Software Engineering Institute

FloCon 2010 Proceedings

January 2010

CERT Program

http://www.sei.cmu.edu



Carnegie Mellon

Copyright 2010 Carnegie Mellon University

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This material has been approved for public release and unlimited distribution except as restricted below.

Internal use:* Permission to reproduce this material and to prepare derivative works from this material for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use:* This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

* These restrictions do not apply to U.S. government entities.

Presentation and Demo:

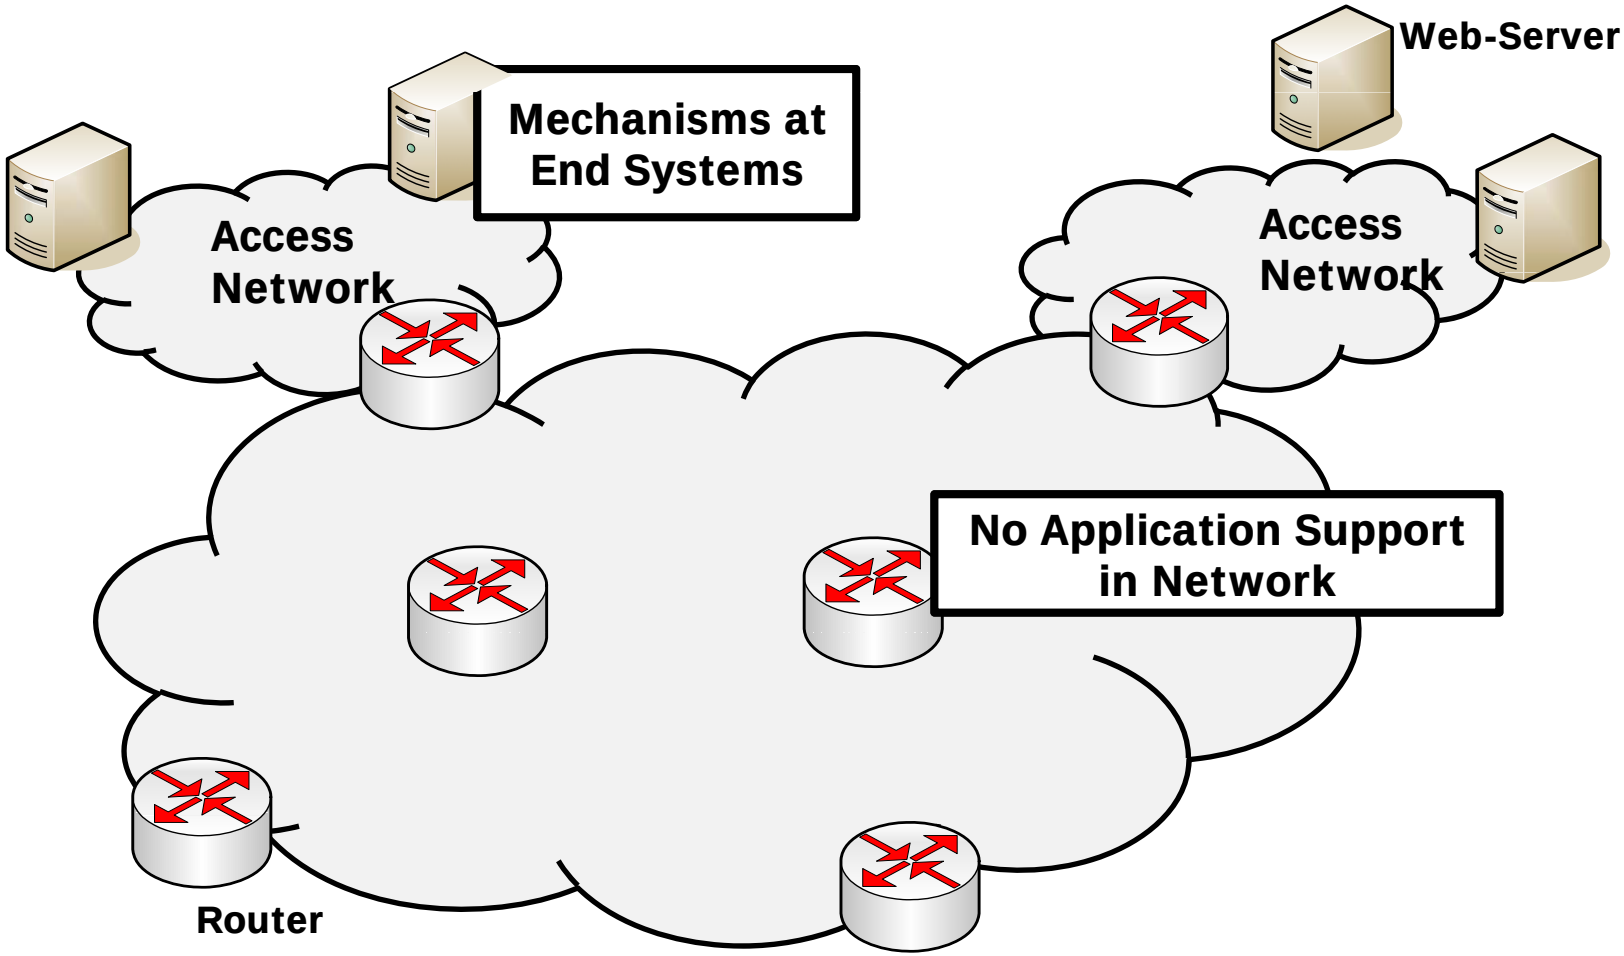
Flow Valuations based on Network-Service Cooperation

Tanja Zseby, Thomas Hirsch
Competence Center Network Research
Fraunhofer Institute FOKUS, Berlin, Germany



© 2010 Fraunhofer FOKUS, T. Zseby

Today: End-to-End Principle



© 2010 Fraunhofer FOKUS, T. Zseby

Internet Problems

- Security Threats
  - Immense monetary loss
  - Fast adapting attack patterns
- Management costs
  - Increasing network size
  - Increasing heterogeneity (technical, administrative)
- Lack of support for users and applications
  - Quality, security levels, route options, privacy, etc.

Reconsider Internet Design Principles

© 2010 Fraunhofer FOKUS, T. Zseby

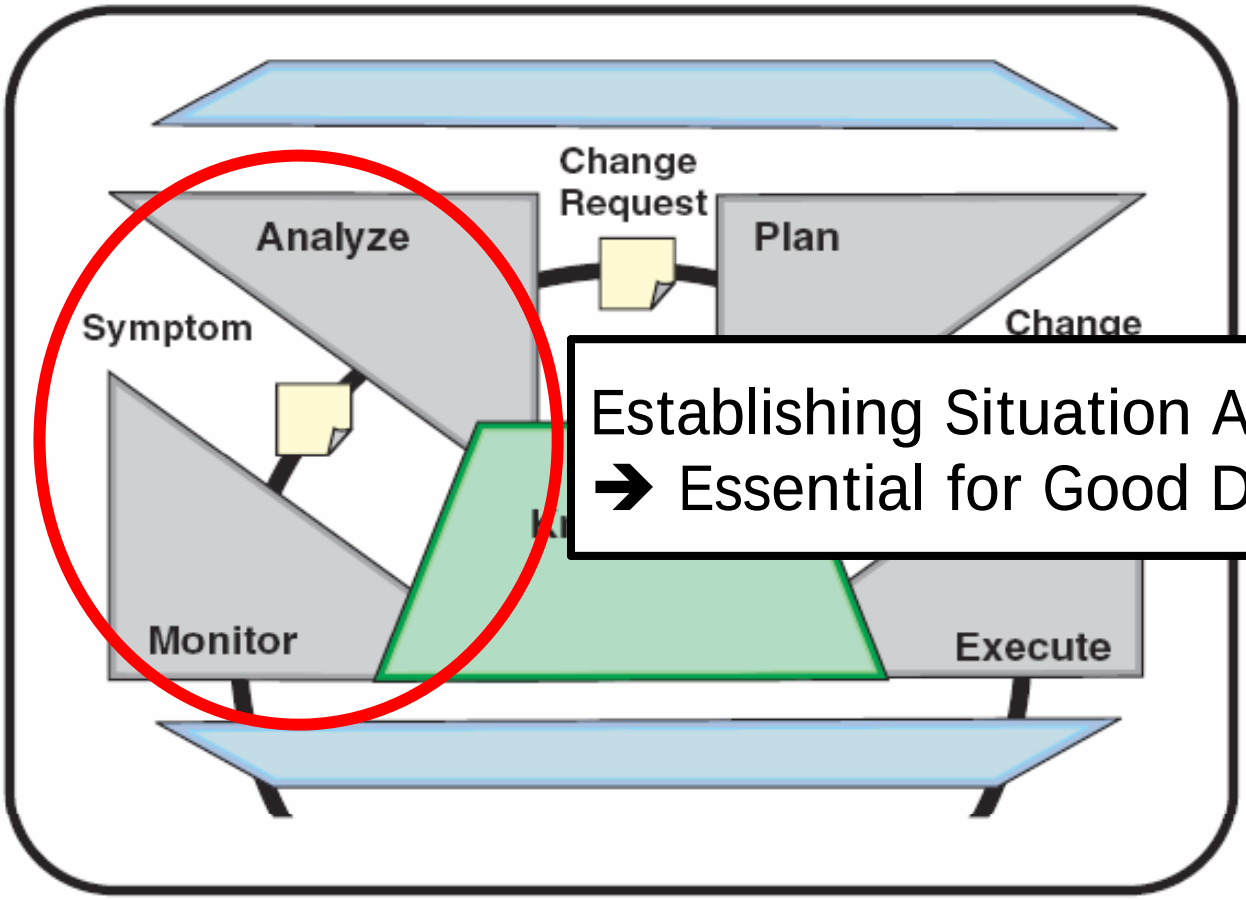
The Idea of Autonomic Networking

- Bring ***decision cycles*** in network nodes
  - Establish situation awareness
  - Enable decision processes beyond routing
- Objectives:
  - Support for Applications in the network
    - Flexible levels for quality, security, etc.
  - Self-Management
    - Reduction of human intervention
  - Self-Protection
    - Protection of network and applications



© 2010 Fraunhofer FOKUS, T. Zseby

Decision Cycle



Establishing Situation Awareness
➔ Essential for Good Decisions

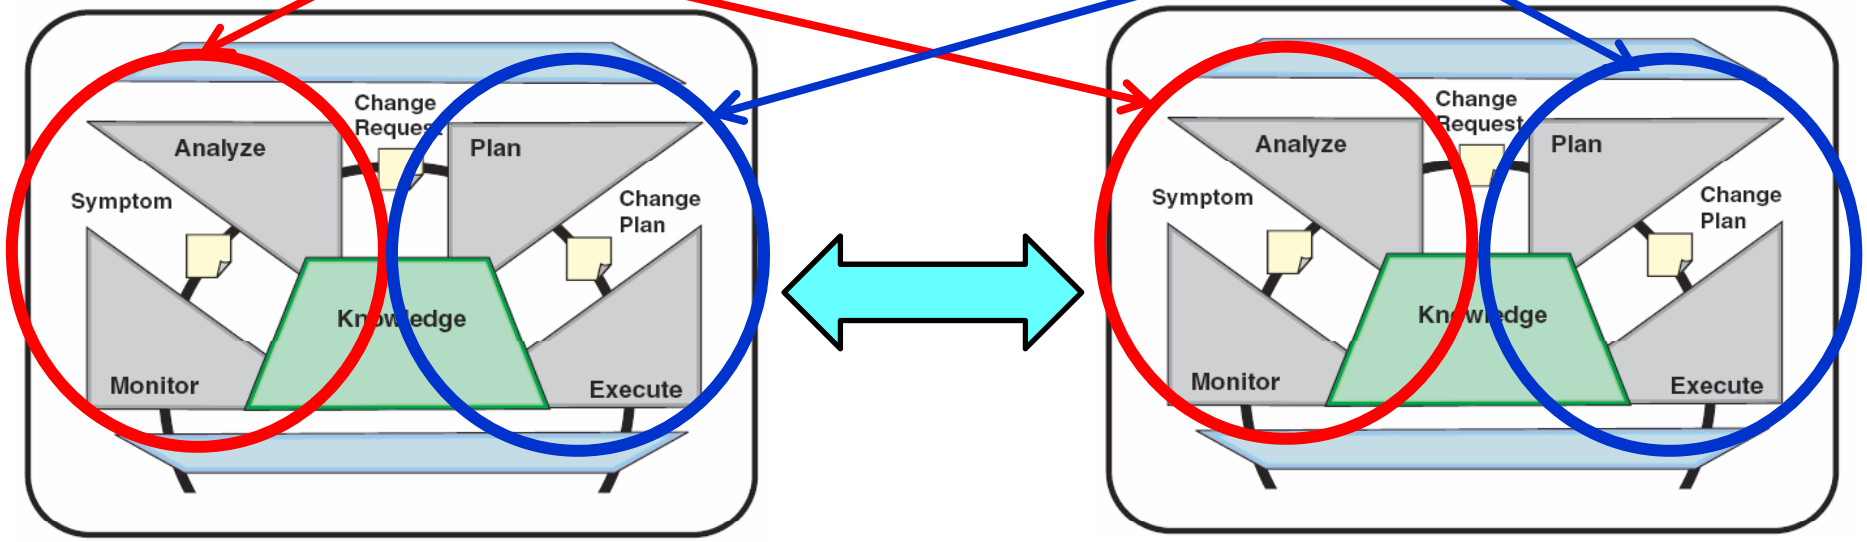
Autonomic Computing [IBM]: MAPE
Monitor-Analyze-Plan-Execute

© 2010 Fraunhofer FOKUS, T. Zseby

Enabler: Node Collaboration

Improve
Situation Awareness

Enforce
Joint Strategy



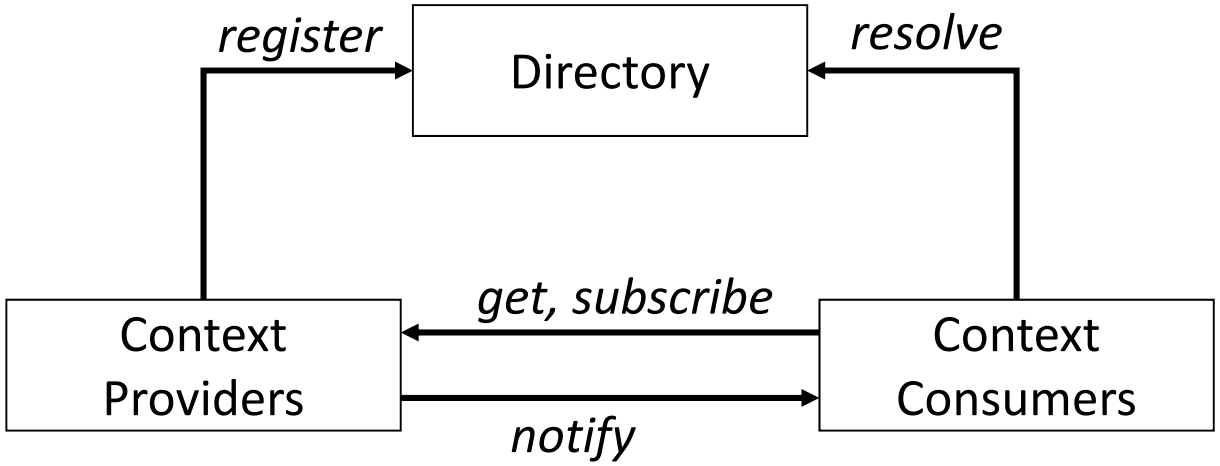
Provide Different Viewpoints

Learn from Others

© 2010 Fraunhofer FOKUS, T. Zseby

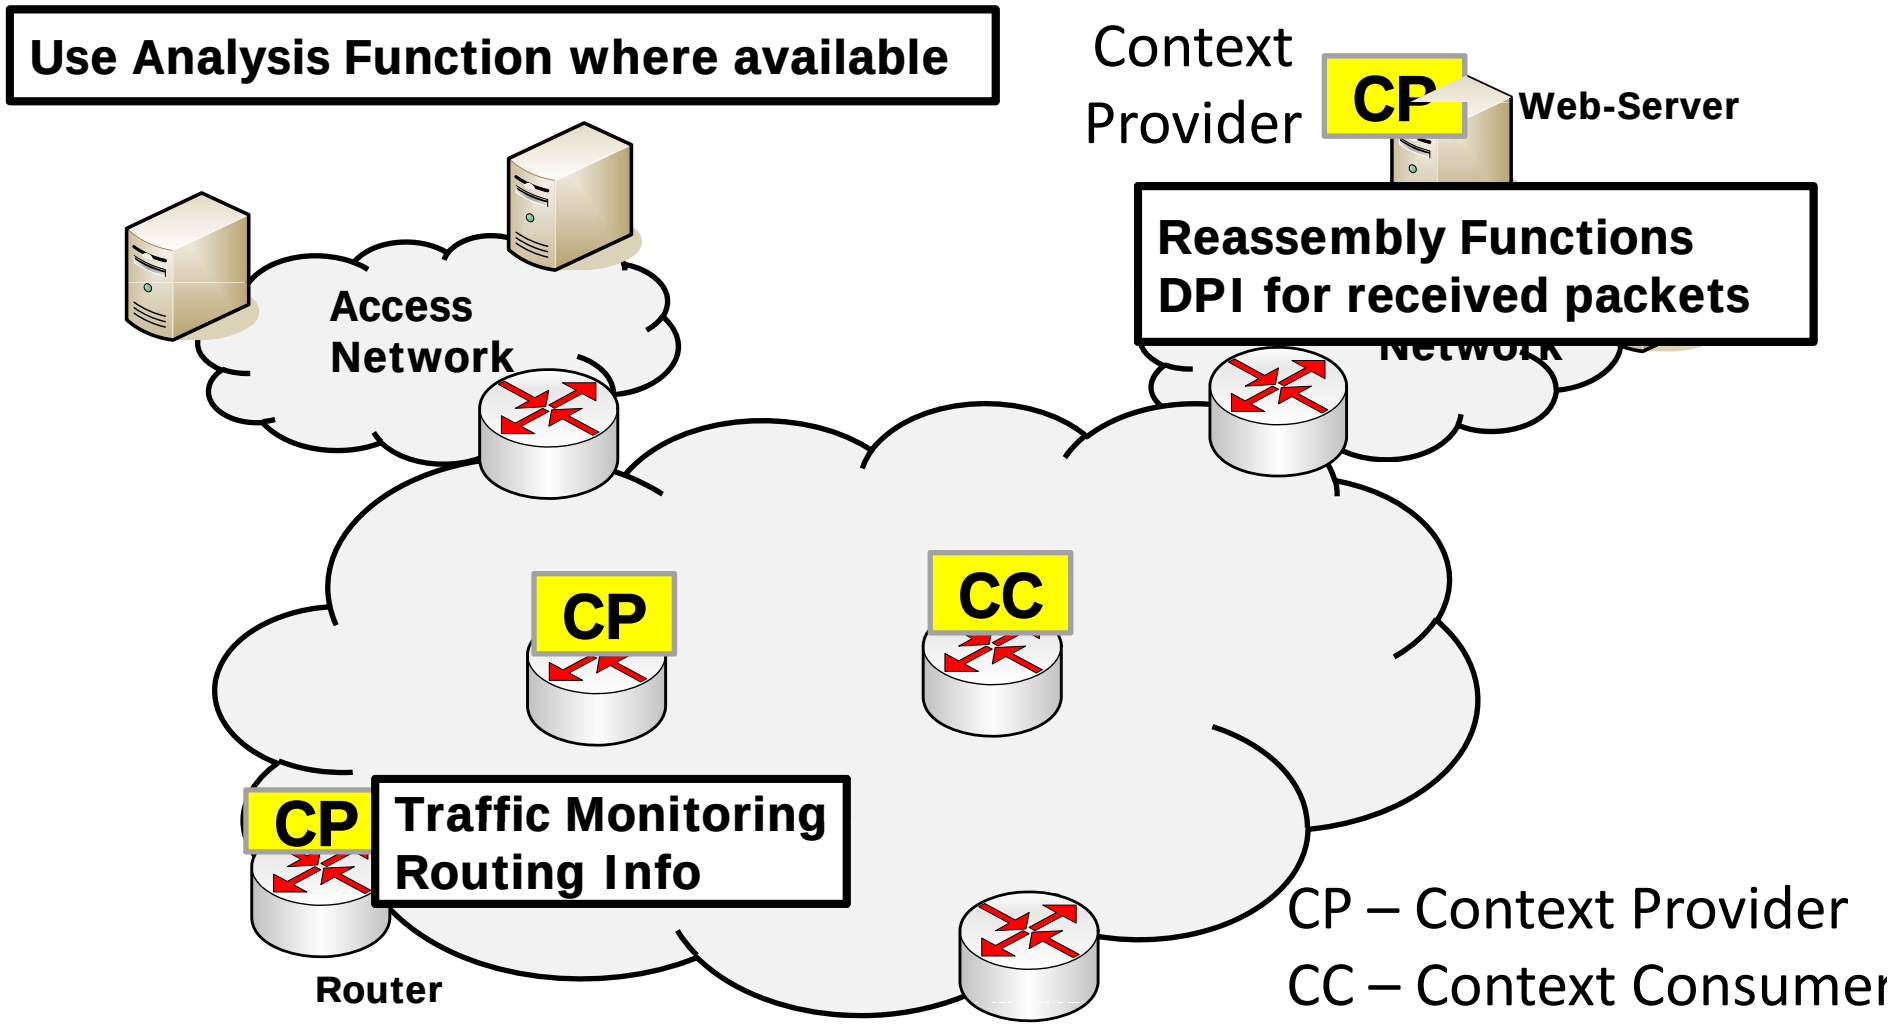
Node Collaboration System (NCS)

- Nodes that want to collaborate
  - Register with directory
  - Provide location of information they can offer (context provider)
  - Can get location of network information from others (context consumer)



© 2010 Fraunhofer FOKUS, T. Zseby

Node Collaboration System



© 2010 Fraunhofer FOKUS, T. Zseby

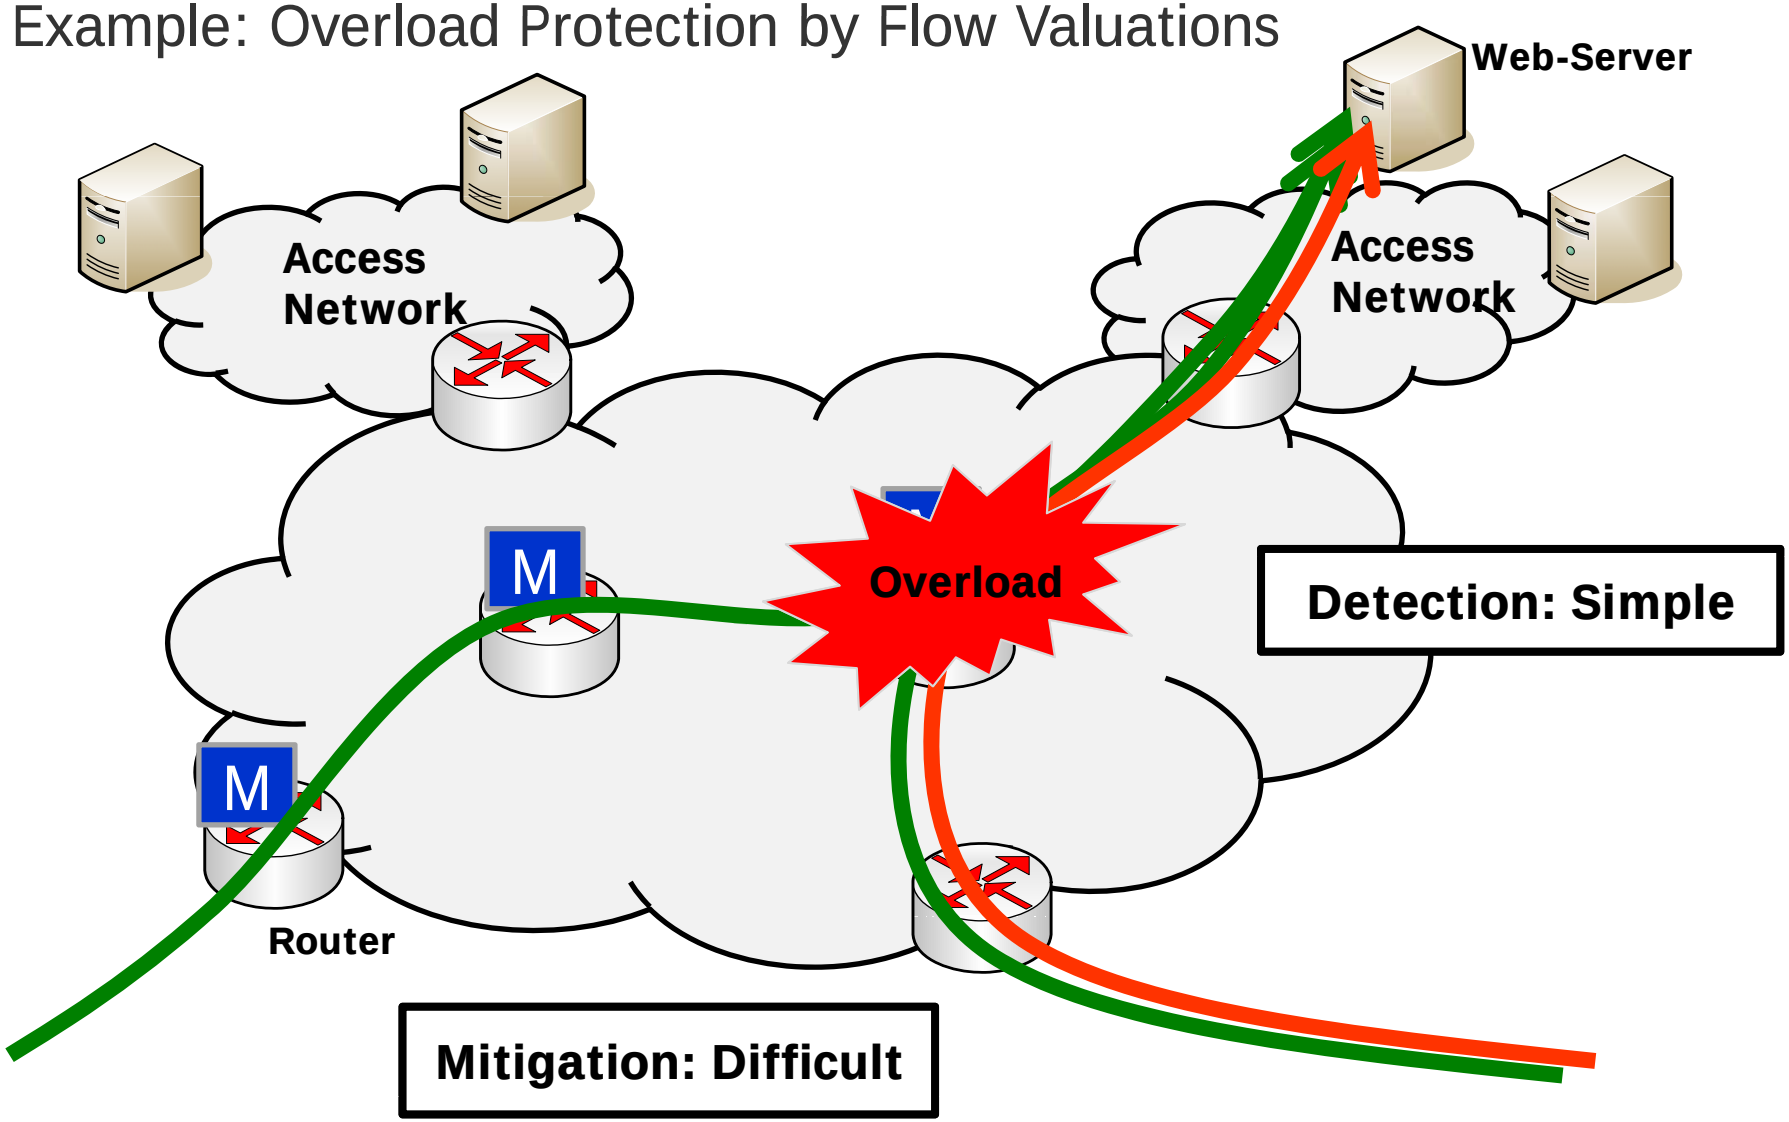
Implementation

- Registration and Resolve
  - HTTP–like Protocol
  - Unified Resource Identifier for addressing info
- Actual Information transfer with different protocols
  - Push (subscribe to CP) or pull (explicitly request info)
  - IP Flow Information Export (IPFIX)
- Additional support for
  - Invoking new measurements
  - Artificial context providers (situation generation for assessing new decision algorithms)

J. Tiemann, D. Witaszek: Context Coordination and Dissemination System – Architecture and Basic Implementation, FOKUS Technical report TR-2008-0303-CCDS-Basics



Example: Overload Protection by Flow Valuations



© 2010 Fraunhofer FOKUS, T. Zseby

Objectives

- Protect access links and servers
  - Detect and mitigate overload situations in the Network
  - Mitigate before access or servers are affected
- Reduce collateral damage
  - Important traffic should be protected
  - Unimportant or malicious traffic can/should be filtered
- Mitigation as close to the originator as possible
  - Protect core from unwanted traffic

- Avoid:
  - Providers to decide about importance of traffic (net neutrality)
  - Deep Packet Inspection (high effort)

© 2010 Fraunhofer FOKUS, T. Zseby

Challenge

- Overload may be originated by
  - Legitimate traffic (Flash Crowds)
  - Abusive DDoS
- Goal: Reduce traffic to prevent overload
  - Cause may be unknown
  - But: reduce collateral damage

Challenge: Which flows should be blocked?

© 2010 Fraunhofer FOKUS, T. Zseby



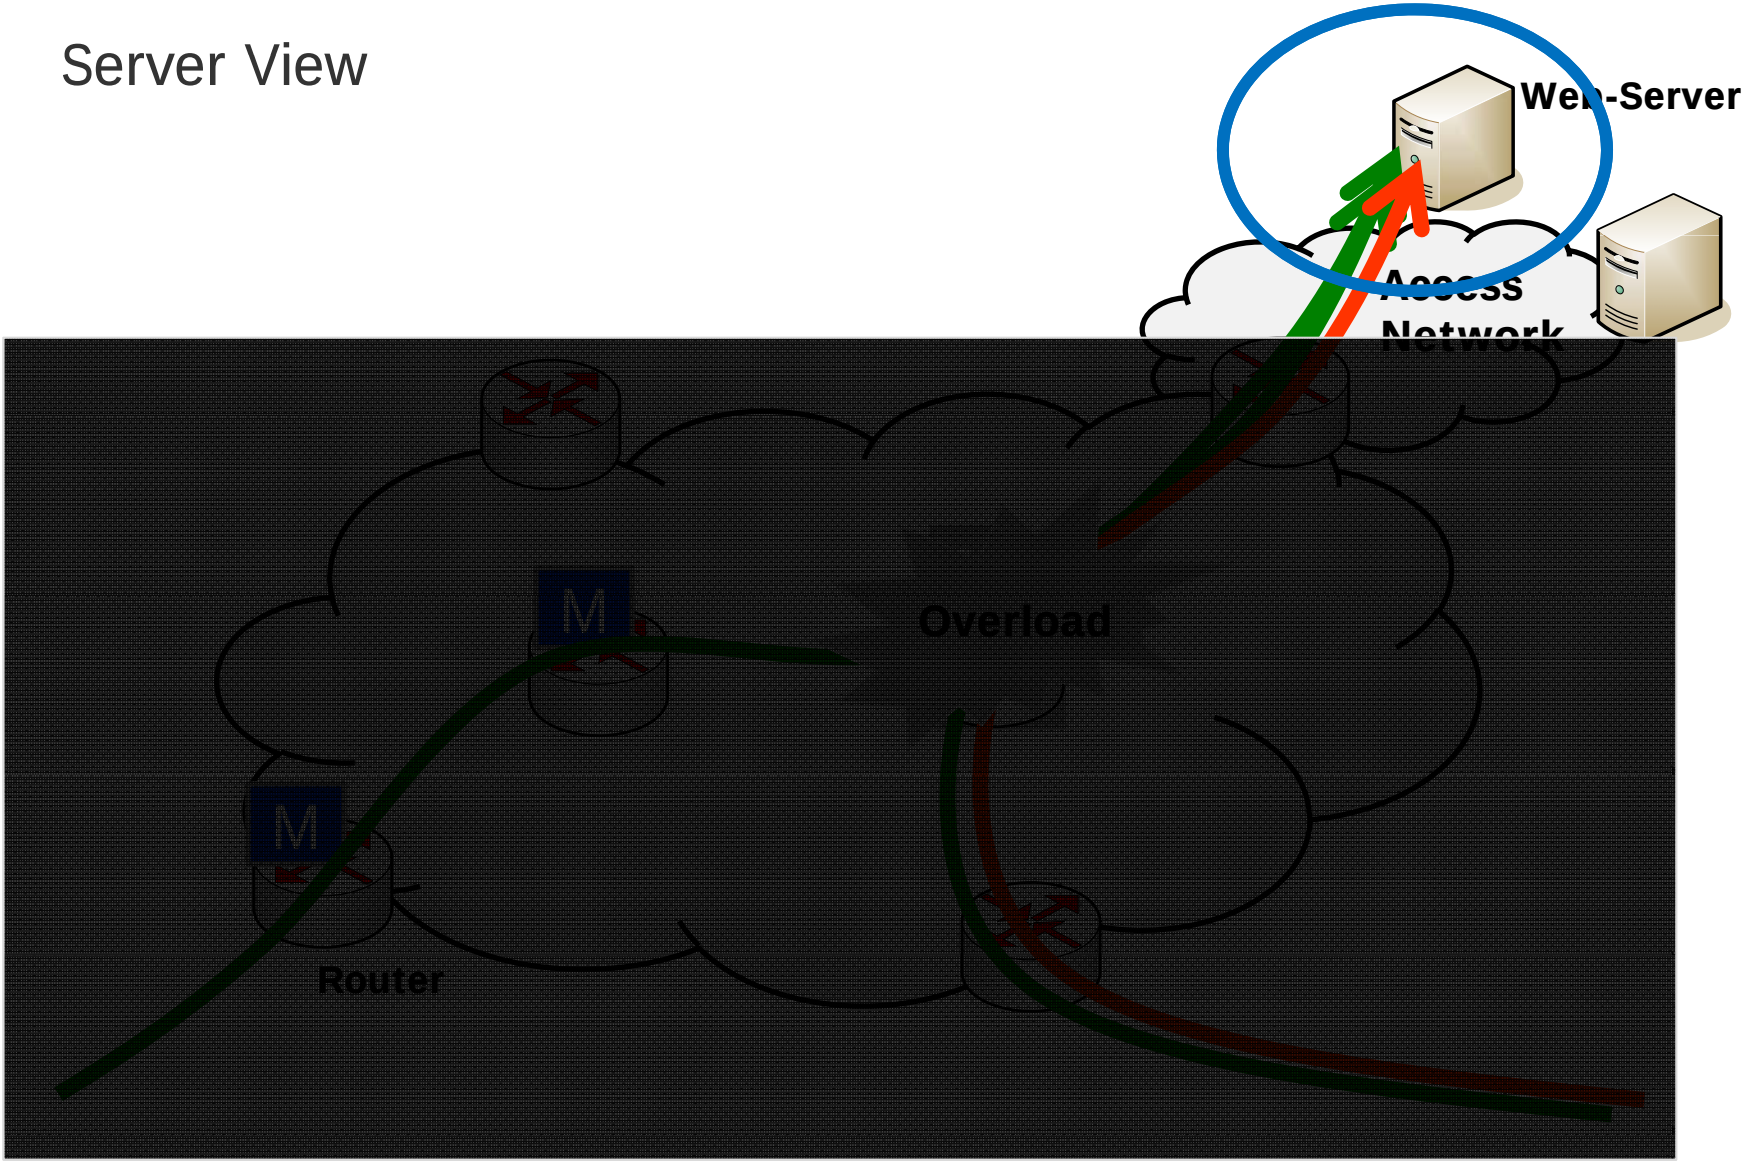
Conventional Solutions

- Random Filtering
  - No separation between good and bad packets
  - Problem: High collateral damage (blocking of legitimate users)
- Intrusion detection systems
  - Attempt to separate good and bad packets by Inspection (DPI)
  - Problem: High Effort in Network

- Alternative: Network-Service Cooperation
  - Cooperate to establish situation awareness
  - Utilize analysis functions at application level
  - Combine information from multiple cooperative sources
  - Combine preferences to form a joint decision

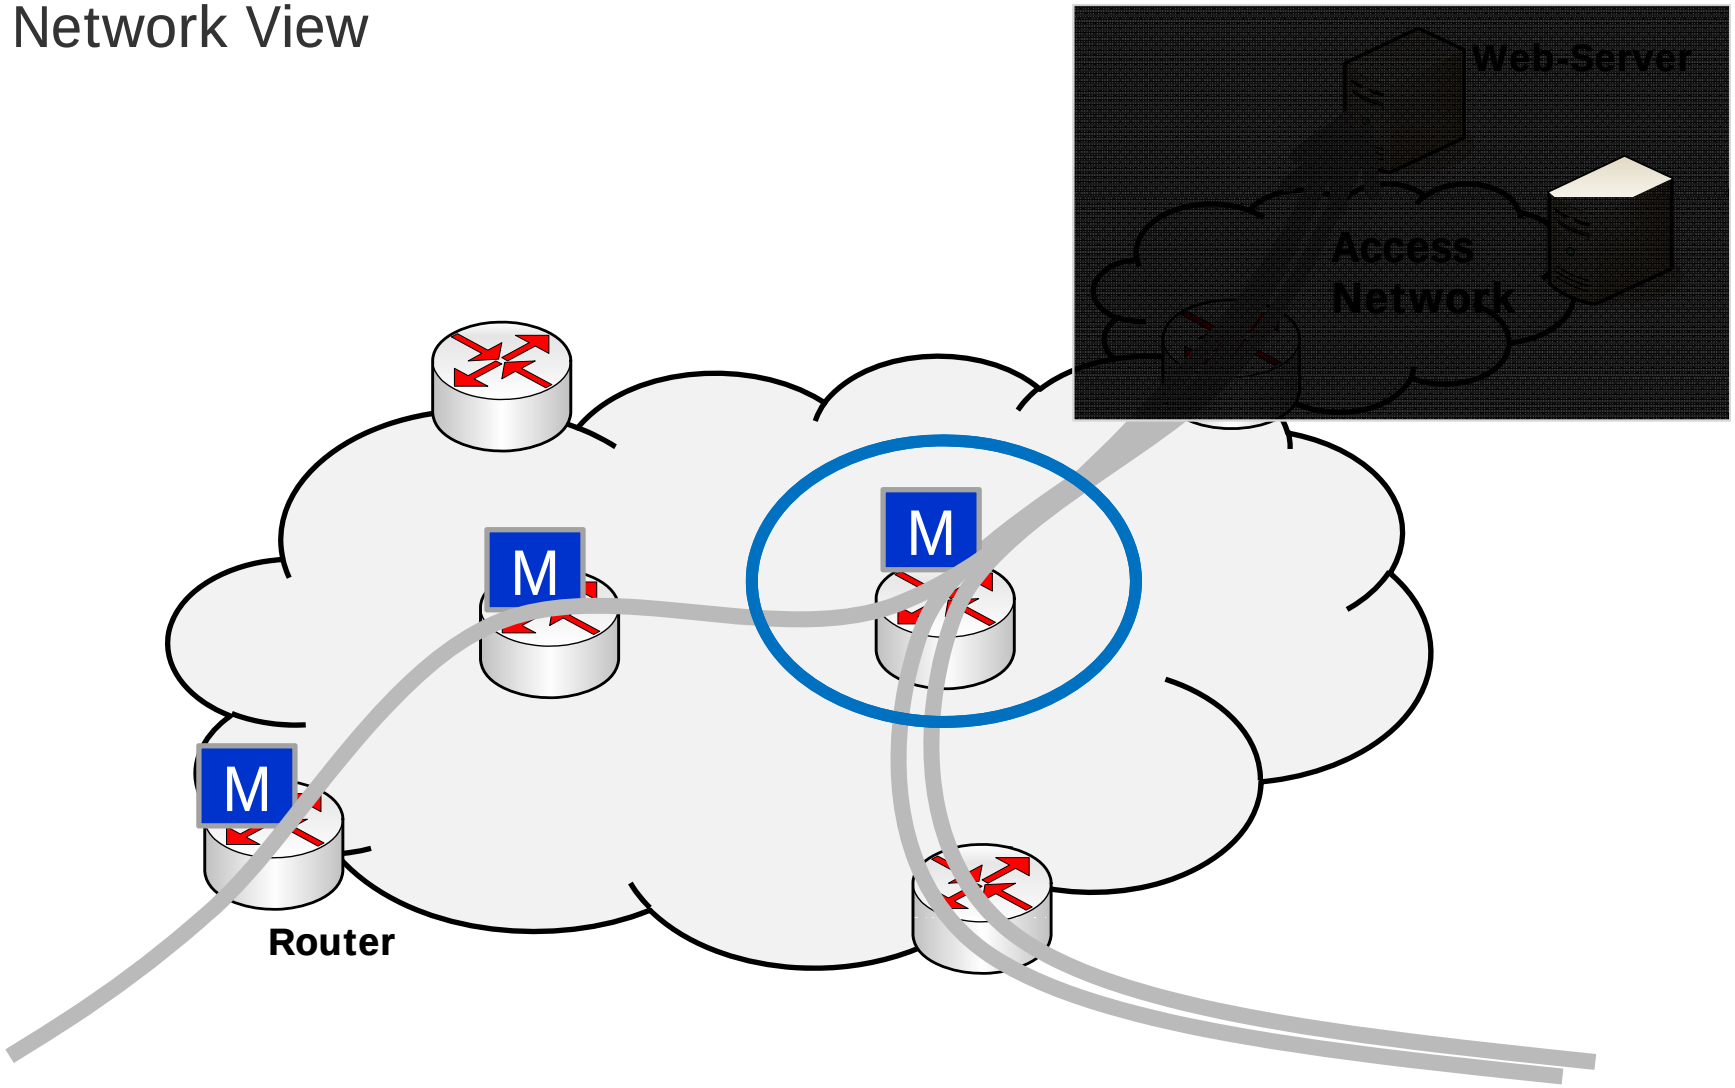
© 2010 Fraunhofer FOKUS, T. Zseby

Server View



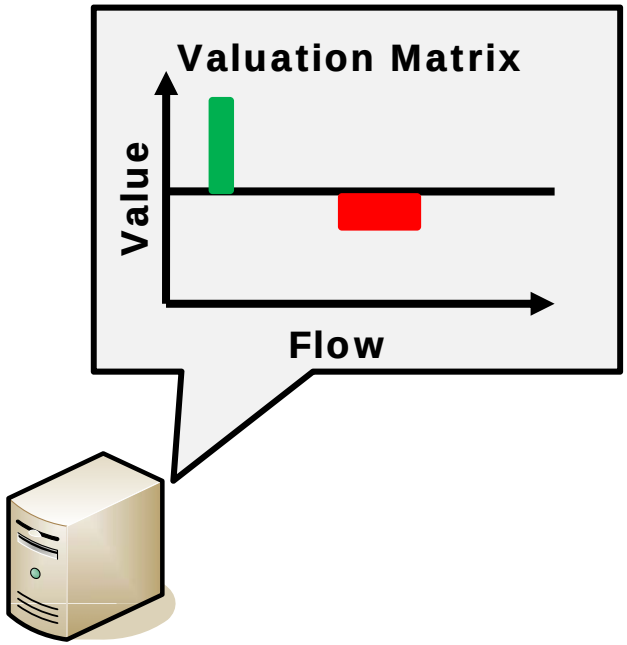
© 2010 Fraunhofer FOKUS, T. Zseby

Network View

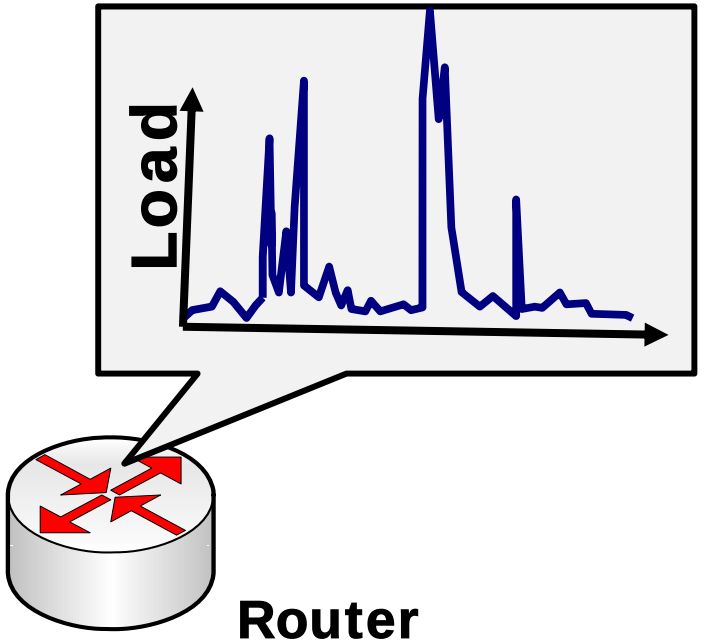


© 2010 Fraunhofer FOKUS, T. Zseby

Approach: Network-Service Cooperation



Server



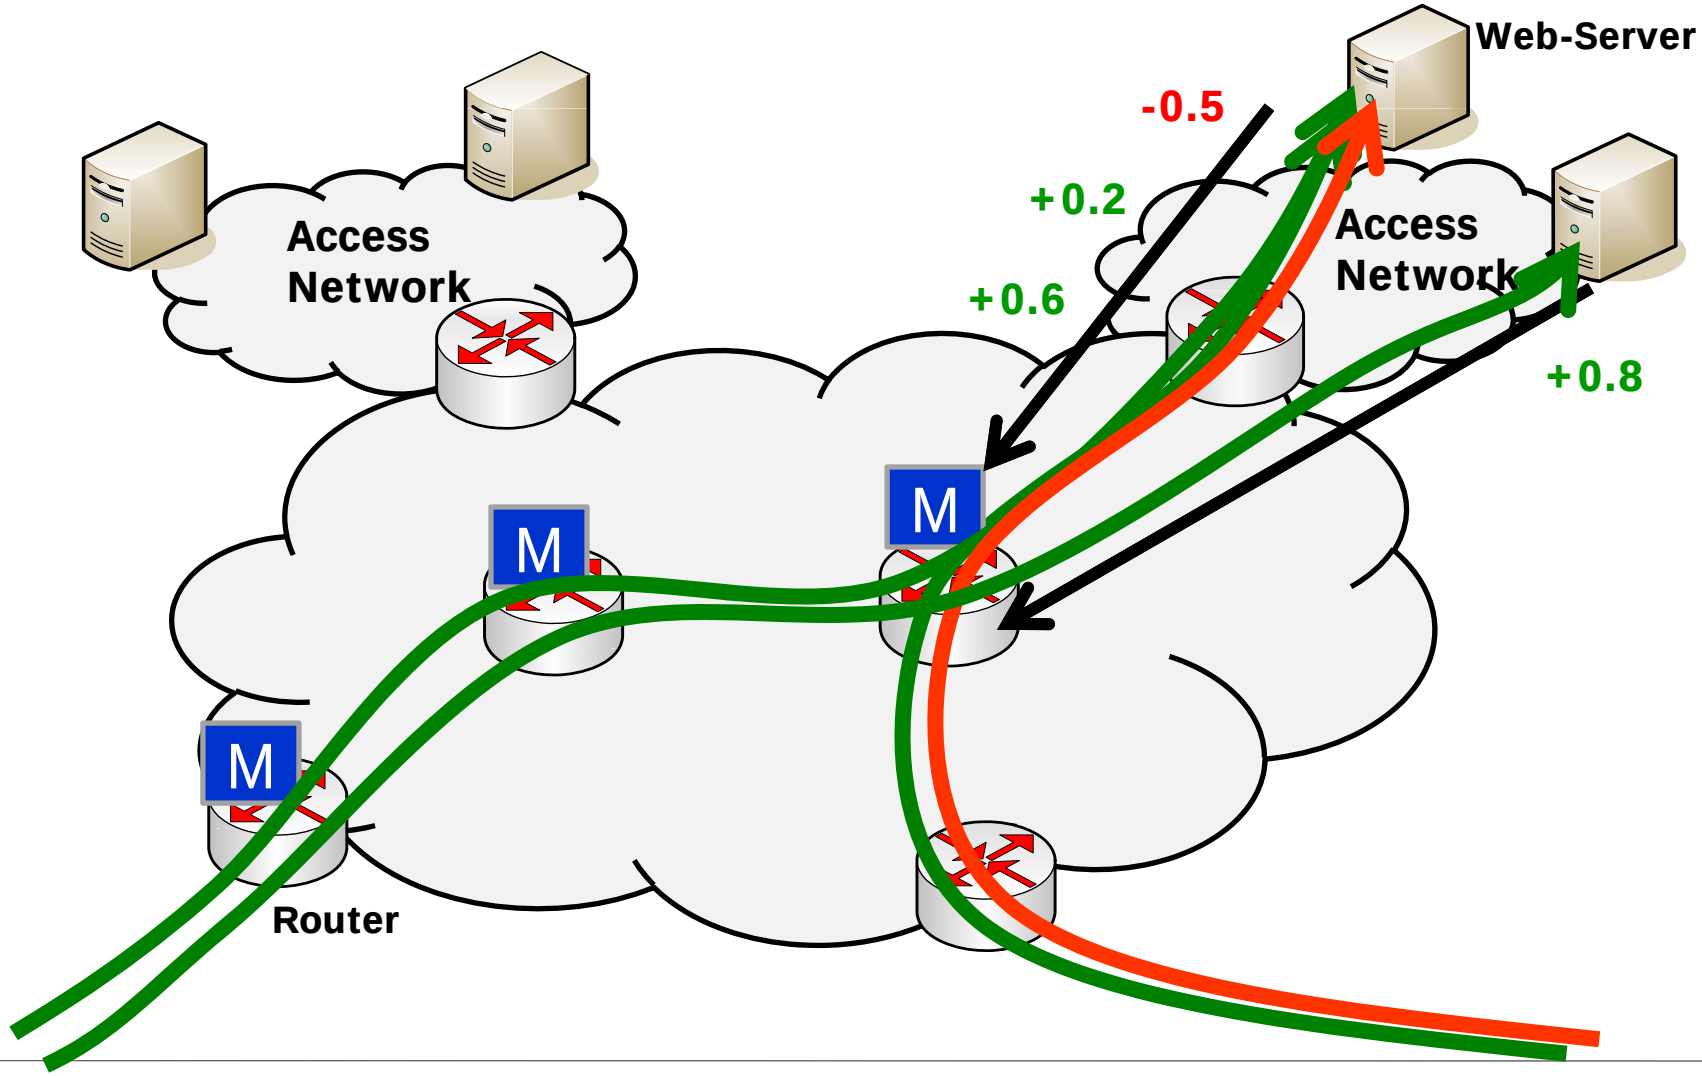
Router

Servers
- know the value of their user traffic

Routers
- Can detect overload
- Can filter before traffic enters access links
- But don’t know the traffic value

© 2010 Fraunhofer FOKUS, T. Zseby

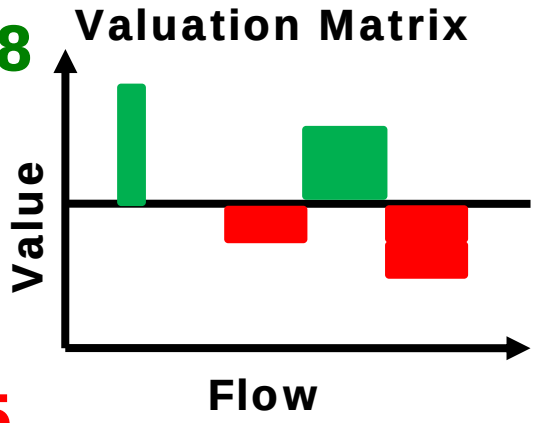
Normal Operation: Servers provide Valuation Reports



© 2010 Fraunhofer FOKUS, T. Zseby

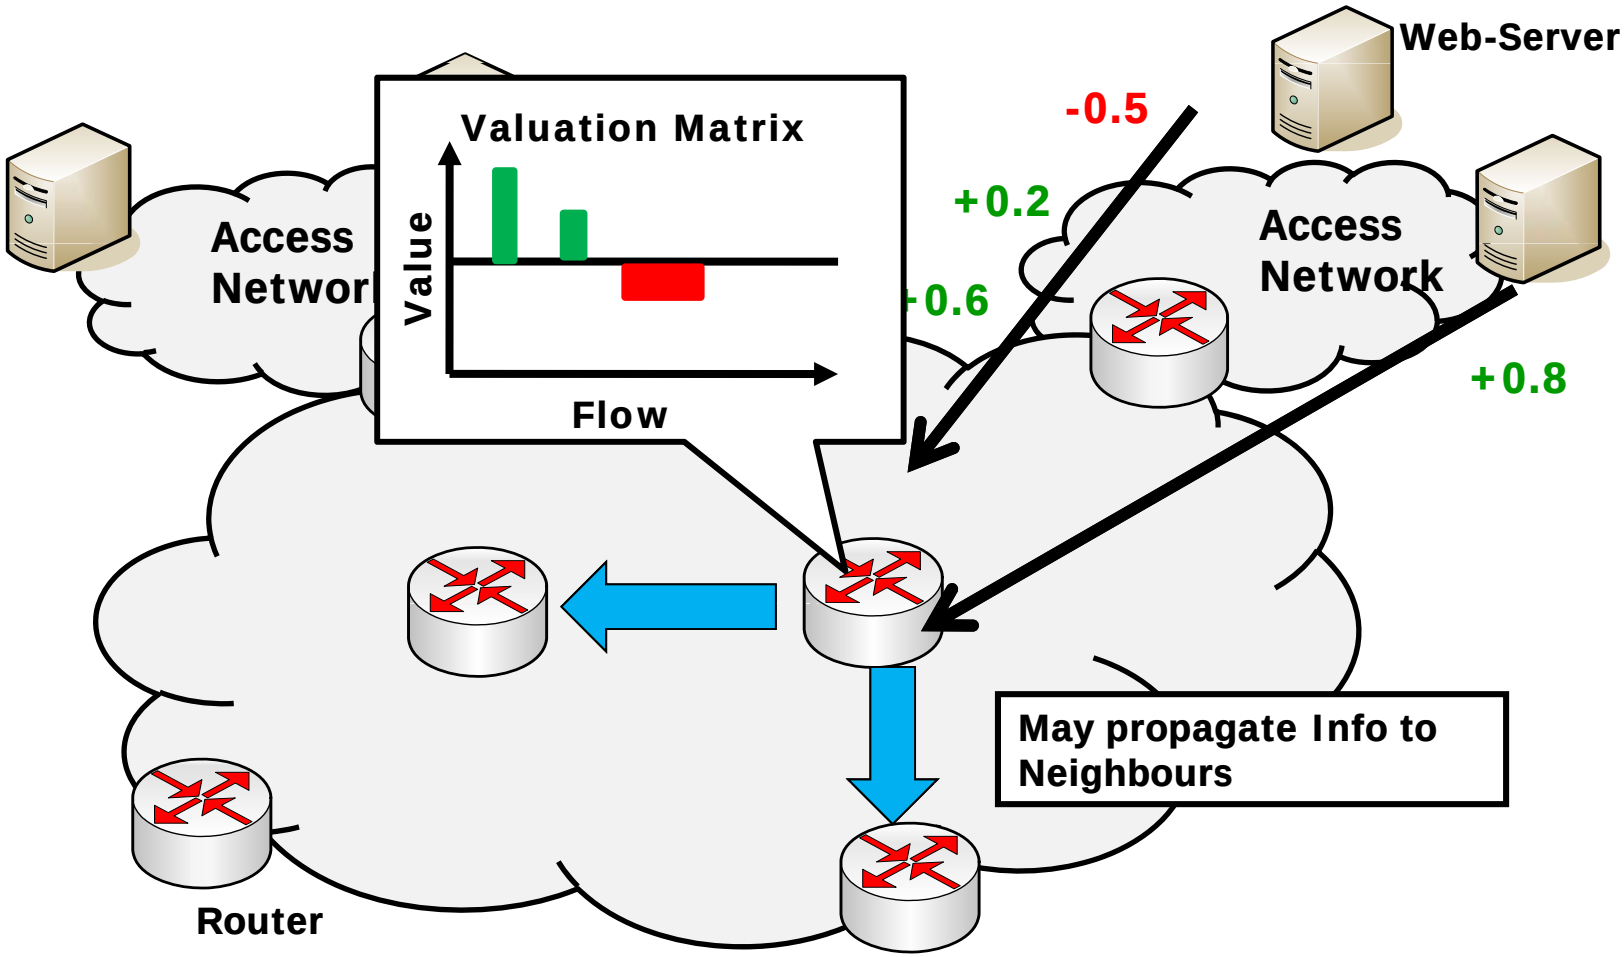
Flow Valuations

- Assessment of user traffic with simple system
- Example valuations
  - User logged in with password → **+0.8**
  - Regular browsing pattern → **+0.5**
  - Customer completed buying → **+1.0**
  - Failed login attempt → **-0.1**
  - Repeatedly loading single site → **-0.5**
  - Attempt to log into database server → **-0.9**



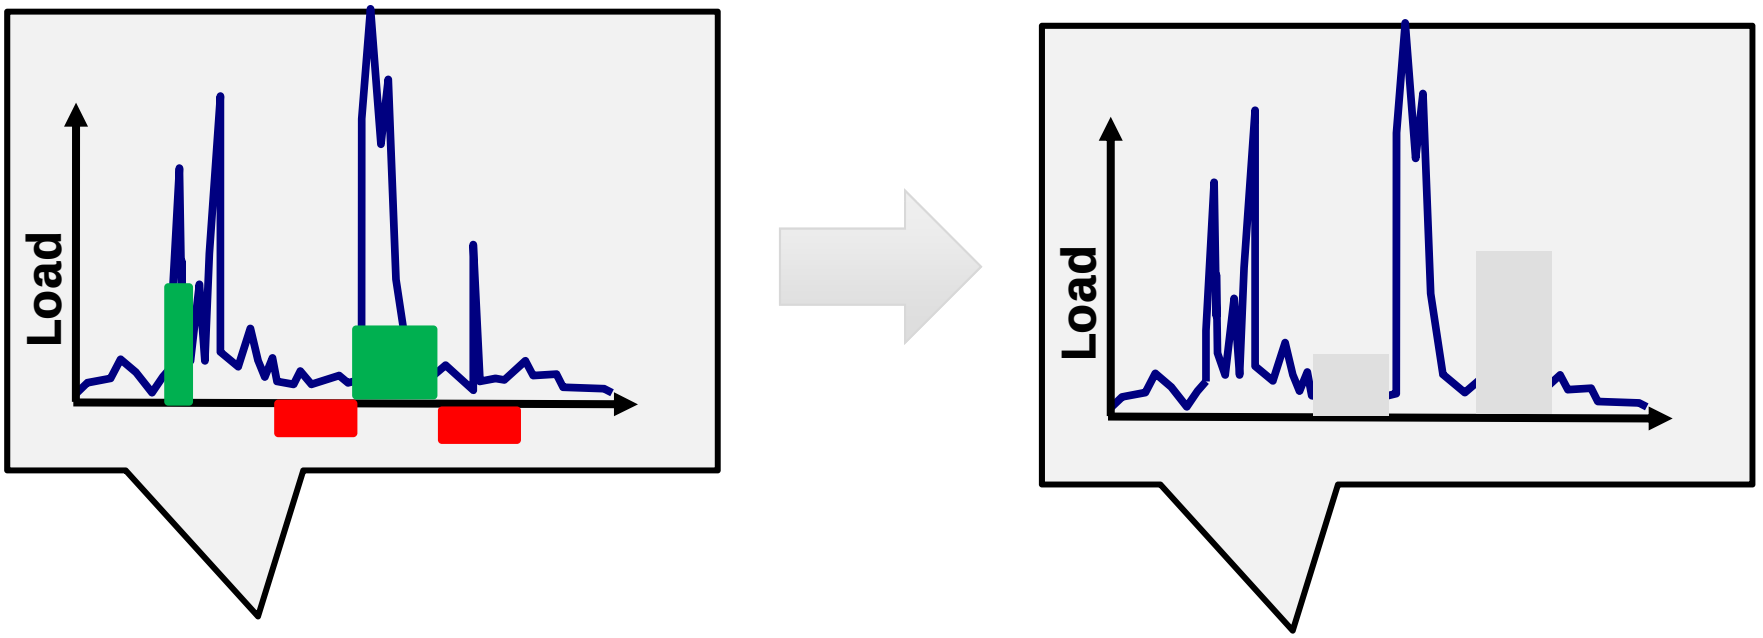
© 2010 Fraunhofer FOKUS, T. Zseby

Normal Operation: Router aggregates Valuations



© 2010 Fraunhofer FOKUS, T. Zseby

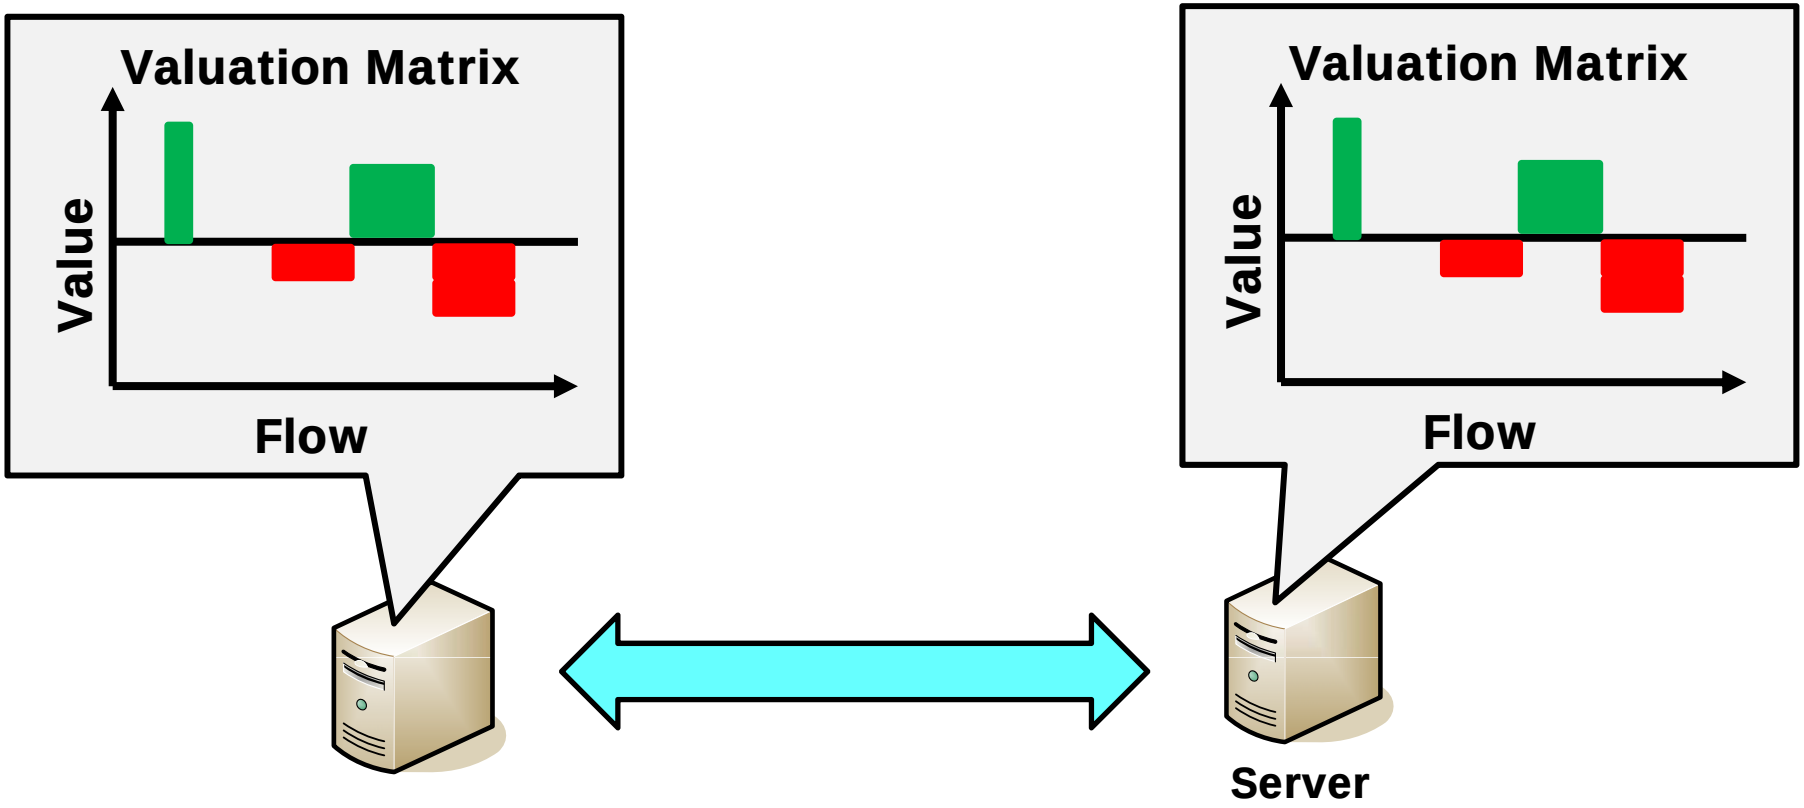
Overload Situation: Correlate Information



Blocking Decision based on aggregated valuations

© 2010 Fraunhofer FOKUS, T. Zseby

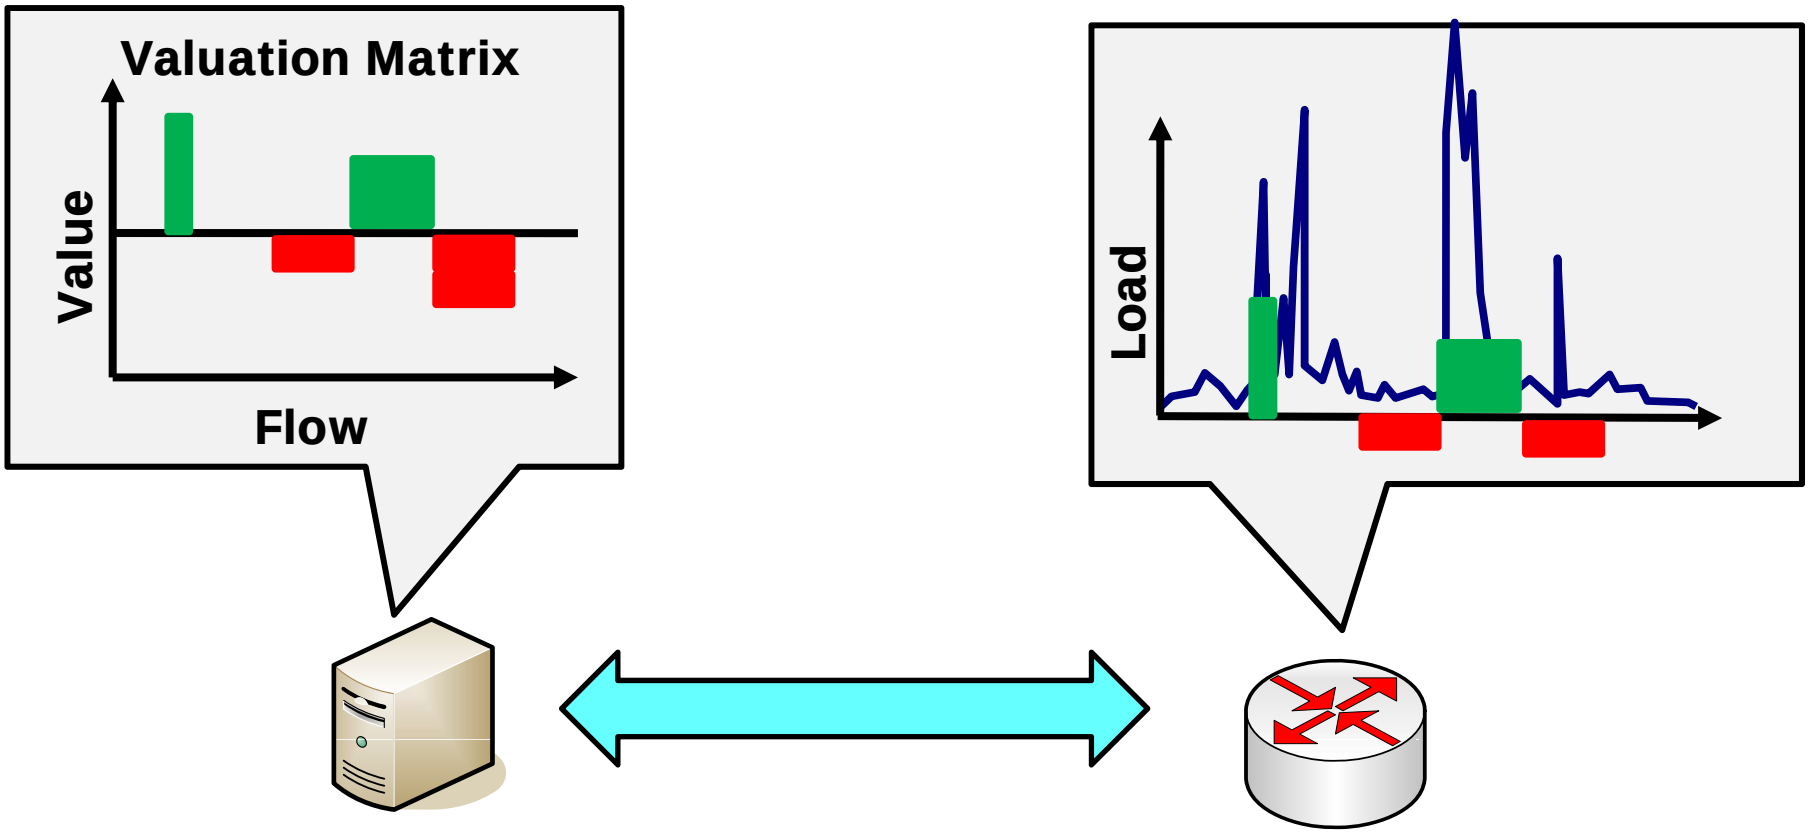
Good: Collaboration of Nodes



Learn about traffic value

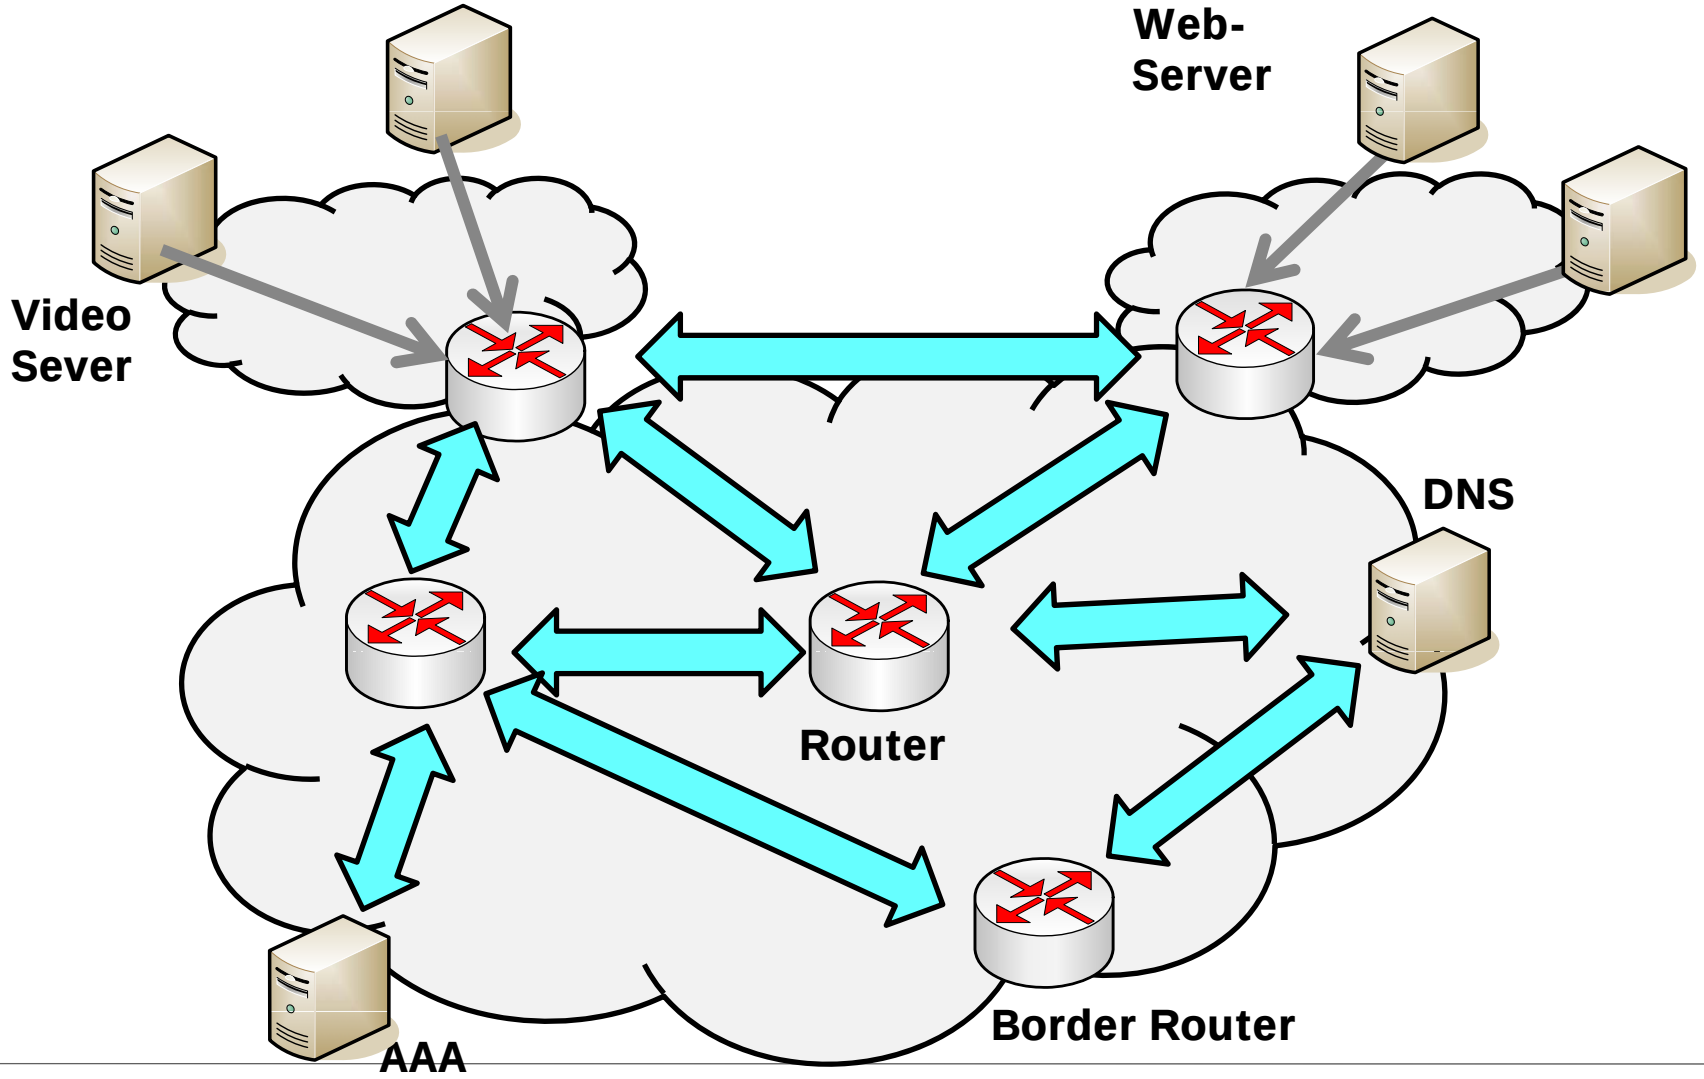
© 2010 Fraunhofer FOKUS, T. Zseby

Better: Collaboration of Different Nodes



© 2010 Fraunhofer FOKUS, T. Zseby

Unbeatable: A Heterogeneous Collaborative Team



© 2010 Fraunhofer FOKUS, T. Zseby

Status and Future Work

- Node Collaboration System for Flow Valuations
  - Implemented on Cisco AXP routers
  - Using IPIFIX for information transfer
- Cooperation Incentives (in progress)
  - Incentives to provide information
  - Inter-domain exchange
- Integration of further sources
  - Information from sophisticated sources (IDS, AAA,...)
  - Worm detection information from enhanced DNS

© 2010 Fraunhofer FOKUS, T. Zseby

Thank You!

Contact: *tanja.zseby@fokus.fraunhofer.de*

Fraunhofer FOKUS

© 2010 Fraunhofer FOKUS, T. Zseby



Flocon2010 – January 12th, 2010

Introduction to SIE (condensed)

Eric Ziegast
<info@sie.isc.org>

Copyright © 2010 Internet Systems Con



Security Information Exchange

Raison d'être

Providing common legal and privacy framework for sharing sensitive data

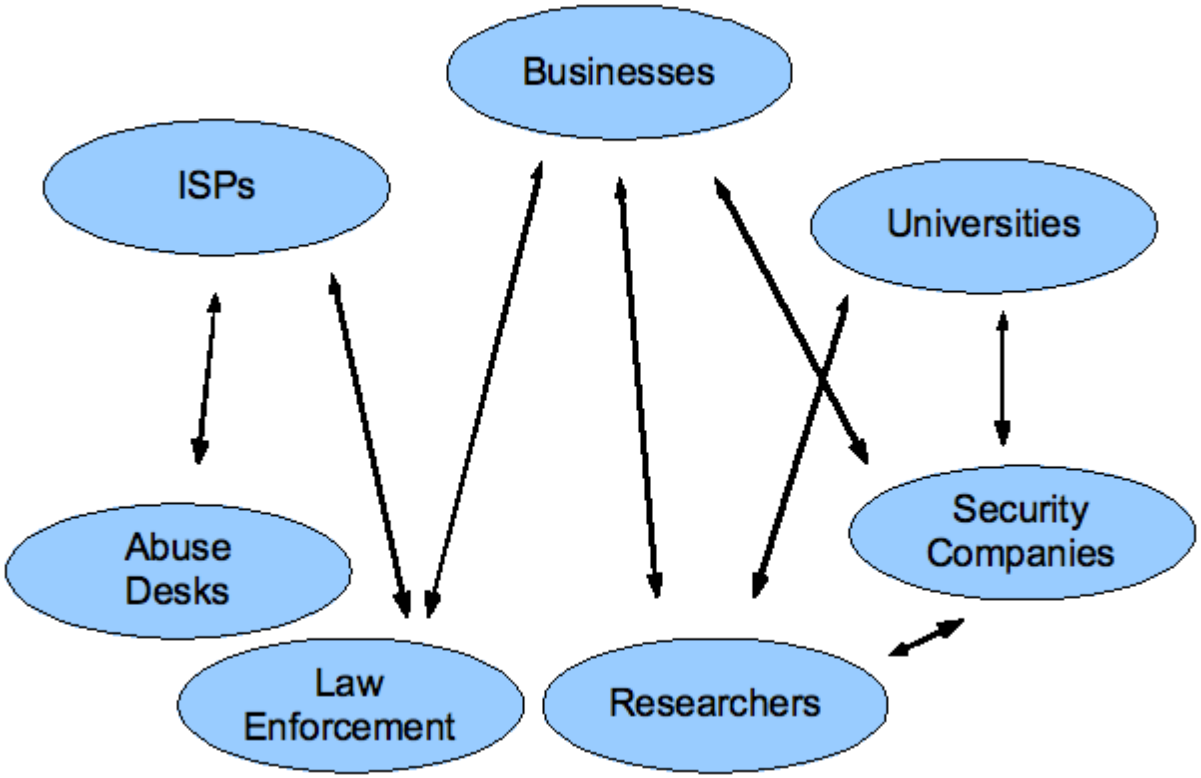
Centralizing security data collection and distribution to bring real-time efficiencies to analysis

Enabling cross-analysis between disparate data sets

Creating network effect between participants (stone soup)

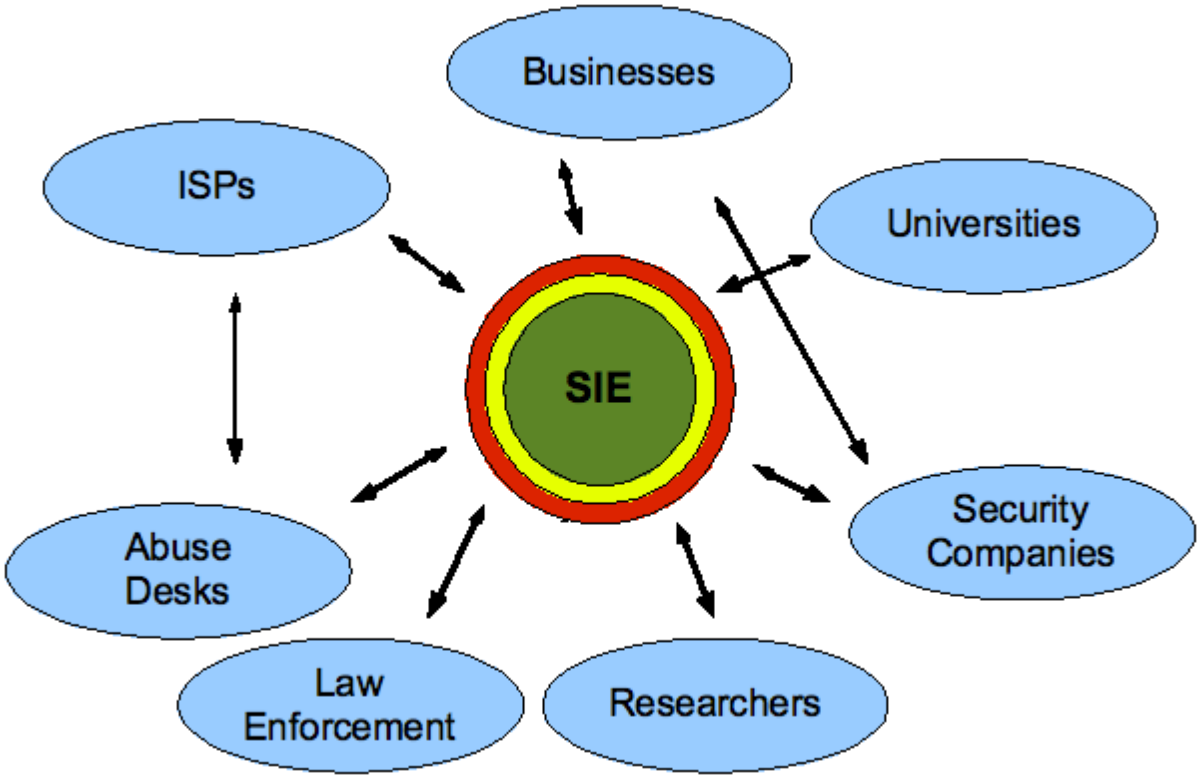
Copyright © 2010 Internet Systems Consortium

Decentralized - bi-lateral



Copyright © 2010 Internet Systems Consortium

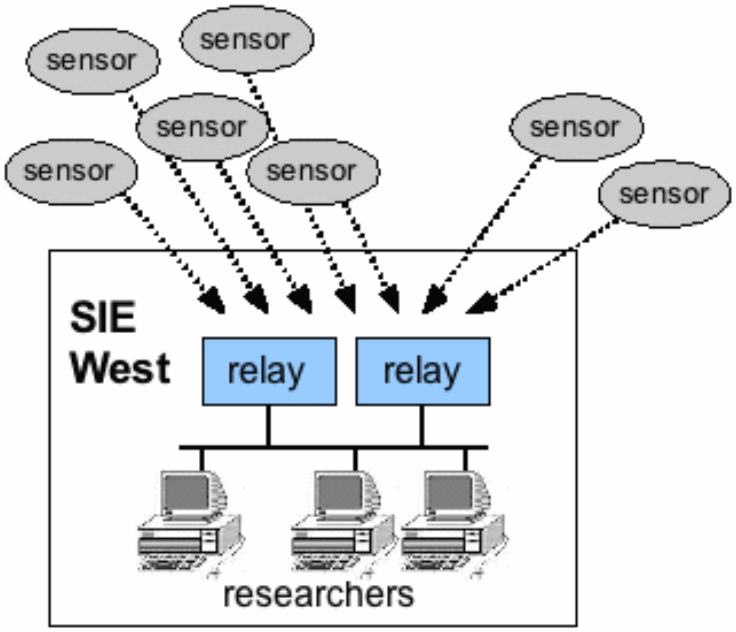
Centralized - multi-lateral



Efficient sharing within common legal/privacy framework

Copyright © 2010 Internet Systems Consortium

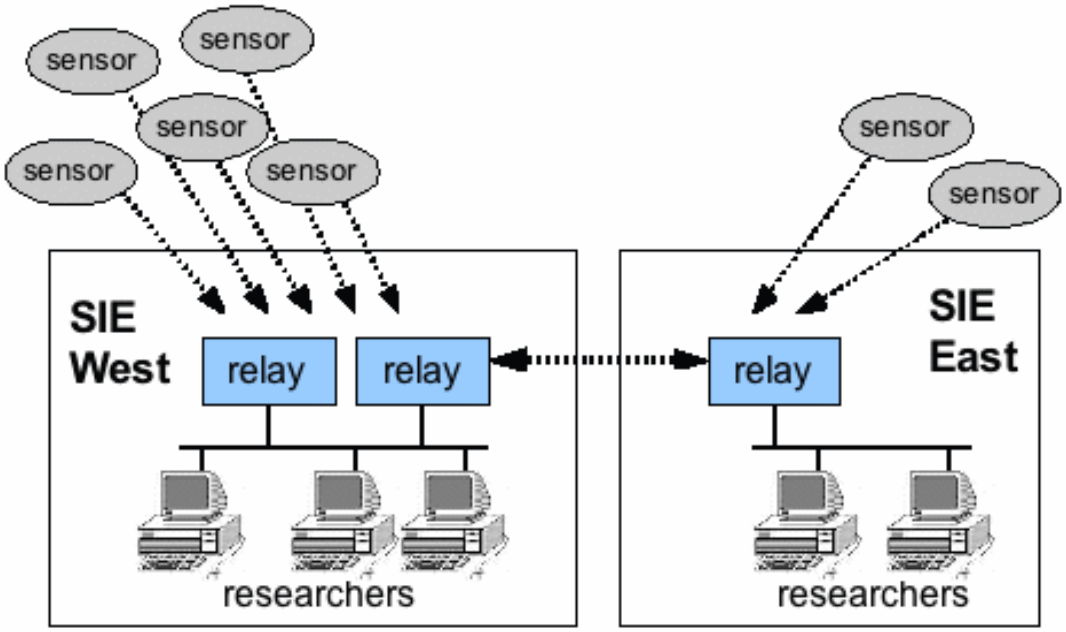
Data distribution model - today



- SF Bay Area, US (PAIX)
- Main sensor relays
- Some researchers getting feeds off switches

Copyright © 2010 Internet Systems Consortium

Data distribution model - east

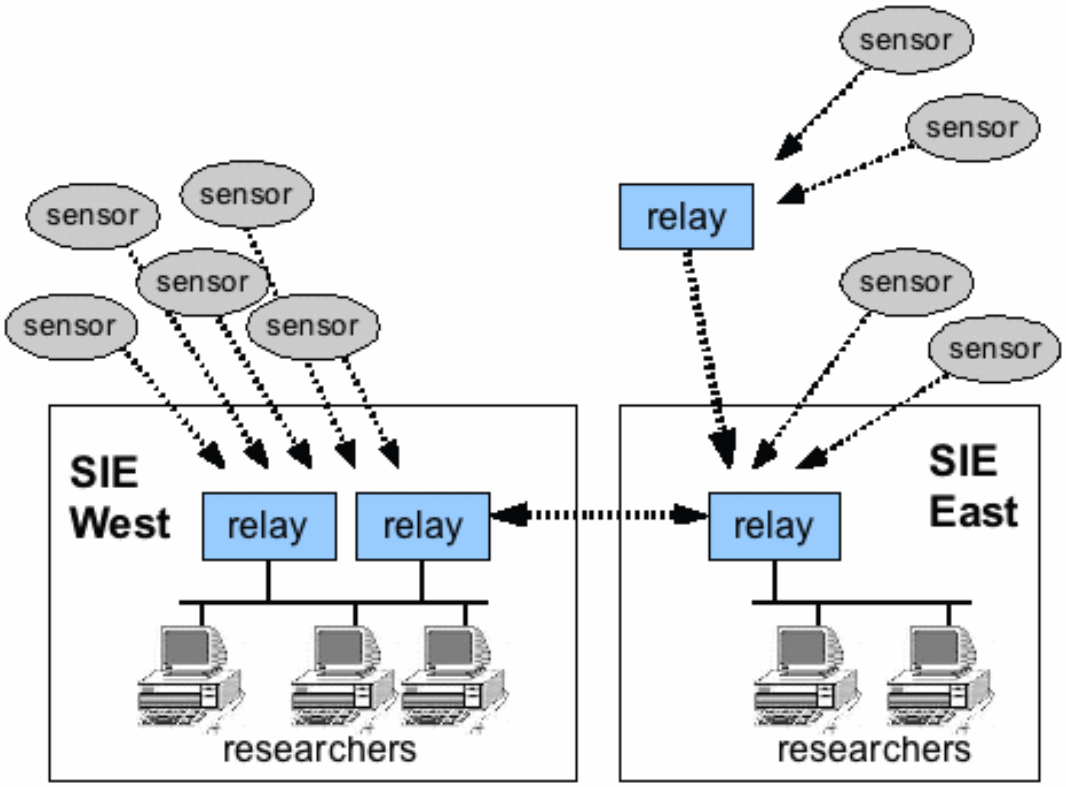


- NYC, US
- Redundant facilities
- More researchers

SRV load balancing

Copyright © 2010 Internet Systems Consortium

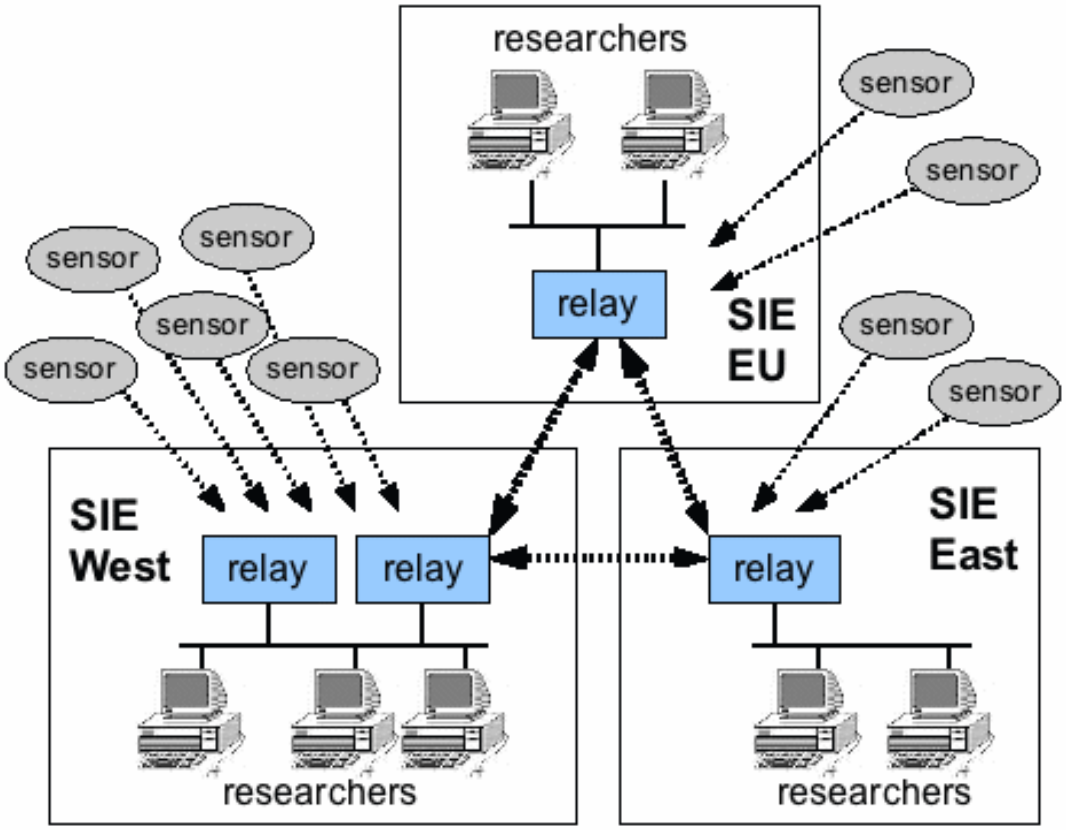
Data distribution model - relay



- Add relays at exchanges in different countries
- Add local sensors
- Local sharing or tools possible within relay

Copyright © 2010 Internet Systems Consortium

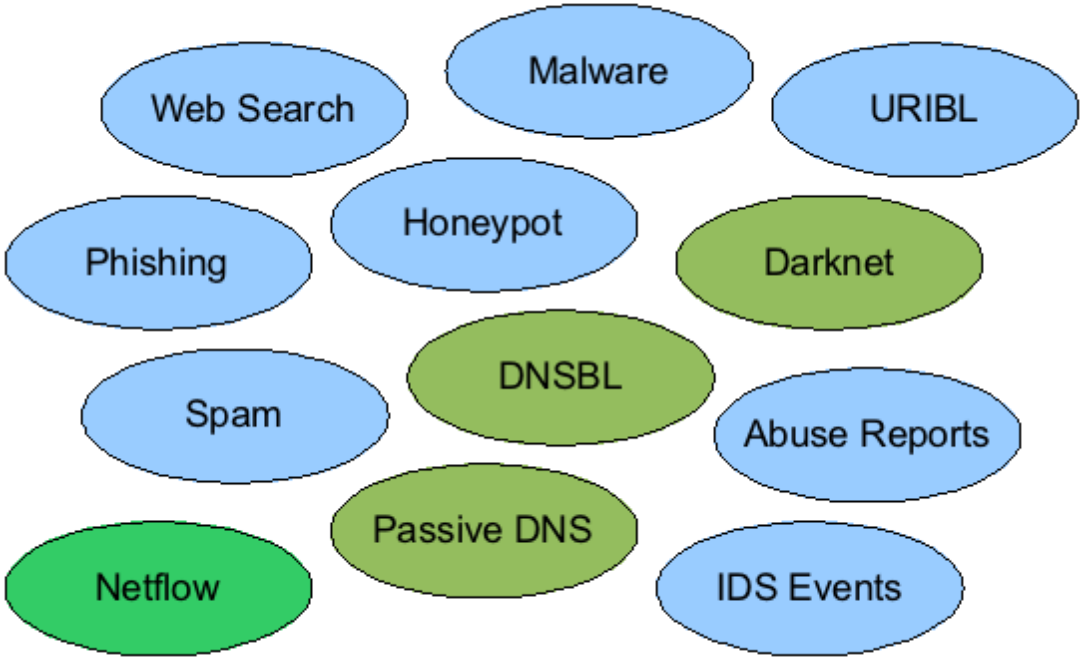
Data distribution model - promote



- Promote node when number of researchers is significant
- Scaling issues

Copyright © 2010 Internet Systems Consortium

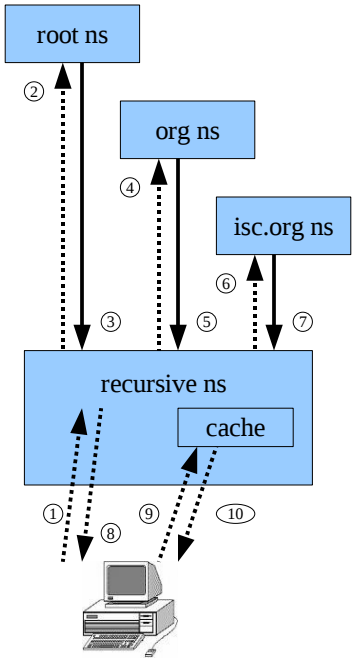
Disparate data



Copyright © 2010 Internet Systems Consortium

pDNS

- SIE started with PassiveDNS in 2007
- Thanks to Florian Weimer (BFK)



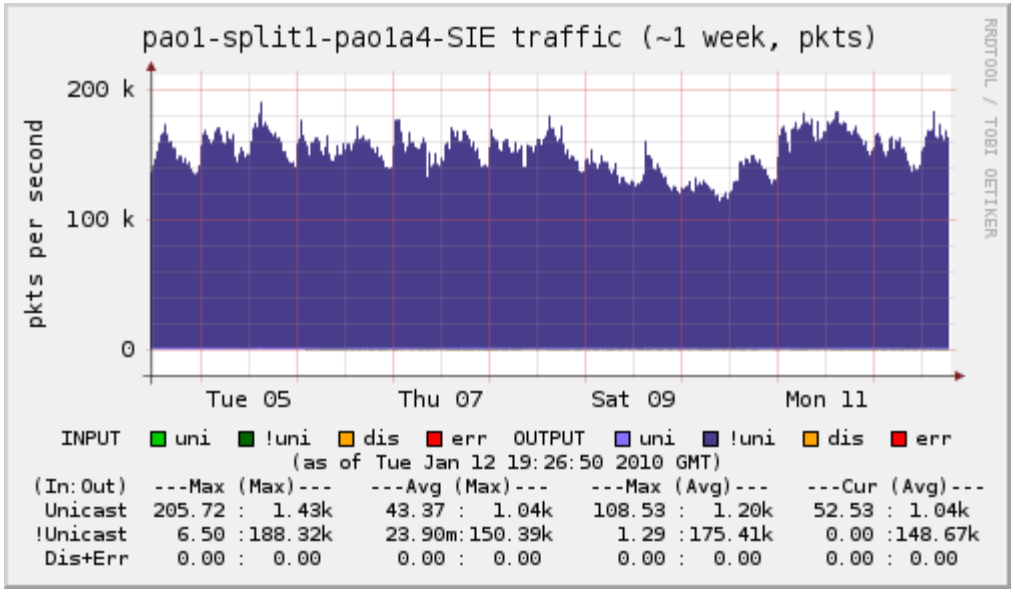
1. client queries for www . isc . org to recursive nameserver
2. recursive nameserver queries root server for “www.isc.org”
3. root nameserver responds with information for “org” nameservers
4. recursive nameserver queries org nameserver for “www.isc.org”
5. org nameserver responds with information for “isc.org” nameservers
6. recursive nameserer asks “isc.org” nameserver for “www.isc.org”
7. isc.org nameserver responds with “www.isc.org” answer
8. answer for “www . isc.org” is returned to client
9. client asks recursive nameserver again for “www.isc.org”
10. the nameserver might serve the answer from cache directly

Only query responses “above the recursive server” are recorded (in steps 3,5,7).

Copyright © 2010 Internet Systems Consortium

Fire hose

150 kpps (not-including sub-packets)



Copyright © 2010 Internet Systems Consortium

Tools

ncap – used primarily for DNS data

plugins – filter data for rebroadcast

nmsg – used to describe any data context

multi-site VPN – services and lookup tols

hardware – high packet rates

fast switch – line rate GigE, jumbo frames

servers – mostly Linux/FreeBSD, 64-bit, multi core, as much RAM as possible

storage – large disk for arrears, SSD

Copyright © 2010 Internet Systems Consortium

ncap

ftp://ftp.isc.org/isc/ncap

Evolution from pcap/dnscap

- Defragmentation
- Drop link layer info
- Normalized network format
- Nanosecond timestamps
- User-defined flags

New features key to SIE

- I/O – File, BPF, Unicast, broadcast, multicast
- Plug-in modules
- dedupe, pattern matching, table lookups

Copyright © 2010 Internet Systems Consortium

Broadcast architecture

Not a database – flow freely

Not a just a packet dump

Real time, not arrears

Loosely coupled multi-processor

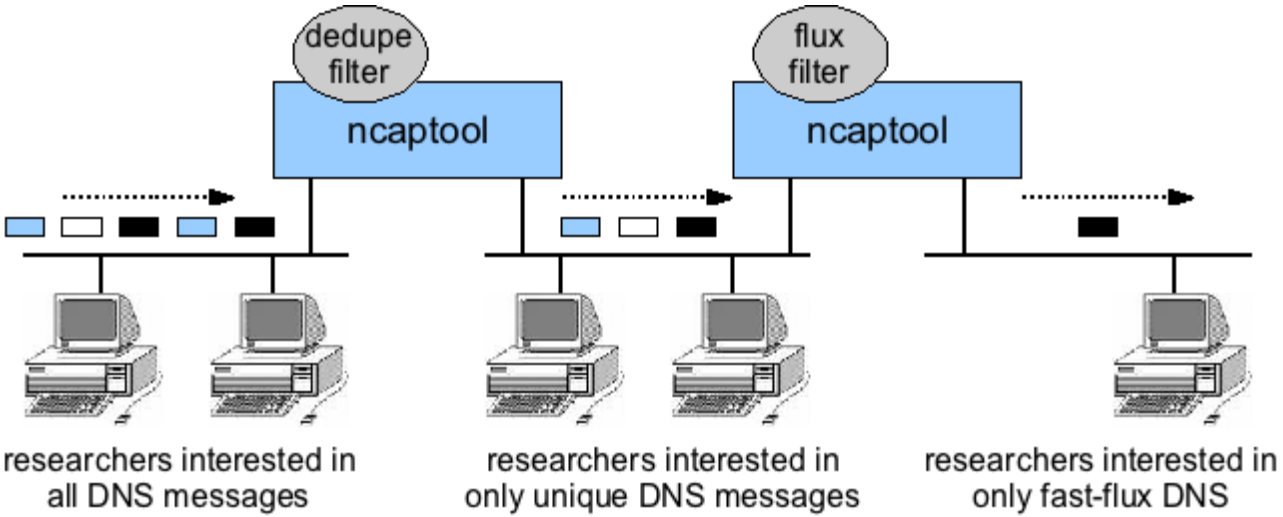
Ethernet switch

Partitions via VLANs or “channels”

Copyright © 2010 Internet Systems Consortium

ncap

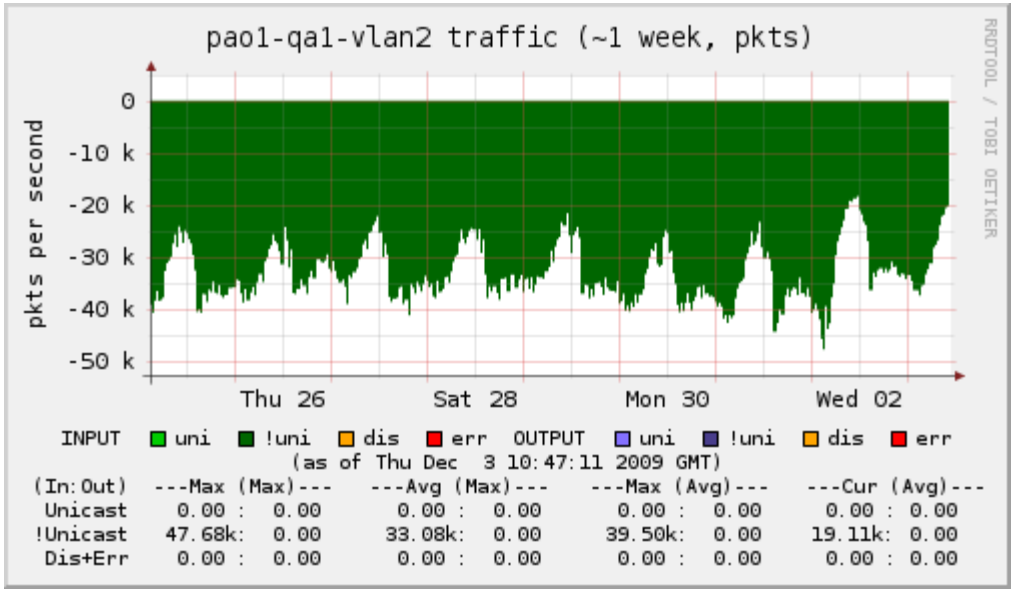
plug-in filters in action



Copyright © 2010 Internet Systems Consortium

filters – raw passive dns

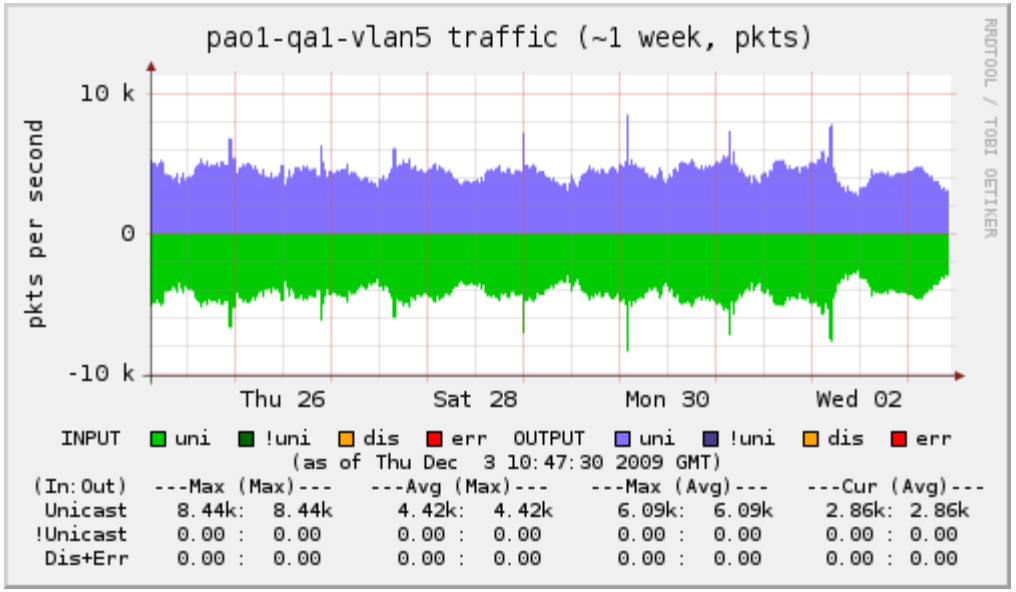
max at 40000 pps – too much



Copyright © 2010 Internet Systems Consortium

filters – deduplicated dns

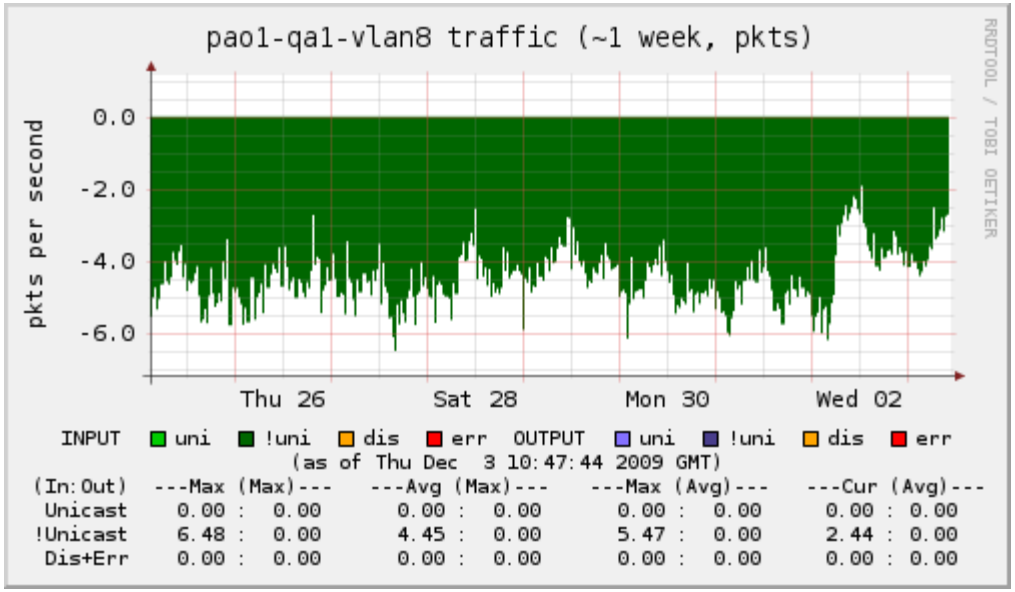
max at 5000 pps - better



Copyright © 2010 Internet Systems Consortium

filters – fast-flux dns

rate is 2-4 pps



Copyright © 2010 Internet Systems Consortium

nmsg

Ncap is great, but what about non-packet data?

Needs:

- represent different data types – extensible
- fast/scalable – jumbo frames
  - Coalescing, Fragmentation
- multiple methods for I/O -
- file, pipe, unicast, broadcast
- filtering methods

Copyright © 2010 Internet Systems Consortium

Channels (nmsg, et.al.)

Spam – yes, we want spam

URL Link Pairs – search engines

Conficker

sinkhole URL, DNS, P2P

https://conficker.sie.isc.org

Mitigation efforts – other botnets

Upcoming

Darknet (/17,/16,/16,/16,....)

Malware – not just hash

Copyright © 2010 Internet Systems Consortium

Channel example - pDNS

$ nmsgtool -l 10.0.202.255/8430 -o - -c 1

[137] [2009-12-03 12:29:57.804048000] [1:1 ISC ncap] [e46032b8] [] []

[192.55.83.30].53 [##.##.##.##].62855 udp [115]

dns QUERY,NOERROR,3324,qr

1 radio.wareznet.net,IN,A

0

2 wareznet.net,IN,NS,172800,ns1.wareznet.net

wareznet.net,IN,NS,172800,ns2.wareznet.net

3 ns1.wareznet.net,IN,A,172800,66.45.225.82

ns2.wareznet.net,IN,A,172800,66.45.225.83

.,CLASS512,TYPE41,32768,[0]

Copyright © 2010 Internet Systems Consortium

Channel example – web sinkhole

$ nmsgtool -l 10.16.80.255/8430 -o - -c 1

[330] 2009-03-02 22:12:27.558313023 [1:4 ISC http] [00000000 00000000]

type: sinkhole

srcip: YYY.YY.YYY.YY

srcport: 64707

dstip: 149.XX.XX.XX

dstport: 80

request:

GET /search?q=0 HTTP/1.0

User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1; .NET CLR 2.0.50727; .NET CLR 3.0.04506.648; .NET CLR 3.5.21022; .NET CLR 1.1.4322)

Host: 149.XX.XX.XX

Pragma: no-cache

p0f data (kudos to Chris Lee)

Copyright © 2010 Internet Systems Consortium

Channel example - spam

$ nmsgtool -l 10.16.25.255/8430 -c 1 -o -

[407] [2009-12-03 11:40:00.195077816] [1:2 ISC email] [0829f21a] [] []

type: spamtrap

srcip: 189.15.60.161

helo: bl15-60-161.dsl.asiatel.tl

from: REDACTED@spamtrapdomain.net

bodyurl: http://dc0ca4266.xivivxt.cn/

bodyurl: http://www.w3.org/1999/xhtml

bodyurl: http://94e433.xivivxt.cn/

bodyurl: http://60436719c5.xivivxt.cn/

bodyurl: http://4229da8a0.xivivxt.cn/

bodyurl: http://2d0a7d68.xivivxt.cn/ff24490.gif

bodyurl: http://08a6e3884b.xivivxt.cn/

bodyurl: http://www.w3c.org/TR/1999/REC-html401-19991224/loose.dtd



Copyright © 2010 Internet Systems Consortium

Expand with pDNS

$ whois -h 10.255.1.16 124.42.113.146,ip | awk '$3 == “A” {print $1}'

faa76.gupicfd.cn

fb6886169.yepekfy.cn

...etc...

ns3.j8w.ru

www.vapagnj.cn

xidisqs.cn

xivivxt.cn

yepekfy.cn

zuvidtn.cn

zuwohxc.cn

zuyimqg.cn

zuzewnp.cn

zuzovgw.cn

Copyright © 2010 Internet Systems Consortium

Combining data

Jose Nazario / Thorsten Holz - Malware08

- Heuristics, Point system, Fast-flux

David Dagon / Wenke Lee

- pDNS + string matching -> FakeAV

Richard Clayton – UKNOF13

- pDNS hosts + active scans => blocking policies

Andrew Fried – Blackhat DC 2010

- Spam, BGP AS, pDNS, TLD zone data
- Zeus/Avalanche (www.irs.gov.dhdkdzg.eu)

Ed Stoner – Flocon 2010

- pDNS/ncap + netflow/silk

Copyright © 2010 Internet Systems Consortium

How organizations can help

Take bi-lateral sharing methods and enable real-time multi-lateral sharing via SIE

Bring servers to SIE and create value-added services

lots of data yet to be analyzed

get familiar with tools

Install sensors – enable researchers connected to SIE to analyze data that would otherwise be lost – your junk is another's treasure

pDNS, spam, netflow, darknet blocks, etc.

Copyright © 2010 Internet Systems Consortium

Questions?

Email: info@sie.isc.org

Web: https://sie.isc.org/

Eric Ziegast – SIE Programme Manager

+1.650.423.1363 (PST8PDT)

Nmsg:

ftp://ftp.isc.org/isc/nmsg

https://lists.isc.org (nmsg-dev)

Ncap:

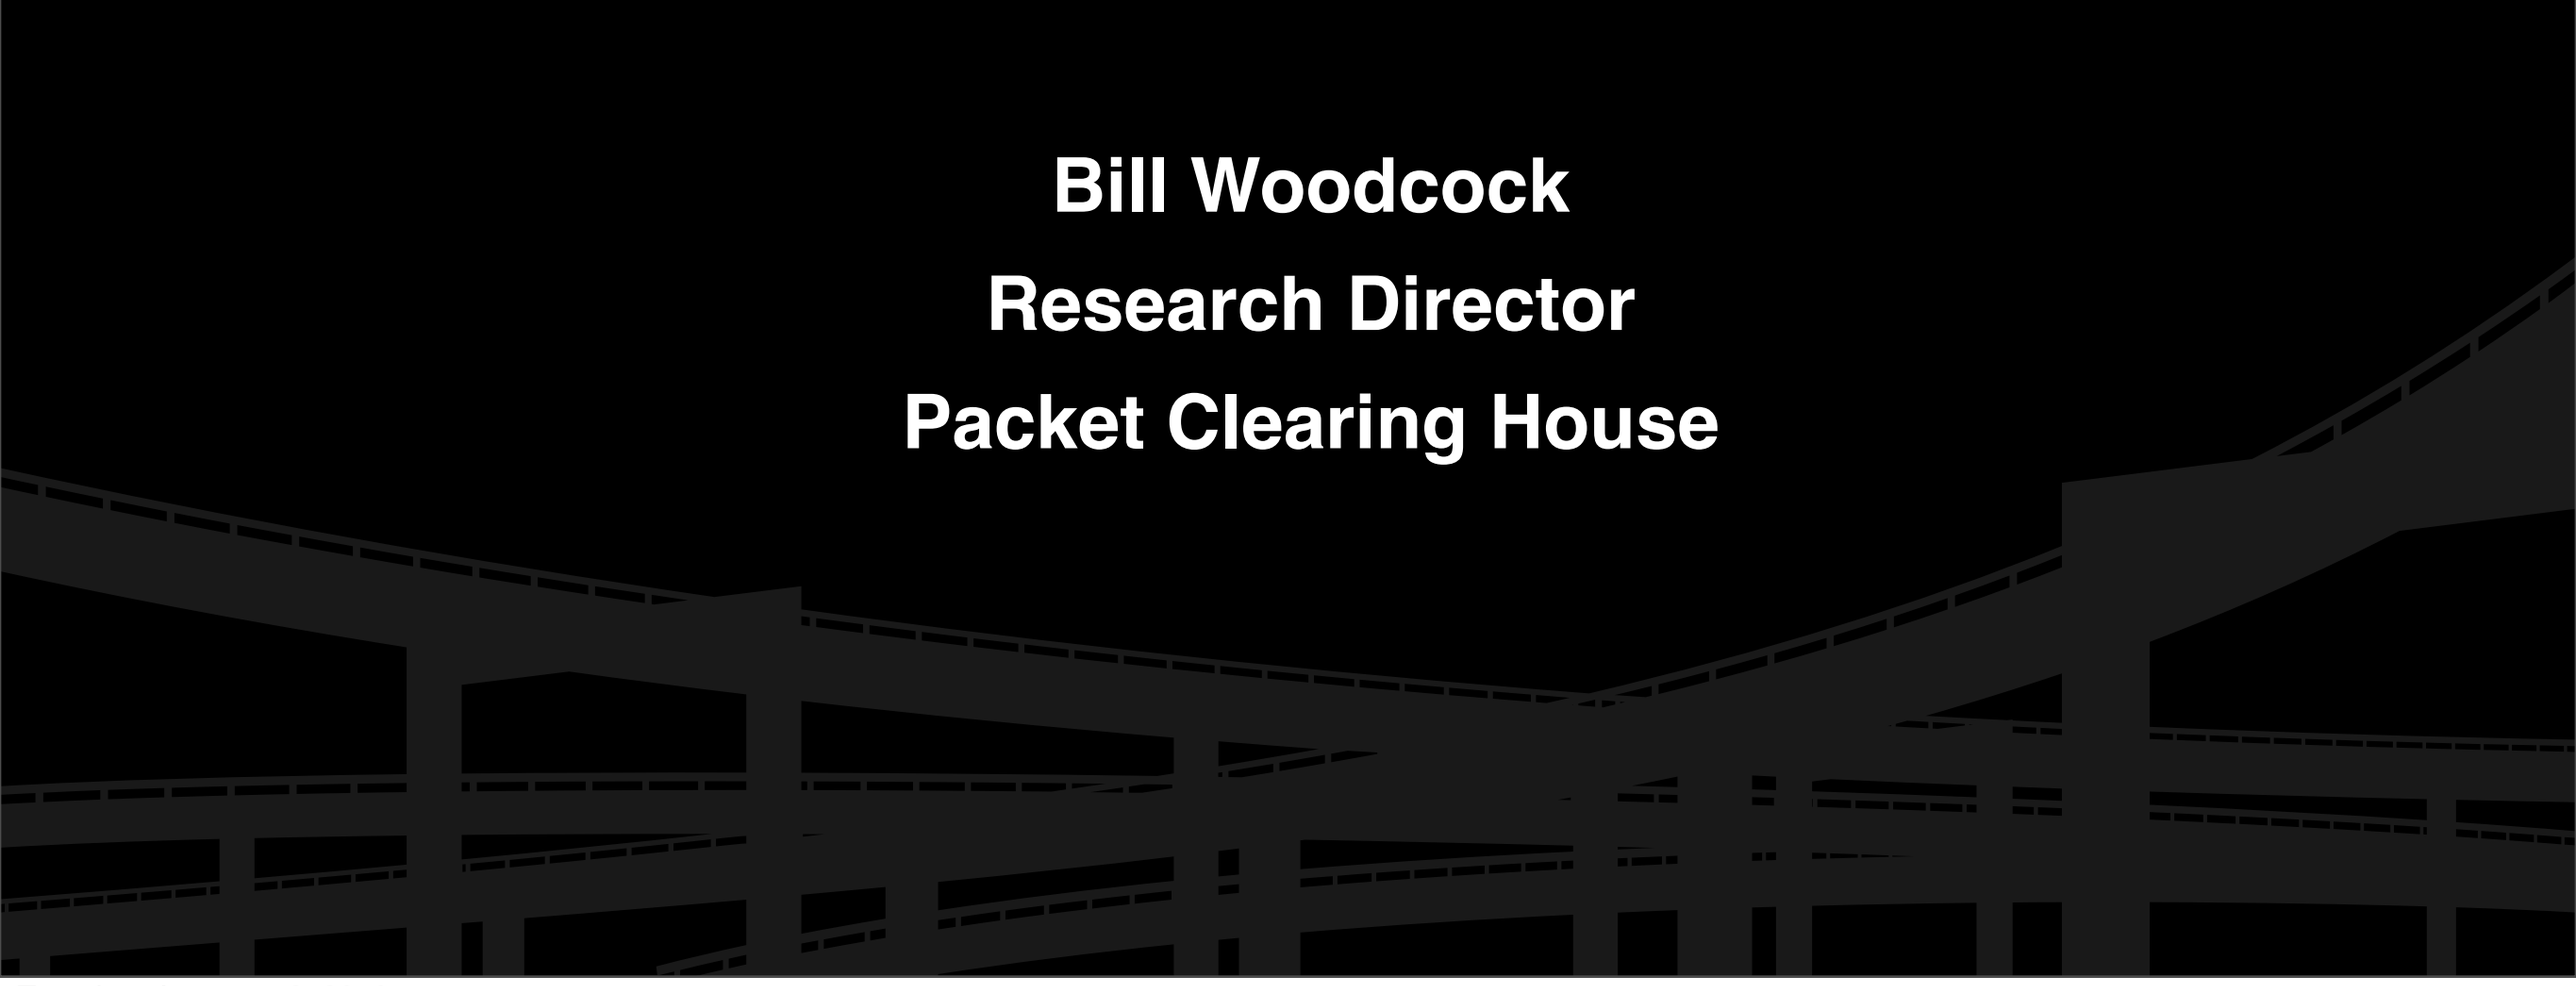
ftp://ftp.isc.org/isc/ncap

Copyright © 2010 Internet Systems Consortium

PCH
Packet Clearing House

FloCon2010
Keynote

Bill Woodcock
Research Director
Packet Clearing House



Tuesday, January 12, 2010

PCH
Packet Clearing House

Most of you are working with flow data to improve security.

PCH

Packet Clearing House



Tuesday, January 12, 2010

Packet Clearing House

But...

Flow data can improve Internet service billing.

Flow data can optimize routing.

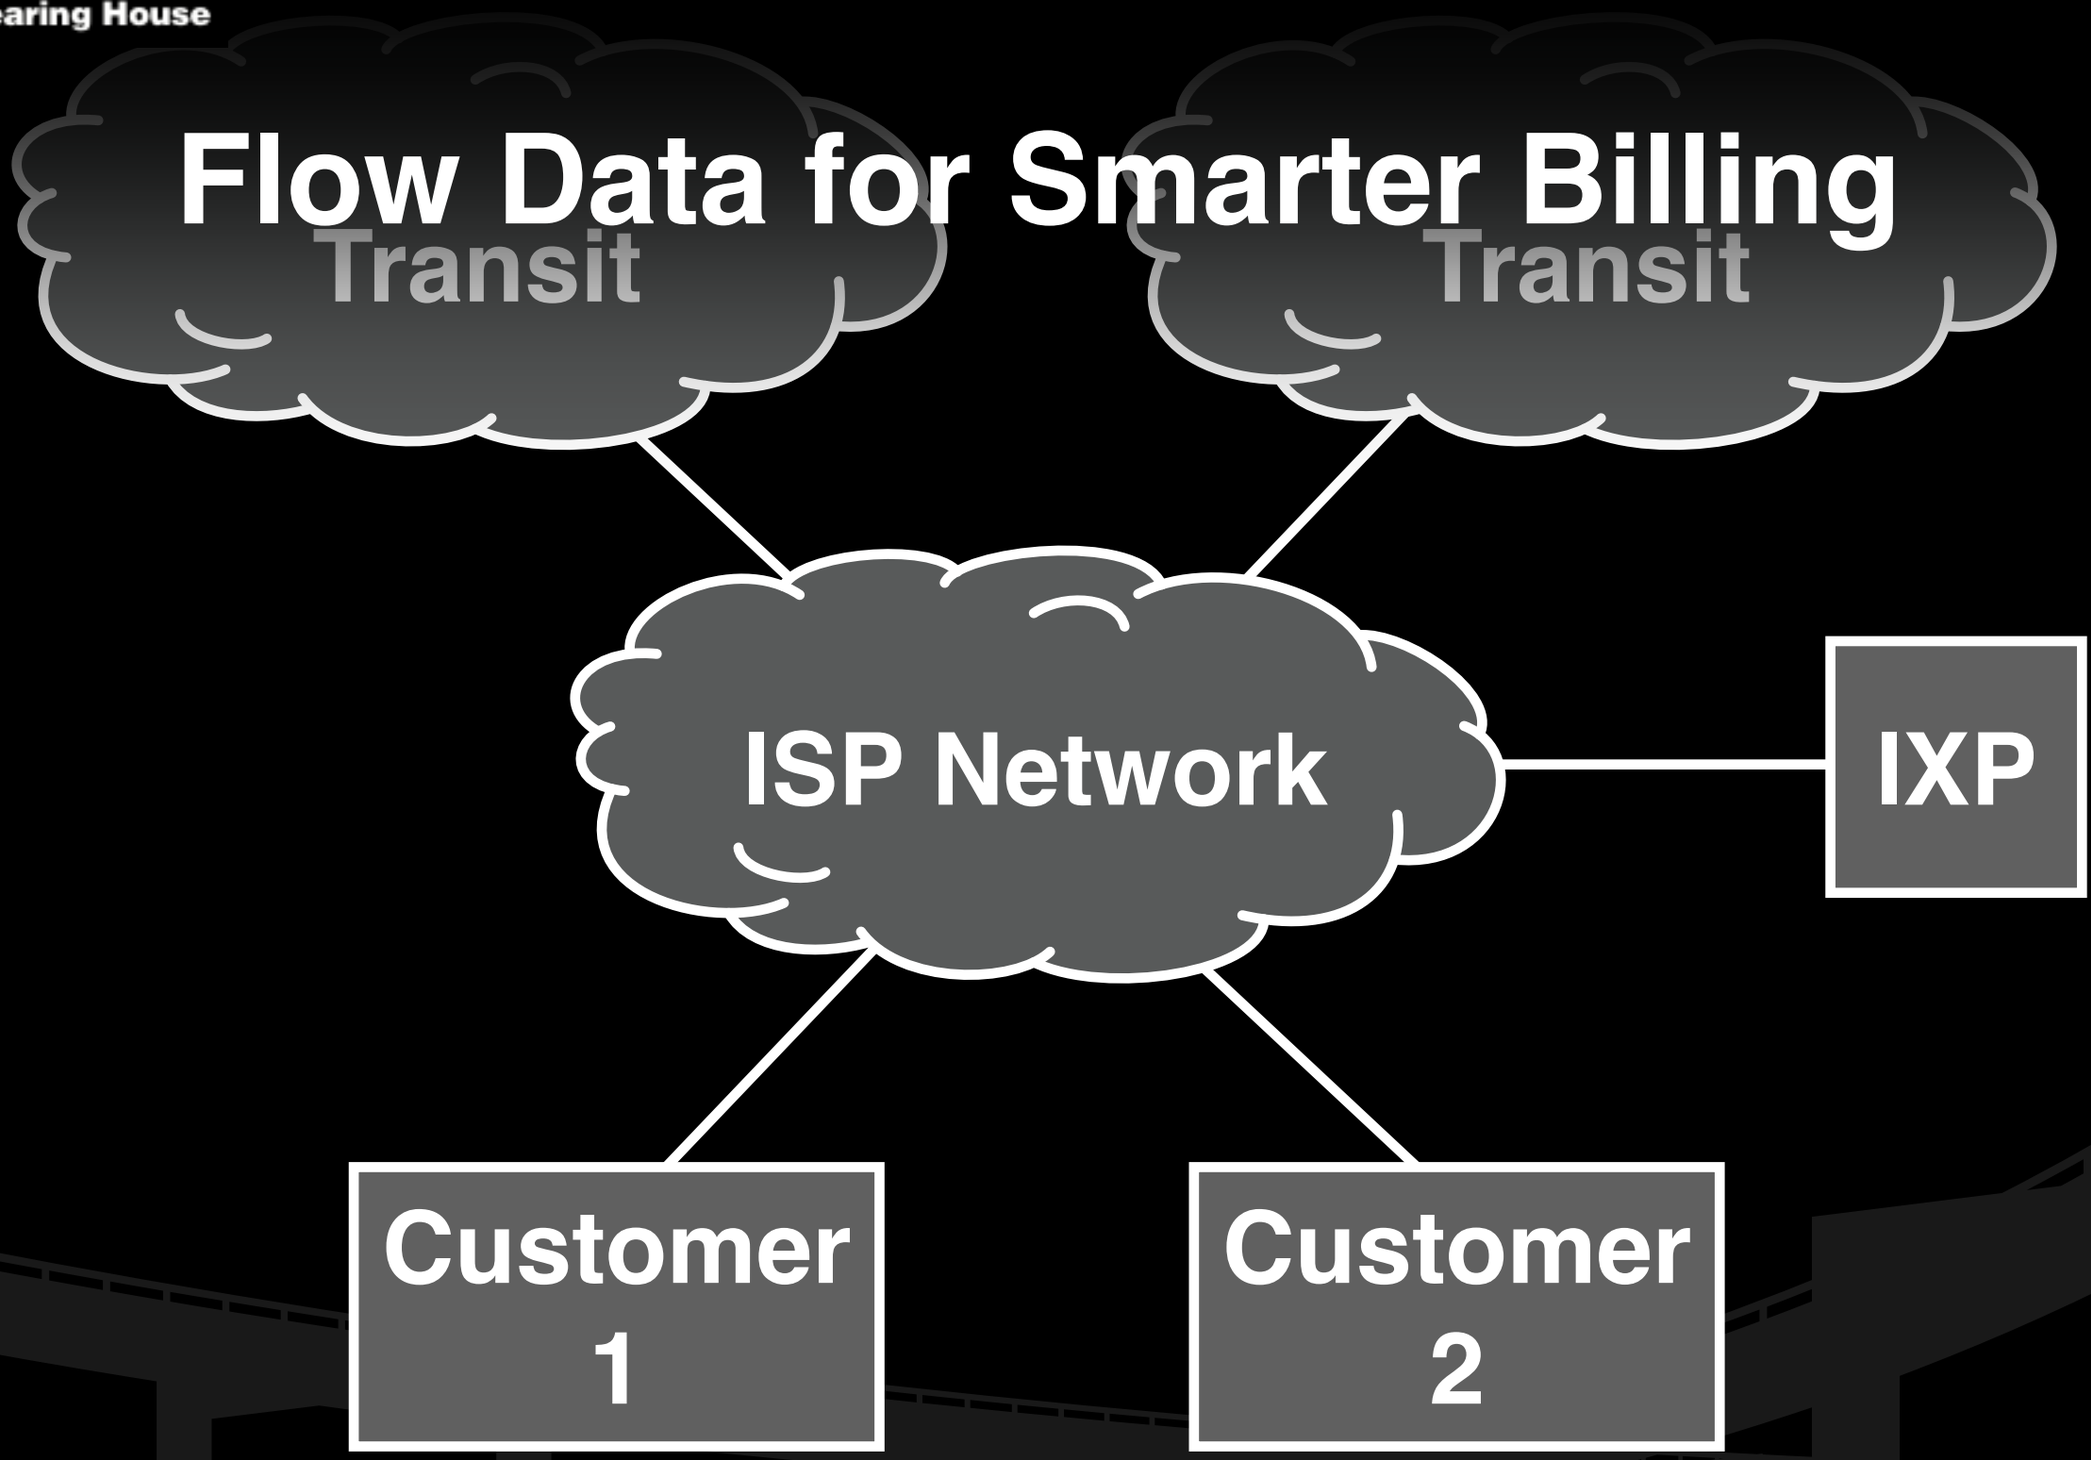
Flow data itself is becoming a tempting target.

Flow Data for Smarter Billing

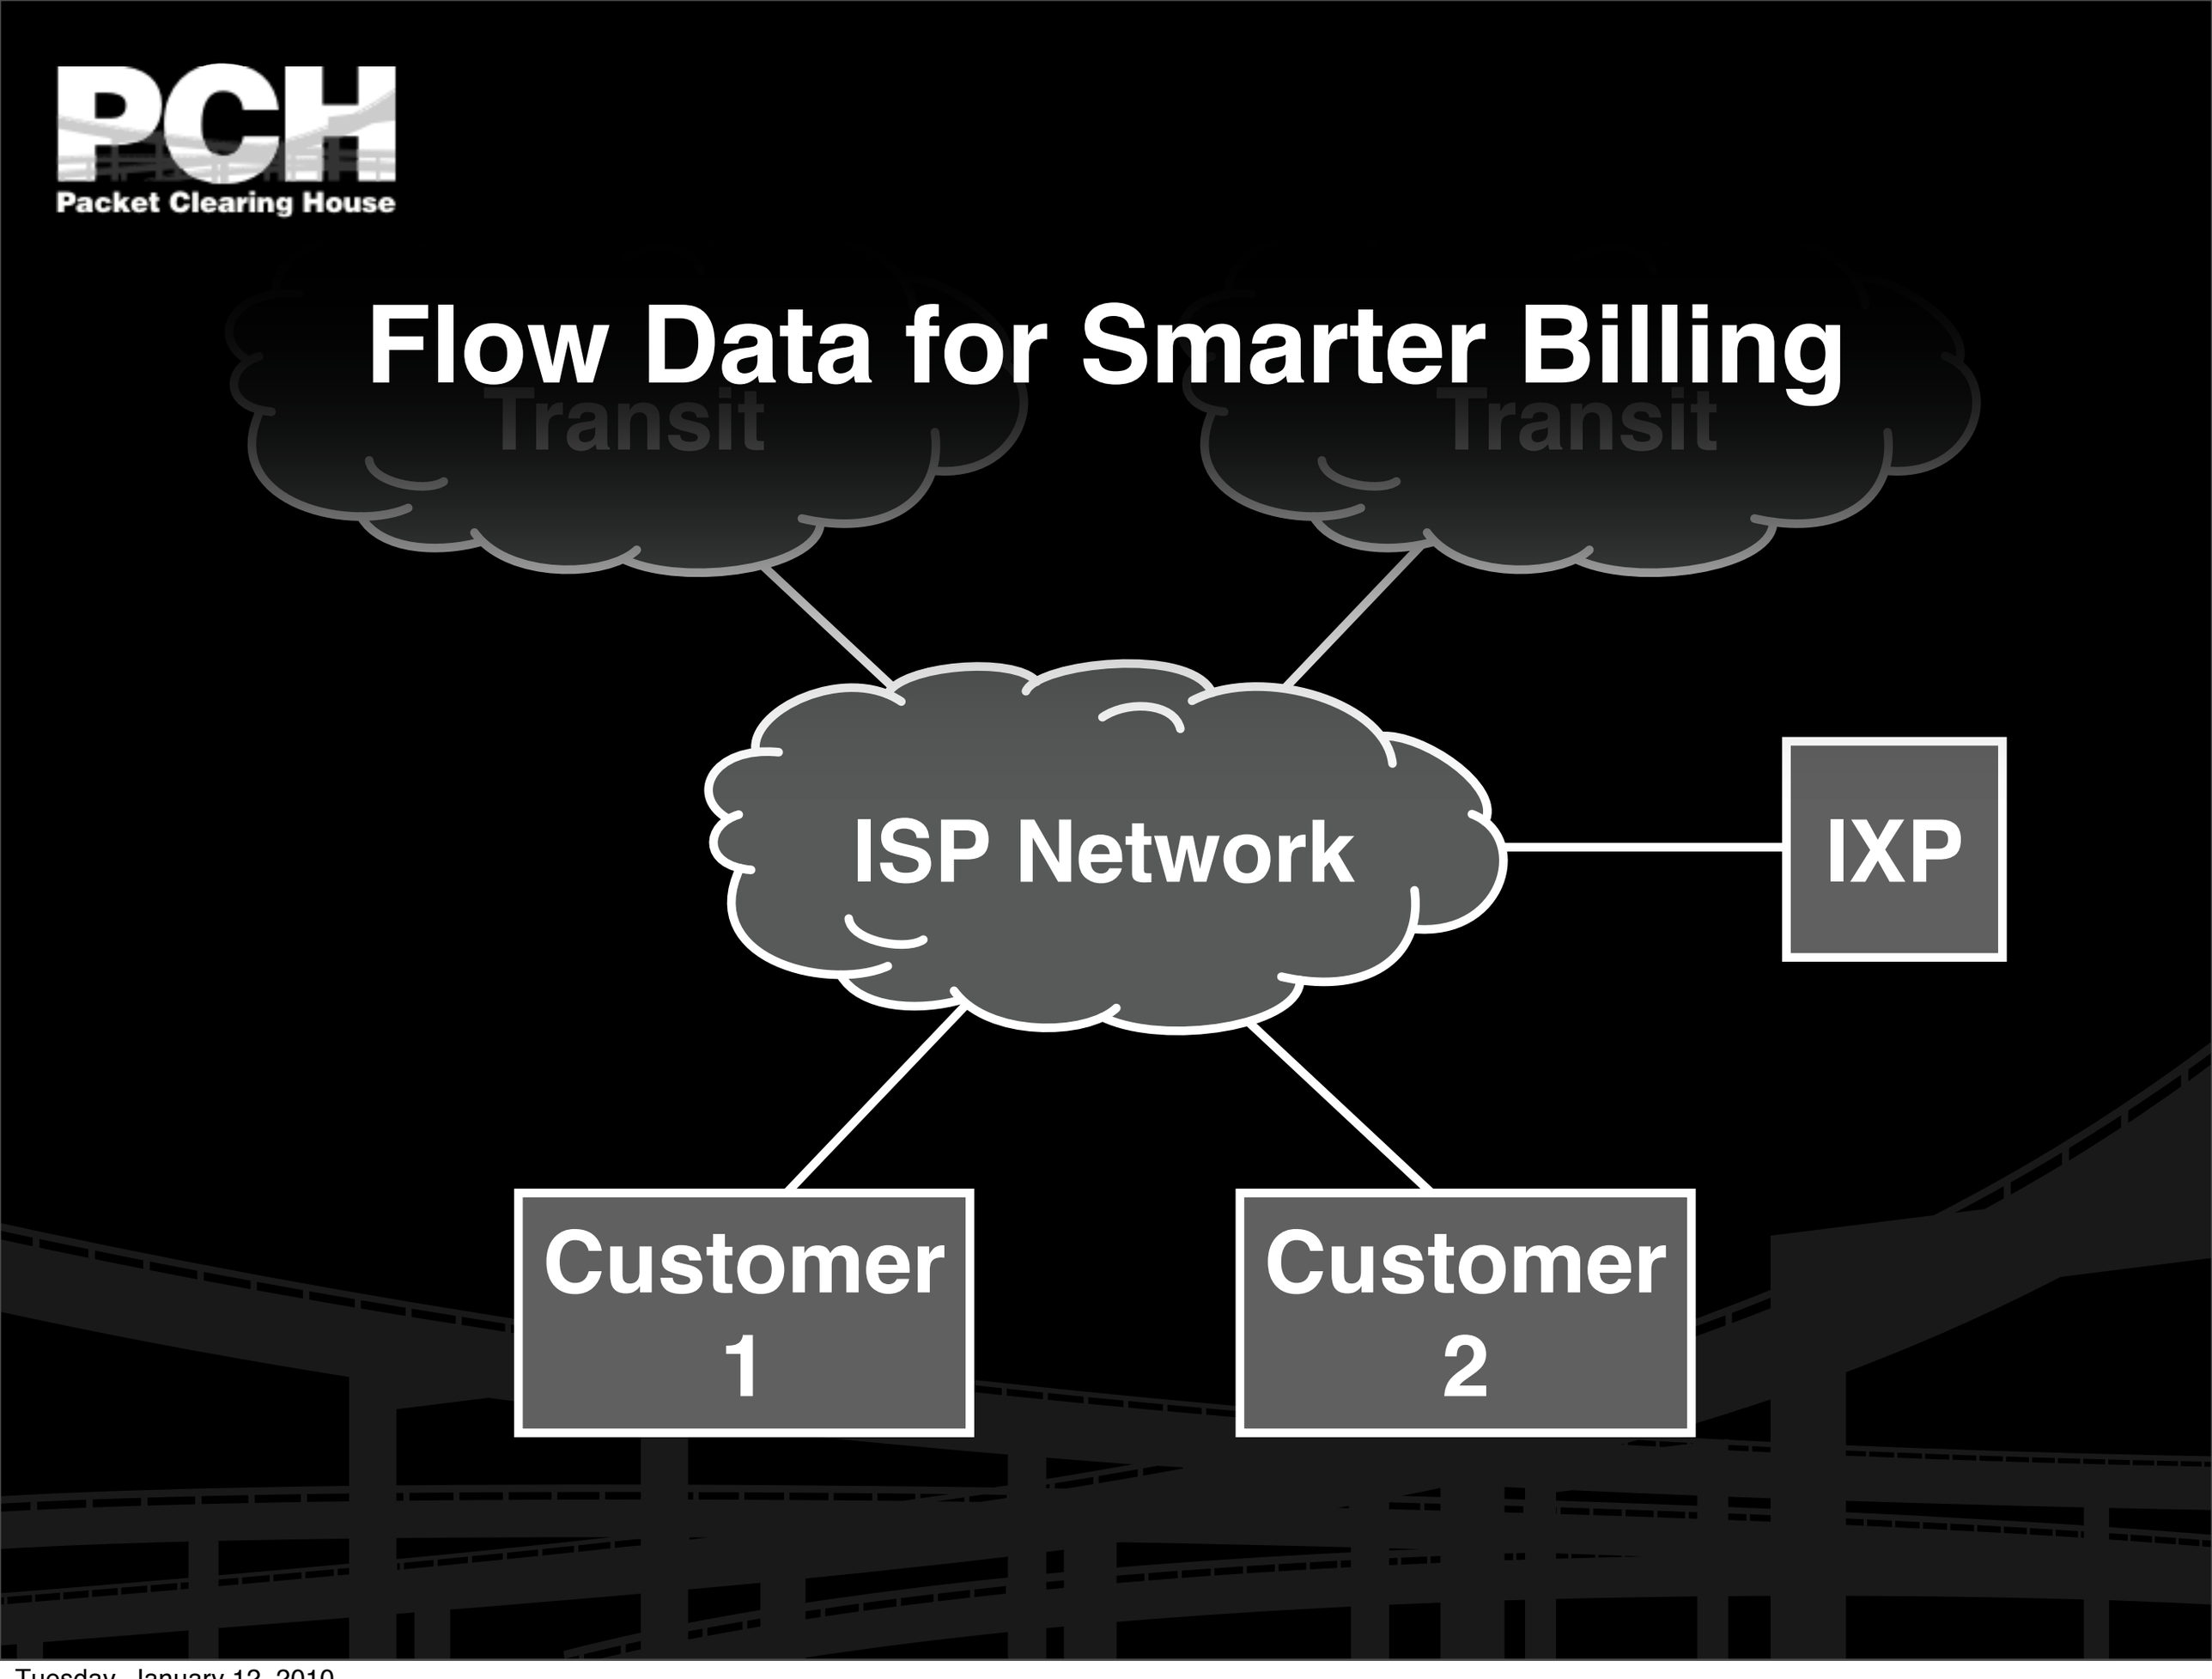
The problem:

Business models for selling Internet access are constrained and crude when ISPs use simple interface bit-counters.

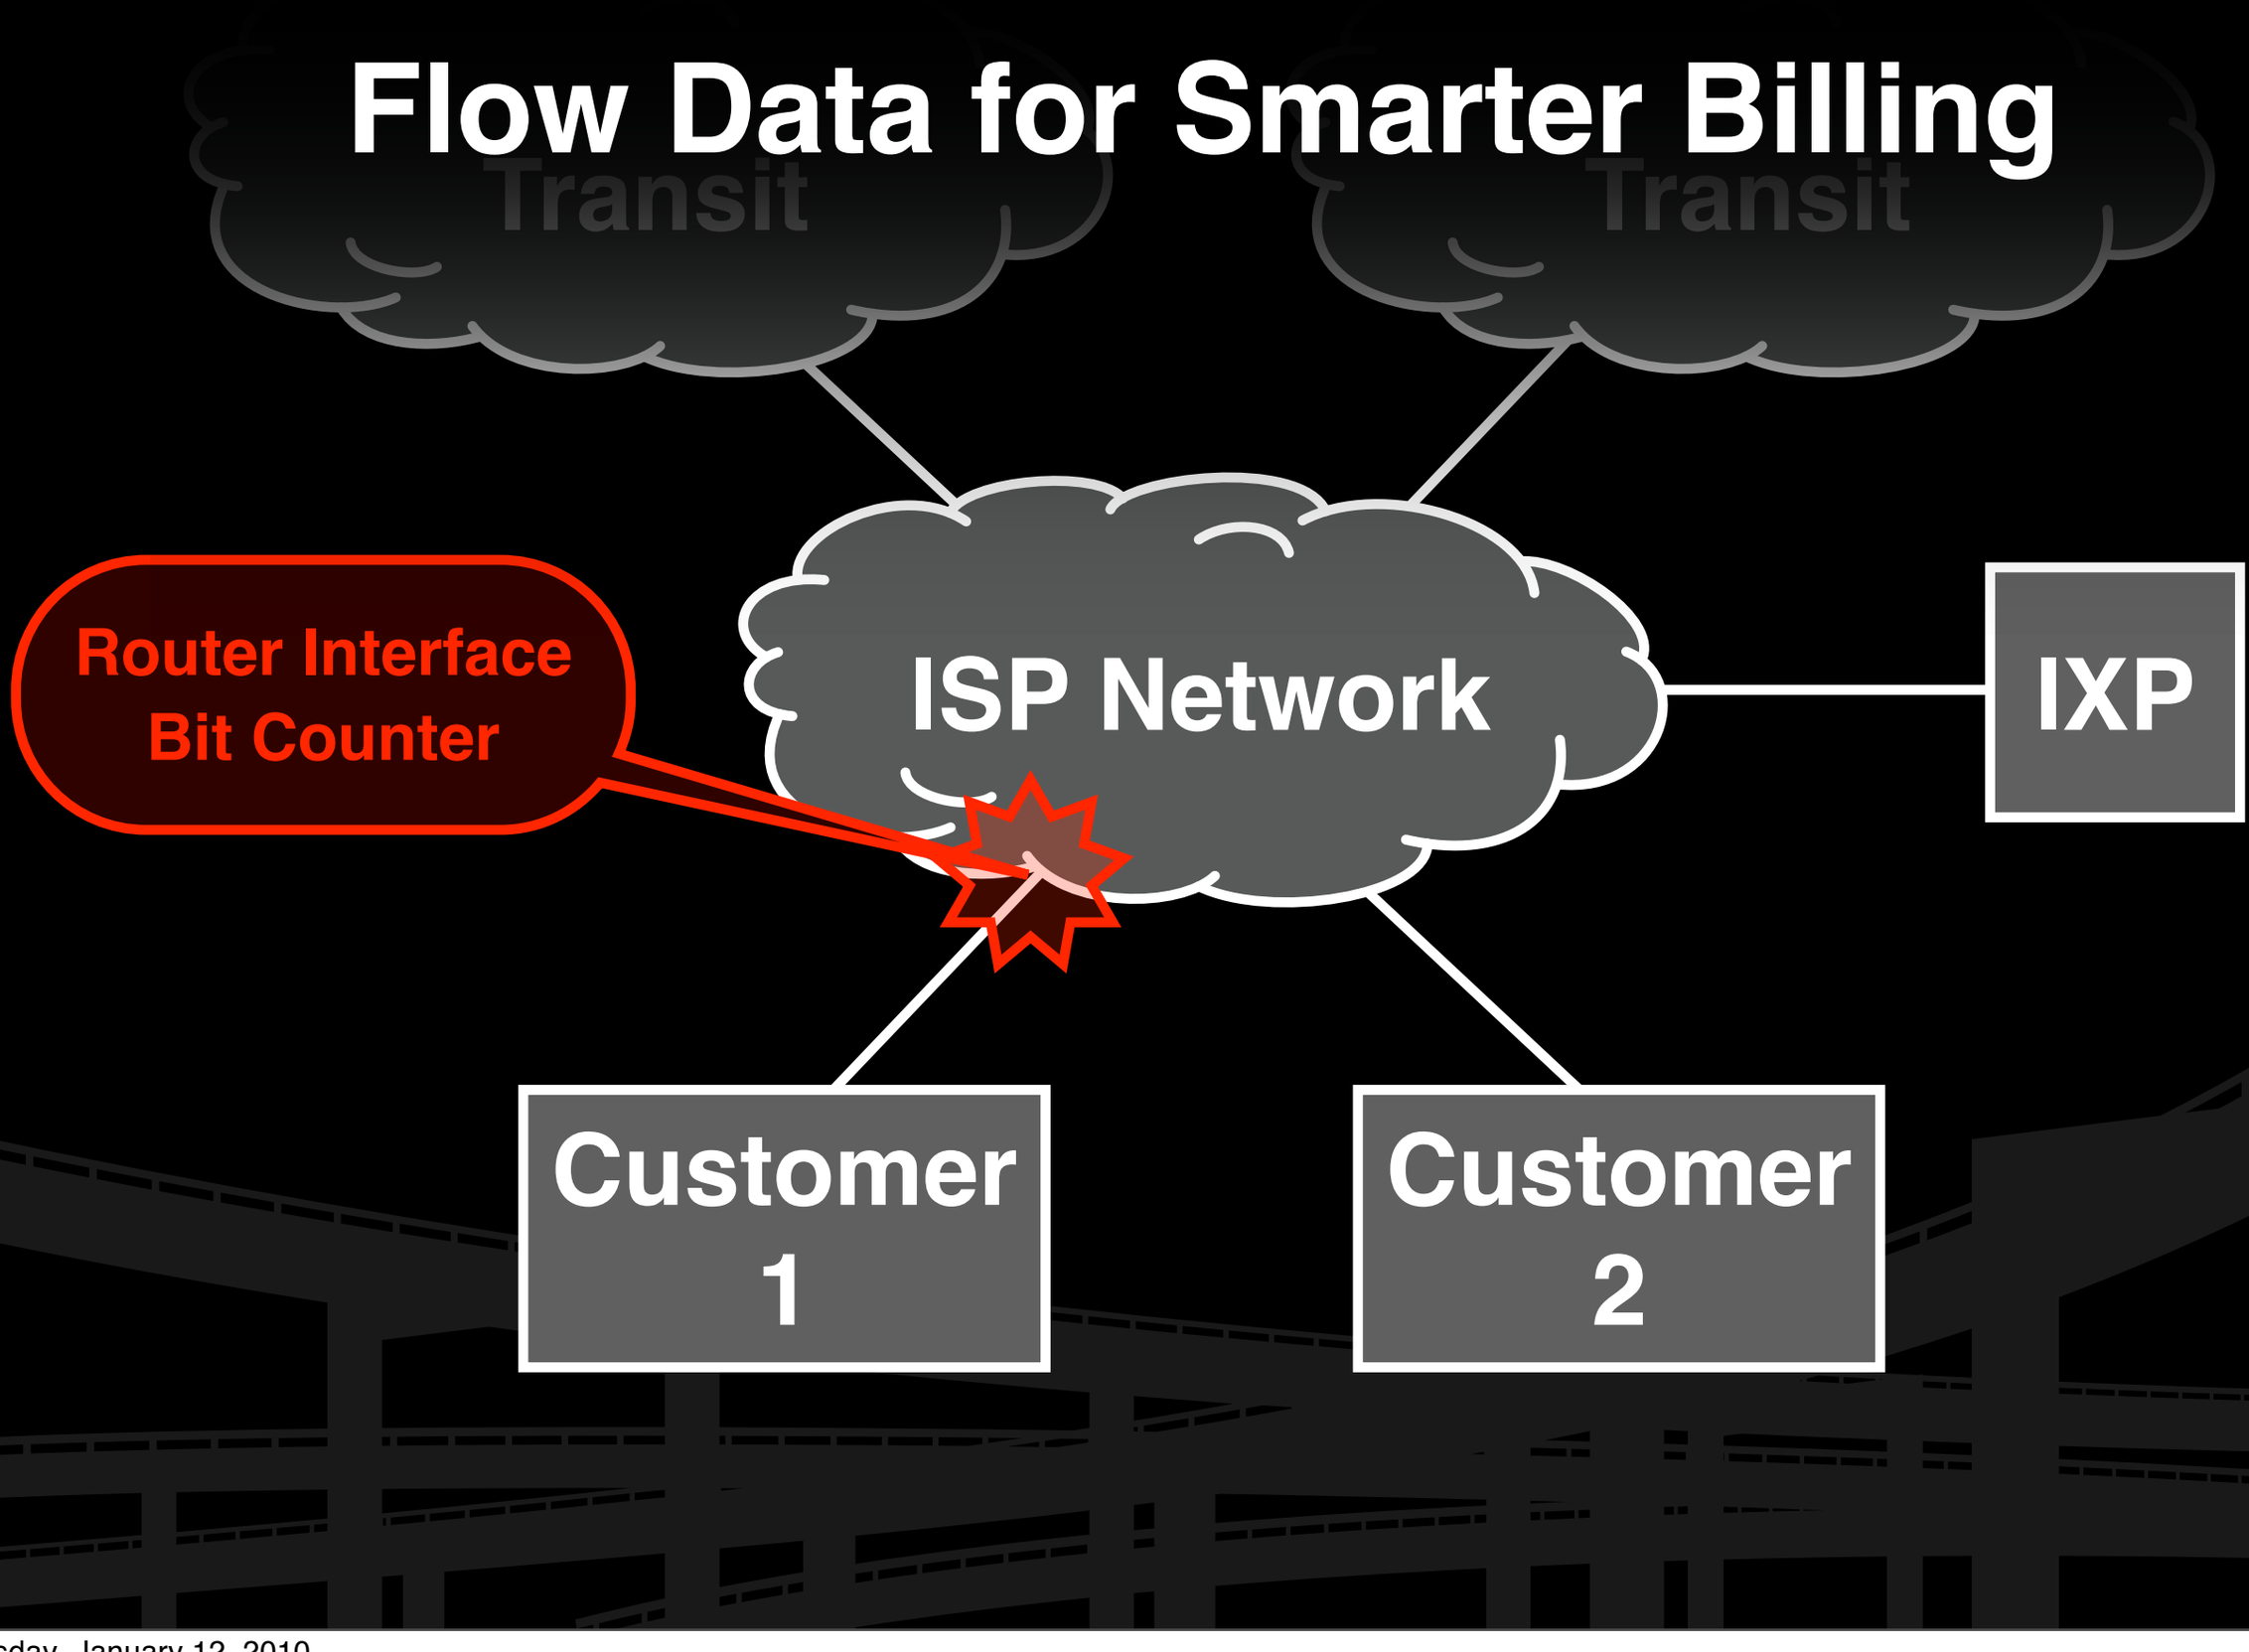
Flow Data for Smarter Billing



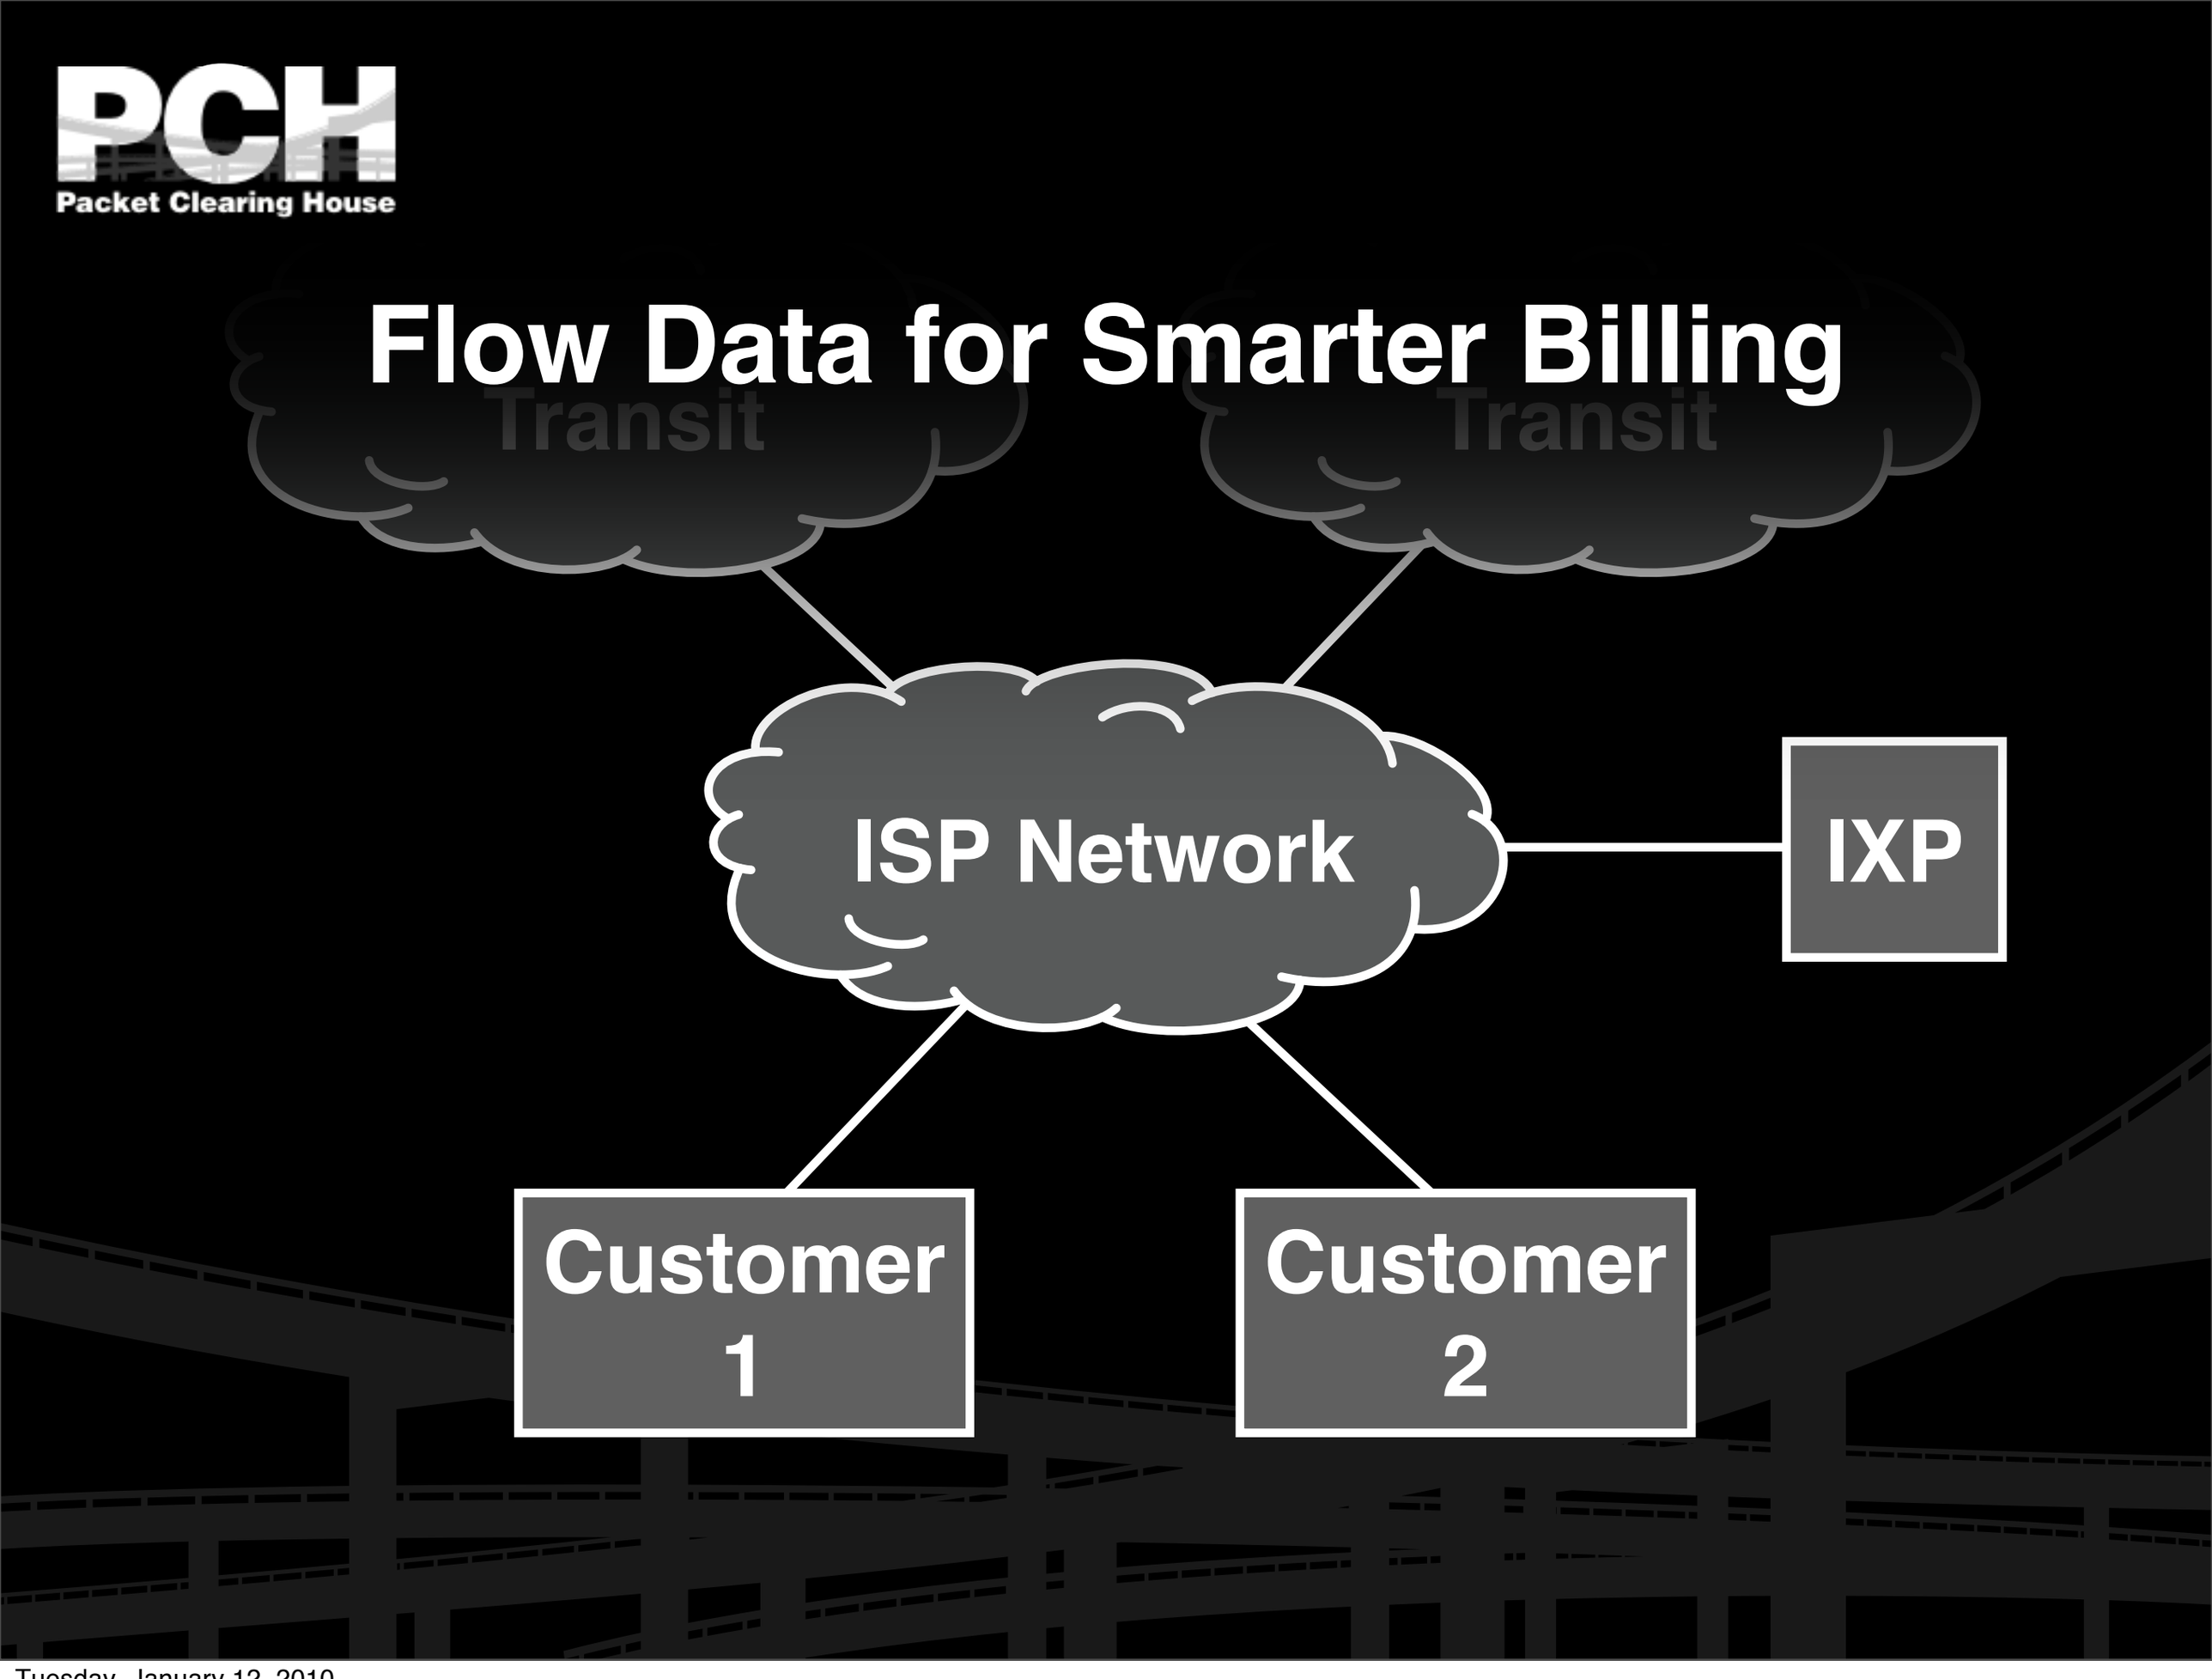
Flow Data for Smarter Billing



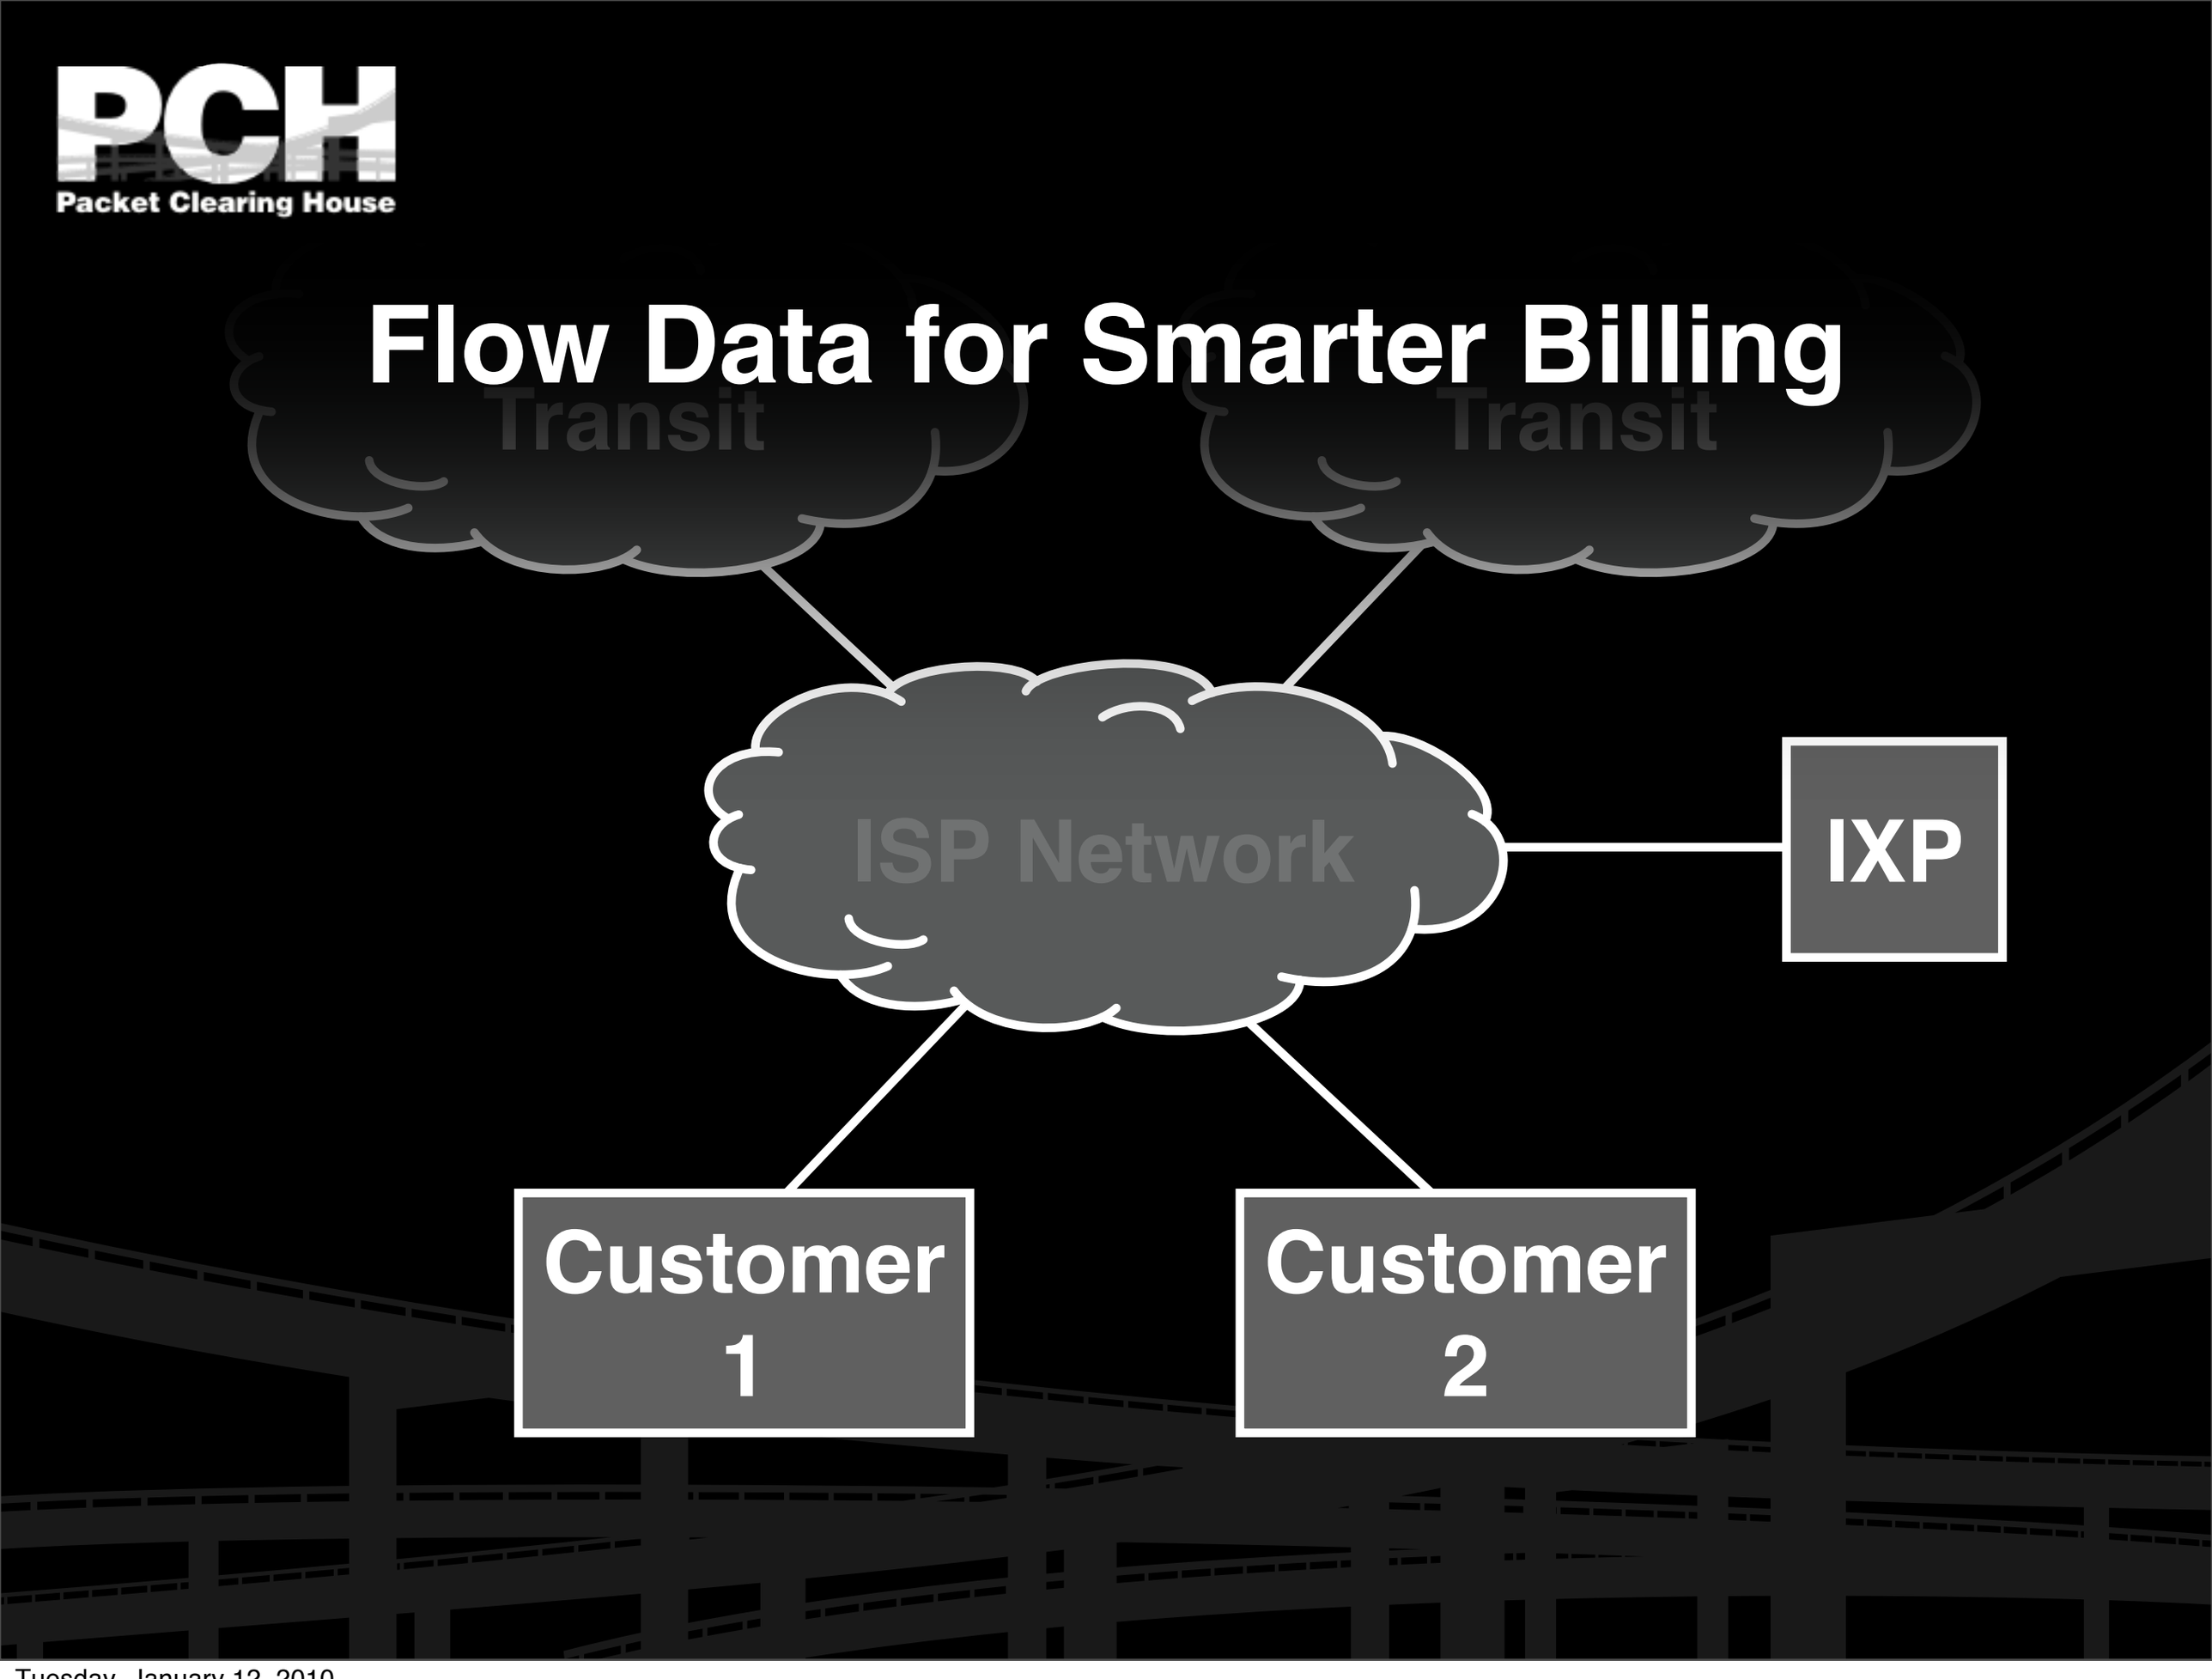
Flow Data for Smarter Billing



Flow Data for Smarter Billing

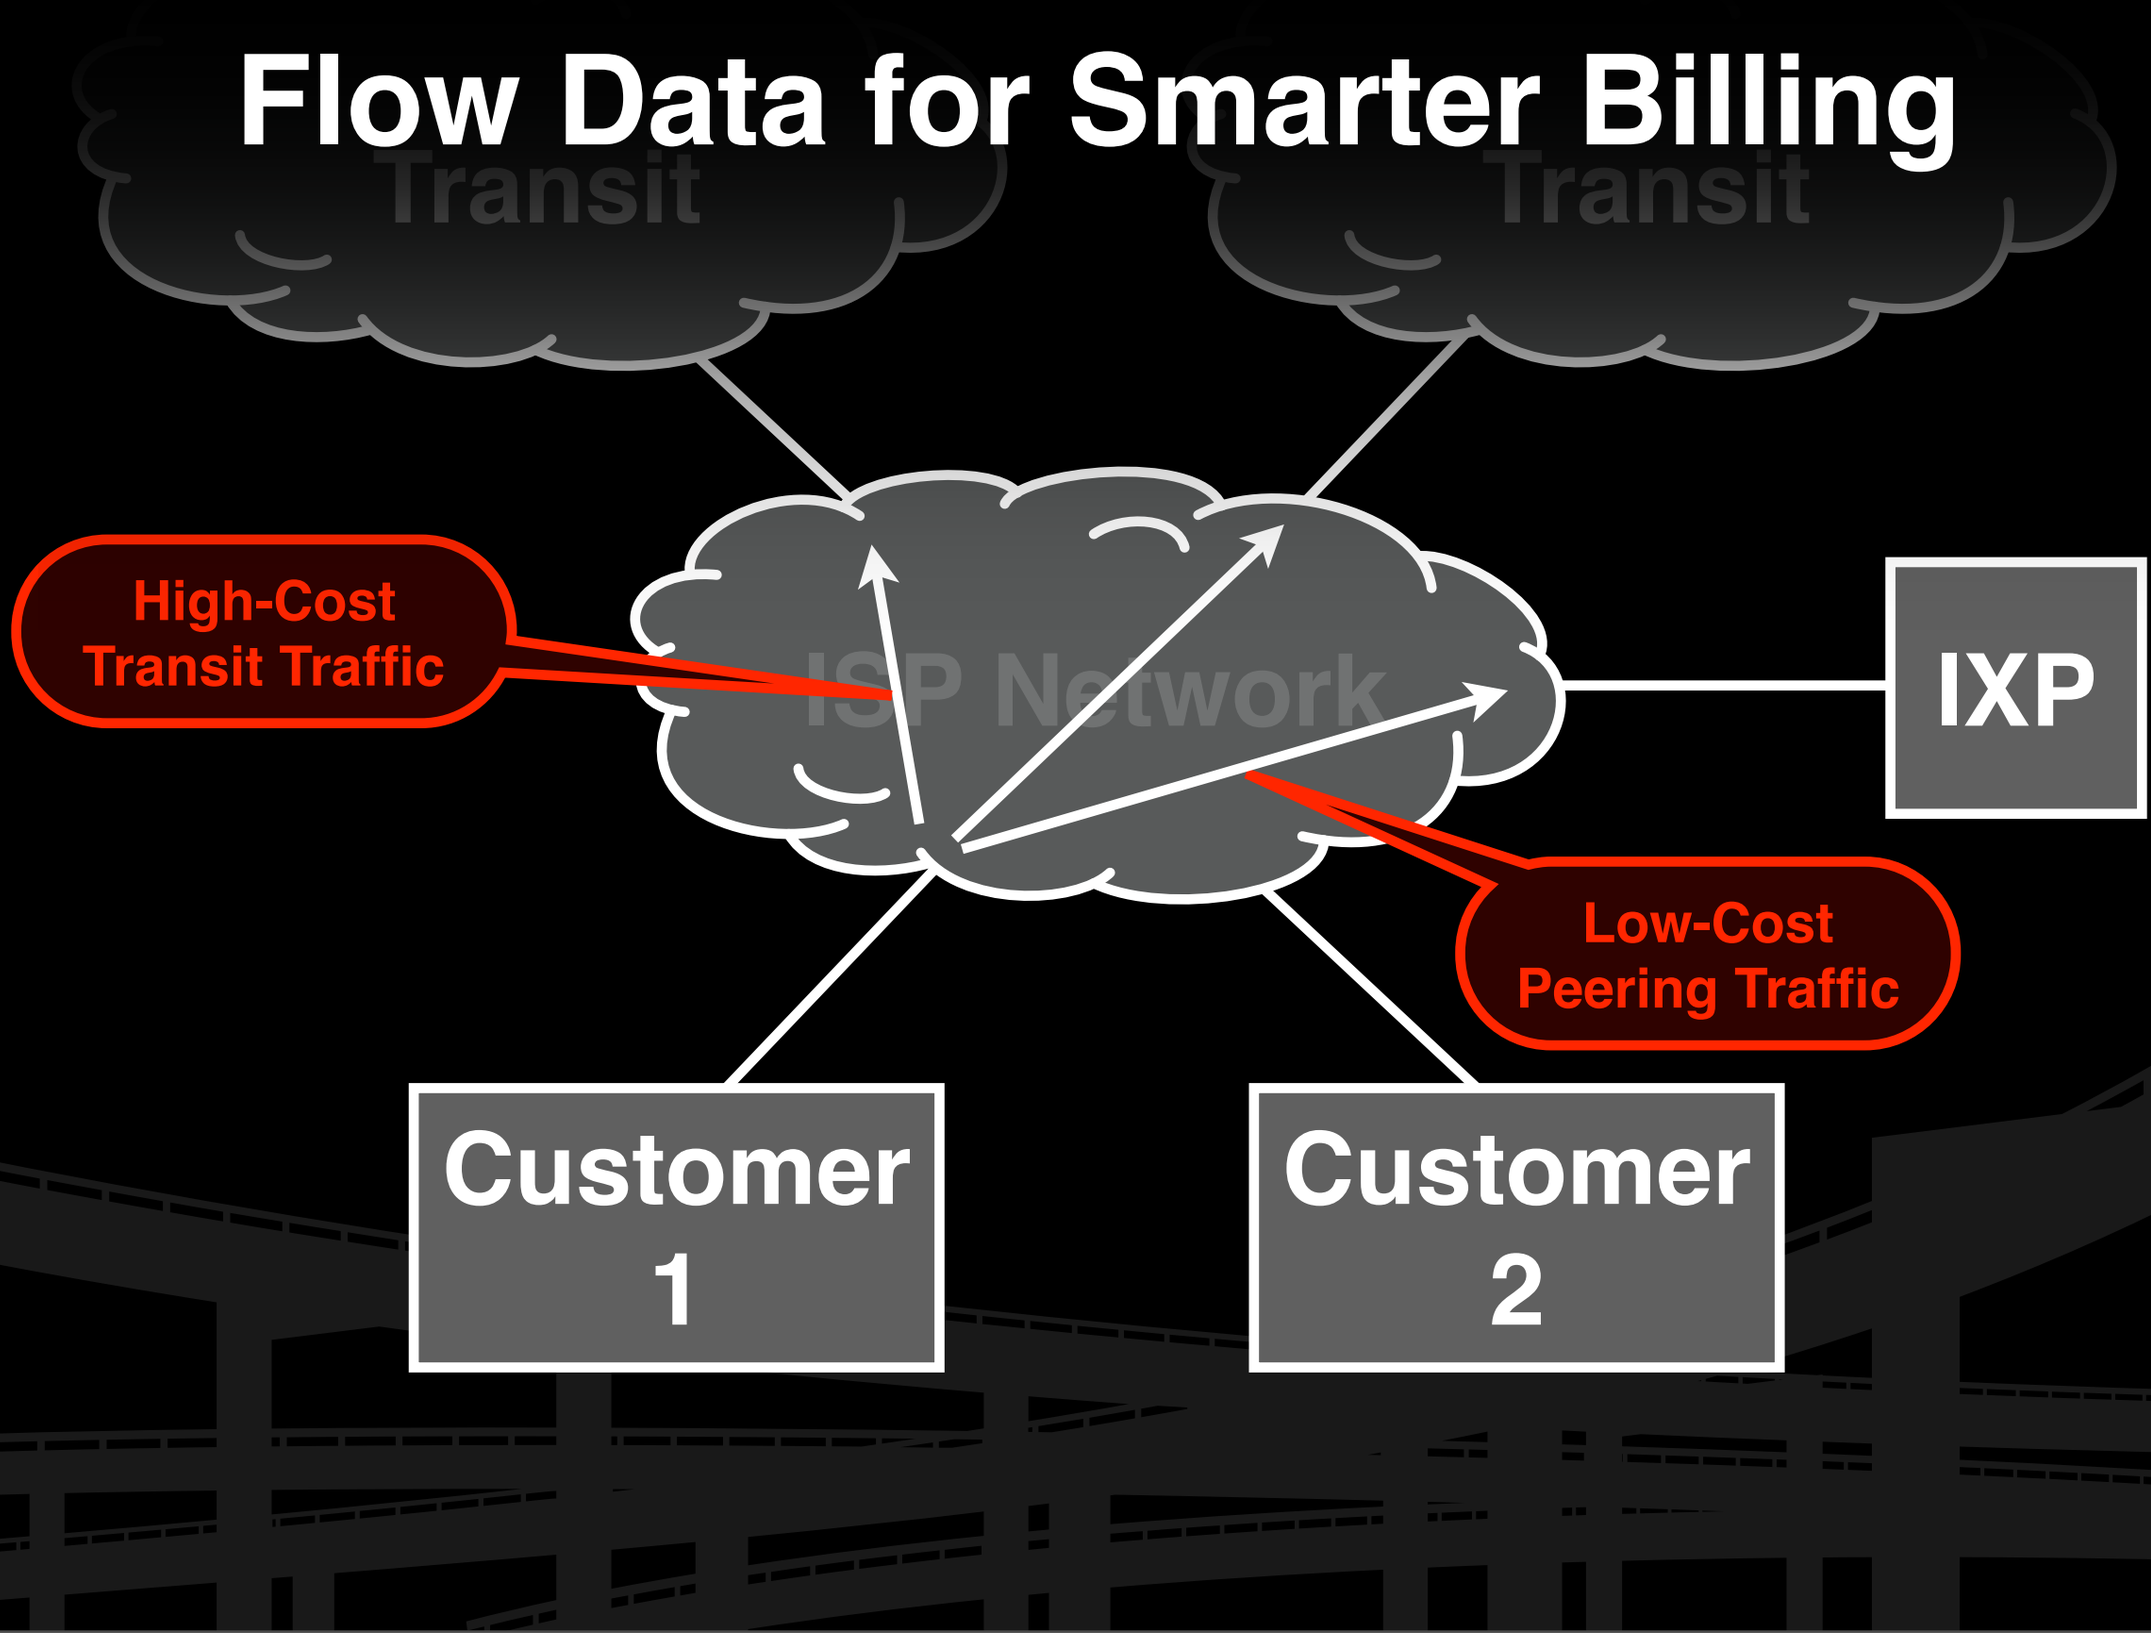


Flow Data for Smarter Billing

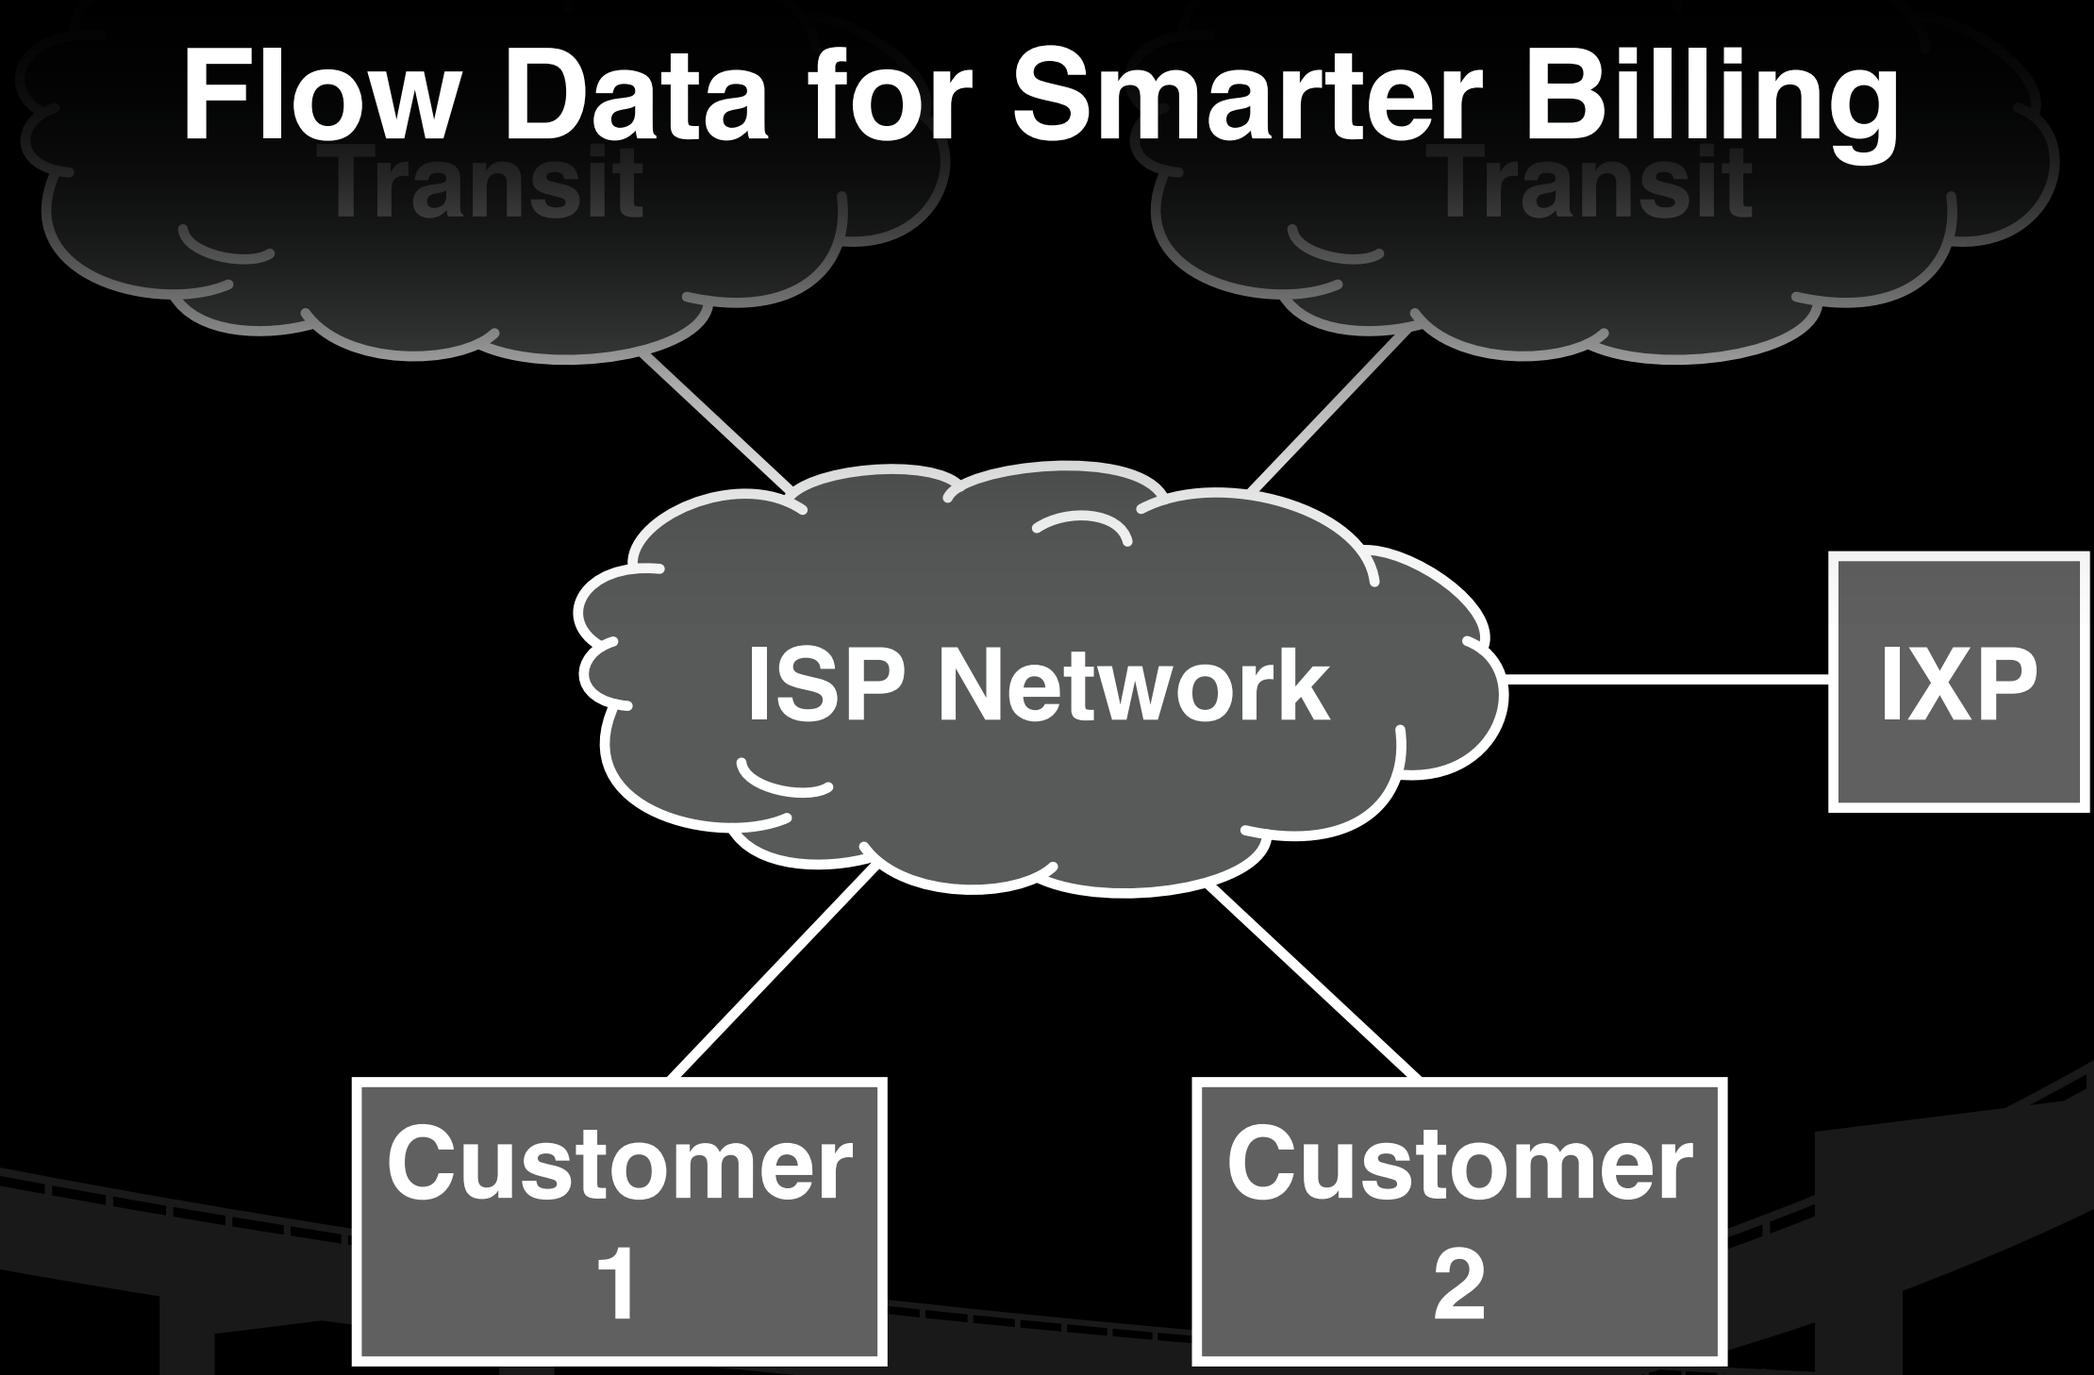


Packet Clearing House

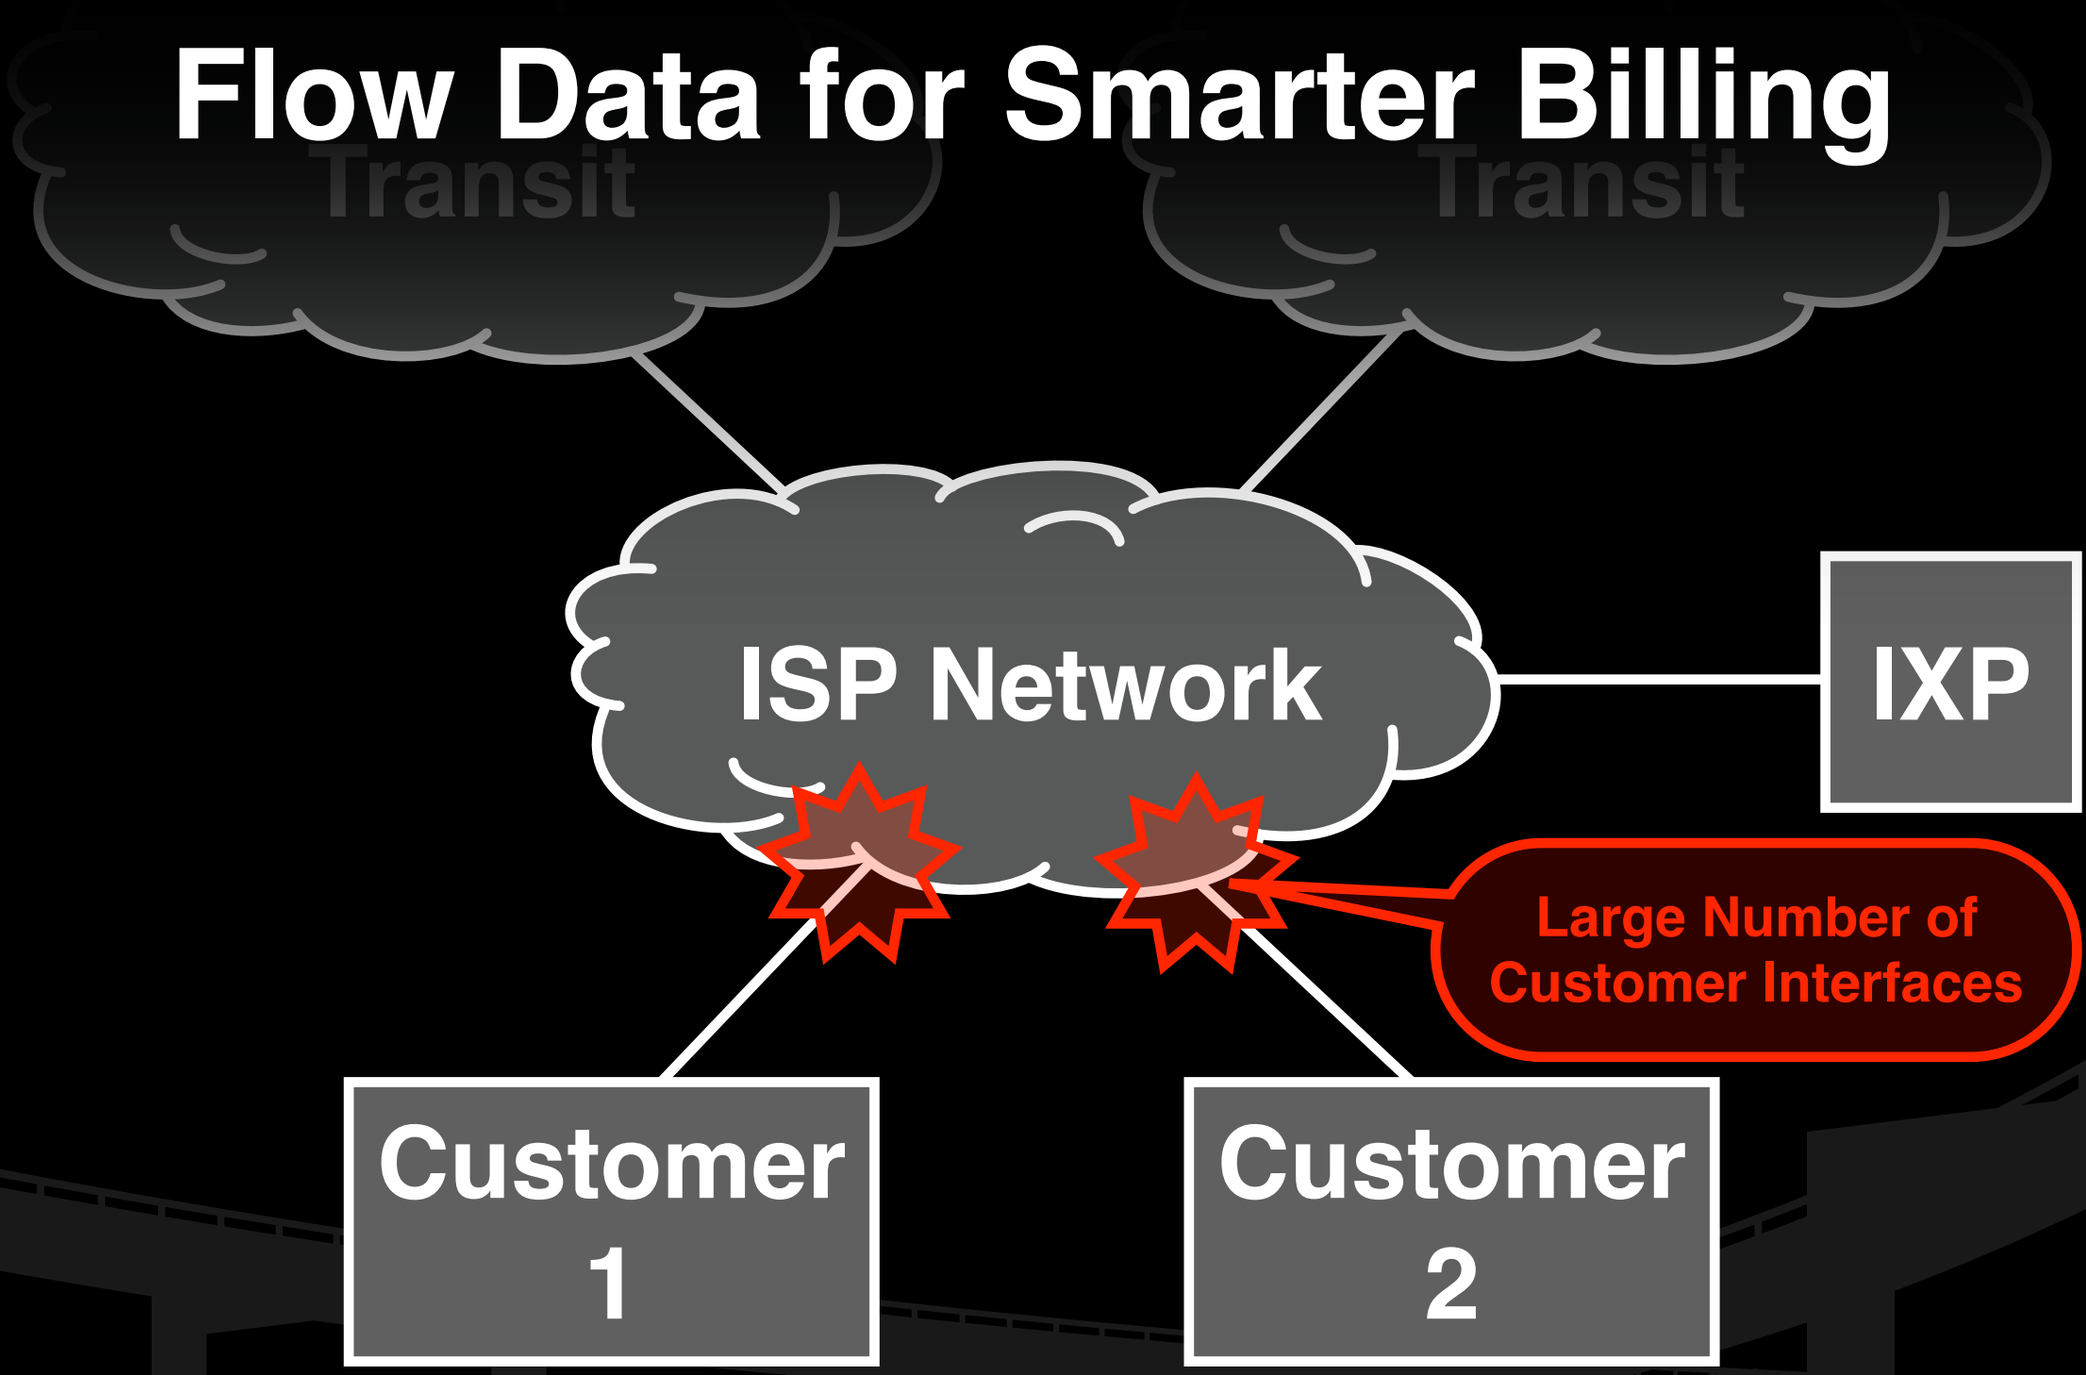
Flow Data for Smarter Billing



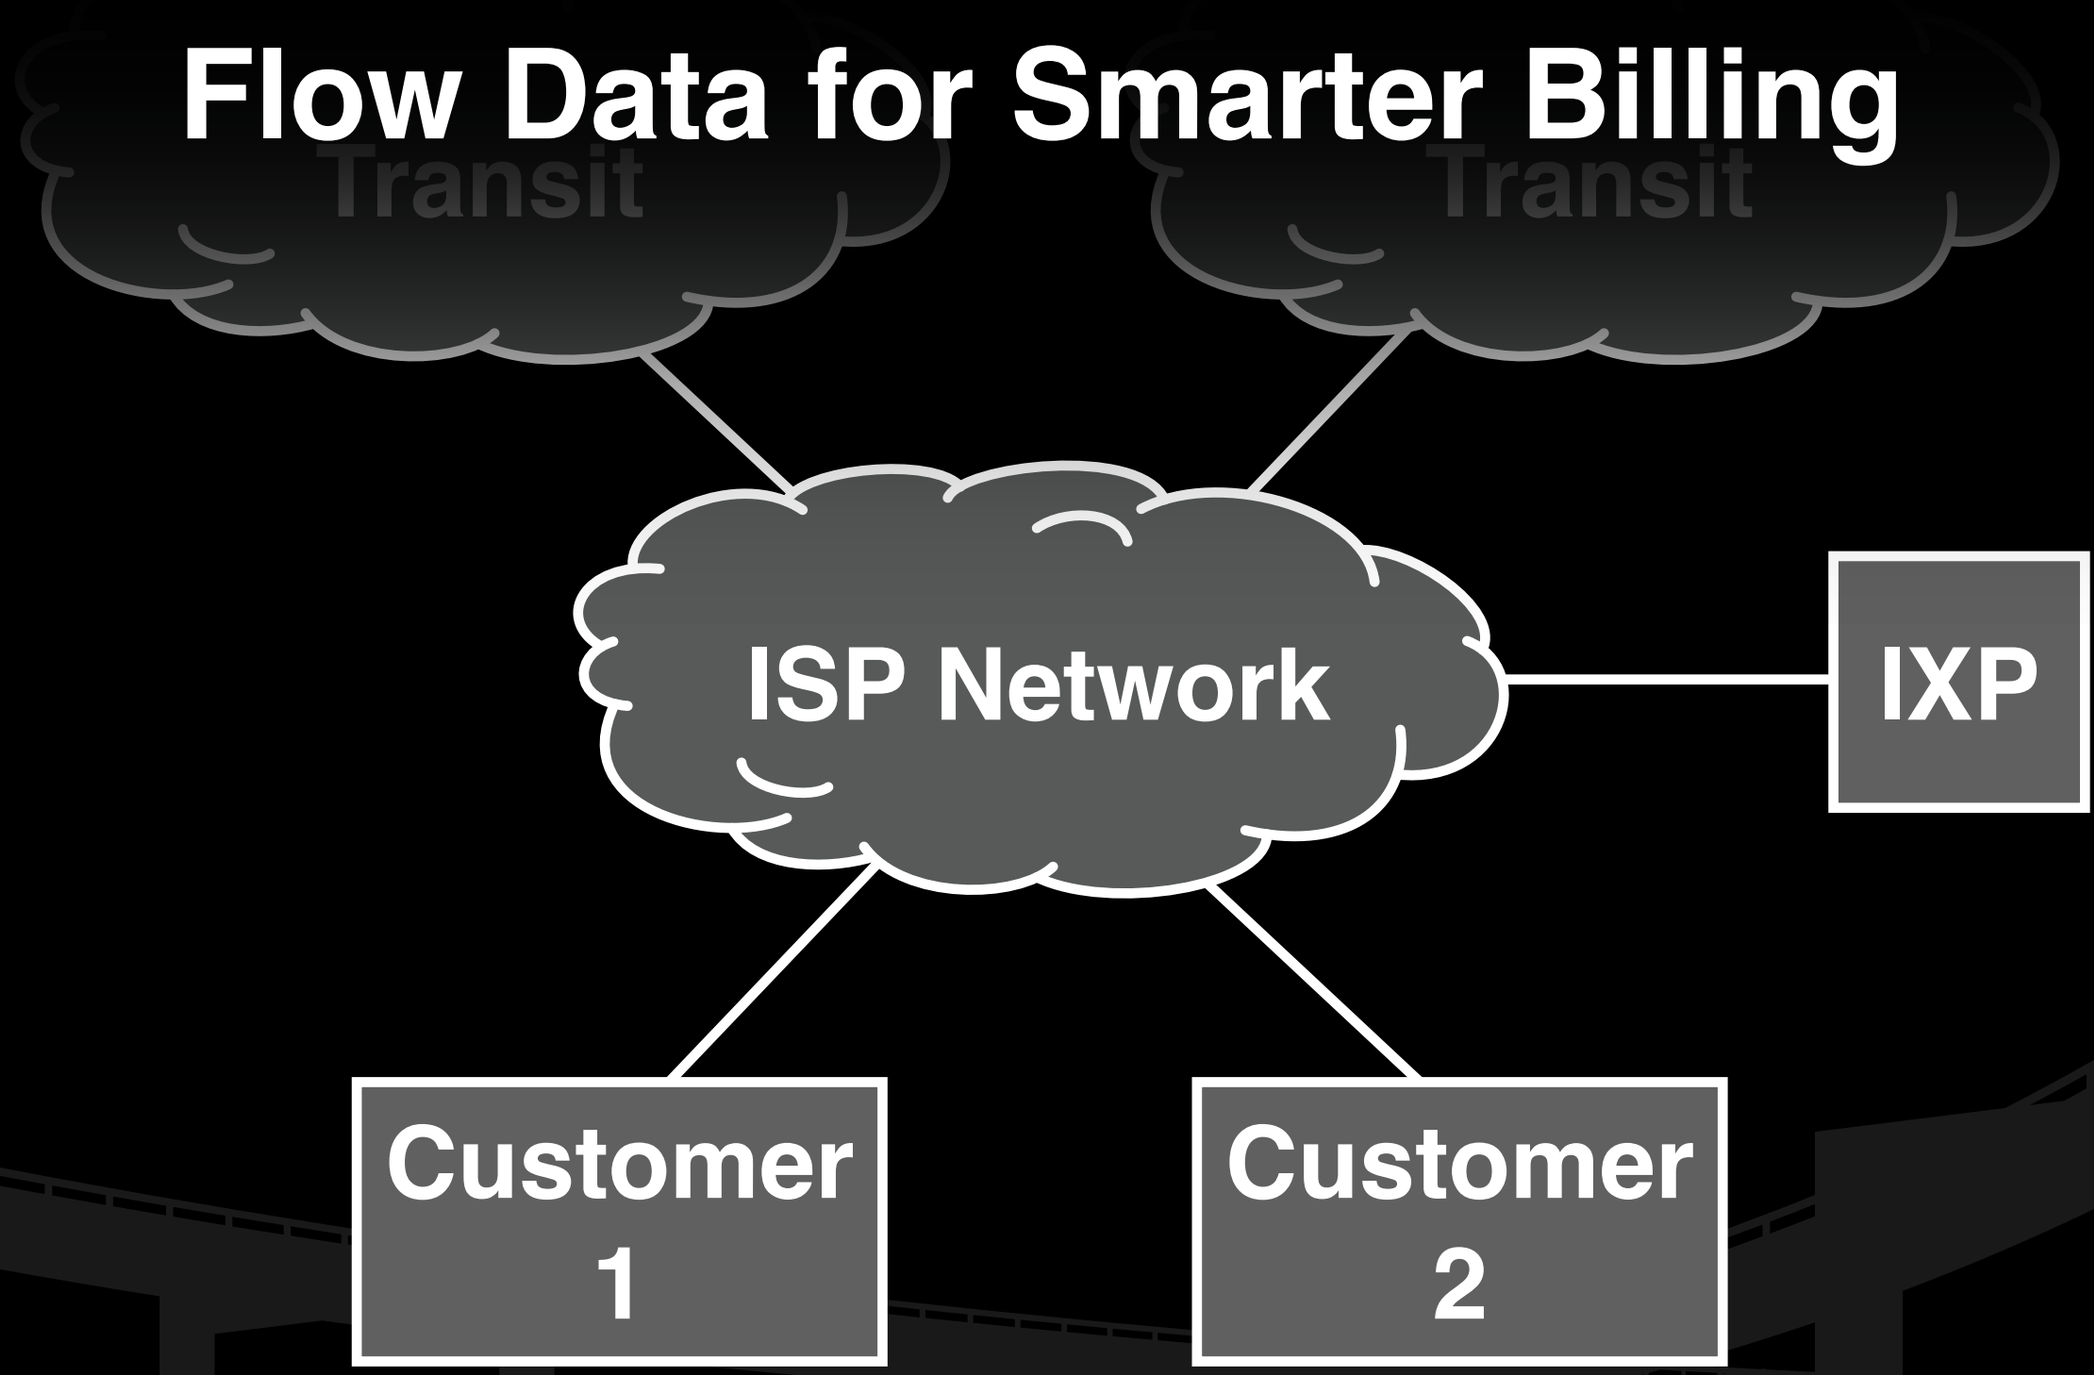
Flow Data for Smarter Billing



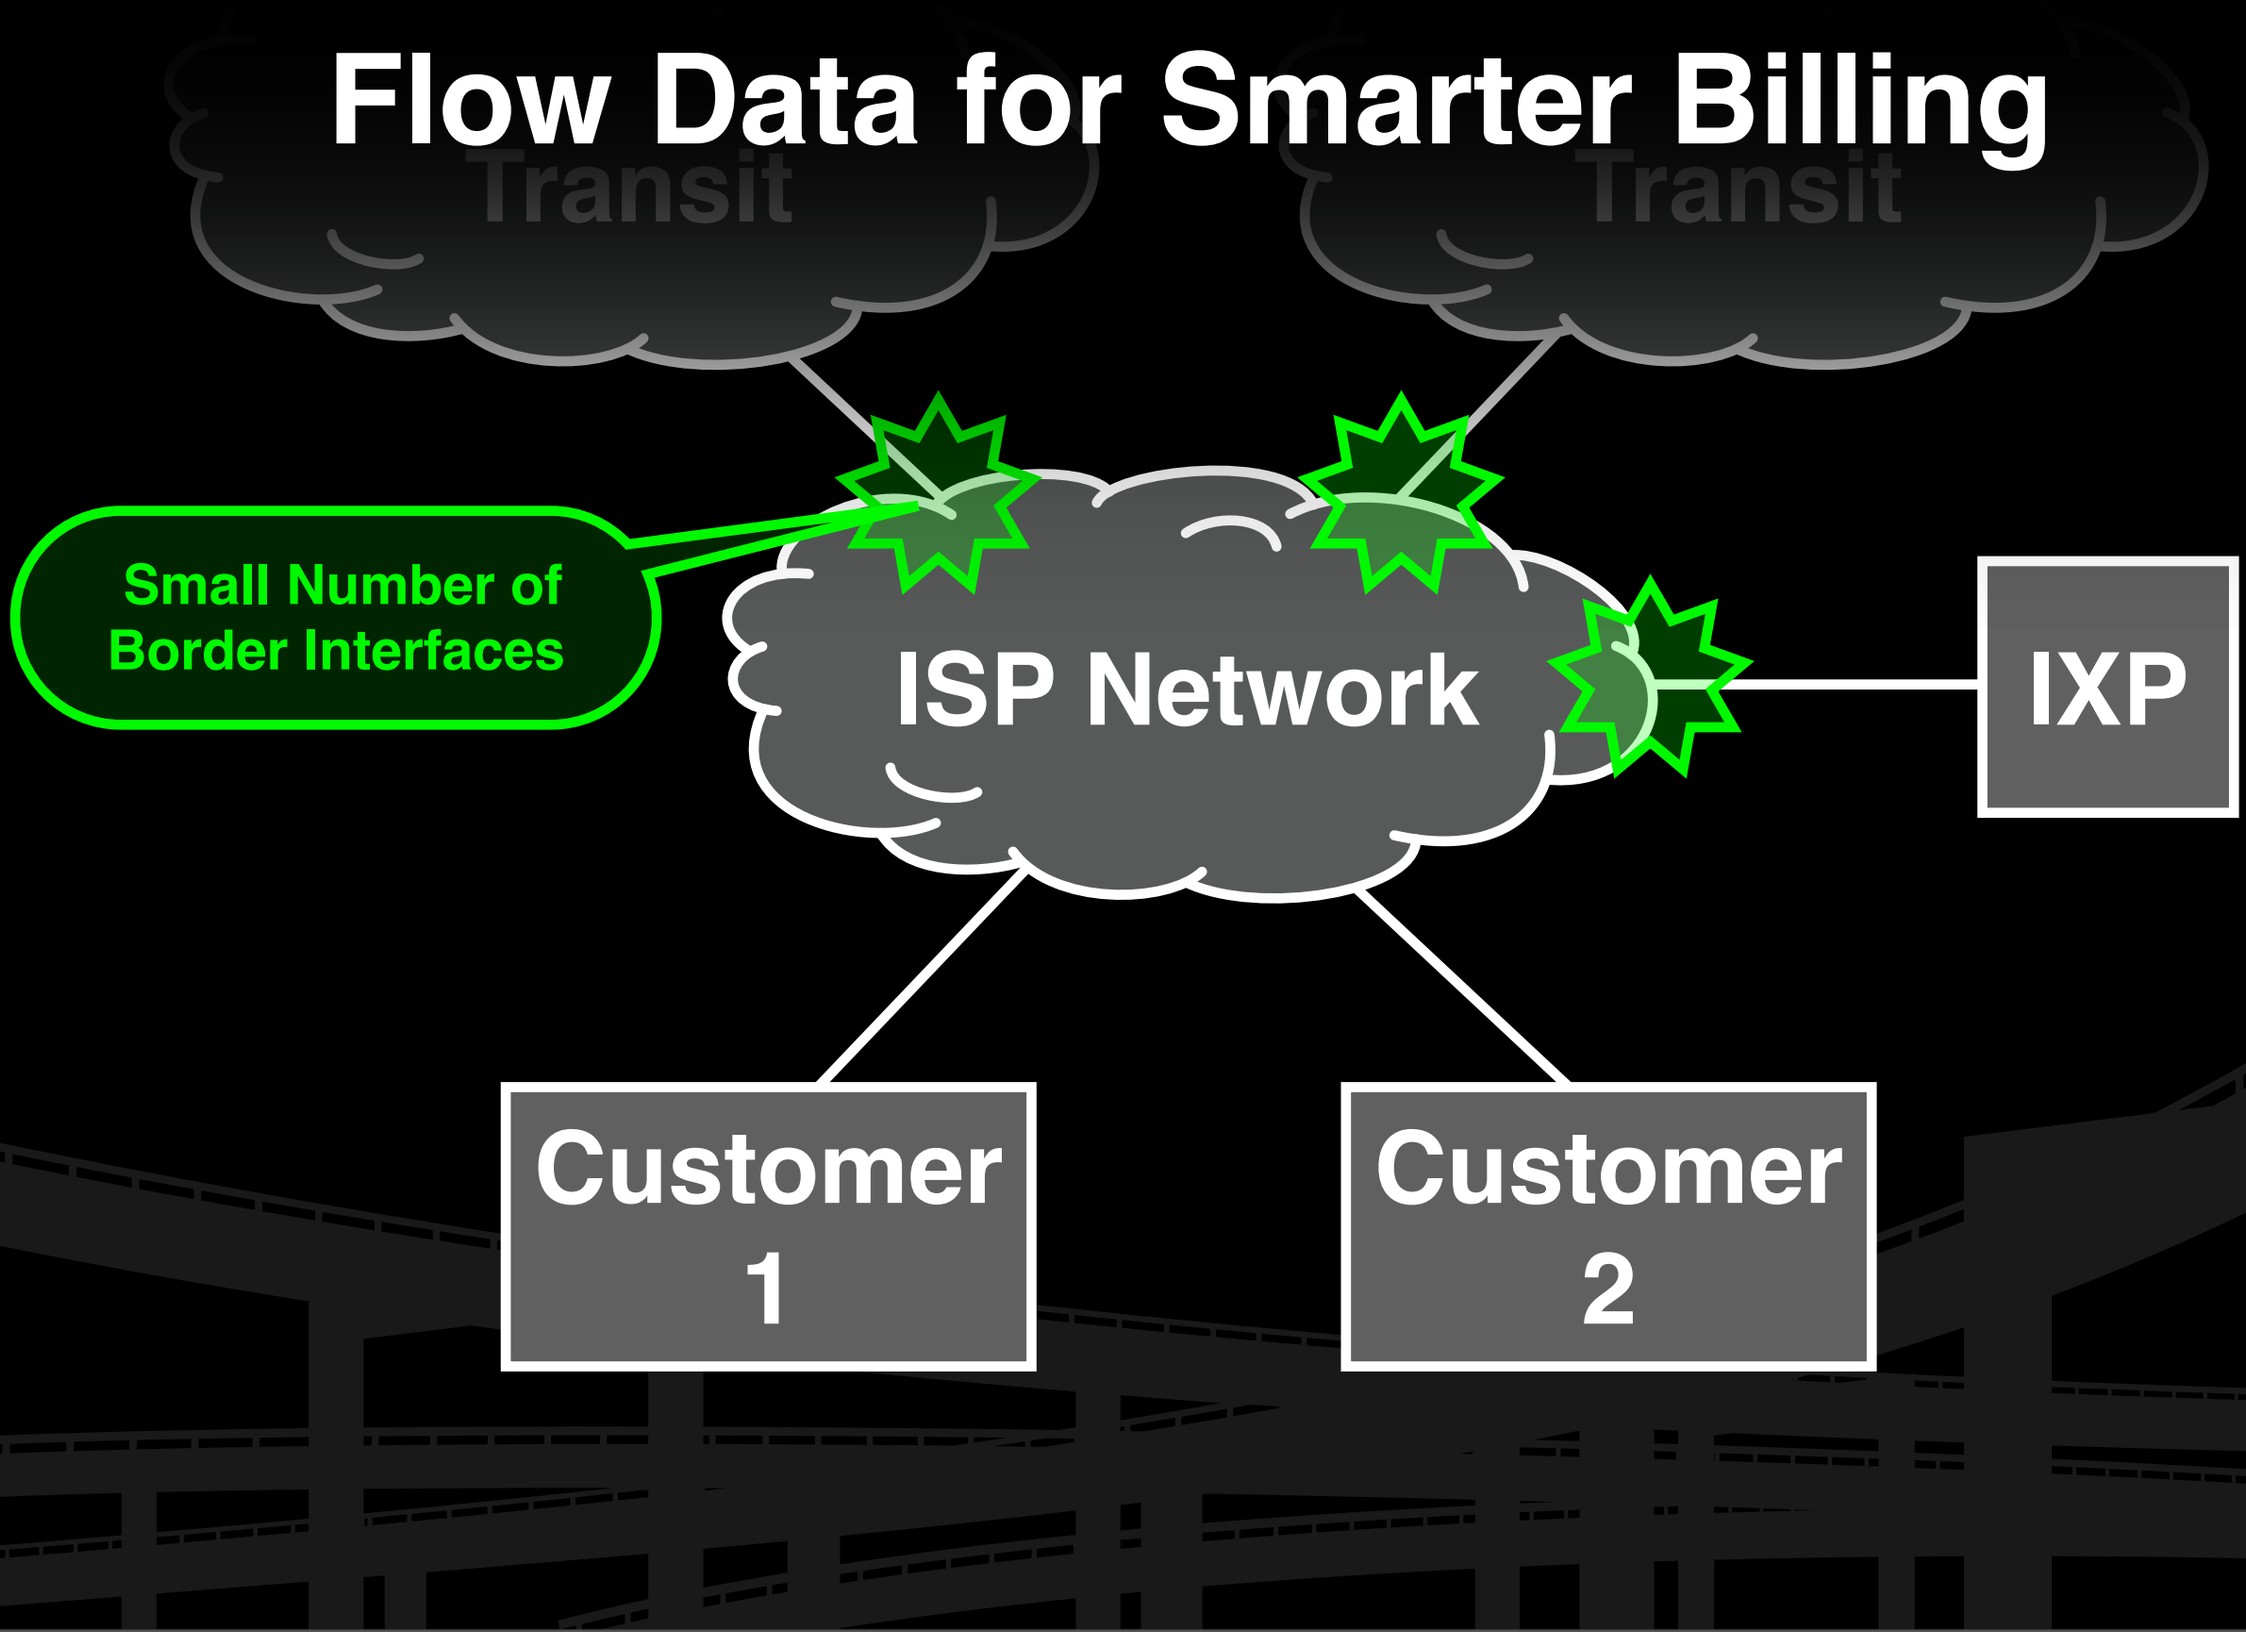
Flow Data for Smarter Billing



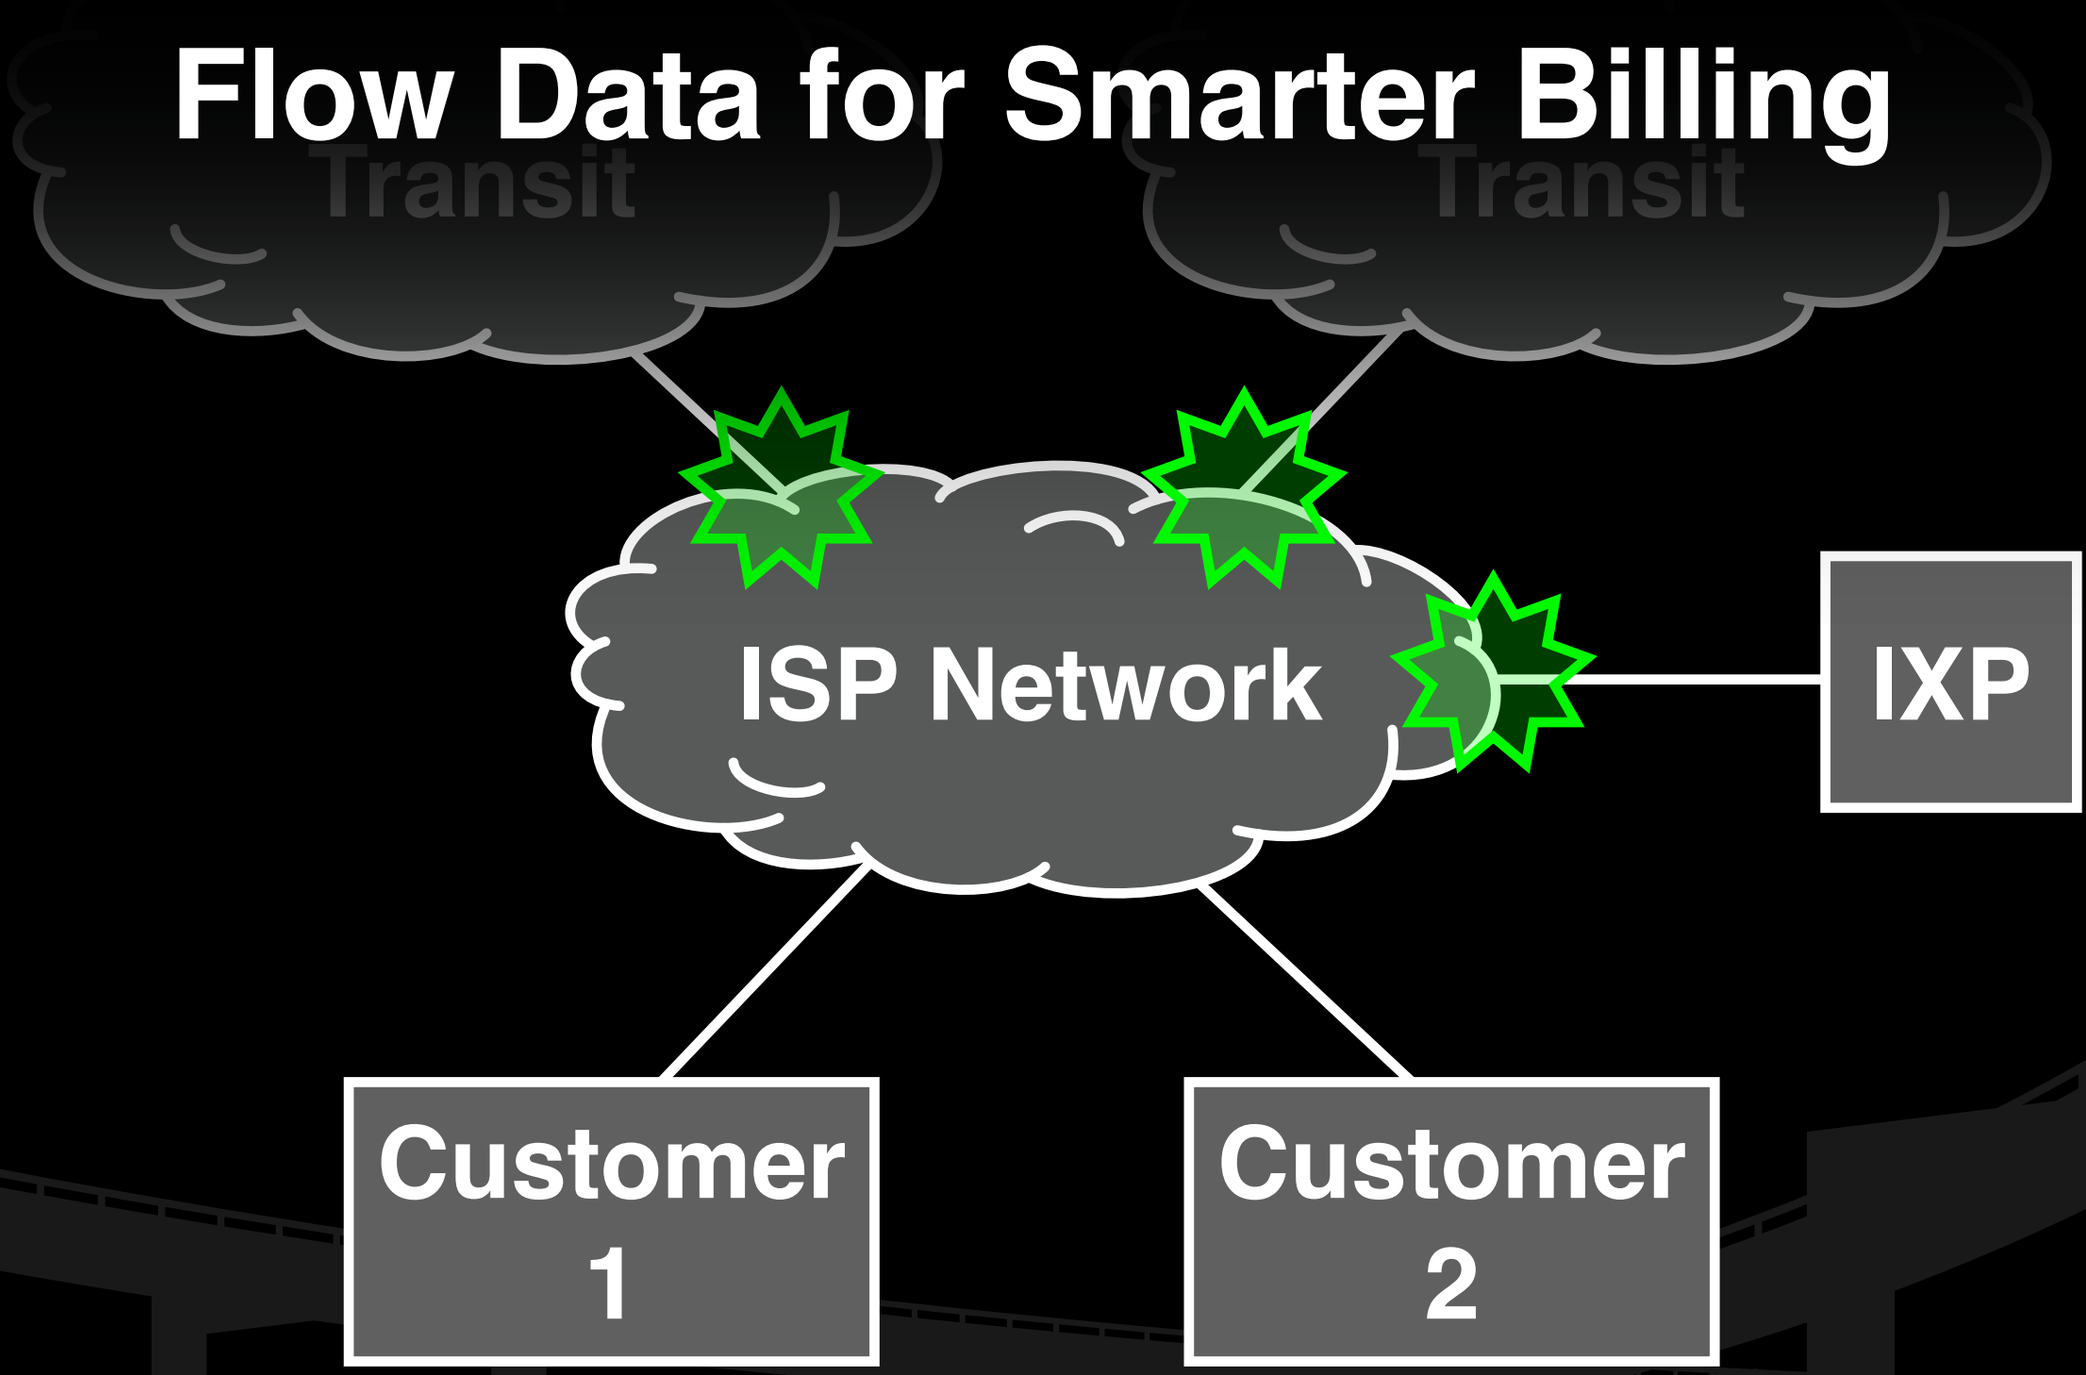
Flow Data for Smarter Billing



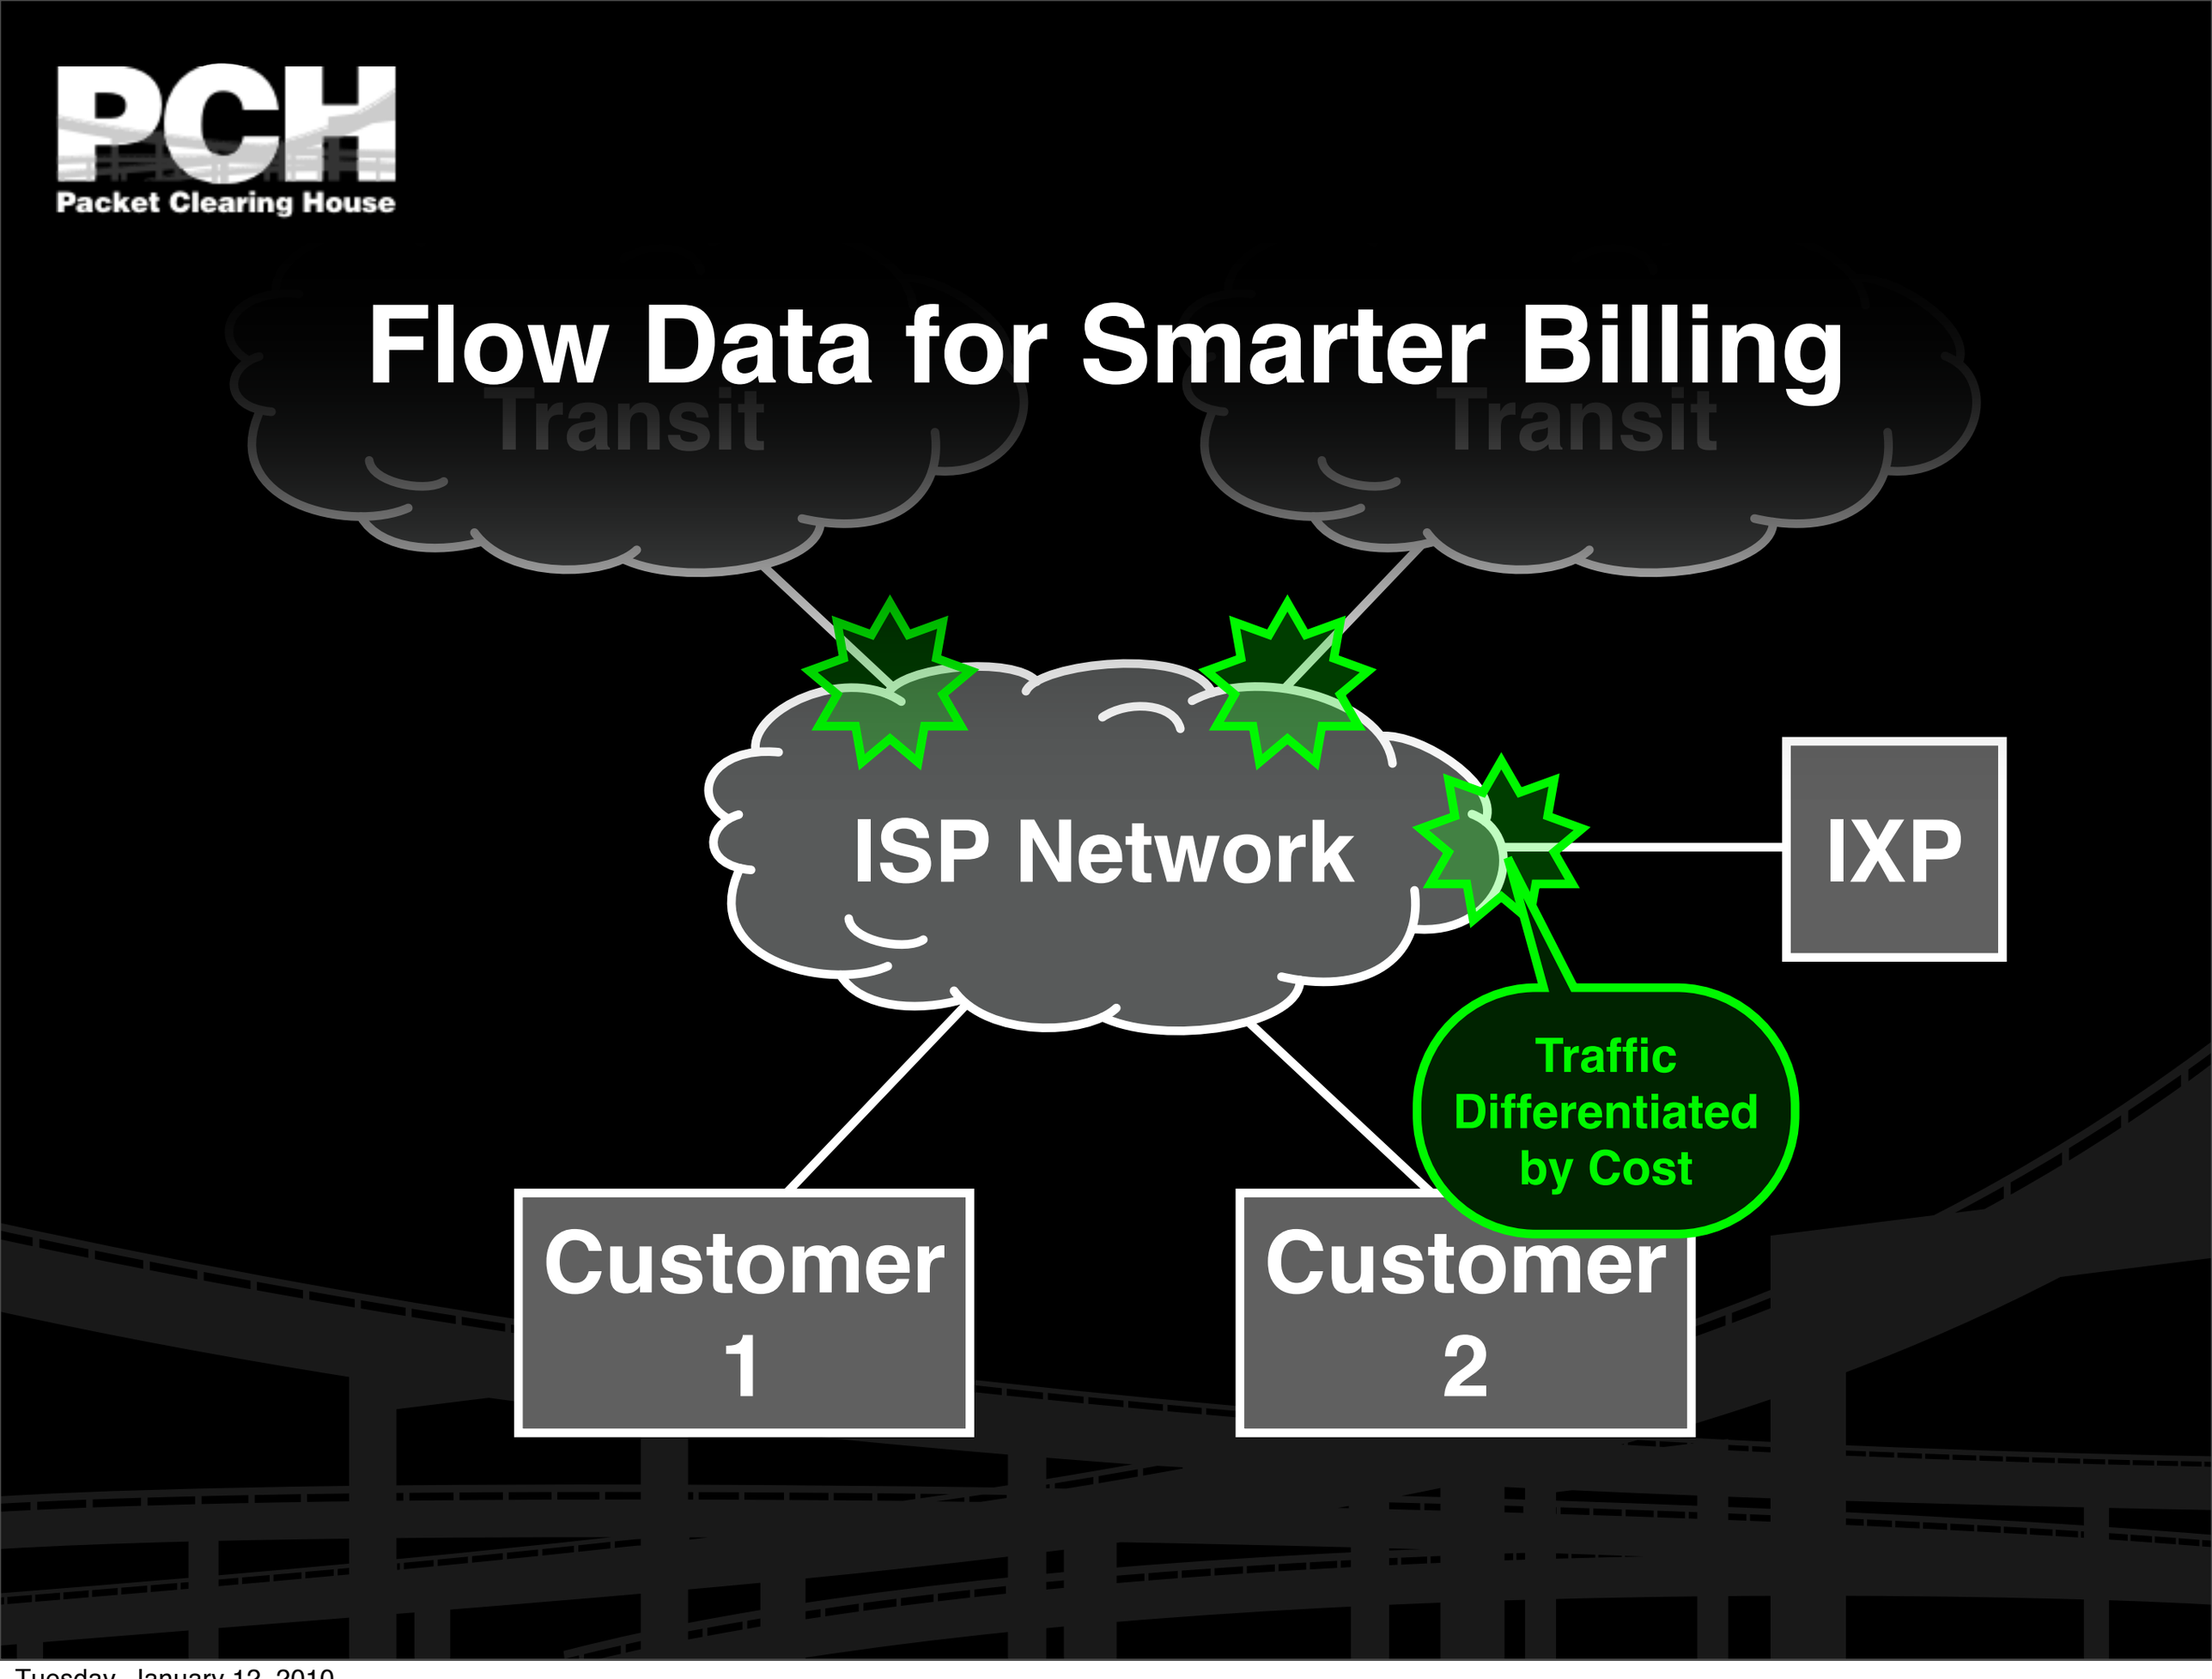
Flow Data for Smarter Billing



Flow Data for Smarter Billing

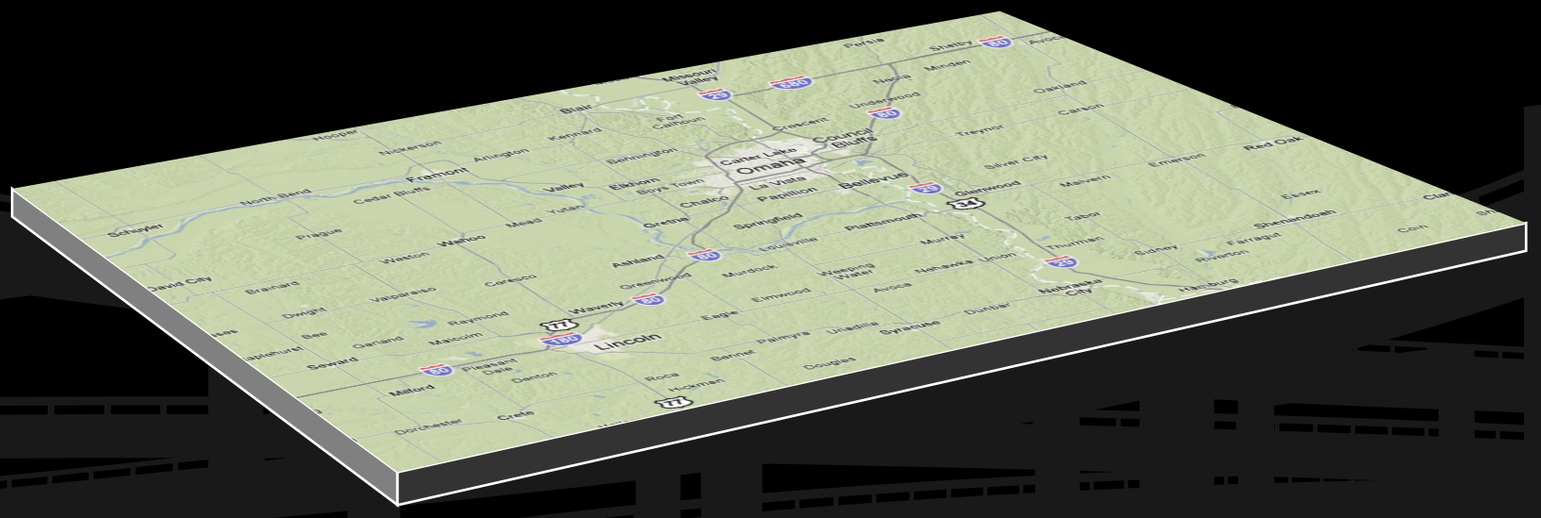


Flow Data for Smarter Billing



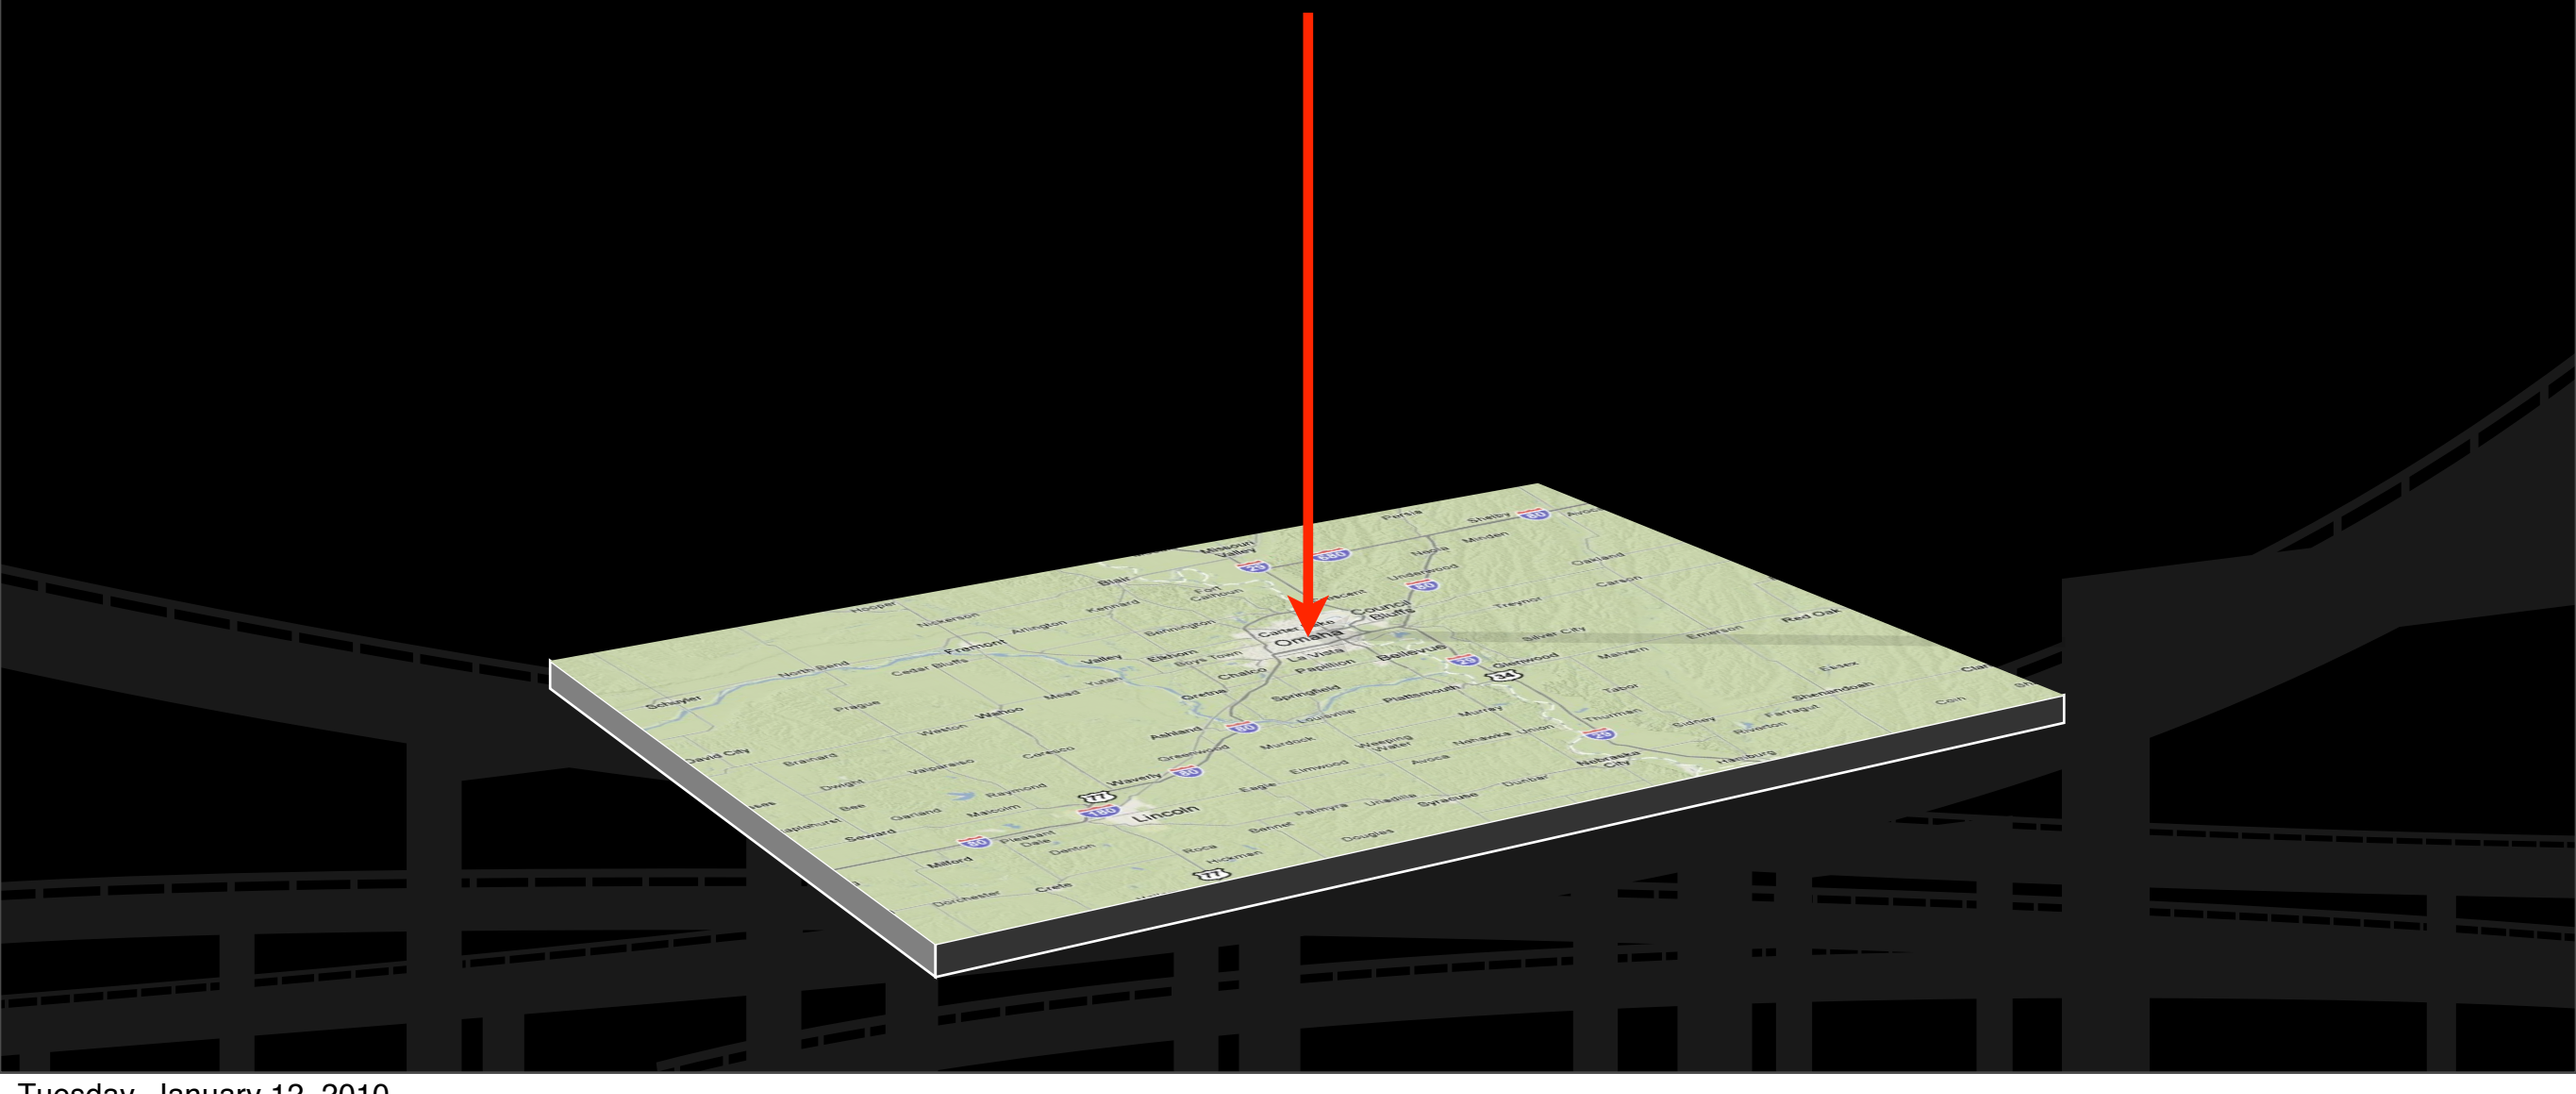
PCH

Packet Clearing House

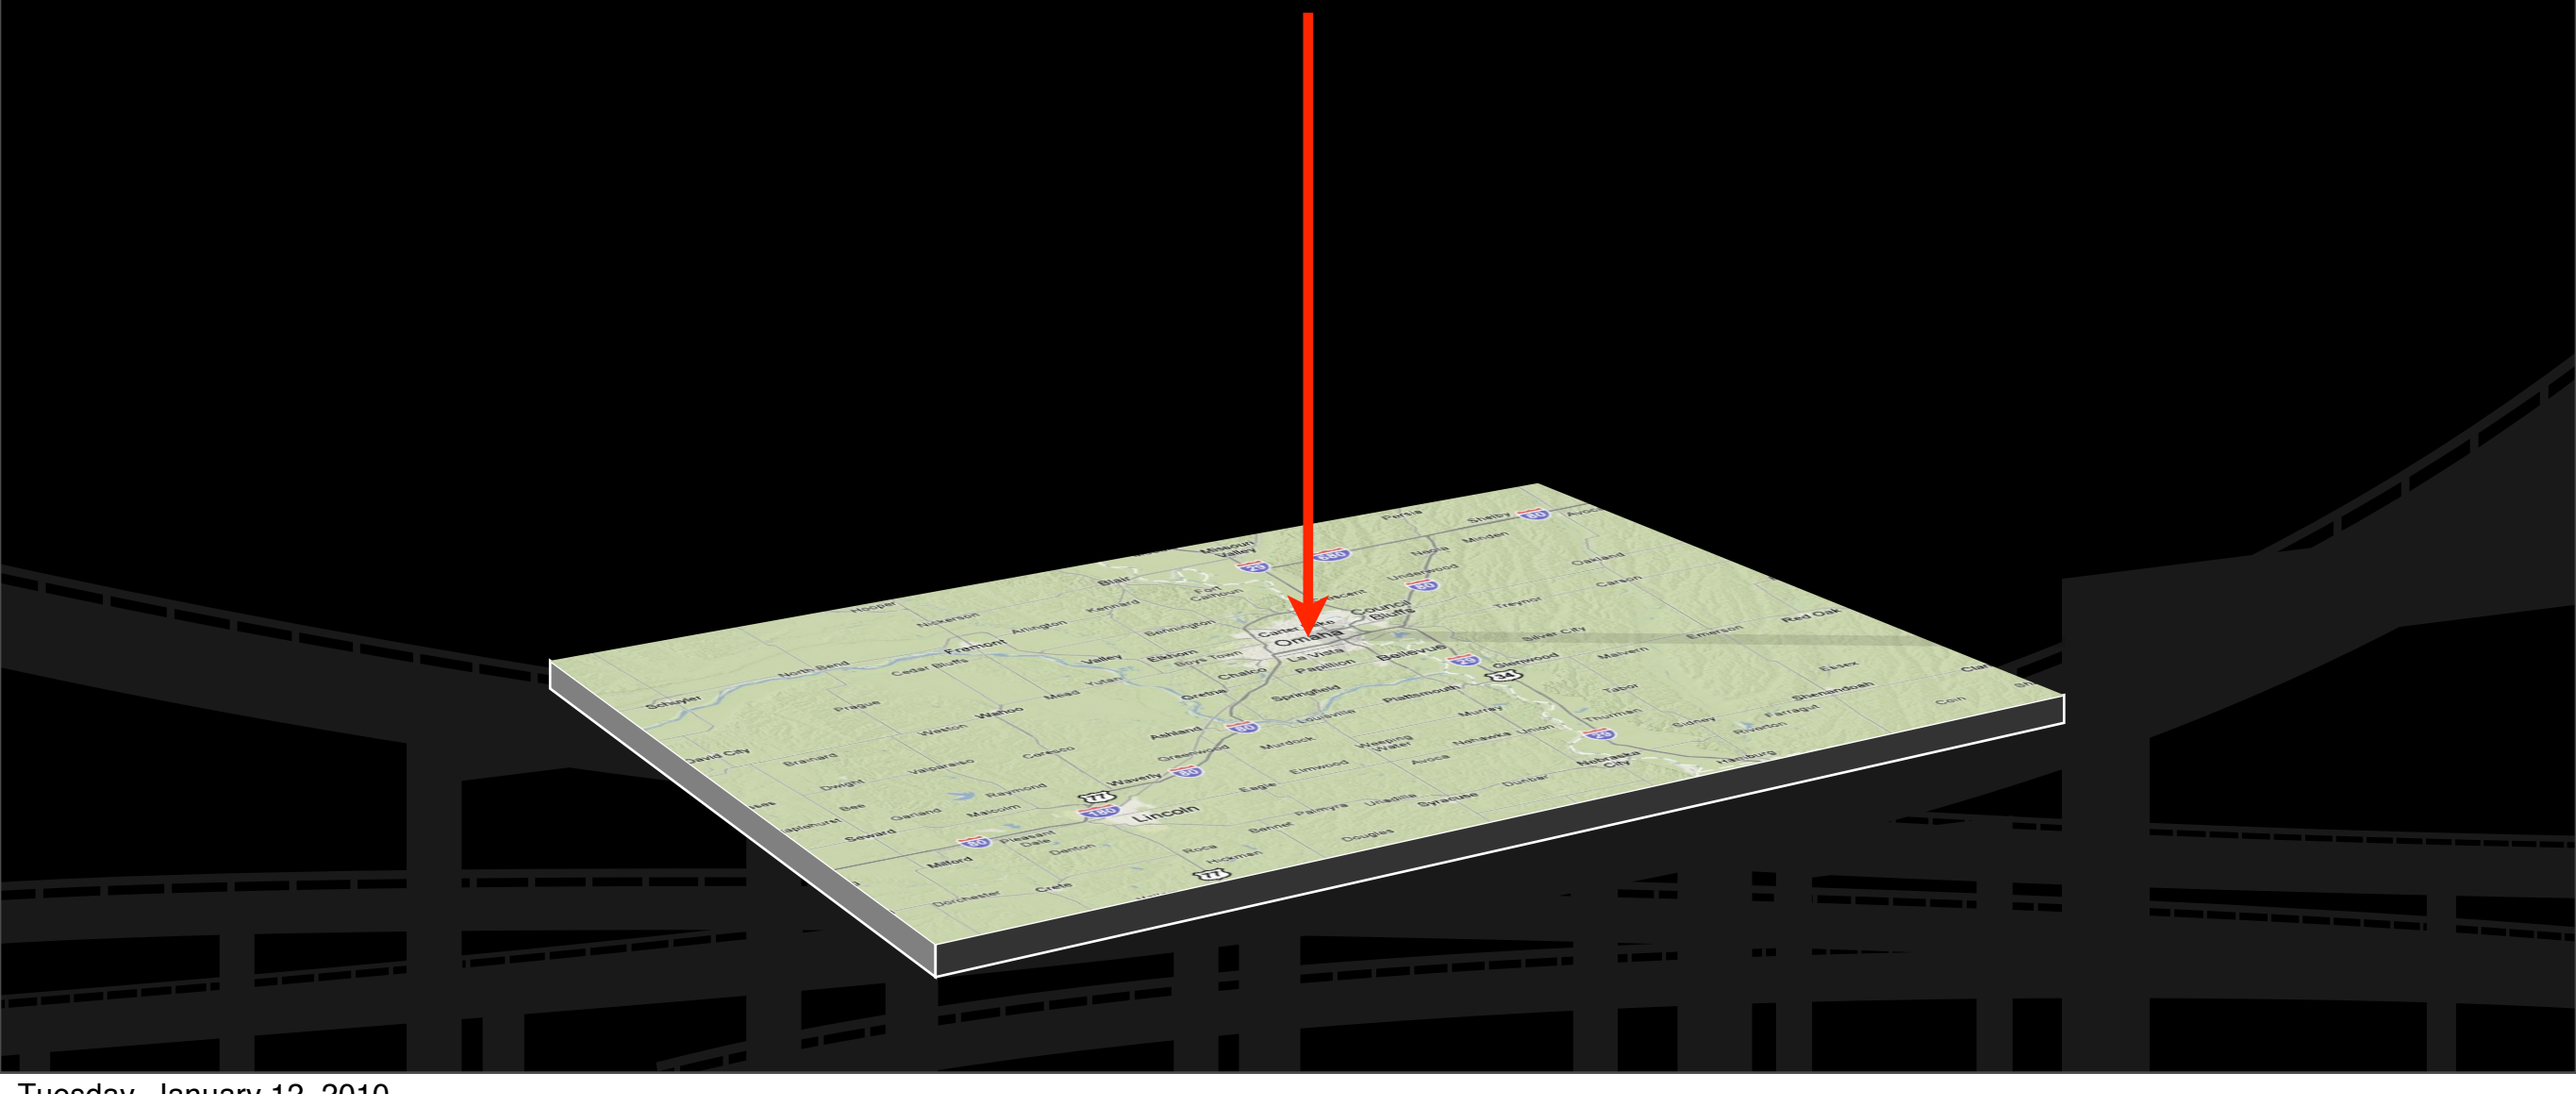


Tuesday, January 12, 2010

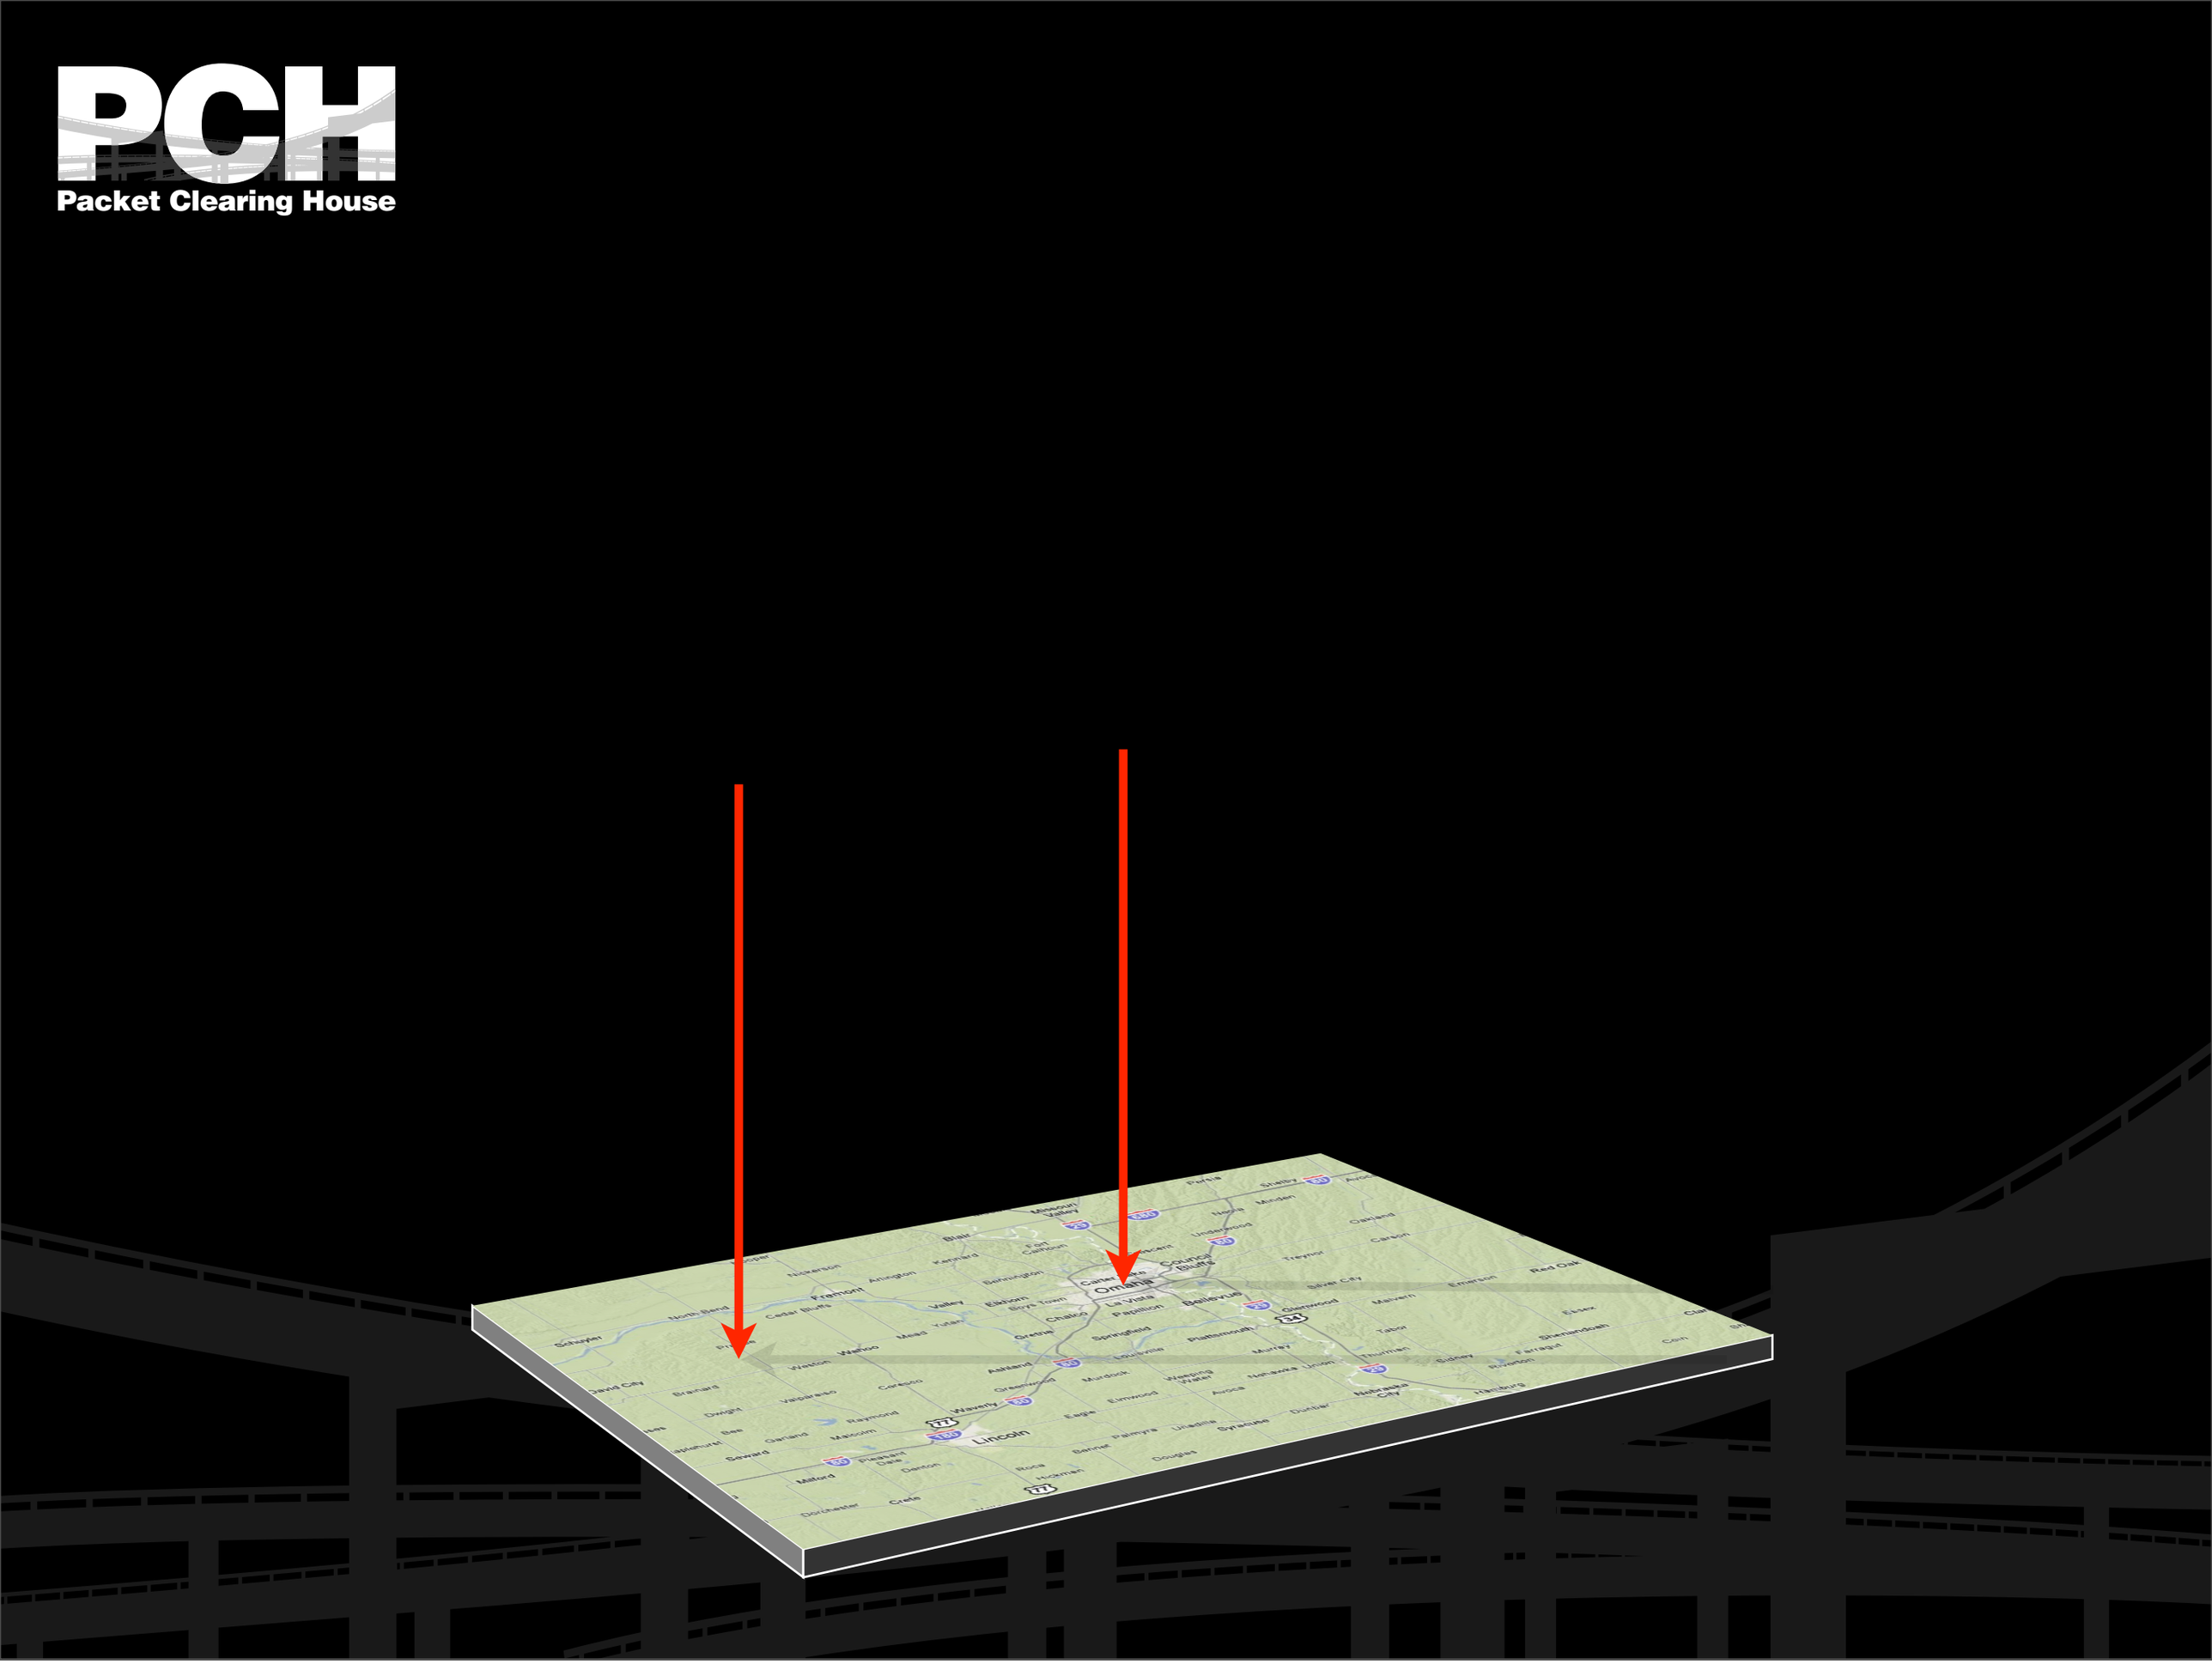
PCH
Packet Clearing House



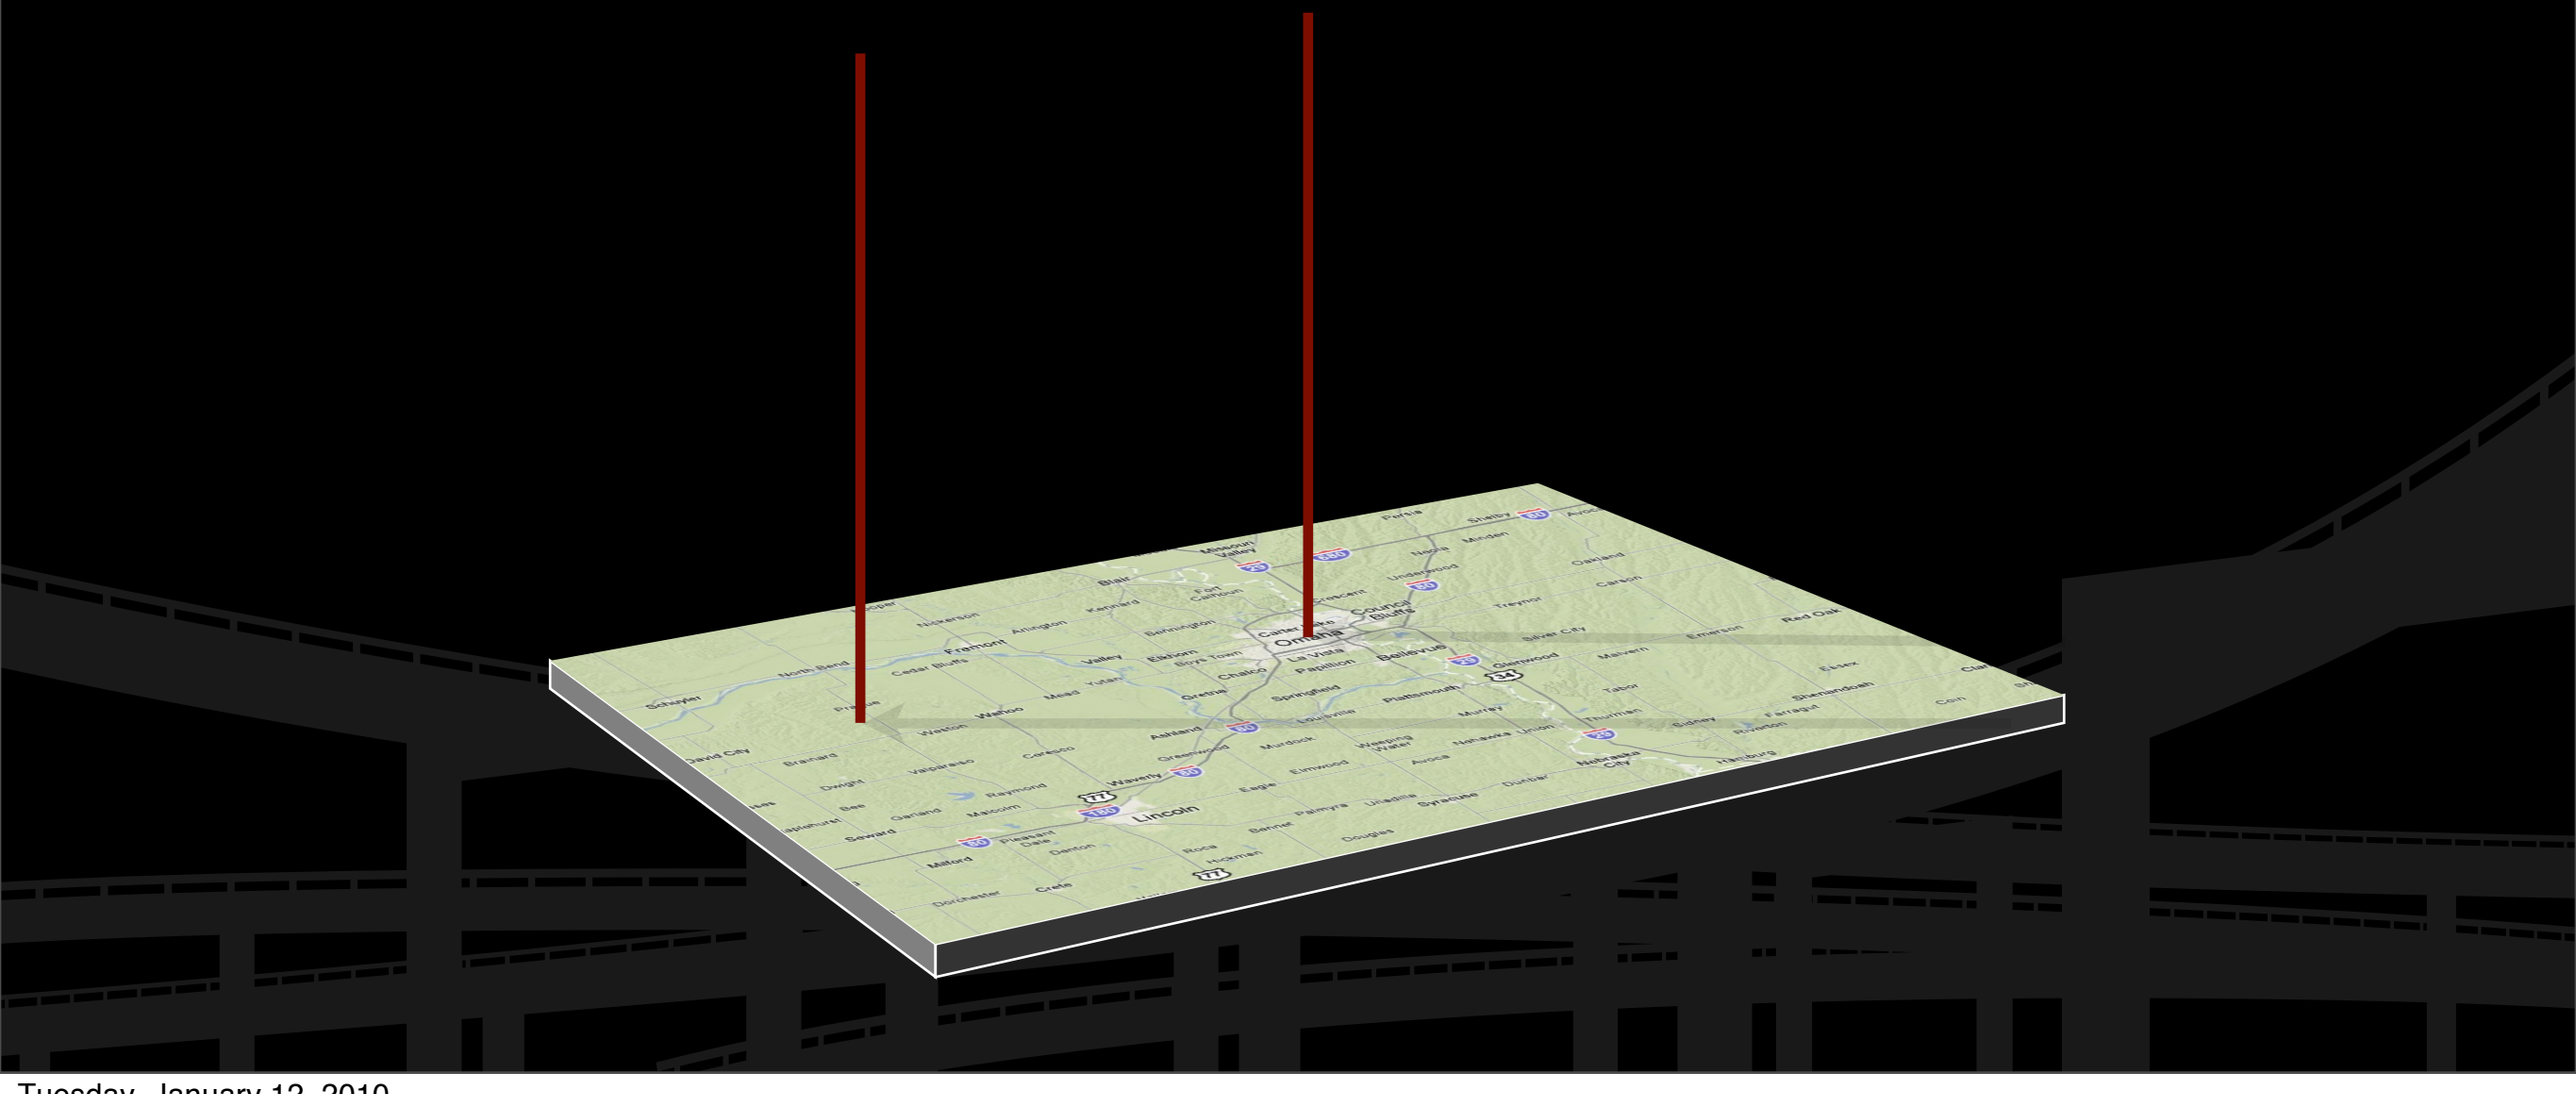
PCH
Packet Clearing House



PCH
Packet Clearing House

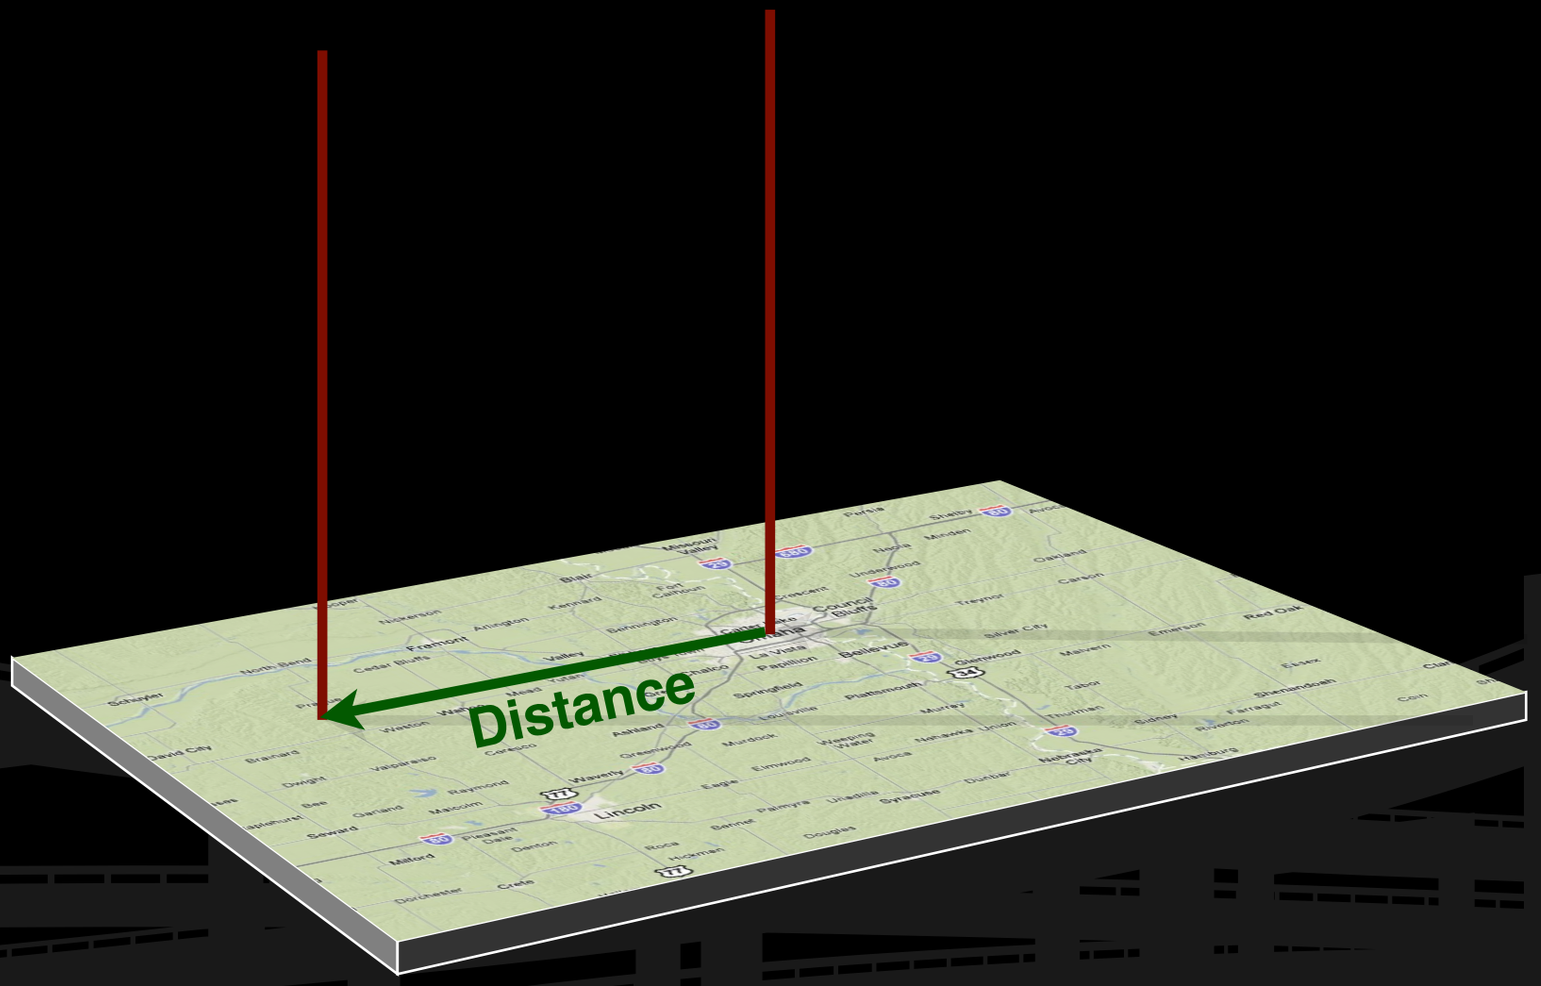

Omaha
Council Bluffs
Bellevue
Lincoln
Fremont
Blair
Papillion
La Vista
Ashland
Wahoo
Valley
Waverly
Seward
Milford
Crete
Beatrice
Nebraska City
Falls City
Red Oak
Shenandoah
Sidney
Tabor
Thurman
Murray
Plattsmouth
Springfield
Gretna
Elkhorn
Bennington
Arlington
Kennard
Fort Calhoun
Missouri Valley
Logan
Modale
Underwood
Treynor
Carson
Silver City
Emerson
Hamburg
Farragut
Auburn
Dawson
Humboldt
Table Rock
Steinauer
Pawnee City
Lewiston
Syracuse
Dunbar
Palmyra
Unadilla
Eagle
Elmwood
Murdock
Greenwood
Weeping Water
Avoca
Cortland
Dewitt
Raymond
Malcolm
Garland
Bee
Exeter
Fairmont
Geneva
Fairfield
Hastings
Juniata
Holstein
Ong
Glenvil
Harvard
Sutton
Clay Center
Fairfield
Edgar
Inland
Trumbull
Deweese
Lawrence
Nelson
Superior
Hardy
Hubbell
Endicott
Daykin
Reynolds
Pickrell
Firth
Hickman
Pleasant Dale
Denton
Roca
Panama
Berks
Douglas
Sarpy
Cass
Otoe
Johnson
Nemaha
Richardson
Pawnee
Gage
Jefferson
Thayer
Saline
Fillmore
Clay
Webster
Nuckolls
Franklin
Harlan
Furnas
Gosper
Phelps
Kearney
Adams
Hall
Hamilton
York
Polk
Butler
Saunders
Lancaster
Seward
Dodge
Washington
Burt
Thurston
Wayne
Cuming
Stanton
Madison
Platte
Colfax
Dakota
Dixon
Cedar
Knox
Antelope
Pierce
Boone
Nance
Merrick
Howard
Greeley
Wheeler
Valley
Sherman
Buffalo
Hall
Hamilton
Clay
Adams
Kearney
Phelps
Gosper
Furnas
Harlan
Franklin
Webster
Nuckolls
Thayer
Jefferson
Gage
Saline
Fillmore
Clay
Nemaha
Richardson
Pawnee
Johnson
Otoe
Lancaster
Saunders
Cass
Sarpy
Douglas

PCH
Packet Clearing House

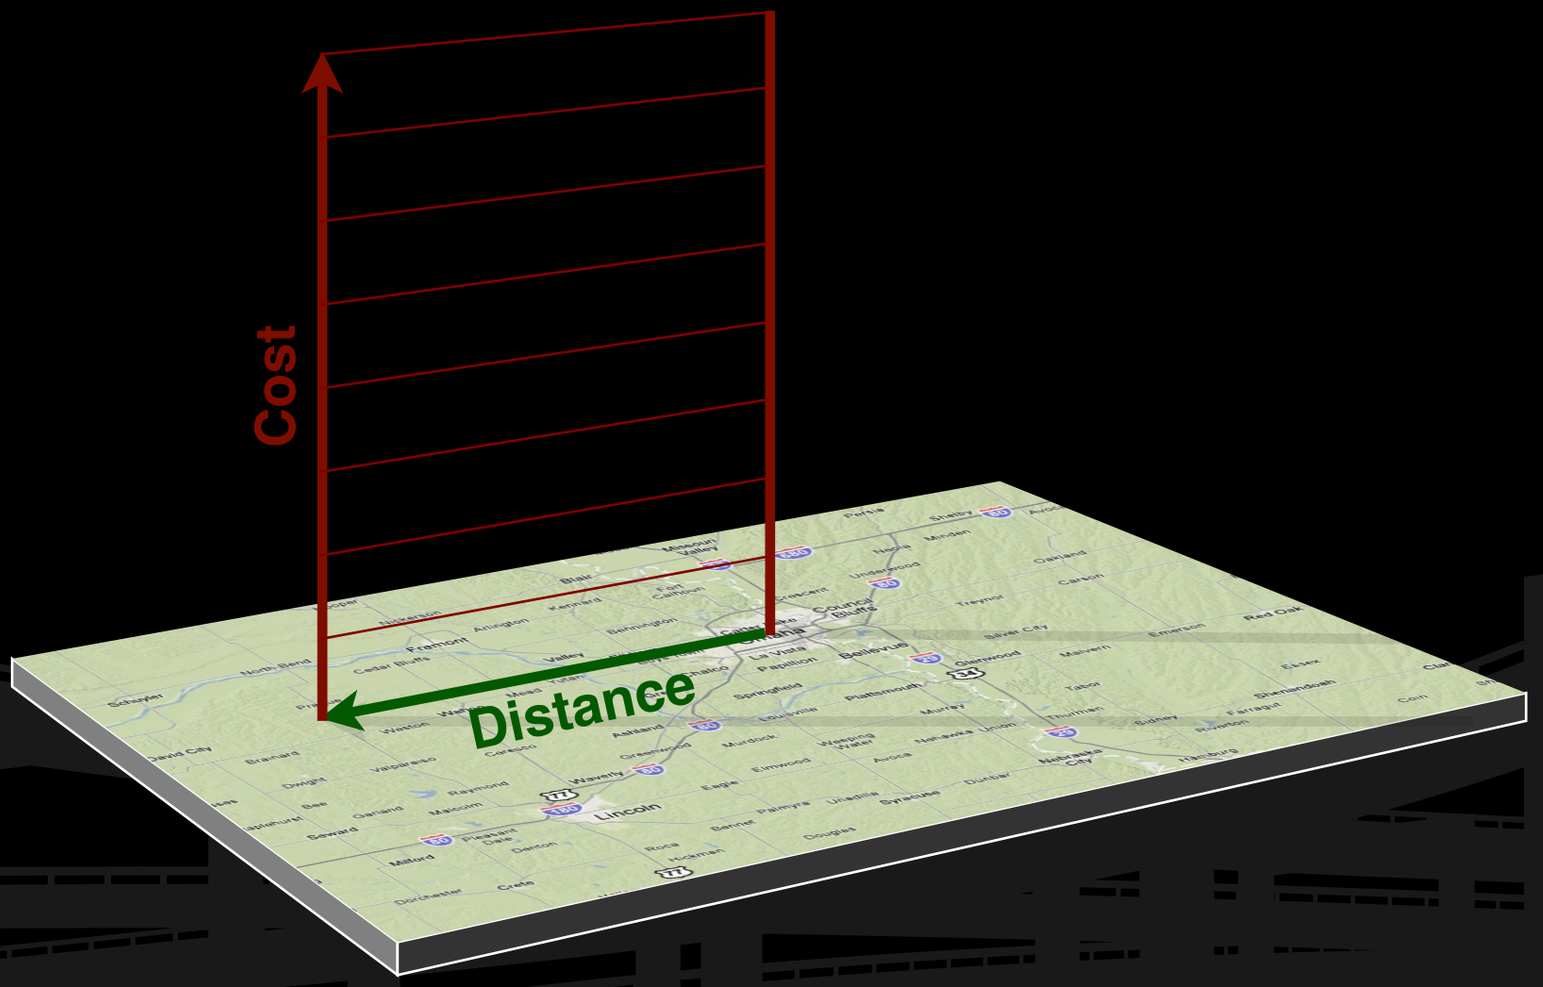


Tuesday, January 12, 2010

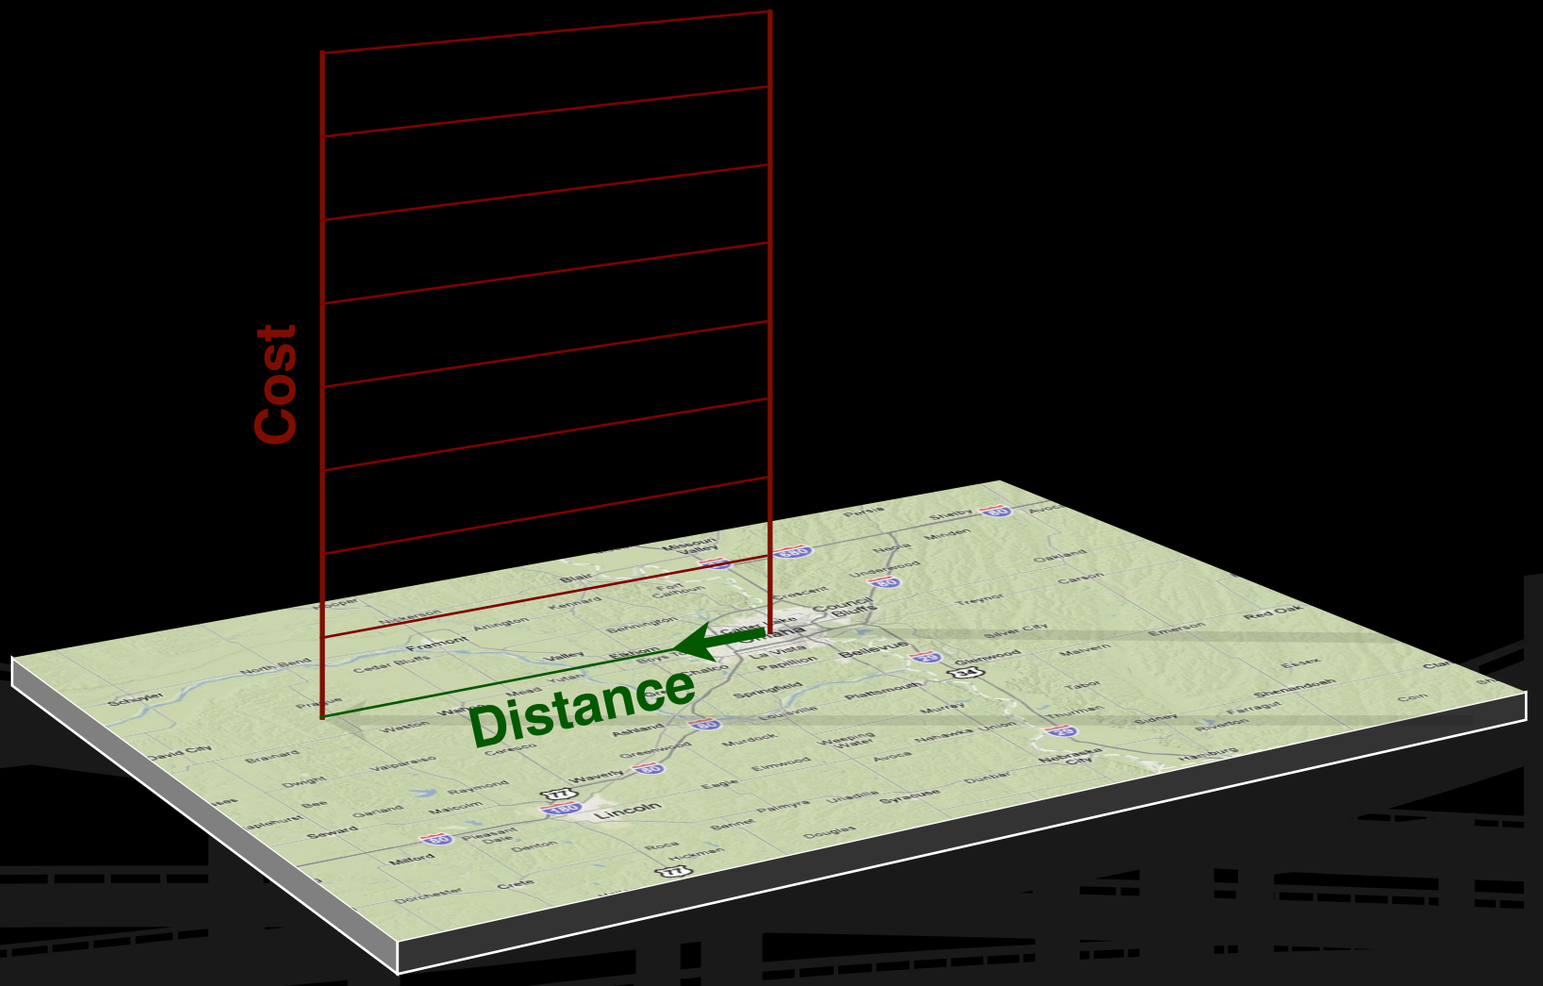
PCH
Packet Clearing House



PCH
Packet Clearing House

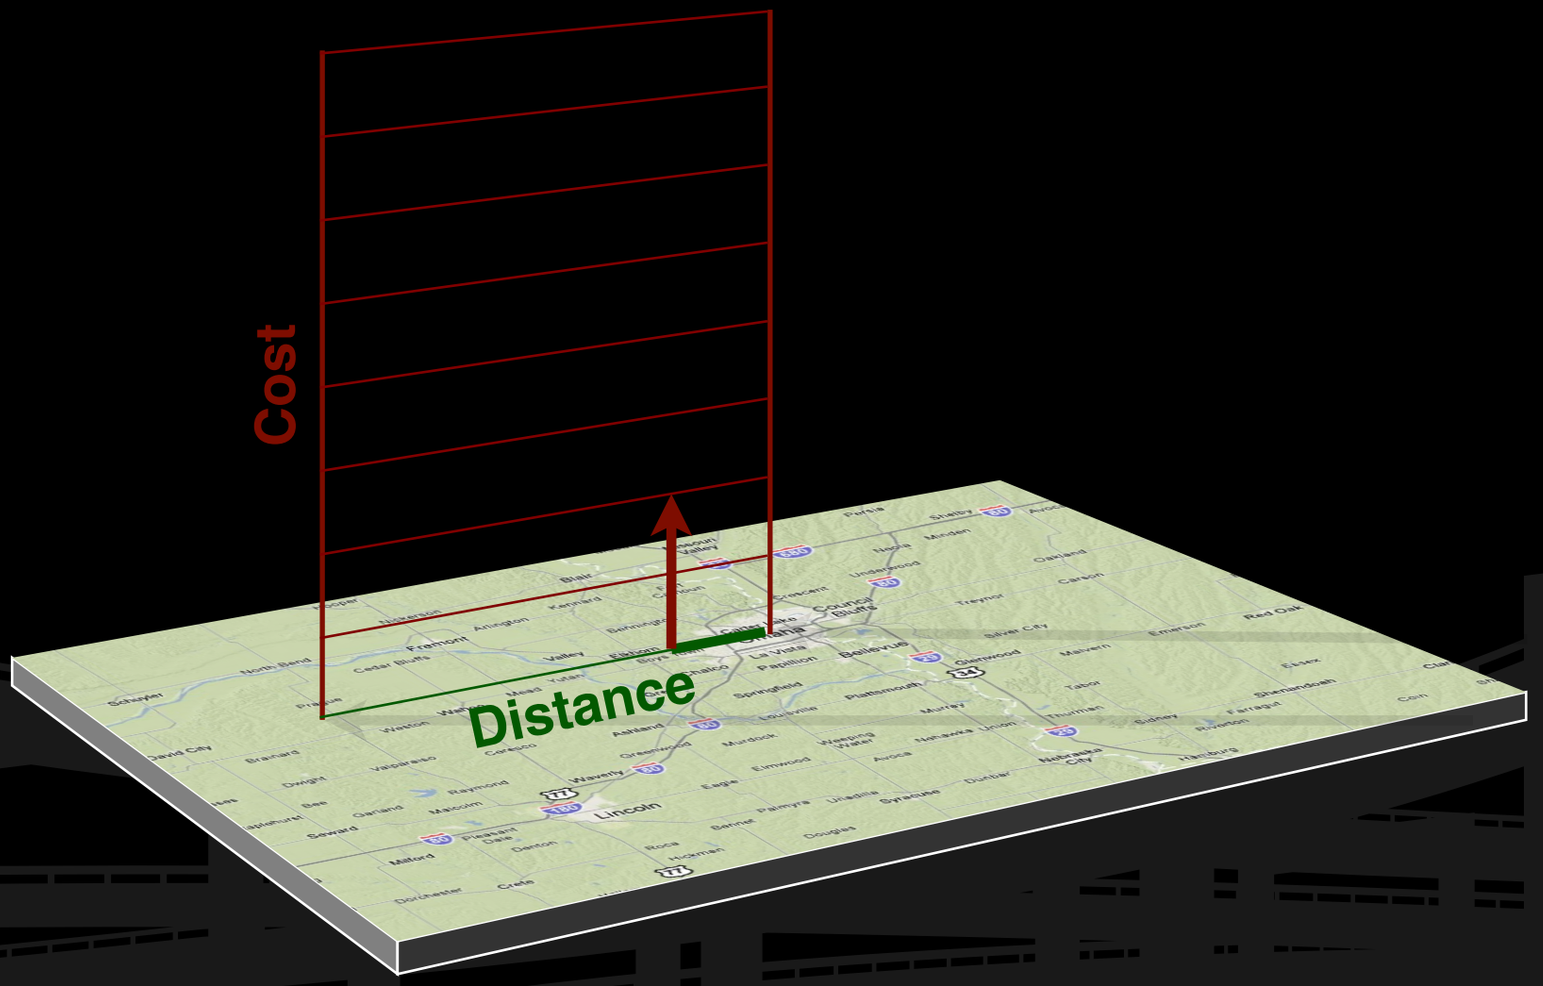


PCH
Packet Clearing House

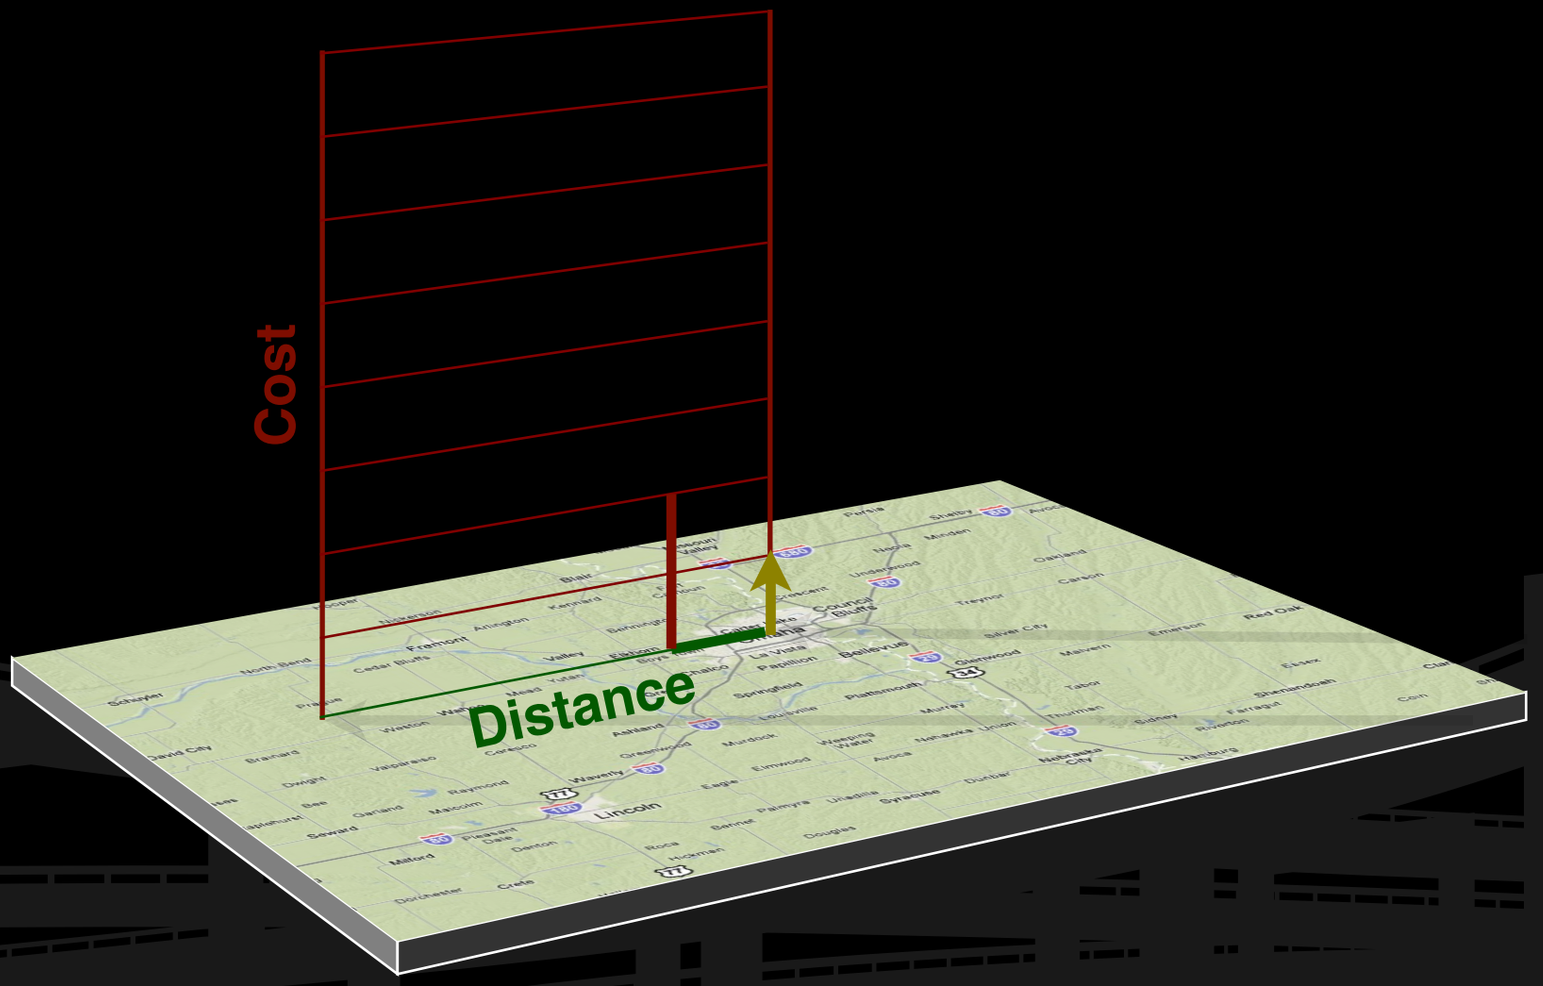


Tuesday, January 12, 2010

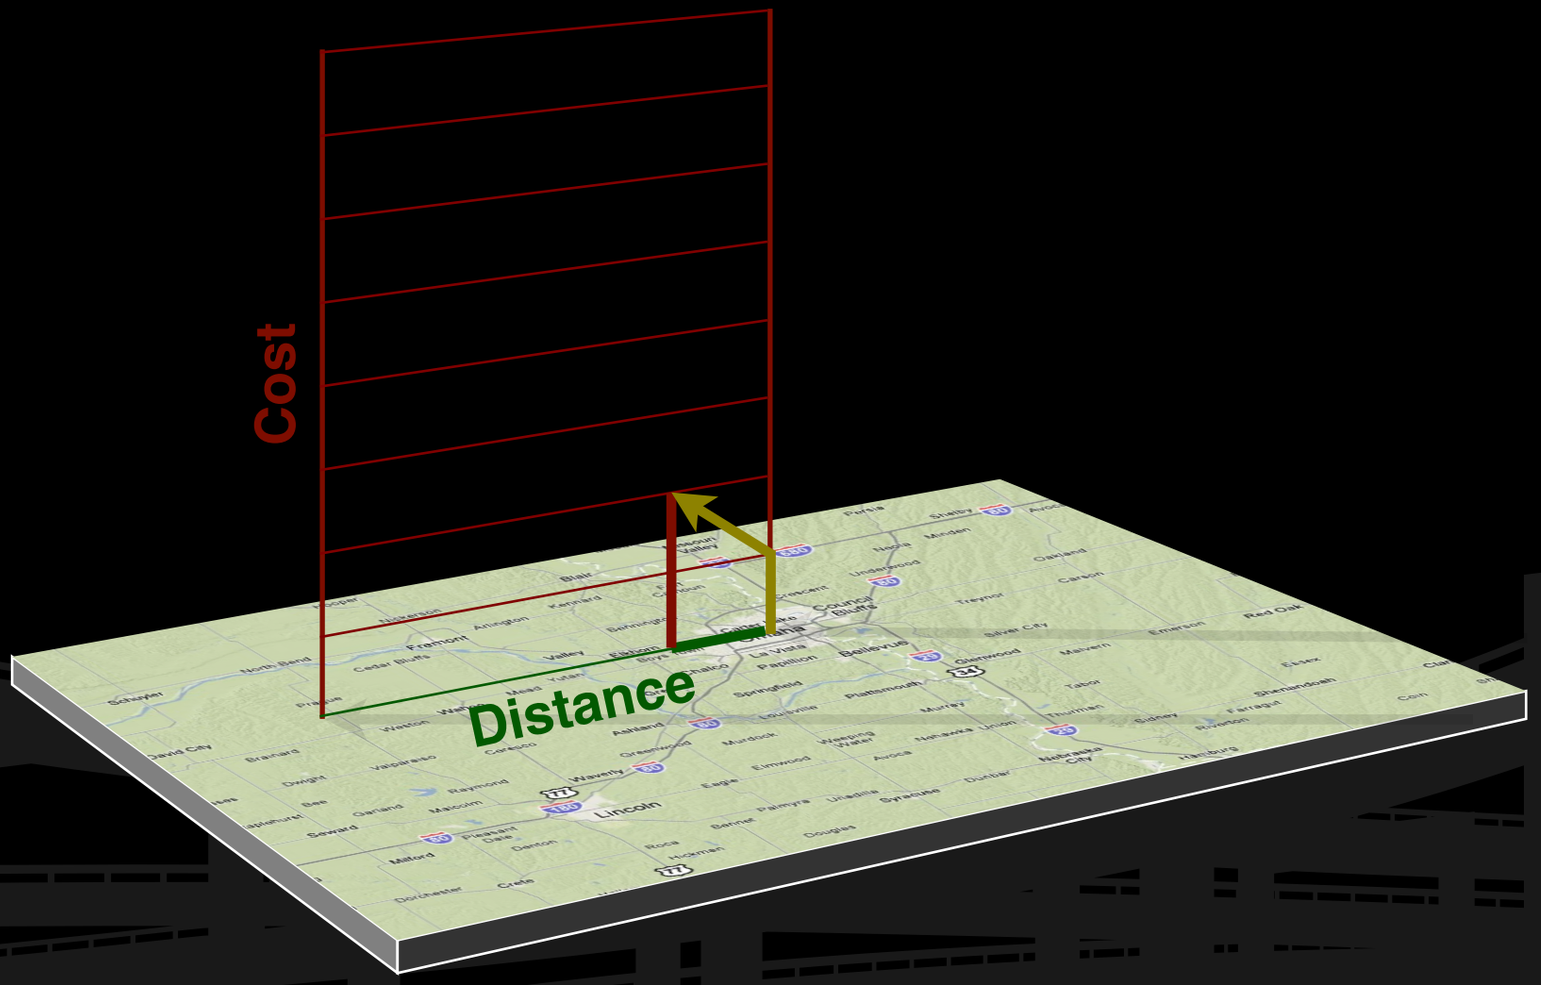
PCH
Packet Clearing House



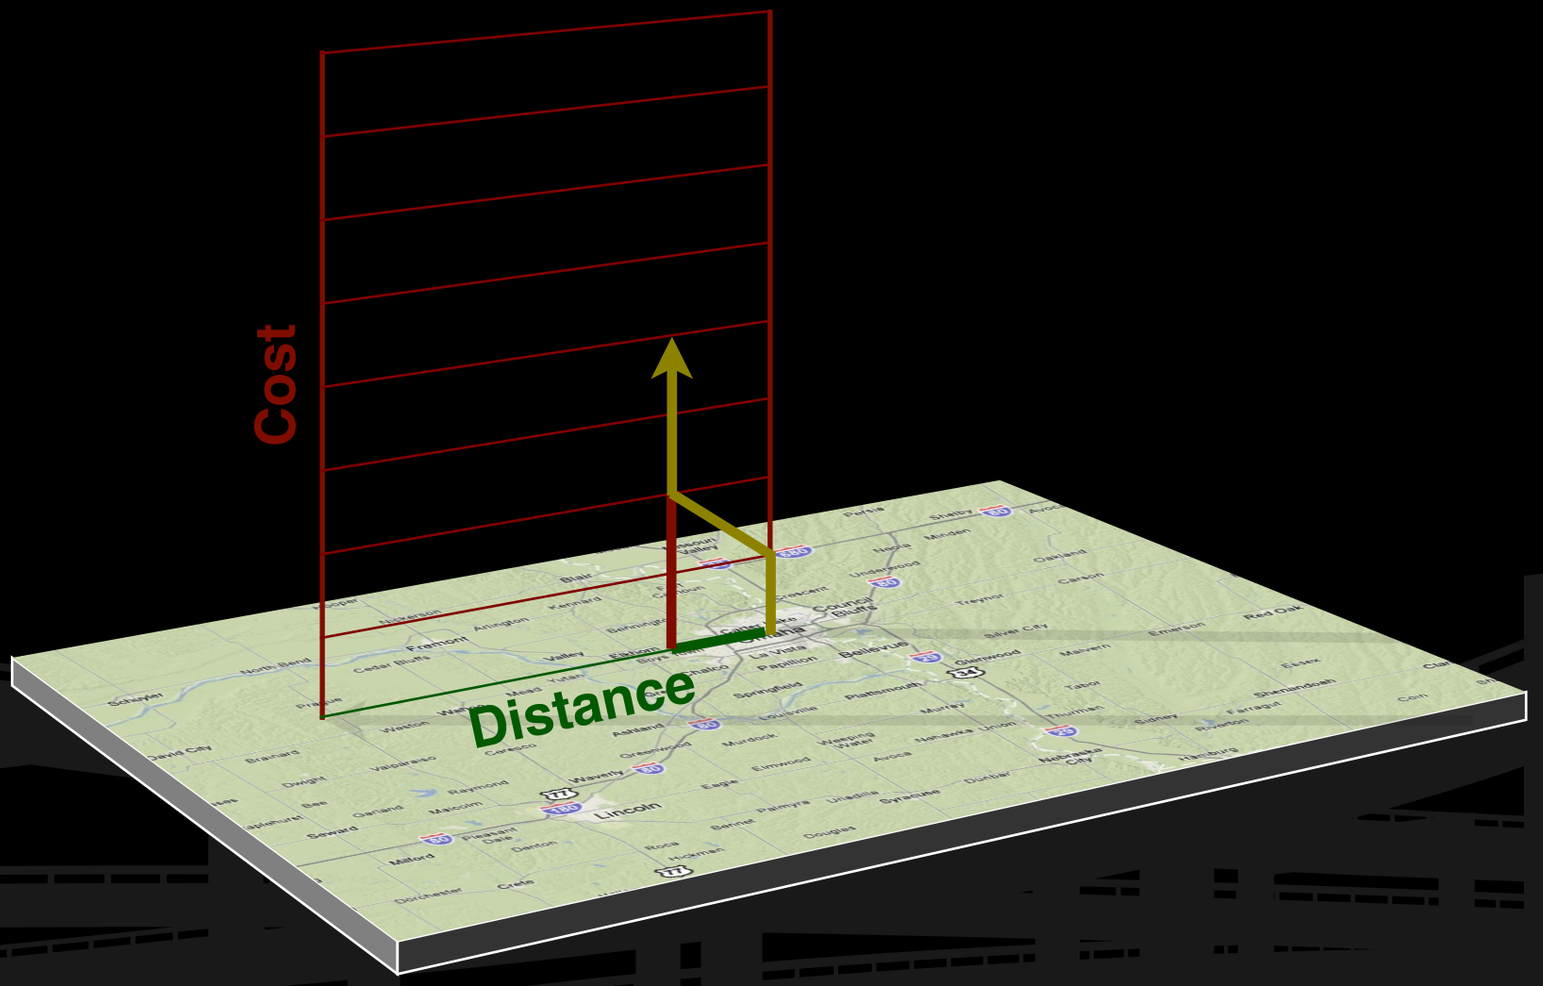
PCH
Packet Clearing House



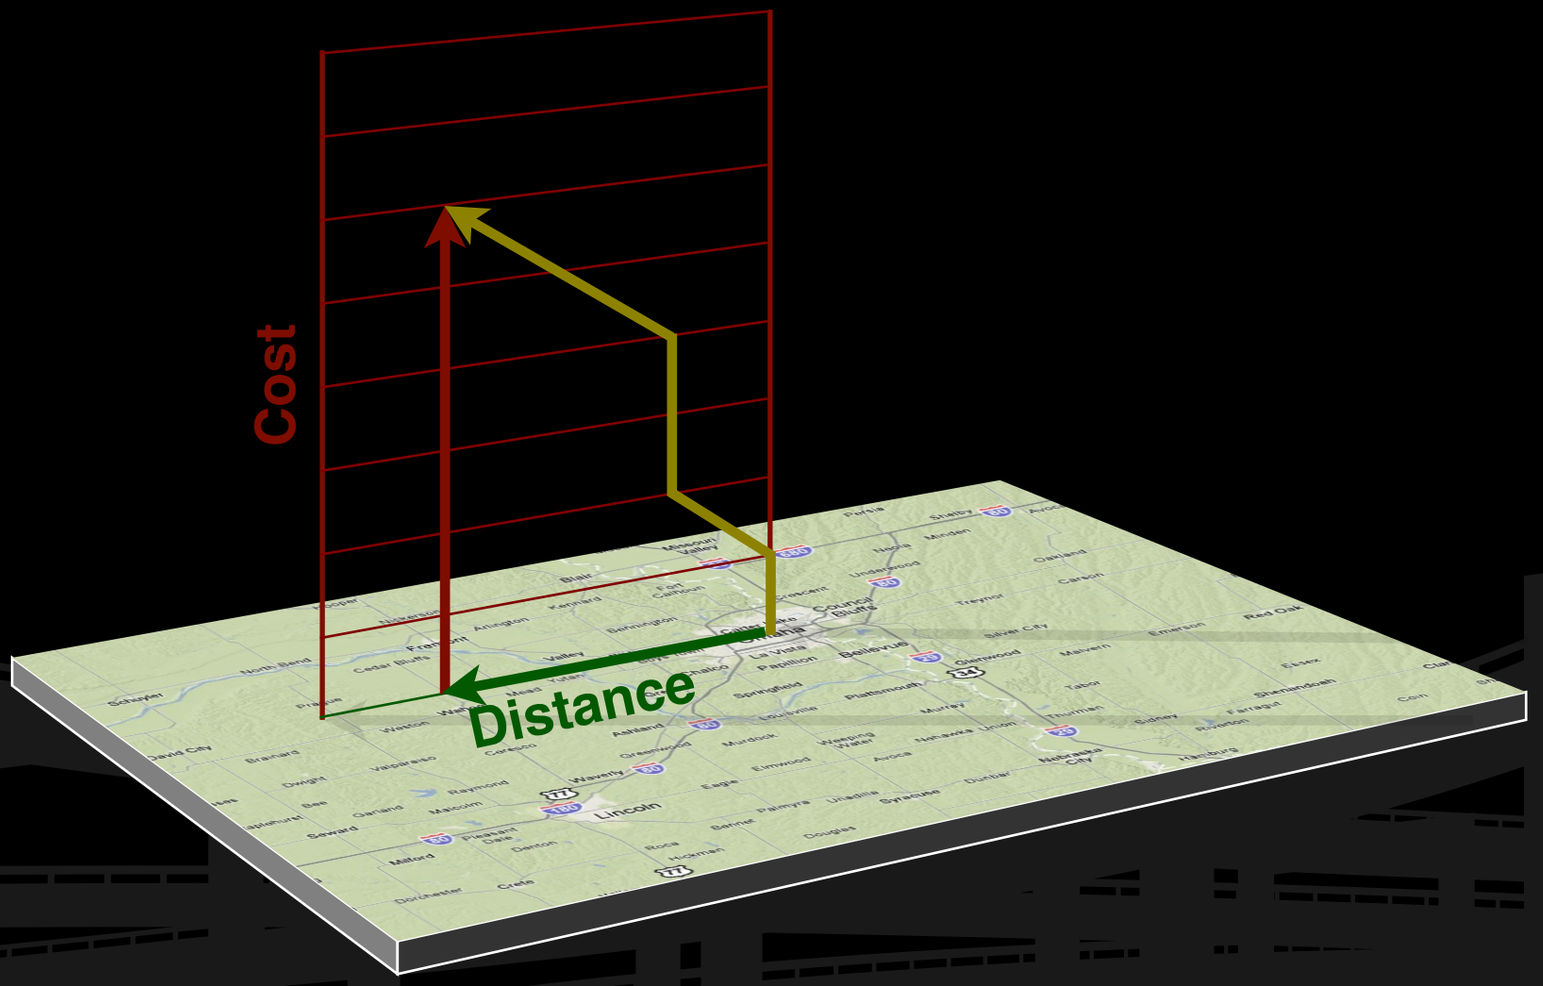
PCH
Packet Clearing House


Cost
Distance
Omaha
Bellevue
Council Bluffs
Lincoln
Fremont
Blair
Papillion
Plattsmouth
Nebraska City
Seward
Wahoo
Norfolk
Schuyler
Columbus
Hastings
Kearney
Grand Island
Beatrice
Falls City
Tecumseh
Auburn
Sidney
Shenandoah
Red Oak
Clarinda
Malvern
Glenwood
Carson
Oakland
Underwood
Treynor
Silver City
Emerson
Tabor
Thurman
Farragut
Riverton
Hamburg
Percival
Atchison
Rock Port
Maryville
Stanberry
Albany
Gallatin
Cameron
Savannah
St Joseph
Weston
Platte Woods
Parkville
Kansas City
Liberty
Excelsior Springs
Smithville
Plattsburg
Lathrop
Trimble
Kearney
Holt
Pleasant Hill
Harrisonville
Belton
Raymore
Peculiar
Grandview
Lee's Summit
Blue Springs
Independence
Sugar Creek
Buckner
Oak Grove
Higginsville
Odessa
Lexington
Wellington
Waverly
Denton
Malcolm
Raymond
Emerald
Bennet
Palmyra
Douglas
Syracuse
Dunbar
Union
Nehawka
Murdock
Elmwood
Eagle
Greenwood
Cedar Creek
Weeping Water
Avoca
Louisville
Springfield
La Vista
Papillion
Ashland
Wahoo
Yutan
Mead
Valley
Waterloo
Elkhorn
Bennington
Arlington
Kennard
Blair
Fort Calhoun
Missouri Valley
Logan
Woodbine
Modale
Pisgah
Persia
Shelby
Minden
Hancock
Underwood
Crescent
Honey Creek
Neola
Mineola
Crescent
Council Bluffs
Carter Lake
Omaha
Bellevue
Offutt AFB
Papillion
La Vista
Springfield
Louisville
Plattsmouth
Murray
Union
Nehawka
Weeping Water
Avoca
Elmwood
Eagle
Greenwood
Murdock
Cedar Creek

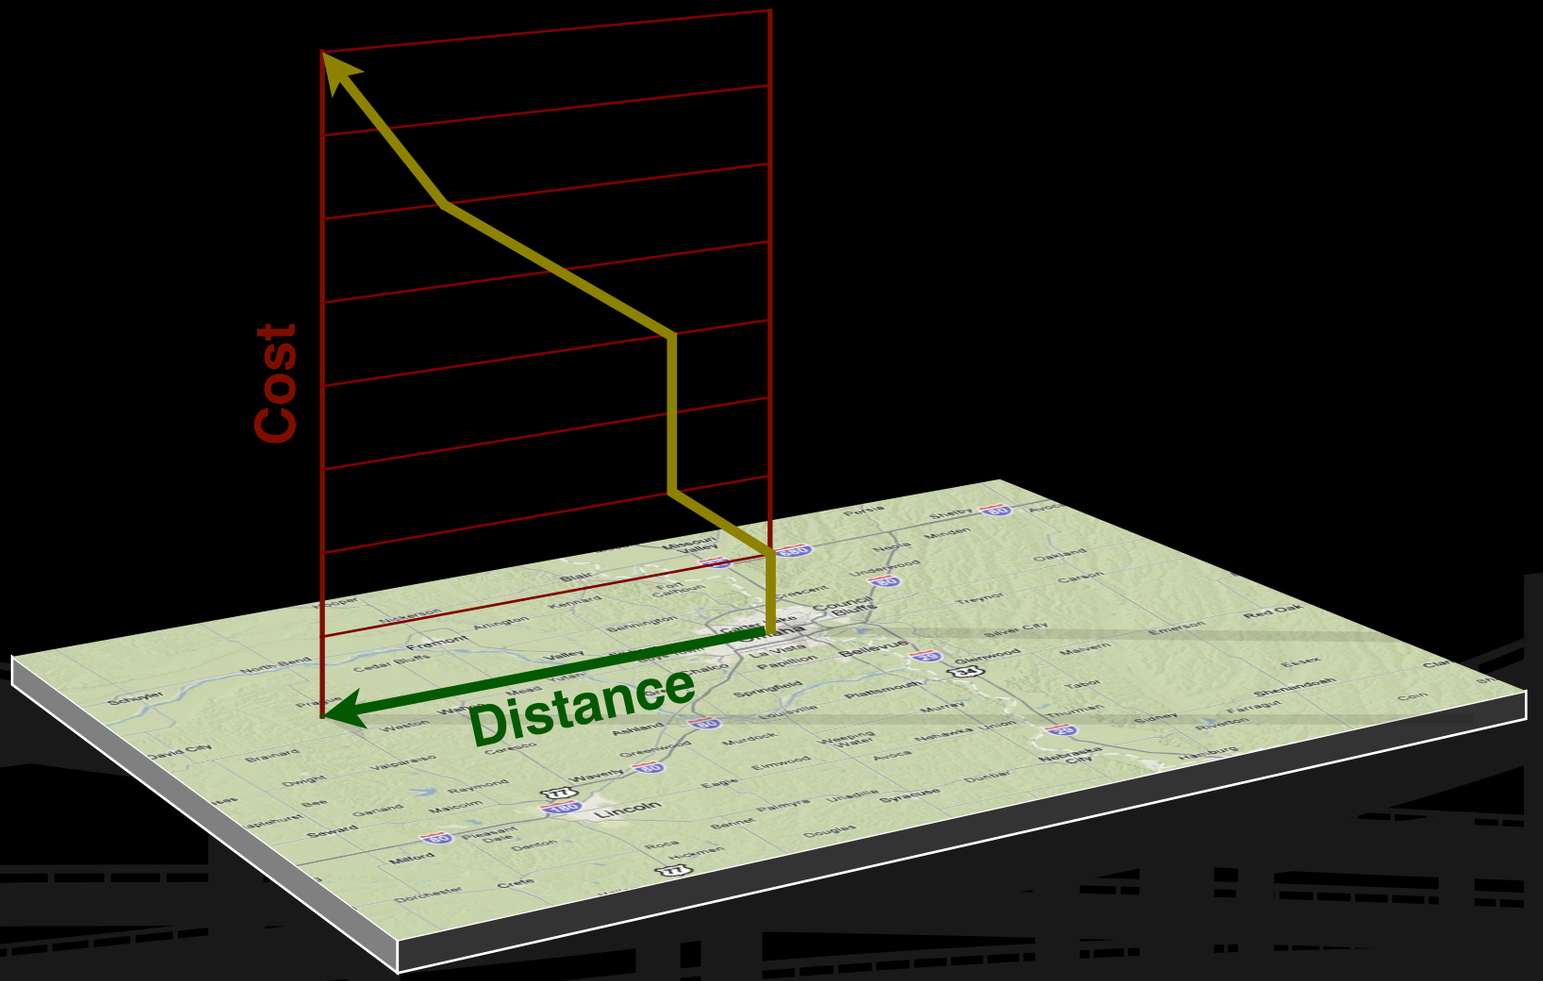
PCH
Packet Clearing House



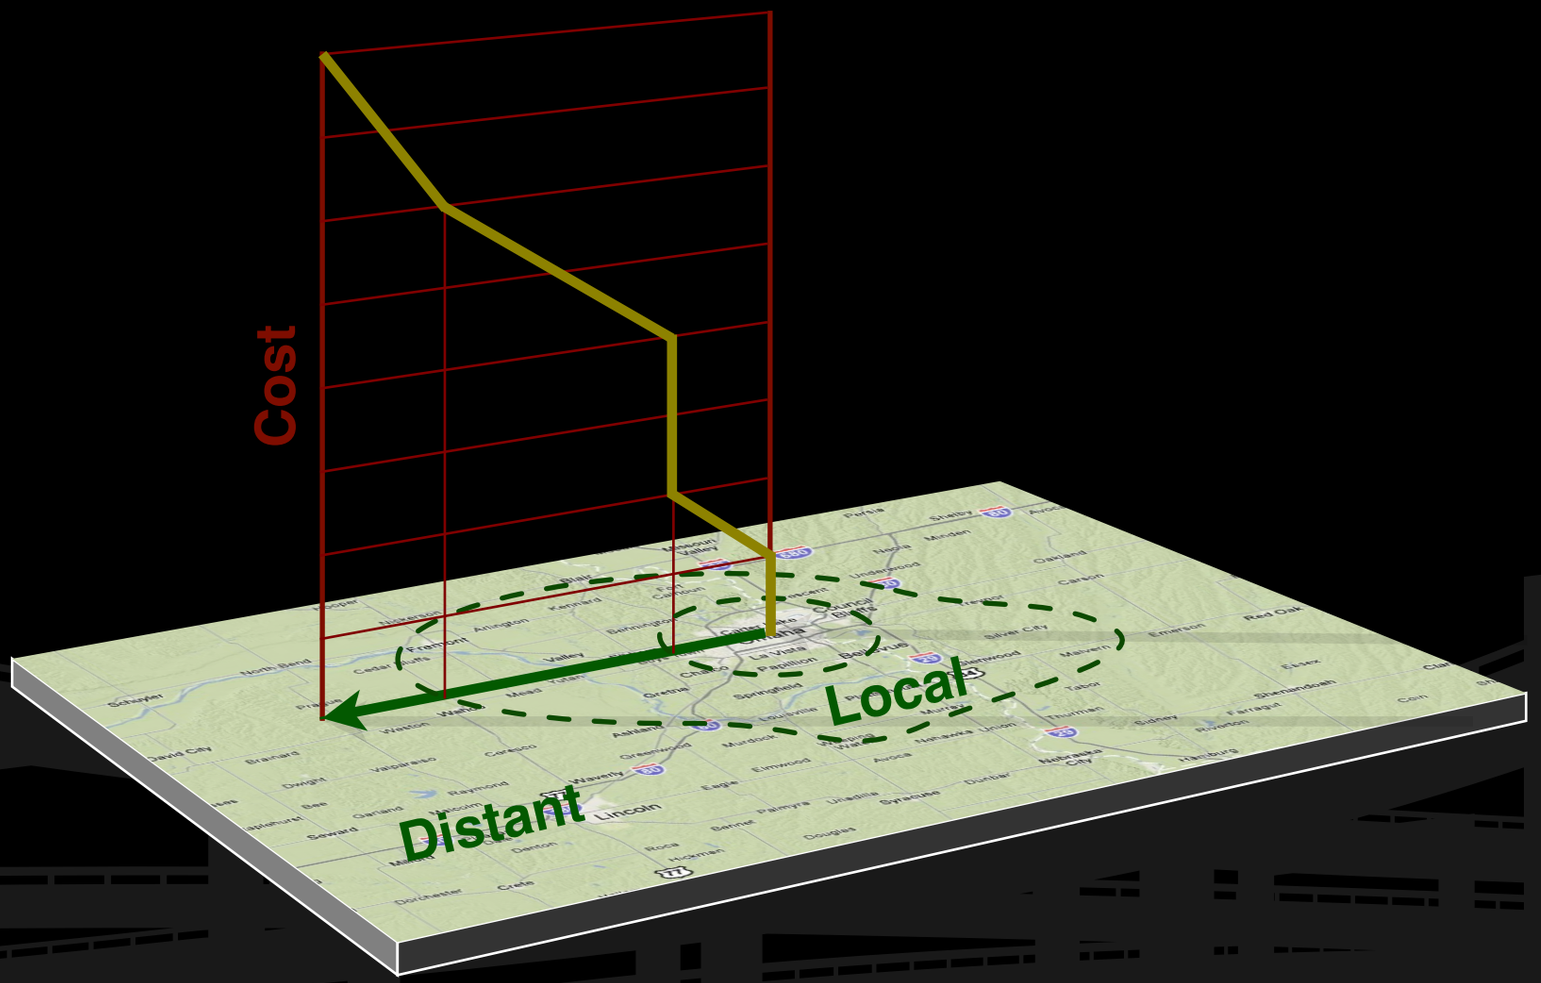
PCH
Packet Clearing House



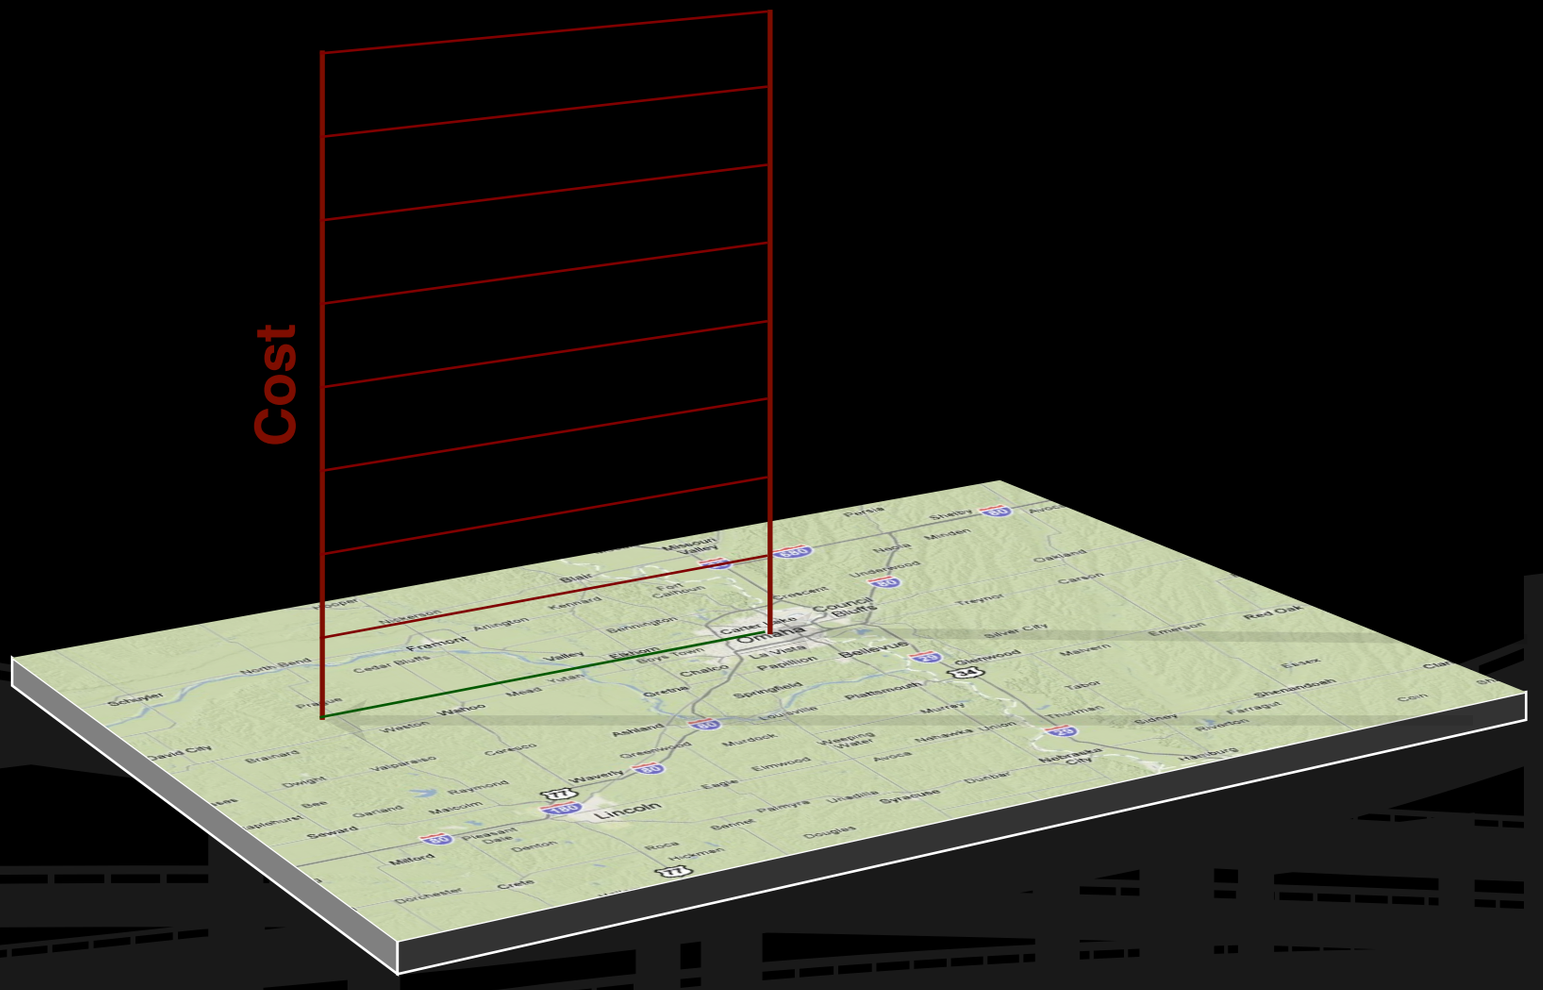
PCH
Packet Clearing House



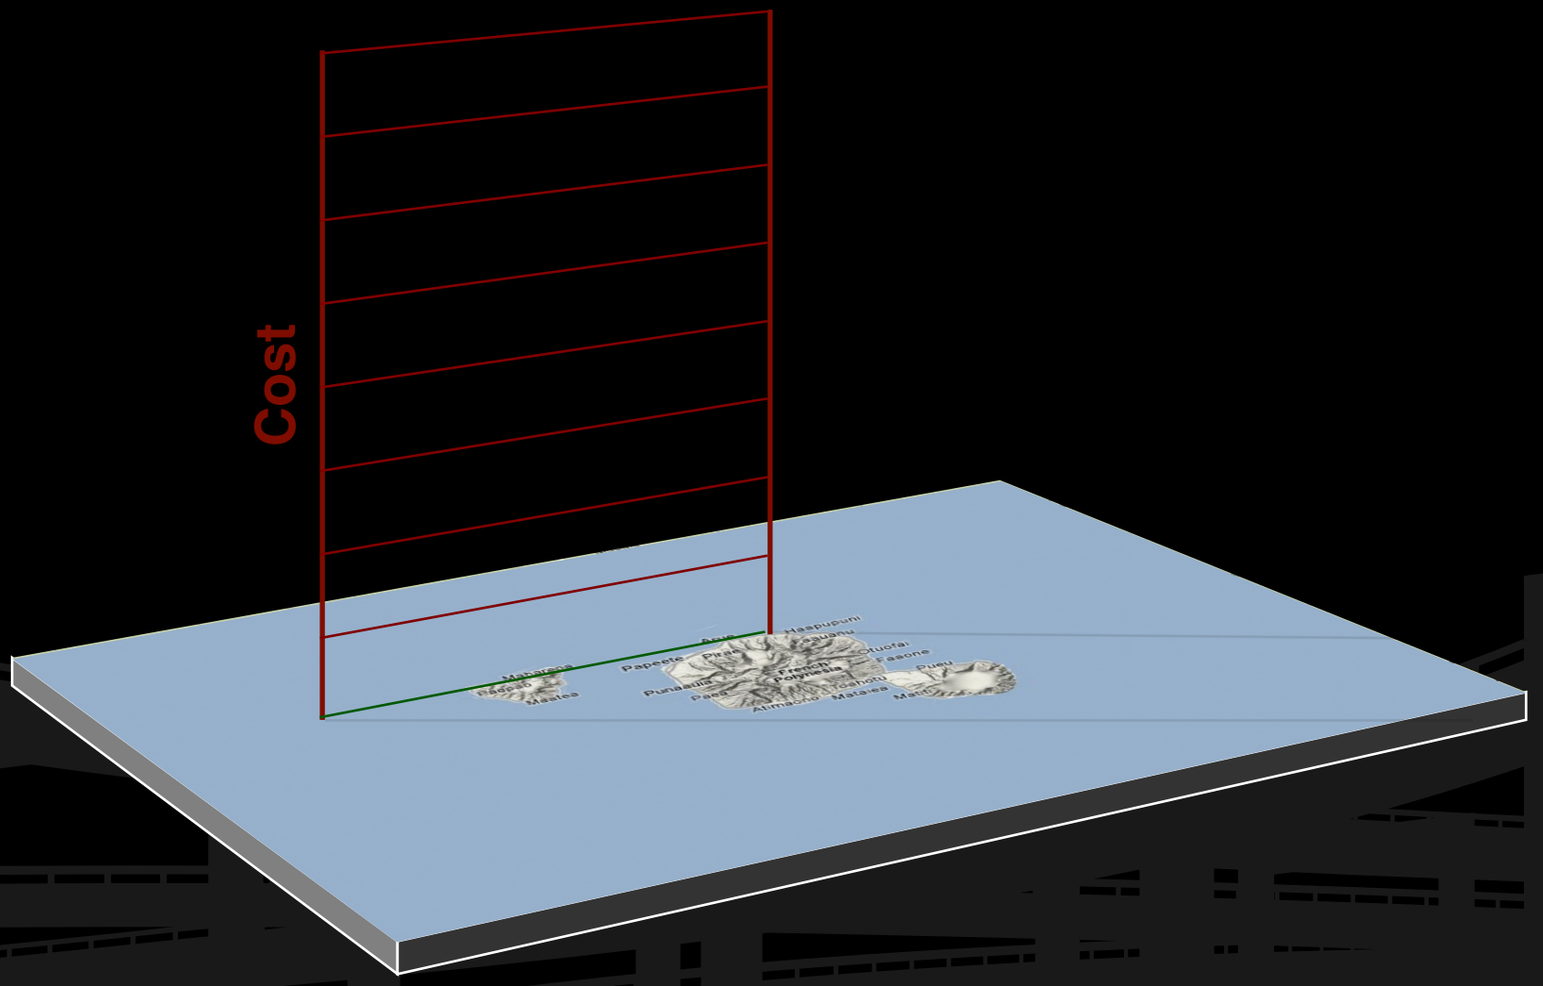
PCH
Packet Clearing House



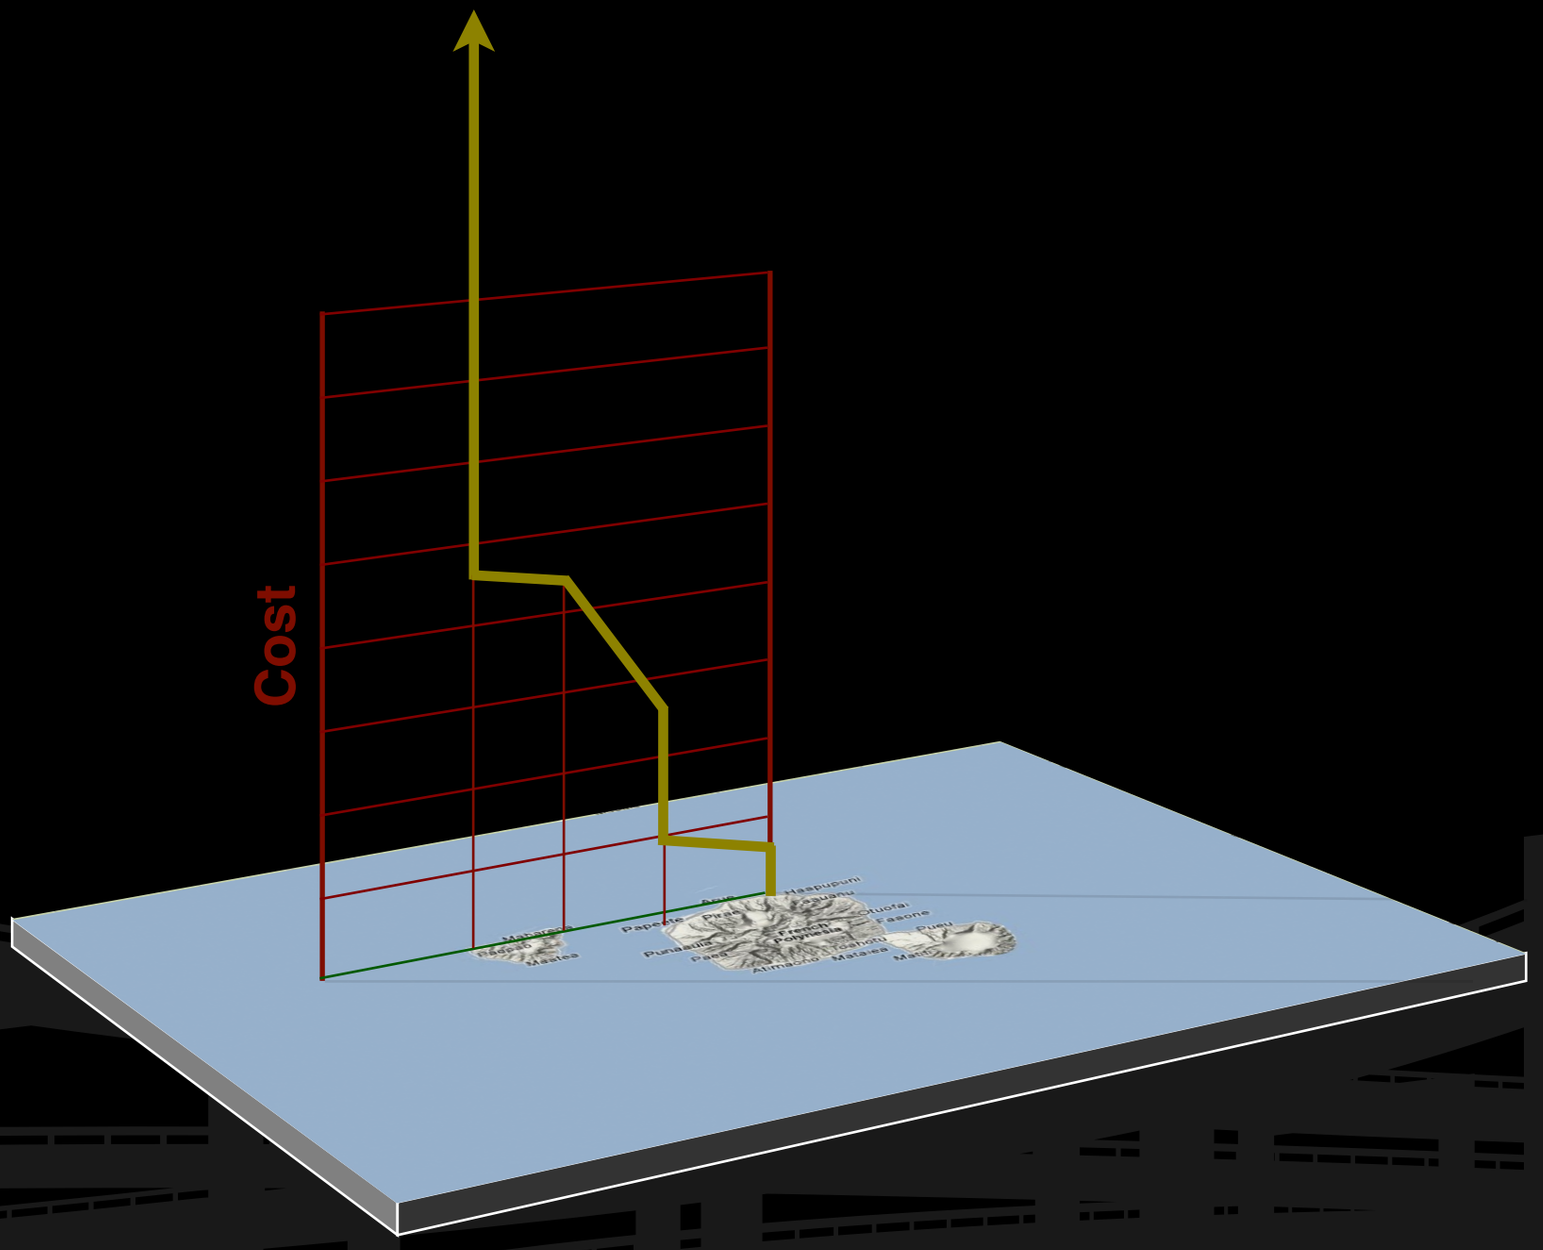
PCH
Packet Clearing House


Cost
Omaha
Lincoln
Council Bluffs
Bellevue
Papillion
La Vista
Fremont
Blair
Plattsmouth
Nebraska City
Seward
Wahoo
Ashland
Waverly
Valley
Chalco
Springfield
Gretna
Elkhorn
Carter Lake
Missouri Valley
Underwood
Oakland
Silver City
Red Oak
Malvern
Glenwood
Tabor
Shenandoah
Sidney
Hamburg
Thurman
Murray
Louisville
Union
Weeping Water
Syracuse
Dunbar
Avoca
Elmwood
Eagle
Palmyra
Douglas
Murdock
Greenwood
Cedar Bluffs
Scribner
North Bend
Arlington
Kennard
Bennington
Irvington
Crescent
Treynor
Carson
Emerson
Hooper
Dodge
Colfax
Platte
Saunders
Cass
Otoe
Johnson
Nemaha
Pawnee
Richardson
Gage
Lancaster
Sarpy
Douglas
Washington
Burt
Thurston
Monona
Harrison
Shelby
Audubon
Pottawattamie
Mills
Fremont
Page
Taylor
Adams
Clay
Fillmore
Saline
Jefferson
Thayer
Nuckolls
Webster
Franklin
Harlan
Furnas
Gosper
Phelps
Kearney
Hall
Hamilton
York
Polk
Dawson
Buffalo
Sherman
Howard
Greeley
Valley
Garfield
Wheeler
Boone
Nance
Merrick
Hall
Hamilton
Clay
Adams
Kearney
Phelps
Gosper
Furnas
Harlan
Franklin
Webster
Nuckolls
Thayer
Jefferson
Saline
Fillmore

PCH
Packet Clearing House



PCH
Packet Clearing House



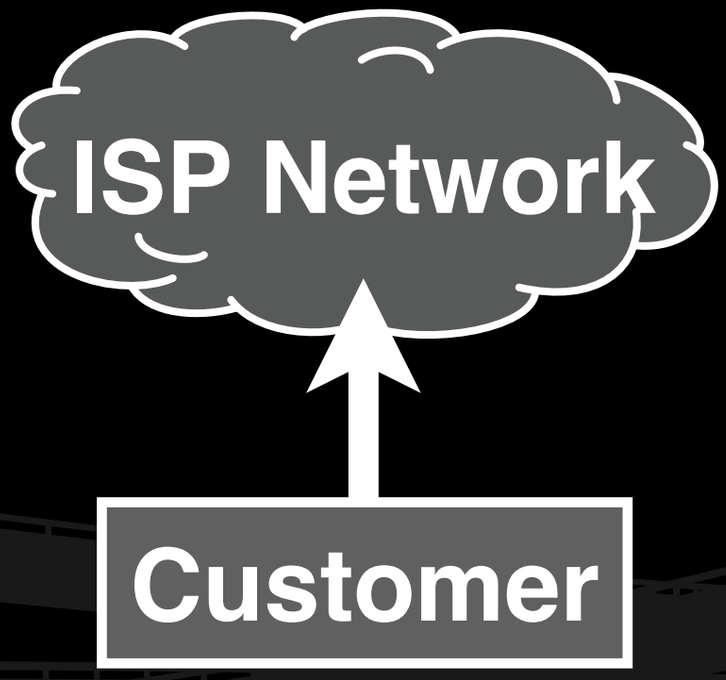
Flow Data for Routing Optimization

The problem:

Internet service providers have complex blends of customer types and customer traffic. They need to continuously reduce their APBDC to remain competitive, but generally try do so by guesswork, rather than by using the data available to them.

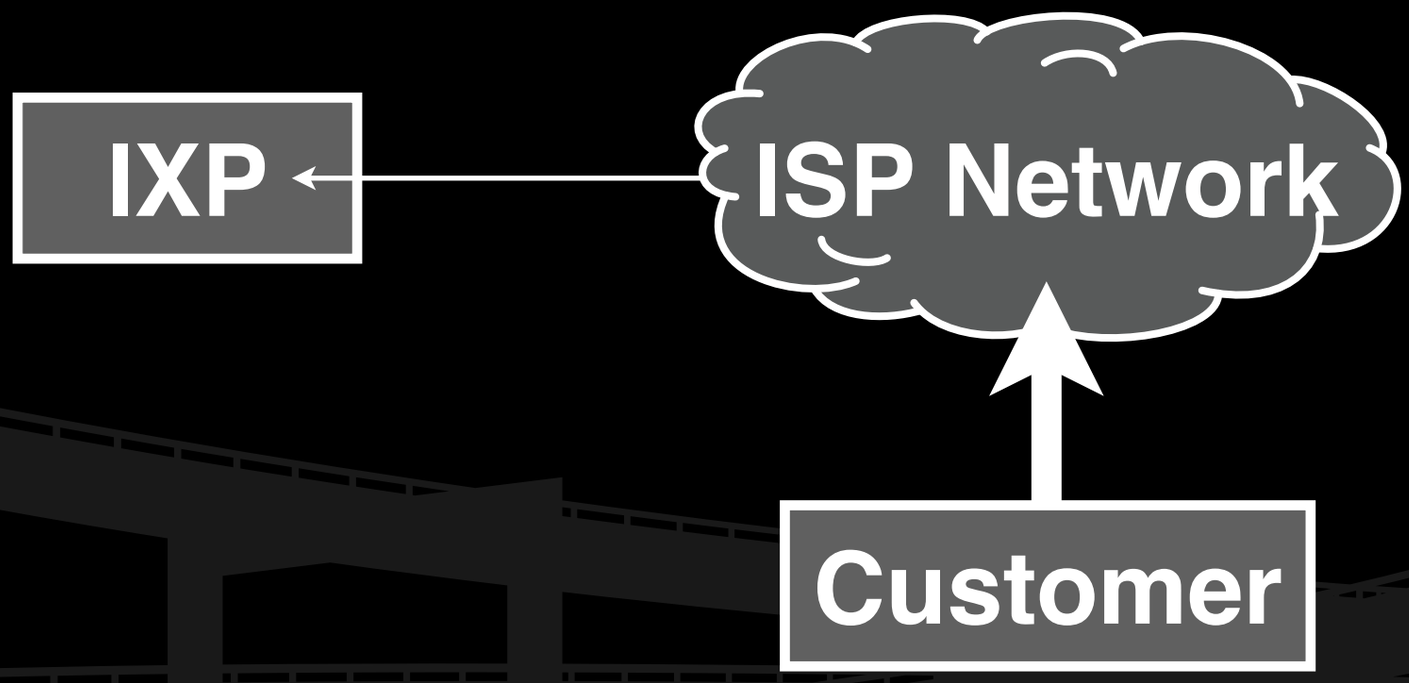
PCH
Packet Clearing House

Flow Data for Routing Optimization



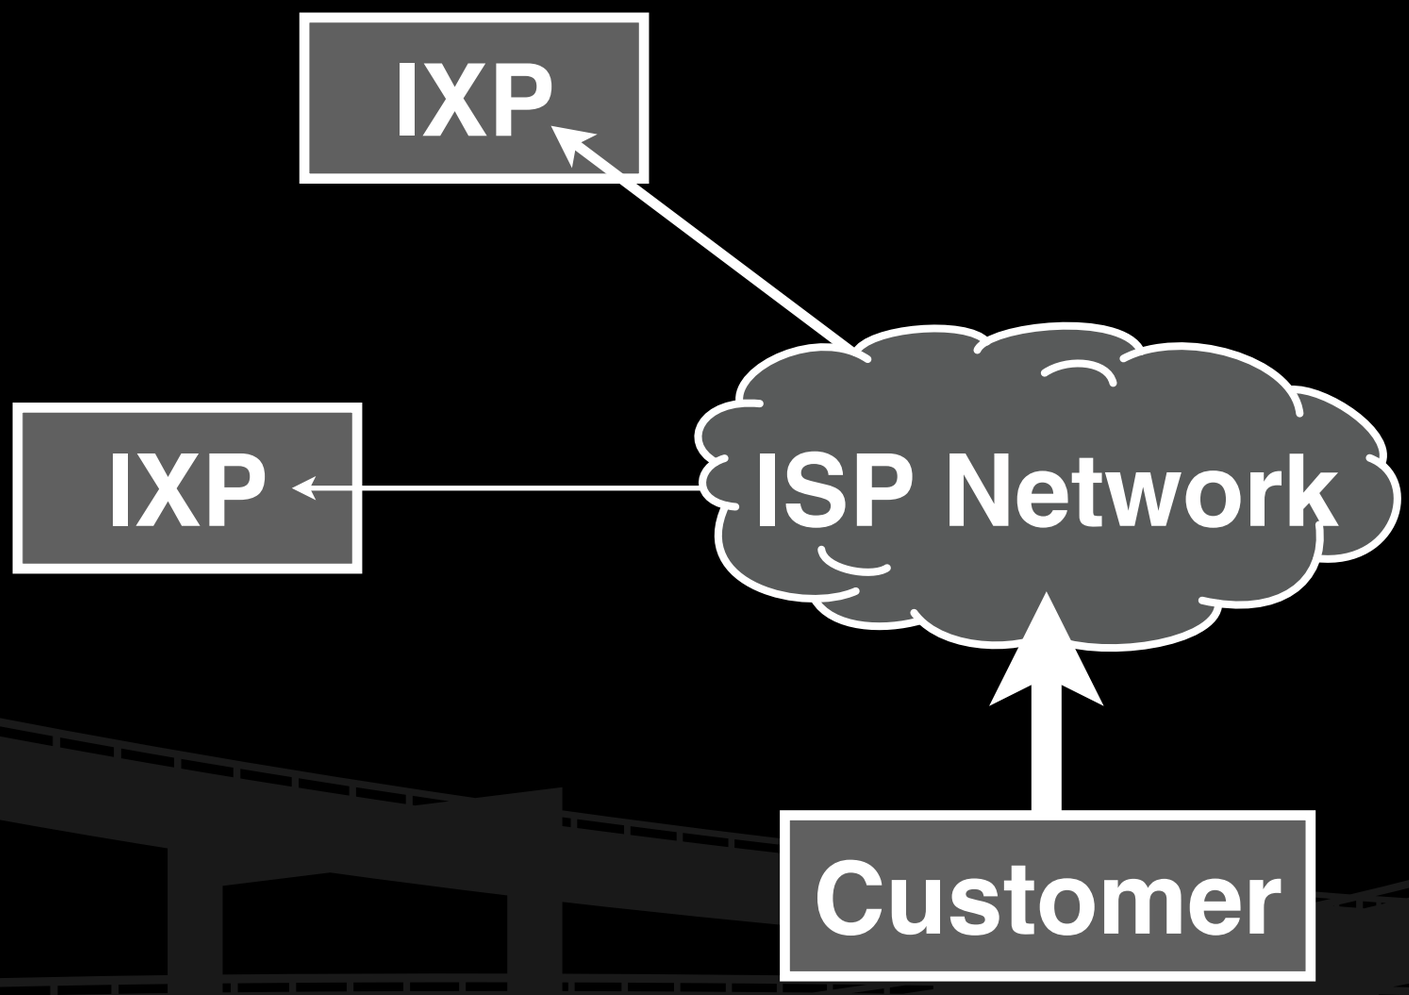
Packet Clearing House

Flow Data for Routing Optimization



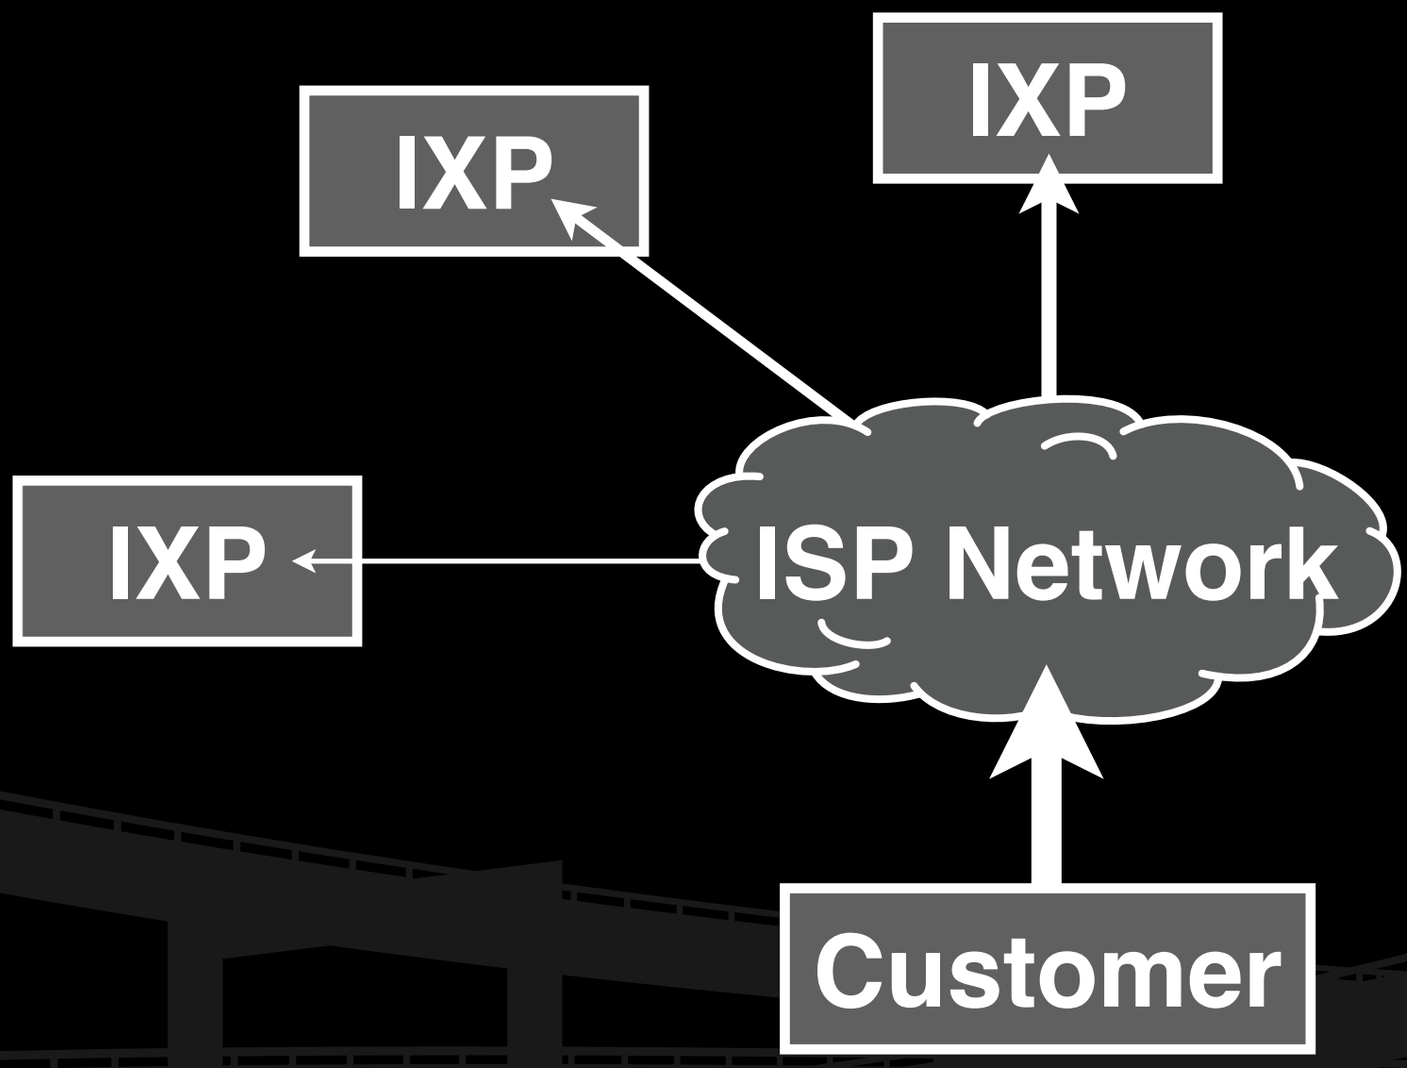
Packet Clearing House

Flow Data for Routing Optimization



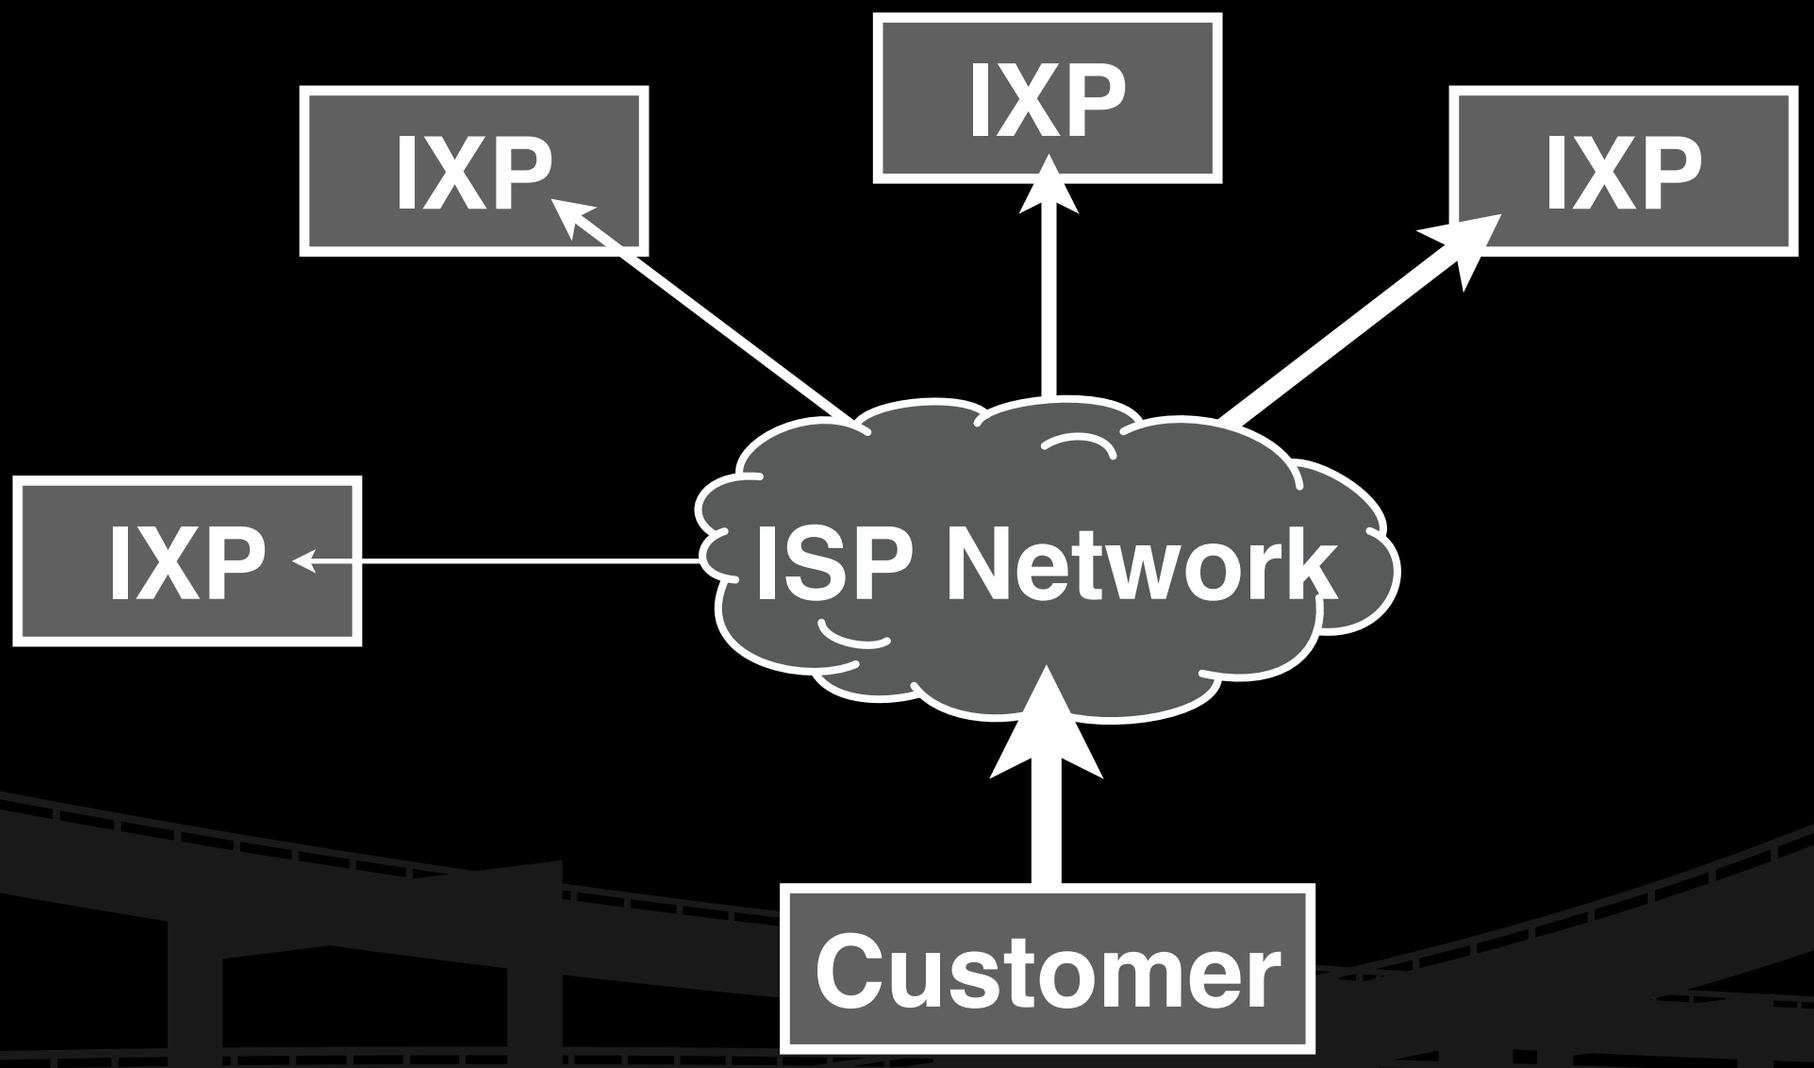
Packet Clearing House

Flow Data for Routing Optimization



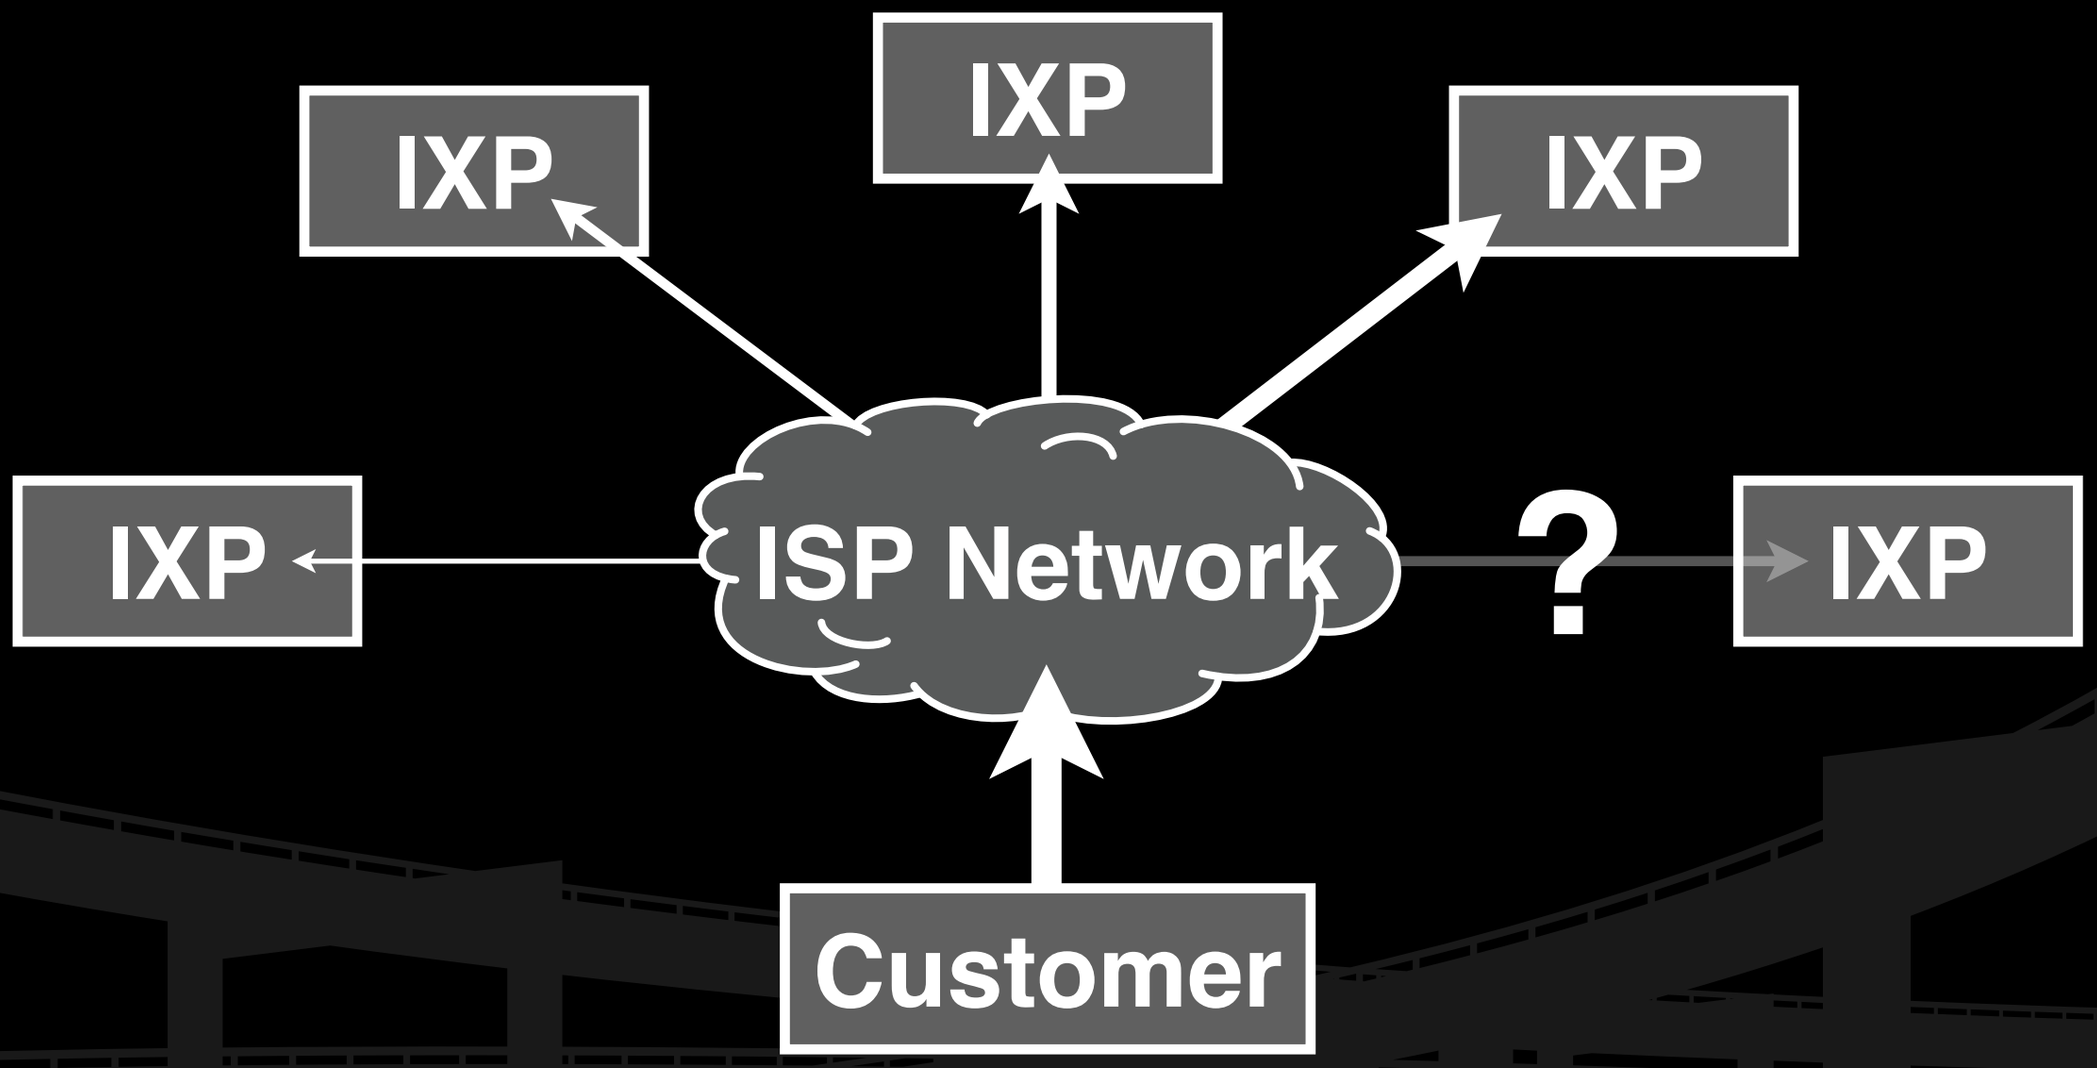
Packet Clearing House

Flow Data for Routing Optimization



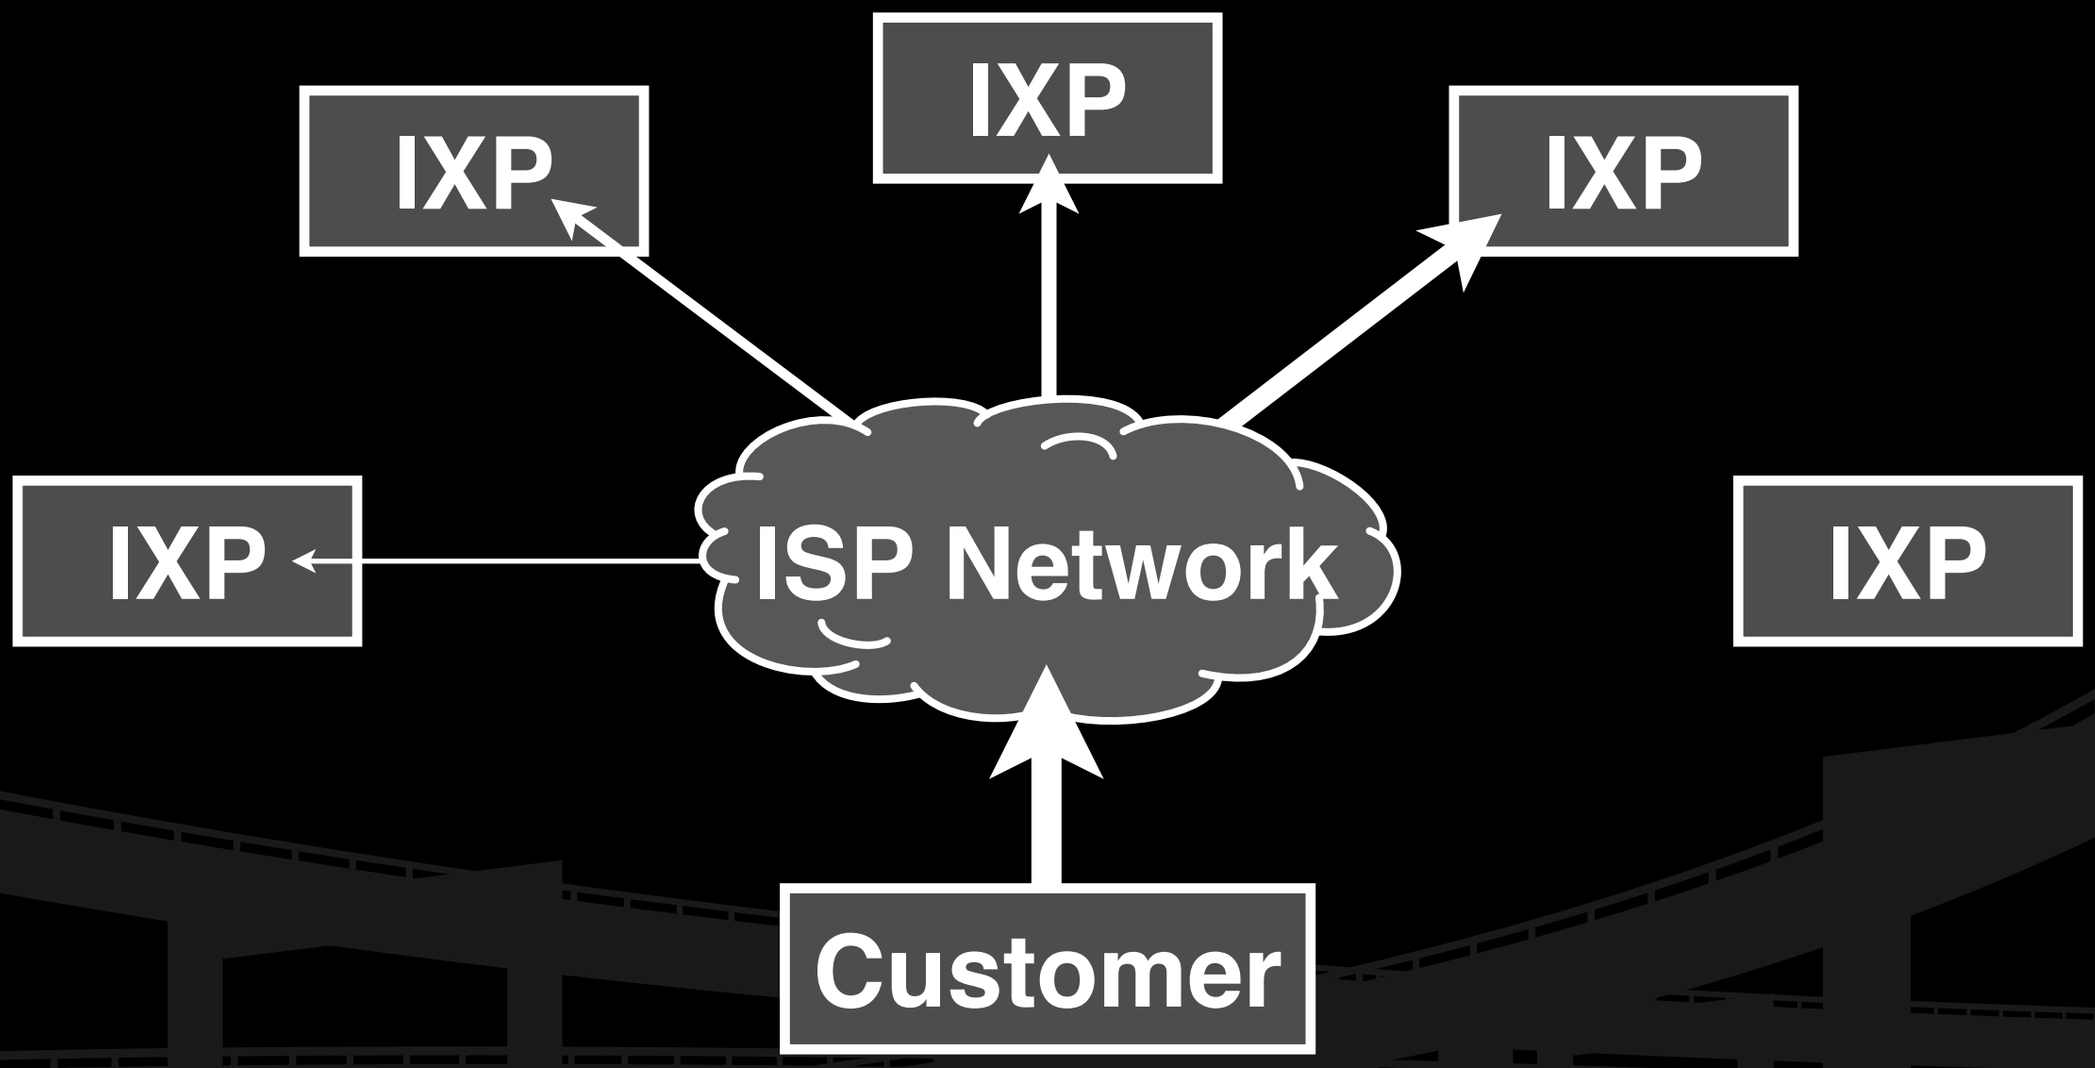
Packet Clearing House

Flow Data for Routing Optimization

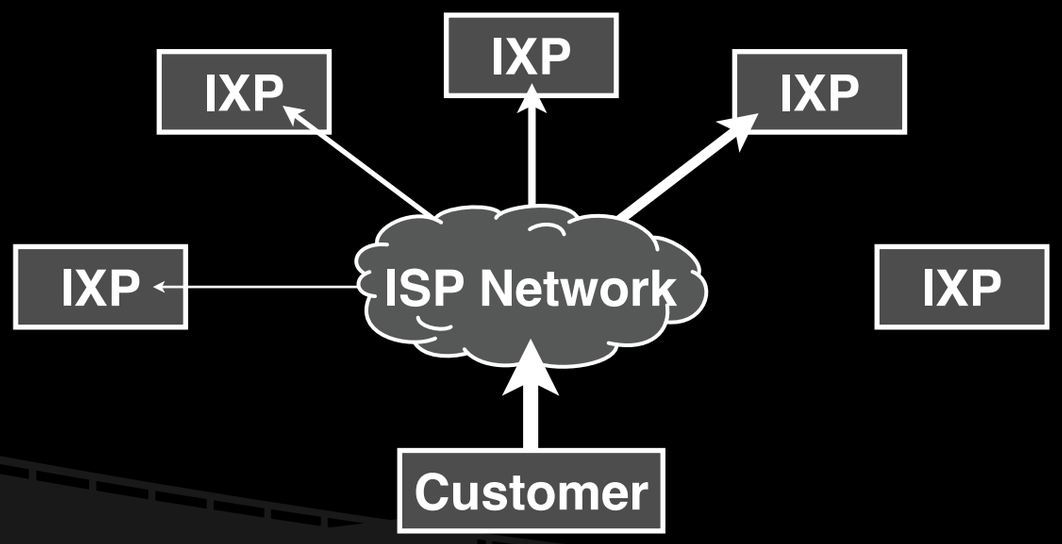


Packet Clearing House

Flow Data for Routing Optimization

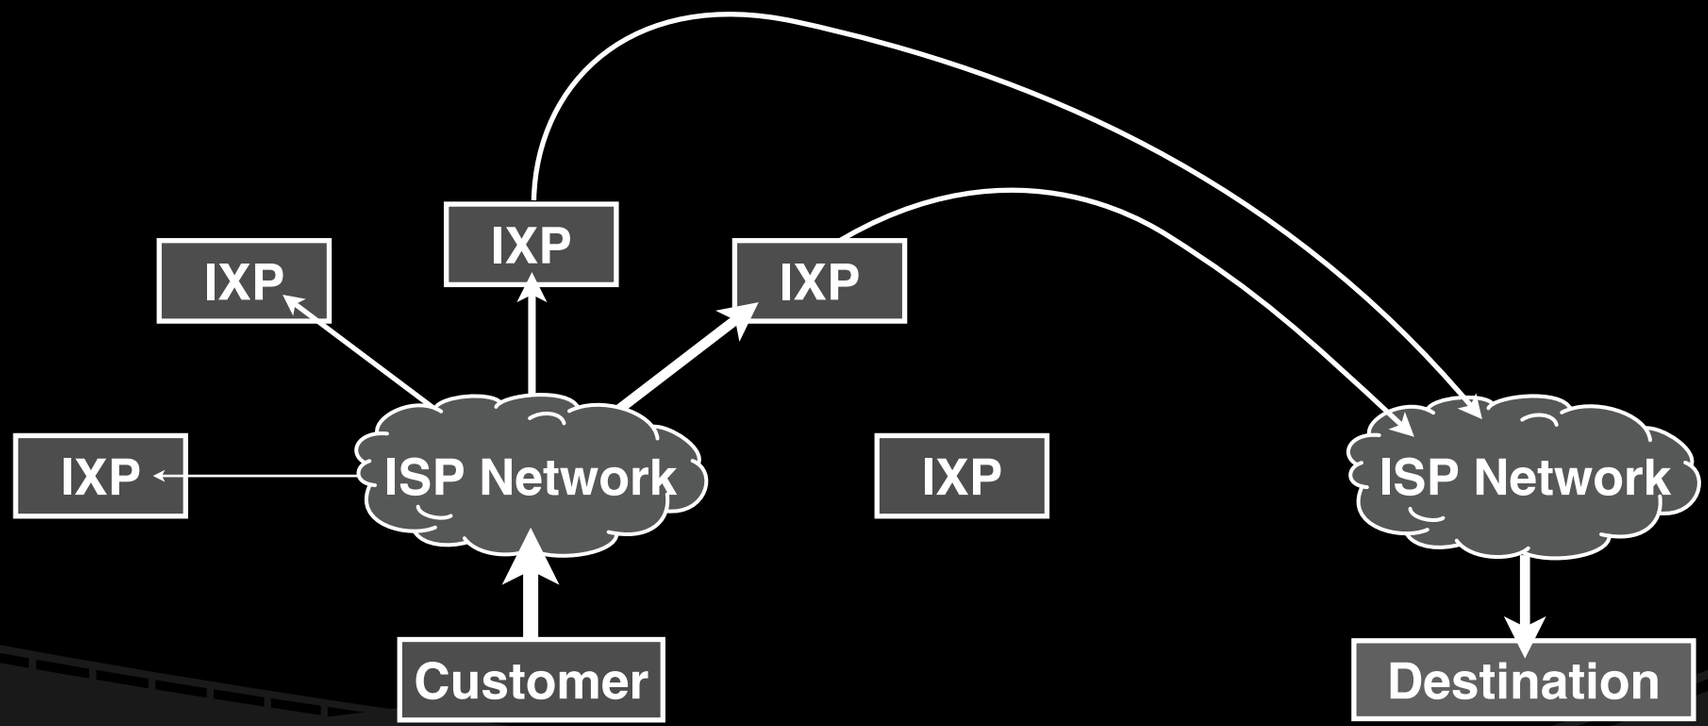


Flow Data for Routing Optimization



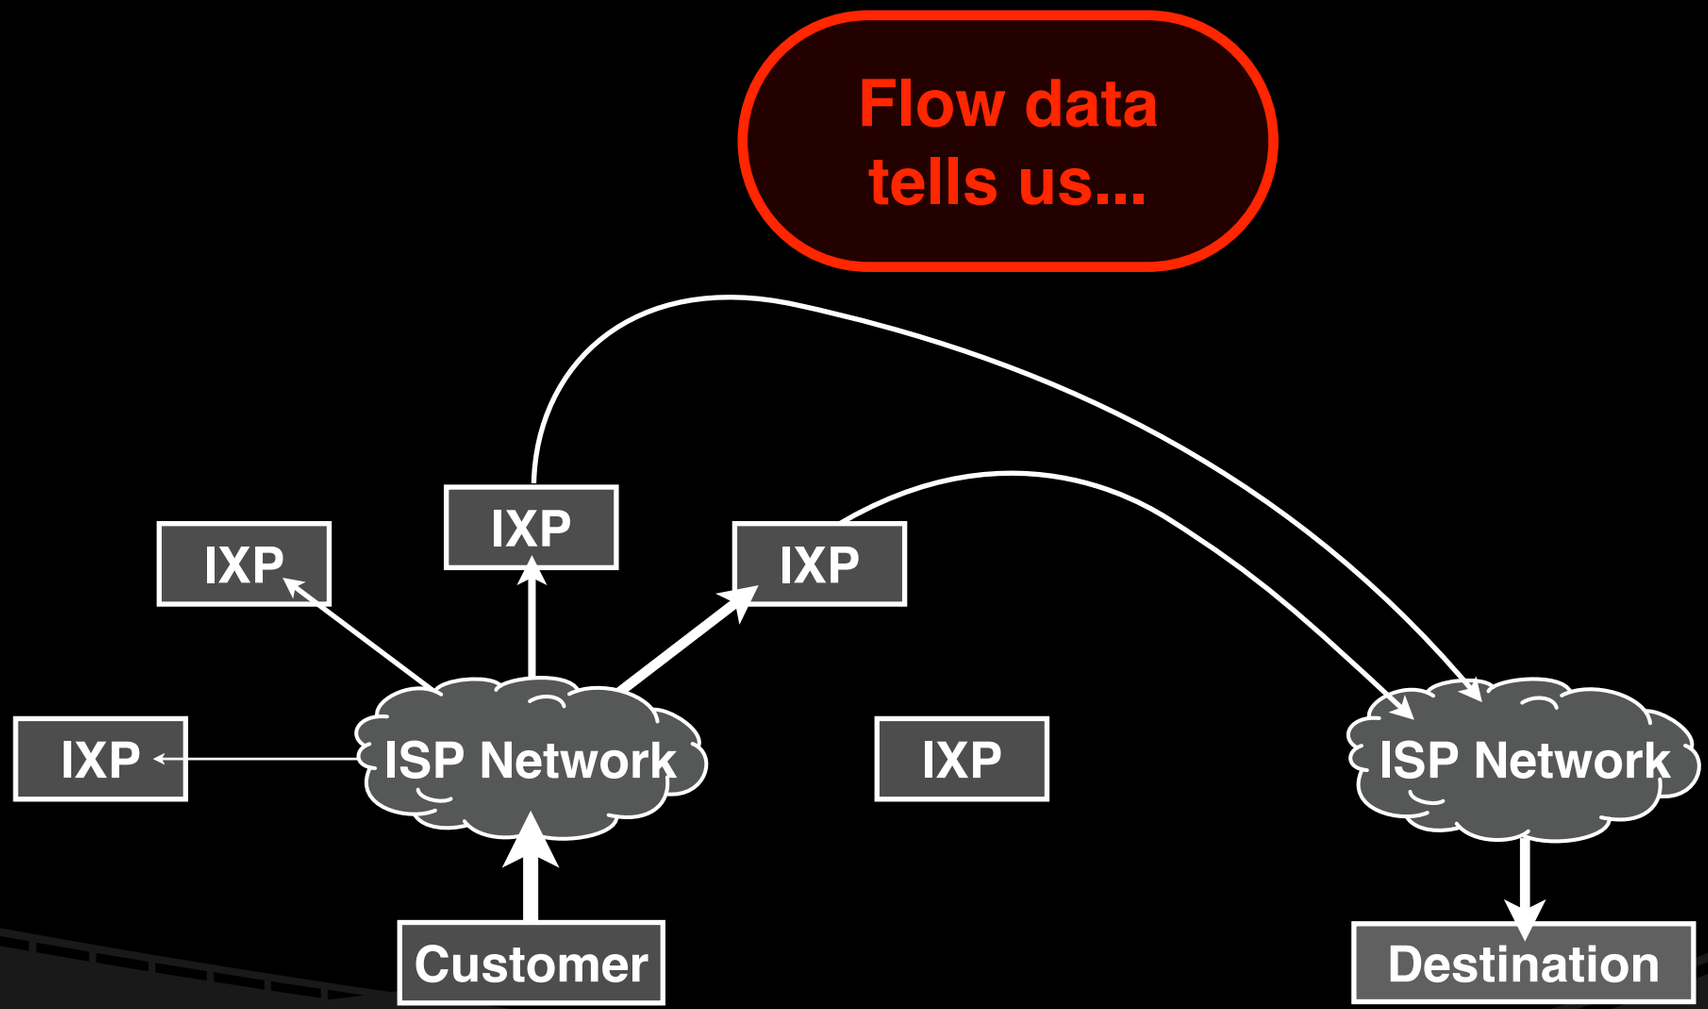
Packet Clearing House

Flow Data for Routing Optimization



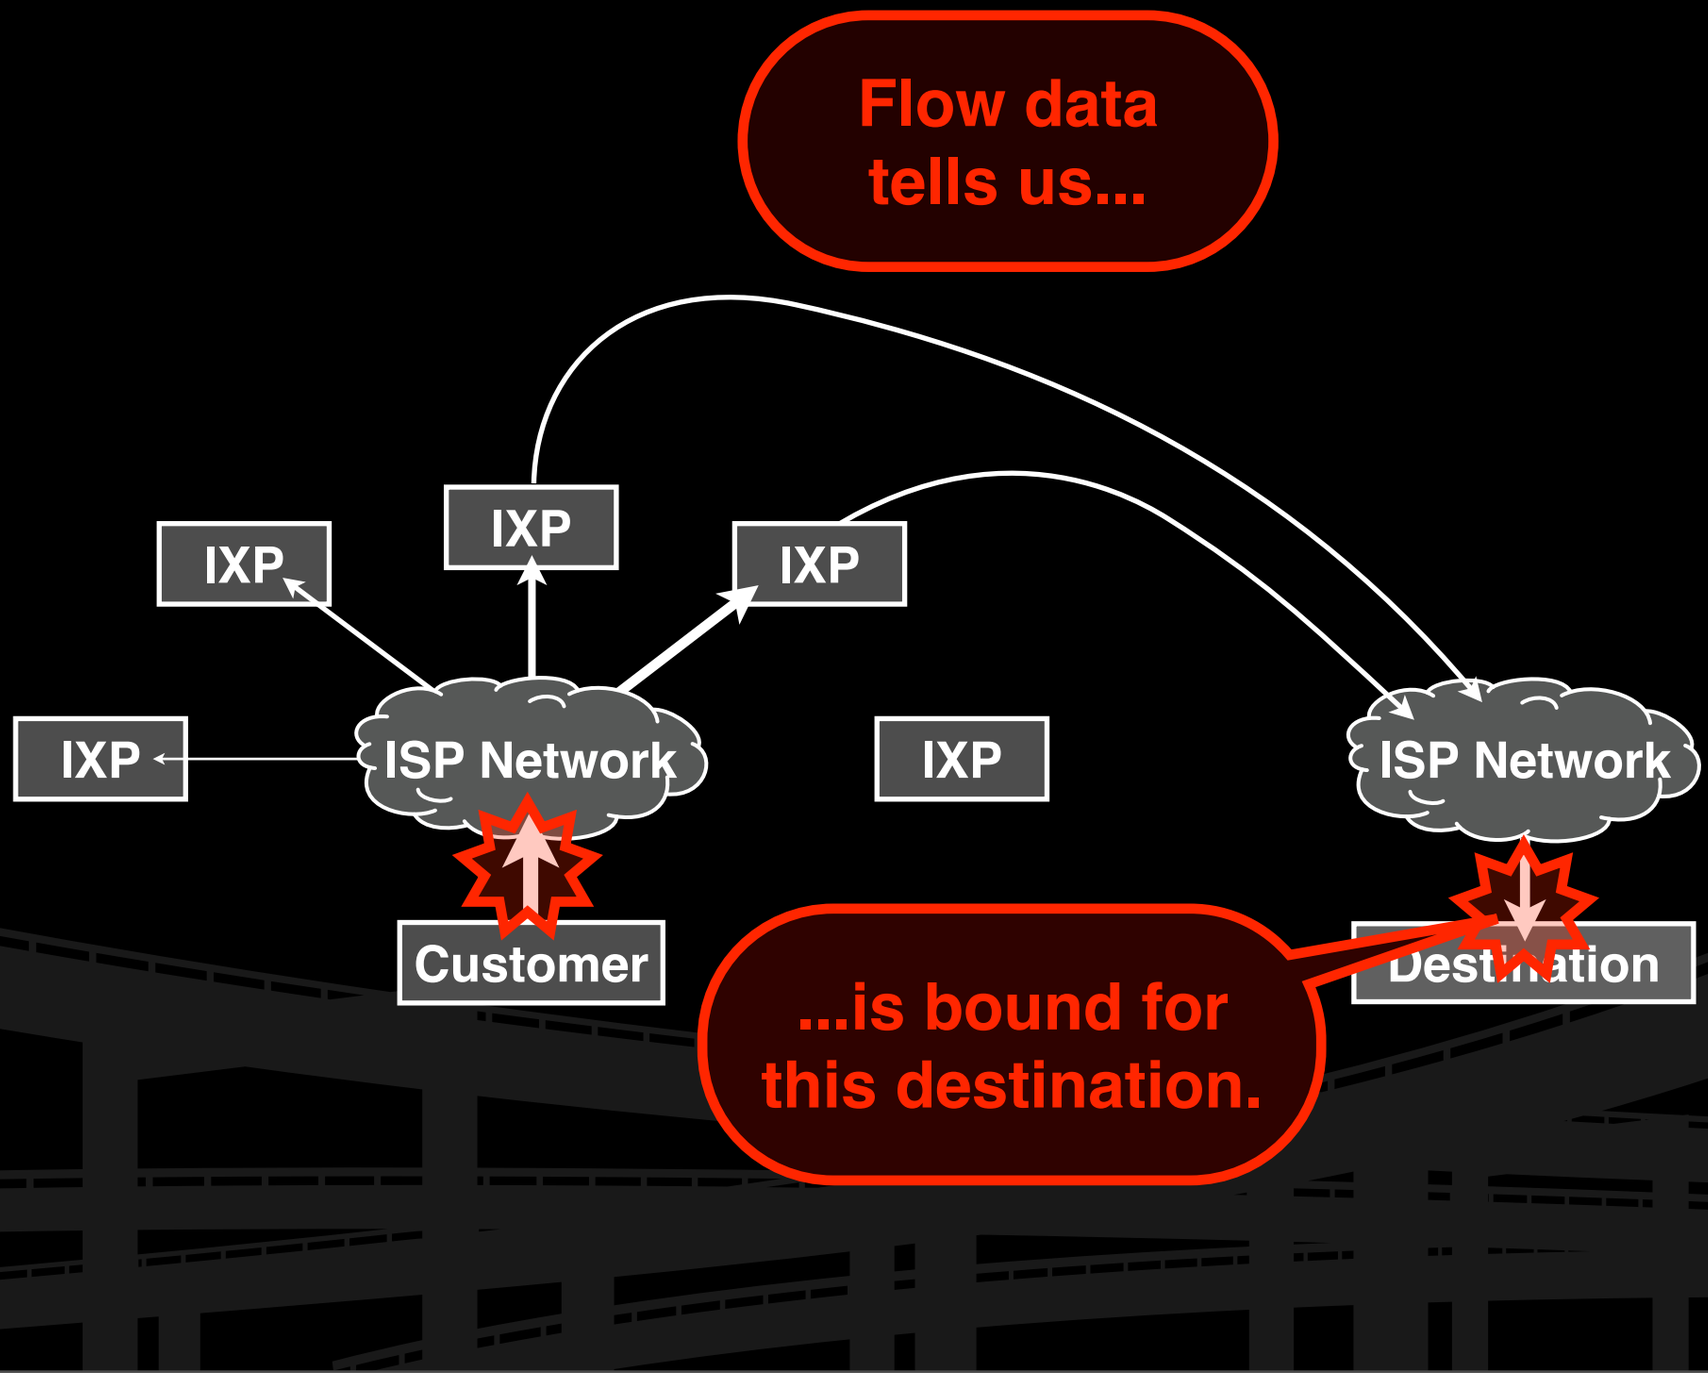
Packet Clearing House

Flow Data for Routing Optimization



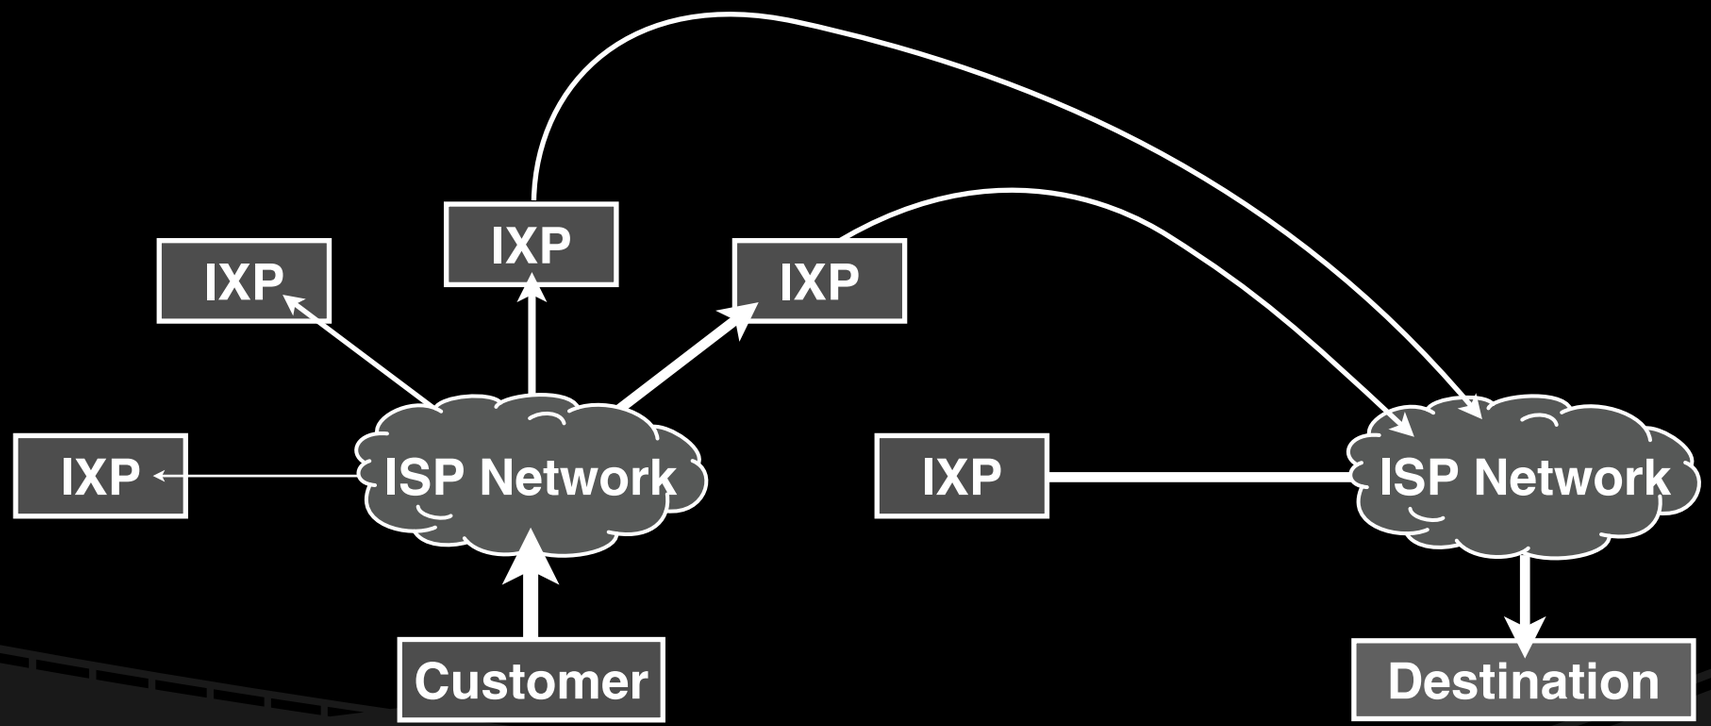
PCH
Packet Clearing House

Flow Data for Routing Optimization



Packet Clearing House

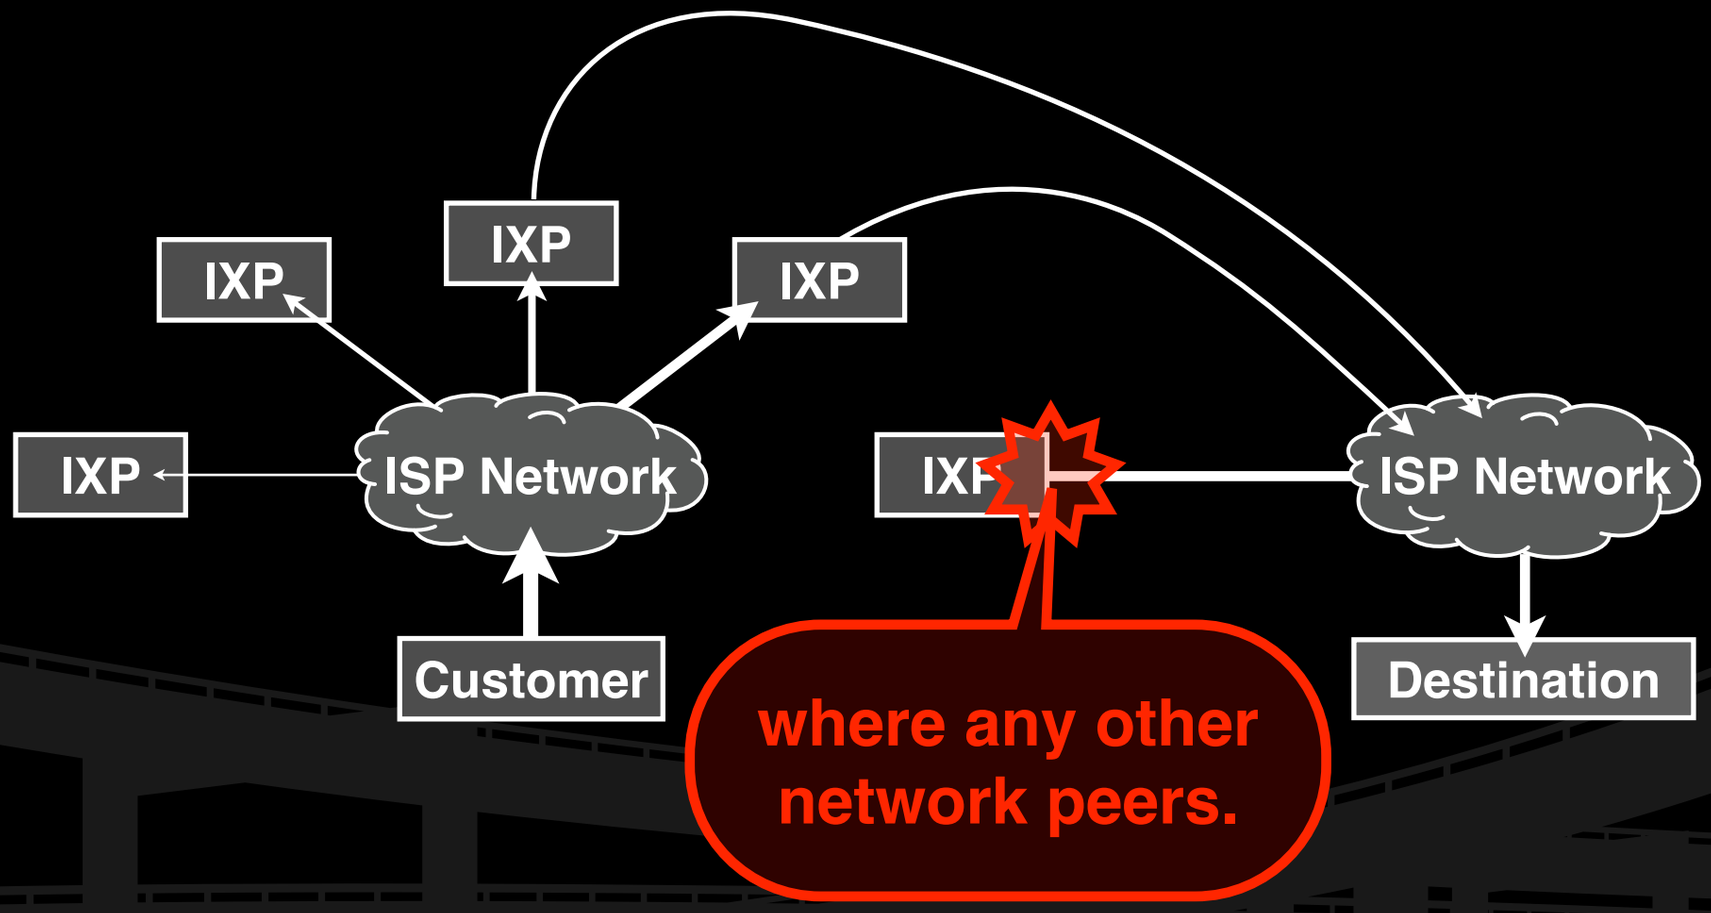
Flow Data for Routing Optimization



Packet Clearing House

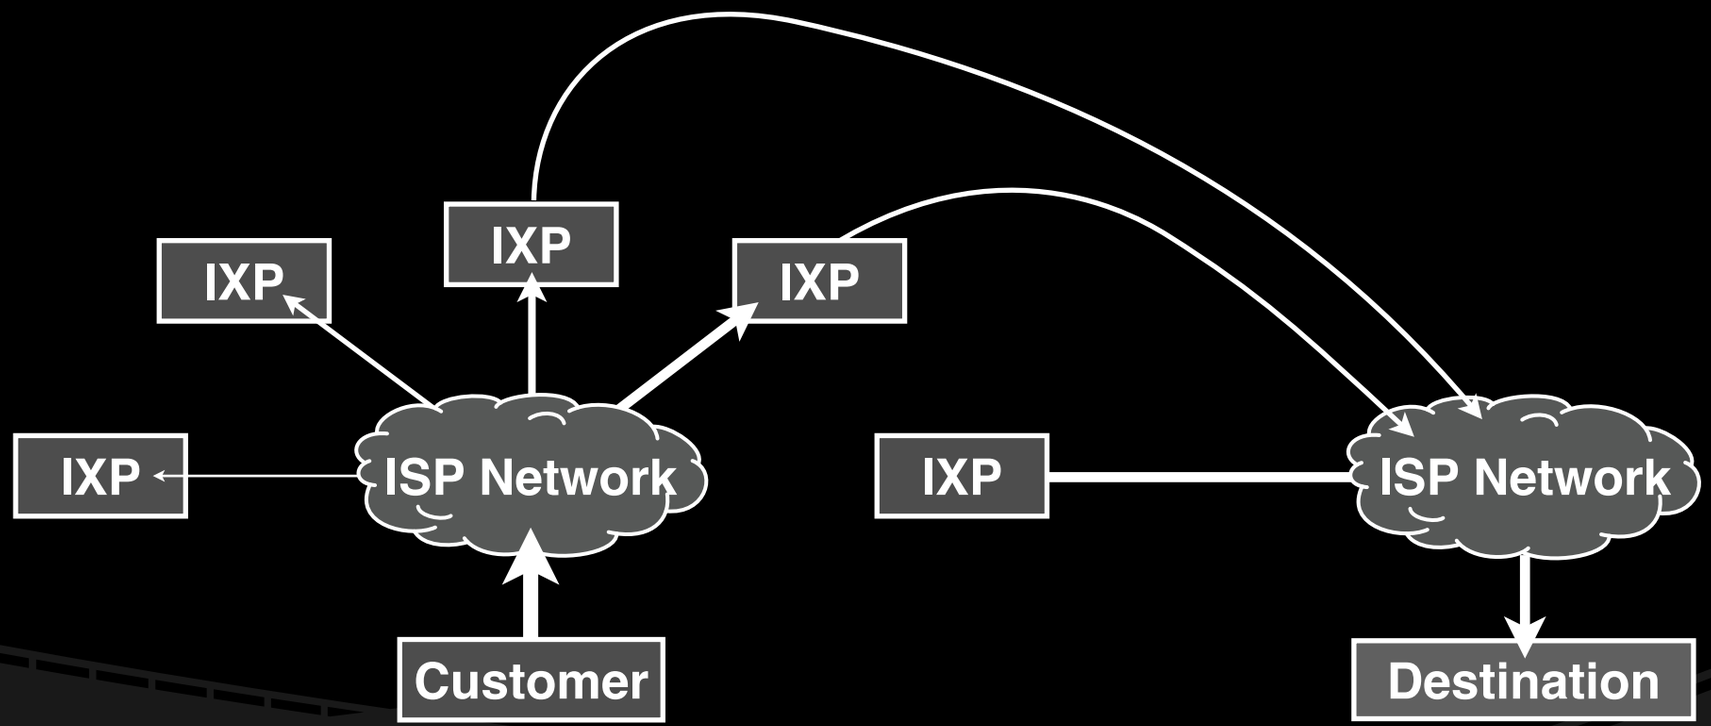
Flow Data for Routing Optimization

IXP participation data tells us...



Packet Clearing House

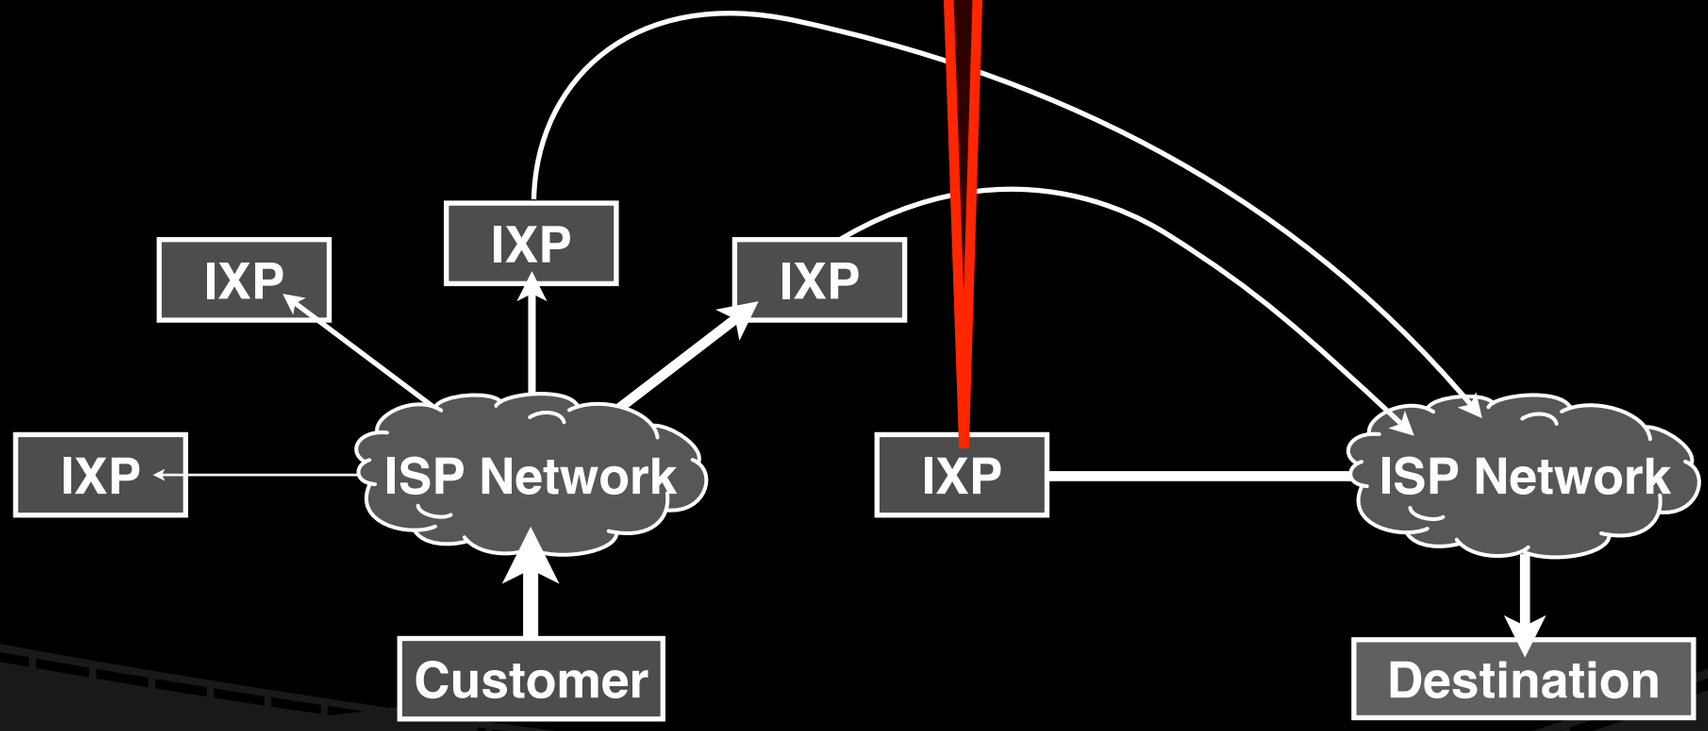
Flow Data for Routing Optimization



Packet Clearing House

Flow Data for Routing Optimization

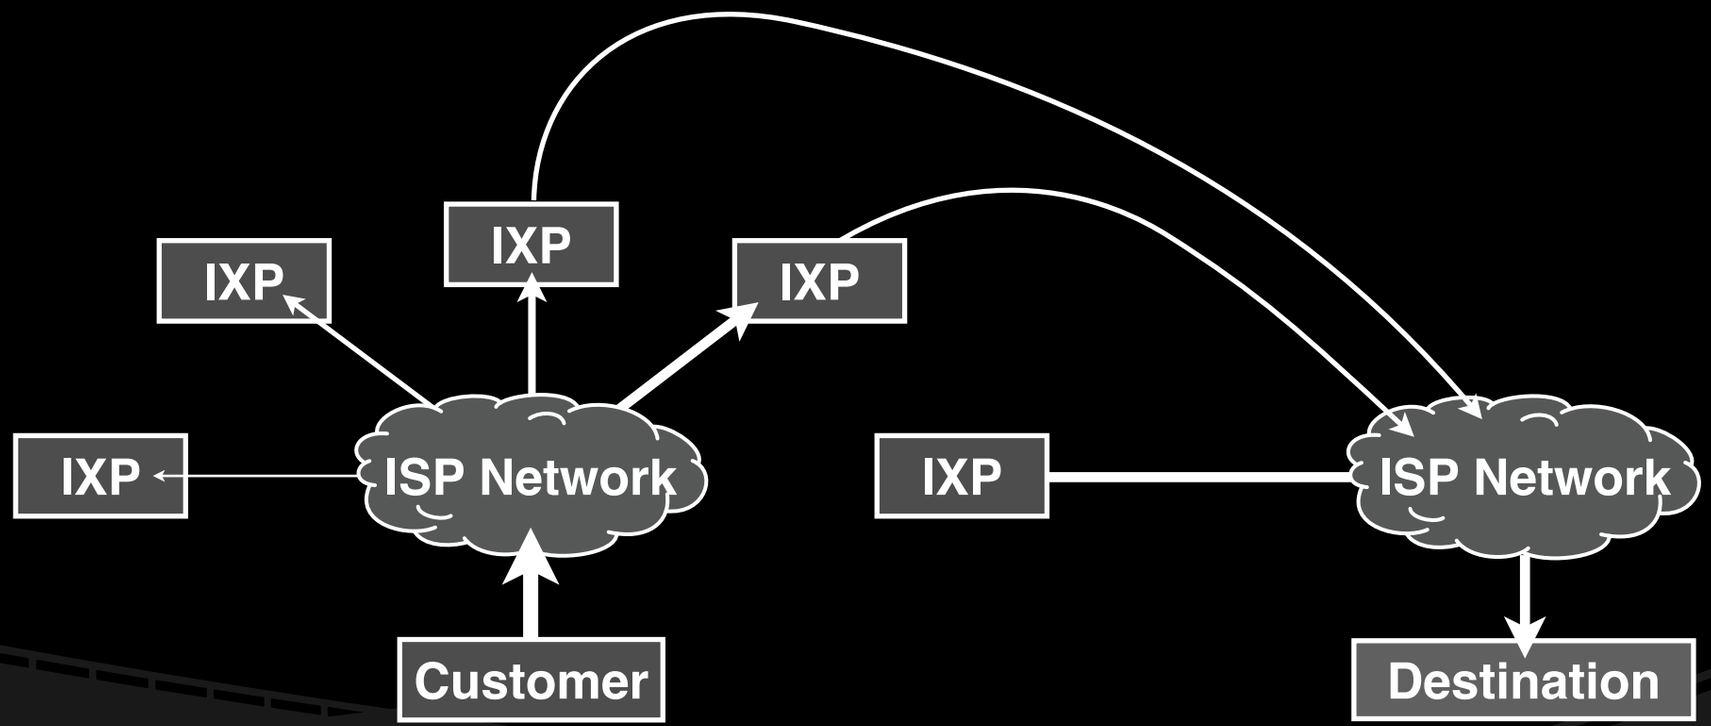
By combining the flow data for ALL participants at an IXP...



...with the routing table changes we'd see if we peered there...

Packet Clearing House

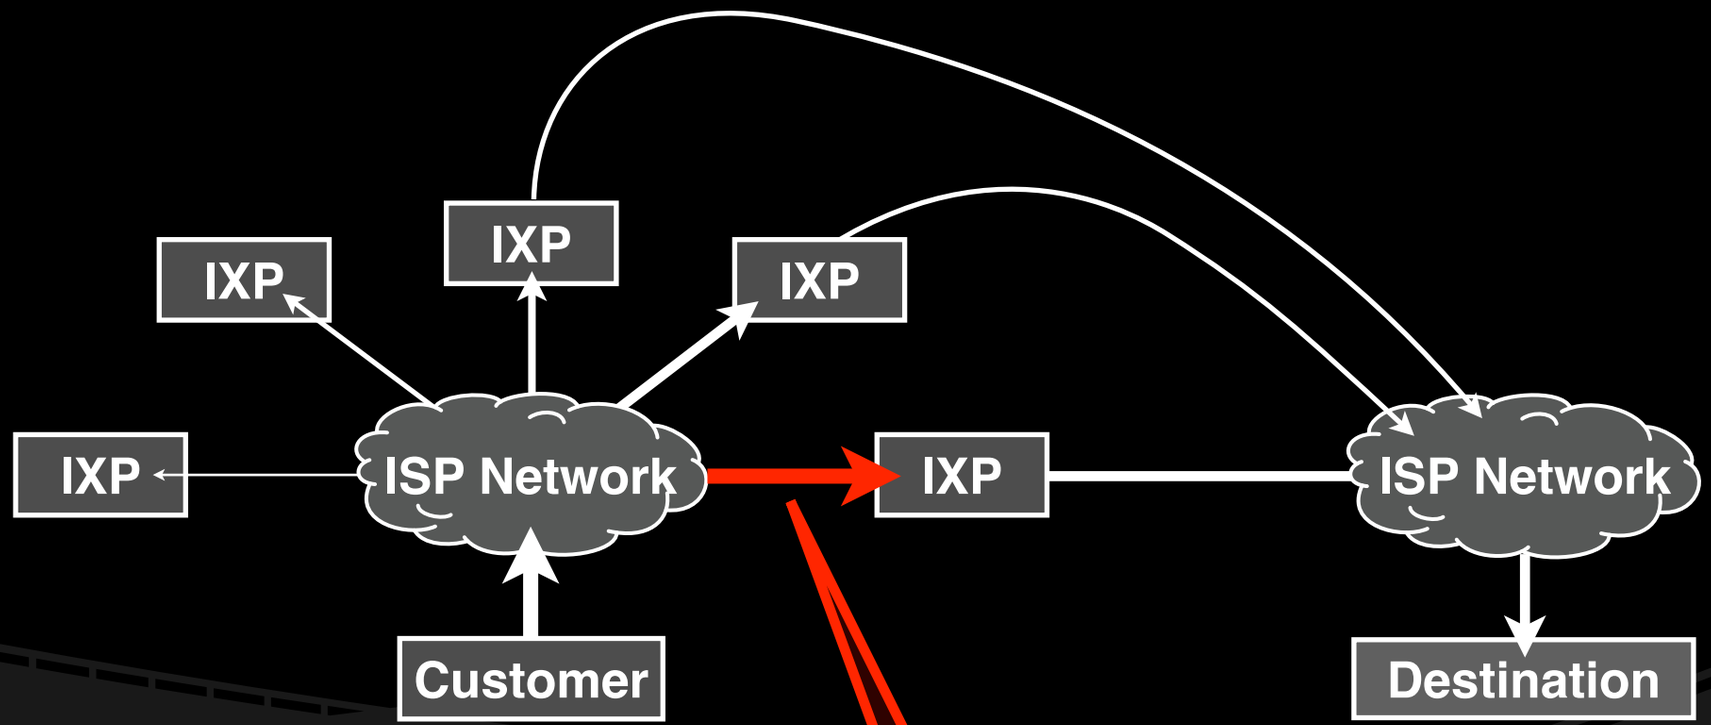
Flow Data for Routing Optimization



Packet Clearing House

Flow Data for Routing Optimization

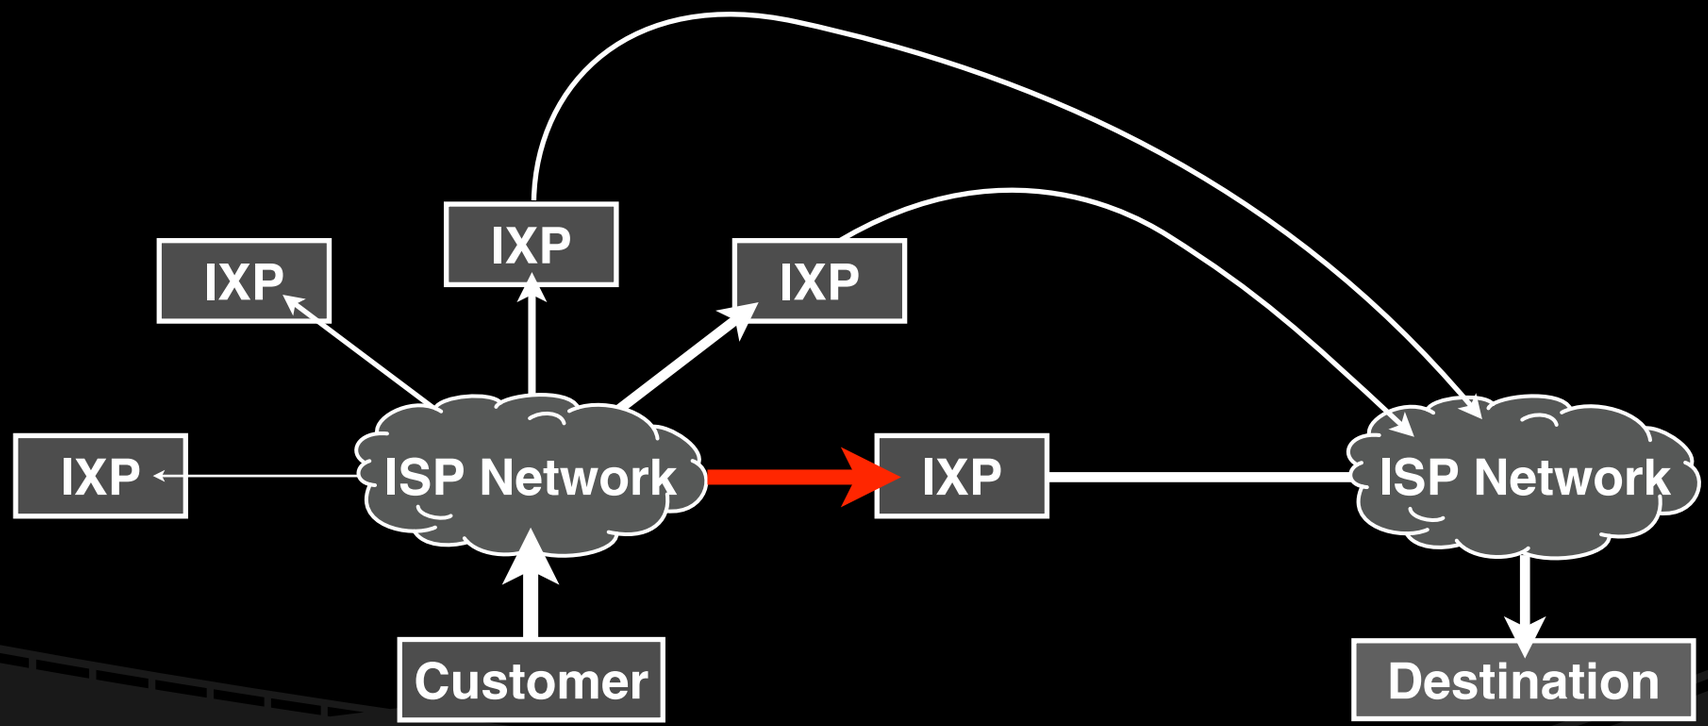
We can tell what portion of our traffic...



...would flow across this link, if it existed.

Packet Clearing House

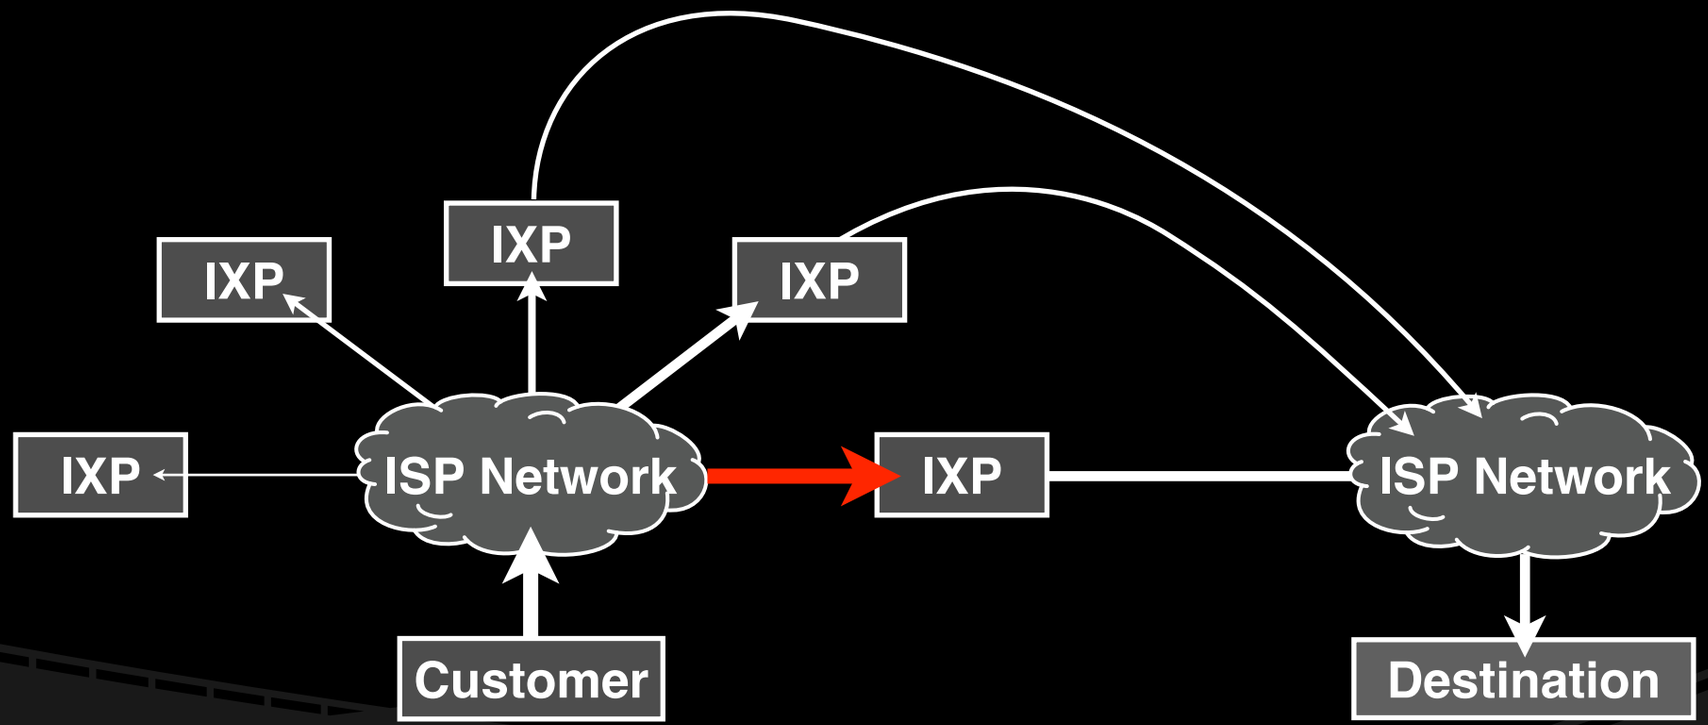
Flow Data for Routing Optimization



Packet Clearing House

Flow Data for Routing Optimization

Combining that with the projected cost of reaching the IXP...



...allows us to calculate the new APBDC that would obtain.

Packet Clearing House

Flow Data is a Tempting Target

The problem:

With all the neat tricks you can do with flow data, it's becoming a more and more tempting target for people who want to analyze user behavior.

PCH
Packet Clearing House

Thanks, and Questions?

Bill Woodcock
Research Director
Packet Clearing House
woody@pch.net



Beyond the Top Talkers

Empirical Correlation of Conficker-C Infected IP Space

Rhiannon Weaver, CERT/NetSA

FloCon 2010

January 12th 2010

© 2010 Carnegie Mellon University

NO WARRANTY

THIS MATERIAL OF CARNEGIE MELLON UNIVERSITY AND ITS SOFTWARE ENGINEERING INSTITUTE IS FURNISHED ON AN “AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this presentation is not intended in any way to infringe on the rights of the trademark holder.

This Presentation may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

Introduction

Volume statistics and Top-N lists

- Often we prioritize loudest talkers by IP Address
- N has to be small for manual analytic workflow
- NAT, DHCP complicate the picture
- Uncleanliness:
  - Bad guys tend to clump together by administered network
  - Net blocks, responsible parties, and WHOIS, oh my!

What does this look like in a case study?

- Conficker-C botnet
- What does dynamic allocation look like?
- Who do we find with an IP focus? What about /24s?
- Show some pretty pictures

Looking at Conficker-C

Network telescope into infections:

```
rwfilter  --start=2009/03/05:00 --end=2009/03/25:00    /
 --type=in --proto=17 --sport=1024- --pass=stdout  |   /
rwfilter –input=stdin --d-conficker --dyn=conficker.so /
–pass=conCtraffic.rw
```

Conficker-C hosts scan the internet randomly

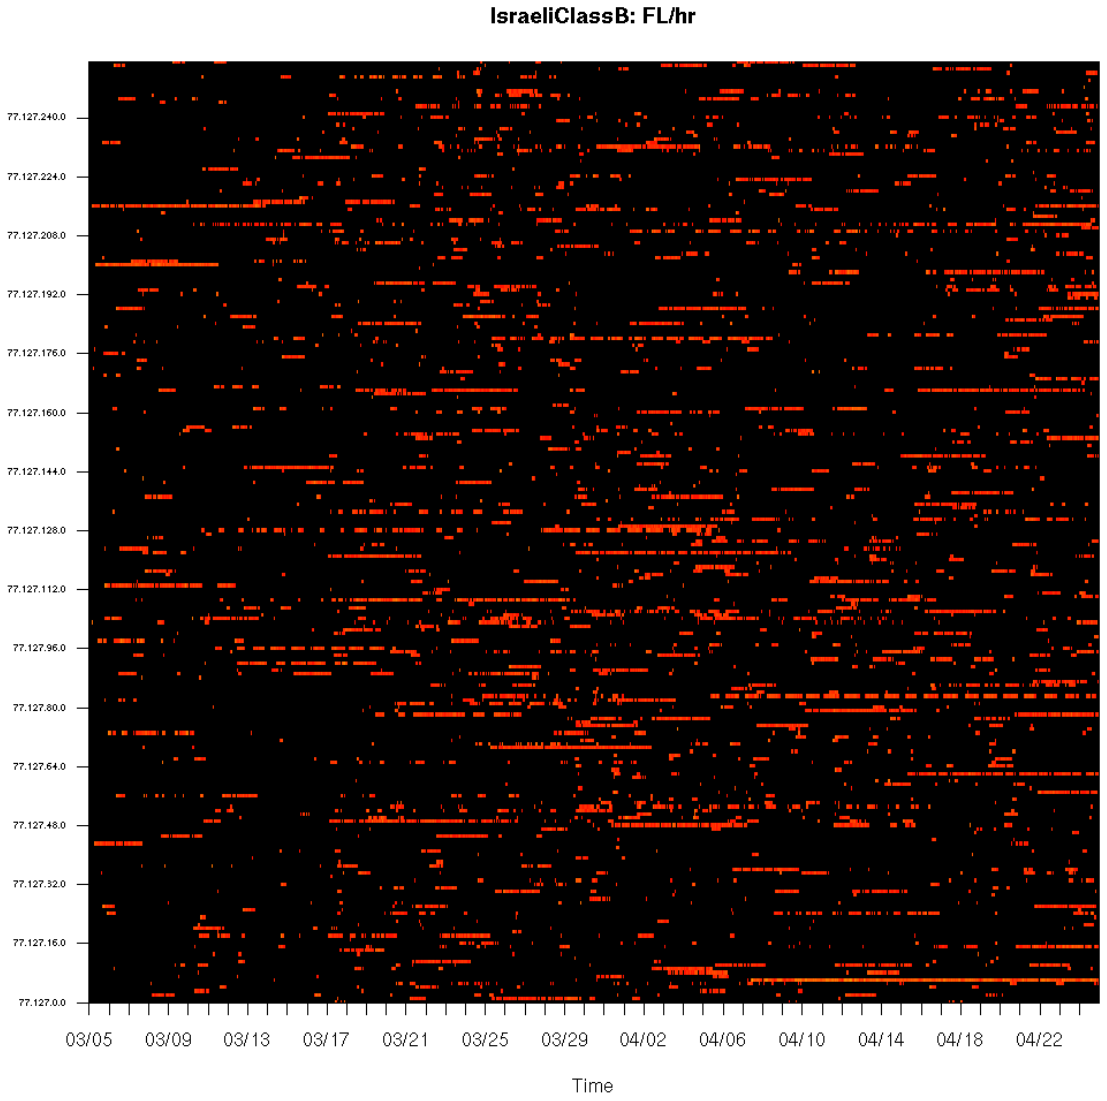
Flow rates vary among IP addresses, /24 blocks

Looking at Conficker-C

Network telescope into infections:

Conficker-C hosts scan the internet randomly

Flow rates vary among IP addresses, /24 blocks

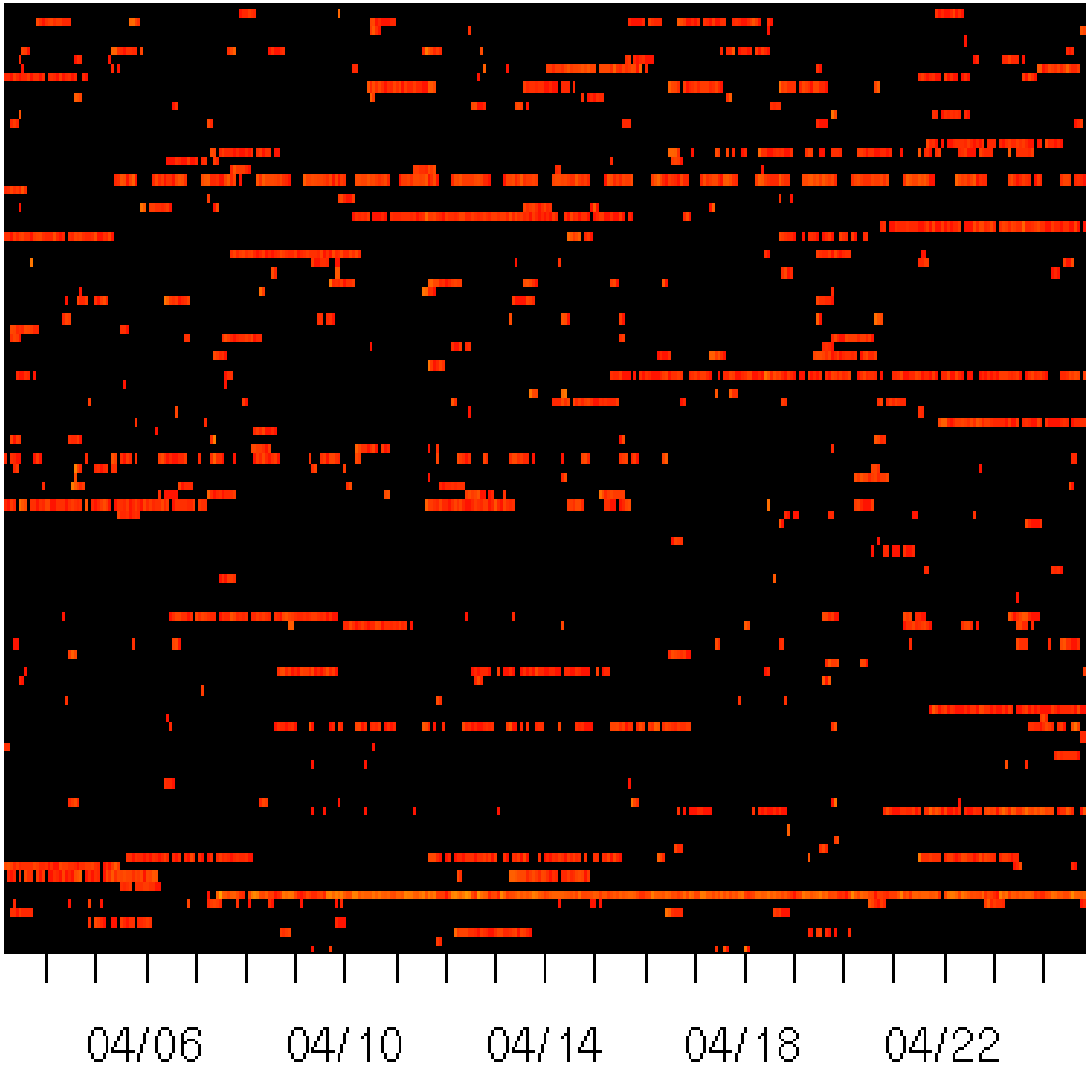


Looking at Conficker-C

Network telescope into infections:

Conficker-C hosts scan the internet randomly

Flow rates vary among IP addresses, /24 blocks

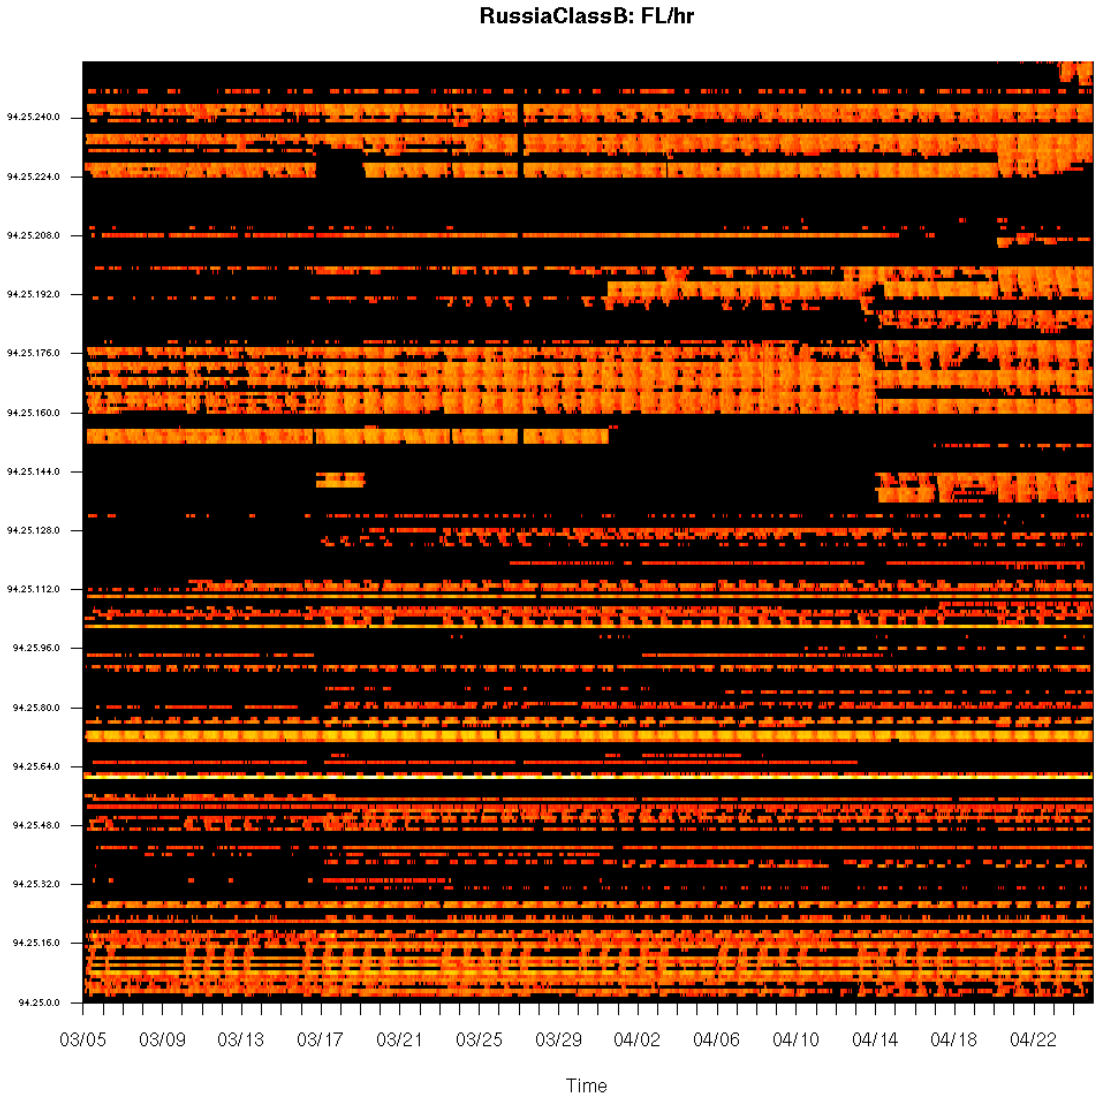


Looking at Conficker-C

Network telescope into infections:

Conficker-C hosts scan the internet randomly

Flow rates vary among IP addresses, /24 blocks



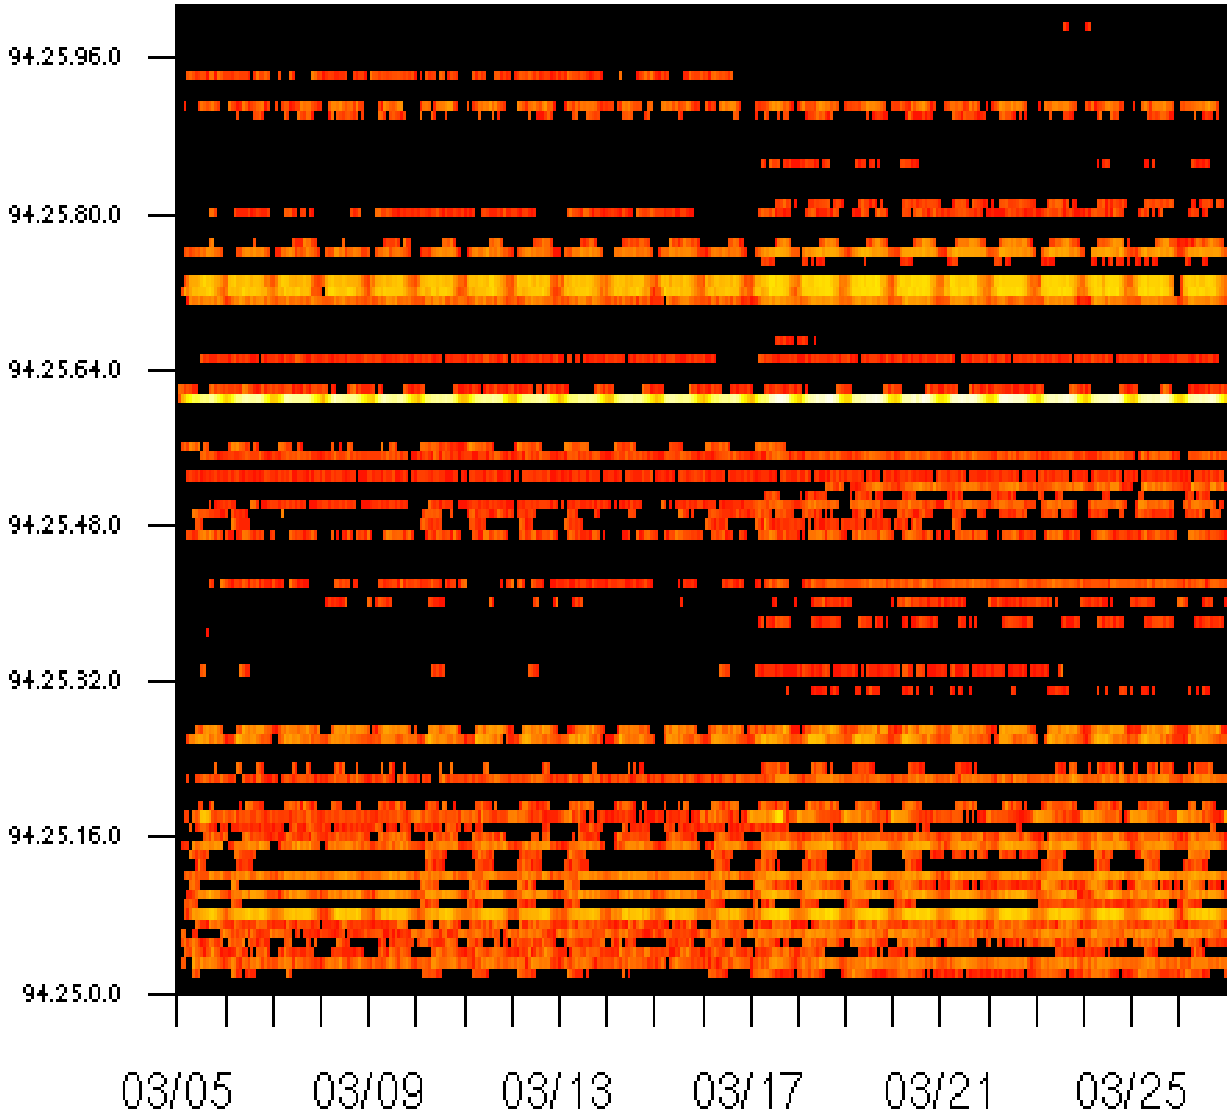
Looking at Conficker-C

Network telescope into infections:

Conficker-C hosts scan the internet randomly

Flow rates vary among IP addresses, /24 blocks

- Some seen for <24 hours, some for every hour
- Some average 2 to 3 pings/hour, some 1000s



Compiling Top-N lists

Data: Top 1000 talkers from March 3 through April 24, 2009

- by day (53 days) and by hour (1272 hours)
- by IP Address and by /24 net block
- Look at blocks appearing in hourly top 20 IP blacklist, /24 blacklist, or both

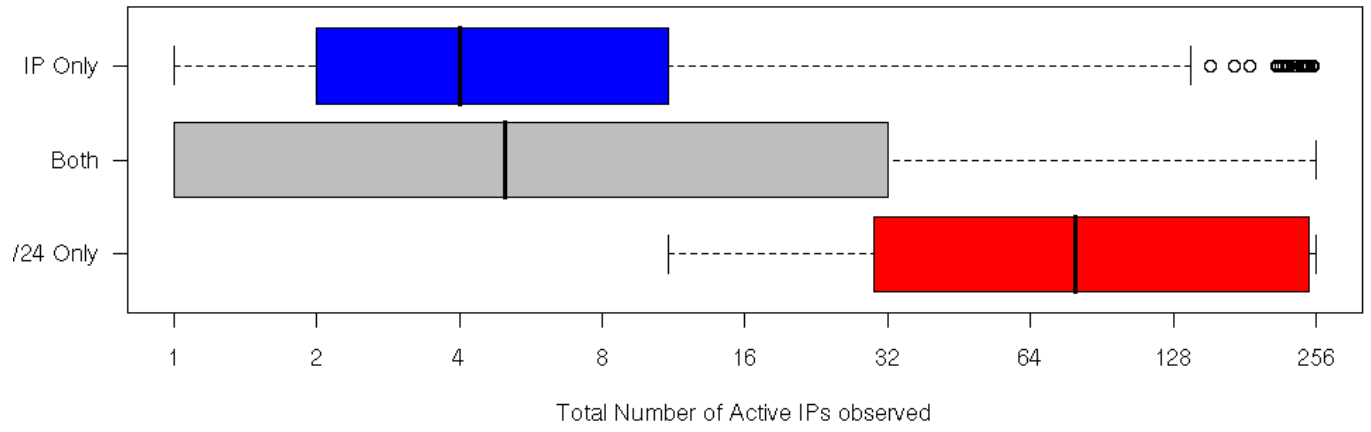
Supplementary data: Flows from all /24s seen performing Con-C scans over the 2 month period.

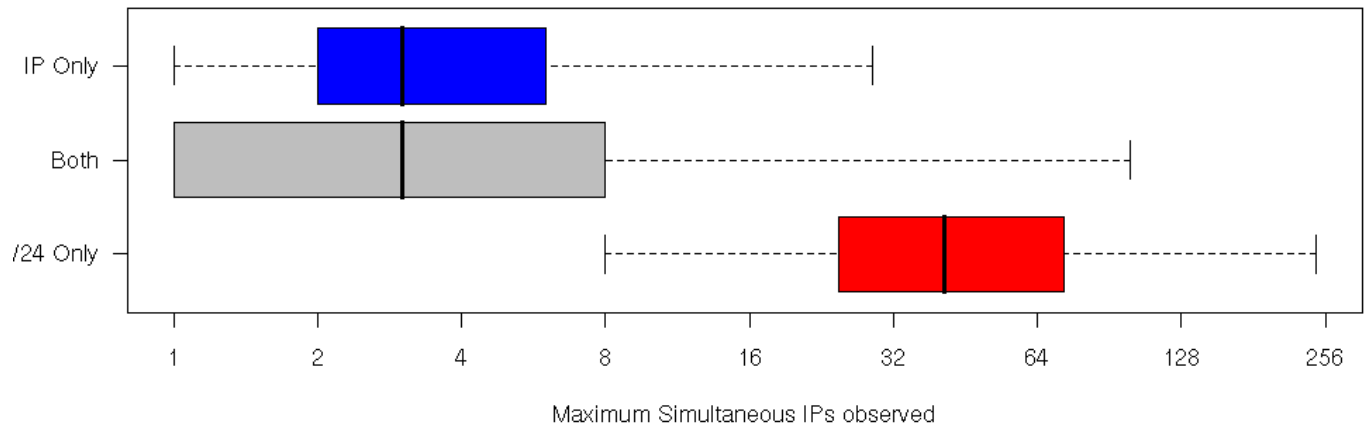
- 1,091,013 blocks.

Summary Information by /24 net block

- TotalIP: Total number of active IP Addresses seen
- Nonzero: Total number of hours observed scanning
- MaxIP: Maximum number of simultaneous IPs per hour
- MaxFL: Maximum number of flows seen per hour
- Mean0FL: Mean number of flows per active hour
- TalkRate: ~Total volume sent (=Mean0FL * Nonzero).
- Country Code

Top 20 Lists by Net Block





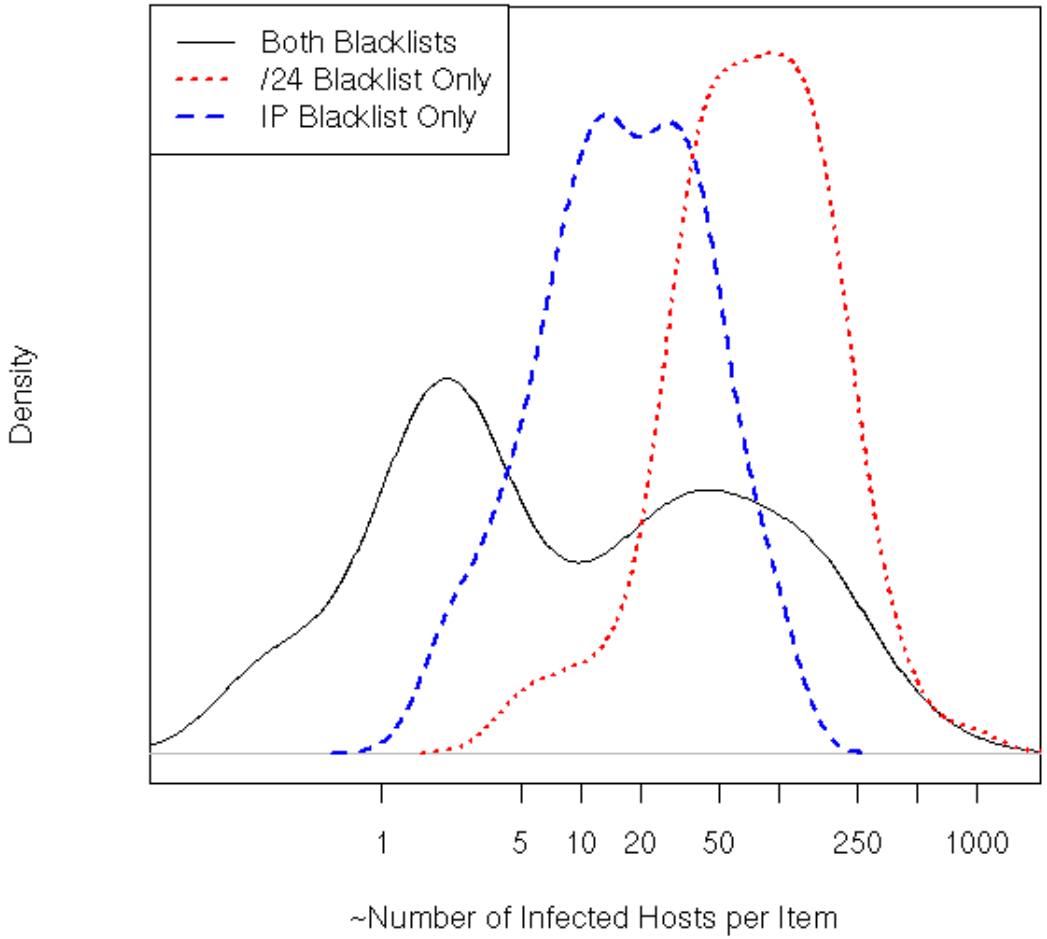
Infected Machines

On average:

approx. 4 flows/hour from one active infected host

Divide Mean Flows by 4 to get approximate mean hosts per hour

(I will affectionately deem this calculation “cowboy statistics”)



Who are the “Plague Dogs”?

Persistently seen in both IP and /24 blacklists:

IP Address	#Top20 IP	#Top20 /24	Avg Hosts/HR	Notes
128.125.179.129	1174	1172	591.8	USC, CA children's hospital
75.141.184.3	1153	972	319.5	Charter Communications, FT Worth TX
94.25.61.240-247 (*)	217	916	485.0	JSC Rostelecom Client Ulyanovsk, Russia (/17)
206.113.142.245	1220	582	245.5	MCI Communications Services, Inc. Ashburn, VA
65.122.8.1	1136	556	243.1	Roosevelt School District
206.160.168.34	1147	215	162.7	Sprint, Reston VA
216.115.160.40	1026	212	159.8	Unibase, Utah, US

(*) NAT from the Russian Class B we saw earlier

Other net blocks are US-based, sparse among neighbors.

Sparse Activity

Slash24	TotalIP	Nonzero	MaxIP	MaxFL	Mean0FL	TalkRate	CC
75.141.129.0	6	28	1	5	2.25	63.00	us
75.141.130.0	2	22	1	23	3.59	79.00	us
75.141.131.0	2	11	1	11	3.09	34.00	us
75.141.132.0	2	4	1	3	2.00	8.00	us
75.141.133.0	1	8	1	14	5.63	45.00	us
75.141.134.0	1	4	1	6	2.50	10.00	us
75.141.135.0	1	3	1	3	1.67	5.00	us
75.141.136.0	2	5	1	4	2.00	10.00	us
75.141.137.0	2	4	1	5	2.25	9.00	us
75.141.138.0	1	2	1	4	2.50	5.00	us
75.141.139.0	1	27	1	6	2.44	65.99	us
75.141.140.0	1	3	1	5	2.33	7.00	us
75.141.152.0	1	8	1	5	1.88	15.00	us
75.141.184.0	3	1222	2	3283	1278.18	1808926.19	us
75.141.187.0	2	9	1	4	2.11	19.00	us

Characteristics (1 phrase or less): Big NATs in small to mid-sized allocations.

Back to Russia

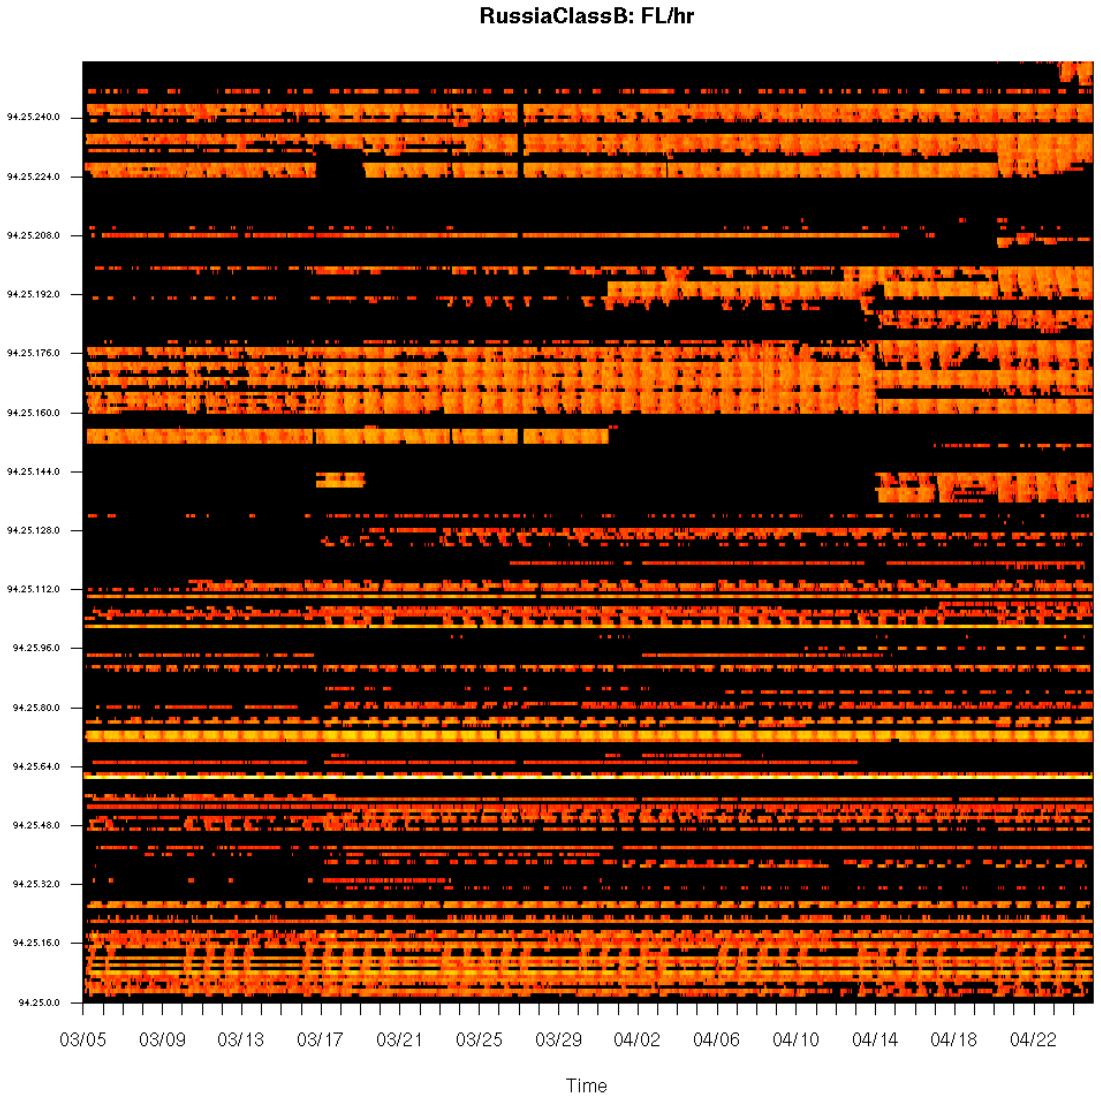
By eye, this /16 has a lot of activity going on

But the top ranks (barring the plague dog block 94.25.61.0):

IP by Hour: 128
/24 by Hour: 103
IP by Day : 361
/24 by Day: 343

Avg Host/HR = 942

Imagine if someone were trying to hide in this network!



Plague Dogs: IP Blacklist

Most visible blocks seen in IP blacklist only

		Highest Rank				
IP Address	#Active in /24	Daily IP	Daily /24	Hourly IP	Hourly /24	Avg Host/HR
77.93.38.7	10	7	43	2	22	129.2
193.239.178.194	13	13	60	7	30	106.2
217.118.82.1	3	14	76	7	30	87.6
195.54.3.58	20	9	61	5	27	87.4
84.22.140.186	2	12	86	6	23	85.6
194.187.148.40	7	11	89	5	44	71.2

- Russian and Ukrainian addresses
- Gaming networks, ISPs
- Allocations /24 through /22

The big NATs are already caught in both IP and /24 blacklists

Plague Dogs: /24 Blacklist

Most visible blocks seen in /24 blacklist only

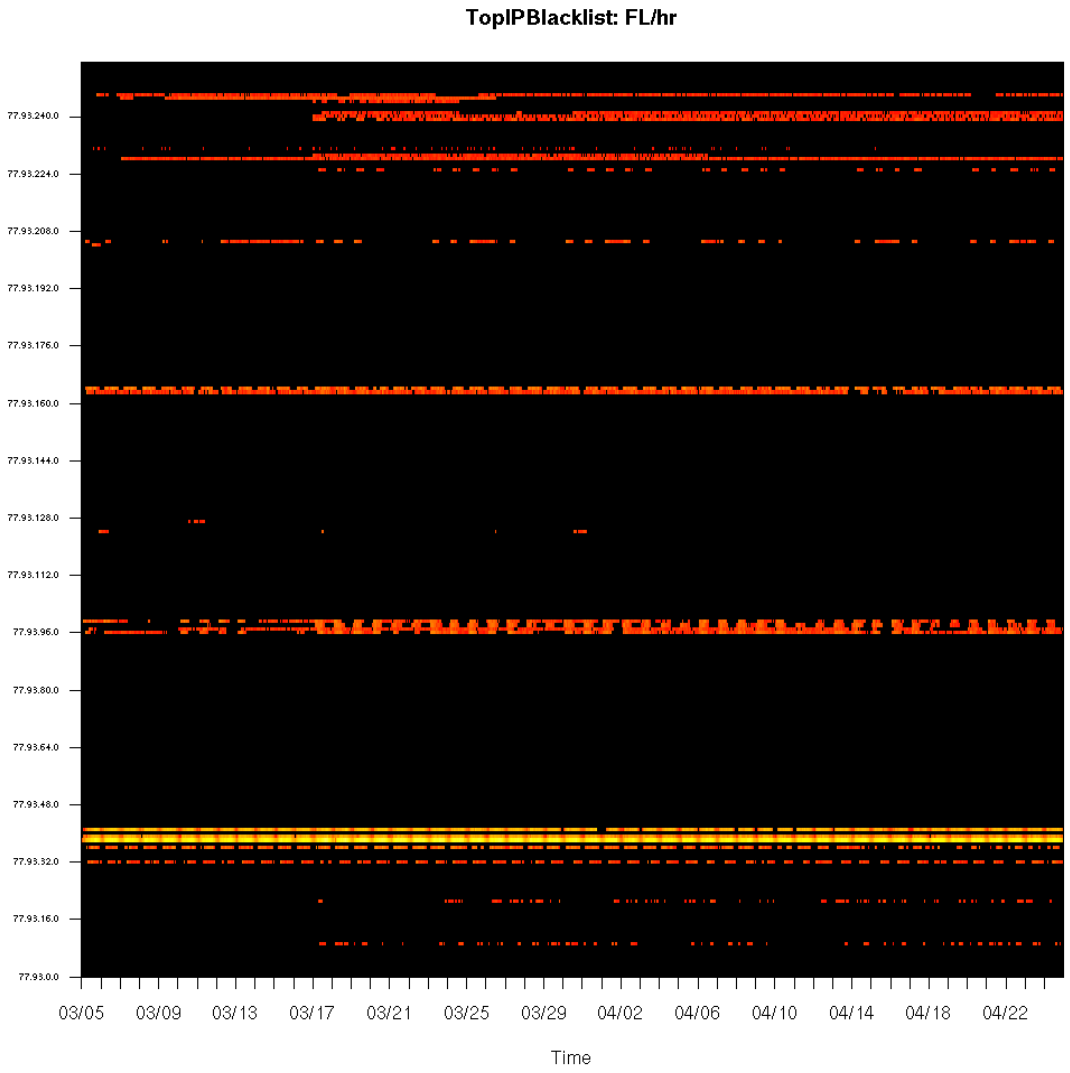
		Highest Rank				
/24 Net block	#Active in /24	Daily IP	Daily /24	Hourly IP	Hourly /24	Avg Host/HR
195.46.34.0	87	219	1	49	1	886.7
168.8.212.0 (*)	97	41	1	13	1	662.1
125.60.241.0	139	355	3	34	1	423.5
83.234.227.0	16	190	6	64	2	300.5
77.120.128.0	256	1000+	6	198	3	291.4
77.120.129.0	256	1000+	5	612	2	280.6

(*) Showed up in IP list top 20 for 2 out of 1272 hours

- Ukraine, Russia, Philippines, US
- ISPs, Telecom, and the Georgia Board of Education
- More “bang for the buck” than IP lists

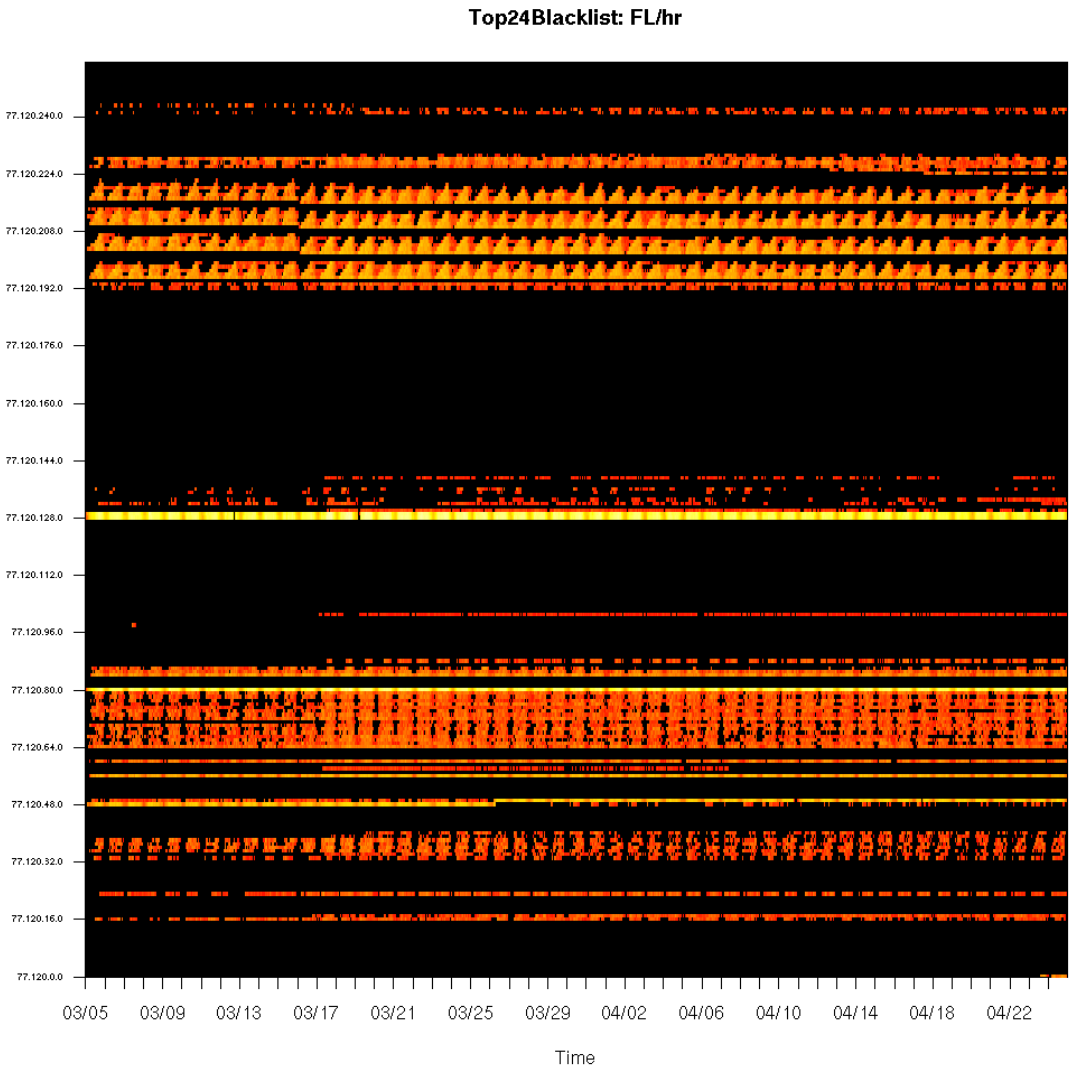
East Europe 77.93.x.x

CC
it: 10
ru: 6
ua: 5
cz: 2
lv: 2
ro: 2



Ukraine Datasvit 77.120.x.x

CC
ua: 83



Digging Deeper

"Spread" in summary statistics

Slash24	TotalIP	Nonzero	MaxIP	MaxFL	Mean0FL	TalkRate	CC
41.221.16.0	254	1222	243	2036	373.60	402491.16	dz
41.221.17.0	254	1222	242	2051	392.71	421225.87	dz
41.221.18.0	254	1221	229	1889	325.68	354943.85	dz
41.221.19.0	254	1222	224	1550	288.70	318780.14	dz
41.221.20.0	15	1104	12	30	5.92	6311.66	dz
41.221.22.0	1	43	1	31	5.74	246.99	dz
41.221.23.0	5	935	4	85	15.30	12083.11	dz
41.221.24.0	7	1142	6	80	9.19	9745.35	dz
41.221.25.0	1	8	1	13	5.63	45.00	dz
41.221.26.0	254	1222	245	2393	395.43	420475.77	dz
41.221.27.0	128	1207	124	1580	280.46	295347.39	dz
41.221.28.0	2	203	2	14	3.74	757.12	dz
41.221.29.0	2	210	2	17	3.93	828.82	dz

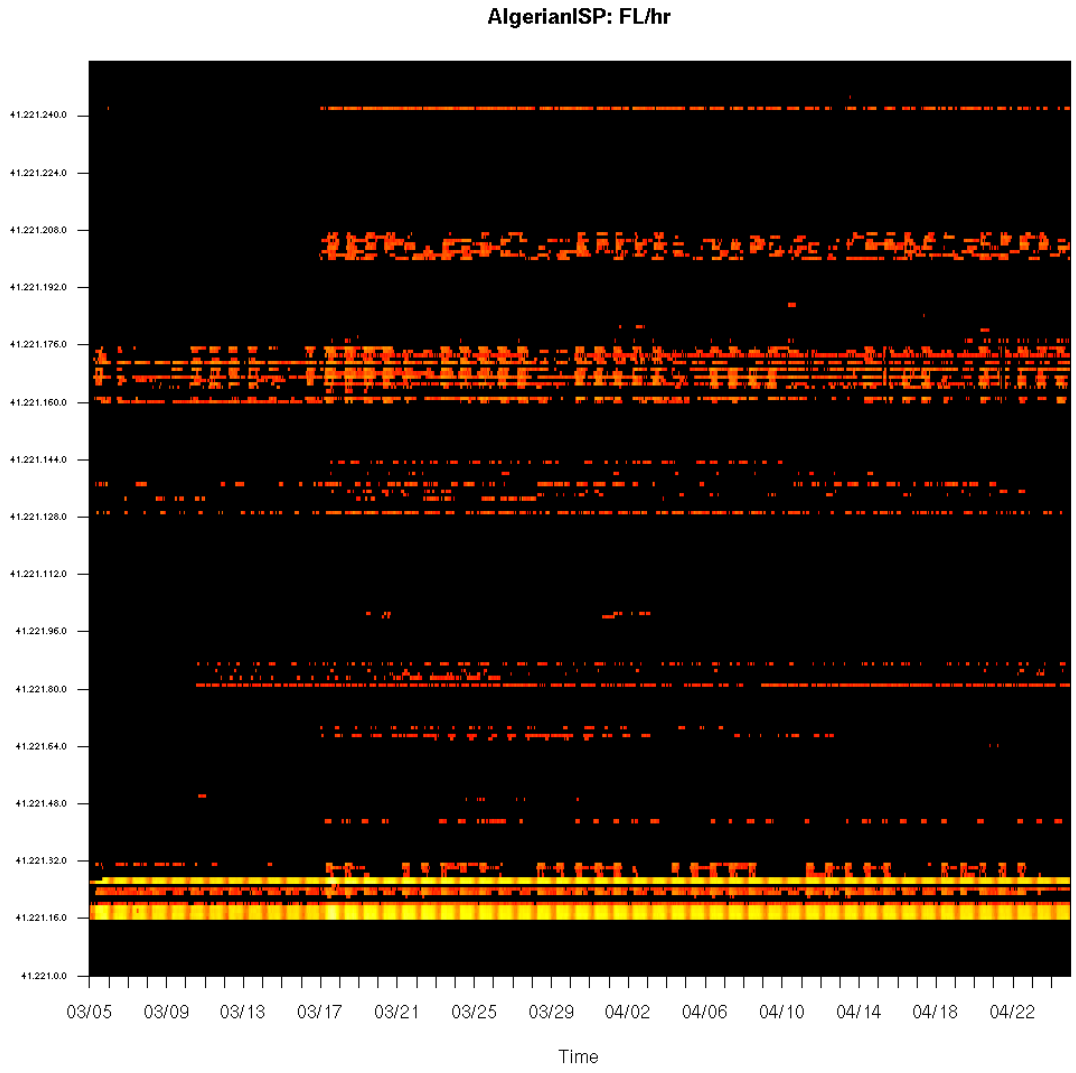
Circled blocks

- ranked between 9-20 in /24 Blacklist
- 70-100 Hosts/HR, each, avg 500 hosts/hr
- none in Top 1000 IPs

Digging Deeper

Slash24	TotalIP	Nonzero	MaxIP	MaxFL	Mean0FL	TalkRate	CC
41.221.160.0	22	491	3	57	7.08	3334.33	ng
41.221.161.0	71	704	10	123	14.76	9832.43	ng
41.221.165.0	41	810	8	88	14.16	8973.12	ng
41.221.166.0	30	340	2	51	6.45	2052.59	ng
41.221.167.0	126	1015	13	90	14.67	14372.03	ng
41.221.168.0	18	387	3	24	6.85	2705.39	ng
41.221.169.0	59	776	8	90	15.84	12033.09	ng
41.221.171.0	28	857	8	70	13.73	10495.70	ng
41.221.172.0	43	243	7	42	7.80	1982.30	ng
41.221.173.0	3	743	2	37	2.67	1936.29	ng
41.221.174.0	16	518	5	40	8.43	4326.92	ng
41.221.175.0	68	437	11	104	13.80	5553.46	ng
41.221.200.0	63	387	3	35	7.24	2798.75	cv
41.221.201.0	59	397	4	27	6.15	2486.93	cv
41.221.202.0	54	346	5	27	6.02	2078.70	cv
41.221.203.0	79	387	5	32	7.70	2927.52	cv
41.221.204.0	76	384	4	69	8.64	3163.38	cv
41.221.205.0	65	384	5	46	6.39	2411.20	cv
41.221.206.0	53	239	4	30	6.40	1453.10	cv
41.221.207.0	41	193	3	23	6.15	1178.31	cv

Europe/Africa 41.221.x.x



Digging Deeper

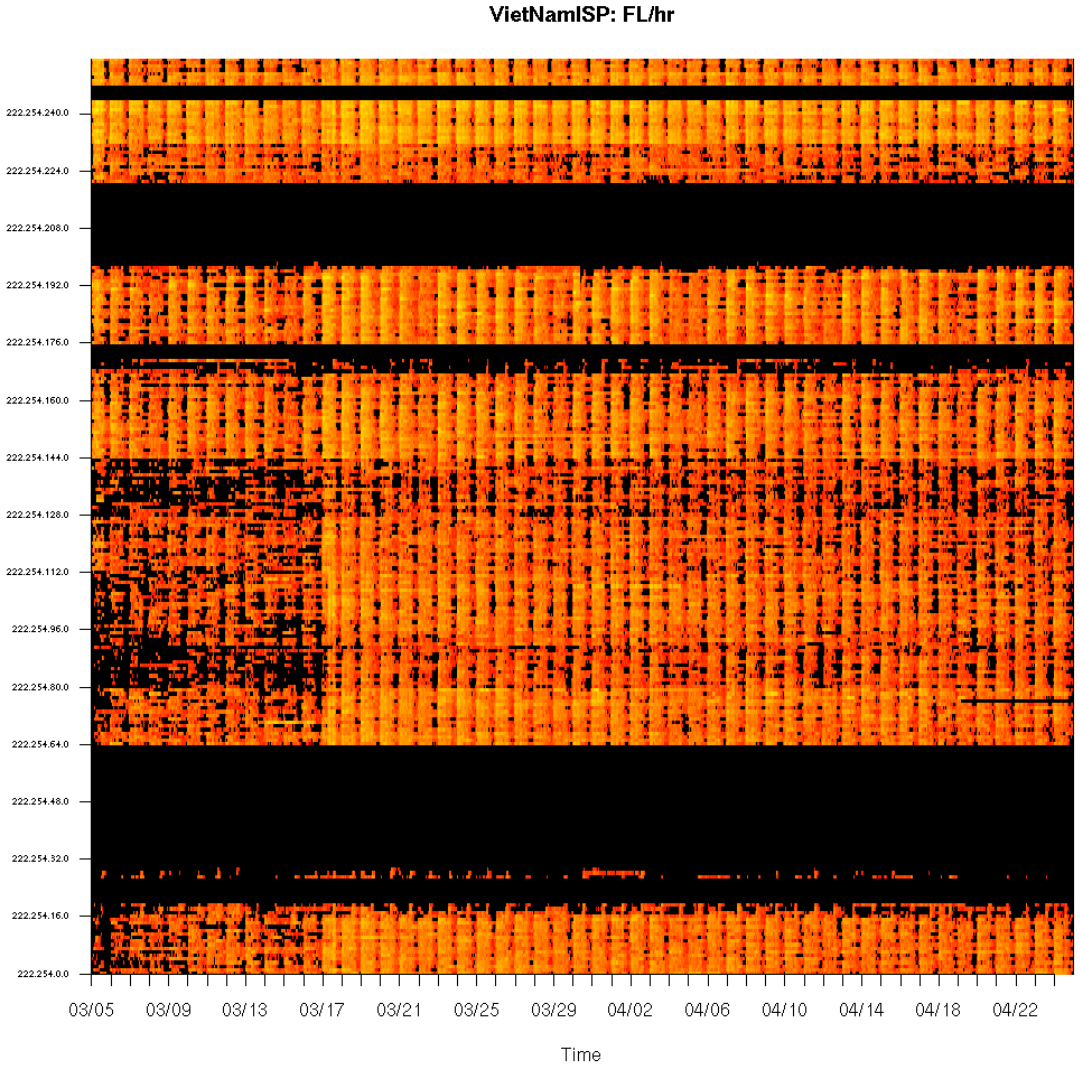
"Spread" in summary statistics

Slash24	TotalIP	Nonzero	MaxIP	MaxFL	Mean0FL	TalkRate	CC
222.254.180.0	229	1132	19	497	35.68	34702.36	vn
222.254.181.0	225	1149	18	550	30.91	31987.85	vn
222.254.185.0	230	1149	16	458	34.88	35522.45	vn
222.254.188.0	246	1108	24	508	44.17	40827.85	vn
222.254.189.0	244	1078	20	549	45.17	39800.97	vn
222.254.190.0	240	1120	19	747	33.12	30691.78	vn
222.254.191.0	243	1153	22	535	48.32	44407.26	vn
222.254.192.0	238	1078	16	477	34.85	34866.95	vn
222.254.194.0	230	1146	16	597	47.48	50603.01	vn
222.254.195.0	232	1045	19	423	44.75	41512.70	vn

All blocks in this grid

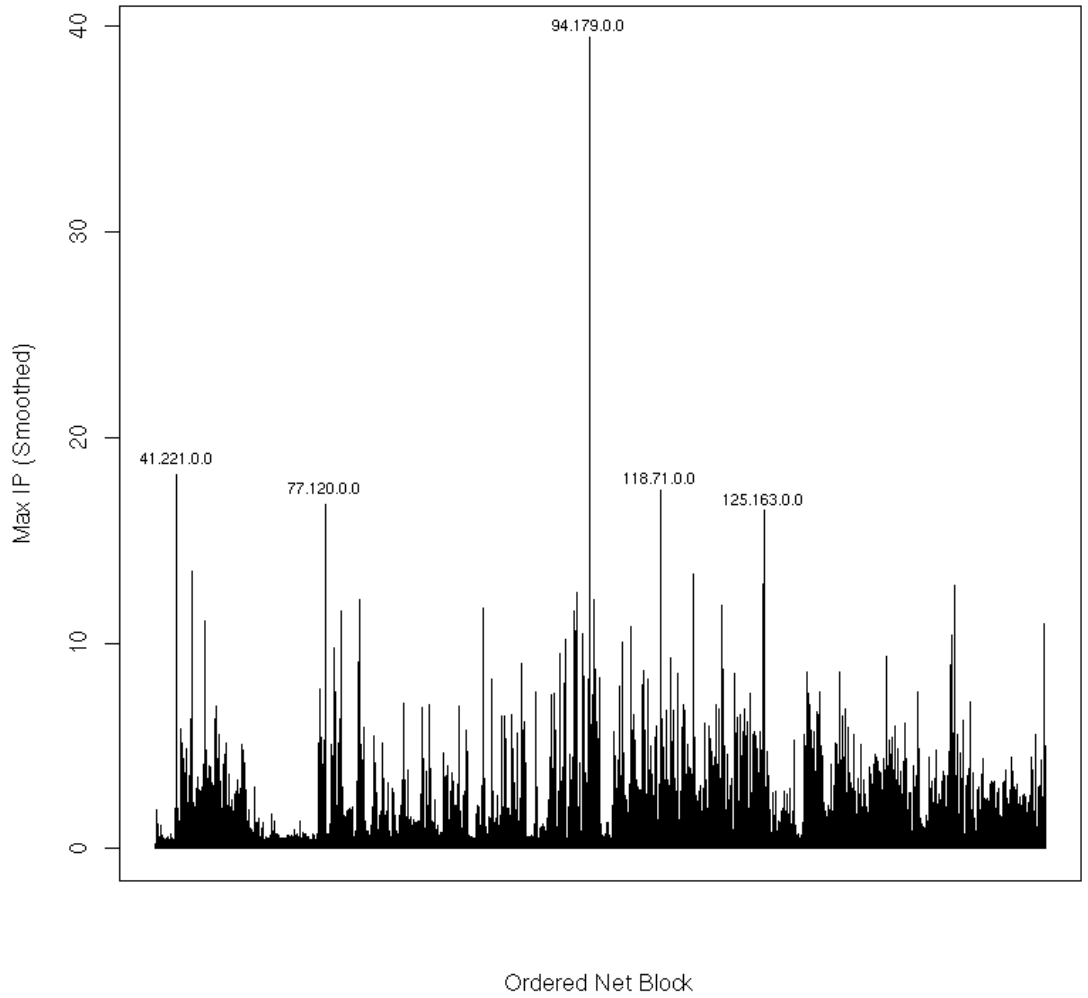
- IP address appeared briefly in at least Top 20 hourly IP rank
- 7-12 Hosts/Hr
- none in Top 20 /24 Blacklist

Viet Nam 222.254.x.x (3 Telecoms)

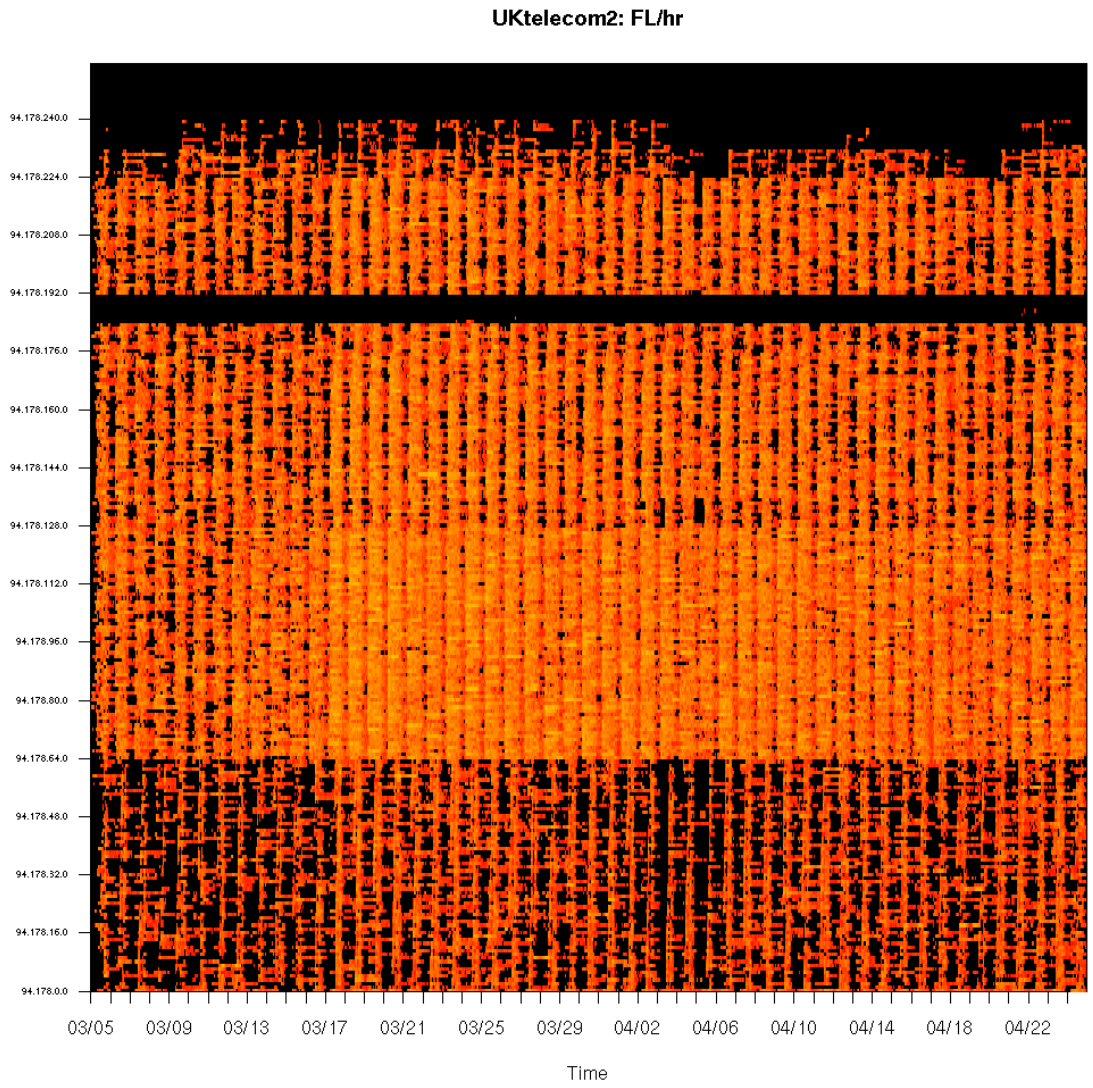


Smoothing Across Net Blocks

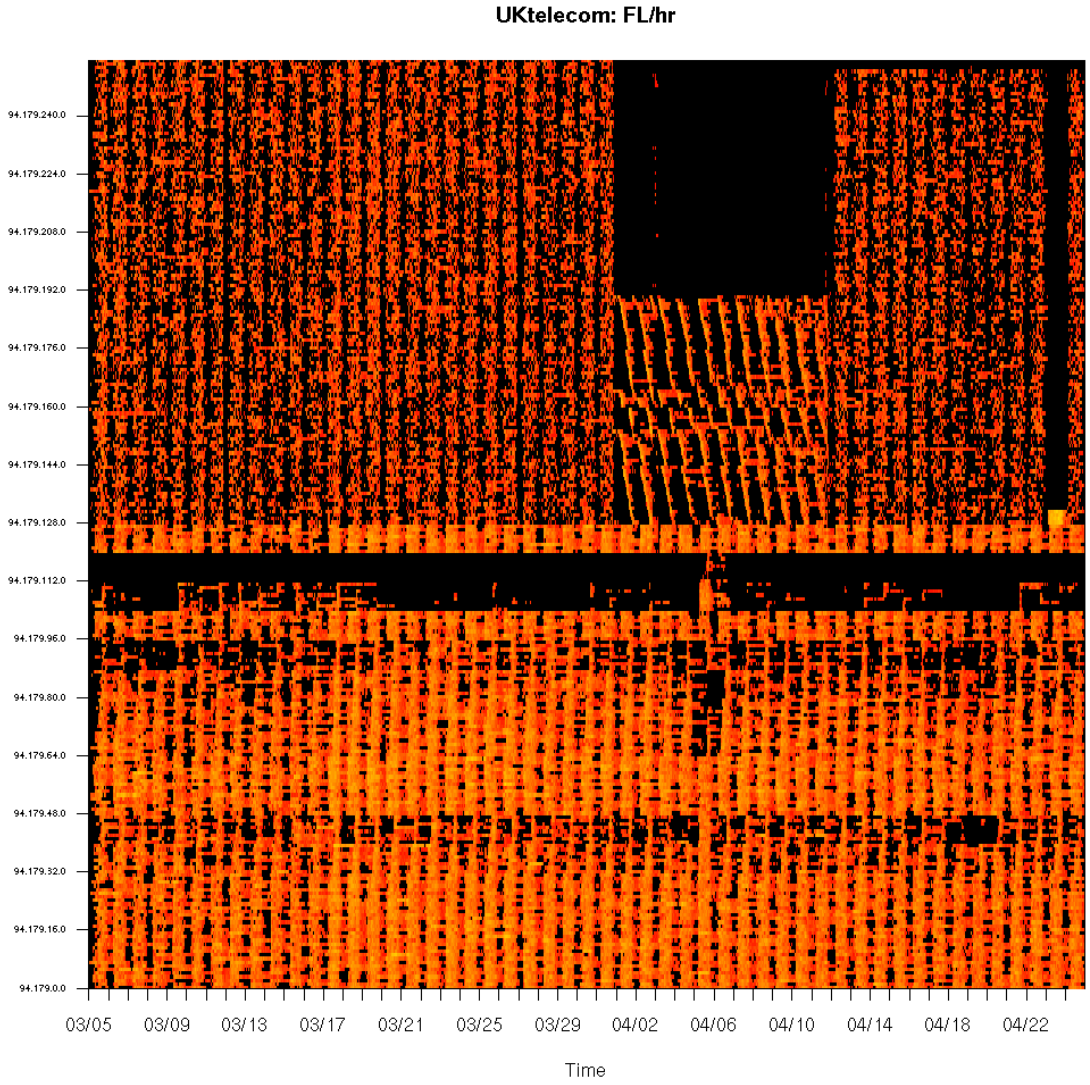
- Highlights contiguous blocks with similar behavior
- Use variable bandwidth for multiple views
- Max # IPs; Bandwidth=75



UKR Telecom 94.179.x.x

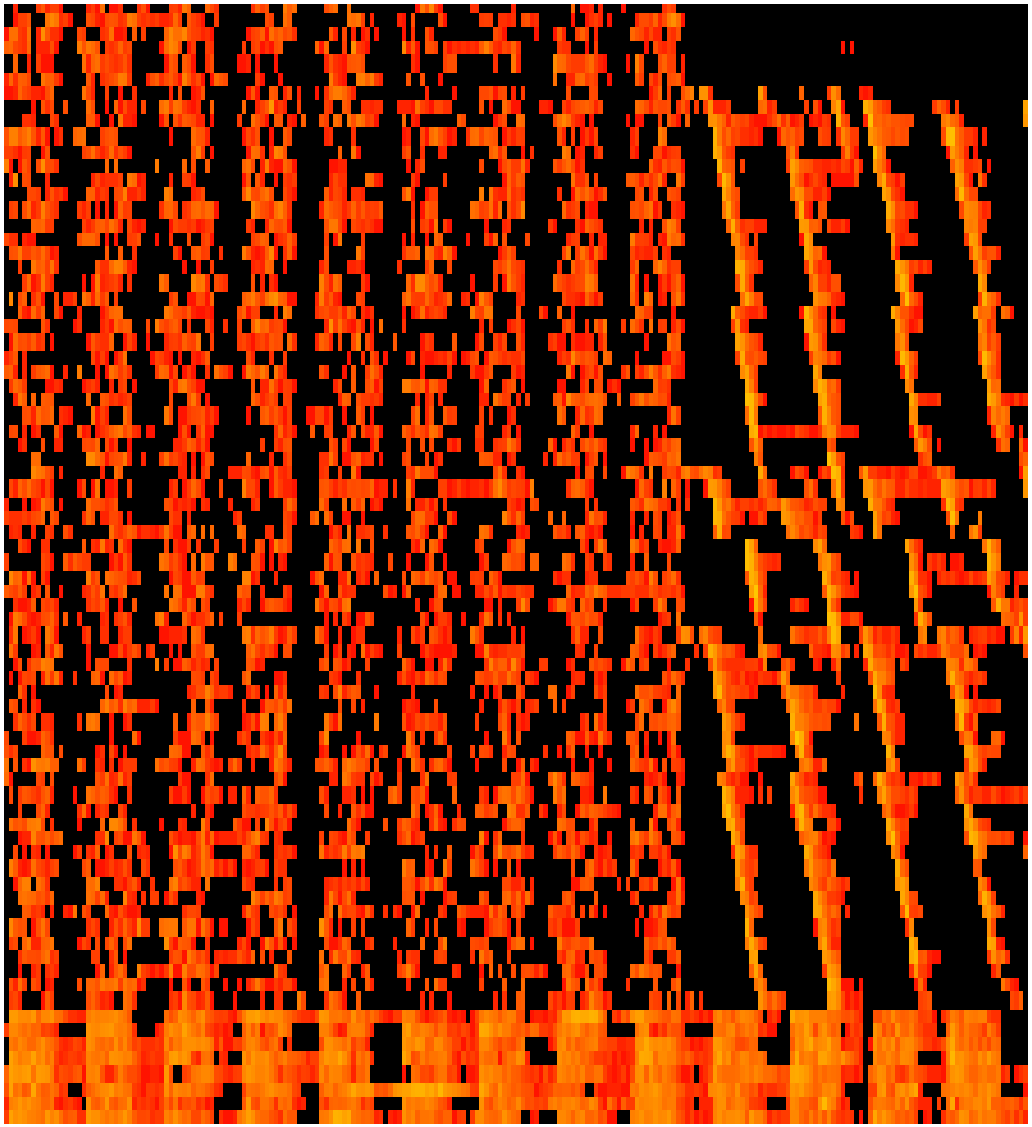


UKR telecom 94.178.x.x

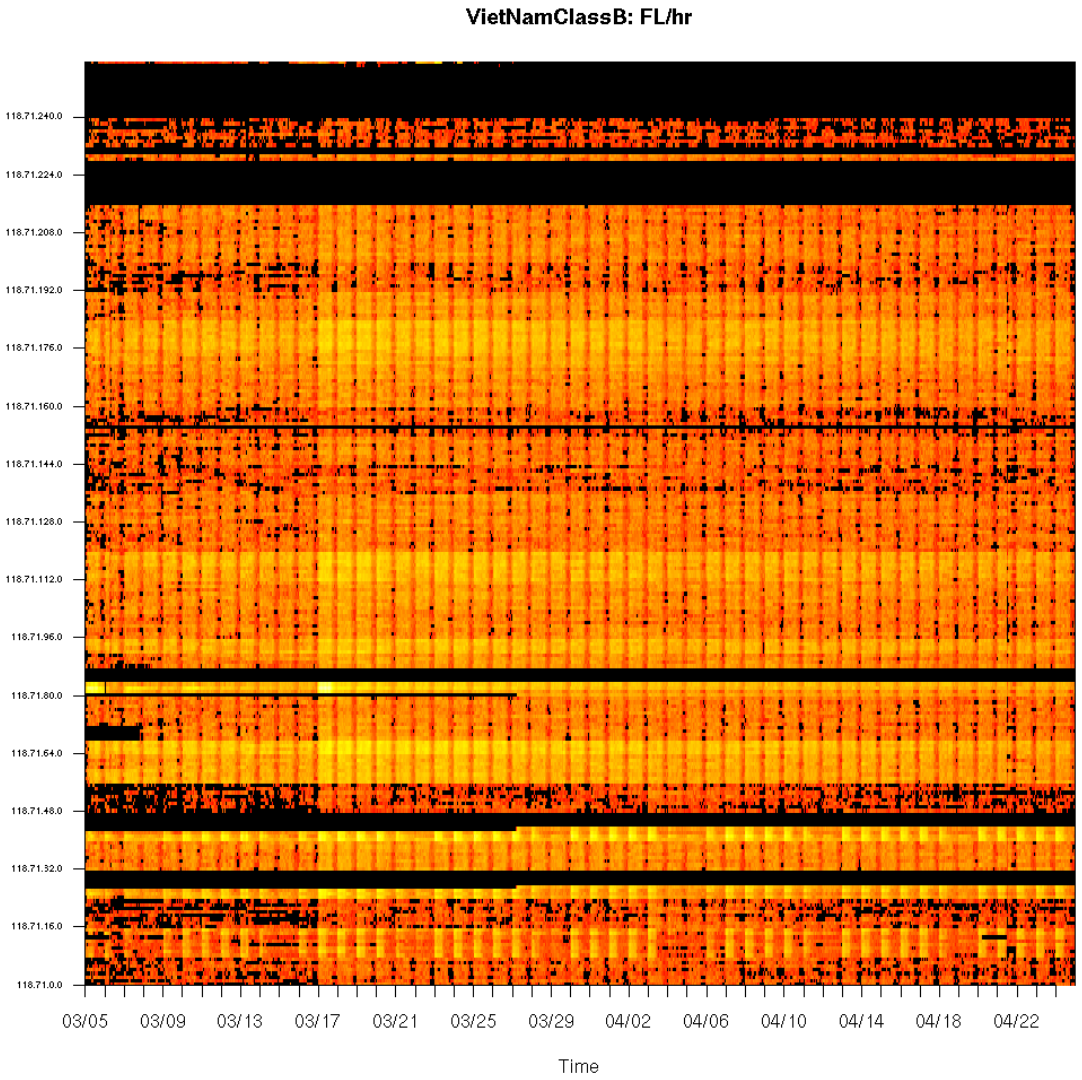


UKR telecom 94.178.x.x

?????



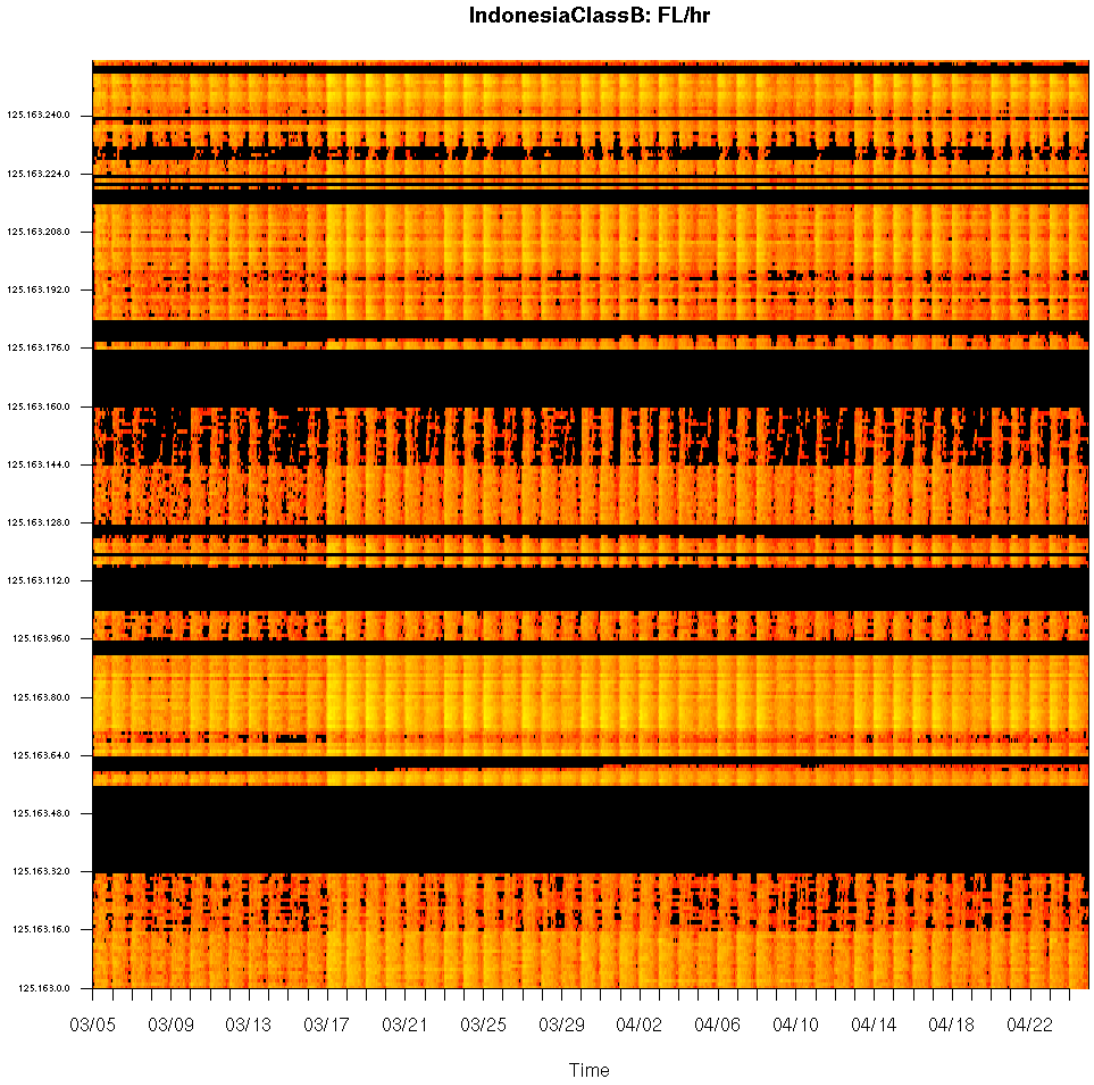
Viet Nam 118.71.x.x (1 Telecom)



Indonesia 125.163.x.x (2 cities)

Semarang

Bandung



Who do we see with Top 20?

Class B (/16)	#Active /24s	#Top 20 IP	#Top 20 /24	Highest Rank Daily IP	Daily /24	Hourly IP	Hourly /24	Avg Host/HR
118.71.0.0	215	8	103	163	4	5	1	2294.2
125.163.0.0	183	2	1	683	270	20	18	2136.6
222.254.0.0	185	29	1	174	172	8	19	1049.9
94.179.0.0	236	0	0	609	669	129	194	760.0
94.178.0.0	256	0	0	693	1000+	216	345	500.4

Conficker Attribution

Who is behind Conficker?

- Conficker A would shut itself down if it detected a Ukranian keyboard setup
- Two IPs were able to interact with both Conficker.B and Conficker.A hosts
  - 200.68.xxx.xx Alternativagratis.com – Argentina
  - 81.23.xxx.xxx Kyivstar.net - Kiev, Ukraine
- Rogue AV Product source is Baka Software (Kiev, UK)
- Two Kiev based ISPs with large netblocks run under the radar
  - Con-C bootstraps a peer list, so it is in the interest of the controllers to have peers available

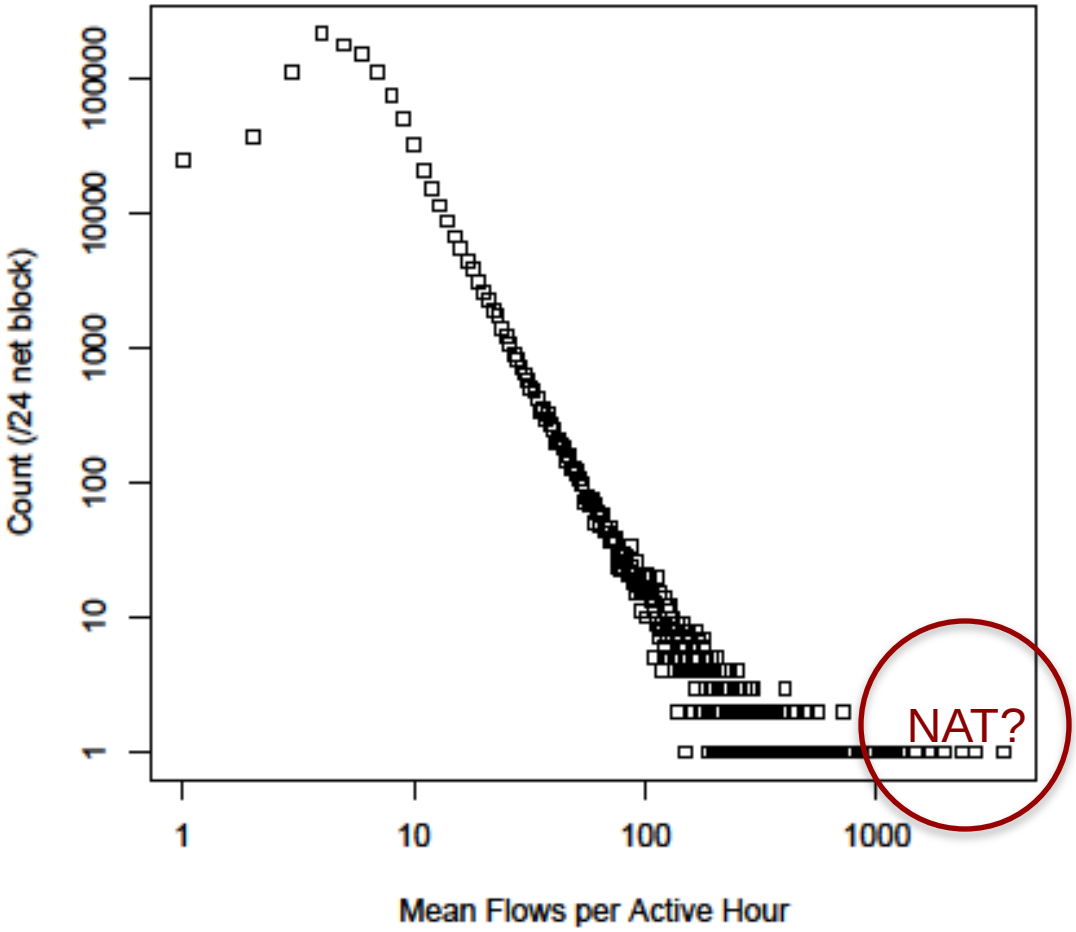
Take-home points/discussion

- Top-N has to go pretty far down to find all the interesting stuff
- Let correlation in summary statistics help us find the big blocks
- Big blocks seem a bit more international (network conventions?)
- Are evil-doers really trying to hide (zippy “bullet-proof” networks), or is it just large scale DHCP?
- Telecom ISPs have abuse contacts, but how useful are they?

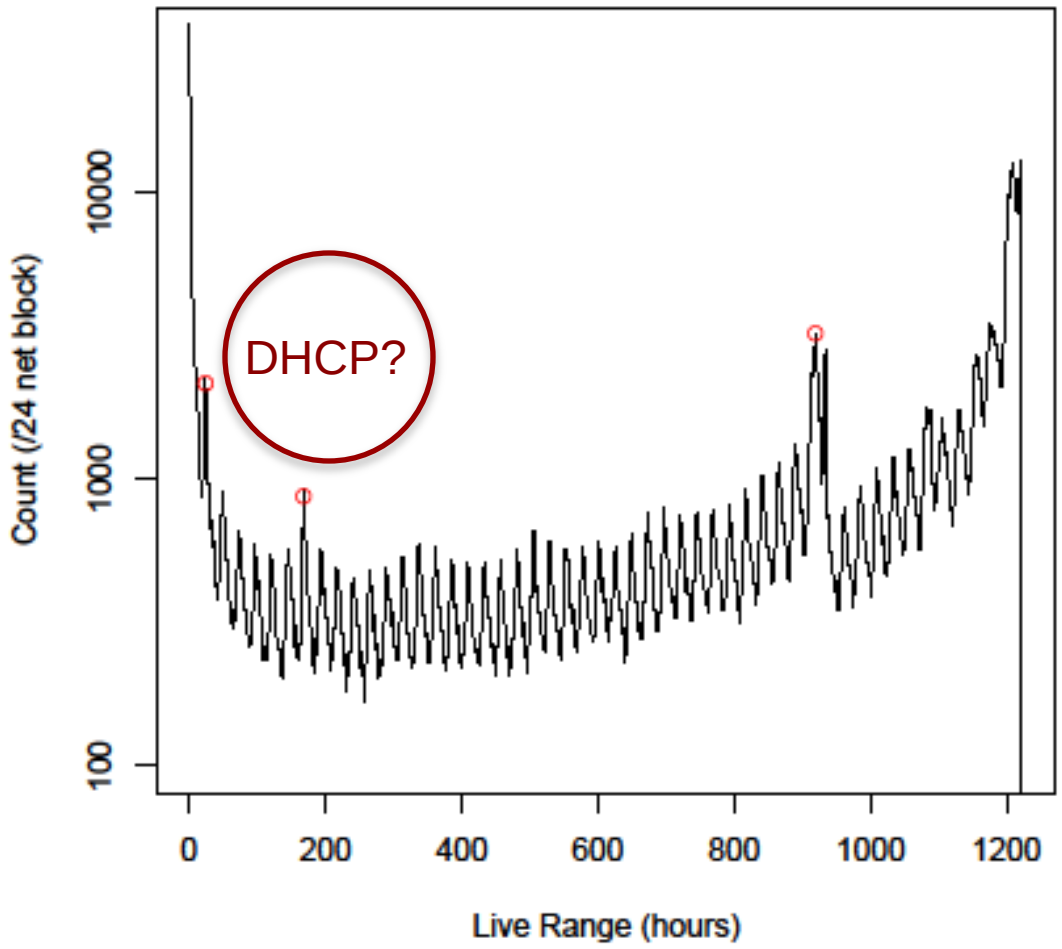
Thank You!

Extra slides

Measurement



Measurement





SiLK and the Virtual Training Environment

George Warnagiris
(Presenter: Markus De Shon)
FloCon 2010

© 2010 Carnegie Mellon University

Goals

Enable remote training delivery

Modularize for maintainability

Standardize across target audiences

Transition as widely as possible

Scale best-practice training to all students

The Virtual Training Environment

Adobe Presenter-based

- Plays using standard Adobe Flash plugin

Slides

- Must be PowerPoint

Audio (MP3)

Video (if available)

Transcription or script

Virtual lab environment

- VMware images
- Virtual network
- Reuse/enhance standard system images where possible

Lessons learned so far

Distributed module development grouped by common theme is possible

- Need project lead to keep organized
- Results can be of uneven quality
- Helps identify best training developers
- Permits rapid development/update

Labs end up decoupled from lectures

Different target audiences require different data sets

Future efforts

Greater utilization of best trainer(s)

Update to keep pace with software releases

Additional tools

- iSiLK
- Visualization

Training access

VTE https://www.vte.cert.org/

SiLK course links http://tools.netsa.cert.org/

My Profile

About this Page

The list of topics and content items in this course are listed to the right. By clicking on the content item, you can retrieve information about and launch the class material.

Functions

-- User --

My VTE
My Profile
My Courses
Using Einstein for Network Traffic Analysis
My Training History

Course View

Using SiLK for Network Traffic Analysis

Description: Using SiLK for Network Traffic Analysis Description
Progress: (100%)
Duration: 3 Hours
Instructor:
Office Hours:
Support Forum:

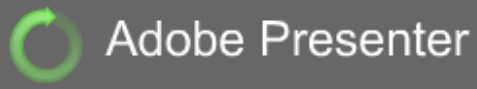
Print Progress Report
Course Certificate
Virtual Office
Access Forum

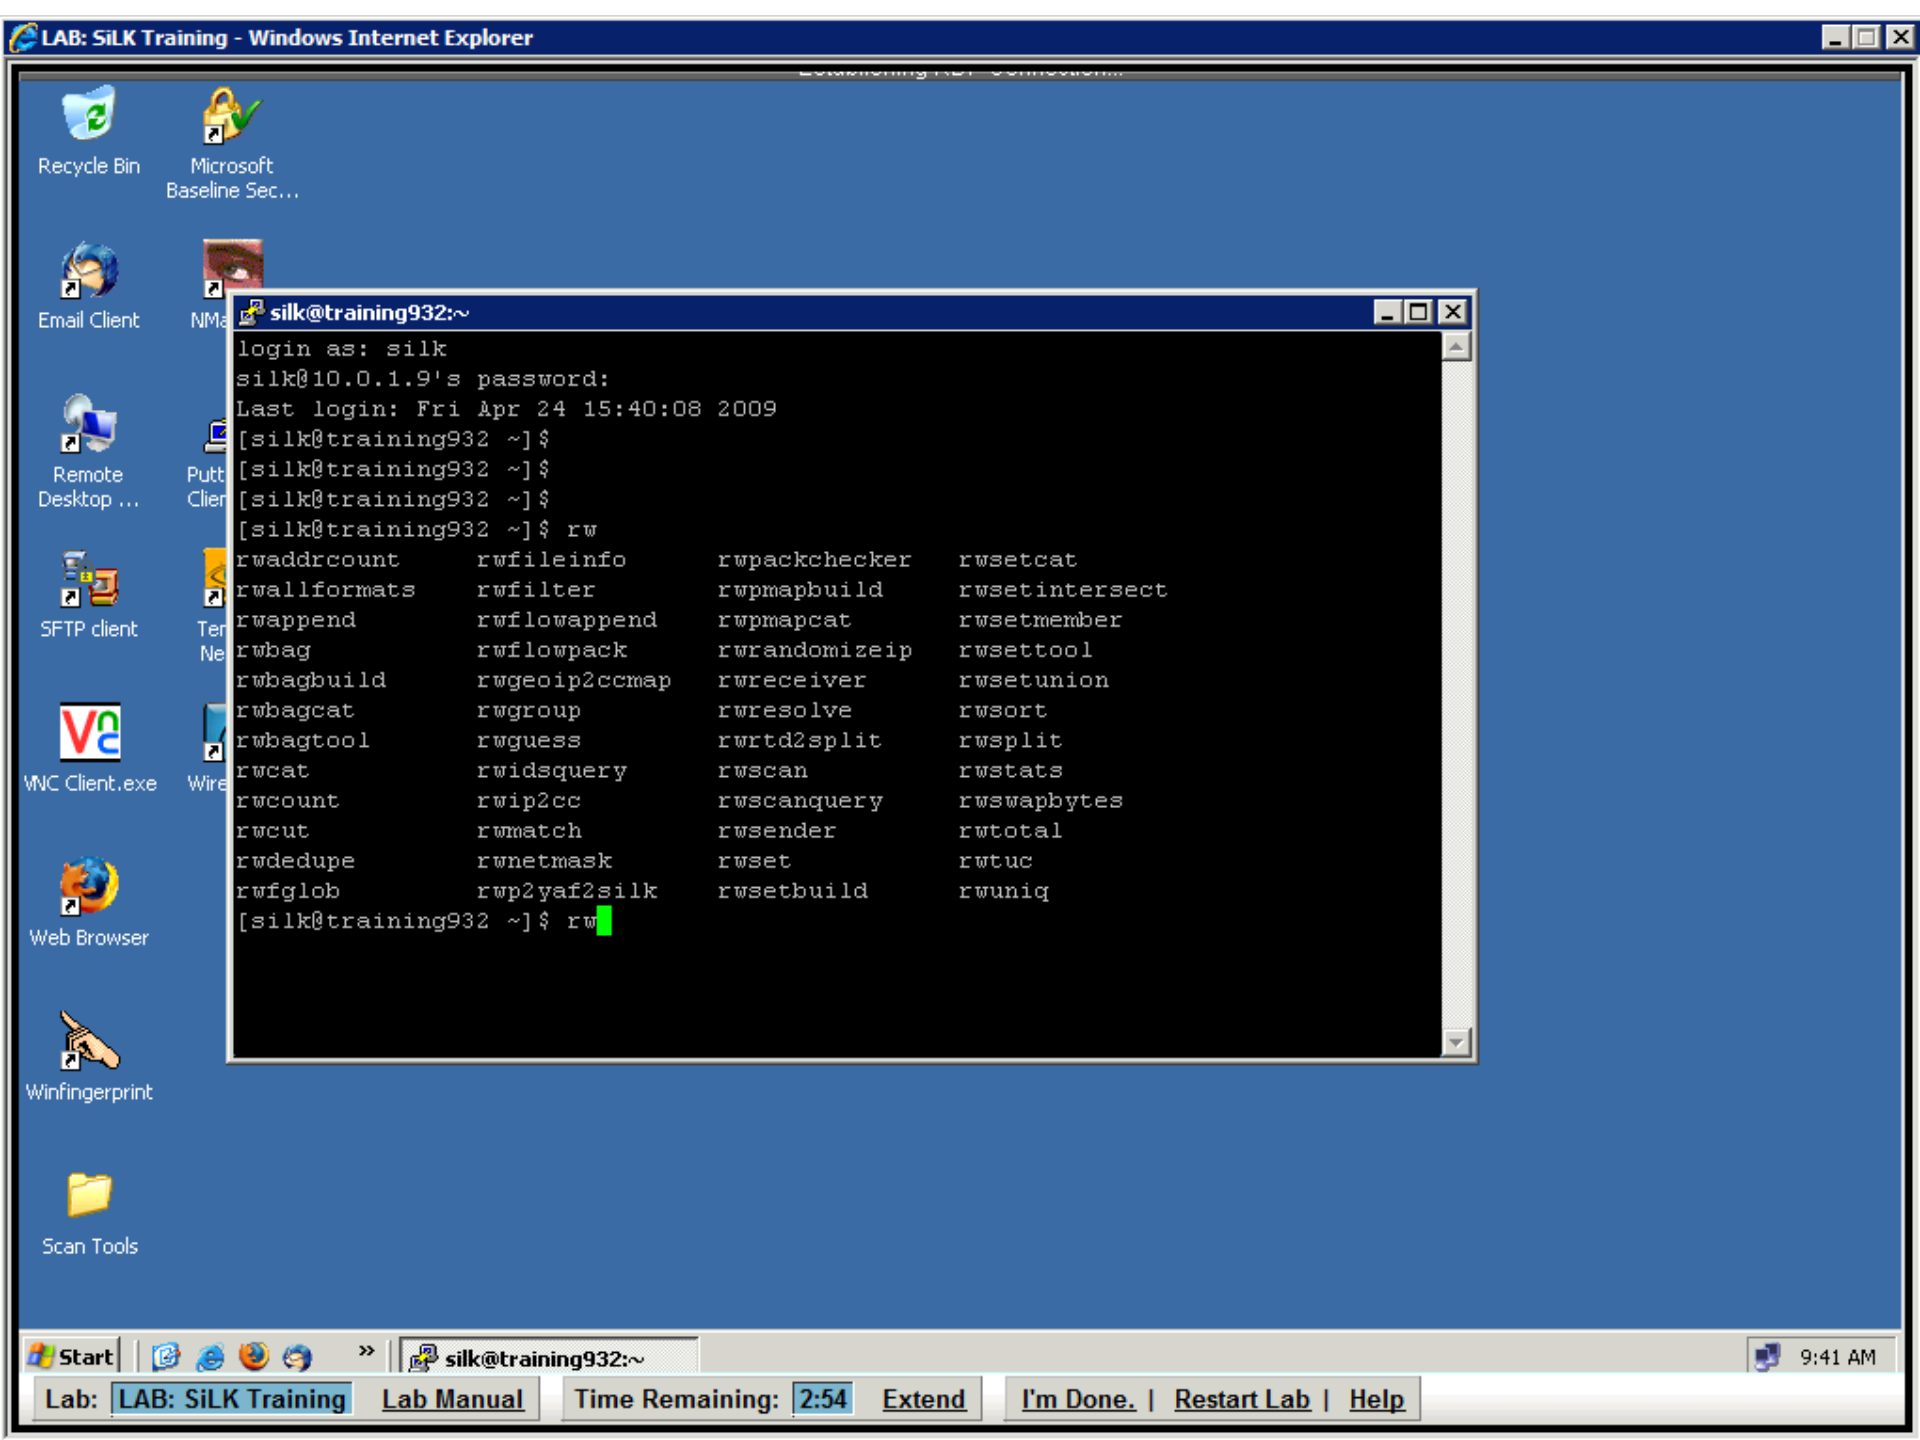
- Using SiLK for Network Traffic Analysis
  - Introduction to Using SiLK
    - Network Flow [C]
    - SiLK Beginning Flow Analysis [C]
  - Basic Tools
    - rwfilter [C]
    - Counting Tools: rwcount, rwstats, rwuniq [C]
    - rwappend-rwsplit [C]
    - rwfileinfo-rwglob [C]
    - rwcut and rwcat [C]
    - rwsort [C]
  - Advanced Tools
    - Sets [C]
    - Prefix Maps (pmaps) [C]
    - Advanced SiLK Tools: Bags [C]
    - Using Tuples with SiLK [C]
  - SiLK Summary

Details

Simulated c

CERT-NetSA - SiLK Training - CountingTools2




LAB: SiLK Training - Windows Internet Explorer
Recycle Bin
Microsoft Baseline Sec...
Email Client
Remote Desktop ...
SFTP client
VNC Client.exe
Web Browser
Winfingerprint
Scan Tools
silk@training932:~
login as: silk
silk@10.0.1.9's password:
Last login: Fri Apr 24 15:40:08 2009
[silk@training932 ~]$
[silk@training932 ~]$
[silk@training932 ~]$
[silk@training932 ~]$ rw
rwaddrcount rwfileinfo rwpackchecker rwsetcat
rwallformats rwfilter rwpmapbuild rwsetintersect
rwappend rwflowappend rwpmapcat rwsetmember
rwbag rwflowpack rwrandomizeip rwsettool
rwbagbuild rwgeoip2ccmap rwreceiver rwsetunion
rwbagcat rwgroup rwresolve rwsort
rwbagtool rwguess rwrtd2split rwsplit
rwcat rwidsquery rwscan rwstats
rwcount rwip2cc rwscanquery rwswapbytes
rwcut rwmatch rwsender rwtotal
rwdedupe rwnetmask rwset rwtuc
rwfglob rwp2yaf2silk rwsetbuild rwuniq
[silk@training932 ~]$ rw
Start
silk@training932:~
9:41 AM
Lab: LAB: SiLK Training
Lab Manual
Time Remaining: 2:54
Extend
I'm Done. | Restart Lab | Help

NO WARRANTY

THIS MATERIAL OF CARNEGIE MELLON UNIVERSITY AND ITS SOFTWARE ENGINEERING INSTITUTE IS FURNISHED ON AN “AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this presentation is not intended in any way to infringe on the rights of the trademark holder.

This Presentation may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT



Network Flow Analysis Using SiLK (v1.1.3)

Training provided by the
CERT Network Situational Awareness Group

January 2009

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community



Course Modules

- Introduction
- Basic SiLK Tools
- Advanced SiLK Tools
- Unix Scripting
- Visualization
- Conclusion

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Module Outline

- About the Training
- Introduction to Flow
- Introduction to Unix
- Beginning Analysis

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

About The Training

This Course Is...

- A "How-To" session on analyzing network flow data
- A tutorial with "What is this" and "Try It" exercises
- Part of a suite of courses being offered

This Course Assumes...

- Fundamental TCP/IP knowledge
- Interest in Network Operations or Security

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Logistics

- Schedule
- Facilities
- Network Connection
- Handouts
- Training Material and Analysts' Handbook
- Quick Reference Guides

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Why Flow Analysis?

What type of analysis has been done with flow?

- Forensics support
- Inventory for large, dynamic networks
- Usage profiling (bandwidth studies)
- Waste (how much traffic is recreational)?
- Identify worm precursors
- Spam Detection

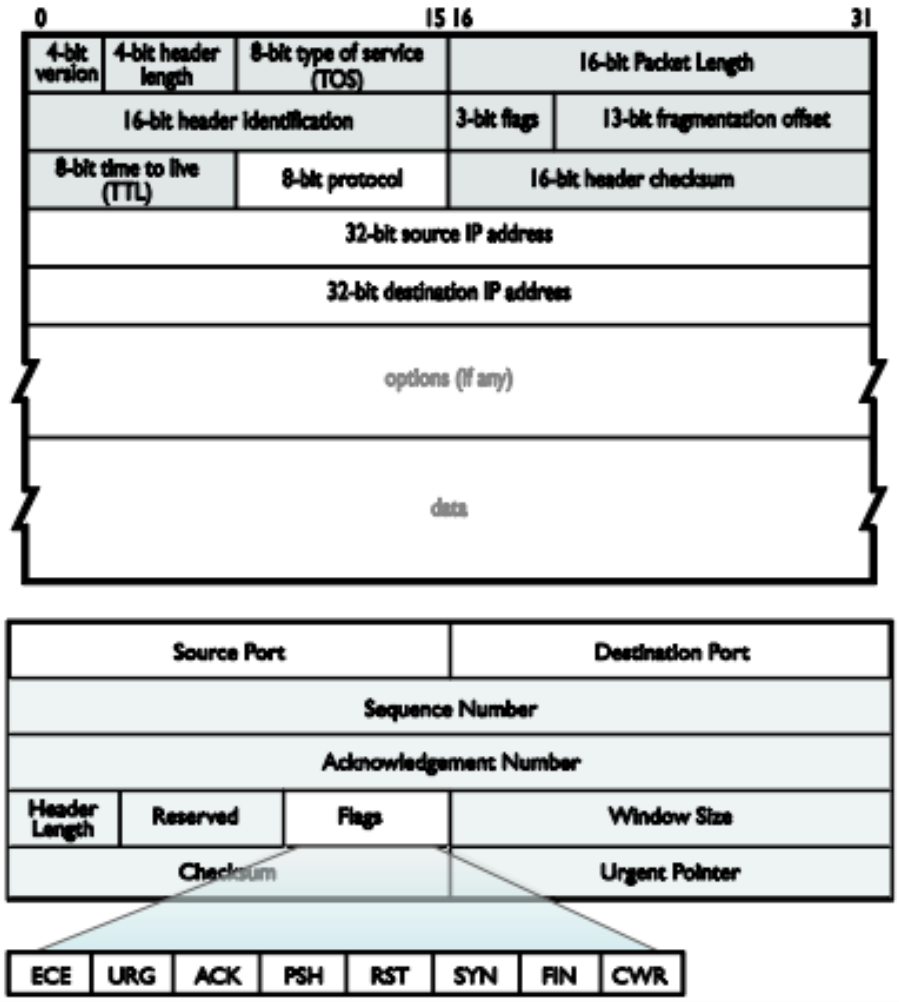
Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

In The Beginning...

First there were packets

- Source & Destination IP
- Size Options
- Routing Information (TTL, fragmentation, QoS, etc) &



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Then Protocols added Ports

Accounting for Packets

The ISP asks "Who Is Using My Bandwidth?"

- Only the routers know
- Very high volume
- Routers know packets, we need them to summarize

NetFlow was developed by Cisco Systems in 1996

- Proprietary
- Evolved into the primary network accounting method
- Supported by most major routers
- IETF Standard on IPFIX (based on Cisco Netflow v9), RFC5101

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What Is a Flow?

A flow is an aggregated record of packets
SiLK flows are defined by five unique attributes

- Source and Destination IP
- IP Protocol (TCP, UDP, ICMP, IPSec, etc)
- Source and Destination Port

These five keys form a “tuple”

- Similar to a “primary key” in a database record

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What does a flow know?

Each unique flow (tuple) has associated attributes

- Timing (start, stop)
- Volume (packets, bytes)
- TCP flags
- Collection Location (sensor, traffic type)
- [Next Hop IP]

Flows get flushed when they close

- Timeouts, TCP FIN/RST, Router resources low

Our flows are *unidirectional*

- The unique key is [sIP, sPort, dIP, dPort, protocol]

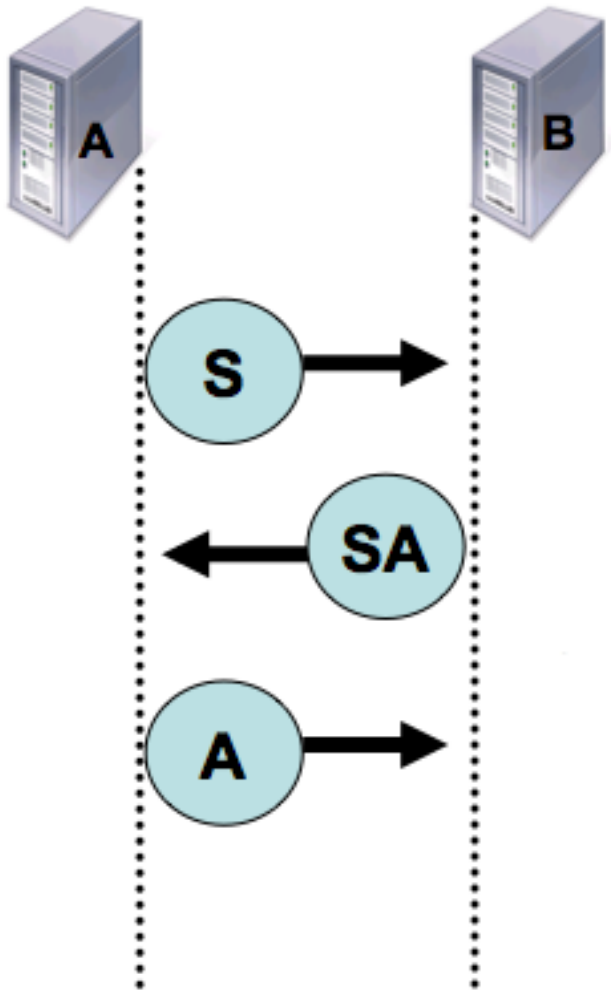
Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Flows are half duplex

For the simple example of a TCP 3-way handshake, consider how flows are counted.

- Flow 1 is created when the sensor observes the first packet between hosts A and B.
- Flow 2 is created with the second packet. Swapped IPs means a new flow.
- With the third packet, and ACK, Flow 1 is updated since the source and destination addresses and ports match. updated.



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What Is This #1

```
           sIP|            dIP|sPort|dPort|pkt| bytes|flags|
63.236.206.174|    72.24.144.5|44800|   25| 21| 19606|FS PA|
   72.24.144.5| 63.236.206.174|   25|44800| 17|  1066|FS PA|
63.236.206.174|    72.24.144.5|44800|   25|  1|    40|  R  |
63.236.206.174|    72.24.144.5|44800|   25|  1|    40|  R  |
63.236.206.174|    72.24.144.5|44800|   25|  1|    40|  R  |
63.236.206.174|   72.24.146.90|44800|   25|  1|    40|  R  |
  72.24.146.90| 63.236.206.174|   25|44800|  1|    49|F  PA|
```


Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What Is This #2

```
           sIP|          dIP|pro|pkts|bytes|       sTime|
66.142.134.179|72.24.150.186|  1|   2|  122|00:00:00.582|
66.142.134.179|72.24.148.123|  1|   2|  122|00:00:00.911|
66.142.134.179| 72.24.146.95|  1|   2|  122|00:00:01.783|
66.142.134.179|72.24.159.123|  1|   2|  122|00:00:01.895|
66.142.134.179|72.24.145.227|  1|   2|  122|00:00:02.220|
66.142.134.179| 72.24.154.87|  1|   2|  122|00:00:02.329|
66.142.134.179|72.24.149.212|  1|   2|  122|00:00:02.550|
66.142.134.179| 72.24.158.18|  1|   2|  122|00:00:02.766|
66.142.134.179| 72.24.150.34|  1|   2|  122|00:00:02.875|
66.142.134.179|72.24.153.102|  1|   2|  122|00:00:02.879|
66.142.134.179| 72.24.144.61|  1|   2|  122|00:00:03.421|
66.142.134.179|  72.24.129.2|  1|   2|  122|00:00:03.530|
66.142.134.179|72.24.129.224|  1|   2|  122|00:00:03.642|
66.142.134.179|72.24.151.196|  1|   2|  122|00:00:04.184|
```


Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What Is This #3

```
          sIP|          dIP|sPort|dPort|pkt|flags|      sTime|
72.24.144.12|   68.8.27.65|63126|   80|  7|FS PA|00:01:31.232|
  68.8.27.65| 72.24.144.12|   80|63126|  5|FS PA|00:01:31.232|
72.24.144.12|   68.8.27.65|63277|   80|  8|FS PA|00:01:42.642|
  68.8.27.65| 72.24.144.12|   80|63277|  8|FS PA|00:01:42.642|
72.24.144.12|   68.8.27.65|63330|   80|  7|FS PA|00:01:51.052|
  68.8.27.65| 72.24.144.12|   80|63330|  5|FS PA|00:01:51.052|
                         [pause]
72.24.144.12|   68.8.27.65|63707|   80|  8|FS PA|00:02:47.722|
  68.8.27.65| 72.24.144.12|   80|63707|  8|FS PA|00:02:47.831|
                         [pause]
72.24.144.12|   68.8.27.65|63957|   80|  8|FS PA|00:03:20.036|
  68.8.27.65| 72.24.144.12|   80|63957|  8|FS PA|00:03:20.036|
                         [pause]
72.24.144.12|   68.8.27.65|64504|   80|  8|FS PA|00:04:12.501|
  68.8.27.65| 72.24.144.12|   80|64504|  8|FS PA|00:04:12.501|
```


Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What Is This #4

```
          sIP|          dIP|sPort|dPort|pkts|flags|       sTime|
72.24.129.20|82.80.30.150|   80| 1220|  152| S PA|00:00:23.602|
82.80.30.150|72.24.129.20| 1220|   80|   90| SRPA|00:00:23.602|
72.24.129.20|82.80.30.150|   80| 1221|1126| S PA|00:00:23.710|
82.80.30.150|72.24.129.20| 1221|   80|  413| SRPA|00:00:23.710|
72.24.129.20|82.80.30.150|   80| 1223|   63| S PA|00:00:26.341|
82.80.30.150|72.24.129.20| 1223|   80|   39| S PA|00:00:26.341|
72.24.129.20|82.80.30.150|   80| 1224|    8| S PA|00:00:26.883|
82.80.30.150|72.24.129.20| 1224|   80|    7| SRPA|00:00:26.883|
82.80.30.150|72.24.129.20| 1223|   80|    1|  R A|00:01:33.068|
```


Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

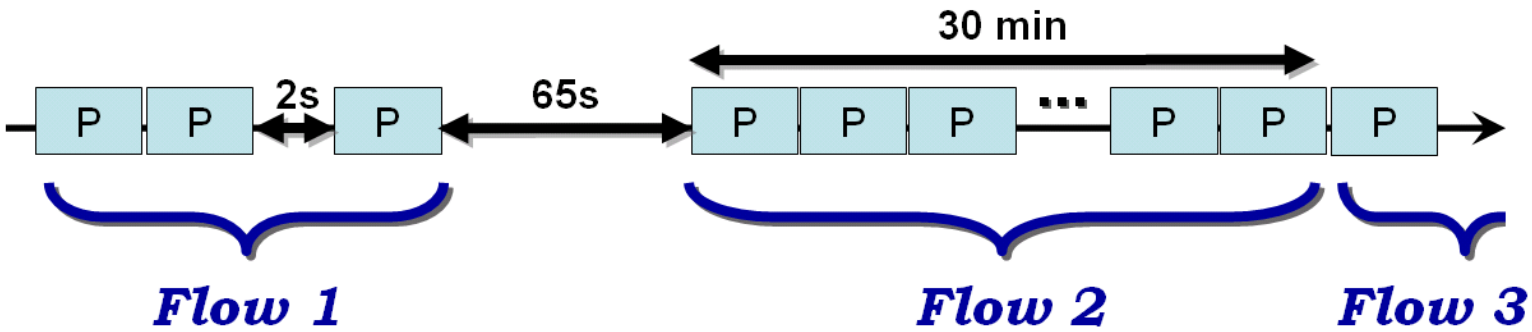
It's all a matter of timing

The flow buffer has to be kept manageable
Inactivity timeout:

- If there's no activity within [30] seconds, flush the flow

Active timeout:

- Flush all flows open for [30] minutes



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What Is This #5

```
          sIP|          dIP|sPort|dPort|pro|Pkt|byte|   sTime|   dur|
8.97.138.194|72.24.145.68|  500|  500| 17|  1| 112|00:02:31| 0.000|
72.24.145.68|8.97.138.194|  500|  500| 17|  1| 112|00:02:31| 0.000|
8.97.138.194|72.24.145.68|  500|  500| 17|  2| 224|00:13:53|40.498|
72.24.145.68|8.97.138.194|  500|  500| 17|  2| 224|00:13:53|40.498|
8.97.138.194|72.24.145.68|  500|  500| 17|  2| 224|00:25:10|45.582|
72.24.145.68|8.97.138.194|  500|  500| 17|  2| 224|00:25:10|45.582|
8.97.138.194|72.24.145.68|  500|  500| 17|  1| 112|00:36:03| 0.000|
72.24.145.68|8.97.138.194|  500|  500| 17|  1| 112|00:36:03| 0.000|
8.97.138.194|72.24.145.68|  500|  500| 17|  1| 112|00:43:19| 0.000|
72.24.145.68|8.97.138.194|  500|  500| 17|  1| 112|00:43:19| 0.000|
8.97.138.194|72.24.145.68|  500|  500| 17|  3| 336|00:47:30|46.088|
72.24.145.68|8.97.138.194|  500|  500| 17|  3| 336|00:47:30|46.088|
72.24.145.68|8.97.138.194|  500|  500| 17|  1| 112|00:53:32| 0.000|
8.97.138.194|72.24.145.68|  500|  500| 17|  1| 112|00:53:32| 0.000|
72.24.145.68|8.97.138.194|  500|  500| 17|  2| 208|00:58:42| 0.000|
8.97.138.194|72.24.145.68|  500|  500| 17| 20|2232|00:58:49|90.095|
```


Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What Is This #6

```
          sIP|          dIP|sPort|dPort| pkts|flg|       sTime|     dur|
 72.24.147.6|58.210.70.72|35282|   22|29640|PA |00:00:11.361|1800.63|
58.210.70.72| 72.24.147.6|   22|35282|29633|PA |00:00:11.911|1800.08|
 72.24.147.6|58.210.70.72|35282|   22|30824|PA |00:26:23.092|1800.82|
58.210.70.72| 72.24.147.6|   22|35282|30825|PA |00:26:23.092|1800.82|
 72.24.147.6|58.210.70.72|35282|   22|29346|PA |00:56:24.020|1800.90|
58.210.70.72| 72.24.147.6|   22|35282|29347|PA |00:56:24.020|1800.90|
 72.24.147.6|58.210.70.72|35282|   22|31107|PA |01:00:10.783|1800.20|
58.210.70.72| 72.24.147.6|   22|35282|31113|PA |01:00:11.301|1800.68|
 72.24.147.6|58.210.70.72|35282|   22|29227|PA |01:26:25.036|1800.95|
58.210.70.72| 72.24.147.6|   22|35282|29228|PA |01:26:25.036|1800.95|
 72.24.147.6|58.210.70.72|35282|   22|30880|PA |01:56:26.096|1800.82|
58.210.70.72| 72.24.147.6|   22|35282|30878|PA |01:56:26.096|1800.82|
 72.24.147.6|58.210.70.72|35282|   22|30302|PA |02:00:11.301|1800.65|
58.210.70.72| 72.24.147.6|   22|35282|30287|PA |02:00:11.843|1800.10|
 72.24.147.6|58.210.70.72|35282|   22|31998|PA |02:26:27.028|1800.90|
58.210.70.72| 72.24.147.6|   22|35282|31999|PA |02:26:27.028|1800.90|
 72.24.147.6|58.210.70.72|35282|   22|32764|PA |02:56:28.040|1800.88|
```


Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What Is This #7

sIP	dIP	sPort	dPort	pkt	flags	sTime	dur
72.24.144.17	10.25.235.38	40395	80	45	S PA	1:59:34.81	1759.18
10.25.235.38	72.24.144.17	80	40395	44	S PA	1:59:34.81	1759.07
10.25.235.38	72.24.144.17	80	40395	40	PA	2:29:39.82	1797.62
72.24.144.17	10.25.235.38	40395	80	40	A	2:29:39.93	1797.51
10.25.235.38	72.24.144.17	80	40395	40	PA	3:00:23.46	1800.17
72.24.144.17	10.25.235.38	40395	80	40	A	3:00:23.57	1800.17
10.25.235.38	72.24.144.17	80	40395	40	PA	3:31:09.83	1797.52
72.24.144.17	10.25.235.38	40395	80	40	A	3:31:09.93	1797.52
10.25.235.38	72.24.144.17	80	40395	40	PA	4:01:53.42	1797.72
72.24.144.17	10.25.235.38	40395	80	40	A	4:01:53.51	1797.64
10.25.235.38	72.24.144.17	80	40395	35	RPA	4:32:37.18	1560.50
72.24.144.17	10.25.235.38	40395	80	34	A	4:32:37.29	1520.89
72.24.144.17	37.52.53.241	40395	80	13	FS PA	5:18:41.57	0.48
37.52.53.241	72.24.144.17	80	40395	18	FS PA	5:18:41.63	0.43
72.24.144.17	42.15.190.19	40395	80	9	FS PA	8:21:01.15	4.14
42.15.190.19	72.24.144.17	80	40395	6	FS PA	8:21:01.15	4.14
42.15.190.19	72.24.144.17	80	40395	1	A	8:21:05.29	0.00
72.24.144.17	10.46.227.72	40395	80	7	FS PA	9:21:24.36	0.22
10.46.227.72	72.24.144.17	80	40395	6	FS PA	9:21:24.47	0.22
72.24.144.17	18.113.57.14	40395	80	6	FS PA	9:39:43.67	0.11
18.113.57.14	72.24.144.17	80	40395	4	FS PA	9:39:43.67	0.21

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

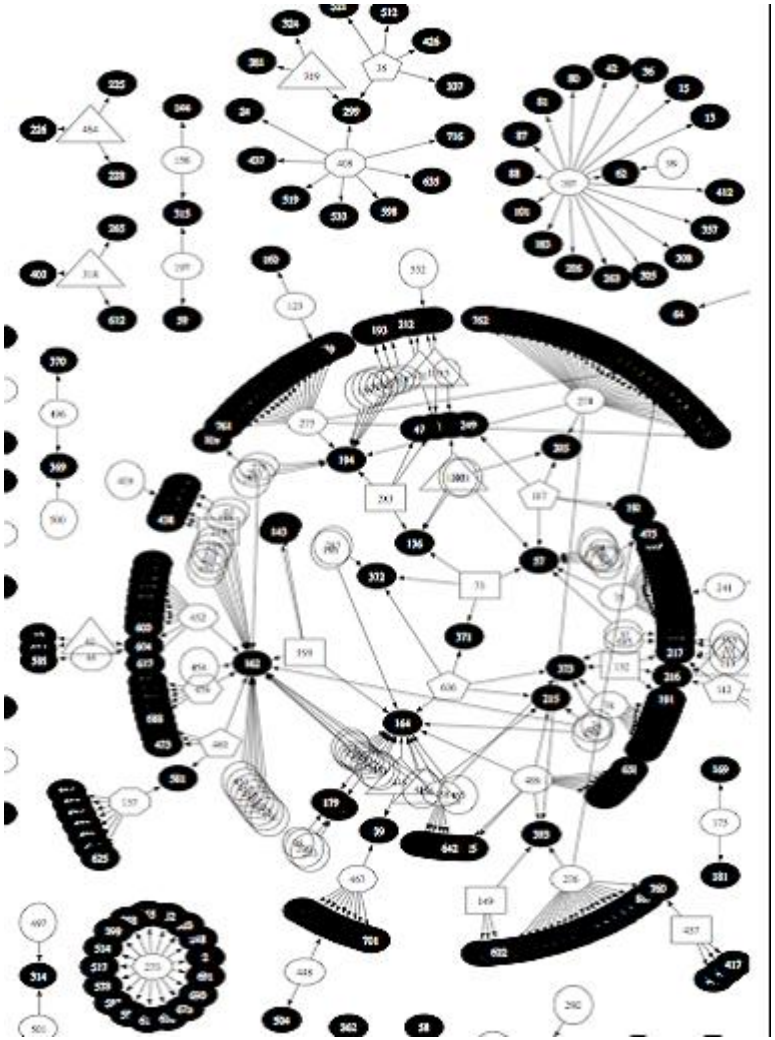
CERT

Where do I collect flows?

Flow is often collected at the border

- Watch internal and external communications
- Identify services on your network
- Identify resources your machines use regularly

Most routers can generate flows



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Flow vs. IDS

IDS

+ Content inspection
  - Presents an interpretation of raw data
  - Tuning means discarding false positive data

Flow

- No content available
+ Gives direct observations
+ No tuning, keep everything

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Flow vs. Firewall

Firewalls

+ Block unwanted traffic
  - Not intended as a historial record; logging is secondary

Flow

- Completely passive
+ Logging is primary
+ Audits the firewall

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Got a question? Flow can help.

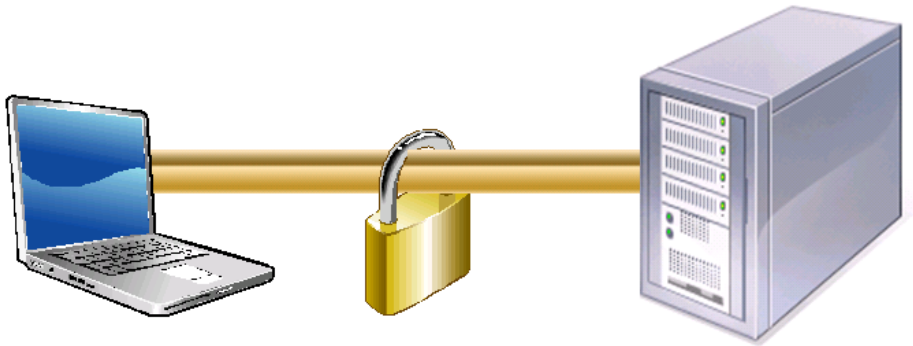
- What's on my network?
- What happened before the event?
- Where are policy violations occurring?
- What are the most popular web sites?
- How much volume would be reduced with a blacklist?
- Do my users browse to known infected web servers?
- Do I have a spammer on my network?
- When did my web server stop responding to queries?
- Who uses my public DNS server?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

About ssh

- ssh creates a secured connection between your computer and the ssh server
- ssh is your primary tool for moving things between you and your analysis server



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #1!

Login
Check for access to data (`ls /data`)
Type "`which rwfilter`"
Type "`rwfilter --help | more`"
Logout (optional)

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #2!

Move a file from the server to your workstation:

```
scp server:/remote/path/to/file.ext
/local/directory/
```

Move a file from your workstation to the server:

```
scp /local/file.ext server:/remote/directory/
```


Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Text Editors

Write only: `echo "blah" > file`
Simple: `vi`
Flexible but not always available: `emacs`
Other simple text file tools

- `cat`: print it out
- `more`, `less`: print it out one page at a time
- `head`, `tail`: print out just the beginning (or end)
- `wc -l`: count the number of lines

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Getting around

Some other commands you may need:

- `cd`: change directory
- `ls`: list the current directory contents
- `mkdir`: make a directory
- `rm`: remove a file
- `cp`: copy a file
- `logout` or `exit`: log out

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #3!

1. Create a sample file on the server.
2. Move the file from the server to your local machine, and open it in the local text editor. Change the file and move it back to the server.
3. Use head and tail to display the second line of a file which contains 5 lines.

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Where's the GUI command prompt?

- It's not quite the same as Windows:
- You'll always be working from a command prompt
- We'll be doing lots of text manipulation
- There's occasional CR-LF messiness
- Data can get big, but that's usually OK

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Get used to using pipes

- Pass output from one command as input to another
- Stop things with ctrl+c
- Also watch out for ctrl+s (suspend), restart with ctrl+q
- Also watch out for ctrl+z (put in background), continue with `fg`

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

About SiLK

- The System for internet Level Knowledge
- `http://tools.netsa.cert.org`
- Packing System
  - Accepts Netflow
  - Stores data in a very space-efficient binary flat file

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Analysis Suite

- Used to query binary flat files from the packing system
- Some mirror Unix text tools for operate on binary flow files e.g., cut, uniq, sort, split
- Some work with large IP data collections sets, bags and prefix maps
- All support ad-hoc analysis needs

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What SiLK Does

- Optimized for extremely large data collections
- Very compact record format
- Large amount of history can stay on line
- Command line interface
- Keep data in the native binary format as long as possible
- Retrospective analysis
- Most useful for analyzing past network events
- May feed an automated report generator
- Good for forensics (what happened before the incident?)

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Flavoring your Flows

Without content data, flows often seem very bland
SiLK flavors flow data with add-ons:

- Address sets e.g., blacklists
- Address Bags give a value to an address
- Prefix Maps give an arbitrary label to a group of addresses e.g., Country Code Mapping
- Hooks for custom libraries

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

About Classes and Types

SiLK assigns each flow record a CLASS and a TYPE

Class

- Duplicates the purpose of the router
- Sample classes might be Border, Internal, Customer
- We will simply use "All"

Type

- Separate inbound from outbound
- Queries often run against a single type to improve performance
- Other types are common also
- in, inweb, out, outweb, null

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

The Flow Repository

The Repository: a directory structure holding binary flow files

Directory structure based on:

- Sensor
- Type
- Year
- Month
- Day

File name based on:

- Type
- Sensor
- YYYYMMDD.HH

All times are GMT

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #4!

We're using anonymized flow in the repository at /data. SSH in to the server and determine:

1. Which dates is data available for?
2. What classes and types of data are available?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

The Training Repository

Based on LBNL Anonymized data set
`http://www.icir.org/enterprise-tracing/Overview.html`
Sensor name and date/time locates data within the repository

- S0 – anonymized general flows
- S1 – anonymized scanning flows, different anonymization
- Selected dates and times in 2004 and 2005
- Avaliable data types:
  - out, outweb: source internal, destination not internal
  - in, inweb: source not internal, destination internal

Timeouts

- 1800s (30 min) active timeout
- 60s inactive timeout

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Dates in Sample Data

- 2004/10/04:20-22
- 2004/12/15:08-23
- 2004/12/16:01-06,16-23
- 2004/12/17:00-03
- 2005/01/06:19-23
- 2005/01/06:00-06,10-23
- 2005/01/08:00-05

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #5!

We've glossed over the nuance of how SiLK handles ICMP flows. Type in the following command and look at the output:

```
rwfilter --type=in --start-date=2004/10/04 \
--protocol=1 --max-pass-records=10 \
--pass-destination=stdout \
| rwcut --fields=sip,sport,dip,dport,icmptypecode
```

1. How does SiLK store ICMP type and code information?
2. What did this command actually do?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What have we done so far?

- An Introduction to Flow
- A Brief Discussion of Unix
- A Flow Analysis Teaser

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Section Outline

Basic SiLK Tools

- rwfilter
- Printing and Sorting Tools
- Counting Tools
- Other Tools

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

So much to do, so little time...

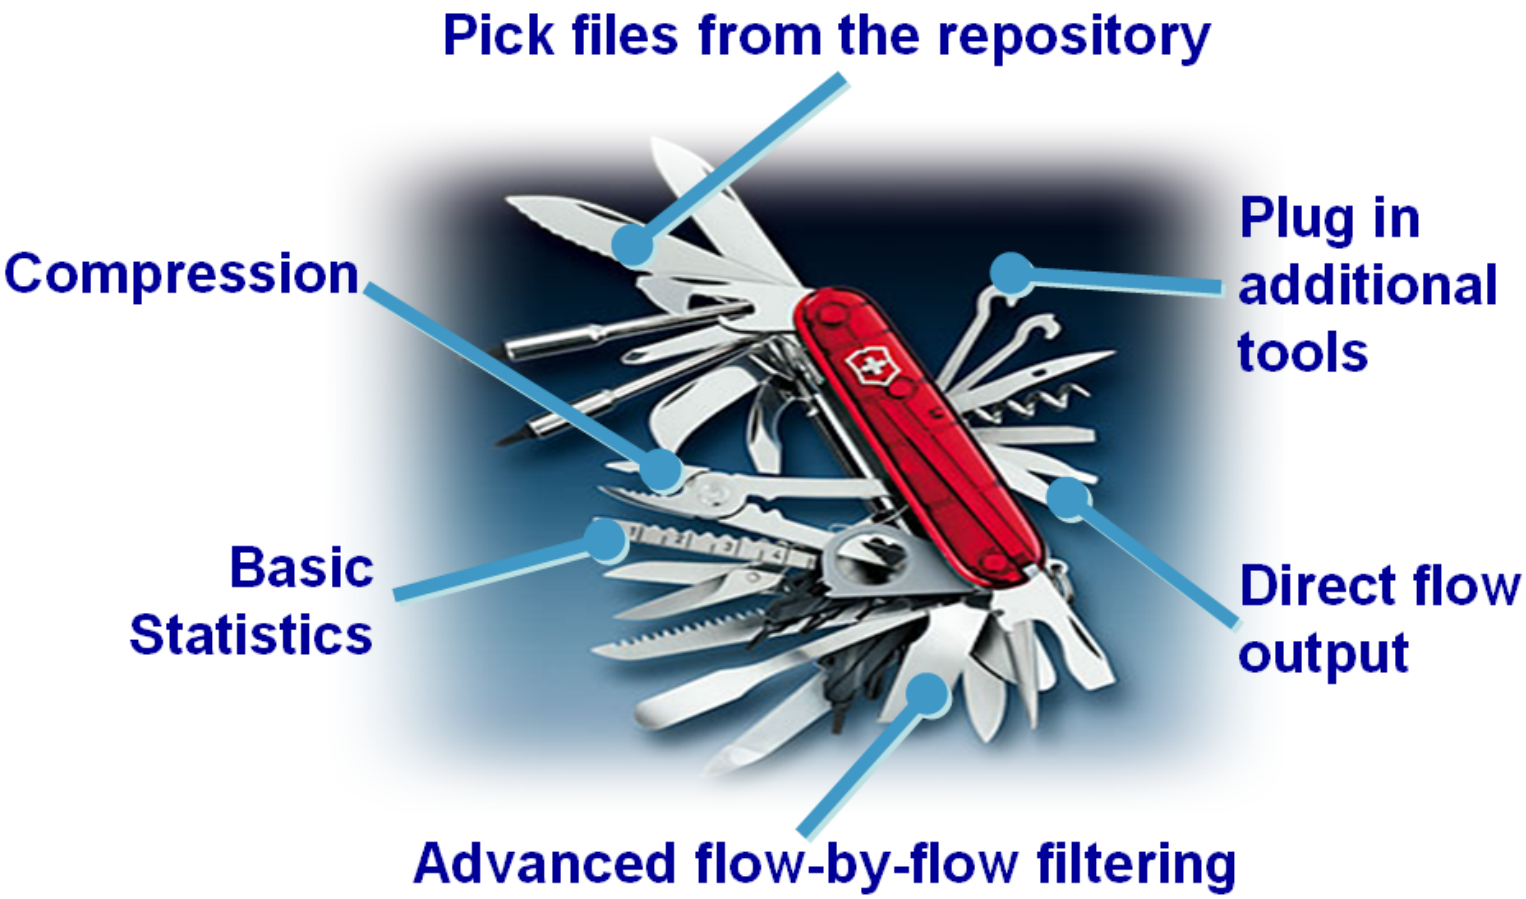
We can't discuss all parameters for every tool Resources

- Analyst's handbook
- SiLK Reference Guide (hardcopy man pages)
- `rw[something] --help`
- `man rw[something]`
- `http://tools.netsa.cert.org`

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

rwfilter



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

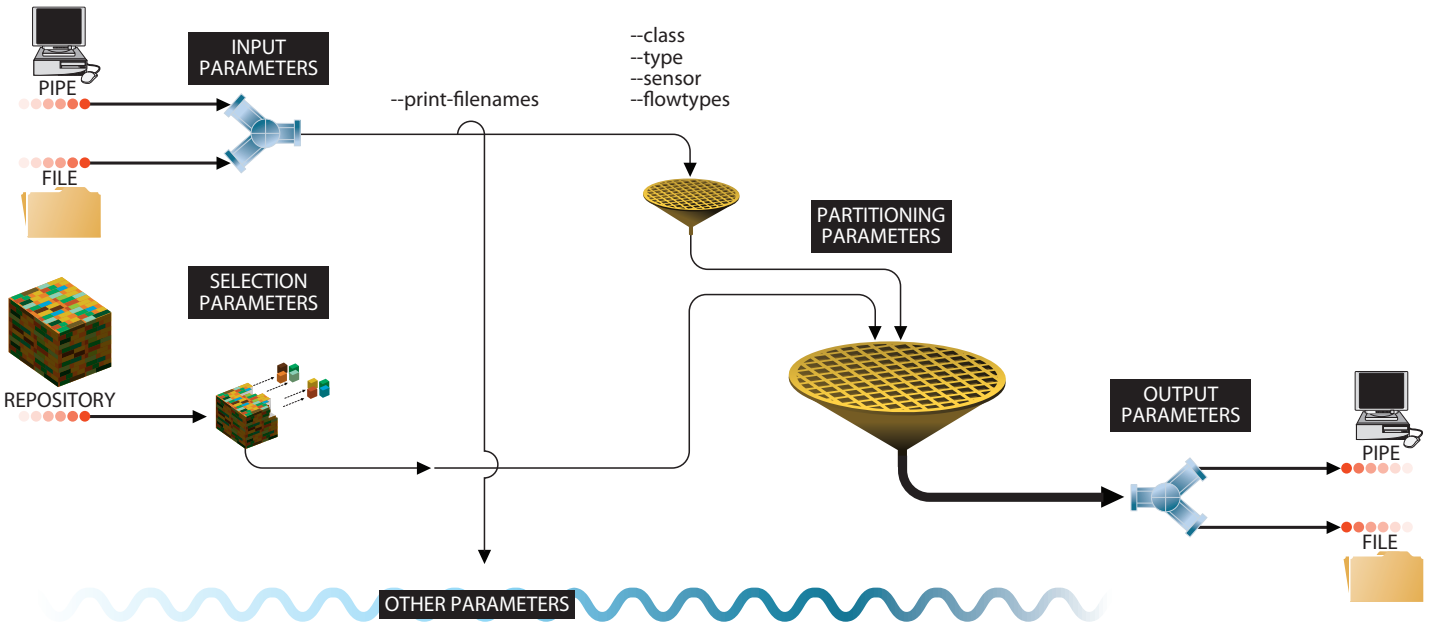
rwfilter Command Structure

- Most of the time: any order of parameters
- Parameters may be abbreviated to unique prefix
- Five different groups of parameters:
  - Input – file, repository, pipe
  - Selection – which part of repositiory
  - Partitioning – which flows among the selected
  - Output – going where (pipe, file)
  - Other – IP version, filter statistics, etc.

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

rwfilter Command Flow



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

rwfilter Requirements

Each rwfilter call must have:

- Somewhere to get records from:
  - File name
  - `--input-pipe=stdin` or other pipe
  - Repository (default or `--data-rootdir=./myarchive`) with selection parameters (type, sensor, start-date, end-date, class)
- Some description of what records are wanted (partitioning parameters)
- Some description of where records should go:
  - `--pass=myfile.rw`
  - `--fail=stdout`
  - `--print-statistics`

```
rwfilter --start-date=2008/12/05:00 \
--end-date=2008/12/05:03 --type=all \
--protocol=6 --packets=1-3 --pass=dec05.rw
```

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Selection Parameters

These options control access to repository files

- `--start-date=2007/10/03:00`
- `--end-date=2007/10/03T03`
- `--sensor=S0`
- `--class=all`
- `--type=in`

Alternatively, use a pipe or a file

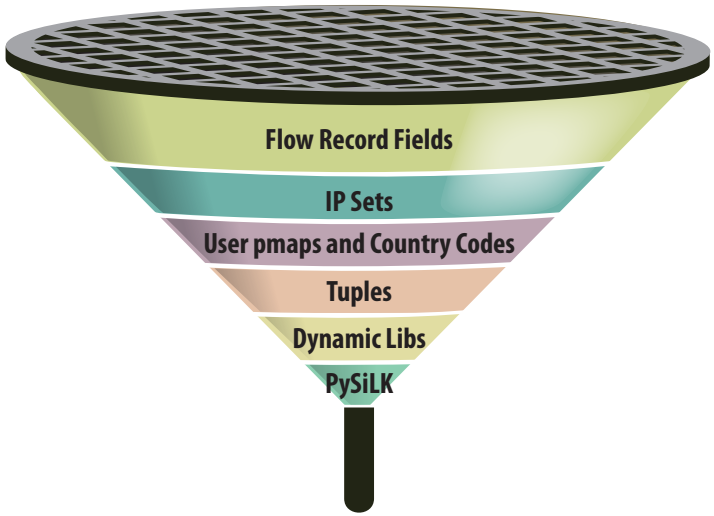
- `--input-pipe=stdin` – Useful for chaining filters through stdin/stdout
- `myfile.rw` – Useful for filtering previous results

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Partitioning Parameters

Partioning is the most complex



- Partitioning parameters form an “and” expression
- Too few parameters means too much output
- Can refine partitioning with another rwfilter call
- Some of these are beyond the scope of this course

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Flow partitioning parameters: Record Fields

Pass records based on flow record fields; one is required

- `--[not-]saddress, --[not-]daddress`: Wildcard like 12.5,7,9.2-250.x
- `--protocol`: IP protocol
- `--sport, --dport, --aport` TCP, UDP ports (caveat: ICMP)
- `--tcp-flags=SF; --flags-all=S/SA; --fin-flag;...`
- `--icmp-type; --icmp-code`
- `--bytes, --packets, --bytes-per-packet`

At least one partitioning parameter is required

- Use `--proto=0-` to pass all

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Flow partitioning parameters: Flow Record Time

start-date, end-date choose repository files, but do not look at the actual flow records

- `--stime`, `--etime`: choose flows that start (or end) within a time range
- `--active-time`: flows active in a time range
- Time format: YYYY/MM/DD:HH:MM:SS
- Time range format: [Time]-[Time]

Duration

- `--duration=1-10`: number of seconds the flow was active

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Flow partitioning Parameters: Flags

- `--tcp-flags=[FSRPAUEC]`
- `--fin-flag`, `--syn-flag`, etc.
- `--flags-all=[FSRPAUEC]/[FSRPAUEC]`
- `--flags-initial=[FSRPAUEC]/[FSRPAUEC]`
- `--flags-session=[FSRPAUEC]/[FSRPAUEC]`

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Flow partitioning Parameters: Advanced

Some of these will be discussed later:

- `--max-pass`: limit the number of records passed
- `--sipset`, `--dipset`, etc: limit to specfici IP addresses
- `--ipport`: IP/port pairs
- `--pmap`; prefix map
- `--dynamic-library`: dynamically loaded library
- `--scc`, `--dcc`: country codes
- compression

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Output Parameters

rwfilter leaves the flows in binary(compact) form

- `--pass`, `--fail`: direct the flows to a file or pipe
- `--all`: destination for everything pulled from the repository
- One output is required but more than one can be used

Other useful output:

- `--print-statistics`
- `--print-volume-statistics`

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Other Parameters

- `--dry-run`: test the command (useful for scripting)
- `--ipversion=6`: process IPv6 data
- `--print-filenames`: print files from which flow records came
- `--help`: print condensed help text
- `--man`: print manual page
- `--version`: print configuration info
- `--threads`: parallelize rwfilter run

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #6!

The time to run an initial query against the repository often depends on the number of files which will be accessed. How many files in the repository will be opened with this command?

```
rwfilter --sensor=s0 \
  --start-date=2004/12/15:19
```

(note: you have to add extra parameters to this command to make it work)

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #7!

Often you will want to track an individual address or address block. Develop a filter command to retrieve:

- Flows to the 131.243.10.0/24 CIDR block,
- Leaving our network,
- On 12/16/2004 at 17:00 GMT,
- And save the flows in the file `netblock.rw`.
  How many packets, bytes and flows were retrieved?

How many packets, bytes, and flows were retrieved?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #8!

Let's look for short, bursty outbound ssh traffic. Develop a filter command that does the following:

- Pulls out all outbound ssh (TCP port 22) flows,
- On 12/17/2004,
- Between 00:00 and 04:00 GMT,
- That lasted less than 60 seconds,
- With an average of more than 60 bytes per packet,
- And store the result in a file named short-ssh.raw

How many records did you retrieve? How many files in the repository were opened?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #9!

Examine traffic trends. What is the change in mail traffic volume between 19:00 and 20:00 hours on 12/15/2004 for the mail server at 128.3.26.249?

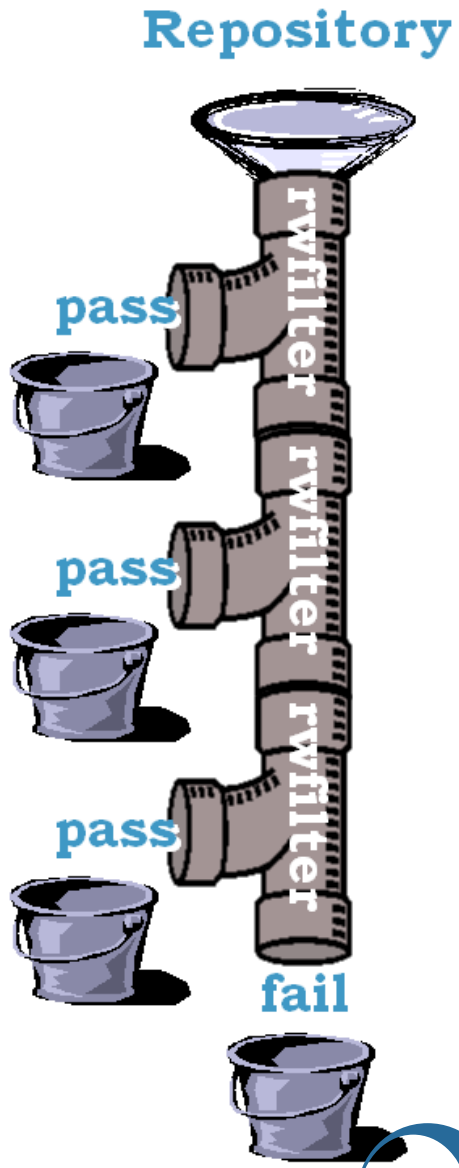
Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Chaining filters

It is often very efficient to chain rwfilter commands together

- Use `--pass` and `--fail` to segregate bins
- Use `--all` so you only pull from the repository once



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What Is This #8

```
rwfilter \
  --start-date=2007/09/30 \
  --type=outweb \
  --bytes=100000- \
  --pass=stdout \
| rwfilter \
  --input-pipe=stdin \
  --duration=60- \
  --pass=long-http.rw \
  --fail=short-http.rw
```


Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #10!

Let's revisit the last example for some more analysis. For the mail server at 128.3.26.249, and looking only at outbound traffic for the 19:00 hour on 12/15/2004, use a single command to find out both:

- The total number of SMTP flows (TCP port 25), and
- The number of flows which were for outbound messages

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Common rwfilter Typos

- `--port` or `--destport`: not an option name
- `--saddress=file`: pointing to a filename; should be an IP
- `--sip=10.1.2.3`: sip specifies an IPSet; use saddr for addresses
- `--start=2005/11/04:06:00:00` start-date and end date use only down to the hour
- `---start-date`: should be only two dashes
- `-- start=2007/05/22`: no space between `--` and the option

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

But I can't read binary...

rwcut provides a way to display binary records as human-readable ASCII

- Useful for printing flows to the screen
- Useful for input to text processing tools
- You'll usually only need the `--fields` argument

sip	packets	sval	flags	application
dip	bytes	dval	initialflags	icmptypecode
sport	Sensor	in, out	sessionflags	attributes
dport	scc	dur	dur+msec	type
protocol	dcc	stime	stime+msec	stype
class	nhip	etime	etime+msec	dtype

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Pretty Printing SiLK Output

Default output is fixed-width pipe delimited data

```
                sIP|            dIP|pro|pkts|bytes|
   207.240.215.71| 128.3.48.203|  1|    1|   60|
   207.240.215.71|  128.3.48.68|  1|    1|   60|
   207.240.215.71|  128.3.48.71|  1|    1|   60|
```

Tools with text output have these formatting options

- `--no-titles`: suppress the first row
- `--no-columns`: suppress the spaces
- `--delimited`; `--column-separator`
- `--legacy-timestamps`: better for import to Excel

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #11!

Create a file ssh.rw that contains all outbound SSH flows from 12/16/2004:17. Experiment with rwcut and Unix text tools to try and sort out records:

1. Can you tell which flows are from internal SSH servers, and which are from external SSH servers?
2. Which flows look like SSH keep-alives?
3. Which flows had the most data transfer?

Try to write rwfilter commands against ssh.rw to query these records, and display them with rwcut

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

rwsort

Why sort flow records?

- Records are recorded as received, not in time order (look at records from the last exercise)
- Analysis often requires finding outliers

rwsort options

- fields (same as rwcut) is required
- in, out (stdin / stdout are defaults)
- For improved sorts, specify a buffer size
- For large sorts, specify a temporary directory

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

I only believe what I see

You'll be tempted to work with text-based records

- It's easy to see the results and postprocess with other tools (e.g., perl)
- It takes a lot of space, and it's much, much slower

Guiding Principle: Keep flows in binary format as long as possible

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #12!

Often HTTP beaconing consists of very small HTTP requests. Let's get a feel for what HTTP data looks like, even before we start to find these beacons.
What do the smallest outbound HTTP web client flows look like on 12/15/2004?

- First, find them using rwsort
- Second, find them using sort
- Which was faster?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Counting Tools

The suite contains several counting tools:

- `rwcount` - count across time
- `rwaddrcount` - count across addresses
- `rwuniq` - count on arbitrary field combinations
- `rwstats` - descriptive statistics and counts

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

rwcount

Basic counting:

- `rwcount myfile.rw > count_file`
- Produces byte, packet and flow totals by time

Common Options:

- `--bin-size`: changes the size of each bin (in seconds)
- `--skip-zeroes`: should empty bins be printed?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

rwaddrcount

Basic counting:

- `rwaddrcount --[print-option] myfile.rw`
- `--use-dest` to work with dIP; default is sIP

Print Options:

- `--print-stat`: Lists total number of addresses found
- `--print-ips`: Just print out the IP address, nothing else
- `--print-recs`: Lists bytes, packets, records, times for each address

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

rwstats

Great for generating top-N, bottom-N lists
Group by (choose one or two):

- Addresses
- Ports
- Protocols

Output Limit

- Count
- Top, Bottom
- Threshold (specific value range)
- Percentage

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

rwuniq

The more general case for rwstats
Mirrors the unix “uniq -c” command

- Creates a giant hash table where you define the key
- Memory is expensive, so we can’t uniq everything

Common Options:

- `--fields`: same as cutting and sorting
- `--all-counts`: collect bytes, packets and flows
- `--bin-time`: size the bins when uniq’ing on time

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #13!

Find scans traffic in the sample data (while the anonymizer removed some of the simple scans, they didn't find them all). When you find one, answer the following questions:

- What type of scan was it?
- When did it start/end?
- How fast was it?
- What did the scanner discover?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #14!

This example showcases the very useful `--dip-distinct` feature of rwuniq:
For 2004/12/15, how many clients connected to the highest volume web servers?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Oops...I forgot where this came from...

rwfileinfo

- Each SiLK file (flows, sets, bags, prefix maps, etc.) has a header which logs data
- rwfileinfo prints out that data
- For flow files, it also (usually) keeps a history of the commands used to generate the file

Try It!

- rwfileinfo *.rw

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

When the files are LARGE

rwcat

- Send a binary rw file to stdout

rwappend

- Join multiple files together

rwsplit

- Carve large files into pieces

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

How long will this really take?

rwfglob

- Find out which files will be pulled from the repository
- Find out whats available and whats missing
- Use the output in other file-processing scripts

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Maintain anonymity

rwnetmask

- Mask off low order bits of source and/or destination addresses

rwrandomizeip

- Randomly replace source and destination addresses

rwtuc

- Change text flow data into binary (opposite of rwcut)

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Who was that?

rwresolve

- Perform a DNS lookup on text output
- Caveat: it uses your analysis host's DNS resolver
- Caveat: DNS is subject to change

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What have we done so far?

Basic SiLK Tools

- rwfilter
- Printing and Sorting Tools
- Counting Tools
- Other Tools

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Section Outline

Advanced SiLK Tools

- Sets
- Bags
- Prefix Maps

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Blacklists, Whitelists, Books of Lists...

Too many addresses for the command line?

- Spam block list
- Malicious web sites
- Arbitrary list of any type of addresses

Create an IP set!

- Individual IP address in dotted decimal or integer
- CIDR blocks, 192.168.0.0/16
- Wildcards, 10.4,6.x.2-254

Use it directly within your filter commands

- `--sip`, `--dip`, `--anyset`

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Set Tools

`rwset`: create sets from binary flows
`rwsetbuild`: create sets from text
`rwsetcat`: print out an IP set into text (**very useful**)
`rwsetmember`: test if IP is in given IP sets
`rwsettool`: perform set algebra (set, union, intersection) on multiple IP sets

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #15!

Flow is also very useful for creating network inventories. What /24 net blocks are populated within my network? Which block has the densest population?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Other uses of IP Sets

Perform set arithmetic on IP data

- What addresses on my spam blacklist are also bot infected?

Randomly select items for sampling

- `rwsettool --sample --size=100`

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Bags: sets with attitude

Bags are generally IPSets with an associated integer

- Usually a count or sum
- Could also be ports or protocols

Bags can make sets

Math operations can be performed on bags

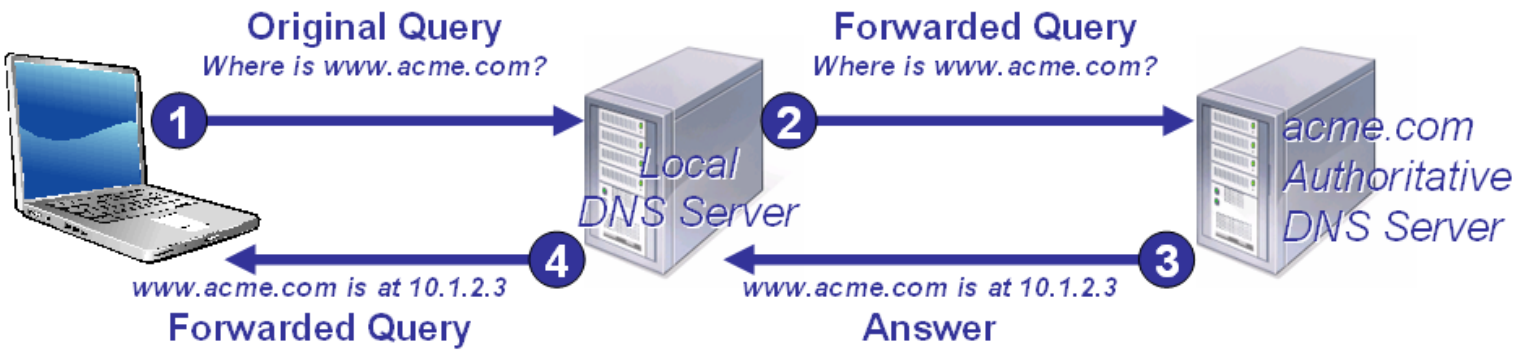
Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #16!

Let's look for DNS clients that are using an external DNS resolver.

- First, let's take a moment to review DNS:
- When a client wants an address, it asks its local DNS server
- The local DNS server does all the work



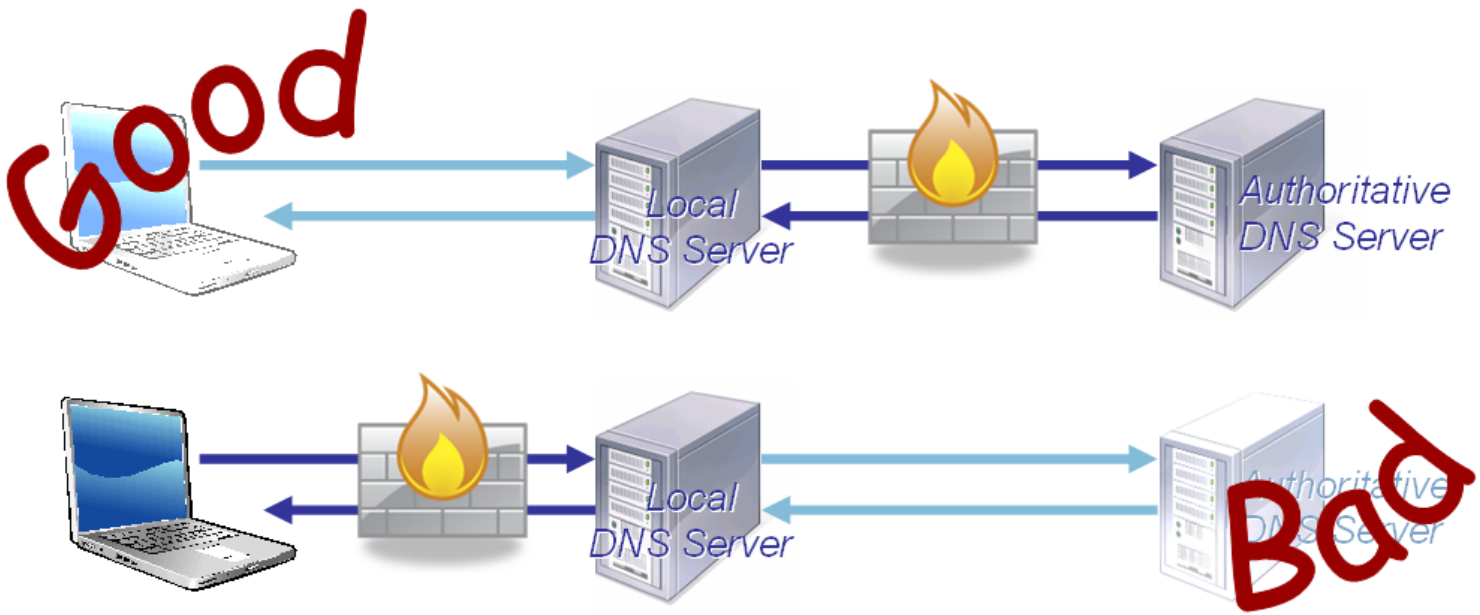
Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #16!(2)

The Local DNS Server should be local

- Can be assigned manually or by DHCP
- Up to three can be assigned, but often only one is used



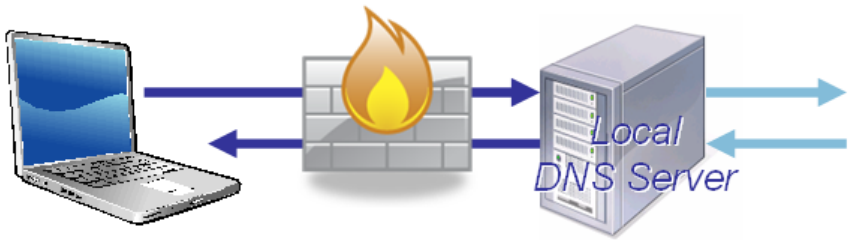
Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #16!(3)

Once again, let's look for DNS clients that are using an external DNS resolver:

- Use bags to count the number of outbound DNS connections per address,
- Create a candidate set from that bag of addresses with more than 100 outbound flows, and
- Find the number of unique destination addresses for the candidates.



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Prefix Maps(pmaps): sets with bling

Assign an arbitrary label to address prefixes

- Start with a text file of IP ranges and labels
- Order from least to most specific
- Compile the text file with rwpmapbuild
- Print out the pmap with rwpmapcat

The input file:

```
10.0.0.0/8                Private Unassigned
192.168.0.0/16            Private Unassigned
172.16.0.0/12             Private Unassigned
10.0.1.100 10.0.1.200 Workstation DHCP
10.0.1.1   10.0.1.50  Servers
10.0.2.1   10.0.2.50  Servers
10.0.3.1   10.0.3.50  DMZ Servers
```

No other pmap tools (?!?)

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Using pmaps

pmaps don't have their own tools, they fit in with existing tools

- rwfilter
- rwsort
- rwcut
- rwuniq

This allows you to add your own fields to flow

- Query all your servers:
  `rwfilter --sval="Servers","DMZ Servers"`

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Port-based pmaps

It's also possible to create prefix maps based on ports.

- Useful for well-known service ports; e.g., IRC, HTTP
- Also useful for ICMP

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #17!

Create an ICMP prefix map from the ICMP types (or types and codes).

- Look at unassigned ICMP type/code values that are in use. Which ICMP types receive the most traffic?
- Note: ICMP type/code values are assigned by RFC792; a summary table is available from IANA at `http://www.iana.org/assignments/icmp-parameters`

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Are we there yet?

Advanced SiLK Tools

- Splitting and merging
- Sets
- Bags
- Prefix Maps

More reliance on examples to demonstrate these concepts

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Scripting

Why script?

- Repeatable analyses
- Encapsulating syntax
- Composing complex commands

How to script?

- Shell scripting (we'll use bash)
  - Good reference: `http://tldp.org/LDP/abs/html/`
- Python (beyond this class, but widely used)
  - Good reference: `http://docs.python.org/tut/`

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Getting started

Put your typed commands into a file

- From our second example, `ls -lR /data |grep "/data"`

Run the file

- `bash script.sh` or
- `sh script.sh` or
- `chmod +x script.sh; ./script.sh`

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Using variables

`name="value"` (quoting is optional)

- Use the variable by adding a "$" prefix
- Everything is case sensitive

Display values with "echo"

- e.g., `echo "Variable var is $var"`

`$1, $2, ...` `$9` are input parameters

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What Is This #9

```
cmd="rwfilter"
#  selection
cmd="$cmd --type=in,out"
cmd="$cmd --start=$1"
#  partition
cmd="$cmd --protocol=1"
#  output
cmd="$cmd --print-volume"
#
# execution
date>>logfile
echo "$cmd">>logfile
$cmd
```


Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Loops

For loops

```
for variable in list
do
  stuff
done
```

`list` is a space-delimited set of values

While loops

```
while [ condition ]
do
    stuff
done
```

Note: The square brackets and spaces are required

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Conditionals

```
if [[ condition ]]
then
  do something
elif [[ condition ]]
then
  do another thing
else
  do final thing
fi
```

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #18!

Let's do some investigation of spam activity.
Re-run the query from the spam exercise in the previous section, and save the results to a file. From the single text file, find:

- What is the worst spammer by byte count and by record count?
- Which unique blacklist entry numbers were seen?
- Did any spammers send to multiple internal mail servers?
- As you run the initial command, send it to the background, kill it, then run it again

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #19!

Trend the buildup of source addresses over the course of a day. Write a script that will:

- Accept an input date
- Create a cumulative set file without any addresses in it
- Then, for each hour in the input date:
  - Create a set of all the outbound source addresses
  - Add the outbound source addresses to the cumulative set
  - Count the addresses in the hourly and cumulative set files

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What have we done so far?

Unix Scripting

- More Unix Commands and Techniques
  - Working with processes
  - Working with text
- Automating Daily Tasks
  - Variables
  - Loops
  - Conditions

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Section Outline

Visualization

- Basic graphs
- The Excel Cookbook
- The Gnuplot Cookbook
- Advanced graphs

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

A picture is worth a thousand data points

Exploratory

- Throw something against the wall
- Is anything going on?

Explanatory

- Find the point inside the big picture
- Place the point in context
- Avoid technical obscurities

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

You've Got Style

Colors, patterns and shapes are important

- Useful for separation and contrast
- Draw attention to important features

Use multiple curves

- Compare and contrast data sets

Avoid unnecessary clutter

- Limit point count or point size
- Limit text size; avoid redundant text

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What I meant, not what I said

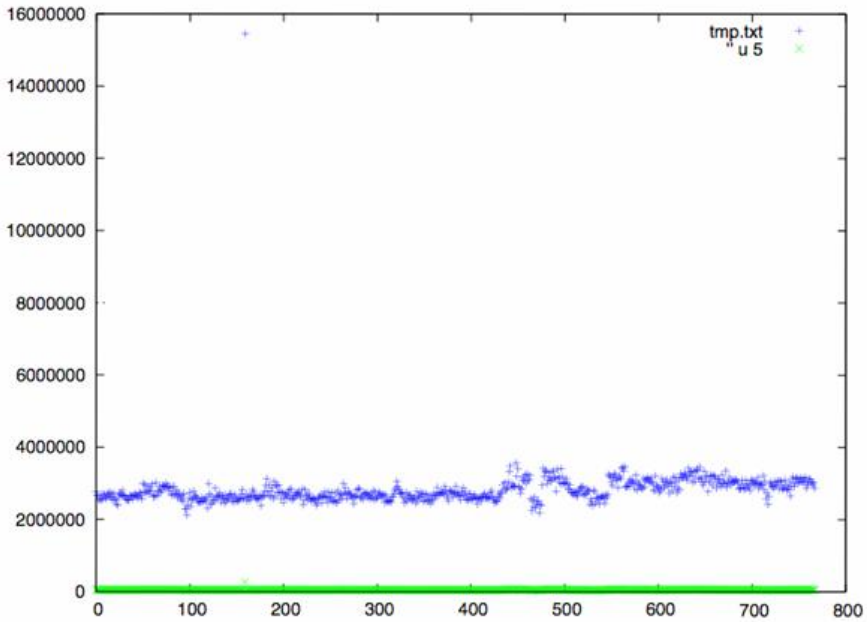
Avoid bad visualizations which

- Look overly complex
- Promote the insignificant
- Lose the context of the issue
- Contain no information or too much information
  - There's a difference between data and information

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

The Worst Plot Ever



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

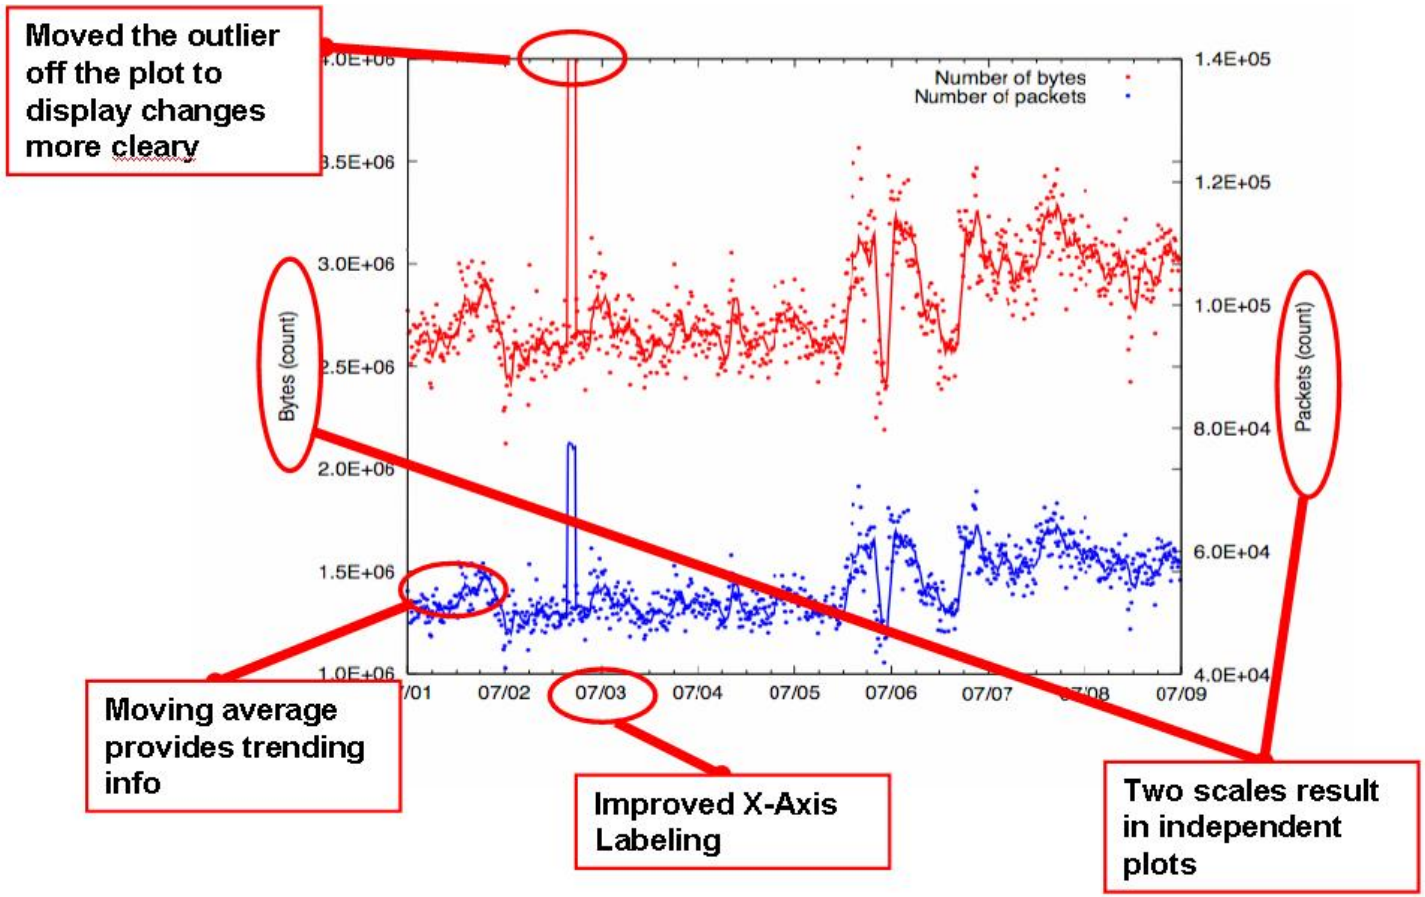
Why is this the worst plot ever?



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Same data, better plot



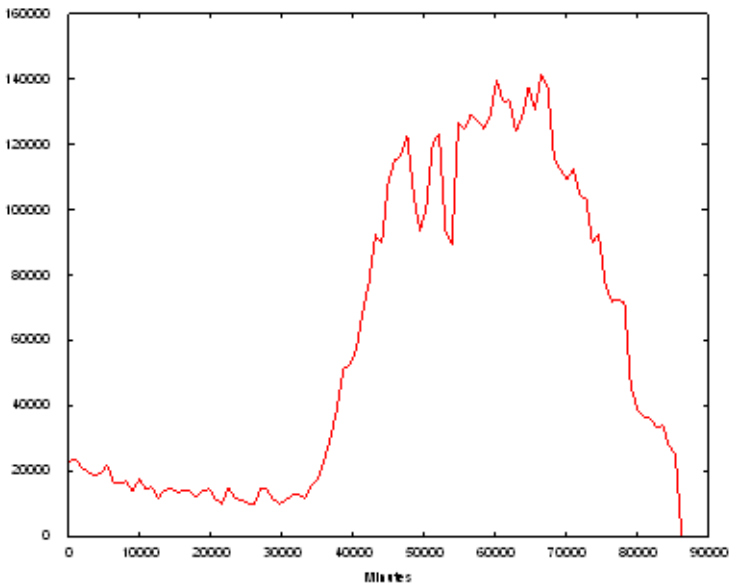
Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Getting started: two primary types

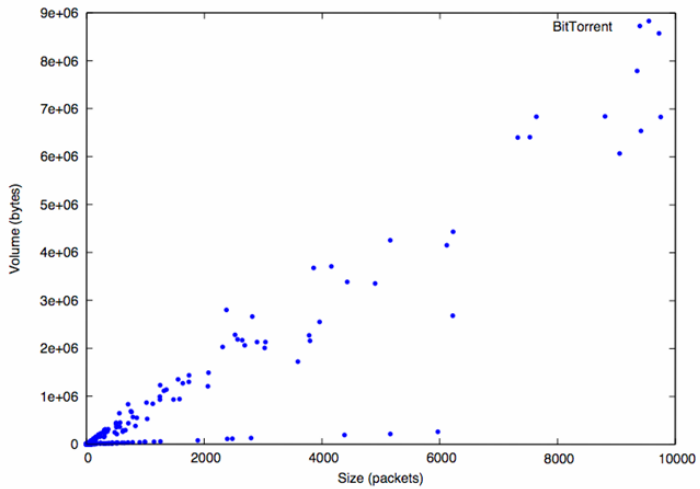
Time-series graphs

- Display how a value changed over time



Scatterplots

- Compare two distinct values



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Scale data consistently!

Byte volumes

- Always use “Megabits per Second” or “Mbits per second”
- Note that Mbps = Mbits per second; MBps = MBytes per second
- Mbps = (Bytes $*8$)/(seconds in the time bin)

Packet counts

- Always count packets per second

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #20!

Our task for the rest of this section will be to get a better understanding of high-port to high-port TCP traffic. Before we begin graphing, we need the data.

- Extract outbound TCP flows,
- for 12/15/2004,
- where both ports are greater than 1024,
- with both SYN and ACK flags set,
- with 5 or more packets,
- and save the output as "`highport.rw`"

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Getting started with Excel

Excel works on the client

- Data must be loaded into a client-side spreadsheet
- Transfer bulk data with scp
- Transfer small data sets with cut and paste

Once the data's on the client, we can graph

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Formatting data for Excel

Column Delimiters

- Excel likes comma separators, csv files
- Excel can turn pipe-delimited data into columns
  - Menu — Data — Text to Columns

Date/Time Data

- Excel likes “legacy-timestamps”

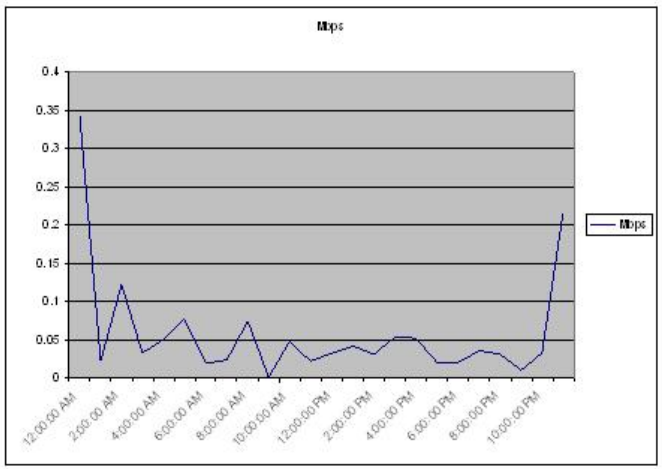
Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

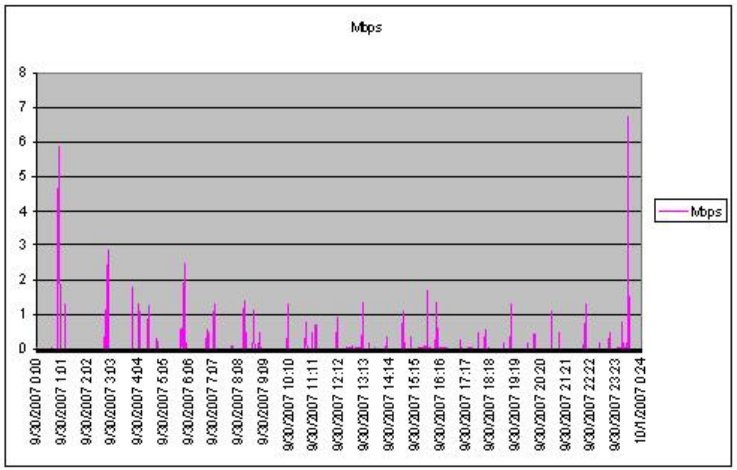
CERT

Try It #21!

Our first Excel visualization shows the volume of high-port to high-port traffic over time.

- Using rwcount, create and plot volume of highport.rw data for 1-hour and 1-minute bins.





Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

Try It #22!

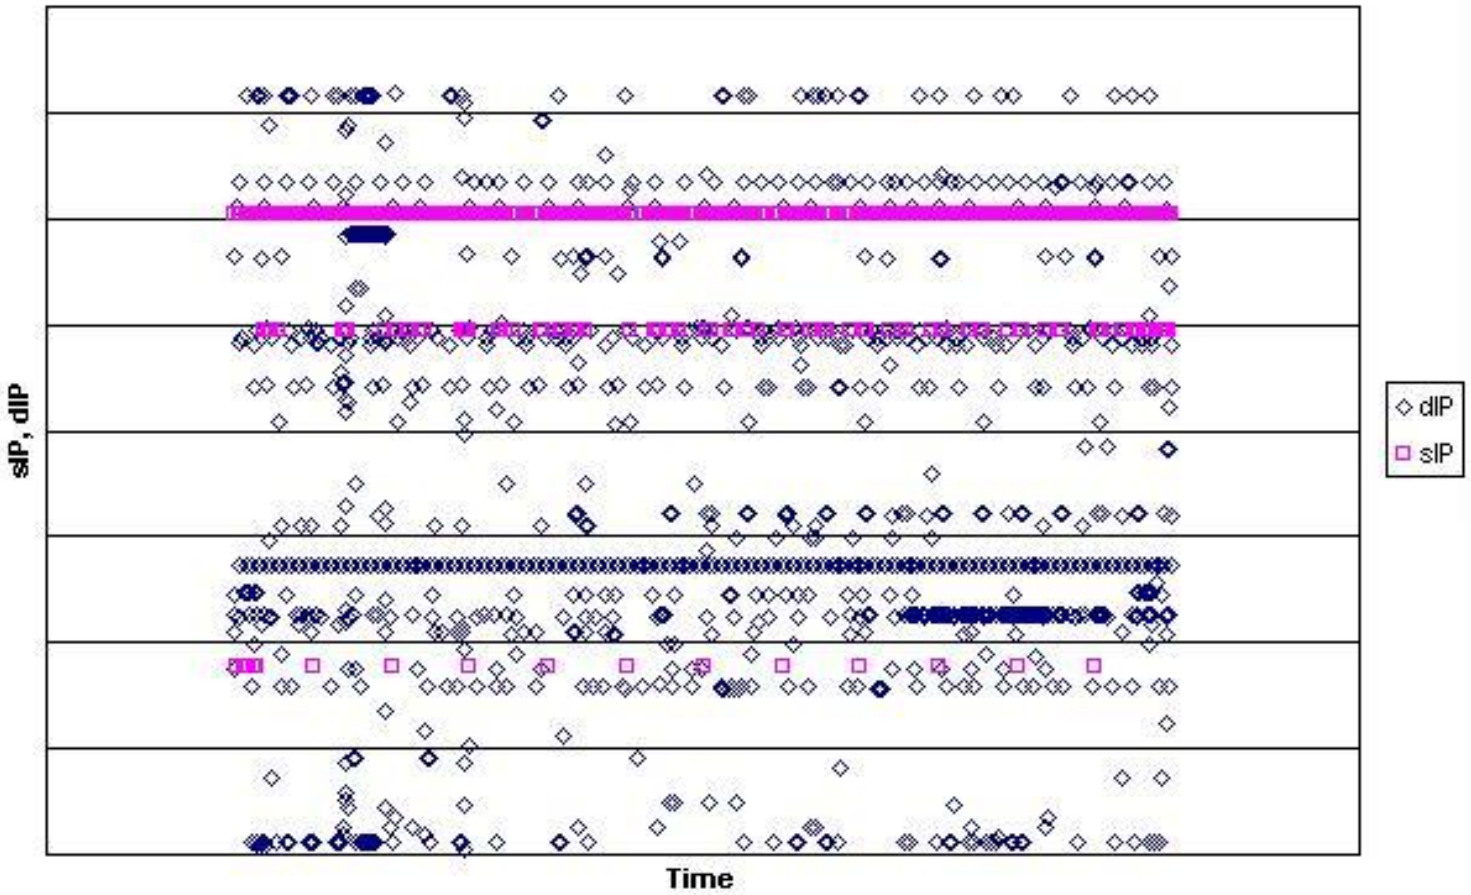
Our second excel visualization uses scatterplots to try and better understand the data

- Use rwcut to dump out the raw flow records with numeric IP, and import them into Excel
- Create one or more scatter plots using two-variable combinations such as
  - source port vs. dest port
  - source port vs. flow size and dest port vs. flow size
  - source port vs. time and dest port vs. time
  - source IP vs. dest IP
  - source IP and dest IP vs. time

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

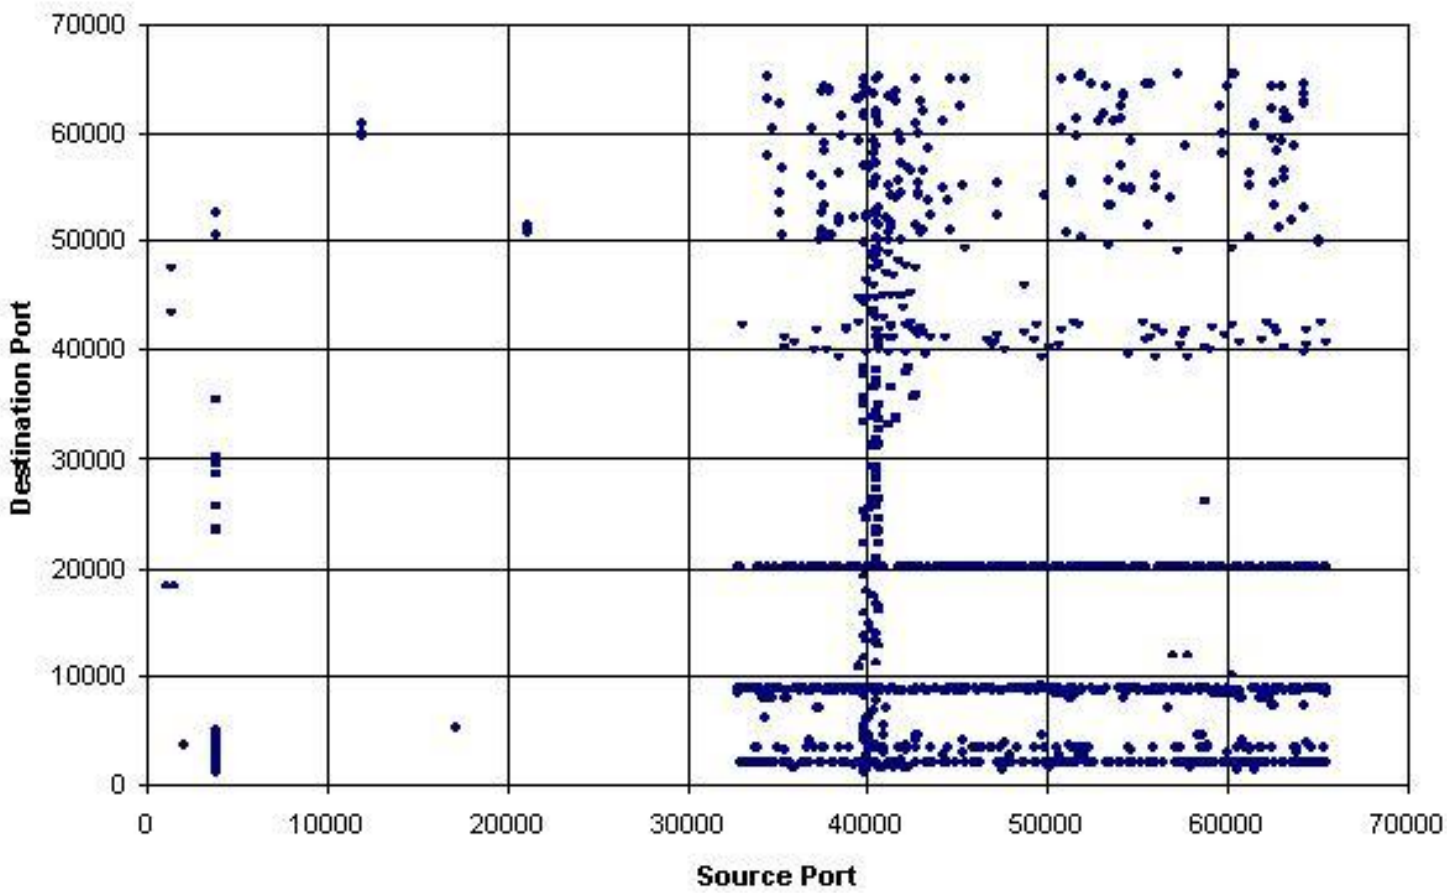
Source IP and Destination IP over time(Existence Plot)



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Source Port vs Destination Port



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Why not stick with Excel?

Excel has some concrete limitations:

- Maximum of 65k rows
- Can't easily automate
- Difficult to regularly update data for a template plot

Gnuplot makes a great complementary tool

- Fully scriptable
- Runs on the server
- Only limits are file sizes
- No GUI

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What is gnuplot

Command-line interactive data plotting utility

- Originally intended for scientists to visualize mathematical functions
- Now supports many non-interactive uses

Multiple output formats

- Interactive (X-windows)
- Direct to printers
- Many file types

We'll use it to plot text data into a postscript file

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

A minimal gnuplot script

```
#! /usr/bin/gnuplot
set terminal postscript enhanced color solid
set output "temp.ps"
set title "Hello, World"
plot sin(x)
```

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #23!

Use this script to reproduce high port traffic as in the previous example:

```
#! /usr/bin/gnuplot
set terminal postscript enhanced color solid
set output "hourly.ps"
set title "Hourly High Port Trends"
set datafile separator ","
set timefmt '%m/%d/%Y %H:%M:%S'
set xdata time
mbps(x)=x*8/3600/1000000
plot 'hourly.csv' using 1:(mbps($3)) with lines \
  title "Mbps"
```


Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Try It #24!

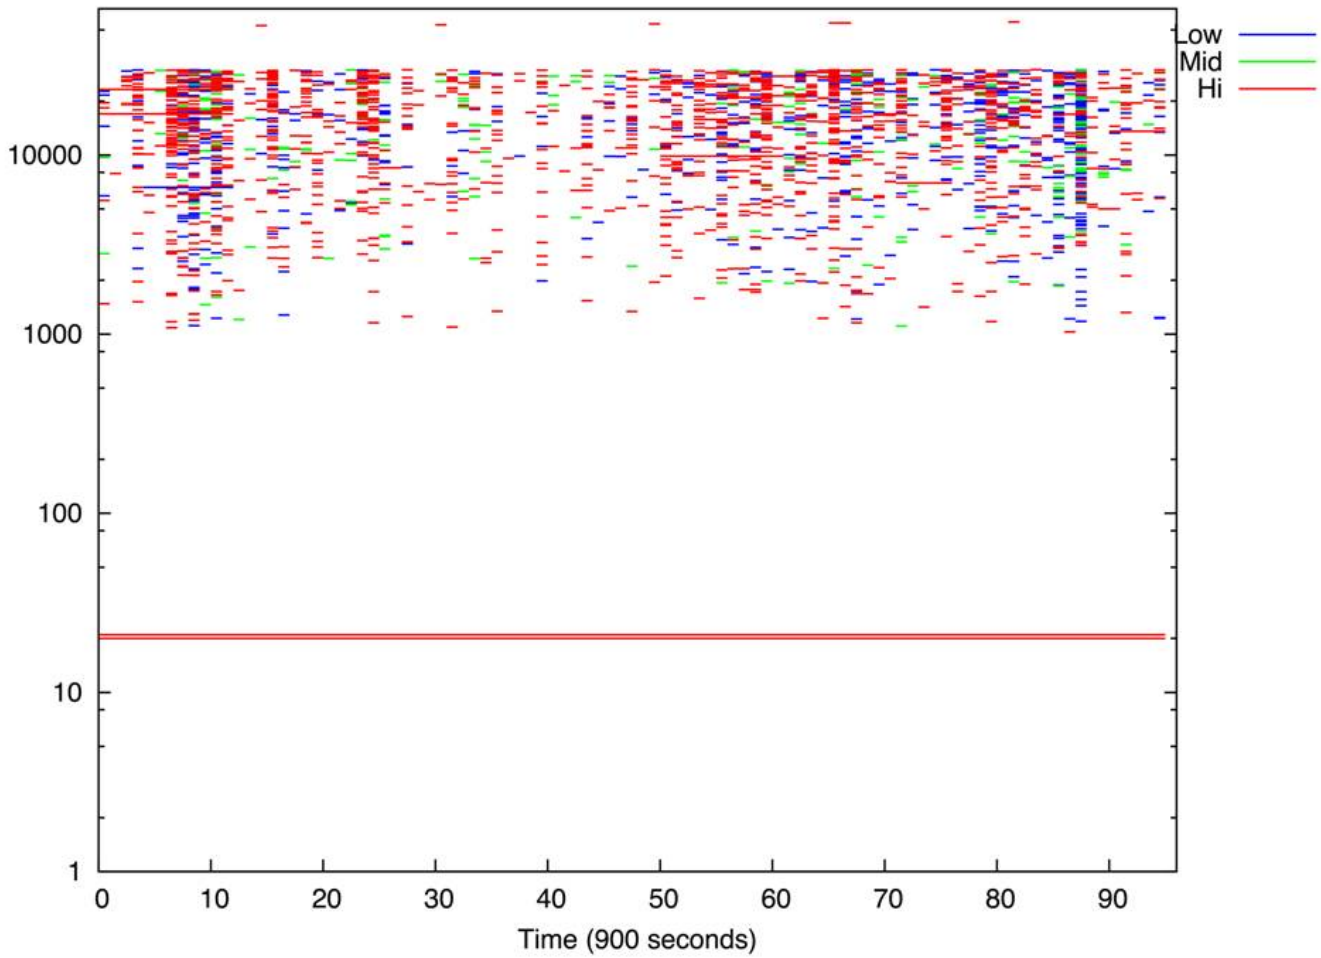
Use gnuplot to reproduce the results of Excel scatterplot exercise above. Start with the script from the previous example.

- Once you've generated some output, try plotting other column variations.

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

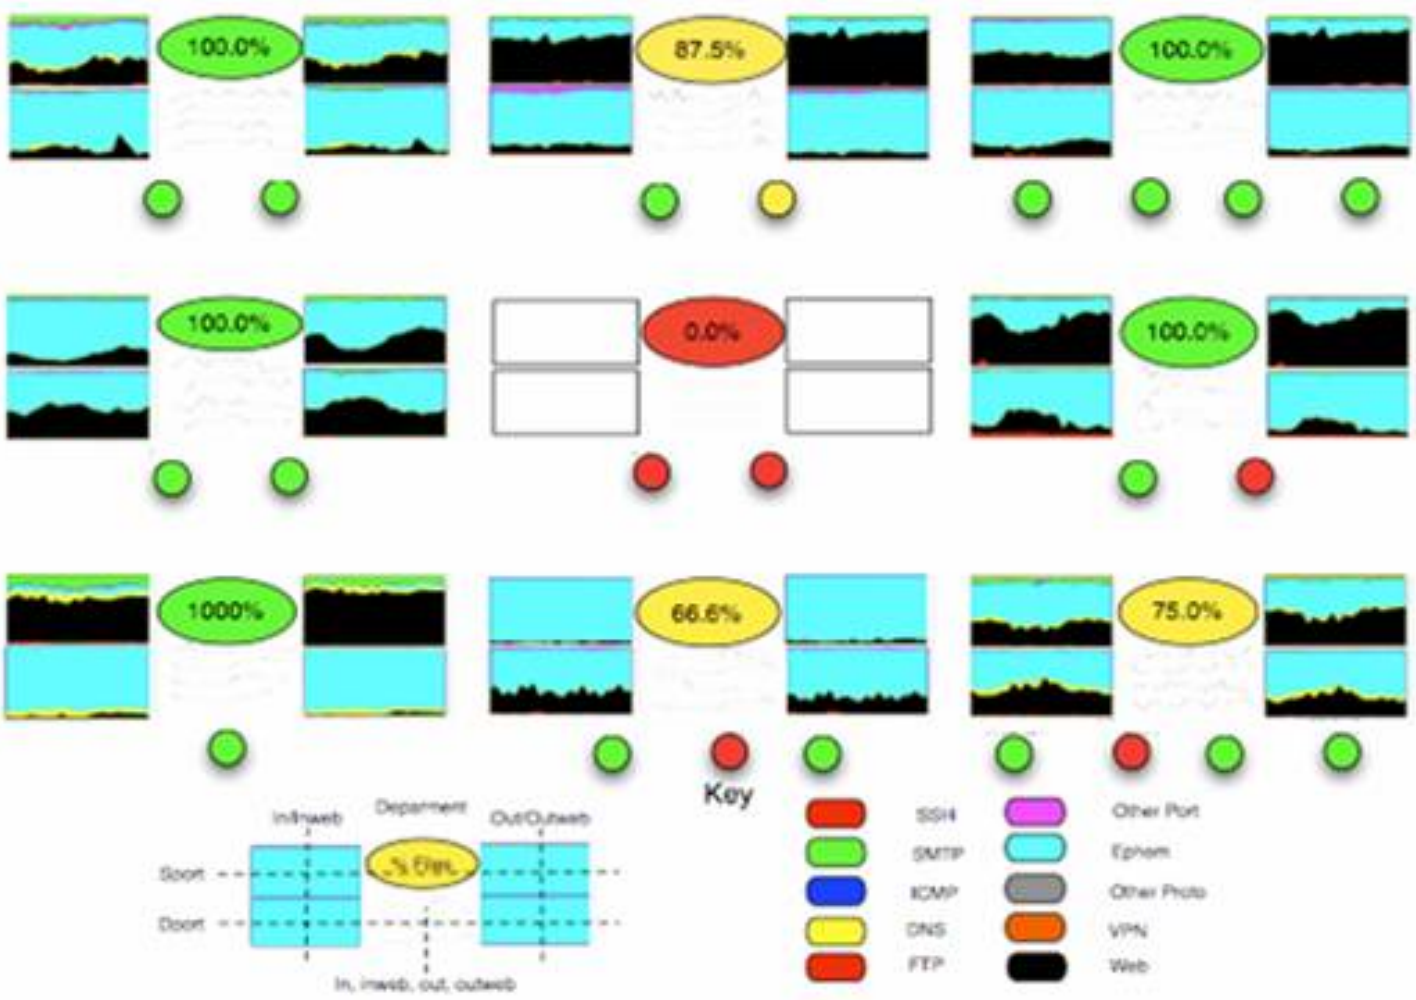
Advanced Graphs: Existence Plots



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

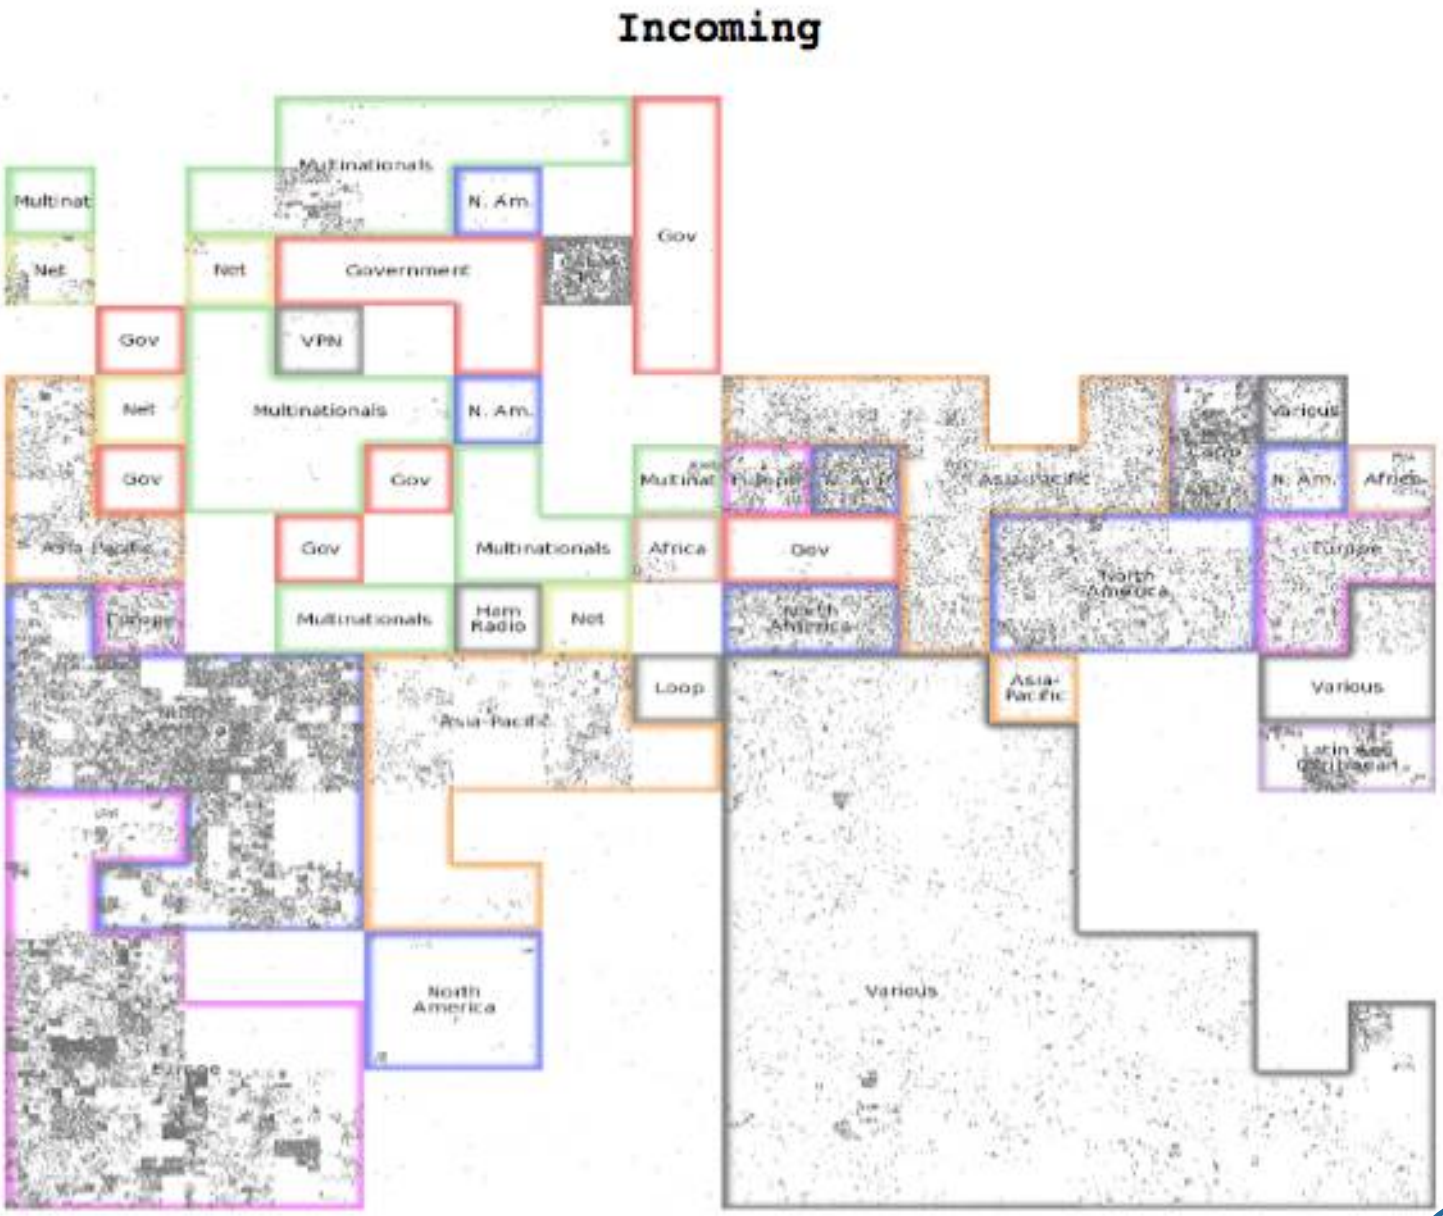
CERT

Advanced Graphs: Information Density



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

Advanced Graphs: Abstract Art



Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

What have we done so far?

Visualization

- Basic graphs
- The Excel Cookbook
- The Gnuplot Cookbook
- Advanced graphs

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Using the SiLK Tool Suite

Many tools not touched on in this training:

- Packet content: rwptoflow, rwpmatch
- Text to Flow: rwtuc
- Address aggregation: rwnetmask
- Newer and broader tools still under development

Explore the tool set for awareness you need to build of your network

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Where now?

Explore Your Data

- Rich data source

Using SiLK for Situational Awareness

- Top-N lists (rwfilter + rwstats)
- Odd behavior frequencies (rwfilter+rwcount)
- Drill-down (rwfilter +rw*)
- Build off of examples presented

Script analyses of interest

- Shell script

Create visualization templates

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

What's coming

Tools manipulating higher-level representations:

- Clusters
- Graphs
- Patterns

Expanded data collection Collaboration with other flow analysis efforts

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

SiLK Support

- `mailto:silk-help@cert.org` for bug reports and general inquiries
- FloCon `http://www.cert.org/flocon/` conference on network security analyses with flow data
- `http://tools.netsa.cert.org/` – open source release of SiLK, with documentation
- `http://www.cert.org/netsa` – papers and presentations by CERT NetSA Group

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Communications Resources

- `silk-help@cert.org`, SiLK tool suite help
- `flocommunity@cert.org`, discussions and announcements for network flow analysts
- FloCon (http://www.cert.org/flocon), Flow analysis conference

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Other References

- The analysts handbook! Great source of how-to info.
- The SiLK Reference Guide.
- Toolname help (or man toolname)
- http://tools.netsa.cert.org

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Other Useful Hints

- Use Toolname –help for help and/or analysts handbook (or me)
- Do as much analysis in the binary as possible: it's faster.
- Use sets (or tuples) whenever possible vice going thru the same data multiple times.
- When pulling together by-address, by-protocol, or by-port counts for later processing, use bags instead of rwuniq

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT

Questions?

Introduction
About
Flow
Unix
Beginning Analysis
Basic SiLK Tools
rwfilter
Printing and Sorting Tools
Counting Tools
Other Tools
Advanced
Sets
Bags
Prefix Maps
Unix Scripting
Visualization
Basic Graphs
Excel
Gnuplot
Advanced Graphs
Conclusion
The Community

CERT



DNS and Flow

Bulk DNS Analysis

Ed Stoner

© 2009 Carnegie Mellon University

DNS and Flow

NO WARRANTY

THIS MATERIAL OF CARNEGIE MELLON UNIVERSITY AND ITS SOFTWARE ENGINEERING INSTITUTE IS FURNISHED ON AN “AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this presentation is not intended in any way to infringe on the rights of the trademark holder.

This Presentation may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

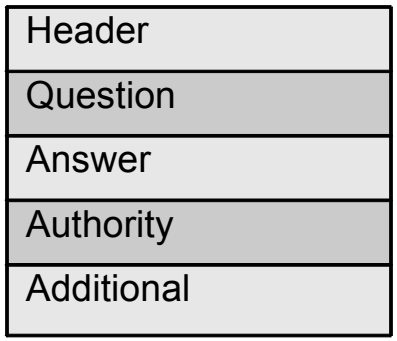
This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

Summary

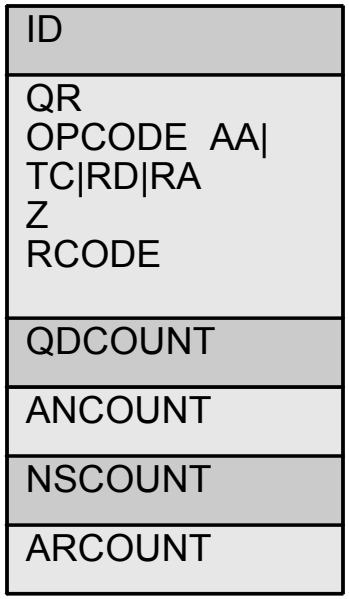
There is a wealth of information in DNS traffic that can add another dimension to flow analysis. We will explore different techniques to analyze DNS traffic and combine that analysis with flow analysis.

DNS packet format

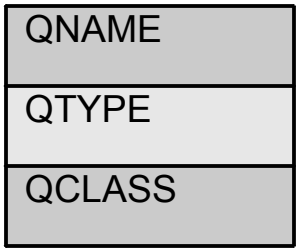
Message Format



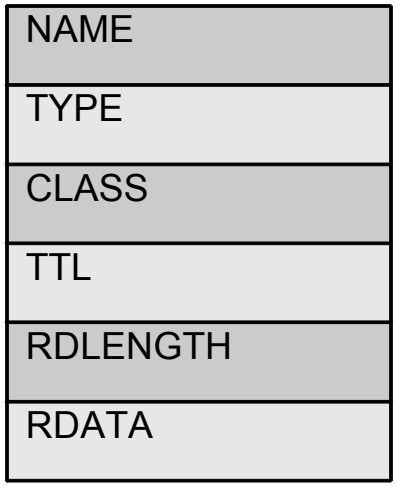
Header



Question



Answer, Authority, and Additional



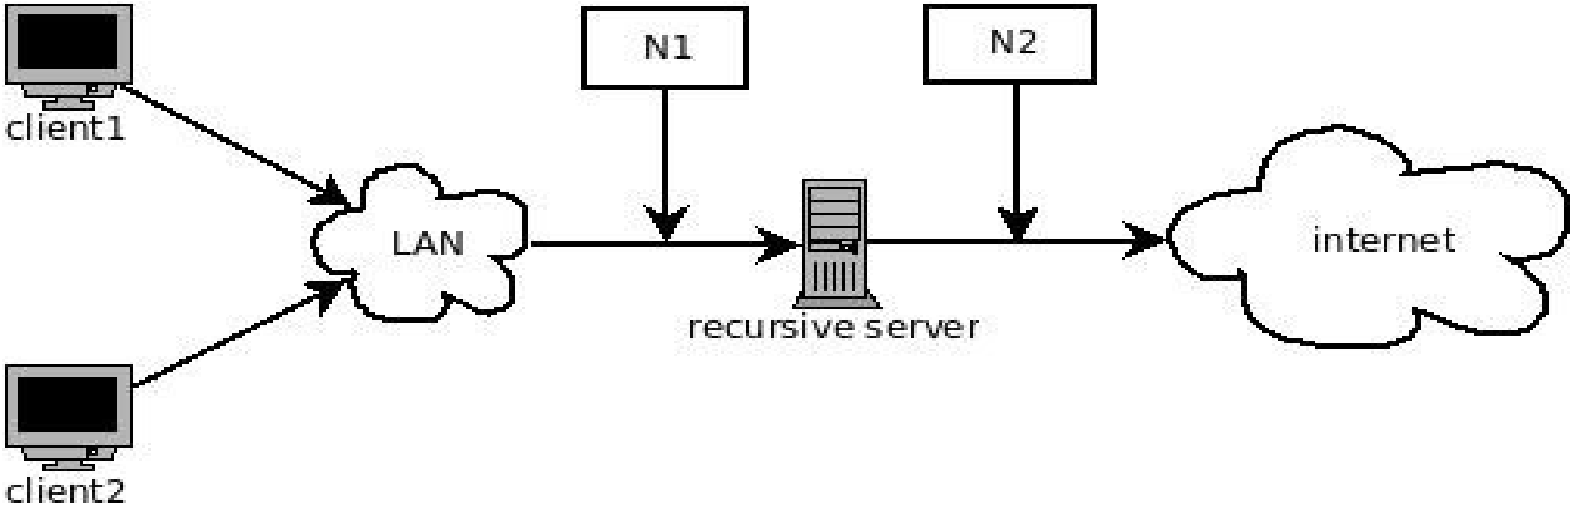
Passive DNS

Why we want to:

- No additional queries for someone to see
- You see more than you otherwise would
- Can detect things you otherwise couldn't
- You see what machines actually used ...

Passive DNS

Why we need to:



Client1: www.goodsite.com, 10.1.2.3
Client 2: www.badsite.com, 10.1.2.3

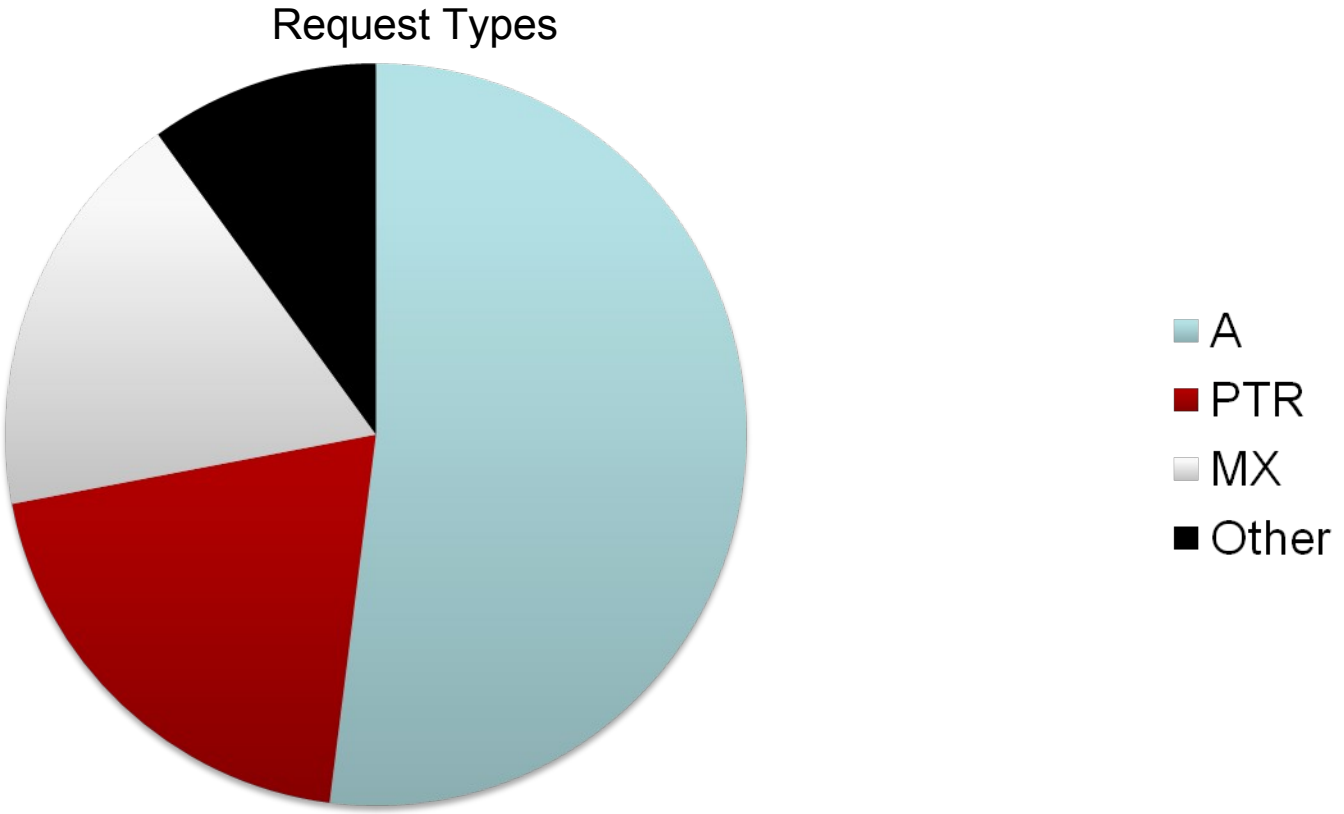
N1: one to one mapping

N2: one to one, one to many, no mapping

Our Setup

- SIE channel 5
  - ~ 260 million packets/day (3100 packets/sec)
  - Represents ~ 370 million packets (de-dup over 4 hours)
  - ~ 200 Bytes/packet
  - ~ 56 GB day raw / ~ 17 GB day with gzip
  - 200,000 msgs per file, ~ 1200 files per day
  - Typical query time between 30 min and 2 hours

Traffic Summary



RBLs account for many millions of A record request per day

For certain networks, up to 80% of lookups are to RBLs

common RBLs seen: ciphertrust.net, vcxde.com, borderware.com, sonicwall.com, fzrbl.org

Fast Flux

- Lots of IP addresses per one domain name
- Provides better uptime for bad sites
  - Load distribution
  - Hard to trace
  - Hard to takedown
- How to find
  - Iterate over message with
    - Low ttl (less than 2000 seconds)
    - Lots of A records per message (10 or more)
  - Iterate by qname of possible messages
    - Total number of uniq A records/IP addresses (25 or more)
    - Total number of ASNS (20 or more)

Fast Flux found (10/20/2009,10/23/2009)

brewers-ca.com.til1tlli.net.
brewers-ca.com.tll1tlii.com.
brewers-ca.com.tll1tlli.net.
cadtrans.net.til1tlli.com.
cpan.cpanel.net.
csajn.com.
dessaxzaa.co.uk.
diff.cpanel.net.
doubleclickr.ru.
ffffefvl.co.uk.
heiiikuv.eu.
httpupdate.cpanel.net.
layer1.cpanel.net.
layer2.cpanel.net.
mgpra.com.
okkkikla.eu.
okkkikll.eu.
rdate.cpanel.net.
rdate.darkorb.net.
rrref1aaz.eu.
sclrz.com.
til1tlli.com.
til1tlli.net.
tll1tlii.com.
tll1tlli.com.
tll1tlli.net.
ttl1lil.com.
ttl1lil.net.
ttl1lll.com.
ttl1lll.net.
towernet.capitalonebank.com.f1iiiti.net.
towernet.capitalonebank.com.f1iiitl.com.
towernet.capitalonebank.com.f1iiitl.net.
towernet.capitalonebank.com.f1liitl.com.
towernet.capitalonebank.com.ltiil1.com.
towernet.capitalonebank.com.ltiil1.net.
towernet.capitalonebank.com.ltlii1.com.
towernet.capitalonebank.com.ltlii1.net.
towernet.capitalonebank.com.ltlil1.com.
towernet.capitalonebank.com.ltlil1.net.
towernet.capitalonebank.com.racder1c.net.
towernet.capitalonebank.com.racder1x.com.
towernet.capitalonebank.com.raeder1f.net.
towernet.capitalonebank.com.rarder1g.com.
towernet.capitalonebank.com.raxsder1.com.
towernet.capitalonebank.com.raxsder1.net.
towernet.capitalonebank.com.rzasder1.com.
towernet.capitalonebank.com.t1fliil.tc.
towernet.capitalonebank.com.yyy1yyrd.co.uk.
towernet.capitalonebank.com.yyy1yyre.co.uk.
towernet.capitalonebank.com.yyy1yyrf.co.uk.
towernet.capitalonebank.com.yyy1yyrg.co.uk.
towernet.capitalonebank.com.yyy1yyrj.co.uk.
towernet.capitalonebank.com.yyy1yyrk.co.uk.
towernet.capitalonebank.com.yyy1yyrl.co.uk.
towernet.capitalonebank.com.yyy1yyrm.co.uk.
towernet.capitalonebank.com.yyy1yyro.co.uk.
towernet.capitalonebank.com.yyy1yyrq.co.uk.
towernet.capitalonebank.com.yyy1yyrr.co.uk.
towernet.capitalonebank.com.yyy1yyrs.co.uk.
towernet.capitalonebank.com.yyy1yyru.co.uk.
towernet.capitalonebank.com.yyy1yyrv.co.uk.
towernet.capitalonebank.com.yyy1yyrx.co.uk.
update.microsoft.com.bbttyak.co.uk.
update.microsoft.com.bbttyak.org.uk.
update.microsoft.com.bbttyam.co.uk.
update.microsoft.com.bbttyam.me.uk.
update.microsoft.com.bbttyap.co.uk.
update.microsoft.com.bbttyap.me.uk.
update.microsoft.com.bbttyaz.co.uk.
update.microsoft.com.bbttyaz.me.uk.
update.microsoft.com.dessaxqaa.co.uk.
update.microsoft.com.dessaxzaa.co.uk.
update.microsoft.com.dessaxzaa.me.uk.
update.microsoft.com.dessaxzaq.co.uk.
update.microsoft.com.dessaxzaq.me.uk.
update.microsoft.com.gerrasawa.eu.
update.microsoft.com.heiiikok.eu.
update.microsoft.com.heiiikoy.eu.
update.microsoft.com.heiiikul.eu.
update.microsoft.com.heiiikum.eu.
update.microsoft.com.heiiikuv.eu.
update.microsoft.com.heiiikuy.eu.
update.microsoft.com.n111sae.eu.
update.microsoft.com.n111sak.eu.
update.microsoft.com.n111sap.eu.
update.microsoft.com.n111saq.eu.
update.microsoft.com.n111say.eu.
update.microsoft.com.n111saz.eu.
update.microsoft.com.okkkikkf.eu.
update.microsoft.com.okkkikkl.eu.
update.microsoft.com.okkkikla.eu.
update.microsoft.com.okkkiklf.eu.
update.microsoft.com.okkkikll.eu.
update.microsoft.com.okkkiklo.eu.
update.microsoft.com.okkkilkf.eu.
update.microsoft.com.okkkulkf.eu.
update.microsoft.com.okktulkf.eu.
update.microsoft.com.okktylkf.eu.
update.microsoft.com.rrref1aaz.eu.
update.microsoft.com.rrref1akz.eu.
update.microsoft.com.rrref1okz.eu.
update.microsoft.com.rrref1ykz.eu.
update.microsoft.com.rrrefjokz.eu.
update.microsoft.com.ujihkei.eu.
update.microsoft.com.ujihkni.eu.
update.microsoft.com.ujihkoi.eu.
update.microsoft.com.ujihkui.eu.
updates.cpanel.net.
www.adbnr.ru.
www.ads-t.ru.
www.adsyndication.ru.
www.adtcp.ru.
www.bannerdriven.ru.
www.bannert.ru.
www.clickmeter.ru.
www.doublebanner.ru.
www.doubleclickr.ru.
www.htmlads.ru.
www.posteonline.it.sclrz.com.
www.yahoosite.ru.
yahoosite.ru.
ztrblg.com.

Malicious Domains

- Registered by bad actor – not compromised
- How to find
  - Cheat by starting with list
    - APWG
    - Maybe won't have to
  - Name has large amount of unique characters (over 20)
  - Name has tld in middle (www.yourbank.com.imbad.com)

Other

- DNS exfiltration/tunneling
  - Over 40 uniq chars in qname
- DNS amplification
  - For DDOS participation
- Outbound connections with no previous resolution
- DNS rebinding
  - www.attacker.com -> some.public.ip, ttl = 2
  - www.attacker.com -> 10.1.mydatabase.ip
- Just plain out of the ordinary
  - ns1.ziyouforever.com (zi you men – “door to freedom”)
    - `784bc3c09961b67b5f3f6f6783a54881b59f5e53680937d7ce281407.6.bnhyj.com`
    - `08f0b06a25a5cf1f9df501bc39306fbc6ff7875646817b4845c17da0.6.ewsxz.com`

On with the flow

- How to use results with flow
- Pysilk

```
import ncap
import sdnslib
import silk
ips = silk.IPSet()
f = ncap.ncapfile('/path/to/my/file')
for msg in f:
    dnsmsg = sdnslib.message(msg.payload)
    for rr in dnsmsg.answers:
        ips.add(rr['address'])
ips.save('/path/to/my/ip.set')
```

IPs to names

```
import ncap
import sdnslib
import silk
lookup = {}
for msg in ncap.ncapfile('/path/to/my/file.ncap')
    dnsmsg = sdnslib.message(msg.payload)
    for rr in dnsmsg.answers:
        lookup[rr['name']] = rr['address']
for rec in silk.SilkFile('/path/to/my/file.rw')
    print lookup[rec.dip]
```



A Temporal Logic For Network Flow Analysis

Tim Shimeall
tjs@cert.org

© 2010 Carnegie Mellon University

NO WARRANTY

THIS MATERIAL OF CARNEGIE MELLON UNIVERSITY AND ITS SOFTWARE ENGINEERING INSTITUTE IS FURNISHED ON AN “AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this presentation is not intended in any way to infringe on the rights of the trademark holder.

This Presentation may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

© 2010 Carnegie Mellon University

Overview

Motivation

Temporal Logic

Application to Flow

Example

Implementations

© 2010 Carnegie Mellon University

Motivation

Clarify timing relationships

Formalize analysis semantics

- Clearer discussions
- Enhance automation & frameworks
- Combining analyses

Avoid over-specification of timing

Support reasoning about analysis tasks

Access temporal logic methods

© 2010 Carnegie Mellon University

Temporal Logic

Logic with explicit inclusion of time

Classically, first-order logic, could be any logic form

Temporal interpretation: Instantiating circumstances

- Linear time with rollback on contradiction
- Branching time with branch termination on contradiction
- Advantage to linear: simpler structure, no worry over paths
- Advantage to branching: can express path-related conditions

© 2010 Carnegie Mellon University

Temporal Logic Operators

Next(*t*,*p*) – *p* is true in the instant after *t*

Global(*p*) – *p* is true independent of time

Following(*t*, *p*) – *p* is true at some instant after *t*

Until(*t*,*p*,*q*) – *p* is true at each instant after *t* until *q* is true

Forall (p) – p is true along all paths

Exists (p) – p is true along at least one path

© 2010 Carnegie Mellon University

Adaptation to Flow

Description first, then reasoning

Iterative semantics – suitable for filter-like processing

Specific semantics:

- 5-tuple
- Ordinal time (inexact comparisons)
- Related flows

© 2010 Carnegie Mellon University

Adapted semantics

$R(f_1,f_2)$ relation – flow-flow connection

p(f,…), q(f,…) – logic predicates on flow records/fields

Enable reasoning using Horn clause resolution and backtracking

© 2010 Carnegie Mellon University

Temporal Operators for Flow

Globally:

G(p): forall(R(f,f') → p(f) and p(f'))

Next:

N(f,f'): iff R(f,f') and f'.stime > f.stime and
does not exists (
R(f,f'') and f'.stime > f''.stime >f.stime)

N*(f,f'): transitive relation on N

X(f,p): forall(N(f,f') →p(f'))

Following:

F(f,p): exists(N*(f,f') and p(f'))

Until:

U(f,p,q):
exists (N*(f,f'') and q(f''),
forall (N*(f,f') and f''.stime>f'.stime → p(f') and not q(f')))

© 2010 Carnegie Mellon University

Descriptive Temporal Example

Spam(s,f):

R(f,f'): f.sip = f'.sip = s and s not on *whitelist*

If and only if

|{f', Following(f,f',f'.stime < f.stime+*5min* and f'.dport=*email)*}|>15 and

|{f', Following(f,f', f'.stime < f.stime+*5min* and f'.dport=*email)*}| ≥ |{f', Following(f,f', f'.stime < f.stime+5*min*)}| * 0.1

© 2010 Carnegie Mellon University

Implementation

Use temporal logic to express analysis criteria

Prolog-based (GNU-Prolog)

- Logic programming, incorporating time in resolution
- Initial prototype to refine semantics
- Construct interface to analysis tools (plugin)

Python-based (PySiLK)

- Declarative programming, incorporate limited resolution mechanism
- Secondary prototype to demonstrate applicability

Eventually construct reasoning rules for analysis relationships or proof

© 2010 Carnegie Mellon University

Conclusions

Temporal logic adaptation of flow analysis offers opportunity to encompass large literature of pre-existing methods

Formalization of time relationships offers opportunity to improve flow analysis methods

More formal reasoning on flow analysis?

© 2010 Carnegie Mellon University



Flow Analysis for Network Situational Awareness

Tim Shimeall

January 2010

© 2010 Carnegie Mellon University

NO WARRANTY

THIS MATERIAL OF CARNEGIE MELLON UNIVERSITY AND ITS SOFTWARE ENGINEERING INSTITUTE IS FURNISHED ON AN “AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this presentation is not intended in any way to infringe on the rights of the trademark holder.

This Presentation may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

© 2010 Carnegie Mellon University

What We Will Cover

Introduction

Your Network

- Fundamentals of networks, flow, and protocols
- Malicious traffic

External Events & Trends

- Malware
- Networks in the Broad

Working Together

- Network dependencies
- Analysis

Summary

© 2010 Carnegie Mellon University

What this class is

Approaches, methods, trends of interest in building network situational awareness

Big picture view of analysis

Assumes you have good handle on analysis tool suites

© 2010 Carnegie Mellon University

What this class is not

Cool tricks with SiLK

Installing and using SiLK

Everything you need to do analysis

Bits and bytes on the wire

© 2010 Carnegie Mellon University

A couple of rules…

ASK!

No smoking/chewing

Cell phones on stun

No smelly food

© 2010 Carnegie Mellon University

Recognition Stances





© 2010 Carnegie Mellon University

Definition

The systematic gathering, analysis and interpretation of data from local and remote networks regarding structure, applications, traffic and resources to produce actionable information for decision making in network operations and defense. *(Shimeall, 2008)*

Network Situational Awareness

© 2010 Carnegie Mellon University

Alternate Definitions

Situation Awareness (SA): “The perception of elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future." (Endsley, 1988)

Network SA: “the operational picture that consolidates all available information that is actually needed for identifying attacks and for selecting and applying appropriate countermeasures.” (Kemmerer et. al., 2008)

© 2010 Carnegie Mellon University

Network Situational Awareness Practice

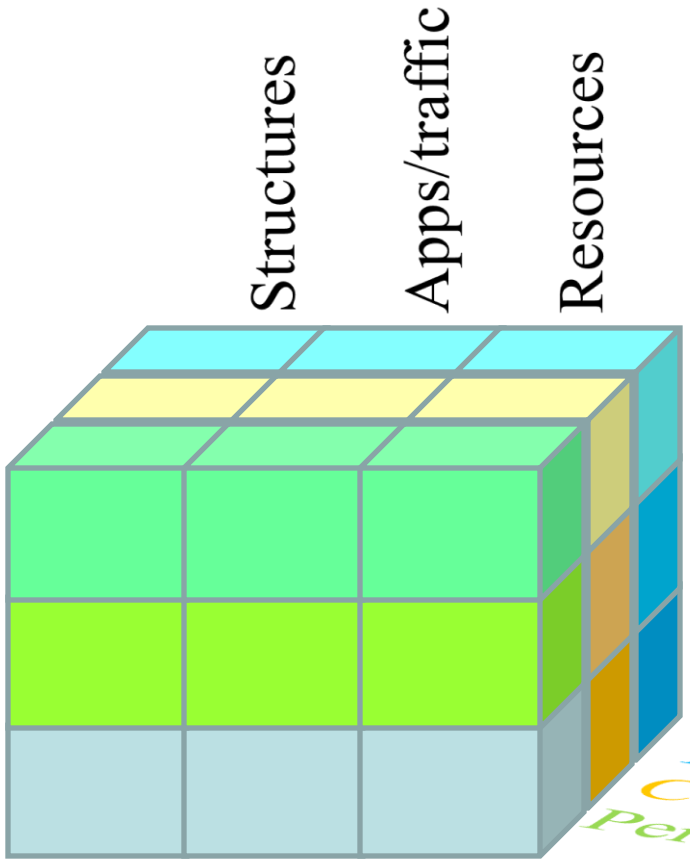
Know your network

Know current external events and trends

Know how they fit together

© 2010 Carnegie Mellon University

How these Definitions Apply



Know Network

Know Events/Trends

Know How They Fit

Predict
Comprehend
Perceive

© 2010 Carnegie Mellon University

Vulnerability Note VU#800113

Multiple DNS implementations vulnerable to cache poisoning

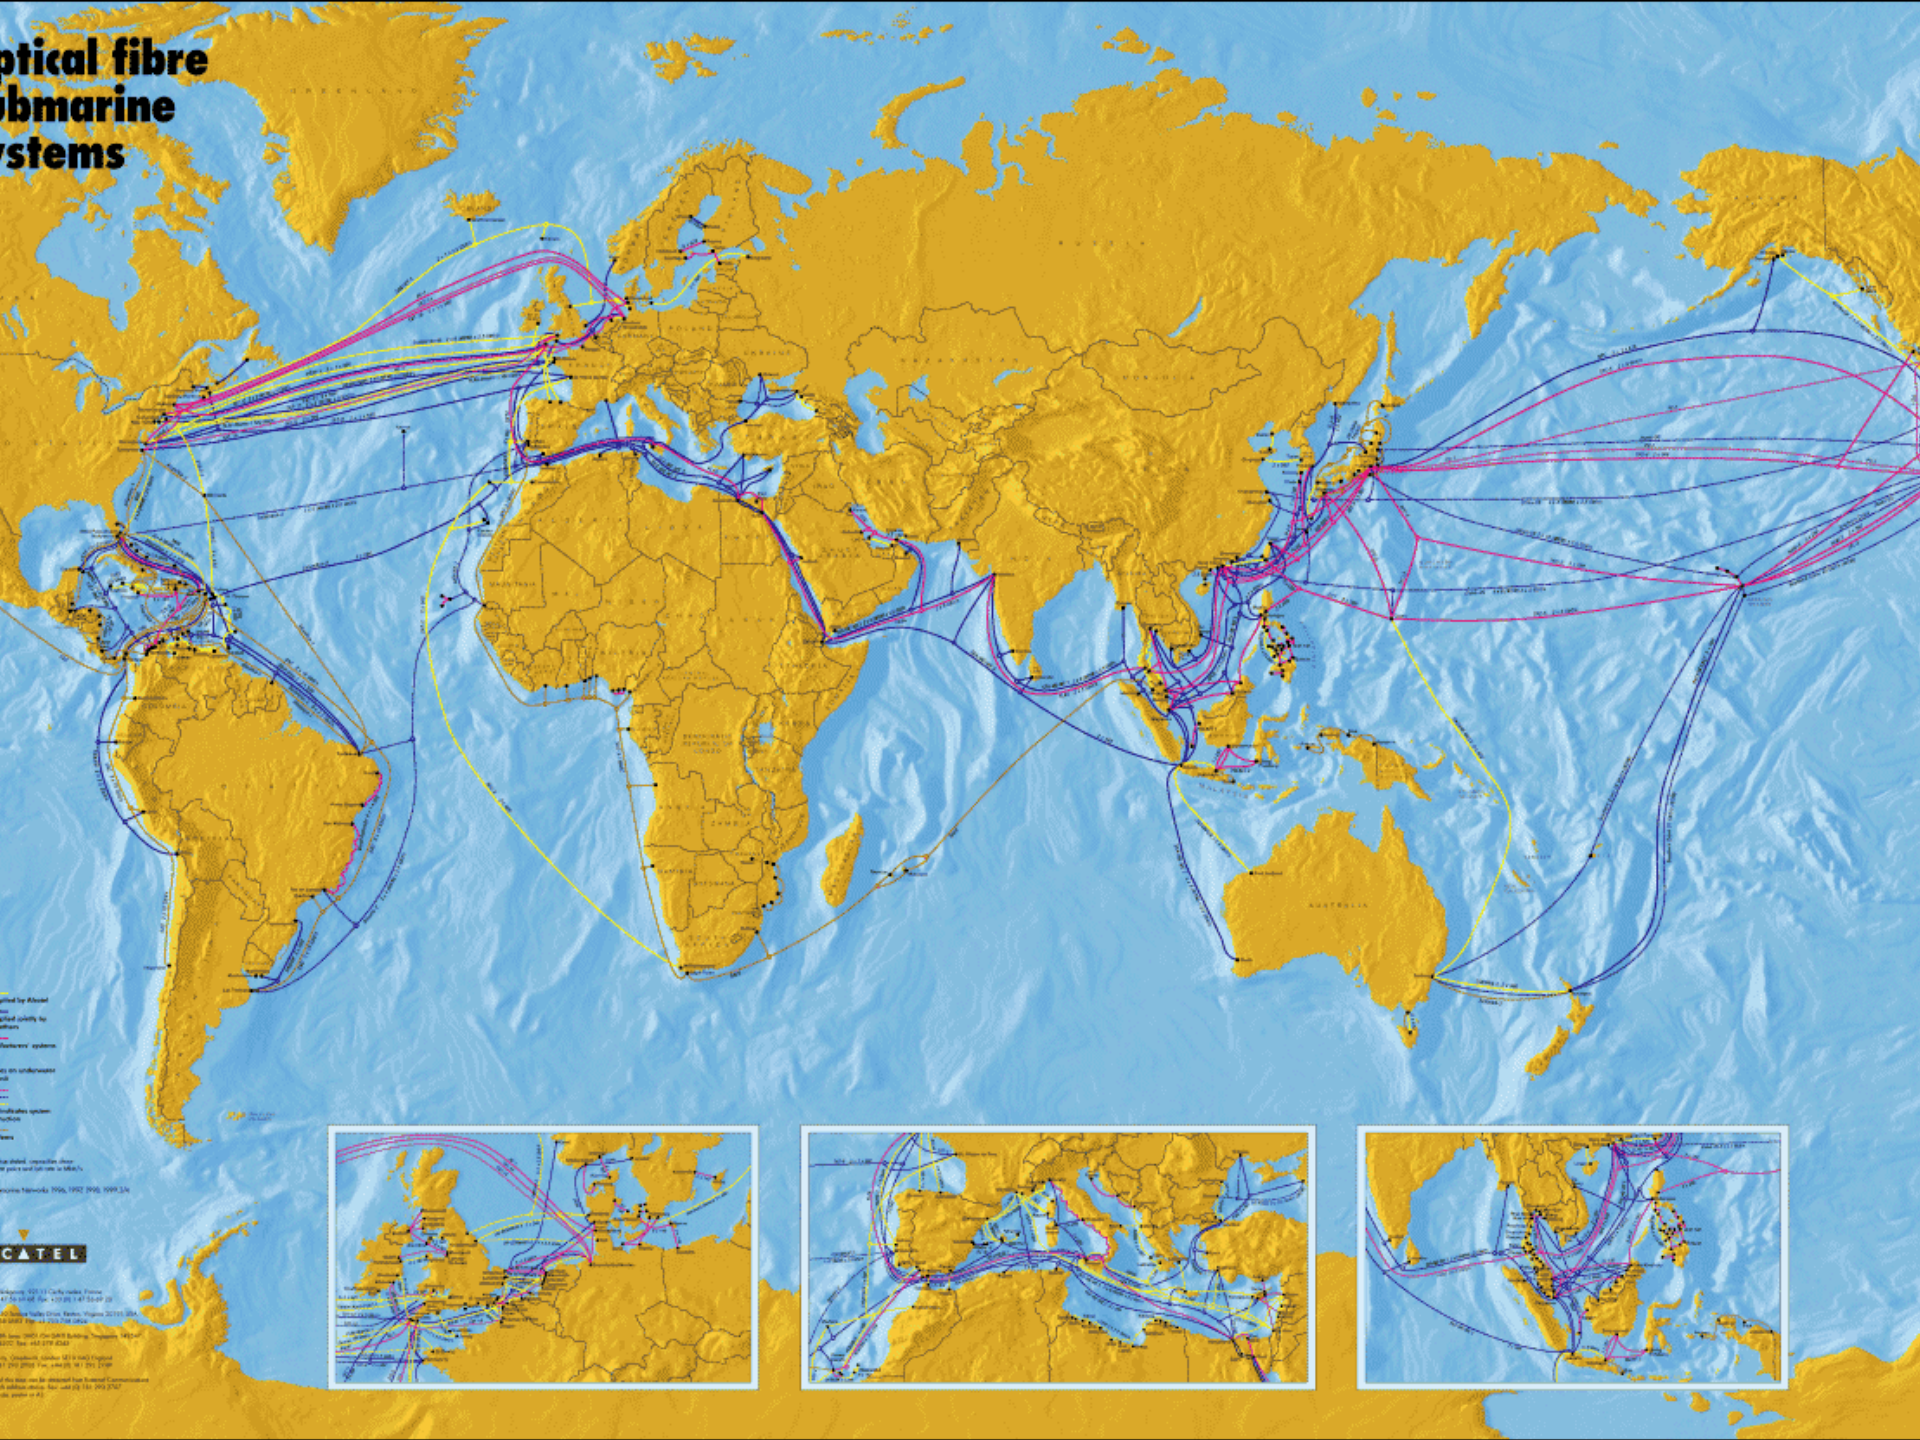
- http://www.youtube.com/watch?v=XDKw8ny6IcM#

July 2008

Sarah on DNS



© 2010 Carnegie Mellon University


ptical fibre
ubmarine
ystems
GREENLAND
RUSSIA
MONGOLIA
AUSTRALIA
CATEL

Estonia, April 2007



© 2010 Carnegie Mellon University

The 2008 Olympics

USA TODAY Home News Travel Money Sports Life Tech

News » World ▪ Troop deaths in Iraq

Internet sites still blocked for Olympic reporters

Updated 7/30/2008 10:34 PM | Comments 45 | Recommend 8 E-mail | Save | Print | RSS



Enlarge By Oded Balilty, AP

A foreign journalist uses Internet services provided at the Main Press Center Tuesday in Beijing. After promising the Internet will be uncensored for journalists during the Olympics, the IOC delivered a stark clarification on Tuesday, saying many Internet sites will be blocked under controls applied by China's communist government.

BEIJING (AP) — Olympic organizers are backtracking on another promise about coverage of the Beijing Games, keeping in place blocks on Internet sites in the Main Press Center and venues where reporters will work.

Mixx it
Other ways to share:
Yahoo! Buzz
Digg
Newsvine
Reddit
Facebook
What's this?

STORY CORRECTION: Ebersol's comments mischaracterized

The blocked sites will make it difficult for journalists to retrieve information, particularly on political and human rights stories the government dislikes. On Tuesday, sites such as Amnesty International or any search for a site with Tibet in the address could not be opened at the Main Press Center, which will house about 5,000 print journalists when the games open Aug. 8.

"This type of censorship would have been unthinkable in Athens, but China seems to have more formalities," said Mihai Mironica, a journalist with ProTV in Romania. "If journalists



© 2010 Carnegie Mellon University

Questions of Interest

Is my bandwidth increasing from business-related activity, or from non-work related activity?

How will my business be impacted by implantation of more stringent security policy?

If my backbone Internet Service Provider chooses to de-peer with another backbone provider, how will I be affected?

How will socio-political uprisings impact my network?

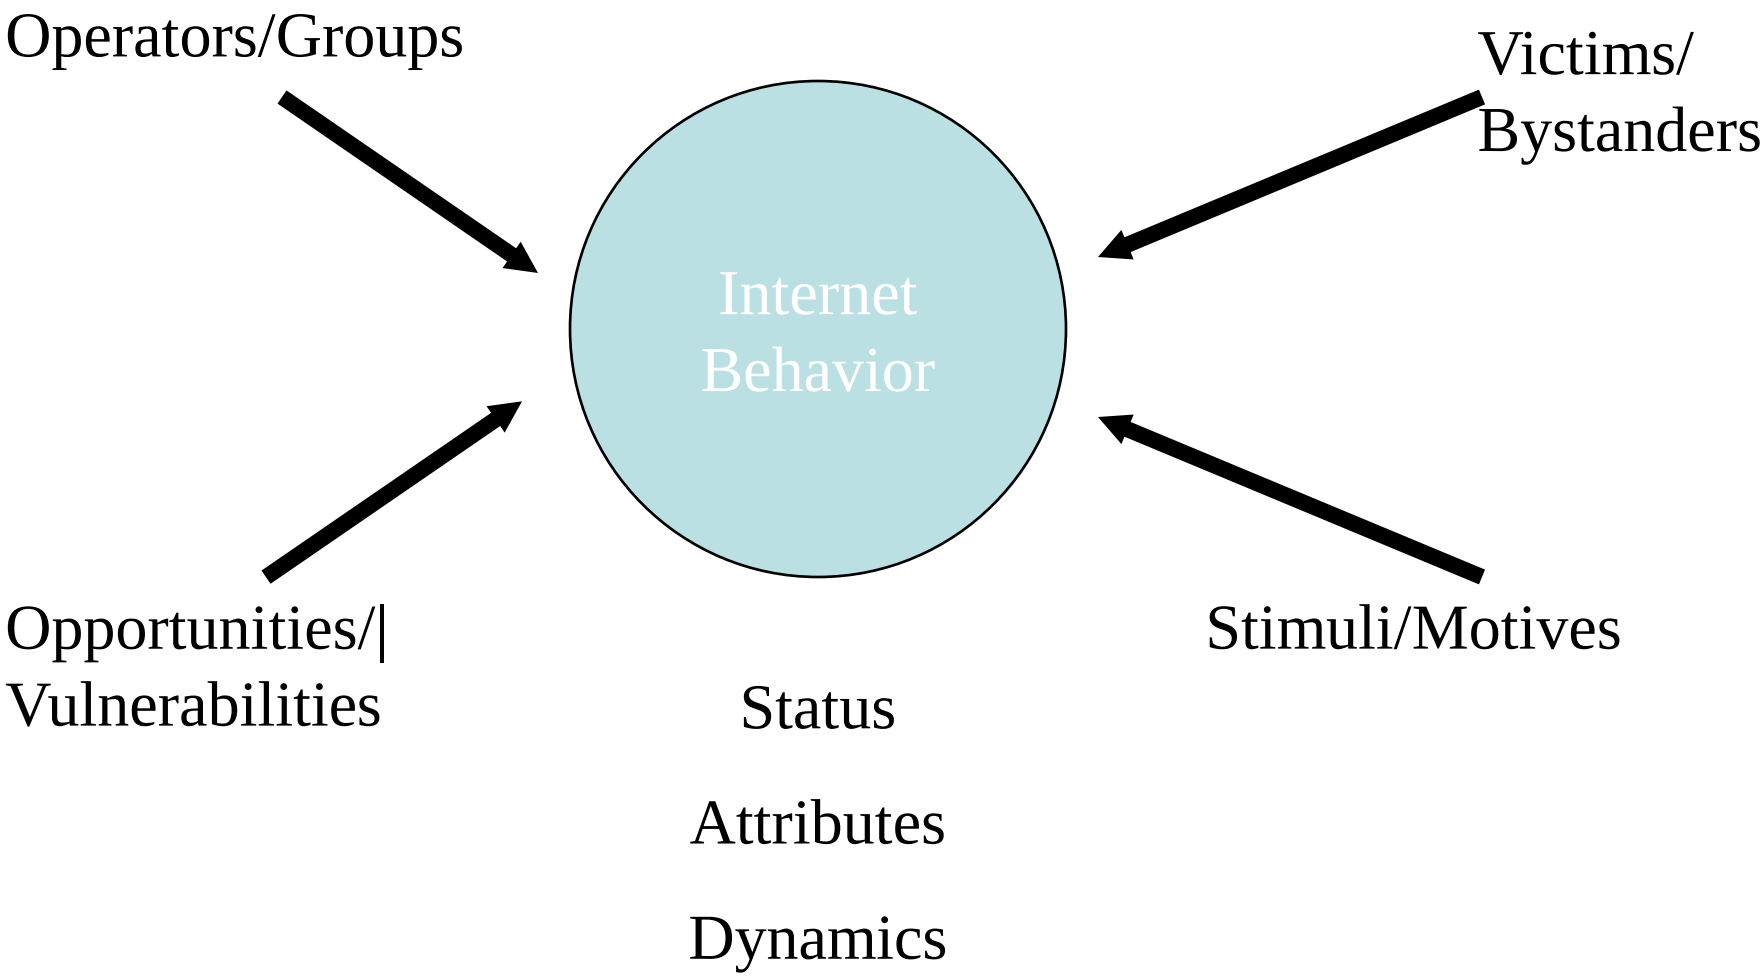
What are the most important dependencies my network has with external resources?

Do computers on my network follow policy?

Can my network survive a distributed denial-of-service attack? How can I prioritize resources during a bandwidth-limiting attack?

© 2010 Carnegie Mellon University

Building Understanding



© 2010 Carnegie Mellon University

Challenges to Analysis

- Gathering sufficient datasets to make statistically valid judgments
- Developing automated technical analysis tools
- Developing a reliable methodology for cyber-analysis
- Overcoming organizational bias against sharing information



© 2010 Carnegie Mellon University



Network Fundamentals

© 2010 Carnegie Mellon University

What We Will Cover

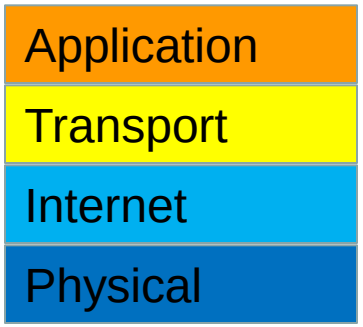
- Introduction
- Your Network
  - Fundamentals of networks, flow, and protocols
  - Malicious traffic
- External Events & Trends
  - Malware
  - Networks in the Broad
- Working Together
  - Network dependencies
  - Analysis
- Summary

© 2010 Carnegie Mellon University

The Internet

Internet: Network of Networks

- Connected by routers, no central control
- Using common set of protocols



Internet Protocol (IP)

- flexible routing of information from source to destination

Transmission Control Protocol (TCP)

- Sequencing of series of packets to transmit data reliably over Internet

Other protocols running on top of IP:

- UDP – one-directional burst of packets
- ICMP – network management protocol
- BGP – Router protocol
- IPSEC – VPN traffic

© 2010 Carnegie Mellon University

How IP Works

Packet switched:

- Flow of information broken into chunks
- Each routed independently by best route to destination
- Destination must reassemble into correct order (offset)

Internet Address:

- Logical network (location) & Logical host (identity)
  - Subnet masking & CIDR blocks
- V4 (1982) -- current version (32 bit addresses)
- V6 (1999) -- forthcoming version (128 bit addresses)

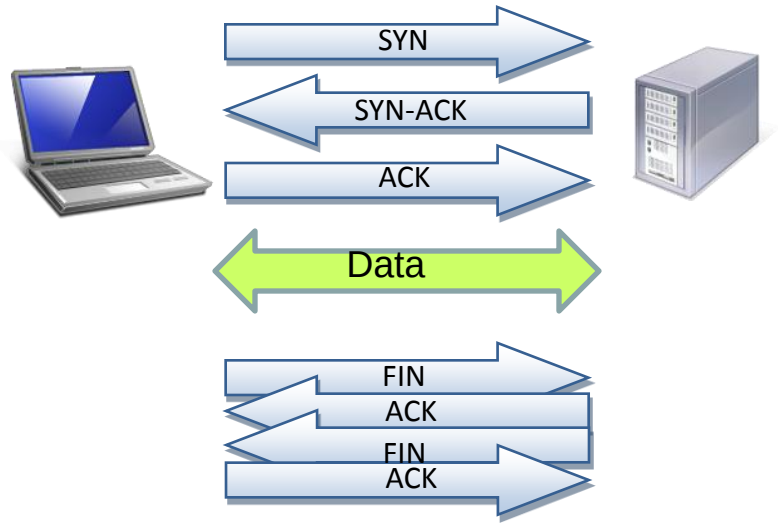
© 2010 Carnegie Mellon University

How TCP Works

Header added to packets

Connection: IPs and ports

- Service port
- Dynamic port



Initial handshake

- SYNch sequence numbers (assembly/retransmit)
- ACKnowledge SYNch
- ACK (with data)

Termination

- FINalize/ ACK FIN
- ReSeT/ no ack

© 2010 Carnegie Mellon University

Application Protocols

On top of TCP and UDP (mostly)

Sequences of packets to support common applications

- SMTP
- HTTP
- HTTPS / TLS
- DNS
- Others?

© 2010 Carnegie Mellon University

Simple Mail Transfer Protocol

- TCP/25 by default
- Transfer-agent based
- Text Protocol
- Single connection, multiple messages (maybe)
- Easily forged

S: 220 smtp.example.com ESMTP Postfix
C: HELO relay.example.org
S: 250 Hello relay.example.org, I am glad to meet you
C: MAIL FROM:<bob@example.org>
S: 250 Ok
C: RCPT TO:<alice@example.com>
S: 250 Ok
C: RCPT TO:<theboss@example.com>
S: 250 Ok
C: DATA
S: 354 End data with <CR><LF>.<CR><LF>
C: From: "Bob Example" <bob@example.org>
C: To: Alice Example <alice@example.com>
C: Date: Tue, 15 Jan 2008 16:02:43 -0500
C: Subject: Test message
C: Hello Alice.
C: Your friend, Bob
C: .
S: 250 Ok: queued as 12345
C: QUIT
S: 221 Bye {The server closes the connection}

Example from wikipaedia.org

© 2010 Carnegie Mellon University

HyberText Transfer Protocol

TCP/80 or TCP/8080

Vast bulk of traffic

Request/Response

Stateless (almost)

Requests:

request-line
headers (host)
empty line
optional message

Request-line:

- GET url
- PUT url
- HEAD url
- DELETE url
- TRACE url
- OPTIONS
- POST url
- CONNECT

Response:

Status line (404, 200)
Message (data)

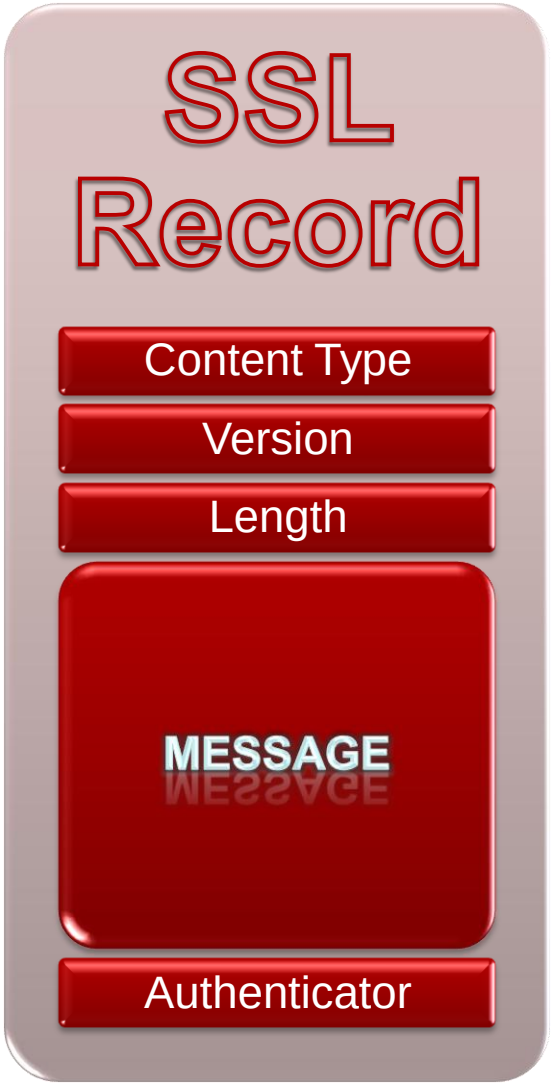
© 2010 Carnegie Mellon University

SSL Record Format

There are four content types:

- Handshake, 0x16
  - Negotiate encryption capabilities
- ChangeCipherSpec, 0x14
  - Flag the recipient that the encryption scheme has changed
  - Prereq is a successful handshake
- Application, 0x17
  - The bulk of the data transfer
- Alert, 0x15
  - Errors and warnings

Regardless of content type, the ***message*** is always encrypted with the current session state

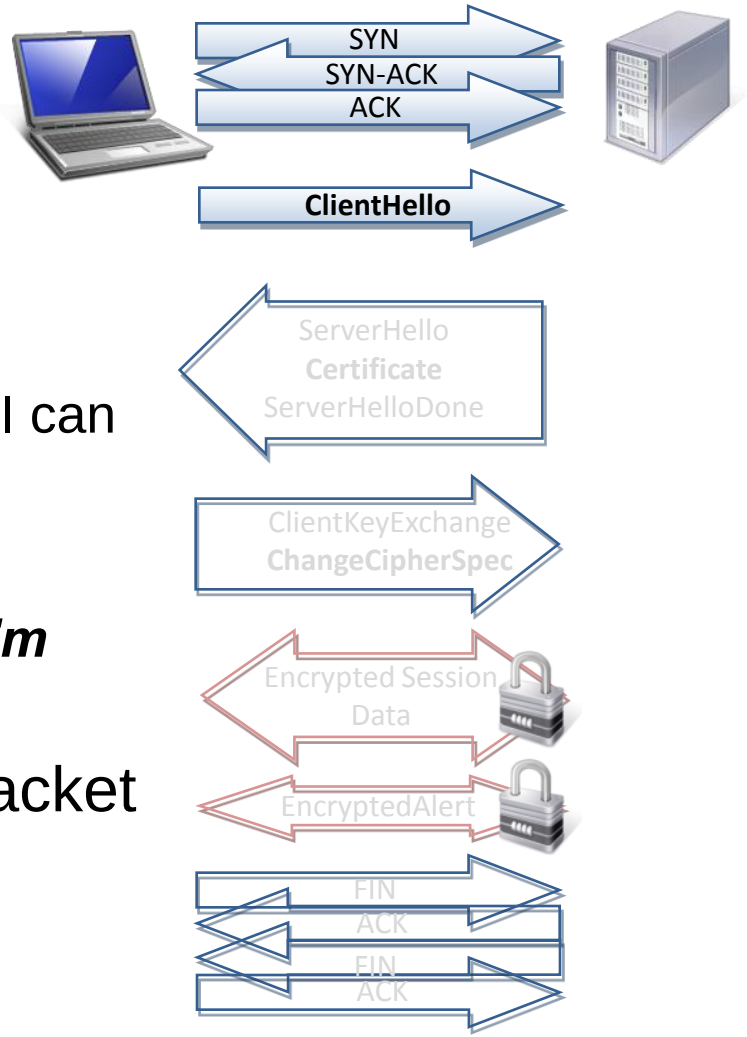


© 2010 Carnegie Mellon University

SSL Client Hello

What happens when a client connects to an SSL server?

- 3-way handshake
- Client sends hello
  - These are the encryption ciphers I can handle
  - Here's a random seed number
  - ***Here's the name of the server I'm trying to connect to***
- The client hello all fits in one packet

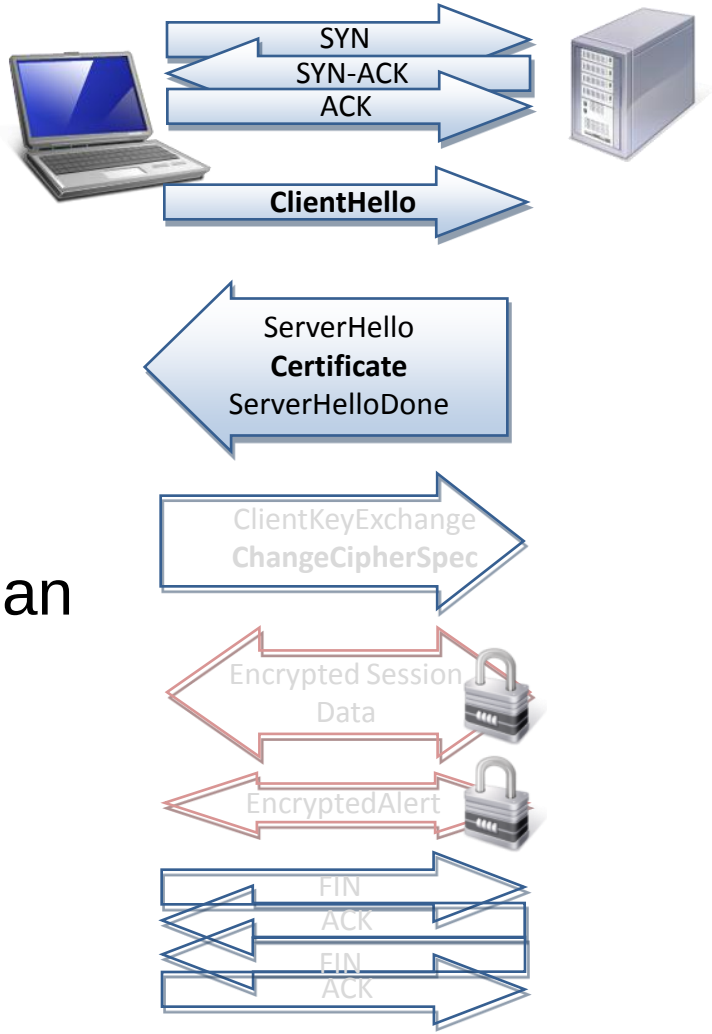


© 2010 Carnegie Mellon University

SSL Server Response

Server replies to the client

- Server sends “Hello”
  - Encryption and seed
- Server sends its certificate
  - public key and usually cert chain
- Server sends “Done”
- Response usually takes more than one packet

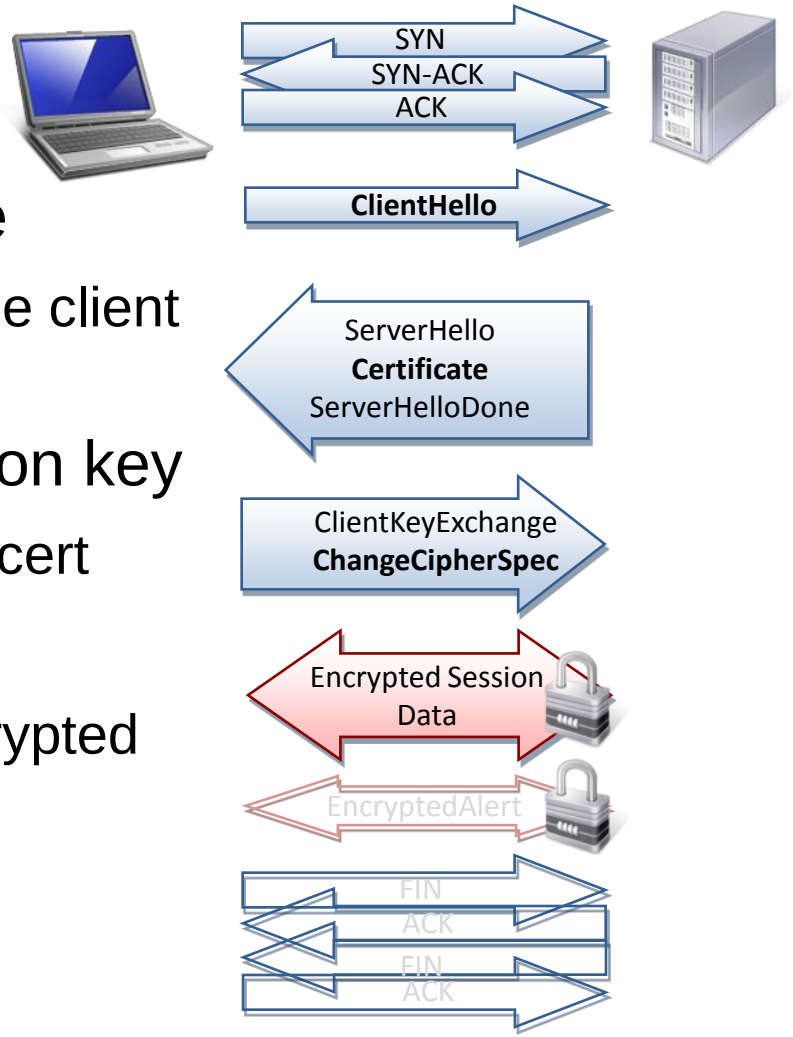


© 2010 Carnegie Mellon University

SSL Client Response

After the server is done:

- Behind-the-scenes, the client validates the certificate
  - This only generates traffic if the client needs an updated CRL
- The client exchanges a session key
  - The key is encrypted with the cert
- `ChangeCipherSpec`
  - From now on, all traffic is encrypted

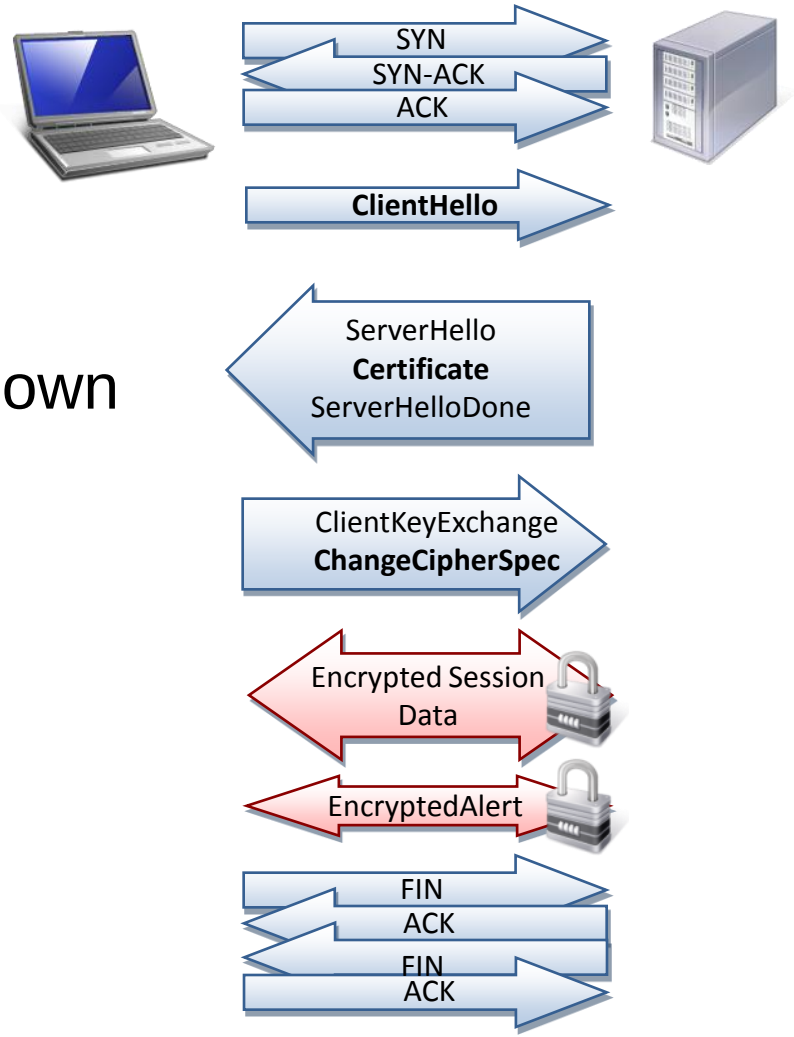


© 2010 Carnegie Mellon University

SSL Session Teardown

When they're done

- Someone issues an alert
  - It's encrypted, so we can't tell anything about it
- The TCP connection is torn down



© 2010 Carnegie Mellon University

Domain Name System

More than just hostname → IP

Query hierarchy of nameservers

- Local nameserver (resolver): answer from cache or preloaded resolutions, may do recursive queries
- Authoritative nameserver: answer based on domains it covers, or recurse
- Root nameserver: answer top-level, delegate, or generate errors

© 2010 Carnegie Mellon University

DNS protocol

UDP/53 or TCP/53

Client queries local (address, ptr, mx, ns, hinfo, any)

Local responds from cache or queries to root

Root responds with referral to TLD or error

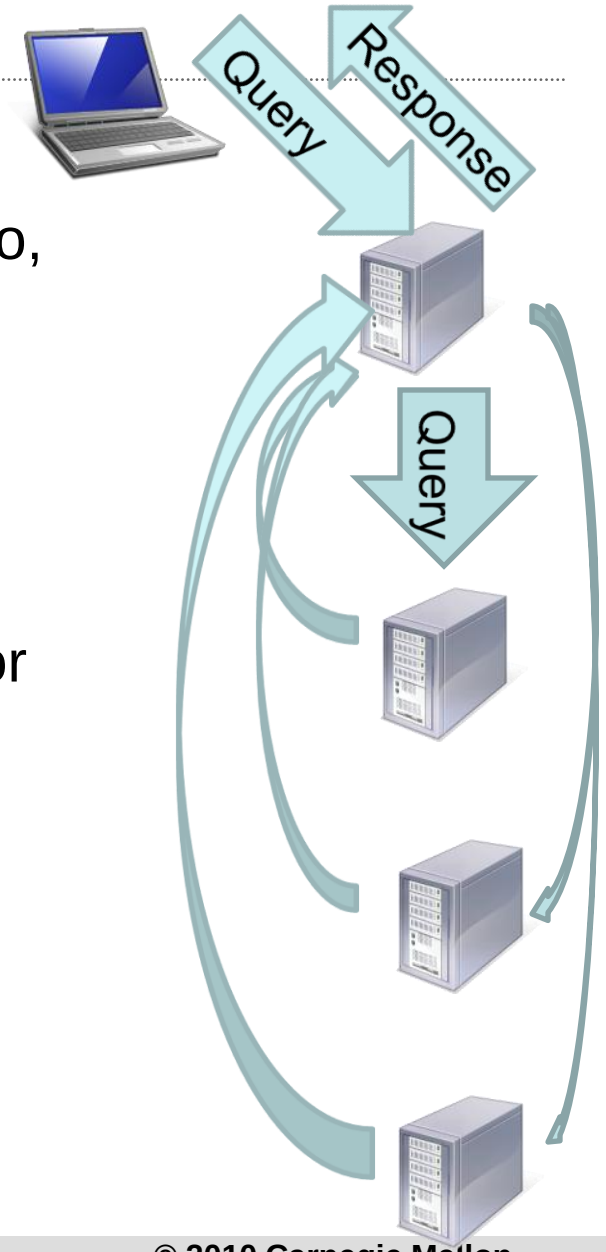
Local queries TLD

TLD responds with referral to authority or error

Local queries authority

Authority sends answer

Local sends answer



© 2010 Carnegie Mellon University

Other Protocols

Wikipaedia is your friend!

Want to know:

- Sequence of messages
- Packets/message
- Size range of packets
- Service ports/protocols
- Other application protocol characteristics

© 2010 Carnegie Mellon University

Peer to Peer

Cloud discovery

- Fumble (past sites)
- Service ports
- Registered sites

Join

Export data (often stream)

Import data (often slow)

© 2010 Carnegie Mellon University

Why do We Care?

Know your network

- Services, servers
- Dependencies

Know your threats

- Open vectors
- Non-protocol on service ports
- Signalling/Beaconing

© 2010 Carnegie Mellon University



Malware and Malicious Traffic

© 2010 Carnegie Mellon University

What We Will Cover

- Introduction
- Your Network
  - Fundamentals of networks, flow, and protocols
  - Malicious traffic
- External Events & Trends
  - Malware
  - Networks in the Broad
- Working Together
  - Network dependencies
  - Analysis
- Summary

© 2010 Carnegie Mellon University

Malicious Traffic

Scanners

Worms/Viruses

DDos/Backscatter

Spam/Response

Backdoors

© 2010 Carnegie Mellon University

A taxonomy of Attributes

Backscatter: Few sources, scattered evenly across network, generally contains RST or ACK flags.

Scans: Single source, usually strikes the same port on many machines, or different ports on the same machine

DoS: Multiple sources, single target, usually homogenous (but no requirement). May be oddly sized

Worms: Scanning from a steadily increasing number of hosts

Key Servers: Identifiable by IP addresses (e.g. Hotmail)

© 2010 Carnegie Mellon University

Step 1: What does scanning look like?

Scanning generates a lot of flows, if there is any scanning, the noisiest is probably the scanner.

So, we want to find the noisiest IP, refilter (or count by IP and/or port)

Use what we know about how the protocol behaves as well. (TCP, for example)

How many packets in TCP prior to sending data?

How many bytes in a TCP packet if there's no data?

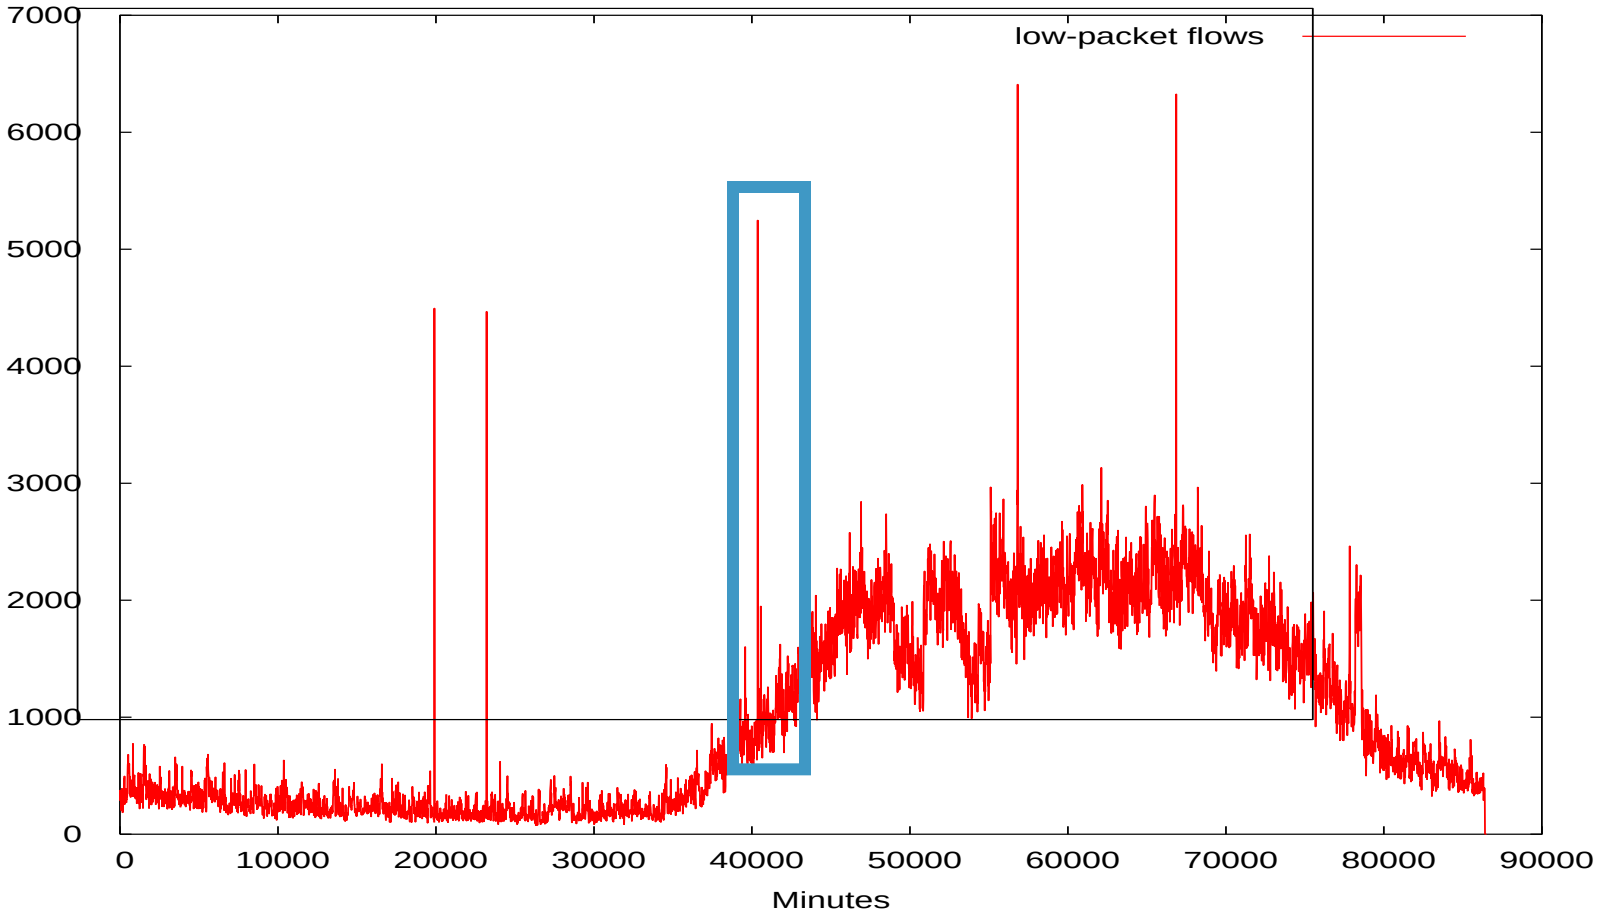
What's TCP do if it can't connect the first try?

What SiLK tools can count by IP?

What traffic goes to addresses with no host?

CERT | Software Engineering Institute | Carnegie Mellon
© 2010 Carnegie Mellon University

Find the Scan



© 2010 Carnegie Mellon University

Step 2: Discerning what spikes are

Use rwstats and look at the busiest source IP addresses

In general, there will be a small number of addresses with a large number of packets

- Those are usually scanning, or worms, or similar
- Although it could be hotmail

© 2010 Carnegie Mellon University

Step 3: Checking vulnerabilities

In general, you are looking at scan data after the fact.

Next step is to determine what the scanner found.

Reverse the query using IP Sets

© 2010 Carnegie Mellon University

Step 3: Checking Vulnerabilities (ii)

Build a set of all the targets that the scanner visits

Use this set as the source IP's and the scanner's IP as the destination IP on the opposite direction of traffic

Profile the flags and the responses

- Who SYN|ACK'd?
- Who RST?

© 2010 Carnegie Mellon University

Measuring Worm Propagation

When looking for a worm you'll usually have a profile

- Target ports hit, packet size, etc.

In those cases, you can use tools like rwfilter and rwstats for accounting rather than discovery

- Worms usually result in an increase in service activity, especially if the service is relatively low-traffic (I.e., not web)

© 2010 Carnegie Mellon University

DDoS Analysis

DDoS = **D**istributed **D**enial **O**f **S**ervice

Crude, mechanical attack where some number of subverted clients overwhelm a target

- Does *not* have to be spoofed
- Does *not* have to be a SYN Flood
- Does *not* have to involve garbage traffic

Can strike at different layers of the stack

- SYN Floods: hit TCP
- Other floods will affect router, IP traffic

© 2010 Carnegie Mellon University

DDoS Analysis

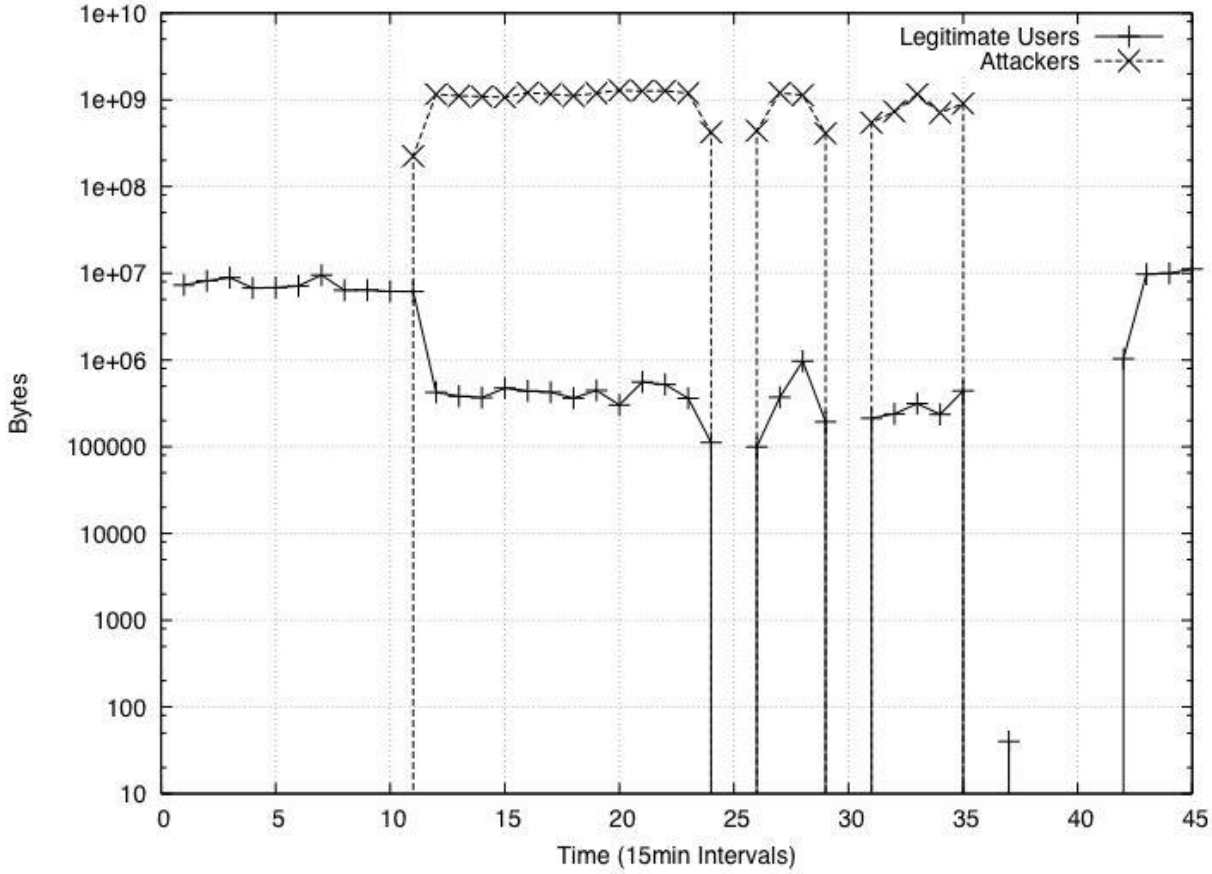
First step will involve identifying and plotting the traffic

- DoS generally is not subtle (can't afford to be in several cases)
- Various plots are useful:
  - Traffic volume over time
  - Number of clients visiting target over time

In an ideal world, should be obvious

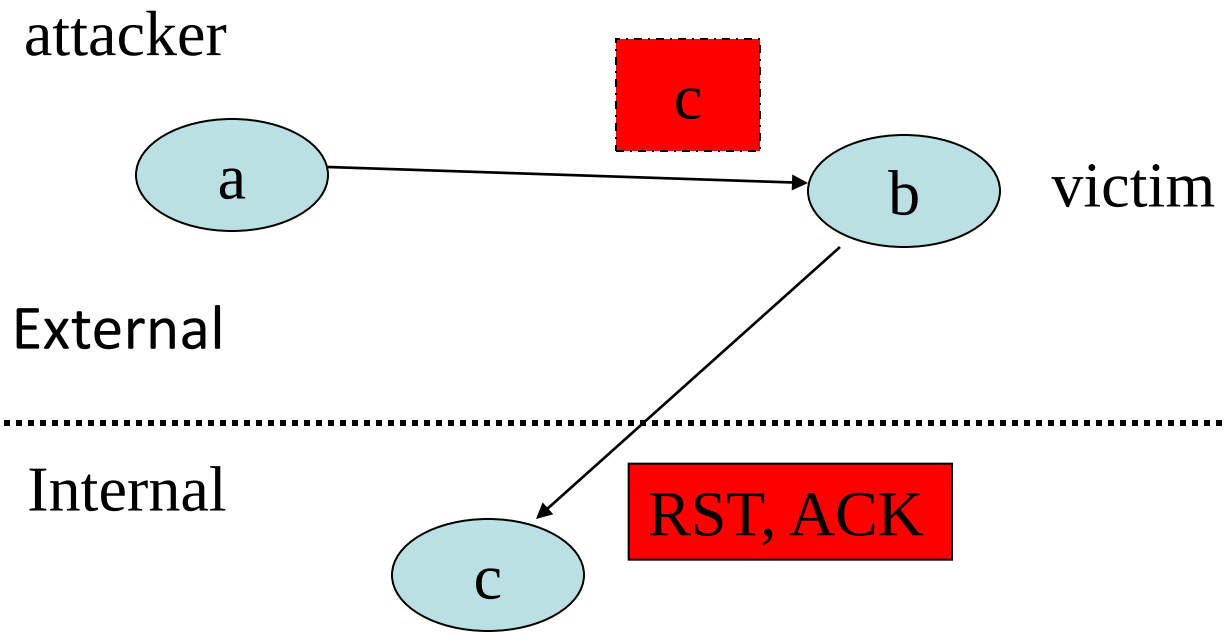
© 2010 Carnegie Mellon University

DDos



© 2010 Carnegie Mellon University

Backscatter



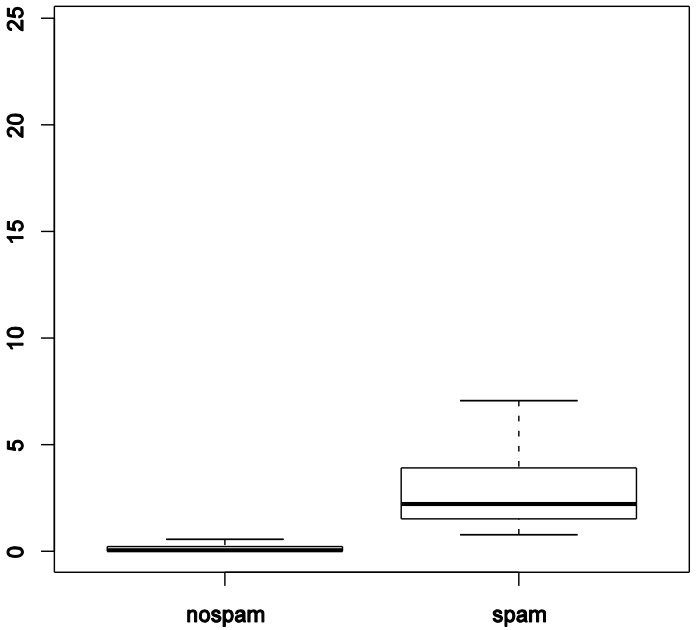
CERT | Software Engineering Institute | Carnegie Mellon
© 2010 Carnegie Mellon University

Spam

Rate (the main one)

Locality

Exclusivity



© 2010 Carnegie Mellon University

Responses to Spam

About 0.1% of spam gets response

Most spam responses not email

Can detect spam flows, then look for non-automated replies to spam flows

Suspicious destinations

- Alleged criminal connections
- DHCP pools
- ADSL pools

© 2010 Carnegie Mellon University

What We Will Cover

- Introduction
- Your Network
  - Fundamentals of networks, flow, and protocols
  - Malicious traffic
- External Events & Trends
  - Malware
  - Networks in the Broad
- Working Together
  - Network dependencies
  - Analysis
- Summary

© 2010 Carnegie Mellon University

Purpose of Modern Malware

Steal information (key logging, screen scraping, DNS redirection, etc.)

Commandeer system resources (launch DOS, relay network traffic, spam, etc.)

Propagate

Locate and kill security programs and competing malware

Hide

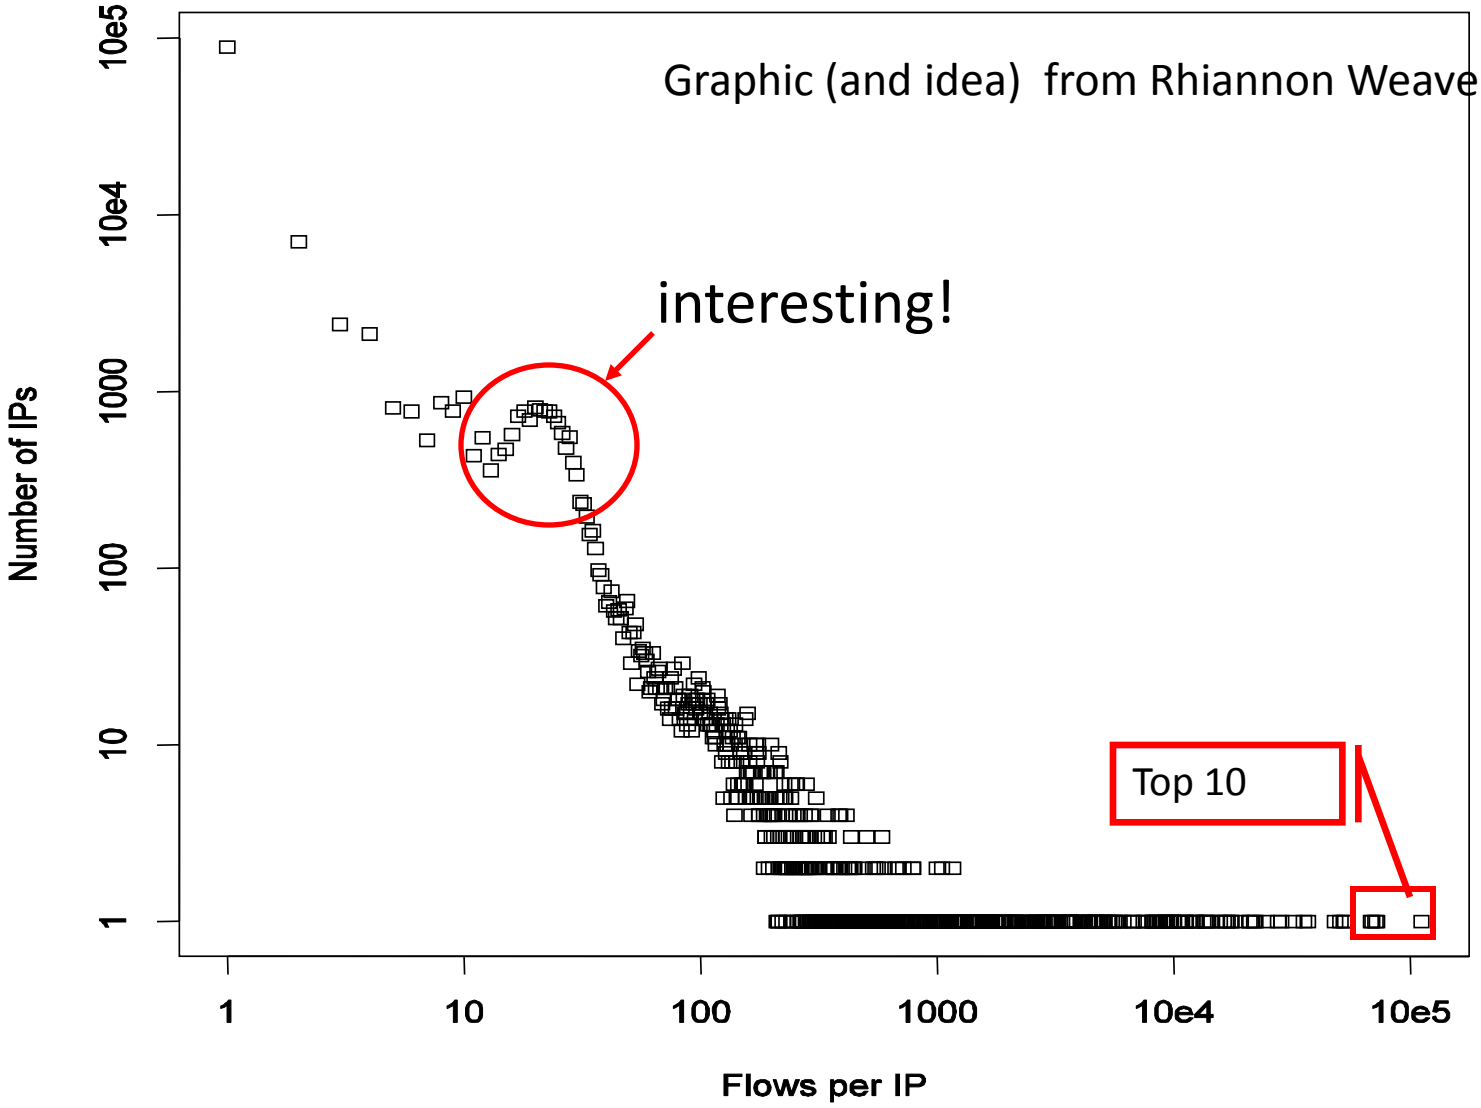
© 2010 Carnegie Mellon University

Methods of Modern Malware

- Most not novel
  - Derived from base of existing code
    - Borrow from other malware
    - Borrow from tutorial examples
  - Change command and control channel
  - Slightly adjust obfuscation scheme
- Where novel, closely held
  - Avoid antivirus
  - Block competitors

© 2010 Carnegie Mellon University

Top N isn't everything



© 2010 Carnegie Mellon University

Malware and the Internet

- Propagation (dealt with previously – violations of locality; finding scans)
- Receiving commands
  - C&C site
  - C&C channel
  - Downloads (loading more malware)
- Executing commands
  - DOS detection
  - Data exfiltration drop site
  - Fraud artifacts

© 2010 Carnegie Mellon University

C&C Site

- Beaconing
  - Regular contact
  - Outside of expected port range
  - Not otherwise explained
- Violations of locality
  - C&C site is a compromised host; not regular contact
- Violations of rationality
  - Not in customer/user area
  - Unexpected server location
  - Warning: Porn, low-end commerce, tourism sites
  - Peer-to-Peer overuse

© 2010 Carnegie Mellon University

C&C Channel

- IRC – not human-originated chat
  - Very short contacts (1-2 packets)
  - Somewhat regular contacts
  - Malware authors rely on encryption and obfuscation
- DNS – no follow-on contact activity
  - Disguise beacon as lookup
  - Disguise C&C message in body of response
  - But domain doesn't actually exist, so follow-on contact doesn't happen

© 2010 Carnegie Mellon University

Downloads

Not usually assorted with normal mechanisms

- Web, ftp, etc.

High-port to high-port connections

- Above 10,000
- With content
- Short duration
- From unusual contact

© 2010 Carnegie Mellon University

Executing Commands

- Outbound scanning
- Outbound Ddos
  - Frequency of contact
  - Diversity of ports/protocols
  - Site outside of locality
- Data exfiltration is file transfer
  - High-to-High ports
  - HTTP outbound data from client
- Fraud artifacts
  - Pay-per-view
  - Pay-per-click
  - Pay-per-install

© 2010 Carnegie Mellon University

External Sources of Situational Information

- CERT: http://www.cert.org
- CVE Database - http://www.cve.mitre.org/
- Internet Storm Center - http://isc.sans.org/
- SecureWorks Research & Threats - http://www.secureworks.com/research/
- DefCon - http://www.defcon.org/
- Microsoft Security - http://www.microsoft.com/security/default.mspx

© 2010 Carnegie Mellon University

Leading Questions

What service is being hit?

What protocol is being used?

Are there several protocols in use?

How many packets?

How large are the packets?

How long is a session?

Any beaconing?

Spreading phase vs. Quiet phase Vs. Active phase?

© 2010 Carnegie Mellon University



Networks in the Broad

© 2010 Carnegie Mellon University

What We Will Cover

- Introduction
- Your Network
  - Fundamentals of networks, flow, and protocols
  - Malicious traffic
- External Events & Trends
  - Malware
  - Networks in the Broad
- Working Together
  - Network dependencies
  - Analysis
- Summary

© 2010 Carnegie Mellon University

Why the Internet Works

Peering and Hosting

Partial compliance with standards

Partial trust

What this all means to network situational awareness (and your network)

© 2010 Carnegie Mellon University

Peering and Hosting

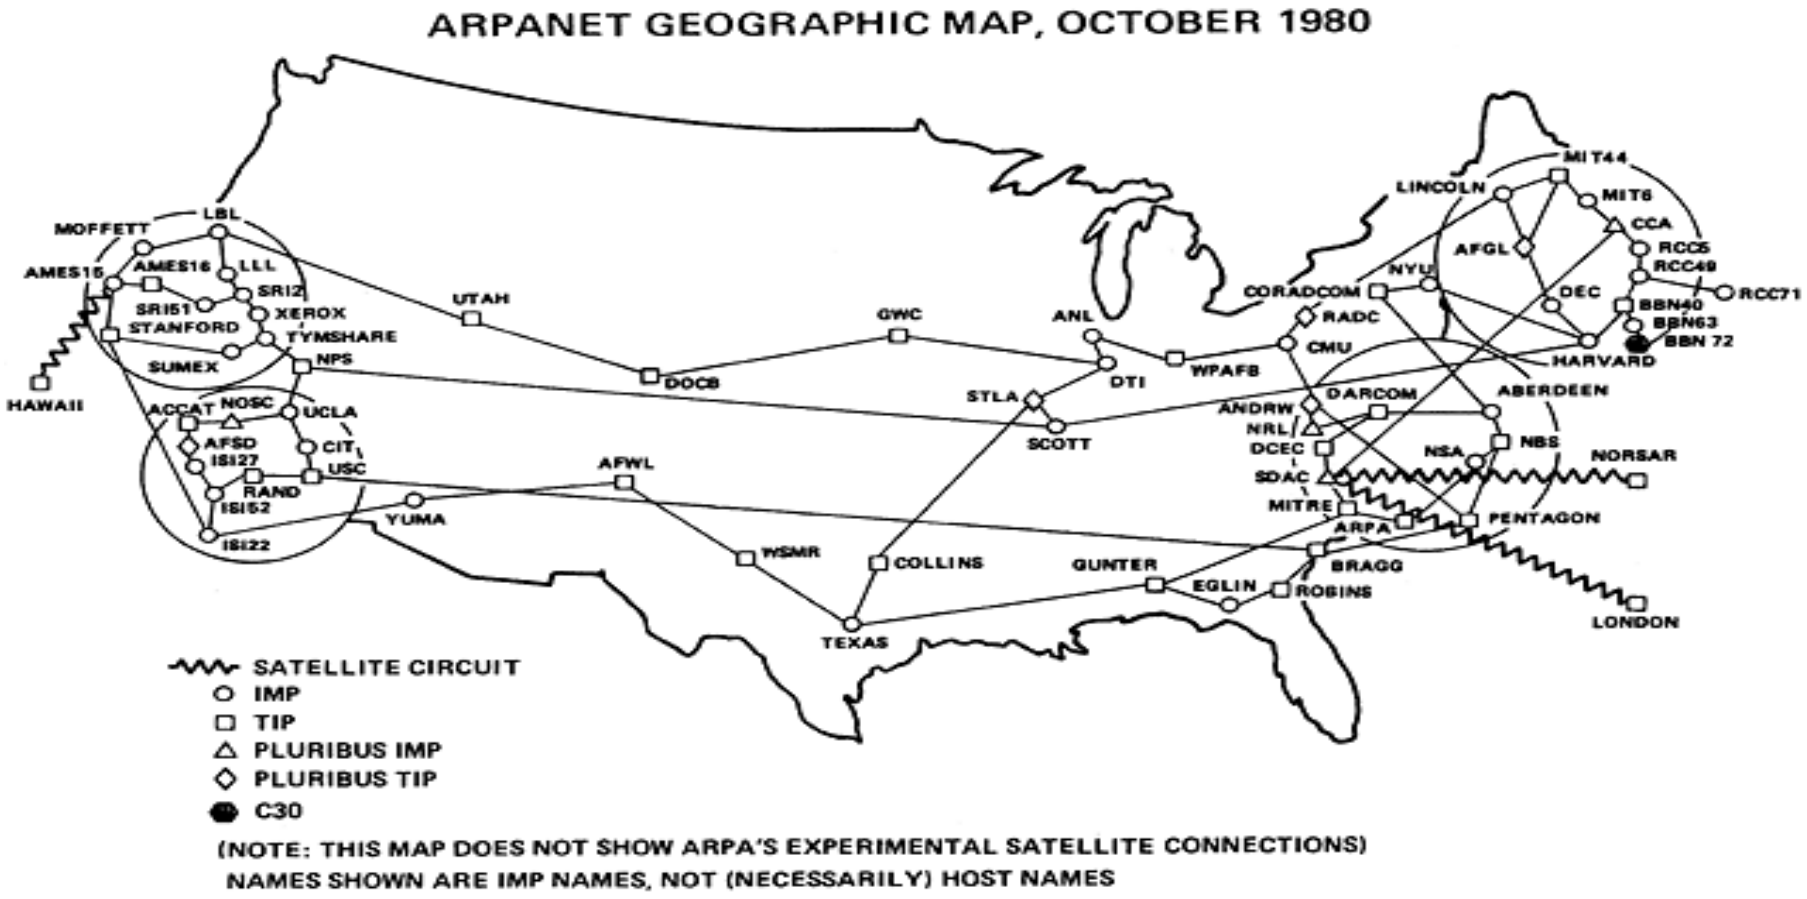
Hosting – contractual agreement to provide access

Peering – voluntary agreement to exchange access

How you contract is largely independent of how your information is routed.

© 2010 Carnegie Mellon University

The Old Net



© 2010 Carnegie Mellon University

Arpanet to Internet

Arpanet: Defense-related institutions and research universities ($$$)

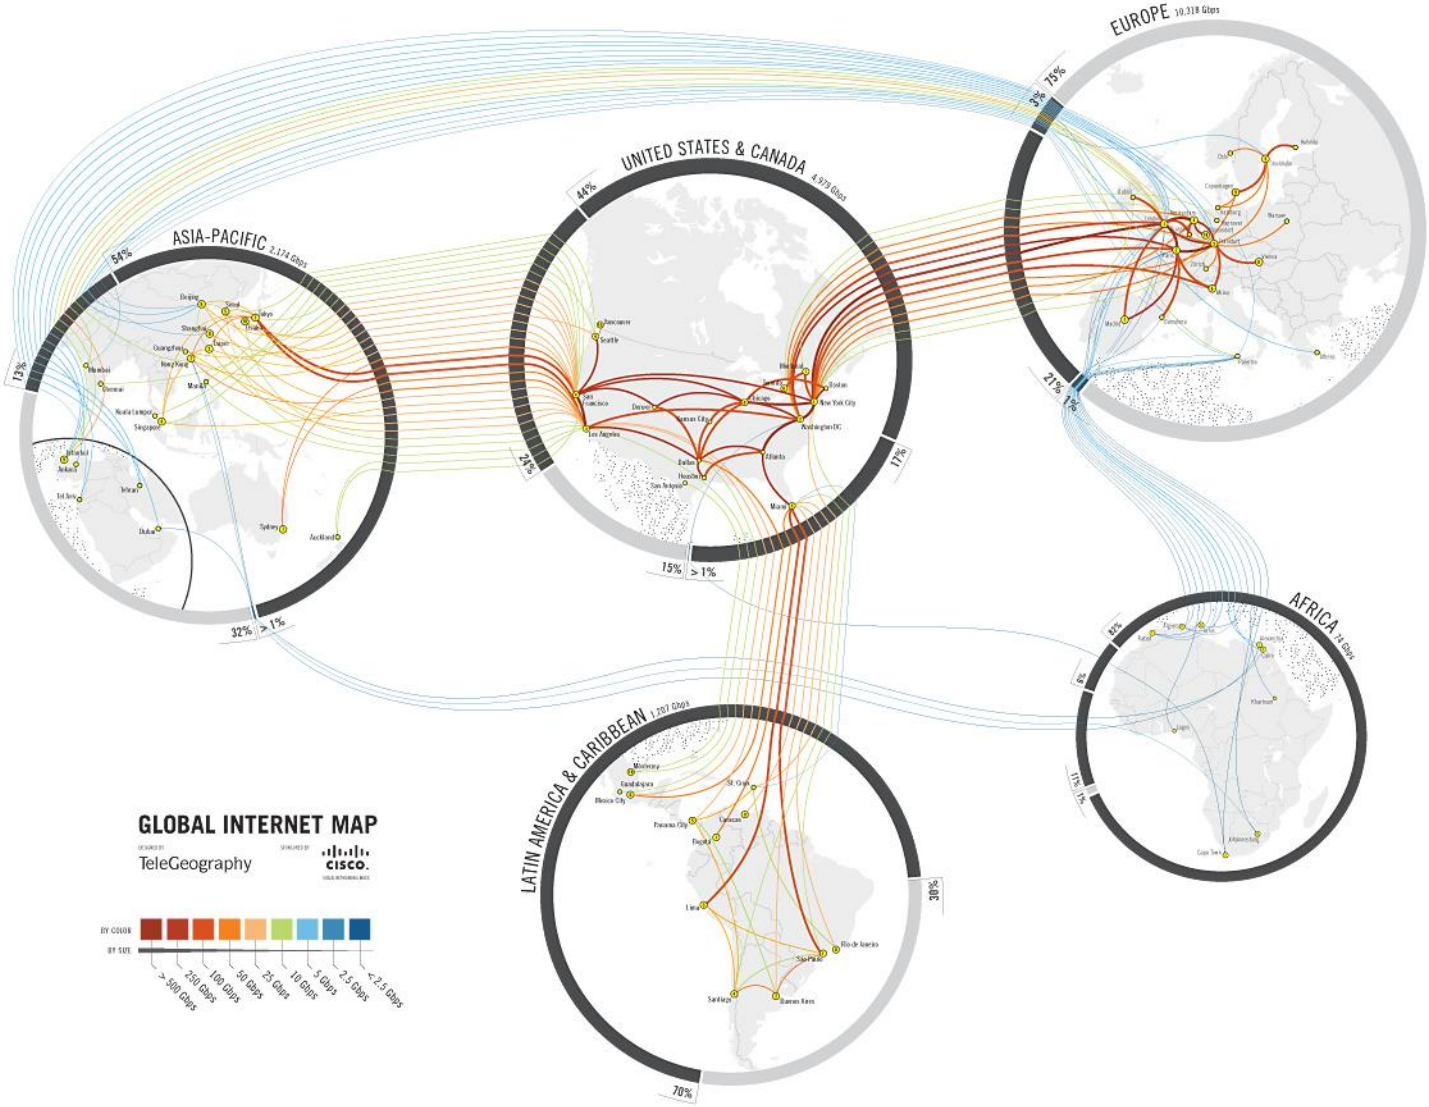
NSFNet: science-related institutions and research universities ($)

Usenet: peer-based institutions

Fleemco!ucvax!berkvax!gateway@berkley.arpa

Addressing schemes still referenced in RFCs

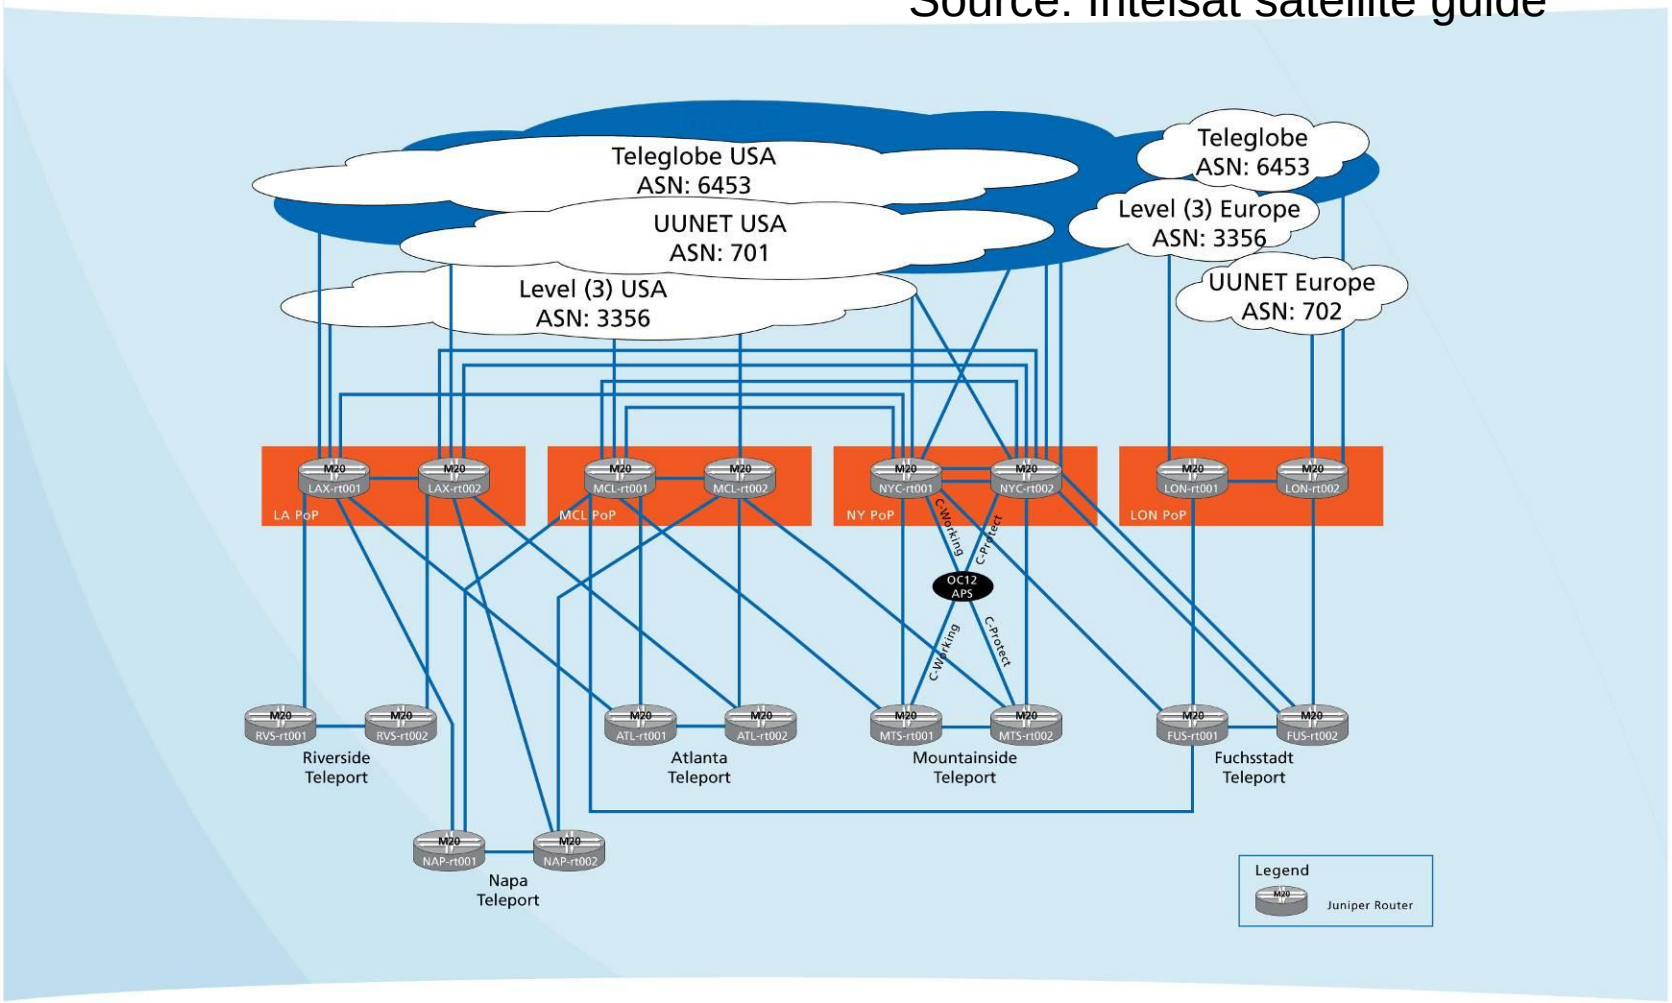
© 2010 Carnegie Mellon University



© 2010 Carnegie Mellon University

Intelsat GXS® IP Network

Source: Intelsat satellite guide



© 2010 Carnegie Mellon University

Partial Compliance with Standards

RFCs are NOT binding documents

- Descriptions of suggested practice
- Acceptance varies widely

NO authoritative body over users

- IANA can pull accreditation to offer names
- But institutions involved readily find other hosting
- NO standardization of host or network configuration

© 2010 Carnegie Mellon University

Partial Trust

Tom trust Freda, and Freda trusts Harry, but Tom doesn't need to trust Harry

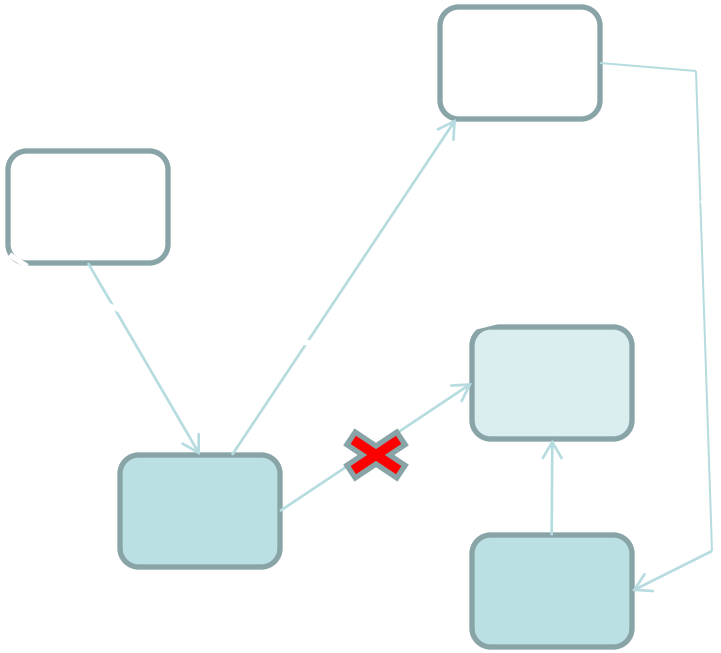
If a host uses an address, it is presumed to have the authority to use that address

Some of this is changing with VPNs, and with IPv6

© 2010 Carnegie Mellon University

What this means for Situational Awareness

- Weird traffic artifacts are common
- Weird routing behavior is common
- Routes change
- Redundant routing is the rule, not the exception
- How A talks to B isn't how B talks to A



© 2010 Carnegie Mellon University

Geolocation of IP Addresses

- “technique of determining a user's geographic latitude, longitude and, by inference, city, region and nation by comparing the user's public Internet IP address with known locations of other electronically neighboring servers and routers” (http://www.linuxjournal.com/article/7856)
- Series of Inference rules
- On-line database
- Registration information

© 2010 Carnegie Mellon University

Rules

Where is this address routed to?

Where is this address routed from?

What is the physical location of those routers?

Where have other hosts on the same network been sited?

Is there location-specific information available on that host? For how large a set of locations?

© 2010 Carnegie Mellon University

Geolocation Methods

- Database lookup
  - http://www.caida.org/tools/utilities/netgeo/ -- OLD
  - http://www.ip2location.com/
  - http://www.ipligence.com/geolocation/
- DNS lookup
  - host -t LOC yahoo.com
    yahoo.com LOC 37 23 30.900 N 121 59 19.000 W 7.00m 100m 100m 2m
- Traceroute, then look at where nearby routers are

© 2010 Carnegie Mellon University

86.140.13.8

(Thoughtfully sent to me by a spammer)

Netgeo: IANA reserved

Ip2location.com: British Telecom home office in London

Ipligence: Somewhere in the United Kingdom

No LOC record in DNS

- host86-140-13-8.range86-140.btcentralplus.com

© 2010 Carnegie Mellon University

Traceroute 86.140.13.8

```
. . .
 7  equinixexchange2.bt.com (206.223.115.140)  6.951 ms  7.176 ms  7.596 ms
 8  t2c1-p5-0-0.uk-eal.eu.bt.net (166.49.164.226)  88.313 ms  87.844 ms  87.264 ms
 9  166-49-168-22.eu.bt.net (166.49.168.22)  88.093 ms  87.684 ms  87.949 ms
10  core1-te0-10-0-0.ealing.ukcore.bt.net (62.6.200.109)  81.990 ms  81.976 ms  82.344 ms
11  core1-pos5-1.reading.ukcore.bt.net (194.74.65.190)  91.785 ms  90.783 ms  91.475 ms
    MPLS Label=23 CoS=5 TTL=1 S=0
12  core1-pos5-4.birmingham.ukcore.bt.net (62.6.204.1)  85.324 ms  85.857 ms  85.146 ms
13  bar4-pos8-0.birmingham2.broadband.bt.net (195.99.120.42)  85.313 ms  85.997 ms  84.909 ms
14  217.47.249.115 (217.47.249.115)  90.754 ms  90.986 ms  90.884 ms
15  217.41.172.45 (217.41.172.45)  85.263 ms  84.725 ms  86.857 ms
16  217.41.172.69 (217.41.172.69)  91.118 ms  90.829 ms  90.920 ms
17  217.41.172.73 (217.41.172.73)  92.357 ms  91.154 ms  91.784 ms
18  217.41.172.10 (217.41.172.10)  90.673 ms  92.228 ms  91.226 ms
19   * * *
```

Birmingham, England -> Hertfordshire England

30 miles to London, 90 miles to Birmingham

© 2010 Carnegie Mellon University

Case Studies

Dependencies are clearest when systems fail

Looking at a variety of causes

- Natural events
- Business decisions
- Nationalism/Political

What are the imlications for network situational awareness?

© 2010 Carnegie Mellon University

Natural Events

Storms

Disasters

Wear/Corrosion

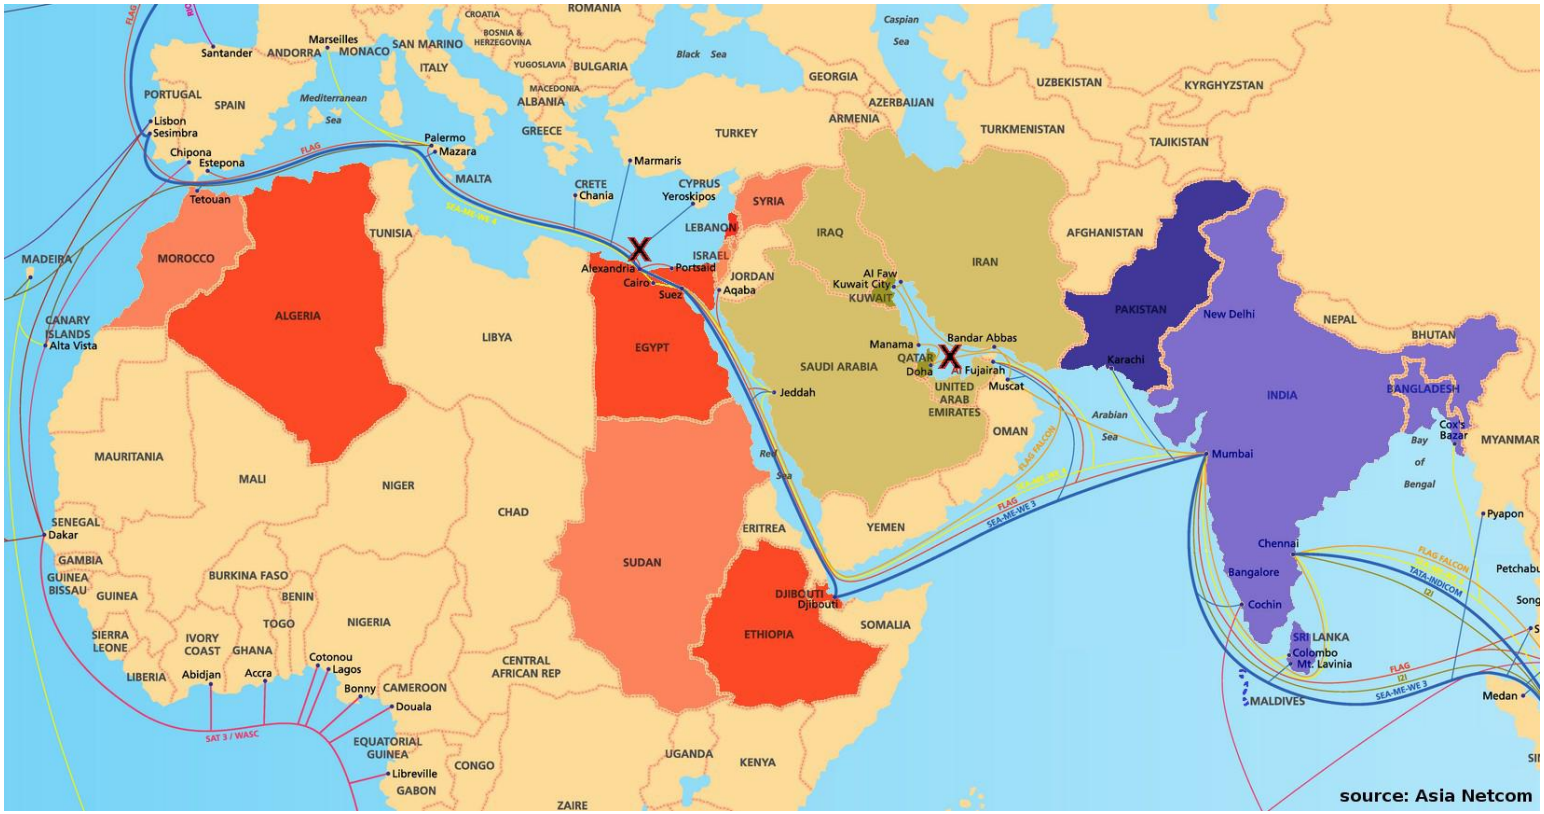
Frequency is not uncommon

Most common defense is redundancy

© 2010 Carnegie Mellon University

Cable Cut

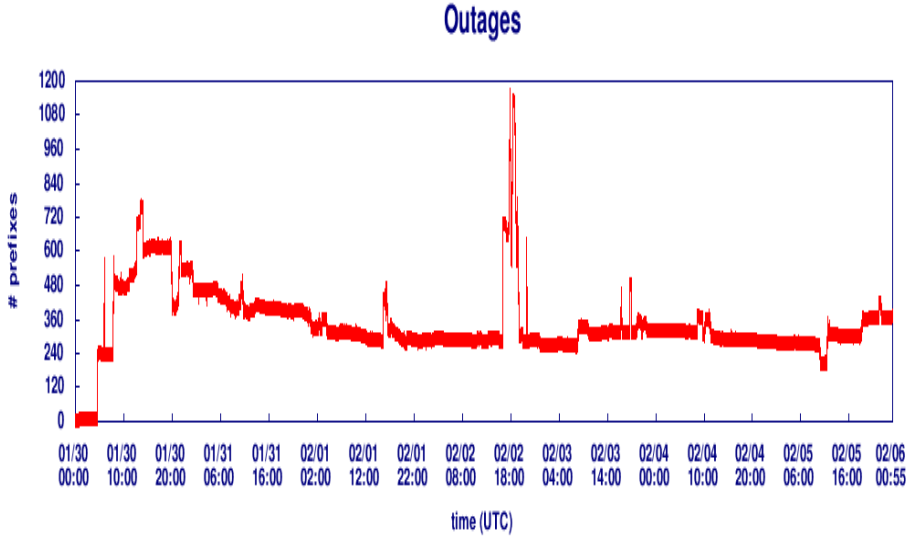
http://www.renesys.com/tech/presentations/pdf/menog3-cablebreaks.pdf



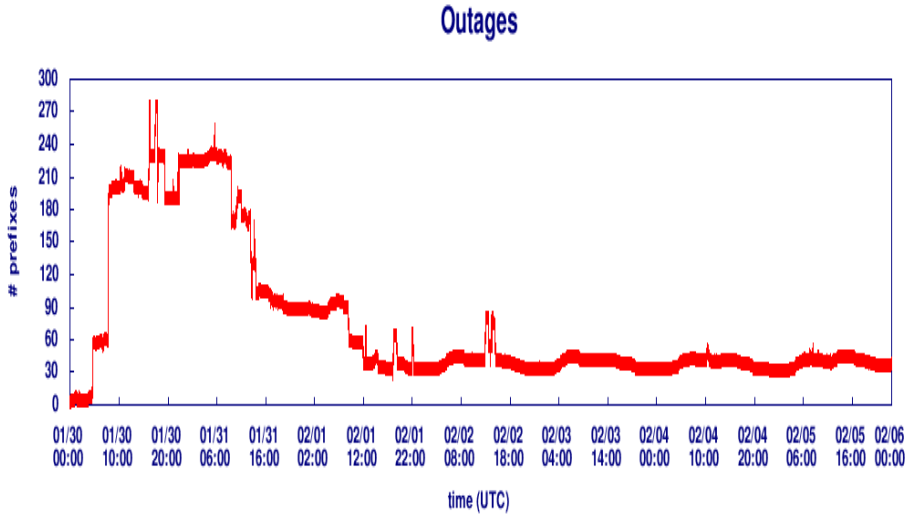
© 2010 Carnegie Mellon University

Impacts

Egypt



Kuwait



Egypt: 96% of prefixes suffered some outage (loss of routing on Internet)

- New Service providers
- Traffic Sharing among providers

Kuwait: 70% of prefixes suffered some outage

© 2010 Carnegie Mellon University

Business Decisions

Does loss of connectivity benefit anyone?

How do the business relationships work?

Where are business relationships most vulnerable?

How is redundancy a business asset?

How is redundancy a business liability?

What does this mean for network situational awareness?

© 2010 Carnegie Mellon University

De-Peering

http://www.renesys.com/tech/presentations/pdf/nanog43-peeringwars.pdf

- Internet is a voluntary network
  - Transit: For-cost service
  - Peering: Voluntary interconnection
  - Tier 1: No transit providers, only peering providers and customers
    - All Tier 1's have to peer, or Internet is divided
- What is volunteered can be withdrawn
  - March 13: Cogent de-peers Telia (restored March 28)
  - Several thousand prefixes affected in multiple nations

© 2010 Carnegie Mellon University

Cyber-Attacks

- Cyber War vs. Net War
- Advantages: Cost, effectiveness, deniability, recoverability
- Disadvantages: Escalation, retaliation, collateral damage
- Examples:
  - Chechnya: Russian hacker groups
  - Estonia: Nationalist gone cyber
  - Georgia: Mentored hackers

© 2010 Carnegie Mellon University

Cyber War Evaluation

- Where does awareness end?
- How does an organization manage its vulnerability?
- Can the organization manipulate the cost model for the attackers?
- Can the attackers manipulate the cost model for the organization?
- What does this mean for situational awareness?

© 2010 Carnegie Mellon University



Mission and Internet Dependencies

© 2010 Carnegie Mellon University

What We Will Cover

- Introduction
- Your Network
  - Fundamentals of networks, flow, and protocols
  - Malicious traffic
- External Events & Trends
  - Malware
  - Networks in the Broad
- Working Together
  - Network dependencies
  - Analysis
- Summary

© 2010 Carnegie Mellon University

What is Mission?

What is the reason your organization exists?

What are the values that motivate your organization?

What activities does your organization need to perform?

What services does your organization need to exhibit?

© 2010 Carnegie Mellon University

Mission and Network Design

Services

Servers

Quality of service

Points of presence

Access points

Acceptable degradation of services

© 2010 Carnegie Mellon University

Mission and Internet Usage

Strategy of Internet usage

Point of presence sizing

External firewall / Proxies

© 2010 Carnegie Mellon University

Mission and Situational Awareness

What do you need to watch?

- Host vs. Network
- Key Servers
- Key Clients
- Key Routers

How do you need to watch?

- Real-time
- Retrospective

© 2010 Carnegie Mellon University

Case Study

Large meeting, held at intervals greater than one year

Involves numerous individuals of national prominence

- Public presentations and votes
- Private meetings / decision-making sessions

Strong social impact at national level

High degree of visibility

Always held in large metropolitan area

Always draws strong protest events

Disruption/failure of event has social, political, and economic impacts

Event operations dependent on local infrastructures (telecomm, transportation, energy)

Heavily IT dependent

© 2010 Carnegie Mellon University

Event Production

Local organizers (task force for event only)

- Staff-up / Set-up / Event / Teardown / Staff-down over 1+ year span
- Little or no contact with previous event task force

Contractors/Consultants/Vendors

Venue (infrastructure owners -- assume Major Event Center)

- operations staff
- Access control
- Security Policy
- Additional contractors/vendors

National Organization

- Staff
- Participant source (event attendees and speakers)
- Funding Sponsors: extra services, priority access, no interest in security
- No IT Security Policy for event
- Views event as source of key decisions, important publicity

© 2010 Carnegie Mellon University

Event Protesters

Opposition/protest activity can be significant and possibly violent

Awareness by opposition groups of the impact of IT as a means of disruption

Violence against supporting infrastructures

Non-violent interference with event (shouting, physical intrusion, computer mischief, dramatics)

Identification of “single points of failure”

Right to peaceful protest must be respected

© 2010 Carnegie Mellon University

Media

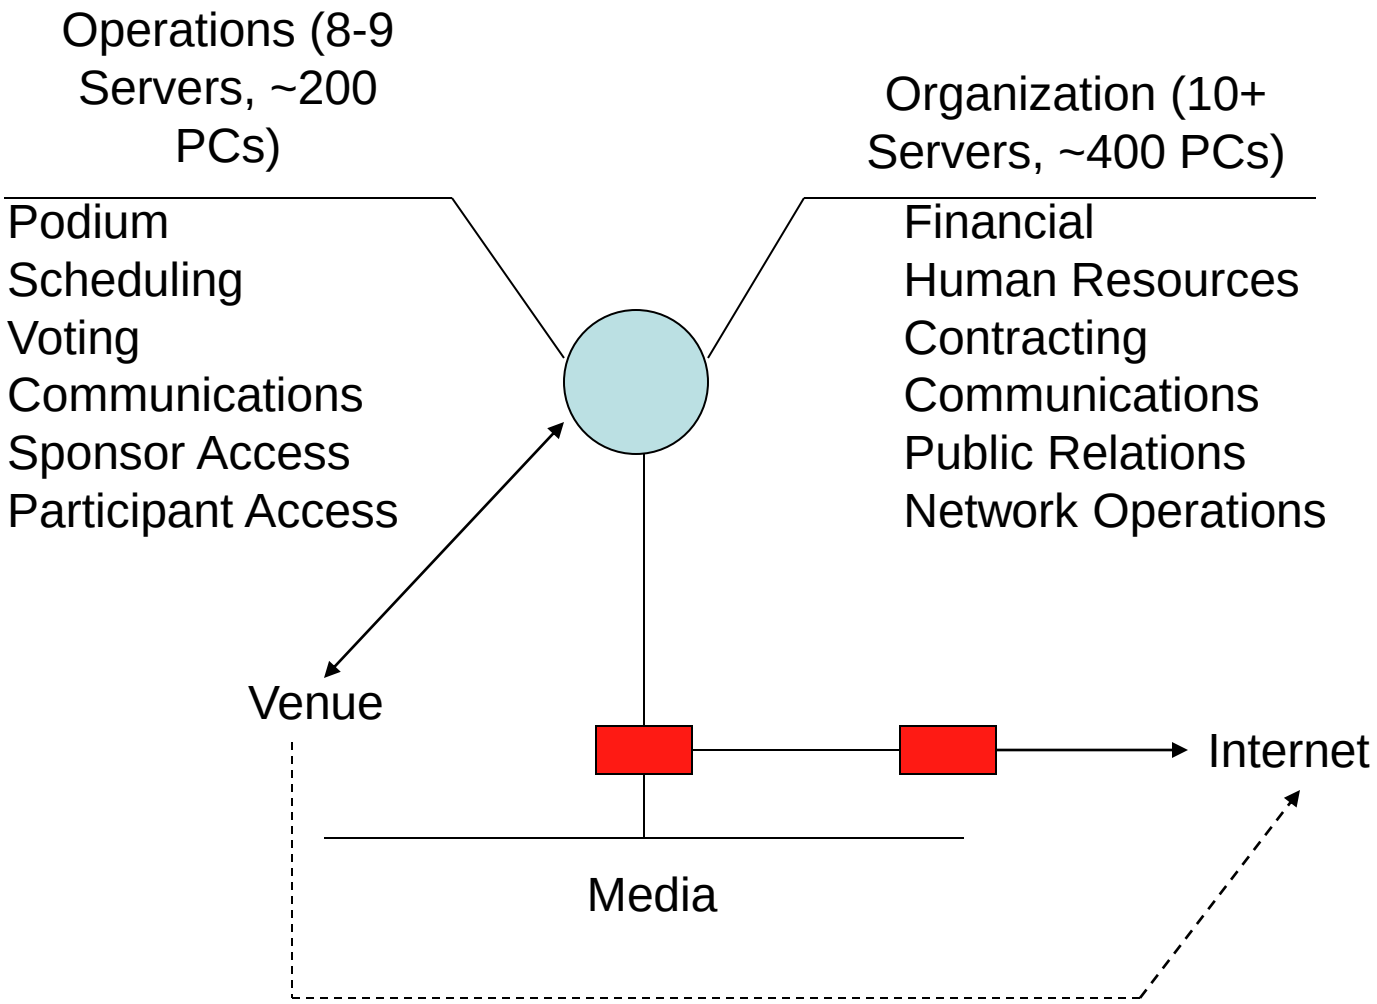
print/radio/television

local/national/international

- Access to event (limited screening: Press freedom)
- Access to information
  - Deadline pressure
  - Privacy compromise
- Source of valued publicity

© 2010 Carnegie Mellon University

Event Network



© 2010 Carnegie Mellon University

High Concern Risks (1)

Loss of Event/Critical time

- Time at event is short, highly contended
- Loss of infrastructure
- Loss of communications

Loss of Critical personnel

- Physical threat
- Embarrassment / Compromise
- Lack of respect for stakeholders

Loss of trust in decisions made at event

- Can include public exposure of sensitive data
- Compromise of voting systems

© 2010 Carnegie Mellon University

High Concern Risks (2)

Loss of privacy, yielding decreased impact on event, decreased participation with organization

- Individuals of prominence can lose privacy – can include a physical security risk (schedules, timetables, etc.)
- Such a loss may not directly impact event – impact delayed
- Can result in loss of Sponsorship, financial support, public perception of competence

Loss of control of convention flow due to disruption (schedule trust loss)

- DDOS can result in loss of critical services – can impact success of meetings
- Intrusion can result in false scheduling, loss of public confidence, loss of trust in IT integrity

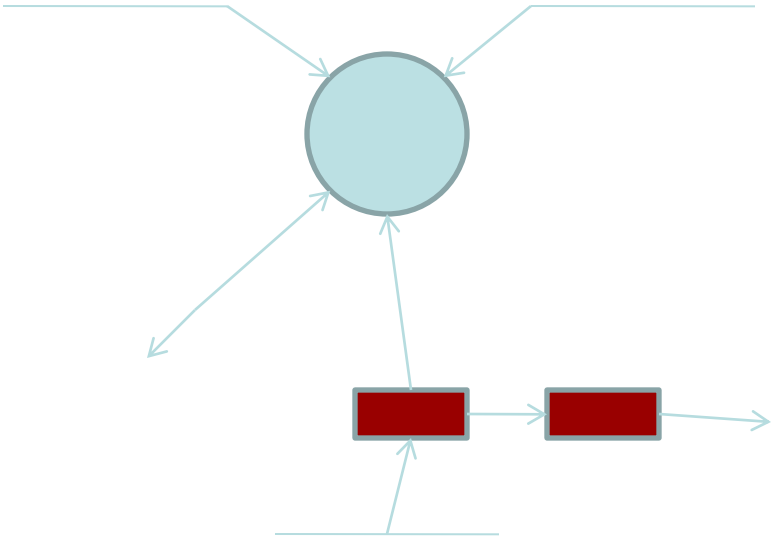
Loss of financial control of convention

- Intrusion or DDOS can result in loss of critical financial sponsorship
- Result can include additional financial loss to recover from malicious intrusion or attack
- Loss of public confidence can limit economic/financial support which is considered critical to ultimate success of meetings

© 2010 Carnegie Mellon University

Situational Awareness

- What is the mission?
- What is the need for situational awareness?
- What needs to be monitored?
- How does it need to be monitored?



© 2010 Carnegie Mellon University

Internet Dependencies

Insourcing

- What don't we use the Internet for?
- What situational awareness needs for internal connectivity?

Outsourcing

- VPN
- Trust relationships
- Exposures

© 2010 Carnegie Mellon University

Exporting and Importing Services

What do we need the Internet for?

What should we make someone else's problem?

When should we take on someone else's problem?

What does this imply for Situational Awareness?

© 2010 Carnegie Mellon University

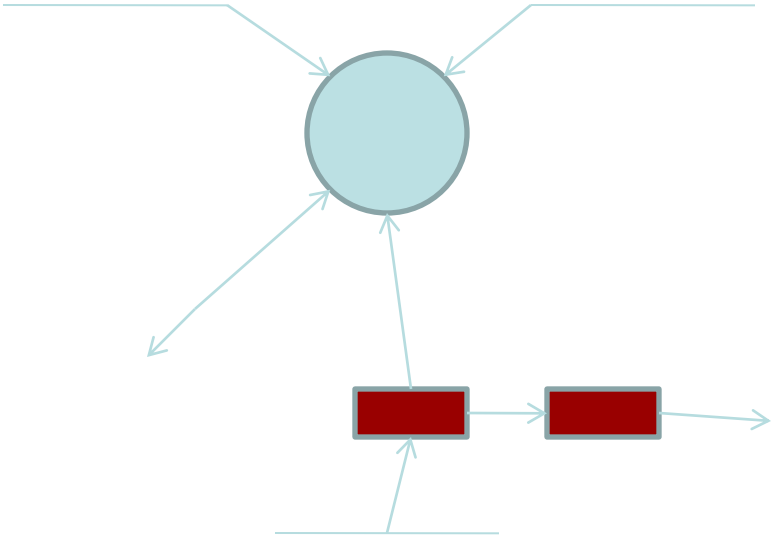
Case Study

What do they need from the Internet?

What does the Internet need from them?

What should they outsource?

What mustn't they outsource?



© 2010 Carnegie Mellon University



Doing Analysis

© 2010 Carnegie Mellon University

What We Will Cover

- Introduction
- Your Network
  - Fundamentals of networks, flow, and protocols
  - Malicious traffic
- External Events & Trends
  - Malware
  - Networks in the Broad
- Working Together
  - Network dependencies
  - Analysis
- Summary

© 2010 Carnegie Mellon University

Building an Analysis Team

Define your mission

- Develop network situational awareness
  - Security
  - Availability
  - Non-traditional threat scenarios
- Know your network,
  Know external events & trends,
  Know how they work together

© 2010 Carnegie Mellon University

Building an Analysis Team (2)

Define your needed skill sets

- Layer 2: Networking
- Layer 3 & 4: Ports and Protocols
- Layer 5 - 7: Applications
- Layer 9: Management

Also remember soft skills

- Team building & Communications
- Research & Record Keeping

© 2010 Carnegie Mellon University

How would they help?

- Security Architect
- Network Engineer
- Programmer
- Statistician
- Tech Writer
- Mathematician
- Manager
- Team Leader
- Mentor

© 2010 Carnegie Mellon University

Performance Metrics

Are they just a management nightmare?

Why do you need metrics?

- Justify existing staff (and staff increases)
- Keep the threat visible
- Senior management sound bites
- Justify expenditures

© 2010 Carnegie Mellon University

Good Performance Metrics

Baseline Inventory ("Know your network")

- Services Offered
- Services Consumed

Noise Level

- High volume events that can be quantified, normalized and compared
- Scan counts; vuln sweeps; brute force attempts
  - Compare with open source metrics for the same event
  - Do not combine discrete event types

© 2010 Carnegie Mellon University

Good Performance Metrics (2)

- Near misses
  - Evaluate your defense-in-depth architecture
  - Virus hits; drive-by-download misses; policy violations
- Hits
  - Should be rare
  - List, but avoid summarizing
    - individual hits can not be compared with each other
- Work load
  - Overtime hours for team members

© 2010 Carnegie Mellon University

Calibrating Your Infrastructure

Can an alarm tell you if a theft has occurred?

Rely on out-of-band management

- Leverage existing solutions built for other needs

© 2010 Carnegie Mellon University

What will you never know?

Total ground truth on a network of any size

How often is bad considered good? (false negative)

What is the next attack?

Why did they attack you?

What are your competitors seeing?

© 2010 Carnegie Mellon University

Inherently Partial Data

Technology shifts

Attacker actions

Defender actions

Managerial decisions

Network bandwidth

© 2010 Carnegie Mellon University

Correlation and Causation

Baseline in dynamic environment

Correlation vs. Causation

Implications

- Need to be cautious in kinds of conclusions
- Consider strategies for dealing with analysis gone wrong

© 2010 Carnegie Mellon University

Indication and Proof

Indication: There is reason to believe

Proof: There is no other logically defensible explanation

How much confidence do you need?

Cost of false positive?

Cost of false negative?

© 2010 Carnegie Mellon University

Clustering and Extrapolation

- Clustering groups reports into meaningful classes
- Similarity metric applied to common features
  - Cohesion function calculates degree of similarity
  - Clustering generates overlapping clusters (clumps)
    - Minimizes cohesion function betweens incident sets
- Extrapolation fills in the reporting gaps
  - Extrapolation criterion establishes when and how
- Generates extrapolated incidents (x-incidents)

© 2010 Carnegie Mellon University

Correlation and Abduction

- Identifies sequences that constitute staged attack
  - Generates x-incident chains
  - Starting context establishes understanding of initial system/network configuration
- Causal relationships through pre-/post-condition chaining
  - Precondition of first incident must satisfy starting context
  - Postcondition of each incident must satisfy precondition of the subsequent incident
- Techniques available (abduction) for filling in gaps
  - Strings together x-incident chains using attack patterns
  - Abduction criterion establishes when and how

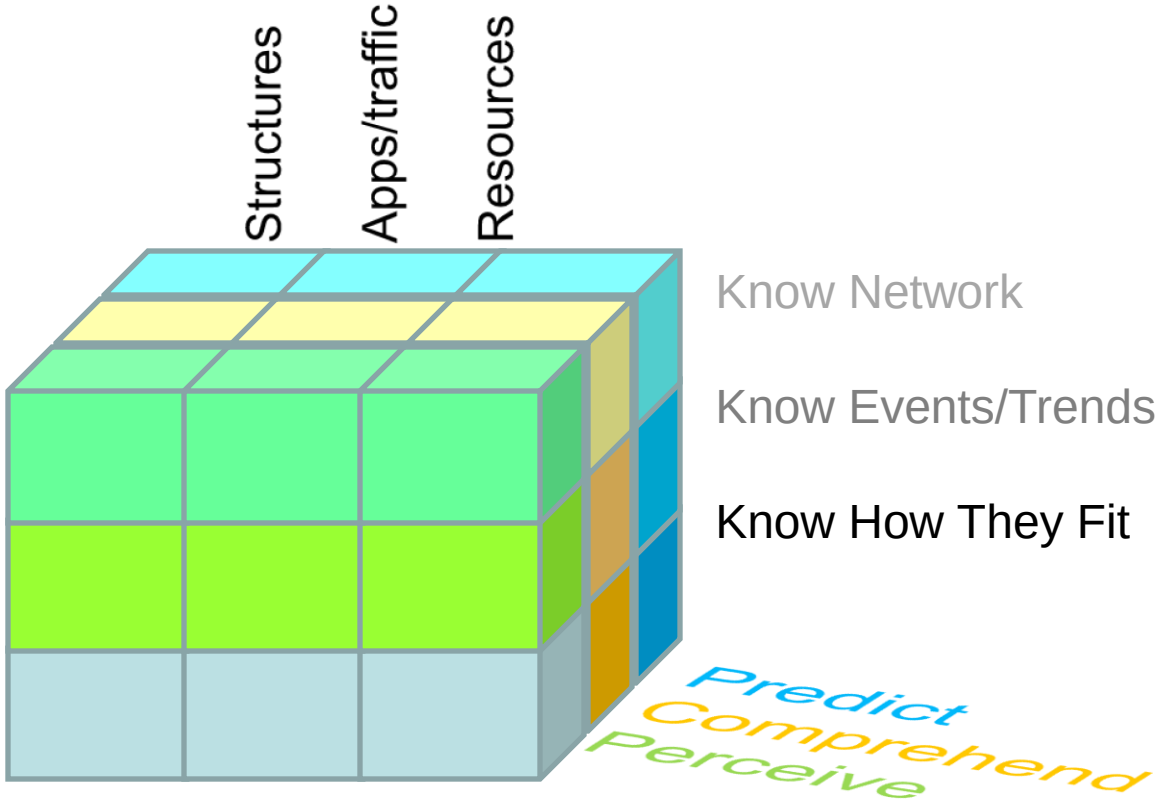
© 2010 Carnegie Mellon University

What We Will Cover

- Introduction
- Your Network
  - Fundamentals of networks, flow, and protocols
  - Malicious traffic
- External Events & Trends
  - Malware
  - Networks in the Broad
- Working Together
  - Network dependencies
  - Analysis
- Summary

© 2010 Carnegie Mellon University

Where we have been



© 2010 Carnegie Mellon University

Governance Questions

How do I know what I'm looking for?

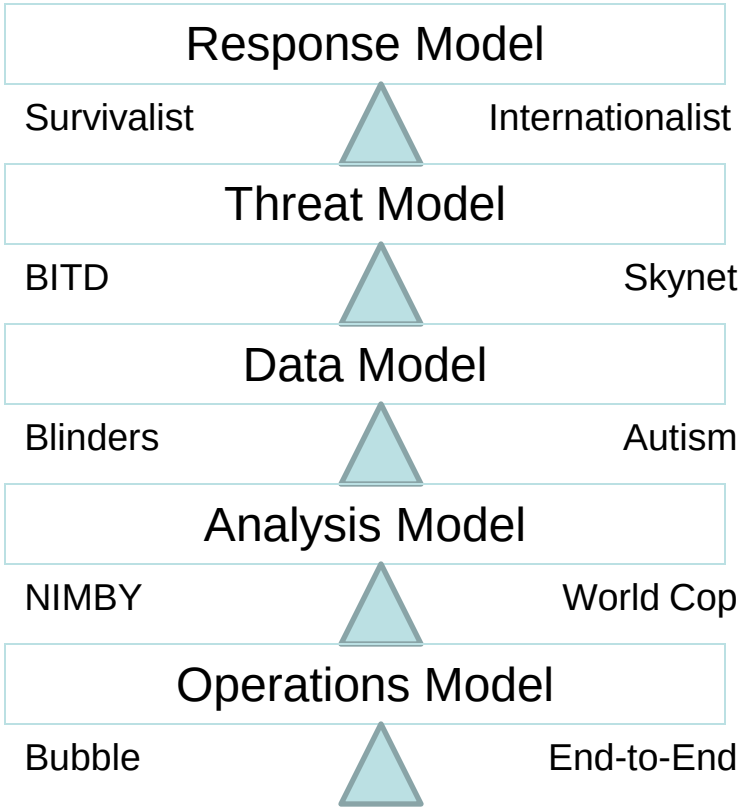
How do I know why I'm looking?

How do I know where to look?

How do I know when it's found?

© 2010 Carnegie Mellon University

NetSA Governance



© 2010 Carnegie Mellon University

Future of NetSA

- Analysis coalitions (FloCon on steroids)
- NetSA as a service
- Diversity of specializations
  - Malware vs. Inventory vs. Architecture
  - Diagnose vs. Track/Trace vs. Profile
  - Packet level vs. Trace level vs. Flow level vs. Organizational
- More approachable formalisms
  - Visualizations with meaning
  - Drill down with extensibility
  - Better handling of time

© 2010 Carnegie Mellon University

Conclusions

NetSA is new as a systematic approach at scale

NetSA is changing (arms race, technology)

- Need awareness WITHOUT total information
- Need awareness WITH confidence and timeliness
- Need awareness WITHOUT global state
- Need awareness WITH extensibility

We live in interesting times…

© 2010 Carnegie Mellon University

Parallel Processing in Netflow Data Fusion

- George Saylor
- Michael Rash
- G2, Inc.

- Flocon, 2010



Agenda

- Setting the stage – a description of factors
- Environment
- Parallel processing to facilitate:
  - Fusing large netflow dataset with other large datasets
  - Joining netflow and other data to enrichment data such as IP reputation and GeoIP lists
  - Incident investigation on long time ranges using fused datasets



Agenda cont'd

- Specific use cases
  - Social Networks for Computers and Networks
  - Long Term Data Exfiltration
  - Comparing netflow, system characteristics, and time to determine impact
  - DNS Fast Flux detection and associated botnet detection
  - SPAM behavior on a given network



Setting the stage

- Targeting very large environments (ISP)
  - > than 2 trillion netflow records per year
  - Existing large siloed data stores of:
    - Large netflow database
    - Assets – with history
    - System characteristics – with history
    - DNS queries/responses
    - Large store of PCAP data



Setting the Stage cont'd

- Goals
  - Fuse this data together
  - Add enrichment data
    - IP reputation data/Black Lists/White Lists
    - Malware domains
    - Google Safe Browsing API
    - Spamhaus
    - Dshield
    - GeoIP



Setting the Stage Cont…

- Goals Cont…
  - Unified query interface (instead of scripting)
  - Ability to run against big datasets (without hurting anyone else)
  - Iterative discovery by pivoting on results of previous queries
  - Large scale summary and computation
  - Beacon detection over long time ranges
    - Correlating with other indicators
  - System Level Social Network Analysis



Environment

- Selected a computing cluster that was optimized for:
  - Very large joins
  - Large data searches
  - Fast indexed access to data
  - Index building does not affect end users
    - Index builds on one system, distributes to queried nodes
  - Ability to partition on values and ranges



Use Case: Social Network Growth

- Social Networks for Computers and Networks
  - Given an IP address or addresses, or network range characterize relationships between those and other computers or networks
- Scenario:
  - An IDS alert provides some indication that a given host is compromised
  - We may want to know who this machine talks to
  - Whether the machine's "friends list" (other systems) is growing rapidly



Use Case: Data Exfiltration Discovery

- Social Network prioritization based on bytes transferred
- Scenario:
  - An IDS alert provides some indication that a given host is compromised.
  - Cross reference attacker's IP with the GeoIP and other reputation databases.
  - Construct the social network of other systems this IP communicates with.
  - Sort by bytes transferred – the top systems in this list may be candidates for data exfiltration nodes.



Use Case: SPAM Node and C2 Discovery

- SPAM nodes and mail server error code thresholds
- Scenario:
  - IDS signatures for thresholding large numbers of recipient unknown error codes sent back from mail servers to local systems.
  - Build system level social network concentrating on command and control communications and other observables.
  - Try to identify command and control connections – may become obvious with overlapping external IP's in SPAM node

Use Case: DNS Fast Flux

- Storing large number of DNS requests and responses
- Scenario:
  - Use thresholding operation for the number of IP addresses given in response to DNS requests for the same hostname.
  - Also use low TTL values in DNS responses as a search metric.
  - Look for “unusual” DNS hostnames that use multiple numbers and strange patterns (e.g. “f4n3upyhqj.com”).



Use Case: Complex Fusion

- Mixing together numerous data sources to gain better situational awareness
- Scenario:
  - Retrieve all flows where:
    - For a given IP list (derived from previous exercises for example).
    - Platform is Windows XP and the patch level is below a minimum level or there is no netflow record to the Windows Update service to indicate that it has been patched.
    - Exhibits communications to disreputable hosts and/or suspicious geo locations.
    - Compute this via a single interface.



Questions?

george.saylor@g2-inc.com
michael.rash@g2-inc.com





CENTER FOR APPLIED CYBERSECURITY RESEARCH

INDIANA UNIVERSITY

Pervasive Technology Institute

Geography of Internet2 Netflow

David A. J. Ripley

Indiana University, Center for Applied Cybersecurity Research.

Tony H. Grubesic

Indiana University, Department of Geography.

Timothy C. Matisziw

University of Missouri, Department of Geography/Department of Civil and Environmental Engineering.

- ***geography***
- **1. a.** The science which has for its object the description of the earth's surface, treating of its form and physical features, its natural and political divisions, the climate, productions, population, etc., of the various countries.

[Oxford English Dictionary, 2nd Edition, 1989.]

- Historically – spatial and temporal movement of goods, money, people and *information*.
- The Internet doesn't move goods or people directly, but it can influence their movement.
- The Internet transports information.
  - Information ≈ Money

- Access and Awareness
  - What kinds of data are available?
  - How are they acquired, from whom?
  - What are their strengths and limitations?

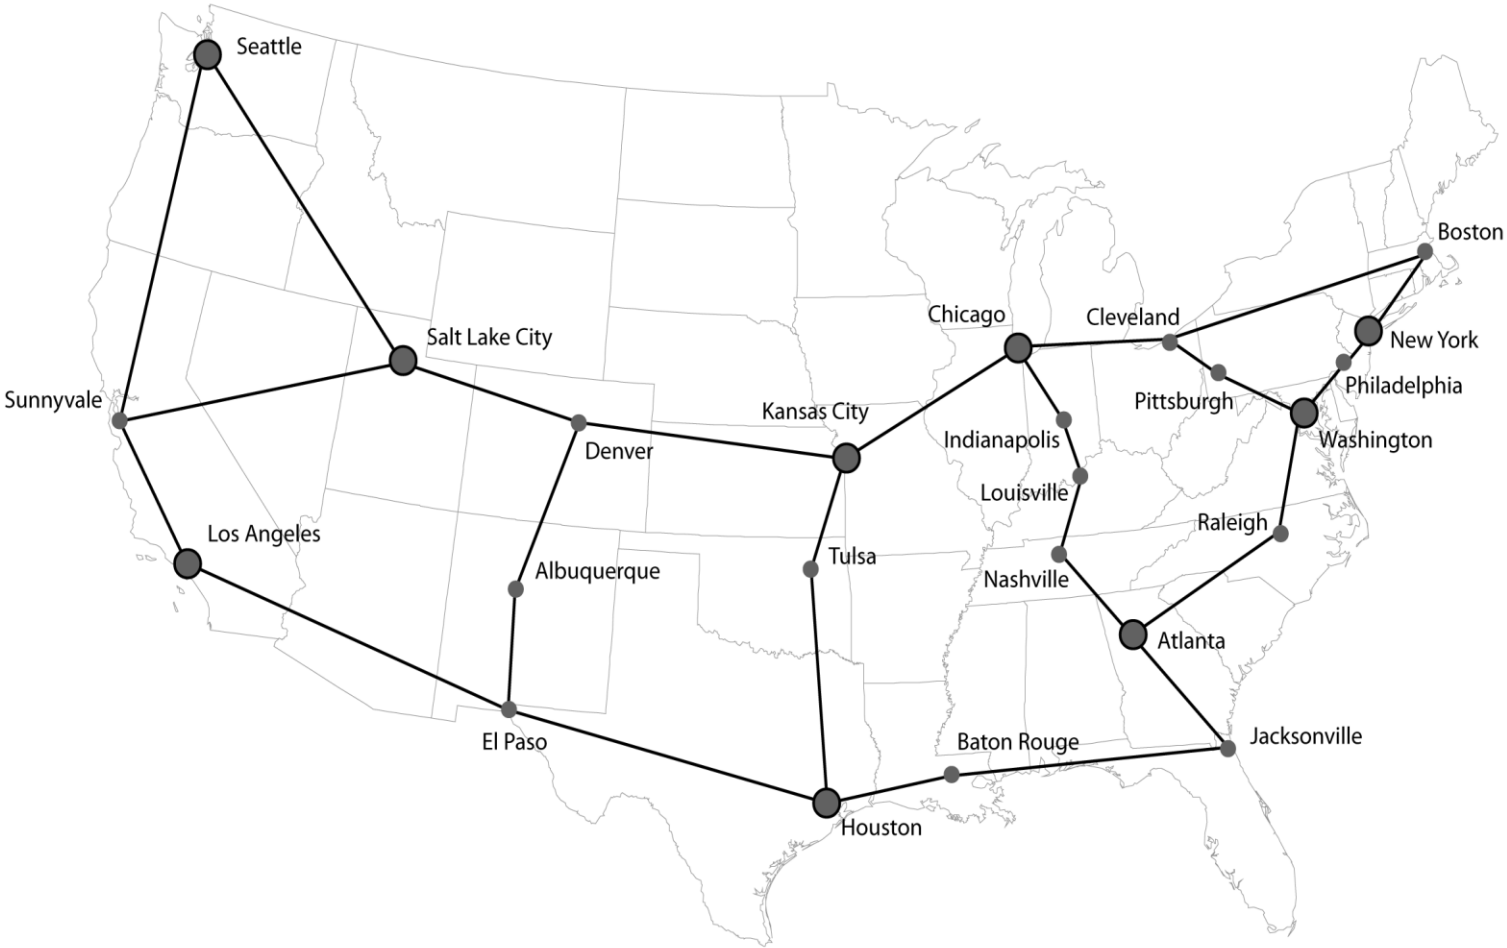
- Unfamiliar territory
  - Opportunities!

- Interest in Internet traffic volumes
  - Backbone traffic levels not interesting enough
  - IP-level a bit too specific
    - Policy/access difficulties
  - Interested in relative connectedness (and volume of traffic) between locations.
    - For certain values of “location.”

- Internet2 Network Netflow records:
  - Sourced from 9 core routers
  - 100:1 Sampling
  - /21 Anonymization required
  - ASN is not present/reliable
  - Includes CPS
    - So it's not just academic traffic.
  - Typically approx 2 billion records per day.

CENTER FOR APPLIED
CYBERSECURITY RESEARCH

INDIANA UNIVERSITY
Pervasive Technology Institute



- Needed some way to identify source and destination locations
  - The policy-mandated /21 anonymization presents a problem.
  - As does the lack of ASN

- How much of a problem is the anonymization?

- IP Anonymization – how much of a problem is it?
  - Surprisingly, not quite as much as we thought.
  - We're not interested in specific hosts.
- 12,000 prefixes announced on I2 Network;
  - 36% are 21 bits or shorter
  - Assuming constant utilization*, that's 98% of addresses unambiguously identified to the level of an AS.

- Admittedly, there are some assumptions that definitely need to be tested.
- But still not as bad as we initially thought it might be.

- ASN is a reasonable identifier for our purposes.
  - It may be a very poor way to determine location.
    - Local vs. national organizations; IU vs. Comcast.
  - But it gives us a location for the flow of information/money/whatever
    - Which may not have anything to do with the location of the end systems.

- Scripted WHOIS lookups
  - Registering entity's name and postal address.
- Geocoding
  - Most GIS packages will do this, with greater precision than we were interested in.
  - Centroid of ZIP code areas.

- 24 hour period
- ~2 billion flow records

- 175,000 unique sIPs, 232,000 unique dIPs
- 29,709 ZIP code areas.

- Aggregate traffic hourly by sZIP/dZIP pair.

- Capturing Connectivity
  - *Degree of node* - simple measure of connectivity – number of connections from one node to another.
  - Weight with traffic volume?
  - Calculate for both ingress and egress
    - Although there are issues here with NetFlow

CENTER FOR APPLIED
CYBERSECURITY RESEARCH

INDIANA UNIVERSITY
Pervasive Technology Institute



Table 3: Top 25 ZIP Code Areas for Averaged Ingress and Egress Flows

ZIP Code Area	City	State	Avg. Flow (Bytes)
98052	Redmond	WA	167253553.9
94043	Mountain View	CA	93077533.8
85281	Tempe	AZ	74199354.0
08054	Mount Laurel	NJ	67033097.2
20171	Herndon	VA	40396656.5
94089	Sunnyvale	CA	38924803.7
02142	Cambridge	MA	22123045.8
20910	Silver Spring	MD	19639697.8
30319	Atlanta	GA	18881723.6
08002	Cherry Hill	NJ	18804546.6
94129	San Francisco	CA	18488762.3
60510	Batavia	IL	15503621.5
95125	San Jose	CA	15380908.5
08542	Princeton	NJ	13317164.7
90025	Los Angeles	CA	12992821.1
53706	Madison	WI	11362748.7
78758	Austin	TX	11270133.0
94301	Palo Alto	CA	10583457.6
94720	Berkeley	CA	9175479.5
65201	Columbia	MO	9161357.3
02139	Cambridge	MA	8800400.8
98105	Seattle	WA	8490026.4
08540	Princeton	NJ	8344733.5
77002	Houston	TX	8334969.0
77843	College Station	TX	8202075.0



CENTER FOR APPLIED
CYBERSECURITY RESEARCH

INDIANA UNIVERSITY
Pervasive Technology Institute

Top 25 ZIP Code Areas for Averaged Degree of Node Measure

ZIP Code Area	City	State	Average Connections
97403	Eugene	OR	3016.5
47907	West Lafayette	IN	2981.5
50011	Ames	IA	2922.0
52242	Iowa City	IA	2894.0
47408	Bloomington	IN	2798.5
65201	Columbia	MO	2718.5
02138	Cambridge	MA	2697.0
61801	Urbana	IL	2633.5
06269	Storrs-Mansfield	CT	2597.5
68588	Lincoln	NE	2579.0
48824	East Lansing	MI	2549.5
43201	Columbus	OH	2524.5
02215	Boston	MA	2518.0
02139	Cambridge	MA	2516.0
97331	Corvallis	OR	2424.0
06105	Hartford	CT	2385.5
43212	Columbus	OH	2366.0
50131	Johnston	IA	2364.0
70803	Baton Rouge	LA	2287.5
46556	Notre Dame	IN	2281.5
46202	Indianapolis	IN	2246.5
45221	Cincinnati	OH	2239.5
40475	Richmond	KY	2220.5
48104	Ann Arbor	MI	2213.5
74078	Stillwater	OK	2208.0



CENTER FOR APPLIED
CYBERSECURITY RESEARCH

INDIANA UNIVERSITY
Pervasive Technology Institute

Top 10 ZIP Code Areas for the Origin-based Degree of Node Measure

ZIP Code Area	City	State	Unique Connections
97403	Eugene	OR	4073
02139	Cambridge	MA	4065
02138	Cambridge	MA	3966
65201	Columbia	MO	3845
47907	West Lafayette	IN	3838
68588	Lincoln	NE	3792
61801	Urbana	IL	3703
47408	Bloomington	IN	3700
50011	Ames	IA	3697
52242	Iowa City	IA	3696

Top 10 ZIP Code Areas for the Destination-based Degree of Node Measure

ZIP Code Area	City	State	Unique Connections
50011	Ames	IA	2147
47907	West Lafayette	IN	2125
48824	East Lansing	MI	2109
52242	Iowa City	IA	2092
97403	Eugene	OR	1960
55455	Minneapolis	MN	1955
06269	Storrs-Mansfield	CT	1952
47408	Bloomington	IN	1897
48104	Ann Arbor	MI	1889
06105	Hardtford	CT	1772

- Repeatable methodology
  - Approximate endpoints by… location/corporate entity.
  - Not ideal, but maybe good enough
  - Working with what's available
- Novel research community & direction
  - Additional efforts in the pipeline



CENTER FOR APPLIED
CYBERSECURITY RESEARCH

INDIANA UNIVERSITY
Pervasive Technology Institute

- Questions, comments?

- daripley@iu.edu

Flow Data at 10 GigE and Beyond

What can (or should) we do ?

Scott Pinkerton
pinkerton@anl.gov
Argonne National Laboratory
www.anl.gov

About me ….

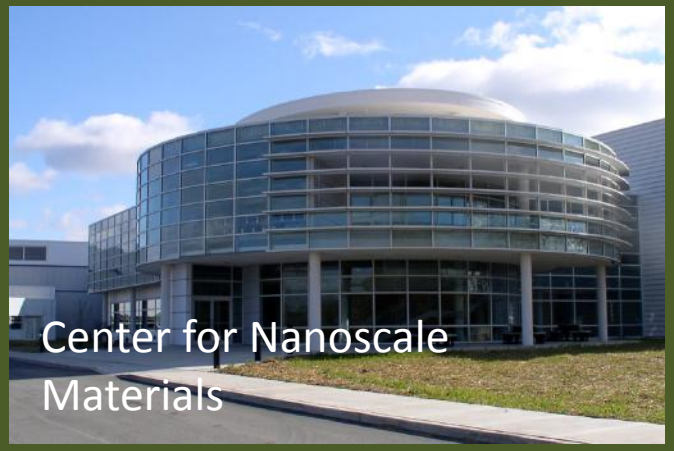
- Involved in network design, network operation & network security for the last 10 years
- Flow data practitioner
- Campus perspective
- Our flow data uses typically include:
  - Real-time anomaly detection
  - Forensic analysis



User facilities











Using Flow Data in a campus environment

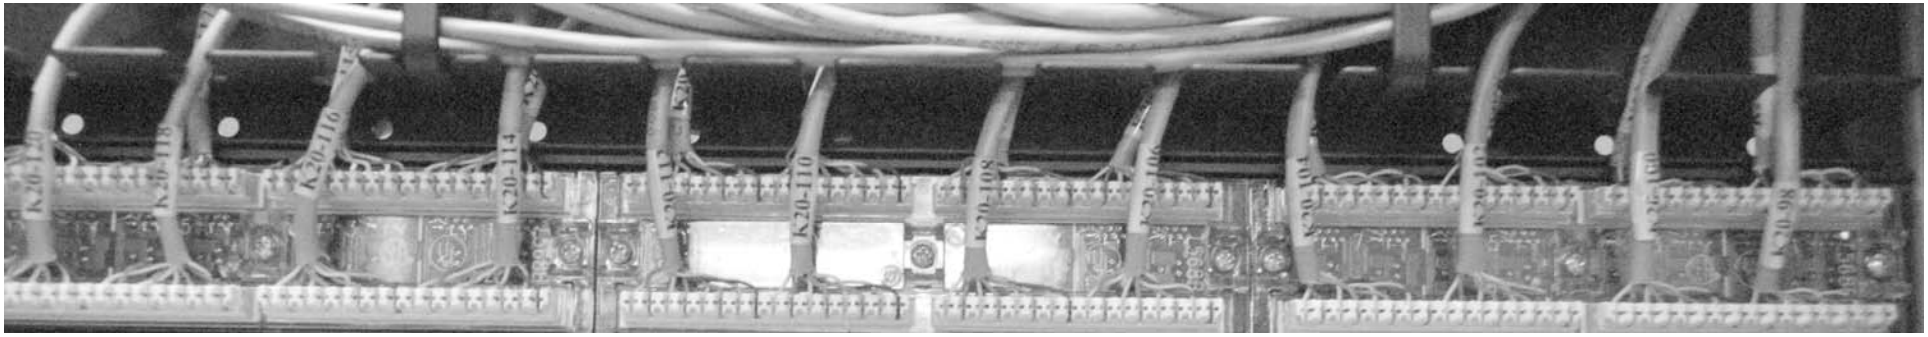
- In ~2000 started collecting Netflow data from all of the core campus network devices using the OSU Flowtools package
- By 2004, we were collecting Netflow data down in the distribution and access layers of the campus network
- Today, still consider flow data to be a critical part of our anomaly detection systems. Goals are to:
  - Protect the Laboratory computers from the Internet
  - Protect the Internet from the Laboratory computers
  - Have visibility into “lateral movement” of compromised hosts
- Campus environments can be large ….



Texas A&M Campus Network

- Wired Network
  - 10 Gbps backbone
  - 50,000 computers
  - 90,000 wired ports
- Gateway to regional and national networks

- Wireless Network
  - 11 million square ft. of wireless access
  - 340+ buildings with wireless access across 5200 acres



U of M Twin Cities Campus Network

- 23 Cisco 6509s
- 4,323 Cisco 3750s
- 1,133 Switch Stacks
- 74,414 Switchports
- Redundant 10-Gigabit Backbone
- Topology: 18 layer-2 switched domains interconnected by a layer-3 MPLS-VPN backbone



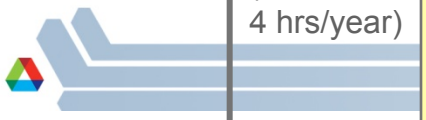
A "big science" Perspective - driving speeds & feeds

- Data networks continue to evolve in support of the scientific mission
- Key drivers include:
  - Large Hadron Collider (LHC), CERN
    - CERN to US Tier1 data rates: 10 Gbps by 2007, 30-40 Gbps by 2010/11
  - Leadership Computing Facilities (LCF), ANL and ORNL
  - Relativistic Heavy Ion Collider (RHIC), BNL
  - Large-scale Fusion (ITER), France
  - Climate Science
    - Significant data set growth is likely in the next 5 years, with corresponding increase in network bandwidth requirement for data movement (current data volume is ~200TB, 1.5PB/year expected rate by 2010)



Science Network Requirements Aggregation Summary

Science Drivers / Science Areas / Facilities	End2End Reliability	Connectivity	2006 End2End Band width	2010 End2End Band width	Traffic Characteristics	Network Services
Advanced Light Source	-	• DOE sites • US Universities • Industry	1 TB/day 300 Mbps	5 TB/day 1.5 Gbps	• Bulk data • Remote control	• Guaranteed bandwidth • PKI / Grid
Bioinformatics	-	• DOE sites • US Universities	625 Mbps 12.5 Gbps in two years	250 Gbps	• Bulk data • Remote control • Point-to-multipoint	• Guaranteed bandwidth • High-speed multicast
Chemistry / Combustion	-	• DOE sites • US Universities • Industry	-	10s of Gigabits per second	• Bulk data	• Guaranteed bandwidth • PKI / Grid
Climate Science	-	• DOE sites • US Universities • International	-	5 PB per year 5 Gbps	• Bulk data • Remote control	• Guaranteed bandwidth • PKI / Grid
High Energy Physics (LHC)	99.95+% (Less than 4 hrs/year)	• US Tier1 (DOE) • US Tier2 (Universities) • International (Europe, Canada)	10 Gbps	60 to 80 Gbps (30-40 Gbps per US Tier1)	• Bulk data • Remote control	• Guaranteed bandwidth • Traffic isolation • PKI / Grid



Science Network Requirements Aggregation Summary

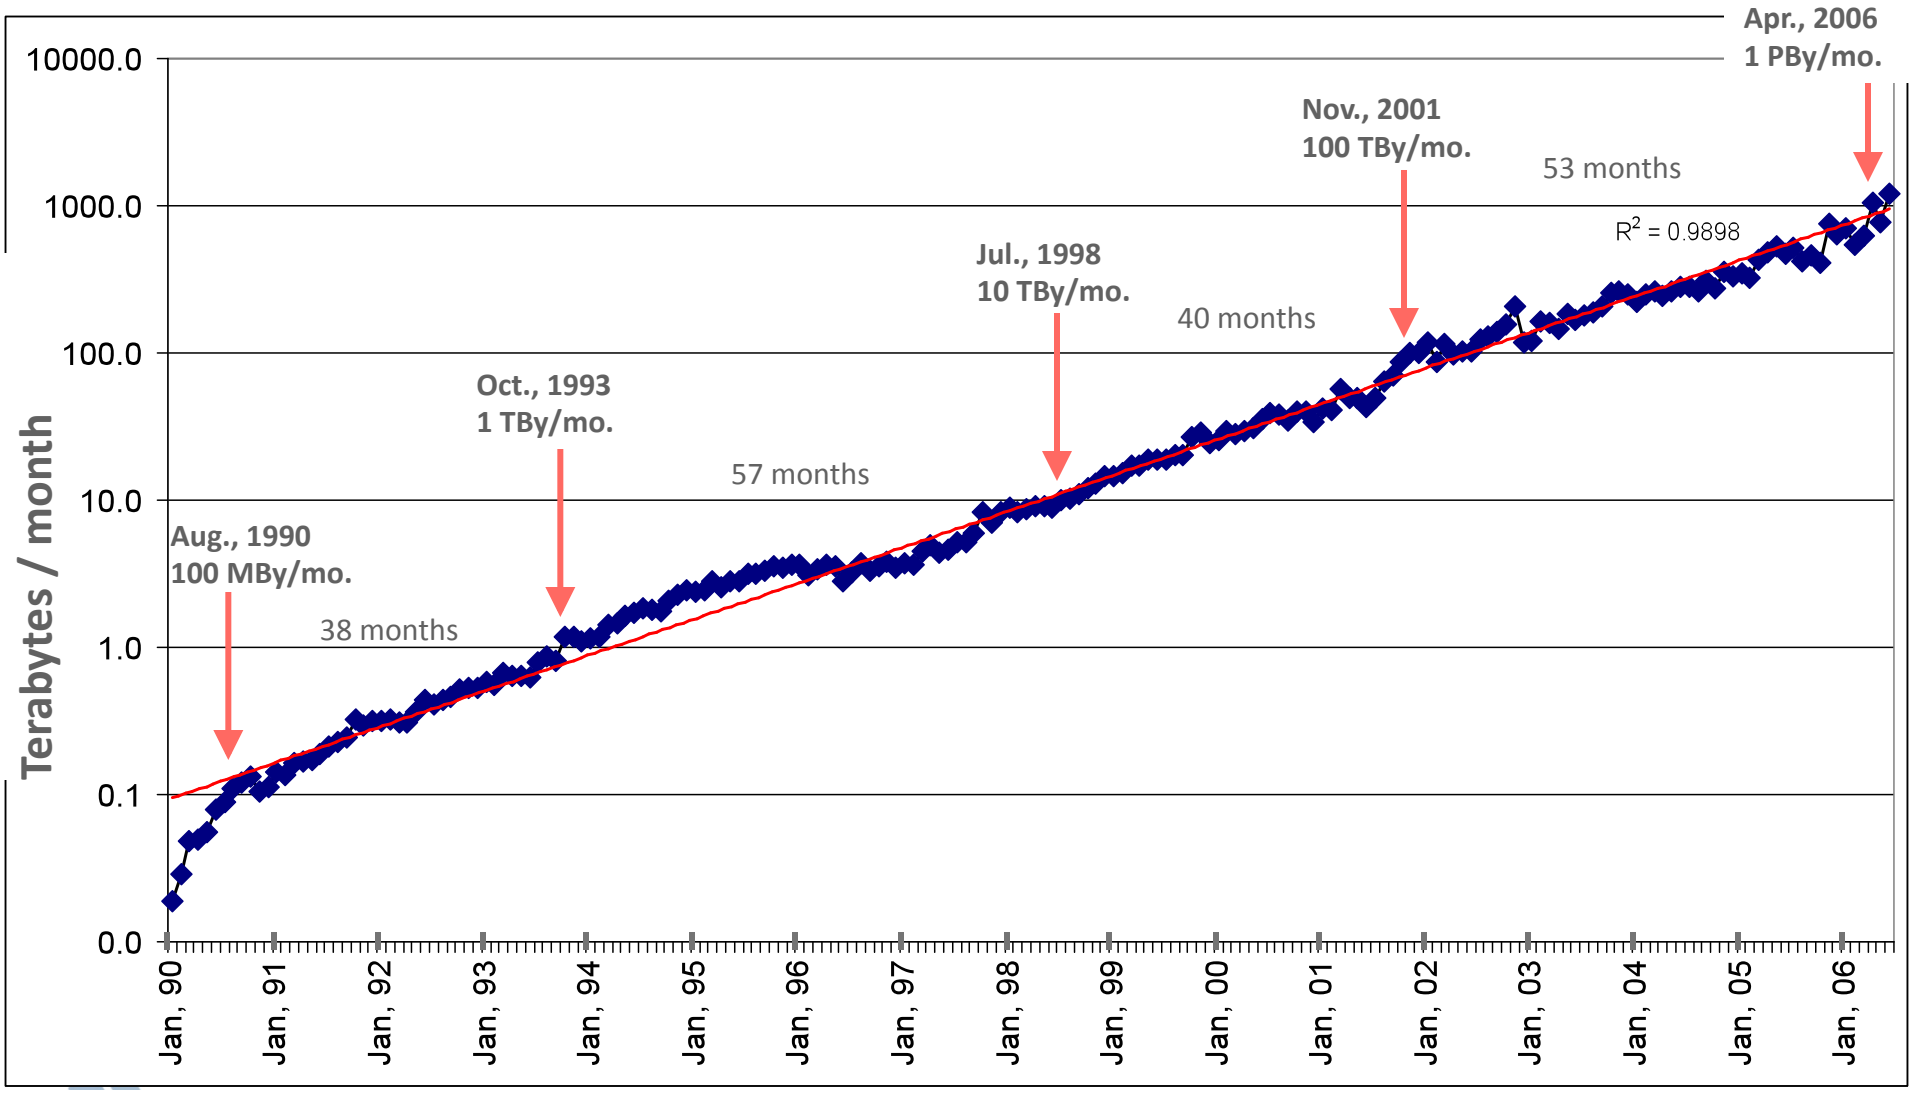
Science Drivers / Science Areas / Facilities	End2End Reliability	Connectivity	2006 End2End Band width	2010 End2End Band width	Traffic Characteristics	Network Services
Magnetic Fusion Energy	99.999% (Impossible without full redundancy)	• DOE sites • US Universities • Industry	200+ Mbps	1 Gbps	• Bulk data • Remote control	• Guaranteed bandwidth • Guaranteed QoS • Deadline scheduling
NERSC	-	• DOE sites • US Universities • Industry • International	10 Gbps	20 to 40 Gbps	• Bulk data • Remote control	• Guaranteed bandwidth • Guaranteed QoS • Deadline Scheduling • PKI / Grid
NLCF	-	• DOE sites • US Universities • Industry • International	Backbone Band width parity	Backbone band width parity	• Bulk data	
Nuclear Physics (RHIC)	-	• DOE sites • US Universities • International	12 Gbps	70 Gbps	• Bulk data	• Guaranteed bandwidth • PKI / Grid
Spallation Neutron Source	High (24x7 operation)	• DOE sites	640 Mbps	2 Gbps	• Bulk data	

Table 1: ALCF requirements summary

Feature	Key Science Drivers		Anticipated Network Requirements	
Time Frame	**Science Instruments and Facilities**	**Process of Science**	**Local Area Network Bandwidth and Services**	**Wide Area Network Bandwidth and Services**
Near-term (0-2 years)	• ALCF production resources (intrepid)	• Large file transfers. Other labs and computing centers are common targets, but it can be any institution based on INCITE users needs. • Some real-time video, computational steering, real-time control apps possible.	• Node to node is handled by proprietary vendor interconnect. 425 MB/s per link. • Node to storage is approx. 1,000 ports of 10 gigabit. • Other local needs are primarily admin-related and are trivial.	• 10s of TB/day • 10-30 Gbps
2-5 years	• Next major machine upgrade	• Large file transfers. Other labs and computing centers are common targets, but it can be any institution based on INCITE users needs. • Real-time video, computational steering, real-time control apps more common, but still relatively small in comparison to file transfers.	• Node to node is handled by proprietary vendor interconnect. 1-5 GB/s per link. • Node to storage is likely InfiniBand-based and on the order of 3K-5K ports • Other local needs are primarily admin-related and are trivial.	• 100s of TB/day • 100-300 Gbps
5+ years	• Push towards exascale computing	• Massive data sets are common. File transfers still dominate, but WAN file systems, distributed databases use grows. • Machines are sufficiently powerful that computational steering, real time simulations are used regularly • Use of collaboration tools continues to grow	• Node to node is probably still handled by proprietary vendor interconnect, but could be standards based, such as InfiniBand. • Node to storage is likely InfiniBand or other standards-based interconnect. • Other local needs are primarily admin-related and are trivial.	• Petabytes per day • Terabit networks



ESnet Traffic has Increased by 10X Every 47 Months, on Average, Since 1990



Log Plot of ESnet Monthly Accepted Traffic, January, 1990 – June, 2006

Key Take Aways

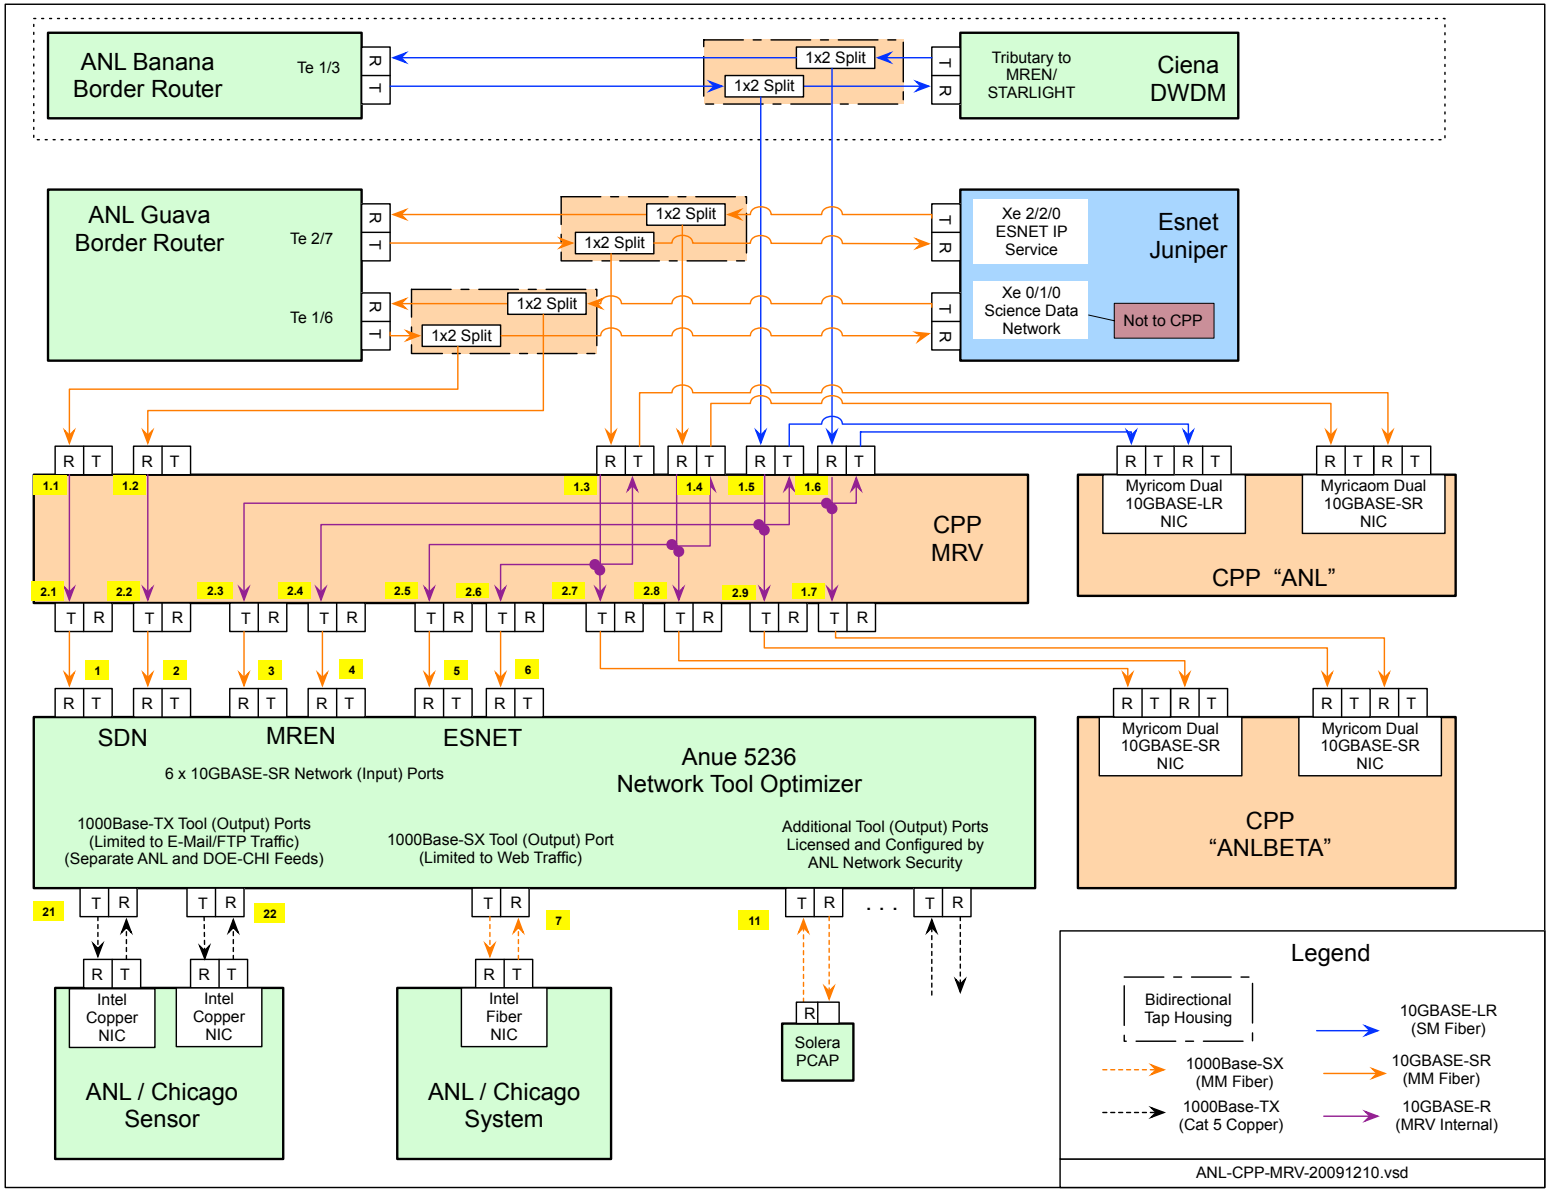
- Building networks for the future – takes a lot of planning
- Or, maybe more importantly it takes a lot of predicting (future requirements)
- Without the planning (and the predicting) how can the vendors gear up to provide the necessary capabilities ?
- Are we doing a good job communicating future requirements for flow data ?

Future of non-sampled Flow data seems bleak (IMHO)

- Speeds and feeds increasing to keep pace with scientific demand
- Many/most vendors are struggling to provide non-sampled flow data directly from the switches or routers just @ 10 Gbps (much less at 40 or 100 Gbps)
- Can optical taps really scale up to provide the needed number of monitor points ?
  - For me, I think the answer is no



Leveraging taps to create monitor points



What Can We Do - Process Perspective ?

- Identify our needs/requirements
- Write it down
- Communicate it to the vendors



What are our needs/requirements/drivers ?

- Strong support for “existence” analysis
- Scalable
  - From a campus perspective (monitoring at the border & internally)
  - From a “big science” perspective (speeds/feeds & large file txfer)
  - Equals – built in to the switches & routers (IMHO)
- Non-sampled data
  - Sampled data has its place (traffic engineering & other)
  - Painful to perform forensic analysis with sampled data



What if ideas ?

- Leveraging “cores” internal to the switches & routers for custom applications
  - Bloom filter ?
  - What info/data types would we want available to an internal app ?
- Adapting to the Very Very Large data txfers
  - Do we need a new scale for active timeout ? Was 30 seconds, now 30 minutes or 3 hours ?
- Notifications at start of a new flow – first packet ?



As a community – what can we do ?

- Should we try and develop future requirements ?
- Do we have enough energy/motivation to do it ?
- Can we agree on requirements ?
- Can we influence the networking equipment purchase decisions ?
- Your thoughts ??





Stager – A Generic Tool for Presenting Network Statistics

FloCon 2010
January 11-14, 2010

Arne Øslebø
arne.oslebo@uninett.no

What is Stager?

- A web based tool for presenting and aggregating most types of network statistics
- Store high level reports in database
- Detailed reports from other sources
- Access control
- Stable version
  - NetFlow
  - Qflow - IPFIX probe with extra QoS attributes
- Development version
  - Qstream, Mping, Asmping, Rude/Crude, SNMP
- http://software.uninett.no/stager/
  - GPL license

UNINETT

Overview report

Setup Back | Protocol | overview | Octets percent | 10 | Show | Add

All interfaces | vpn-gw.uio.no | In | none | 1

Protocol

Thursday 15. October 2009, 21:00

All observation points (in)

Pie chart | Other | Plot graph

Select	Observation Point	Octets - percent							
		TCP	UDP	GRE	ESP	IPv6	ICMP	PIM	IPv6-ICMP
	elverum-gw.elverum-gw2-1	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%
	oslo-trd2	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%
	trd-tromso2	0.00%	0.00%	0.00%	0.00%	0.00%	75.68%	24.32%	0.00%
	elverum-gw.elverum-gw2-2	0.00%	36.36%	0.00%	0.00%	0.00%	63.64%	0.00%	0.00%
	hoytek-niva	0.64%	65.75%	0.00%	0.00%	0.00%	33.61%	0.00%	0.00%
	trd-niva	3.86%	66.29%	0.00%	0.00%	0.00%	29.84%	0.00%	0.00%
	oslo-oslomsh	19.42%	53.90%	0.00%	0.00%	0.00%	26.68%	0.00%	0.00%
	teknobyen2-teknobyen	5.89%	3.82%	0.00%	0.00%	0.00%	23.70%	66.59%	0.00%
	nhh2-nhh	7.41%	69.95%	0.00%	0.00%	0.00%	22.64%	0.00%	0.00%
	nhh-gw2.nhh-gw	7.41%	69.95%	0.00%	0.00%	0.00%	22.64%	0.00%	0.00%

Stager 4.0, 2004-2009 © UNINETT AS

Processing the report took 486.82ms

UNINETT

Global report

Setup Back | Source port | global | Standard | 10 | Show | Add | Bookmarks..

All interfaces | uninett.uninett-wlan_ans | In | none | 1

Source port

Tuesday 12. January 2010, 3:00
All observation points (in)

Select	Port number	Port name	Observation point	Octets		Packets		Flows	
				bit/s	Percent	Pkts/s	Percent	flows/s	Percent
	80	http	uninett-nordunet	159M	20.52%	$14.9 \cdot 10^3$	9.87%	25.7	9.63%
	80	http	uninett-nordunet4	127M	17.12%	$12.1 \cdot 10^3$	9.93%	22.3	11.04%
	80	http	trd-oslo	117M	19.04%	$11.2 \cdot 10^3$	10.45%	23.7	11.90%
	80	http	oslo-dnmi	92.3M	92.03%	$8.92 \cdot 10^3$	66.50%	13.1	88.40%
	33661	33661	uninett.uninett-nett	85.5M	93.45%	$7.38 \cdot 10^3$	81.17%	$5.83 \cdot 10^{-3}$	0.03%
	80	http	oslo-gw.stolav-gw1	85.0M	14.92%	$8.08 \cdot 10^3$	8.93%	19.2	12.24%
	49152	49152	nix.nix	78.6M	44.92%	$7.31 \cdot 10^3$	24.25%	0.188	0.39%
	80	http	oslo-trd	73.4M	9.06%	$6.61 \cdot 10^3$	5.97%	10.6	5.24%
	80	http	bergen-oslo	68.6M	19.68%	$7.05 \cdot 10^3$	12.18%	17.8	21.07%
	48959	48959	trofast-gw1	65.3M	54.92%	$5.79 \cdot 10^3$	54.89%	$4.17 \cdot 10^{-3}$	0.05%

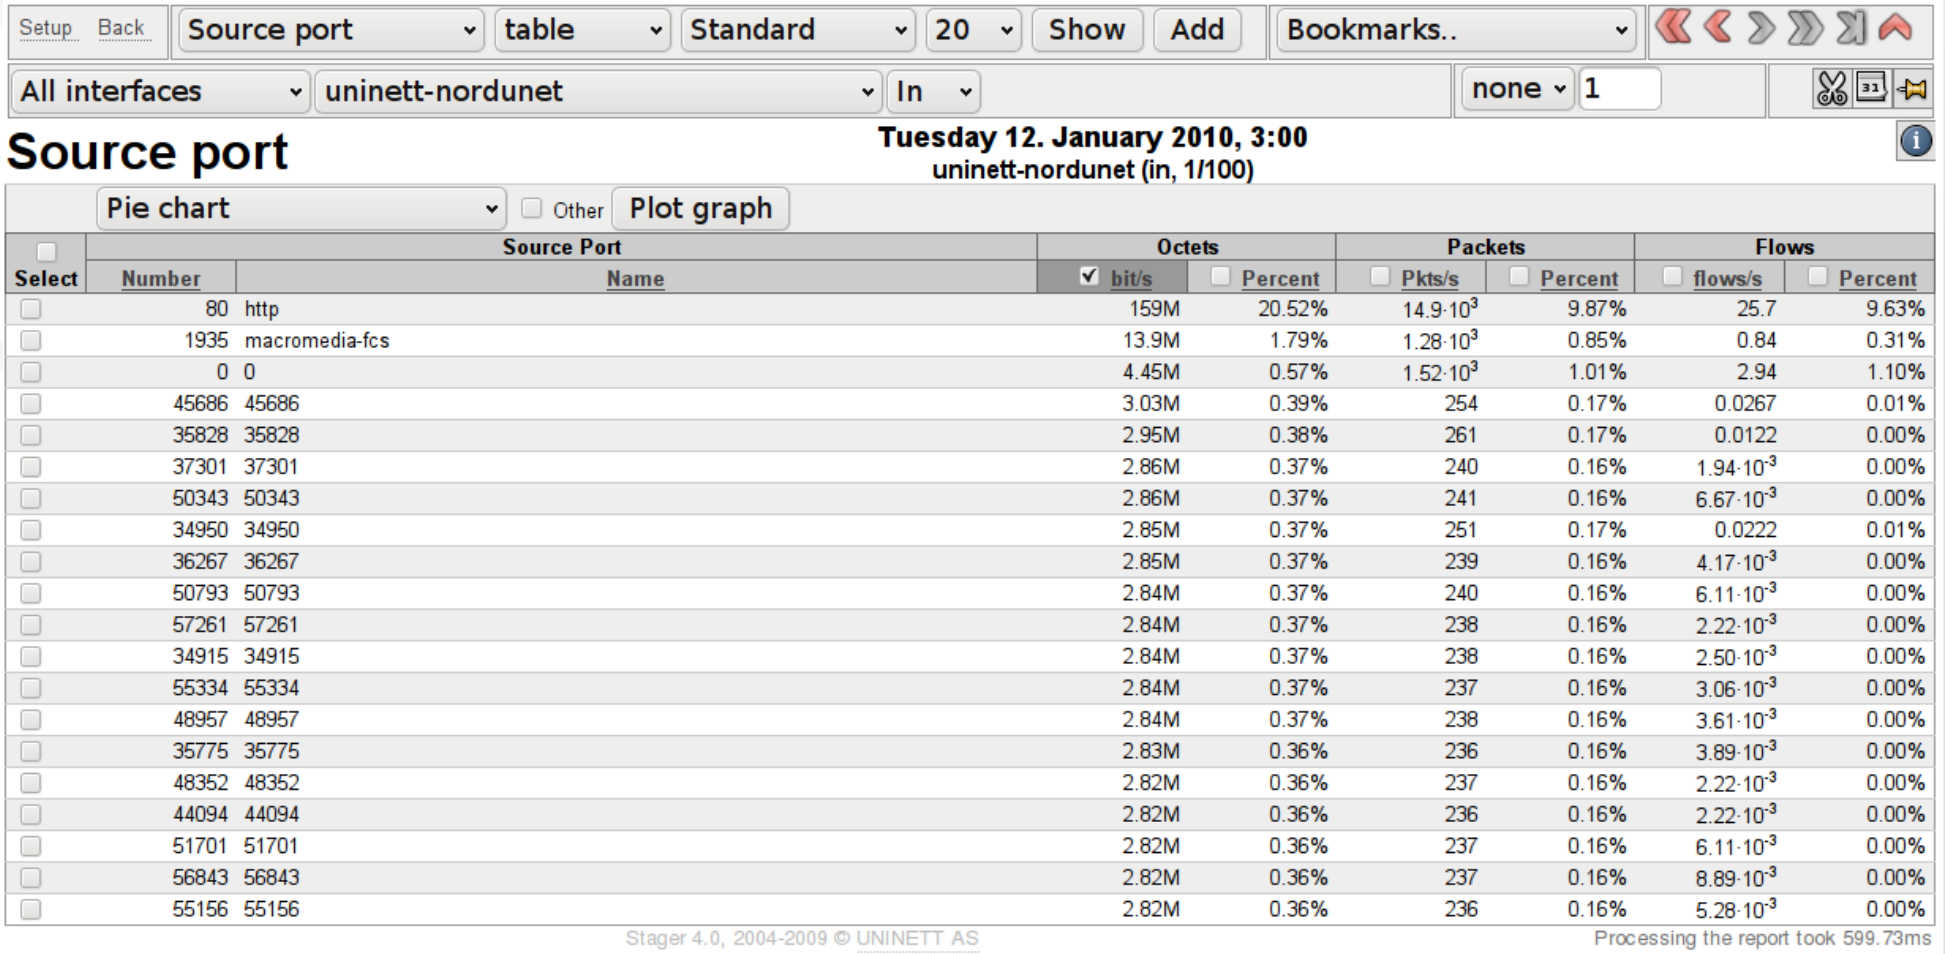
Stager 4.0, 2004-2009 © UNINETT AS

Processing the report took 681.74ms

4

UNINETT

Detailed report



Select	Source Port		Octets		Packets		Flows	
	Number	Name	bit/s	Percent	Pkts/s	Percent	flows/s	Percent
	80	http	159M	20.52%	$14.9 \cdot 10^{3}$	9.87%	25.7	9.63%
	1935	macromedia-fcs	13.9M	1.79%	$1.28 \cdot 10^{3}$	0.85%	0.84	0.31%
	0	0	4.45M	0.57%	$1.52 \cdot 10^{3}$	1.01%	2.94	1.10%
	45686	45686	3.03M	0.39%	254	0.17%	0.0267	0.01%
	35828	35828	2.95M	0.38%	261	0.17%	0.0122	0.00%
	37301	37301	2.86M	0.37%	240	0.16%	$1.94 \cdot 10^{-3}$	0.00%
	50343	50343	2.86M	0.37%	241	0.16%	$6.67 \cdot 10^{-3}$	0.00%
	34950	34950	2.85M	0.37%	251	0.17%	0.0222	0.01%
	36267	36267	2.85M	0.37%	239	0.16%	$4.17 \cdot 10^{-3}$	0.00%
	50793	50793	2.84M	0.37%	240	0.16%	$6.11 \cdot 10^{-3}$	0.00%
	57261	57261	2.84M	0.37%	238	0.16%	$2.22 \cdot 10^{-3}$	0.00%
	34915	34915	2.84M	0.37%	238	0.16%	$2.50 \cdot 10^{-3}$	0.00%
	55334	55334	2.84M	0.37%	237	0.16%	$3.06 \cdot 10^{-3}$	0.00%
	48957	48957	2.84M	0.37%	238	0.16%	$3.61 \cdot 10^{-3}$	0.00%
	35775	35775	2.83M	0.36%	236	0.16%	$3.89 \cdot 10^{-3}$	0.00%
	48352	48352	2.82M	0.36%	237	0.16%	$2.22 \cdot 10^{-3}$	0.00%
	44094	44094	2.82M	0.36%	236	0.16%	$2.22 \cdot 10^{-3}$	0.00%
	51701	51701	2.82M	0.36%	237	0.16%	$6.11 \cdot 10^{-3}$	0.00%
	56843	56843	2.82M	0.36%	237	0.16%	$8.89 \cdot 10^{-3}$	0.00%
	55156	55156	2.82M	0.36%	236	0.16%	$5.28 \cdot 10^{-3}$	0.00%

Stager 4.0, 2004-2009 © UNINETT AS

Processing the report took 599.73ms

UNINETT

Graph

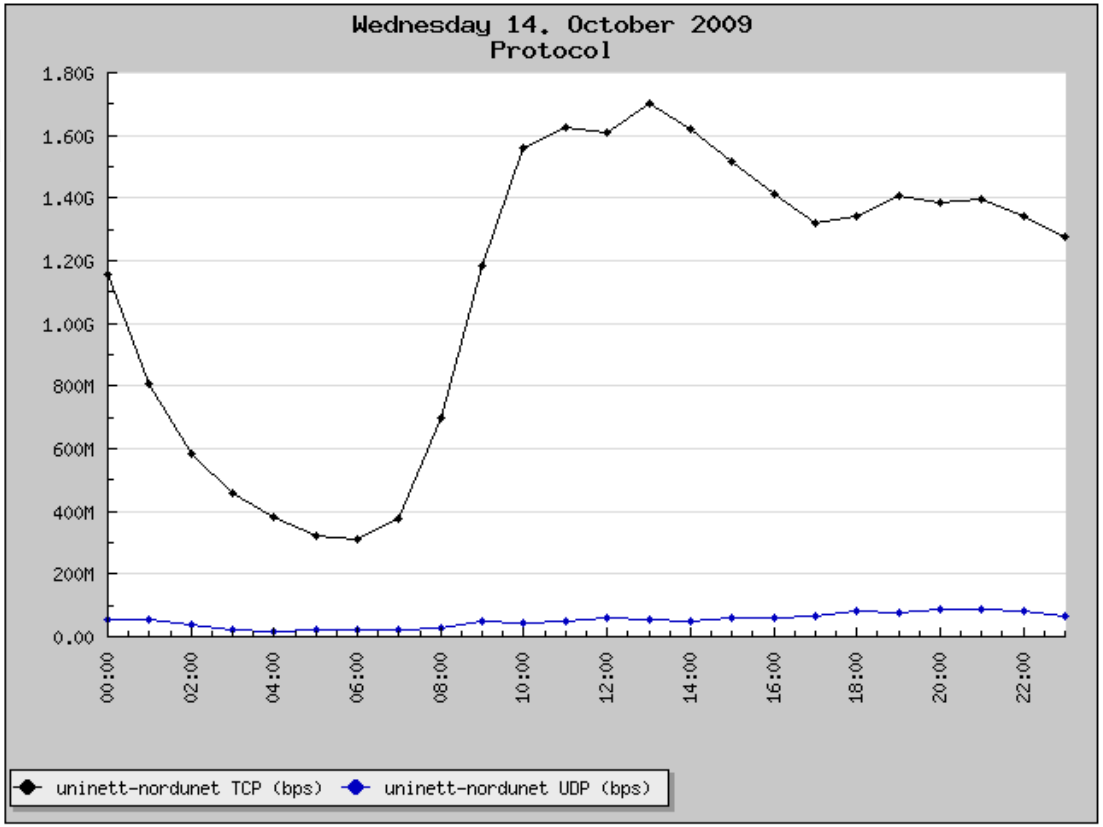
Setup Back | Protocol | table | Octet details | 20 | Show | Add | hour...

oslo-gw | uninett-nordunet | In | none | 1

Protocol

Wednesday 14. October 2009

[Return to table] | Line plot | Linear | Other | Replot

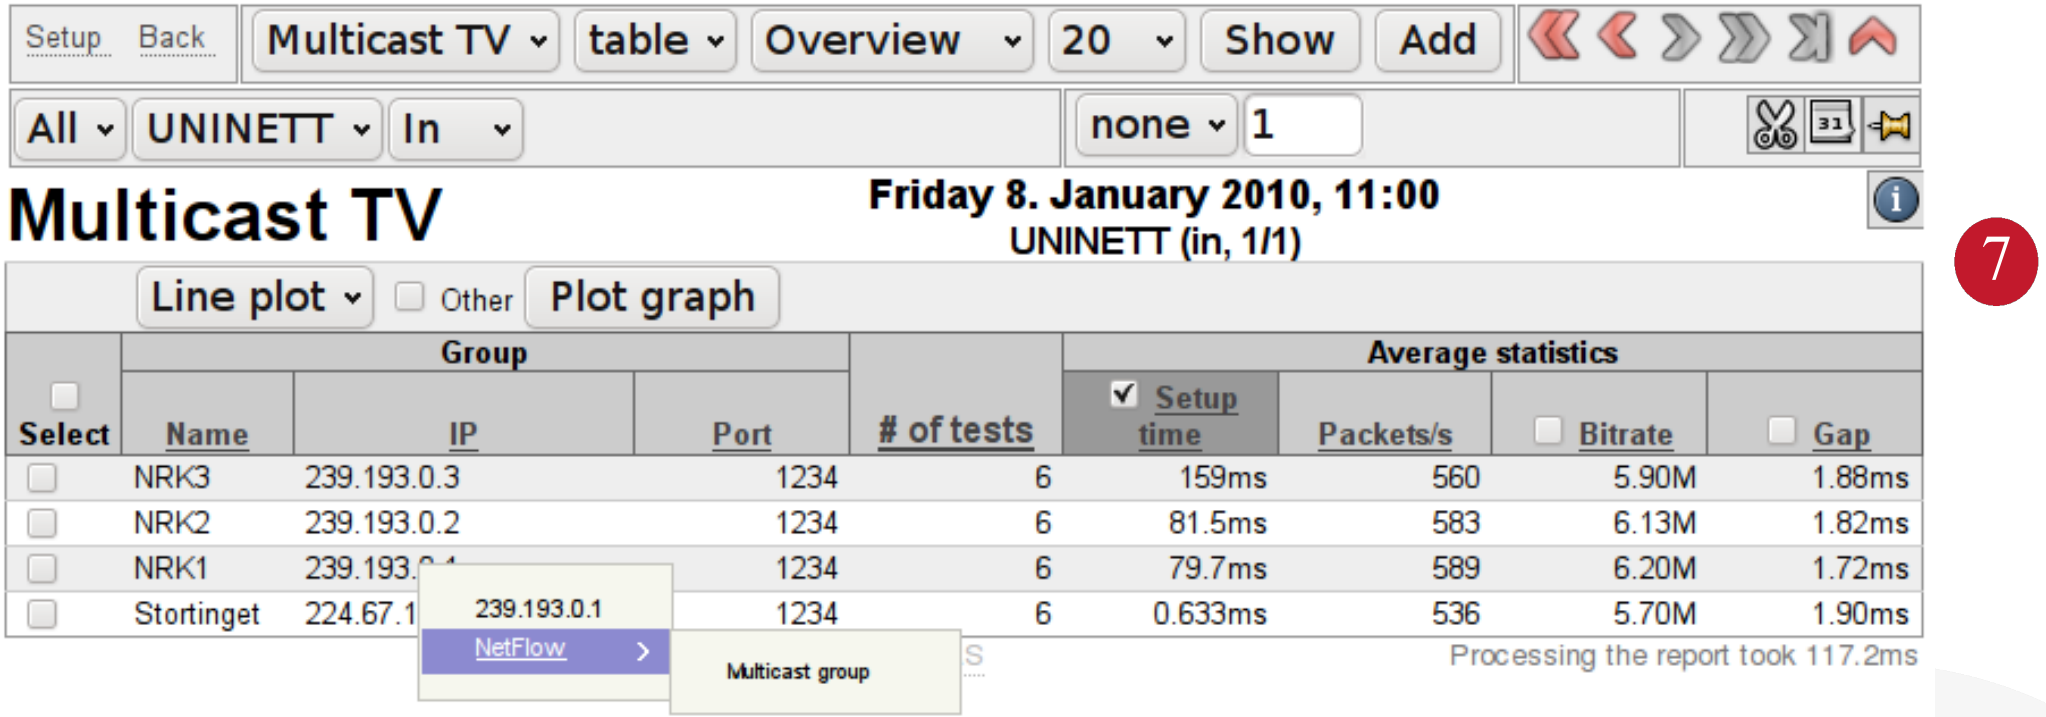


Stager 4.0, 2004-2009 © UNINETT AS

Processing the report took 168.55ms

UNINETT

Context menus



Multiple reports

Setup Back | Protocol | table | Standard | 5 | Show | Add | Bookmarks.. | hour...

All interfaces | trd-teknobyen | In | none | 1

Protocol

Wednesday 14. October 2009
trd-teknobyen (in, 1/100)

Line plot | Other | Plot graph

Select	Protocol Number	Protocol Name	Octets bit/s	Octets Percent	Packets Pkts/s	Packets Percent	Flows flows/s	Flows Percent
	6	TCP	75.9M	98.70%	$7.62 \cdot 10^3$	93.33%	4.95	64.15%
	17	UDP	915k	1.19%	442	5.41%	2.04	26.42%
	1	ICMP	52.0k	0.07%	71.4	0.87%	0.691	8.96%
	47	GRE	28.2k	0.04%	26.3	0.32%	0.0178	0.23%
	103	PIM	2.46k	0.00%	5.22	0.06%	0.0178	0.23%

Protocol

Tuesday 13. October 2009
trd-teknobyen (in, 1/100)

Line plot | Other | Plot graph

Select	Protocol Number	Protocol Name	Octets bit/s	Octets Percent	Packets Pkts/s	Packets Percent	Flows flows/s	Flows Percent
	6	TCP	152M	99.10%	$14.3 \cdot 10^3$	95.46%	5.94	66.18%
	17	UDP	1.17M	0.76%	545	3.63%	2.23	24.79%
	50	ESP	92.8k	0.06%	14.1	0.09%	$3.98 \cdot 10^{-3}$	0.04%
	47	GRE	62.4k	0.04%	37.8	0.25%	0.0183	0.20%
	1	ICMP	57.3k	0.04%	79.4	0.53%	0.77	8.57%

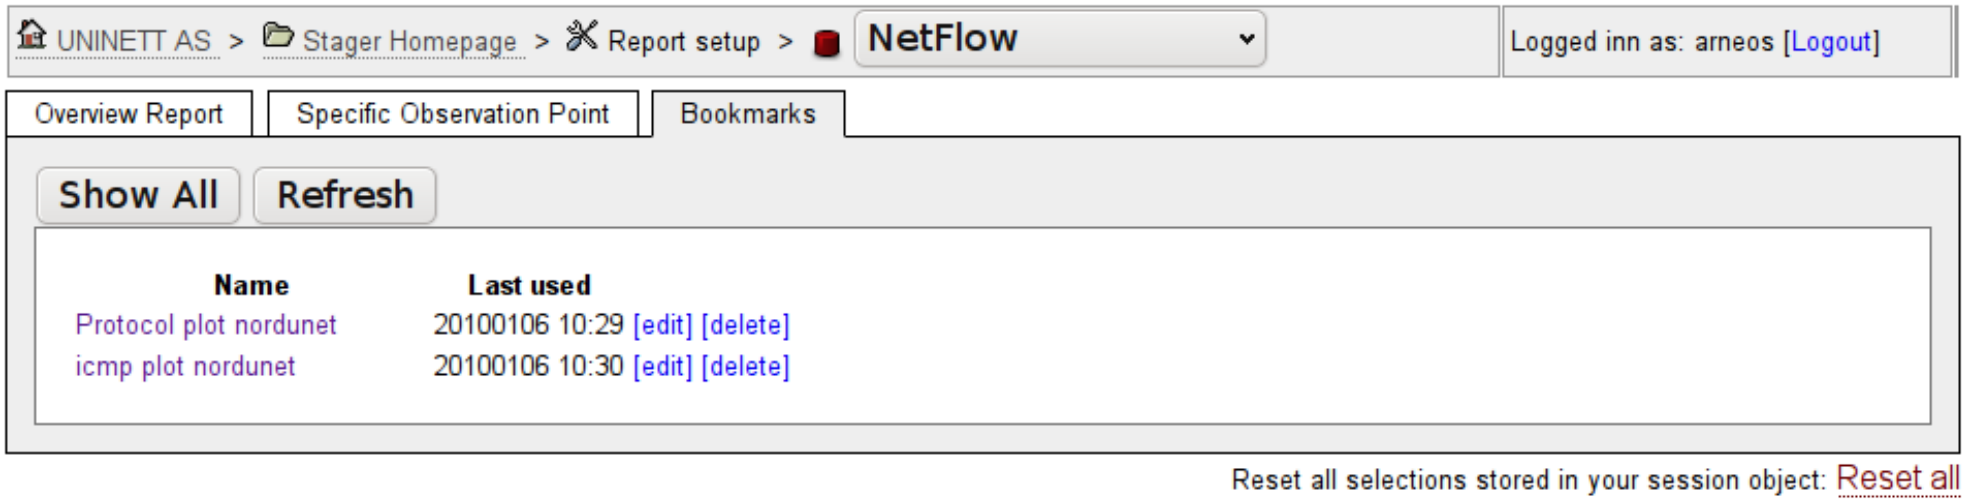
Stager 4.0, 2004-2009 © UNINETT AS

Processing the report took 265.77ms

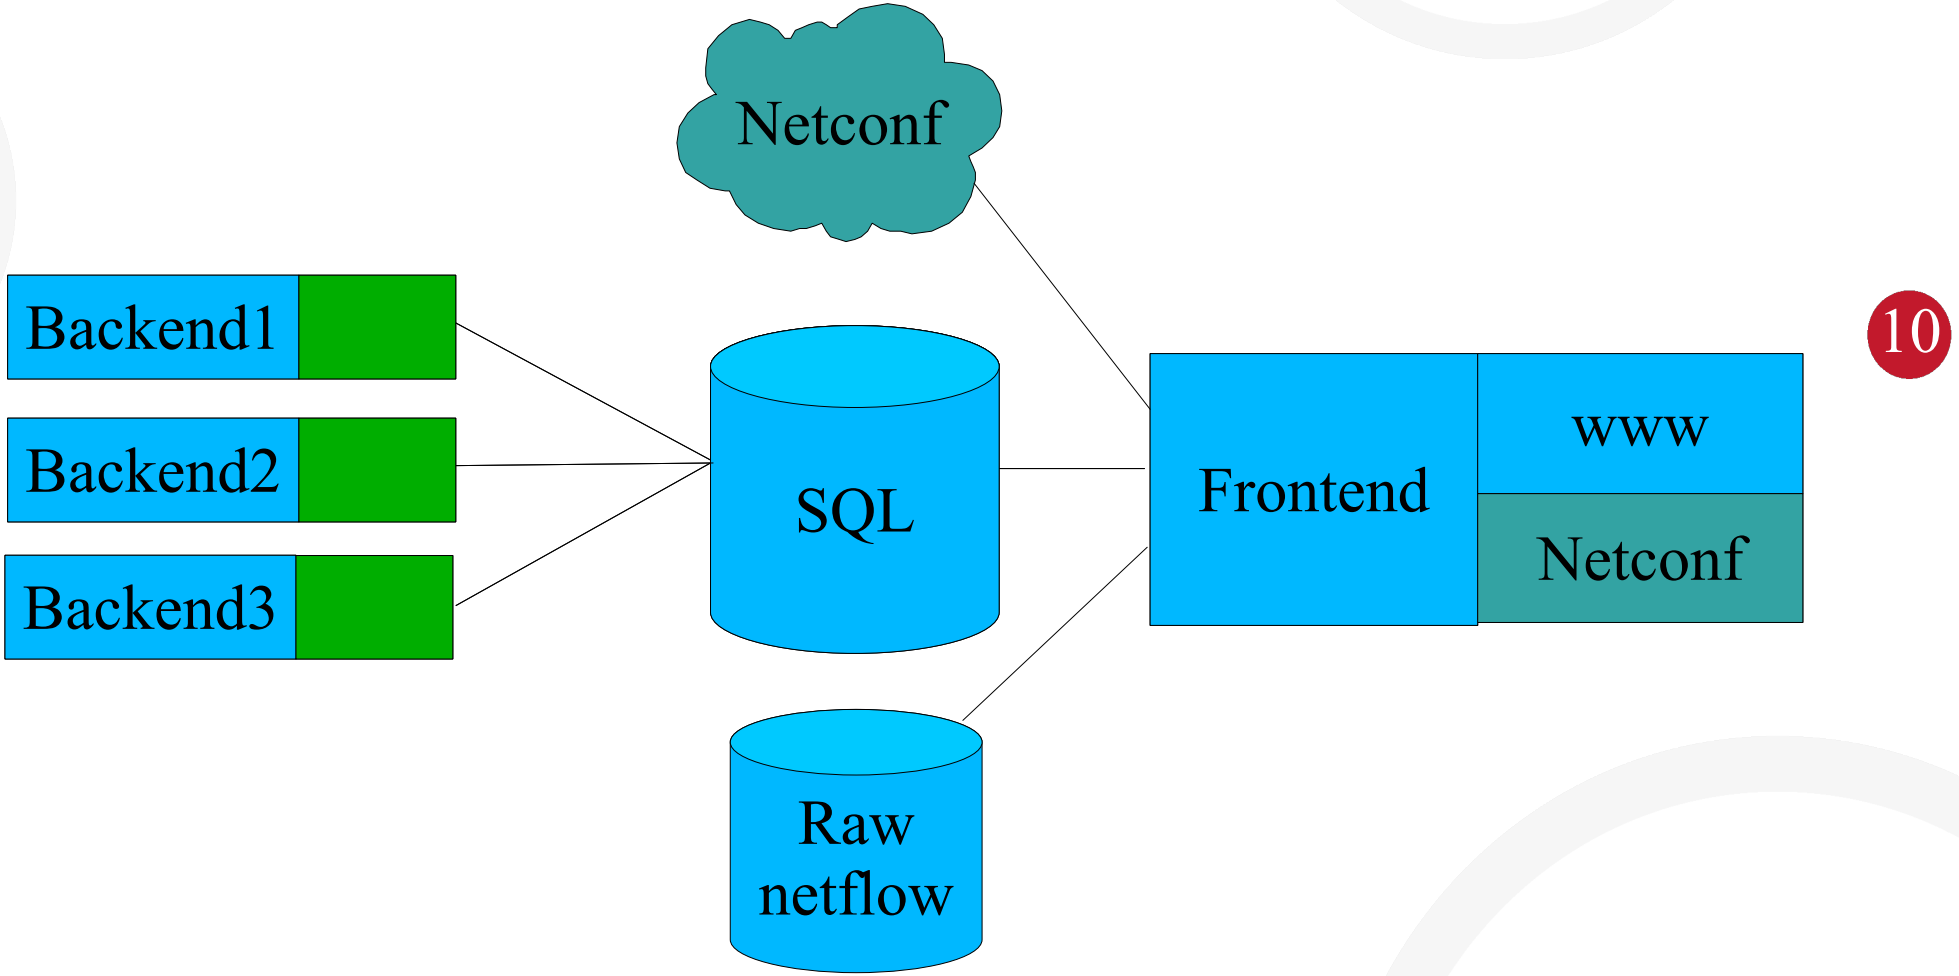
UNINETT

Bookmarks

- Setting up multiple reports can involve a lot of clicking
- Possible to create bookmarks with relative time
  - Always show last hour, day, week etc.



Stager architecture



UNINETT

Template based

- All data for a report is stored in a database or collected from other sources
- XML based template specifies how the frontend should present the data
  - bps, pps, %, temperature, IP address etc.
- Built in functions
  - divideandmultiply, persec, bitpersec, concat etc.
- Transformations
  - Table, overview, global (and matrix).
- Multiple views for each transformation.
  - Specifies which data types should be visible
- New reports can be added without code changes to the frontend

UNINETT

Backend robustness

- Graceful handling of database downtime
  - Store reports to files
  - Insert into database when possible
- Avoid multiple instances of the same backend
  - Wait for previous instances that are still processing the raw data
  - Detect dead locks and memory starvations
- Supports anycast for NetFlow
  - Raw NetFlow data automatically processed by active collector

UNINETT

Anycast support

- All collectors are configured to collect data from all routers
- Script automatically detects which routers a collector actually collects data from
  - Development version: raw NetFlow data automatically copied to one server
- Collectors run Linux
- Loopback interface configured with anycast IP address
- Runs Quagga routing software
  - BGP peering
- Sanity check scripts
  - Takes down BGP peering if problems

UNINETT

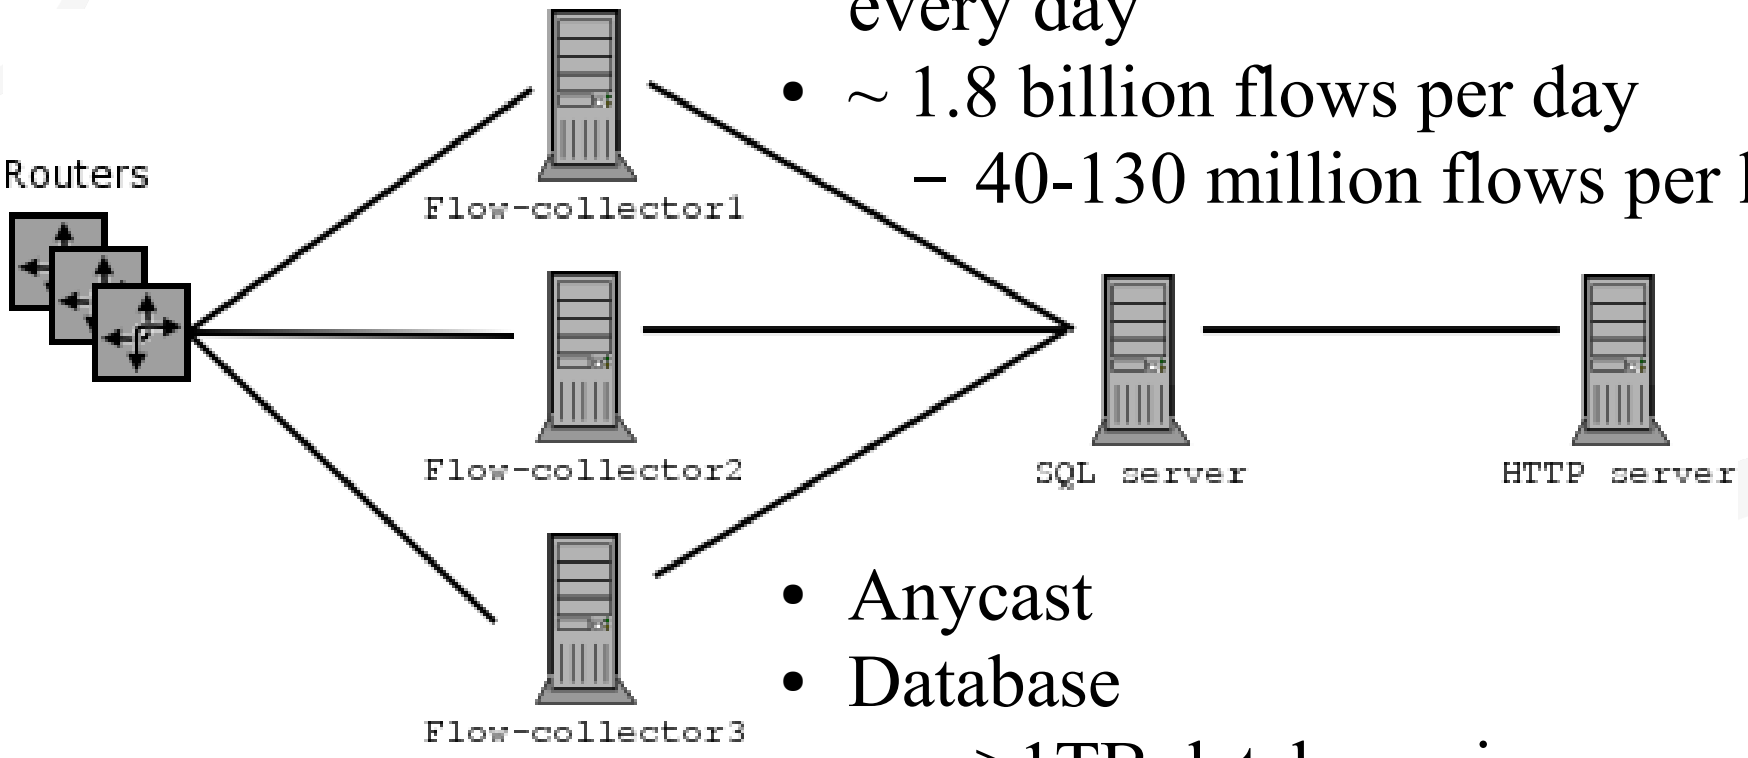
Netflow in Stager

- Runs nfdump
- Reports generated every hour
- Aggregated to day, week, month and year
- Available reports
  - Summary
  - Interface
  - Src/Dst AS/IP/Port (Top X)
  - Protocol
  - Mulicast
    - Src AS
    - Src IP
    - Groups (Dst IP)
- Possible to view raw NetFlow data for more details
- Reports follow links, not interfaces

UNINETT

Our NetFlow setup

- 34 routers
- 439 interfaces
- 1/100 sampling rate on most routers
- ~25Gb of compressed raw Netflow data every day
- ~ 1.8 billion flows per day
  - 40-130 million flows per hour



- Anycast
- Database
  - >1TB database size
  - >800 millions entries in a single table

UNINETT

NetFlow performance

Backend

Data from January 5 between 12:00-13:00

	Col1	Col2	Col3	Total
Netflow size	470MB	62MB	654MB	1186MB
Sequentially	3min 33s	1min 36s	6min 5s	11min 14s
Simultaneously	3min 36s	1min 38s	6min 17s	11min 31s

Frontend

Report	Time
IP Protocol overview	1540ms
IP Protocol overview (previous timeperiod)	487ms
IP Protocol detailed report	223ms
Top Src Port overview	1068ms
Top Src Port overview (previous timeperiod)	507ms
Top Src Port details	442ms
Top Src IP address from raw NetFlow	3220ms

UNINETT

Upcoming features

- Improved anycast support
- Improved visualization
  - Today only simple graphs and pie charts are supported
- Better integration between different data sources
  - Shows reports from different sources on the same screen
- Anomaly detection
  - Today we only have simple threshold alerts
- Improved NetFlow reports
  - Top X + specified ports/IP addresses/AS numbers

UNINETT

Questions?

Arne Oslebo
arne.oslebo@uninett.no

http://software.uninett.no/stager

UNINETT

REDJACK

First experiences with Cuckoo bags

John McHugh - RedJack, LLC and The University of North Carolina

Jeff Janies - Redjack LLC

Teryl Taylor - Dalhousie University

FloCon 2010

New Orleans

January 2010

REDJACK

What is a cuckoo bag?

- SiLK sets and bags have single index field
  - chosen from subset of SiLK record fields
  - bags have single volume data field: flows, pkts, bytes
  - pointer tree implementation limits key to 32 bits
- Cuckoo bags have multiple index fields
  - all meaningful SiLK record fields plus
    - derived fields such as country code, and
    - key fields can be masked or reduced in precision
  - multiple data fields, volume, plus “span”, plus TBD
  - efficient, hash based indexing

REDJACK

Why Cuckoo?

- Cuckoo bags use multiple hash functions, so there are several places to put an object.
- If these are all full, their occupants alternates are checked and if there is a space, the occupant is kicked out to the alternate space.
  - This is likened to the European Cuckoo bird which lays its eggs in the nests of other birds, dumping one or more existing eggs.
  - The search for an entry to move is done recursively until a space is found, or we give up.

REDJACK

Give Up?

- At every level, the search expands.
  - Takes longer to find a hole
  - above about 90% table occupancy it is better to reallocate and rehash.
  - Since the new table is less than 50% full, no searching is required on the rehash
  - If you know how big the table needs to be, you can avoid searching altogether.
    - First search typically occurs at 65%+ occupancy

REDJACK

Advantages and disadvantages

- Works with IPv6 keys and multiple keys
- A set is a bag with no data
  - Can treat a bag as a set for set operations
  - Disk representation is similar to rwbags
- Key is explicitly part of memory representation
  - can require more space; depends on locality
- Constant time lookup for filter applications
  - does not grow with size as with R/B trees
  - can use multiple cores to speed hashing

What do we have?

REDJACK

- cubag program
  - like rwbag / rwset but more general

    --bag-file=<path>:<key>..<key>:<data>..<data>

    --set-file=:<path>:<key>..<key>
  - Can be repeated for multiple bags / sets
  - key fields: {s,d,nh}IP, v{4,6}{s,d,nh}IP, protocol, {s,d}Port, {s,e}Time, duration, sensor, input, output, {s,d}cc, {,initial,session}flags, attributes, application, typeclass, ICMPtypecode, IPversion, bytes, pkts
  - data fields: flows, bytes, packets, duration, span, counts
  - Times to second only
  - Span is minimum sTime, maximum eTime for key
  - Count is derived data field during projection

REDJACK

What else?

- Command options for `rw{set, bag}` superset
- Key modifiers
  - masking IPs and flags (&, 255.255.0.0) or (&,SAFR)
  - reduction of times (*,3600) or (*,86400)
    - hourly, daily grouping by start or end time
    - will build plugin for rwcount style binning
  - example
    - hourly volumes between /16s and hosts in a /16
    - v4sip(&,255.255.0.0),v4dip(&,0.0.255.255),sTime(/*,3600)
    - TCP Initial state flags per IP
    - v4sip,initialflags(&,SAFR)

REDJACK

cubagcat

- Simple listing of cubag
  - Count entries, describe bag
  - With or without headers (cubags are self describing)
  - epoch and clock time formats (times, duration, span)
  - zero padding of IPs, integer IPs for IPv4
  - No network structure (have to limit to IPv4, single key)
  - No binning (moves to bag tool)
  - Per field statistics

REDJACK

Example: Mixed IPv4, IPv6 Bag

sourceIP	protocol	IPVer	Flows
::	58	6	194
64.86.88.116	41	4	20
128.237.230.30	17	4	1
128.237.238.167	1	4	10
128.237.238.167	41	4	20
128.237.243.180	17	4	8
128.237.247.204	17	4	11
128.237.248.255	17	4	2
128.237.254.83	17	4	10
2001:200::8002:203:47ff:fea5:3085	58	6	1
2001:5a0:300::5	58	6	1
2001:5a0:300:100::2	58	6	1
2001:5a0:300:200::2	58	6	1

REDJACK

cubagtool (under construction)

- Everything rw{set,bag}tool does, cubagtool does better (or right)
- Additional operations for projection, binning
  - user defined field names for “count” field(s)
- Mix of unary, binary, n-ary operations
  - some unary ops combine w. others in one pass
- Stream operations allow arbitrary size growth
  - If inputs and outputs maintain sort order, memory representation of output not needed
    - set union, intersection, bag addition, subtraction

REDJACK

cubagtool hacks

- Work with text from `cubagcat`
- We need set prefix projection now
  - script to drop trailing set key fields and merge/count
- We also need set intersection and difference
  - script runs through 2 set listings, similar keys
  - 3 outputs (common to both, in first and not second, in second and not first) Could add set union, as well
- Finally, need to join bags on common key
  - output has key, selected data fields

REDJACK

Coming soon!!

- plugin for `rwfilter` that will filter flow records in the manner of the current tuples using a cuckoo set (will automatically extract the cover set of a cuckoo bag)
- `cubagbuild` to construct cuckoo sets and bags from text records.
- plugin for `cubag` for time distributed binning volume fields in the manner of `rwcount`.
- plugin for `cubag` to do sums of squares of data

REDJACK

Case studies

- We present 3 examples
  - Web activity profiling
    - looking for repeated connection patterns: host pairs, temporal regularity, consistent volumes
  - Client Server activity
    - Feeds FloVis activity viewer
  - Dark Space analysis
    - Characterizing traffic in empty network segments or the space between hosts

REDJACK

Web Profiling

- Demonstrate a clear, consistent communication pattern for a given host over a time interval.
- Patterns provide evidence:
  - Of similar activity.
  - User/process preference for external hosts
- *Note, here we only discuss the detection of the initial pattern and avoid discussion of the verification process of a candidate web profile.*

Cubags: Represent Trends

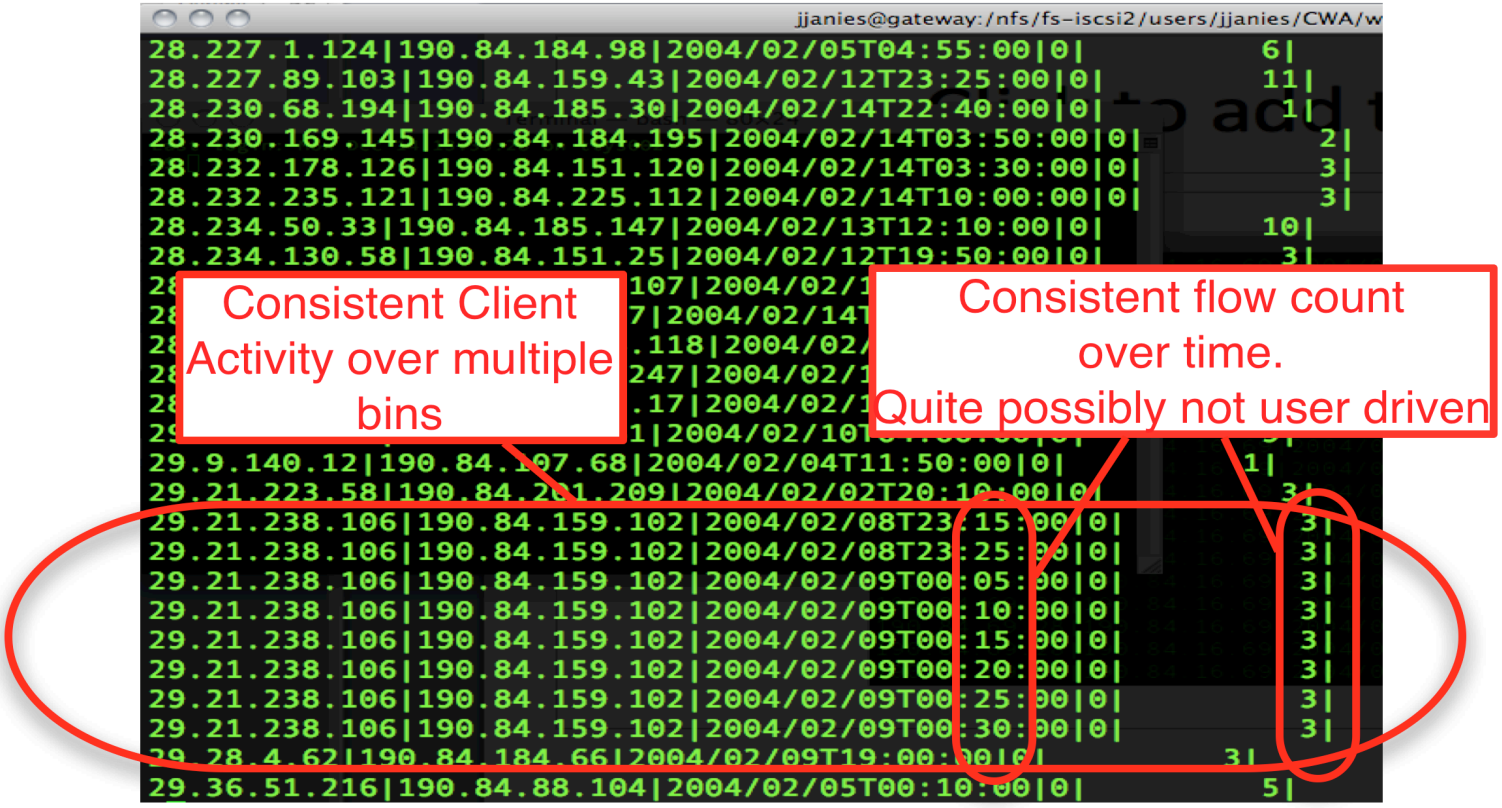
REDJACK

- Understanding common elements in client web activity.
  - Destination IP/Port
  - Intermittent/continuous
  - Size
- Trend of web client activity over time with 5 minute bins.

```
rwfilter --start=2004/02/01 --end= 2004/02/14 \
        --proto=6 --sport=1024- --dport=80,443 --pass=stdout | \
cubag --bag-file:clientActivity.cub:sip,dip,stime(/*,300):flows,bytes
```

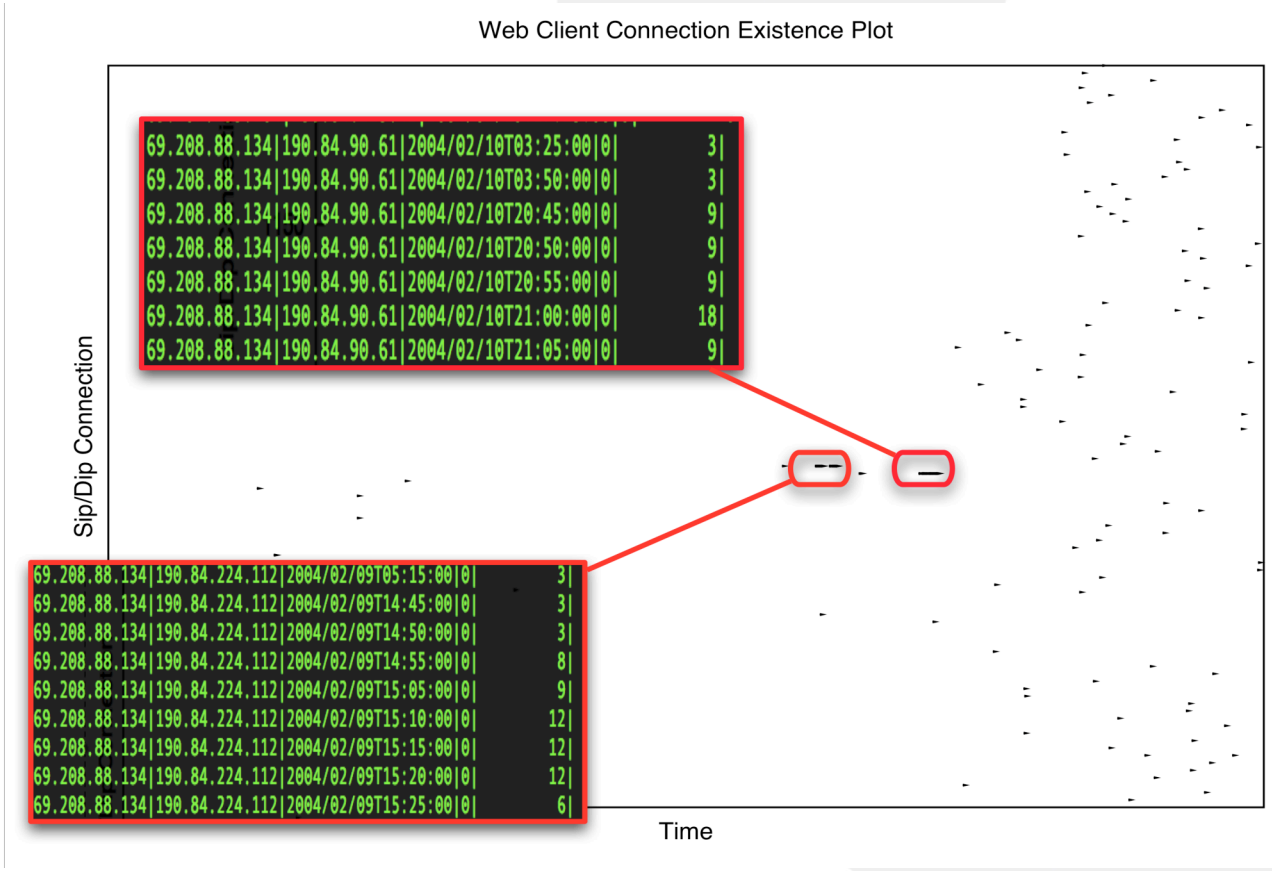
REDJACK

Cubag: Organized Raw Data with Meaning



REDJACK

Showing Consistent Patterns in Communication



REDJACK

Client / Server Characterization

- 5 categories: Idle, C, S, C/S-diff, C/S-same
  - Hosts that are client and server may be questionable
  - Look at changes over time - 1 hour bins
    - sudden changes suspicious
    - plot a week or more using FloVis Activity viewer
- Client starts conversations (TCP initial SYN)
- Server replies (TCP initial SYN/ACK)

REDJACK

Computing sets

- Client and server sets, with and without ports

```
rwfilter ... --flags-init=S/SAFR ... | \
cubag --set=cp.cus:v4sip,stime(/*,3600),dport \
        --set=c.cus:v4sip,stime(/*,3600)
```

- Server similar with SA/SAFR and sport
- Intersecting gets C/S, differencing gets C only and S only

```
cubagtool --intersect --output=cssp.cus cp.cus sp.cus
cubagtool --difference --output=cop.cus cp.cus cssp.cus
```

etc.

REDJACK

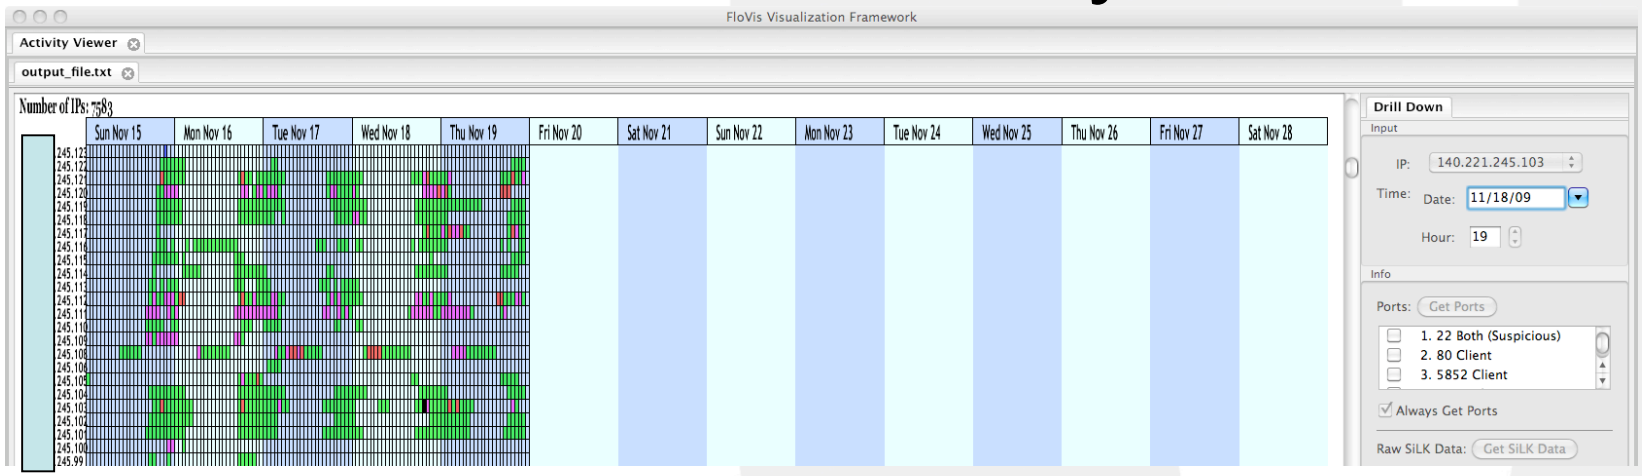
Two kinds of client / servers

- For a few services, it is normal for a host to be client and server (SMTP, DNS, etc.)
- For others, this may be suspicious
- We have sets of C, S, CS, with ports
  - the later are the CS on the same port
- We also have CS without port information
- Extract IPs from CS same port and difference with all CS to get CS on different ports

```
cubagtool --project:v4sip,stime --output=css.cus cssp.cus
cubagtool --difference --output=csd.cus cs.cus css.cus
```

REDJACK

Selected C / S activity results



What is it?

sIP	dIP	sPort	dPort	pro	pkts	bytes	initF	flags	sTime	dur
xxx.yyy.245.103	aaa.bbb.88.194	34359	22	6	725	55417	S	S PA	2009/11/18T19:28:09.845	163.961
aaa.bbb.88.194	xxx.yyy.245.103	22	34359	6	495	94839	S A	S PA	2009/11/18T19:28:09.894	163.912
ccc.ddd.118.175	xxx.yyy.245.103	15912	22	6	2	88	S	SR	2009/11/18T19:56:58.285	0.172
xxx.yyy.245.103	ccc.ddd.118.175	22	15912	6	1	48	S A	S A	2009/11/18T19:56:58.285	0.172
and later										
ccc.ddd.118.175	xxx.yyy 245.103	60076	22	6	3	132	S	S	2009/11/18T20:29:13.204	94.197
xxx.yyy.245.103	ccc.ddd.118.175	22	60076	6	8	352	S A	S A	2009/11/18T20:29:13.204	94.197

Harmless in this case, but worrisome nonetheless.

REDJACK

Dark Space

Dark space is unoccupied address space. Some organizations own large blocks of it. It is also the space between addresses in allocated space. The /22 that we observe has 117 active addresses, 899 that are dark (8 invisible). By filtering out the active addresses, we can look at the residue.

Note that the fact that there is legitimate activity in the space may provoke some of the dark space activity. Barford observed this a few years ago when he added activity to a previously dark /8. This data is from Feb. 2006 - Mar. 2007. Large scale collection failure in Aug. and Nov.

REDJACK

Who is there? What are they doing?

- TCP Scanners; Outside to dark for SYN only
  - Sets v4sIP,hour; v4sIP,dPort; v4sIP,v4dIP; dPort,v4sIP; dport,v4dIP; dPort,hour; hour,v4sIP; hour,v4dIP; hour,dPort; sCC,hour; sCC,dPort, sCC,v4dIP
  - Bags v4sIP:flows,pkts,bytes,span; dPort:f,p,b,span hour:f,p,b; sCC:f,p,b,span
- Project second field off sets; add count -> bag
- Join all bags with same key
  - gives bags by v4sIP, dPort, hour, sCC
  - counts from rollups and volumes, span
  - sort by field(s) of interest
- We present results by dPort only here

REDJACK

Dark space results - Sources/dPort

dPort	SrcCnt	DstCnt	Hours	Span	Flows	packets	bytes
445	92314	899	6609	397T15:07:29	454525	2510187	122636688
4662	34602	12	696	345T15:39:41	87837	261776	12980056
139	16896	899	4347	397T13:50:27	334937	1513998	73964832
80	16422	899	6065	347T05:43:03	266119	935422	45768460
1433	13676	899	1315	386T09:14:32	434305	1961423	94348040
9272	12452	1	151	7T00:31:15	22995	66551	3270628
47190	5098	1	18	0T16:24:17	15406	46316	2281044
34001	4551	1	46	2T15:18:13	18196	51396	2516076
135	3593	899	1550	389T20:54:06	111191	262077	12612116
14662	3568	1	53	3T02:39:44	7593	22776	1124904
2967	3042	899	506	82T02:48:06	38554	86082	4165340
23	2962	894	105	390T09:40:04	9077	28373	1426464
5900	2579	899	774	395T03:18:28	367558	1647276	79555252

REDJACK

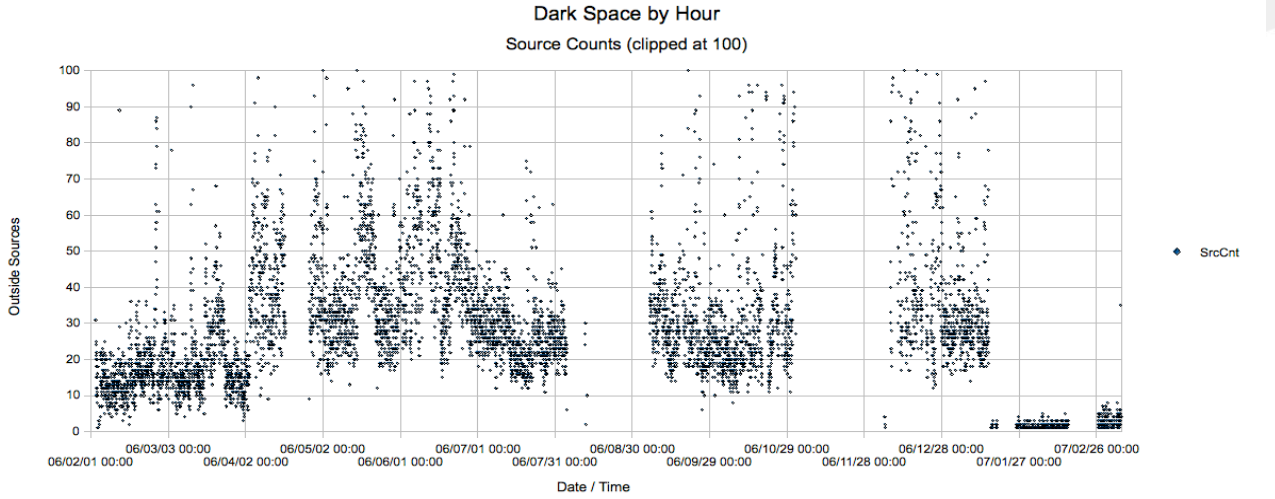
Dark space results - Hours/dPort

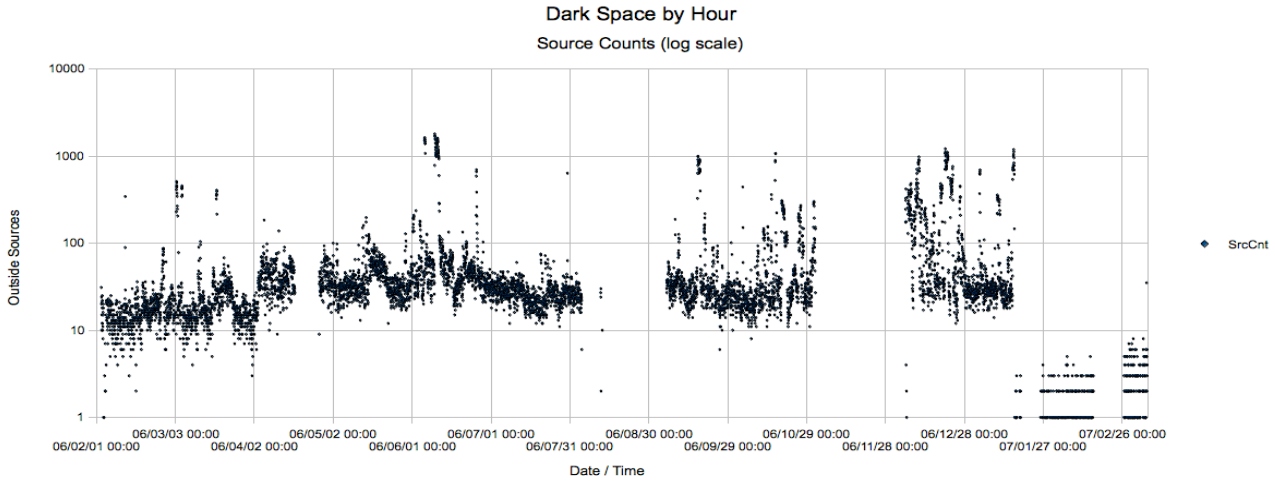
dPort	SrcCnt	DstCnt	*Hours*	Span	Flows	packets	bytes
445	92314	899	*6609*	397T15:07:29	454525	2510187	122636688
80	16422	899	*6065*	347T05:43:03	266119	935422	45768460
139	16896	899	*4347*	397T13:50:27	334937	1513998	73964832
3157	1612	1	*2683*	264T07:00:52	12206	35463	1707028
12879	1516	1	*2511*	264T06:54:34	11255	29923	1446792
1080	135	899	*2244*	376T16:22:58	19659	48549	2256972
43631	1381	1	*2200*	263T23:36:29	7396	19761	975604

Almost 1800 ports scanned. Hosts/port and hours/port vary widely
Cases with large number of sources, small number of targets unexplained. Port *du jour* effect also visible with short span.
SrcCnt = 100 at rank 54, 10 at rank 138, 1 at rank 540
DstCnt = 100 at rank 136, 10 at rank 171, 1 at rank 329
Hours = 100 at rank 62, 10 at rank 179, 1 at rank 661
Most activity is low frequency in some or all dimensions!

REDJACK

Dark Space Sources per hour





Sources per hour are in the 10-100 range most of the time

Bursts of up to about 2000 sources per hour occur irregularly and persist for several hours.

Conclusions

REDJACK

- Multikey sets and bags support complex analysis tasks
  - Time keys simplify multiperiod analysis
    - Eliminate false zeros of `rwcount`
  - Rich key set allows (almost) arbitrary viewpoints
  - Projection and join bring disparate sources together
  - Simple scripts can serve as report generation drivers
  - Stream operations remove memory size based limits
- Predictable space / time performance for real time
  - Can take advantage of multicore processors
  - Constant time lookup for filter applications
  - Arbitrary key fields
  - Adaptable to packet or other streams

REDJACK

Acknowledgements

- Funding, in part via a grant from Cisco Systems to the University of North Carolina
  - Thanks especially to Henry Stern of Ironport
- The FloVis project at Dalhousie for providing applications and data
- Ron McLeod of TARA
- RedJack for encouragement, support, and ongoing interest.
  - Greg Virgin, Michael Collins, Doug Creager

REDJACK

Questions?



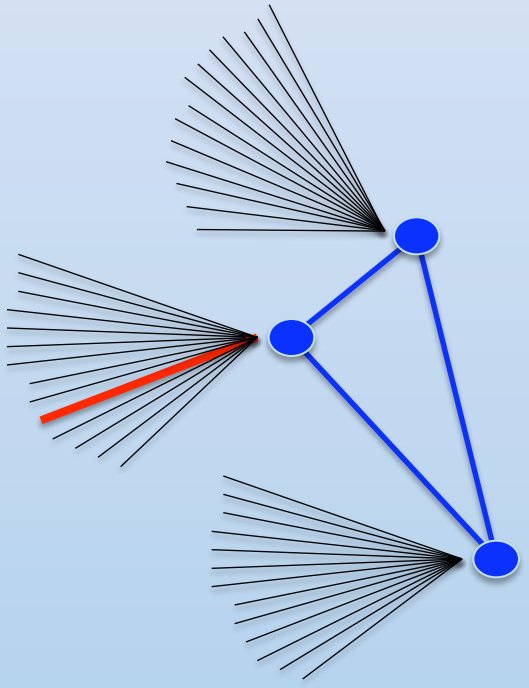
DMnet: Detection Mitigation Network

A Behavioral Analysis System Supporting Trust Measurements

Owen McCusker, Scott Brunza,
Sonalysts, Inc.
{mccusker,brunza}@sonalysts.com

Dr. Carrie Gates,
CA Labs
carrie.gates@ca.com

Joel Glanfield, Diana Paterson,
Dalhousie U.
{joelglanfield, paterson}@cs.dal.ca



An ADAPTIVE DISTRIBUTED DYNAMIC Approach

To an ADAPTIVE DISTRIBUTED DYNAMIC Threat

SBIR DATA RIGHTS Contract No.: NBCHC070124
Contractor Name: Sonalysts, Inc.
Contractor Address: 215 Parkway N., Waterford, CT. 06815
Expiration of SBIR Data Rights Period: 14 Jan 2011

The Government's rights to use, modify, reproduce, release, perform, display, or disclose technical data or computer software marked with this legend are restricted during the period shown as provided in paragraph (b)(4) of the Rights in Noncommercial Technical Data and Computer Software--Small Business Innovative Research (SBIR) Program clause contained in the above identified contract. No restrictions apply after the expiration date shown above. Any reproduction of technical data, computer software, or portions thereof marked with this legend must also reproduce the markings

COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

SONALYSTS

Agenda

- The Need
- Behavior-based Classification
- Trust Derived from Behavior
- Where We are Today

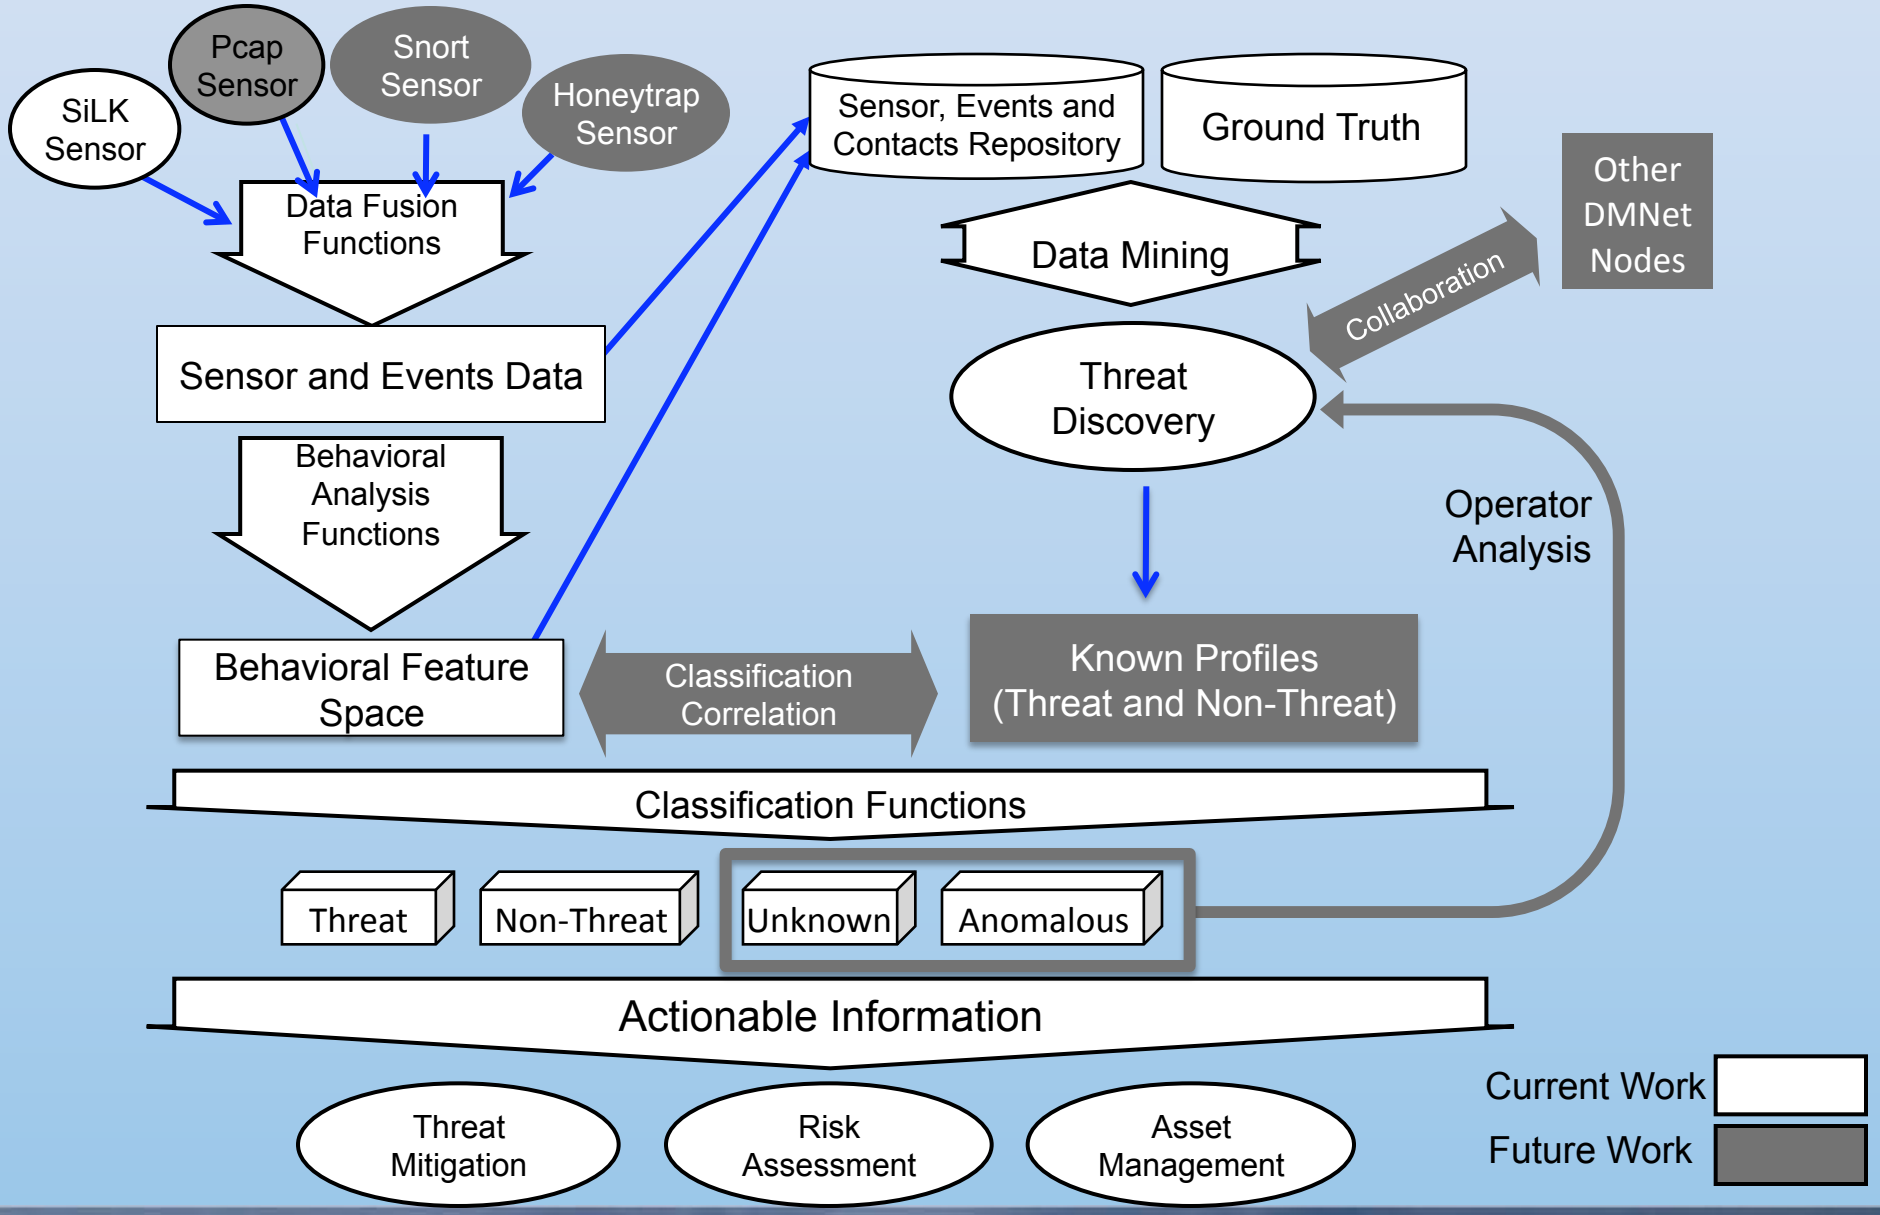
COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

The Need

- The cyber threat is distributed, dynamic, and multi-scale in time
  - IDS technology is focused on "**single source**" solutions, "**single time-scales**".
  - Threats are buried in the noise of everyday traffic
  - Cyber defense technologies adapt mostly through the use of signatures (exception: Anomaly Detection)
- We need enabling technologies that facilitate the creation of **adaptive** and **open** distributed defense technologies
- Our Contribution:
  - Creation of an **aggregated behavioral feature space**
  - **Separation** of **trust** from **behavior aggregation**
  - Initial use of **ontology** to map behavior to threat
  - Share behaviors between COIs to break through privacy barriers

COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Approach: DMnet Architecture



COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Behavior-based Classification

- Ingest events and data from multiple sensor types
  - Architecture supports different sensor types
  - Currently using SiLK
- Derive features from capture events and data
  - Create a rich feature space used for behavioral analysis
  - Leverage primitive features during analysis
- Identify Behaviors
  - The Goal is to create a **behavioral language** used to describe and identify cyber threats
  - Based on analyzing feature space using different n-tuple sets.

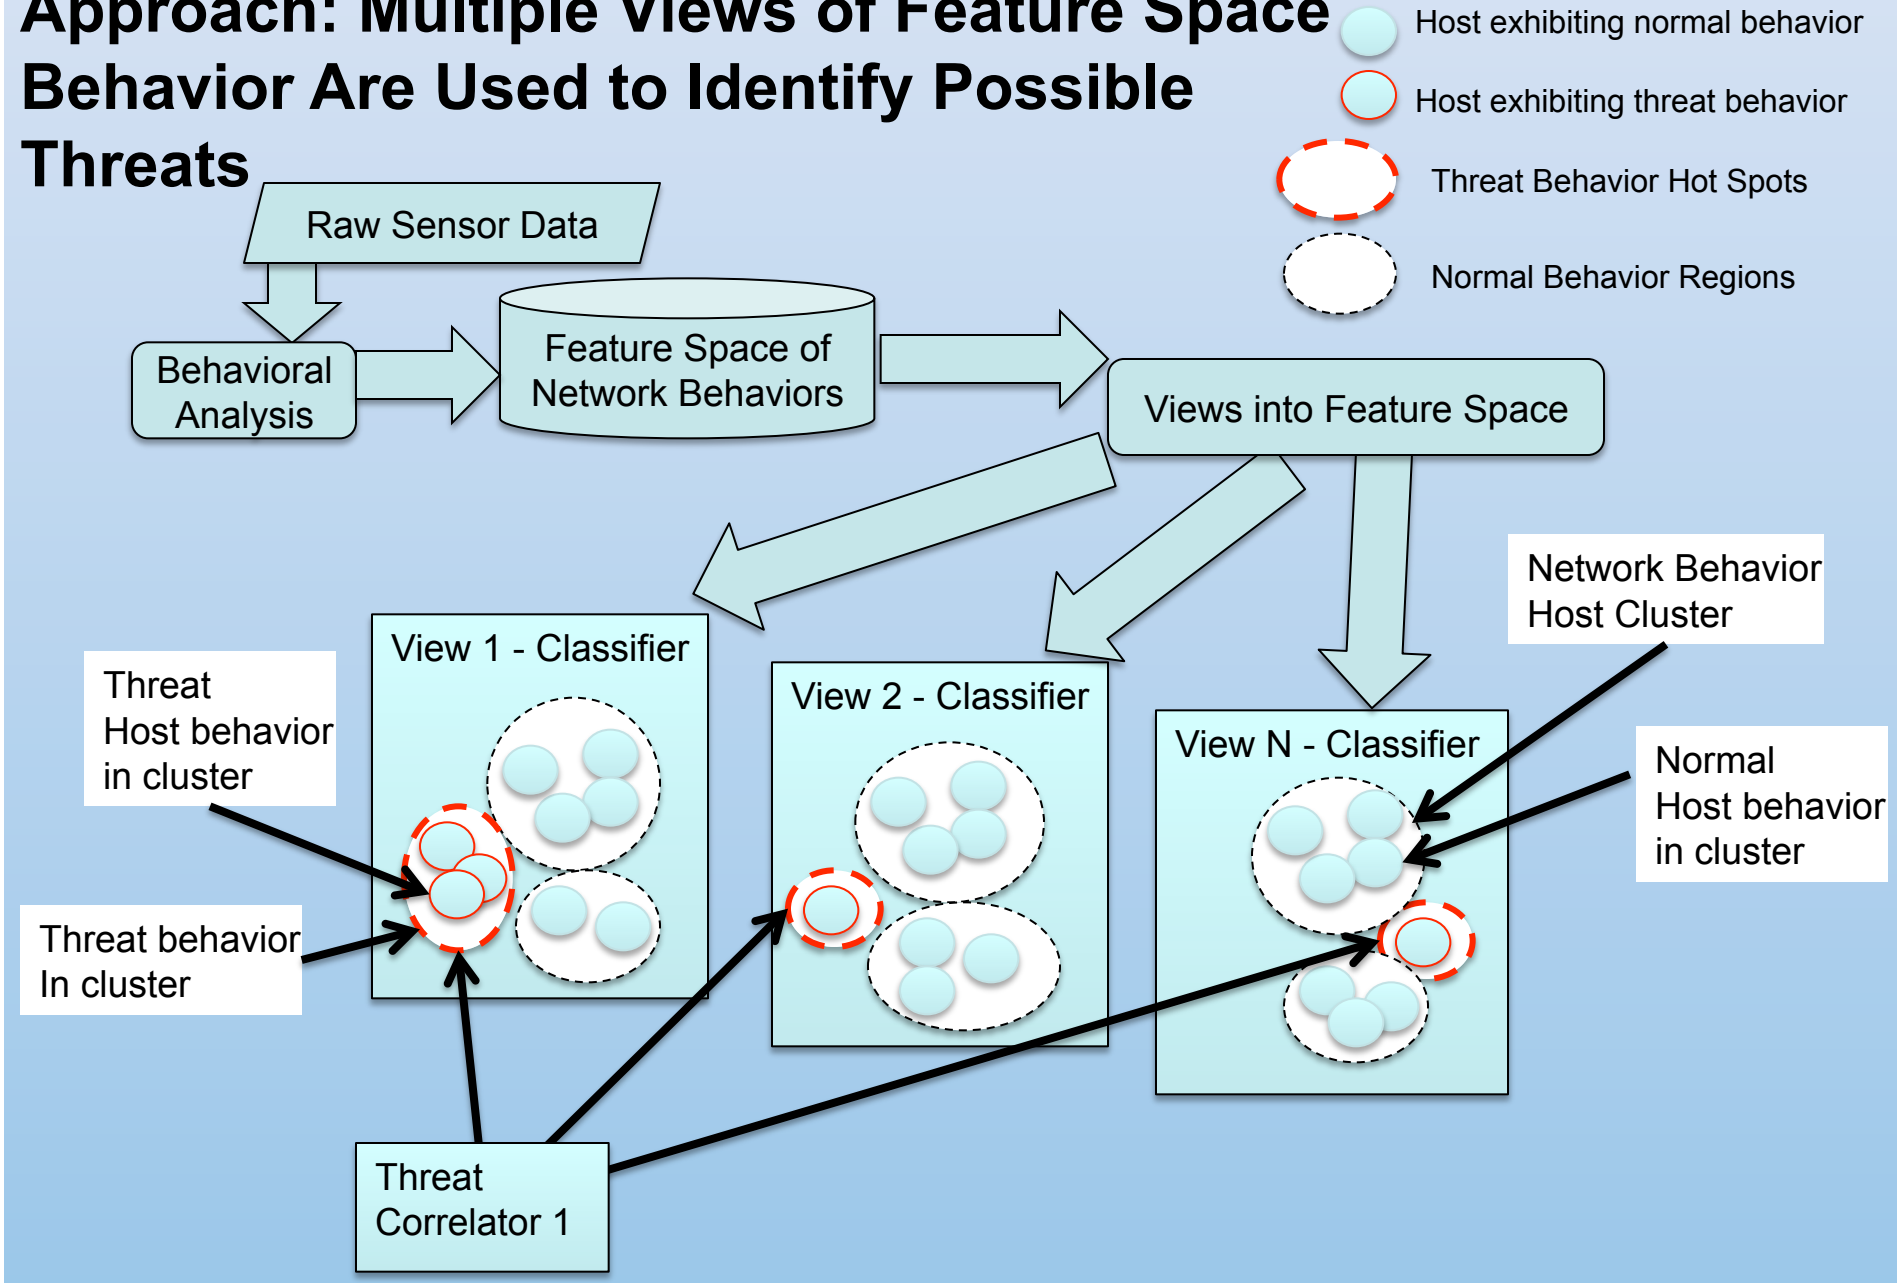
COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Behavior-based Classification

- Threats are detected by the **identification of multiple behavioral patterns**
  - Has an “analog” to OCR and voice recognition
- Threat behavior can be characterized / detected using **adaptive heuristics**.
  - Success despite primitive state of current rule set.
  - Architecture will support concurrent use of more complex and adaptive heuristics.
- Analysis can be **enhanced** by:
  - **Increased Community of Interest size** (number of correlated network sensors),
- **Automation** necessary to improve / expand analysis.
  - There is a need for a common behavioral **Ontology**
- **Application of “Hyperplaning” for Botnet Detection**
  - University of Connecticut

COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

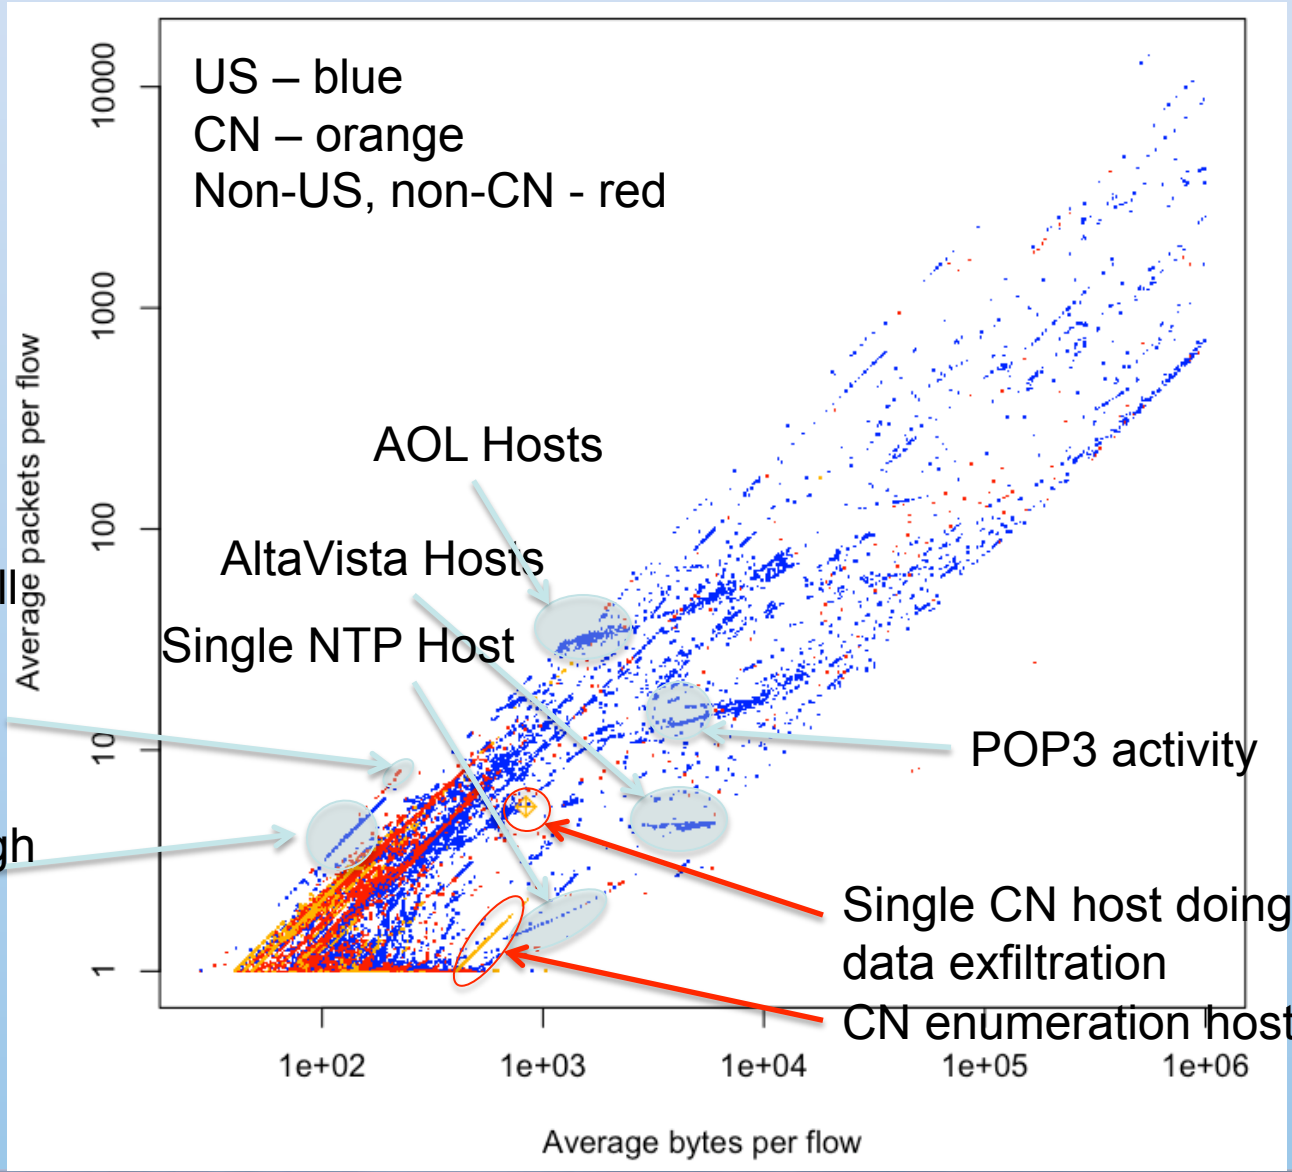
Approach: Multiple Views of Feature Space Behavior Are Used to Identify Possible Threats



COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Benefit: Weak Signal Cyber Detection (Threat signals stand out of the Noise)

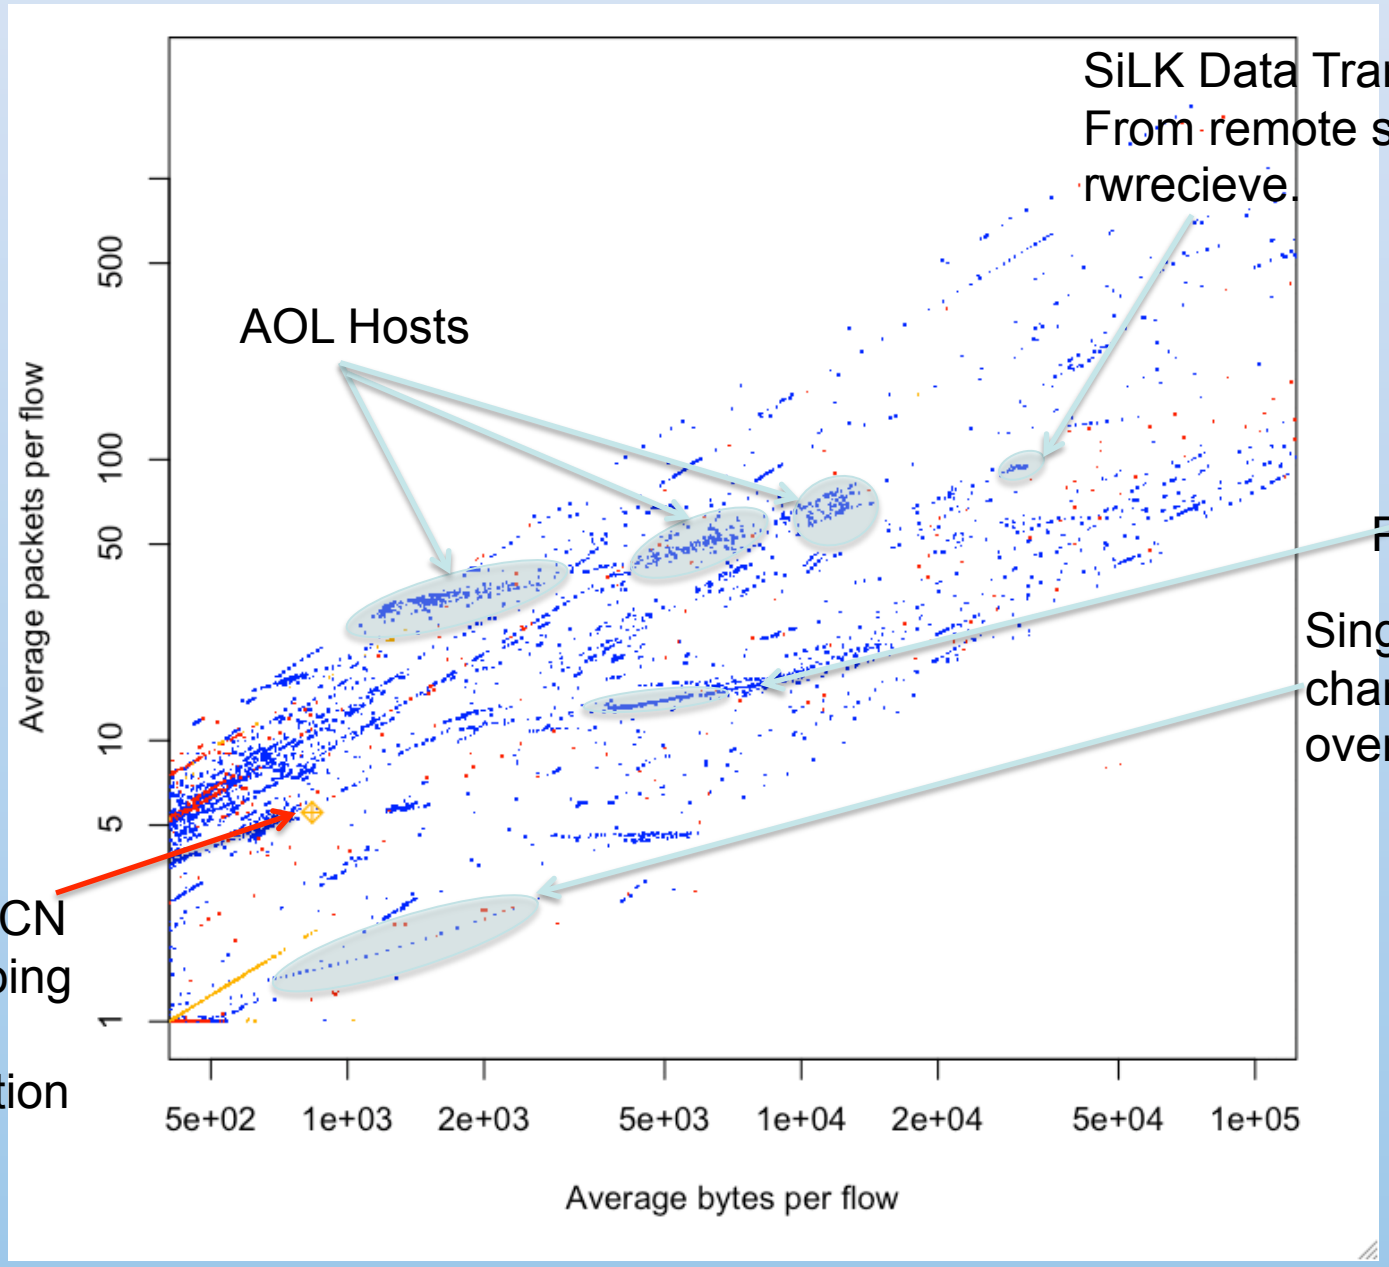
CN Data Exfiltration Case Post-Event Comparison of Selected Host



Pings hitting firewall from around globe. Panther Express included.

Hosts NAT'd through firewall

COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

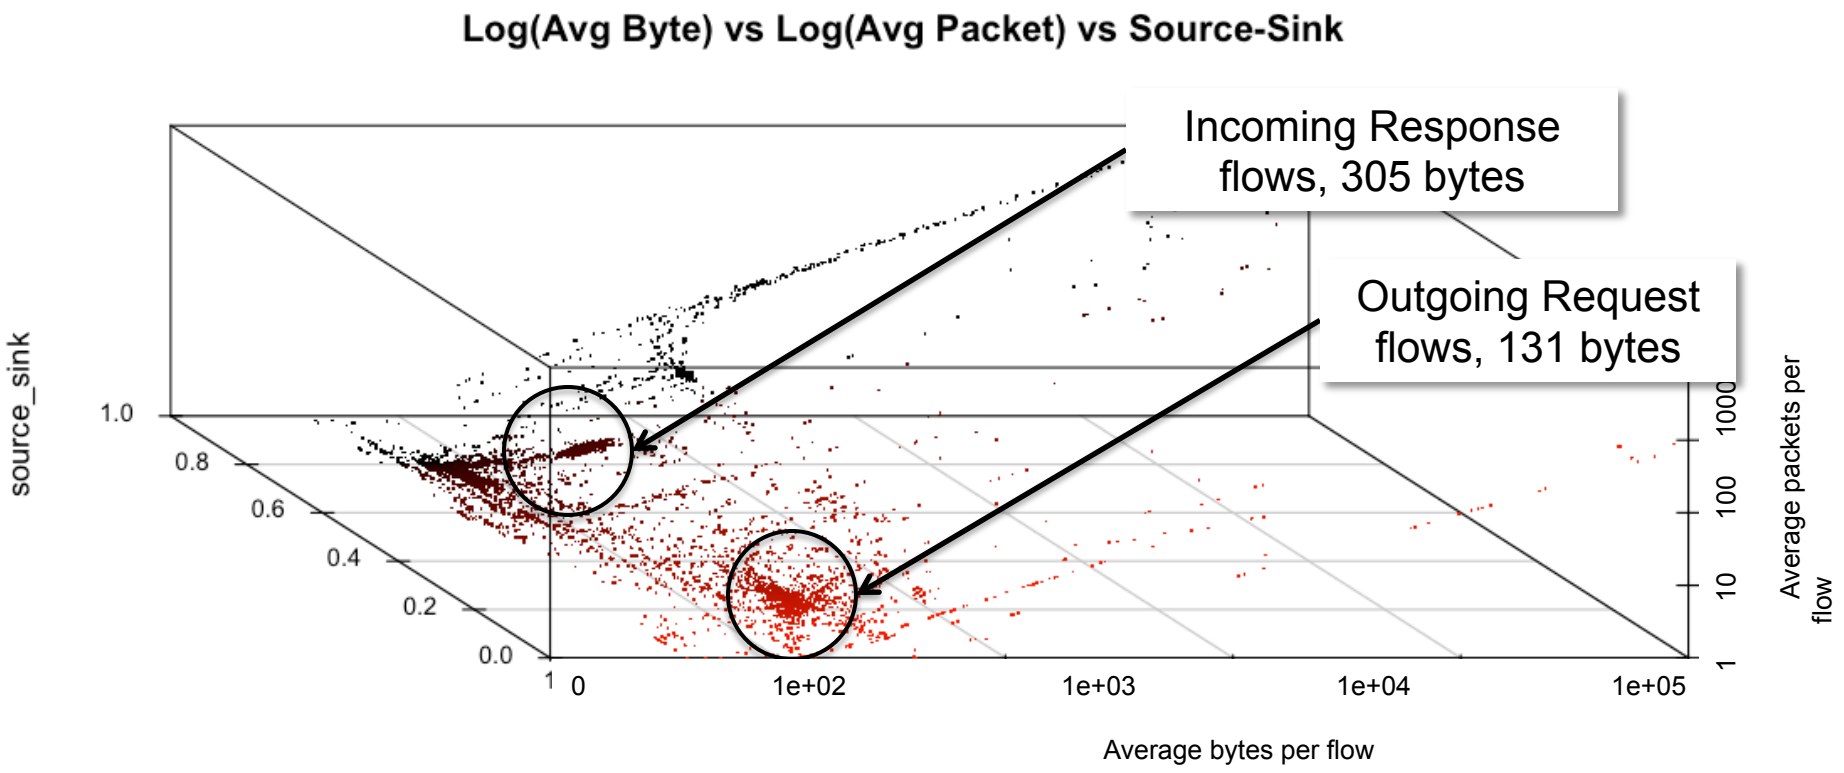


COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Differentiator: Adding in another visual dimension (source/sink) to separate out host behaviors

October (ingested 5 days)
Total Contacts: 251K

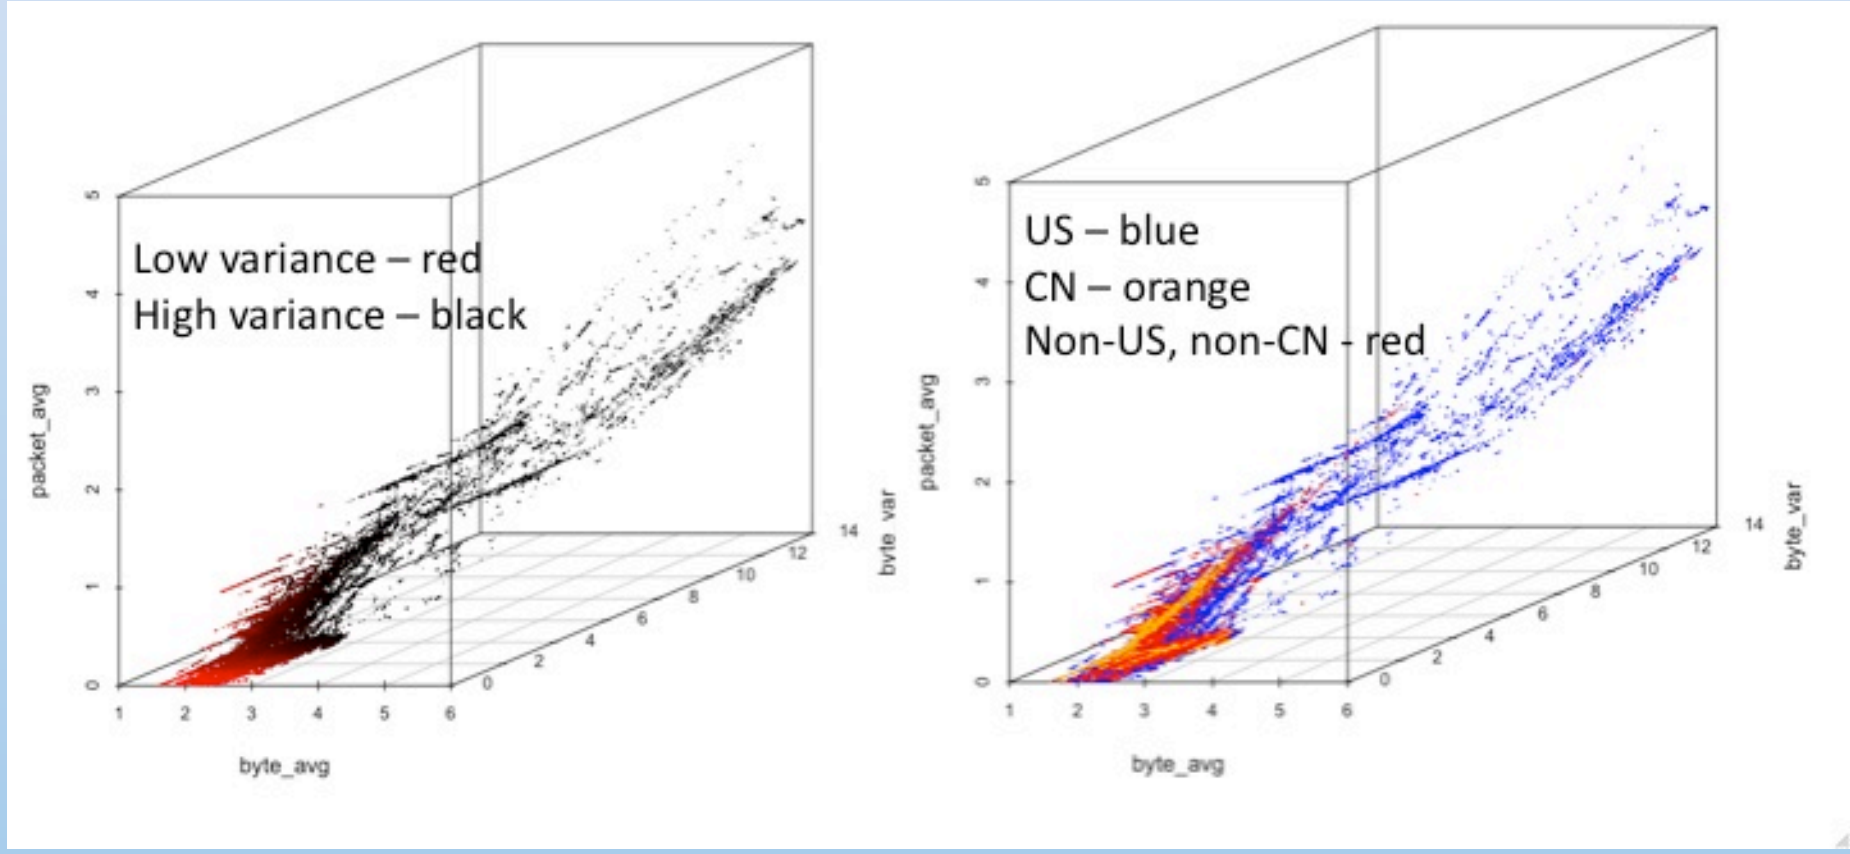
Red is source, Black is Sink
DHT Protocol



COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

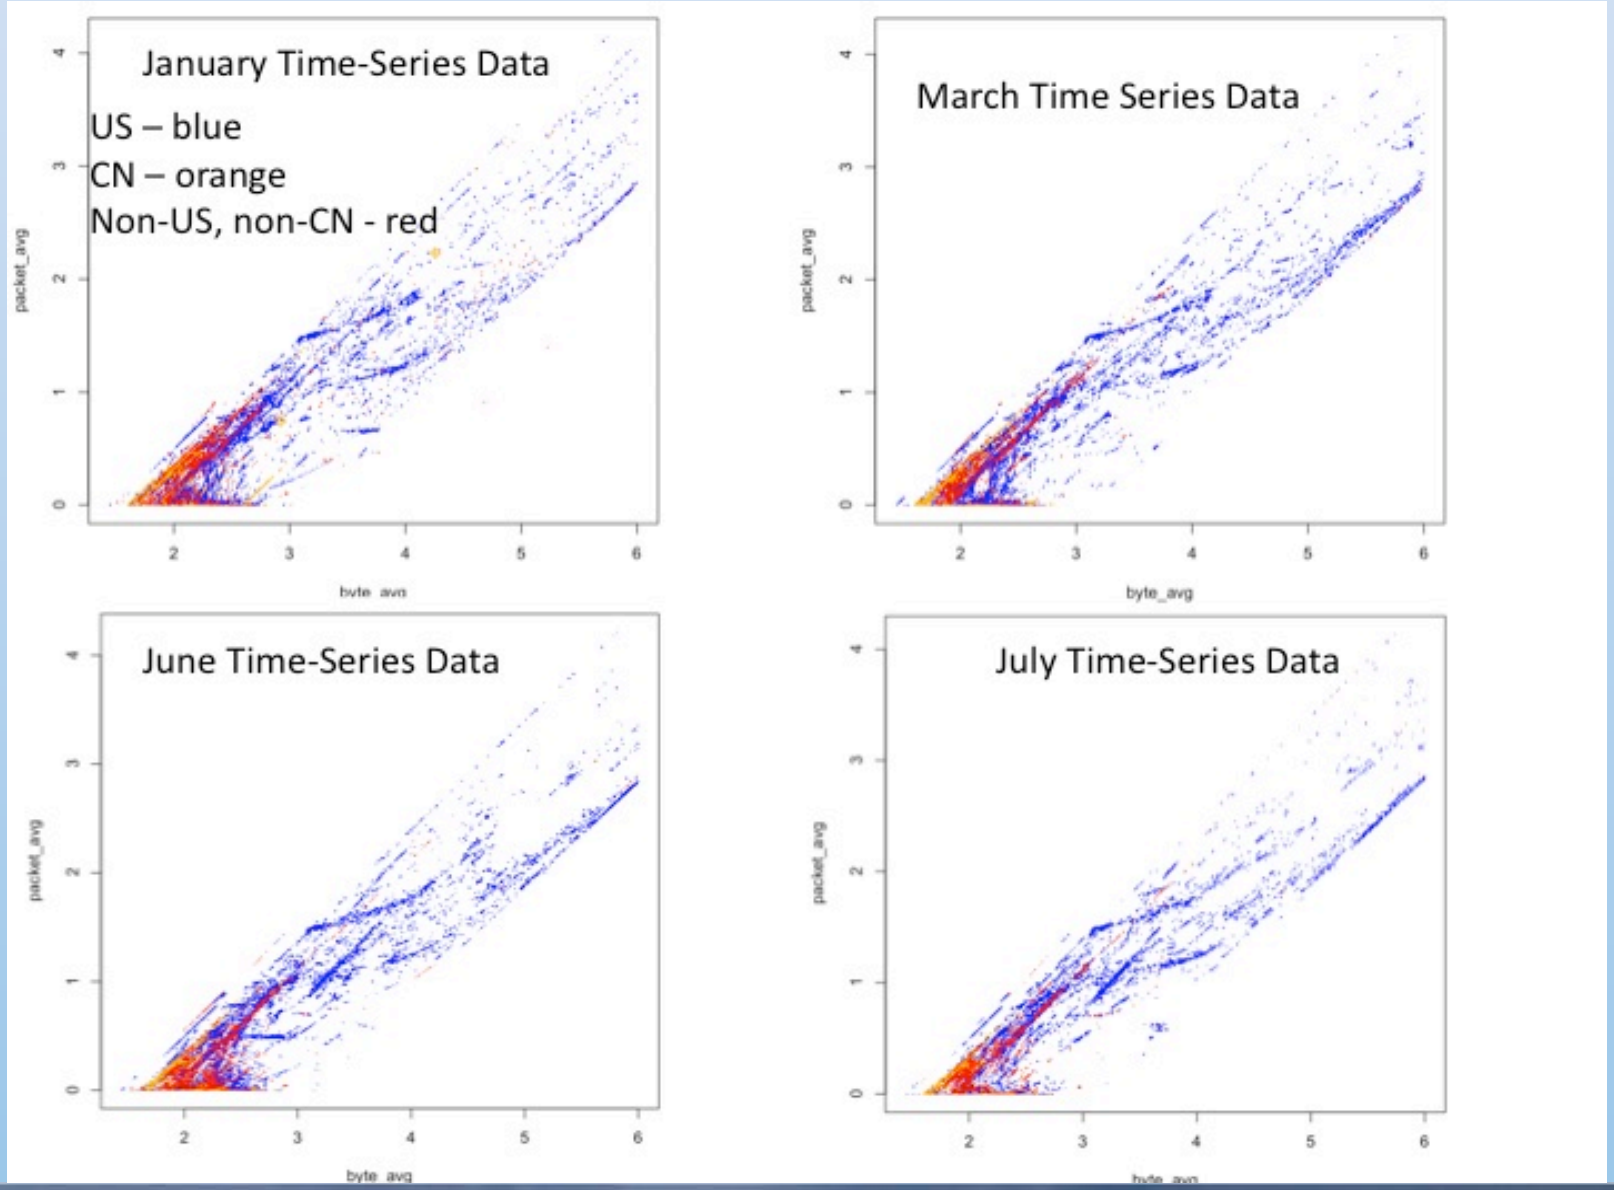
SONALYSTS

Differentiator: Adding in another visual dimension (byte variance) to separate out host behaviors



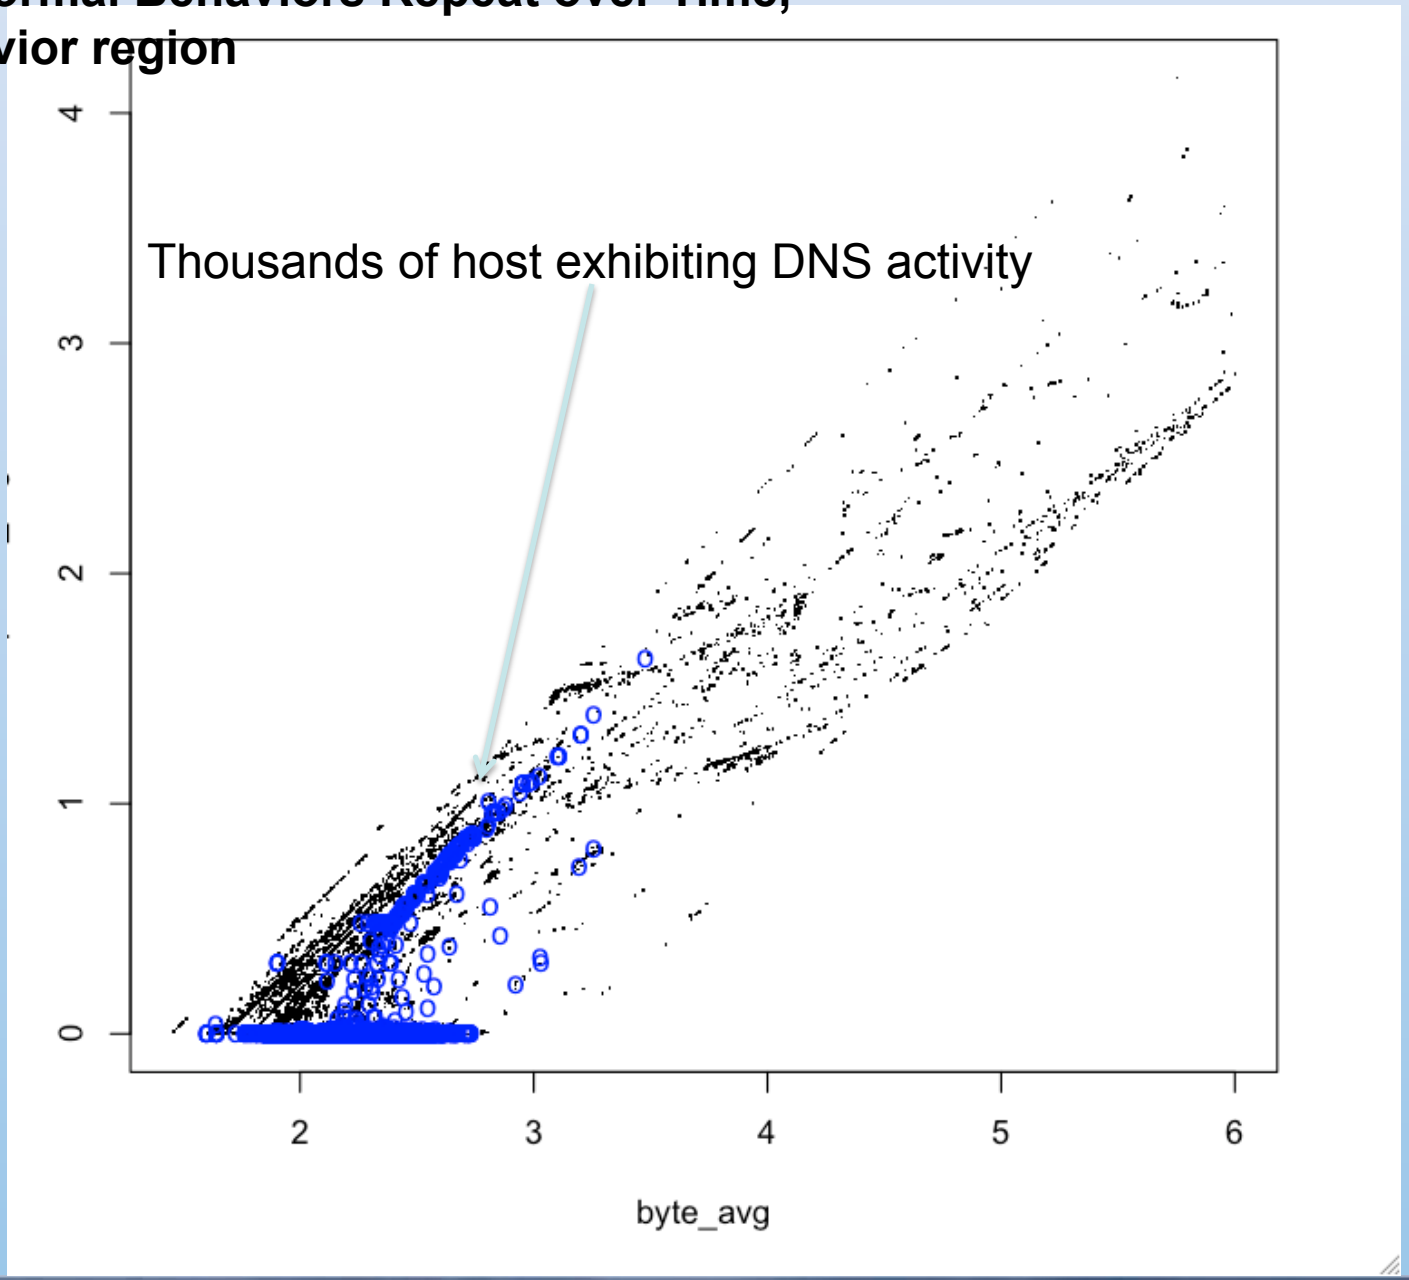
COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Benefit: Normal Behaviors Repeat over Time, DNS behavior region



COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

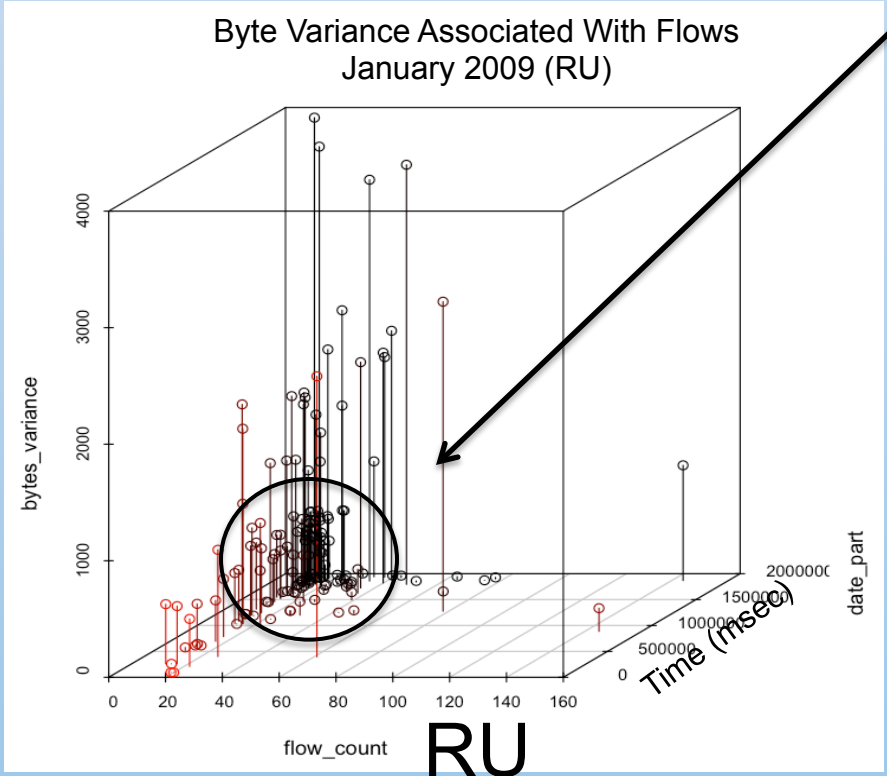
Benefit: Normal Behaviors Repeat over Time, DNS behavior region

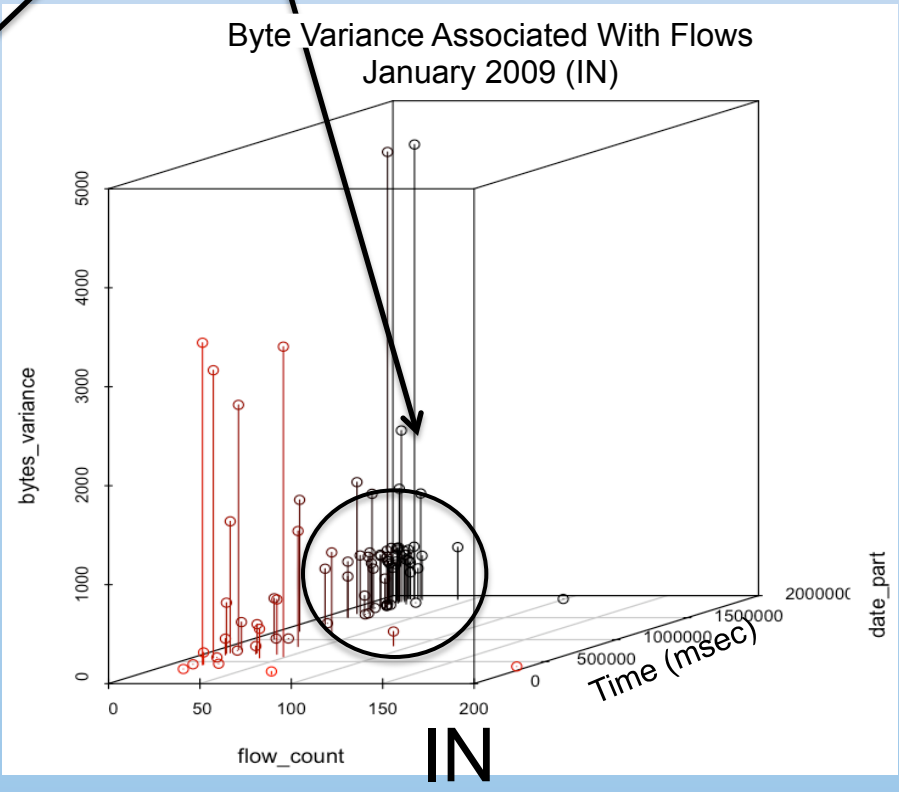


COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Benefit: Threat Anticipation Regional Byte Variance Over Time

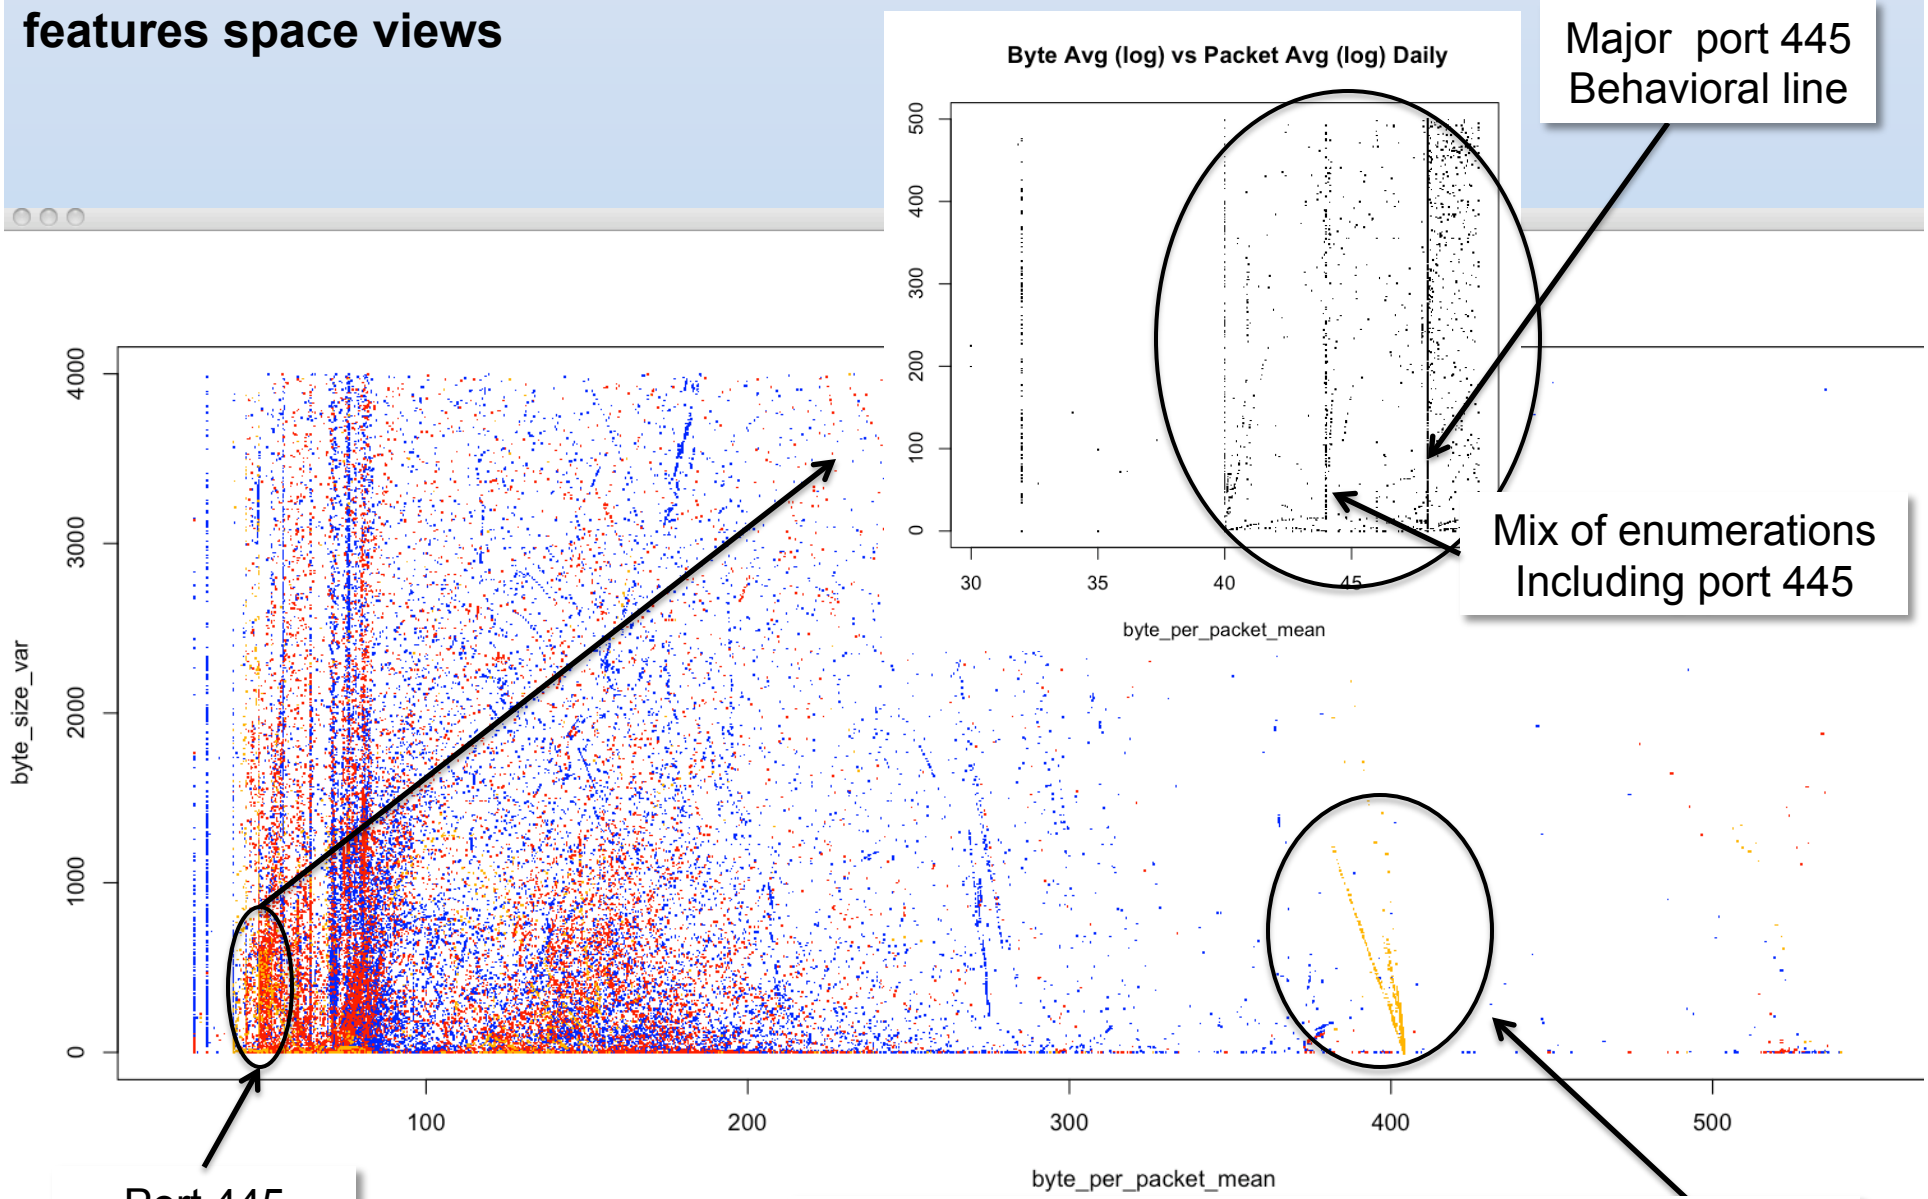
Unexplained / suspect behavior originating from two different regions at the same time.





COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Benefit: Consistent Behaviors within features space views



COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERV

Trust Derived from Behavior

- Trust is **subjective** driven by the security policies of the institution
  - Host network behaviors are **objective**
- Trust is too difficult to share without a common understanding of risk
- The overall trust of a host is a weighted sum of all trust behaviors
  - Each **measured behavior** is given a value of **trust**
- The change in behaviors can be used as a measure in trust
  - Use of multiple protocols, compared to single protocol
  - High variance in byte usage
  - High variance in entropy of payload

COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

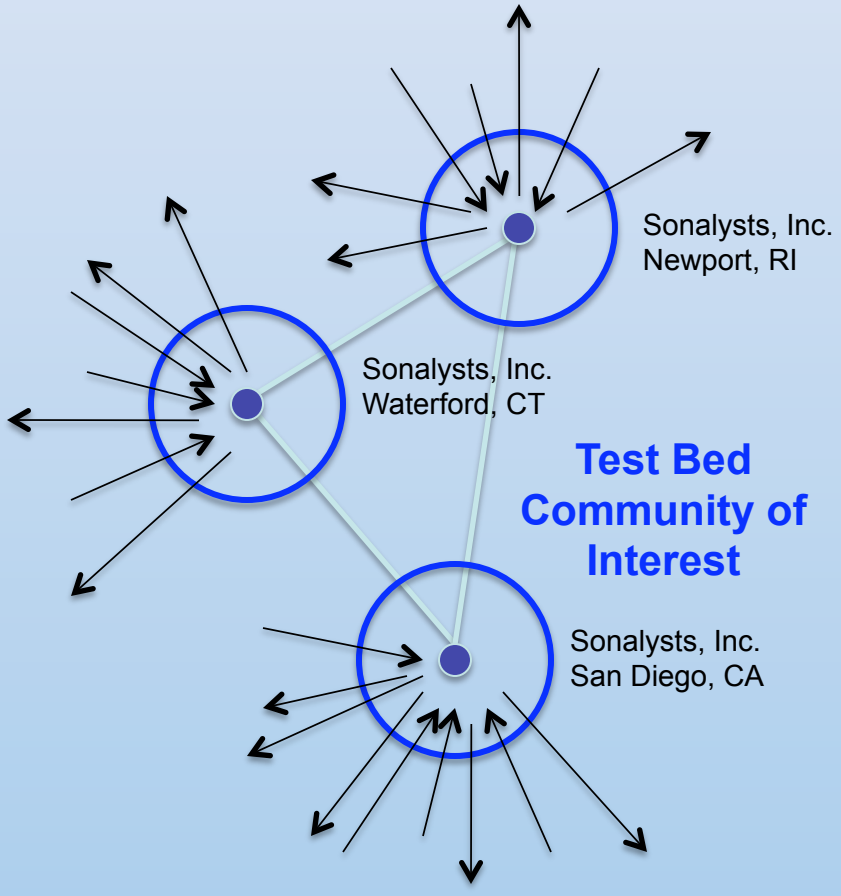
Where are we Today?

- ITT – **3rd Party evaluation**
  - Web-base interface, CLI tool
  - Identified threats using a simplified set of heuristics
  - Solid system, more work needed false positives/negatives
- Researching the application of **Biological Immune System (BIS)** concepts to system
  - **Self, Non-self concepts** combined with **Computational Trust**
- Created a **commercial** service for network analysis
  - **Behavioral Analysis leading to Situational Awareness**
- HPC based architecture
  - Created a small cluster of nodes using **OpenMPI** to test scaling our system
- High Bandwidth
  - Just starting integrating and **Endace DAG** for network capture

COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Approach: Test Bed

- Using:
  - Networked sensors leveraged across a trusted community of interest,
  - **Host-centric behavioral aggregation**
  - Multiple data fusion and mining methodologies, and
  - Concurrent classification and correlation algorithms.
- To:
  - Connect the dots over long time periods (e.g. months).
  - Detect and characterize threat behavior in near real time.
  - Perform weak signal cyber detection.
- Sharing **just** behaviors **minimizes** impact of **user data privacy**

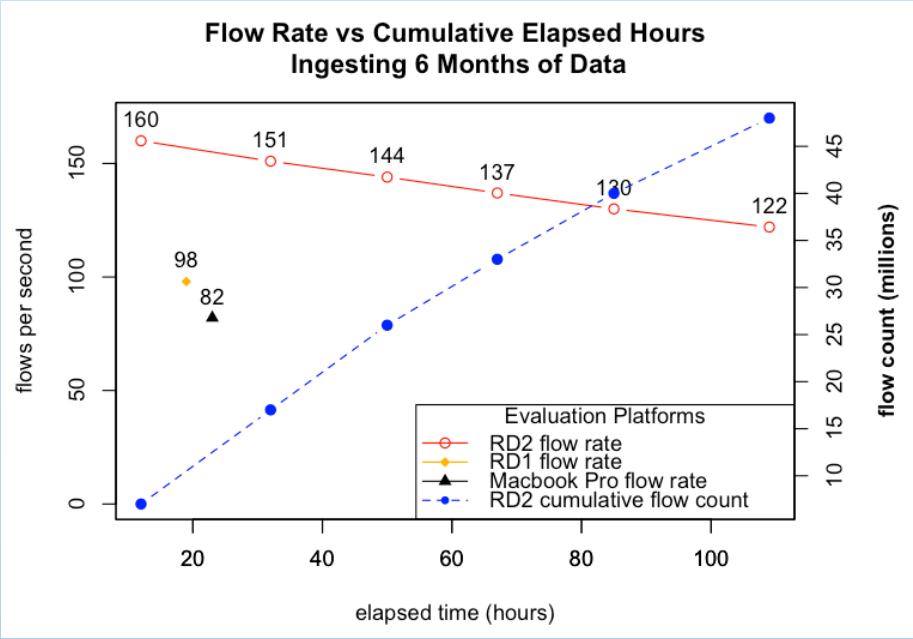


Today:

- Monitor incoming & outgoing traffic outside the perimeter.
- ~400,000 host contacts / month.
- Characterize host behavior and look for changes over time that suggest threat behavior.
- 72 basic characteristics extracted and synthesized.
- Correlate with country of origin.

COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

DMnet – RD2 Prototype





Prototype Evaluation

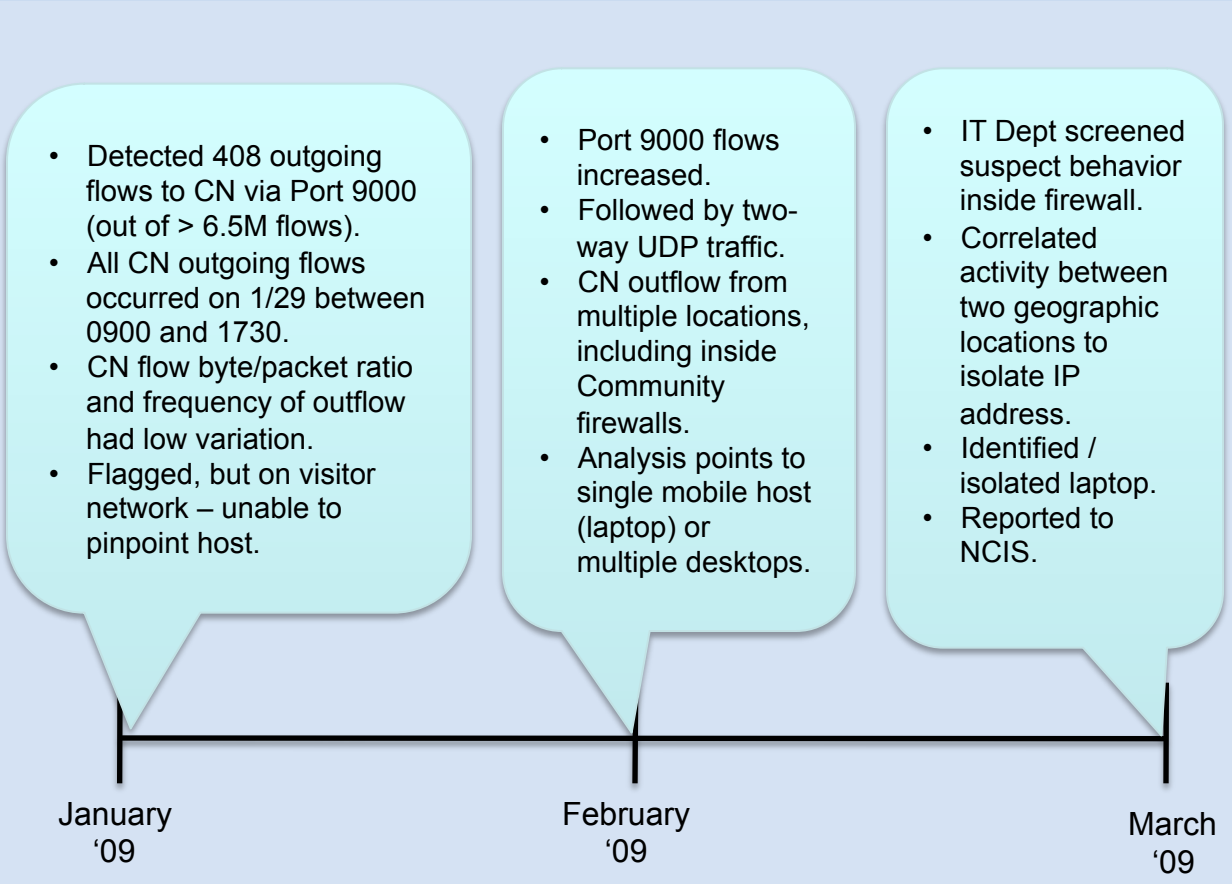
- Processed 6 months of data on site in 4 days
  - 30 times real time for Sonalysts
  - Improved performance from last year
- TODO:
  - Integrate behavioral learning
  - Integrate classifiers and correlators

COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Benefit: Successful Detection of Data Exfiltration to a Sophisticated "CN Host" via Company Laptop

Monthly, we conduct semi-automated regional analysis of outgoing network flow.

- Analysis script looks at all incoming and outgoing data to community of interest.
- Heuristics set used to parse data into manageable subsets.
  - Based on location, port usage, direction, port pair bandwidth utilization, IP address bandwidth usage, client-server behavior, protocol analysis.
- Manually review subsets looking for anomalies
  - Trend analysis of past reports.



CN exfiltration was not and likely could not have been detected by existing firewall technology.

COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Thank you!

Owen McCusker, Scott Brunza,
Sonalysts, Inc.
{mccusker,brunza}@sonalysts.com

Dr. Carrie Gates,
CA Labs
carrie.gates@ca.com

Joel Glanfield, Diana Paterson,
Dalhousie U.
{joelglanfield, paterson}@cs.dal.ca

COPYRIGHT 2009 SONALYSTS, INC ALL RIGHTS RESERVED

Simply Top Talkers

Jeroen Massar, Andreas Kind and Marc Ph. Stoecklin



© 2009 IBM Corporation

Motivation and Outline

- Need to understand and correctly handle dominant aspects within the overall flow of traffic
  - Top-k Problem
  - Optimize peering relationships (top autonomous systems)
  - Analyze congestion cases (top hosts)
- Technical challenge
  - Compute sorted top-n views from very large numbers of flow information records
  - No possibility to store individual counters per aspect components
- This presentation
  - Propose and analyze simple techniques to compute top-k listings for single and composed traffic aspects
  - Show that the techniques are suitable for in-memory aggregation databases

© 2010 IBM Corporation

Related Work

- E. D. Demaine, A. Lopez-Ortiz, and J. I. Munro. Frequency estimation of internet packet streams with limited space. In *Proceedings of the 10th Annual European Symposium on Algorithms*, pages 348–360, 2002.

- C. Estan and G. Varghese. New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice. *ACM Trans. Comput. Syst.*, 21(3):270–313, 2003.

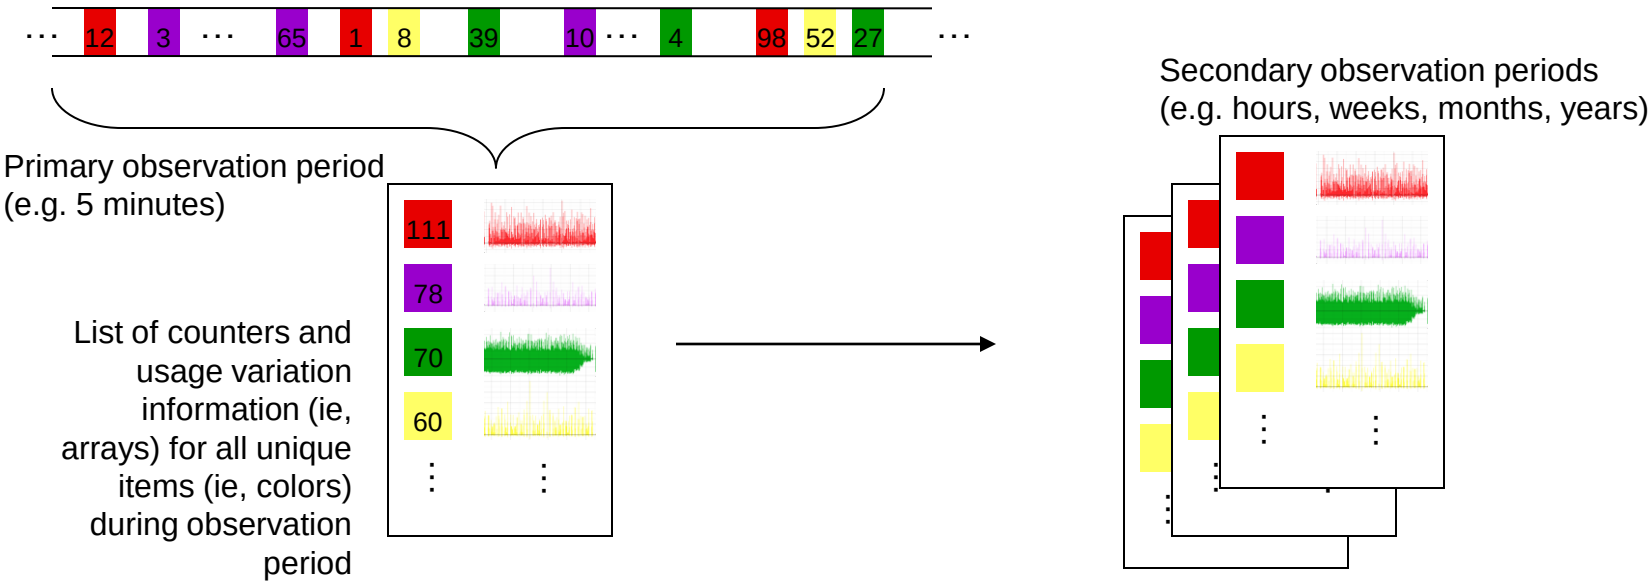
- Metwally A, Agrawal D, El Abbadi A (2005) Efficient computation of frequent and top-k elements in data streams. In: Proceeding of the 2005 international conference on database theory (ICDT'05), Edinburgh, UK, pp 398–412

- X. Dimitropoulos, P. Hurley, and A. Kind: Probabilistic Lossy Counting: An Efficient Algorithm for Finding Heavy Hitters," ACM SIGCOMM Computer Communication Review, Jan. 2008.

- X. Dimitropoulos, M. Stoecklin, P. Hurley, and A. Kind: The Eternal Sunshine of the Sketch Data Structure," Elsevier Computer Networks, 2008.

© 2010 IBM Corporation

Top-k Problem

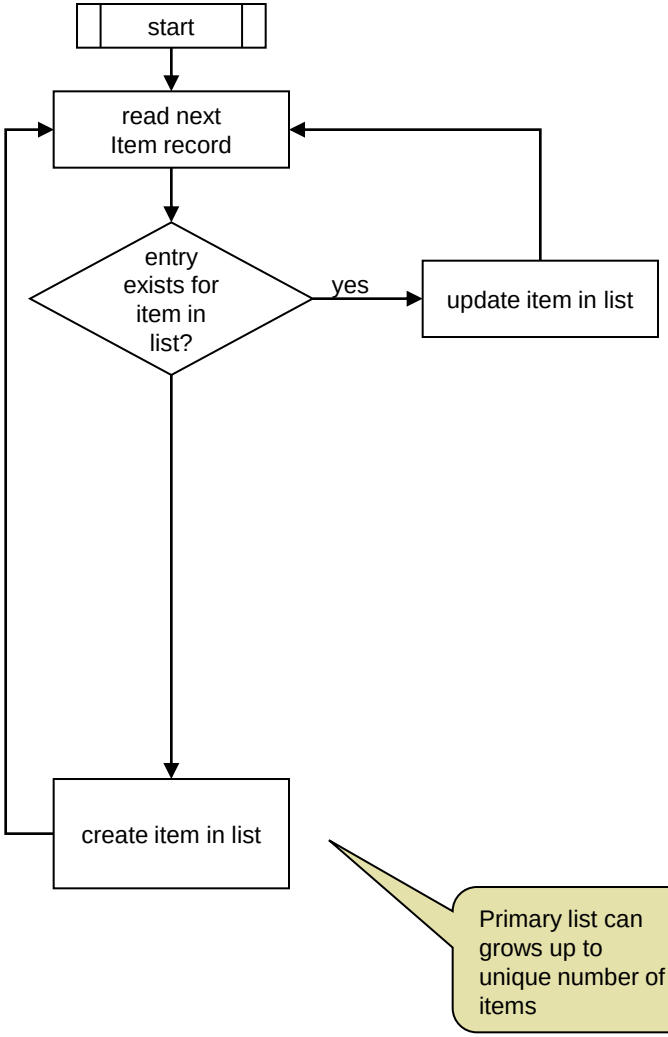


- Large number of unique items per observation period
  - Large alphabet
  - Counters and usage arrays don't fit into memory

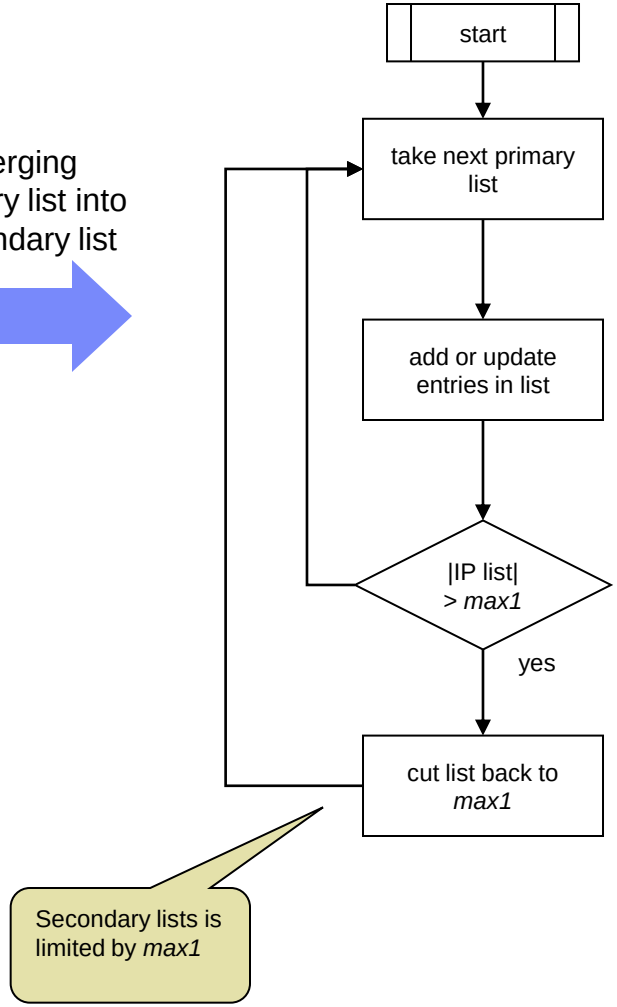
- How to cut back sorted list?
  - Correct ordering
  - Correct volume information for each item
  - Reduce the cost of inserting new items (that anyway would not make it into the sorted top-k list)

© 2010 IBM Corporation

One Threshold Scheme (*max1*)

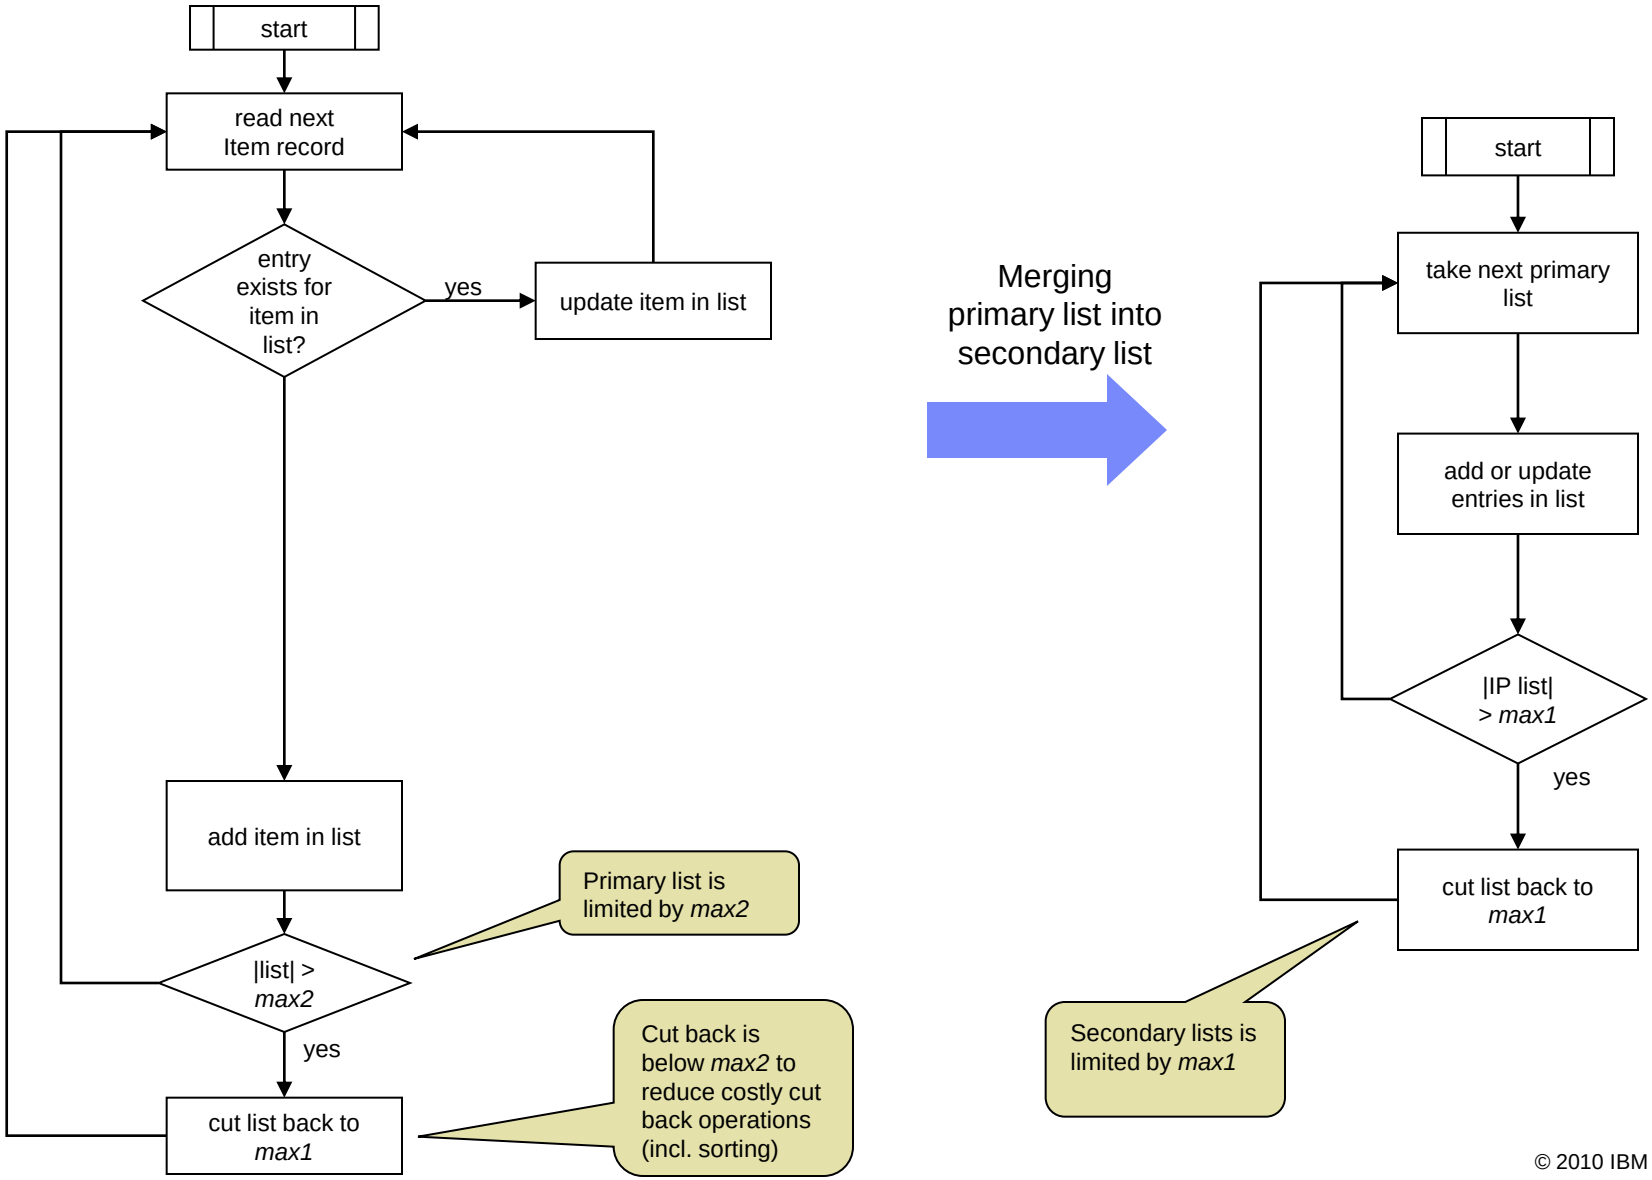


Merging primary list into secondary list



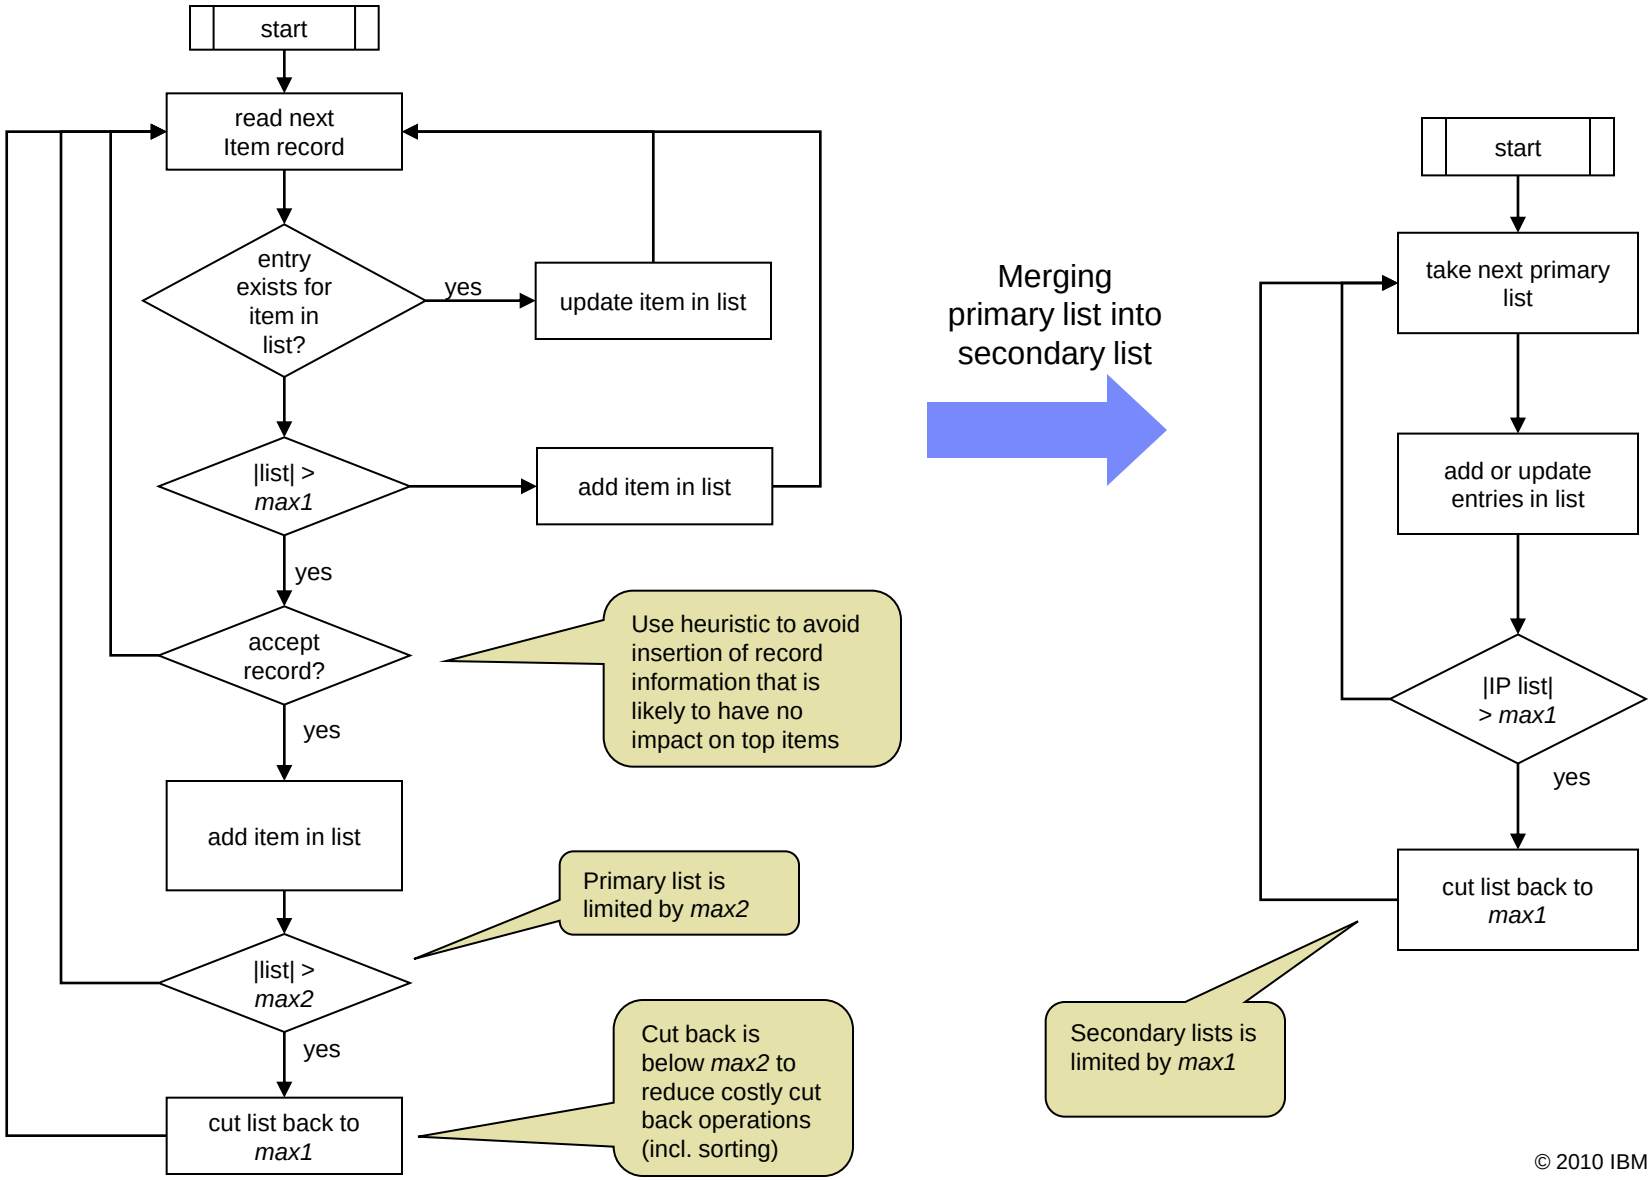
© 2010 IBM Corporation

Two Threshold Scheme (*max1*, *max2*)



© 2010 IBM Corporation

Two Threshold Scheme (*max1*, *max2*) with Heuristic



© 2010 IBM Corporation

Heuristics

- 0. Exact

- max1,max2 style:
- No Heuristic
- Short duration (< 1 second)
- Low Packet count in one flow (< 4 pkts)
- Bytes < total_bytes_seen/total_flows_seen (averaging)

- These heuristics are very simple and more importantly ‘cheap’ to implement memory and cpu wise.

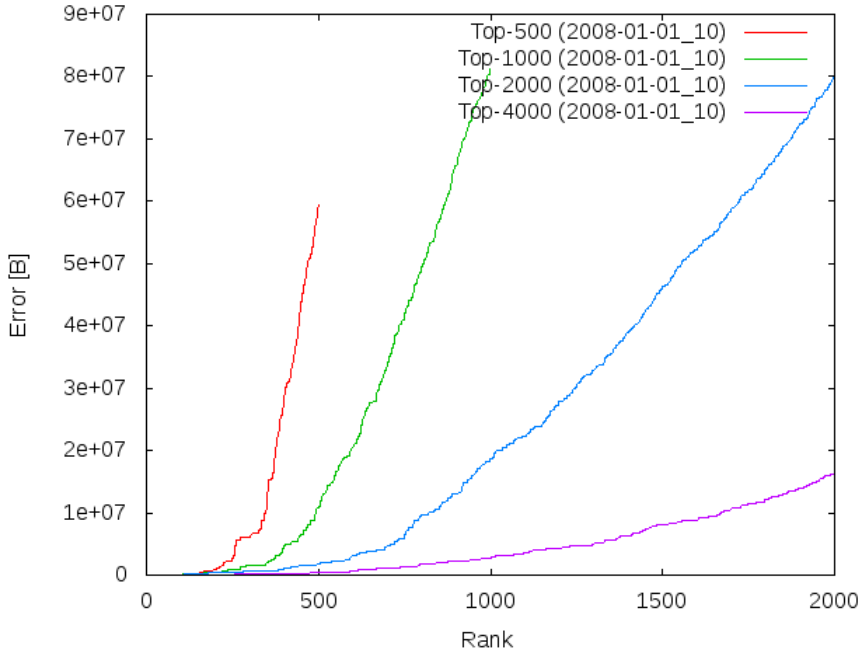
© 2010 IBM Corporation

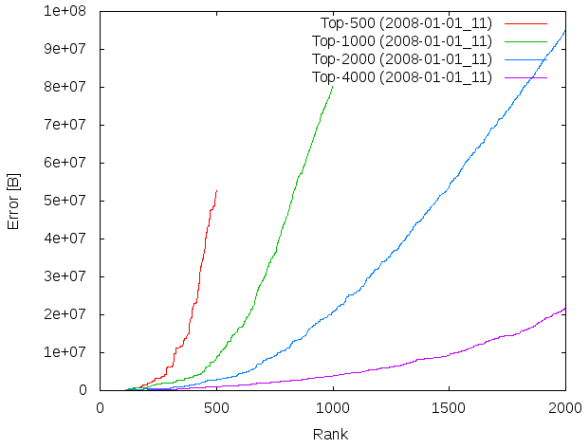
Data Set

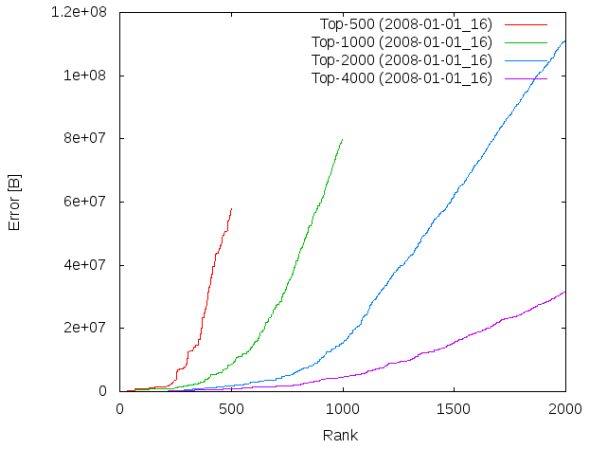
- We picked three days of data collected at one of the IBM datacenters
- Did an expensive query: calculate the exact top K entries
- Fortunately there are machines with large amounts of memory and number of cores…
- Then we applied the defined heuristics to the data set, which run in near real-time as we don't need a lot of memory or cpu power to store all the entries
- One of the biggest advantages of course is that one can almost stay in-cache when one has proper cpu's (4 or 8 MiB cache per core)
- results on next slides….
- Note that the dataset represents several Petabytes of network traffic, the 'error' rate in comparison is thus effectively very low…..

© 2010 IBM Corporation

Error rate for Top K without heuristic



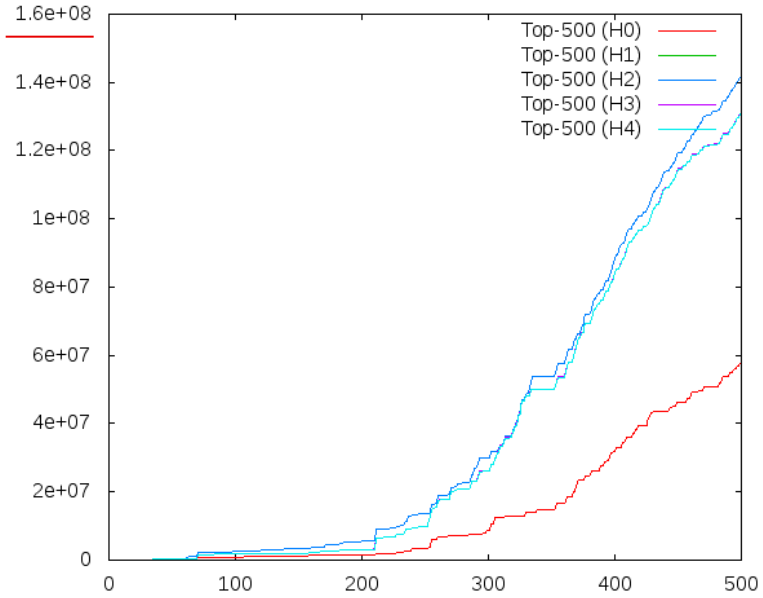


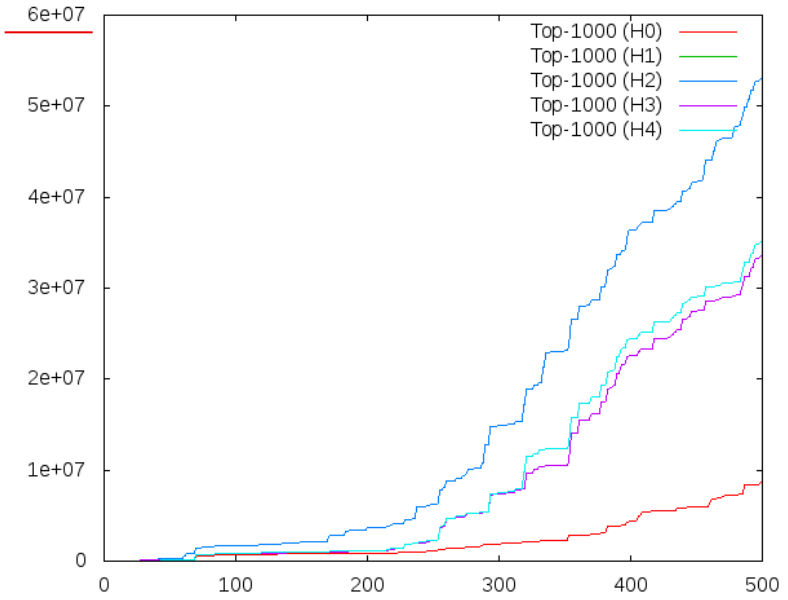


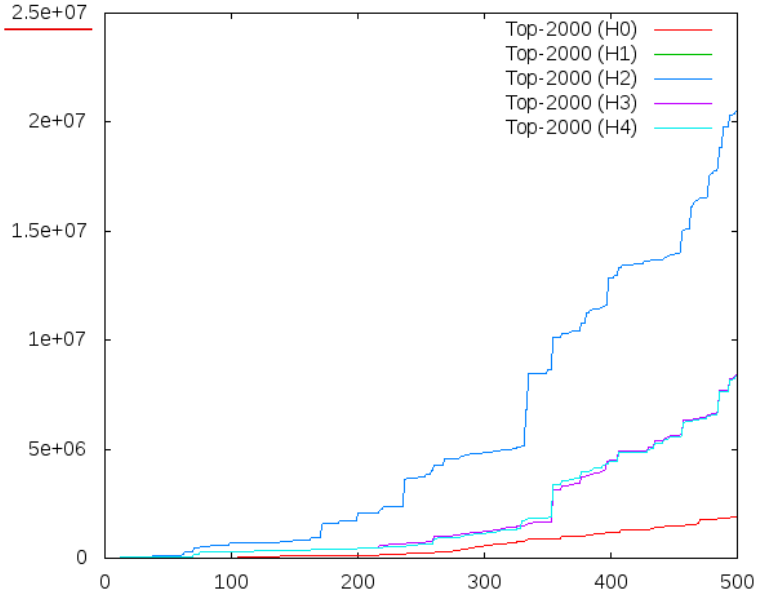
Amount of bytes that are not seen for top X items

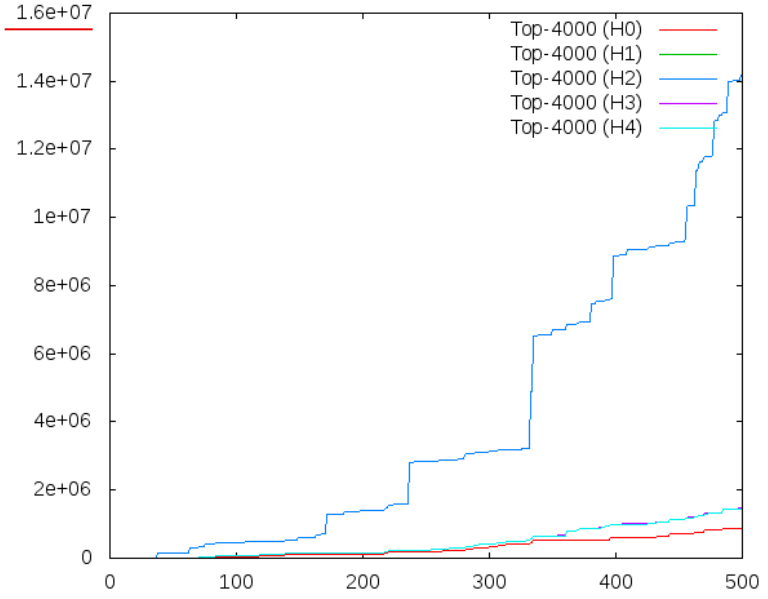
© 2010 IBM Corporation

Error rate for Top K when applying various heuristics









tion

Conclusion

- Simple heuristics and thus cheap processing:
  - Differences are minimal for this ‘datacenter’ dataset, other datasets, eg for an enterprise network seem to prefer the ‘bytes < avg’ heuristic a bit more, which makes sense as it is hard to get in the top K unless you can at least beat the average.

- Two thresholds
  - max2 = 2 * max1
  - Correct order up to k = X * max1

© 2010 IBM Corporation

AURORA

http://www.zurich.ibm.com/aurora/



© 2010 IBM Corporation

Towards Reliable Traffic Classification Using Visual Motifs

Wilson Lian[1] John McHugh[1,2] Fabian Monrose[1]

[1]University of North Carolina at Chapel Hill

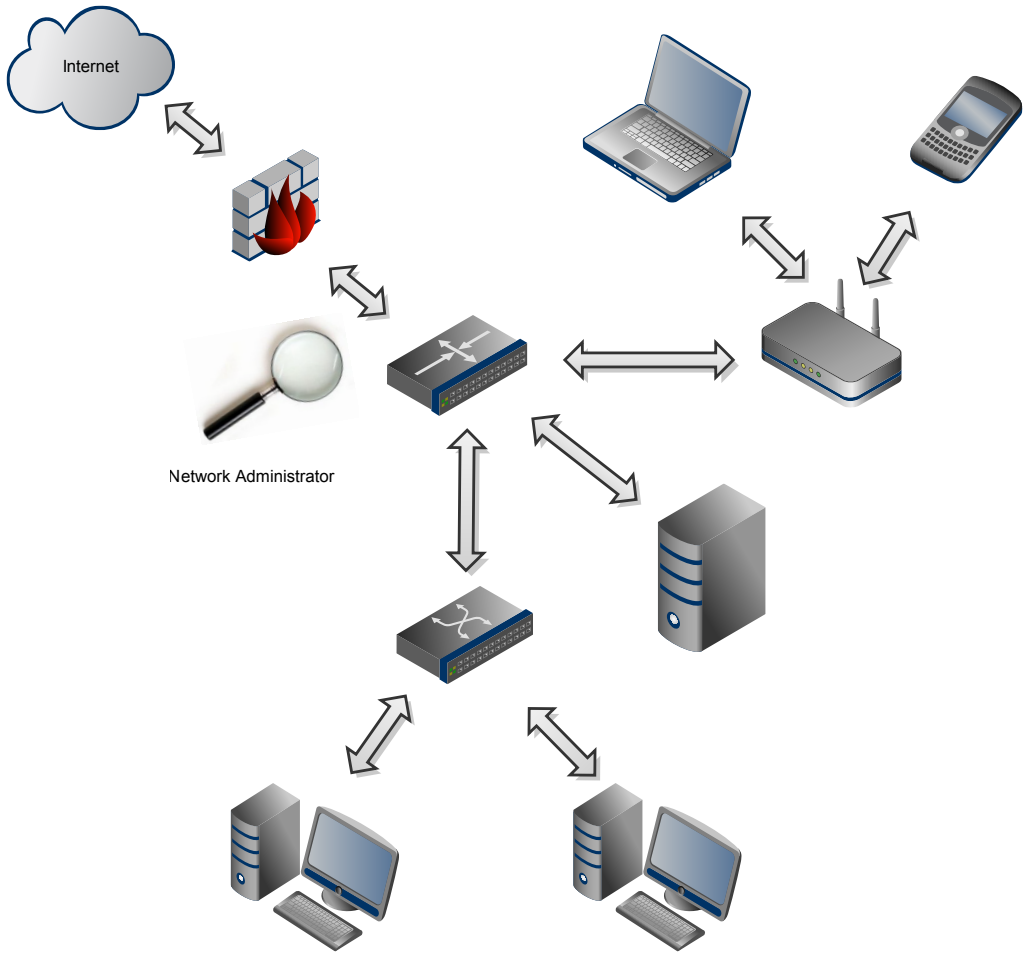
[2]RedJack, LLC

FloCon 2010

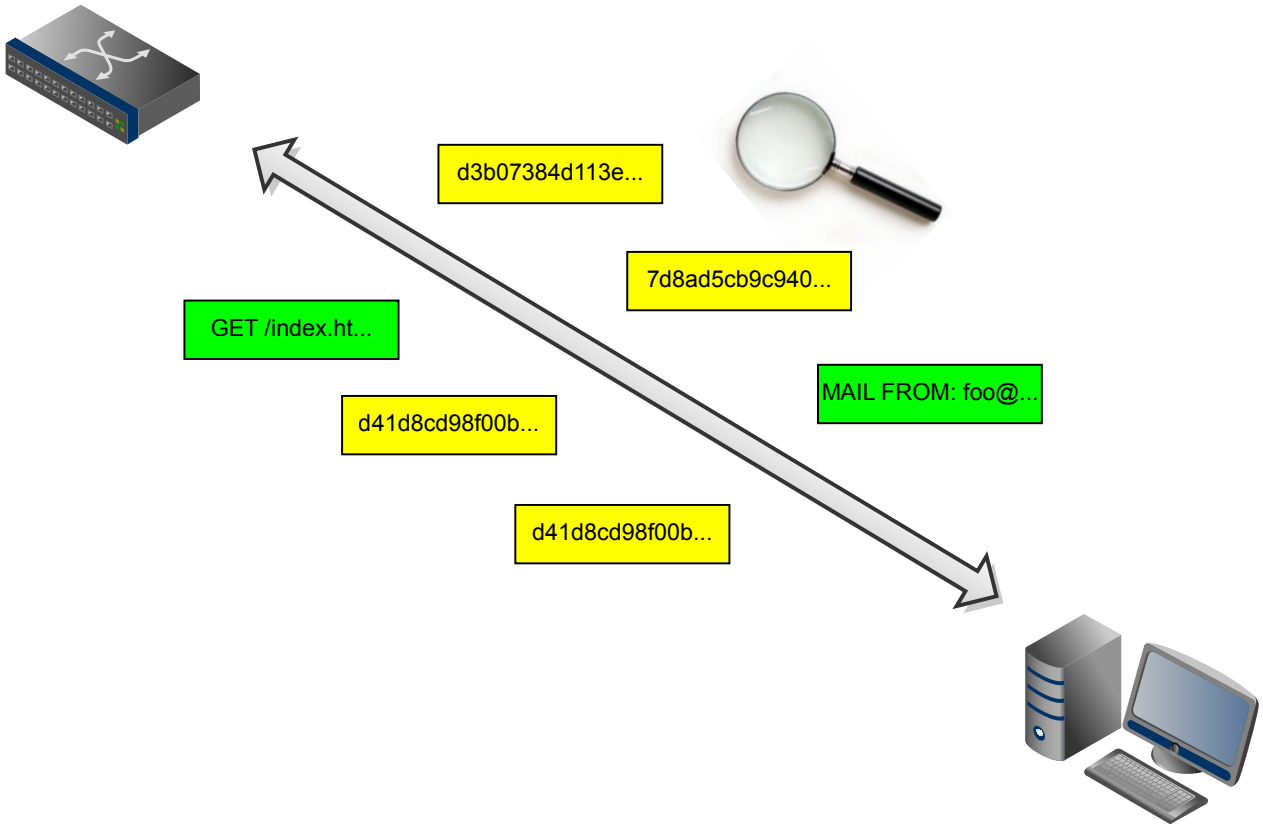
Overview

- Background
- Visual Motifs
- Traffic Classification
- Evaluation

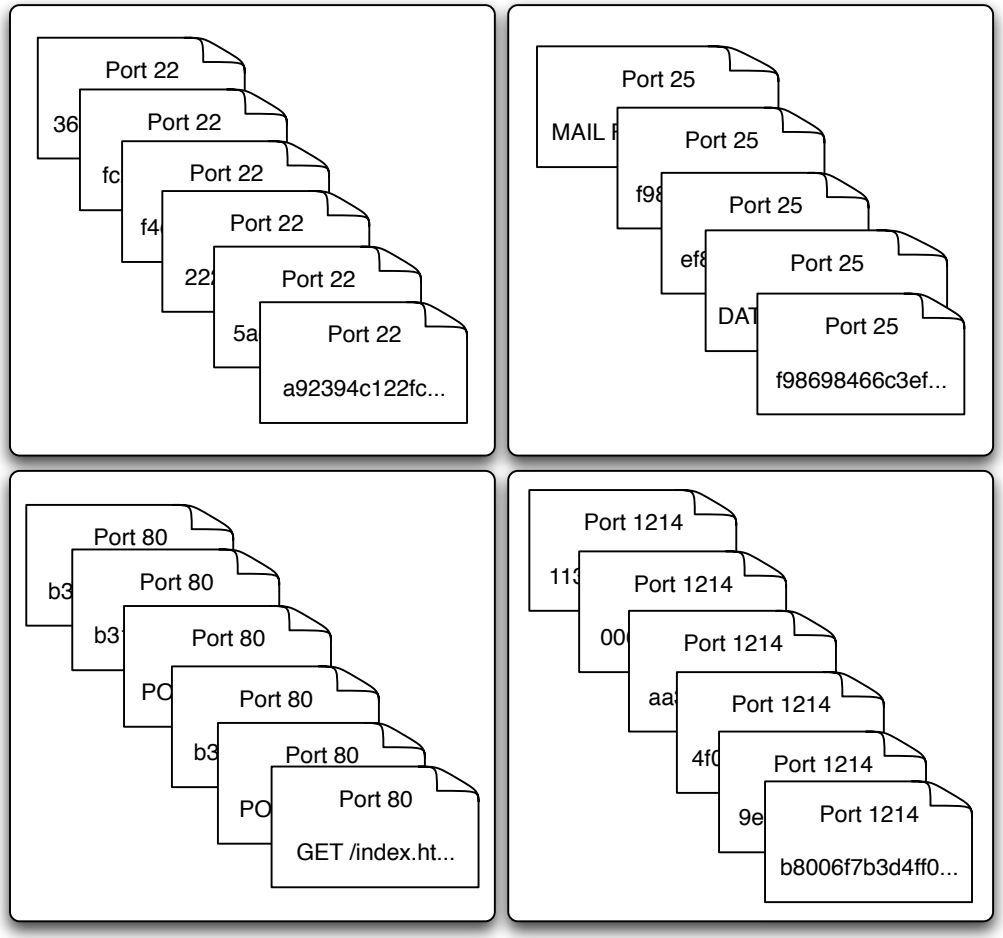
Motivation



Motivation


d3b07384d113e...
7d8ad5cb9c940...
GET /index.ht...
MAIL FROM: foo@...
d41d8cd98f00b...
d41d8cd98f00b...

Goals


Port 22
Port 22
Port 22
Port 22
Port 22
Port 22
a92394c122fc...
Port 25
Port 25
Port 25
Port 25
Port 25
f98698466c3ef...
Port 80
Port 80
Port 80
Port 80
Port 80
Port 80
GET /index.ht...
Port 1214
Port 1214
Port 1214
Port 1214
Port 1214
Port 1214
b8006f7b3d4ff0...

Assumptions

- Reliable transport via TCP
- Stream Cipher
  - No access to payload
  - Length preservation
- Negligible packet loss & retransmission

Related Work

- *Scatter (and other) Plots for Visualizing User Profiling Data and Network Traffic*, Goldring 2004.
- *Using Visual Motifs to Classify Encrypted Traffic*, Wright *et al.* 2006
- *Intelligent Classification and Visualization of Network Scans* Muelder *et al.* 2008.
- *FloVis: A Network Security Visualization Framework*, Taylor 2009.

Timeline Heatmaps

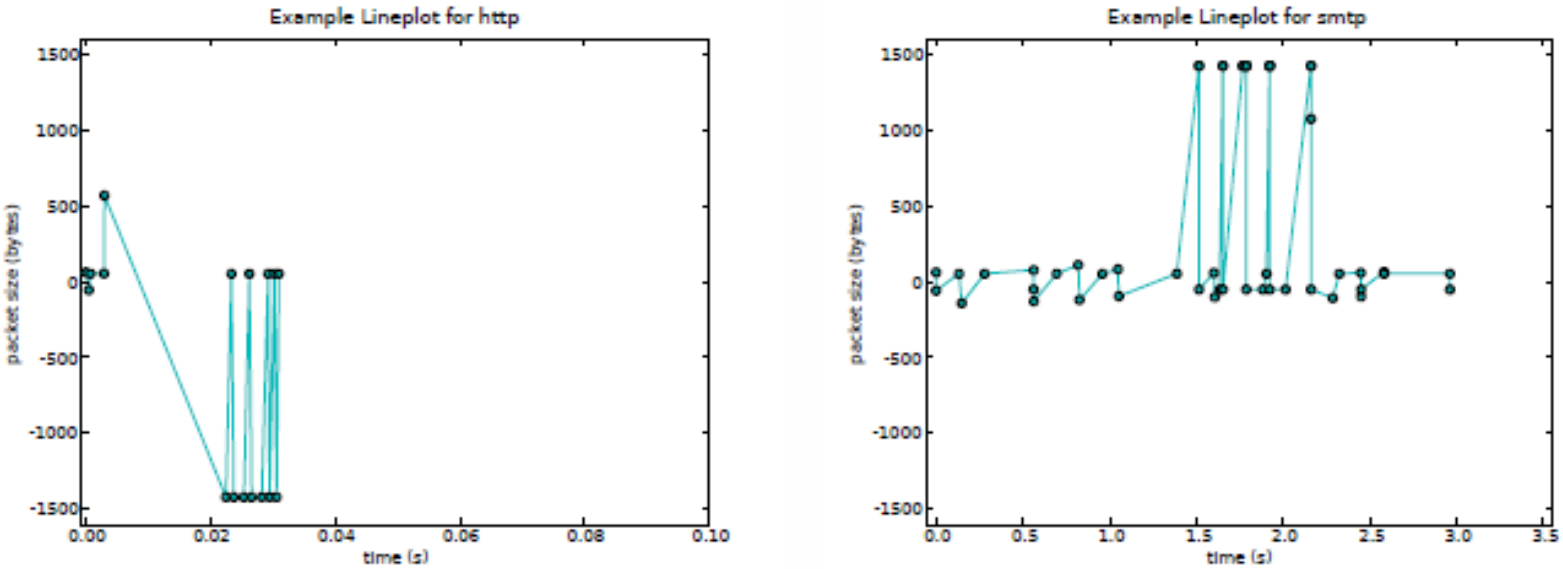


Image credit: Wright *et al.* 2006

Timeline Heatmaps

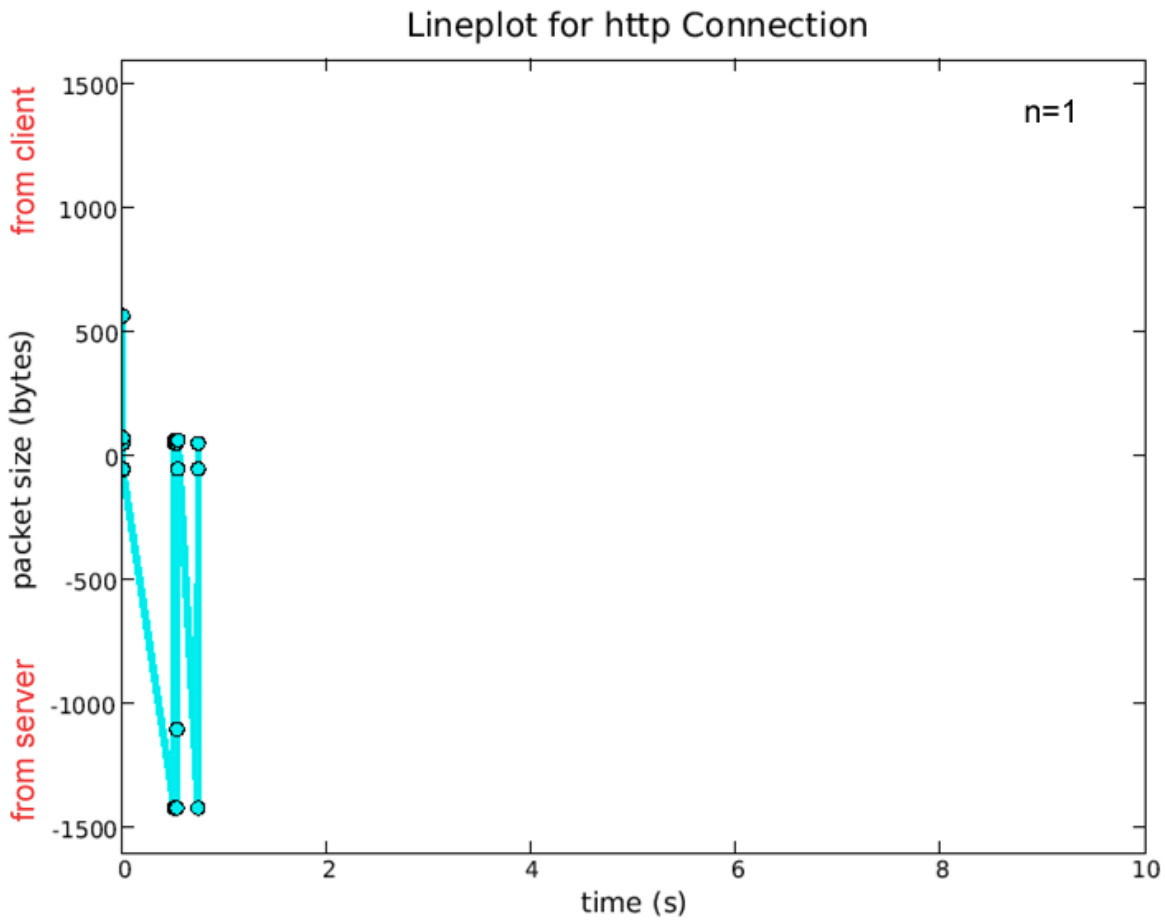


Image credit: Wright *et al.* 2006

Timeline Heatmaps

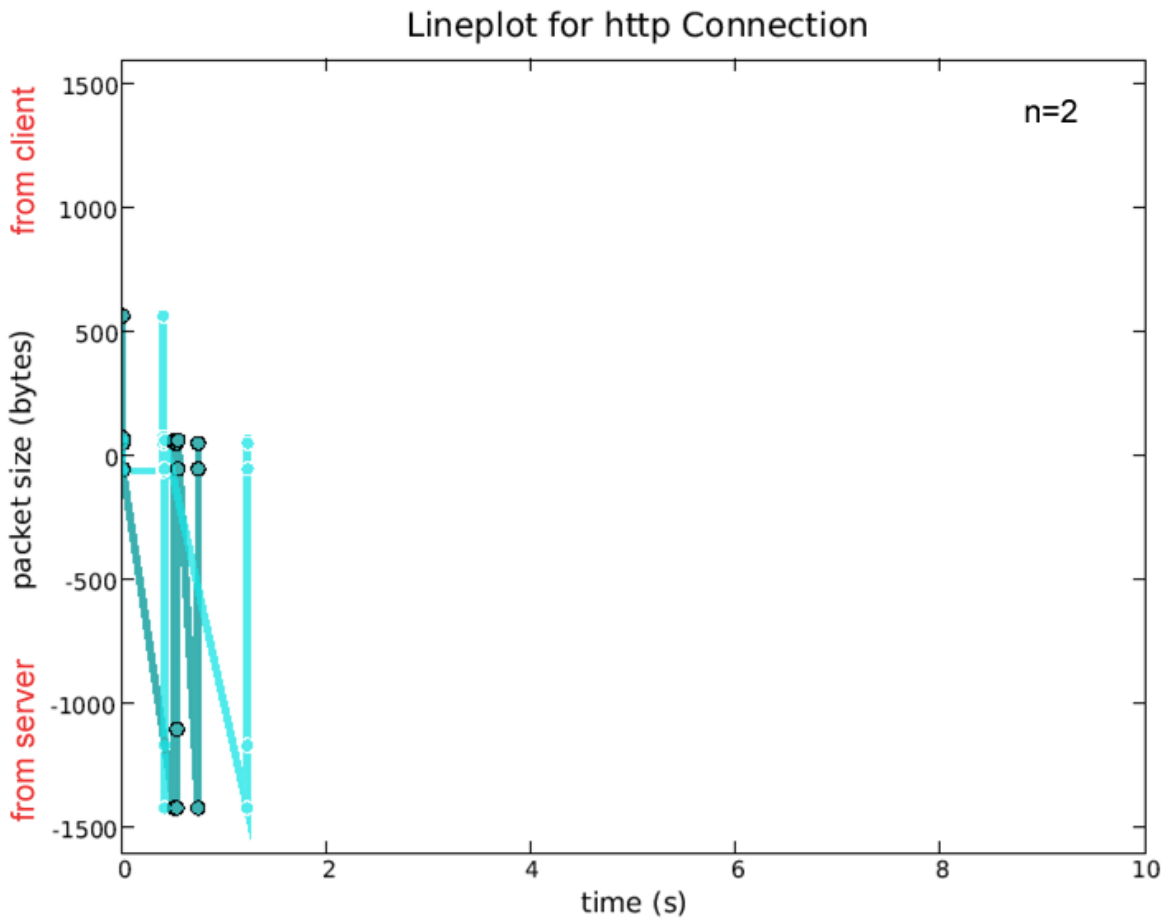


Image credit: Wright *et al.* 2006

Timeline Heatmaps

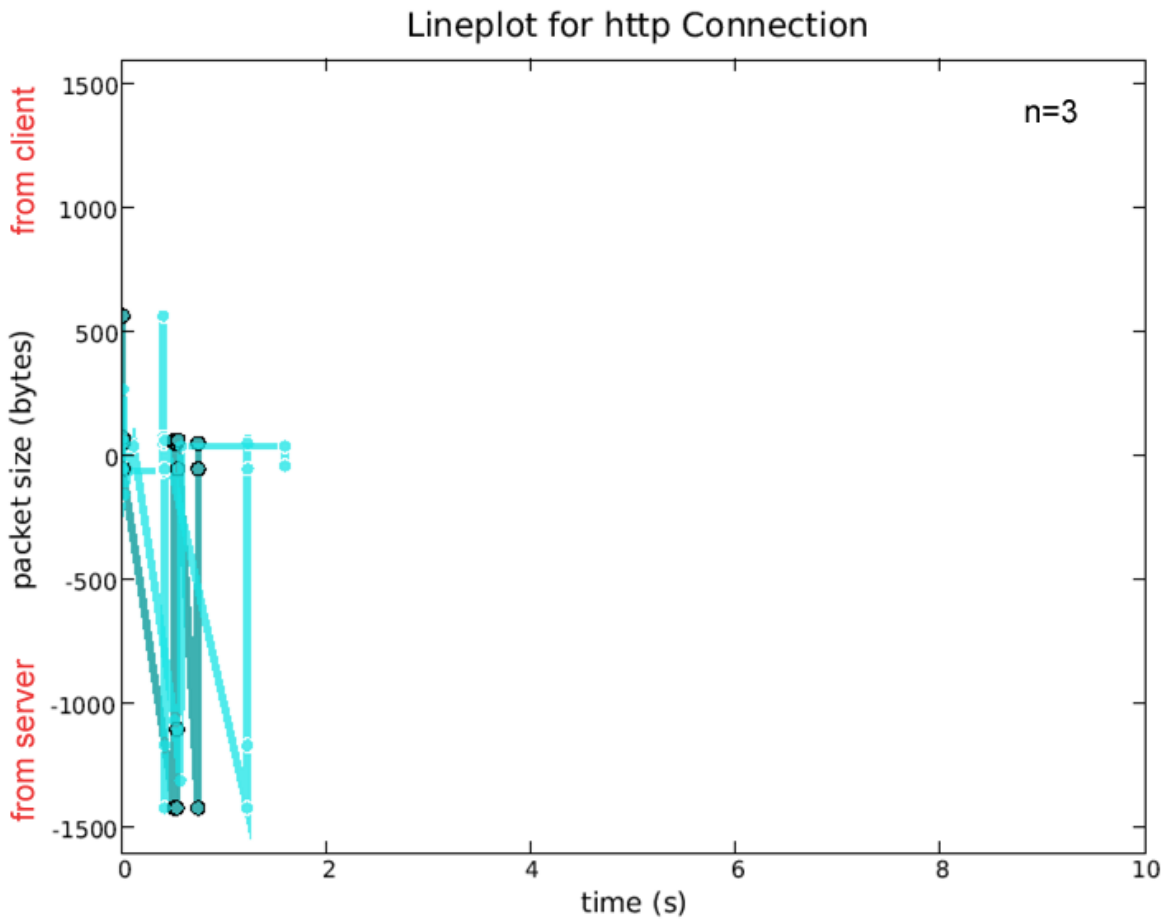


Image credit: Wright *et al.* 2006

Timeline Heatmaps

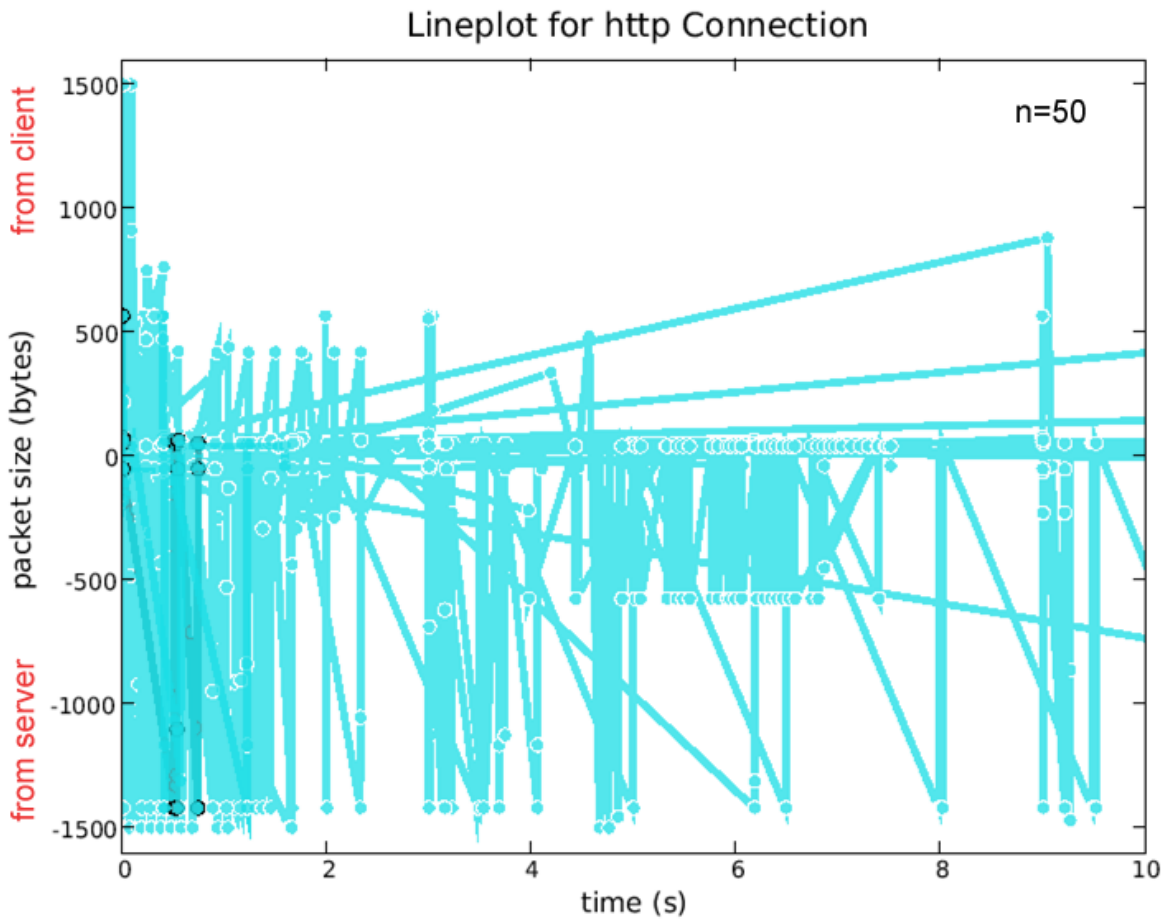


Image credit: Wright *et al.* 2006

Timeline Heatmaps

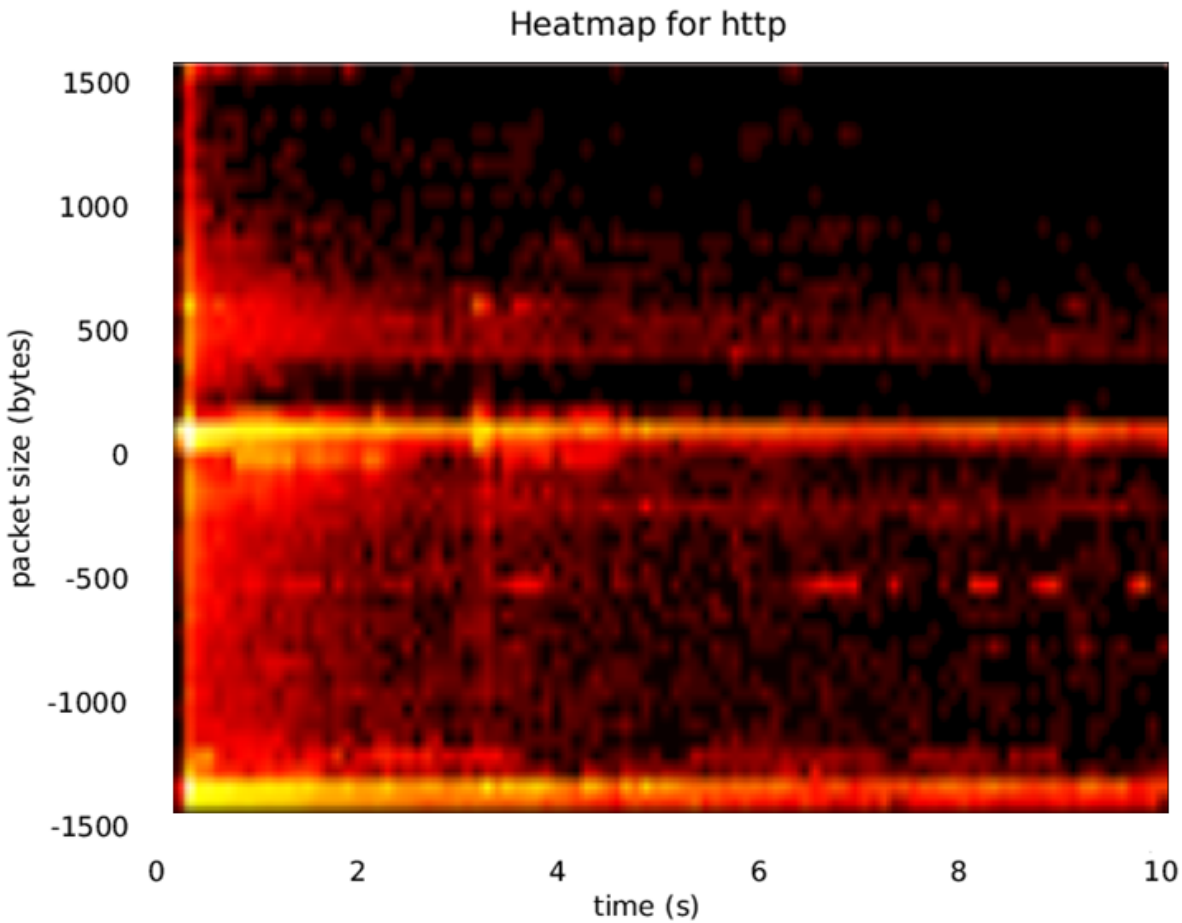
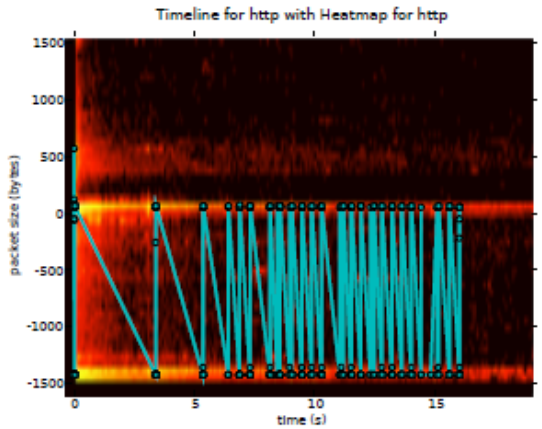
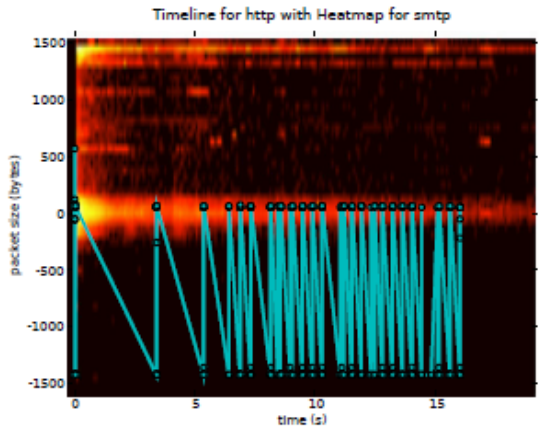


Image credit: Wright *et al.* 2006

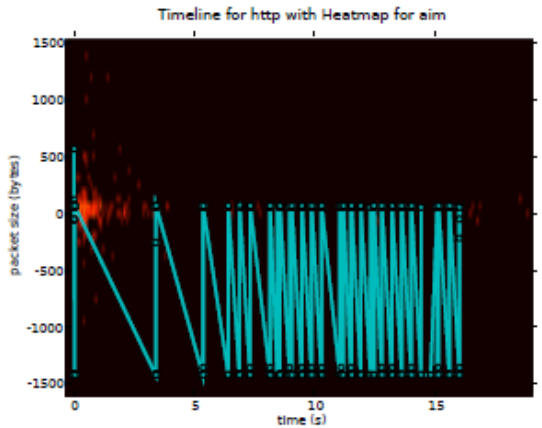
Timeline Heatmaps



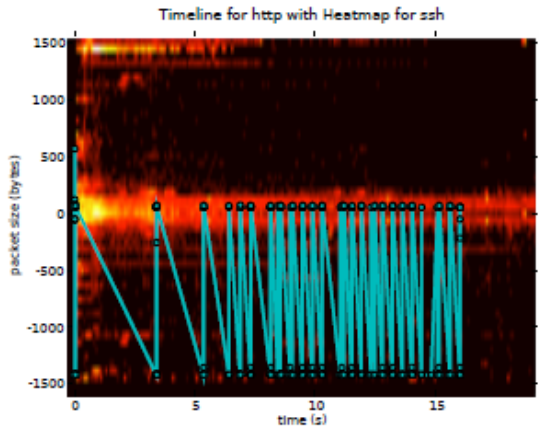
(a) HTTP



(b) SMTP



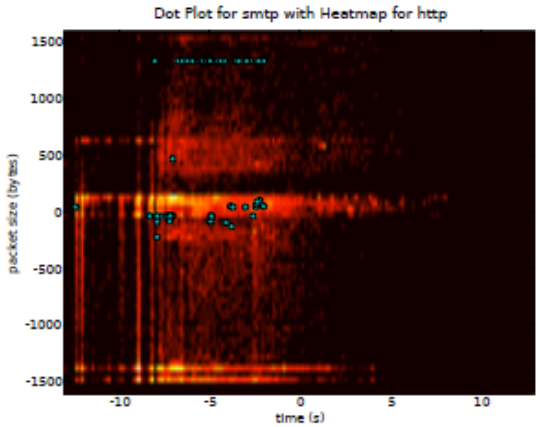
(c) AIM



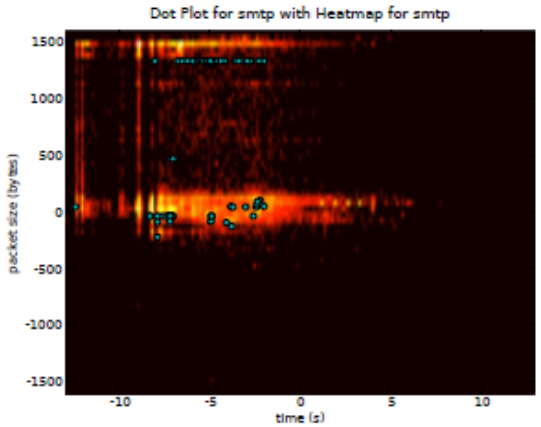
(d) SSH

Image credit: Wright *et al.* 2006

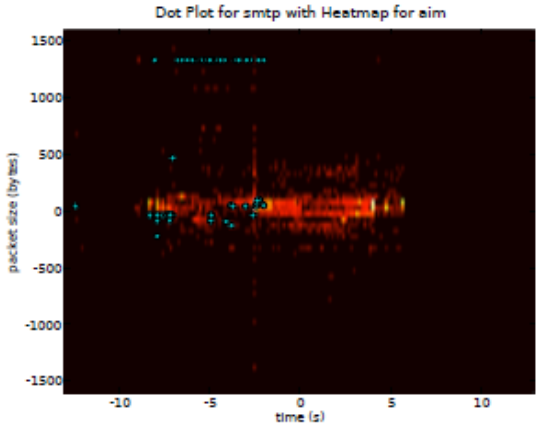
Unigram Heatmaps



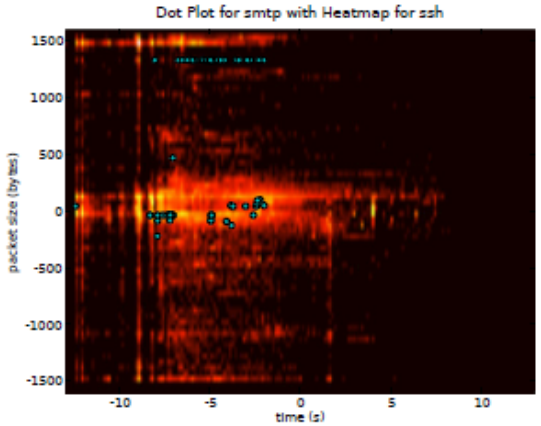
(a) HTTP



(b) SMTP



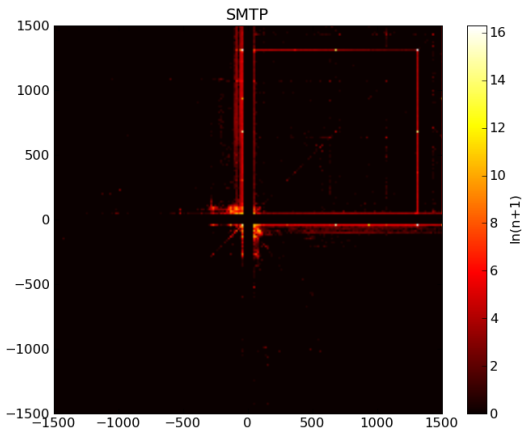
(c) AIM

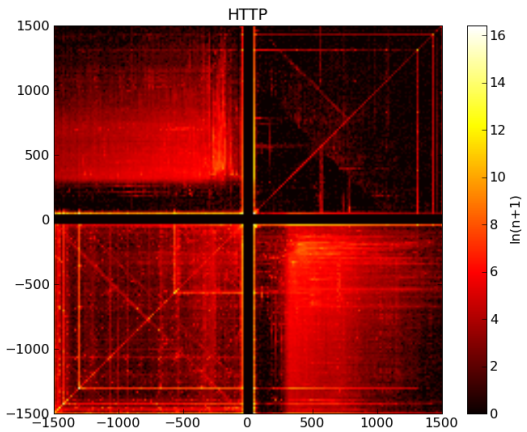


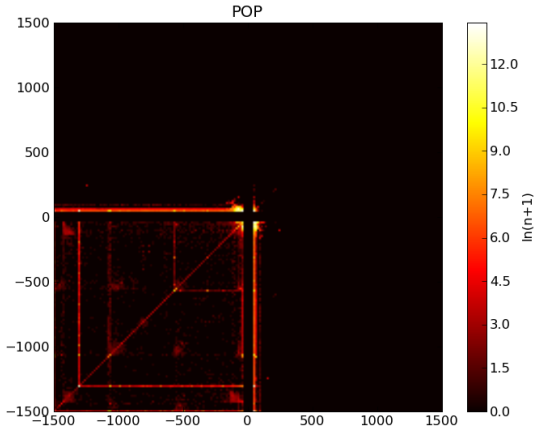
(d) SSH

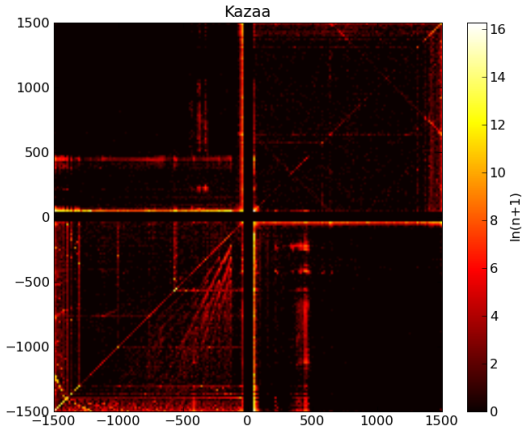
Image credit: Wright *et al.* 2006

Bigram Heatmaps

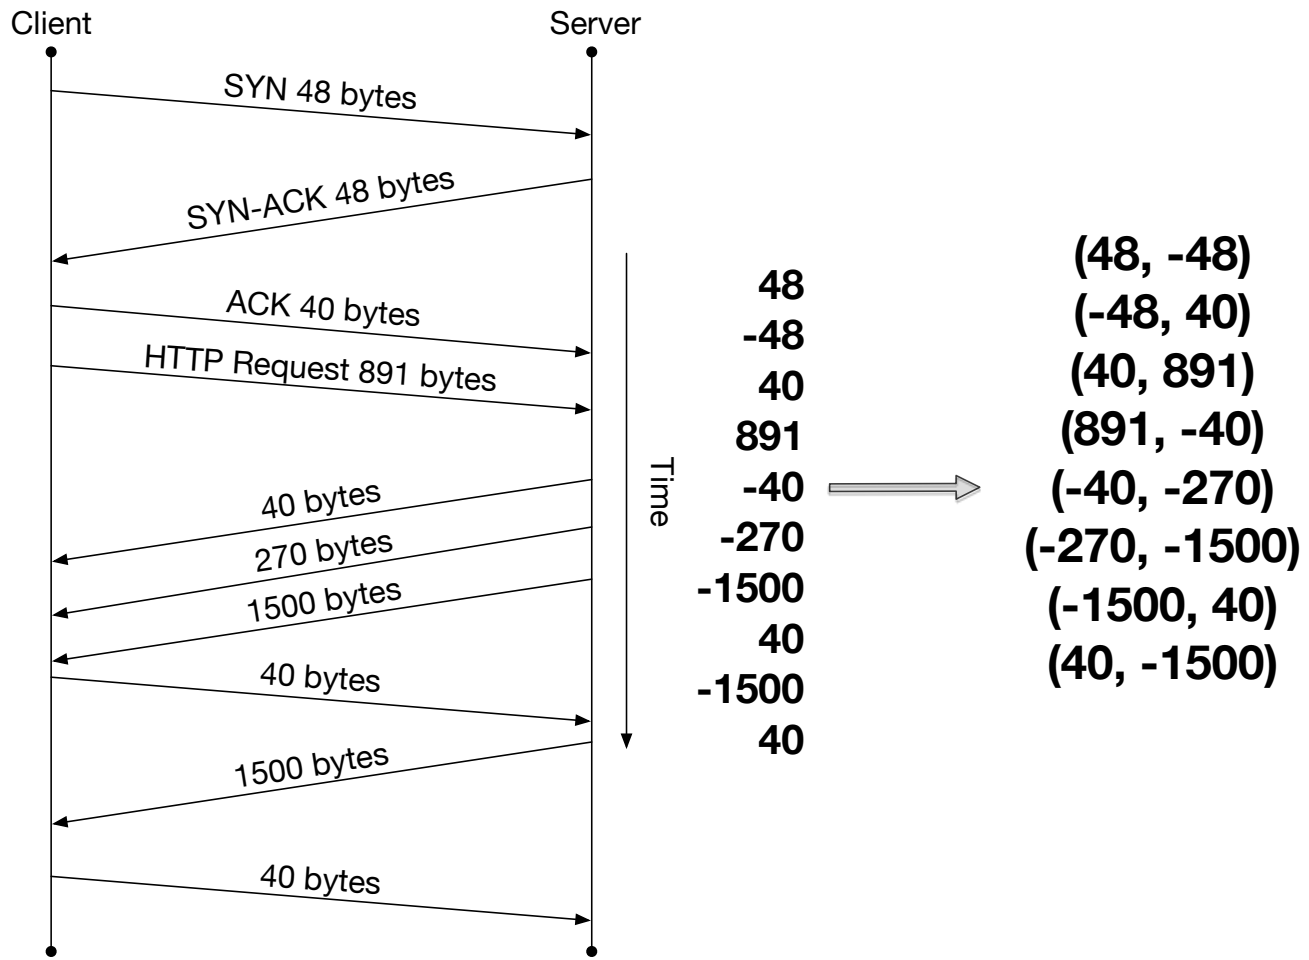








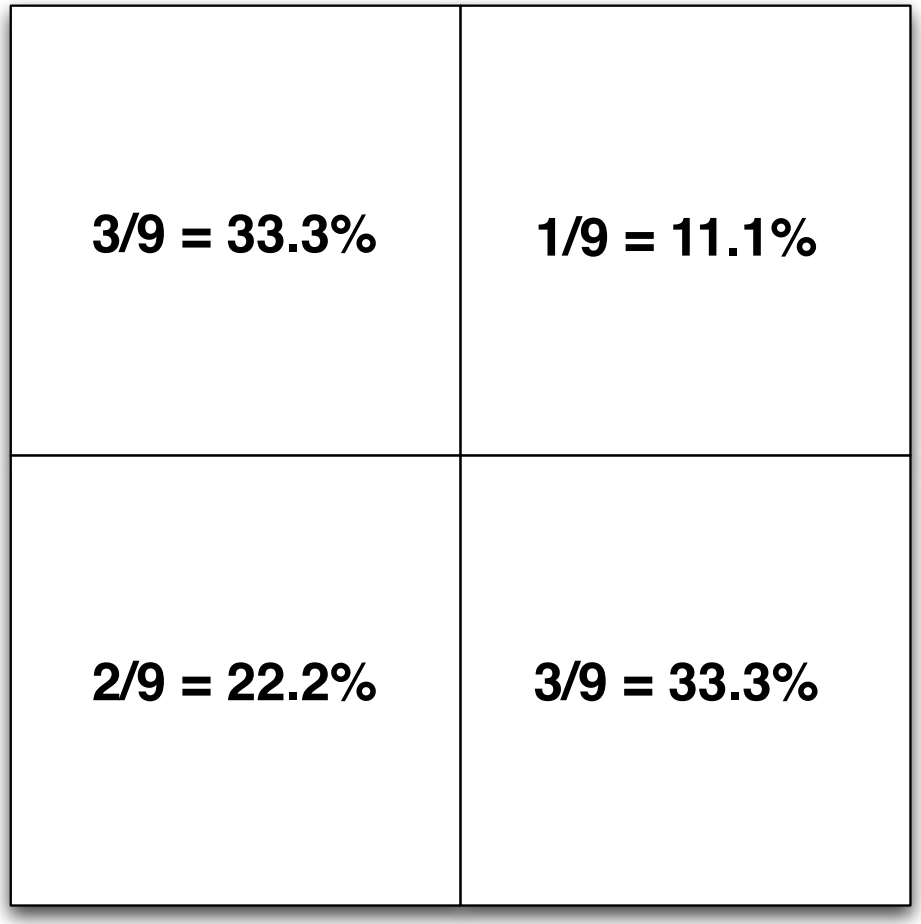
Heatmap Construction



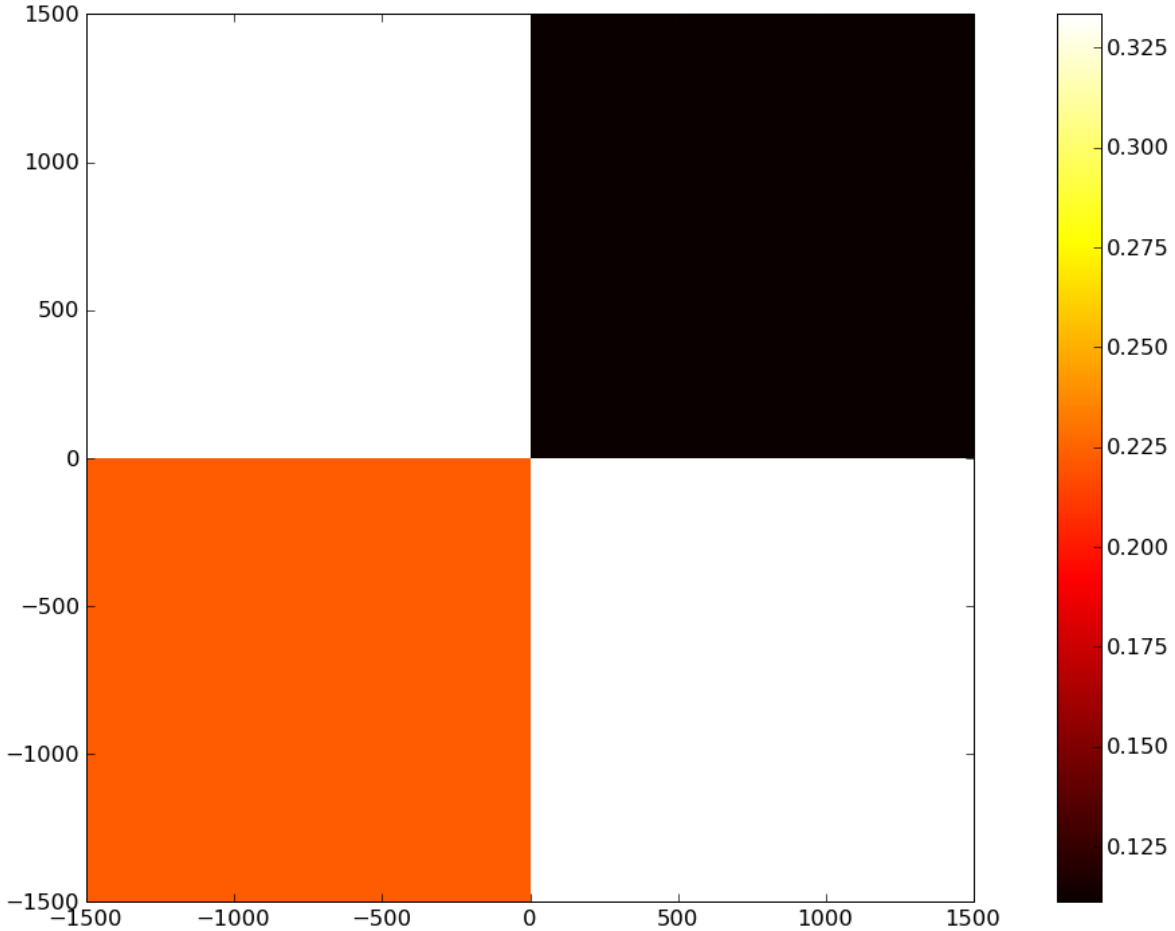
Heatmap Construction

(-48, 40) **(-1500, 40)** **(-1500, 40)**	**(40, 891)**
(-40, -270) **(-270, -1500)**	**(48, -48)** **(891, -40)** **(40, -1500)**

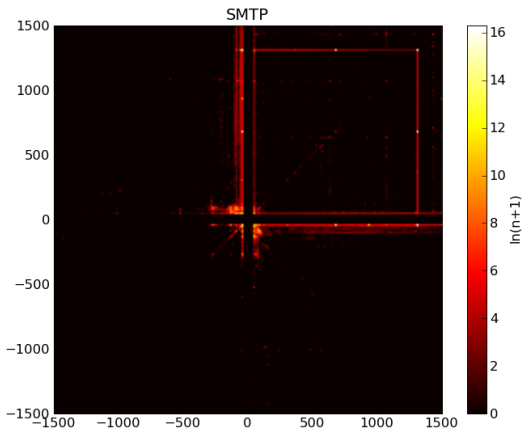
Heatmap Construction

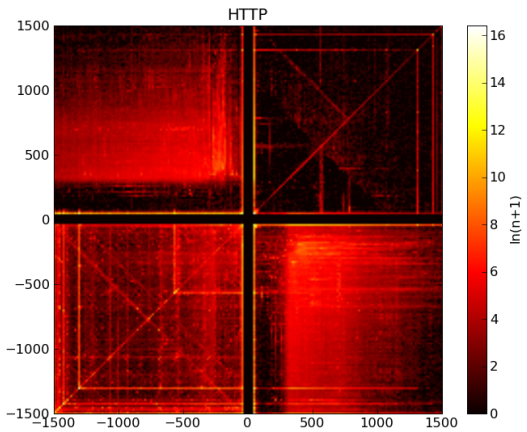


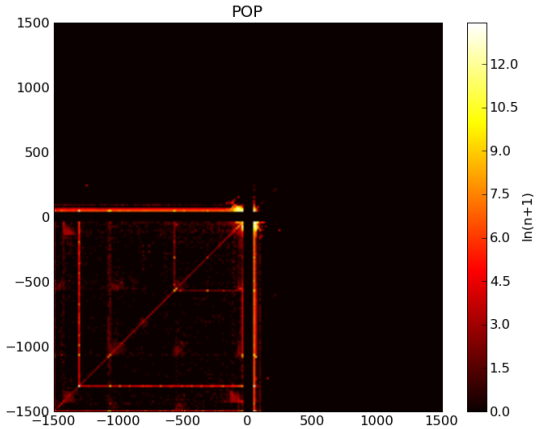
Heatmap Construction

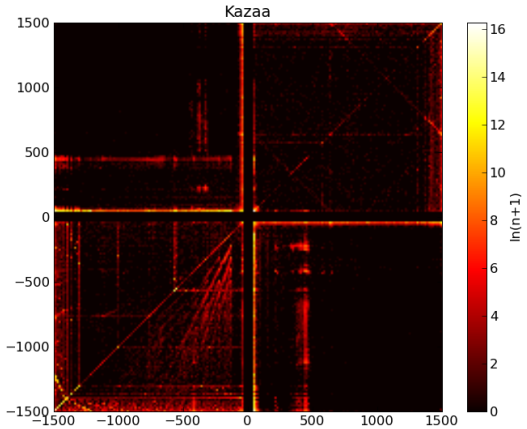


Bigram Heatmaps





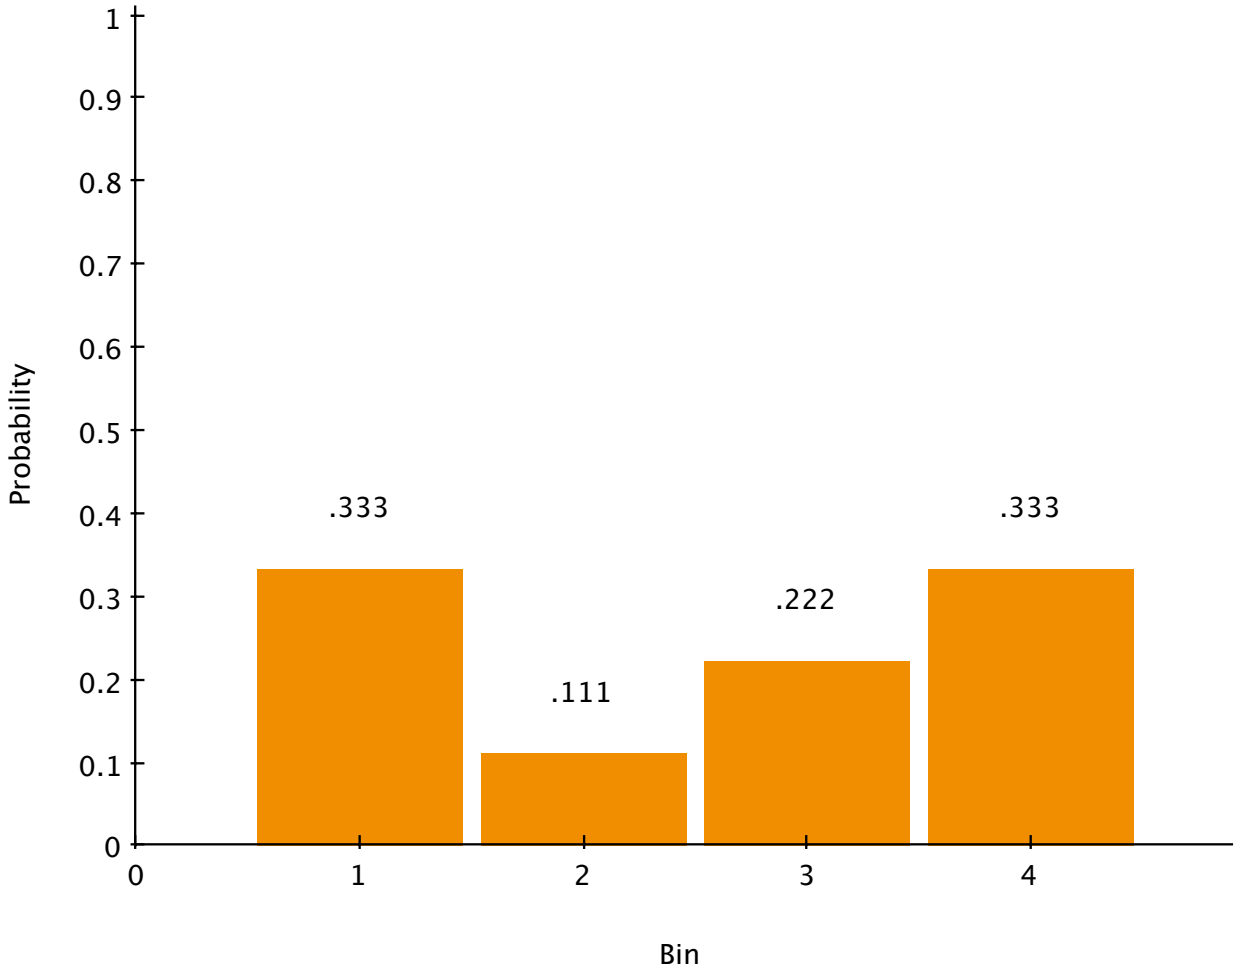




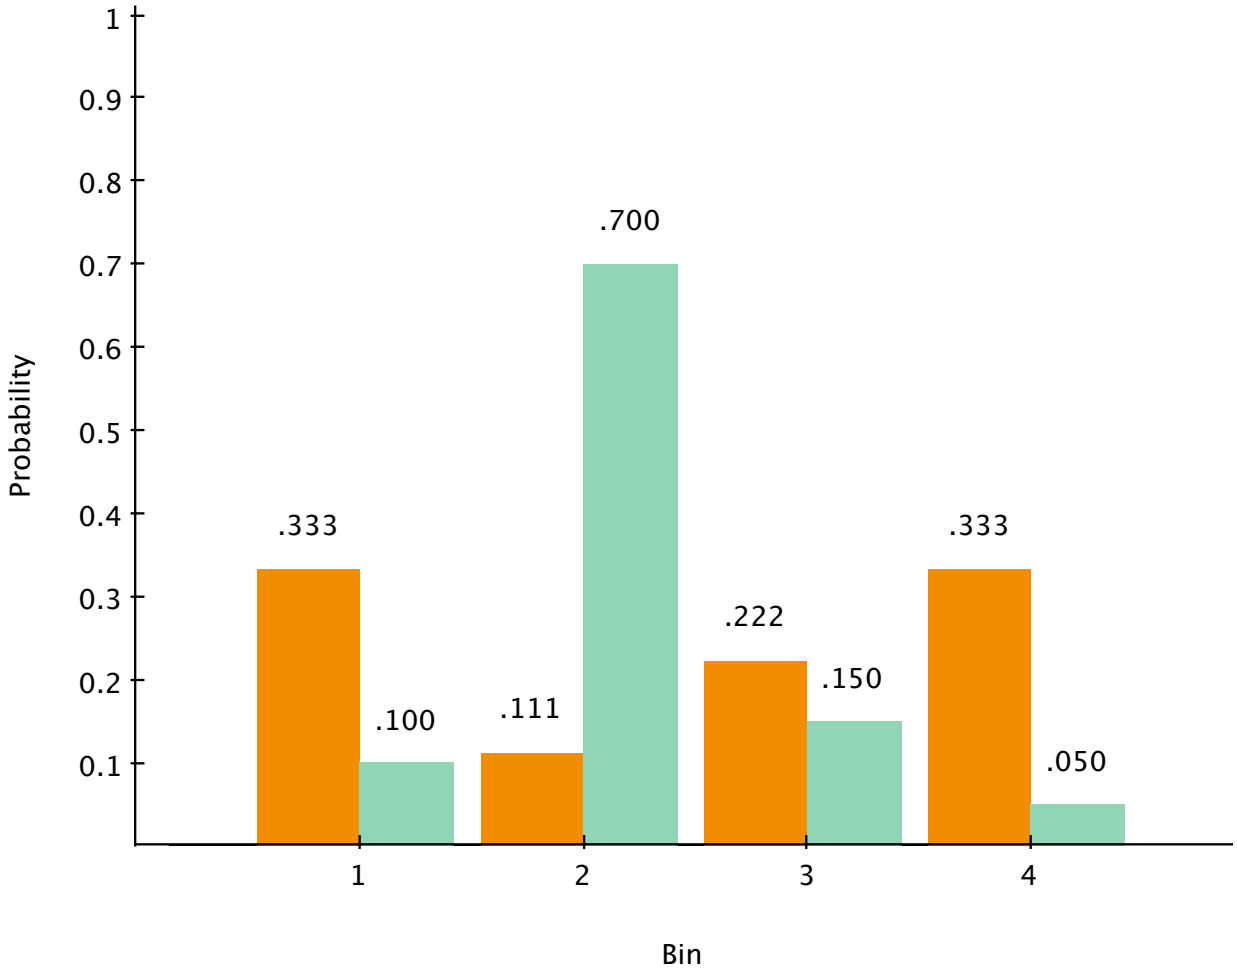
Modeling Protocol Behavior

3/9 = 33.3% **1**	**1/9 = 11.1%** **2**
2/9 = 22.2% **3**	**3/9 = 33%** **4**

Modeling Protocol Behavior



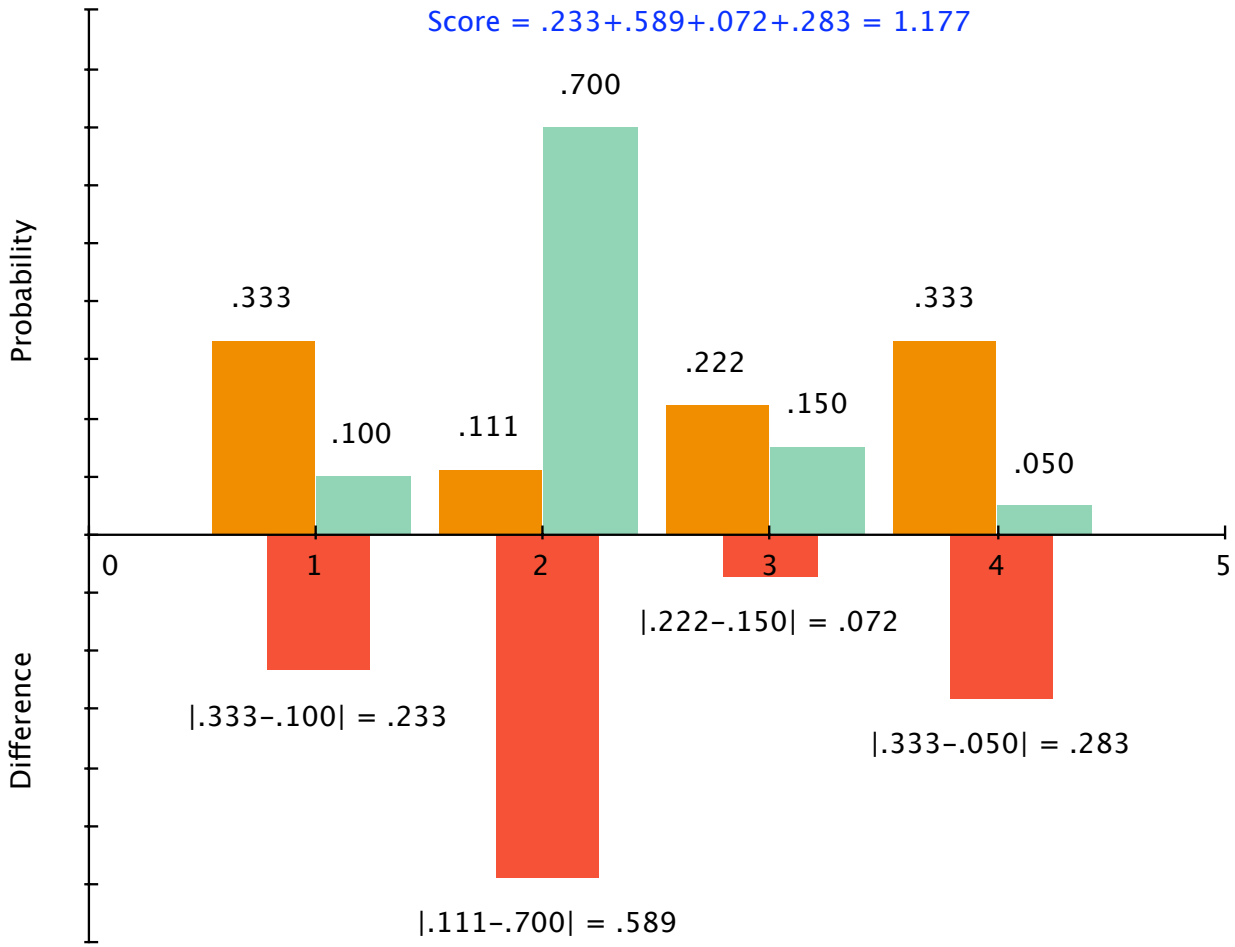
Comparing Protocol Models



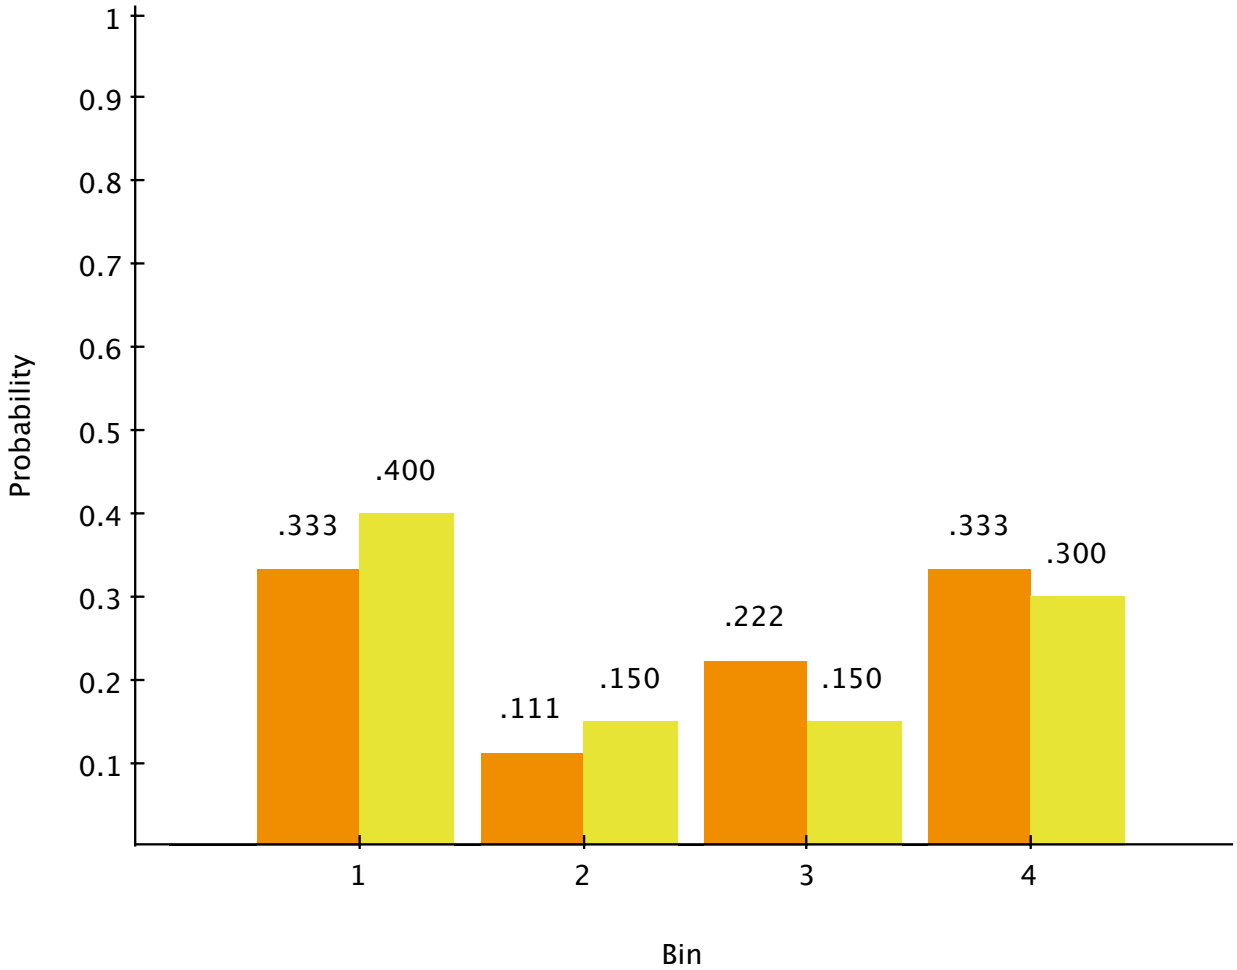
Comparing Protocol Models

$$
\begin{aligned}
A_{total} &= \sum_{k=1}^{n} A_k \\
B_{total} &= \sum_{k=1}^{n} B_k \\
Score_{A \leftrightarrow B} &= \sum_{i=1}^{n} \left| \frac{A_i}{A_{total}} - \frac{B_i}{B_{total}} \right| \\
&= \frac{1}{A_{total} \cdot B_{total}} \sum_{i=1}^{n} |A_i \cdot B_{total} - B_i \cdot A_{total}|
\end{aligned}
$$

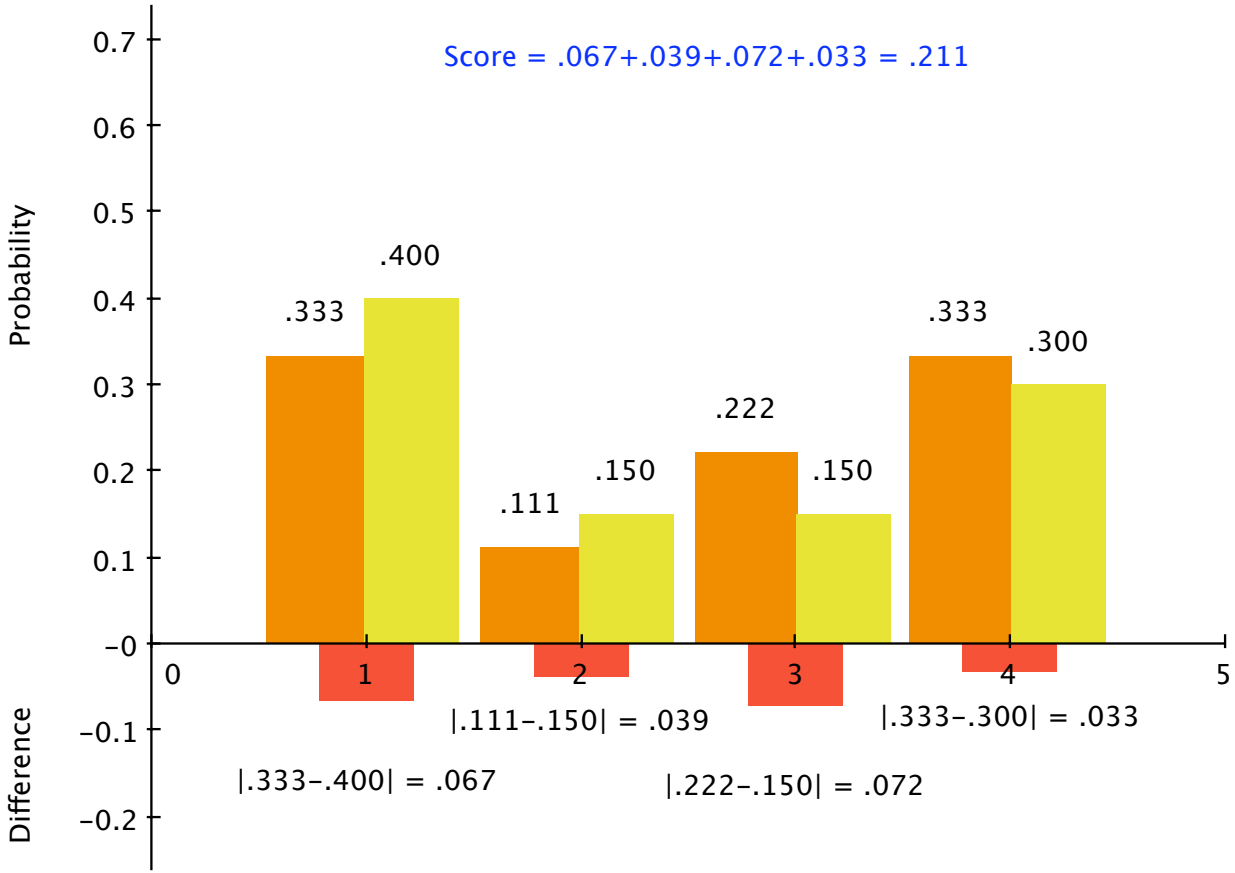
Comparing Protocol Models



Comparing Protocol Models



Comparing Protocol Models



Classifying Samples: Easy as 1-2-3

1. Create training models for desired protocols
2. Build distribution for sample network trace
3. Find training model with lowest difference score

$$
\begin{aligned}
Score_{A \leftrightarrow B} &= \sum_{i=1}^{n} \left| \frac{A_i}{A_{total}} - \frac{B_i}{B_{total}} \right| \\
&= \frac{1}{A_{total} \cdot B_{total}} \sum_{i=1}^{n} |A_i \cdot B_{total} - B_i \cdot A_{total}|
\end{aligned}
$$

Evaluation

- How much traffic must be collected for:
  - Training
  - Testing
- Precision?

$$\frac{true\ positives}{true\ positives + false\ positives}$$

- Recall?

$$\frac{true\ positives}{true\ positives + false\ negatives}$$

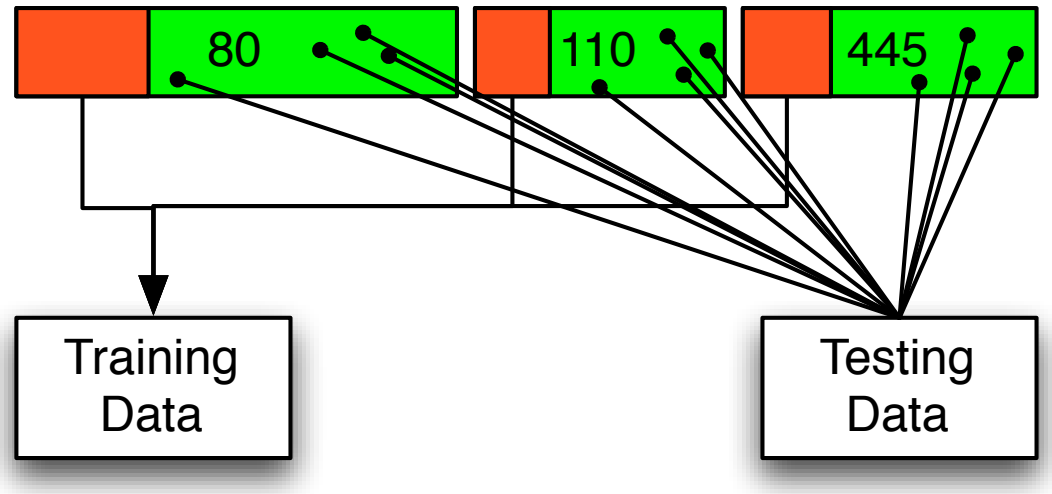
Data

- CRAWDAD Dataset
- Weekdays: January 19, 2004 – February 6, 2004
- Ports with sufficient traffic
  - $\geq$ 1M packets
  - 0.3% of ports $\rightarrow$ 95.21% of packets
- Keep top 10 ports by number of sessions observed
- No ground truth

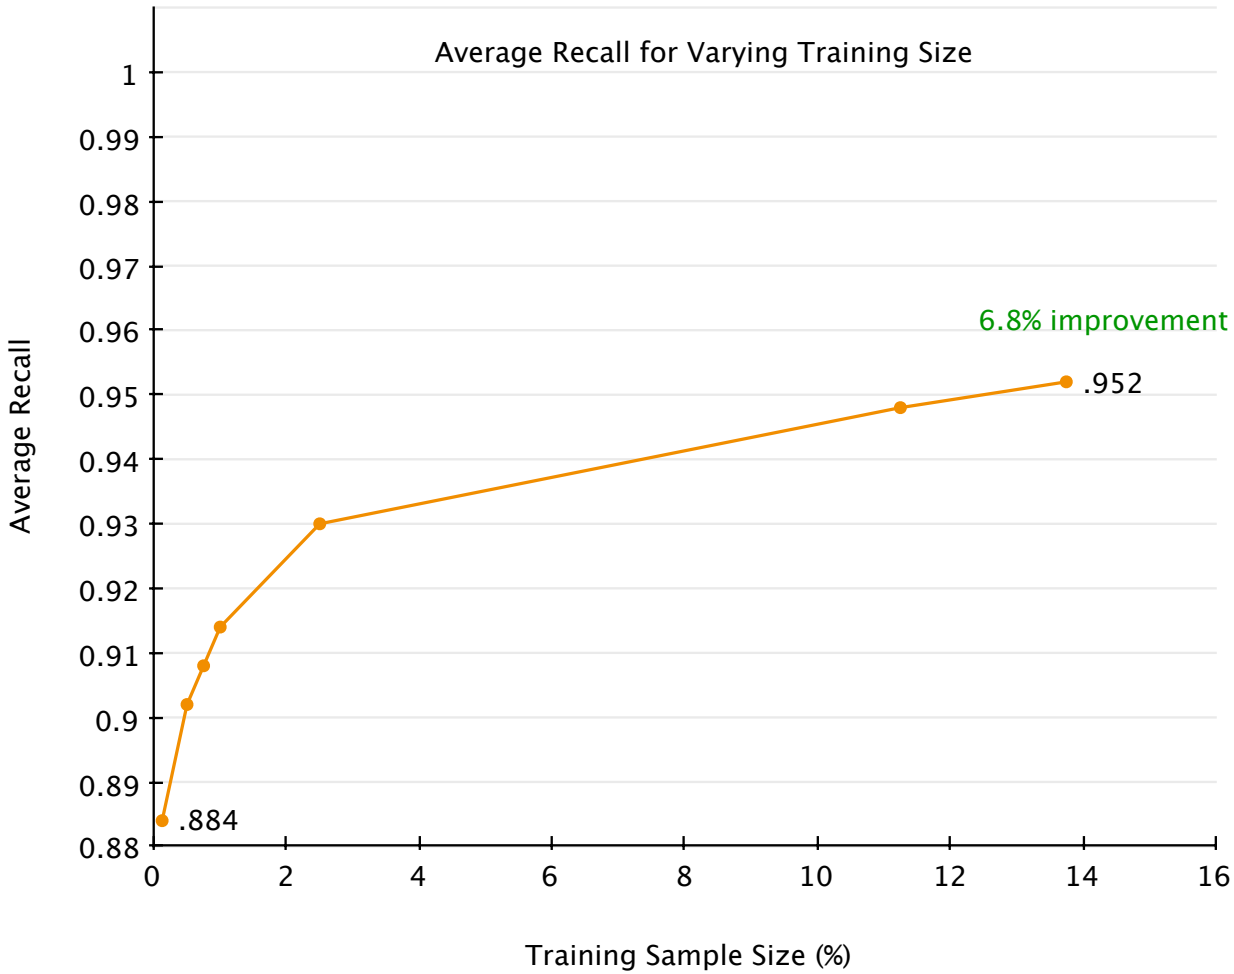
Total Packets	1.3 Billion
Traffic Volume	707 GB
Observed Ports	64,214
Sessions	5.2 Million
Port 80 Sessions	1.7 Million

Methodology

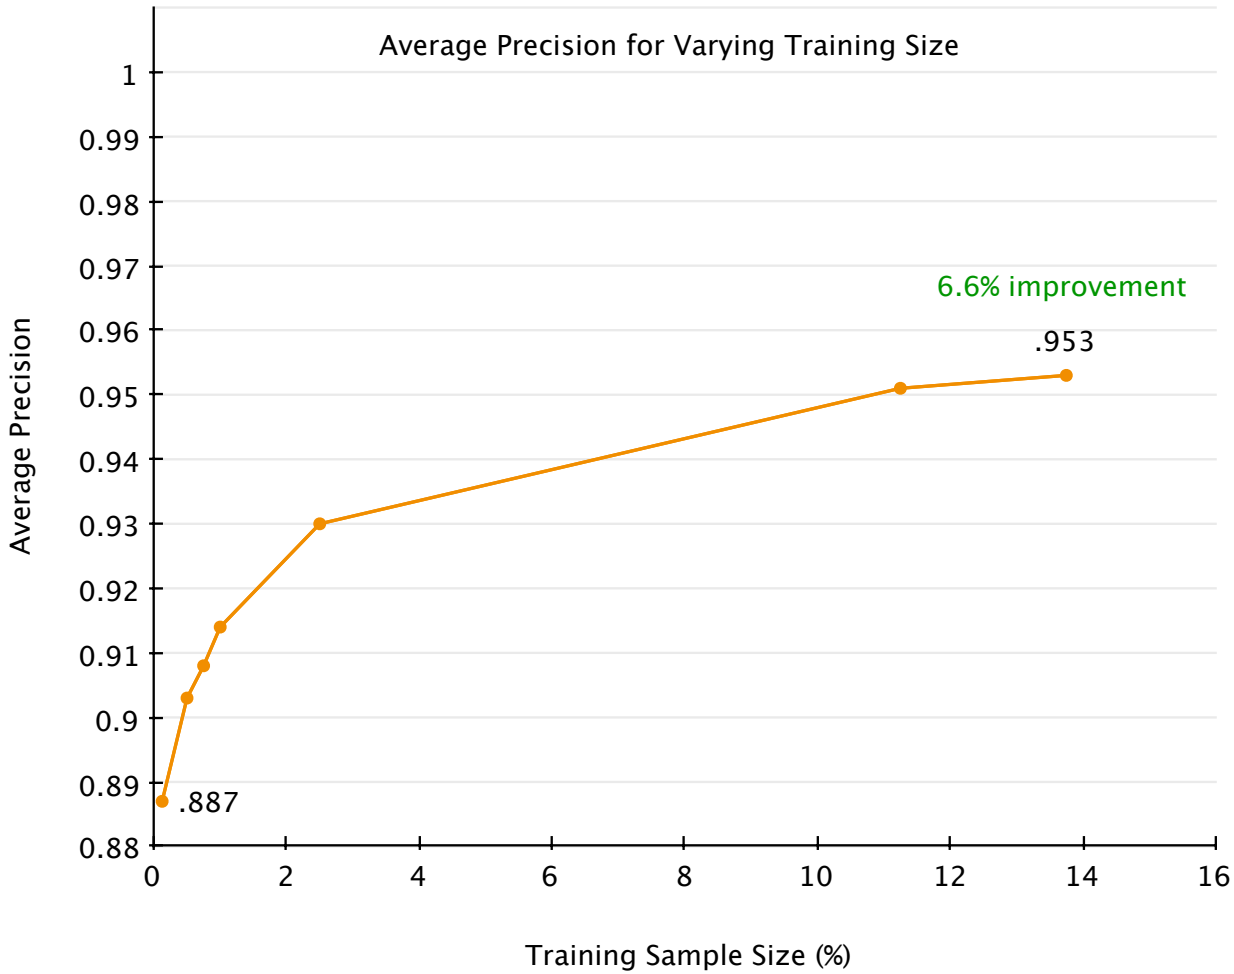
- Trial :=
  1. Randomly sample some percentage of available data for each port and train classifier
  2. Randomly sample some number of the remaining data points for each port and create testing samples
  3. Classify testing samples
- 50 Trials



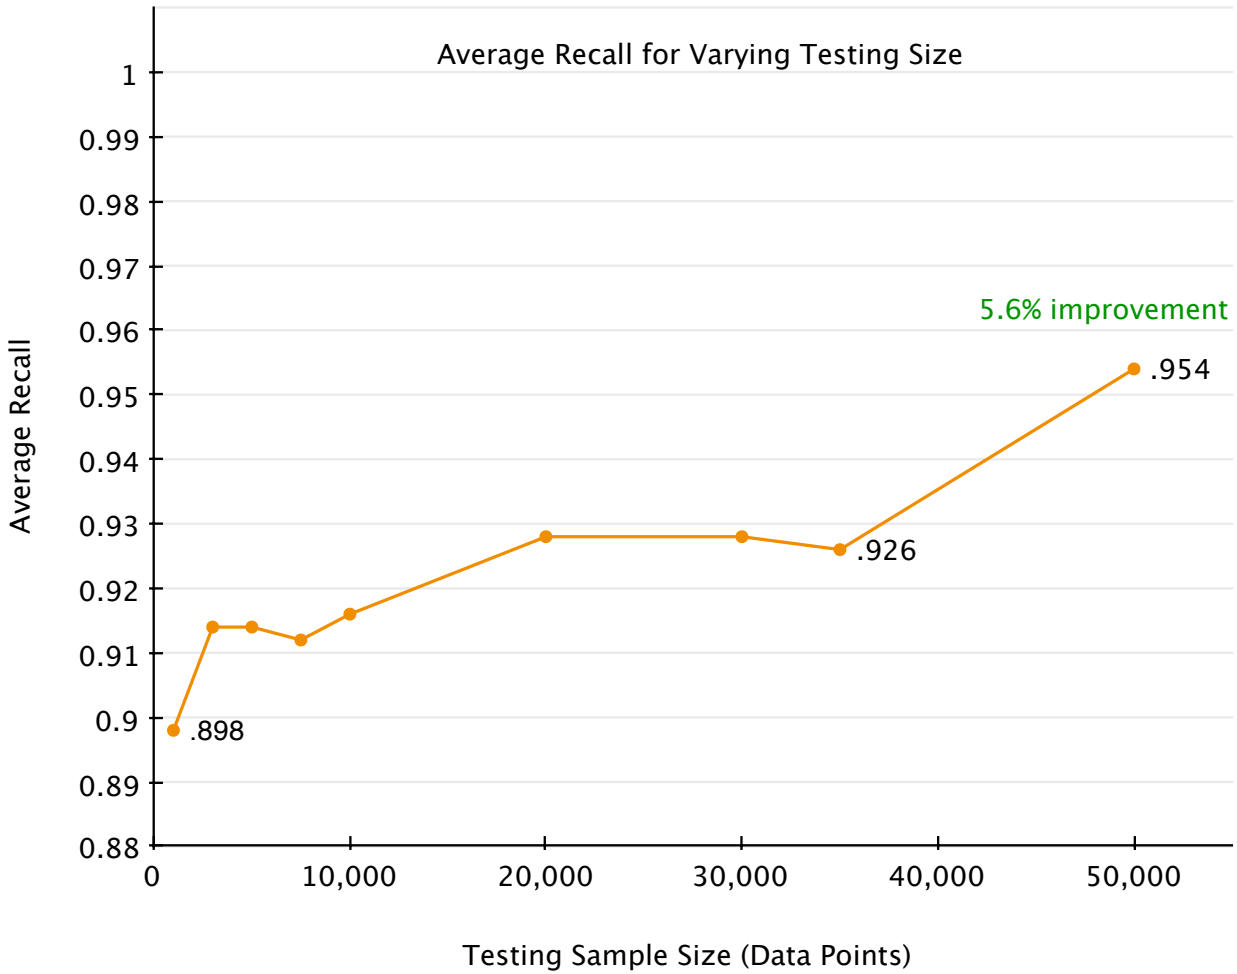
Training Size Selection



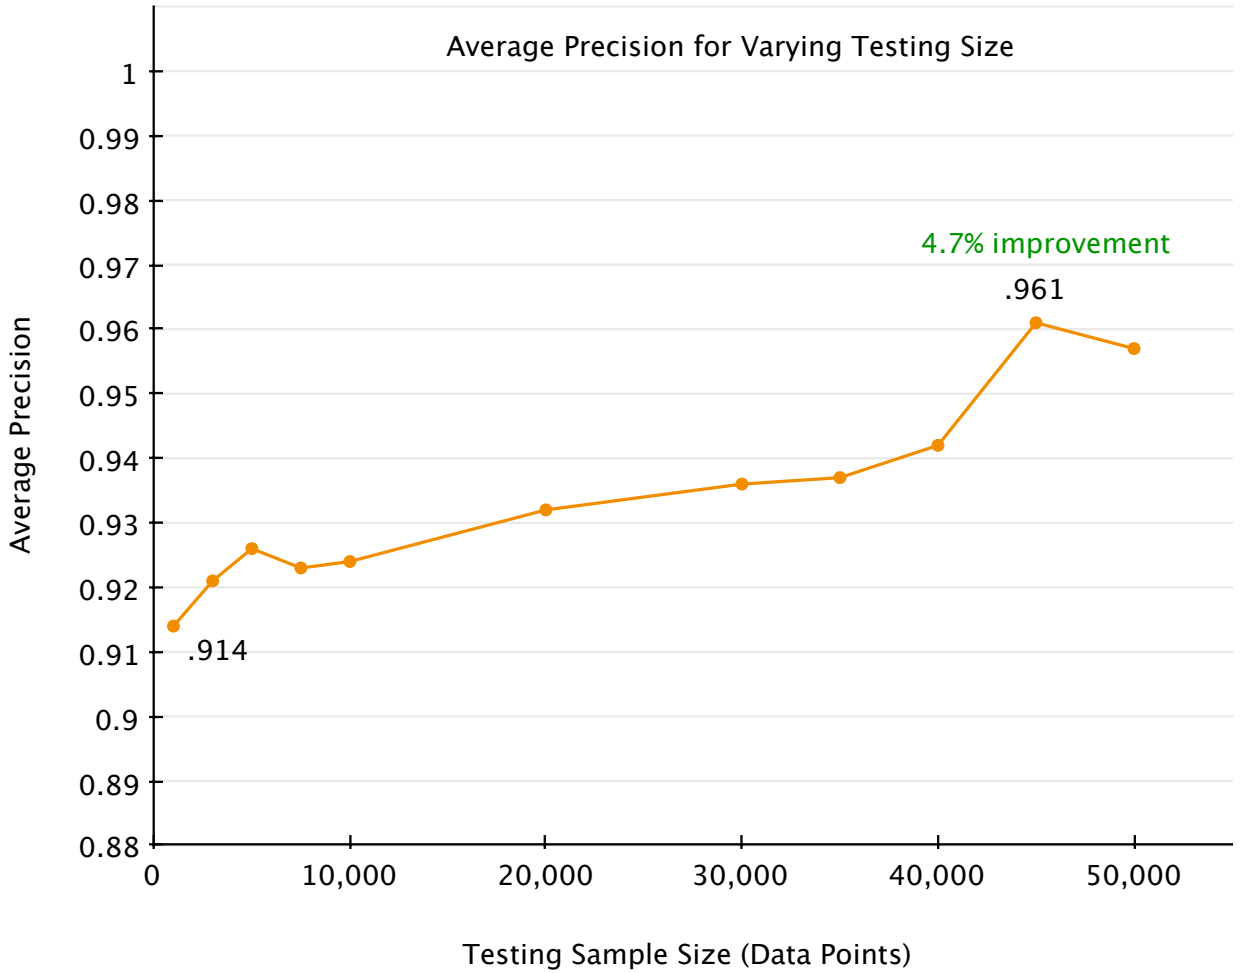
Training Size Selection



Testing Size Selection



Testing Size Selection



Results

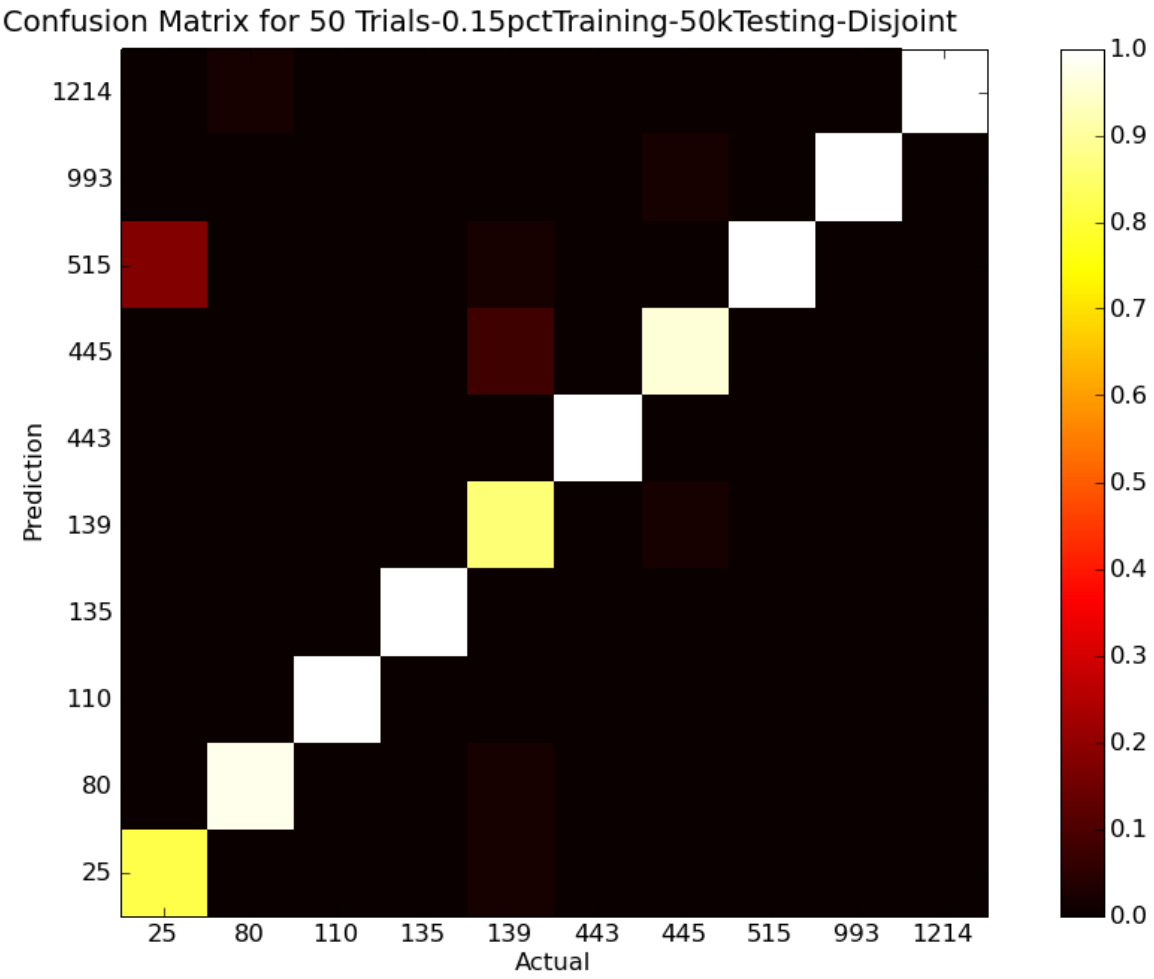
50 Trials

15% Training Set Size

50,000 Data Points Testing Set Size

96.5% Precision

96.0% Recall



Classification Confidence Threshold

- Goal: Eliminate close calls
- Require 1st place candidate to lead 2nd place by certain amount to make decision
- Standard deviation of scores

Methodology v2.0

- Randomly sample some percentage of available data for each port and train classifier
- Randomly sample some number of the remaining data points for each port and create testing samples
- Attempt to classify testing samples
  - If all testing samples reach threshold, done.
  - If any testing sample fails, rebuild testing samples and try again.

Classification Confidence

50 Trials

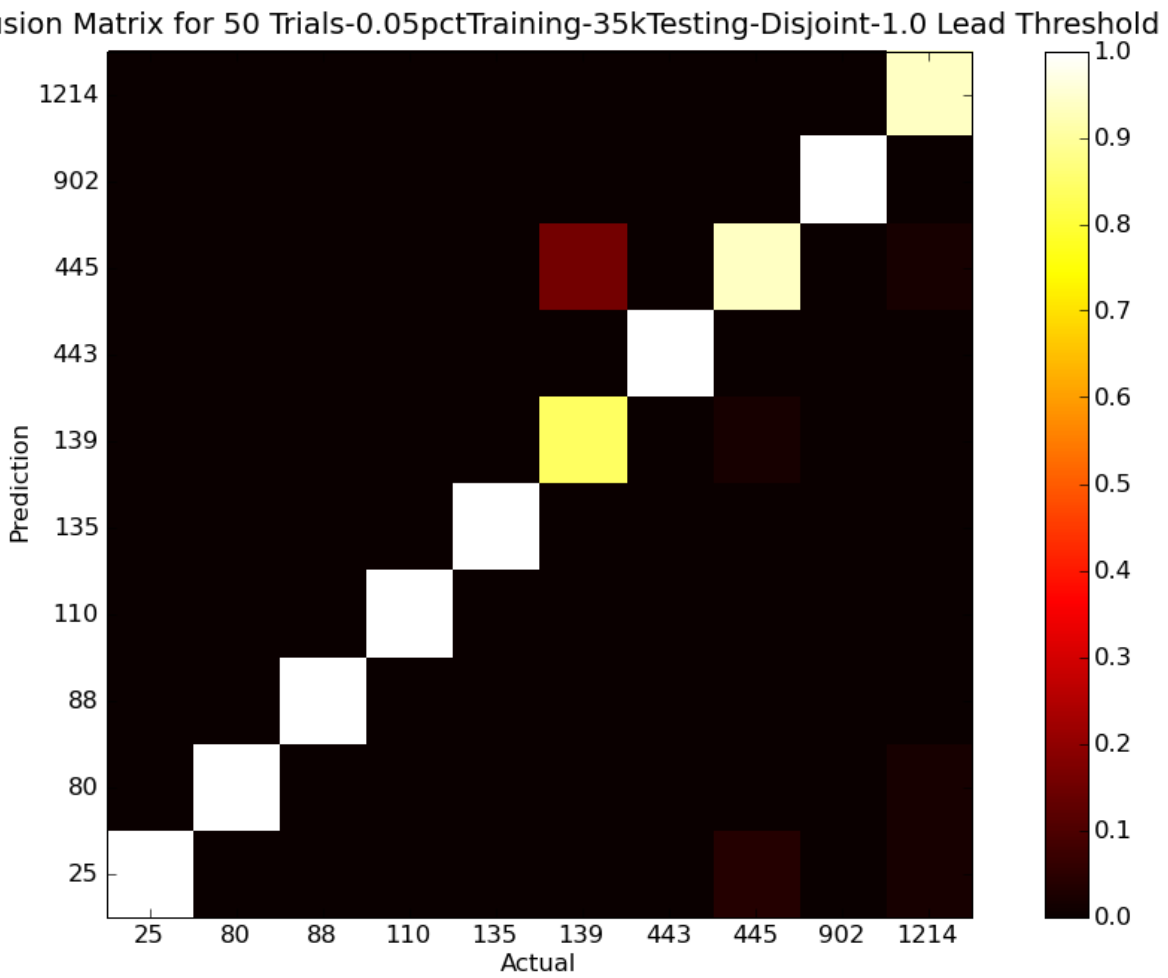
5% Training Set Size

35,000 Data Points Testing Set Size

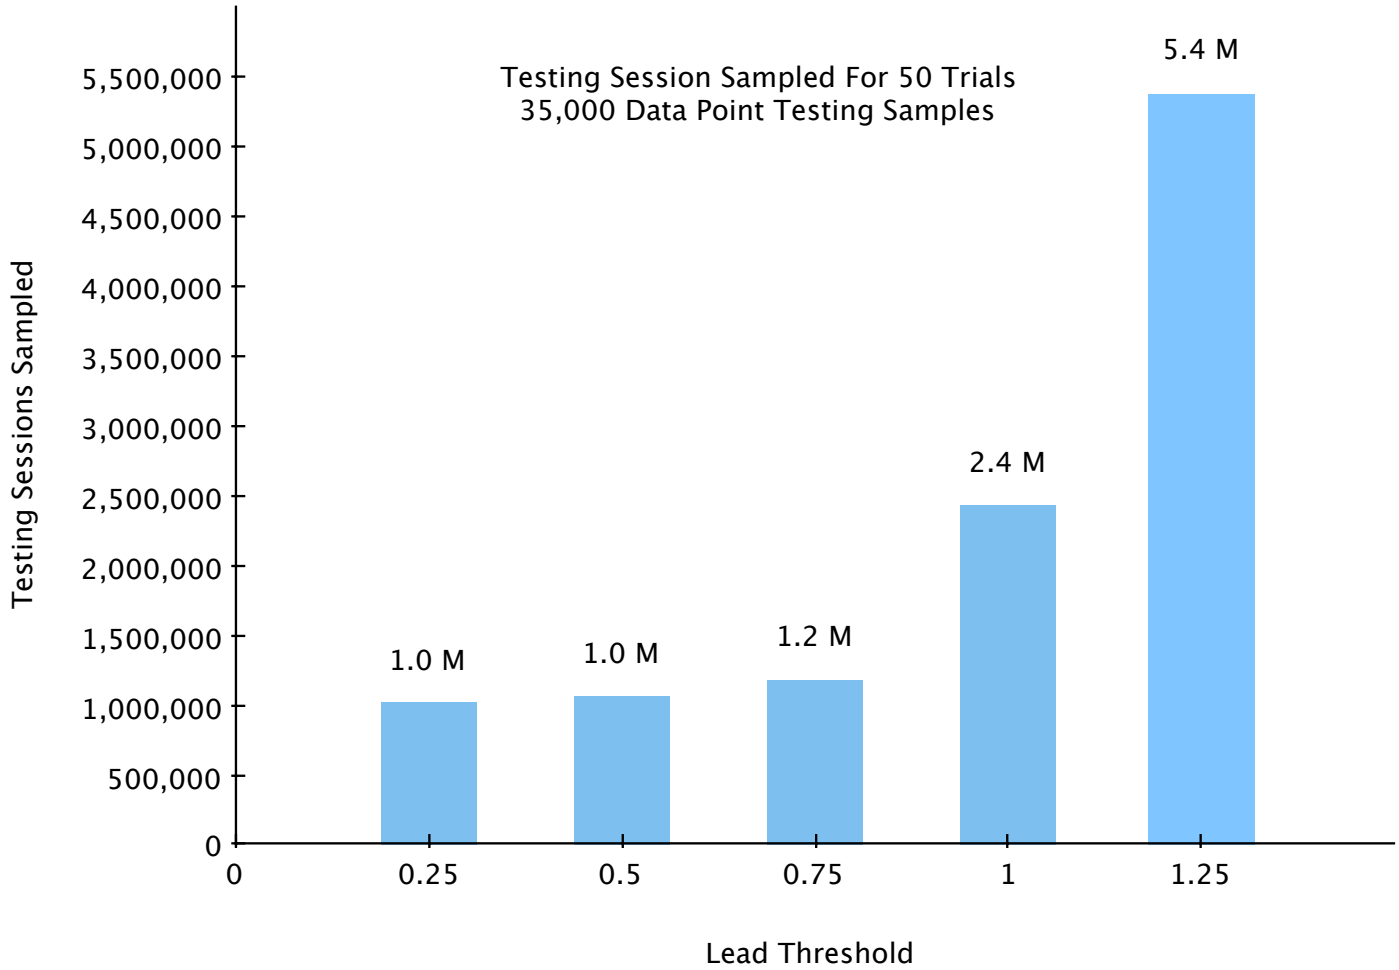
1.0 Lead Threshold

96.9% Precision

96.6% Recall



Classification Confidence



Ground Truth Testing

MIT Lincoln Labs DARPA Data
50 trials, 5% training sample size, 35,000 data point testing sample size, 1.25 lead threshold

- Precision: 98.3%
- Recall: 98.0%

Results

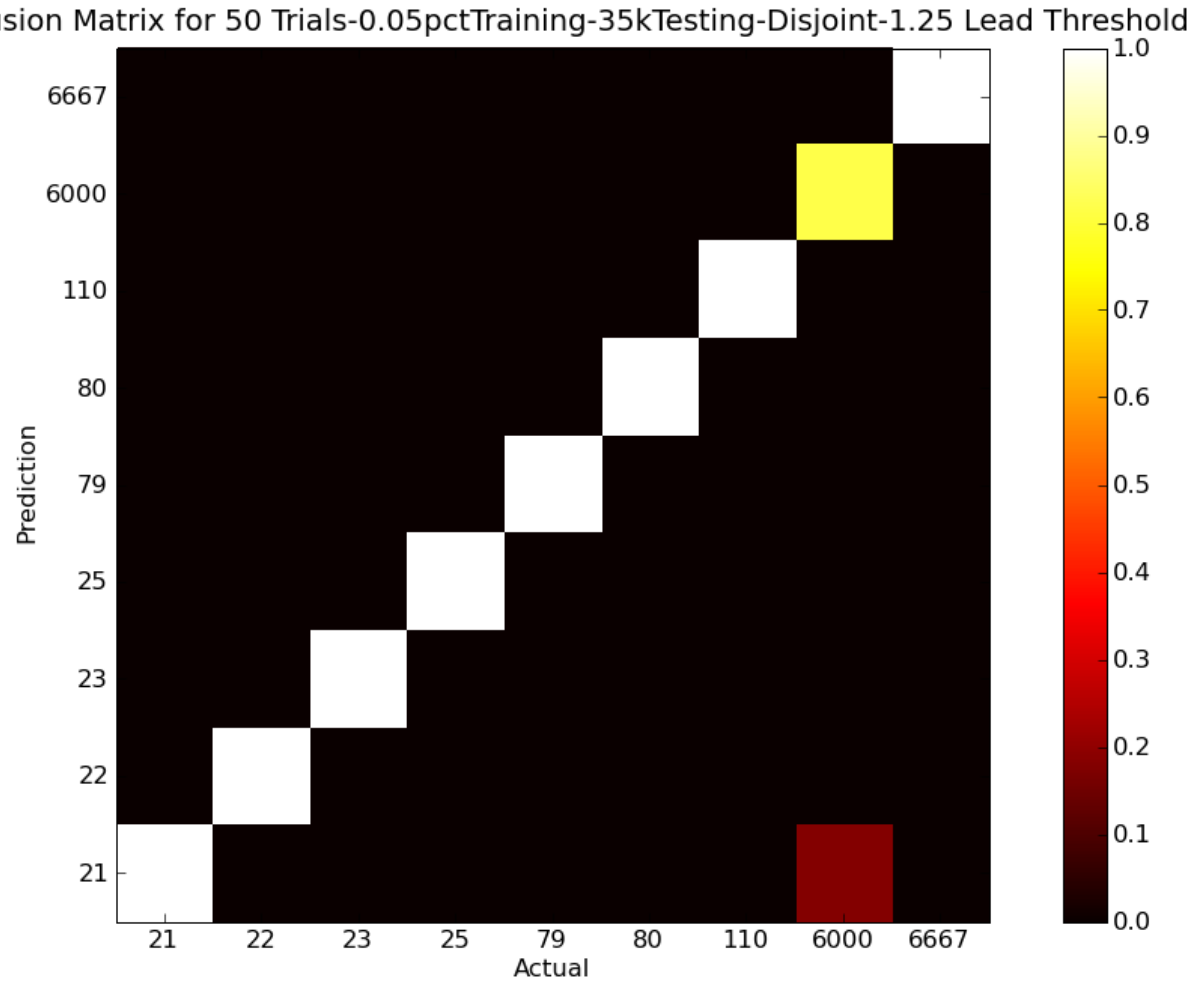
50 Trials

5% Training Set Size

35,000 Data Points Testing Size

1.25 Lead Threshold

Ground Truth

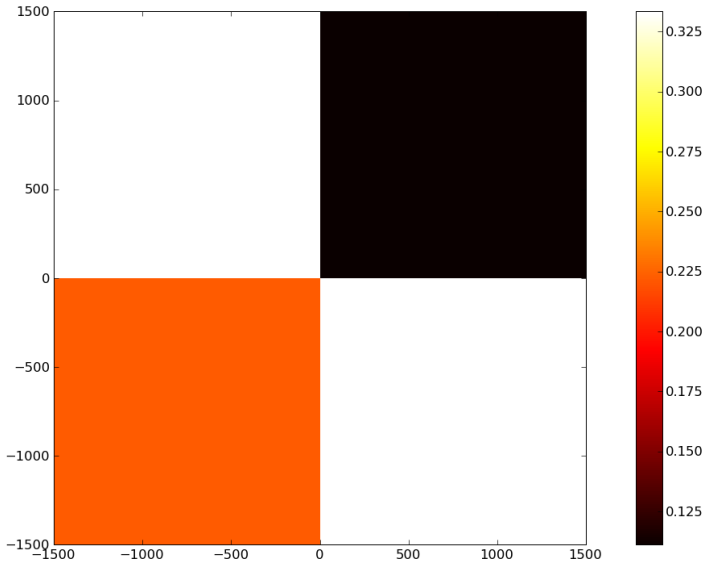


Evasion

One might attempt to thwart our technique by padding all packets to MTU.

Reduces problem to 4-quadrant problem.

Can still make decisions based on relative prevalence of each quadrant.



Current/Future Work

- Packet loss/re-transmission may cause unpredictable results
- On-line classification
- Training and testing from separate datasets
- UDP
- Subcategorization

Conclusion

- Modeling protocol behavior using only packet size, direction, and order
- Resistant to encryption and padding
- Average precision and recall $> 97\%$
- Quick and reliable traffic inspection
- Useful for pre-screening traffic for deeper analysis

Questions?

Thanks for listening.
Q & A

wwlian@gmail.com

"SASUKE" Traffic Monitoring Tool
Traffic Shift Monitoring Based on Correlation between BGP Messages and Flow Data

Atsushi Kobayashi
Yutaka Hirokawa
Hiroshi Kurakami
NTT Information Sharing Laboratories

Outline

- **Introduction**
  - Background
  - Motivation
  - Challenge
- **System Architecture**
  - BGP Collection
  - Flow Collection
  - Visualization
- **Traffic Change Detection Method**
- **Conclusion**

Background

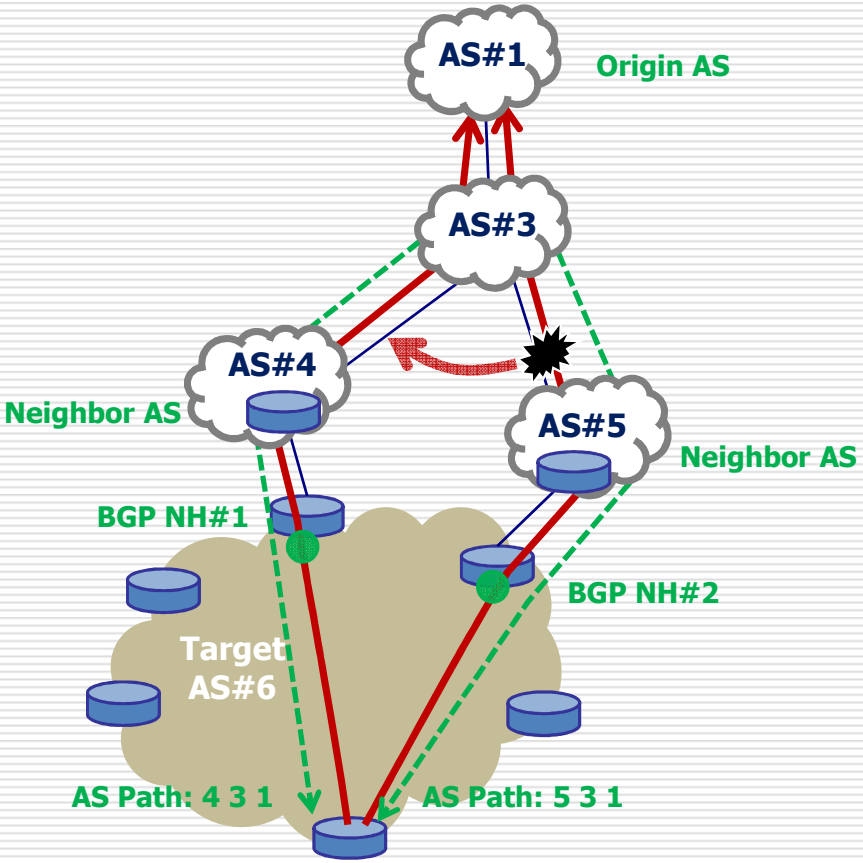
- **Announcement of unwanted or invalid BGP route suddenly leads to traffic diversions.**
  - Cutting of submarine cable, route hijacking, misconfiguration, …
- **Moreover, it disrupts traffic or causes congestion on other backbone links.**





Motivation

- **Our goal is to reduce the load for troubleshooting.**
- **Our tool detects a traffic change and then identifies BGP route announcements involved.**
  - Monitors traffic volume for BGP attributes that have an impact on the traffic change:
    - ✓ Origin ASN
    - ✓ Neighbor ASN (peer ASN)
    - ✓ AS Path
    - ✓ BGP Next Hop
    - ✓ Community.
  - Identifies route changes that have an impact on the traffic change.



Related Work

- **Flow records from border routers can be utilized for origin or neighbor ASN traffic analysis.**
  - However, border gateway router cannot export both origin and neighbor ASNs.
  - Difficult to collect BGP Next Hop and AS Path info.
- **Some commercial collectors with BGP sessions can sum up traffic on the basis of BGP attributes.**
  - There are few tools for analyzing the interrelation of BGP and Flow data.
- **BGP and Flow analysis system have been proposed by several groups[1)].**
  - Simpler method and its visualization are required.

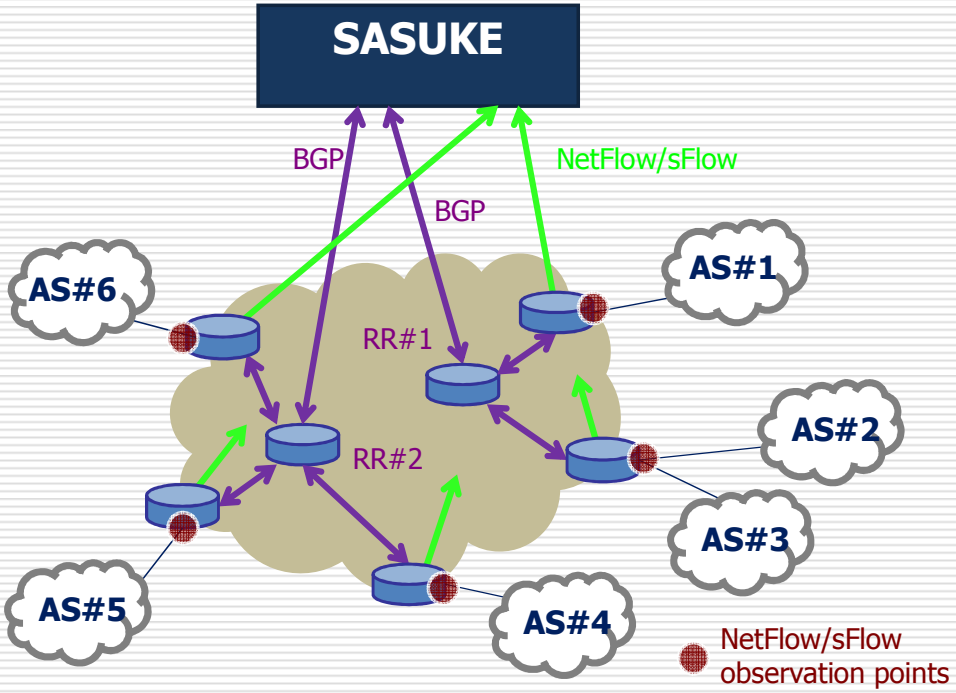
1) For example, J. Wu, Z. M. Mao, J. Rexford, and J. Wang, "Finding a needle in a haystack: Pinpointing significant BGP routing changes in an IP network," in Proc. NSDI, May 2005.

Challenge

- **The challenge is to identify route changes from a huge number of BGP route announcements.**
  - Hundreds of thousands of route announcements per day
- **Handle the huge load of flow records.**
  - Thousands of flow records per second
- **Explore a simple detection method and its real-time visualization.**

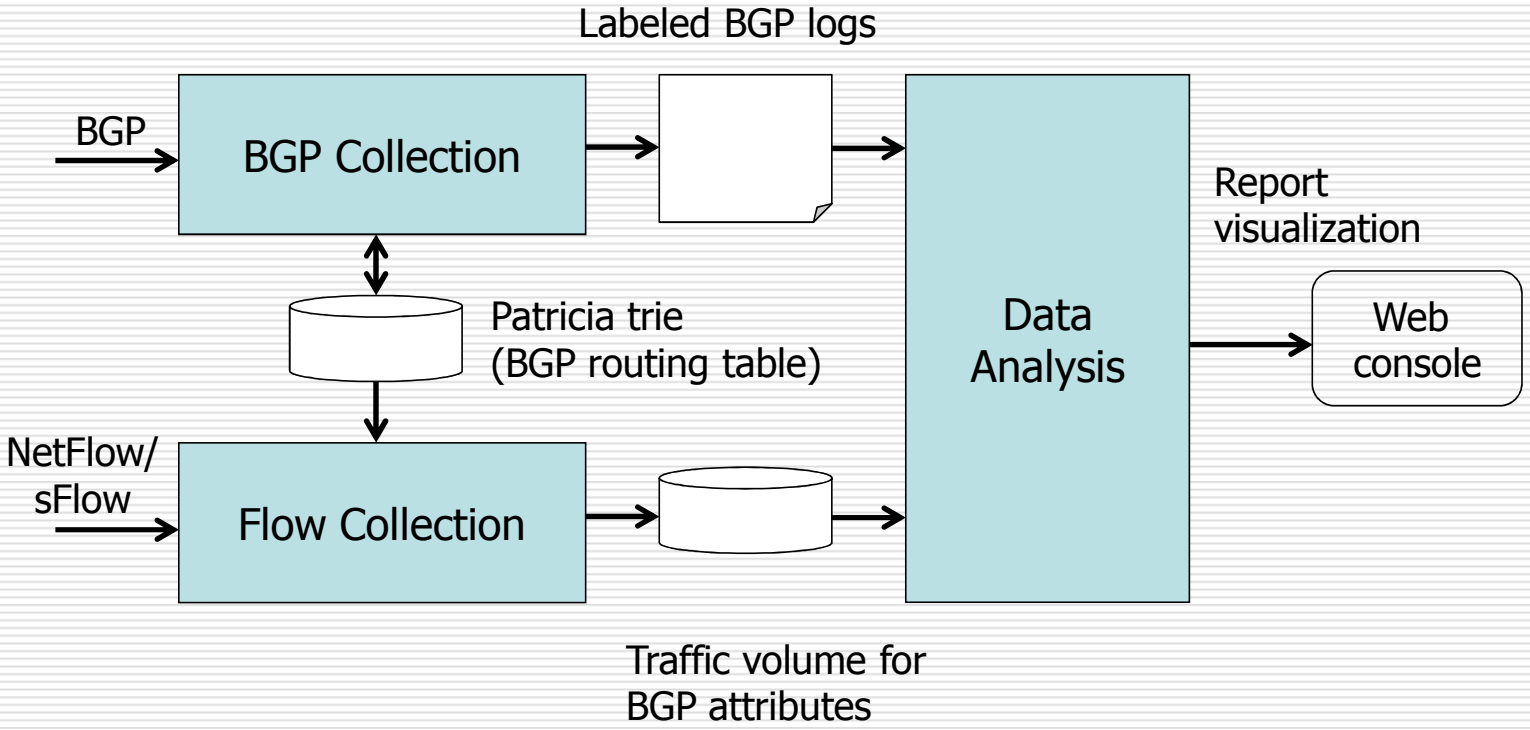
Data Source

- **Captures BGP data from BGP sessions to border routers or route reflectors as a BGP route reflector client.**
  - Border router feeds best routes to SASUKE tool.
- **Sets NetFlow/sFlow observation points at the periphery of the target AS.**



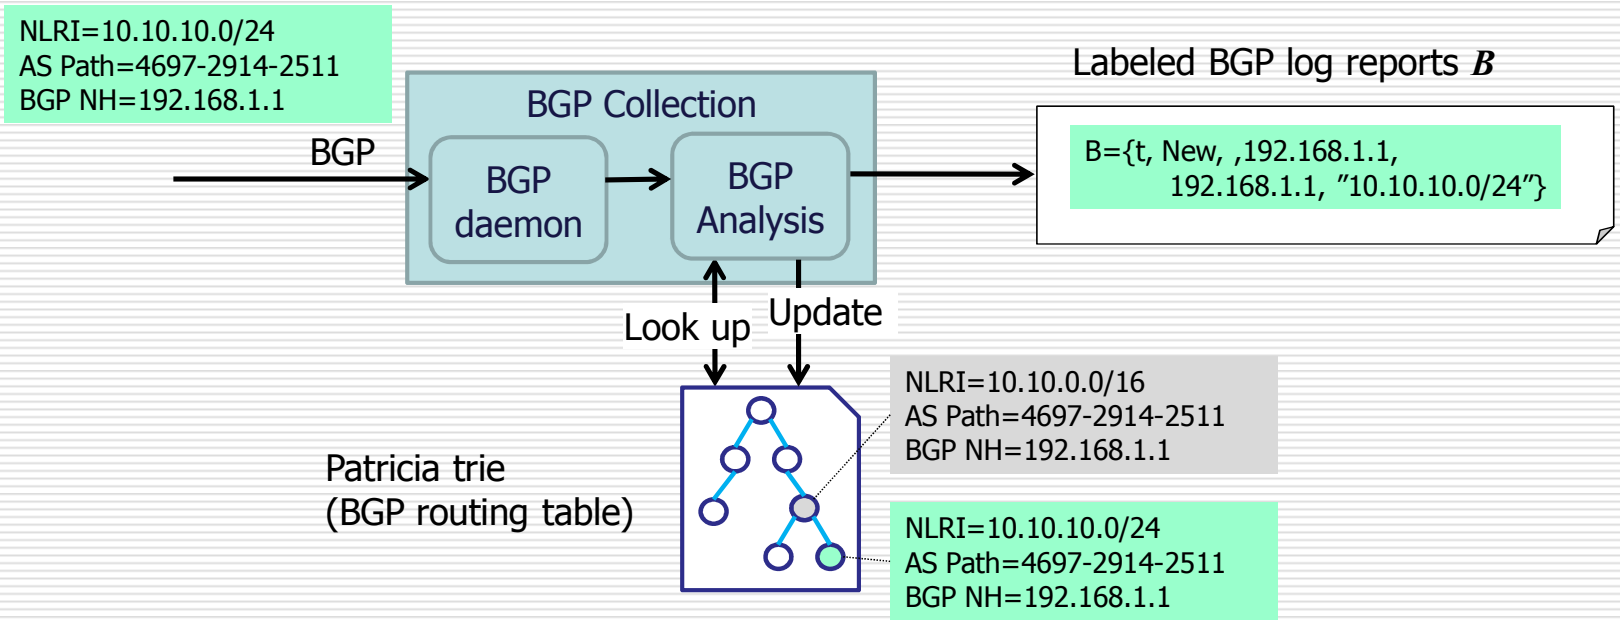
System Architecture

- 3 system components:
  - BGP Collection
  - Flow Collection
  - Data Analysis: correlation between BGP and Flow data



BGP Collection

- **Builds BGP routing tables as Patricia trie.**
  - Maintains tables for each BGP peering session.
- **Creates a BGP log report B to identify BGP messages that may cause a traffic change by comparing against the Patricia trie.**
  - Identifies BGP message type and BGP attributes that have changed from the old ones in the Patricia trie.



BGP Log Reports

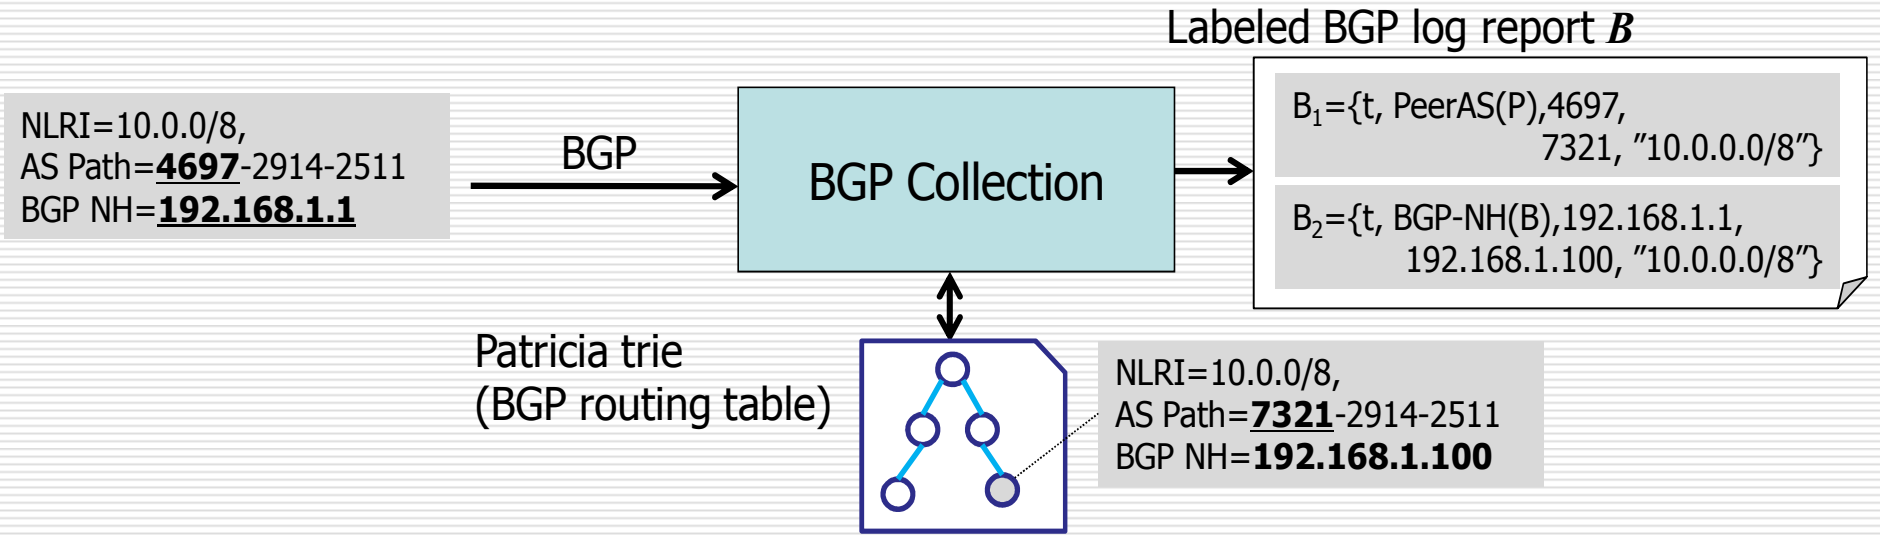
- **BGP log report is represented as follows.**

$$B=\{t, c_{type}, a_{type}, a_{new}, a_{old}, prefix, id\}$$

  - t is timestamp of when the BGP message arrived.
  - c_{type} is change type:
    - "*New*", "*Withdraw*", "*Change*", or "*Duplicate*".
  - a_{type} indicates the changed BGP attribute type:
    - *"Origin ASN", "Neighbor ASN", "AS Path", "BGP NH"*, or *"Community"*
      - BGP community often gives route categories: region, peering type.
  - a_{new}/a_{old} are new/old BGP attribute values.
    - When c_{type} is *"New"* or *"Withdraw"*, a_{old} or a_{new} is a "null" value.
  - ***Prefix*** is network address in NLRI.
  - ***Id*** is an identifier to correlate with traffic data.
    - At this stage, the value is "null".

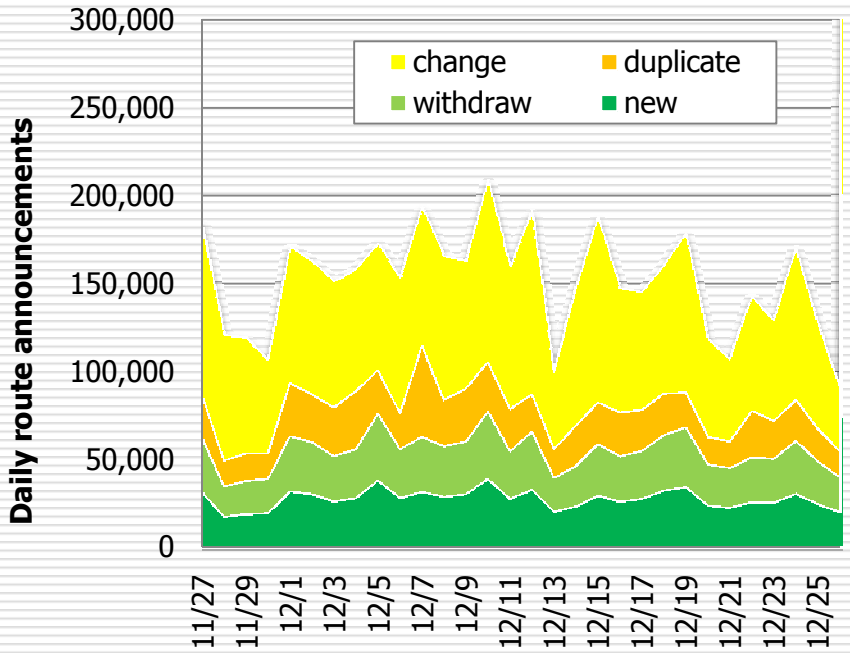
BGP Log Reports

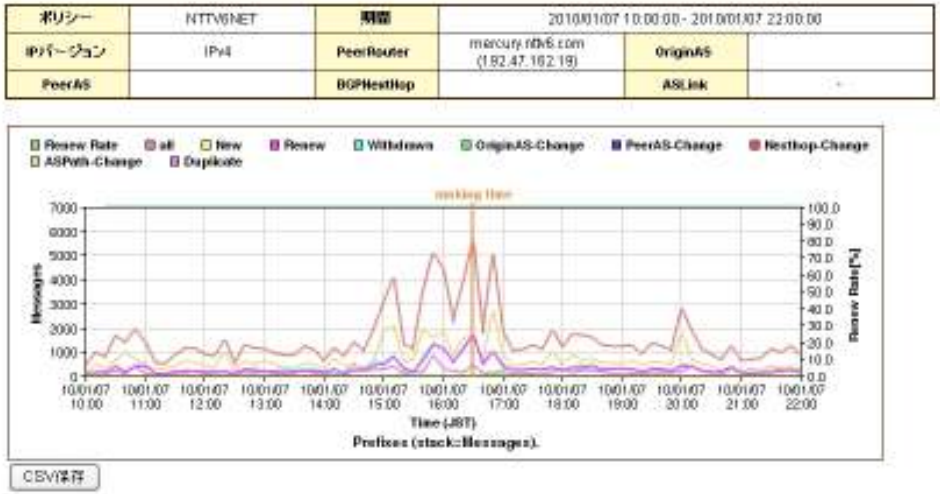
- **Creates multiple BGP log reports when multiple BGP attributes are changed.**



Visualization for BGP Log Reports

- **Labeled BGP logs are presented in time-series.**
  - Top-N origin/neighbor ASNs involved in the most BGP messages are represented when the spike happens.



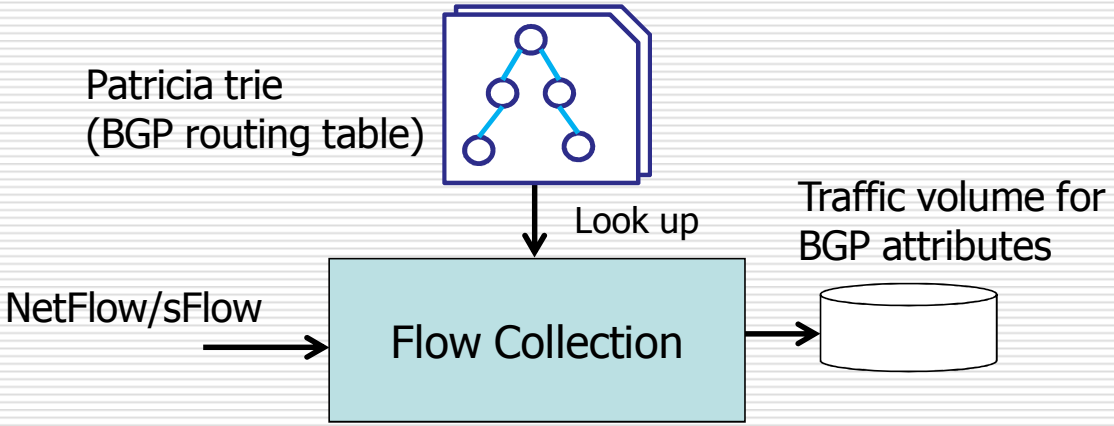


Top20 All Ranking

No.	OriginAS	Count	OriginAS		
			Country	Region	Organization
1	7545	3532		apnic	TPG-INTERNET-AP TPG Internet Pty Ltd
2	14522	425		lacnic	Satnet
3	5800	371		arin	DNIC-ASBLK-05800-06055 - DoD Network Information Center
4	20858	200		afrinic	EGYNET-AS
5	8474	102		ripencc	AS8474 Riscom AS
6	5536	87		afrinic	Internet-Egypt
7	47869	60		ripencc	SMARTLINK-AS Smart Links of Telecommunication Services
8	12524	35		ripencc	INDEX Autonomous System
9	6822	35		ripencc	SUPERONLINE-AS SuperOnline autonomous system
10	3912	34		arin	NMSU-AS - Cherry-net

Flow Collection

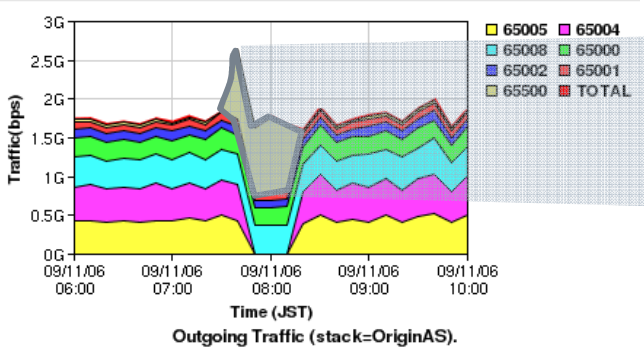
- **Selects an appropriate Patricia trie**.
  - Compares the Flow Record and peering data.
    - Are Exporter and BGP router the same device or not?
    - Are Exporter and BGP router located in the same region or not?
- **Sums up traffic on the basis of BGP attributes:**
  - Origin ASN, Neighbor ASN, AS Path, BGP NH, Community, and Prefix
  - These BGP attributes are retrieved from the Patricia trie by a longest match based on source/destination IP addresses.



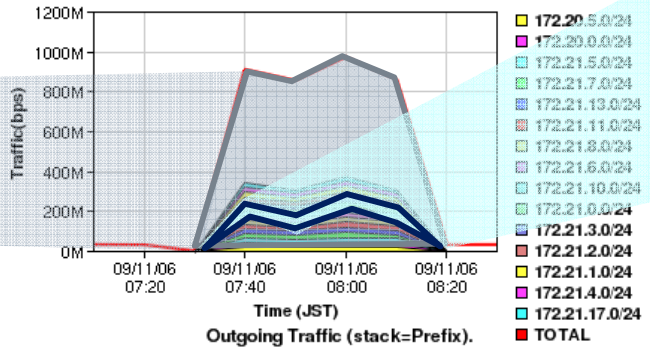
Visualization for Traffic Data

- Drill down into the detailed traffic data step by step from stacked area chart.

Origin ASNs stacked



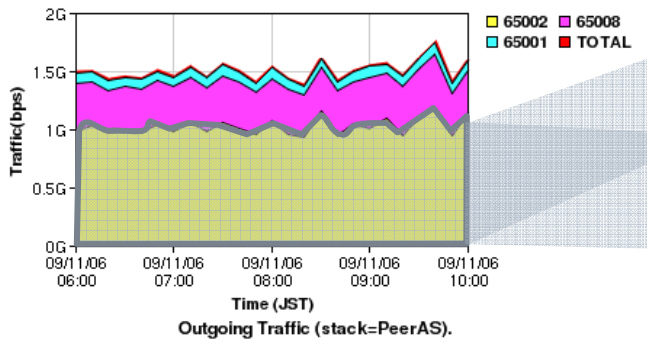
Prefix stacked



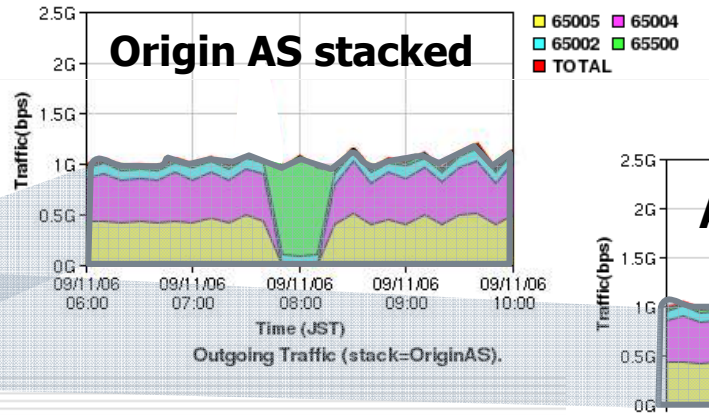
Detailed Flow record list



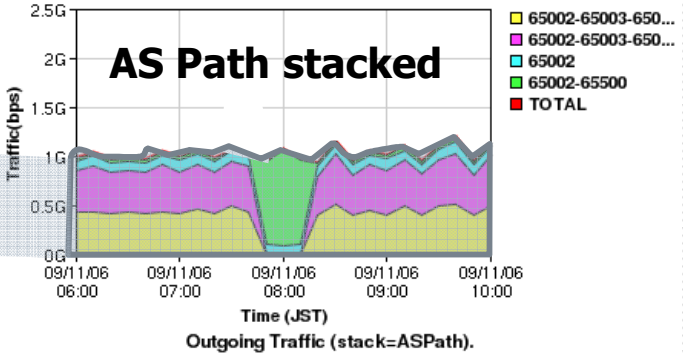
Neighbor ASNs stacked



Origin AS stacked

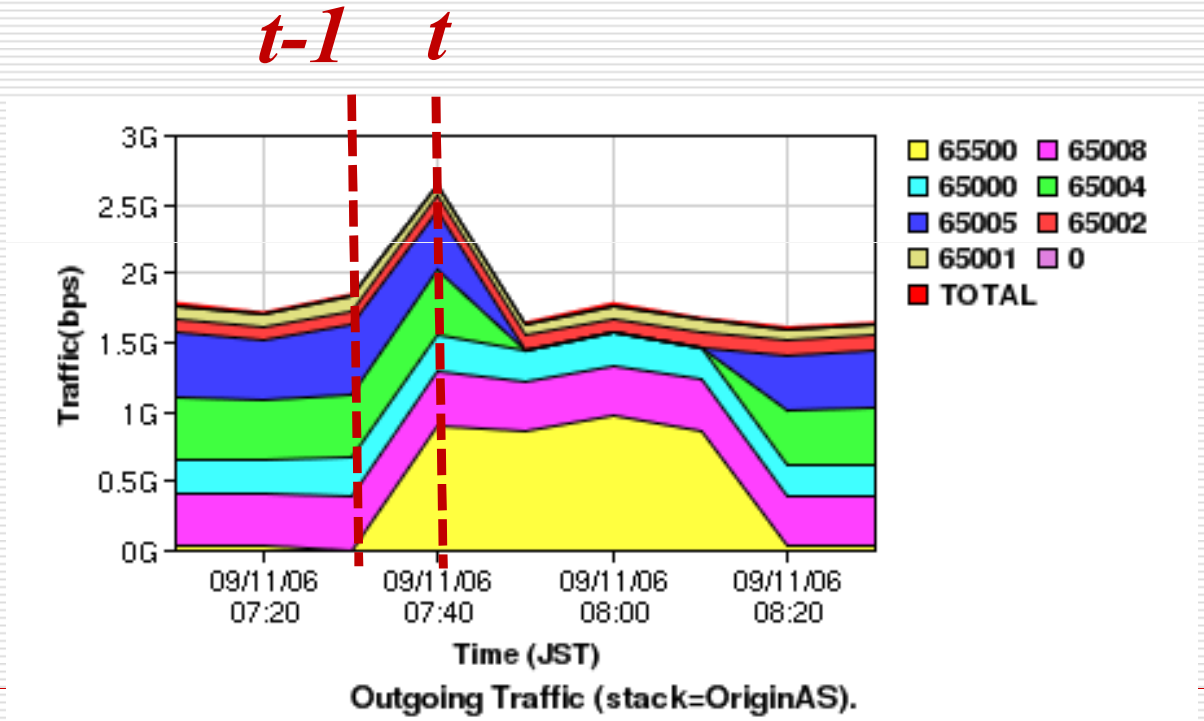


AS Path stacked



Traffic Change Detection Method

- **Focuses on a Top-N ranked by traffic volume on BGP attributes:**
  - Origin ASN, Neighbor ASN, AS Path, BGP NH, and Community.
- **Evaluates similarities of Top-N ranks between time slot t and t-1.**
  - Traffic volume weights the evaluation results.



Traffic Change Detection Method

- **Calculates the correlation coefficient *r(t,t-1)* between the ranks of time slots *t* and *t-1*.**

$$r(t,t-1)=\frac{\sum_{i=1}^{N}\left(c(f(i,t),t)-c_{avg}(t)\right)\left(c(f(i,t),t-1)-c_{avg}(t-1)\right)}{\sqrt{\sum_{i=1}^{N}\left(c(f(i,t),t)-c_{avg}(t)\right)^2}\sqrt{\sum_{i=1}^{N}\left(c(f(i,t),t-1)-c_{avg}(t-1)\right)^2}}$$

  - ***f(i,t)*** is defined as a BGP attribute value ranked ***i*** by traffic volume of time slot ***t*.**
  - ***c(f(i,t),t)*** is traffic volume of ***f(i.t)*.**
  - Top-***N*** statistics data set $C(t)_i$ is presented as an array:
    - $C(t)_i=\{c(f(1,t),t), c(f(2,t),t), \ldots, c(f(N,t),t)\}$ $(i=1, 2, 3, \ldots, N)$
  - $C_{avg}(t)$ gives the average of $C(t)_i$.
- **Evaluates whether *r(t,t-1)* exceeds a threshold.**

Identify Most Affected Traffic

- **Investigates which $f(i,k)$ has the greatest impact on traffic change, as follows.**
  - Top-N rank $C(k)(t)_i$ indicates $C(t)_i$ except for $c(f(k,t),t)$ ranked k.

    $C(k)(t)_i=\{c(f(1,t),t), \ldots, c(f(k-1,t),t), c(f(k+1,t),t), \ldots, c(f(N,t),t)\}$
  - Calculates the correlation coefficient $r(k)(t,t-1)$.
  - Selects the greatest values $r(k)(t,t-1)$ from

    $\{r(1)(t,t-1), r(2)(t,t-1), \ldots, r(N)(t,t-1)\}$.
- **Then, we recognize that $f(k,t)$ is the most affected BGP attribute.**

Traffic Change Reports

- **Finally, it creates a traffic change report *R*.**

$$R=\{t, type, f(i,k), r(t,t-1), \delta, id\}$$

  - *type* gives traffic volume type:
    - *"Origin ASN", "Neighbor ASN", "AS Path", "BGP NH"*, or *"Community"*
  - δ gives the traffic volume difference between *t* and *t-1*.
  - *id* is an identifier to correlate with BGP log reports.
    - At this stage, the value is set to a unique value.

Correlation between BGP and Flow

- **Correlates between BGP log report B and traffic change report R.**
  - Looks for the BGP log report B involved with traffic change report R.
  - Then, R and B are given the same id value to link them.

$R=\{t,\ type,\ f(i,k),\ r(t,t\text{-}1)\ ,\ \delta,\ id\}$
$B=\{t,\ c_{type},\ a_{type},\ a_{new},\ a_{old},\ prefix,\ id\}$

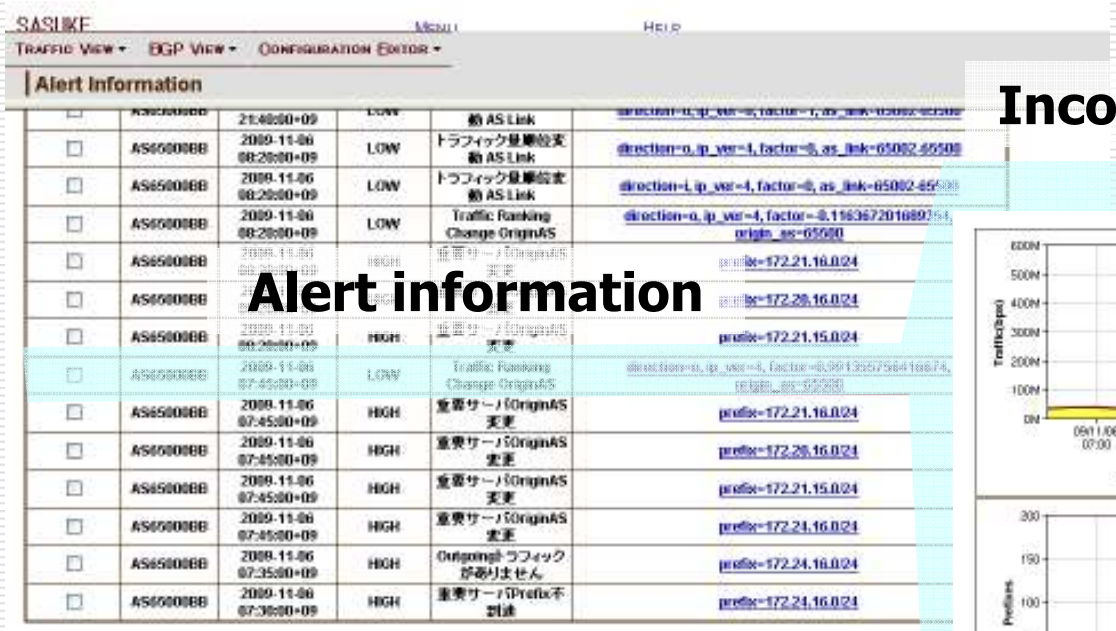
for all BGP log reports B where $t\text{-}Tw < B.t < t+Tw$ do

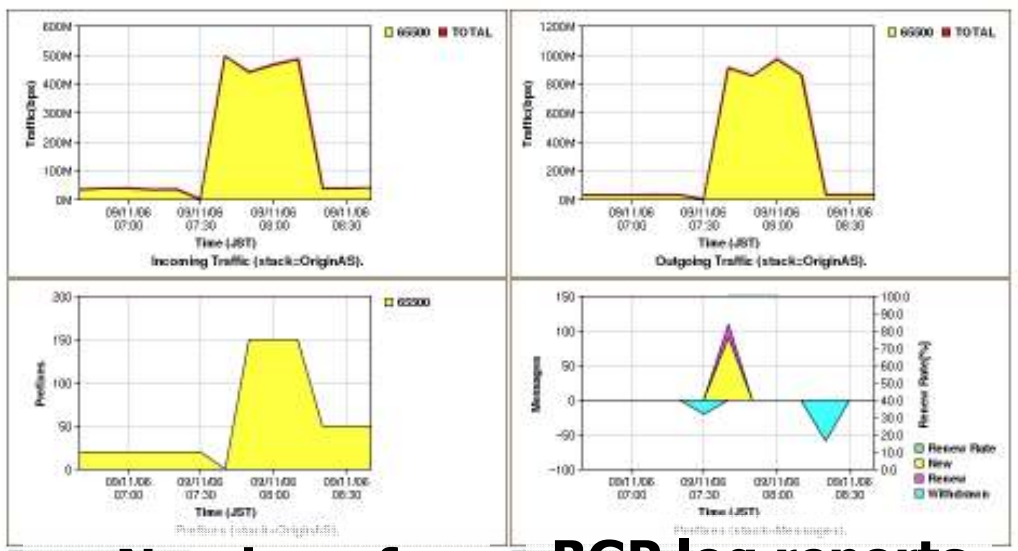
if $R.\delta > 0$ and $B.a_{type} = R.type$ and $B.a_{new} = R.f(i,k)$ *then*
$B.id = R.id$;
else if $R.\delta < 0$ and $B.a_{type} = R.type$ and $B.a_{old} = R.f(i,k)$ *then*
$B.id = R.id$;
end if

end for

Visualization for BGP and Flow

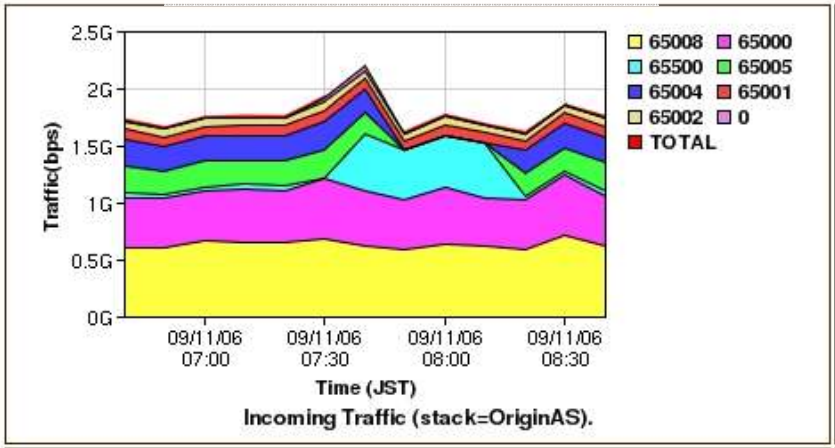
- **Creates traffic-change-related alert for operators.**
  - Alert links the graphs of traffic volume area chart and of BGP log reports involved.



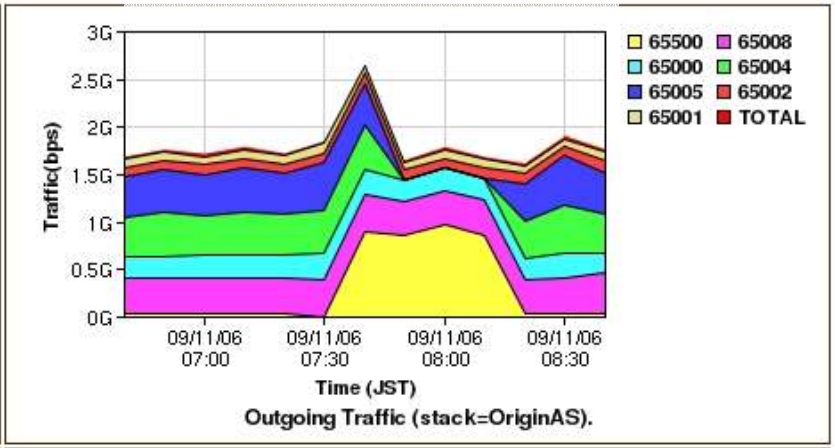


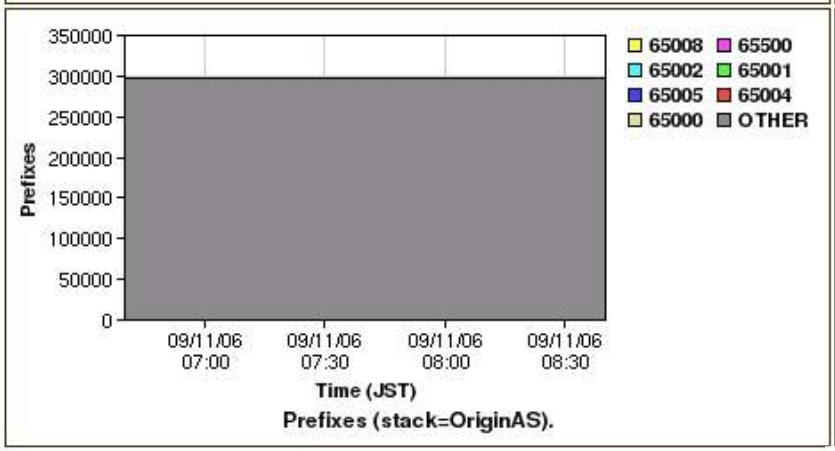
Overall Traffic View

Incoming Traffic

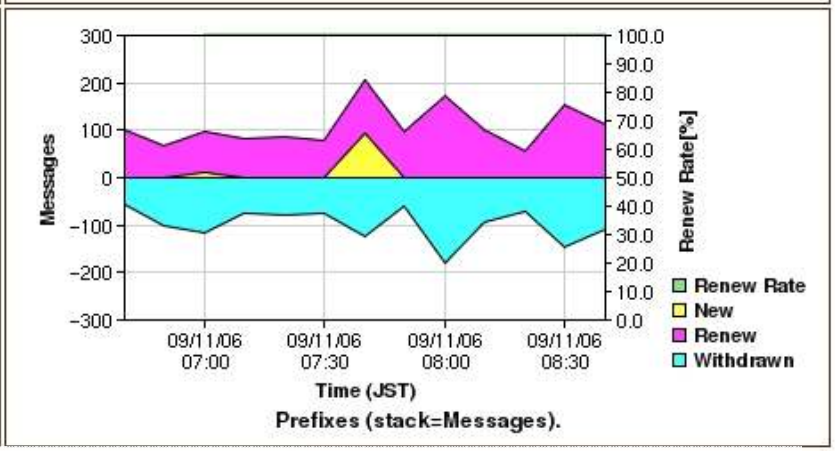


Outgoing Traffic





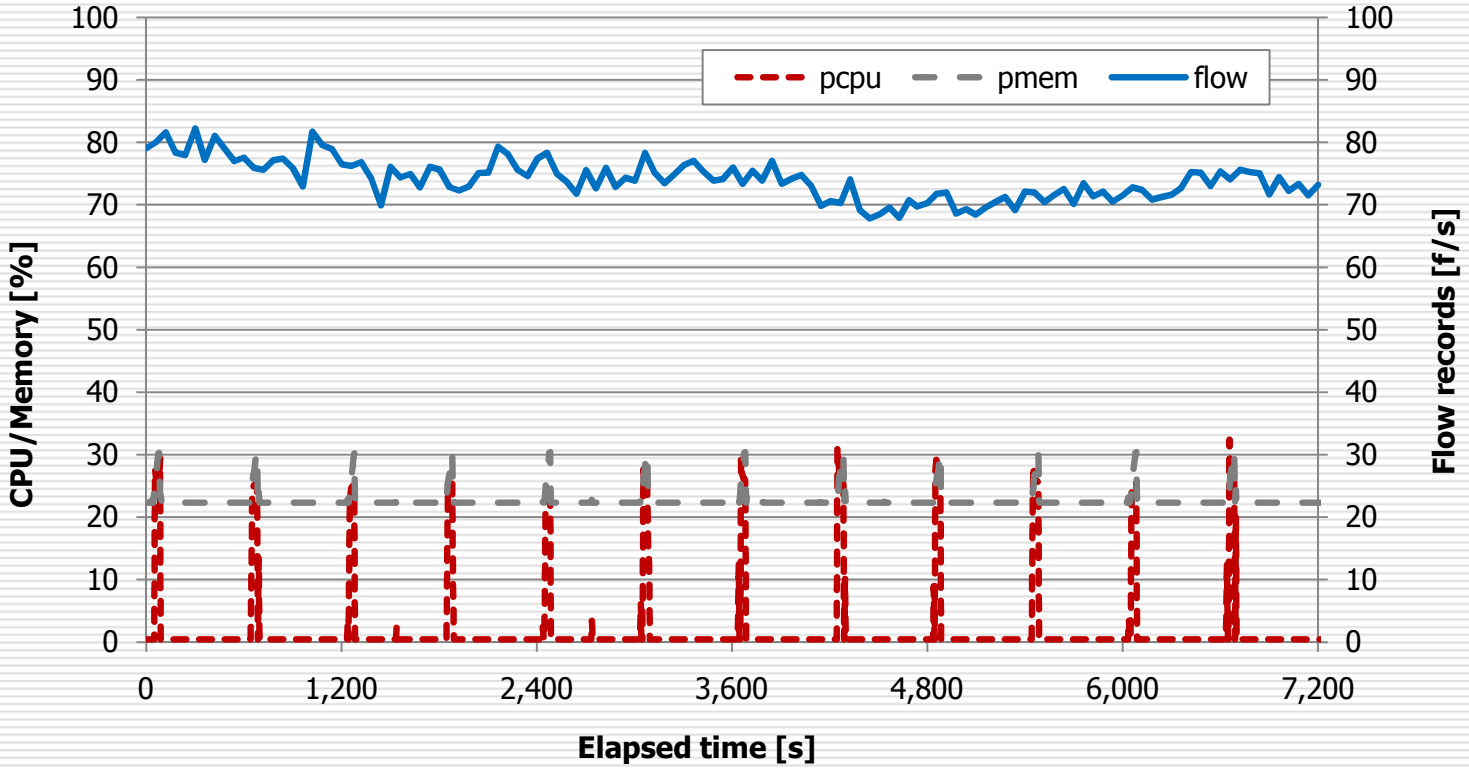
Number of Prefix



BGP log reports

Evaluation of SASUKE Tool

- **"SASUKE" has been introduced in some commercial networks as an experimental phase.**
  - Much more performance evaluation is needed.



Conclusion

- **I demonstrated the traffic change detection method implemented in the "SASUKE" tool.**
  - Focuses on the similarities between time consecutive Top-N ranks in time-series.
  - Correlates between BGP log reports B and traffic change reports R.
  - Alleviates the troubleshooting load for network operators.
    - Visualizes BGP log and traffic data.
    - Links multidimensional traffic data related to BGP attributes.
- **More evaluation is needed for performance and accuracy.**

Thank you very much.

This study was supported by the Ministry of Internal Affairs and Communications of Japan.

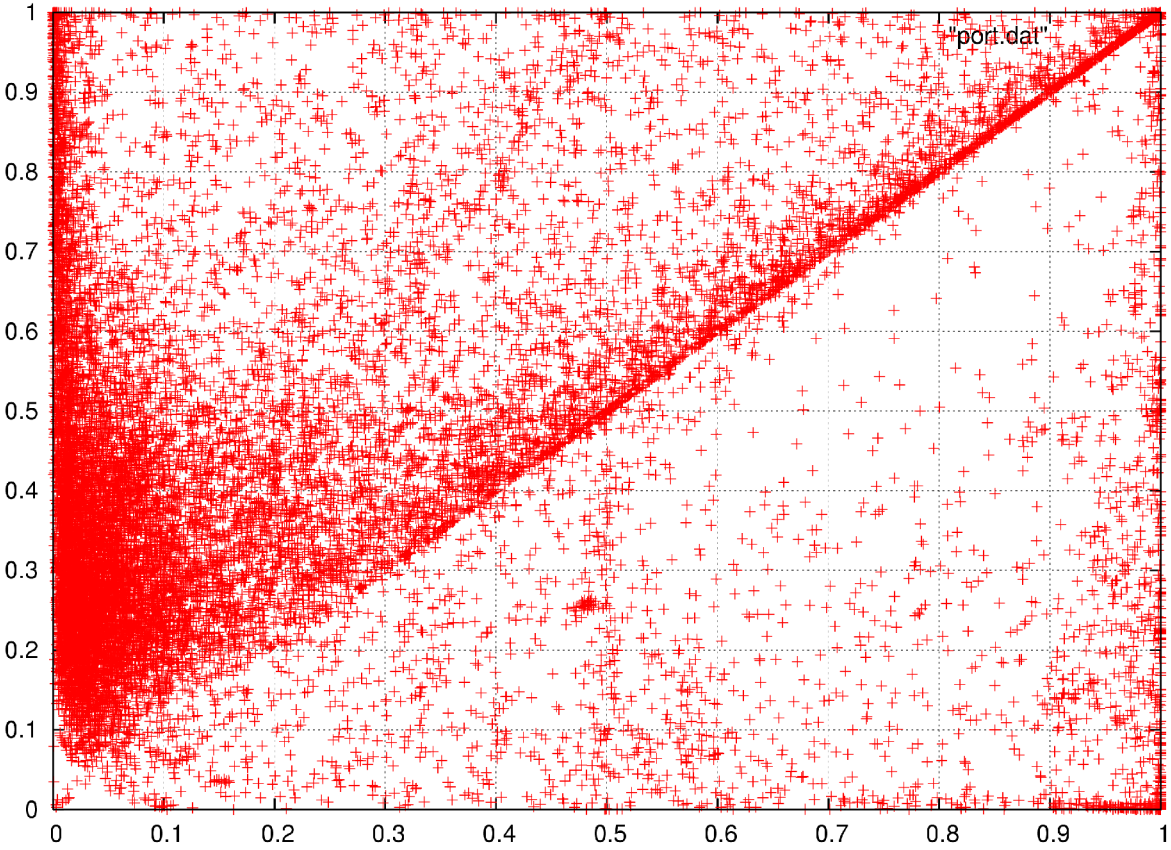
Network Host Classification Using Statistical Analysis of Flow Data

Alex Kent, Mike Fisk, Eugene Gavrilov
Los Alamos National Laboratory

Overview and Objectives

- Host/IP address profiling based on flow data over some time interval
  - 10 minutes to 7 days have been examined with 24 hours providing repetitively stationary results
  - Generate histograms of peer hosts, source ports, and destination ports over the time interval
  - Compute Shannon entropy values for the 3 dimensions
- Outcomes:
  - Provide IP behavior “snapshots” of individual hosts
  - Allow comparison of behavior through clustering
  - Build models over large host sets in real-time

Source / Destination Port Variance



Does not provide an effective representation of categorical data sets

Sample (simplified) Histogram of Flow Data

		Cumulative bytes	Cumulative packets	Cumulative sessions
Host A				
	Peer A	34958	324	54
	Peer B	3948	132	13
	Peer C	231	43	9
	Peer D	5675	123	29
	Src Port 1	2358	77	32
	Src Port 2	13246	345	67
	Src Port 3	1231	75	12
	Dst Port 1	54467	5653	199
	Dst Port 2	563	345	1
Host B	Peer X	842	347	23
	Peer Y	23879	3452	874
	Peer Z	9463	232	78
	...	...	...	...
...				

Shannon Entropy
Using Packet Count Histograms

$$Entropy = -1 \times \sum_{i=1}^{N} p_i \times log_2(p_i)$$

$$\text{where} \quad p_i = \frac{Packets_i}{TotalPackets}$$

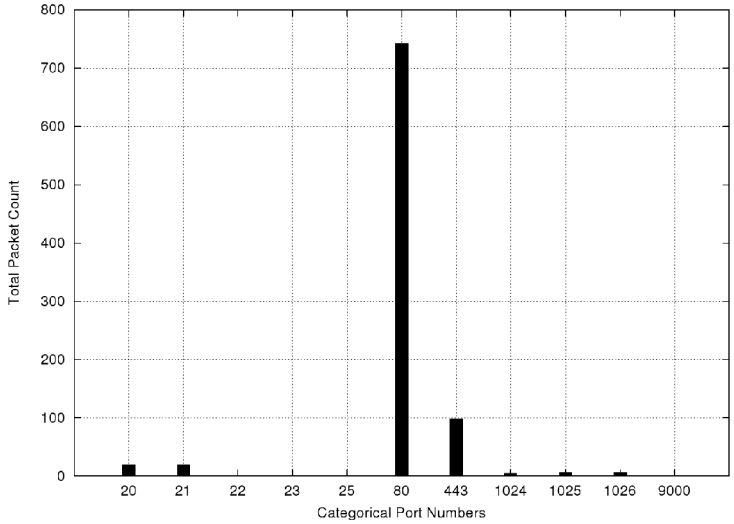
- Computed for host peers, source ports, and destination ports time-delineated histograms leveraging byte, packet, and session totals
  - Packet histograms/entropy calculation favored in final analysis
  - Since base 2: port entropy will be 0-16, peer 0-32 (IPv4)

Sample Entropy Calculation

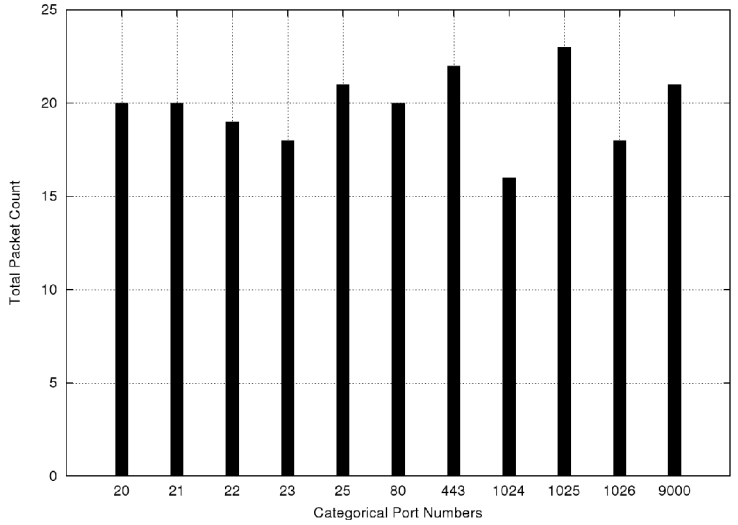
For Host A:

	Cumulative packets	*Pi* Packets	*Pi* * log2(*Pi*)	Entropy
Peer A	324	0.52	-0.49	
Peer B	132	0.21	-0.47	
Peer C	43	0.07	-0.27	
Peer D	123	0.20	-0.46	1.69
Src Port 1	77	0.15	-0.42	
Src Port 2	345	0.69	-0.37	
Src Port 3	75	0.15	-0.41	1.19
Dst Port 1	5653	0.94	-0.08	
Dst Port 2	345	0.06	-0.24	0.32

Visual Example of Low/High Entropy



Low Entropy Example

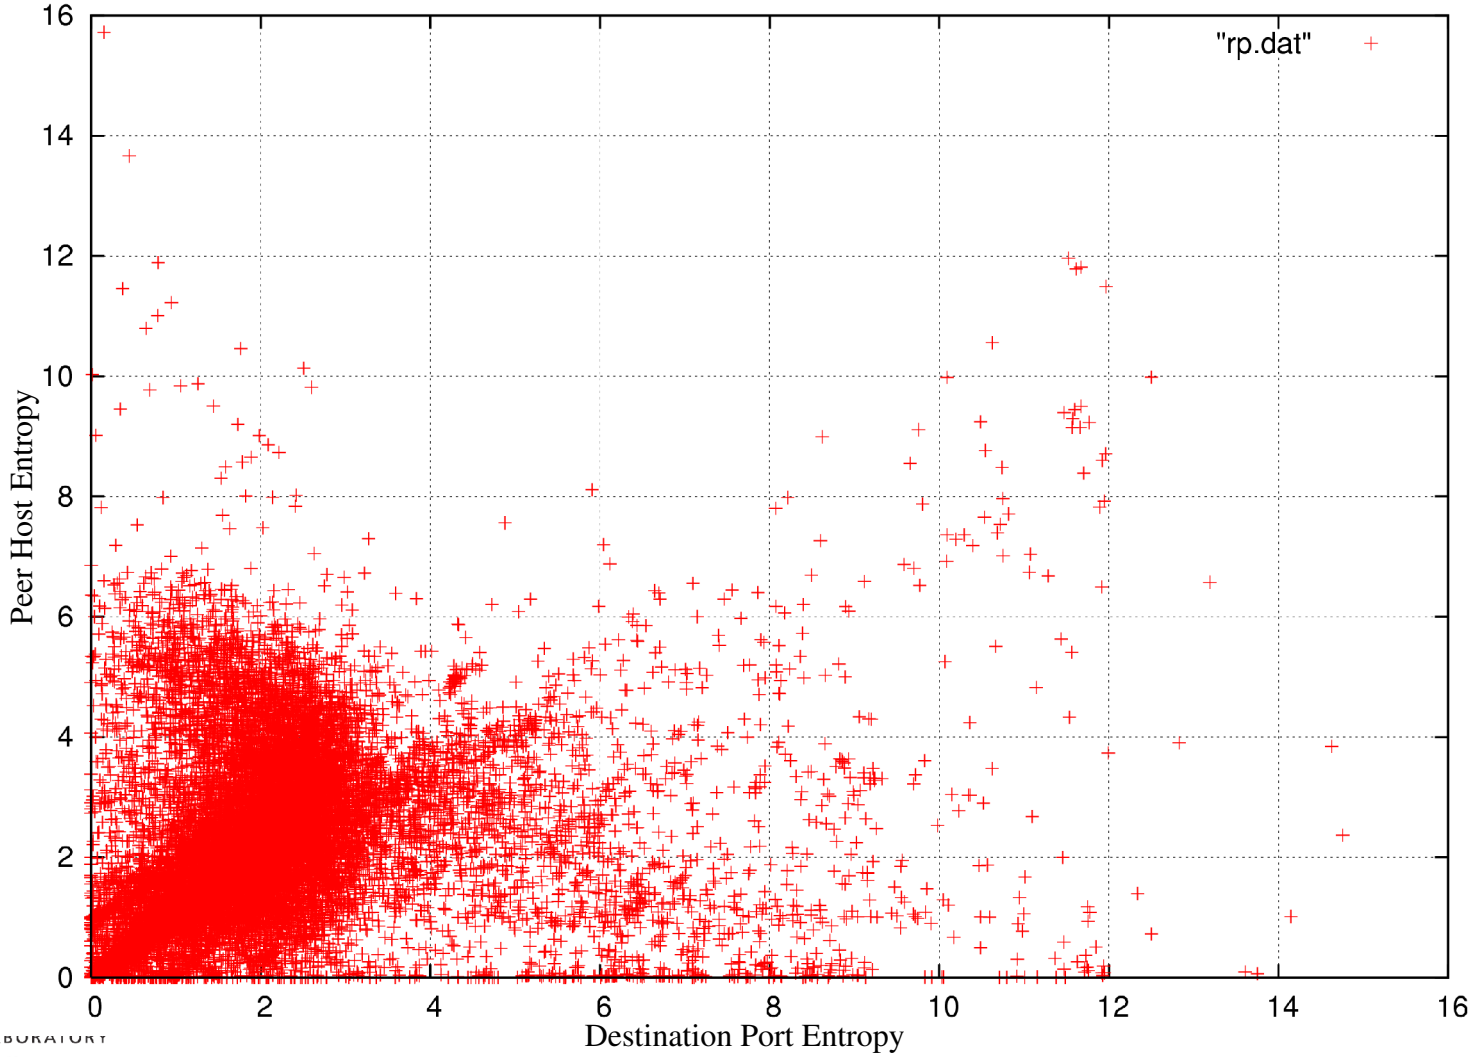


High Entropy Example

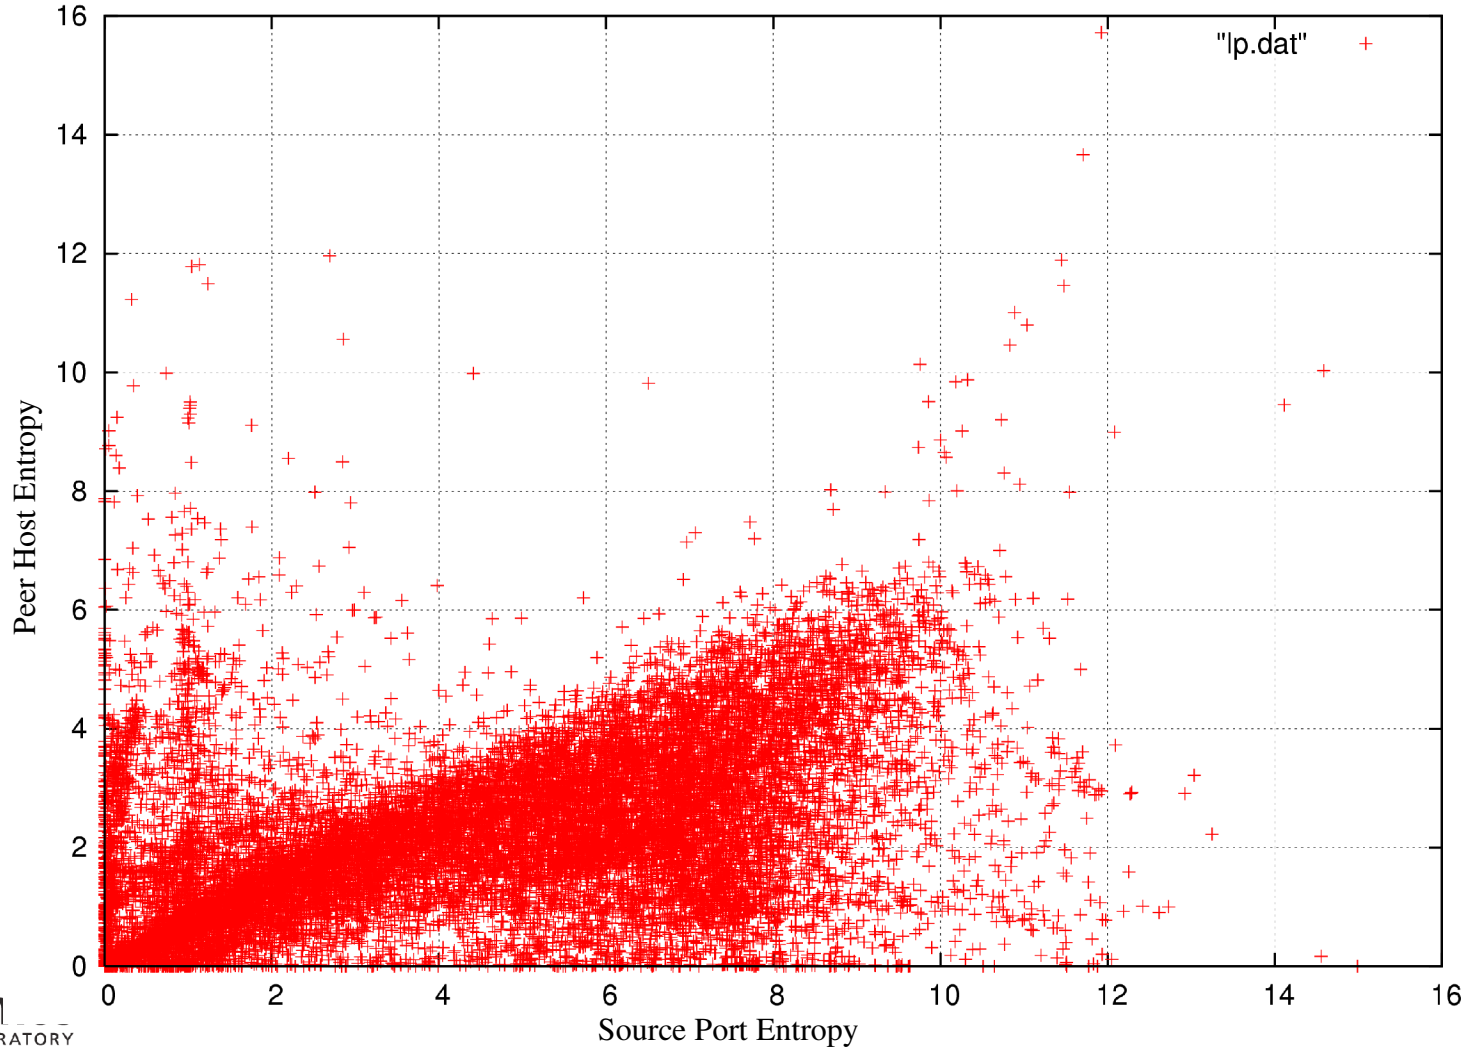
Data Overview

- Uses generic flow data
  - Required fields:
    - SRC IP, DST IP, SRC Port, DST Port, Protocol, Packets, Bytes
- Los Alamos unclassified network primarily over 24 hours (inclusive of a work day)
  - Approximately 200 million flows analyzed
  - 17,326 unique internal (Los Alamos) IP's observed
  - Day-to-day traffic very consistent

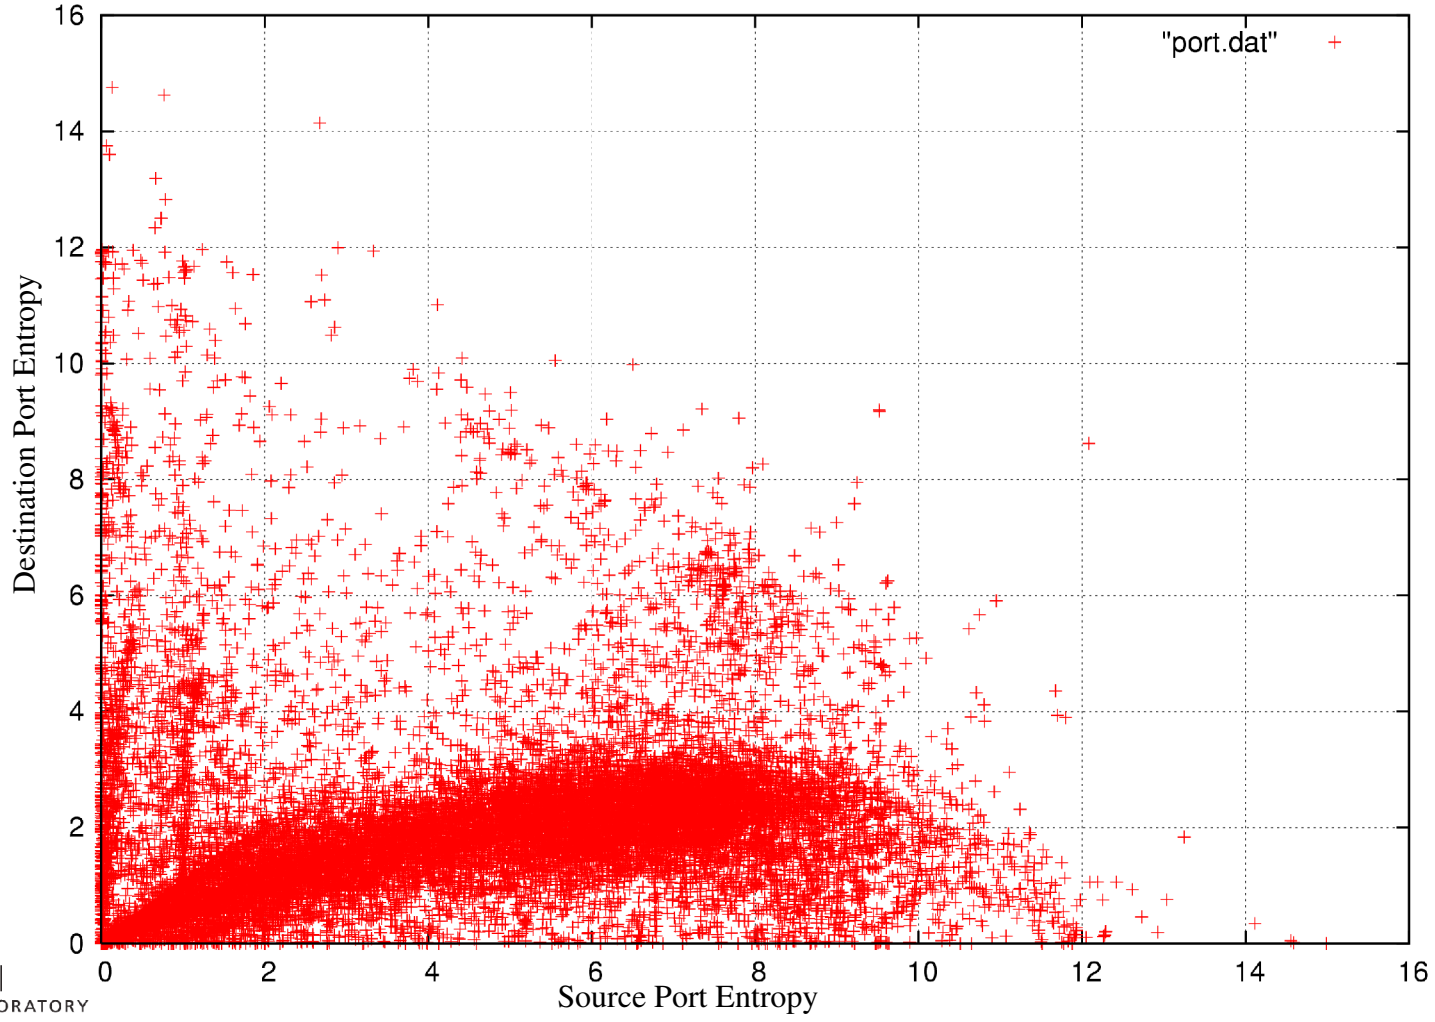
Destination Port / Peer Entropy



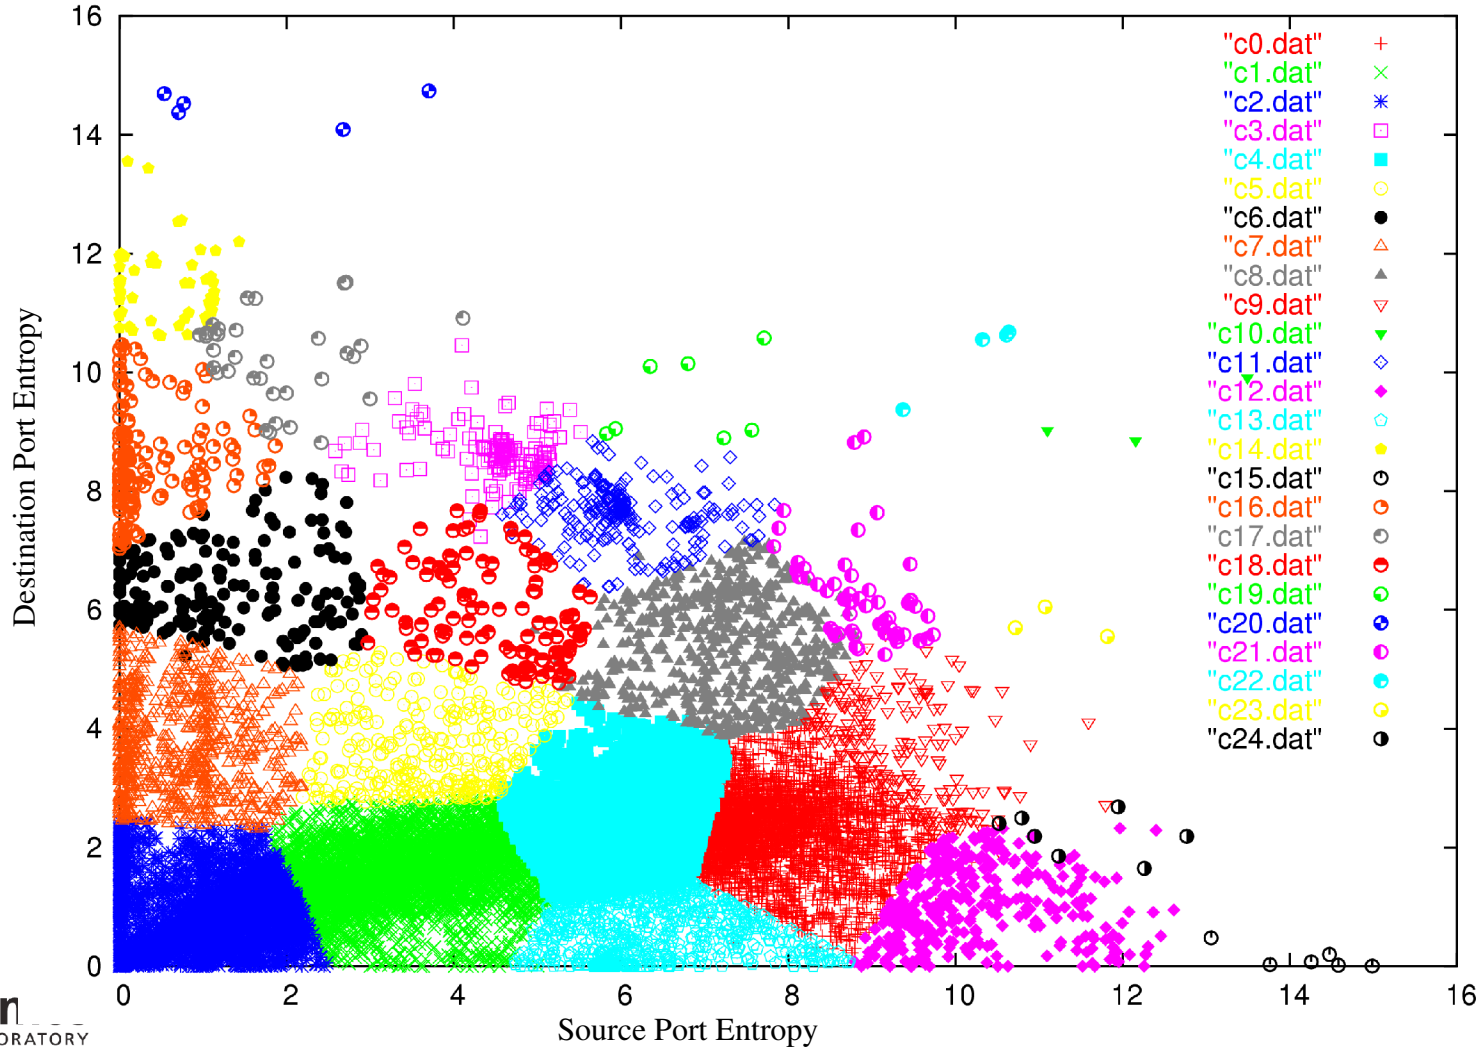
Source Port / Peer Entropy



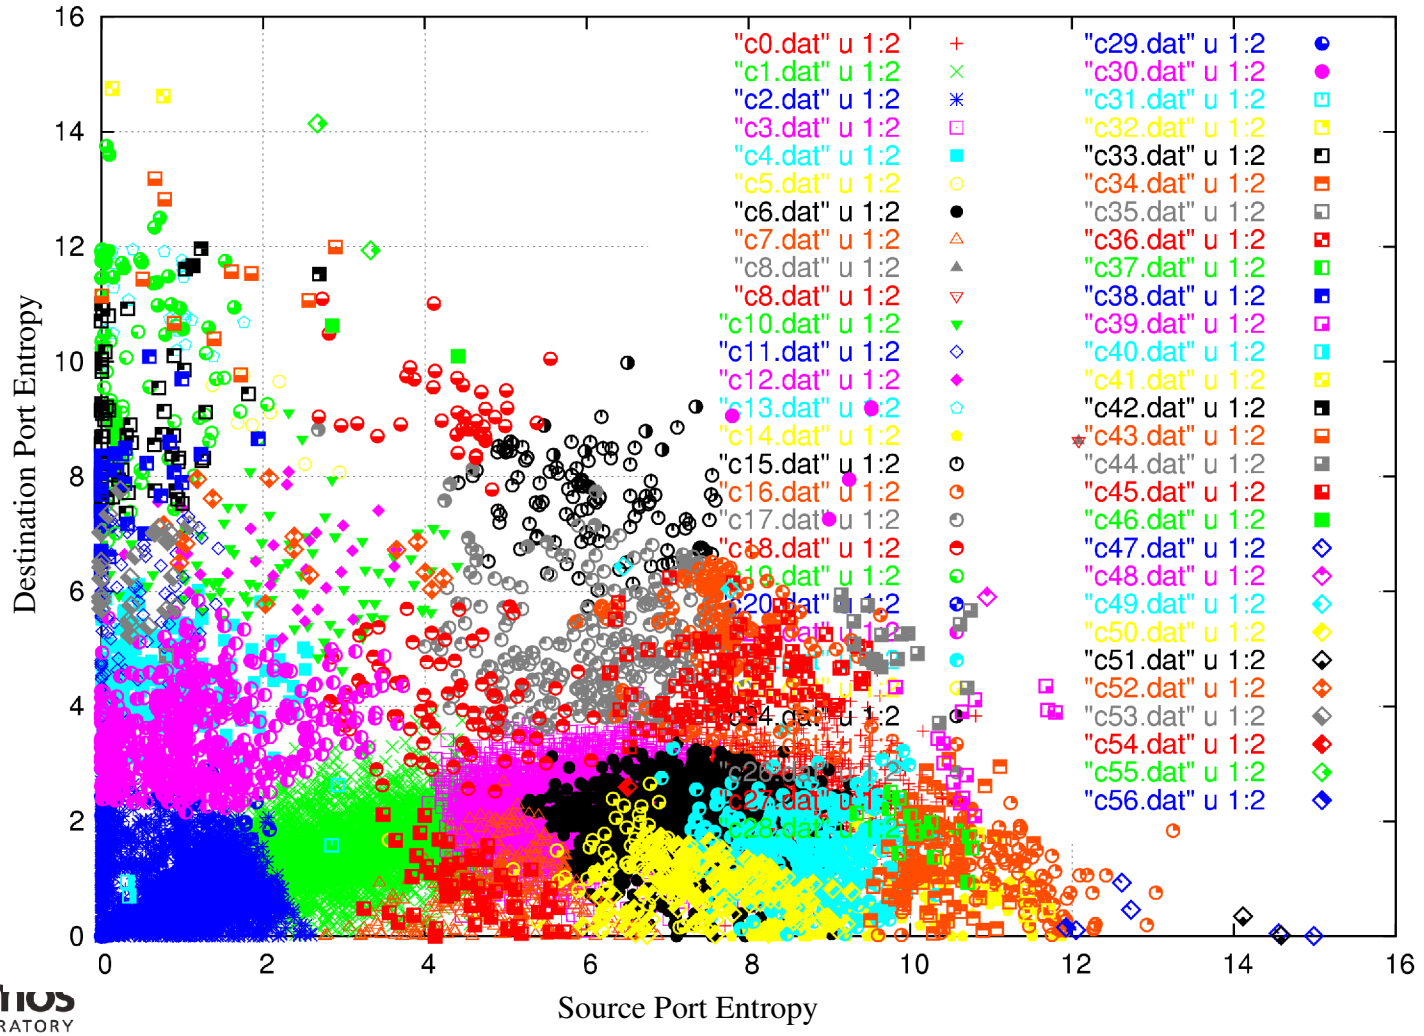
Source Port / Destination Port Entropy



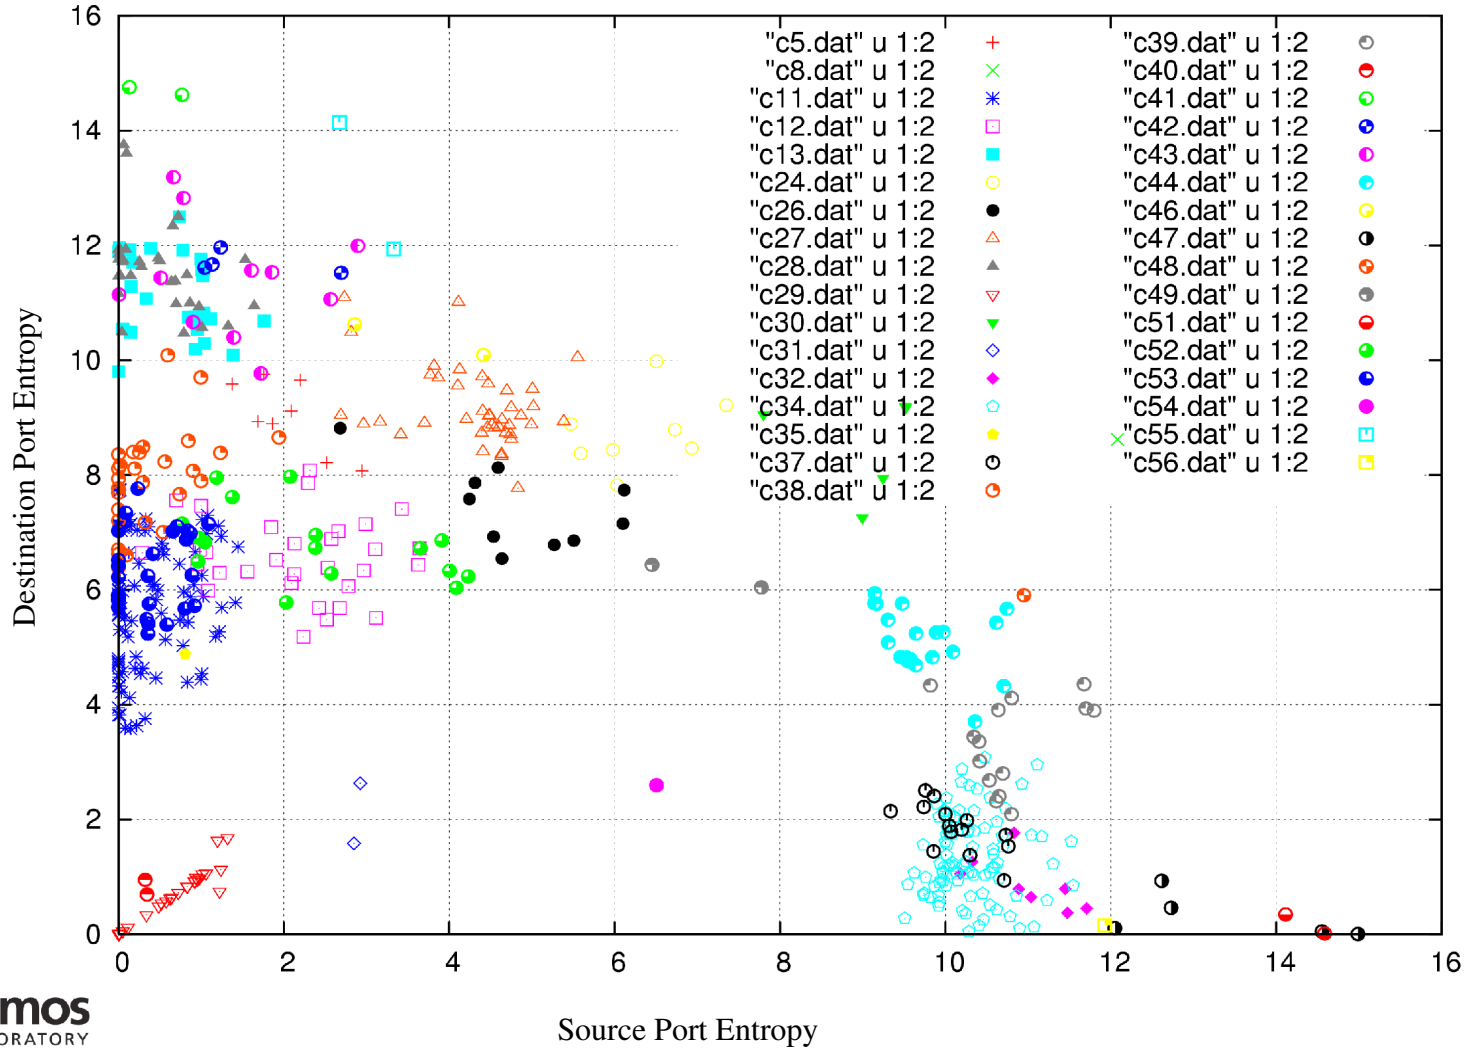
Source/Destination Port Entropy Clustering



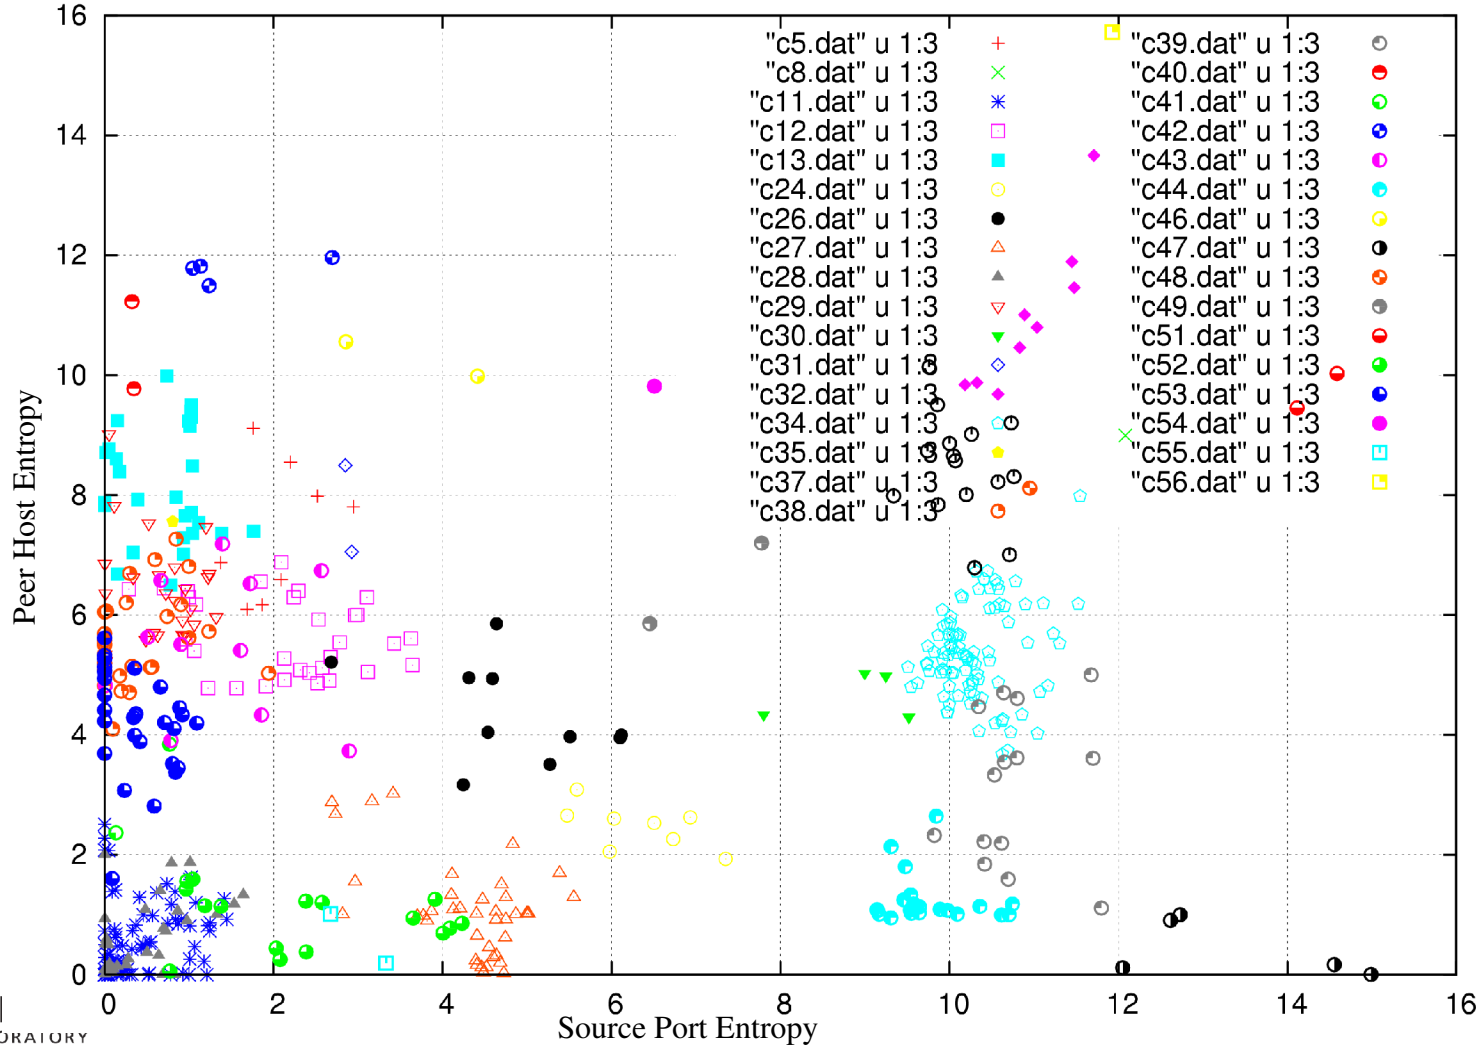
Entropy Clustering w/ 3 Dimensions



Interesting Clusters (<50)
749 Hosts (4.3% of total)



Interesting Clusters (b)
Source port versus Peers

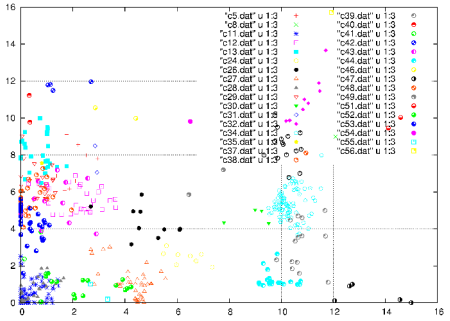


Major Servers

Host	Local Port Ent	Remote Port Ent	Peer Ent	Cluster
SMS 1	1.04	11.61	11.78	42
SMS 2	1.13	11.67	11.82	42
Int WWW	1.24	11.97	11.49	42
ActiveDir 1	2.7	11.52	11.96	42
ActiveDir 2	4.41	10.09	9.98	46
ActiveDir 3	2.86	10.62	10.56	46
DNS 1	1.76	9.76	9.11	5
DNS 2	0.74	12.5	9.99	13
VulnScanner	12.08	8.62	8.99	8
MailRelay 1	7.4	5.12	4.31	36
MailRelay 2	7.42	5.05	4.29	36
MailRelay 3	7.58	5.13	4.45	36

Clusters C32 & C56
Bad Behavior (worm variants)

Host IP	Local Port Ent	Remote Port Ent	Peer Ent
Remote Host A	11.45	0.79	11.89
Remote Host B	10.18	1.06	9.84
Internal Host A	10.32	1.26	9.87
Internal Host B	10.89	0.79	11.01
Remote Host C	10.83	1.77	10.46
Remote Host D	11.71	0.45	13.67
Remote Host E	11.04	0.65	10.8
Remote Host F	11.93	0.15	15.72



Current, On-going Work

- Demonstrated 1 million+ flows/minute processing on single system
  - Redesigning, porting system to map/reduce architecture for improved scaling and distributed processing
- Integrating additional network flow data types (e.g. custom perimeter collected flows)
- Static centroids for comparing host movements between k-means clusters
  - Enable predefined clusters, cluster definitions, and host movement between clusters
- Histogram merging that allows graceful data aging for continuous data feed and anomaly detection
- Application of novel change detection and machine learning across time series output

NTT

NTT Information Sharing Platform Laboratories

IPTV Traffic "Qcast": ~~IP Multicast Traffic~~ Monitoring System with IPFIX/PSAMP

Shingo Kashima and Atsushi Kobayashi

NTT Information Sharing Platform Laboratories

© 2009 NTT Information Sharing Platform Laboratories

NTT

Outline

- **Introduction**
  - Motivation
  - IP Multicast Streaming Traffic
- **Issues in Existing Multicast Monitoring**
- **Requirements**
  - Requirements
  - Difficult Requirements for Current NetFlow
  - Why IPFIX/PSAMP?
- **Our System: Qcast**
  - System Overview
  - System Architecture
  - System Evaluation
- **Conclusion**

© 2009 NTT Information Sharing Platform Laboratories

Motivation

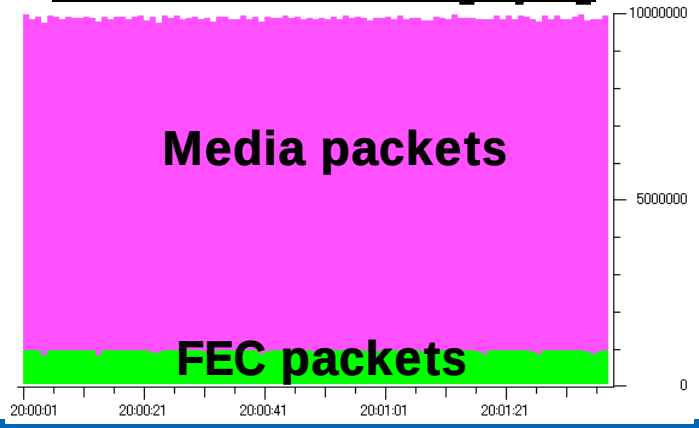
- **Multicast service has started in several provider networks.**
  - Large number of broadband users leads to heavy demand for IP multicast streaming services, such as IPTV.
- **Existing multicast tools work, but not well enough to monitor streaming services in large-scale networks.**
  - Multicast ping, trace route, and multicast MIB.
- **Easy troubleshooting tools are required.**
  - IPFIX/PSAMP seems helpful.

© 2009 NTT Information Sharing Platform Laboratories

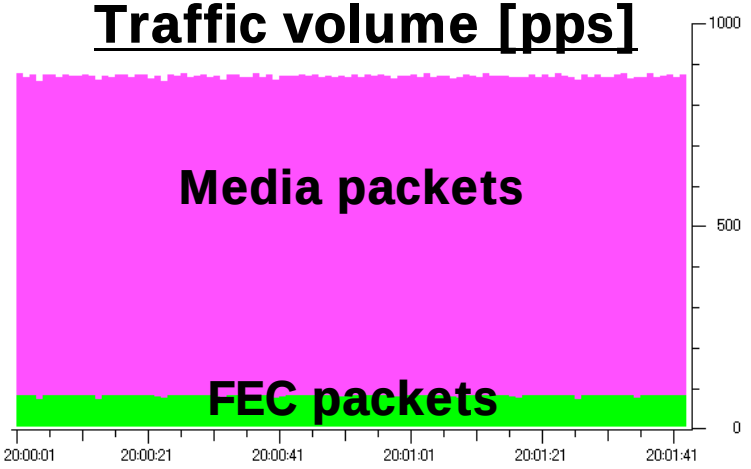
IP Multicast Streaming Traffic

- **Traffic volume of an IPTV channel: 10 Mbps, 0.9 kpps.**
  - Packet size: from 1300 to 1400 bytes.
- **IP multicast stream traffic includes two kinds of packets.**
  - Media packets and FEC packets.
- **IP multicast stream traffic includes RTP headers.**
  - Packet loss can be easily detected by keeping track of RTP seq. number.
- **More than 50 channels pass through an ISP network.**

Traffic volume [bps]



Traffic volume [pps]



© 2009 NTT Information Sharing Platform Laboratories

NTT

Issues in Existing Multicast Monitoring

- **Multicast ping and trace route**
  - Detect fault point and check continuity by using test packets.
  - Do not observe real packets.

  → Inadequate for detecting service quality deterioration and confirming service quality.

- **Mirroring + packet capture**
  - Last resort for confirming service quality.
  - But requires great care and is not suitable for always-on monitoring.

© 2009 NTT Information Sharing Platform Laboratories

Requirements

- **Requirement #1: Detect service quality deterioration and confirm service quality.**
  - Detect packet loss, disorder, and duplicates within 1 minute while there is continuous packet loss at 1/1000.
  - Monitor packet delay variation.
- **Requirement #2: Perform always-on monitoring of traffic volume and service quality of each IPTV-channel and each customer.**
  - Always monitor per {S,G}.
  - Always monitor per VLAN in access network.
- **Requirement #3: Localize failure point.**
  - Localize failure point not only in a service suspension but also when service quality deterioration occurs.
- **Requirement #4: Use at low cost as soon as possible.**
  - Necessary because multicast streaming service has already started.

© 2009 NTT Information Sharing Platform Laboratories

NTT

Difficult Requirements for Current NetFlow

- **Requirement #1: For the current NetFlow exporter implementation, it is impractical to detect service quality deterioration and monitor service quality.**
  - Flow records in NetFlow cannot include packet loss, disorder, and duplicates.
  - In general, many operators use random sampling to introduce NetFlow.

- **IPFIX/PSAMP seems helpful in meeting this requirement.**

© 2009 NTT Information Sharing Platform Laboratories

Why IPFIX/PSAMP?

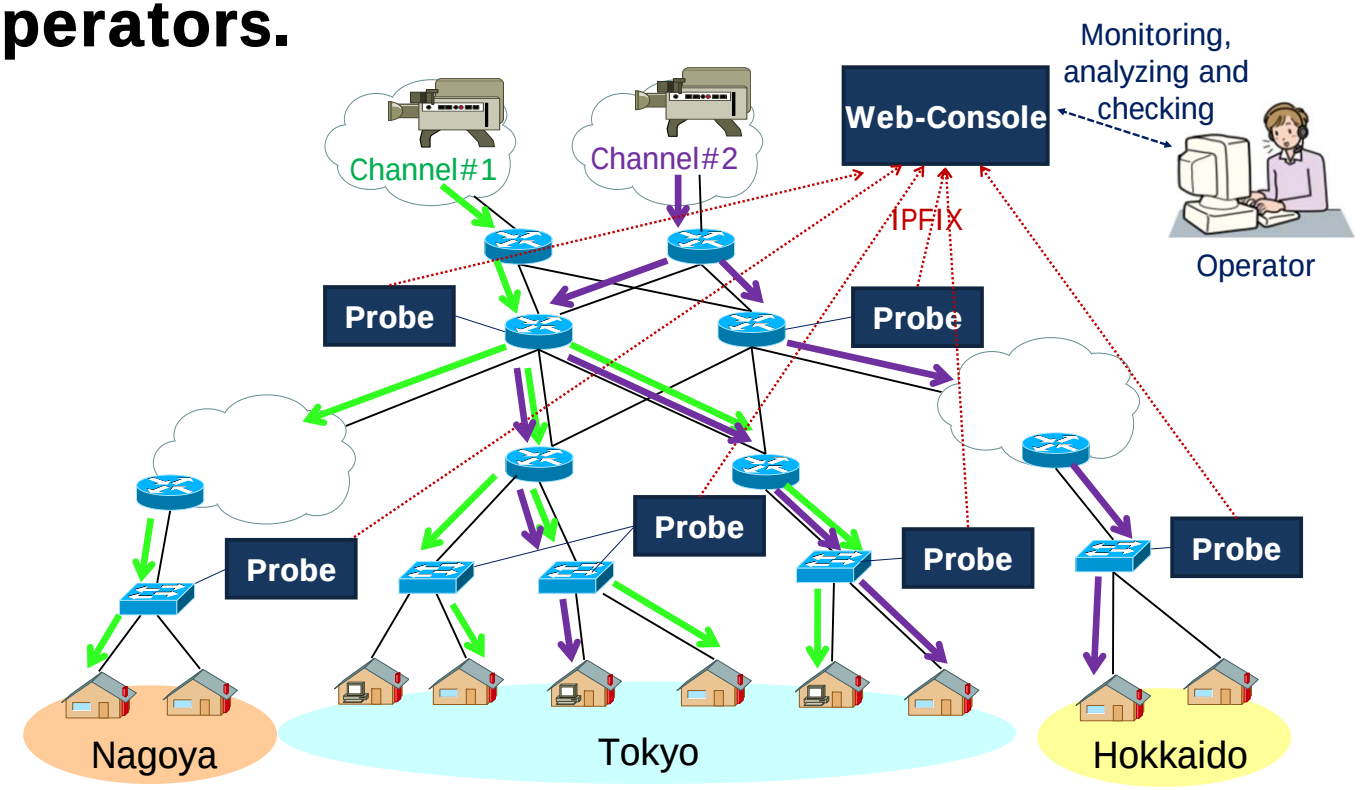
- **PSAMP (RFC 5475)**
  - "Property Match Filtering" can focus the monitoring on IPTV traffic by selecting on the basis of packet header value.
  - "Systematic Time-based Sampling" can detect packet loss and packet interval time by selecting continuous packets.

- **IPFIX (RFC 5153)**
  - "Enterprise-specific Information Elements" can export not only UDP/IP header information but also application header information.

© 2009 NTT Information Sharing Platform Laboratories

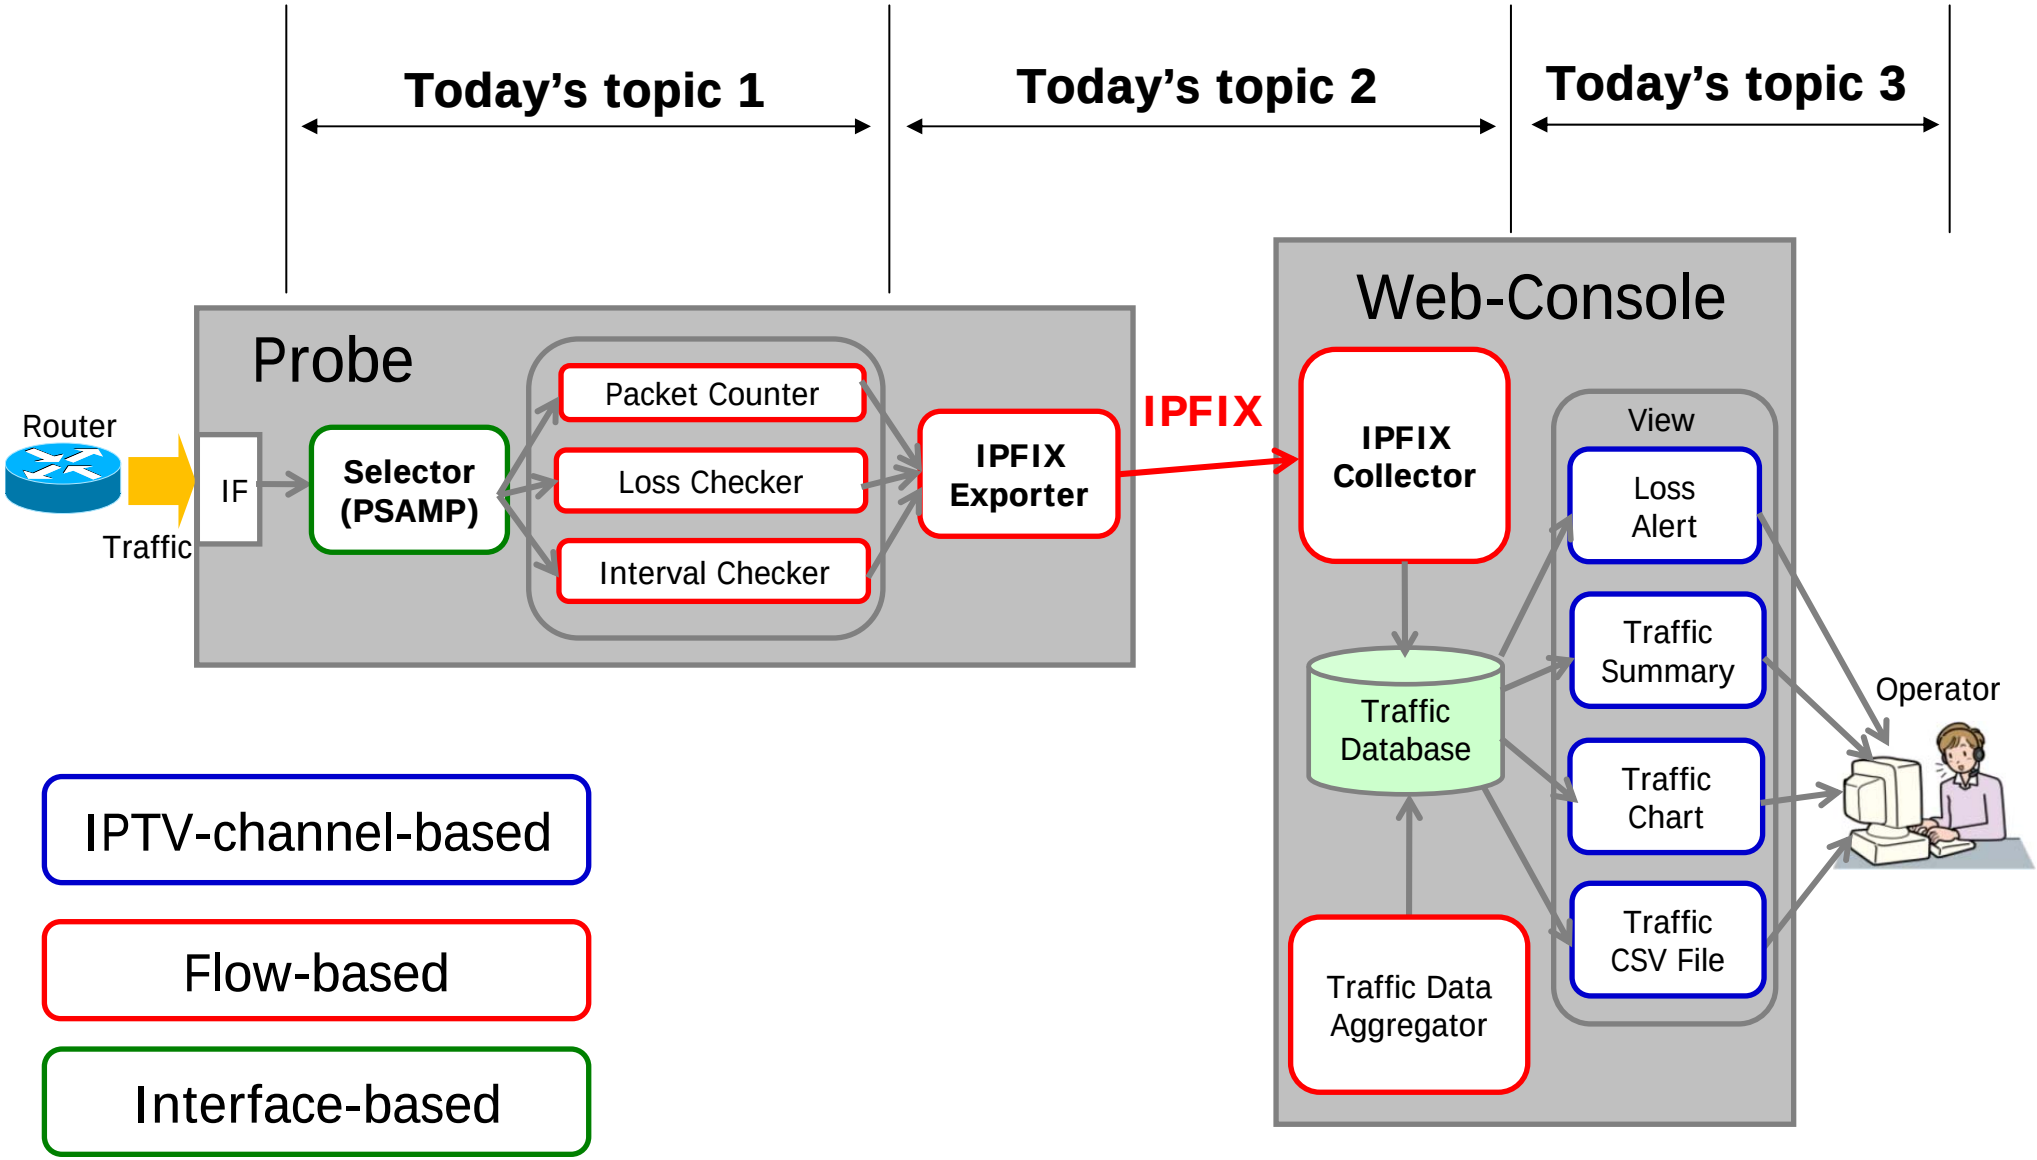
Our System: Qcast Overview

- **Probe: Captures traffic from mirror port, monitors IPTV traffic with PSAMP, and exports with IPFIX.**
  - Runs on a general-purpose personal computer.
- **Web-Console: Collects IPFIX information and shows it to operators.**



© 2009 NTT Information Sharing Platform Laboratories

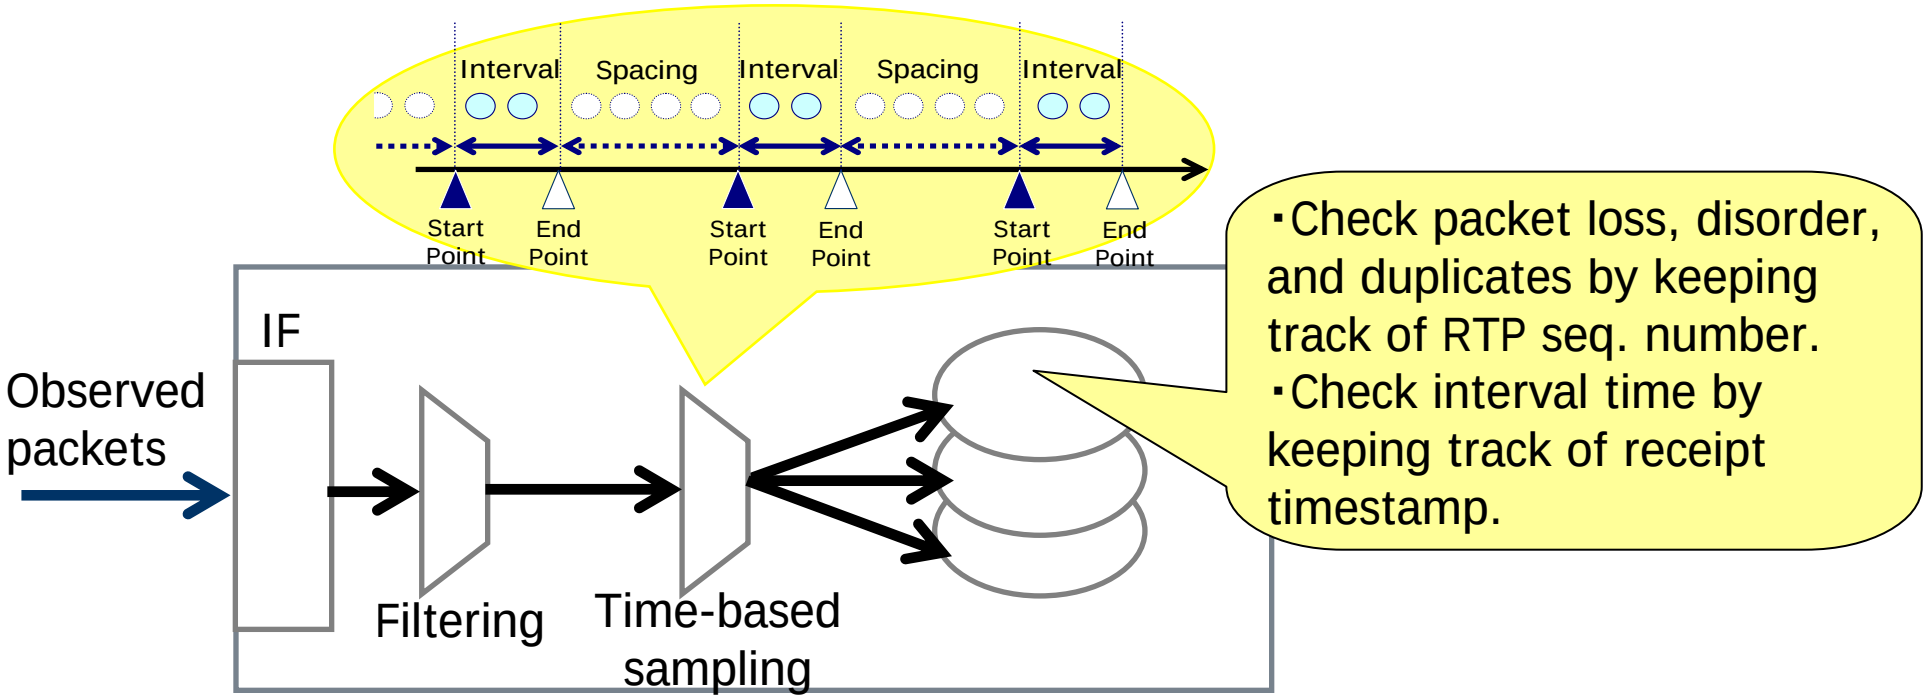
Qcast Architecture



© 2009 NTT Information Sharing Platform Laboratories

Topic 1: Packet Loss and Interval Time

- **Combination of PSAMP techniques in Probe**
  - Observe packets at input interface.
  - Select multicast packet by “Property Match Filtering”.
    - Example: “Destination IPv6 Address == FF38::/16”
  - Extract them using “Systematic Time-based Sampling”.
    - All input packets during the interval period are selected.



© 2009 NTT Information Sharing Platform Laboratories

Topic 2: Exporting Traffic Data

- **IPv6 template**
  - Flow key information
    - Includes VLAN ID
  - Traffic volume information
  - Packet loss information
  - Interval time information
    - Uses “Enterprise-specific Information Elements”

	0 1 2 3 4 5 6 7 0 1 2 3 4 5 6 7	0 1 2 3 4 5 6 7 0 1 2 3 4 5 6 7
0	Set ID (0x0002)	Length
4	Template ID (0x0103)	Field Count = 18 (0x0012)
8	octetDeltaCount (0x0001)	Field Length (0x0004)
12	packetDeltaCount (0x0002)	Field Length (0x0004)
16	srcTransportPort (0x0007)	Field Length (0x0002)
20	dstTransportPort (0x000B)	Field Length (0x0002)
24	srcIPv6Address (0x001B)	Field Length (0x0010)
28	dstIPv6Address (0x001C)	Field Length (0x0010)
32	droppedPacketDeltaCount (0x0085)	Field Length (0x0004)
36	flowStartMilliseconds (0x0098)	Field Length (0x0004)
40	flowEndMilliseconds (0x0099)	Field Length (0x0004)
44	vlanId (0x003A)	Field Length (0x0002)
48	protocolIdentifier (0x0004)	Field Length (0x0001)
52	ipVersion (0x003C)	Field Length (0x0001)
56	rtpIntervalAvgTime (0x8001)	Field Length (0x0004)
60	ENTERPRISE NUMBER (0x000000D2)	
64	rtpInternalMaxTime (0x8002)	Field Length (0x0004)
68	ENTERPRISE NUMBER (0x000000D2)	
72	rtpIntervalMinTime (0x8003)	Field Length (0x0004)
76	ENTERPRISE NUMBER (0x000000D2)	
80	rtpIntervalVariance (0x8004)	Field Length (0x0004)
84	ENTERPRISE NUMBER (0x000000D2)	
88	rtpPayloadType (0x8005)	Field Length (0x0002)
92	ENTERPRISE NUMBER (0x000000D2)	

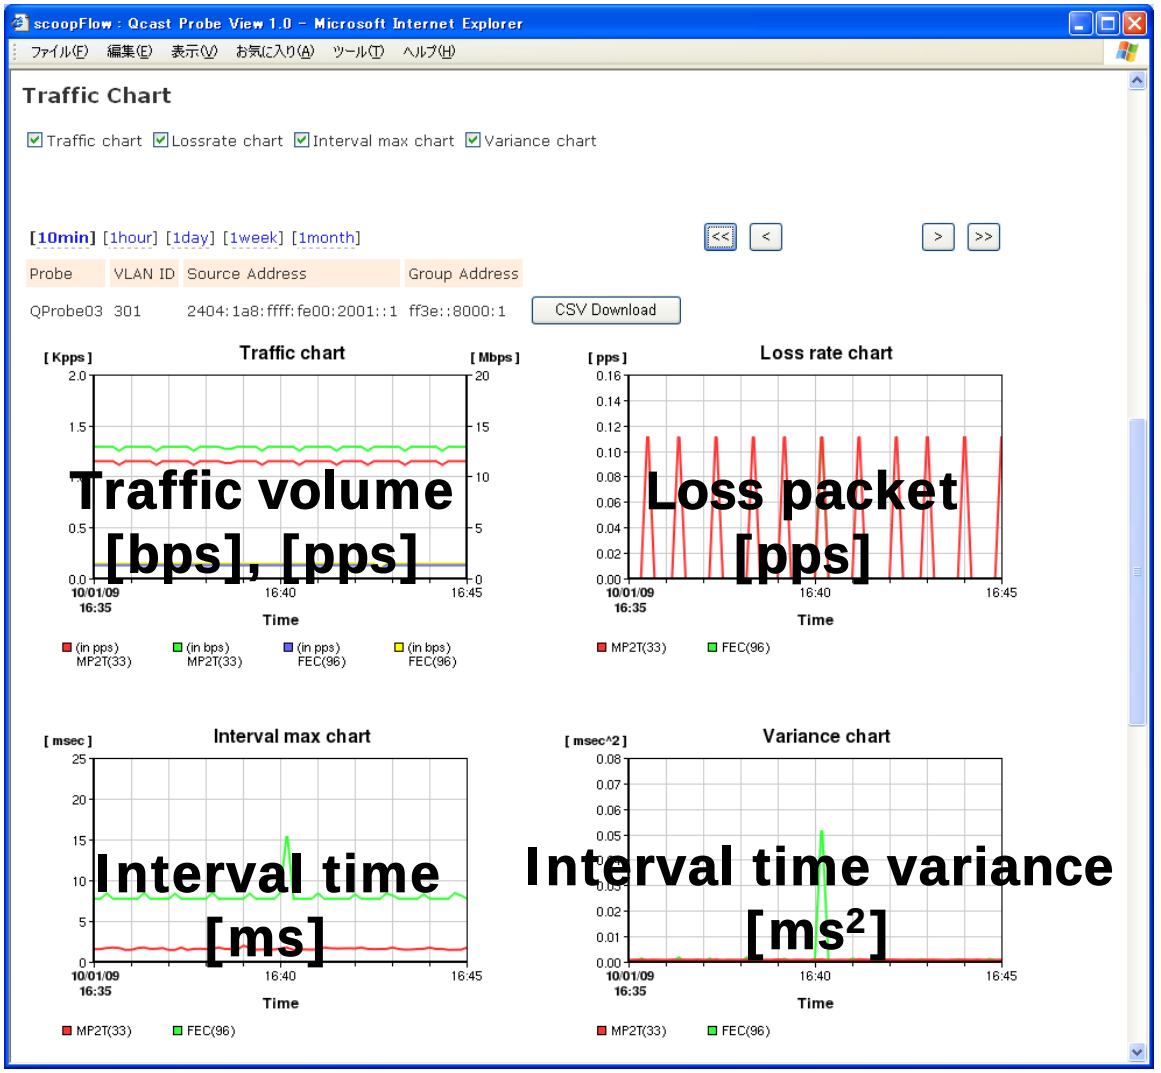
- **Option template**

	0 1 2 3 4 5 6 7 0 1 2 3 4 5 6 7	0 1 2 3 4 5 6 7 0 1 2 3 4 5 6 7
0	Set ID (0x0003)	Length
4	Template ID (0x0106)	Field Count (0x0003)
8	Scope Field Count (0x0001)	exporterIPv4Address (0x0082)
12	Scope 1 Length (0x0004)	samplingTimeInterval (0x0133)
16	Field Length (0x0004)	samplingTimeSpace (0x012E)
20	Field Length (0x0004)	flowActiveTimeout (0x0024)
24	Field Length (0x0002)	

© 2009 NTT Information Sharing Platform Laboratories

Topic 3: View [1/6]

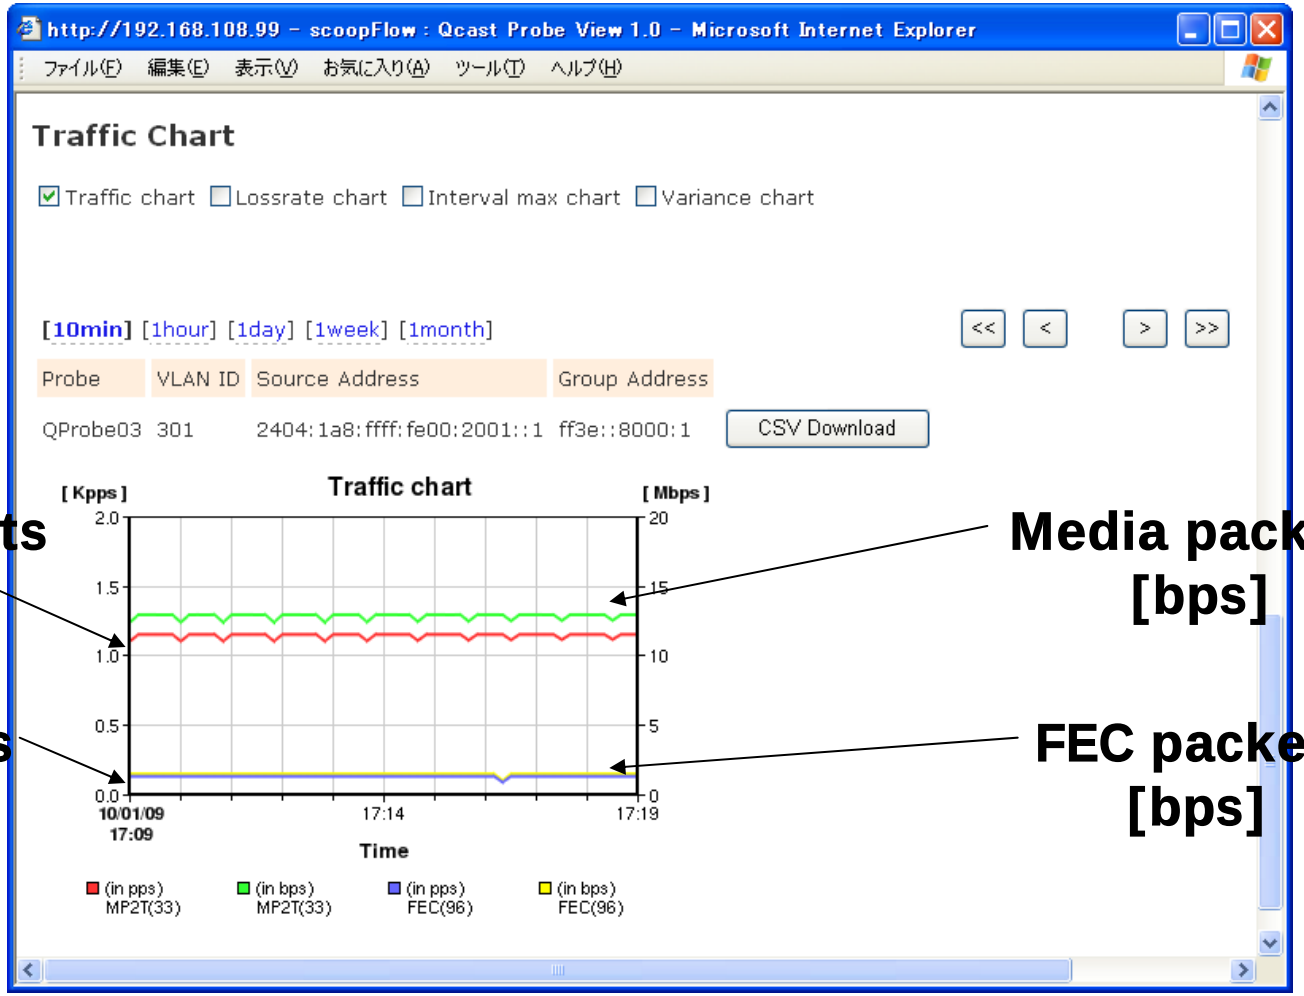
- **Shows four kinds of traffic information**



© 2009 NTT Information Sharing Platform Laboratories

Topic 3: View [2/6]

- **Shows a traffic chart separated by RTP payload type.**
  - {Media packets, FEC packets}



Media packets [pps]

FEC packets [pps]

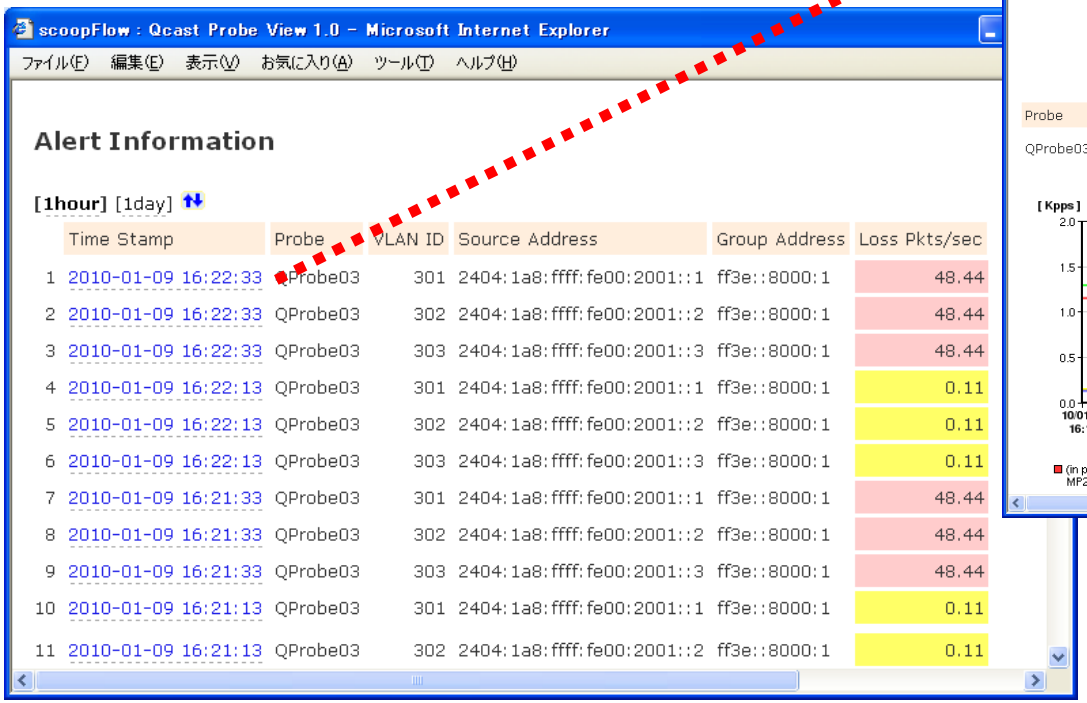
Media packets [bps]

FEC packets [bps]

© 2009 NTT Information Sharing Platform Laboratories

Topic 3: View [3/6]

- **Shows packet loss alert**
  - Two-level alert
    - Red: dangerous level
    - Yellow: warning level
  - Link to traffic chart



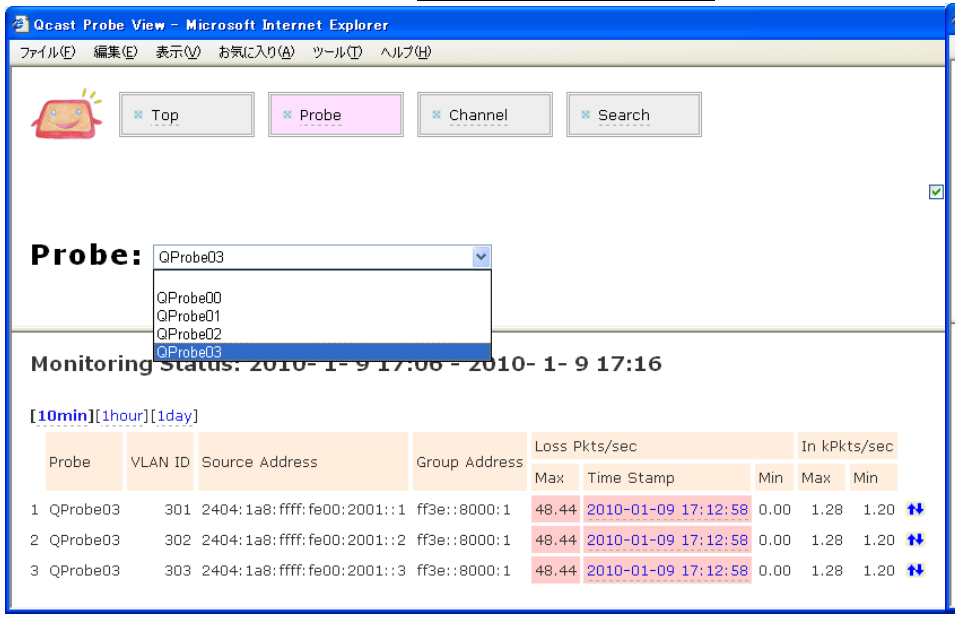


© 2009 NTT Information Sharing Platform Laboratories

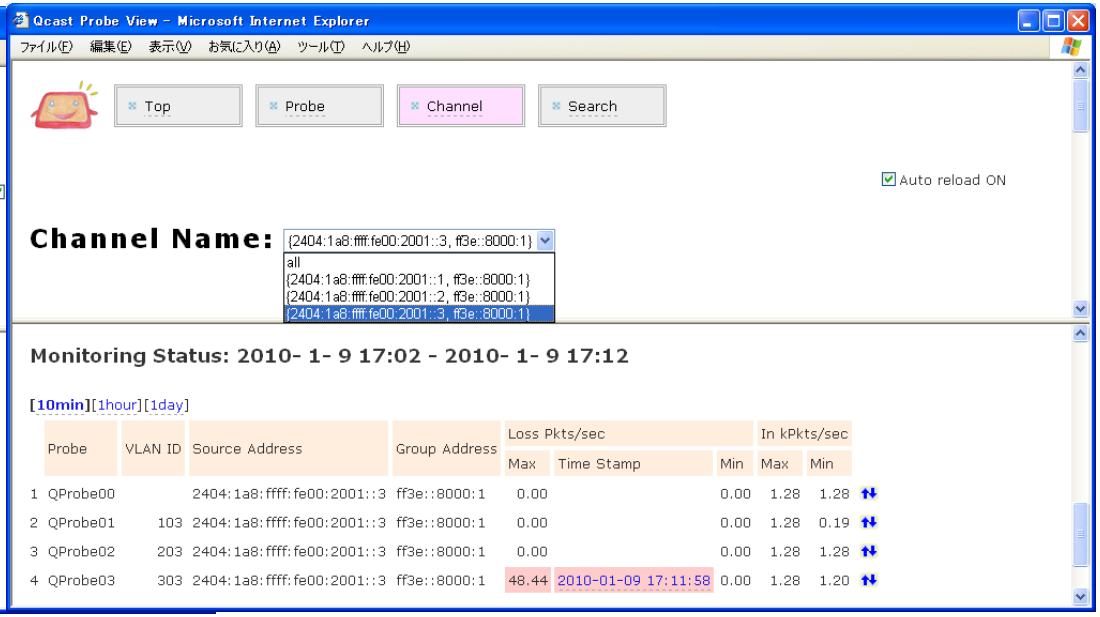
Topic 3: View [4/6]

- **Probe view**
  - When you select a probe from the Probe pull-down menu, the specified interface information is shown.
- **Channel view**
  - When you select {S,G} from the Channel Name pull-down menu, the specified IPTV-channel information is shown.

Probe view



Channel view



© 2009 NTT Information Sharing Platform Laboratories

Topic 3: View [5/6]

- **Probe view**
  - Also shows traffic volume chart built up from all channels.
  - Useful for capacity planning.

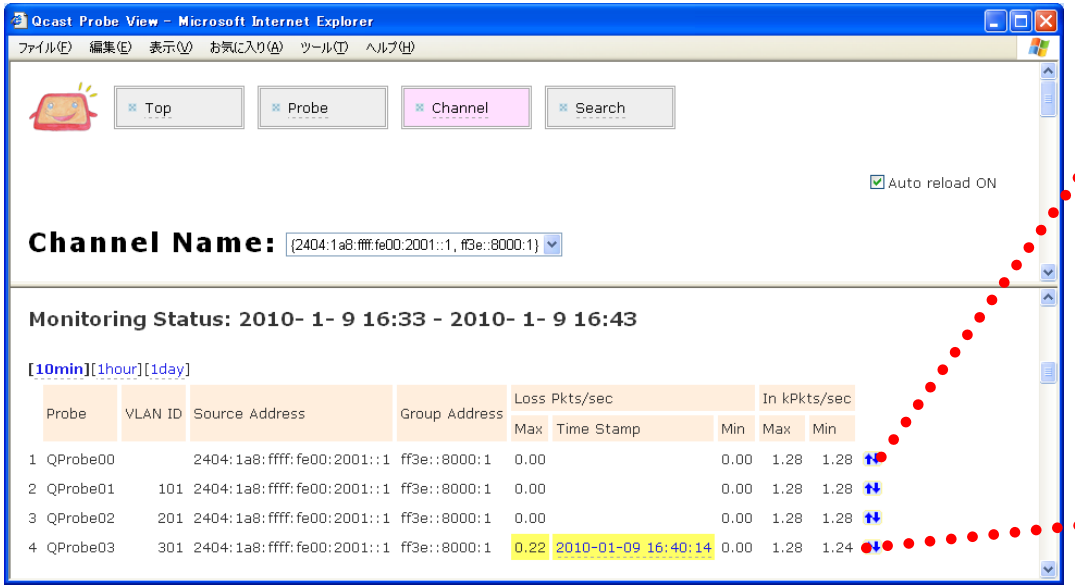


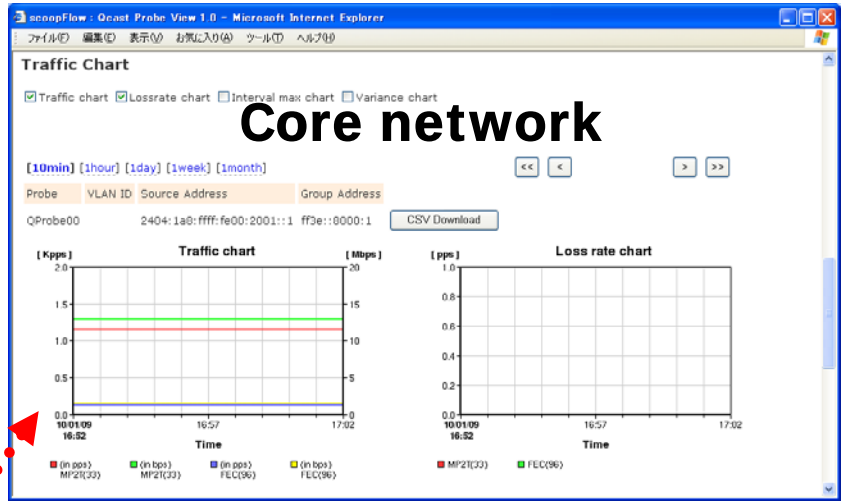
{2404:1a8:ffff:fe00:2001::3, ff3e::8000:1}
{2404:1a8:ffff:fe00:2001::2, ff3e::8000:1}
{2404:1a8:ffff:fe00:2001::1, ff3e::8000:1}

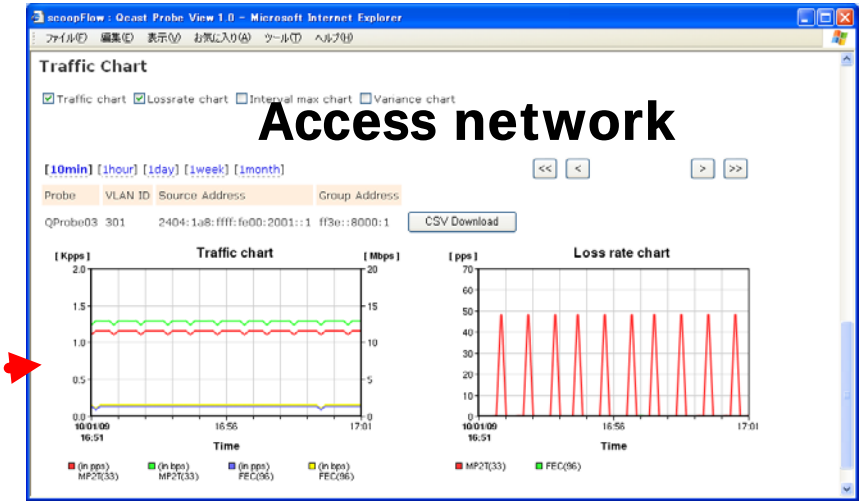
© 2009 NTT Information Sharing Platform Laboratories

Topic 3: View [6/6]

- **Channel view**
  - Also localizes fault point by comparing traffic charts of multiple probes.



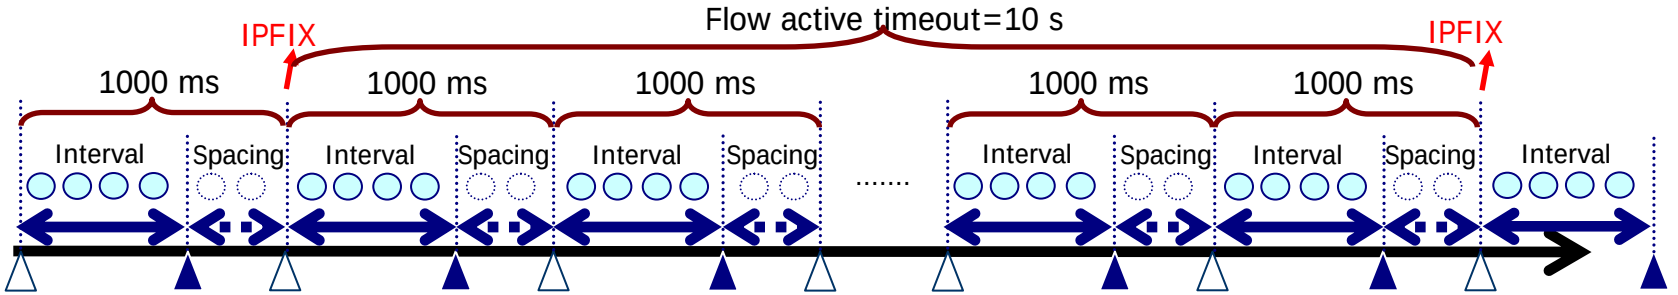




© 2009 NTT Information Sharing Platform Laboratories

Evaluation of Probe

- **Experimental assumptions:**
  - Packets of 20—100 IPTV channels pass through a probe.
    - Traffic volume of an IPTV channel: 10 Mbps, 0.9 kpps.
    - IPTV channel includes two kinds of packets (Media, FEC).
  - We evaluated the probe by varying the sampling interval period.
    - Sum of the sampling and spacing interval period was kept at a fixed value of 1000 ms.
    - Flow active timeout had a fixed value of 10 s.



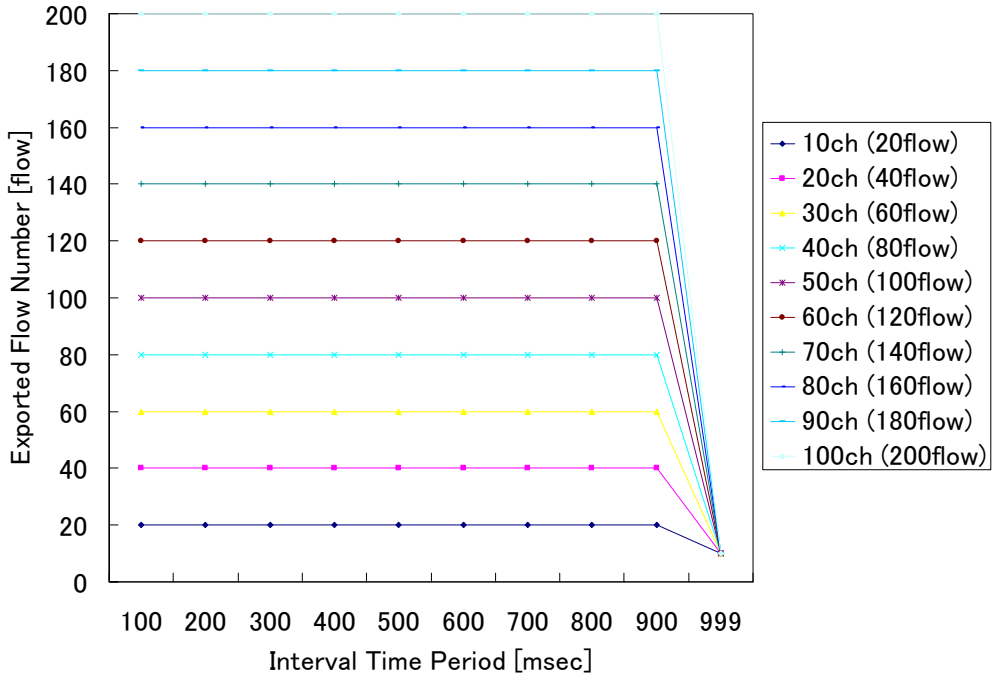
- **Experimental environment:**
  - CPU: Core 2 Duo 2.0GHz, Memory: 3.5 GB, Interface: Intel 1GbE NIC, OS: CentOS 5.3

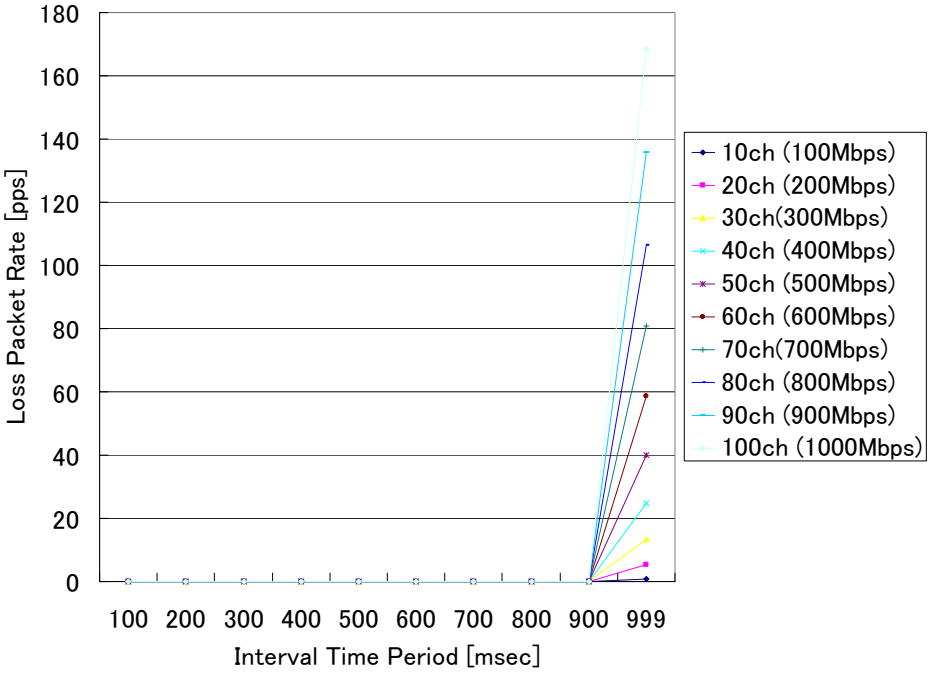


© 2009 NTT Information Sharing Platform Laboratories

Experimental Results for Probe

- **Exported flow number and packet loss rate obtained by changing the interval period from 100 to 999 ms.**
- **Performance had no limitation for interval period from 100 to 900 ms.**



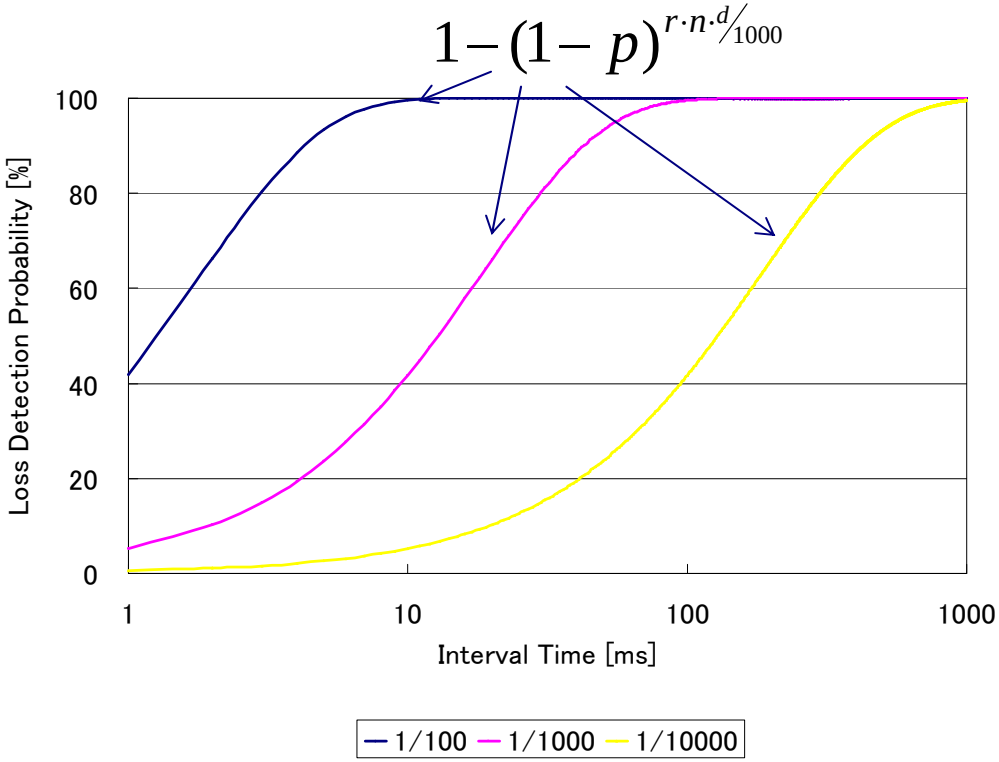


© 2009 NTT Information Sharing Platform Laboratories

Packet Loss Detection Probability

- **We evaluated the detection probability for packet loss within the given monitoring interval (*n*) by changing packet loss rate (*p*) and sampling interval time (*d*).**
  - On the condition that *n* is 1 min and packet rate (*r*) is 0.9 kpps, the experimental results are shown.

- **Detection probability is almost 100% if the sampling interval period is 900 ms and packet loss is 1/1000 for over 1 min.**

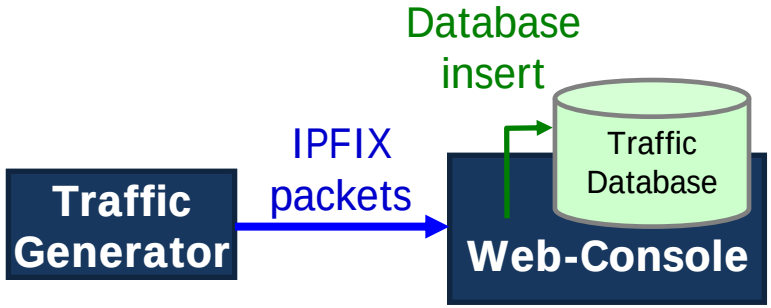


© 2009 NTT Information Sharing Platform Laboratories

Evaluation of Web-Console

- **Experimental assumptions:**
  - IPFIX packets were input to a Web-Console.
    - A IPFIX packet includes 10 flow records.
  - We evaluated the Web-Console by varying the flow rate and number of flow.
    - Flow rate: 10-200 flow/sec (1-20 pps)
    - No. of flow: 200-1000 flow (20-100 packets)

- **Experimental environment:**
  - CPU: Xeon 3.6Hz, Memory: 3.5 GB,
    OS: CentOS 5.3, DB: PostgresSQL 8.37

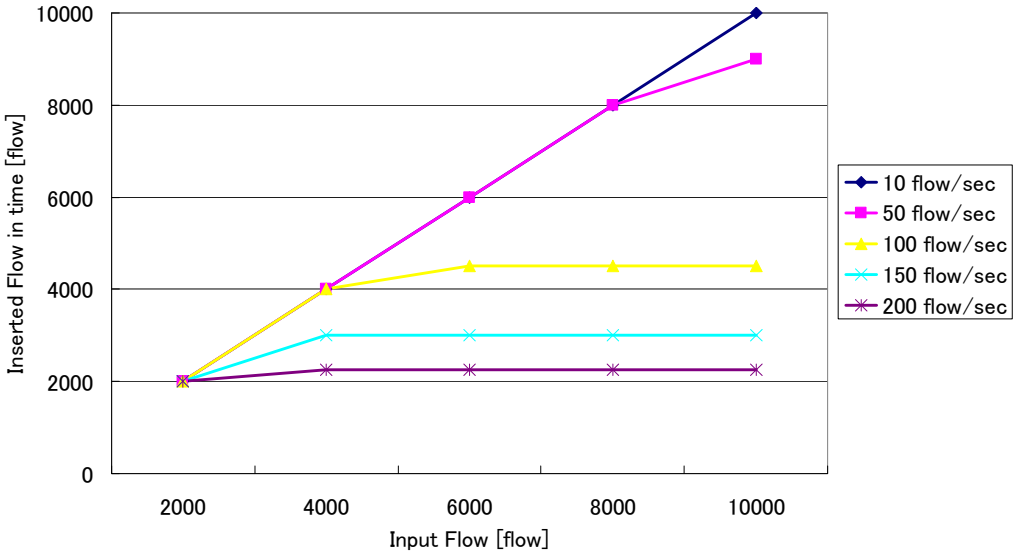


© 2009 NTT Information Sharing Platform Laboratories

Experimental Results for Web-Console

- **We measure the number of inserted flow into traffic database in time by changing the following conditions:**
  - No. of input flow from 2,000 to 10,000
  - Input flow rate from 10 to 200 flow/sec.

- **The performance limit seems to be as follows:**
  - 50 flow/sec: 9,000 flow
  - 100 flow/sec: 4,500 flow
  - 150 flow/sec: 4,000 flow
  - 200 flow/sec: 2,200 flow



- **That is to say:**
  - When flow active timeout is 180 sec, 90 probe (9,000 flow) is supported.
  - When flow active timeout is 45 sec, 45 probe (4,500 flow) is supported.
  - When flow active timeout is 26 sec, 40 probe (4,000 flow) is supported.
  - When flow active timeout is 11 sec, 22 probe (2,200 flow) is supported.

© 2009 NTT Information Sharing Platform Laboratories

Summary

- **We presented a new traffic monitoring method for IP multicast streaming services, such as IPTV, and the implemented system using IPFIX/PSAMP (Qcast).**

- **We showed the feasibility of the Qcast.**

© 2009 NTT Information Sharing Platform Laboratories

NTT

Thank you very much.

Please come and see our demonstration.

This study was supported by the Ministry of Internal Affairs and Communications of Japan.

© 2009 NTT Information Sharing Platform Laboratories

A Case Study – Using Flow to Identify Specific Malware Characteristics



United States Computer Emergency Readiness Team (US-CERT)

Office of Analysis

Network Analysis Branch

January 12, 2010



Homeland Security

Caveats

- Portions of the malware will not be briefed
- All IPs have been changed
  - 192.168.x.x – ‘Good Guys’
  - 10.1.1.x, 5.1.1.x, 2.1.1.x – ‘Bad Guys’ from reported traffic
  - 172.16.2.x, 14.1.2.x, 2.1.2.x – ‘Bad Guys’ from discovered traffic

Homeland Security

Background

- In the early part of 2009, a set of malware was received from multiple United States Government Agencies (USGA's)
- US-CERT analyzed malware in an attempt to successfully identify its behavior within flow
- Utilizing SiLK and PERL, it was possible to narrow an initial large data set to just a few suspicious IP addresses

Initial Data Pull

sIP	dIP	sPort	dPort	pro	packets	sTime	bytes	flags	dur	type	initialF
10.1.1.126	192.168.1.19	1088	2496	6	46	2009/06/01T01:28:02.448	51015	FS PA	5.367	IN	S A
10.1.1.126	192.168.1.19	443	2493	6	3	2009/06/01T01:23:14.775	129	SRPA	0.001	INWEB	S A
10.1.1.126	192.168.2.6	8080	56105	6	4	2009/06/01T01:29:35.274	168	FS A	30.981	INWEB	S A
192.168.1.19	10.1.1.126	2493	443	6	3	2009/06/01T01:23:14.774	176	S PA	0.001	OUTWEB	S
192.168.2.6	10.1.1.126	56105	8080	6	5	2009/06/01T01:30:04.602	312	S PA	30.85	OUTWEB	S
192.168.2.6	10.1.1.126	56120	8088	6	1	2009/06/01T03:35:28.723	40	R	0	OUT	
192.168.2.6	10.1.1.126	56192	8088	6	1	2009/06/01T03:42:20.452	40	R	0	OUT	
10.1.1.126	192.168.3.228	8088	3833	6	4	2009/06/01T08:29:37.230	168	FS A	1.969	IN	S A
10.1.1.126	192.168.3.228	8088	3850	6	19	2009/06/01T08:30:32.540	783	S PA	18.024	IN	S A
10.1.1.126	192.168.3.228	1088	3917	6	46	2009/06/01T08:43:56.022	51015	FS PA	2.86	IN	S A
10.1.1.126	192.168.3.228	8088	3850	6	193	2009/06/01T08:34:13.861	8552	PA	784.708	IN	PA
10.1.1.126	192.168.3.228	8088	3850	6	3	2009/06/01T08:49:10.106	125	F PA	0.11	IN	PA
192.168.3.228	10.1.1.126	3833	8088	6	5	2009/06/01T08:29:07.170	309	S PA	1.93	OUT	S
192.168.3.228	10.1.1.126	3850	8088	6	27	2009/06/01T08:30:02.393	7404	S PA	17.866	OUT	S
192.168.3.228	10.1.1.126	3917	1088	6	28	2009/06/01T08:43:25.914	1176	FS PA	2.774	OUT	S
192.168.3.228	10.1.1.126	3850	8088	6	172	2009/06/01T08:33:43.839	27445	PA	784.601	OUT	PA
192.168.3.228	10.1.1.126	3850	8088	6	3	2009/06/01T08:48:40.018	125	F PA	0.109	OUT	PA
192.168.3.228	10.1.1.126	3833	8088	6	1	2009/06/01T09:29:16.300	40	R	0	OUT	
10.1.1.126	192.168.2.6	8088	56119	6	8	2009/06/01T01:31:01.975	332	FS PA	9.32	IN	S A
10.1.1.126	192.168.2.6	8088	56120	6	26	2009/06/01T01:31:23.275	1182	S PA	110.45	IN	S A
10.1.1.126	192.168.2.6	8088	56120	6	19	2009/06/01T01:35:18.703	937	PA	120.253	IN	PA
10.1.1.126	192.168.2.6	1088	56136	6	18	2009/06/01T01:39:43.473	17128	FS PA	1.959	IN	S A
10.1.1.126	192.168.2.6	1088	56138	6	20	2009/06/01T01:40:26.589	19768	FS PA	1.871	IN	S A
10.1.1.126	192.168.2.6	1088	56139	6	46	2009/06/01T01:40:39.582	51015	FS PA	4.92	IN	S A
10.1.1.126	192.168.2.6	8088	56120	6	27	2009/06/01T01:39:14.034	1410	PA	145.437	IN	PA
10.1.1.126	192.168.3.55	8088	2453	6	12	2009/06/01T01:41:38.099	503	S PA	16.342	IN	S A
10.1.1.126	192.168.3.55	8088	2453	6	4	2009/06/01T01:43:27.786	170	PA	0.465	IN	PA
10.1.1.126	192.168.2.6	8088	56120	6	18	2009/06/01T01:45:25.114	840	PA	62.663	IN	PA
10.1.1.126	192.168.3.55	8088	2453	6	1	2009/06/01T01:48:30.394	40	F A	0	IN	F A
10.1.1.126	192.168.2.6	1088	56147	6	43	2009/06/01T01:49:41.930	47311	FS PA	4.209	IN	S A
10.1.1.126	192.168.2.6	1088	56153	6	28	2009/06/01T01:55:27.435	29816	FS PA	2.523	IN	S A
10.1.1.126	192.168.2.6	1088	56153	6	1	2009/06/01T01:55:30.066	40	A	0	IN	A
10.1.1.126	192.168.2.6	8088	56120	6	86	2009/06/01T01:49:41.855	3729	PA	634.678	IN	PA



Homeland Security

Approach

1. Ports of Interest
2. Flag Combinations
3. Port Jumping
4. Multiple-use Source Ports
5. Guilty by Association (Thanks Jay Brown!)

Homeland Security

Ports of Interest

- Malware analysis showed that Command and Control (C2) traffic consistently attempted to communicate over TCP ports 80, 443, 8080,1088,8088,8099

```
rwfilter --threads=2 --type=out,in,outweb,inweb --aport=<suspicious ports> --proto=6
--sensors=<list of sensors> --start-date=<date of query> --pass=$temp
```

- Only the initial outbound sessions SIP, DIP and DPort are needed

```
rwfilter $temp --type=out,outweb --sport=1024- --flags-initial=s/spearfuc --
pass=stdout |rwuniq --fields=1,2,4
```

- Filters:
  - Remove FTP Sessions
- After this filter, a set file is created of all suspicious IPs



Homeland Security

TCP Flag Combinations

- A large amount of traffic of interest can be found when looking at specific flag combinations
- By applying various filters, several queries are executed to identify specific traffic patterns that match the specific malware behavior
  - Beaconing:
    - S/S, S/SR, S/FA, RA/RA
  - Data Transfers/Keep Alives:
    - PA/PA, A/A

```
rwfilter <path to bin> --dipset=<set file from initial pull>
--type=outweb,out --sport=1024- --flags-all=pa/spearfuc --
pass=stdout | /usr/bin/rwuniq --fields=1,2,4
```

- Intersections are then applied throughout to create the final IP address list of interest



Homeland Security

ORIGINAL

sIP	dIP	sPort	dPort	pro	packets	sTime	bytes	flags	dur	type	initialF
10.1.1.126	192.168.1.19	1088	2496	6	46	2009/06/01T01:28:02.448	51015	FS PA	5.367	IN	S A
10.1.1.126	192.168.1.19	443	2493	6	3	2009/06/01T01:23:14.775	129	SRPA	0.001	INWEB	S A
10.1.1.126	192.168.2.6	8080	56105	6	4	2009/06/01T01:29:35.274	168	FS A	30.981	INWEB	S A
192.168.1.19	10.1.1.126	2493	443	6	3	2009/06/01T01:23:14.774	176	S PA	0.001	OUTWEB	S
192.168.2.6	10.1.1.126	56105	8080	6	5	2009/06/01T01:30:04.602	312	S PA	30.85	OUTWEB	S
192.168.2.6	10.1.1.126	56120	8088	6	1	2009/06/01T03:35:28.723	40	R	0	OUT	
192.168.2.6	10.1.1.126	56192	8088	6	1	2009/06/01T03:42:20.452	40	R	0	OUT	
10.1.1.126	192.168.3.228	8088	3833	6	4	2009/06/01T08:29:37.230	168	FS A	1.969	IN	S A
10.1.1.126	192.168.3.228	8088	3850	6	19	2009/06/01T08:30:32.540	783	S PA	18.024	IN	S A
10.1.1.126	192.168.3.228	1088	3917	6	46	2009/06/01T08:43:56.022	51015	FS PA	2.86	IN	S A
10.1.1.126	192.168.3.228	8088	3850	6	193	2009/06/01T08:34:13.861	8552	PA	784.708	IN	PA
10.1.1.126	192.168.3.228	8088	3850	6	3	2009/06/01T08:49:10.106	125	F PA	0.11	IN	PA
192.168.3.228	10.1.1.126	3833	8088	6	5	2009/06/01T08:29:07.170	309	S PA	1.93	OUT	S
192.168.3.228	10.1.1.126	3850	8088	6	27	2009/06/01T08:30:02.393	7404	S PA	17.866	OUT	S
192.168.3.228	10.1.1.126	3917	1088	6	28	2009/06/01T08:43:25.914	1176	FS PA	2.774	OUT	S
192.168.3.228	10.1.1.126	3850	8088	6	172	2009/06/01T08:33:43.839	27445	PA	784.601	OUT	PA
192.168.3.228	10.1.1.126	3850	8088	6	3	2009/06/01T08:48:40.018	125	F PA	0.109	OUT	PA
192.168.3.228	10.1.1.126	3833	8088	6	1	2009/06/01T09:29:16.300	40	R	0	OUT	

DISCOVERED

sIP	dIP	sPort	dPort	pro	packets	sTime	bytes	flags	dur	type	initialF
192.168.4.12	2.1.2.149	4042	80	6	6	2009/06/18T03:11:02.194	280	SR	29.893	OUTWEB	S
14.1.2.97	192.168.4.12	443	4042	6	1	2009/06/18T03:11:02.388	40	R A	0	INWEB	R A
192.168.4.12	172.16.2.194	4024	8080	6	51	2009/06/18T03:11:12.283	3639	RPA	457.766	OUTWEB	PA
172.16.2.194	192.168.4.12	8080	4024	6	68	2009/06/18T03:11:12.495	2873	PA	457.764	INWEB	PA
192.168.4.12	14.1.2.97	4062	443	6	2	2009/06/18T03:31:57.166	96	S	6.276	OUTWEB	S
14.1.2.97	192.168.4.12	443	4062	6	1	2009/06/18T03:31:58.164	40	R A	0	INWEB	R A
192.168.4.12	2.1.2.149	4062	80	6	6	2009/06/18T03:32:04.211	280	SR	29.97	OUTWEB	S
14.1.2.97	192.168.4.12	443	4062	6	1	2009/06/18T03:32:04.414	40	R A	0	INWEB	R A
192.168.4.12	2.1.2.149	4062	80	6	2	2009/06/18T03:32:34.189	88	SR	1.124	OUTWEB	S
192.168.4.12	14.1.2.97	4068	443	6	2	2009/06/18T03:52:51.495	96	S	6.803	OUTWEB	S
14.1.2.97	192.168.4.12	443	4068	6	1	2009/06/18T03:52:52.467	40	R A	0	INWEB	R A
192.168.4.12	2.1.2.149	4068	80	6	6	2009/06/18T03:52:59.059	280	SR	29.946	OUTWEB	S
14.1.2.97	192.168.4.12	443	4068	6	1	2009/06/18T03:52:59.270	40	R A	0	INWEB	R A
192.168.4.12	2.1.2.149	4068	80	6	2	2009/06/18T03:53:29.036	88	SR	1.235	OUTWEB	S
192.168.4.12	14.1.2.97	4071	443	6	2	2009/06/18T04:13:47.256	96	S	6.85	OUTWEB	S
14.1.2.97	192.168.4.12	443	4071	6	1	2009/06/18T04:13:48.229	40	R A	0	INWEB	R A
192.168.4.12	2.1.2.149	4071	80	6	6	2009/06/18T04:13:54.361	280	SR	29.94	OUTWEB	S
14.1.2.97	192.168.4.12	443	4071	6	1	2009/06/18T04:13:55.080	40	R A	0	INWEB	R A
192.168.4.12	2.1.2.149	4071	80	6	2	2009/06/18T04:14:24.336	88	SR	0.071	OUTWEB	S



Homeland Security

Port Jumping

- Non-USGA IPs are hosting services on multiple ports
- The initial data pull is filtered using the set file of suspicious IP's to identify which serve data on multiple ports
- Filters:
  - Remove IPs hosting only 80 and 443
  - Remove IPs hosting on only 1 port

ORIGINAL

sIP	dIP	sPort	dPort	pro	packets	sTime	bytes	flags	dur	type	initialF
10.1.1.126	192.168.1.19	1088	2496	6	46	2009/06/01T01:28:02.448	51015	FS PA	5.367	IN	S A
10.1.1.126	192.168.1.19	443	2493	6	3	2009/06/01T01:23:14.775	129	SRPA	0.001	INWEB	S A
10.1.1.126	192.168.2.6	8080	56105	6	4	2009/06/01T01:29:35.274	168	FS A	30.981	INWEB	S A
192.168.1.19	10.1.1.126	2493	443	6	3	2009/06/01T01:23:14.774	176	S PA	0.001	OUTWEB	S
192.168.2.6	10.1.1.126	56105	8080	6	5	2009/06/01T01:30:04.602	312	S PA	30.85	OUTWEB	S
192.168.2.6	10.1.1.126	56120	8088	6	1	2009/06/01T03:35:28.723	40	R	0	OUT	
192.168.2.6	10.1.1.126	56192	8088	6	1	2009/06/01T03:42:20.452	40	R	0	OUT	
10.1.1.126	192.168.3.228	8088	3833	6	4	2009/06/01T08:29:37.230	168	FS A	1.969	IN	S A
10.1.1.126	192.168.3.228	8088	3850	6	19	2009/06/01T08:30:32.540	783	S PA	18.024	IN	S A
10.1.1.126	192.168.3.228	1088	3917	6	46	2009/06/01T08:43:56.022	51015	FS PA	2.86	IN	S A
10.1.1.126	192.168.3.228	8088	3850	6	193	2009/06/01T08:34:13.861	8552	PA	784.708	IN	PA
10.1.1.126	192.168.3.228	8088	3850	6	3	2009/06/01T08:49:10.106	125	F PA	0.11	IN	PA
192.168.3.228	10.1.1.126	3833	8088	6	5	2009/06/01T08:29:07.170	309	S PA	1.93	OUT	S
192.168.3.228	10.1.1.126	3850	8088	6	27	2009/06/01T08:30:02.393	7404	S PA	17.866	OUT	S
192.168.3.228	10.1.1.126	3917	1088	6	28	2009/06/01T08:43:25.914	1176	FS PA	2.774	OUT	S
192.168.3.228	10.1.1.126	3850	8088	6	172	2009/06/01T08:33:43.839	27445	PA	784.601	OUT	PA
192.168.3.228	10.1.1.126	3850	8088	6	3	2009/06/01T08:48:40.018	125	F PA	0.109	OUT	PA
192.168.3.228	10.1.1.126	3833	8088	6	1	2009/06/01T09:29:16.300	40	R	0	OUT	

DISCOVERED

sIP	dIP	sPort	dPort	pro	packets	sTime	bytes	flags	dur	type	initialF
172.16.2.194	192.168.4.12	8088	4032	6	29	2009/06/18T02:48:02.303	29860	FS PA	1.081	IN	S A
192.168.4.12	172.16.2.194	4033	8088	6	14	2009/06/18T02:48:28.217	612	FS PA	0.919	OUT	S
172.16.2.194	192.168.4.12	8088	4033	6	20	2009/06/18T02:48:28.488	19772	FS PA	0.919	IN	S A
192.168.4.12	14.1.2.97	4034	443	6	2	2009/06/18T02:49:59.061	96	S	6.346	OUTWEB	S
14.1.2.97	192.168.4.12	443	4034	6	1	2009/06/18T02:50:00.087	40	R A	0	INWEB	R A
192.168.4.12	2.1.2.149	4034	80	6	6	2009/06/18T02:50:06.301	280	SR	29.92	OUTWEB	S
192.168.4.12	14.1.2.97	4038	443	6	3	2009/06/18T02:52:30.227	136	SR	36.824	OUTWEB	S
14.1.2.97	192.168.4.12	443	4038	6	1	2009/06/18T02:52:31.202	40	R A	0	INWEB	R A
192.168.4.12	2.1.2.149	4038	80	6	6	2009/06/18T02:52:49.247	280	SR	29.157	OUTWEB	S
192.168.4.12	2.1.2.149	4038	80	6	2	2009/06/18T02:53:19.281	88	SR	30.026	OUTWEB	S
192.168.4.12	14.1.2.97	4039	443	6	2	2009/06/18T02:54:05.203	96	S	0.284	OUTWEB	S
14.1.2.97	192.168.4.12	443	4039	6	1	2009/06/18T02:54:12.252	40	R A	0	INWEB	R A
192.168.4.12	2.1.2.149	4039	80	6	6	2009/06/18T02:54:19.050	280	SR	30.005	OUTWEB	S
192.168.4.12	2.1.2.149	4039	80	6	2	2009/06/18T02:54:49.115	88	SR	0.281	OUTWEB	S
192.168.4.12	14.1.2.97	4042	443	6	2	2009/06/18T03:10:55.109	96	S	6.306	OUTWEB	S
14.1.2.97	192.168.4.12	443	4042	6	1	2009/06/18T03:11:00.472	40	R A	0	INWEB	R A
192.168.4.12	2.1.2.149	4042	80	6	6	2009/06/18T03:11:02.194	280	SR	29.893	OUTWEB	S
192.168.4.12	172.16.2.194	4024	8080	6	51	2009/06/18T03:11:12.283	3639	RPA	457.766	OUTWEB	PA

Multiple-use Source Ports

- Malware is observed communicating over the same source port multiple times for outgoing communications to multiple IP addresses
- A clear pattern is observed in this communication dependent on the malware variant

ORIGINAL

sIP	dIP	sPort	dPort	pro	packets	sTime	bytes	flags	dur	type	initialF
192.168.3.55	5.1.1.60	2506	443	6	3	2009/06/01T02:13:52.435	191	S PA	0.099	OUTWEB	S
192.168.3.55	5.1.1.60	2506	443	6	1	2009/06/01T02:18:53.755	40	A	0	OUTWEB	A
192.168.3.55	5.1.1.60	2506	443	6	1	2009/06/01T02:23:52.466	40	F A	0	OUTWEB	F A
192.168.3.55	5.1.1.60	2530	443	6	3	2009/06/01T02:43:52.715	191	S PA	0.321	OUTWEB	S
192.168.3.55	5.1.1.60	2530	443	6	4	2009/06/01T02:45:03.612	261	F PA	7.647	OUTWEB	A
192.168.3.55	2.1.1.38	2549	80	6	6	2009/06/01T03:04:47.949	280	SR	29.787	OUTWEB	S
192.168.3.55	5.1.1.60	2549	443	6	2	2009/06/01T03:05:10.920	96	S	6.773	OUTWEB	S
192.168.3.55	2.1.1.38	2549	80	6	2	2009/06/01T03:05:18.122	88	SR	0.725	OUTWEB	S
192.168.3.55	2.1.1.38	2549	443	6	4	2009/06/01T03:05:30.105	184	SR	29.846	OUTWEB	S
192.168.3.55	2.1.1.38	2569	80	6	6	2009/06/01T03:26:04.184	280	SR	29.814	OUTWEB	S
192.168.3.55	5.1.1.60	2569	443	6	2	2009/06/01T03:26:20.894	96	S	0.641	OUTWEB	S
192.168.3.55	2.1.1.38	2569	80	6	2	2009/06/01T03:26:34.206	88	SR	1.101	OUTWEB	S
192.168.3.55	2.1.1.38	2569	443	6	4	2009/06/01T03:26:46.148	184	SR	31.426	OUTWEB	S
192.168.3.55	2.1.1.38	2582	80	6	6	2009/06/01T03:47:08.291	280	SR	29.554	OUTWEB	S
192.168.3.55	5.1.1.60	2582	443	6	3	2009/06/01T03:47:36.982	144	S	1.107	OUTWEB	S
192.168.3.55	2.1.1.38	2582	80	6	2	2009/06/01T03:47:38.287	88	SR	0.741	OUTWEB	S
192.168.3.55	2.1.1.38	2582	443	6	4	2009/06/01T03:47:50.386	184	SR	29.702	OUTWEB	S
192.168.3.55	2.1.1.38	2601	80	6	6	2009/06/01T04:08:30.247	280	SR	29.323	OUTWEB	S

DISCOVERED

sIP	dIP	sPort	dPort	pro	packets	sTime	bytes	flags	dur	type	initialF
192.168.4.12	14.1.2.97	4034	443	6	2	2009/06/18T02:49:59.061	96	S	6.346	OUTWEB	S
192.168.4.12	2.1.2.149	4034	80	6	6	2009/06/18T02:50:06.301	280	SR	29.92	OUTWEB	S
192.168.4.12	14.1.2.97	4038	443	6	3	2009/06/18T02:52:30.227	136	SR	36.824	OUTWEB	S
192.168.4.12	2.1.2.149	4038	80	6	6	2009/06/18T02:52:49.247	280	SR	29.157	OUTWEB	S
192.168.4.12	2.1.2.149	4038	80	6	2	2009/06/18T02:53:19.281	88	SR	30.026	OUTWEB	S
192.168.4.12	14.1.2.97	4039	443	6	2	2009/06/18T02:54:05.203	96	S	0.284	OUTWEB	S
192.168.4.12	2.1.2.149	4039	80	6	6	2009/06/18T02:54:19.050	280	SR	30.005	OUTWEB	S
192.168.4.12	2.1.2.149	4039	80	6	2	2009/06/18T02:54:49.115	88	SR	0.281	OUTWEB	S
192.168.4.12	14.1.2.97	4042	443	6	2	2009/06/18T03:10:55.109	96	S	6.306	OUTWEB	S
192.168.4.12	2.1.2.149	4042	80	6	6	2009/06/18T03:11:02.194	280	SR	29.893	OUTWEB	S
192.168.4.12	172.16.2.194	4024	8080	6	51	2009/06/18T03:11:12.283	3639	RPA	457.766	OUTWEB	PA
192.168.4.12	14.1.2.97	4062	443	6	2	2009/06/18T03:31:57.166	96	S	6.276	OUTWEB	S
192.168.4.12	2.1.2.149	4062	80	6	6	2009/06/18T03:32:04.211	280	SR	29.97	OUTWEB	S
192.168.4.12	2.1.2.149	4062	80	6	2	2009/06/18T03:32:34.189	88	SR	1.124	OUTWEB	S
192.168.4.12	14.1.2.97	4068	443	6	2	2009/06/18T03:52:51.495	96	S	6.803	OUTWEB	S
192.168.4.12	2.1.2.149	4068	80	6	6	2009/06/18T03:52:59.059	280	SR	29.946	OUTWEB	S
192.168.4.12	2.1.2.149	4068	80	6	2	2009/06/18T03:53:29.036	88	SR	1.235	OUTWEB	S



Homeland Security

Guilty By Association (GBA)

- A system that communicates with a known or suspected malicious IP may communicate with others
- For all systems that communicated with the set of suspicious IP's
  - Find out whom else they communicated with
  - Intersect these IP's and look for commonalities
- It is possible that no common IP's are found because we are ***intersecting*** sets
  - 27.1.1.x – Previously identified malicious IP's
  - 23.1.2.x, 50.1.2.x – 'Bad Guys' discovered using GBA tool



Homeland Security

ORIGINAL

sIP	dIP	sPort	dPort	pro	packets	sTime	bytes	flags	dur	type	initialF
192.168.3.35	27.1.1.69	3567	443	6	1	2009/07/02T02:33:32.354	48	S	0.084	OUTWEB	S
192.168.3.35	27.1.1.69	3567	443	6	1	2009/07/02T02:33:39.114	48	S	0.085	OUTWEB	S
27.1.1.69	192.168.3.35	443	3567	6	1	2009/07/02T02:33:32.438	40	R A	0	INWEB	R A
27.1.1.69	192.168.3.35	443	3567	6	1	2009/07/02T02:33:39.199	40	R A	0	INWEB	R A
27.1.1.69	192.168.5.16	1088	3214	6	254	2009/07/02T02:40:59.107	11307	S PA	1787.146	IN	S A
192.168.5.16	27.1.1.69	3214	1088	6	205	2009/07/02T02:40:57.452	40100	S PA	1787.555	OUT	S
27.1.1.69	192.168.5.16	8080	3228	6	50	2009/07/02T02:47:45.050	63459	FS PA	0.221	INWEB	S A
27.1.1.69	192.168.5.16	8080	3244	6	226	2009/07/02T02:53:41.460	287585	FS PA	0.851	INWEB	S A
27.1.1.69	192.168.5.16	8080	3249	6	1831	2009/07/02T02:55:08.362	73240	FS PA	4.884	INWEB	S A

DISCOVERED

sIP	dIP	sPort	dPort	pro	packets	sTime	bytes	flags	dur	type	initialF
192.168.3.110	50.1.2.250	3738	443	6	3	2009/07/02T02:02:27.443	192	S PA	0.657	OUTWEB	S
192.168.3.110	50.1.2.250	3738	443	6	2	2009/07/02T02:12:27.418	80	F A	0.034	OUTWEB	F A
192.168.3.35	50.1.2.250	2696	8080	6	3	2009/07/02T02:07:32.207	176	S PA	0.026	OUTWEB	S
192.168.3.35	50.1.2.250	3051	443	6	3	2009/07/02T02:18:50.057	176	S PA	0.026	OUTWEB	S
192.168.3.35	50.1.2.250	3094	8080	6	3	2009/07/02T02:20:26.228	176	S PA	0.026	OUTWEB	S
192.168.3.35	50.1.2.250	3164	8080	6	3	2009/07/02T02:21:39.120	172	S PA	0.026	OUTWEB	S
192.168.3.35	50.1.2.250	2448	443	6	51	2009/07/02T02:00:39.141	4245	S PA	1782.976	OUTWEB	S
192.168.3.110	50.1.2.250	3836	443	6	3	2009/07/02T02:32:28.076	192	S PA	0.137	OUTWEB	S
192.168.3.110	50.1.2.250	3836	443	6	2	2009/07/02T02:42:28.051	80	F A	0.029	OUTWEB	F A
192.168.3.35	50.1.2.250	3699	8080	6	24	2009/07/02T02:37:43.138	1016	FS PA	0.216	OUTWEB	S
192.168.3.35	50.1.2.250	2448	443	6	44	2009/07/02T02:31:37.480	6051	F PA	824.724	OUTWEB	PA
192.168.5.16	50.1.2.250	3228	8080	6	29	2009/07/02T02:47:43.396	1210	FS PA	0.72	OUTWEB	S
192.168.5.16	50.1.2.250	3244	8080	6	117	2009/07/02T02:53:40.305	4728	FS PA	0.852	OUTWEB	S
192.168.5.16	50.1.2.250	3249	8080	6	3984	2009/07/02T02:55:07.207	5041254	FS PA	4.91	OUTWEB	S
192.168.3.35	23.1.2.71	3608	443	6	8	2009/07/02T02:35:09.290	368	FS PA	0.114	OUTWEB	S
50.1.2.250	192.168.5.16	1088	3214	6	20	2009/07/02T03:11:14.009	837	PA	99.326	IN	PA
50.1.2.250	192.168.5.16	1088	3214	6	27	2009/07/02T03:18:40.436	1094	F PA	212.886	IN	PA
192.168.5.16	50.1.2.250	3214	1088	6	15	2009/07/02T03:11:12.380	962	PA	99.771	OUT	PA
192.168.5.16	50.1.2.250	3214	1088	6	21	2009/07/02T03:18:39.307	1034	F PA	212.886	OUT	PA



Homeland Security

Future

- Continually update white list
- Further testing and integration of GBA

Homeland Security

Contact

- Technical comments or questions
  - US-CERT Security Operations Center
  - Email: soc@us-cert.gov
  - Phone: +1 888-282-0870
- Media inquiries
  - US-CERT Public Affairs
  - Email: media@us-cert.gov
  - Phone: +1 202-282-8010
- General questions or suggestions
  - US-CERT Information Request
  - Email: info@us-cert.gov
  - Phone: +1 703-235-5110
- * Information available at http://www.us-cert.gov/contact.html

Homeland Security

Questions?



Homeland Security

Project Bloom: Empowering the Security Research Community Through Data Products and Computing

Minaxi Gupta — School of Informatics and Computing
Greg Travis, David Ripley — ANML
Doug D. Pearson — REN-ISAC

Indiana University, Bloomington

The Security Landscape

- Cybercrime involves a thriving multi-billion dollar underground economy
- Websites connected to phishing, malware, and scams spring up by the millions every day
- Creative spam, social engineering, and search-engine optimization techniques lure users to malicious websites
- Bots are a major enabler of cybercrime
- Security researchers need good data to study trends and devise effective defenses

Impediments to Data Availability

- Good data sets are either not available or available only within a closed group
- This hurts new progress and verification of previous results
- Even when data is available, it is rarely documented well or available for long durations
- Data procurement process is long or riddled with legal hassles
- Even when data is available, storing and computing on it, and indexing it for the future, requires resources rarely available at a single institution

Goals of Project Bloom

- Goal 1: to provide well-curated, long-term, raw data sets to the security research community in an efficient manner using an established federation that can make the data acquisition process almost instantaneous
- Goal 2: to provide researchers compute power close to the data sources so they can avoid moving large data sets across the Internet and can focus only on taking back the key derived results
- Goal 3: to offer rich data products derived from raw data so researchers can bypass common data-processing hurdles and be more productive

Partnerships

- Initial focus on three commonly-used data types: NetFlow, Darknet, Passive DNS
- To ensure quality and long-term availability, they are not community driven as of now
- Raw data will be kept for 10 years on a rolling basis
- Partners:
  - Indiana University:
    - REN-ISAC: 266 higher-education member institutions, some offer their darknet data
  - Internet2: NetFlow data from 300 member institutions
  - ISC/SIE: Passive DNS data from 15 ISPs around the world

Data Set 1: NetFlow

- Information contained in each NetFlow record: IPs, ports, protocol, bytes, packets of a *flow*, TCP flags, neighboring ASNs, routing information
- Widely used for traffic engineering, network provisioning, and for identifying security and performance problems
- A commonly-used source of high-quality NetFlow data is the Internet2 Observatory
- Limitations of the Observatory's offering:
  - Lengthy proposal process required to access data
  - Data is masked due to privacy concerns
  - Data is sampled
  - No data products or computing facility to enhance utility

NetFlow Under Bloom

- Data sources:
  - Internet2 Observatory
  - Indiana University
- A comprehensive technical, policy, and federation framework to cut down access times
- Time-limited unmasked data so conclusions about individual IPs can be drawn by trusted researchers
- Unsampled flow records from Indiana University
- Rich data products and parallel-computing facilities to process data

Data Set 2: Darknet

- Darknet: Allocated, advertised, but unused IP address space
- Packets coming to a darknet are unsolicited and may indicate malicious activity. Ex: Backscatter, scanning packets
- State of data availability:
  - IMS: not available publicly
  - Team Cymru: offers software to allow organizations to monitor data
  - UCSD's /8: CAIDA plans to offer it under PREDICT, offers limited view
  - Others: not widely known or available

Darknet Under Bloom

- Data sources:
  - Midwest: six /24s
  - East coast: 2 /24s
  - South: four /24s
  - West coast: UCSD's /8
  - Australian subcontinent: two /24s
  - More to be added
- Rich data products and parallel-computing facilities to process data

Data Set 3: Passive DNS

- Passive DNS: Responses obtained by the local DNS servers from (authoritative) DNS servers around the world
- When collected globally, passive DNS data offers a window into access patterns for hosts and domains around the world, including malicious ones
- State of data availability:
  - ISC/SIE: lengthy data procurement process, expensive
  - Others: mostly local, do not offer a global view

Passive DNS Under Bloom

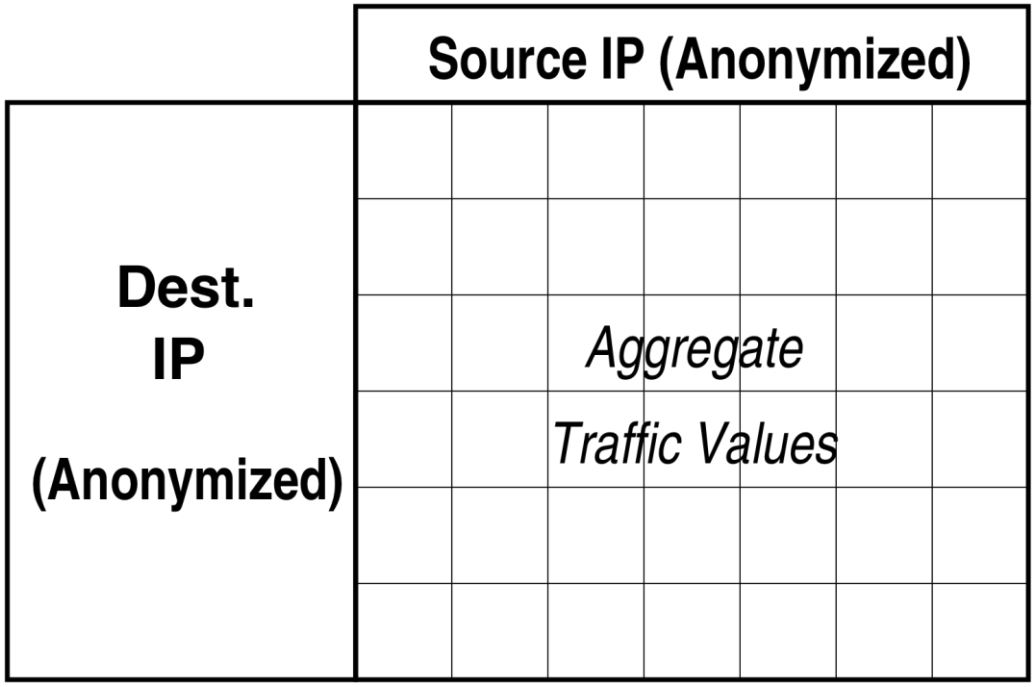
- Data sources:
  - ISC/SIE: collects passive DNS data from 15 large ISPs and commercial DNS service providers around the world, including a Tier1 ISP, two US Cable/DSL providers, and four US-based universities
  - DNS History Database Project (DHDB): collects passive DNS data from 8 data sources, including University of Auckland (New Zealand), France, and Norway
- Rich data products and parallel-computing facilities to process data

Data Products

- Fall in two categories:
  - based on individual data types
  - derived by combining data types
- Initial offering is composed of 12 products
- To derive them, we surveyed:
  - how researchers have used NetFlow data available through the Internet2 Observatory
  - CAIDA's offering of various data sets
  - research papers using darknet and passive DNS data
  - DHS, I3P, and Internet2 workshops that discussed cybersecurity data needs

Data Products: Example 1

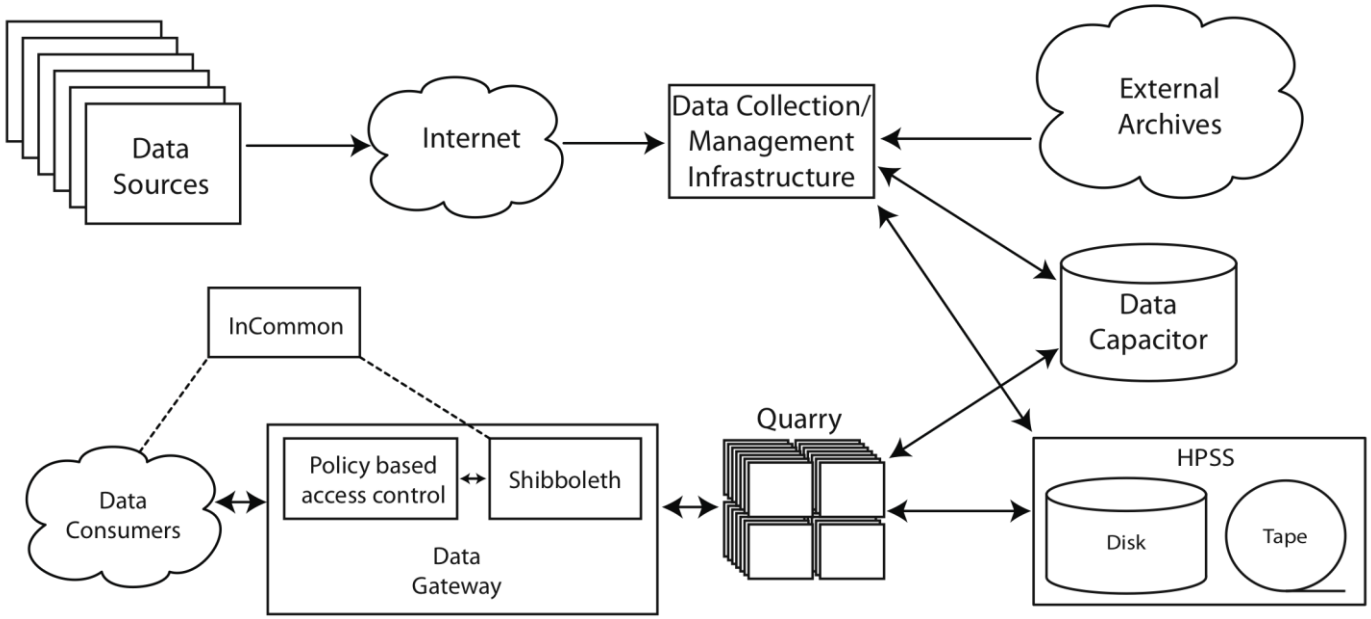
- Aggregate TCP traffic between each source and destination IP at a 15-minute granularity
- Can be combined to infer aggregate traffic for an IP address or BGP prefix at a daily, weekly, or monthly granularity
- 30 days worth of matrices would require 3TB storage



Data Products: Example 2

- Internet domains fall under gTLDs (.com, .net, etc.) and ccTLDs (.kr, .cn, etc.)
- Zone files enumerating domains within a TLD are a useful tool to investigate DNS aspects of malicious websites
- Prominent gTLD zone files are available but they cover only about 50% of Internet domains
- ccTLD zone files are not available, preventing researchers from obtaining a global view of malicious domains
- Using passive DNS and darknet data, we will generate nightly zone files for all TLDs, including ccTLDs, and archive them for 10 years on a rolling basis
- This data product would require 100TB of storage!

Architecture and Indiana University Resources



- Quarry: IBM e1350 distributed shared-memory supercomputer cluster with 1120 processor cores sharing 1.2TB of RAM
- Data Capacitor: 1PB of inline mass storage connected
- Massive Data Storage Service (MDSS): 2.8PB of magnetic tape and disk storage
- All connections are 20Gb/s Ethernet
- Access to Bloom will be through a Shibboleth instantiation, *InCommon*, already in use by Internet2

Conclusions

- Raw data and data products combined warrant 1PB of storage and this will increase with Internet traffic
- Access to Indiana University resources is not exclusive
- Financial support is needed to
  - provision adequate compute and storage for the research community
  - instantiate and maintain Bloom infrastructure
  - program and offer data products
  - evolve data products based on community feedback
  - administrative maintenance
- Your feedback and support is crucial!

Questions/Thoughts?

Contact Gregory Travis at:

greg@iu.edu

812 855 5091

Information Security:
It's Everyone's
Responsibility.

MITS

Cybersecurity

Know Your Network

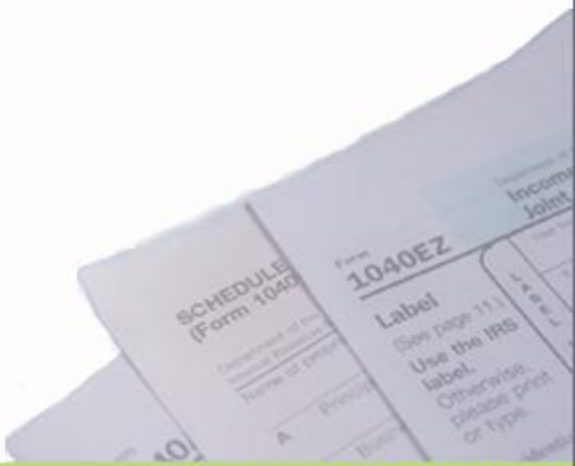
Computer Security Incident Response Center

"Tracking Compliance…Identifying and Mitigating Threats"

Teaming for Results

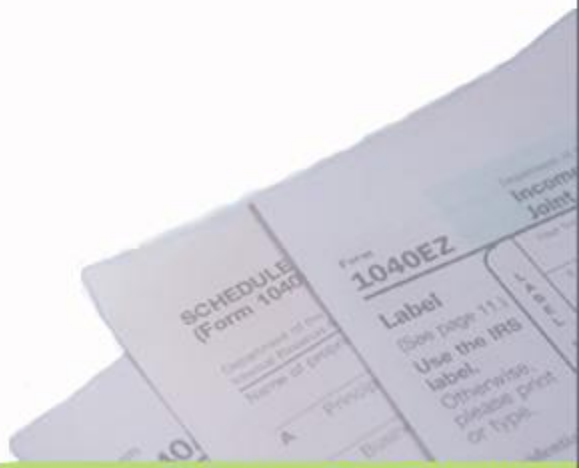
Agenda

- Background
- What does "Know Your Network" mean?
- About me
- Breadth, then depth
- Jumping off points approach
- Tiers of joy
- Tier 1 analysis
- Tier 2 analysis
- Tier 3 analysis
- Some interesting examples
- Case study
- Q&A



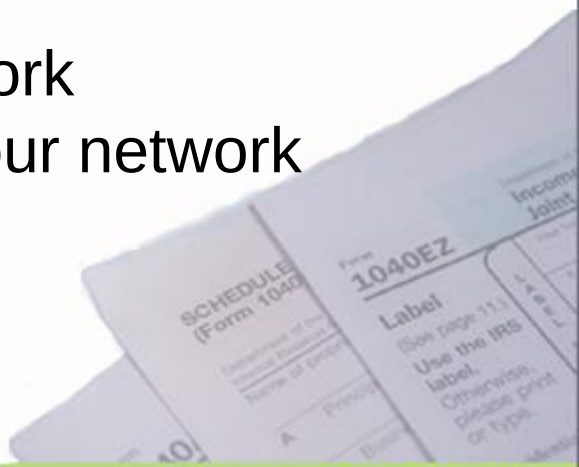
Background

- Networks are large, busy, and complicated (a challenge)
- Data is voluminous (another challenge)
- Need to know what belongs on a network to know what doesn't
- Profiling/modeling an enterprise network is not feasible
- Need to approach the problem in a different way
- Consider a sampling/best approximation approach



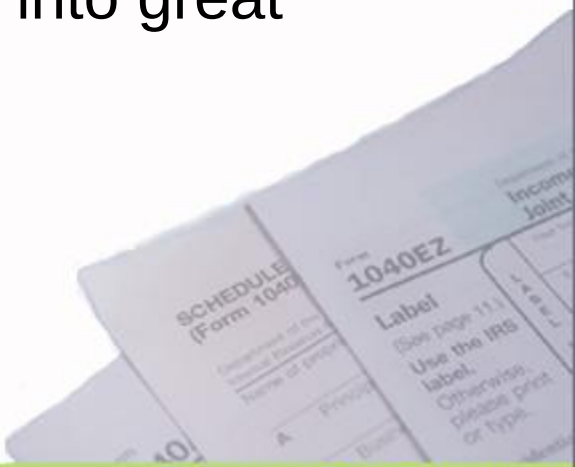
What Does "Know Your Network" Mean?

- Need to know what belongs on a network to know what doesn't
- In practicality, this isn't possible
- Using an organized, well-structured approach to analysis, you can come close
- Close is usually good enough for starters (you're likely to find things you're not finding now)
- Continuously tune and re-evaluate analytical approach as time progresses
- Along the way, you'll be learning your network
- Eventually, you'll come close to knowing your network



About Me

- I don't really like talking about myself
- Briefly:
  - Former Chief of Analysis for US-CERT
  - Currently a consultant focused on helping organizations build and enhance their network traffic analysis programs
  - Also a USAF Reservist focused on network traffic analysis and training analysts
  - Started out as programmer, then transitioned into infosec
  - Lots of experience turning smart people into great analysts

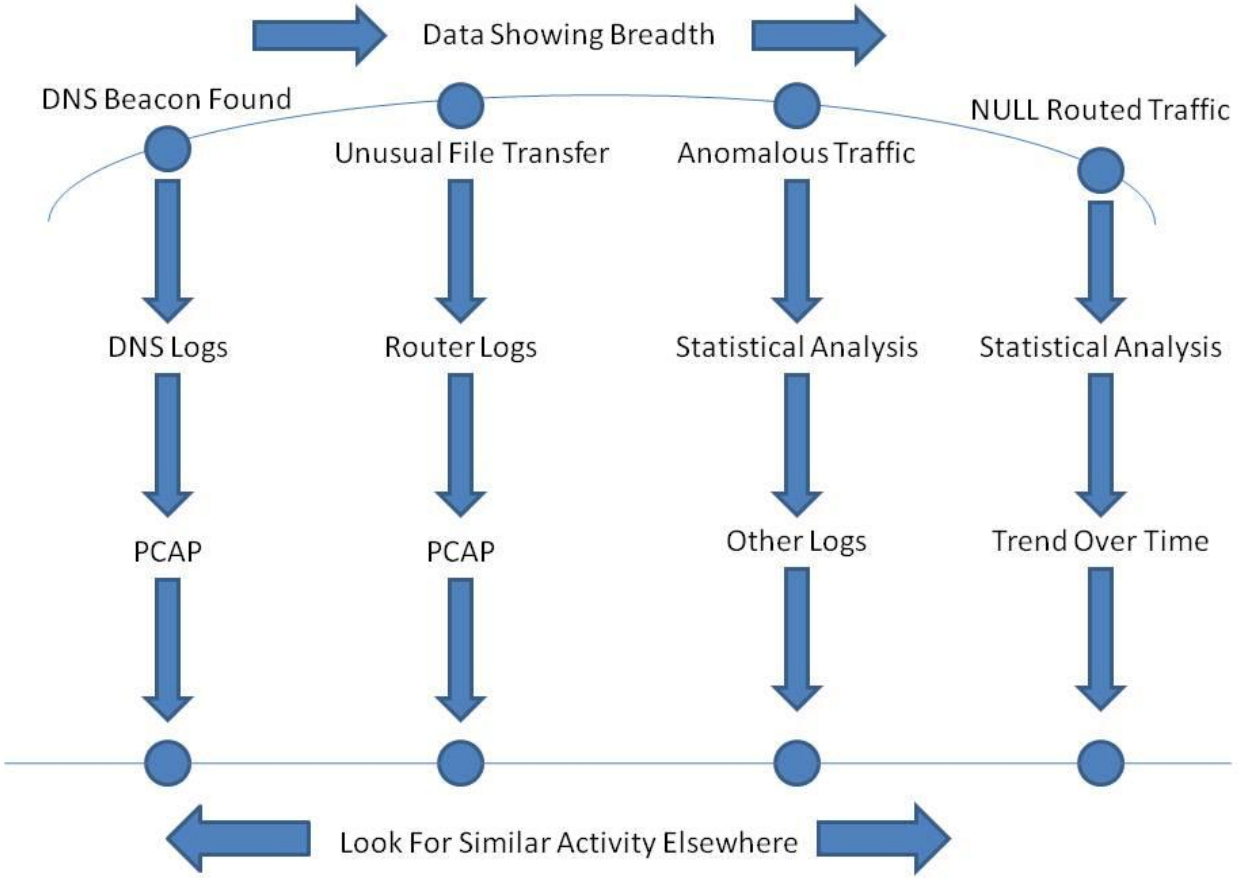


Breadth, Then Depth

- Seek first to understand, then to be understood

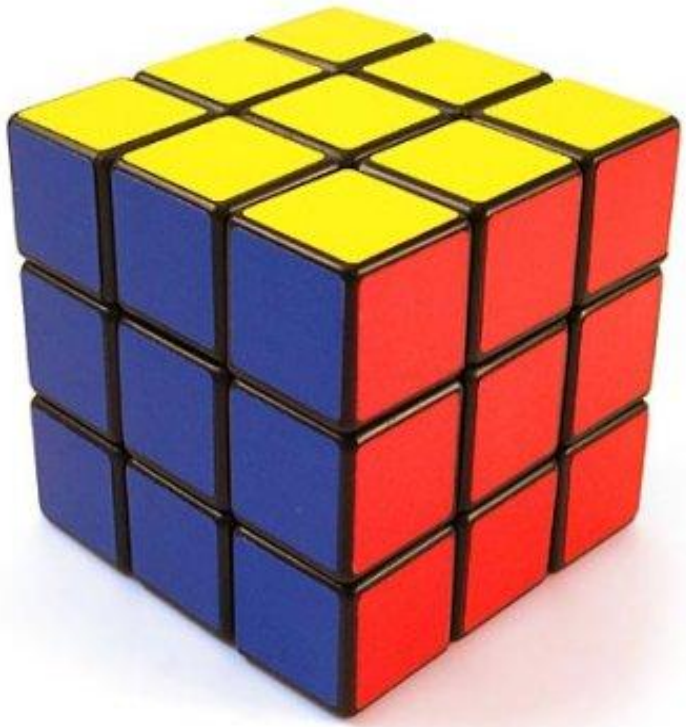


Breadth, Then Depth (Continued)



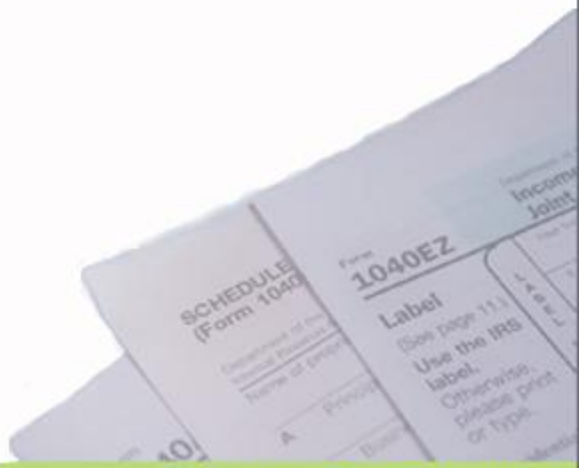
Jumping Off Points Approach

- A large network is essentially a very complicated black box
- Observing the network from many different jumping off points enables us to better understand what may actually be going on inside and facilitates analyst workflow



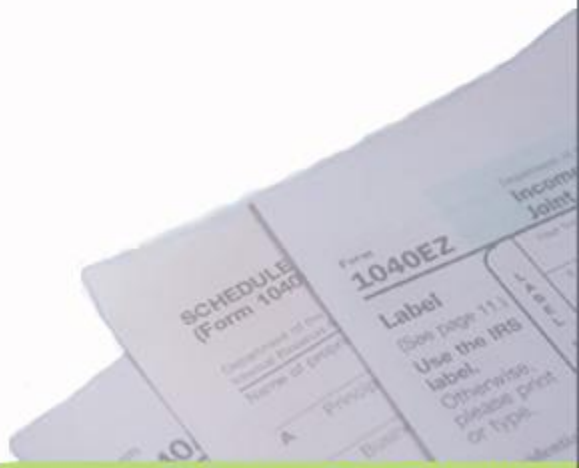
Tiers of Joy

- Tier 1 -- basic analysis used for situational awareness
- Tier 2 -- mid-level analysis used for ad hoc queries and moderately complex queries
- Tier 3 -- advanced analysis (looking for the needle in a pile of needles)



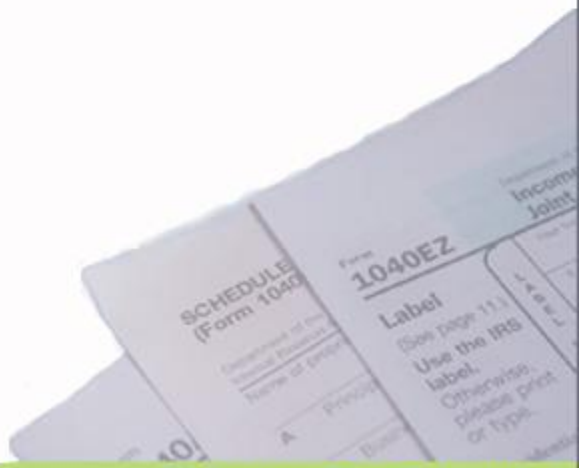
Tier 1 Analytical Methods

- IP watch lists
- Queries keyed on IP address
- Traffic volume
- Sensor up/down status
- Counts
- etc...



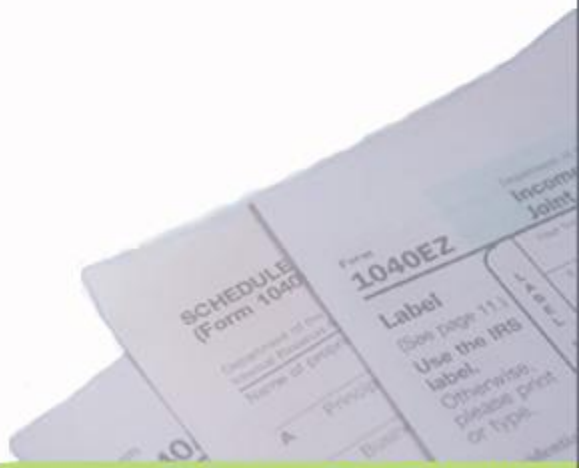
Tier 2 Analytical Methods

- Queries keyed on fields other than IP address
- Queries keyed on more than one field
- Queries looking for anomalies that are moderately difficult to spot
- Correlation with other data sources
- Automation/scripting
- etc...



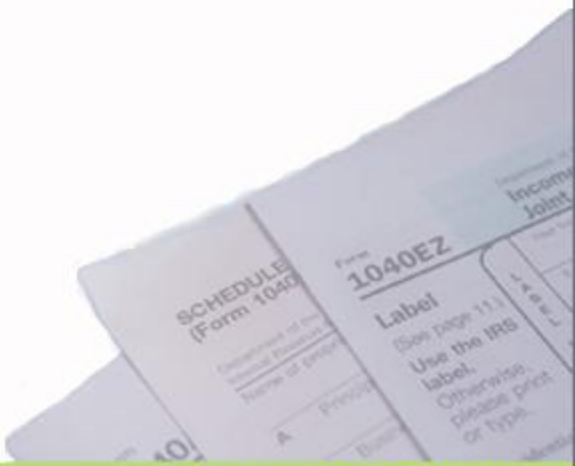
Tier 3 Analytical Methods

- Performing statistical analysis
- Examining how the data inter-relate
- Trending over time
- Trending over other fields
- Reviewing data in 3-tuples not involving time
- etc...



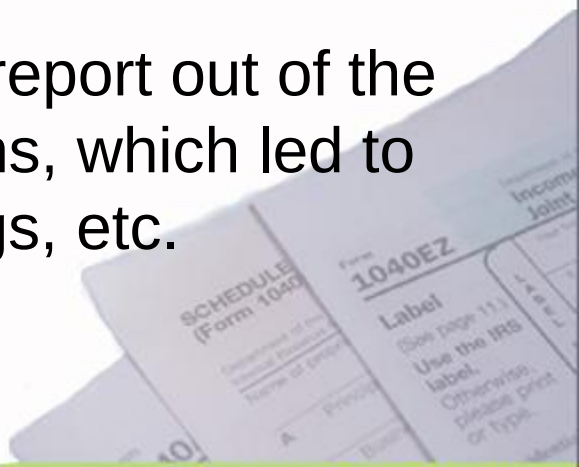
A Few Interesting Examples

- Source port, destination port, number of bytes 3-tuple
- Same number of bytes/packets/flows every N minutes
- Encrypted traffic over unencrypted protocols
- Unencrypted traffic over encrypted protocols
- Large transfers outbound from desktops (not servers -- know your network!)
- Packets not conforming to IETF standards
- And many more!



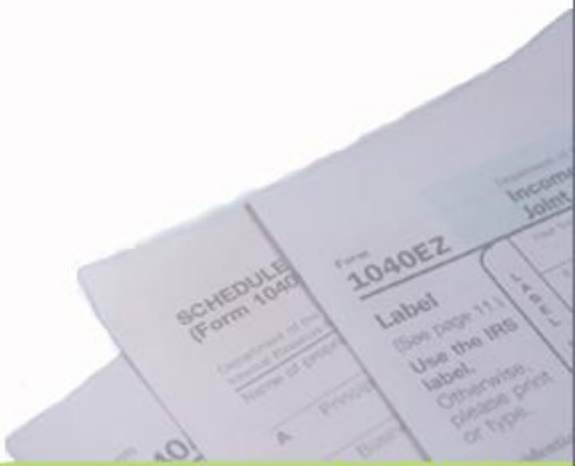
Case Study

- Studied a large, enterprise network over a week's worth of data using the approach described here
- Interesting findings included (but weren't limited to):
  - The network routed far more protocols than most people realized it did
  - Like it or not, the network was already routing IPv6
  - Some of the traffic found was designed to fly under the radar and would not be caught by even the most current signatures (Donald Rumsfeld)
  - Continuous tuning is necessary -- the report out of the findings generated interest and questions, which led to more analysis, which led to more findings, etc.



Acknowledgments

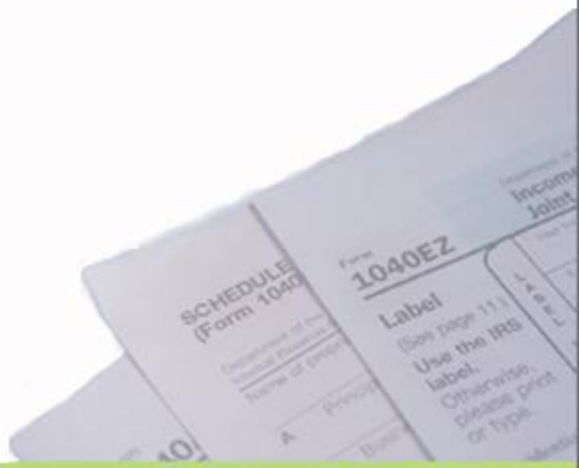
Thank you IRS CSIRC!



Q&A

Question/Comments/Produce?

Josh Goldfarb
President
NetflowData LLC
(240) 393-1314 (m)
(240) 235-3744 (f)
josh@netflowdata.com
http://www.netflowdata.com/



ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Abstracting and Visualizing Host Behaviour through Graphs

Eduard Glatz
Computer Engineering and Networks Laboratory
ETH Zurich (Switzerland)
eglatz@tik.ee.ethz.ch



Motivation

- Research in behavioural host profiling
  - Dominant and new session structures/application mixes
  - Evolution of host profiles over time
- Investigation of security incidents
  - Is this IP address a server or a client?
  - What services is this IP address providing?
- Teaching
  - Explain how Berkeley sockets work
  - Show complex communication patterns

Idea

- Common tools focus on traffic as a whole
- Browsing through flow lists might be a solution - but is unattractive when lists get very long
- Summarization techniques for flow lists exist, but are specialized on anomaly detection

- Idea: *develop your own tool*

Host Behaviour seen in Traffic Data

Traffic types

- Application traffic (user-triggered)
- Basic lookup traffic (application-triggered), e.g. DNS
- Infrastructure traffic (system-triggered), e.g. DHCP

Host profiling

1. What application mixes are prevalent?
2. Which roles are incorporated by hosts? (e.g. client, server, P2P role)
3. How do these two properties depend on each other?

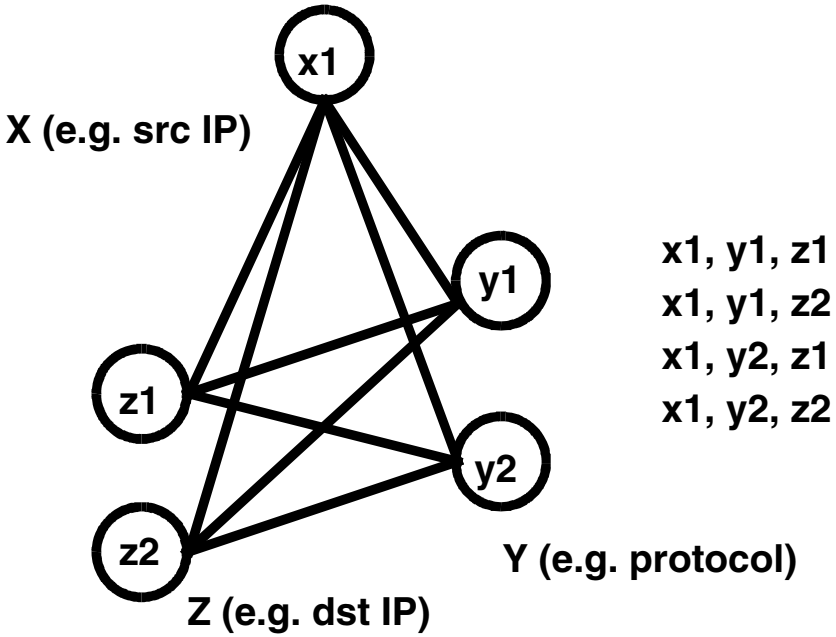
How to represent Host Traffic?

Idea: *use graphs*

- Nodes correspond to flow attributes
- Links show flow attributes that appear together
- Result: very dense/noisy graph

Problem:

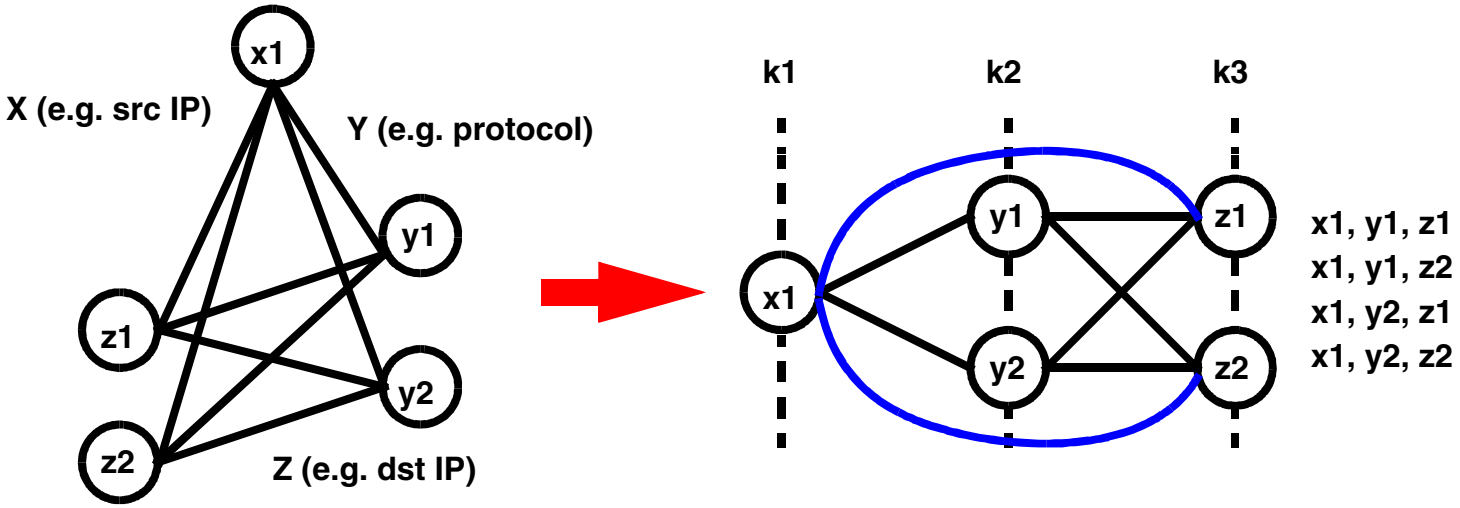
- Which relationships are most interesting to illustrate?



Transaction Visualization by k-Partite Graphs

Approach:

- K-partite graphs plus abstraction, e. g.

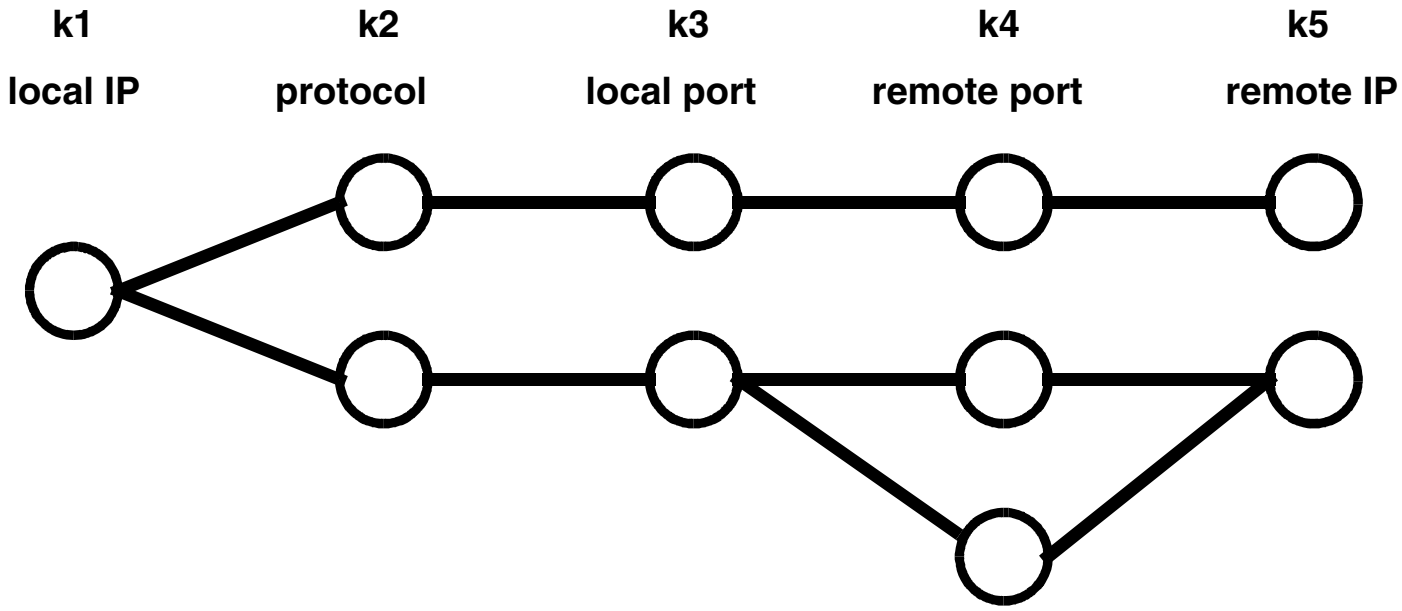


Abstraction:

- Purge blue lines and re-arrange partitions as needed to keep most interesting edges

Host Application Profile (HAP) Graphlet

We propose: Host traffic visualized by 5-partite graph



- Terminology: local/remote instead of source/destination
- Optionally annotate nodes with attribute values (not shown)

HAP Graphlet

- Visualizes BSD socket based communication in a straight forward way:
  - IP addresses assigned to first/last partition (k1, k5) show layer 3 connectivity
  - Central partitions (k2..k4) show layer 4 connectivity
  - Respects port number uniqueness (per protocol, per IP address)

Properties of HAP Graphlets

- Near-planar structure
- Shows remote IPs and ports associated with local port numbers (flows grouped per application)
- Host roles are appearant:
  - Server role
  - Client role
  - Peer role

What Graph Structures can we expect?

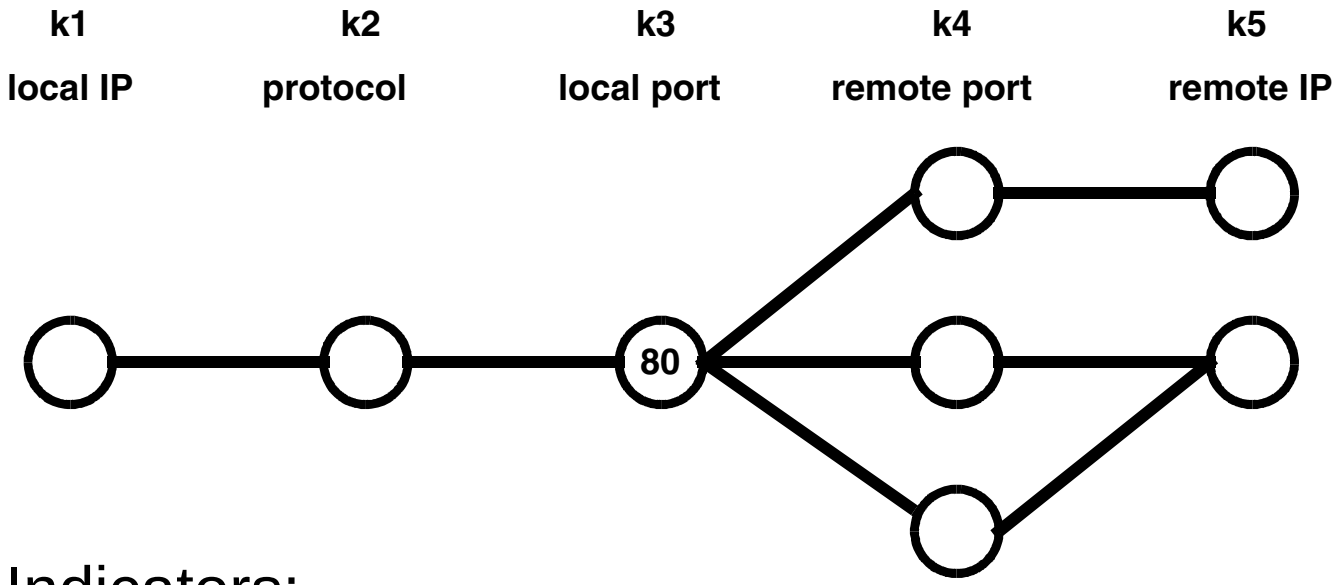
One session per application per peer:

- Client/server host roles
- P2P host roles

Complex sessions (applications) that use one or more connections:

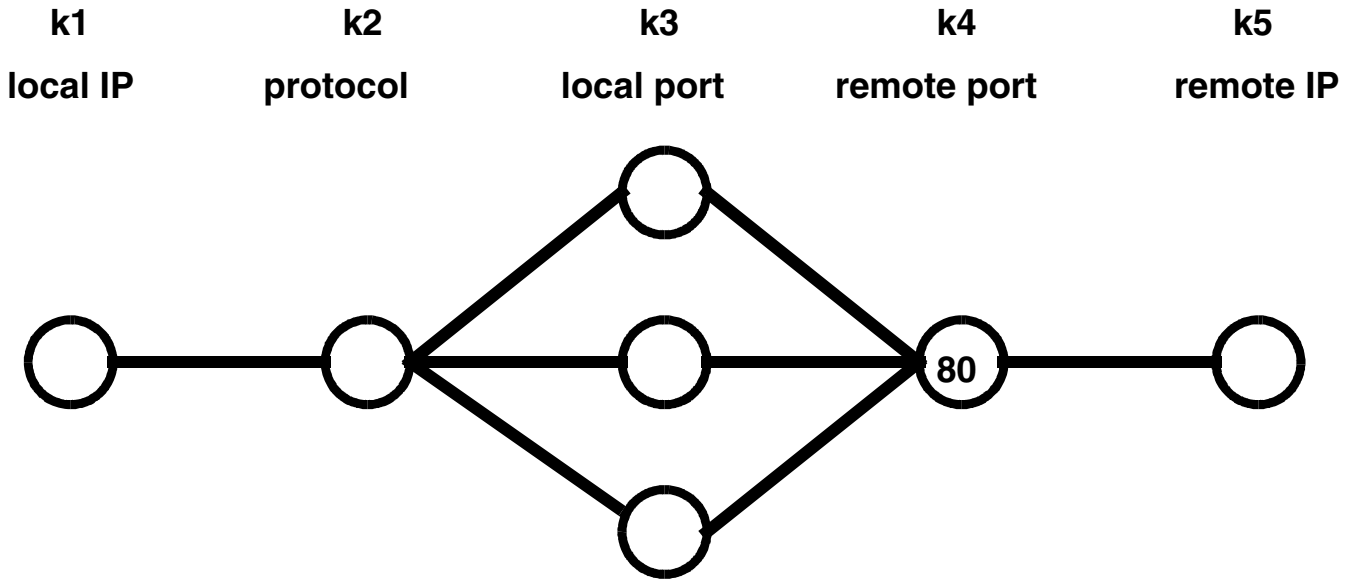
- To handle different tasks in parallel (e. g. control and data exchange)
- To improve performance (parallel flows to same remote endpoint)

Server Role



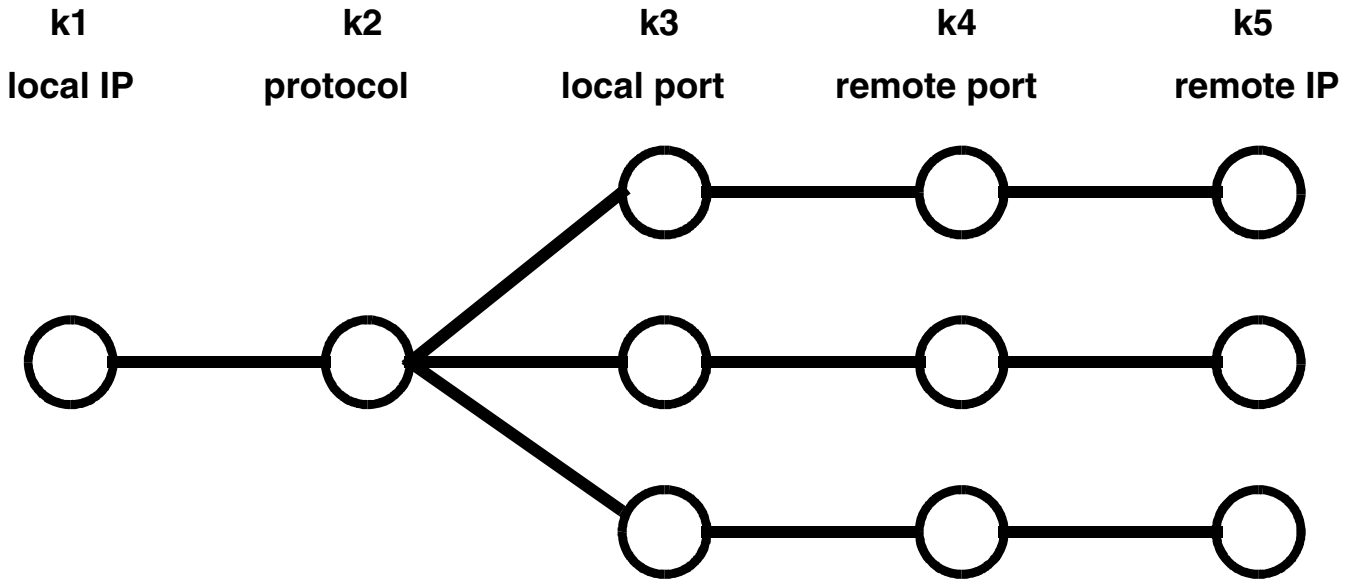
- Indicators:
  - K3 -> k4 out-degree d1 > 1 (multiple remote connections)
  - K3 -> k5 (virtual) out-degree d2 > 1 (multiple clients)
  - Often: d1 > d2 (parallel connections)

Client Role



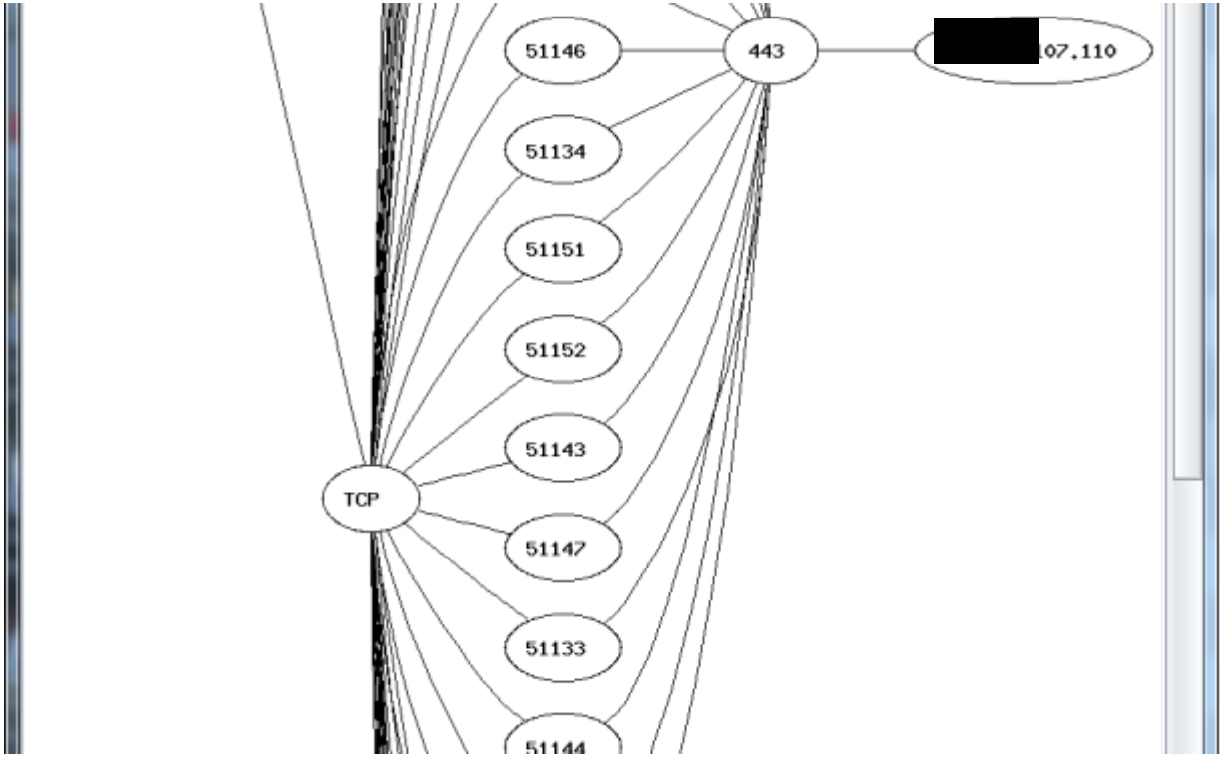
- Indicators:
  - K4 <- k3 in-degree d3 > 1
  - Often: d3 = 1
    (work around: local port > 1024, remote port < 1024)

Peer-to-Peer Role



- Indicators:
  - Many remote peers connected, and involved port numbers all above 1024
  - Hard to confirm: needs additional data sources

- Ideally, HAP graphlet fits into available screen area
- But …



Role Summarization

Idea:

- Compress *per-role subgraphs*

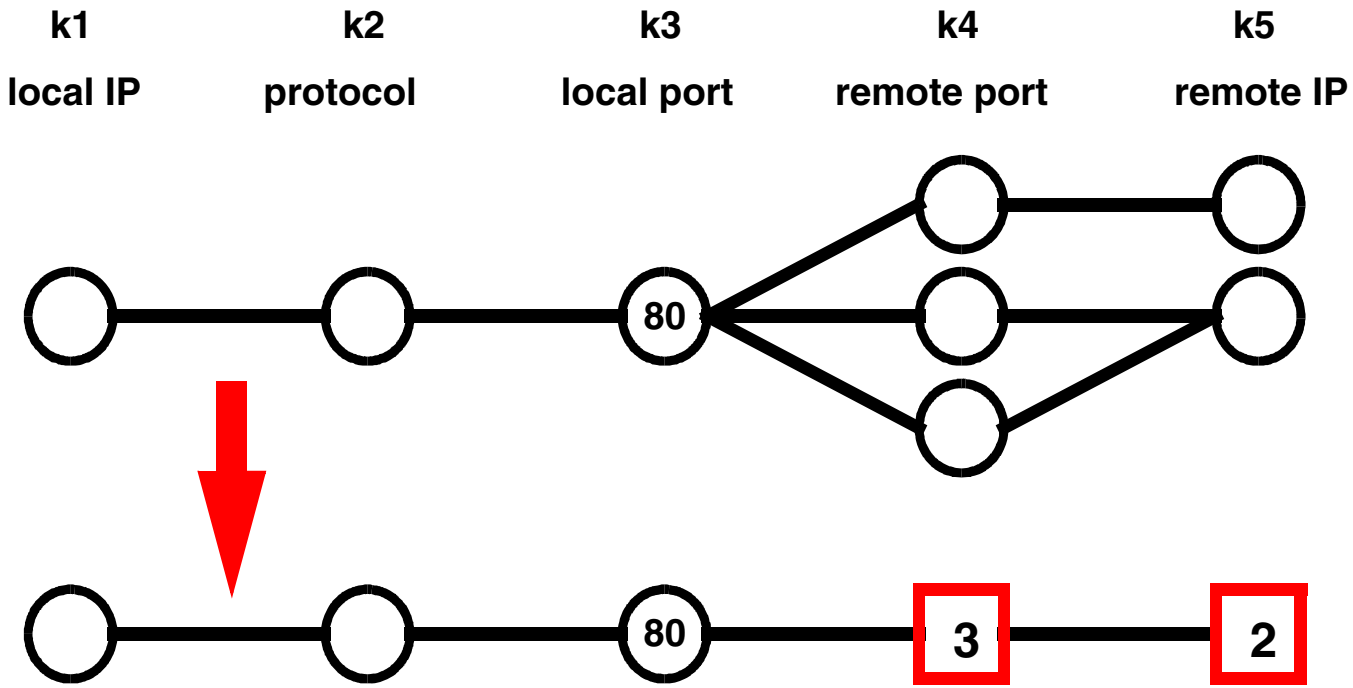
Prequisite:

- Roles can be associated with sub-graphs

Methodology:

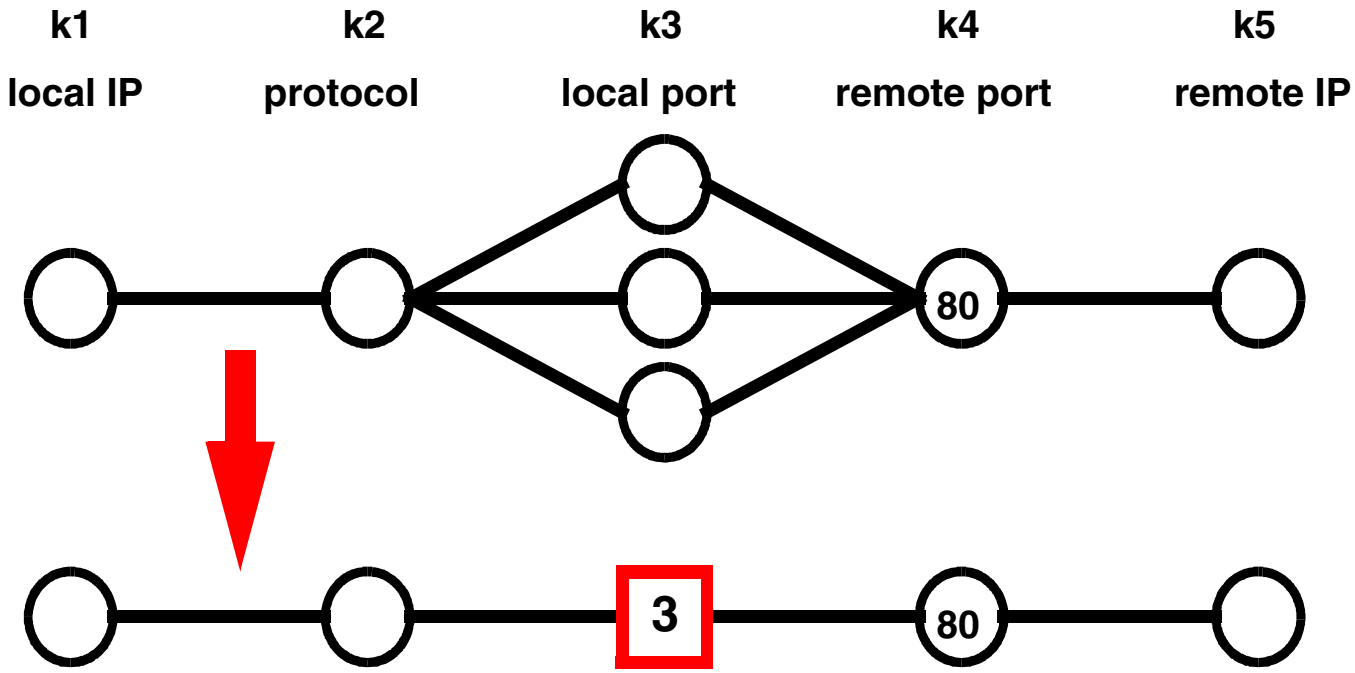
- Decompose graphlet into role-related subgraphs
- Replace such sub-graphs by summary sub-graphs
- Ignore graph partitions without role assignment
- ➢ Decomposition and replacement algorithm depends on role types (server/client/p2p roles)

Server Role Summary



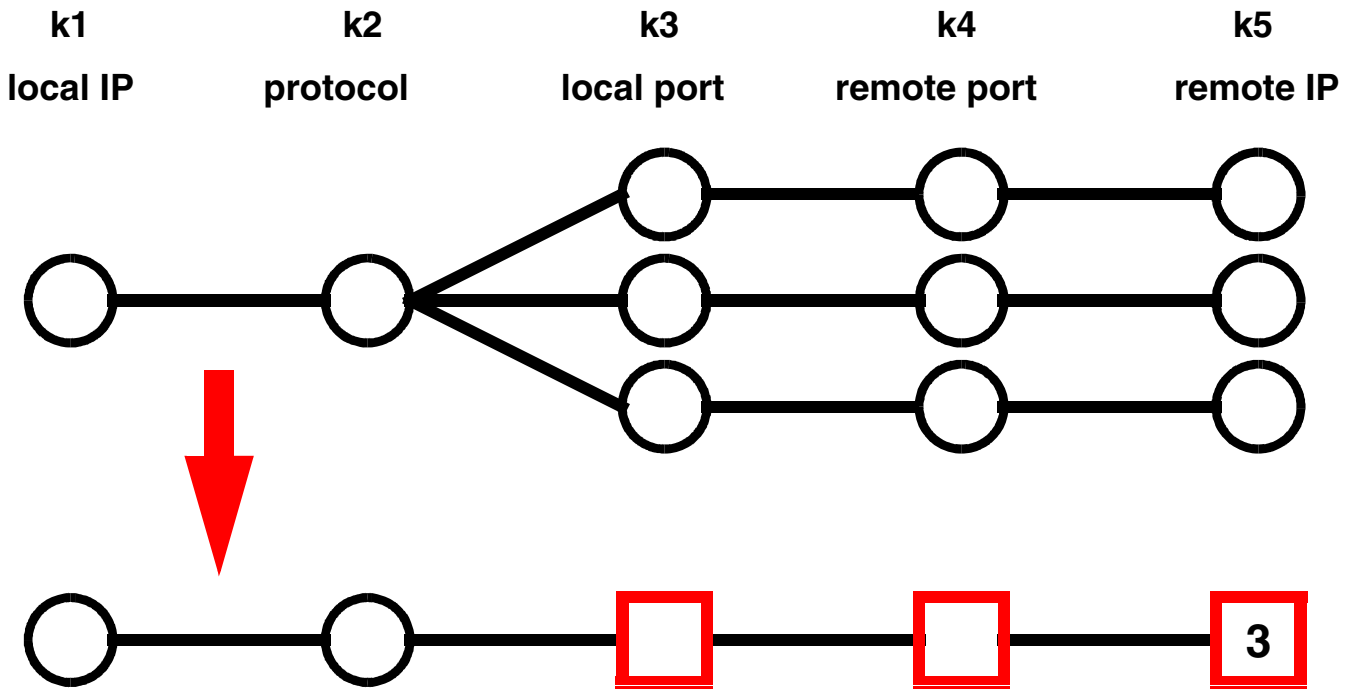
- Replace server-role related sub-graph
- Node annotations mark #connections and #clients

Client Role Summary



- Replace client-role related sub-graph
- Node annotation marks number of connections

Peer-to-Peer Role Summary

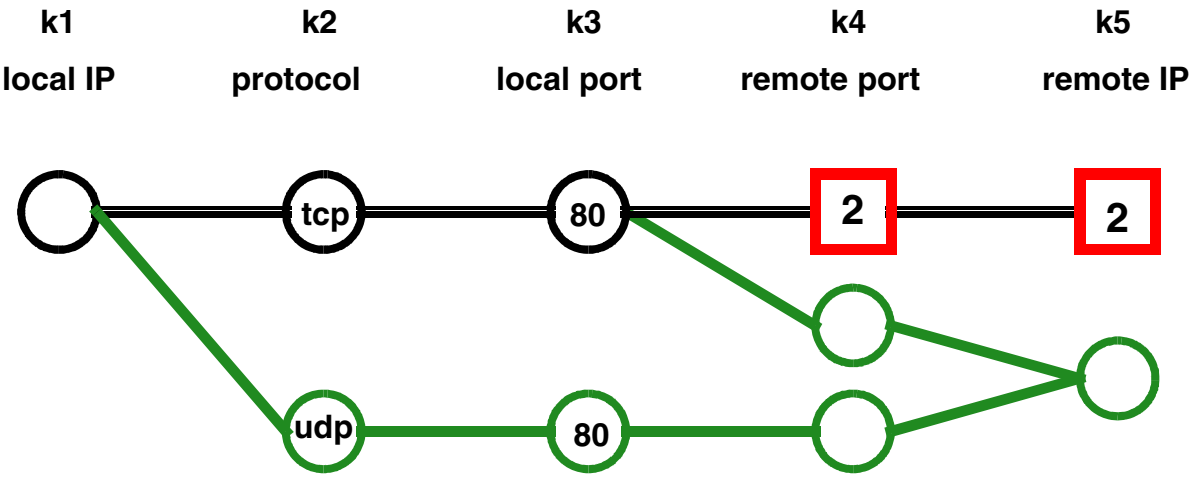


- Replace P2P-role related sub-graph
- Node annotation marks number of peers

Minimizing Information Loss

Scenario 1 (server role):

- One or more clients use double server connectivity through two protocols (e. g. control and data connections)
- Full summarization cannot include both connection paths



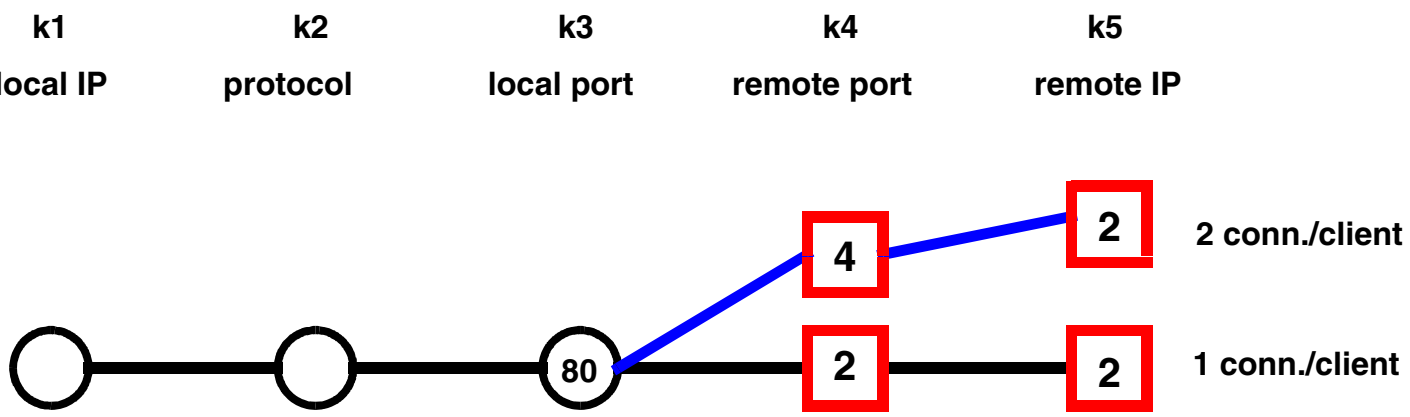
Approach:

- Do not summarize affected client(s)
- Use available screen height as a constraint

Minimizing Information Loss

Scenario 2 (server role):

- One or more clients use parallel connections to server
- Full summarization gives average parallelization degree only



Approach:

- Split summary into suitable parallelization groups
- Use available screen height as a constraint

Productive and Unproductive Traffic

Fact:

- A considerable part of Internet traffic is unproductive (e. g. scanning, misdirected flows)

But:

- We are mainly interested in productive traffic to characterize host behaviour
- When scan traffic enters the picture, then we want to identify it as such

Scan Traffic Filtering

- Mark and optionally remove scan traffic from visualization

- How to distinguish scan from productive traffic?
  - Hypothesis:
    productive traffic is bidirectional,
    e. g. involves bilateral interaction on the transport layer

- Methodology:
  - Pair unidirectional flows in opposite direction that use identical endpoints
  - Look „over the fence“ (i. e. observation interval borders) when searching a buddy for a within-interval unidirectional flow

Evaluation of Filtering Approach

- We inspected two dark IP ranges maintained by our ISP for scan traffic monitoring (over ~2400 h)
  - Captured at least traffic observed by network telescopes
  - Our advantage: we can do it over all address ranges

- Limitations:
  - NTP: in symmetric mode NTP source sends periodically unacknowledged messages to peers subscribed (RFC 958)
  - Multicast Source Discovery Protocol (MSDP) : works unidirectional
  - Discard protocol (tcp port 9)
  - Situations of multi-connected applications which run over different interfaces (and one connection is unidirectional)

The Tool: HAPviewer

- Unix/Linux application with graphical user interface

- Typical use cases:
  - Qualitative studies of roles incorporated by hosts
  - Investigations of complex connection structures (e.g. of P2P applications)
  - Identifying unknown service ports
  - Evaluations of scan traffic
  - Teaching of Berkeley socket model

Tool Architecture and Technologies

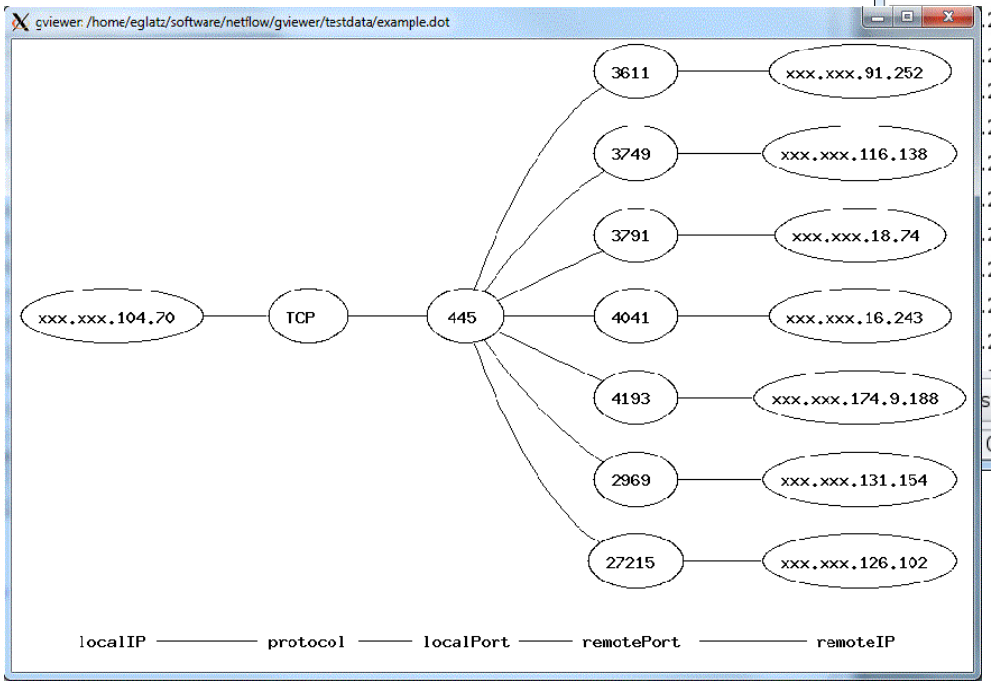
- Technology overview:
  - Unix/Linux based
  - C++
  - Gtk over gtkmm (C++ wrapper)
  - Graphviz
  - Pcap++

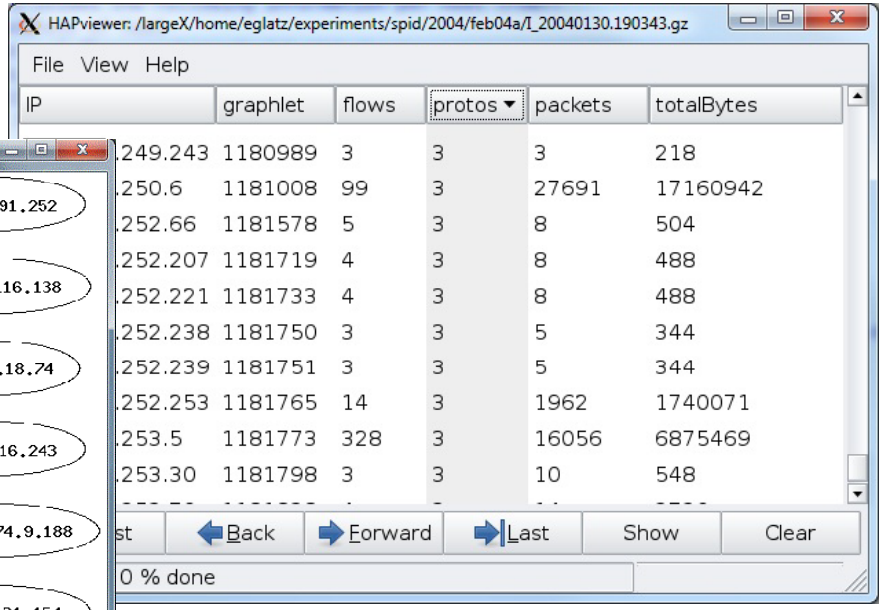
- Availability:
  - Available under dual licence GPL/BSD

Contributions

- Graph-based host traffic visualization
  - Illustrates used protocols and the functional role of local and remote port numbers
- Techniques for filtering unwanted traffic
  - Identification of scan traffic or misdirected flows
  - Prerequisite for proper service port identification
- Open-source visualization tool
  - Processes efficiently several millions of flows
  - Provides techniques for summarizing our graph-based profiles
  - Options for manipulating profiles and displaying graphs

Questions?





Tool will be demonstrated

Traffic Analysis Using Streaming Queries

Mike Fisk
Los Alamos National Laboratory

mfisk@lanl.gov

Los Alamos
NATIONAL LABORATORY
EST. 1943

Outline

- Intro to Continuous Query Systems
  - a.k.a Streaming Databases
  - Relevance to data networks
- Optimizing the evaluation of multiple Boolean queries
  - Counting Algorithm
  - Snort
  - Static Dataflow Optimization
    - Common Subexpression
    - Vector Algorithms
- Performance Comparisons

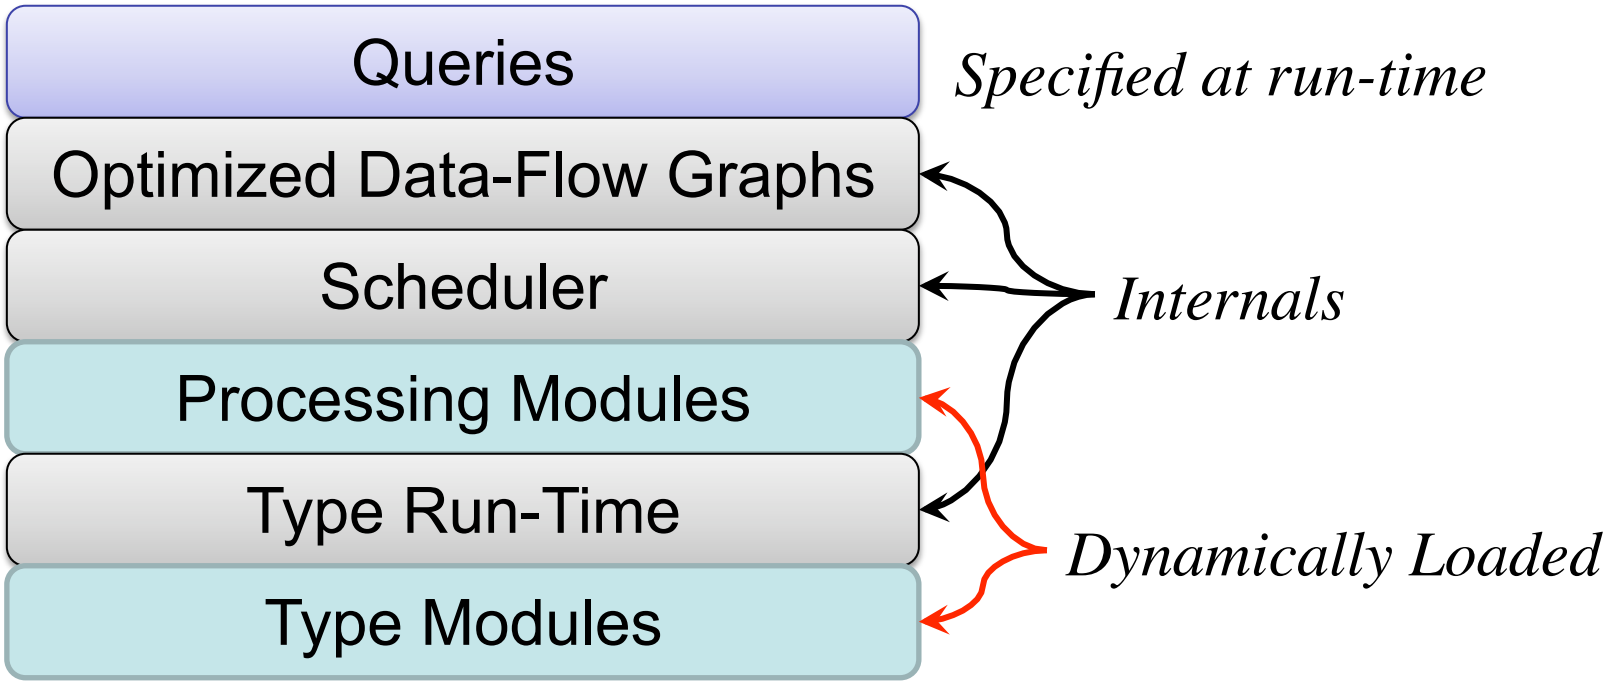
Observations

- Traffic analysis tools are data-type-specific
  - Flowtools – netflow
  - Snort – pcap
  - Psad – iptables logs
  - …
- Most analysis systems lack a framework for optimizing rules/queries
  - Reordering boolean expressions
  - Grouping (common sub-expressions)
  - Vector/set operations

Continuous Query Systems

- Continuous Query systems are to streaming data what Relational Database systems are to stored data
  - Filtering, summarization, aggregation
- Example datasets:
  - Sensor data (temperature, traffic, etc)
  - Stock exchange transactions
  - Packets, flows, logs
- Inefficient and high latency to load data into a traditional database and query periodically.
  - How often could you afford to re-execute the query?
- Example systems:
  - NiagraCQ (Wisc), Telegraph (Berkeley), SMACQ, etc.
  - Commerical: StreamBase, etc.
- Example systems in disguise:
  - Snort, router ACLs, firewall filters, packet classification, egrep

System for Modular Analysis & Continuous Queries



Los Alamos NATIONAL LABORATORY EST.1943

Type Model

- Stream of dynamically & heterogeneously typed objects
  - Each object can have different type
  - Types need not be statically defined in advance
- Objects refer to storage locations
  - Internal to the object, or references into other objects or external memory
- Objects have fields
  - Fields are (indifferently) struct elements, enums, unions, casts, string conversions, etc.
  - Fields are first-class objects
  - Fields can be dynamically attached to objects
- Objects are immutable
  - Enables parallelism without locking

Type Module Definition

- There are no fundamental types
- Pcap packet example

```
struct dts_field_spec dts_type_packet_fields[] = {
//Type          Name                Access Function if not fixed
{ "timeval",    "ts",               NULL },  // Fixed-length, fixed-location
{ "uint32",     "caplen",           NULL },
{ "uint32",     "len",              NULL },
{ "ipproto",    "ipprotocol",       dts_pkthdr_get_protocol },  // Function-pointer
{ "string",     "packet",           dts_pkthdr_get_packet },
{ "macaddr",    "dstmac",           dts_pkthdr_get_dstmac },
{ "nuint16",    "ethertype",        dts_pkthdr_get_ethertype },
{ "ip",         "srcip",            dts_pkthdr_get_srcip },
…
```

SMACQ Processing Modules

- Modules are the atoms of query optimization
- Written in C++ or Python
- Take arbitrary flags and arguments
  - Unix command-line style
- Introspection: Can ask runtime to identify downstream invariants
  - When module can do eager pre-filtering (e.g. hardware prefilter on NIC, database query, etc.)
- Event-driven (produce/consume) API
  - Can use “threaded” wrapper if lazy (really co-routines)
- Can embed other query instantiations
  - Can instantiate new scheduler, or share primary (preferred)

Example Processing Module (Python)

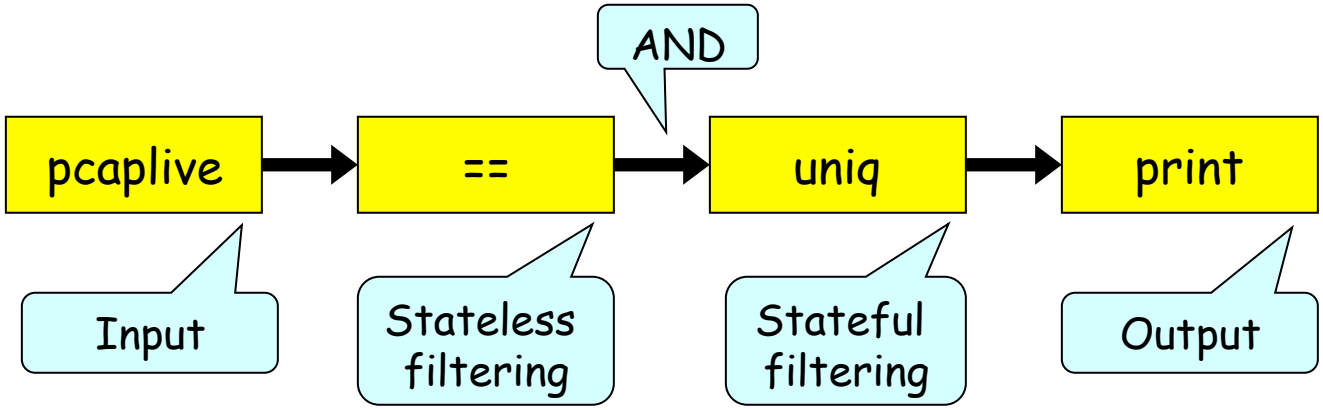
```
Class Dumper:
    """Print a few elements of each datum and pass every 5th"""
    def __init__(self, smacq, *args):
        print ('init', args)
        self.smacq = smacq      #Save reference to runtime
        self.buf = []           #List of objects received

    def consume(self, datum):
        for i in 'srcip', 'dstip', 'ipprotocol', 'len':
            v = datum[i].value
            print (i, datum[i].type, type(v), v)
        self.buf.append(datum)
        if len(self.buf) == 5:
            self.smacq.enqueue(datum)   # Output object downstream
            self.buf = []
```

Los Alamos NATIONAL LABORATORY EST.1943

Query Model: Dataflow Graphs

- Queries are dataflow graphs

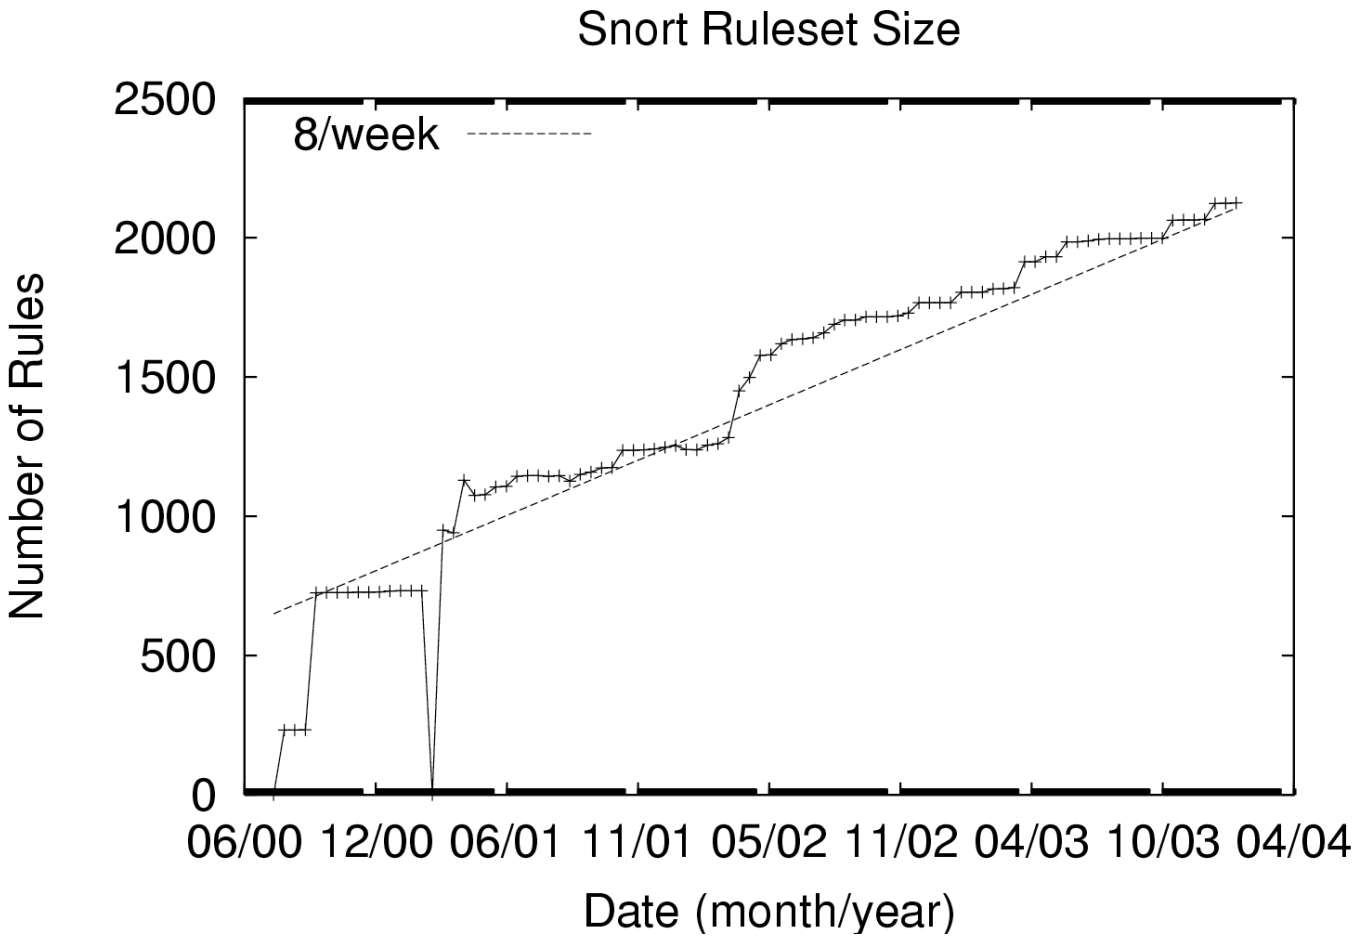


- Modules declare algebraic properties:
  - *stateless (map), annotation, vector, demux, (associative)*
  - Enables optimization, rewriting, parallelization, map/reduce
- Static optimizer applies all data-flow optimizations permitted by algebraic properties of the involved modules

Optimizing Continuous Queries

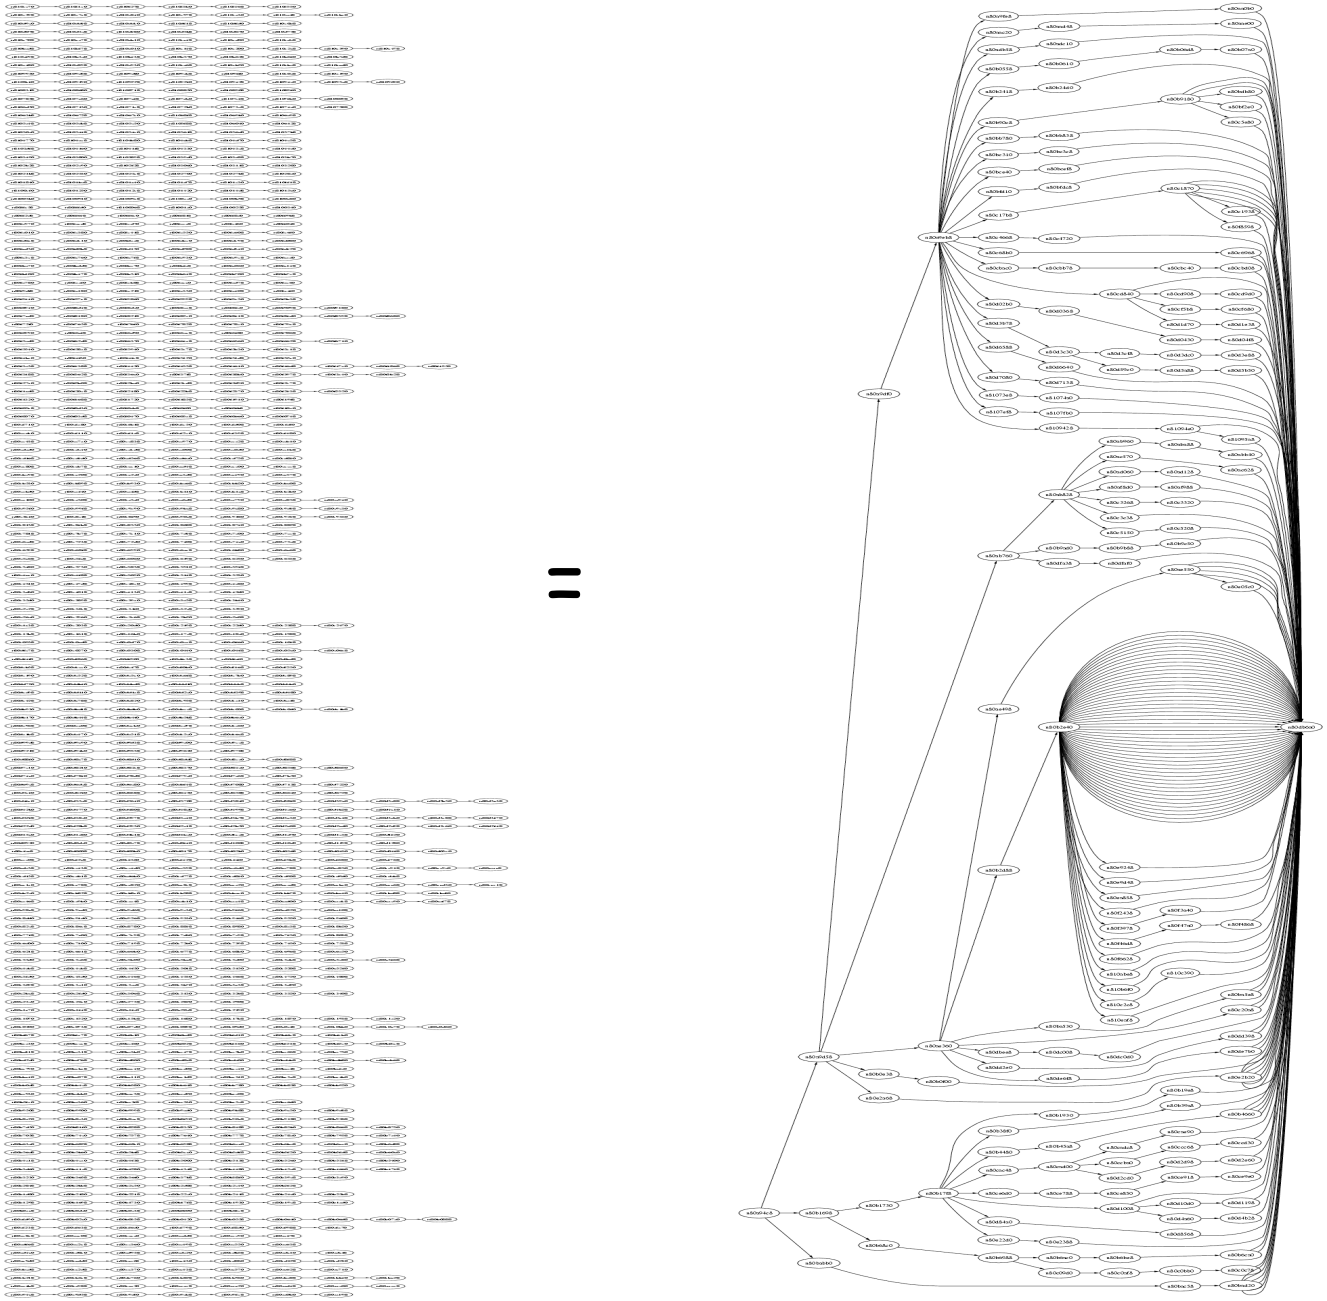
- Traditional database query optimization:
  - Uses data indexes
  - Minimizes individual query times
- Continuous-query optimization:
  - Executing many queries simultaneously
  - Minimize resource consumption per unit of data input
    - Maximize data throughput

Why is multiple query processing important? Approximately 8 new rules each week



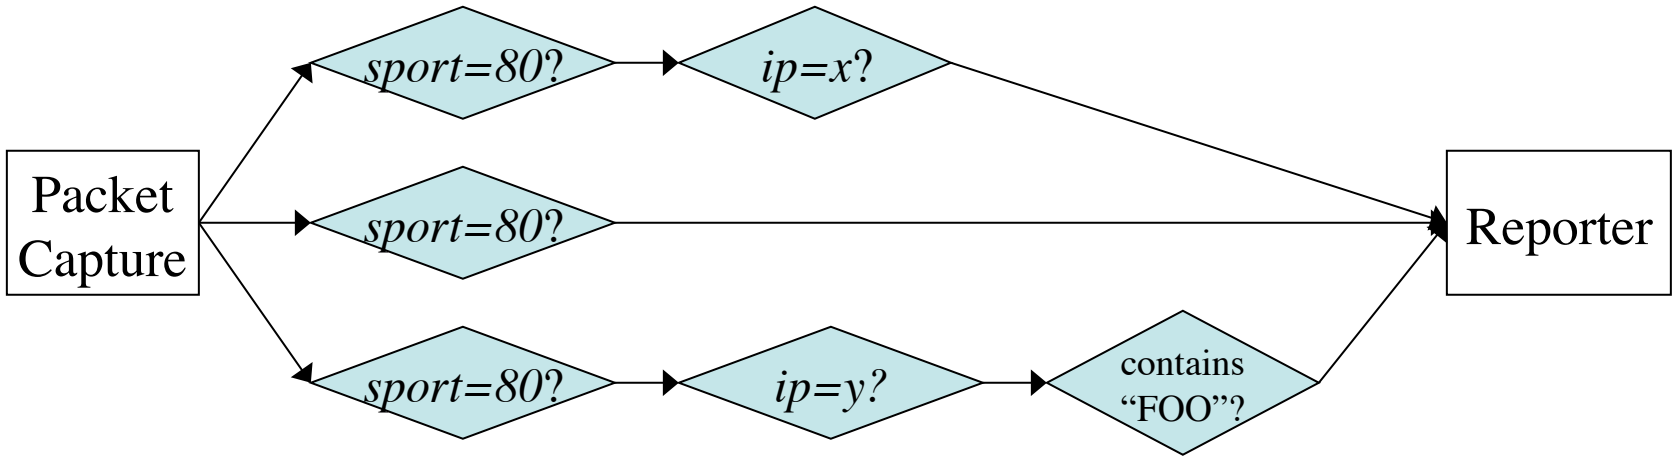
Los Alamos NATIONAL LABORATORY EST.1943

Optimization of 150 Snort Rules



Example Queries

6 Tests in 3 Rules

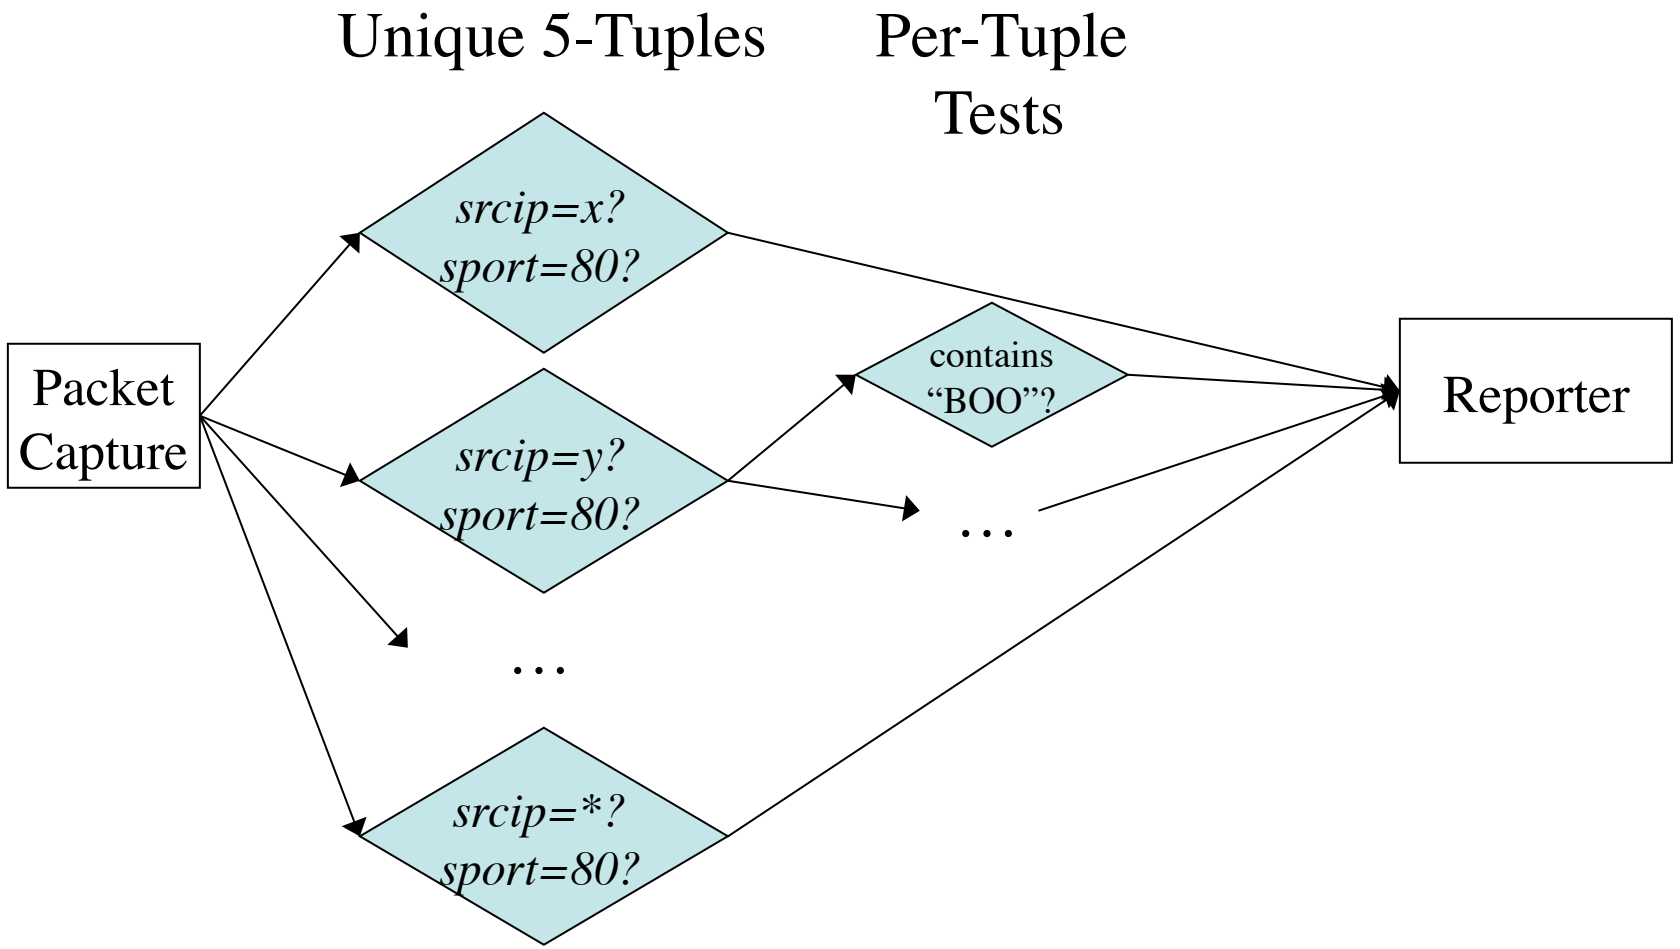


Los Alamos NATIONAL LABORATORY EST. 1943

Snort Approach

[Roesh, LISA 99]

Example: 6-7 Tests

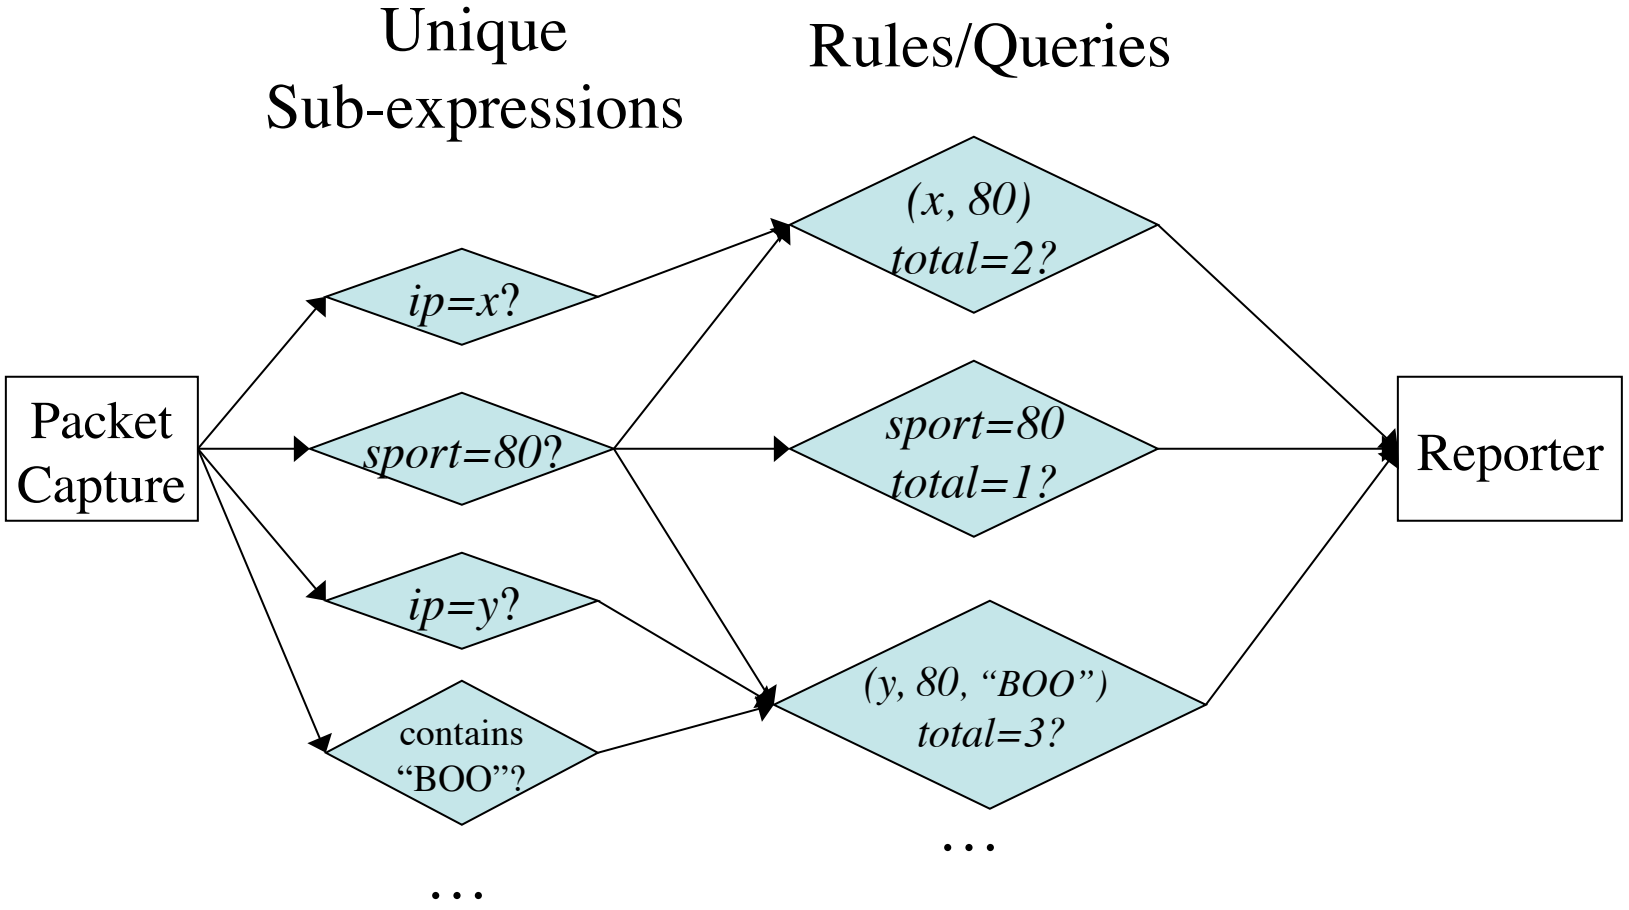


Los Alamos NATIONAL LABORATORY EST.1943

Counting Approach

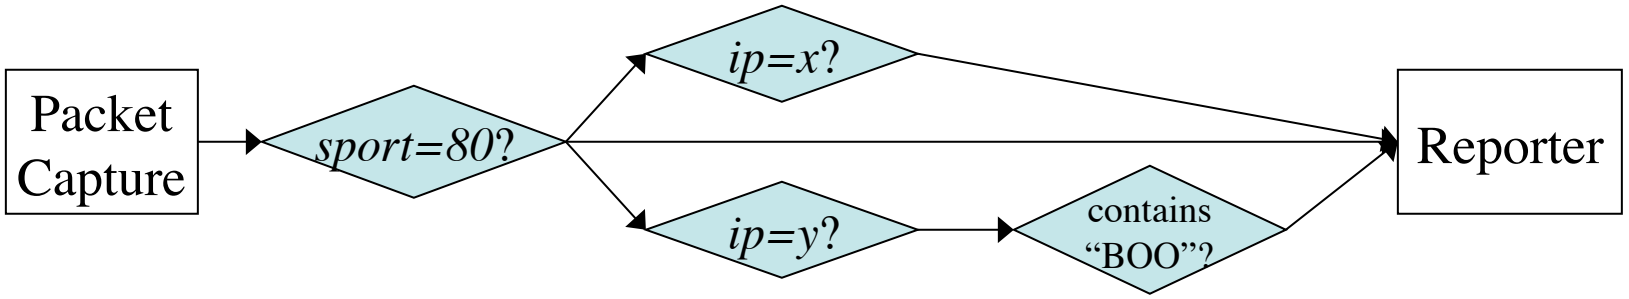
[*Carzaniga & Wolf, SIGCOMM 03*]

Example: 7 Tests



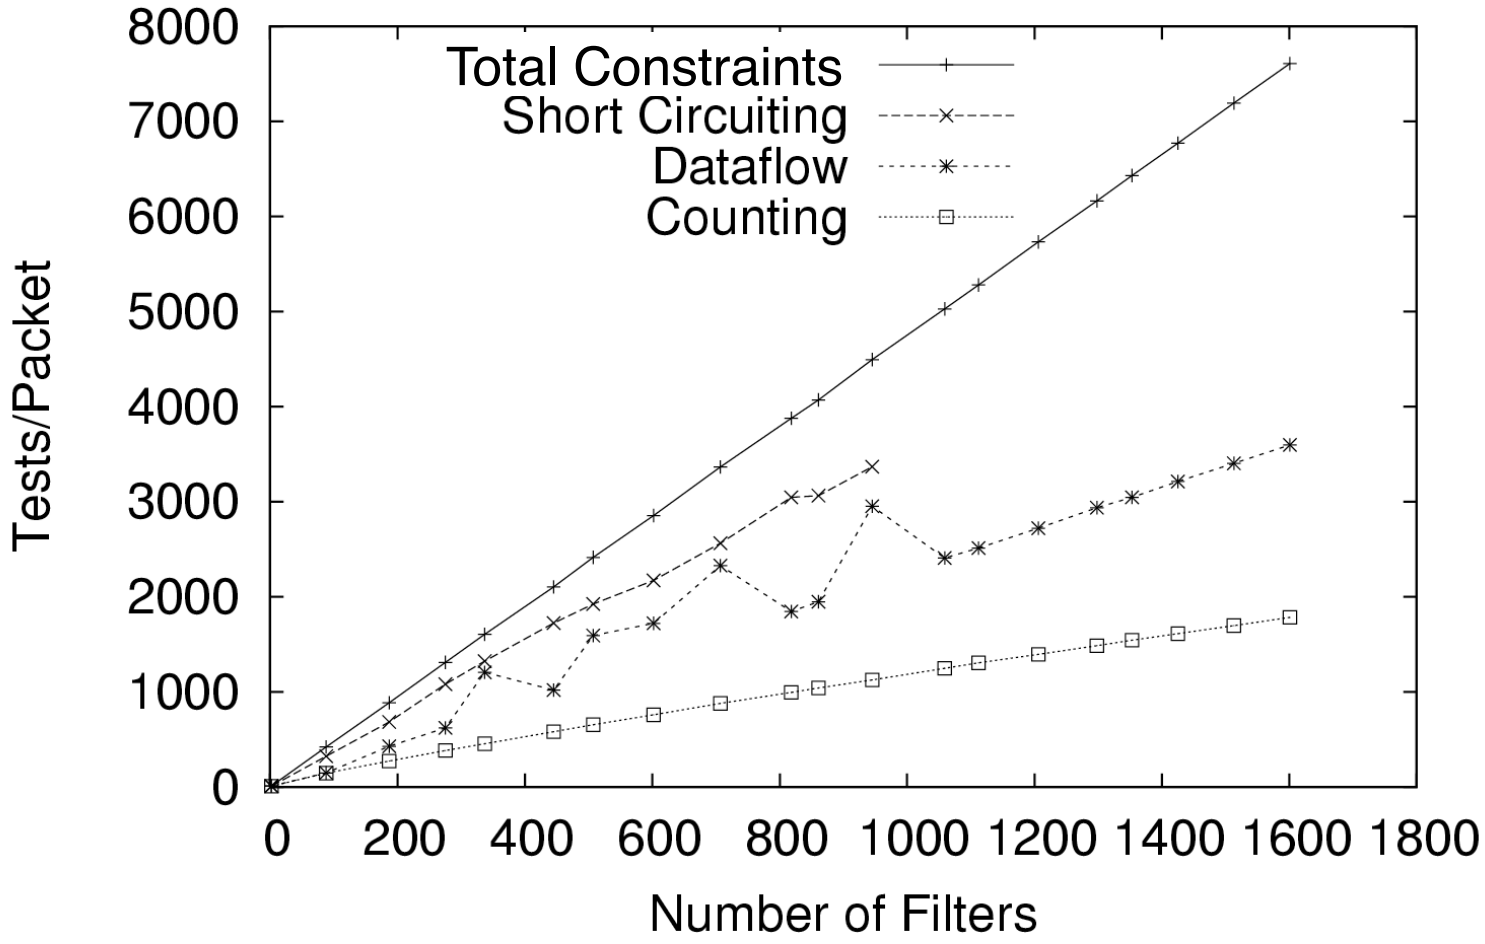
Data-Flow Approach

Example: 1-4 Tests



1. Common roots
2. Common leaves
3. Common upstream graphs
4. Common downstream graphs

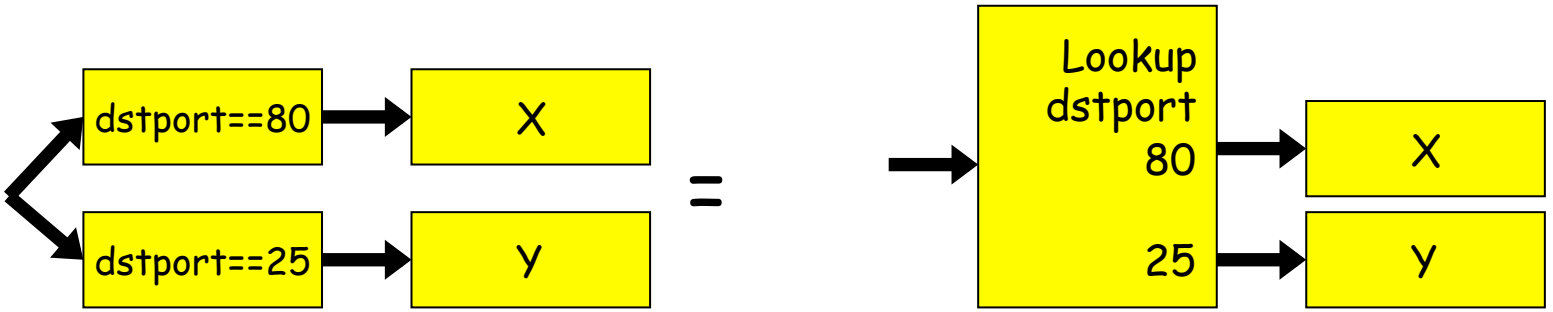
Performance Comparison



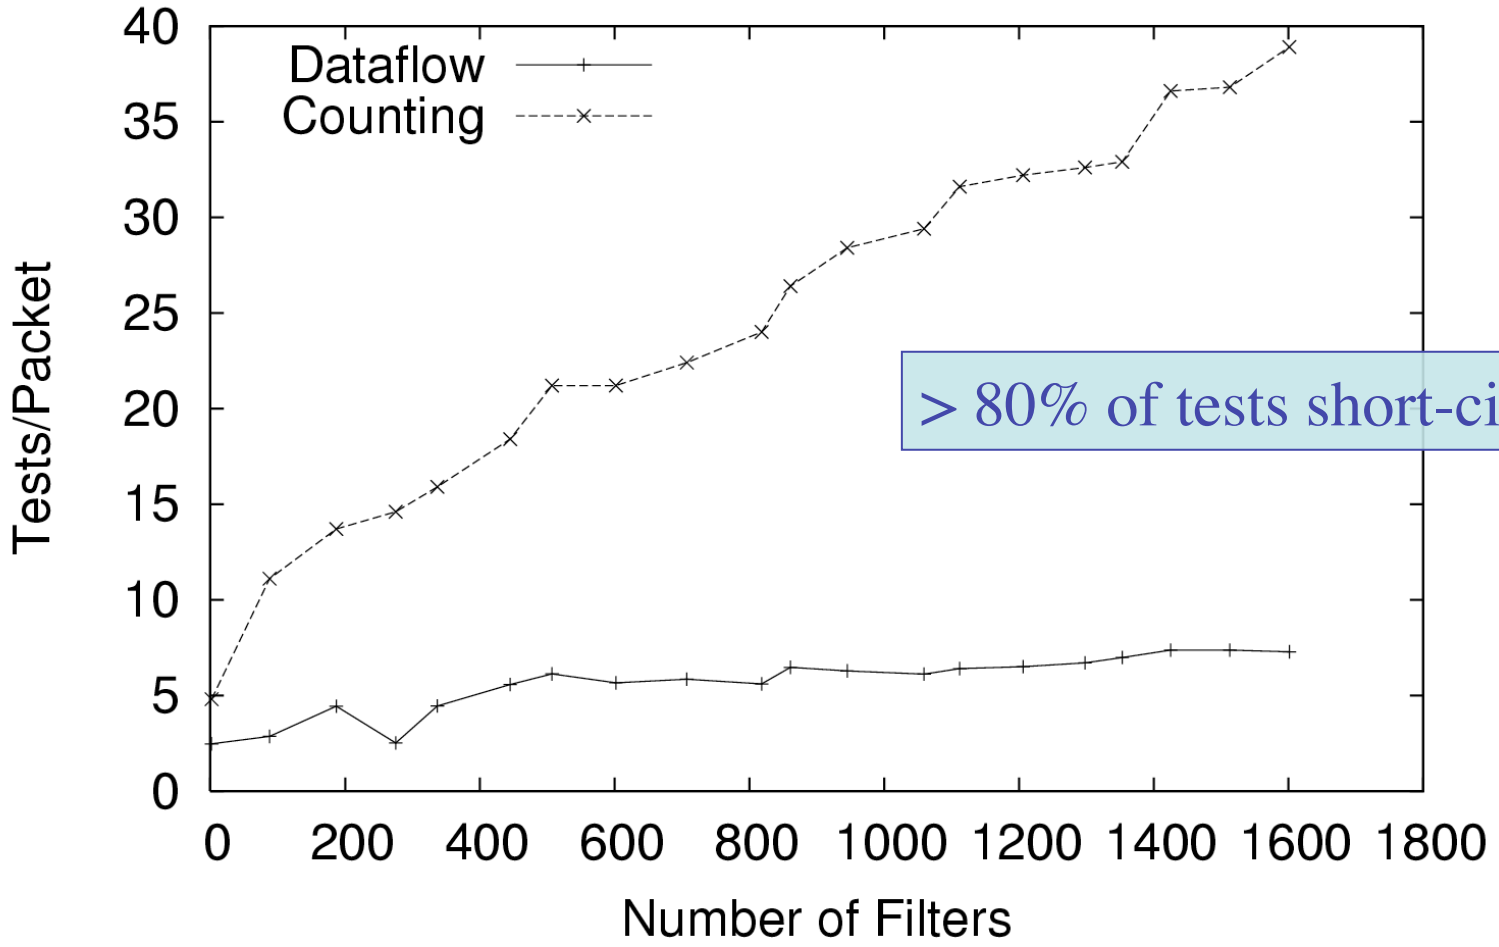
Los Alamos NATIONAL LABORATORY EST.1943

Vector Functions

- Most optimizations in stream analysis have employed a class of algorithms that can be characterized as *vector functions:*
  - $f(x, v) = f(x, v_1), f(x, v_2), \ldots$
  - Vector version is typically O(1) or O(log n) instead of O(n)
- Examples
  - Set of equality tests becomes a single lookup in a hash-table
  - Set of string matches becomes a single DFA to traverse



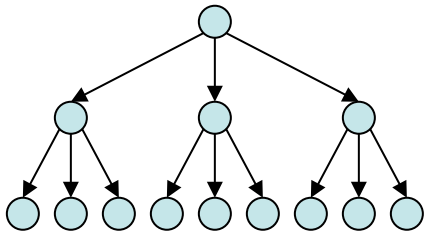
Performance Comparison with Vector Functions



> 80% of tests short-circuited

Los Alamos NATIONAL LABORATORY EST.1943

Analysis: Why was Counting better only without vectors?



- Assume that each test results in *p* more tests
  - $p = fanout \bullet short\text{-}circuiting$
  - $p \leq fanout$
  - $0 \leq short\text{-}circuiting \leq 1$
- Assume data-flow of tests is a balanced tree of depth *d*
  - *d* is an integer ≥ 1
- Expected number of evaluations:

  $1 + p + p^2 + p^3 + \ldots + p^{d-1} \quad = \quad (1 - p^d) / (1 - p)$
- Let *u* = number of unique tests = Counting's performance

  $s(1 - p^d) / (1 - p) < u$ if $(d > 1, p < 1)$
- For IDS test: *d* = 6
  - With Vectors (u=39): *p* < 1.7 is desired. Actual *p* = 1
  - Without Vectors (u=1782): p < 4.2 is desired. Actual *p* = 5.8

Los Alamos NATIONAL LABORATORY EST.1943

Supported Query Languages

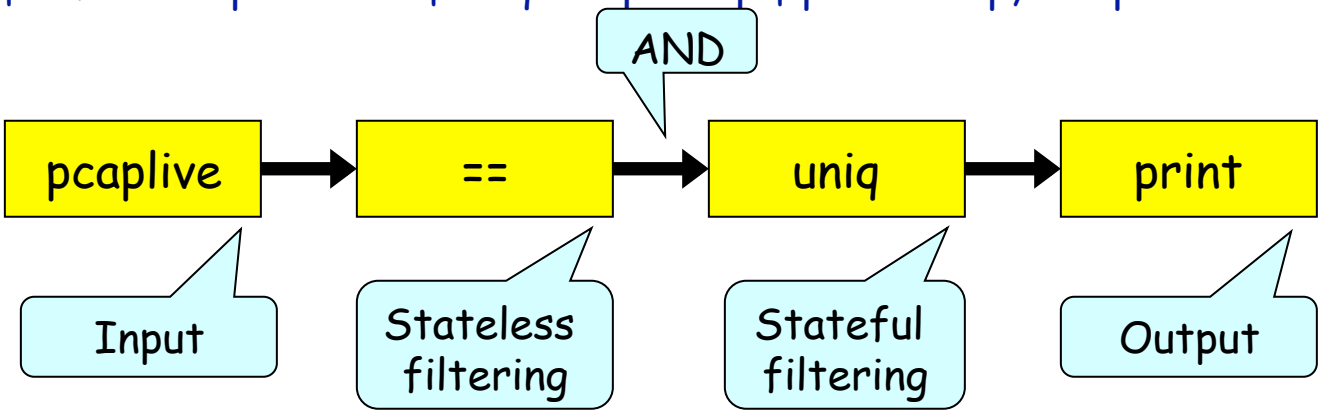
- SQL style:

```
print srcip, dstip from
    (cflow where dstport==80 and uniq(srcip, dstip))
```

  - Misplaced belief that since SQL is well defined, people can just use it
  - Deeply nested queries make you wish you were merely nested in s-expressions

- Unix pipe style:

```
cflow | where dstport==80 | uniq srcip dstip | print srcip, dstip
```



Los Alamos NATIONAL LABORATORY EST. 1943

Supported Query Languages

- Datalog

```
Pairs    :- cflow | uniq(srcip dstip)
SrcCount :- count() group by ipprotocol srcport
DstCount :- count() group by ipprotocol dstport
Pdf      :- filter(count) | pdf
Print    :- sort(-r probability) | print(type ipprotocol port probability)

Pairs | SrcCount | const(-f type src) | Pdf | rename (srcport port) | Print
Pairs | private | DstCount | const(-f type dst) | Pdf | rename (dstport port) | Print
```

  - Clean, allows named subexpressions

Los Alamos NATIONAL LABORATORY EST.1943

Join Models

- DFA module
  - Define a state machine where transitions specified as Booleans on new inputs
- SQL style
  - Example: print running cross-product

```
print a.ipid b.ipid from
    pcapfile(0325@1112-snort.pcap) a, b where a.ipid != b.ipid
```

  - New keyword UNTIL defines when state can be removed
    - "NEW" refers to newly input data for comparison
  - Example: print retransmissions within the same second

```
print expr(b.ts - a.ts) from pcaplive() a until(new.a.ts.sec > a.ts.sec), b until(new)
   where b.ts > a.ts and a.srcip == b.srcip and a.srcport == b.srcport
           and a.seq == b.seq and a.payload != "" and b.payload != ""
```

Usage Experience

- Online detection & automated response systems
- Ad-hoc queries for forensic analysis and data exploration
- Feature extraction for other software

Conclusions

- Continuous Queries provide a common query syntax, software infrastructure, and optimization framework for traffic analysis
- CQ necessary for streaming applications, sufficient for ad-hoc forensic analysis

Open source at SMACQ.SF.NET

Los Alamos NATIONAL LABORATORY EST.1943

Conclusions

Open source at
SMACQ.SF.NET

- Continuous Queries provide a common query syntax, software infrastructure, and optimization framework for traffic analysis
- Two identified strategies for static optimization of multiple queries
  - Remove (Counting) or Reduce (Data-flow) redundant tests
  - Boolean (Data-flow) short-circuiting removes need for some subsequent tests
- Performance Analysis:
  - Counting is preferable when short-circuiting is rare
  - Data-flow out-performs counting when short-circuiting is significant
    - When breadth of graph is reduced with vector functions, actual IDS workload benefits significantly from short-circuiting
- Data-flow approach can also benefit from additional, dynamic reordering of tests to maximize early short-circuiting



Strip Plots: A Simple Automated Time-Series Visualization

Sid Faber
sfaber@cert.org

© 2010 Carnegie Mellon University

Overview

Motivation & Goals

Sample Output

The Basics

Special Features

© 2010 Carnegie Mellon University

NO WARRANTY

THIS MATERIAL OF CARNEGIE MELLON UNIVERSITY AND ITS SOFTWARE ENGINEERING INSTITUTE IS FURNISHED ON AN “AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this presentation is not intended in any way to infringe on the rights of the trademark holder.

This Presentation may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

© 2010 Carnegie Mellon University

Motivation and Goals

© 2010 Carnegie Mellon University

Caveat

This is analyst code, not engineering code

- Your mileage may vary

© 2010 Carnegie Mellon University

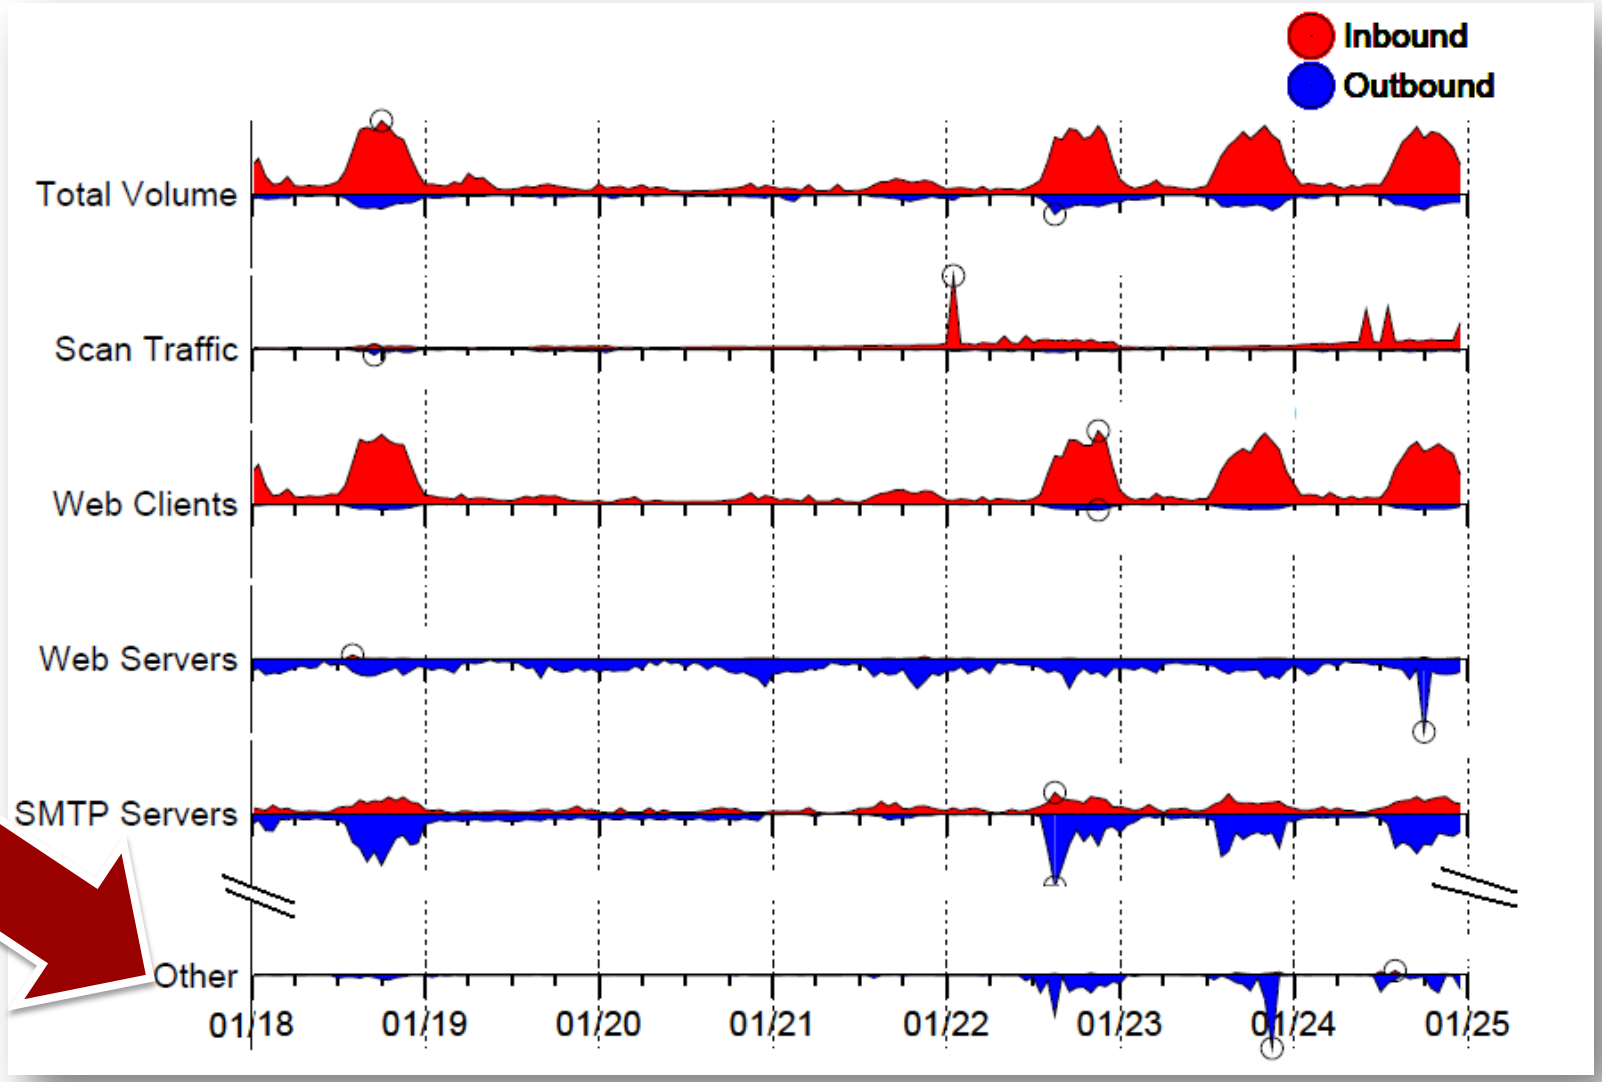
Motivation

Support network profiling for Situational Awareness

- I know most of what's on my network
  - Based exclusively on past observations
- I can filter / categorize out routine traffic
- What can I do with "leftover" traffic?
  - Is this something new to add to my profile?
  - Is something odd happening on my network?

© 2010 Carnegie Mellon University

Example



© 2010 Carnegie Mellon University

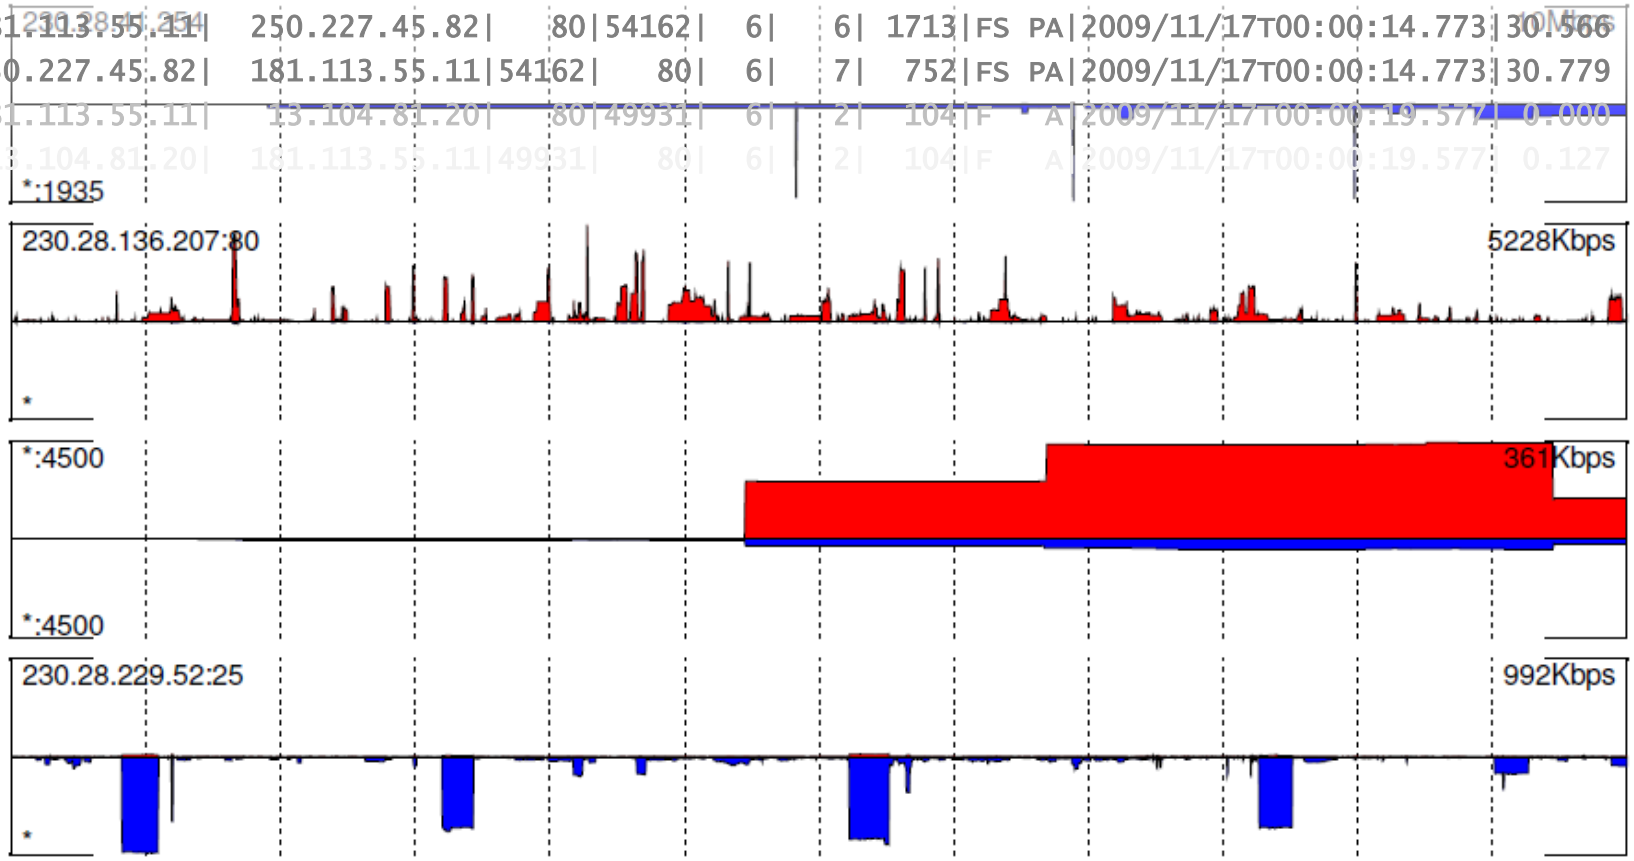
Goals

Nearly self-maintaining network profile

- Batch processing
- Email delivery
- ***Quick*** triage of “leftovers”
  - Add to my profile?
  - Something odd?
  - <5 minutes per report
- Self-sufficient description: ***No Additional Explanation Necessary***

© 2010 Carnegie Mellon University

sIP	dIP	sPort	dPort	pro	pkts	bytes	flags	sTime	dur
0.33.6.97	181.113.55.11	63220	80	6	2	104	F A	2009/11/17T00:00:01.236	0.071
181.113.55.11	0.33.6.97	80	63220	6	2	104	F A	2009/11/17T00:00:01.236	0.000
181.113.55.11	94.145.72.124	80	50151	6	8	9016	S PA	2009/11/17T00:00:07.233	0.232
92.4.182.32	181.113.55.11	1067	80	6	20	1822	PA	2009/11/17T00:00:09.827	2.125
181.113.55.11	92.4.182.32	80	1067	6	25	32910	PA	2009/11/17T00:00:09.840	1.713
92.4.182.32	181.113.55.11	1056	80	6	20	1868	PA	2009/11/17T00:00:09.841	2.684
181.113.55.11	92.4.182.32	80	1056	6	24	30819	PA	2009/11/17T00:00:09.842	2.355
139.173.53.191	181.113.55.11	9429	80	6	1	40	R A	2009/11/17T00:00:12.299	0.000
181.113.55.11	184.199.238.201	80	33559	6	33	41996	FS PA	2009/11/17T00:00:12.866	1.222
184.199.238.201	181.113.55.11	33559	80	6	28	1878	FS PA	2009/11/17T00:00:12.866	1.222
181.113.55.11	250.227.45.82	80	54162	6	6	1713	FS PA	2009/11/17T00:00:14.773	30.566
250.227.45.82	181.113.55.11	54162	80	6	7	752	FS PA	2009/11/17T00:00:14.773	30.779
181.113.55.11	13.104.81.20	80	49931	6	2	104	F A	2009/11/17T00:00:19.577	0.000
13.104.81.20	181.113.55.11	49931	80	6	2	104	F A	2009/11/17T00:00:19.577	0.127

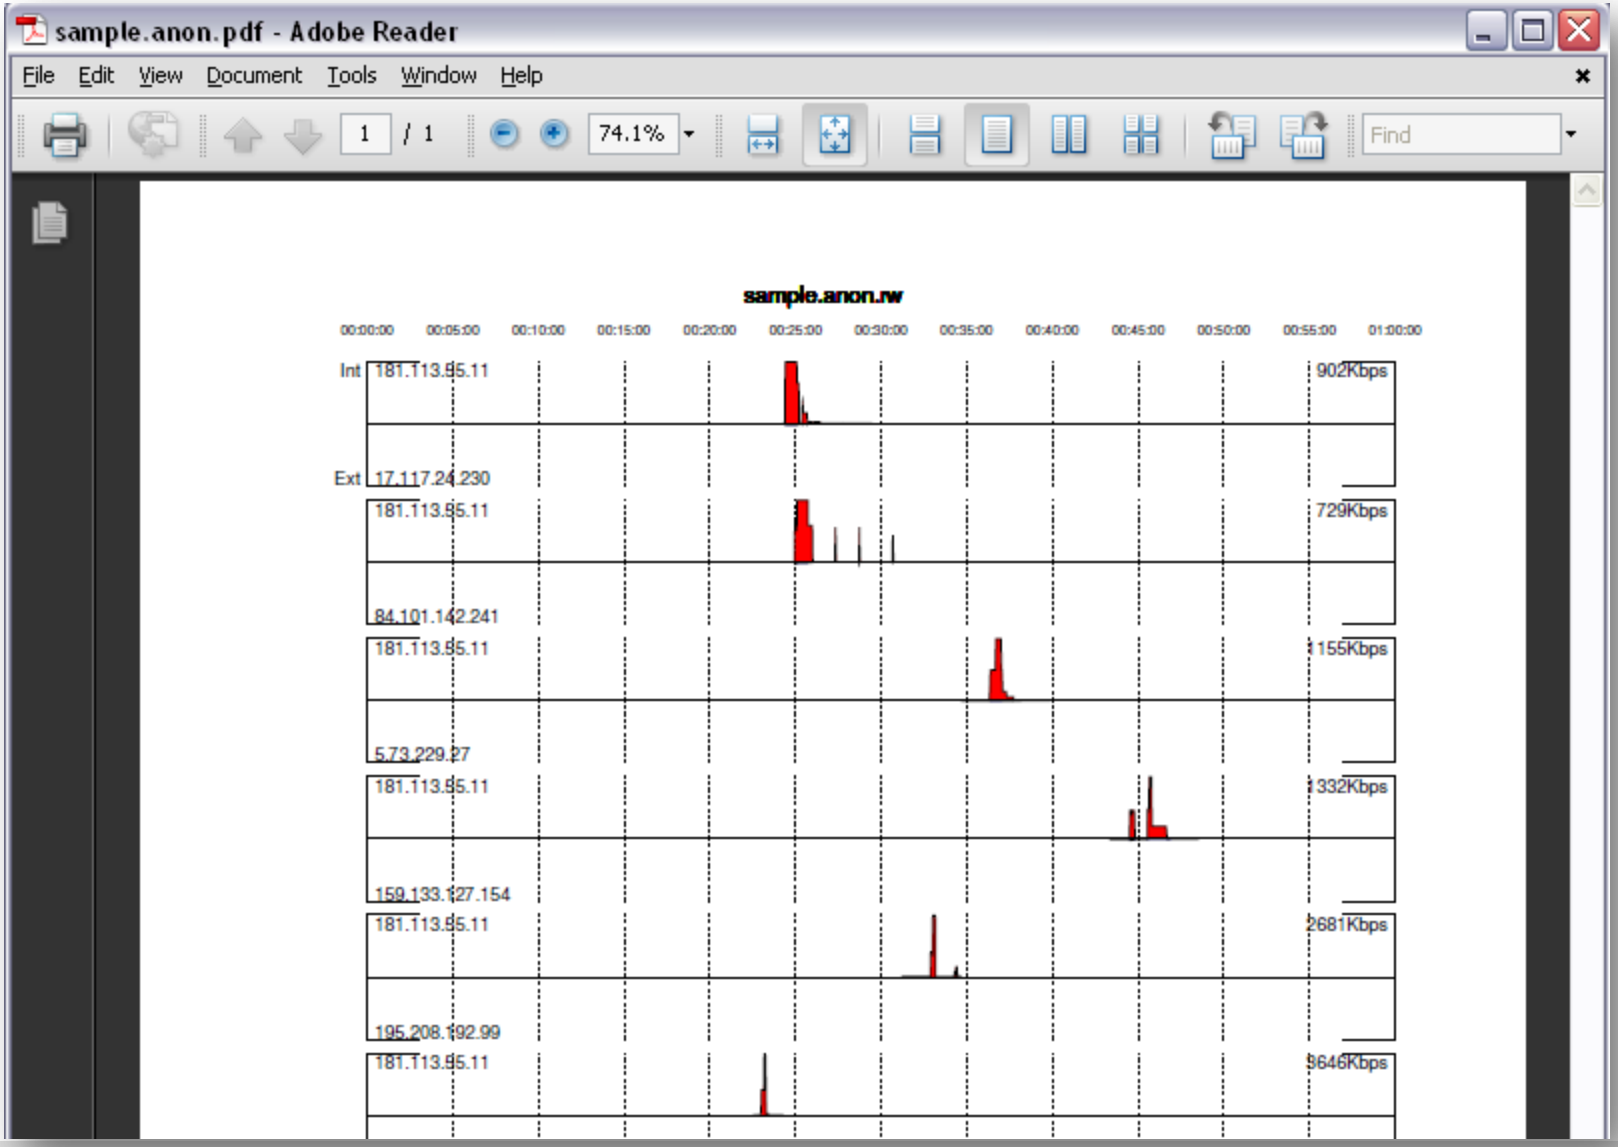


© 2010 Carnegie Mellon University

Sample Output

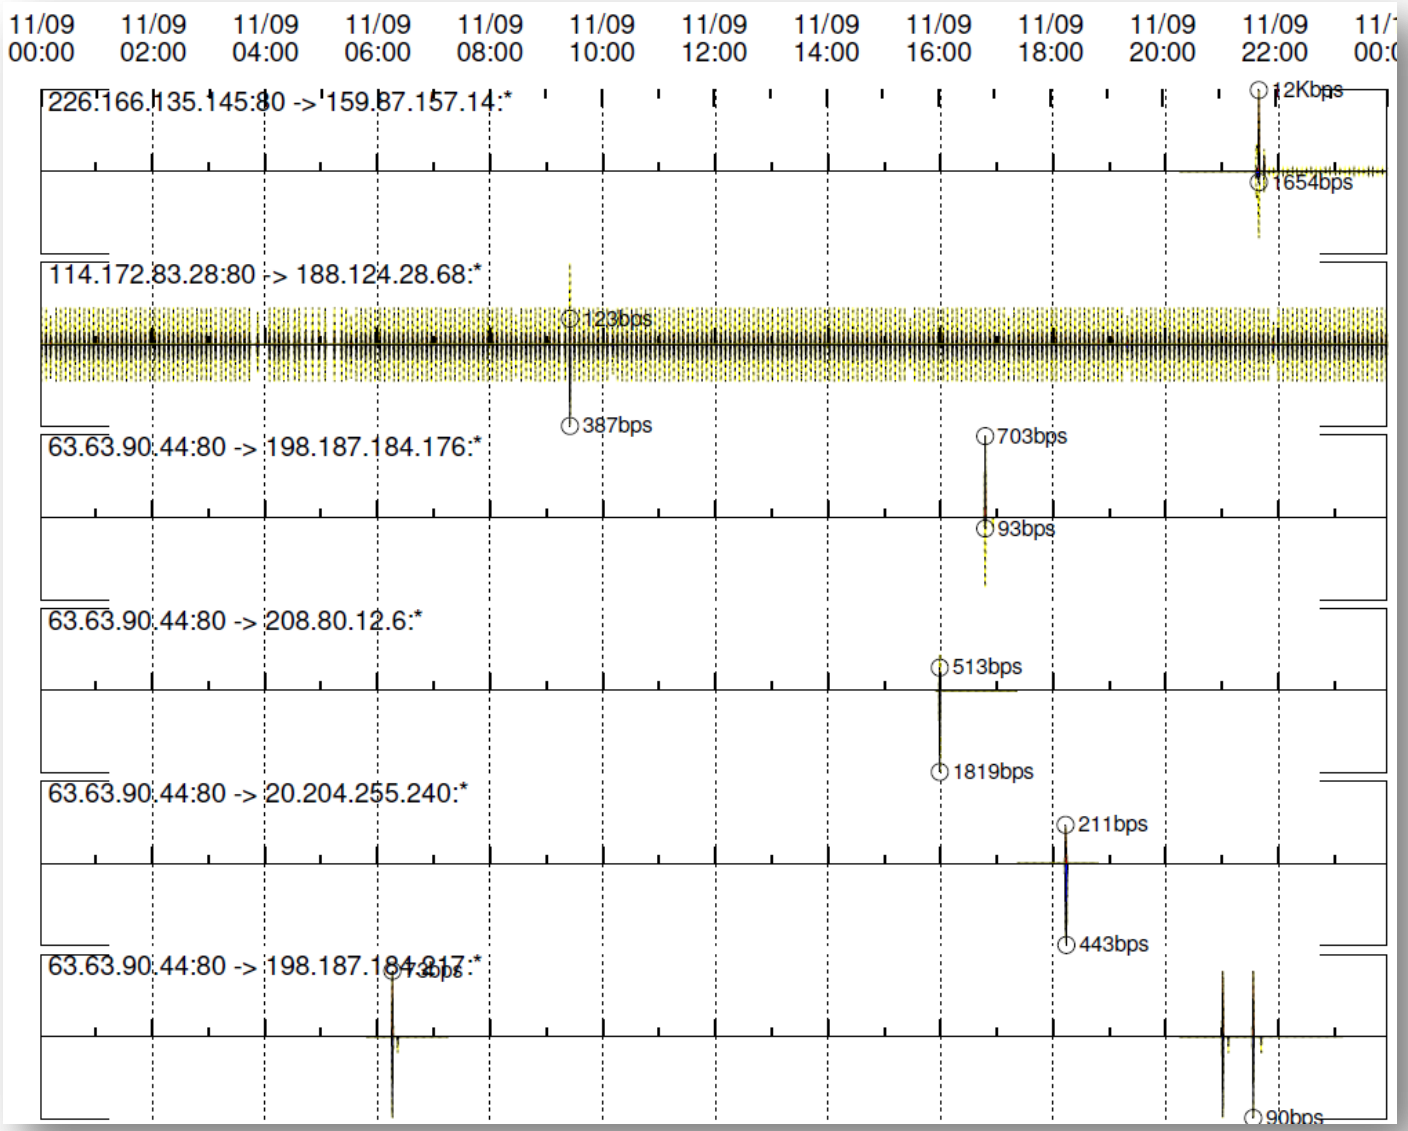
© 2010 Carnegie Mellon University

One Hour of Standard HTTP Traffic



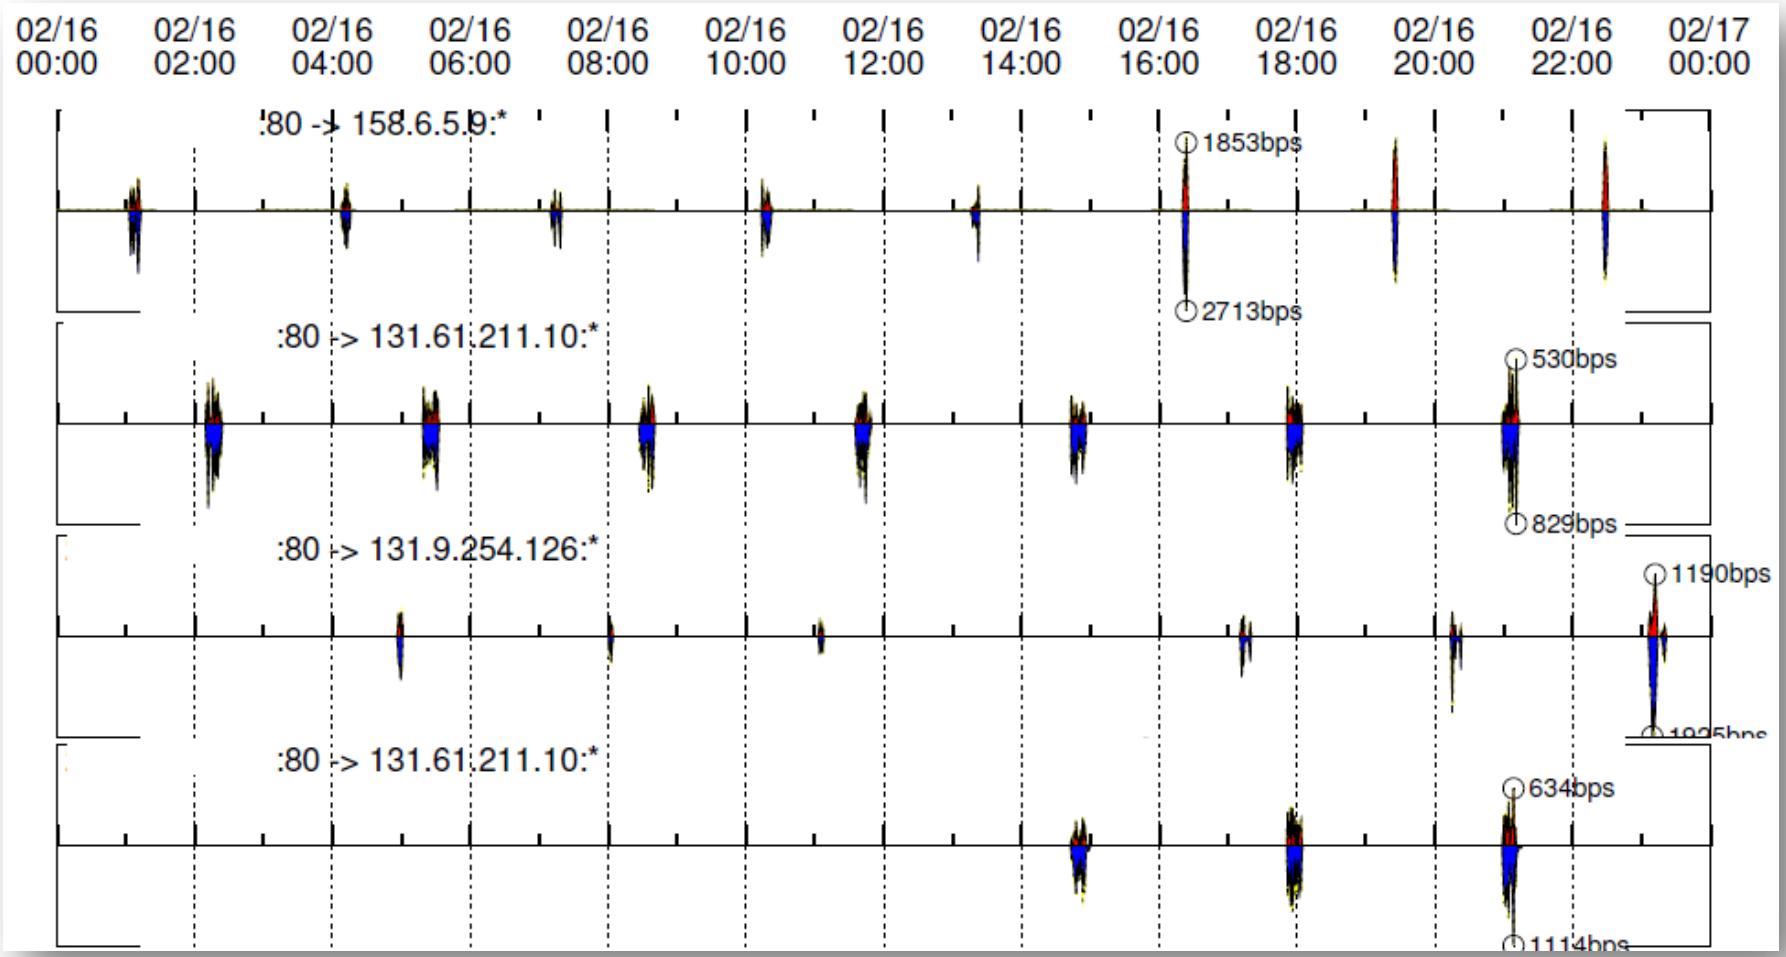
© 2010 Carnegie Mellon University

More http: can you spot the beacon?



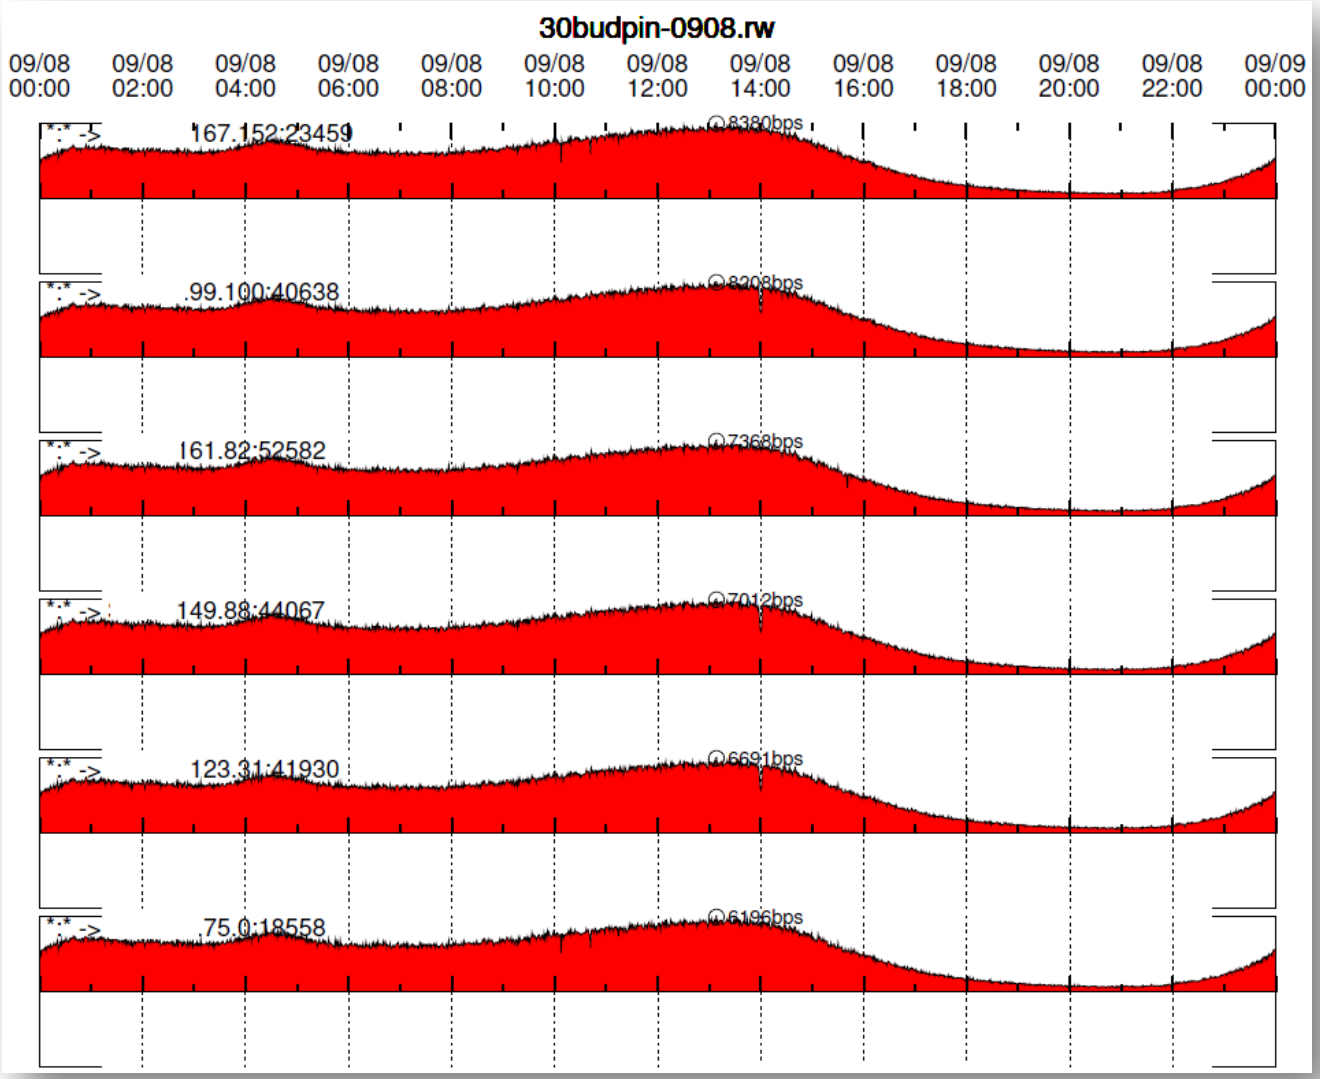
© 2010 Carnegie Mellon University

Conficker



© 2010 Carnegie Mellon University

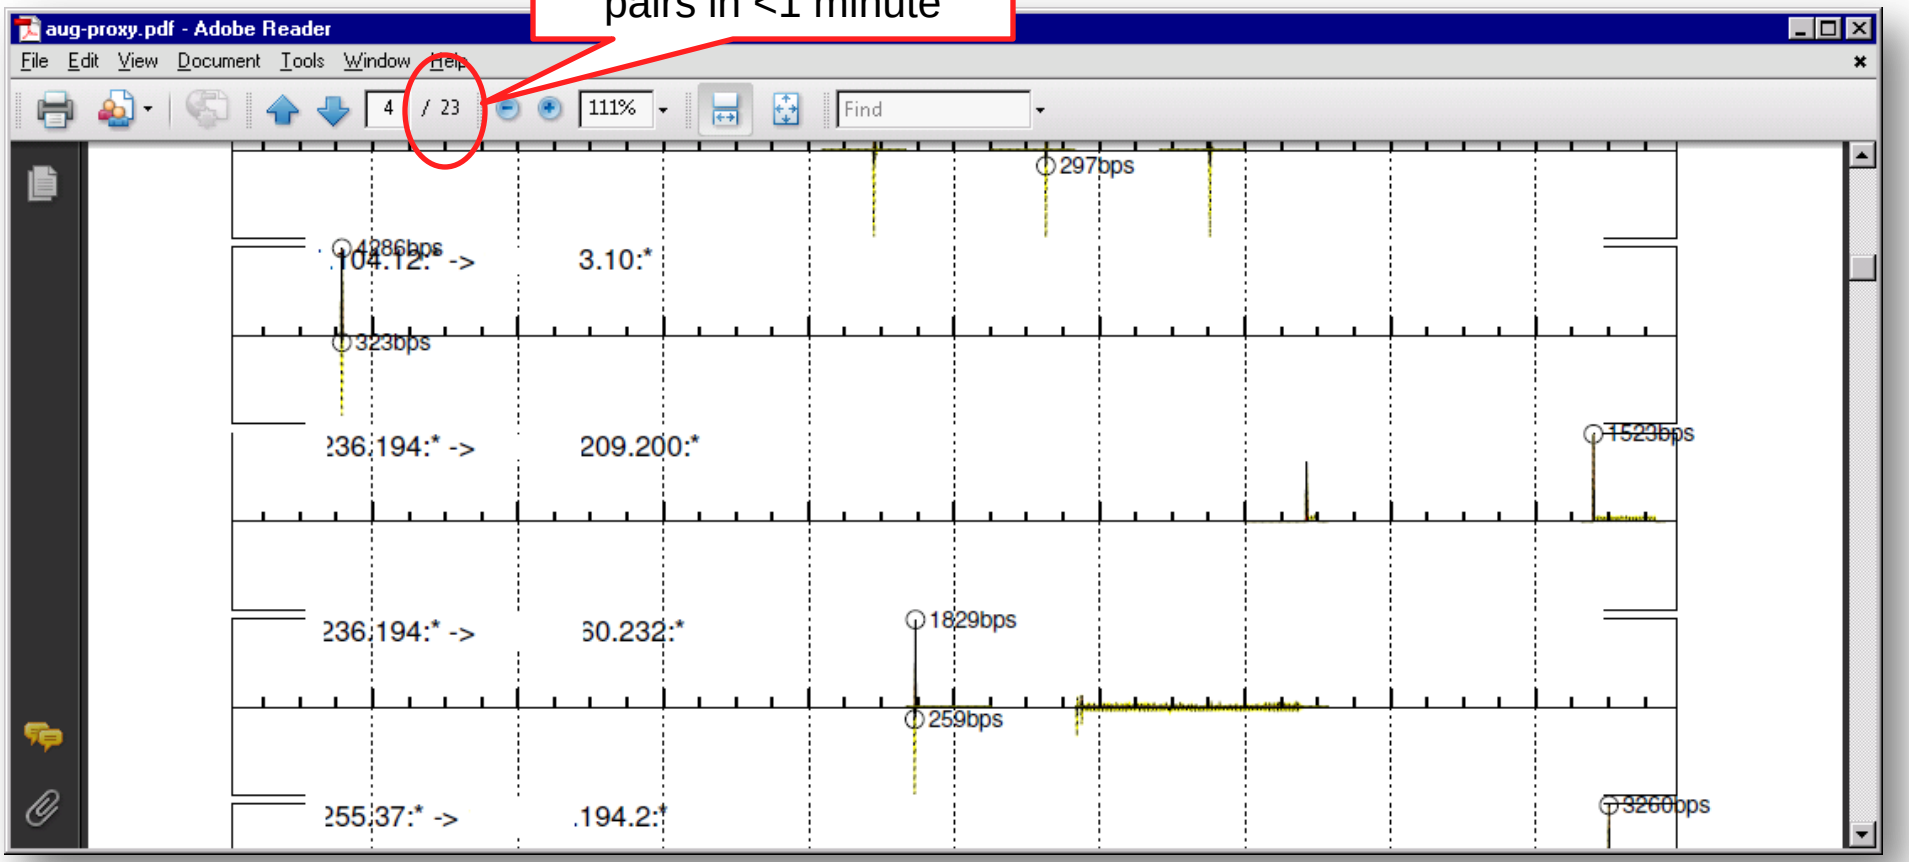
30-byte UDP packets



© 2010 Carnegie Mellon University

More web…

23 pages: Review traffic for 230 sip/dip pairs in <1 minute



© 2010 Carnegie Mellon University

The Basics

© 2010 Carnegie Mellon University

Basic script workflow

Find the top talkers

Count in & out traffic for each talker

Plot traffic for each

Compile plots into a single .pdf

© 2010 Carnegie Mellon University

Dependencies

SiLK toolset

gnuplot

ghostscript (gs)

Python (but not pySiLK)

© 2010 Carnegie Mellon University

./stripplot.py

```
 $ stripplot.py
Usage: stripplot.py [options] FILE
Creates a strip-plot of the most significant traffic contributors
within a raw SiLK data file

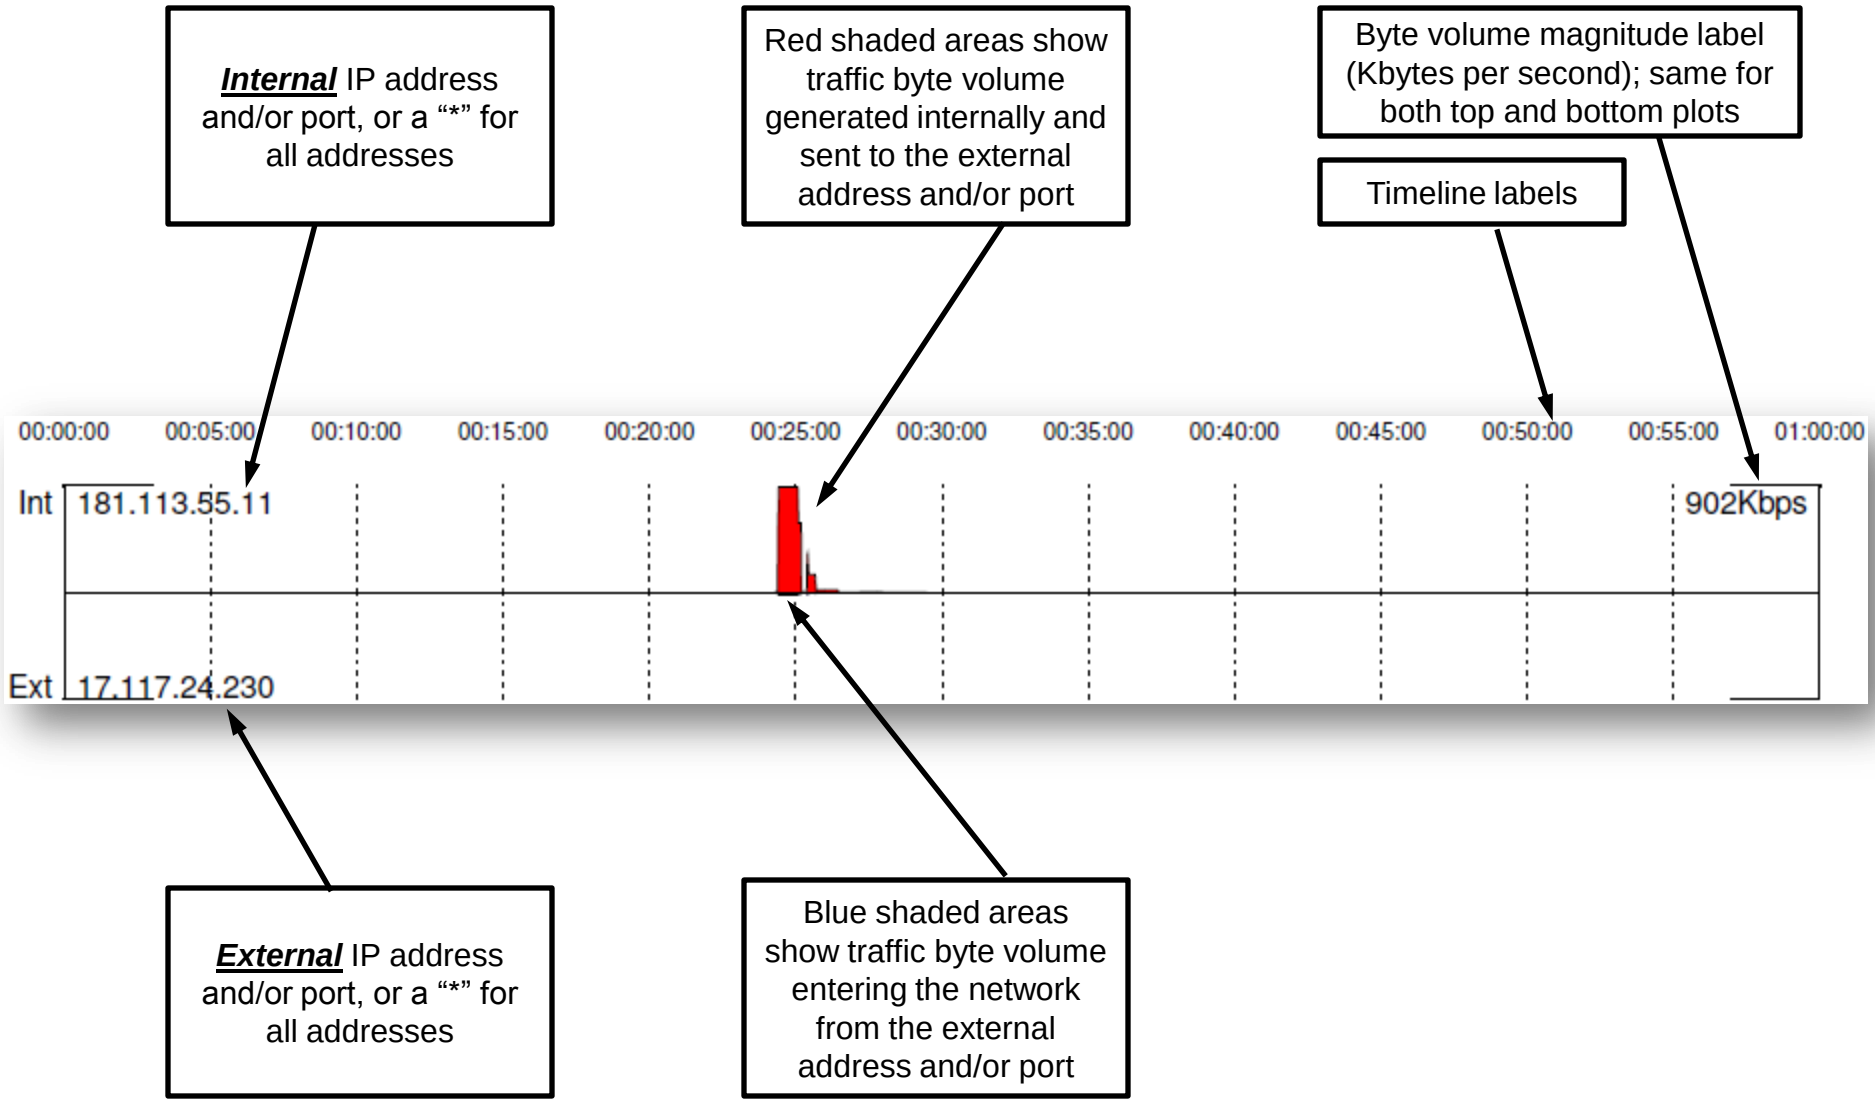
  FILE   (required) The binary SiLK file to analyze

  Options:
  --binsize   Default counting bin size (seconds) (def=auto)
  --bottomleft Bottom left tag; allows %(substitution)s; use '-' for none (def
              =-)
  --bottommiddle Words on the bottom right of the page, allows %(substitution)
              s; use '-' for none (def=FOUO)
  --bottomright Words on the bottom middle of the page, allows %(substitution)
              s; use '-' for none (def=Page %(page)i of %(pagecount)s)
  --count     Number of plots to output (def=5)
  --endtime   Plot end time, YYYY/MM/DDTHH:MM:SS (def=auto)
  --fields    rwuniq-style list of fields to group IN traffic on, or '*' for a
              utomatic (NOT ALL FIELDS WORK) (def=sip,dip)
  --flags     Include this option to add a plot of TCP flags to the strips (de
              f=0)
  --help      Print this output (def=)
  --types     Inbound and outbound types; these are used to make sure the IN a
              ddress is on the top plot; must be in the form [in-type/out-type
              ]; unspecified types work fine but either address may end up on
              the top. (def=in/out,inweb/outweb,inicmp/outicmp)
  --pdffilepath PDF final output file (def=tmp.pdf)
  --plotfile  Temporary gnuplot script file to create (def=tmp.plot)
  --plotsperpage Number of plots per page (def=10)
  --prefilter rwfilter expression to apply to flow file before selecting what
              to plot.  NOTE:  this filter is NOT applied to the trends themse
              lves, only to the selection routine (def=--proto=0-)
  --psfilepath Post-script file to generate (def=tmp.ps)
  --selectionval Choose the top [count] combinations to plot based on this val
              ue; must be either 'bytes', 'packets', 'flows' or 'none'); if 'n
              one' then output is in rwuniq (random) order (def=bytes)
  --topleft   Words on the upper left of the page, allows %(substitution)s; us
              e '-' for none (def=-)
  --topright  Words on the upper right of the page, allows %(substitution)s; u
              se '-' for none (def=-)
  --trendline Highlighted and dotted trendline to add to plot; f for flows, p
              for packets (def=b)
  --starttime Plot start time in YYYY/MM/DDTHH:MM:SS format (def=auto)
  --maintitle Title for this plot, allows %(substitution)s; use '-' for none (
              def=auto)
  --verbose   Print out debugging info, use twice for more info and to print d
              ebugging info on the plot itself (def=0)

Fields with string substitution support the following:
  %(page)i    Current page number
  %(date)s    Date the report was printed
  %(time)s    Time the report was printed
  %(pagecount)i Total number of pages in the report
  %([setting])s Any of the report configuration settings
              (run with -v option to see settings
```


© 2010 Carnegie Mellon University

An individual strip



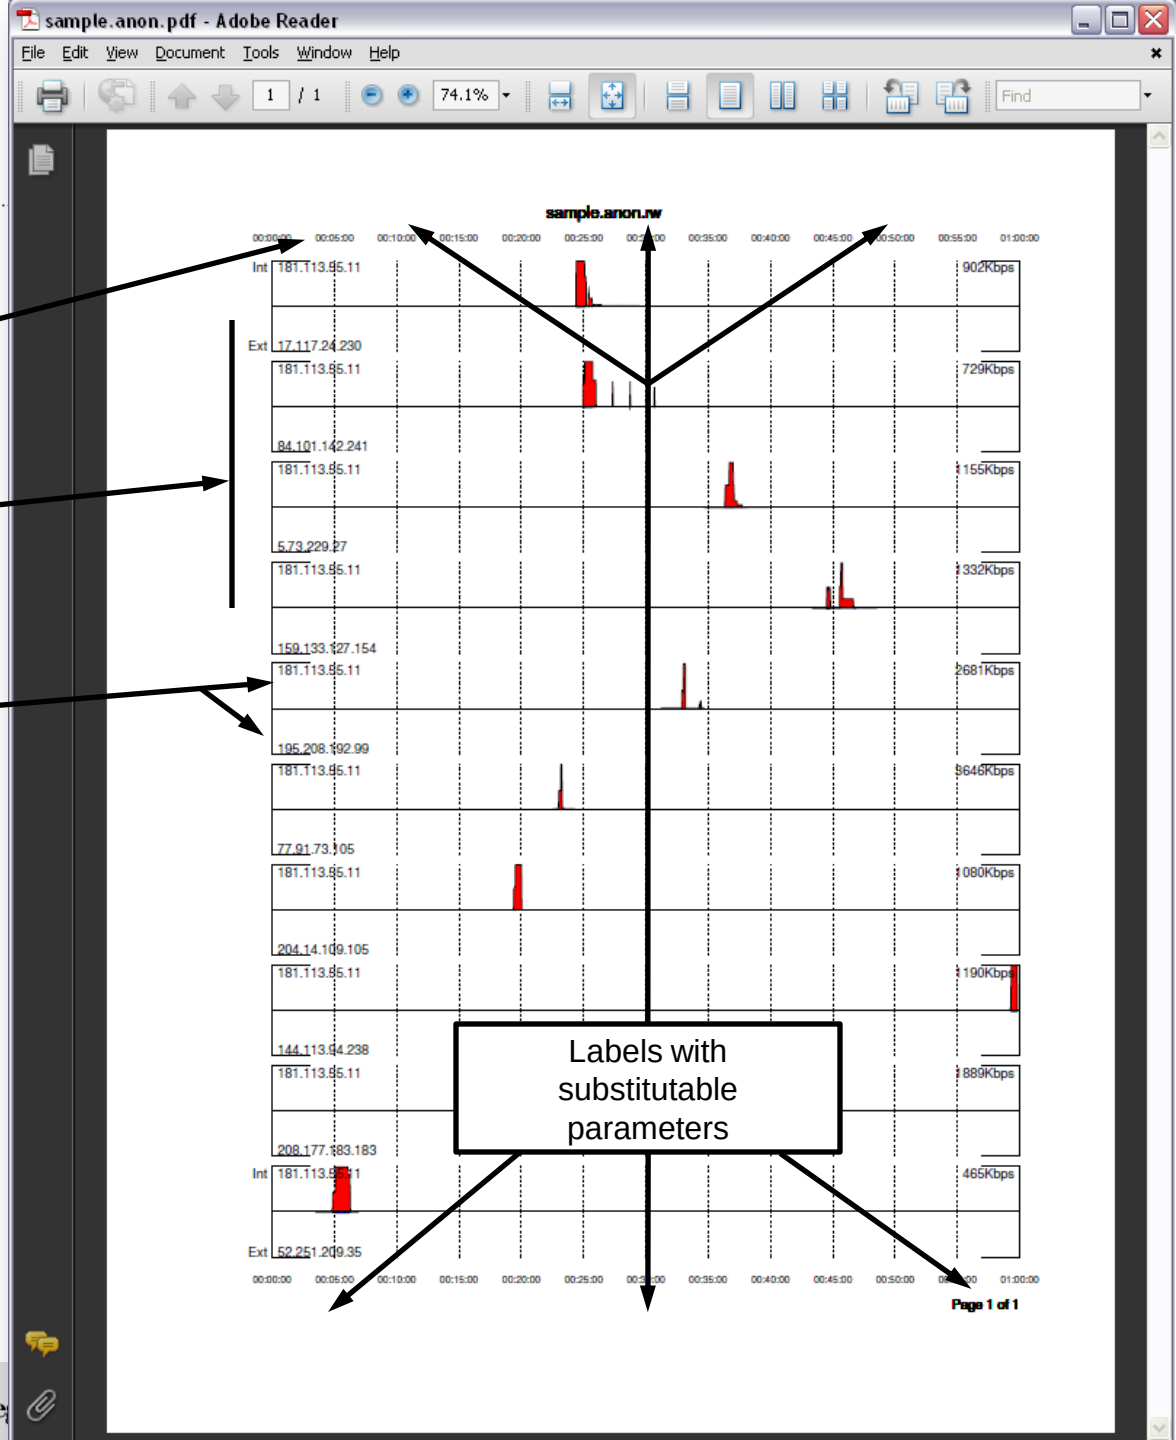
© 2010 Carnegie Mellon University

The Full Page

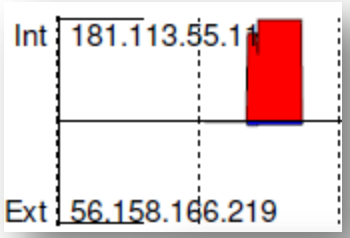
Time scale/formatting courtesy of gnuplot

Consistent timeline throughout

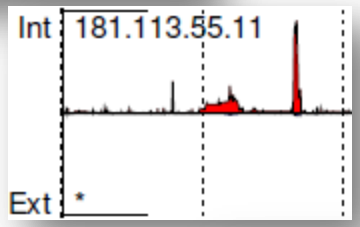
Each strip has an independent magnitude, but top and bottom are the same magnitude



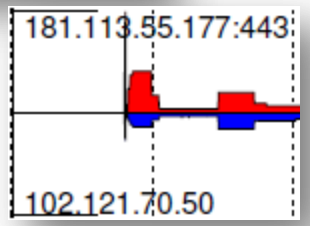
--fields=[rwuniq field set] defines groupings



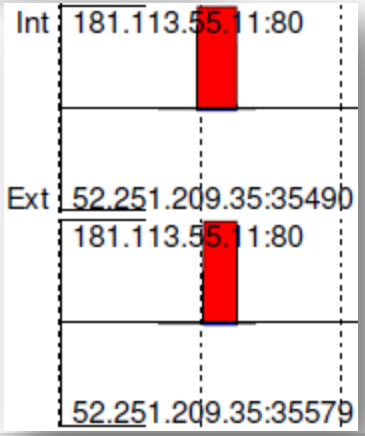
--fields=sip/dip (default): group on client/server pairs



--fields=dip: group on dest (my) address



--fields=sip,dip,dport: group on external services



--fields=sip,dip,sport,dport: group on TCP sessions

© 2010 Carnegie Mellon University

Q: How does it choose what strips to plot?

A: It uses rwuniq to group records, then chooses the largest by byte volume:

```
rwuniq [file] \
   --fields=[field list] \
   | sort -nr -k [bytes column]
```


© 2010 Carnegie Mellon University

Tweaking the selection criteria

Sometimes you want to show the top packets or flows

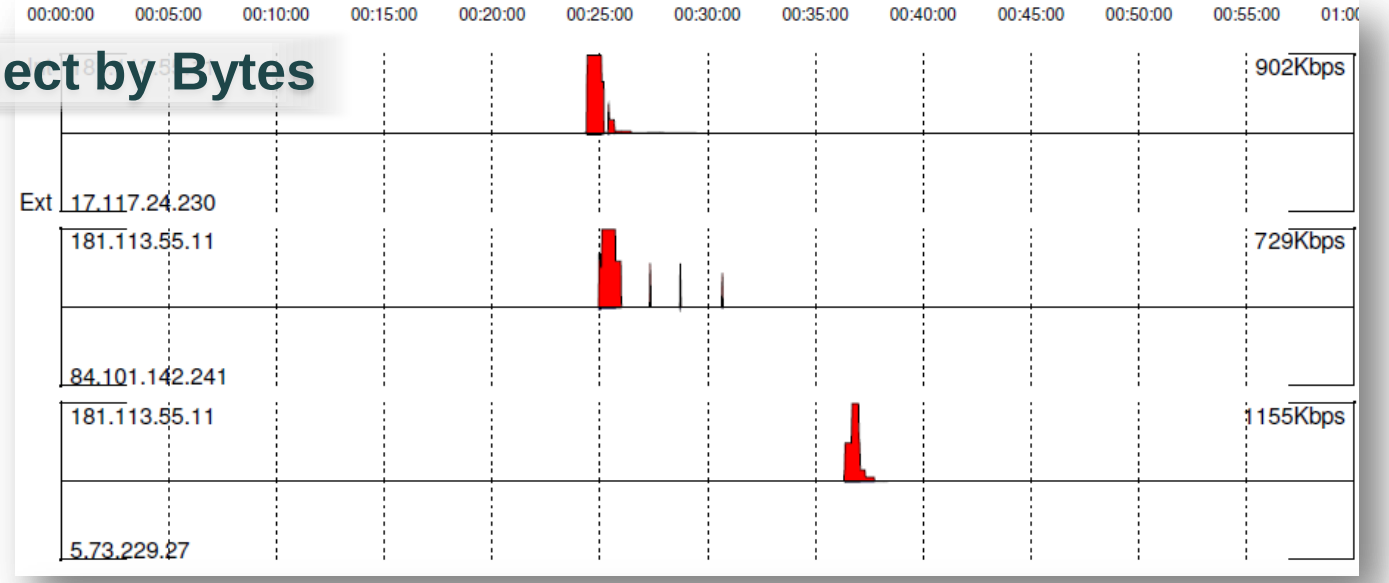
- Repeated failed connection attempts
- Bot phone-home

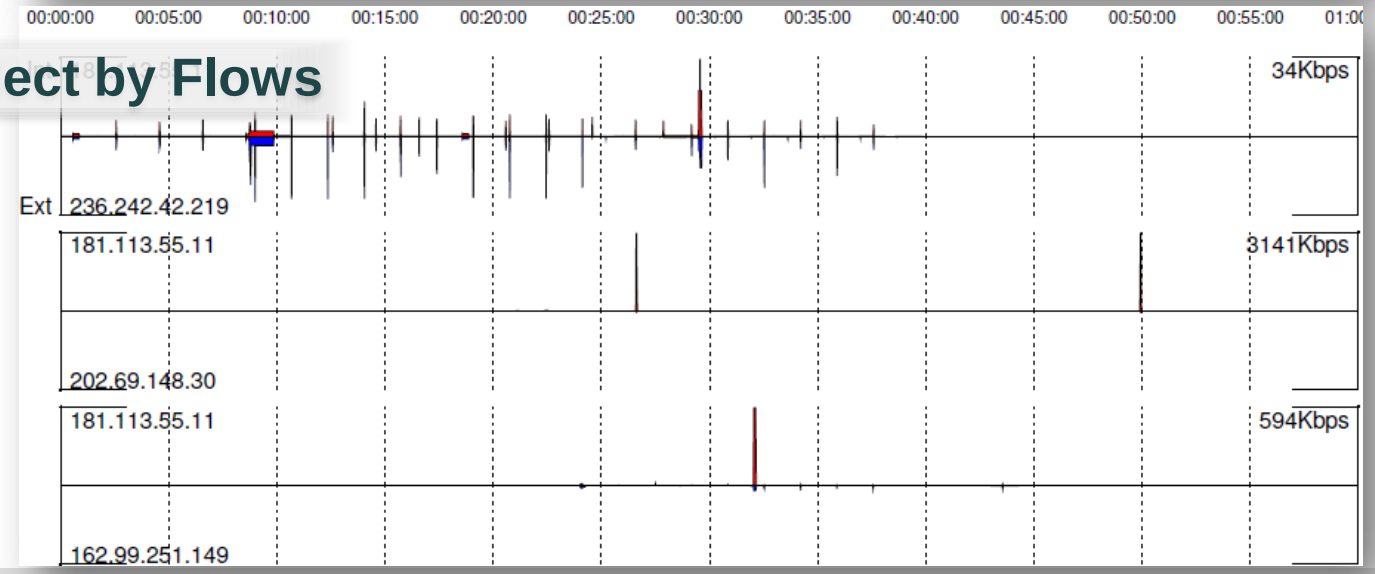
`--selectionval=flows|packets`

- Sort by top flow count or packet count

© 2010 Carnegie Mellon University

selectionval example





© 2010 Carnegie Mellon University

Other Common Options

--count

- Defines how many strips to plot, defaults to 5

--prefilter

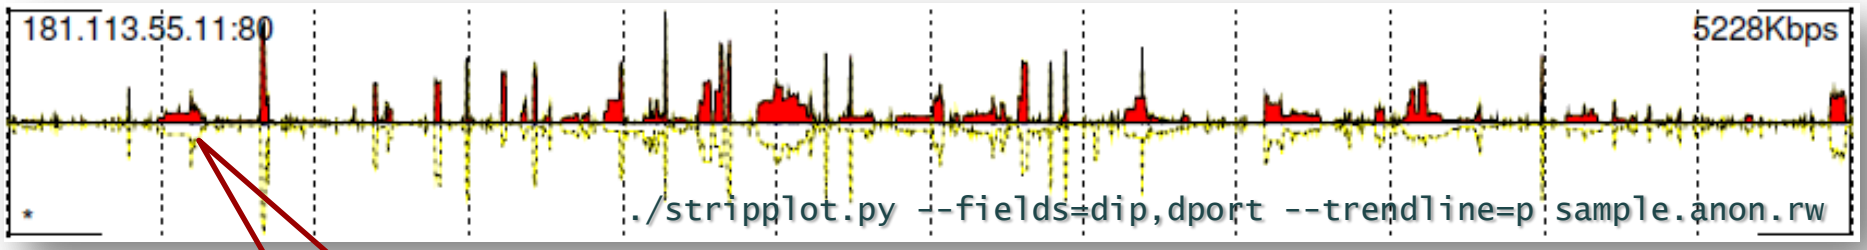
- Filter out the .rw file before plotting it

© 2010 Carnegie Mellon University

Trend lines

--trendline=f [or p]

- Adds a highlighted dotted line for flows (or packets)
- No labeling, no agnitudes, not symmetric
- Good for drawing out low-volume data



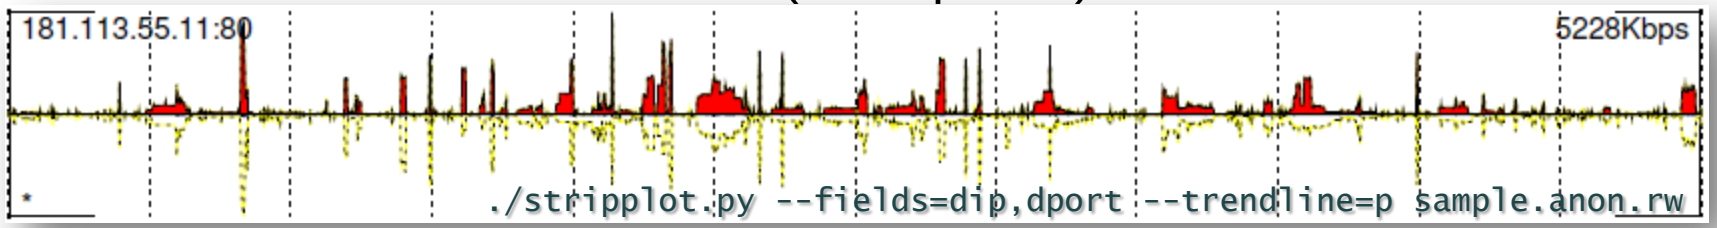
The trend line show ACK packet counts generally match data packets

© 2010 Carnegie Mellon University

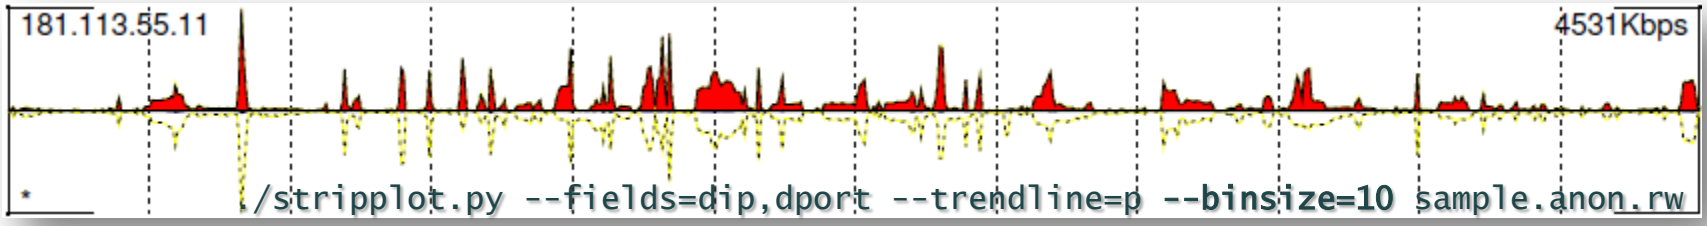
Bin size

Points on the time axis:

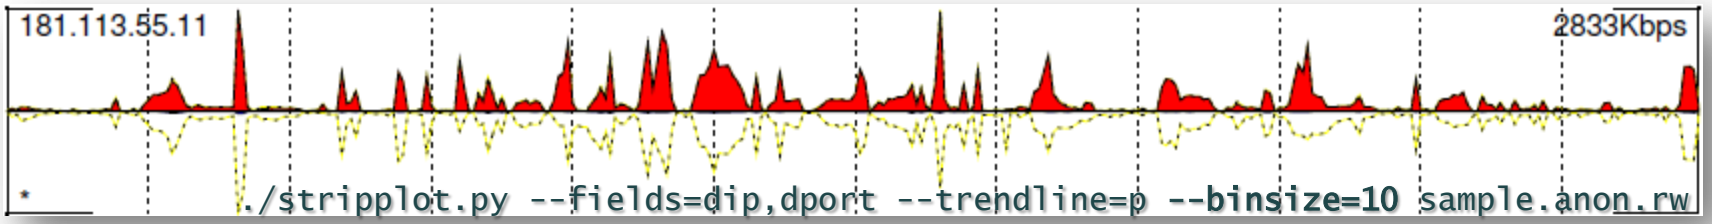
One hour defaults to 1 second bins (3600 points):



5 second bins makes little difference:

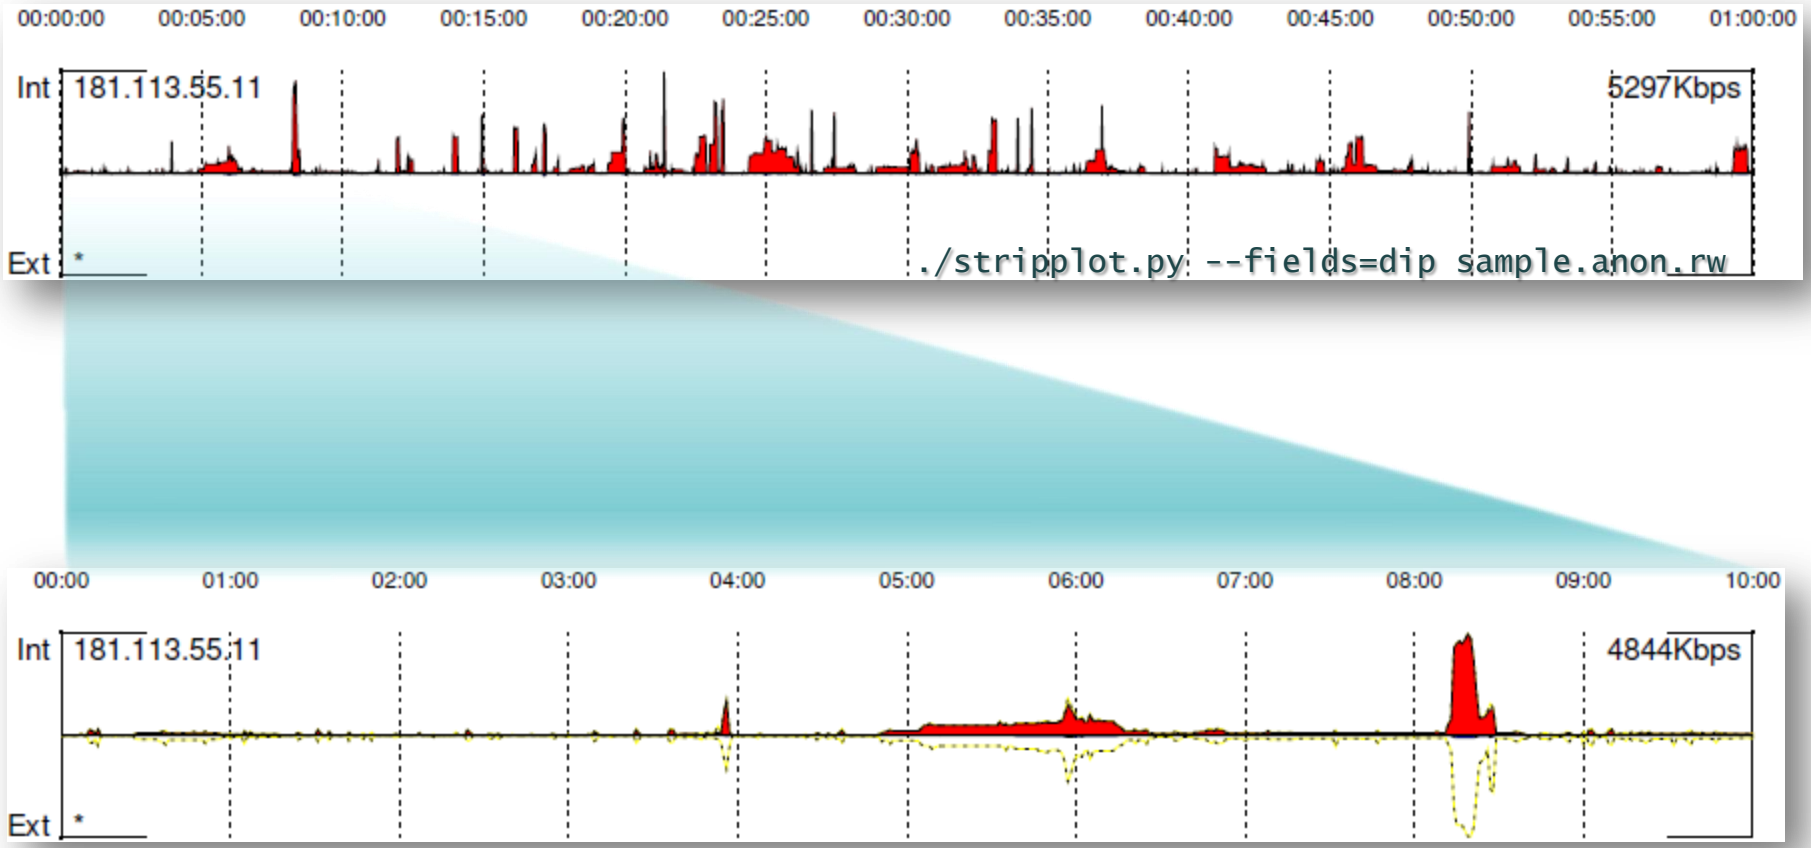


At 10 seconds (360 bins), the picture gets fuzzy:



Starttime, endtime

Zoom in on a particular time frame:




```
./stripplot.py --fields=dip --starttime=2009/11/17T00:00:00 --endtime=2009/11/17T00:10:00 \
                --trendline=p sample.anon.rw
```


© 2010 Carnegie Mellon University

Special Features

© 2010 Carnegie Mellon University

Automatically selecting fields

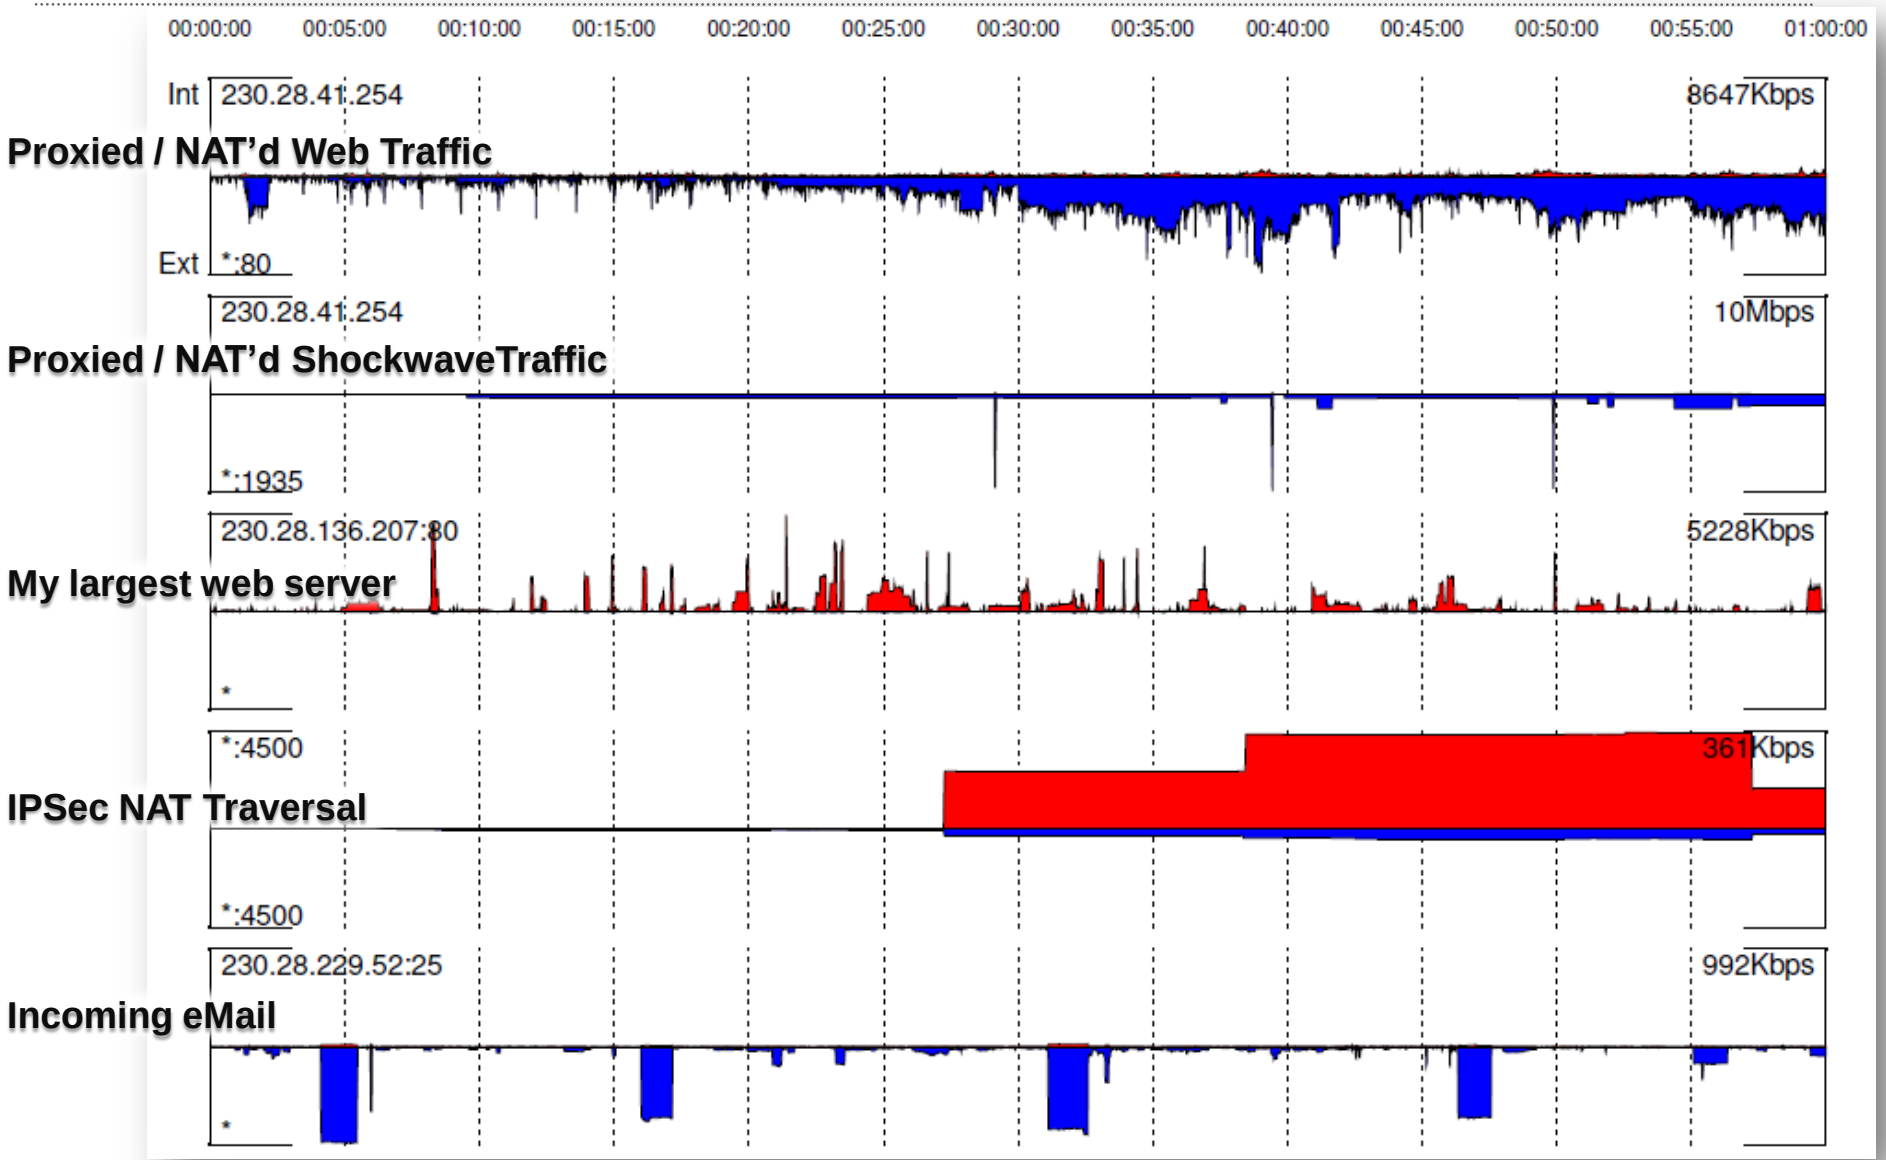
--fields=*

- Selects the best combination for 2, 3, 4 or 5 of sip,dip,sport,dport,proto
- For example, looking at all traffic for one network:

```
*:80 -> 230.28.41.254:*
*:1935 -> 230.28.41.254:*
*:* -> 230.28.136.207:80
*:4500 -> *:4500
*:* -> 230.28.229.52:25
45.178.111.132:61296 -> 230.28.41.74:56066
239.213.117.254:8080 -> 230.28.41.254:53474
*:* -> 230.28.229.75:25
45.178.111.132:50769 -> 230.28.41.74:55722
*:443 -> 230.28.41.254:*
```


© 2010 Carnegie Mellon University

Auto-select Example



© 2010 Carnegie Mellon University

Auto-select: how does that work?

Simple, but slow. Find the top value for rwuniq on:

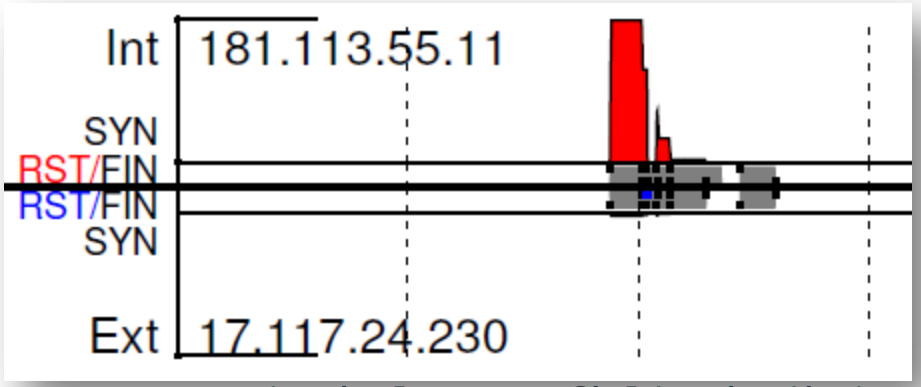
- sip,sport,dip,dport
- sip,sport,dip
- sip,sport,dport
- sip,sport
- sip,dip,dport
- sip,dip
- sip,dport
- sport,dip,dport
- sport,dip
- sport,dport
- dip,dport

© 2010 Carnegie Mellon University

Plotting TCP Flags

--flags

- Display points for TCP flags
- Works fine, just can't find a good generic use case




```
./stripplot.py --fields=sip,dip \
               --flags sample.anon.rw
```

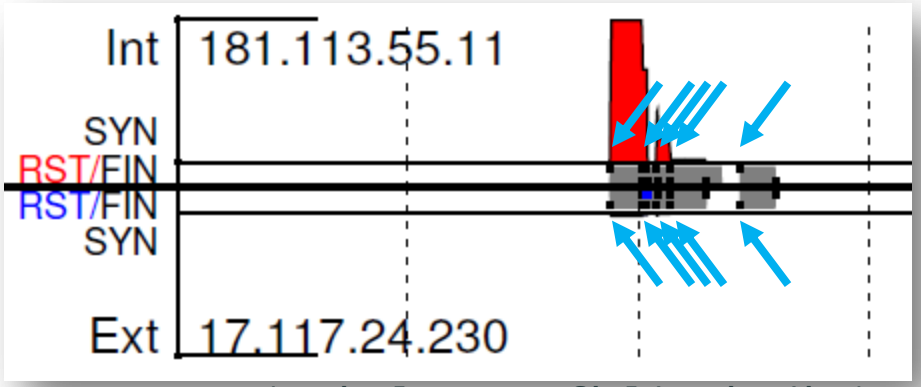
Activity: Grey bar

© 2010 Carnegie Mellon University

Plotting TCP Flags

--flags

- Display points for TCP flags
- Works fine, just can't find a good generic use case




```
./stripplot.py --fields=sip,dip \
    --flags sample.anon.rw
```

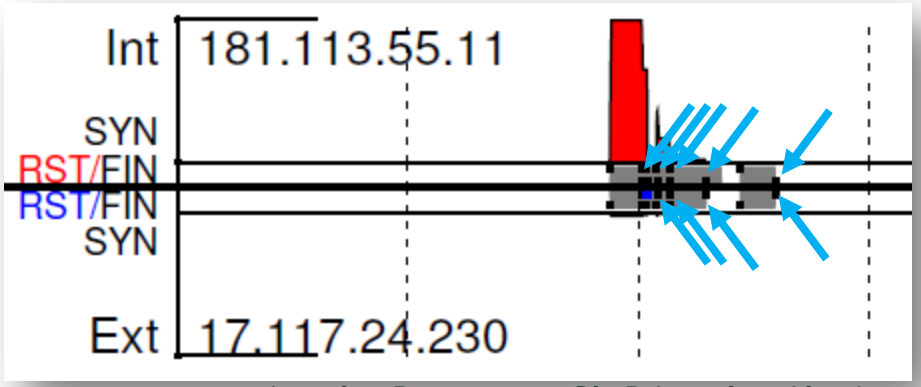
SYN Packets (outside, black)

© 2010 Carnegie Mellon University

Plotting TCP Flags

--flags

- Display points for TCP flags
- Works fine, just can't find a good generic use case




```
./stripplot.py --fields=sip,dip \
               --flags sample.anon.rw
```

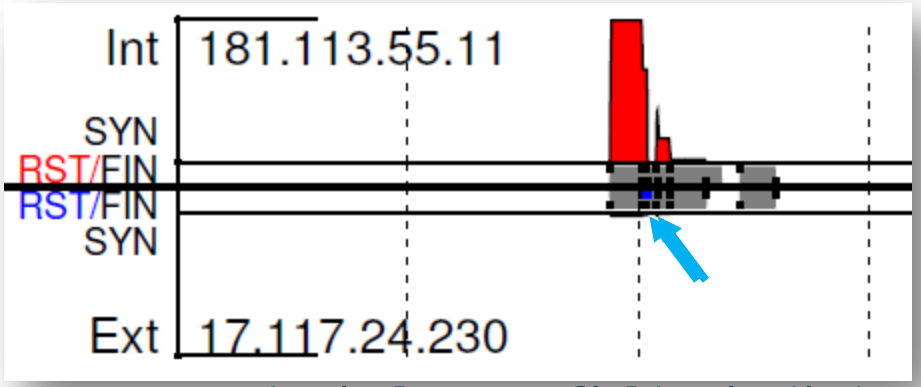
FIN Packets (middle, black)

© 2010 Carnegie Mellon University

Plotting TCP Flags

--flags

- Display points for TCP flags
- Works fine, just can't find a good generic use case




```
./stripplot.py --fields=sip,dip \
    --flags sample.anon.rw
```

RST Packets (middle, blue or red)

© 2010 Carnegie Mellon University

Verbose output, -v

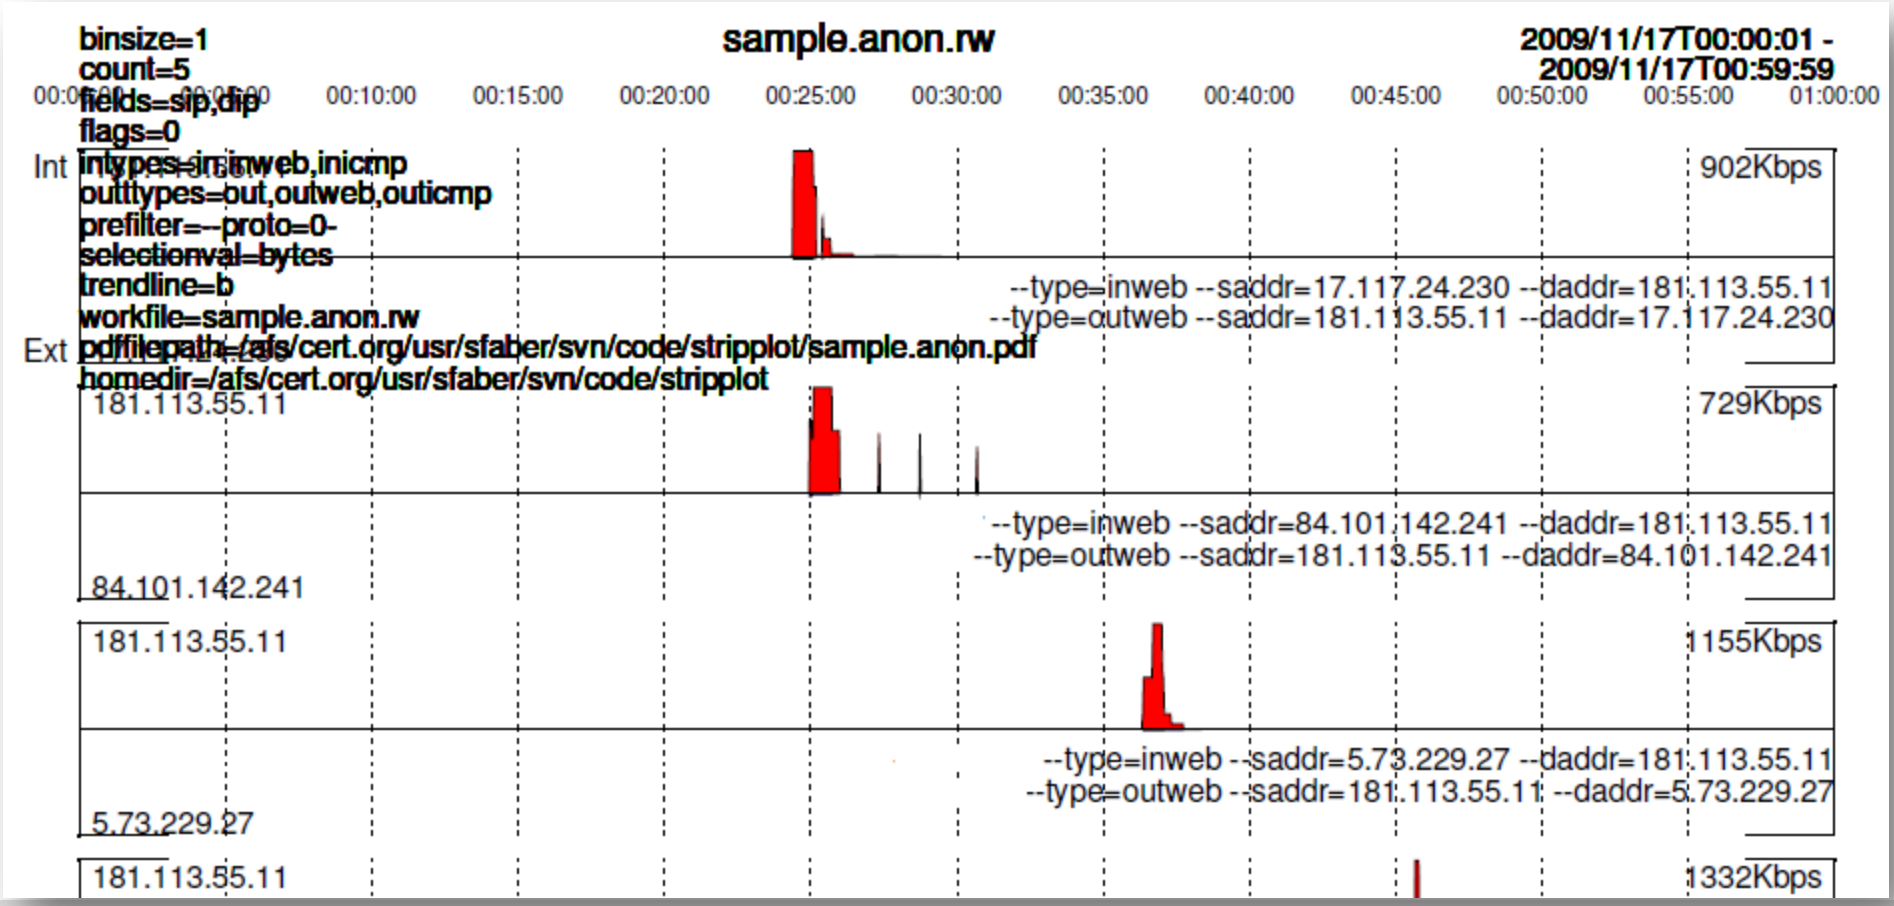
```
$ ./stripplot.py --fields=sip,dip --flags -v sample.anon.rw
Found 17.117.24.230:* -> 181.113.55.11:*
Found 84.101.142.241:* -> 181.113.55.11:*
Found 5.73.229.27:* -> 181.113.55.11:*
Found 159.133.127.154:* -> 181.113.55.11:*
Found 195.208.192.99:* -> 181.113.55.11:*
#  Settings:
#  binsize    1
#  bottomMargin 0.08
...
#  usableWidth 0.9
#  verbose    1
#  workfile   sample.anon.rw
(000): Page 001, Plot 000
(001): Page 001, Plot 001
(002): Page 001, Plot 002
(003): Page 001, Plot 003
(004): Page 001, Plot 004
```


© 2010 Carnegie Mellon University

Very Verbose Output, -vv

Echos all rw* commands

Adds lots of info to the .pdf output



© 2010 Carnegie Mellon University

Overriding default types

--types=in/out,inweb/outweb

- It's OK if the type doesn't actually exist (ie, multiple installations)
- Has to match rwcut type field

© 2010 Carnegie Mellon University

Open Issues

Better error checking

- Most inputs are passed directly to the rw* tools
- Occasional errors trying to plot empty data sets

Doesn't work well for transit traffic

- Assumes "in" and "out" traffic

Problems if you have outbound traffic only

- Trend selection doesn't seem to work well

TCP flags

- Tech is there, but visualization needs lots of improvement

© 2010 Carnegie Mellon University



Strip Plots: A Simple Automated Time-Series Visualization

Sid Faber
sfaber@cert.org

© 2010 Carnegie Mellon University

Defense Information Systems Agency
Department of Defense

Lessons Learned while providing SiLK Training

Jim Downey, DISA

Disclaimer

The views are the authors alone and do not necessarily reflect the Defense Information Systems Agency. This is for information purposes only. No warranty of any kind is expressed or implied.



About me

- **Worked in a CSIRT as: Incident Analysis and Coordination, Penetration Testing, Antivirus support. Worked with SiLK on and off since 2002.**

- **Most of what I'm saying is first hand, but we've had other trainers too.**

Some variables and some constants

- Topics stayed largely the same over this discussion. Trained analysts on SiLK along with a few other tools (visualization, scan database, iSiLK) built using SiLK
- Started with real data and context, migrated towards obfuscated data
- Used different facilities and gear at different times
- Only recently started pre-testing for Unix proficiency
- No pre-testing for TCP/IP familiarity

Facilities

- Our customers are in roughly six different time zones
- Our own facility:
  - Having a CAC doesn't mean having a certificate on that CAC.
  - Live training via teleconferencing difficult to do and almost impossible to do effectively with in person students at the same time
- Customer provided facilities:
  - Insist on testing full connection including usernames and passwords (at least one!) before you jump on that plane. Then quietly check against your data server's logs.
  - Just because it works on one workstation doesn't mean all the workstations on a network have the same browser configuration.
  - Just because a commercial version of SSH client is installed doesn't mean the customer's firewall allows tcp/22 outbound to the training server.
  - If you've never seen pictures or been before, and the customer says there's an overhead projector, bring a pointer in case it's beyond arms reach.

The tools

- **SiLK installed easily on Linux environment. iSiLK installed easily under Windows without System Admin privileges (your mileage may vary).**

- **At least one Unix (Ubuntu) environment where iSiLK client doesn't work. Open source software doesn't necessarily mean it'll work on an Open Source Operating System.**

- **rwrandomize could potentially be replaced with a process vice a simple tool, as others have suggested**

Training Data

- If you can train on real data, with context, that's the most compelling to a student.
- If you obfuscate your data, save the map! We later added asset data and had to grab new netflow data to have them similarly obfuscated.
- Keep the general types of use cases, even if you change the names, locations, dates, etc. Otherwise, looses interest.
- For obfuscated data, we used a couple of /16s plus added some additional flow data covering a Denial of Service attack, and changed the IPs and times.

Students & Training Material

- A fair number of students had forgotten some important details about TCP like the 3 way handshake.
- A fair number had minimal experience with Unix, a few were software developers.
- Highlights found in SEI (or other existing or easily created material) are sufficient for the above cases.

I've been trained, now what?

- Users may not know what to do with flow data other than scope and impact based on IDS alerts.
  - Know the users overall mission and business processes
    - Integrate with this where it makes sense
    - Advocate rescoping where it makes sense
    - You may have to do a lot of the reworking
- Users may not have the skills to build out processes.
  - You may find it useful to write scripts, provide mappings, etc yourself.
  - Someone has to maintain these


REDJACK

REDJACK

Outline

- Core Concept and Goal
  - Expressiveness *vs* Representation
- Towards a Behavioral Dictionary
- Example Behavior: Fumbling
  - What is Fumbling?
  - Why Fumbling?
  - Who Fumbles?
- Current Study: ID Crawlers Via Fumbling

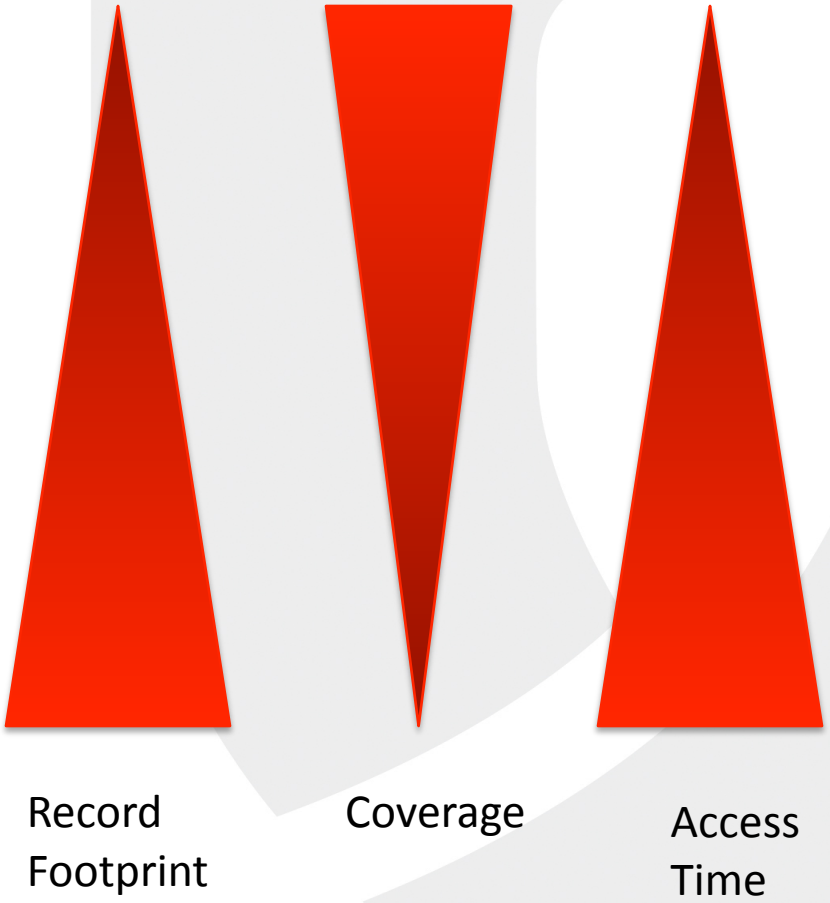
REDJACK

CORE CONCEPTS

REDJACK

Expressiveness Vs. Representation

- "If I just had **TCPDUMP OF EVERYTHING** all my problems would be solved"
- Most traffic, on a flow-by-flow basis is either garbage or uninteresting



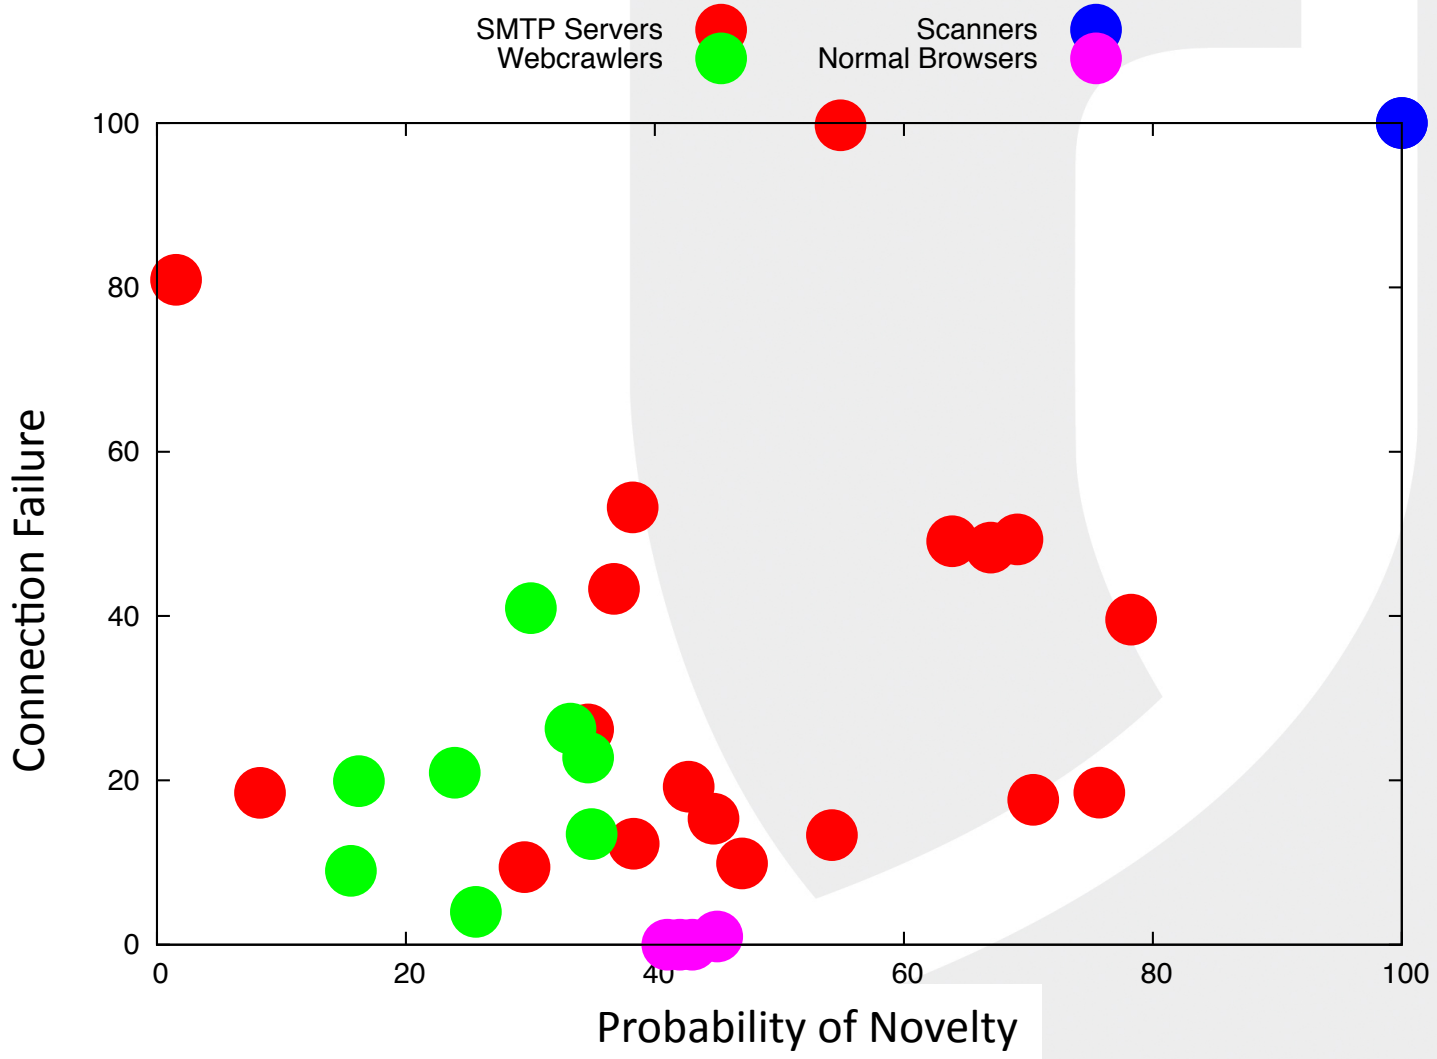
REDJACK

Goal

- Develop narratives which describe activity between hosts in a more abstract fashion
  - "This is fumbly"
  - "This is chatty"
- Ideally, these attributes will be
  - Intuitive (an analyst can grasp them by looking at a log)
  - Rigorous (derived from some model of behavior)
- Partly identification applications by behavior

REDJACK

A Crude Picture



REDJACK

Attributes For Narratives…

- This is clustering, just on different axes
- Possible attributes:
  - Probability of connection failure
  - *Locality*
  - Probability of file transfer
  - Packet size

REDJACK

CASE STUDY: FUMBLING

REDJACK

What is Fumbling?

- Intuitively, fumbling is a *consistent* failure to connect with a host
  - Previously used to identify BitTorrent [Collins06,Bartlett07]
- Challenge: differentiating fumbling
  - From scanning, where clients probe 'a lot' [Jung04]
  - From normal surfing, where clients get bored and move on

REDJACK

What Fumbles?

- Routed/automated lookup
  - SMTP
  - P2P
  - NNTP
- Search bots
- Scanners don't fumble – they seek out everything
- Users don't fumble – they lose patience

REDJACK

Why Care About Fumbling?

- Scanning false positives
  - Uncleanliness Data – don't mark yahoo unclean
  - Differentiate scanners in a naturally noisy set -- SMTP
- Identify applications that require blind lookup
  - Internal p2p applications (unknown ports)
  - Google doesn't publish crawler IP addresses

REDJACK

CHARACTERIZING FUMBLING

REDJACK

Source Data

- Task: differentiate crawlers by quantifying fumbling
- 4 days of crawling data
- Crawlers identified by IP space
  - Cuil: Google “competitor”, embarrassing launch last year
  - Yeti: Naver.com (Korean search engine)
  - “Twiceler”: Some searchbots use twiceler as an ID, refers to twiceler from domains cuil says are not its domains
  - Voila: Voila.fr search engine (French)

REDJACK

Basic Numbers

Engine	ID#	Flows	IPs	Failed Flow	Failed IP
Cuil	1	3760	189	504	45
	2	4945	170	195	42
	3	3128	204	1033	43
Yeti	4	2635	247	84	28
“Twiceler”	5	5338	185	829	51
Voila	6	12808	680	2745	75
	7	12506	679	2306	73
“Twiceler”	8	2252	172	101	45

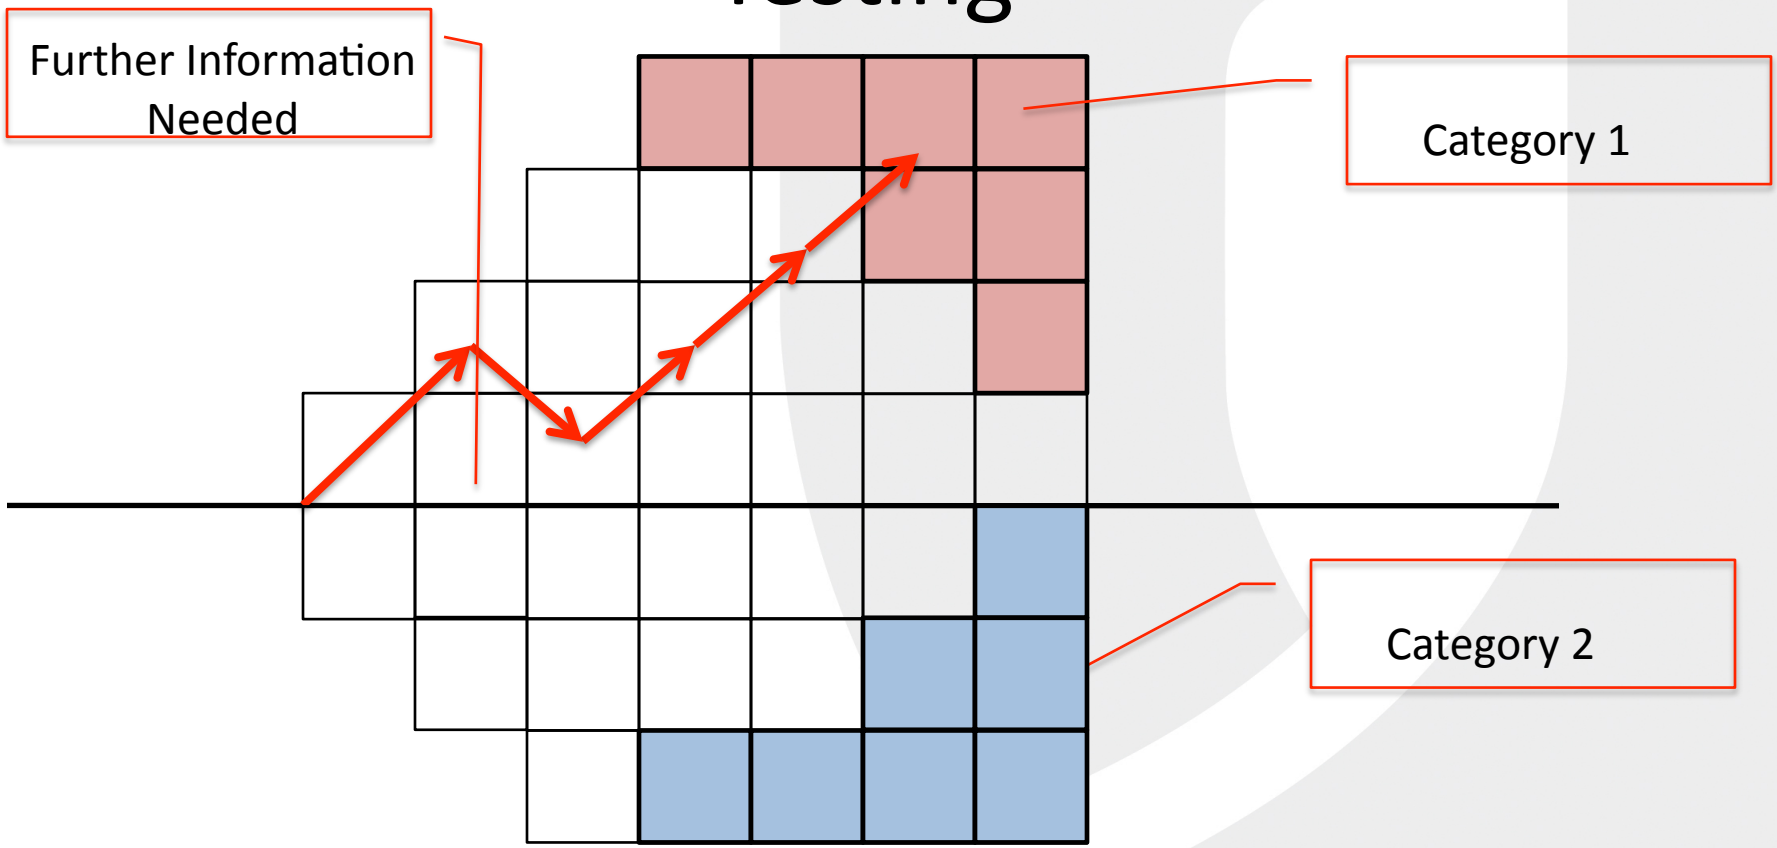
REDJACK

Consecutive Failure Rate

- Number of times that a failure occurs repeatedly
- Used in darkspace analysis – scans are marked as such when > 3-5 consecutive failures [Jung, 2004]
- Fumblers are different because they have a nontrivial success rate

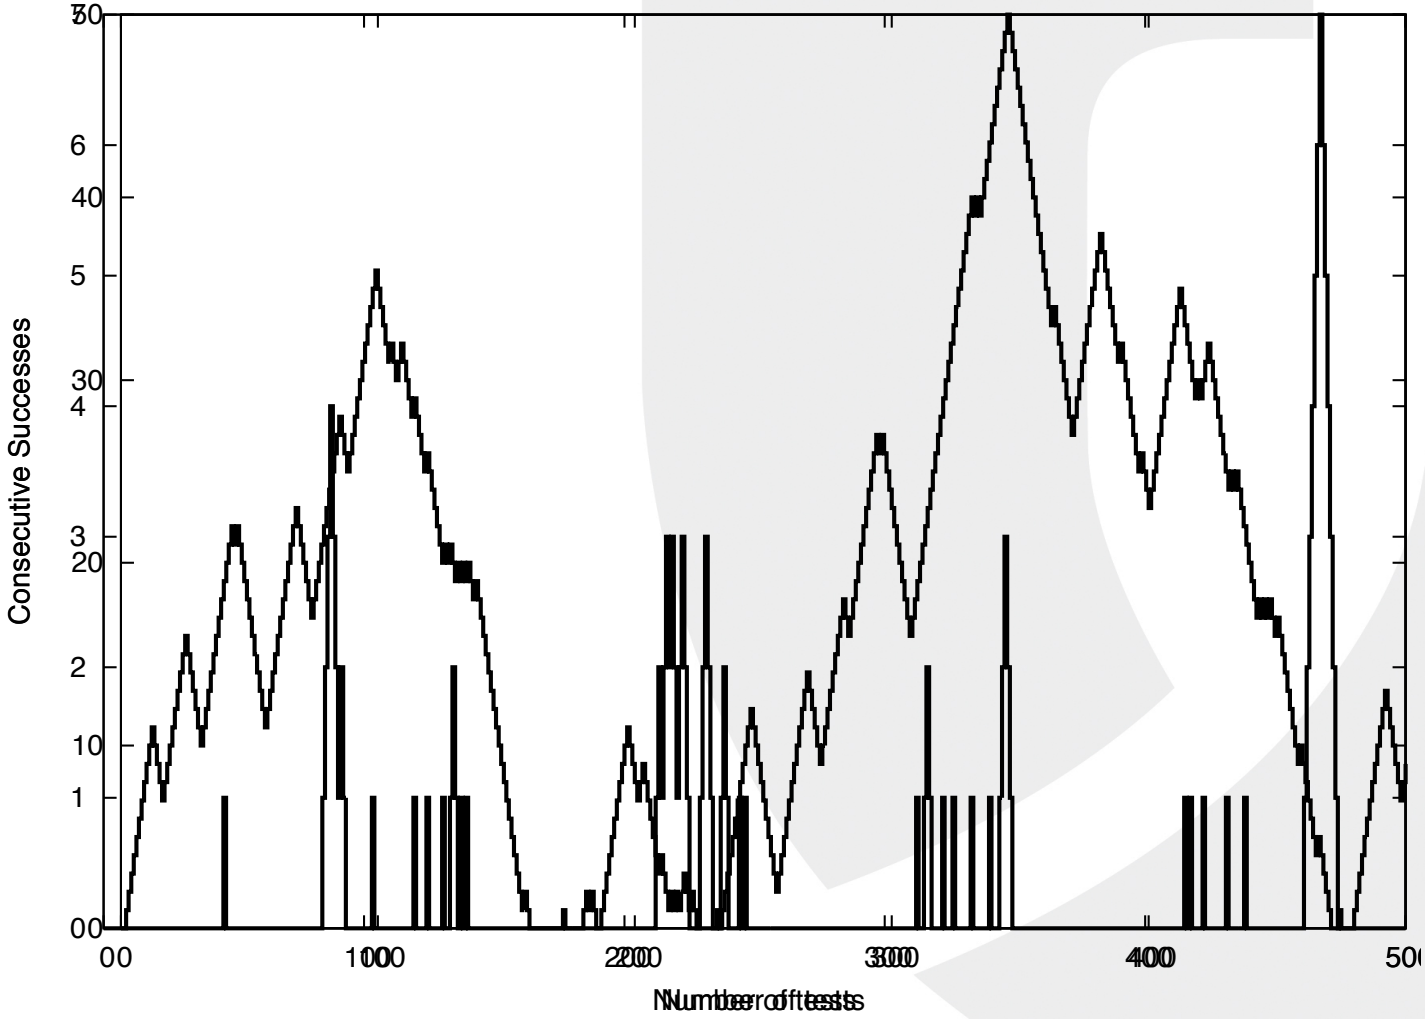
REDJACK

Visualizing Sequential Hypothesis Testing



REDJACK

Some Failure Plots



REDJACK

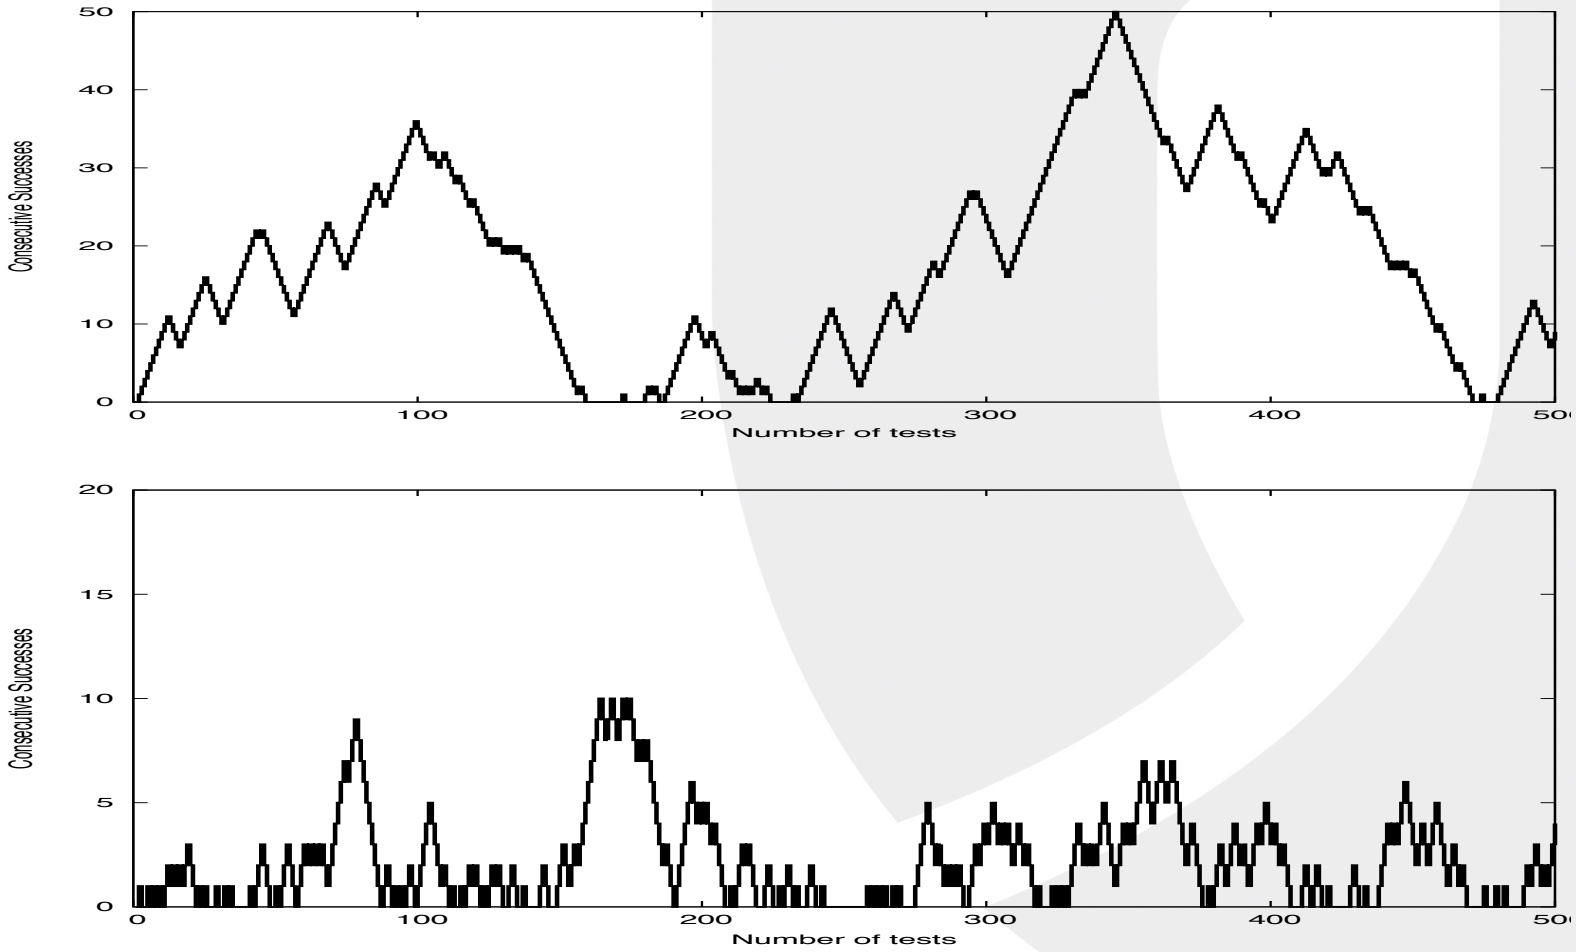
Why The Mountains?

- In the long run, high success rate:
  - 0.5% failure to 70% in the worst case
  - Vs. 99.95% failure rate for scanners
- However, failures are common mode
  - IP address X is down
  - IP address X is hit repeatedly

Engine	ID#	FPR (4 failures)
Cuil	1	9.10%
	2	1.50%
	3	34.4%
Yeti	4	10.3%
“Twiceler”	5	17.9%
Voila	6	13.9%
	7	1.00%
“Twiceler”	8	1.00%

REDJACK

Permuting Addresses



REDJACK

Results of Permutation

- Drops expected values down
- “realtime” detection is no longer valid
- Fumbling requires both a presence and absence…

Engine	ID #	Norm FPR	Random FPR
Cuil	1	9.10%	0.00%
	2	1.50%	0.00%
	3	34.4%	15.5%
Yeti	4	10.3%	0.00%
“Twiceler”	5	17.9%	0.00%
Voila	6	13.9%	0.00%
	7	1.00%	0.00%
“Twiceler”	8	1.00%	0.00%

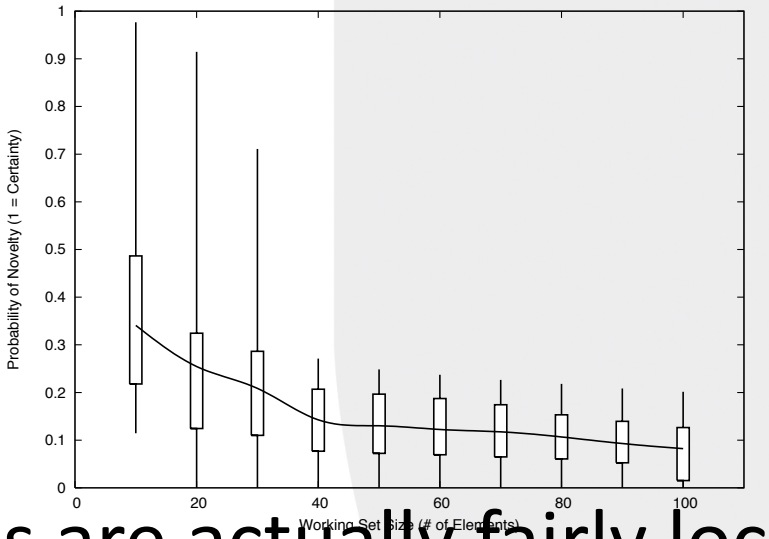
REDJACK

Locality

- Propensity of users to sit around a set of common hosts [McHugh03]
- Modeled using a working set:
  - LRU stack, fixed size
  - Locality is then the probability, when an address is presented, of not replacing an address in the working set

REDJACK

Searchbots Are Local



- Searchbots are actually fairly local
  - Slightly *more* localized than humans, it turns out
  - CDNs? Single-page sites?
- *Much* more local than scanners

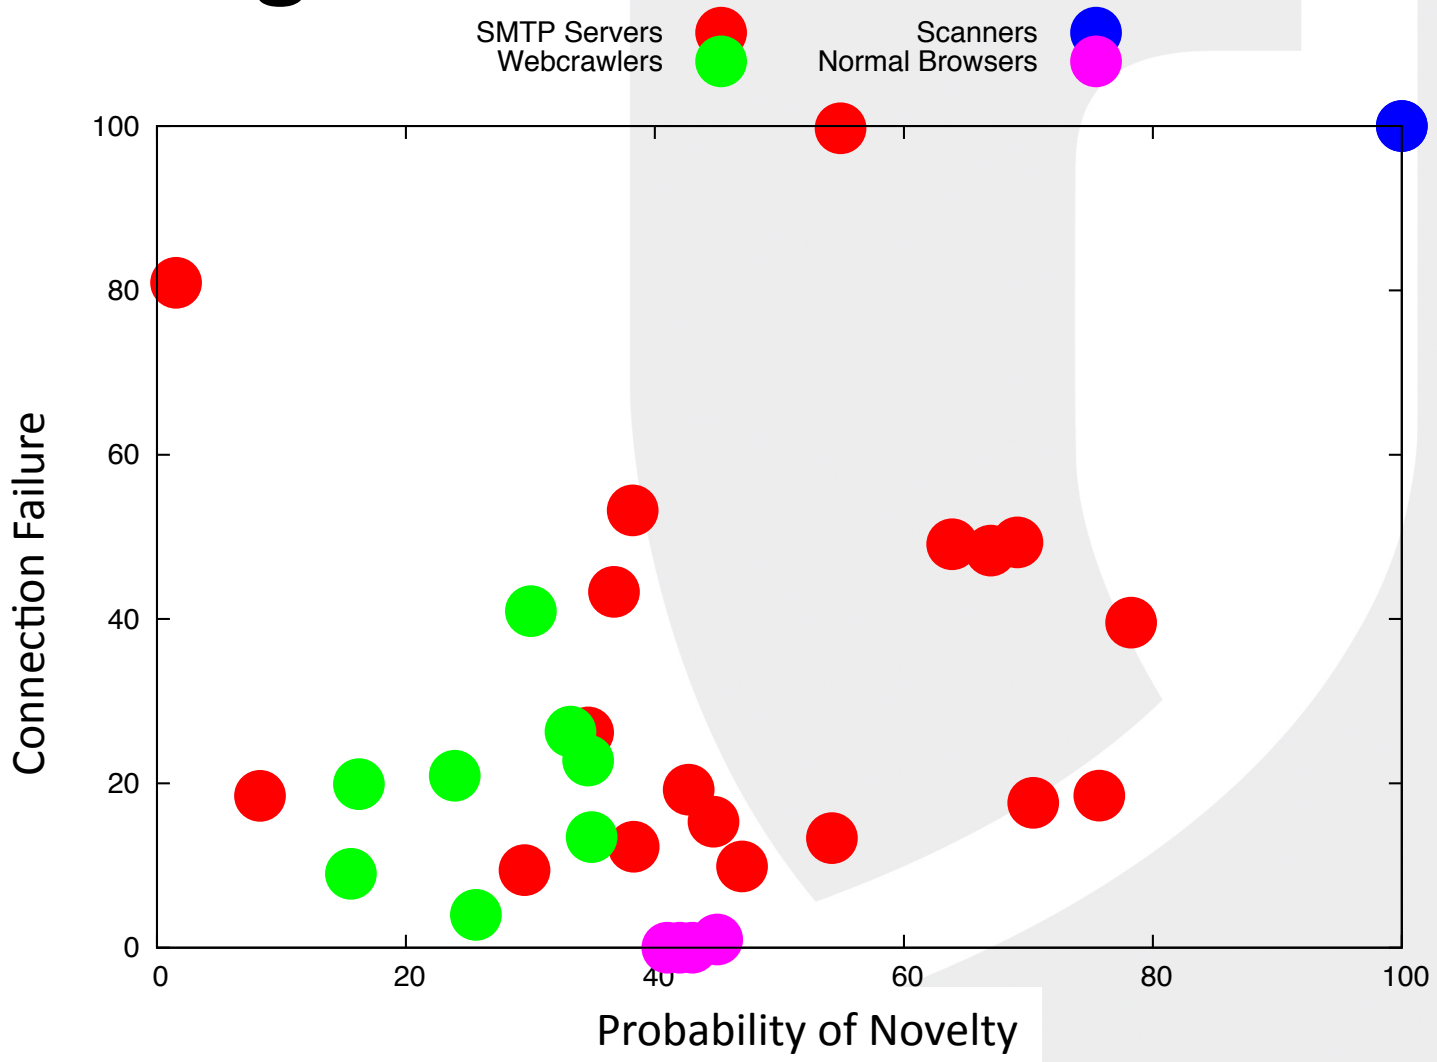
REDJACK

Start Classifying

Local?	Connects?	
	Yes	No
Yes	Surfer	Searchbot
No	Hitlist Scanner?	Scanner

REDJACK

Leading Us Back To This Picture…



REDJACK

Conclusions

- Combining locality with detection failure may provide an indicator of fumbling
  - Have to develop a suitable *n* (working set size)
  - *N* also changes over time
- A false positive is an indicator your IDS isn't done yet
  - We can differentiate searchbots from scanners with more information, but it may cost us 'realtime'
  - Whatever 'realtime scan detection' is worth…

Introduction to Argus

http://qosient.com/argus

FloCon 2010
New Orleans, La
Jan 11, 2010

Carter Bullard
QoSient, LLC

carter@qosient.com

Carter Bullard carter@qosient.com

- QoSient - Research and Development Company
  - Naval Research Laboratory (NRL), GIG-EF, JCTD-LD, DISA, DoD Network Performance and Security Research
- Inventor/Developer Argus http://qosient.com/argus
- FBI/CALEA Data Wire-Tapping Working Group
- QoS/Security Network Management - Nortel/Bay
- Security Product Manager – FORE Systems
- CMU/SEI CERT
  - Network Intrusion Research and Analysis
  - NAP Site Security Policy Development
  - Network Security Incident Coordinator
- NFSnet Core Administrator (SURAnet)
- Standards Efforts
  - Editor of ATM Forum Security Signaling Standards
  - IETF Working Group(s) Contributor
  - Internet2 Security WG
  - NANOG





Argus

- Argus is a network utilization audit system

  Argus was officially started at the CERT-CC as a tool in incident analysis and intrusion research. It was recognized very early that Internet technology had very poor usage accountability, and Argus was a prototype project to demonstrate Red Book strategies for LAN and CAN network auditing.

- Composed of
  - Real-time Network flow monitor
  - Network flow data collection system
  - Network flow data processing programs
  - Audit data repository tools





Argus History

- Georgia Tech (1986)

  Argus was the first network flow data system. Started at Georgia Tech, Argus was used as a real-time network operations and security management tool. Argus monitored the Morris Worm, and was instrumental in discovery for the "Legion of Doom" hacking investigations.

- CERT/SEI/Carnegie Mellon University (1991)

  Argus was officially supported by the CERT as a tool in incident analysis and intrusion research. Used to catalog and annotate any packet file that was provided to the CERT in support of Incident Analysis and Coordination, it was a focal point for research in intrusion analysis and Internet security.

- Argus Open Source (1995)

  Transitioned into public domain in 1995. Supported by CMU and CERT/SEI at many levels including argus developers mailing list.

  Used now by a large number of educational, commercial and governmental sites for network operations, security and performance management.





Argus Licensing

- GNU GPL 3
- US DoD open source license
  - Gargoyle
    - Available from https://software.forge.mil
    - Growing development community
- Alternate licensing available





Disclaimers

- Argus is a proof-of-concept project.
  - Argus and its clients are examples of what can be done
  - Any concept, any time, as long as it fits
- Argus is NOT Netflow™
  - There are a lot of network flow data methodologies
  - Community needs to realize Netflow is Cisco's approach
  - Let's stop using Netflow as term for Network Flow Data
- Argus is NOT IPFIX
  - Lots and lots and lots of issues with IPFIX
  - Argus does attempt to avoid or resolve basic IPFIX problems





Introduction to Argus

- Argus Design (20m)
- Data Generation (25m)
- Client Programs (45m)
- Data Collection and Archiving (45m)
- Situational Awareness (45m)





Argus Design





Argus Design

- Comprehensive Network Accountability
- Based on Tried and True Methodology
- High Utility/Applicability
- High Performance
- Deployable / Scalable





Comprehensive Network Accountability

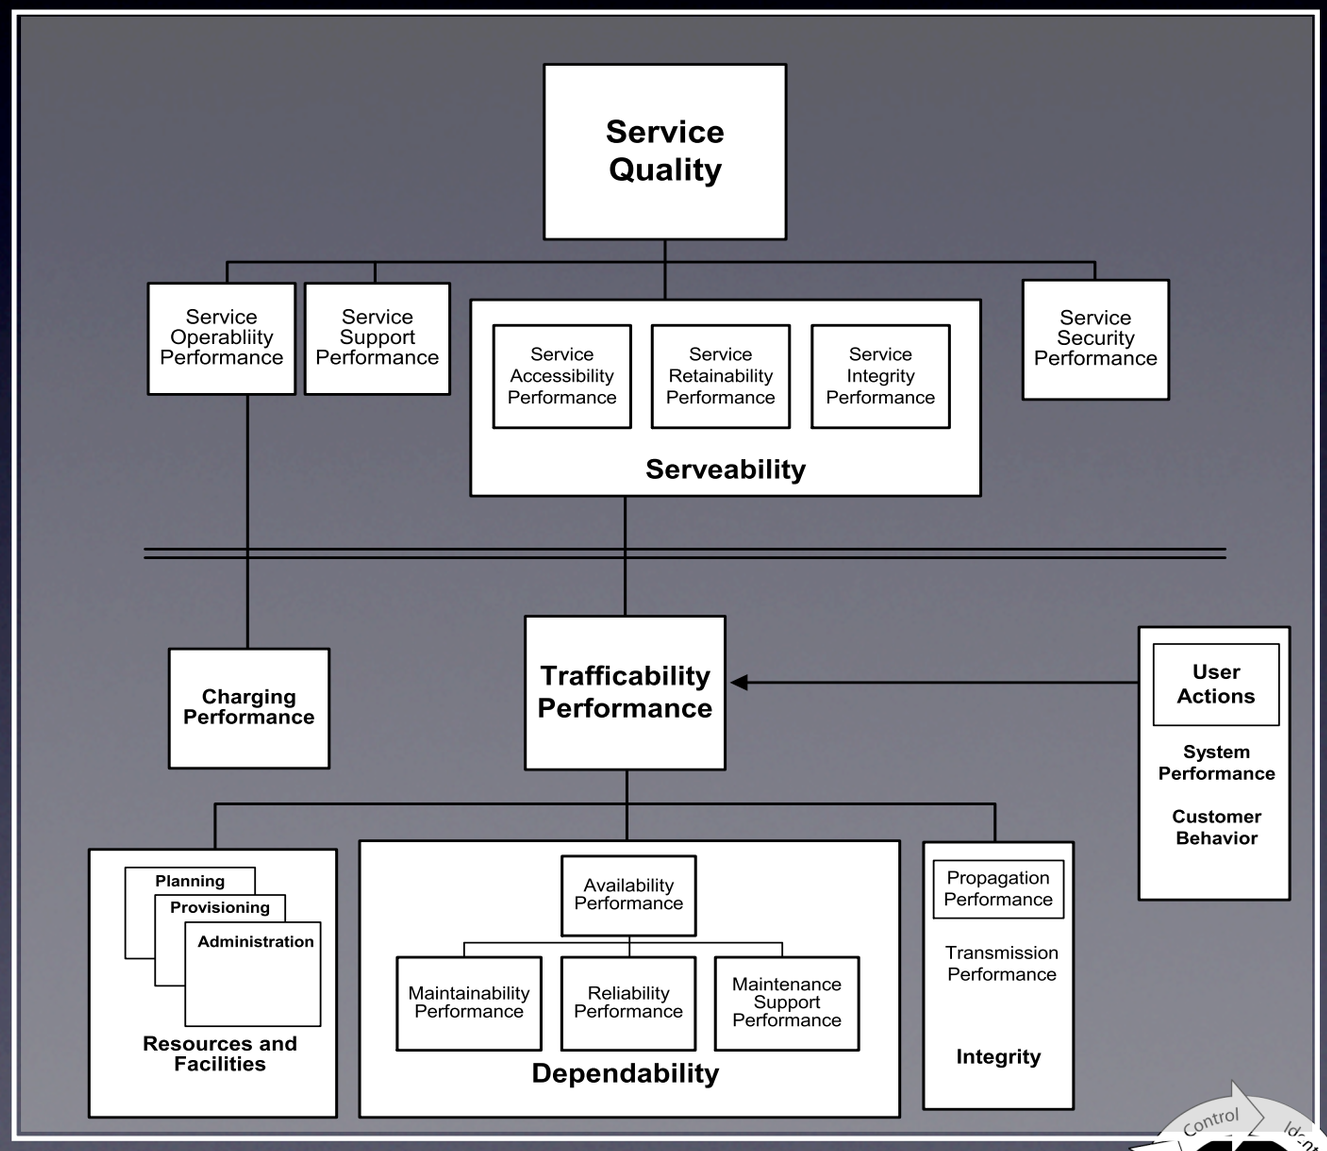
- Ability to account for all/any network use
  - Red Book prescribed method for trusted networking
- At a level of abstraction that is useful
  - Network Service Functional Assurance
    - Was the network service available?
    - Was the service request appropriate?
    - Did the traffic come and go appropriately?
    - Did the traffic get the treatment it was support to?
    - Did the service start and end normally?
  - Network Control Assurance
    - Is the control plane operational?
    - Was the service request appropriate?
    - Did the traffic come and go appropriately?





Tried and True Methodology

- ITU Network Service Quality and Usage strategies
- Support mature network service measurement architectures
- Passive, comprehensive, service oriented, integrated, shareable, extensible, accessible



From ITU-T Recommendation E.800 Quality of Service, Network Management and Traffic Engineering

Network Measurement Architectures

- Examples
  - ITU TMN, X.700, OSI, SPAND, EMA, NIMI, Optivity, NetView, UNMA, PRIMA, MFN, SNMP to name just a very few.
  - Strategies involve either pure passive, or a bundle of active and passive methods.
- Component Design
  - All measurements, whether passive, active, extractive, or injective, involve a passive measurement component
  - Argus is designed to be that passive component
- Systems Design
  - Data Generation, Collection, Disposition, Access
  - May need to provide it all.





Utility/Applicability

Feedback Directed Network Optimization



Function	Description	
Identify	Discover and Identify comprehensive network behavior	Collect and process network behavioral data
Analyze	Collect and transform data into optimization metrics, establish baselines occurrence probabilities and prioritize events.	
Plan	Establish optimization criteria, both present and future and implement actions, if needed	Provide information and feedback internal and external to the project on the optimization outcomes as events.
Track	Monitor network behavioral indicators to realize an effect.	
Control	Correct for deviations from criteria.	

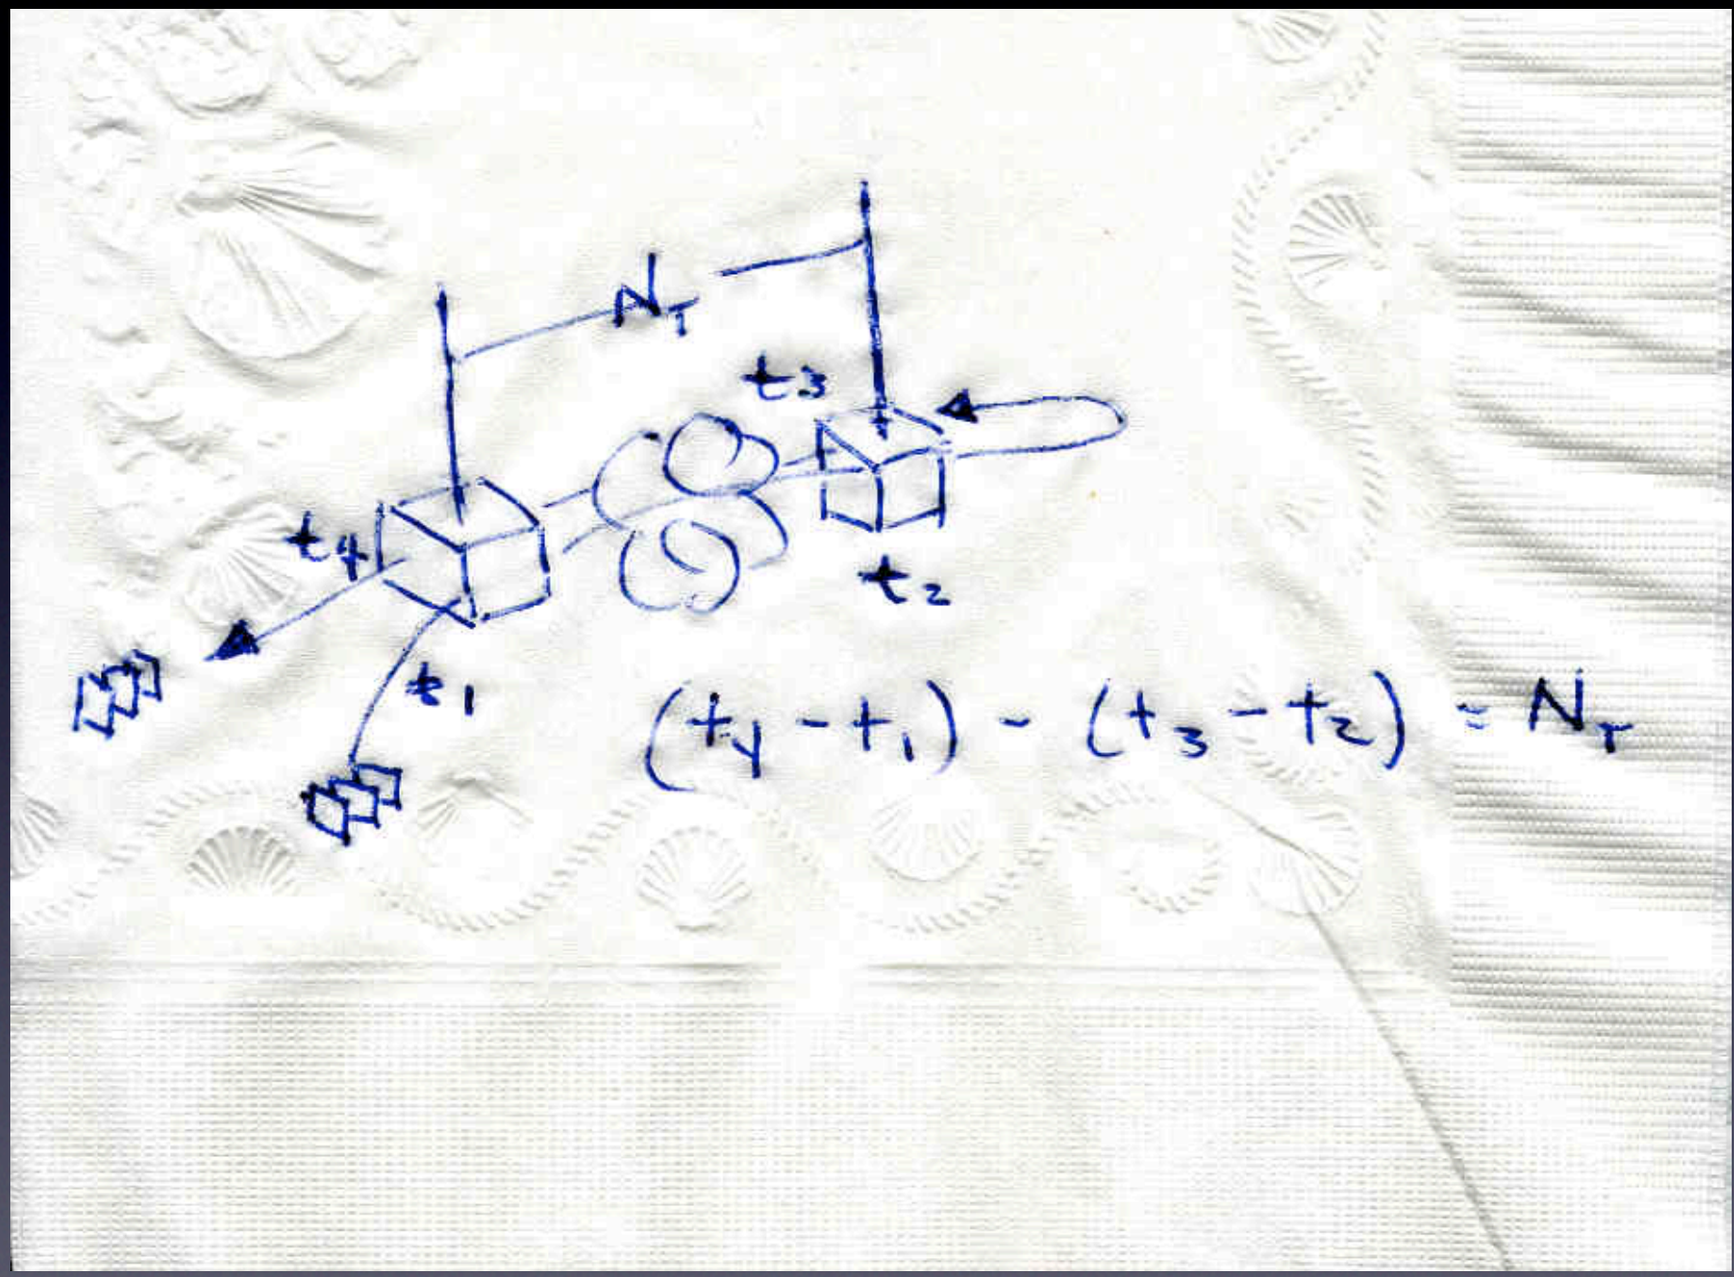


Utility/Applicability

- Extensive Identifiability Methods
  - Multi-Layer Accountability
    - End Point Object Identifiers
      - Arbitrary encapsulation parsing and reporting
      - Need to support many, many formats
    - Service Oriented Object and State Identifiers
- Standards Based Traffic Metrics
  - Availability, Reachability, Connectivity (RFC 2678)
    - Bi-Directional
  - Service Oriented Usage/Performance Measurements
  - Advanced Performance Measurements
    - Loss, Jitter, Delay, Power, Distance, Performance
- Relational Algebraic Constraints






N_T
t_3
t_4
t_2
t_1
$(t_4 - t_1) - (t_3 - t_2) = N_T$

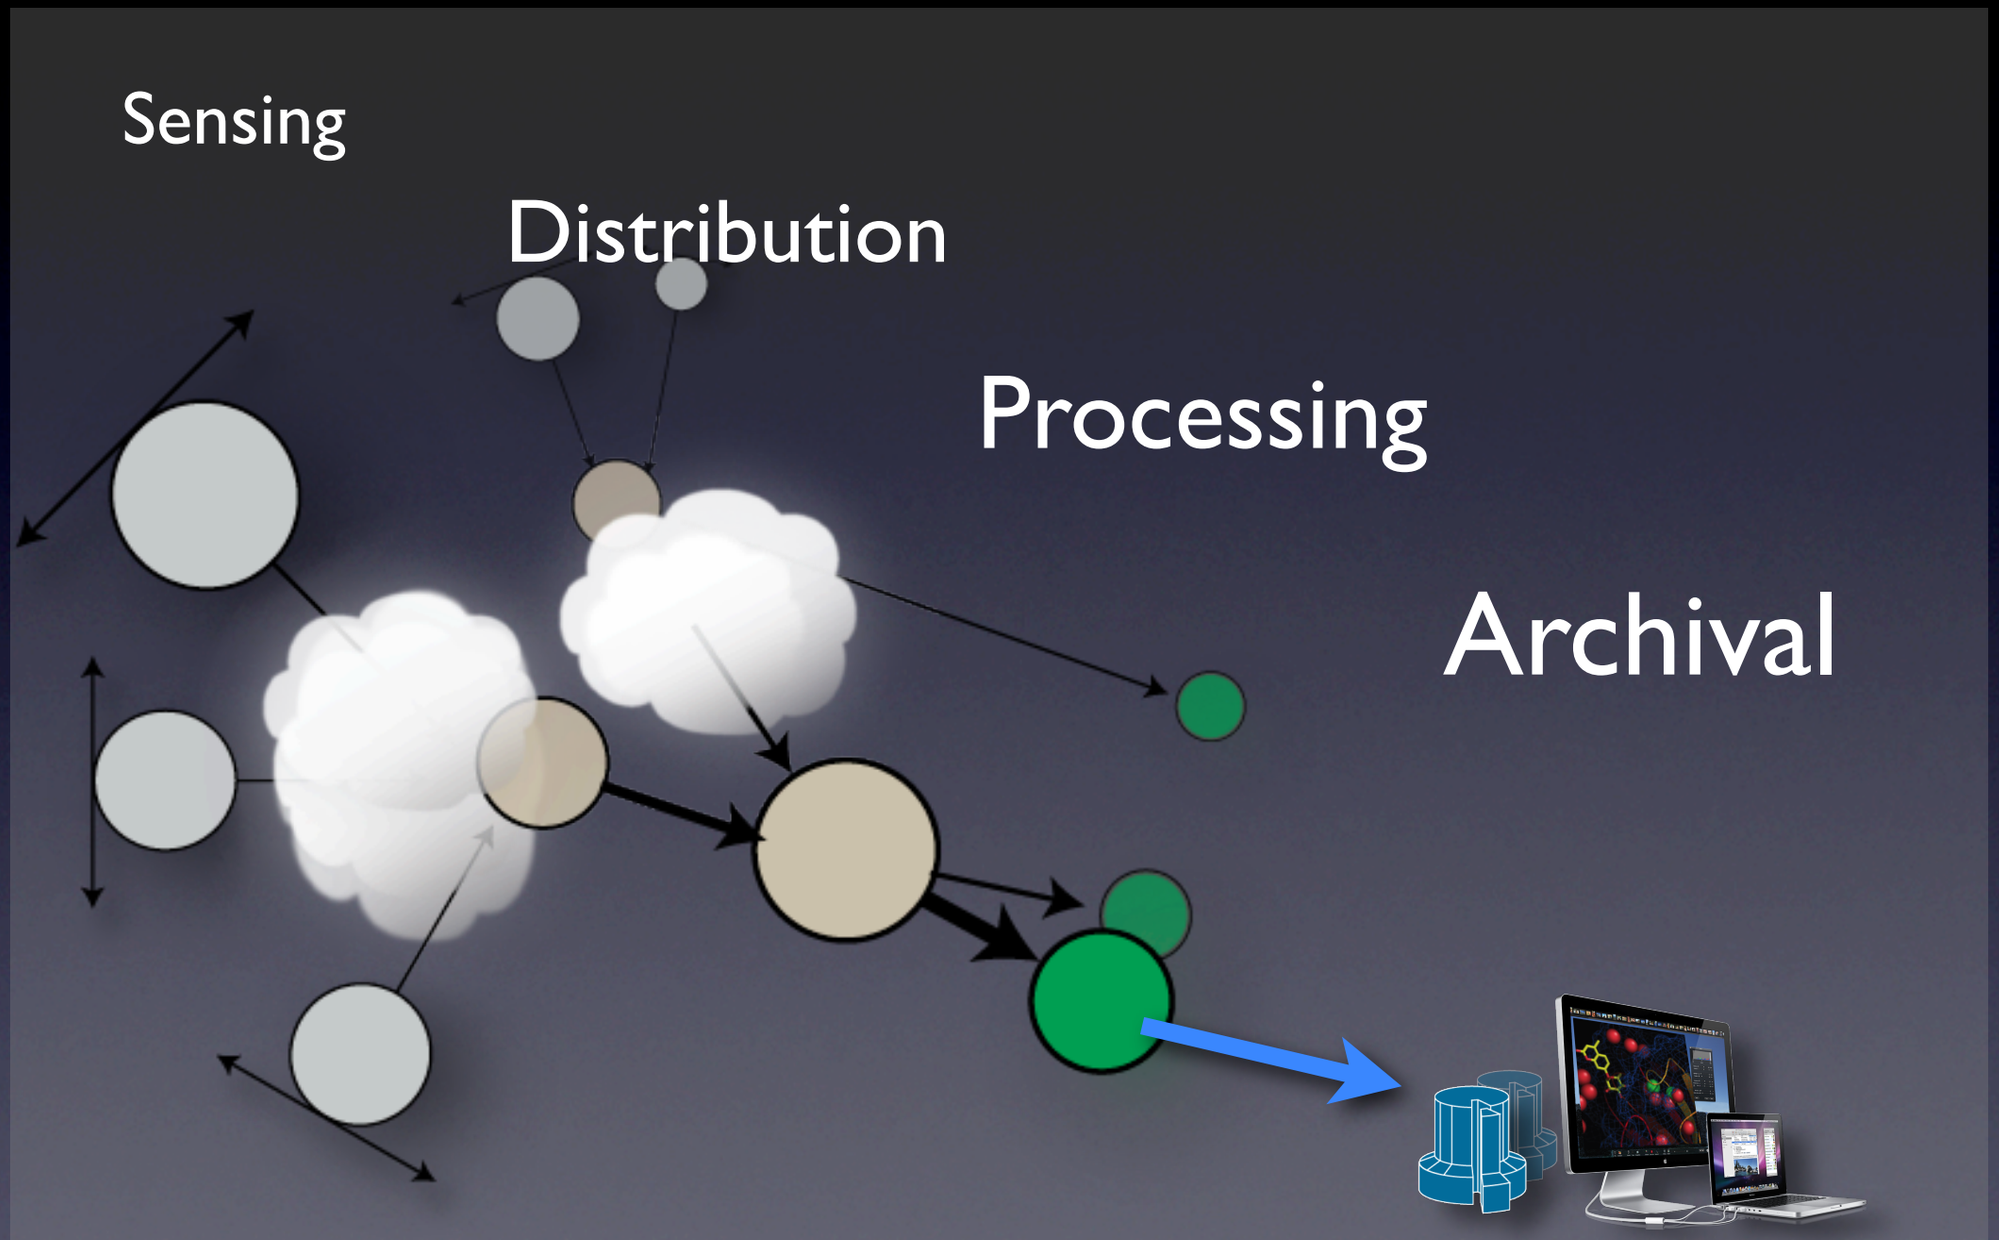
So what does all that mean?

- We need an Internet transaction concept
  - Network flow data record (Net CDR)
  - Network service oriented
    - Initiation, status, termination state indications for 5-tuple flows
    - All services - ARP, DHCP, DNS, OSPF, TCP....
- Data generation must be timely/deterministic/non-statistical/relevant/comprehensive
- Approach needs to perform and scale
  - Formal generation/consumption architectures
  - Collect/transport/process/correlate/join/select/search/store data
- Need to convey as much about the network traffic as possible
  - Support Layer 2/3/4/5+ semantics (including non-IP traffic)
- Needs to really solve some problems

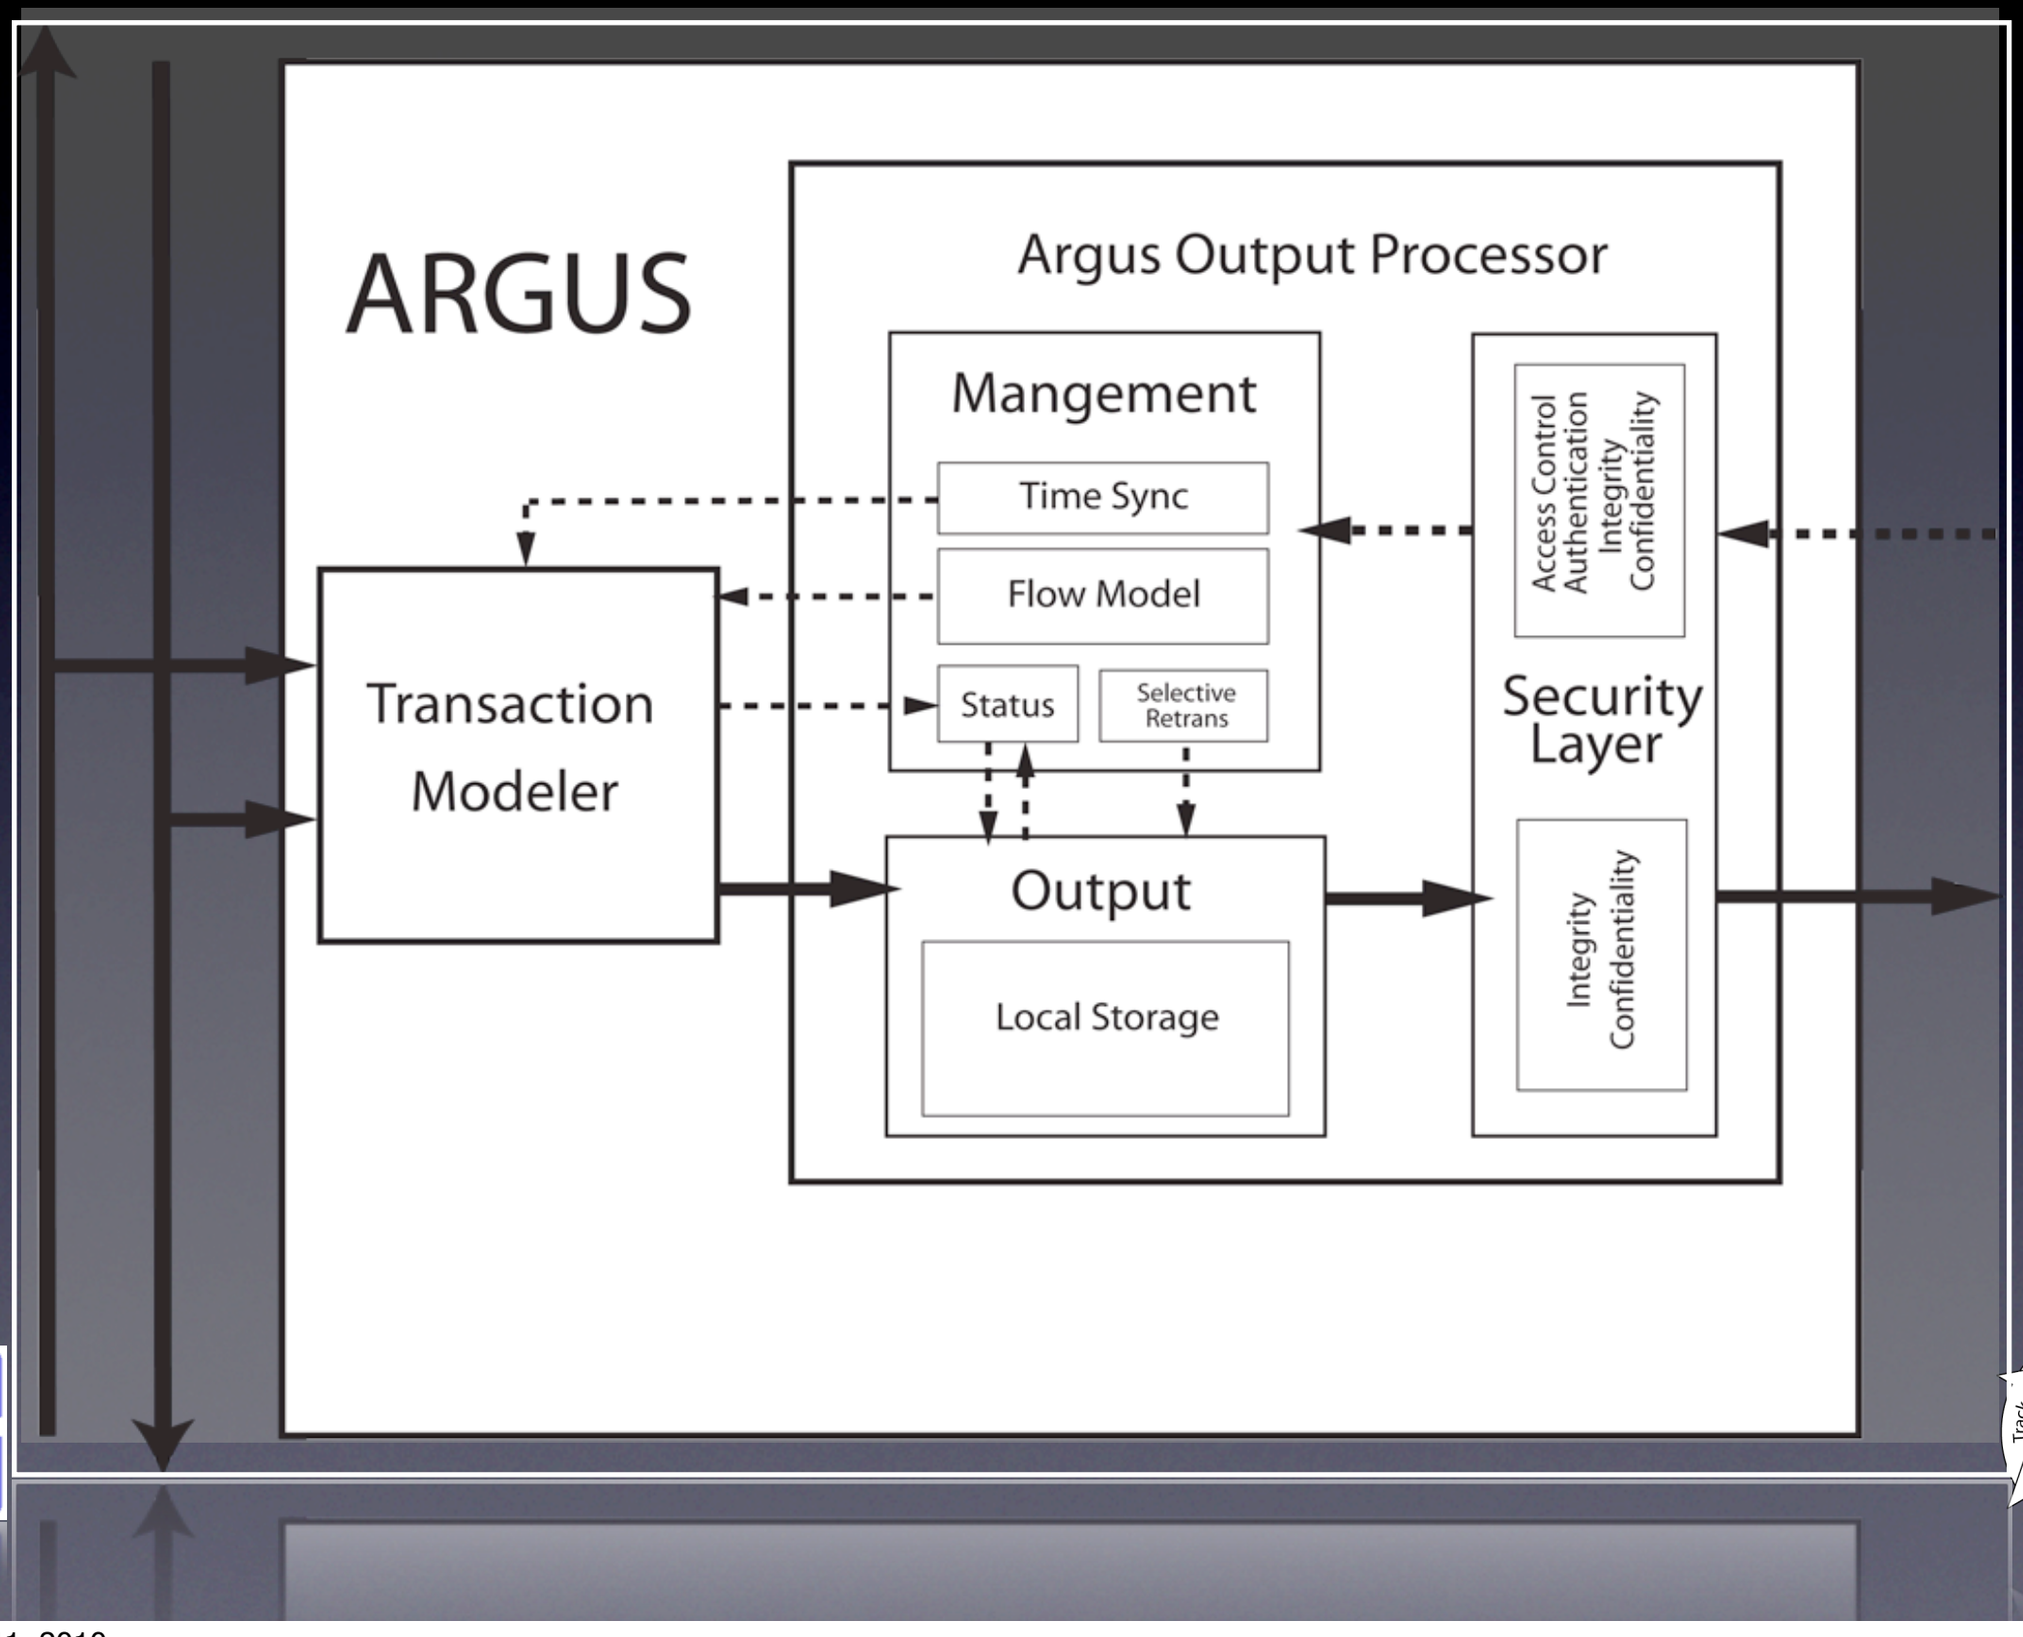




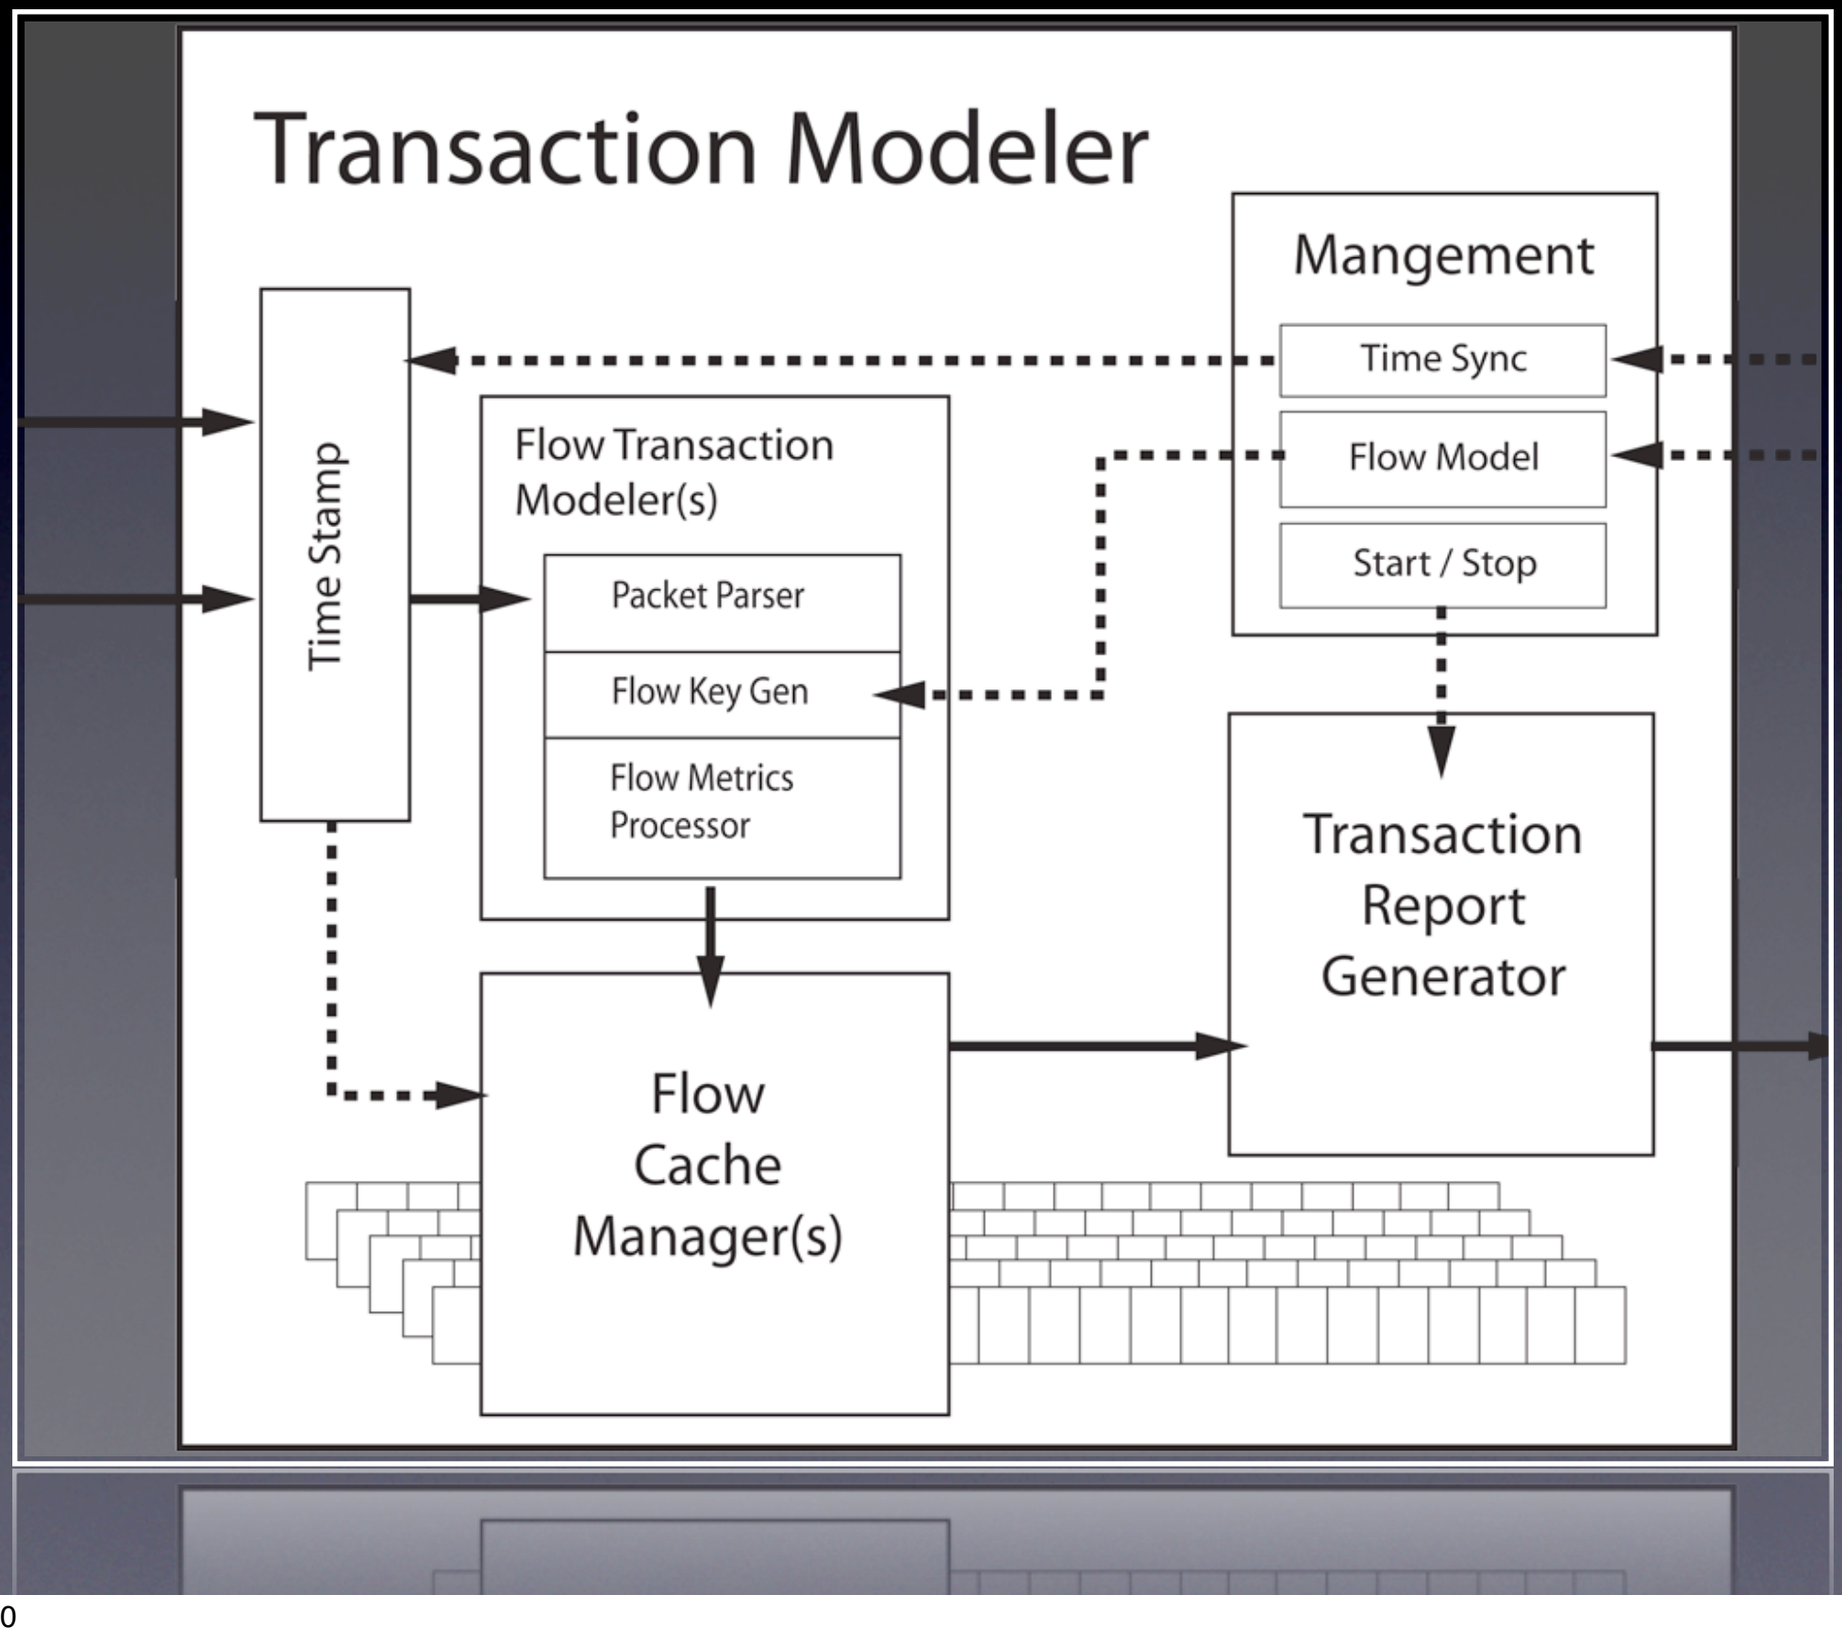
Argus System Design



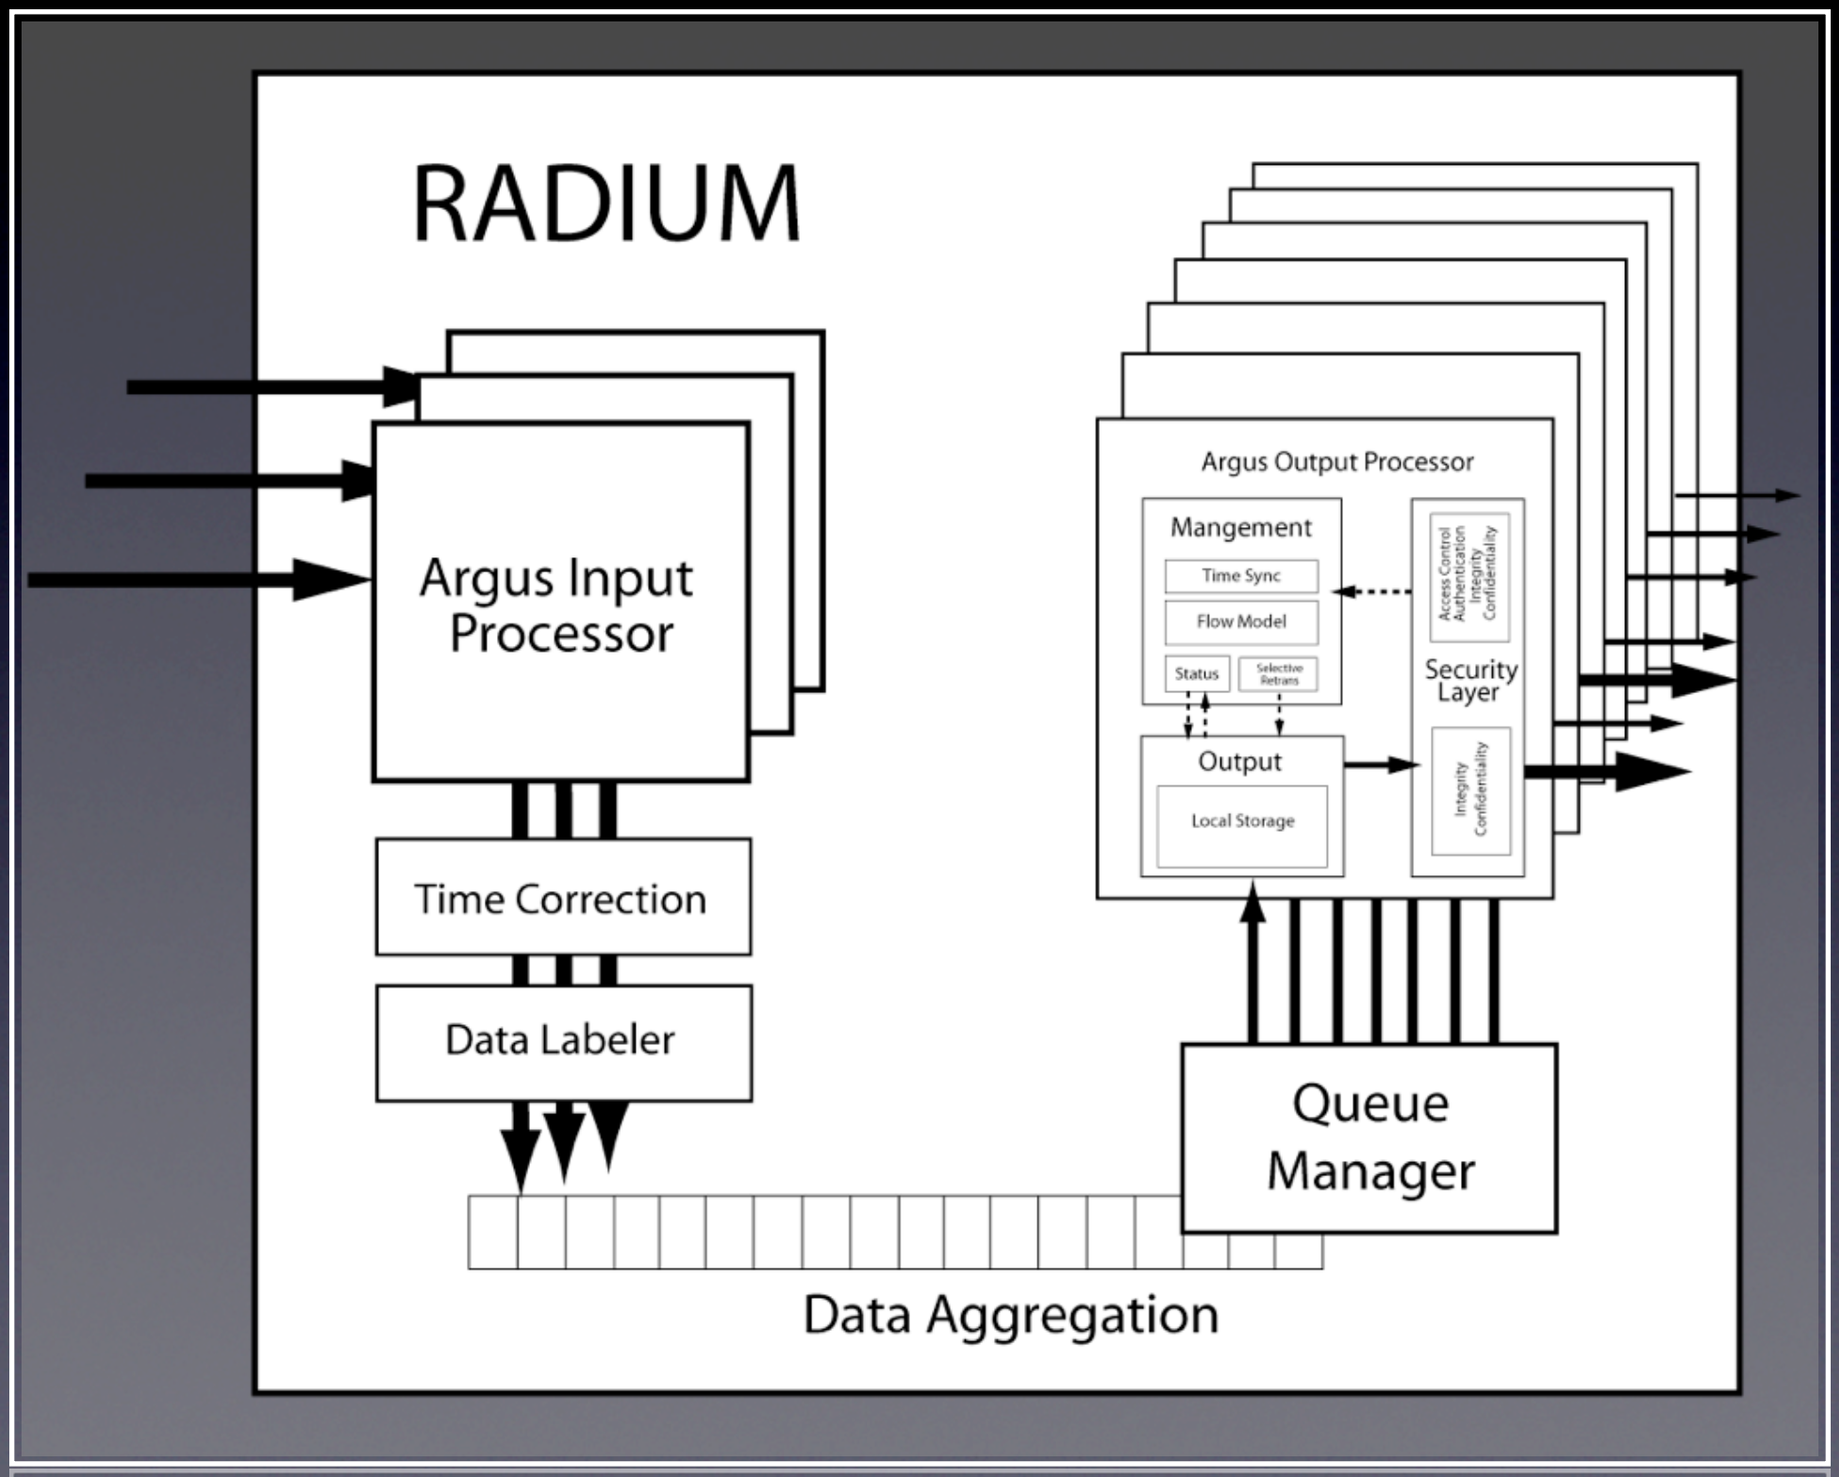
Argus Sensor Design



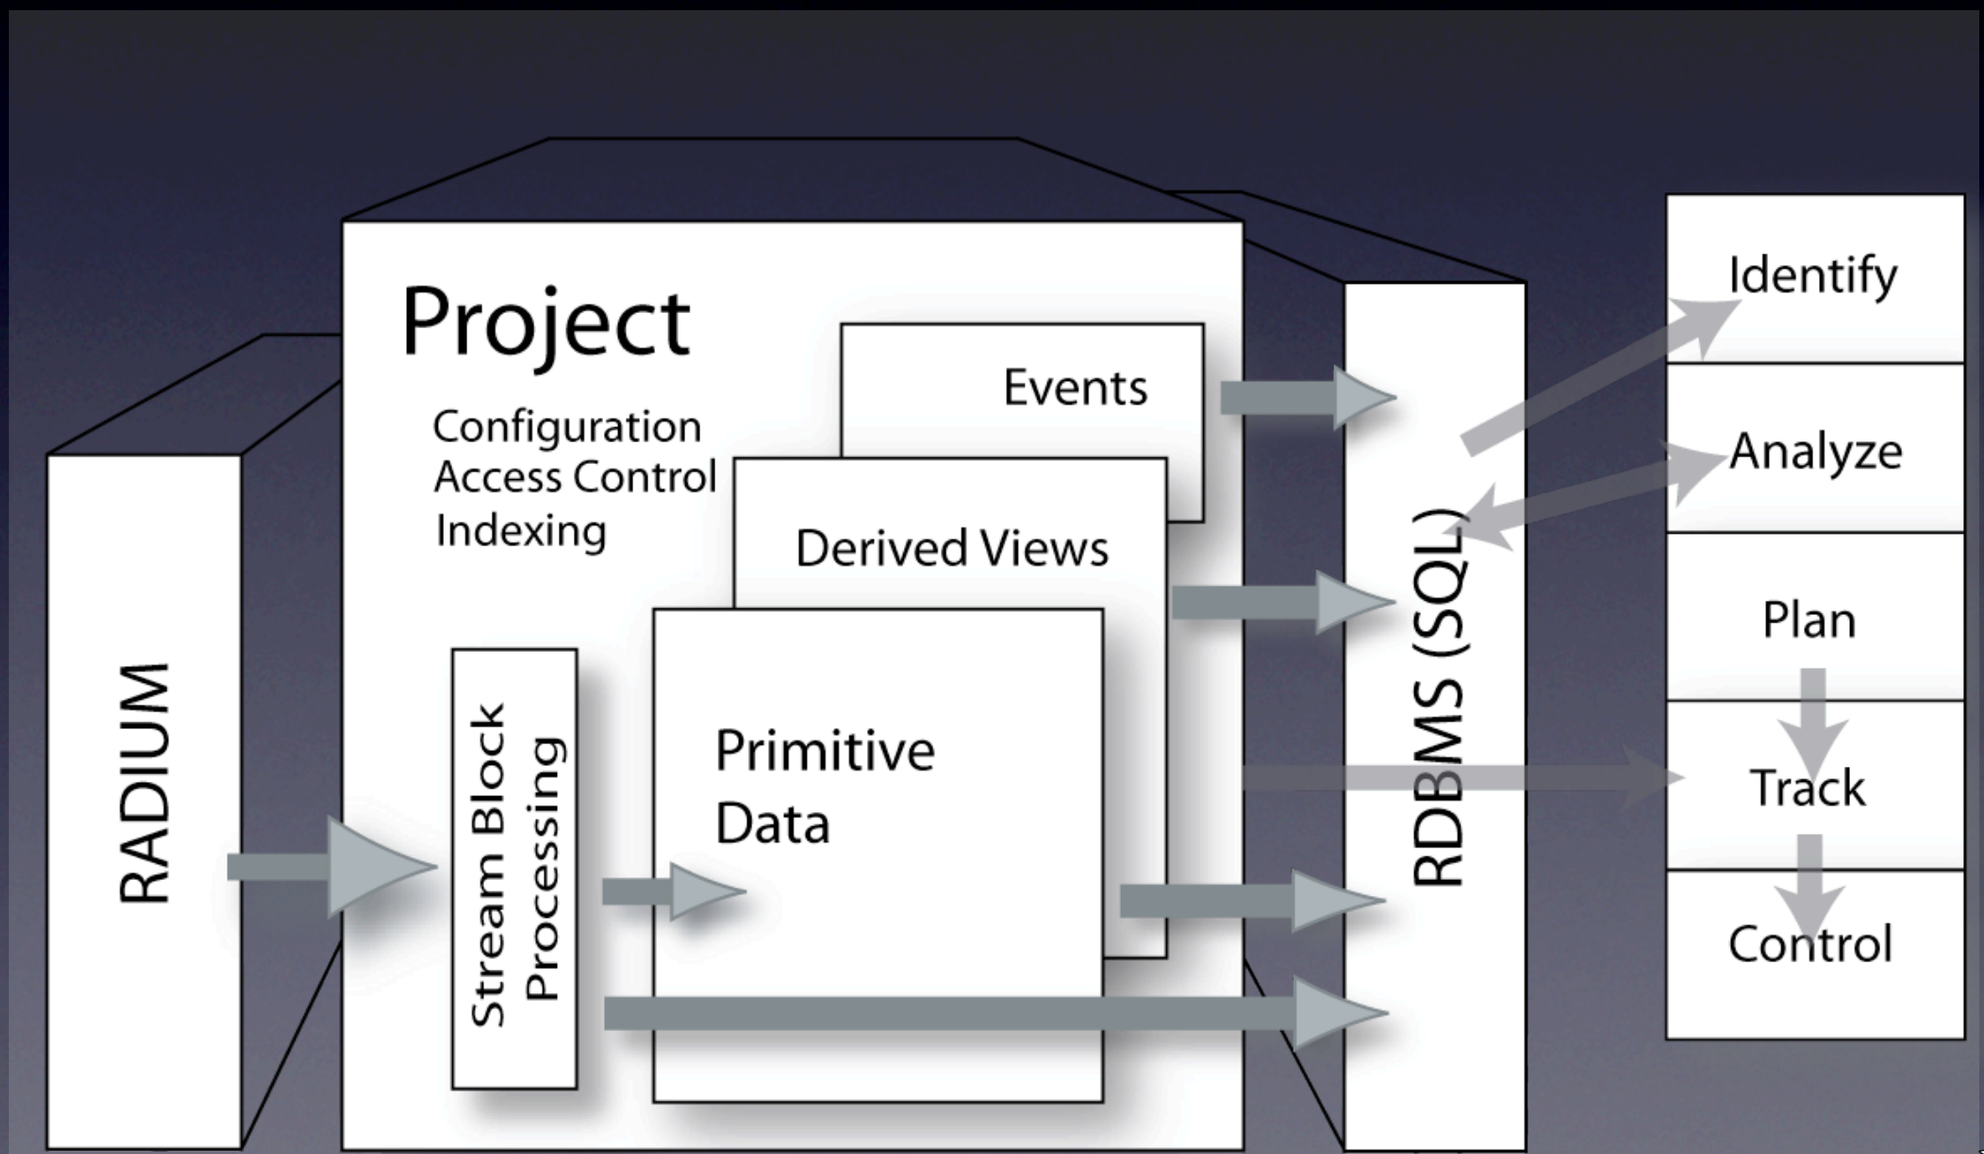
Argus Sensor Design



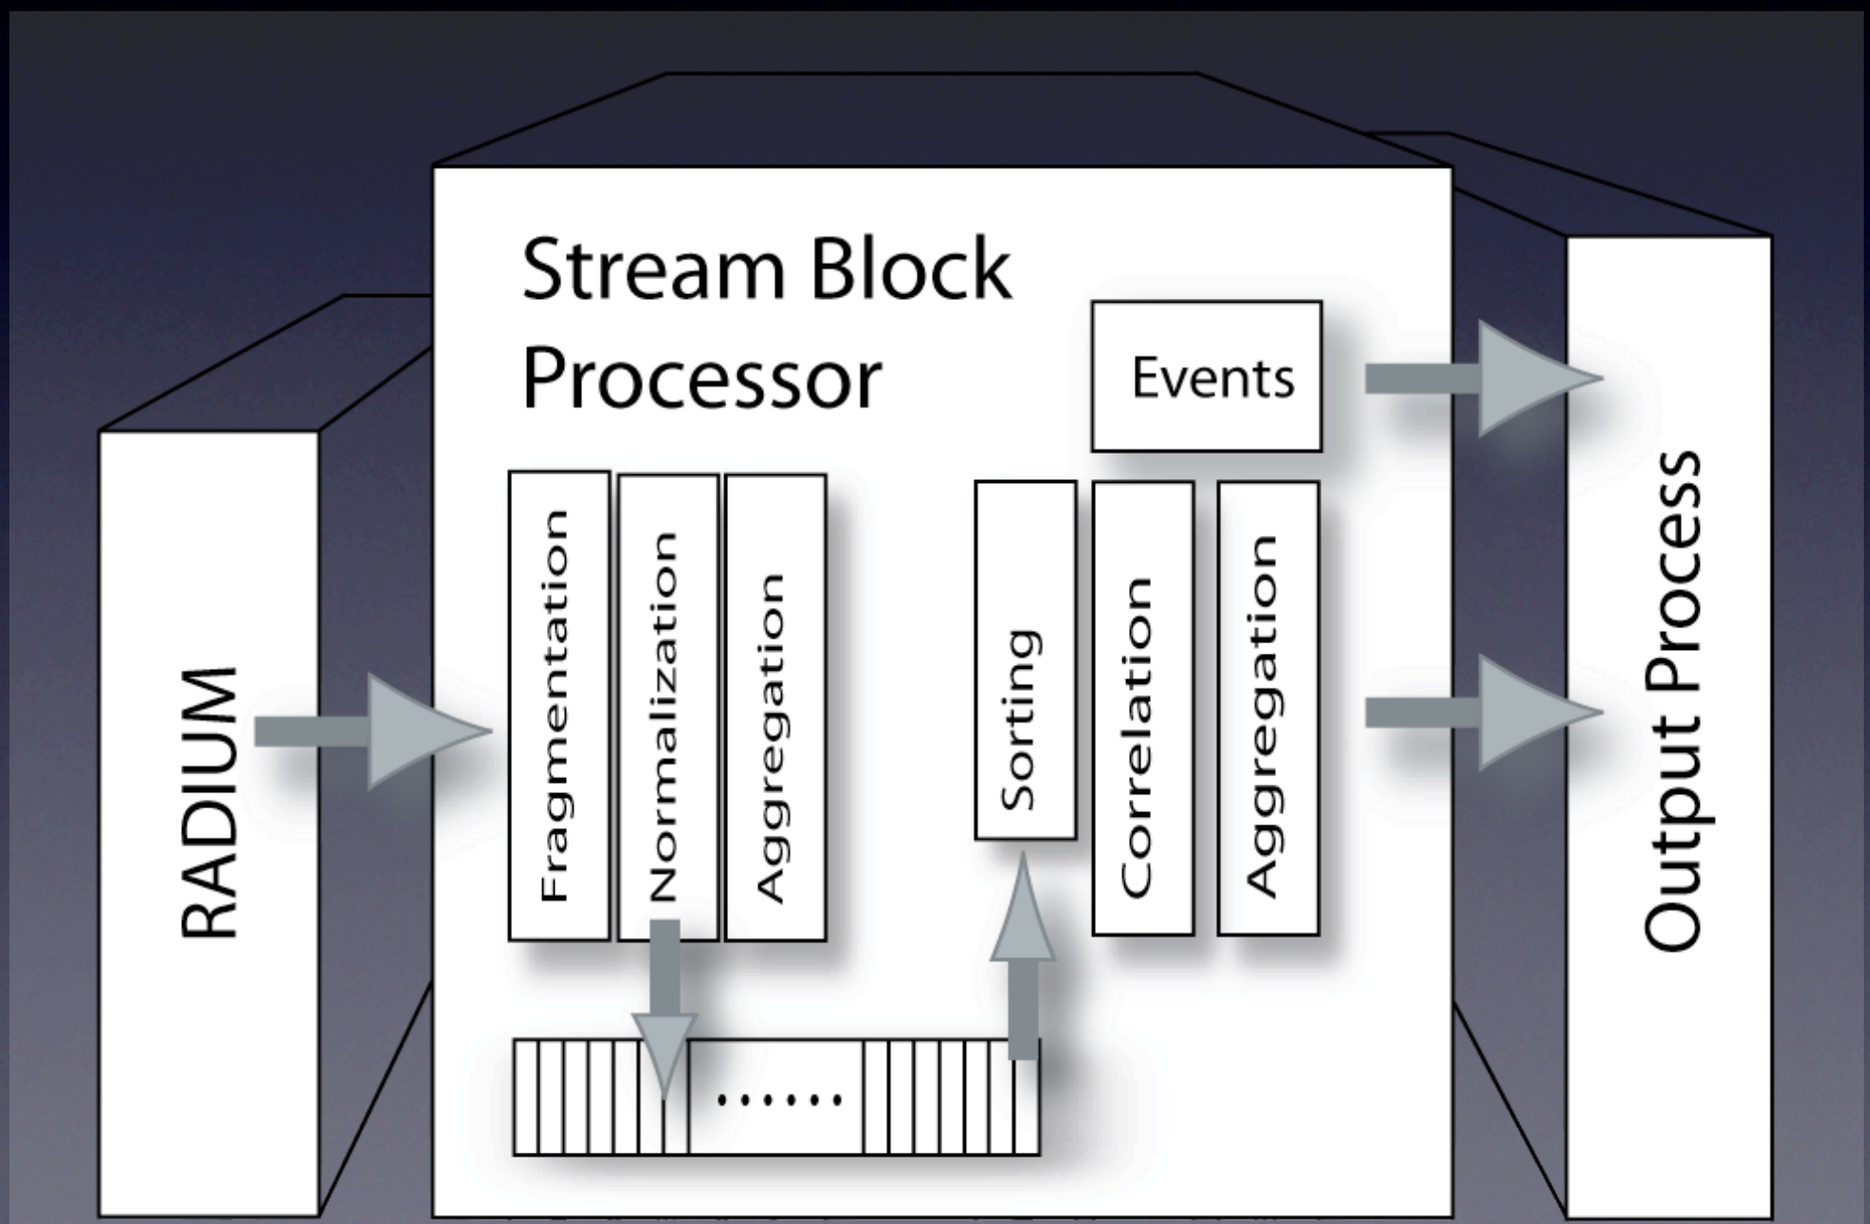
Radium Distribution



Argus Processing Design



Stream Block Processor Design



Data Generation





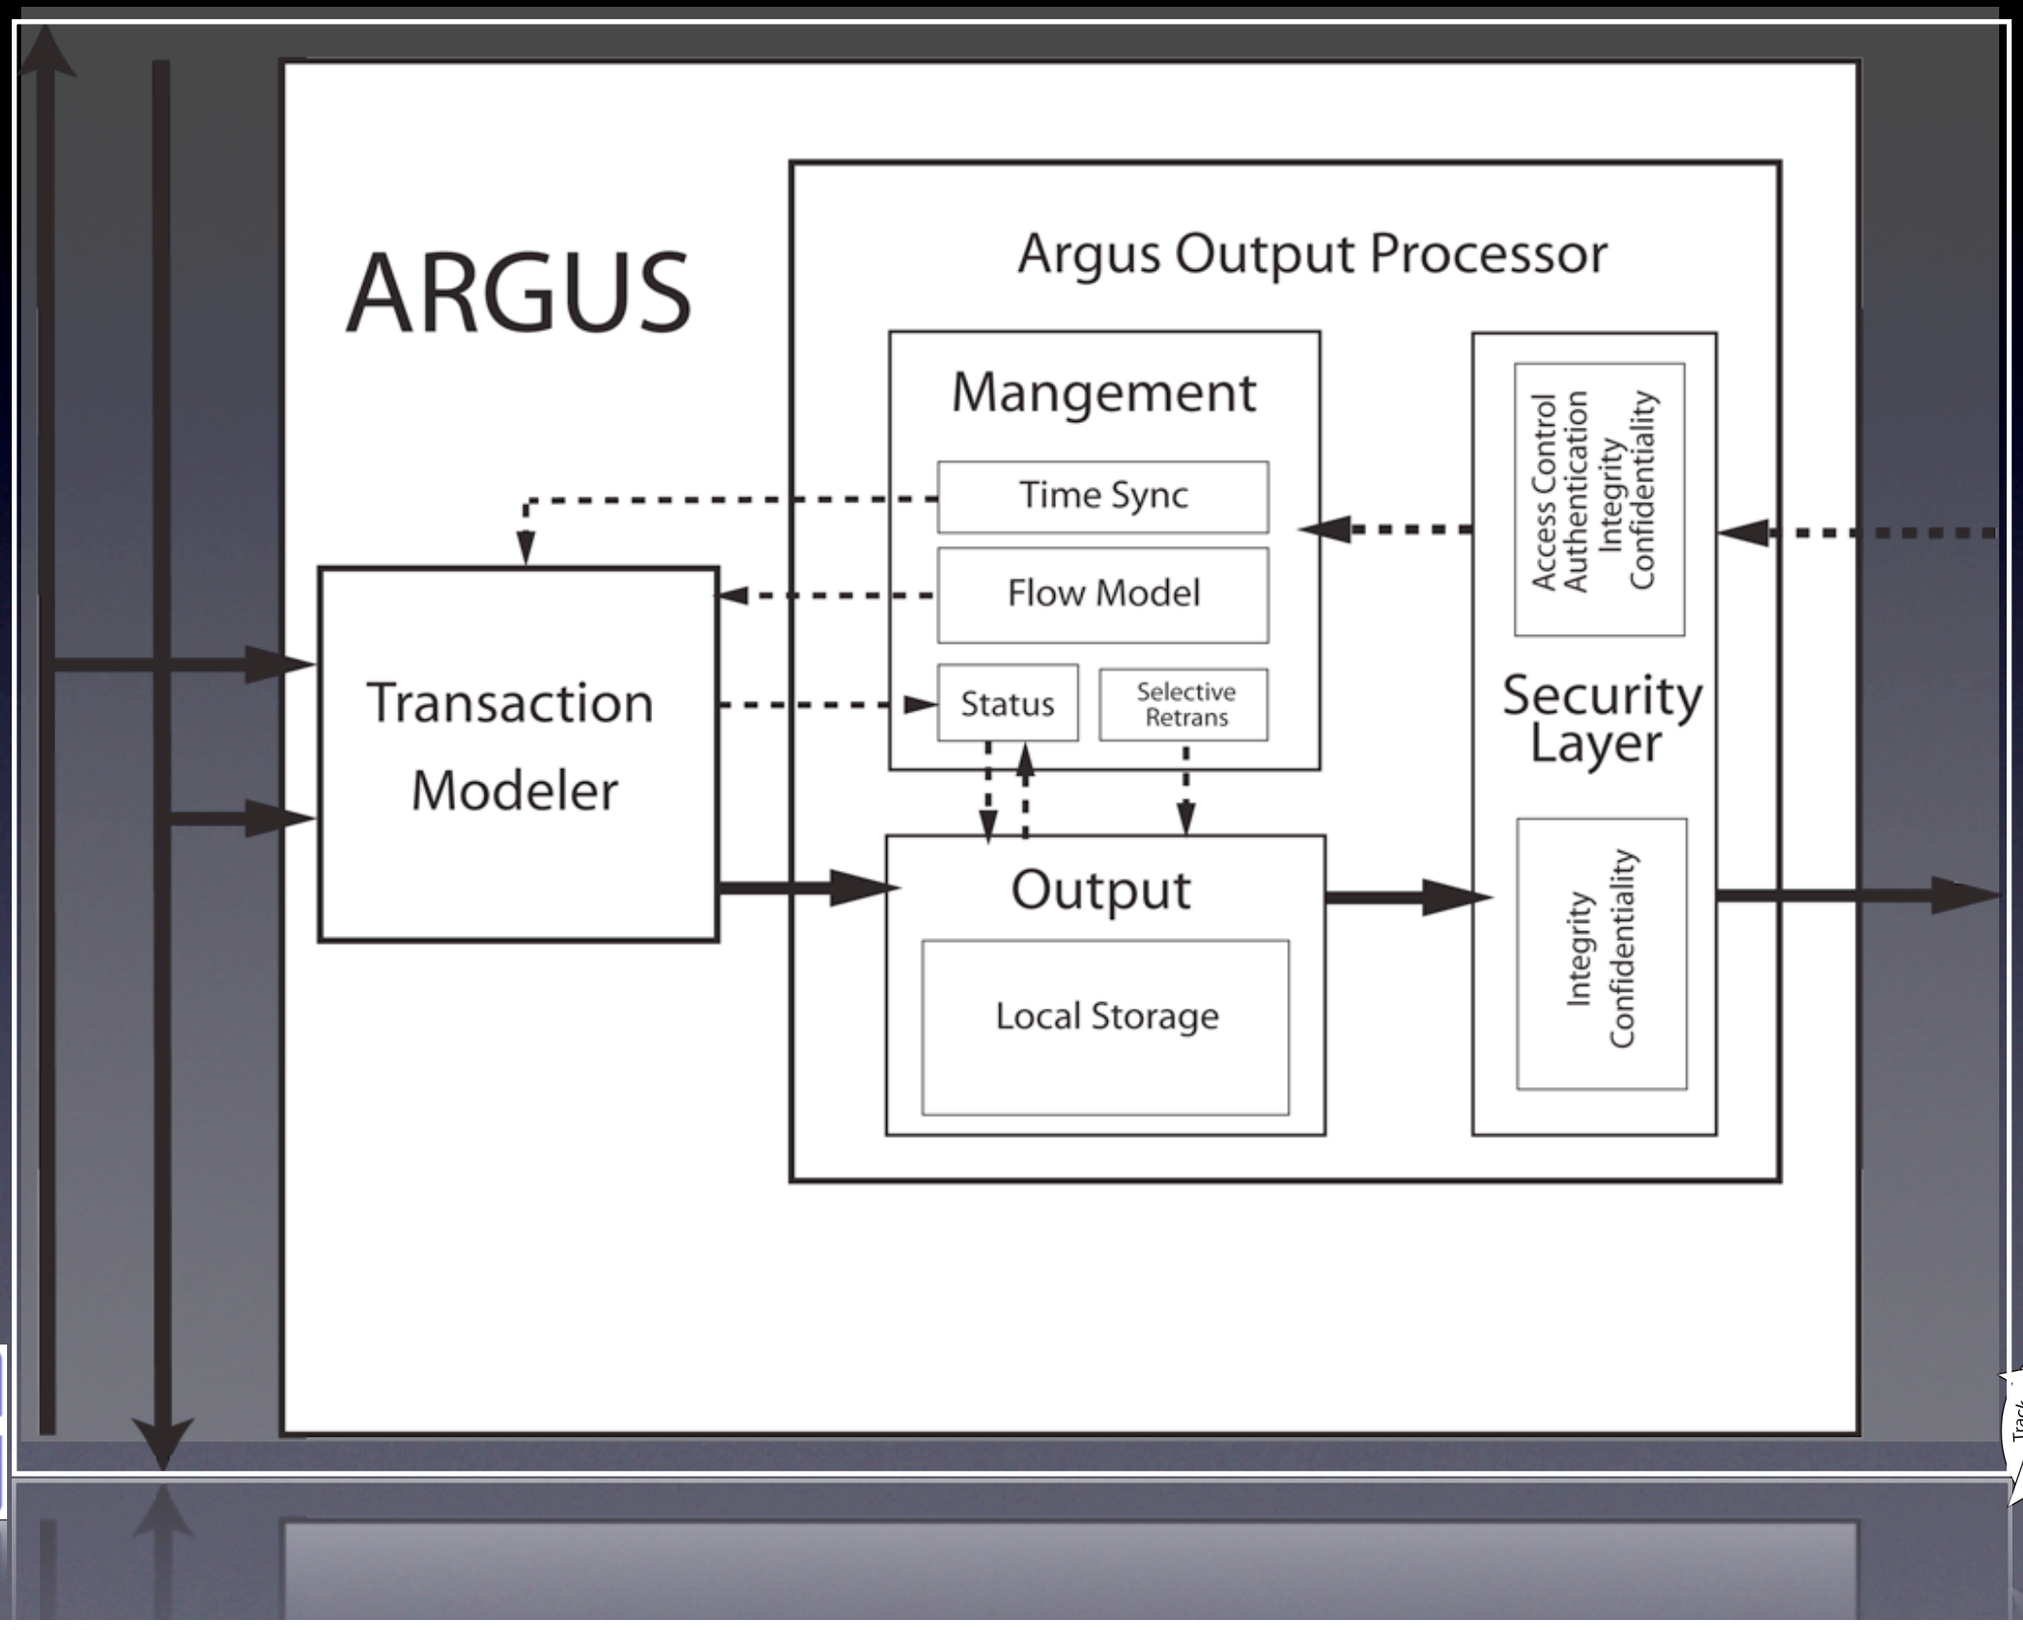
Data Generation

- Packets to Flows
- Getting Started with Argus
- Argus Deployment
- Configuration
- Running Argus

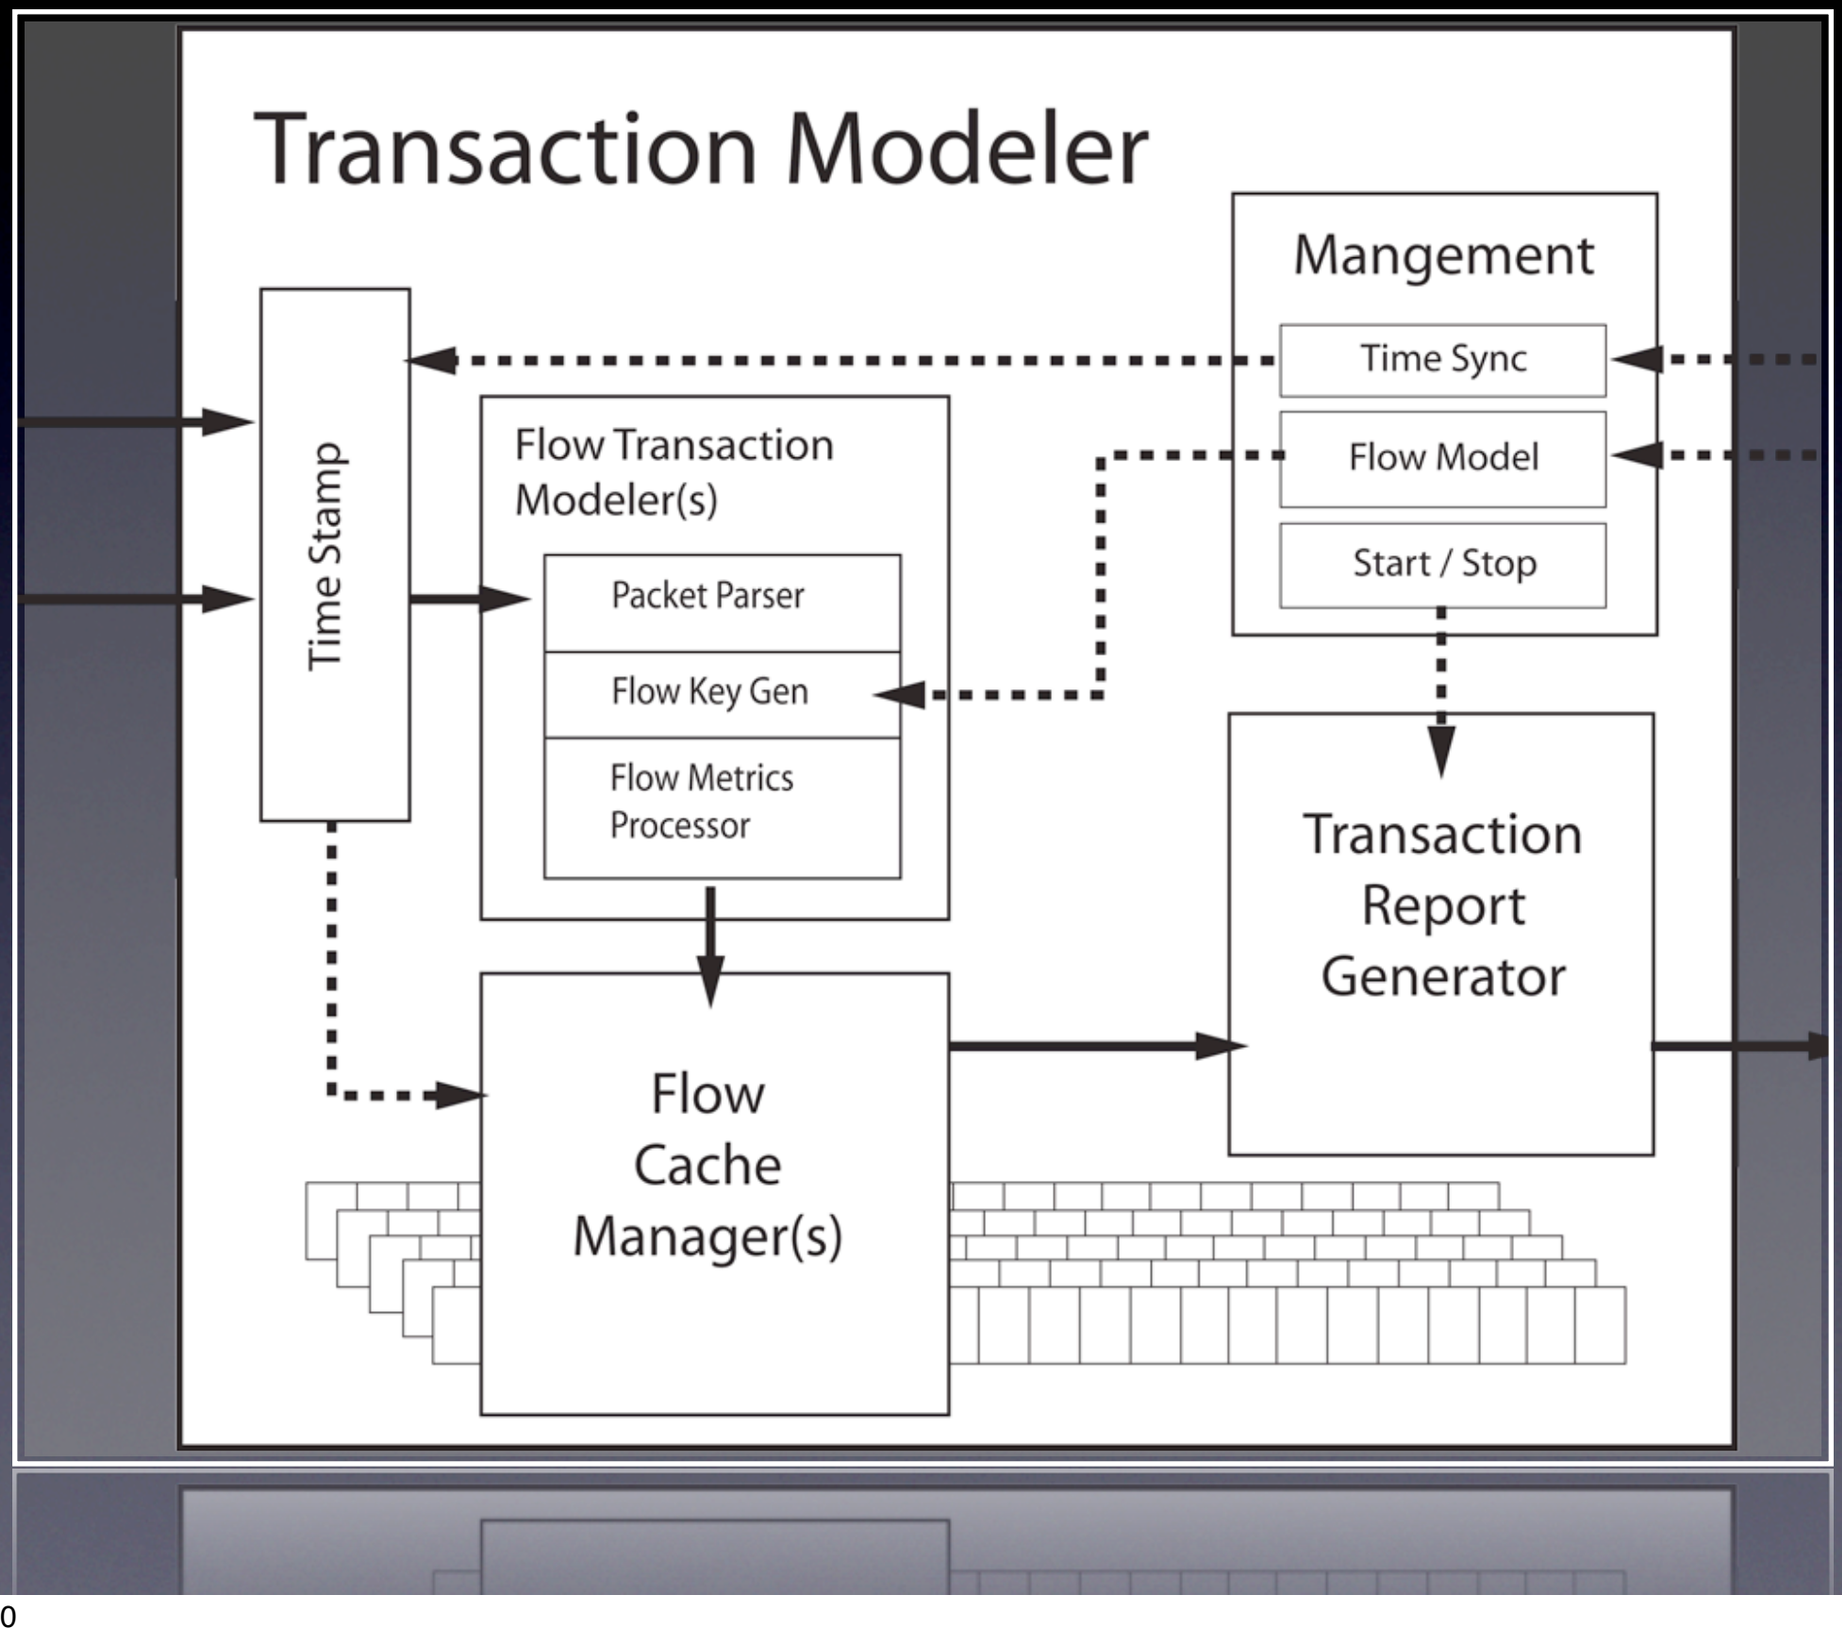




Packets to Flows



Argus Sensor Design



Packets to Flows

- Packet Timestamping
  - Methodology, Time Synchronization and resolution
- Packet Header Parser
  - Multiple flow tracking strategies determines parser
  - Supports OSI, IEEE, IP and Infiniband packet formats
  - Innermost Layer 3 target header (service layer)
    - Complex encapsulation stacking
      - L2 -> L3 -> L2 -> L3 -> L4 -> L3 -> L3
      - Support protocol discovery
  - Limited by packet snap size
    - Argus supports complex packet capture support
    - Privacy issues
    - Control plane vs data plane parsing





Packets to Flows

- Flow Key Generation
  - All packets are classified into a flow of some kind
  - Argus supports 14 fundamental flow types
    - Not protocols, flow types (P, P1-P2, Multicast/Unicast, etc....)
    - Bi-directional support for all flow types (when they exist)
    - Bi-direction flow keys for all supported encapsulations
- Flow Key is "key" to all flow tracking
- One packet one flow rule
  - Simplify flow machine call structure
  - Control plane is the exception
    - ICMP packet accounted for in ICMP flow
    - ICMP state mapped to flow identified in contents





Packets to Flows

- Flow Metrics Processor
  - Metric and attribute generation
    - Some metrics can be derived from packet itself
      - Packet size, application demand, reachability
    - Others require state
      - connectivity, availability, RTT, rate, loss, jitter
  - Flow attribute (re)assignments
    - Flow state machine tracking
    - Dynamic attribute tracking
- Flow Cache Manager
  - Controls reporting of flow status
  - Controls dynamic flow redefinitions/reassignments





Getting Started

- http://qosient.com/argus
- ‘Using Argus’ and ‘Getting Argus’ Links
- Argus documentation
  - Man pages provided in distribution
  - HOW-TO and FAQ on the web site.
  - Argus developers mailing list
    - argus-info@lists.andrew.cmu.edu.
    - Most questions are answered here
    - Email carter@qosient.com





Getting Argus

- http://qosient.com/argus/downloads.htm
- Many linux distributions have a port
- Current stable version is argus-3.0.2
- Provided as tarball source package

- Depends on:
  - libpcap - http://tcpdump.org/release
  - bison - http://www.gnu.org/software/bison/bison.html
  - Client ragraph() requires perl & rrd_tool





Making Argus

- Simple installation
  - ./configure; make
- Complex environments
  - Read ./README and ./INSTALL
  - Cygwin/OpenWRT
- Support standard autoconf options
  - ./configure --help
  - Common variations
    - prefix=/your/destination/directory
    - SASL Support
    - Native compiler options





Installing Argus

- Simple installation
  - make install
- ./INSTALL describes some complex examples
- /etc/argus.conf
- System startup configuration
  - Linux chkconfig.1 support
  - MacOS X /Library/StartupItems support
- RPM support - ./lib/argus.spec





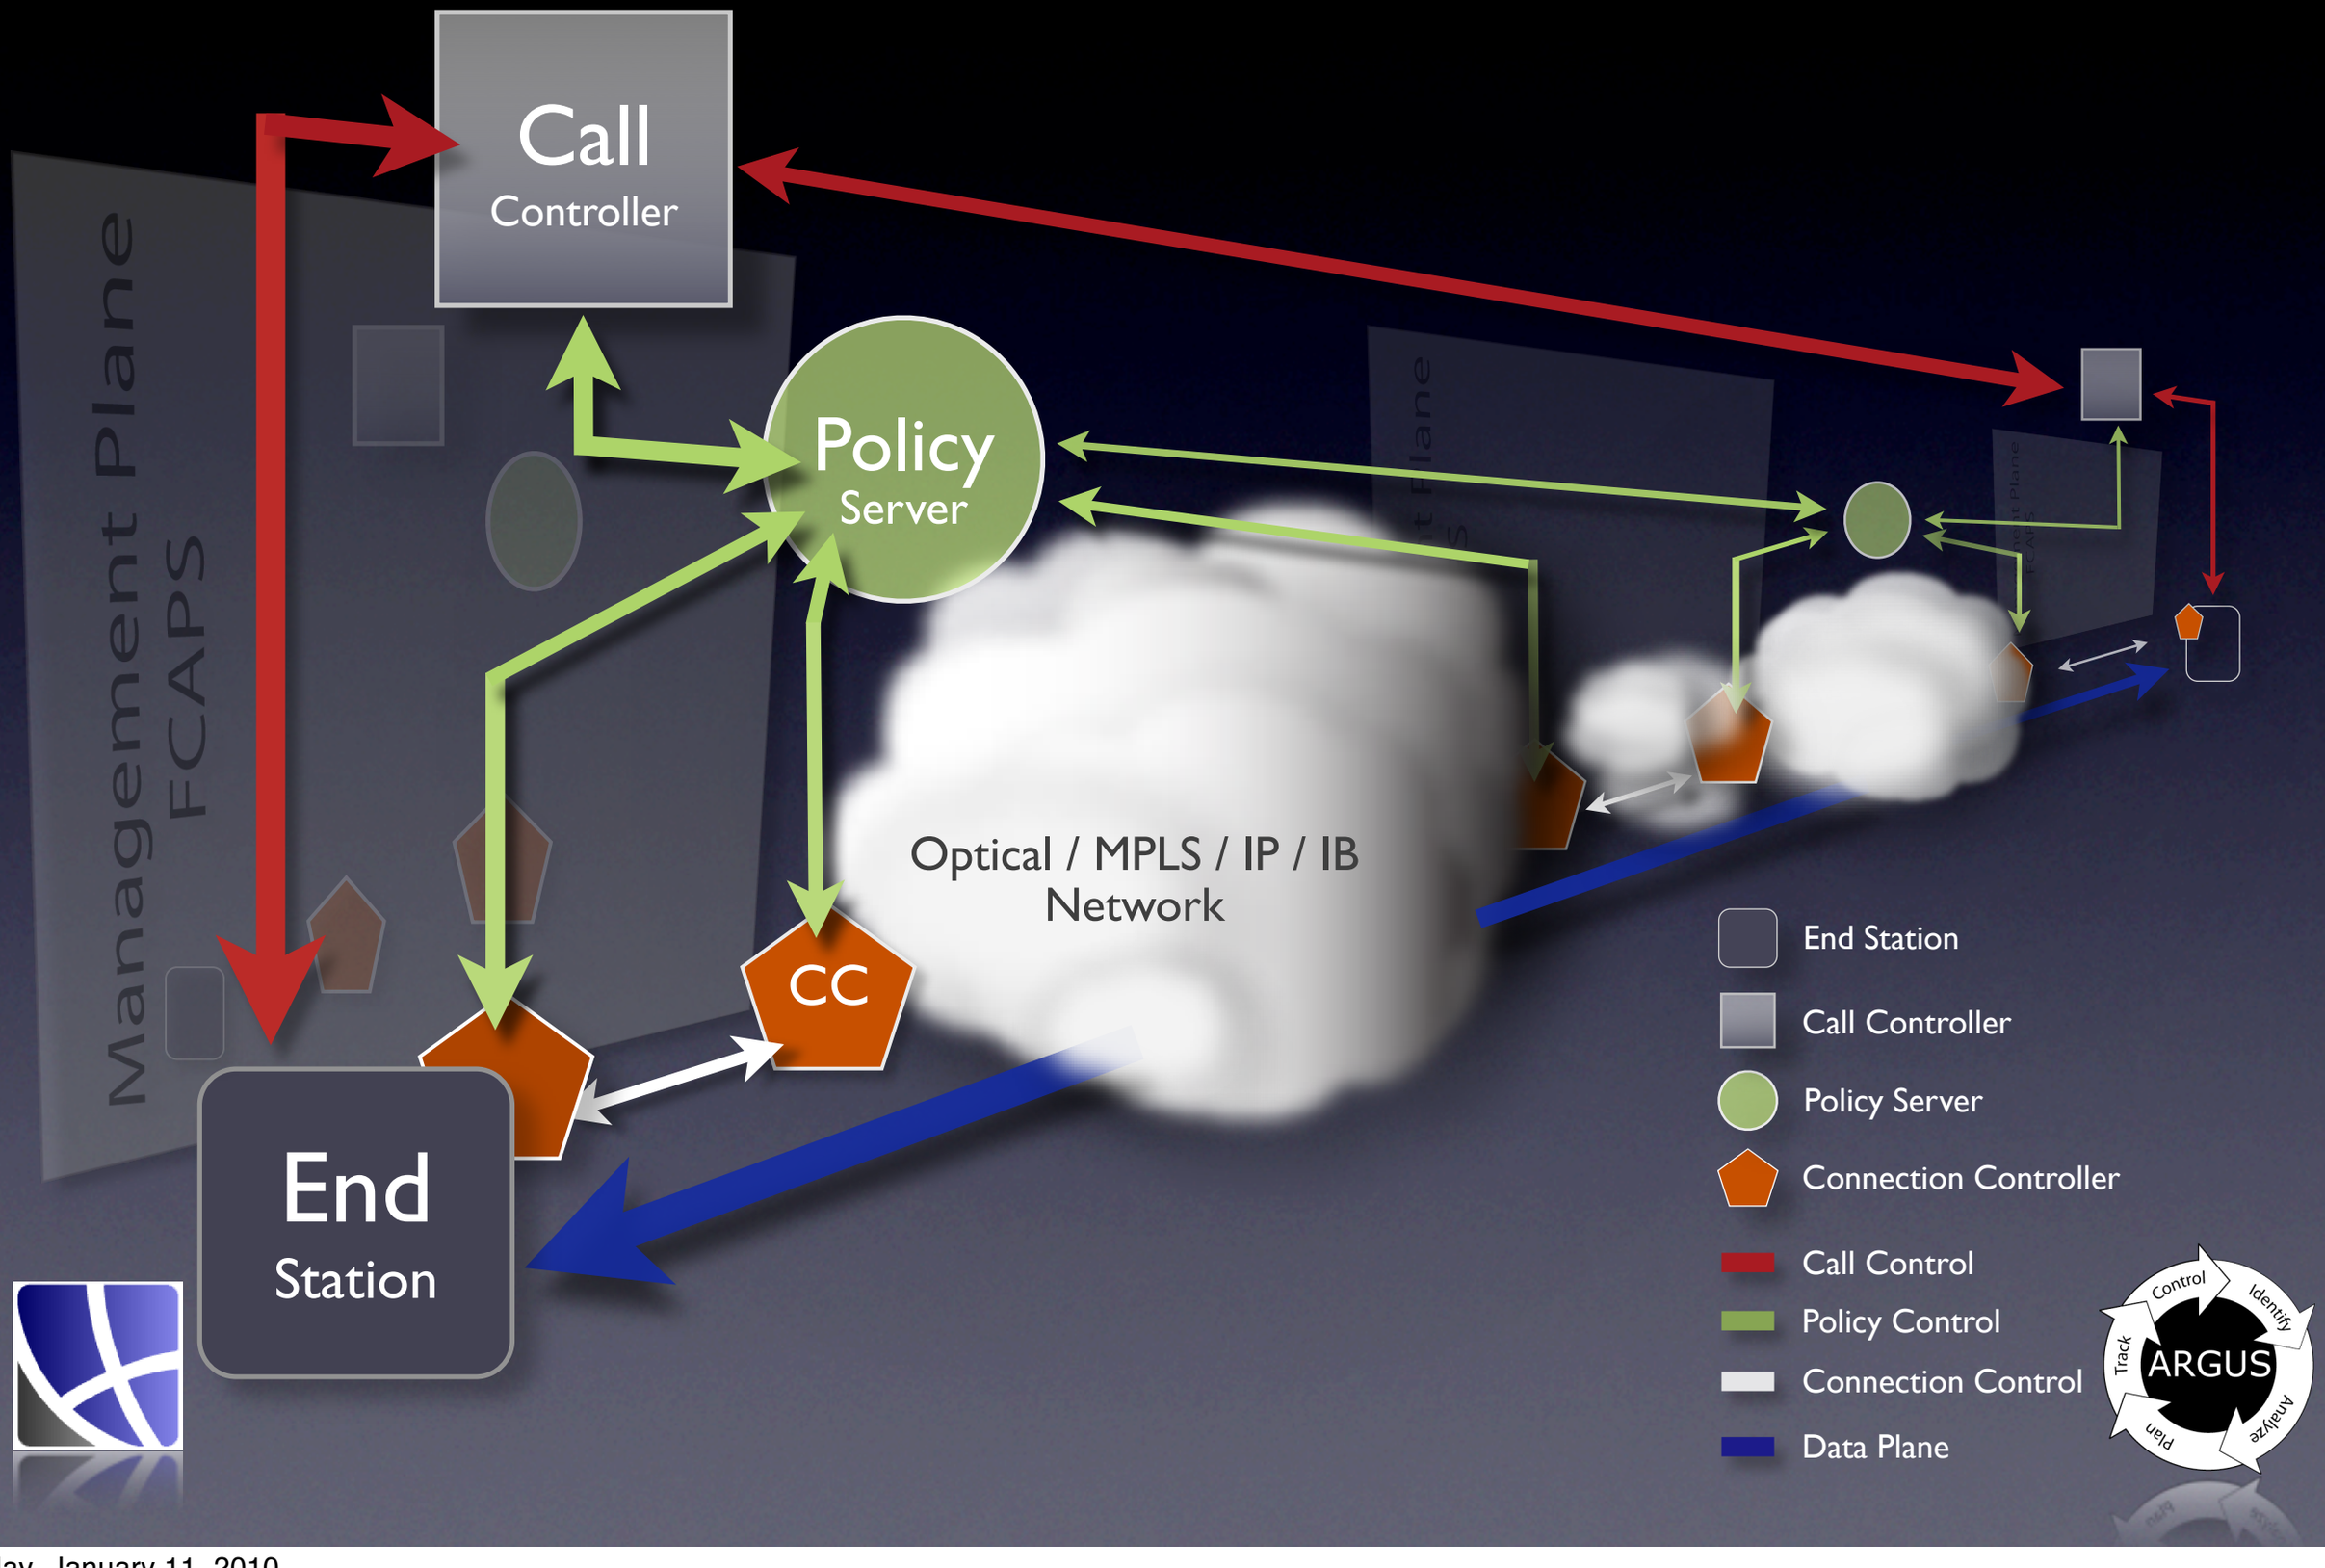
Deployment

- Monitoring Strategies
  - Enterprise Border Monitoring
  - Subnet Monitoring
  - End System Monitoring
  - Complex/Comprehensive Monitoring

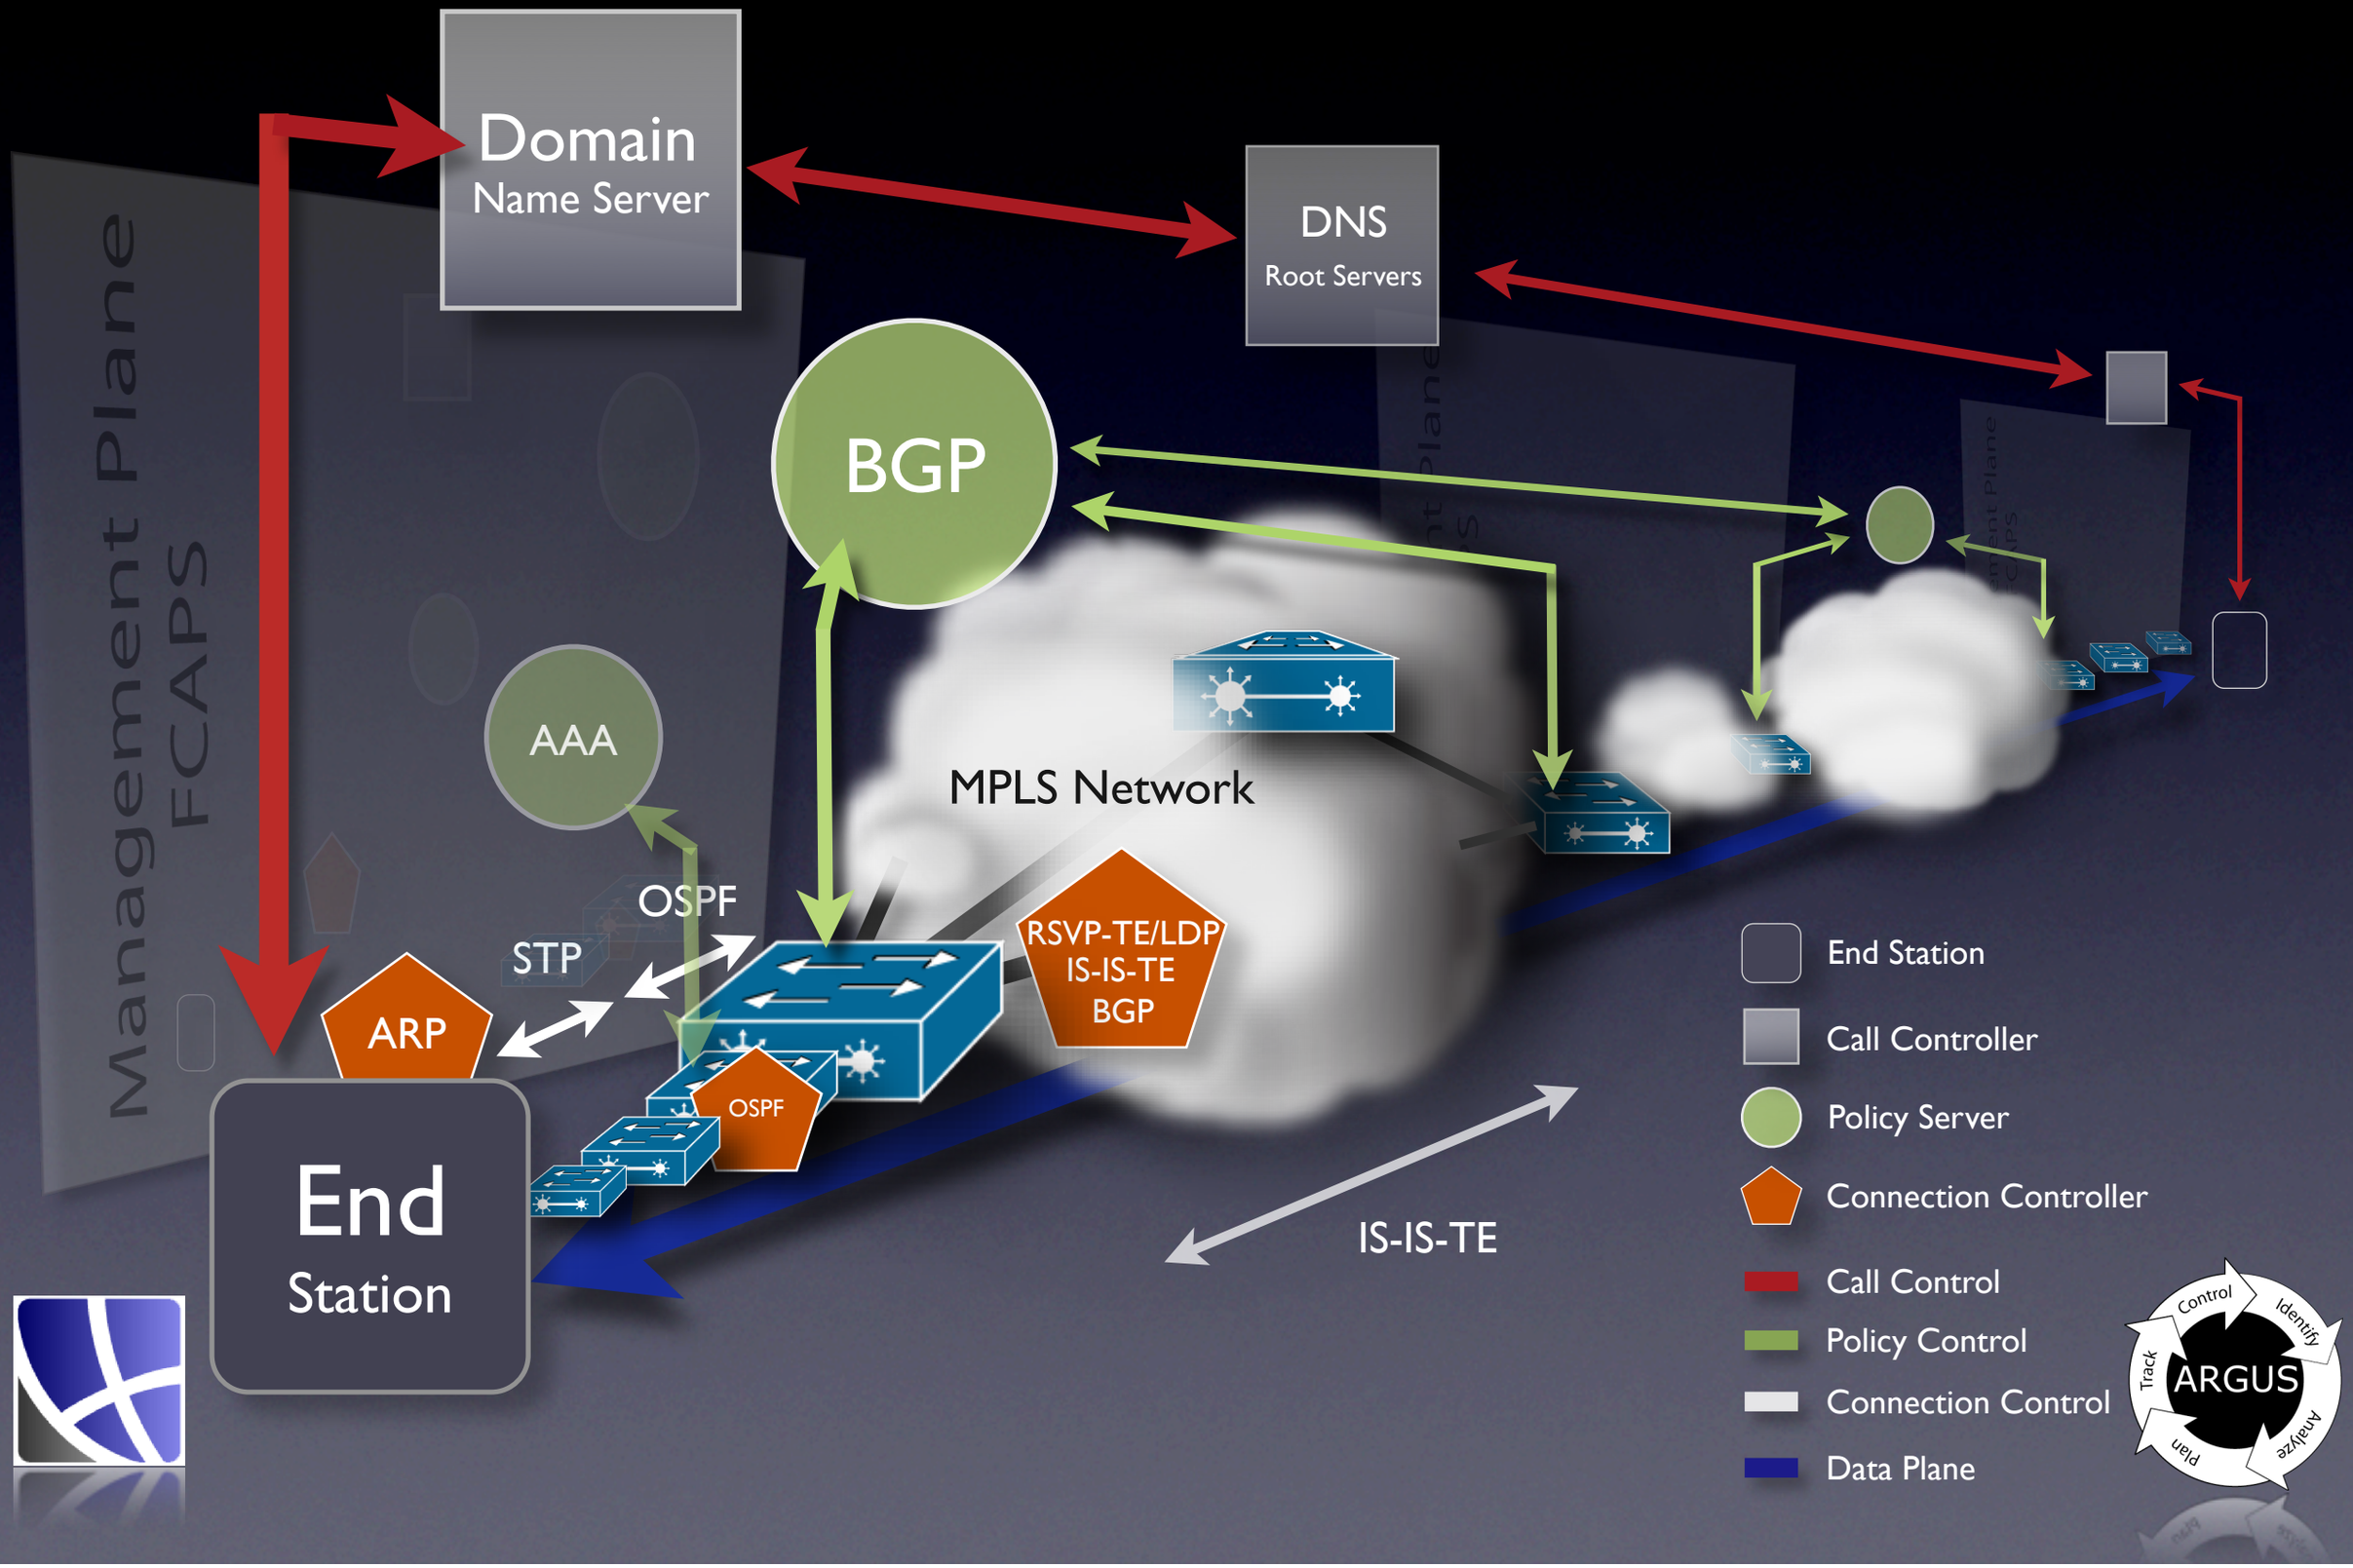




Network Reference Model

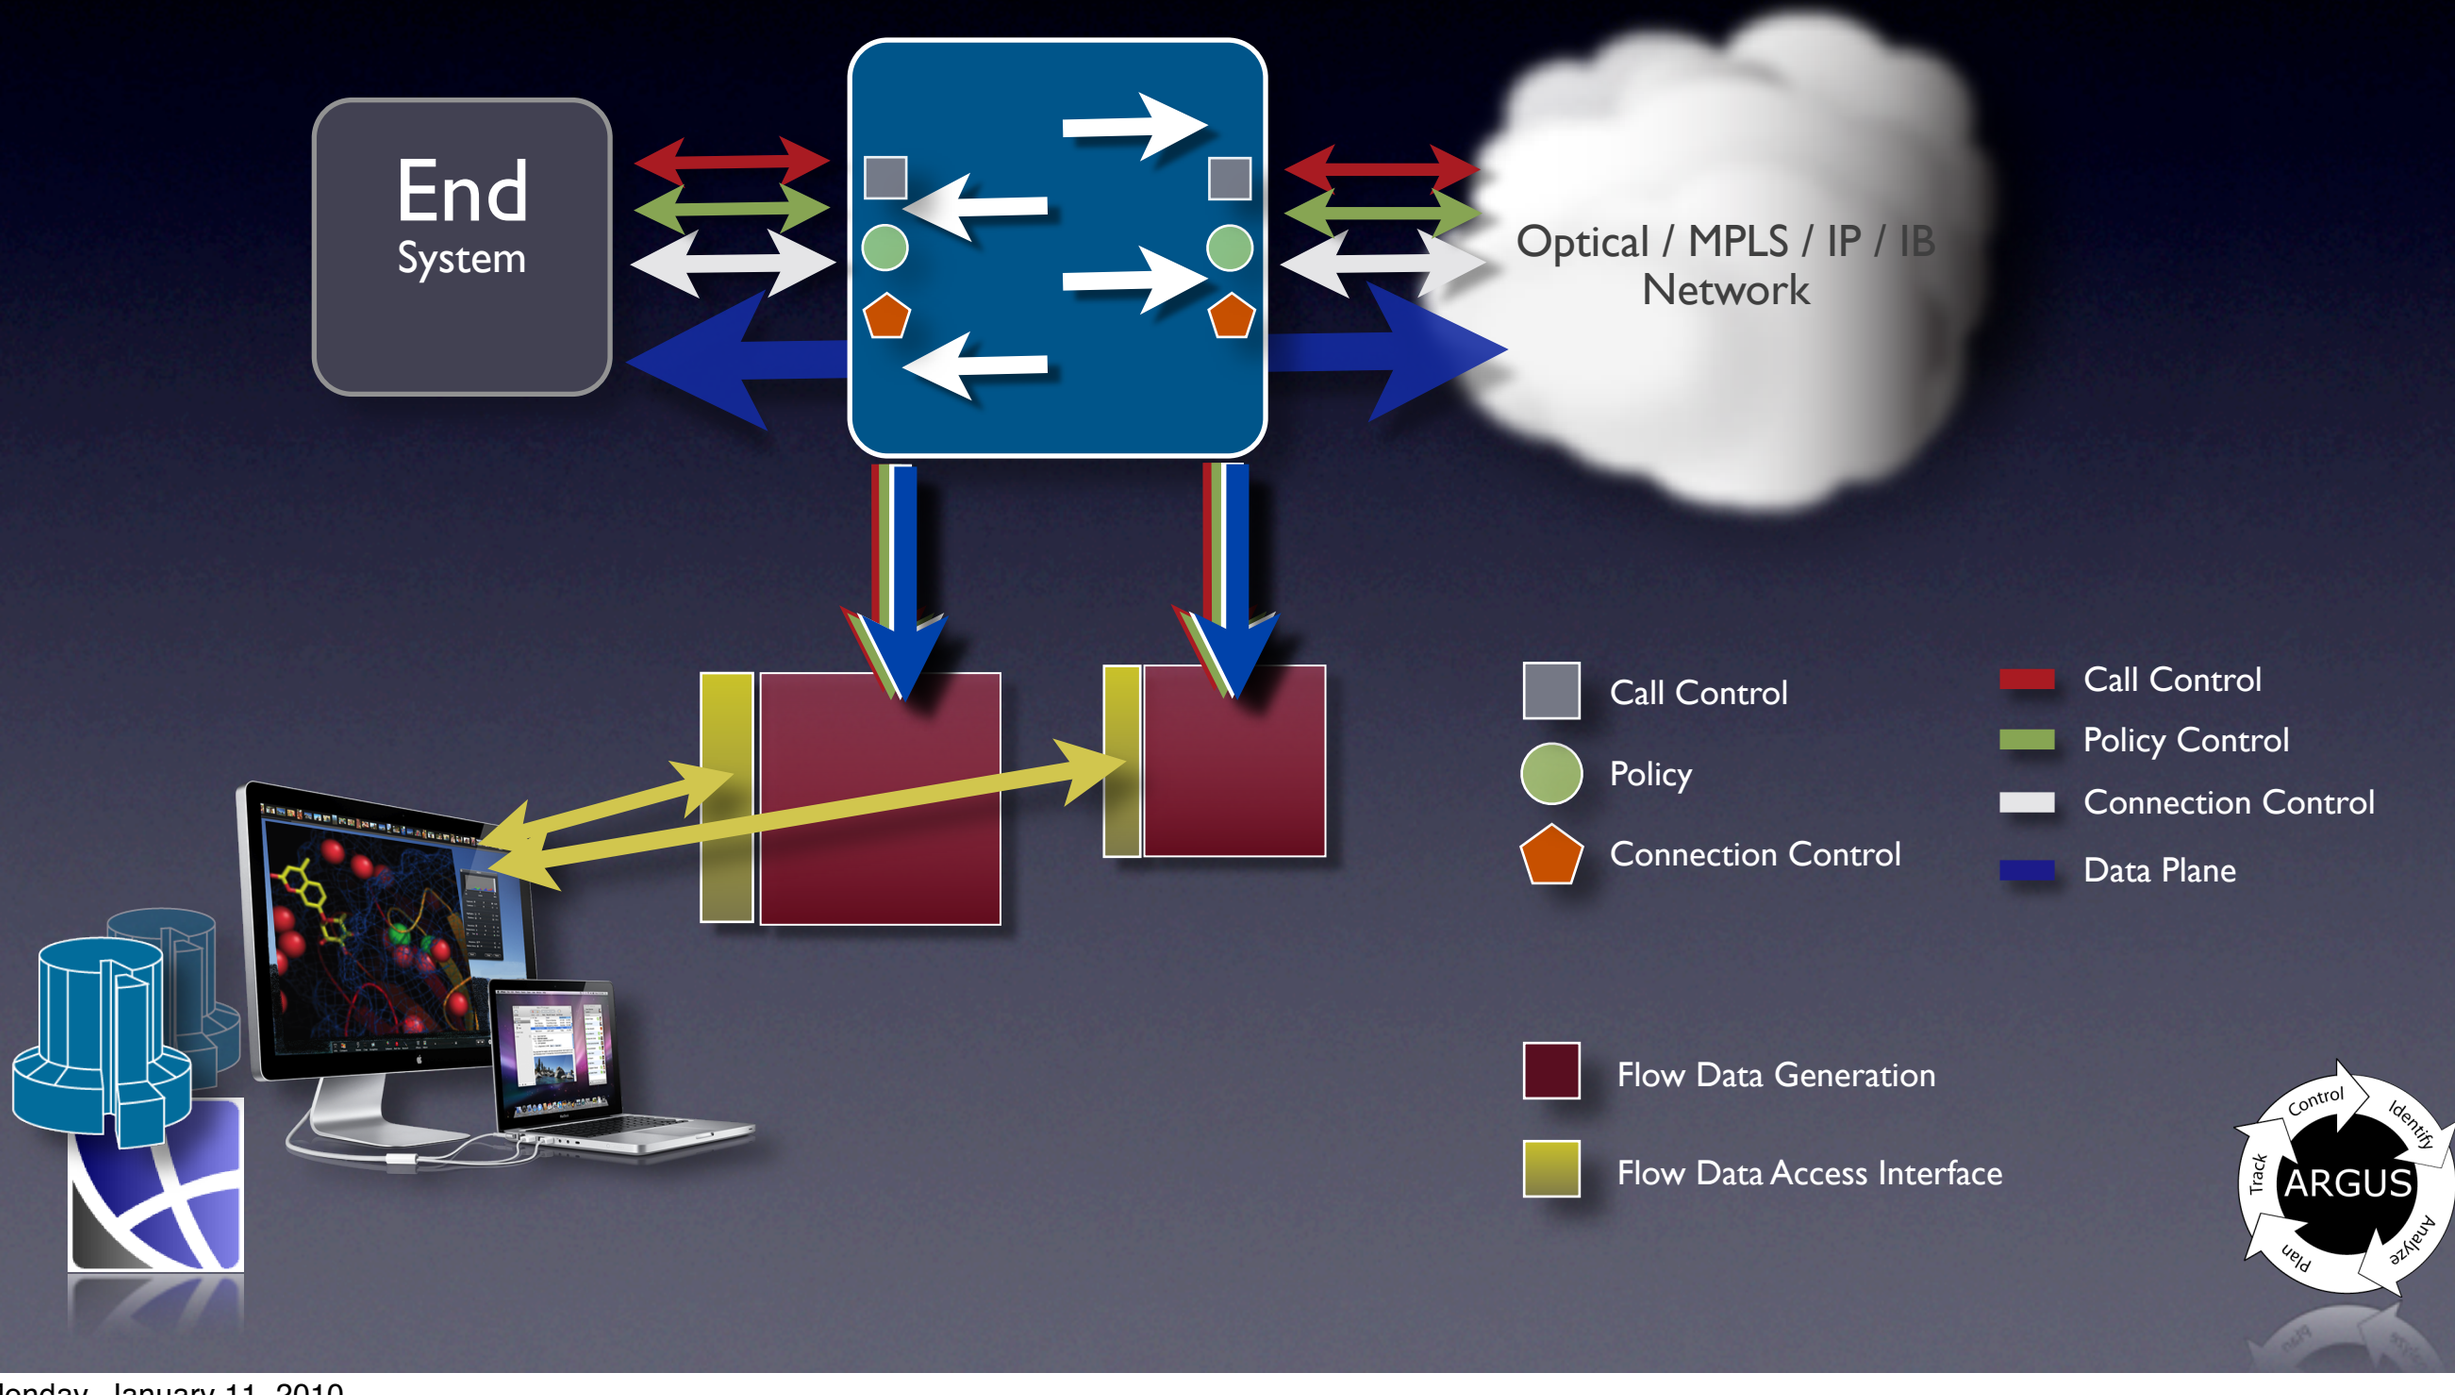


Internet Networking Model



Enterprise Border Monitoring

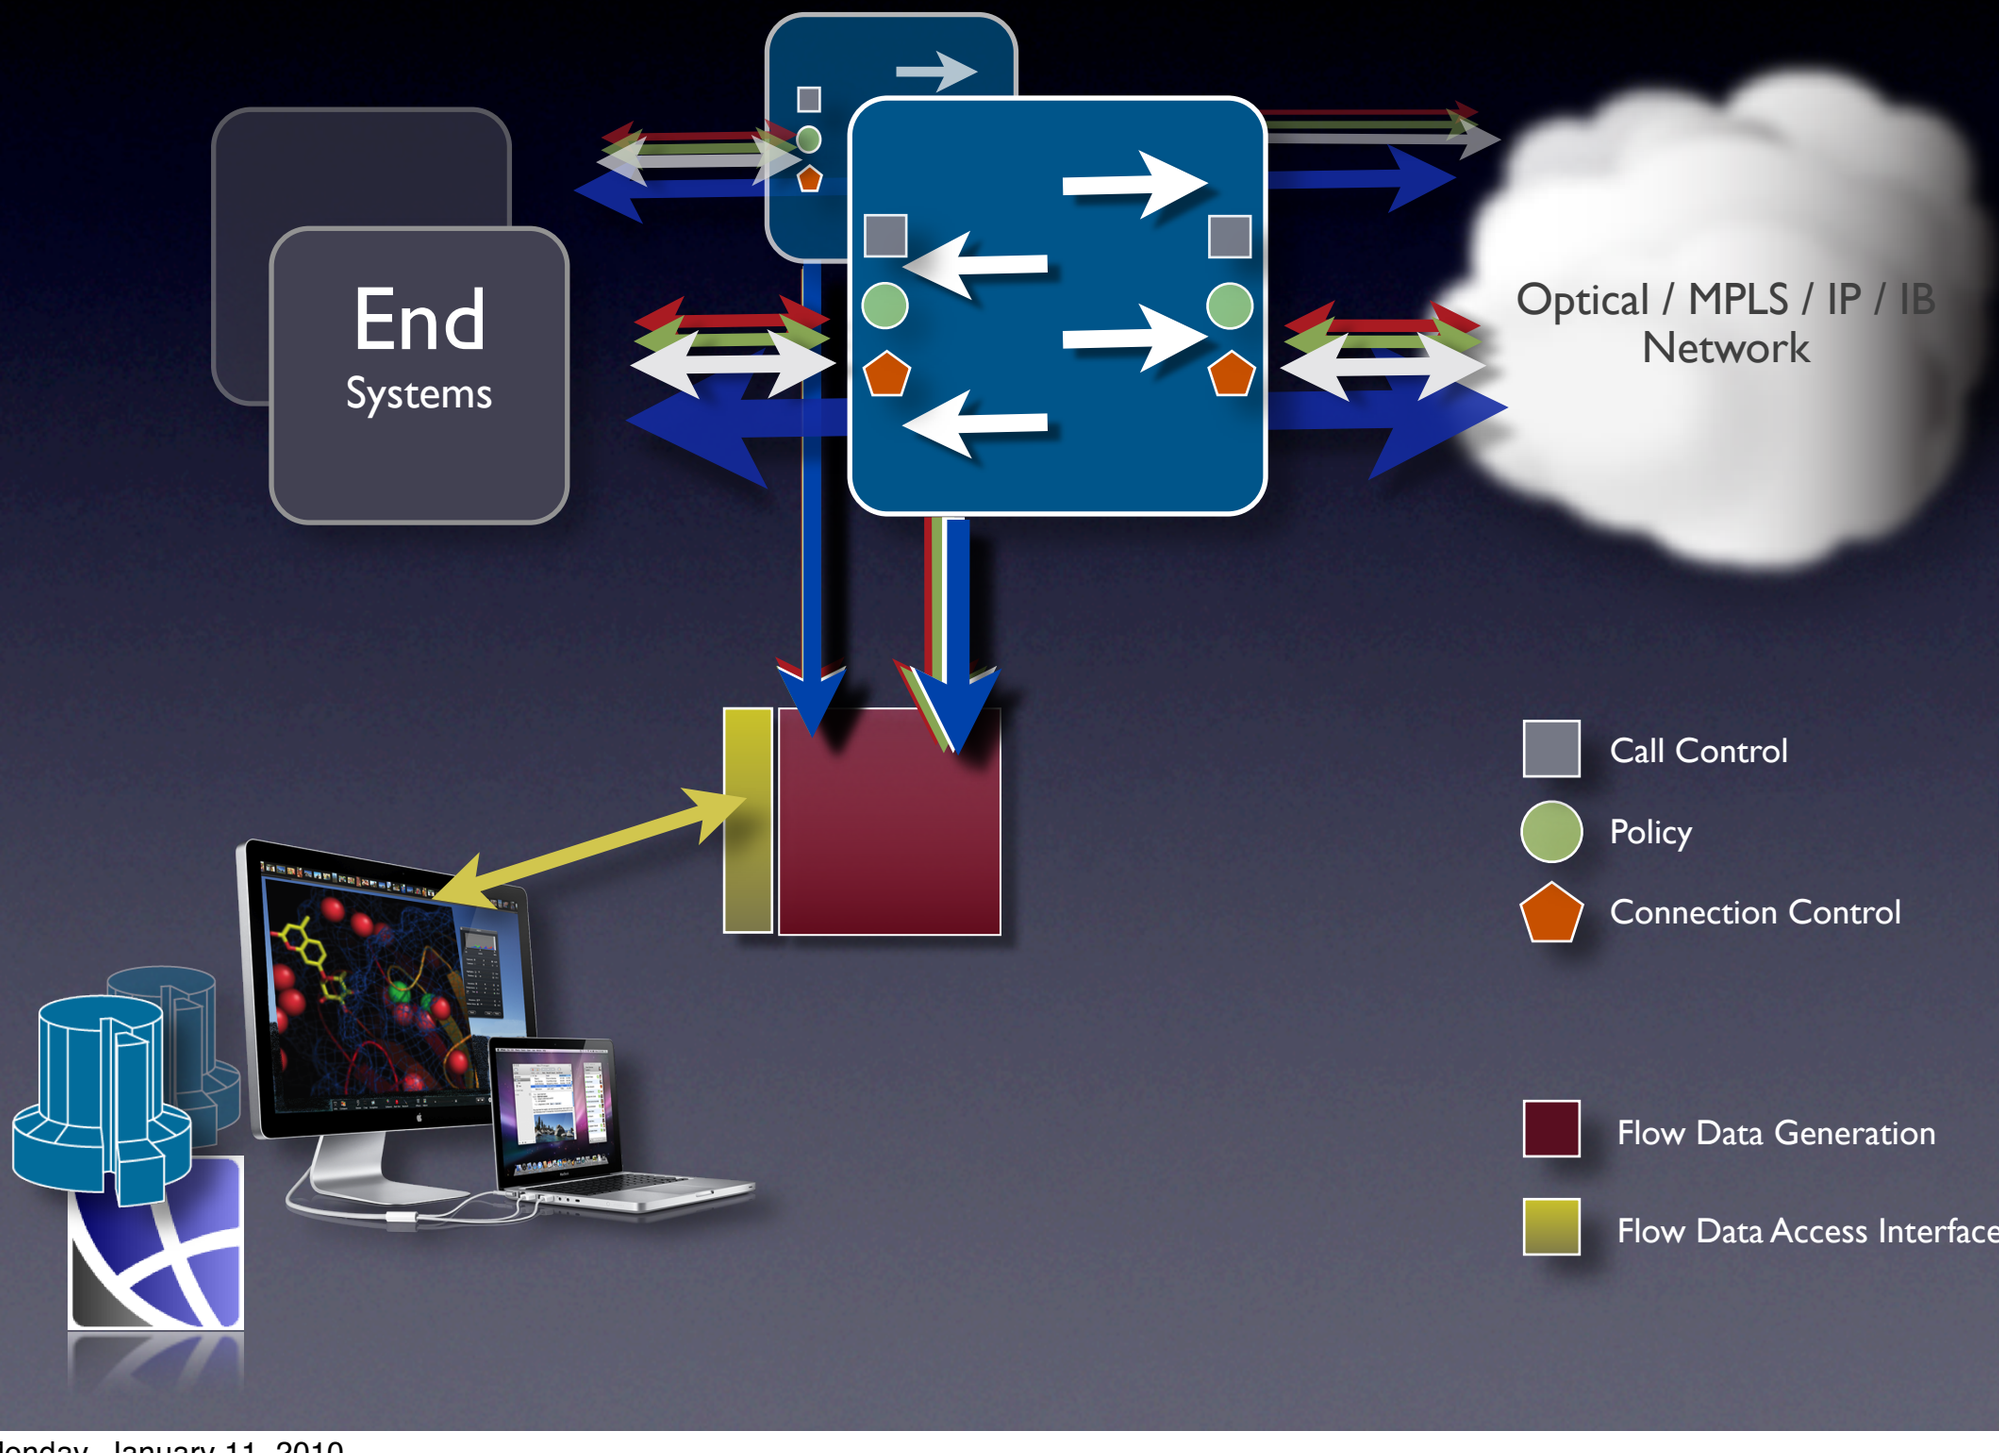
Internal/External Strategies



Enterprise Border Monitoring

Asymmetric Routing Strategies

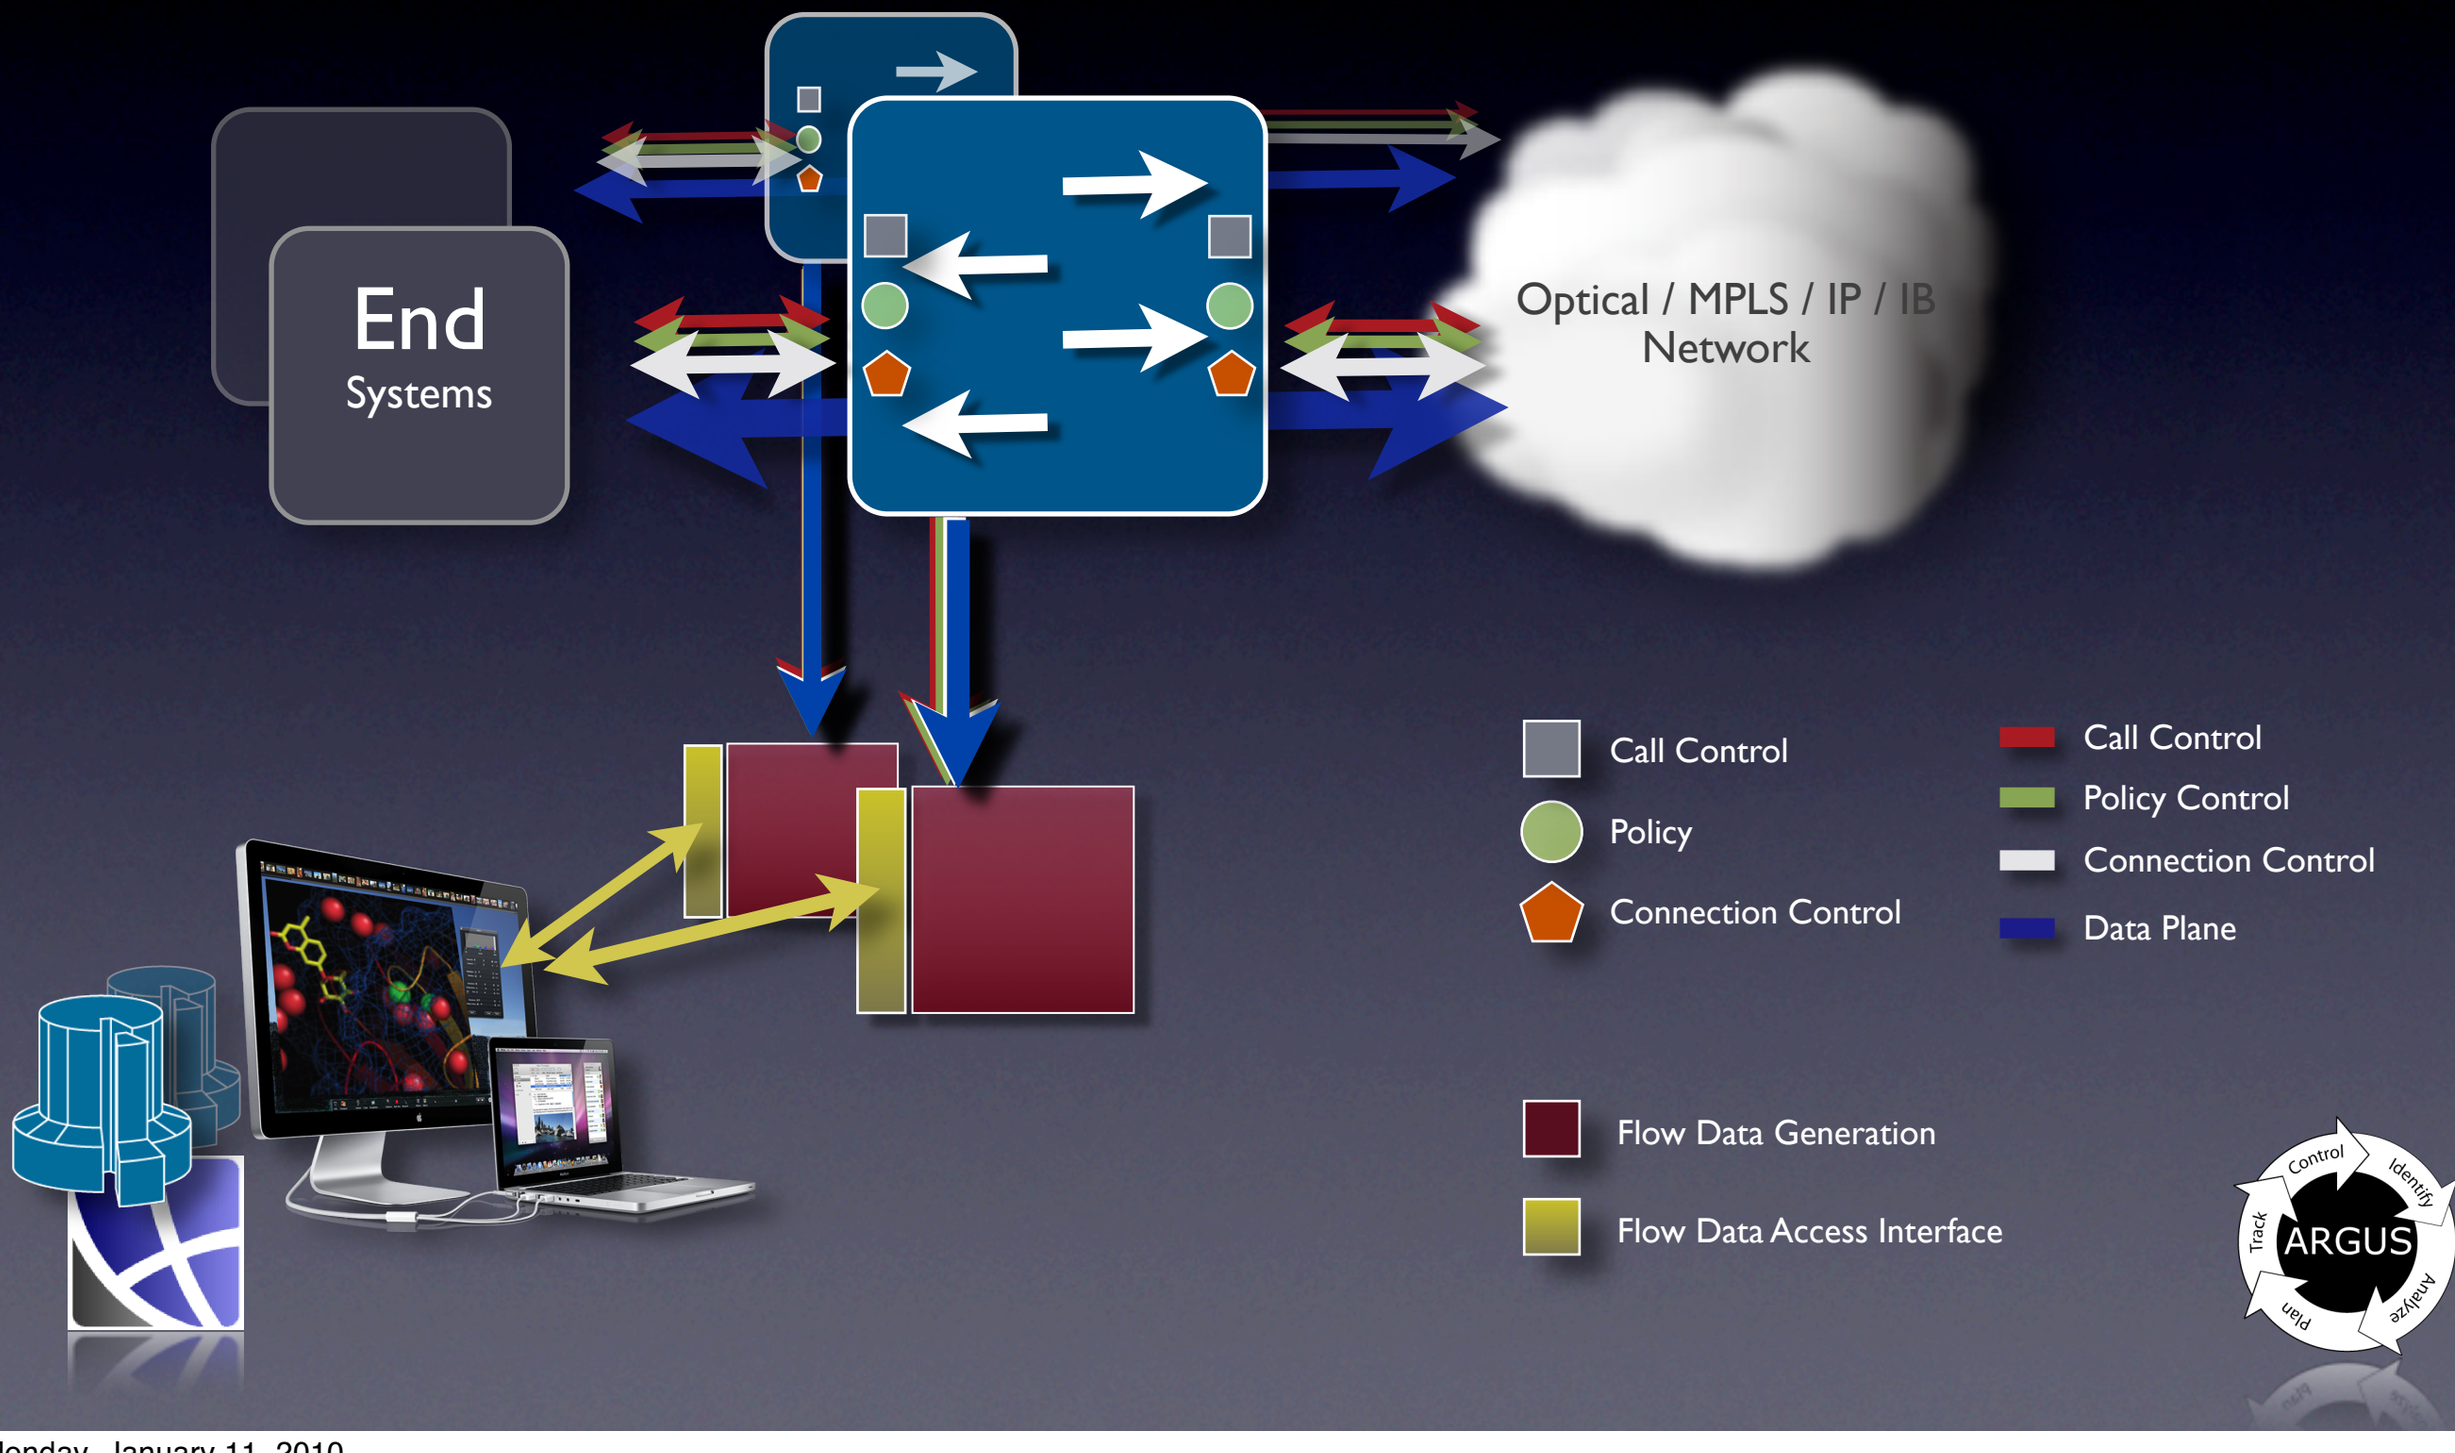
Single Probe



Enterprise Border Monitoring

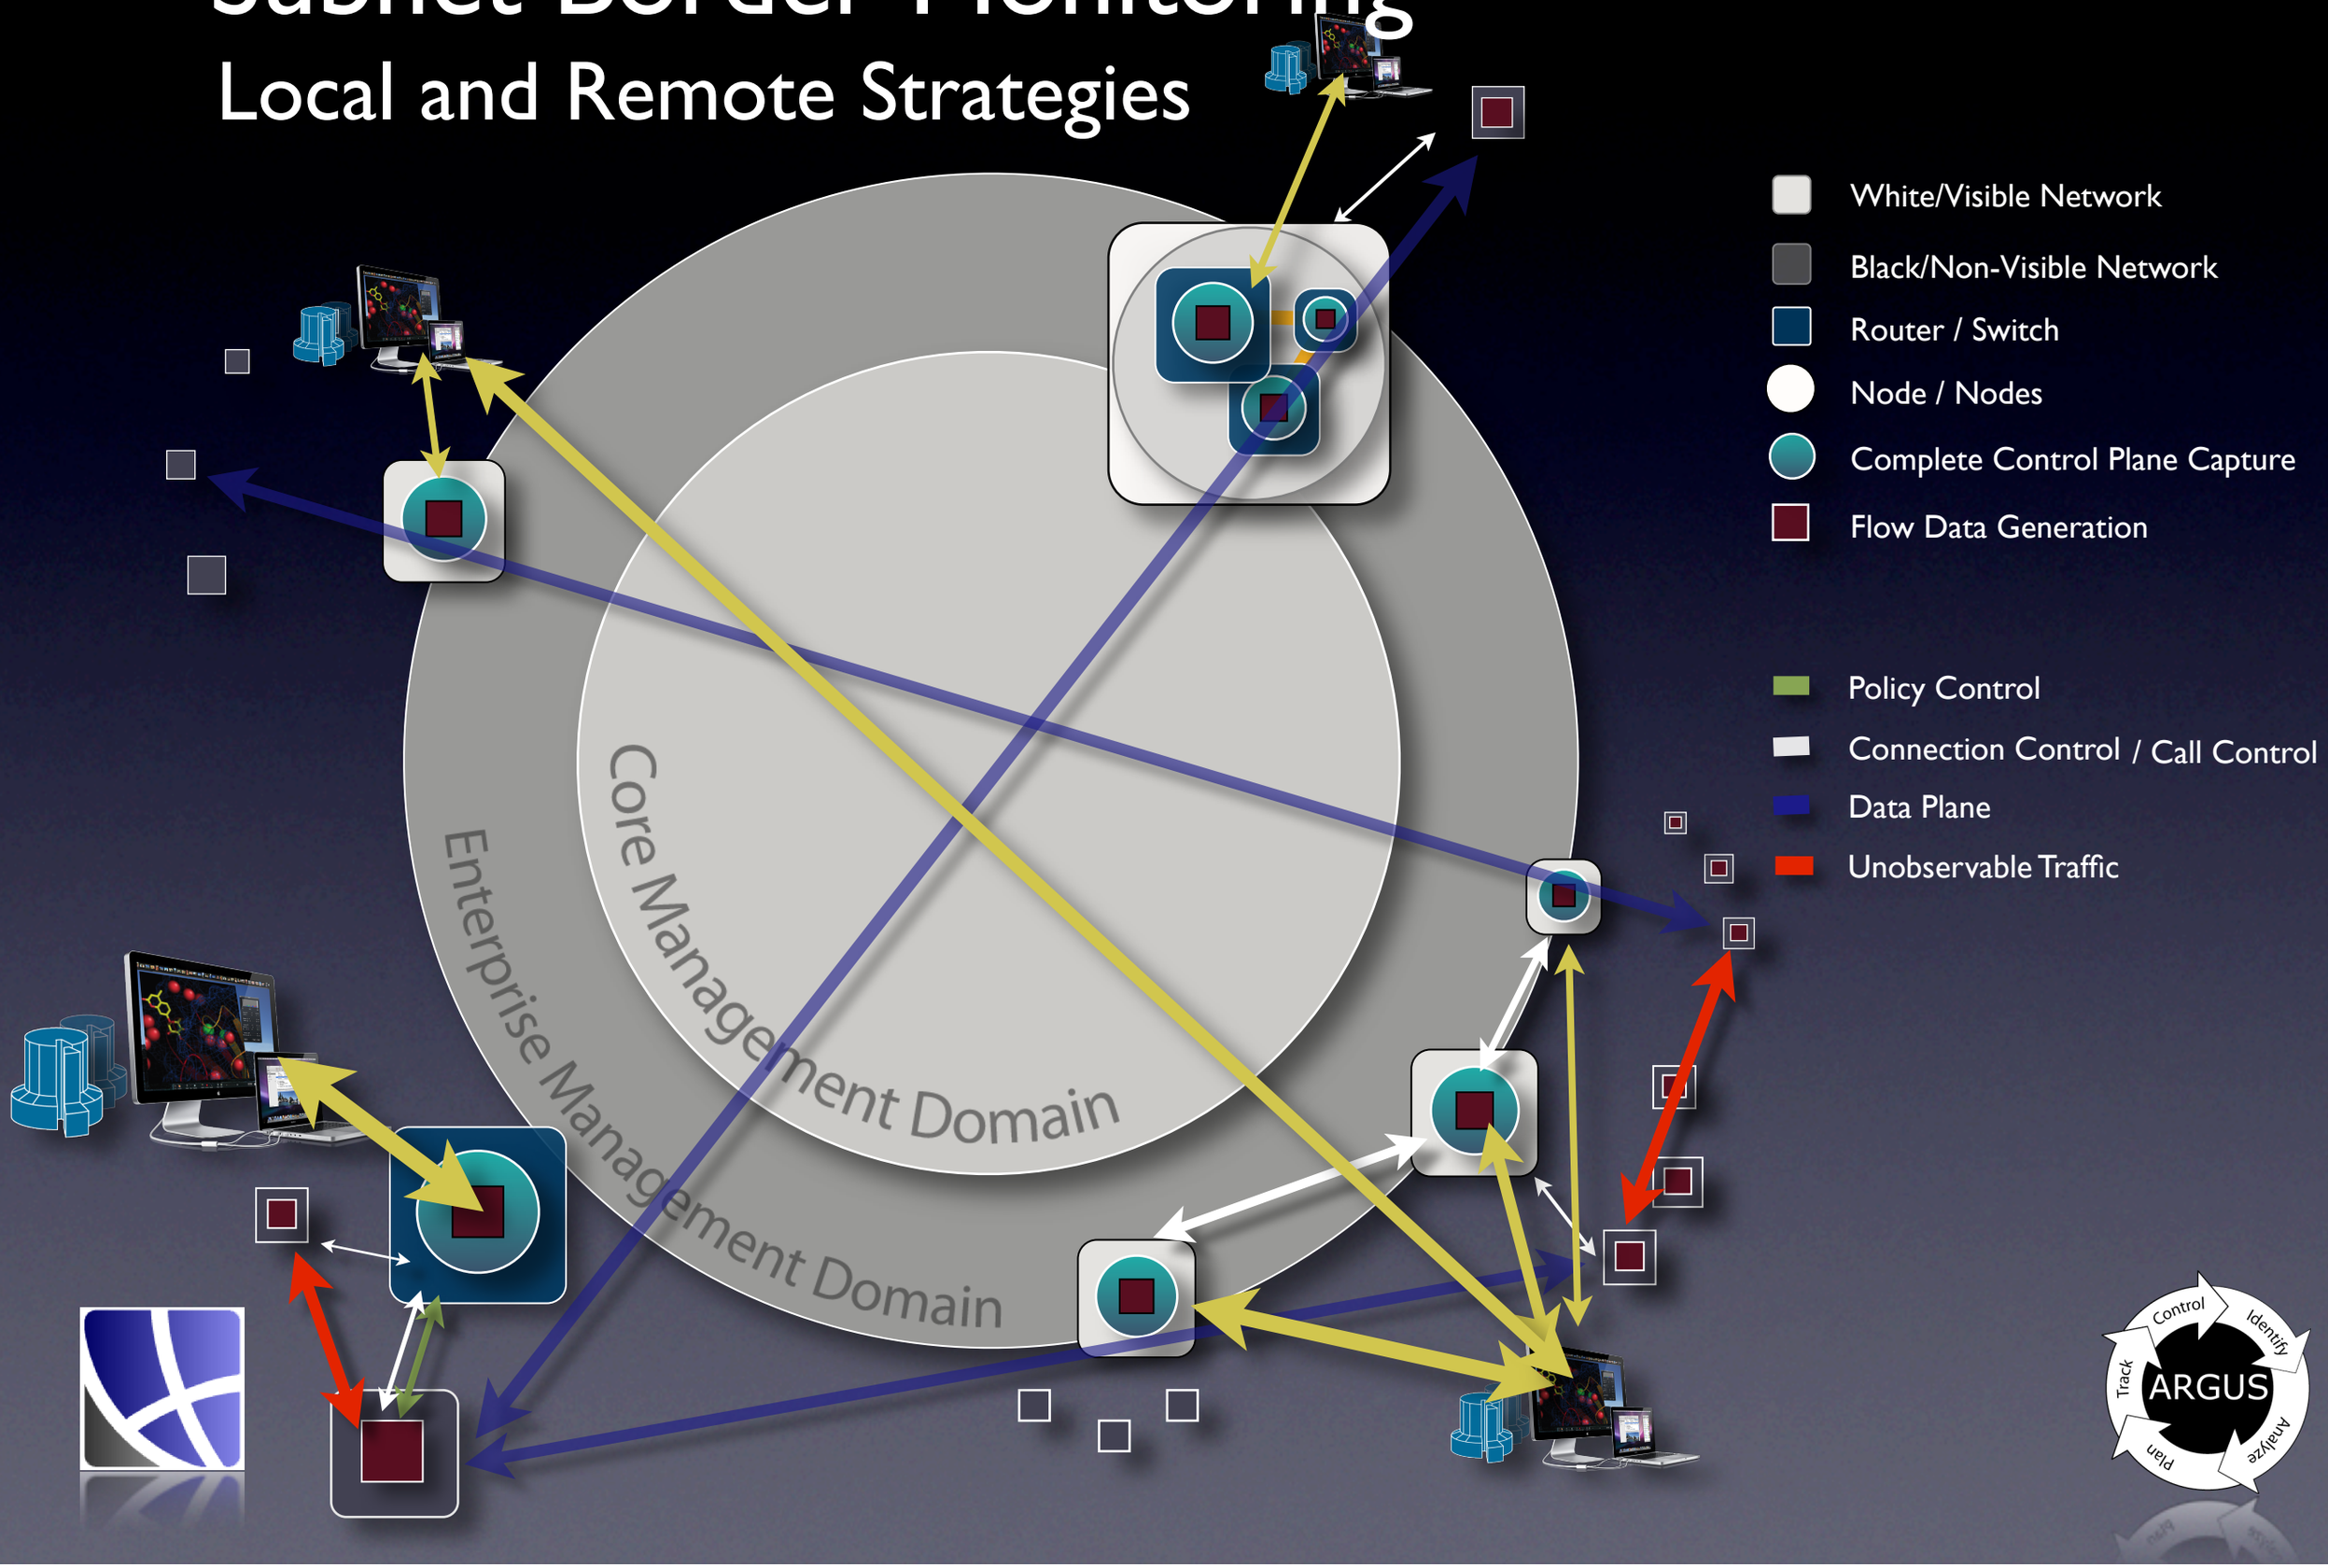
Asymmetric Routing Strategies

Multiple Probes

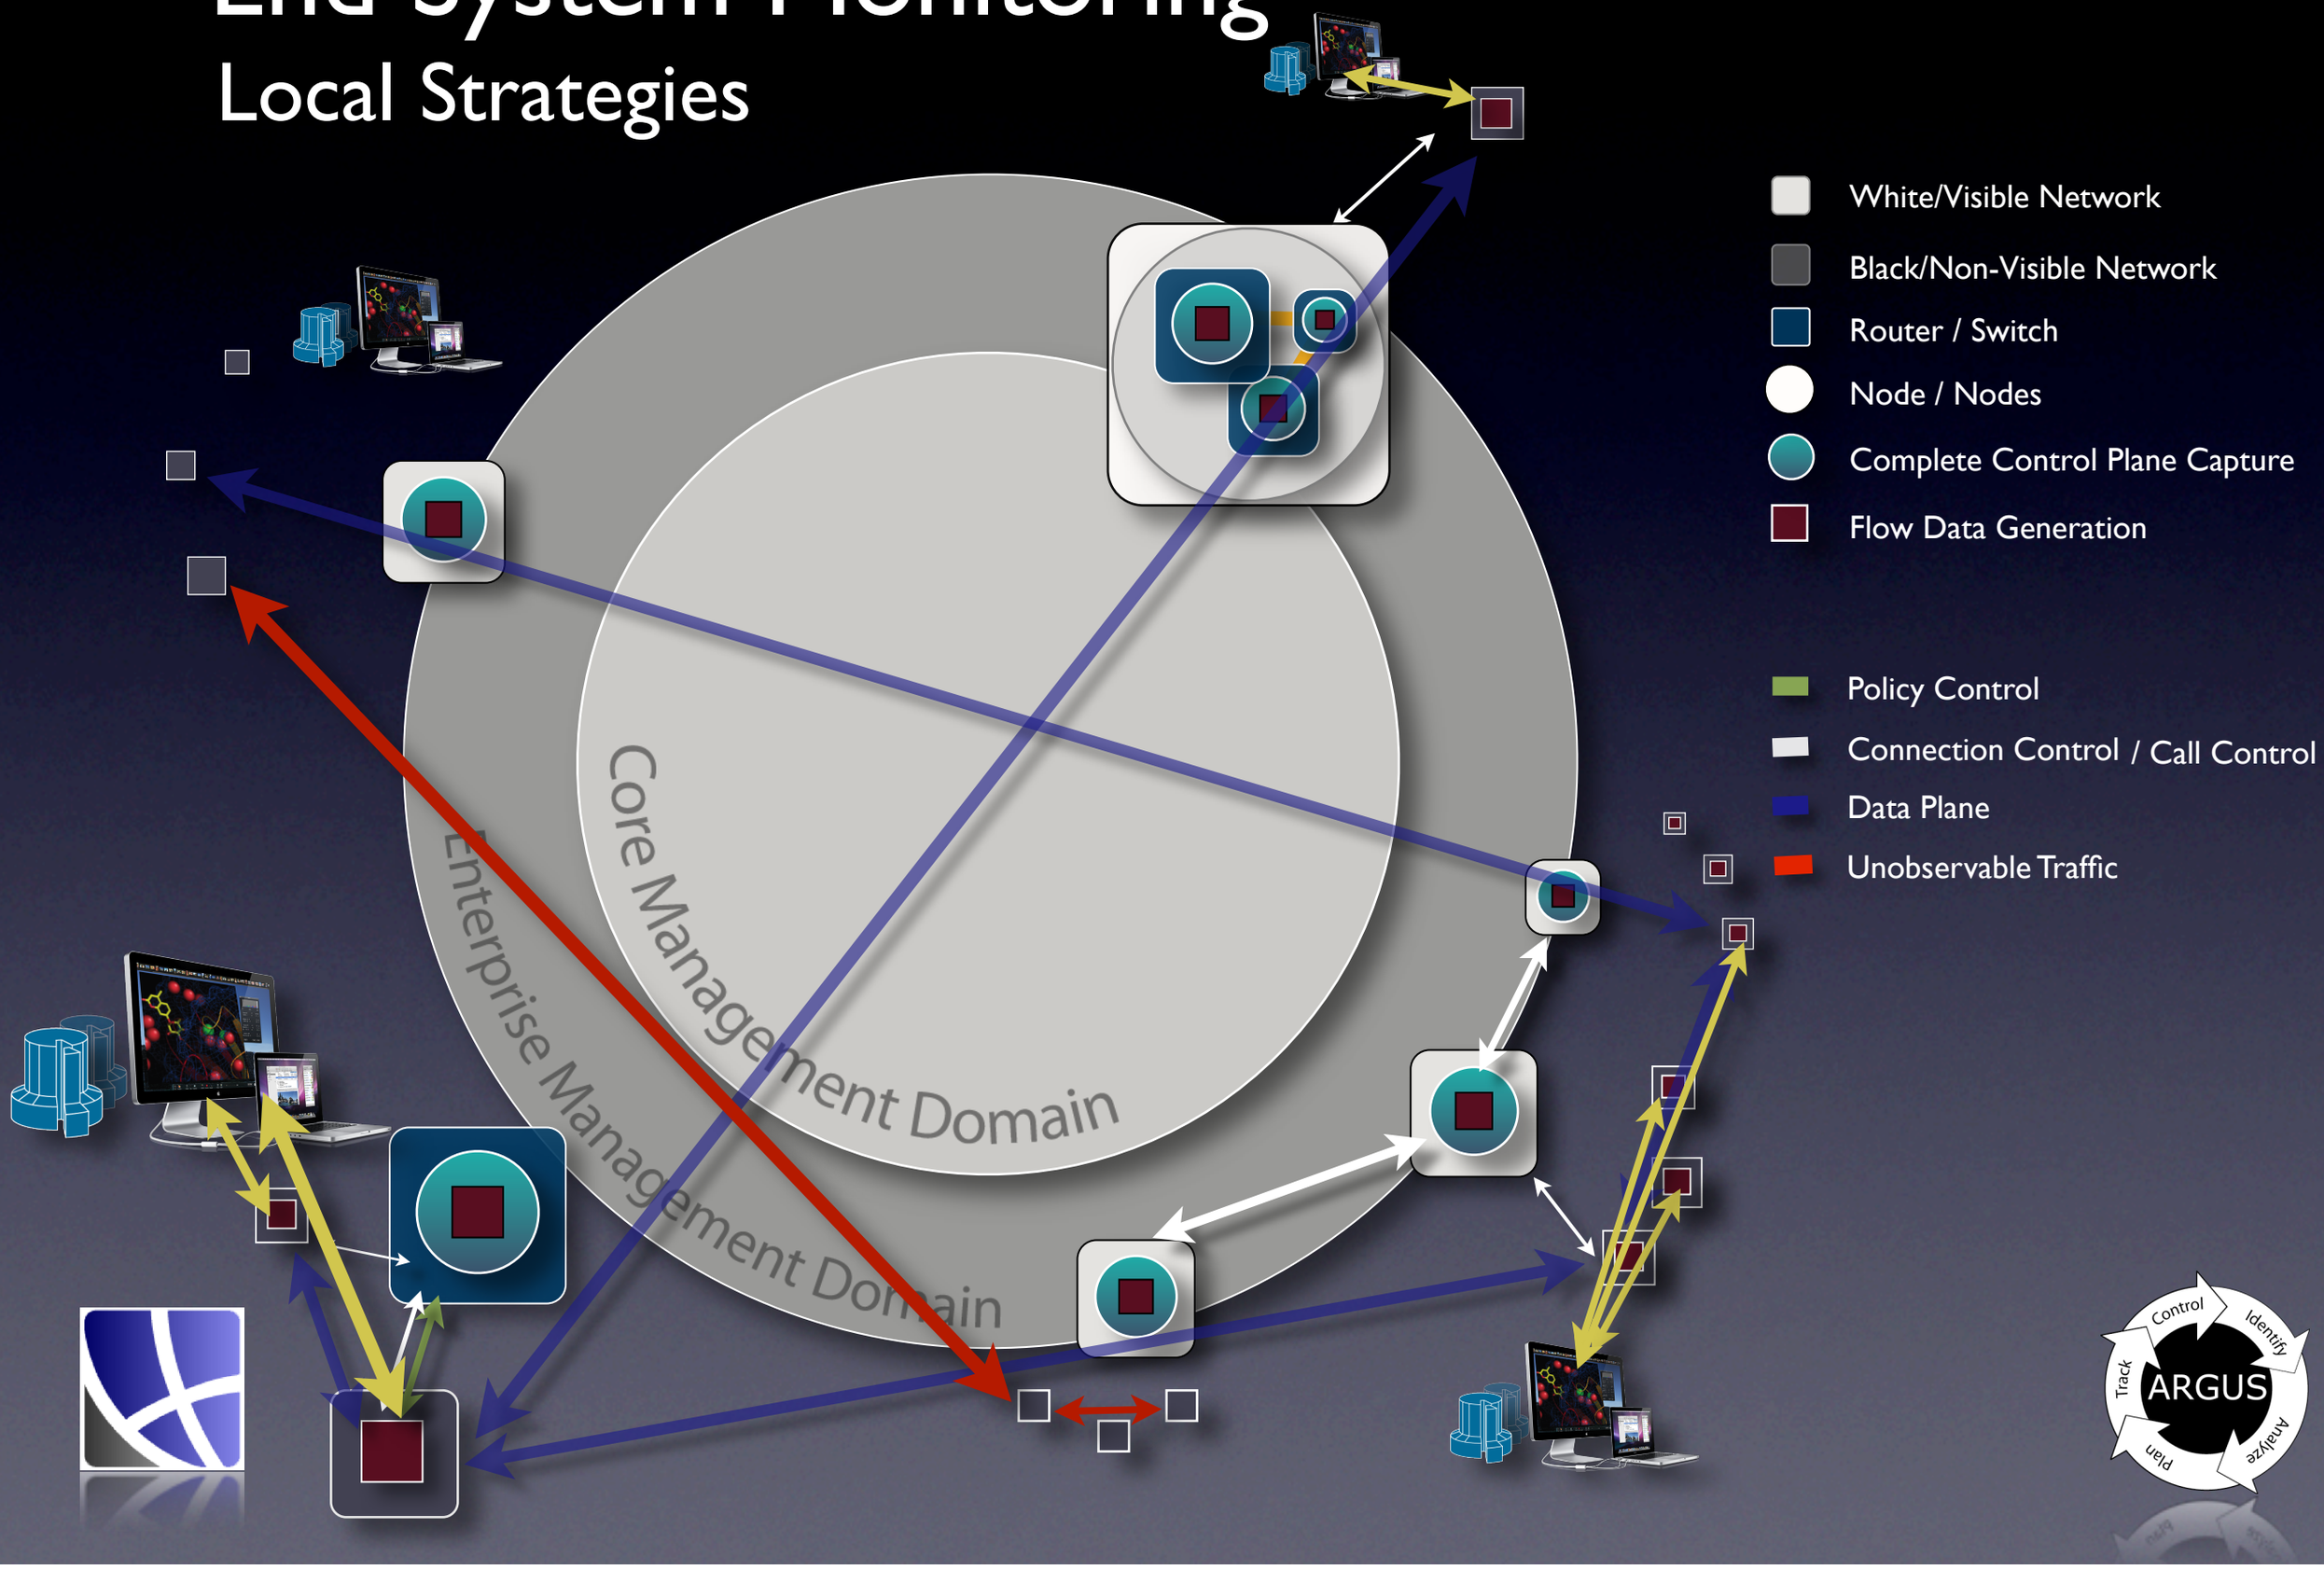


Subnet Border Monitoring

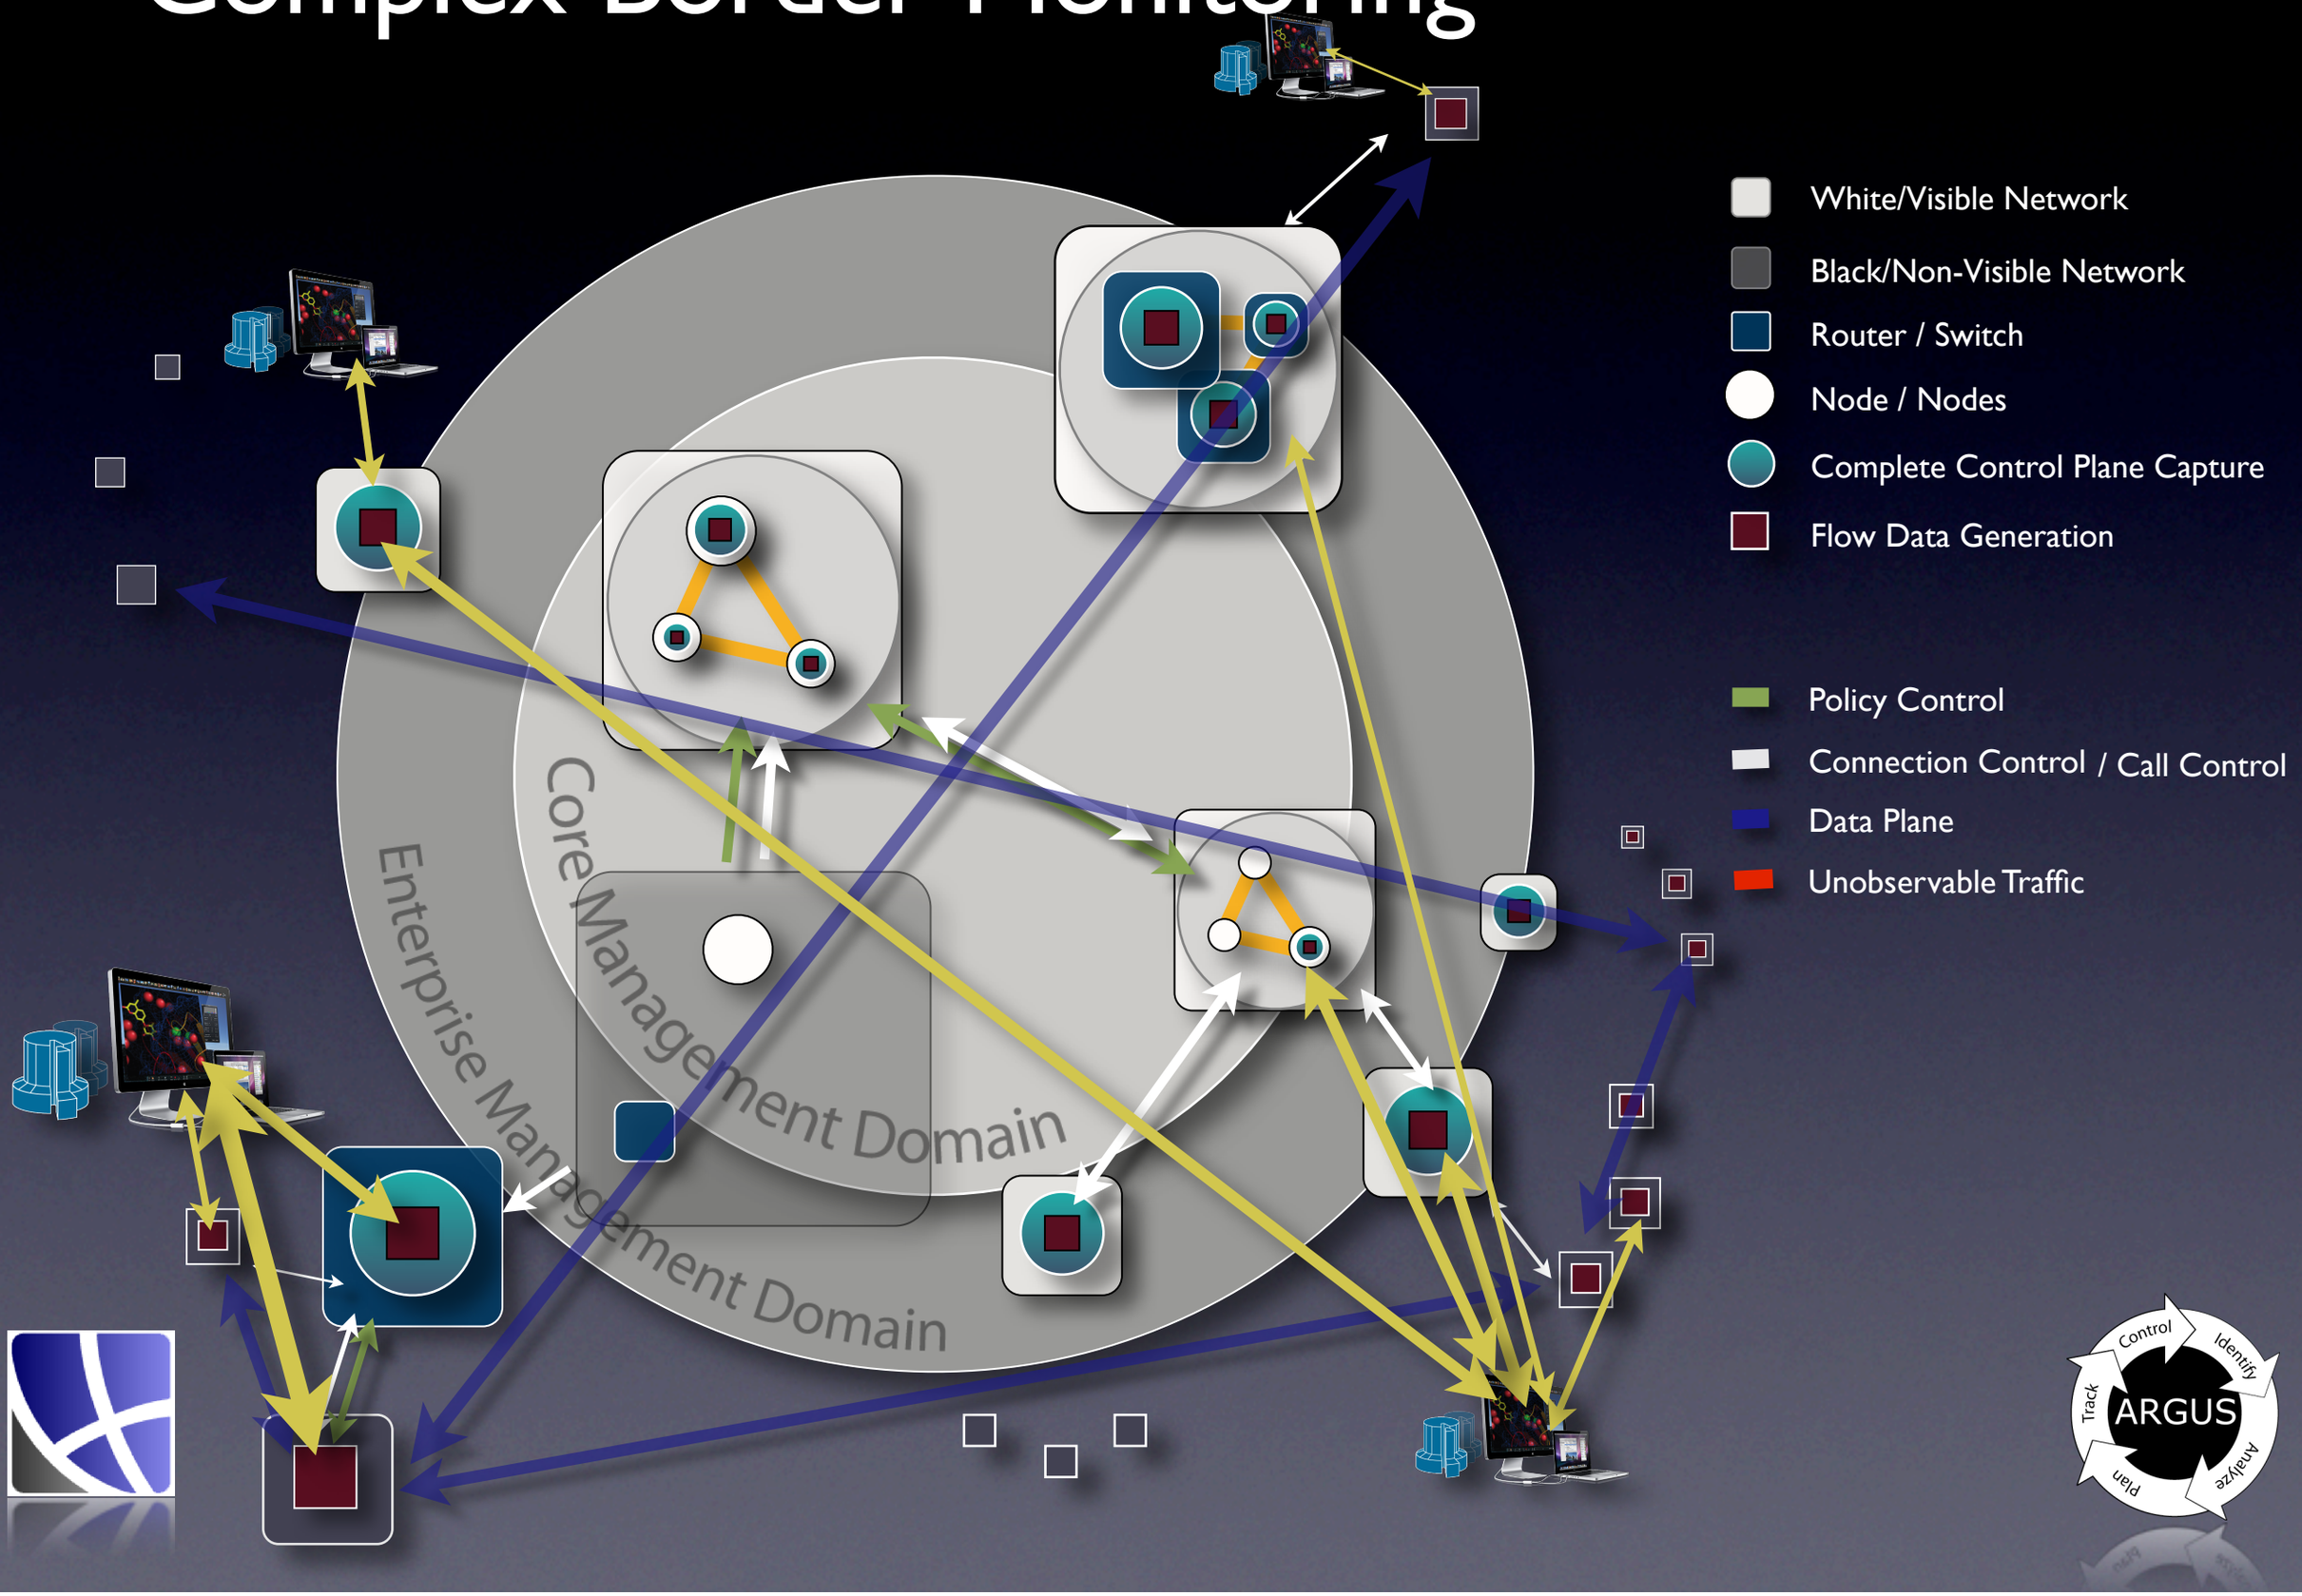
Local and Remote Strategies



End System Monitoring
Local Strategies



Complex Border Monitoring



Complex Monitoring

- Critical elements
  - Time synchronization
  - Comparable flow key models
  - Observation Domain ID Allocations
    - argus sourceID
- Best Common Practices
  - Time synchronization
  - Comparable flow key models
  - Observation Domain ID Allocations
    - argus sourceID





Running Argus

- Packet File Processing
- Multiple Interface Strategies
- Security
- Local Storage





Packet File Processing

- Supported formats
  - Historically supported all of them
  - Now, there are many packet capture conversion tools
  - Currently supports libpcap and ERF
- argus -r packet.file -w packet.file.argus
  - Default options are a good start
  - There are hundreds of options
  - Should utilize argus.conf when possible
- Best Common Practices
  - Don’t throw away the original packet capture file
  - Keep your data together; packets and flow data
  - Structure your data around time





Packet Streams

- Boils down to Access to Packets!!
- Host/End System Deployment
  - Integrated NIC Interfaces
  - Special purpose NIC/FPGA capture
- Network Deployment
  - Integrated Deployment
    - Linux based routers / OpenWRT
    - Experimental hardware
  - Packet Copy Techniques
    - Multiport Repeaters
    - Span Ports
    - UTP Network Taps (including Regen Taps)
    - Optical Network Taps
    - Packet Demultiplexors





Configuration

- argus.conf
  - Running Environment
  - Monitor Characteristics
  - Flow Data Metrics
  - Security Mechanisms





Running Environment Configuration

- Argus Daemon
- Argus Access Port
- Argus Bind IP Address
- Argus Interface
- Argus Go Promiscuous
- Argus Collector
- Argus Chroot Directory
- Argus Set User ID / Group ID
- Argus Output File
- Argus Set PID & PID Path
- Argus Packet Capture File
- Argus Environment Variables





Monitor Characteristics Configuration

- Argus Monitor ID
- Argus Flow Status Interval
- Argus MAR Status Interval
- Argus Debug Level
- Argus Filter Optimizer





Flow Data Metrics Configuration

- Argus Flow Type
- Argus Generate Response Time
- Argus Generate Packet Size
- Argus Generate Jitter Data
- Argus Generate MAC Data
- Argus Generate Application Byte Metrics
- Argus Generate TCP Performance Metrics
- Argus Generate Bi-Directional Time Stamps
- Argus Capture Data Length





Security Mechanisms Configuration

- Argus support the use of SASL to provide strong authentication and confidentiality protection.
- The policy that argus uses is controlled through the use of a minimum and maximum allowable protection strength. Very SASL specific.
  - RA_MIN_SSF
    - This is the minimum security strength factor for the connection. An SSF of 0 allows for no protection. An SSF of 1 will supply integrity protection without privacy.
  - RA_MAX_SSF
    - The MAX_SSF is normally used to specify the strength of encryption. 56, as an example, specifies 56-bit DES. This value should not be less than the MIN_SSF.





Common Problems

- Can't start argus
  - Permissions (interface/filesystem)
    - run as root
- Can't connect to running Argus
  - Tcp_wrappers getting in the way
    - check syslog()
- Argus closes connection after a while
  - Client doesn't read data fast enough
    - improve resources between argus and client





Client Programs





Disclaimers

- Argus clients do indeed have bugs
  - Argus and its clients are prototypes
  - It really is just me doing this work
- If it doesn't do what its suppose to do
  - More than likely its a bug.
  - Sometimes its because we haven't finished it.
  - Sometimes the documentation is wrong.
- Please report any problems!!!
  - Send email to the list or to me
  - Sometimes, it gets fixed that day.





Argus Client Programs

- Basic Operations
- Aggregation
- Data Splitting
- Graphing/Visualization
- User Data Processing
- Semantic Enhancement
- Anonymization





Basic Operations

- ra - provides data file creation, printing, filtering, stripping
- ratop - curses() based argus data GUI
- rasort - provides in memory sorting
- racount - simple record counting
- rahosts, raports, routers - Perl script examples of simple data processing.





Argus Client Program Configuration

- Extensive command line options
- .rarc configuration
- Runtime Environment
- Data Access
- Printing Support
- Security





Argus Client Program Configuration

- Runtime Environment
  - RA_SET_PID
  - RA_PID_PATH
  - RA_OUTPUT_FILE
  - RA_TIME_RANGE
  - RA_TZ
  - RA_DEBUG_LEVEL
  - RA_TIMEOUT_INTERVAL
  - RA_UPDATE_INTERVAL
  - RA_DELEGATED_IP
  - RA_RELIABLE_CONNECT





Argus Client Program Configuration

- Data Access
  - RA_ARGUS_SERVER
  - RA_SOURCE_PORT
  - RA_CISCONETFLOW_PORT
  - RA_TIME_RANGE
  - RA_RUN_TIME
  - RA_FILTER





Argus Client Program Configuration

- Printing Support
  - RA_PRINT_MAN_RECORDS
  - RA_PRINT_LABELS
  - RA_FIELD_SPECIFIER
  - RA_FIELD_DELIMITER
  - RA_FIELD_WIDTH
  - RA_PRINT_NAMES
  - RA_PRINT_RESPONSE_DATA
  - RA_PRINT_UNIX_TIME
  - RA_TIME_FORMAT
  - RA_USEC_PRECISION
  - RA_USERDATA_ENCODE





Argus Client Program Configuration

- Security
  - RA_USER_AUTH
    - This is the user name and, depending on the MECH, a group name, for the SASL account to be used for authentication.
  - RA_AUTH_PASS
    - This is the password for the SASL account to be used for authentication. This plain-text entry does pose some issues, so be sure and protect your .rarc file if this method is used.
  - RA_MIN_SSF
    - This is the minimum security strength factor for the connection. An SSF of 0 allows for no protection. An SSF of 1 will supply integrity protection without privacy.
  - RA_MAX_SSF
    - The MAX_SSF is normally used to specify the strength of encryption. 56, as an example, specifies 56-bit DES. This value should not be less than the MIN_SSF.





Aggregation

Network data aggregation is a MASSIVE topic. It drives most of the data analysis and report generation and is the heart of all the interesting programs.

- racluster, rabins, ratop
  - Each field has its own aggregation methods
  - Semantic preservation





ra()

ra is the basis of all ra* programs, in that it is the simplest of the client programs written against the client library. It is simply, "read the data source and print each record, one record at a time".

- All ra* programs do what ra() does
- Use ra to inspect individual records
- ra.1 is your primary documenation
- If you want to develop argus clients
  - ra.c should be your first example





ratop()

ratop is the top() equivalent for argus data. It is becoming the argus data editor as time goes on.

ratop is also an argus data aggregator.

- ratop is the 1° argus client
- it incorporates all basic functions into a single program.
- it is the best program example for near-realtime situational awareness





Aggregation

- Process of merging records
  - Isoflow Aggregation
    - Time extension / data reduction schemes
      - Consolidate status reports to single record
  - Anisoflow aggregation
    - Reduced flow key
      - Select/Join style data consolidation
    - Expanded flow key
      - Addition of flow attributes to key (i.e. DSBytes)
      - Can be used to generate new metrics
- Each field has its own aggregation rules
  - Data accumulation, reduction, replacement, f(x)
  - Must be statistically or relationally sound





Aggregation

- racluster
  - Default options create isoflow aggregation
  - simple command-line options
    - “-m field [field ...]” option defines anisoflow
      - aggregation scope is duration of data stream
      - storage problems as each flow is cached
- racluster.conf
  - compound aggregation rule structures
  - control of cache holding times and idle times
  - option for semantic preservation





Aggregation

- Configuration

```
RACLUSTER_MODEL_NAME=Test Configuration
RACLUSTER_PRESERVE_FIELDS=yes
RACLUSTER_REPORT_AGGREGATION=yes
RACLUSTER_AUTO_CORRECTION=yes

filter="tcp or udp"  model="saddr daddr proto dport"       status=120 idle=3600 cont
label ="Class-Video" model="srcid saddr daddr proto dport" status=5   idle=10
filter="tcp or udp"  model="saddr daddr proto dport"       status=30  idle=120
filter="icmp"        model="saddr daddr proto dport sport" status=60  idle=30
filter="arp"         model="saddr daddr proto dport sport" status=120 idle=60
filter=""            model="saddr daddr proto"             status=300 idle=3600
```





Data Splitting

Argus can generate a lot of data. Tools that help in data disposition are very, very helpful. Here we are providing basic file processing tools, like the unix command split().

- rasplit
  - First designed to support sorting
  - Split based on time, size, count, and contents
  - Time based splitting basis for 1° archive methods
- rastream
  - rasplit + file processing on close





Graphing/Visualization

Network data graphing is a powerful communication tool for report generation, etc..., but it is also the best way to verify and validate the correctness of data processing.

- ragraph
- Perl script processing rabins() output
  - Generating rrd and running rrd_graph
    - Single object / multi metric graphing
    - Strict time-series representations
    - 1 second minimum resolution
  - Variations use GNU Plot, Mathmatica, MatLab
- CSV file generation
  - Data support for many 3rd party systems
    - Numbers, Excel, AfterGlow, PicViz, etc...
    - Issues with date formats - ./support/Config/excel.rc





User Data Processing

- All ra* programs can grep() user data.
- radump
  - Port of tcpdump decoding logic to argus user data buffers.
    - Use this as a prototype for content sensitive analytics.
- rauserdata/raservices
  - Upper protocol discovery, classification, identification, verification
  - Will provide guess for unknown protocols.
  - Experimental, but works very well.





Semantic Enhancement

- ralabel
  - Geospatial Information Merging
    - Country Codes, City, Area Code, Postal Codes, Physical Address, Lat/Lon
  - Netspatial Information Fusion
    - Origin AS Number
    - Domain Name
    - Path Information
  - Flow Classification
  - Tagging
  - ralabel.conf





Anonymization

Network data anonymity is a big topic when considering sharing for research and collaboration.

The strategies used by argus are intended to preserve the information needed to convey the value of the data, and change or throw everything else away.

- ranonymize
  - User Data Capture
  - Time
  - Network Object Identifiers
    - Network Addresses
    - Service Access Point Identifiers (ports)
  - Sequence Numbers





Data Collection





Data Collection

All ra* programs can read data from any Argus data source, files, stream, encrypted, and/or compressed, and can write current file structure.

Making an argus data repository needs just a little bit more.

- File Distribution
- Radium Distribution
- Argus Repository Establishment
  - cron
  - rasplit/rastream
  - rasqlinsert/rasql

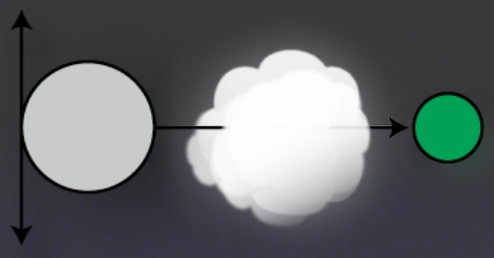




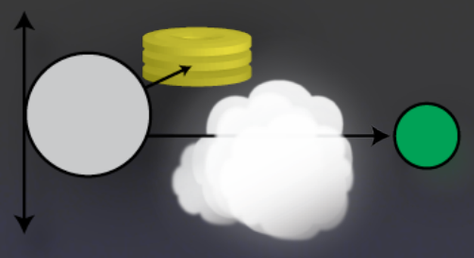
Data Collection



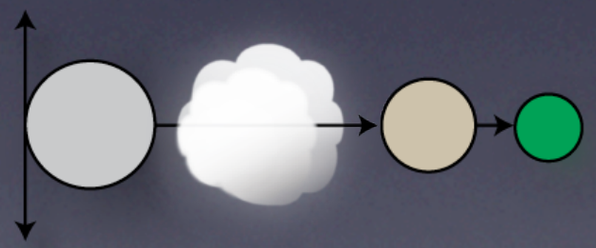
Argus reading from packet files or network and writing directly to disk



Argus reading from the network and writing directly to network based client



Argus reading from the network and writing directly to disk and network based client



Argus reading from the network and writing directly to a network Radium, writing to a client



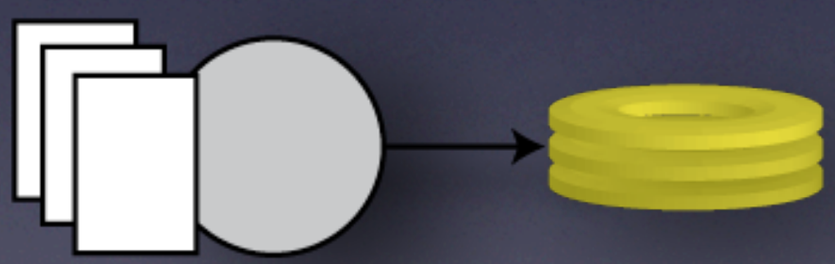
Argus writing to local Radium which is writing directly to disk and to network based clients



Many Argi writing directly to a Radium based distribution network, which is providing data to a set of clients.

Data Collection

- Local Generation and Storage
  - Basis for argus-2.0 argusarchive.sh
  - Direct argus support for renaming files
  - Normally cron mediated
  - Issues with time and record spans
- System designer has most control !!!



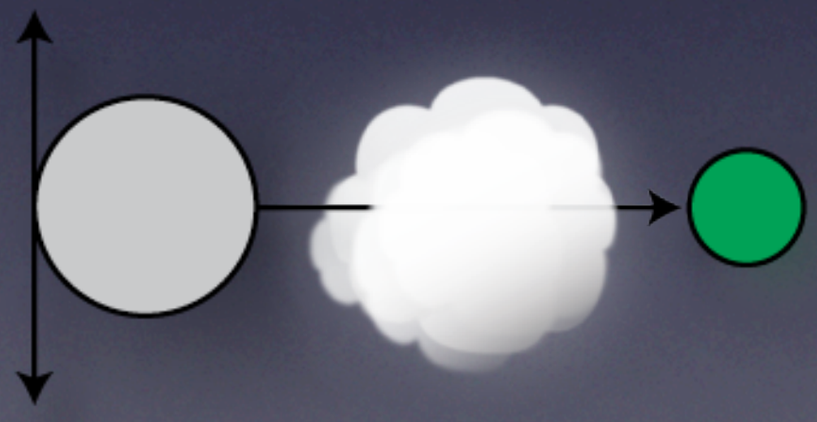


Argus reading from packet files or network and writing directly to disk



Data Collection

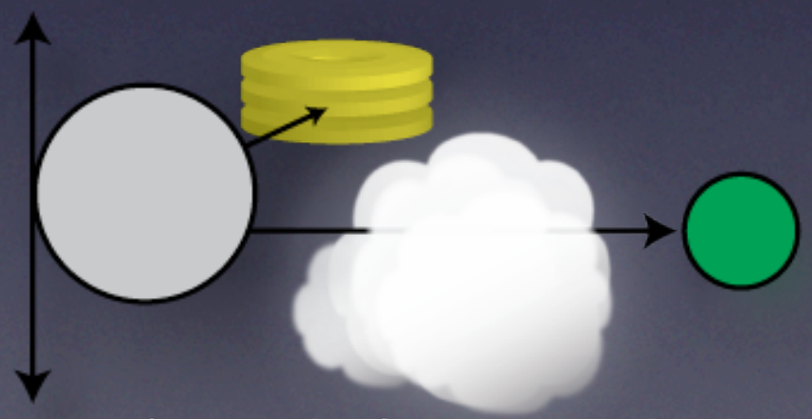
- Local Generation Remote Collection
  - Most high performance systems use this strategy
    - Provides explicit scalability and performance capabilities
    - Relieves argus from physical device blocking
    - Network interfaces generally faster than local storage devices
- Introduces network transport issues
  - Reliability, connection vs. connection-less, unicast vs multicast, congestion avoidance, access control and confidentiality



Argus reading from the network and writing directly to network based client

Data Collection

- Local Storage and Remote Collection
  - Used when data reliability is most critical
    - Local storage provides explicit data recovery
    - File collection provides additional distribution flexibility
    - Scheduled transport
- Reduces ultimate sensore performance
  - Argus itself is doing a lot of work
  - Packet processing is really the ultimate limit



Argus reading from the network and writing directly to disk and network based client

Data Collection

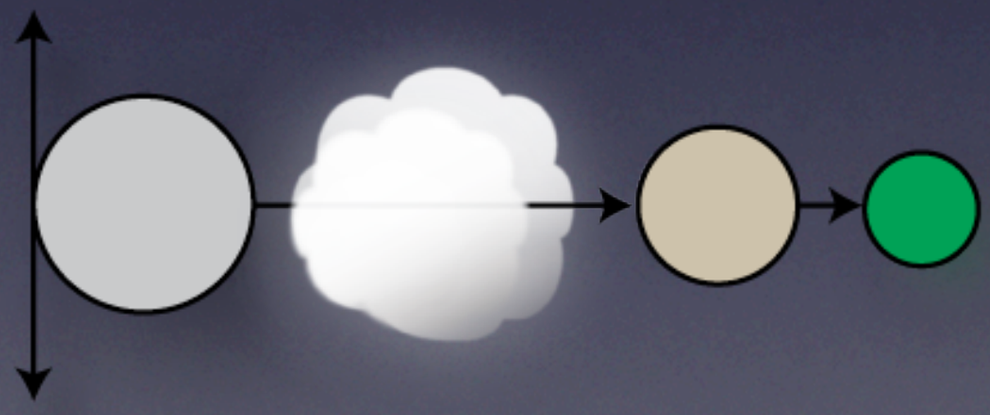
- Radium
  - Primary argus data distribution technology
  - Radium is a ra* program with an argus output processor.
    - Read from many sources
    - Write to many clients
    - Serve up argus data files
    - Process/transform data
    - Configuration is combo of argus() and ra()
- Supports very complex data flow machine architectures.





Data Collection

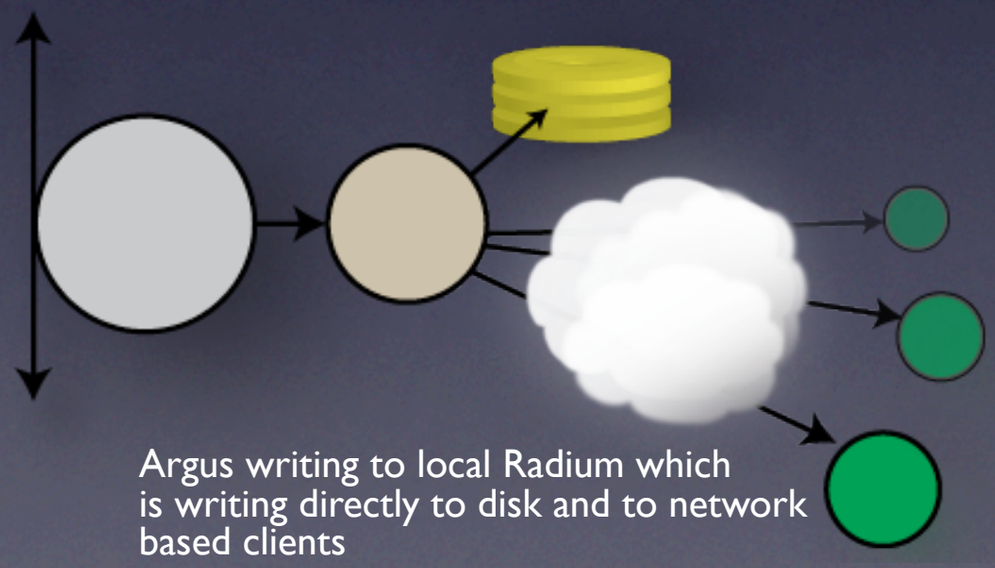
- Local Generation Remote Distribution
  - Most prevalent strategy used in argus-3.0
    - Provides explicit scalability and performance capabilities
    - Provides most stable collection architecture from client perspective
    - Single point of attachment for complete enterprise
  - Least reliable of ‘advanced’ strategies
    - Radium failure interrupts continuous stream collection, with no opportunity for recovery



Argus reading from the network and writing directly to a network Radium, writing to a client

Data Collection

- Local Distribution and Storage
  - Best methodology
    - Provides explicit scalability and performance capabilities
    - Provides most reliable collection architecture
    - Multiple points of attachment, multiple points of control
- Most expensive strategy at data generation
  - Radium deals with device and remote client requests for data which does come with a processor and memory cost



Argus writing to local Radium which is writing directly to disk and to network based clients

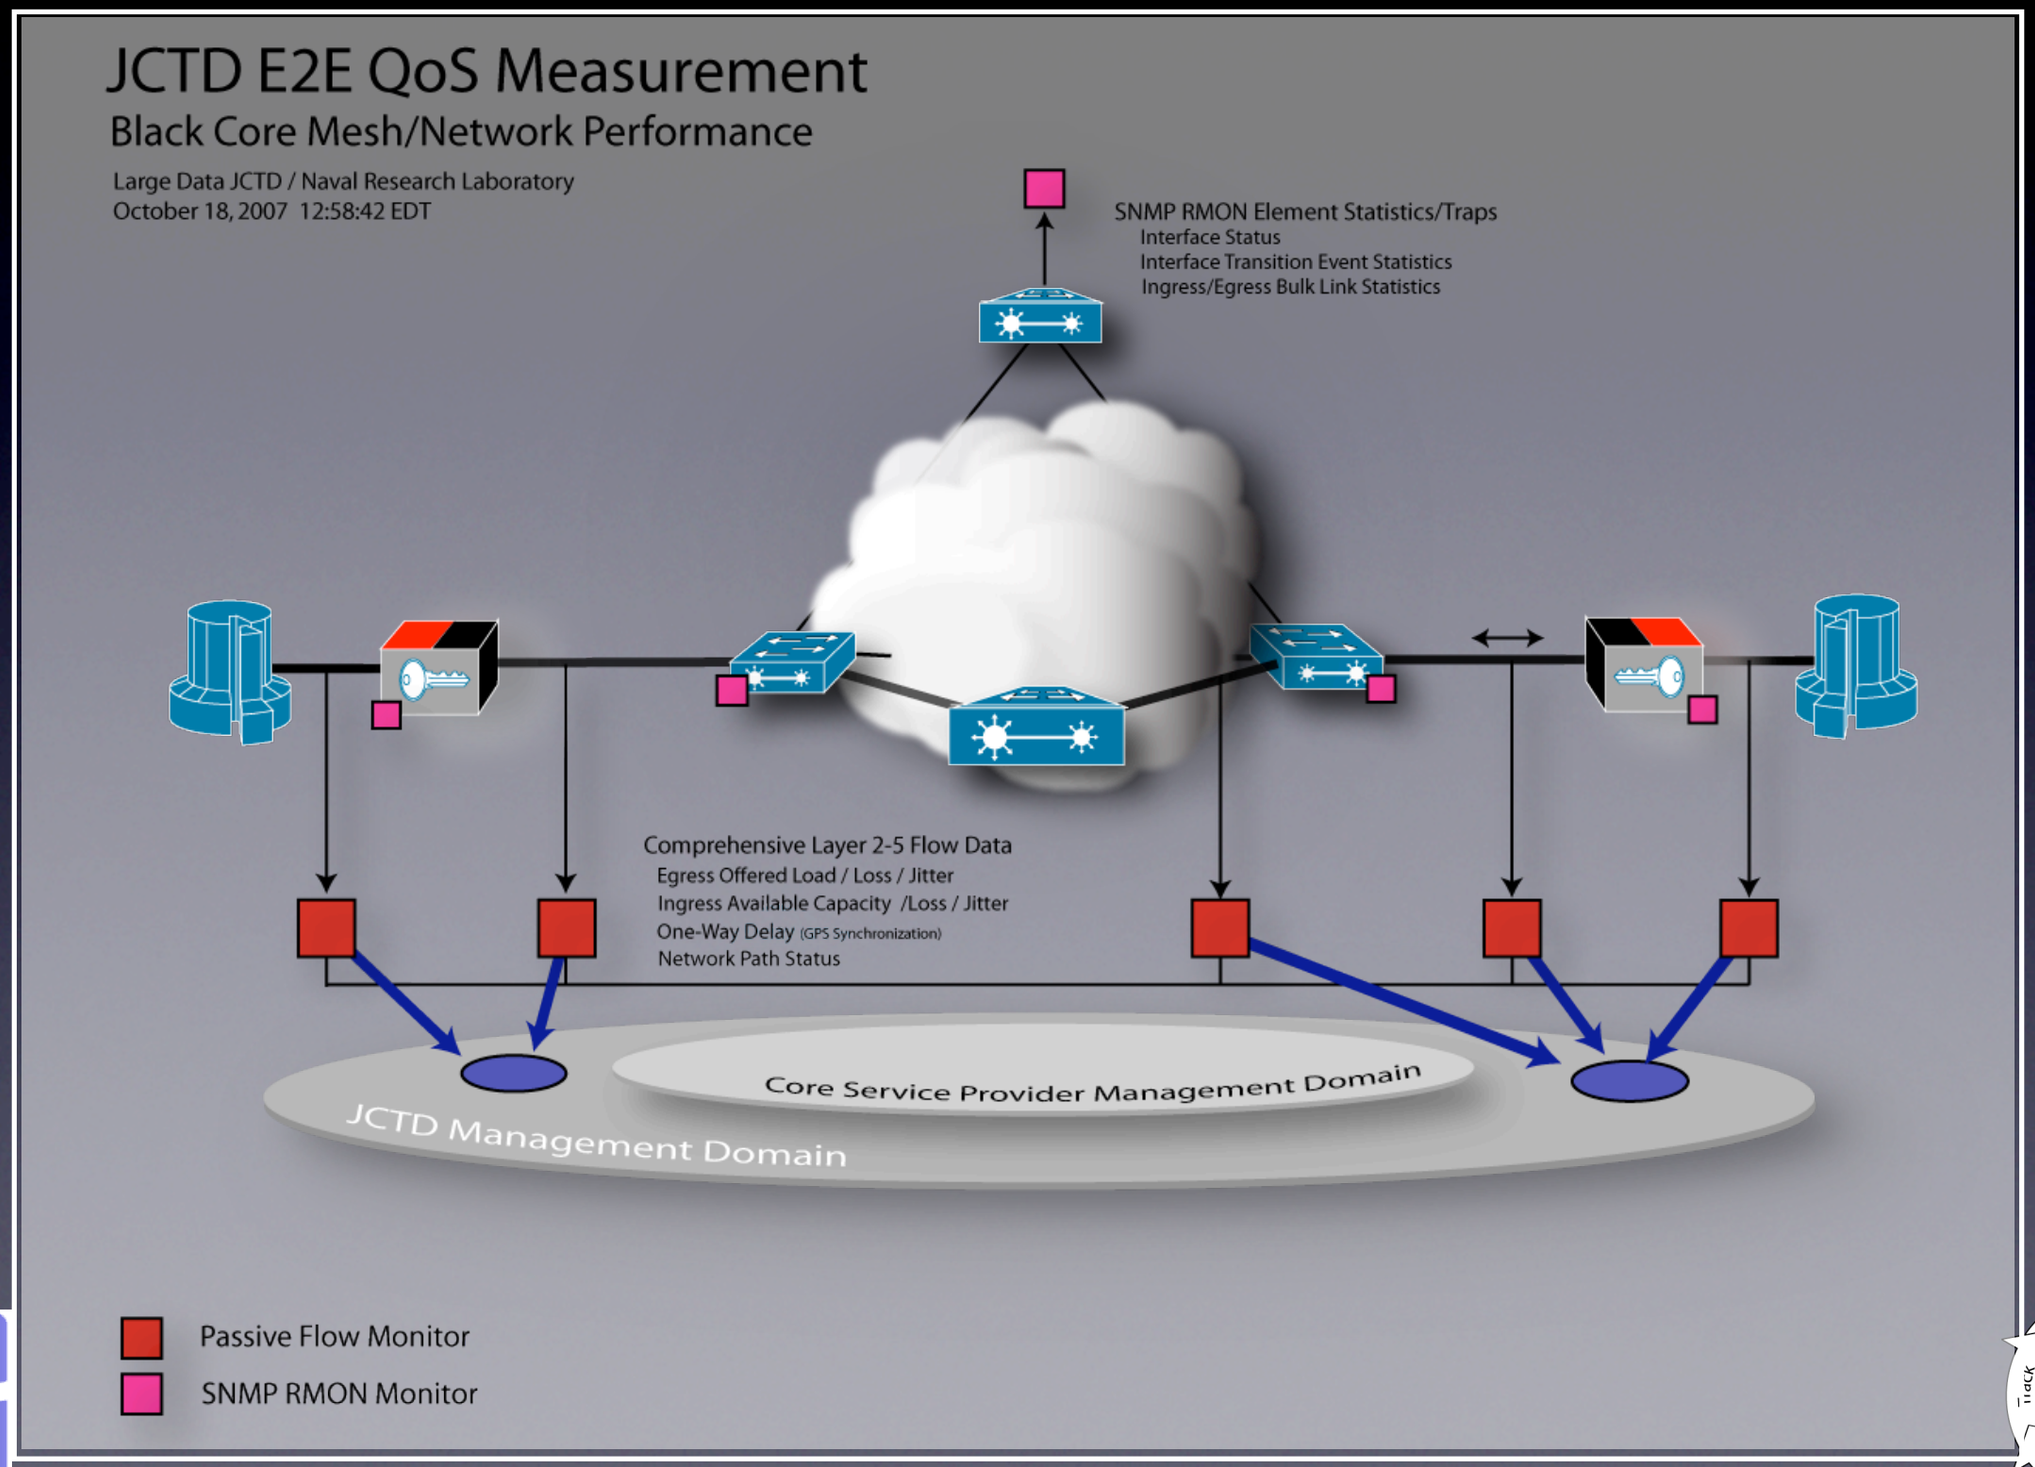
Data Collection

- Complex data flow machine architectures
  - Architecture of choice for scalability
    - Provides explicit scalability and performance capabilities
    - Provides most parallelism
    - Multiple points of attachment, multiple points of control
- Can get a little too complex
  - Merging of multiple flows, multiple times, introduces complex data duplication issues, and allows for complex, incompatible data schemas to co-exist



Many Argi writing directly to a Radium based distribution network, which is providing data to a set of clients.

Multi-Point Monitoring



Argus Repositories

- Argus Repository Establishment
  - Formal Ingest/Disposition
- Repository Function
  - Primitive Data Repository
    - General Archive
    - Access Control
    - Retention Policies
    - Modification Policy (Compression)
  - Derived Data Repositories





Argus Repositories

- Native File System
  - Simplicity
  - Performance
  - Compatibility
- Relational Database System (RDBMS)
  - Extensive Data Handling Capabilities
  - Complex Management Strategies
  - Performance Issues





Data Ingest

- Simple Record Storage
  - File generation
  - MySQL table insertion
- Stream Block Processing
- Primitive Data
  - rasplit
  - rastream
- Dervied Data
  - rabins - controlled context rasplit with data aggregation





Best Common Practices

- File system archives
  - Primitive and derived data file systems
  - RDBMS managed complex indexing
  - rastream
    - /sourceId/year/month/day file structure
    - 5 minute files
    - 288 entries per day
    - Matches native file system performance for searching
    - Analogous to Google's Big Table filesystem
- RDBMS based archives
  - Short term data held in RDBMS
  - Rolled into file based system after N days.
  - Binary data inserted into database
  - Primitive data schema includes 'autoid'
  - Table names provide explicit partitioning





Situational Awareness





Situational Awareness

Level 1 SA - Perception

- The perception of elements in the environment within a volume of time and space
- Involves timely sensing, data generation, distribution, collection, combination, filtering, enhancement, processing, storage, retention and access.

Level 2 SA - Comprehension

- Understanding significance of perceived elements in relation to relevant goals and objectives.
- Involves integration, correlation, knowledge generation.

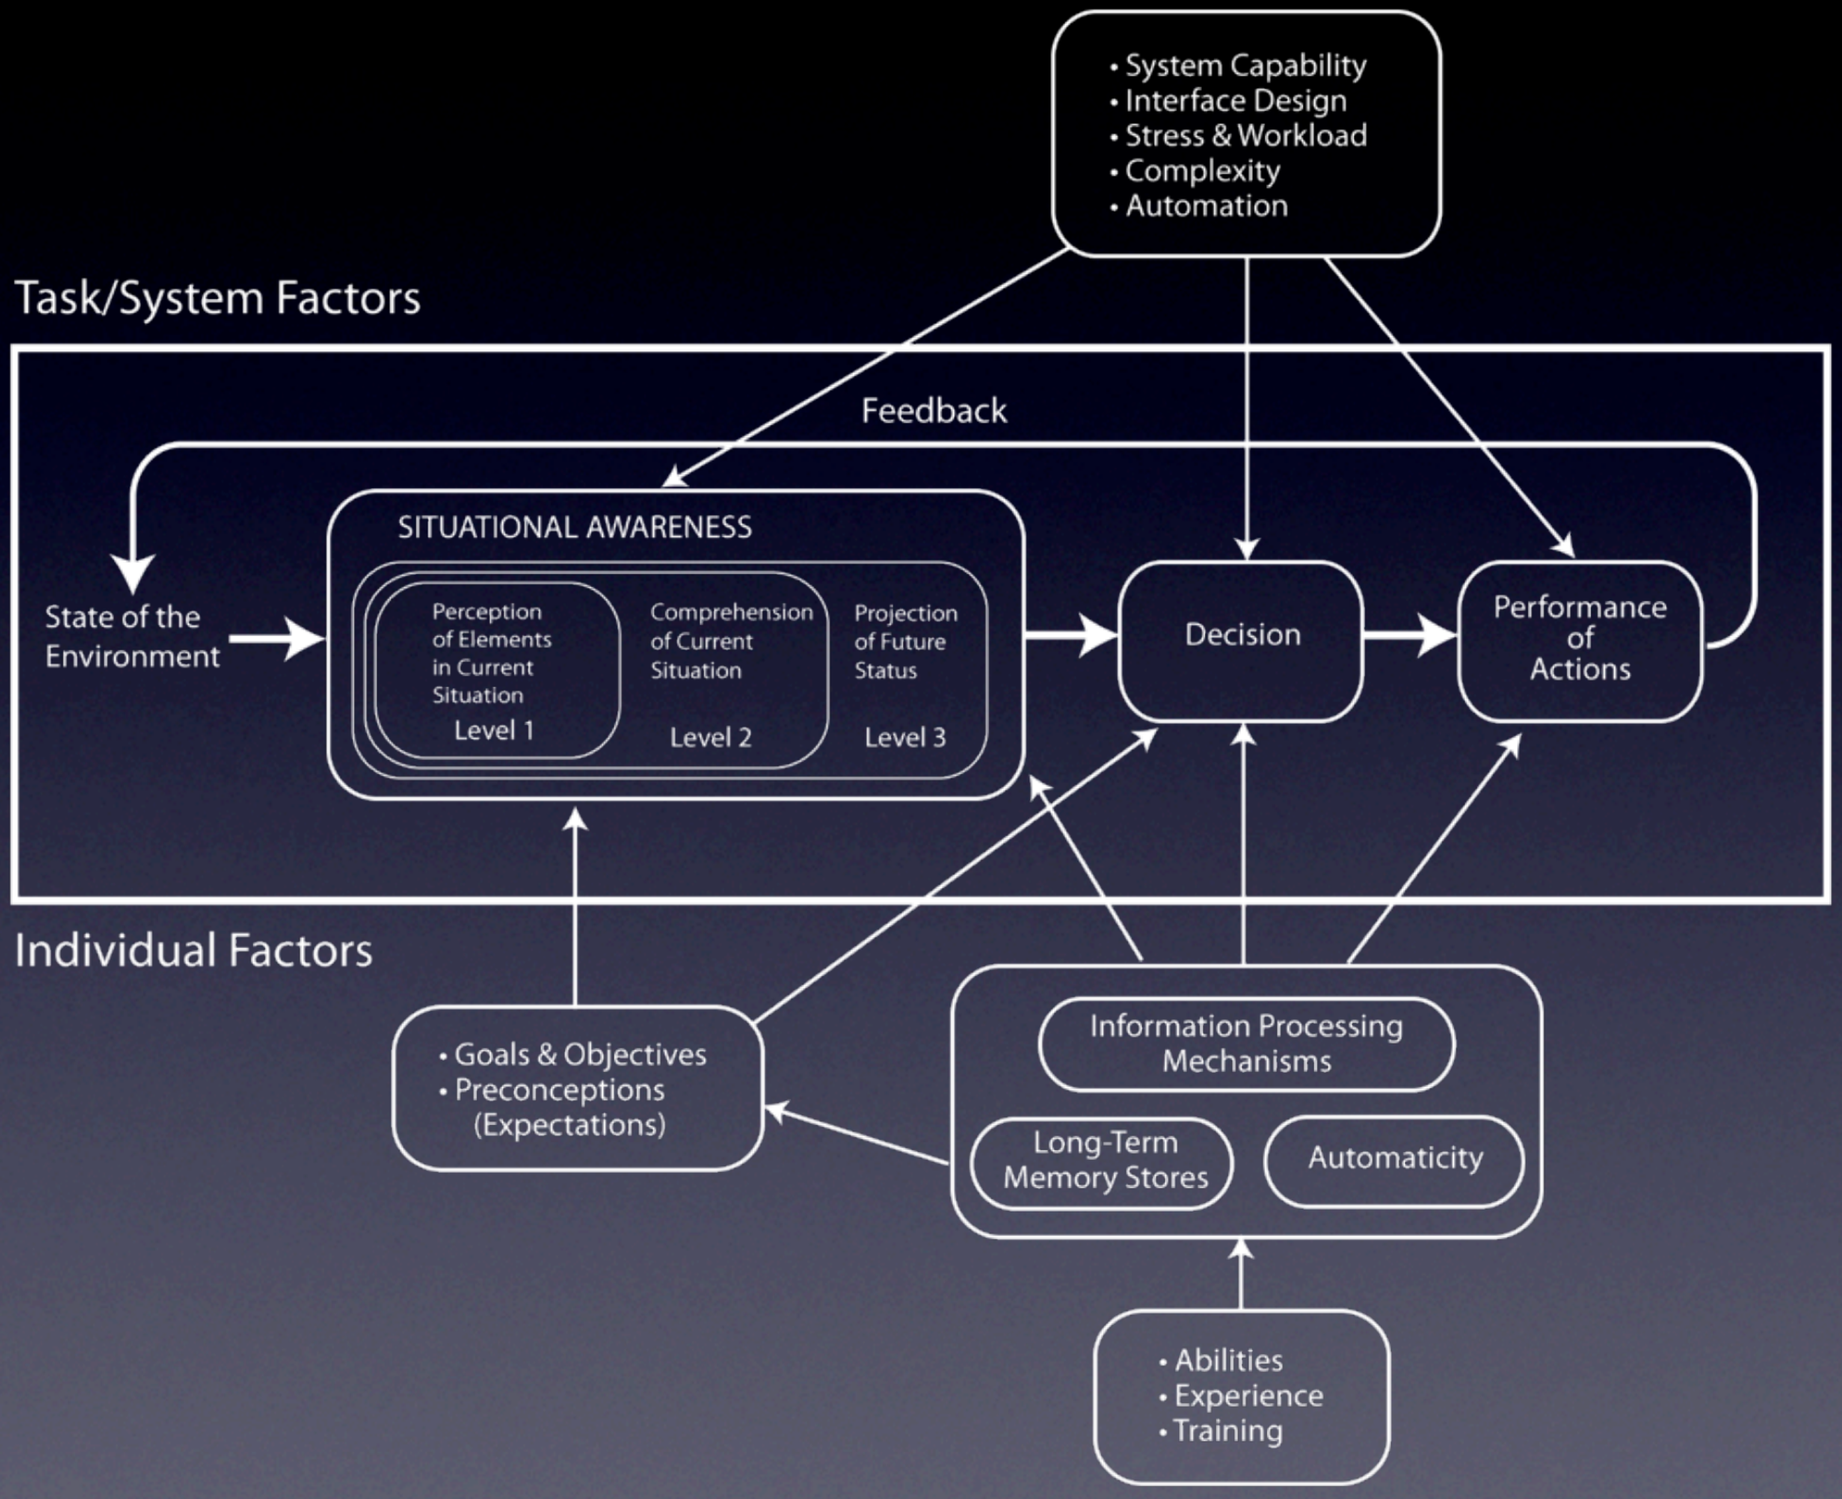
Level 3 SA - Projection of Future Status



Endsley, M. R. (1995b). Toward a theory of situation awareness in dynamic systems. Human Factors 37(1), 32-64.



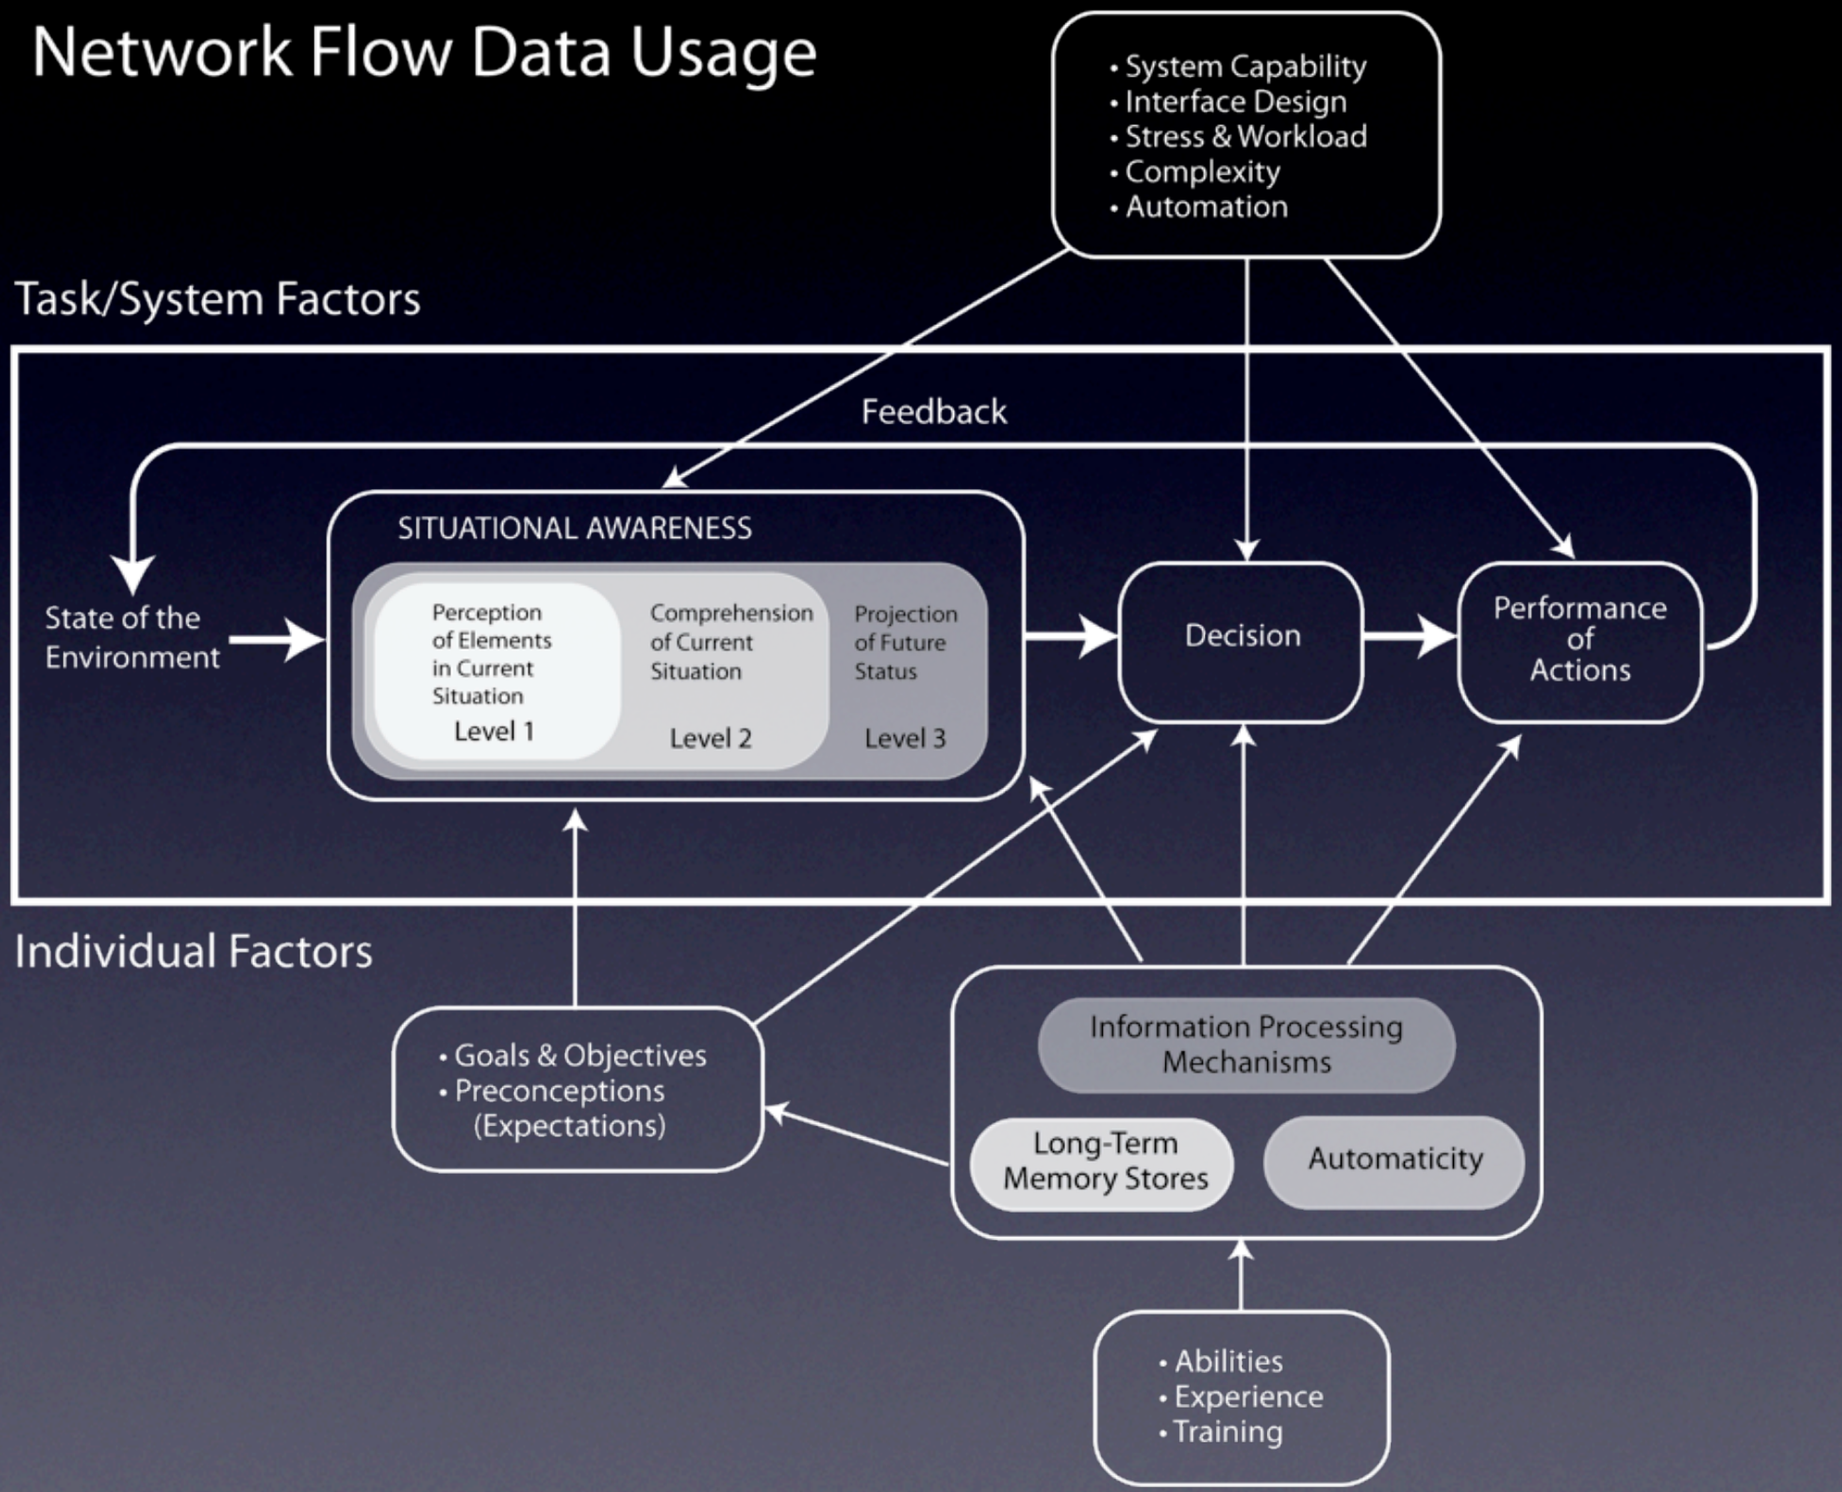
Model of Situational Awareness in Dynamic Decision Making



Endsley, M. R., ***Human Factors***, Vol. 37, No. 1, 1995. Copyright 1995 by the Human Factors and Erogonomics Society. All rights reserved.

Model of Situational Awareness in Dynamic Decision Making

Network Flow Data Usage



Endsley, M. R., ***Human Factors***, Vol. 37, No. 1, 1995. Copyright 1995 by the Human Factors and Erogonomics Society. All rights reserved.

Situational Awareness System

Basic design is local sensing, data collection and management, with local near real time data processing and large scale data sharing to support multi-dimensional network activity comprehension.

- Federated Database Model
  - Access controlled by local administrative domain (scoping)
  - Cloud-like distributed processing and query support
  - Flexible data management strategies
  - Large numbers of simultaneous users
- Near real-time information availability
  - Register for information of interest
  - Complex data processing / aggregation / enhancement / advertisement
  - Large scale data correlation processing
  - Anonymization
- Substantial historical data access





Situational Awareness

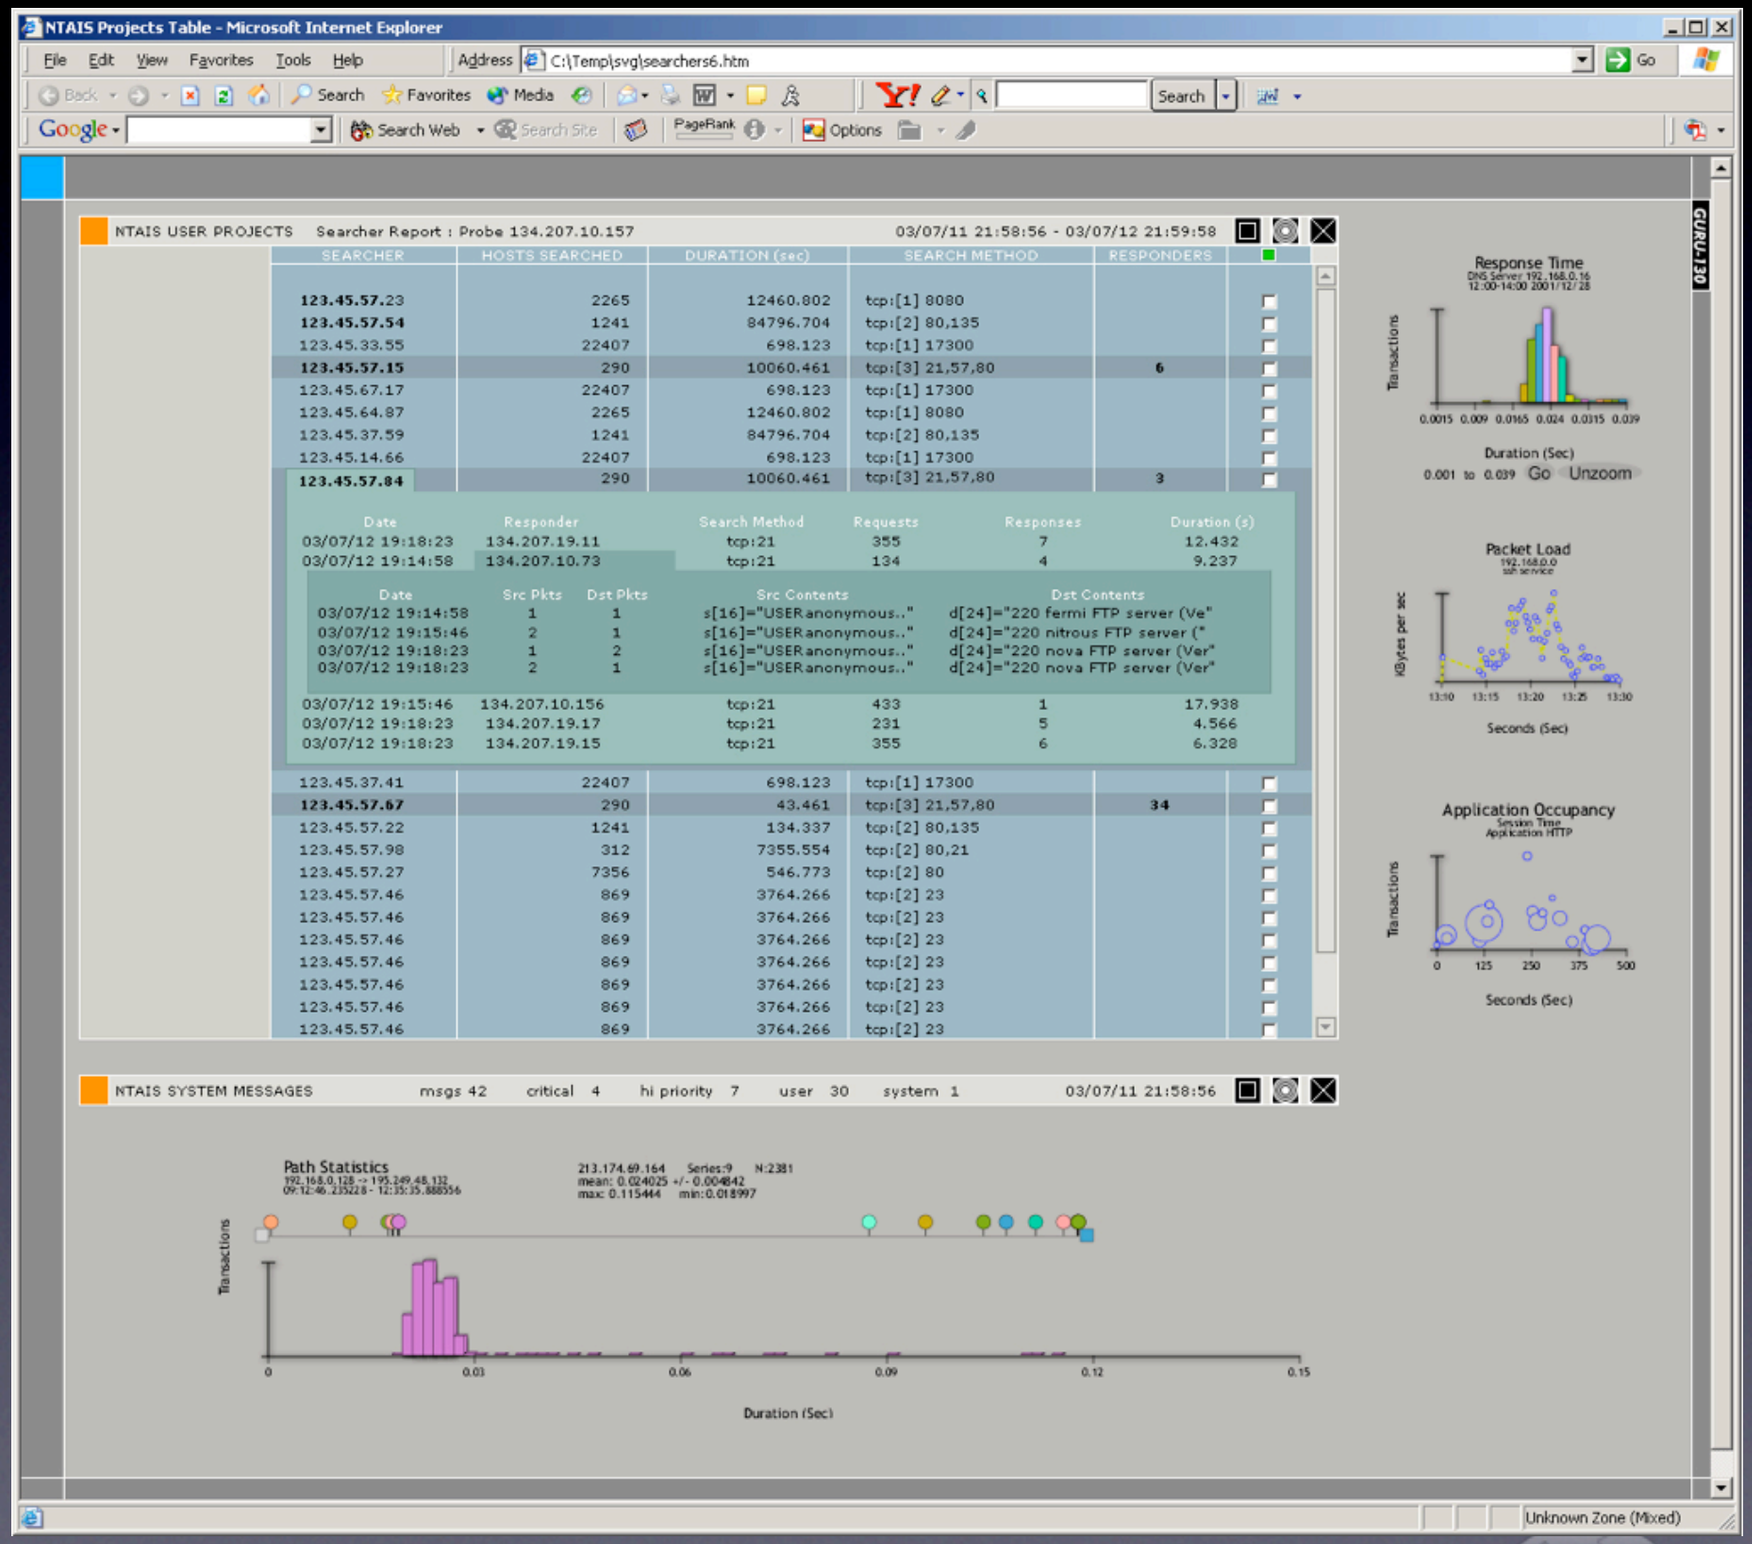
- ratop
  - Develop near-realtime view of what is going on in this network, right now.
  - Complex aggregation rule set for data representation
  - Allows for historical and current data stream comparisons
- rasqlinsert
  - ratop functions with state cached in mysql
  - goal is to have the last X minutes available
  - interface is an argus data stream





Situational Awareness

- GUI Strategies
  - Real Time Access
  - Web Based vs Native
  - Drill Down
  - Complex Data Methods
    - Time Series
      - Traditional
    - Spatial Information
  - Complex Visualizations
  - Generally not well used



Network Perception Goals

- Total Semantic Capture (Comprehension)
  - State initializations and transitions
  - Policy dissemination and enforcement
  - Topology, resource allocations, error conditions
- Context Awareness
  - Multi-Layer Identifiers (ethernet, MPLS labels, etc....)
  - Globally sychronized uSec timestamps
- Enable Near Real-Time State Awareness
  - Large scale access and data sharing
  - Multi-dimensional Correlation
- Complete Historical Reconstruction





Network Sensing Strategy

- Third-party Control/Data Plane monitor/sensors
  - Can't rely on the network switch/router vendors to do this.
- Each network device must provide complete Control Plane packet capture
  - Any packet that originates from or terminates on the device must be captured in its entirety.
  - Data must include port of origination/transmission, direction and UTC time stamp.
  - Before and after any encryption/decryption.
- Data plane flow data available for sharing, status reporting, and archival.
- Now we have the data we need to drive Data/ Control Plane Situational Awareness.





Network Information Model

- Multi-tiered Information Model
  - Not every application needs the same type of information
  - System needs to allow "customer" to define what it wants
- And, as conditions change, level of detail and frequency of status reports also needs to change

1. Data/Control Plane Service Existence Flow Strategy
   1.1. Matrix Flow with Service Identifiers
   1.2. Operational/Security Fault Status Flow Records
2. Data/Control Plane Service Performance Strategy
   2.1. Transactional Flow with Ops and Performance Attributes
   2.2. Operational Fault Status Flow Records
3. Total Packet Content Flow Strategy
   3.1. Transactional Flow with Aggregated Content
   3.2. Complete Remote Packet Capture





Packet / Flow Strategies

- Packet data for complete comprehension
- Flow data provides multi-tiered data model.
  - Data Reduction / Semantic Preservation
    - Service Oriented Transactional Abstractions
    - Complex Data Representations
    - Flexible Compression Strategies
    - Multiple Flow Content Representations
  - Semantic Access Control Schemes
    - Inter/Intra Domain Data Sharing
    - Complex Data Aggregation Scoping
    - Anonymization
  - Cross Domain/Dimensional Correlation
    - Unified Object Specifications
    - Self-Synchronization Methodologies
    - High Resolution Timestamping





Discovery Detection

Network scan detection is not as important as it was decades ago, but understanding who responds to scans, and what they respond with, is still a very important thing to know.

- radark.pl
  - Track IP addresses that attempt to connect to non-existent hosts (a network explorer)
  - If these network addresses ever get a response from existing nodes on non-public service SAPs, then report these accesses.
  - Include what the responder responded with.





Conclusions

- http://qosient.com/argus
- argus-info@lists.andrew.cmu.edu
- carter@qosient.com





Network Flow Data Fusion
GeoSpatial and NetSpatial Data Enhancement

FloCon 2010
New Orleans, La

Carter Bullard
QoSient, LLC

carter@qosient.com









Carter Bullard carter@qosient.com

- QoSient - Research and Development Company
  - Naval Research Laboratory (NRL), GIG-EF, JCTD-LD, DISA, DoD Network Performance and Security Research
- Inventor/Developer Argus http://qosient.com/argus
- FBI/CALEA Data Wire-Tapping Working Group
- QoS/Security Network Management - Nortel/Bay
- Security Product Manager – FORE Systems
- CMU/SEI CERT
  - Network Intrusion Research and Analysis
  - NAP Site Security Policy Development
  - Network Security Incident Coordinator
- NFSnet Core Administrator (SURAnet)
- Standards Efforts
  - Editor of ATM Forum Security Signaling Standards
  - IETF Working Group(s) Contributor
  - Internet2 Security WG
  - NANOG





FloCon 2010 Flow Innovation

This year's conference will focus on flow data analysis within the context of other data sources. Presenters are encouraged to consider how flow is a piece of the puzzle.

Which Puzzle?

- Cyber-Situational Awareness and Network Defense (CND)
  - Near real-time awareness of threats, status, and performance, with awareness of external attacks and insider abuse/misuse.
- Assured Enterprise Management and Control
  - Critical infrastructure must operate as intended, with management, control and information assured.





Cyber-Situational Awareness

Level I SA - Perception

- The perception of elements in the environment within a volume of time and space
- Involves timely sensing, data generation, distribution, collection, combination, filtering, enhancement, processing, storage, retention and access.

Level 2 SA - Comprehension

- Understanding significance of perceived elements in relation to relevant goals and objectives.
- Involves integration, correlation, knowledge generation.

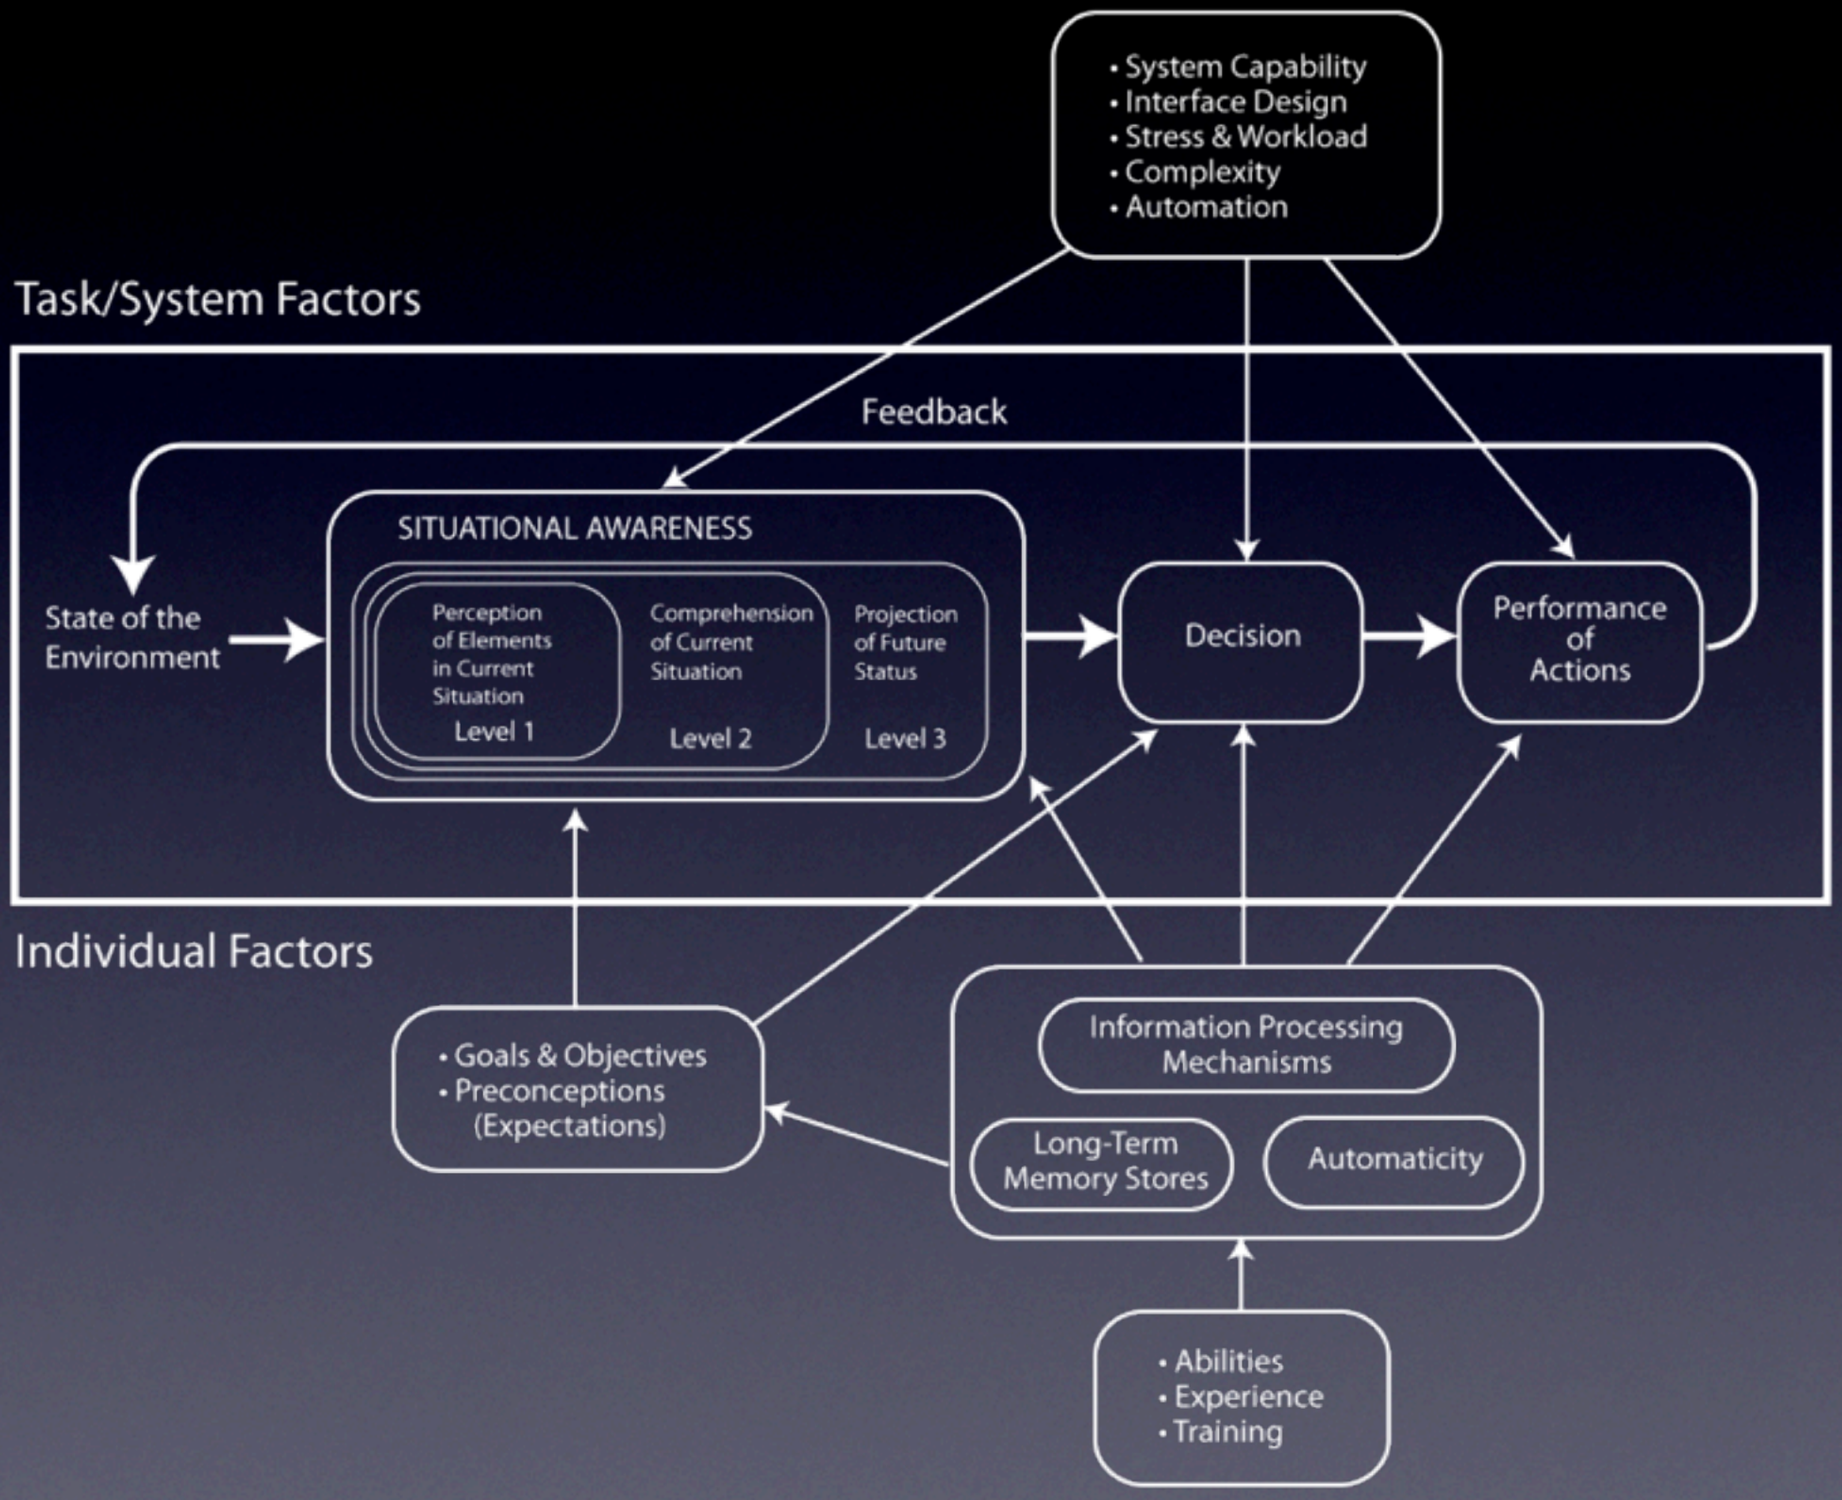
Level 3 SA - Projection of Future Status





Endsley, M. R. (1995b). Toward a theory of situation awareness in dynamic systems. Human Factors 37(1), 32-64.

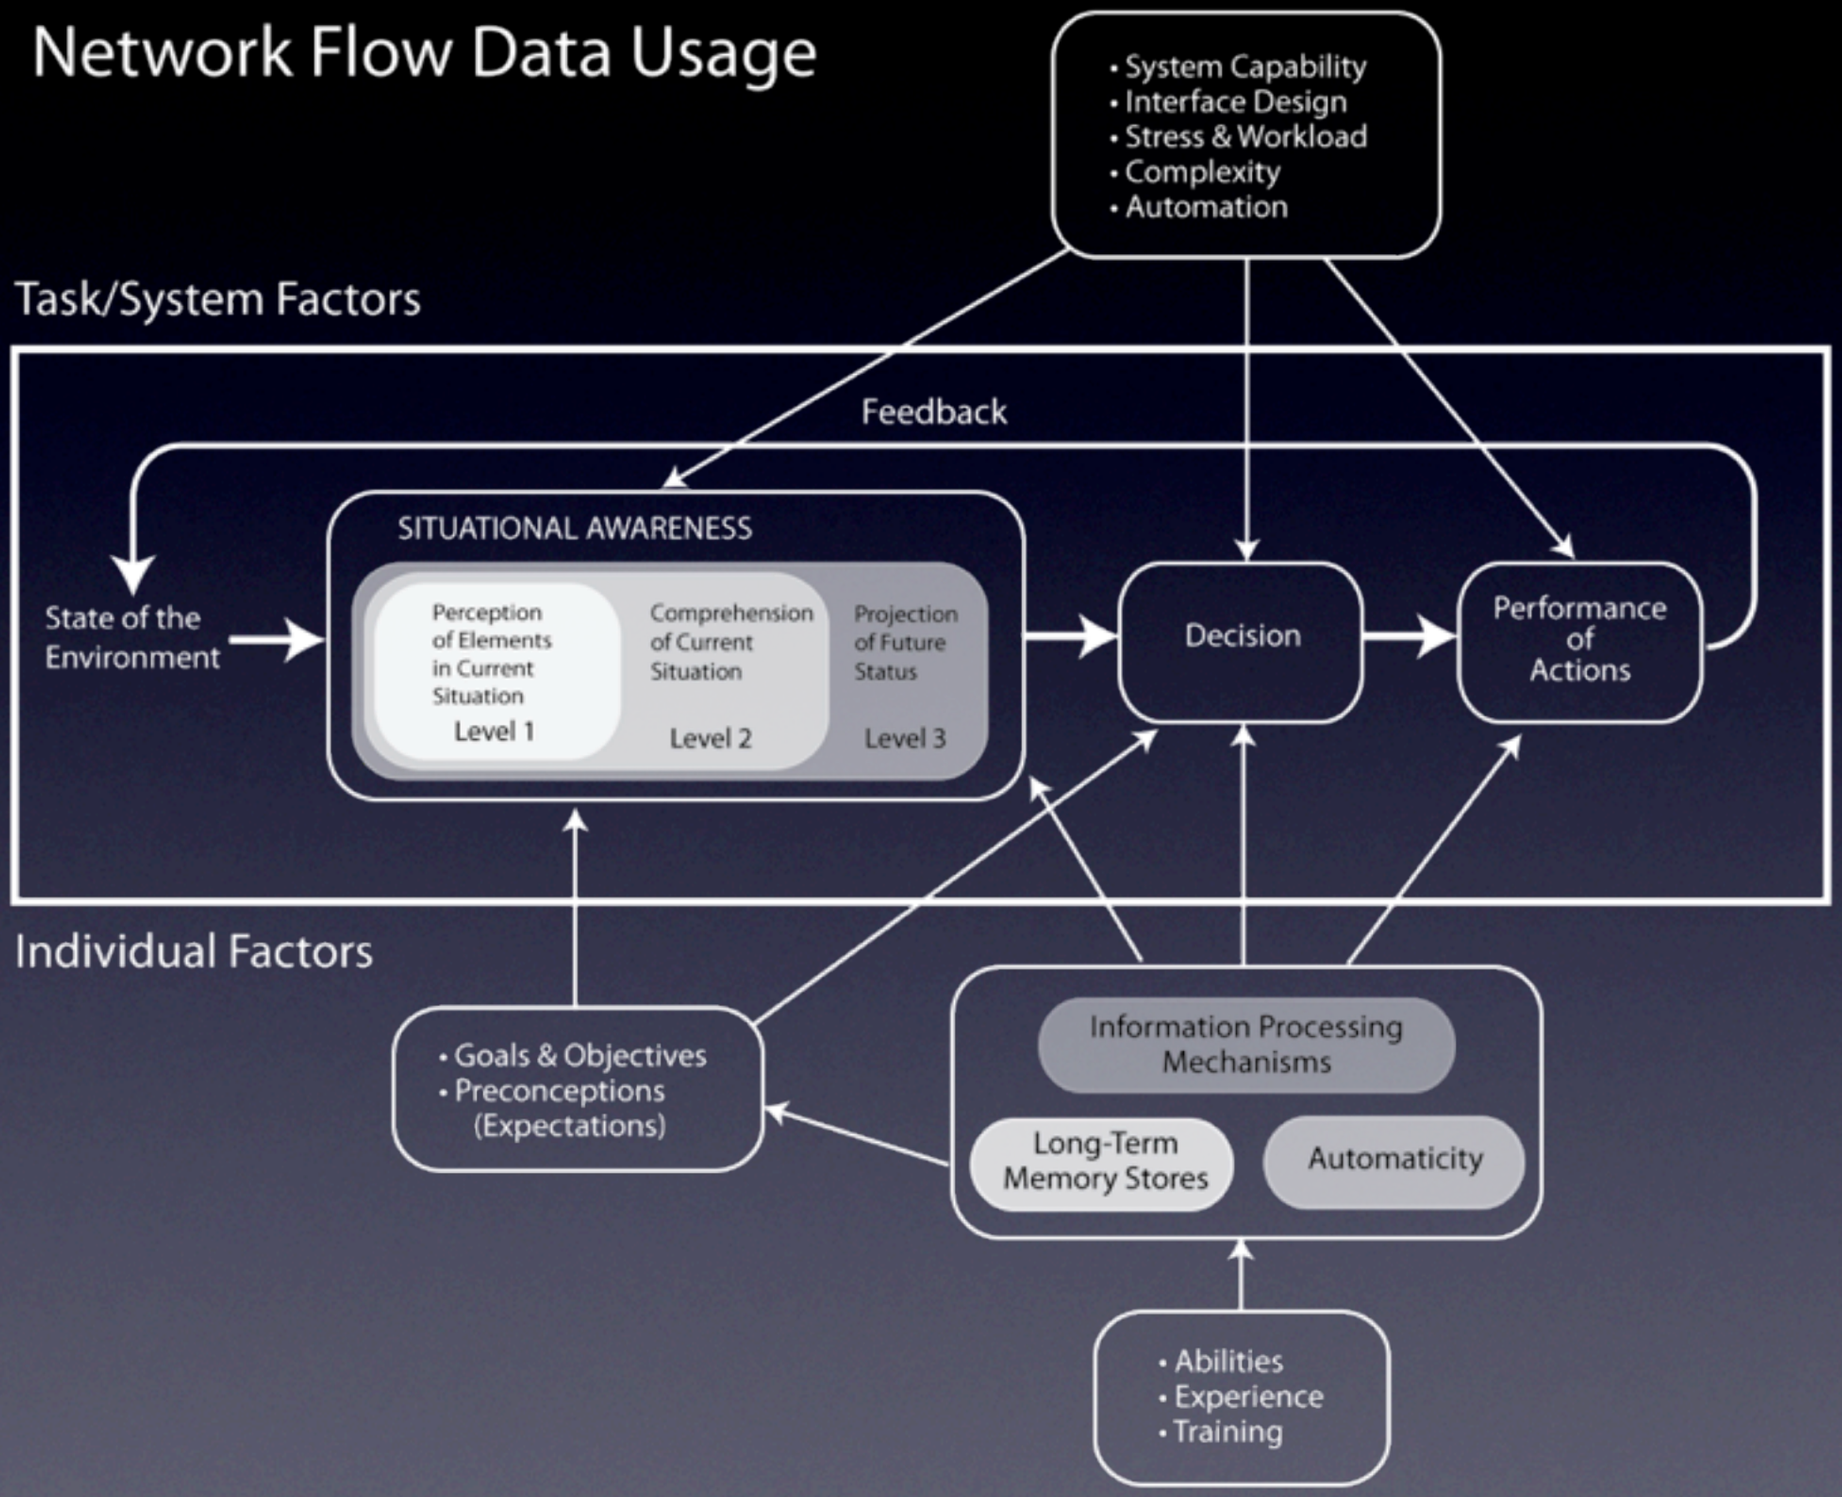
Model of Situational Awareness in Dynamic Decision Making



Endsley, M. R., *Human Factors*, Vol. 37, No. 1, 1995. Copyright 1995 by the Human Factors and Erogonomics Society. All rights reserved.

Model of Situational Awareness in Dynamic Decision Making

Network Flow Data Usage



Endsley, M. R., *Human Factors*, Vol. 37, No. 1, 1995. Copyright 1995 by the Human Factors and Erogonomics Society. All rights reserved.

Who/What/When/Where

Sometimes, 'Where' is the only criteria for comprehending that there is a problem.

- Data isn't suppose to be coming from there.
- Data isn't suppose to be going that way.
- Data should to be coming from there but .....
- Where is this data coming from !!!!!!

Network flow data can be used in perception and comprehension of some of these very complex concepts, but the data needs to have some specific qualities in order to successively support 'where' functions.





Who/What/When/Where

- GeoSpatial Information
  - Association with geographic information, GIS and Geomatics.
    - GeoLocation
      - Identification of 'real-world' geographic location information
      - Generally IP Address, MAC, RFID, Triangulation Based (rarely GPS based).
      - Time Zone, Country Codes, Region, City, Postal/Zip Codes, Lat/Lon
      - Commercial/Open Source Data Sources
        - Regional Internet Registries
        - ISP Provided Information

- Used primarily for Marketing and Directed Advertisement.
- Applications to E-commerce are emerging (taxation).
- VoIP/SIP Based Emergency Services.
- Lots of Standards (OGC, IEEE, W3C, ITU, IETF)
- Little guidelines for privacy protection issues

- Important Issue in Mobile Ad-hoc Networking
  - Gradient, aspect and visibility
  - Distance optimizations for power minimization and path length

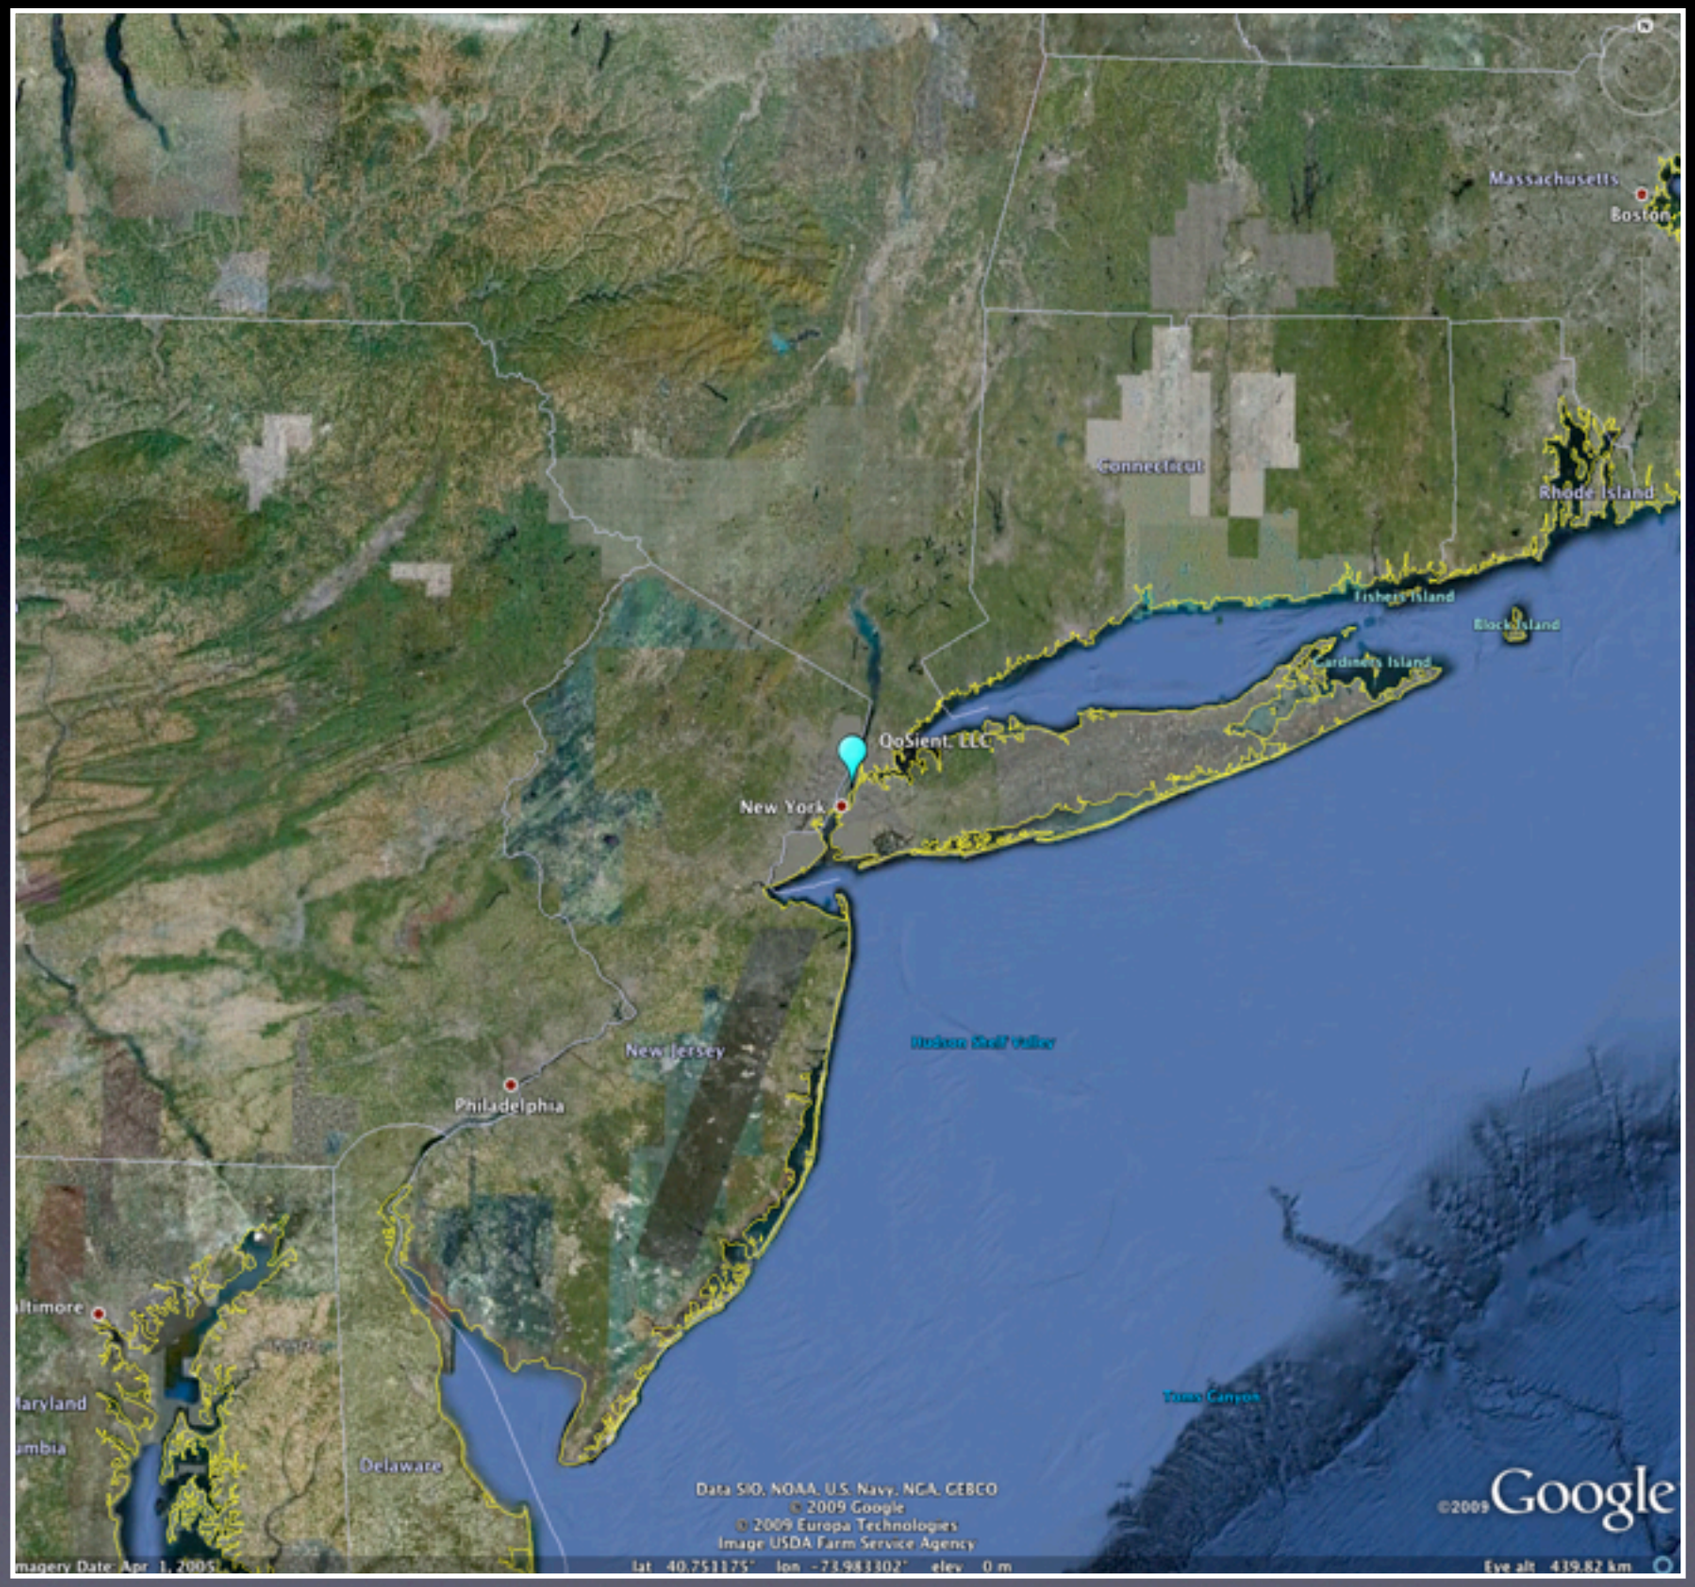




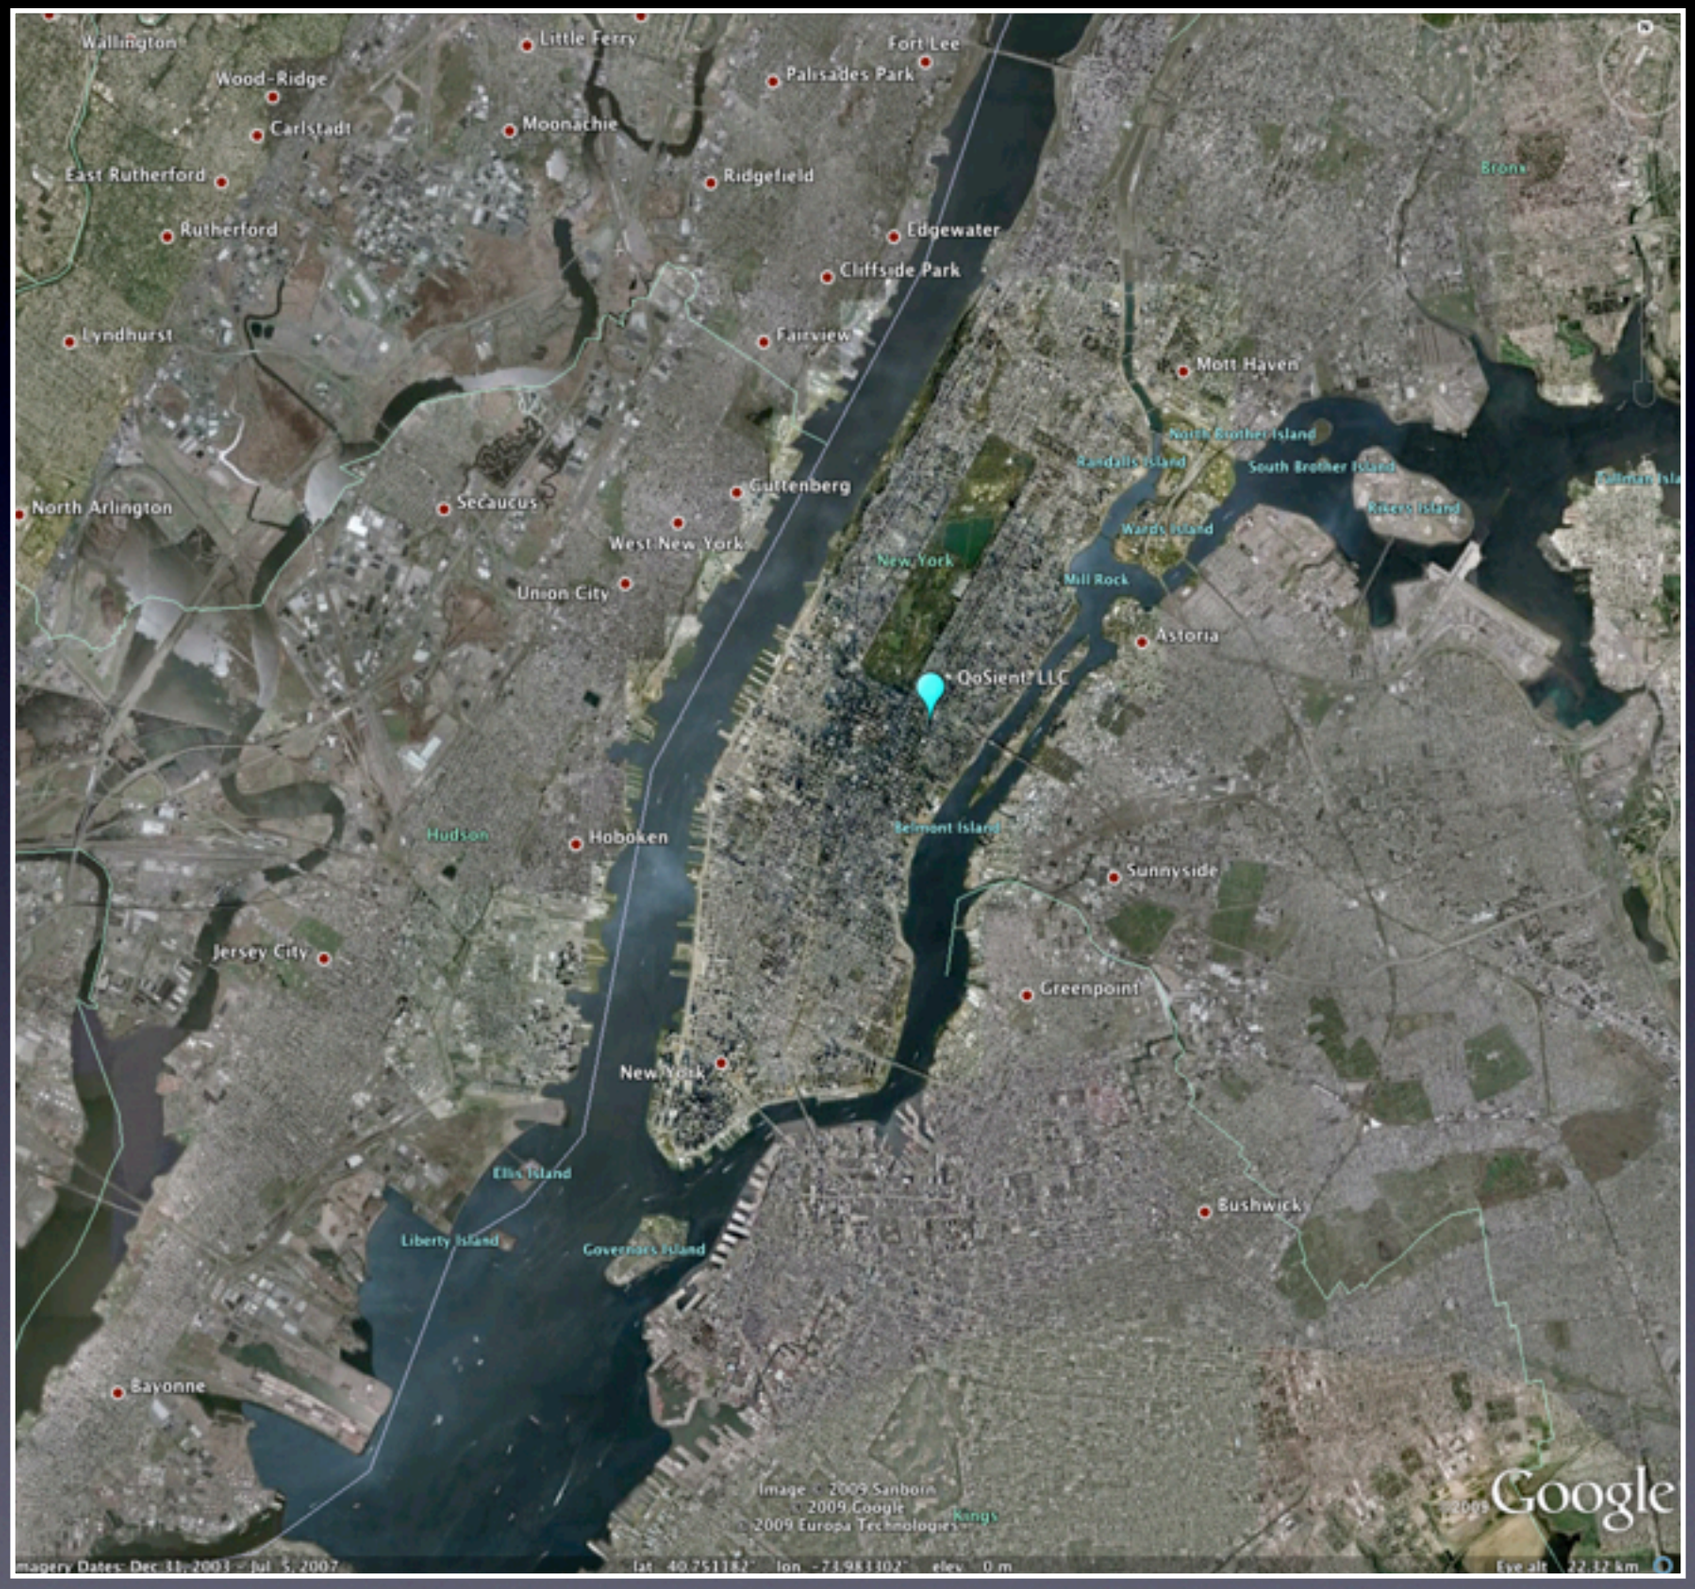
Where is QoSient.com?



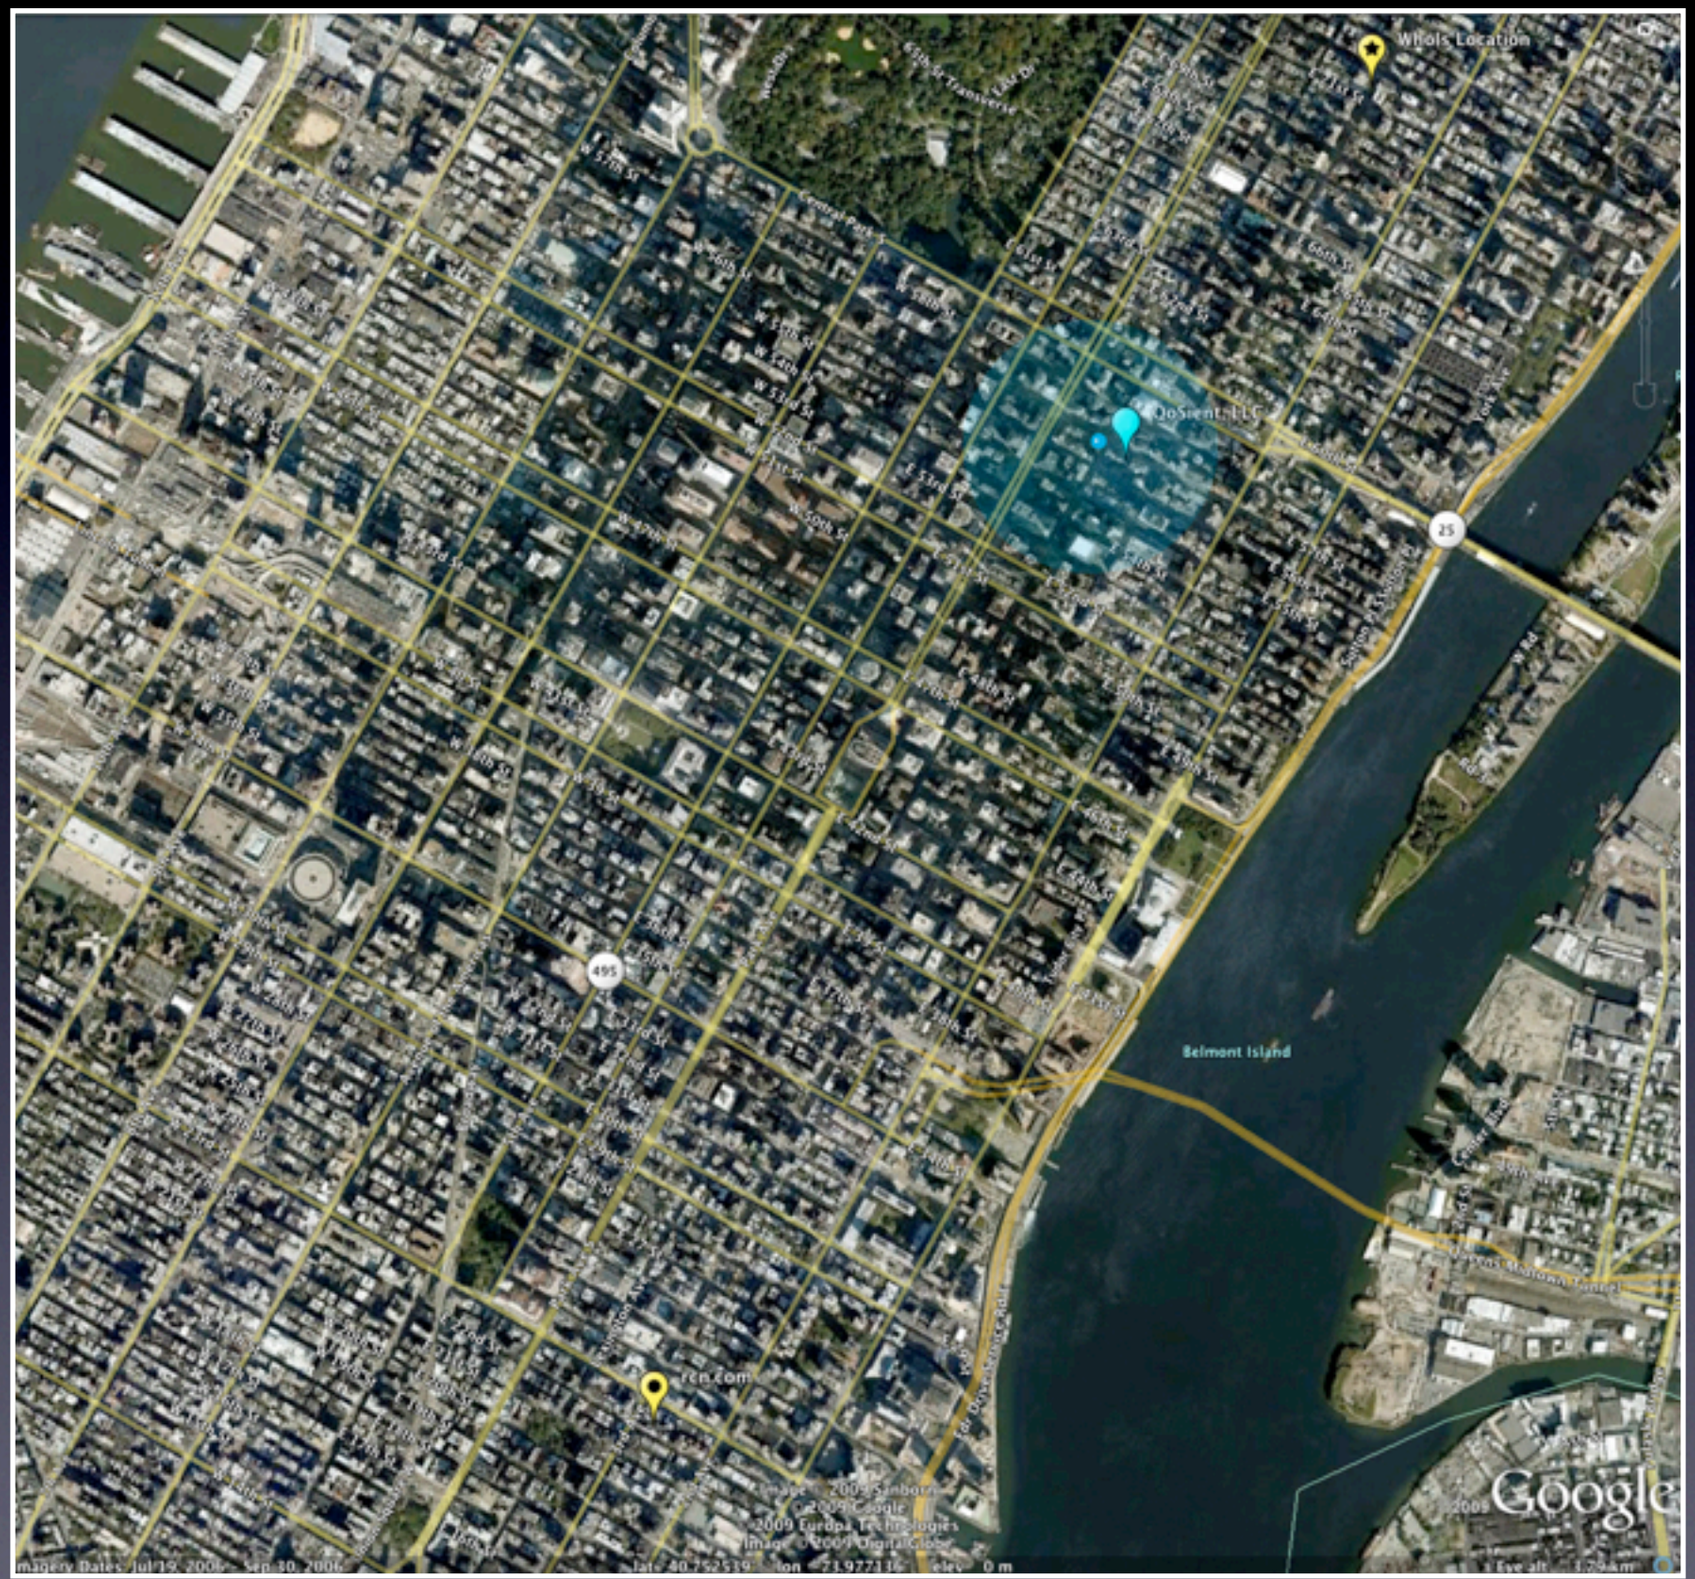
Where is QoSient.com?



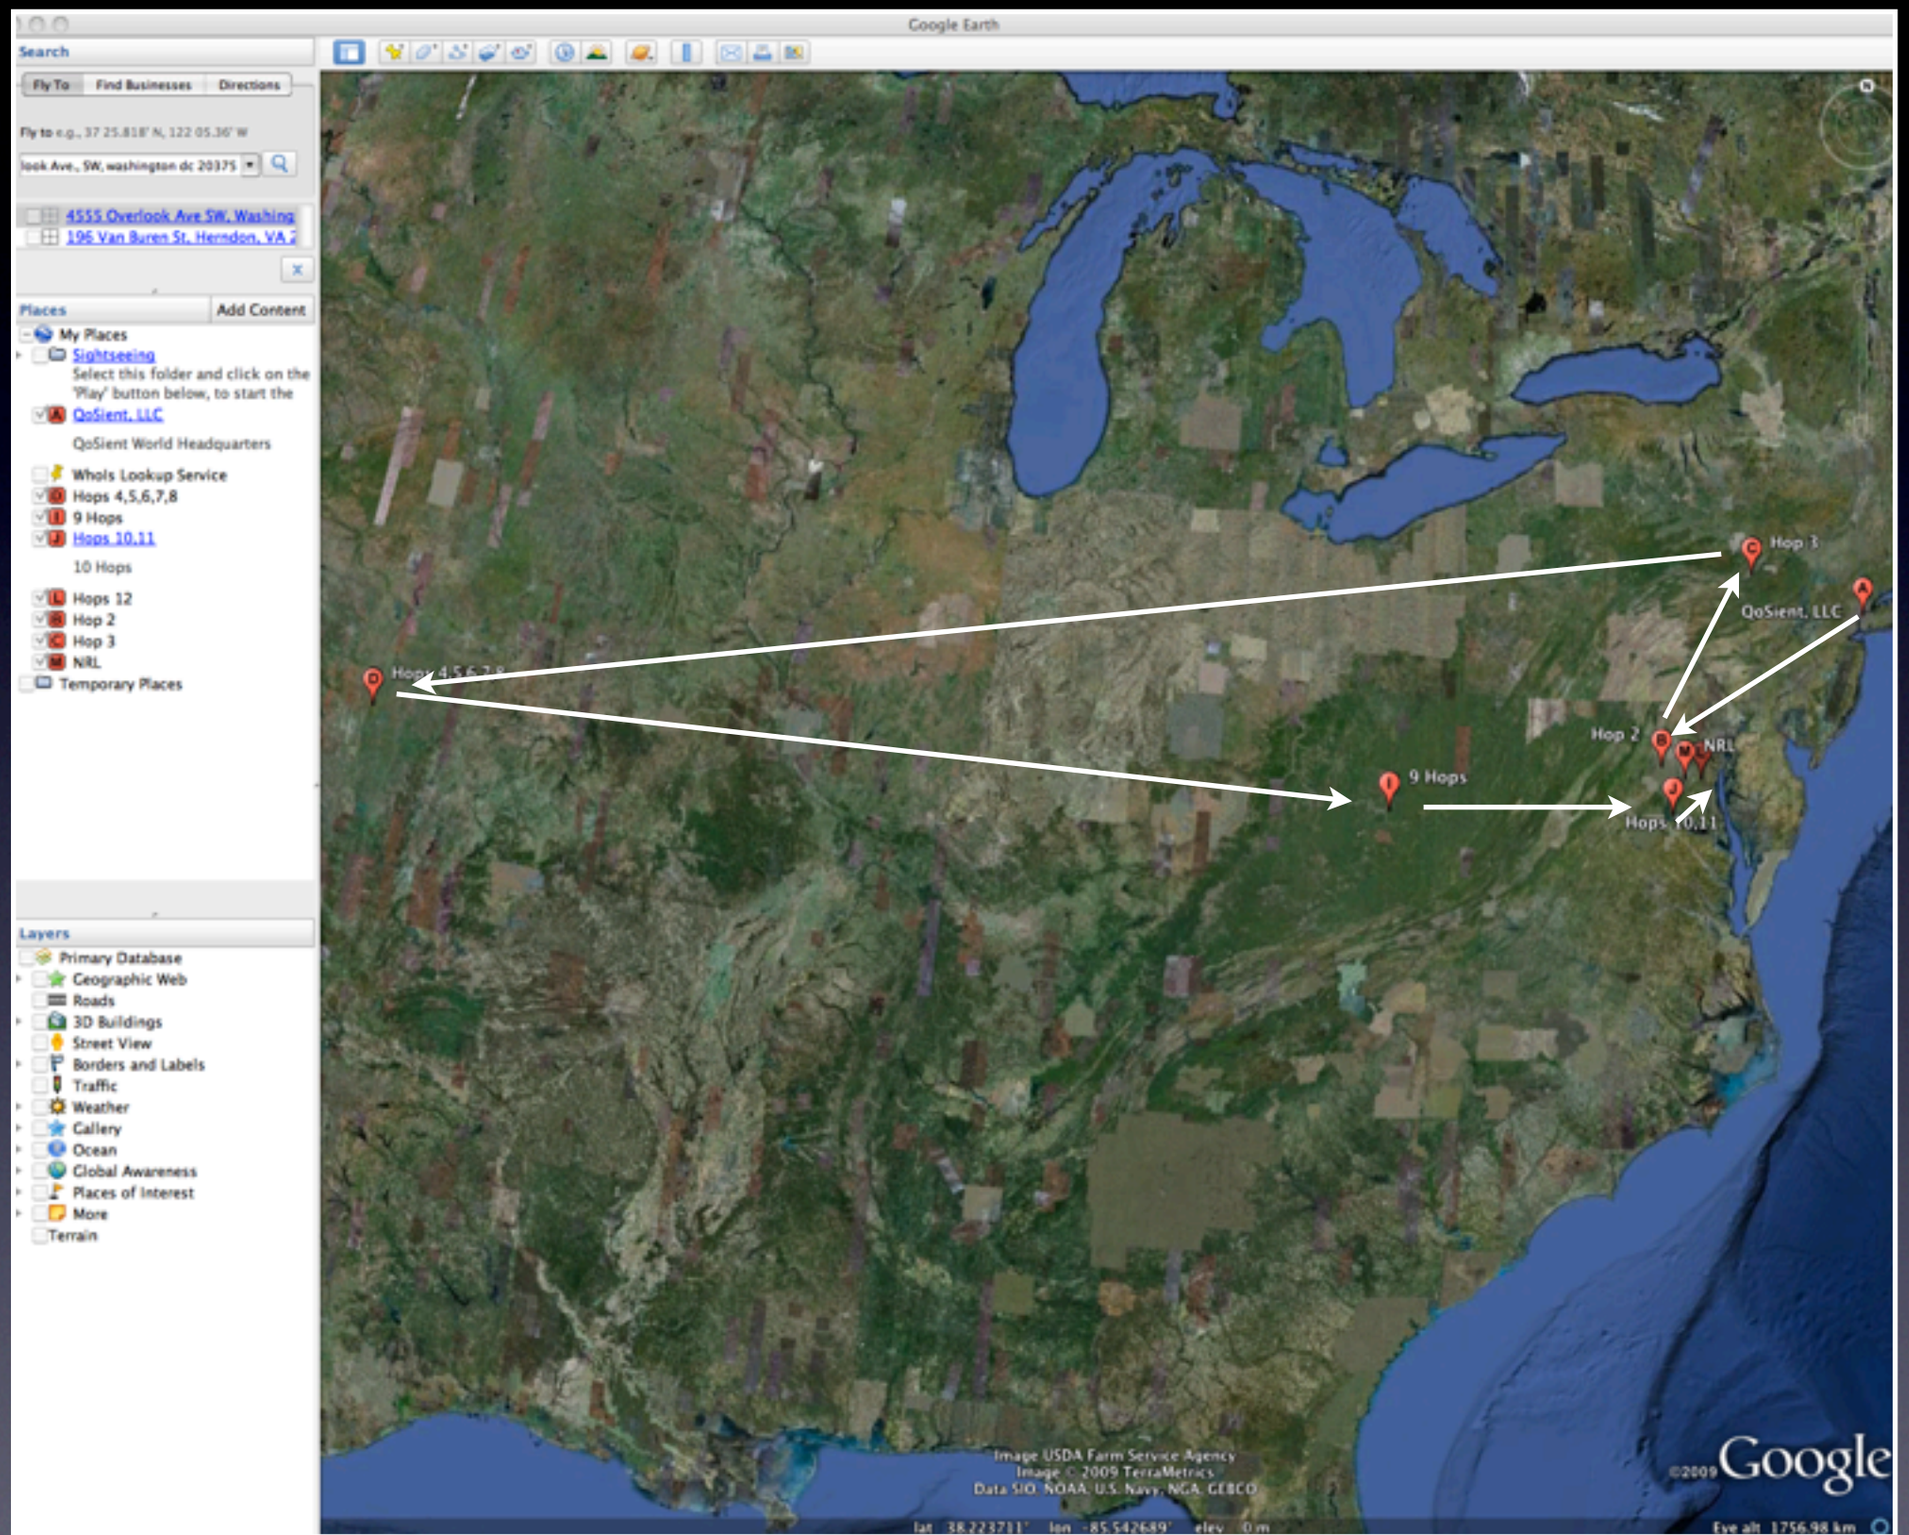
Where is QoSient.com?



Where is QoSient.com?



Network Path Information?



Who/What/When/Where

- Issues using geospatial information and flow data
  - There is no GeoSpatial Information in data packets
  - Most network flow data must be enhanced external to sensor
    - Flow data enhancement during/after data collection or distribution
  - No relational algebraic constraints on geospatial identifiers
    - IP addresses are not globally unique.
    - IP Address / Geolocation mappings are not formally managed/maintained.
  - Issues involve accuracy, relevancy, dynamism and time
  - IPFIX has not discussed geospatial/netspatial data support.





Who/What/When/Where

- Argus geospatial support
  - Flow Data Semantic Enhancement
    - radium() - collection based enhancement
    - ralabel() - post collection enhancement
  - metadata insertion strategy
    - saddr:lat=42.246532, lon=18.345261
      - geospatial information embedded in each record
  - direct GPS data insertion when available
- Support for printing, graphing, filtering, aggregation, and anonymization.
  - aniso lat/lon aggregation generates bounding box
  - lat/lon anonymization
    - constant offset projected onto either poles or ocean/land boundaries





China Syndrome

- But not all is as it appears to be.
  - QoSient.com constantly scanned by IP addresses from CN
    - Using ARIN databases for country codes.
    - Not a bother at all really.
  - One flow presented with estimated hop-count of 4 hops
    - Argus uses TTL to estimate hops (nearest 2^x - observed TTL).
  - Modified ping to source yielded RTT of less than 5 mSecs
    - Speed of light distance is 465.71 miles one-way.
    - Network distance estimates usually put 5 mSecs close, 0-20 miles.
    - So, what's up with this.

- Source address spoofing?
- Router root-kit attack?
- Routing infrastructure attack?





Domain
Name Server

DNS
Root Servers

BGP

AAA

MPLS Network

OSPF

STP

ARP

RSVP-TE/LDP
IS-IS-TE
BGP

OSPF

End
Station

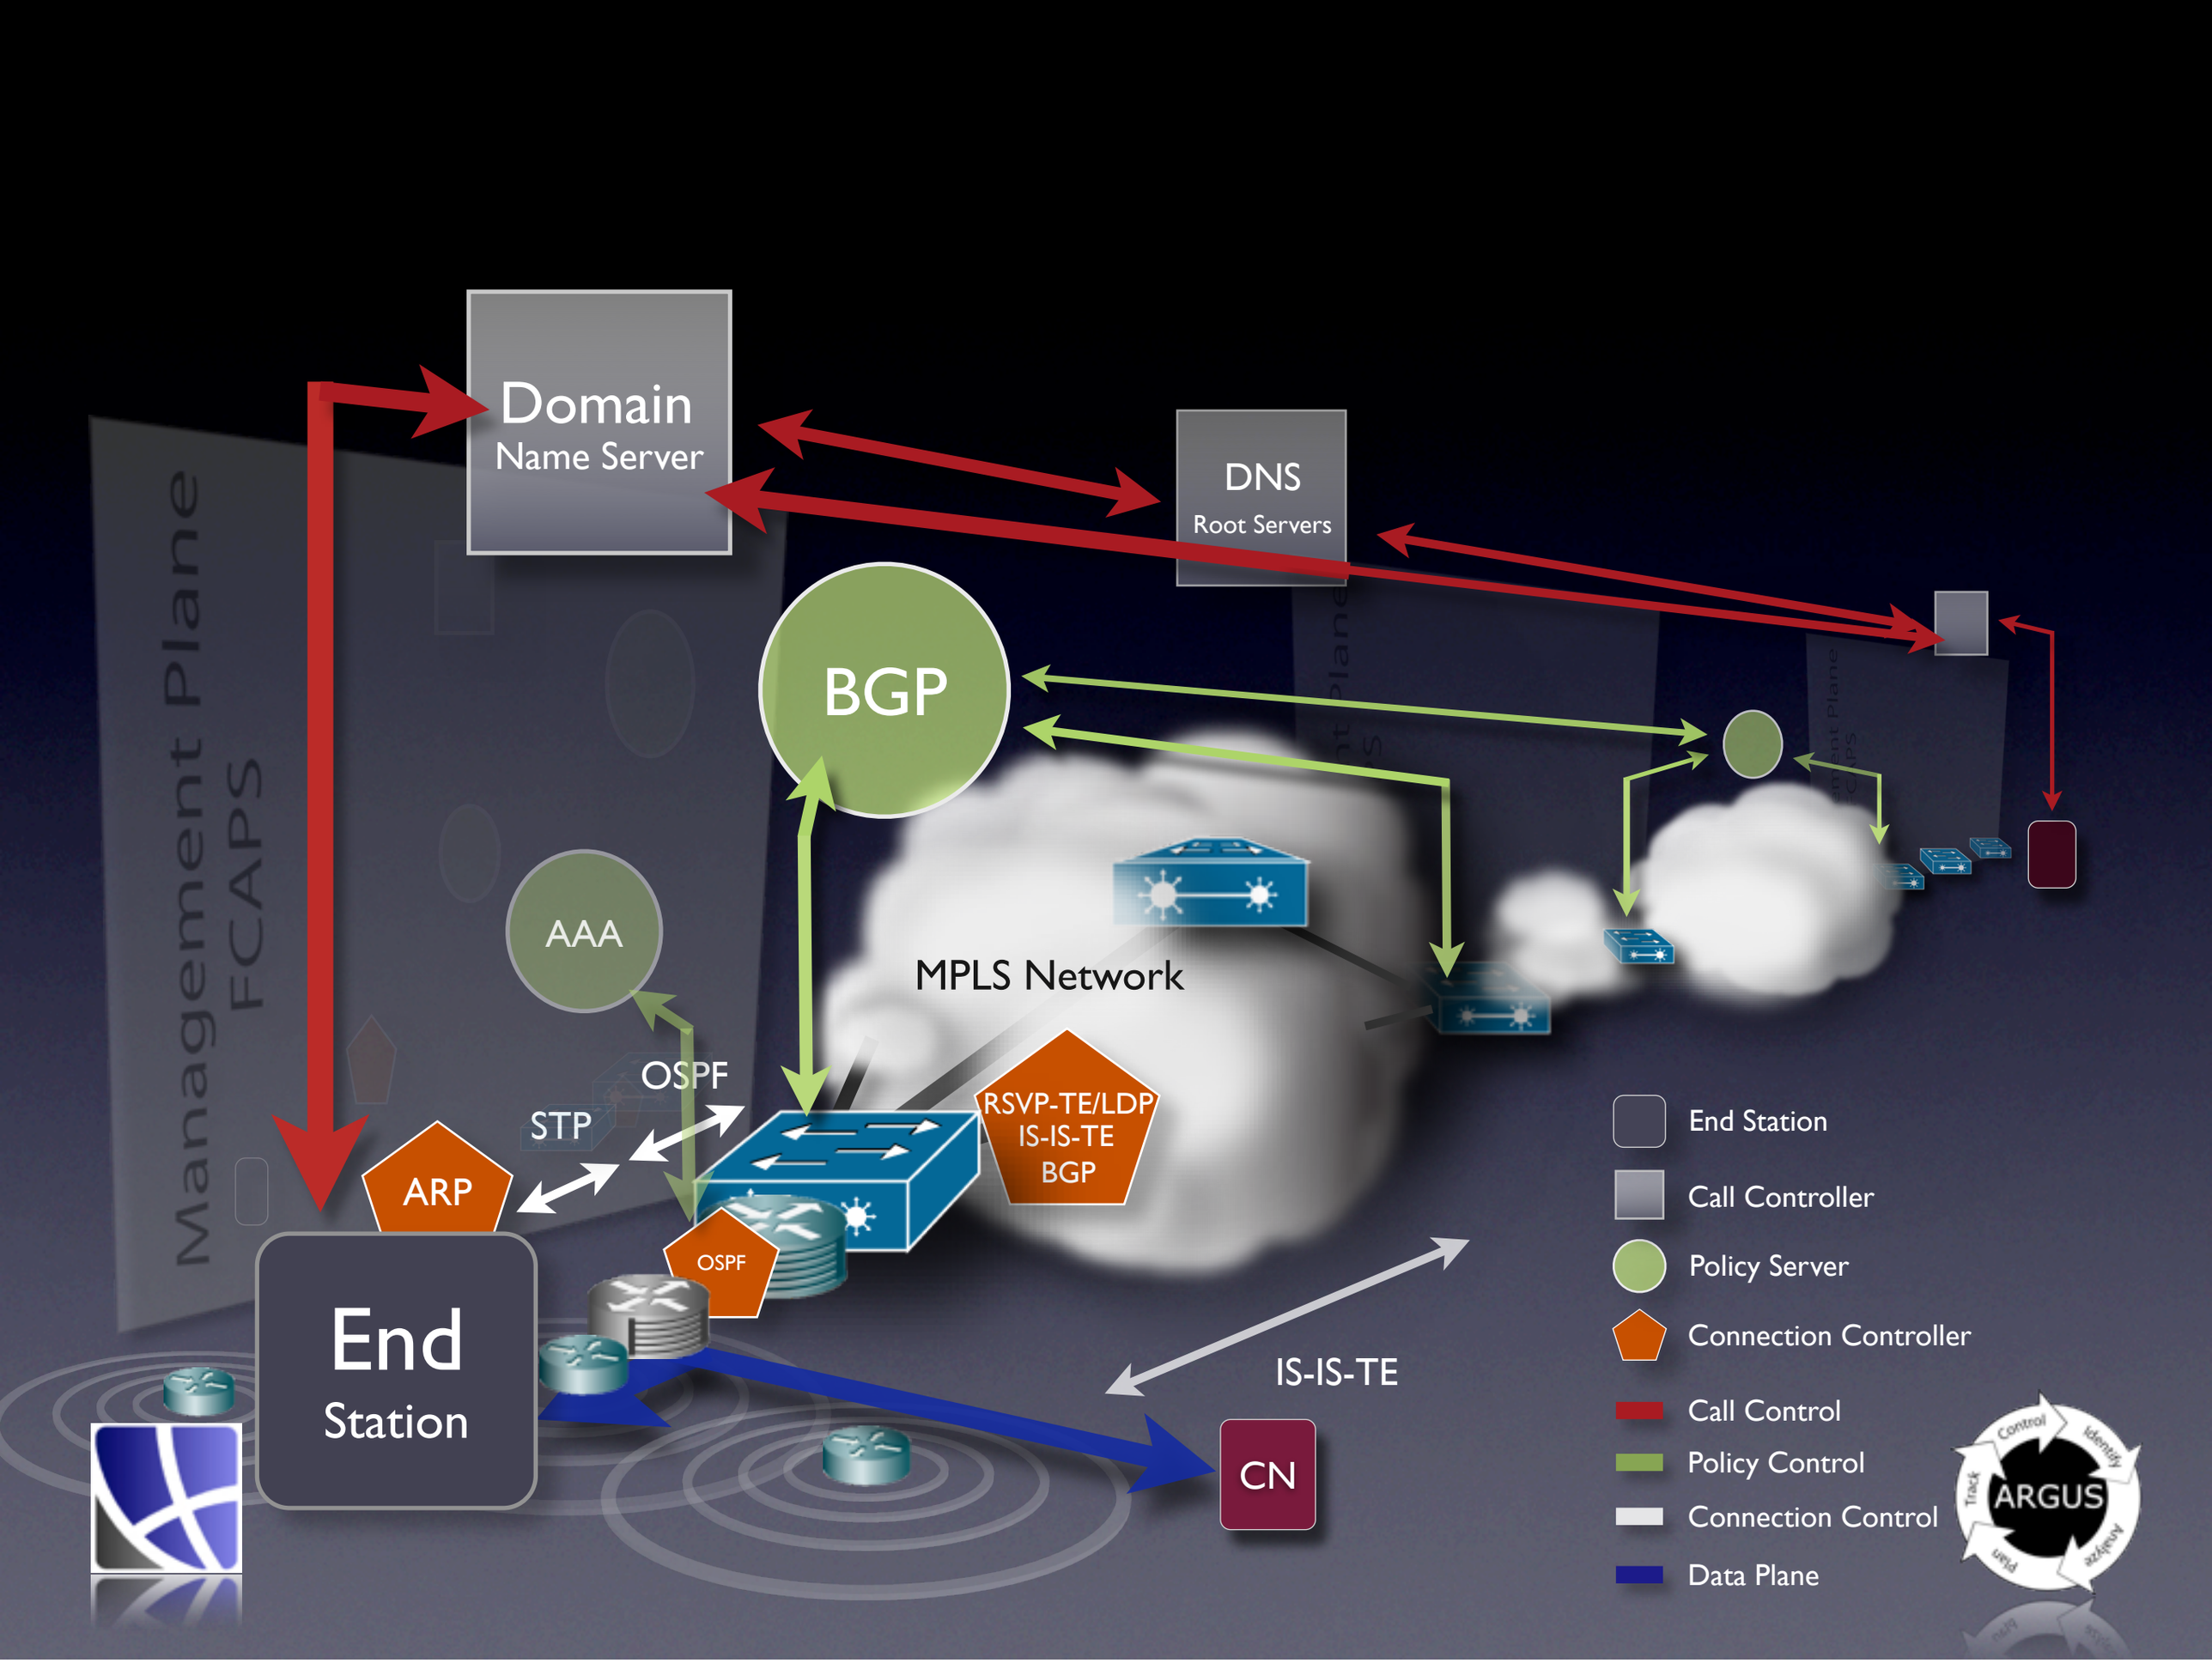
IS-IS-TE

CN

Management Plane
FCAPS

End Station
Call Controller
Policy Server
Connection Controller
Call Control
Policy Control
Connection Control
Data Plane

ARGUS


Domain
Name Server
DNS
Root Servers
BGP
AAA
MPLS Network
OSPF
STP
ARP
RSVP-TE/LDP
IS-IS-TE
BGP
OSPF
End
Station
IS-IS-TE
CN
End Station
Call Controller
Policy Server
Connection Controller
Call Control
Policy Control
Connection Control
Data Plane
ARGUS

Who/What/When/Where

- How Do You Detect This?
  - Geospatial /Netspatial Incongruity
  - Network Distance Estimation and Correlation
    - Service Discovery, Service Usage Optimization, Group Join Optimizations, Shortest Path Routing
  - Methods
    - Global Network Positioning (GNP and NPS), CDN (Akamai), Internet Iso-bar, Internet Distance Maps (IDMaps), Vivaldi, Dynamic Distance Maps (DDM), RON, Landmark Clustering, Dynamic Landmark Triangles, Netvigator
  - All Network Distance Estimation Methods use simple active RTT metrics such as ping() and traceroute(), differentiations involve sampling strategies and statistical analysis.





Need For Active Elements?

- Timeliness of Determination/Validation
  - See a packet from some interesting IP address
    - Need to timely propagate the perception of address
    - Make GeoSpatial assessment
    - Decide to make some form of network estimation
    - Schedule ping/traceroute/probe....

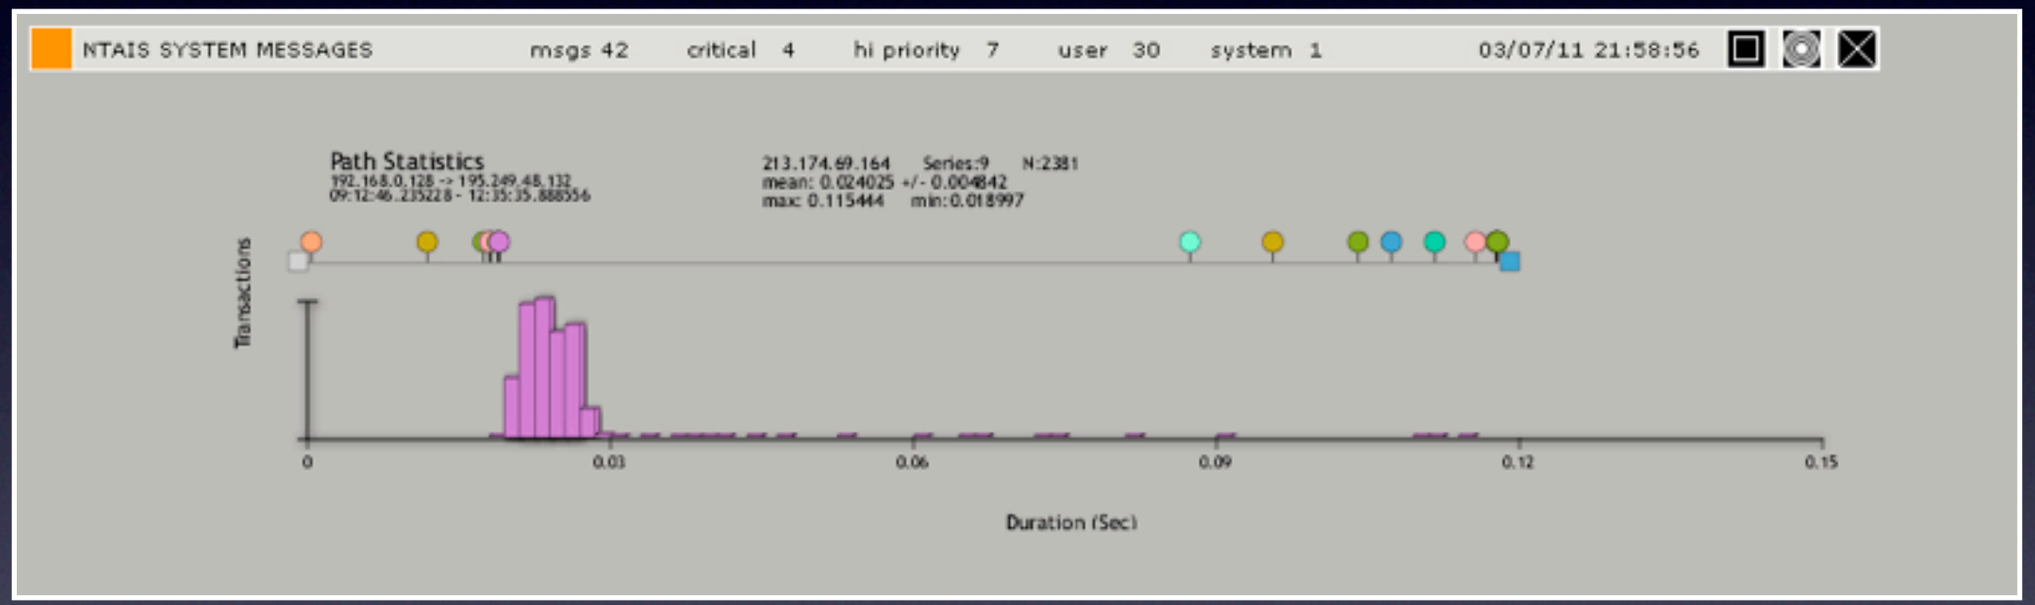
- Flow sensors passively capturing network distance estimation metrics
  - Bi-directional flow monitors
    - Capture RTT regardless of protocol type
    - P1-P2 flow tracking captures traceroute information
  - Billions of Location Metrics Per Day
    - 8-10K Host Associations Per Work Group
    - 25% Infrequent, 35% Transient, 40% Persistent



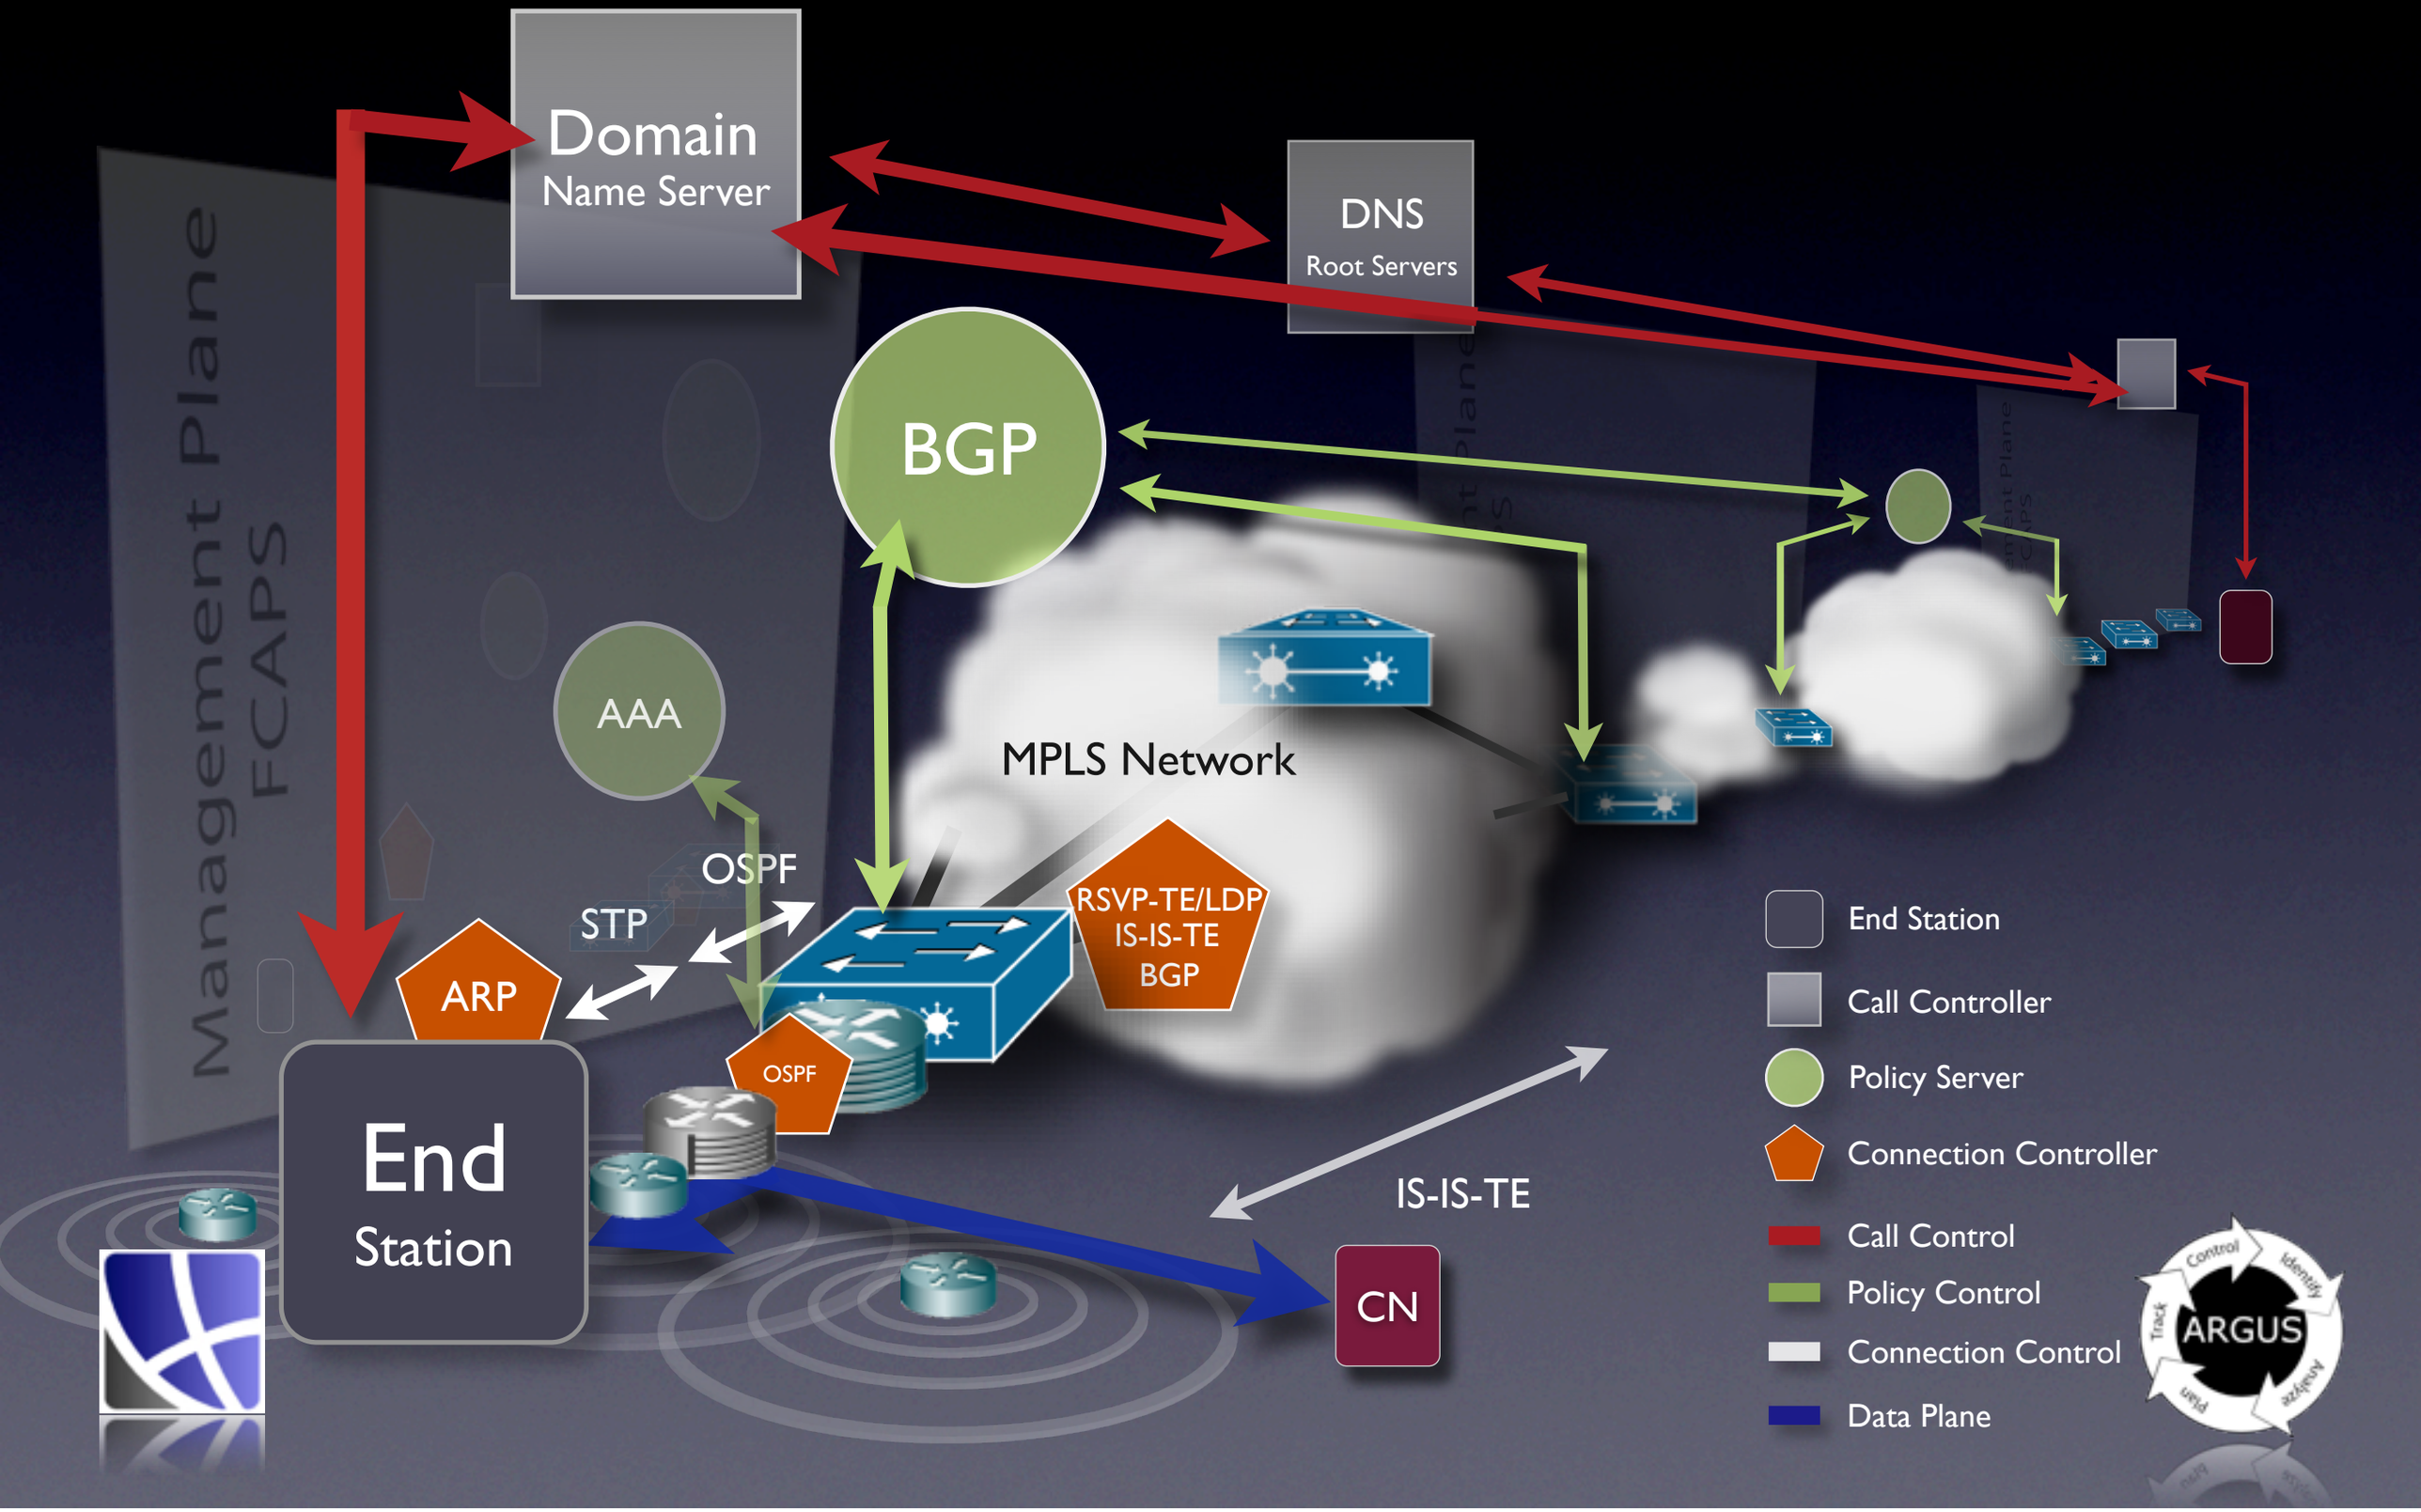


Geo/NetSpatial Correlation

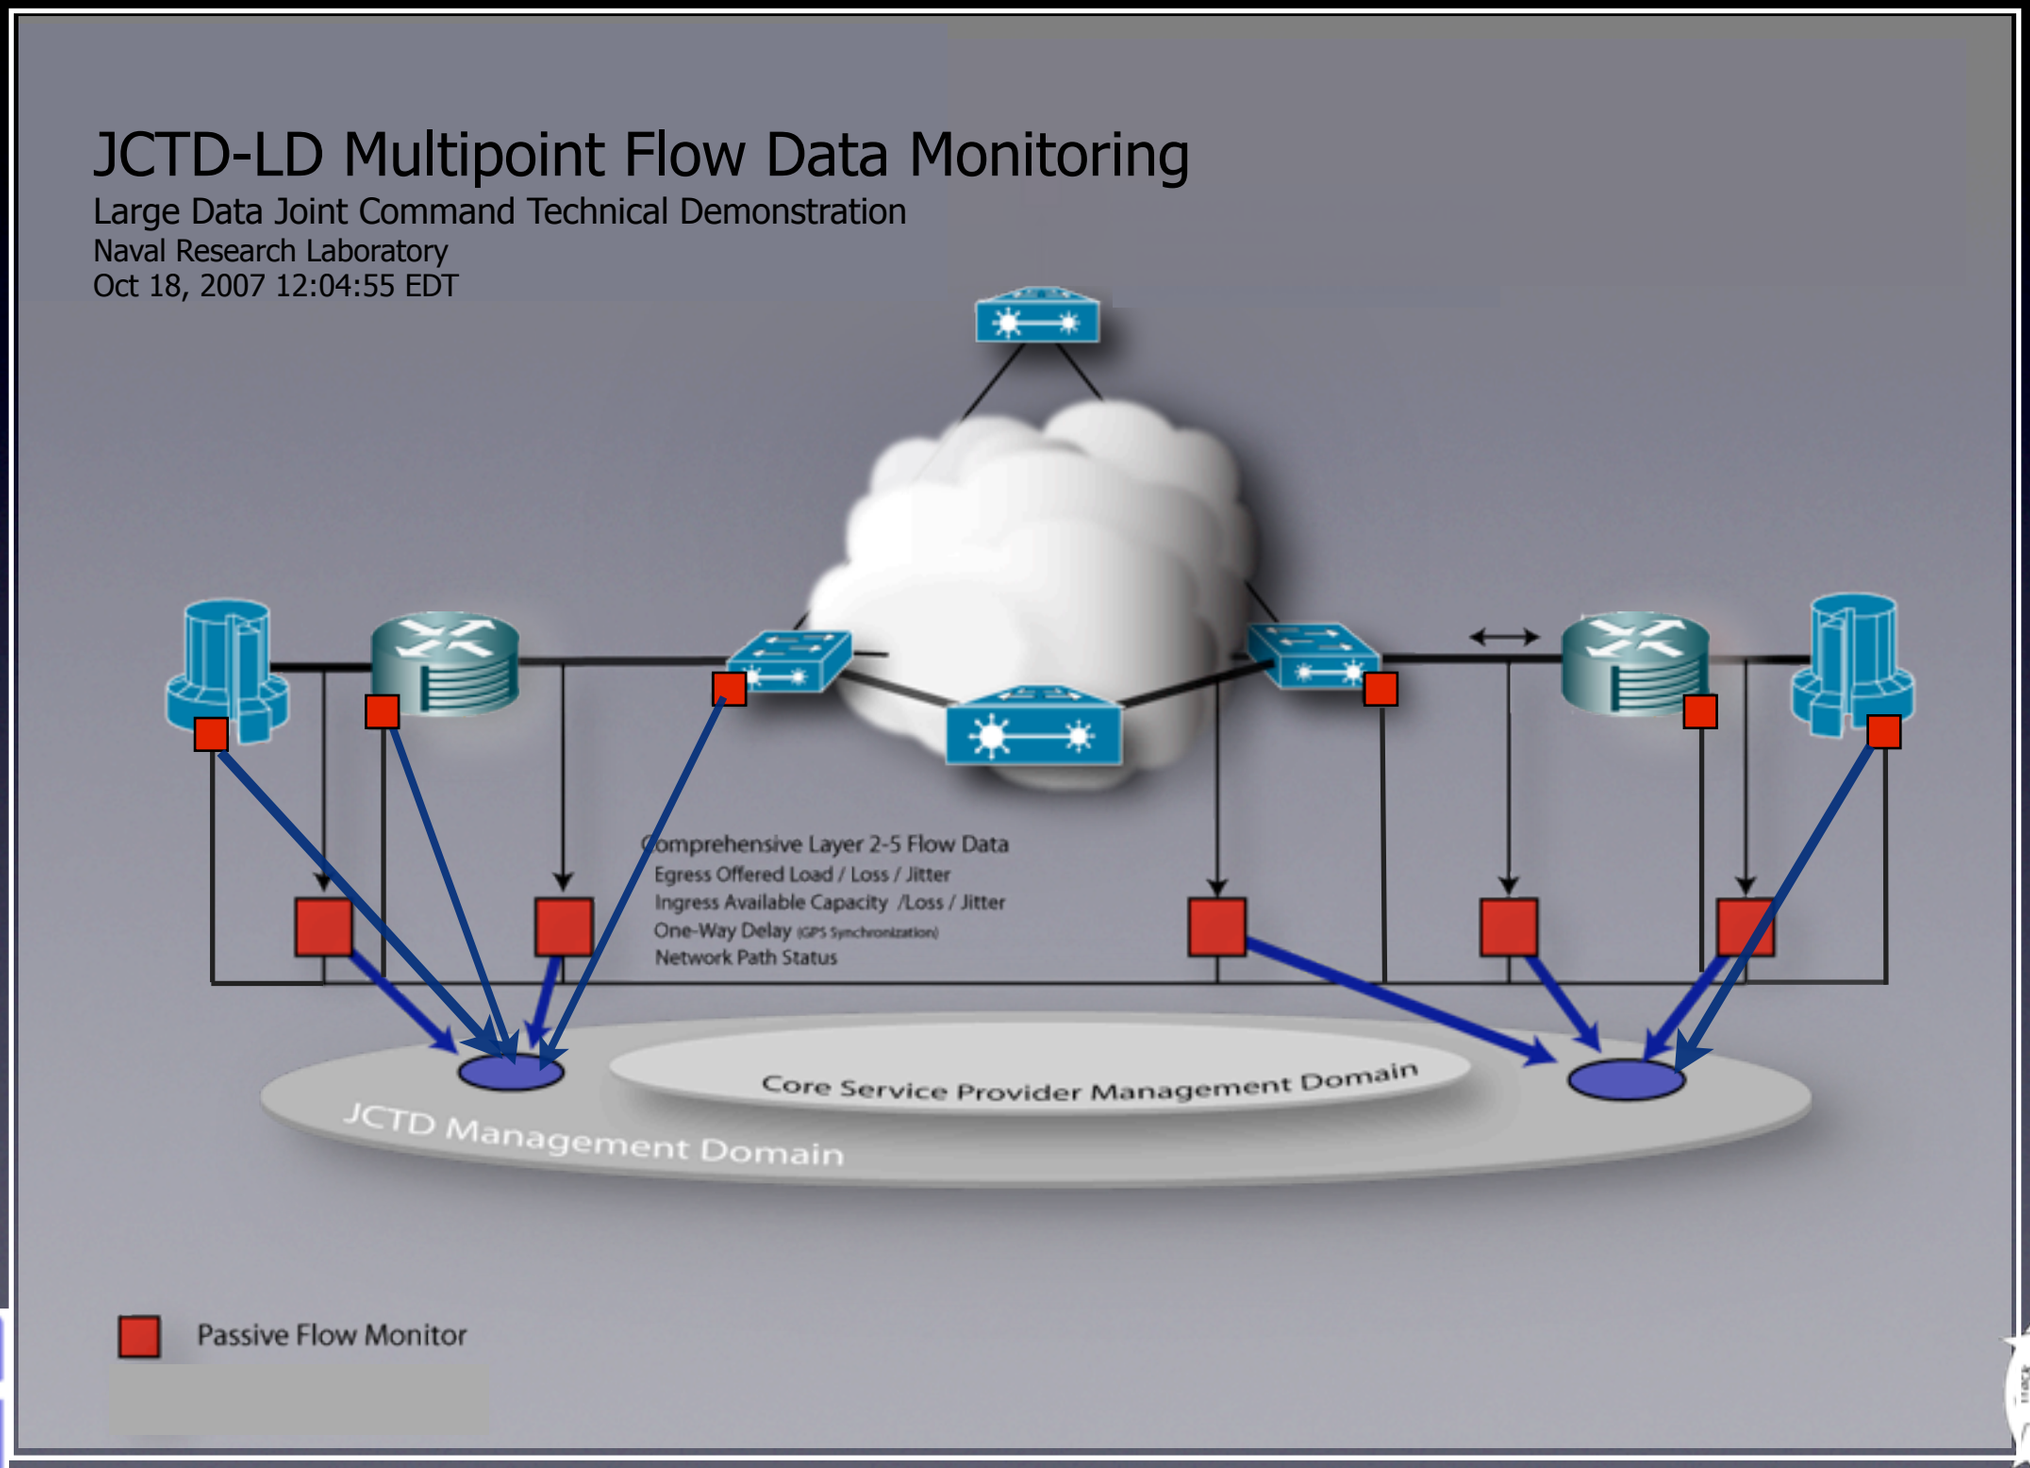
- Network Distance Estimation Accuracy
  - Large samples generate good results



- But, network distance does not relate to physical distance.
- At least you can get a sense that something is up


Management Plane
FCAPS
Domain
Name Server
DNS
Root Servers
BGP
AAA
MPLS Network
OSPF
STP
ARP
OSPF
RSVP-TE/LDP
IS-IS-TE
BGP
End
Station
IS-IS-TE
CN
End Station
Call Controller
Policy Server
Connection Controller
Call Control
Policy Control
Connection Control
Data Plane
ARGUS
Control
Identify
Analyze
Plan
Track

Multi-Point Monitoring



MultiProbe Correlation

- Look for flows at multiple points
  - Differential analysis
    - One-way delay
    - Loss statistics
  - Path assurance

- Sensor placement provides utility
  - Exploit geospatial nature of observation domain
  - Validate explicit compartmentalizaiton
    - Exterior / Interior verification

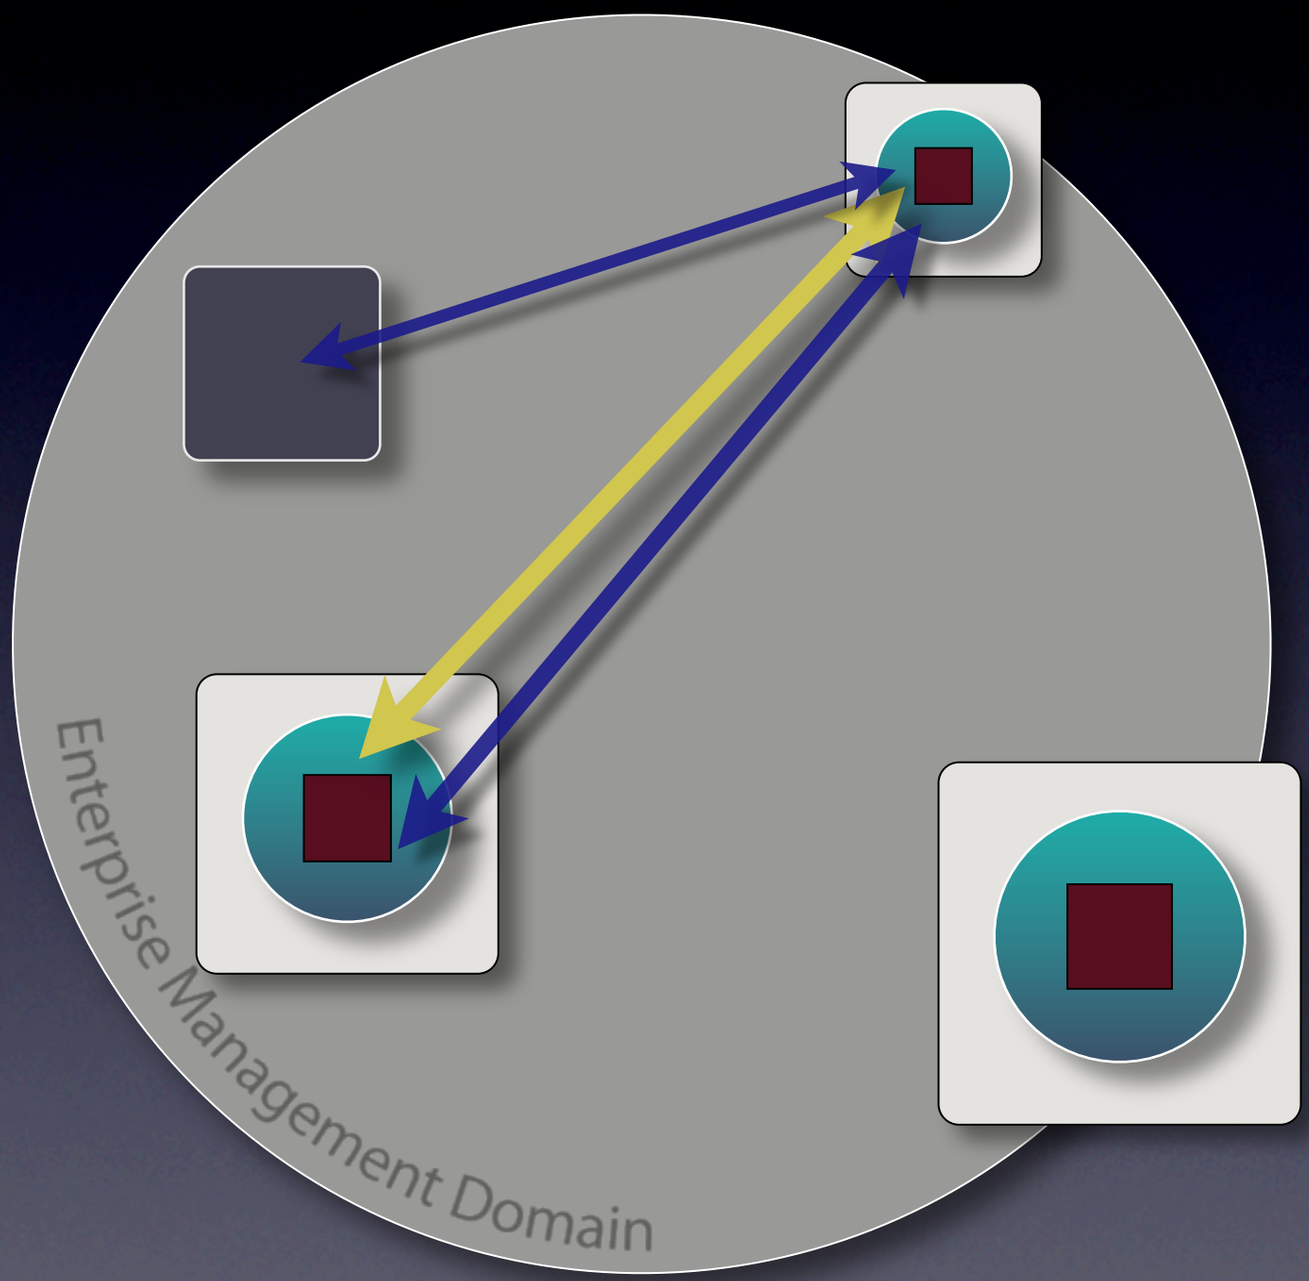




GeoSpatial Situational Awareness System

Mixed Black-box White-box Approach

Local Area Network Implementation



White/Visible Node

Black/Non-Visible Node

Comprehensive Flow IS

Flow Data Generation

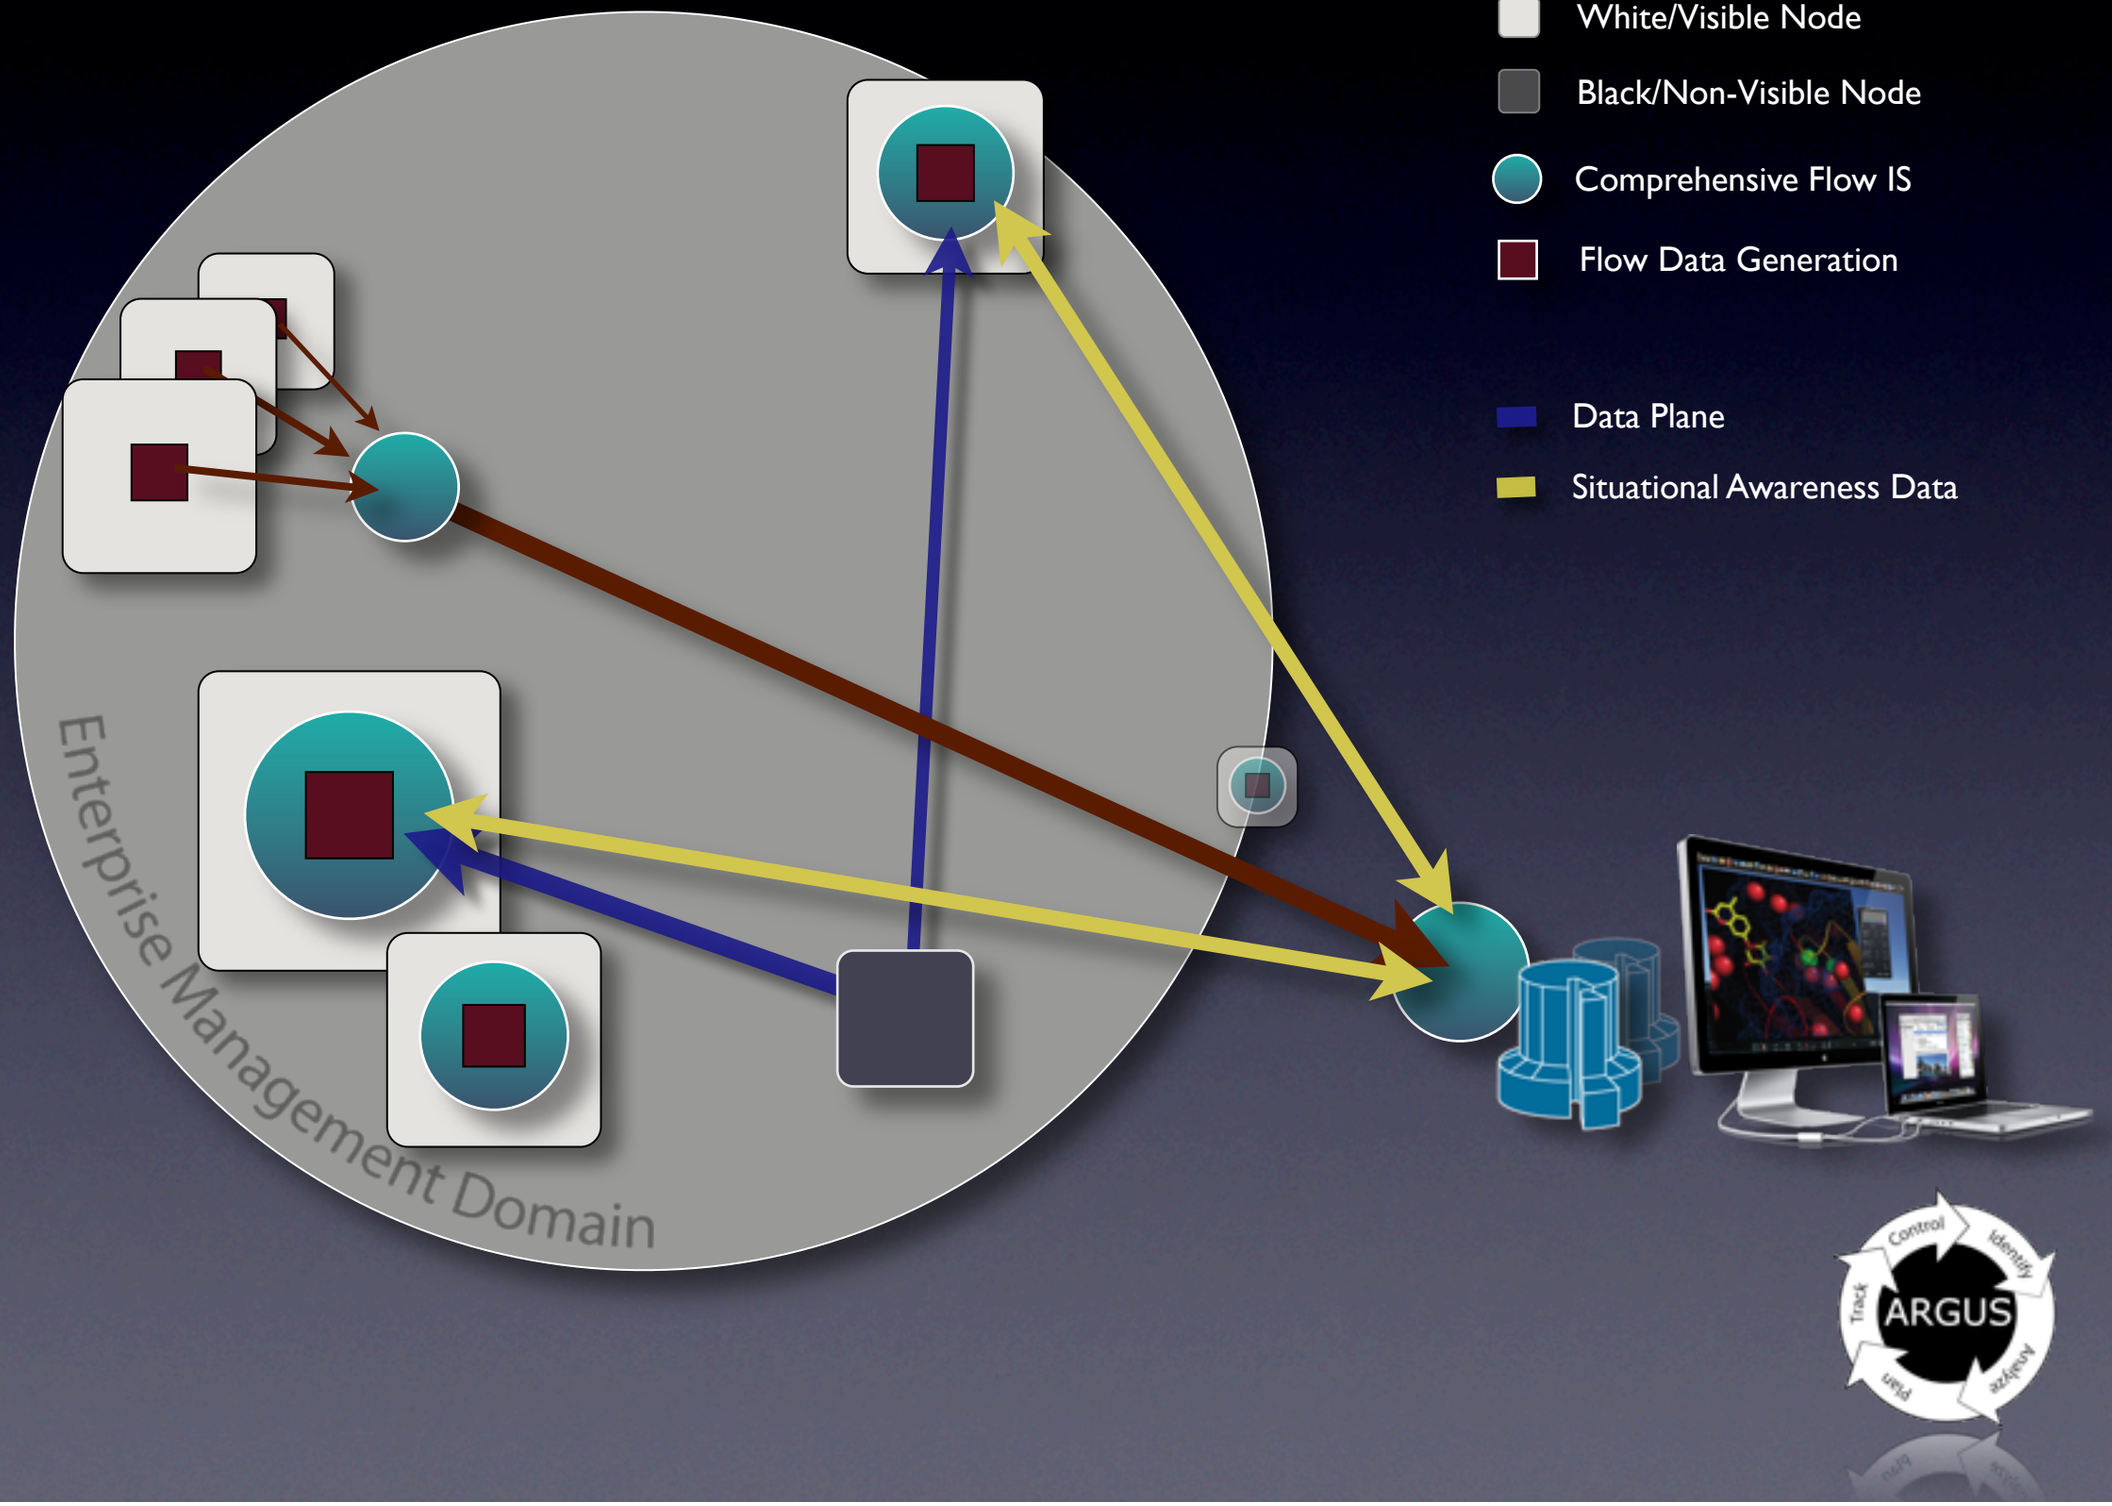
Data Plane

Situational Awareness Data

GeoSpatial Situational Awareness System

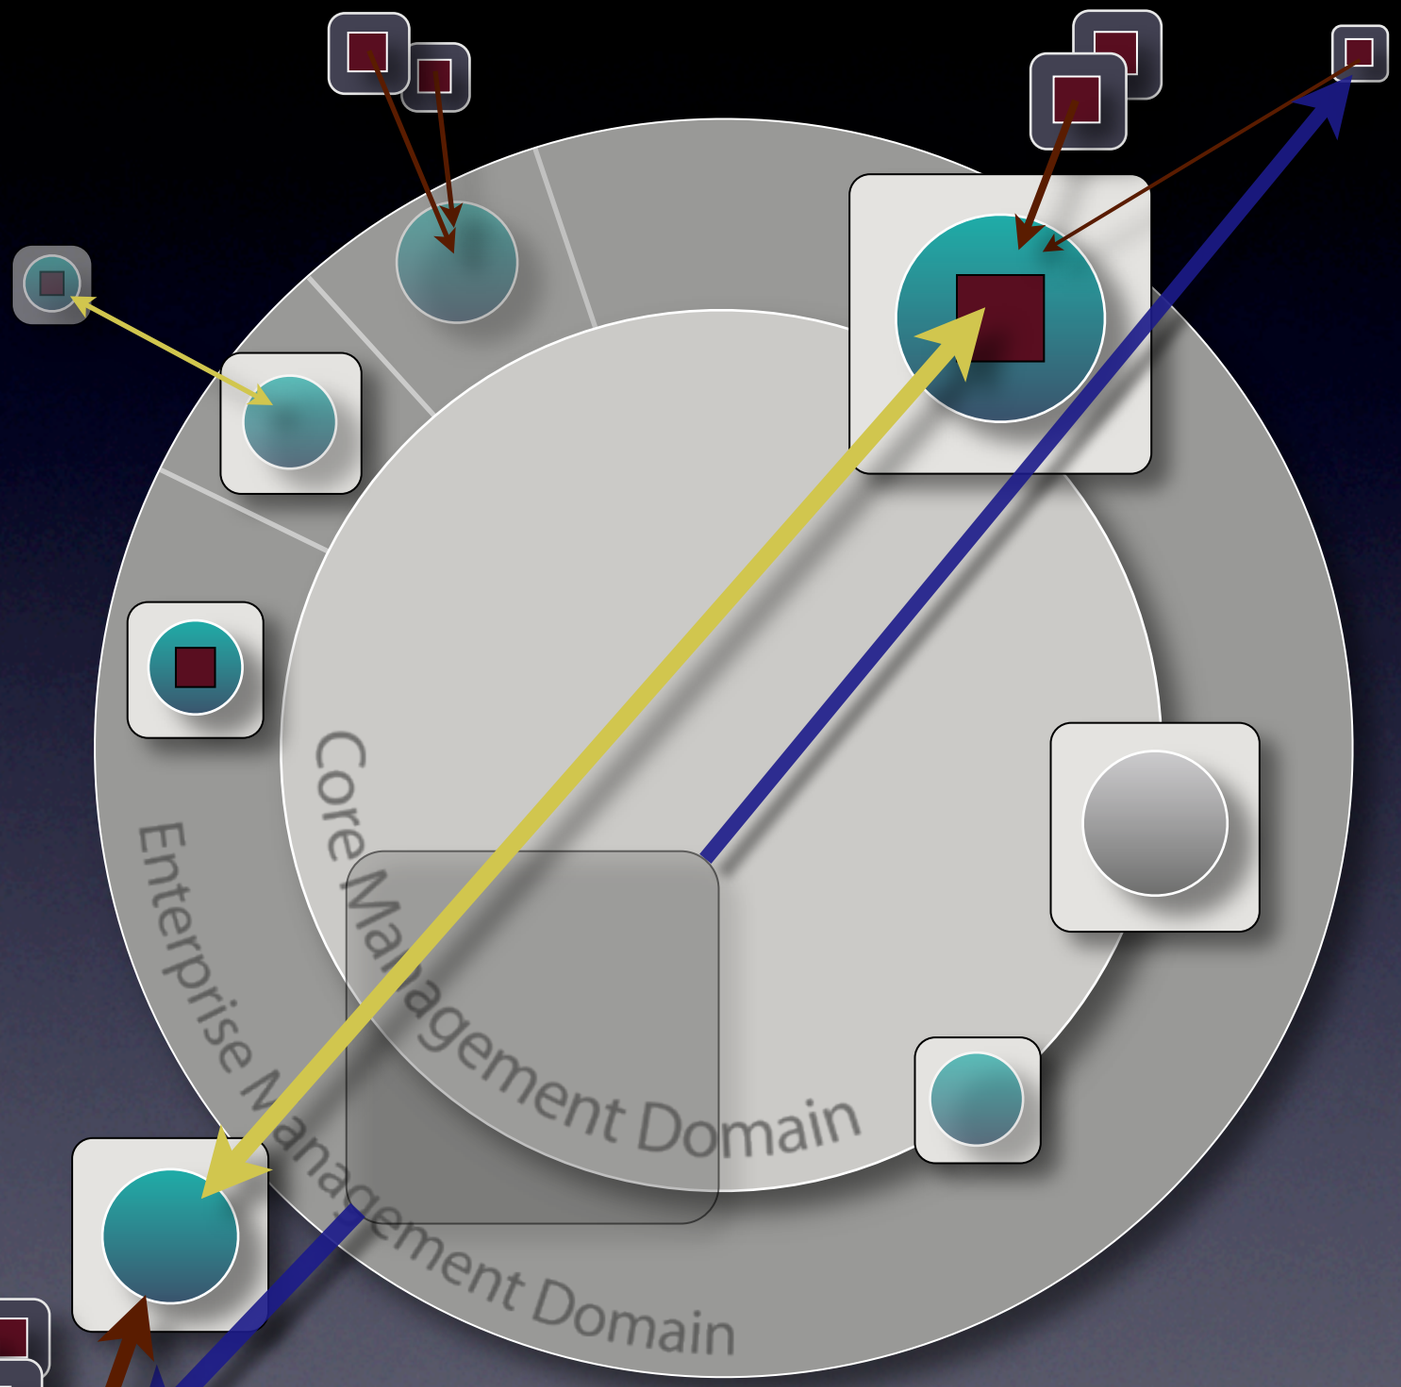
Mixed Black-box White-box Approach

Local Area Network Implementation



GeoSpatial Situational Awareness System

Mixed Black-box White-box Approach



White/Visible Node

Black/Non-Visible Node

Comprehensive Flow IS

Flow Data Generation

Data Plane

Situational Awareness Data

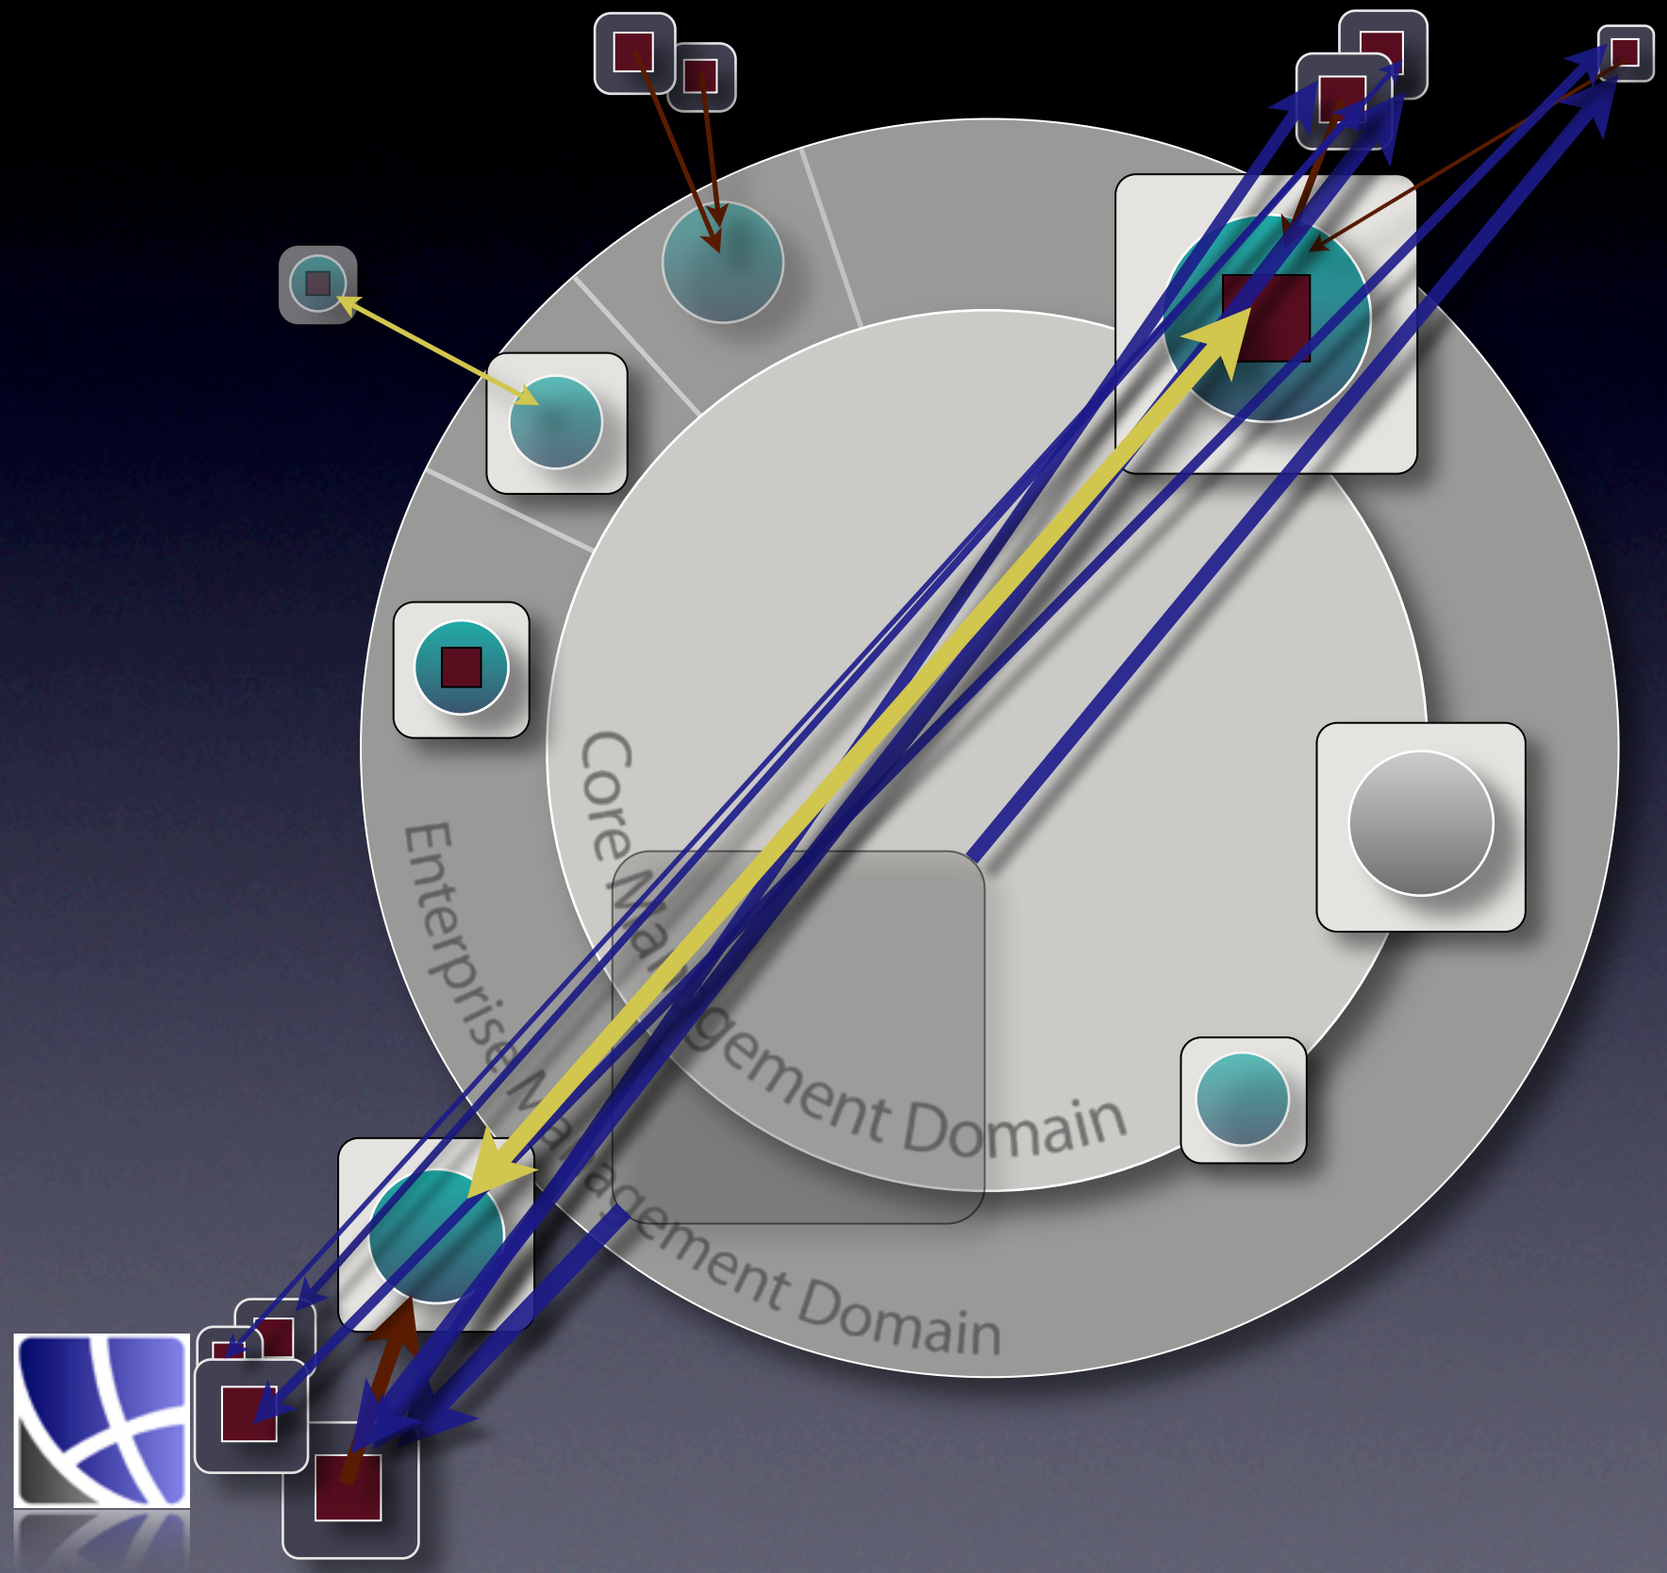






GeoSpatial Situational Awareness System

Mixed Black-box White-box Approach



White/Visible Node

Black/Non-Visible Node

Comprehensive Flow IS

Flow Data Generation

Data Plane

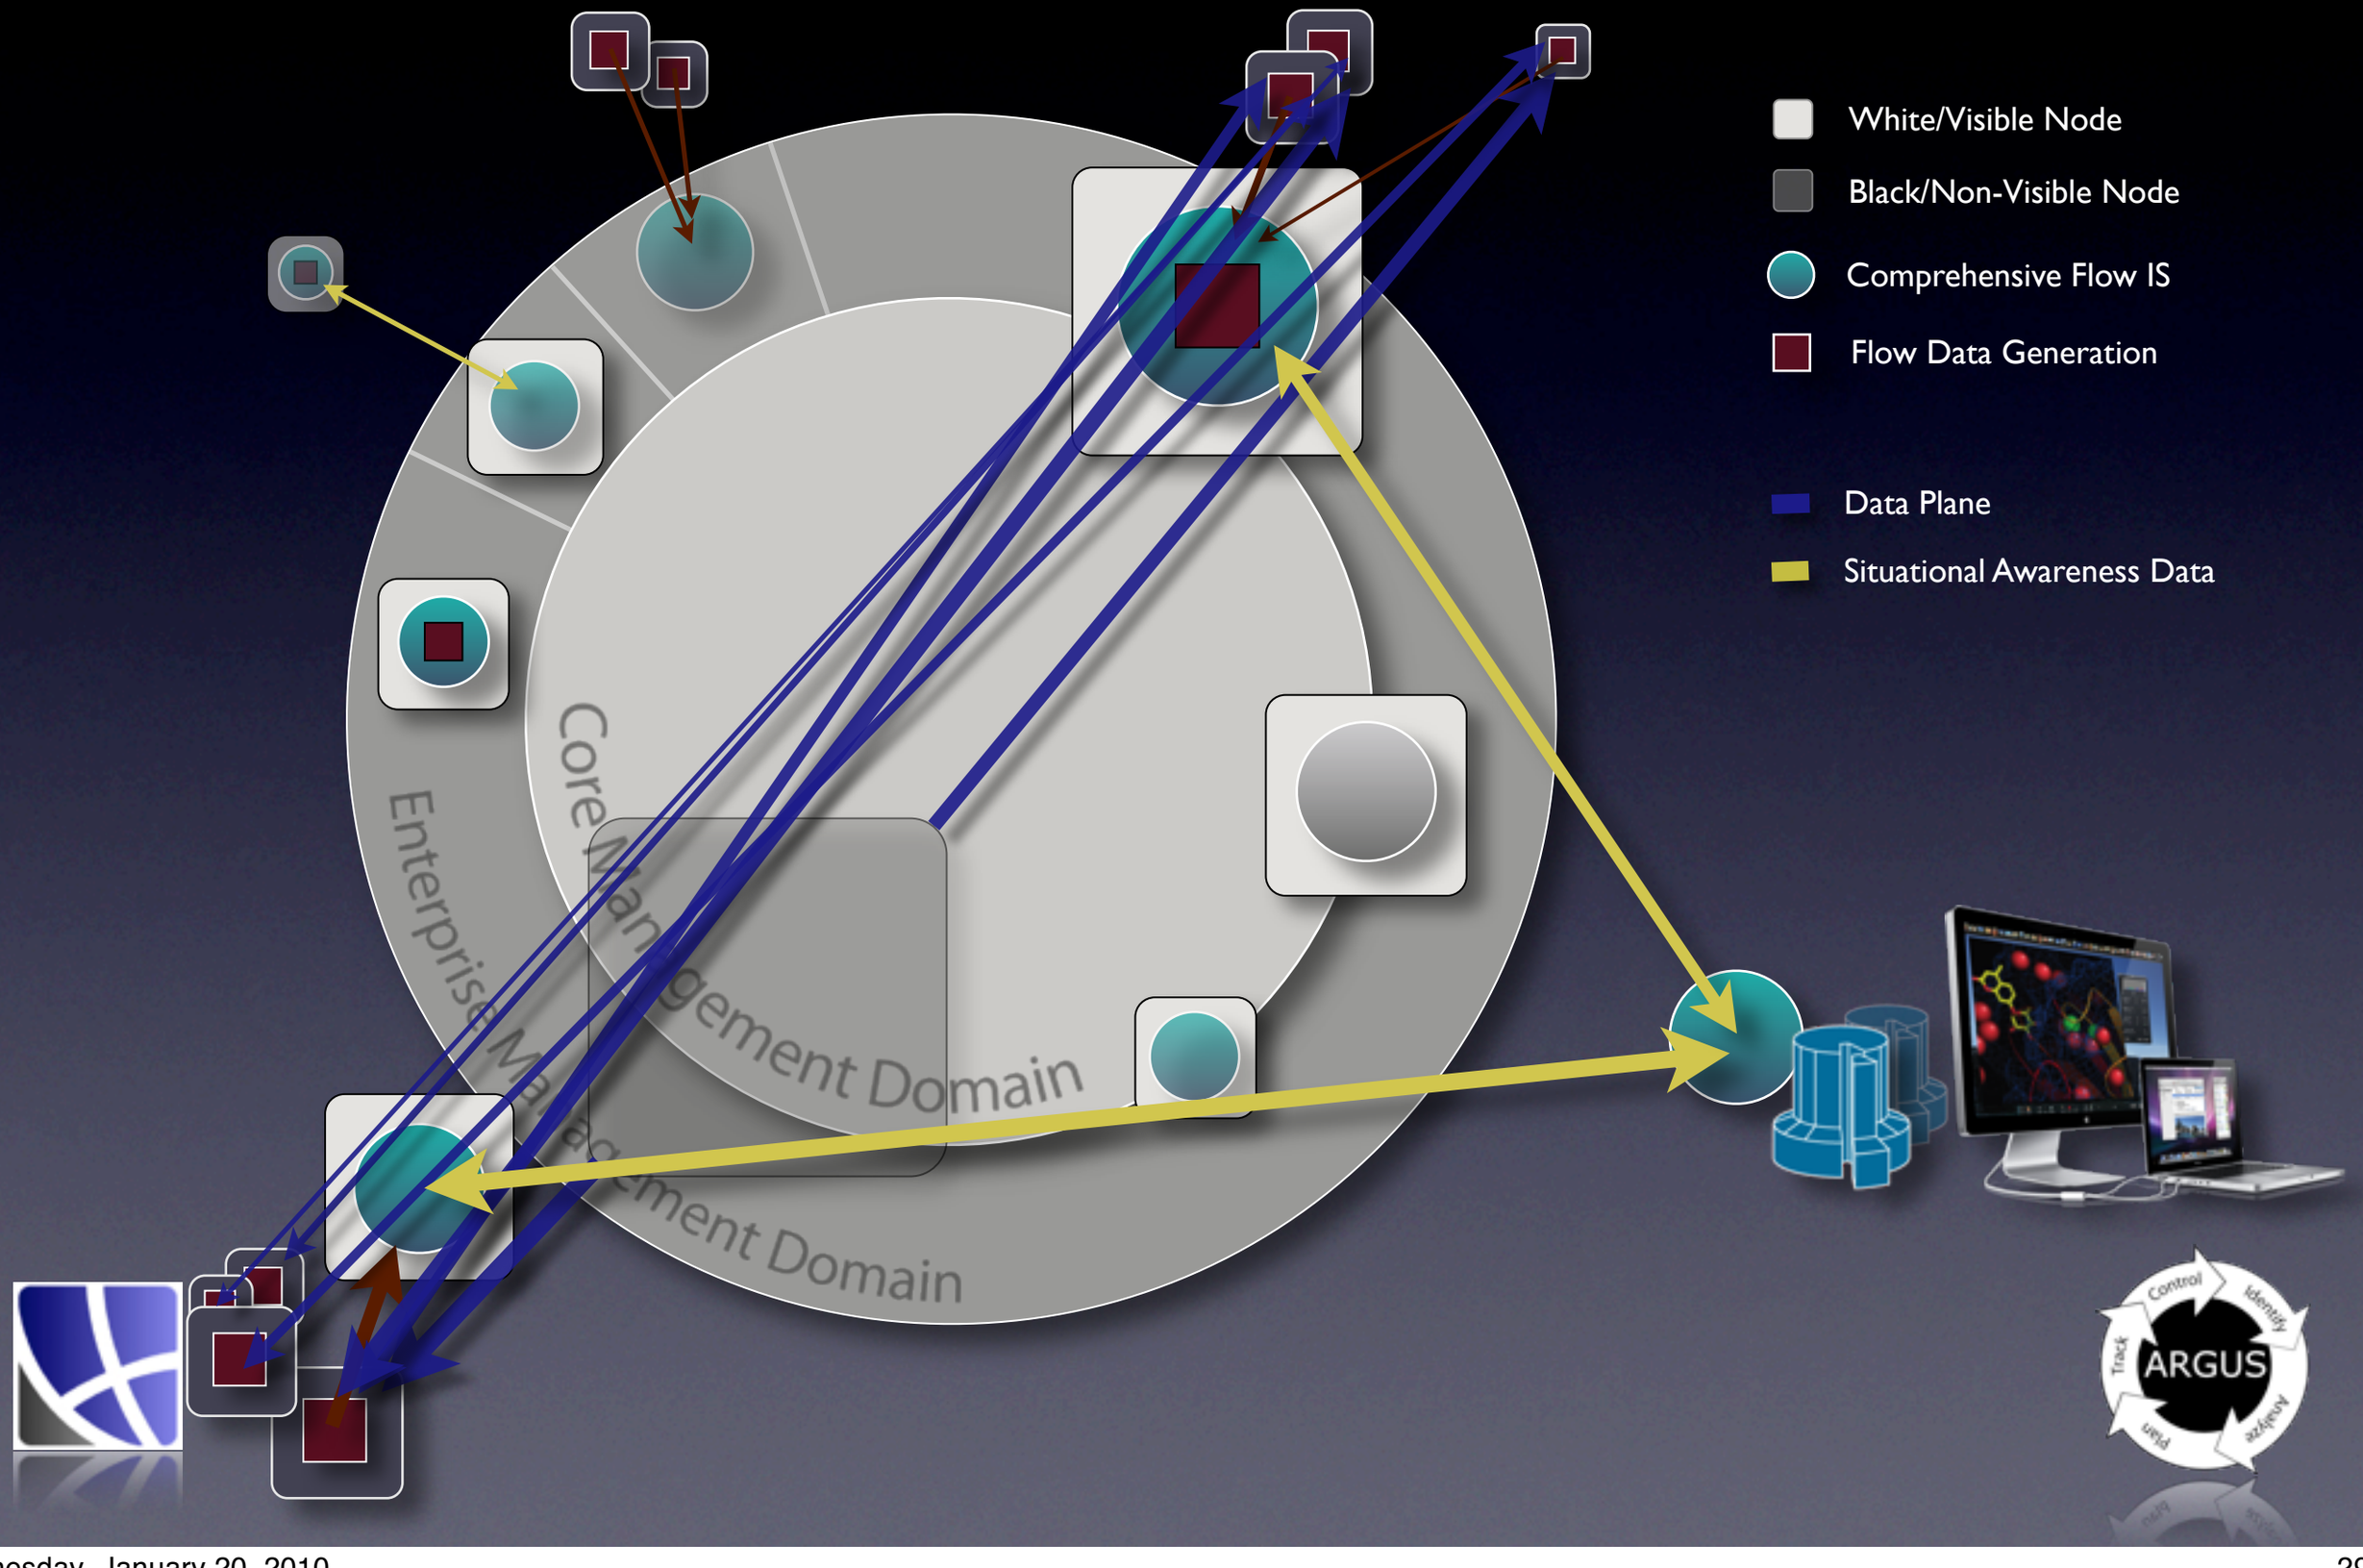
Situational Awareness Data





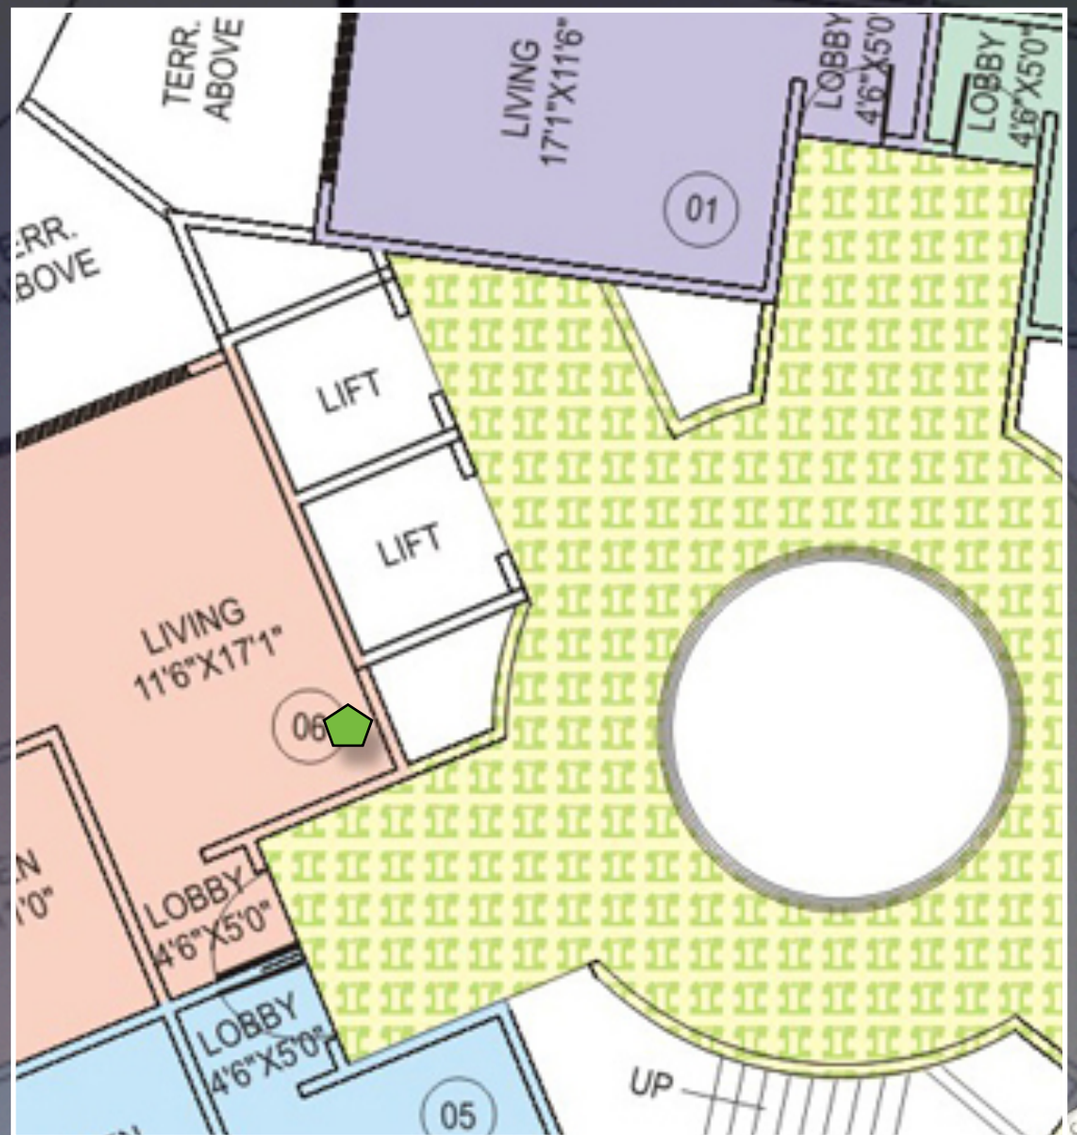
GeoSpatial Situational Awareness System

Mixed Black-box White-box Approach



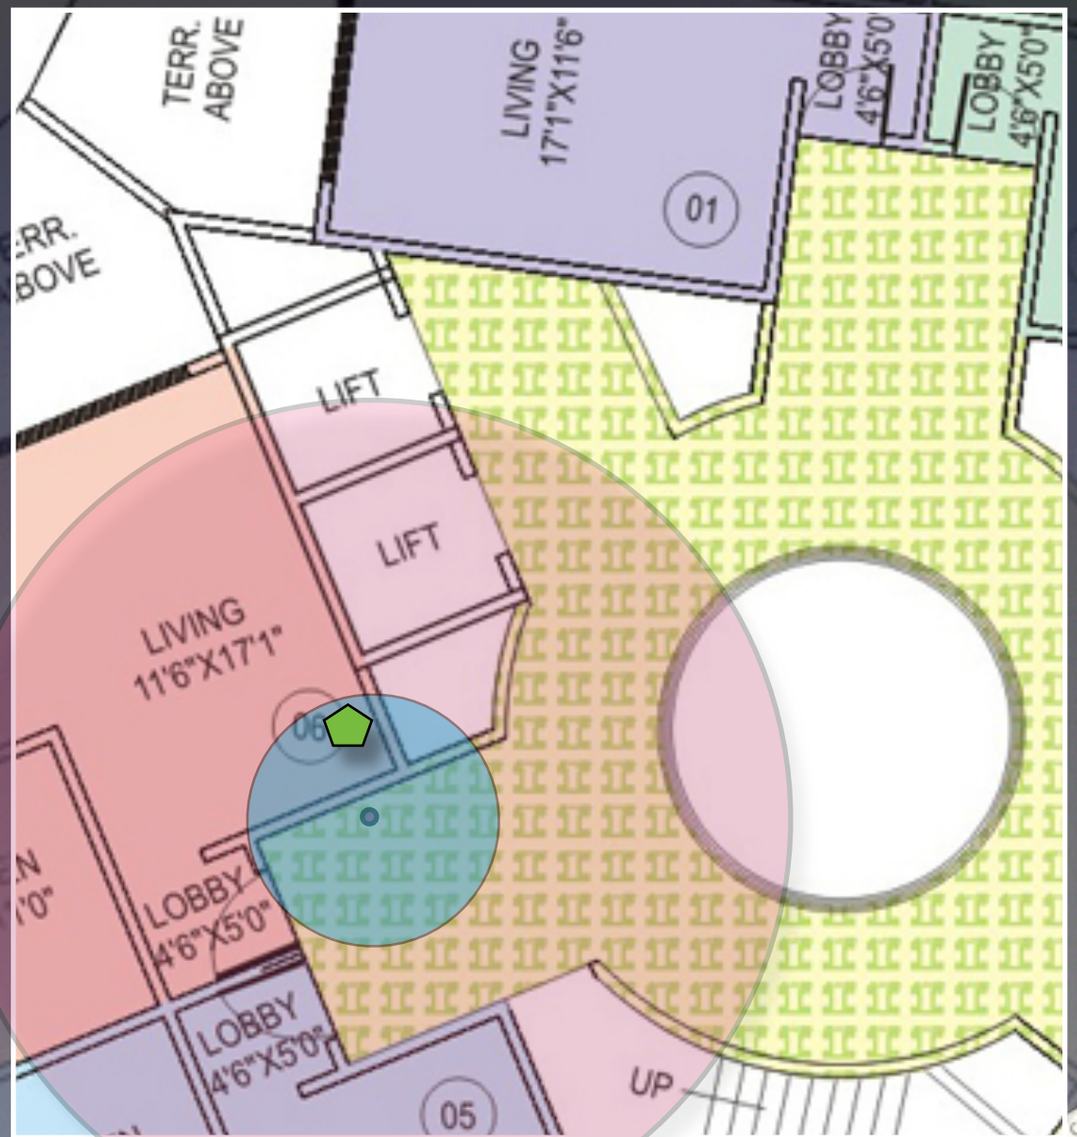
Threat Distance Assessment

802.11 a/b/g/n
Bluetooth



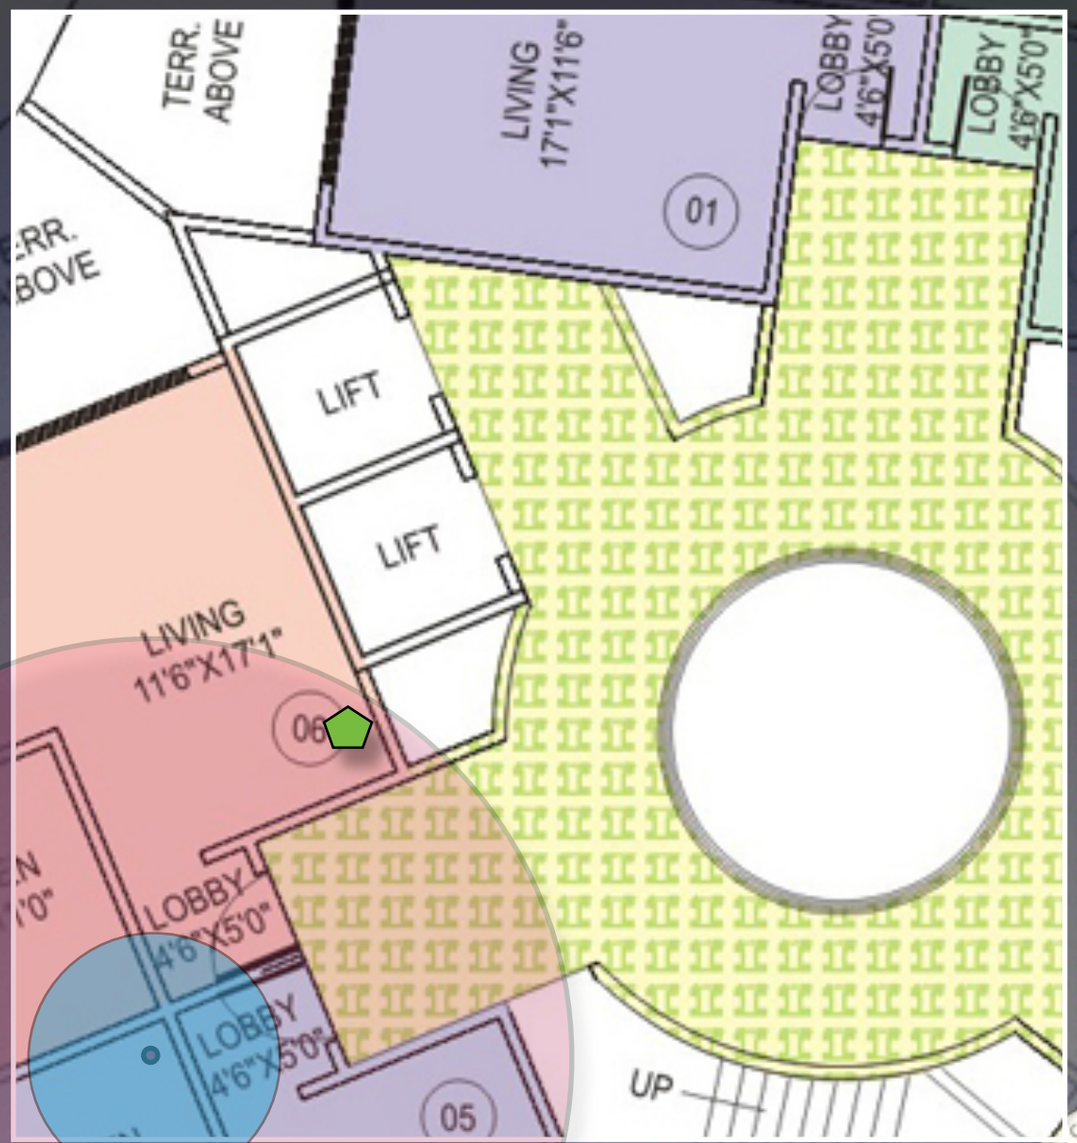
Threat Distance Assessment

- Attack distances can be a matter of inches



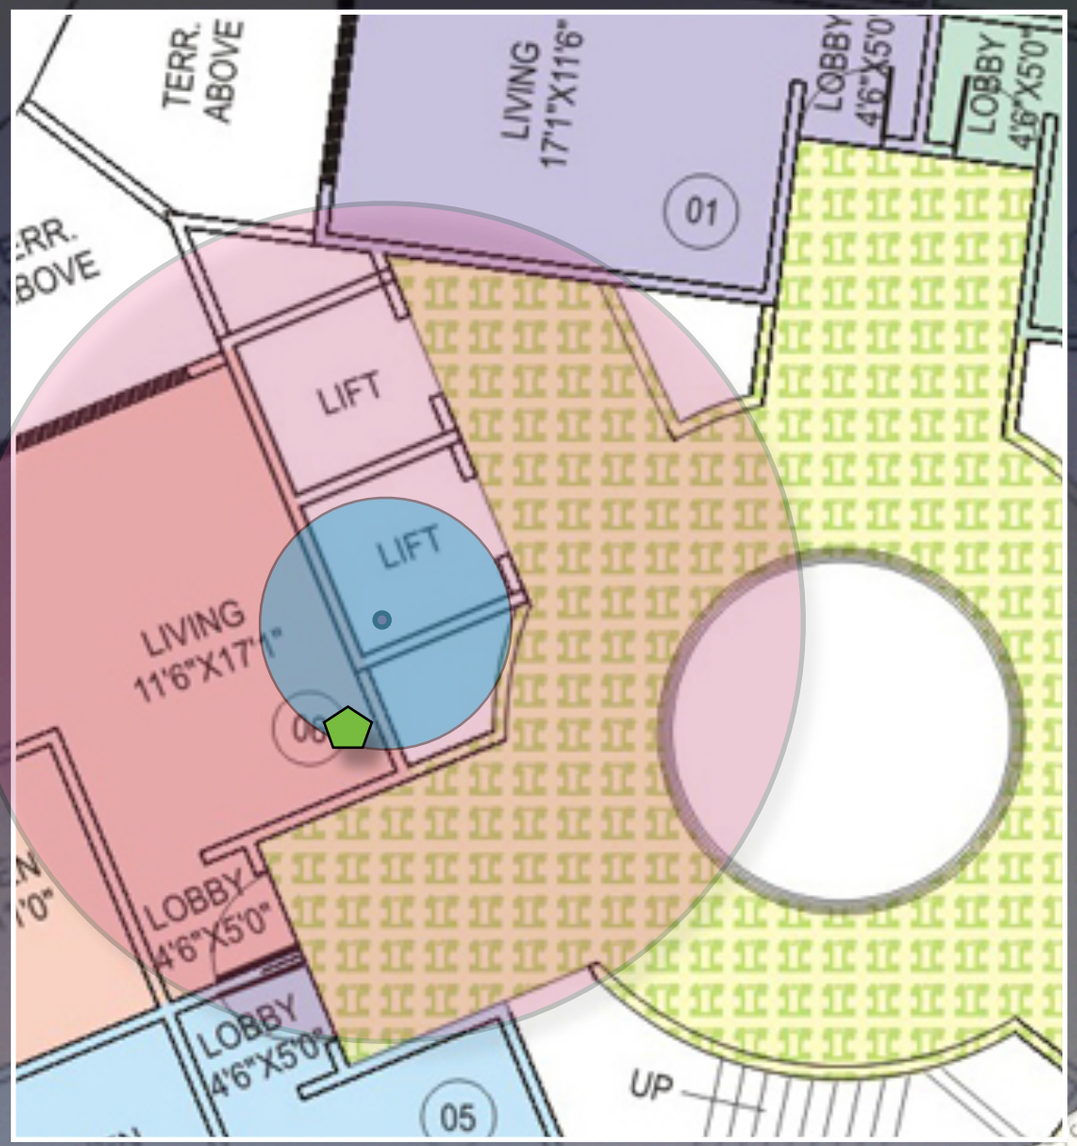
Threat Distance Assessment

- Attack distances can be a matter of inches
- Attack positioning can be rather surprising



Threat Distance Assessment

- Attack distances can be a matter of inches
- Attack positioning can be rather surprising



Now Inches are Important

- How do we get the sensitivity?
  - Flow Data Generation in End Systems
    - Very dynamic situation
    - Commercial GPS is not accurate enough
      - Predominately an indoor problem
    - Unfortunately, the end system may not be cooperative
  - Flow Data Generation in Wireless Routers
    - Need to perform triangulation for 3D space
    - May require 4 signal strength values for complete discrimination.
- Fuse signal strength information
  - Average/max/min receive signal strength vector over life of the flow.
  - Provide result of triangulation computations
  - Correlate with signal strength of 802.11 abgn beacon information





Conclusions

- Network flow data is a/the primary forensics and security data source for many large scale security systems.
- Fusion of non-packet derived information and flow data can be a very very nice thing.
- Some data requirements do exist
  - Time synchronization is critical
  - Semantic similarity
  - Relational algebraic constraints exist
- Flow data / flow data fusion is a big deal.
- Lots of work to be done in this area!!!



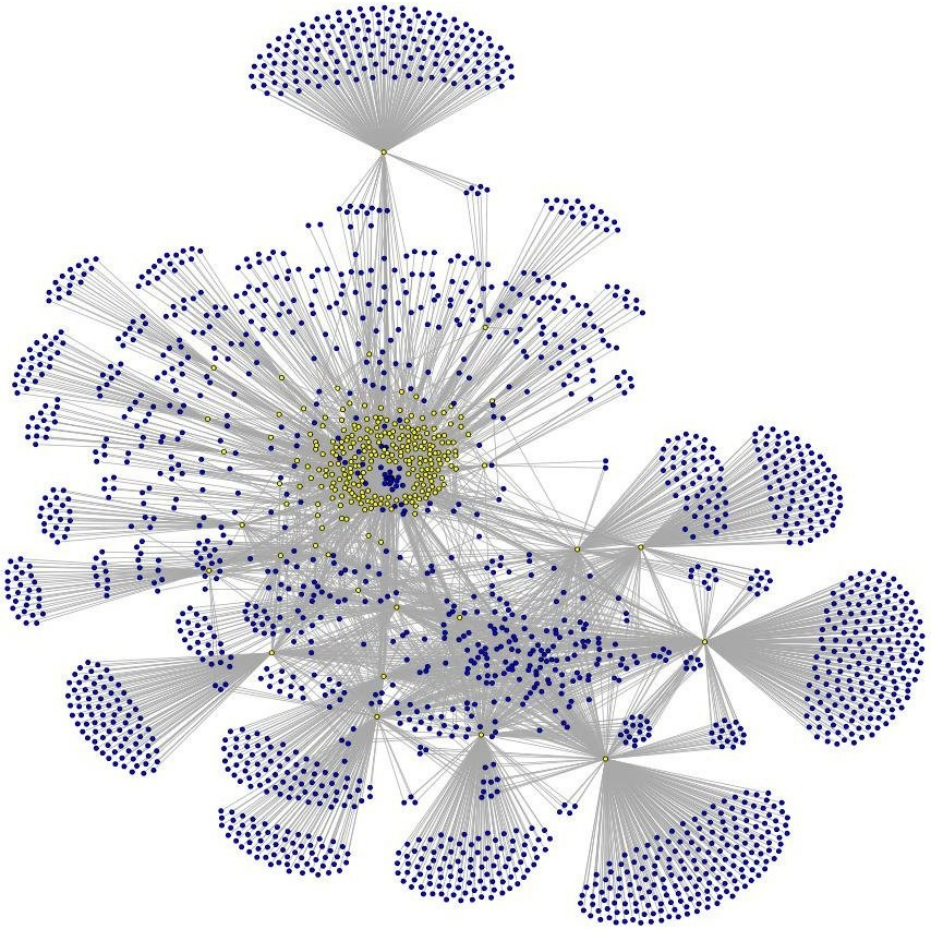


Realtime Change Detection & Automatic Network Response

Alex Brugh (presenter)
Mike Fisk, Josh Neil, Paul Ferrell, Scott Miller, Danny Quist
Los Alamos National Laboratory
Advanced Computing Solutions
LA-UR 09-08132

Outline

- Use of Flow in Change Detection
  - Current Methods
  - Areas of research
- Response
  - Automated framework
  - Methods



Change Detection

- We'd like to detect misuse of the network
- Most researchers have attempted to train systems to detect malicious traffic
  - But there is no realistic, representative training data
  - Over-fitting the training data is always a risk
- Our approach: focus on detecting change in the status quo
  - Change can be detected without training data
  - Statistically significant changes tend to be interesting to security and/or operations
  - Finding only future attacks isn't perfect but isn't bad

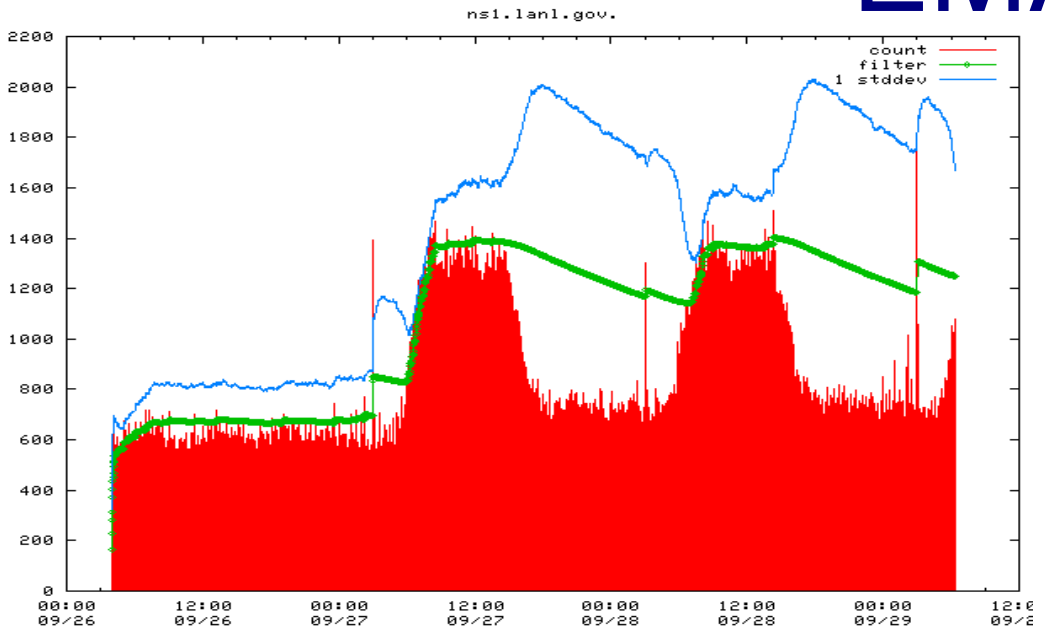
Flows

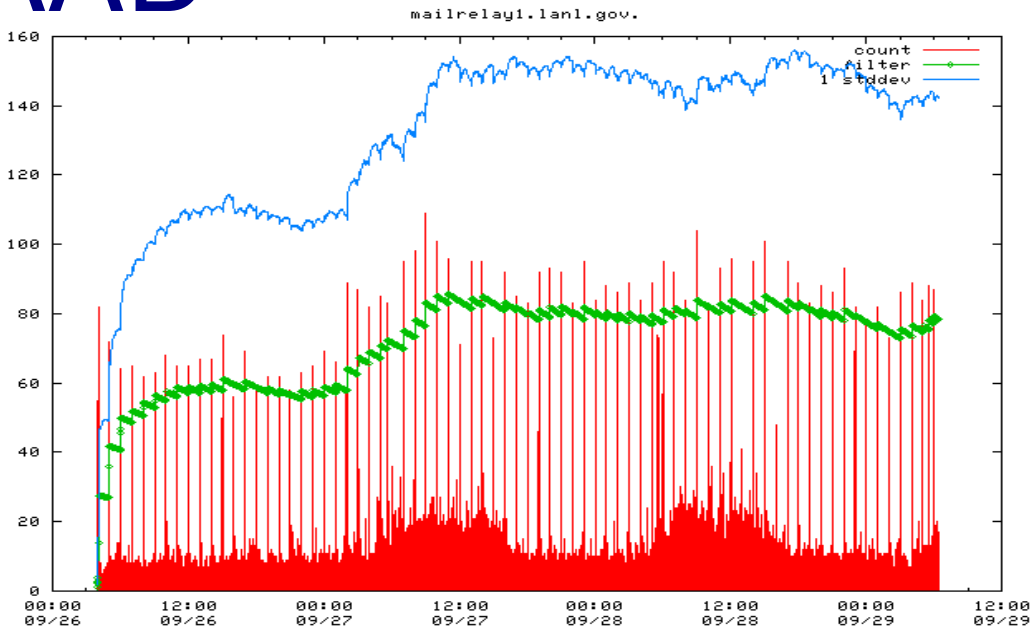
- Flows are binned by time
  - e.g., X counts per minute
- Many possible features
  - Connections from 1 host to many, from many to 1
  - Protocol distribution (e.g, TCP/UDP)
  - Ports
  - Bytes or packets sent/received
  - Graphs

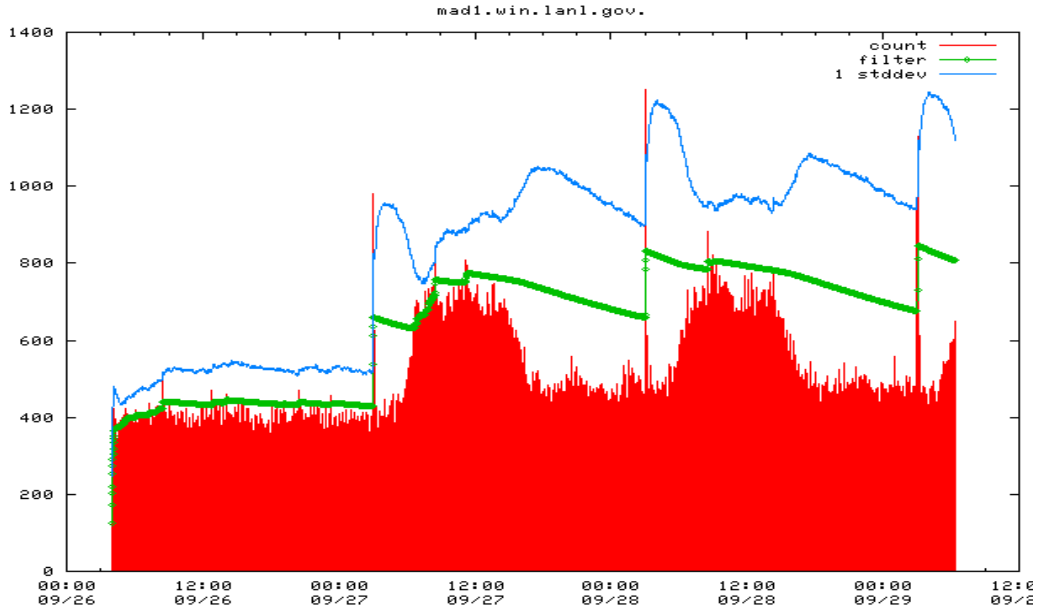
EMAAD

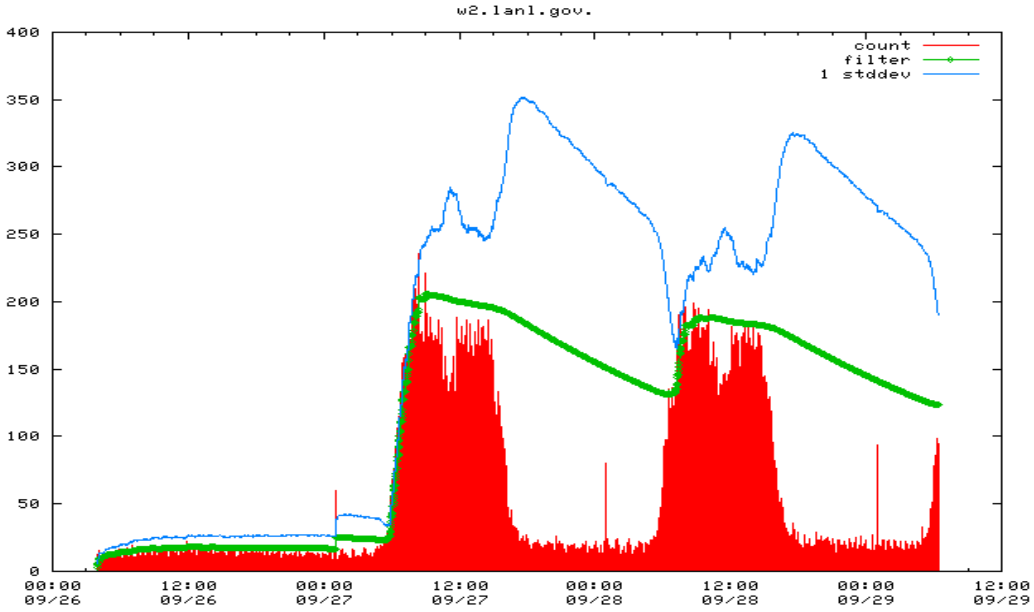
- Exponential Moving Average Anomaly Detector
- Models number of unique connections originating from a given host per unit time
  - IP to IP, ignores ports
- How it works
  - Assumes Gaussian model
  - Asymmetric moving average
    - Adapt quickly to increases
    - Decay slowly with decreases
  - Model built for each host (little state)

EMAAD


ns1.lanl.gov.
count
filter
1 stddev
2200
2000
1800
1600
1400
1200
1000
800
600
400
200
0
00:00 09/26
12:00 09/26
00:00 09/27
12:00 09/27
00:00 09/28
12:00 09/28
00:00 09/29


mailrelay1.lanl.gov.
count
filter
1 stddev
160
140
120
100
80
60
40
20
0
00:00 09/26
12:00 09/26
00:00 09/27
12:00 09/27
00:00 09/28
12:00 09/28
00:00 09/29
12:00 09/29


mad1.win.lanl.gov.
count
filter
1 stddev
1400
1200
1000
800
600
400
200
0
00:00 09/26
12:00 09/26
00:00 09/27
12:00 09/27
00:00 09/28
12:00 09/28
00:00 09/29


w2.lanl.gov.
count
filter
1 stddev
400
350
300
250
200
150
100
50
0
00:00 09/26
12:00 09/26
00:00 09/27
12:00 09/27
00:00 09/28
12:00 09/28
00:00 09/29
12:00 09/29

EMAAD

- Sub-minute response times
  - Continually recalculates as a bin updates
  - Handles 10s of thousands without
- Alarm causes
  - Host discovery (nmap, ping sweeps, etc)
  - Misconfiguration
  - HP printer administration software

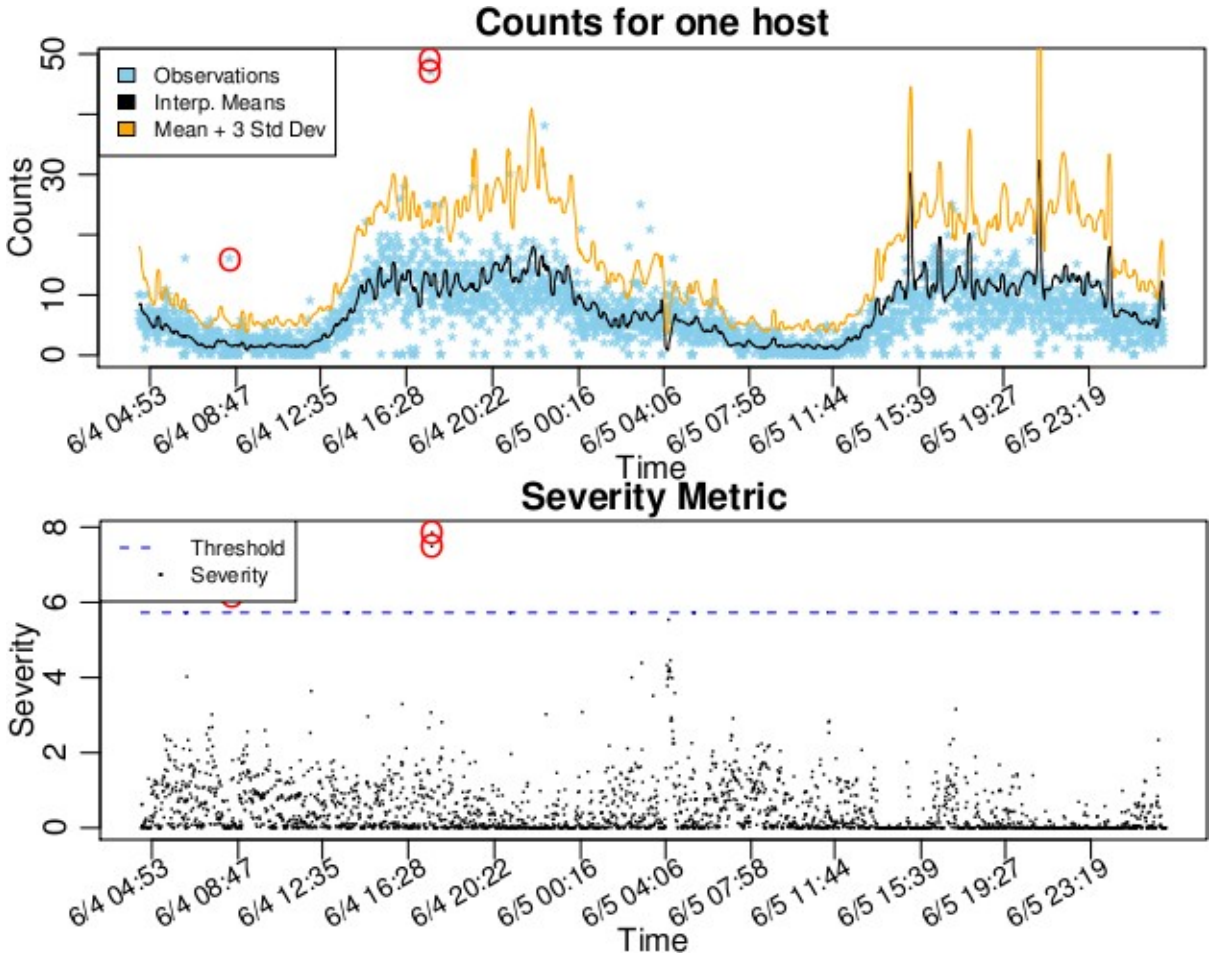
EMAAD Mark II

- We applied research by Lambert and Liu of Bell Labs [1]
  - Focuses on network degradation
  - Assumes a negative binomial model for changes
  - Also an EWMA
  - Quadratic mean interpolation
  - Has a normalized “Severity” metric
  - Outlier handling
    - Replace with random 'abnormal' draw
  - Able to detect gradual changes

EMAAD Mark II

- Modified to detect increases in activity
- Models 10 minute intervals independently
- Different cycles exist throughout the week
  - Monday-Thursday != Friday != Saturday,Sunday
  - More state than EMAAD Mk I
- Needs to train on a several cycles of data
  - Can't detect an abnormal Monday without seeing several normal ones first
  - Can use historical data, of which we have plenty

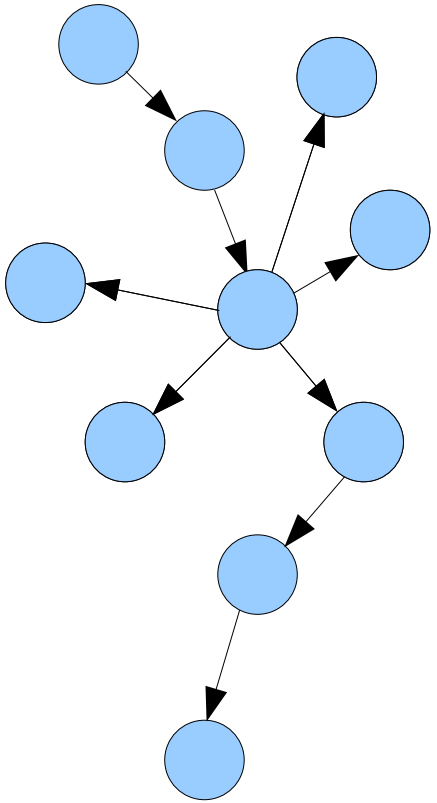
EMAAD Mark II



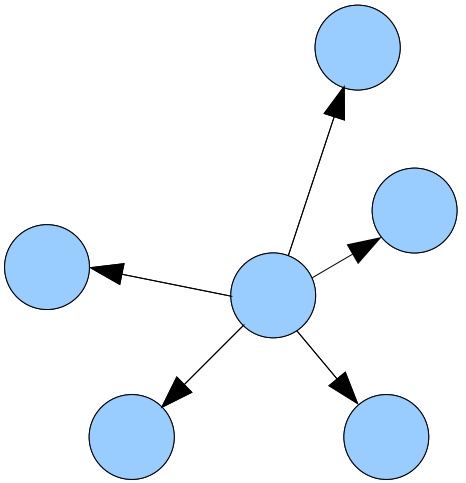
Next Generation Change Detection

- IP graphs based on flow records
  - Allows examination of subgraphs in a network
  - Looks for anomalies in a “neighborhood”
  - Assumes hackers behave locally
- Handles daily periods in the data
- Hidden Markov model used on edges of the graph
  - Distinct states of activity in flows
  - Probabilities in changing between states varies across time

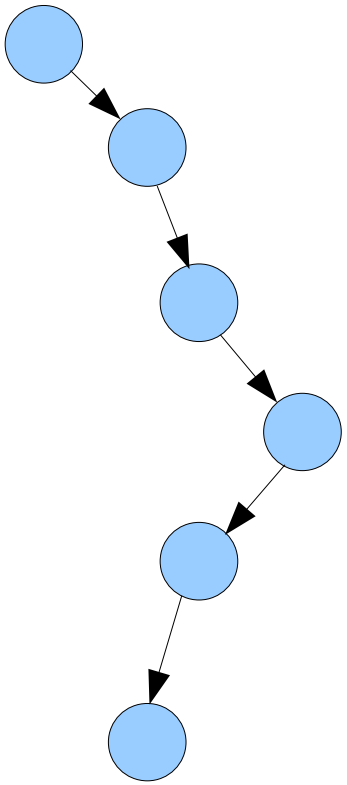
Next Generation Change Detection



Sample graph



Scan



Directed Attack

Response System

- Change detected, now what
  - Send e-mail to an analyst
  - Log this event somewhere
  - Quarantine a host
  - A combination of actions
- Not all events are equal
  - The responses shouldn't be either
  - A dynamic system needed

FRNSE

- Framework for Responding to Network Security Events
- “Replacing lowest level analyst work with a script”
- Takes alerts from Agents
  - Formated in XML
  - Transmitted securely
  - Alerts are queued and retried if unsuccessfully sent
- Uses rules to implement policy
- Can respond in a configurable way
- Exemptions can be specified
  - We still need Name Servers, pwn3d or not

FRNSE

- Python API for easy rule creation and response interfacing
- It has web interface for configuration
  - Rules
  - Exceptions
  - Data types
  - Viewing alerts

Agents

- EMAAD
- TippingPoint
- Snort
- AirDefense
- MassAV

FRNSE Production Results

- In 2007:
  - 919,737 alerts handled by FRNSE
  - 283,192 automatic firewall blocks
  - 2,293 analyst tickets generated
  - 179 automatic internal host quarantine events

Of 919,737 alerts generated in 2007, 99.75% were directly addressed by technical implementation of policy and required no analyst intervention.

Response Arsenal

- Internal Host Quarantine
- Perimeter block on external host
- Blackhole DNS name
- Open a Ticket for analysts
- Send email or page

Host-Based Quarantine

- Get a host off the network quickly
- Impact as few collateral systems as possible
- Handle the dynamic nature of a large network
  - Laptops move
  - Switches, ugh
- Deal with resourceful users
  - Jump ports
  - Change IP

Host-Based Quarantine

- Three lists need to be maintained by hand
  - The MAC vendor codes of switches
  - The IP addresses of routers
  - Every SNMP community string ever
- Harvest ARP tables from routers
- Switches found in ARP data
- Harvest forwarding tables from switches
- Crunch data
- Store in a database

Host-Based Quarantine

- Host location mapped and stored by MAC
- MAC/IP mappings stored in a separate table
  - MACs might map to multiple IPs, also stored
  - IPs could have multiple MACs, ugh
- Complete history* of every host's location
  - Useful in forensics/investigations
- Web interface
  - Can view a switch
  - Search by MAC or IP

Host-Based Quarantine

- Our current environment
  - Aprox. 1600 Switches, multiple vendors, vintages, settings
  - 20 Routers
- Runs every half hour
  - Parallelism in device harvesting and crunching needed to achieve this
  - Port jumpers handled after every run
  - Aprox. 40,000 database records generated per run
    - Aprox 2 million a day.
- Time from request to quarantine between 10 and 30 seconds depending on switch

Conclusion

- Flows used in change detection
- Detected changes feed into larger response system
- Of our responses we're pretty proud of Host-Based Quarantine

Questions

Alex Brugh

abrugh@lanl.gov

[1]D. Lambert and C. Liu. Adaptive Thresholds: Monitoring Streams of Network Counts Online. Journal of the American Statistical Association, 101(473):78–88, 2006.

High-Throughput Real-Time Network Flow Visualization

Daniel Best

Research Scientist | Information Analytics

daniel.best@pnl.gov

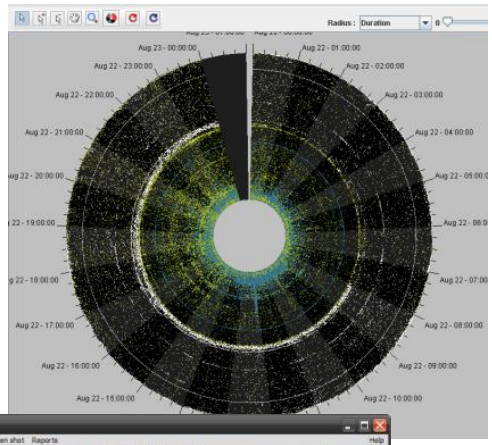
Douglas Love, Shawn Bohn, William Pike



Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Tools and a Pipeline to Provide Defense in Depth





- Traffic Circle
  - Visualization for situational awareness
- Correlation Layers for Information Query and Exploration (CLIQUE)
  - Network behavior visualization using LiveRac interface
- Middleware for Data-Intensive Computing (MeDiCi)
  - Data pipeline

Pacific Northwest
NATIONAL LABORATORY

Motivation

- Improve upon current analysis capabilities
  - Provide a mechanism for multiple tools to feed off the same data
  - Move away from batch processing of flow data
  - Support both forensic and real-time monitoring capabilities
- Implement a streaming analytics tool set
  - Handle large volumes of flow data (millions to billions) per view
  - Help analysts gain situational understanding of current state of a network
- Provide engaging flow visualizations
  - Visualization of both raw data and aggregates
  - Automated identification of off normal conditions

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

What our users want

- They want to know what normal looks like – from an individual host, to a group, to an enterprise
- They want ways to overcome limitations in analyzing raw transactions
  - Lots of data (billions of transactions/day)
  - Lots of unique actors (IPv6: 6.67 * 1027 IP addresses per square meter on Earth)
  - Lots of noise
- If they know what they're looking for, they can build a signature to detect it. **But what's in the data that they don't already know to look for?**
  - Need to link data reduction techniques with exploratory analysis interfaces
- They want to know where to focus!

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Middleware for Data-Intensive Computing (MeDICi)

- Provides a high-throughput data communication and processing pipeline
  - Creates the substrate for real-time information sharing
- Mechanism to hand information to multiple tools
  - Multiple tools “subscribe” to MeDICi data, so that tools can be combined for defense in depth
- Capable of streaming data transformations
  - Handles data changes needed by a client prior to being transmitted to that client
    - Offloads computational cycles to the MeDICi server

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

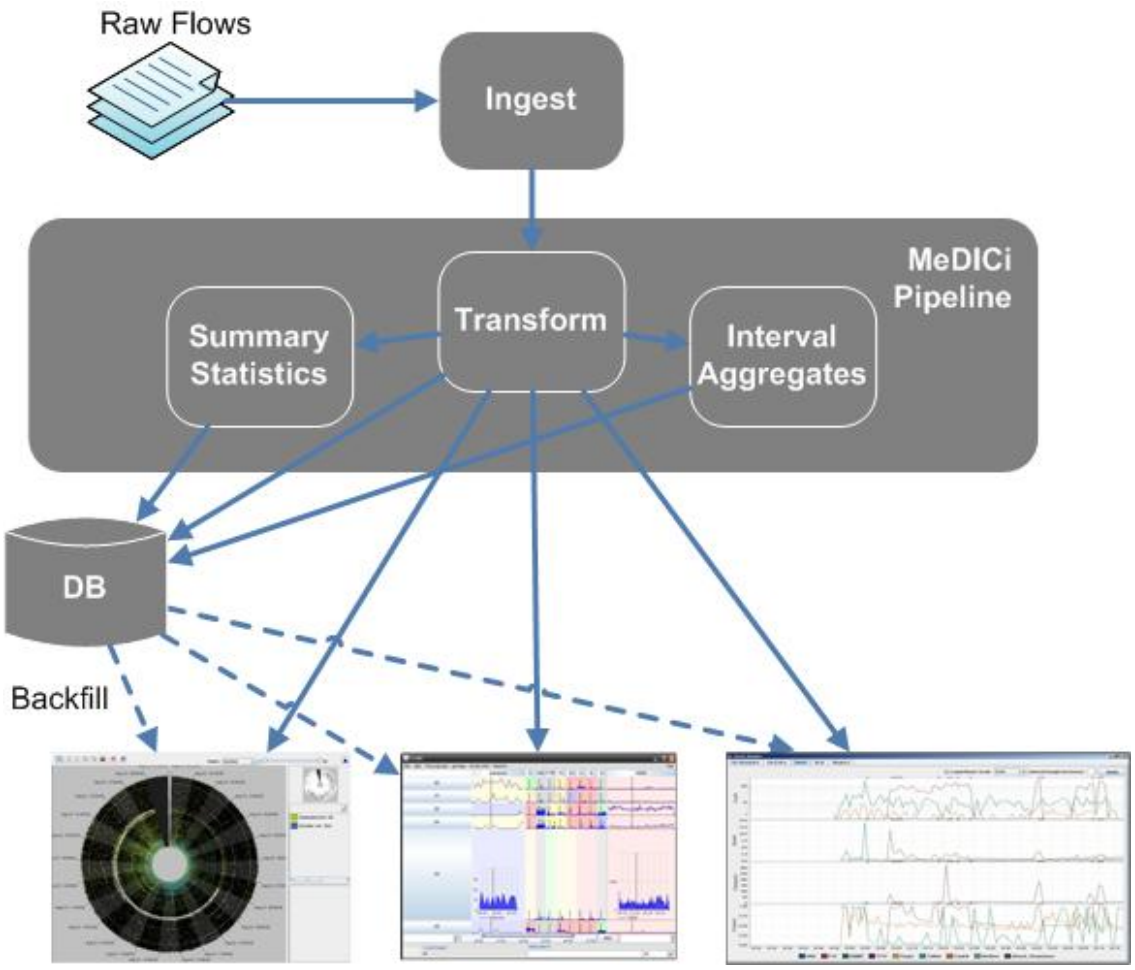
Using MeDICi for real-time analytics

- Components can be created by different developers using various languages
  - Traffic Circle and CLIQUE use Java, but components from other languages can be wrapped easily
- Information is passed between components using a producer / listener mechanism
  - Apache ActiveMQ message broker
  - JMS messaging standard
- Components are chained together to create a pipeline
  - Aggregates
  - Summary Statistics
  - Others

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

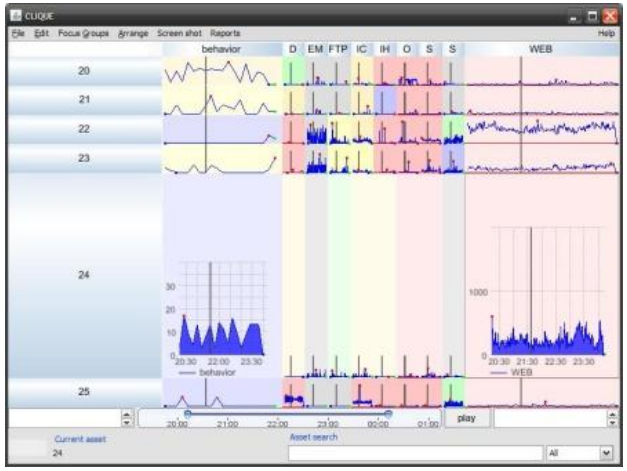
A sample MeDICi pipeline



Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle *Since 1965*

CLIQUE

- Produces behavioral summaries based on a live sensor feed
  - A behavior is a model of predicted activity based on past activity
  - CLIQUE helps visualize the deviation of an actor from its predicted activity
  - Working at the level of behaviors helps cope with large data volumes
  - Display “walks” along with incoming flow information to show current state
- Helps highlight trends and patterns in high-volume flows
  - Capability to compare behavior deviations with category activity



Pacific Northwest
NATIONAL LABORATORY

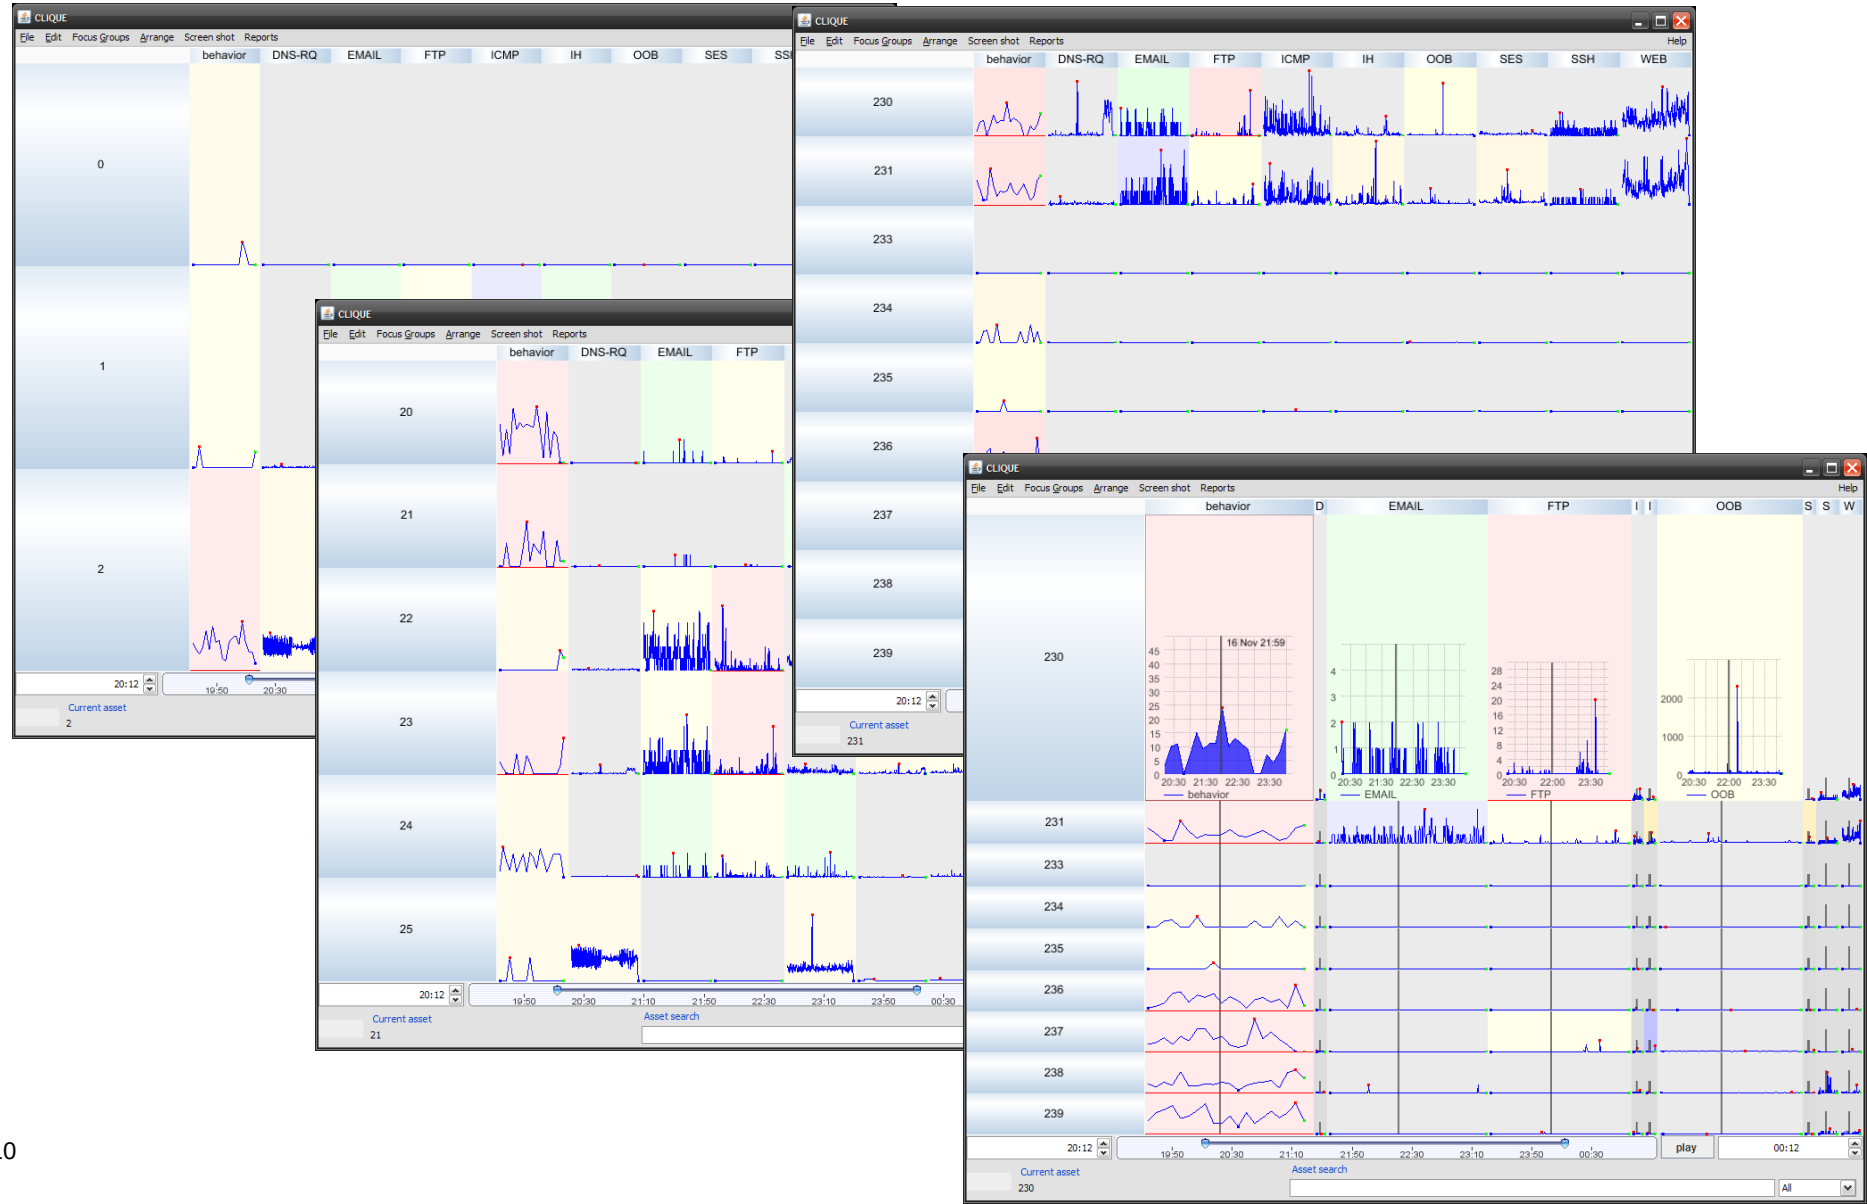
Proudly Operated by Battelle *Since 1965*

CLIQUE behavior analysis

- Utilizes Symbolic Aggregate approXimation (SAX)
  - To deal with the scaling issue in the temporal dimension (scales by summarization/aggregation)
- Creates a SAX representation across all categories for a given actor
- Converts SAX representation to a glyph
  - Produces a language
- Creates matrix of glyphs and temporal sub-segments
- Compares current matrix and historic matrix to yield behavior deviation plot

Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by Battelle *Since 1965*

CLIQUE Walkthrough



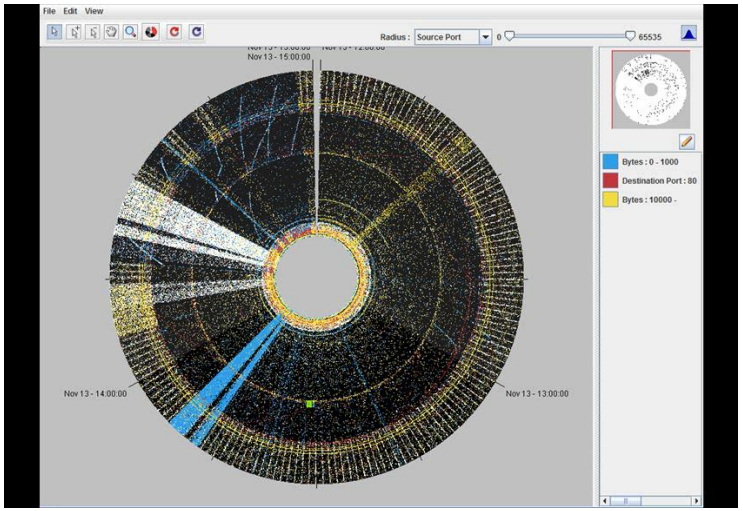
Traffic Circle

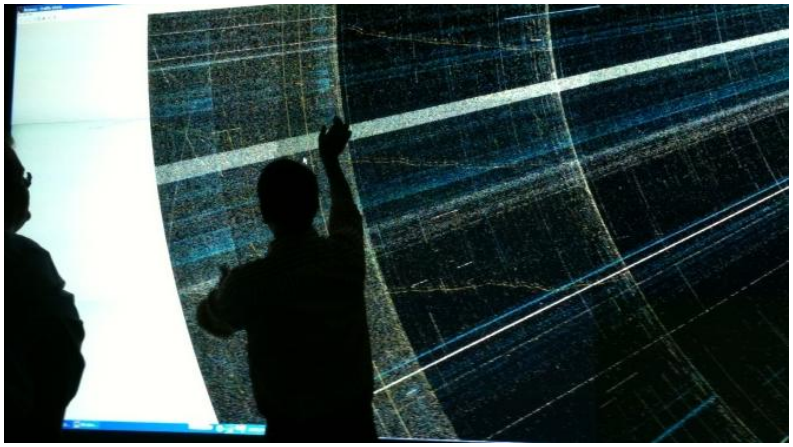
- Interactive visualization of patterns in high volume flow data
  - When you can see more data you can see patterns previously hidden when examining by subsets
- Visualizes raw flow information
  - Drill-through from CLIQUE
  - Manages throughput by utilizing a threaded architecture that distributes query and rendering
- Can operate in forensic mode or real-time animation mode

Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by Battelle *Since 1965*

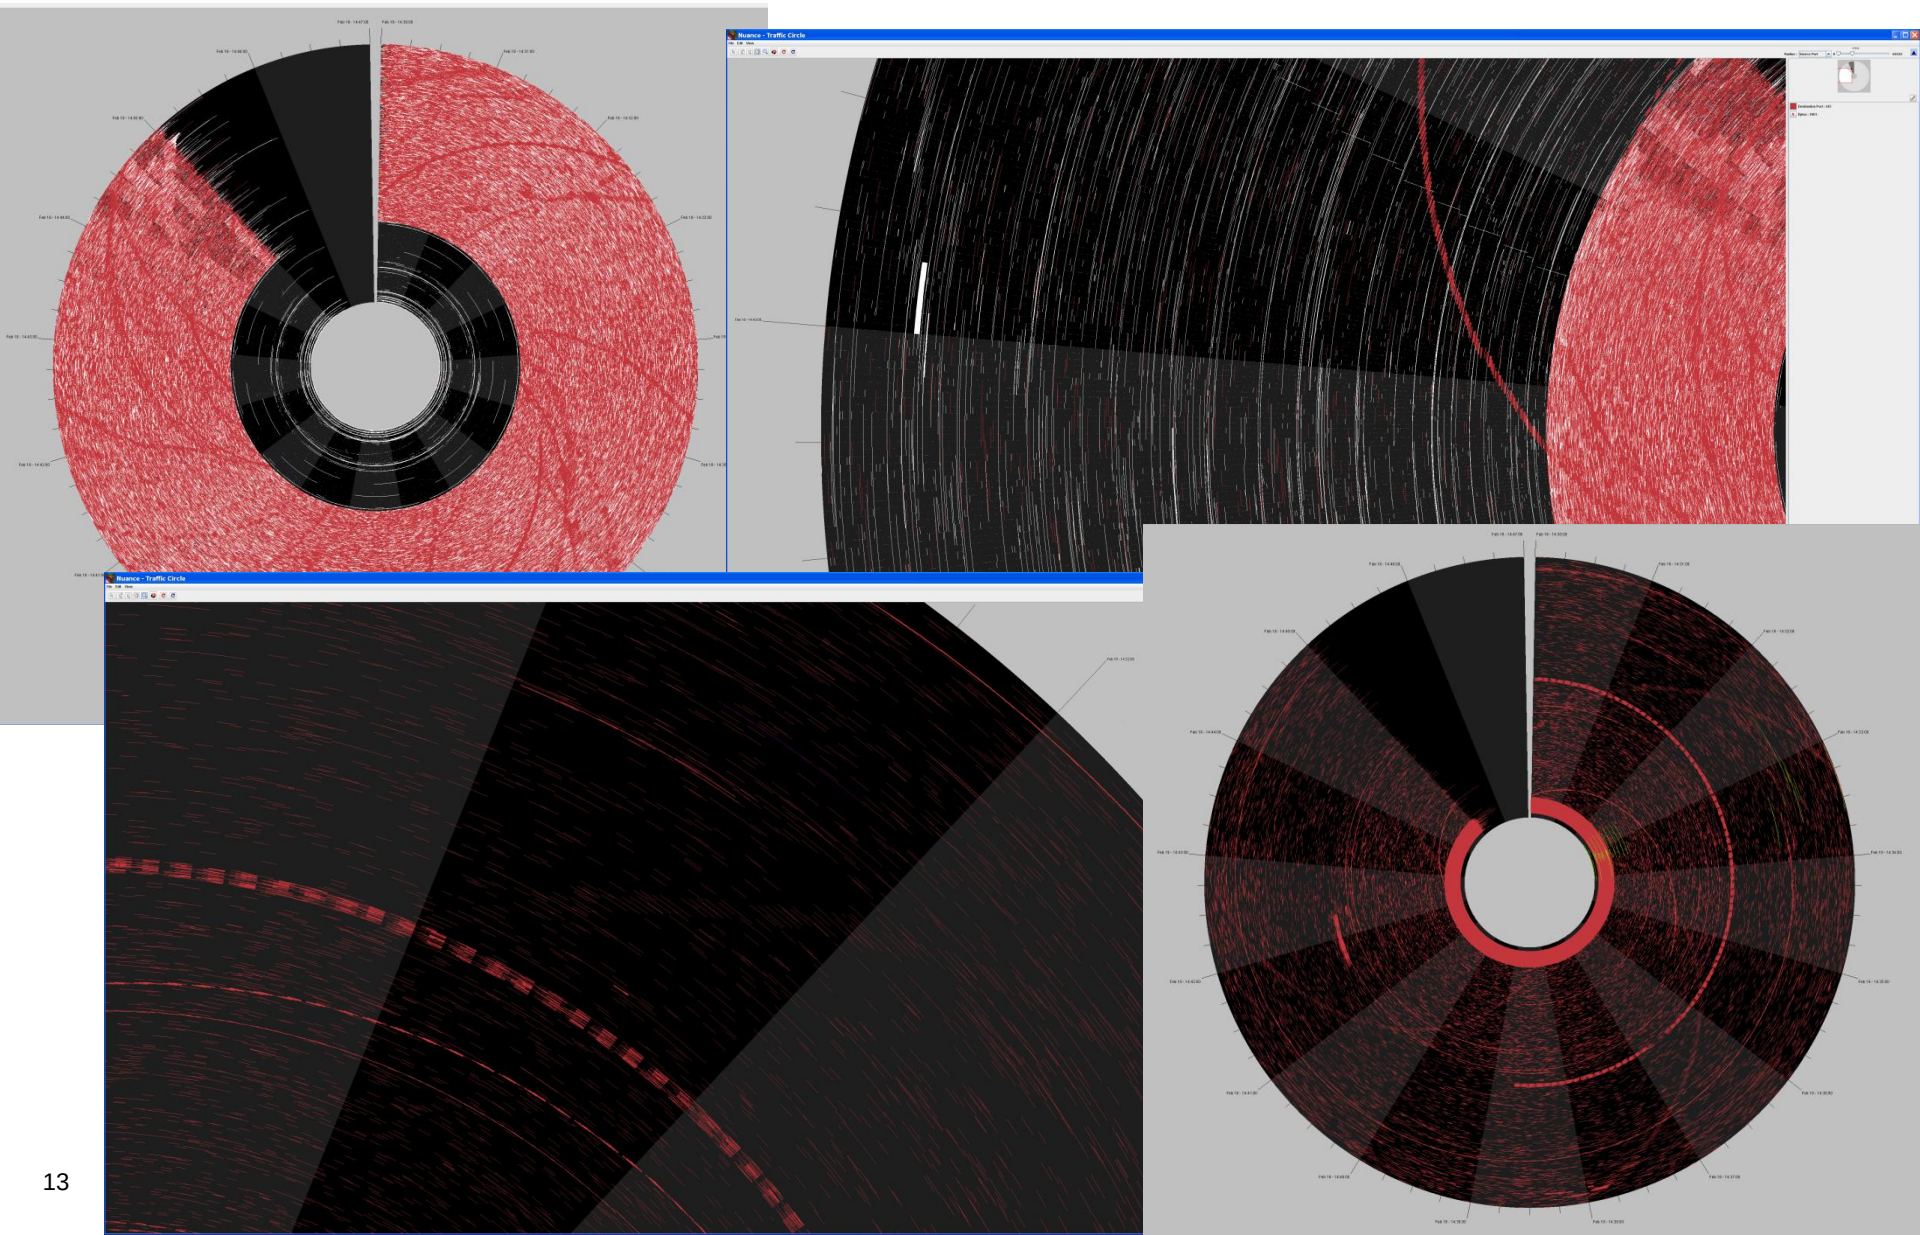
Traffic Circle

- Circular “time wheel” metaphor
  - Flows ordered by start time
  - Arc length corresponds to duration
  - Spinnable interface
- Filters can be added
  - Color coding
  - Hide / show capability
- Operationally demonstrated at data volumes upward of 125 million flows
  - Using high-performance backfill database





Traffic Circle Walkthrough



Conclusions

- Visualization at different levels of abstraction supports situational awareness in large data sets
- High-throughput pipelines are necessary to scale visual analysis to operational data volumes
- Modeling helps analysts baseline normal activities and quickly identify off-normal conditions

Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

Future Directions

- Develop a predictive capability
  - Nascent behavioral changes can be detected visually in real-time in Traffic Circle and CLIQUE
  - CLIQUE classifier enables **sequence detection** for proactive threat identification

- Explore extensions to other domains
  - Financial fraud detection
  - SCADA system reliability and security

Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by Battelle *Since 1965*