The Modern Cybersecurity Stack

Data-Driven Network Monitoring with Bro

Robin Sommer

Corelight, Inc. /

International Computer Science Institute /

Lawrence Berkeley National Lab

robin@icsi.berkeley.edu

https://www.icir.org/robin

INTERNATIONAL COMPUTER SCIENCE INSTITUTE

corelight

Network Security Monitoring with Bro

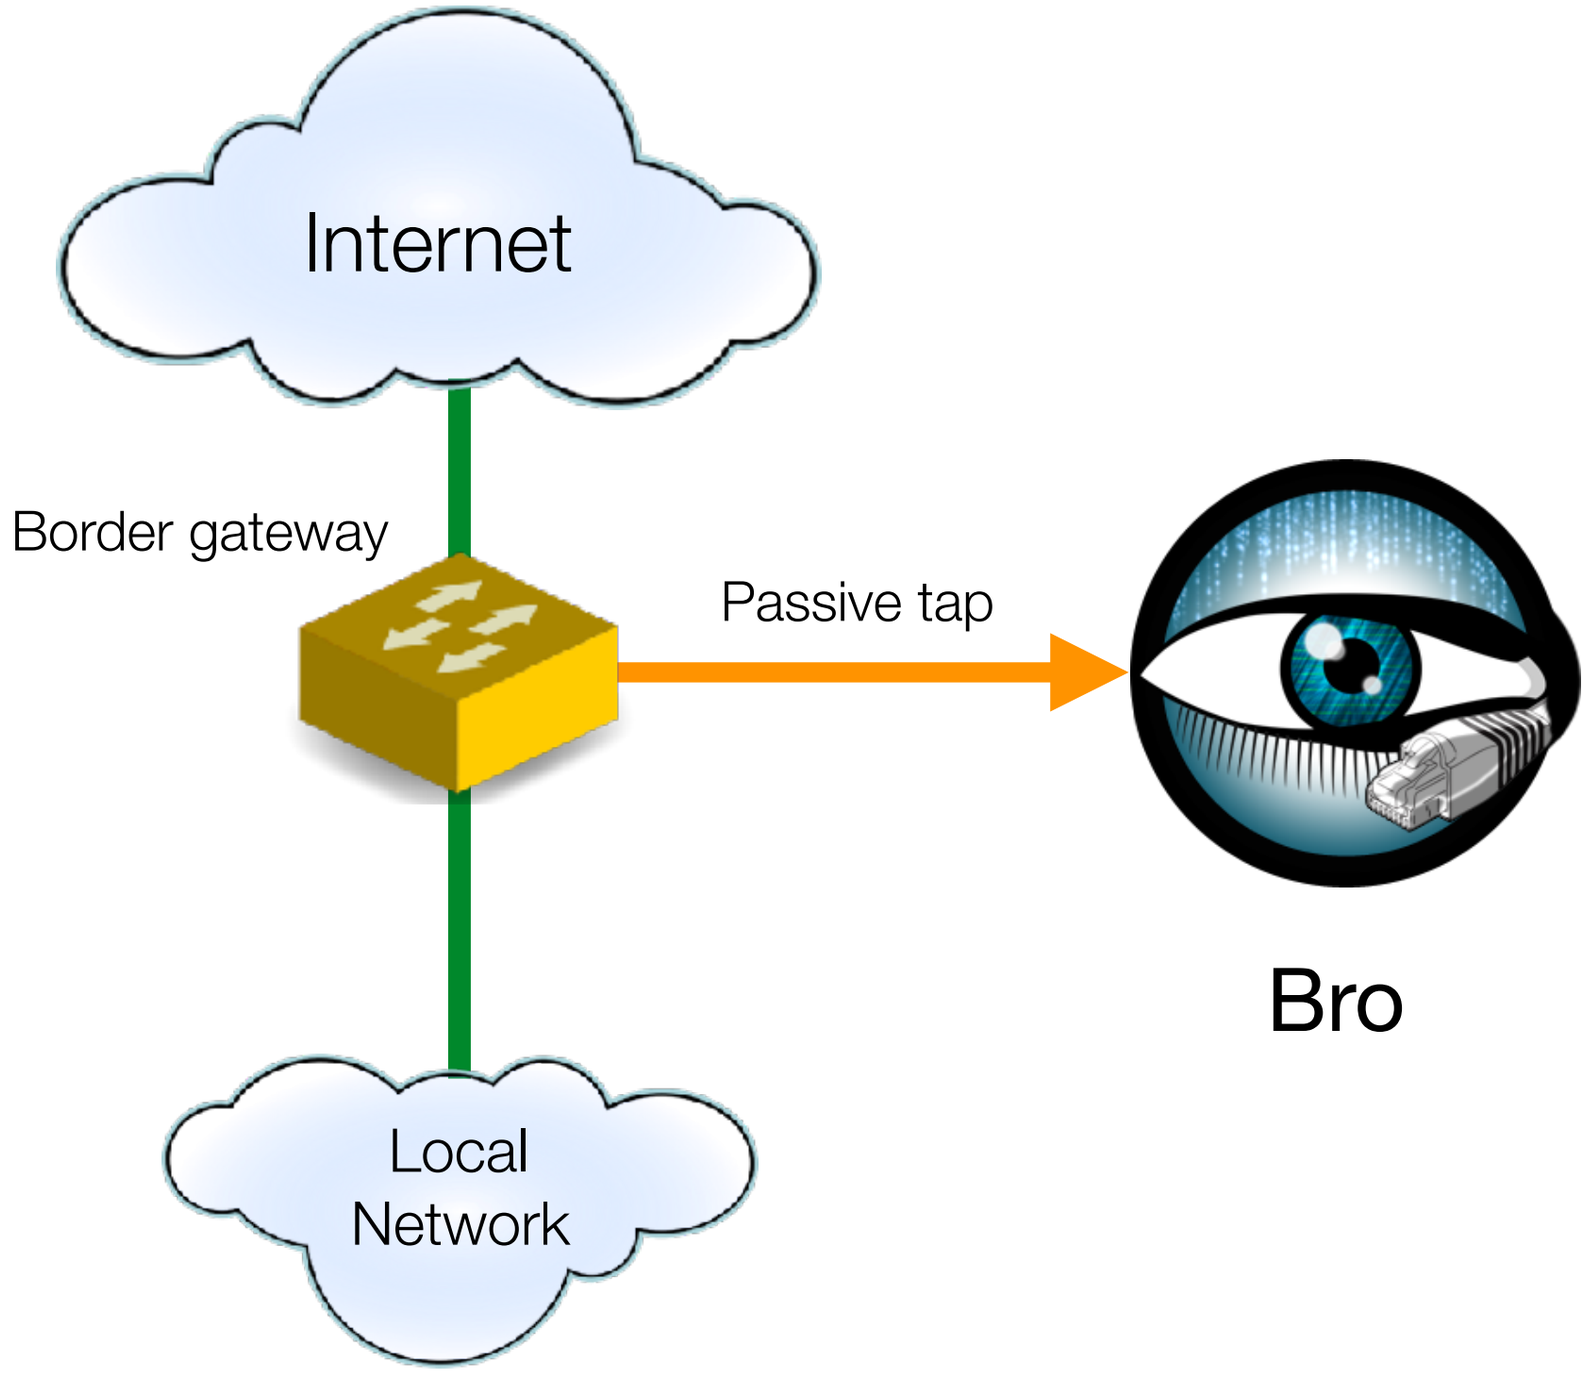

The Bro Platform

Open Source
BSD License

Analysis

Intrusion Detection | Vulnerabilit. Mgmt | File Analysis | Traffic Measure-ment | Traffic Control | Compliance Monitoring

Platform

Programming Language | Standard Library

Packet Processing

Tap

Network

Bro's been around for a while …

It took two decades for Bro to become one of the most popular open-source network security tools.

1995

Vern writes first line of code.

1998

Best Paper Award at USENIX Security

2006

1st Bro Workshop, SC in Tampa, FL

2016

BroCon '16, TACC, Austin, TX

Bro Today

"The best-kept secret in security"

Tremendous deployment base

Amazon, Facebook, GE, Mozilla, Salesforce, Target.
Department of Energy, Department of Defense, White House.
Most National Labs, many EDUs, many HPC facilities.

Bro skills in high demand

PepsiCo, Booz Allen Hamilton,
Radian, USAA, John Hopkins, BAE, Yahoo,
GDIT, Raytheon.
(Source: monster.com)

Community

180 attendees at BroCon'16
100 organizations at BroCon '16
6,500 Twitter followers
1,200 mailing list subscribers
1,800 stars on GitHub
Downloads from 150 countries

Industry funding

$350,000 in 2016

Recognition

InfoWorld Bossie Award
GitHub Security Showcase
Mozilla Open-Source Award
NSF Highlight to Congress 2016

Why has Bro become popular?

The legacy cyber security stack

Opaque, proprietary,

fueled by fear

The modern cyber security stack

Open-source, based on science,

fueled by data & analytics

Creating Visibility

Rich, structured, real-time data
for incident response, forensics, & analytics.

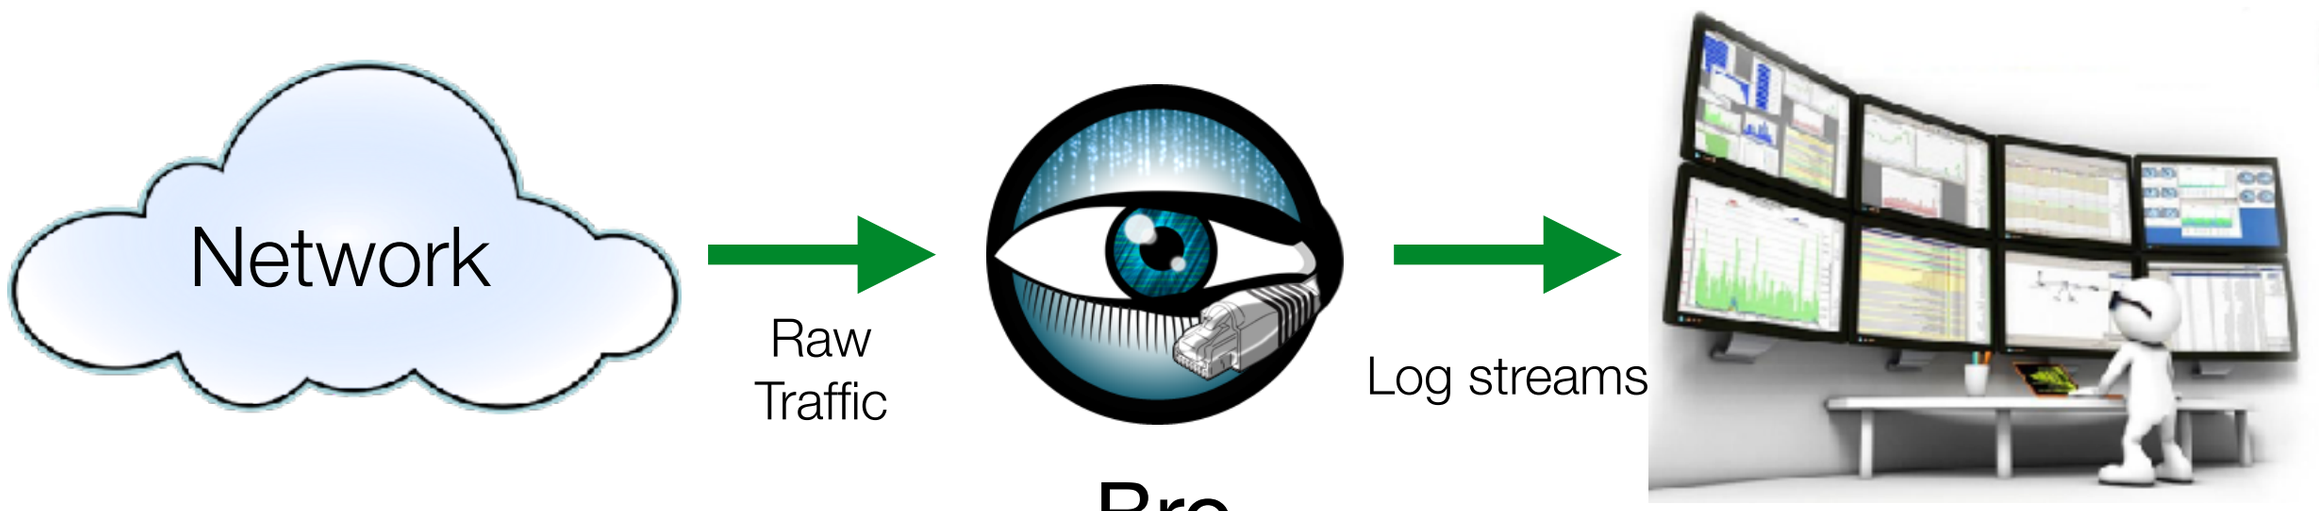

This data is what draws people to using Bro.

They have the analytics tools already, but they need high-quality input.

Connection Logs

conn.log

ts	1393099415.790834	Timestamp
uid	CSoqsg12YRTsWjYbZc	Unique ID
id.orig_h	2004:b9e5:6596:9876:[…]	Originator IP
id.orig_p	59258	Originator Port
id.resp_h	2b02:178:2fde:bff:[…]	Responder IP
id.resp_p	80	Responder Port
proto	tcp	IP Protocol
service	http	App-layer Protocol
duration	2.105488	Duration
orig_bytes	416	Bytes by Originator
resp_bytes	858	Bytes by Responder
conn_state	SF	TCP state
local_orig	F	Local Originator?
missed_bytes	0	Gaps
history	ShADafF	State History
tunnel_parents	Cneap78AnVWoA1yml	Outer Tunnel Connection

Understand Your Network (1)

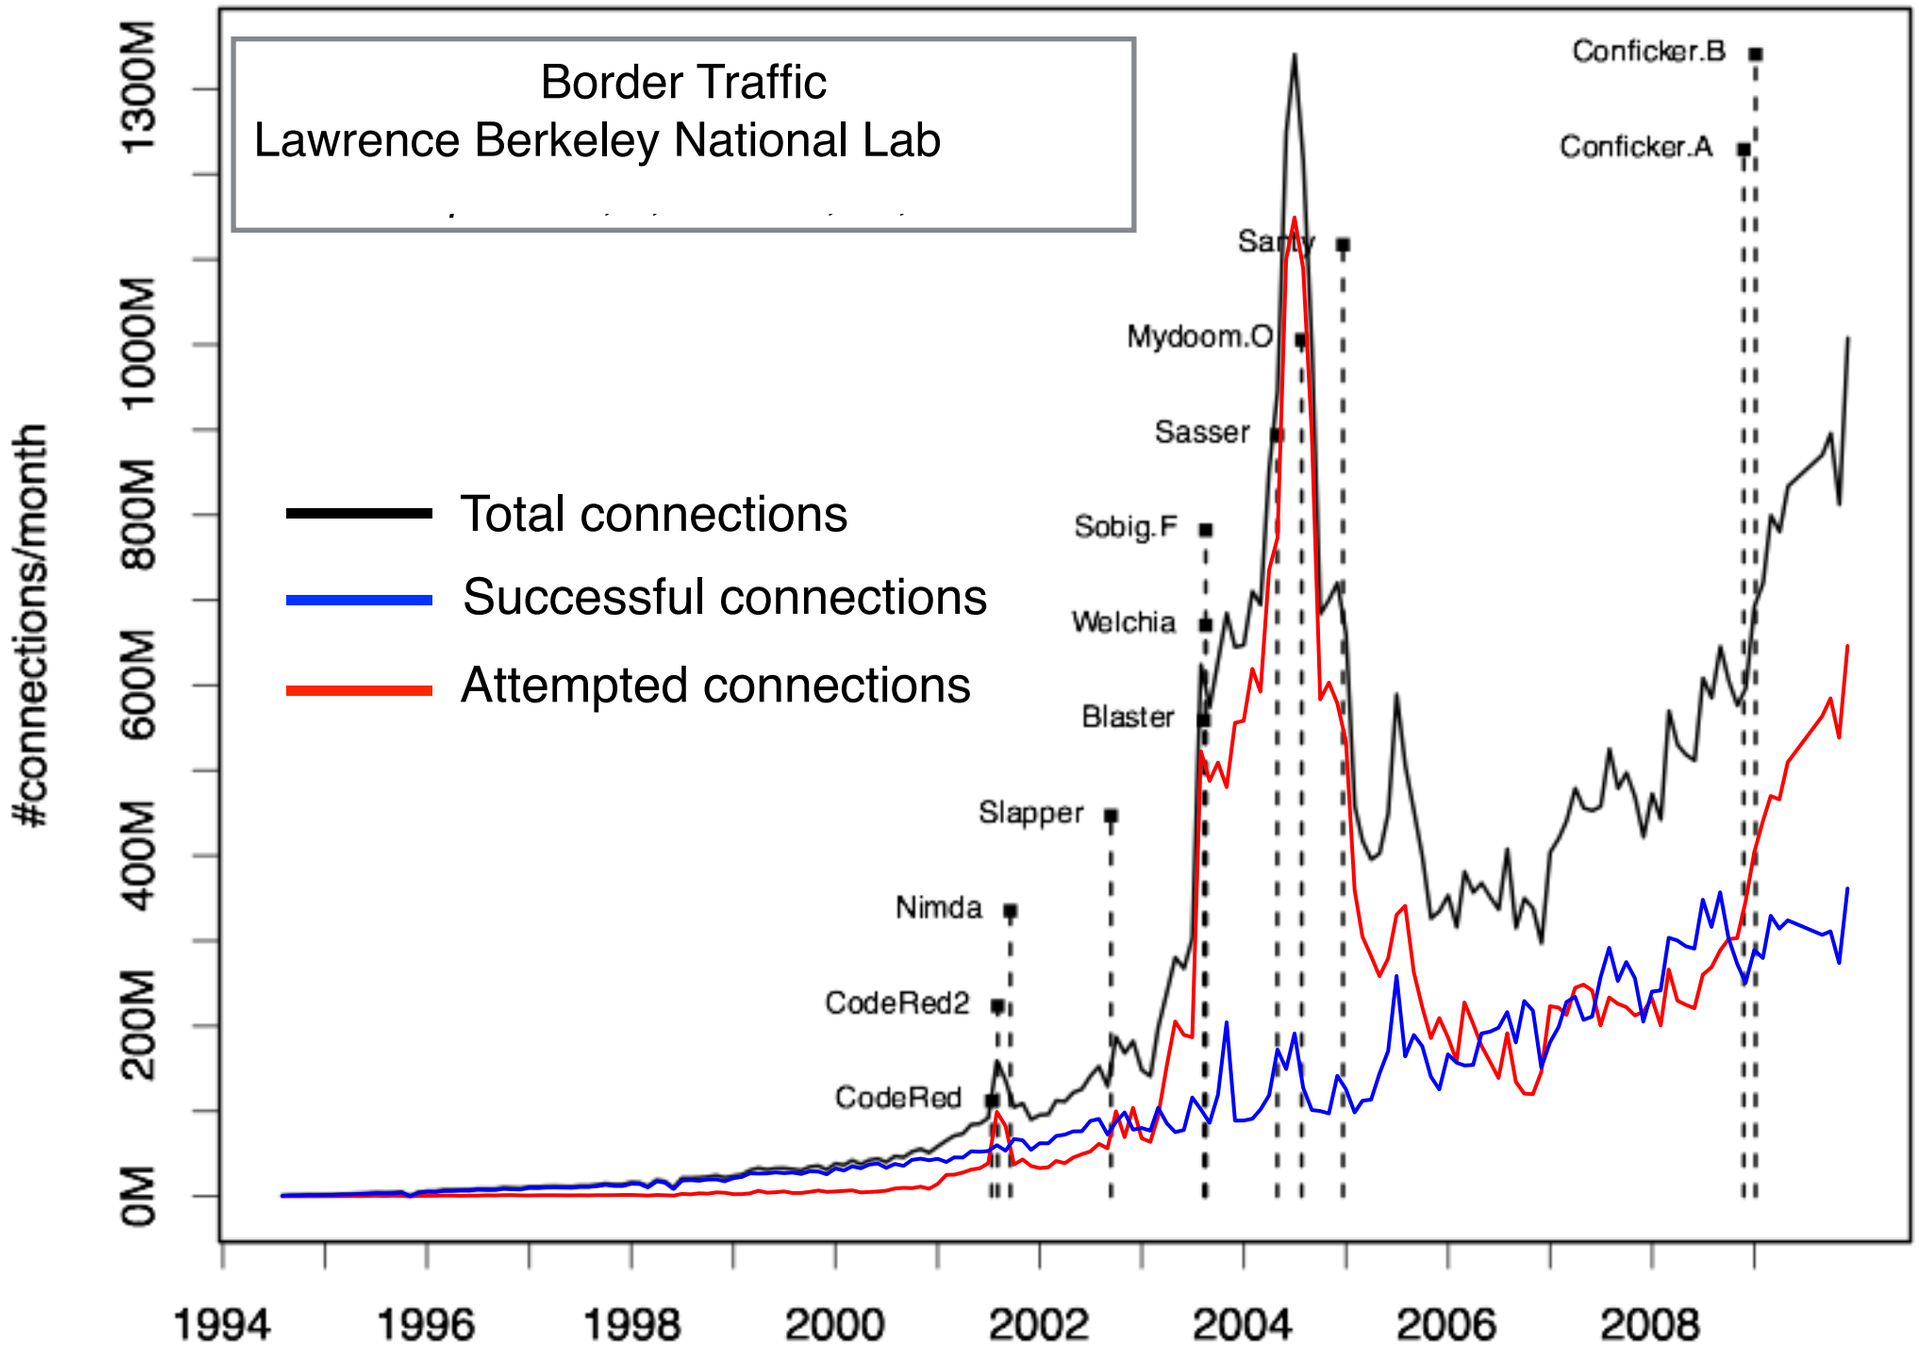

HTTP

http.log

ts	1393099291.589208
uid	CKFUW73bIADw0r9pl
id.orig_h	2a07:f2c0:90:402:41e:c13:6cb:99c
id.orig_p	54352
id.resp_h	2406:fe60:f47::aaeb:98c
id.resp_p	80
method	POST
host	com-services.pandonetworks.com
uri	/soapservices/services/SessionStart
referrer	-
user_agent	Mozilla/4.0 (Windows; U) Pando/2.6.0.8
status_code	200
username	anonymous
password	-
orig_mime_types	application/xml
resp_mime_types	application/xml

Understand Your Network (2)

Top HTTP servers by IP addresses vs host headers

```
a198-189-255-200.deploy.akamaitechnolgies.com
a198-189-255-216.deploy.akamaitechnolgies.com
a198-189-255-217.deploy.akamaitechnolgies.com
a198-189-255-230.deploy.akamaitechnolgies.com
a198-189-255-225.deploy.akamaitechnolgies.com
a198-189-255-206.deploy.akamaitechnolgies.com
a198-189-255-201.deploy.akamaitechnolgies.com
a198-189-255-223.deploy.akamaitechnolgies.com
                72.21.91.19
a198-189-255-208.deploy.akamaitechnolgies.com
a198-189-255-207.deploy.akamaitechnolgies.com
          nuq04s07-in-f27.1e100.net
 a184-28-157-55.deploy.akamaitechnologies.com
a198-189-255-224.deploy.akamaitechnolgies.com
a198-189-255-209.deploy.akamaitechnolgies.com
a198-189-255-222.deploy.akamaitechnolgies.com
a198-189-255-214.deploy.akamaitechnolgies.com
          nuq04s06-in-f27.1e100.net
       upload-lb.pmtpa.wikimedia.org
          nuq04s08-in-f27.1e100.net
```

```
          ad.doubleclick.net
          ad.yieldmanager.com
       b.scorecardresearch.com
          clients1.google.com
    googleads.g.doubleclick.net
       graphics8.nytimes.com
             l.yimg.com
liveupdate.symanteclive update.com
          mt0.google.com
       pixel.quantserve.com
       platform.twitter.com
       profile.ak.fbcdn.net
             s0.2mdn.net
  safebrowsing-cache.google.com
       static.ak.fbcdn.net
          swcdn.apple.com
       upload.wikimedia.org
          www.facebook.com
    www.google-analytics.com
          www.google.com
```

corelight

Software

software.log

ts	1392796839.675867
host	10.209.100.2
host_p	-
software_type	HTTP::BROWSER
name	DropboxDesktopClient
version.major	2
version.minor	4
version.minor2	11
version.minor3	-
version.addl	Windows
unparsed_version	DropboxDesktopClient/2.4.11 (Windows; 8; i32; en_US; Trooper 5694-2047-1832-6291-8315)

Understand Your Network (3)

Top Software by Number of Hosts

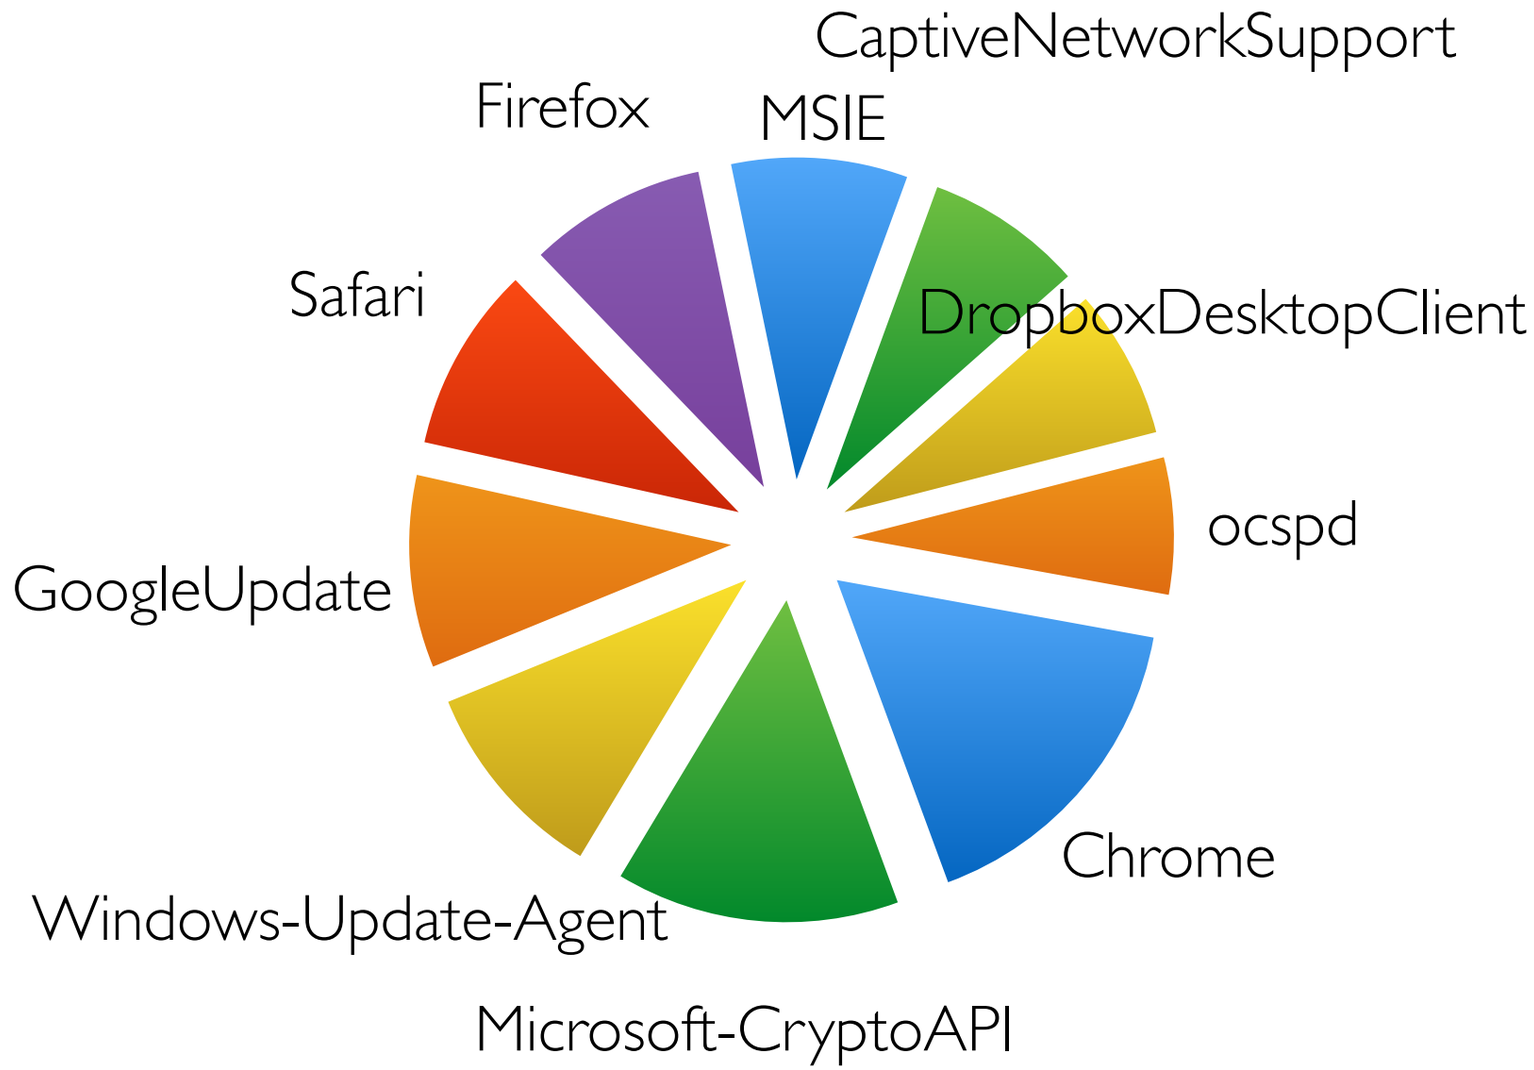

Files

files.log

ts	1392797643.447056
fuid	FnungQ3TI19GahPJP2
tx_hosts	191.168.187.33
rx_hosts	10.1.29.110
conn_uids	CbDgik2fjeKL5qzn55
source	SMTP
analyzers	SHA1,MD5
mime_type	application/x-dosexec
filename	Letter.exe
duration	5.320822
local_orig	T
seen_bytes	39508
md5	93f7f5e7a2096927e06e[…]1085bfcfb
sha1	daed94a5662a920041be[…]a433e501646ef6a03

Understand your Malware


```
# cat files.log | bro-cut mime_type sha1 | awk '$1 ~ /x-dosexec/‘
application/x-dosexec   5fd2f37735953427e2f6c593d6ec7ae882c9ab54
application/x-dosexec   00c69013d34601c2174b72c9249a0063959da93a
application/x-dosexec   0d801726d49377bfe989dcca7753a62549f1ddda
[…]
```

```
# dig +short 733a48a9cb4[…]2a91e8d00.malware.hash.cymru.com TXT
"1221154281 53"
```

SSL & X.509

ssl.log

ts	1392805957.927087
uid	CEA05l2D7k0BD9Dda2
id.orig_h	2a07:f2c0:90:402:41e:c13:6cb:99c
id.orig_p	40475
id.resp_h	2406:fe60:f47::aaeb:98c
id.resp_p	443
version	TLSv10
cipher	TLS_DHE_RSA_WITH_AES_256_CBC_SHA
server_name	www.netflix.com
subject	CN=www.netflix.com,OU=Operations, O=Netflix, Inc.,L=Los Gatos, ST=CALIFORNIA,C=US
issuer_subject	CN=VeriSign Class 3 Secure Server CA, OU=VeriSign Trust Network,O=VeriSign, C=US
not_valid_before	1389859200.000000
not_valid_after	1452931199.000000
client_subject	-
client_issuer_subject	-
cert_hash	197cab7c6c92a0b9ac5f37cfb0699268
validation_status	ok

Understand the (SSL) World

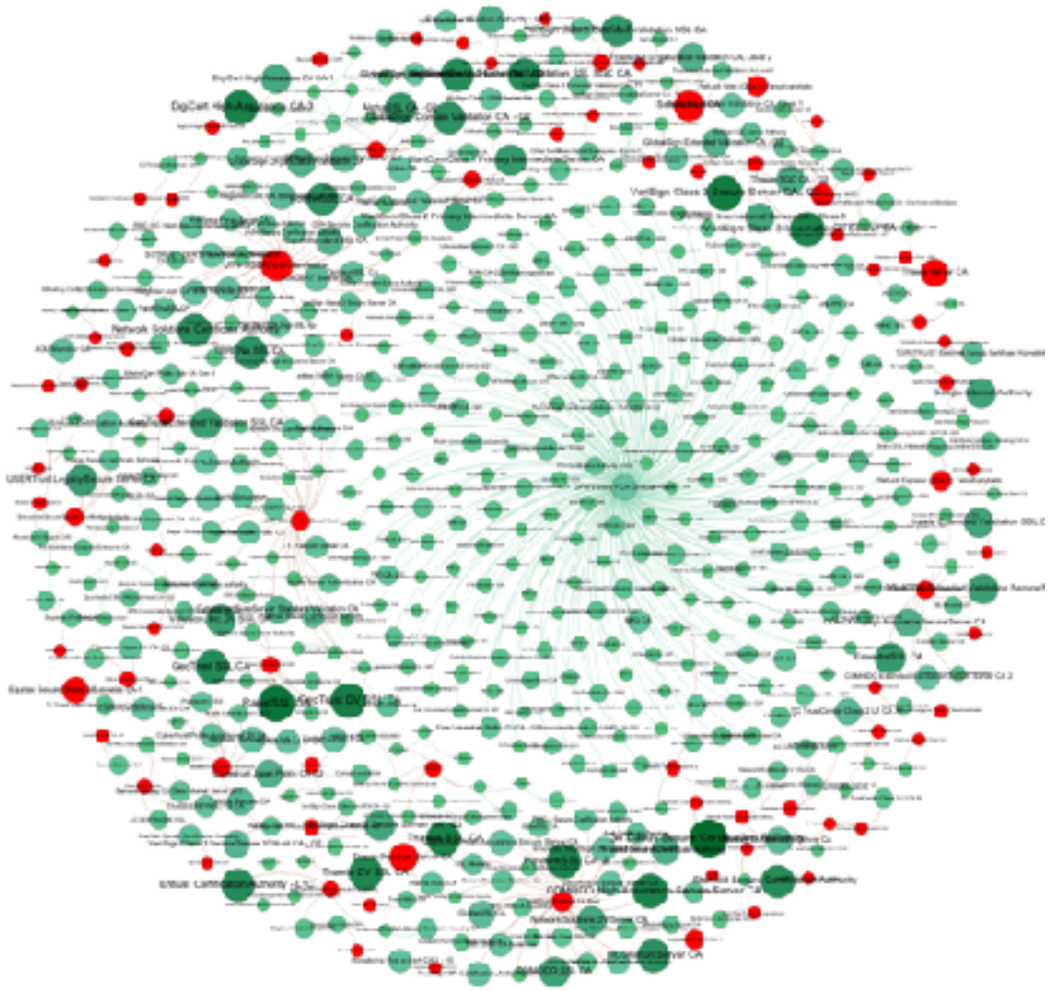

The ICSI Certificate Notary

Four years of passive data: 14M SSL certificates, 240B sessions

`https://notary.icsi.berkeley.edu`

All This Data is Invaluable For Incident Response

If you're compromised, you want to know:

What happened?

How did it happen?

Is anybody else affected?

Has it happened before?

How did a bunch of academics get there?

Bro History

1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015 2016 2017

About 20 academic publications presenting Bro-related research.

Initial Bro versions are addressing an operational need at LBNL	Basic research at ICSI drives continuous innovation

Feedback loop crucial for both sides

Operational deployment in large-scale open-science networks

Example: Processing performance

Back in the days …

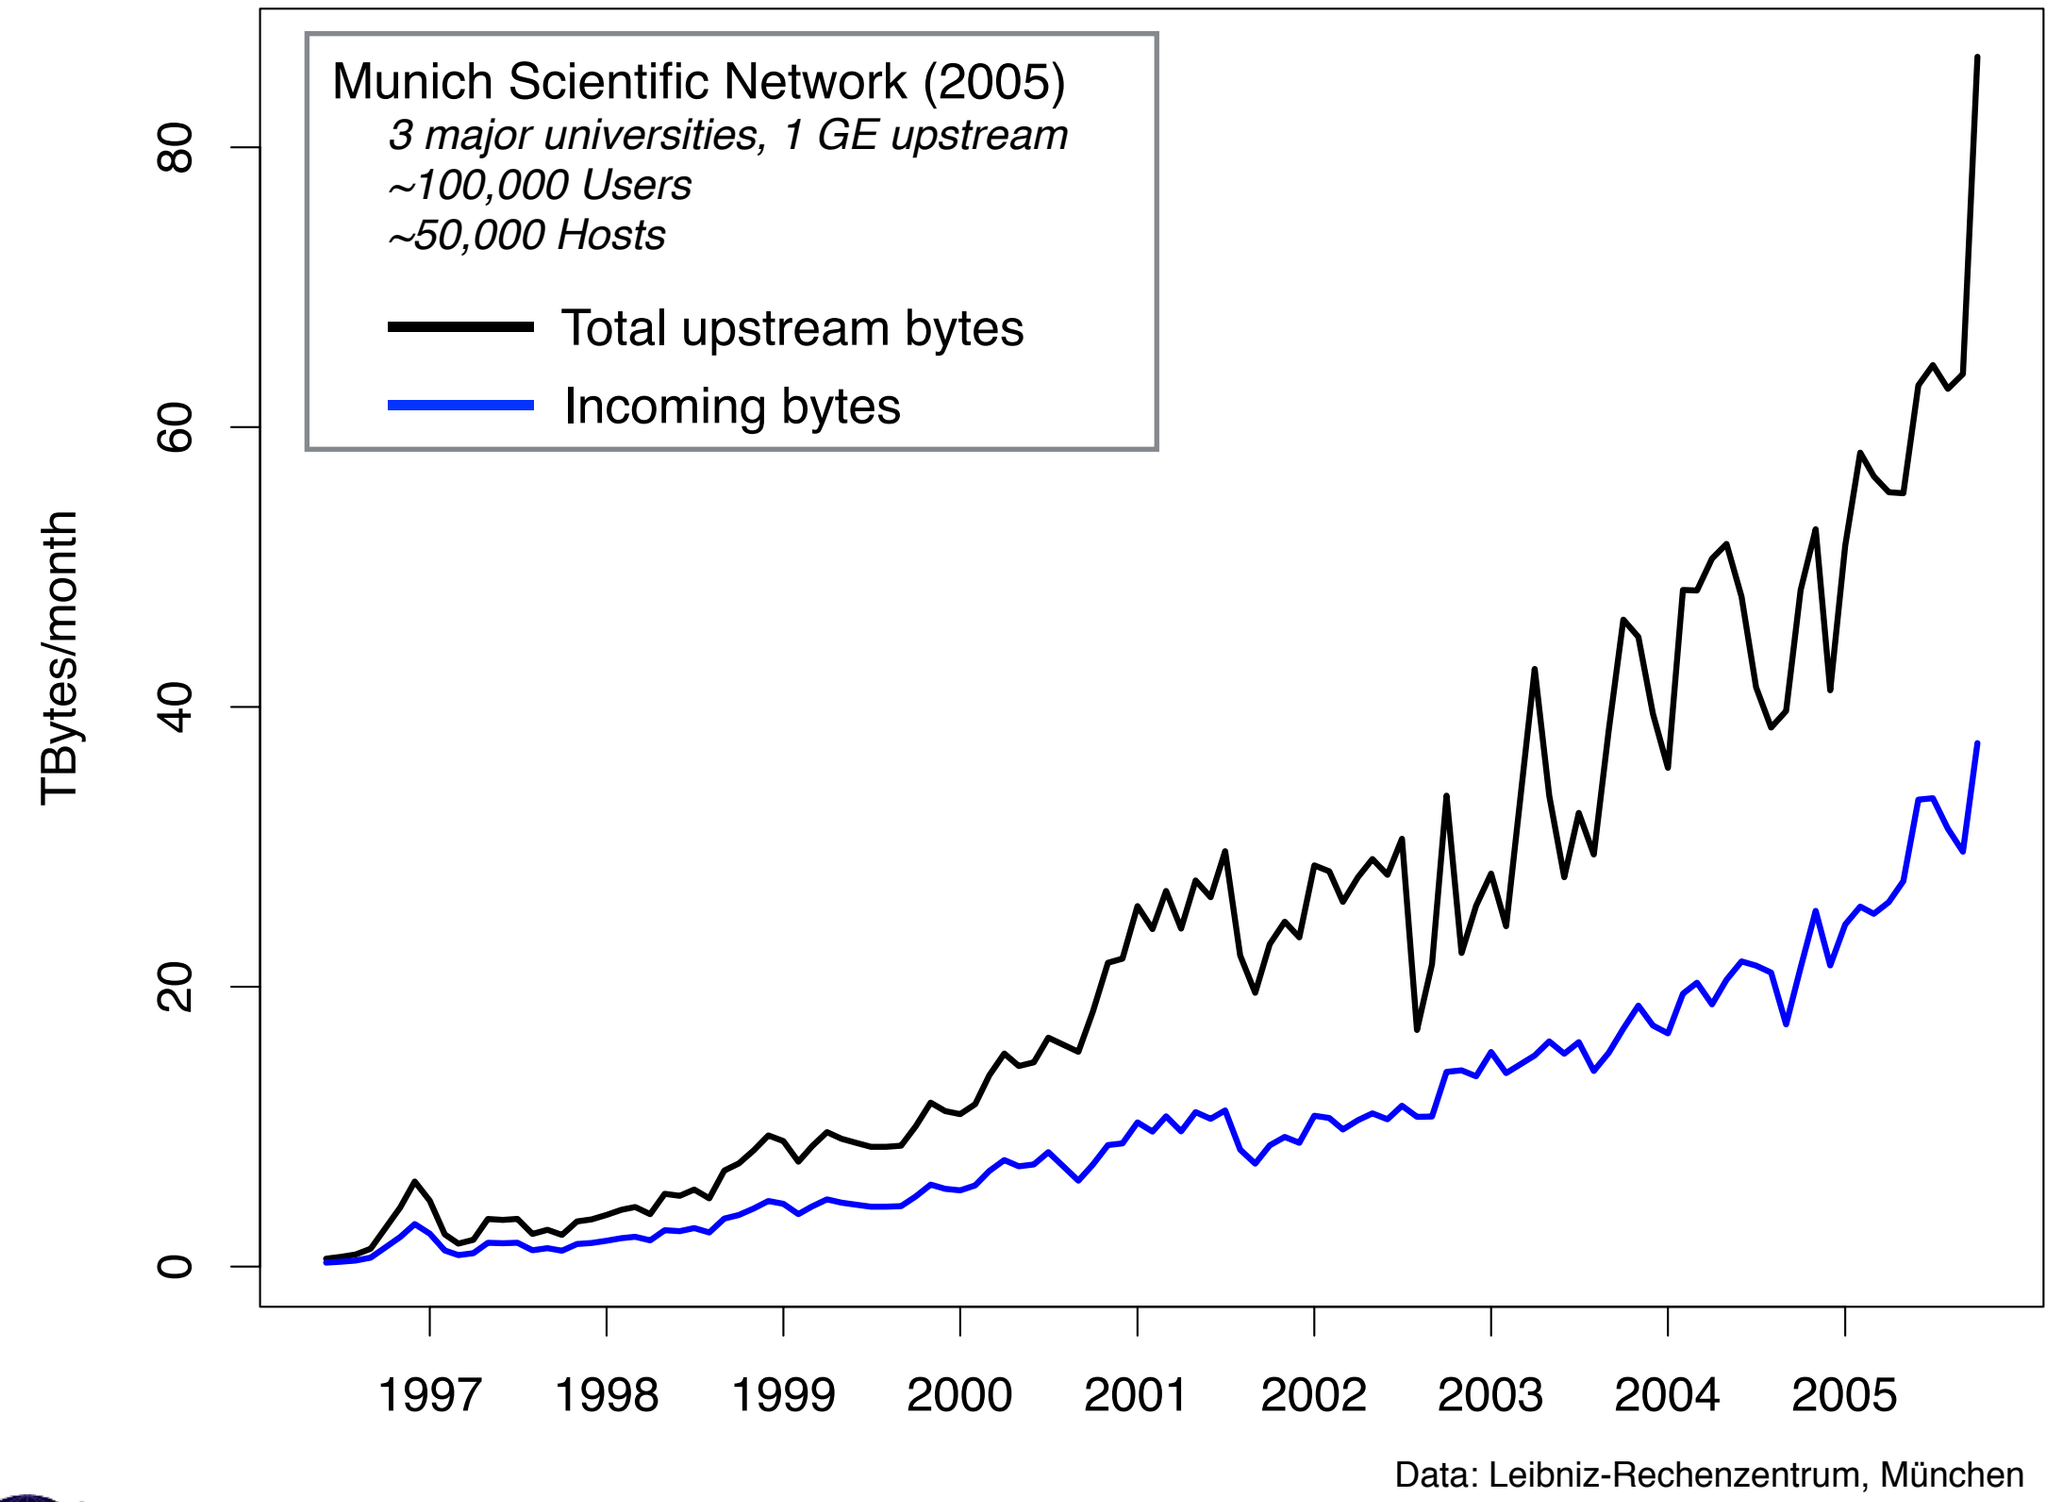

Data: Leibniz-Rechenzentrum, München

And in 2014 …

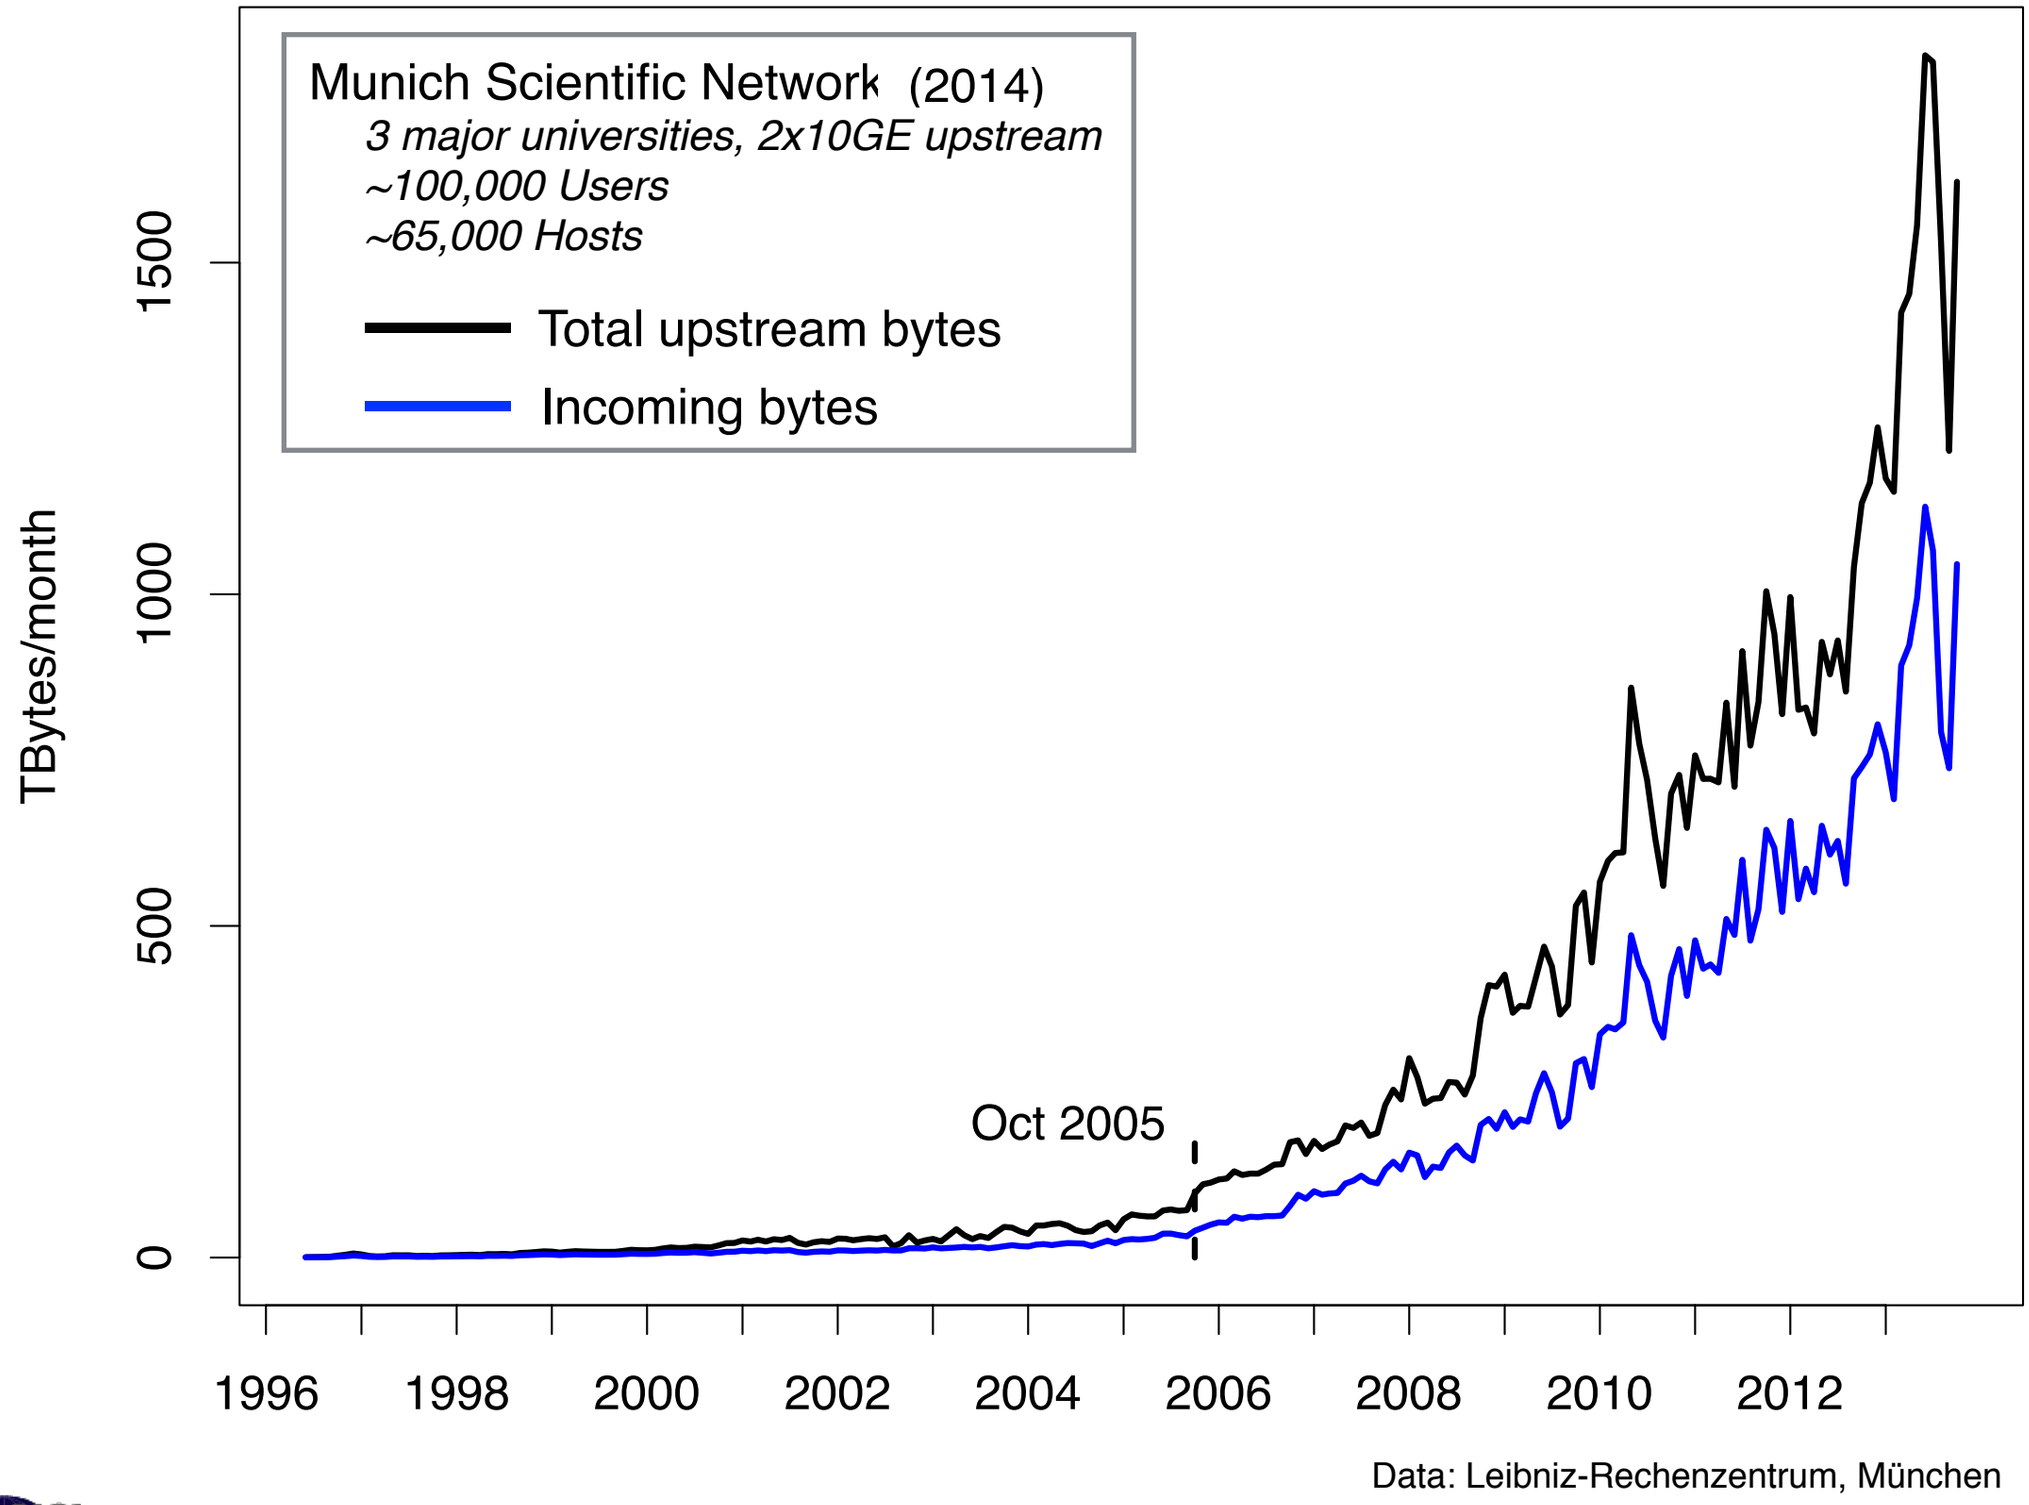

Data: Leibniz-Rechenzentrum, München

LBNL in 2006

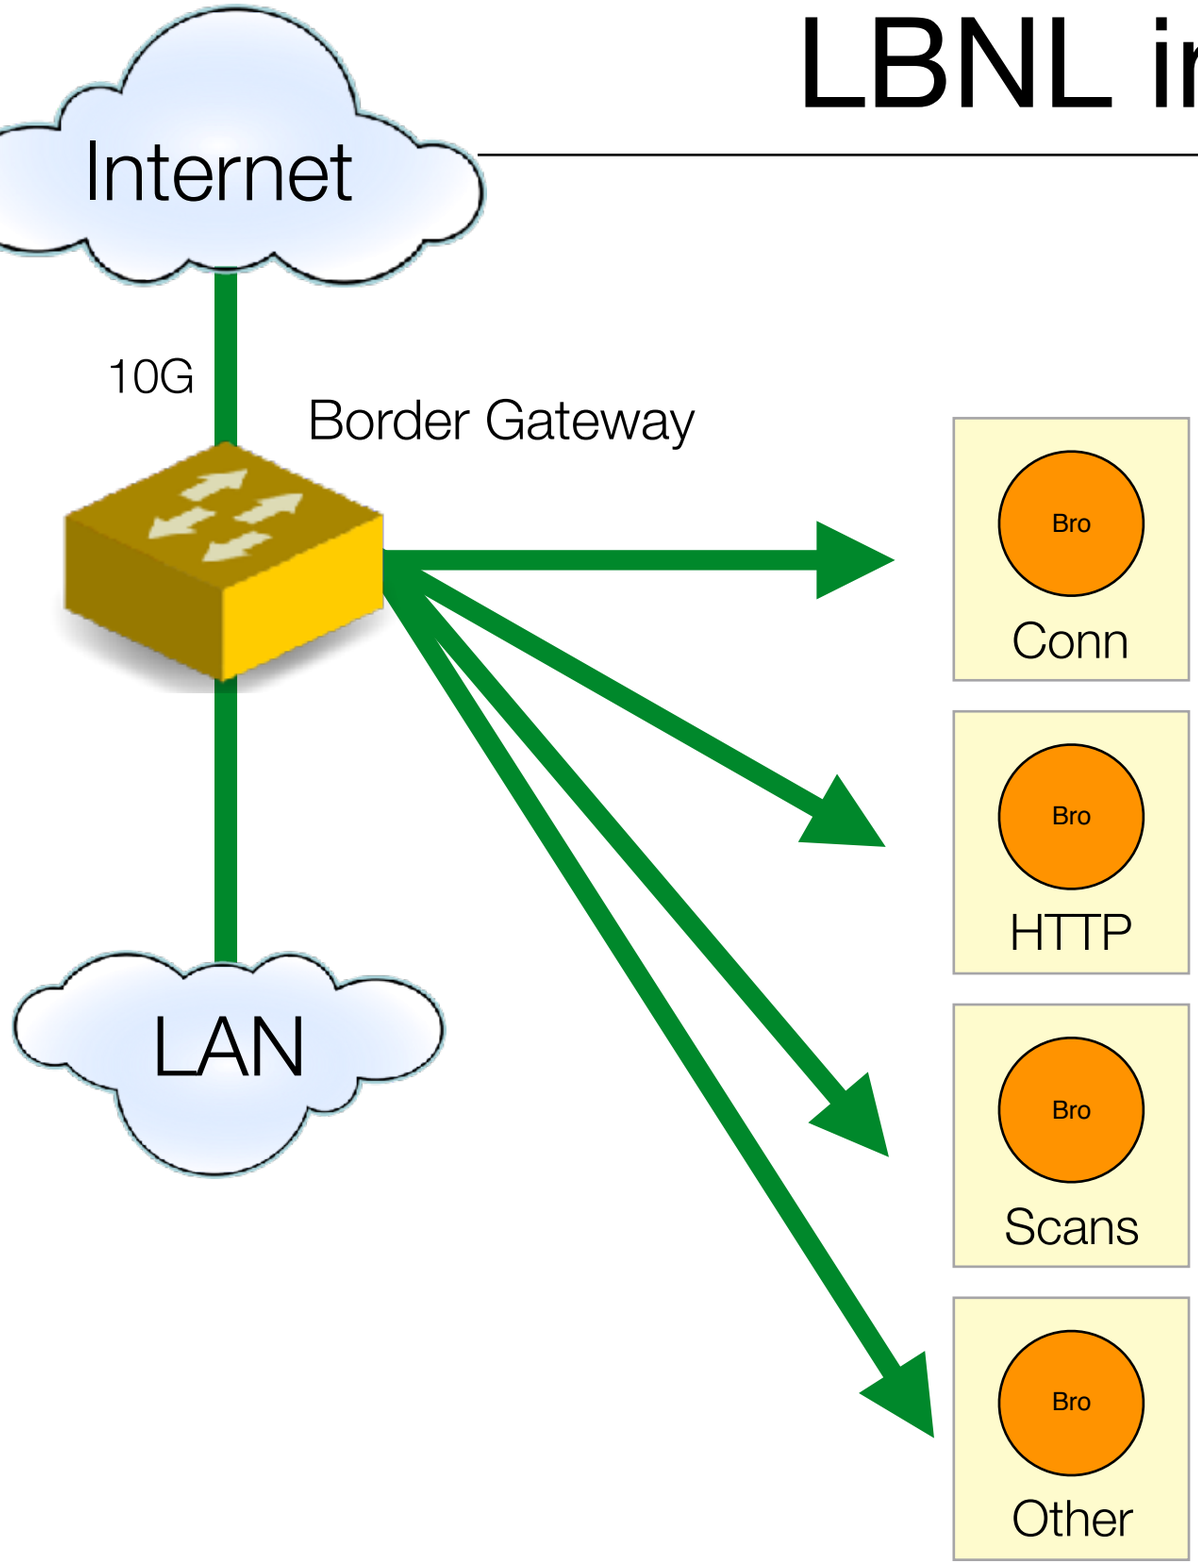
Internet
10G
Border Gateway
LAN
Bro
Conn
Bro
HTTP
Bro
Scans
Bro
Other

Bro Cluster

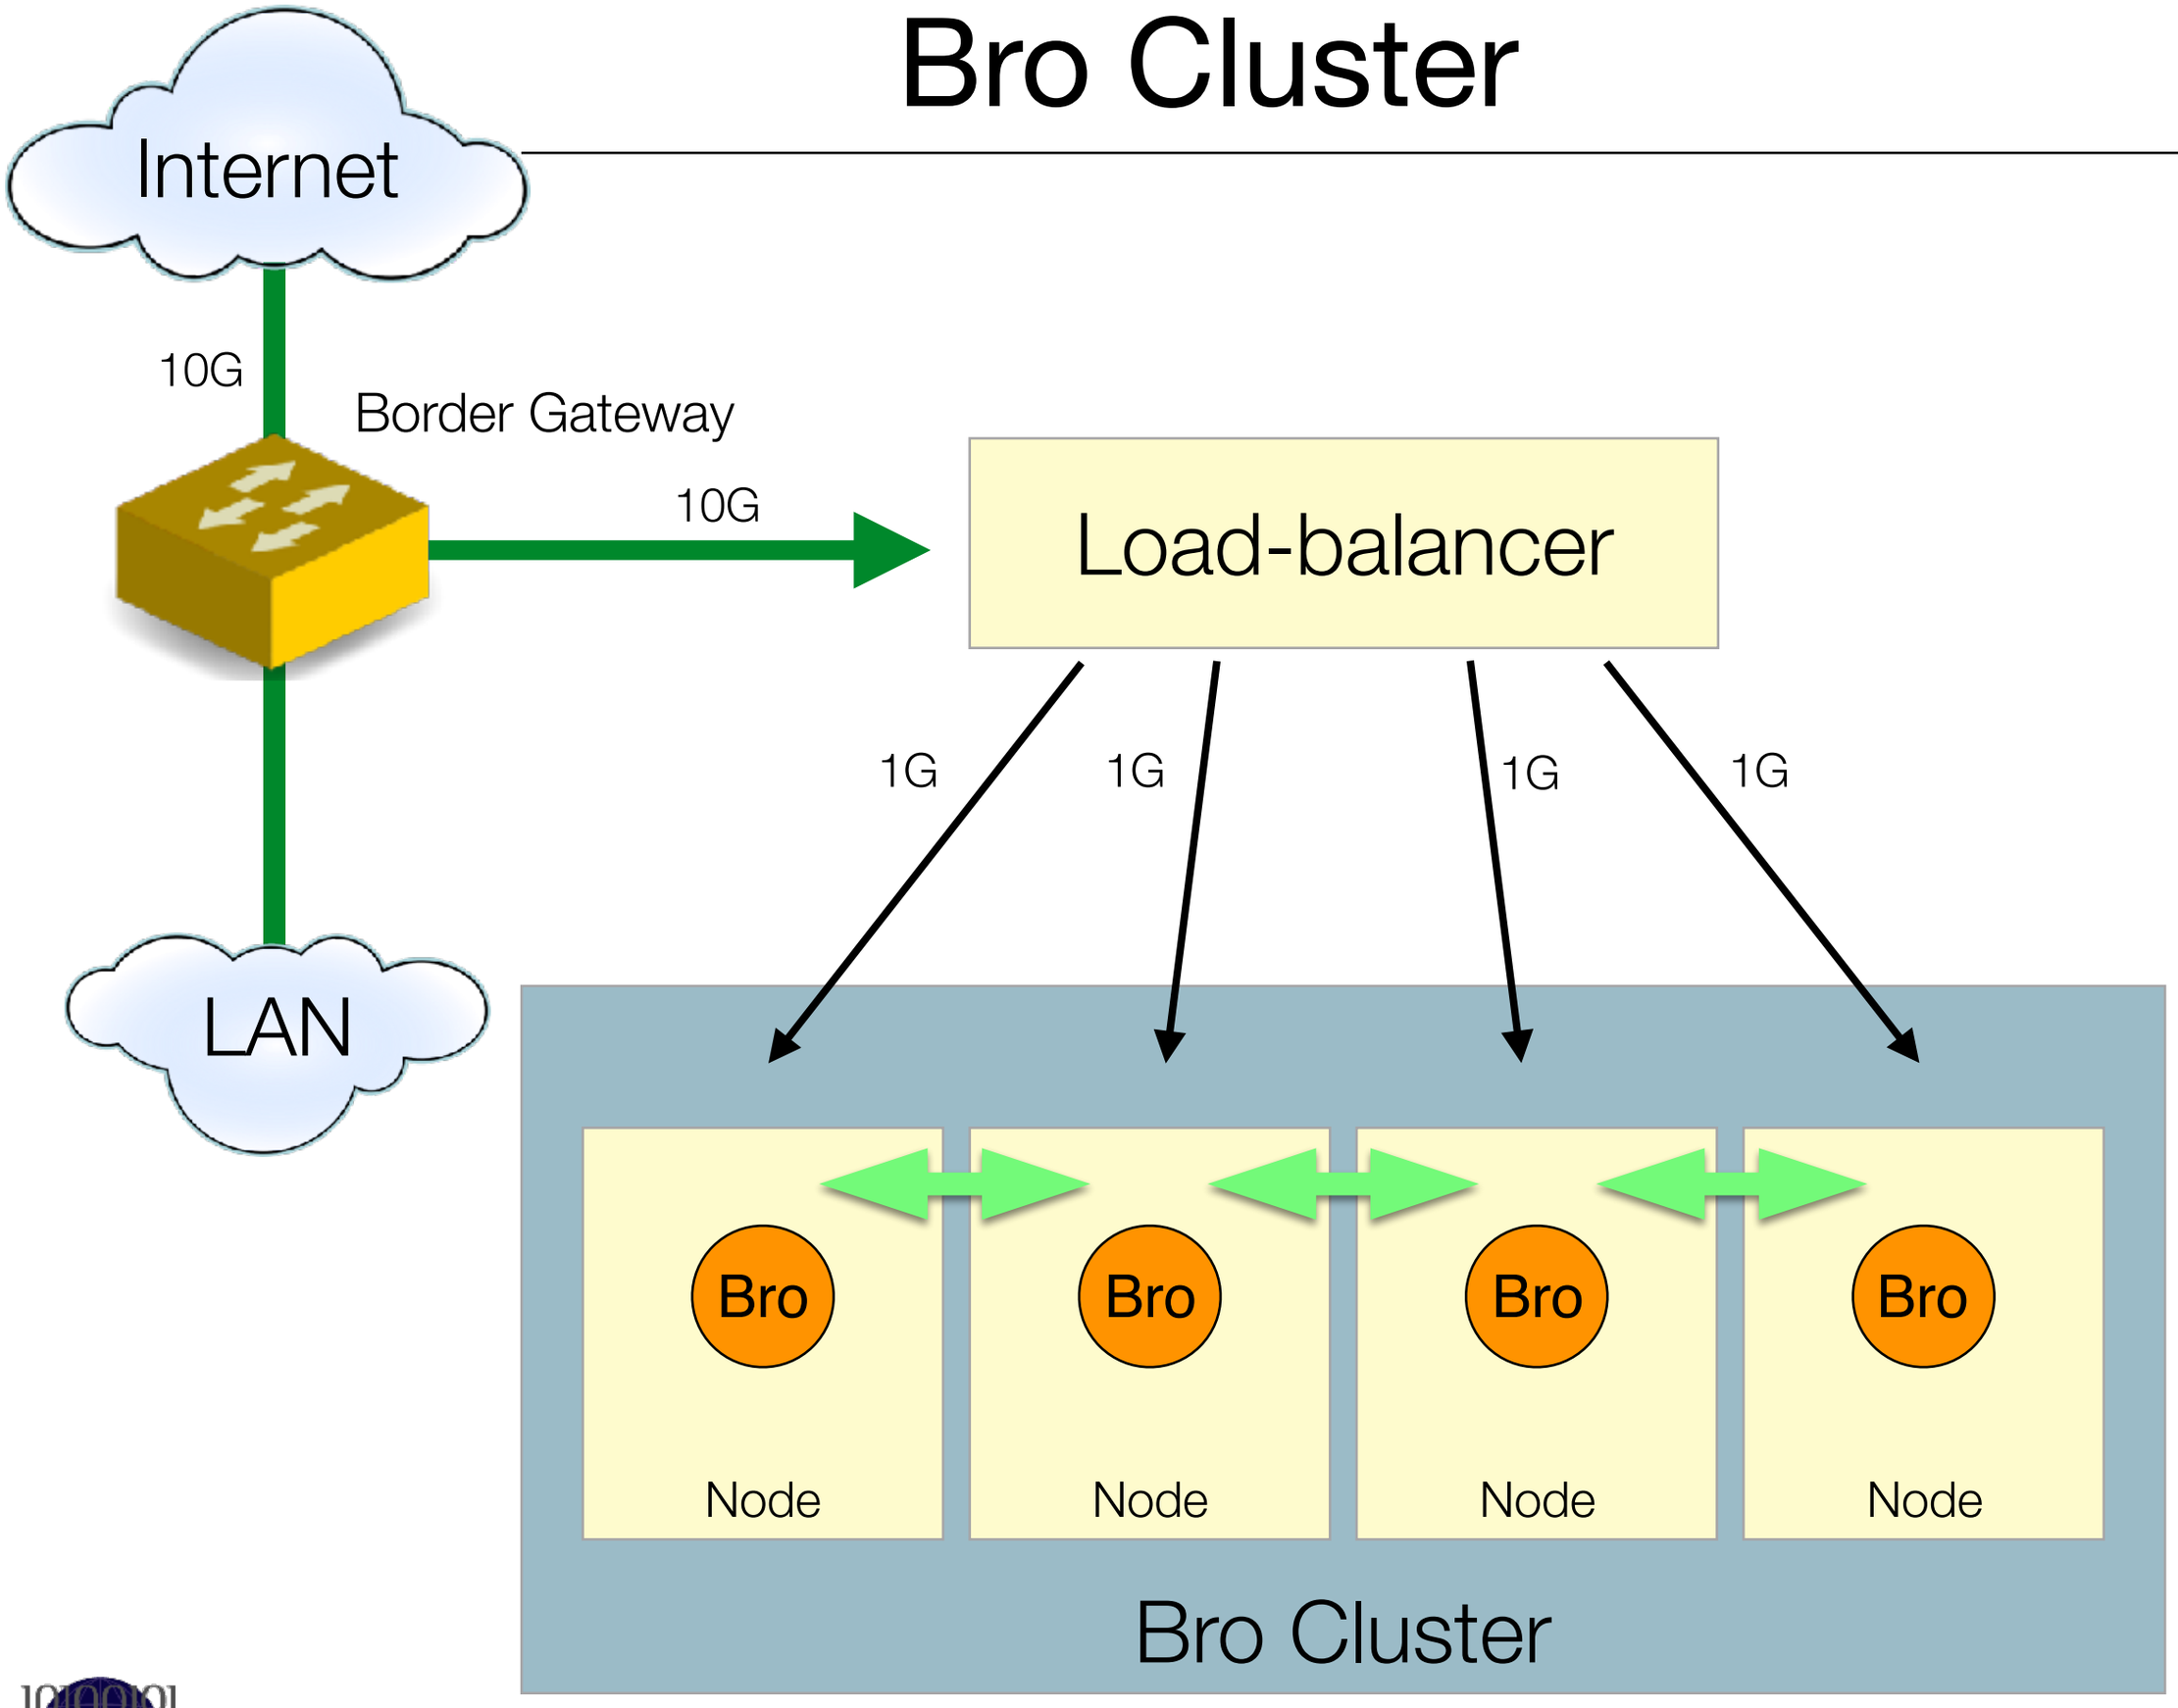
Internet
10G
Border Gateway
10G
Load-balancer
1G
1G
1G
1G
LAN
Bro
Bro
Bro
Bro
Node
Node
Node
Node
Bro Cluster

Bro Cluster

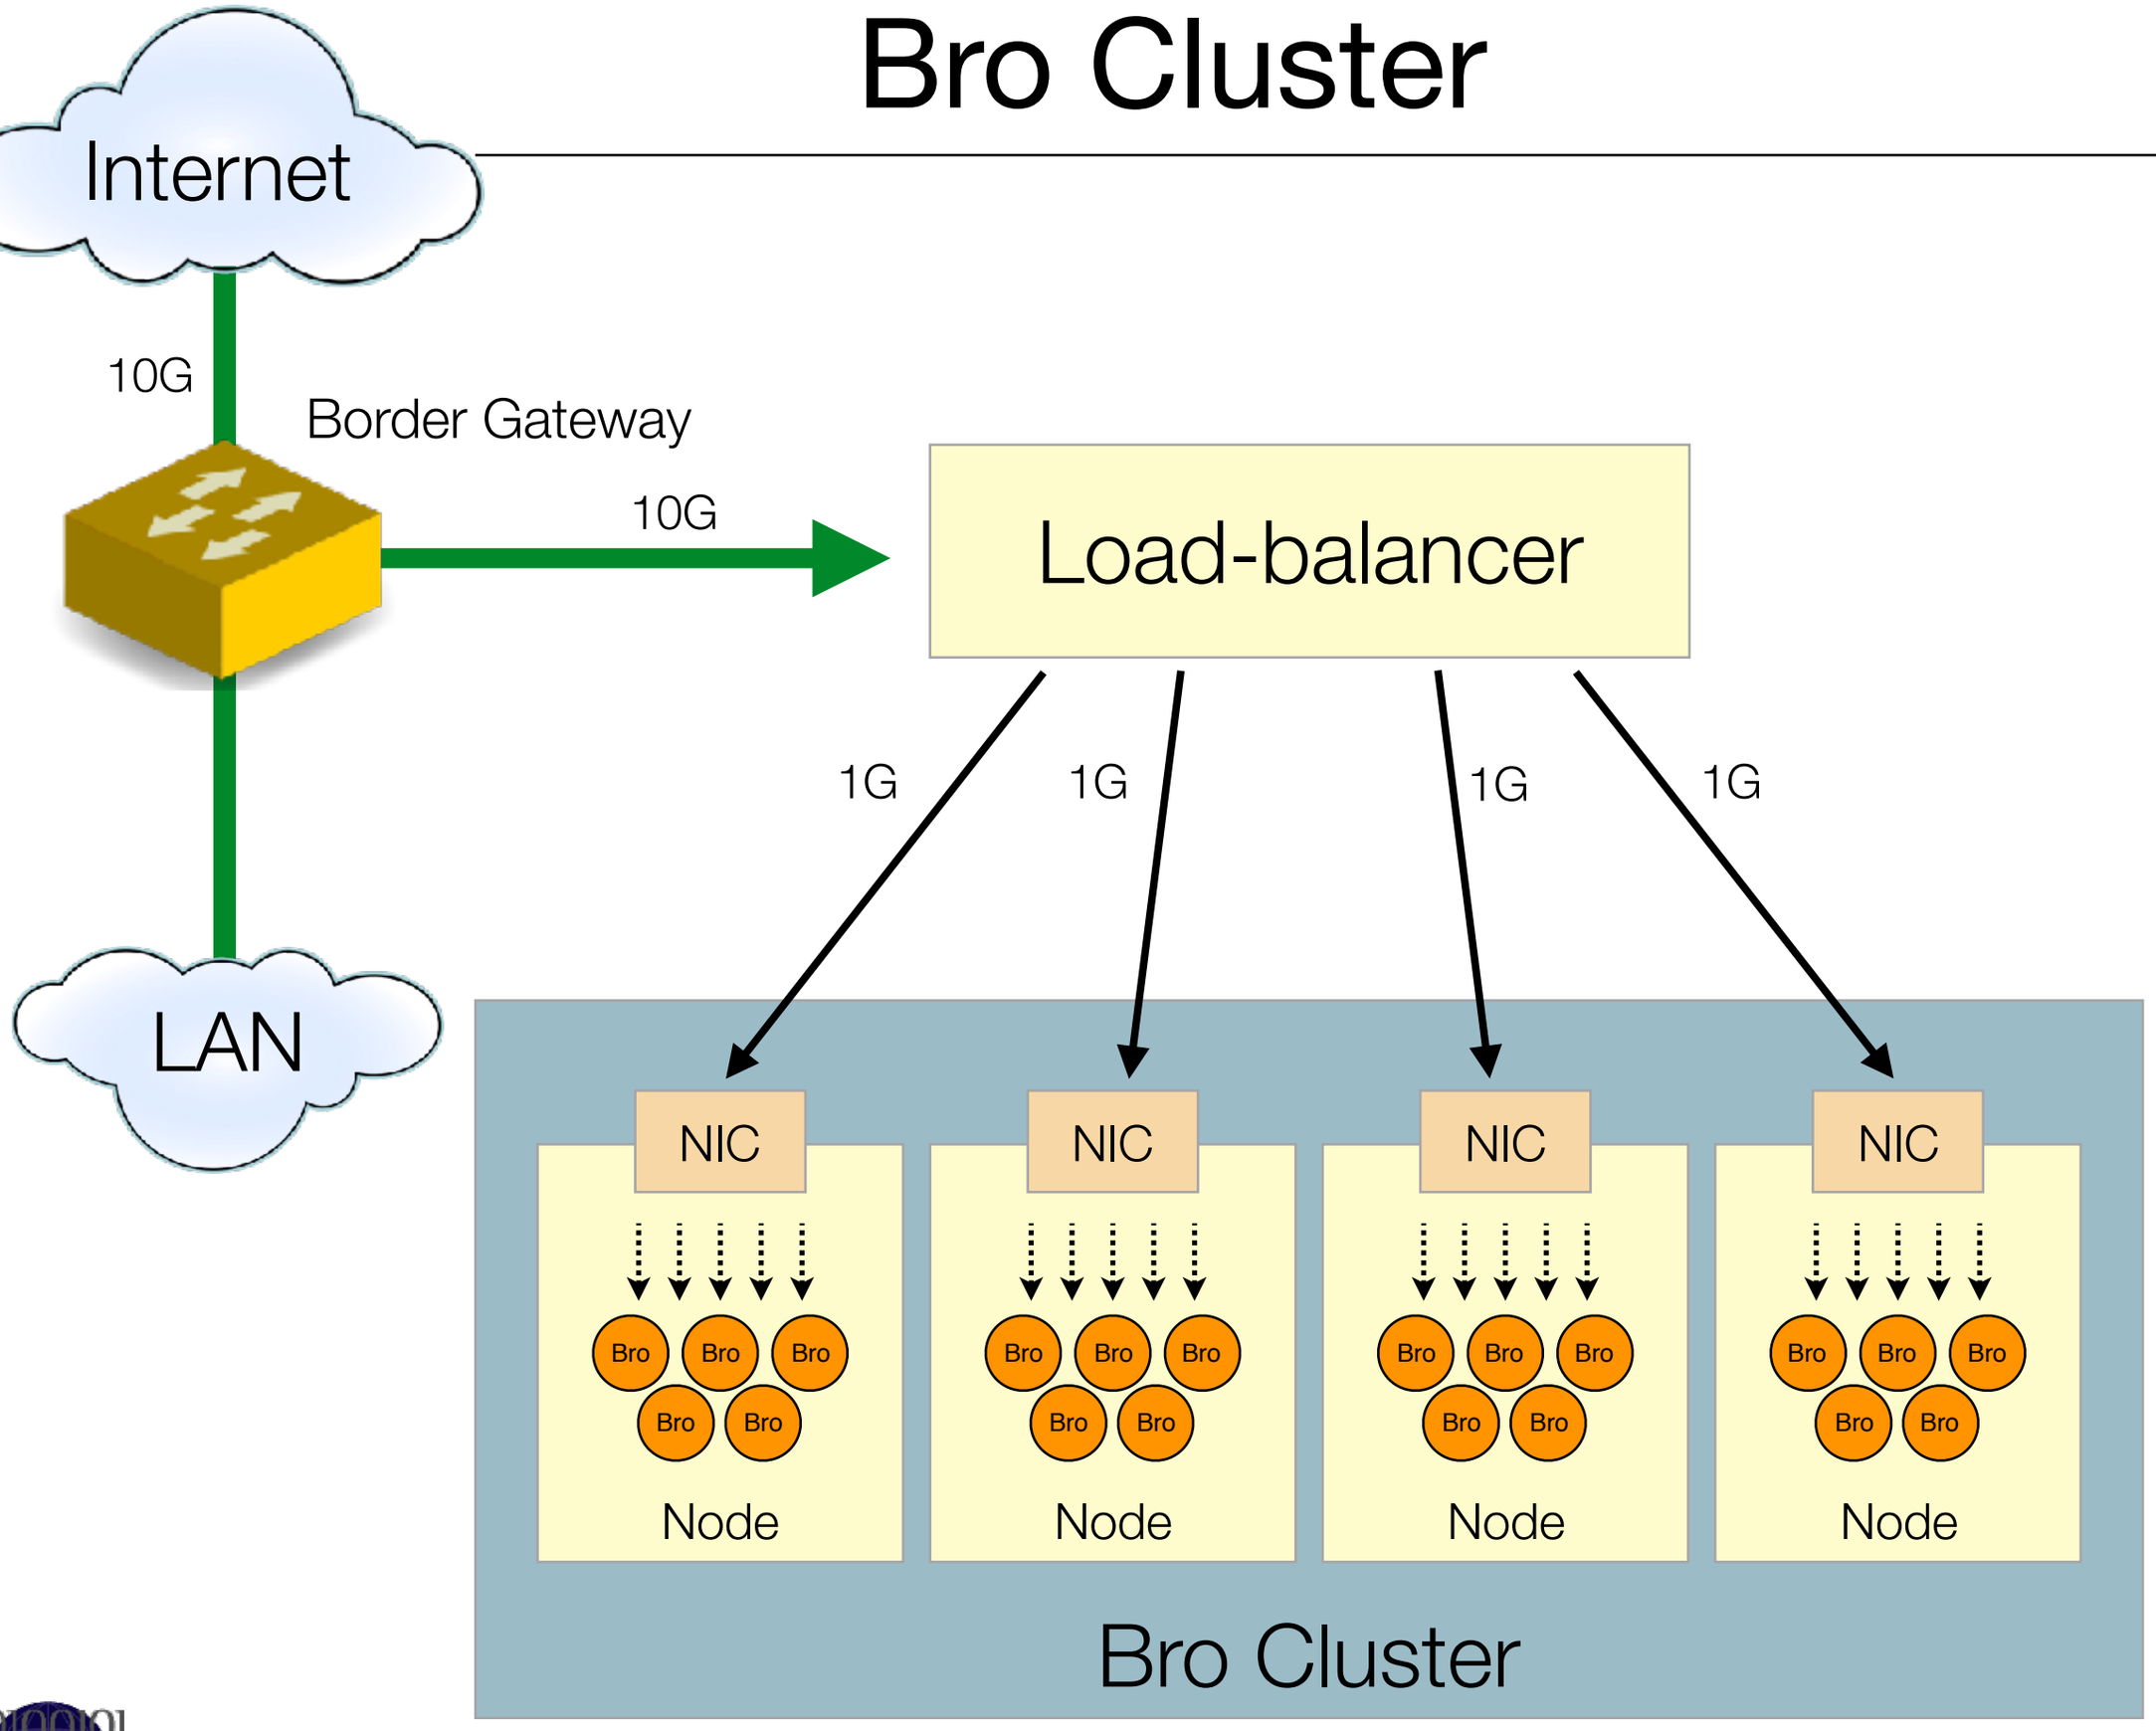
Internet
10G
Border Gateway
10G
Load-balancer
1G
1G
1G
1G
LAN
NIC
NIC
NIC
NIC
Bro
Node
Bro Cluster

A Production Load-Balancer

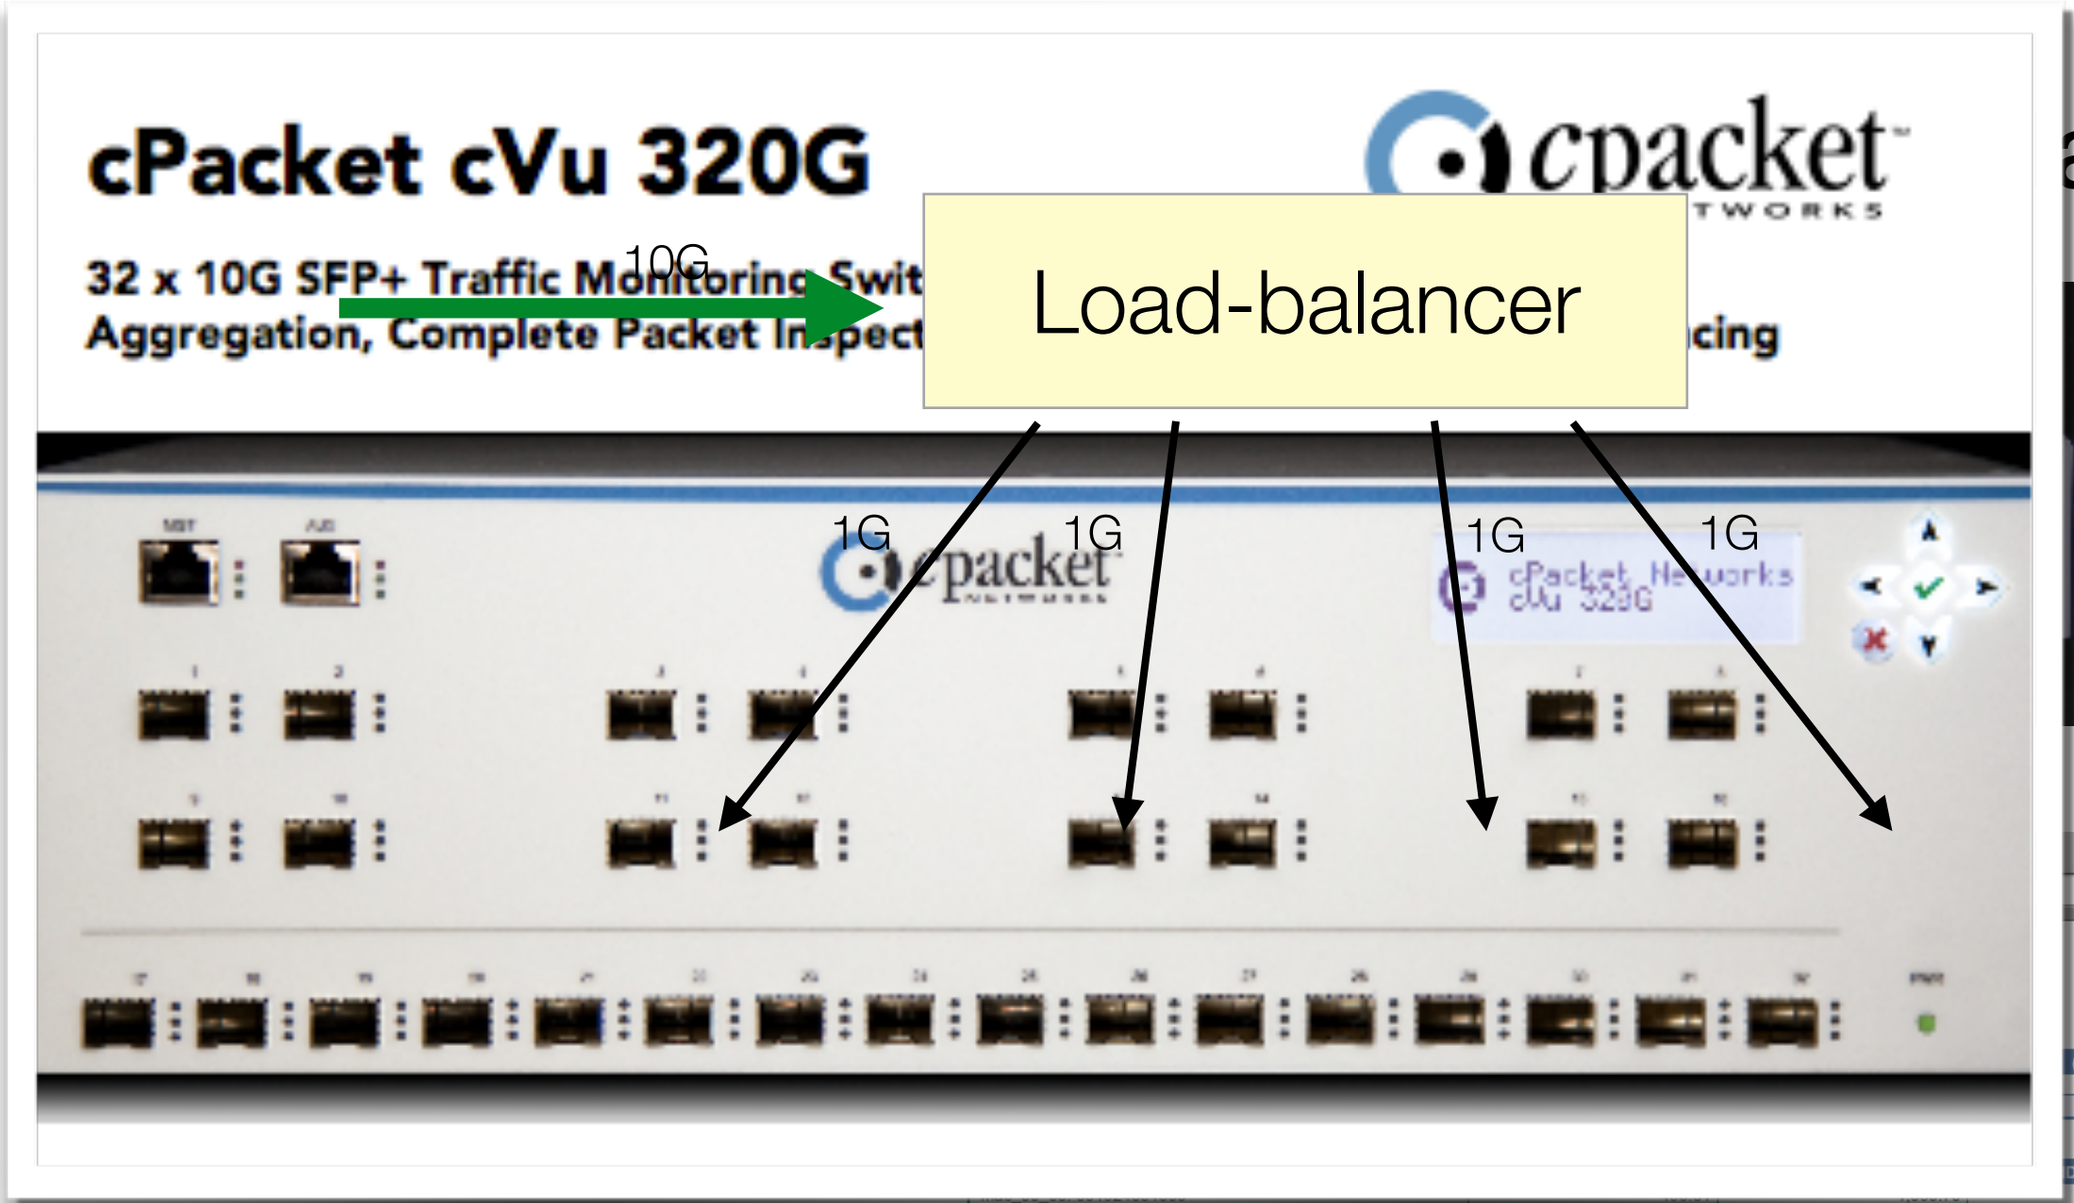

(pps)	Max: (bps)	(pps)
1,345.96	101,256,079	17,629.8
1,345.96	101,256,079	17,629.8

			Dev: (pps)	Max: (pps)
			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

Today: 100G Bro Cluster at LBNL

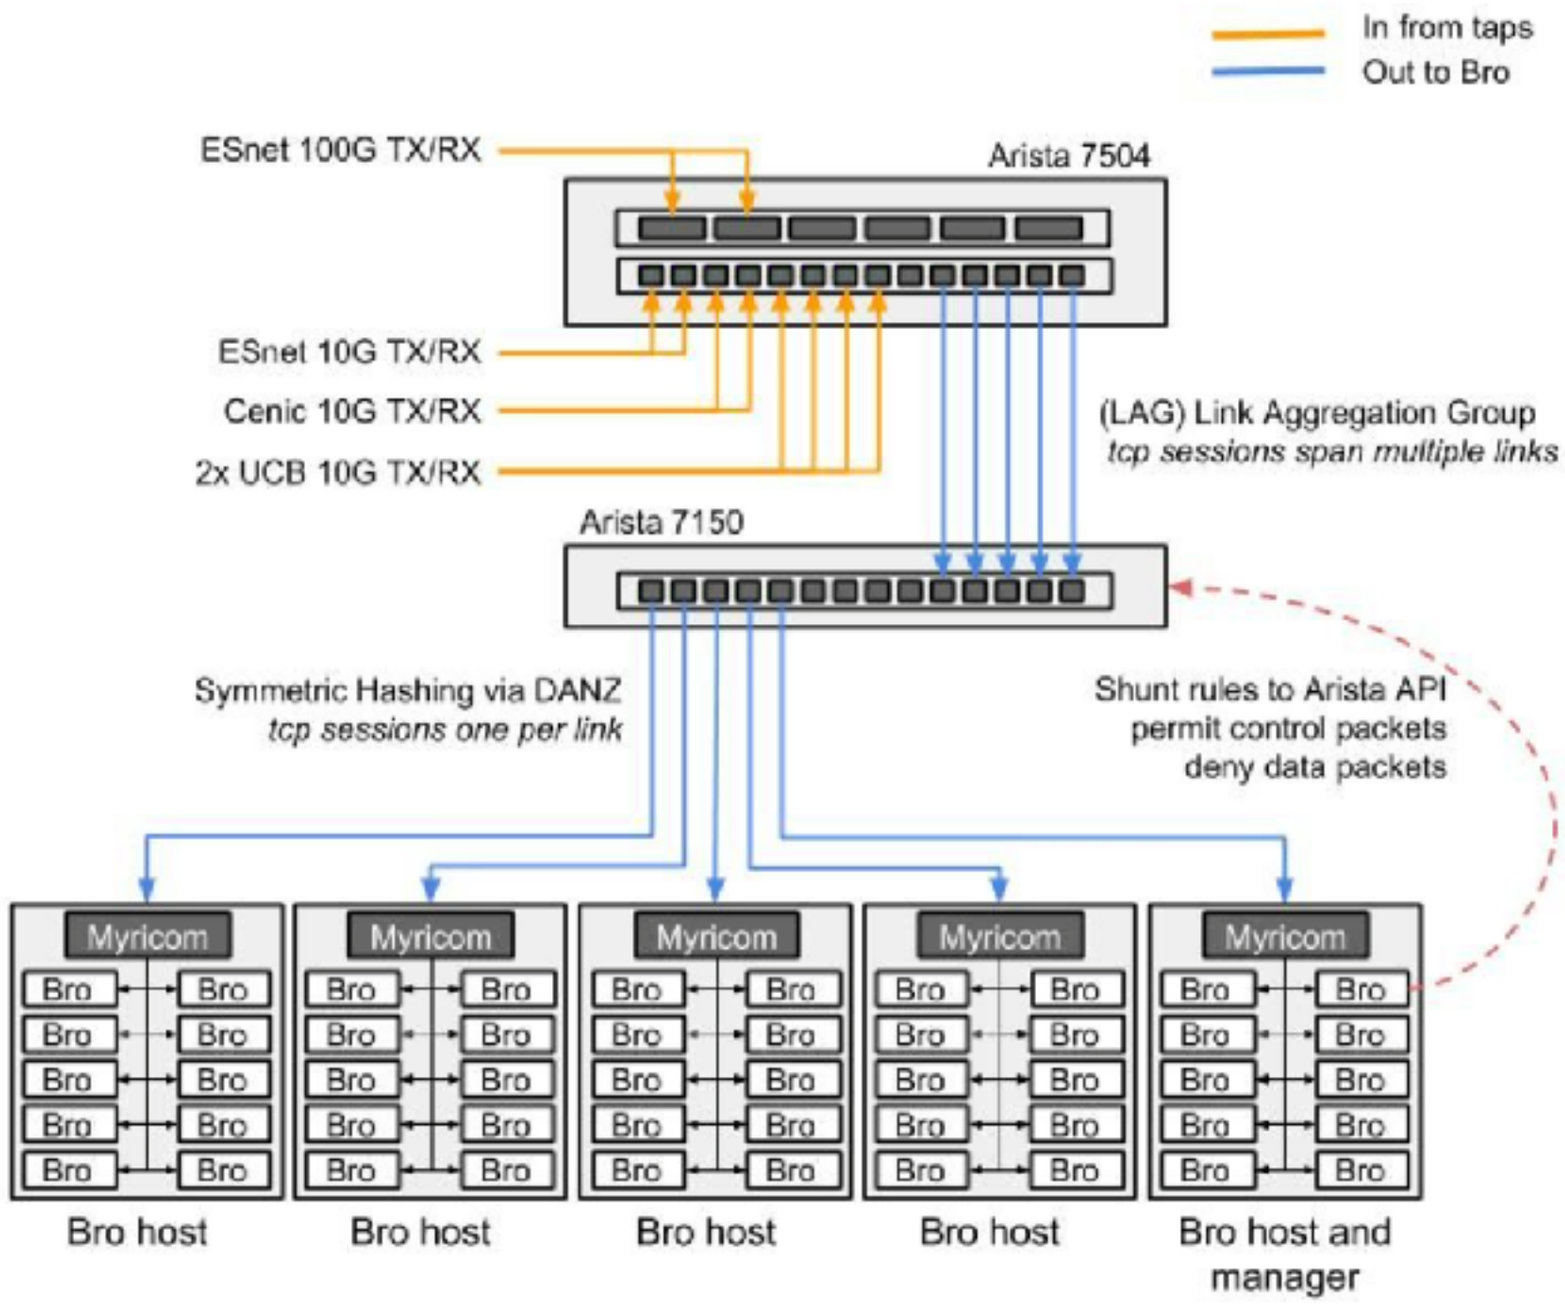

http://go.lbl.gov/100g

Bro History, Part 2

1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015 2016 2017

About 20 publications presenting Bro-related research.

Initial Bro versions are addressing an operational need at LBNL

Basic research at ICSI drives continuous innovation

Feedback loop crucial for both sides

Operational deployment in large-scale open-science networks

Powerful, stable open-source technology that scales to large HPC networks.

Bro paper becomes one of the most cited publications in the space.

However, nobody else even knows about it.

Those who do, find it too difficult to use.

And we don't have the resources to maintain it anymore.

Bro History

1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015 2016 2017

About 20 publications presenting Bro-related research.

Initial Bro versions are addressing an operational need at LBNL

Basic research at ICSI drives continuous innovation

NSF funds tech-transfer through Bro Center at ICSI and NCSA

Domain experts revamp the user interface
Engineers improve code base & documentation
Outreach focuses on community building

Security teams are looking for new solutions

Bro deployment skyrockets

Operational deployment in large-scale open-science networks

"Everybody uses it."

Now we have a new challenge: A different kind of demand that the open-source team cannot satisfy anymore …

A Tale of Two Users

Science & Higher Education

Happy to experiment.

Used to open-source software.

Driven by skilled individuals.

Limited funding.

Bro Center of Expertise

Enterprises & Governments

Used to purchasing solutions.

Require reliable point of contact.

Avoid dependence on individuals.

More flexible budgets.

ICSI Spin-off

Corelight, Inc.

Enterprise-grade Bro solutions, from the creators of Bro.

Bootstrapping: Commercial Bro support

Today: Fully-supported, turn-key Bro appliances

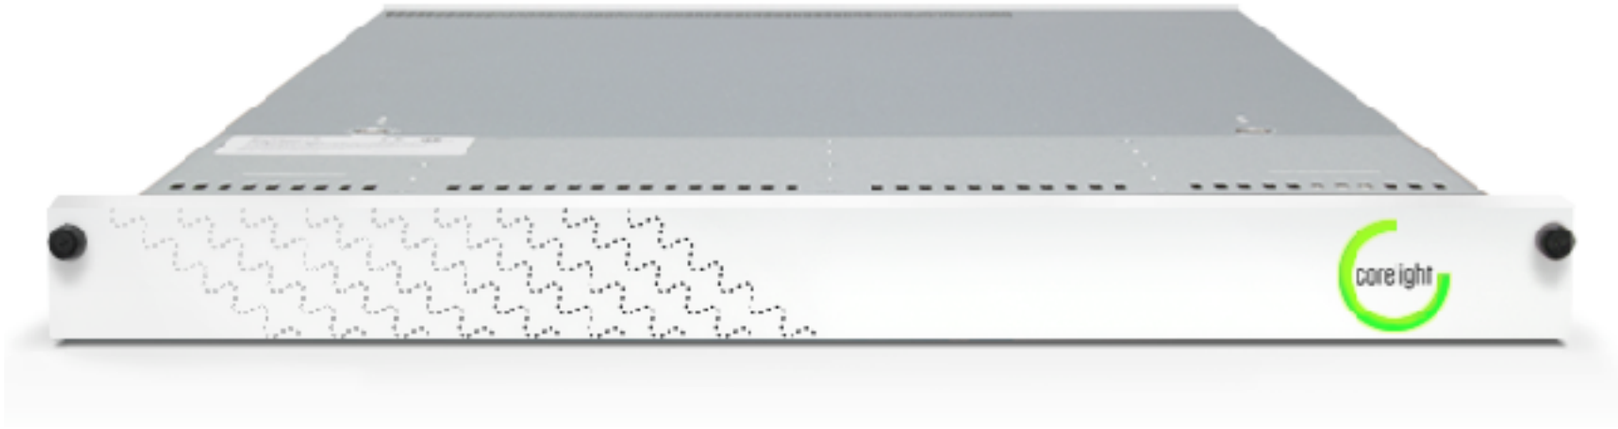

BroBox One

Visibility, made elegantly simple.

Advantage: Integration

With BroBox One we are controlling the full stack.

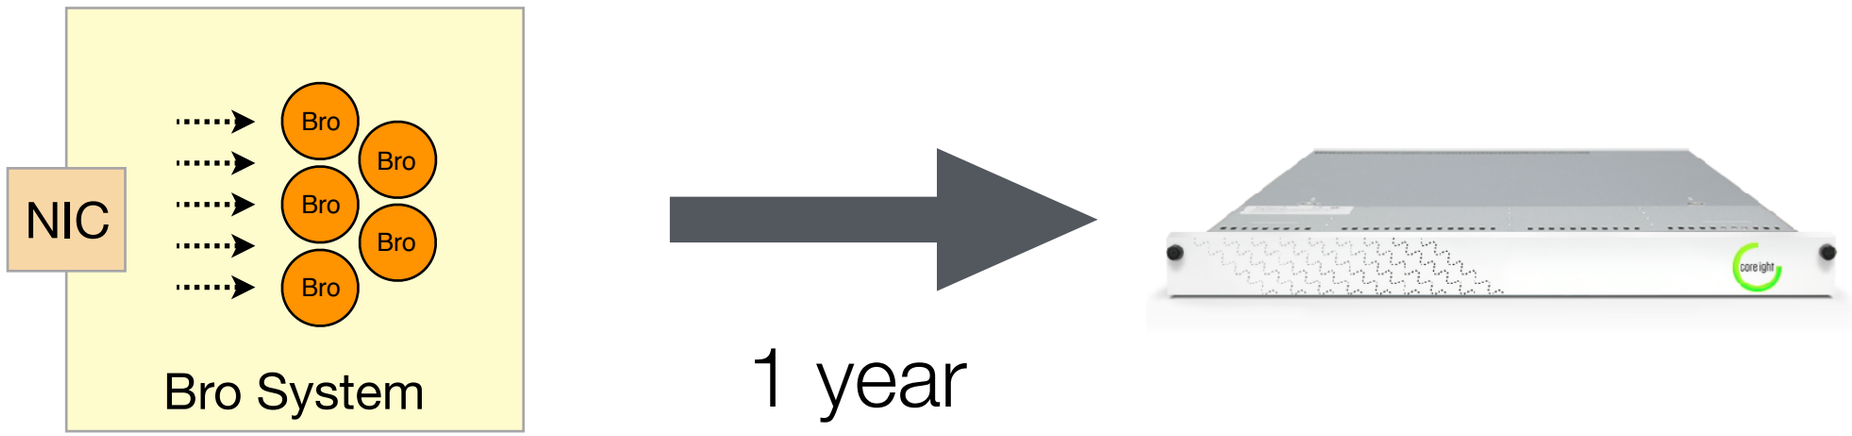

We can take integration much further,
while maintaining the open-source spirit.

Bro History

1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015 2016 2017

About 20 publications presenting Bro-related research.

Initial Bro versions are addressing an operational need at LBNL

Basic research at ICSI drives continuous innovation

NSF funds tech-transfer through Bro Center at ICSI and NCSA

15 more ICSI publications presenting Bro-related research

Operational deployme

Large-scale real-world networks

Bro deployments skyrocket

Close interaction with real environments continues to open up new research opportunities

corelight

Protecting Enterprise Environments

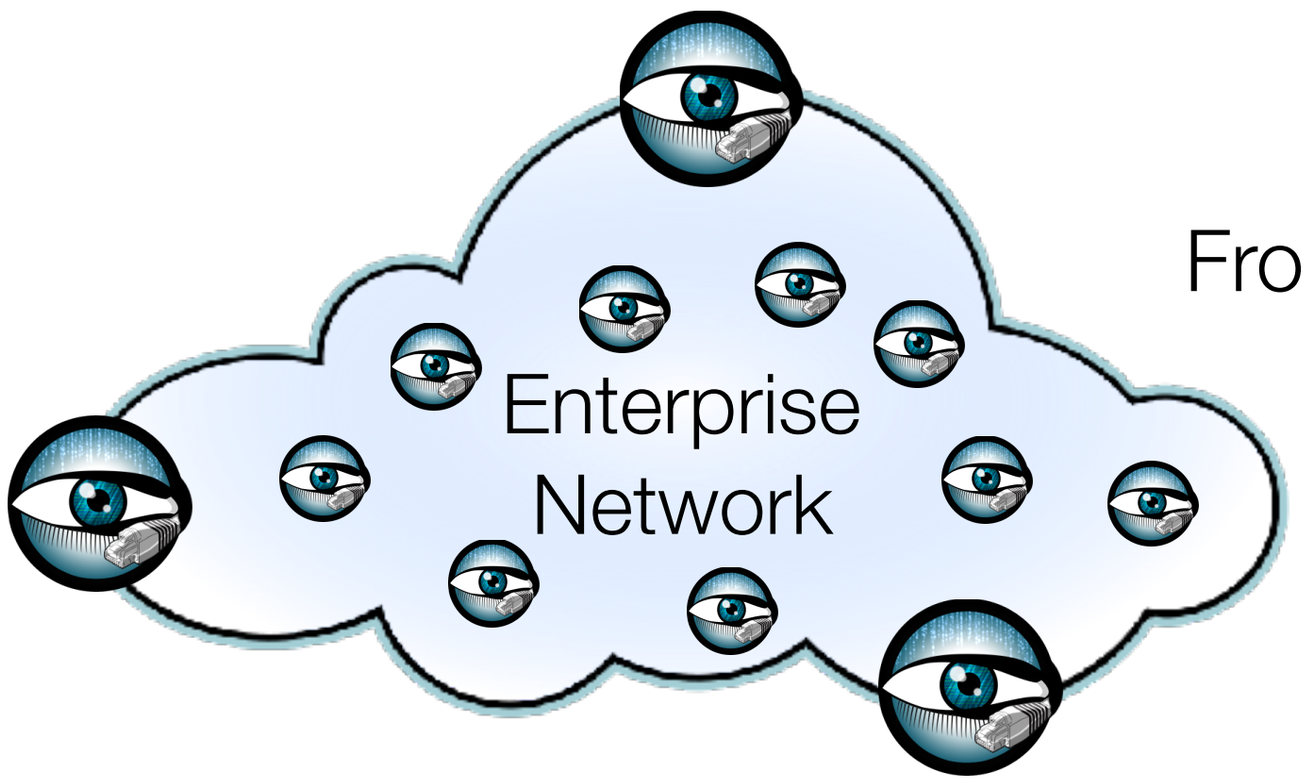

From perimeter to internal.

From standalone to coordinated.

From passive to active.

We are working on functionality to support all of this.

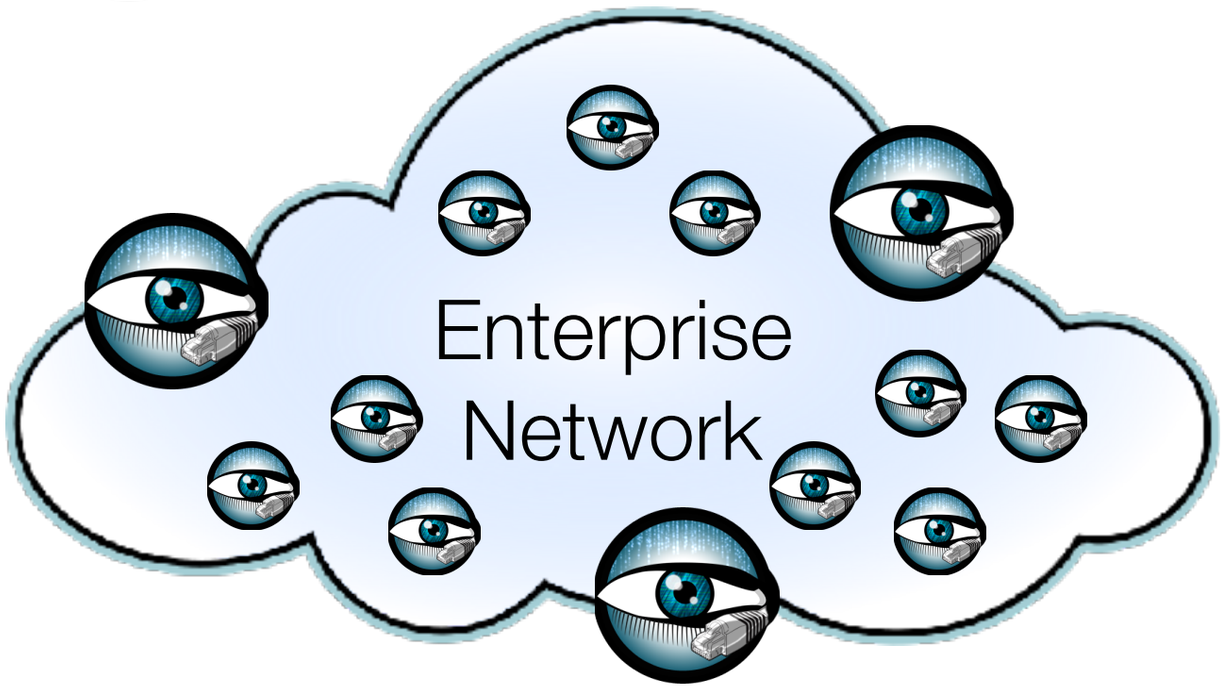

Foundation: Broker

Bro's new unified communication library.

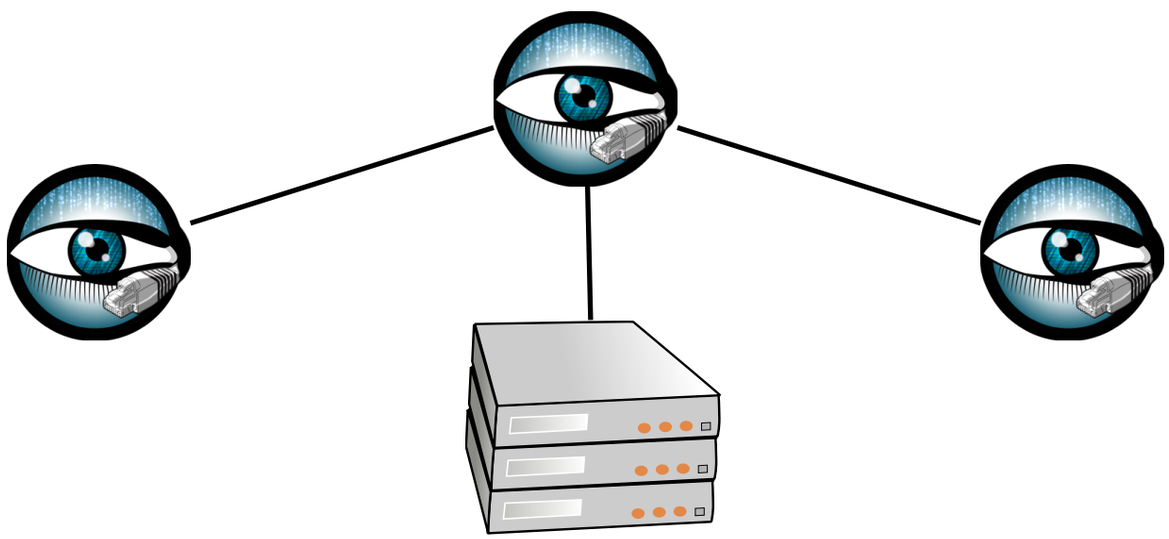

Log forwarding.
Event exchange.
Global key/value stores.

Public/subscribe.
APIs for Bro, C++, C, Python.
BSD license.

`https://github.com/bro/broker`

ICSI

corelight

Integrating Host Monitoring

Leverage control over end hosts.

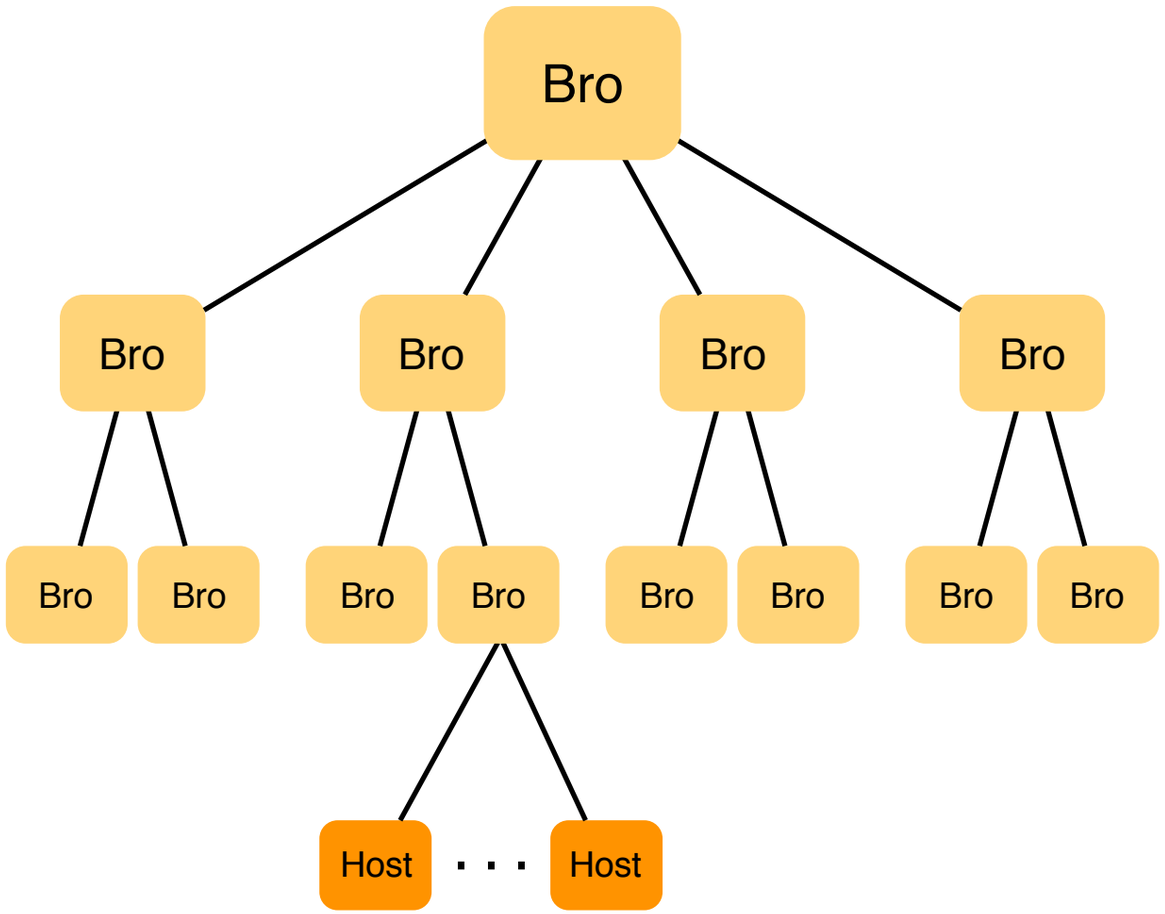

osquery

processes

All running processes on the host system.

Column	Type	Description
pid	INTEGER	Process (or thread) ID
name	TEXT	The process path or shorthand argv[0]
path	TEXT	Path to executed binary
cmdline	TEXT	Complete argv
cwd	TEXT	Process current working directory
root	TEXT	Process virtual root directory
uid	BIGINT	Unsigned user ID
gid	BIGINT	Unsigned group ID
euid	BIGINT	Unsigned effective user ID
egid	BIGINT	Unsigned effective group ID
parent	INTEGER	Process parent's PID

```
select * from processes where pid = 1
```

Source: Facebook

https://github.com/bro/bro-osquery

Protecting ICS and IoT

Goal: Detect attacks as unexpected process deviations

PLC

Controller /
Human Machine Interface

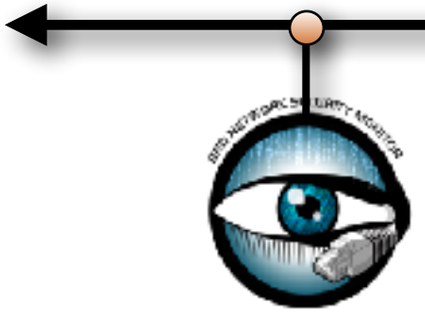

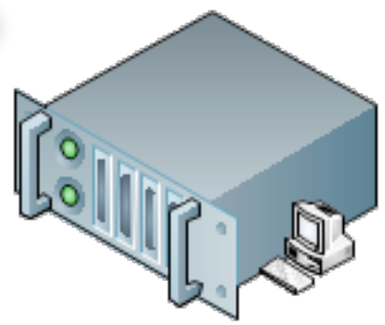

Testbed setup: Water tank with heater

Reg.	Name	Type	Desc
HR0010	V1On	bool	Status of valve 1
HR0011	V2On	bool	Status of valve 2
HR0012	HeaterOn	bool	Status of the heater
HR0020	TankLevelSP	fixpoint	SP tank level (L)
HR0021	TankLevel	fixpoint	Level of the tank (L)
HR0022	TempSP	fixpoint	SP water temp.
HR0023	Temp	fixpoint	Water temp (celsius)
HR0030	TankLevelAl	enum	Alarms tank level
HR0031	TempAl	enum	Alarms water temp.

Protecting Science DMZs

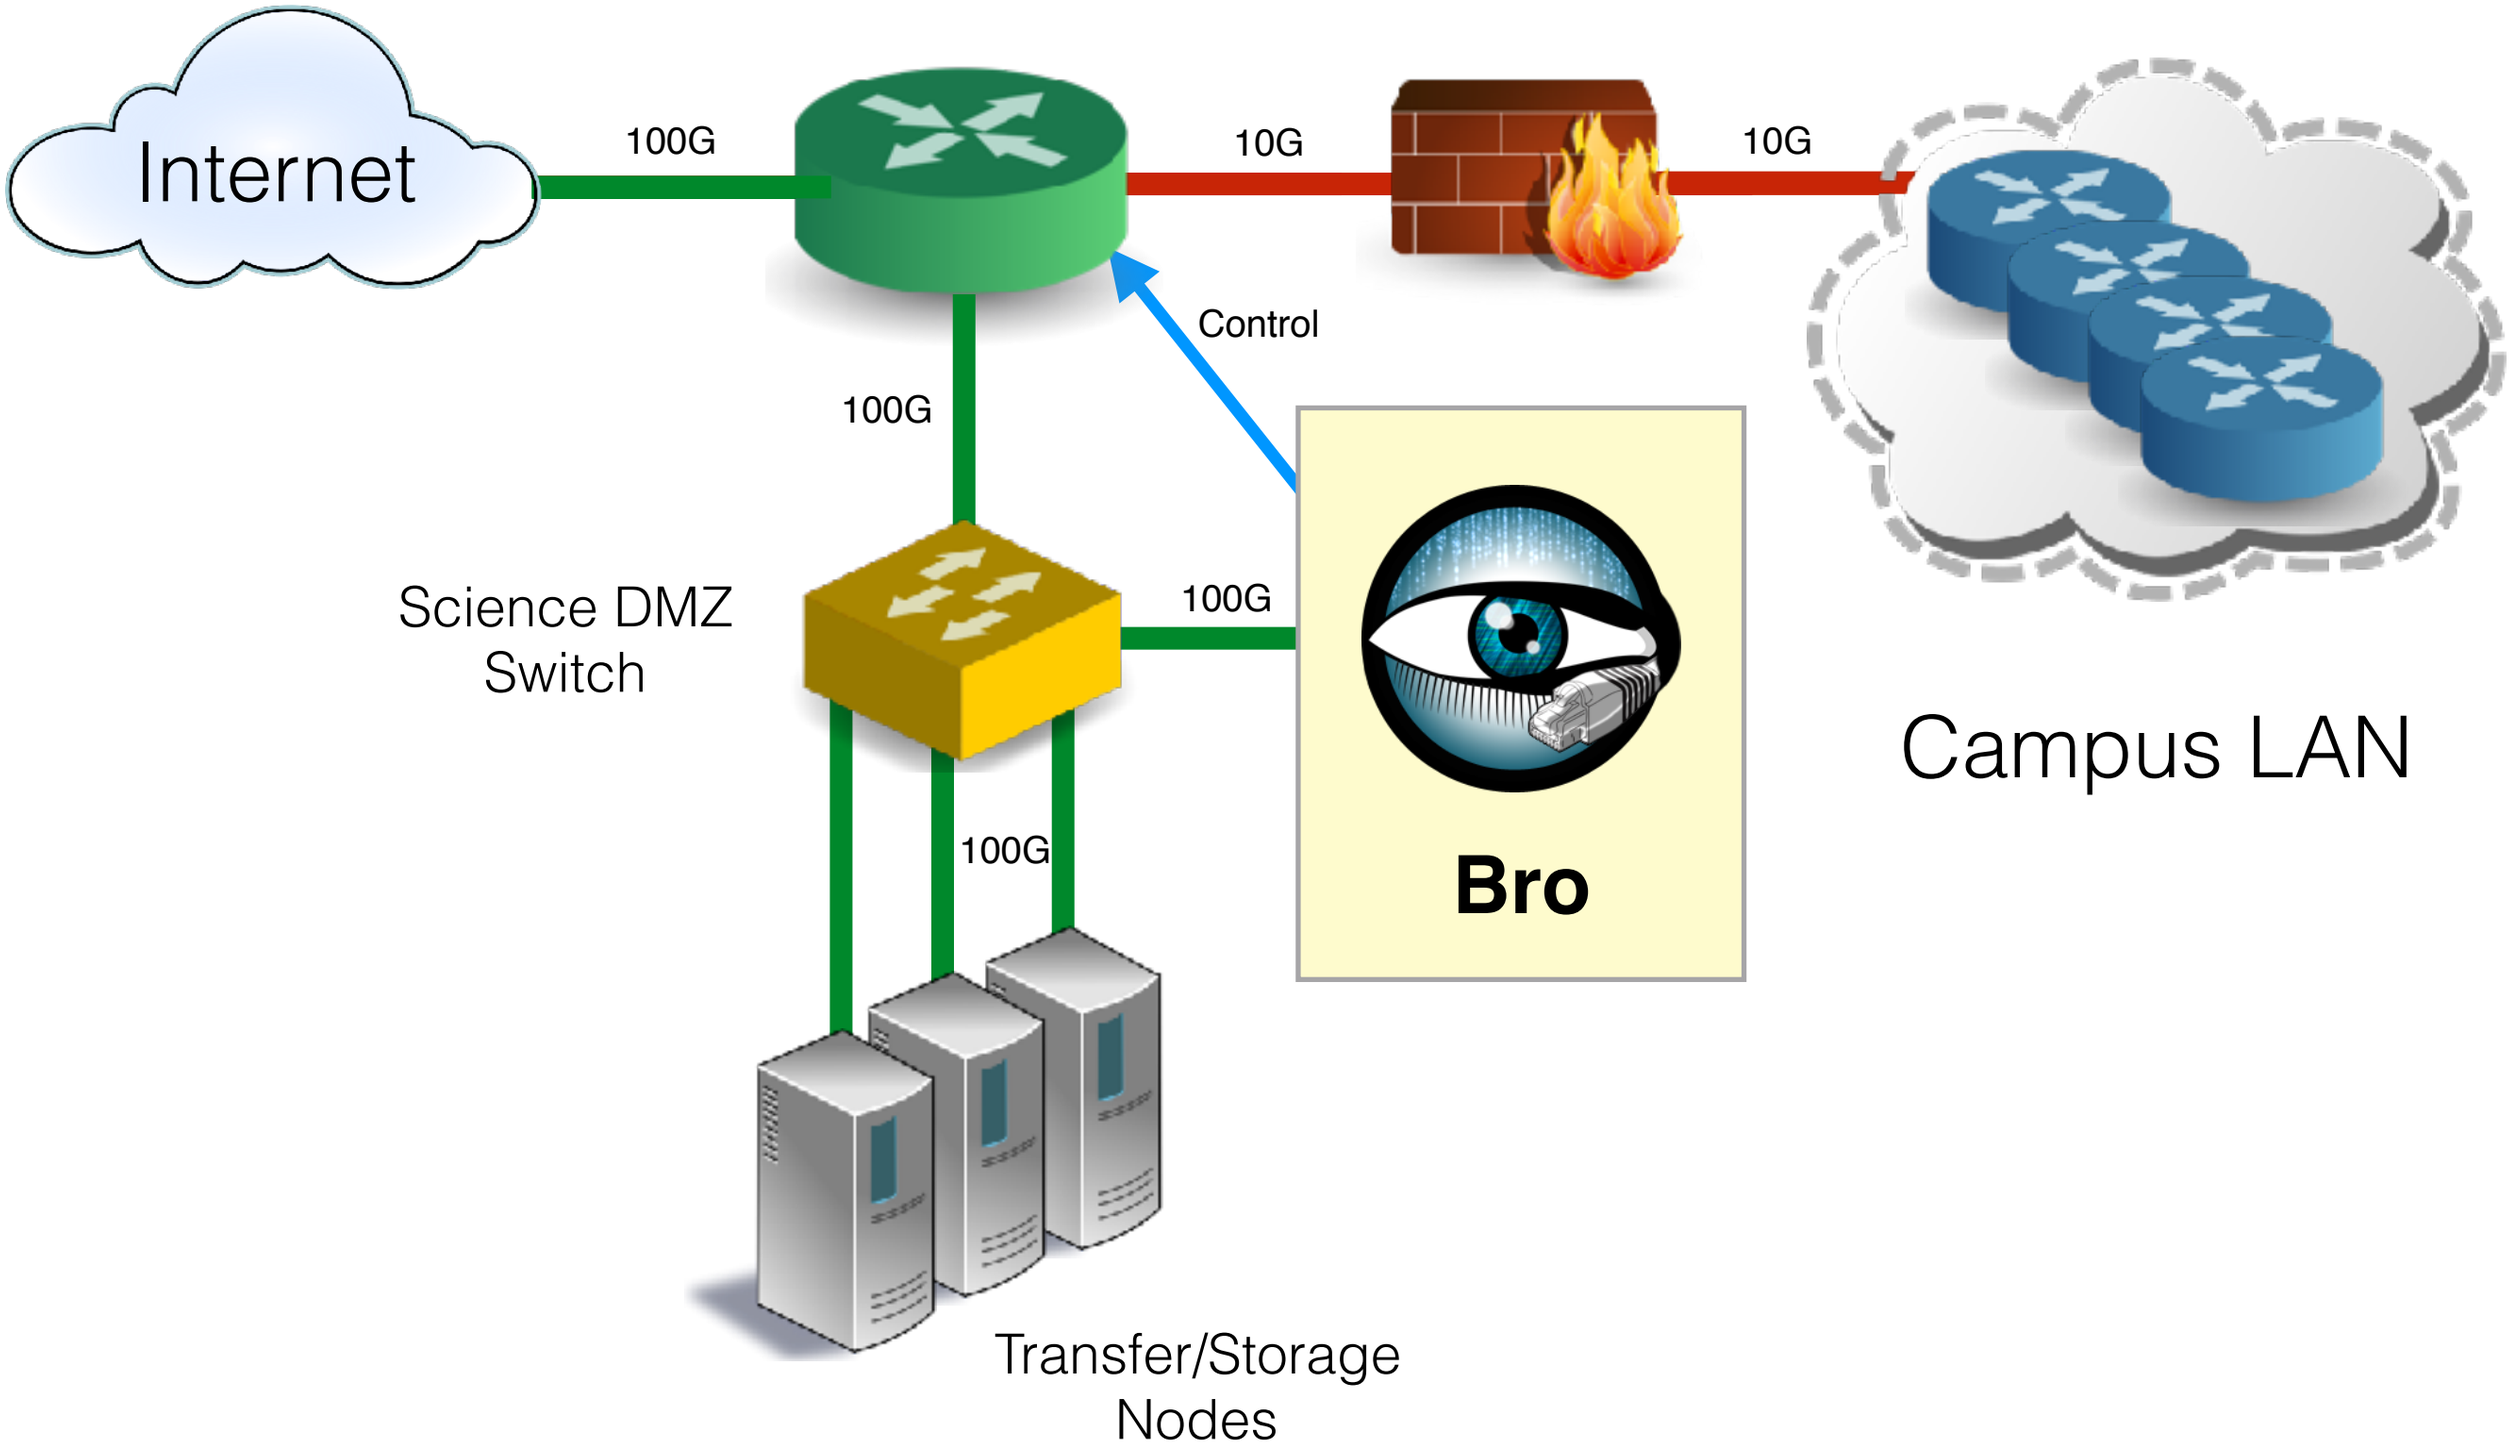

The Modern Cyber Security Stack

Open-source, based on science,
fueled by data & analytics

Basic research

Feedback loop is
crucial for innovation

Real-world operations

The U.S. National Science Foundation has enabled much of Bro.

Bro is coming out of two decades of academic research, along with extensive transition to practice efforts. NSF has supported much of that, and is currently funding the Bro Center of Expertise at the International Computer Science Institute and the National Center for Supercomputing Applications.

The Bro Project is a member of Software Freedom Conservancy.

Software Freedom Conservancy, Inc. is a 501(c)(3) not-for-profit organization that helps promote, improve, develop, and defend Free, Libre, and Open Source Software projects.

The Bro Project
www.bro.org
info@bro.org

Enterprise Solutions
www.corelight.com
info@corelight.com