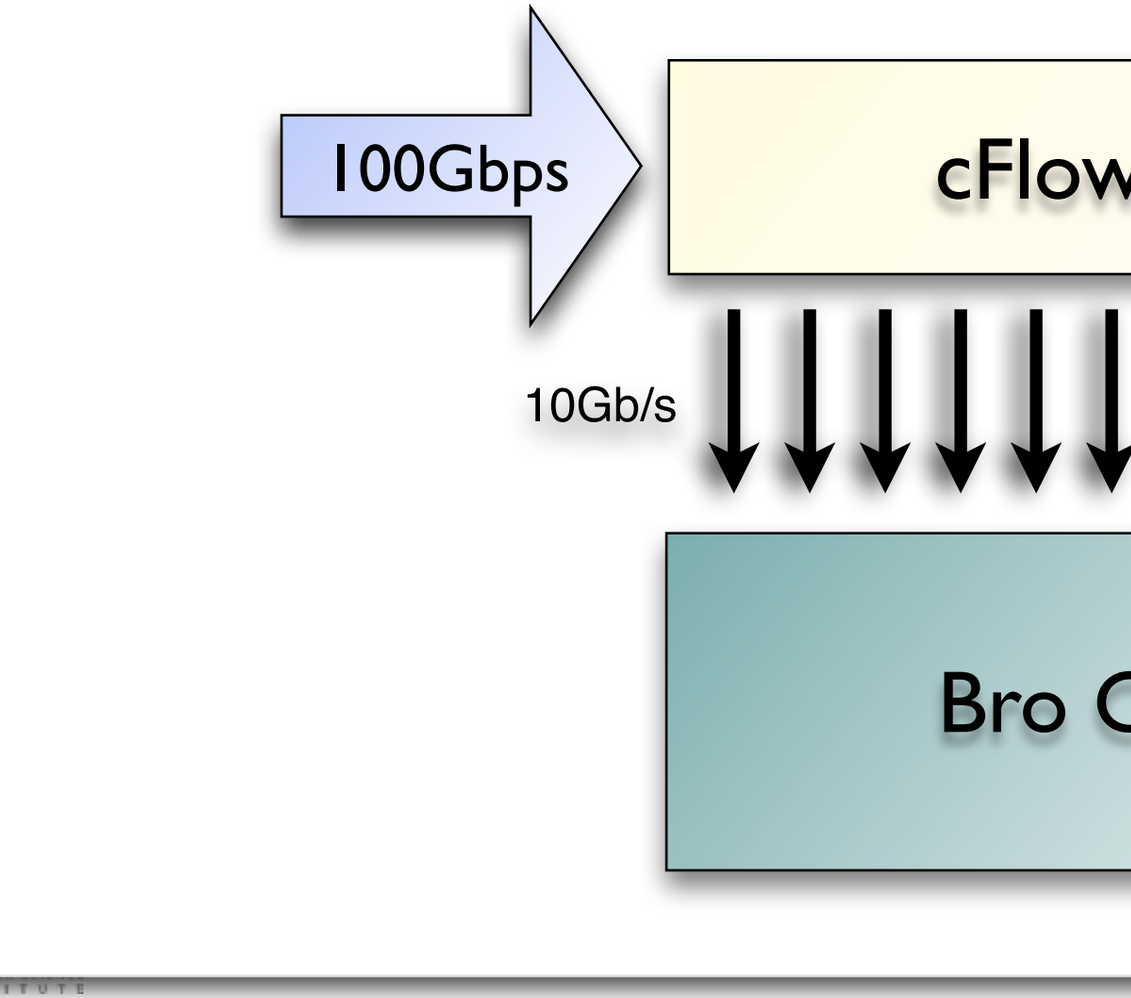

100 Gb/s Load-balancer

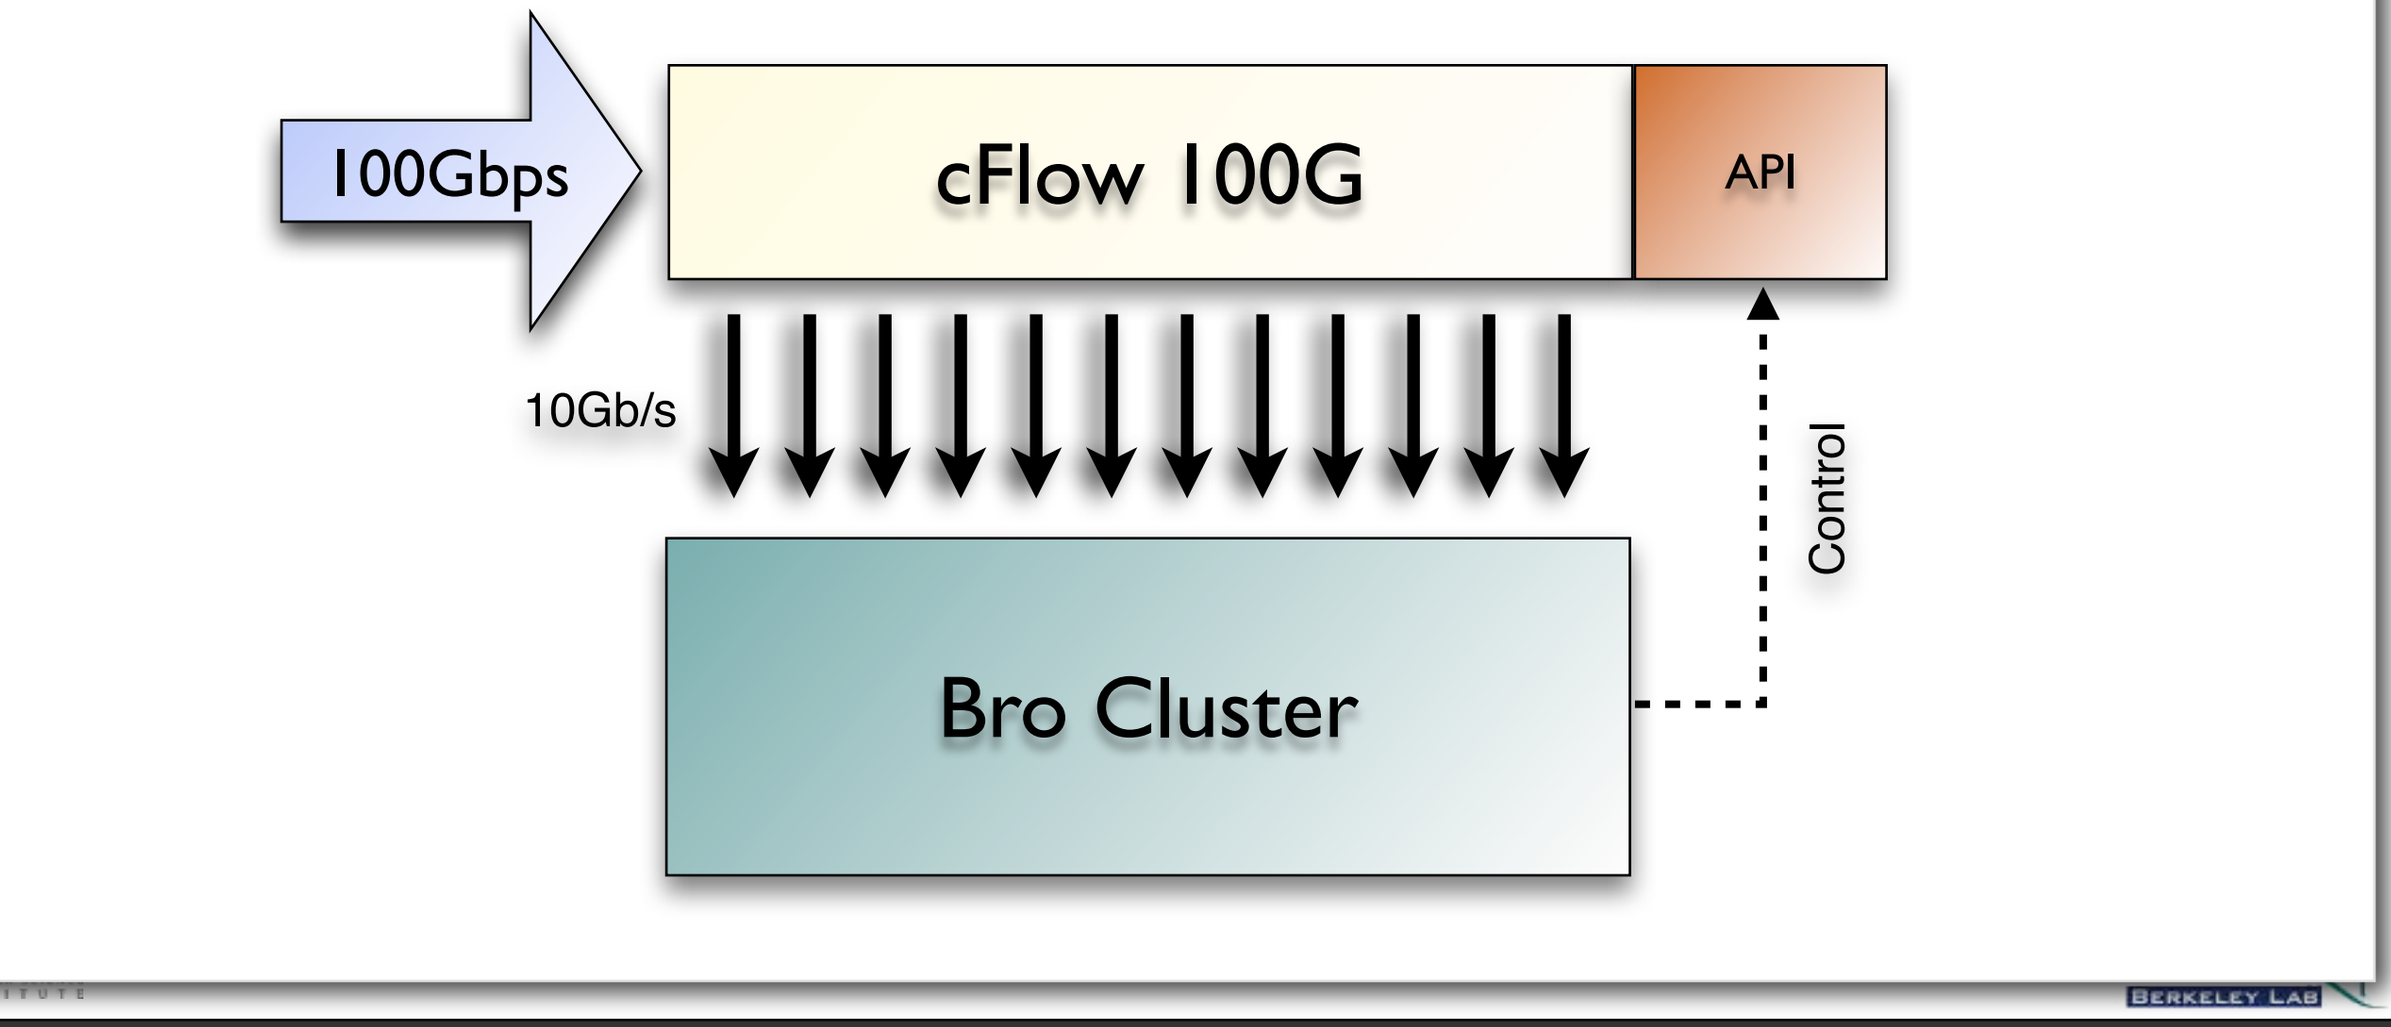

Going Multi-Core

Going Multi-Core

Bro is single-threaded

- Backends have multiple cores, which are mostly idling.
- Work-around: “Cluster in a box”

Going Multi-Core

Bro is single-threaded

- Backends have multiple cores, which are mostly idling.
- Work-around: “Cluster in a box”

We *really* want multi-threading.

- Must scale well with increasing numbers of cores.
- Must be transparent to the operator.

Going Multi-Core

Bro is single-threaded

- Backends have multiple cores, which are mostly idling.
- Work-around: “Cluster in a box”

We *really* want multi-threading.

- Must scale well with increasing numbers of cores.
- Must be transparent to the operator.

For some IDS, that’s not so hard.

- For others, it is ...

Research: Multi-Threaded DPI

Research: Multi-Threaded DPI

Traffic is “almost” embarrassingly parallel.

- Most activity is independent.

Research: Multi-Threaded DPI

Traffic is “almost” embarrassingly parallel.

Most activity is independent.

Analysis can be structured around *units.*

Simulations predict excellent scalability.

Research: Multi-Threaded DPI

Traffic is “almost” embarrassingly parallel.

- Most activity is independent.

Analysis can be structured around *units.*

- Simulations predict excellent scalability.

Incorporate architecture-level properties.

- Memory hierarchy, non-standard CPUs / bus systems.

Working Together
Collective Intelligence

REN-ISAC's Security Event System

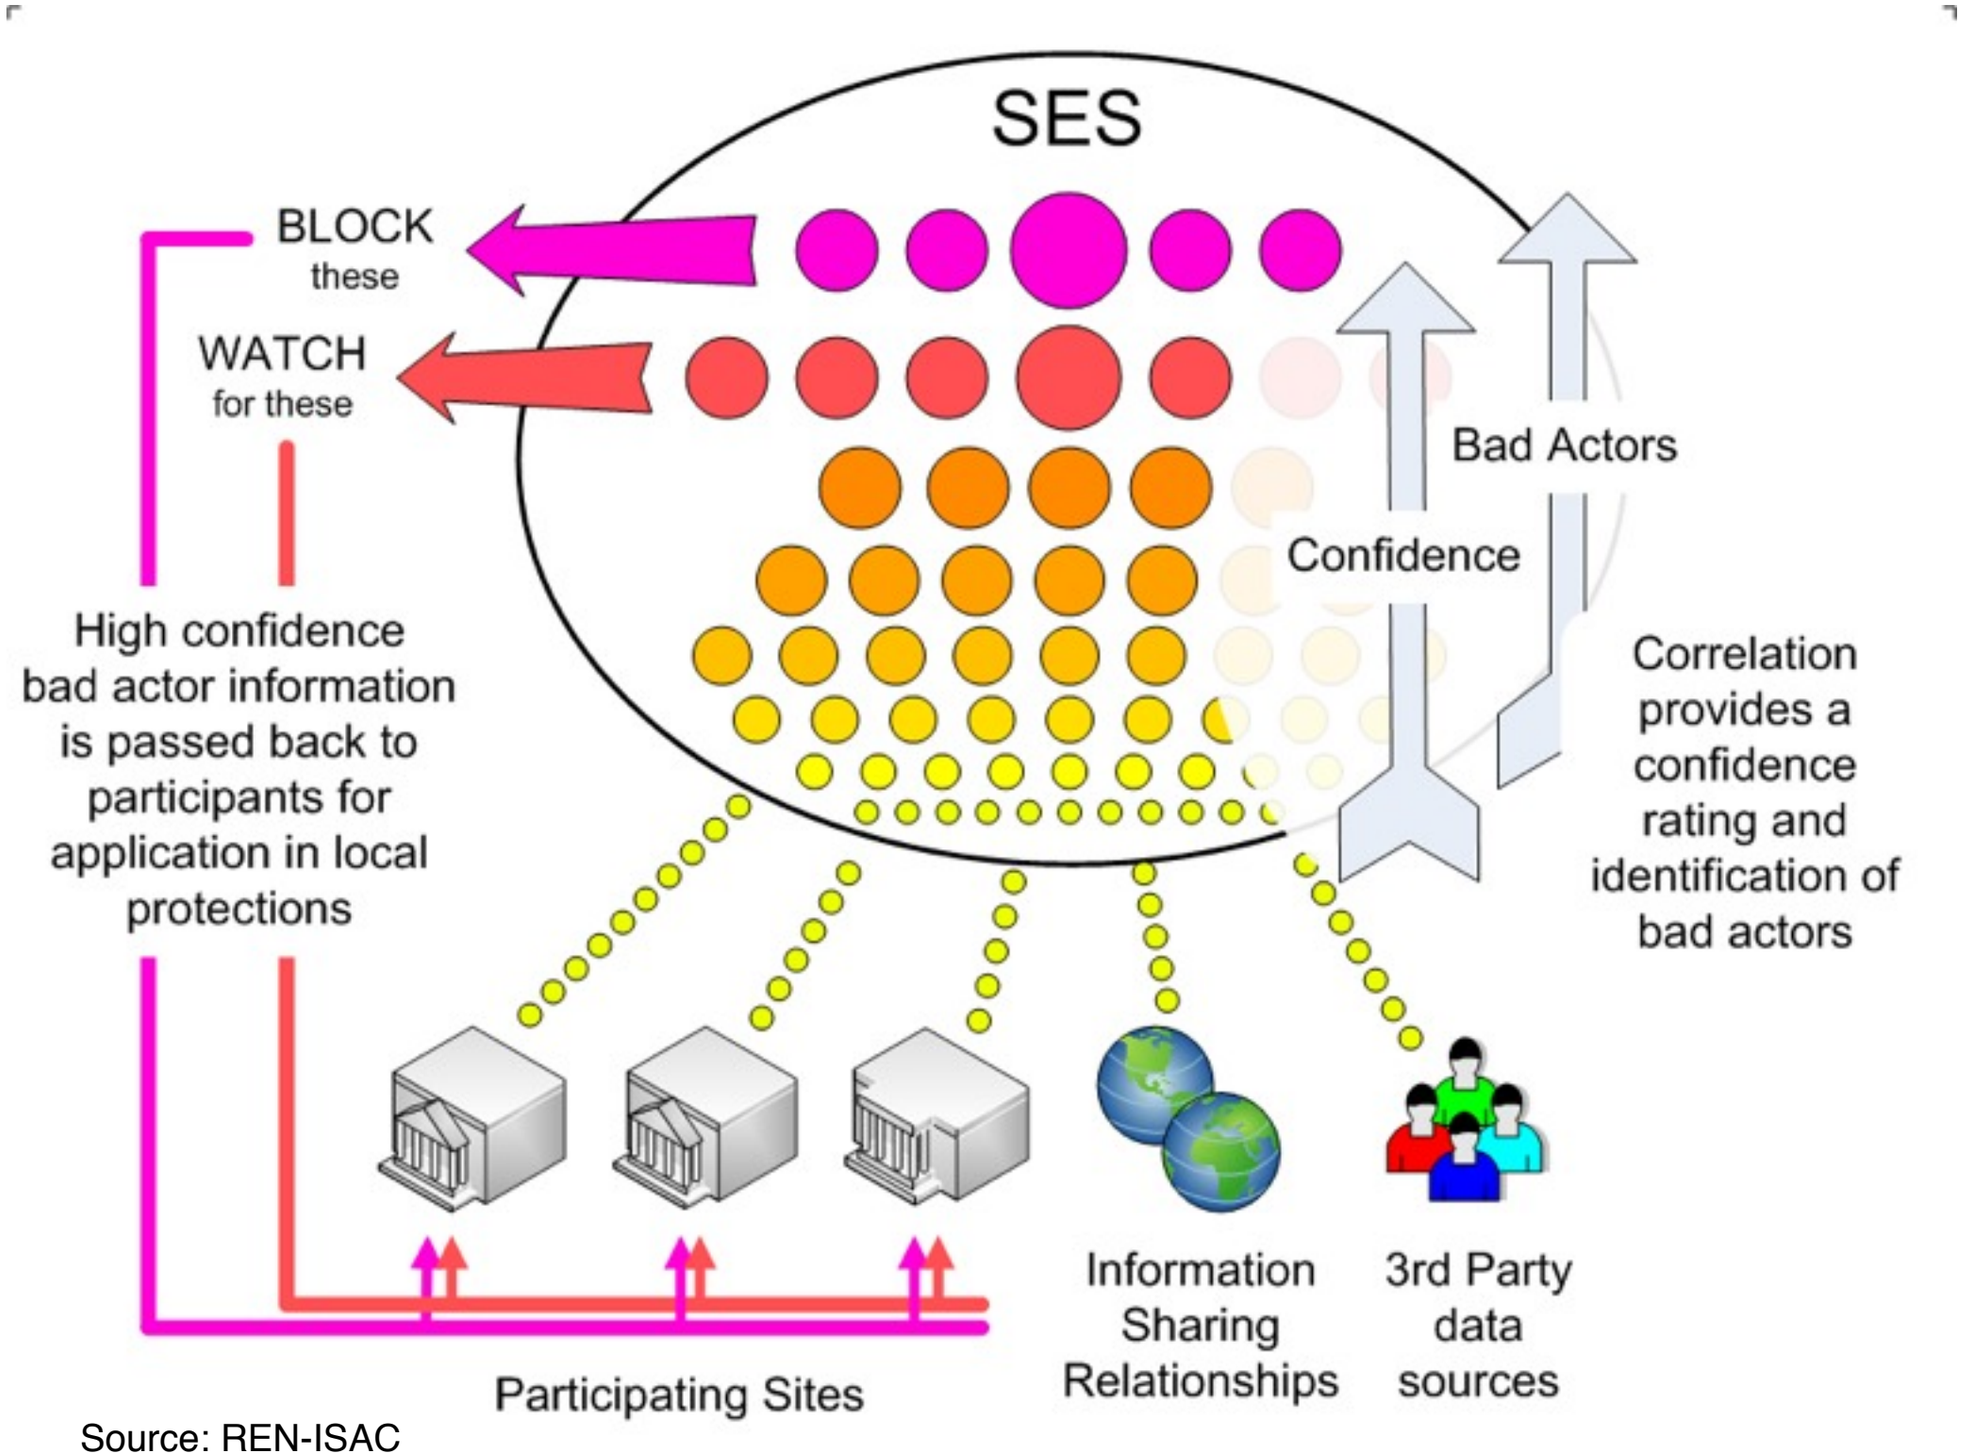

Source: REN-ISAC

REN-ISAC's Security Event System

Example: Basic blacklist

IP	Reason	Timestamp
66.249.66.1	Connected to honeypot	1333252748
208.67.222.222	Too many DNS requests	1330235733
192.150.186.11	Sent spam	1333145108

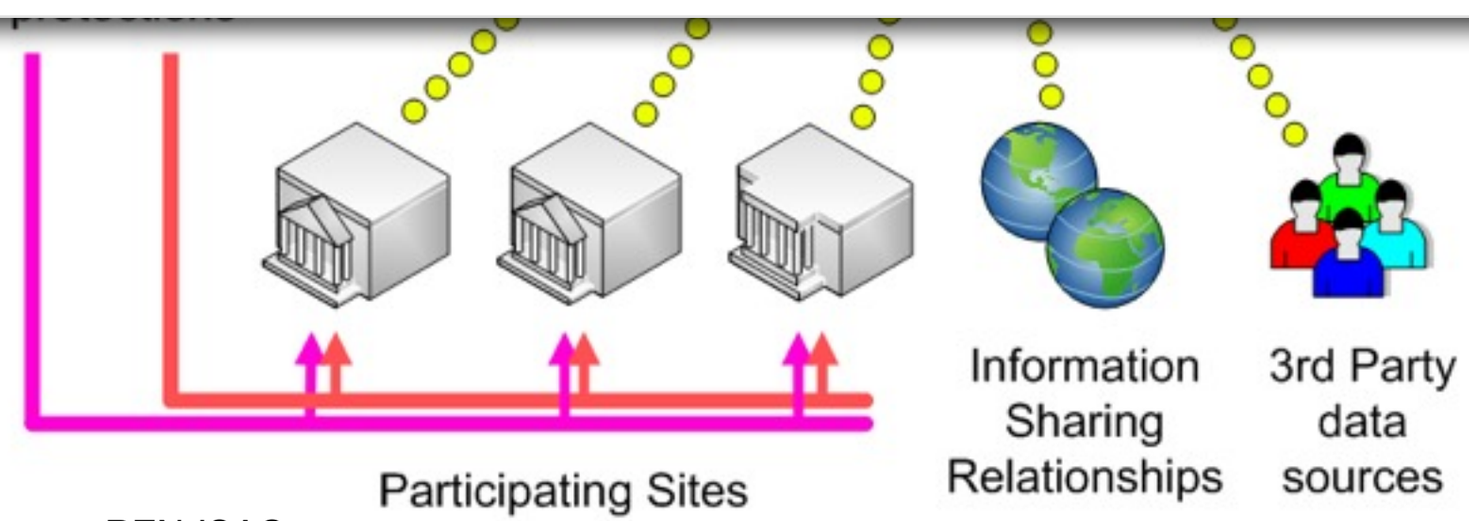

Source: REN-ISAC

How to Leverage Intelligence?

How to Leverage Intelligence?

IDS State

	Network State		Configuration
Volume	High		Low
Update frequency	High		Static
Examples	*Connection State*		*IDS rules*

How to Leverage Intelligence?

IDS State

	Network State	Intelligence	Configuration
Volume	High	Low/Medium	Low
Update frequency	High	Low/Medium	Static
Examples	*Connection State*	*Blacklists, Network Configuration*	*IDS rules*

Getting Intelligence Into Bro

Bernhard Amann, Robin Sommer, Aashish Sharma, Seth Hall
A Lone Wolf No More: Supporting Network Intrusion Detection with Real-Time Intelligence
Proc. Symposium on Research in Attacks, Intrusions and Defenses (RAID), September 2012

Getting Intelligence Into Bro

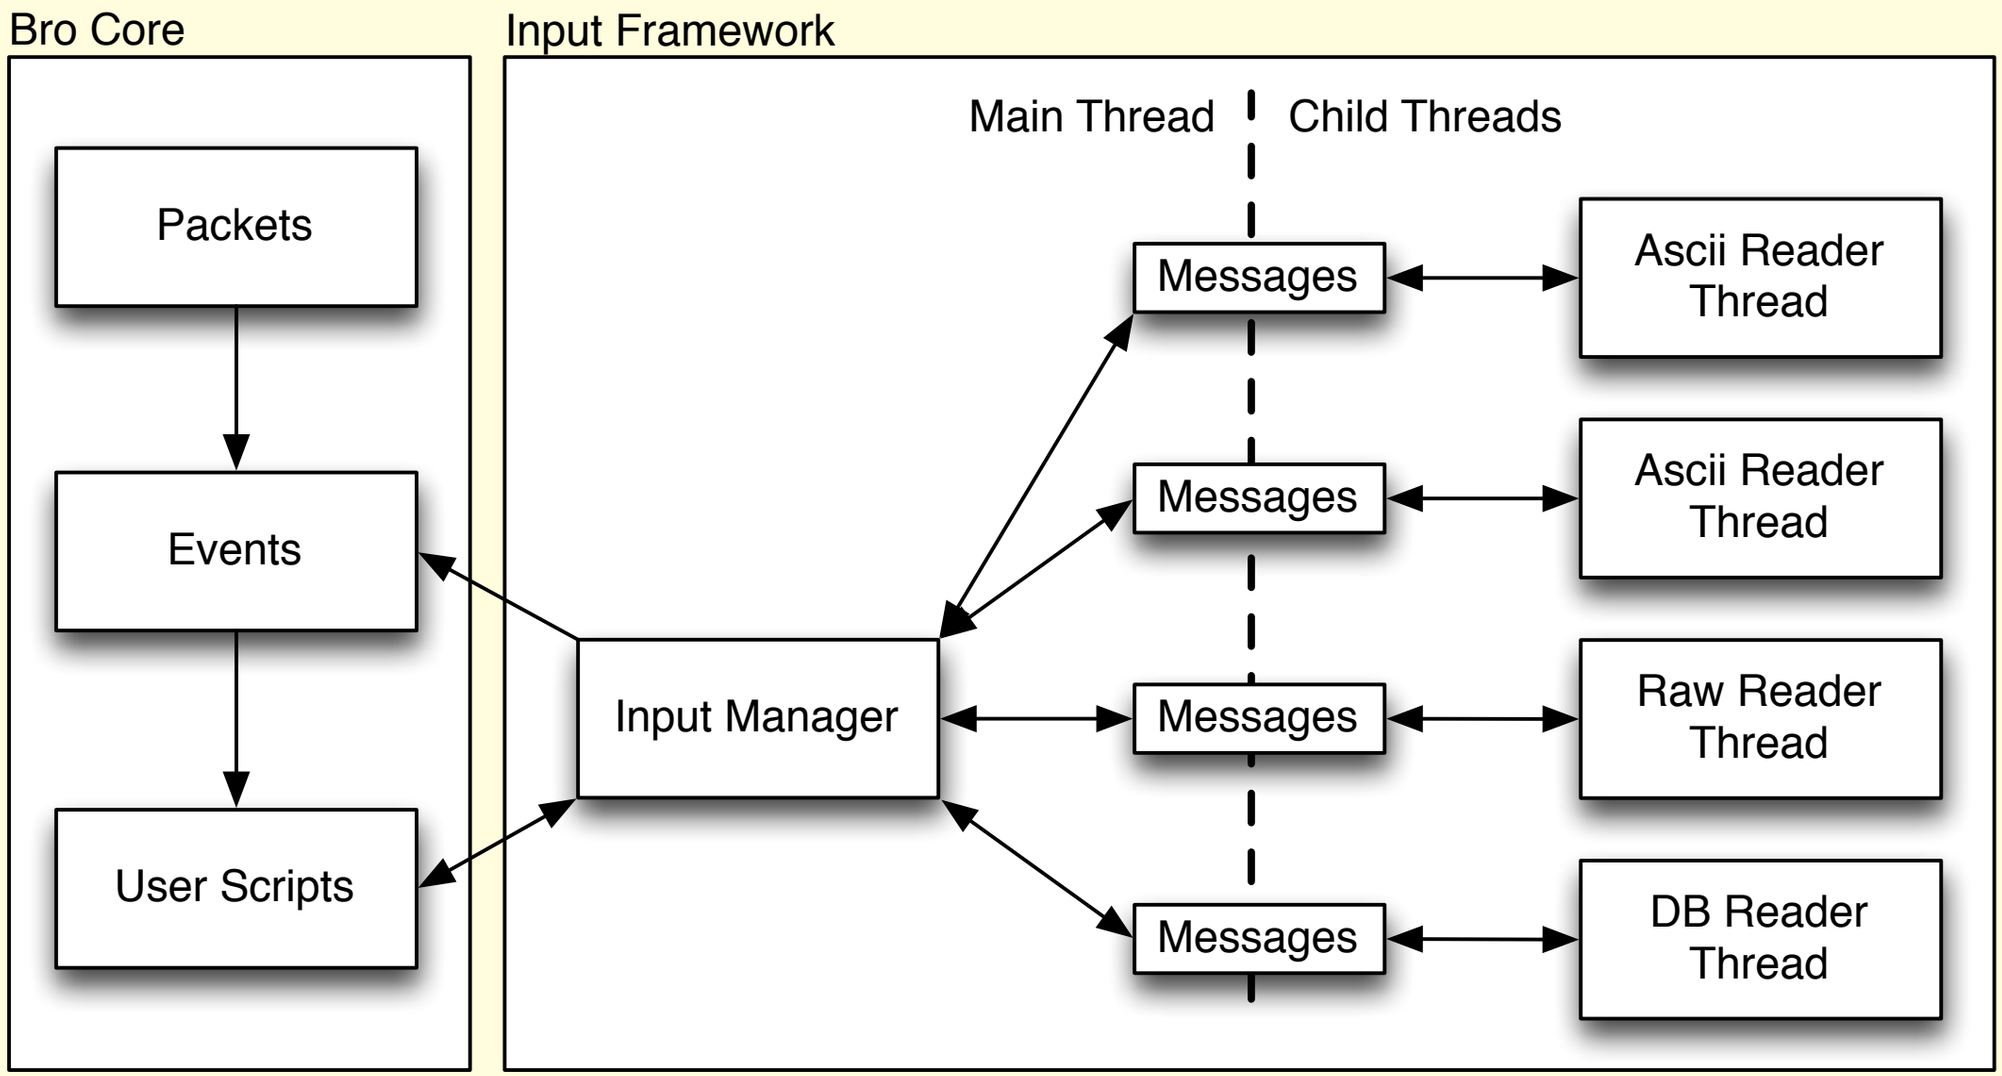

Bernhard Amann, Robin Sommer, Aashish Sharma, Seth Hall
A Lone Wolf No More: Supporting Network Intrusion Detection with Real-Time Intelligence
Proc. Symposium on Research in Attacks, Intrusions and Defenses (RAID), September 2012

Getting Intelligence Out of Bro

SSL Notary: Independent Perspective on Certificates

Getting Intelligence Out of Bro

SSL Notary: Independent Perspective on Certificates

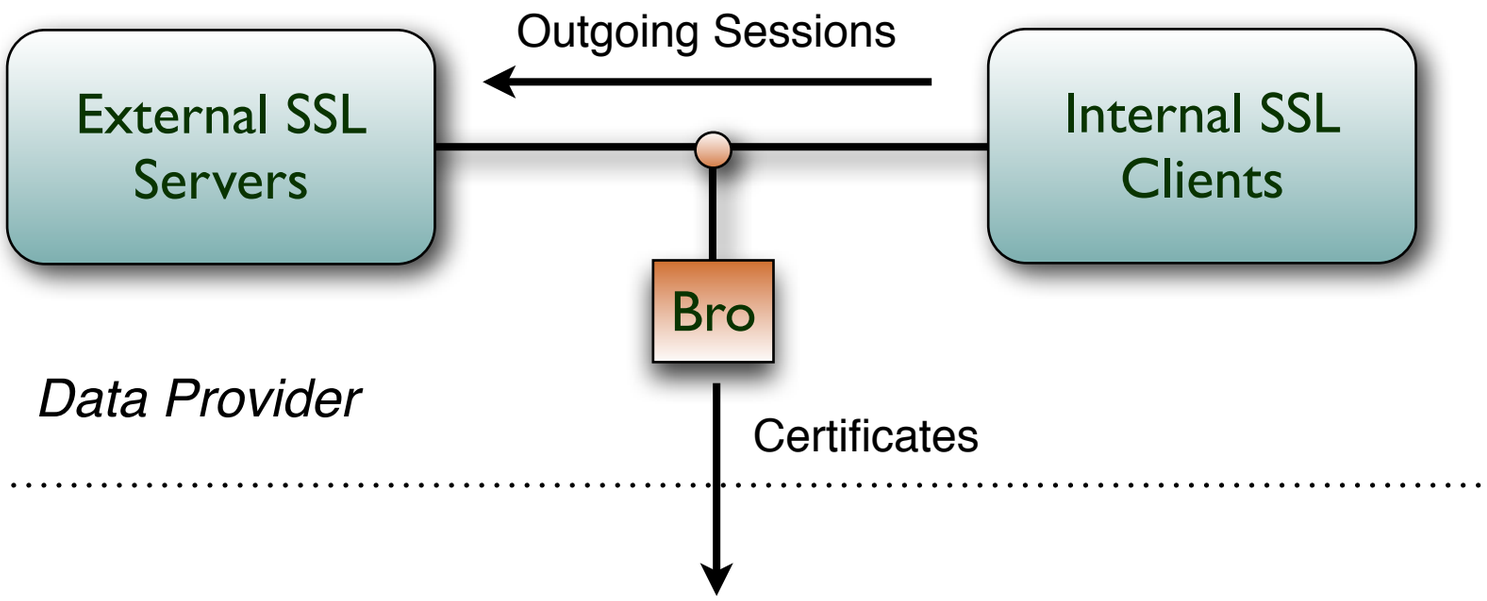

Getting Intelligence Out of Bro

SSL Notary: Independent Perspective on Certificates

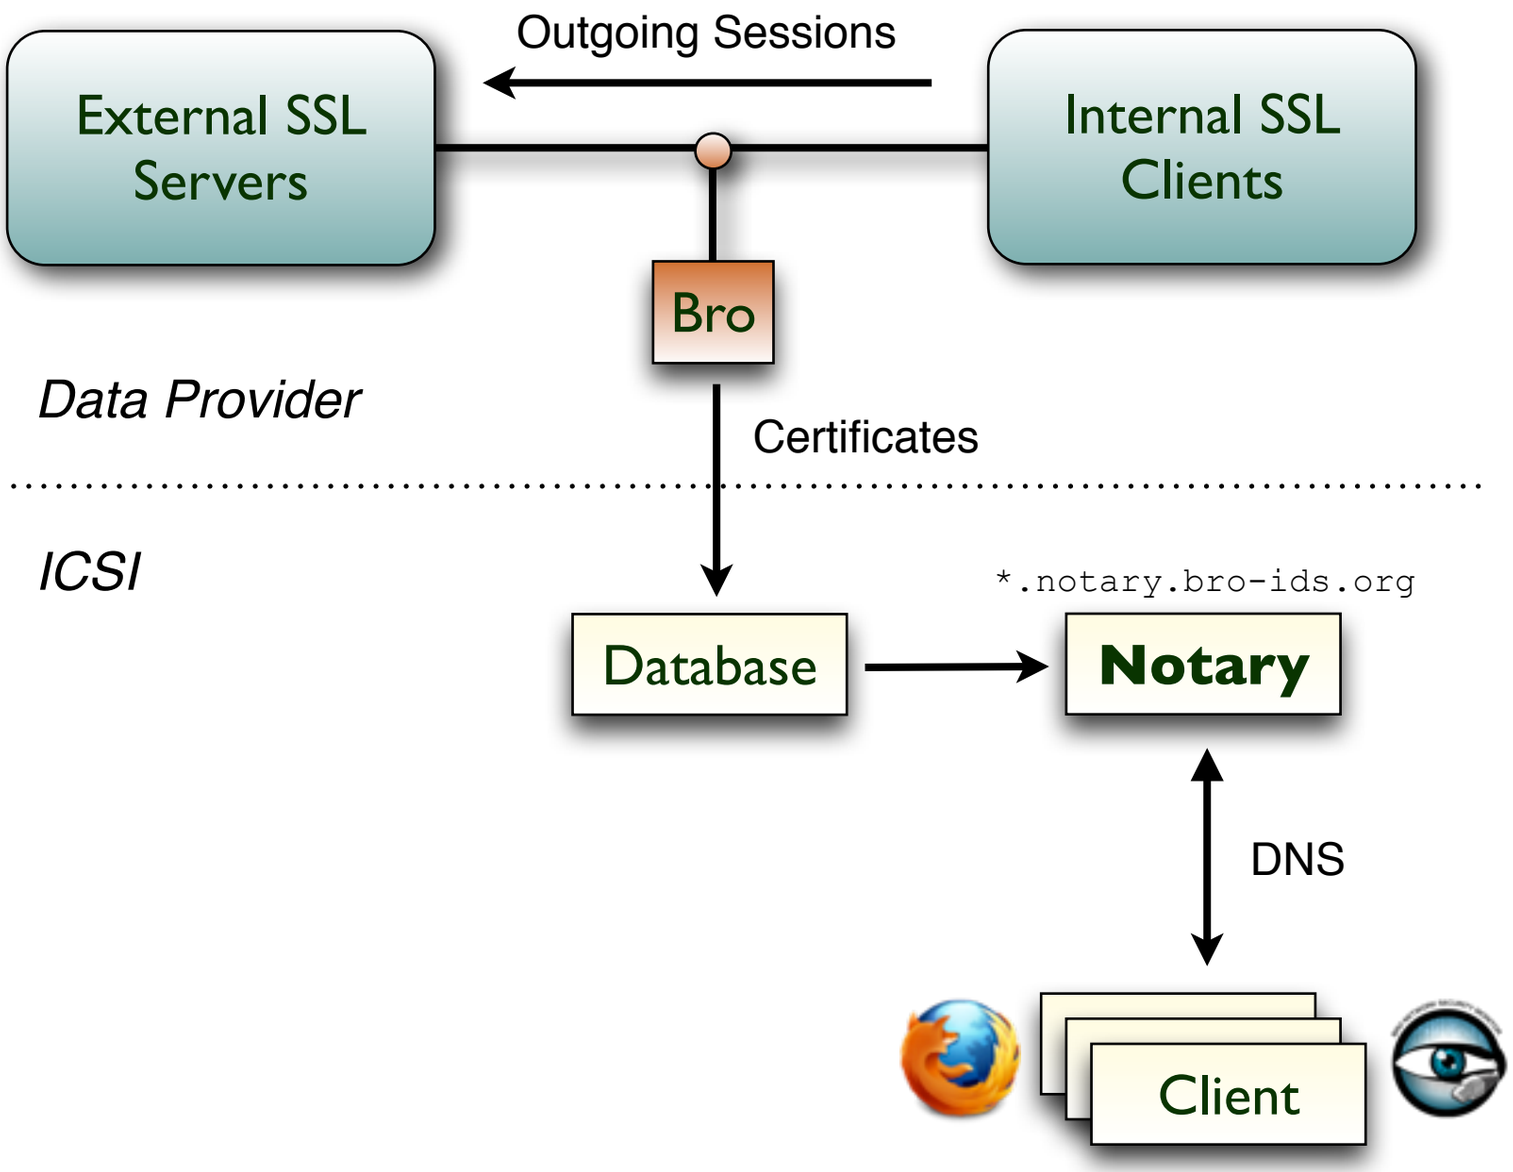

Notary: Data Providers

Notary: Data Providers

Site	Users	Certificates Total	Certificates Notary	Sessions	Duration (days)
University 1	60000	17.6M	222K	2.8B	162
University 2	50000	328K	185K	2.4B	170
University 3	3000	13K	9K	13M	138
University 4	90000	19K	17K	13M	3
Research Lab 1	250	155K	22K	40M	191
Research Lab 2	4000	93K	64K	420M	170
Government Network	50000	92K	90K	250M	151
Total (unique)	**257250**	18M	**340K**	5.6B	

September, 2012

Notary: Data Providers

Site	Users	Certificates Total	Certificates Notary	Sessions	Duration (days)
University				2.8B	162
University 2				2.4B	170
University 3				13M	138
University 4				13M	3
Research Lab				40M	191
Research Lab				420M	170
Governmen Network				250M	151
Total (uni				5.6B	

Collected Features

- Server Certificate
- Available ciphers
- Client SSL Extensions
- Server SSL Extensions
- Hash(Client, Server)
- Hash(Client, SNI)
- Hash(Client Session ID)
- Hash(Server Session ID)
- Selected Cipher
- Server Name Indication
- Ticket Lifetime Hint
- Timestamp
- SSL Protocol Version

Using it for Measurements, too ...

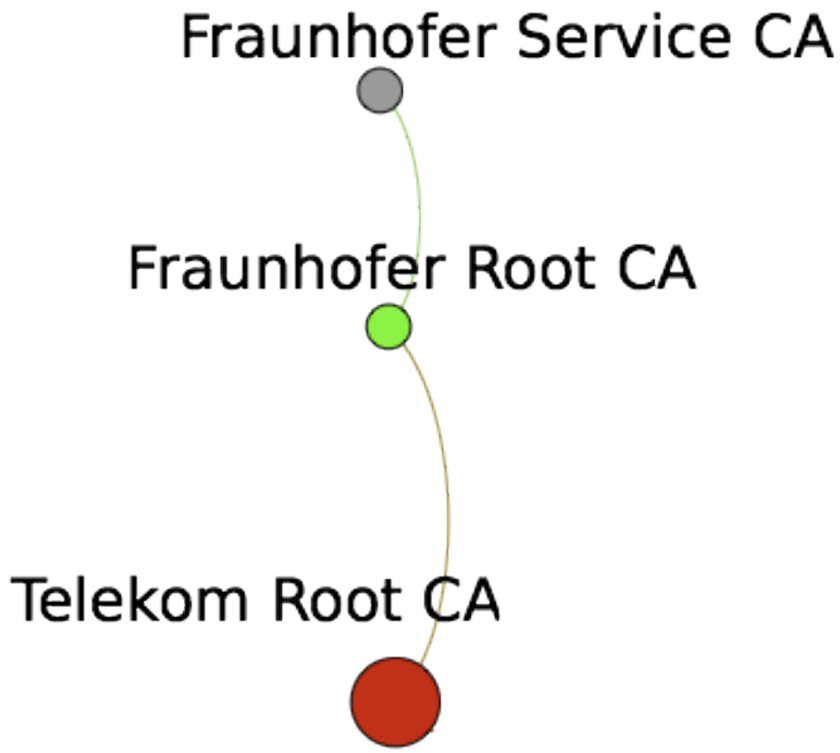

Using it for Measurements, too ...

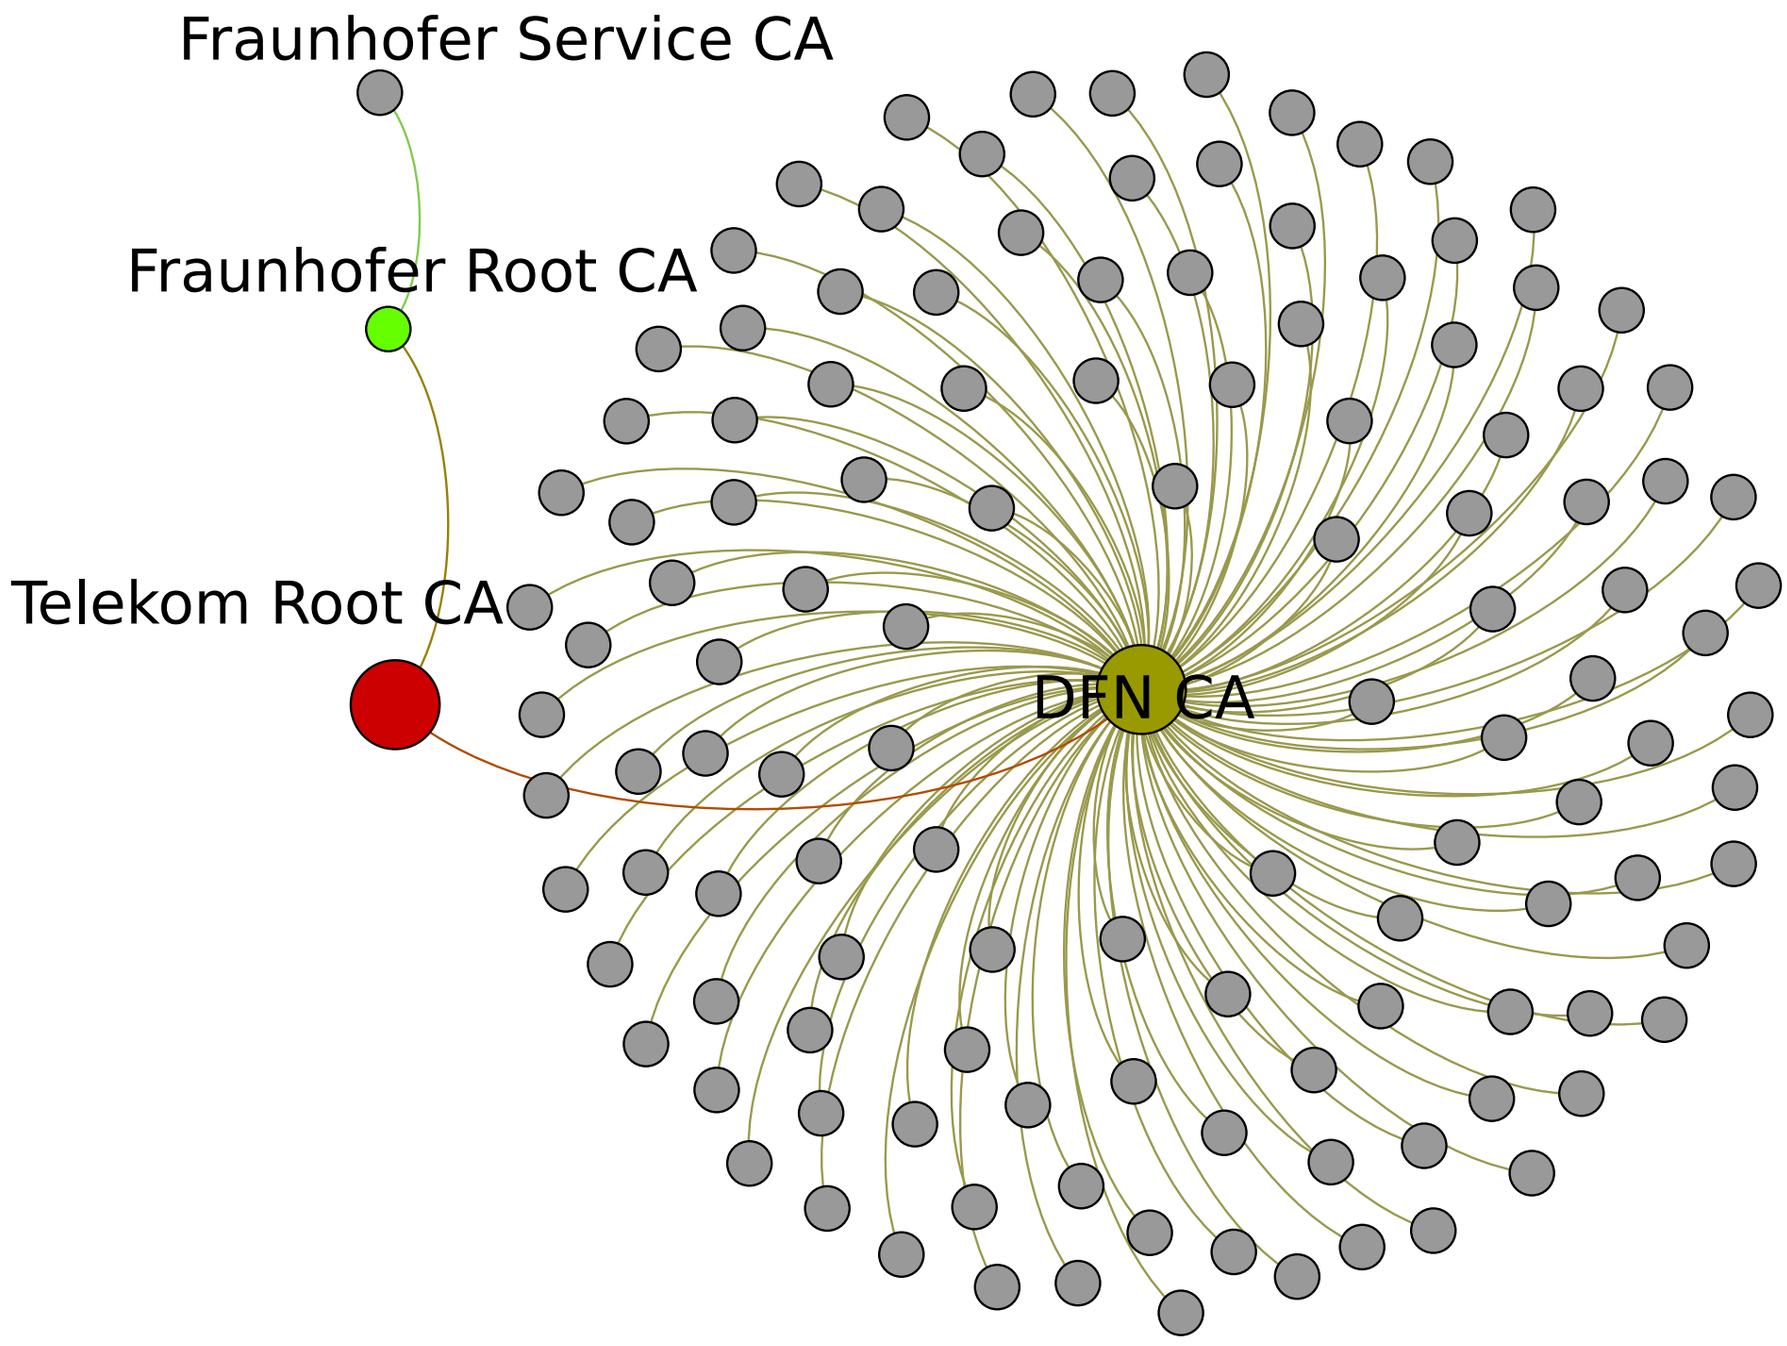

Summary

Summary

New Attack Trends.

Underground economy; targetted attacks.

Summary

New Attack Trends.

Underground economy; targetted attacks.

Bro.

From research to operations.

Summary

New Attack Trends.

Underground economy; targetted attacks.

Bro.

From research to operations.

Performance.

Scaling Bro Clusters to 100 Gbits/sec.

Summary

New Attack Trends.

Underground economy; targetted attacks.

Bro.

From research to operations.

Performance.

Scaling Bro Clusters to 100 Gbits/sec.

Collective Intelligence.

Sharing information in real-time.

Thanks for your attention.

Robin Sommer

International Computer Science Institute, &
Lawrence Berkeley National Laboratory

robin@icsi.berkeley.edu
http://www.icir.org/robin