Software Engineering Institute

FloCon 2008 Proceedings

January 2008

CERT Program

http://www.sei.cmu.edu



Carnegie Mellon

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
JAN 2008		**00-00-2008 to 00-00-2008**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
FloCon 2008 Proceedings	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Carnegie Mellon University ,Software Engineering Institute (SEI),Pittsburgh,PA,15213**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
FloCon 2008 held in Savannah, GA, from January 7-10, 2008

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**Same as Report (SAR)**	**1156**	

Copyright 2008 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.

A flexible DDoS detection System using IPFIX

Thomas Hirsch, **Tanja Zseby**
Flocon Workshop 2008
January 07-10, 2008
Fraunhofer Institute FOKUS





This work was done in the context of the NetCentric Security project. NetCentric Security is a project of Deutsche Telekom Laboratories supported by the Fraunhofer Institute for Open Communication Systems (FOKUS).

© 2007-2008 FhG FOKUS

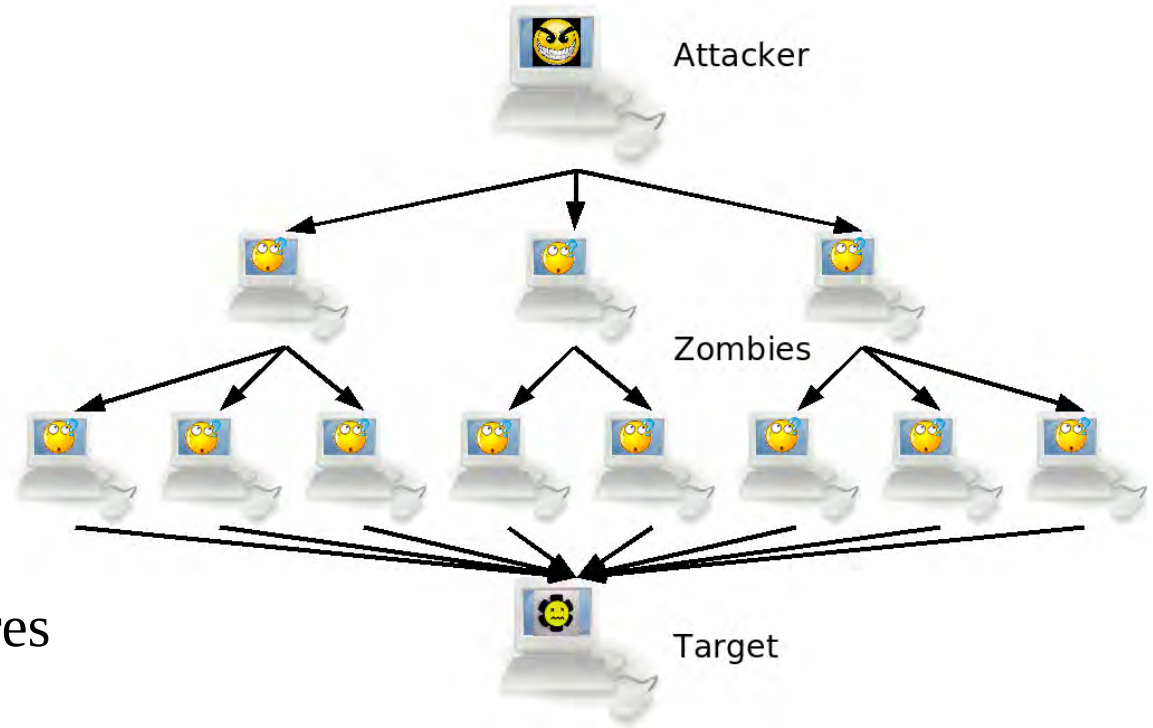
Outline

- Introduction:
  - Denial of Service – The Internet Bottleneck problem
- The Architecture
  - System Architecture
  - OpenIMP platform
  - DDos Detection Metrics
  - Detection using Latent Semantic Indexing and Clustering
- Conclusion:
  - How does IPFIX support the integration of new metrics

© 2007-2008 FhG FOKUS

The DDoS Problem

- DDoS Flooding attacks saturate the final link(s)
- Filters are only effective before the bandwidth becomes scarce
- Hence, the end user can hardly take effective measures

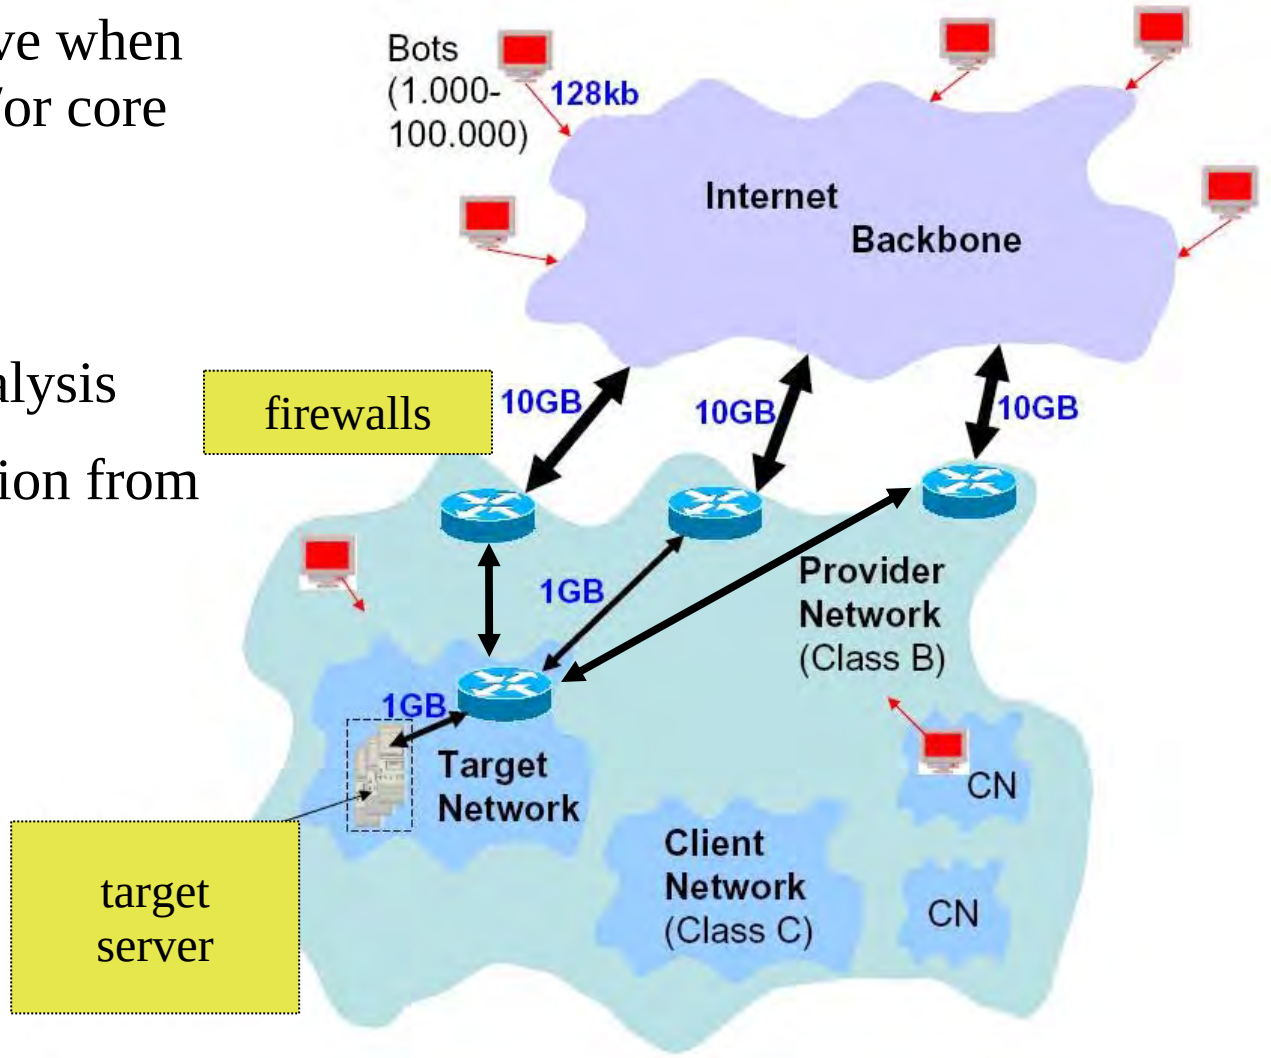


© 2007-2008 FhG FOKUS

Fraunhofer Institute for Open Communication Systems

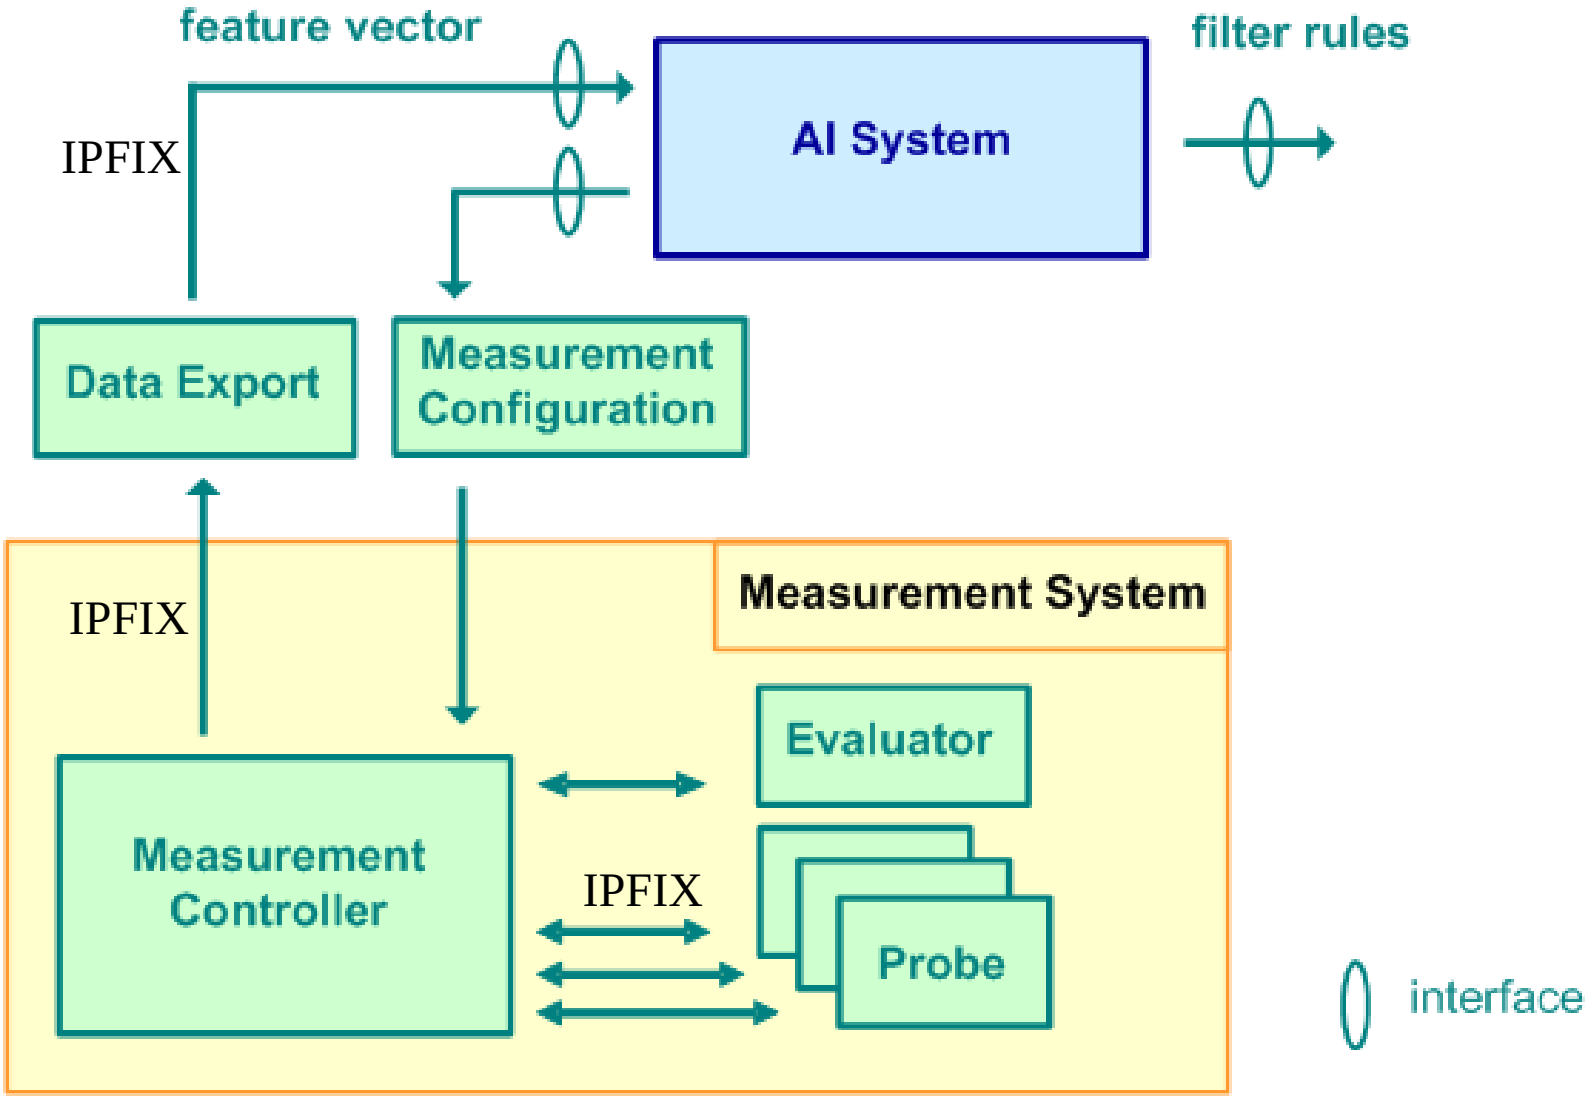
Mitigating DDoS at ISP level

- Mitigation can be effective when implemented on ISP and/or core routers
- This requires
  - high-speed traffic analysis
  - Information aggregation from various sources



© 2007-2008 FhG FOKUS

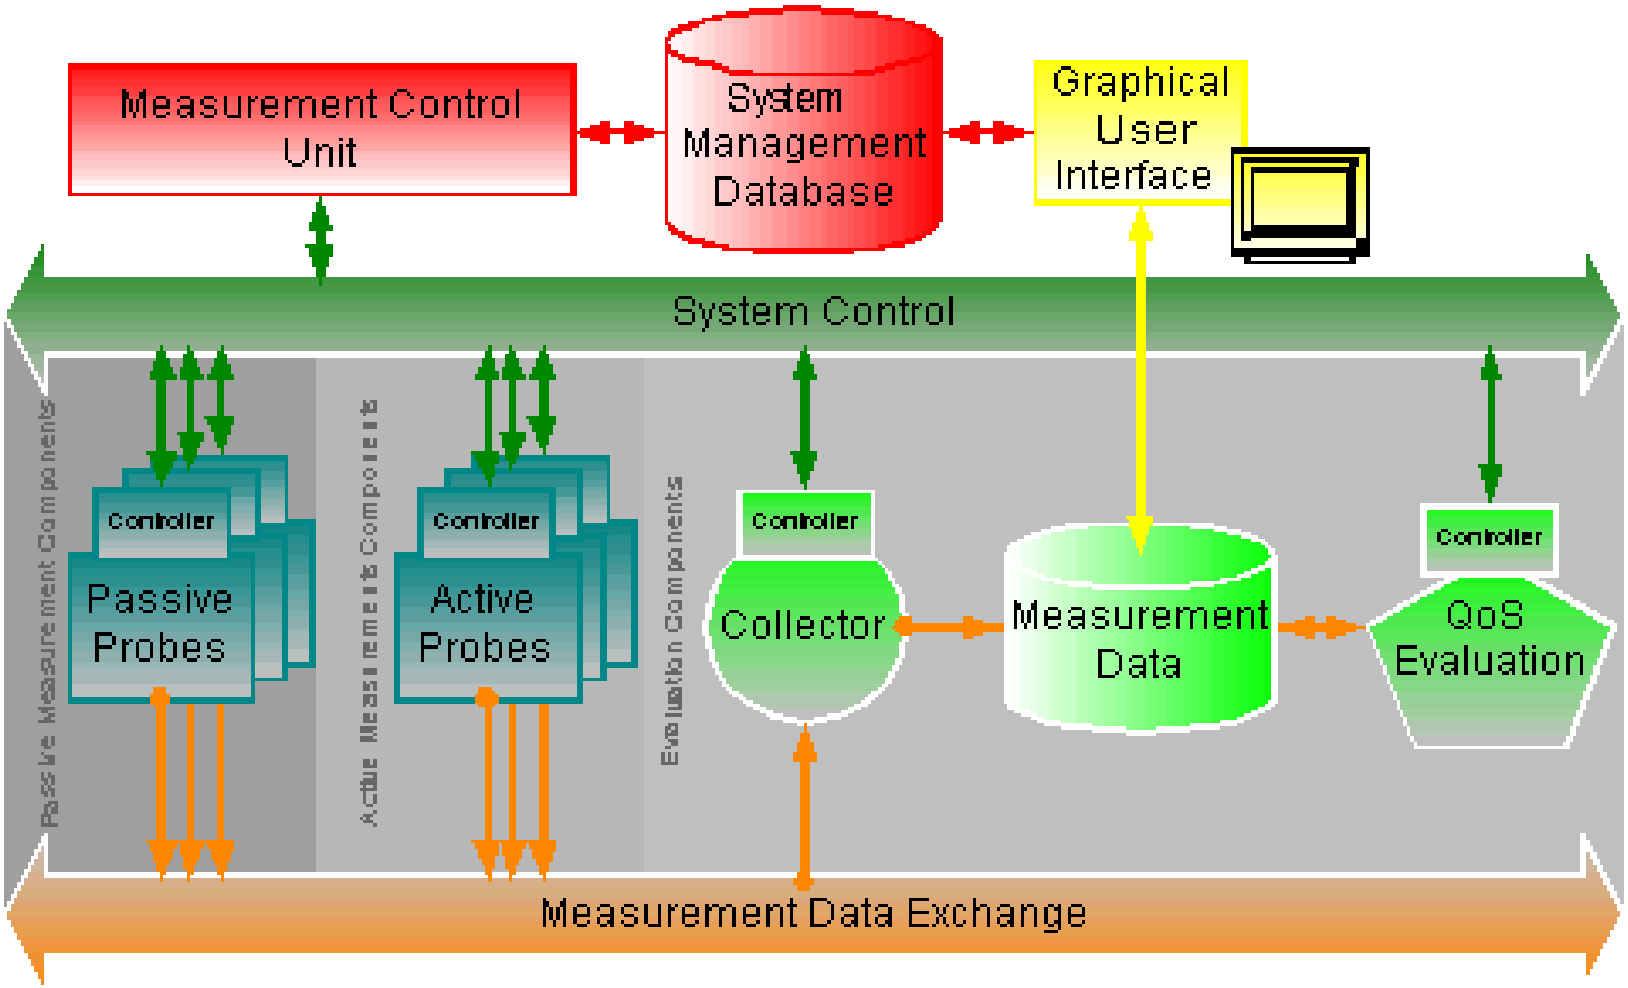
System Overview



© 2007-2008 FhG FOKUS

Fraunhofer Institute for Open Communication Systems

OpenIMP

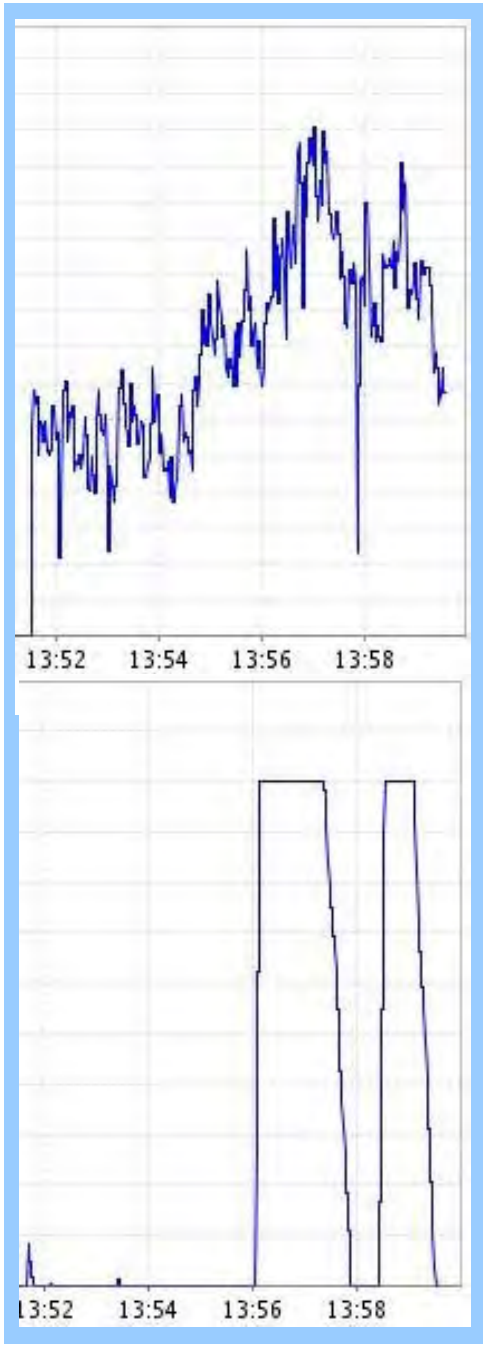


© 2007-2008 FhG FOKUS

Fraunhofer Institute for Open Communication Systems

DDoS Detection Metrics

- Some examples
  - Packet Count (above)
  - Byte Count
  - Packet count per flow / flag / message type
- Transformations
  - CUSUM (below)
  - Wavelet
  - Entropy



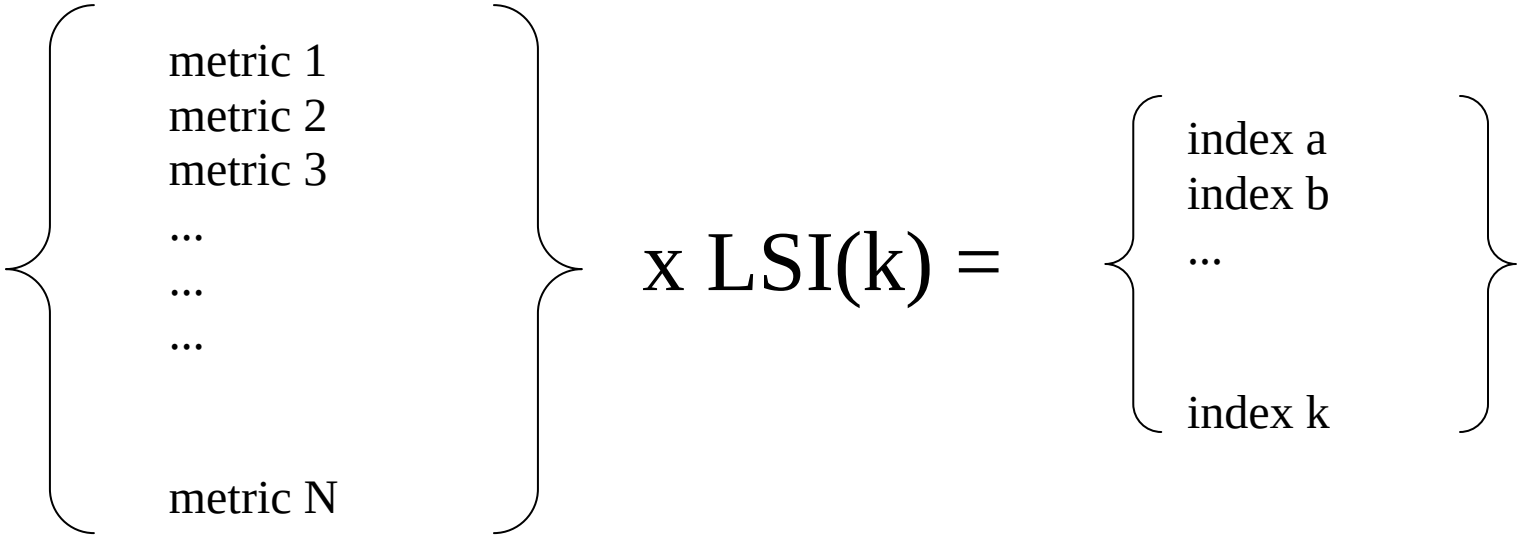
- A multitude of proposals in different papers!
- Which ones to implement?

© 2007-2008 FhG FOKUS

Fraunhofer FOKUS Institute for Open Communication Systems

Latent Semantic Indexing

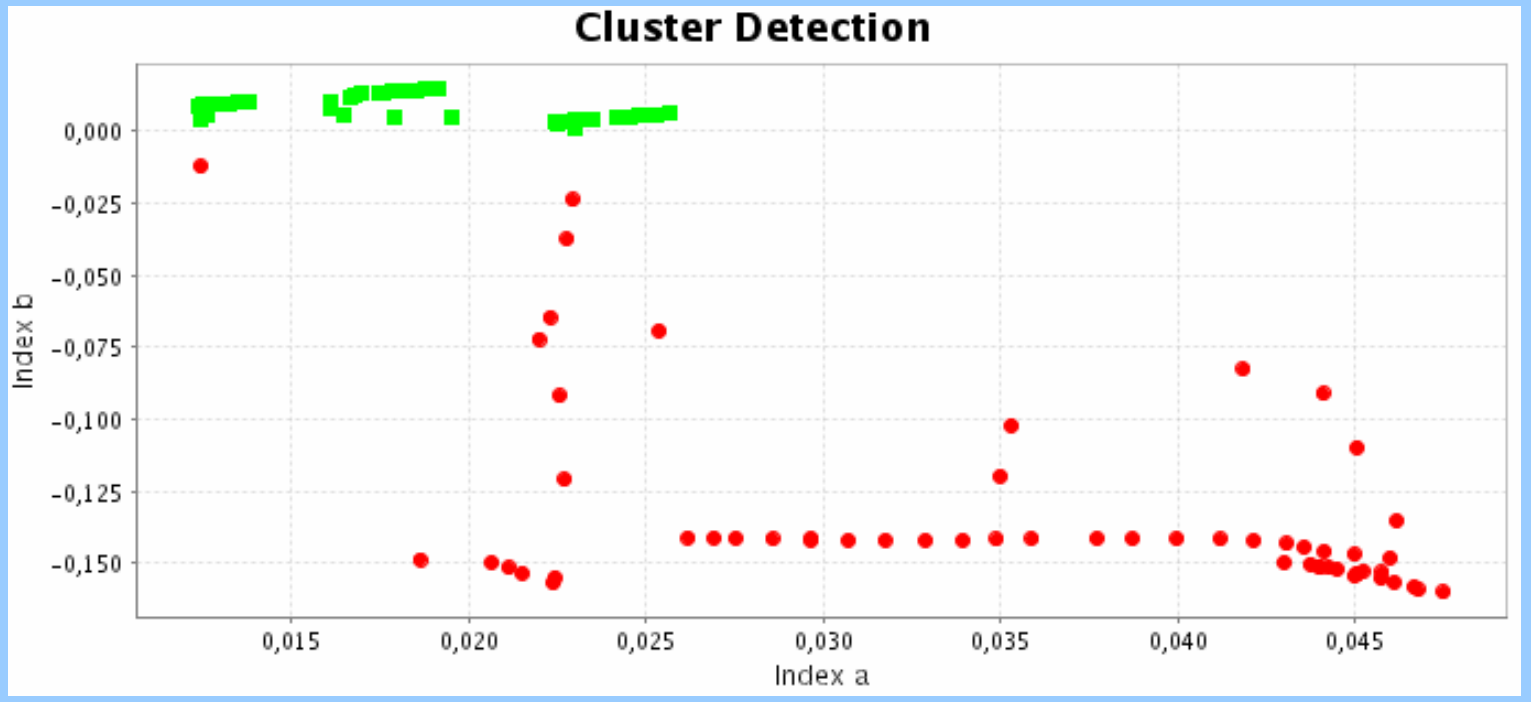
- allows to reduce a multi-dimensional feature vector
- into a lower-dimensional feature vector (easier to process)
- information preserving (principle components)
- maps all metrics into one uniformly sized *feature vector*

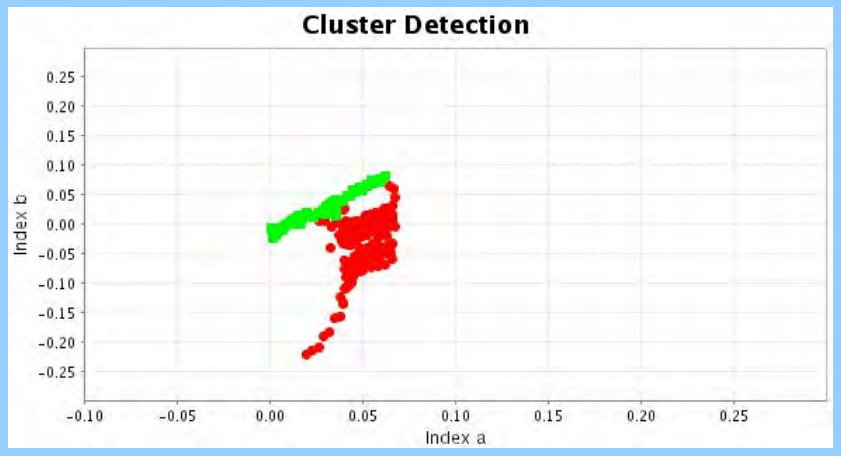


© 2007-2008 FhG FOKUS

Fraunhofer Institute for Open Communication Systems

Cluster Detection







© 2007-2008 FhG FOKUS

Fraunhofer Institute for Open Communication Systems

Cluster Detection

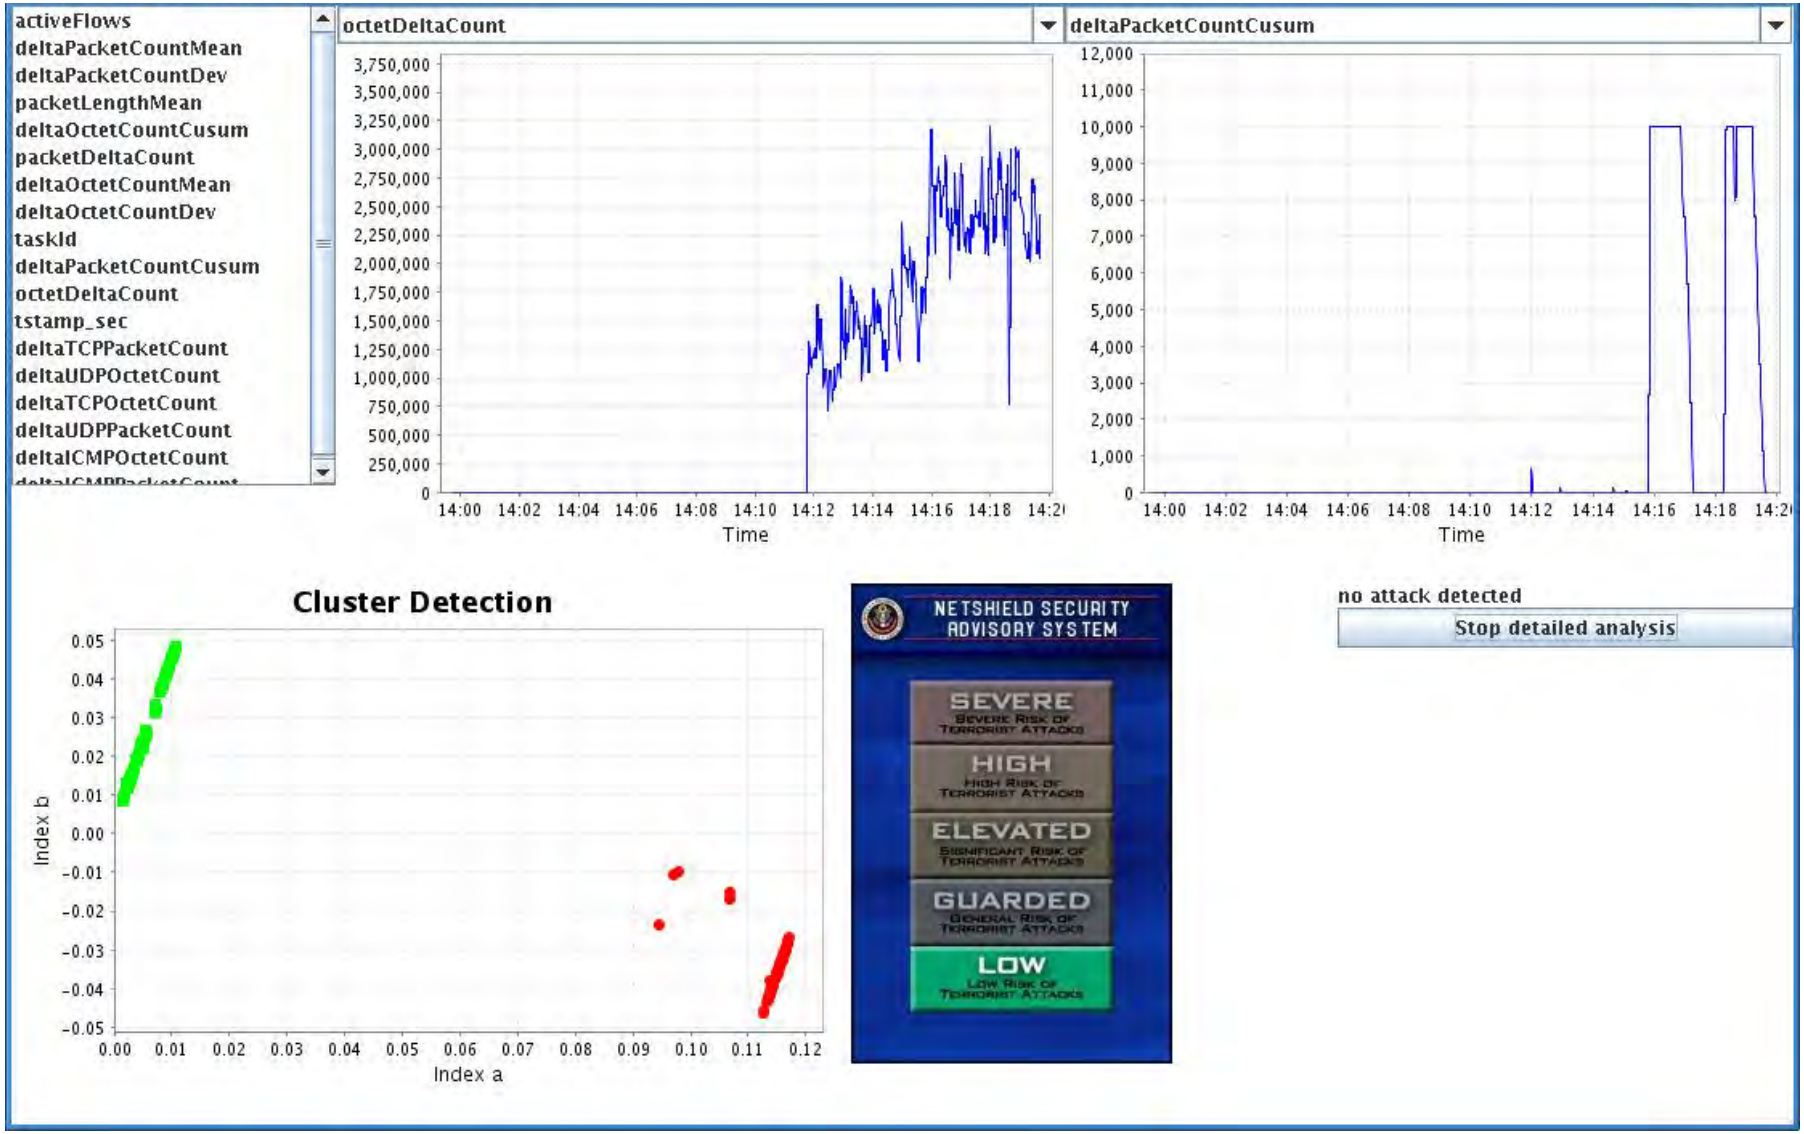
- Unknown Clusters are a possible threat
- Reactions include
  - Filtering, if bandwidth is scarce anyway
  - Detailed analysis of identified anomalies



© 2007-2008 FhG FOKUS

Fraunhofer FOKUS Institute for Open Communication Systems

What it looks like...



FOKUS

Fraunhofer Institute for Open Communication Systems

© 2007-2008 FhG FOKUS

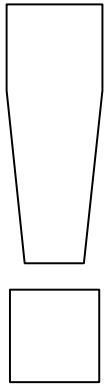
The advantage of using IPFIX

- Established standard for network metrics
- New probes/metrics can be added into the system
  - They immediately speak the language of the system
  - Standard components (routers) may provide the data
  - A training phase is needed for new information sources
- Latent Semantic Indexing reduces any number of metrics
- Cluster Detection operates on the same feature space size

- Detection seamlessly integrates new IPFIX information sources

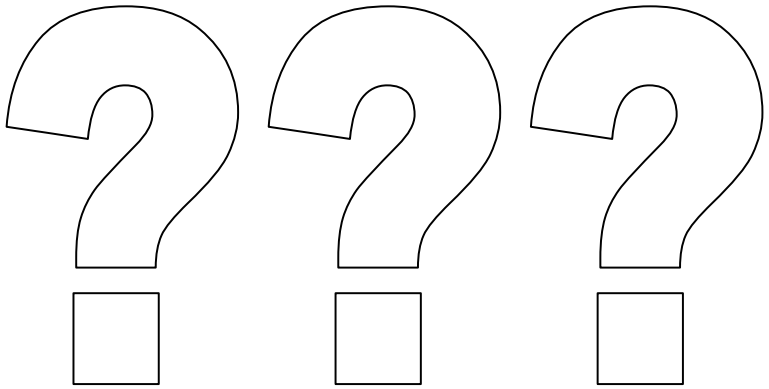
© 2007-2008 FhG FOKUS

Thank You!



© 2007-2008 FhG FOKUS

Questions?



© 2007-2008 FhG FOKUS



A Software Tool for Multi-Field Multi-Level NetFlows Anonymization

<http://scrub-netflows.sourceforge.net/>

William Yurcik

Clay Woolam, Latifur Khan, Bhavani Thuraisingham

University of Texas at Dallas



Motivation: Anonymization?

Anonymization enables entities to share types of data that would otherwise not be shared

(1) Private Data

- User-identifiable information
  - user content (Email messages, URLs)
  - user behavior (access patterns, application usage)
- Machine/Interface addresses
  - IP and MAC addresses

(2) Secret Data

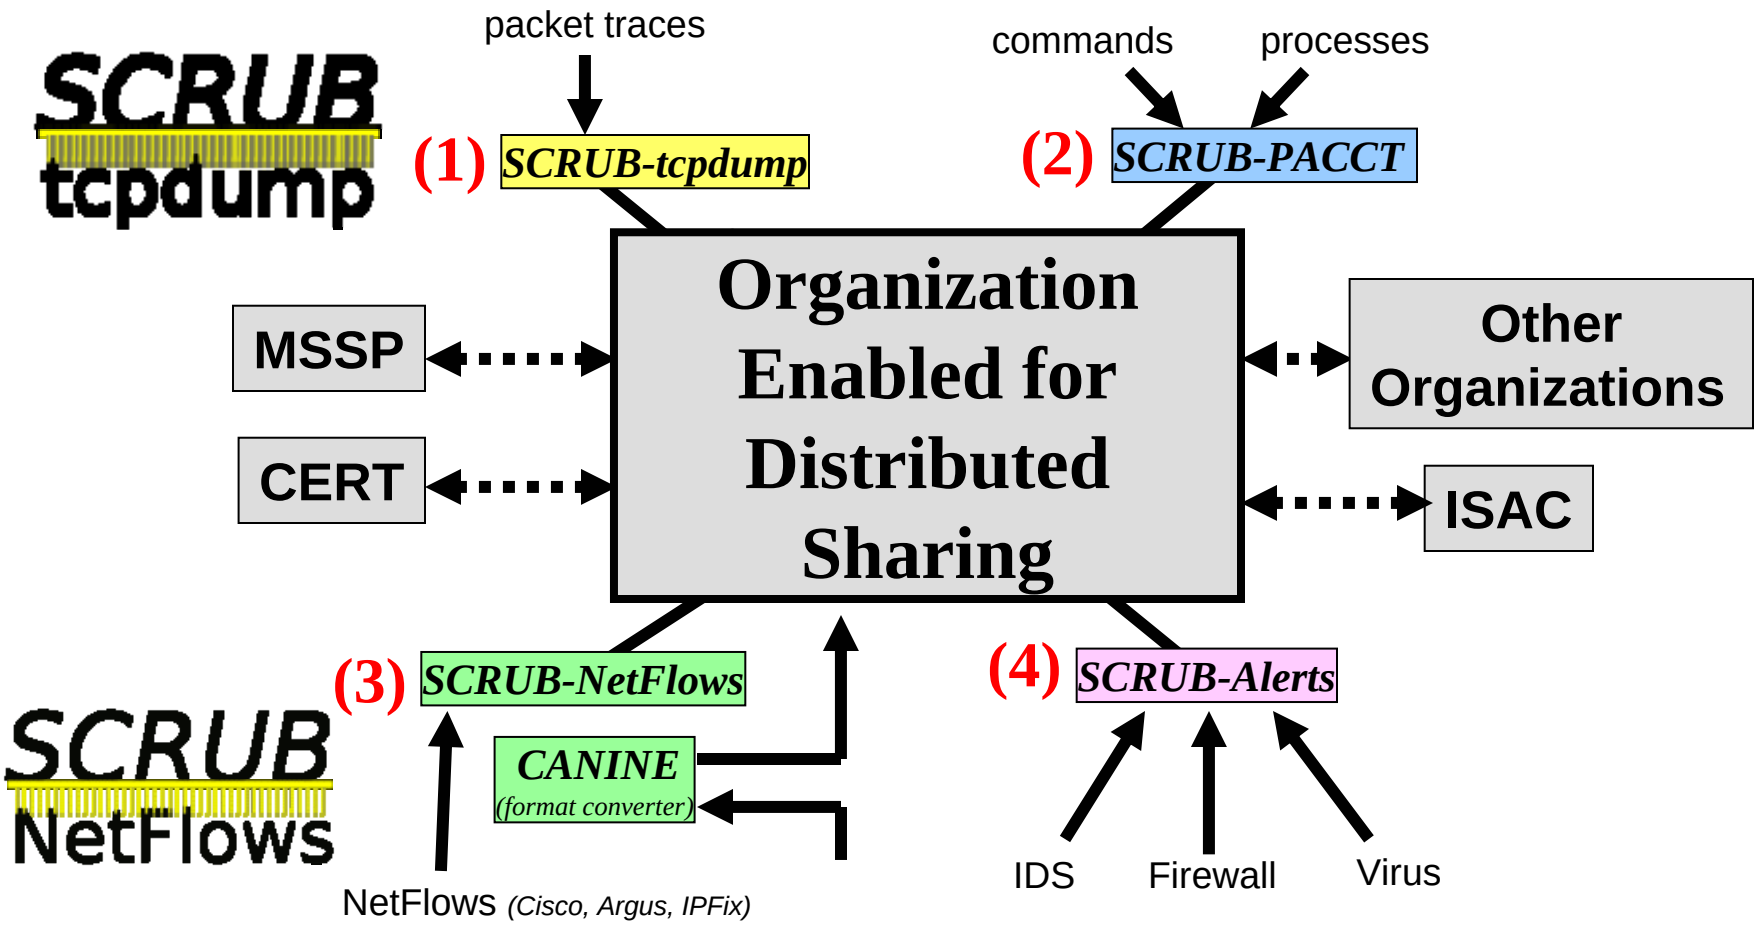
- System configurations (services, topology, routing)
- Traffic patterns (connections, mix, volume)
- Security defenses (firewalls, IDS, routers)
- Attack impacts

Motivation: Sharing?

- **Chasing attackers away (to other organizations) does not improve security**

- **Security data is needed between organizations to correlate events across administrative domains (cumulative learning between organizations)**
  - **Detect attacks**
  - **Blacklist attackers and attacker techniques**
  - **Distinguishing between normal and suspicious network traffic patterns**

*SCRUB** Infrastructure



CANINE (Flocon'05) a NetFlows Converter/Anonymizer



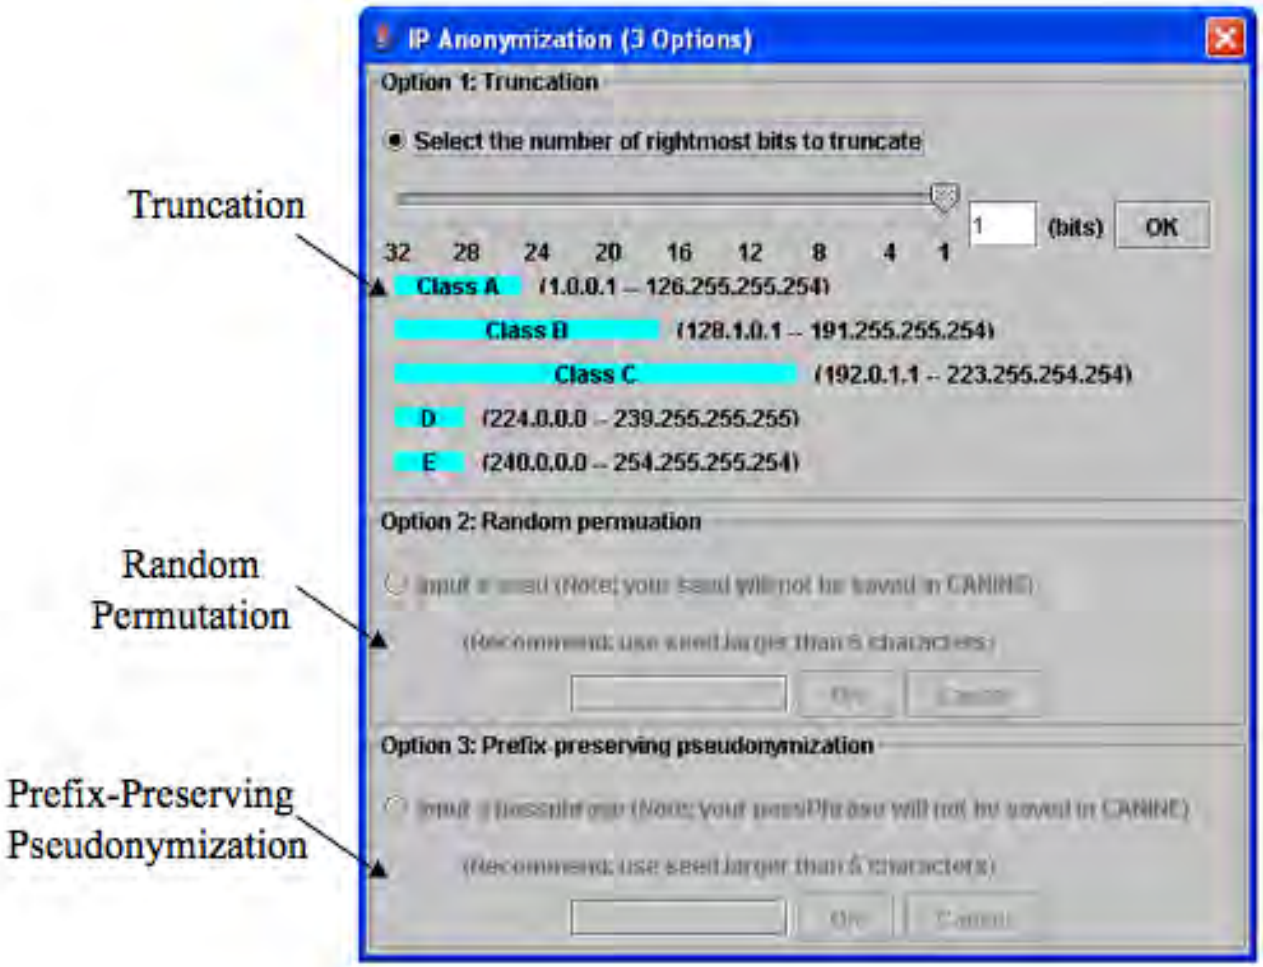
- **CANINE: Converter and ANonymizer for Investigating Netflow Events**

<http://security.ncsa.uiuc.edu/distribution/CanineDownLoad.html>

- **Converter**
  - Cisco V5 & V7, ArgusNCSA, CiscoNCSA, NFDump
- **Anonymizer**
  - 5 NetFlow fields (multi-field)

  (1) IP, (2) Timestamp, (3) Port, (4) Protocol, (5) Byte Count
  - Multiple options for each field (multi-level anonymization)
- **Java GUI – easy to use point-and-click**

IP Address Anonymization in CANINE



SCRUB NetFlows (Flocon'08)
New & Improved NetFlows Anonymizer

- **ASCII-based PERL code**
  - works on any NetFlows format converted to ascii
  - optimized code (multi-threaded parallelization)
- **Anonymizes more NetFlow fields (10>5)**
  - adding support for additional fields is minimal
  - (6) TimeStamp (first/last pkt) (7) TOS (8) TTL (9) TCP Flags (10) Packet Count
- **Improved/More anonymization options per field**
  - Fixes Crypto-PAn IP address anonymization flaw
  - Working on tailoring semantics to low/medium/high
- **Command line operation**
  - UNIX friendly, consistency with other SCRUB* tools
  - cascaded streaming operation available via piping

SCRUB-NetFlows
Multi-Level Anonymization Options

- Black Marker (filtering/deletion)
- Pure Randomization (replacement)
- Keyed Randomization (replacement)
- Annihilation/Truncation (accuracy reduction)
- Prefix-Preserving Pseudonymization (IP address)
- Grouping (accuracy reduction)
  - Bilateral Classification
- Enumeration (time, adding noise)
- Time Shift (time, adding noise)

Example: Timestamp Field (First/Last Pkt)

- Black Marker
  - replacement of field with a predefined constant (0)
- Random Time Shift
  - increments given time by a random value within a user defined window
- Enumeration
  - sorts entries by timestamp, applies black-marker
- Distance-preserving pseudonymization
  - preserve distance between two timestamps
- More
  - including pure/keyed randomization, truncation, unit annihilation

Addressing Crypto-PAn Flaw in SCRUB-NetFlows

- Crypto-PAn is widely used for prefix-preserving pseudonymization
  - flaw discovered – attacker can reverse-engineer the original prefix mapping in a given dataset
- Our use of Crypto-PAn
  - Begin with two separate instances of Crypto-PAn with two distinct keys: Crypt1 and Crypt2
  - Determine network and host portion of IP address
  - Run Crypt1 and Crypt2 on the IP address
  - Return the network of Crypt1 concatenated with the host given by Crypt2

Example usage

- Anonymizations done on one line of an Argus NetFlow
  - The program is told to black marker the source IP, randomize the destination IP, and black marker the first timestamp

```
$ ./scrub-netflow.pl -r ArgusData_146_78 -w AnonData -o "srcip bm dstip rand firsttimestamp bm"
Anonymizing ARGUS format
$ tail -n 1 AnonData
01 Jan 71 01:01:01  02 Oct 03 14:00:50   udp 10.10.10.11.1118   ->    39.7.114.87.55525 6        0
   4856         0           INT

$ tail -n 1 ArgusData_146_78
02 Oct 03 14:00:00  02 Oct 03 14:00:50   udp 132.156.189.139.1118   ->    228.154.76.120.55525 6
0         4856         0           INT

$
```

Anonymization for Sharing: The Privacy vs. Analysis Tradeoff



while anonymization protects against information leakage it also destroys data needed for security analysis

- Zero-Sum? (more privacy <> less analysis & vice versa)
- We are now making measurements of the tradeoff
  - another story but we can talk off-line

Summary

- Critical need for security data sharing between organizations
- Anonymization can provide safe security data sharing
  - Multi-Field: prevent information leakage
  - Multi-Level: no one-size-fits-all anonymization solution
- *SCRUB-NetFlows* as part of a data sharing infrastructure *(SCRUB*)* supporting multiple data sources
  - NetFlows is not the only data source of interest
- No “One-Size-Fits-All” anonymization policy
  - multi-level anonymization options can/should be tailored to requirements of sharing parties to optimize tradeoffs
  - privacy/analysis anonymization tradeoffs need to be characterized

*SCRUB** References

Background on Using Anonymization to Safely Share Security Data

A.J. Slagell and W. Yurcik, "Sharing Computer Network Logs for Security and Privacy: A Motivation for New Methodologies of Anonymization," *1st IEEE Intl. Workshop on the Value of Security through Collab. (SECOVAL),* 2005.

A.J. Slagell and W. Yurcik, "Sharing Network Logs for Security and Privacy: A Motivation for New Methodologies of Anonymization," *ACM Computing Research Repository (CoRR) Technical Report cs.CR/0409005*, September 2004.

X. Yin, K. Lakkaraju, Y. Li, and W. Yurcik, "Selecting Log Data Sources to Correlate Attack Traces For Computer Network Security: Preliminary Results," *11th Intl. Conf. on Telecomunications,* 2003.

W. Yurcik, James Barlow, Yuanyuan Zhou, Hrishikesh Raje, Yifan Li, Xiaoxin Yin, Mike Haberman, Dora Cai, and Duane Searsmith, "Scalable Data Management Alternatives to Support Data Mining Heterogeneous Logs for Computer Network Security," *SIAM Workshop on Data Mining for Counter Terrorism and Security*, 2003.

J. Zhang, N. Borisov, and W. Yurcik, "Outsourcing Security Analysis with Anonymized Logs," *2nd IEEE Intl. Workshop on the Value of Security through Collab. (SECOVAL),* 2006.

J. Zhang, N. Borisov, W. Yurcik, A.J. Slagell, and Matthew Smith, "Future Internet Security Services Enabled by Sharing of Anonymized Logs," *Workshop on Security and Privacy in Future Business Services held in conjunction with International Conference on Emerging Trends in Information and Communication Security (ETRICS),* University of Freiburg Germany, 2006.

SCRUB* Tool (1) SCRUB-tcpdump <http://scrub-tcpdump.sourceforge.net/>

W. Yurcik, C. Woolam, G. Hellings, L. Khan, and B. Thuraisingham, *"SCRUB-tcpdump*: A Multi-Level Packet Anonymizer Demonstrating Privacy/Analysis Tradeoffs," *3rd IEEE Intl. Workshop on the Value of Security through Collab. (SECOVAL),* 2007.

SCRUB* Tool (2) SCRUB-PACCT <http://security.ncsa.uiuc.edu/distribution/Scrub-PADownLoad.html>

C. Ermopoulos and W. Yurcik, "NVision-PA: A Process Accounting Analysis Tool with a Security Focus on Masquerade Detection in HPC Clusters," *IEEE Intl. Conf. on Cluster Computing (Cluster),* 2006.

K. Luo, Y. Li, C. Ermopoulos, W. Yurcik, and A.J. Slagell, "*SCRUB-PA*: A Multi-Level Multi-Dimensional Anonymization Tool for Process Accounting," *ACM Computing Research Repository (CoRR) Technical Report cs.CR/0601079,* January 2006.

W. Yurcik and C. Liu, "A First Step Toward Detecting SSH Identity Theft in HPC Cluster Environments, Discriminating Masqueraders Based on Command Behavior," *1st Intl. Workshop on Cluster Security (Cluster-Sec)* in conjunction with *5th IEEE Intl. Symposium on Cluster Computing and the Grid (CCGrid),* 2005.

SCRUB* Tool (3) SCRUB-NetFlows <http://scrub-netflows.sourceforge.net/>>

Y. Li, A.J. Slagell, K. Luo, and W. Yurcik, "CANINE: A Combined Converter and Anonymizer Tool for Processing NetFlows for Security," *13th Intl. Conf. on Telecomunications Systems*, 2005.

K. Luo, Y. Li, A.J. Slagell, and W. Yurcik, "*CANINE:* A NetFlows Converter/Anonymizer Tool for Format Interoperability and Secure Sharing," *FLOCON – Network Analysis Workshop (Network Flow Analysis for Security Situational Awareness),* 2005.

A.J. Slagell, J. Wang, and W. Yurcik, "Network Anonymization: The Application of *Crypto-PAn* to Cisco NetFlows," IEEE/*NSF/AFRL Workshop on Secure Knowledge Management (SKM),* 2004.

SCRUB-NetFlows

<http://scrub-netflows.sourceforge.net/>

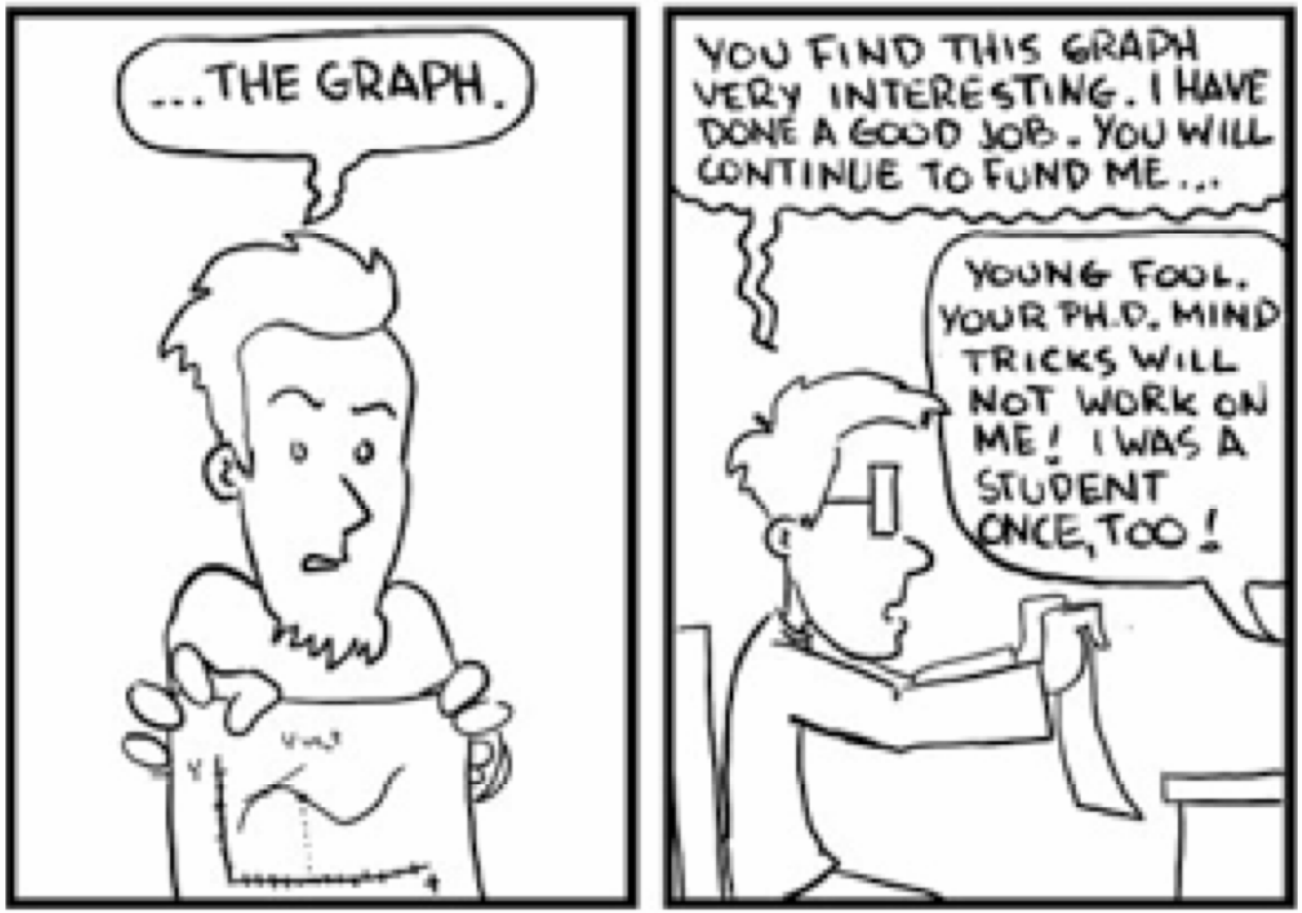


Visualizations of Flow and Analytical Results

Presentation by: Phil Groce and Jeff Janies
Network Situational Awareness group SEI/CERT

© 2007 Carnegie Mellon University

Visualizations are Tools



Visualizations are Tools



Visualizations are Tools



>



Visualizations are Tools





Visualizations are Tools

What we **Want**



What we **think** we have



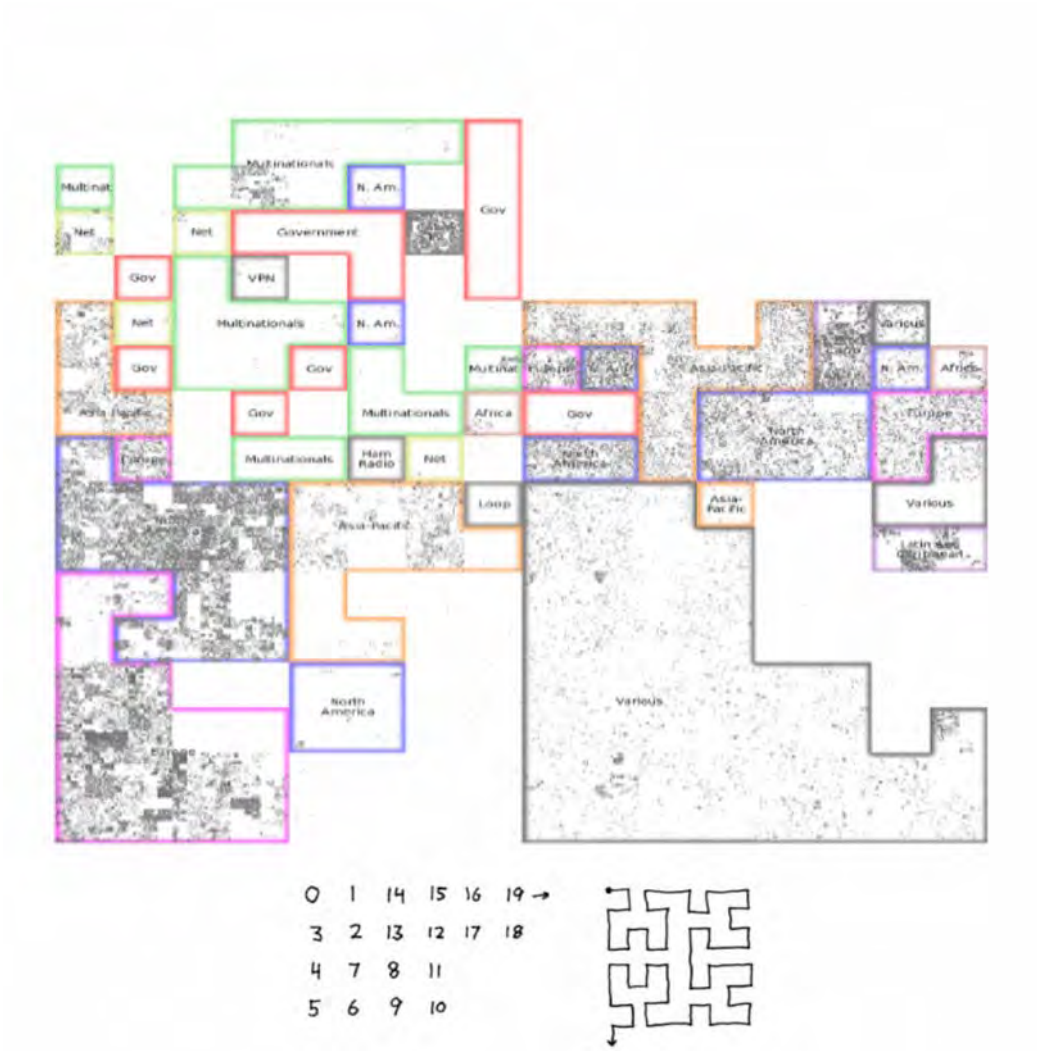
What we **Have**



Hilbert Curve:

Broad Scale Visualization

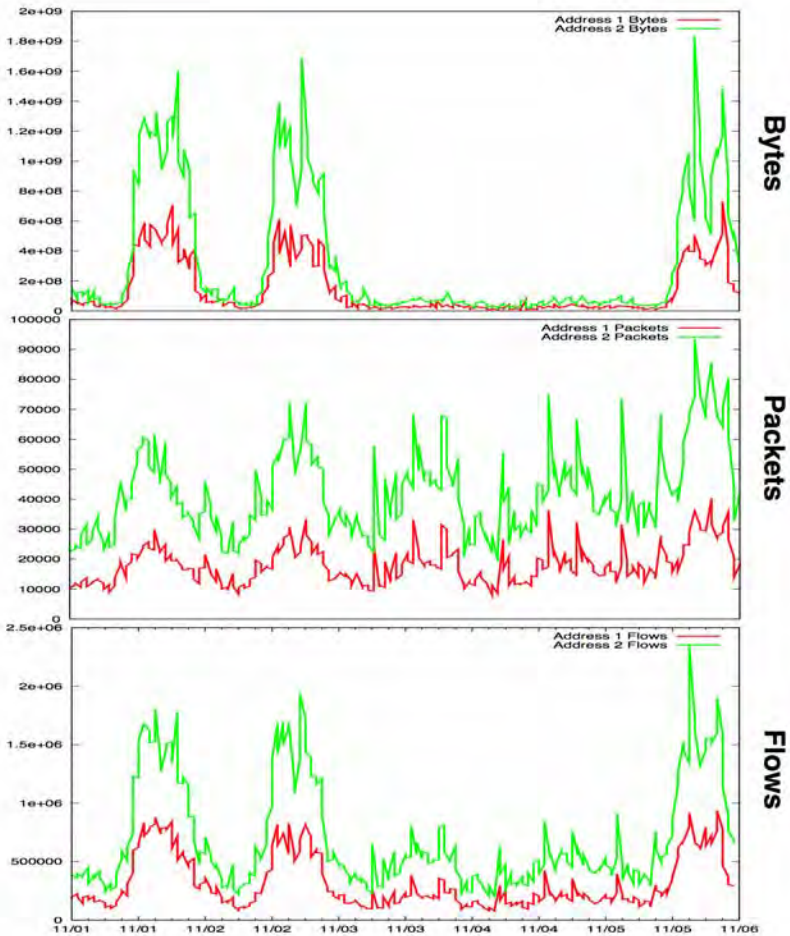
Hilbert Curve



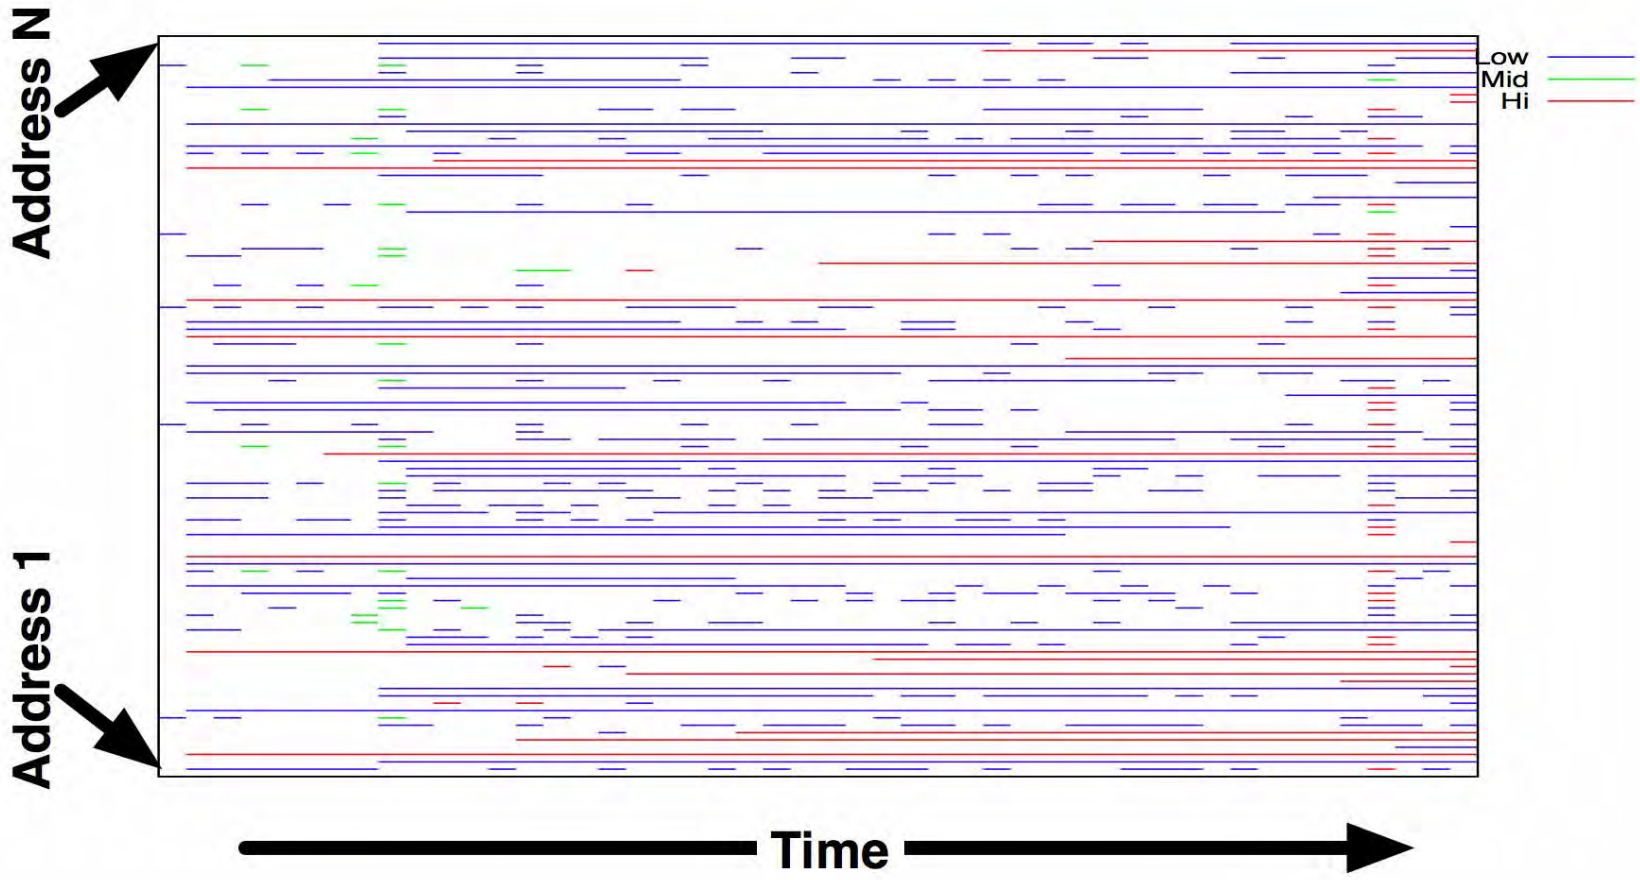
Time Series:

The Tried and True Hammer

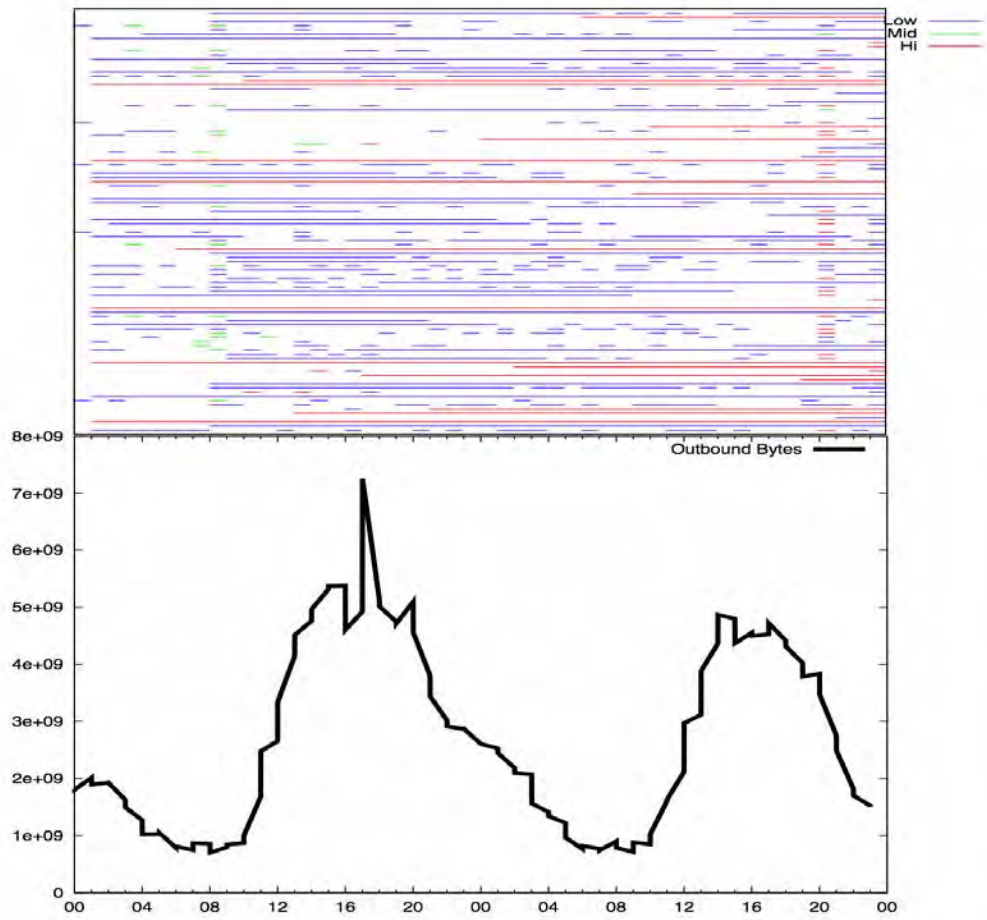
Time Series


Address 1 Bytes
Address 2 Bytes
Bytes
Address 1 Packets
Address 2 Packets
Packets
Address 1 Flows
Address 2 Flows
Flows
11/01
11/01
11/02
11/02
11/03
11/03
11/04
11/04
11/05
11/05
11/06

Existence Plots



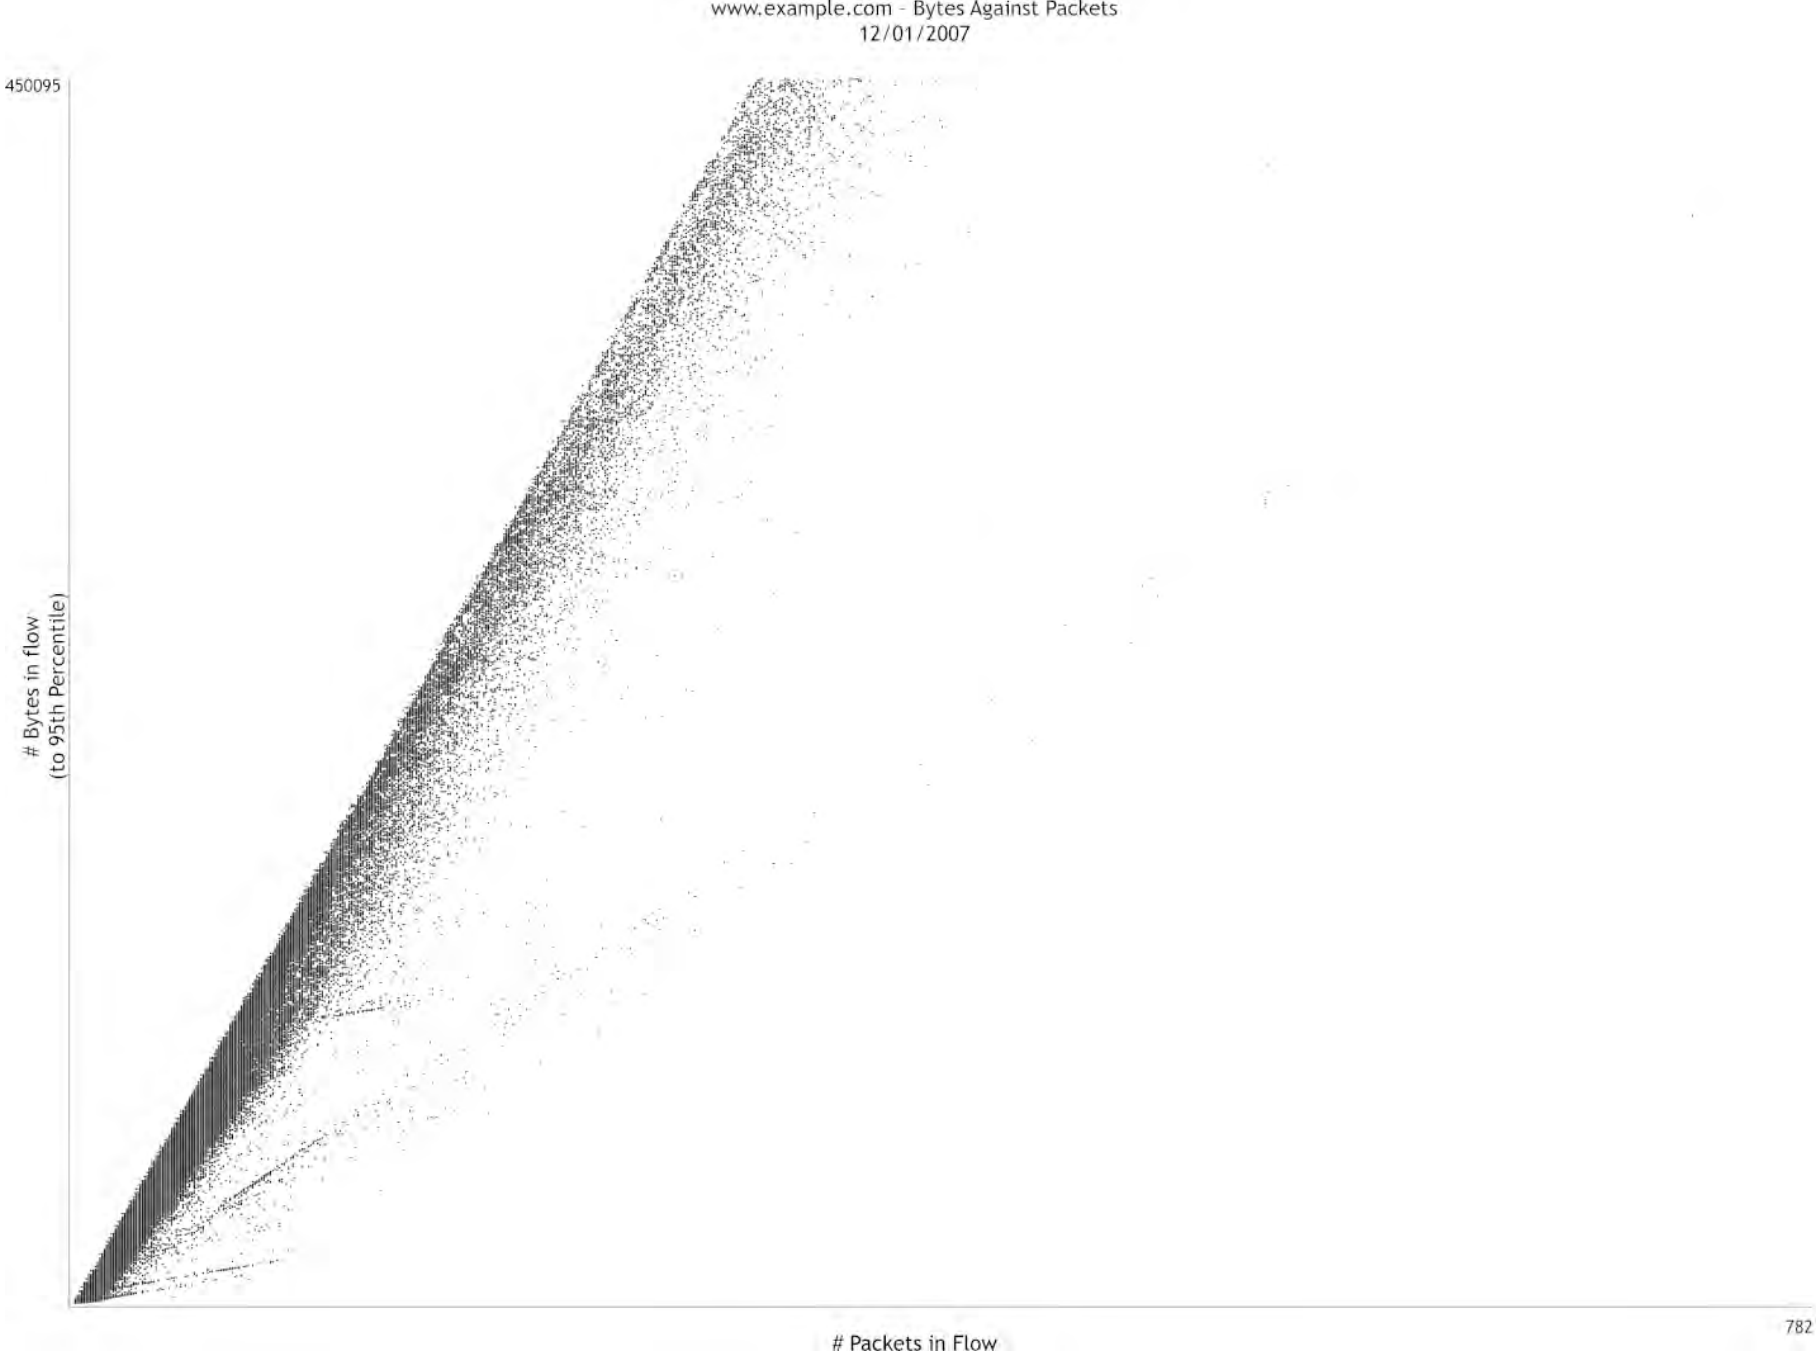
Existence Plots

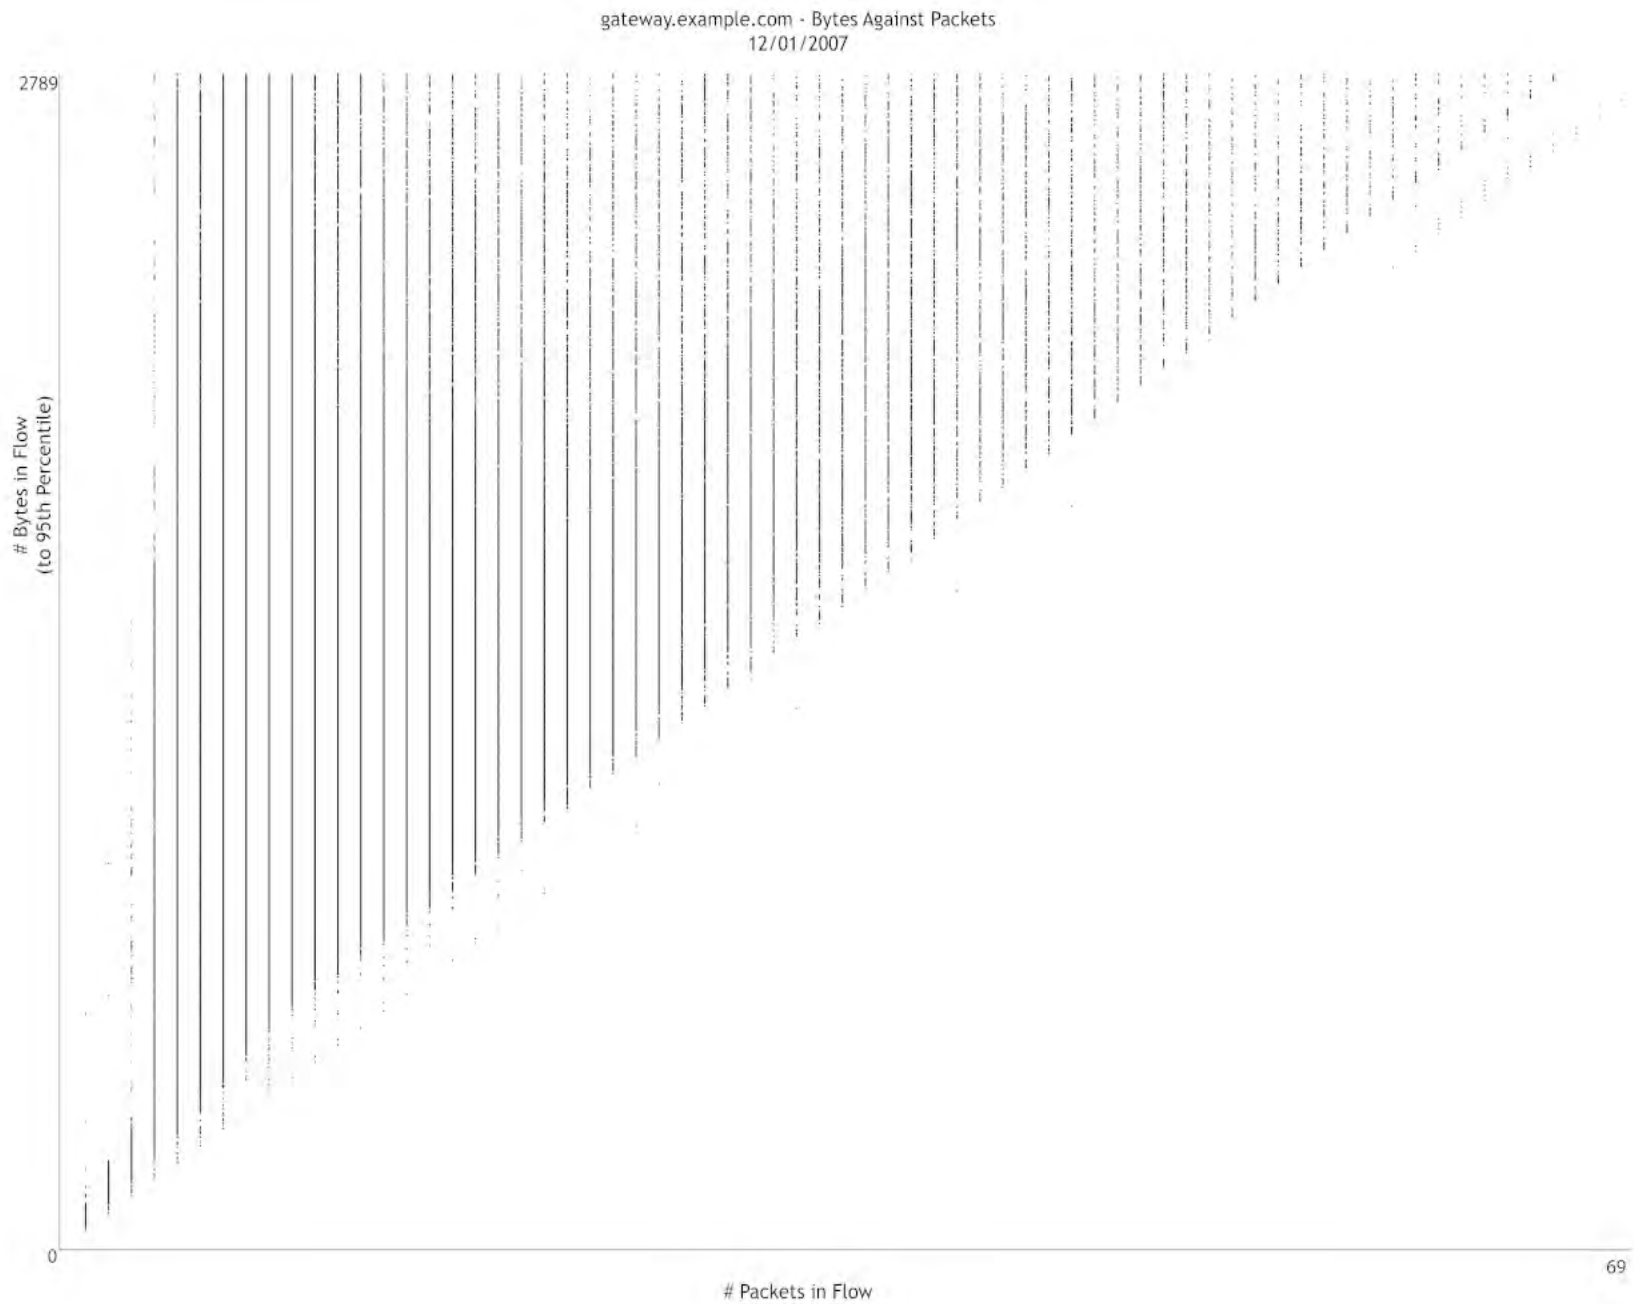


Scatter Plots:

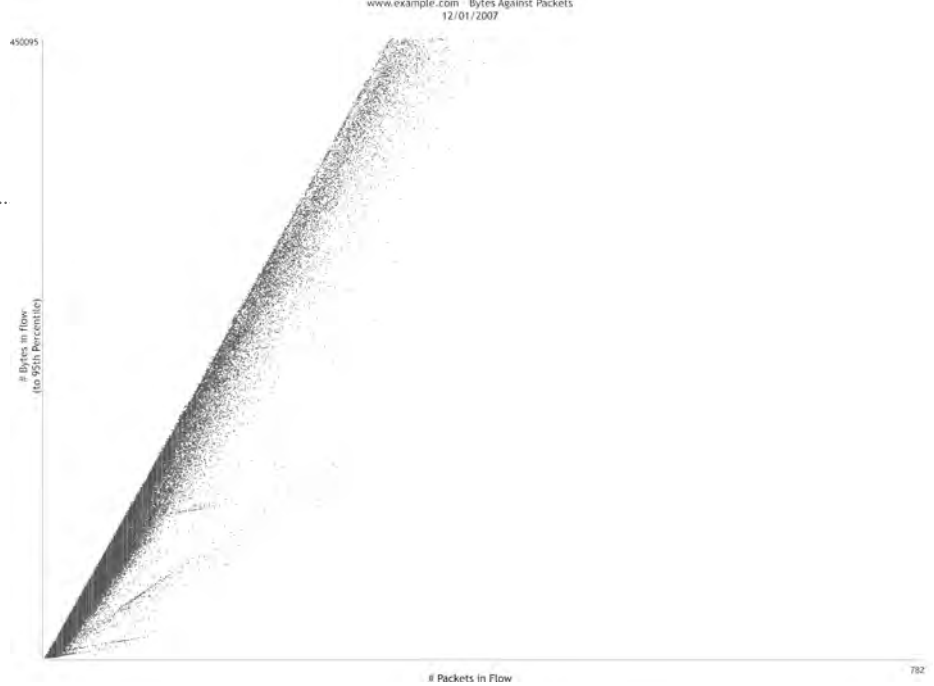
How variables relate to each other

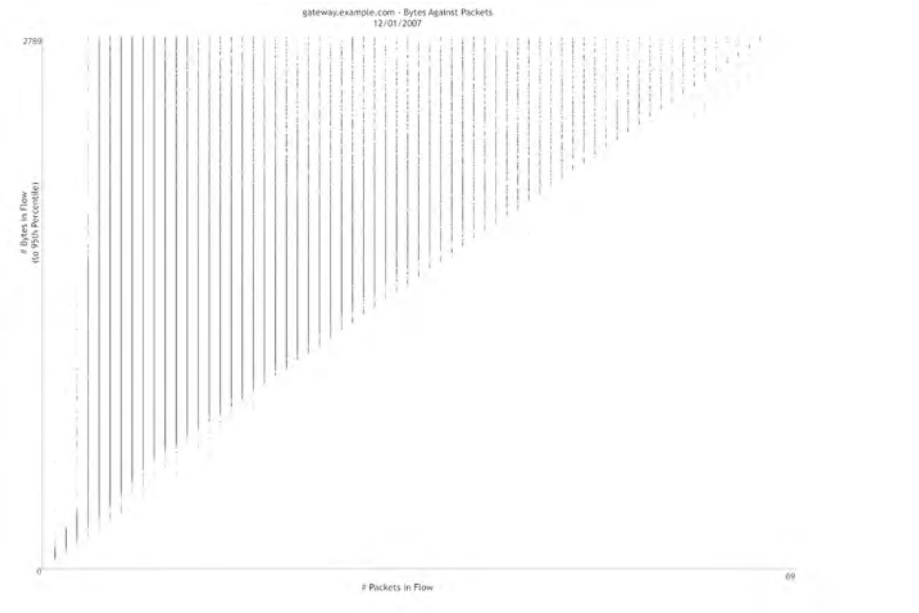

smtp.example.com - Bytes Against Packets
12/01/2007
742
Bytes in flow
(to 95th Percentile)
0
14
Packets in Flow

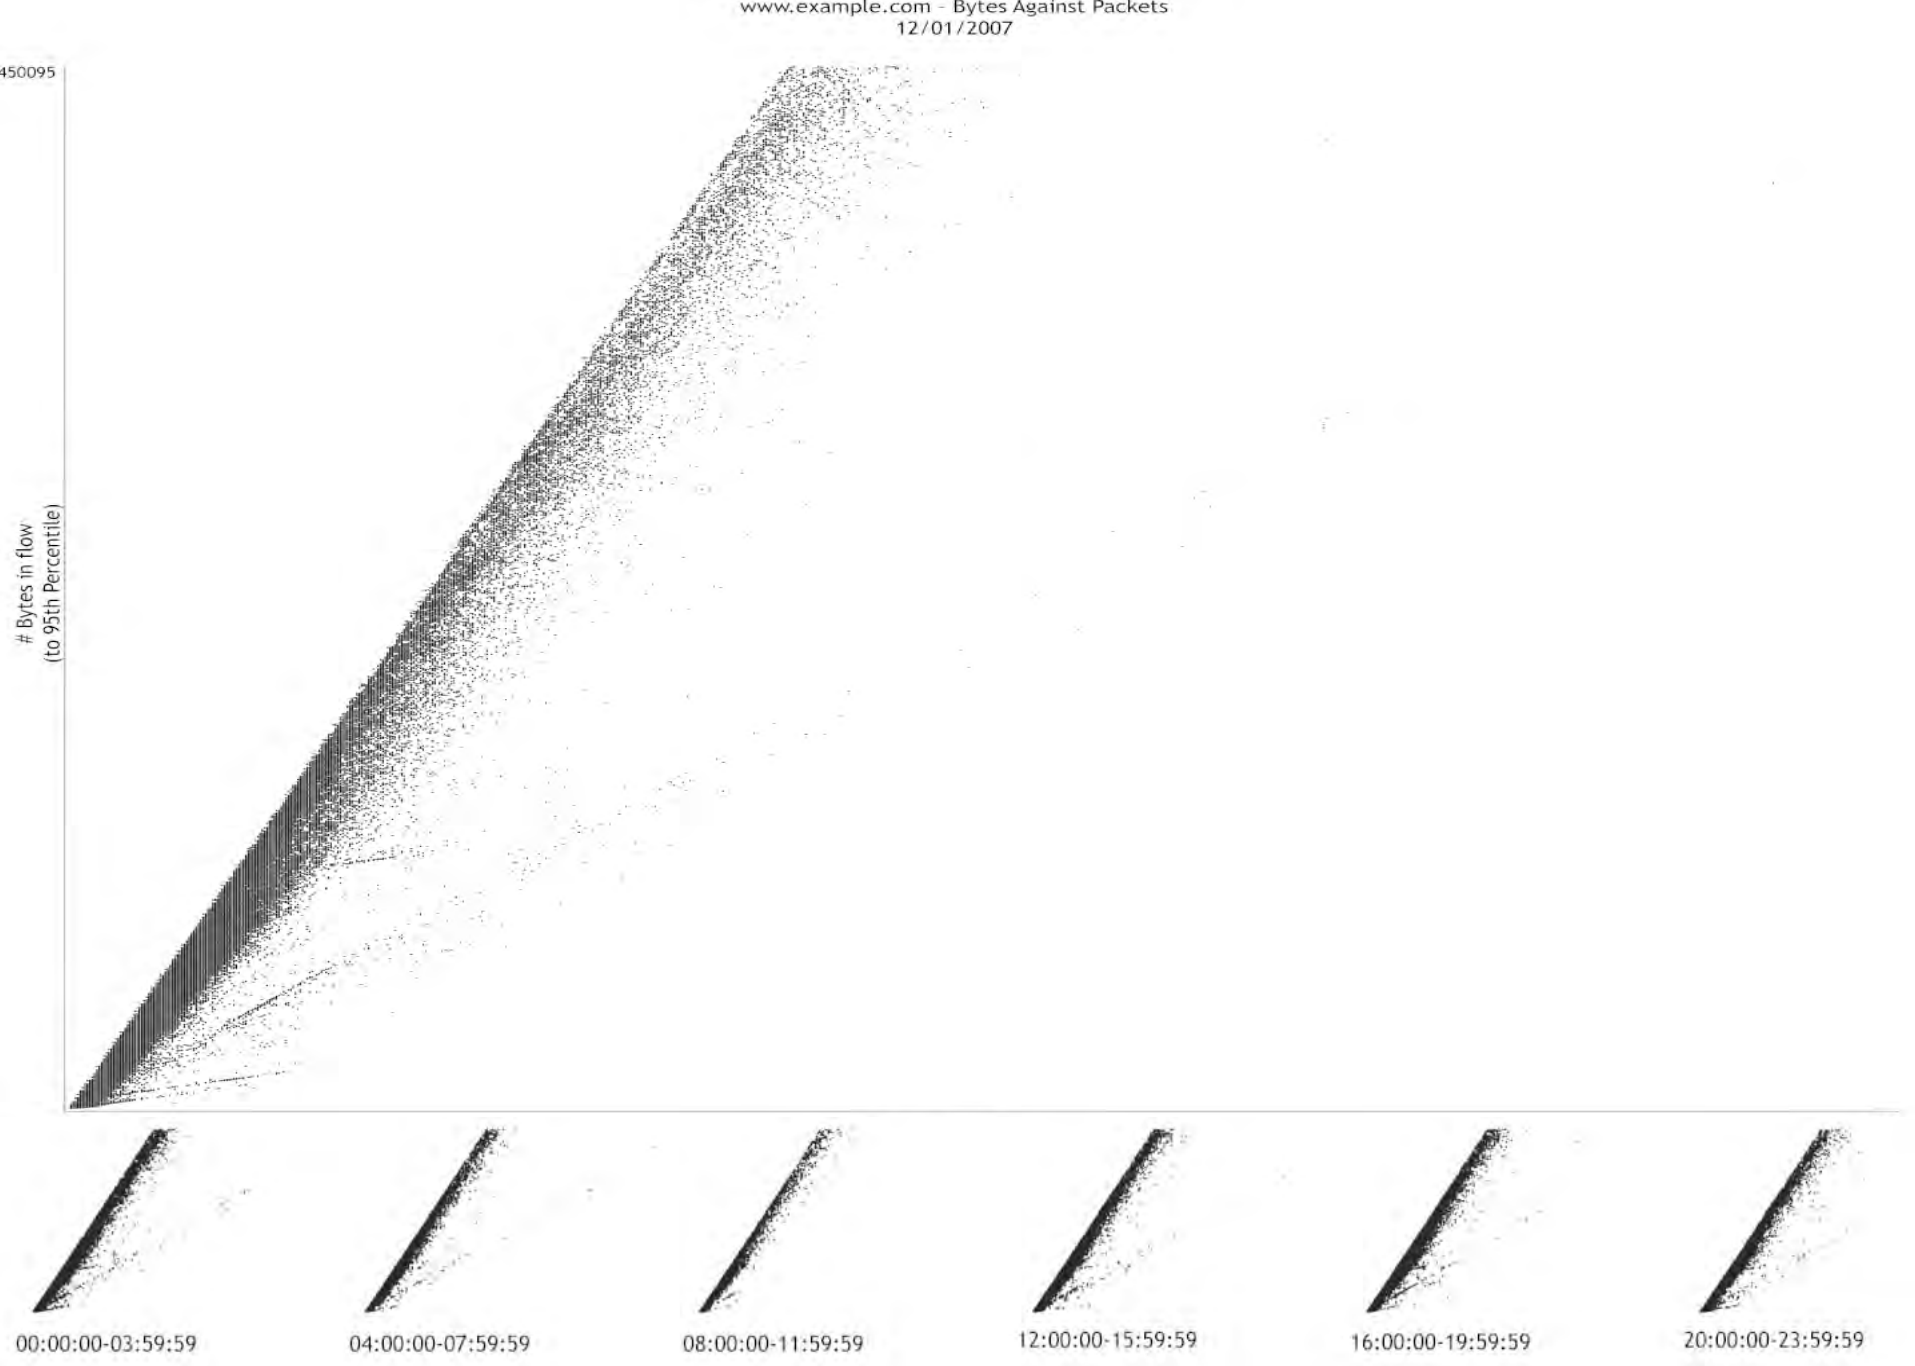

www.example.com - Bytes Against Packets
12/01/2007
450095
Bytes in flow
(to 95th Percentile)
782
Packets in Flow


gateway.example.com - Bytes Against Packets
12/01/2007
2789
Bytes in Flow
(to 95th Percentile)
0
69
Packets in Flow


smtp.example.com - Bytes Against Packets
12/01/2007
742
0
Bytes in Flow
(to 95th Percentile)
Packets in Flow
14


www.example.com - Bytes Against Packets
12/01/2007
450095
Bytes in Flow
(to 95th Percentile)
Packets in Flow
782


gateway.example.com - Bytes Against Packets
12/01/2007
2789
Bytes in Flow
(to 95th Percentile)
0
Packets in Flow
69


www.example.com - Bytes Against Packets
12/01/2007
450095
Bytes in flow
(to 95th Percentile)
00:00:00-03:59:59
04:00:00-07:59:59
08:00:00-11:59:59
12:00:00-15:59:59
16:00:00-19:59:59
20:00:00-23:59:59

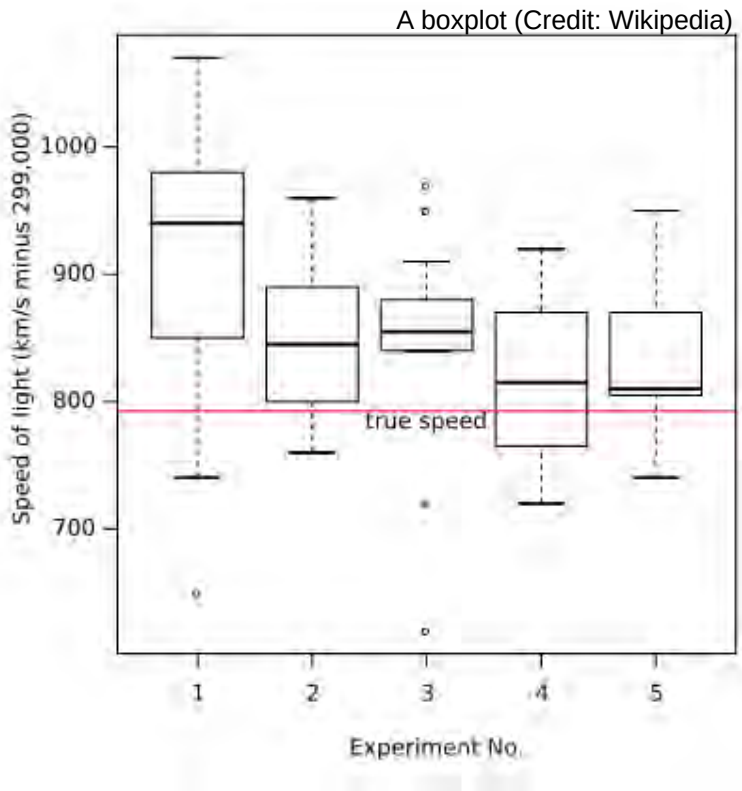


Distributions:

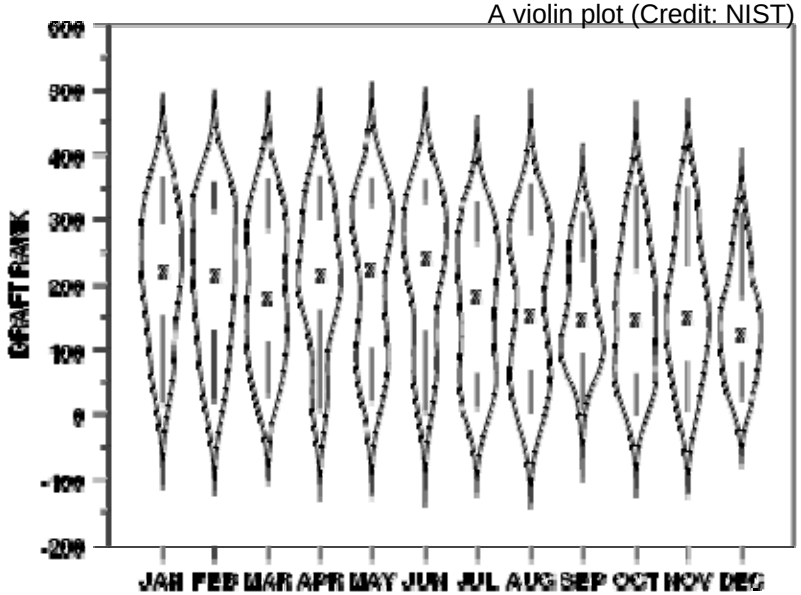
How a variable relates to itself

Box and Violin Plots

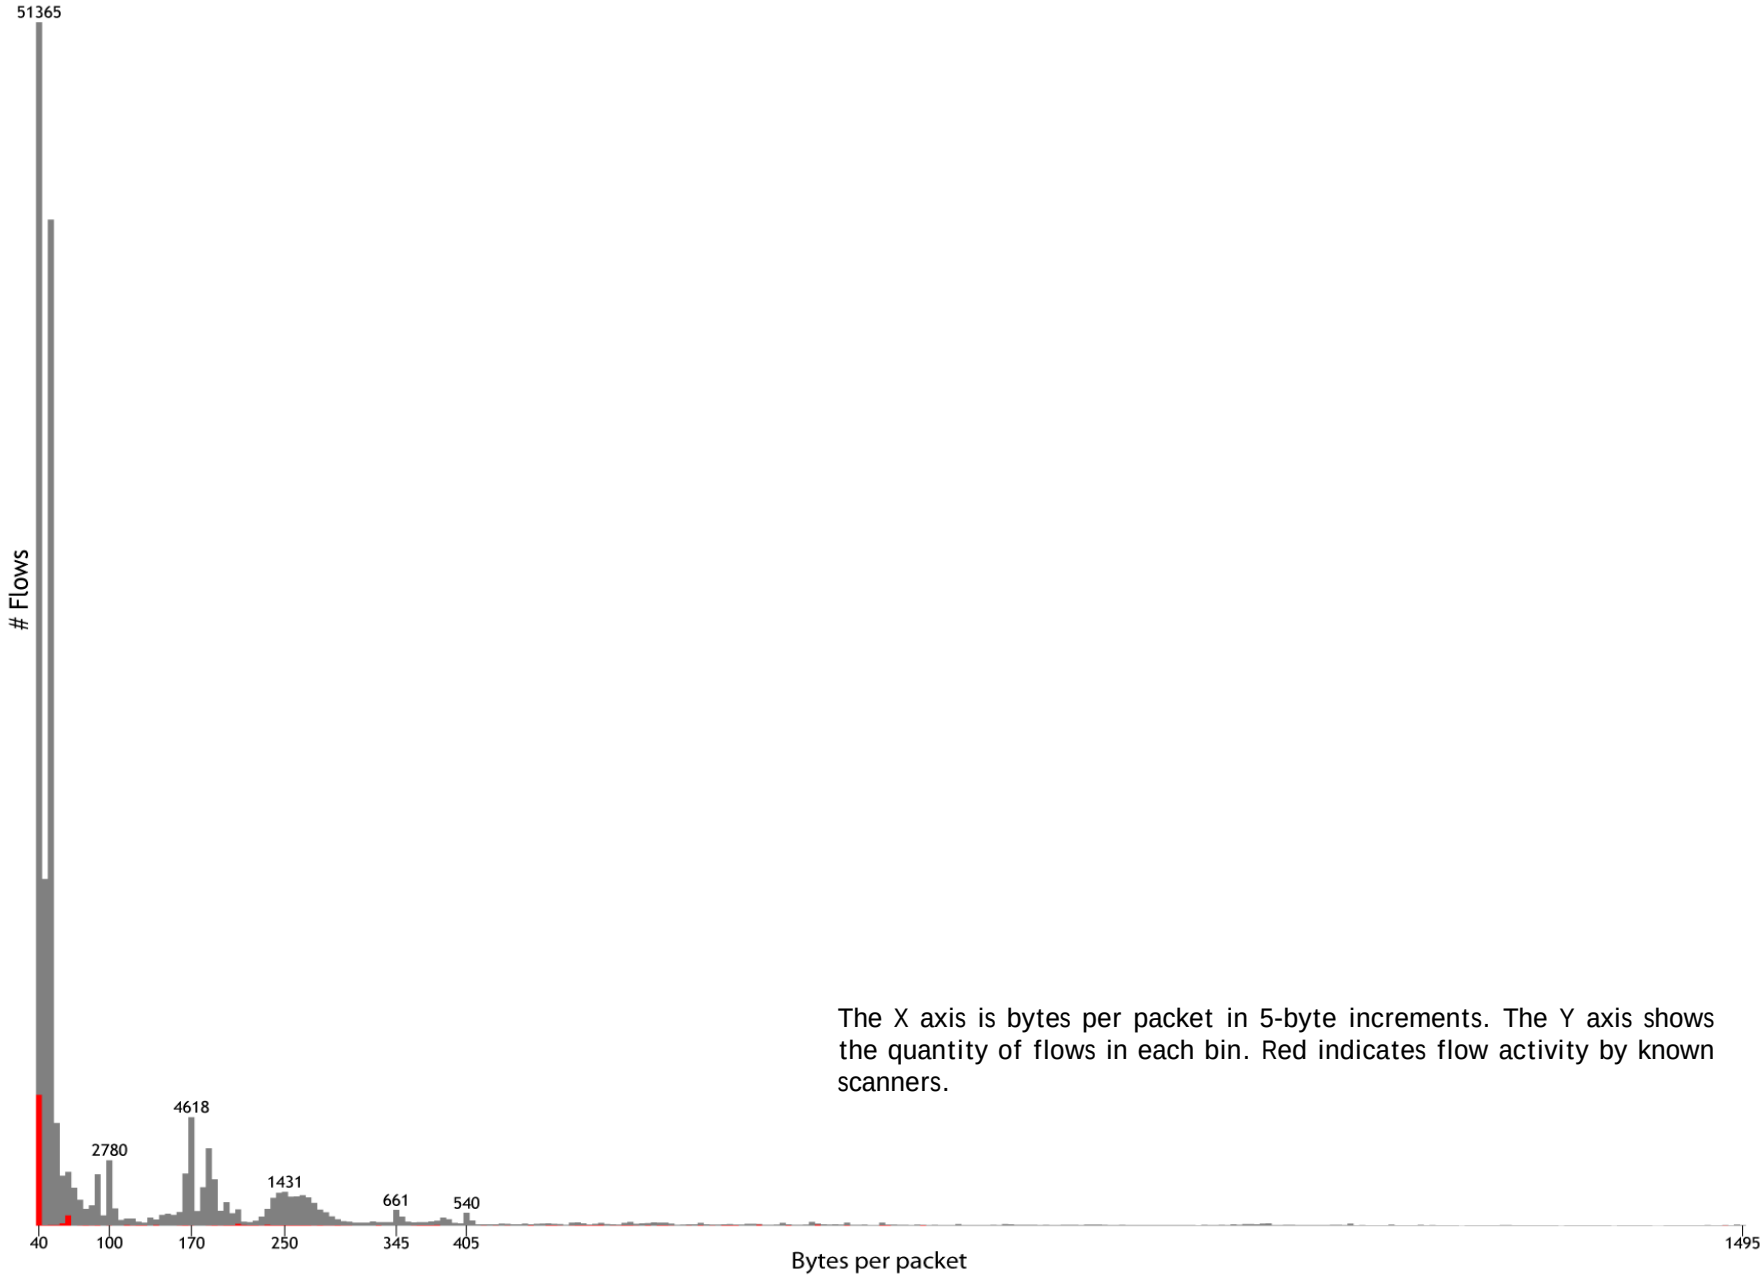
A boxplot (Credit: Wikipedia)



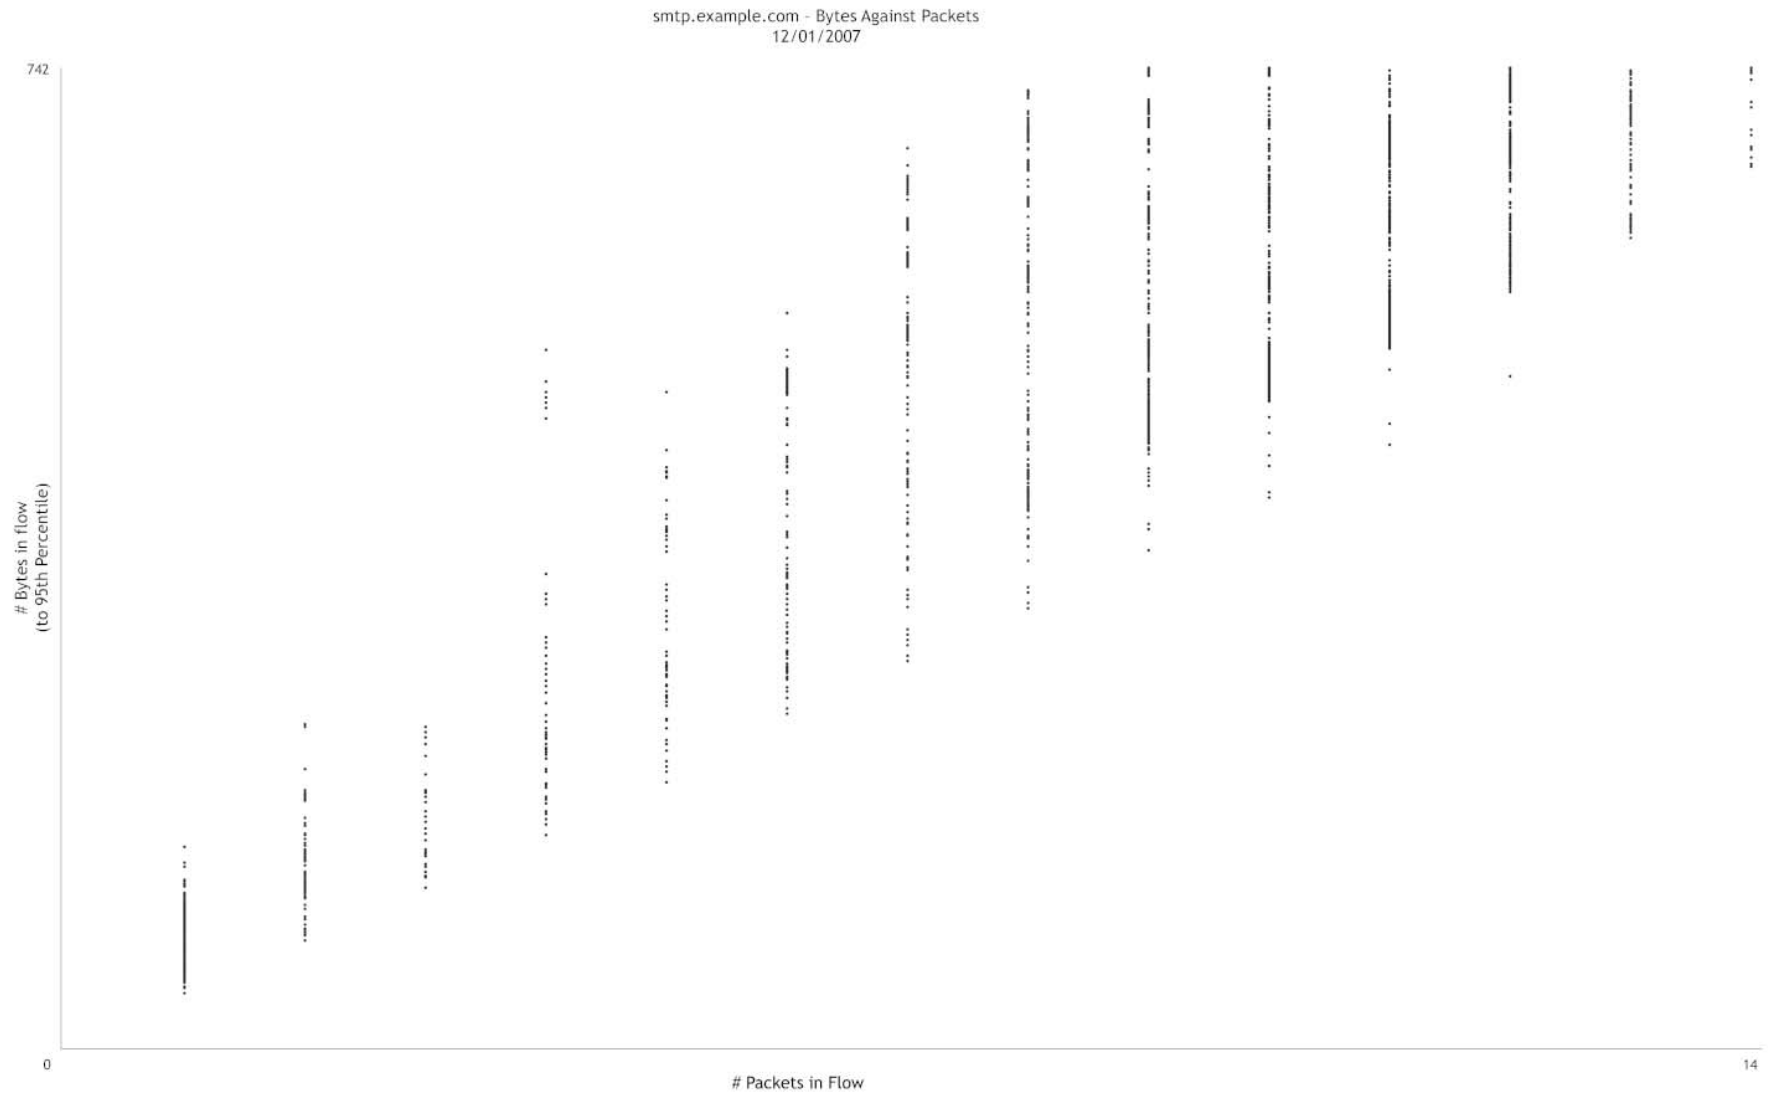
A violin plot (Credit: NIST)



example.com Flow Volume, Binned by Bytes per Packet
2007/12/01

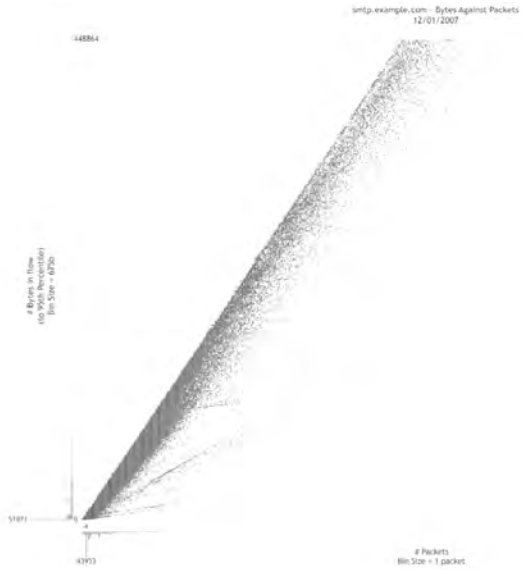


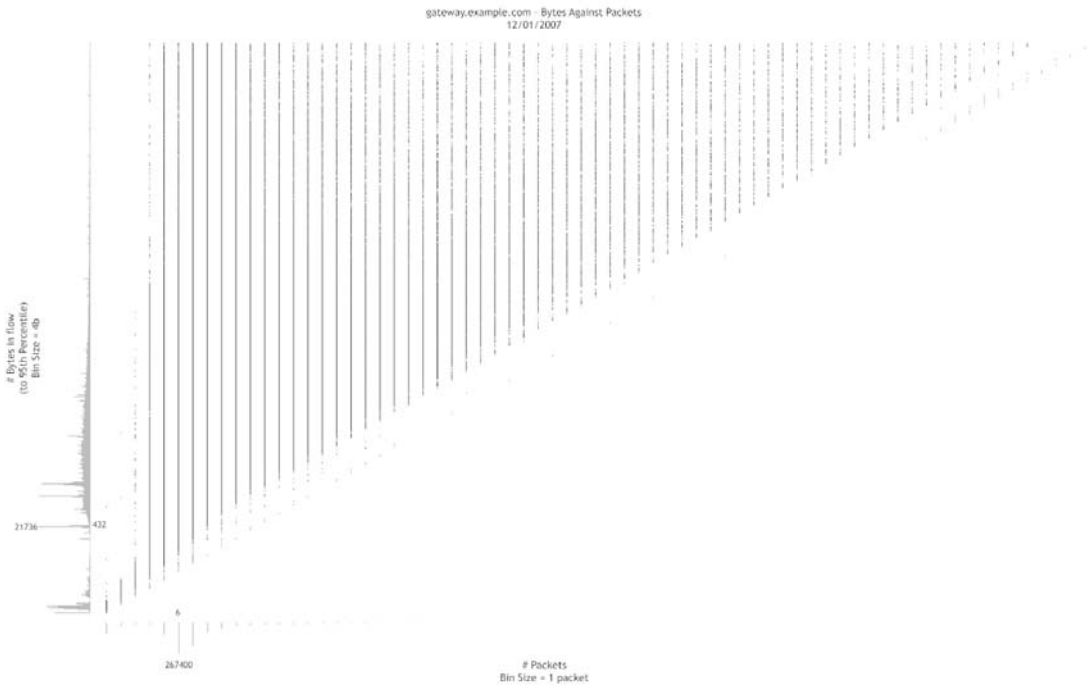
The X axis is bytes per packet in 5-byte increments. The Y axis shows the quantity of flows in each bin. Red indicates flow activity by known scanners.


smtp.example.com - Bytes Against Packets
12/01/2007
742
Bytes in flow
(to 95th Percentile)
0
14
Packets in Flow


smtp.example.com - Bytes Against Packets
12/01/2007
742
Bytes in flow
(to 95th Percentile)
0
Packets in Flow
14


smtp.example.com - Bytes Against Packets
12/01/2007
664
642
606
449
85
86635
56
35
Bytes in flow
(to 95th Percentile)
Packets in Flow

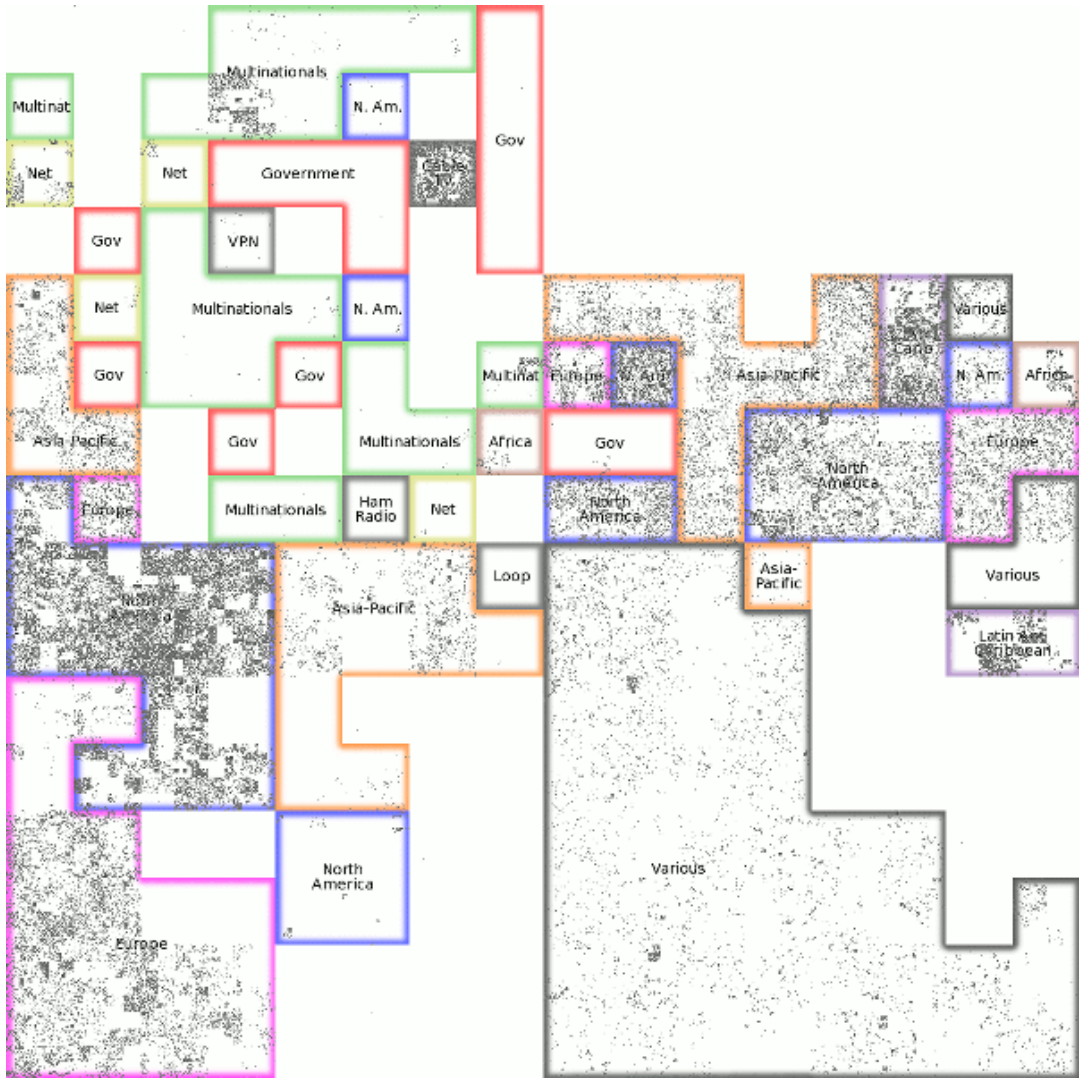

smtp.example.com - Bytes Against Packets
12/01/2007
Bytes in flow
(to 95th Percentile)
Bin Size = 672b
Packets
Bin Size = 1 packet


gateway.example.com - Bytes Against Packets
12/01/2007
Bytes in flow
(to 95th Percentile)
Bin Size = 4b
21736
432
6
267400
Packets
Bin Size = 1 packet

Movies:

“Escaping the Flat Land”

Hilbert Curve (The Movie)



Divider Slide:

Slide format can be used to begin new sub-section

Revisiting the Threshold Random Walk Scan Detector

Vagishwari Nagaonkar and
John McHugh
Faculty of Computer Science
Dalhousie University

Presented at FLOCON 2008

DALHOUSIE UNIVERSITY
Inspiring Minds

Objectives

- Analyze year long trace collected using SiLK tools from a /22 enterprise network for scanning activity.
- Use Threshold Random Walk (TRW), one of the most effective algorithms for early scan detection, to detect the scanning activity.
- Find out, if using Bloom filters along with TRW, sequentially, can we detect the scanners, that went undetected using only TRW ?
- And we shall soon be enlightened …. ☺

DALHOUSIE UNIVERSITY
Inspiring Minds

Rationale

- TRW is very effective, but has some problems:
  - In case of slow or stealthy scanning?
  - In case of UDP or ICMP?
  - In case of repetitive scanning?
- Using Bloom filter to eliminate repetitive input to TRW and look for reverse matches in time ordered data.
  - Can we detect the slow scans?
  - Can we detect UDP and ICMP scans?
  - Can we score ICMP responses to non ICMP?
- Lets see ….

DALHOUSIE UNIVERSITY
Inspiring Minds

Threshold Random Walk

- Scan Detection Algorithm based on sequential hypothesis testing.
- Uses a positive reward based scan detection.
  - For a given host, keeps a ratio which
    - In case of successful connection, is decreased
    - In case of unsuccessful connection, is increased.
  - This ratio is compared with two thresholds
    - If it goes above one, then it's a scanner
    - If it goes below the other, then it's benign
    - If it goes neither way, i.e., is in between the two thresholds, then can't say

DALHOUSIE UNIVERSITY
Inspiring Minds

Threshold Random Walk

- **The ratio is calculated as :**

$$\Lambda(Y) \equiv \frac{\Pr[Y|H_1]}{\Pr[Y|H_0]} = \Pi_{i=1}^{n} \frac{\Pr[Y_i|H_1]}{\Pr[Y_i|H_0]}$$

- **Where the probabilities are :**

$$\Pr[Y_i = 0|H_0] = \theta_0, \quad \Pr[Y_i = 1|H_0] = 1 - \theta_0$$
$$\Pr[Y_i = 0|H_1] = \theta_1, \quad \Pr[Y_i = 1|H_1] = 1 - \theta_1$$

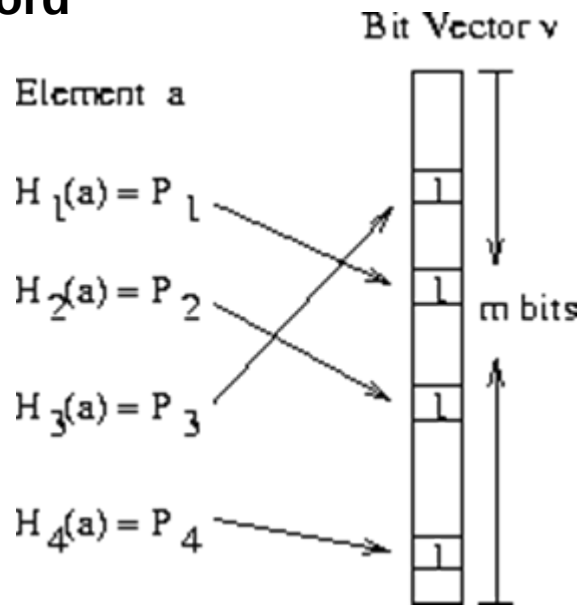
  - **Y = success (0) or failed (1) connection attempt**
  - **H0 = benign hypothesis**
  - **H1 = scanner hypothesis**
  - **Θ0 = probability that the source is benign, for a successful connection attempt**
  - **Θ1 = probability that the source is scanner for a successful connection attempt**
- **The thresholds are calculated based on**
  - **desired true positive (β = 0.99)**
  - **desired false positive (α = 0.01)**

$$\eta_1 \leftarrow \frac{\beta}{\alpha} \qquad \eta_0 \leftarrow \frac{1-\beta}{1-\alpha}$$

DALHOUSIE UNIVERSITY
Inspiring Minds

Bloom Filter

- **Data structure used to test the membership of an element for a given set. Uses bit array to record multiple hash values per element.**
- **Basic properties of these filters are :**
  - **False positives possible, but no false negatives.**
  - **Elements can be added to the set, but cannot be removed.**
  - **The higher the percentage of set bits, higher the probability of false positives.**
  - **Space efficient as compared to other set membership testing methods.**
  - **Cannot be reverse engineered to find the set of elements present in it.**



DALHOUSIE UNIVERSITY
Inspiring Minds

TRW + Bloom

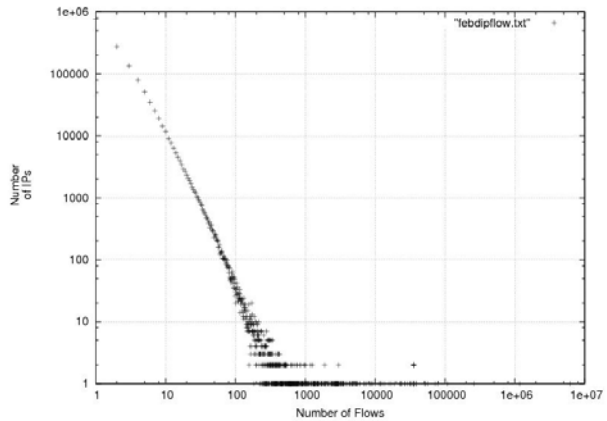
- **TRW hit or miss definition modified**
  - **For a given tuple in the flow record eg {sip, dip}**
    - **HIT = if a corresponding response entry {dip,sip} is found within a specified timeout period**
    - **MISS = if a corresponding entry {dip,sip} is not found within a specified timeout period**
- **Bloom Filter uses 10 hash functions and a bit vector of size 2^32**
- **Simple Set up :**
  - **Pass the flow records through the bloom filter to get unique entries for a given specified tuple**
  - **Different tuple combinations used are {sip,dip}, {sip,dip,protp}, {sip,dip,sport}, {sip,dip,dport}, {sip,dip,sport,dport,proto}**
  - **Then analyze the output using the TRW scanning algorithm.**

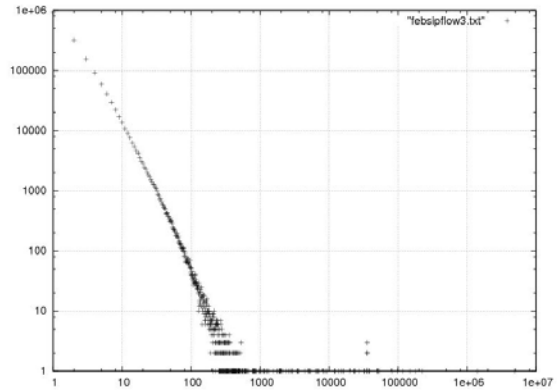
DALHOUSIE UNIVERSITY
Inspiring Minds

The Dataset

	Out IPs Seen		Non-responsive Out IPs	
	EtoO	OtoE	Count	Percent
Feb	26680	7270	19410	72.75
Mar	30232	3866	26366	87.21
Apr	56126	14576	41550	74.03
May	2355612	106893	2248719	95.46
June	2847371	283270	2564101	90.05
July	2601834	246312	2355522	90.53
Aug	30181	29097	1084	3.59
Sept	126913	126549	364	0.29
Oct	330740	277438	53302	16.124
Nov	4050	2932	1118	27.60
Dec	2226535	254484	1972051	88.57
Total	**10636274**	**1352687**	**9283587**	**87.28**

- #Inside Hosts
  - Total Address Space = 1024
  - #Active hosts in a given day = varies between 60-70
  - Active Address Space ~ 6%
- Plot shows flows per number of IP addresses.
- How easy it is to defeat TRW in this network?

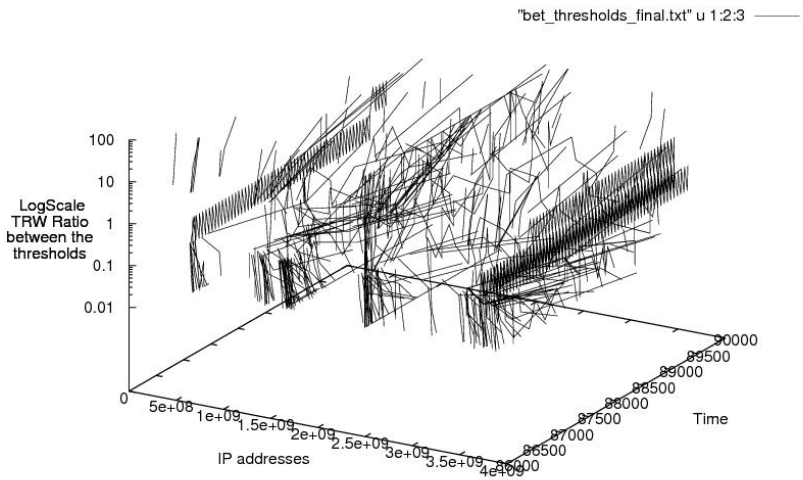




DALHOUSIE UNIVERSITY
Inspiring Minds

Issues to keep in mind

- Number of packets per flow record >= 1
- The time granularity is only till sec, millisecond not available.
- For a packets received in the same sec, the order of the flow records is the outside to inside seen first always, irrespective of the actual order.
- Background noise in the traffic.
- ICMP ping traffic causes false detection.



sIP	dIP	sPort	dPort	pro	packets	bytes	flags	sTime	dur
24.222.0.20	134.190.64.203	0	0	1	5	420		2006/03/01T00:00:46.000	4.000
134.190.64.203	24.222.0.20	0	2048	1	5	420		2006/03/01T00:00:46.000	4.000
24.222.0.20	134.190.64.203	110	33636	6	6	373	FS PA	2006/03/01T00:00:46.000	0.000
134.190.64.203	24.222.0.20	33636	110	6	8	394	FSRPA	2006/03/01T00:00:46.000	0.000
24.222.0.20	134.190.64.203	0	0	1	5	420		2006/03/01T00:01:46.000	4.000
134.190.64.203	24.222.0.20	0	2048	1	5	420		2006/03/01T00:01:46.000	4.000
24.222.0.20	134.190.64.203	110	33647	6	6	373	FS PA	2006/03/01T00:01:46.000	0.000
134.190.64.203	24.222.0.20	33647	110	6	7	354	FSRPA	2006/03/01T00:01:46.000	0.000
24.222.0.20	134.190.64.203	0	0	1	5	420		2006/03/01T00:02:46.000	4.000
134.190.64.203	24.222.0.20	0	2048	1	5	420		2006/03/01T00:02:46.000	4.000
24.222.0.20	134.190.64.203	110	33661	6	6	373	FS PA	2006/03/01T00:02:46.000	0.000
134.190.64.203	24.222.0.20	33661	110	6	8	394	FSRPA	2006/03/01T00:02:46.000	0.000

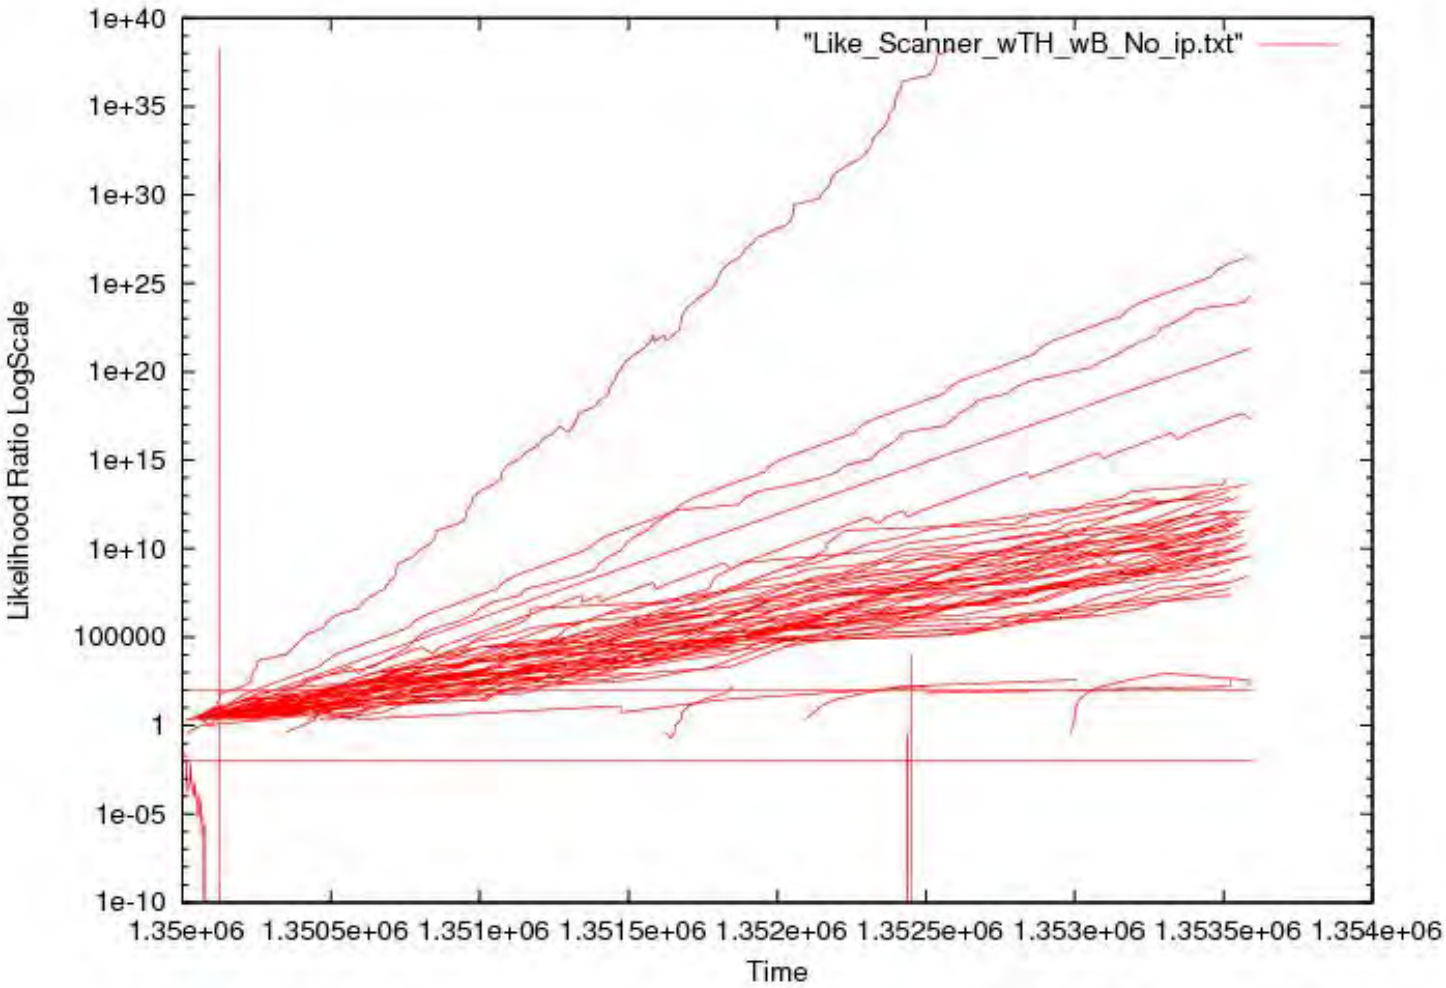
DALHOUSIE UNIVERSITY
Inspiring Minds

Preliminary Results

- Used theta0=0.7 and theta=0.3 for the TRW algorithm.
- Kept a timeout period of 10 sec.
- Max number of false positives due to ICMP
  - Is this a ping tunnel ???
  - Lots of out ips contacting single in ip with lots of ping requests and getting responses, effectively lots of bytes being transferred.
- The SD and SDP options using bloom detect horizontal scans.
- The SDSP and SDDP options detect vertical scans.
- The SDSDP covers both.

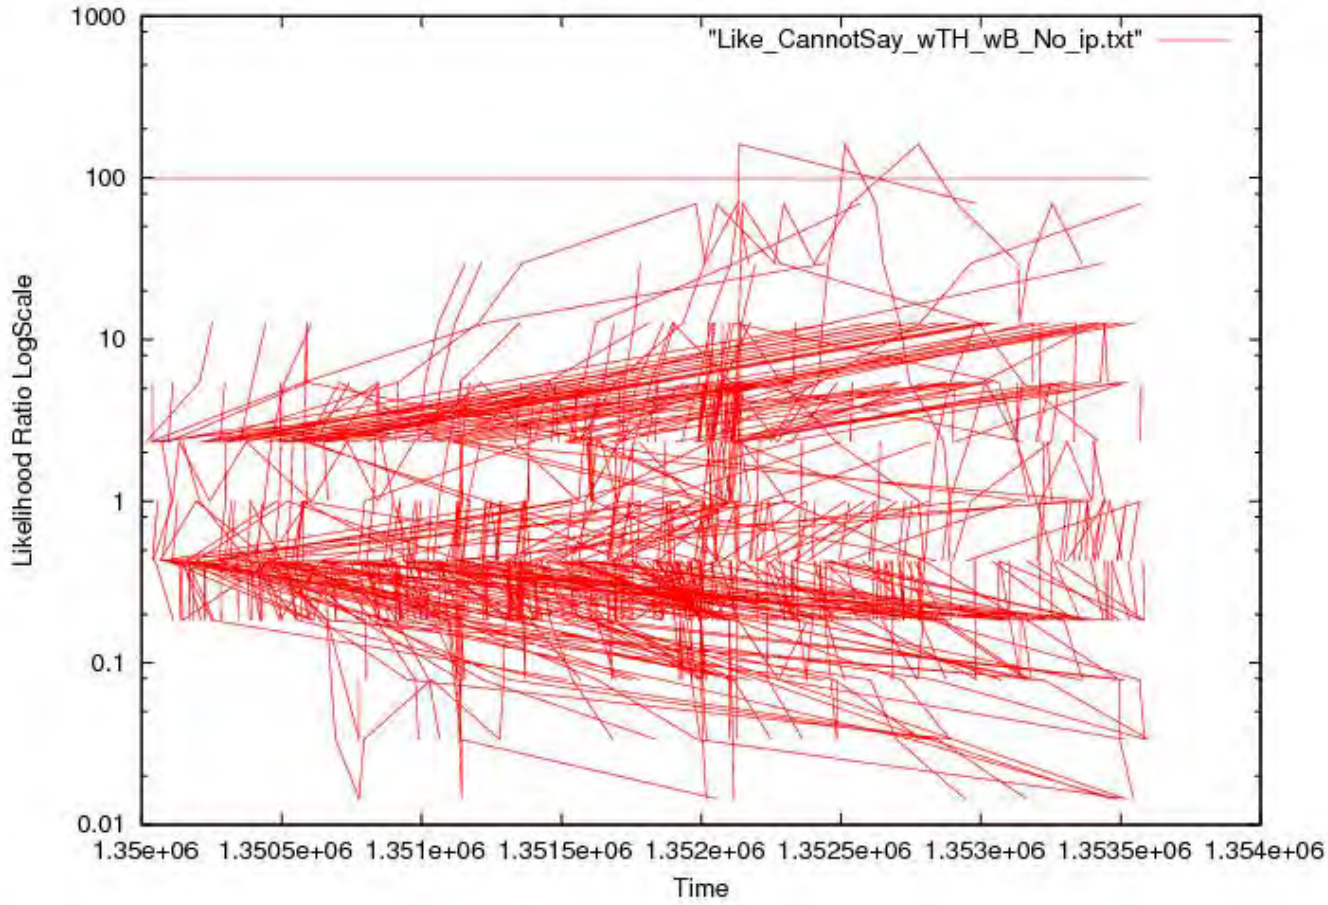
For Hourly File	#Scanners Detected
All	53
All_no_icmp	14
All_SD	3
All_SDP	3
All_SDSP	16
All_SDDP	4
All_SDSDP	5

Preliminary Results
Plot of Likelihood ration for Scanners



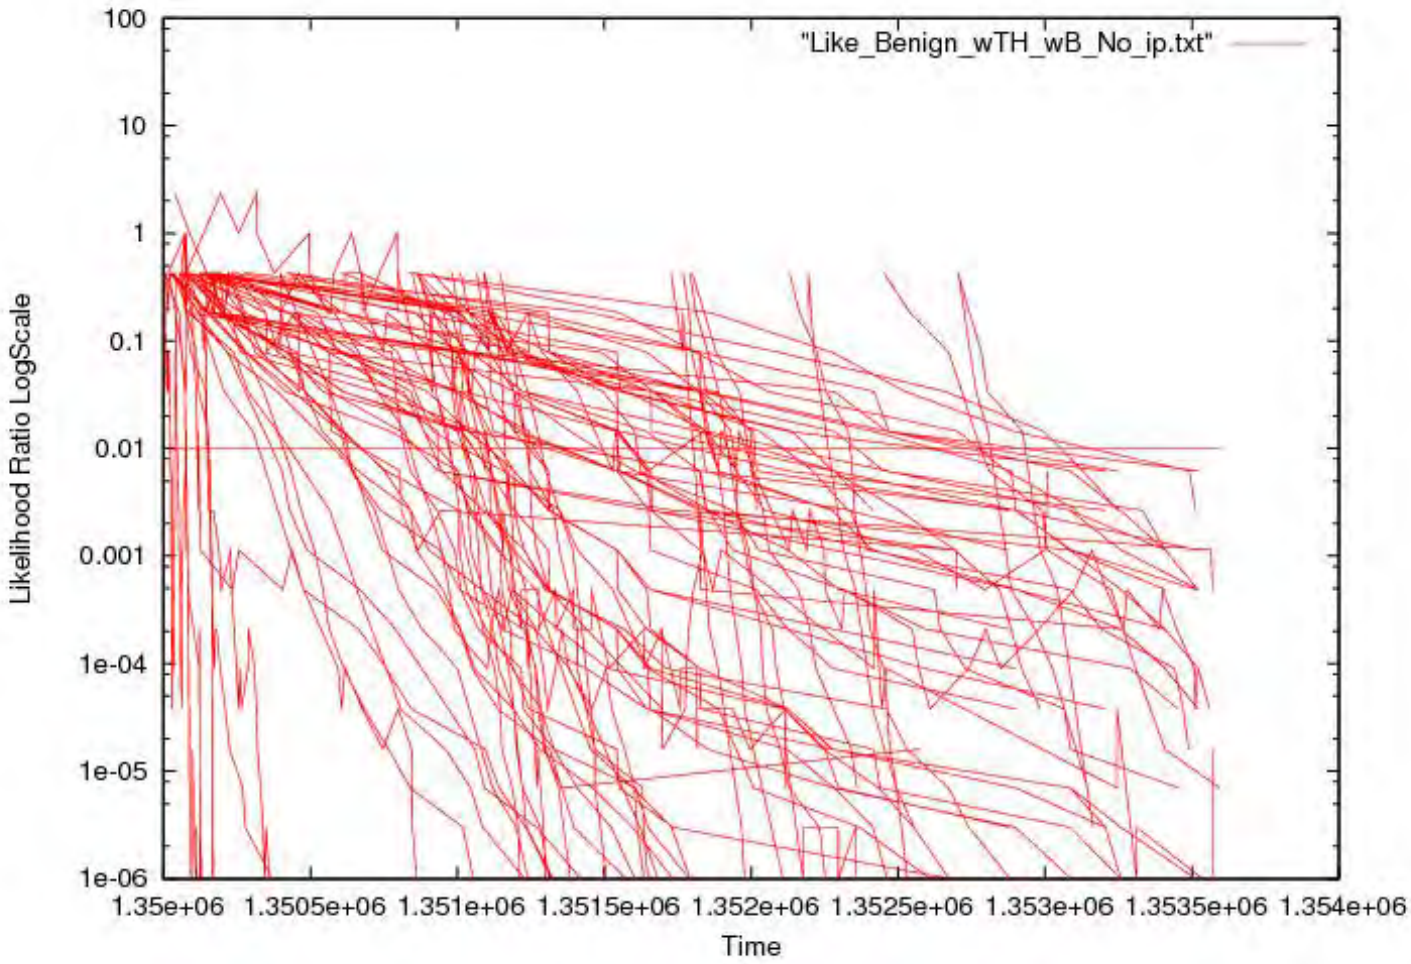
DALHOUSIE UNIVERSITY
Inspiring Minds

Preliminary Results
Plot of Likelihood ration for “Can’t Say”s



DALHOUSIE UNIVERSITY
Inspiring Minds

Preliminary Results
Plot of Likelihood ration for Benign



DALHOUSIE UNIVERSITY
Inspiring Minds

Initial Conclusions

- The Bloom filter ▪ ▪ ▪ somewhat reduces the false positives
  - only unique entries for given filter criteria considered by TRW.
- Using specific options for the bloom filter it is faster to detect vertical or horizontal scanning
- Need to improve the technique by
  - Checking for change in the theta0 and theta1 values effecting the overall results.
  - Check for real time scenarios.
- Some IPs go Scanner; then return to Can't Say.
- Still more data is left to be analysed (In progress)
- Certain issues mentioned earlier need to be taken care of e.g dealing with the number of packets per flow

DALHOUSIE UNIVERSITY
Inspiring Minds

References

- http://pages.cs.wisc.edu/~cao/papers/summary-cache/node8.html
- http://en.wikipedia.org/wiki/Bloom_filter
- Distributed Evasive Scan Techniques and Countermeasures - Min Gyung Kang[1] , Juan Caballero[1] and Dawn Song[1]
- **The Limits of Global Scanning Worm Detectors in the Presence of Background Noise -** David W. Richardson, Steven D. Gribble, and Edward D. Lazowska
- Scan Detection: A Data Mining Approach - GyÄorgy J. Simon, Hui Xiong, Eric Eilertson, Vipin Kumar
- **DNS-based Detection of Scanning Worms in an Enterprise Network -** David Whyte, Evangelos Kranakis, P.C. van Oorschot
- **Very Fast Containment of Scanning Worms - Nicholas Weaver, Stuart Staniford, Vern Paxson**
- **Fast Portscan Detection Using Sequential Hypothesis Testing -** Jaeyeon Jung, Vern Paxson, Arthur W. Berger, and Hari Balakrishnan
- Many more …

Acknowledgments

- Ron McLeod
- The Faculty of Computer Science
- Dalhousie University

DALHOUSIE UNIVERSITY
Inspiring Minds

Network Analysis of Point of Sale System Compromises

Operation Terminal Guidance

Chicago Electronic & Financial Crimes Task Force

U.S. Secret Service

U.S. Secret Service

Outline

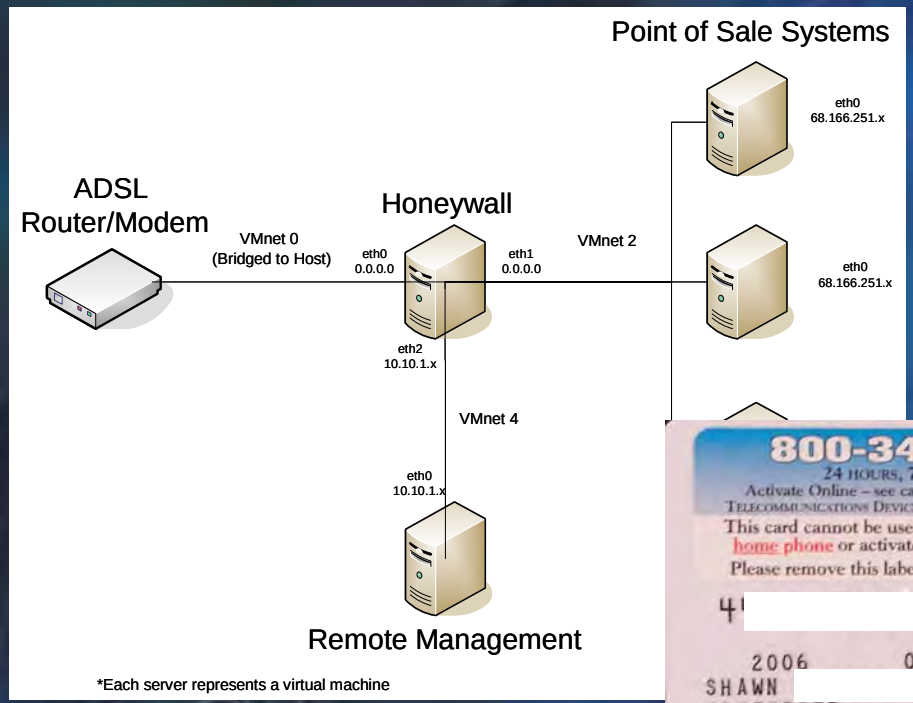
- Background
- Hypothesis
- Deployment Methodology
- Data Analysis
- Findings
- Discussion

Investigative Goals

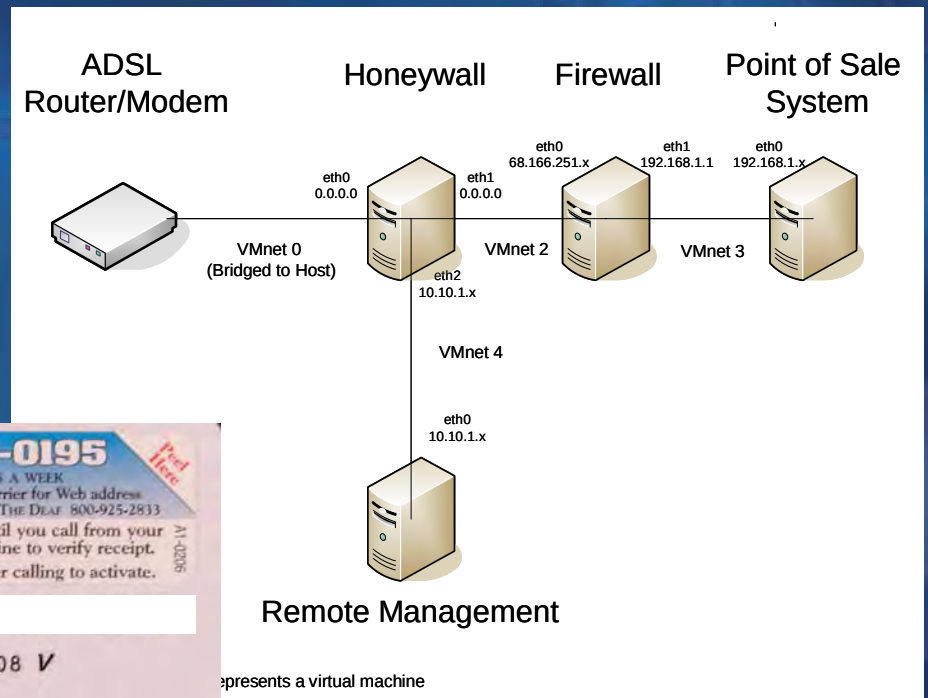
- Hypothesis: Remote attackers were not targeting point of sale (POS) system software.
  - The underlying operating system and installed applications are not deployed in accordance with Payment Card Industry Data Security Standard
  - POS system compromises are a result of automated scanning and vulnerability exploitation

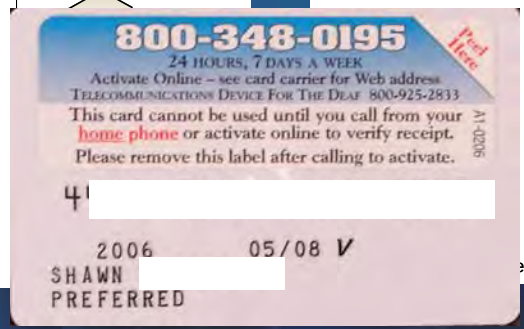
Deployment Methodology

Test Group Honeynet



Control Group Honeynet



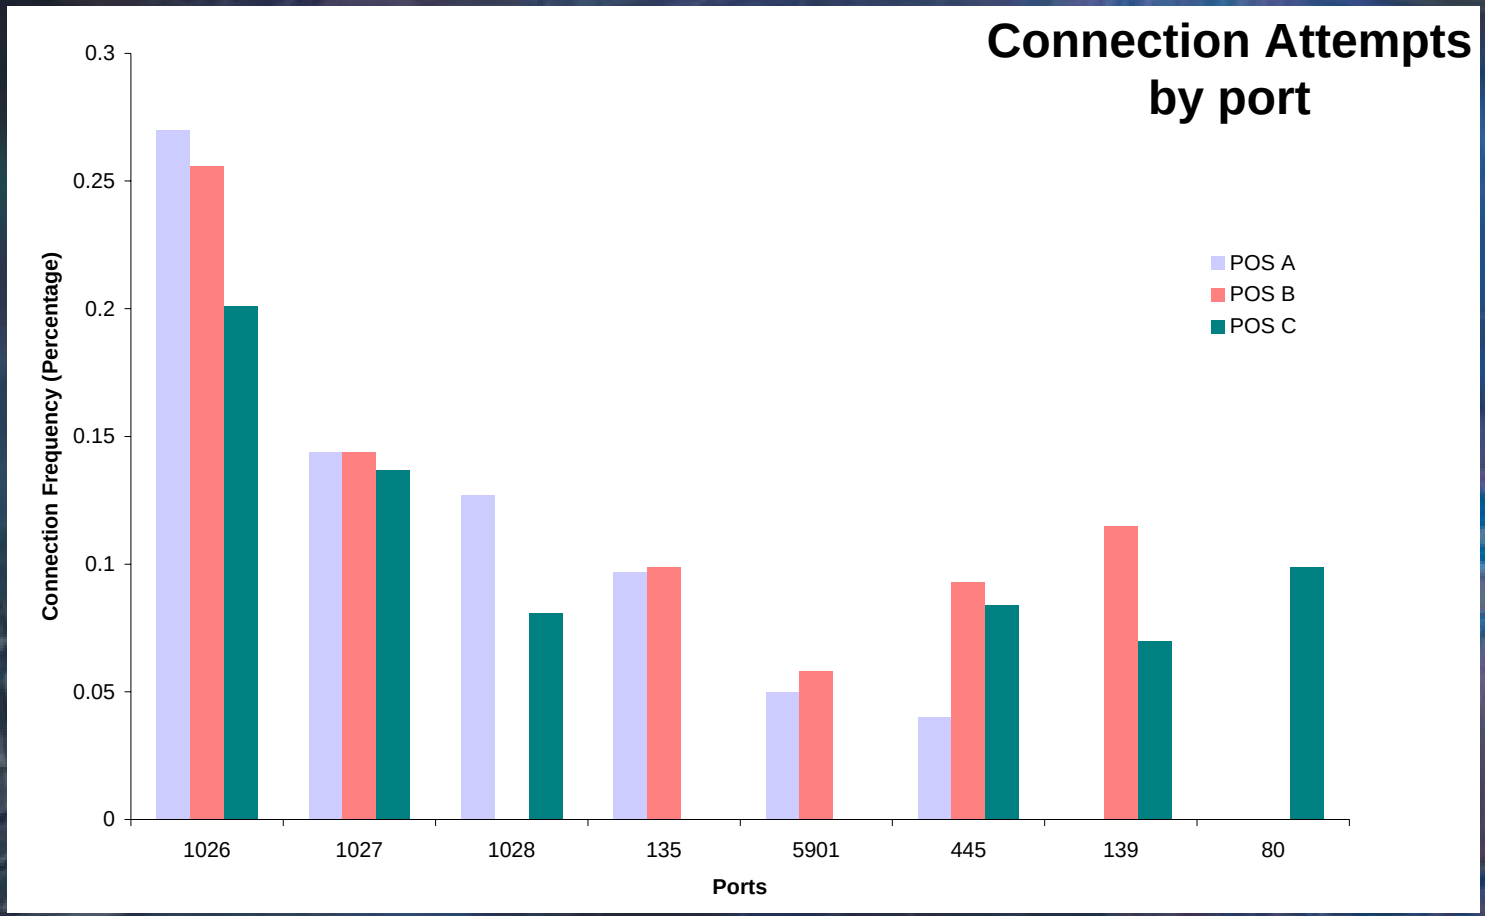


Honeytoken

U.S. Secret Service

Data Analysis

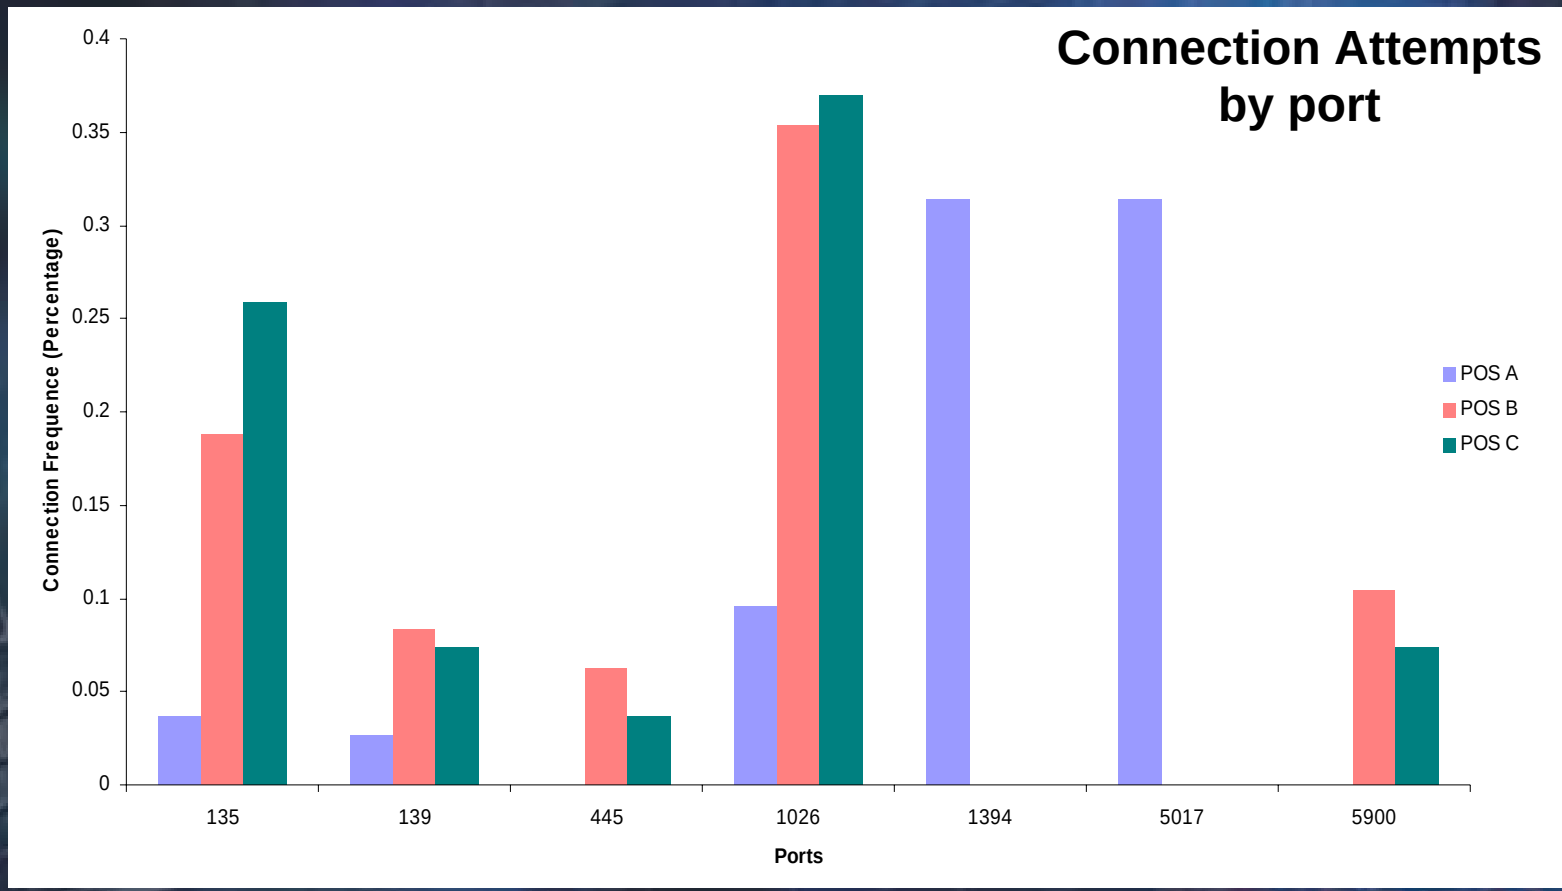
Control Group



U.S. Secret Service

Data Analysis

Test Group



U.S. Secret Service

Data Analysis

- Association rules
  - Clustering
    - T: Number of virtual POS systems with connection attempts from a single source
    - n_i: Number of packets from a source to a virtual POS system
    - N: Total number of packets from a source to all three POS systems
    - $N=\sum n_i$

$$\text{Support}(R) = \frac{\#\text{ connections (POS system A, B, and C)}}{\#\text{connections}}$$

Data analysis methodology from F. Pouget and M. Dacier. "Honeypot Based Forensics."

U.S. Secret Service

Data Analysis

Control Group Clusters

Port	Item Sets	Support %	Support % > 1%
80	Cluster 1: T=1, N=3 Cluster 2: T=1, N=1 Cluster 3: T=2, N=8 (n=5, n=3)	43.5% 10.9% 4.3%	1
135	Cluster 4: T=1, N=1 Cluster 5: T=1, N=2	54.5% 22%	2
139	Cluster 6: T=1, N=2 Cluster 7: T=1, N=3	75% 10.1%	1
445	Cluster 8: T=1, N=1 Cluster 9: T=1, N=2 Cluster 10: T=1, N=3	20% 70% 7.1%	2
1026	Cluster 11: T=1, N=1	53.5%	1
1027	Cluster 12: T=1, N=1	98%	1
1028	Cluster 13: T=1, N=1	83%	1
5901	Cluster 14: T=1, N=2	90.9%	1

Data Analysis

Test Group Clusters

Port	Item Sets	Support %	Support % > 1%
445	Cluster 1: T=2, N=34	22.2%	0
1026	Cluster 2: T=2, N=3 Cluster 3: T=3, N=3 (n=1,n=1, n=1) Cluster 4: T=1, N=1	1.8% 20% 50.9%	2
1394	Cluster 5: T=1, N=12 Cluster 6: T=1, N=15 Cluster 7: T=1, N=6 Cluster 8: T=1, N=9	20% 16.7% 1.7% 16.7%	3
2967	Cluster 9: T=3, N=8 (n=2, n=3, n=3) Cluster 10: T=3, N=30 (n=10, n=10, n=10)	10% 10%	0
5900	Cluster 11: T=3, N=3	20%	0

U.S. Secret Service

Data Analysis

- Edit Distance Analysis
  - Extract TCP payloads from previous identified cluster members
  - Compare packets from each IP address against all others identified through clustering

Source A	Source B
<mss E..0..@.o.A.;W\. D..s.]..... p...^2..... <mss E..0..@.o.A.;W\. D..s.]..... p...^2.....	<mss E..0.{@.k.l\=.y. D..s.....jd..... p..... <mss E..0.{@.k.l\=.y. D..s.....jd..... p.....

Attack Phrases

U.S. Secret Service

Data Analysis

Control Group Phrase Distance

Cluster	Port	Phrase Distance (Lines)	Std Deviation
Cluster 6	139	2	9
Cluster 7	139	1	5
Cluster 8	445	3	10
Cluster 9	445	5	8
Cluster 10	445	4	18
Cluster 11	1026	86	169
Cluster 13	1028	12	65
Cluster 14	5901	32	12

***Clusters 1,2, 3,4,5, and 12 were discarded as not statistically significant

U.S. Secret Service

Data Analysis

Test Group Phrase Distance

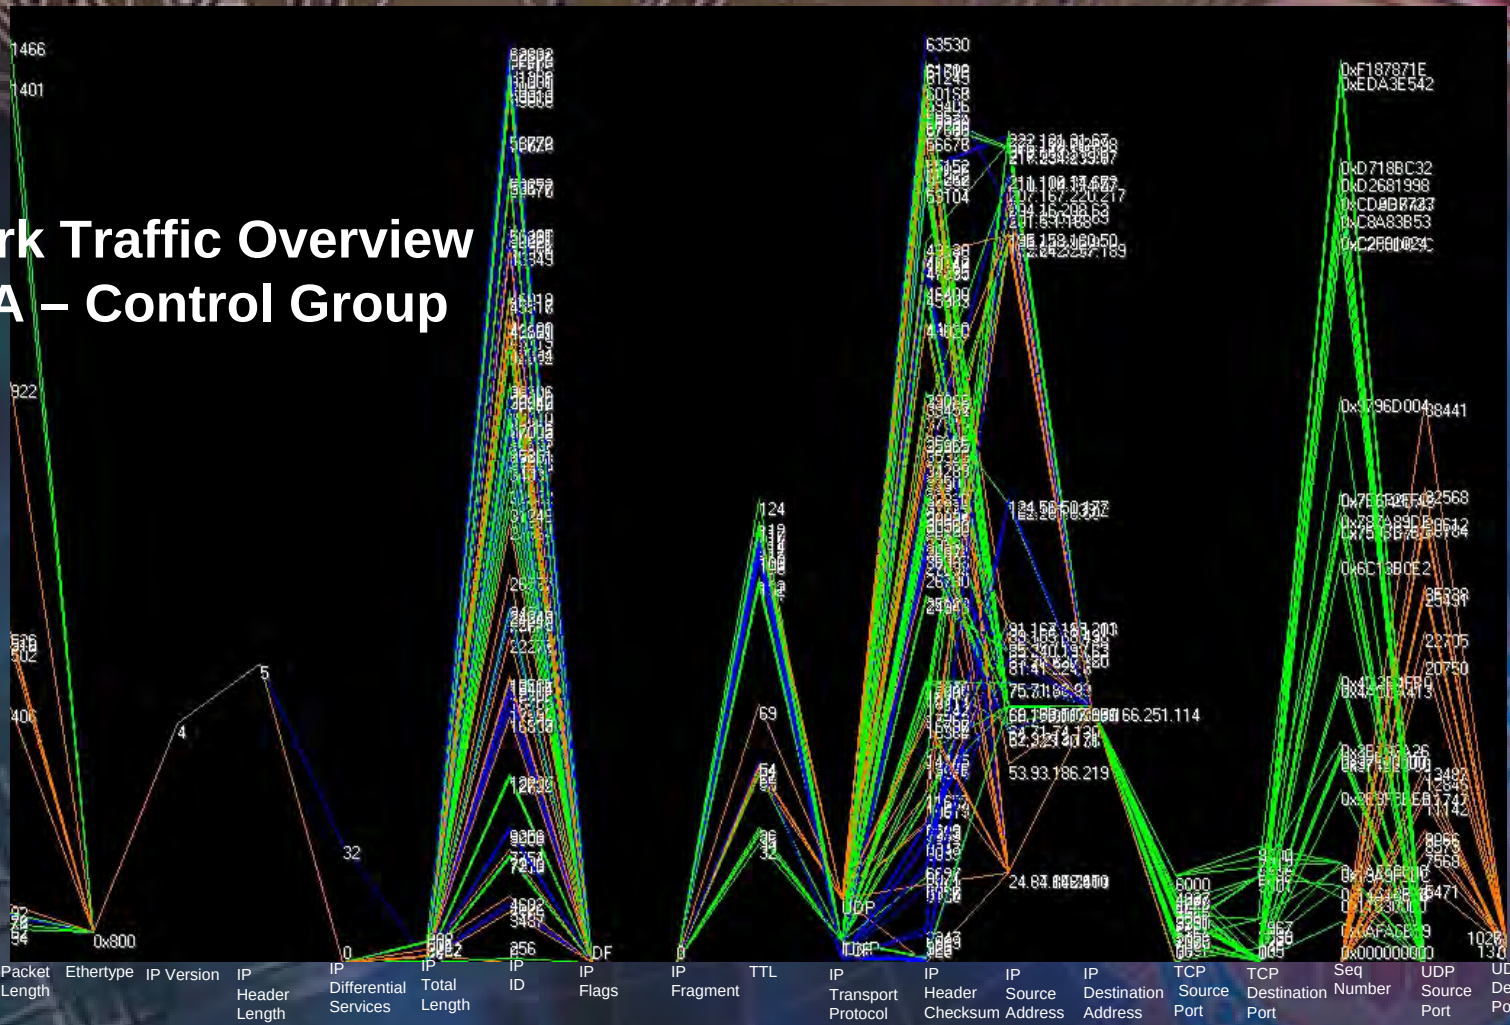
Cluster	Port	Phrase Distance (Lines)	Std Deviation
Cluster 2	1026	324	238
Cluster 5	1394	360	85
Cluster 6	1394	280	170
Cluster 7	1394	529	136
Cluster 8	1394	1422	1143
Cluster 11	5900	240	257

***Clusters 1,3,4,9,10 were discarded as not statistically significant

U.S. Secret Service

Data Analysis

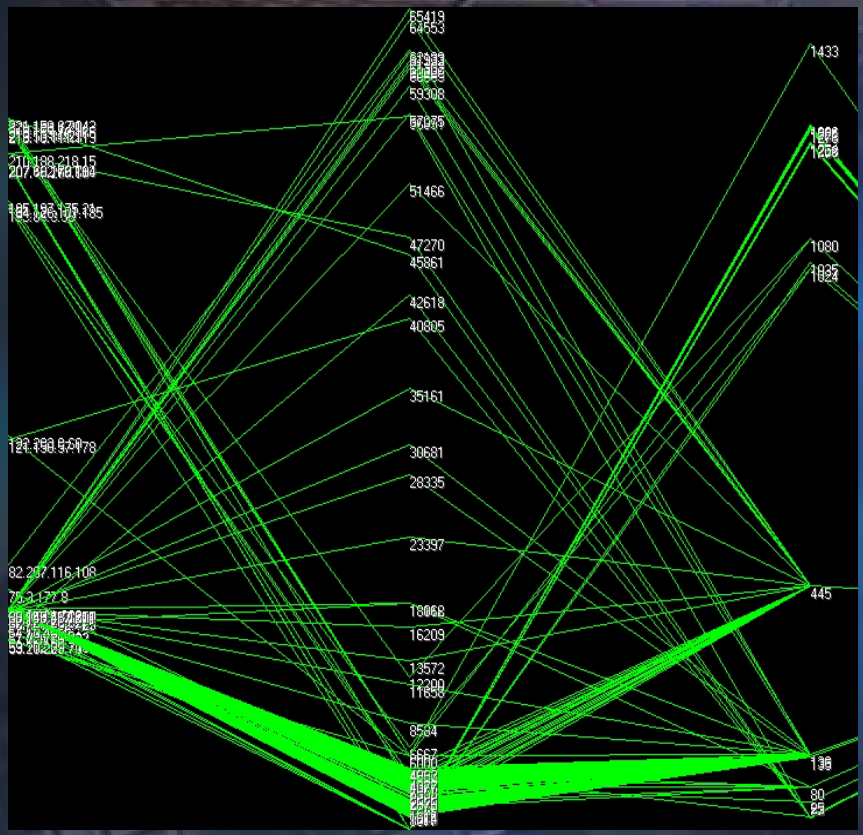
Network Traffic Overview
POS A – Control Group



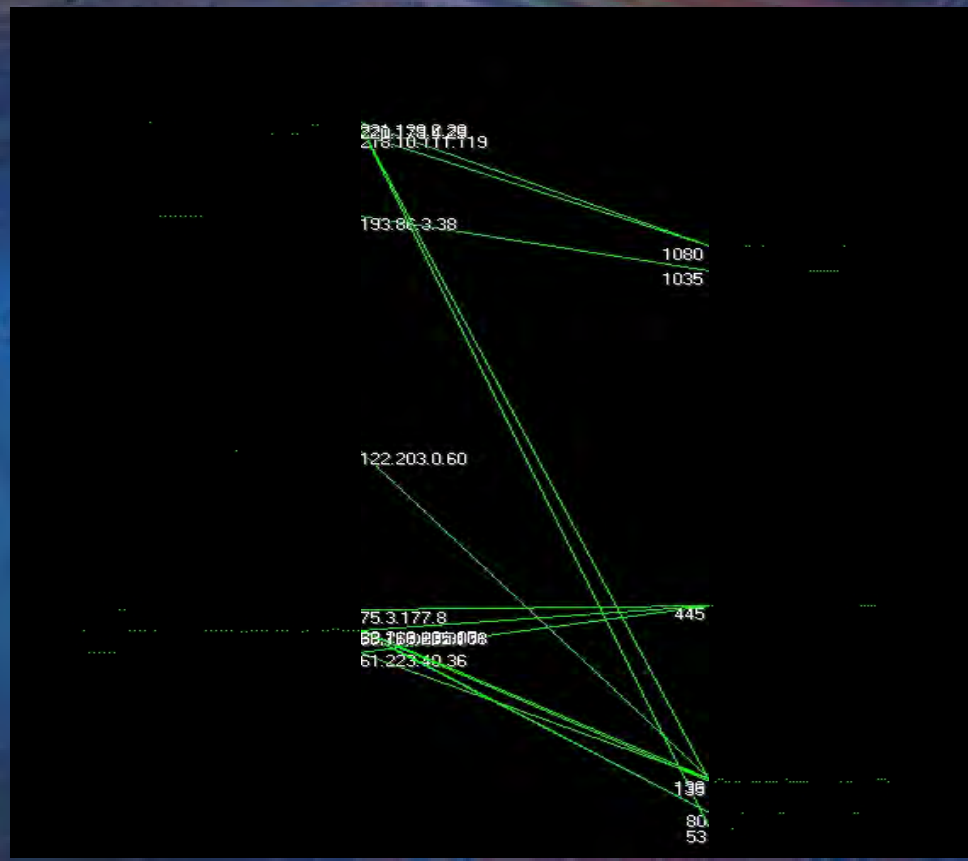
Visualization methodology from Greg Conti's. "Security Data Visualization."

U.S. Secret Service

Data Analysis



Source IP | TCP Source Port | TCP Destination Port

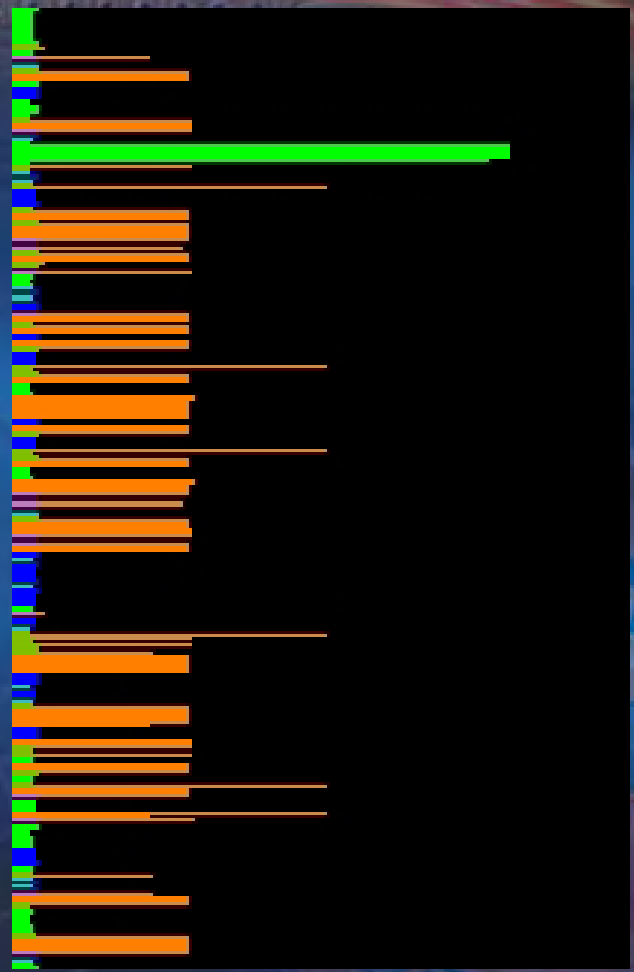


Source IP | TCP Destination Port

U.S. Secret Service

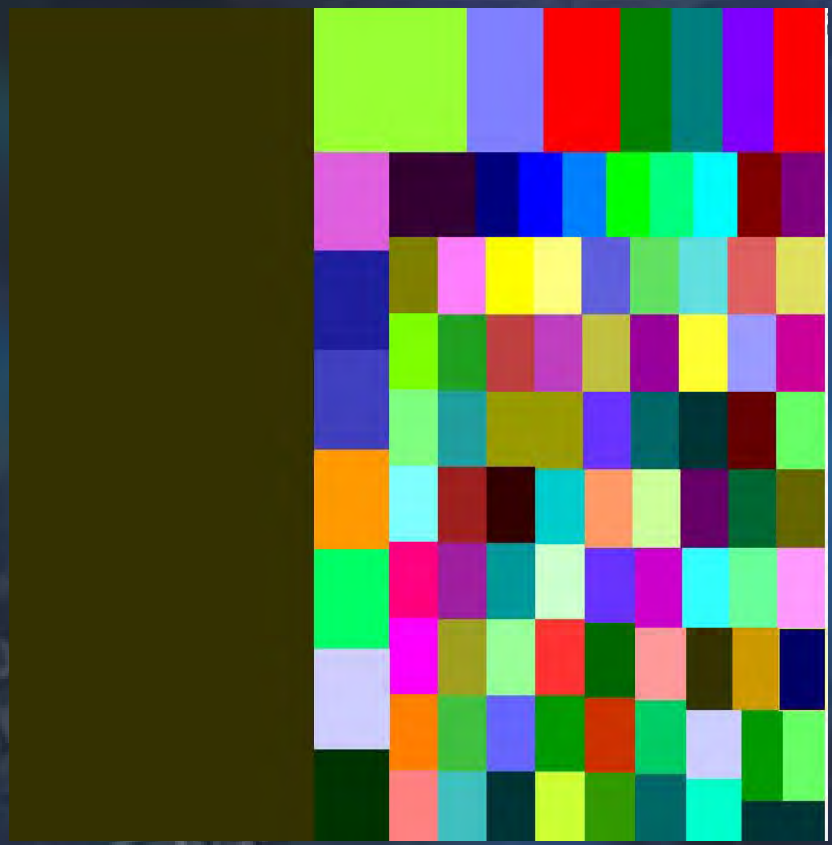
Data Analysis

- The TCP outlier is associated with browsing public web site to ensure connectivity
- Uniform length of packets

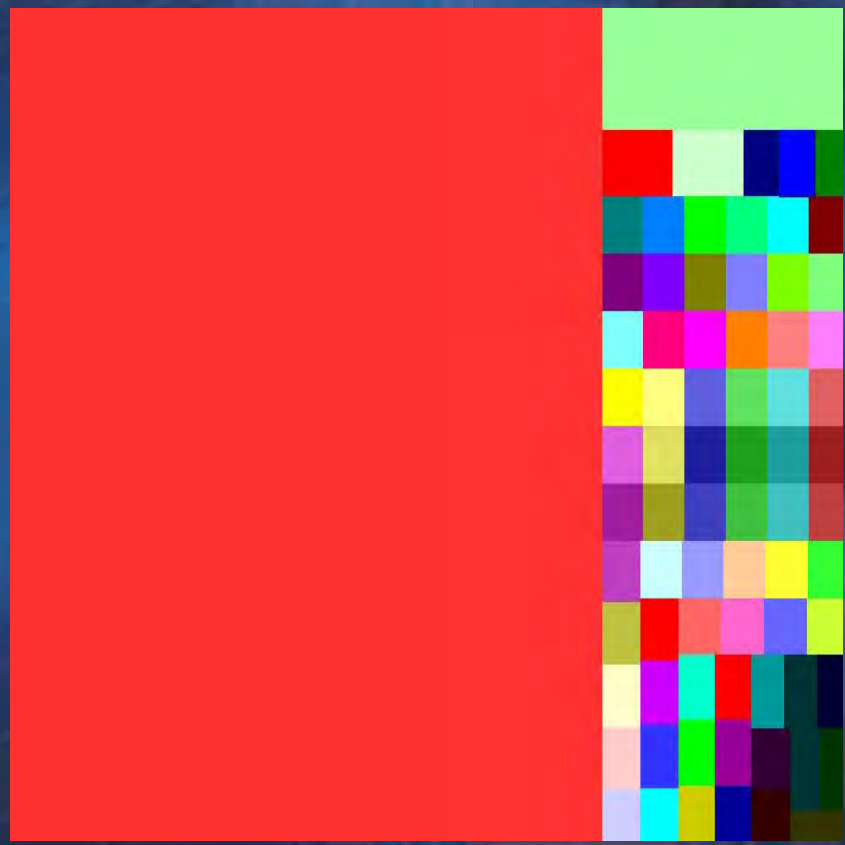


Data Analysis

TCP Packet Tree Map

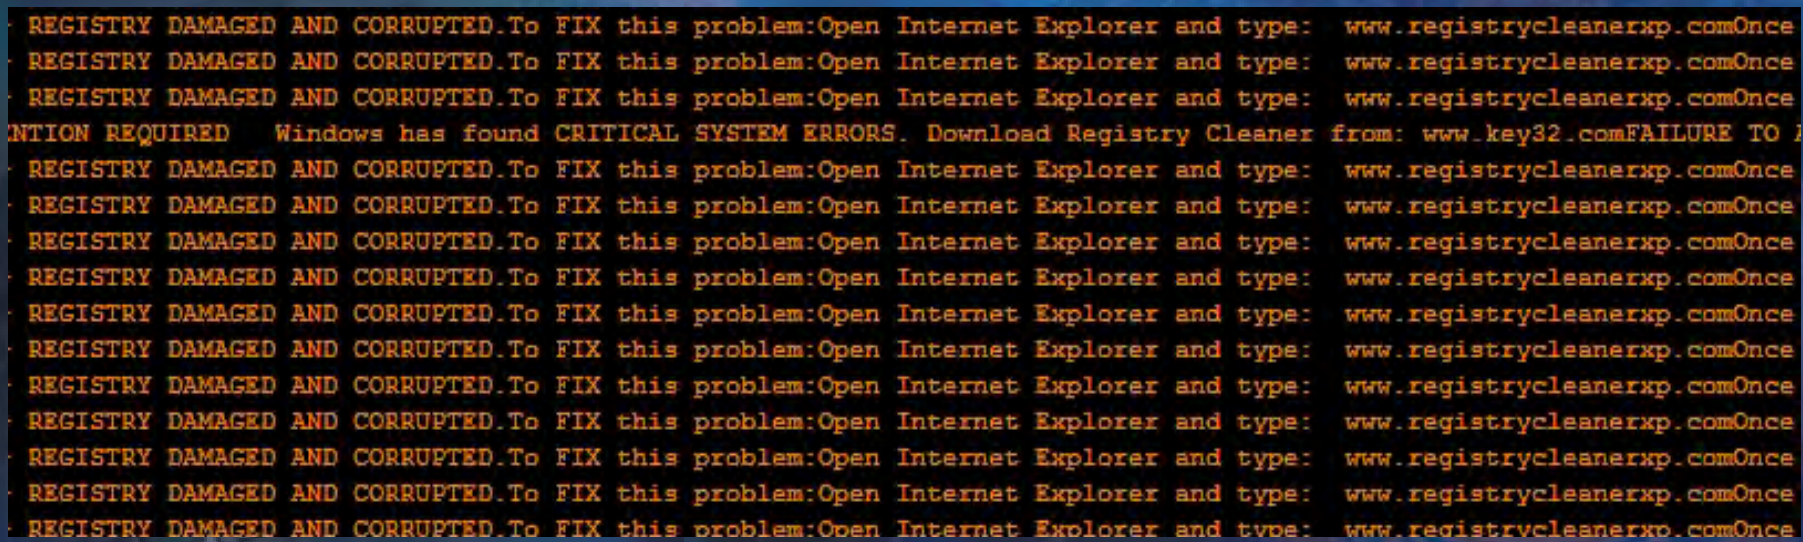


UDP Packet Tree Map



Data Analysis

- Examination of the UDP packets identified in the previous tree map revealed them to be spam targeting messenger applications



U.S. Secret Service

Findings

- Automated scanning of select set of ports
- Multiple exploits targeting multiple OS's from single source IP address
- Attackers not aware compromised system is a POS system until after compromise and exploit
- Insecure installation of operating system and applications lead to compromise

Discussion

All references available upon request

Ryan E. Moore
Special Agent
U.S. Secret Service
312-353-5431
ryan.moore@usss.dhs.gov

U.S. Secret Service



Anonymizing Network Flow Data

Timothy J. Shimeall (tjs@cert.org)
January 2007
FloCon 2008

© 2008 Carnegie Mellon University

Overview

The balance of anonymization

Subnet-preserving

Subnet-collapsing

Host-preserving

Host-collapsing

Ports & Other issues

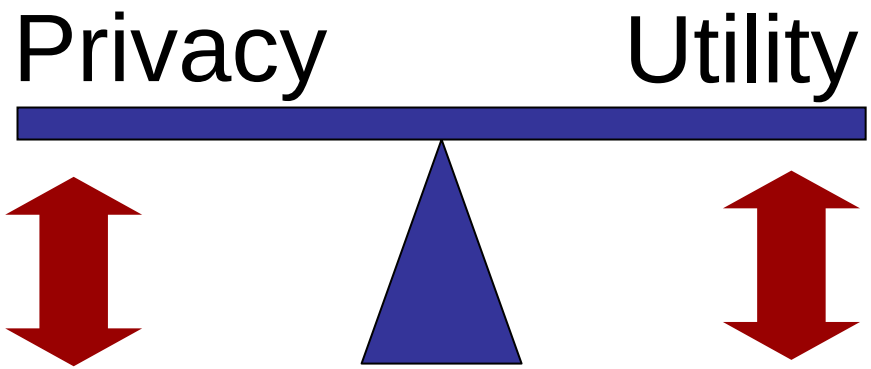
Conclusion

The Balance of Anonymization

Flow itself preserves some privacy by aggregation and eliding content.

Anonymization is to aid in preserving the privacy of organizations represented in the data

- Data owner
- Partner or Customer
- Incidental
- Attacker



The more you anonymize the data, the less analyses can be done with it.

Need to explore a range of options

Subnet Preserving

Preserve host identity while concealing network.

How:

- Prepare list of networks
- Assign random substitution for network prefix
- Mask and replace prefix on each address
- Associative array works well for substitutions

Balance:

- Enables analysis down to host identity, but not organization identity
- Can be reversed by outside knowledge (server suffixes)

~~248.204~~.5.3

010.005.5.3

Subnet Collapsing

Conceal network structure and host identity, but preserve commonality of network

How:

- Reduce all address to the network
- Prepare random substitution for network
- Replace address with network substitutions

Balance:

- Allows network-level behavior analysis
- Might be reversed by organizations with lots of contact with data source

248.204.5.3

248.204.0.0

010.005.0.0

Host Preserving

Preserve host identity while concealing network commonality

How:

- Generate list of addresses
- Generate random substitution for each address
- Replace each occurrence with same substitution

Balance:

- Allows host-specific analysis
- Difficult to reverse

~~248.204.5.3~~

10.2.3.9

~~248.204.5.12~~

192.168.12.7

Host Randomizing

Do not preserve host or network identity (a.k.a., remove address content in any useful way)

How:

- Replace each occurrence of each address with random value
- Allow repetition of random values

Balance:

- Only permit analysis that does not involve address information
- Extremely difficult to reverse

~~248.204.5.3~~ ~~128.0.3.2~~

10.2.3.7 192.168.7.12

~~248.204.5.3~~ ~~0.5.4.1~~

192.168.17.37 10.2.3.7

Ports and Other Issues

There's more to anonymization of flow than addresses

- Network ports can be very revealing (OS fingerprinting)
- Timing information might be revealing
- TCP flags might be revealing (odd patterns)

Can anonymize this information:

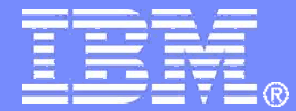
- Ports: reduce to service, substitute; reduce to common/reserved/dynamic
- Timing: restart epoch; rescale timing; collapse interval
- TCP flags: reduce to function; remove OS-dependencies

Conclusion

Data sharing is difficult

Anonymization can be useful, but limiting

Anonymized does not mean private or irreversible



Zurich Research Laboratory

Dynamic Adaptation of Flow Information Granularity for Incident Analysis

Marc Ph. Stoecklin <mtc@zurich.ibm.com>
Andreas Kind <ank@zurich.ibm.com>
Jean-Yves Le Boudec <jean-yves.leboudec@epfl.ch>

© 2008 IBM Corporation

Outline

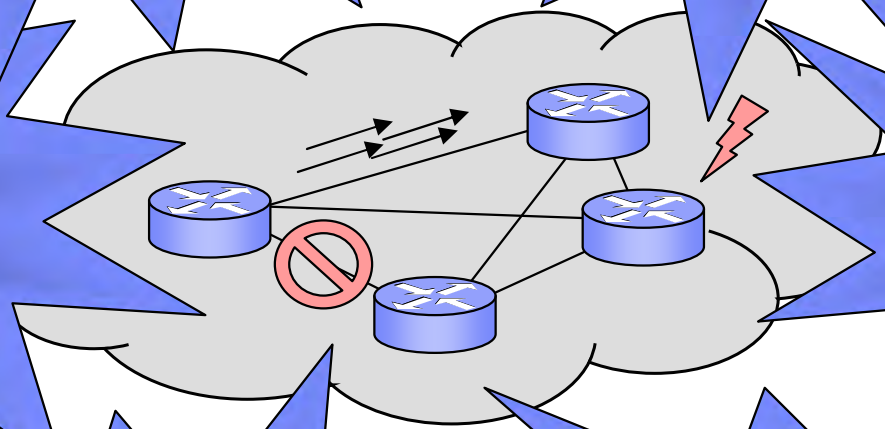
- Problem statement and objectives
- Adapting flow information granularity
  - Increasing granularity with zoom monitors
  - Decreasing granularity with relevance-sensitive compression
- Implementation
- Results
- Conclusion and outlook

© 2008 IBM Corporation

What is currently going on in my network?

Available sampled information is not sufficient to under-stand this incident!

What caused the surge of traffic to the mail server last Tuesday at 2pm?



What content is in the packets of this flow?

What was going on in my network before the incident?

What is causing this abnormal network activity?

© 2008 IBM Corporation

Problem Statement

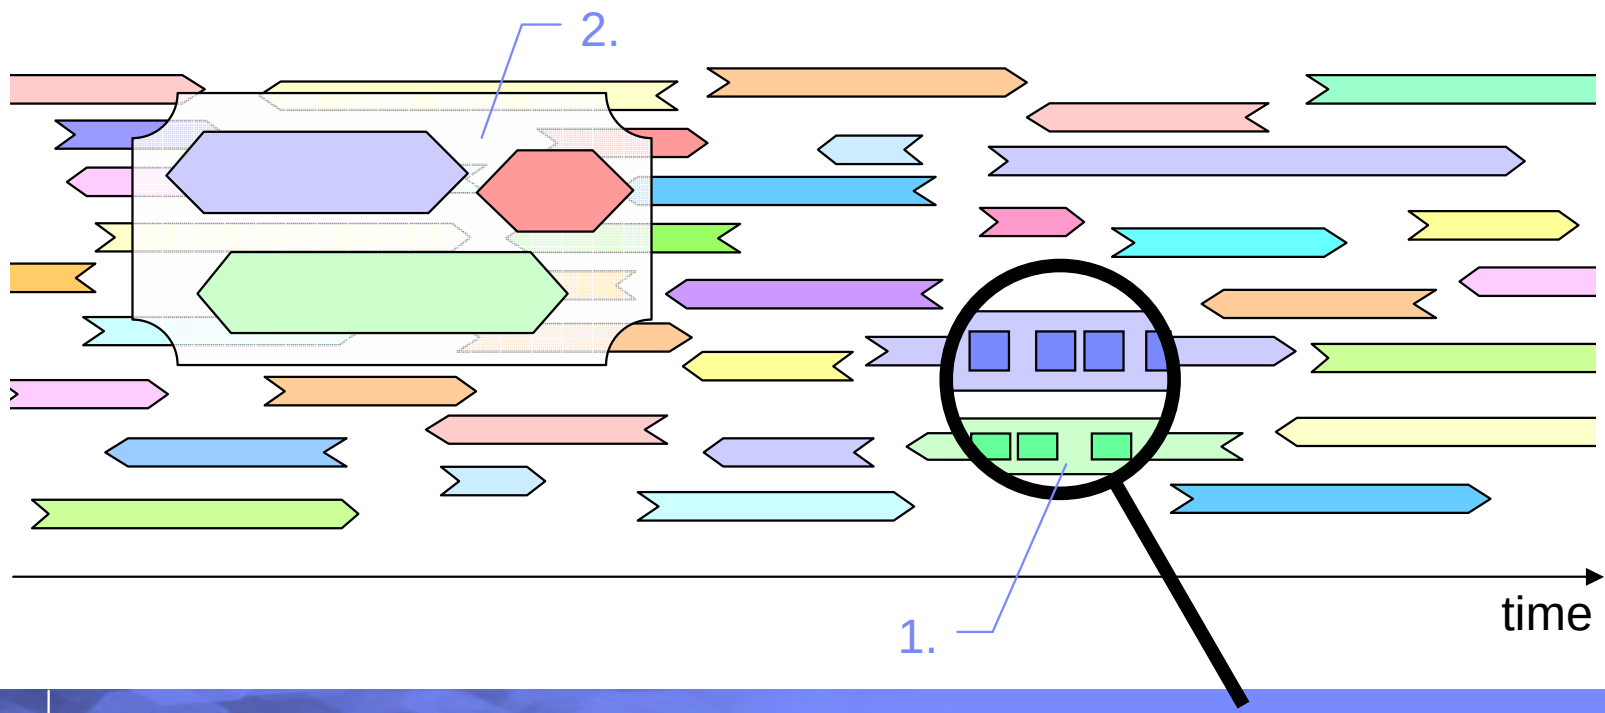
- Trade-off in network **traffic information collection** for **incident analysis**
  - **Raw packet traces**: finest level of detail but impractical to manage and search
  - **Flow traces**: high-level traffic abstraction but aggregated
- Traditional flow exports may **not provide traffic details required** to understand causes of incidents
  - Sampling on metering device
  - Aggregated IP addresses (prefixes) or AS level information in exports
  - Missing layer 2 and layer 3 header information
  - No packet content information
- Flow-level information is **often redundant** for incident analysis
  - Limited additional value on the flow level when given a set of prior traffic observations
  - Sequences of similar flows (streams, remote sessions, web/mail traffic, file transfers)
  - Flow record collections are still tedious to search, store, and analyze

© 2008 IBM Corporation

Objectives and Goals

- Extend a collector system to enable more accurate incident analysis
- Adapt information granularity depending on relevance of the traffic:

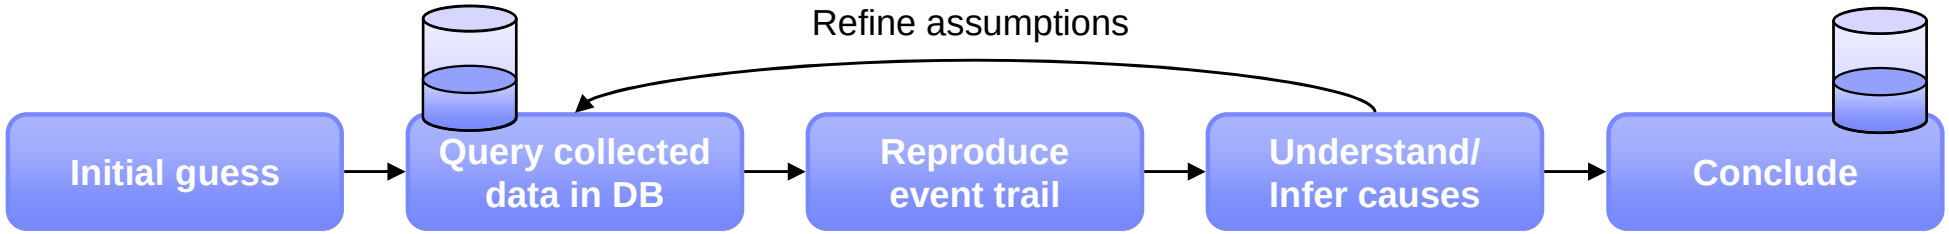
1. Focus in on particular traffic events to obtain more details
2. Compress known/less relevant traffic events (conserve a meaningful abstraction)



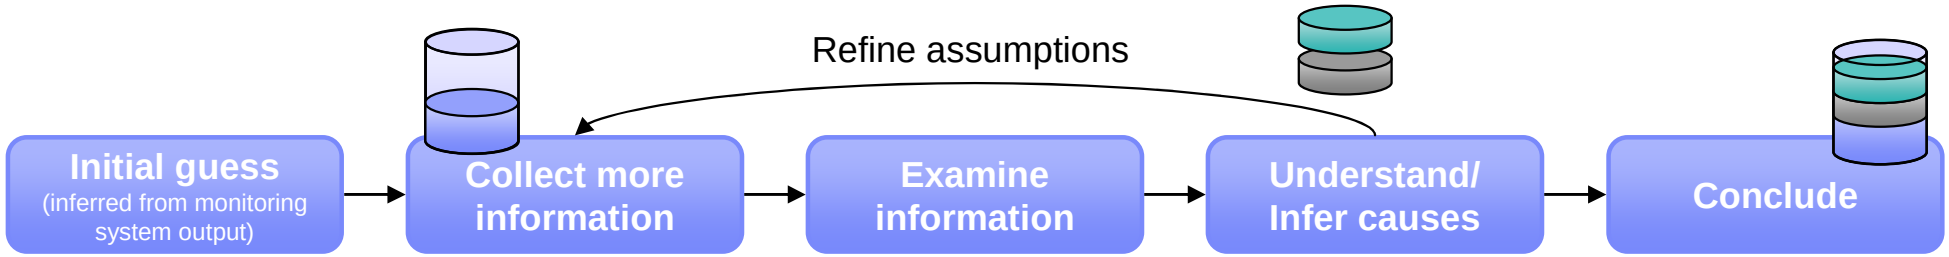
© 2008 IBM Corporation

Traffic Collection for Incident Analysis

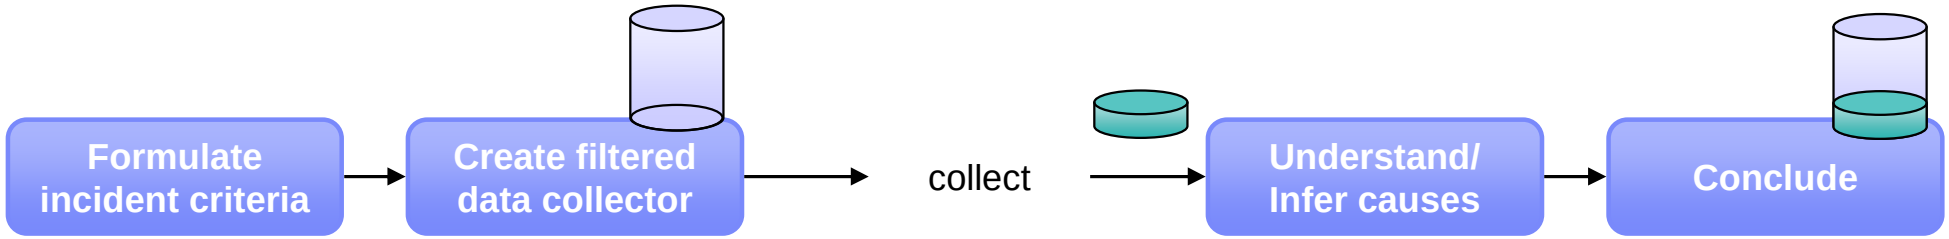
- **After-the-fact analysis**



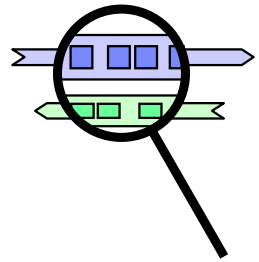
- **Real-time analysis**



- **Future incident trap**



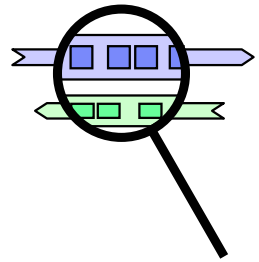
© 2008 IBM Corporation



Increasing Traffic Information Granularity

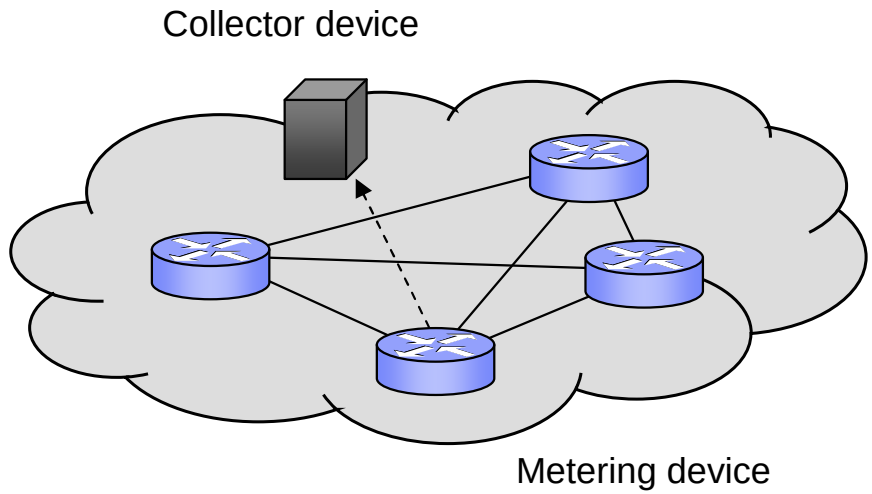
- **Problem**
  - Collecting detailed traffic information is cumbersome
  - Fixed and limited amount of information in traditional flow exports (e.g., NetFlow 5)
  - Management of information load collected
- **Traditional approach**
  - Physically attach a probe or packet dumping device at router (e.g., tcpdump with filtering)
  - Collection of rigid traffic information (e.g., entire packets)
  - No aggregation of data (e.g., bytes), analyze collected data: manual scripting
- **How to simplify data collection? Create zoom monitors!**
  - Dynamically controlled collection of relevant traffic information at desired level of detail
  - Centralized management of data collection campaigns
  - Make use of capabilities of network device inventory (routers, switches)
    - e.g., Cisco IOS Flexible NetFlow
  - Off-load aggregation and filtering to network devices

© 2008 IBM Corporation



Zoom Monitors

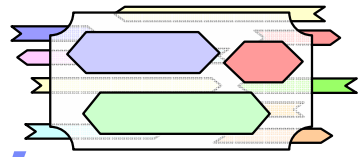
- **Specification**
  - Metering point and collector device
  - Zoom monitor lifespan
  - Filter criteria
  - Traffic aspects to be exported



- **Export collection and display**
  - Reconfigure metering device to create specific exports
  - Prepare collector device to store exported traffic information
  - Visualization of stored information (user interface)

- **Examples**
  - Show me the payload of all DNS requests of host 10.3.4.5 in the next 10 minutes
  - Look for all internal hosts scanning on TCP service port 9996 (e.g., candidate worm traffic)
  - Account all traffic flows using a particular service type (e.g., Voice over IP)
  - Export unsampled flow measurements from subnet 10.9.3.1/24

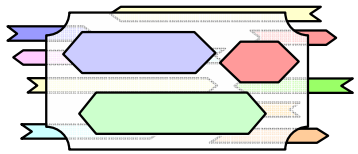
© 2008 IBM Corporation



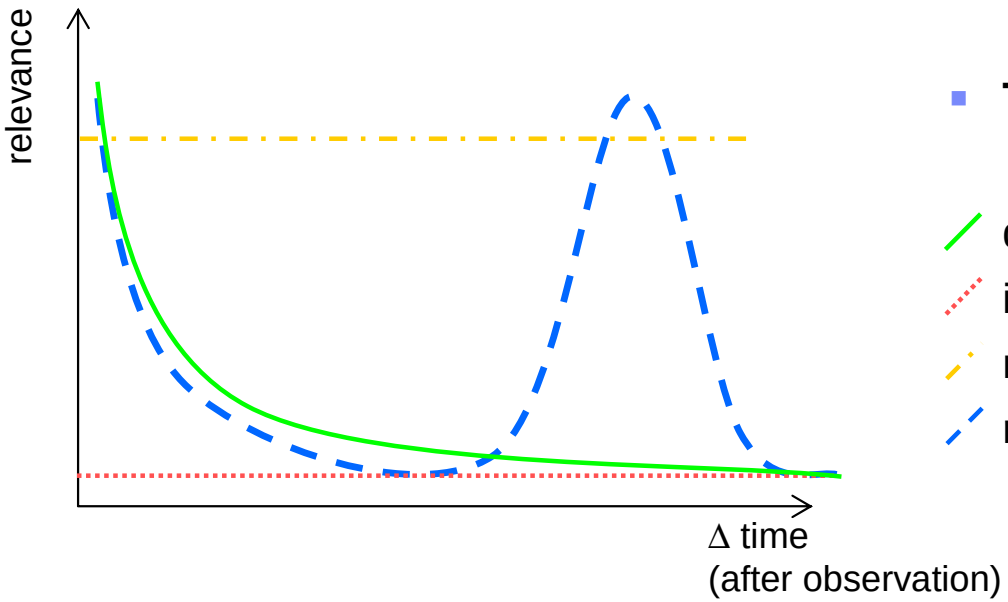
Decreasing Traffic Information Granularity

- **Problem**
  - Most stored traffic information is irrelevant for incident analysis (never accessed/requested)
  - Increased storage overhead and search complexity
- **Traditional approaches**
  - Rolling database: keep all flow records up to a limit (e.g., #entries, age): information removal
  - Uniform compression: adapt resolution of flow information (hourly, daily, weekly)
  - Keep top-k entries (according to some aspect)
- **How can we do better?**
  - Gradually compress information of irrelevant traffic events in a lossy fashion
    - With minimal impact on incident analysis tasks
  - Summarize similar events (coarse-grained representation)

© 2008 IBM Corporation



How should the traffic be compressed?



- **Typical information relevance patterns**

decreasing with time after collection (normal)

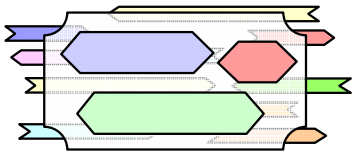
irrelevant

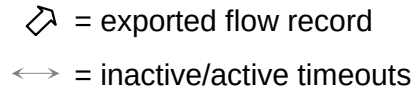
non-decreasing (keep for later analysis)

regaining relevance in posterior analysis

- **Multi-staged granularity reduction over time and with relevance**
  - We model information relevance with a “temperature” value: “hot” for latest events
  - Temperature decreases gradually over time: temperature ≈ interestingness of data
  - Degression of granularity takes place as a function of the temperature
  - Temperature can be increased for abnormal events: keep fine-grained representation

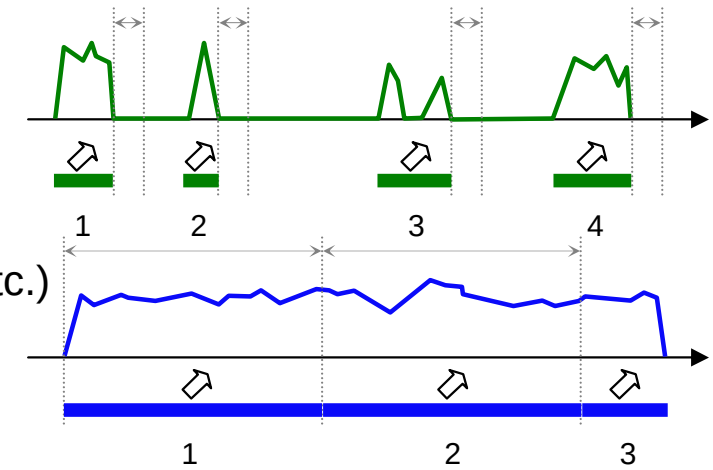
© 2008 IBM Corporation





Observations

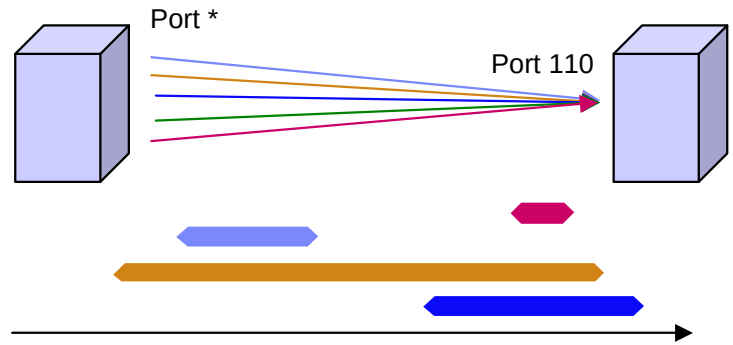
- **Flow export problem**
  - Multiple exports for a single connection
  - Examples:
    - Long-lived connections (streams, remote sessions, etc.)
    - Timeouts on routers (inactive/active timeout)
    - Change in service type (ToS field)



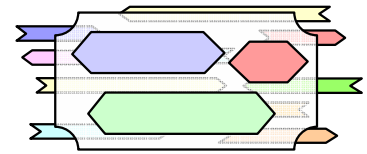
- **Bi-directionality problem**
  - Most flows have a reverse counterpart (= redundancy)



- **Information similarity problem**
  - Sets of records with limited added value on the flow level
  - Examples:
    - Groups of flows with similar properties (Web, mail, printer traffic, polling)
    - Known short-lived flows (DNS queries, etc.)
    - Typical similarity properties: <IP,port>, <IP,application>, <subnet,application>



© 2008 IBM Corporation



Compression Model[1]

	Abstraction models			
	Flow record	**Flow**	**Conversation**	**Session**
Raw exports	Yes	No	No	No
Flow definition (5-tuple)	Yes	Yes	Yes	No (subset thereof)
Direction	Uni-directional	Uni-directional	Bi-directional	Bi-directional
# Flow records	1	≥ 1	≥ 1	≥ 1
# Flows	1	1	1 or 2	≥ 1 or ≥ 2
# Conversations	1	1	1	≥ 1

[1] without prior knowledge such as domain or application specific information

© 2008 IBM Corporation

Implementation

- **Metering device configuration for zoom monitors**
  - Reconfiguration of metering devices
  - Management console
- **Export collector**
  - Storage and querying
- **Traffic information compression**
  - Aggregation technique

© 2008 IBM Corporation

Metering Device Configuration

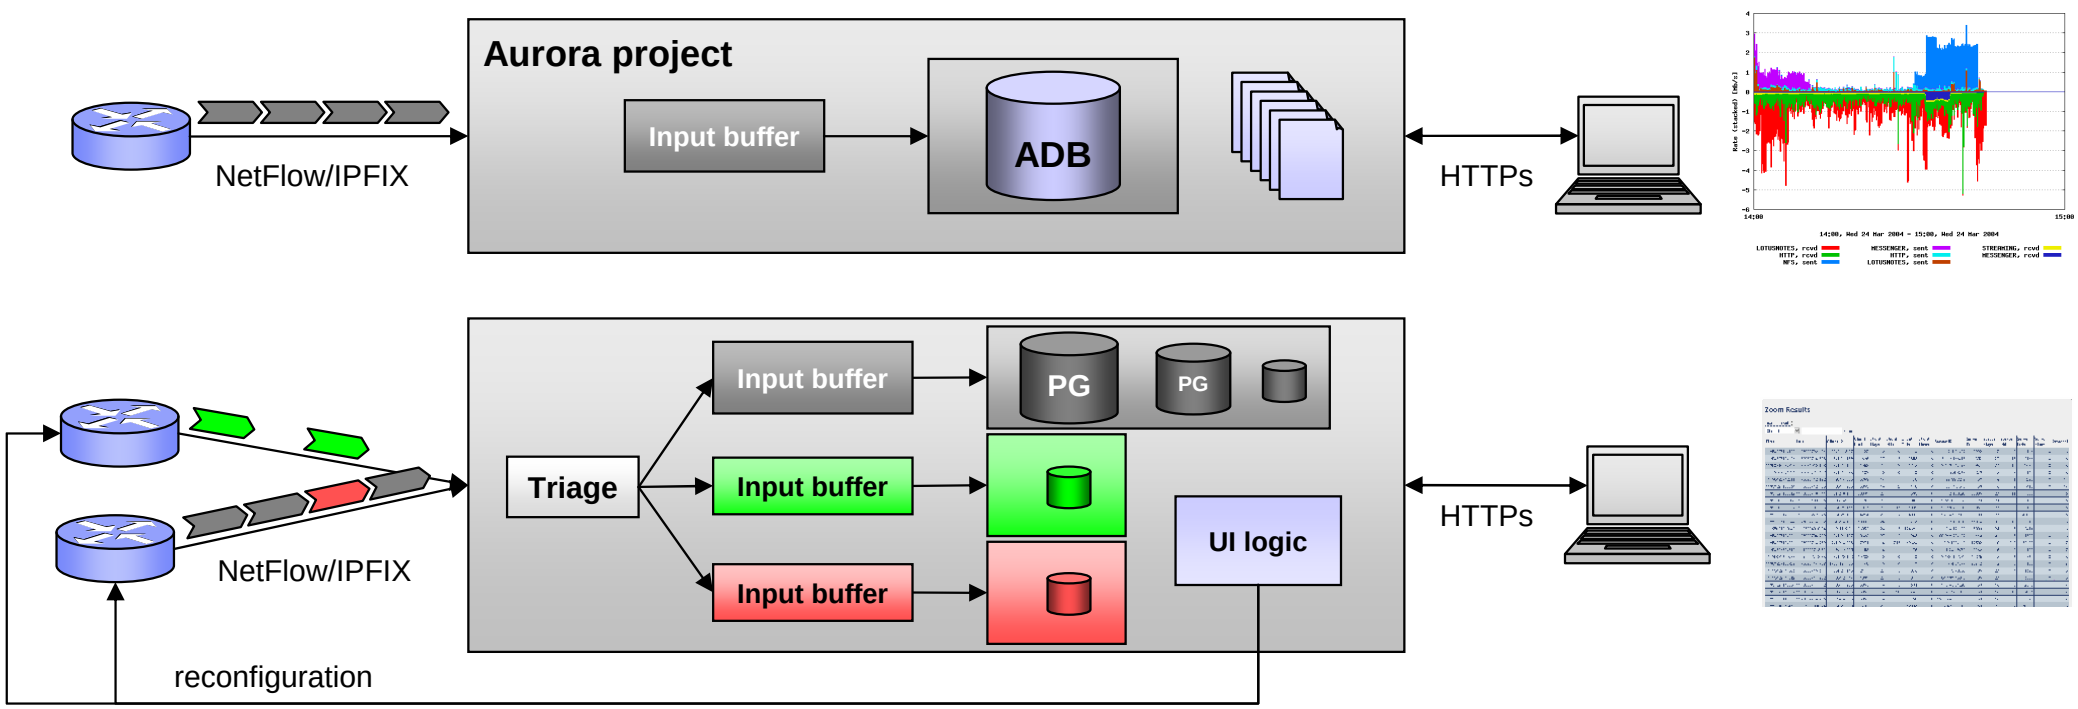
- **Technologies**
  - **Cisco IOS Flexible NetFlow (FNF)**
    - Configuration of multiple customized monitors
    - Currently: input filtering for FNF monitors not available (input filters needed at collector)
  - **Hespera Traffic Meter (IBM Research)**
    - Software-based flow monitor supporting NetFlow v5 and v9, IETF IPFIX exports
    - Customized flow exports (variable templates), CLI-based reconfiguration
    - Filtering with tcpdump syntax

- **User-based creation of dynamic zoom monitors**
  - Web-based specification of zoom monitors
  - Deployment on metering device (CLI-based) and management (e.g., lifespan)
    - Future: XML-based configuration (cf. [Dimitropoulos/Kind] or [NetConf])
  - Registering the zoom monitor at collector device (for disambiguation/triage)
  - Pre-defined zoom monitor templates from library

© 2008 IBM Corporation

Export Collector

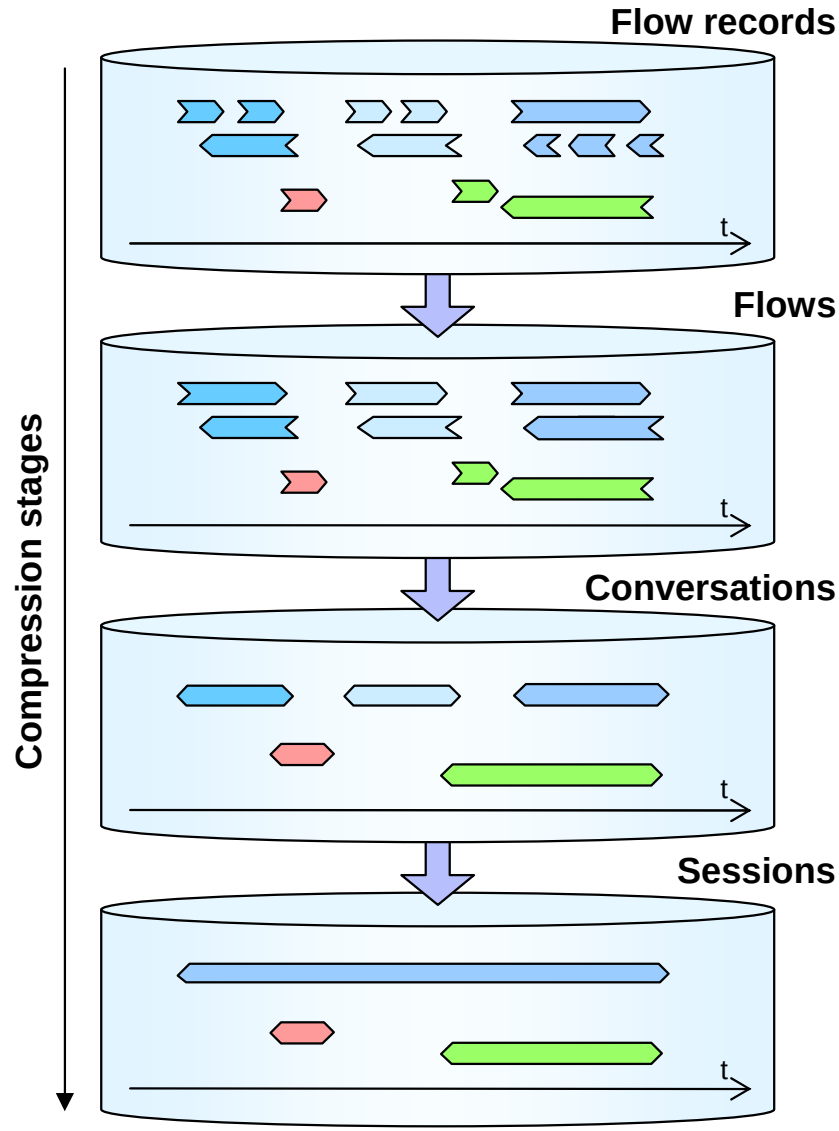
- **Prototype based on the Aurora flow analyzing system (IBM Research)**
  - Replaced existing aggregation database (ADB) with PostgreSQL (PG) backend
  - Input triage according to zoom monitors
  - Relevance-sensitive compression for default flow exports
  - Extension of the web user interface



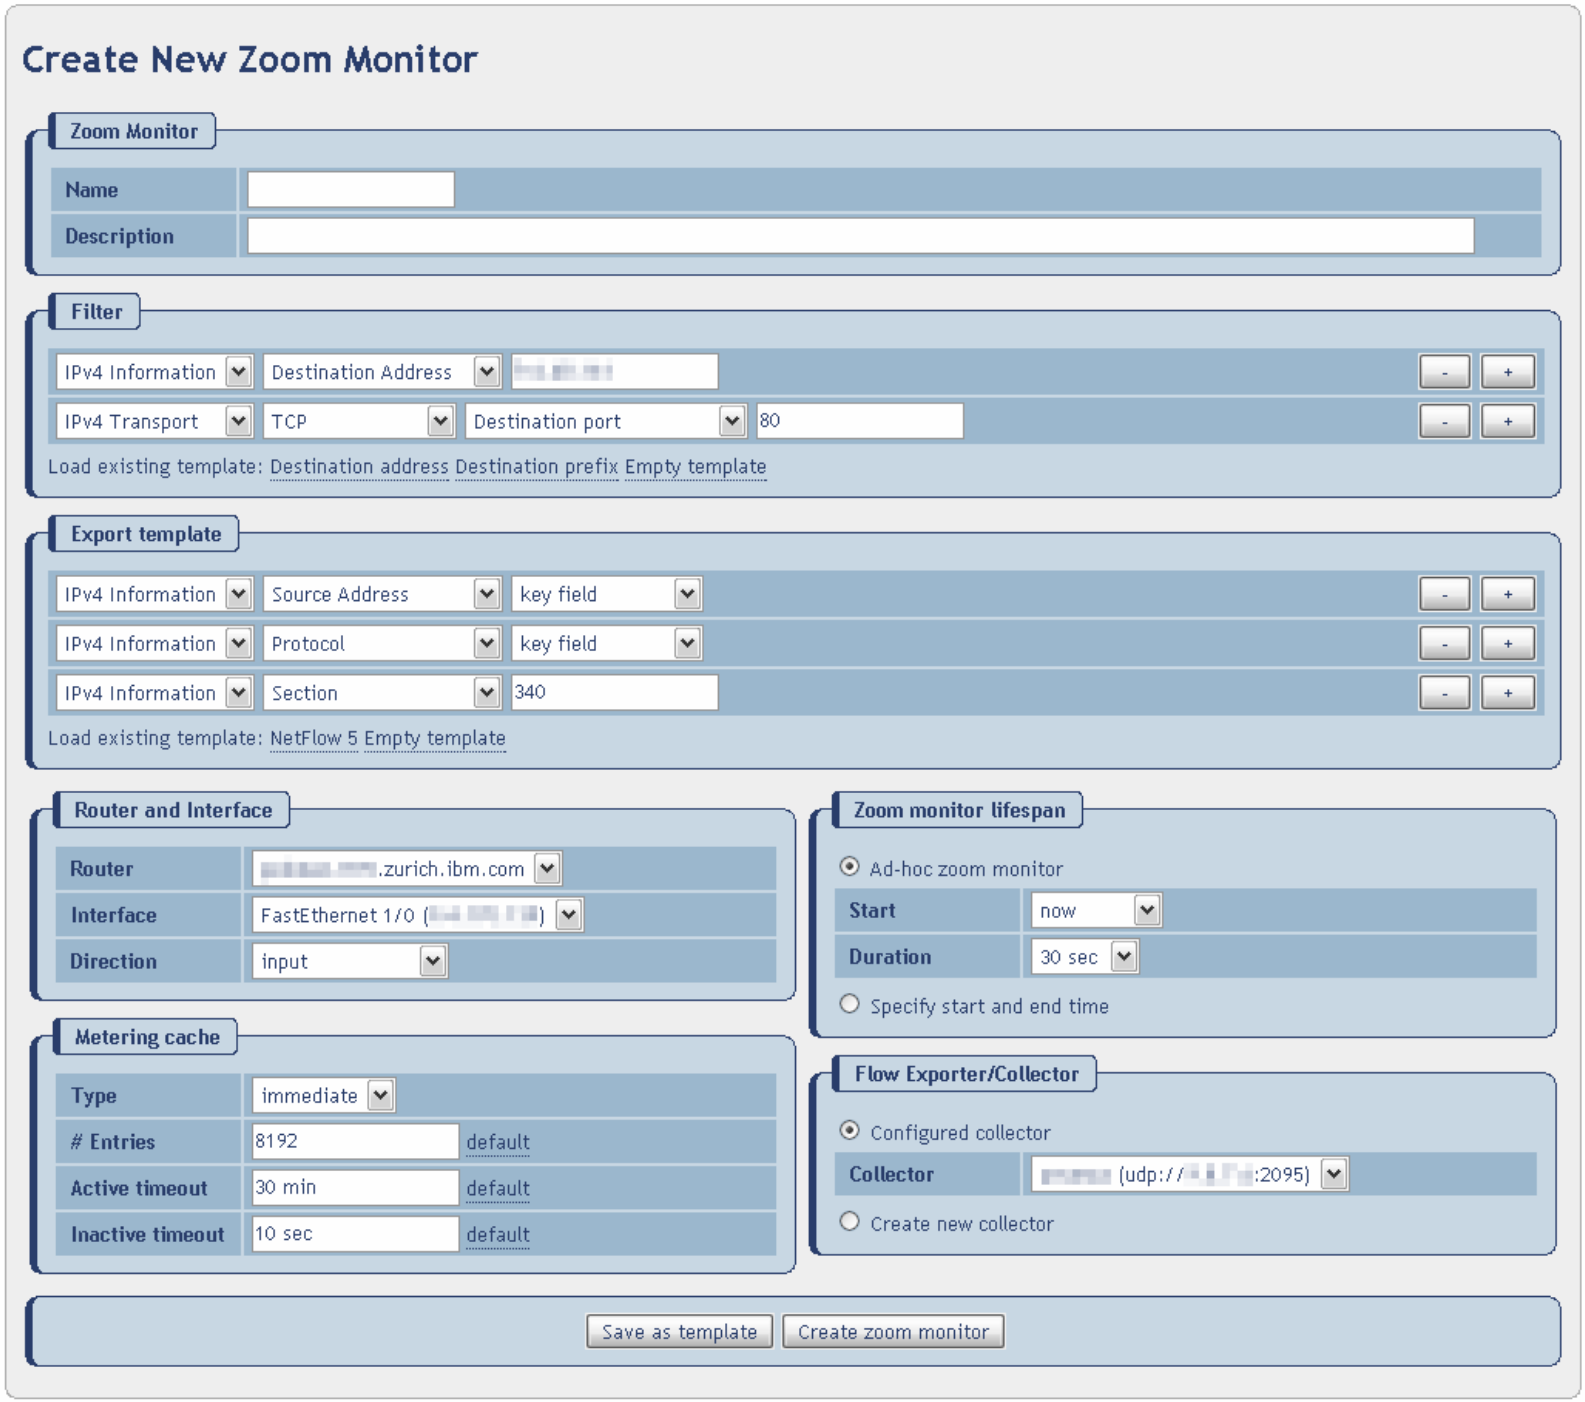
© 2008 IBM Corporation

Traffic Information Compression

- **Incrementally populate the databases**
  - Update entries in databases
  - Remove entries based on temperature values from finer-grained databases
  - Keep “Session” database
- **Considerations**
  - 1-by-1 inserts/updates are generally slow
  - Prepare entry sets and use bulk imports
  - Partitioning and indexing



© 2008 IBM Corporation



Filter definition

Export information

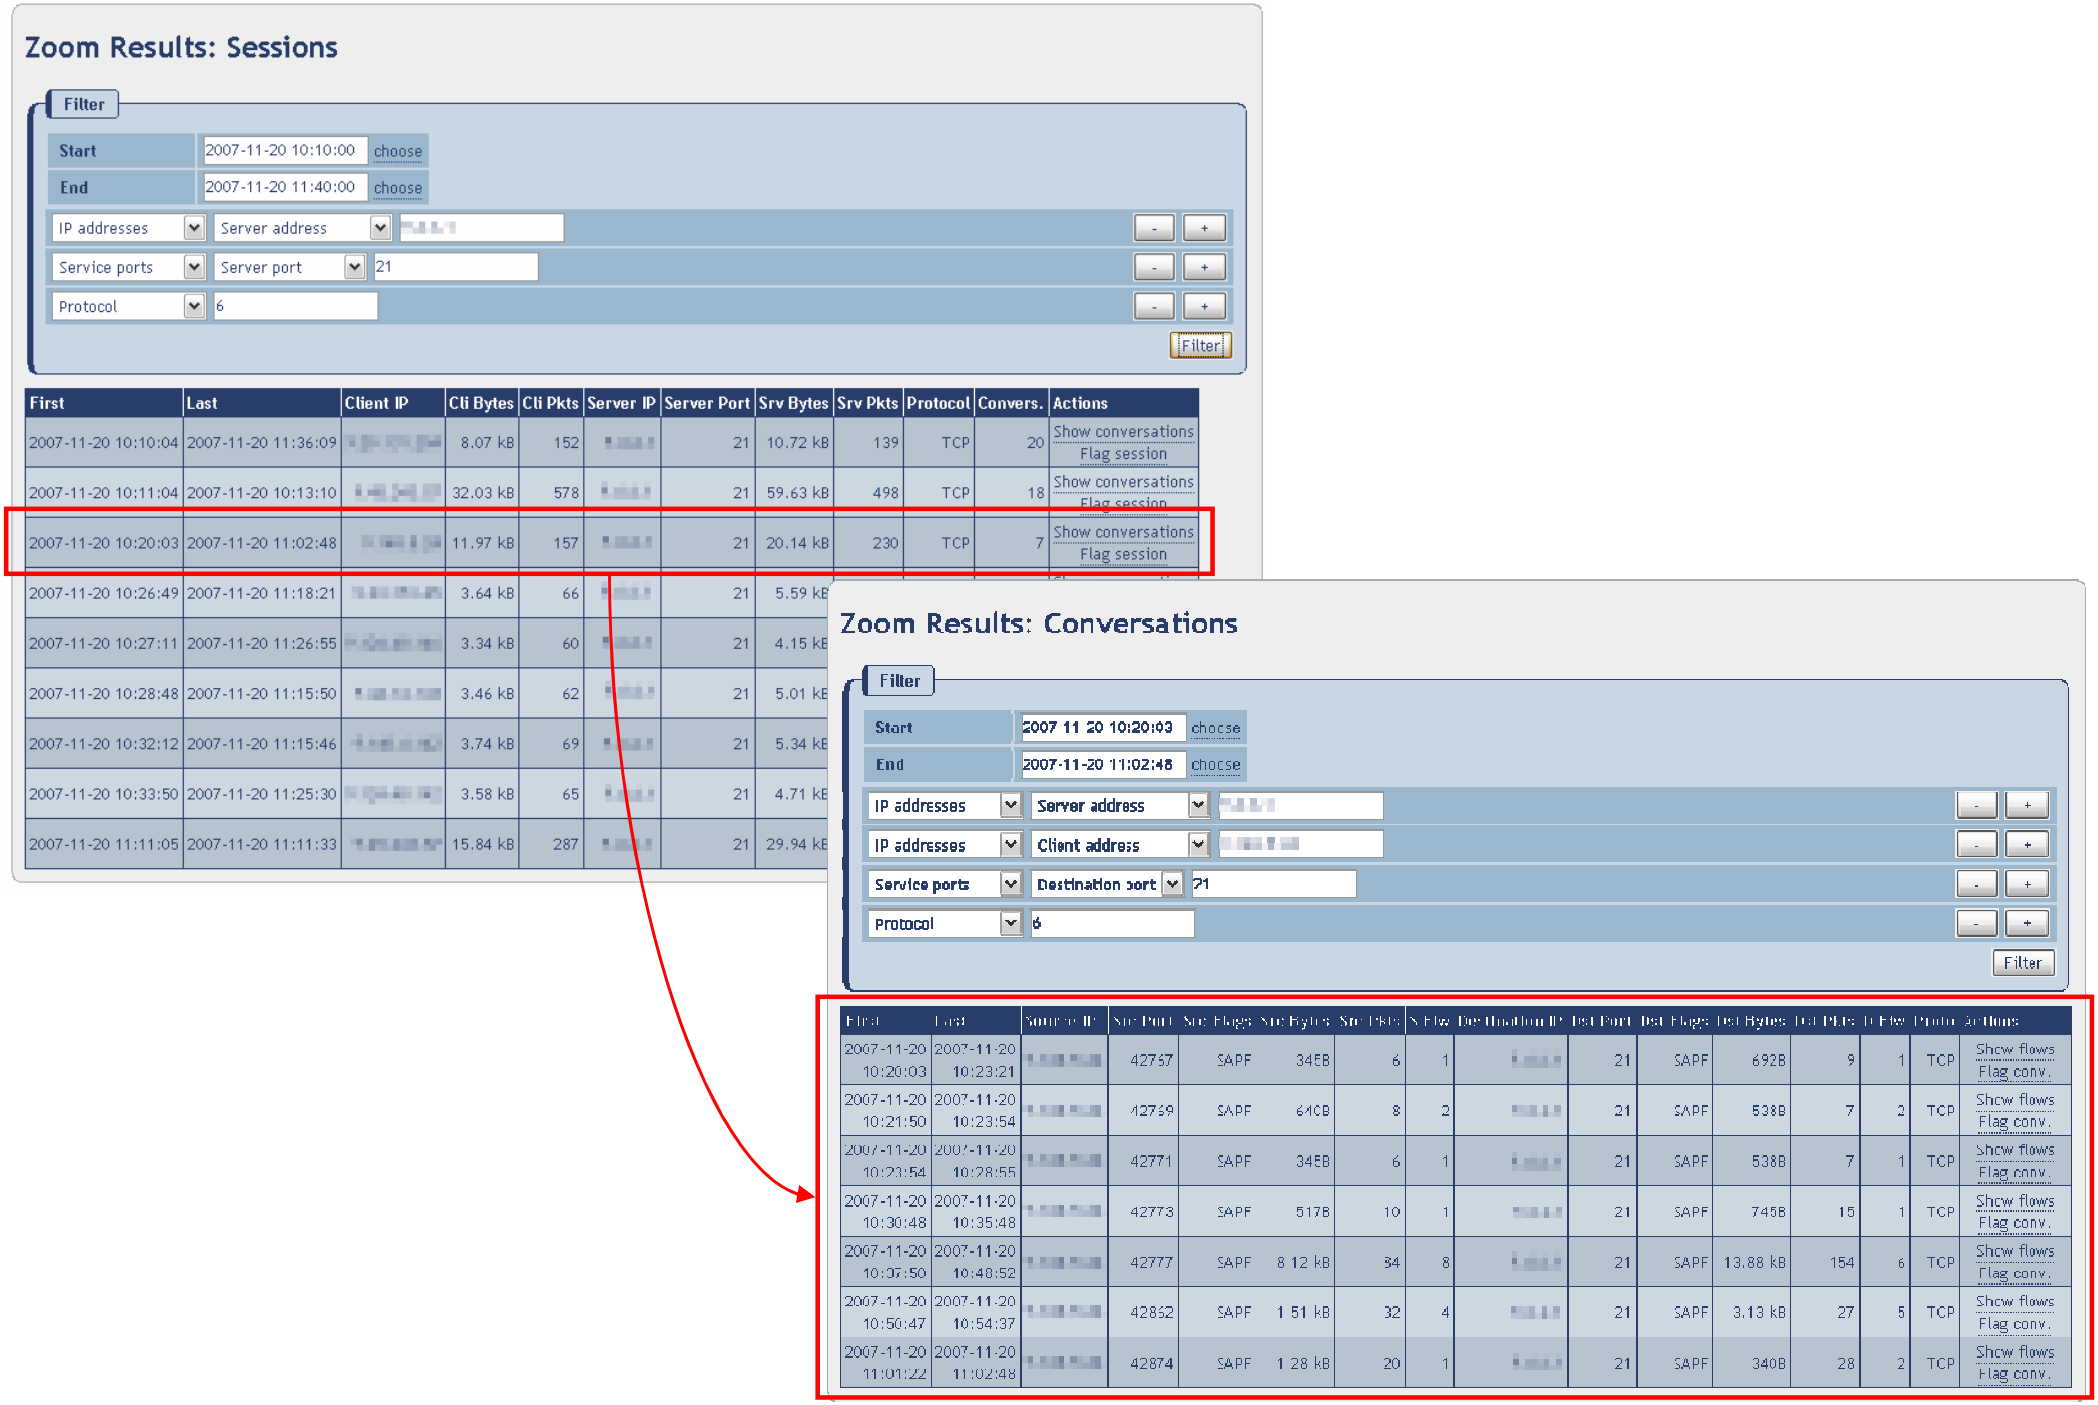
Router/Interface

Lifespan

Collector

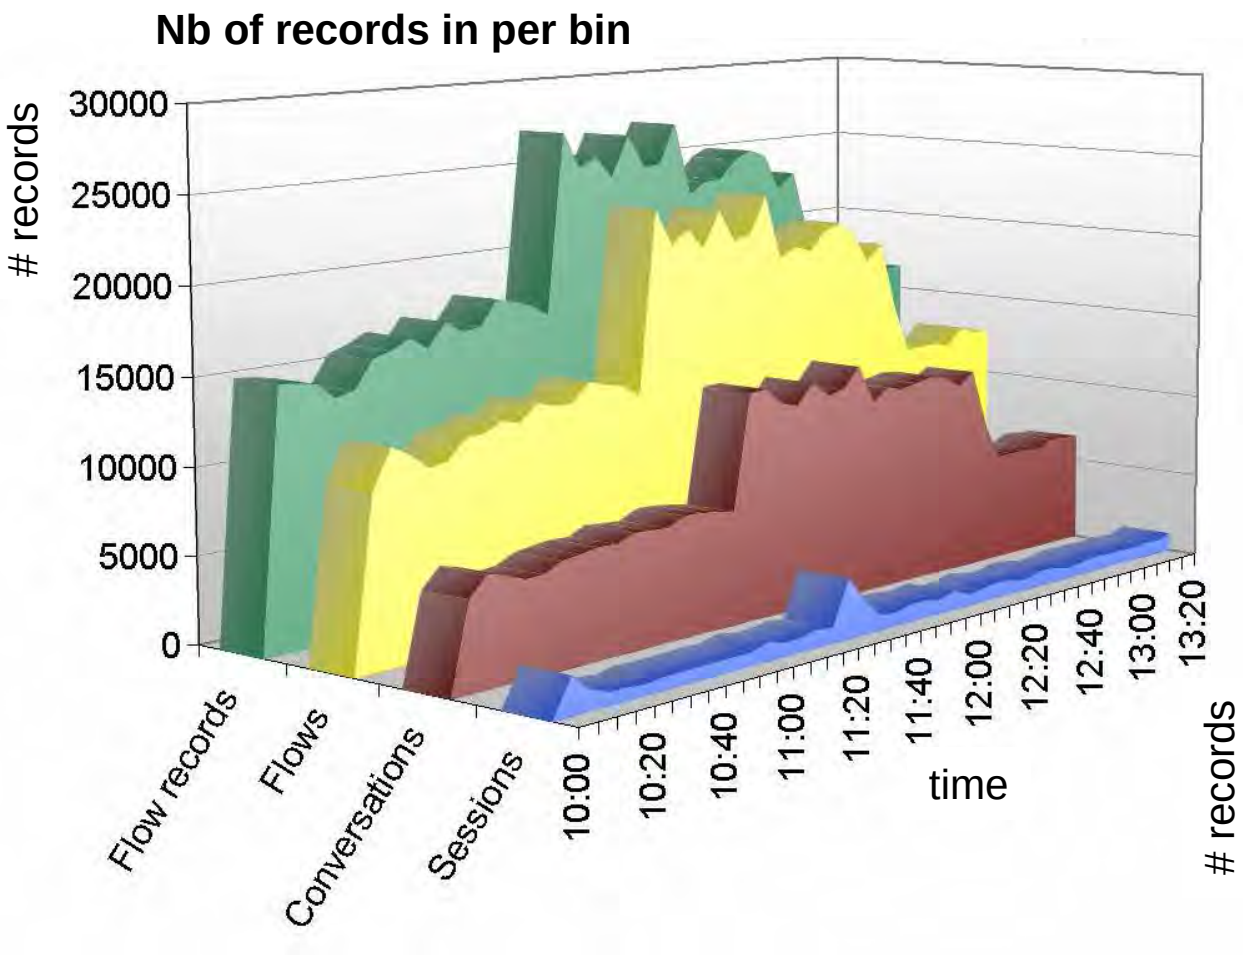
Cache

© 2008 IBM Corporation



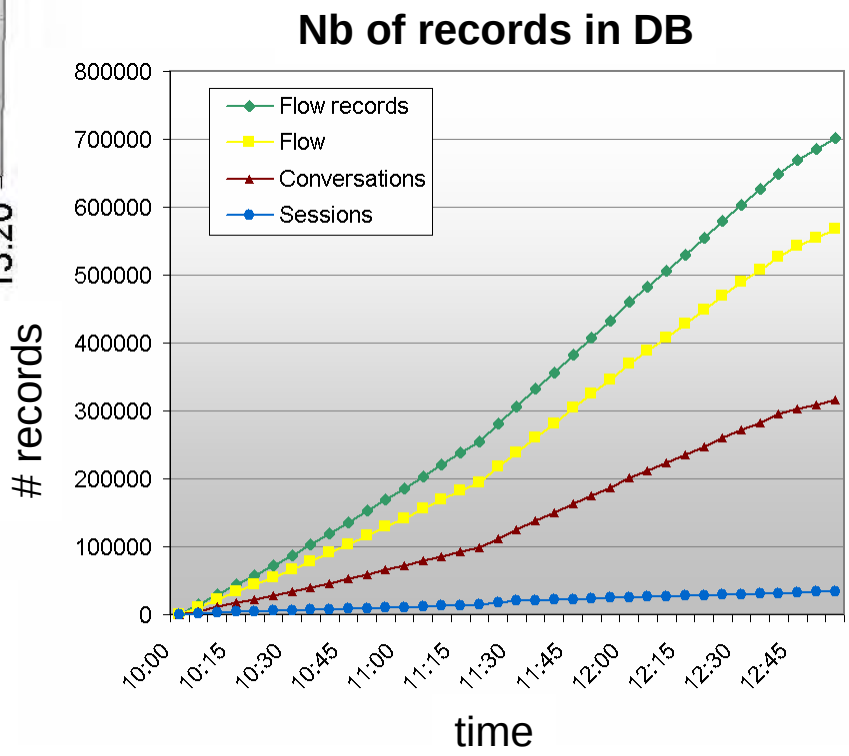
© 2008 IBM Corporation

Results: Compression (WAN traffic)



- **Average compression ratio**

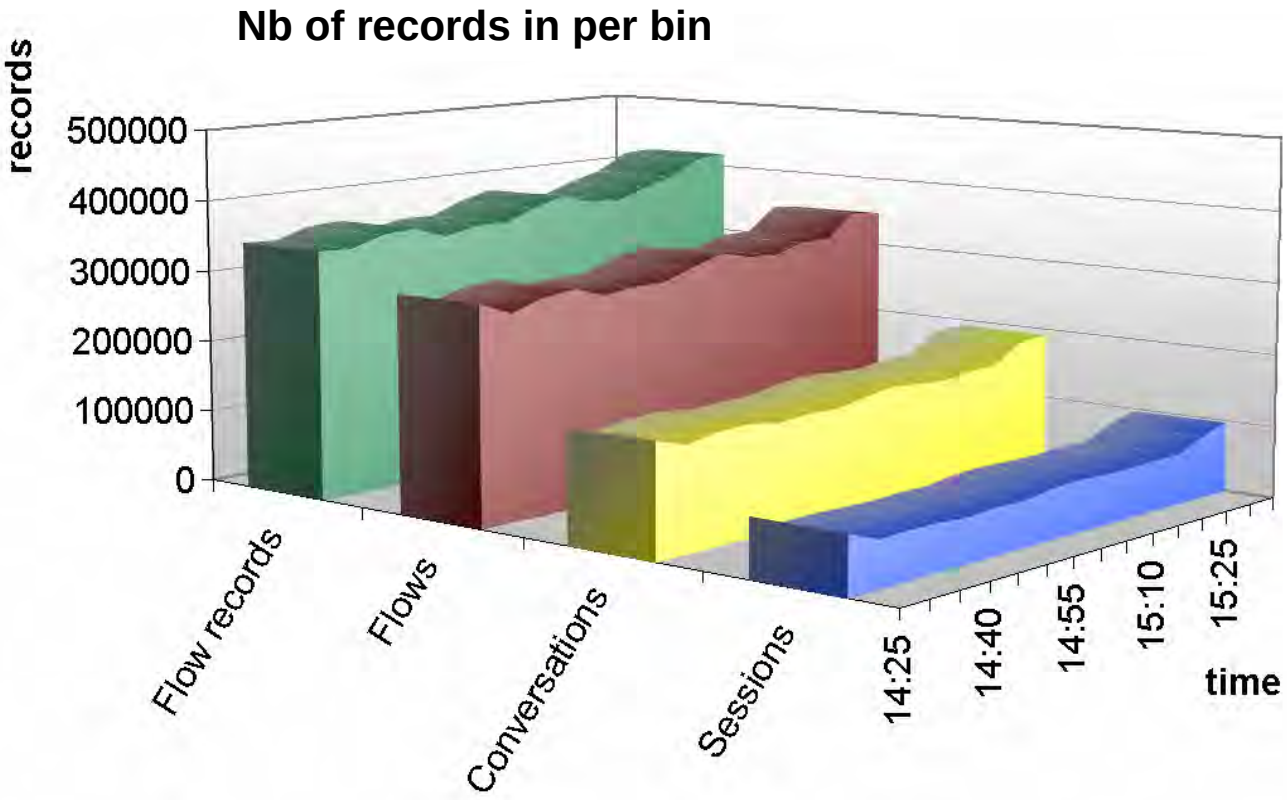
#flow records : #flows	**1.26**	σ = 0.07
#flow records : #conversations	**2.34**	σ = 0.28
#flow records : #sessions	**22.80**	σ = 7.00



- Session inactive timeout: 20min

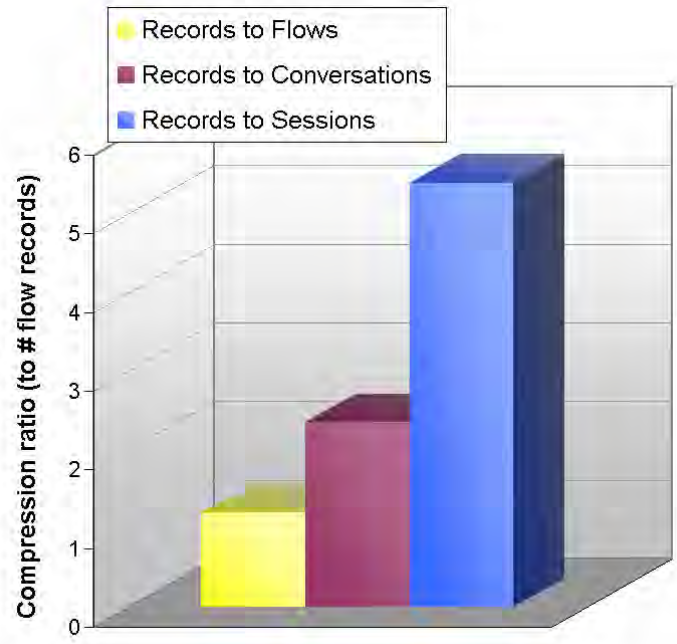
© 2008 IBM Corporation

Results: Compression (datacenter traffic)



- **Average compression ratio**

#flow records : #flows	**1.19**	σ = 0.02
#flow records : #conversations	**2.35**	σ = 0.07
#flow records : #sessions	**5.39**	σ = 0.46



- Session inactive timeout: 20min

© 2008 IBM Corporation

Future Work and Visions

- **Automated zoom monitor creation**
  - Interface to a behavior-based network anomaly detection system
  - Proactive collection of proofs for posterior forensic analyses of abnormal events
- **Distributed collector infrastructure**
  - Distributed collectors, e.g., at multiple sites (scalability)
  - Transfer required information to central reporting system on demand
- **Enhance compression technique**
  - Meta-data representation using anomaly sensor input
  - Application-sensitive compression
- **Cisco IOS Flexible NetFlow with input filters**
  - Perform filtering on routers to replace software-based metering (and filtering)

© 2008 IBM Corporation

Conclusion

- **Incident analysis tool adapting flow information granularity**
  - Increase level of detail of relevant/unknown traffic events
  - Decrease level of detail (compress) of less relevant events
  - Keep a meaningful abstraction of all traffic events
- **Creation of customized zoom monitors**
  - Zoom in on specific traffic to gain additional information about its properties and behavior
  - Centralized management of metering devices for traffic detail collection
- **Reduced long-term storage requirements**
  - Encouraging test results with multiple flow information granularity levels

© 2008 IBM Corporation

References

- IBM Research. “Aurora – Network Traffic Analysis and Visualization”. http://www.zurich.ibm.com/aurora/
- Xenofontas Dimitropoulos and Andreas Kind. “Configuration of Monitors”. FloCon2008.
- NETCONF IETF Working Group. http://www.ops.ietf.org/netconf/
- Cisco Systems, Inc. “Cisco IOS Flexible Netflow”. Product website: http://www.cisco.com/en/US/products/ps6965/products_ios_protocol_option_home.html

© 2008 IBM Corporation



Anonymizing Network Flow Data

Timothy J. Shimeall (tjs@cert.org)
January 2007
FloCon 2008

© 2008 Carnegie Mellon University

Overview

The balance of anonymization

Subnet-preserving

Subnet-collapsing

Host-preserving

Host-collapsing

Ports & Other issues

Conclusion

The Balance of Anonymization

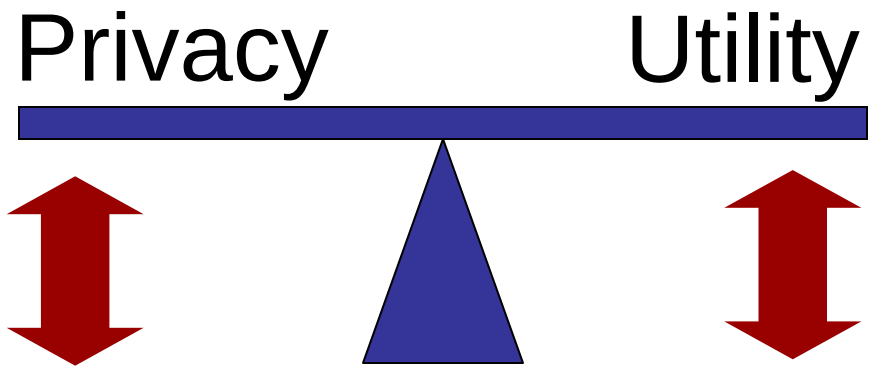
Flow itself preserves some privacy by aggregation and eliding content.

Anonymization is to aid in preserving the privacy of organizations represented in the data

- Data owner
- Partner or Customer
- Incidental
- Attacker

The more you anonymize the data, the less analyses can be done with it.

Need to explore a range of options



Subnet Preserving

Preserve host identity while concealing network.

How:

- Prepare list of networks
- Assign random substitution for network prefix
- Mask and replace prefix on each address
- Associative array works well for substitutions

Balance:

- Enables analysis down to host identity, but not organization identity
- Can be reversed by outside knowledge (server suffixes)

~~248.204~~.5.3

010.005.5.3

Subnet Collapsing

Conceal network structure and host identity, but preserve commonality of network

How:

- Reduce all address to the network
- Prepare random substitution for network
- Replace address with network substitutions

Balance:

- Allows network-level behavior analysis
- Might be reversed by organizations with lots of contact with data source

248.204~~.5.3~~

~~248.204~~.0.0

010.005.0.0

Host Preserving

Preserve host identity while concealing network commonality

How:

- Generate list of addresses
- Generate random substitution for each address
- Replace each occurrence with same substitution

Balance:

- Allows host-specific analysis
- Difficult to reverse

~~248.204.5.3~~

10.2.3.9

~~248.204.5.12~~

192.168.12.7

Host Randomizing

Do not preserve host or network identity (a.k.a., remove address content in any useful way)

How:

- Replace each occurrence of each address with random value
- Allow repetition of random values

Balance:

- Only permit analysis that does not involve address information
- Extremely difficult to reverse

~~248.204.5.3~~ ~~128.0.3.2~~

10.2.3.7 192.168.7.12

~~248.204.5.3~~ ~~0.5.4.1~~

192.168.17.37 10.2.3.7

Ports and Other Issues

There's more to anonymization of flow than addresses

- Network ports can be very revealing (OS fingerprinting)
- Timing information might be revealing
- TCP flags might be revealing (odd patterns)

Can anonymize this information:

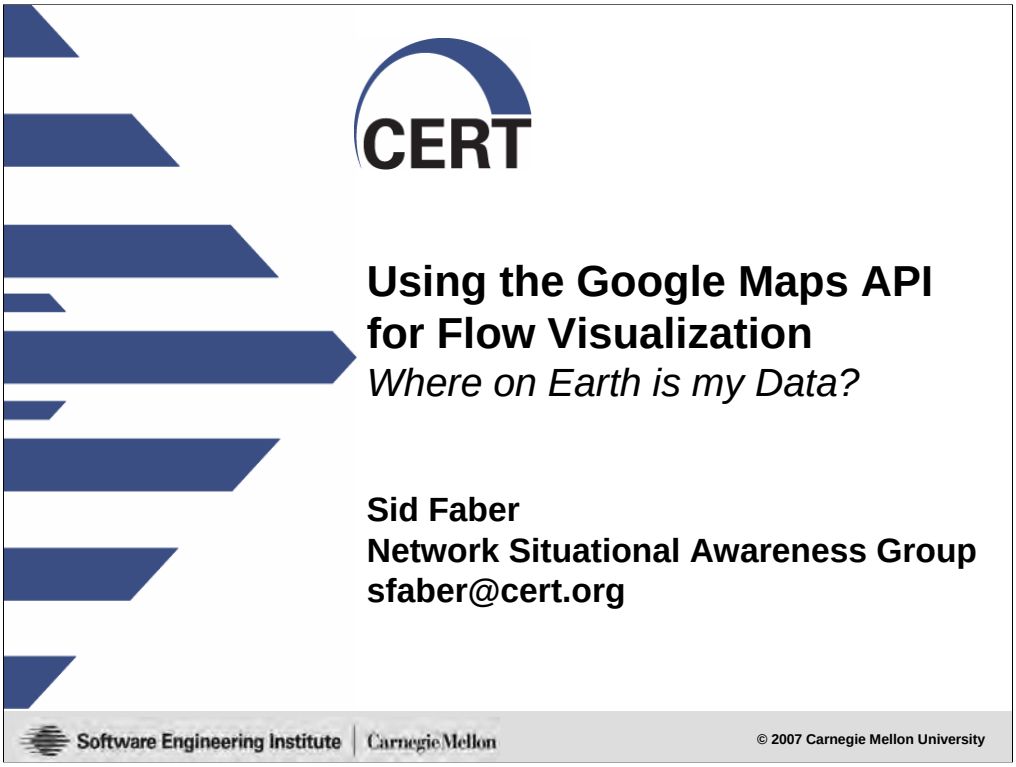
- Ports: reduce to service, substitute; reduce to common/reserved/dynamic
- Timing: restart epoch; rescale timing; collapse interval
- TCP flags: reduce to function; remove OS-dependencies

Conclusion

Data sharing is difficult

Anonymization can be useful, but limiting

Anonymized does not mean private or irreversible



© 2007 Carnegie Mellon University

Agenda

- Step 1: Extracting Flow Data
- Step 2: Geolocation
- Step 3: Convert to XML
- Aside: The Google Maps API
- Step 4: The HTML Page

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
2

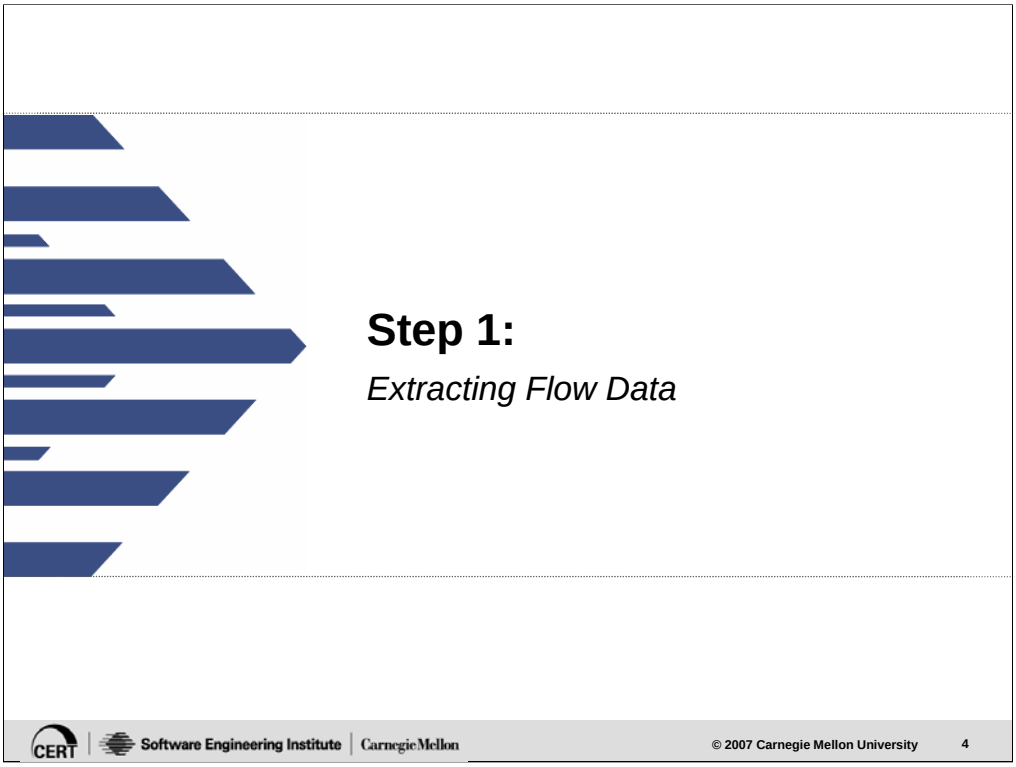
Data Used for Demo

SC06 Data Set

- November 14, 2006
- Goal is to look at who talked to whom

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
3



© 2007 Carnegie Mellon University

Extracting Flow Data

What story do you want to tell with geolocation?

- Traffic source or destination
  - Data record = one value per address
- Relations between addresses
  - Data record = one value per source, destination address pair

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
5

Extracting Flow Data: SiLK Example

Traffic destination

```
$ rwfilter
    --start=2006/11/14
    --proto=0-255
    --class=all --pass=stdout
    | rwuniq
    --fields=dip --bytes  > dst.txt
140.221.159.103  12568504471655
    172.30.5.11  11381325217792
    172.30.6.11   7397483692032
```


© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University

Step 1: Summary

Extract Flow Data

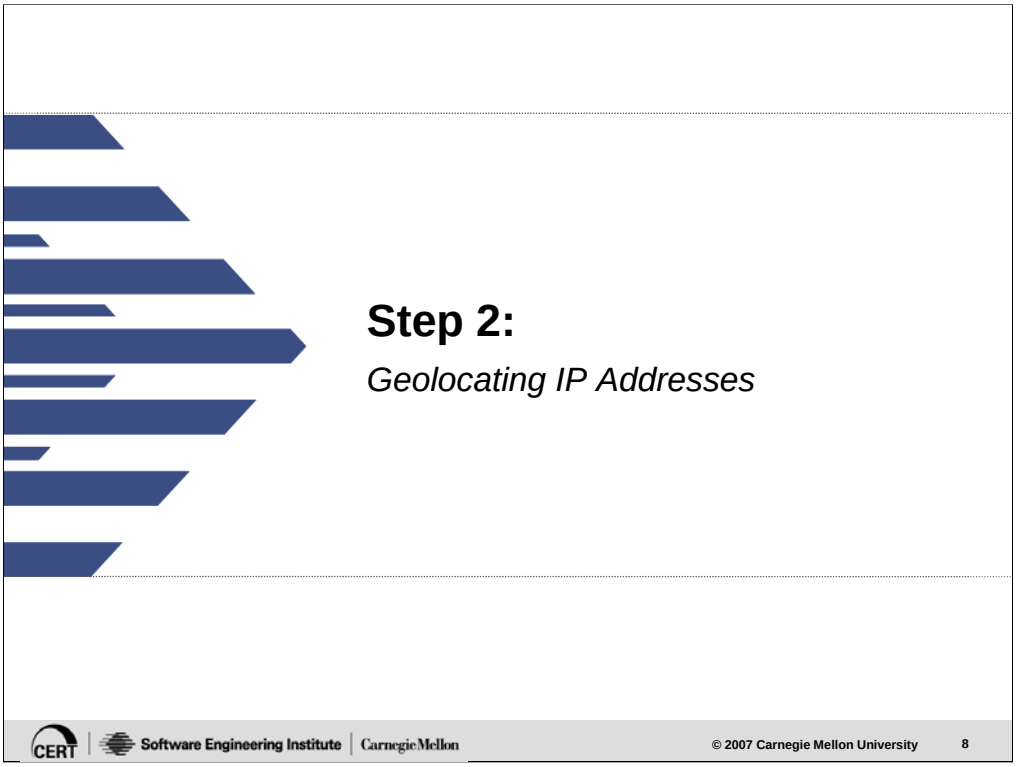
- Start with raw flow data
- End with summarized flow data (2 columns)
  - Destination IP, value
  - Space delimited

- For Example:

```
140.221.159.103  12568504471655
    172.30.5.11  11381325217792
    172.30.6.11   7397483692032
```


© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University



© 2007 Carnegie Mellon University

Geolocating by Country

Map IP to Country: IPligence, http://www.ipligence.com

```
"0000000000","0033554431","US","UNITED STATES","NA"...
"0033554432","0050331647","DE","GERMANY","EU","EUROPE"
```

Map Country to Lat/Long: MaxMind, http://www.maxmind.com/app/country_latlon

Numeric IP

```
US,38.0000,-97.0000
DE,51.0000,9.0000
```

Combine IP-to-Lat/Long Mapping

```
0000000000 0033554431 US 38.0000 -97.0000
0033554432 0050331647 DE 51.0000   9.0000
0050331648 0067108863 HK 22.2500 114.1667
```


© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 9

Geolocating by Addresses

DNS LOC

```
$ host -t LOC cmu.edu
   cmu.edu LOC 40 26 39.000 N 79 56 36.200 W 283.00m ...
```

Caida Netgeo

```
$ wget http://netgeo.caida.org/perl/netgeo.cgi \
  ?target=128.2.10.162
  ...
  TARGET:         128.2.10.162<br>
  NAME:           CMU-NET<br>
  NUMBER:         128.2.0.0 - 128.2.255.255<br>
  LAT:            40.44<br>
  LONG:           -79.95<br>
  ...
```

Hostip.info, http://www.hostip.info/dl/index.html

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 10

Sample Commercial Data: Quova

1	start_ip_int	50331648	67272896
2	end_ip_int	50378239	67272959
3	cidr	24	26
4	continent	north america	north america
5	country	united states	united states
6	country_iso2	us	us
7	country_cf	80	97
8	region	northeast	northeast
9	state	connecticut	massachusetts
10	state_cf	10	87
11	city	fairfield	woburn
12	city_cf	10	77
13	postal_code	06825	01888
14	phone_number_prefix	203	781
15	timezone	-5	-5
16	latitude	41.1753	42.4867
17	longitude	-73.2812	-71.1543
...		...	...

Numeric IP

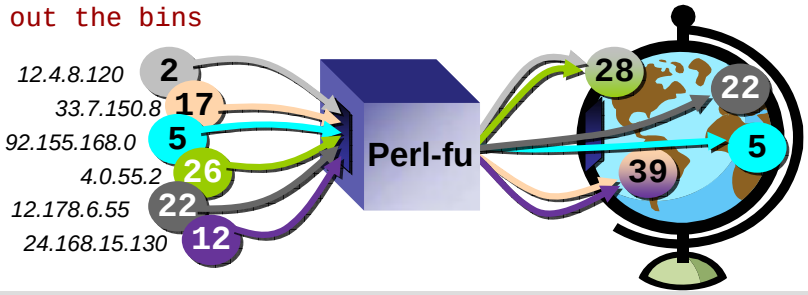
© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 11

Add location to data and regroup

Perl-fu pseudocode:

```
Read location data into a lookup table
For each line of data {
    Extract IP and [value]
    Find lat,long coordinates for IP
    Create a bin for the coordinates and add [value]
}
Print out the bins
```



© 2007 Carnegie Mellon University

Geolocating with SiLK pmaps

Prefix maps associate a value with an IP address prefix

- Text based pmap:

```
#Start-IP  End-IP     CC Lat     Long
0033554432 0050331647 DE 51.0000 9.0000
0050331648 0067108863 HK 22.2500 114.1667
```

pmap value

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 13

Building the Geolocation pmap

Some perl-fu:

```
read countrylatlng.txt into a hash
foreach line in the ipligence data set {
  look up the countrylatlng.txt line for
  the code
  print out the ip range, country code and
  coordinates
}
```

- See *make-geo-cc-pmap.pl* in the sample code

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 14

Using the Geolocation pmap

Use the pmap with rwuniq:

```
$ rwfilter \
  --start=2006/11/14 \
  --proto=0-255 \
  --class=all --pass=stdout \
  | rwuniq \
  --pmap-file=geo-cc.pmap \
  --fields=dval --bytes --delimited=" " --no-titles \
> geo-dst.txt

US 38.0000 -97.0000 102372319236580
JP 36.0000 138.0000   9965004709495
CA 60.0000 -95.0000    569989239278
```


© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 15

Step 2: Summary

Geolocate Flow Data

- Start with summarized flow data
- End with location data (4 columns)
  - Destination label, latitude, longitude, value
  - Space delimited
  - SiLK pmaps combine steps 1 and 2
- For example:

```
US 38.0000 -97.0000 102372319236580
JP 36.0000 138.0000 9965004709495
CA 60.0000 -95.0000 569989239278
```


© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 16



© 2007 Carnegie Mellon University

XML Data

Convert to XML

- The GoogleMaps routine we'll be using takes XML input
- We define the schema
- We'll process Step 2 data with a simple awk command

```
$ cat geo-dst.txt | \
  awk ' BEGIN {print "<markers>"} \
  { printf "<marker lbl=\"%s\" lat=\"%s\" lng=\"%s\" \
    val=\"%s\"/> \n", $1, $2, $3, $4} \
  END { print "</markers>"} ' \
  > geo-dst.xml
```


© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 18

Step 3: Summary

Convert to XML

- Start with labels, coordinates and values
- End with XML document with the same data

- For example:

```
<markers>
<marker lbl="CN" lat="35.0000" lng="105.0000“ val="704206"/>
<marker lbl="MR" lat="20.0000" lng="-12.0000" val="200"/>
<marker lbl="KN" lat="17.3333" lng="-62.7500" val="646"/>
</markers>
```

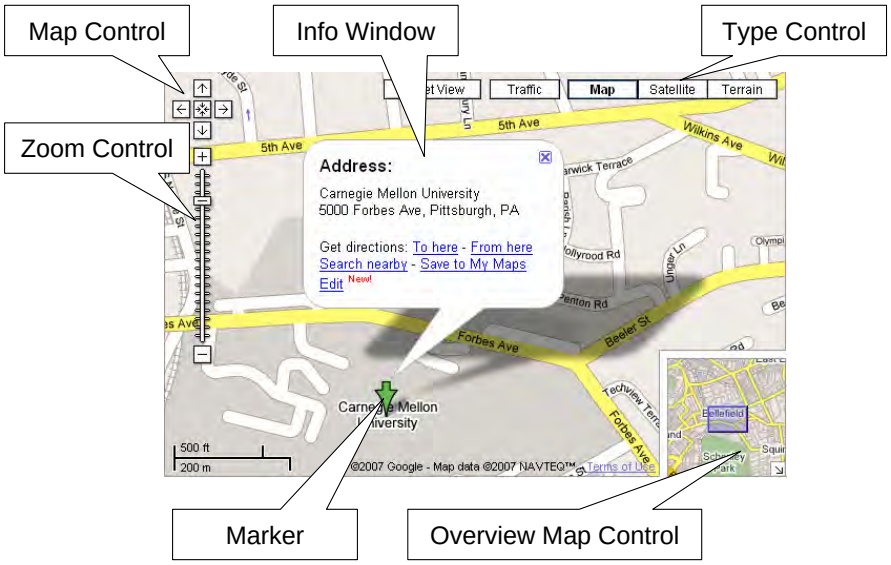

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
19



© 2007 Carnegie Mellon University

Google Maps Widgets



© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
21

Google Maps API Fundamentals

http://code.google.com/apis/maps/documentation/

- Very well documented, lots of examples
- Start simple (like this demo)
- Requires very basic javascript and HTML knowledge

General flow:

- Include the source code
- Create the map
- Drop markers onto the map

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 22

About keys and data

In order to include the library source, you need a key

- The key uniquely identifies your URL
- Not necessary when serving via a file:// URL

Doesn't the data get posted up to Google?

- No, Google only sees you requests for the underlying map images
- All marker placement and labeling is done local to the client with overlays

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
23



© 2007 Carnegie Mellon University

geo-dst.html (part 1)

```
<html><head><title>IP Geolocation Example</title>
<script src="http://maps.google.com/maps?file=api&amp;v=2&amp;key="
    type="text/javascript"></script>

<script type="text/javascript">
//  This is the file that contains the point data
var map;
var xmlFile = "geo-dst.xml";
//  Called when the map is loaded.  This function
//  creates the map, adds controls to it, and then
//  the points are laid on top of the map
function load() {
  if (GBrowserIsCompatible()) {
    map = new GMap2(document.getElementById("map"));
    map.addControl(new GLargeMapControl());
    map.addControl(new GOverviewMapControl());
    map.addControl(new GMapTypeControl());
    map.setCenter(new GLatLng(38, -97), 1);
    loadpoints();
  }
}
```


© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 25

geo-dst.html (part 2)

```
//  http://code.google.com/apis/maps/documentation/services.html#XML_Requests
function loadpoints() {
  GDownloadUrl(xmlFile, function(data, responseCode) {
    var xml = GXml.parse(data);
    var markers = xml.documentElement.getElementsByTagName("marker");
    for (var i = 0; i < markers.length; i++) {
      var point = new GLatLng(parseFloat(markers[i].getAttribute("lat")),
                              parseFloat(markers[i].getAttribute("lng")));
      descr = markers[i].getAttribute("lbl")+"; "+markers[i].getAttribute("val");
      map.addOverlay(new GMarker(point, {title:descr, clickable:false }));
    }
  });
}
</script></head>

<body onload="load()" onunload="GUnload()"><h2>IP Geolocation Example</h2>
  <div id="map" style="width: 640px; height: 480px"></div>
</body>
</html>
```


© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 26

The Results...

IP Geolocation Example



© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
27

Customizing Marker Icons

Two modifications needed

- Define the different icons upon initialization
- Choose the icon when points are added

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
28

geo-dst-v2.html (part 1)

```
...
function load() {
  if (GBrowserIsCompatible()) {
    map = new GMap2(document.getElementById("map"));
    map.addControl(new GLargeMapControl());
    map.addControl(new GOverviewMapControl());
    map.addControl(new GMapTypeControl());
    map.setCenter(new GLatLng(38, -97), 1);

    //create different pins
    sredicon.image = "green-s.png";
    sredicon.shadow = "shadow-s.png";
    sredicon.iconSize = new GSize(8, 13);
    sredicon.shadowSize = new GSize(14, 13);
    sredicon.iconAnchor = new GPoint(4, 12);
    sredicon.infoWindowAnchor = new GPoint(5, 1);

    mredicon.image = "red-m.png";
    mredicon.shadow = "shadow-m.png";
    mredicon.iconSize = new GSize(12, 20);
    ...
    loadpoints();
  }
}
```


© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
29

geo-dst-v2.html (part 2)

```
//  http://code.google.com/apis/maps/documentation/services.html#XML_Requests
function loadpoints() {
  GDownloadUrl(xmlFile, function(data, responseCode) {
    var xml = GXml.parse(data);
    var markers = xml.documentElement.getElementsByTagName("marker");
    for (var i = 0; i < markers.length; i++) {
      var point = new GLatLng(parseFloat(markers[i].getAttribute("lat")),
                              parseFloat(markers[i].getAttribute("lng")));
      ...

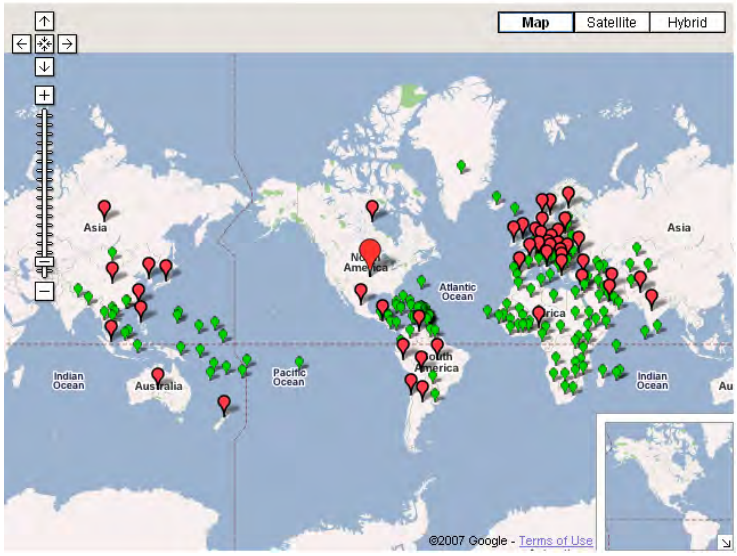
      var ratio = Math.log (  parseFloat(markers[i].getAttribute("val")) /
        minval) / Math.log (maxval / minval) ;
      //
      //  Plot the pin corresponding to the logarithmic ratio
      //
      if (ratio < 0.2) {
        map.addOverlay(new GMarker(pointlist[i], {icon:sredicon, title:de...
      } else if (ratio < 0.9) {
        map.addOverlay(new GMarker(pointlist[i], {icon:mredicon, title:de...
      } else {
        map.addOverlay(new GMarker(pointlist[i], {icon:lredicon, title:de...
      }
    }
   ...
```


© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 30

The Results...

IP Geolocation Example



© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
31

Adding Links

Need a new data set

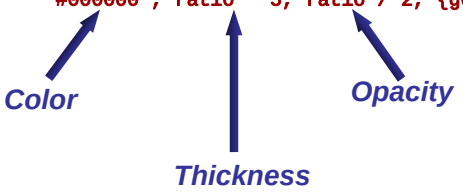
- Create an XML file with source location, destination location and value
- Add a new function to read and plot the data file

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
32

geo-dst-v3.html

```
function loadlinks() {
  GDownloadUrl(xmlFile, function(data, responseCode) {
    ...
    var slink = new GLatLng(parseFloat(links[i].getAttribute("slat")),
                            parseFloat(links[i].getAttribute("slng")));
    var elink = new GLatLng(parseFloat(links[i].getAttribute("elat")),
                            parseFloat(links[i].getAttribute("elng")));
    map.addOverlay (new GPolyline ([slink, elink],
                    "#000000", ratio * 5, ratio / 2, {geodesic:true}));
    ...
```



© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon © 2007 Carnegie Mellon University 33

The Results...

IP Geolocation Example



© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
34

Where to go from here

Make it your own

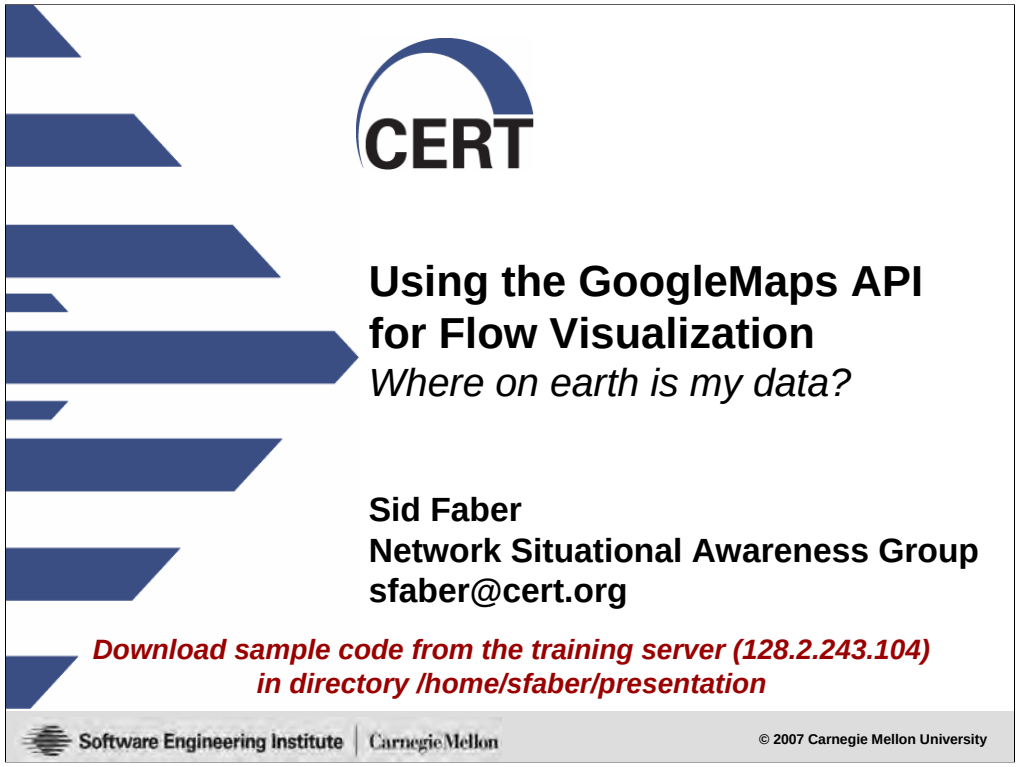
- Generate info window popups
- Drag markers
- Add driving directions

See http://code.google.com/apis/maps/

Download sample code from the training server (128.2.243.104) in /home/sfaber/presentation

© 2007 Carnegie Mellon University

CERT | Software Engineering Institute | Carnegie Mellon
© 2007 Carnegie Mellon University
35



© 2007 Carnegie Mellon University

Flow Analysis in a Wireless Environment with short DHCP Leases

Sanket Parikh

John McHugh

Dept. of Computer Science

Dalhousie University

DALHOUSIE UNIVERSITY
Inspiring Minds

Project Objectives

- Analysis of Wireless Network Data from University of Dartmouth (Crawdad Archive)
- Adding MAC Layer information in Net Flow tools for identification of nodes and Activities performed by a node.
- Return converted flow data to the Crawdad archive.

DALHOUSIE UNIVERSITY
Inspiring Minds

Project Rationale

- The main issue in analyzing wireless network data from many environments is the assignment of temporary IP Addresses using DHCP with short leases.
- The total user population often exceeds the available address space, and a given user may connect to the network for short sessions from a number of different locations making complicating per platform analyses.
- Work to date has concentrated on mobility rather than platform behaviour.

The Data

- 160 GB of compressed tcpdump packet headers.
- Collected continuously from 2 Nov 04 - 28 Feb 04
- 18 collection points academic, library, residence
- Nothing beyond IP Headers except TCP ports and flags, UDP ports.
- Anonymized with prefix preserving technique
  - Usage agreement precludes attacking anonymization to determine user identity.
  - Low order 24 bits of MAC also anonymized
    - List of known wireless MAC addresses provided

DALHOUSIE UNIVERSITY
Inspiring Minds

Technical Approach - 1

- Tried to use vlan tag fields to avoid altering YAF record format.
- Use the Forward and Reverse vlan tag fields to get source and destination MAC addresses into the yafscii
- Since these are 16 bits use perfect hash of MAC
- Problems:
  - vlan tag is in unidirectional extension of flow. Need both, even for unidirectional flows.
  - would like to use with real time and when MAC set not completely known

DALHOUSIE UNIVERSITY
Inspiring Minds

Technical Approach

- We added MAC to the bidirectional flow root in yaf, with both source and destination MAC addresses.
- There are a number of subtleties here, including the use of memcopy that introduces field order dependencies (an IPv4 optimization) and the assumption that MAC flag implies vlanid not zero.
- Once the MAC addresses are into the yafscii output, we started converting it into SiLK for further data analysis
- Shortly after we finished, CERT added MAC address support to YAF and we will use it in the future.

Technical Approach

- We created a module *yafscii2tuc.c*
  - Inserts minimal perfect hash index of MAC in in / out
  - Adds sensor id from command line to identify the sniffers.
- We split the output of the *yafscii2tuc* into separate hourly streams and use *popen* to send each one to a separate invocation of *rwtuc* so that the resulting files are in a proper date hierarchy.
- We also use *rwsort* on the *rwtuc* output to ensure time order and because *rwtuc* does not compress.

DALHOUSIE UNIVERSITY
Inspiring Minds

Minimal perfect hashes

- A Minimal Perfect Hash maps a set of N unique strings into integers in [0...N-1]
  - Packages available on internet designed for null terminated strings
  - Modified for counted strings
  - Extracted all MACS from Dartmouth packet data
  - Grouped to bring common usages together, e.g. known wireless, gateways, etc. then created MPH
  - 17000+ MACs, 11,000+ with IP packets.
- Lookup is constant time, collision free

DALHOUSIE UNIVERSITY
Inspiring Minds

Remaining problems

- yaf does not deal with decreasing time well
  - In live capture, packets are always in increasing time order no matter what the clock says
  - In playback the same holds unless the file has been reordered.
  - Several Dartmouth sensors exhibit decreasing time, probably due to ntp or other clock adjustments.
- Data from one of the sensors “breaks” the pipe
  - This may be related to the time problem above or may be due to another problem
  - Truncated packets may lead to other pathologies in yaf

Next steps

- We want to reassign the IPs currently used to a consistent IP that is related to the MAC index.
- First we need to determine if any wireless IPs are associated with gateway MACs.
  - This would occur if a wireless unit talked to another wireless unit via a routed connection, e.g. units connecting via separate sniffers.
  - Start by creating sets for each MAC type and looking for intersections
  - May have to explore DHCP strategy in more detail.
- This is currently underway.

DALHOUSIE UNIVERSITY
Inspiring Minds

MAC types

- There are 5 categories of MACS actively involved
  - Known Wireless MACs with IP traffic
  - Other MACs with IP packets
  - Multi cast MACs
  - Gateway MACs
  - Broadcast MACs
- A large number of MACs have no IP traffic
  - Some appear only at link layer, others in MAC list but not seen
- We used rwfilter to build sets for each type of MAC address based on the input and output field values

DALHOUSIE UNIVERSITY
Inspiring Minds

Project Outcomes

- We found some interesting information during analysis of the datasets. There are traces which shows some IP addresses appeared in two different sniffers located to different locations.
- The reason may be the physical location of sniffers for collecting data. Though sniffers were not located at proper distance from each other, there might be the chances for getting same IP traces in two different sniffers.
- This seems improbable and needs further study

DALHOUSIE UNIVERSITY
Inspiring Minds

Next Steps

- With the technique we used for this research should prove useful for similar data from wireless “hot spots”, airport, hotels and convention center networks and more.
- Same approach can be used to analyze data by using MAC layer information in Flow Analysis tools to identify the activities and movements of nodes in Wireless Networks.

DALHOUSIE UNIVERSITY
Inspiring Minds

Flow Visualization Using MS-Excel

Visualization for the Common Man

Presented by Lee Rock and Jay Brown
US-CERT Analysts
Einstein Program

Background

- US-CERT Mission
- Einstein Program
  - > Large volumes of traffic
  - > Architecture limitations
- Proactive vs. Reactive analysis
- Slow application certification process

Pro's and Con's

- Pro's:
  - Visualization allows for rapid analysis
  - Patterns are easy to identify
  - Flexibility in analysis
  - Most enterprises have MS Office (Excel)
- Con's:
  - Excel plotting engine is limited
  - Max of 65K records (recommend <= 50K)
  - Data must be imported and formatted
  - Memory management is an issue

Data Preparation Steps

- Data Pull
- Data Reduction
- Importing Data
- Data Formatting
- Sample analysis slides

Data Pull

Analysts have several options when trying to pull interesting datasets. Several methods we find useful are:

- Collecting data during non-business hours
  - Reduces traffic from users; helps expose automated sessions
- Search for outbound traffic only
  - Reduces noise from scanning, etc.
- Filtering for packets with the PSH/ACK flags set in the initial flags field
  - Focuses the traffic on sessions where data is actually transferred
- Filtering for packets with the SYN flag set in the initial flags field
  - Focuses on sessions initiated by your organization
- Limit traffic to records under 5K bytes
  - Most cyclical sessions (beaconing) happen in this range

Traffic should be refined to provide the best possible dataset for analysts to work with.

Data Reduction

To further enhance the concentration of suspicious data, analysts should:

- Remove replies from servers (responses to inbound server requests)
  - Looking for genuine outbound traffic
- Remove loud, common talkers (instant messenger, web crawlers, etc)
  - Reduces the noise, especially in web traffic
- “Whitelists” and “blacklists” are helpful for filtering

This is an iterative approach – Analyze, Research, Remove.

Importing Data

Data is imported from a pipe delimited text file

sIP	dIP	sPort	dPort	pro	packets	bytes	flags	sTime	dur	eT
10.147.82.96	10.130.166.158	80	5516	6	11	8258	FS PA	2007/10/29T15:07:38.807	0.290	2007/10/29T15:07:39.
10.140.165.218	10.54.98.176	80	5705	6	86	120321	FS PA	2007/10/29T15:07:40.875	0.552	2007/10/29T15:07:41.
10.95.46.146	10.34.134.191	80	5705	6	1	40	A	2007/10/29T15:07:42.473	0.000	2007/10/29T15:07:42.
10.94.132.147	10.168.141.231	80	49094	6	10	7348	FS PA	2007/10/29T15:16:19.666	0.825	2007/10/29T15:16:20.
10.226.143.219	10.162.254.83	80	49297	6	33	43588	FS PA	2007/10/29T15:16:23.020	0.498	2007/10/29T15:16:23.
172.25.4.165	10.161.142.75	80	47356	6	10	7047	FS PA	2007/10/29T15:28:21.421	0.859	2007/10/29T15:28:22.
10.120.9.241	10.36.140.83	80	47489	6	7	1839	FS PA	2007/10/29T15:28:23.285	0.326	2007/10/29T15:28:23.
10.192.192.130	10.124.26.9	5516	80	6	9	945	FS PA	2007/10/29T15:07:38.771	0.292	2007/10/29T15:07:39.
10.8.58.141	10.254.147.27	5705	80	6	55	3589	FS PA	2007/10/29T15:07:40.843	1.591	2007/10/29T15:07:42.
10.215.49.170	10.30.5.168	49094	80	6	8	969	FS PA	2007/10/29T15:16:19.638	0.825	2007/10/29T15:16:20.
10.207.158.173	10.15.150.60	49297	80	6	22	1666	FS PA	2007/10/29T15:16:23.004	0.488	2007/10/29T15:16:23.
192.168.45.69	10.88.159.210	47356	80	6	9	1796	FS PA	2007/10/29T15:28:21.381	0.874	2007/10/29T15:28:22.
10.227.193.146	10.237.117.172	47489	80	6	7	888	FS PA	2007/10/29T15:28:23.261	0.313	2007/10/29T15:28:23.
10.115.234.230	10.144.241.122	80	24503	6	10	6735	FS PA	2007/10/29T16:01:28.698	0.220	2007/10/29T16:01:28.
10.52.224.171	10.232.170.176	80	24601	6	7	1475	FS PA	2007/10/29T16:01:29.421	0.236	2007/10/29T16:01:29.
10.144.199.78	10.208.138.229	80	64021	6	10	6437	FS PA	2007/10/29T16:09:08.791	0.161	2007/10/29T16:09:08.
10.9.152.19	10.233.152.178	80	64124	6	7	1310	FS PA	2007/10/29T16:09:09.883	0.247	2007/10/29T16:09:10.
10.116.235.116	10.27.192.234	80	64021	6	1	40	A	2007/10/29T16:09:08.951	0.000	2007/10/29T16:09:08.
10.0.158.212	10.131.10.198	80	40079	6	11	6378	FS PA	2007/10/29T16:16:40.586	0.247	2007/10/29T16:16:40.
10.40.145.167	10.229.195.82	80	40167	6	15	15095	FS PA	2007/10/29T16:16:41.718	0.317	2007/10/29T16:16:42.
10.40.157.25	10.12.36.164	80	18275	6	10	6242	FS PA	2007/10/29T16:24:32.546	0.235	2007/10/29T16:24:32.
10.33.232.60	10.224.241.212	80	18385	6	21	21877	FS PA	2007/10/29T16:24:34.100	0.409	2007/10/29T16:24:34.
10.75.204.191	10.52.57.127	24503	80	6	8	1439	FS PA	2007/10/29T16:01:28.654	0.223	2007/10/29T16:01:28.
10.6.83.30	10.218.84.41	24601	80	6	7	888	FS PA	2007/10/29T16:01:29.393	0.225	2007/10/29T16:01:29.
10.123.207.187	10.211.245.126	64021	80	6	9	2544	FS PA	2007/10/29T16:09:08.762	0.153	2007/10/29T16:09:08.
10.239.204.27	10.140.242.63	64021	80	6	1	40	R	2007/10/29T16:09:08.951	0.000	2007/10/29T16:09:08.

Data Formatting

Columns within the spreadsheet should be aligned to each field of the flows, Einstein data is formatted to encompass:

- Source IP
- Destination IP
- Source Port
- Destination Port
- Protocol
- Packets
- Bytes
- Flags
- Start Time
- Duration
- End Time
- Sensor
- Type
- Initial Flags

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	sIP	dIP	sPort	dPort	pro	packets	bytes	flags	sTime	dur	eTime	sensor	type	initialFlag
2	x.x.x.x	y.y.y.y	52611	22	6	2569	3385741	FS PA	2007/10/17T00:10:40.722	9.5	2007/10/17T00:10:50.222	X	out	S
3	x.x.x.x	y.y.y.y	7774	22	6	136	10750	PA	2007/10/17T00:08:28.293	1795.691	2007/10/17T00:38:23.984	X	out	A
4	x.x.x.x	y.y.y.y	7774	22	6	106	9046	PA	2007/10/17T00:38:36.714	1800.05	2007/10/17T01:08:36.764	X	out	PA
5	x.x.x.x	y.y.y.y	1688	22	6	1	41	A	2007/10/17T00:01:18.787	0	2007/10/17T00:01:18.787	X	out	A
6	x.x.x.x	y.y.y.y	1688	22	6	1	41	A	2007/10/17T00:06:18.690	0	2007/10/17T00:06:18.690	X	out	A
7	x.x.x.x	y.y.y.y	1688	22	6	1	41	A	2007/10/17T00:11:18.598	0	2007/10/17T00:11:18.598	X	out	A
8	x.x.x.x	y.y.y.y	1688	22	6	1	41	A	2007/10/17T00:16:18.514	0	2007/10/17T00:16:18.514	X	out	A
9	x.x.x.x	y.y.y.y	1688	22	6	1	41	A	2007/10/17T00:21:18.433	0	2007/10/17T00:21:18.433	X	out	A
10	x.x.x.x	y.y.y.y	1688	22	6	1	41	A	2007/10/17T00:26:18.349	0	2007/10/17T00:26:18.349	X	out	A
11	x.x.x.x	y.y.y.y	1688	22	6	1	41	A	2007/10/17T00:31:18.257	0	2007/10/17T00:31:18.257	X	out	A
12	x.x.x.x	y.y.y.y	1688	22	6	1	41	A	2007/10/17T00:36:18.164	0	2007/10/17T00:36:18.164	X	out	A
13	x.x.x.x	y.y.y.y	1281	22	6	956	40630	PA	2007/10/17T00:09:32.281	1798.994	2007/10/17T00:39:31.275	X	out	A
14	x.x.x.x	y.y.y.y	1688	22	6	1	41	A	2007/10/17T00:41:18.068	0	2007/10/17T00:41:18.068	X	out	A
15	x.x.x.x	y.y.y.y	1688	22	6	1	41	A	2007/10/17T00:46:17.971	0	2007/10/17T00:46:17.971	X	out	A

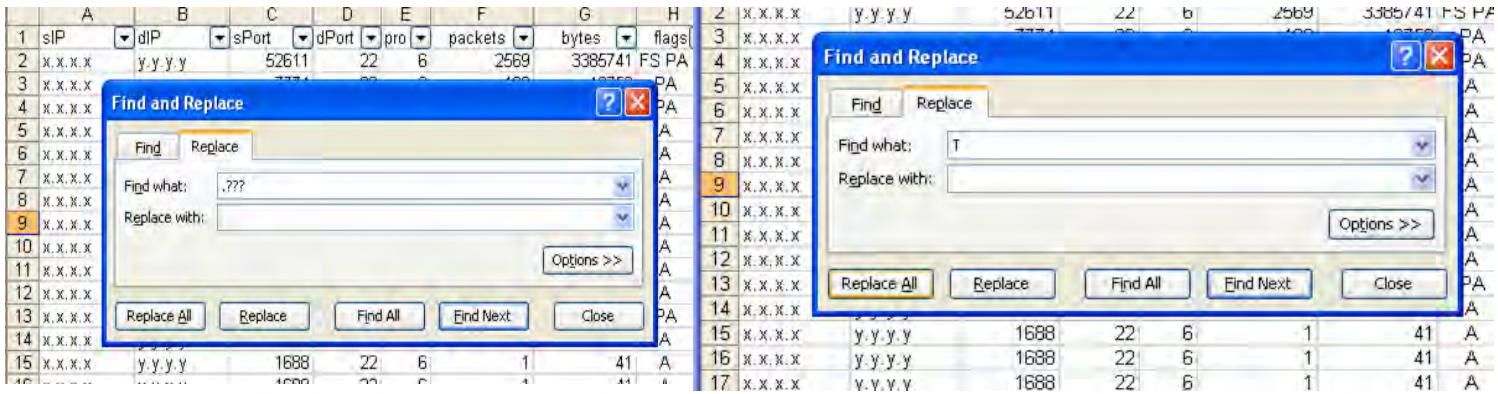
Data Formatting Cont.

US-CERT analysts use two methods to format the Einstein time fields into a format that is able to be plotted:

A: Use the - - legacy-timestamps switch to place the time in a

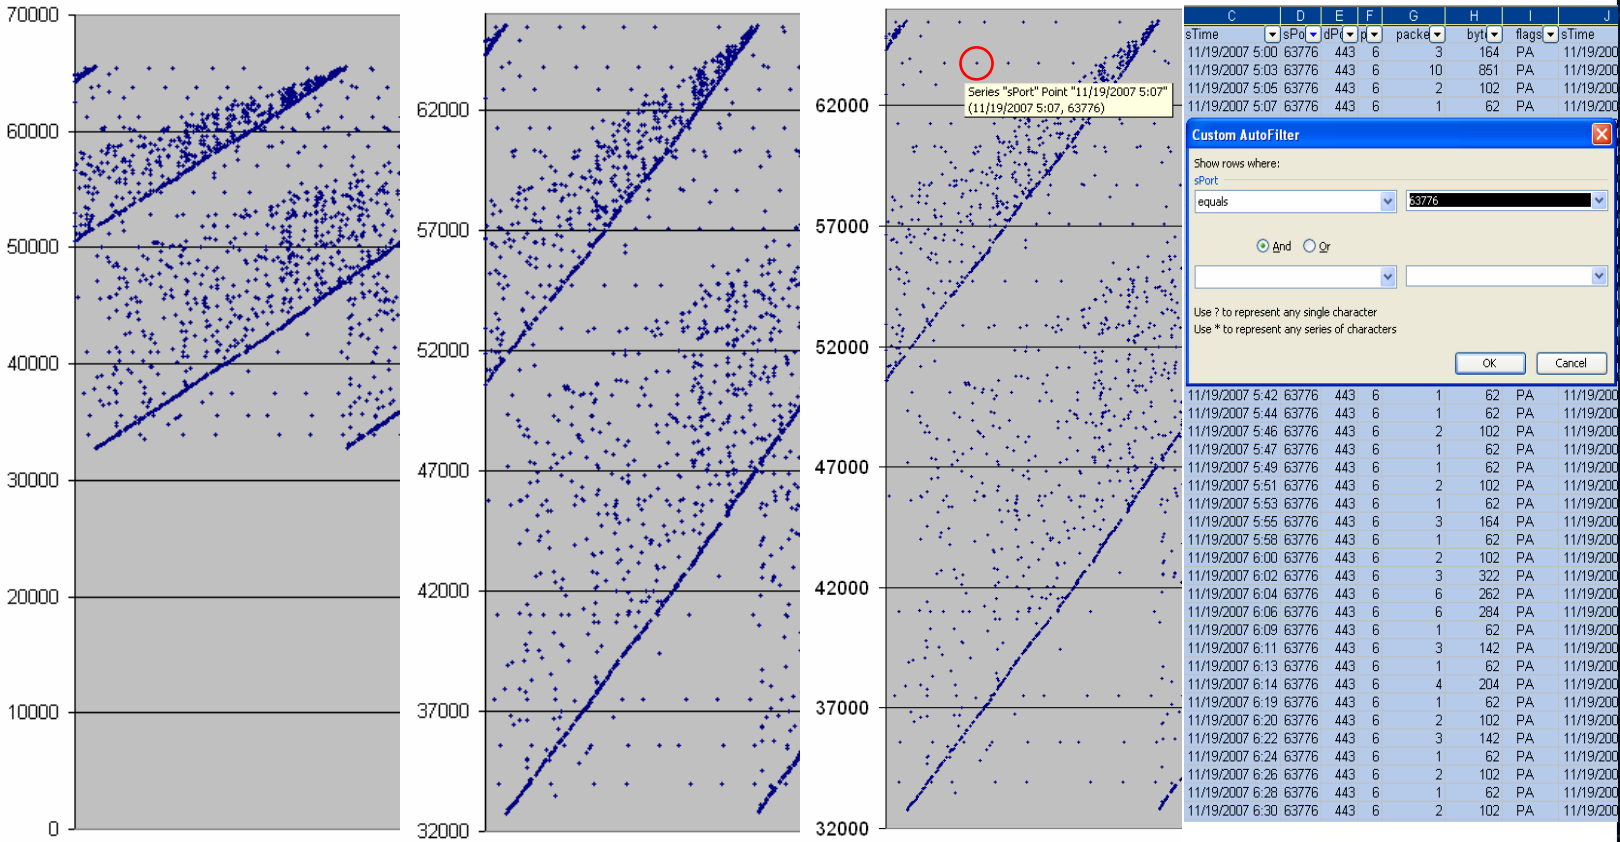
MM/DD/YYYY HH:MM:SS format from the default MM/DD/YYYYTHH:MM:SS.MMM

B: Utilize the replace function in excel to remove the milliseconds from the time and replace the T placeholder with a space:



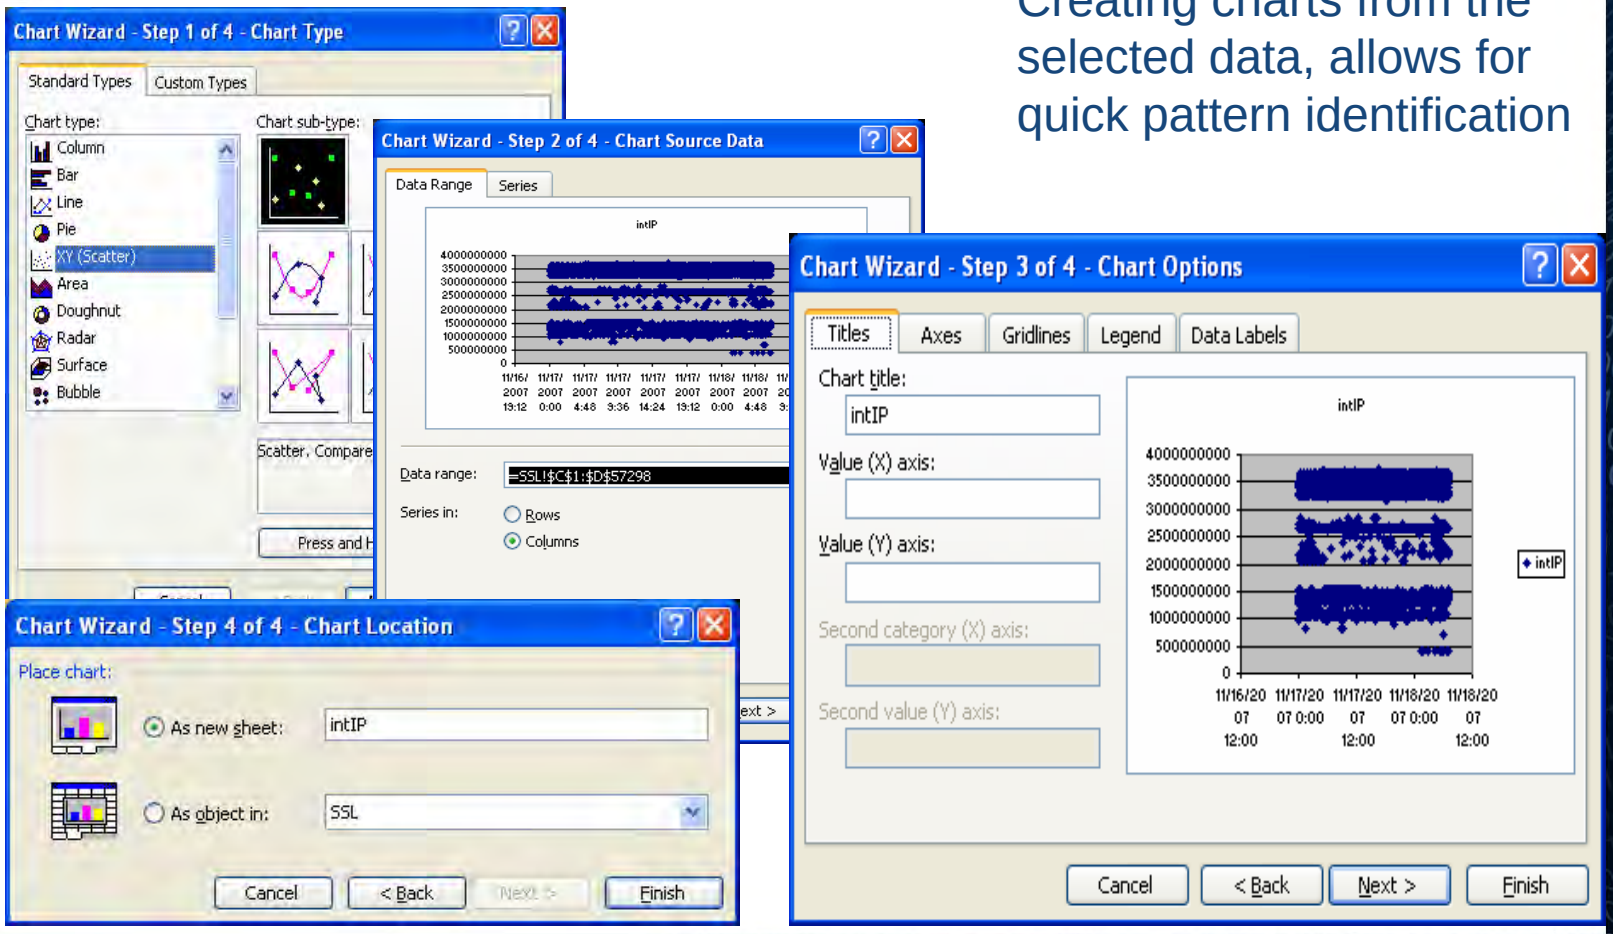
Analysis Workflow

Plot Zoom Highlight AutoFilter



Plot

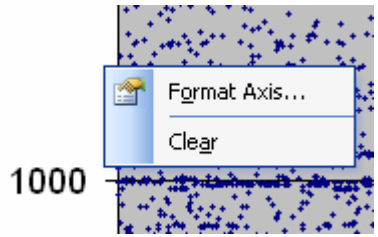
Creating charts from the selected data, allows for quick pattern identification

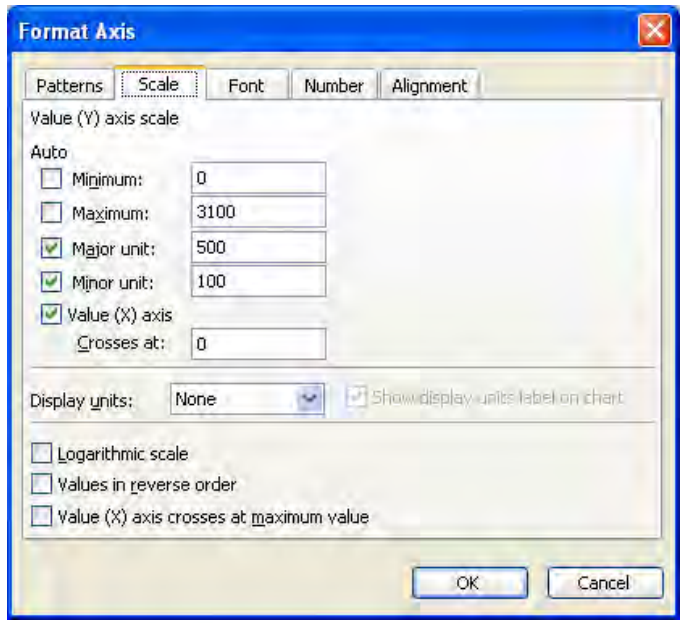


Zoom

You can “zoom” in to specific data points, by changing the scale of the axis

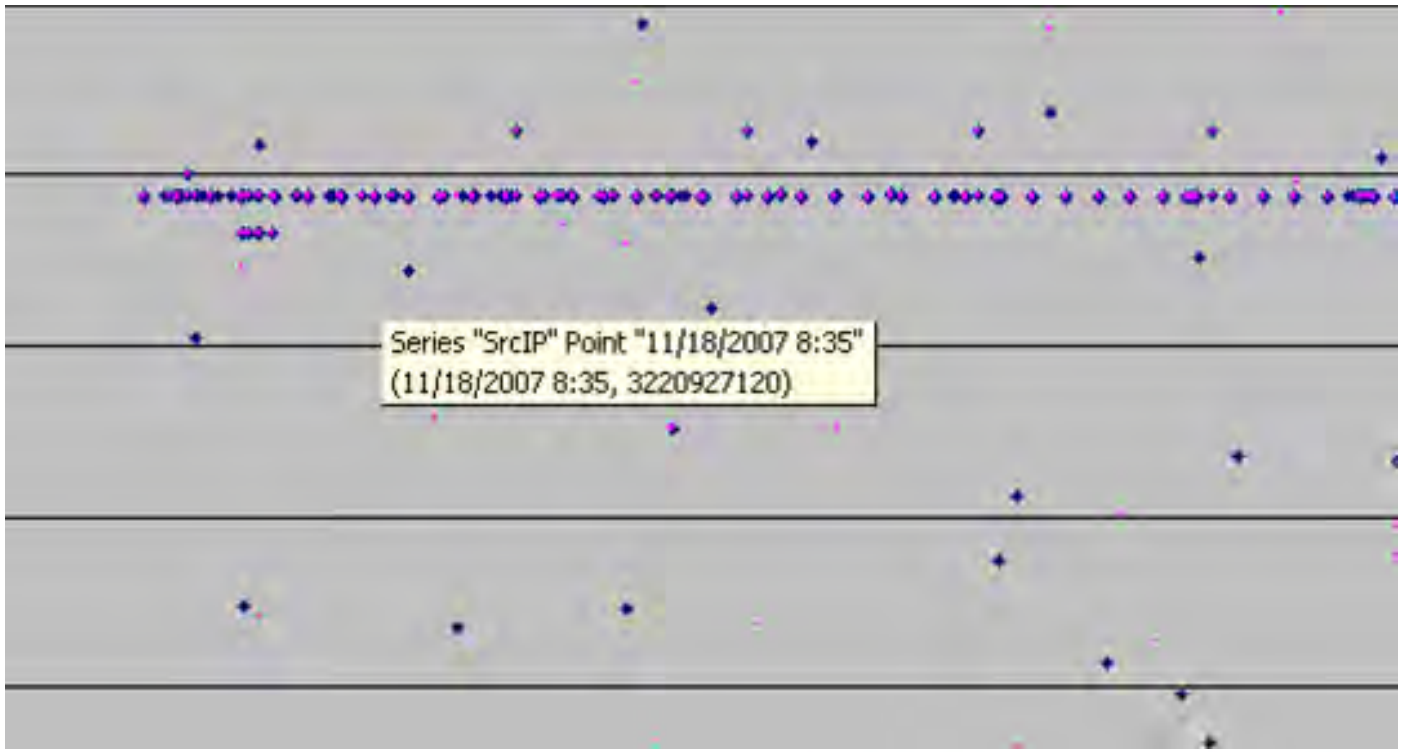
- Right click on the axis
- Select “Format Axis”
- Click on the “Scale” tab
- Adjust scale as desired
- Works for both axis
- Remember to remove





Highlight

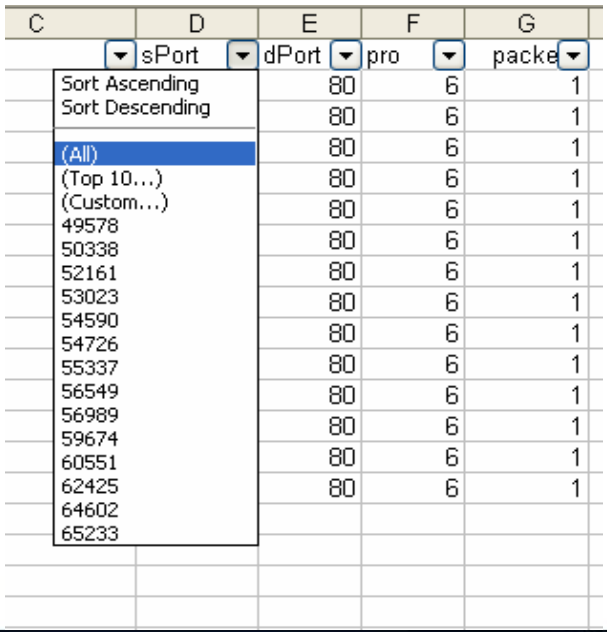
By hovering over a data point in the series an analyst can locate the point in the rest of the records by filtering for the displayed information



AutoFilter

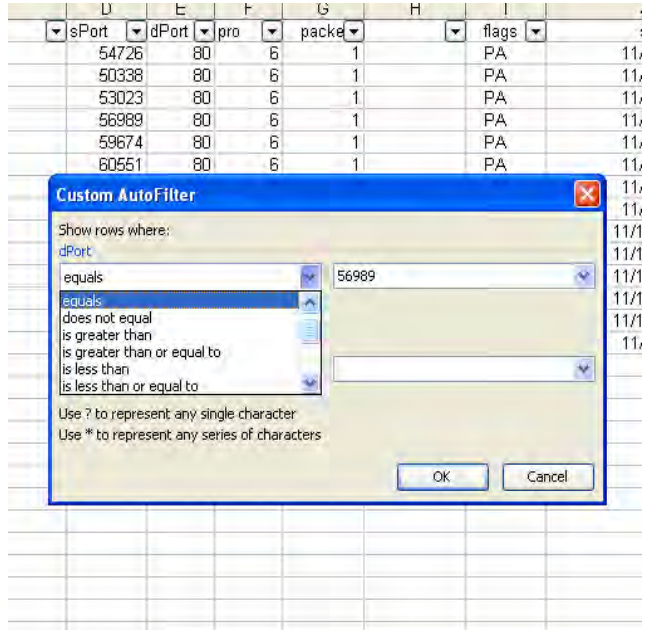
Method A – Drop down list:

Select the desired value from the drop down list



Method B – Custom Filter:

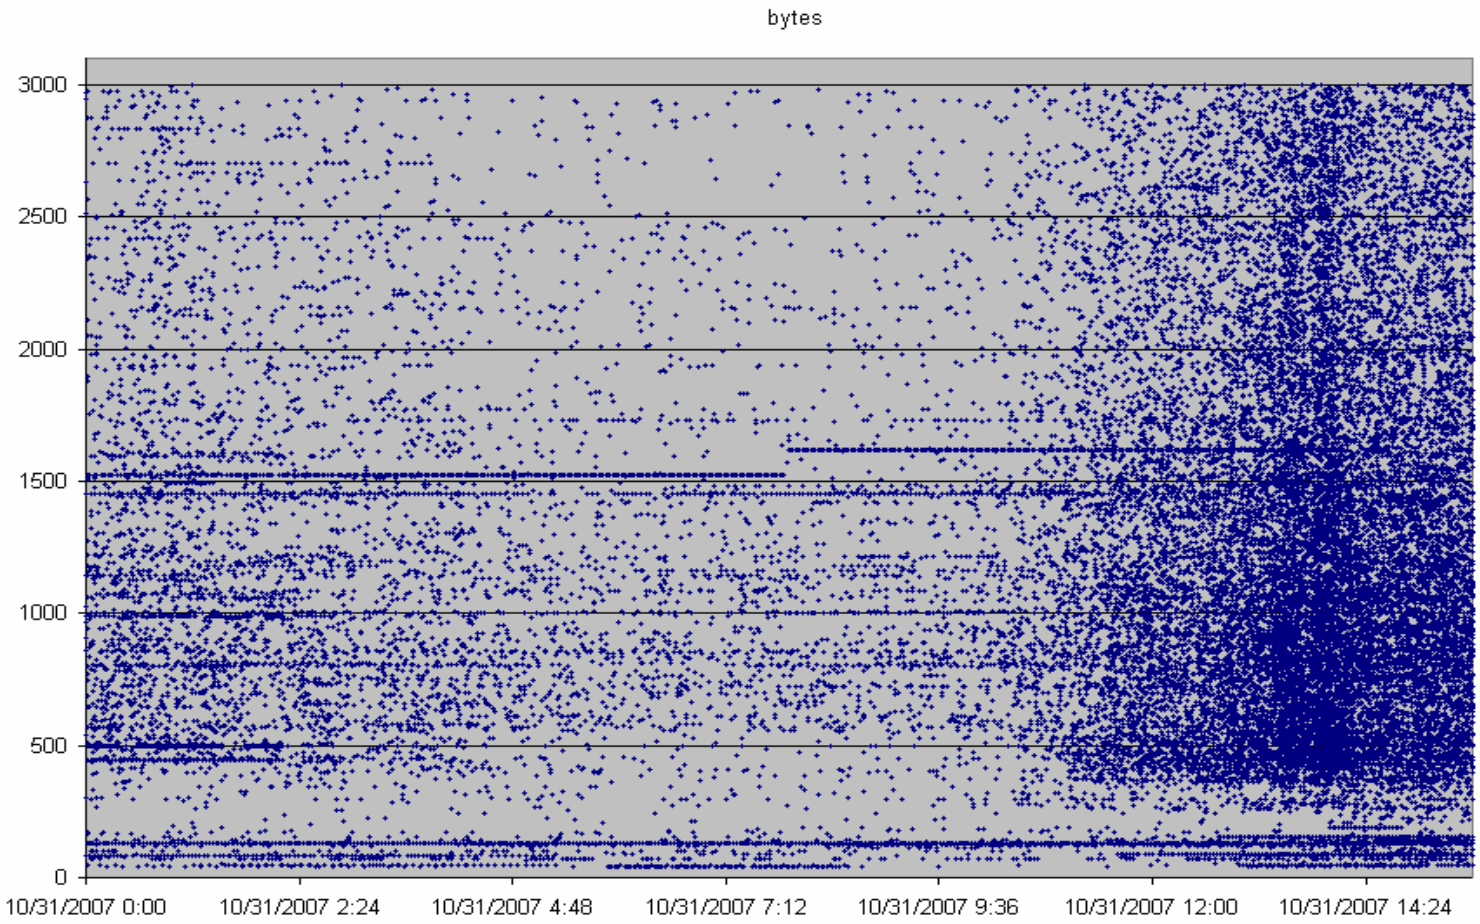
Select data by using Excel's built in boolean logic search functions



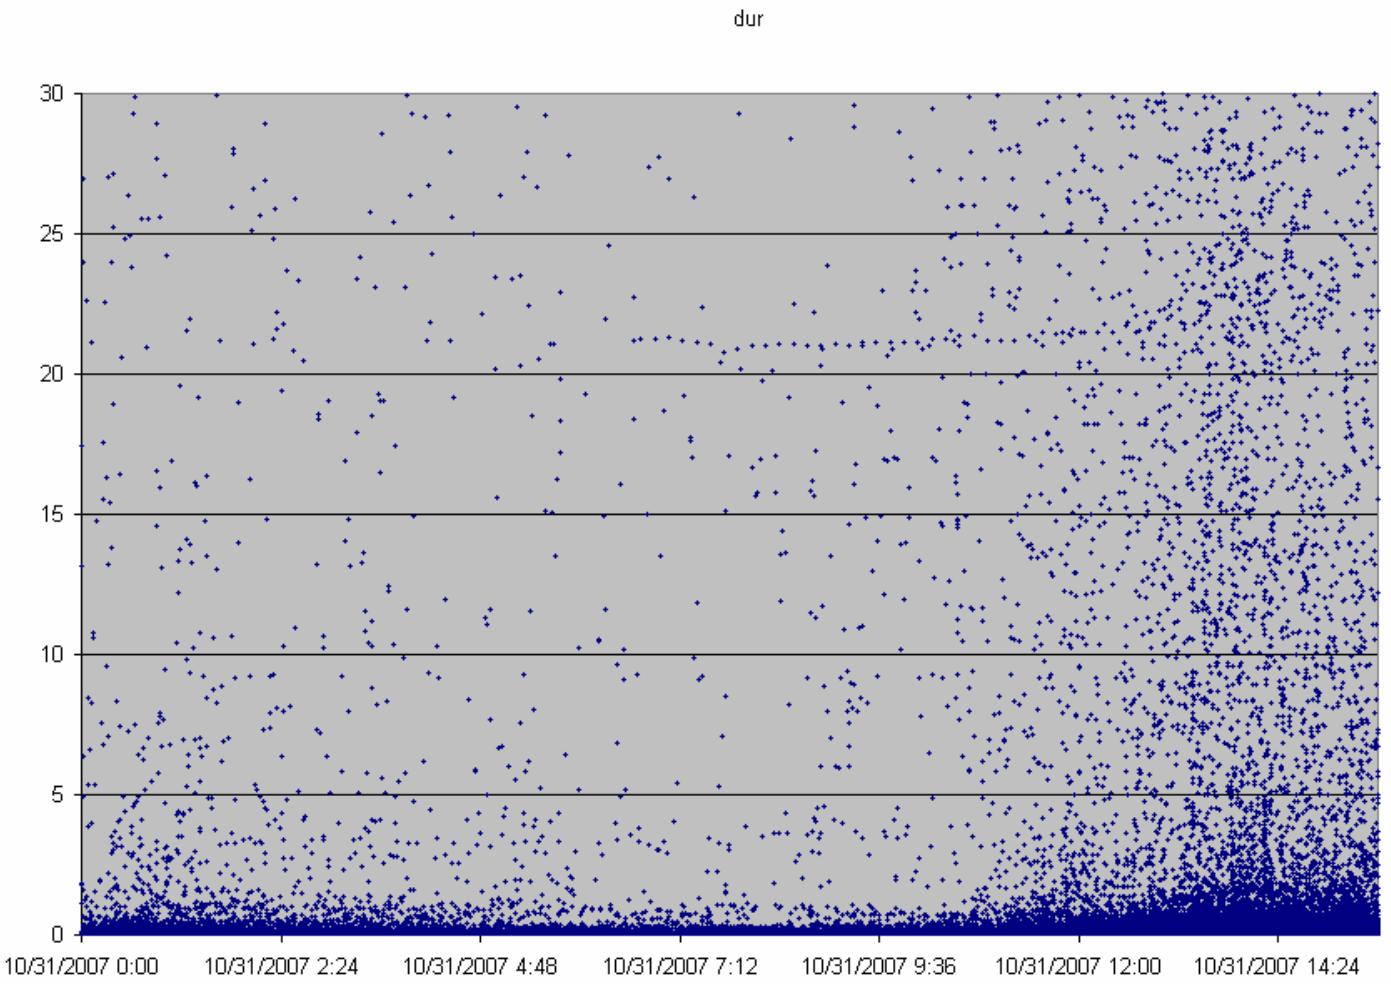
Sample Analysis Slides

- Scatter Plot Analysis
  - Byte Based Patterns
  - Duration Based Patterns
  - sPort vs. dPort Patterns
  - IP Based Patterns
  - Application Pattern

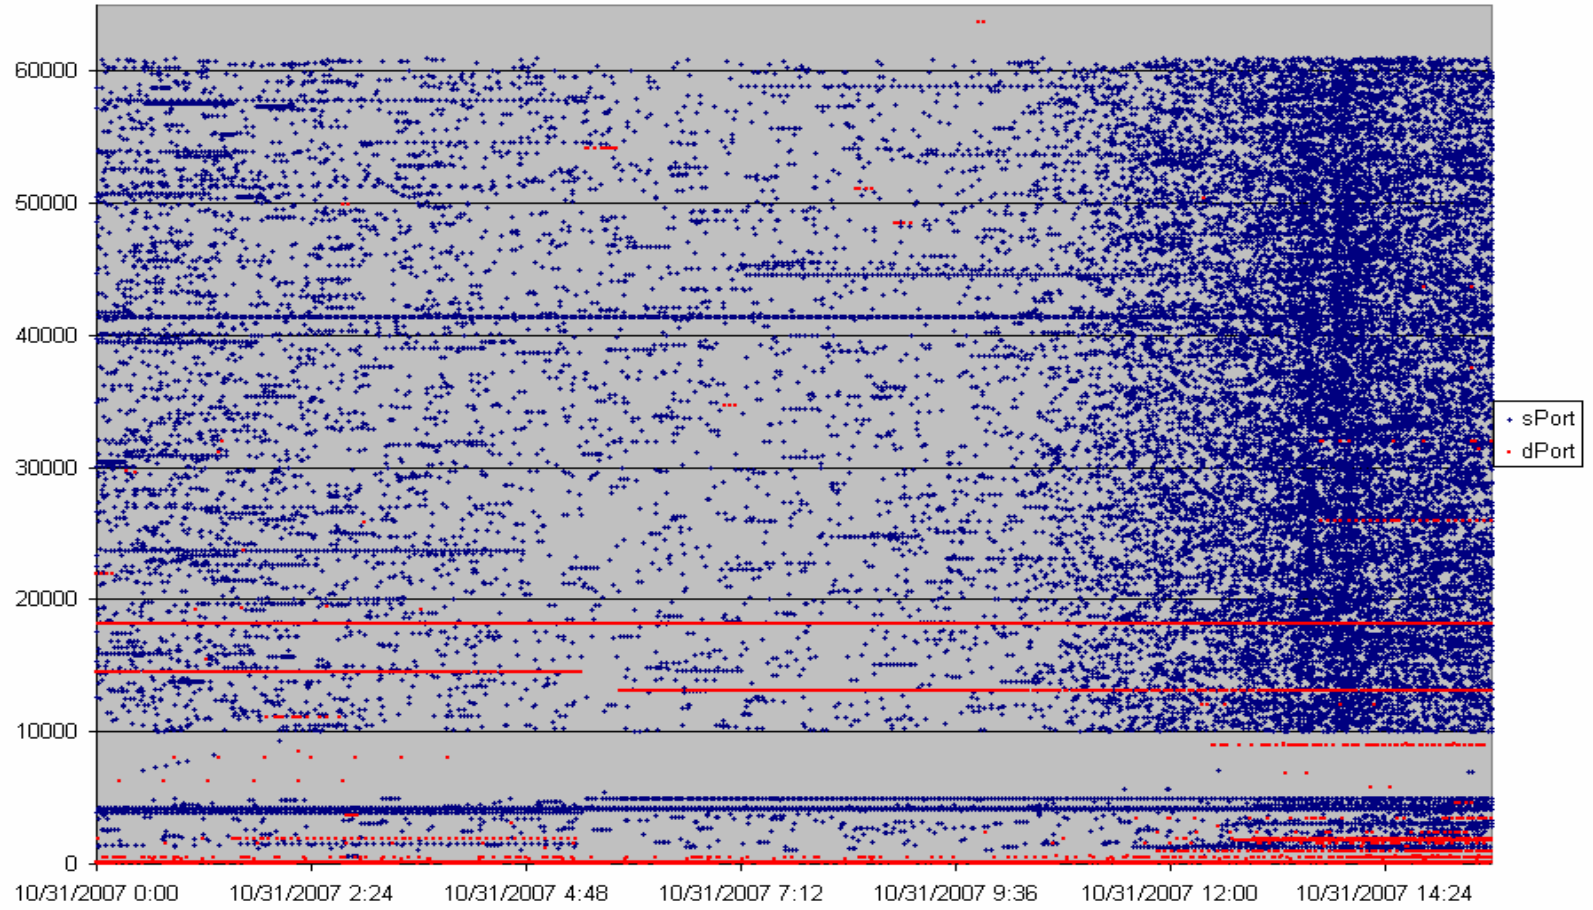
Byte Based Patterns



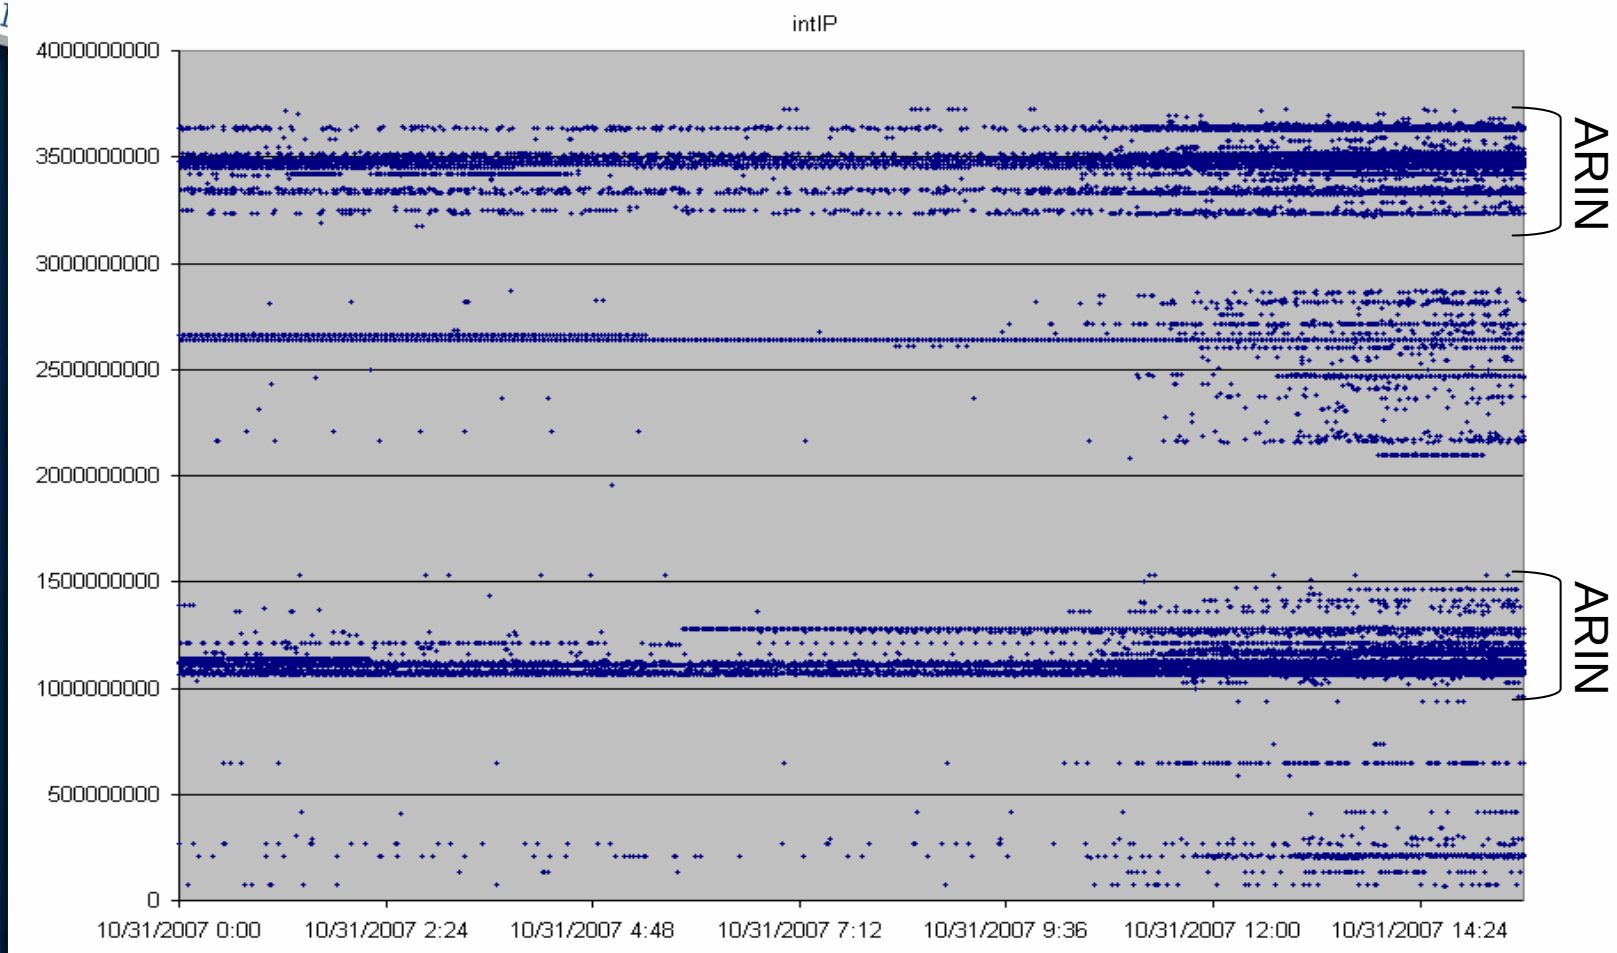
Duration Based Patterns



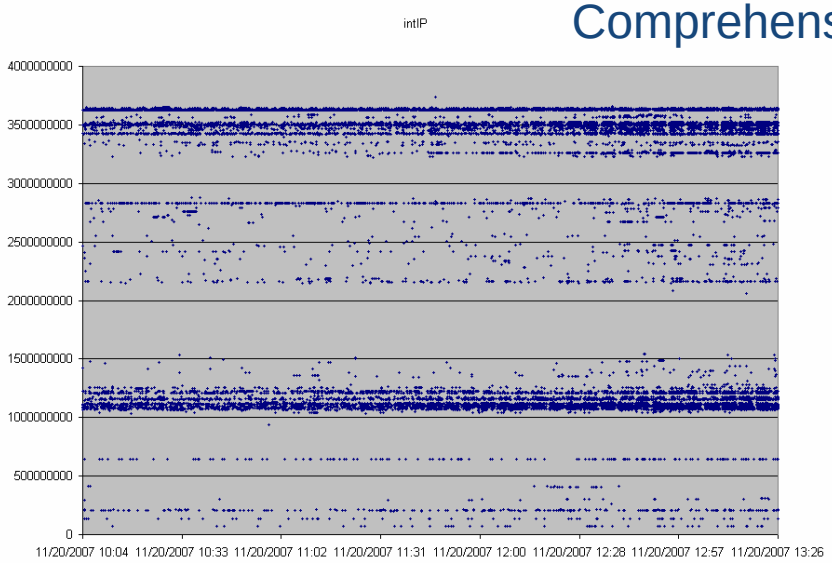
sPort vs. dPort

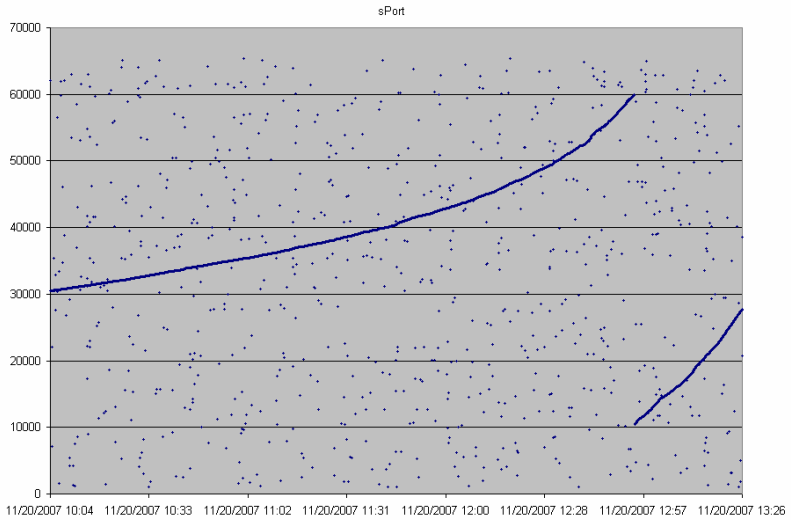


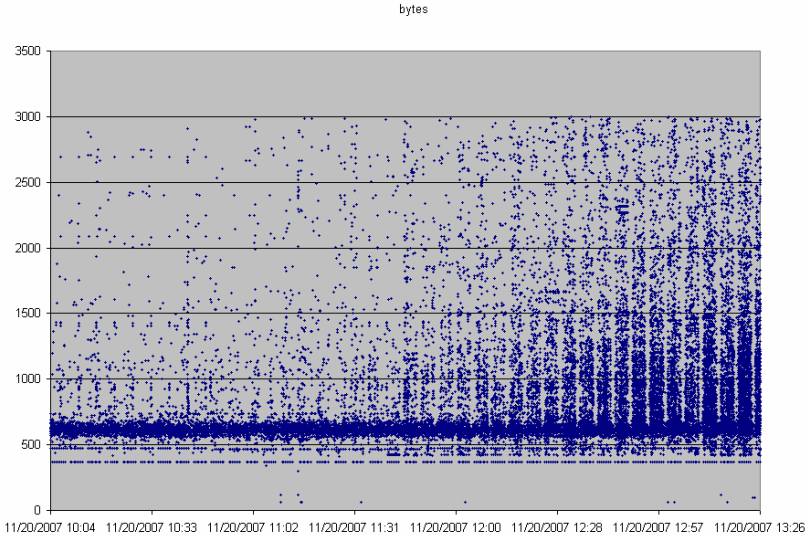
IP Integer Patterns

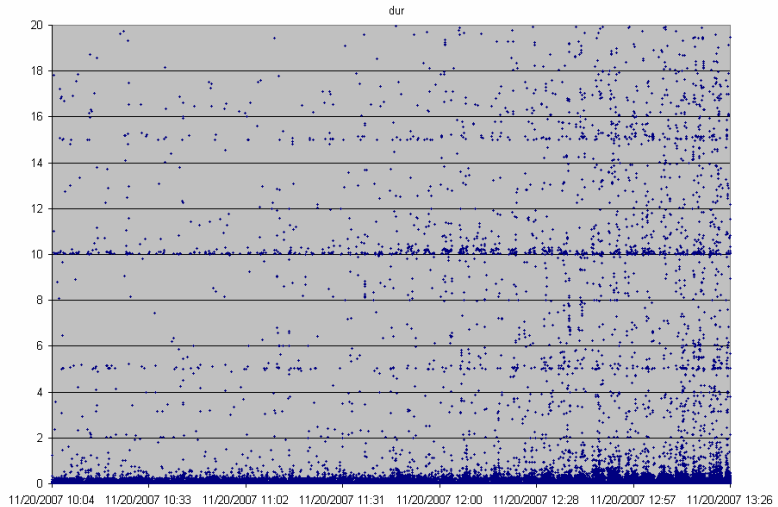


Comprehensive View

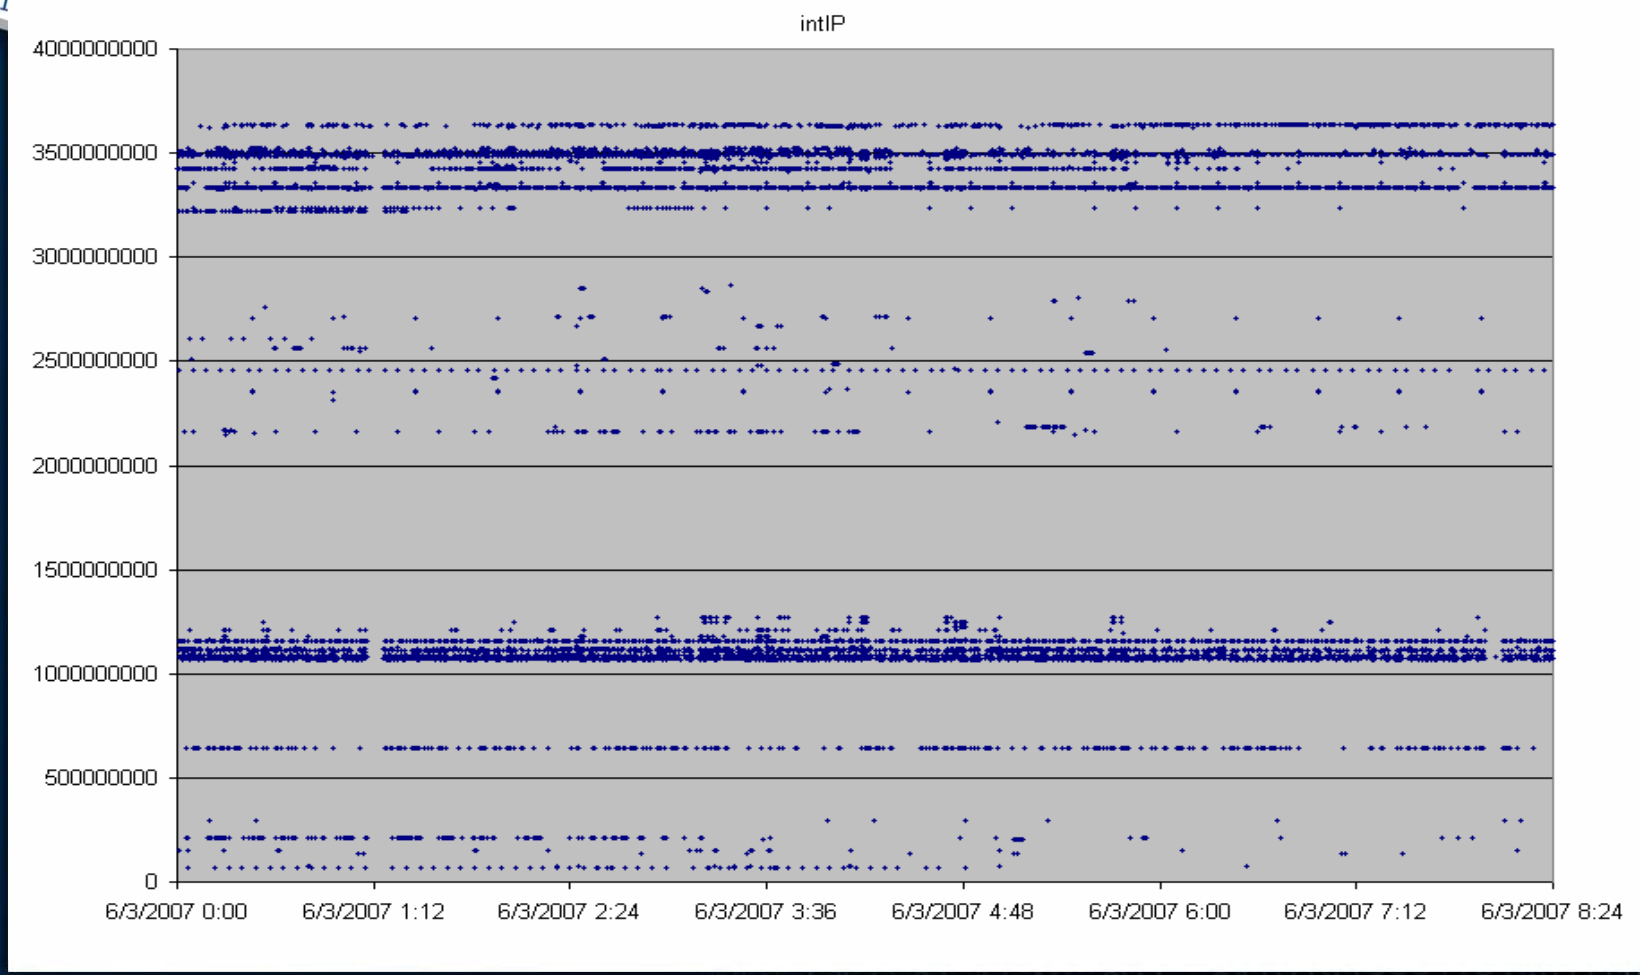




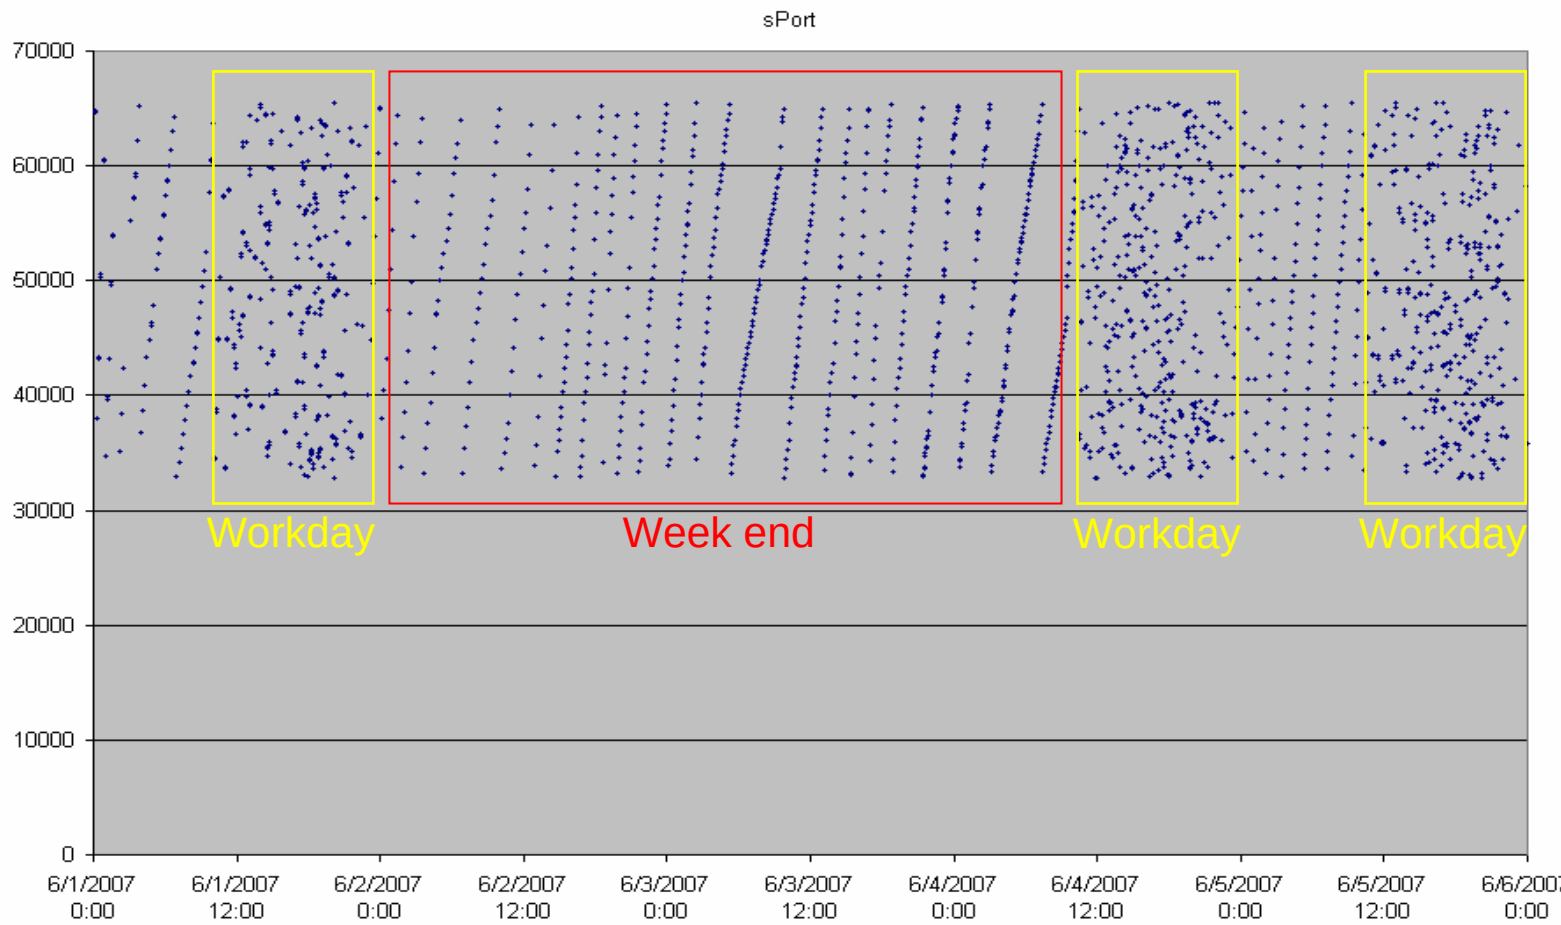




Case Study


intIP
4000000000
3500000000
3000000000
2500000000
2000000000
1500000000
1000000000
500000000
0
6/3/2007 0:00
6/3/2007 1:12
6/3/2007 2:24
6/3/2007 3:36
6/3/2007 4:48
6/3/2007 6:00
6/3/2007 7:12
6/3/2007 8:24

Multi-day View



Case Study Conclusion

After notifying the agency in question, the machines that were generating this traffic were found and forensically examined. The malware turned out to be a keystroke logger that posted data to a specific website and retrieved commands embedded on the same site. Prior to this incident, there was no malware associated with this site.

Additional Analysis

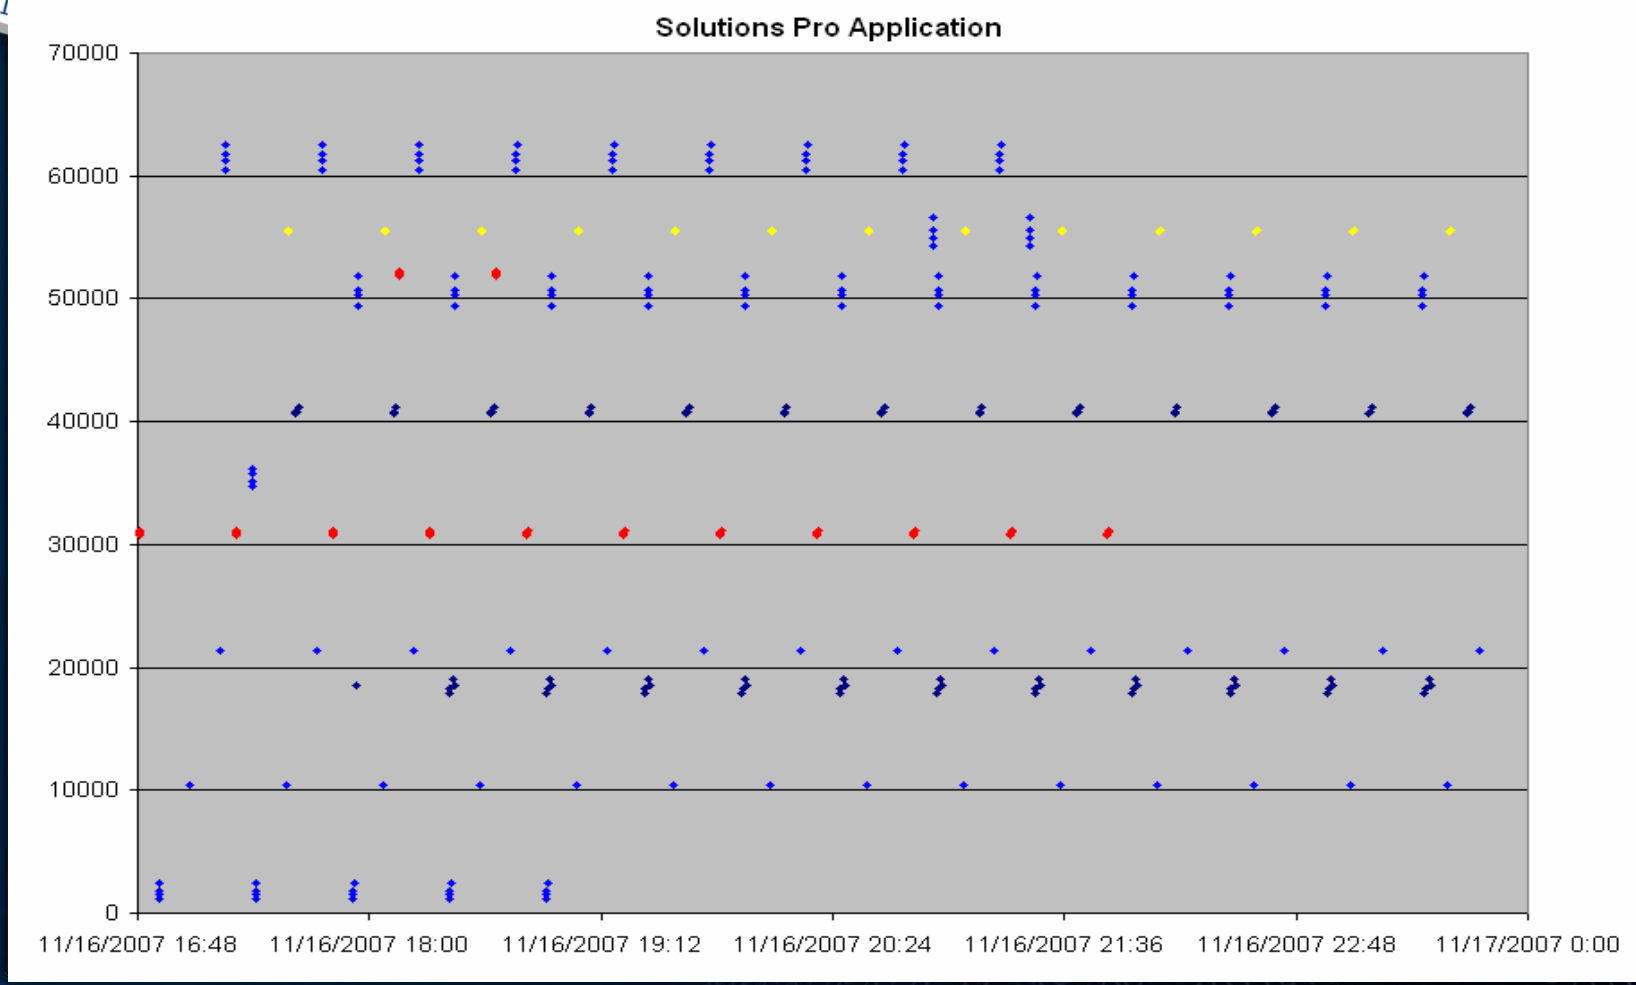
Determining application patterns

- Identifying specific applications

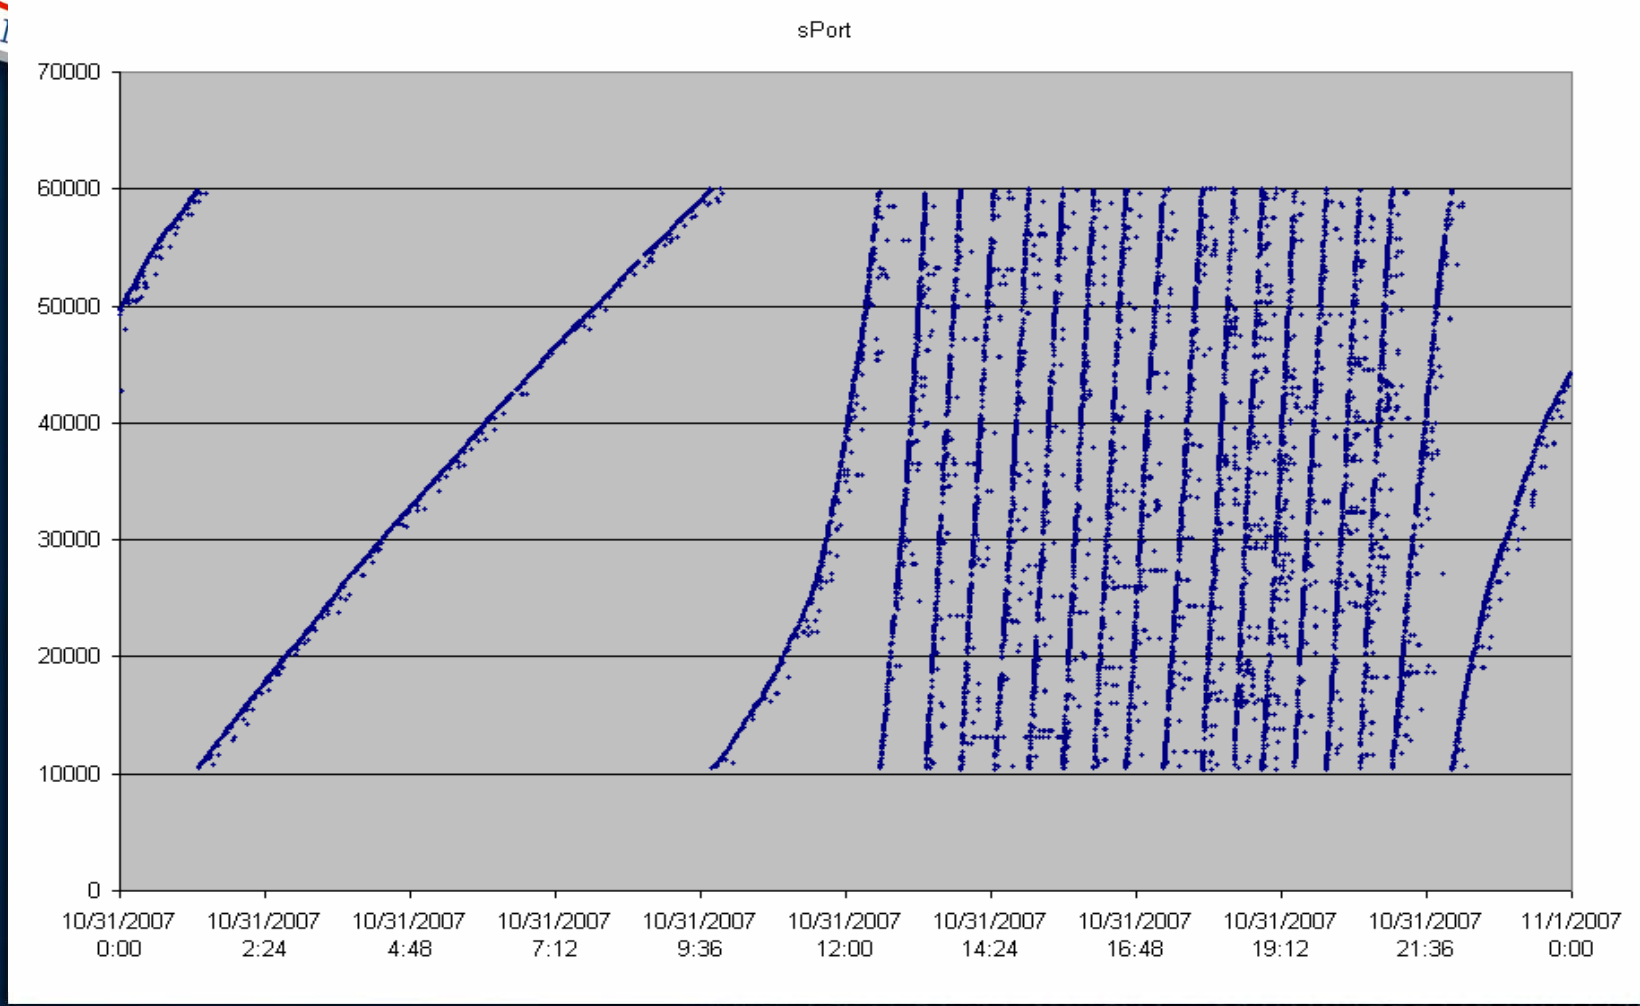
Working with gateway traffic

- Structured gateway
- Proxy gateway
- Gateway mannerisms

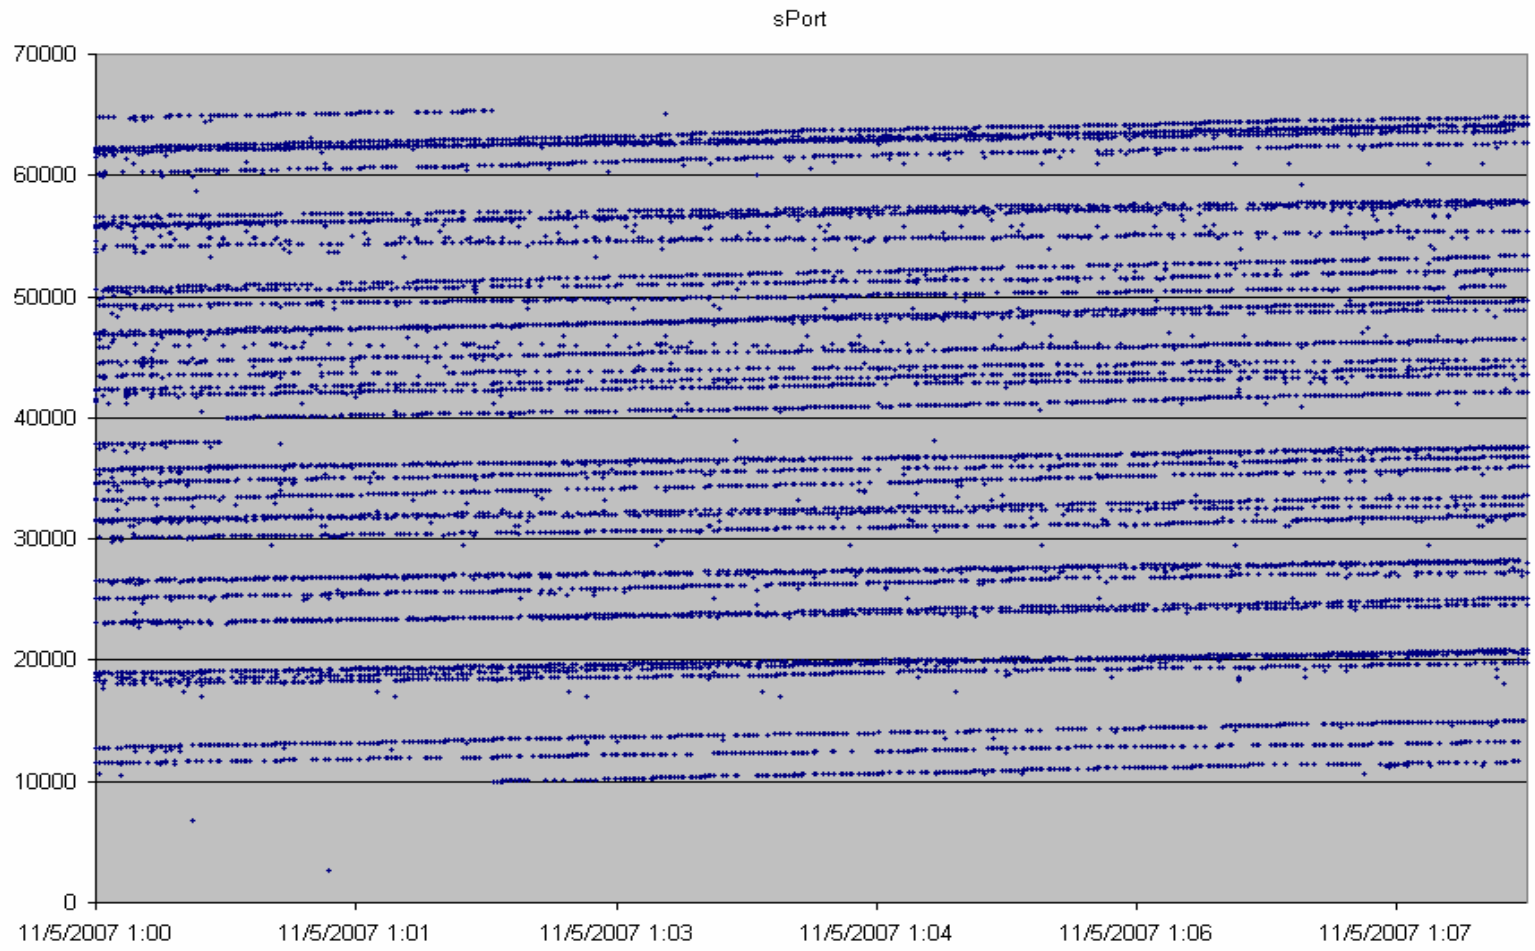
Application Patterns



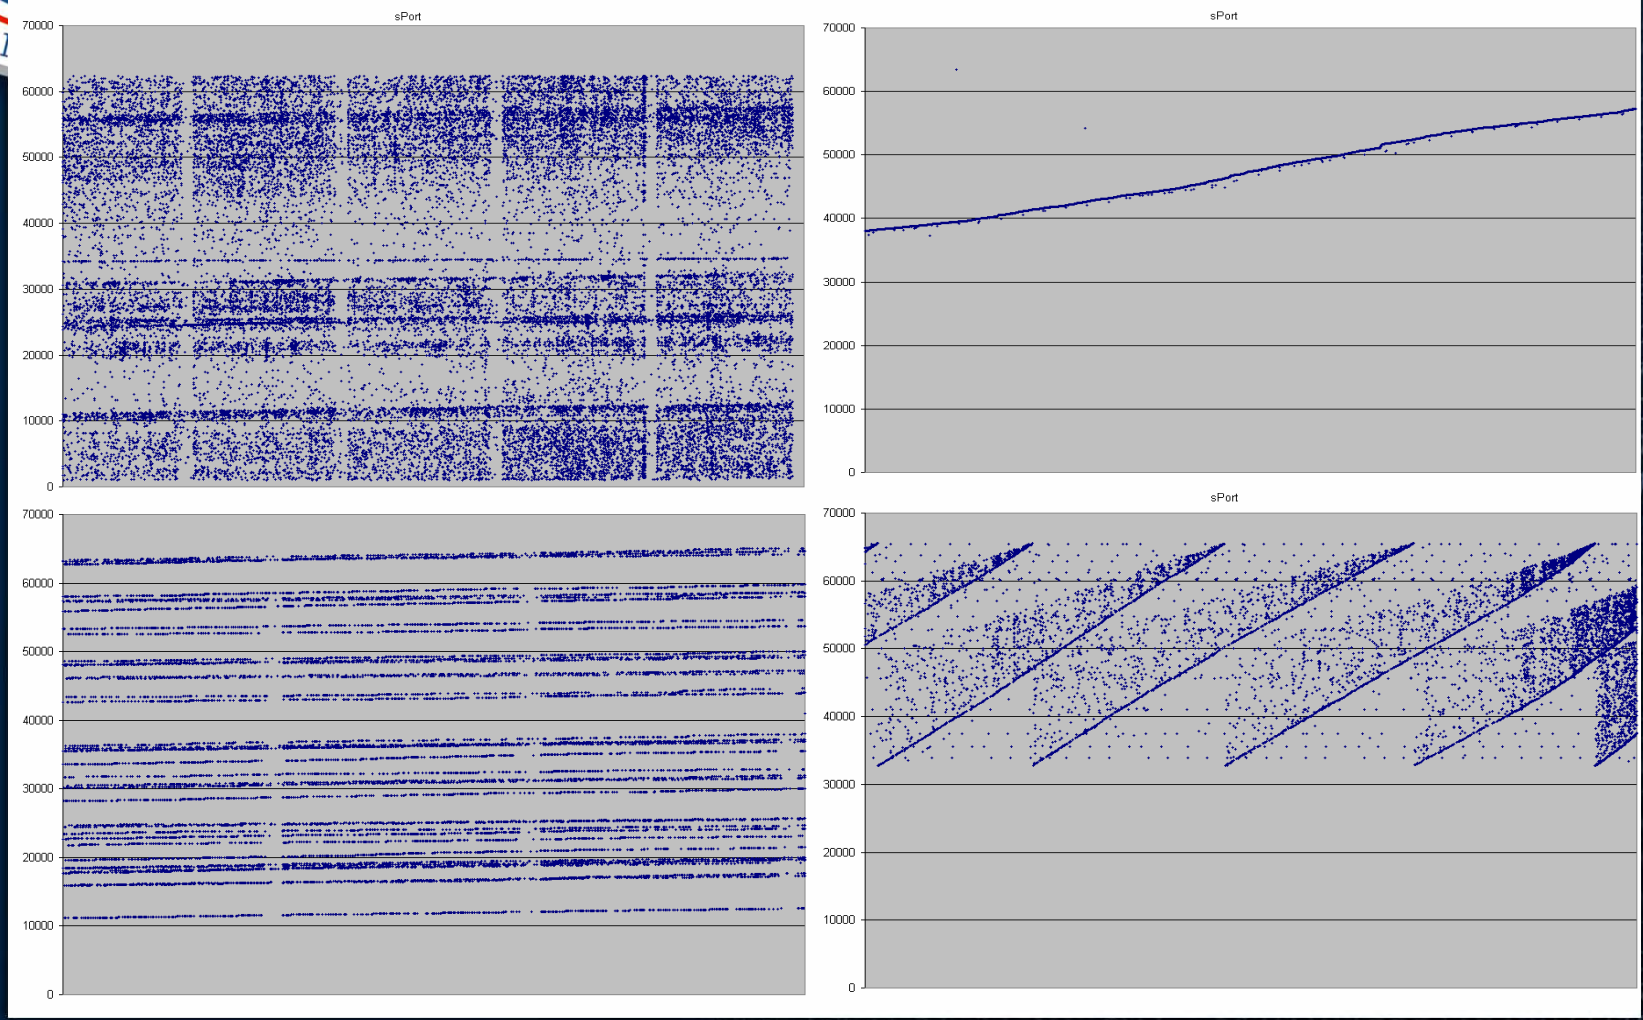
Structured Gateway



Proxy Gateway



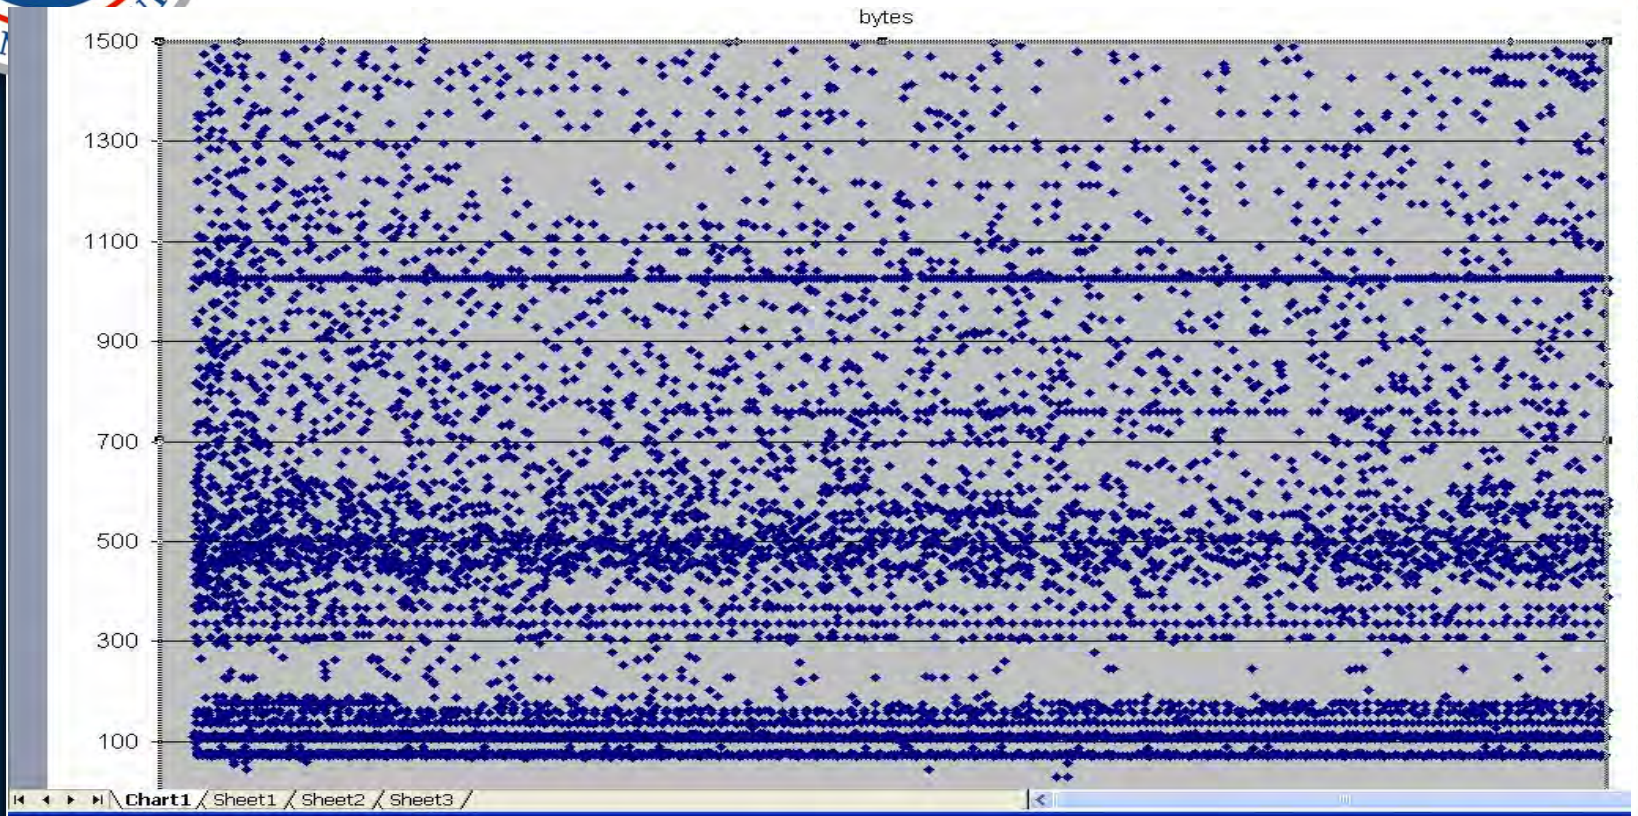
Gateway Mannerisms



Future Directions

- Split view analysis
- Coloring data
- Application coloring
- sPort colored by app
- Gateway coloring to IP

Split View

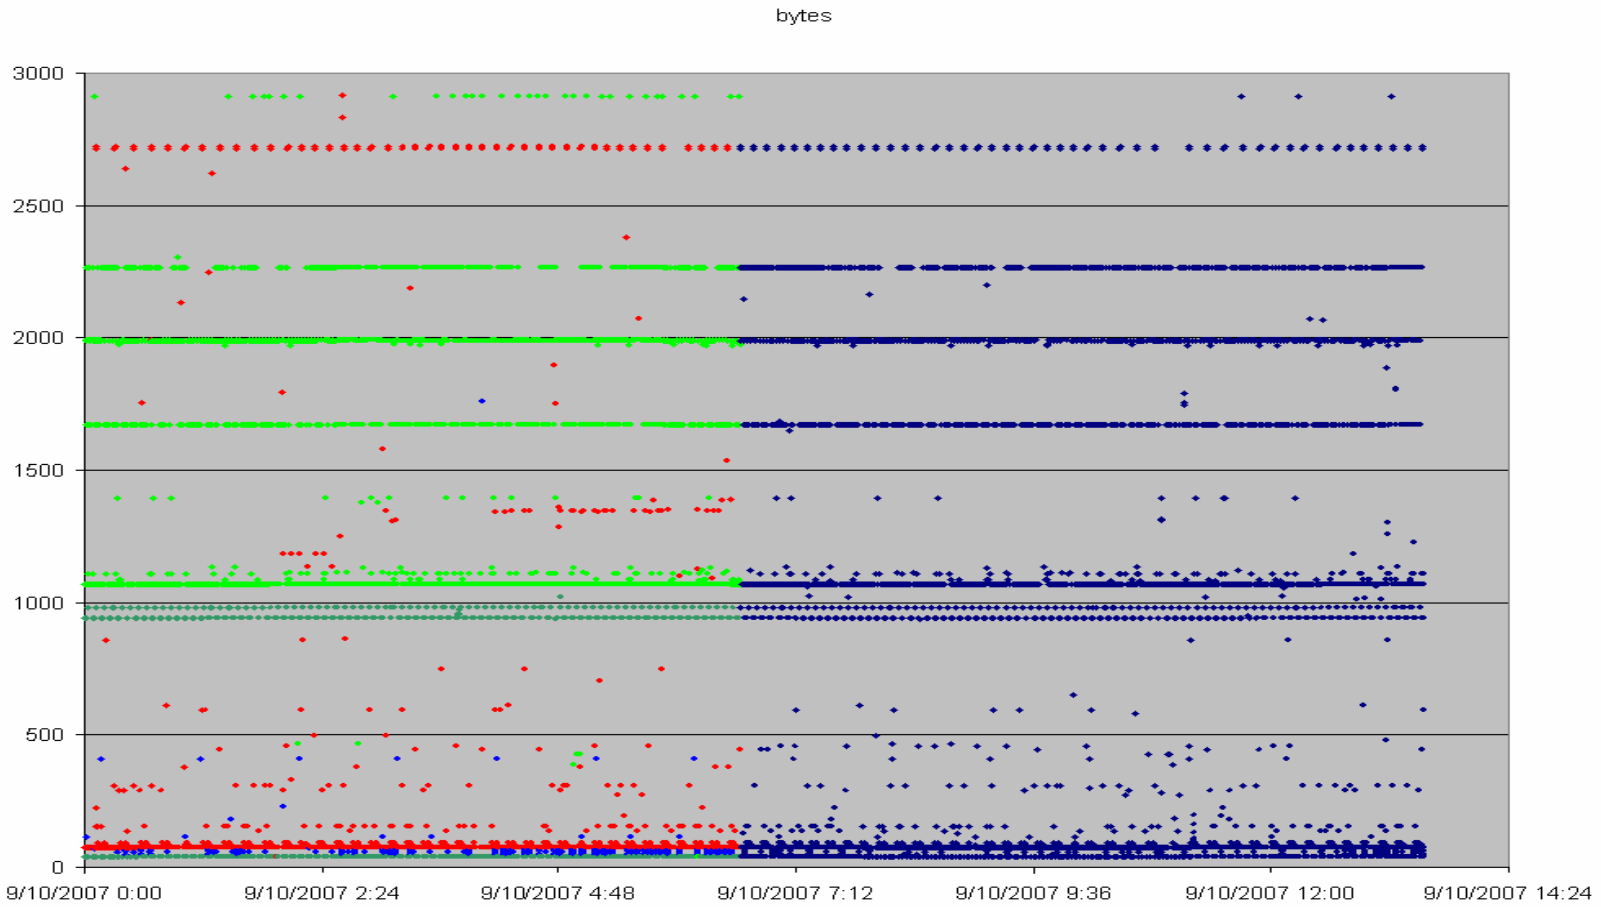


Book1

	A	B	C	D	E	F	G	I	J	K	L	M	N	O
1	sIP	dIP	sTime	bytes	eTime	sPo	dPo	type	initial					
187	xxx.xxx.219.191		9/19/2007 0:01	86	9/19/2007 0:01	4313	21	out	PA					
188	xxx.xxx.180.1		9/19/2007 0:00	108	9/19/2007 0:00	10023	45365	out	PA					
189	xxx.xxx.180.1		9/19/2007 0:00	108	9/19/2007 0:00	10023	1143	out	PA					
190	xxx.xxx.180.1		9/19/2007 0:00	108	9/19/2007 0:00	10023	63549	out	PA					
191	xxx.xxx.208.97		9/19/2007 0:00	116	9/19/2007 0:00	1352	51191	out	PA					

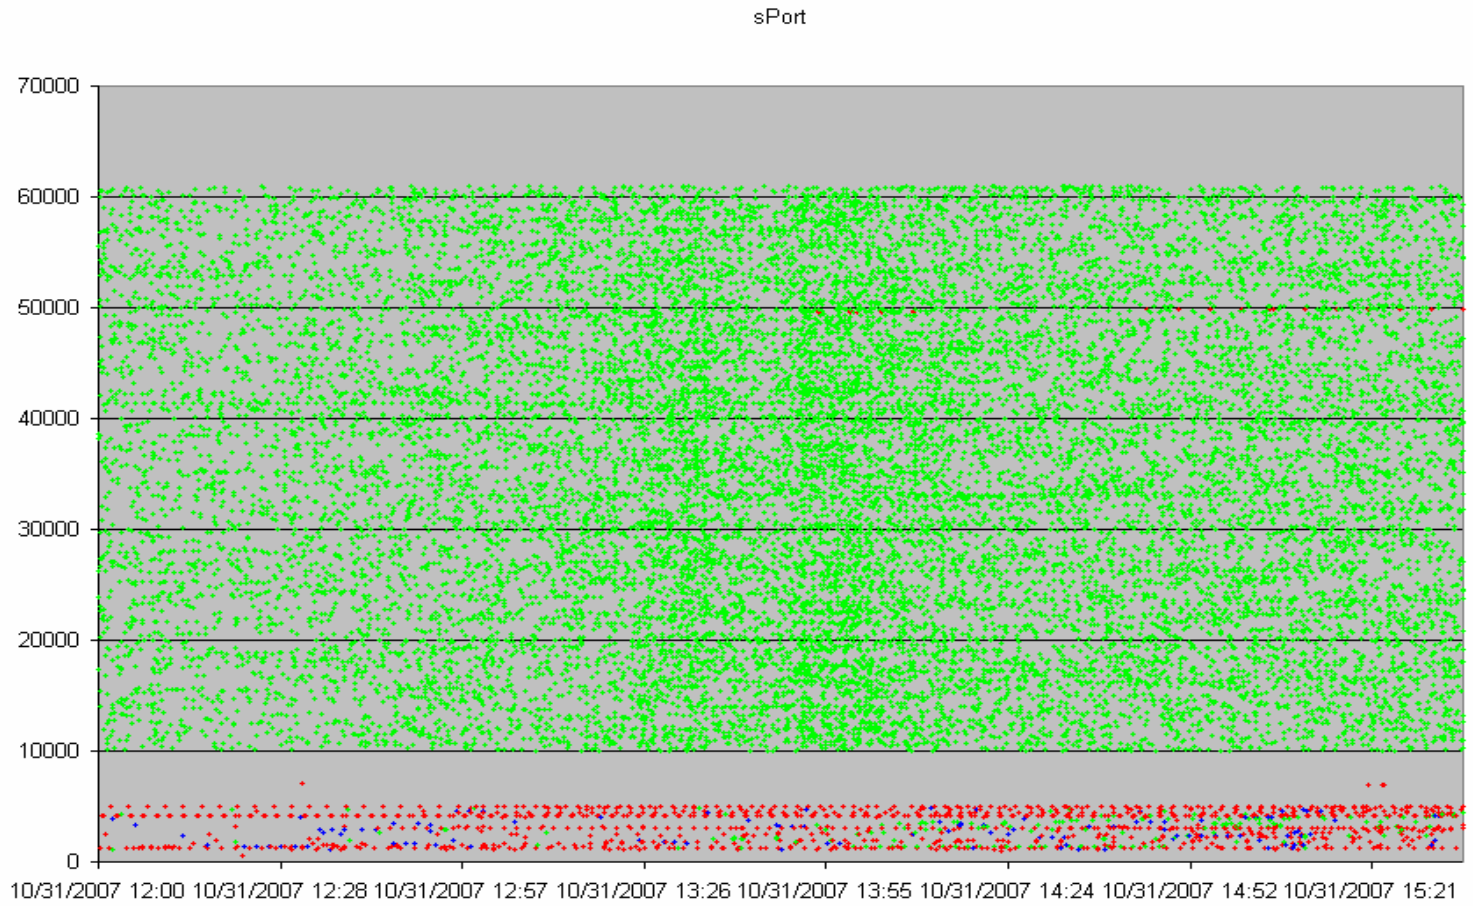
Coloring Example

Green = HTTP, Dark Green = HTTPS, Blue = DNS, Red = Other

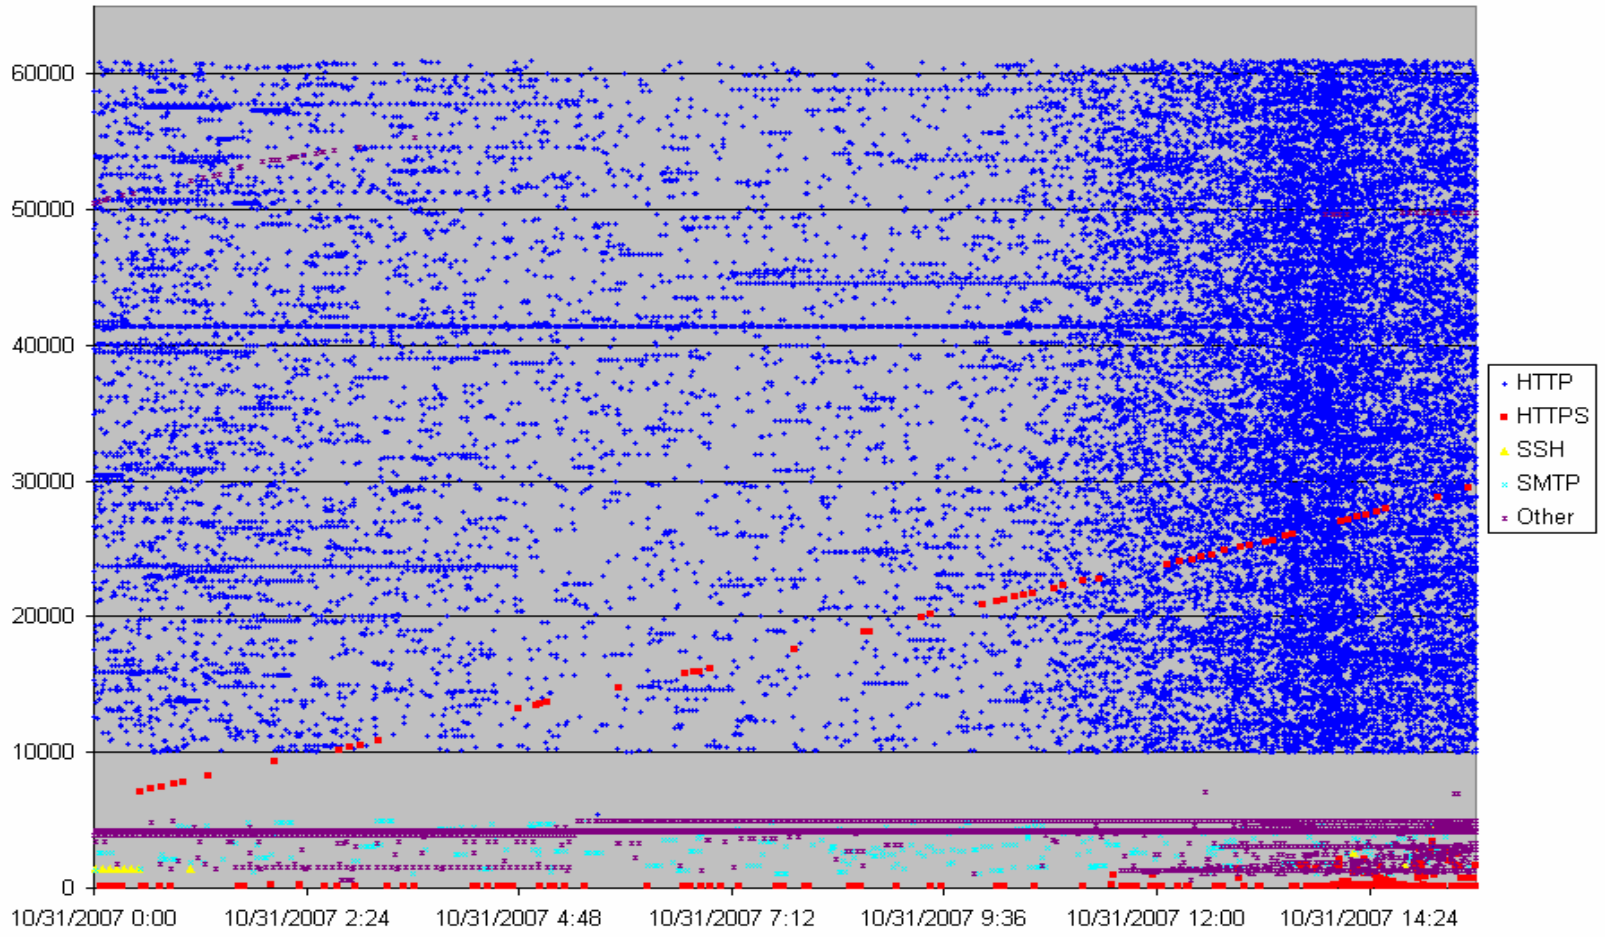


Application Coloring

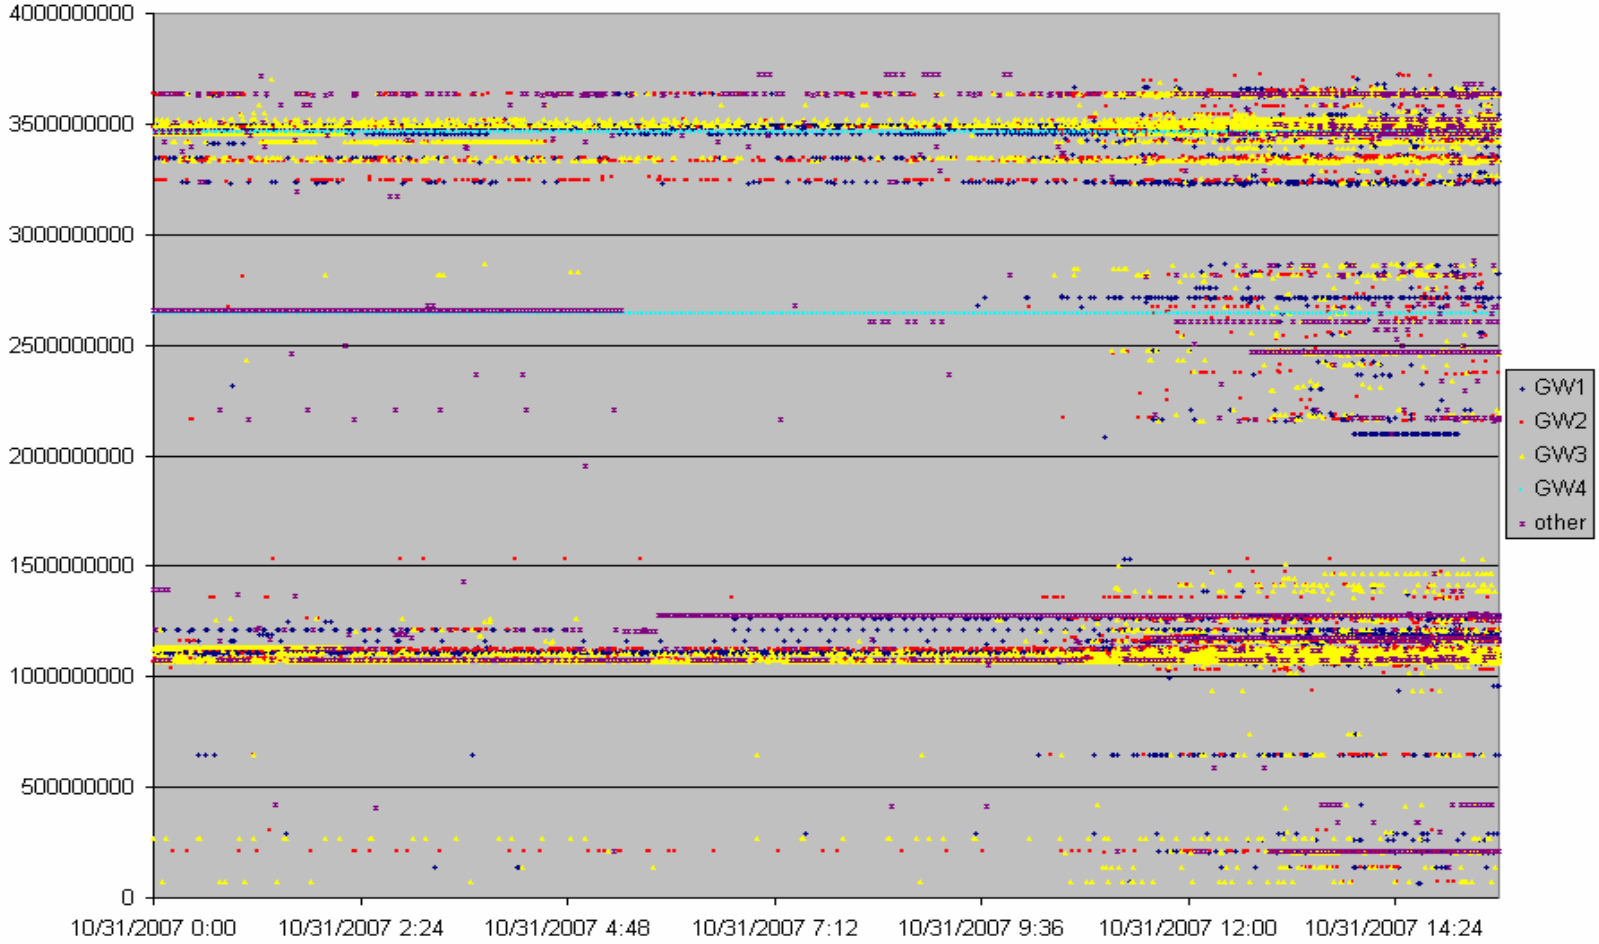
Green = HTTP, Blue = DNS, Red = Other



Color sPort vs Application



Colorization Example – GW2IP



Contact Info

- **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: +1 703-235-5111
- **For more information, visit** http://www.us-cert.gov

Questions?

Lawrence Livermore National Laboratory

Hierarchical Bloom Filters: Accelerating Flow Queries and Analysis

January 8, 2008

FloCon 2008



Chris Roblee

DOE Computer Incident Advisory Capability (CIAC)

Lawrence Livermore National Laboratory, P. O. Box 808, Livermore, CA 94551

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344

UCRL-PRES-236738

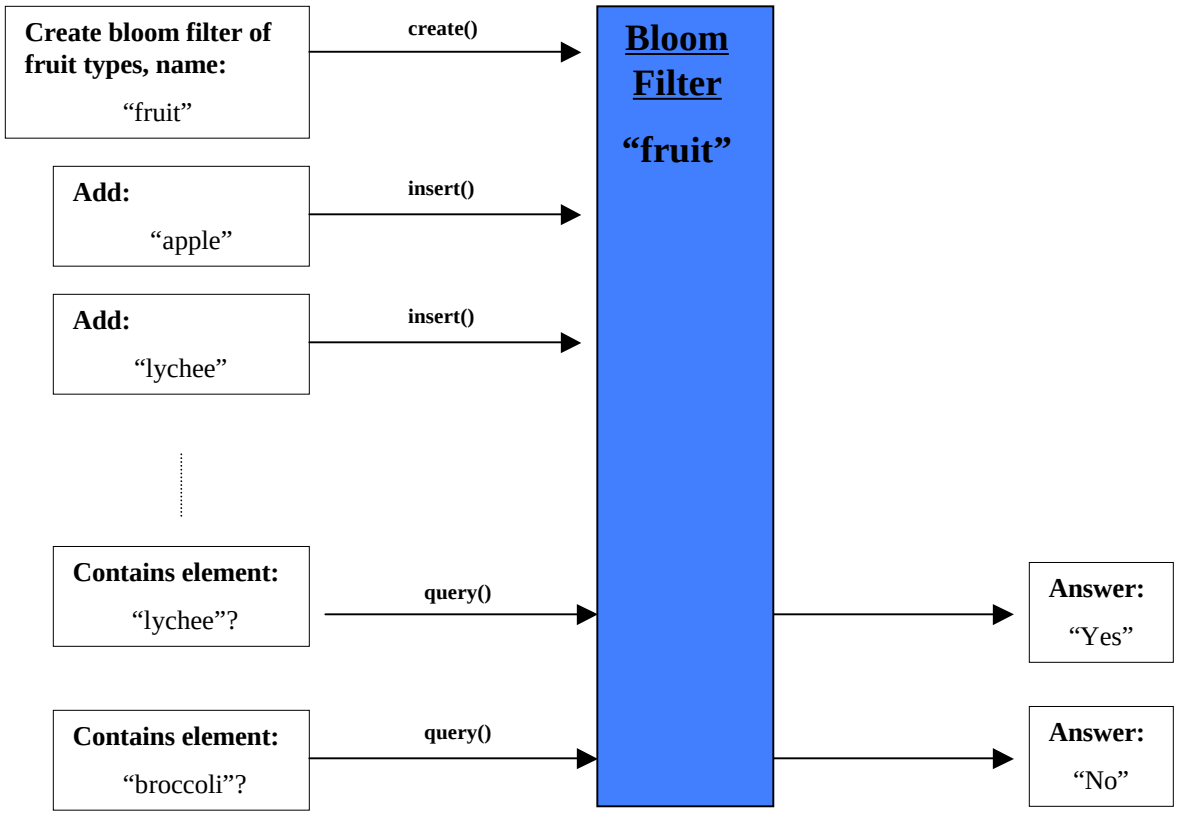
Overview

- Introduction to Bloom Filters
- Overview of CIAC's Bloom Filter-Based indexing System
- Approach's Applicability for CIAC & other CERTs
- Performance on Actual Flow Data
- Applications of Approach in Conjunction With Analytical Tools
  - Facilitating incident detection and analysis with flow visualization tools.

A Very Brief Introduction to Bloom Filters

Introduction to Bloom Filters

- High-level Functionality – trivial



http://www.eecs.harvard.edu/~michaelm/NEWWORK/postscripts/BloomFilterSurvey.pdf
http://en.wikipedia.org/wiki/Bloom_filter

Introduction to Bloom Filters

- **The Concept**
  - Efficient, probabilistic data structure, providing extremely light-weight string lookups, or “approximate membership queries”.
  - Invented by Burton Bloom in 1970 to optimize spellchecking.
  - Trade-off small probability of **false positives** for massive gains in **space and time efficiency**.
  - Popular for various large-scale network applications (e.g., web caches, query routing).

References:

http://www.eecs.harvard.edu/~michaelm/NEWWORK/postscripts/BloomFilterSurvey.pdf

http://en.wikipedia.org/wiki/Bloom_filter

How Bloom Filters Work

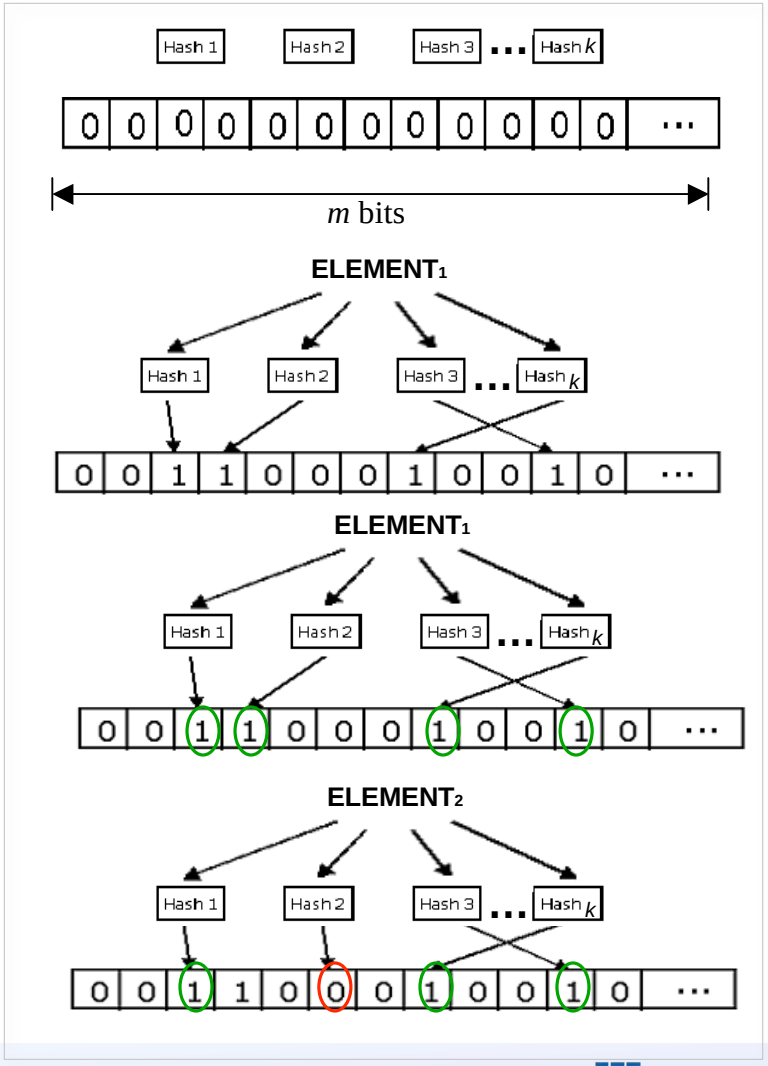
1. **Empty bloom filter is a bit array of *m* '0'- bits.**

2. **Introduce *k* different hash functions, each maps key value to one of m array positions.**

3. **Insert element by feeding it to each hash function, to obtain *k* array positions. Set these bits to '1'.**

4. **Query element (check its existence) by re-feeding into each hash function, and checking corresponding bit positions. If all bits are '1', then element is either in the filter or it's a false positive.**

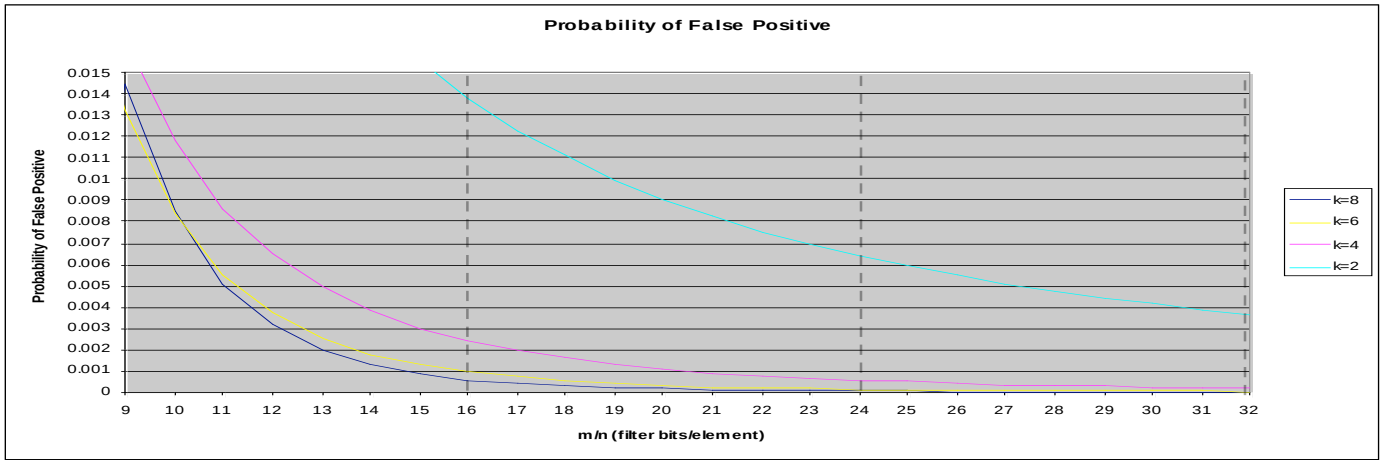
5. **If bit positions of hashes of an element contain a '0', then that element is definitely not in filter (no false negatives).**



Introduction to Bloom Filters

- False Positives
  - Probability of false positive for a <u>populated</u> bloom filter is:

$$p(\text{FP} \approx \left(1 - e^{-kn/m}\right)^k$$



- k - number of hash functions used
- n – number of elements inserted
- m – size of bloom filter (bit array)

Bloom Filters - Summary

- **Quick test of element membership:**
  - **0 likelihood of false negatives**
  - **Tunable false positive rates**
- **Probability of collisions proportional to the number of elements in set & inversely proportional to filter size.**
- **Enforce maximum false positive threshold by tuning filter size:**
  - **Often require as little as one byte per element**

Functionality

- **Significant space and time advantages over many standard, deterministic indexing structures:**
  - **Self-balancing trees**
  - **Tries**
  - **Hash-Tables**
  - **Arrays, Linked Lists**
- **Query time is O(k), independent of number of items in set.**
- **Many open source implementations available.**

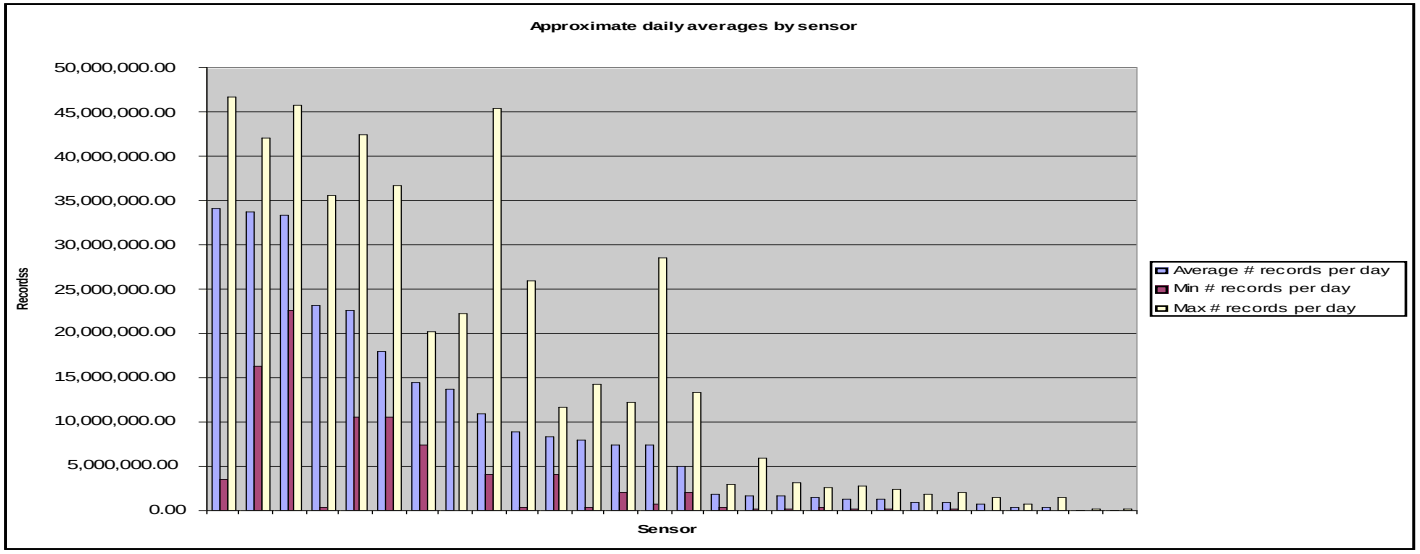
Practicality

Inexpensive, easy to deploy and maintain

Bloom Filters: Operational Viability for CIAC and the CERT Community

CIAC's Flow Collection Review

- CIAC collects massive volumes of biflow data from 29 sensors across the DOE complex:
  - 300-500 million biflows daily (~4600/s)
  - ~14GB/94GB compressed/uncompressed daily



CIAC's Flow Collection Review

- biflow feed:
  - Session summary
  - Fields:
    - Date/Time & Duration
    - Source/Destination IP and Port
    - Protocol Information
    - Bidirectional Byte and Packet Counts
    - Bidirectional Protocol Options
    - Subset of TCP/ICMP flags

Example Biflow Record

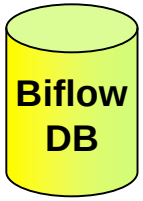
1171066191.997532,20070210000951.997532,site3,flo30,6,192168081021,192,168,81,21,IT,010000001008,10,0,1,8,US,53,1024,0,0,0.0000,0,0,54,0,1,0,0,0,0,0,0,60,0,60,0,,,14,00,+14,0,0,0,0

CIAC Analysis - Legacy Search Methodologies

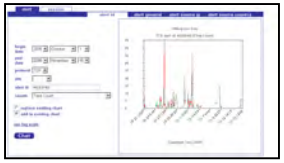
- File *grep*
  - Search sensors and hours for range of interest (e.g., “site3, site12, site21 from 10/1/06 through 12/31/06”).
  - Requires reading/decompressing and combing through GBs of data (from disk) for every day searched.



- RDBMS - Oracle
  - SQL+
  - Perl/JDBC
  - Typically limited* to past ~25 days of bi-directional sessions (~15%)

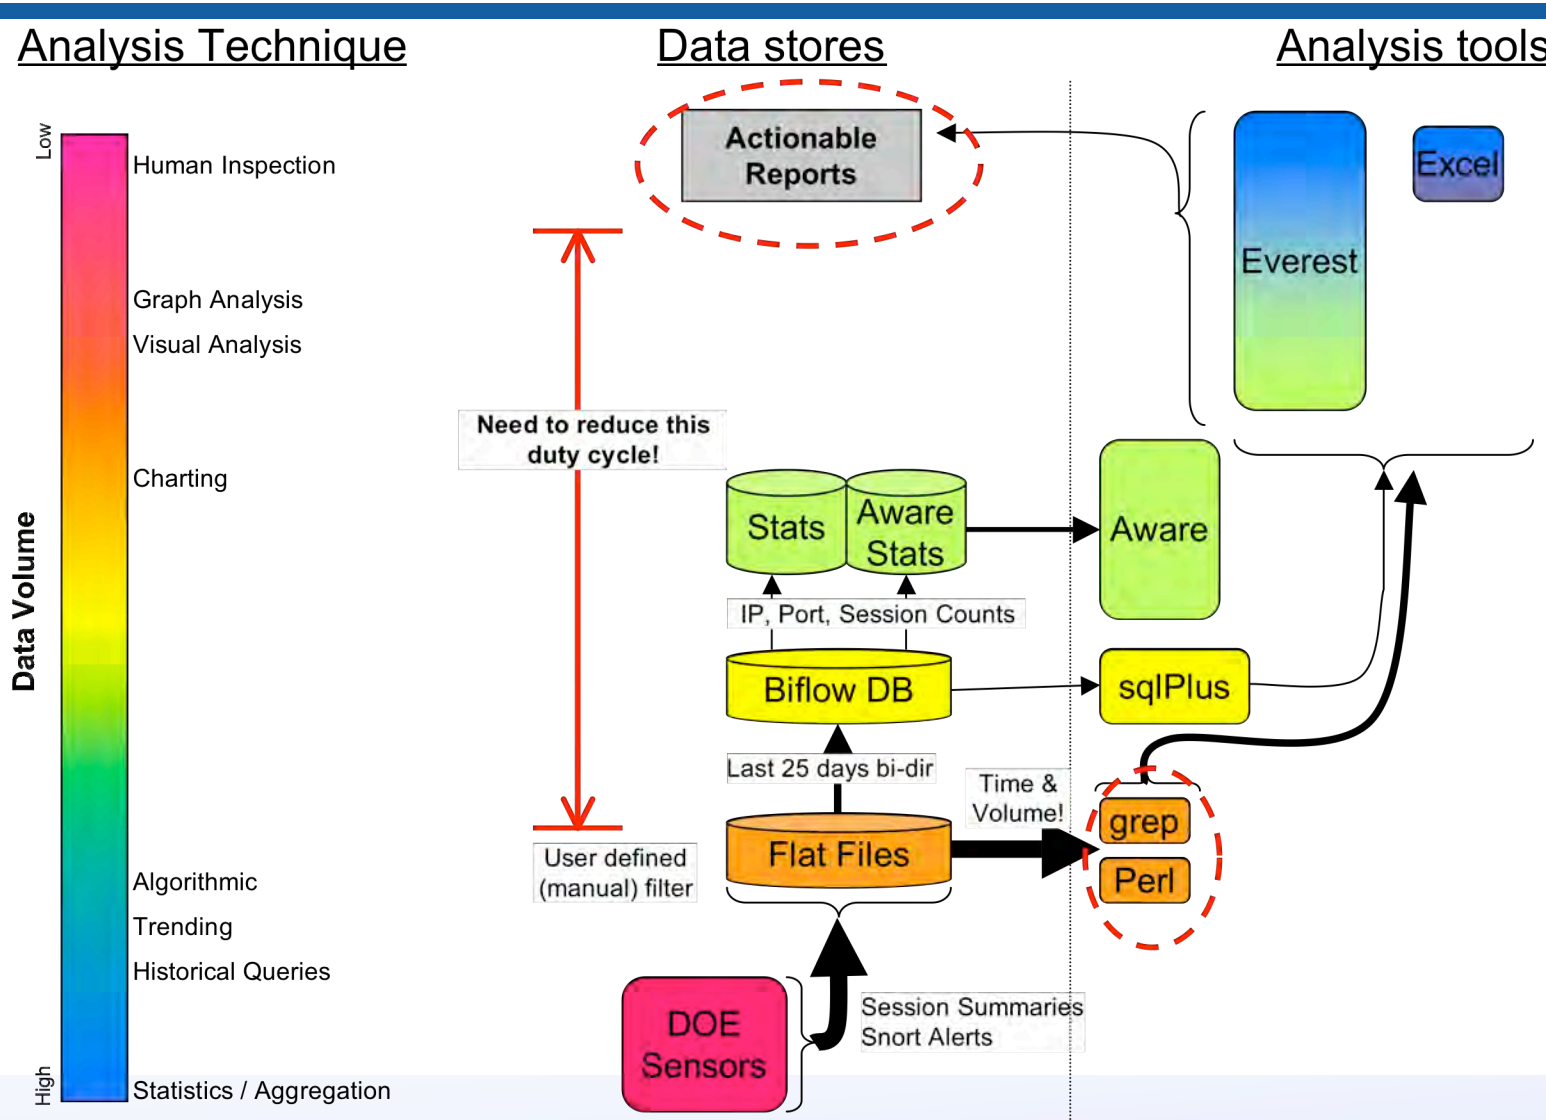


- AWARE web portal
  - High-level charting and statistics (session counts, etc.)



Many mission-critical searches can take several hours or days to complete

Current CIAC Analysis Data Flow



Watch and Warn Query Needs and Issues

- Rapidly search all flow data over long periods of time:
  - Analysts typically search on IP address:
    - Watch list (suspicious, known-bad, etc.)
    - Nodes of interest
    - Compromised internal nodes
  - Various time (hours, days, months) and space (single site, all sites) scales.
  - Require quick turnaround (minutes) to respond to site requests:
    - e.g. “Have you seen these IPs at my site in the past 3 weeks?”

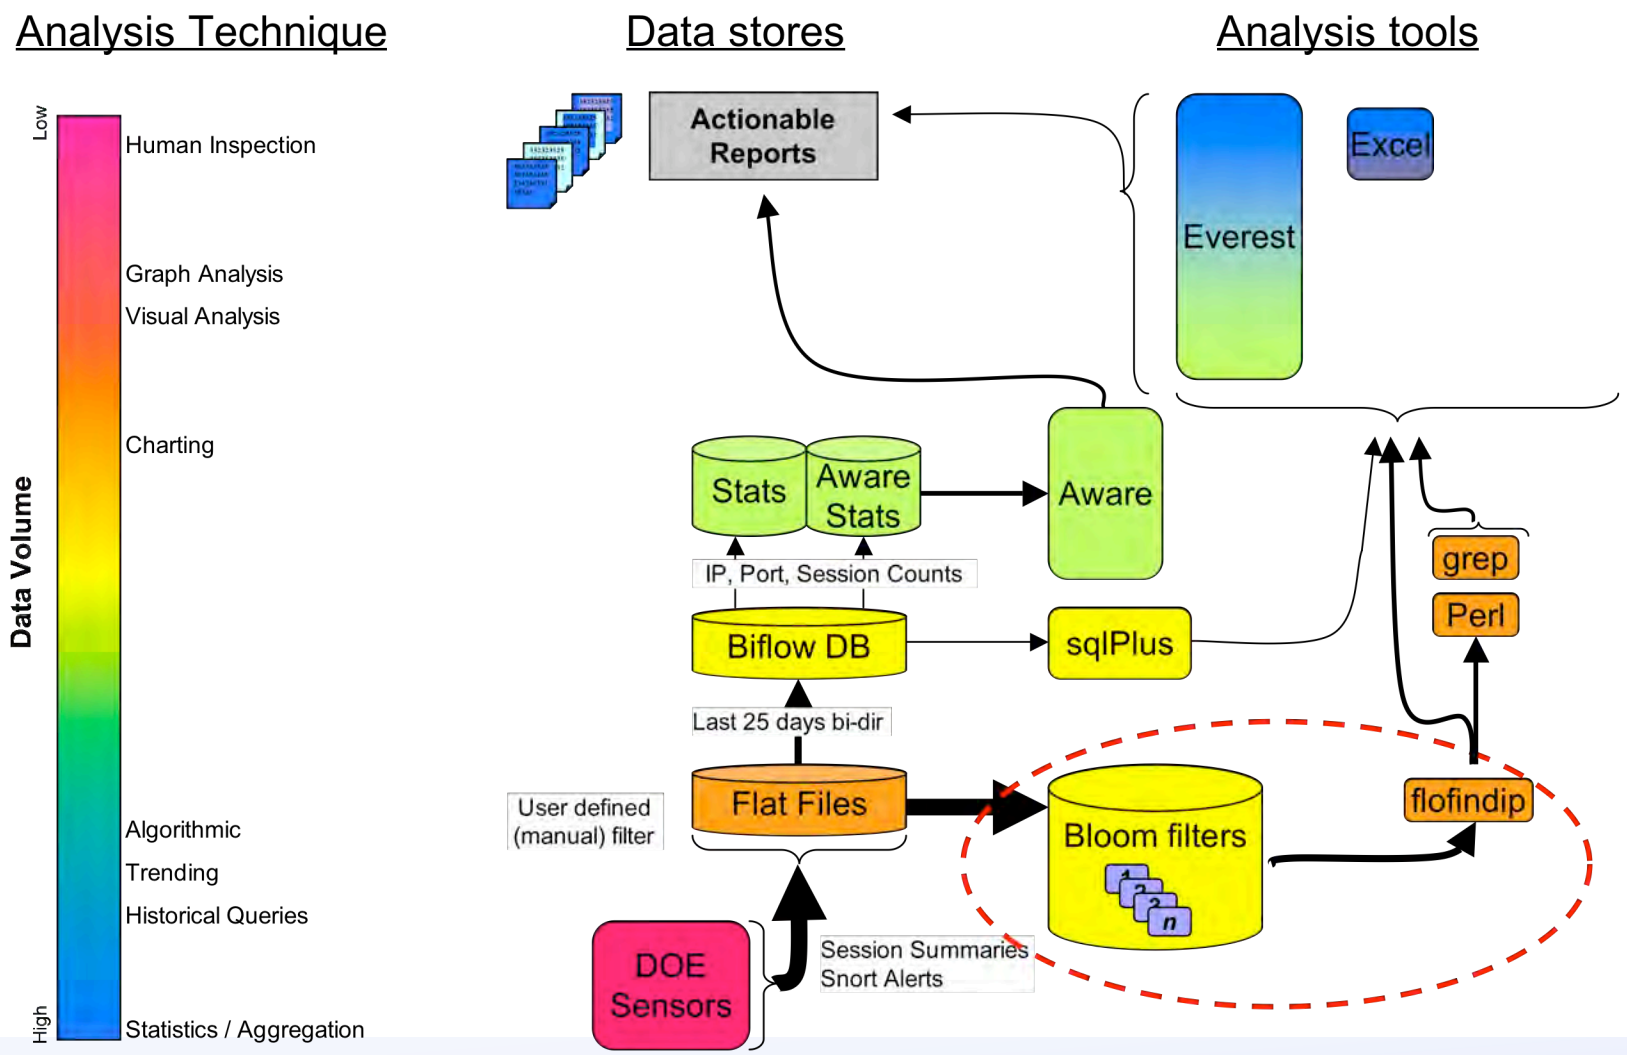
- IP-based searches often yield relatively small result sets:
  - “Interesting” IP might only have been seen in 30 site-hours, whereas 21,600 hours (~1 DOE-month) might have been searched.
    - ➔ 99.9% wasted duty cycle!
  - Need to reduce the search space (raw flow files) through better cataloging of data as it arrives.

Bloomdex:
CIAC's Bloom Filter-based Indexing System for Network Flow Analysis

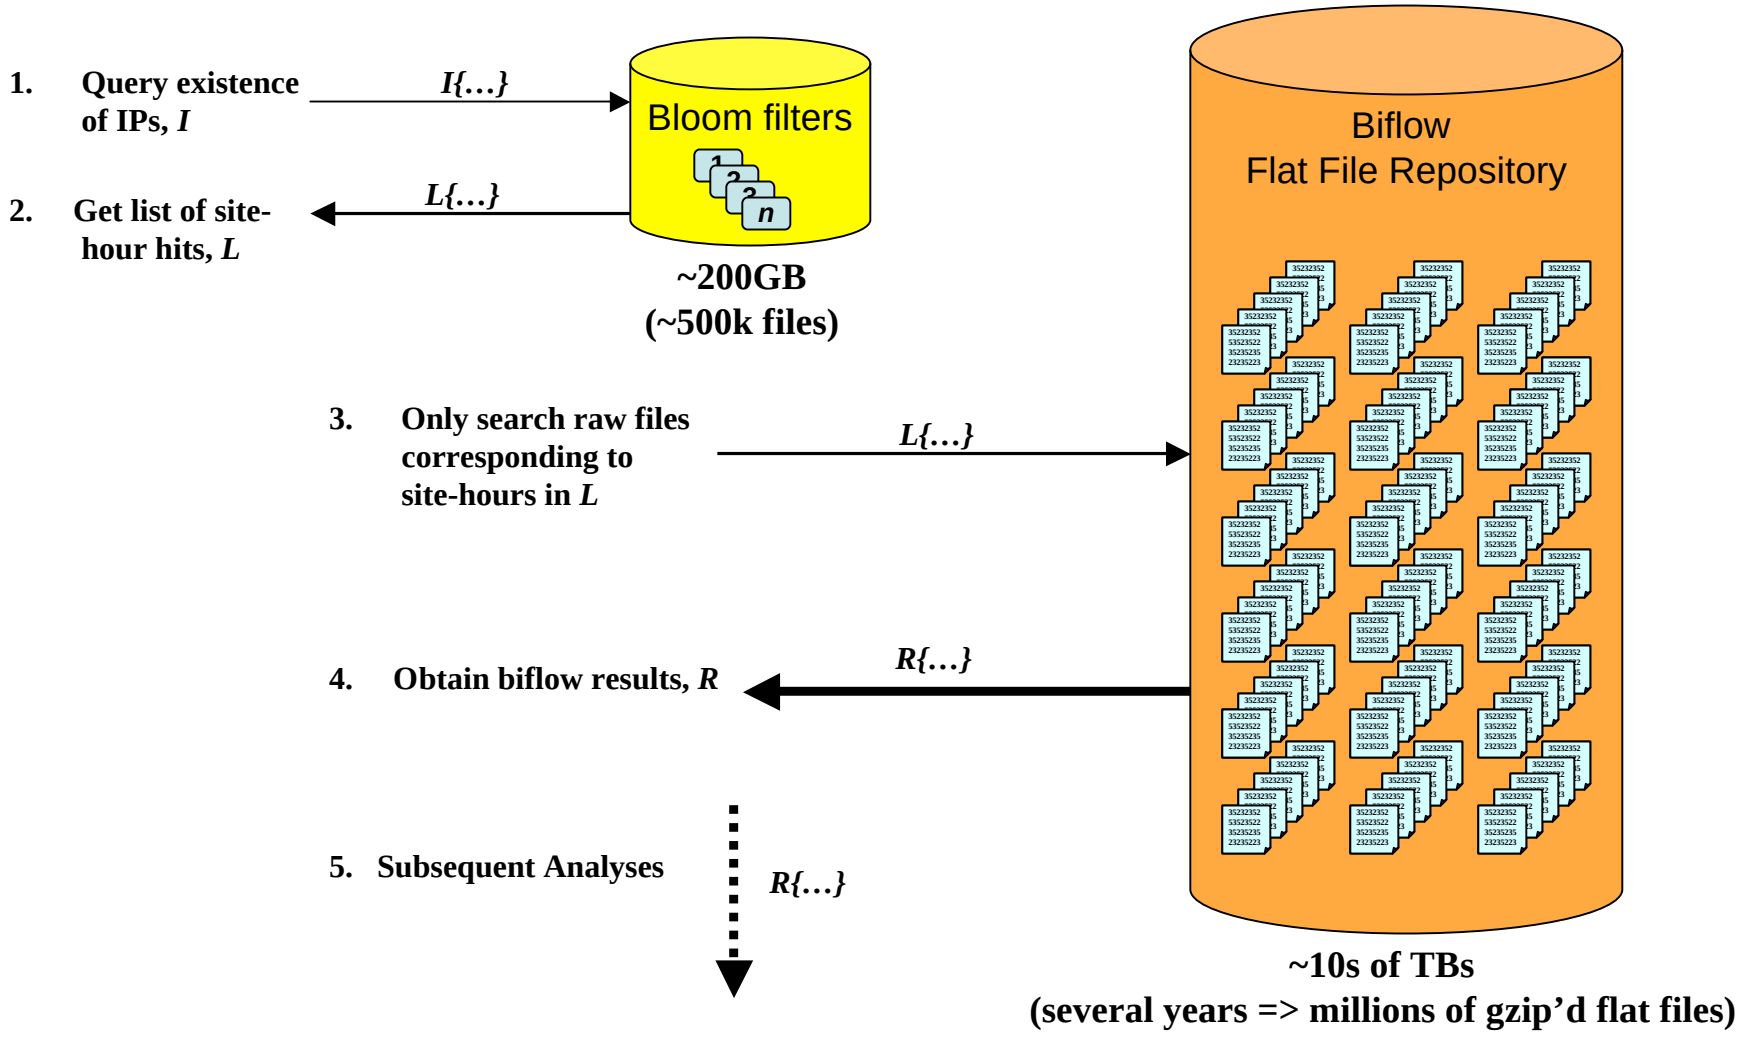
Solution: Bloomdex

- ***Bloomdex***
  - A hybrid hierarchy/file-based Bloom filter system to index CIAC's biflow records.
  - Currently indexed by source or destination IP.
  - Index partitioned by:
    - Site-month (e.g., "SITE8 11/2006")
    - Site-day (e.g., "SITE8 11/5/2006")
    - Site-hour (e.g., "SITE8 11/5/2006 13:00")
  - Uses intuitive directory tree structures and multi-scale bloom filters to accelerate IP-based searches.
  - max(FP rate) ≈ 2×10^{-4} → **3 bytes of storage per unique IP**

Blooomdex - CIAC Analysis Data Flow



Reducing the Biflow Search Space



Bloomdex:
Performance Profile

Bloomdex: Comparative Performance Profiles

Typical analyst IP-based queries:

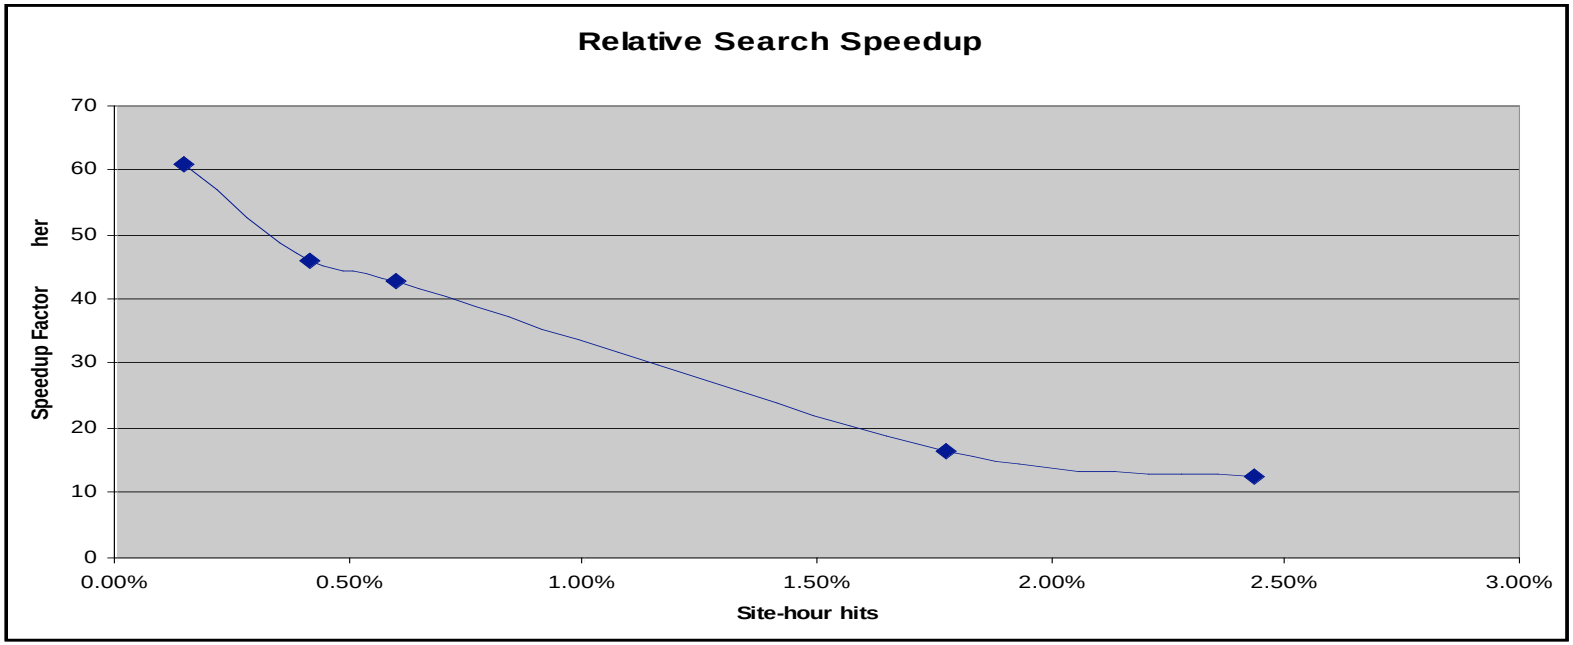
IPs searched	Date-range searched	Site-hours searched	Site-hour hits	% Site-hour hits	Session hits	Raw biflow file hits	Search time (conventional)	Search time (bloomdex)	Relative Speedup
8	12/13/06 - 1/9/07	19,140	466	2.43%	10,594	600	**16.29 hours**	**1.3 hours**	**12.5**
13	10/15/06 - 1/17/07	65,888	1,166	1.77%	158,345	1,667	**57.52 hours**	**3.45 hours**	**16.7**
13	1/22/07 - 1/29/07	4,959	31	0.60%	78	39	**4.16 hours**	**5.82 minutes**	**42.9**
4	1/1/07 - 1/2/07	725	3	0.41%	3	3	**21.5 minutes**	**28 seconds**	**46.1**
9	1/23/07 - 1/24/07	725	1	0.14%	1	1	**41.7 minutes**	**41 seconds**	**61**



- Expect >10x speedup
- Strong dependency on site-hour hit ratio
- Future optimizations to search tools could make it even faster

Bloomdex: Performance Profile

- Comparative Performance:



  - **Strong relationship between speedup and site-hour hit ratio**
  - **Ideal for searches on sparsely-occurring IPs**

Bloomdex: Performance Profile

- Bloom filter generation performance:
  - **Average site-day filter generation rate:**
    - ~ 33/hour = 792/day (current incoming rate: 29/day)
  - **Average site-hour filter generation rate:**
    - ~ 390/hour = 9360/day (current incoming rate: 696/day)

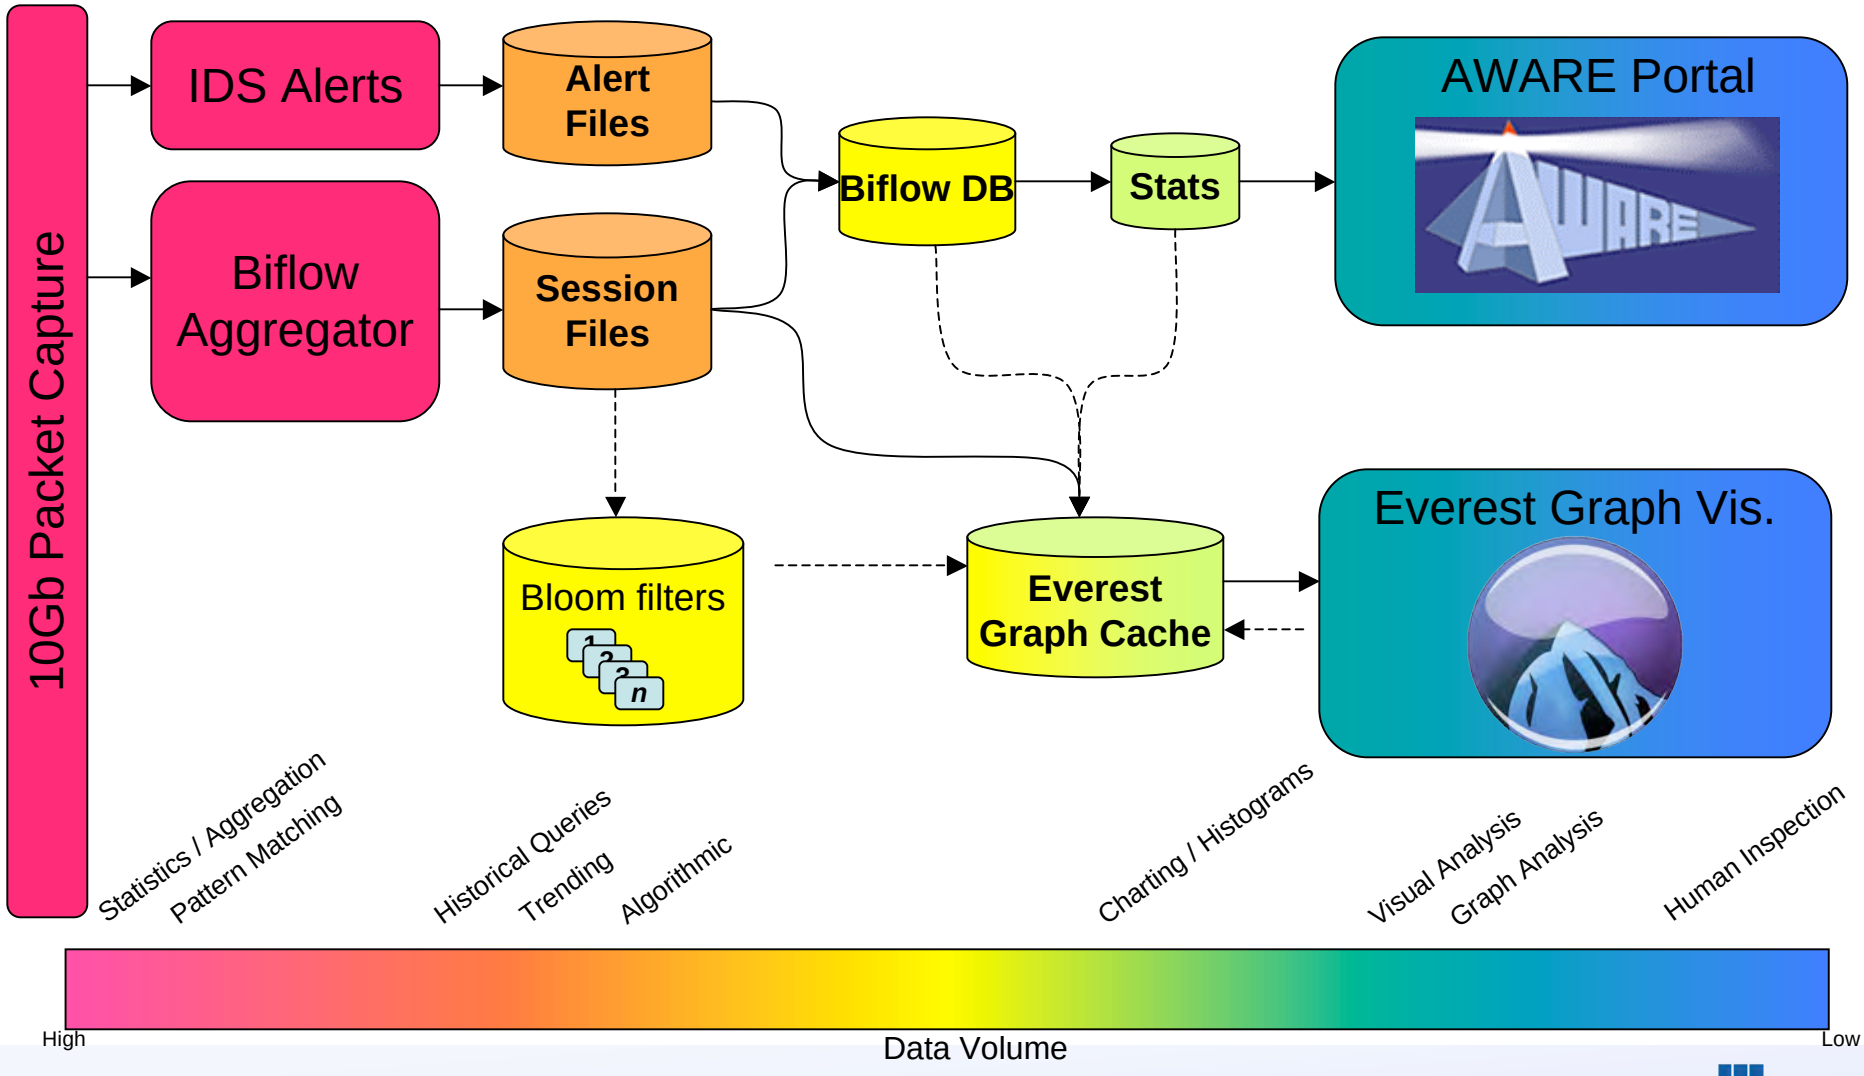
<u>Will scale well to 100+ sites (cheaply)</u>

Bloomdex: Status

- **Coverage**
  - 2.5 years of biflow records indexed.
- **Storage footprint**
  - 3 bytes per unique IP at the site-hour, site-day and site-month levels.
  - Bloom filters currently using ~200GB of shared storage.
- **Exploring additional space and performance-based optimizations**
  - Other dimensions (e.g., port, ip-port, srcip-dstip pairs)
  - Counting Bloom filters
  - Different hashing functions
  - Parallelization

Bloomdex:
Analyst Workflow Integration

Analyst Workflow Integration



Facilitating Incident Analysis with Bloomdex and Everest Flow Visualization

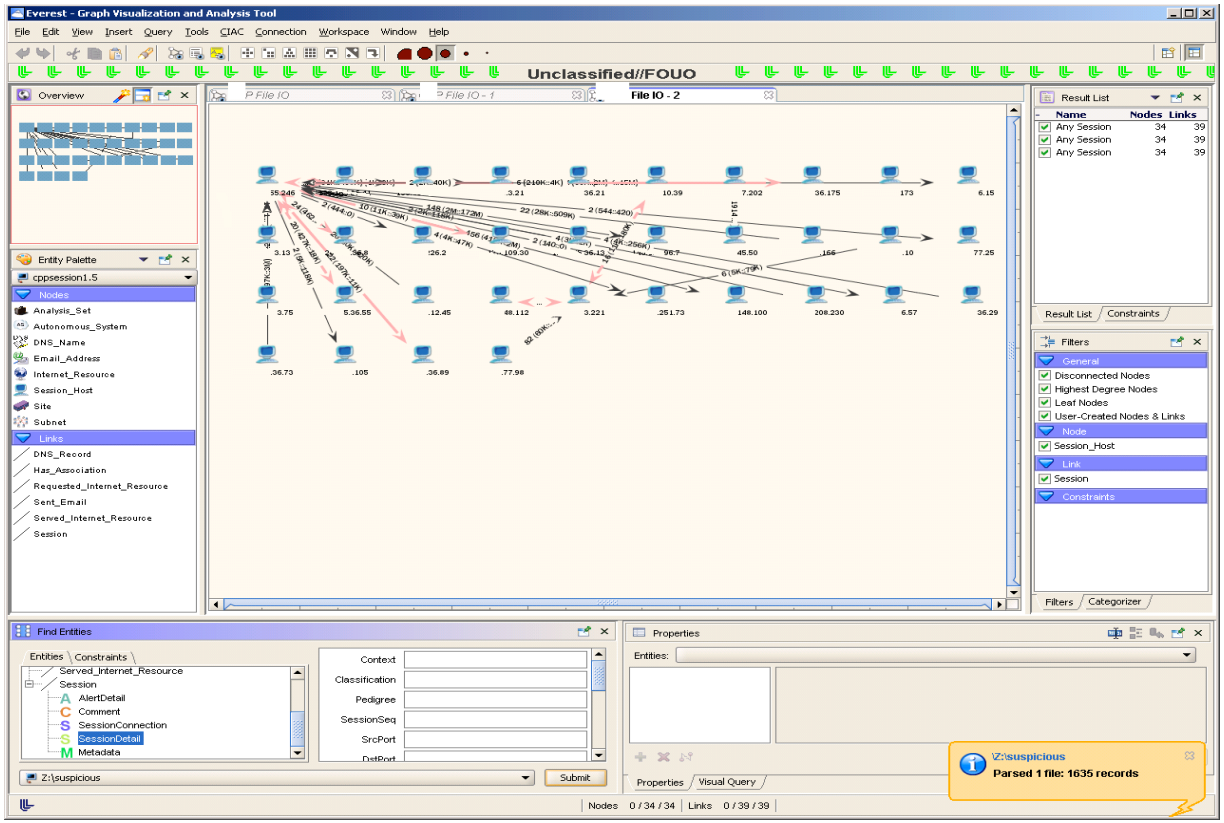
- **Example Use Scenario:**

1. Site reports compromise
   - Supplies 4 suspect IPs to CIAC.

2. CIAC queries biflow data for suspect IPs using *Bloomdex* query tool:
   - Search all sensors over a sufficient time range (perhaps a full year).
   - Quickly identify several other sites with hosts exhibiting similar behaviors.
   - Analysis set narrowed down to just 1,635 sessions.

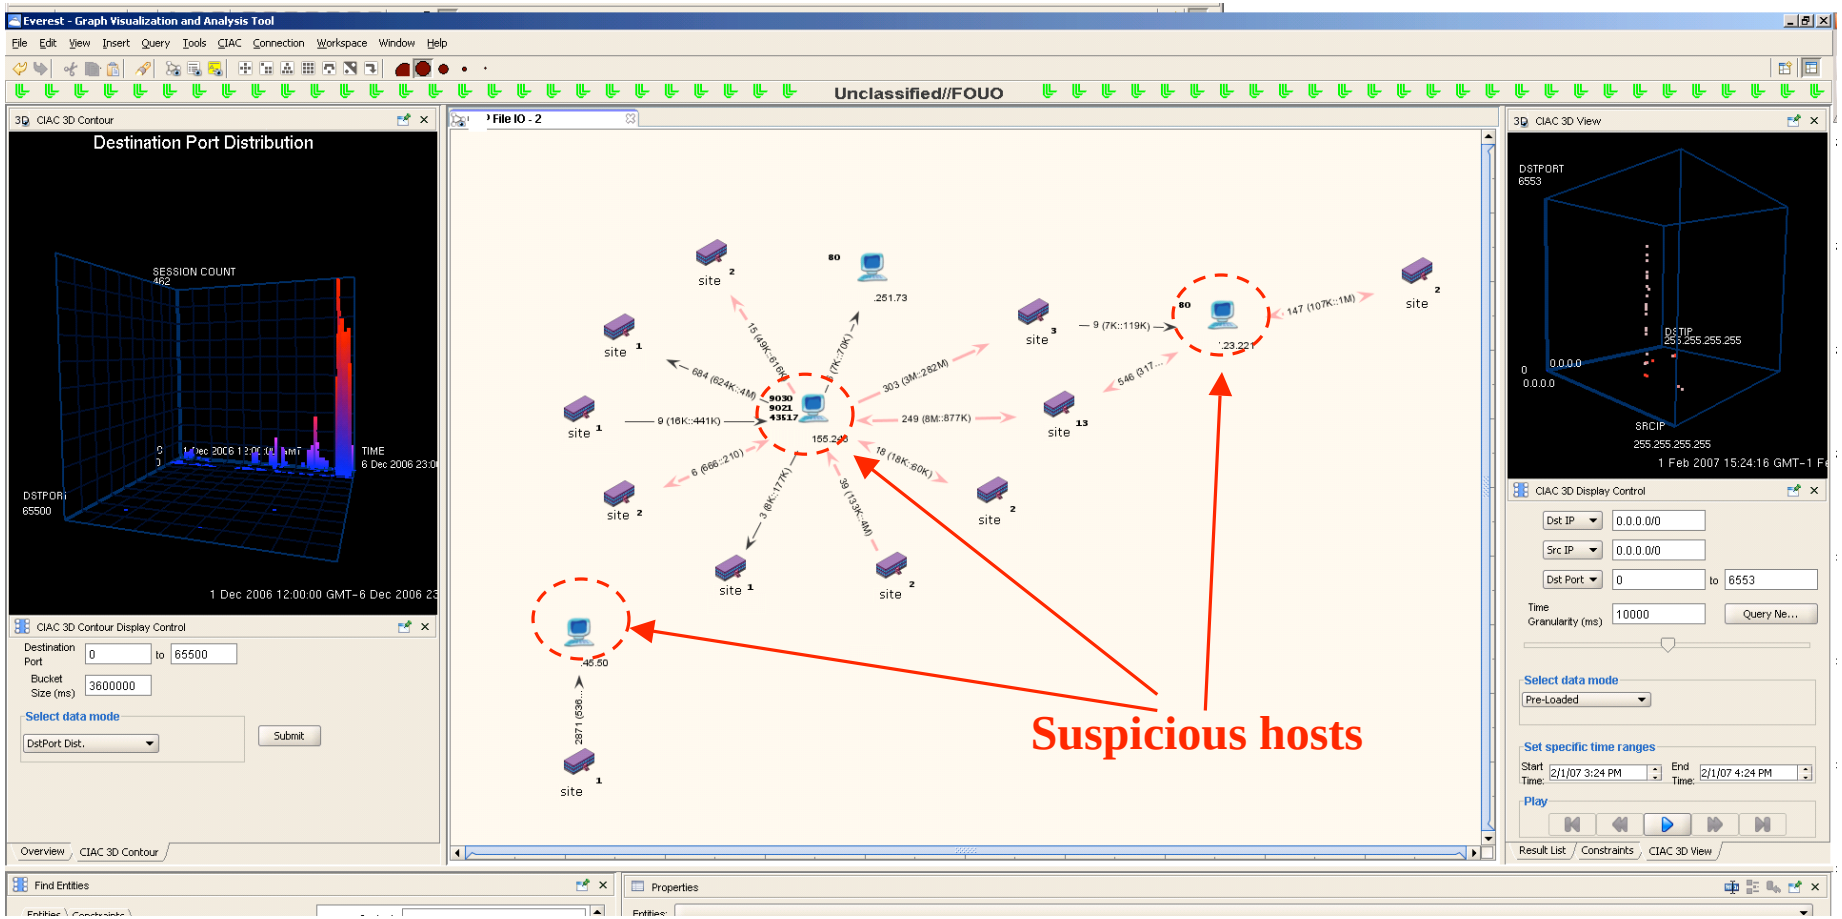
Analysis Using Bloomdex and Everest (2)

3. Launch Everest graph visualization tool, point to *Bloomdex* output file containing result set (1,635 biflow records).
4. Issue general query to generate session graph:



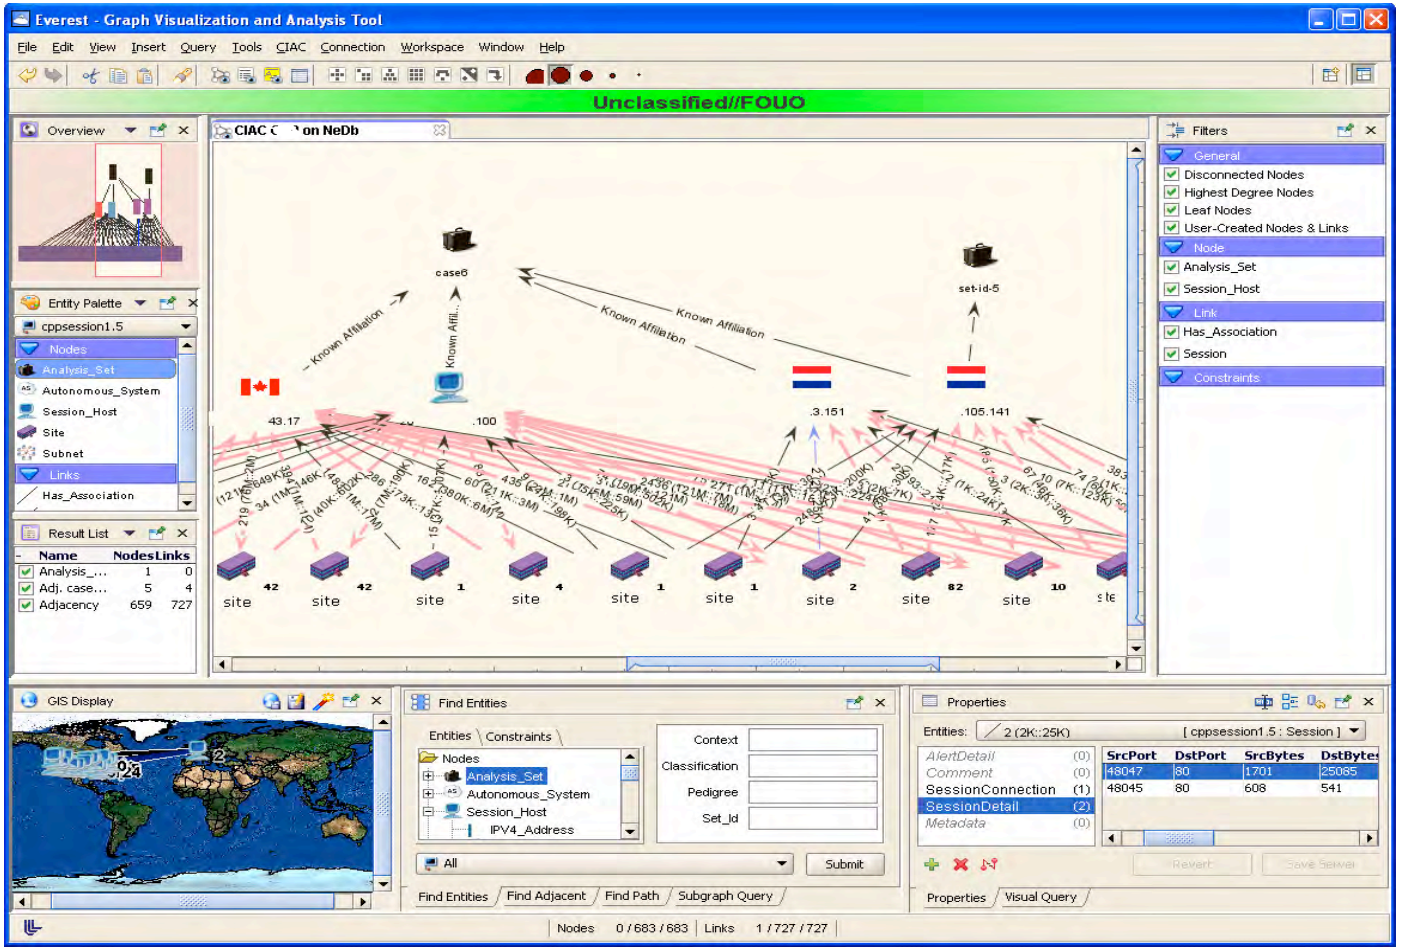
Analysis Using Bloomdex and Everest (3)

5. Perform drill-down or aggregate analysis



Analysis Using Bloomdex and Everest (4)

6. Perform in-depth or summary analysis



Conclusion

- **The *Bloomdex* suite enables significantly faster turnaround times on analyst IP-based queries:**
  - It does this by drastically narrowing the search space through Bloom filter pre-queries.
  - Facilitates use of other analytic tools, such as Everest.
  - Provides significant space savings.
  - **Very straightforward and inexpensive to deploy and maintain.**

- **Future**:
  - Utilize compressed bitmap indexes as an integrated indexing/retrieval solution.

Questions

cdr [at] llnl.gov

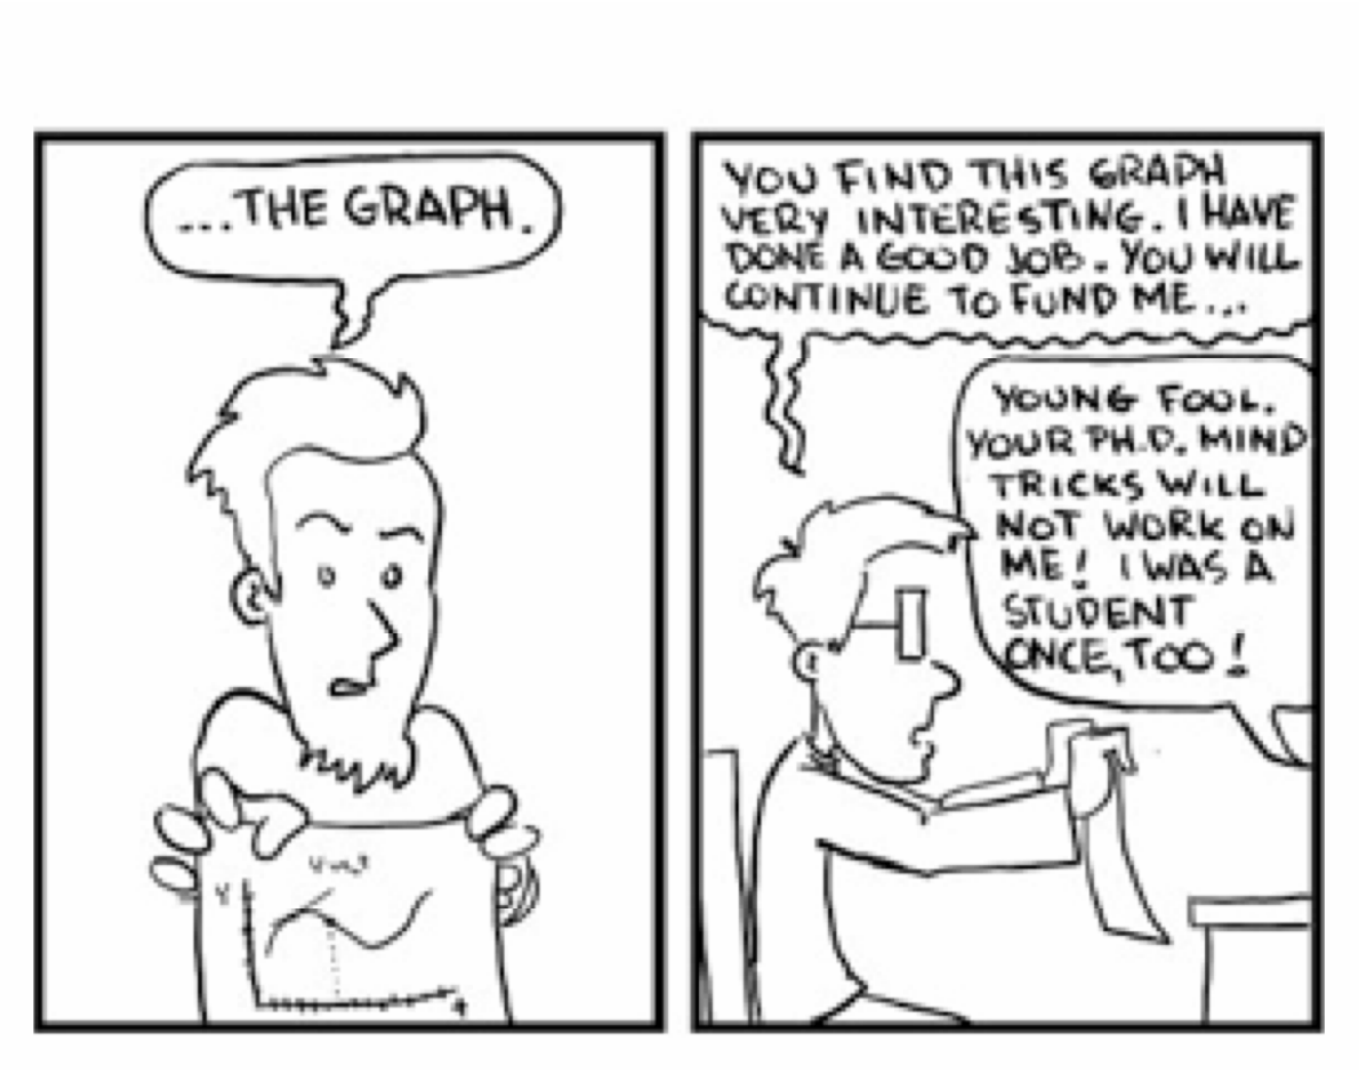


Visualizations of Flow and Analytical Results

Presentation by: Phil Groce and Jeff Janies
Network Situational Awareness group SEI/CERT

© 2007 Carnegie Mellon University

Visualizations are Tools



Visualizations are Tools



Visualizations are Tools



>



Visualizations are Tools



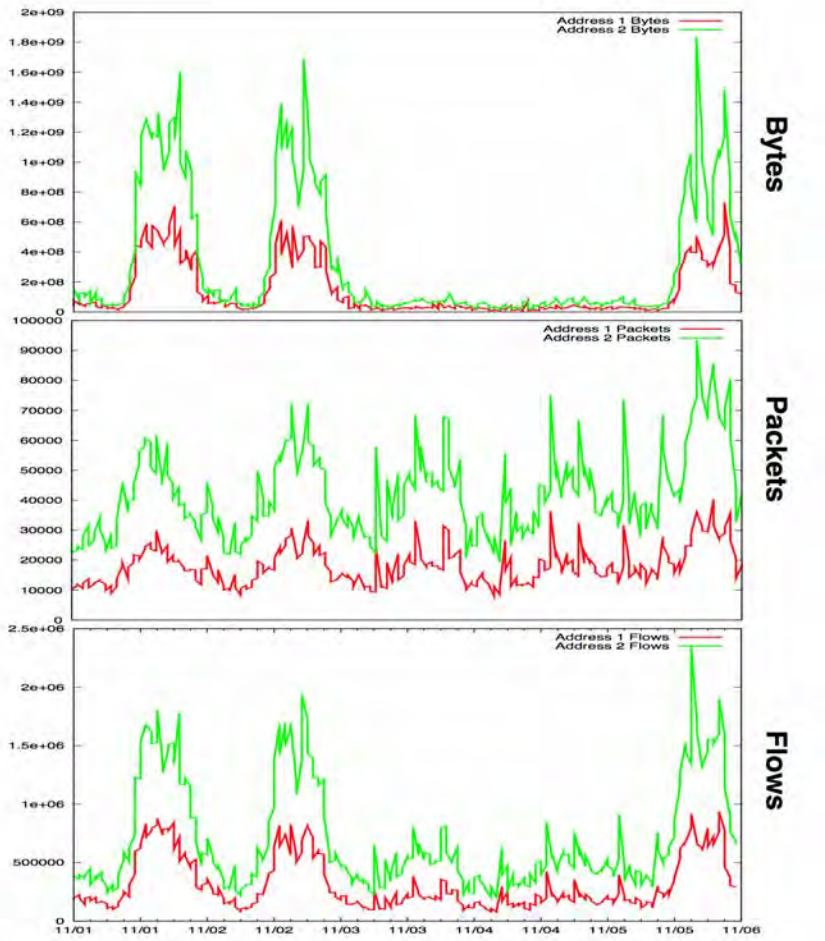




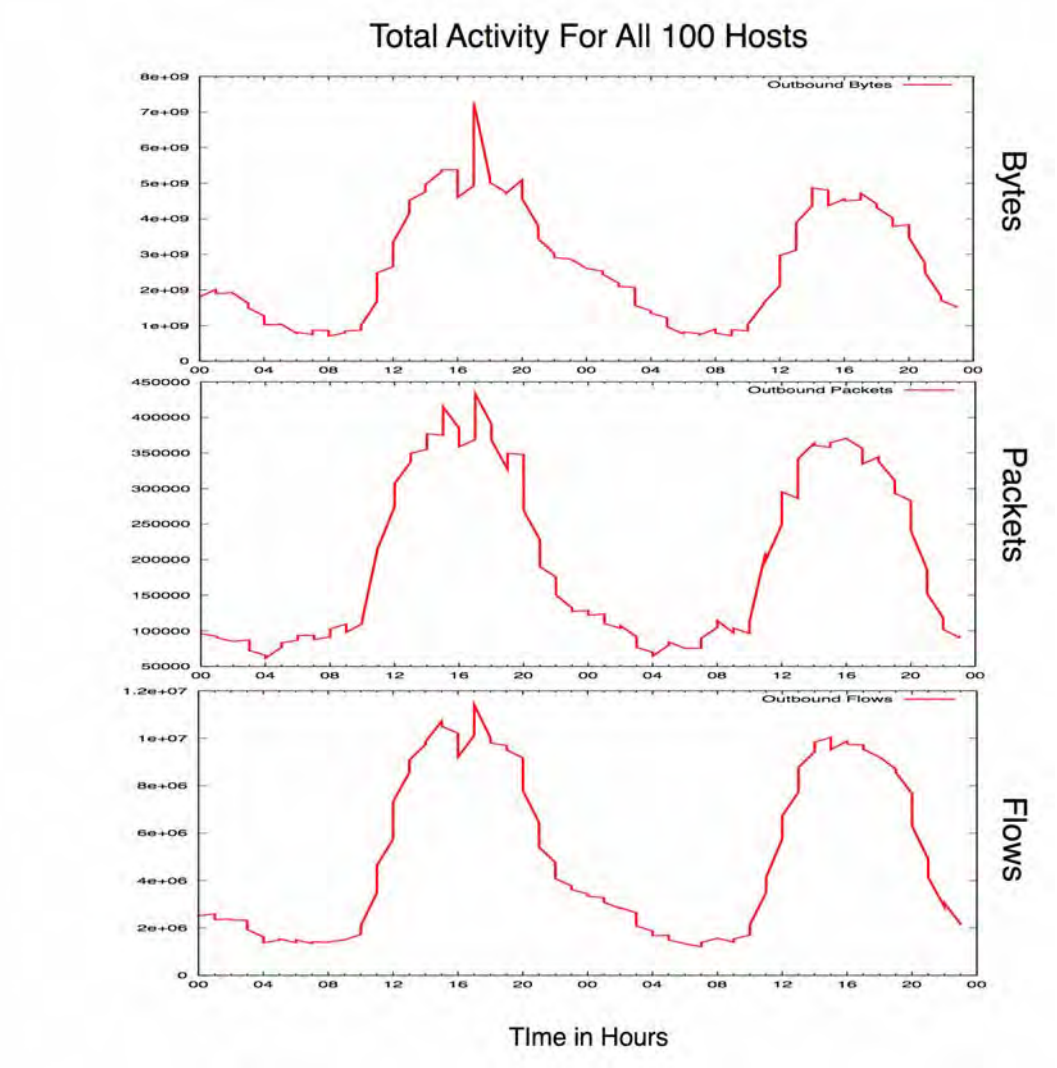
Time Series:

The Tried and True Hammer

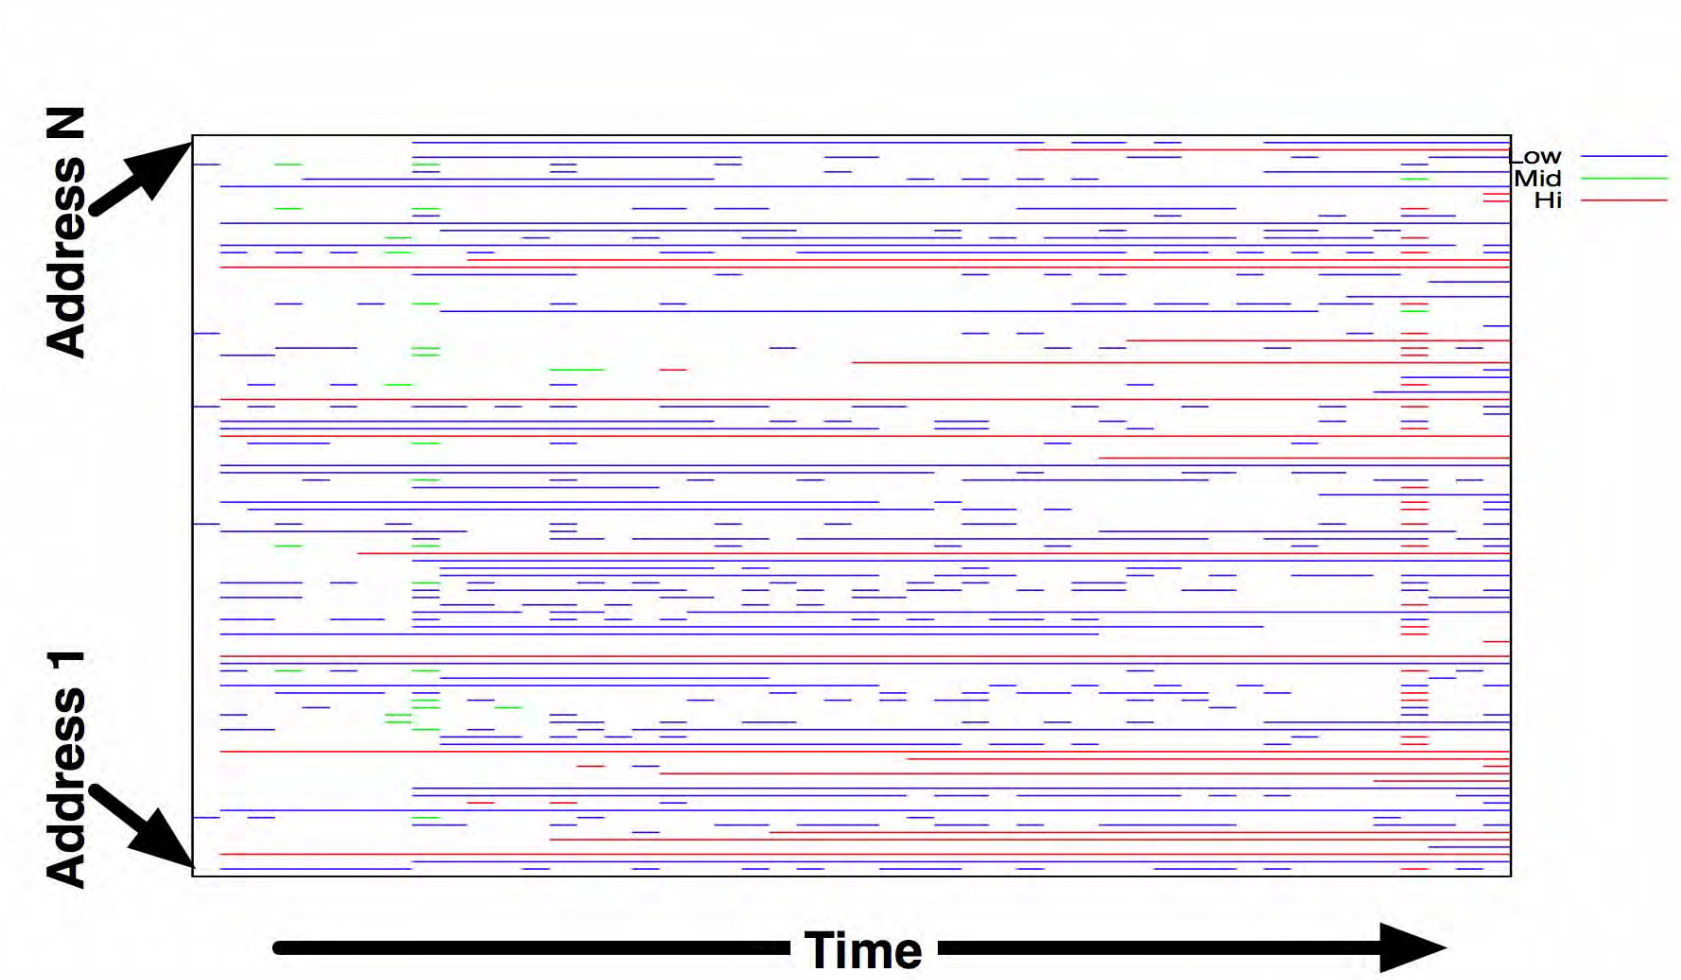
Time Series


Address 1 Bytes
Address 2 Bytes
Bytes
Address 1 Packets
Address 2 Packets
Packets
Address 1 Flows
Address 2 Flows
Flows
11/01
11/01
11/02
11/02
11/03
11/03
11/04
11/04
11/05
11/05
11/06

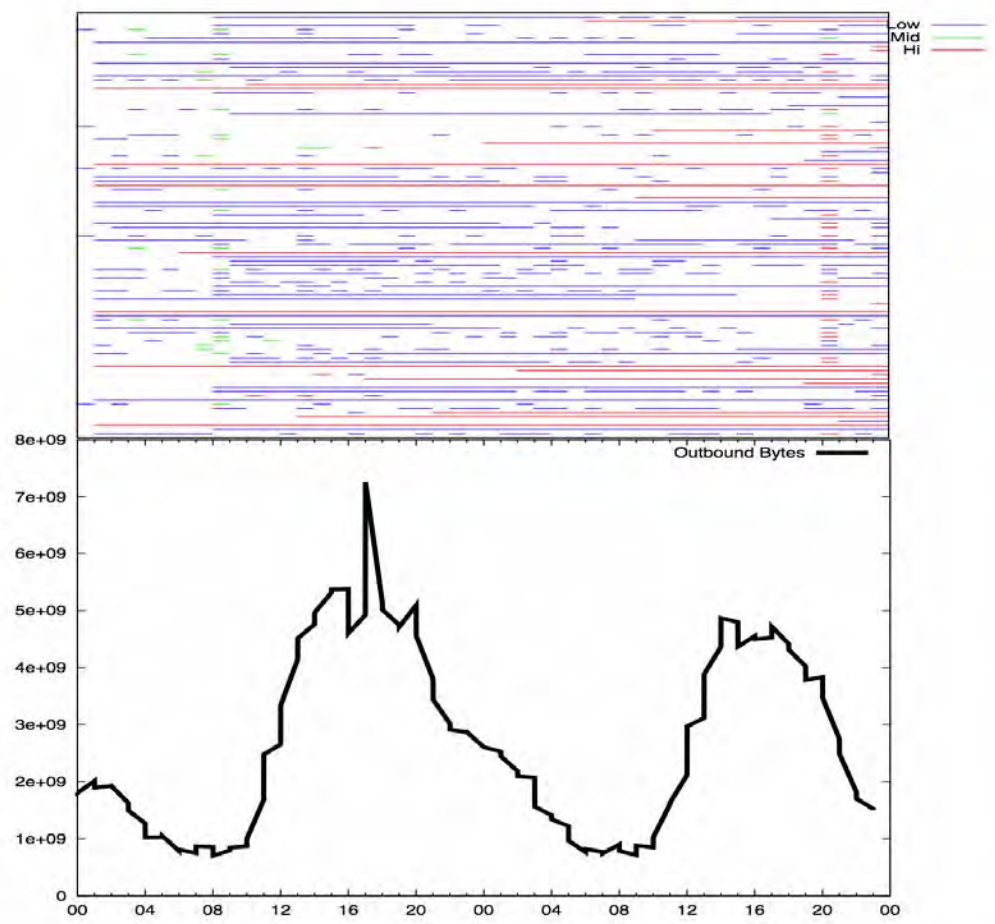
Time Series


Total Activity For All 100 Hosts
Outbound Bytes
Bytes
Outbound Packets
Packets
Outbound Flows
Flows
TIme in Hours

Existence Plots



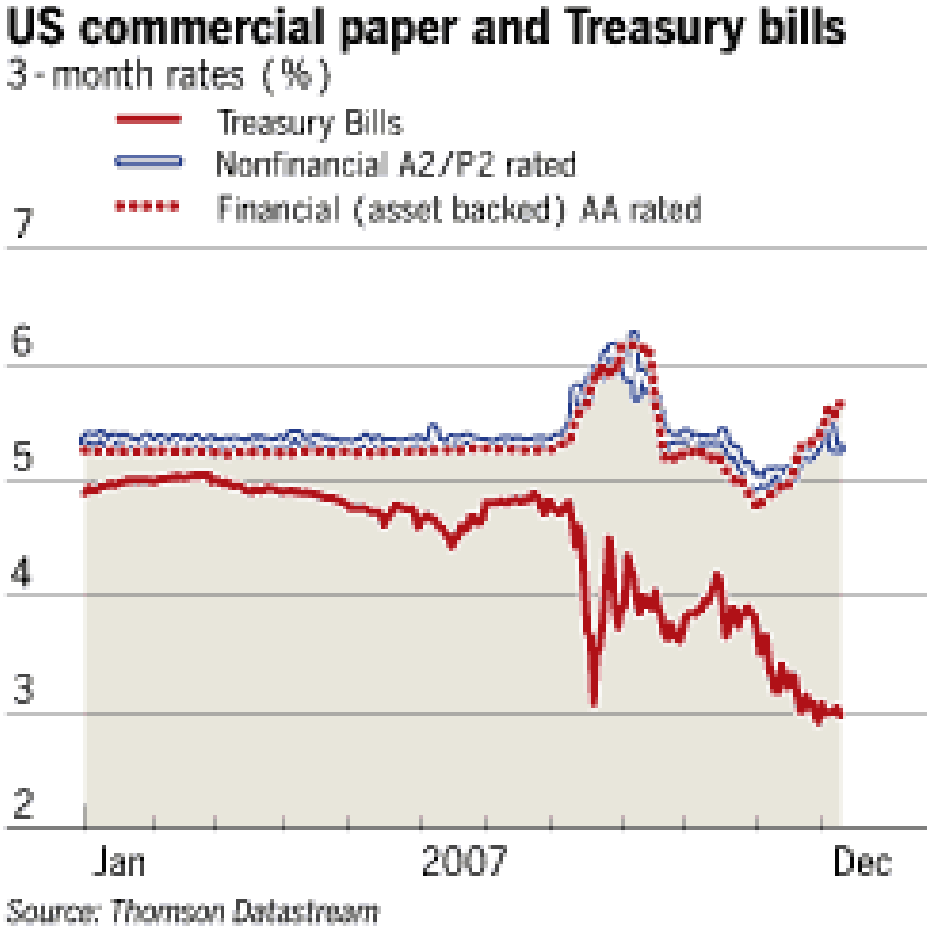
Existence Plots

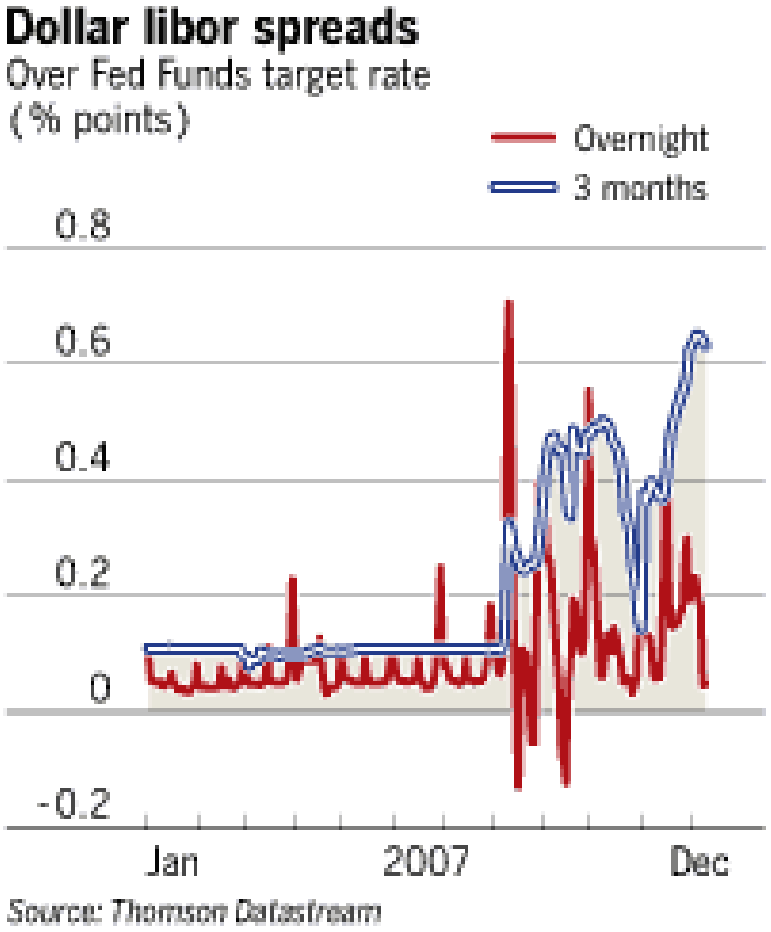


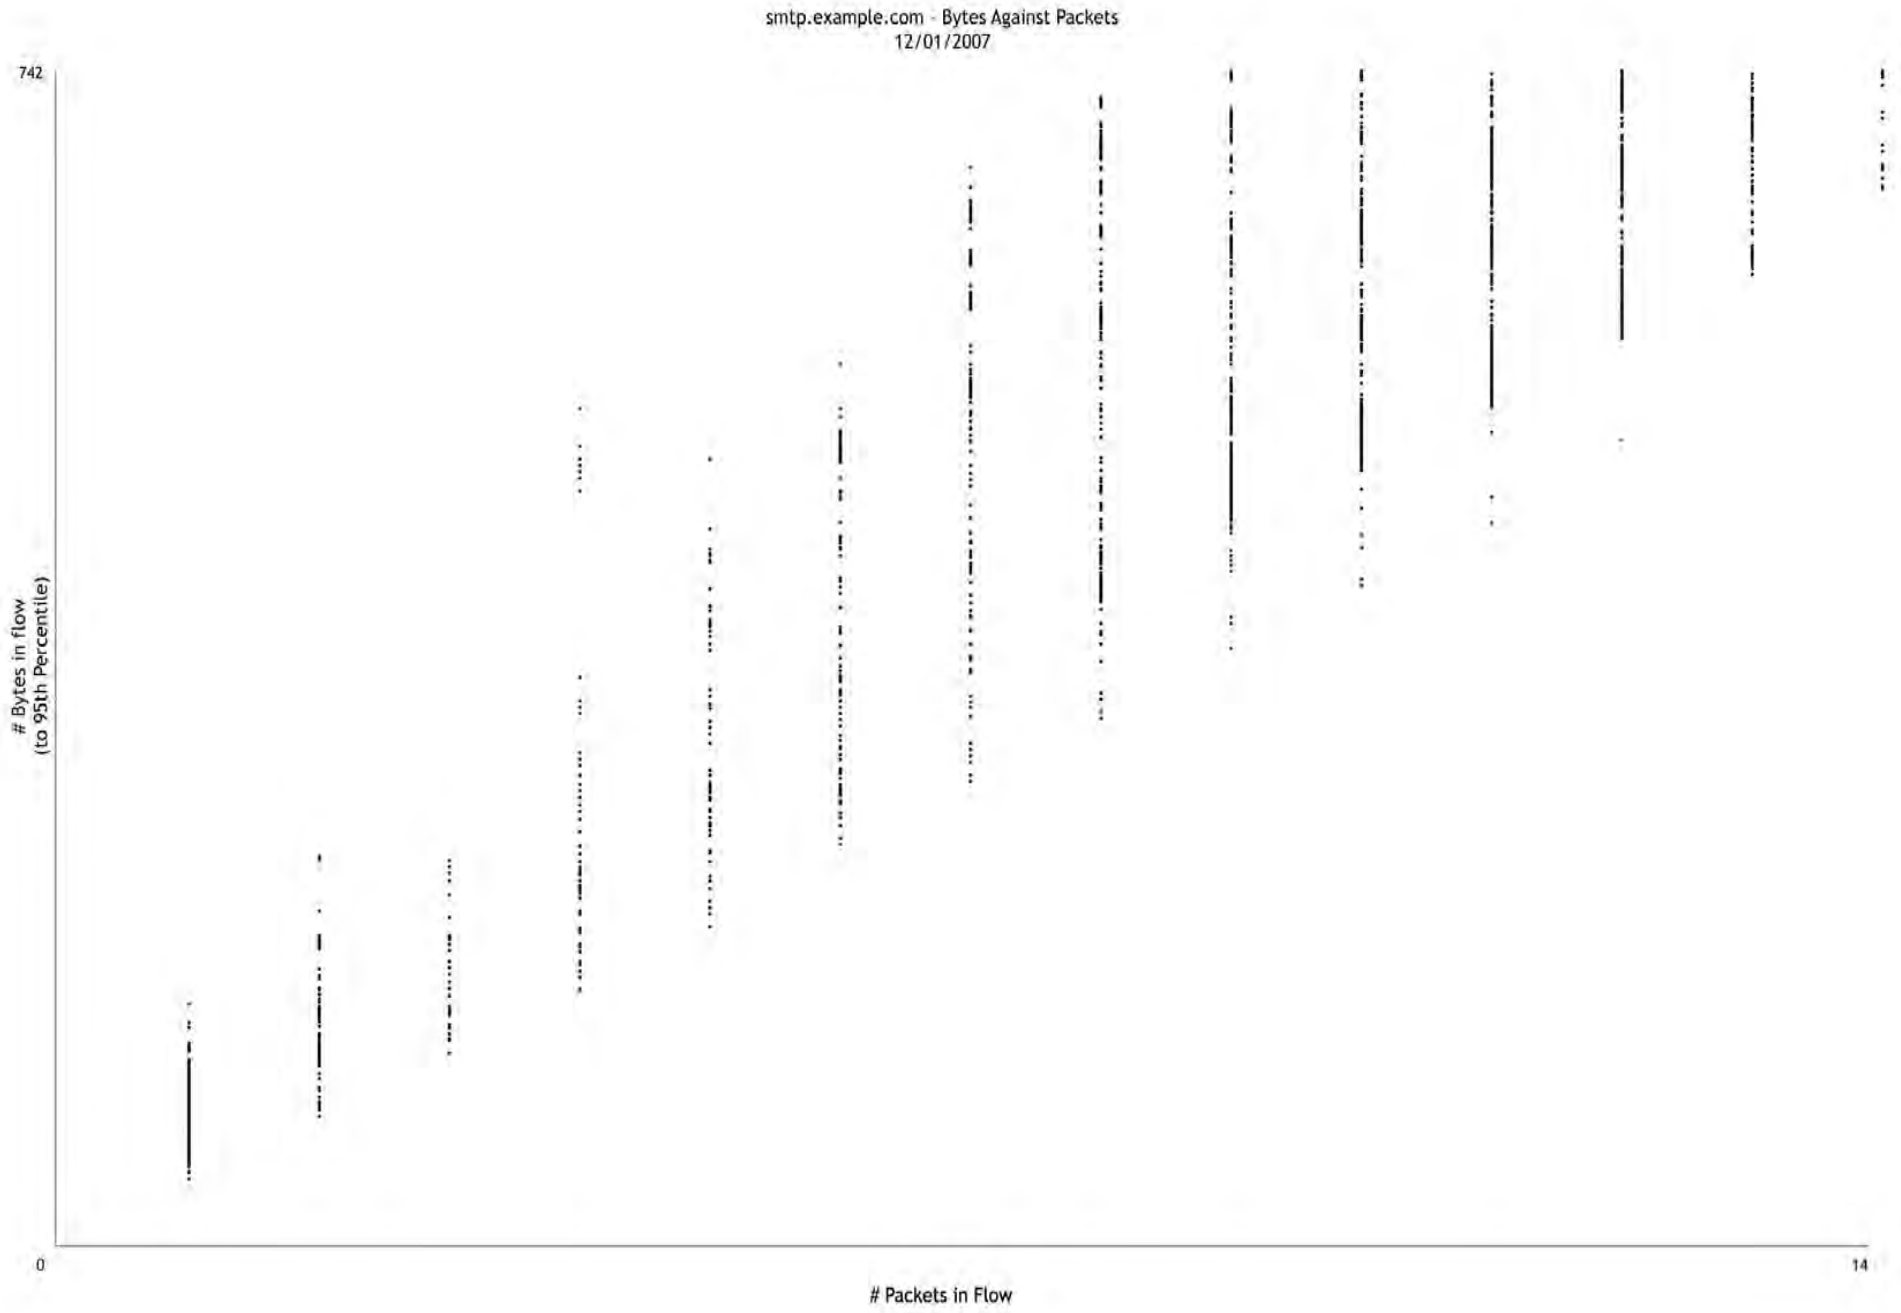


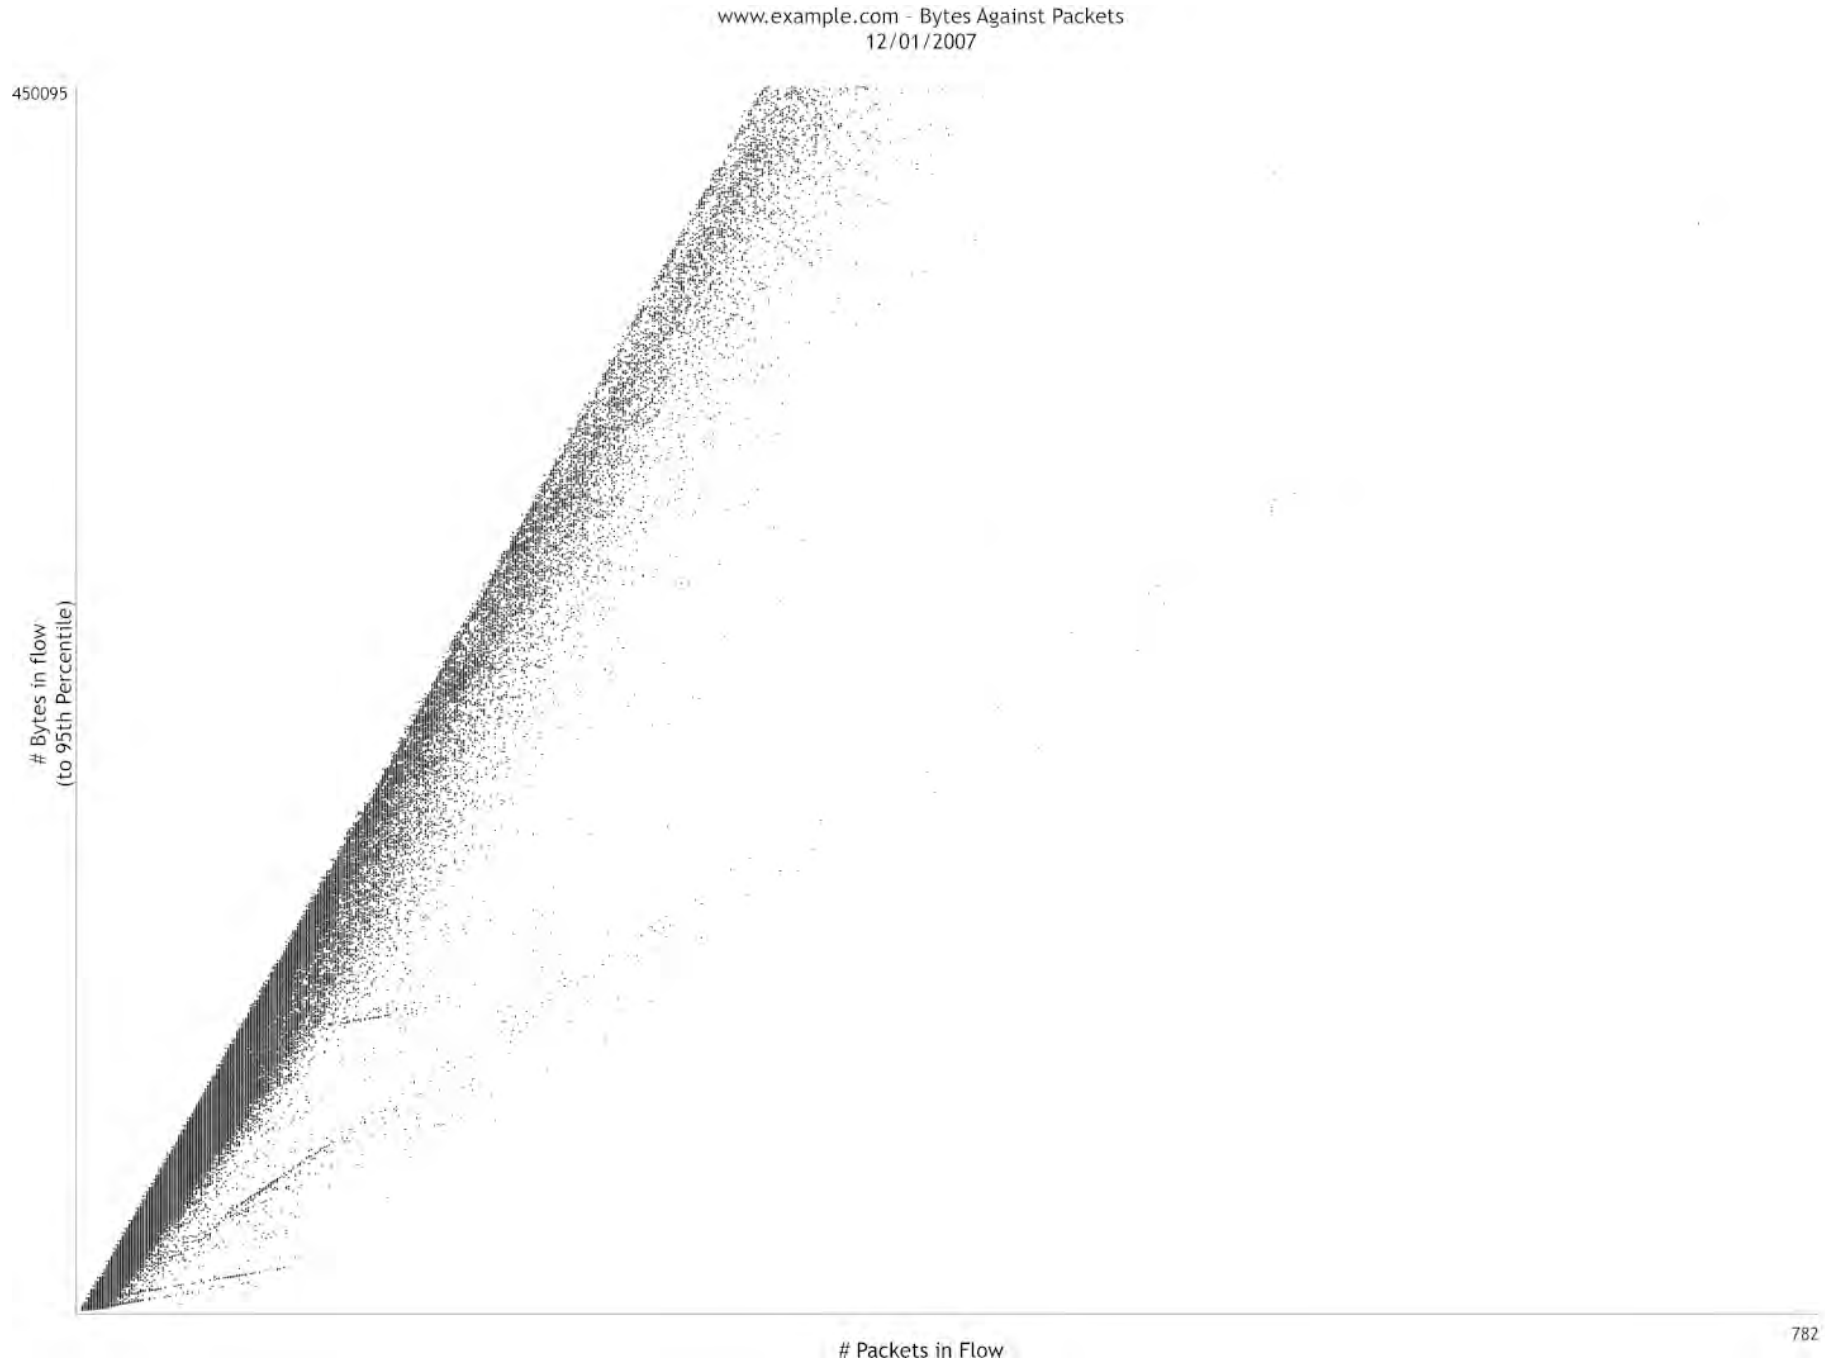
Plotting Relationships

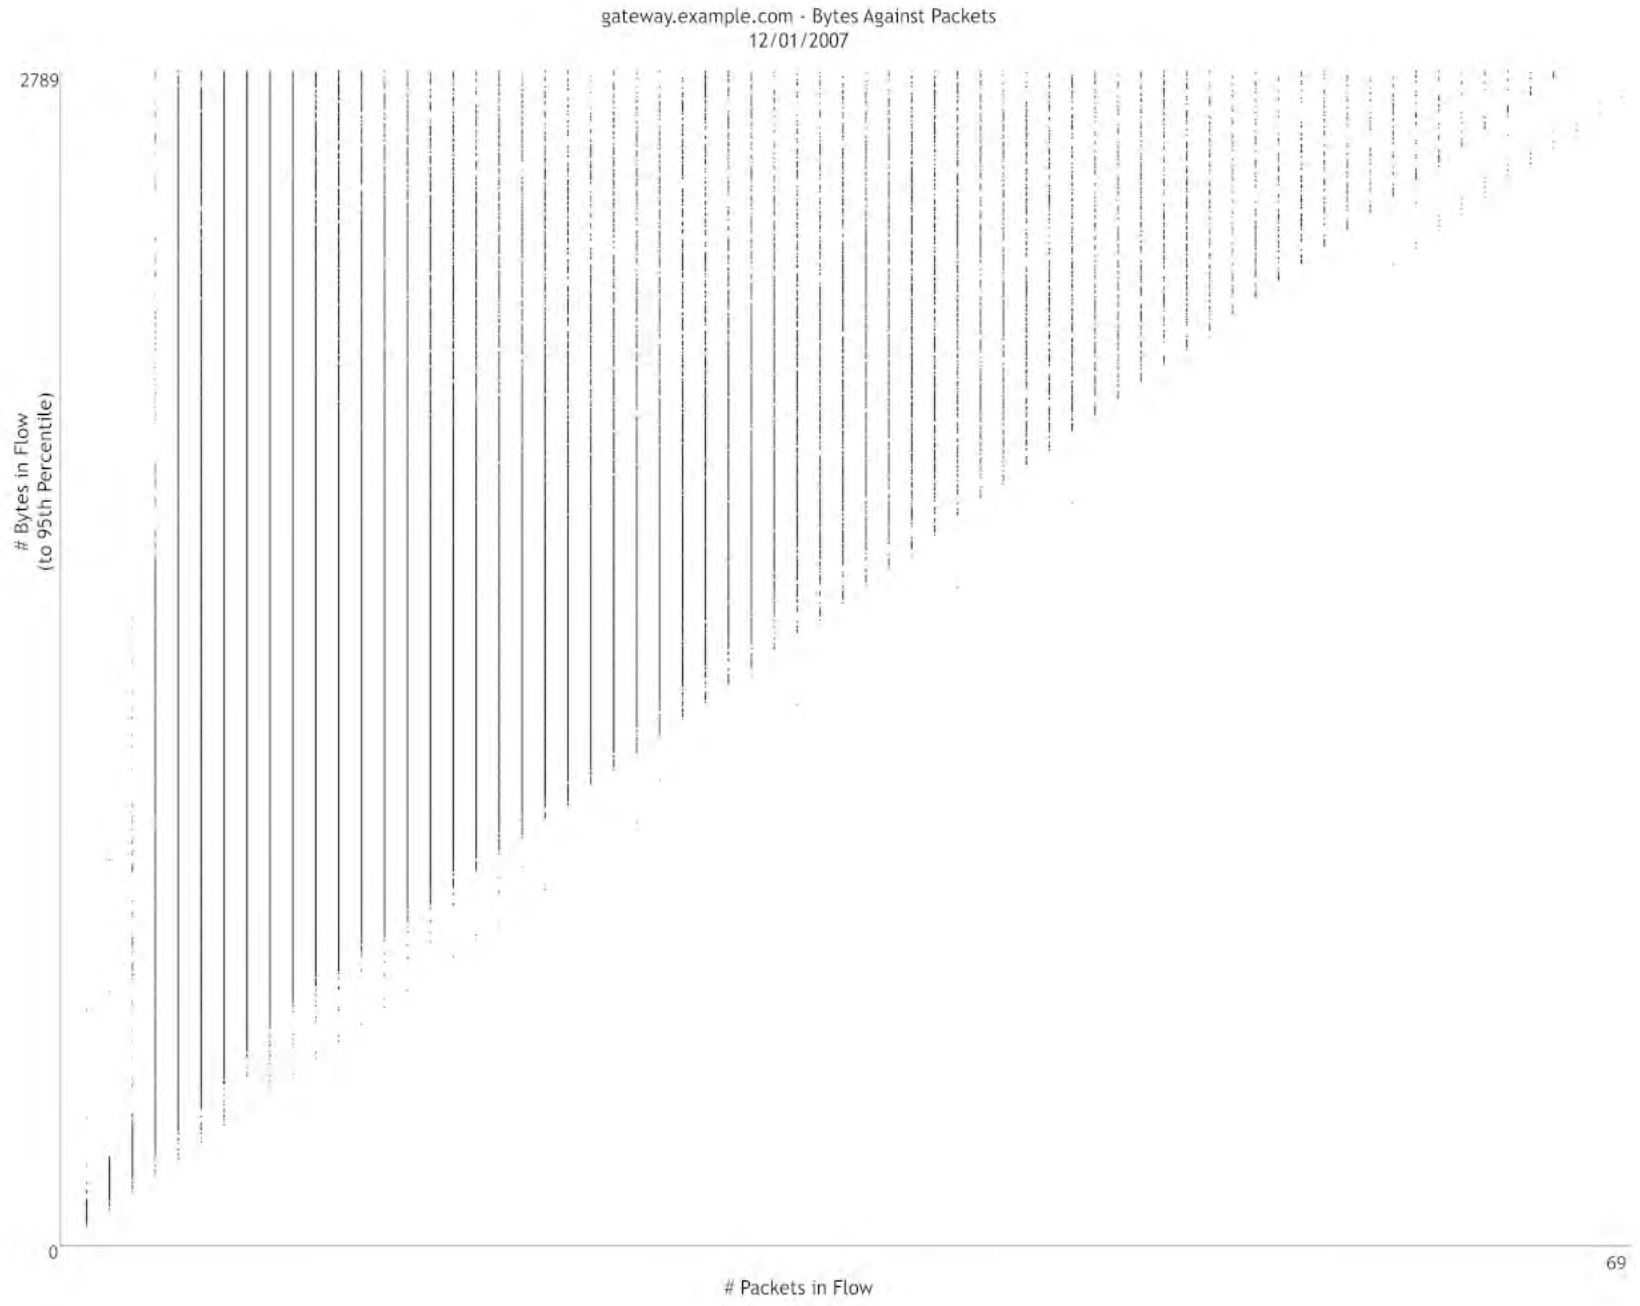
Plotting Relationships

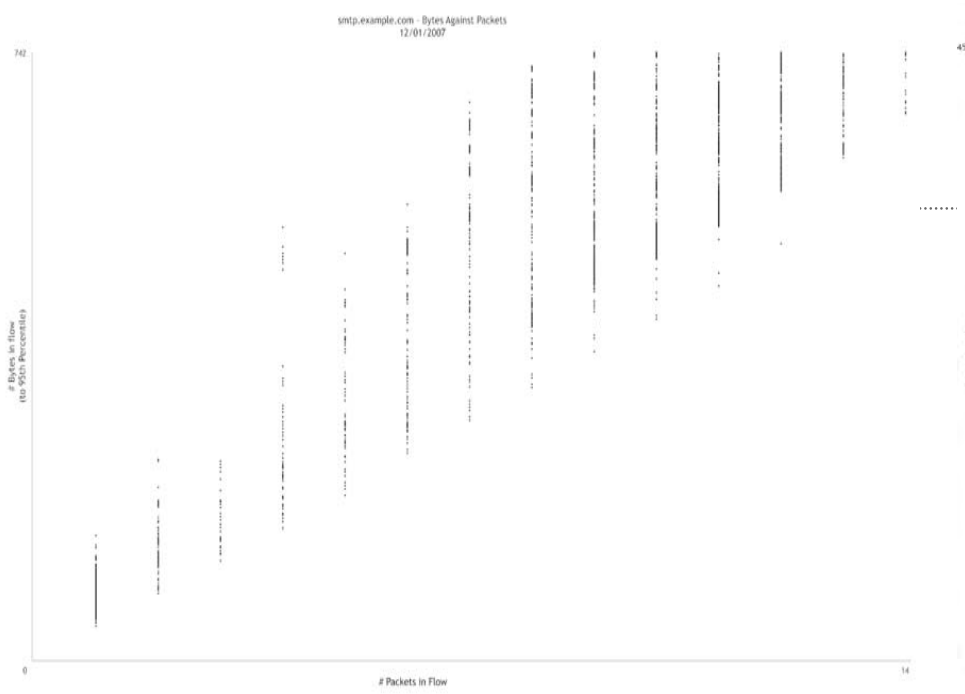


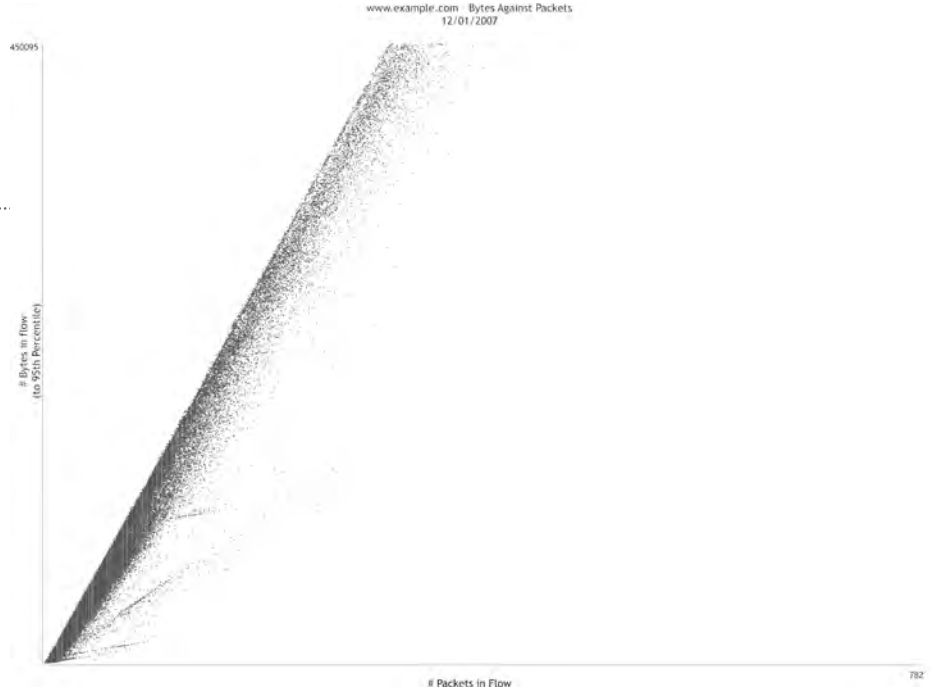


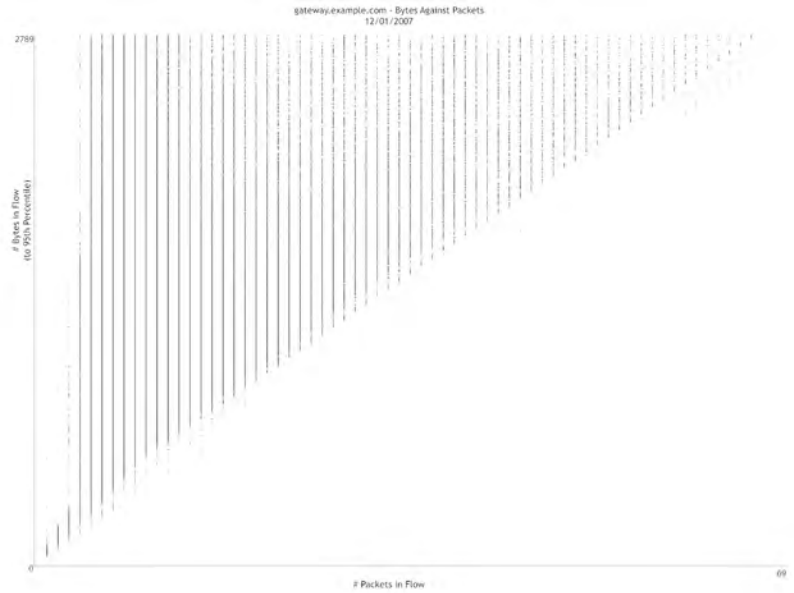

smtp.example.com - Bytes Against Packets
12/01/2007
742
Bytes in flow
(to 95th Percentile)
0
14
Packets in Flow

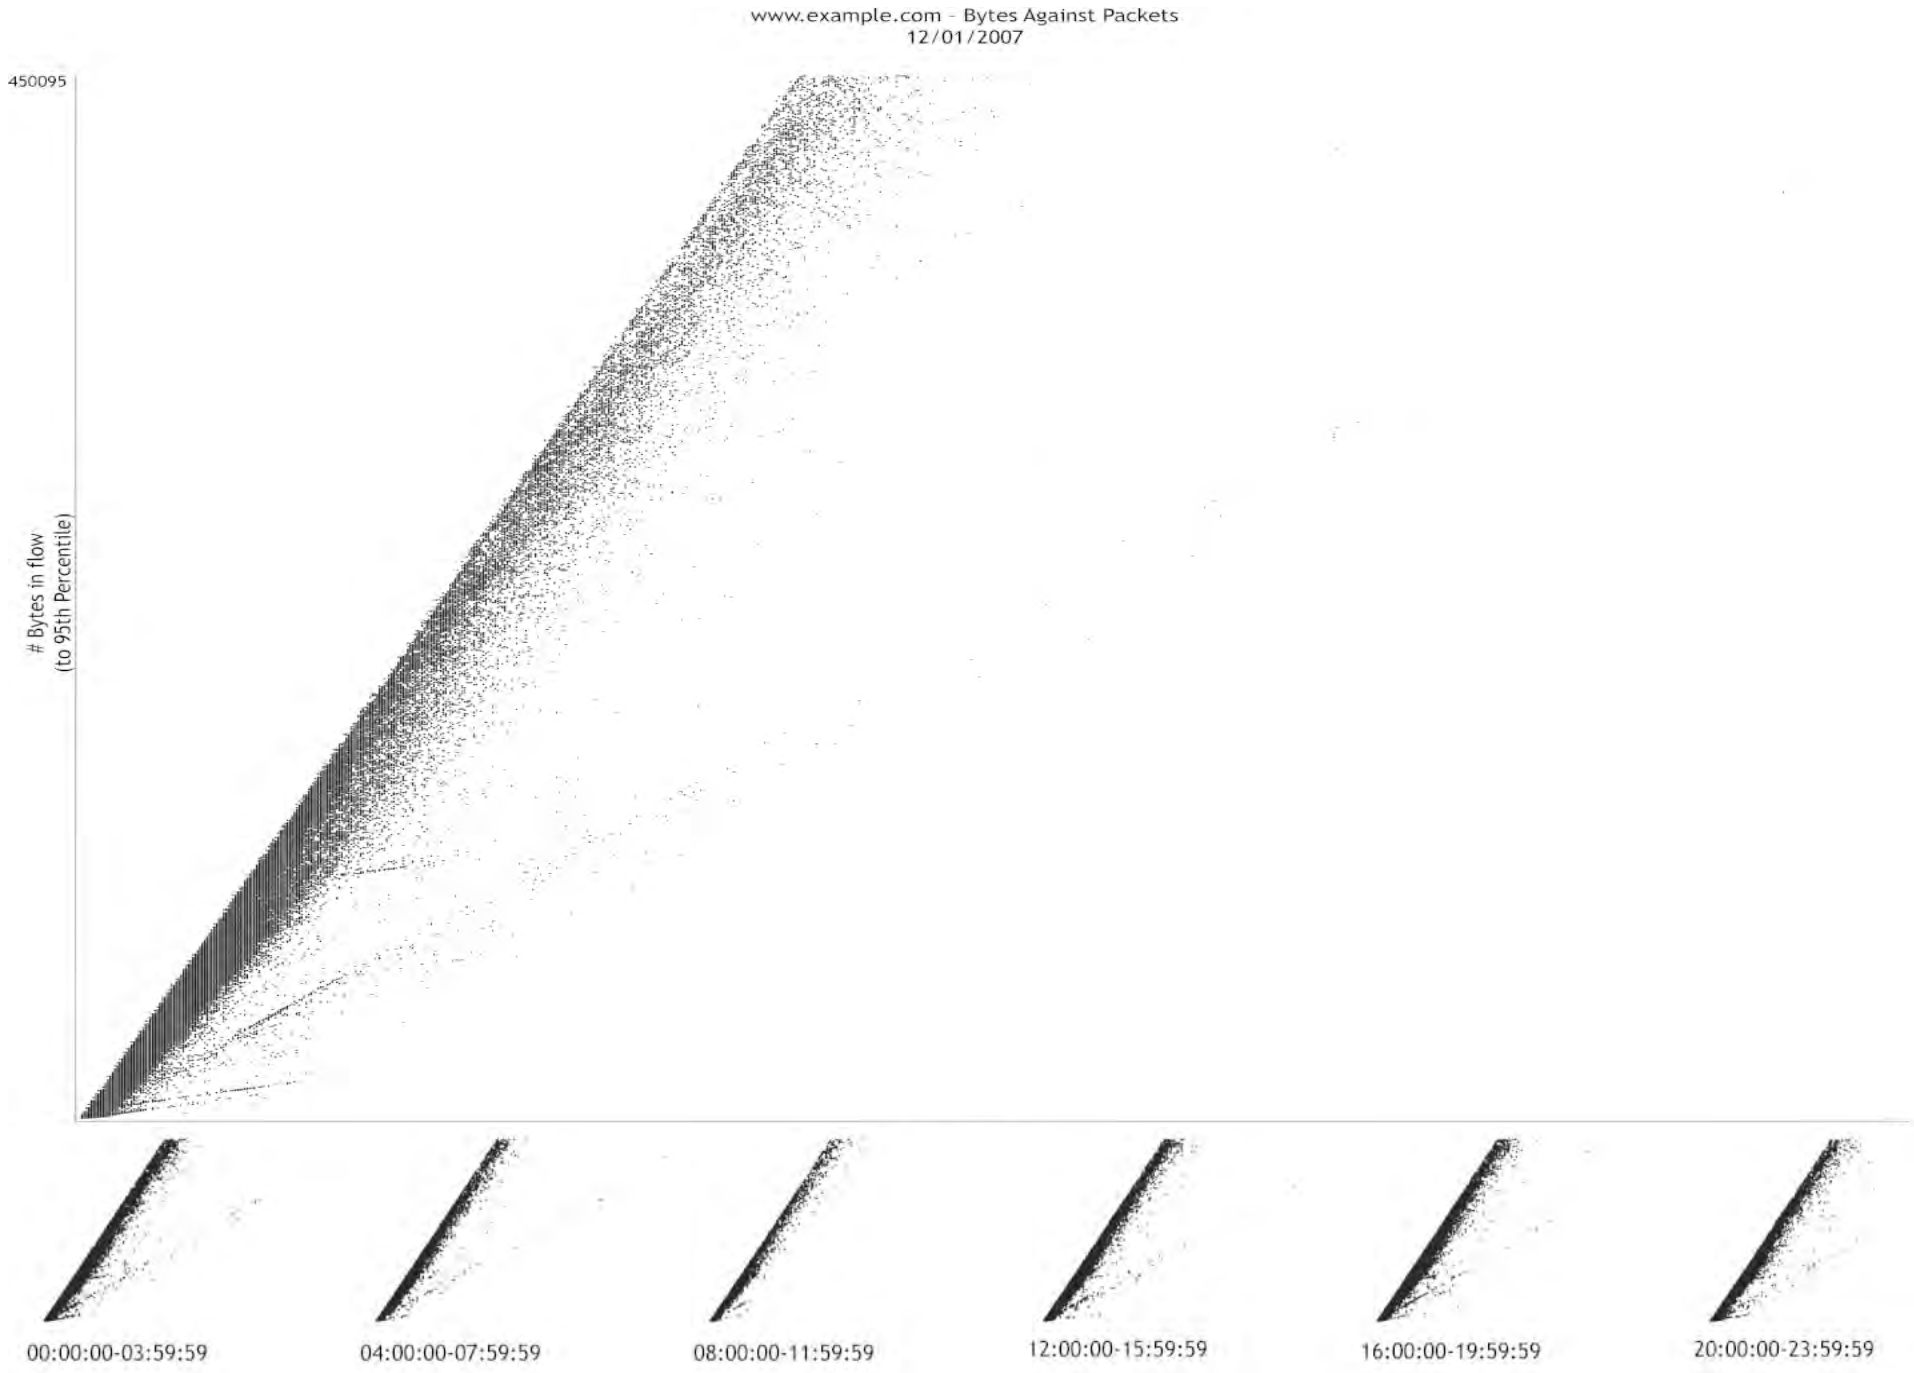

www.example.com - Bytes Against Packets
12/01/2007
450095
Bytes in flow
(to 95th Percentile)
782
Packets in Flow


gateway.example.com - Bytes Against Packets
12/01/2007
2789
Bytes in Flow
(to 95th Percentile)
0
69
Packets in Flow


smtp.example.com - Bytes Against Packets
12/01/2007
742
Bytes in flow
(to 95th Percentile)
0
Packets in Flow
14


www.example.com - Bytes Against Packets
12/01/2007
450095
Bytes in flow
(to 95th Percentile)
Packets in Flow
782


gateway.example.com - Bytes Against Packets
12/01/2007
2789
Bytes in Flow
(to 95th Percentile)
0
Packets in Flow
69


www.example.com - Bytes Against Packets
12/01/2007
450095
Bytes in flow
(to 95th Percentile)
00:00:00-03:59:59
04:00:00-07:59:59
08:00:00-11:59:59
12:00:00-15:59:59
16:00:00-19:59:59
20:00:00-23:59:59

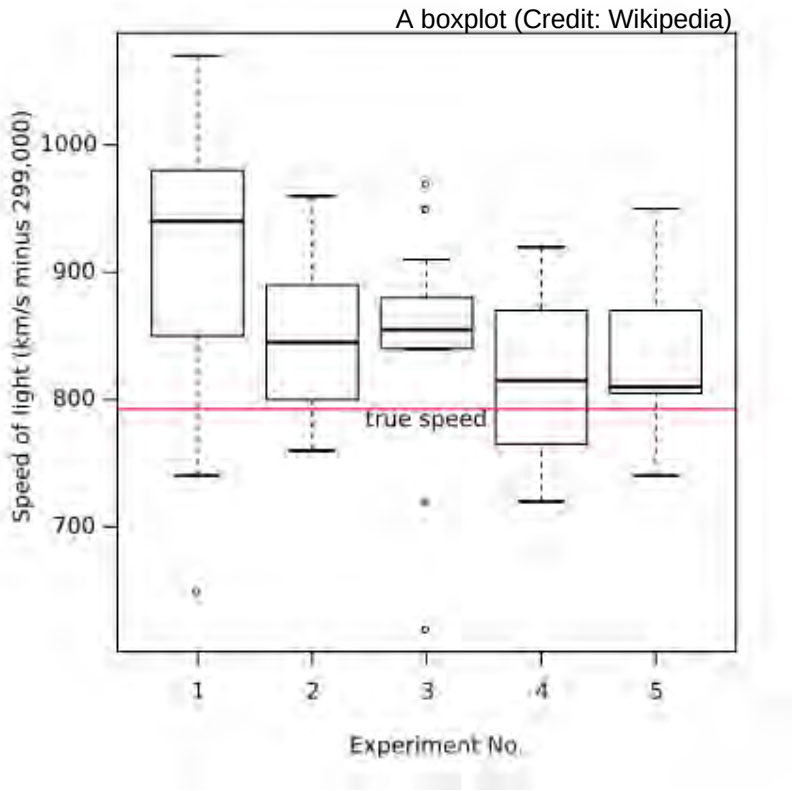


Plotting Distributions:

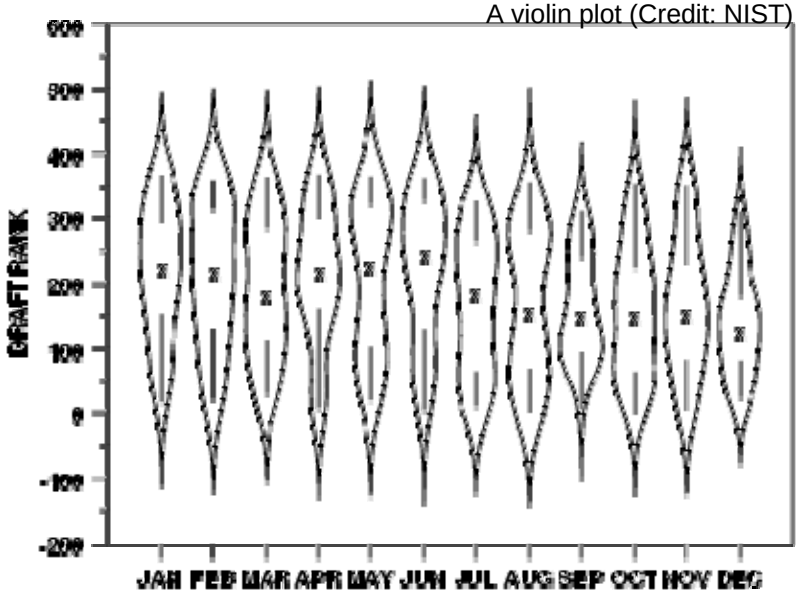
How a variable relates to itself

Box and Violin Plots

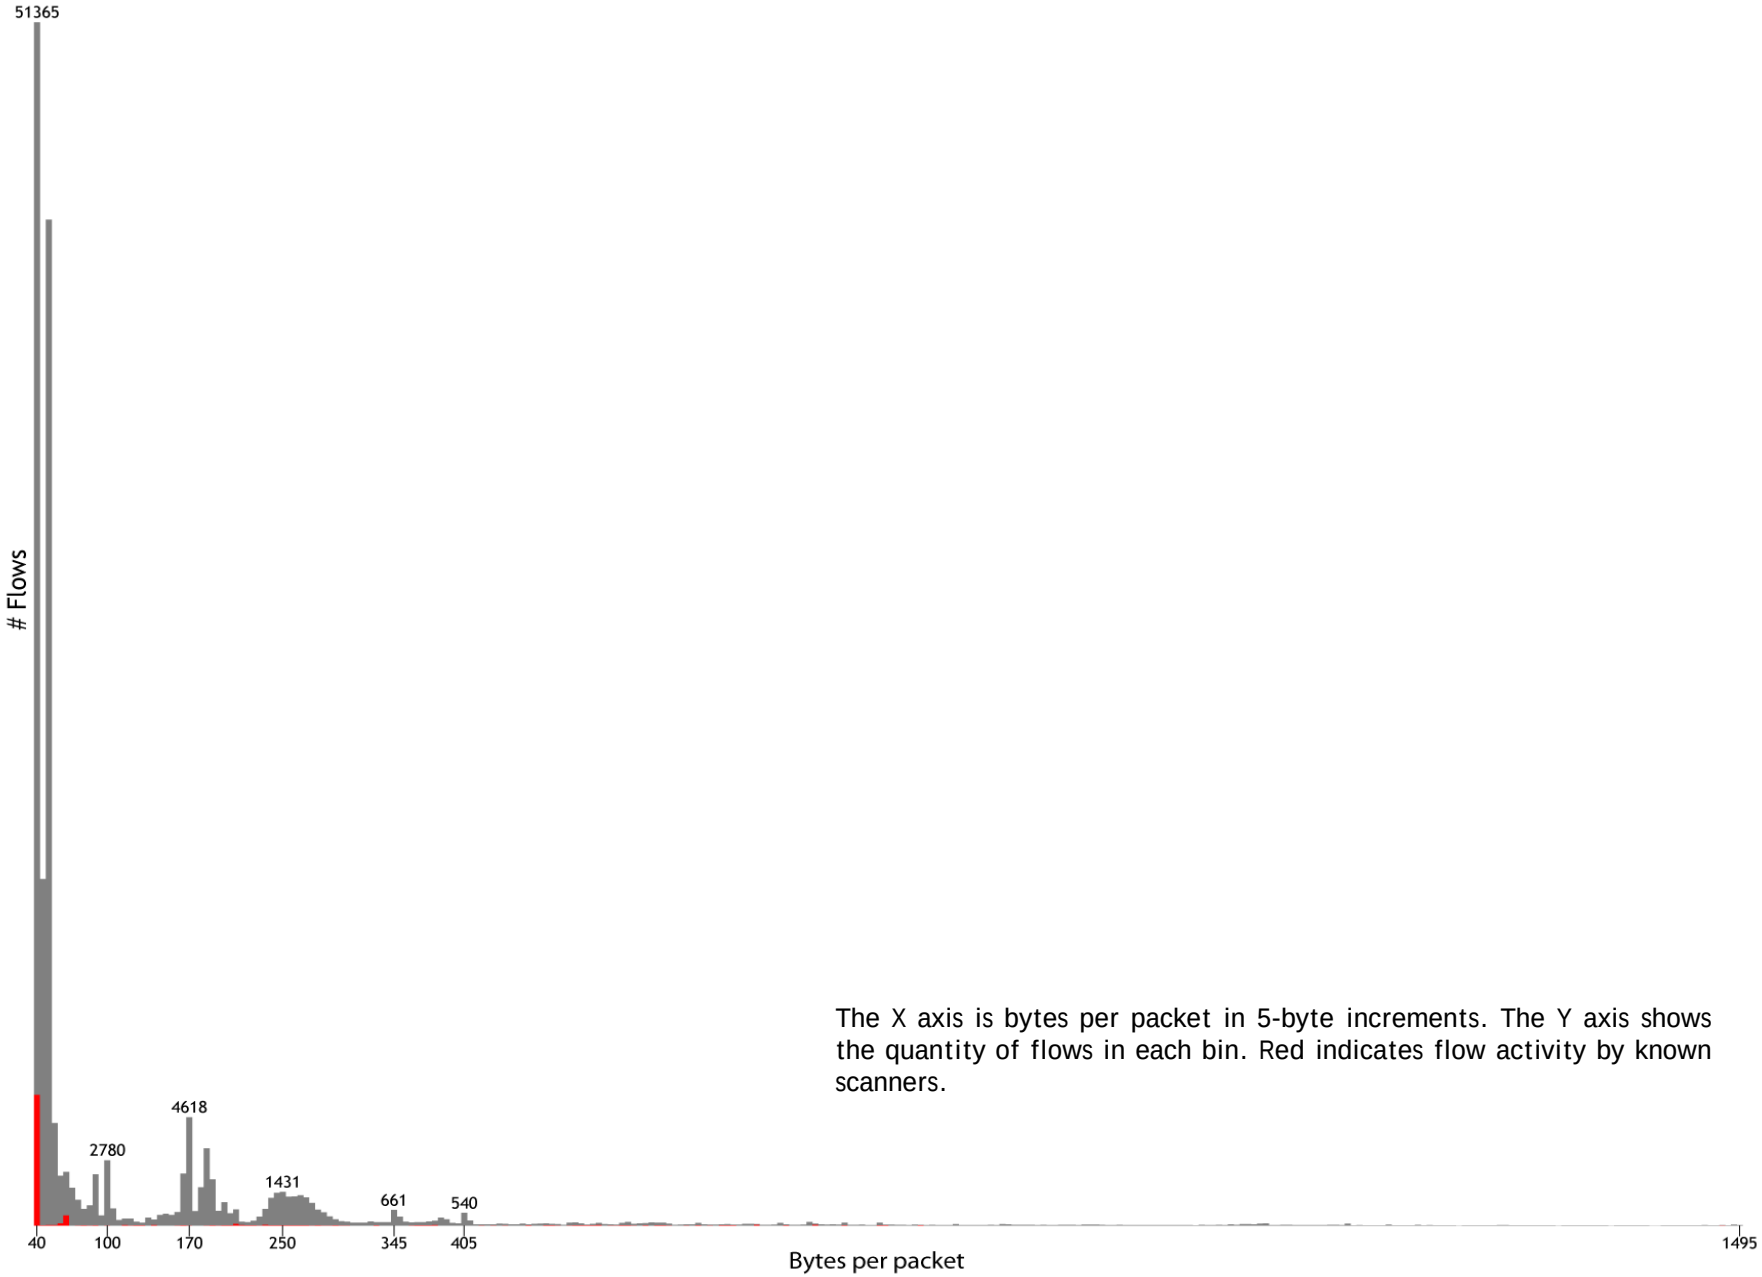
A boxplot (Credit: Wikipedia)



A violin plot (Credit: NIST)



example.com Flow Volume, Binned by Bytes per Packet
2007/12/01

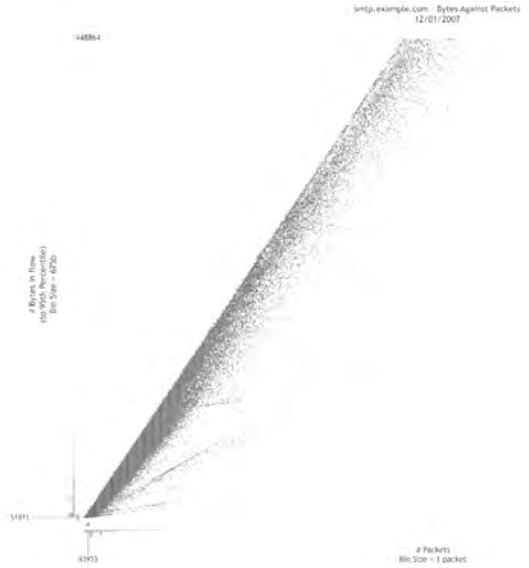


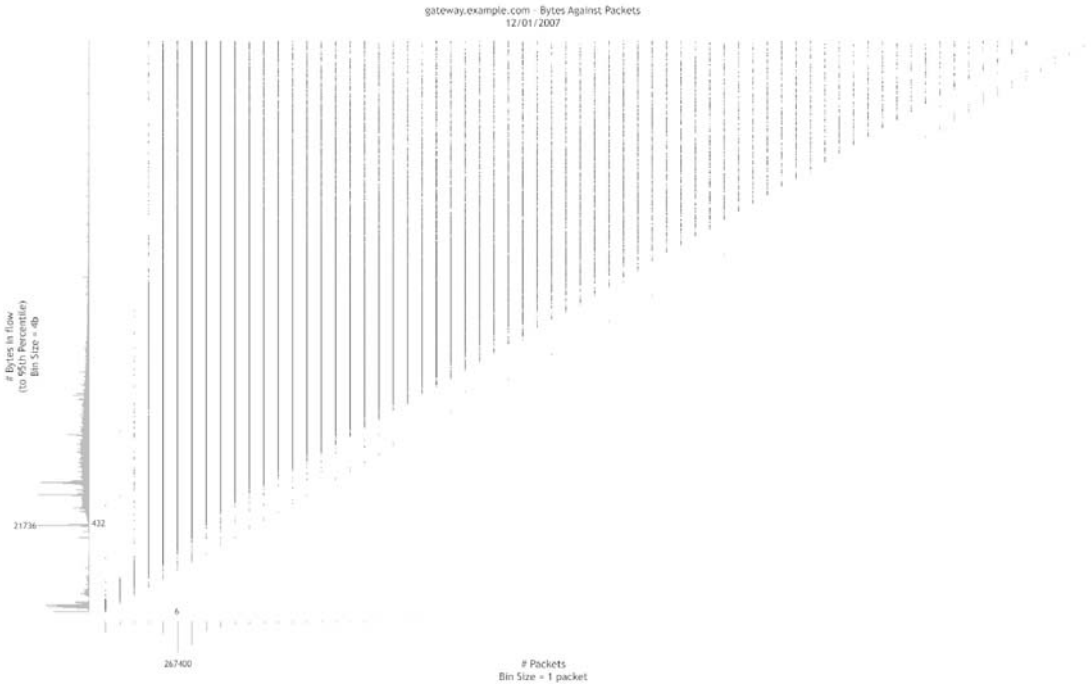
The X axis is bytes per packet in 5-byte increments. The Y axis shows the quantity of flows in each bin. Red indicates flow activity by known scanners.


smtp.example.com - Bytes Against Packets
12/01/2007
742
Bytes in flow
(to 95th Percentile)
0
14
Packets in Flow


smtp.example.com - Bytes Against Packets
12/01/2007
742
Bytes in flow
(to 95th Percentile)
0
Packets in Flow
14


smtp.example.com - Bytes Against Packets
12/01/2007
664
642
606
449
85
86635
56
35
Bytes in flow
(to 95th Percentile)
Packets in Flow

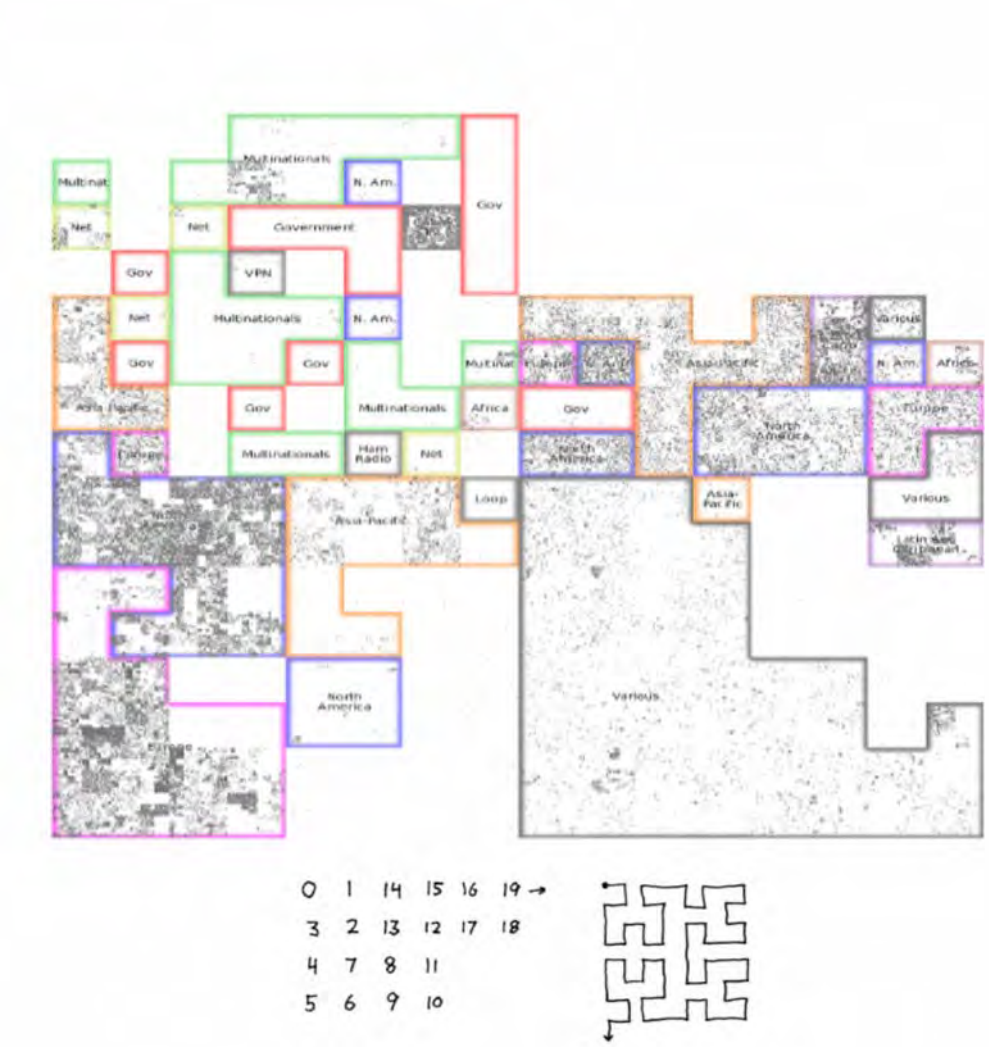

smtp.example.com - Bytes Against Packets
12/01/2007
44884
Bytes in Flow
(to 95th Percentile)
Bin Size = 679b
51871
0
43933
Packets
Bin Size = 1 packet


gateway.example.com - Bytes Against Packets
12/01/2007
Bytes in flow
(to 95th Percentile)
Bin Size = 4b
21736
432
6
267400
Packets
Bin Size = 1 packet

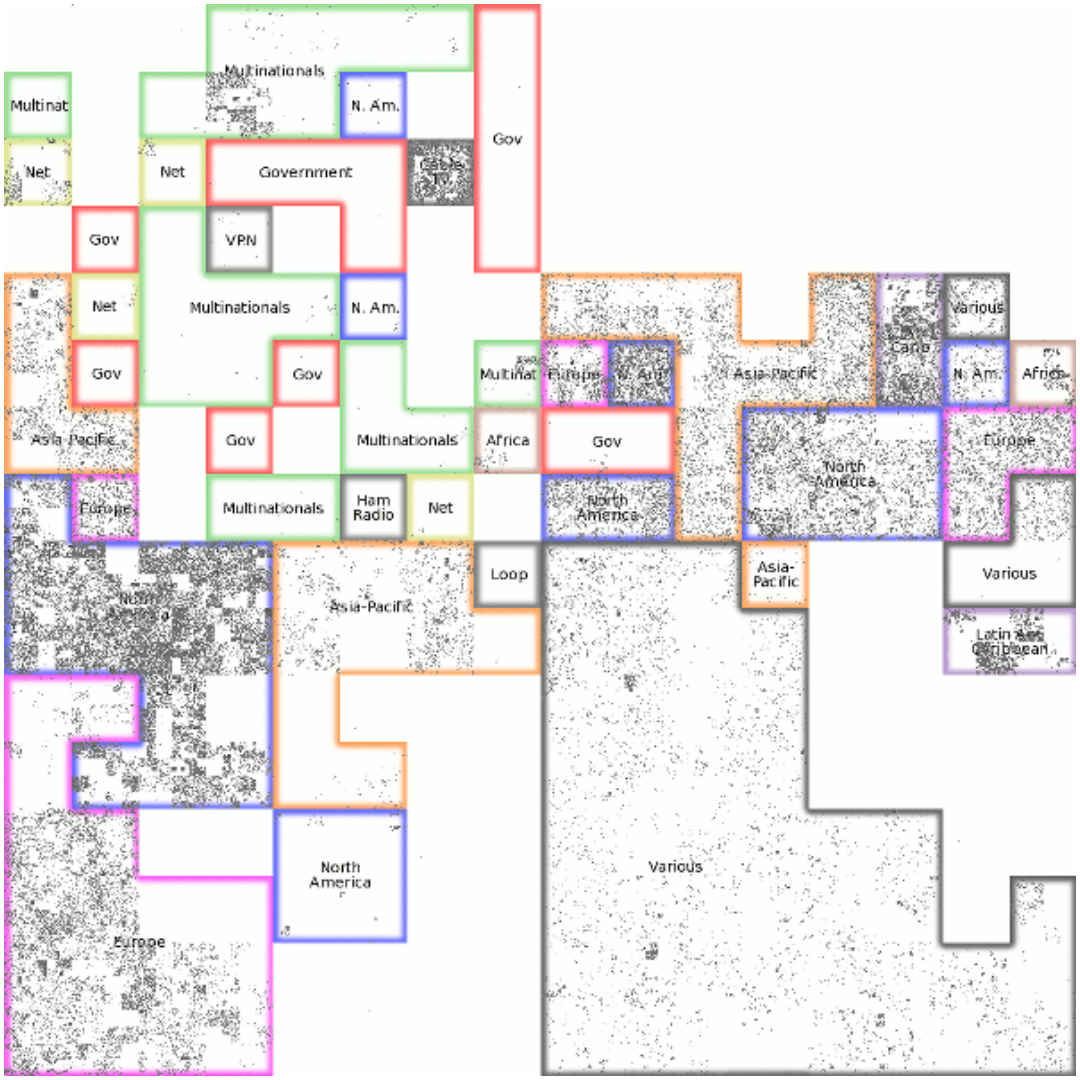
Hilbert Curve:

Broad Scale Visualization

Hilbert Curve



Hilbert Curve (The Movie)



Additional Resources

The R Project. *Introduction to R.* Chapter 13: Graphics. http://cran.r-project.org/doc/manuals/R-intro.html#Graphics

Tufte, E. R. *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press, 1983.

Tufte, E. R. *Envisioning Information.* Cheshire, CT: Graphics Press, 1990.

Tufte, E. R. *Visual Explanations: Images and Quantities, Evidence and Narrative*. Cheshire, CT: Graphics Press, 1997.

Tufte, E. R. *Beautiful Evidence.* Cheshire, CT: Graphics Press, 2006.

Wilkinson, L., et al. *The Grammar of Graphics*. New York: Springer-Verlag, 1999.

Visualization as an Analysis Tool:
Presentation Supplement

This document is a supplement to the presentation "Visualization as an Analysis Tool" given by Phil Groce and Jeff Janies on January 9, 2008 as part of FloCon 2008. The intent of the presentation was to demonstrate how simple tools used together can provide significant insight into network behavior.

This supplement includes annotations to the presentation slides and a set of key points in the presentation, as well as supplementary points that were not included for the sake of brevity. Except where noted, all examples are drawn from real network flow data.

Key Points

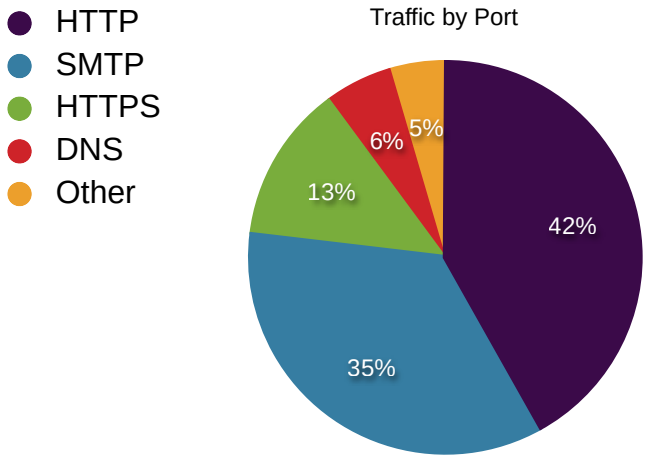
Always tell the truth. Many of the same techniques for making the eye aware of important distinctions in the data can also highlight unimportant distractions, or even create false impressions that the data do not support. Always ensure that the visualization presents an accurate picture of the data.

Learn how to use your tools. It's better to have a limited set of tools you know how to use than a whole bag of tools you don't understand. If nothing else, understand the limitations of the tools so you know when your question demands a new one.

Facilitate direct comparison. Use human perception to your advantage. Put comparable visualizations on the same page so people can see them in peripheral vision and switch back and forth quickly. Align visualizations along common axes. Make sure similar things look similar (e.g., by using the same scales).

Combine complementary visualization techniques. Use visualizations whose insights complement each other in ways that facilitate easy comparison between them.

Tables are visualizations, too. For small sets of data, a table may be the best way for people to consume the data. Compare the following (sample) data formatted as a pie chart and a table:



Port	Traffic (% Volume)
HTTP	42
SMTP	35
HTTPS	13
DNS	6
Other	5

The table takes less space to communicate the same information, and it is easier to map the type of traffic to the value. In the pie chart, the reader must consult a key to map a color to a traffic type, then find the color on the chart. (Putting the type label directly on the pie slice causes problems with the smaller pieces; using callout-style labels makes the pie chart even larger.)

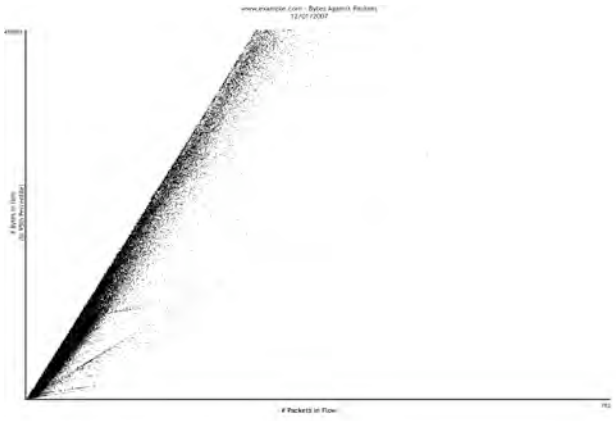
Distinguish between useful and useless precision. The right numbers in the right place are critical to understanding a visualization. Examples include maximum and minimum observed values; selected local maxima and minima; and relevant baselines such as the start and end values of the scale, medians or other useful "average" values, and specific "important" data points (e.g., a point representing an important host or a flow thought to be the source of a compromise).

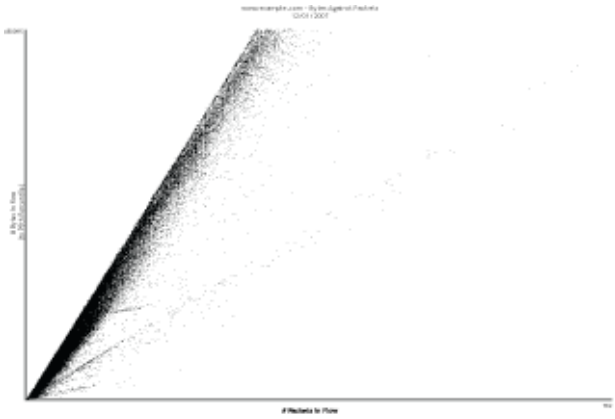
In general, if the exact numbers don't provide an important perspective on the whole visualization, it's probably a distraction. If all the numbers are important, the best visualization may be a table.

Consider sampling. For constant-magnitude data, a visualization over a sample of the data may tell the story as convincingly as a visualization over the full set of data. If the generation of the visualization takes significant time, sampling may improve performance.

Choose your display media wisely; don't underestimate paper. The primary display media available to most analysts are paper and computer monitors. In our experience, analysts overwhelmingly use monitors over paper. Both have advantages, however. Computer screens allow types of interactivity that paper cannot, and digital copies of visualizations can be easily sent electronically. Visualizations on paper can be read without special equipment, and annotated with only a pencil. Moreover, paper resolution ranges from 300 to 1000 dpi. Screen resolution typically ranges from 72-100 dpi.

For comparison, here is the same scatterplot rendered at 300 dpi and 100 dpi:





Use the appropriate number of dimensions. Paper and screens are naturally two-dimensional; color, perspective and motion can provide some additional dimensionality, but will never be as effective as length and width on a flat screen or piece of paper. For example, decorating a scatterplot axis with a histogram of the data along that axis communicates an additional dimension of data density often plotted to lesser effect with color or an isomorphic rendering of a three-dimensional plot.

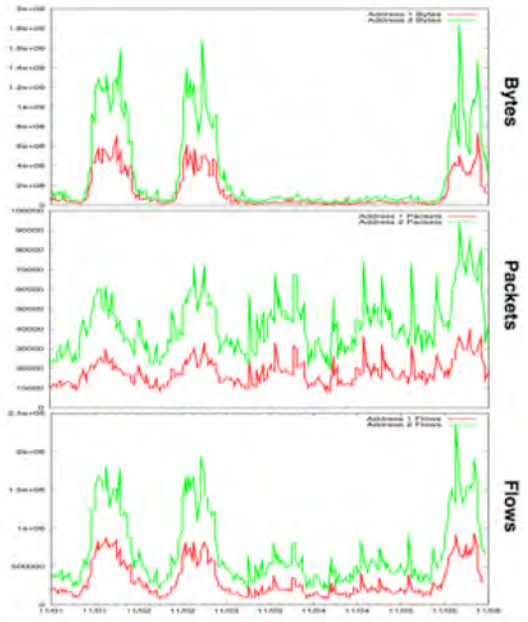
Make everything explicit. When generating visualizations for personal consumption, the most important thing is finding an insightful view on the data. When passing visualizations to others (including yourself in the future), annotate the visualization with everything required to understand where the data came from and what processing has been done on it. (E.g., the command used to extract the data, how the data may have been trimmed for readability, what other transformations have been done on the data.) In particular, if your scale doesn't start at zero for linear scale plots or one for log scale plots, make a special point of noting it.

Basic Visualization Types

Time series

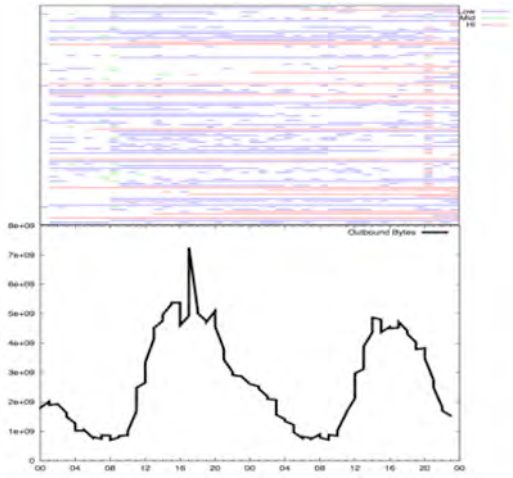
Relating data to time is so intuitive and useful that it risks being used to the exclusion of other visualization tools. It makes sense, then, to optimize time series visualizations as much as possible.

A common problem in time series data is relating multiple independent series (e.g., volume measurements of bytes, packets and flows) by time to give a clear picture of an event. Often, the scales of each series diverge too widely to plot on the same scale without losing detail. Using different scales for each series is an option if the designer can communicate this decision clearly to the reader. Another (often simpler) solution is to plot the series independently and align them on a shared scale, as in the figure to the right.



When comparing very large numbers of series, the above approach breaks down, unless the data is very tightly coupled (e.g., EKG or seismic data). The existence plot trades measurement resolution for scale by defining value ranges and plotting each series as a single line, colored by the range in which the value for that time resides.

To mitigate the loss of resolution, it is often useful to pair an existence plot with a traditional time series that relates to all the series in the existence plot.

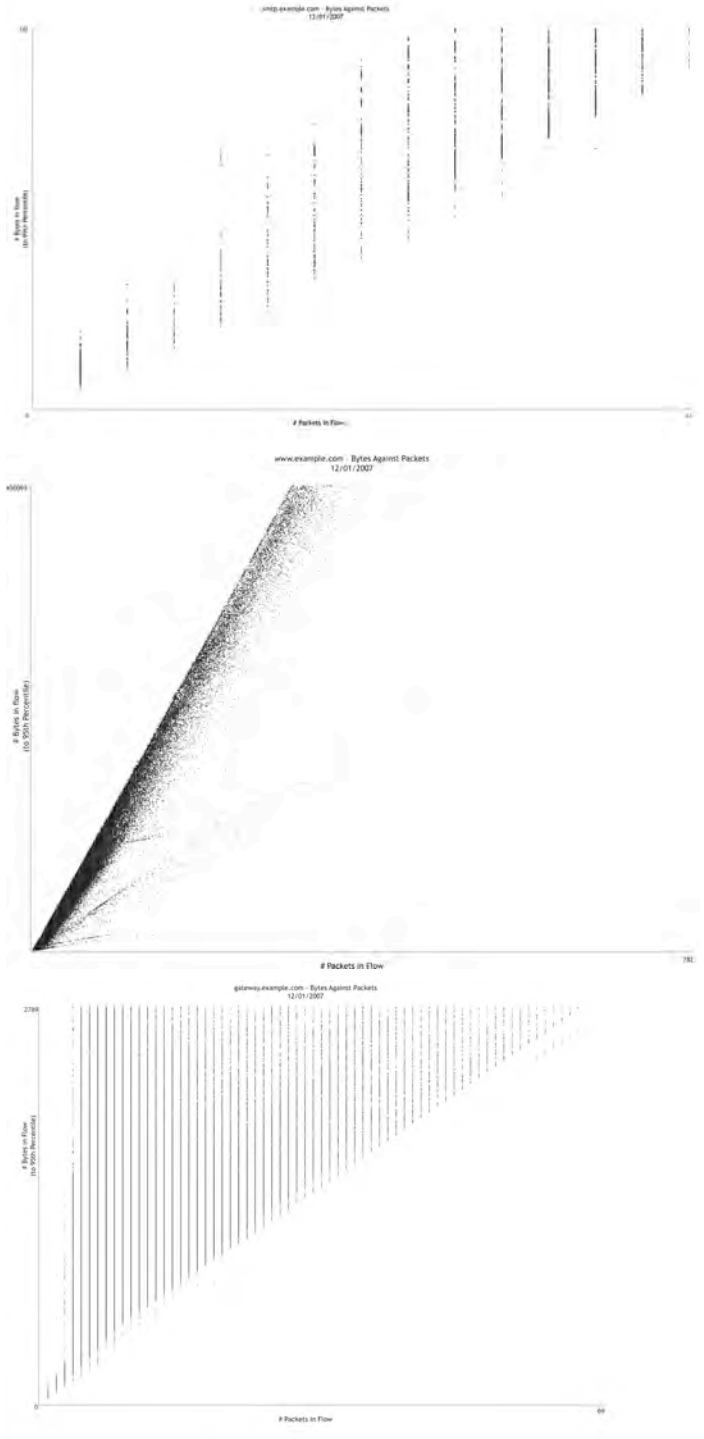


Visualizing Relationships

Time series plots are a specific instance of relating one dimension with another (time). Scatterplots are generalizations of this approach. By plotting points in space, the eye can perceive relationships between these dimensions as lines, shapes or other patterns.

These three machines serve three very different network roles. (Mail server, web server and gateway, respectively.) This is reflected in the different (but consistent) ratios between bytes per flow and packets per flow in their traffic. This, in turn, is visible in the very distinctive "shapes" their traffic takes when these dimensions are plotted against each other.

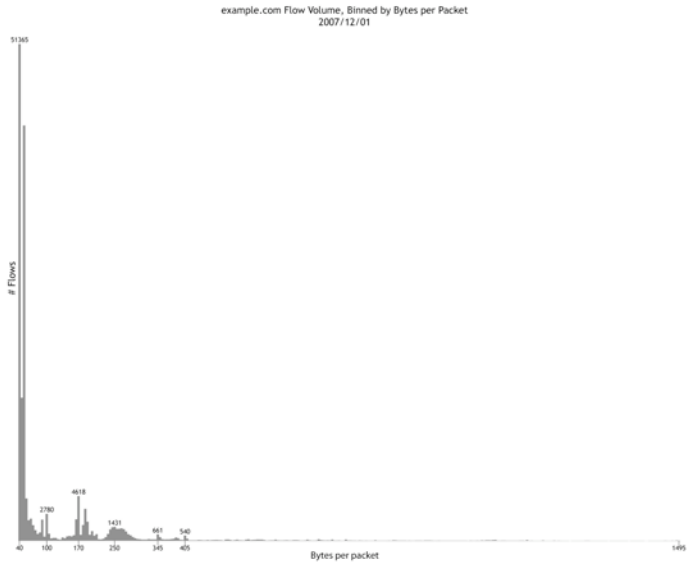
In this plot, multiple points that share exactly the same values show up as a single point. The next section addresses this deficiency.



Visualizing Distributions

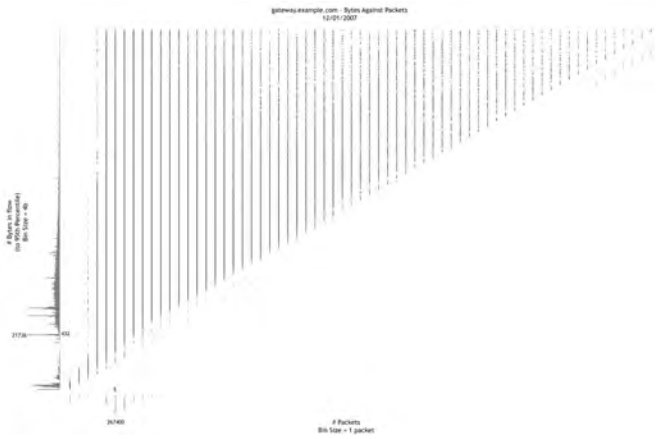
There are many ways to visually analyze the distribution of sets of single values—this document focuses on histograms, but box plots, whisker plots and violin plots; CDF (cumulative distribution function) plots are all in common use.

The distribution of bytes per packet for flows associated with a given host. Most connections fall within a few narrow ranges of values. As with the scatterplots above, the "shape" of the histogram is characteristic of the host behavior.



Because of the physical dimensions of a histogram, and the fact that it operates on a single (data) dimension, it complements scatterplots well to indicate density of data.

In the example at right, the distribution of values gives no indication that a disproportionate number of flows visualized have low packet and byte counts. (267,400 flows have 6 packets; 21,700 flows contain only 432 bytes.)



Additional Reading

The R Project. *Introduction to R.* Chapter 13: Graphics. http://cran.r-project.org/doc/manuals/R-intro.html#Graphics

Tufte, E. R. *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press, 1983.

Tufte, E. R. *Envisioning Information.* Cheshire, CT: Graphics Press, 1990.

Tufte, E. R. *Visual Explanations: Images and Quantities, Evidence and Narrative*. Cheshire, CT: Graphics Press, 1997.

Tufte, E. R. *Beautiful Evidence.* Cheshire, CT: Graphics Press, 2006.

Wilkinson, L., et al. *The Grammar of Graphics*. New York: Springer-Verlag, 1999.

One Year of Peer to Peer

Ron McLeod, BCSc, MCSc.
Director - Corporate Development Telecom Applications Research Alliance
Doctoral Student, Faculty of Computer Science, Dalhousie University

Presentation Summary

This presentation will profile the result of the growth in peer-to-peer applications on a sample network and describe the resultant massive increase in the diversity of traffic. This diversity impacts the ability to profile baseline normative behaviour using Blind Flow Analysis.

I will also briefly discuss the application of SiLKtools, Neural Networks and Bioinformatic strategies to Blind Flow Analysis of real world security problems and how that analysis is affected by the growth in recreational/user driven applications.

What began as a basic design principal of end-to-end management with popular applications in recreational computing is quickly becoming a dominant evolutionary force in network traffic patterns.

Traffic patterns are becoming emergent properties influenced by the voluntary adoption of new systems by individuals without any collective intent.

The network is evolving at the edges.

“Peer-to-Peer is the basic design of the Internet” – Christian Huitema

Sample Network Description

- A Multi-tenant Commercial Network consisting of:
  - ~ 40 user assigned hosts, actual number subject to minor fluctuations over time.
  - ~40 special hosts not assigned to individual users. These hosts form parts of various temporary development and experimental environments.
  - Users were apprised that Network flow data was now being captured for experimental and management reasons.
  - Payload data was neither collected nor examined.
  - Analysts did not have access to the content of specific hosts for further investigation.
  - For confidentiality reasons the identity of the Network is not specified in this Presentation.

A Review of Blind Flow Analysis

The Need for Classification Based on Minimal Information (the extreme case in the world of tomorrow)

- Capturing and examining payload contents is widely viewed as a potential violation of privacy and placed in a category similar to listening in on a telephone call.
- Even attempts to use information derived from the payload (such as ngrams) do little to alleviate the fundamental concern of the user surrounding access to the payload.
- In multi-tenant commercial environments this user concern may be based in protection of commercial confidentiality.
- There is less (although not zero) concern among the user community with regard to the capture and investigation of packet header data (some concern for Source and Destination IP's and MAC's).
- Therefore, the network analyst may be limited to examining a severely reduced subset of the packet header information in an attempt to determine if the system under their management (or monitoring) is operating properly or experiencing anomalous behavior.
- The loss of access to the originating address information means that the analyst no longer has access to a unique field in the data that identifies the individual hosts in the traffic (i.e. they cannot tell one computer from another by looking at the remaining flow record traffic alone).
- In such an environment, what is required is a method of classification that relies on minimal information and the development of traffic flow behaviour models that use only this information.

One Strategy for Comparing A Suspicious Host to a Standard Workstation Using Blind Flow Analysis

Local Baseline Workstation Behaviour (BWB)	**Suspicious Host**
Bytes Transferred in one month < 20 million per month	45 billion per month
Internal DIPs < 10 per month	3 per month
External DIPs < 20 per month	1.74 million per month
Protocols: 1 < 2 %	1 1 %
6 > 70 %	6 9 %
17 < 30 %	17 90 %
Number of Protocols < 5	3

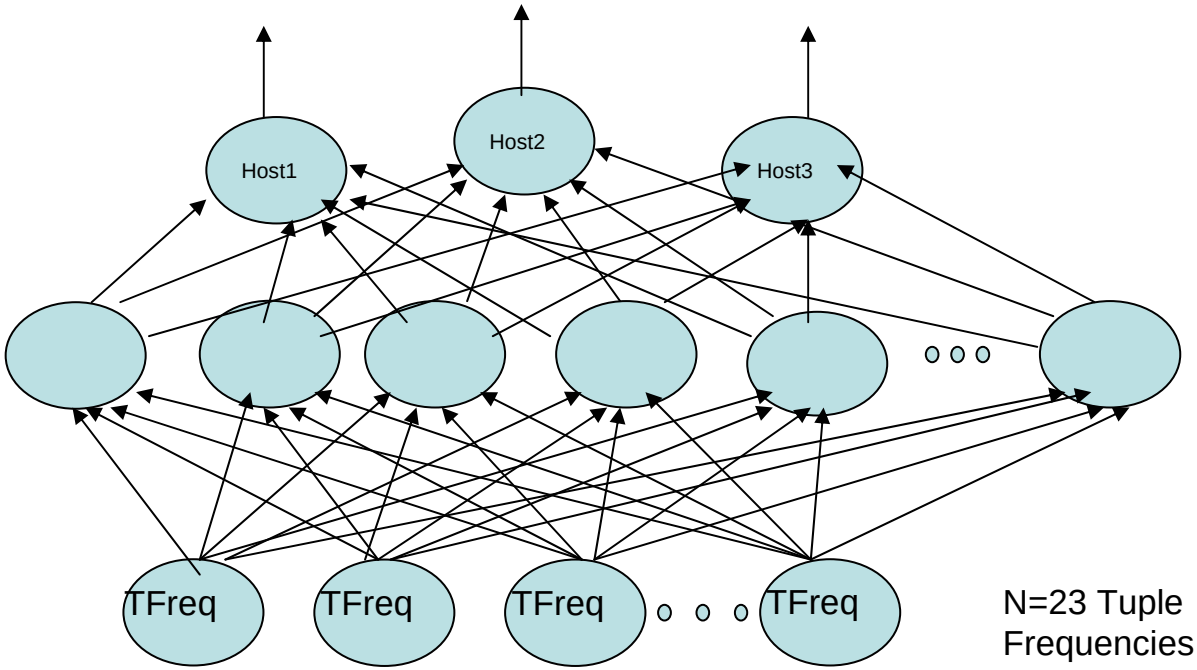
Port Number Range	# of Ports Accessed	%of Ports Accessed	%of Total Bytes Traffic	# of Ports Accessed	%of Ports Accessed	%of Total Bytes Traffic
<1024	< 7	20-50%	<1%	45	0.07%	
1024-5000	< 10	>30%	>90%	3,976	6%	1%
>5000	< 5	<20%	<9%	60,059	93%	99%

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information

- In early 2006 Neural Network was used to classify workstation traffic based on a localized “Workstation Genome”.
- It was found workstation behaviour could be fully described by a set of 23 unique 3-tuples formed by the combination of Protocol, Destination Port, and Byte Range ID – Where Byte Range ID was one of five levels given by:

Bytes	Range
0 – 100	1
100 – 999	2
1000 – 9,999	3
10,000 – 49,999	4
50,000 +	5

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information



Each input frequency vector contains an observed frequency for each 3-tuple for a 24 hour period.
Each 3-tuple is defined as Protocol, Destination Port, Byte Range.
All observed Workstations could be described by a 23 element Vector.

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information

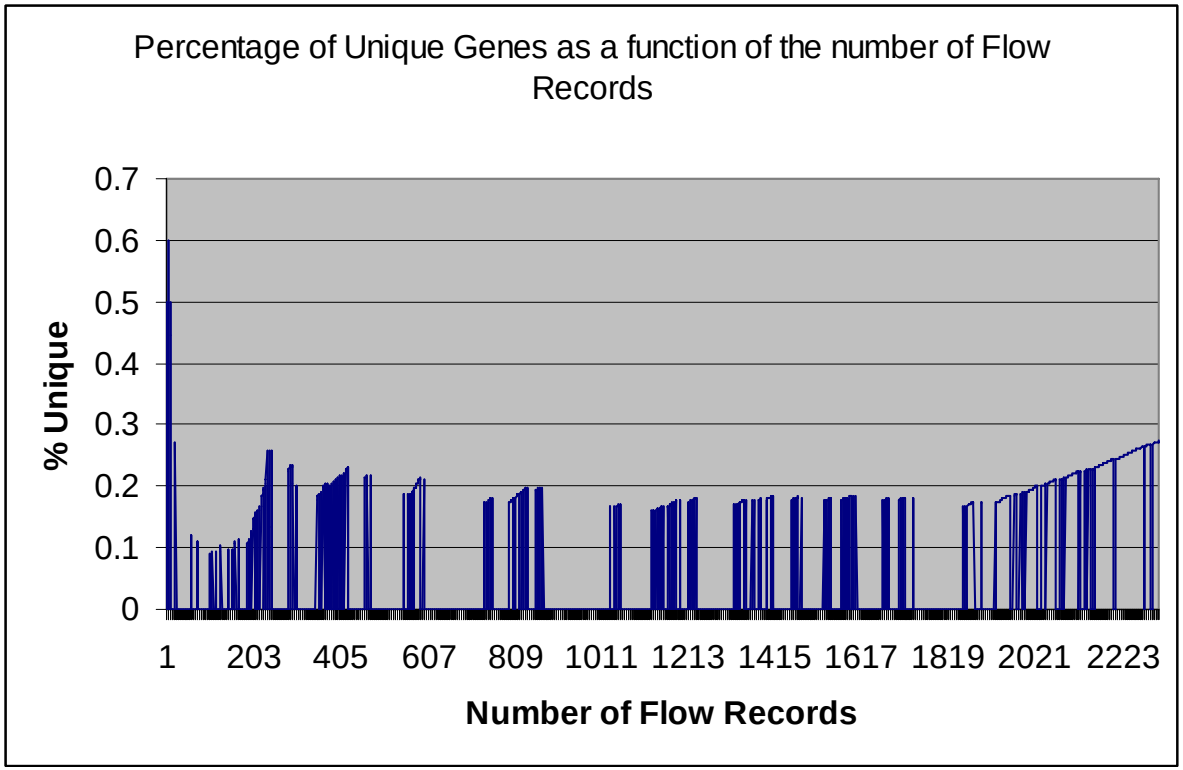
Host ID	Day	Output Vector	Classification (Hit/Miss/Unknown
1 [0 1 0]	1	[0.04 0.86 0.08]	HIT
	2	[0.17 0.97 0.00]	HIT
	3	[0.10 0.91 0.02]	HIT
	4	[0.09 0.95 0.01]	HIT
2 [1 0 0]	1	[0.95 0.06 0.00]	HIT
	2	[0.96 0.04 0.00]	HIT
	3	[0.95 0.06 0.00]	HIT
	4	[0.95 0.07 0.00]	HIT
3 [0 0 1]	1	[0.00 0.09 0.92]	HIT
	2	[0.00 0.00 0.99]	HIT
	3	[0.00 0.12 0.92]	HIT
	4	[0.00 0.00 0.99]	HIT

100% Success rate on uniquely classifying a small sample of the population

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information

- In early 2007 a similar population of workstations was chosen with the goal of testing a Support Vector Machine approach to classification.
- *To the great surprise of the author, the number of unique 3-tuples required to uniquely describe the Workstation Genome had risen from 23 to over 600 in 16 months.*
- Subsequent investigation showed that the diversity of the observed behaviour increased as a function of both population size as well as the length of the sampling period.

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information



By limiting the traffic to ICMP and TCP flow records, the number of unique tuples required to adequately describe the population reached a steady state of approximately 18% of the total number of all expressed tuples.

When UDP traffic was introduced into the sample, the percentage of unique tuples in the population did *not* reach a steady state in proportionality but rather the number of the unique tuples increased in linear proportion to the number of total tuples observed.

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information

- What happened to the network traffic to create such diversity in such a short period of time?
- It is not yet possible to accurately comment on the nature of the change in traffic volume.
- Two fundamental behaviours changed.
  - Protocol Ratio
    - From TCP 70% UDP 30%
    - To TCP 50% UDP 50%
  - Use of Unique Destination Ports by Workstations now parallels Server behaviour.

One Year of Peer-to-Peer

Much has been written lately of the growth and deployment of Peer-to-Peer Protocols

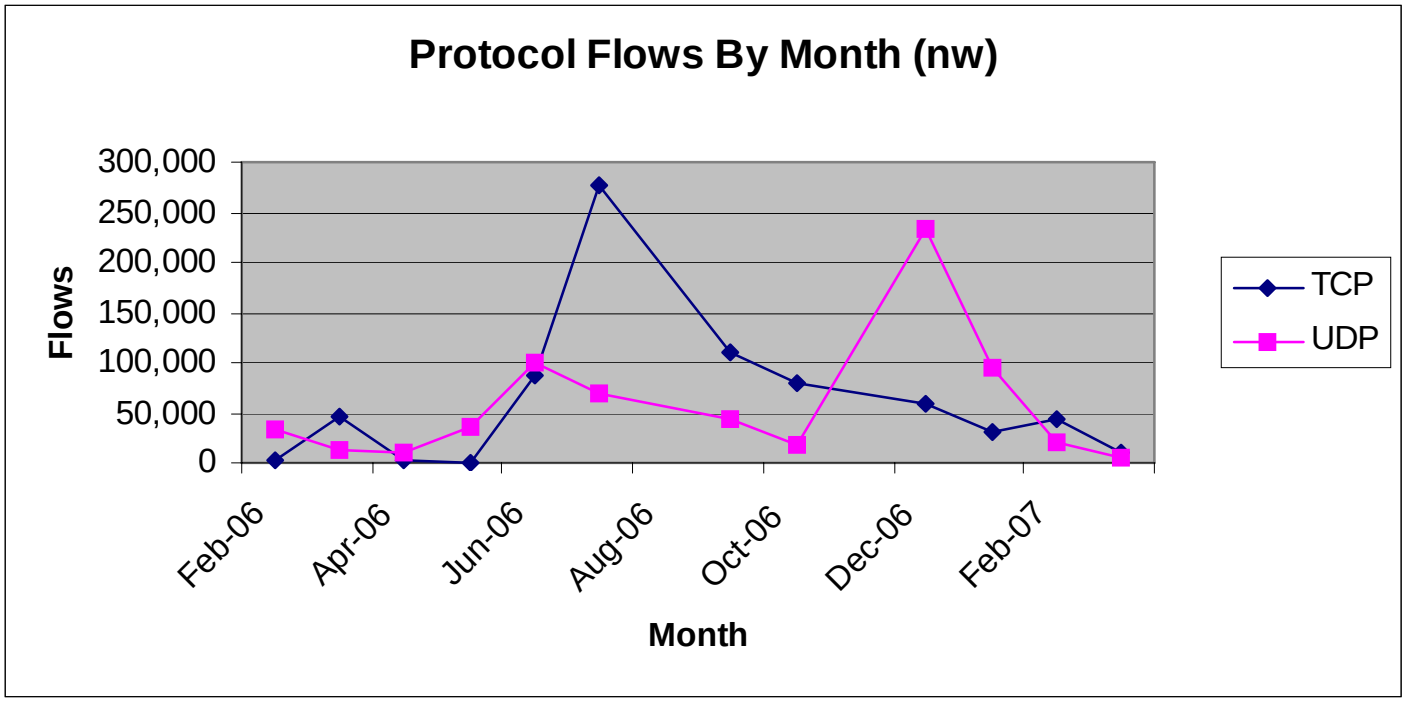
Recommended reading *“Transport Layer Identification of P2P Traffic”, Thomas Karagiannis, et al, IMC’ 04, 2004, Taorimina, Italy.*

Perhaps Peer-to-Peer is the culprit.

Decided to check for the presence of known P2P in the traffic

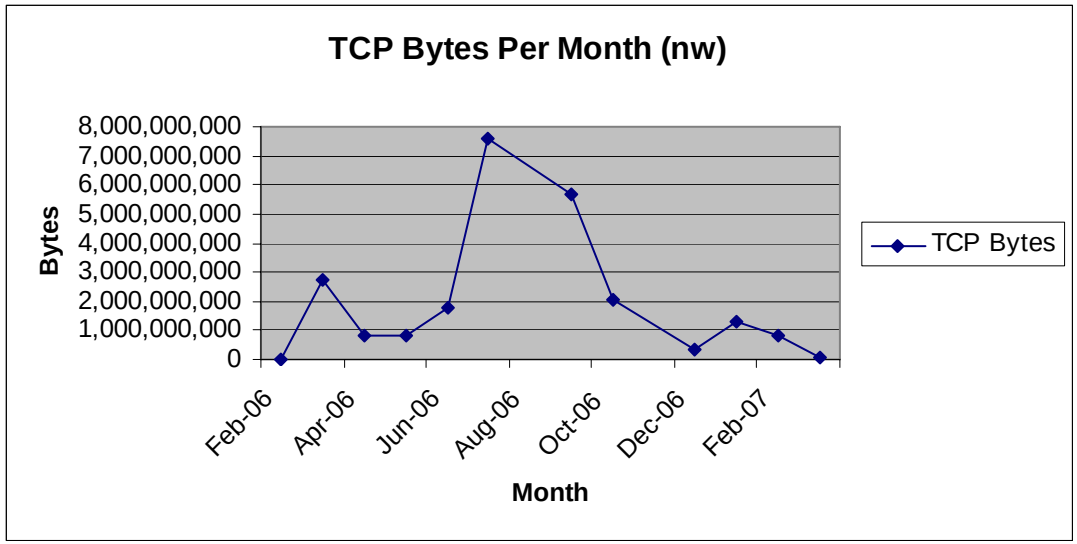
eDonkey2000
Fasttrack
Bittorent
Gnutella
MP2P

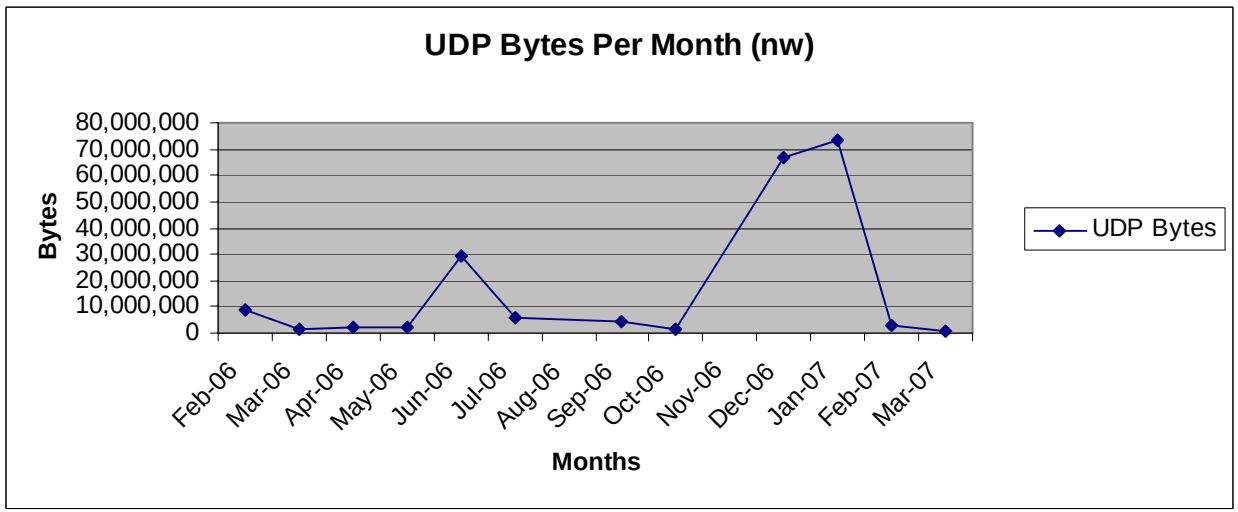
One Year of Peer-to-Peer



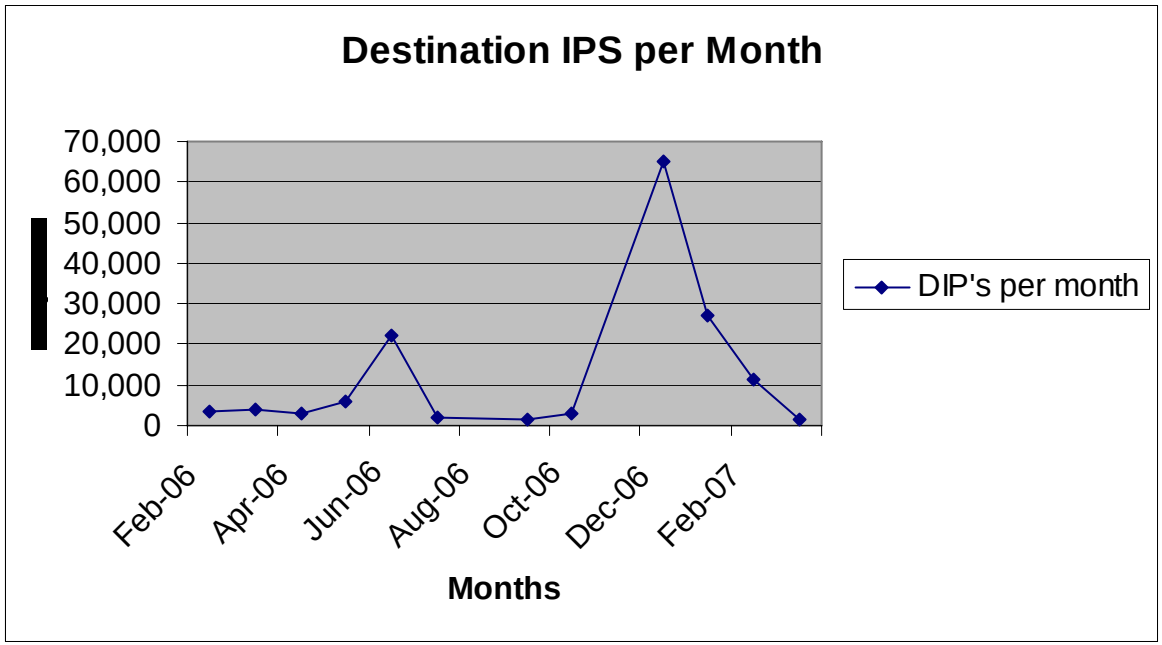
The graph above shows the pattern of flows by protocol for one year for the Target network.

One Year of Peer-to-Peer



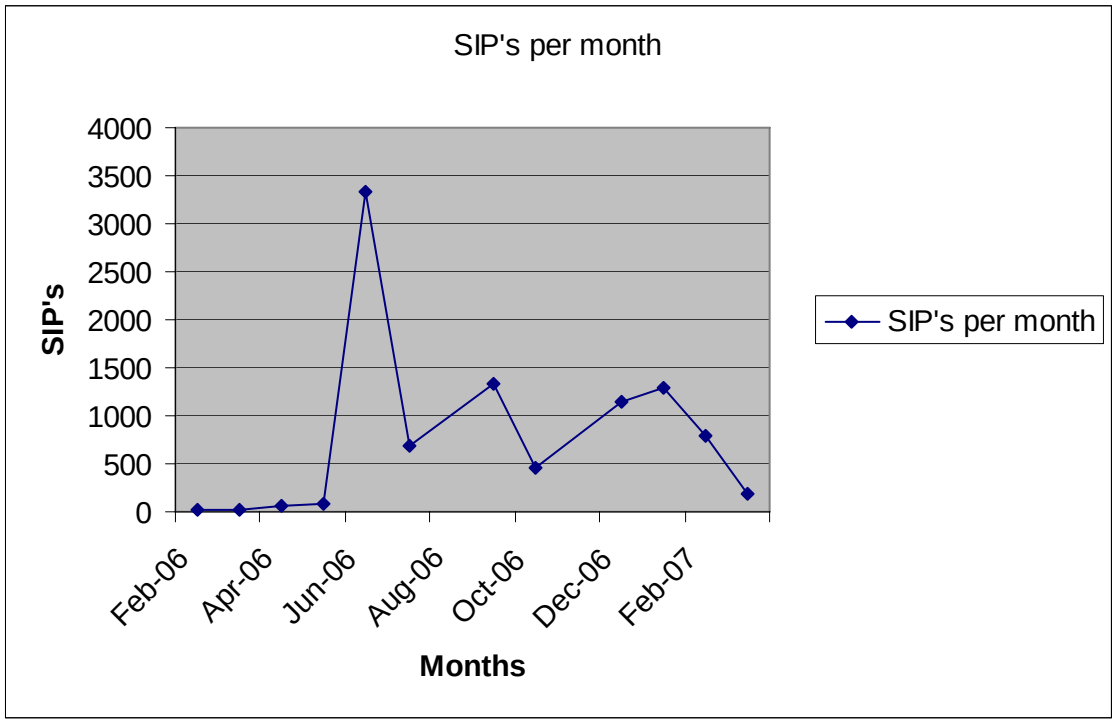


One Year of Peer-to-Peer



For a small network they talked to quite a few friends.

One Year of Peer-to-Peer

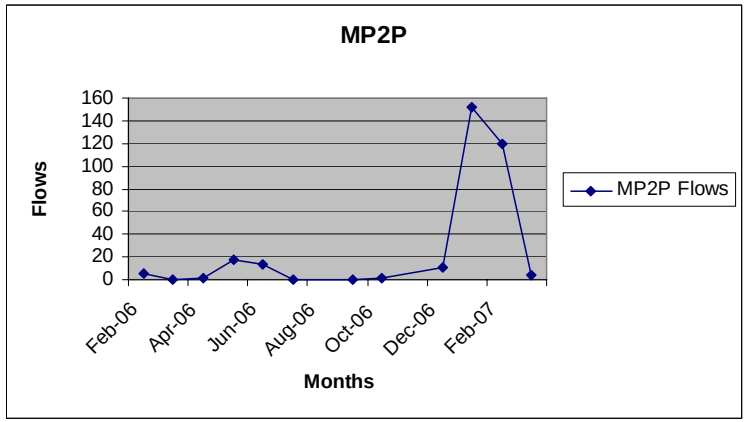


The feeling was mutual.

One Year of Peer-to-Peer

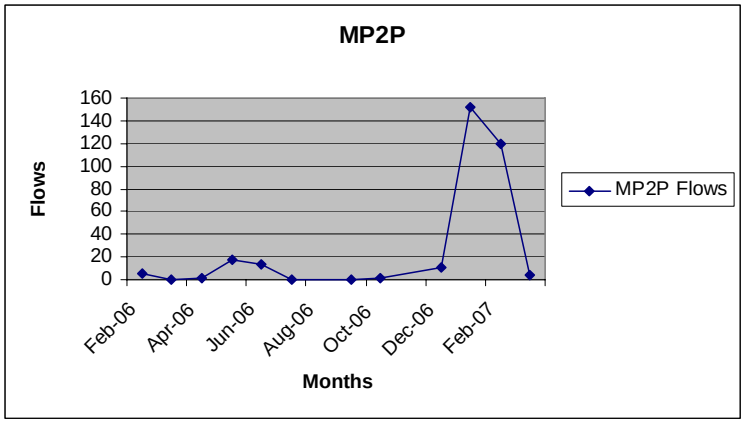
Let’s consider the traffic contribution for each P2P Application in the table.

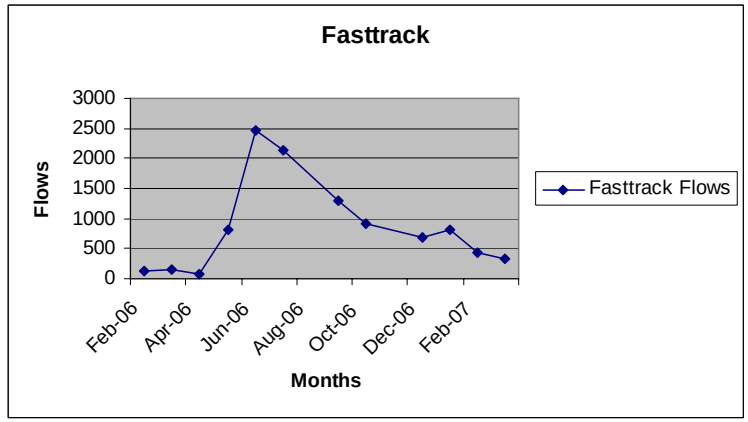
One Year of Peer-to-Peer



MP2P, or Manolito, is a P2P system primarily used to share music files. MP2P traffic was the least contributor to the overall network traffic among the observed systems. This traffic reached a peak flow count of just under 160 in January 2007.

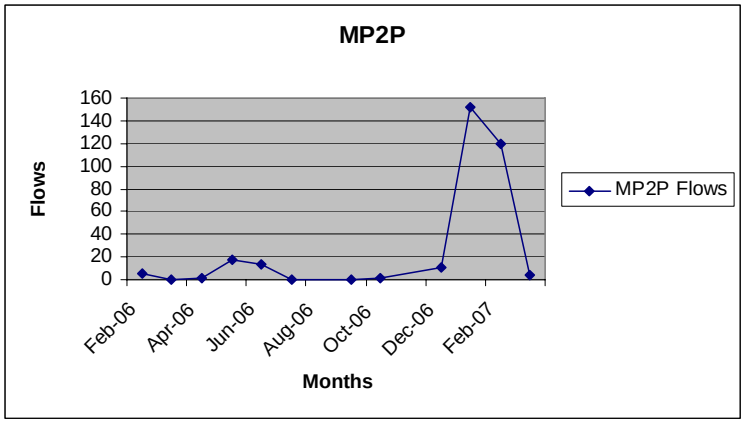
One Year of Peer-to-Peer

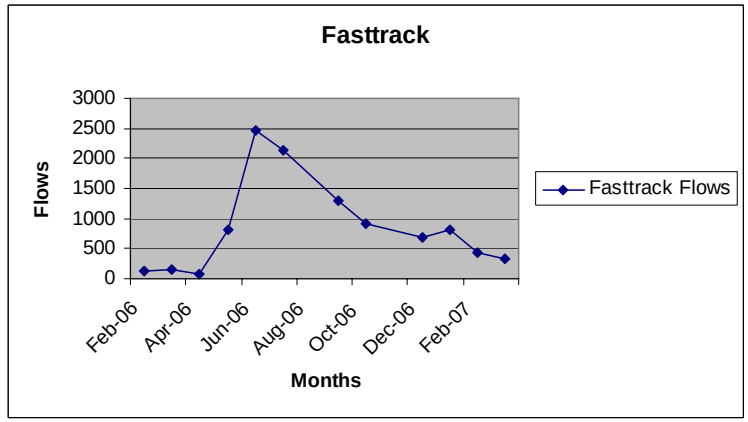


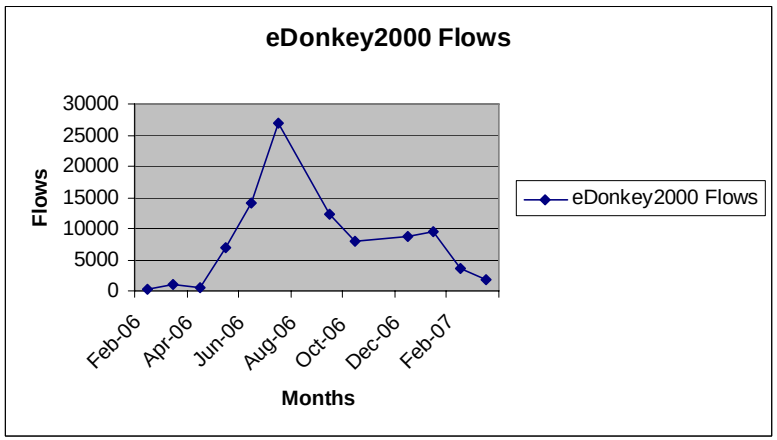


The Fasttrack P2P system is primarily used by Kazaa and its variants to exchange mp3 music files. Fasttrack traffic reached a peak flow count of 2,500 in July 2006.

One Year of Peer-to-Peer

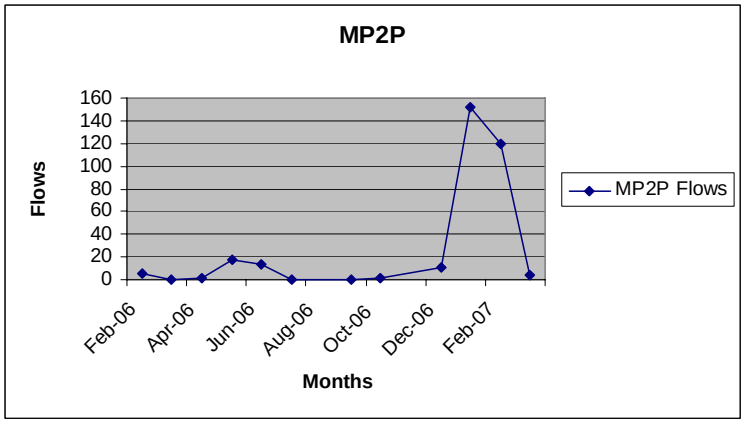


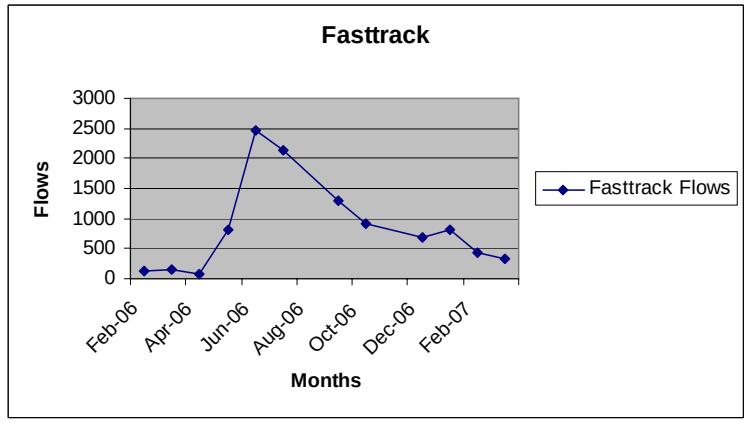




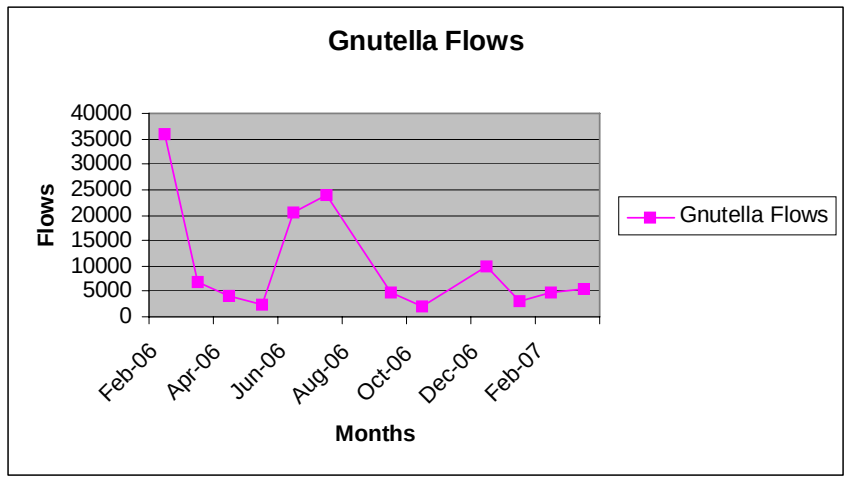
EDonkey2000 was a peer-to-peer system primarily used to distribute large images, video games and software. Although officially discontinued in September 2005 due to legal action brought by the Recording Industry Association of America (RIAA), we speculate, based on our profiling, that we observed eDonkey2000 communication during 2006. EDonkey traffic passed 25,000 flows in July 2006.

One Year of Peer-to-Peer

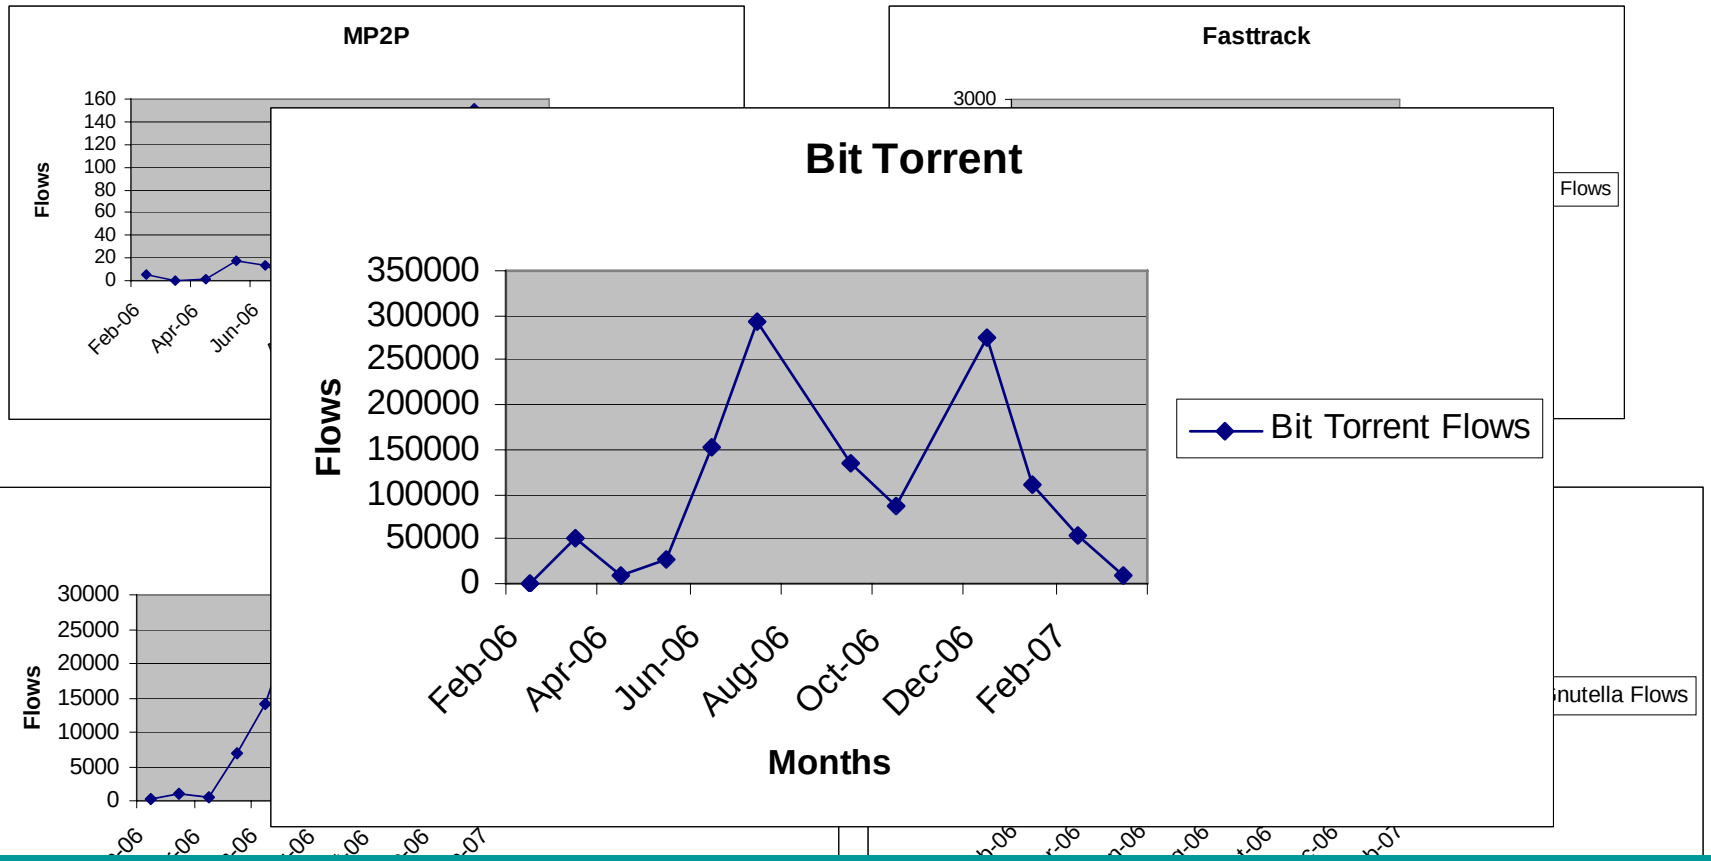




Gnutella is a multi-tier Peer based file exchange system. Traffic from Gnutella ranged from 5,000 to 35,000 flows per month.

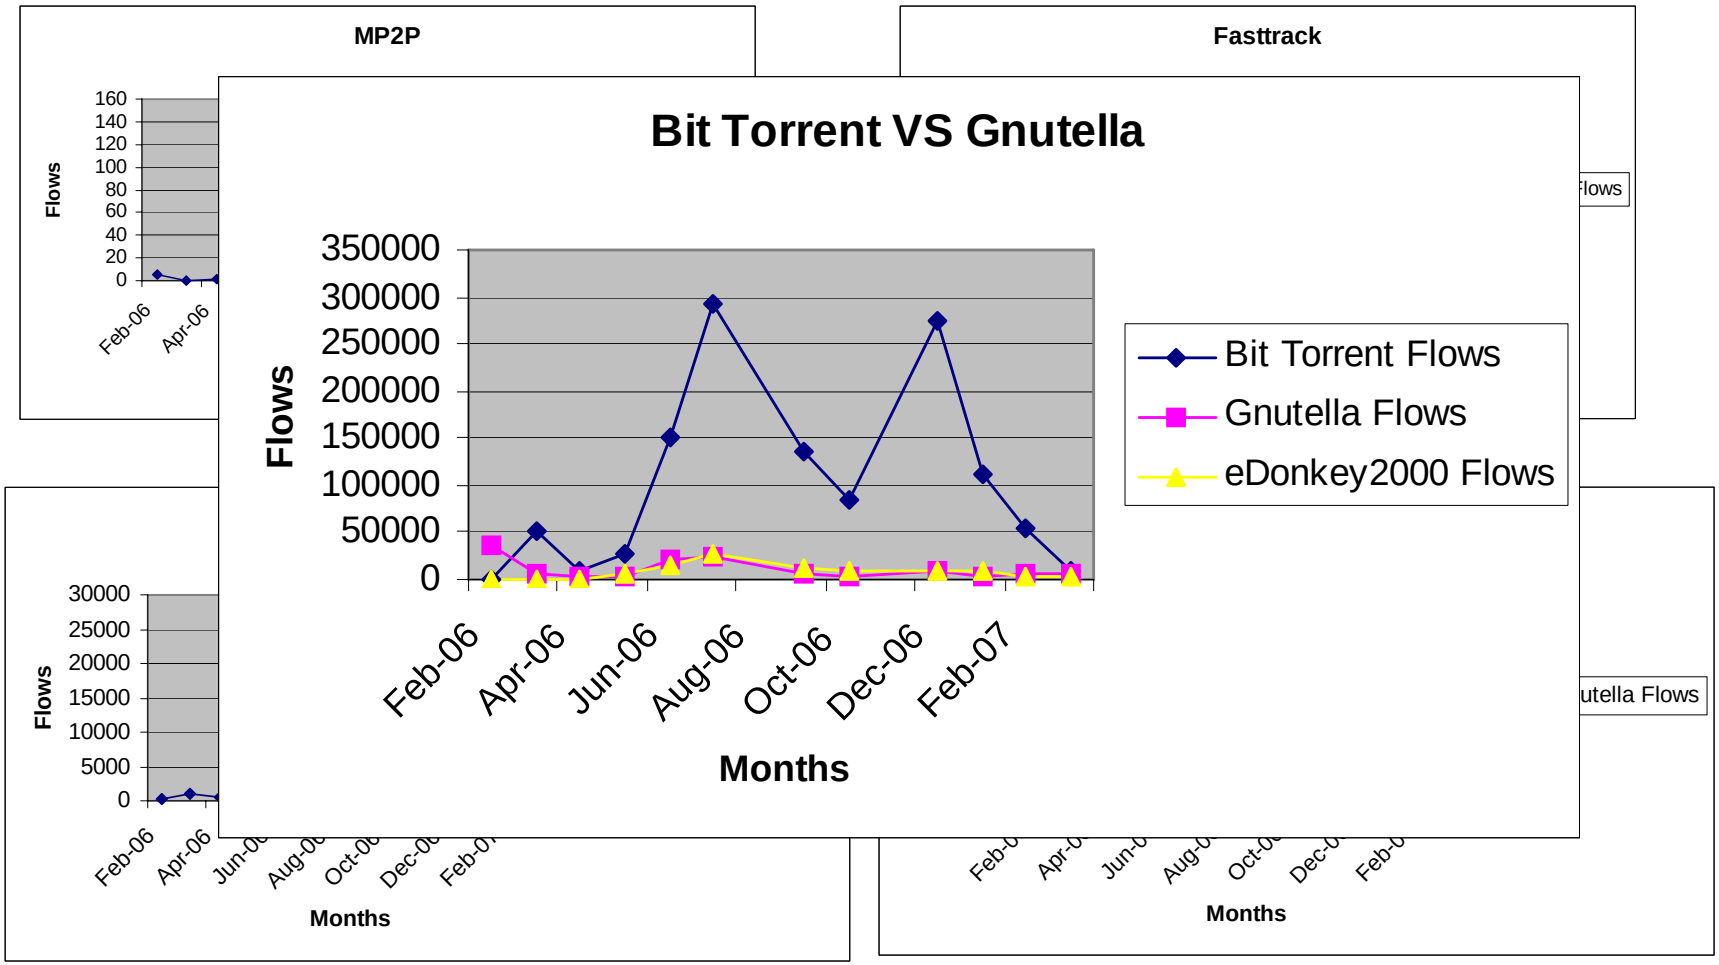


One Year of Peer-to-Peer



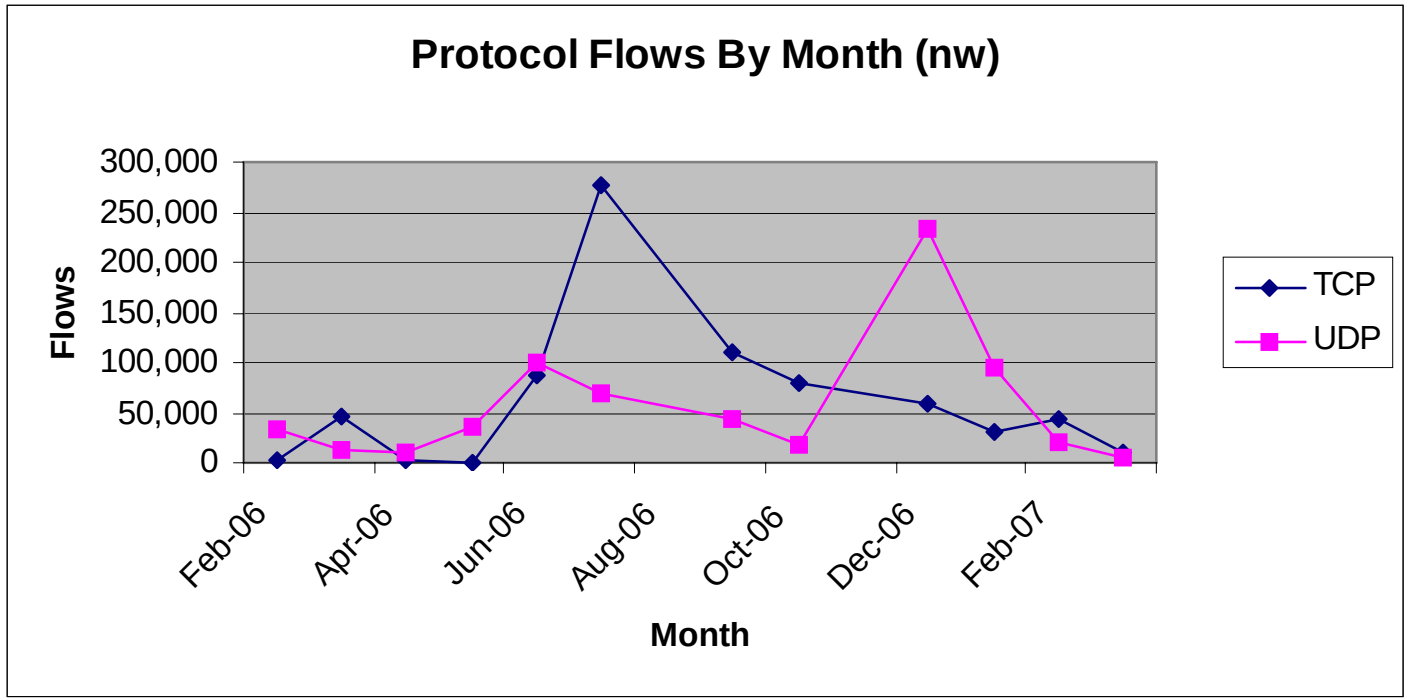
BitTorrent is an ever increasing popular P2P system used for exchanging large data files. Many open source software releases are distributed using BitTorrent. It is also used to distribute legal movie and music downloads. BitTorrent traffic eclipsed most P2P traffic at 300,000 flows.

One Year of Peer-to-Peer

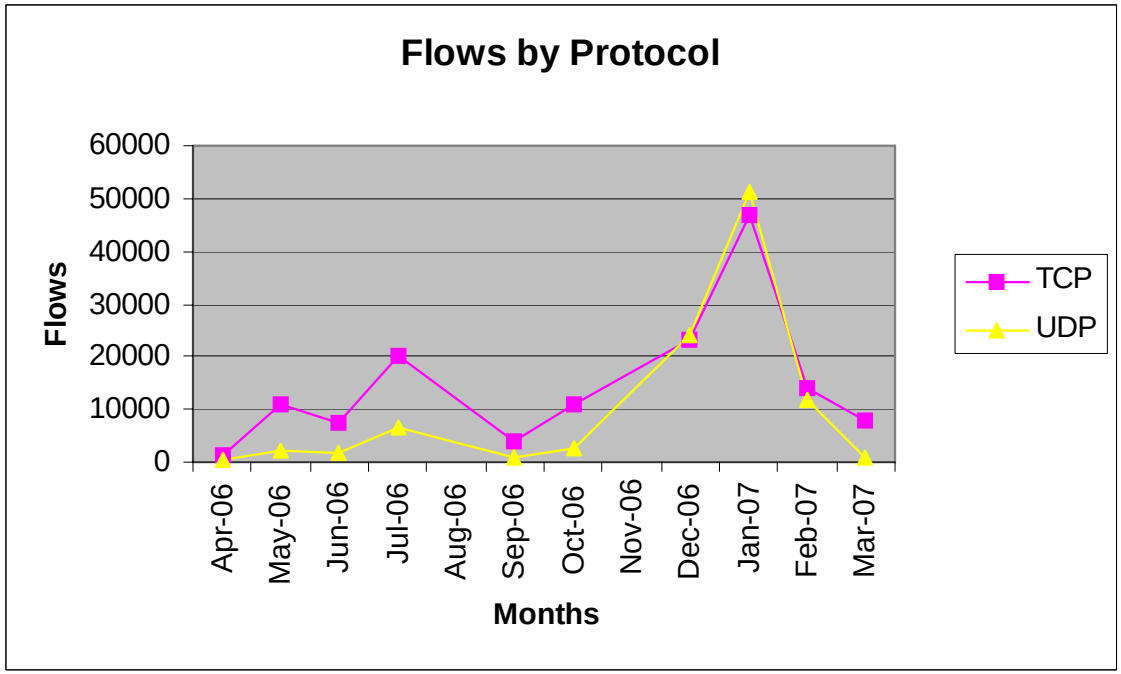


One Year of Peer-to-Peer

Unfortunately the overall Peer-to-Peer flow pattern did not match the pattern that we were seeking. That being a 50/50 ratio of TCP to UDP.

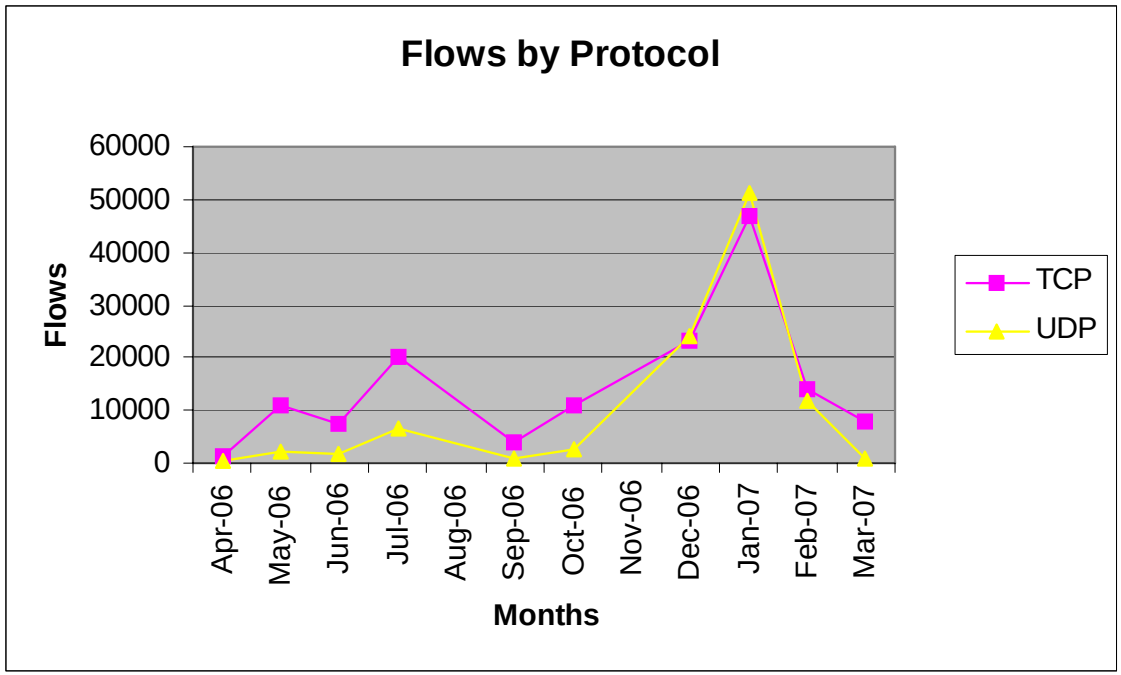


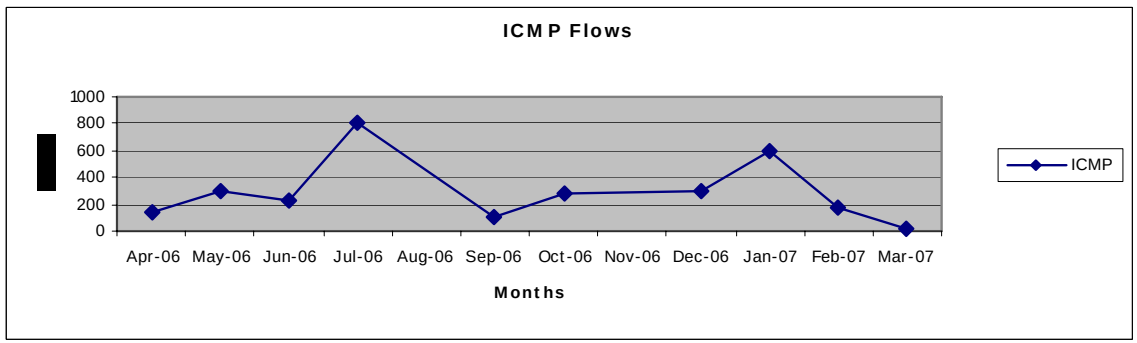
One Year of Peer-to-Peer



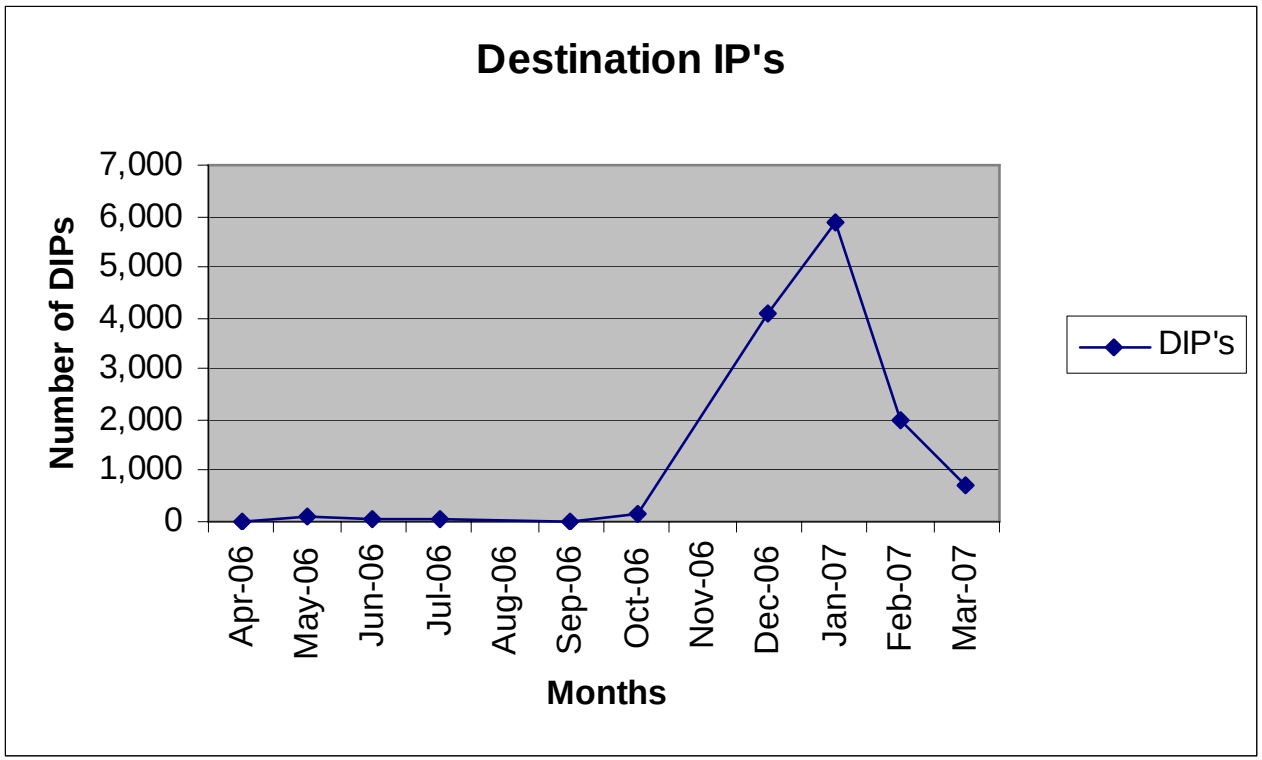
The graph above shows the pattern for which we were searching. This is the traffic from a single user workstation, with a peak flow count of 50,000 flows per month.

One Year of Peer-to-Peer



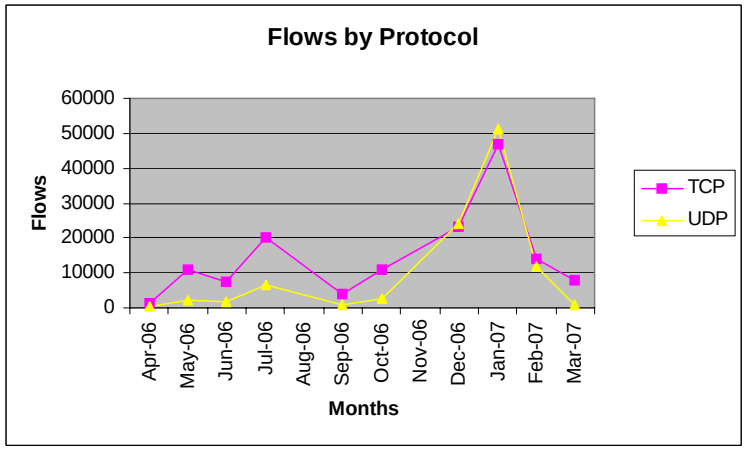


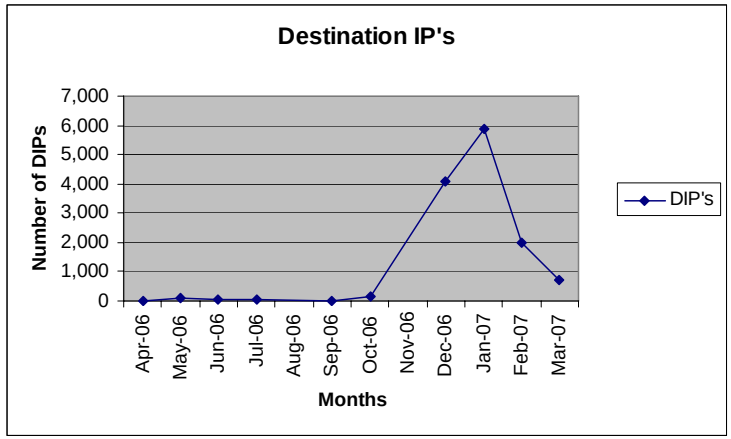
One Year of Peer-to-Peer



This workstation changed its behaviour in late fall 2006 from talking to less than 100 DIPs per month to 6,000 DIPs per month.

One Year of Peer-to-Peer





Who am I ?

One Year of Peer-to-Peer

SKYPE

This traffic pattern is driven by the adoption of Voip by a single user in the target network.

Disclaimer: It is important to point out that since the experimenter had no access to the actual machine or payload data this conclusion is simply conjecture based on known user Behaviour within the target network. (Skype is a wonderful App)

Observations on Traffic for Clients and Peers

- Consumes considerable Resources.
- Represents an Application Level WAN Network for Communication.
- Provides a channel to hide Malicious Activity.

"*McAfee suggested hackers were likely to create malicious software to target instant messaging services, Voice over Internet Protocol (VoIP) telephony services and online gaming sites.*" **Hackers will target social networking sites: security firms - Thursday, November29, 2007, CBC News http://www.cbc.ca**

Evidence that all is not as it Appears

- One day in February a conversation took place between a user host on the Network and a host compromised by an on-line game server.
- Two hours later the user host was attempting to contact a few friends….

Sequentially….

Destination IP	sPort	dPort	Proto	bytes
XXX.XXX.026.000	0	2048	1	56
XXX.XXX.026.000	0	2048	1	168
XXX.XXX.026.001	0	2048	1	56
XXX.XXX.026.001	0	2048	1	168
XXX.XXX.026.002	0	2048	1	56
XXX.XXX.026.002	0	2048	1	168
XXX.XXX.026.003	0	2048	1	56
XXX.XXX.026.003	0	2048	1	168
XXX.XXX.026.004	0	2048	1	56
XXX.XXX.026.004	0	2048	1	168
XXX.XXX.026.005	0	2048	1	56
XXX.XXX.026.005	0	2048	1	168
XXX.XXX.026.006	0	2048	1	56
XXX.XXX.026.006	0	2048	1	168
XXX.XXX.026.007	0	2048	1	56
XXX.XXX.026.007	0	2048	1	168
XXX.XXX.026.008	0	2048	1	56
XXX.XXX.026.008	0	2048	1	168
XXX.XXX.026.009	0	2048	1	56
XXX.XXX.026.009	0	2048	1	168
XXX.XXX.026.010	0	2048	1	56
XXX.XXX.026.010	0	2048	1	168
XXX.XXX.026.011	0	2048	1	56
XXX.XXX.026.011	0	2048	1	168
XXX.XXX.026.012	0	2048	1	56
XXX.XXX.026.012	0	2048	1	168
XXX.XXX.026.013	0	2048	1	56
XXX.XXX.026.013	0	2048	1	168
XXX.XXX.026.014	0	2048	1	56
XXX.XXX.026.014	0	2048	1	168
XXX.XXX.026.015	0	2048	1	56
XXX.XXX.026.015	0	2048	1	168
XXX.XXX.026.016	0	2048	1	56

We Need to Re-Consider our Willingness to be a Peer

- Users willingly download and install client/peer/server software.
- They even participate in strategies to avoid barriers and impediments (like Nat'ing).
- There is an implied trust that the communication is exclusively what it claims to be.
- "When they thought they were playing at war craft, they were actually playing at war craft."

Concluding Notes

- The network is evolving at the edges
- This means that network architectures, management and provisioning strategies are now more responsive then ever.
- Global communication resources are primarily influenced by the uncoordinated activities of individuals.
- Traffic patterns are emergent properties without intent.

Future Work

- Study the growth in diversity of patterns in traffic.
- Study the form and distribution of applications and participants.
- Track Unidentified Anomalies.

- February 2008, TARA will announce the InTARA project

Intelligent Network Traffic Analyzers for Reconstructive and Real Time Analysis

- InTARA will be a multi-million dollar, multi-year project to develop intelligent traffic analysis capabilities for the good guys.
- We are seeking global collaborative research and commercialization partners. Early stage interest from Australia, India, Switzerland, Canada.

Identifying Anomalous Traffic Using Delta Traffic

Tsuyoshi KONDOH and Keisuke ISHIBASHI
Information Sharing Platform Labs.
NTT

Flocon2008, January 7–10, 2008, Savannah GA

Outline

- Background and Motivation
  - Identifying anomalous traffic is the missing piece.
- Our Technique: DELTAA
  - Concepts
    1. Extract anomalous traffic as the delta of normal and anomalous time periods.
    2. Auto-aggregate extracted anomalous traffic.
  - Operation of our technique
    - How to implement the above concepts.
- Evaluation
  - Evaluation using synthesized DDoS traffic.
- Summary

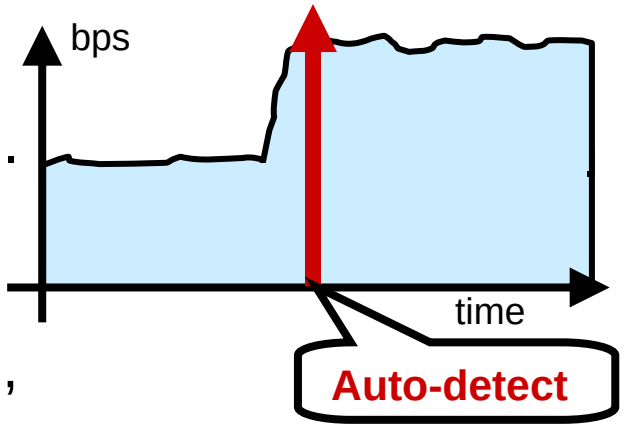
Copyright NTT Information Sharing Platform Laboratories, 2007

Background and Motivation

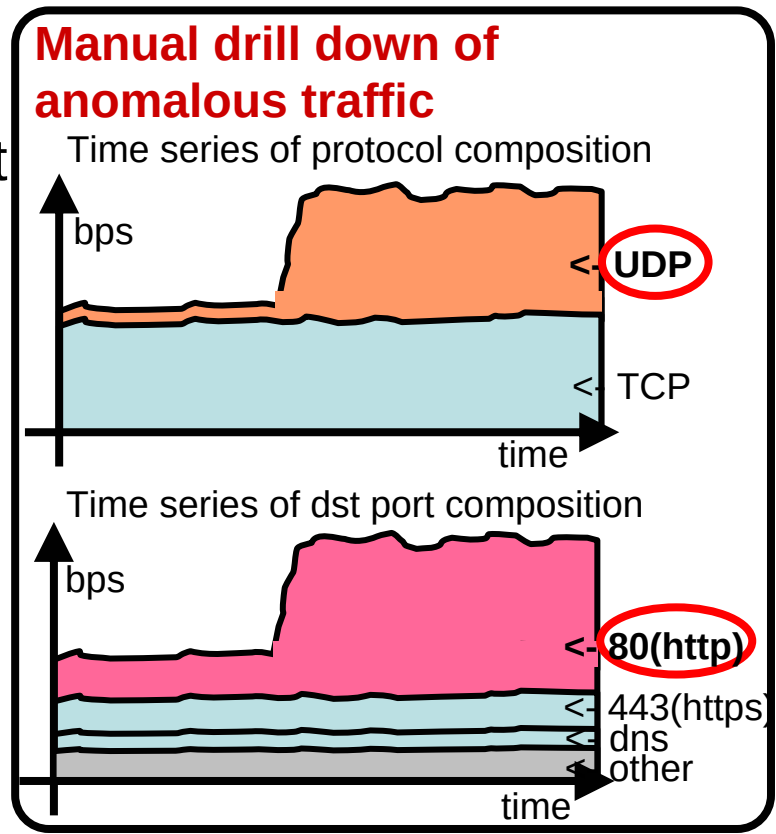
- Monitoring of traffic volumes is widely used for network operation (e.g. MRTG).
- Many techniques for detecting anomalous volume change have been proposed (NBAD, Holt-winters in MRTG, … etc.).
- Some tools to mitigate damage from anomalous traffic. (e.g. drop/rate limit at router, detour to Cisco Guard, etc.)

- However, accurate mitigation needs accurate ACL sets.
- Generating accurate ACL sets requires manual drill down by operator.
  - **Too costly.**

Time series of total traffic by bps



Manual drill down of anomalous traffic



Our Technique: DELTAA

- **DELTAA** outputs ACL sets for filtering or rate limiting to mitigate the damage from anomalous traffic.
  - DELTAA: Delta Traffic Automatic Aggregator

Today, I will focus on two concepts

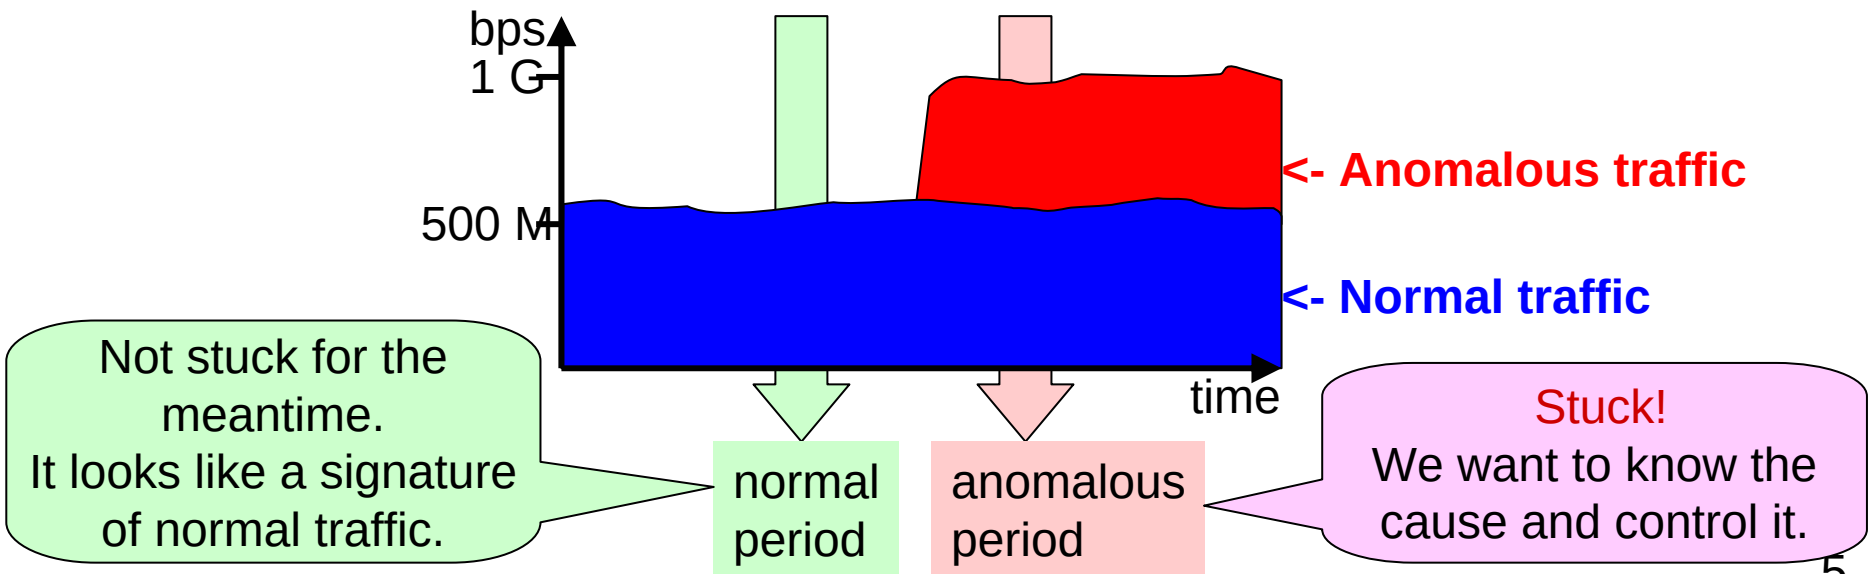
- Three concepts of DELTAA:
  1. Reveal anomalous traffic using delta traffic.
  2. Aggregate delta traffic and generate optimized ACL sets on a single dimension (e.g. source IP address dimension).
  3. Generate multi-dimensional ACL sets by integrating each dimensional anomalous traffic range.

Copyright NTT Information Sharing Platform Laboratories, 2007

Concept #1:
(1) Definition of "Normal" and "Anomalous" Traffic

Throughout this presentation, I use the following definitions.

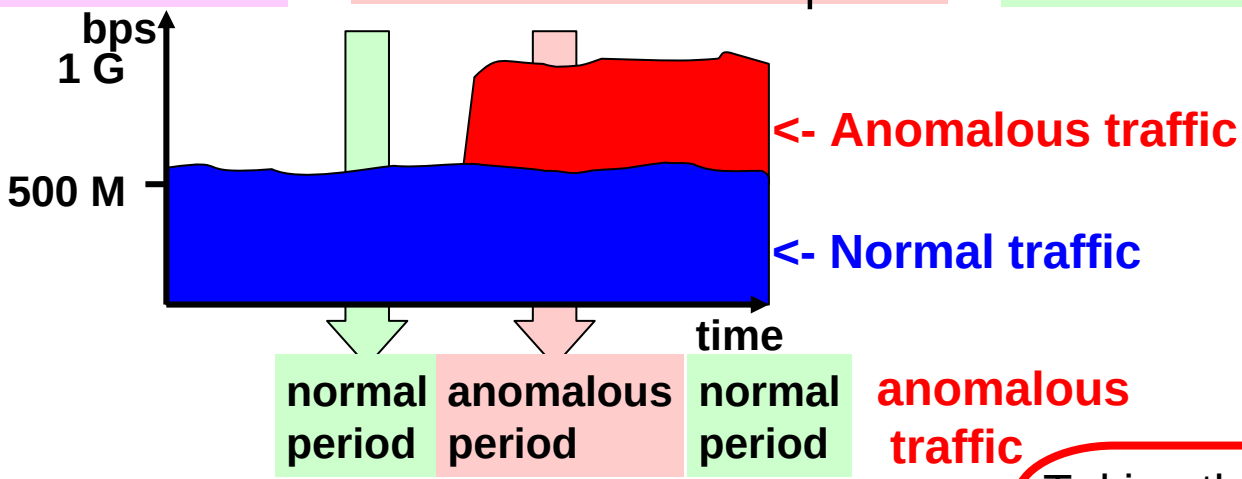
- Anomalous traffic: Traffic that causes a change in traffic volume (bps/pps/fps).
  - BitTorrent and server intrusion are out of scope because they always exist or do not cause a volume change.
- Normal period: Period when traffic volume is normal.
- Anomalous period: Period when traffic volume is anomalous.



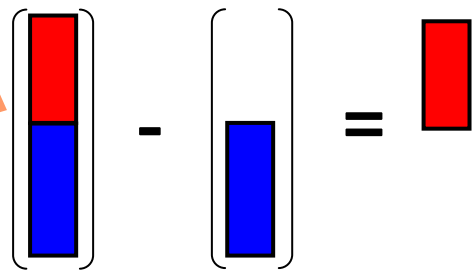
Concept #1 :
(2) Reveal Anomalous Traffic

- Make two assumptions
  1. traffic of normal period = normal traffic
  2. traffic of anomalous period = normal traffic + anomalous traffic
- We can then extract anomalous traffic as the delta of the above two periods.

anomalous traffic = traffic of anomalous period – traffic of normal period



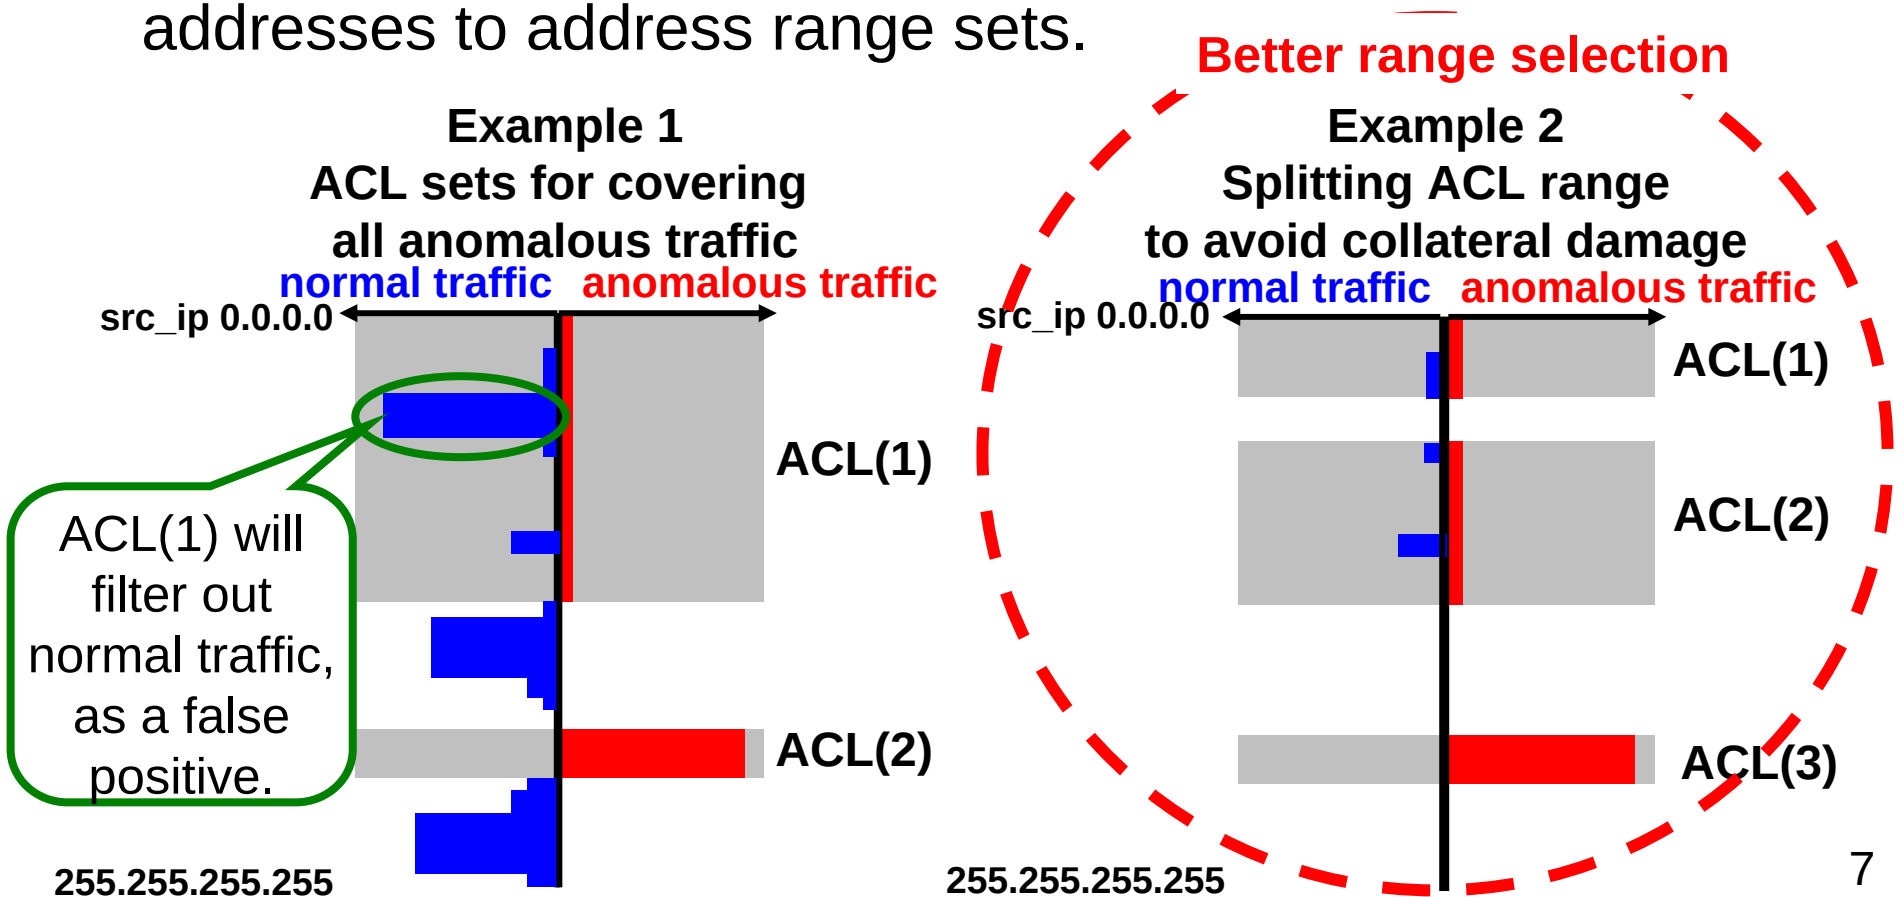
Extracting anomalous traffic from "traffic of anomalous period" is difficult because it is a mixture of normal and anomalous traffic.



Taking the delta between "traffic of normal period" and that of anomalous period, we can effectively extract anomalous traffic.

Concept #2:
Auto-aggregate Delta Traffic

- In aggregation, **optimize a trade-off** (false negative, false positive, number of ACLs) by **using the best range-selection algorithm**.
- Aggregation example: Aggregate from distinct source IP addresses to address range sets.



Explanation of Our Technique

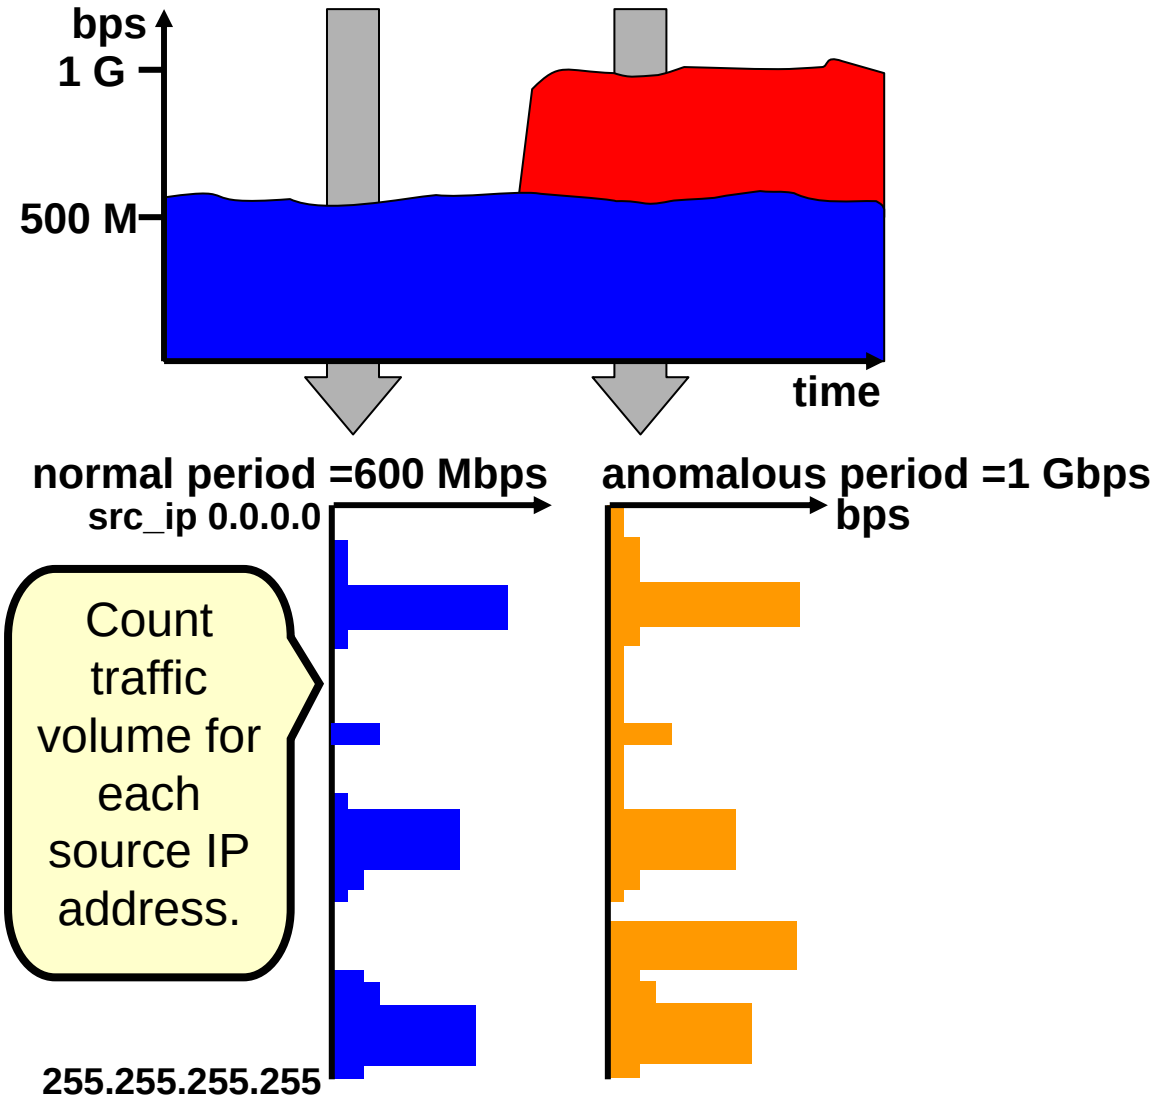
- Our technique can generate multi-dimensional ACL sets.
  - e.g. source/destination IP address, source/destination port, protocol, flow exporter, and router interface
  - Multiple dimensions do not mean independent of above information sets.
  - Our technique merges above information to make multi-dimensional ACL sets.

- In this presentation, I focus on source IP dimension identification as an example and explain step by step.

Copyright NTT Information Sharing Platform Laboratories, 2007

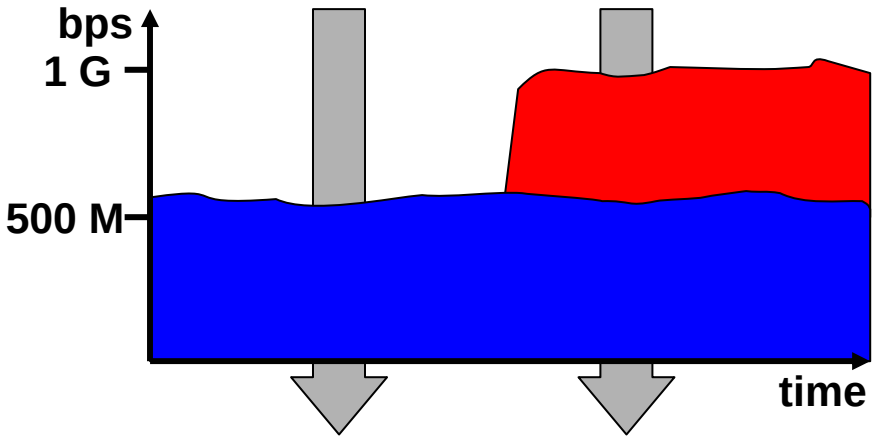
Step 1: (1) Counting Up

Count normal and anomalous periods of traffic for each source IP address.

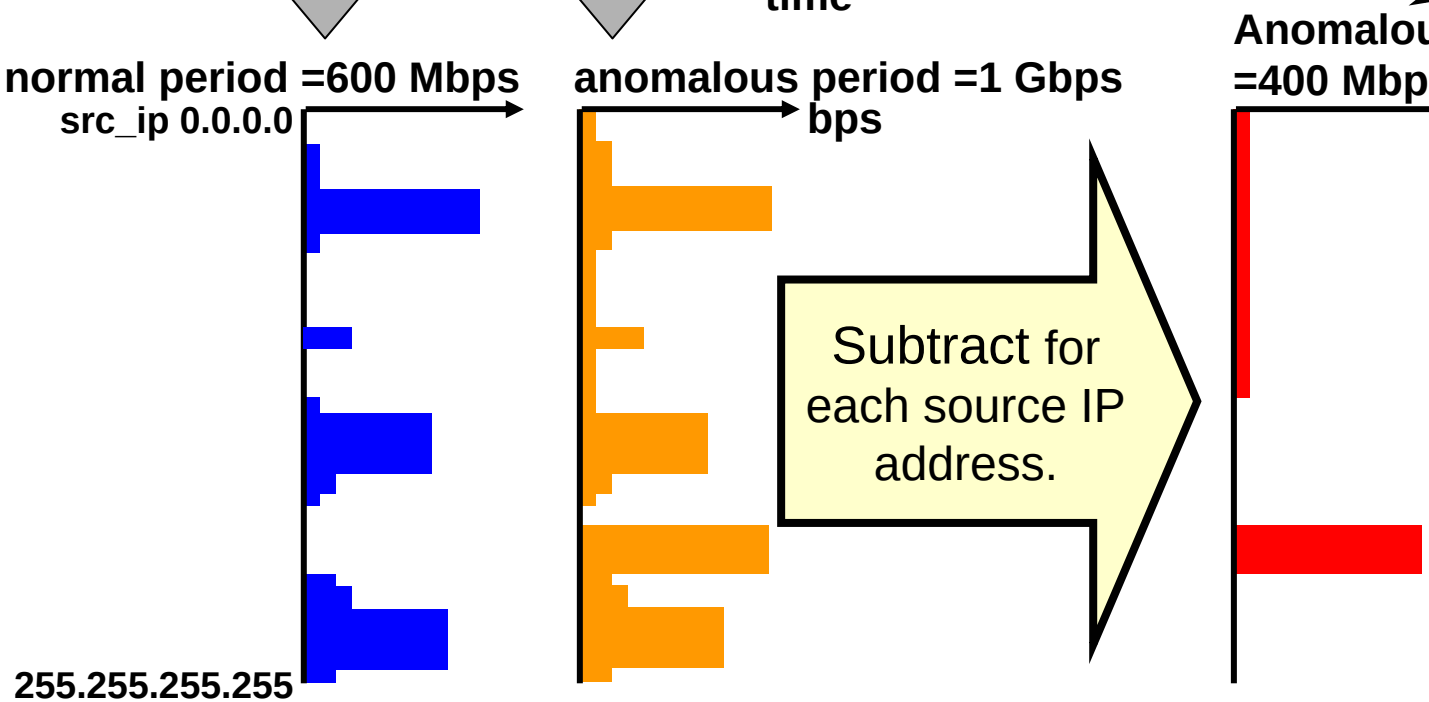


Step 1: (2) Making Delta Traffic

Make delta traffic by subtracting traffic of normal period from that of anomalous period.

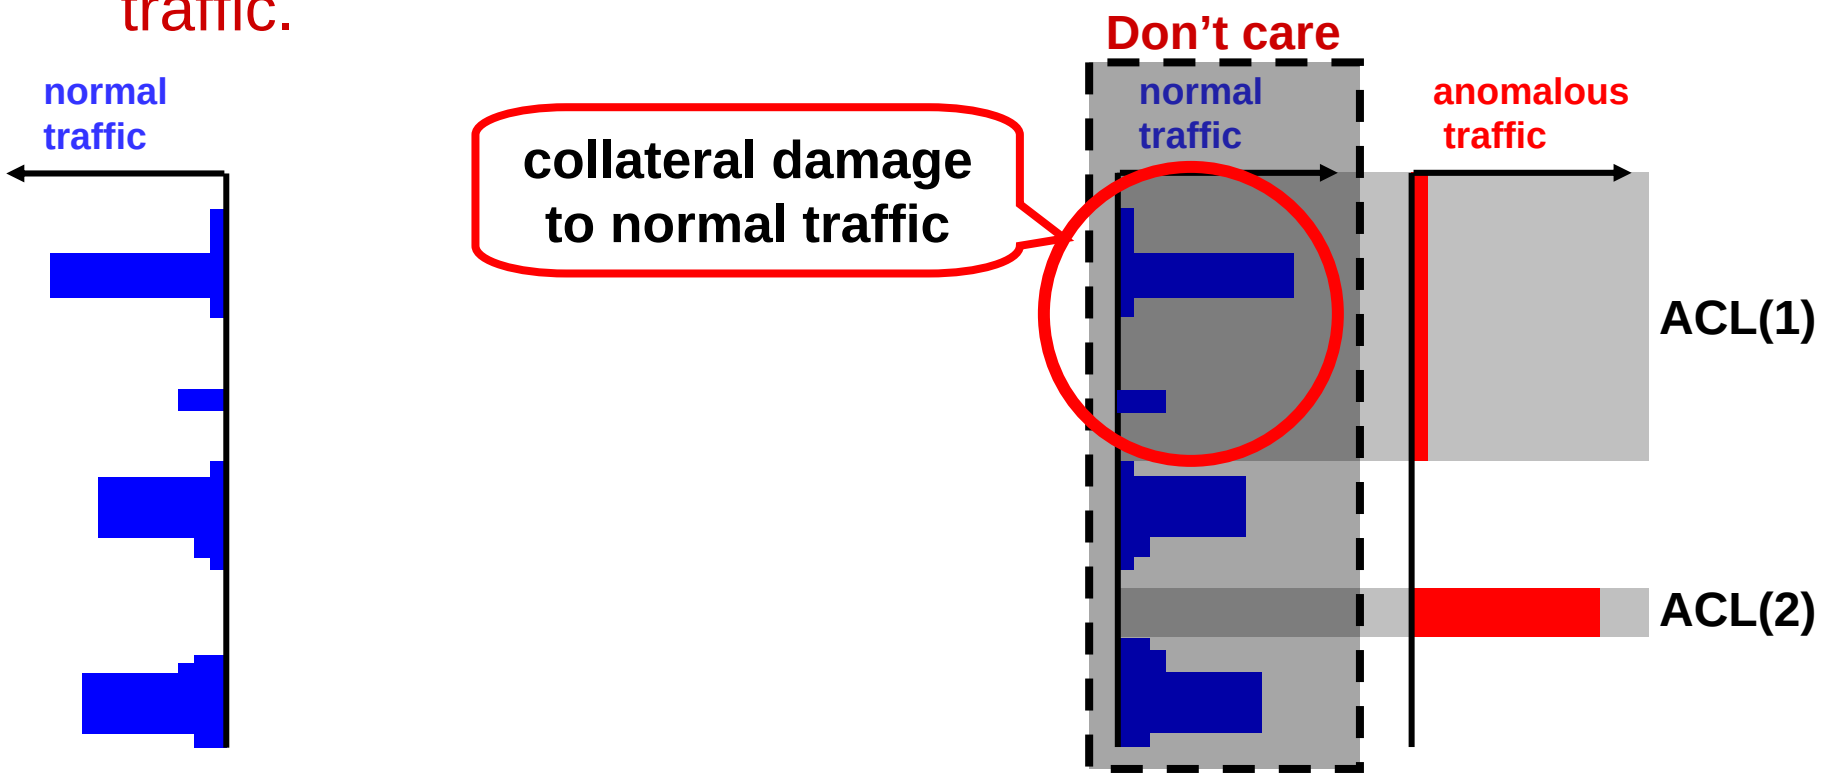


DELTAA obtains anomalous traffic with granularity of source address as delta traffic.



Step 2: (1) Building Tree of Normal and Anomalous Traffic

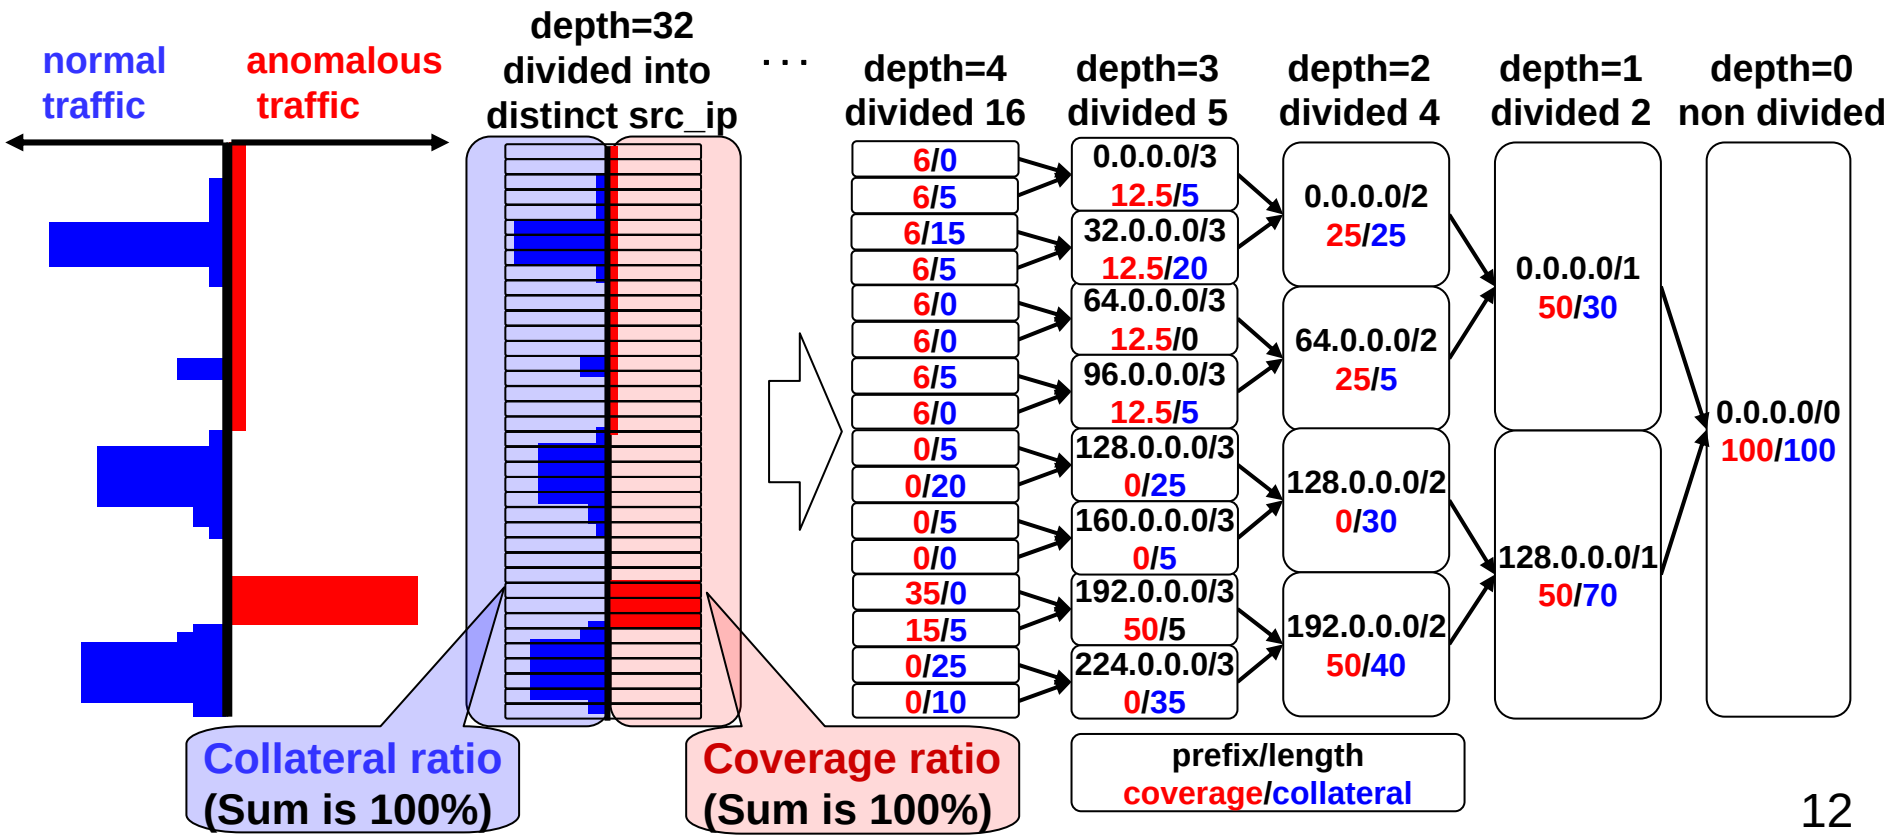
- Example: When we use only anomalous traffic information, collateral damage cannot be avoided.
  - Causes mis-filtering of normal traffic.
- So, build a traffic tree using both normal and anomalous traffic.



Copyright NTT Information Sharing Platform Laboratories, 2007

Step 2: (2) Building Tree of Normal and Anomalous Traffic

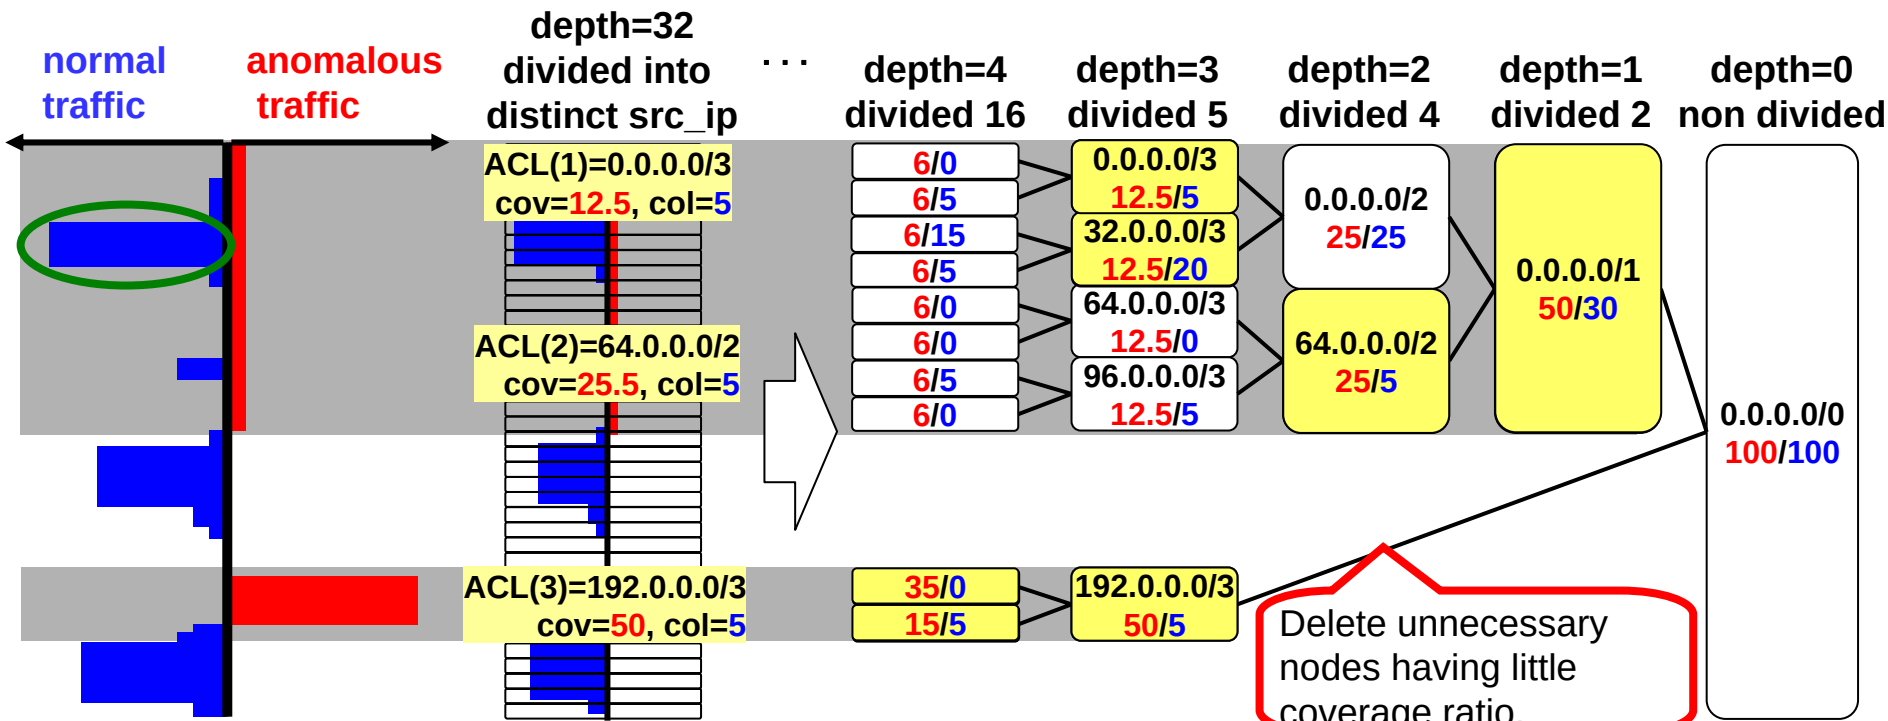
- **Traffic tree making**
  - Build up from individual source IP addresses (depth=32).
  - Each node has information about coverage and collateral ratio.
    - Collateral ratio: normal traffic of the node ÷ total normal traffic
    - Coverage ratio: anomalous traffic of the node ÷ total anomalous traffic
  - Make parent nodes by merging child node information.



Step 3: Selecting Best Node Sets (ACL sets)

1. To reduce search space, delete unnecessary nodes.
   - Unnecessary node: node having little coverage ratio (little anomalous traffic) or little difference from its descendant nodes.
2. Search for best node sets by evaluating goodness of every node combination.
   - Best node combination = **Best ACL sets for source IP dimension**

- But, how to decide goodness of the node sets?

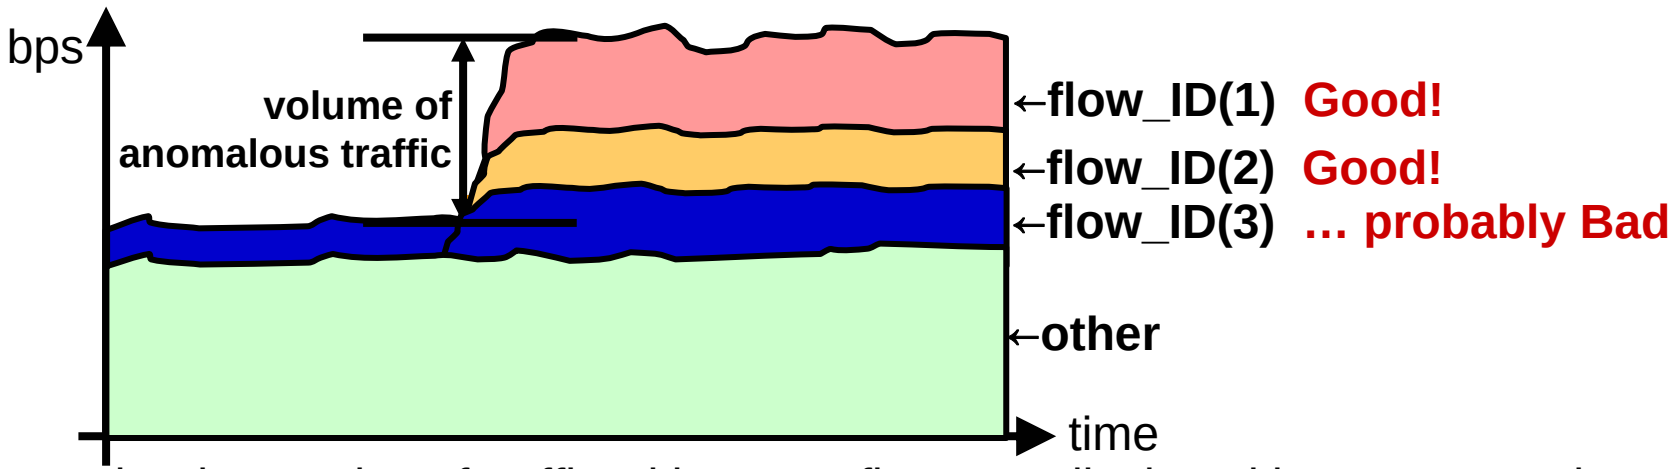


Example 1. Can filter almost all anomalous traffic, but suggested /3

Example 2. Best node sets: Can filter almost all anomalous traffic with little collateral

Criteria of “Goodness”

- Three criteria of identification
  1. Coverage ratio:
     Maximize filtered anomalous traffic = (1 – FNR)
  2. Collateral (damage) ratio:
     Minimize filtered (normal) legitimate traffic = (FPR)
  3. **Number of ACLs**:
     ACL entry budget is limited, so having few ACLs is better.
- But, these three criteria have a trade-off relationship with each other.



Dummy graph: Time series of traffic with output flow_IDs displayed in separate colors

Evaluation Formula for Goodness

- To evaluate goodness of best ACL sets, we use the formula:

  coverage : cov, collateral ratio : $coll$, no. of ACLs : n

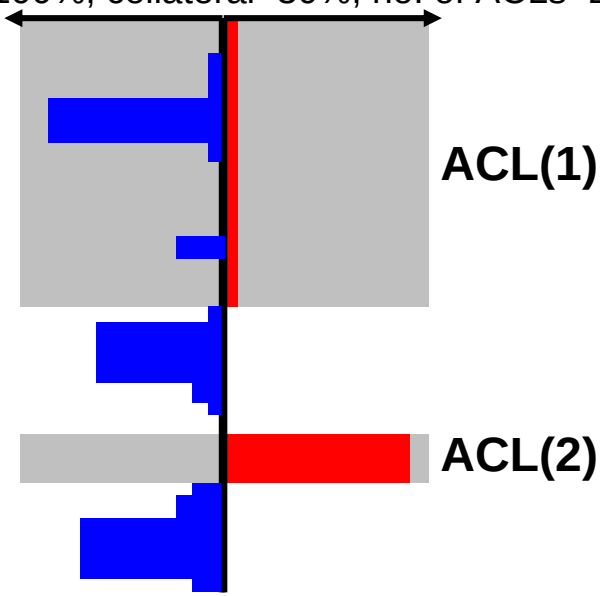
  $$\text{rate} = \frac{(\beta - \alpha) + \alpha \cdot cov - \beta \cdot coll}{n^{\gamma}}$$ (α, β, γ : weighting coefficients)

  - Weighting coefficients can be tuned to reflect network policy or customer requirements.

ACL sets for covering all anomalous traffic

rate= 2.61

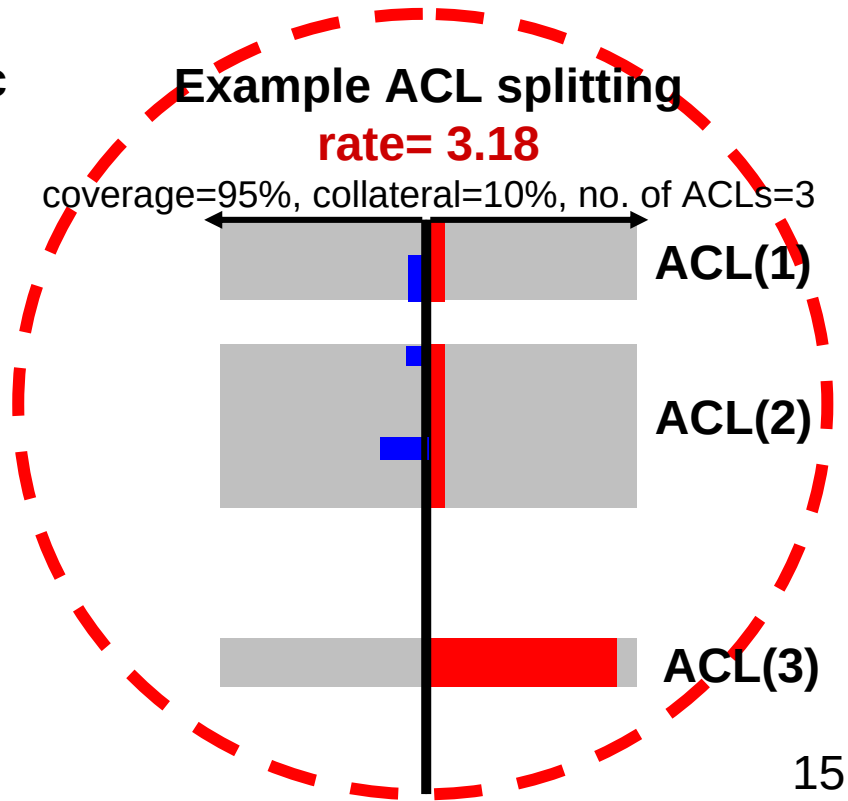
coverage=100%, collateral=30%, no. of ACLs=2



Example ACL splitting

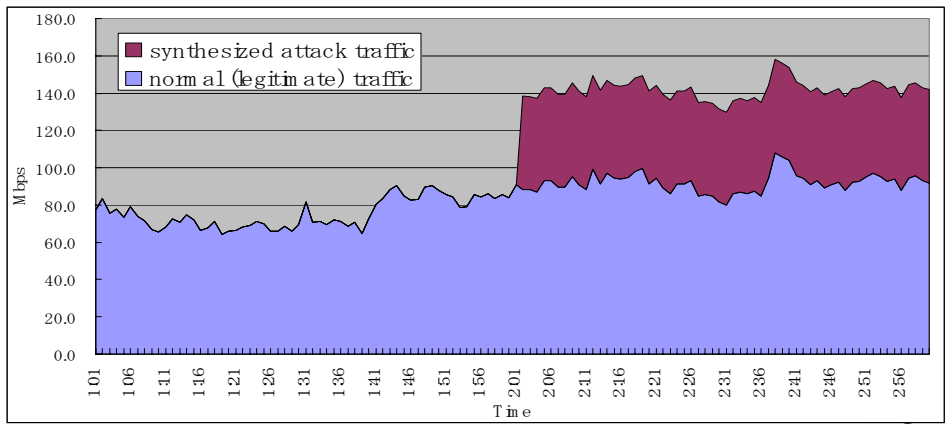
rate= 3.18

coverage=95%, collateral=10%, no. of ACLs=3



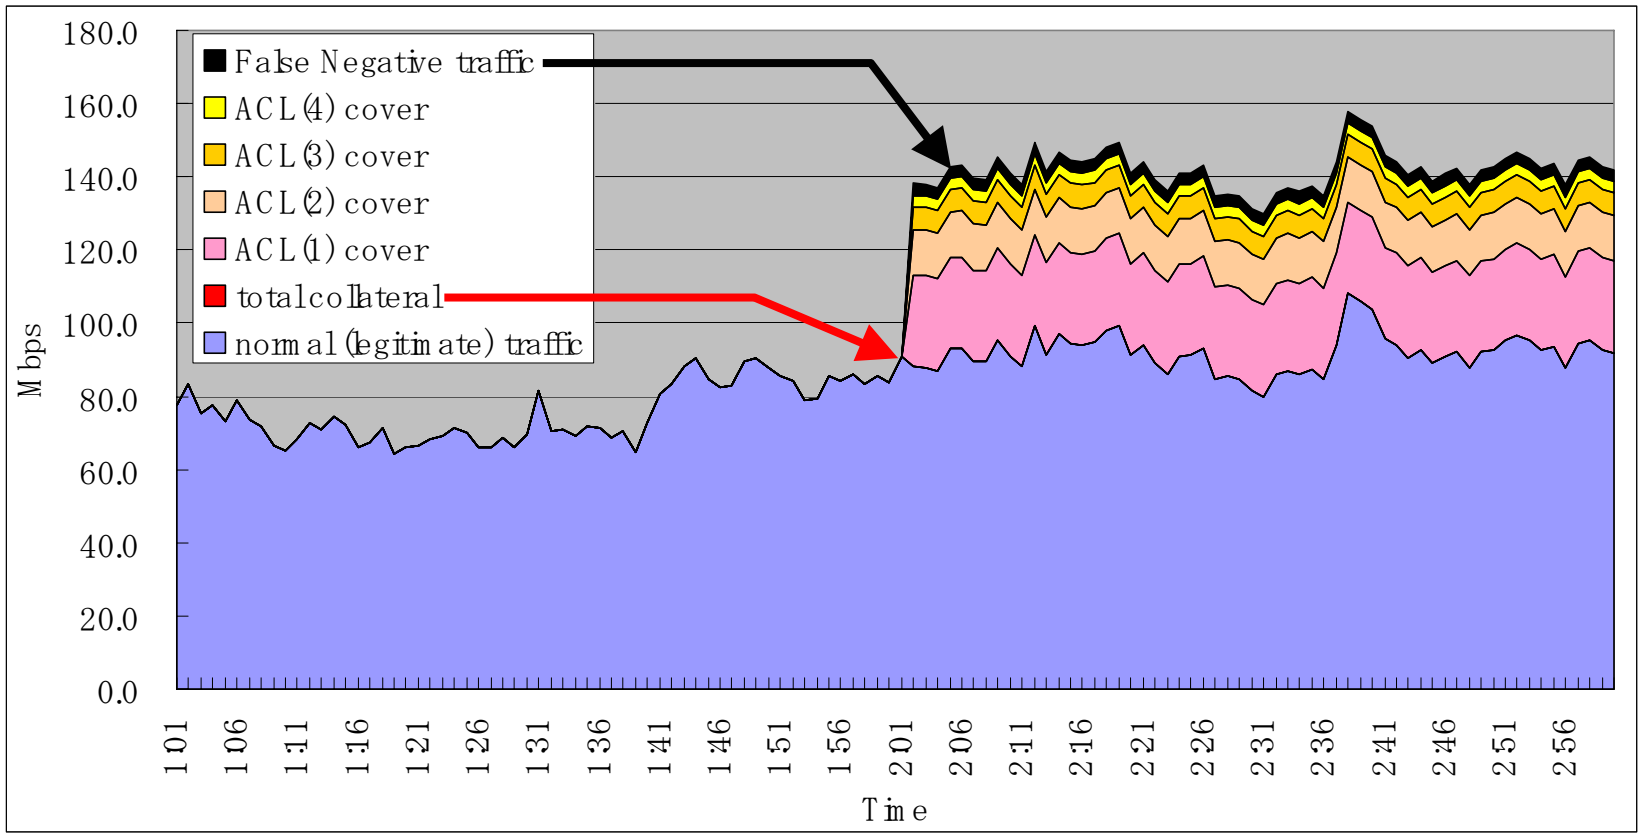
Evaluation and Results: Test Data Set

- Normal traffic: publicly available traffic data captured on transpacific line (100 Mbps)
- Anomalous traffic: injected synthesized DDoS attack traffic
  - Mimic large DDoS attack
    - We choose source/destination addresses that have large normal traffic because simple identification would cause collateral.
  - Destination: Popular server appeared in normal traffic
  - Source: Choose IP address blocks (/16) from which volume of normal traffic to the destination is largest.
  - Port numbers and protocol of attack traffic are the same as those of normal traffic.



Evaluation and Results: Results (1)

- Results: We get four ACL sets with below conditions
  - coverage: 93.75%
  - collateral: 0.00%
  - no. of ACL sets: 4



Copyright NTT Information Sharing Platform Laboratories, 2007

Evaluation and Results: Results (2) OUTPUT

basetime_len= 60.0 (sec) : (1168362060.0 - 1168362120.0) **basic information**
anomtime_len= 60.0 (sec) : (1168362180.0 - 1168362240.0)
base_total_bps= 89,121,539.5
anom_total_bps= 137,729,812.7
diff_total_bps= 48,608,273.2 +54.5 %

1-D_OUTPUT: PROTOCOL= 6 coverage= 100.42 collateral= 95.52 **single dimension**
1-D_OUTPUT: SRC_PORT= high coverage= 108.27 collateral= 33.42 **identification**
1-D_OUTPUT: DST_PORT= high coverage= 100.09 collateral= 96.40 **results**
1-D_OUTPUT: **SRC_IP** coverage= 96.43 collateral= 0.00
119.170.0.0/17 coverage= 51.43 collateral= 0.00
119.170.128.0/18 coverage= 25.72 collateral= 0.00
119.170.192.0/19 coverage= 12.86 collateral= 0.00
119.170.240.0/20 coverage= 6.43 collateral= 0.00
1-D_OUTPUT: **DST_IP** coverage= 102.93 collateral= 2.17
134.45.182.70/32 coverage= 102.93 collateral= 2.17

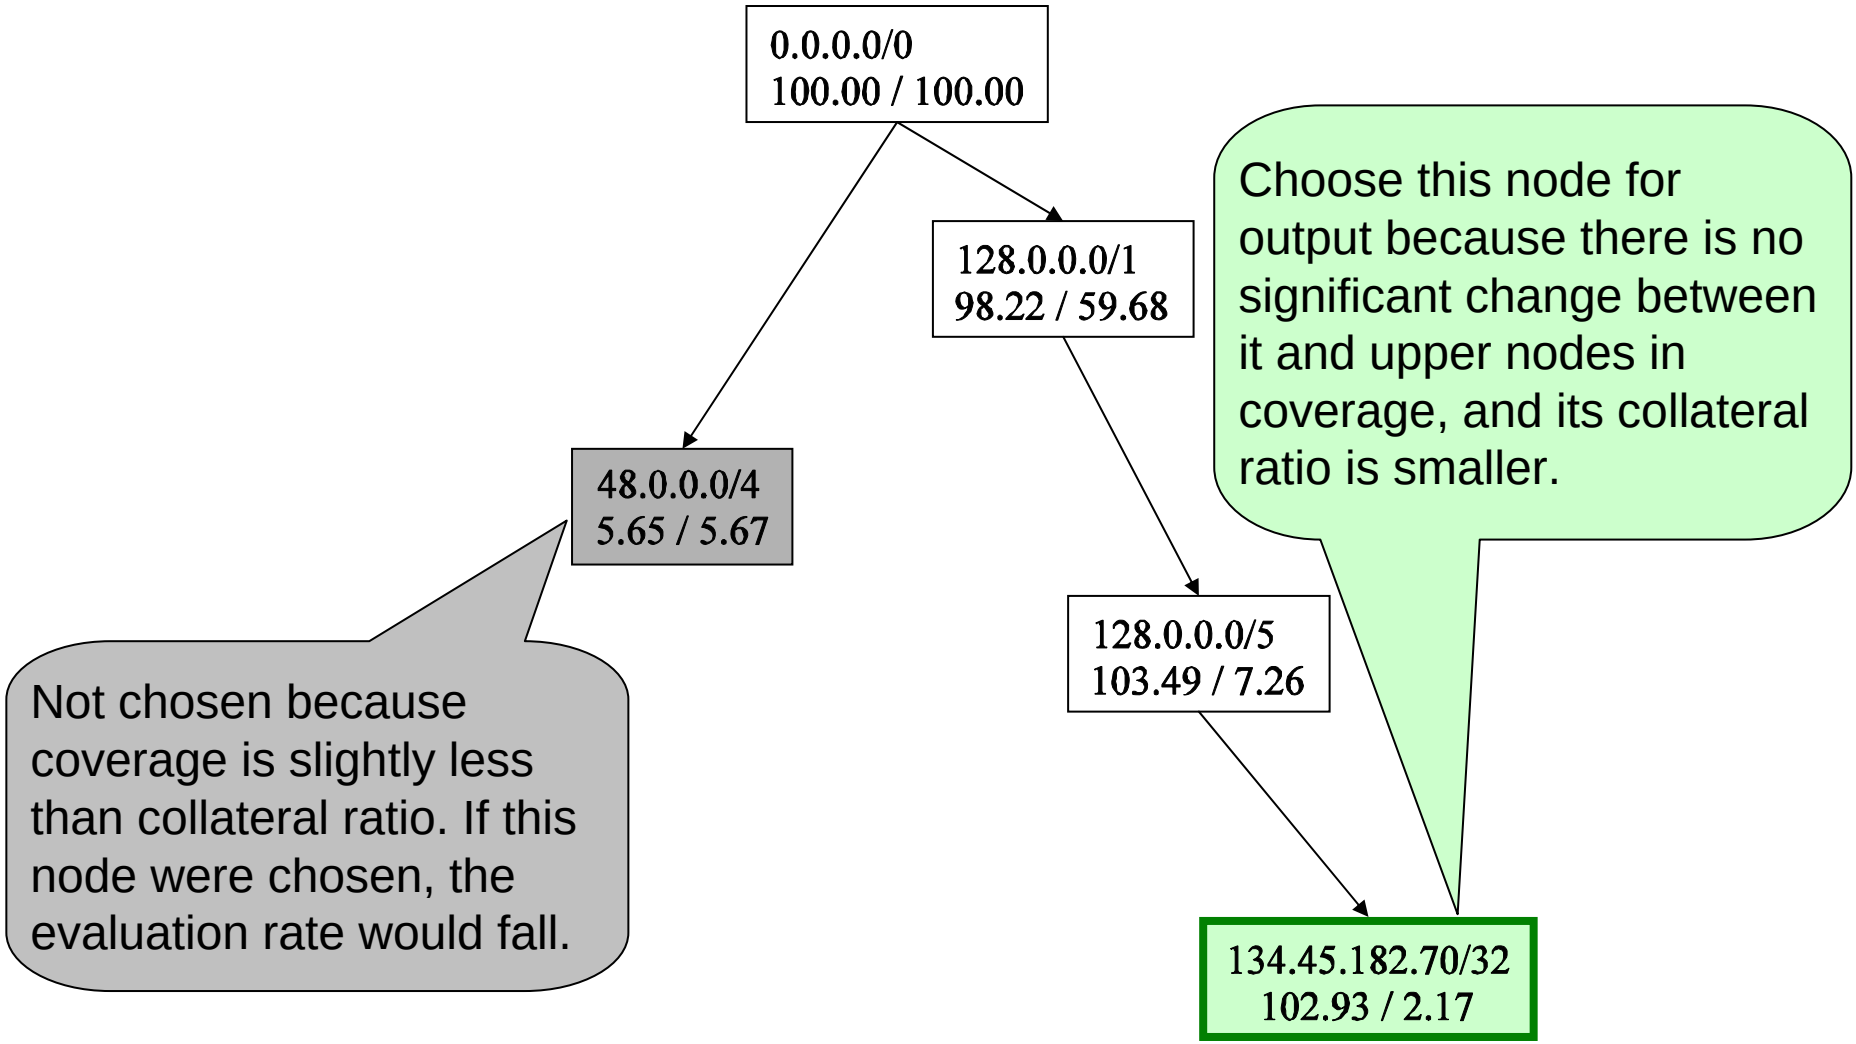
MULTI-DIMENSION_FLOW_OUTPUT coverage= 96.43 collateral= 0.00

flowID_0: cov= 51.43 col= 0.00:	**119.170.0.0/17**	**134.45.182.70/32**	6	high	high
flowID_1: cov= 25.72 col= 0.00:	**119.170.128.0/18**	**134.45.182.70/32**	6	high	high
flowID_2: cov= 12.86 col= 0.00:	**119.170.192.0/19**	**134.45.182.70/32**	6	high	high
flowID_3: cov= 6.43 col= 0.00:	**119.170.240.0/20**	**134.45.182.70/32**	6	high	high

↑ **coverage** ↑ **collateral**: ↑ src_ip ↑ dst_ip protocl scr_port dst_port

Evaluation and Results (3): Destination IP Tree

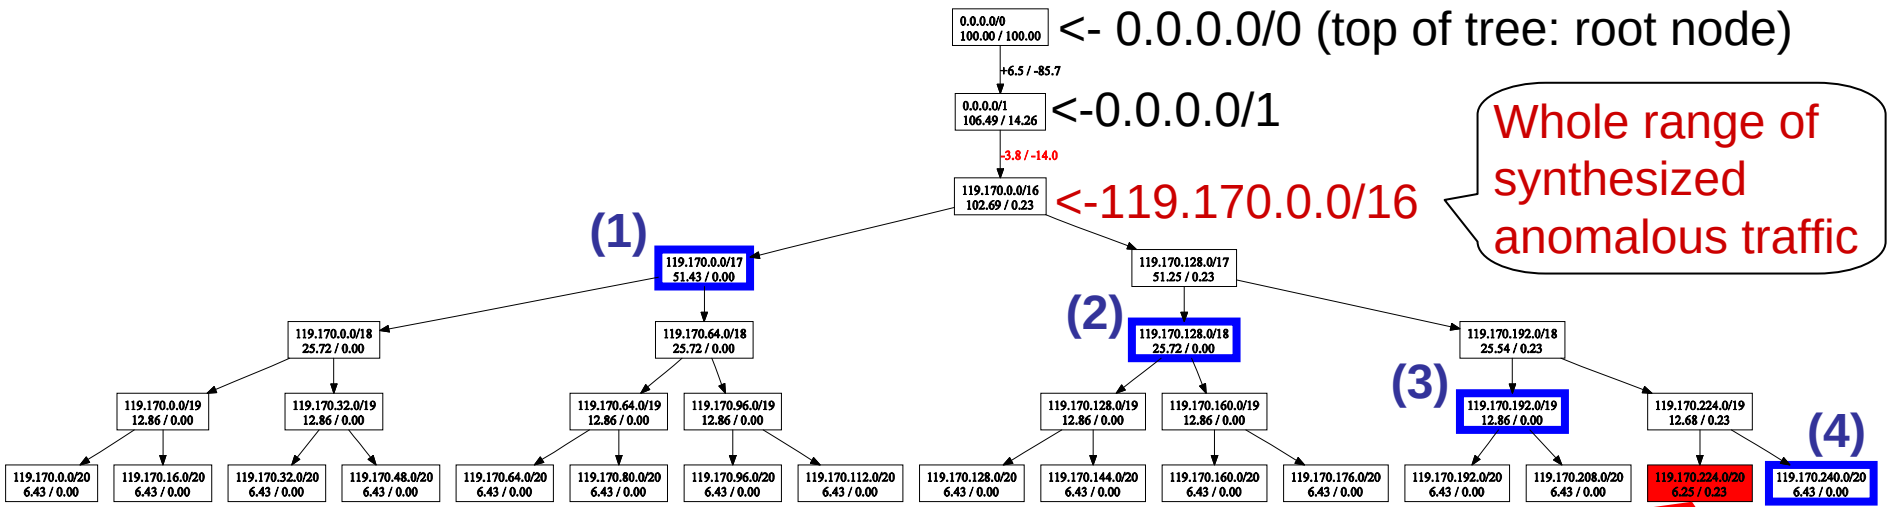
1-D_OUTPUT: **DST_IP** coverage= 102.93 collateral= 2.17
134.45.182.70/32 coverage= 102.93 collateral= 2.17



Copyright NTT Information Sharing Platform Laboratories, 2007

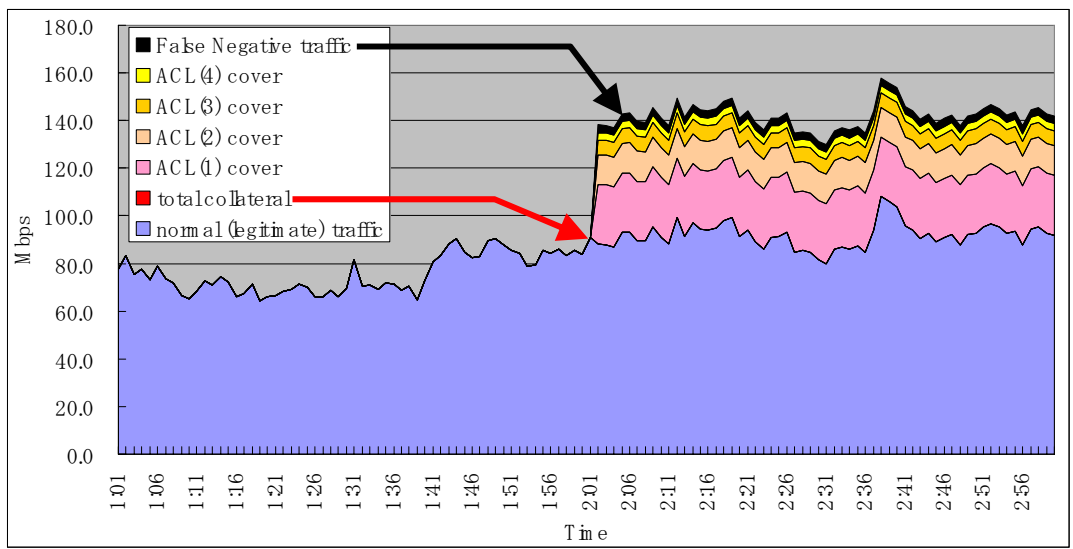
Evaluation and Results (4): Source IP Tree

1-D_OUTPUT: **SRC_IP**		coverage=	96.43	collateral=	0.00
(1)	**119.170.0.0/17**	coverage=	51.43	collateral=	0.00
(2)	**119.170.128.0/18**	coverage=	25.72	collateral=	0.00
(3)	**119.170.192.0/19**	coverage=	12.86	collateral=	0.00
(4)	**119.170.240.0/20**	coverage=	6.43	collateral=	0.00



Summary

- Revealed three criteria of optimal ACL sets.
  - for mitigating DDoS attacks on router
- Proposed DELTAA technique: Optimizes trade-off among the these criteria, using normal and anomalous traffic.
- Showed effectiveness of DELTAA.
  - Evaluation results using prototype and synthesized data sets:
    - coverage: 93.75%
    - collateral: 0.00%
    - no. of ACL sets: 4



Thank you.

Any questions are welcome.

This study was supported by
the Ministry of Internal Affairs and Communications of Japan.

Design for Large-Scale Collection System Using Flow Mediators

Atsushi Kobayashi, Tsuyoshi Kondoh, and
Keisuke Ishibashi
NTT Information Sharing Laboratories

Outline

- Introduction
  - Why do we need a large-scale collection system?
  - What is Flow Mediator?
- Requirements
  - I tried to explore the possibility of a large-scale collection system for large networks.
- Heuristic method of designing traffic collection system
  - Estimate number of flow records after aggregation or sampling
  - Adjust several parameters based on this result
- Summary

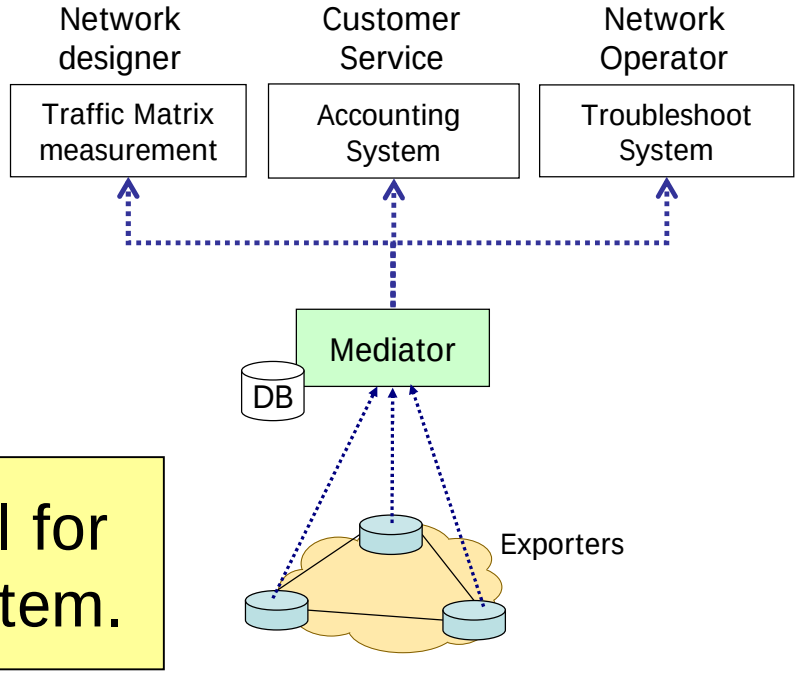
Introduction

- Traffic volumes in ISP networks are becoming huge in the last few years.
  - The number of exported flow records is becoming so huge that a single collector cannot handle them.
- A smaller sampling rate makes small flows invisible.
  - Even if traffic grows, network operators would like to maintain the same sampling rate as much as possible.
- Aggregated flow records from router make port number or IP address invisible.
  - Exporting 5-tuple flow records from router is better.

The demand for a large-scale traffic-collection system is growing.

What is Flow Mediator?

- Flow Mediator† is a device that "mediates" flow records and has the following functions:
  - collects Flow Records from various exporters
  - stores original flow records
  - aggregates flow records flexibly
  - distributes appropriate flow records for collectors/analyzers

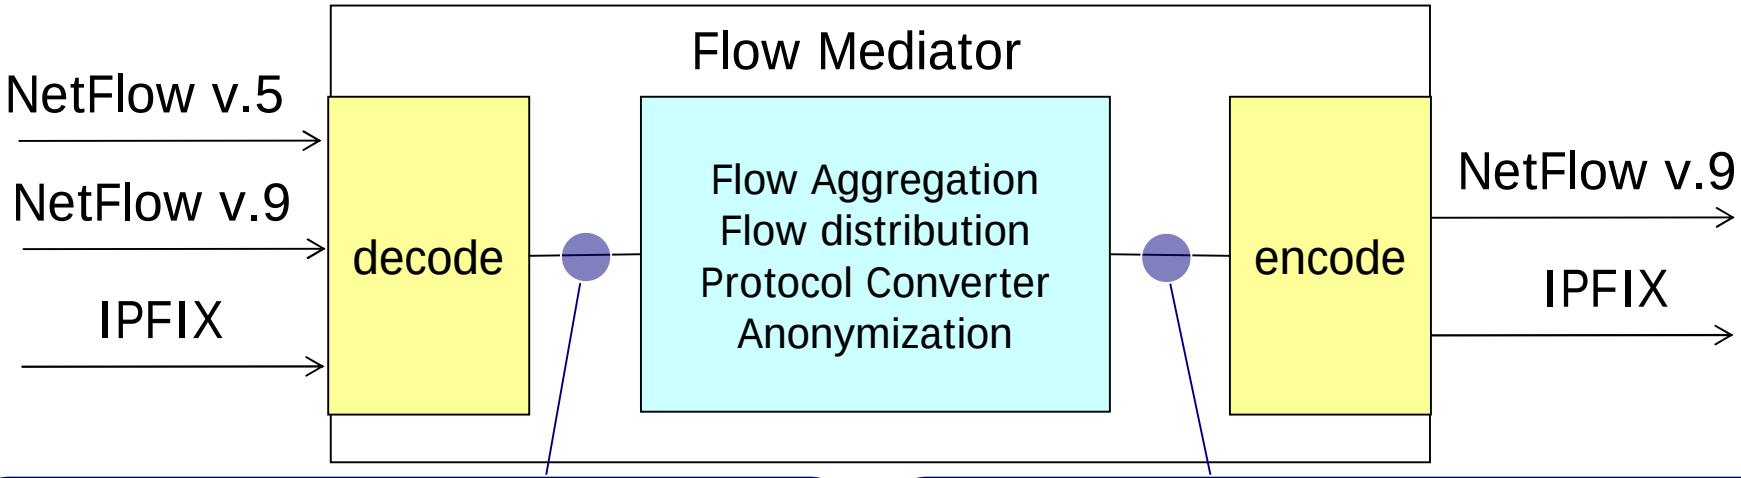


Flow mediator ought to be useful for making large-scale collection system.

† draft-kobayashi-ipfix-mediator-model-01.txt

You can easily make Flow Mediation code

- Net::Flow perl module is available on CPAN.
  - http://search.cpan.org/~akoba/Net-Flow-0.02/
  - The module can encode and decode NetFlow/IPFIX packets.
  - The encoding and decoding functions have a similar IF.




```
my ( $HeaderHashRef,
     $TemplateArrayRef,
     $FlowArrayRef,
     $ErrorsArrayRef) =
     Net::Flow::decode(
                    ¥$packet,
                    $TemplateArrayRef  ) ;
```

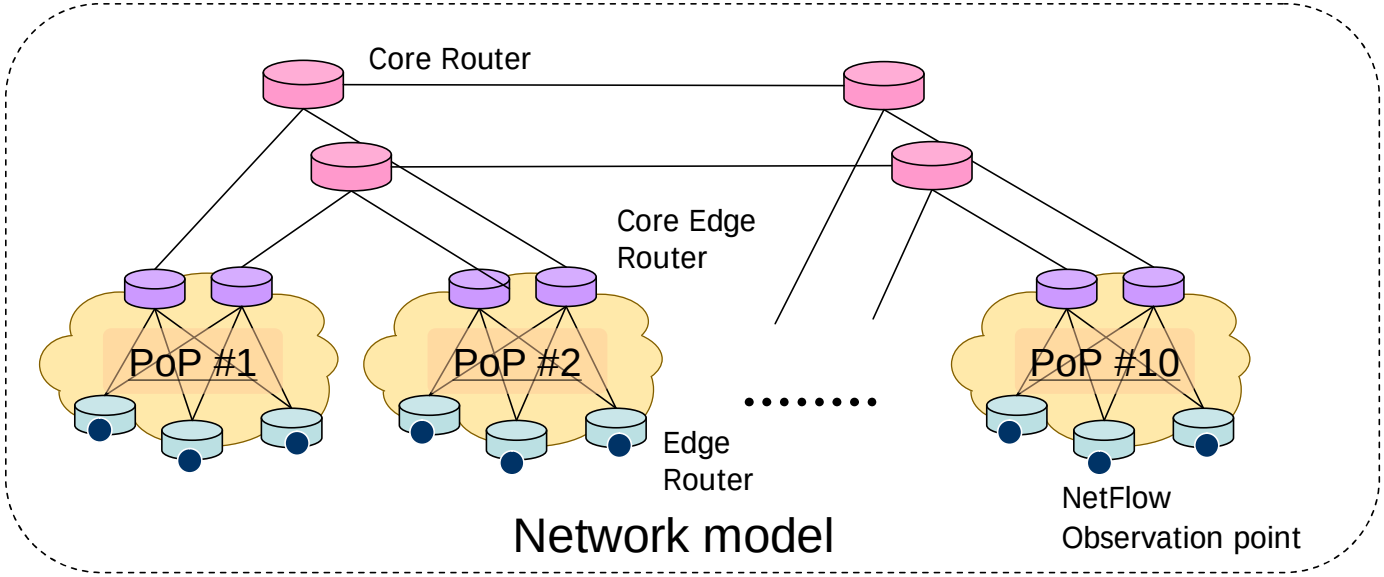
```
my ( $EncodeHeaderHashRef,
     $PktsArrayRef,
     $ErrorsArrayRef) =
     Net::Flow::encode(
                    $EncodeHeaderHashRef,
                    ¥@MyTemplates,
                    $FlowArrayRef,
                    1400 ) ;
```

Requirements

- Make traffic-collection system to meet following requirements
  - Requirement 1: measure traffic flow of entire networks
    - measure traffic matrices PoP by PoP and router by router
  - Requirement 2: store received 5-tuple flow records from router
    - When traffic incident happens, allow inspection of traffic.
  - Requirement 3: design scalable architecture to accommodate large ISP traffic volume

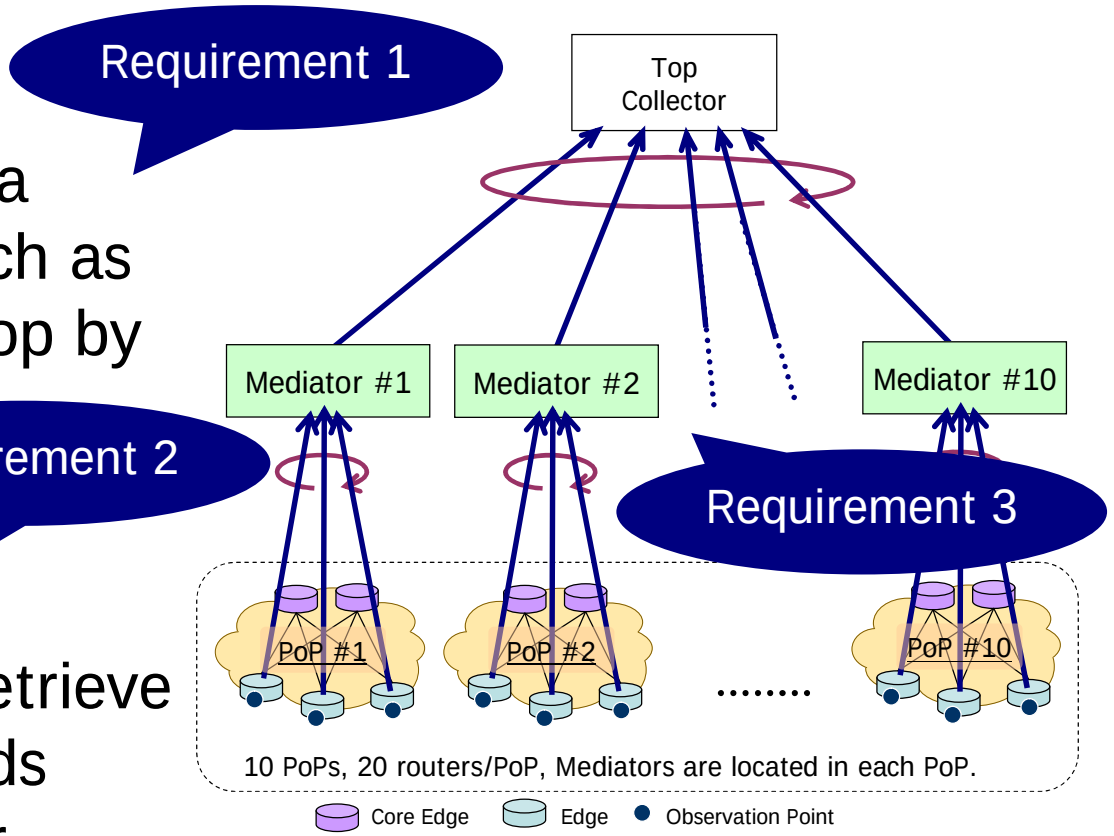
Goal

- Explore heuristic method of designing collection system for introduction into actual network.
- Proposed collection system needs to accommodate following network model.
  - Total traffic volume 500 Gb/s, 100 Mp/s
    - Edge Router 20/PoP×10 PoP = 200
    - NetFlow is enabled on IngressIF of Edge router.



Hierarchical Collection System

- Mediators are allocated in each PoP.
  - They store all flow records, aggregate them, and export them to next collector.
- Top Collector
  - measures wide-area traffic matrices, such as router by router, pop by pop.
- Inspection
  - If traffic incident happens, we can retrieve detailed flow records from Flow Mediator.

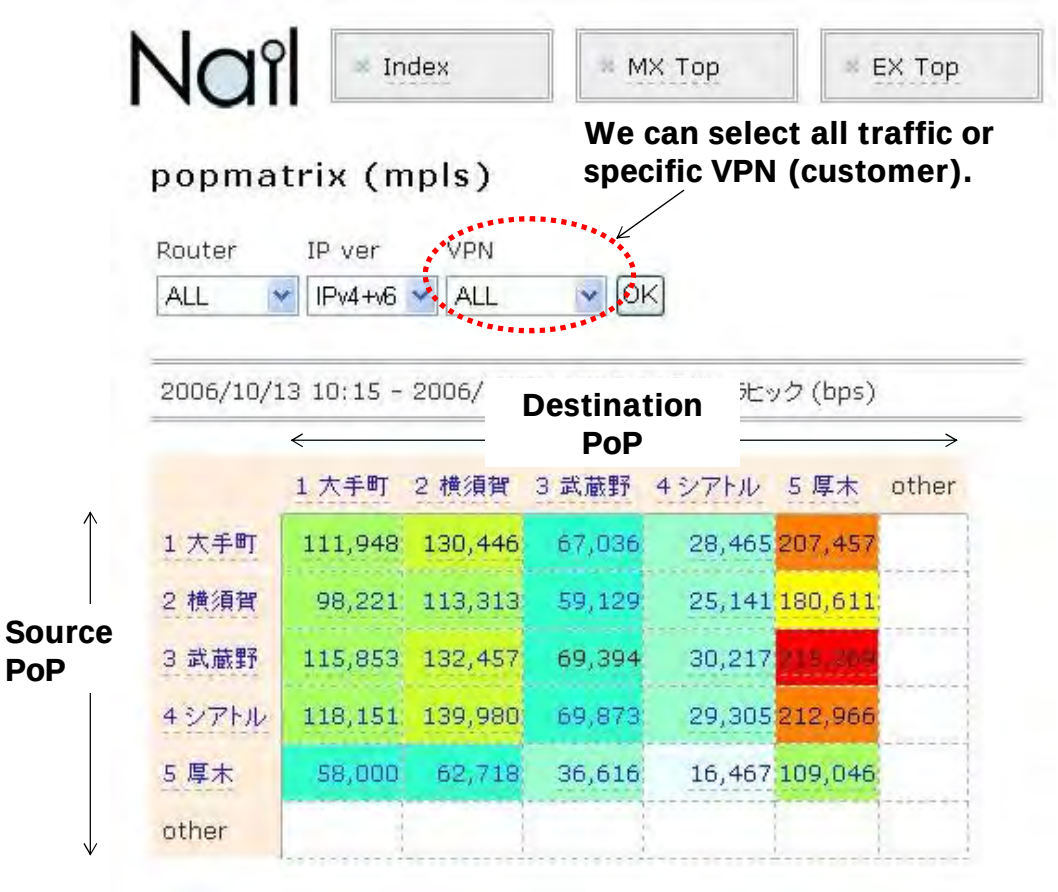


Visualize Traffic Matrices

- Top collector can visualize Router/PoP/AS Traffic Matrixes.

Nail is the name of our traffic matrix visualizer.

Color indicates traffic volume of Source/ Destination pair.

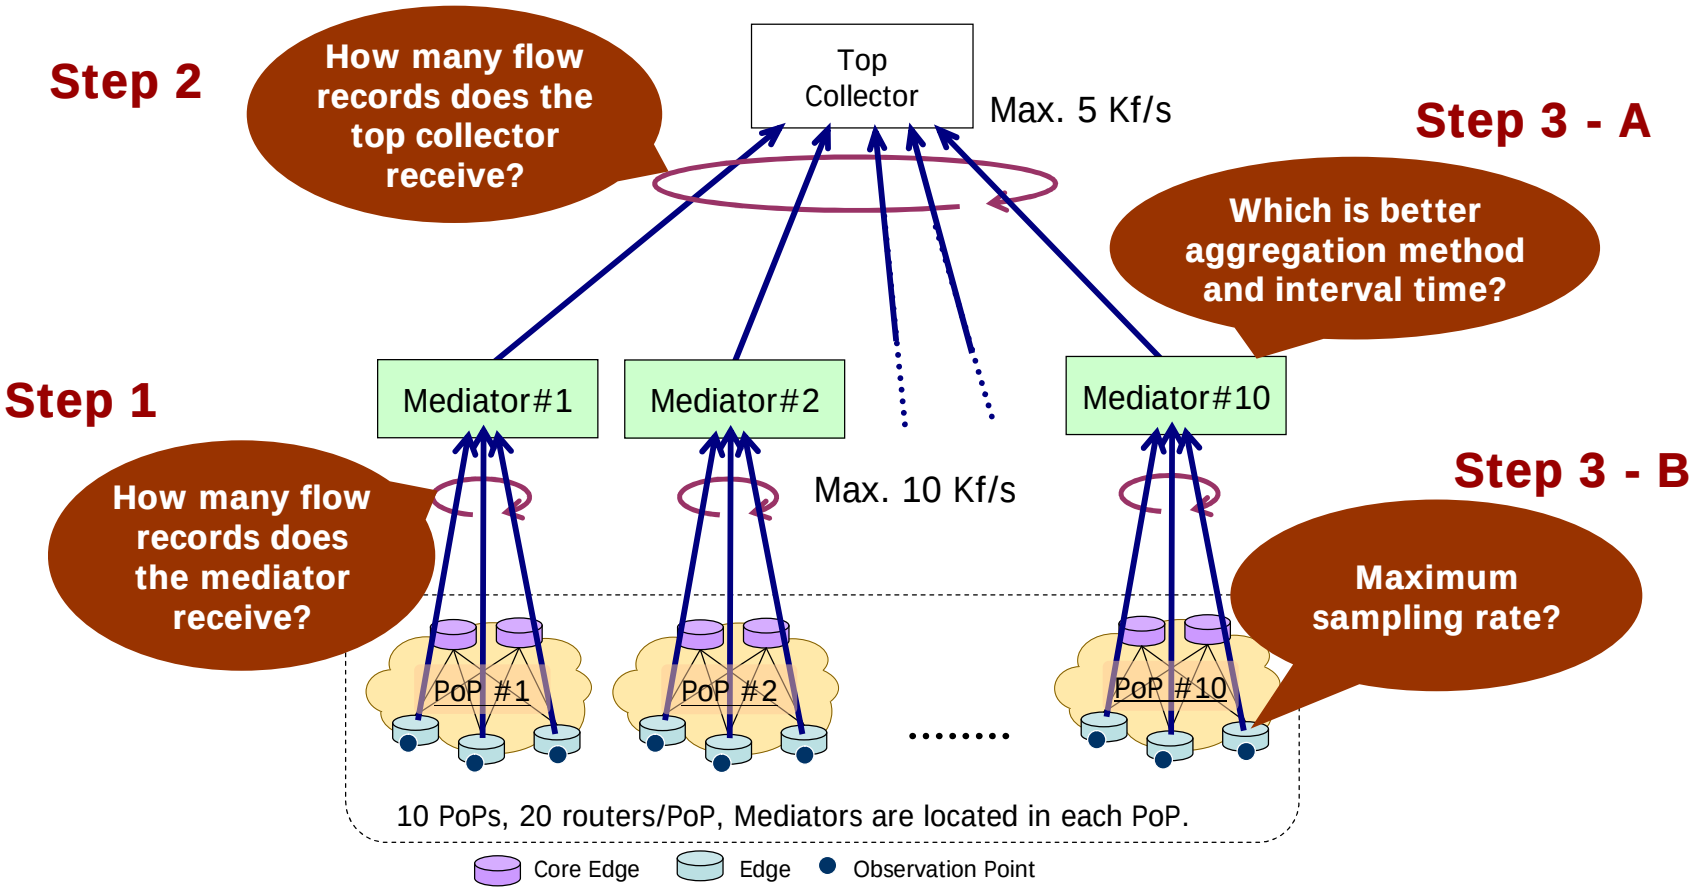


Heuristic Design Method

- Suitable values of several parameters are decided by the following steps.
  - Step 0: measure performance limit of flow mediator and top collector.
  - Step 1: reveal relation between number of flow records and packet sampling
  - Step 2: reveal relation between number of flow records and aggregation that depends on several factors.
    - Aggregation methods (BGP Next-Hop, Prefix, host)
    - Aggregation interval time (20 s, 60 s, 90 s...)
  - Step 3: select suitable value within performance limit.
    - Large sampling rate is preferable.
    - Small granularity of aggregation is preferable.

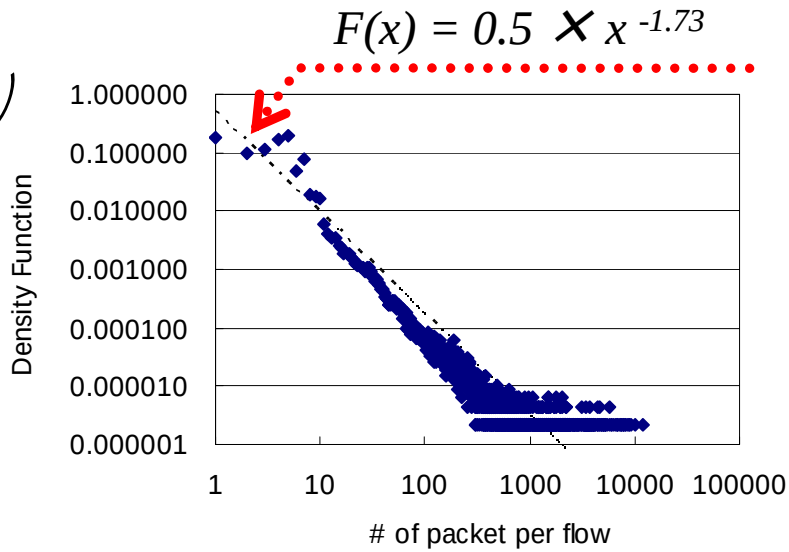
Consideration Points

- List several considerations, as follows.
  - Maximum performances of the top collector and mediators are 5 Kf/s and 10 Kf/s.



Step 1: estimate flow records after sampling

- Estimate number of flow records based on density function of packets per flow.
  - # of packets per flow: x
  - Packets per flow density function: $F(x)$
  - Sampling rate: $1/r$
  - Total number of unsampled flow: f_{all}



$$f_{sampled} = \sum_{x=1}^{\infty} \left(1-(1-1/r)^x\right) \times F(x) \times f_{all}$$

Extraction probability

$0.5x^{-1.73}$

Roughly estimate as follows.
100 Mpps ÷ 20 packets = 5 Mf/s

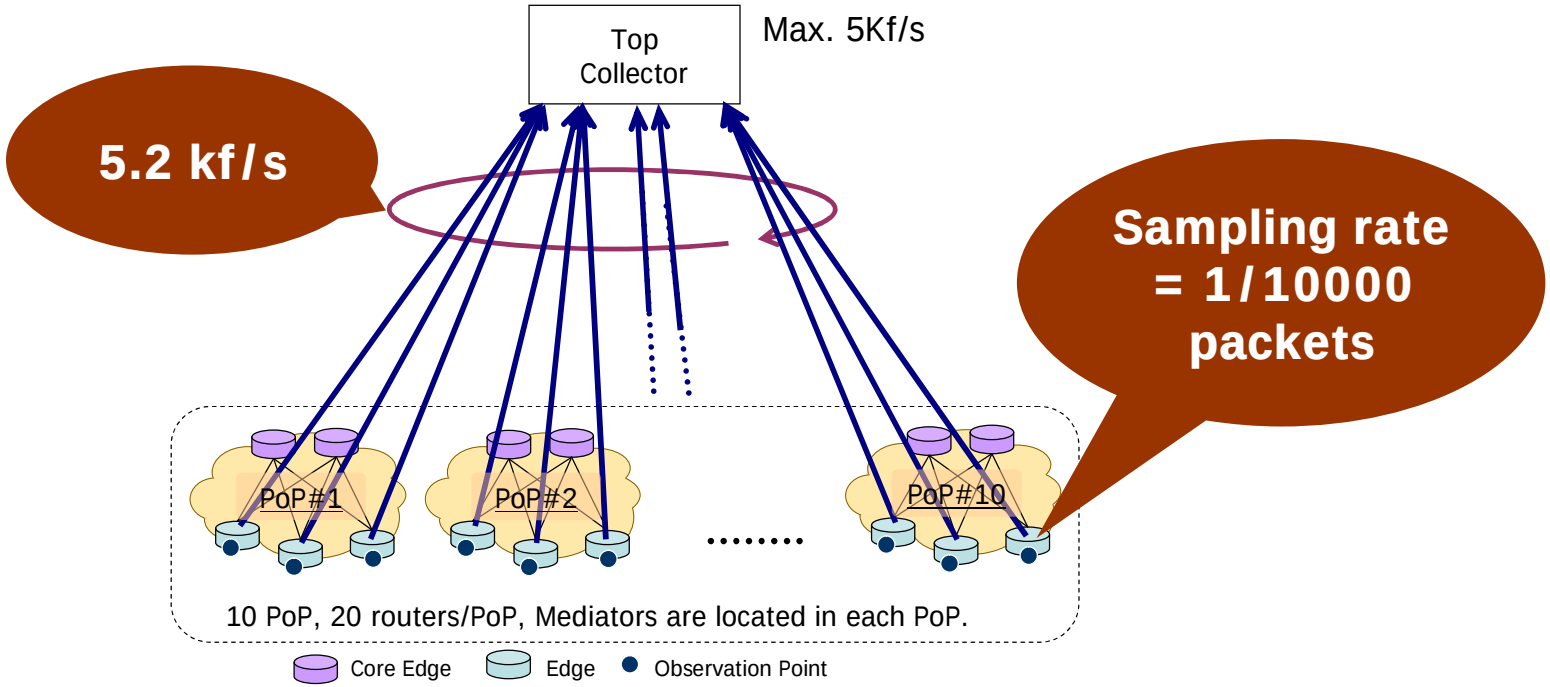
Approximate # of flows when total traffic volume is 500 Gb/s.

Sampling rate	1/100	1/1000	1/10000
$f_{sampled}$	*305 kf/s*	*43 kf/s*	*5.2 kf/s*

Too many flow records without mediator

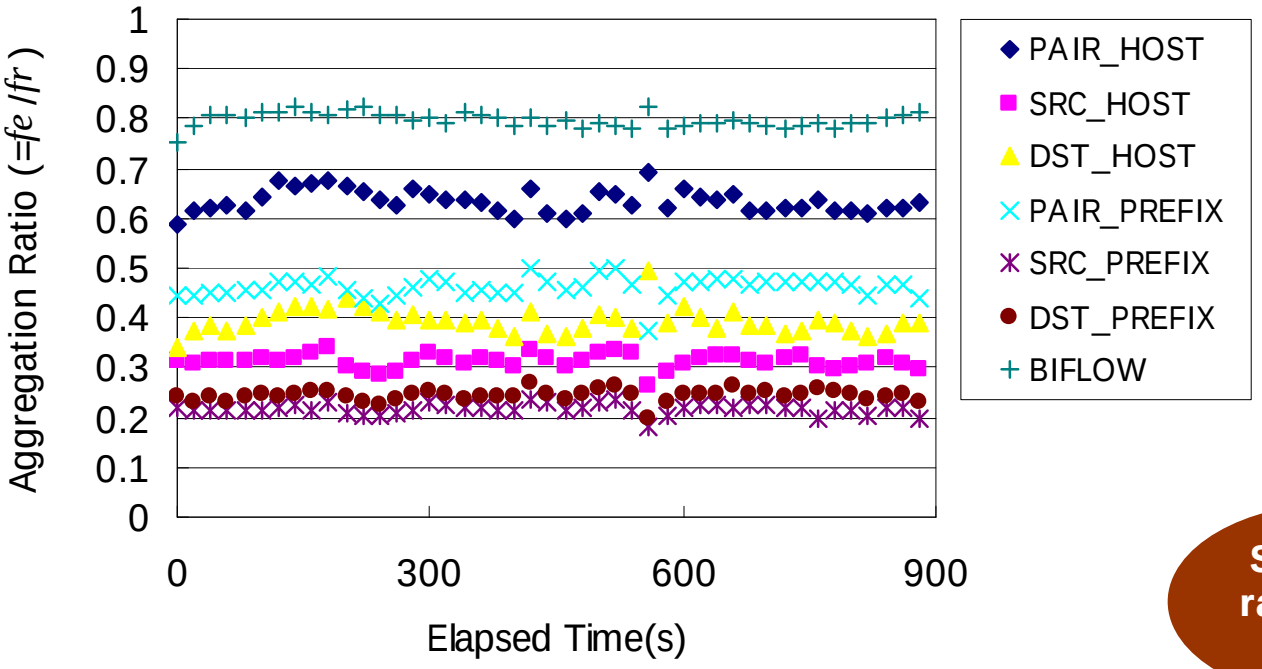
- Even if sampling rate is 1/10,000 packets, the number of flow records exceeds performance limit.

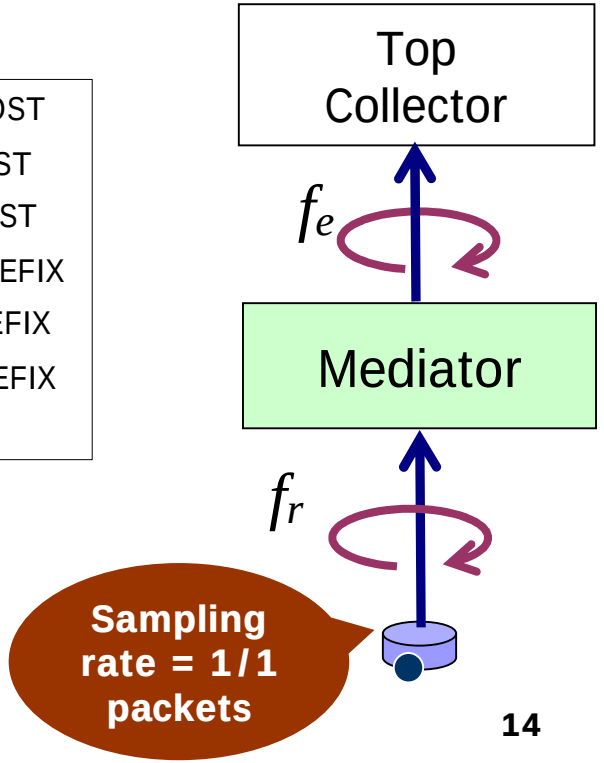
Sampling rate	1/100	1/1000	1/10000
$f_{sampled}$	*305 kf/s*	*43 kf/s*	*5.2 kf/s*



Step 2: flow records after aggregation

- What is the # of flow records after aggregation?
- Mediator aggregates unsampled flow records at 20-second interval.
  - Aggregation efficiency: Prefix > HOST > Pair Prefix > Pair HOST > Bi-Flow
    - The prefix length "/24" is uniformly applied to Prefix Aggregation.
    - Bi-flow is aggregated from two flow directions.

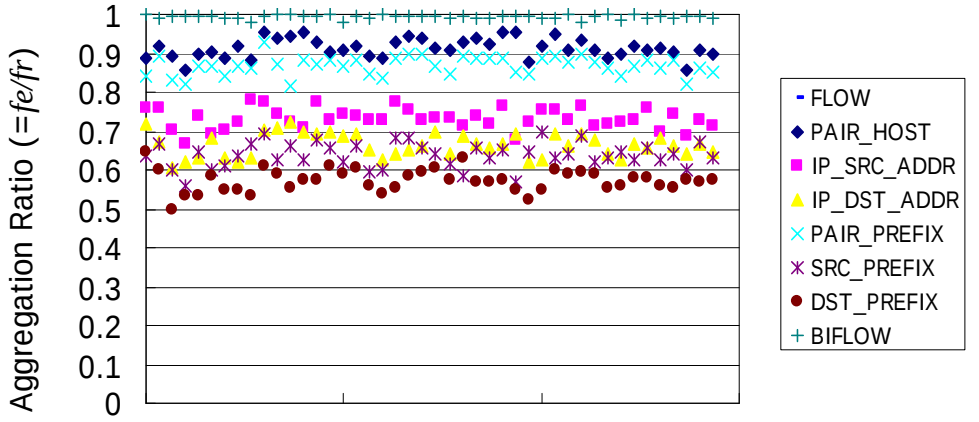




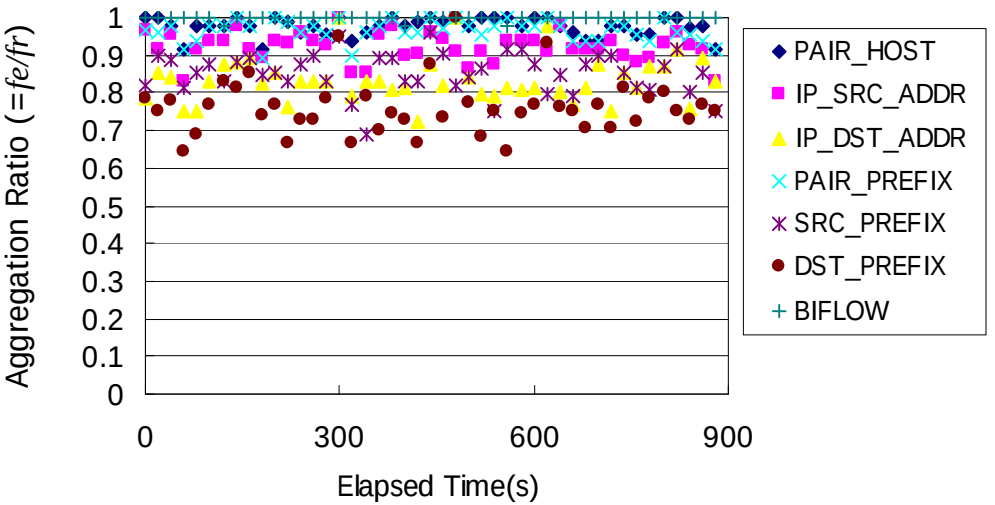
Step 2: Flow records after aggregation, sampling

- Each aggregation method becomes ineffective gradually.
- Bi-flow becomes ineffective immediately.
  - sensitive to sampling rate.

Sampling rate 1/128



Sampling rate 1/1024



Step 2: Which factor influences aggregation?

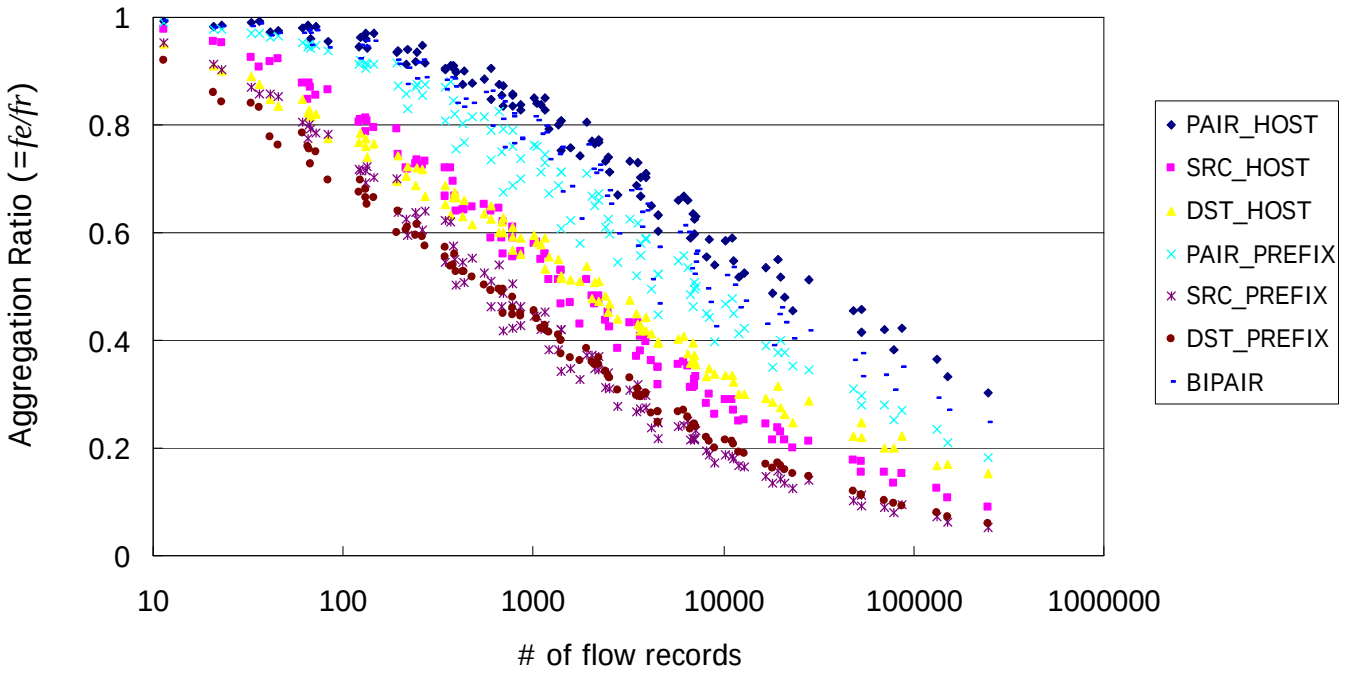
- Aggregation ratio depends on several factors.
  - Traffic Volume through observation point.
  - Sampling rate
  - Aggregation interval time

I guess that the aggregation ratio depends on the number of flow records received in interval time.

Received Flows	3450	3562
Aggregation Interval Time (s)	10	300
Sampling rate (1/r)	1	128
DST_HOST Aggregation ratio	45%	43%
DST_PREFIX Aggregation ratio	30%	32%

Step 2: Which factor influences aggregation?

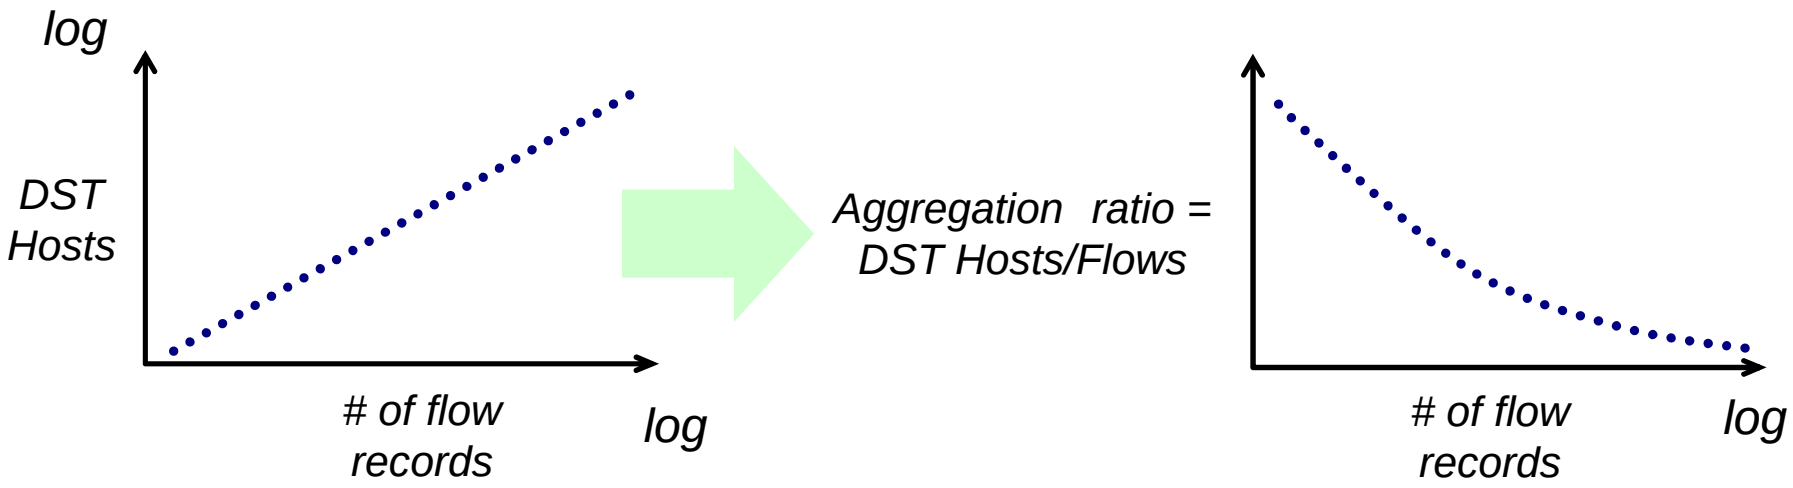
- I plotted all experimental data into one graph.
  - Three MAWI traffic data samples have different volumes.
  - Aggregation Interval time：5 – 300s
  - Sampling rate：1/1 – 1/1024



Aggregation ratio depends on number of received flow records.

Step 2: Formulation of Aggregation Ratio

- Aggregation ratio (R) can be estimated from number of flow records (f_r), as follows.
  - DST Host aggregation: $R_{dsthost} = 1.80 \times f_r^{-0.18}$
  - DST Prefix aggregation: $R_{dstprefix} = 2.34 \times f_r^{-0.26}$
- After all, the aggregation ratio depends on the # of unique hosts or prefixes versus # of flows.



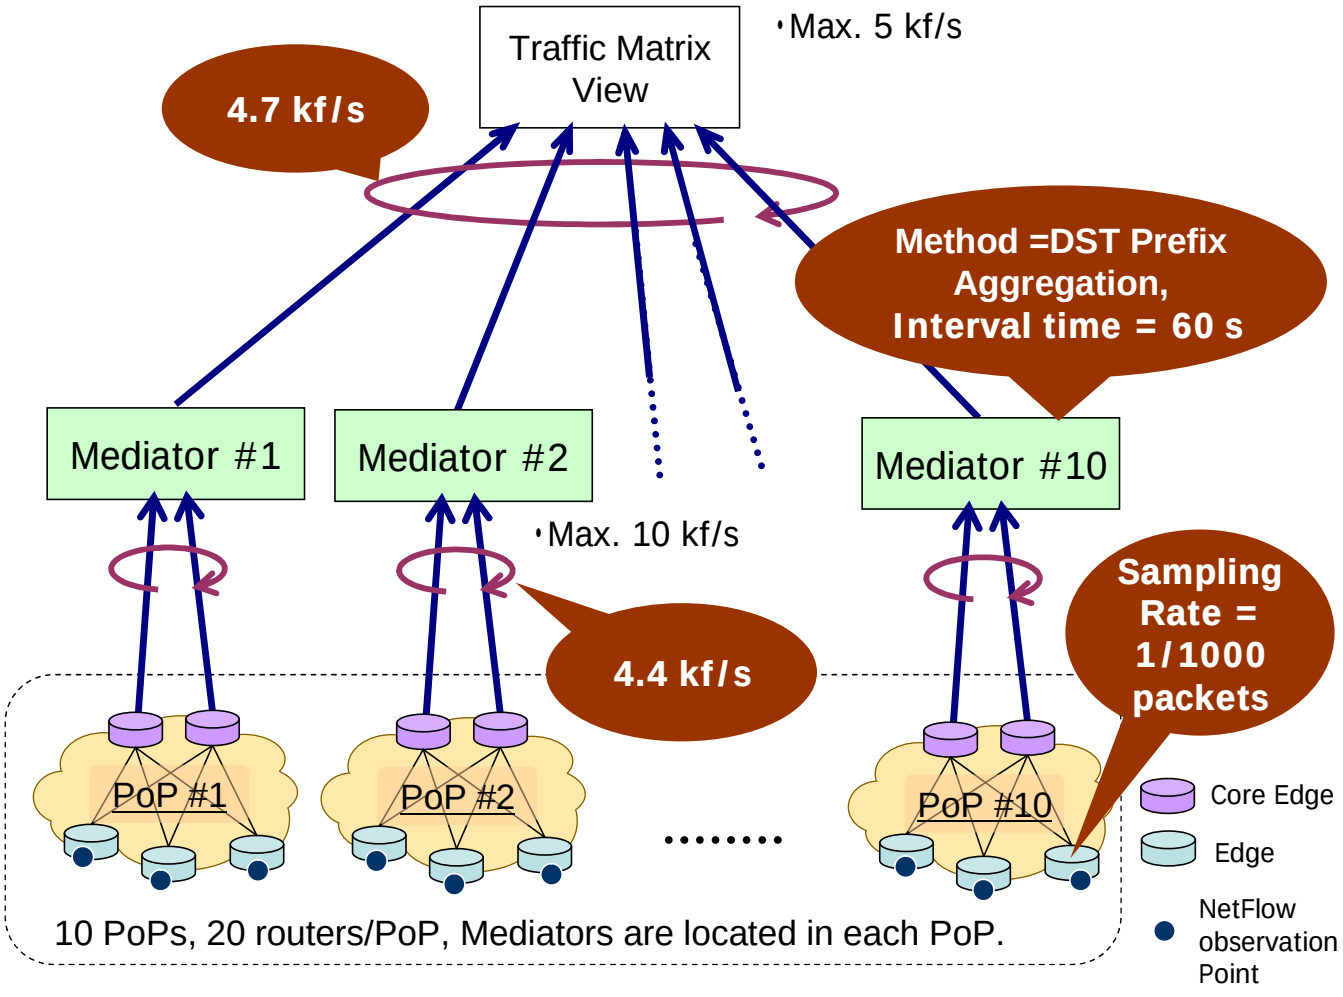
Step 3: Selection of Suitable Values

- I selected suitable value within performance limit.

		Sampling Rate	1/100	1/1000	1/10000
# of received flow records in top collector ($= \Sigma f_e$)	DST_HOST aggregation	Interval time = 60s	*45 kf/s*	9.0 kf/s	***1.6 kf/s***
	DST_Prefix aggregation	Interval time = 60s	*21 kf/s*	***4.7 kf/s***	***0.94 kf/s***
	DST_HOST aggregation	Interval time = 300s	*34 kf/s*	7.0 kf/s	***1.2 kf/s***
	DST_Prefix aggregation	Interval time = 300s	*12 kf/s*	***3.0 kf/s***	***0.62 kf/s***
# of received flow records in mediator (f_r)	—	—	*30 kf/s*	***4.4 kf/s***	***0.6 kf/s***

Example of collection system

- Sampling Rate: 1/1000
- Aggregation Interval time: 60 s



Conclusion

- To make large scale traffic collection system, flow mediator is efficient.
- Revealed relation between number of flow records and several factors:
  - Traffic volume
  - Sampling rate
  - Aggregation method
  - Aggregation interval time
- Demonstrated that traffic collection system using mediator can be introduced into actual large-scale networks.

Thank you for your attention.

This study was supported by the Ministry of Internal Affairs and Communications of Japan.

Design for a Large-Scale Collection System using Flow Mediators

Atsushi Kobayashi, Tsuyoshi Kondoh, and
Keisuke Ishibashi
NTT Information Sharing Laboratories

Outline

- Motivation
  - Approach to the scalability in Large NW
  - What is Flow Mediators?
  - Introduce Hierarchical model
- Design method of Collection system
  - Estimation received Flows after sampling
  - Estimation received Flows after aggregation
  - Results of reference model
- Conclusion

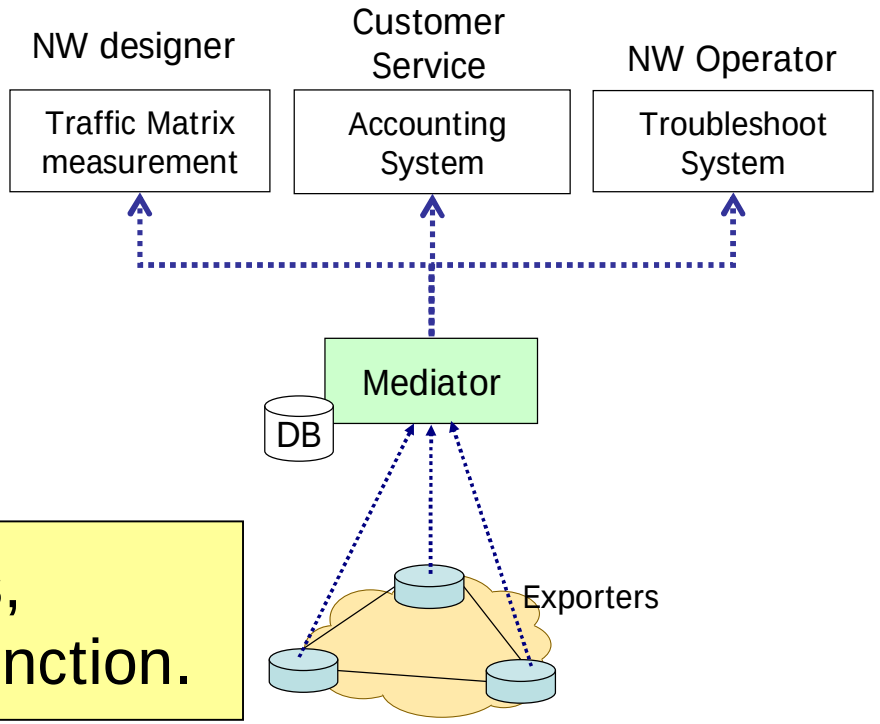
Motivation

- Try to measure wide-area Traffic Matrix and in-depth inspection.
  - Single collector cannot accomplish both requirements.
- Especially, difficult to maintain the scalability of collection system in large NW.
  - Number of exported Flows becomes huge.
    - 100Gb/s traffic creates approximately 50Kf/s with sampling 1/1000.
    - Adjusting sampling rate could cause small Flows to become invisible.
- Approach to the scalability by using Flow Mediators
  - To make Flow collection scalable, efficient, useful.

Please refer to our draft draft-kobayashi-ipfix-large-ps-00.txt

What is Flow Mediator?

- Flow Mediator is a system that "mediates" Flow Records and has the following functions:
  - collects Flow Records from various exporters
  - stores original Flow Records
  - **aggregates Flow Records flexibly**
  - distributes appropriate Flow Records for dedicated collectors/analyzers

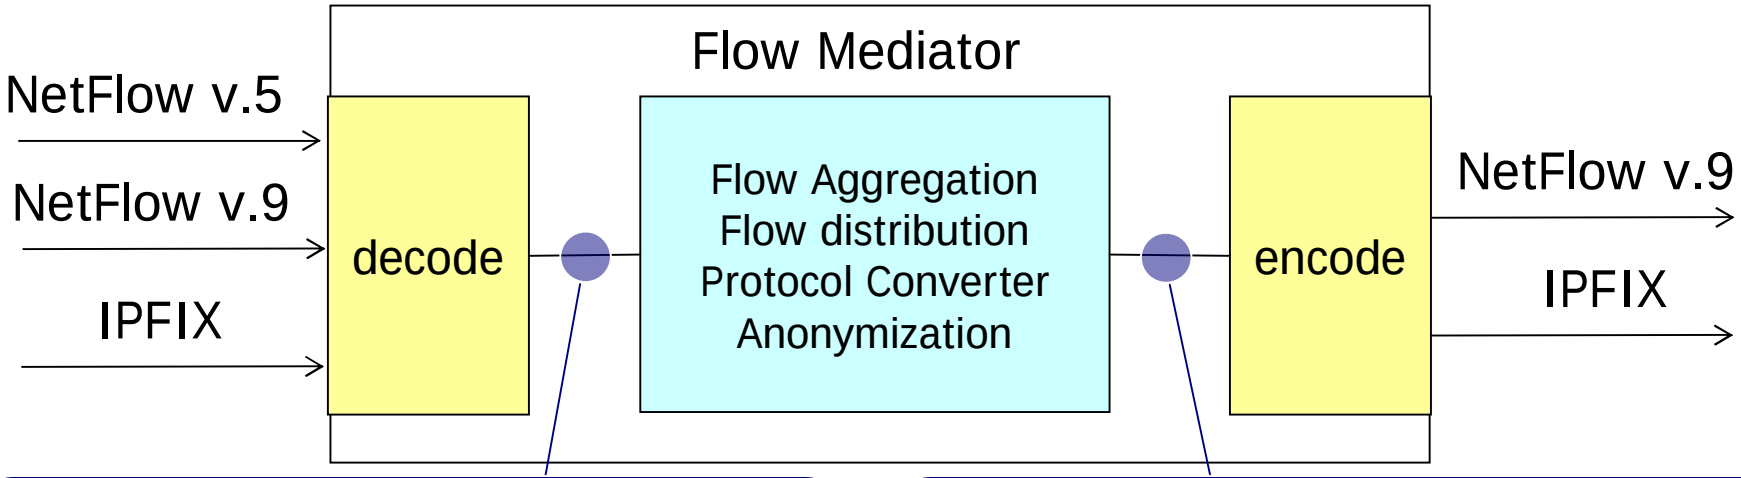


To reduce the number of Flows,
we focus on the aggregation function.

Please refer to our draft draft-kobayashi-ipfix-mediator-model-01.txt

You can feel Flow Mediation

- Net::Flow perl module is available on CPAN.
  - http://search.cpan.org/~akoba/Net-Flow-0.02/
  - The module can decode and encode NetFlow/IPFIX packets.
  - The decoding and encoding functions are similar IF.

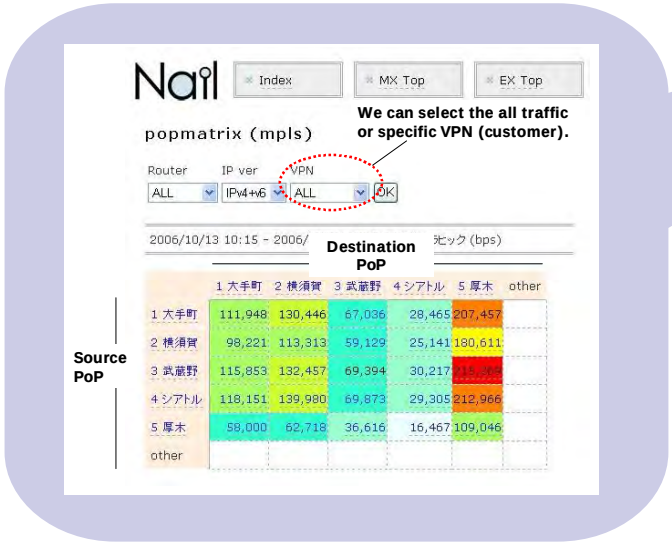


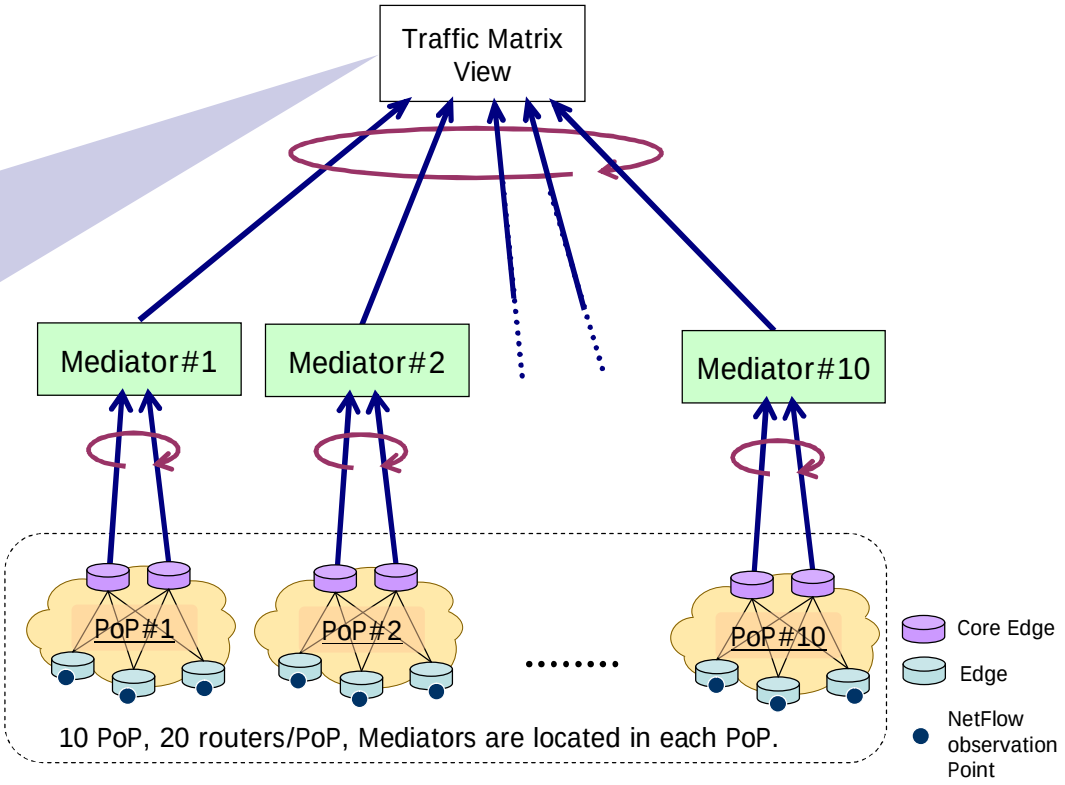

```
my ( $HeaderHashRef,
     $TemplateArrayRef,
     $FlowArrayRef,
     $ErrorsArrayRef) =
     Net::Flow::decode(
                    ¥$packet,
                    $TemplateArrayRef  ) ;
```

```
my ( $EncodeHeaderHashRef,
     $PktsArrayRef,
     $ErrorsArrayRef) =
     Net::Flow::encode(
                    $EncodeHeaderHashRef,
                    ¥@MyTemplates,
                    $FlowArrayRef,
                    1400 ) ;
```

Approach to Hierarchical Model

- Mediators
  - store the whole Flows, aggregate them and export to next collector.
- Top Collector
  - measures wide-area Traffic Matrices.
- If traffic incident happen, we can retrieve the detail Flow data from Mediator.

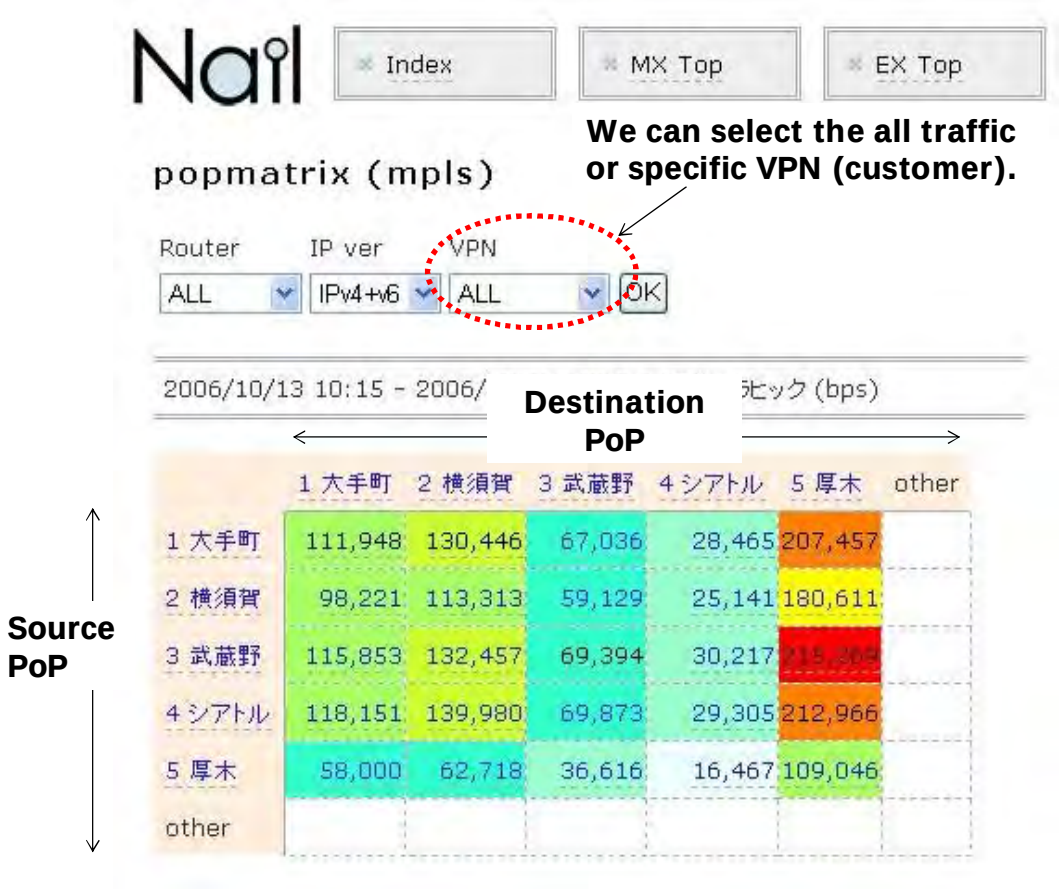




Wide-area Traffic Matrices

- Top collector can visualize the Router/PoP/AS Traffic Matrixes.

The color means the traffic volume of the Source/Destination pair.



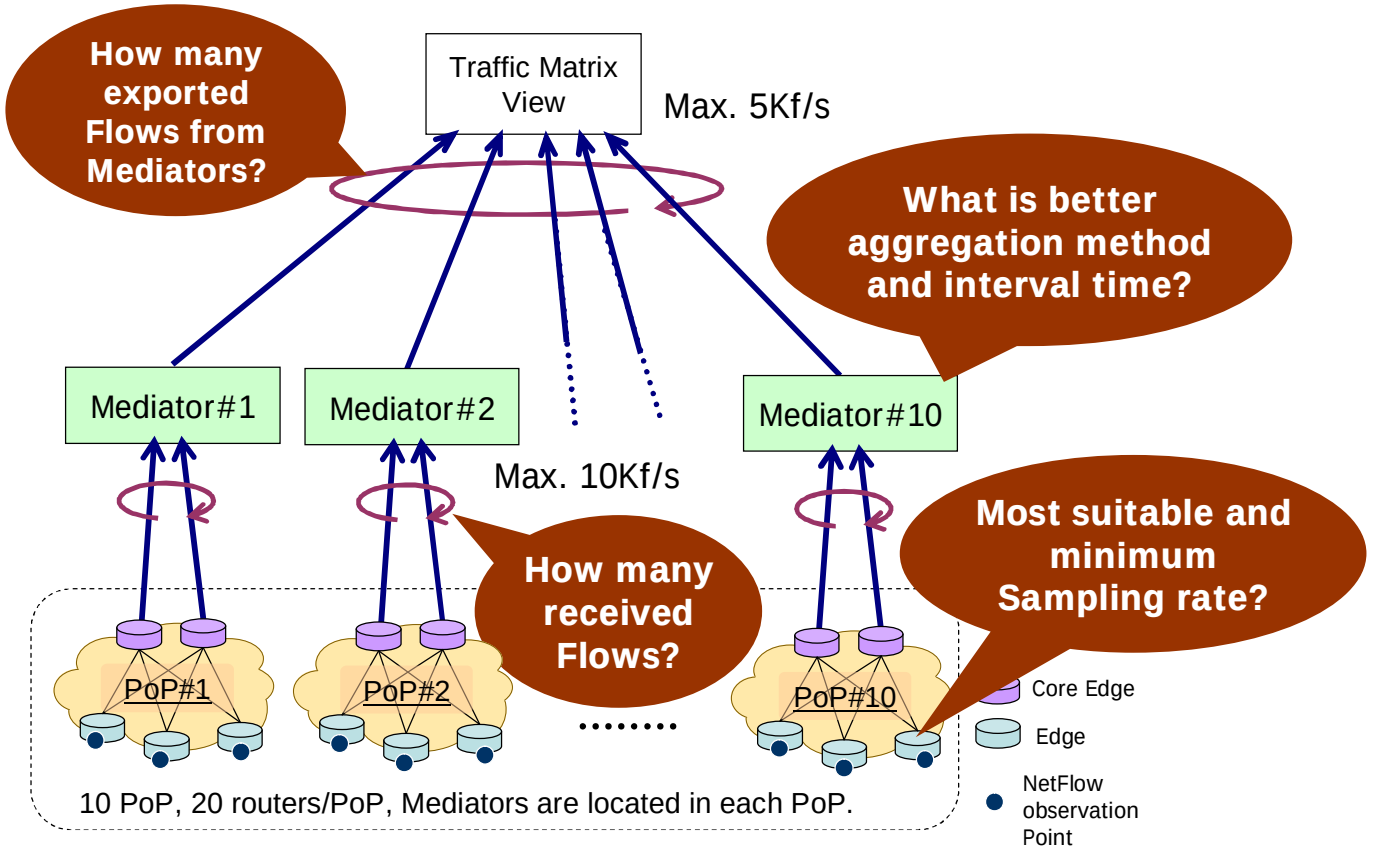
Approach to Design for Hierarchical Model

- To introduce into the real NW, we explored the designing method using Mediator.
  - To design the model, we need to estimate # of Flows roughly.
- Estimate # of the exported Flows from router in sampled flow.
  - How many Flows are reduced by packet sampling?
- Estimate the effect for the aggregation.
  - Aggregation methods (BGP Next-Hop, Prefix, host)
  - Aggregation interval time (20s,60s,90s...)

We tried to estimate # of the received Flow and aggregated Flows based on the MAWI traffic observed international GW.

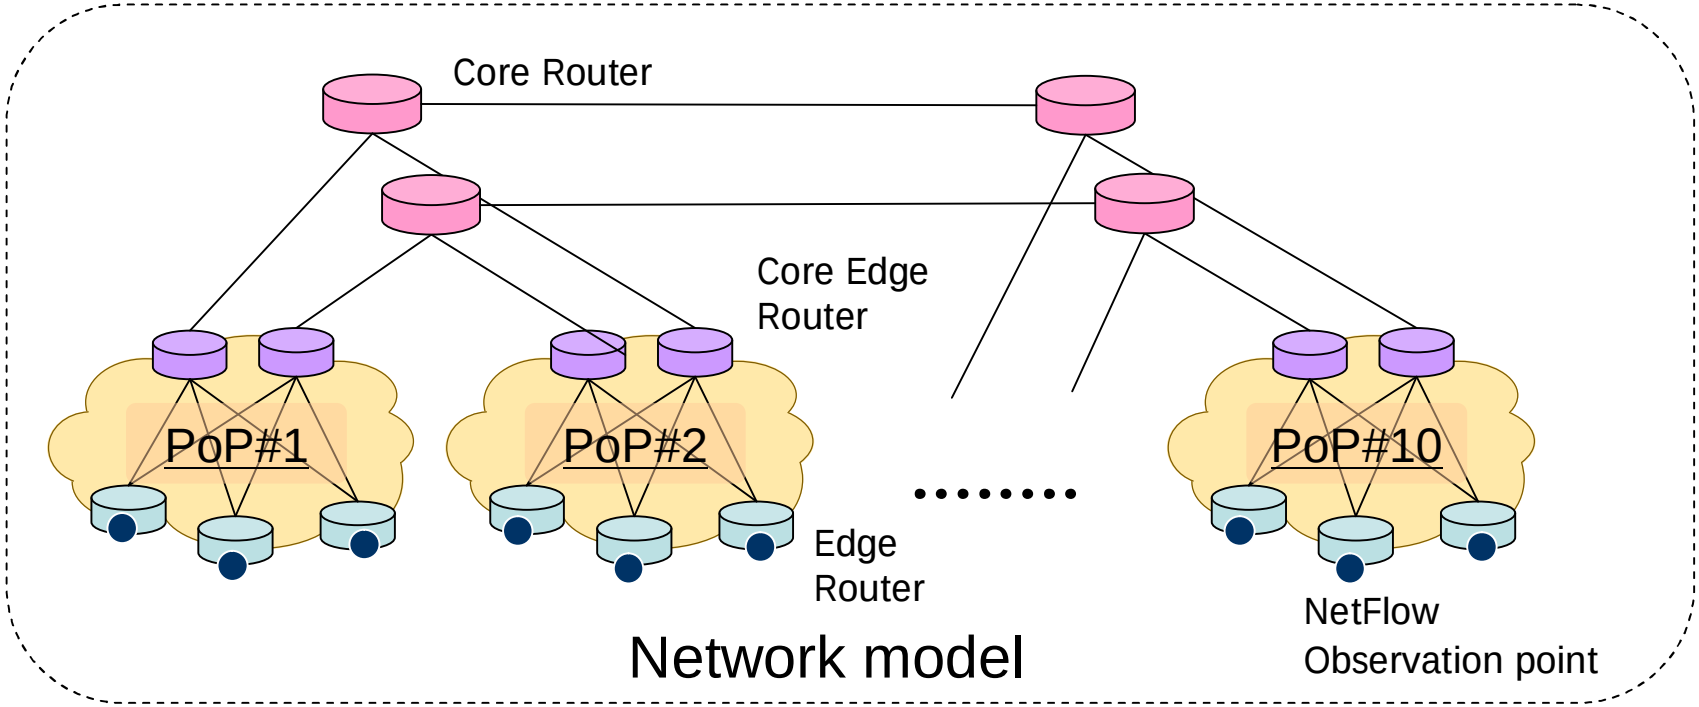
Considerations for Design

- List up the several considerations, as follows.
  - Maximum performances of Top Collector, Mediators are 5 Kf/s, 10 Kf/s.
- Try to explore the aggregation granularity, sampling rate to meet the each performance data.



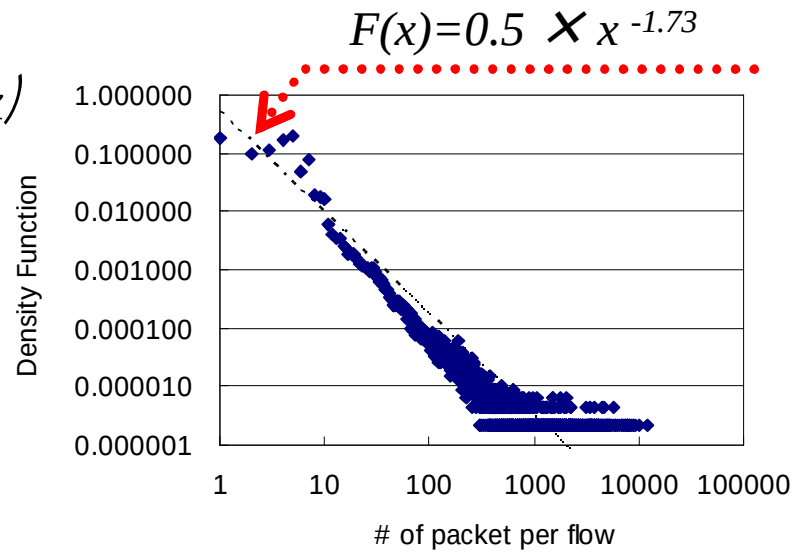
Network Scale Model

- Total traffic volume 500 Gb/s, 100 Mp/s
  - Edge Router 20/PoP×10 PoP=200
  - Core Edge Router 2/PoP×10 PoP =20
  - NetFlow is enabled on the IngressIF of Edge router.



Estimate Flows after Packet Sampling

- Estimate # of the exported flow from router according to the density function of packets per flow.
  - # of packet per flow: x
  - Packets/per flow density function: $F(x)$
  - Sampling rate: $1/r$
  - Total number of unsampled flow: f_{all}



$$f_{sampled} = \sum_{x=1}^{\infty} \left(1-(1-1/r)^x\right) \times F(x) \times f_{all}$$

Extraction probability

$0.5x^{-1.73}$

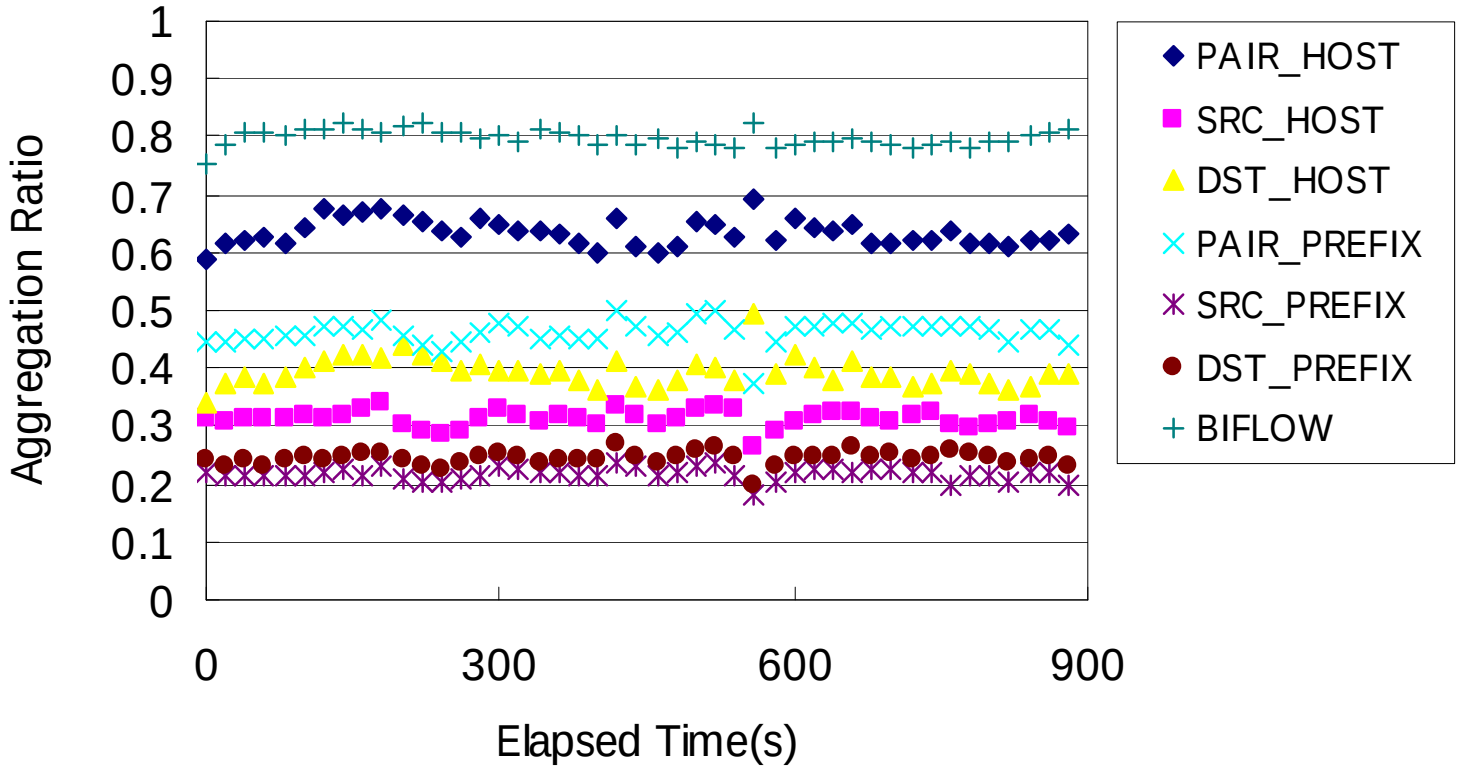
Roughly estimate as follows.
100Mpps ÷ 20packets= 5Mf/s

Approximate # of flow in case of total traffic volume are 500 Gb/s.

Sampling rate	1/100	1/1000	1/10000
$f_{sampled}$	*305 kf/s*	*43 kf/s*	*5.2 kf/s*

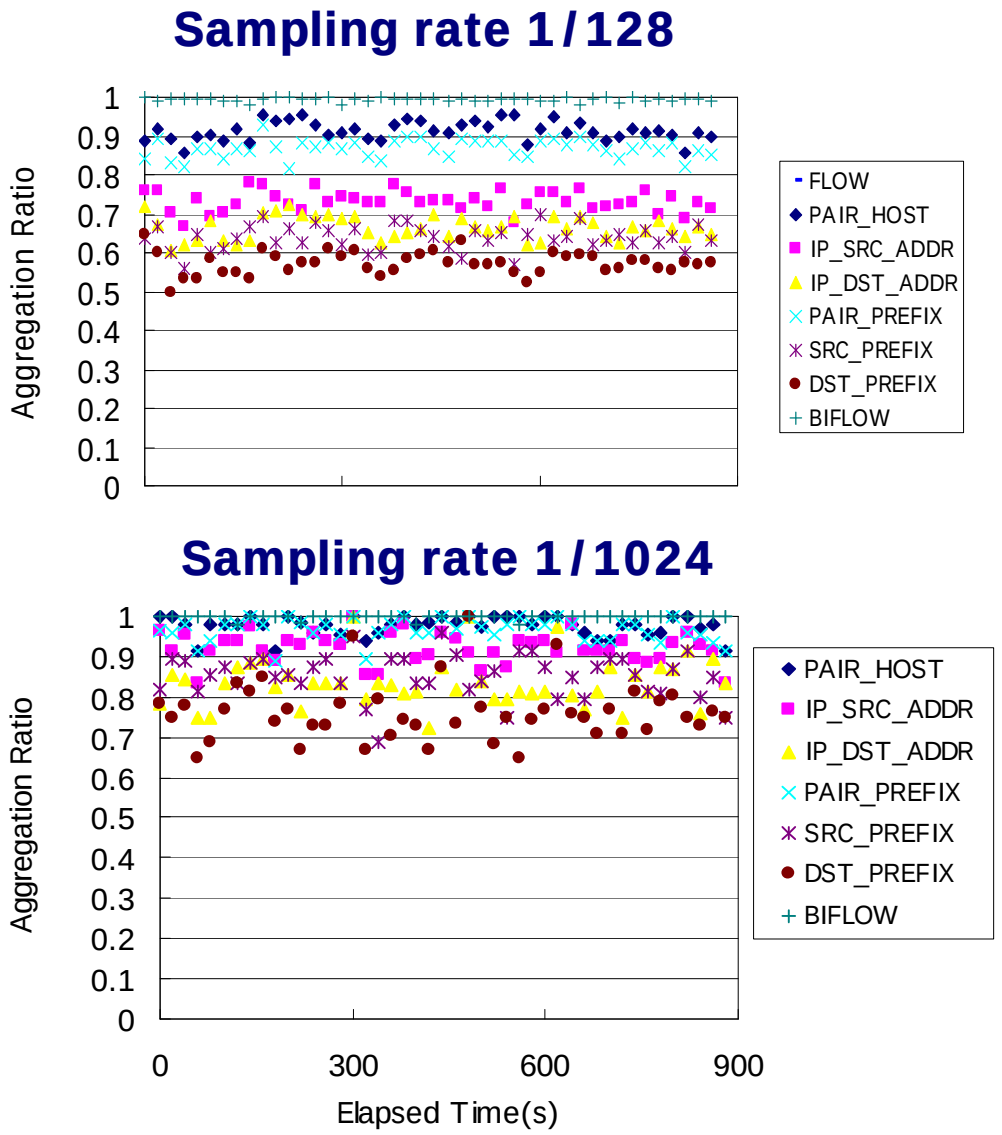
Aggregation Effect in non-sampled NetFlow

- Aggregate non-sampled flows at 20 second interval.
  - The prefix length "/24" is uniformly applied to Prefix Aggregation.
  - Bi-Flow is aggregated from both direction flows.
  - Aggregation Effect: Prefix > HOST> Pair Prefix> Pair HOST> Bi-Flow



Aggregation Effect in sampled NetFlow

- Each Aggregation method become ineffective gradually.
- Bi-Flow becomes ineffective immediately.
  - It is sensitive to sampling rate.



Which factor influences aggregation?

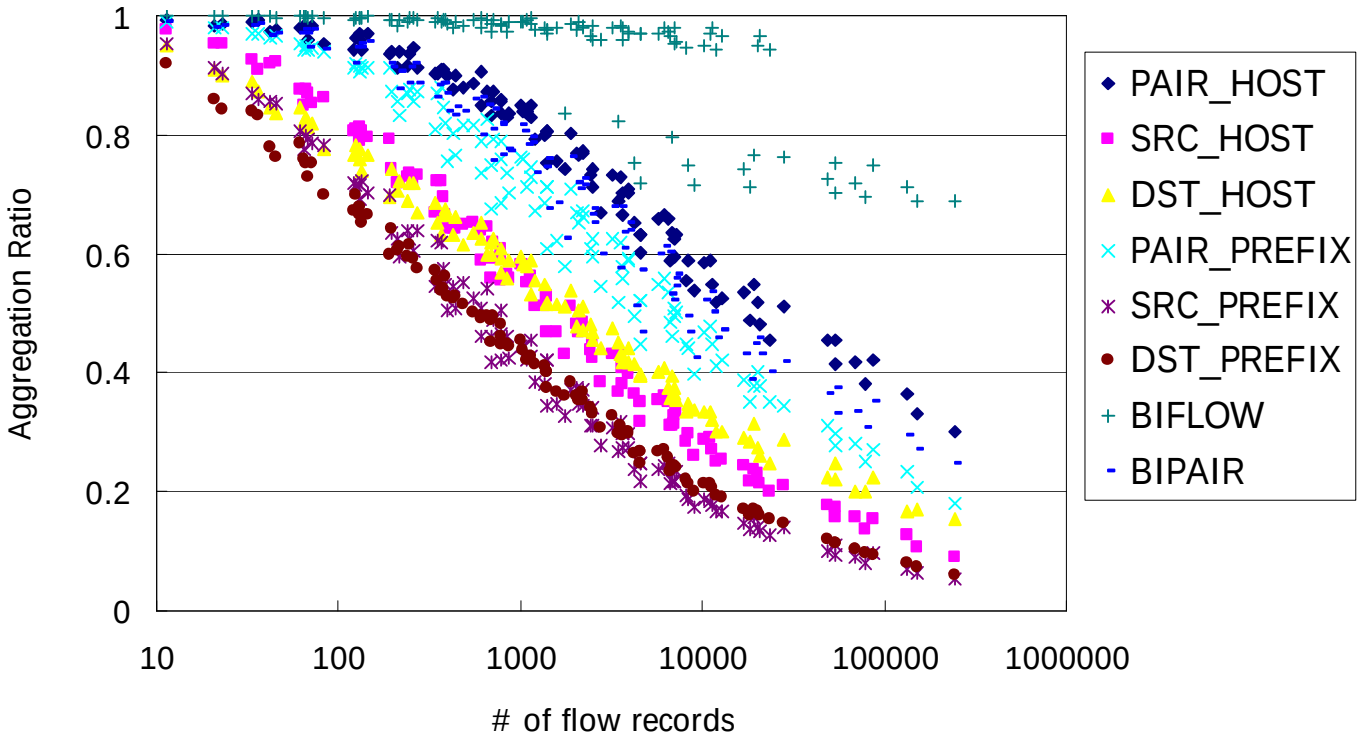
- Aggregation effect depend on the several factor.
  - Traffic Volume through the observation point.
  - Sampling rate
  - Aggregation interval time

But, roughly and simply it depends on # of the received flow between aggregation interval time.

Received Flows	3450	3562
Aggregation Interval Time (s)	10	300
Sampling rate(1/r)	1	128
DST_HOST Aggregation ratio	45%	43%
DST_PREFIX Aggregation ratio	30%	32%

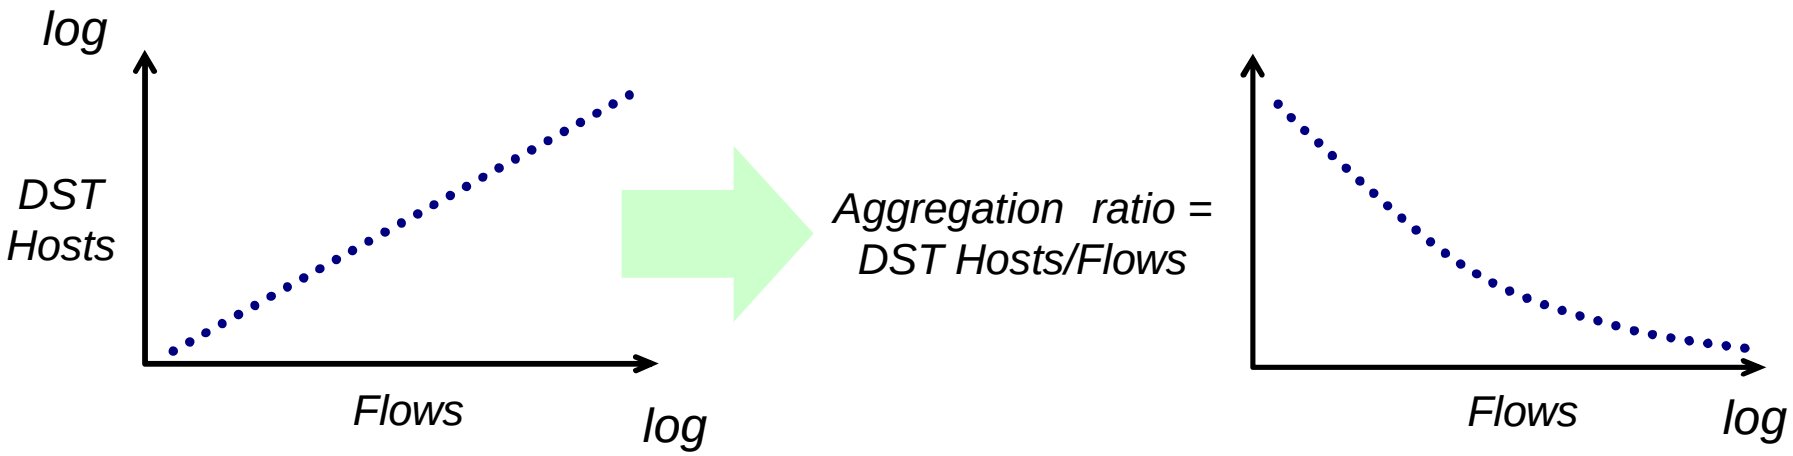
Aggregation Effect against # of Received Flows

- Whole traffic data put in the one graph, such as follows.
  - Traffic data from 3 samples which are different volume.
  - Aggregation Interval time：5s ～ 300s
  - Sampling rate ：1/1 ～ 1/1024



Formulation of Aggregation Effects

- Aggregation ratio(R) can be estimated from # ($f_{received}$) of received flow, as follows.
  - DST Host aggregation: $R_{dsthost} = 1.80 \times f_{received}^{-0.18}$
  - DST Prefix aggregation: $R_{dstprefix} = 2.34 \times f_{received}^{-0.26}$
- After all, it depends on # of unique host or prefix against # of flow.



Flow rate after aggregation (interval=1m)

SamplingRate	1/100	1/1000	1/10000
$f_{sampled}$	305 kf/s	43 kf/s	5.2 kf/s
Received Flow rate per Mediator	30 kf/s	4.4 kf/s	0.6 kf/s

1m Interval

Received Flows at interval time(1m)	918 kflows	132 kflows	18 kflows
DST_HOST aggregation $R = 1.80 \times f_{total}^{-0.18}$	ratio:15% 305×0.15 =45 kf/s	ratio:21% 43 × 0.21 =9.0 kf/s	ratio:31% 5.2 × 0.31 =1.6 kf/s
DST_Prefix aggregation $R = 2.34 \times f_{total}^{-0.26}$	ratio:7% 305×0.07 =21 kf/s	ratio:11% 43 × 0.11 =4.7 kf/s	ratio:18% 5.2 × 0.18 =0.94 kf/s

Flow rate after aggregation (interval=5m)

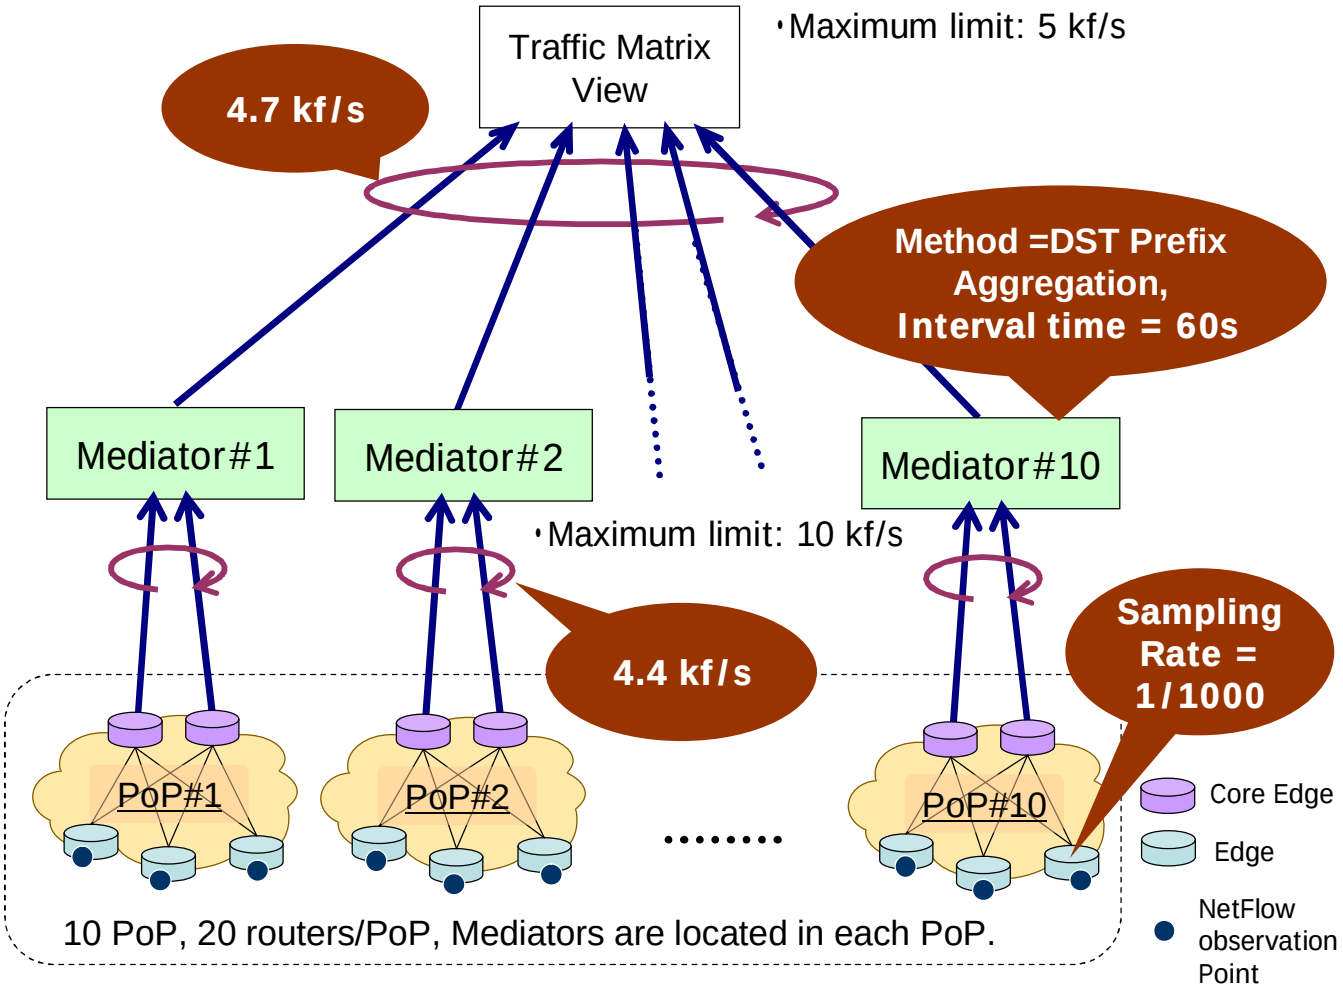
SamplingRate	1/100	1/1000	1/10000
$f_{sampled}$	305 kf/s	43 kf/s	5.2 kf/s
Received Flow rate per Mediator	30 kf/s	4.4 kf/s	0.6 kf/s

5m Interval

Received Flows at interval time(5m)	4.6 Mfows	660 kflows	90 kflows
DST_HOST aggregation $R = 1.80 \times f_{total}^{-0.18}$	ratio:11% 305×0.11 =34 kf/s	ratio:16% 43 × 0.16 =7.0 kf/s	ratio:23% 5.2 × 0.23 =1.2 kf/s
DST_Prefix aggregation $R = 2.34 \times f_{total}^{-0.26}$	ratio:4% 305×0.04 =12 kf/s	ratio:7% 43 × 0.07 =3.0 kf/s	ratio:12% 5.2 × 0.12 =0.62 kf/s

Design of Collection System

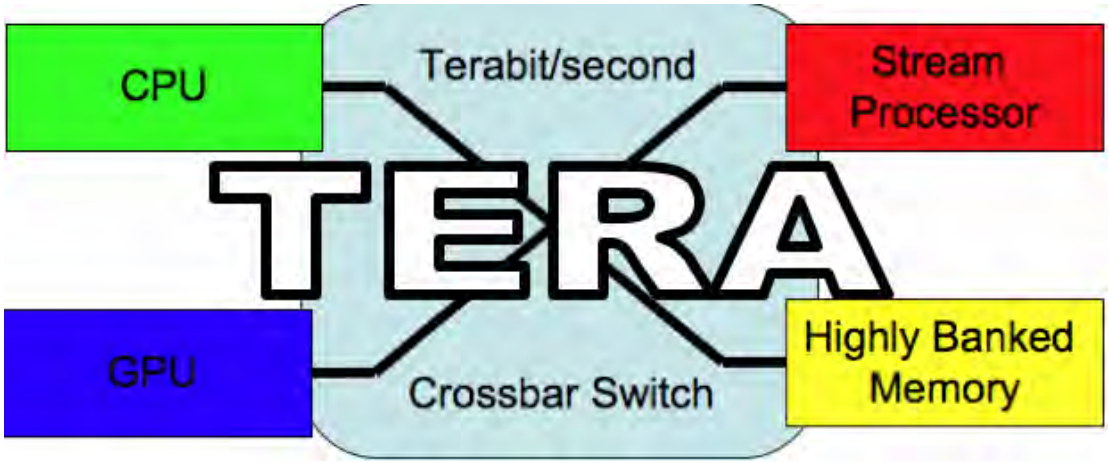
- Sampling Rate : 1/1000
- Aggregation Interval time : 60 s



Conclusion

- Flow Mediation could be easily introduced your NW.
- To maintain the scalability for traffic grows, Flow Mediation is efficient.
- To utilize furthermore the flexibility of aggregation and sampling, using Flow Mediators can control these parameters.
- We can design the hierarchical collection system in large-scale NW.

Thank you for your attention.



On Terabit Flow Analysis

FloCon 2008, Savannah

Jonathan M. Smith

CIS Department, U. Penn

Penn
UNIVERSITY of PENNSYLVANIA



Terabit Network Applications

- Full-fidelity remote visualization and interactive simulation for 80fps HD / 3D HD and beyond, support for holographic visualization
- High-speed sensor data from science experiments
- Immersive simulations and high-fidelity massively multiplayer virtual worlds
- Receive and analyze many concurrent high-fidelity streams of video and/or sensor data - multiple uses in public safety, financial services and other domains

Penn
University of Pennsylvania



Challenges for Flow *Analysis?*

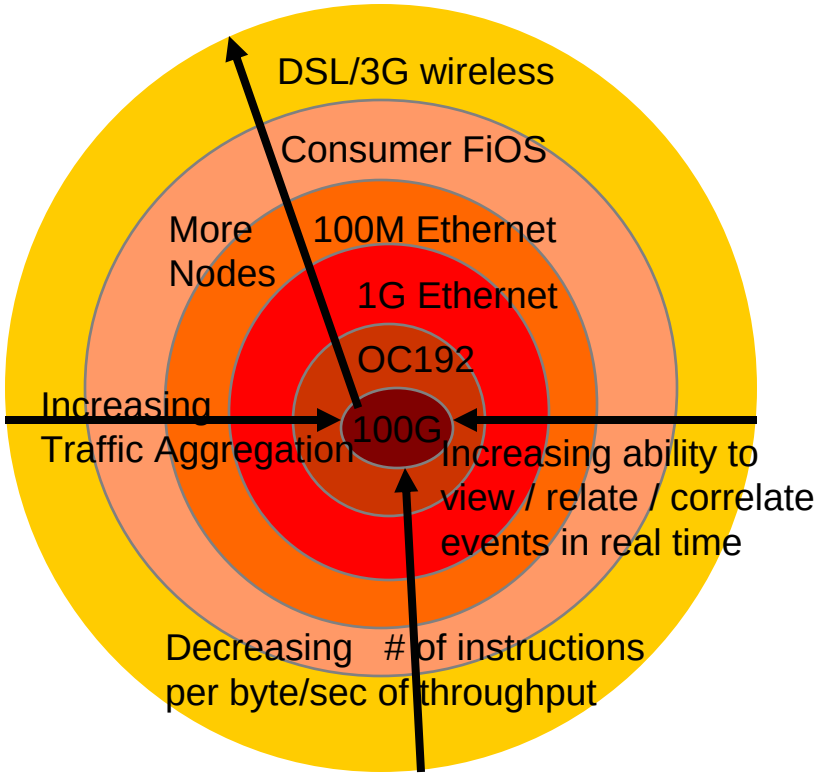
- New kinds of traffic:
  - Extremely High Data Rates
  - Long flows
  - New patterns with P2P and sensors
- Correlation - obtaining the “high ground”
  - Rare events *vs.* attenuated sampling?
- New analysis possible with DPI
- Goal: ingest, record and analyze it *all!*

Penn
University of Pennsylvania



Tradespace: data rates vs. analysis

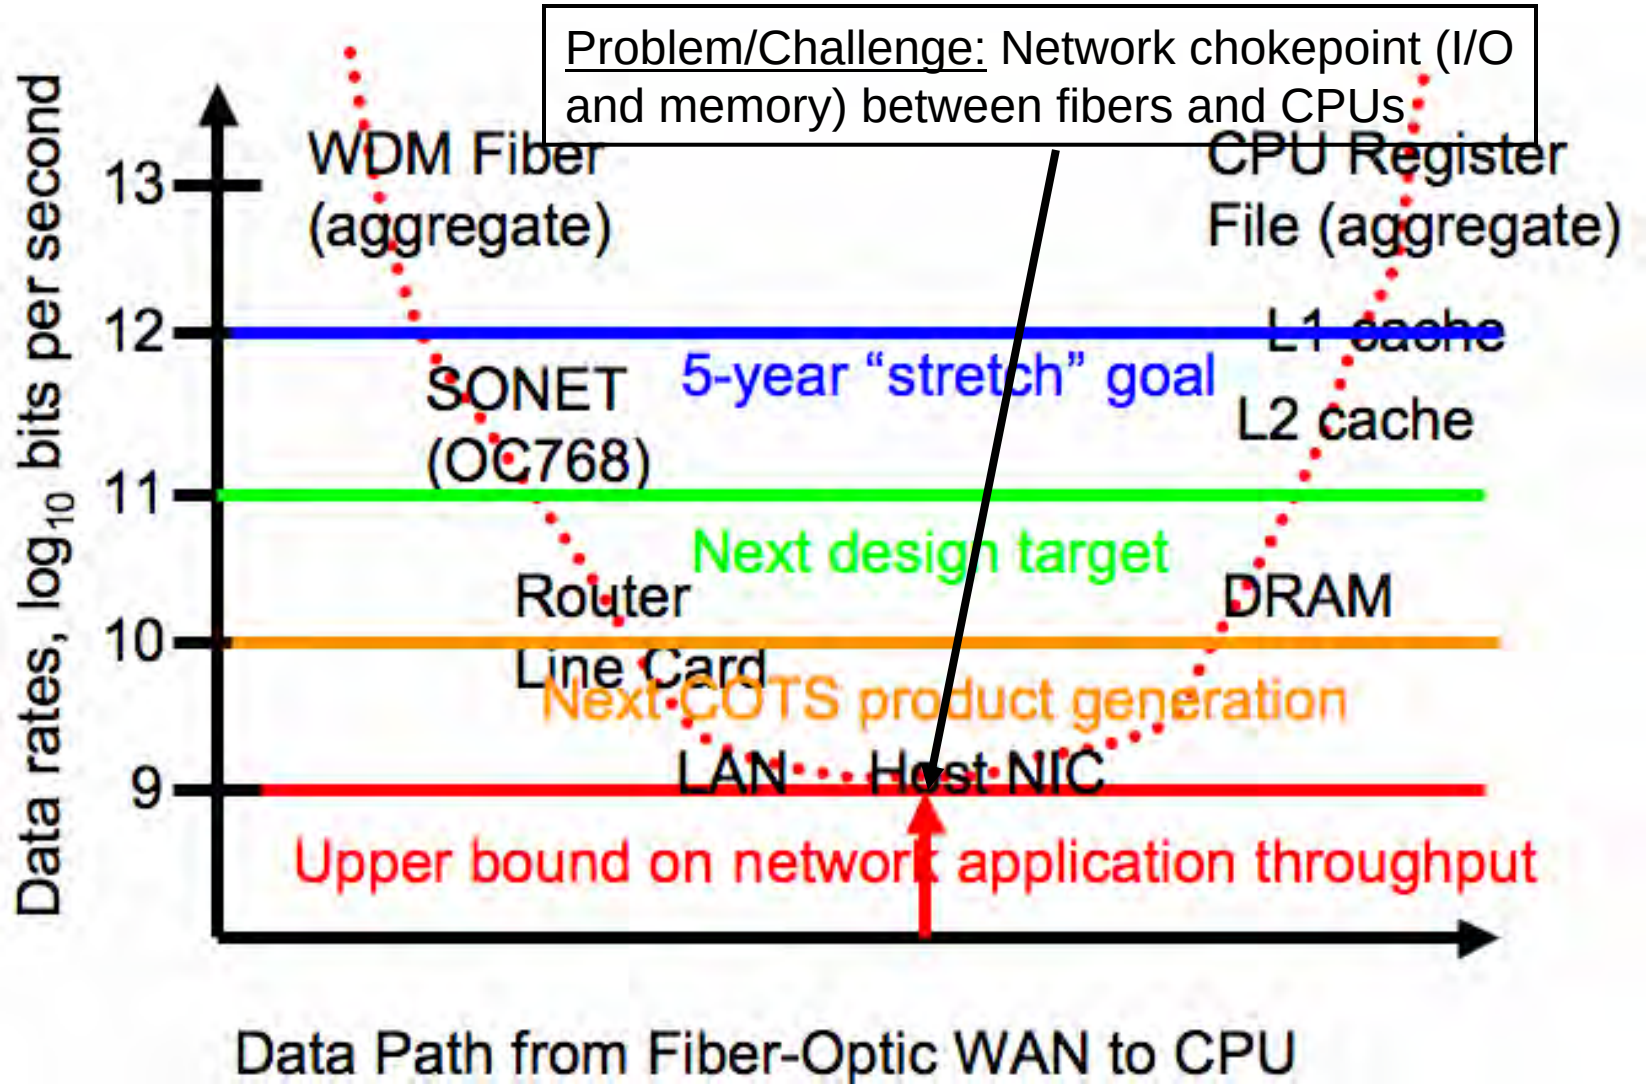
The “high ground”: high aggregation *plus* high data processing rates



Penn
University of Pennsylvania

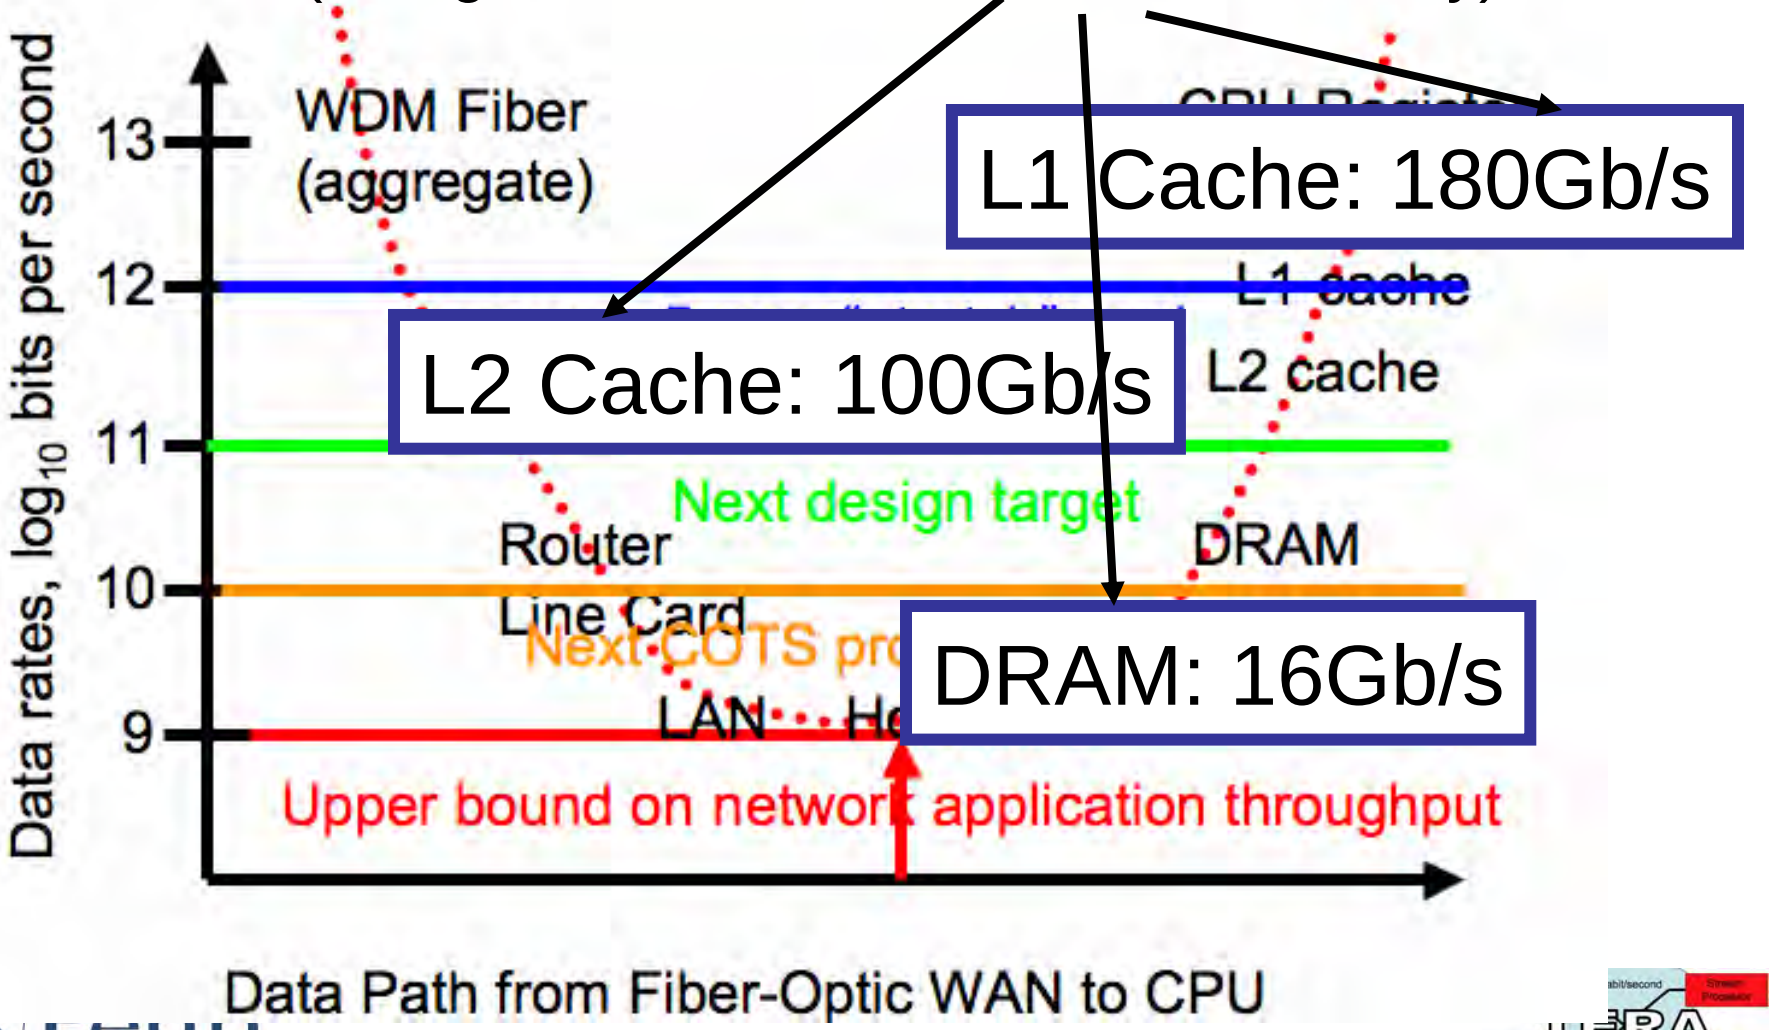
CPU
Terabit/second
Stream Processor
TERA
Crossbar Switch
Highly Banked Memory

The Terabit Chokepoint



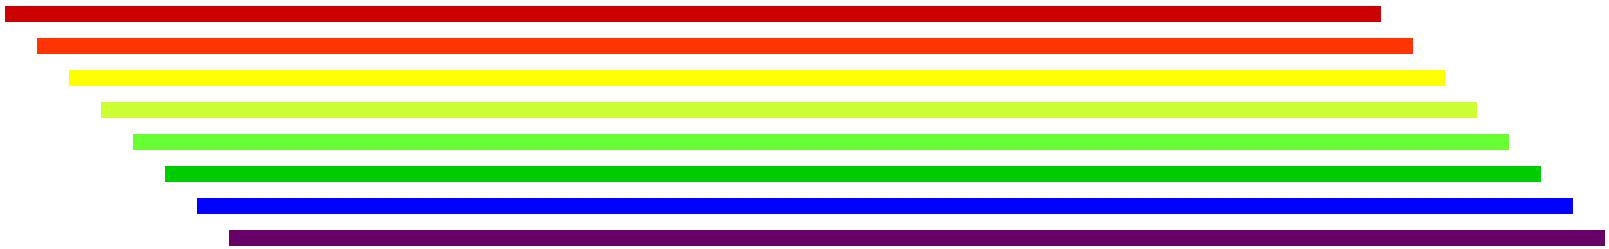
Today's Single-Core PC Performance Measurements

(Using UBUNTU Linux "MEMTEST" utility)



Challenge of Dense Wavelength Division Multiplexing (DWDM)

- Fiber bandwidth is serial bit rate multiplied by number of wavelengths
- E.g., 128*40Gbps in deployed systems (128 lambdas of OC768c SONET)

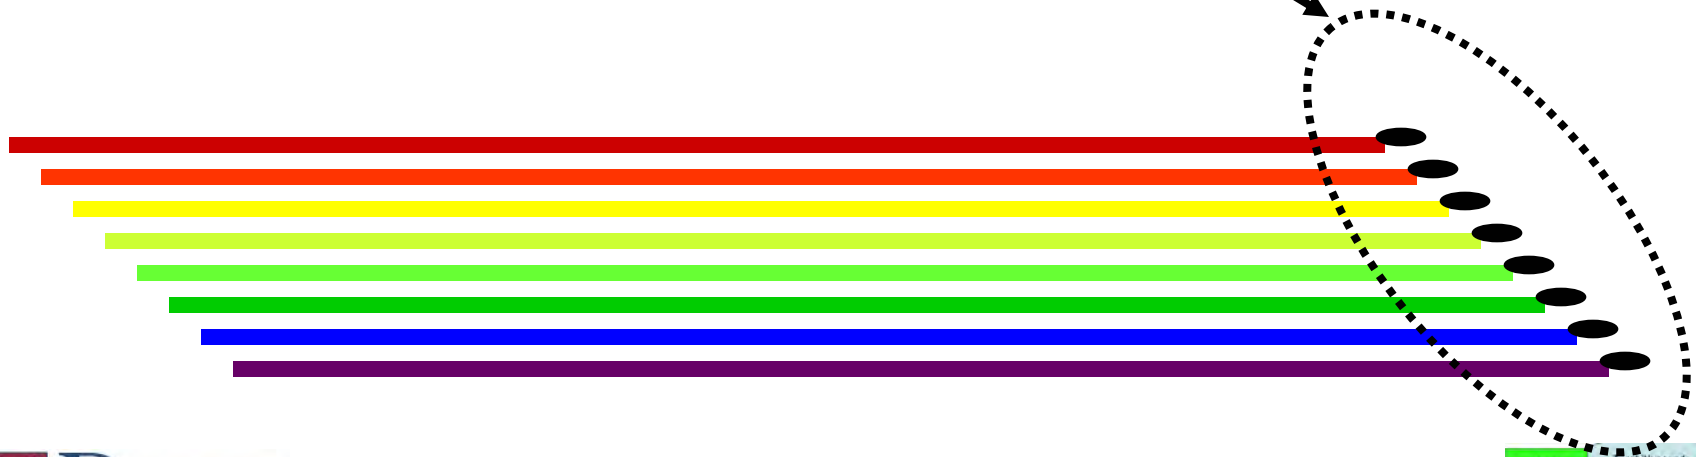


Penn
University of Pennsylvania



Processing Must Scale with Fiber Capacity

- Parallel processing seems necessary
- Memory/processing elements to track line rates and number of wavelengths?

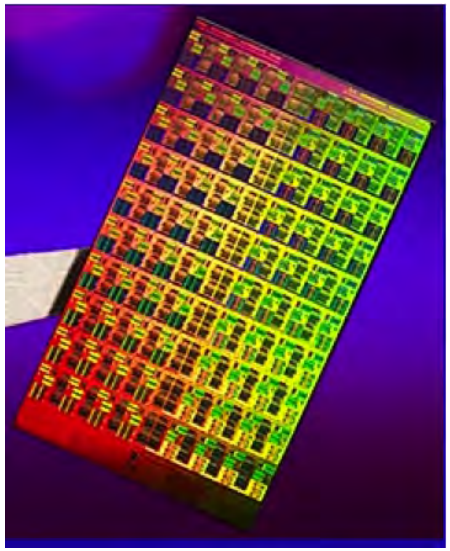


Penn
University of Pennsylvania



Many-Core CPU/GPU Future

- Parallelism floodgate unleashed
  - GPUs and CPUs converging
- Teraflop+ performance in 2009
  - E.g., 32 cores @ 2Ghz
  - 16-element “short” vectors
  - 100 terabit/sec aggregate register bandwidth
  - 1 terabit/sec GDDR3 memory bandwidth
- How do we feed it?

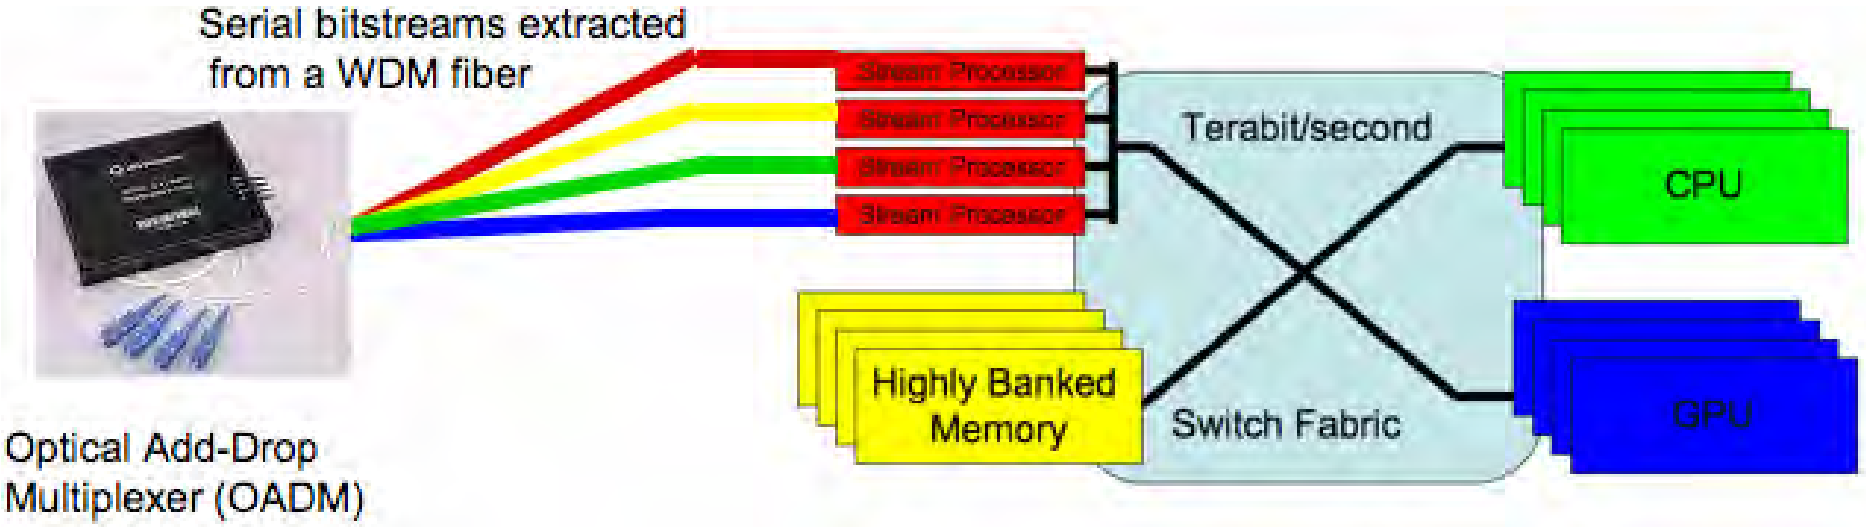


80-core Intel test chip

Penn
University of Pennsylvania



Technical Approach



- Constraints: pins, power, cost
- Switch-based interconnects, parallel paths
  - Direct network/processor interface?
- Stream/graphics engines, banked memories
  - Special high-end pool of DRAMs for NICs?
- New software structures for multicore

Penn
University of Pennsylvania



Components looking good - *architecture* needed

- 1 TB (8 Tbps) memory technologies announced. Fiber good to >10Tb/sec
- 80-1000 cores @ 1-10 Gbps each
- Major challenges: fiber/electronic boundary, data distribution, interconnection network architectures (see, *e.g.*, Dally+Towles)

Penn
University of Pennsylvania



Even *more* processing to scale with fiber capacity?

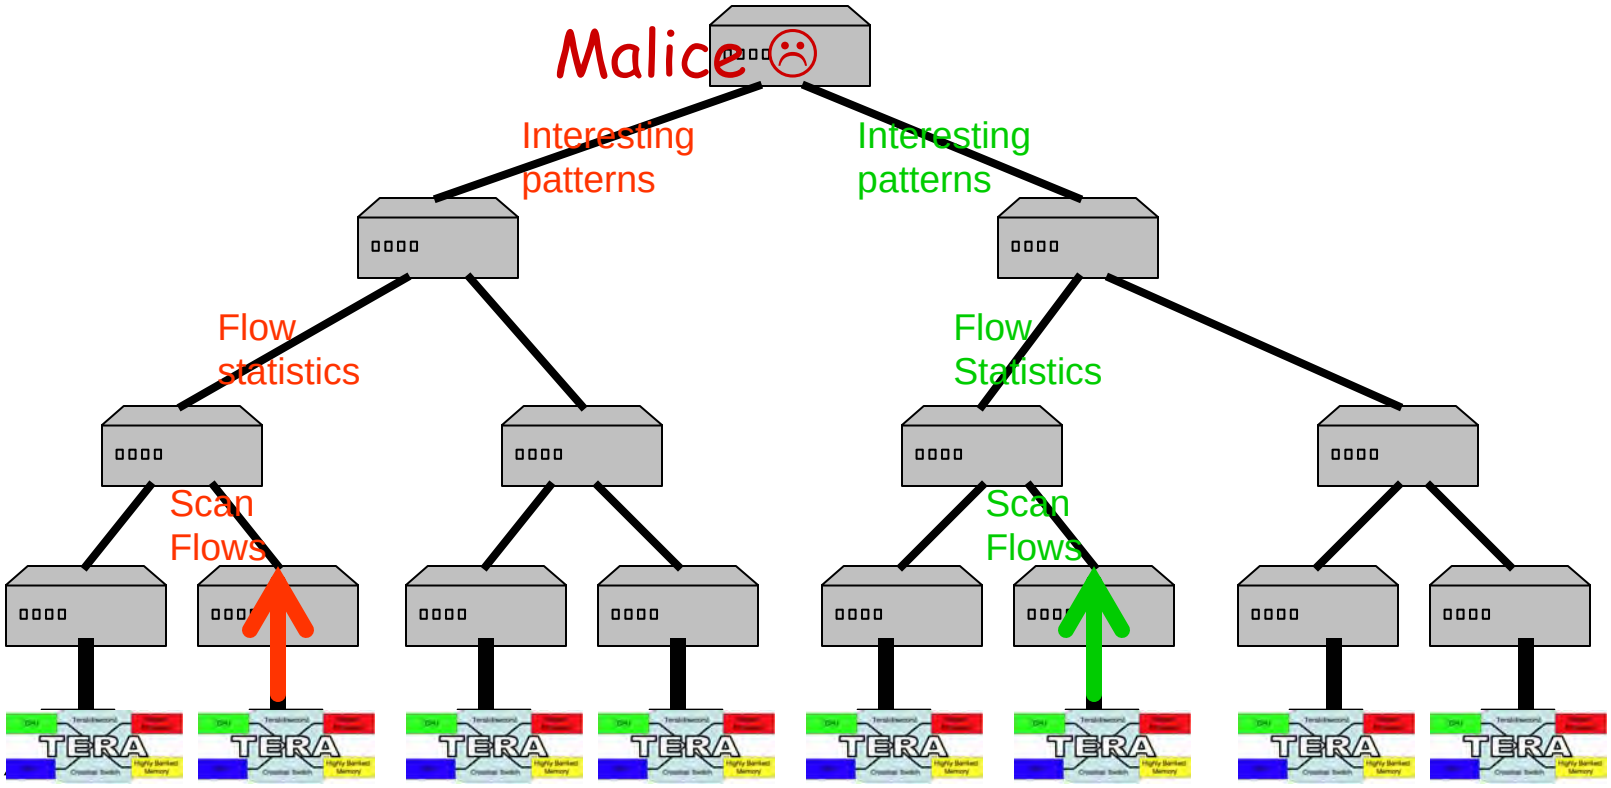
- *Parallel processing* at both multicore (perhaps NPUs?) and "box" level
- Cores track line rates, while degree of "box" parallelism matched against grosser units of wavelengths, *e.g.*, 8:



Penn
University of Pennsylvania



Advanced Broadband Intrusion Detection Engine (ABIDE)



Penn
University of Pennsylvania

TERA
CPU
Terabit/second
Stream Processor
Crossbar Switch
Highly Banked Memory

Help architects to help you

- Computer architects (see Proc. ISCA, Micro, ASPLOS, HPCA, …) evaluate proposals with benchmarks
- Media benchmarks are being developed

  http://euler.slu.edu/~fritts/mediabench/
- Flow analysis needs benchmarks for flow analysis tasks - input side, not just netflow outputs (this is after the fact)

Penn
University of Pennsylvania



Summary

- The future is in parallelism
  - Dense Wavelength Division Multiplexing (DWDM)
  - On-chip networks for multicore
  - Trees for “box”-scale parallelism
- Huge challenges remain
  - Software for new parallelism / media stream analysis; topological choices (*e.g.*, Batcher-Banyan + Crossbar?); load-balancing algorithms
- Need to get flow analysis workloads on computer architecture radar

Penn
University of Pennsylvania

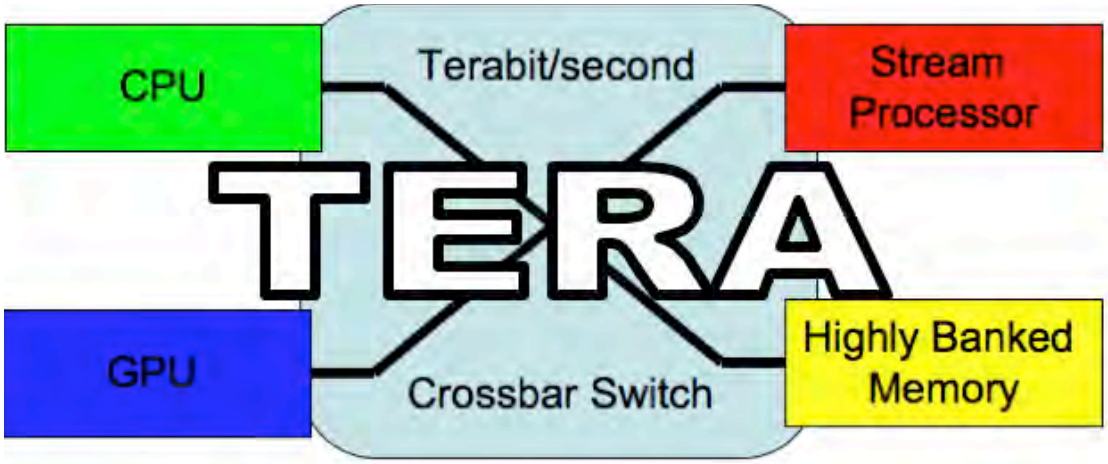


Acknowledgments

- “Terabit Edge Research Activity” (TERA), joint work with Milo M. K. Martin of U. Penn, supported by DARPA/IPTO
- “Advanced Broadband Intrusion Detection Engine” (ABIDE), joint work with M. B. Greenwald and E. Lewis, supported by ARO

Penn
University of Pennsylvania





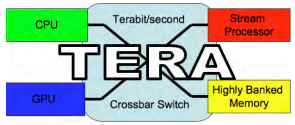
On Terabit Flow Analysis

FloCon 2008, Savannah

Jonathan M. Smith

CIS Department, U. Penn

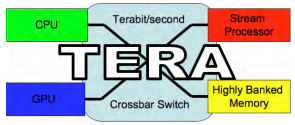
Penn
University of Pennsylvania



Terabit Network Applications

- Full-fidelity remote visualization and interactive simulation for 80fps HD / 3D HD and beyond, support for holographic visualization
- High-speed sensor data from science experiments
- Immersive simulations and high-fidelity massively multiplayer virtual worlds
- Receive and analyze many concurrent high-fidelity streams of video and/or sensor data - multiple uses in public safety, financial services and other domains

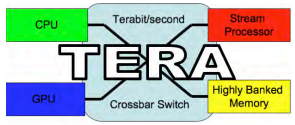
Penn
University of Pennsylvania



Challenges for Flow *Analysis?*

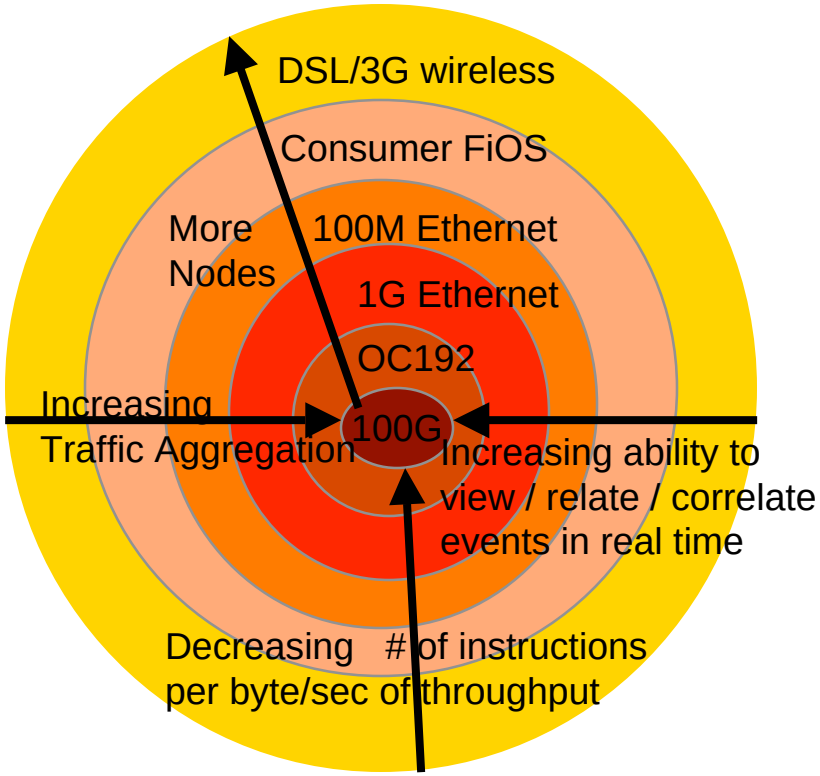
- New kinds of traffic:
  - Extremely High Data Rates
  - Long flows
  - New patterns with P2P and sensors
- Correlation - obtaining the “high ground”
  - Rare events *vs.* attenuated sampling?
- New analysis possible with DPI
- Goal: ingest, record and analyze it *all!*

Penn
University of Pennsylvania



Tradespace: data rates vs. analysis

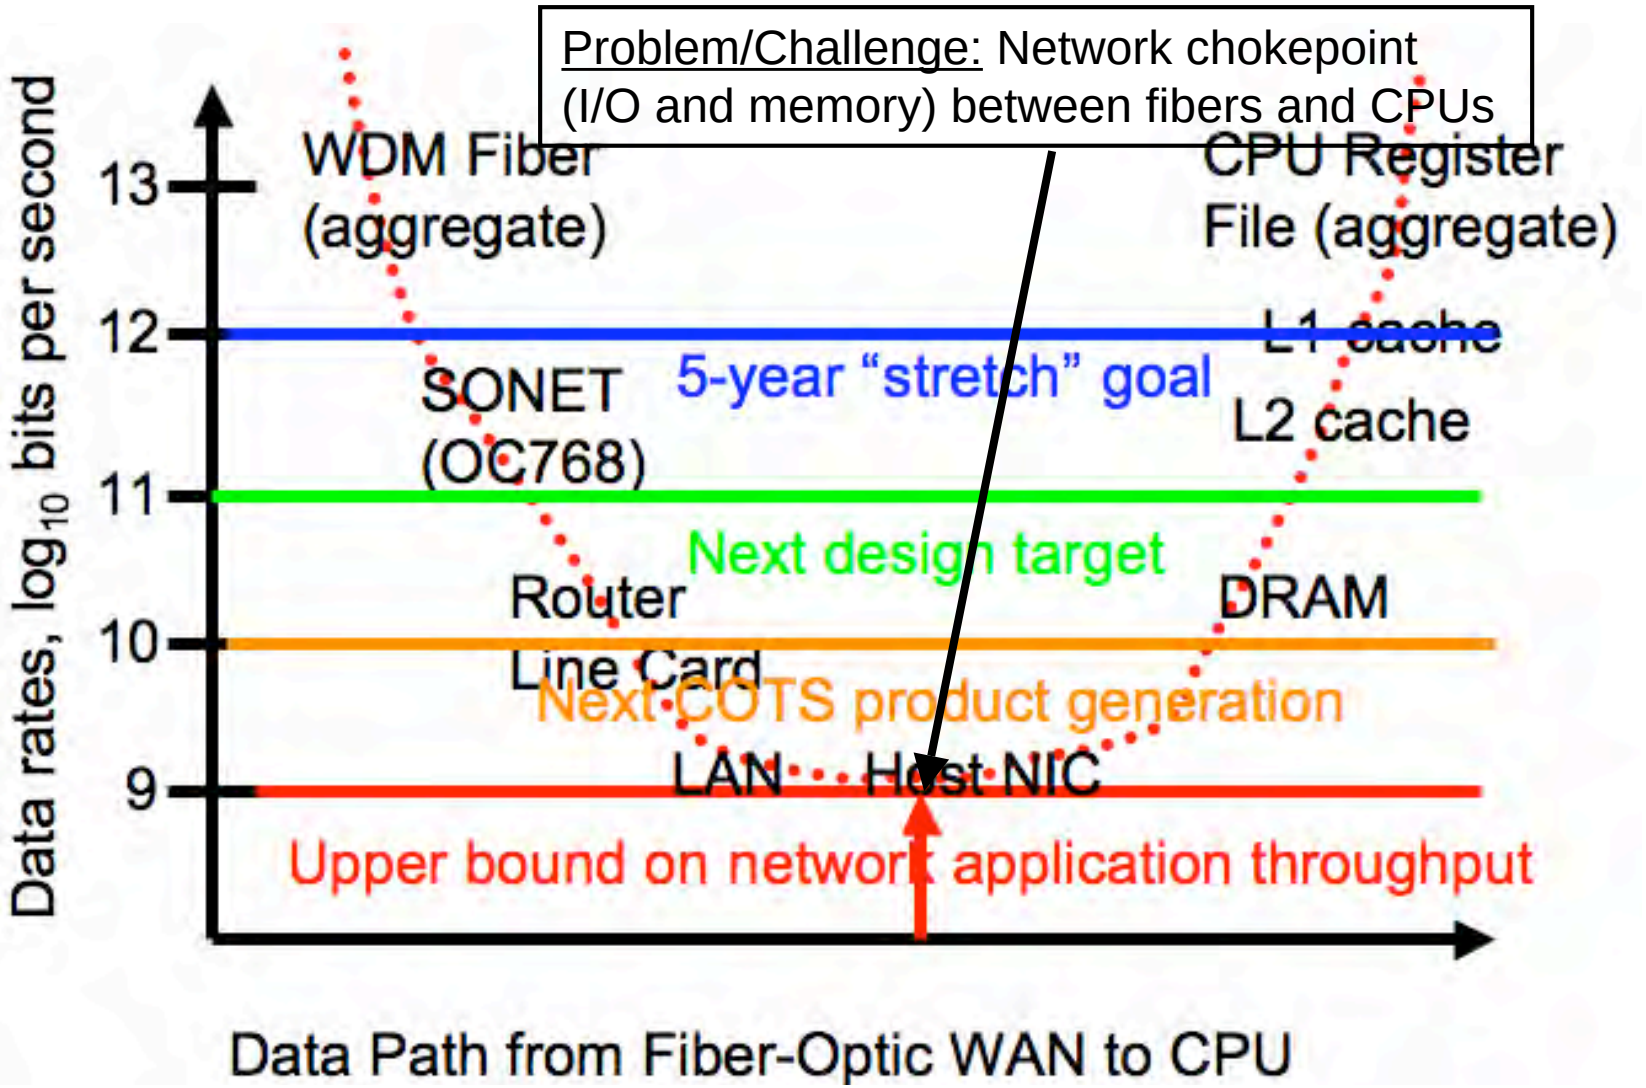
The “high ground”: high aggregation *plus* high data processing rates



Penn
University of Pennsylvania

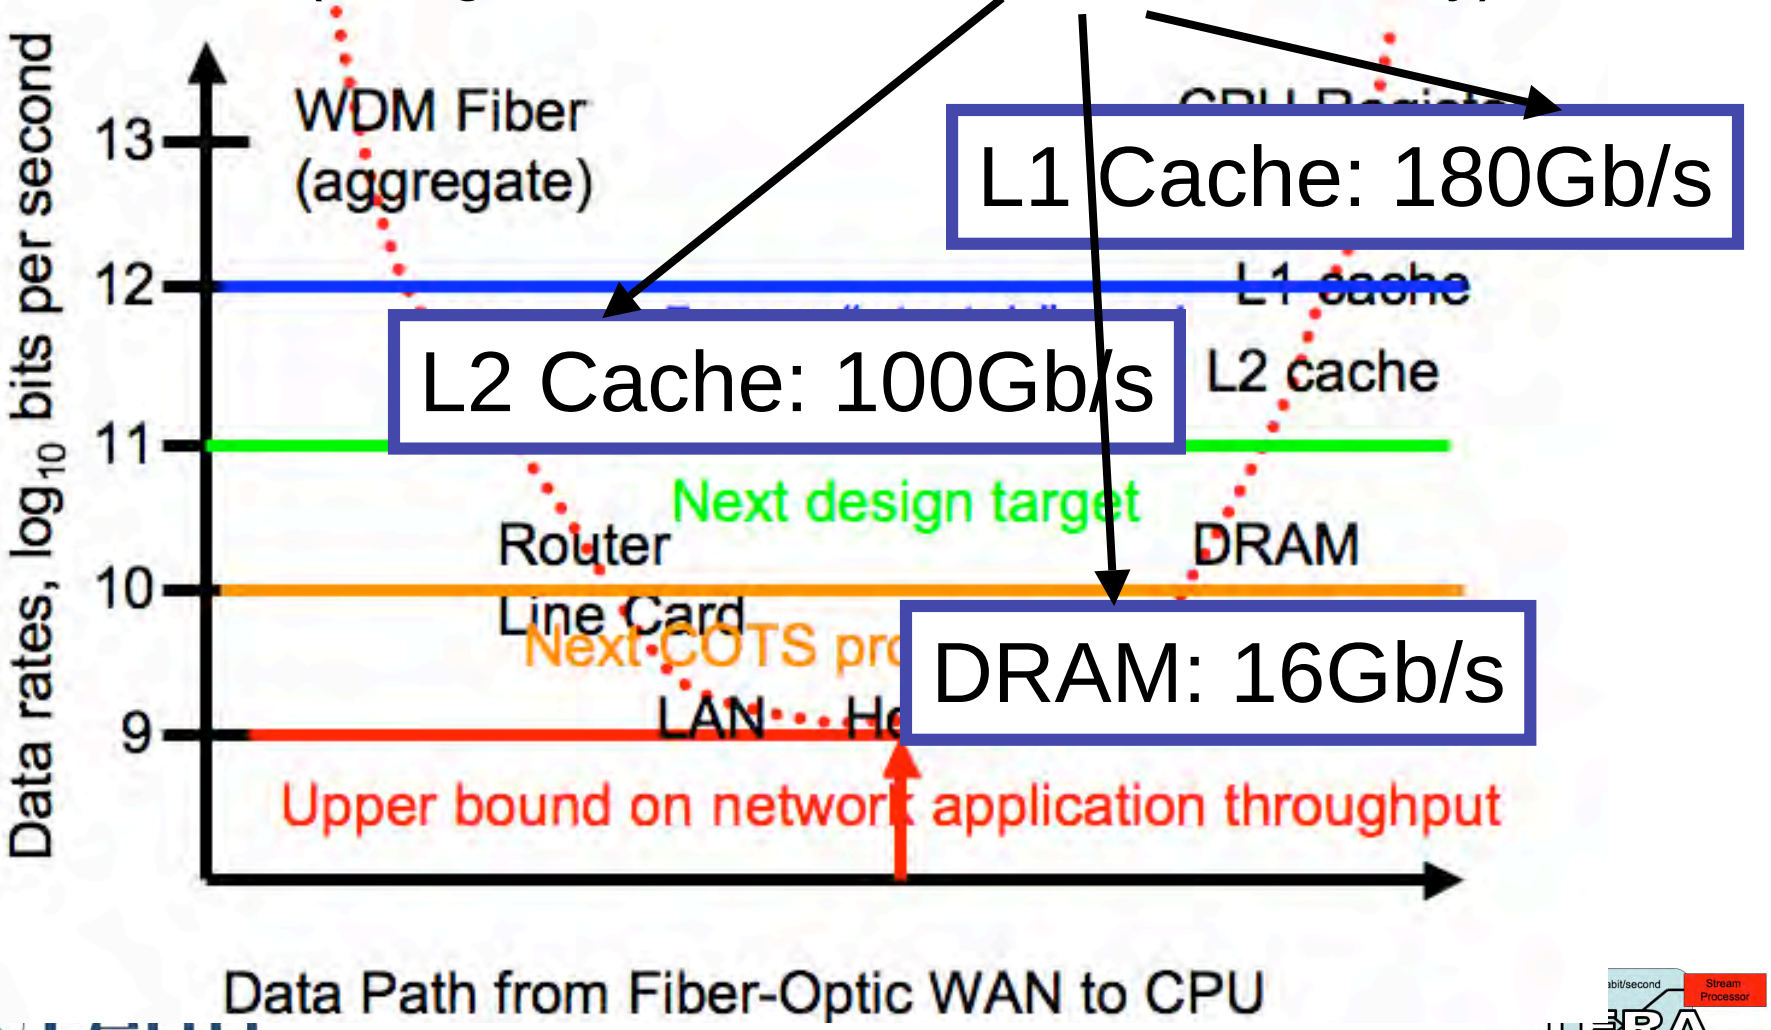
TERA
CPU
Terabit/second
Stream Processor
GPU
Crossbar Switch
Highly Banked Memory

The Terabit Chokepoint



Today's Single-Core PC Performance Measurements

(Using UBUNTU Linux "MEMTEST" utility)

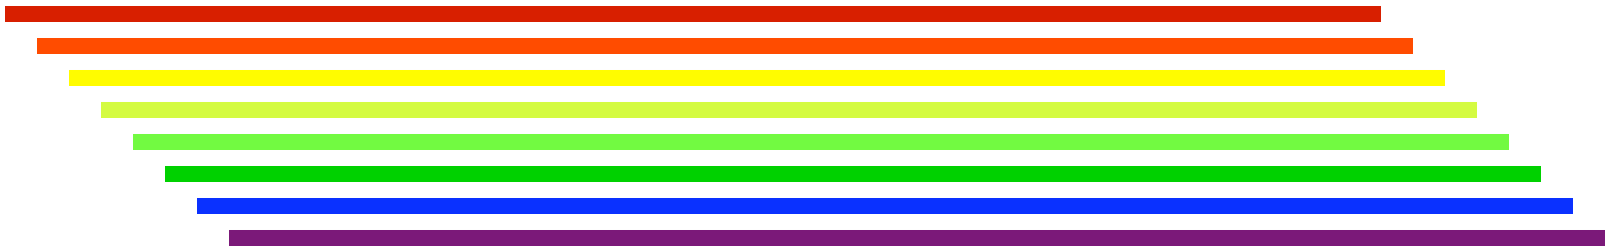


PENN
University of Pennsylvania

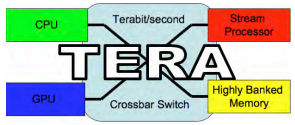
Stream Processor
TERA
Highly Banked Memory
GPU
Crossbar Switch

Challenge of Dense Wavelength Division Multiplexing (DWDM)

- Fiber bandwidth is serial bit rate multiplied by number of wavelengths
- E.g., 128*40Gbps in deployed systems (128 lambdas of OC768c SONET)

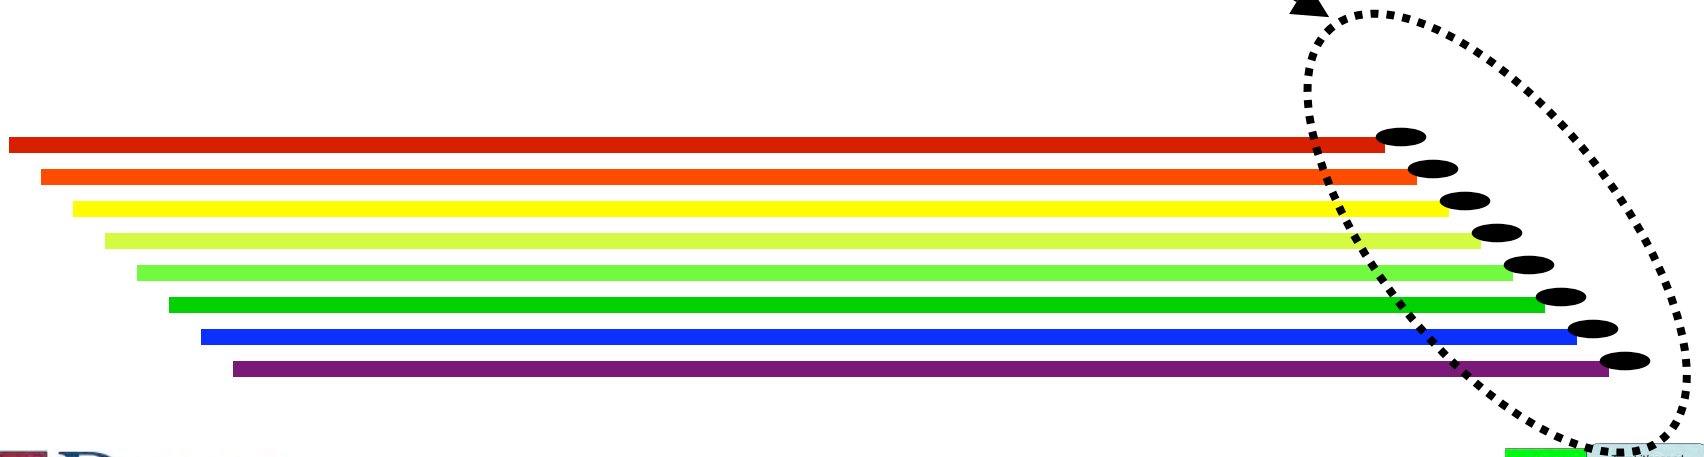


Penn
University of Pennsylvania

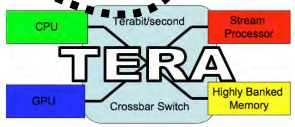


Processing Must Scale with Fiber Capacity

- Parallel processing seems necessary
- Memory/processing elements to track line rates and number of wavelengths?

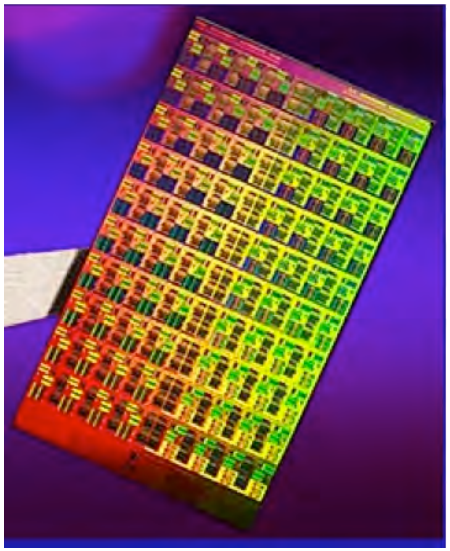


Penn
University of Pennsylvania



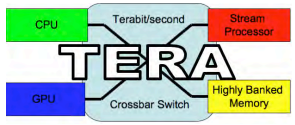
Many-Core CPU/GPU Future

- Parallelism floodgate unleashed
  - GPUs and CPUs converging
- Teraflop+ performance in 2009
  - E.g., 32 cores @ 2Ghz
  - 16-element “short” vectors
  - 100 terabit/sec aggregate register bandwidth
  - 1 terabit/sec GDDR3 memory bandwidth
- How do we feed it?

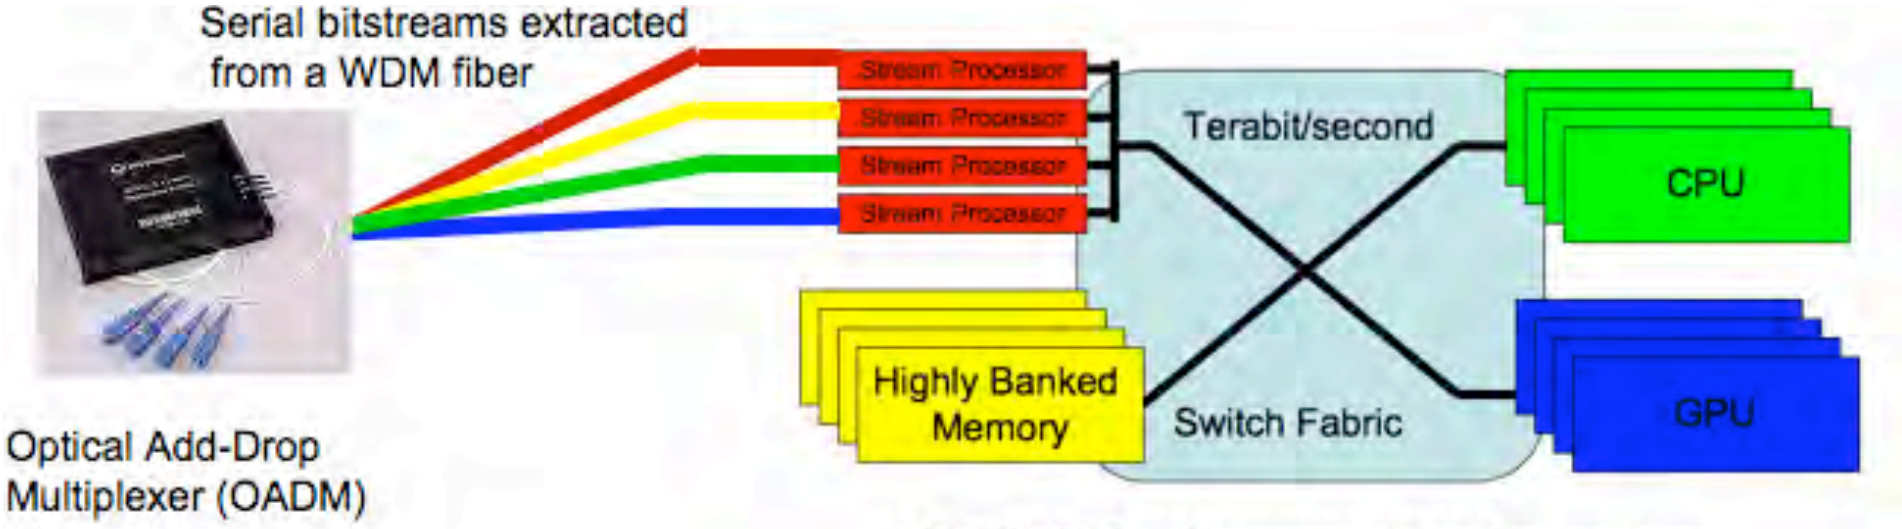


80-core Intel test chip

Penn
University of Pennsylvania

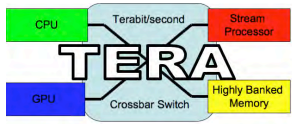


Technical Approach



- Constraints: pins, power, cost
- Switch-based interconnects, parallel paths
  - Direct network/processor interface?
- Stream/graphics engines, banked memories
  - Special high-end pool of DRAMs for NICs?
- New software structures for multicore

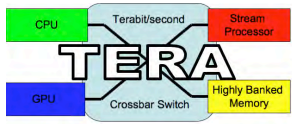
Penn
University of Pennsylvania



Components looking good - *architecture* needed

- 1 TB (8 Tbps) memory technologies announced. Fiber good to >10Tb/sec
- 80-1000 cores @ 1-10 Gbps each
- Major challenges: fiber/electronic boundary, data distribution, interconnection network architectures (see, *e.g.,* Dally+Towles)

Penn
University of Pennsylvania

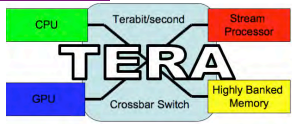


Even *more* processing to scale with fiber capacity?

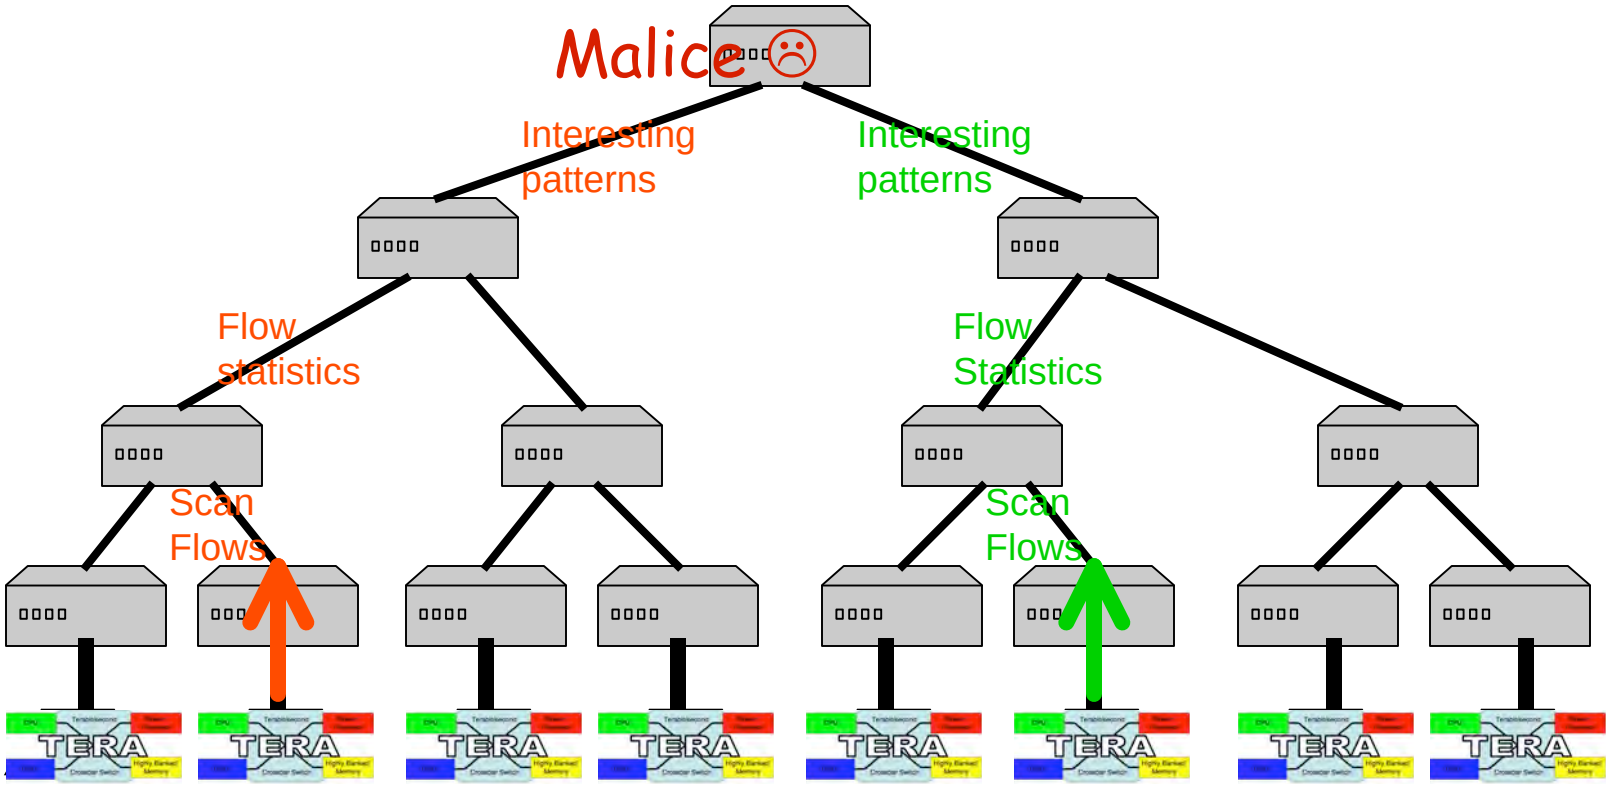
- *Parallel processing* at both multicore (perhaps NPUs?) and “box” level
- Cores track line rates, while degree of “box” parallelism matched against grosser units of wavelengths, *e.g.*, 8:



Penn
University of Pennsylvania



Advanced Broadband Intrusion Detection Engine (ABIDE)



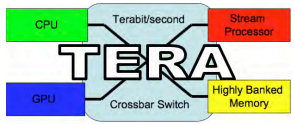
Penn
University of Pennsylvania

CPU
Terabit/second
Stream Processor
TERA
GPU
Crossbar Switch
Highly Banked Memory

Help architects to help you

- Computer architects (see Proc. ISCA, Micro, ASPLOS, HPCA, …) evaluate proposals with benchmarks
- Media benchmarks are being developed `http://euler.slu.edu/~fritts/mediabench/`
- Flow analysis needs benchmarks for flow analysis tasks - input side, not just netflow outputs (this is after the fact)

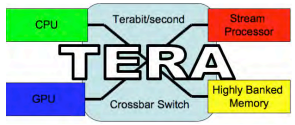
Penn
University of Pennsylvania



Summary

- The future is in parallelism
  - Dense Wavelength Division Multiplexing (DWDM)
  - On-chip networks for multicore
  - Trees for “box”-scale parallelism
- Huge challenges remain
  - Software for new parallelism / media stream analysis; topological choices (*e.g.*, Batcher-Banyan + Crossbar?); load-balancing algorithms
- Need to get flow analysis workloads on computer architecture radar

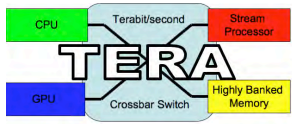
Penn
University of Pennsylvania



Acknowledgments

- “Terabit Edge Research Activity” (TERA), joint work with Milo M. K. Martin of U. Penn, supported by DARPA/IPTO
- “Advanced Broadband Intrusion Detection Engine” (ABIDE), joint work with M. B. Greenwald and E. Lewis, supported by ARO

Penn
University of Pennsylvania



Improvement of Processes for Flow Information

NTT Network Service System Laboratories, NTT Corporation

Hitoshi Irino,
Masaru Katayama

NTT Network System Laboratories

Abstract of this presentation

- Ideas for increasing (optimizing) performances of processes in IPFIX
- Ideas based on all processes using **an order rule of Information Elements/fields**
- These ideas are introduced:
  - Method for **reducing the number of comparisons** between an existing flow and an incoming new packet **in Metering Processes** (MPs) (**Comparison method for multiple fields in MPs**)
  - Method for **reducing the number of copies** of flow records from Metering Process to **Exporting Processes** (EPs) with a predefined order of fields (**Copy method for multiple fields in EPs**)
  - Method for **increasing processing speed for storing** data in incoming packets to file with a predefined format of **Collecting Processes** (CPs) (**Copy method for multiple fields in CPs**)

These are basically the same.

Motivation of this research

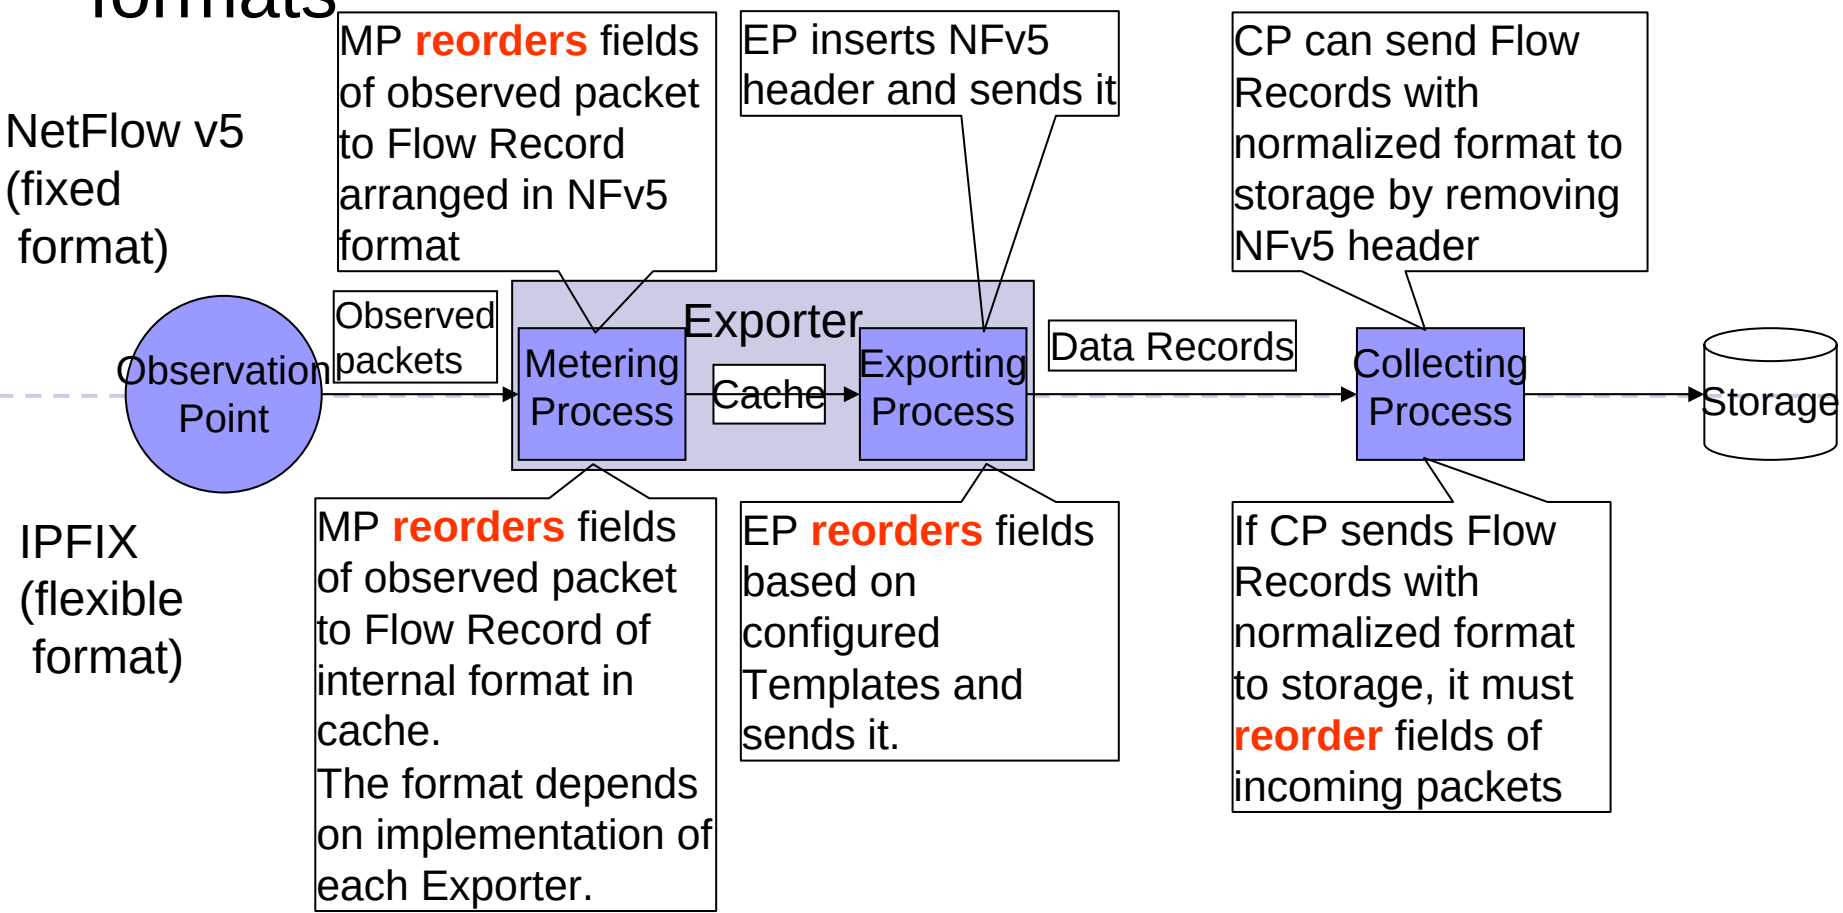
- Background
  - Network bandwidth will continue to increase.
  - IPFIX will be a standard protocol for flow information exchange.
    - Network bandwidth will become broader-band.
      - Use a lower sampling rate.
      - Use fewer Flow Keys.

→ However, flow information will become less accurate.

→ Research on increasing (optimizing) the performances of IPFIX processes

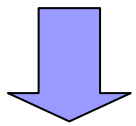
IPFIX features

- IPFIX
  - Advantage: Uses Template-based flexible flow export
  - Disadvantage: More complex than fixed-format protocol
- Comparison of processes between flexible and fixed formats

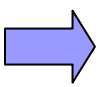


Our approach: **Making the order rule for Information Elements**

- Processes of IPFIX have a high possibility of reordering fields.
  - Reducing the cost of reordering fields can improve their performance.



- Our approach
  - Make the order rule for Information Elements
    - Order rule gives IPFIX processes chances to process multiple fields.
    - Processing multiple fields at a time achieves higher performance than processing one field at a time.
    - The rule does not influence the flexibility of IPFIX.



If a unified order rule of fields/IEs is defined, reordering costs can be reduced.

Idea of order

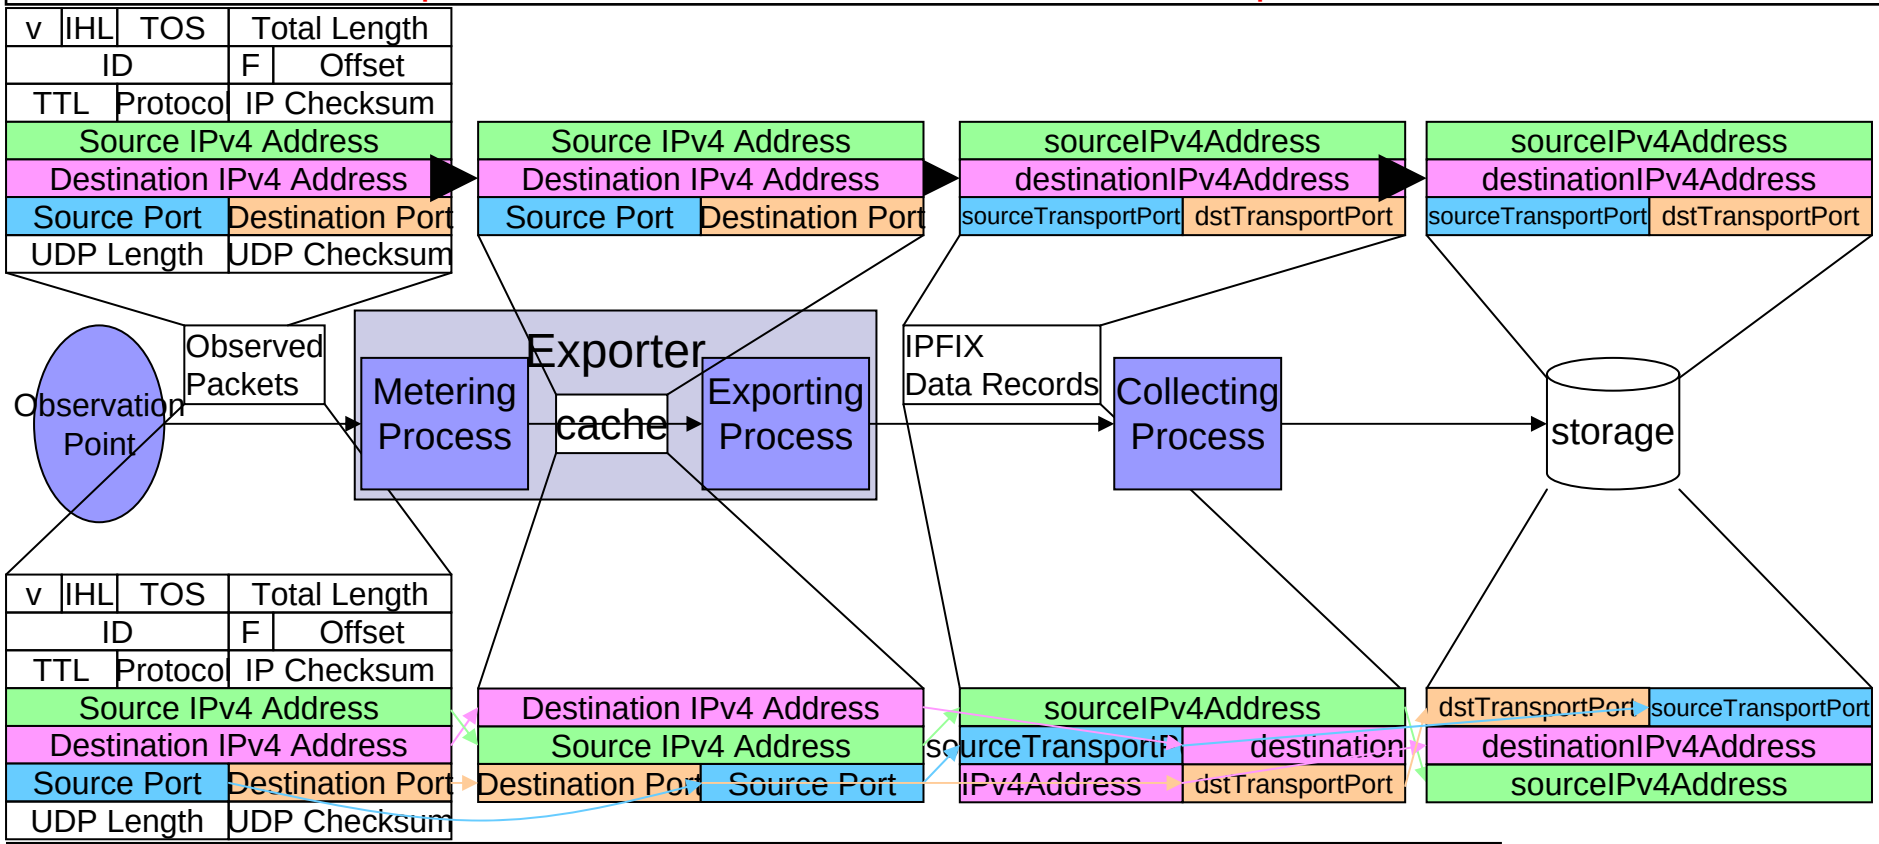
- Idea of order:
  - MPs, EPs and CPs place fields (IEs) in the same order, so it is highly likely that multiple fields will be processed at a time.
    - This reduces reordering costs.
- Order recommended in this presentation
  - Place fields in observed packets in order of protocol header.
  - **Therefore, order of IEs that refer to packets and header fields is recommended.**

	Metering Processes	Exporting Processes	Collecting Processes
Input	Observed packets (network byte order)	Their caches	IPFIX Data Record (network byte order)
Output	(Storing) their caches	IPFIX Data Record (network byte order)	(Storing) files, their DB (real-time analysis)

Example of using same order in MP, EP and CP

Flow Keys: sourceIPv4Address, destinationIPv4Address, sourceTransportPort, destinationTransportPort

Good (ideal) case: Same suggested order,
which refers order of packet header fields used in the cache in Exporter and IPFIX data records



Bad case: Different order used in the cache in Exporter and IPFIX data records

- **If the referential order, which refers to the order of packet fields, is defined, it could, in some cases, lead to increased performance.**
- **If a referential order is undefined, there is no possibility of increased performance.**

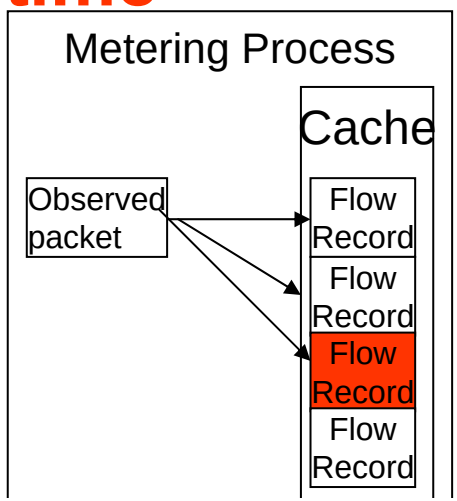
1st idea to improve performance
in environment in which MP, EP, and CP use the same order

Comparison method for multiple fields in Metering Processes (MPs)

NTT Network Service System Laboratories, NTT Corporation

Comparison method for multiple fields in MP (1)

- MP must repeat comparison between existing Flow Records in its cache and new observed packet.
  - To judge whether the new packet belongs to a new flow or an existing one.
- Basically, in this comparison, all fields (IEs) serving as Flow Keys are compared every time.
- **If fields of Flow Records are placed in the same order as packet header fields, MP can compare multiple fields at a time**

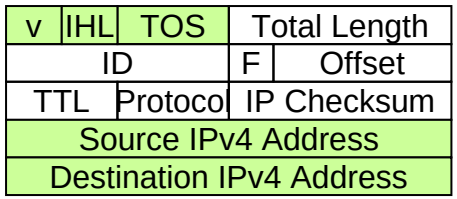


MP repeats comparisons and finds a flow.

Comparison method for multiple fields in MP (2)

Example: Flow Key: Version, IHL, TOS, source Address, destination Address

All fields are compared every time (general approach)



an observed packet

Compare

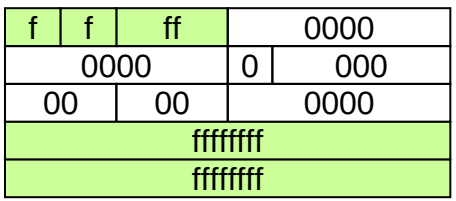
Any format

A Flow Record in cache

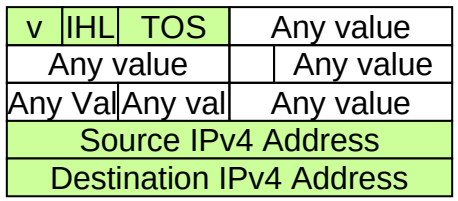
When a packet arrives:
5 comparisons
1. ip version
2. IHL
3. TOS
4. Source Address
5. Destination Address

Multiple field comparison (our approach)

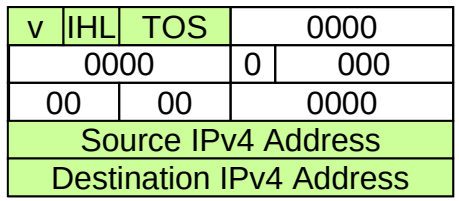
Premise: Fields of Flow Records are placed in the referring order as packet header fields



Mask created when template is defined



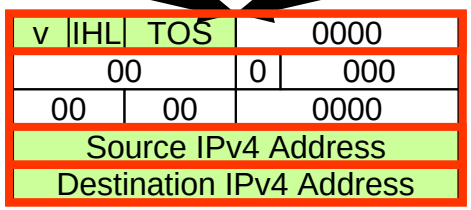
Observed packet



A Flow Record in cache

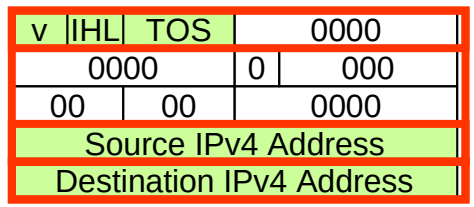
When Template is defined:
Create a Mask

When a packet arrives:
Mask the packet
And
compare these memory areas at the same time
(e.g., memcmp in C language)



Masked observed packet

compare



A Flow Record in cache

Or
1. v + IHL + TOS
2. Source Address
3. Destination Address
(32-bit architecture)

Comparison method for multiple fields in MP (3)

- Number of operations in this method
  - Mask costs smaller than comparison costs.
  - Therefore, this method is effective at increasing performance by reducing the number of comparisons, although it increases mask operations.

	Mask creation	Mask	Comparison
Number of operations	Once in an IPFIX session (when Template is defined)	Depends on the number of observed packets (when packet arrives)	Depends on the number of observed packets and number of flow records in cache

less ⟷ more

- Effective and ineffective cases

v	IHL	TOS	Total Length
ID		F	Offset
TTL	Protocol	IP Checksum	
Source IPv4 Address			
Destination IPv4 Address			

Effective case:
Flow Keys are placed densely

v	IHL	TOS	Total Length
ID		F	Offset
TTL	Protocol	IP Checksum	
Source IPv4 Address			
Destination IPv4 Address			

Ineffective case:
Flow Keys are placed sparsely.

2nd idea to improve performance
in environment in which MP, EP, and CP use the same order

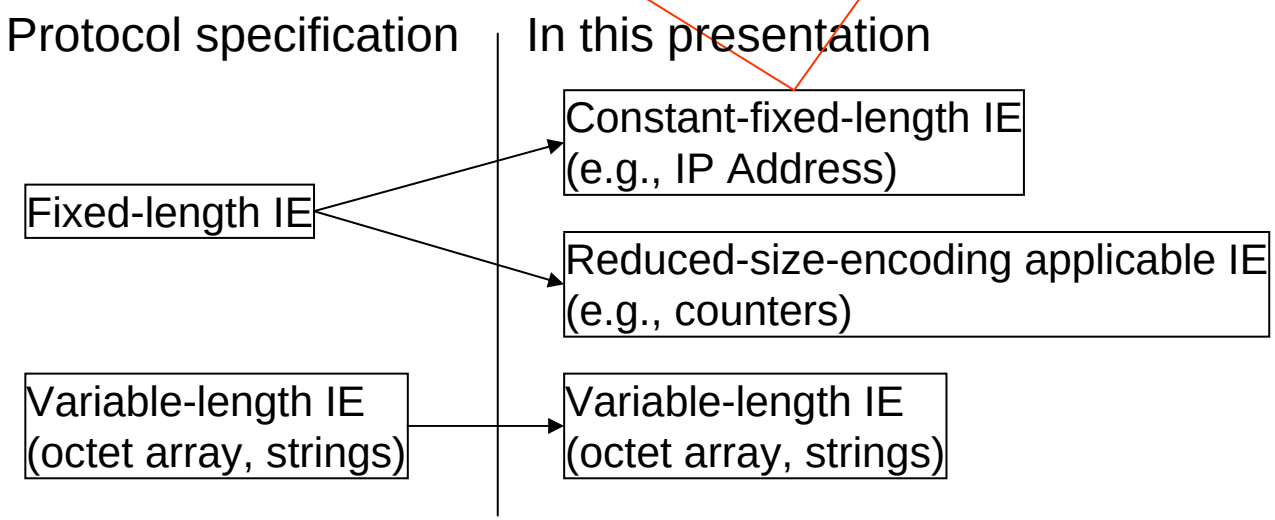
Copy method for multiple fields in Exporting Processes (EPs) and Collecting Processes (CPs)

NTT Network Service System Laboratories, NTT Corporation

Overview of copy method for multiple fields

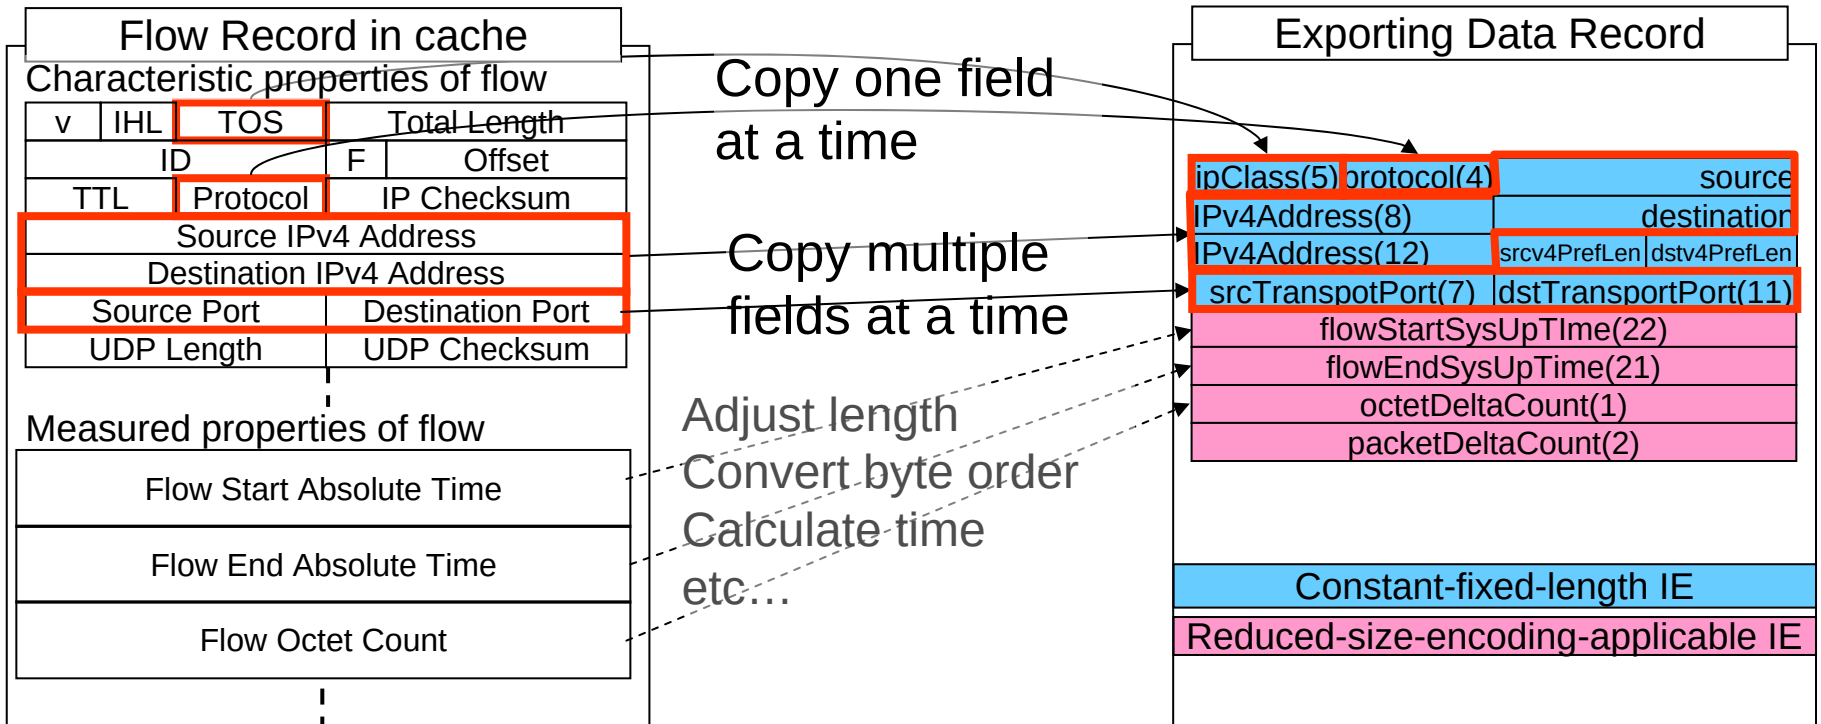
- It is a very simple method.
  - If fields in the format of cache and IEs in exporting Data Records are placed in the same order, EPs have a chance to copy multiple adjacent constant-fixed-length IEs at a time.
  - If IEs in received Data Records and fields in Collectors' internal format to store Flow Records are placed in the same order, CPs have a chance to copy multiple adjacent constant fixed-length IEs at a time too.

- IE size classification of IPFIX (terminology in this presentation)



Example of copy method for multiple fields in EP

- Conditions for copying multiple fields
  - Flow Record in cache and Exporting Data Record must use the same order.
  - IEs must have a constant fixed length.
    - Almost all IE characterizing properties of flow are constant fixed length.
  - Byte-orders must be the same.
    - Observed packet and Exporting Data Records use network byte order.
  - IEs for copying multiple fields must be adjacent.



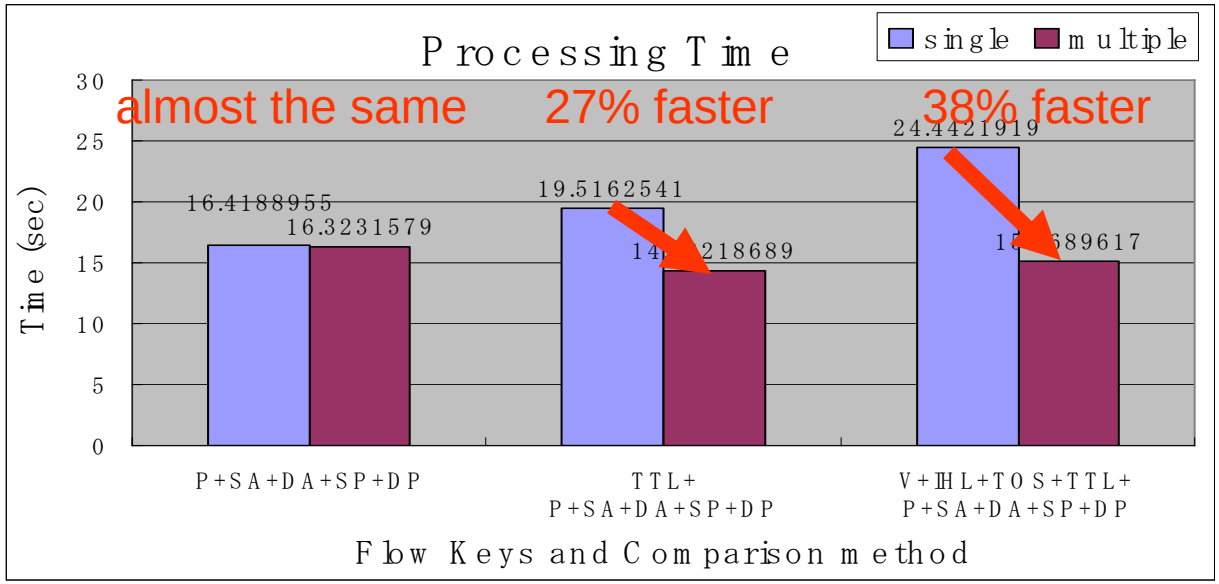
Evaluation & Conclusion

NTT Network Service System Laboratories, NTT Corporation

This material contains an evaluation about only comparison method.

If you want to see an evaluation about copy method, please see a material I talked in past IETF, http://www3.ietf.org/proceedings/07jul/slides/ipfix-10.pdf.

Evaluation of comparison method for multiple fields



v	IHL	TOS	Total Length
ID		F	Offset
TTL	Protoco		IP Checksum
Source IPv4 Address			
Destination IPv4 Address			
Source Port			Dst Port

P+SA+DA+SP+DP

v	IHL	TOS	Total Length
ID		F	Offset
TTL	Protoco		IP Checksum
Source IPv4 Address			
Destination IPv4 Address			
Source Port			Dst Port

TTL+
P+SA+DA+SP+DP

v	IHL	TOS	Total Length
ID		F	Offset
TTL	Protoco		IP Checksum
Source IPv4 Address			
Destination IPv4 Address			
Source Port			Dst Port

V+IHL+TOS+TTL+
P+SA+DA+SP+DP

- **When the density of Flow Key fields is higher, this method works faster.**

Computing environment for the evaluation

- Software Exporter program
  - runs on Intel Xeon 3.06 GHz HT architecture
  - runs on Linux (debian/gnu Linux 4.0)
  - compiled by gcc4
    - optimized option: -O3
- Data used as observed packets:
  - PCAP data published by WIDE project.
  - contains 6,906,333 packets.
  - ftp://mawi.nezu.wide.ad.jp/pub/mawi/samplepoint-B/20060303/200603030100.dump.gz

Conclusion

- Introduced ideas to improve performances of IPFIX processes
  - **Comparison method for multiple fields in MPs**
  - **Copy method for multiple fields in EPs, and CPs**
- These ideas are based on **defining the order rule of IEs/fields**
  - Our recommendation: IEs/f**ields are placed in the order referring to the packet header fields.**
- The order rule is published as an individual Internet Draft
  - http://tools.ietf.org/id/draft-irino-ipfix-ie-order-03.txt
  - If you agree with these ideas, work with us.











Abnormal traffic detection and alert

Yiming Gong
XO Communications

Flocon 2008

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

The problem and request











- XO network
  - OC-192 IP backbone with OC-12 uplinks in our markets and data centers, AS 2828
- Backbone level abnormal traffic detection
  - netflow

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

The problem and request











- Commercial product not good enough
  - You get what GUI gives you
  - Very likely to miss low volume traffic attack
    - (storm worm, scans)
  - By default, alert based on thresholds
  - Lacking data mining ability
  - Cost
- Free flow-based tool
  - Powerful but you need tell them what to do

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

So what we want











- Detect network abnormal traffic
  - both low and high volume
- Non-threshold based
- Automatically
- Fully controlled and customized
- Data mining
- Better be free

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Perfect world



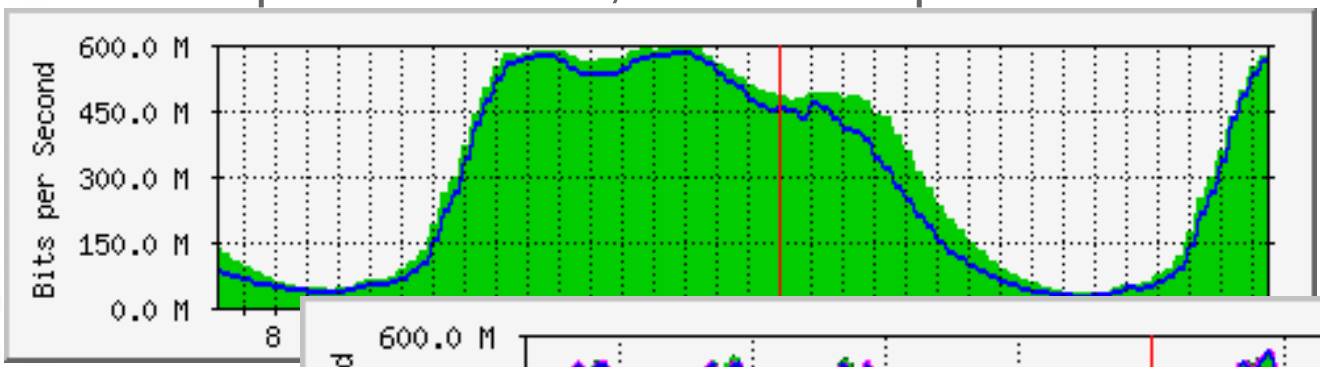


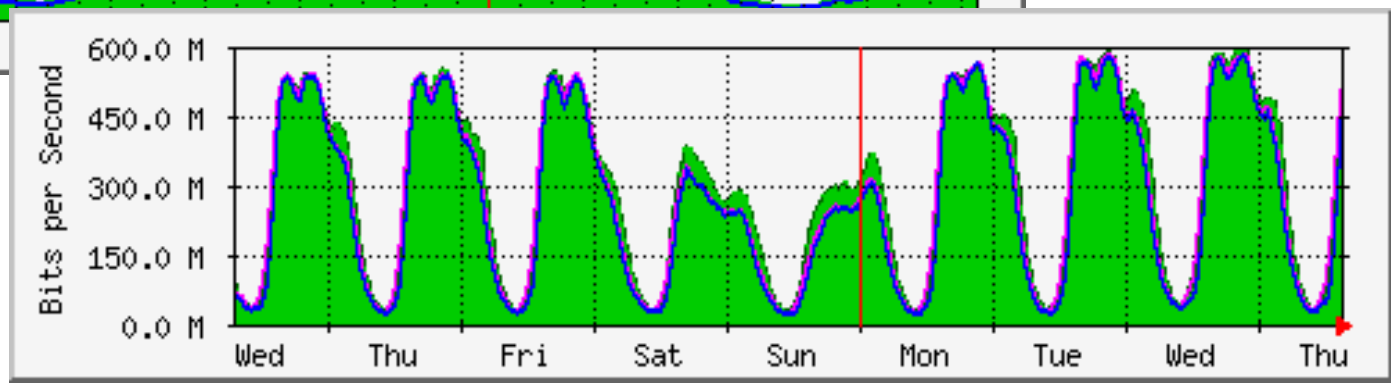






- In a perfect world, traffic shape should be very smooth





- Spike means……?

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Detection at traffic level is not good











- Granularity is too coarse
  - real attack hides behind the huge traffic
- Not easy to tell what is going on
  - SYN attack? ICMP ping flood?

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Our thoughts







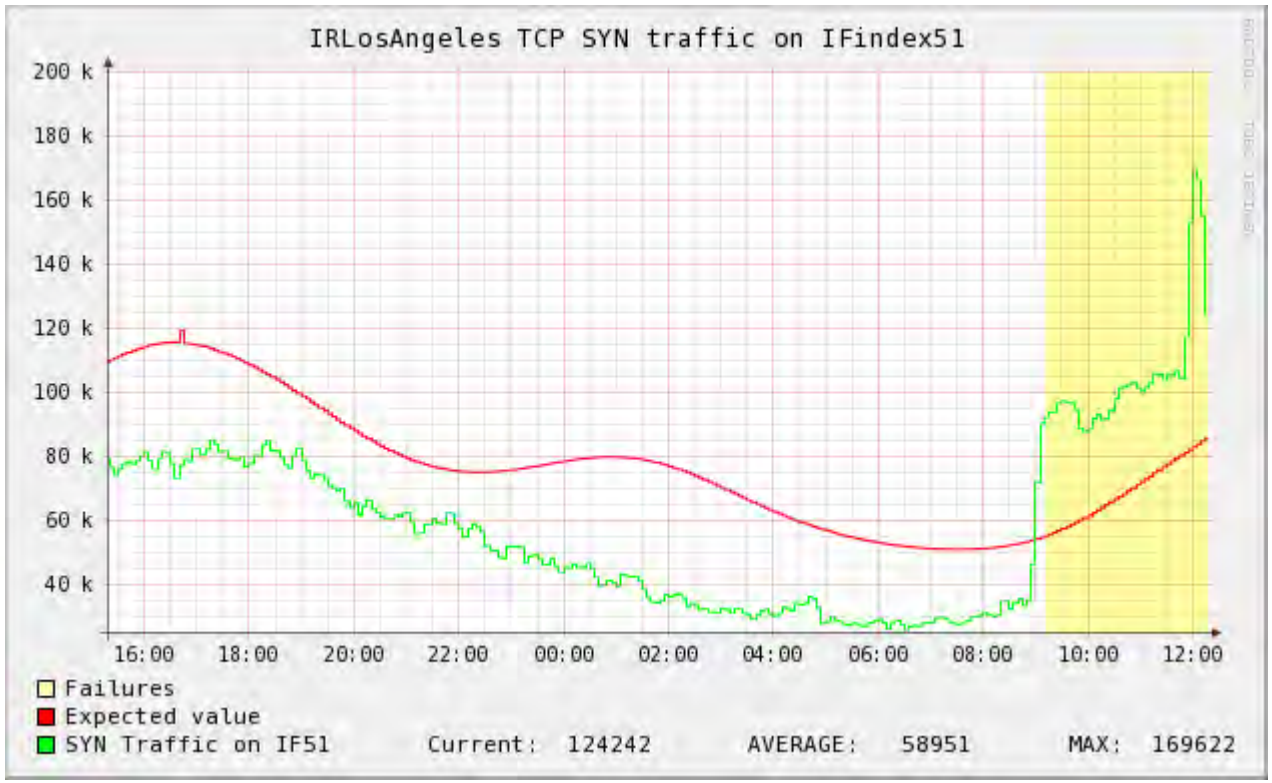




- Netflow based
- Break down raw netflow records to
  - TCP SYN, UDP total, ICMP type|code, protocol on each IFIndex of each router
    - Session
    - Traffic
- For each element, establish a dynamic profile
- When there is spike, something is going on

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Dynamic profile



CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Dynamic profile







- Establishing a profile
  - Using NFDUMP receive, store and process netflow data
  - rrdtool with aberrant behavior module
  - rrdtool (http://oss.oetiker.ch/rrdtool/)
  - aberrant behavior module
    - Learns from past values and uses them to predict the future
    - Tolerance band

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Dynamic profile







- Nfdump

```
yiming> more IR-syn-Amsterdam
13 1864
9 144
21 85
```

- Rrdtool

```
rrdtool create IR-syn-Amsterdam.rrd -s 300
DS:13:GAUGE:1200:0:U        \
DS:9:GAUGE:1200:0:U        \
DS:21:GAUGE:1200:0:U        \
RRA:HWPREDICT:2016:0.001:0.0035:288
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Failure







- Only an entry
  - IR-syn-Amsterdam:[1196800800]RRA[FAILURES][1]DS[13] = 1.0000000000e+00
  - Need script do the trace back work
  - Every 10 minutes, scans the rrd output for failures
  - Short-life spike
    - window-length and failure-threshold
    - rrdtool tune x.rrd -–window-length 5 –failure-threshold 3

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Failure











- Tracking down the failure
  - Nfdump + netsnmp + mysql + whois…
  - Narrowing down from flow and getting the suspicious host(s)

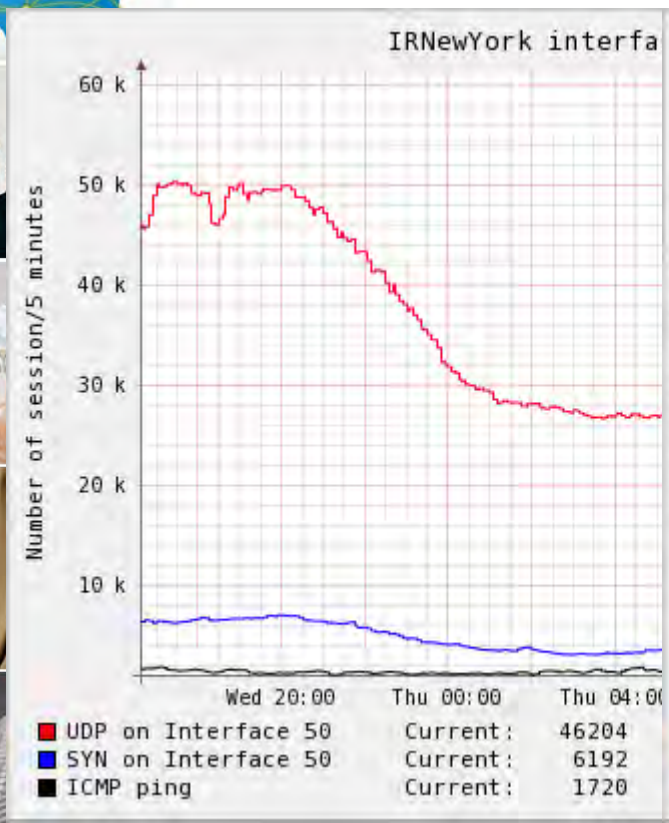
```
yiming> more IR-syn-Amsterdam
13 1864
9 144
21 85
```

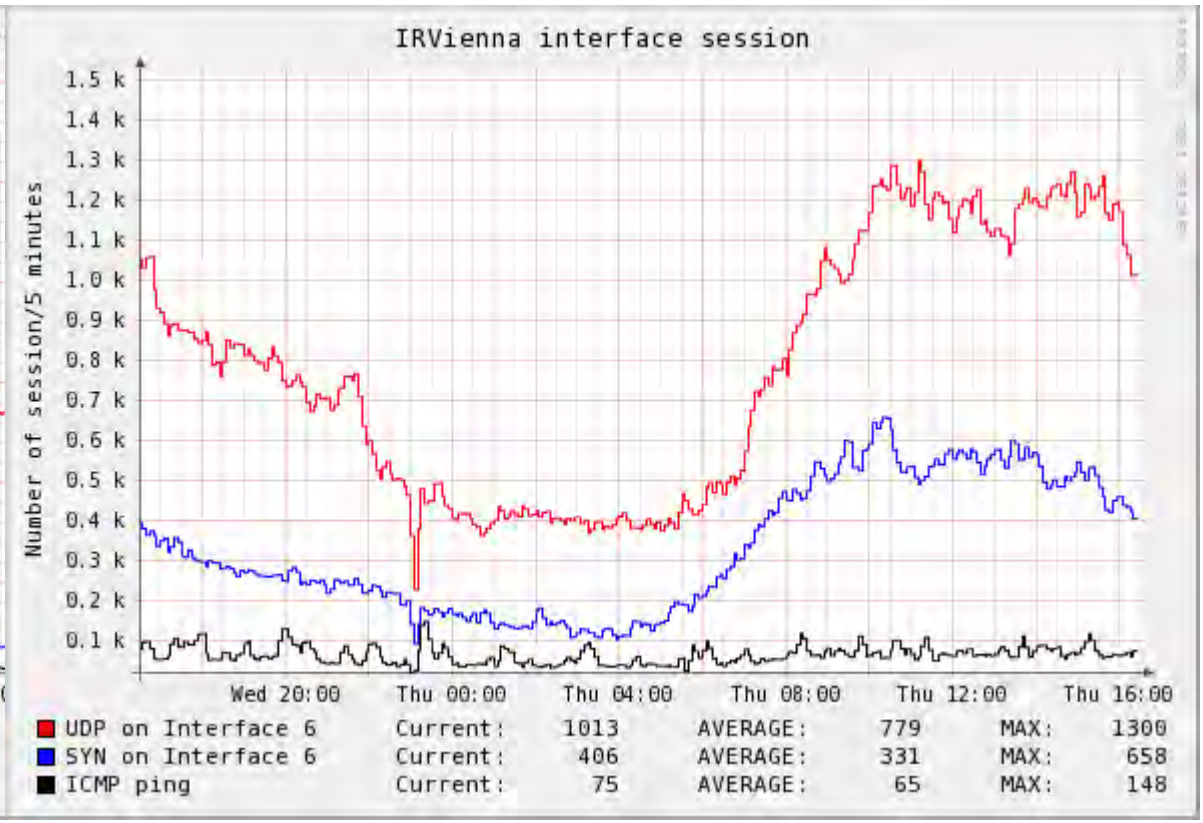
  - Flow of “TCP + SYN bit only + IFindex 13 + router Amsterdam”
  - Finding the most ACTIVE host(s)
    - What is the definition of active?
      - Session number
      - Traffic volume

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Finding active host

- Differences between these two pics?





CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Finding active host







- Different criterion

```
session-icmp*)
        total-number="500";
        flowfilter="proto icmp and port 2048 and if $if";
        trigger-number="280";
;;
 session-syn*)
        total-number="2000";
        flowfilter="proto tcp and flags 2 and if $if";
        trigger-number="600“;
```

- Things we ignored
  - TCP SYN is supposed to be 1, but is 10 now
  - Low volume UDP spikes

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Netflow records











- Pull out necessary data
- Generate alert
  - Picture, email

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert











- Scan alert

```
>IR LosAngeles has 5462 sessions on proto tcp and flags 2 and if 50 in 5 minutes

50 = STRING: [redacted]
50 = STRING: [redacted]

>Snapshot picture

http://[redacted]LosAngeles-50-abnormal.png

>One week|month picture

http://[redacted]LosAngeles-50-abnormal-week.png

http://[redacted]LosAngeles-50-abnormal-month.png

>Top IPs in 10 minutes

Date first seen          Duration Proto          IP Addr         Flows   Packets     Bytes      pps      bps
2007-12-05 08:51:02.520   289.718 any      218.233.1[redacted]    2114      2114     84560        7     2334
2007-12-05 08:51:20.493   130.413 any        218.234[redacted]     605       605     24200        4     1484

>Top IP info

* AS          IP              AS name                                          FQDN
[redacted]  | 218.233.[redacted]  | [redacted] Telecom Inc.  |
[redacted]  | 218.234.[redacted]  | [redacted] Telecom Inc.  |
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert



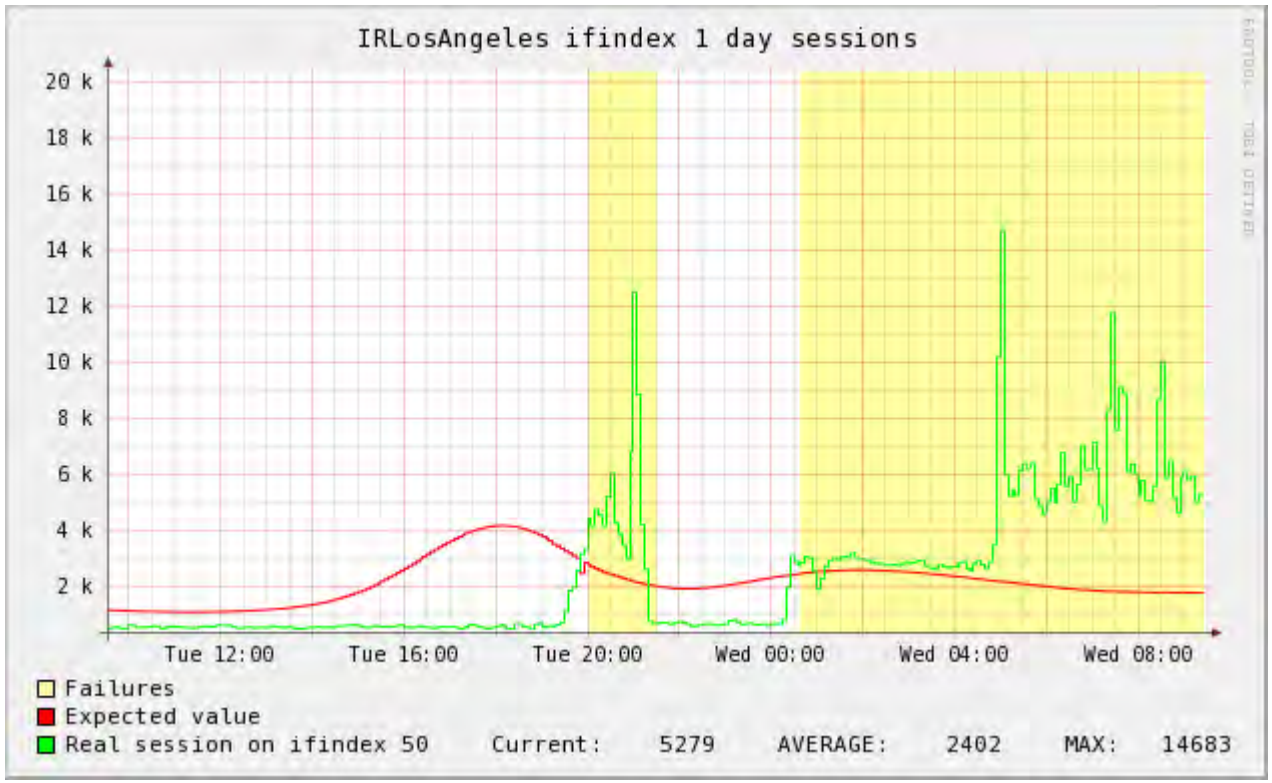








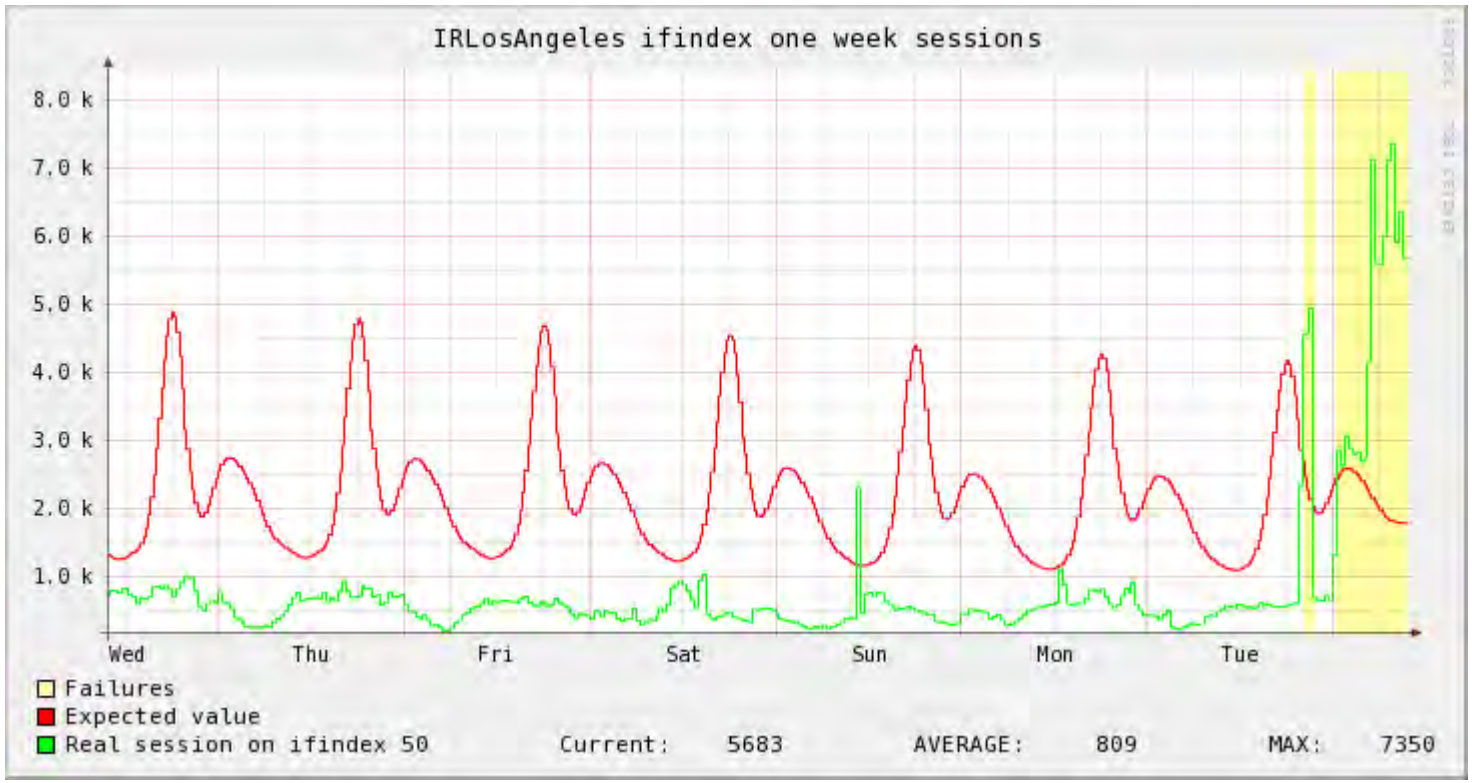
- Day



CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert

- Week



CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert











- Scan alert

```
>Top IP detail

/ ip 218.233.[redacted]

**Traceroute (from hop 5 to 9)

 5   [redacted]
 6   [redacted]gin.ntt.net (206.223.[redacted]19.12)  7.089 ms
 7   [redacted]ntt.net ([redacted]  7.317 ms
 8   [redacted].gin.ntt.net ([redacted]5.20)  73.034 ms
 9   [redacted]ntt.net ([redacted]  73.922 ms

**Protocol summary for 218.233.198.25

Proto  Flows  Packets    Bytes      pps      bps   bpp
  6     2116     2116    84640        7     2337    40
 17        1        1      257        0        0   257

**Sampled netflow records

TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000    65.99.[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000    65.99.[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000    65.99.[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000    65.99.[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000    65.99.[redacted]:7212  ....S.      40   0  [redacted]  50   73
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert











- Scan alert

```
/ ip 218.234.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted] 106.6.[redacted] [redacted] ms
 6  [redacted] n.ntt.[redacted] (206.223.[redacted]119.12)  7.135 ms
 7  [redacted] .net ([redacted])  7.268 ms
 8  [redacted] n.ntt.[redacted] (129.250.[redacted].5.20)  79.209 ms
 9  [redacted] .net ([redacted])  74.073 ms

**Protocol summary for 218.234.[redacted]

Proto  Flows  Packets    Bytes      pps      bps   bpp
 6       605      605    24200        4     1484    40

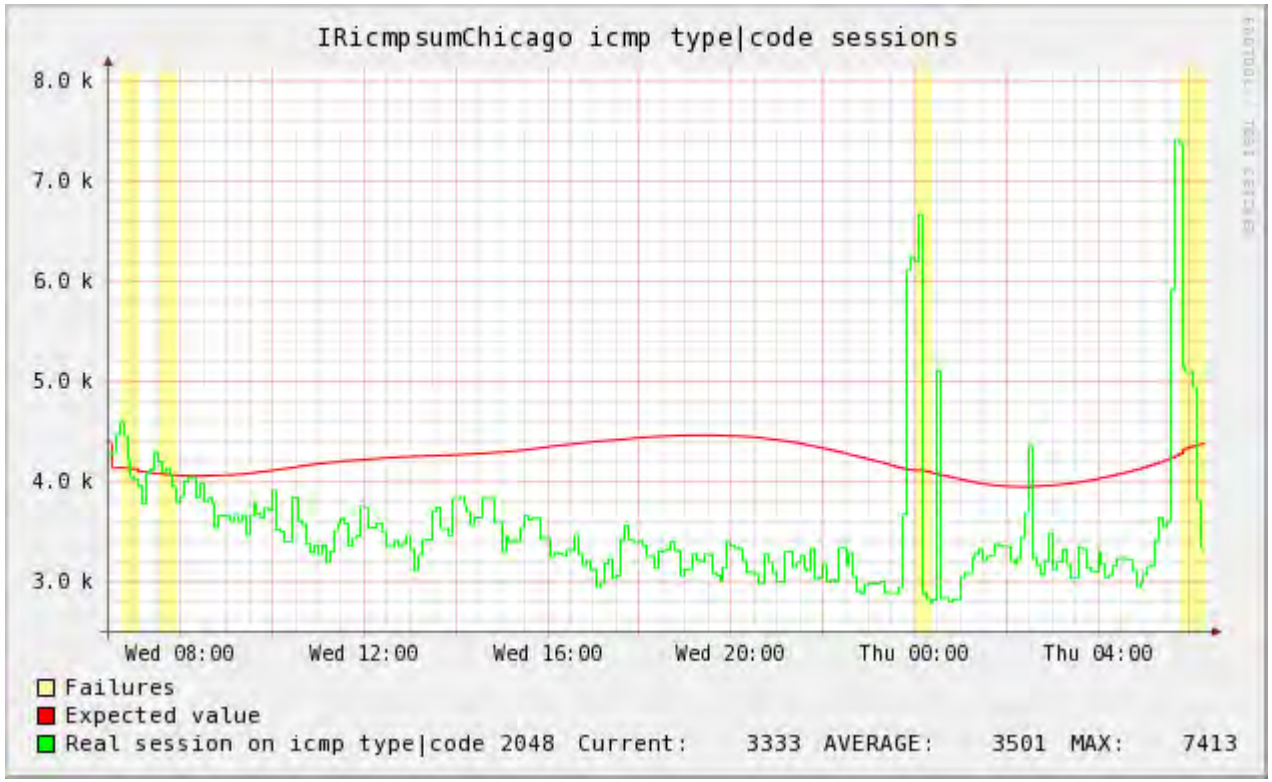
**Sampled netflow records

TCP        218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP        218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP        218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP        218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP        218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP        218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP        218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP        218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP        218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP        218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

alert

- Storm worm



CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert - one week later











- DDos

```
IR LosAngeles has 177002 sessions on proto tcp and flags 2 and if 50 in 5 minutes

50 = STRING:[redacted]
50 = STRING:[redacted]

>Snapshot picture

http://[redacted]LosAngeles-50-abnormal.png

>One week|month picture

http://[redacted]LosAngeles-50-abnormal-week.png

http://[redacted]LosAngeles-50-abnormal-month.png

>Top IPs in 10 minutes

Date first seen          Duration Proto         IP Addr         Flows   Packets    Bytes      pps      bps
2007-12-12 09:00:23.705   320.317 any       89.144.[redacted]   173554    176183   10.0 M      550   261507
2007-12-12 09:00:28.361   297.573 any       211.211[redacted]     1048      1056    50688        3     1362
2007-12-12 09:00:43.293   282.273 any       211.206[redacted]      692       708    33984        2      963
2007-12-12 09:00:43.269   291.093 any       211.44.[redacted]      658       667    42688        2     1173
2007-12-12 09:00:43.401   289.437 any       218.48.[redacted]      633       684    32832        2      907
2007-12-12 09:00:37.353   288.445 any       123.214[redacted]      627       640    30720        2      852
2007-12-12 09:00:23.705   311.869 any       58.127.[redacted]      603       618    39552        1     1014

>Top IP info

* AS          IP              AS name                                                    FQDN
[redacted]  | 89.144.[redacted]  | [redacted]Autonomus System number for [redacted]Net  |  ;; connectior
[redacted]  | 211.211[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.  |
[redacted]  | 211.206[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.  |
[redacted]  | 211.44.[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.  |
[redacted]  | 218.48.[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.  |
[redacted]  | 123.214[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.  |
[redacted]  | 58.127.[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.  |
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert - one week later











- Traceroute returns nothing

```
>Top IP detail

/ ip 89.144.[redacted]

**Traceroute (from hop 5 to 9)  <----------  no traceroute info here

**Protocol summary for 89.144.[redacted]

Proto  Flows  Packets    Bytes      pps      bps   bpp
 6   173974   176610   10.0 M      551   261950    59

**Sampled netflow records

TCP     219.254.[redacted]:2391      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  11
TCP    123.212.[redacted]:2735      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP    219.233.[redacted]:3878      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP     58.124.[redacted]:4375      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP    219.251.[redacted]:4049      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP     58.123.[redacted]:3642      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  11
TCP     123.21[redacted]:2340      89.144.[redacted]:80    ....S.    48  0  [redacted]  50  10
TCP     211.200.[redacted]:4256      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  11
TCP    221.143.[redacted]:4313      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP     218.234.[redacted]:4353      89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert - one week later











```
/ ip 211.211.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted]L.us.xo[redacted].85)  7.113 ms
 6  [redacted]et (65.[redacted]521 ms
 7  [redacted]net (20[redacted]6.505 ms
 8  [redacted]et (65.[redacted].604 ms
 9  [redacted].net (6[redacted]6.511 ms

**Protocol summary for 211.211.[redacted]

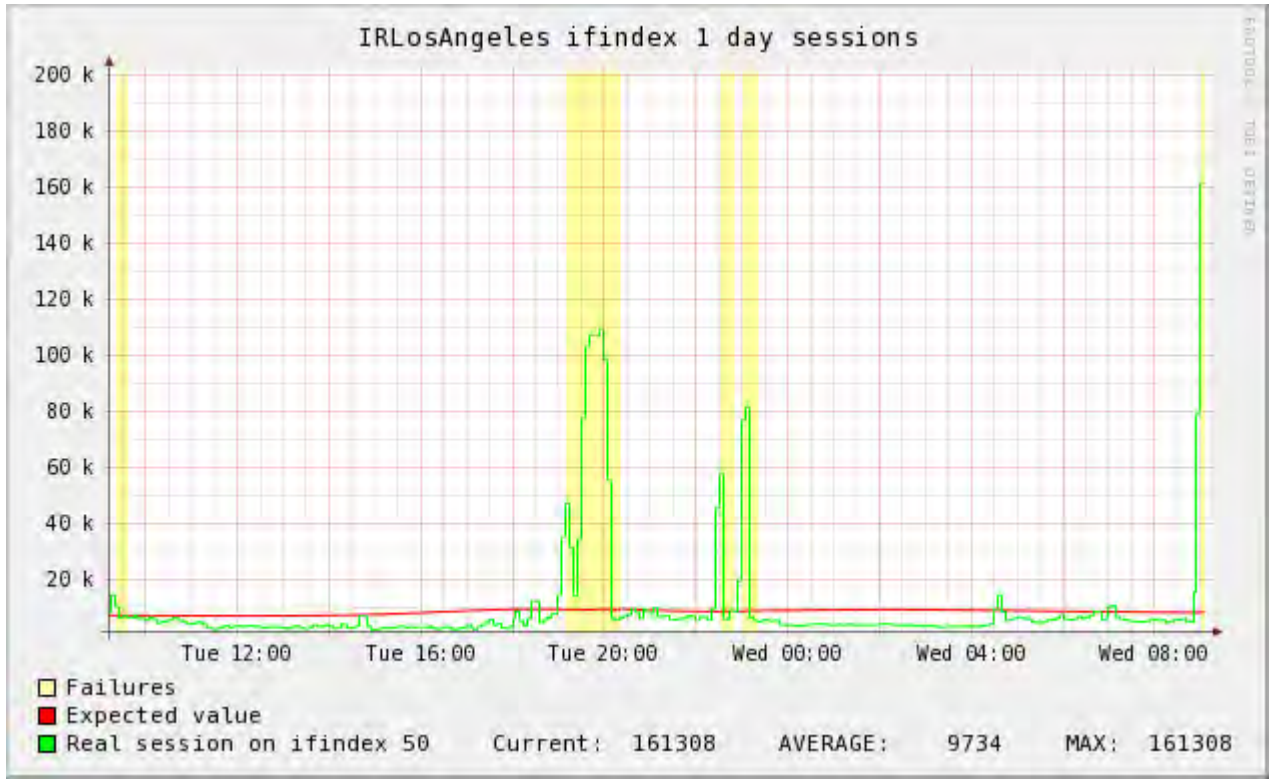
Proto  Flows  Packets    Bytes      pps      bps   bpp
 6     1057     1065    51120        3     1374    48

**Sampled netflow records

TCP    211.211.[redacted]:29937    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:32301    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:30596    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:35573    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:26497    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:31263    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:27378    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:34829    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:28267    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:59695    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert - one week later



CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert - whitelist











- Special customers

```
>Top IP info

* AS            IP              AS name                                  FQDN
[redacted]  | 64.39.[redacted]  | [redacted]Inc.   |   scanner[redacted].com.
[redacted]  | 63.245[redacted]  | [redacted] Corporation   |   core2.[redacted].com.

>Top IP detail

/ ip 64.39.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted]  6.633 ms
 6  [redacted].com (154.54.12.229)  6.814 ms
 7  [redacted].com (154.54.5.17)  66.692 ms
 8  [redacted]com (154.54.7.138)  66.836 ms
 9  [redacted]com (154.54.5.217)  66.824 ms

**Protocol summary for 64.39.[redacted]

Proto  Flows  Packets    Bytes      pps      bps   bpp
  6      182      200     8584        0      231    42
 17        1        1       58        0        0    58

**Sampled netflow records

TCP       64.39.[redacted]:2681     63.245.[redacted]:25     ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:37672    63.245.[redacted]:35459  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:38206    63.245.[redacted]:47123  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:38318    63.245.[redacted]:2870   ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:39700    63.245.[redacted]:34739  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:40293    63.245.[redacted]:26733  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:40210    63.245.[redacted]:53606  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:40603    63.245.[redacted]:63822  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:41626    63.245.[redacted]:55450  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:41565    63.245.[redacted]:2361   ....S.      40   0  [redacted]  31   8
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert – whitelist and misc







- Whitelist <cont>
  - Email servers
  - We don't want to miss real attack even if an IP is on whitelist
- Alert email
  - Suppression period
  - Subject
    - 12-05 abnormal sessions at LosAngeles proto tcp and flags 2 and if 50

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Data mining







- Database
  - 3 tables
    - IP,FQDN,AS
    - Summary
    - Raw netflow data
  - Data mining
    - Which peering neighbor sends out most attack traffic, who is the most attacked, which port is the most popular being scanned…etc.

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Data mining







- Database
  - 3rd party outside data
    - Dshield TOP 10000
    - Dshield AS
    - CBL data
    - Mynetwatchman
    - Our own darknet project output
    - Other private outside data
  - If XO host involved, we will go through these table

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

problem











- Problem
  - Peering neighbor
  - Alert correlation
    - But you can do it in database.

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

What you need











- Nfdump, rrdtool, mysql, net-snmp, apache, some unix commands
- A box with linux installed

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.











- For more info
  - yiming.gong@xo.com

- Thanks!

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Incorporating Network Flows in Intrusion Incident Handling and Analysis

John Gerth
Stanford University

gerth@stanford.edu

EE/CS Network Infrastructure

- Three buildings with one router
  - (Gates) Computer Science
  - (Packard) Electrical Engineering
  - (Allen) Center for Integrated Systems
- Composition
  - 25 VLANs controlled by disparate groups
  - 10,000 IP addresses (about half are active)
  - Eclectic mix of Windows, Linux, Solaris, OS-X, …
  - No firewall beyond minor university filters
- Analysts
  - A half-dozen people with network (and other) responsibilities

Incident Investigation Process

- Find answers to a set of classic questions…
  - Who
  - What
  - When
  - Where
  - Why
  - How
- …using an iterative process
  - Inspect events of a focus node
  - Augment, refine, filter data
  - Compare events of related nodes, looking for correlation
  - Pivot on an “interesting” node to refocus

Network Data Sources
(each step is orders of magnitude more volume)

- **Traffic counters** (SNMP, MRTG, ….)
  - Configurable in network devices
- **Event/Alert logs** (Syslog, HTTPD, SNORT, ...)
  - Collected by firewalls, IDS, individual machines and services
- **Flows** (Netflow, YAF, Argus, ….)
  - Typically collected at border routers or taps
- **Packet Headers / Traces** (tcpdump, wireshark, …)
  - Collected at switches, routers, or taps

Network Flows

- Advantages
  - Relatively uniform and increasingly available
  - Hard to subvert
  - Mitigate privacy concerns
  - Largely insensitive to encryption
- Disadvantages
  - Still voluminous compared to event logs
  - Aggregate measure
  - Lack content

Flow Capture and Data Management

- Sensor
  - Span ports from two Cisco backbone switches
  - See all layer 3 traffic for three buildings (not just external)
  - Argus capture of bidirectional ICMP, UDP, TCP flows
- Collector
  - Raw flows from sensor are multicast locally in realtime
  - Hourly files from sensor compressed and archived
  - 20-30M (peak 70M) Argus flows/day (~1G compressed)
  - Retain several months of data online for analysts to access

Support flat files and database tables

- Flat text files
  - Familiar and familiar tools
  - Extracts useful for exchange and reporting
  - Straightforward sequential processing
  - Import to other tools for aggregation and analysis
- Relational databases
  - No longer exotic
  - Suitable for large data volumes
  - Greater expressibility for queries
  - Built-in support for aggregation and analysis

Database Infrastructure

- MySQL server running on collector
  - Live flows from sensor inserted in real-time
  - Daily tables recreated from archived raw flows
  - Monthly “merge” tables
  - Anonymize extracts for research with CryptoPAN
- Flow schema tuning
  - Transform src/dst to local/remote
  - Add ASN (routeviews.org) and local VLAN metadata
  - Convenience columns for locality, local role, dst port
  - Index most dimensions (adds about 50%)
  - Tables + indices ~2G/day

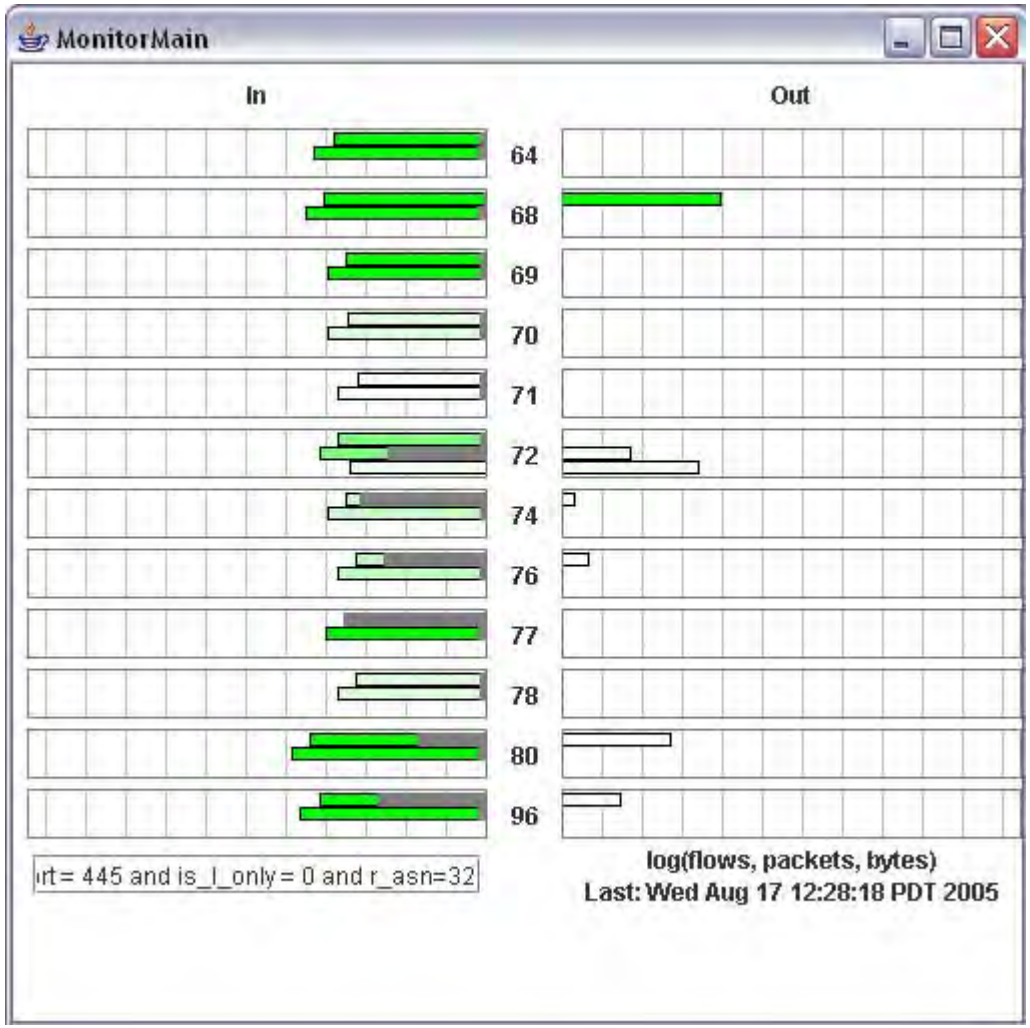
Flows in Incident Handling

- Worms and Trolls
  - Volume and promiscuity
- Immaculate Intrusions
  - Scrubbers, Keyloggers, and Remote Tunnels
- Botnets
  - Beaconing to Command+Control Hosts

Traffic Volume

- Windows Esbot worm circa 2005
  - Spread via PNP buffer overflow
  - Installed backdoor trojan
  - Victim turns into attacker
- Report
  - Overall traffic suddenly increased an order of magnitude
- Analysis
  - Flow distribution showed port 445 at 500-1000 flows/sec
  - Keyed on 445 traffic to identify attackers
  - Used “flow monitor” to reveal local compromises

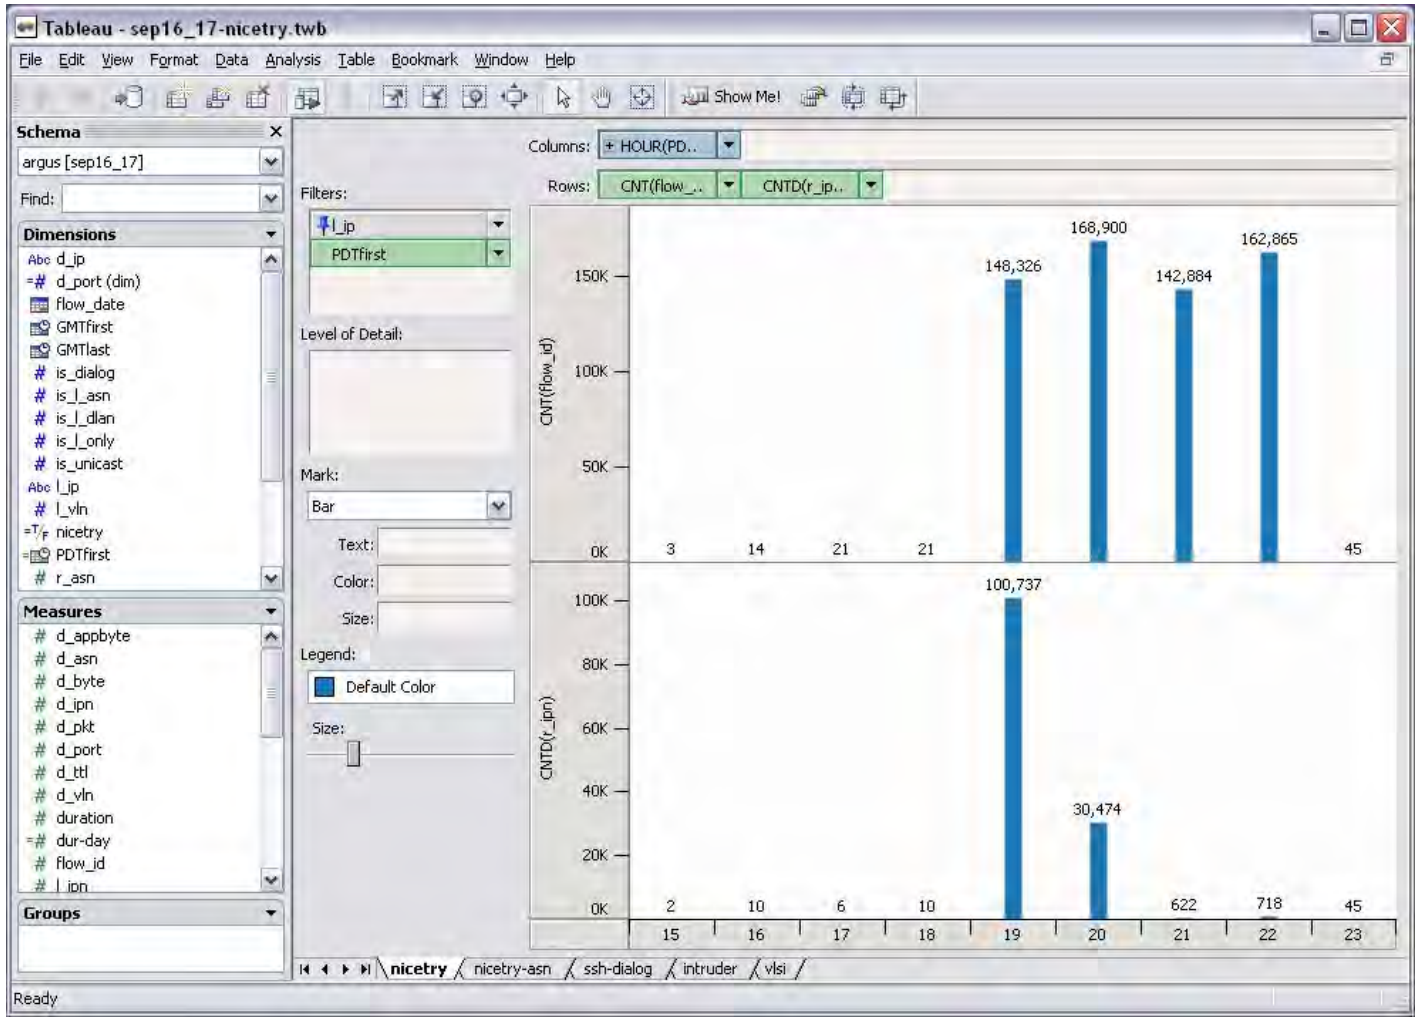
Esbot on the Flow Monitor



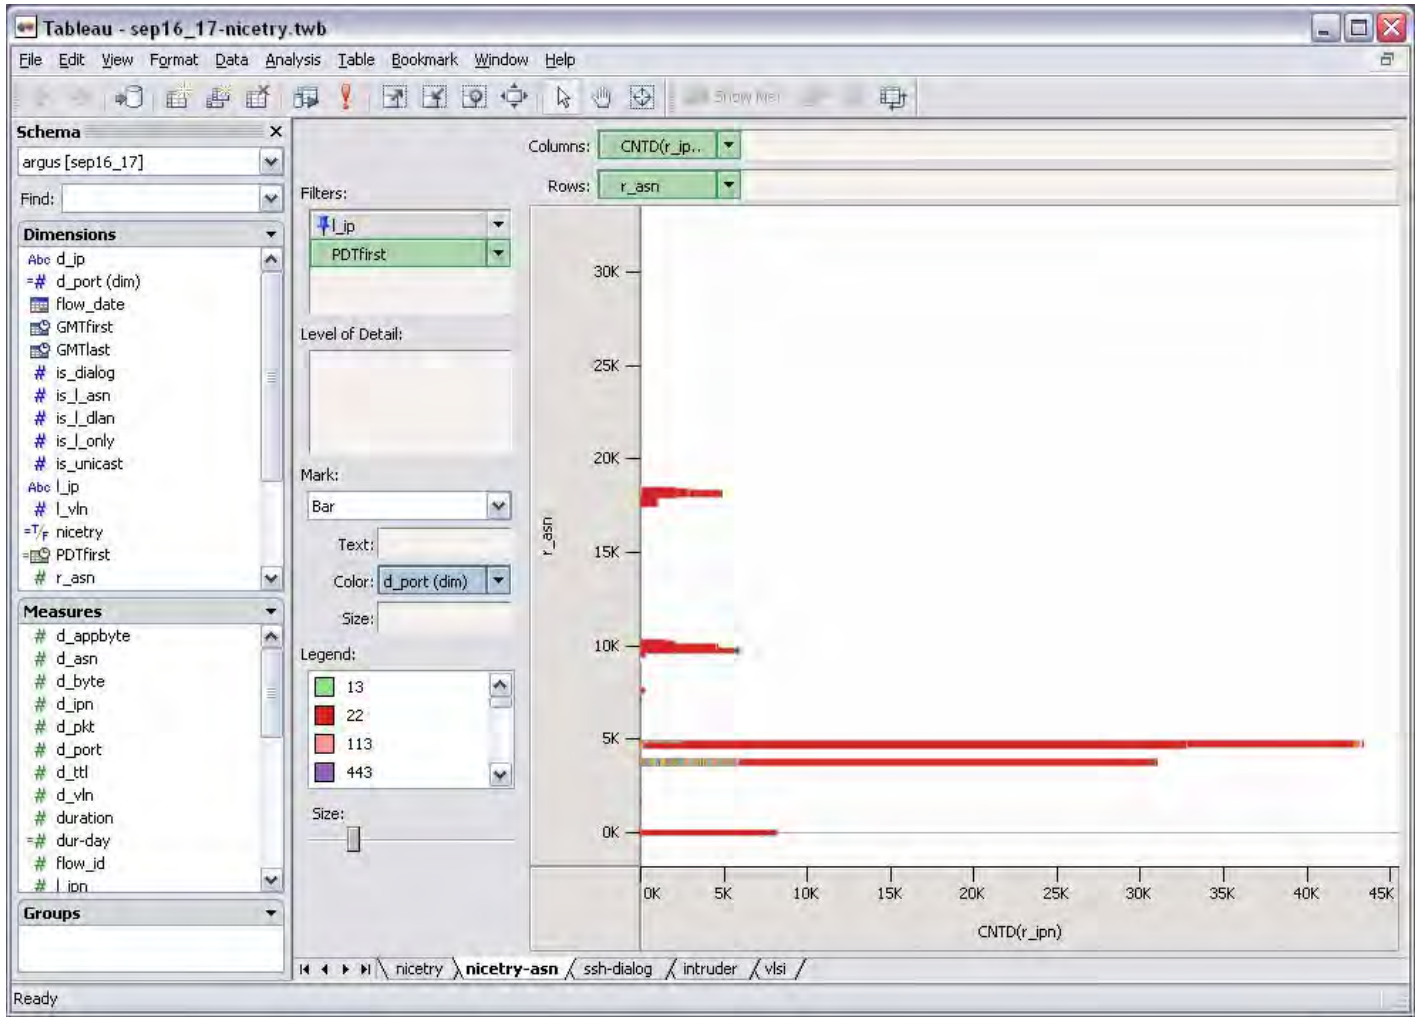
Promiscuity

- SSH Troll
  - Intruder gains access to local machine
  - Installs SSH troll
  - Launches attack on remote networks
- Report
  - Odd outbound traffic spike from local IP
- Analysis
  - Flow distribution showed many IPs, few ASNs, single port
  - Backtrack in time to find initial SSH compromise
  - Pivot reveals other victims

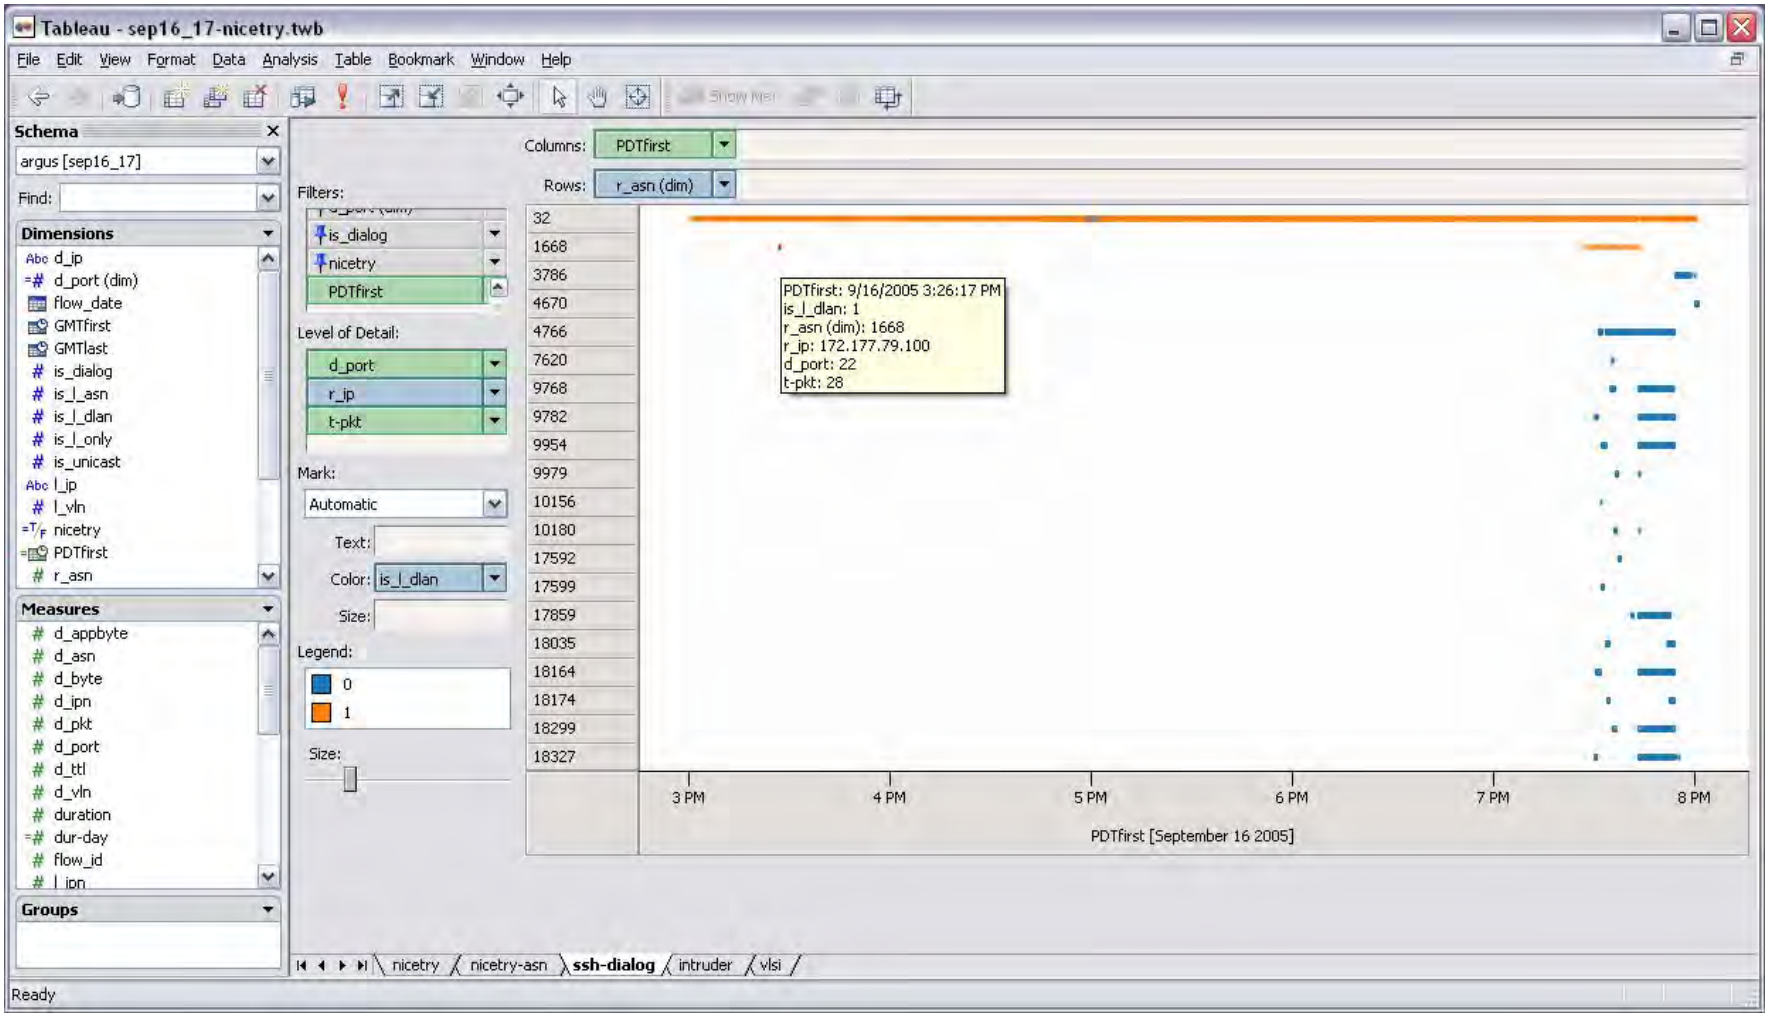
SSH Troll: Volume + Promiscuity



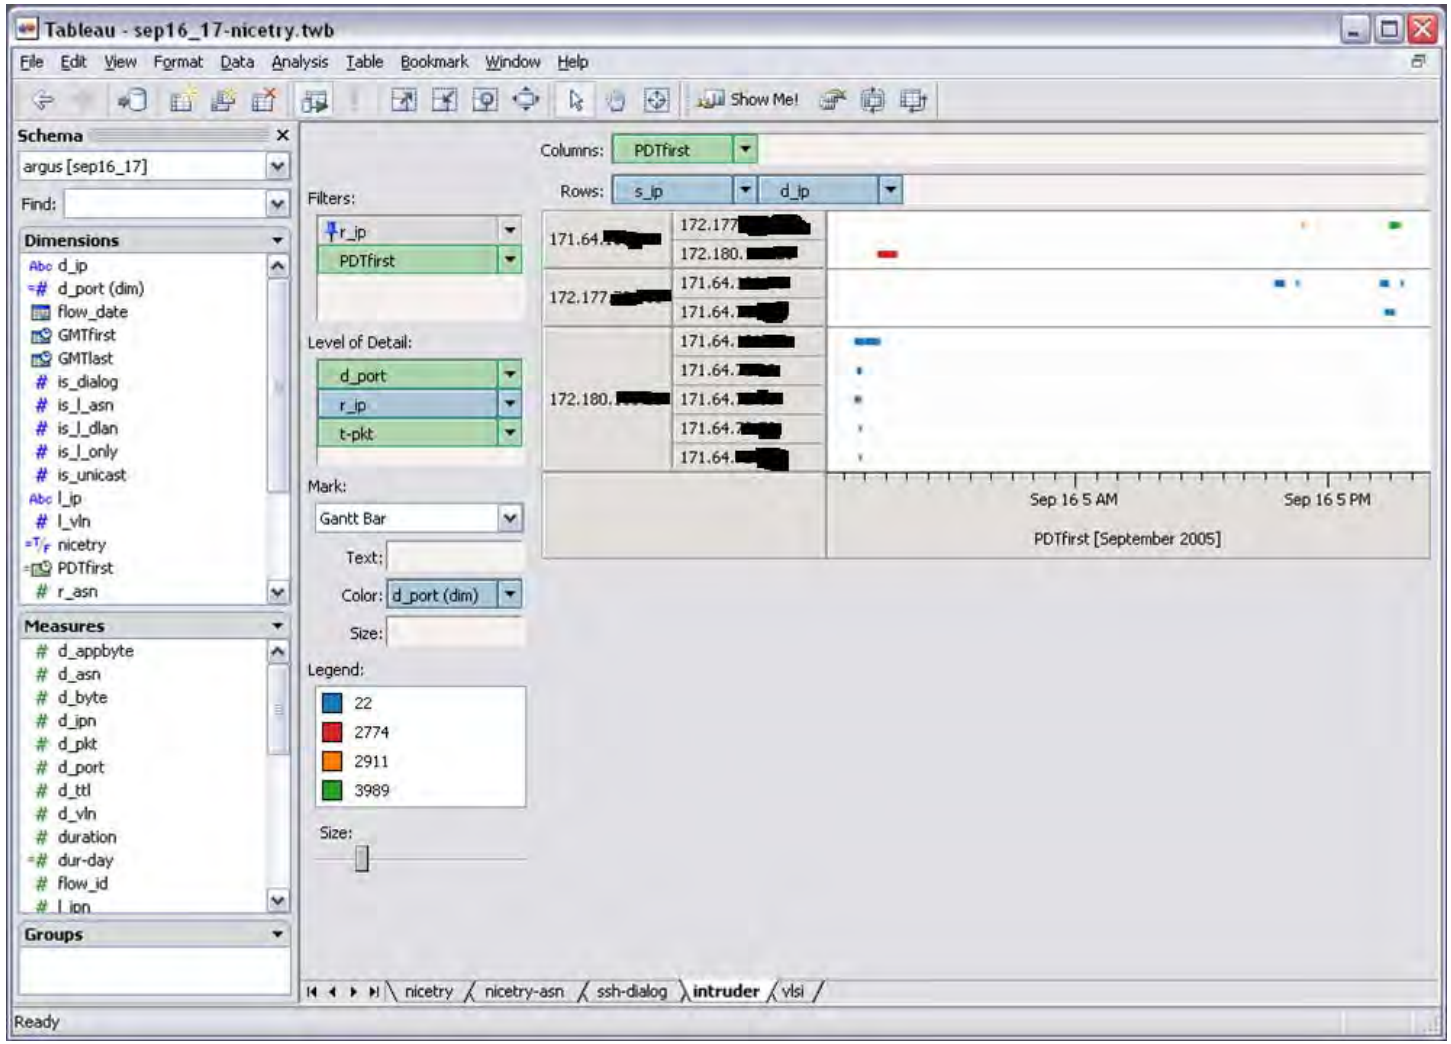
SSH Troll: Identifying targets



SSH Troll: Locate Compromise



SSH Troll: Pivot to identify other victims



Immaculate Intrusions - Keyloggers

- Unprotected X-Window server
  - Intruder maps 0x0 pixel client and signs up for keypress events
  - Steals credentials for other machines from local user
  - Uses credentials to login to experimental machine
- Report
  - Experimental machine crashes when intruder’s tools fail
- Analysis
  - Local user logged in when user not present
  - Discover open X-server on user’s desktop machine
  - Backtrack in time to find keylogger flows
  - Pivot reveals other victims

Immaculate Intrusions - Scrubbers

- Unpatched Linux machine
  - Unpatched server vulnerable to remote root compromise
  - Intruder installs backdoor, trojan binaries, and scrubs logs
  - Uses trojan ssh to steal credentials of local users
  - Uses ssh known_hosts data to attack other local machines
- Report
  - Local machine two hops away found sending spam
- Analysis
  - Backtrack of login sessions leads to compromised machine
  - Trojan binaries found, but no plausible root logins
  - Flow logs show original compromise and backdoor logins
  - Pivot reveals other victims

Immaculate Intrusions - Tunnels

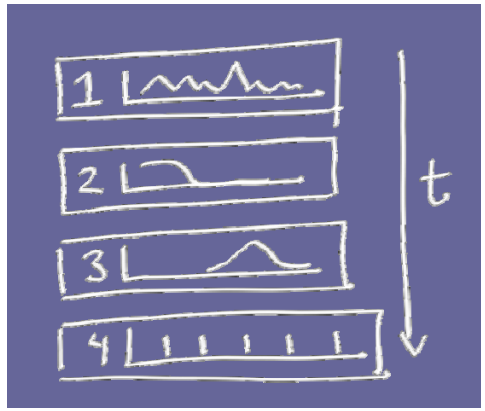
- Tunnels
  - Intruder compromises desktop machine running VNC client
  - Desktop machine has forwarded ports over ssh-tunnel
  - Intruder's traffic is tunnelled and reparented inside cluster
- Report
  - Apparent Nessus scan of ***isolated*** cluster machine
- Analysis
  - System logs of head node show no logins
  - Flow logs show massive ssh traffic from compromised machine

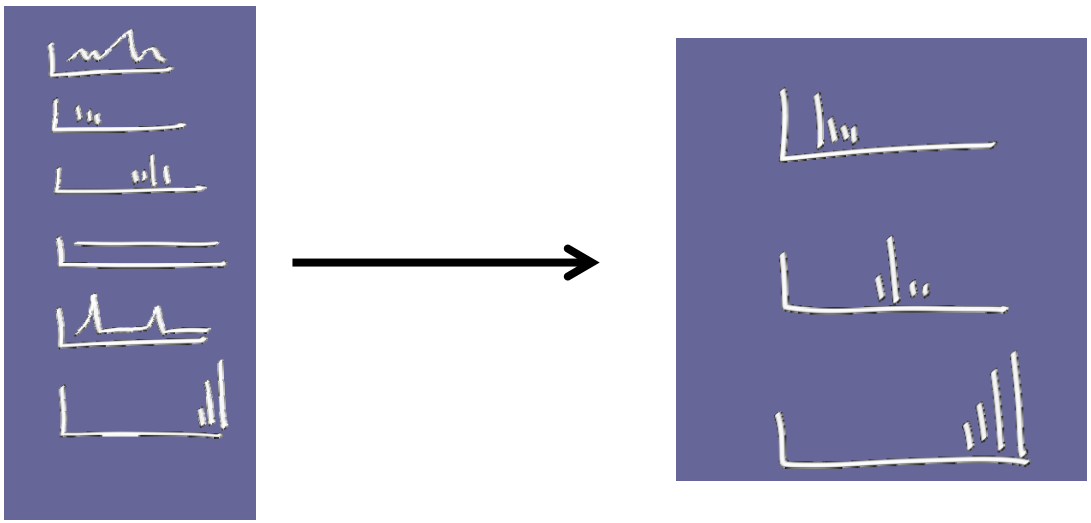
Isis:Visual Analysis of Flow Data

(see paper by Phan et al in VizSec 2007)

Progressive Multiples

- Make exploration history visible
- Reorder rows to reveal structure and event sequencing

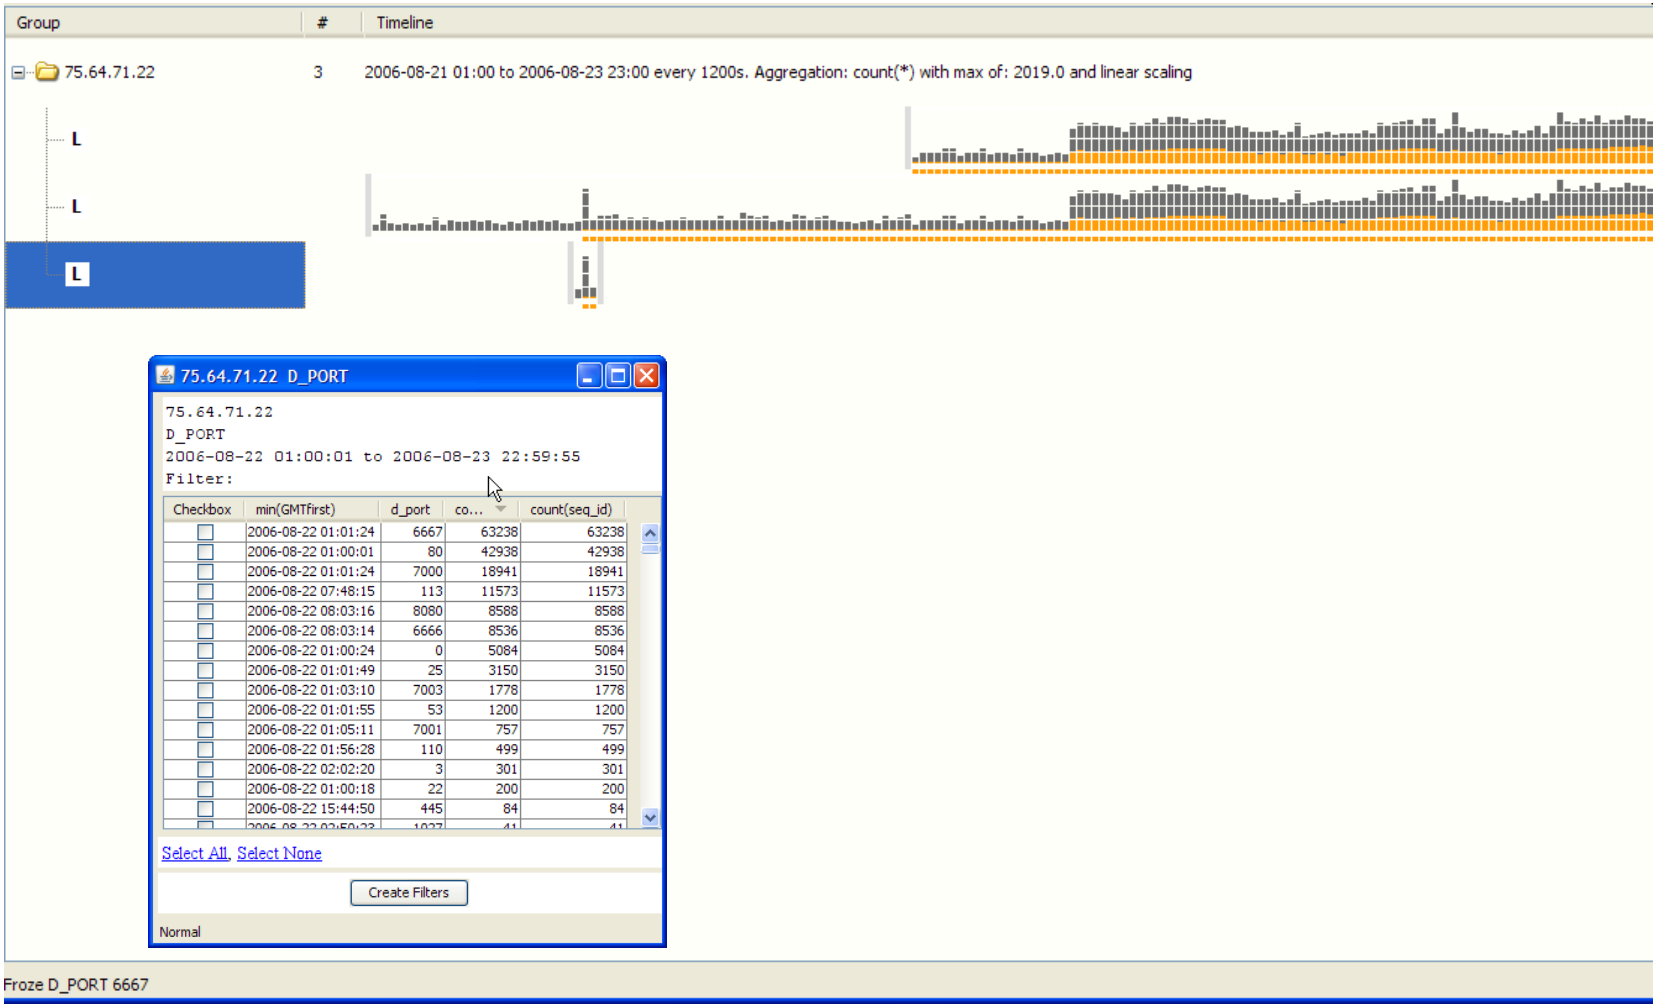




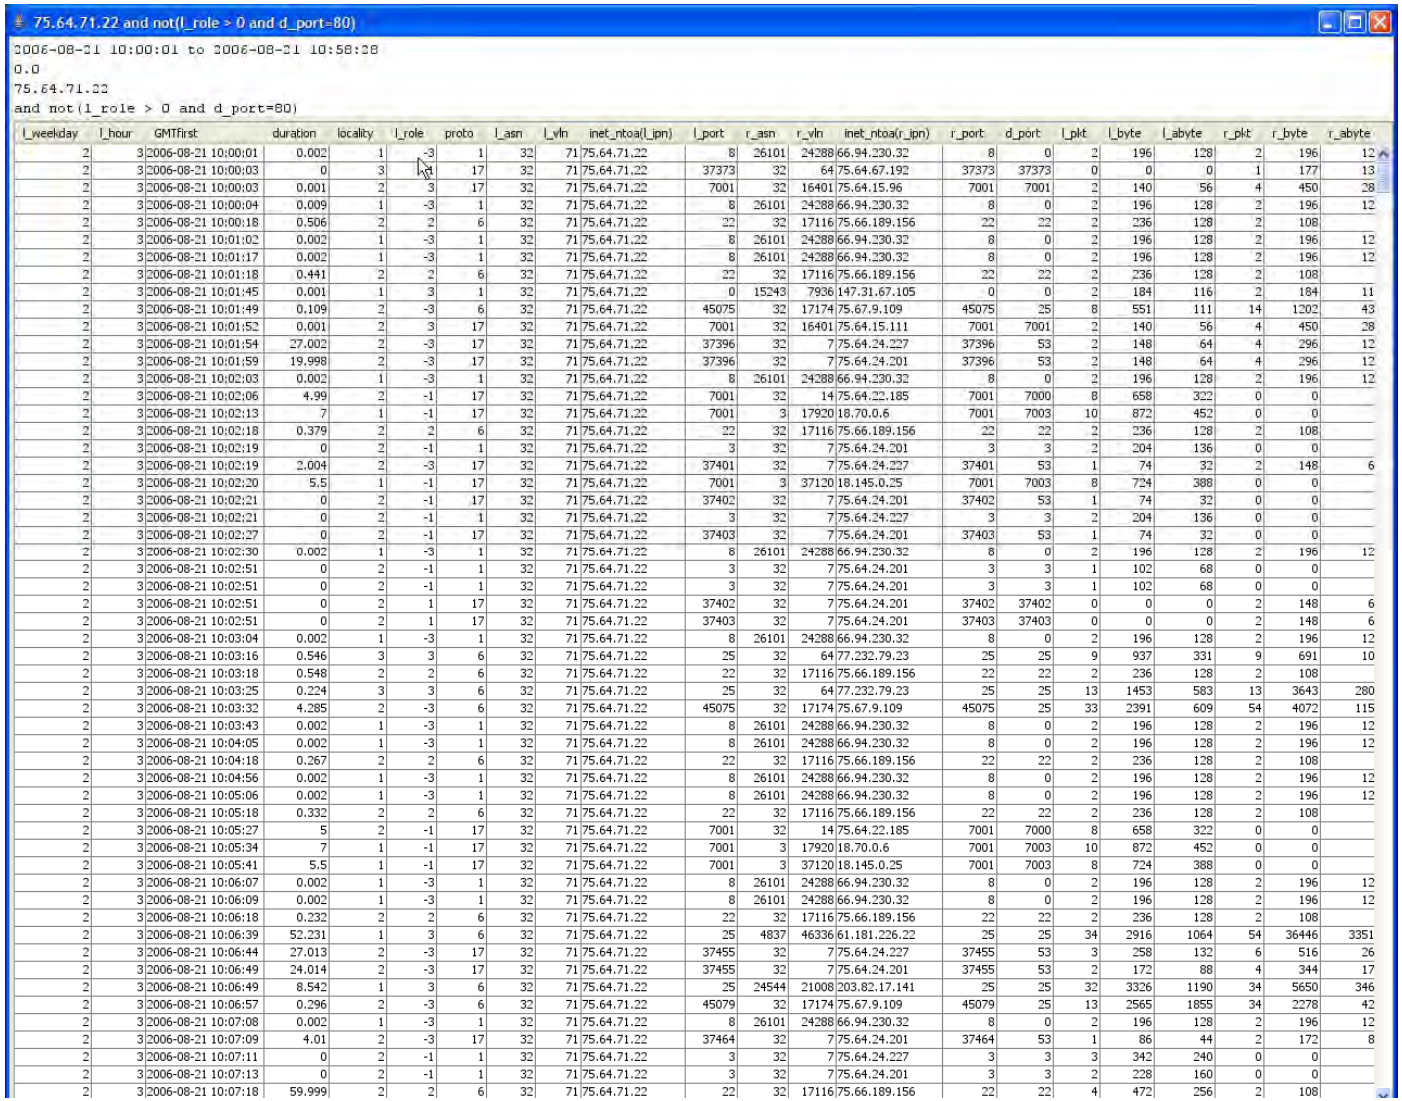
Beaconing

- Botnet zombie
  - Intruder gains access to local machine
  - Installs IRC client bot
  - zombie bot “calls home” periodically
- Report
  - Recurrent traffic to suspect IRC servers
- Analysis
  - Backtrack in time to find initial compromise
  - Observe tool download and installation
  - Pivot …

IRC bot: Timeline Investigation



The Event Table



75.64.71.22 and not(l_role > 0 and d_port=80)

2006-08-21 10:00:01 to 2006-08-21 10:58:28
0.0
75.64.71.22
and not(l_role > 0 and d_port=80)

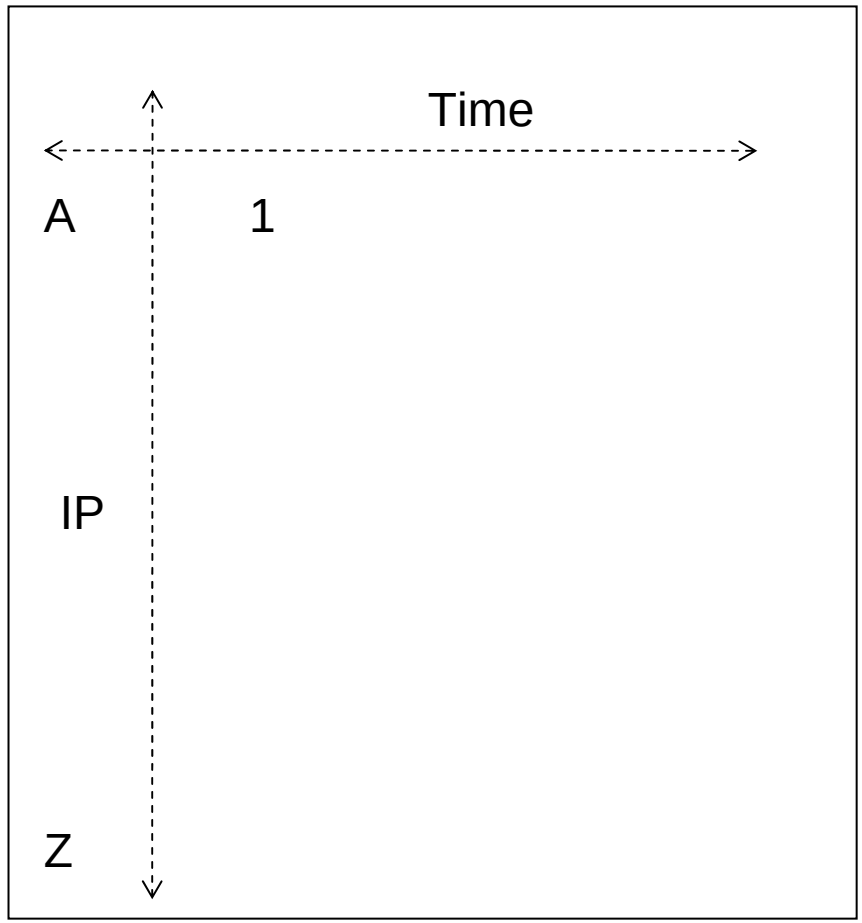
l_weekday	l_hour	GMTfirst	duration	locality	l_role	proto	l_asn	l_vln	inet_ntoa(l_ipn)	l_port	r_asn	r_vln	inet_ntoa(r_ipn)	r_port	d_port	l_pkt	l_byte	l_abyte	r_pkt	r_byte	r_abyte
2	3	2006-08-21 10:00:01	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:00:03	0	3	3	17	32	71	75.64.71.22	37373	32	64	75.64.67.192	37373	37373	0	0	0	1	177	13
2	3	2006-08-21 10:00:03	0.001	2	3	17	32	71	75.64.71.22	7001	32	16401	75.64.15.96	7001	7001	2	140	56	4	450	28
2	3	2006-08-21 10:00:04	0.009	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:00:18	0.506	2	2	6	32	71	75.64.71.22	22	32	17116	75.66.189.156	22	22	2	236	128	2	108	
2	3	2006-08-21 10:01:02	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:01:17	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:01:18	0.441	2	2	6	32	71	75.64.71.22	22	32	17116	75.66.189.156	22	22	2	236	128	2	108	
2	3	2006-08-21 10:01:45	0.001	1	3	1	32	71	75.64.71.22	0	15243	7936	147.31.67.105	0	0	2	184	116	2	184	11
2	3	2006-08-21 10:01:49	0.109	2	-3	6	32	71	75.64.71.22	45075	32	17174	75.67.9.109	45075	25	8	551	111	14	1202	43
2	3	2006-08-21 10:01:52	0.001	2	3	17	32	71	75.64.71.22	7001	32	16401	75.64.15.111	7001	7001	2	140	56	4	450	28
2	3	2006-08-21 10:01:54	27.002	2	-3	17	32	71	75.64.71.22	37396	32	7	75.64.24.227	37396	53	2	148	64	4	296	12
2	3	2006-08-21 10:01:59	19.998	2	-3	17	32	71	75.64.71.22	37396	32	7	75.64.24.201	37396	53	2	148	64	4	296	12
2	3	2006-08-21 10:02:03	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:02:06	4.99	2	-1	17	32	71	75.64.71.22	7001	32	14	75.64.22.185	7001	7000	8	658	322	0	0	
2	3	2006-08-21 10:02:13	7	1	-1	17	32	71	75.64.71.22	7001	3	17920	18.70.0.6	7001	7003	10	872	452	0	0	
2	3	2006-08-21 10:02:18	0.379	2	2	6	32	71	75.64.71.22	22	32	17116	75.66.189.156	22	22	2	236	128	2	108	
2	3	2006-08-21 10:02:19	0	2	-1	1	32	71	75.64.71.22	3	32	7	75.64.24.201	3	3	2	204	136	0	0	
2	3	2006-08-21 10:02:19	2.004	2	-3	17	32	71	75.64.71.22	37401	32	7	75.64.24.227	37401	53	1	74	32	2	148	6
2	3	2006-08-21 10:02:20	5.5	1	-1	17	32	71	75.64.71.22	7001	3	37120	18.145.0.25	7001	7003	8	724	388	0	0	
2	3	2006-08-21 10:02:21	0	2	-1	17	32	71	75.64.71.22	37402	32	7	75.64.24.201	37402	53	1	74	32	0	0	
2	3	2006-08-21 10:02:21	0	2	-1	1	32	71	75.64.71.22	3	32	7	75.64.24.227	3	3	2	204	136	0	0	
2	3	2006-08-21 10:02:27	0	2	-1	17	32	71	75.64.71.22	37403	32	7	75.64.24.201	37403	53	1	74	32	0	0	
2	3	2006-08-21 10:02:30	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:02:51	0	2	-1	1	32	71	75.64.71.22	3	32	7	75.64.24.201	3	3	1	102	68	0	0	
2	3	2006-08-21 10:02:51	0	2	-1	1	32	71	75.64.71.22	3	32	7	75.64.24.201	3	3	1	102	68	0	0	
2	3	2006-08-21 10:02:51	0	2	1	17	32	71	75.64.71.22	37402	32	7	75.64.24.201	37402	37402	0	0	0	2	148	6
2	3	2006-08-21 10:02:51	0	2	1	17	32	71	75.64.71.22	37403	32	7	75.64.24.201	37403	37403	0	0	0	2	148	6
2	3	2006-08-21 10:03:04	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:03:16	0.546	3	3	6	32	71	75.64.71.22	25	32	64	77.232.79.23	25	25	9	937	331	9	691	10
2	3	2006-08-21 10:03:18	0.548	2	2	6	32	71	75.64.71.22	22	32	17116	75.66.189.156	22	22	2	236	128	2	108	
2	3	2006-08-21 10:03:25	0.224	3	3	6	32	71	75.64.71.22	25	32	64	77.232.79.23	25	25	13	1453	583	13	3643	280
2	3	2006-08-21 10:03:32	4.285	2	-3	6	32	71	75.64.71.22	45075	32	17174	75.67.9.109	45075	25	33	2391	609	54	4072	115
2	3	2006-08-21 10:03:43	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:04:05	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:04:18	0.267	2	2	6	32	71	75.64.71.22	22	32	17116	75.66.189.156	22	22	2	236	128	2	108	
2	3	2006-08-21 10:04:56	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:05:06	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:05:18	0.332	2	2	6	32	71	75.64.71.22	22	32	17116	75.66.189.156	22	22	2	236	128	2	108	
2	3	2006-08-21 10:05:27	5	2	-1	17	32	71	75.64.71.22	7001	32	14	75.64.22.185	7001	7000	8	658	322	0	0	
2	3	2006-08-21 10:05:34	7	1	-1	17	32	71	75.64.71.22	7001	3	17920	18.70.0.6	7001	7003	10	872	452	0	0	
2	3	2006-08-21 10:05:41	5.5	1	-1	17	32	71	75.64.71.22	7001	3	37120	18.145.0.25	7001	7003	8	724	388	0	0	
2	3	2006-08-21 10:06:07	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:06:09	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:06:18	0.232	2	2	6	32	71	75.64.71.22	22	32	17116	75.66.189.156	22	22	2	236	128	2	108	
2	3	2006-08-21 10:06:39	52.231	1	3	6	32	71	75.64.71.22	25	4837	46336	61.181.226.22	25	25	34	2916	1064	54	36446	3351
2	3	2006-08-21 10:06:44	27.013	2	-3	17	32	71	75.64.71.22	37455	32	7	75.64.24.227	37455	53	3	258	132	6	516	26
2	3	2006-08-21 10:06:49	24.014	2	-3	17	32	71	75.64.71.22	37455	32	7	75.64.24.201	37455	53	2	172	88	4	344	17
2	3	2006-08-21 10:06:49	8.542	1	3	6	32	71	75.64.71.22	25	24544	21008	203.82.17.141	25	25	32	3326	1190	34	5650	346
2	3	2006-08-21 10:06:57	0.296	2	-3	6	32	71	75.64.71.22	45079	32	17174	75.67.9.109	45079	25	13	2565	1855	34	2278	42
2	3	2006-08-21 10:07:08	0.002	1	-3	1	32	71	75.64.71.22	8	26101	24288	66.94.230.32	8	0	2	196	128	2	196	12
2	3	2006-08-21 10:07:09	4.01	2	-3	17	32	71	75.64.71.22	37464	32	7	75.64.24.201	37464	53	1	86	44	2	172	8
2	3	2006-08-21 10:07:11	0	2	-1	1	32	71	75.64.71.22	3	32	7	75.64.24.227	3	3	3	342	240	0	0	
2	3	2006-08-21 10:07:13	0	2	-1	1	32	71	75.64.71.22	3	32	7	75.64.24.201	3	3	2	228	160	0	0	
2	3	2006-08-21 10:07:18	59.999	2	2	6	32	71	75.64.71.22	22	32	17116	75.66.189.156	22	22	4	472	256	2	108	

From Event Table to Event Plot

Event Table

1	Time	A	…	Measures

Event Plot

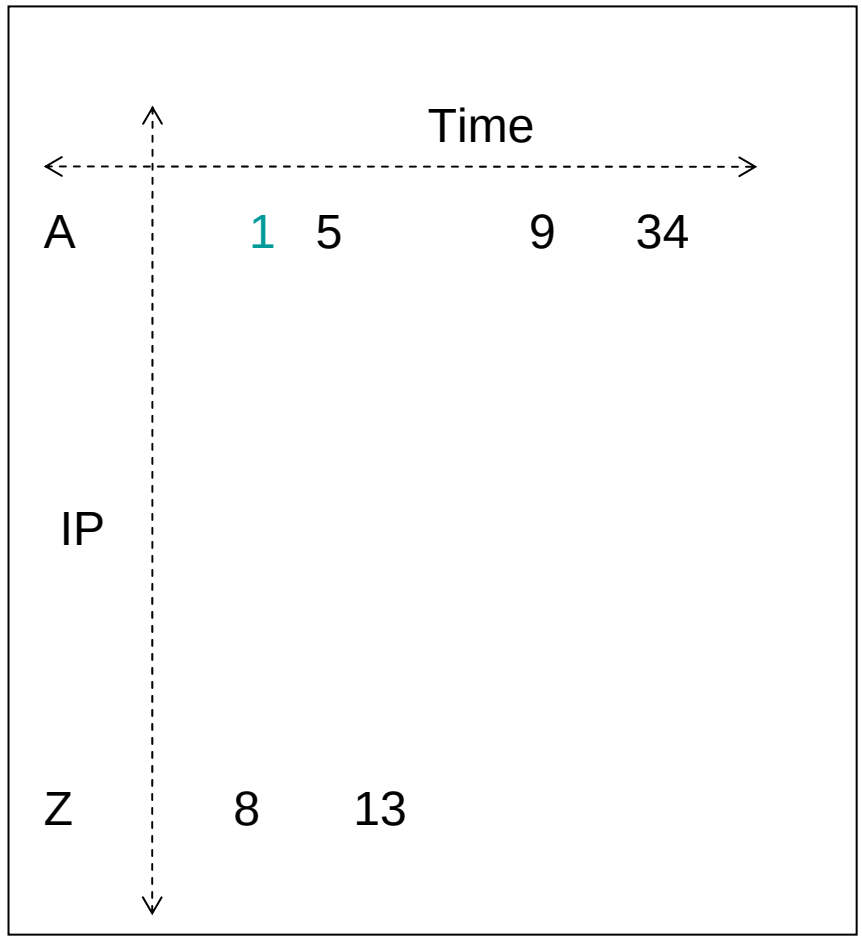


From Event Table to Event Plot

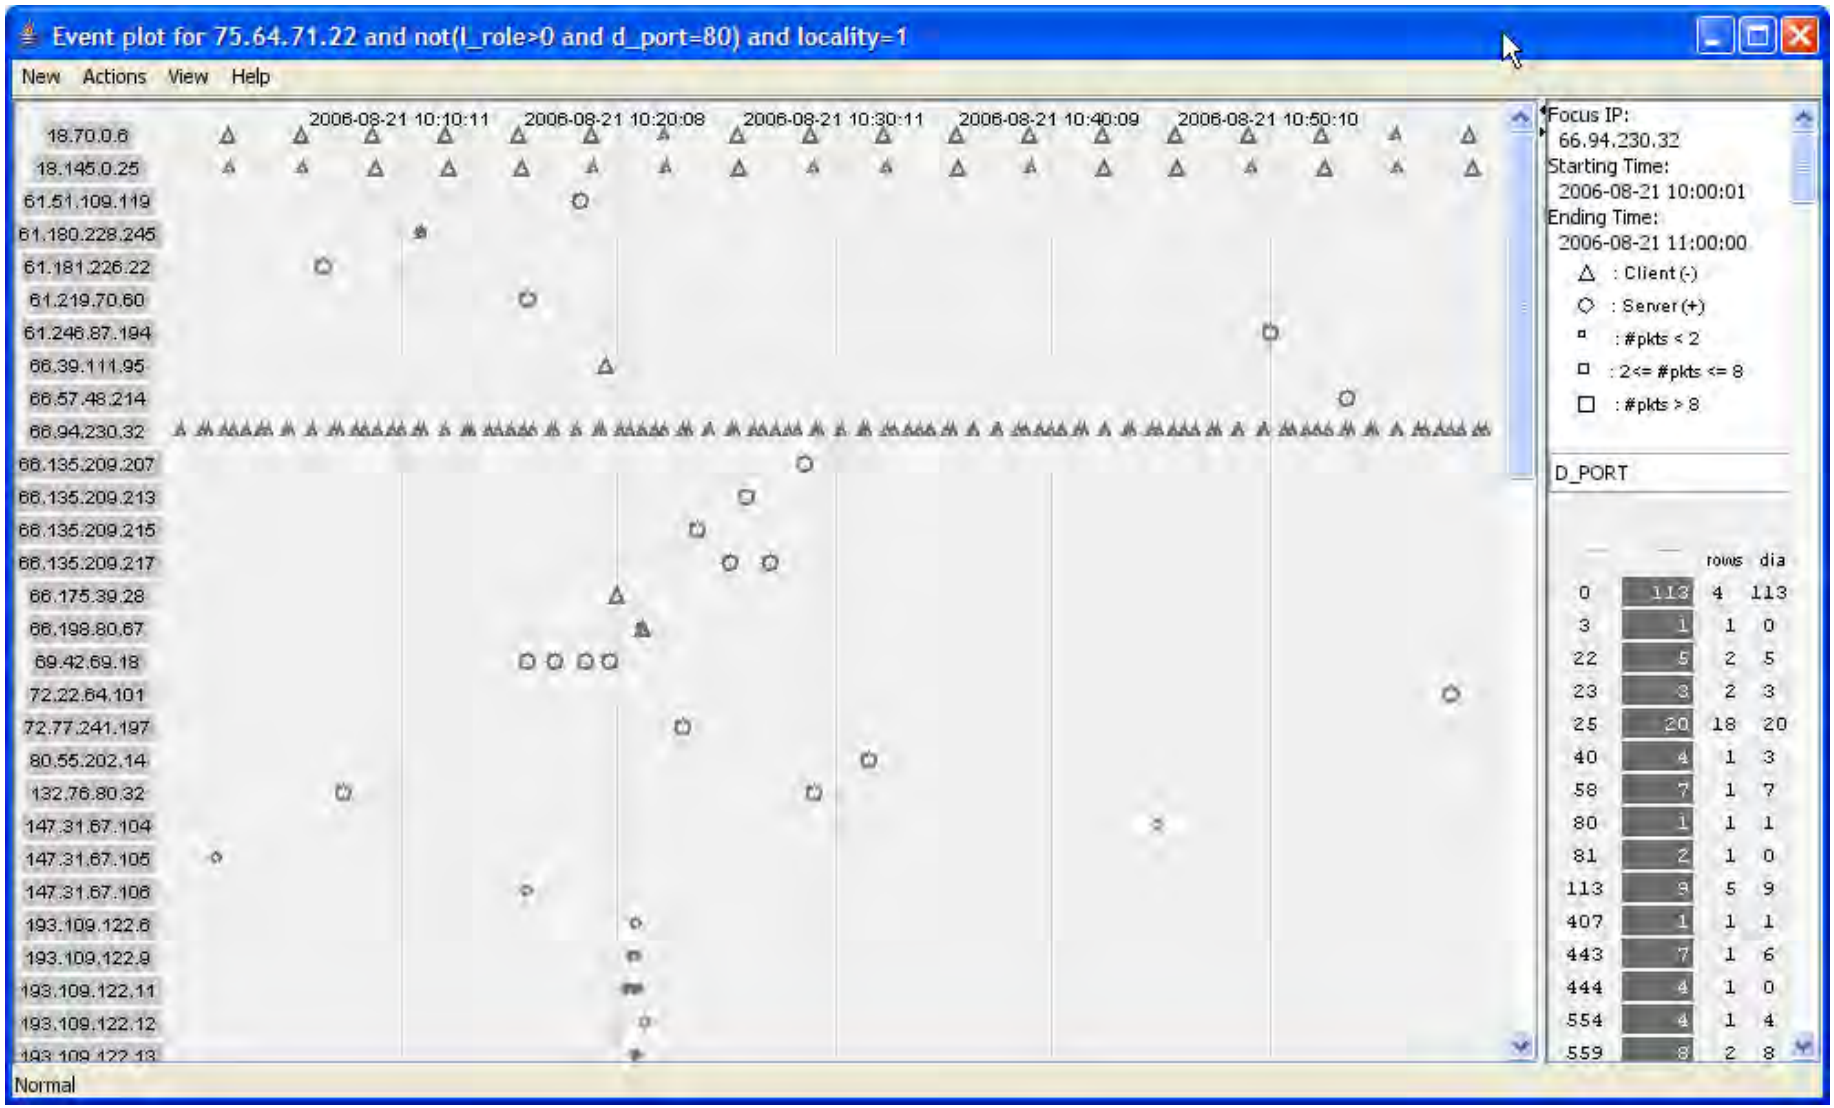
Event Table

1	Time	A	…	Measures
. .				
#	Time	IP	…	Measures
. .				
n	Time	Z	…	Measures

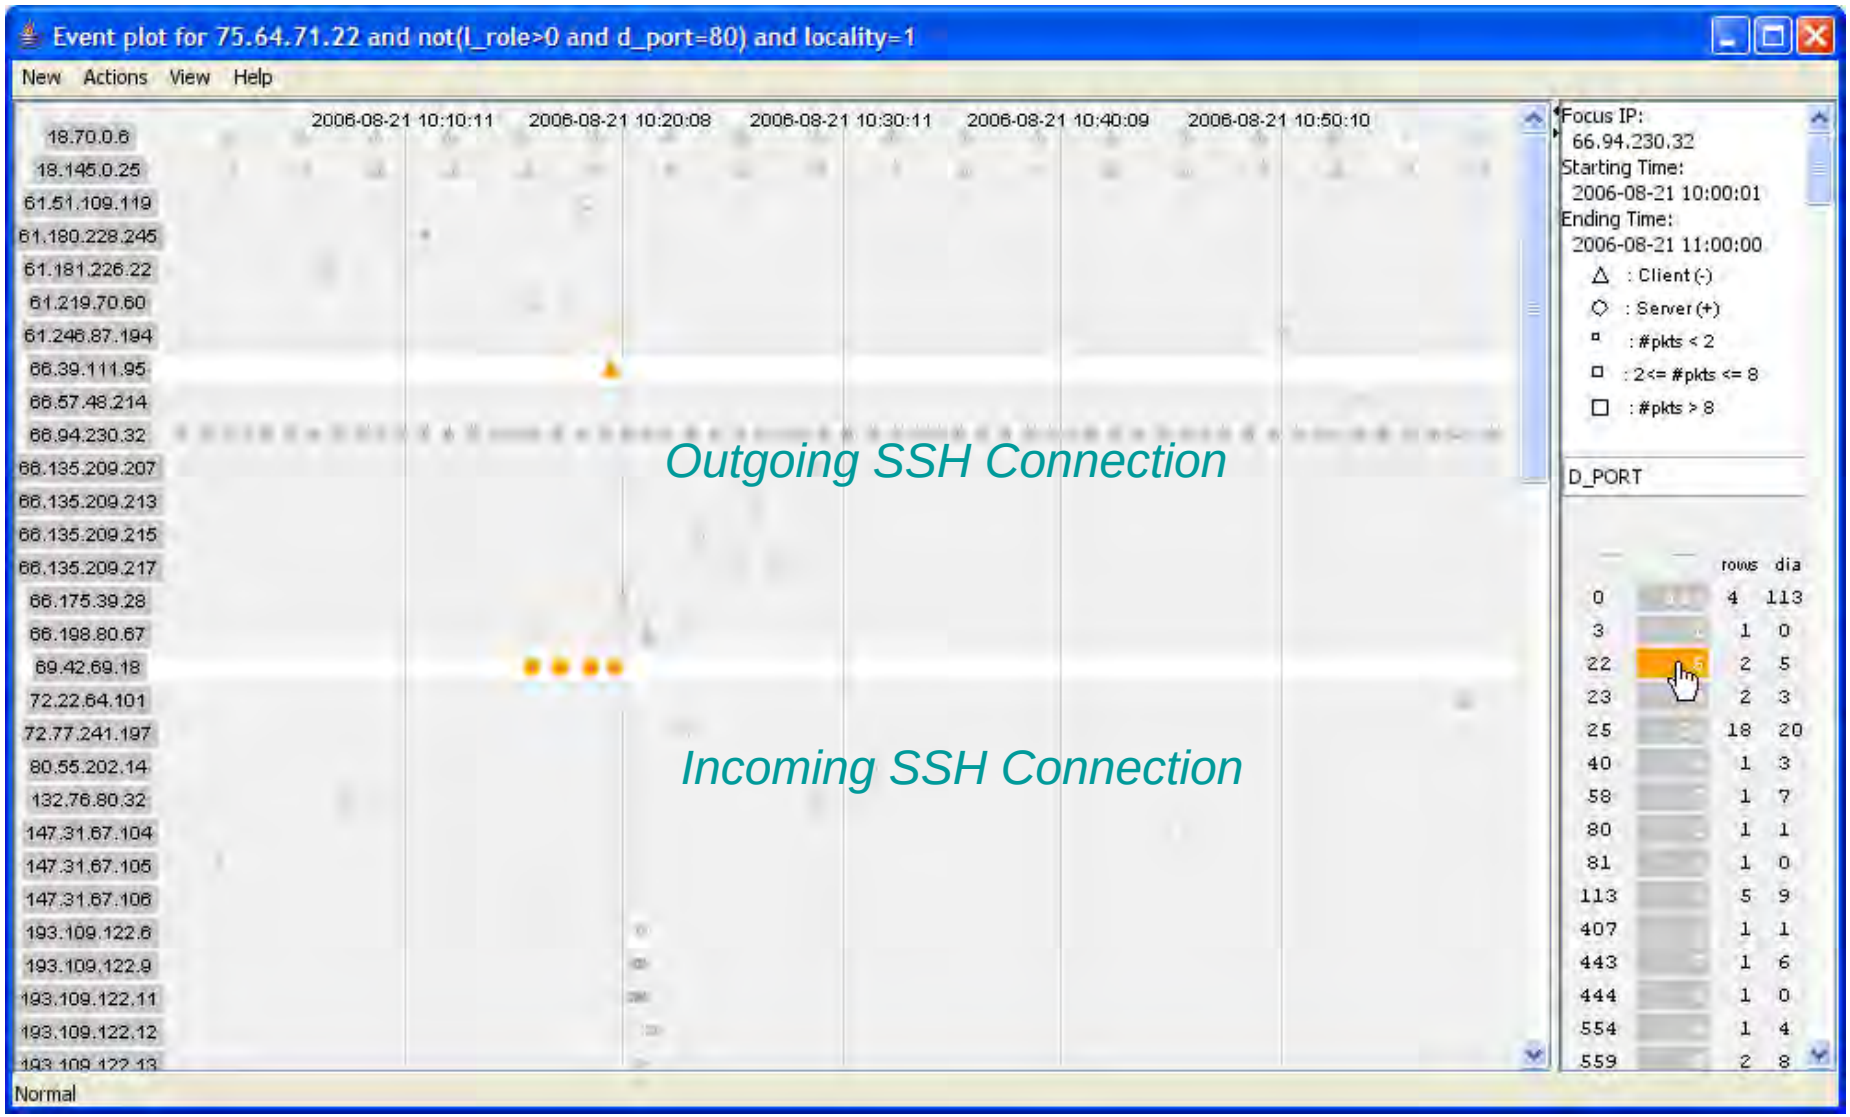
Event Plot



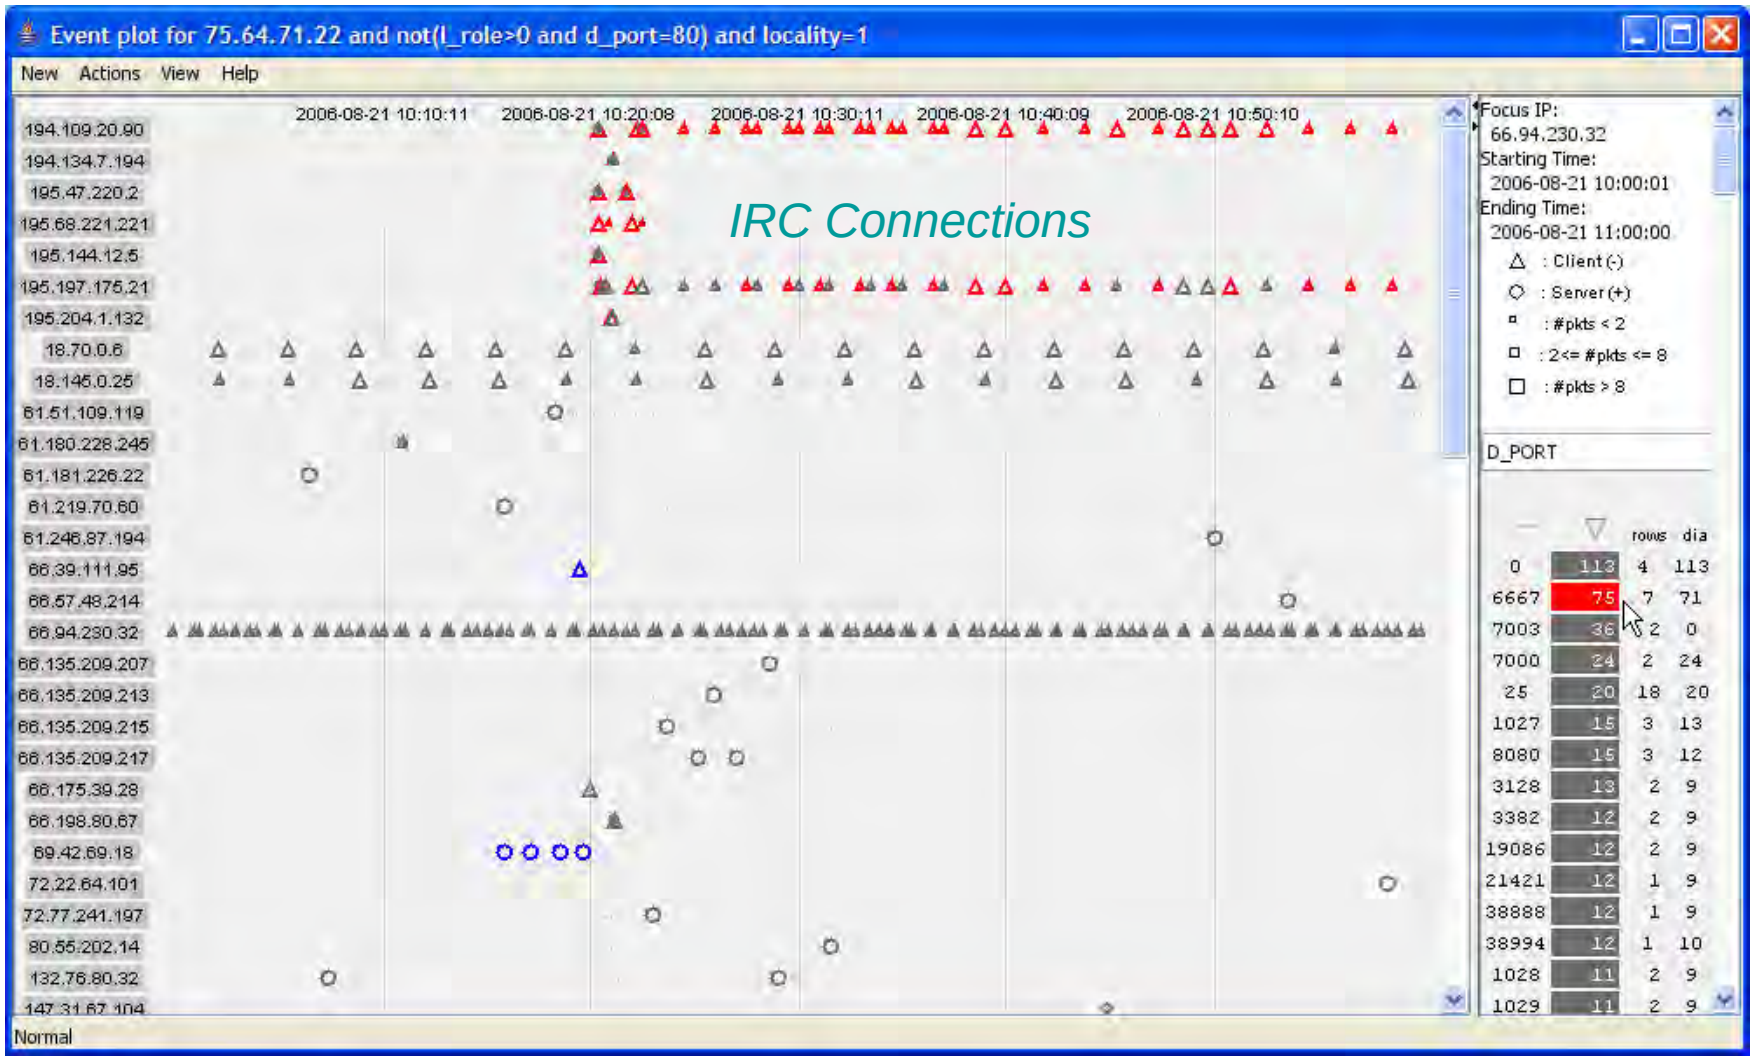
Event Plot


Event plot for 75.64.71.22 and not(l_role>0 and d_port=80) and locality=1
New
Actions
View
Help
2006-08-21 10:10:11
2006-08-21 10:20:08
2006-08-21 10:30:11
2006-08-21 10:40:09
2006-08-21 10:50:10
18.70.0.6
18.145.0.25
61.51.109.119
61.180.228.245
61.181.226.22
61.219.70.60
61.246.87.194
66.39.111.95
66.57.48.214
66.94.230.32
66.135.209.207
66.135.209.213
66.135.209.215
66.135.209.217
66.175.39.28
66.198.80.67
69.42.69.18
72.22.64.101
72.77.241.197
80.55.202.14
132.76.80.32
147.31.67.104
147.31.67.105
147.31.67.106
193.109.122.6
193.109.122.9
193.109.122.11
193.109.122.12
193.109.122.13
Focus IP:
66.94.230.32
Starting Time:
2006-08-21 10:00:01
Ending Time:
2006-08-21 11:00:00
: Client (-)
: Server (+)
: #pkts < 2
: 2<= #pkts <= 8
: #pkts > 8
D_PORT
rows
dia
0 113 4 113
3 1 1 0
22 5 2 5
23 3 2 3
25 20 18 20
40 4 1 3
58 7 1 7
80 1 1 1
81 2 1 0
113 9 5 9
407 1 1 1
443 7 1 6
444 4 1 0
554 4 1 4
559 8 2 8
Normal

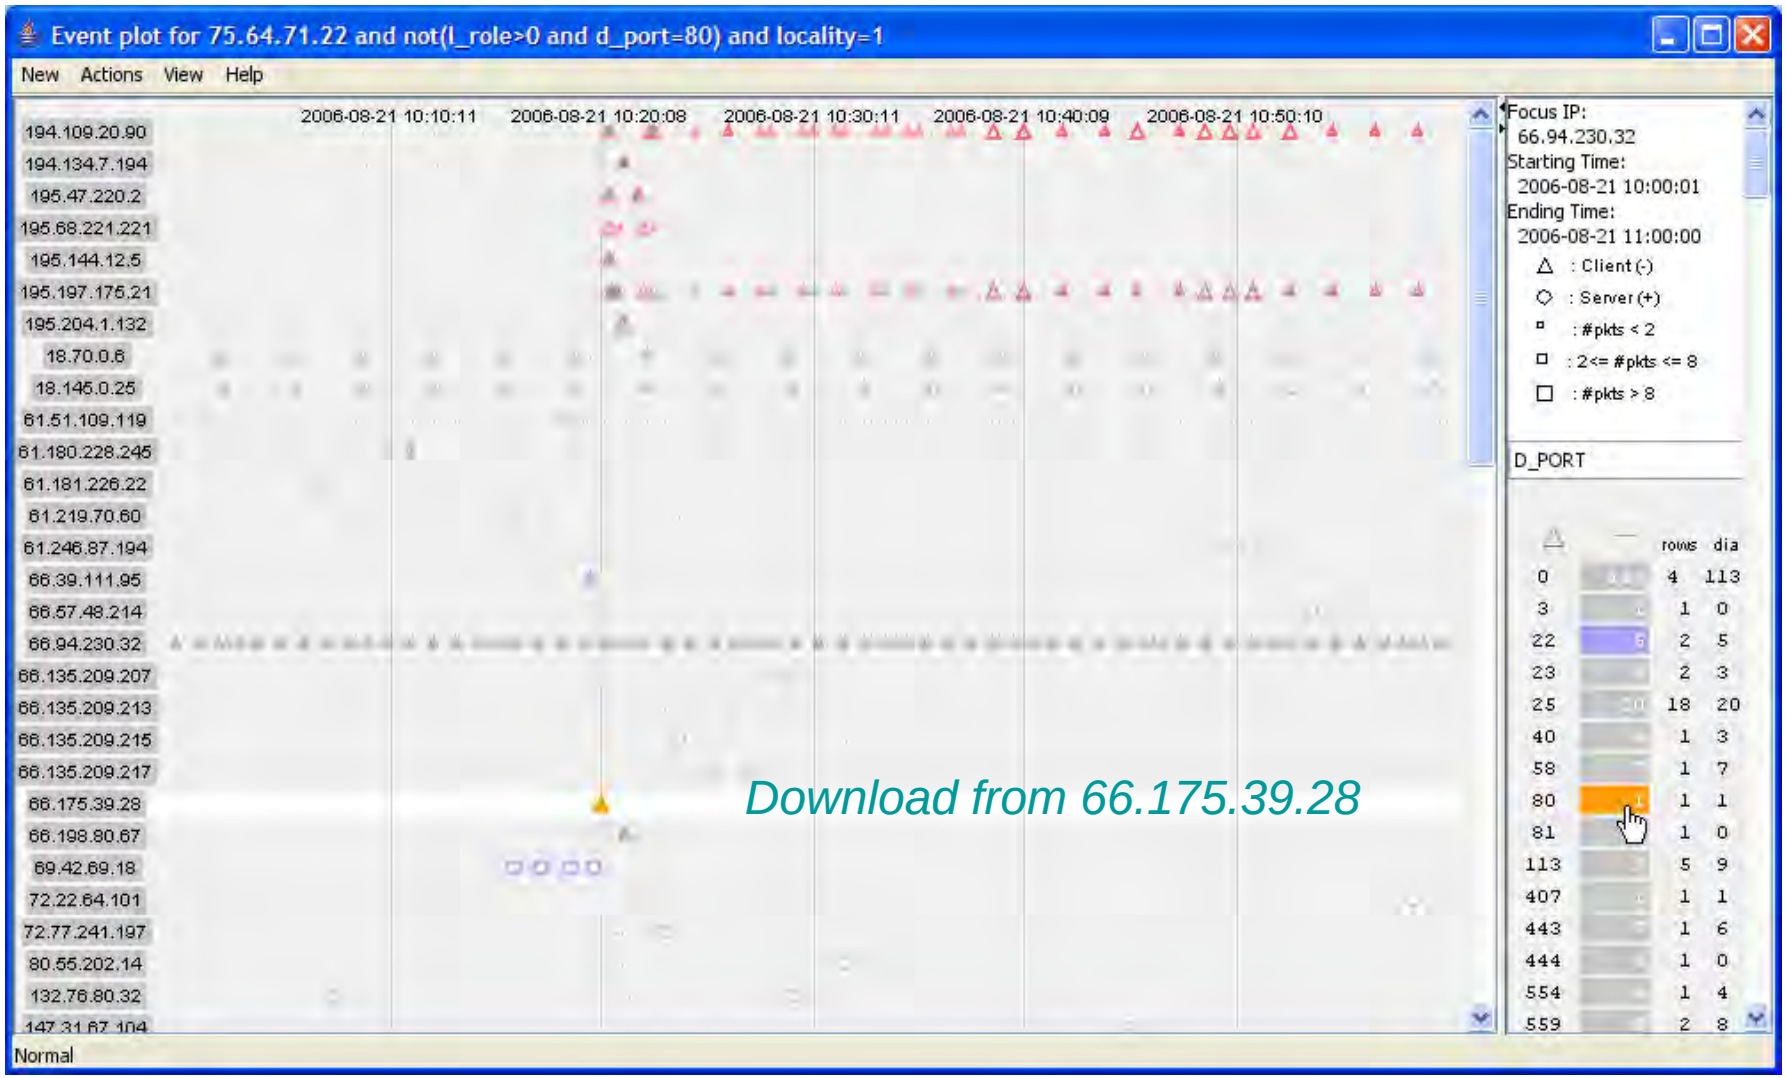
IRC Bot: Initial SSH Connection



IRC Traffic on port 6667



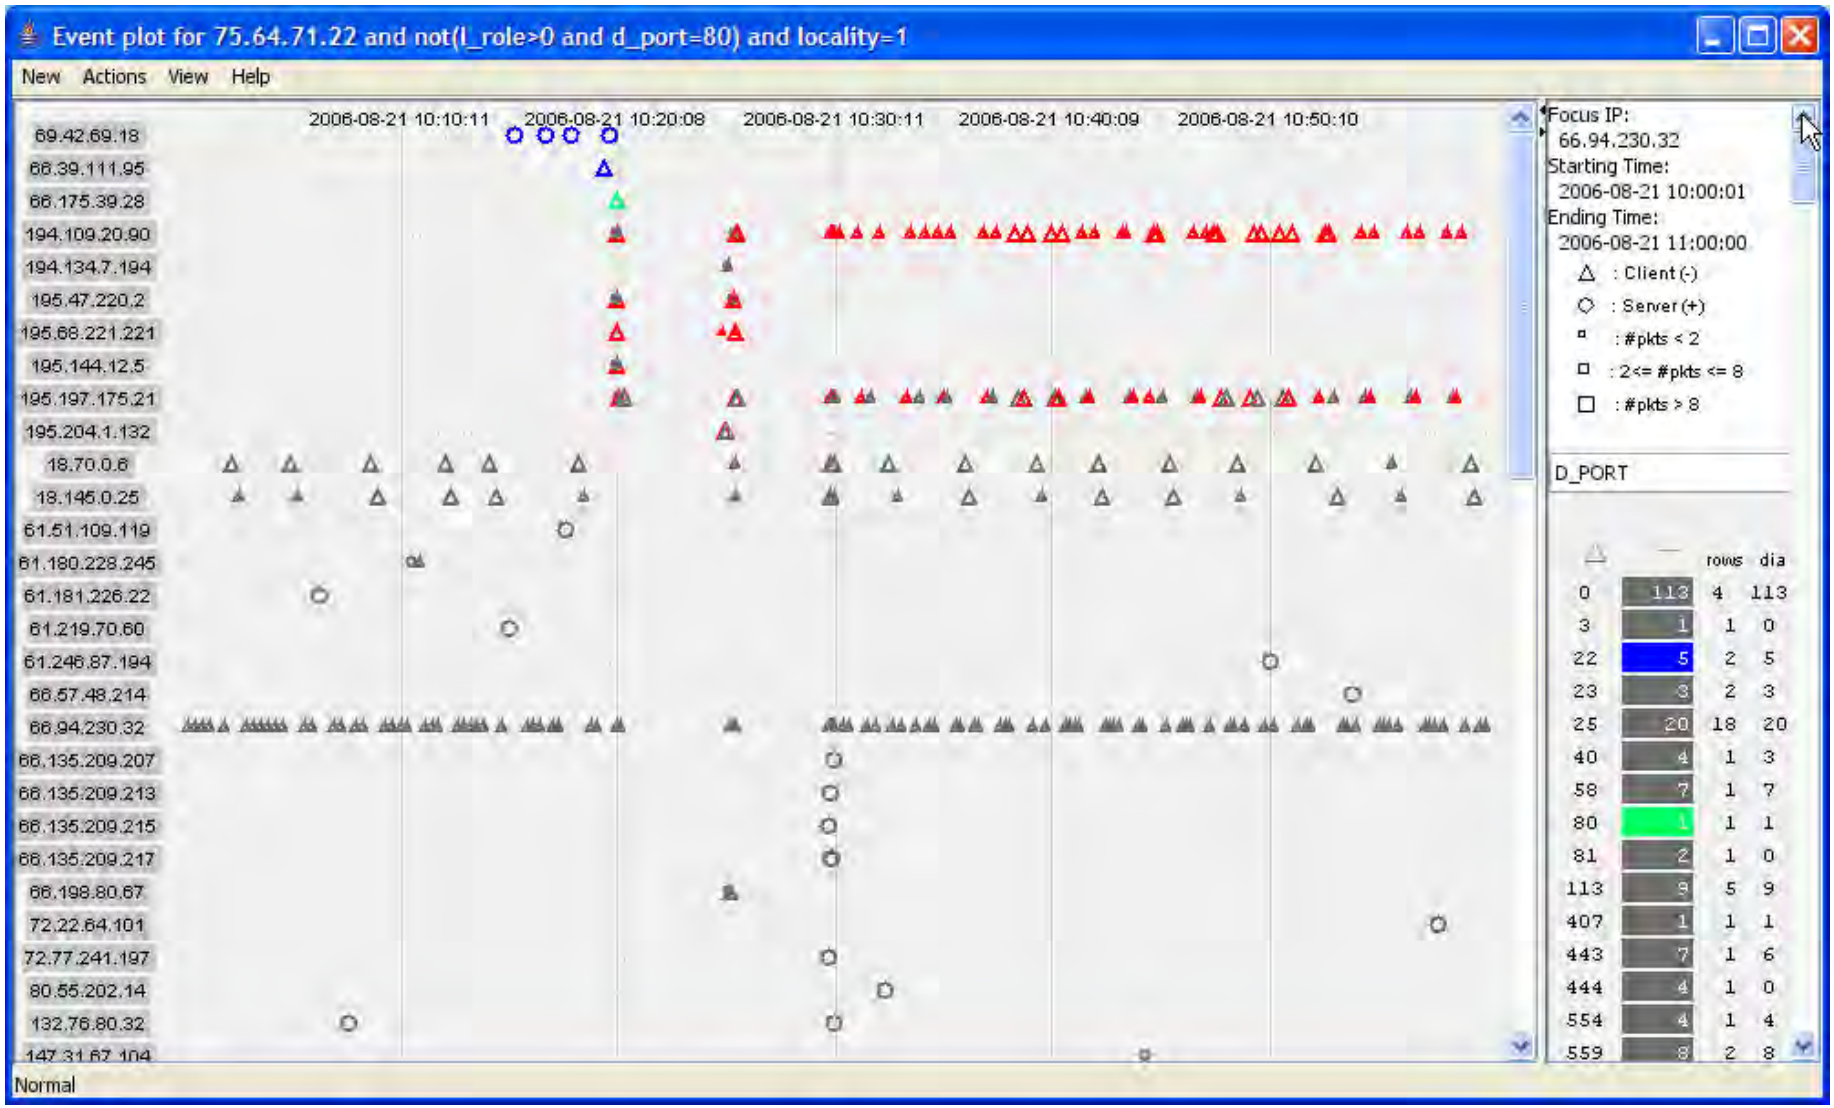
Download of Intrusion Tools



Reordered Rows



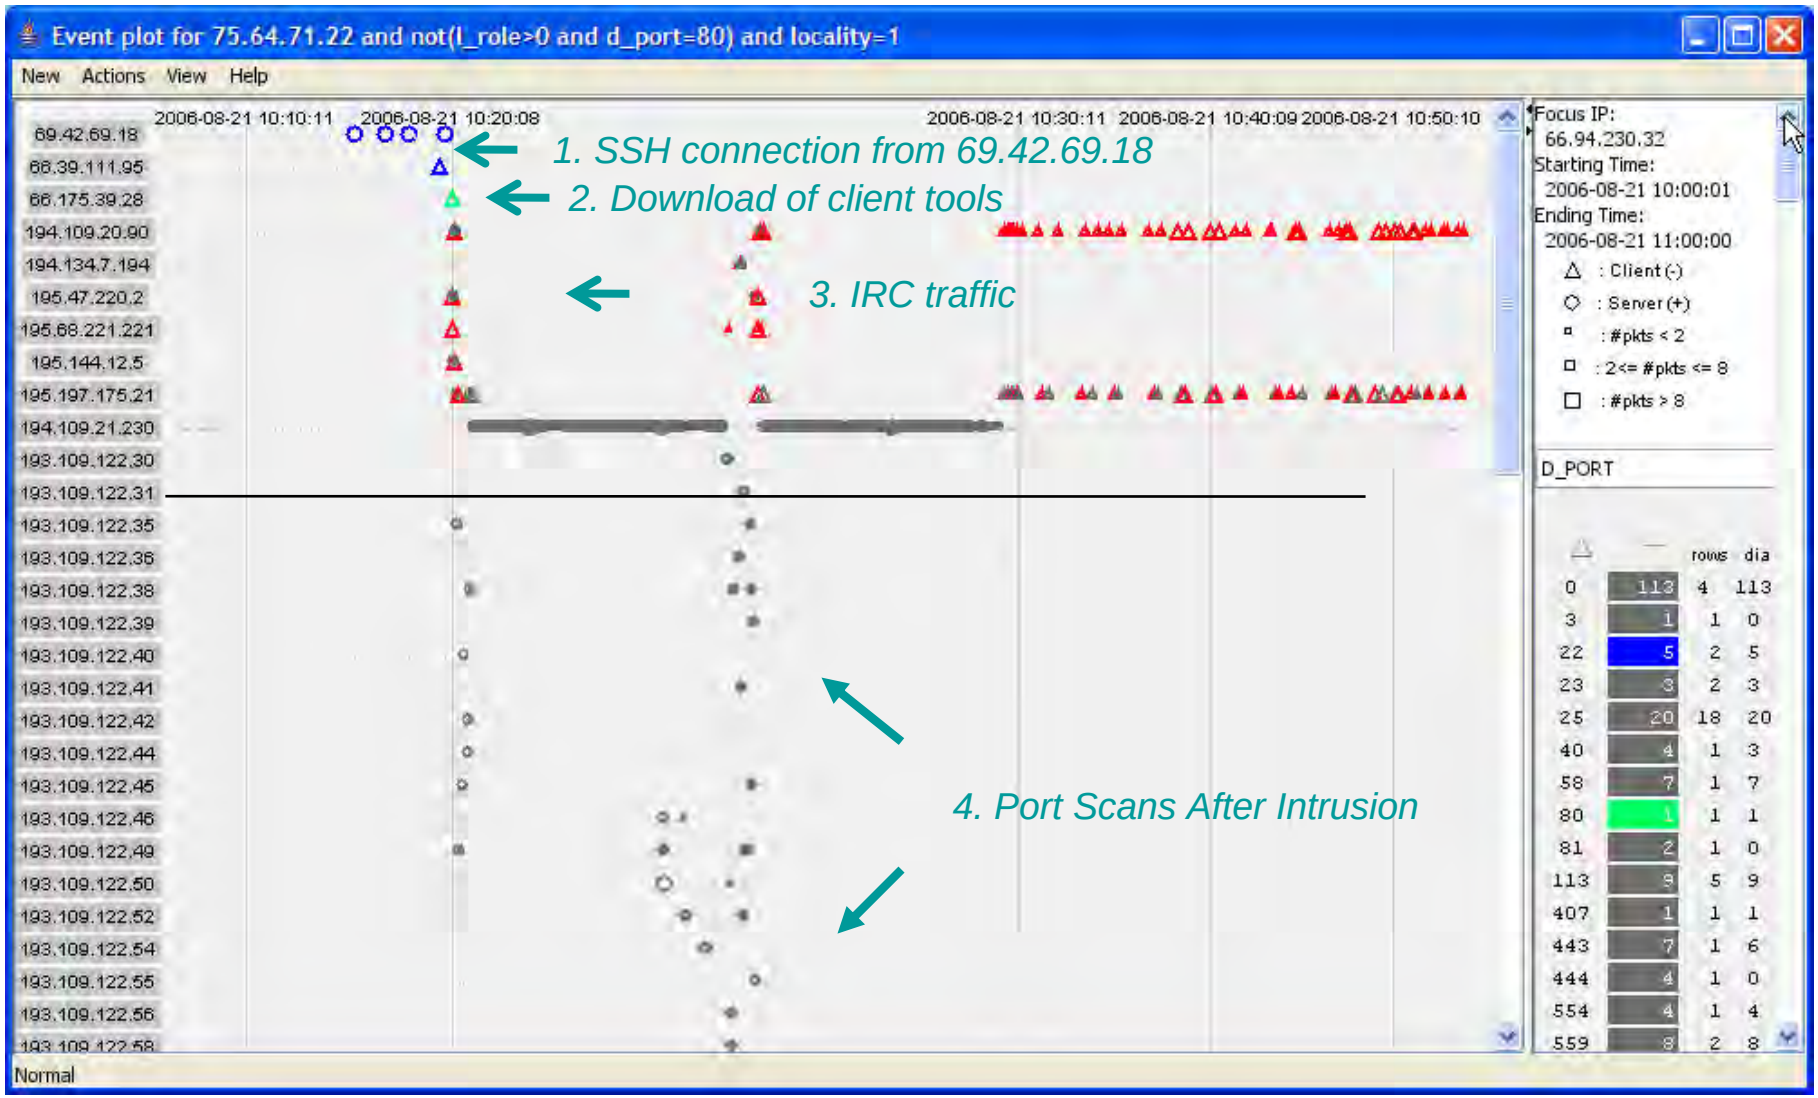
Switch to Ordinal Time



Mine the Gap


Event plot for 75.64.71.22 and not(l_role>0 and d_port=80) and locality=1
New Actions View Help
193.109.122.21
193.109.122.22
193.109.122.23
193.109.122.24
193.109.122.25
193.109.122.26
193.109.122.28
193.109.122.30
193.109.122.31
193.109.122.35
193.109.122.36
193.109.122.38
193.109.122.39
193.109.122.40
193.109.122.41
193.109.122.42
193.109.122.44
193.109.122.45
193.109.122.46
193.109.122.49
193.109.122.50
193.109.122.52
193.109.122.54
193.109.122.55
193.109.122.56
193.109.122.58
193.109.122.59
194.109.21.230
203.82.17.141
2006-08-21 10:20:56 6.193130016326904 local: 2 3
d_port: 10130 role: 2 locality: 1 proto: 6
IP vln asn port pkt byte abyte
l: 75.64.71.22 71 32 10130 1 54 0
r: 194.109.21.230 27904 3265 10130 6 444 0
10080 1 3
10099 2 9
10130 1 6
10242 1 4
10267 1 0
10777 2 9
11225 1 6
11904 1 0
12000 1 0
12047 1 0
12487 1 0
12488 1 0
12489 1 0
12490 1 0
12491 1 0
12492 1 0
12493 1 0
12494 1 0
12495 1 0
12496 1 0
12497 1 0
12499 1 0
12500 1 0
12501 1 0
12502 1 0
12503 1 0
12504 1 0
12505 1 0
12506 1 0
Normal

Sequence of Intrusion



Future Work

- Scalable query performance
  - Want to query billion row tables at interactive speeds
  - Column-oriented database
  - Distribute across commodity cluster
- Finding network signatures
  - Bottom up capture of analyst domain knowledge
    (see our paper by Xiao in VAST 2006)
  - Top down search for frequent patterns
  - Build disparate flows into behaviors (boot, logon, mail, print, surf, …)
- Modeling Local Machine Behavior
  - Shift the burden to the attacker?

The *Ripple* decoded

Carrie Gates

CA Labs

John McHugh

Canada Research Chair in Privacy and Security

Dalhousie University

mchugh@cs.dal.ca

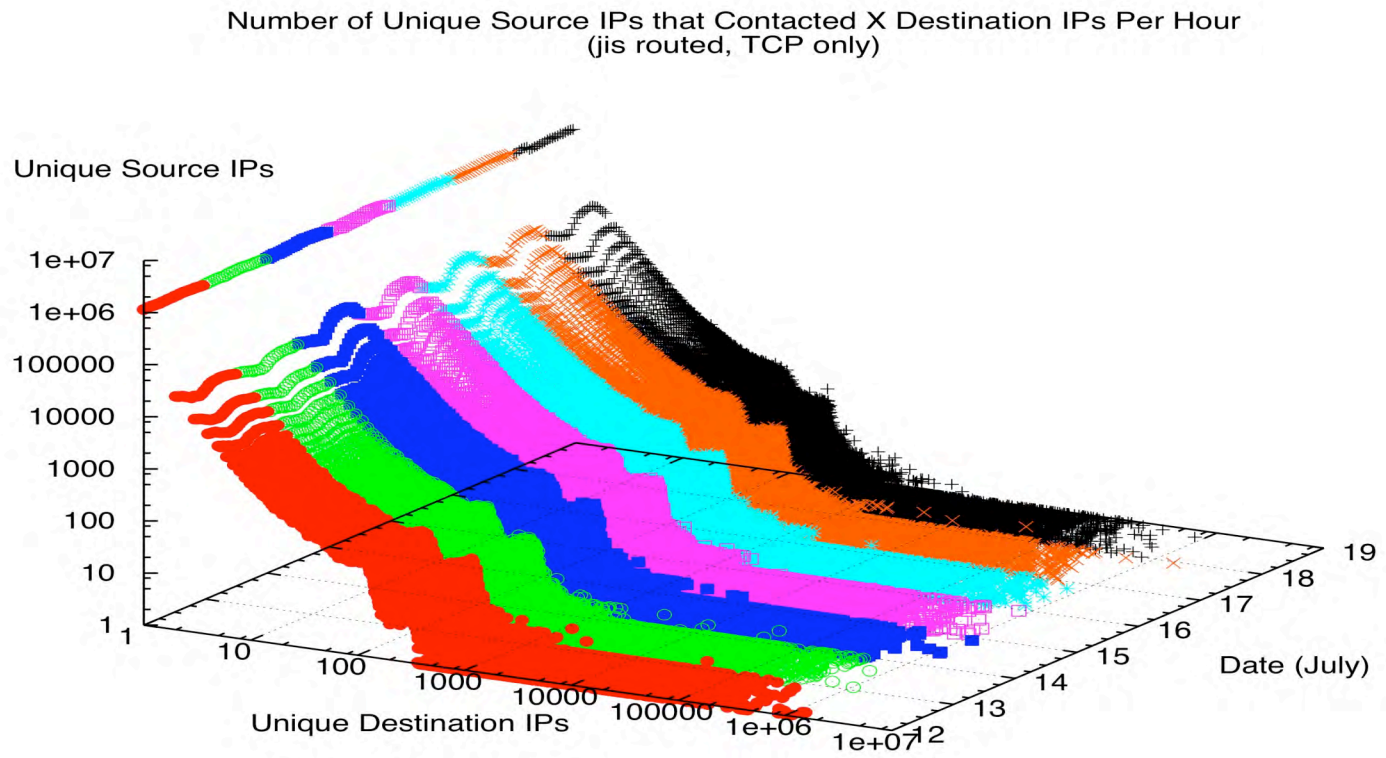
© 2007 by John McHugh, Carrie Gates

Very large scale observation

- Carrie Gates was interested in the degree of fan out from outside to inside for her scan detection work.
- How many outside hosts use exactly one inside host / service pair. (unique destination address/port)
- In the beginning, we did it the hard way, but Bloom filters can be used to find unique sIP,dIP,dport exemplar flows
- If we make a source IP bag from the exemplar flows, the counts will be the number of different host / service pairs contacted by a given source host.
- Invert the bag to determine how many entries have a count of 1, 2, 3, ... Plot hourly results for a week

© 2005,6,7 by John McHugh, Carrie Gates

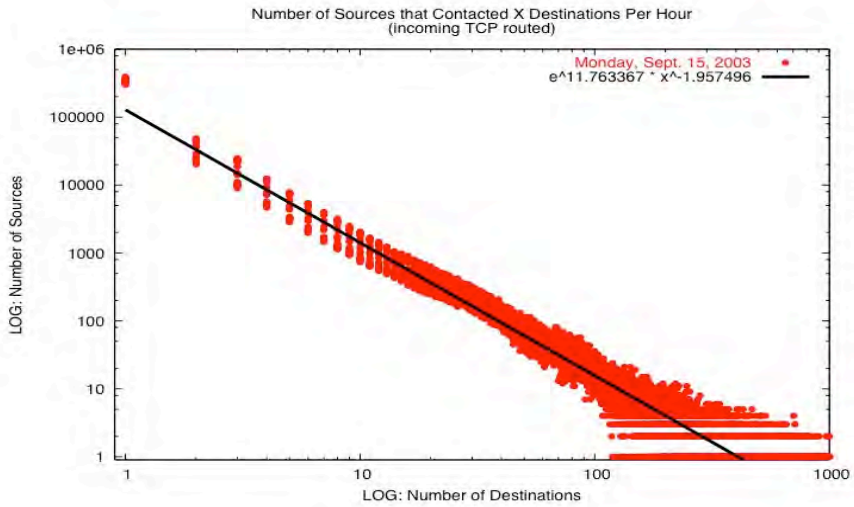
Outside to inside - July 2003



© 2005,6,7 by John McHugh, Carrie Gates

Developing the contact surface

- In the absence of the disturbance seen on the previous page, contact lines seem to follow a power law type of distribution
  - or do they[1].
  - We think this is really at least 3 separate processes
    - VLF noise
    - "normal activity"
    - Bulk scanning



[1] everything is a straight line on log/log paper, especially if you use a fat marker

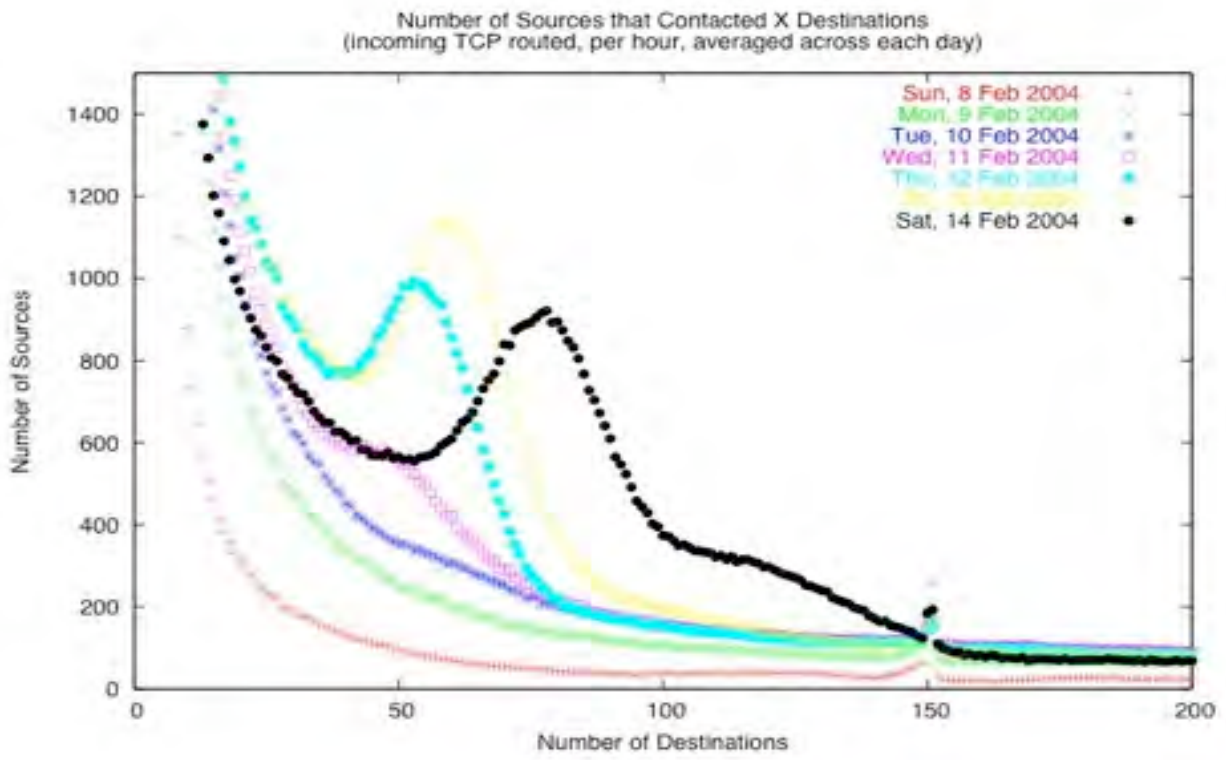
© 2005,6,7 by John McHugh, Carrie Gates

Internet wide disturbance

- The ripple in what would otherwise be a fairly straight log/log plot of connectivity was observed from at least Jan - Aug 2003.
- It went away when Blaster appeared in Aug 2003.
- A similar ripple existed from Feb 11 to May 31 2004 coinciding with the lifetime of Welchia-B
  - In this case, the ripple is due to a few hundred machines scanning at a low, fixed, rate induced by a loop with a “sleep” system call.
- In both cases, they persisted until killed, not patched.
- We have been told that the ripple is back.

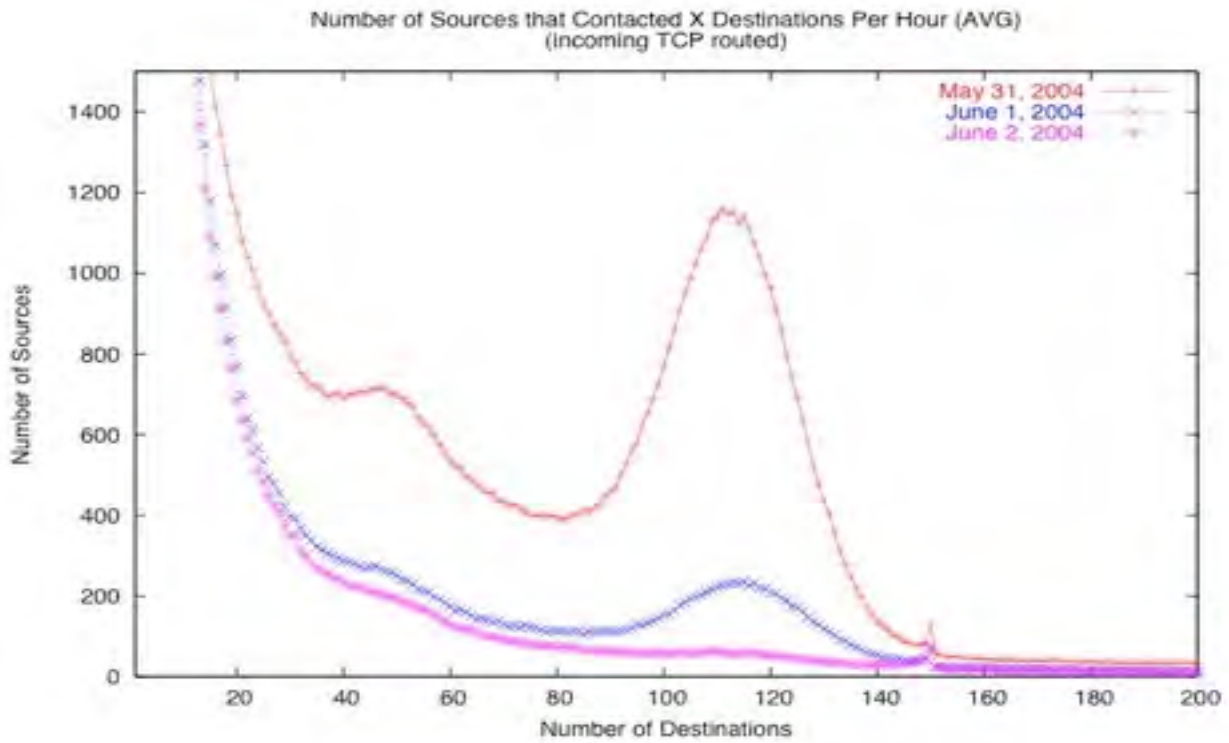
© 2005,6,7 by John McHugh, Carrie Gates

Details of the Welchia.B event - onset



© 2005,6,7 by John McHugh, Carrie Gates

Details of Welchia.B - demise



© 2005,6,7 by John McHugh, Carrie Gates

Design Time Coordination

- The sleep in the scan loop of Welchia.B points to a form of loose, design time, coordination.
- All members of the cohort scan at approximately the same rate, using the same random generation scheme but with a different random seed.
- If we captured all the scans from each member of the cohort, we would expect to see a small, tight, cluster of scanners all contacting nearly the same number of targets.
- We observe only a small portion of the address space and see a small percentage of the scans from each host with substantial interhost variation.

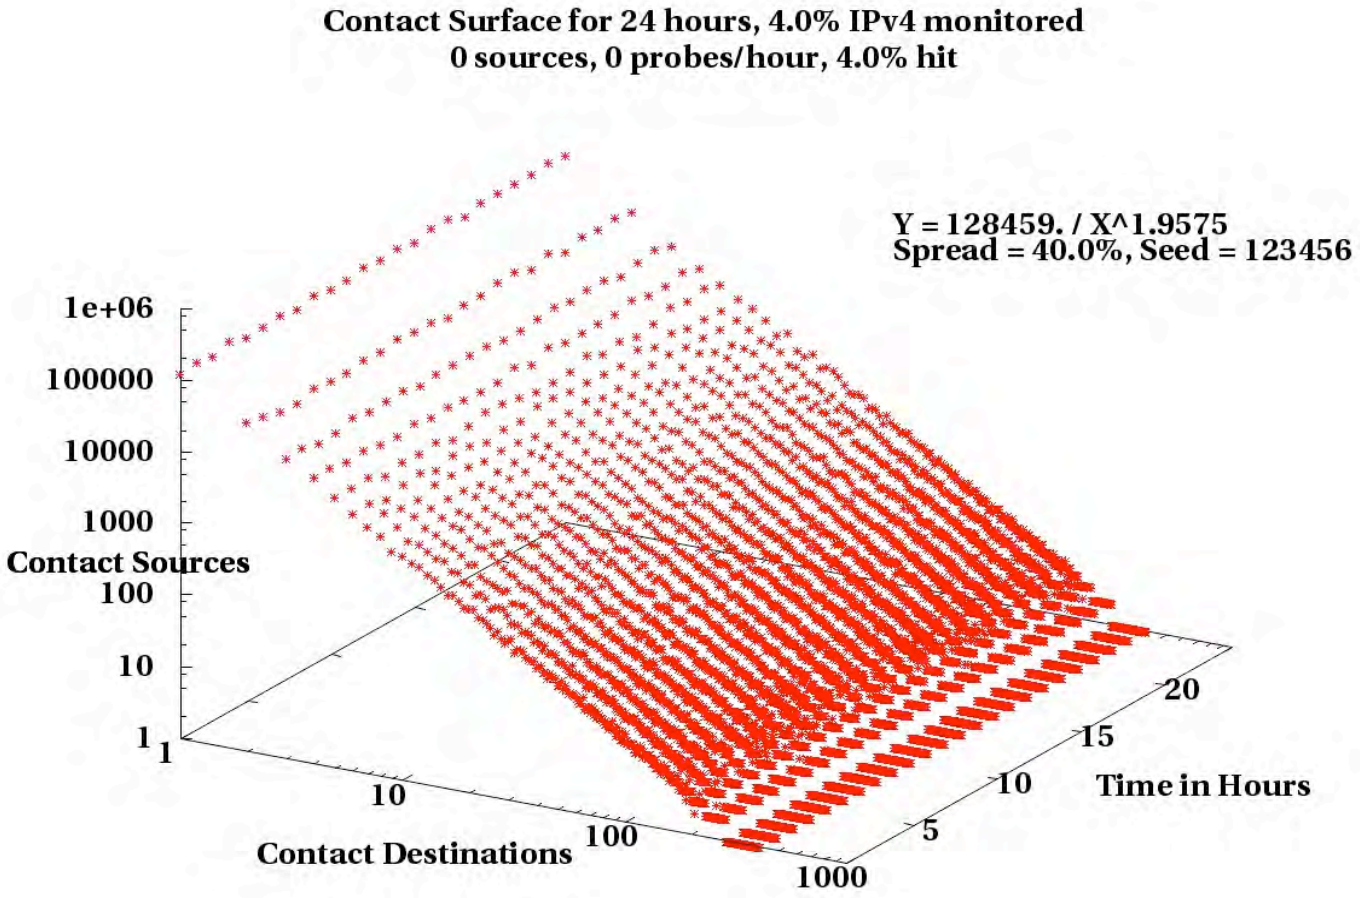
© 2005,6,7 by John McHugh, Carrie Gates

This fall, we simulated the perturbations

- Generated approximation of unperturbed background
  - Dont care about process, only appearance
- Simulated perturbation process parameterized on:
  - Number of sources
  - Probe rate / source
  - % of IPv4 monitored
  - % of probes intercepted
    - For ripple or wave, % monitored = % intercepted
    - For scans targeting monitored network they are different
- Looked at observability as a function of parameters.

© 2005,6,7 by John McHugh, Carrie Gates

Background only - main line process

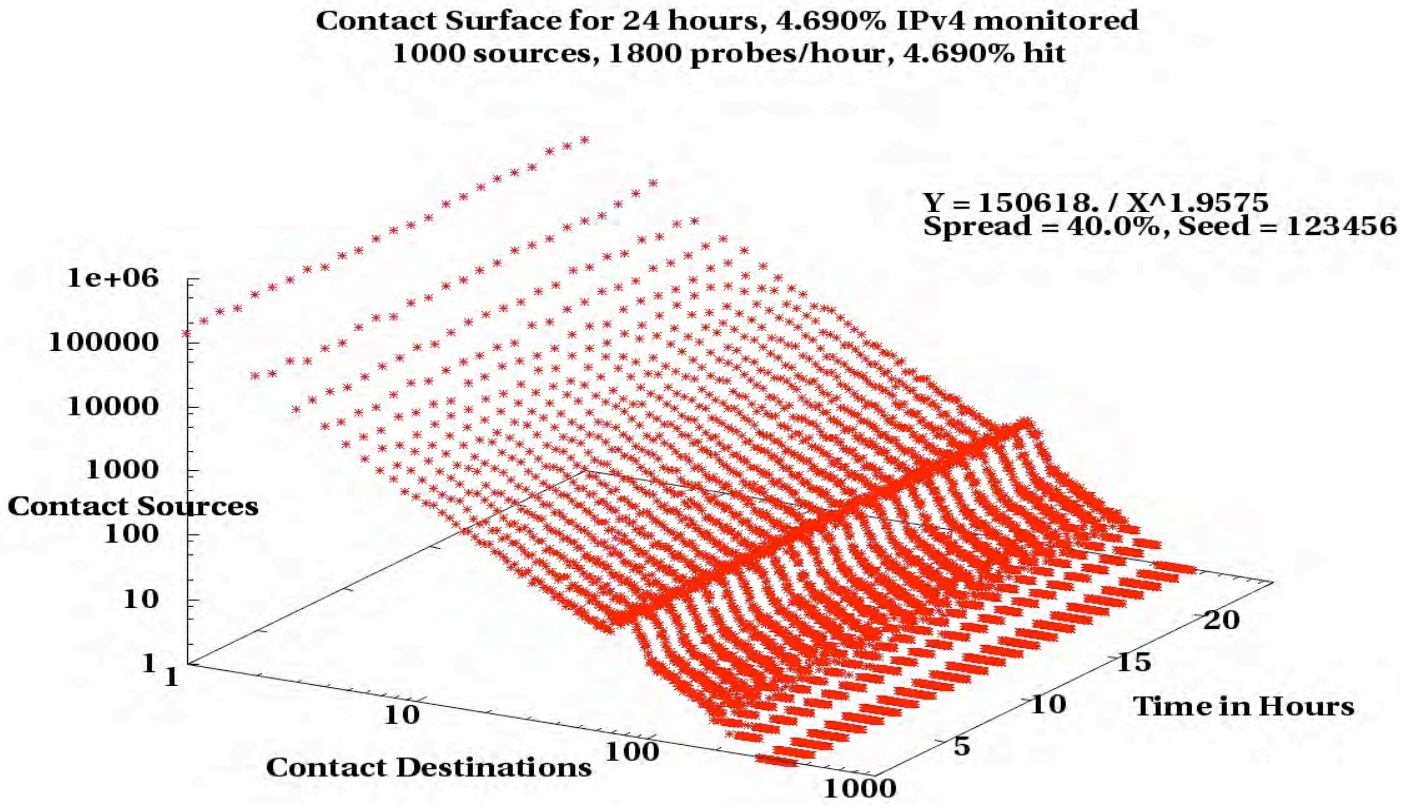


Simulating the ripple

- For each source, S_i; for each probe, j emitted during an observation period; we generate a random $R_{i,j}$ in {0..1.0}.
- If $R_{i,j}$ is < the % of IPv4 monitored, it is a hit.
- Use the hit count to select the appropriate cell in the background traffic contact line and add 1 to it.
  - source S_i hit that number of destinations during the simulated observation period period.
- Plot the modified contact line in either 2D or as part of a 3D contact surface.

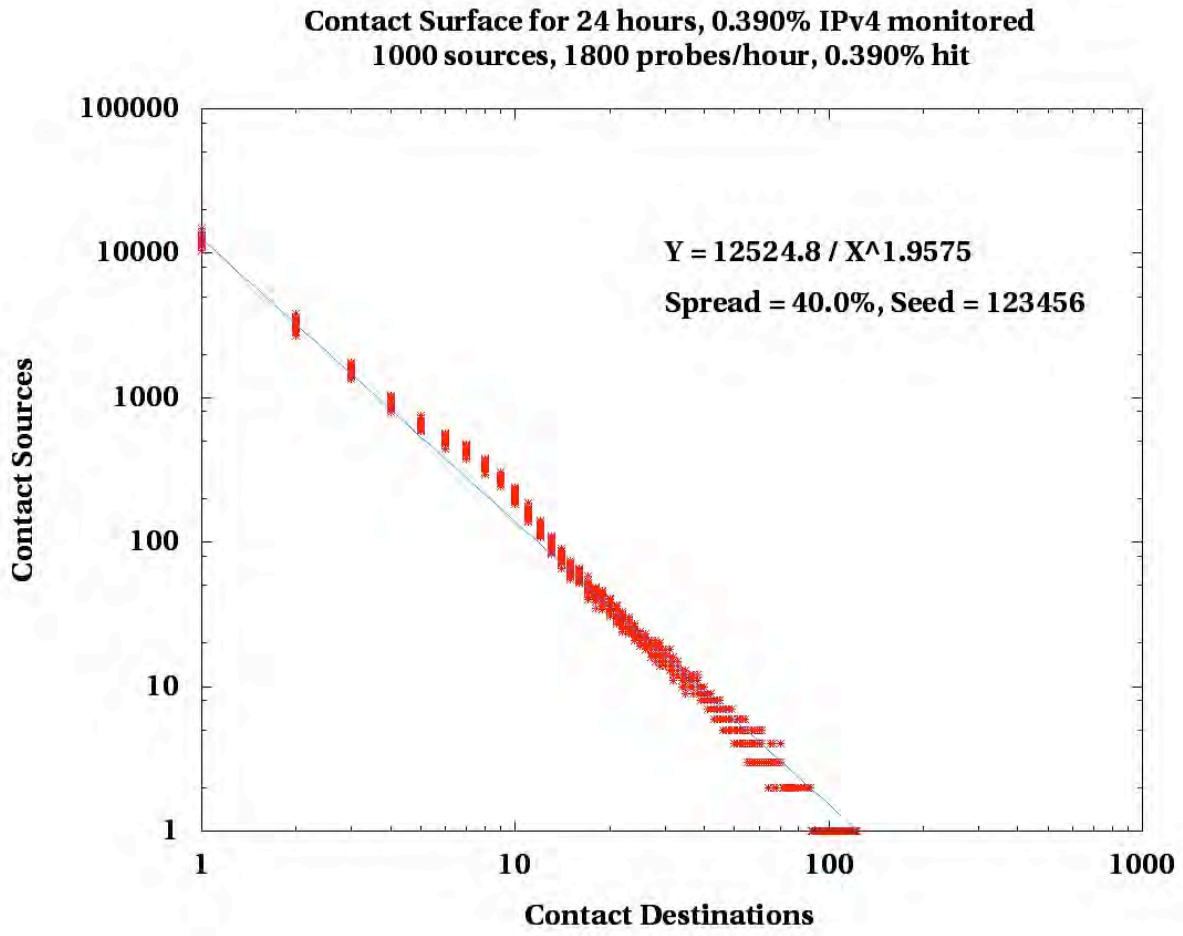
© 2005,6,7 by John McHugh, Carrie Gates

A plausible ripple



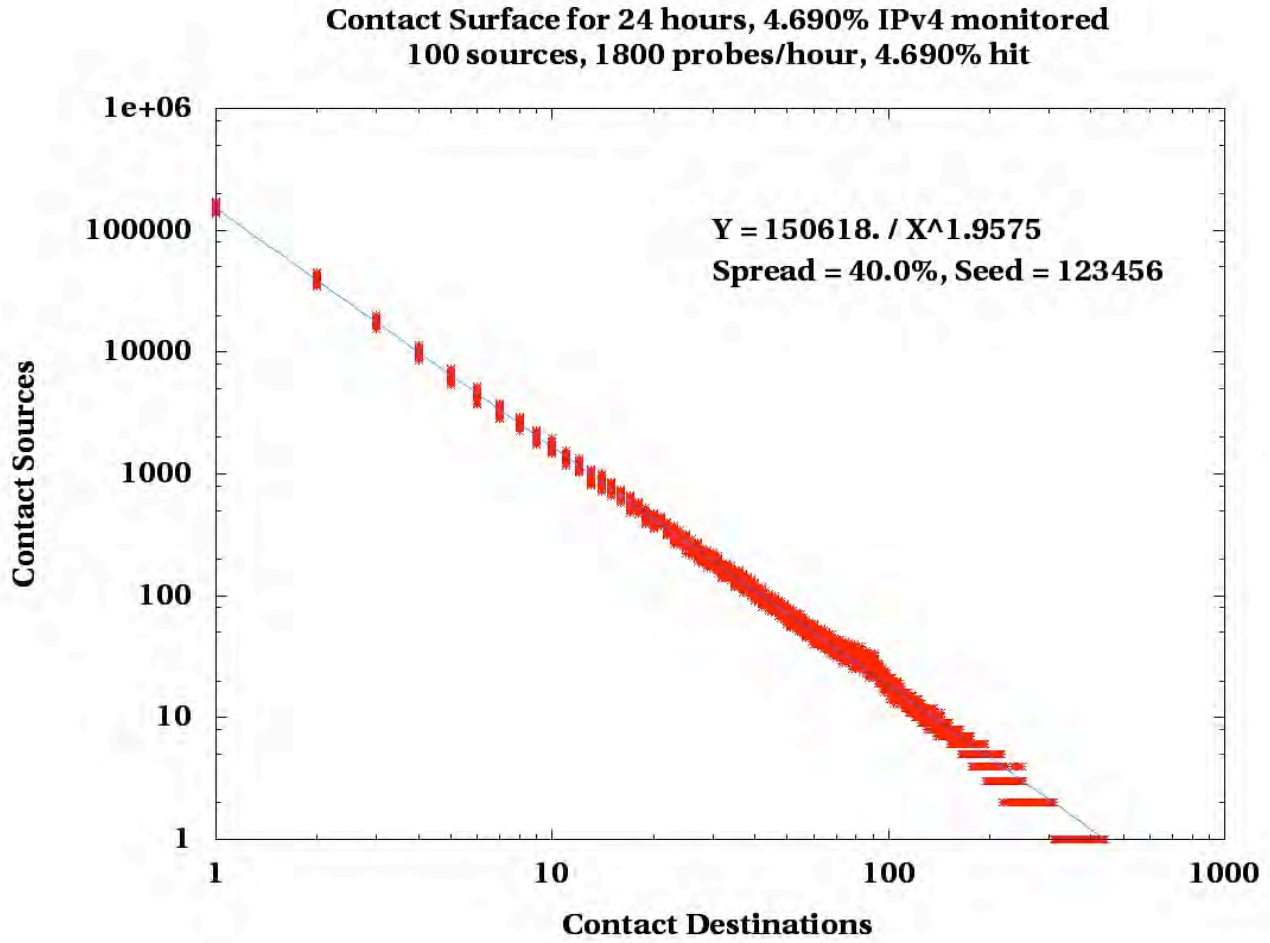
© 2005,6,7 by John McHugh, Carrie Gates

Observability: 1000 probers /16 coverage



© 2005,6,7 by John McHugh, Carrie Gates

Observability: 100 probers 12 X /8 cover



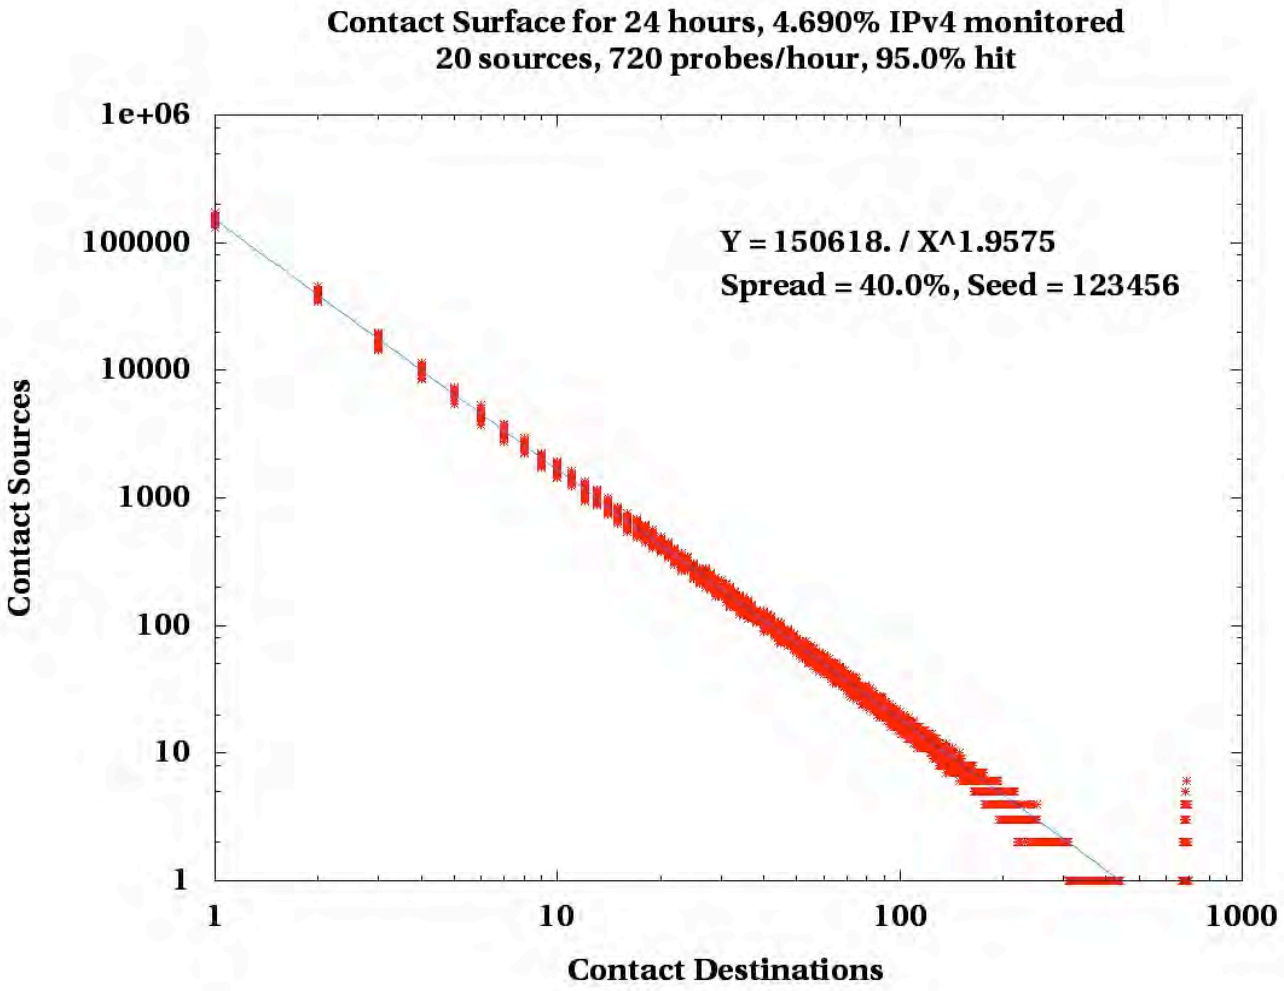
© 2005,6,7 by John McHugh, Carrie Gates

Simulated and real spikes.

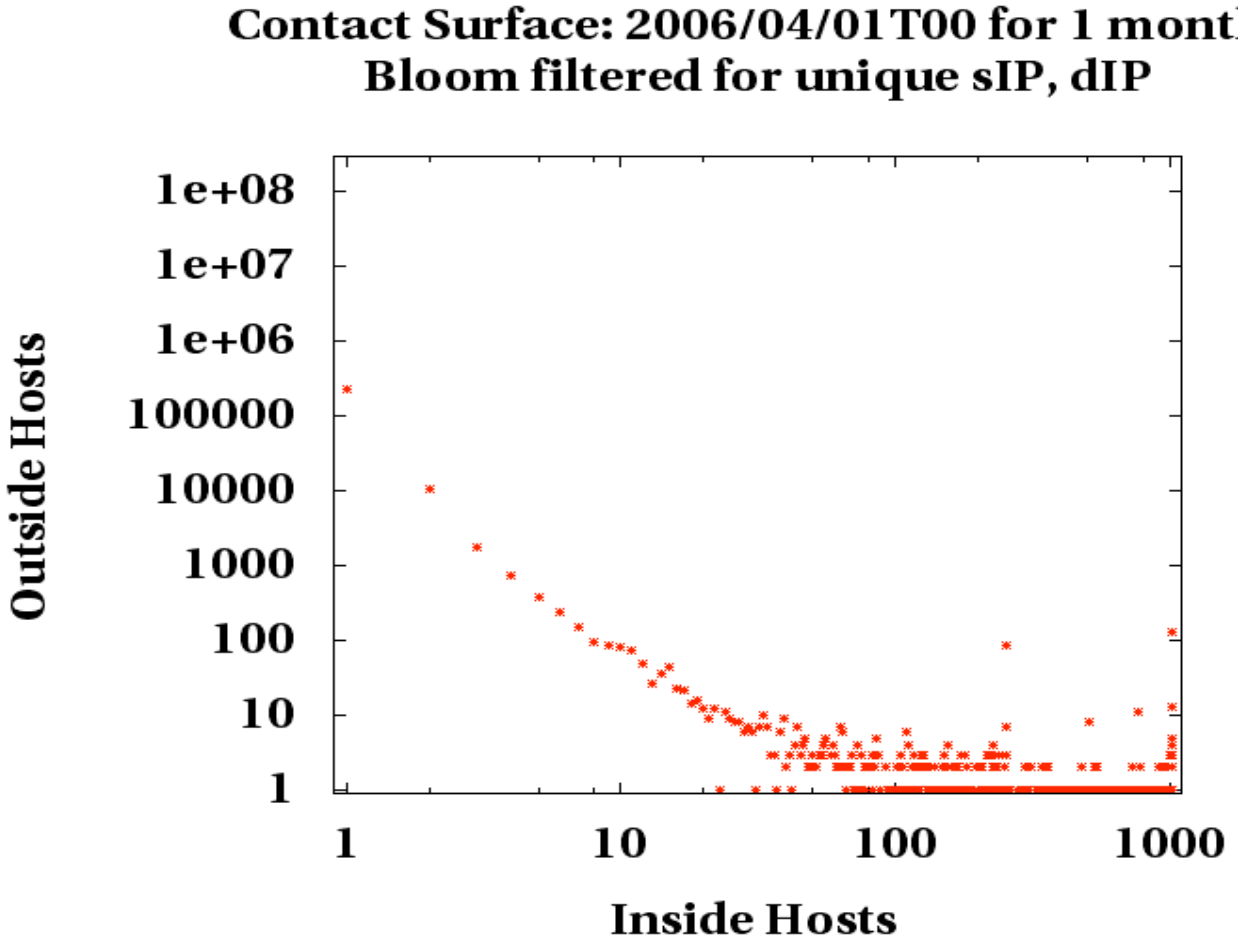
- The spikes appear when the percentage of intercepted probes is high.
  - Occurs when the probes fall mostly, 95%+, in the monitored address space.
  - At 100%, the spike becomes a point
- First, we simulate the spike.
- Next is a one month contact line for our /22, based on Bloom filtering for unique sIP, dIP pairs.
  - Note points at 254, 508, 762 and 1016 addresses.

© 2005,6,7 by John McHugh, Carrie Gates

The spike in the Welchia.B displays



Contact line for April 2006 for a /22



© 2005,6,7 by John McHugh, Carrie Gates

Future work

- We would like to visit or revisit the data for current and past perturbations.
- Develop analytical techniques for identifying cohorts of players exhibiting abritrary, but similar characteristics.
- Explore other regions of the contact surface
- Link visualization to source / cohort identification in the visualization tool we are developing for DHS.

© 2005,6,7 by John McHugh, Carrie Gates

Visual Representations of Flow Data

and the Value of Visual Language

Presented by Sunny Fugate
Space and Naval Warfare Systems Center, San Diego



Copyright 2007 Sunny Fugate

Human-Machine Efficiency

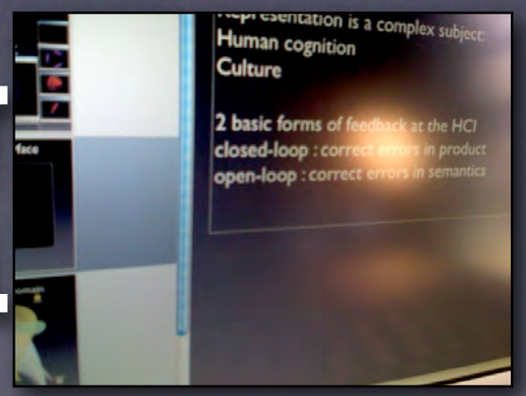
Over-Learned: **Feedback**



haptic

closed-loop : correct errors in **production**

open-loop : correct errors in **semantics**

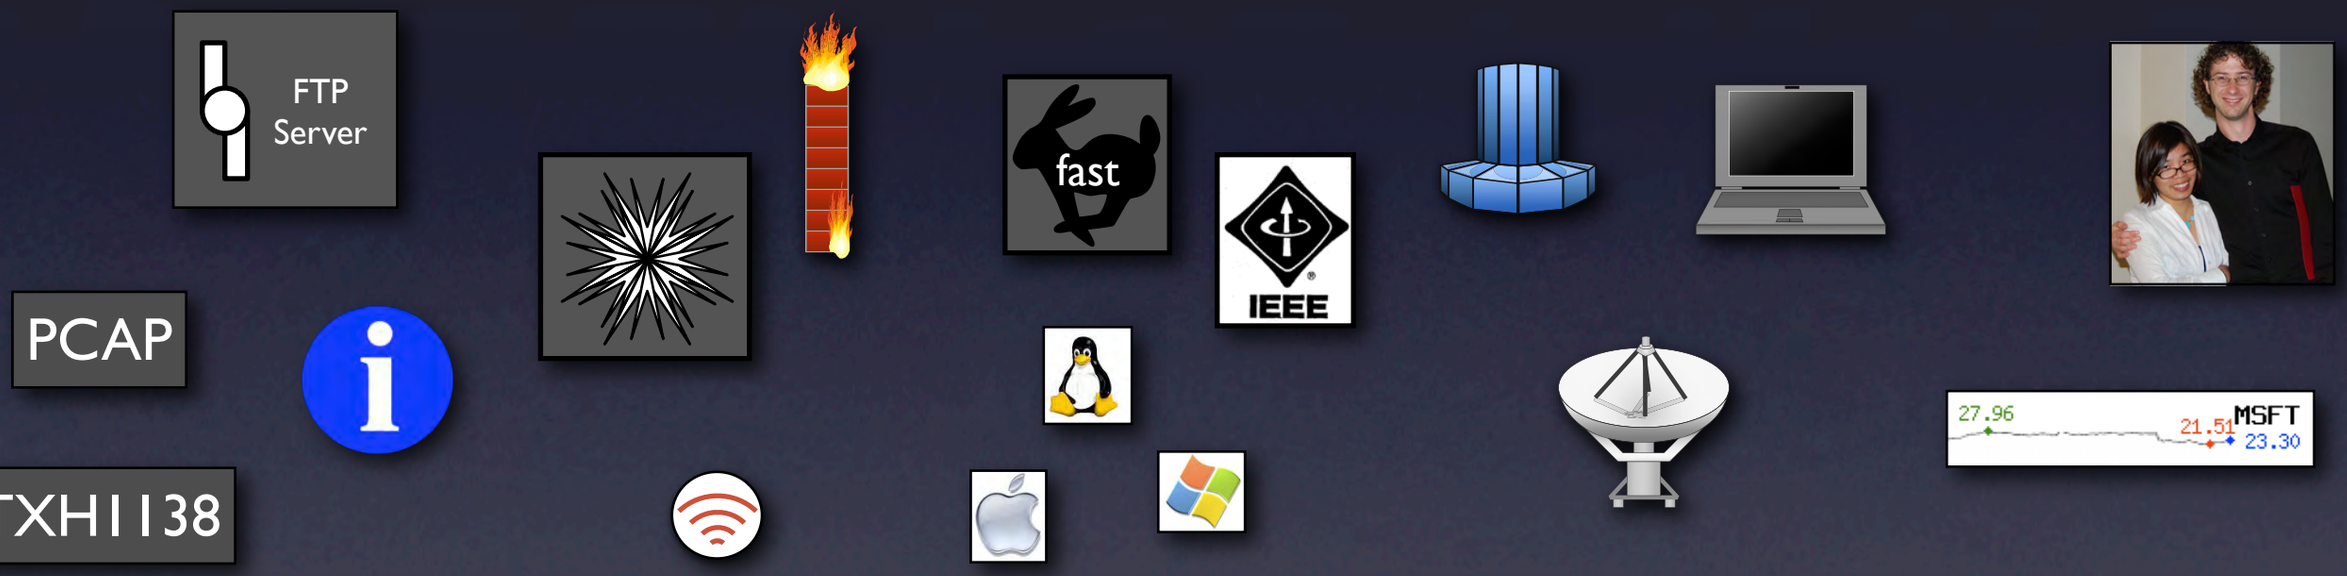


visual / aural

© 2007 Sunny Fugate

Human-Machine Efficiency

Under-Learned: **Representation**



arbitrary | metaphor | association | representational | indexical

Culture/Domain Specificity

© 2007 Sunny Fugate

Human-Machine Efficiency

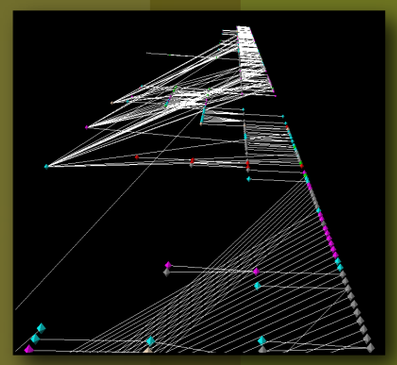
Over-Learned: **Feedback** - haptic vs visual/aural

Haptic Feedback

Visual / Aural Feedback

Linear access

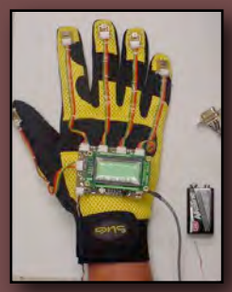
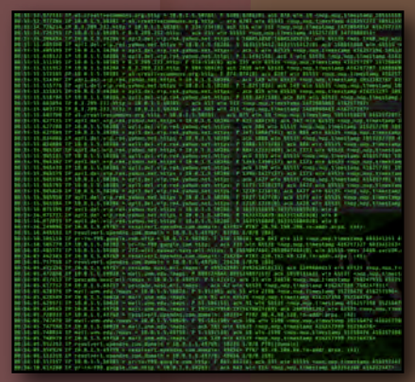
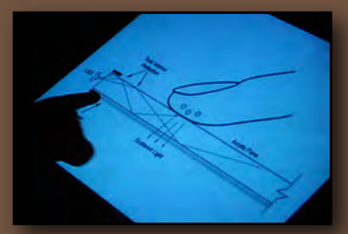




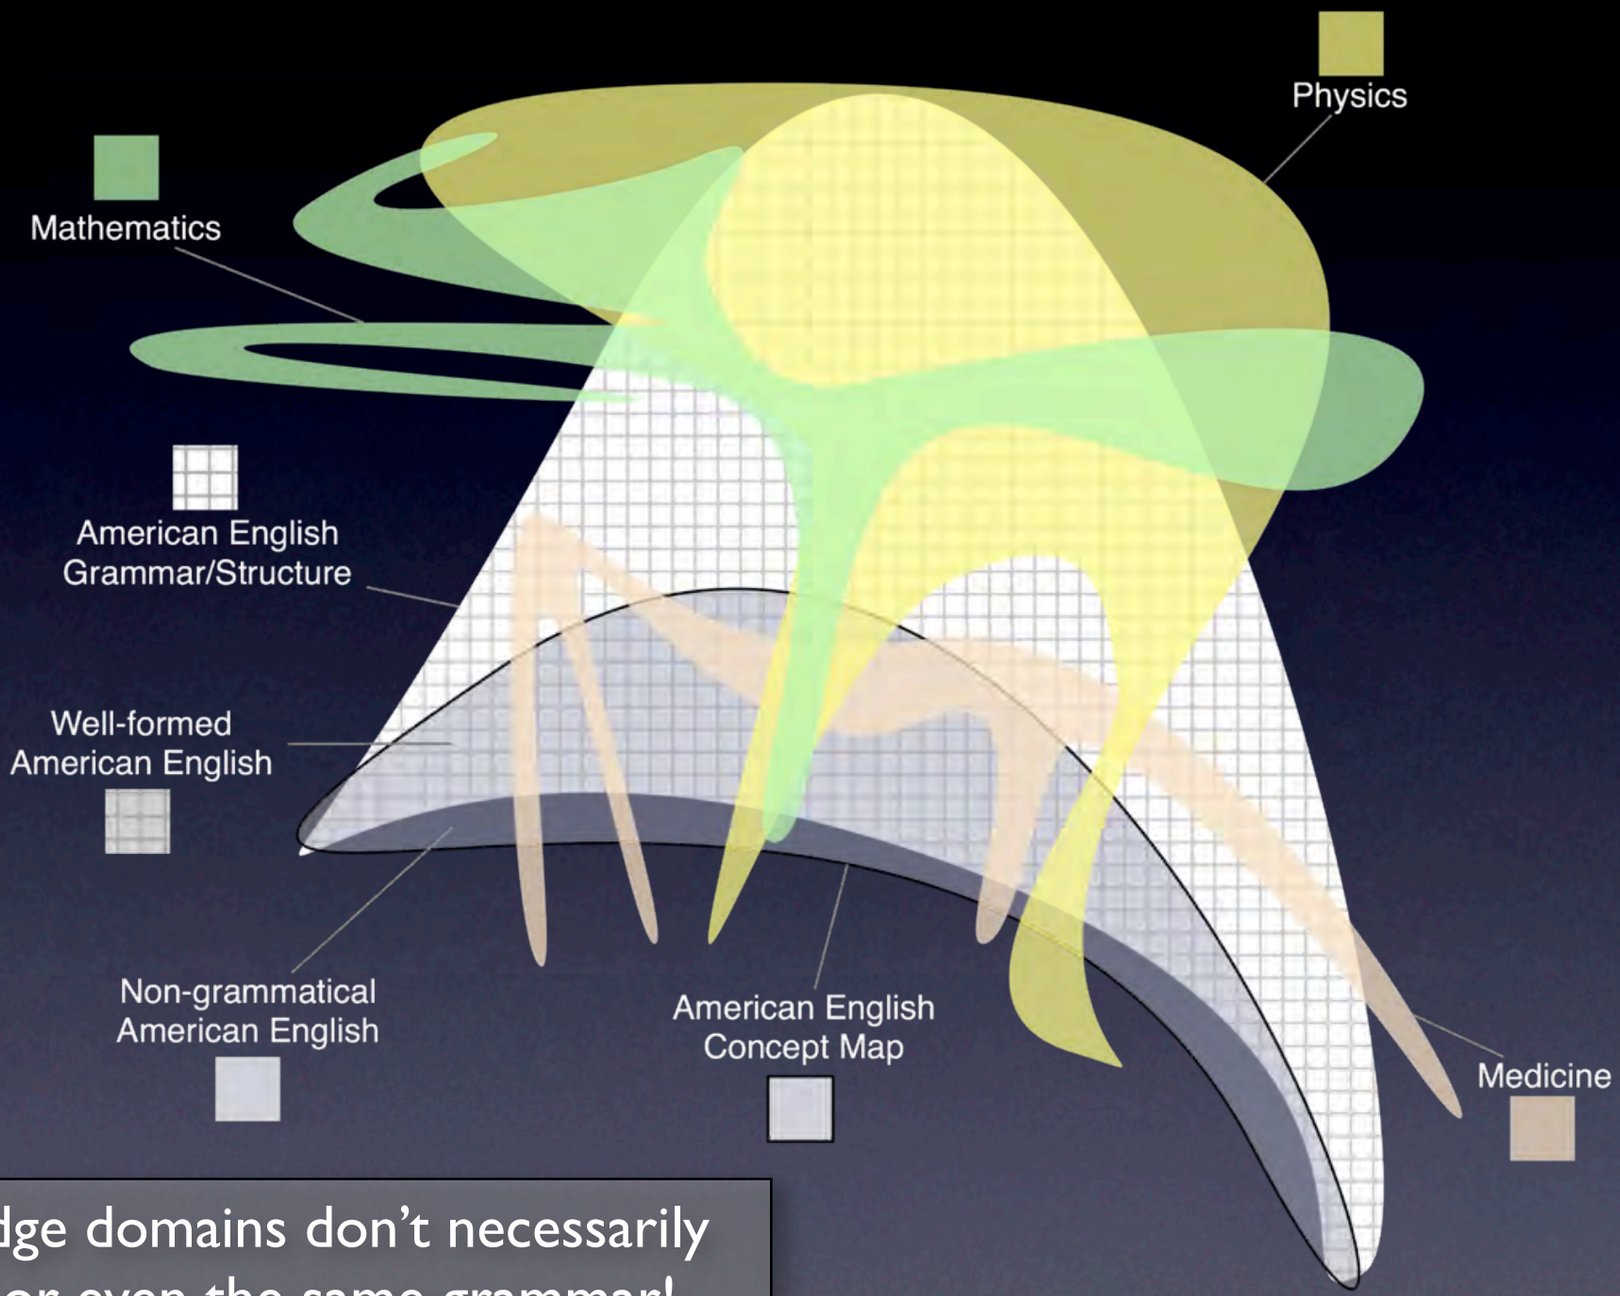
?

Random access


© 2007 Sunny Fugate

Language Domains



Cultures and knowledge domains don't necessarily use the same lexicon or even the same grammar!

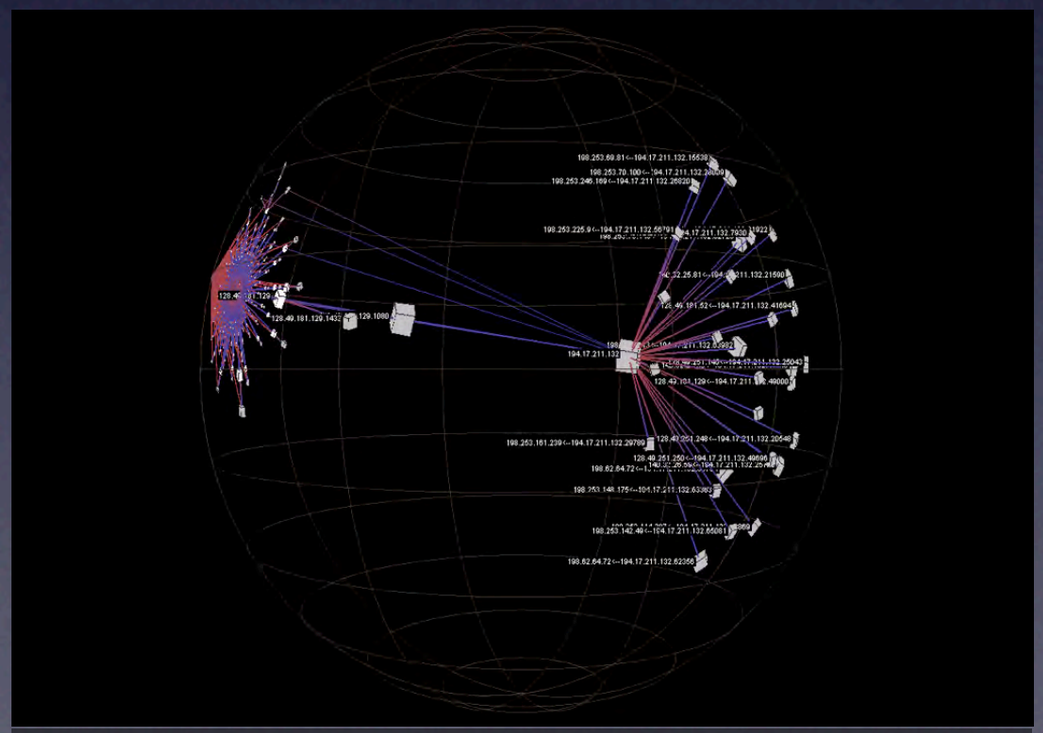
How does the CND lexicon map to common language? Technical language? Military/tactical language?

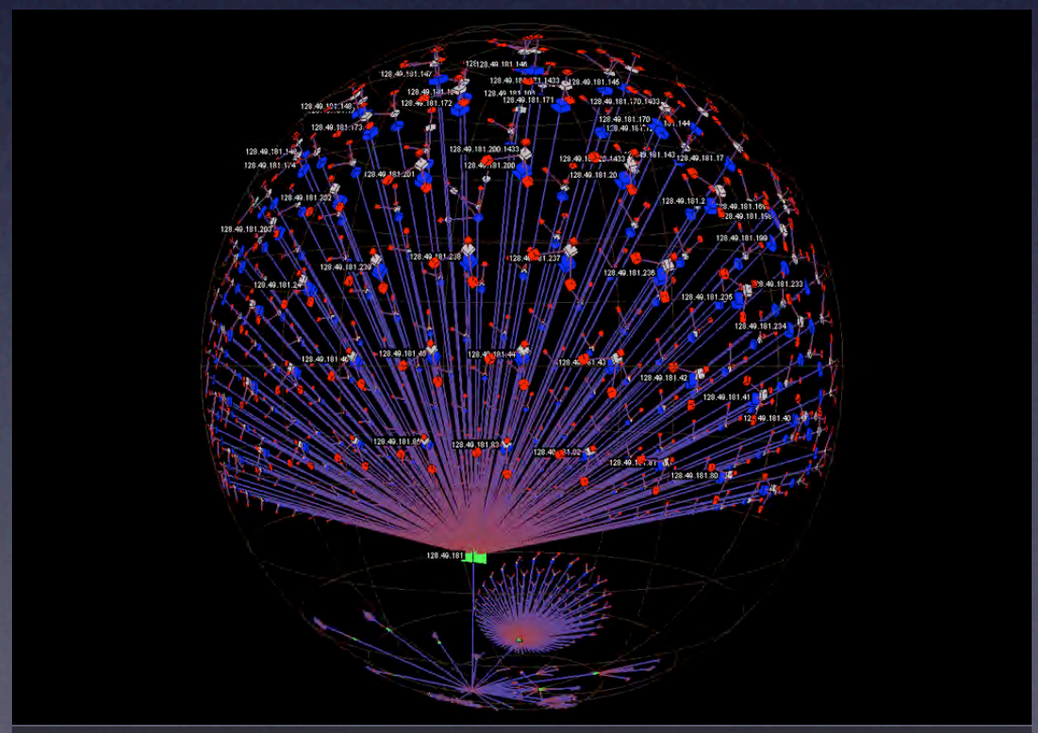
© 2007 Sunny Fugate

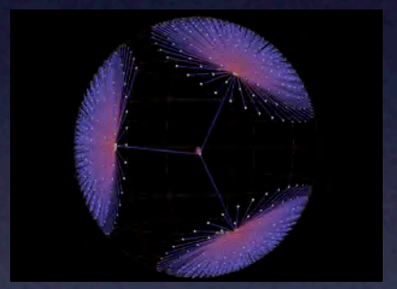
Flow in hyperbolic space

- 3 month SSC project in 2002
- discover and apply network visualization tools
- **Hyperviewer**: quasi-hierarchical hyperbolic space
- '**fish-eye**' 3-d
- Created by Stanford researcher **Tamara Munzner**

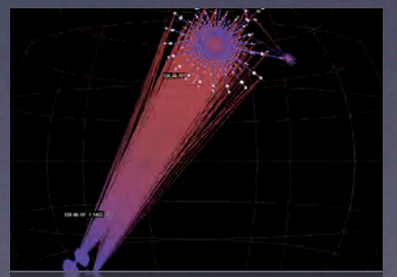










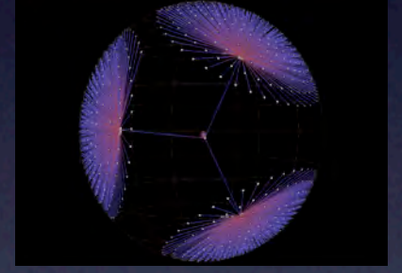


© 2007 Sunny Fugate

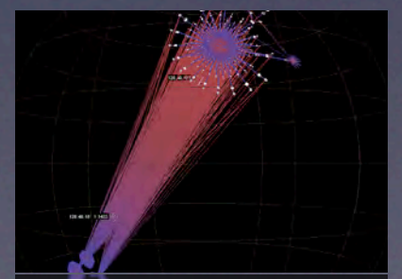
Flow in hyperbolic space

- Easily adapted to a forced-hierarchy view of flow
- **Opensource** C++ library and UI
- Experimented with visual methods

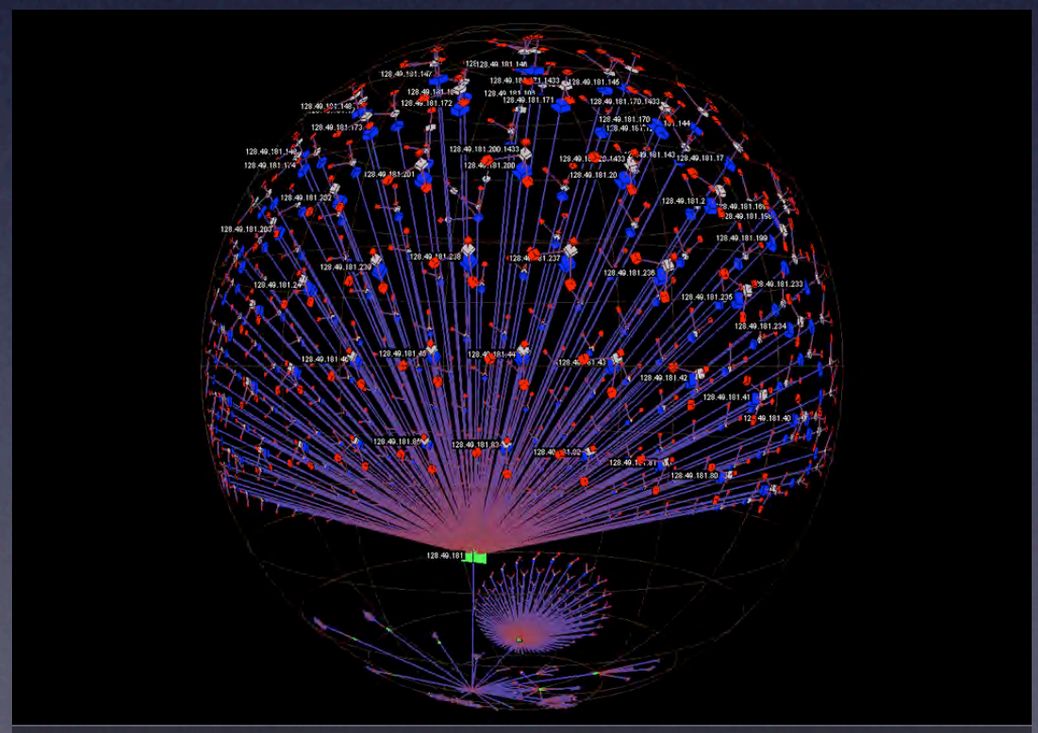




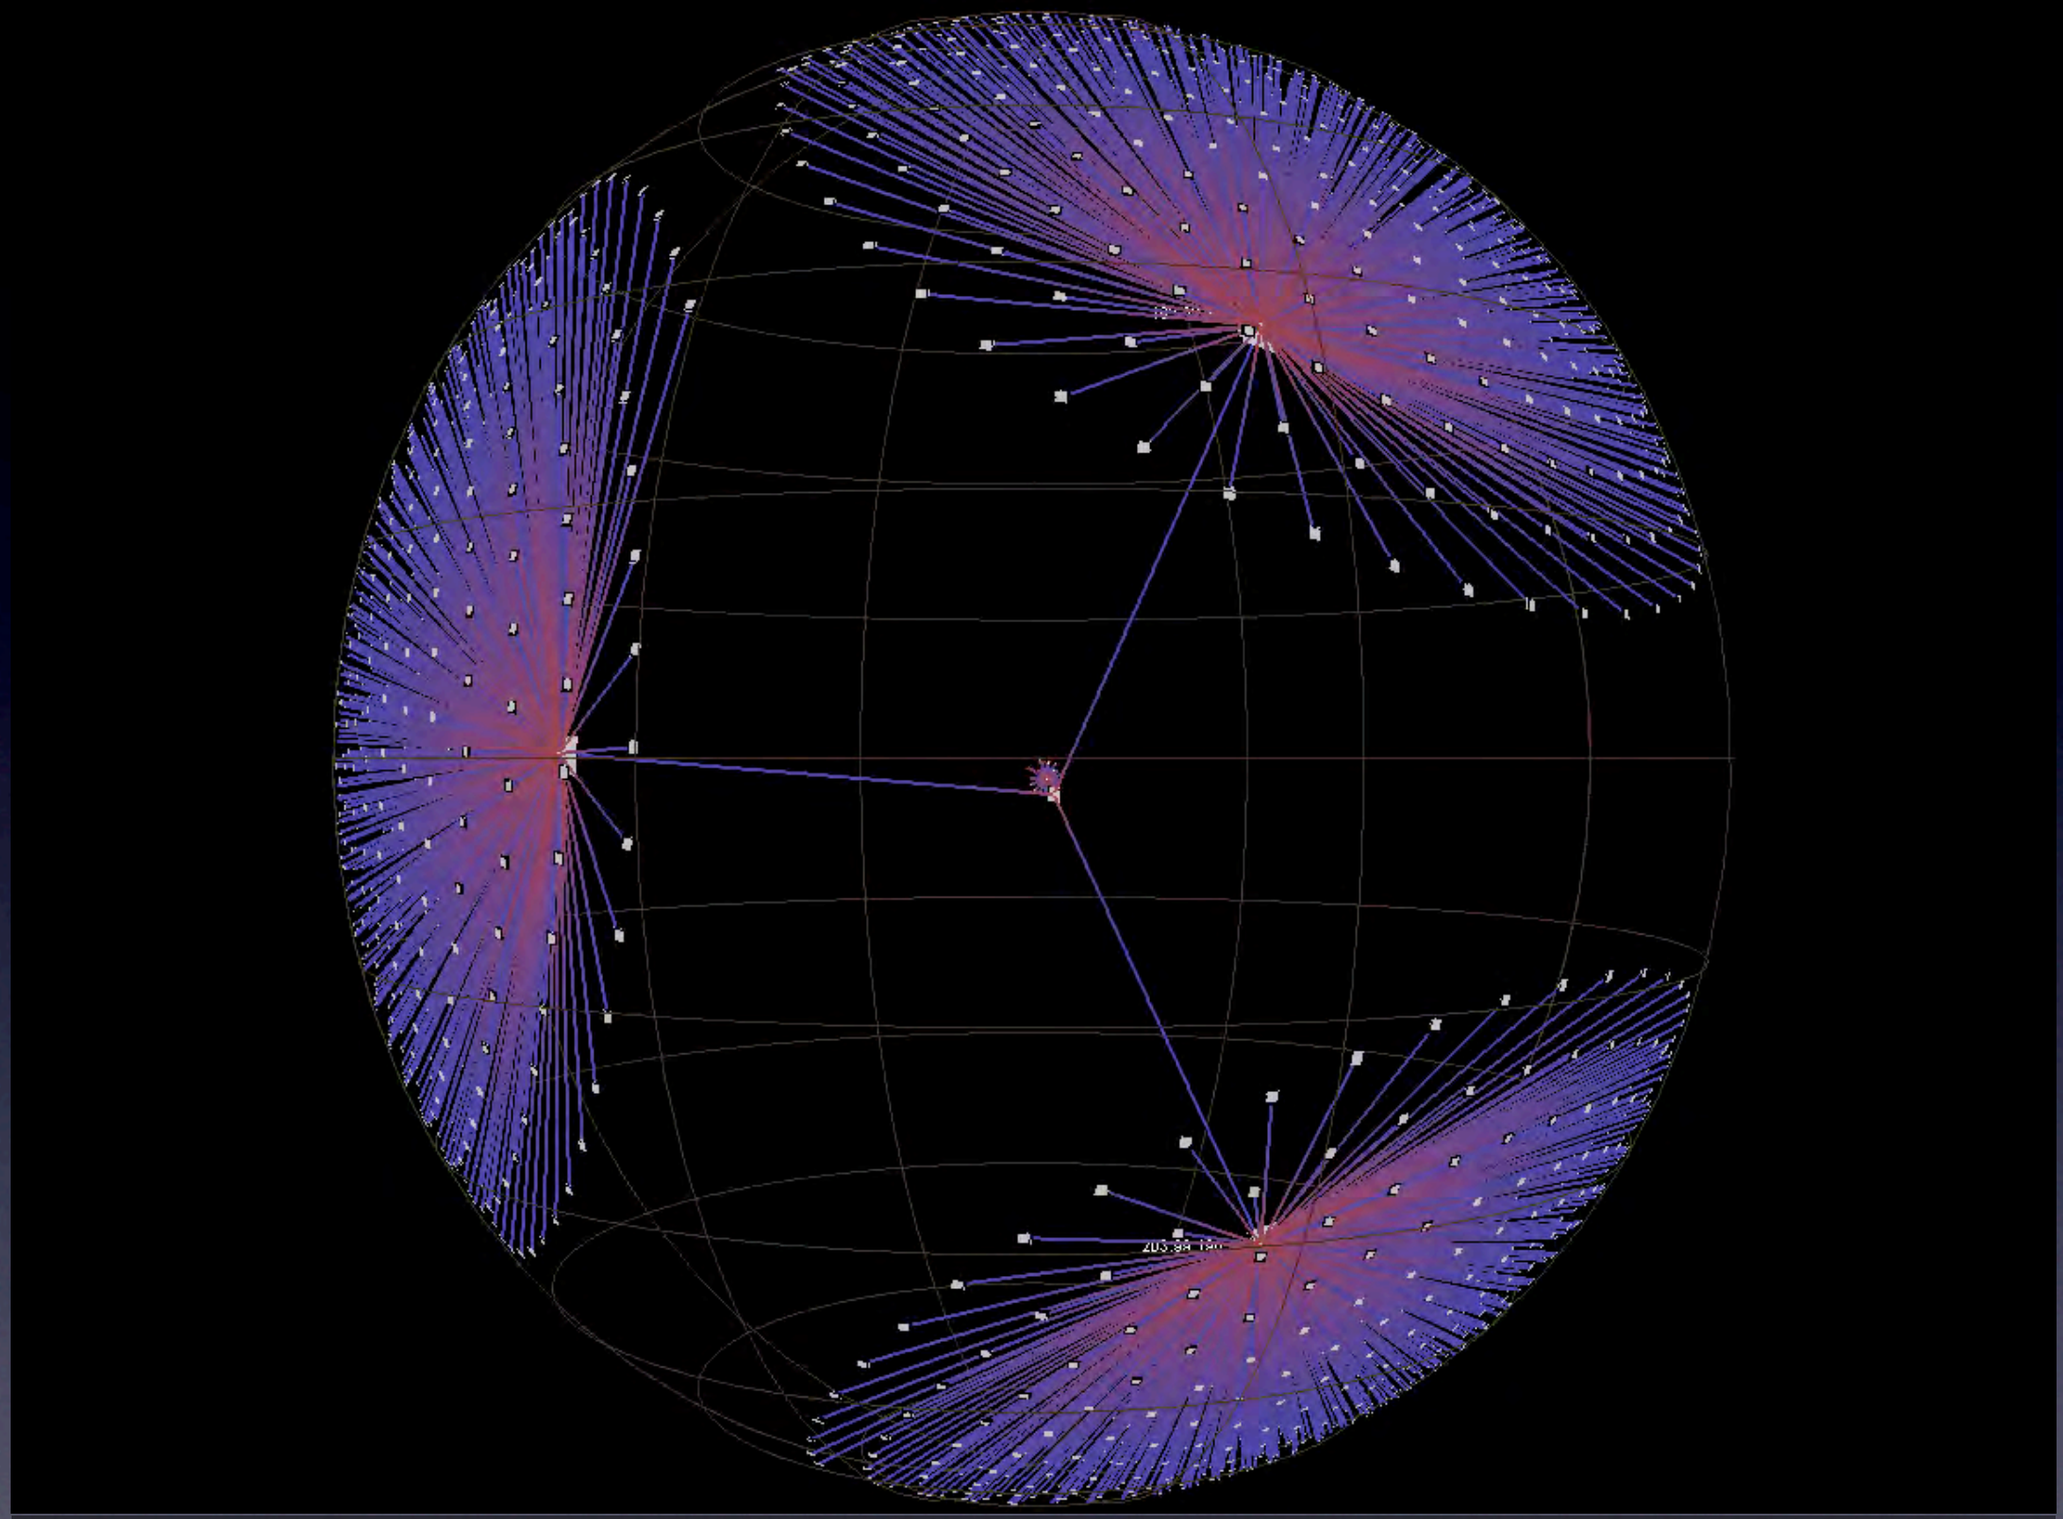




- colors
- graph cycles
- scaling
- text labels
- **graph size**
- search automation

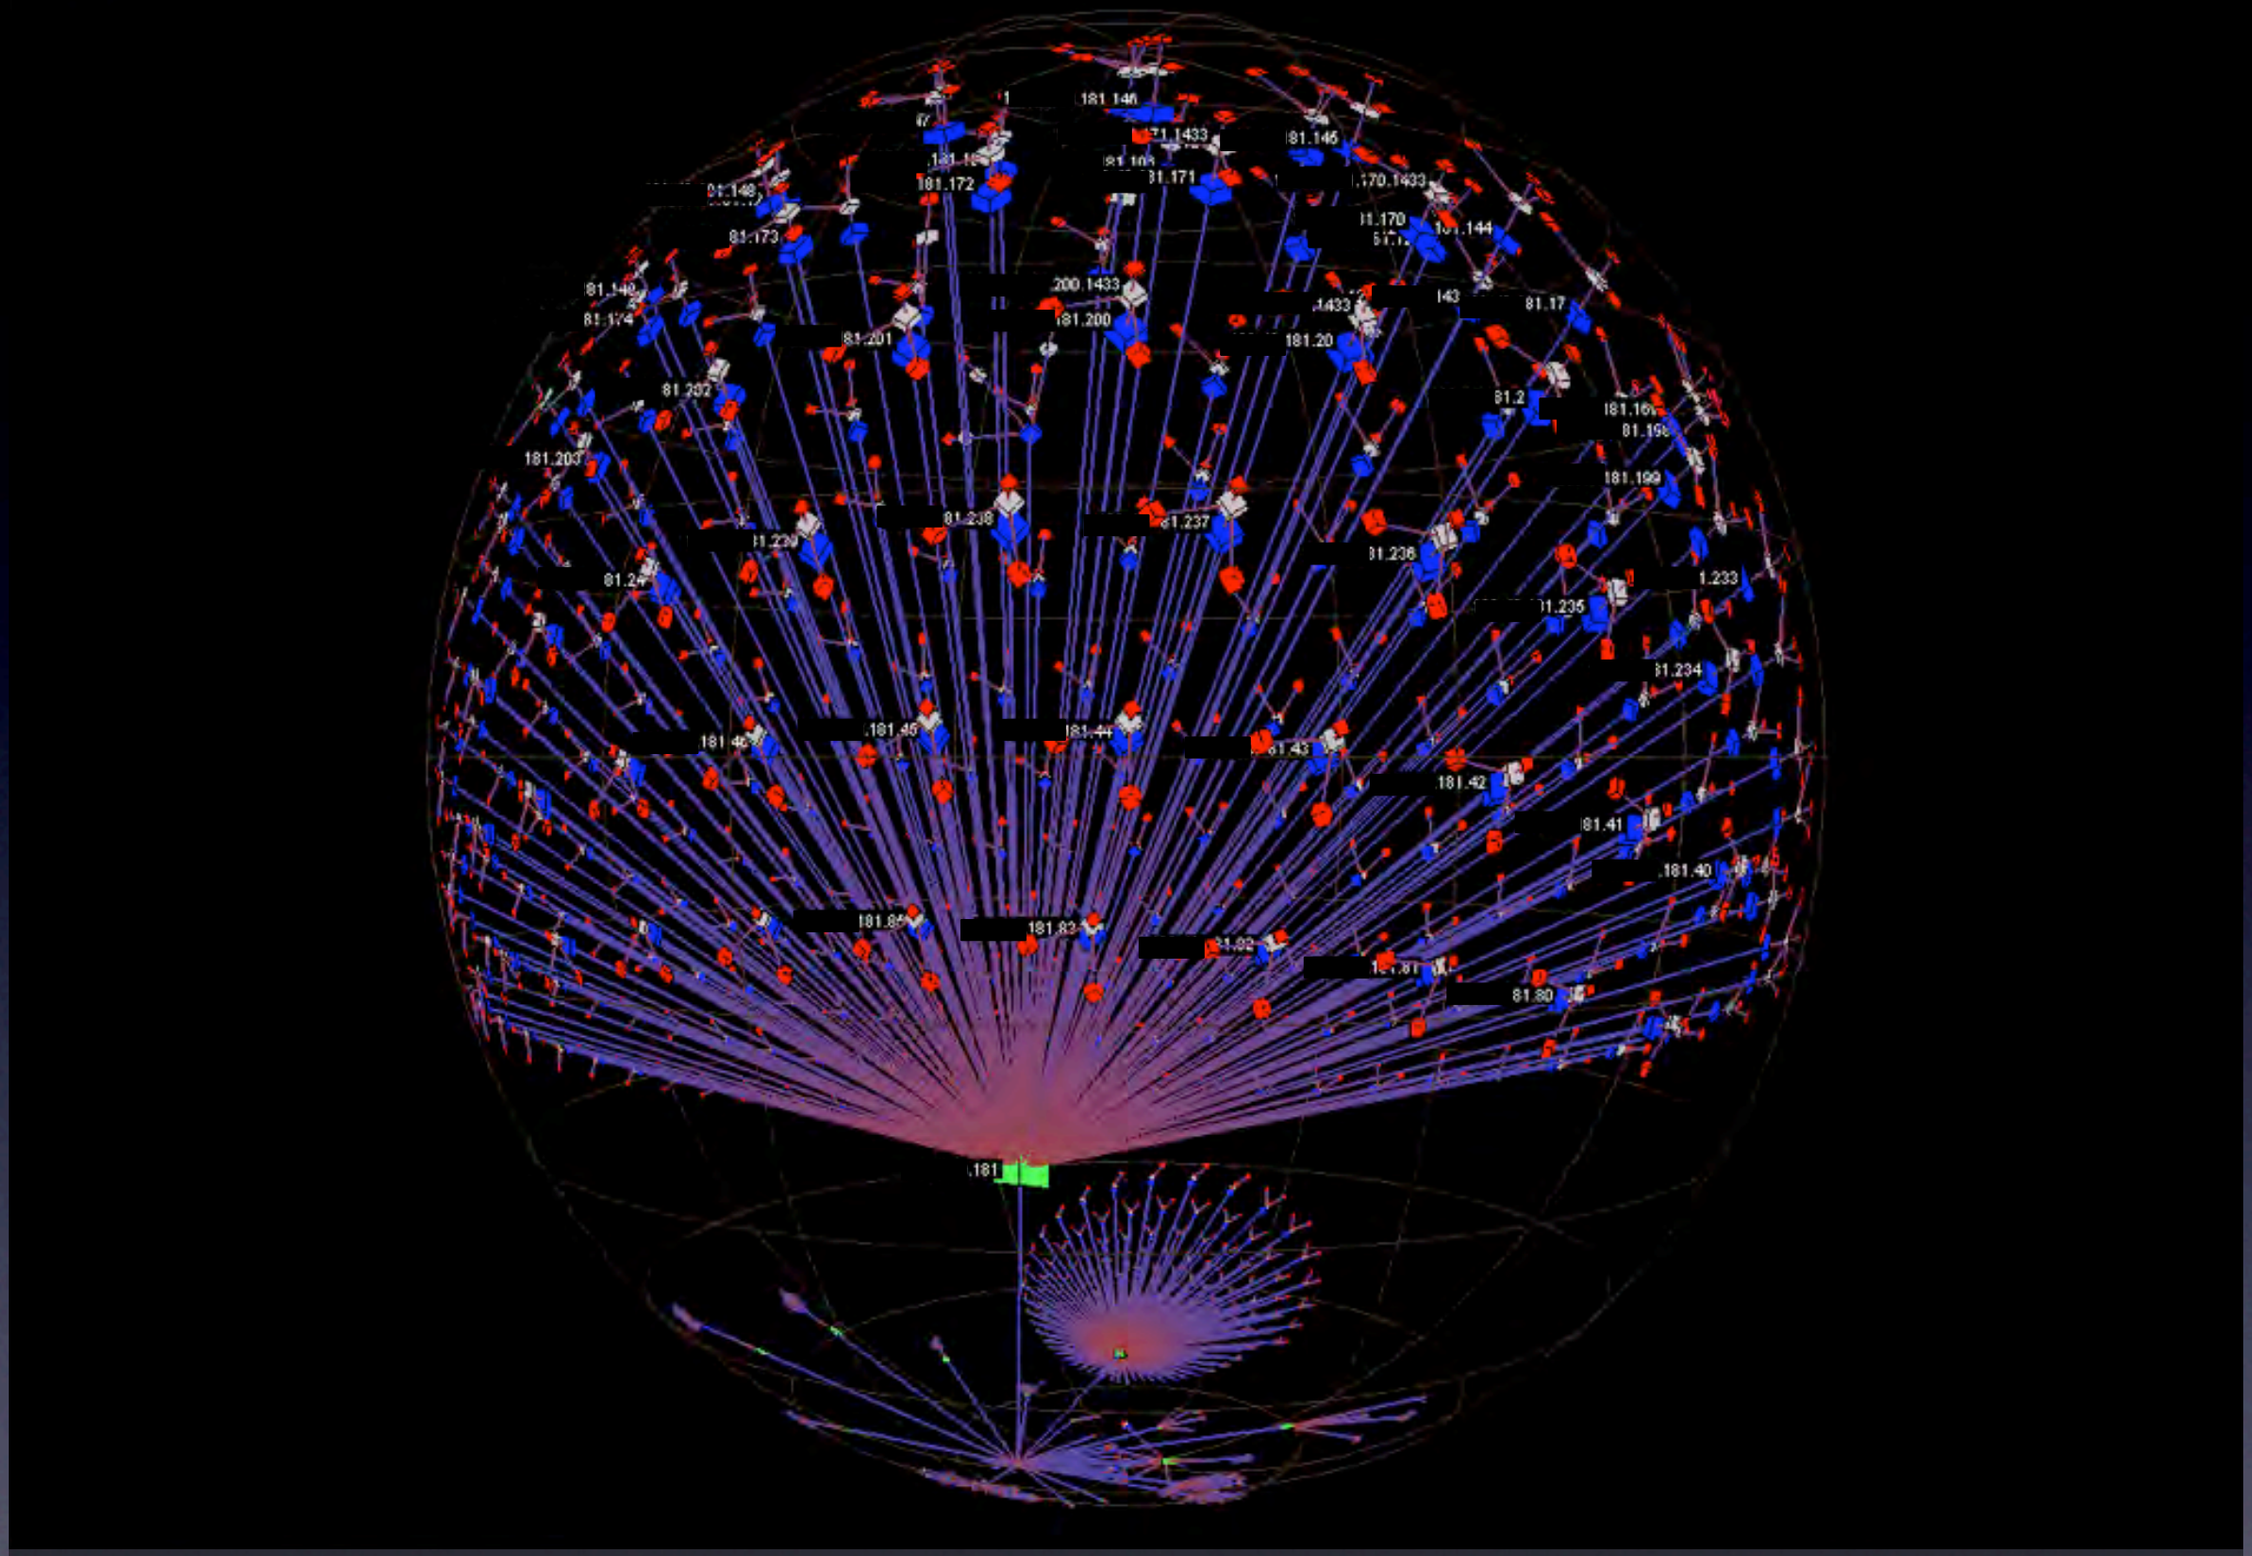


© 2007 Sunny Fugate



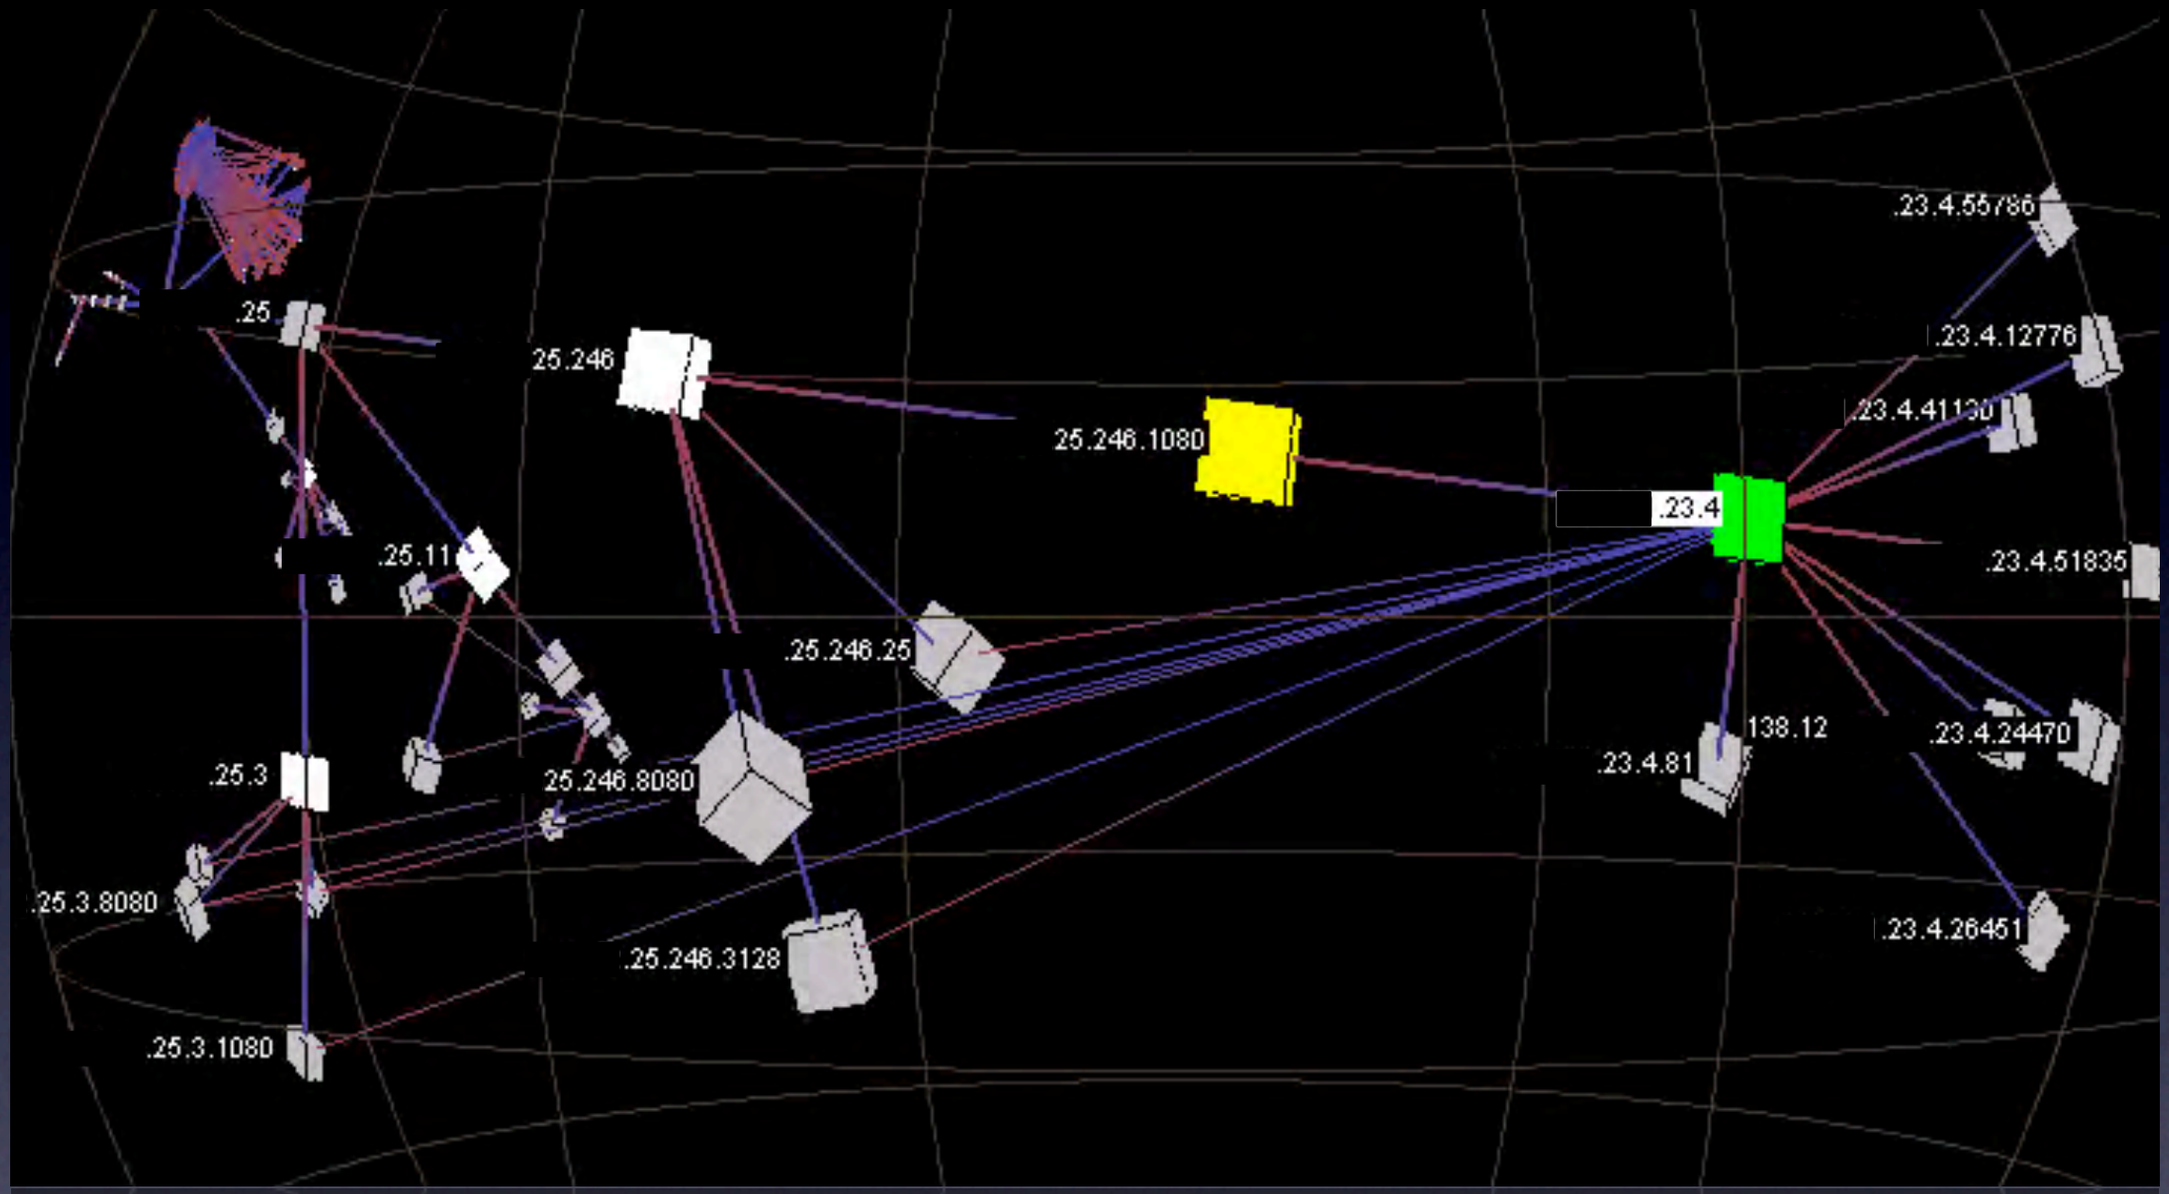
Symmetry in port access from 3 separate clients.

© 2007 Sunny Fugate



src/dst ports colored red/blue

© 2007 Sunny Fugate



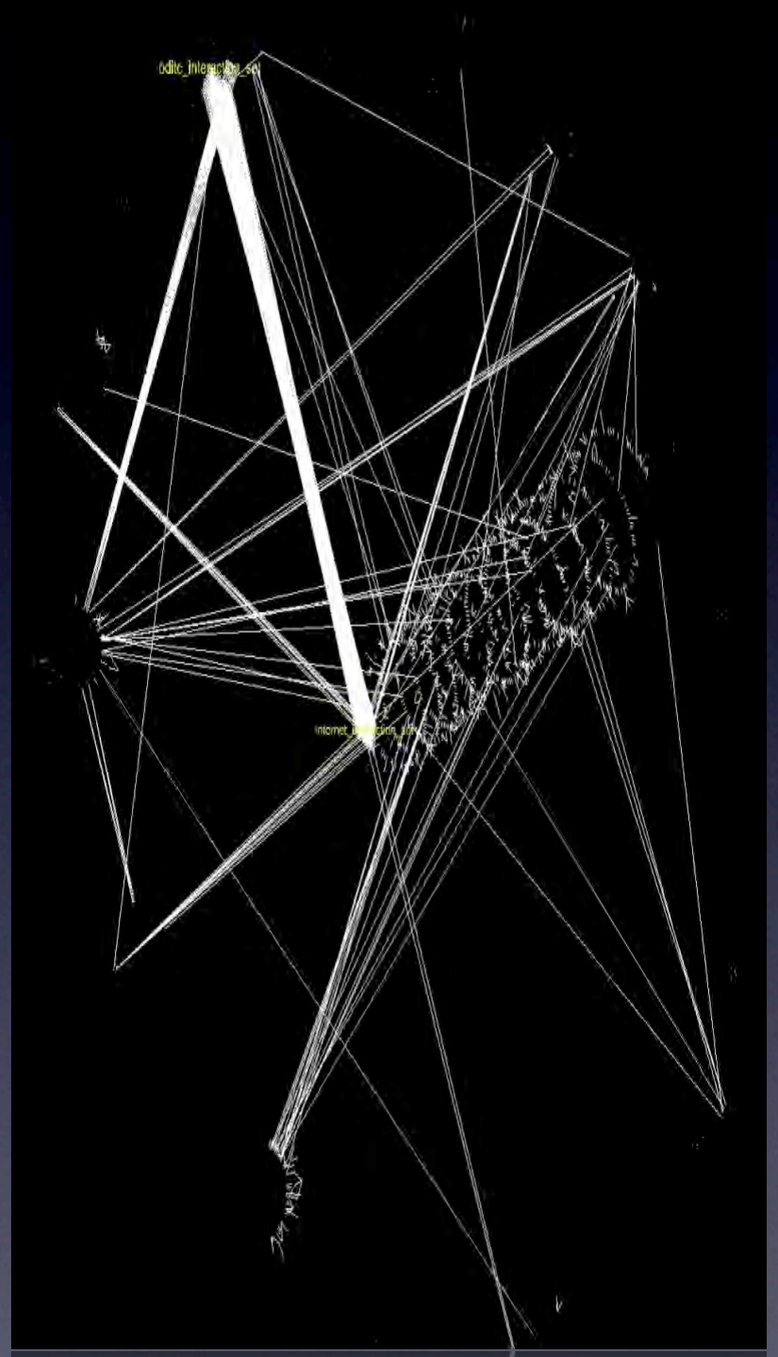
Hierarchy showing client subnet and server ports

© 2007 Sunny Fugate

Shapes Vector

- Acquired by DARPA in 2002
- Developed by Australian DSTO

  (Defence Science Technology Organisation)
- JTF-GNO pilot program from 2003-2006



What is it?

- Intelligent Agents gather information and produce inferences
- Gathers information from multiple sources
  - pcap, flow, Snort, syslog, etc
- IAs performs automated data correlation & knowledge extraction
- Integrates visual and command-line analysis
- Integrated visualization makes use of human vision
- Supports visual analysis and decision-making

© 2007 Sunny Fugate

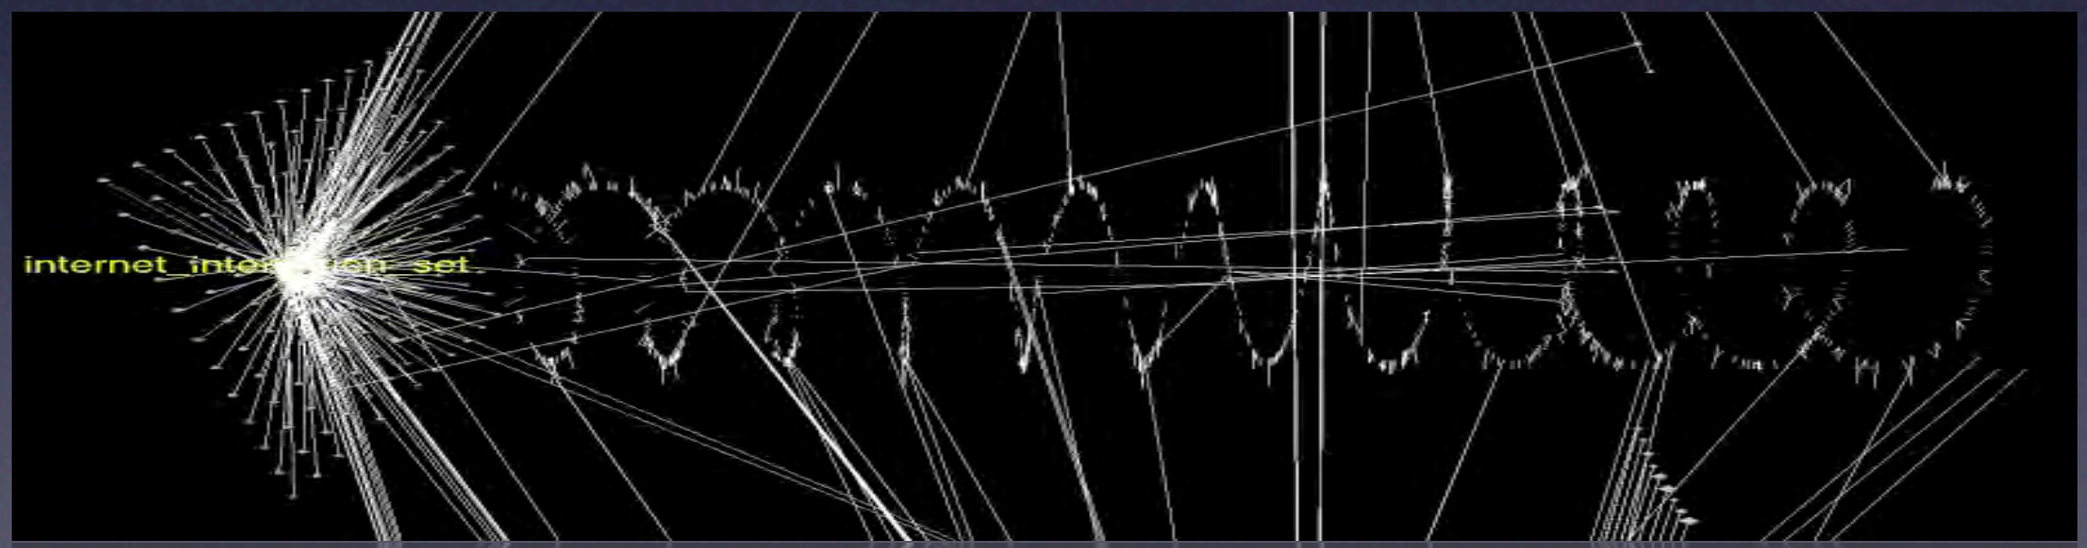
Shapes Vector

Contextual spatial, temporal, social, topological

Spatial physical geography or metaphor

Temporal sequences in time, correlated

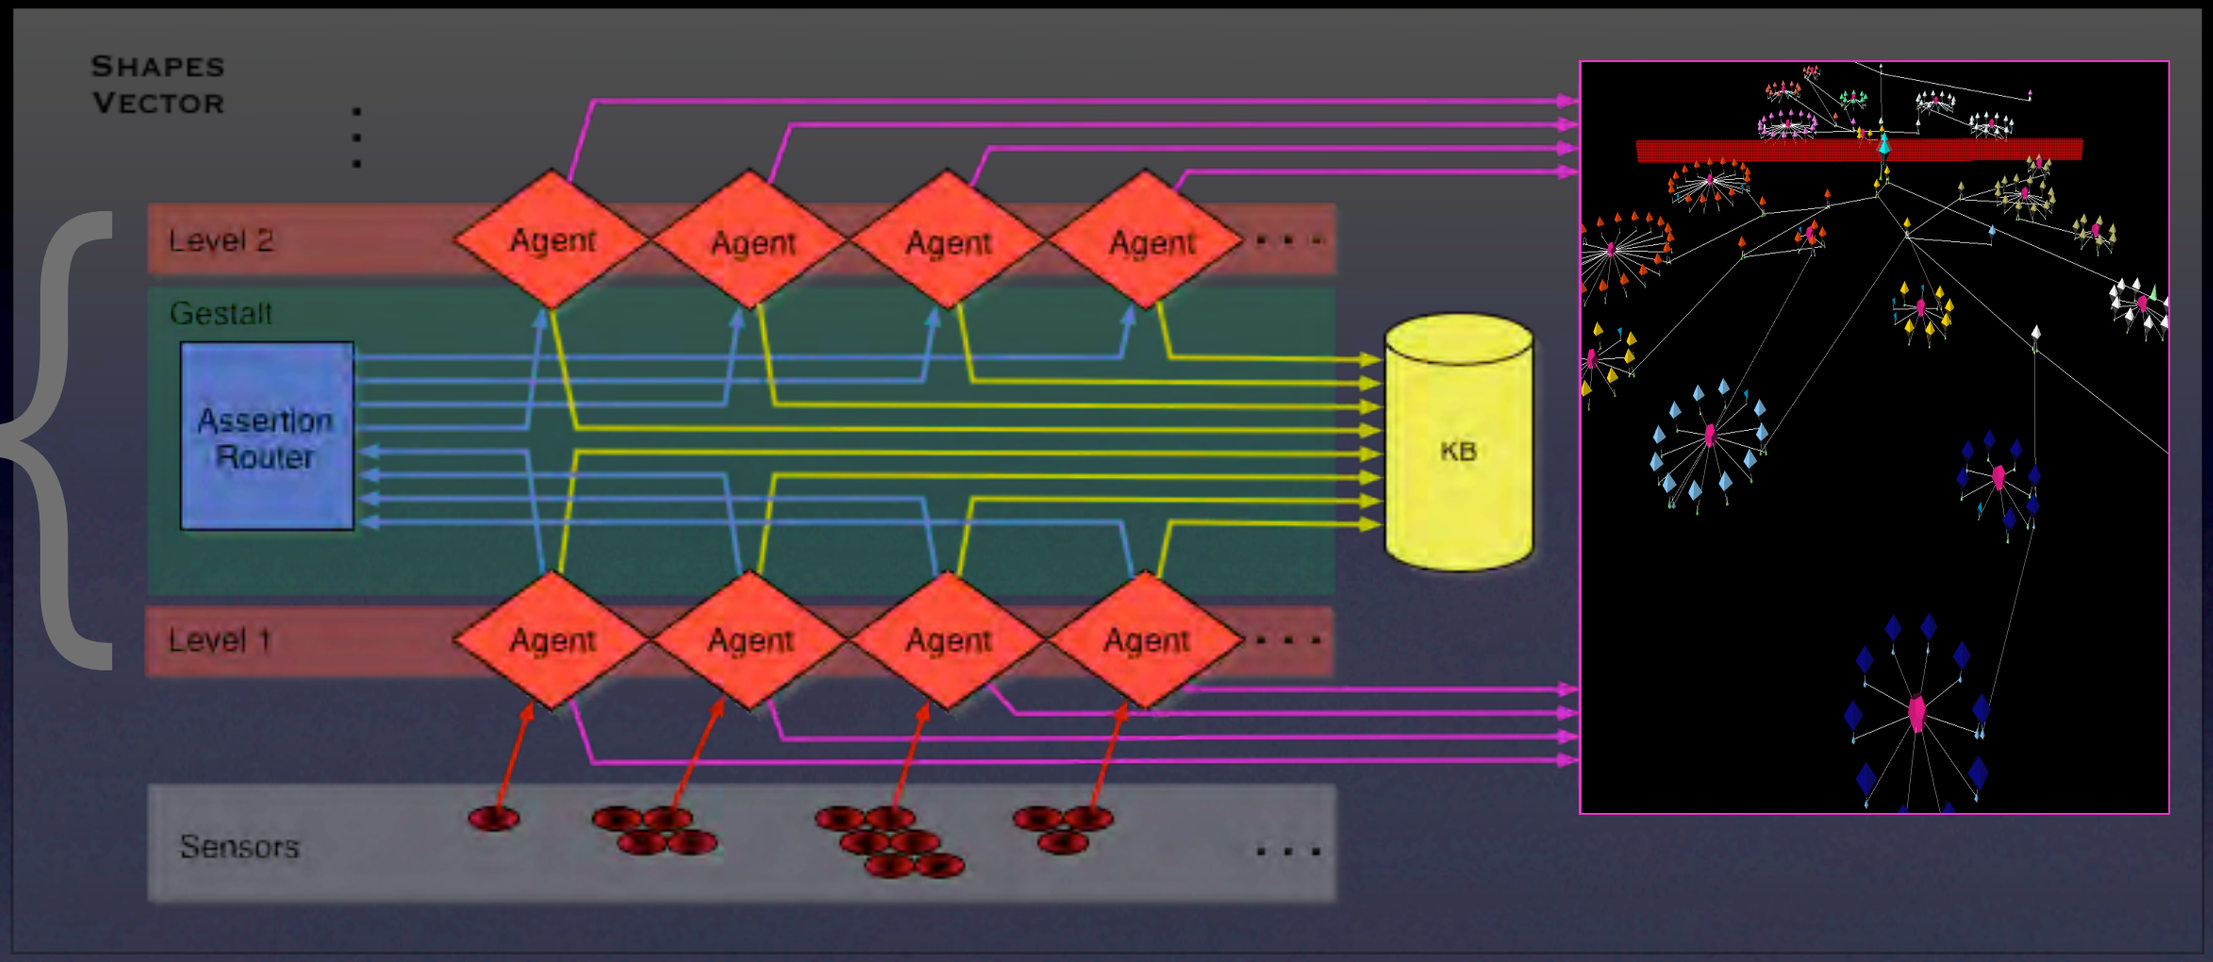
Visual use visual language to depict objects & events



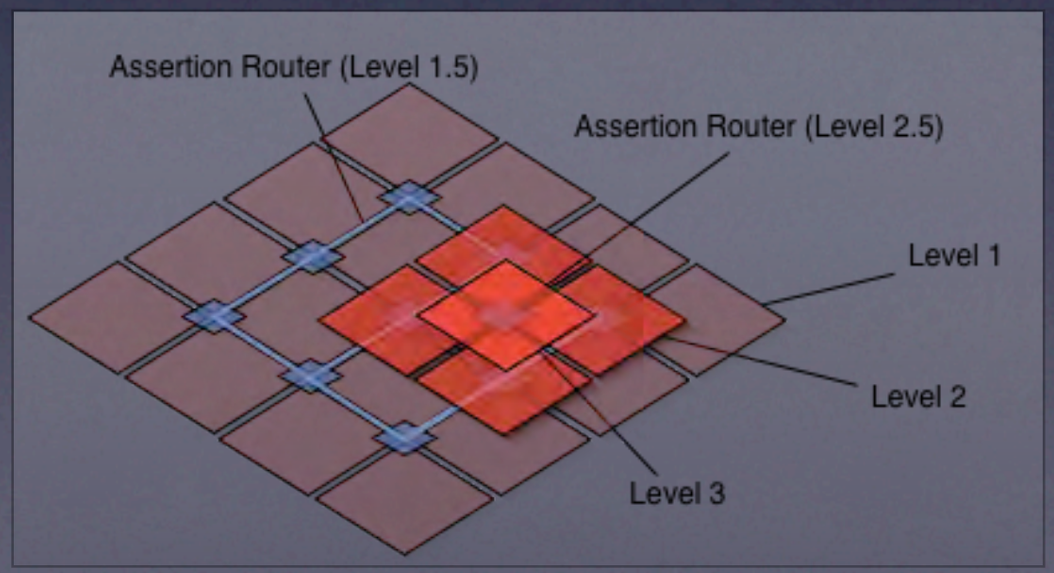
© 2007 Sunny Fugate

Architecture

SVKA



SVKA



- Agents can be written in any language - must conform to the SV ontology and knowledge architecture (SVKA) specification
- Sensors can be built to wrap nearly any information source - must produce SV ontology
- SV ontology is a knowledge description language for network defense

© 2007 Sunny Fugate

Shapes Vector - Visual Language

- Easily defined visual mappings
- No applied theory of visual language

shape/color/scale 

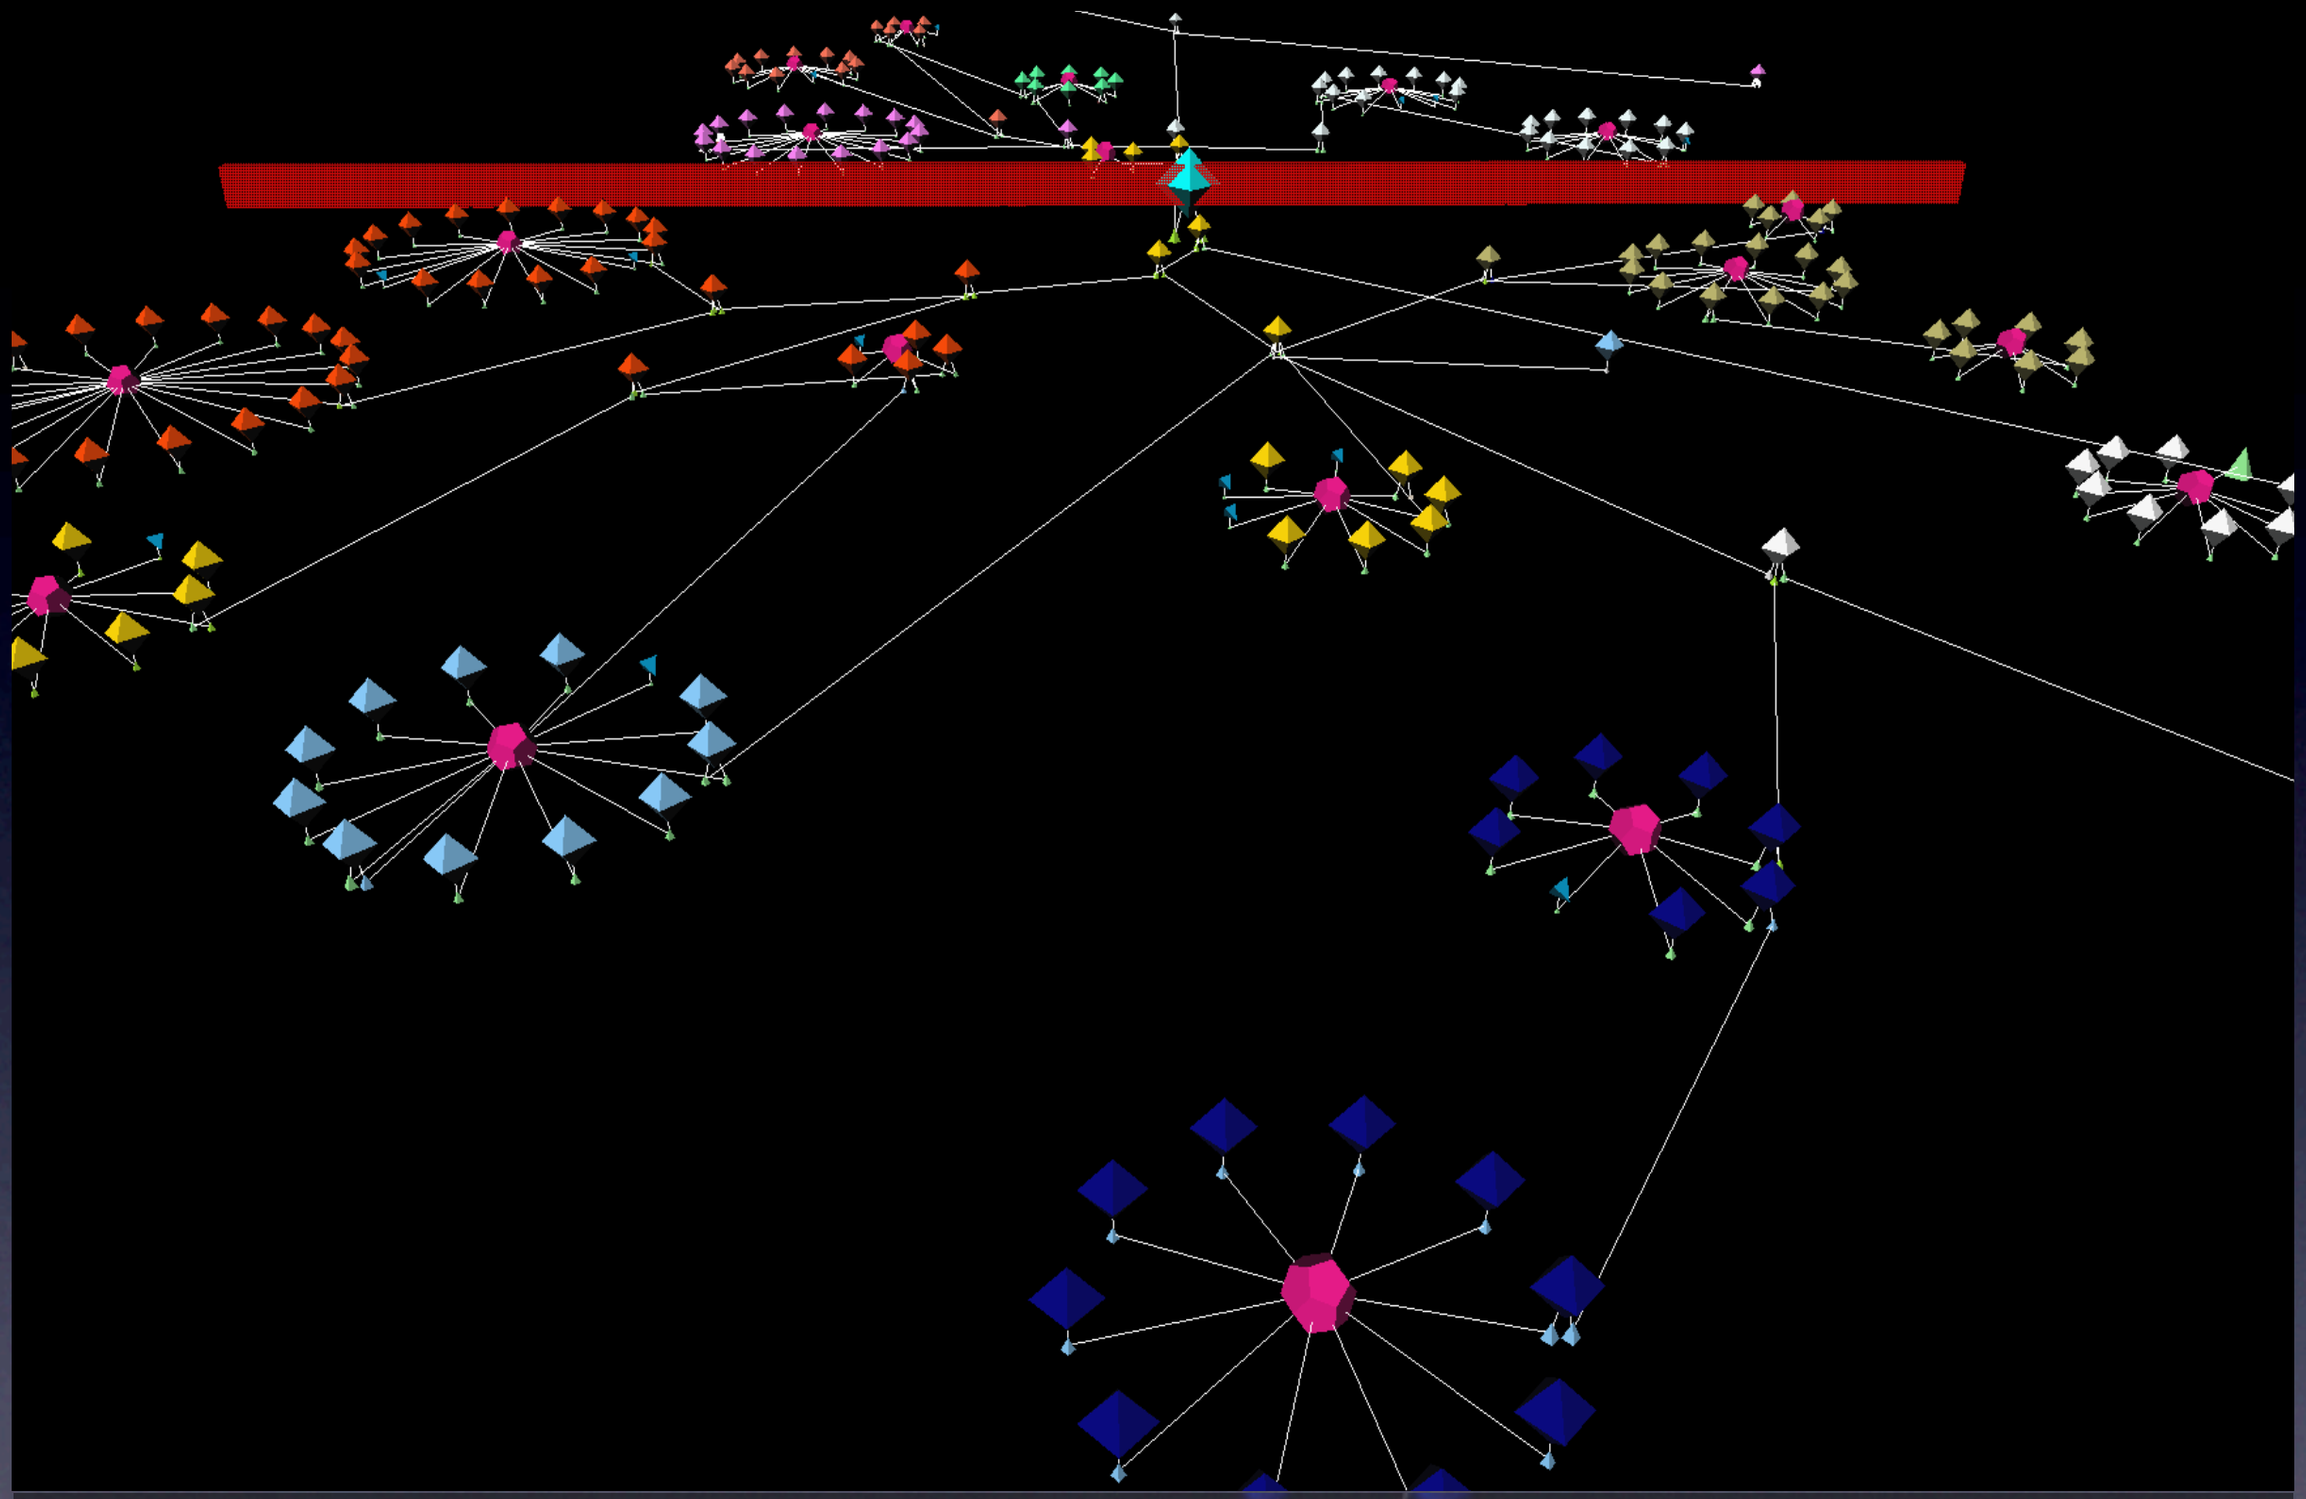
texture/icon 

connection / topology 

movement 

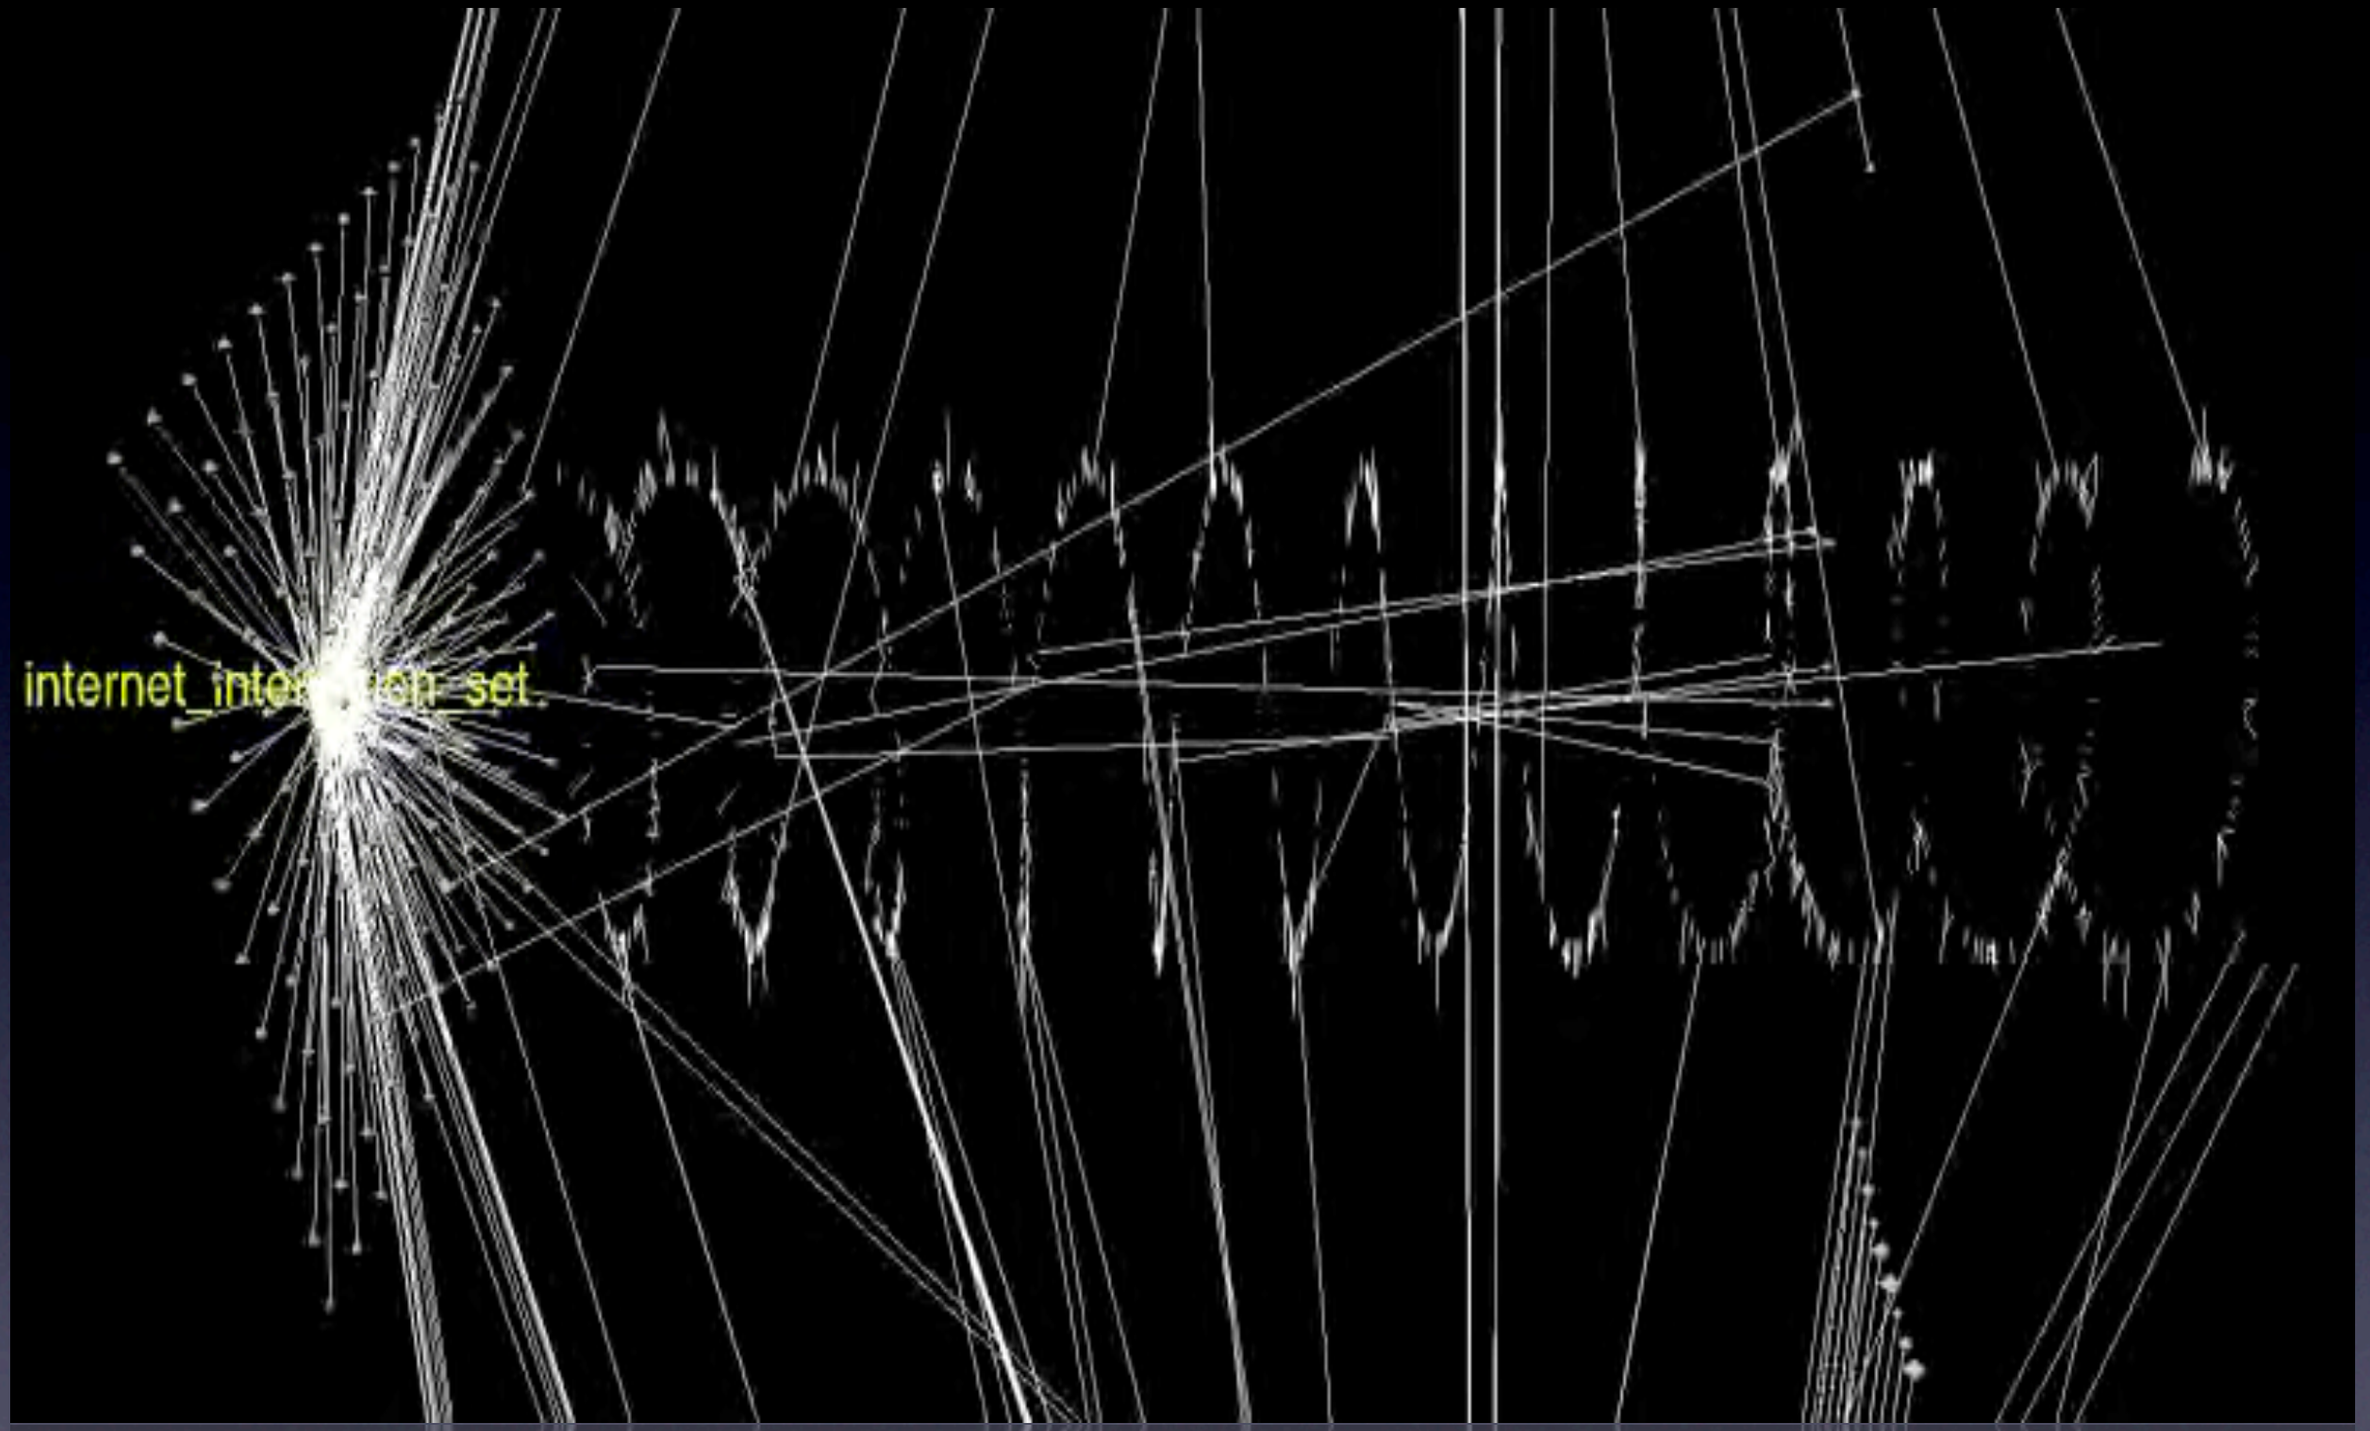
packet events, information exchange, attribute changes, attribute values, host id, software, processes, machine purpose, network topology, social topology, intrusion events, event type, event priority, client vs server, routing, ...

© 2007 Sunny Fugate



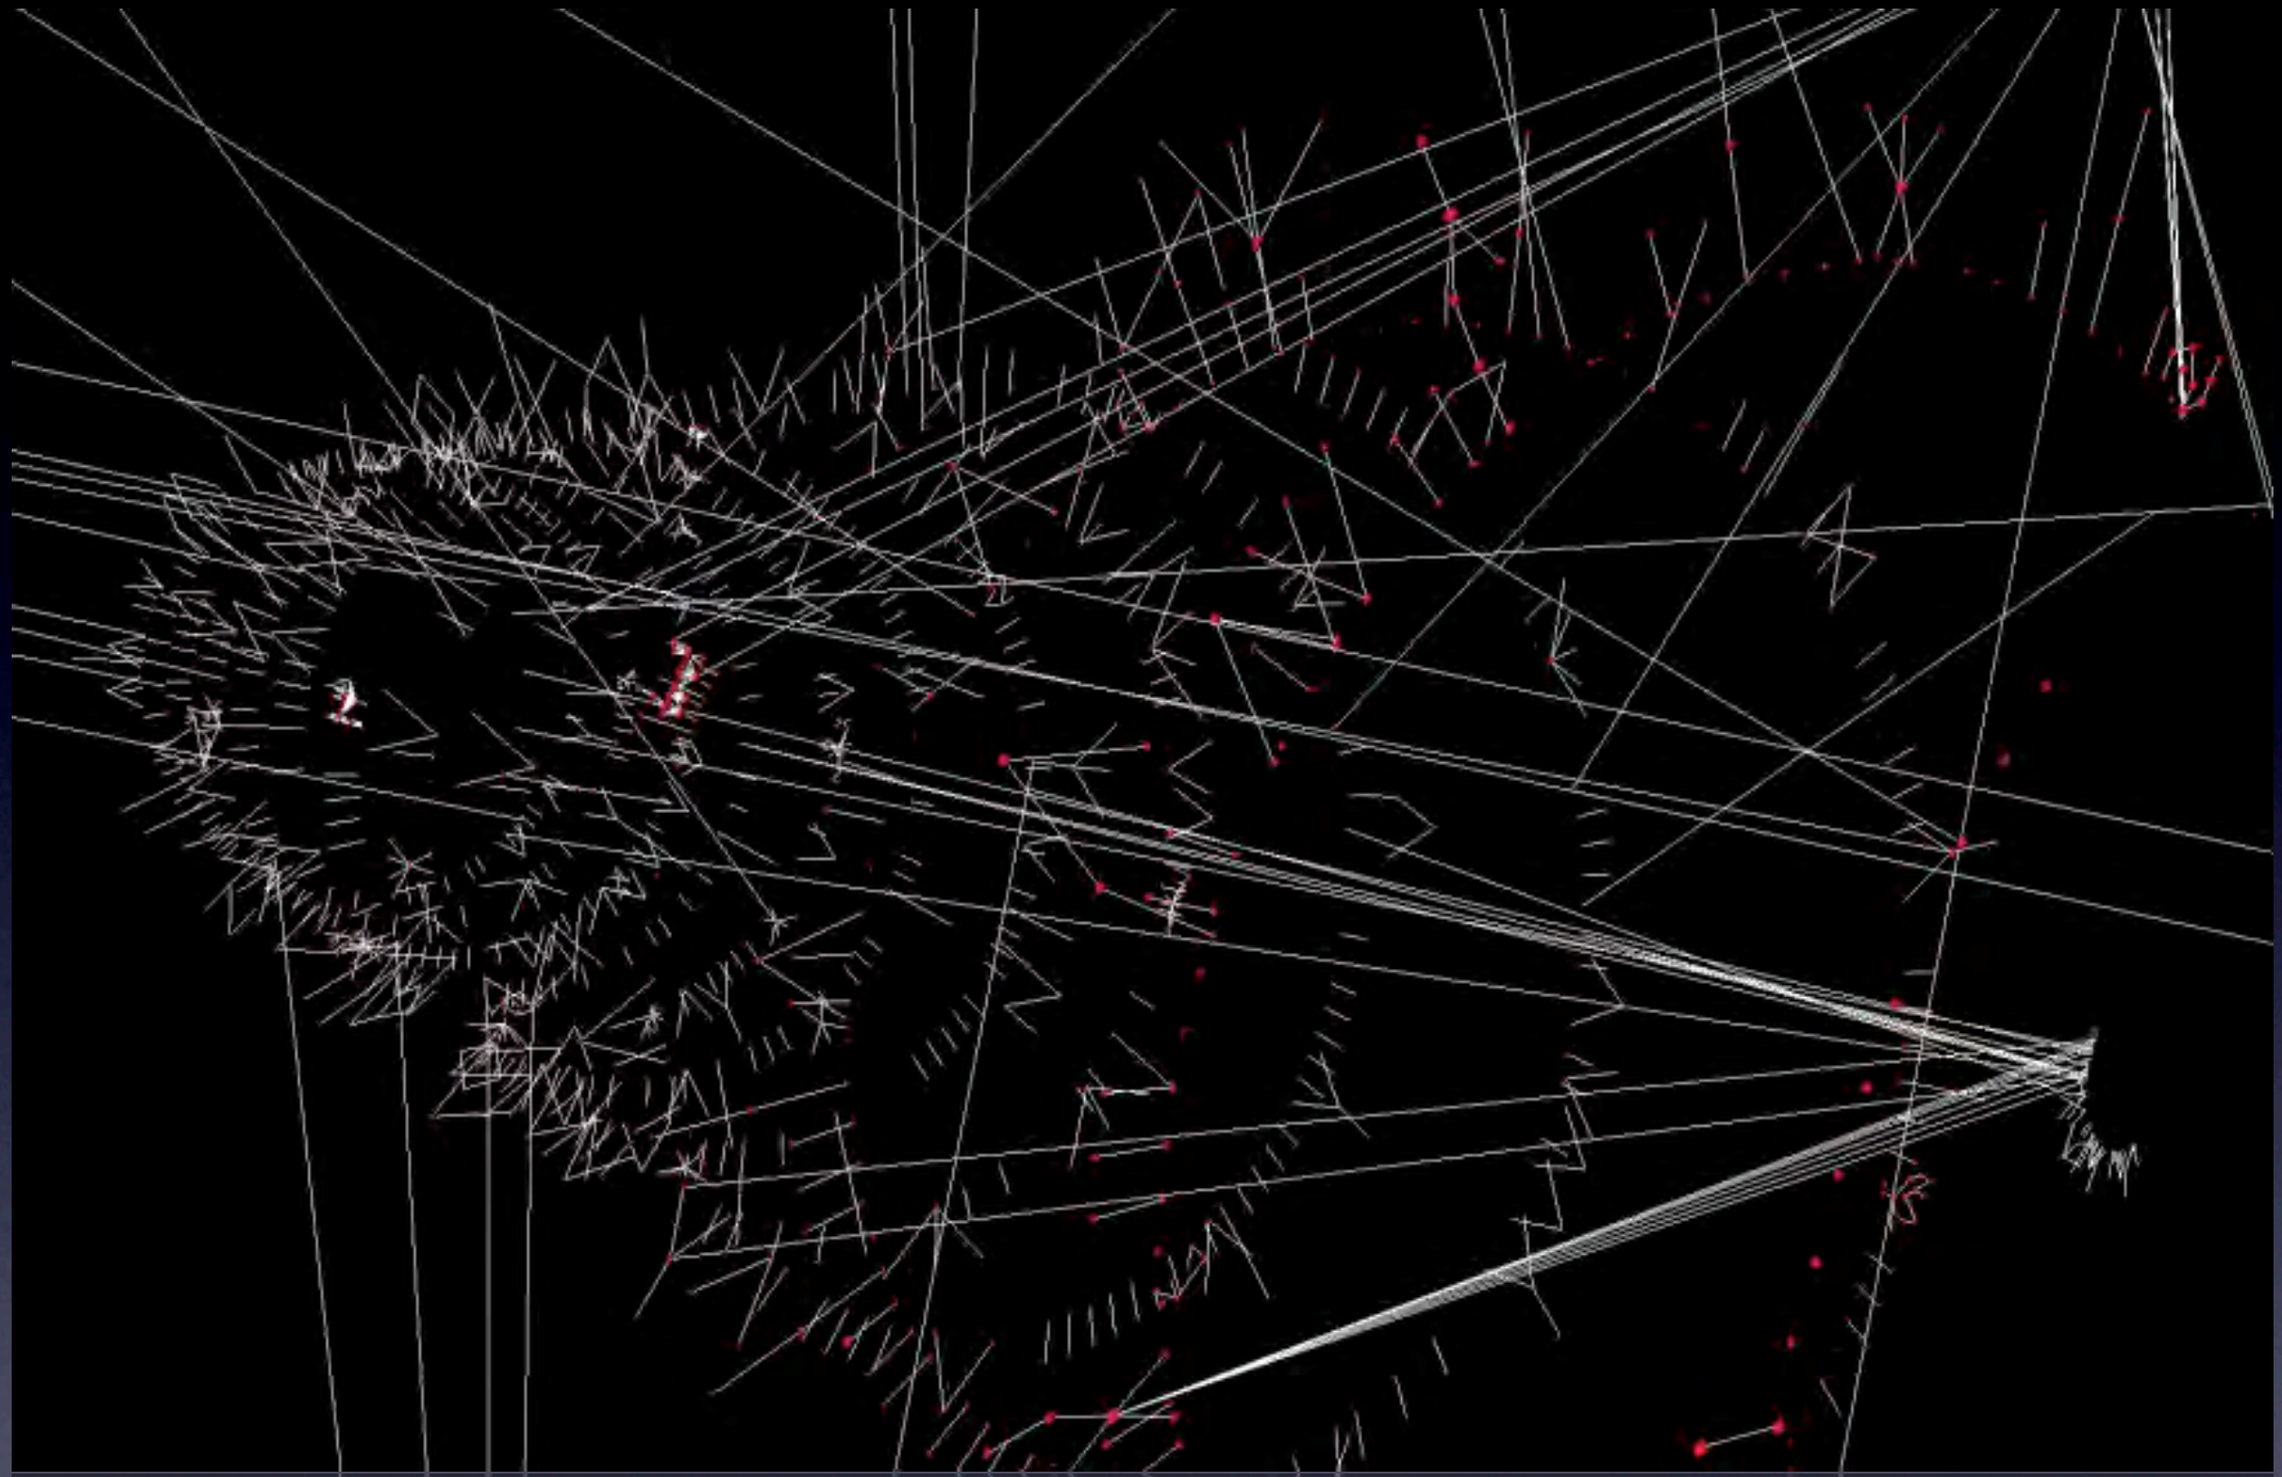
Toplogical layout using visual demarcations
(e.g. firewall, network segment, physical layout)

© 2007 Sunny Fugate

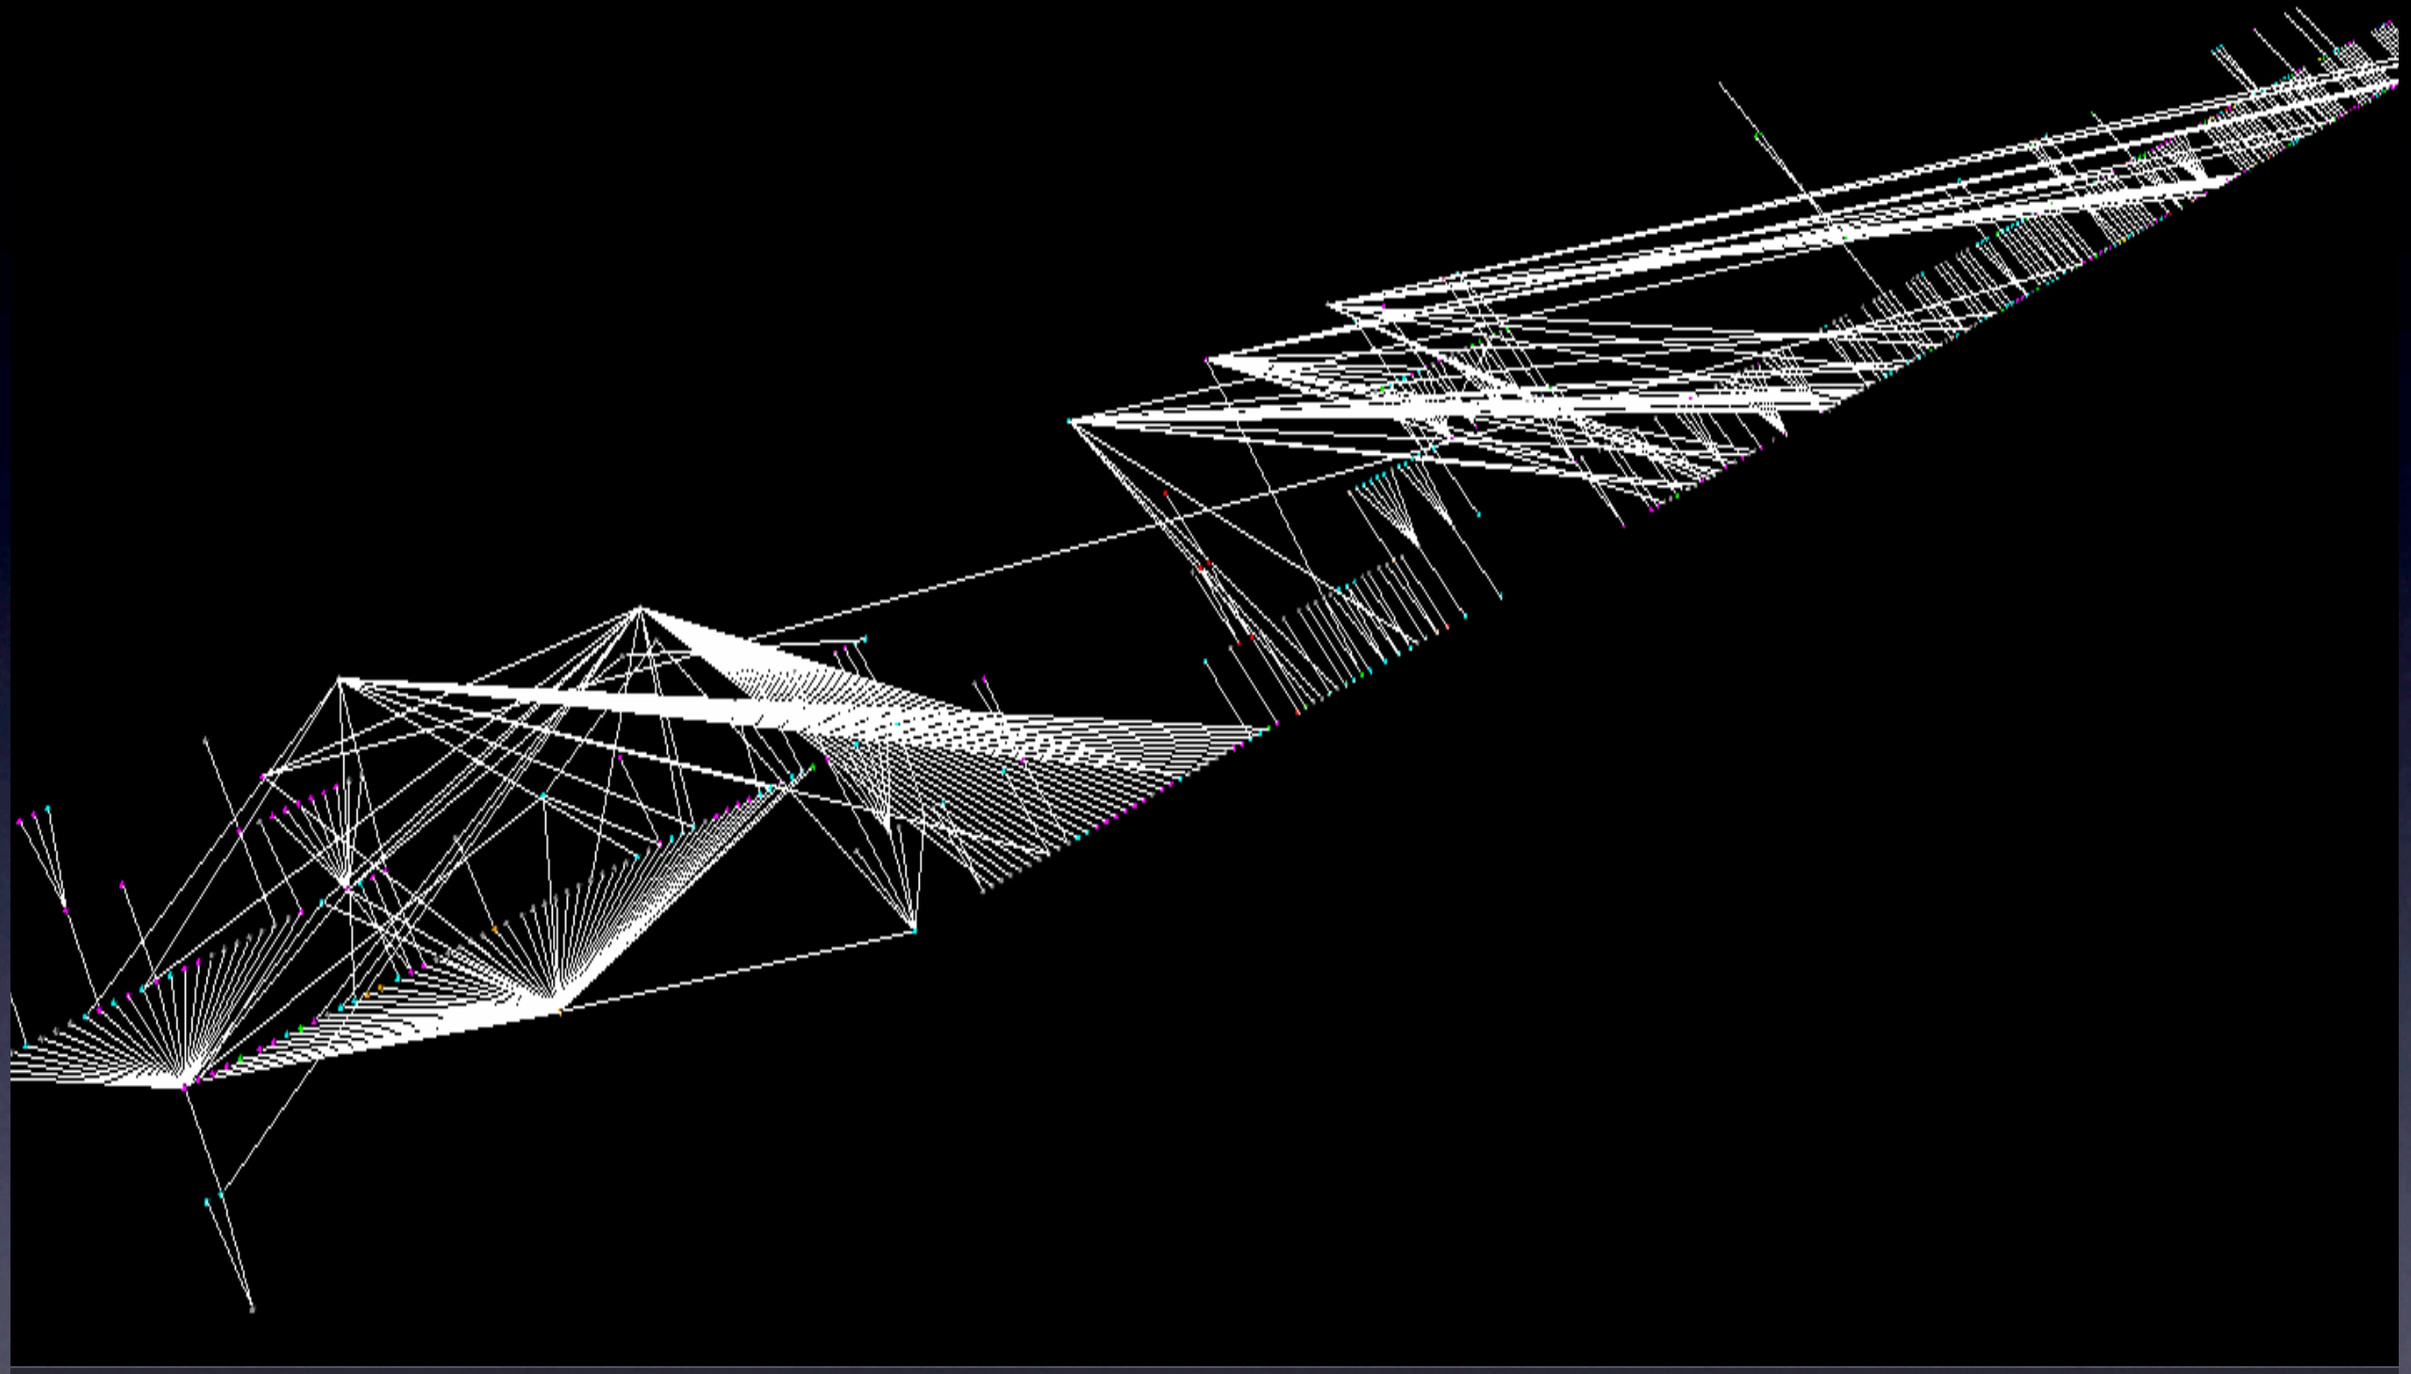


Automated layout to arrange hundreds of sub-graphs in a non-overlapping manner.

© 2007 Sunny Fugate

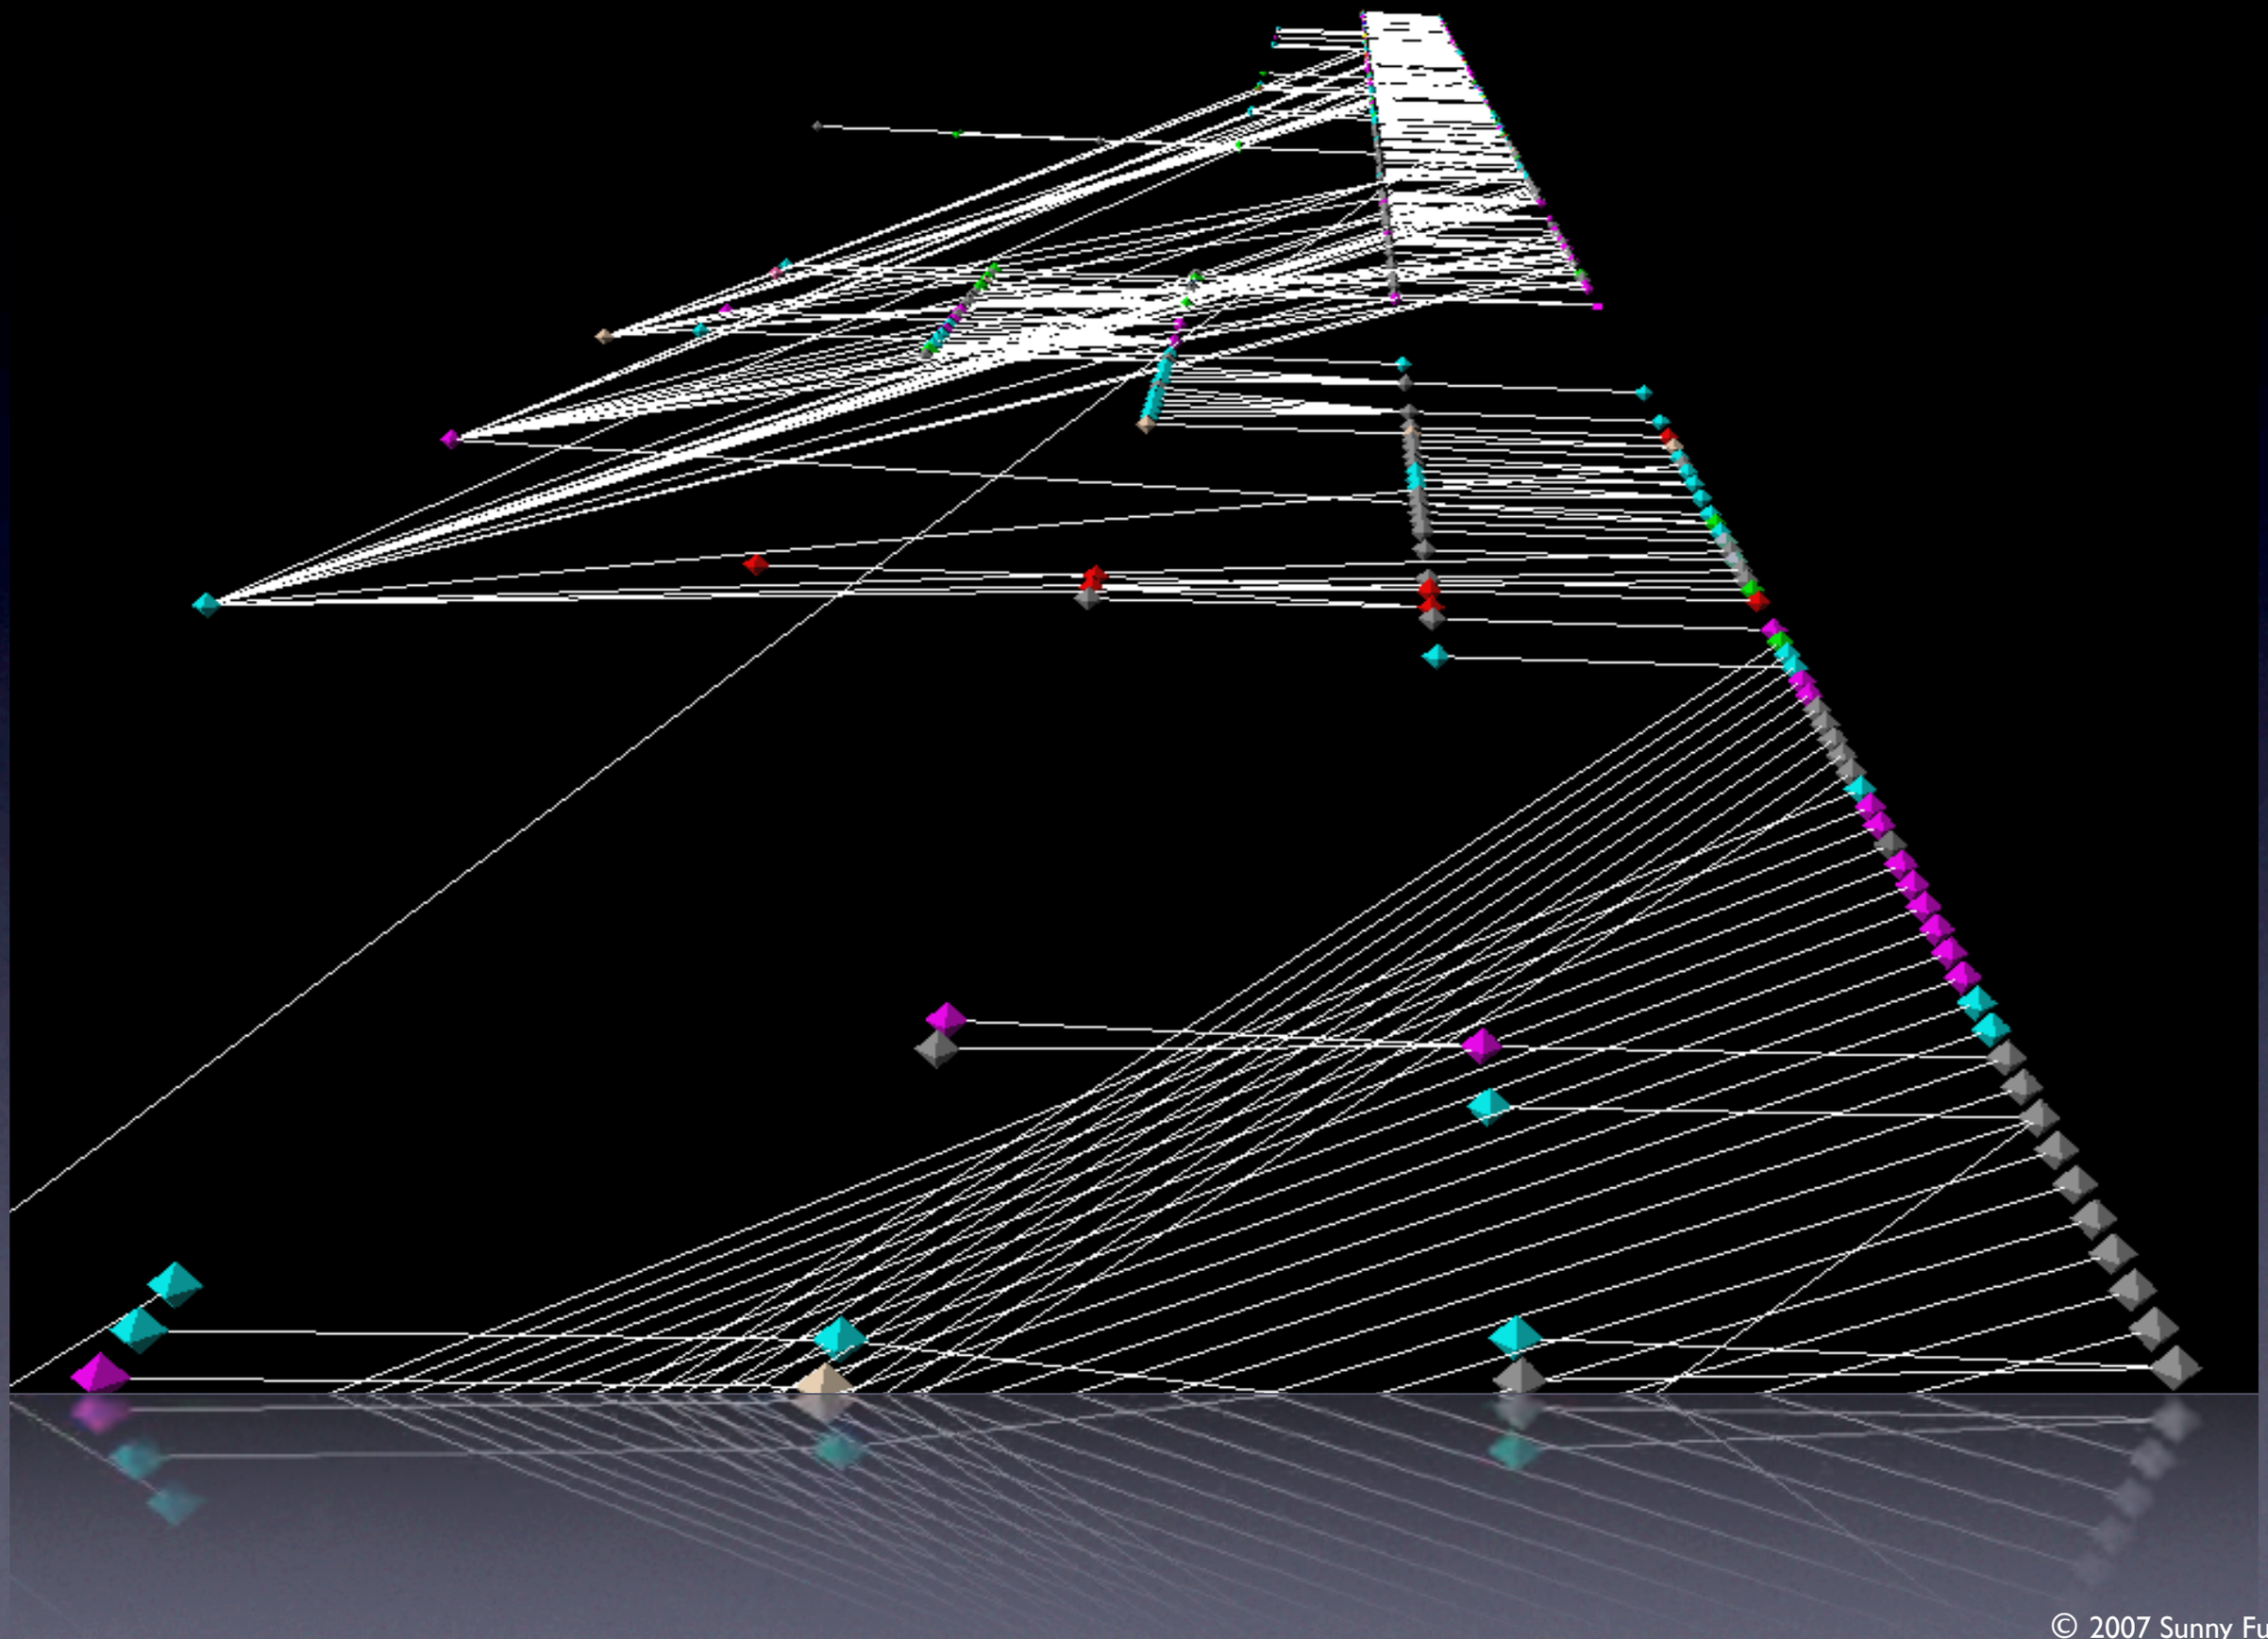


© 2007 Sunny Fugate

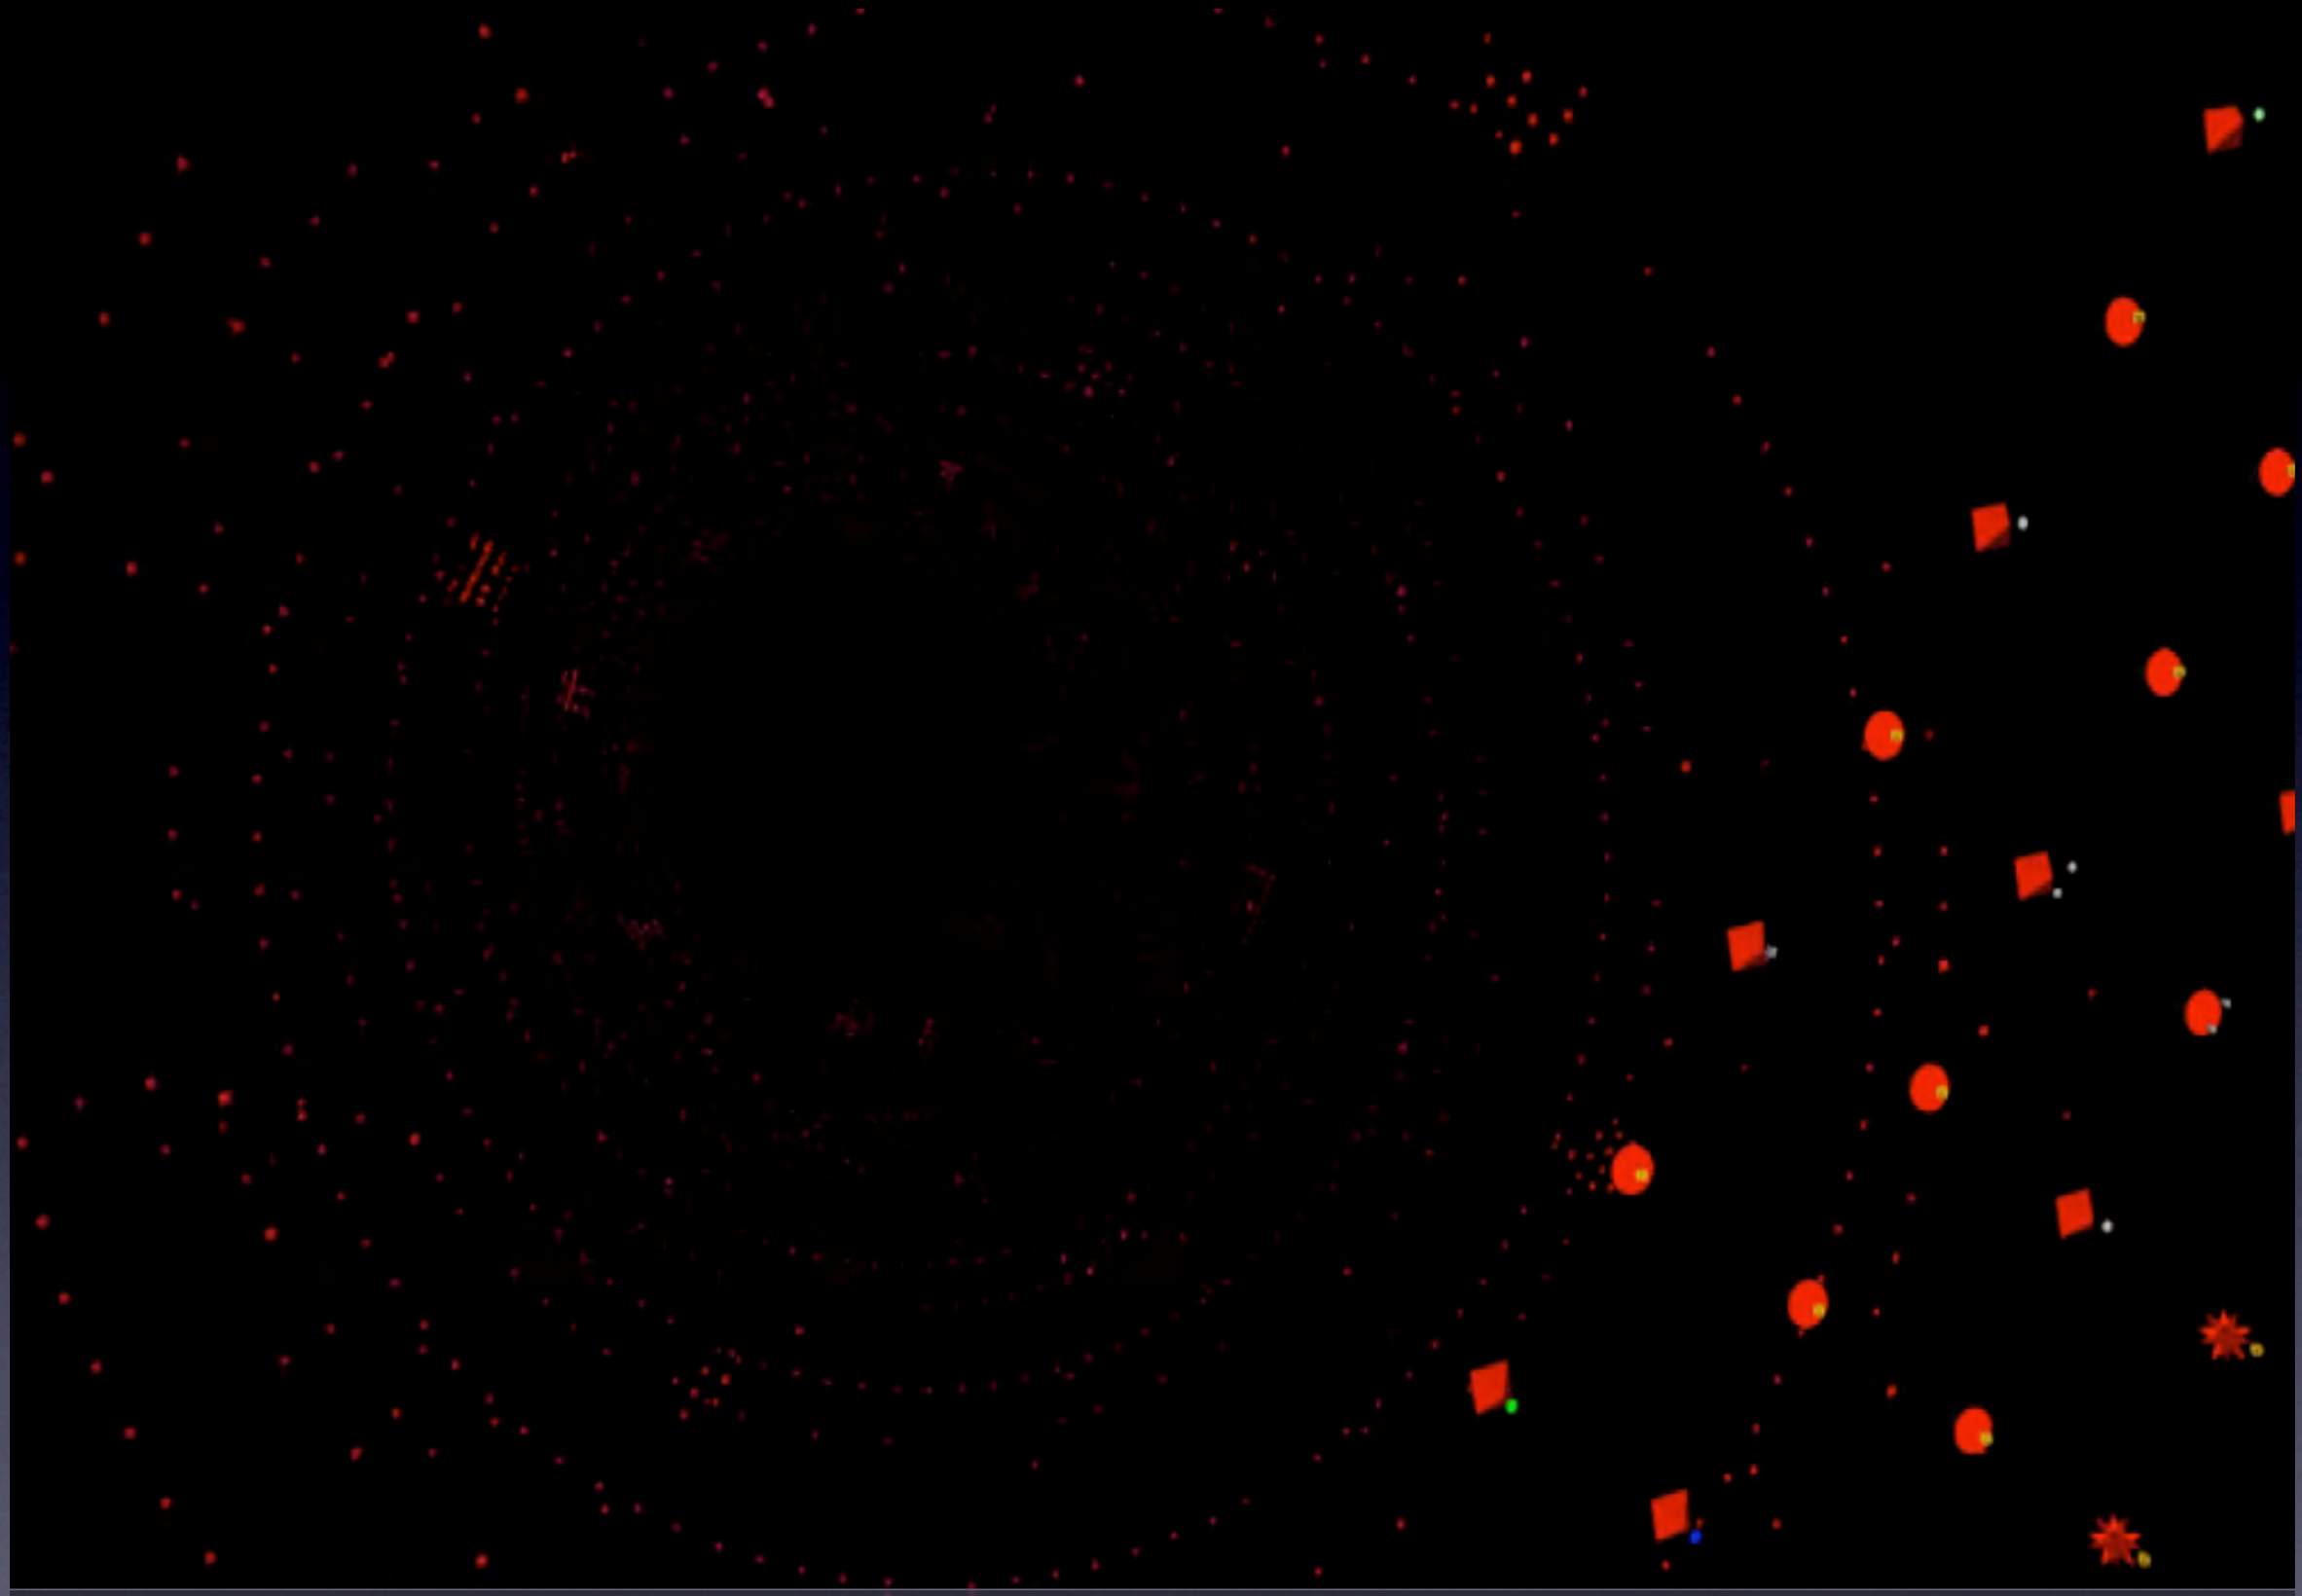


Topological layout discovered using hints in the data (e.g. TTL)

© 2007 Sunny Fugate

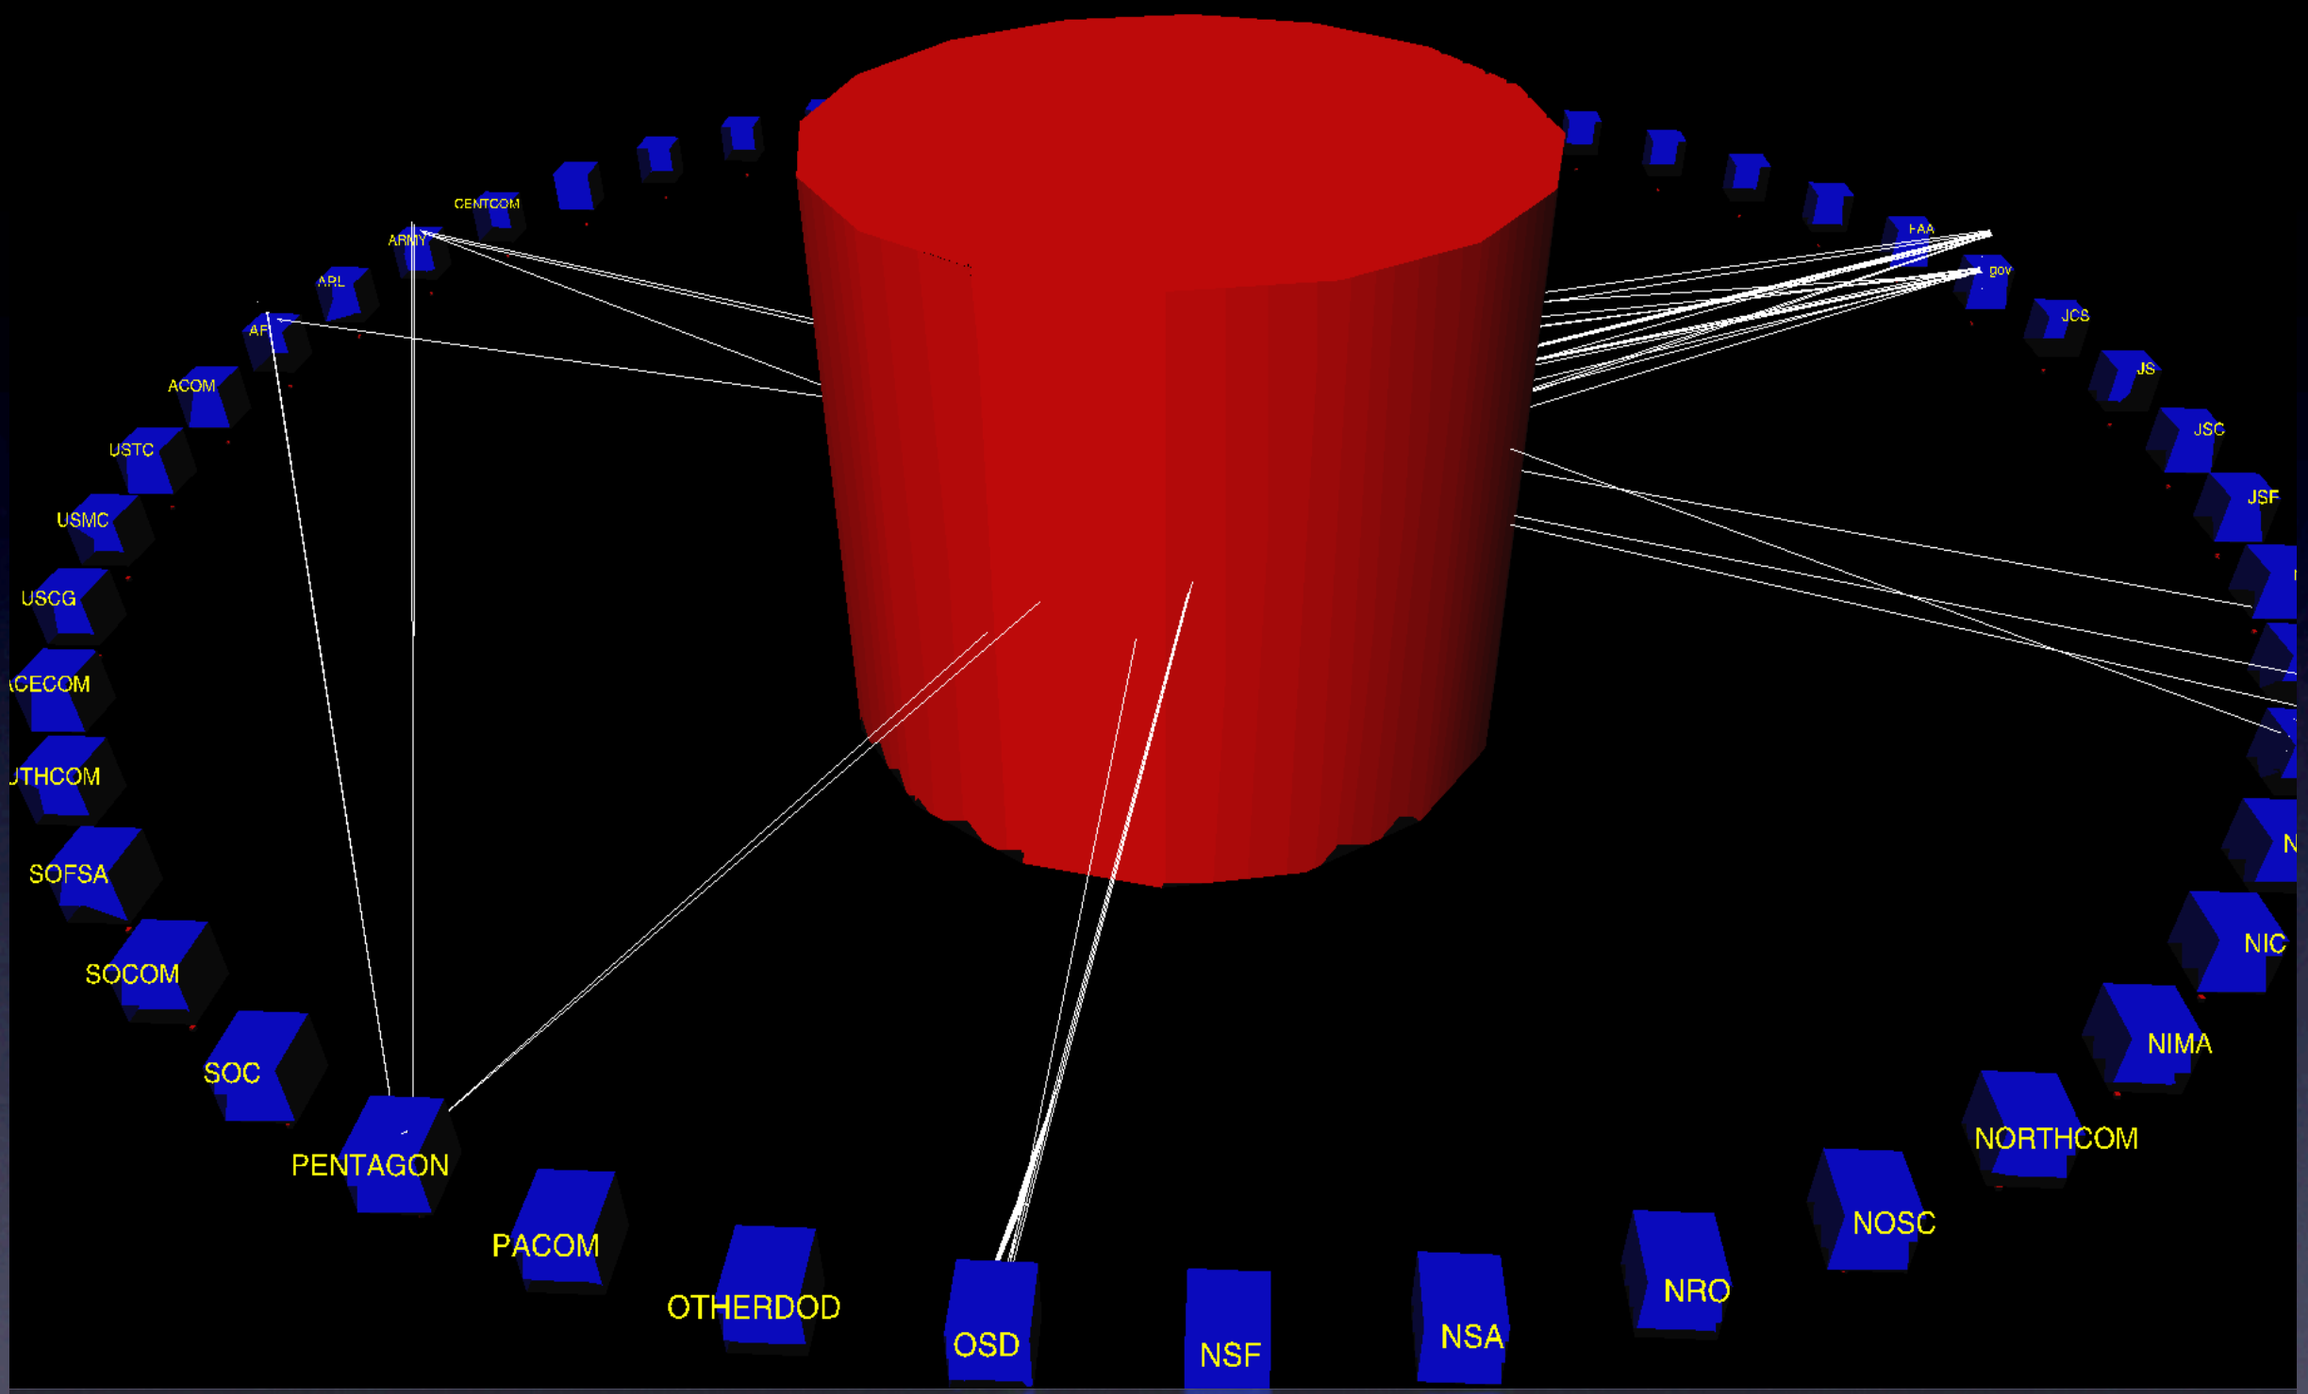


© 2007 Sunny Fugate



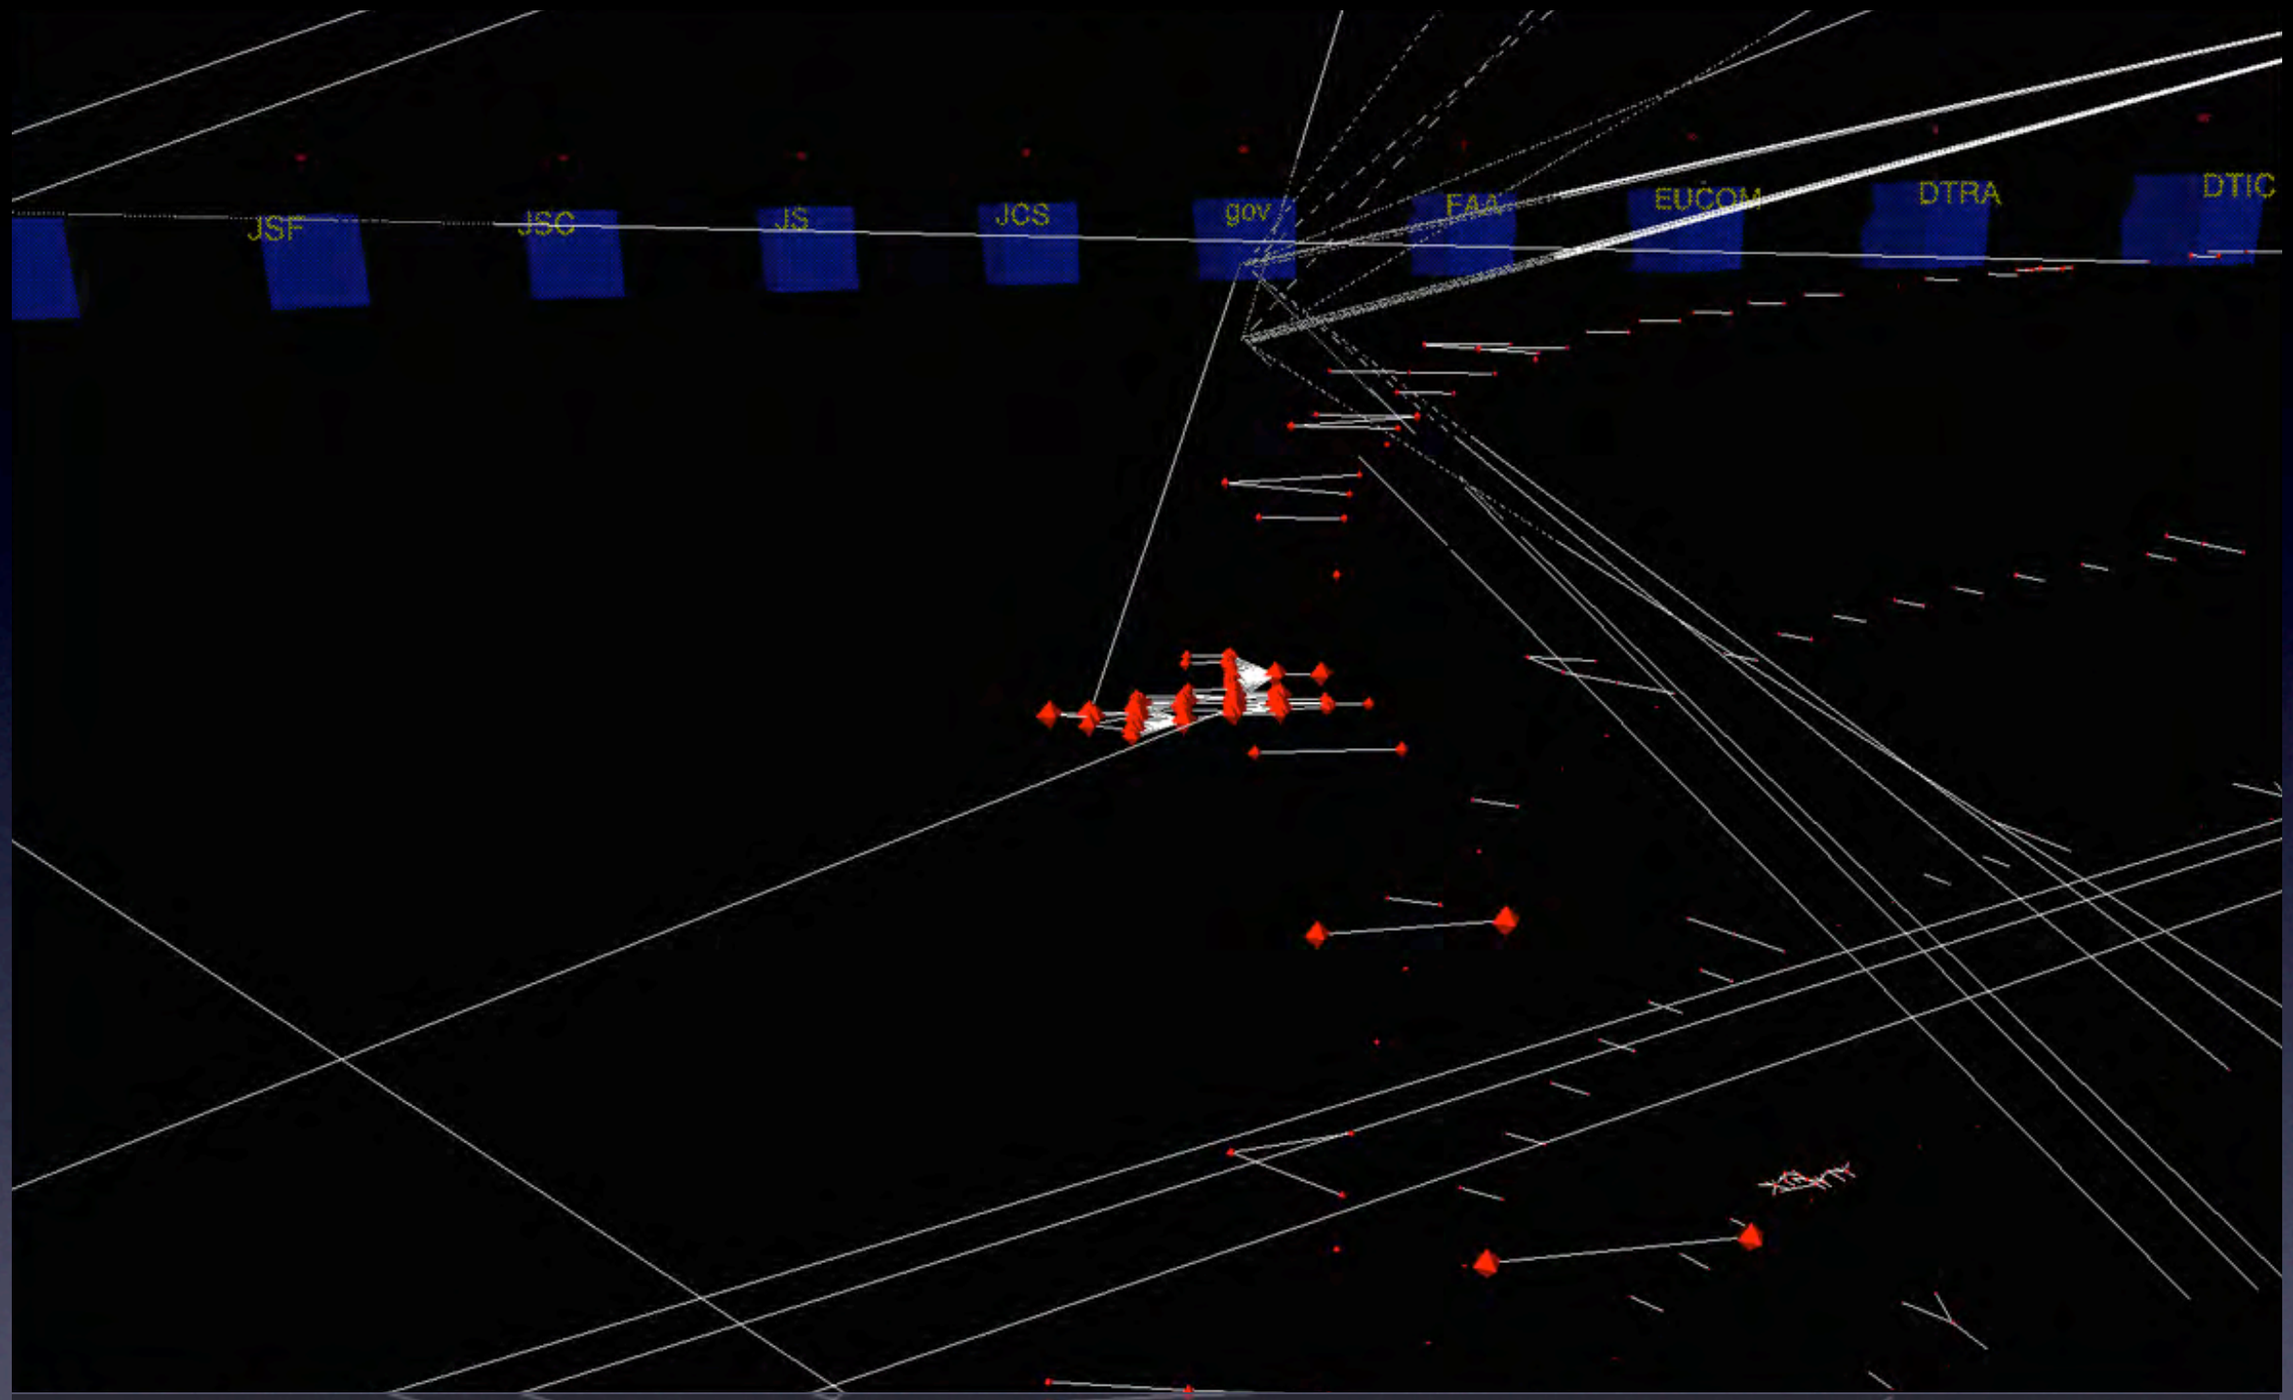
Color, shape, texture, icon, location, arrangement

© 2007 Sunny Fugate



Visual grouping, demarcation, and detail-hiding

© 2007 Sunny Fugate



Expansive vantage points for network analysis

© 2007 Sunny Fugate

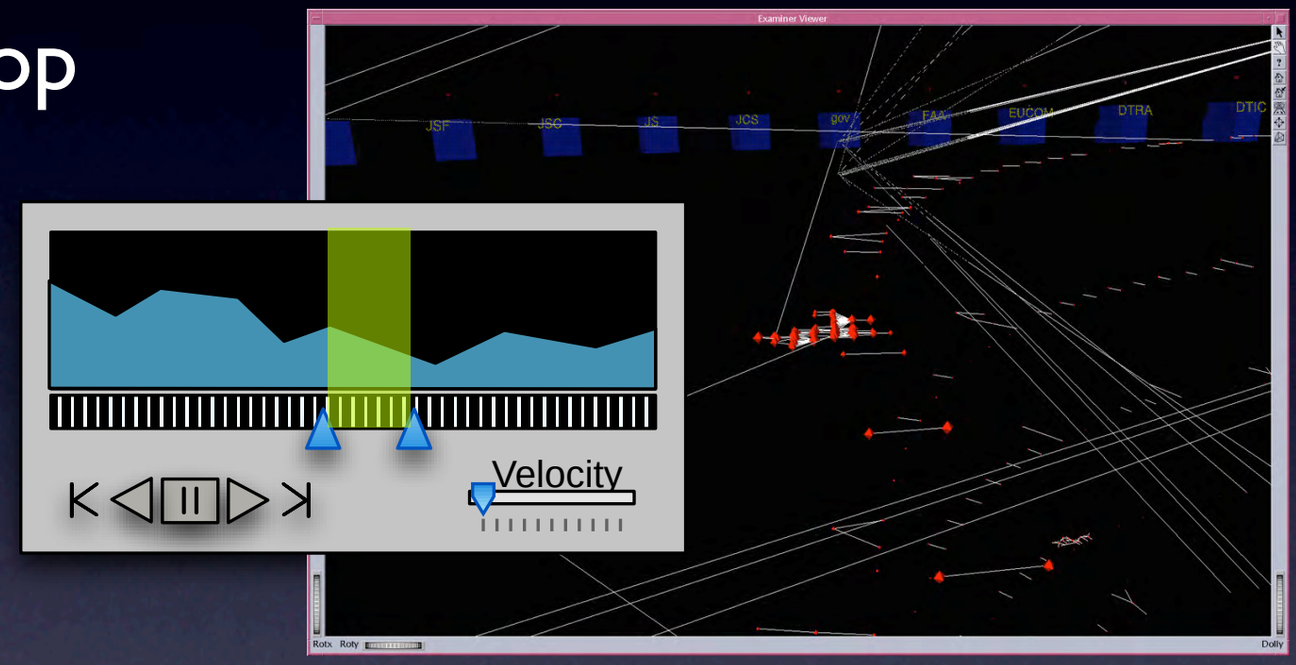
Shapes Vector Flow Viewer

- JTF-GNO funded effort to implement SV
  - Use SV architecture and components
  - DARPA demo system > operational system
  - New scripts, sensors, agents, and GUI
- Results
  - A visual augmentation of CLI
  - Produces a view of social topology
  - Intuitive view of gobs of data
    - static topology and event replay
  - Links statistical views and topology view

© 2007 Sunny Fugate

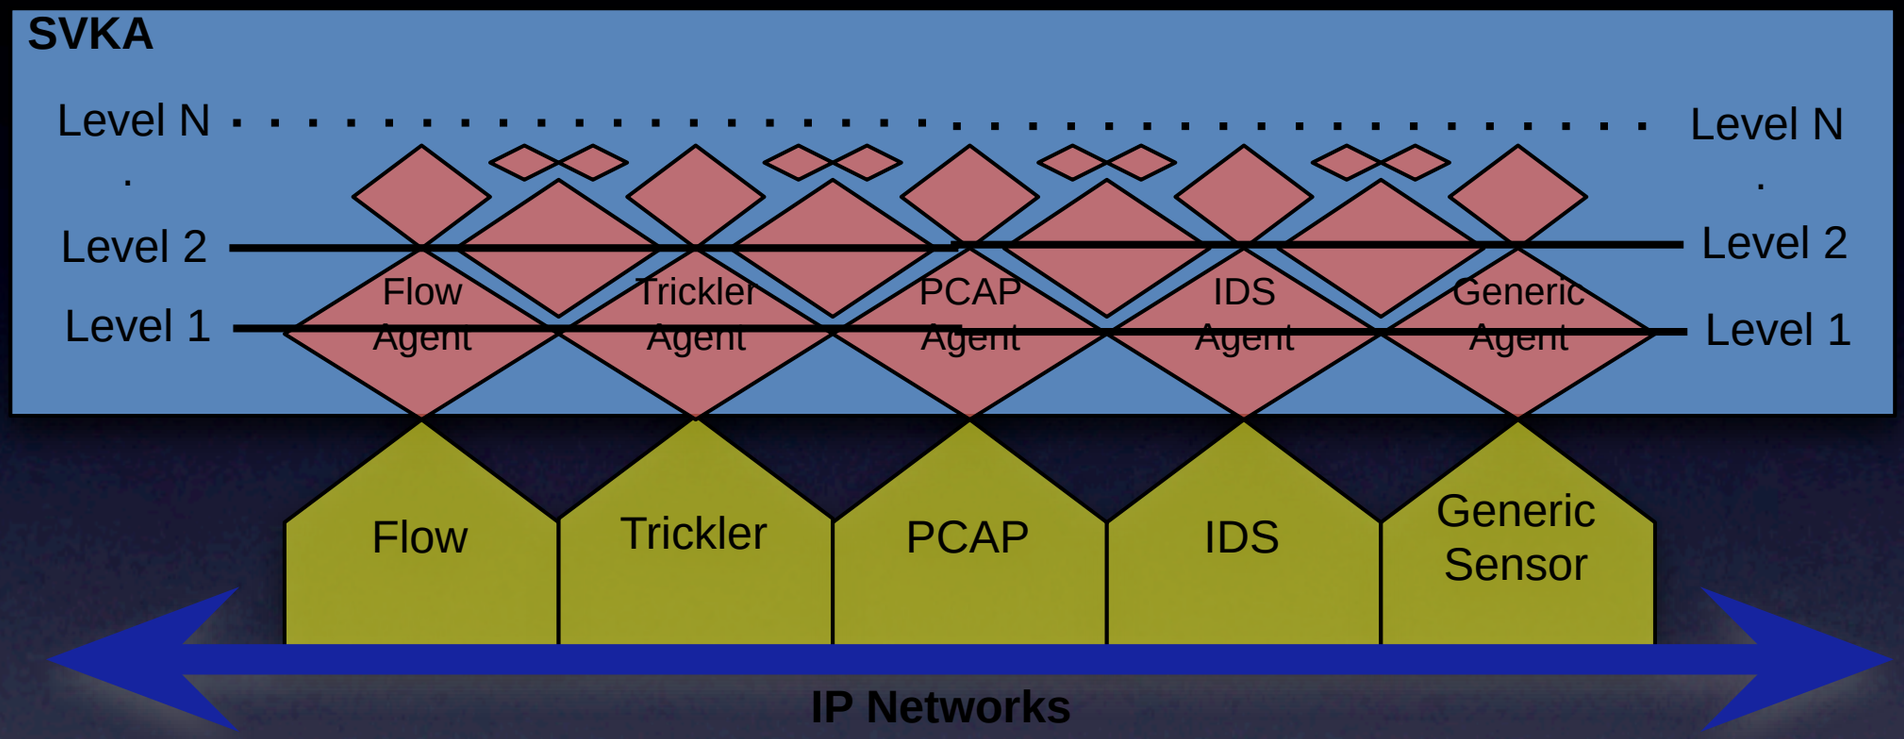
Flow Viewer GUI

- multiple stats views linked to visuals
- playback specific ranges & loop
- adjust replay velocity
- time-skip
- IP and attribute hotlists
- dynamic filtering controls
  - GUI managed rwfilter
  - filter using SV ontology
- integration between flow, Trickler, IDS, & PCAP



© 2007 Sunny Fugate

Flow Viewer Sensors



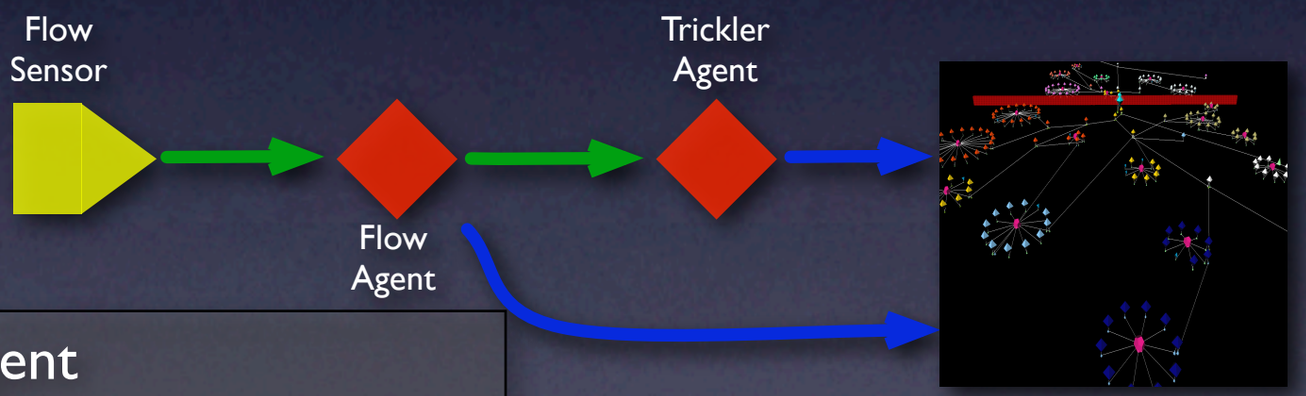
Flow Agent	consumes rwf & rwcut data ✓
Trickler Agent	queries database for most recent attributes ✓
PCAP Agents	queries & reconstructs TCP sessions
IDS Agents	processes IDS logs

© 2007 Sunny Fugate

Flow Viewer
Intelligent Agents

- FlowSensor
  - Converts flow into ontology
  - produces facts

- TricklerAgent
  - uses correlations from FlowAgent
  - query made on every unique IP seen
  - produces visual events

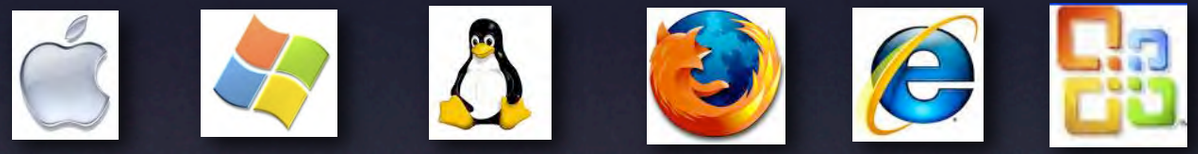


- FlowAgent
  - correlates records
  - counts and corroborates
  - produces inferences
  - produces visual events

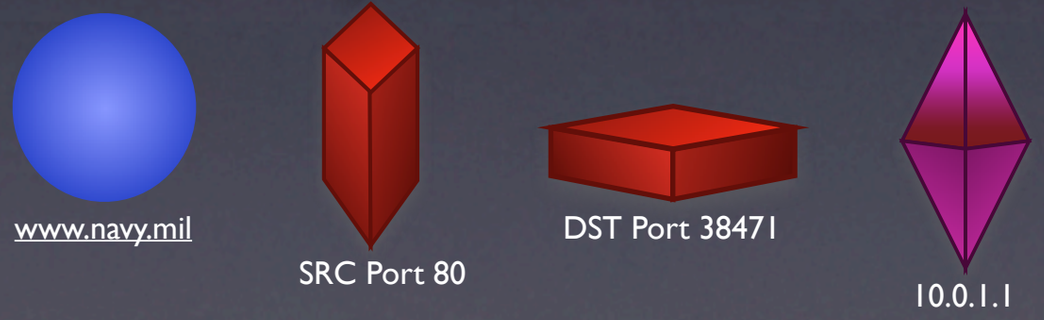
© 2007 Sunny Fugate

Flow Viewer Visual Language

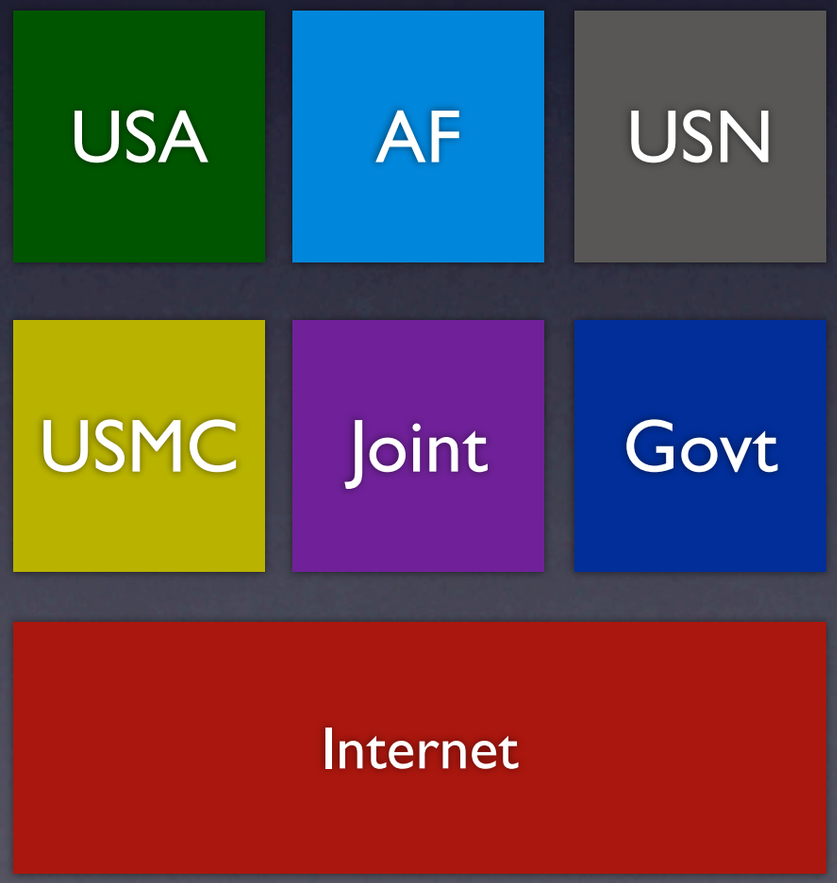
Leverage cultural knowledge



Use metaphors for abstract

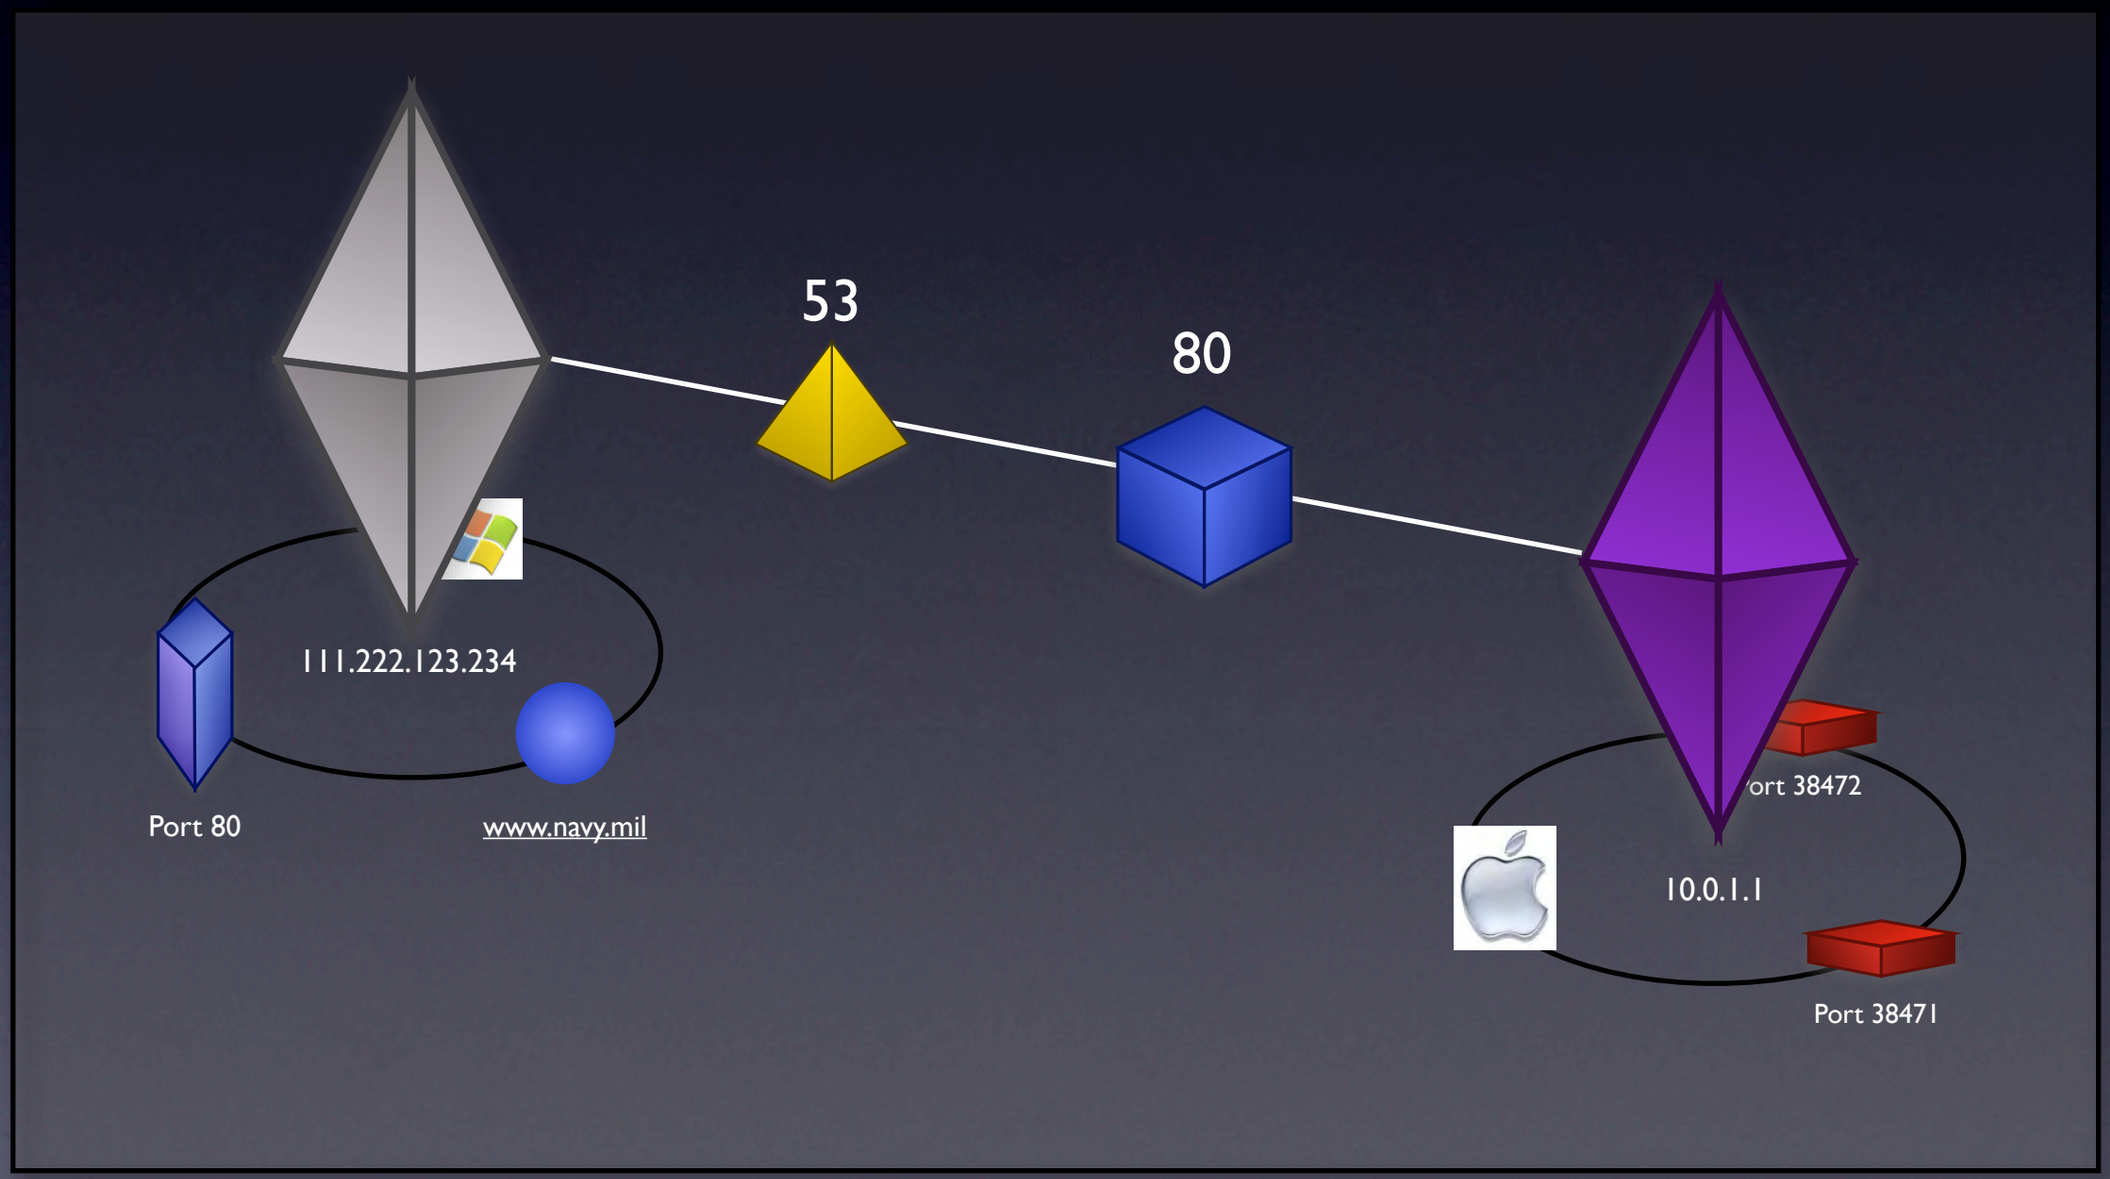


Color by ownership



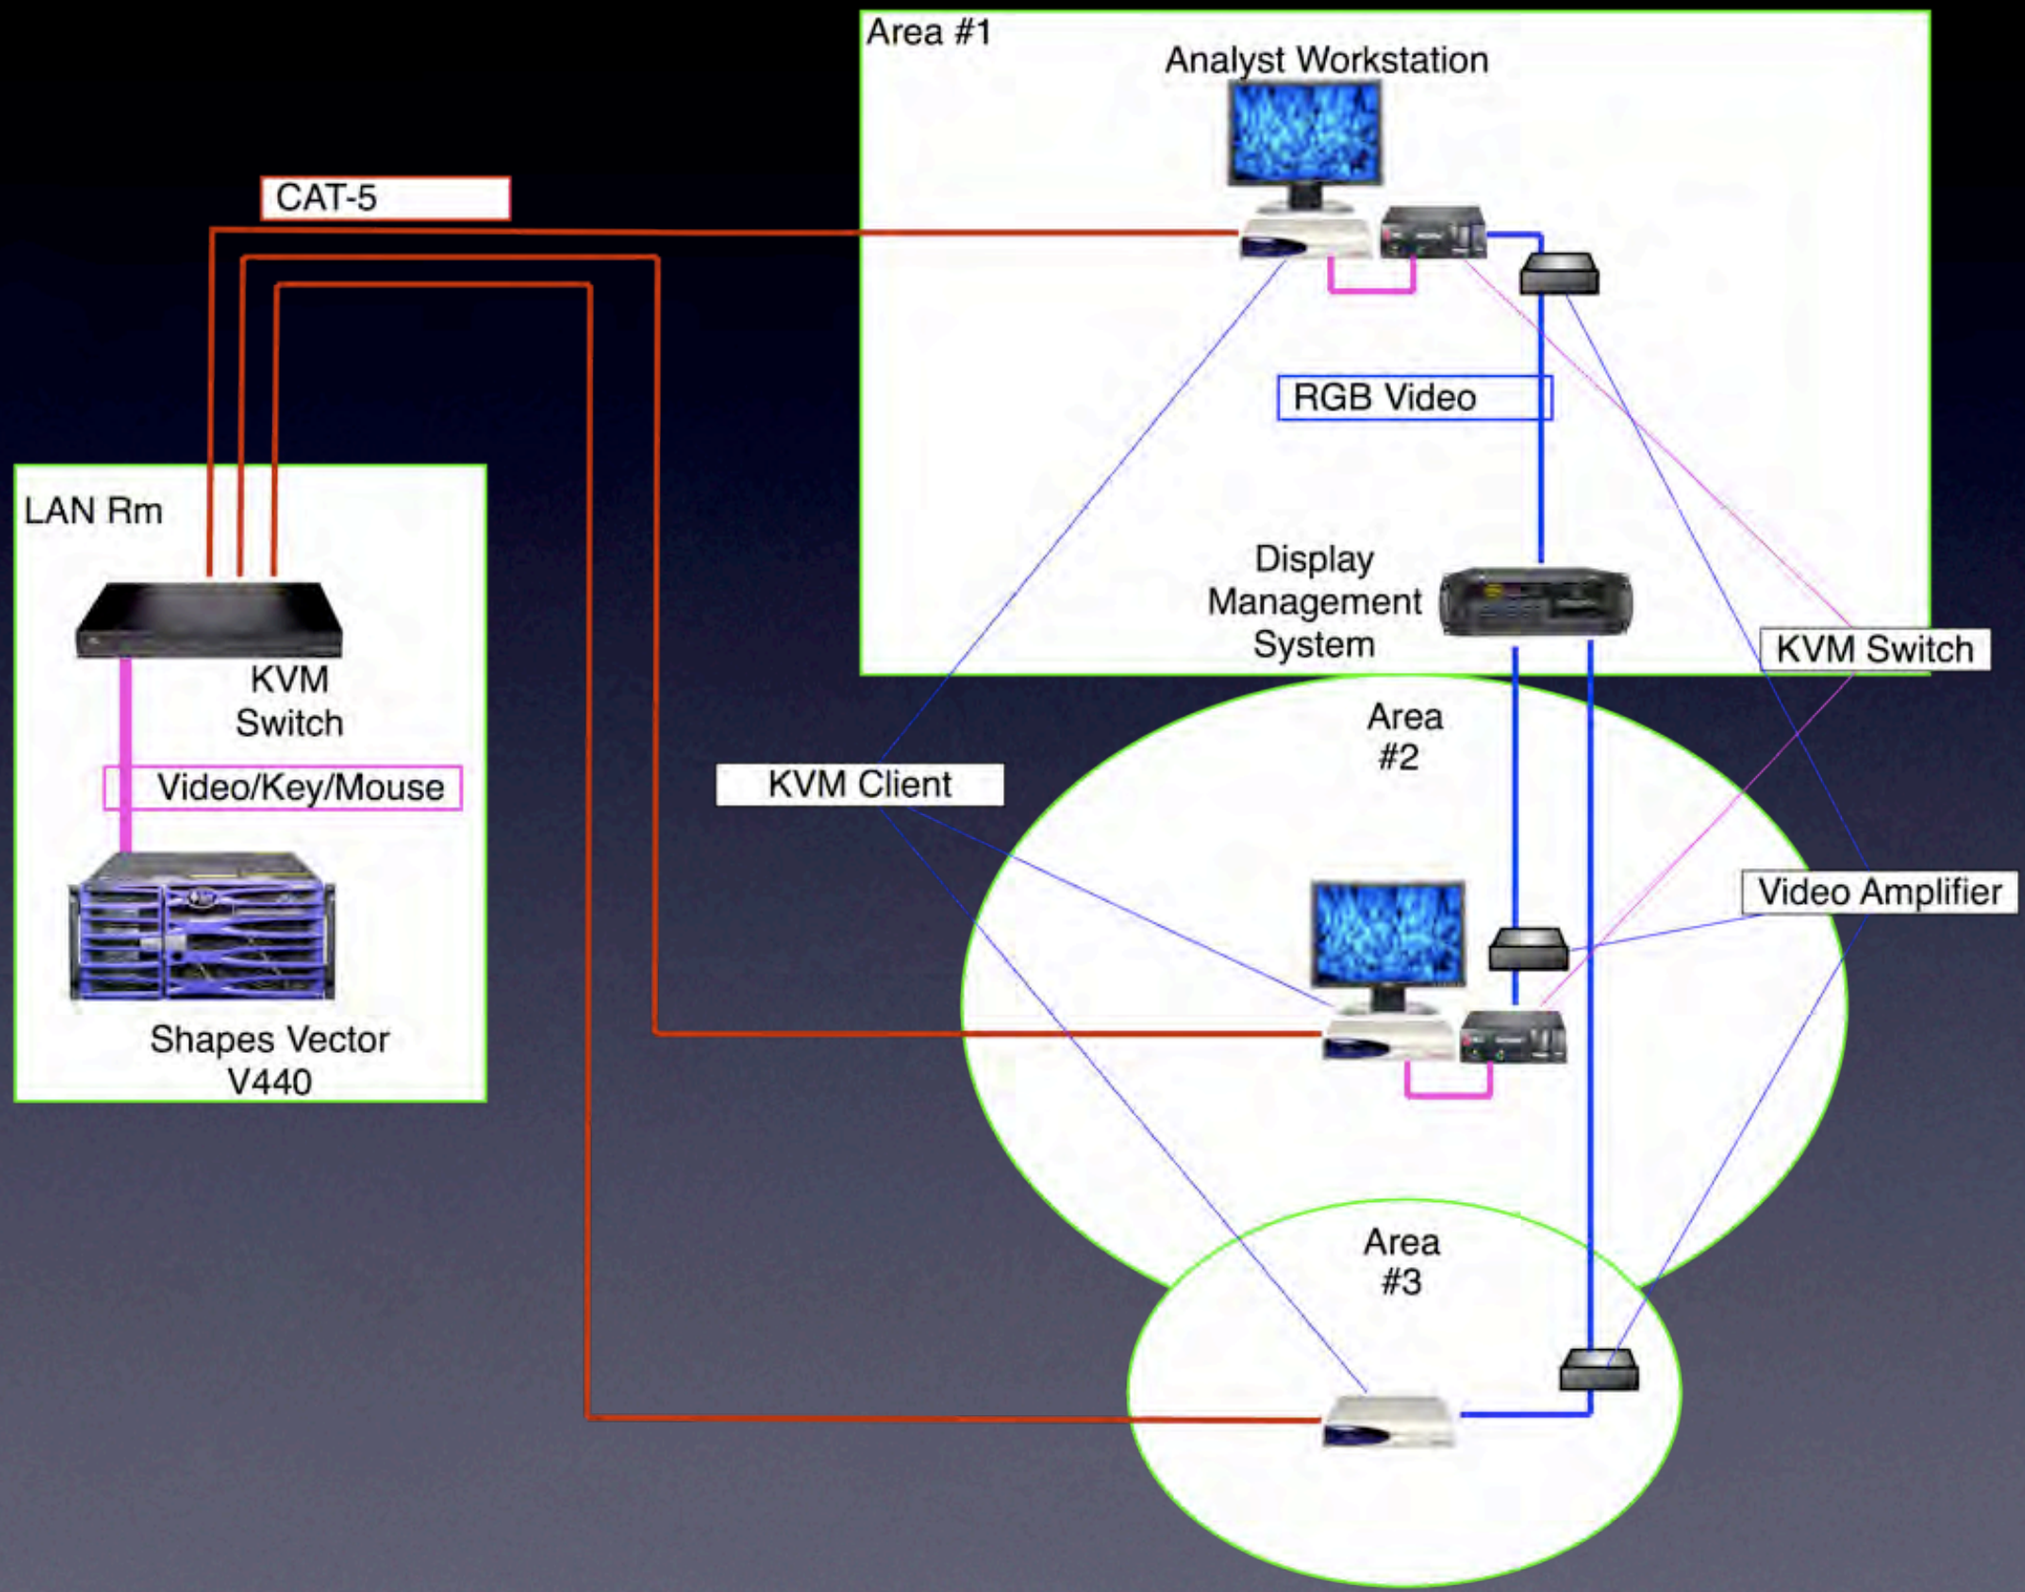
© 2007 Sunny Fugate

Flow Viewer
Visual Language



© 2007 Sunny Fugate

Test installation



© 2007 Sunny Fugate

Flow Viewer

Visualization

- **Tested** using:
  - 100-5000 nodes
  - 1M-3M flows
  - 10K-300K flows per hour
- Integrated **filtering** (rwfilter, SVKA filtering, visual)
  - Visual ID
  - **Queries**
  - **Grouping** (e.g. domain, netblock, vulnerability)
- **Replay**-mode or Real-time
- Historic **visual context**
  - Replay 'on top of' known incident

© 2007 Sunny Fugate

Flow Viewer data prep

Include

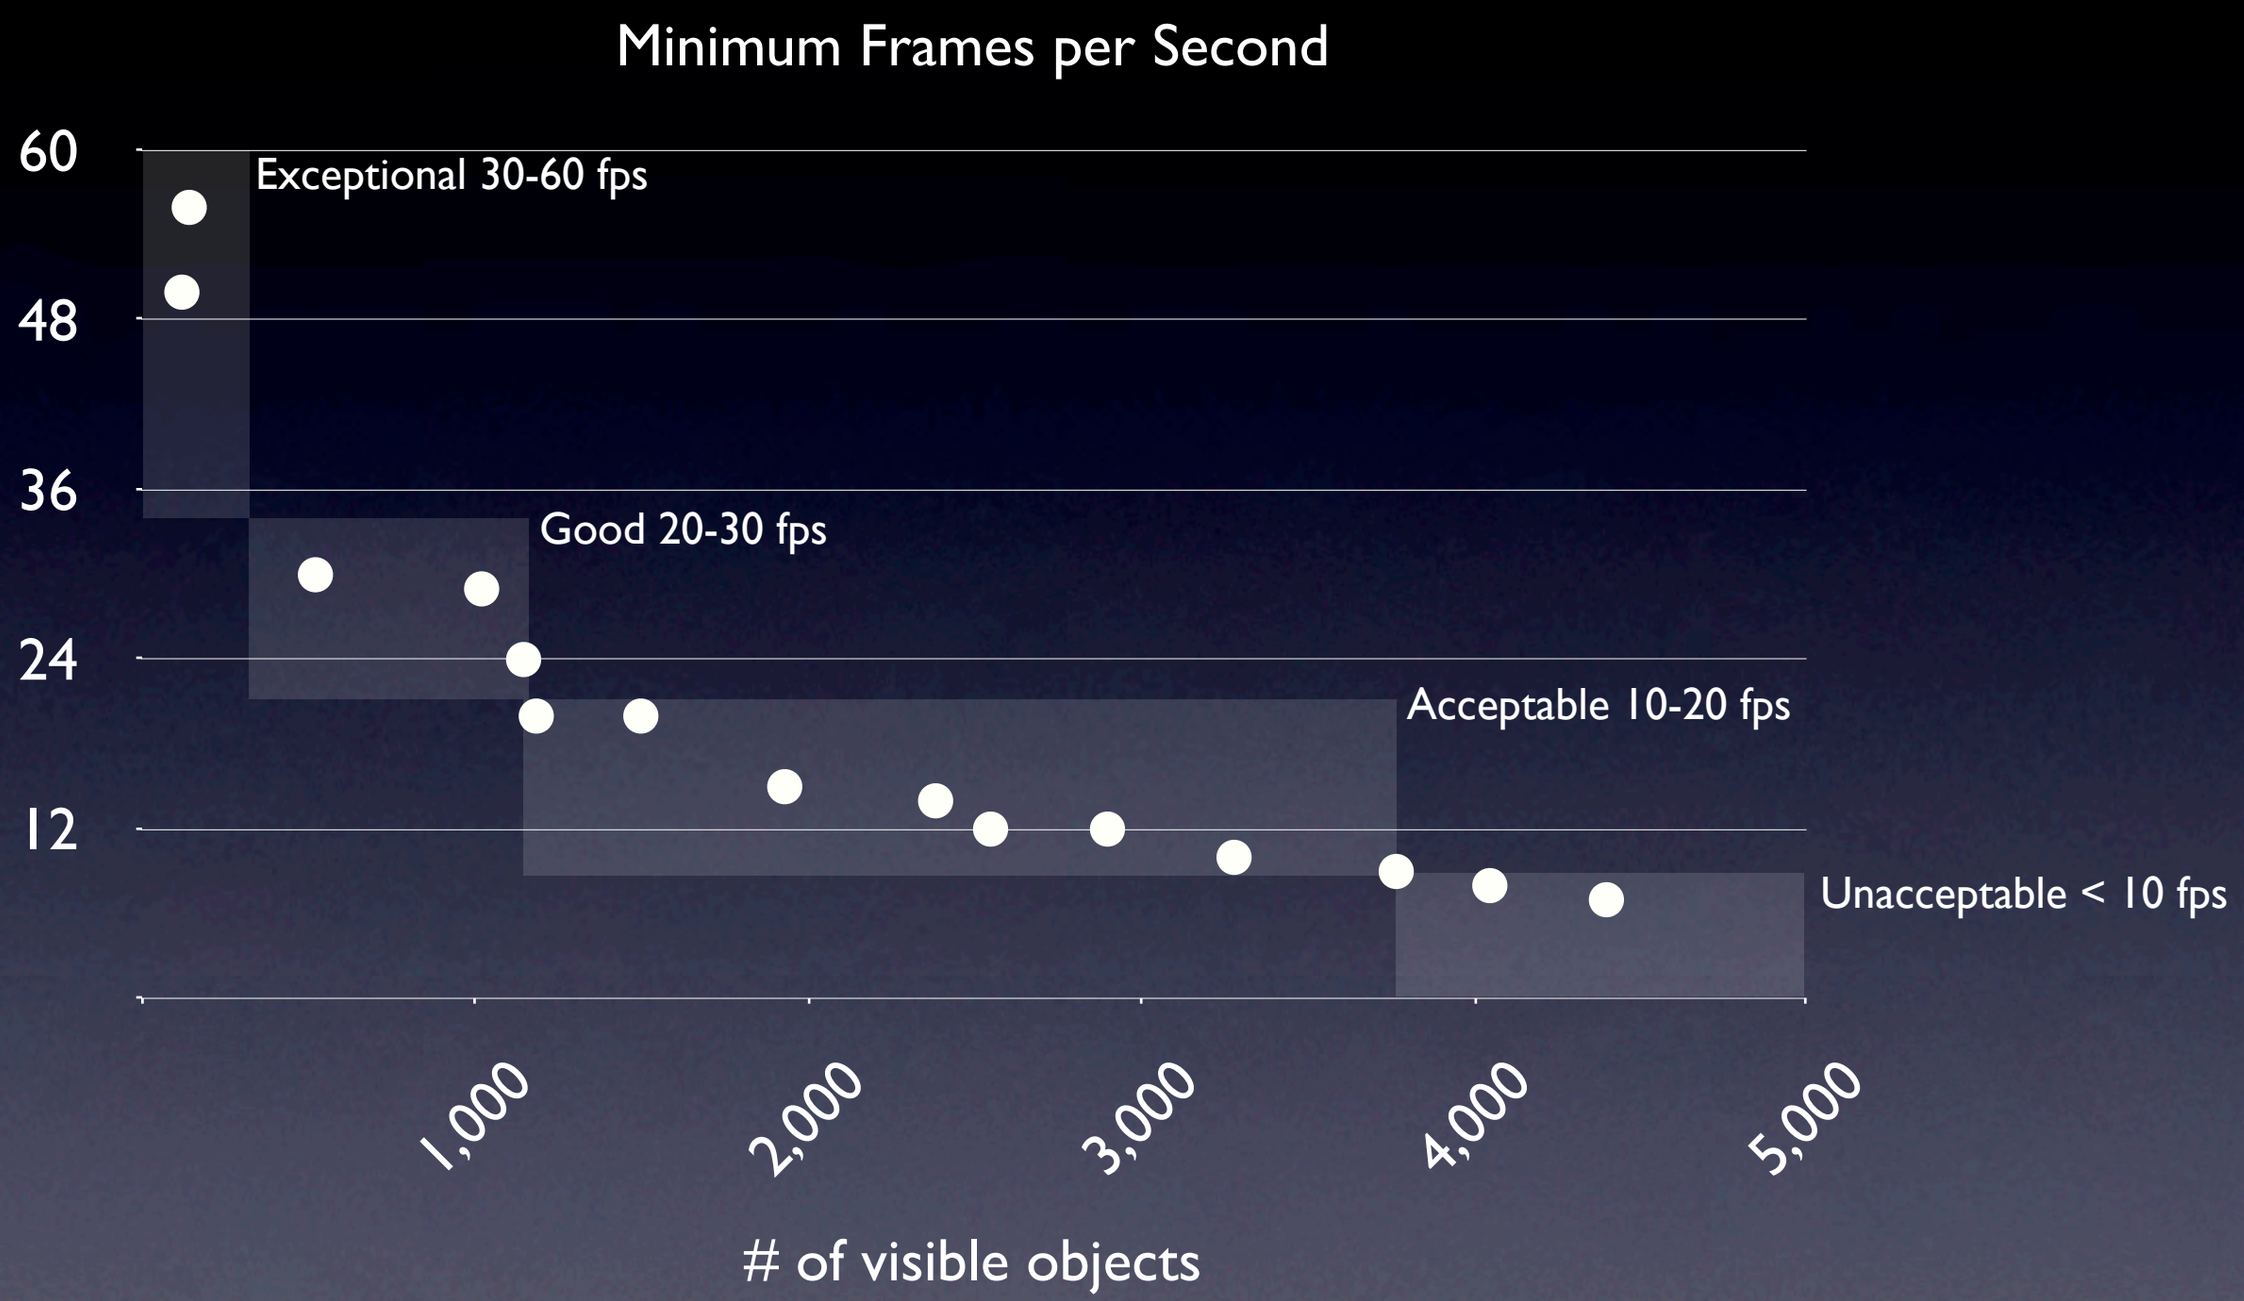
- Incoming & outgoing
- Hub & core-to-core traffic
- Widest possible port ranges
- Time-span wider than the activity (minutes to hours)
- Suspect IPs and ranges

Filter

- Superfluous port traffic (e.g. 80, 53, 25)
- IPs that are unrelated to the incident

© 2007 Sunny Fugate

Flow Viewer Performance



**Graphics performance on dual 1.5GHz SPARC SunFire v440 with Sun XVR 1200

© 2007 Sunny Fugate

Flow Viewer Performance

Real-time Performance	Real-time Records / Hour	Optimal playback rate
Optimal	10K-30K/hour	10X Real-time
Acceptable	40K-100K/hour	Real-time
Poor	100K-300K/hour	1/10 X Real-Time

Sparse data sets can be viewed quickly
e.g. months of data in minutes

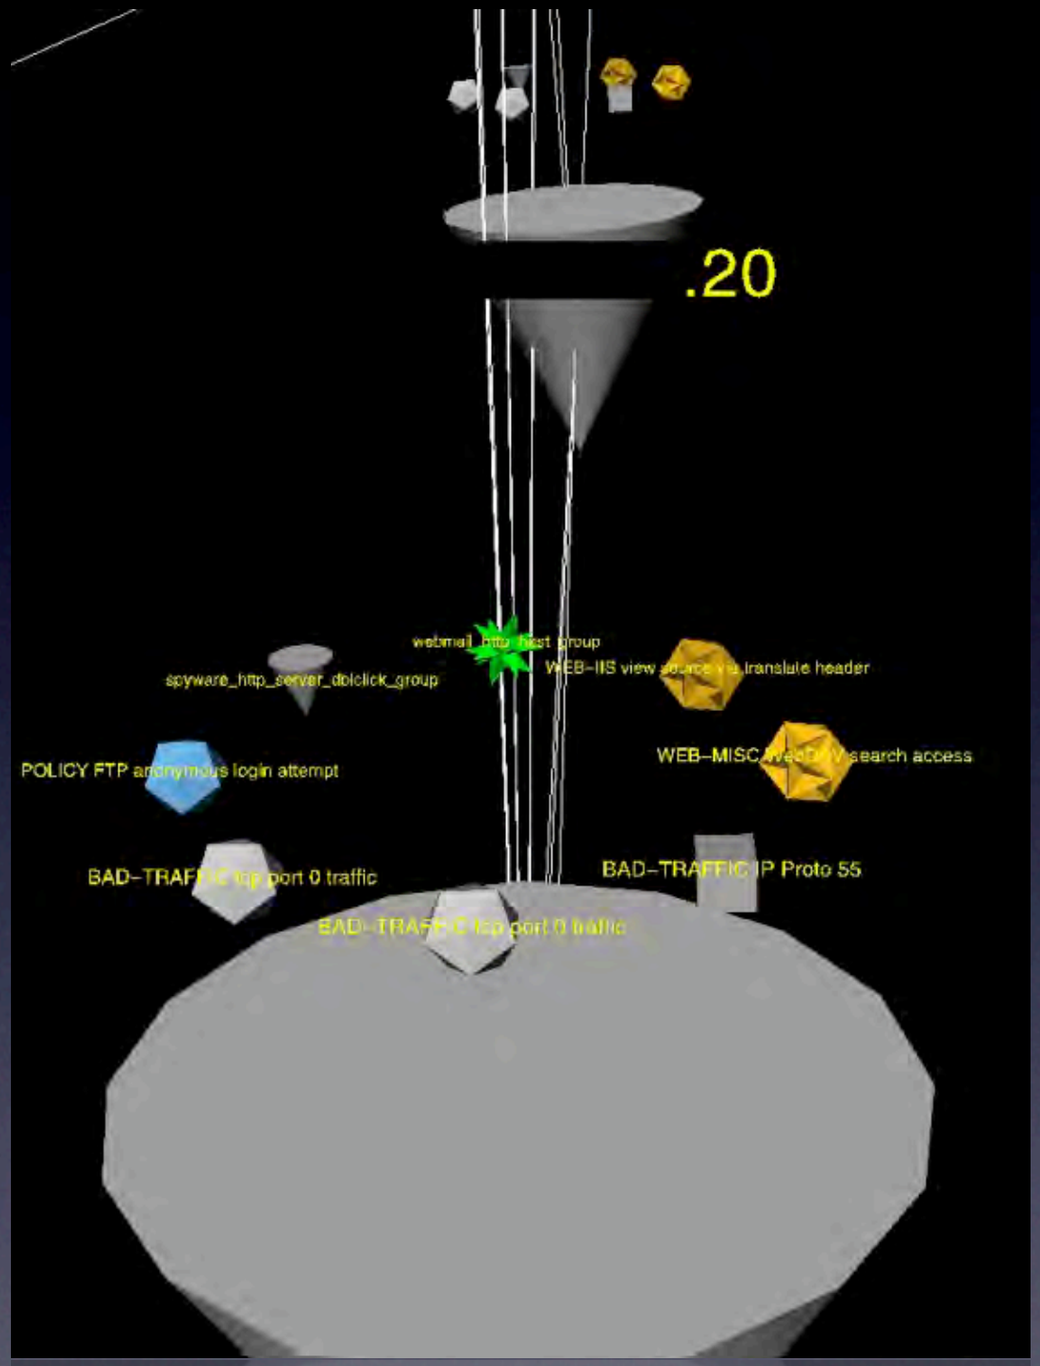
Dense data sets can be viewed slowly or filtered
e.g. seconds of data in minutes

© 2007 Sunny Fugate

Knowledge Depth vs Breadth

What trade-offs are we making?

- UI Feedback?
  - Haptic vs visual feedback
- Data access?
  - Random access vs linear access
- Training?
  - Under-learned vs over-learned
  - Tool complexity
- Meaning?
  - Visual semantic vs text
  - Intuitive/Iconic vs cryptic/coded



© 2007 Sunny Fugate

י

© 2007 Sunny Fugate

References

- [1] T. Abraham, Electronics, and Surveillance Research Laboratory (Australia). Information Technology Division. IDDM: Intrusion Detection Using Data Mining Techniques. DSTO Electronics and Surveillance Research Laboratory, 2001.
- [2] Mark Anderson, Dean Engelhardt, Damian Marriott, and Suneel Randhawa. Data processing and observation system, August 1 2006.
- [3] Mark Anderson, Dean Engelhardt, Damian Marriott, and Suneel Randhawa. Data view of a modelling system, April 11 2006.
- [4] H. H. Clark and W. G. Chase. On the process of comparing sentences against pictures. Cognitive Psychology, 3:472–517, 1972.
- [5] Herbert A. Colle and Gary B. Reid. The room effect: Metric spatial knowledge of local and separated regions. Presence: Teleoperators and Virtual Environments, 7(2):116–128, 1998.
- [6] Science Applications International Corporation. Intrusion Detection System System Protection Profile. National Security Agency, 9800 Savage Road, Fort Meade MD, 20755, version 1.4 edition, February 2002.
- [7] Stephen W. Draper and Donald A. Norman. User Centered System Design: New Perspectives on Human-computer Interaction. CRC, 1 edition, 1986.
- [8] D. Engelhardt and M. Anderson. A distributed multi-agent architecture for computer security situational awareness. Information Fusion, 2003. Proceedings of the Sixth International Conference of, 1, 2003.
- [9] Sunny Fugate. Visual language for tactical communication. In Proceedings of the First Annual Visual and Iconic Language Conference. SPAWAR Systems Center, San Diego, August 2007.
- [10] Sunny Fugate, Emily W. Medina, LorRaine Duffy, Dennis Magsombol, Omar Amezcua, Gary Rogers, and Marion Ceruti. Next-generation tactical-situation-assessment technology (tsat): Iconic language. In Sunny Fugate, editor, Proceedings of the First Annual Visual and Iconic Language Conference. SPAWAR Systems Center, August 2007.
- [11] David Gamon and Allen D. Bragdon. Brains That Work A Little Bit Differently: Recent Discoveries About Common Brain Diversities. Barnes and Noble, 2000.
- [12] James K. Hahn, Hesham Fouad, Larry Gritz, and Jong Won Lee. Integrating sounds and motions in virtual environments. Presence: Teleoperators and Virtual Environments, 7(1):67–77, 1998.
- [13] T. Munzner. INTERACTIVE VISUALIZATION OF LARGE GRAPHS AND NETWORKS. PhD thesis, STANFORD UNIVERSITY, 2000.
- [14] Jakob Nielsen. Usability Engineering (Interactive Technologies). Morgan Kaufmann, 1st edition, 1993.
- [15] CM Reed and NI Durlach. Short paper: Note on information transfer rates in human communication. Presence: Teleoperators and Virtual Environments, 7(5): 509–518, 1998.
- [16] Walter Shepherd. Shepherd's glossary of graphic signs and symbols;. Dent, London,, 1971.
- [17] Edward R. Tufte. Visual Explanations: Images and Quantities, Evidence and Narrative. Graphics Press, Cheshire, Conn., 1997.
- [18] TS TULLIS. An evaluation of alphanumeric, graphic, and color information displays. Human Factors, 23:541–550, 1981.
- [19] D.J. Ward, A.F. Blackwell, and D.J.C. MacKay. Dasher—a data entry interface using continuous gestures and language models. Proceedings of the 13th annual ACM symposium on User interface software and technology, pages 129–137, 2000.
- [20] G.J. Wills. Nicheworks-interactive visualization of very large graphs. Graph Drawing: 5th International Symposium, GD'97, Rome, Italy, September 18-20, 1997. Proceedings, 1997.

Images

- Jeff Han's Multi-Touch Screen Interface, Jeff Kubina, Flickr.com, license: http://creativecommons.org/licenses/by-sa/2.0/deed.en
- Atari joystick, duncan, Flickr.com, license: http://creativecommons.org/licenses/by-nc/2.0/deed.en
- Headphones, daxtoor, Flickr.com, license: http://creativecommons.org/licenses/by-sa/2.0/deed.en

© 2007 Sunny Fugate

SPAWAR
Systems Center
San Diego

Next Generation Tactical Situation Assessment Technology (NG-TSAT)

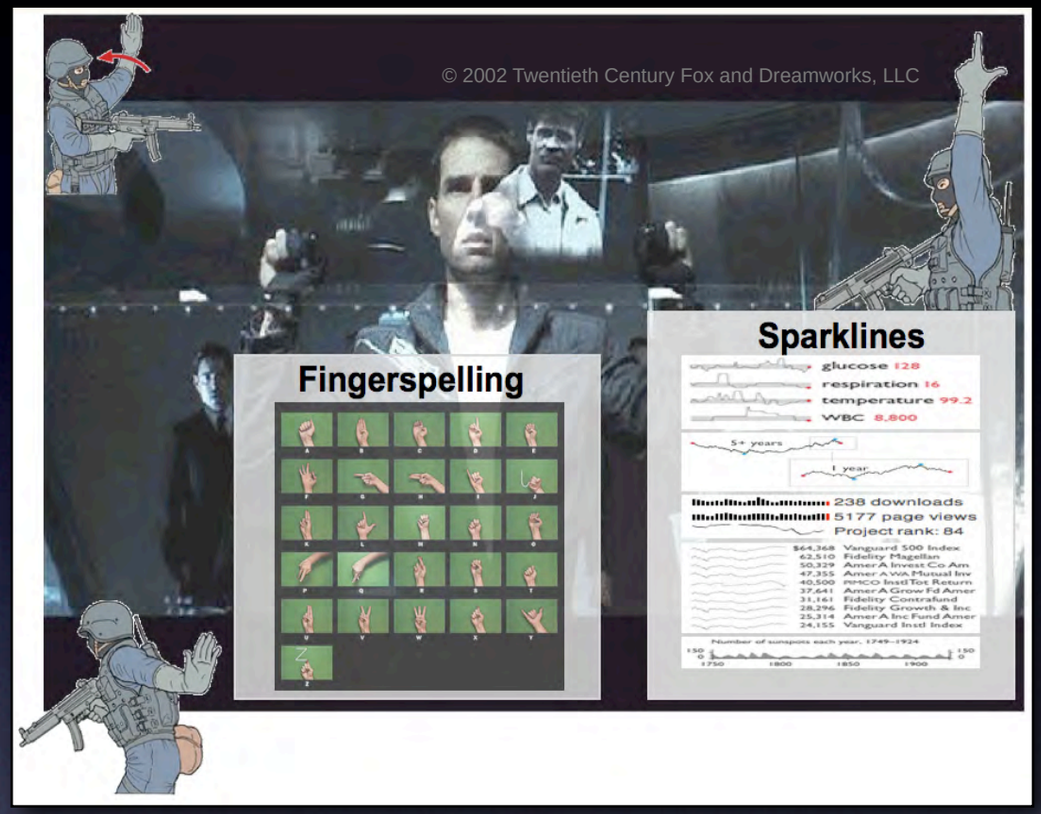
Objective: Next-generation Tactical Chat. Icon-based situation assessment (SA)language supported by wireless gesture-recognition gloves used in hostile or noisy (silence-mandated) environments

Description of Effort:

1.Linguistic Analysis: Analysis of current C² chat logs to determine speech patterns and repetitive SA concepts/themes

2. Iconic Language Development: Output of linguistic analysis determines candidate icons representing most prevalent SA "themes;" development of prototype C² iconic SA language

3. Wireless, Gesture-Recognition Gloves: Develop wireless gloves that recognize C² icons/gestures which can transmit across network to distributed warfighters (replacing keyboard input when in MOPP)



Benefits of TSAT:
Compressed Chat (25% ↓content; 50% ↓reduction in production time) for rapid SA dissemination.
Gesture-recognition in very noisy, distributed ops, or in very austere environments (e.g., the moon)

Challenges:
1. No current method or theory for chat-meaning compression; currently done in prose; computer linguistic analysis of unstructured text still neoteric.
2. Wireless gesture recognition glove technology still in infant stages of development; focused on commercial animation support, not on disciplined language support

TRL: Chat: TRL 1-2; Gesture-recognition: TRL 1-4

Major Milestones FY06:
Linguistic analysis discovery of common C² SA themes
Development of icon/symbols for candidate SA themes
Development of proof-of-concept wireless gesture-recognition glove

Period of Performance: 2007-2012

PI contact info: Dr. LorRaine Duffy, (619) 553-9222, LorRaine.Duffy@navy.mil, SSC San Diego, CA

© SSC SD LorRaine Duffy 03042007

Synaesthesia

Synaesthesia: "a neurological condition in which two or more senses are coupled."

"loud color" "sharp laugh" "bitter wind"

grapheme color synesthesia - letters or numbers are perceived as inherently colored

© 2007 Sunny Fugate

How many numbers contain the digit 6?

9910 9972 3292 7602 82 9054
5636 2710 1944 6330 6560 8101
5177 1955 7029 4083 4643 5710
4935 2256 1495 1025 8375 8518
80 797 2610 3008 8784 1854 2383
9728 4523 573 5914 7975 281
6664 2682 7689 7753 273 5597
799 9960 1437 4534 8601 4563
6734 647 9409 6543 4827 2398
1532

© 2007 Sunny Fugate

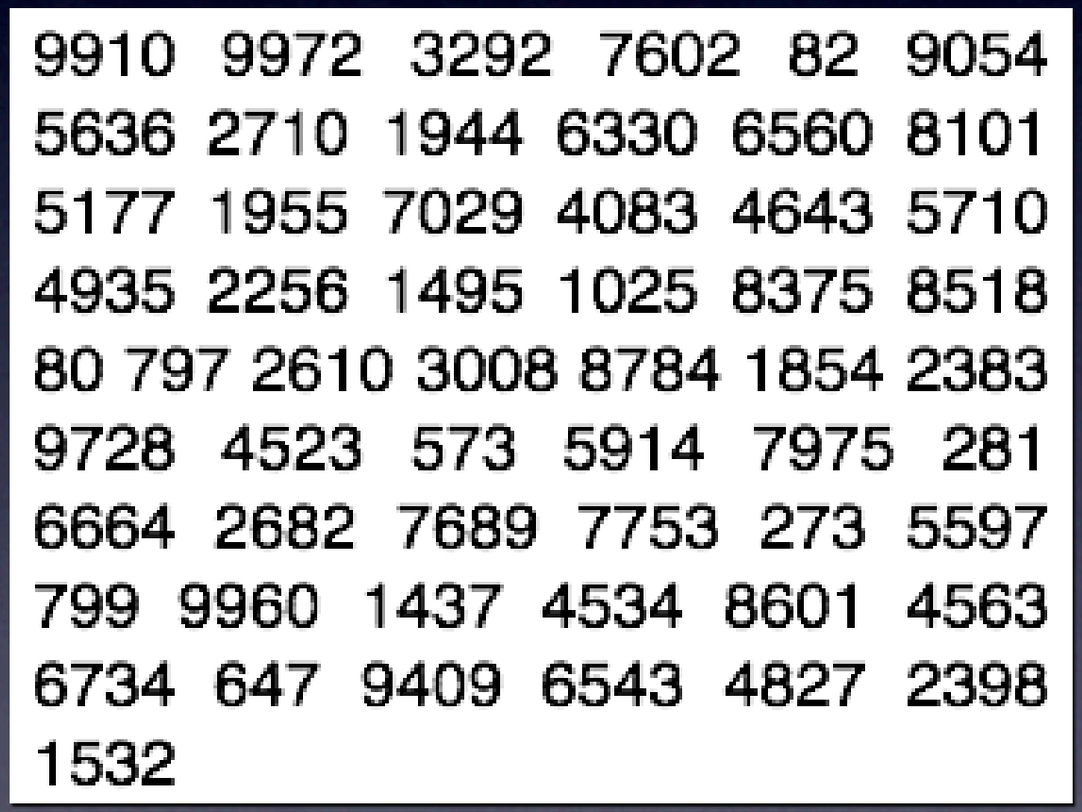
Is this easier?

9910 9972 3292 7602 82 9054 5636
2710 1944 6330 6560 8101 5177
1955 7029 4083 4643 5710 4935
2256 1495 1025 8375 8518 80 797
2610 3008 8784 1854 2383 9728
4523 573 5914 7975 281 6664 2682
7689 7753 273 5597 799 9960 1437
4534 8601 4563 6734 647 9409
6543 4827 2398 1532

© 2007 Sunny Fugate

Emulating Synaesthesia

These methods can be used achieve **sequence disambiguation** and





© 2007 Sunny Fugate

Emulating Synaesthesia

192.168.1.232
129.168.1.233

192.168.1.232
129.168.1.233

192.168.1.232
129.168.1.233

© 2007 Sunny Fugate

The *Ripple* decoded

Carrie Gates

CA Labs

John McHugh

Canada Research Chair in Privacy and Security

Dalhousie University

mchugh@cs.dal.ca

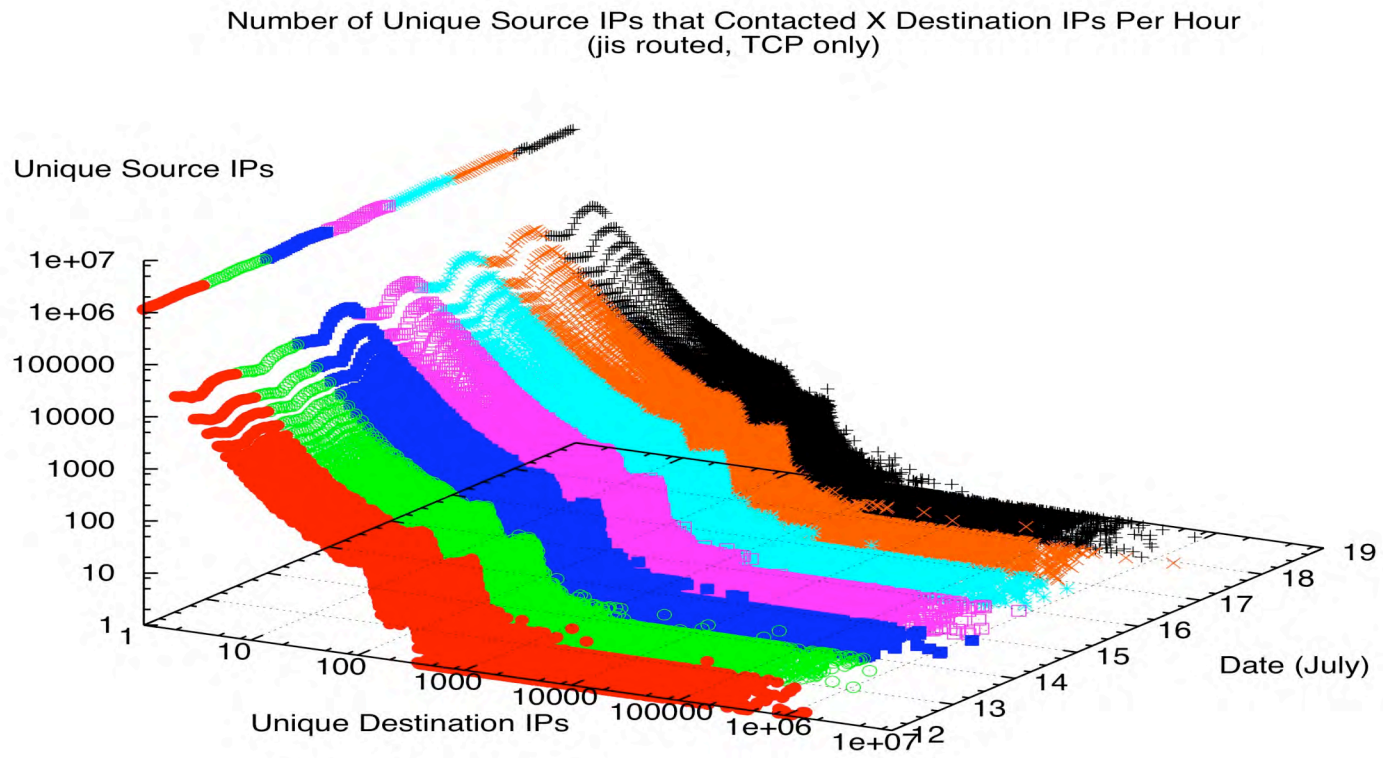
© 2007 by John McHugh, Carrie Gates

Very large scale observation

- Carrie Gates was interested in the degree of fan out from outside to inside for her scan detection work.
- How many outside hosts use exactly one inside host / service pair. (unique destination address/port)
- In the beginning, we did it the hard way, but Bloom filters can be used to find unique sIP,dIP,dport exemplar flows
- If we make a source IP bag from the exemplar flows, the counts will be the number of different host / service pairs contacted by a given source host.
- Invert the bag to determine how many entries have a count of 1, 2, 3, ... Plot hourly results for a week

© 2005,6,7 by John McHugh, Carrie Gates

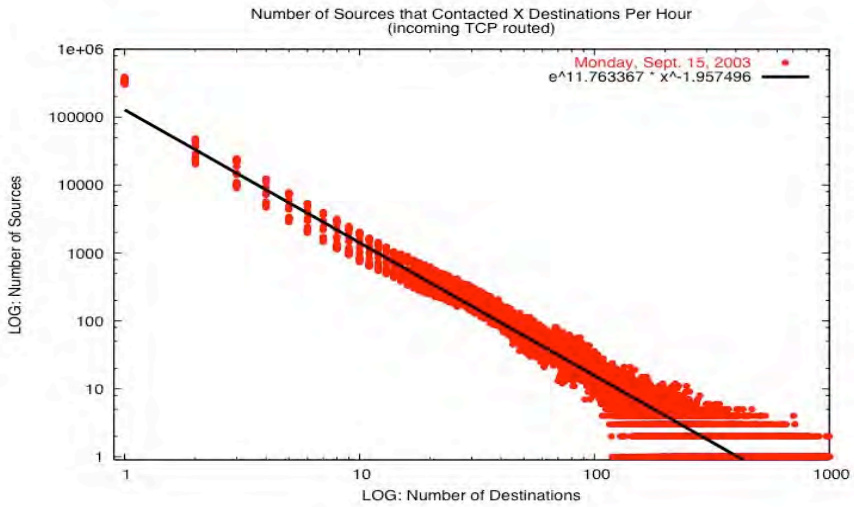
Outside to inside - July 2003



© 2005,6,7 by John McHugh, Carrie Gates

Developing the contact surface

- In the absence of the disturbance seen on the previous page, contact lines seem to follow a power law type of distribution
  - or do they[1].
  - We think this is really at least 3 separate processes
    - VLF noise
    - "normal activity"
    - Bulk scanning



[1] everything is a straight line on log/log paper, especially if you use a fat marker

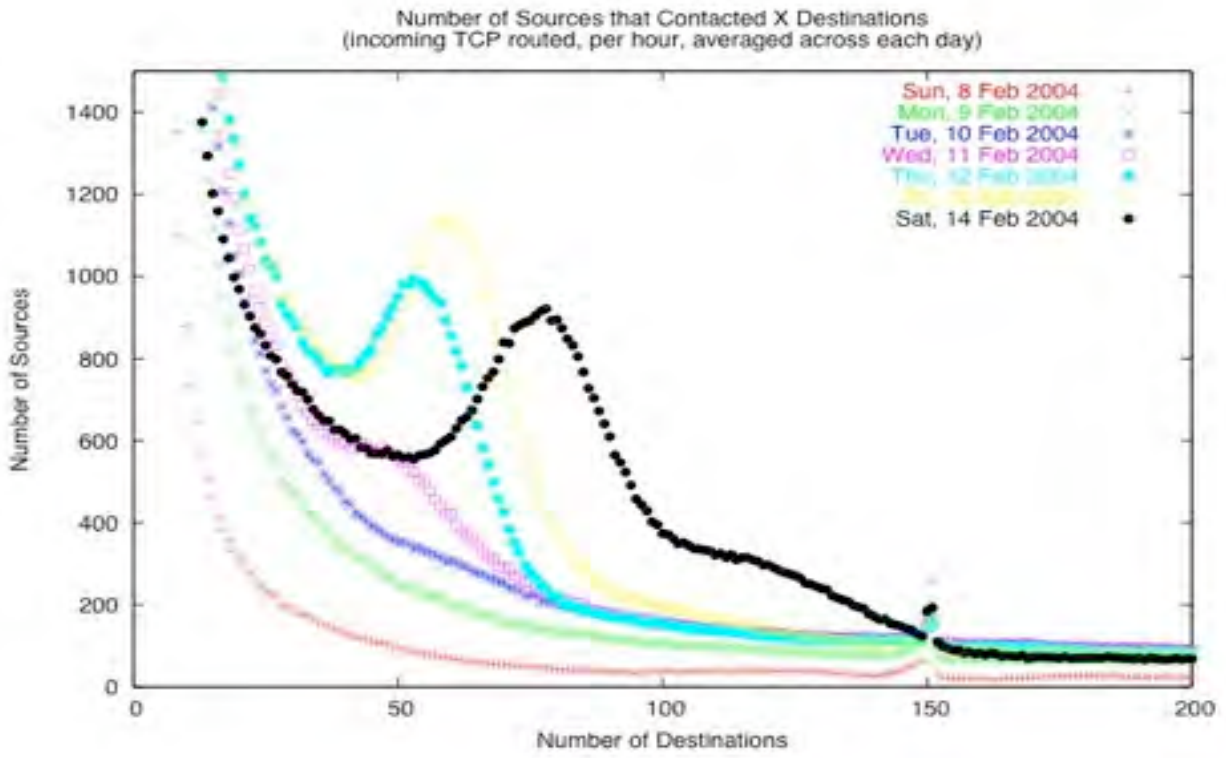
© 2005,6,7 by John McHugh, Carrie Gates

Internet wide disturbance

- The ripple in what would otherwise be a fairly straight log/log plot of connectivity was observed from at least Jan - Aug 2003.
- It went away when Blaster appeared in Aug 2003.
- A similar ripple existed from Feb 11 to May 31 2004 coinciding with the lifetime of Welchia-B
  - In this case, the ripple is due to a few hundred machines scanning at a low, fixed, rate induced by a loop with a “sleep” system call.
- In both cases, they persisted until killed, not patched.
- We have been told that the ripple is back.

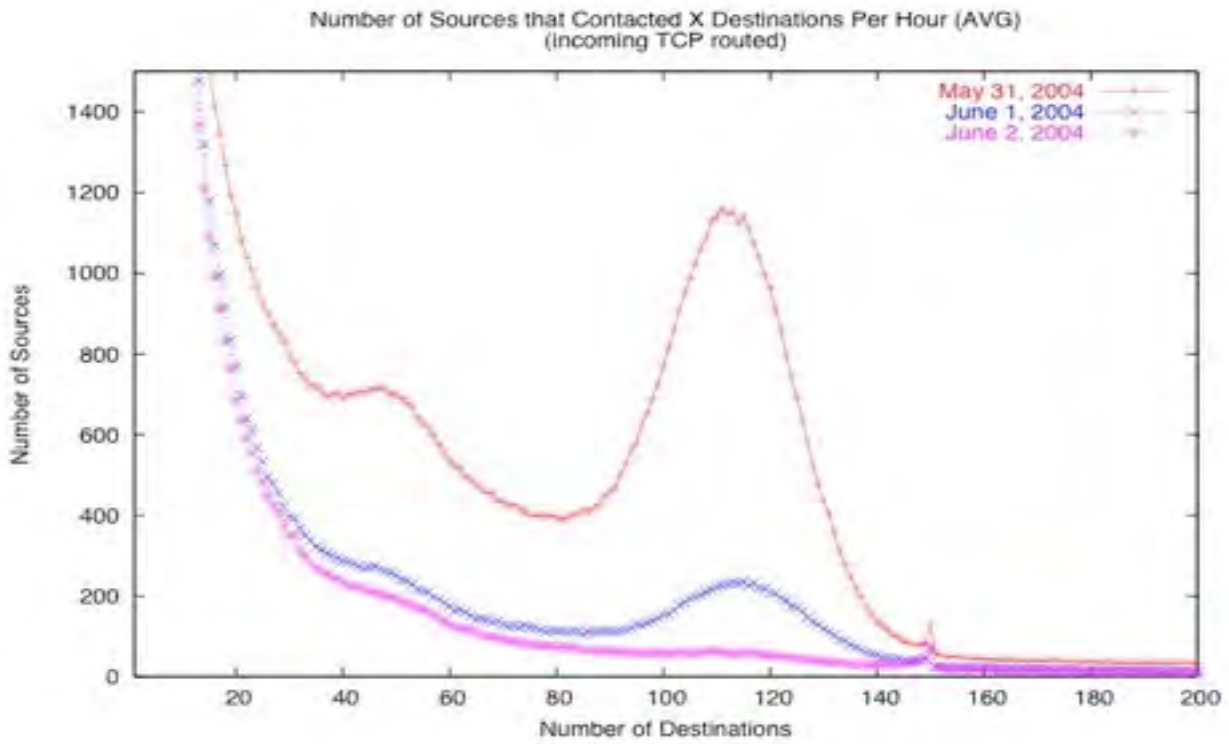
© 2005,6,7 by John McHugh, Carrie Gates

Details of the Welchia.B event - onset



© 2005,6,7 by John McHugh, Carrie Gates

Details of Welchia.B - demise



© 2005,6,7 by John McHugh, Carrie Gates

Design Time Coordination

- The sleep in the scan loop of Welchia.B points to a form of loose, design time, coordination.
- All members of the cohort scan at approximately the same rate, using the same random generation scheme but with a different random seed.
- If we captured all the scans from each member of the cohort, we would expect to see a small, tight, cluster of scanners all contacting nearly the same number of targets.
- We observe only a small portion of the address space and see a small percentage of the scans from each host with substantial interhost variation.

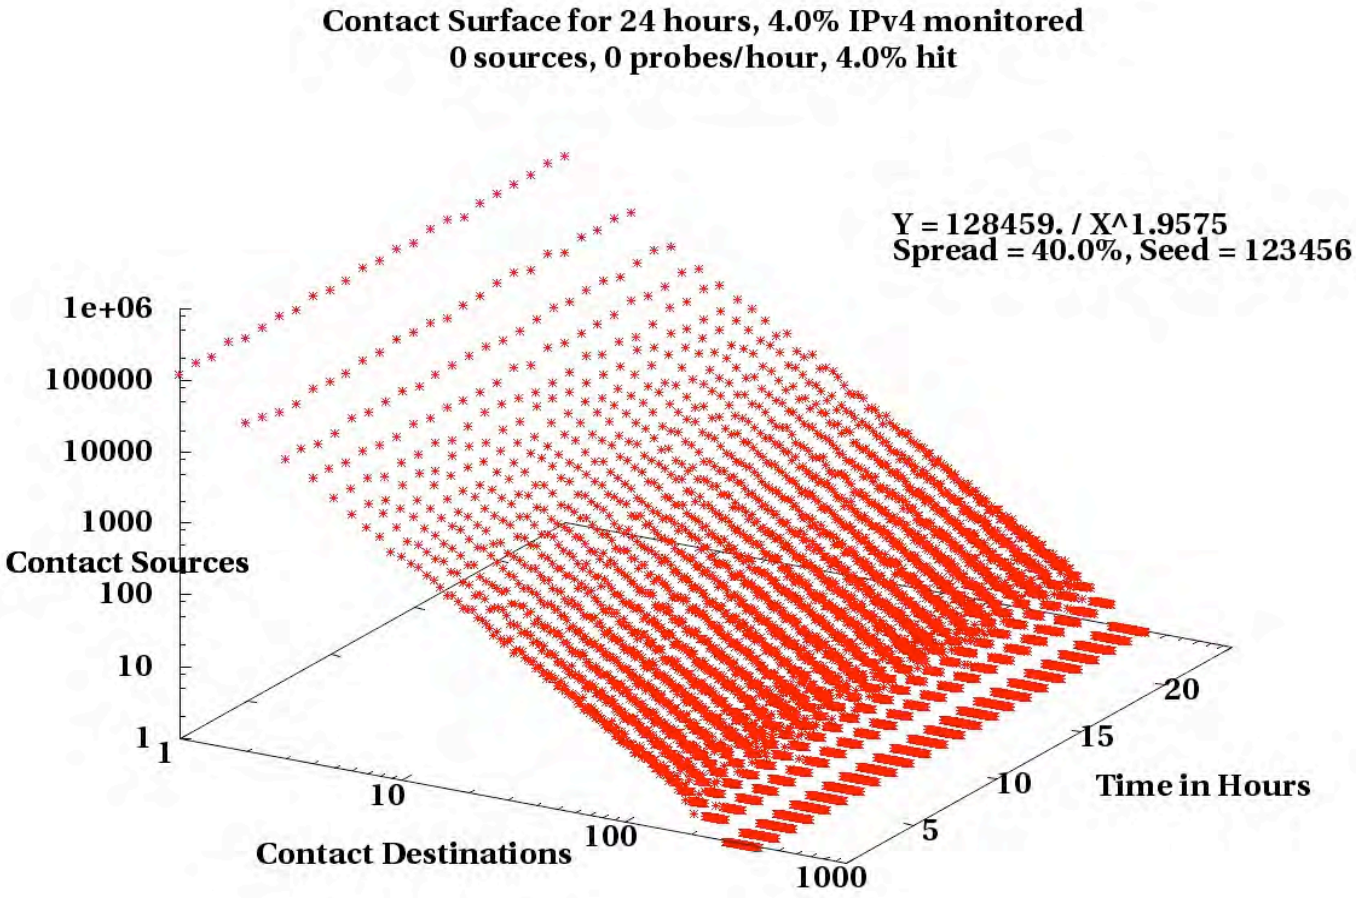
© 2005,6,7 by John McHugh, Carrie Gates

This fall, we simulated the perturbations

- Generated approximation of unperturbed background
  - Don’t care about process, only appearance
- Simulated perturbation process parameterized on:
  - Number of sources
  - Probe rate / source
  - % of IPv4 monitored
  - % of probes intercepted
    - For ripple or wave, % monitored = % intercepted
    - For scans targeting monitored network they are different
- Looked at observability as a function of parameters.

© 2005,6,7 by John McHugh, Carrie Gates

Background only - main line process



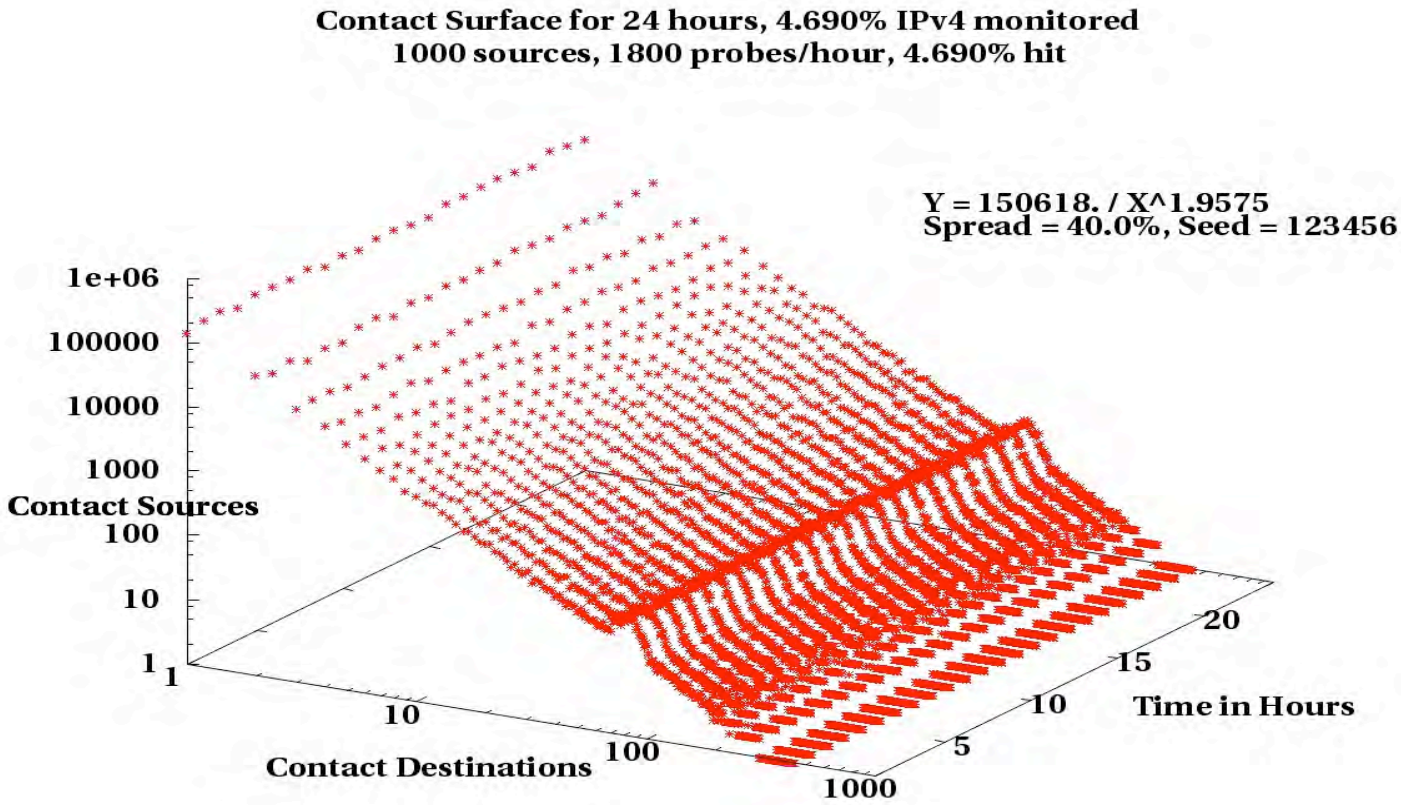
Simulating the ripple

- For each source, S_i; for each probe, j emitted during an observation period;

  we generate a random $R_{i,j}$ in {0..1.0}.
- If $R_{i,j}$ is < the % of IPv4 monitored, it is a hit.
- Use the hit count to select the appropriate cell in the background traffic contact line and add 1 to it.
  - source S_i hit that number of destinations during the simulated observation period period.
- Plot the modified contact line in either 2D or as part of a 3D contact surface.

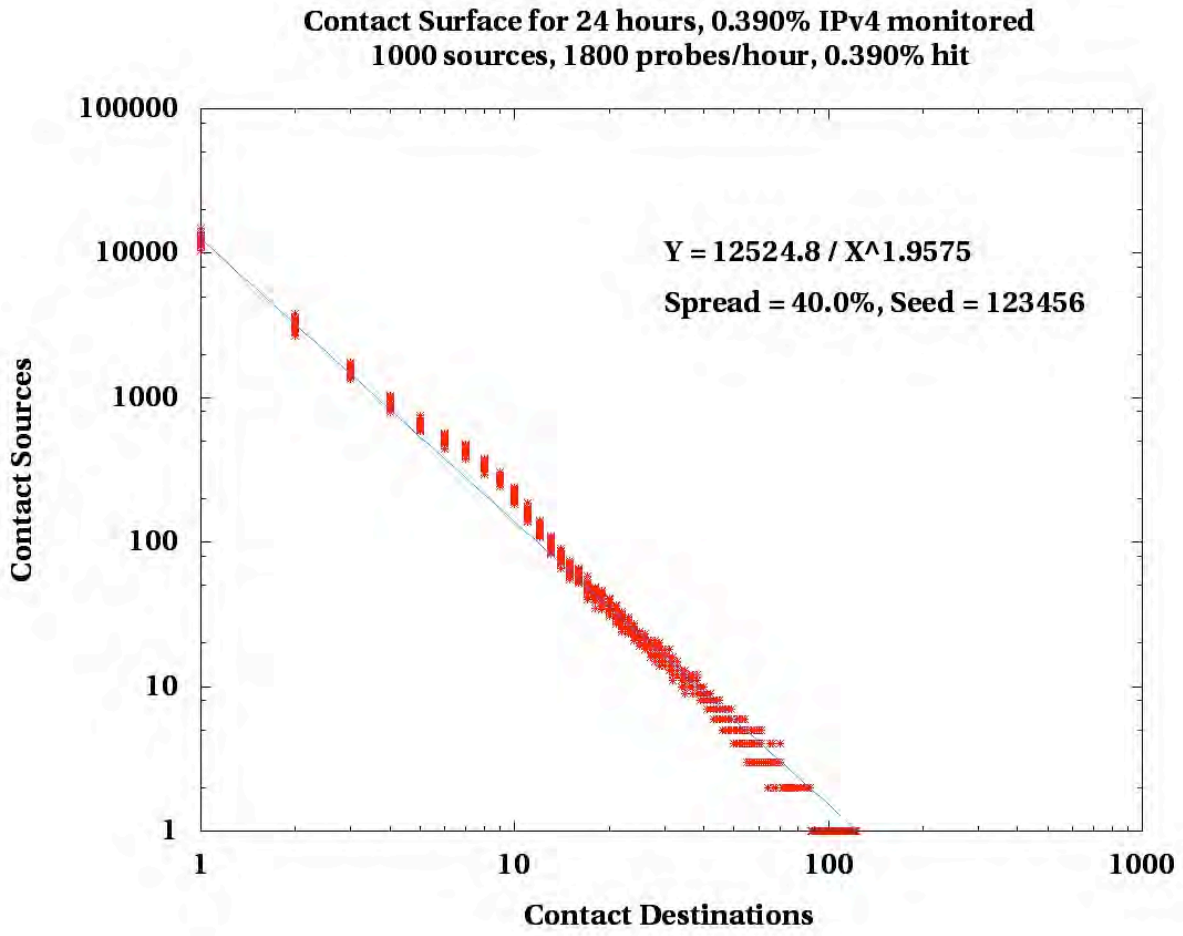
© 2005,6,7 by John McHugh, Carrie Gates

A plausible ripple



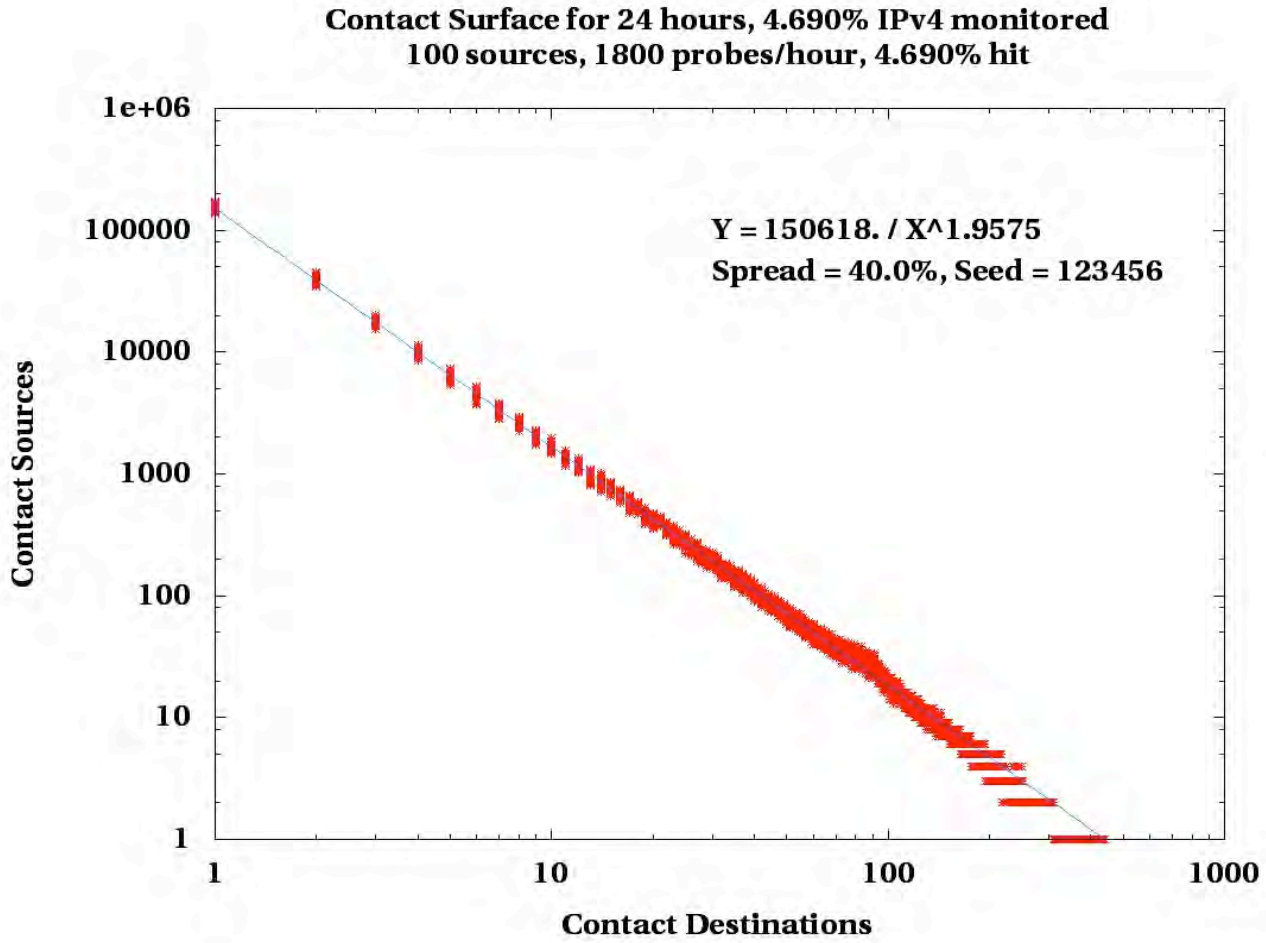
© 2005,6,7 by John McHugh, Carrie Gates

Observability: 1000 probers /16 coverage



© 2005,6,7 by John McHugh, Carrie Gates

Observability: 100 probers 12 X /8 cover



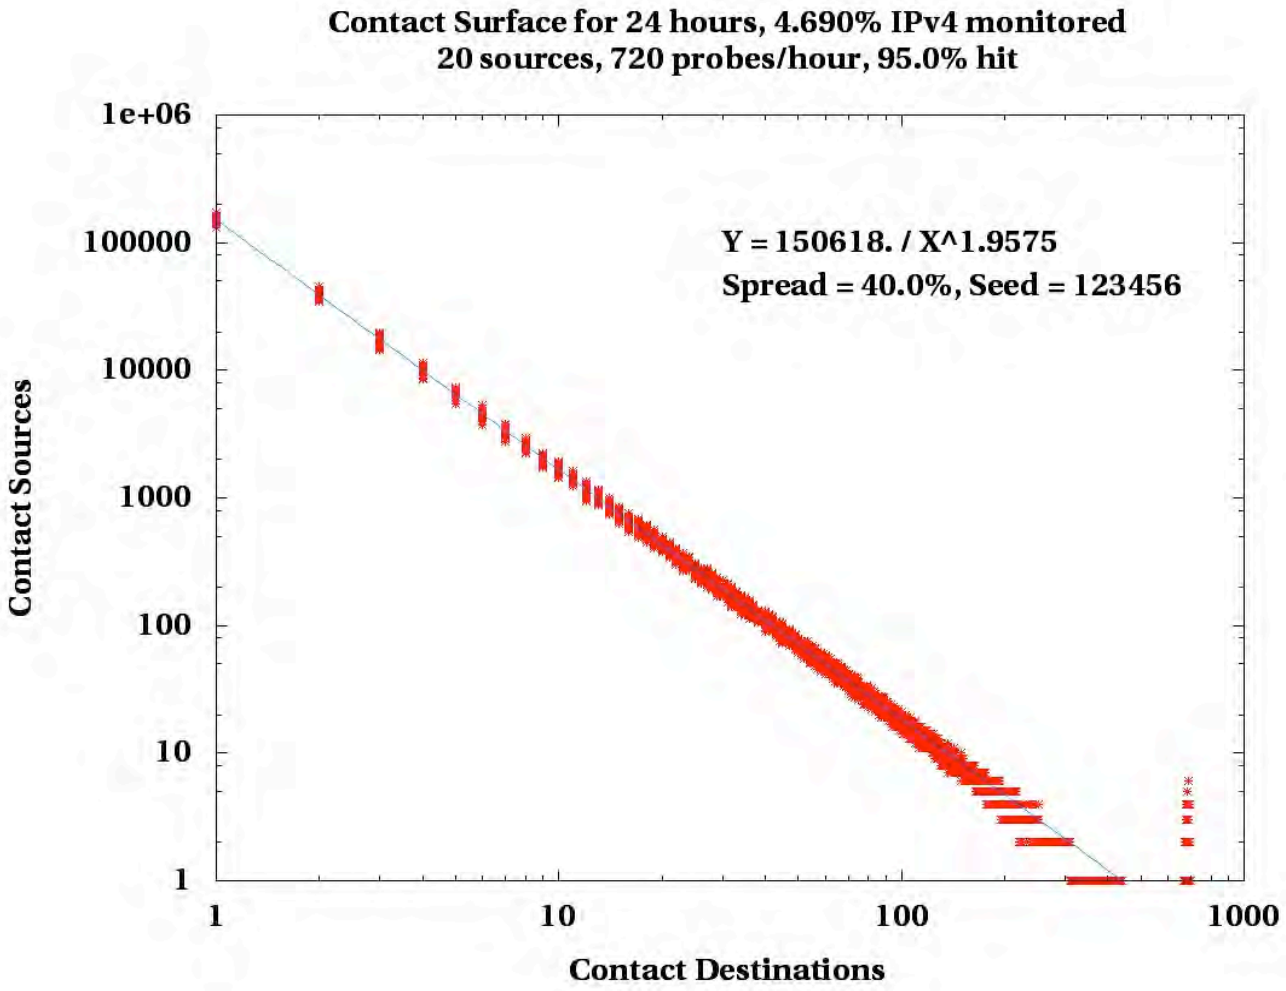
© 2005,6,7 by John McHugh, Carrie Gates

Simulated and real spikes.

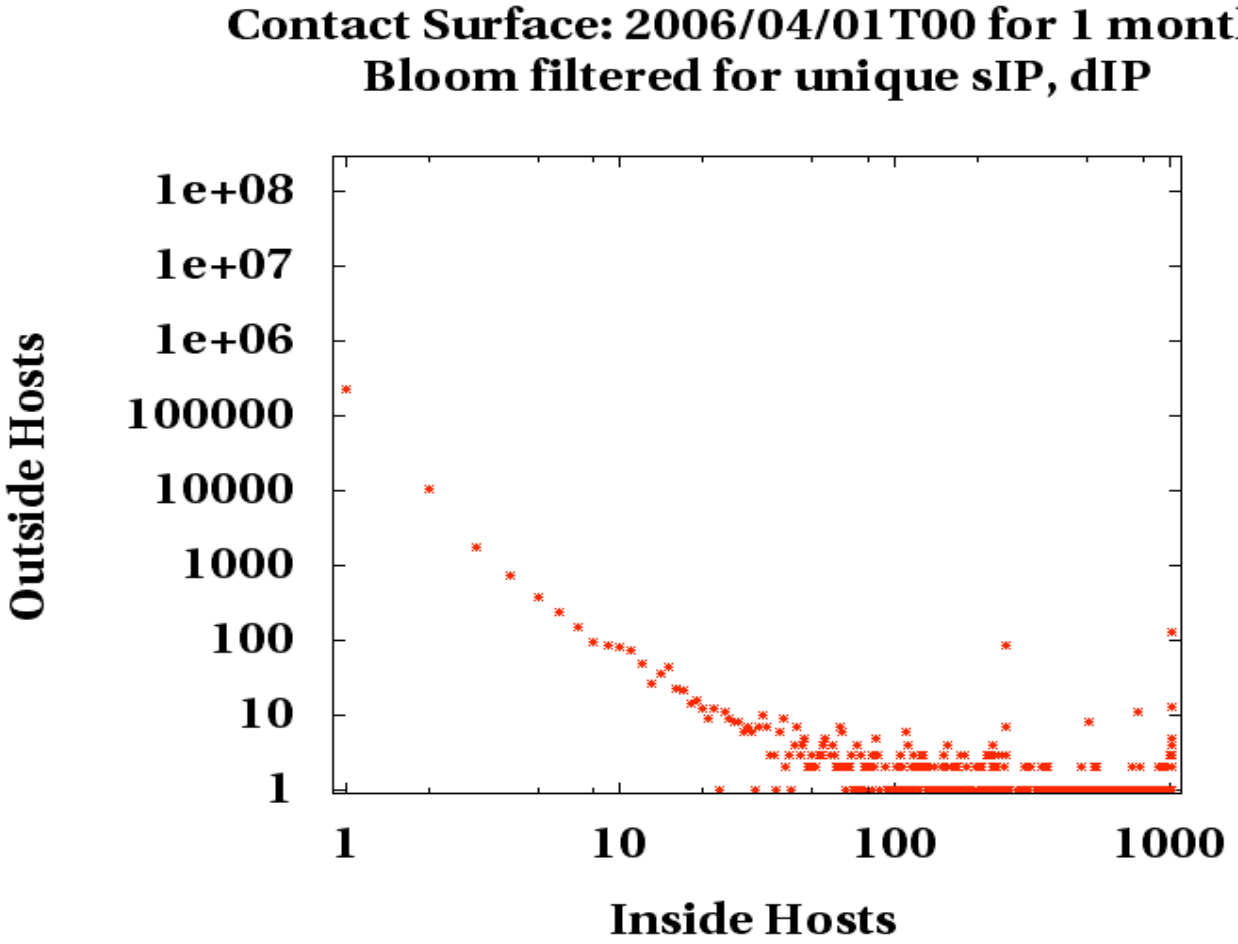
- The spikes appear when the percentage of intercepted probes is high.
  - Occurs when the probes fall mostly, 95%+, in the monitored address space.
  - At 100%, the spike becomes a point
- First, we simulate the spike.
- Next is a one month contact line for our /22, based on Bloom filtering for unique sIP, dIP pairs.
  - Note points at 254, 508, 762 and 1016 addresses.
- Then we will look at a movie for 14 months on the /22

© 2005,6,7 by John McHugh, Carrie Gates

The spike in the Welchia.B displays



Contact line for April 2006 for a /22



© 2005,6,7 by John McHugh, Carrie Gates

Future work

- We would like to visit or revisit the data for current and past perturbations.
- Develop analytical techniques for identifying cohorts of players exhibiting arbitrary, but similar characteristics.
- Explore other regions of the contact surface
- Link visualization to source / cohort identification in the visualization tool we are developing for DHS.
- and always remember ...

© 2005,6,7 by John McHugh, Carrie Gates

Greetings from Canada



© 2005,6,7 by John McHugh, Carrie Gates

Improvement of Processes for Flow Information

NTT Network Service System Laboratories, NTT Corporation

Hitoshi Irino,
Masaru Katayama

NTT Network System Laboratories

Abstract of this presentation

- Ideas for increasing (optimizing) performances of processes in IPFIX
- Ideas based on all processes using **an order rule of Information Elements/fields**
- These ideas are introduced:
  - Method for **reducing the number of comparisons** between an existing flow and an incoming new packet **in Metering Processes** (MPs) (**Comparison method for multiple fields in MPs**)
  - Method for **reducing the number of copies** of flow records from Metering Process to **Exporting Processes** (EPs) with a predefined order of fields (**Copy method for multiple fields in EPs**)
  - Method for **increasing processing speed for storing** data in incoming packets to file with a predefined format of **Collecting Processes** (CPs) (**Copy method for multiple fields in CPs**)

These are basically the same.

Motivation of this research

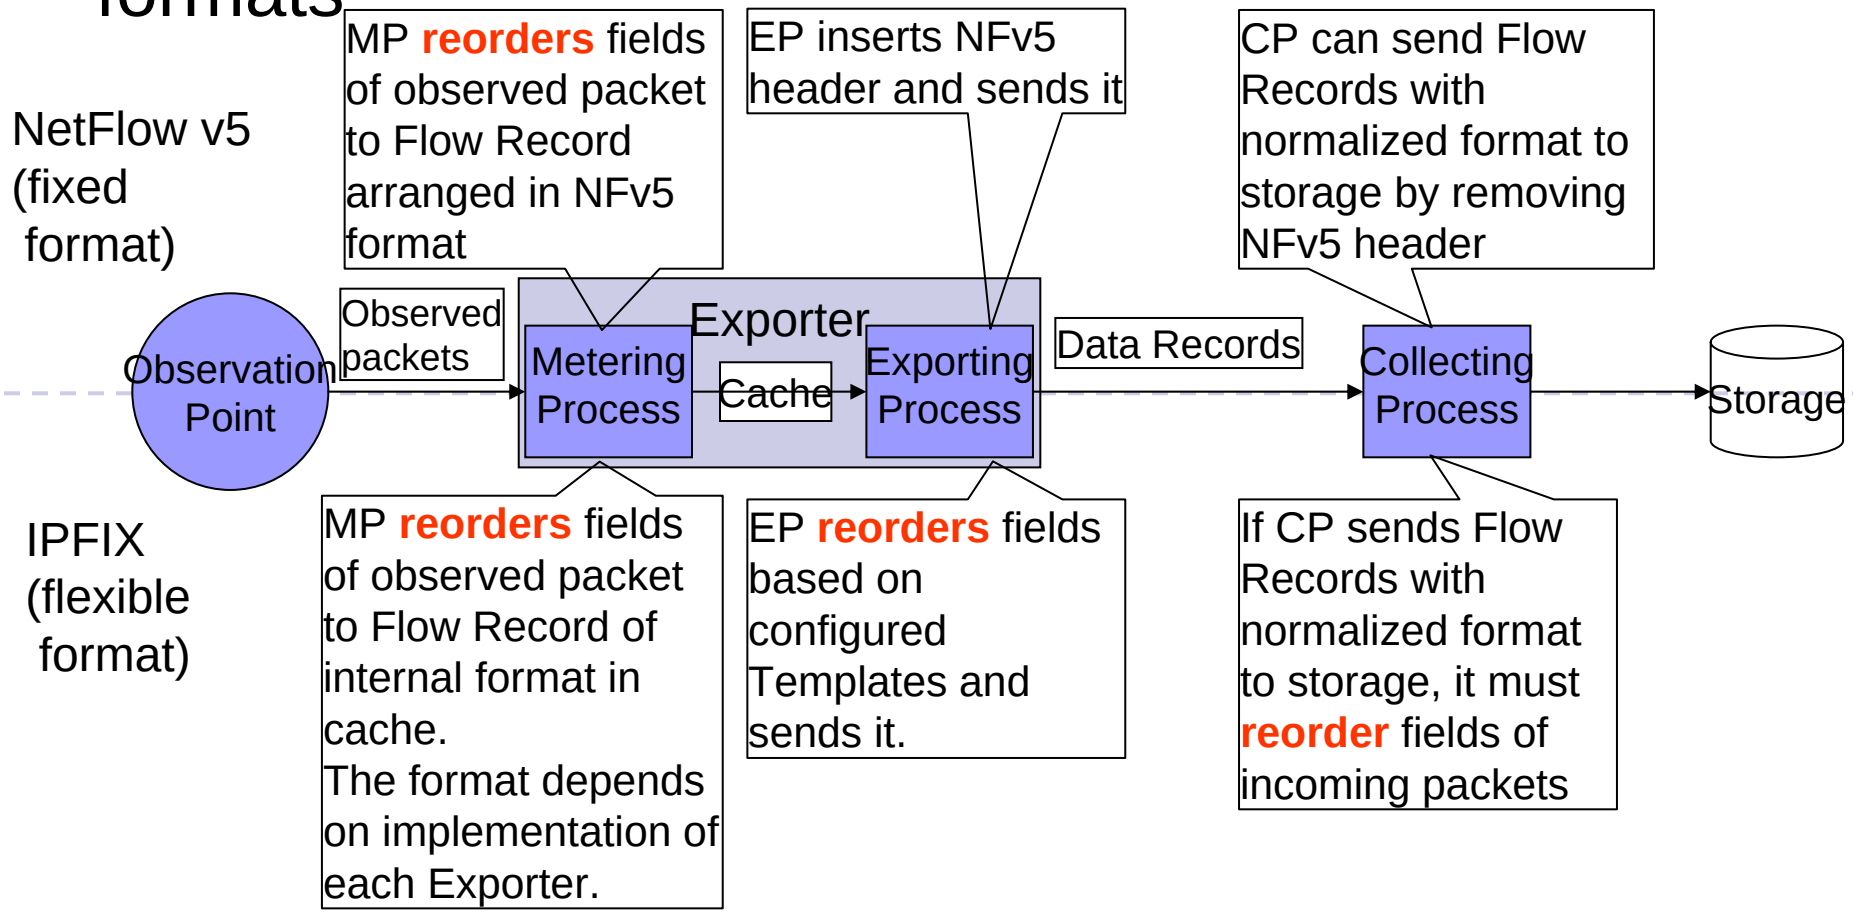
- Background
  - Network bandwidth will continue to increase.
  - IPFIX will be a standard protocol for flow information exchange.
    - Network bandwidth will become broader-band.
      - Use a lower sampling rate.
      - Use fewer Flow Keys.

→ However, flow information will become less accurate.

→ Research on increasing (optimizing) the performances of IPFIX processes

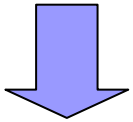
IPFIX features

- IPFIX
  - Advantage: Uses Template-based flexible flow export
  - Disadvantage: More complex than fixed-format protocol
- Comparison of processes between flexible and fixed formats

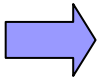


Our approach: **Making the order rule for Information Elements**

- Processes of IPFIX have a high possibility of reordering fields.
  - Reducing the cost of reordering fields can improve their performance.



- Our approach
  - Make the order rule for Information Elements
    - Order rule gives IPFIX processes chances to process multiple fields.
    - Processing multiple fields at a time achieves higher performance than processing one field at a time.
    - The rule does not influence the flexibility of IPFIX.



If a unified order rule of fields/IEs is defined, reordering costs can be reduced.

Idea of order

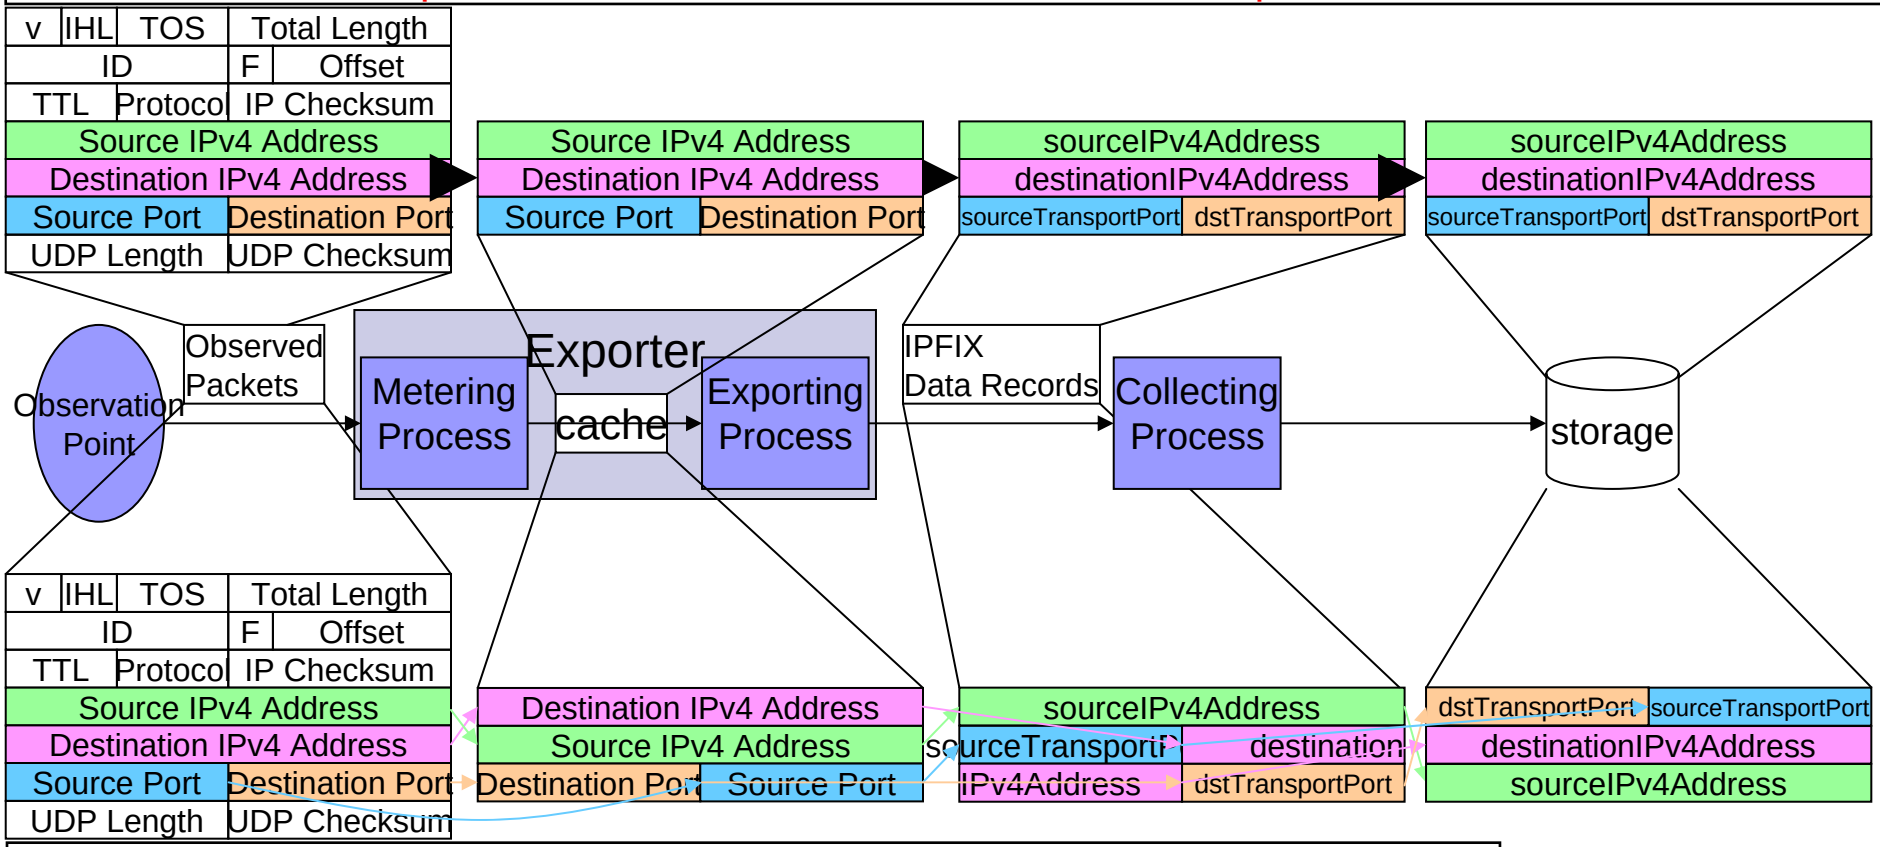
- Idea of order:
  - MPs, EPs and CPs place fields (IEs) in the same order, so it is highly likely that multiple fields will be processed at a time.
    - This reduces reordering costs.
- Order recommended in this presentation
  - Place fields in observed packets in order of protocol header.
  - **Therefore, order of IEs that refer to packets and header fields is recommended.**

	Metering Processes	Exporting Processes	Collecting Processes
Input	Observed packets (network byte order)	Their caches	IPFIX Data Record (network byte order)
Output	(Storing) their caches	IPFIX Data Record (network byte order)	(Storing) files, their DB (real-time analysis)

Example of using same order in MP, EP and CP

Flow Keys: sourceIPv4Address, destinationIPv4Address, sourceTransportPort, destinationTransportPort

Good (ideal) case: Same suggested order,
which refers order of packet header fields used in the cache in Exporter and IPFIX data records



Bad case: Different order used in the cache in Exporter and IPFIX data records

- **If the referential order, which refers to the order of packet fields, is defined, it could, in some cases, lead to increased performance.**
- **If a referential order is undefined, there is no possibility of increased performance.**

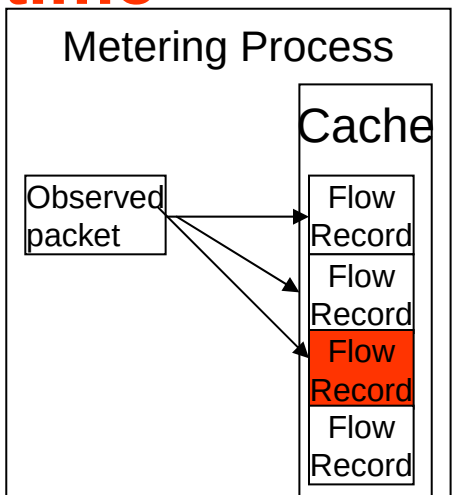
1st idea to improve performance
in environment in which MP, EP, and CP use the same order

Comparison method for multiple fields in Metering Processes (MPs)

NTT Network Service System Laboratories, NTT Corporation

Comparison method for multiple fields in MP (1)

- MP must repeat comparison between existing Flow Records in its cache and new observed packet.
  - To judge whether the new packet belongs to a new flow or an existing one.
- Basically, in this comparison, all fields (IEs) serving as Flow Keys are compared every time.
- **If fields of Flow Records are placed in the same order as packet header fields, MP can compare multiple fields at a time**

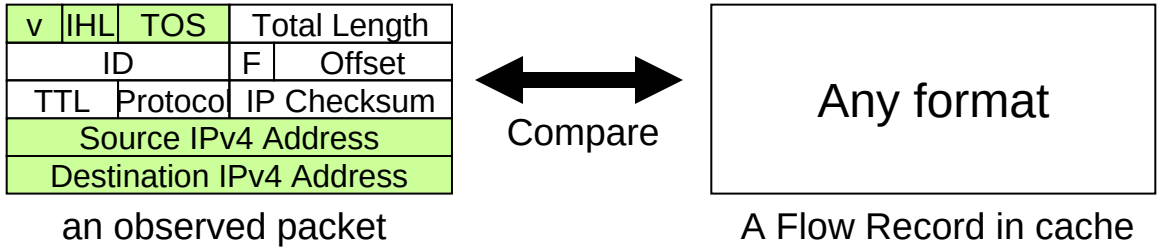


MP repeats comparisons and finds a flow.

Comparison method for multiple fields in MP (2)

Example: Flow Key: Version, IHL, TOS, source Address, destination Address

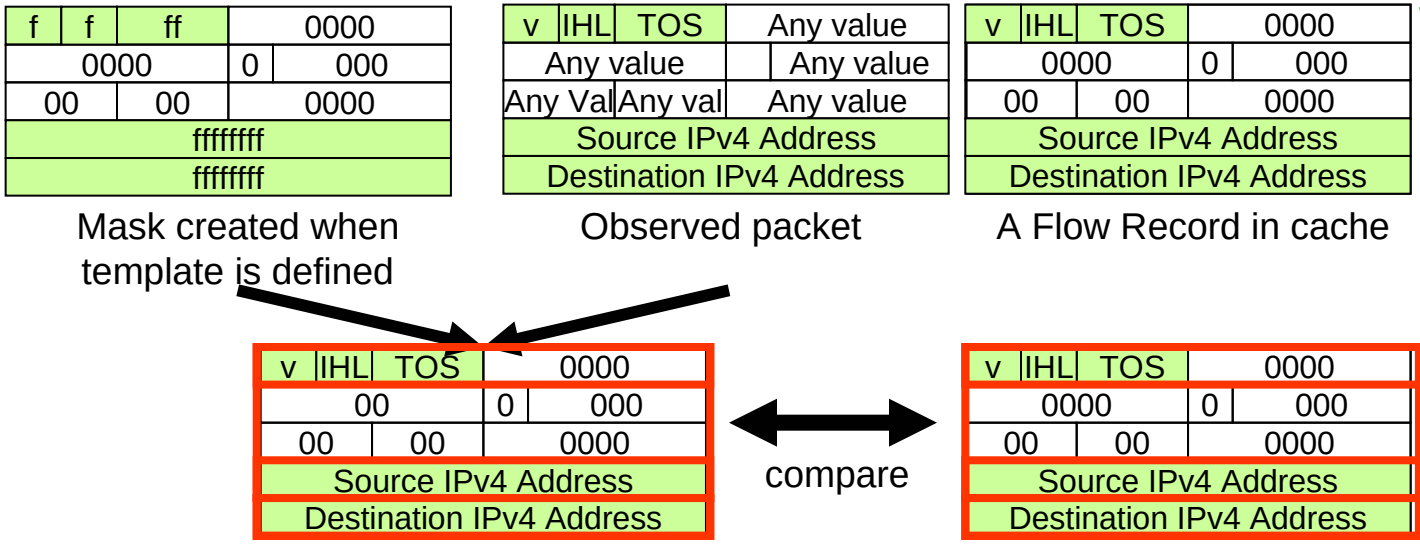
All fields are compared every time (general approach)



When a packet arrives:
5 comparisons
1. ip version
2. IHL
3. TOS
4. Source Address
5. Destination Address

Multiple field comparison (our approach)

Premise: Fields of Flow Records are placed in the referring order as packet header fields



When Template is defined:
Create a Mask

When a packet arrives:
Mask the packet
And
compare these memory areas at the same time
(e.g., memcmp in C language)

Or
1. v + IHL + TOS
2. Source Address
3. Destination Address
(32-bit architecture)

Comparison method for multiple fields in MP (3)

- Number of operations in this method
  - Mask costs smaller than comparison costs.
  - Therefore, this method is effective at increasing performance by reducing the number of comparisons, although it increases mask operations.

	Mask creation	Mask	Comparison
Number of operations	Once in an IPFIX session (when Template is defined)	Depends on the number of observed packets (when packet arrives)	Depends on the number of observed packets and number of flow records in cache

less ⟷ more

- Effective and ineffective cases

v	IHL	TOS	Total Length
ID		F	Offset
TTL	Protoco	IP Checksum	
Source IPv4 Address			
Destination IPv4 Address			

Effective case:
Flow Keys are placed densely

v	IHL	TOS	Total Length
ID		F	Offset
TTL	Protoco	IP Checksum	
Source IPv4 Address			
Destination IPv4 Address			

Ineffective case:
Flow Keys are placed sparsely.

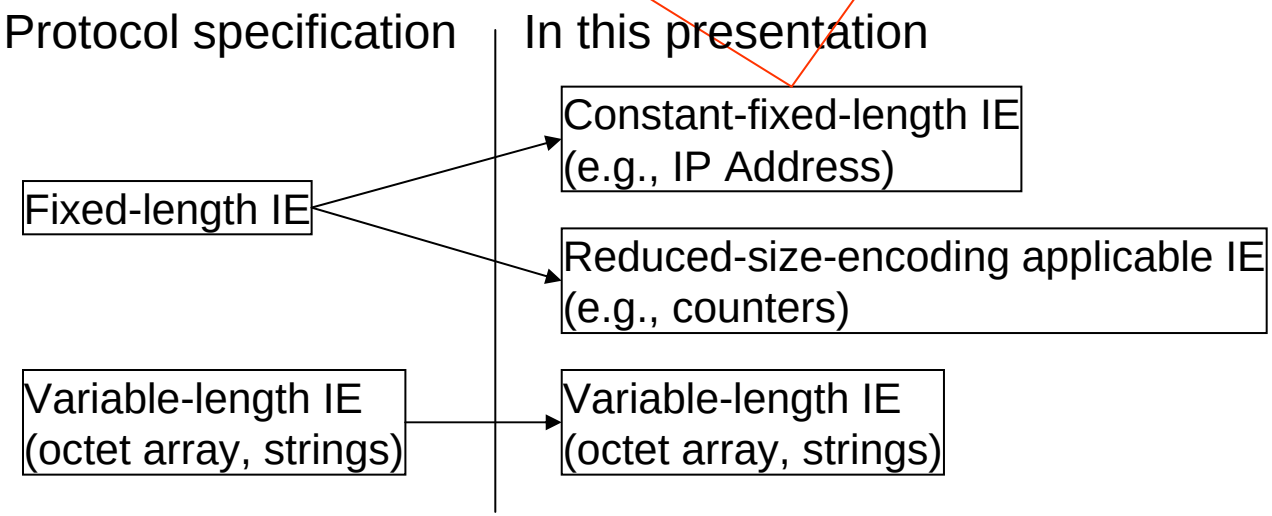
2nd idea to improve performance
in environment in which MP, EP, and CP use the same order

Copy method for multiple fields in Exporting Processes (EPs) and Collecting Processes (CPs)

NTT Network Service System Laboratories, NTT Corporation

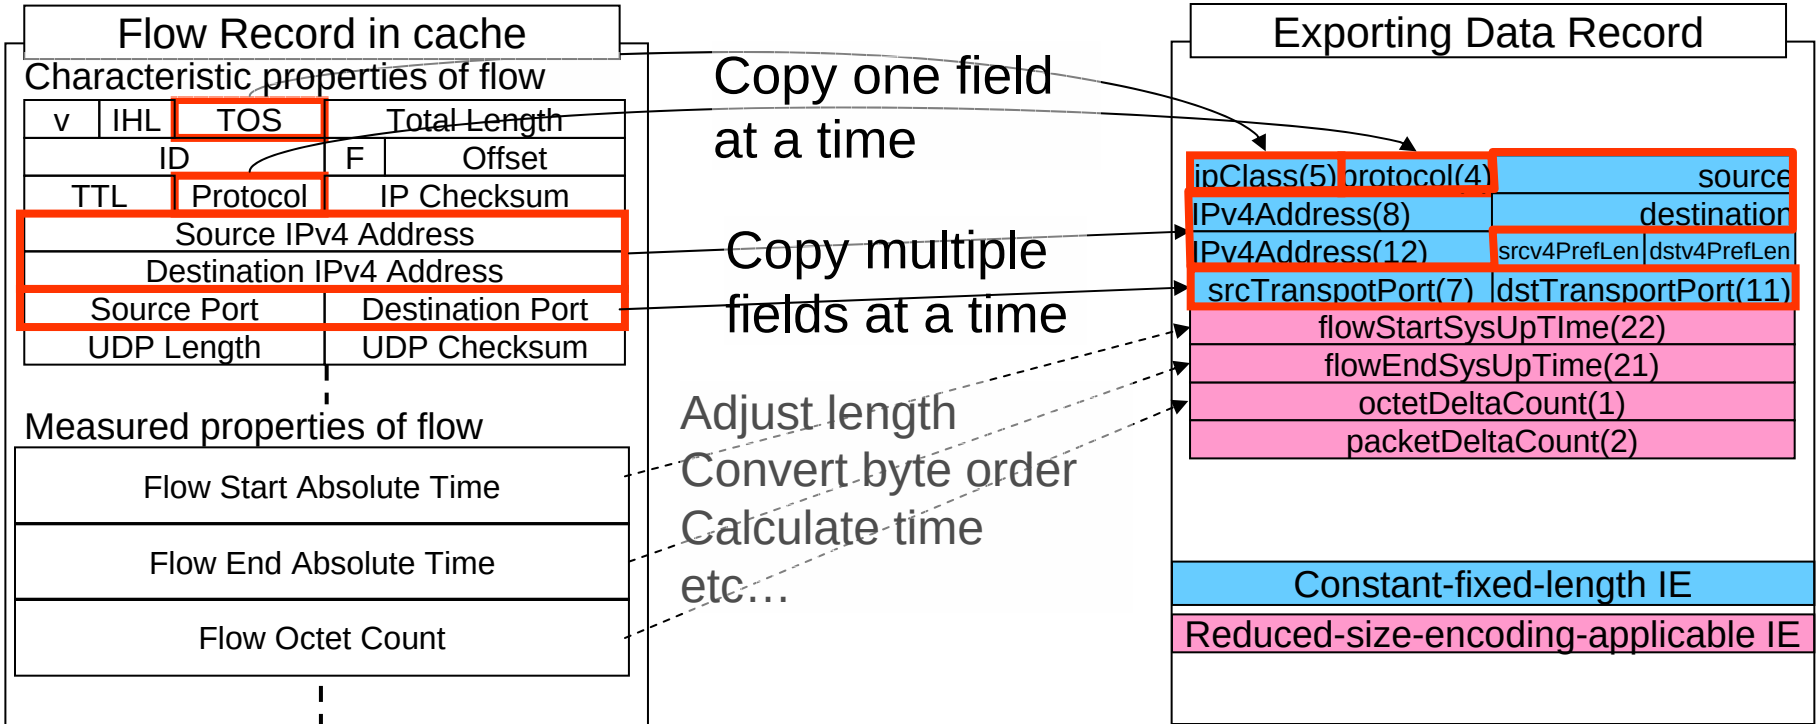
Overview of copy method for multiple fields

- It is a very simple method.
  - If fields in the format of cache and IEs in exporting Data Records are placed in the same order, EPs have a chance to copy multiple adjacent constant-fixed-length IEs at a time.
  - If IEs in received Data Records and fields in Collectors' internal format to store Flow Records are placed in the same order, CPs have a chance to copy multiple adjacent constant fixed-length IEs at a time too.
- IE size classification of IPFIX (terminology in this presentation)



Example of copy method for multiple fields in EP

- Conditions for copying multiple fields
  - Flow Record in cache and Exporting Data Record must use the same order.
  - IEs must have a constant fixed length.
    - Almost all IE characterizing properties of flow are constant fixed length.
  - Byte-orders must be the same.
    - Observed packet and Exporting Data Records use network byte order.
  - IEs for copying multiple fields must be adjacent.



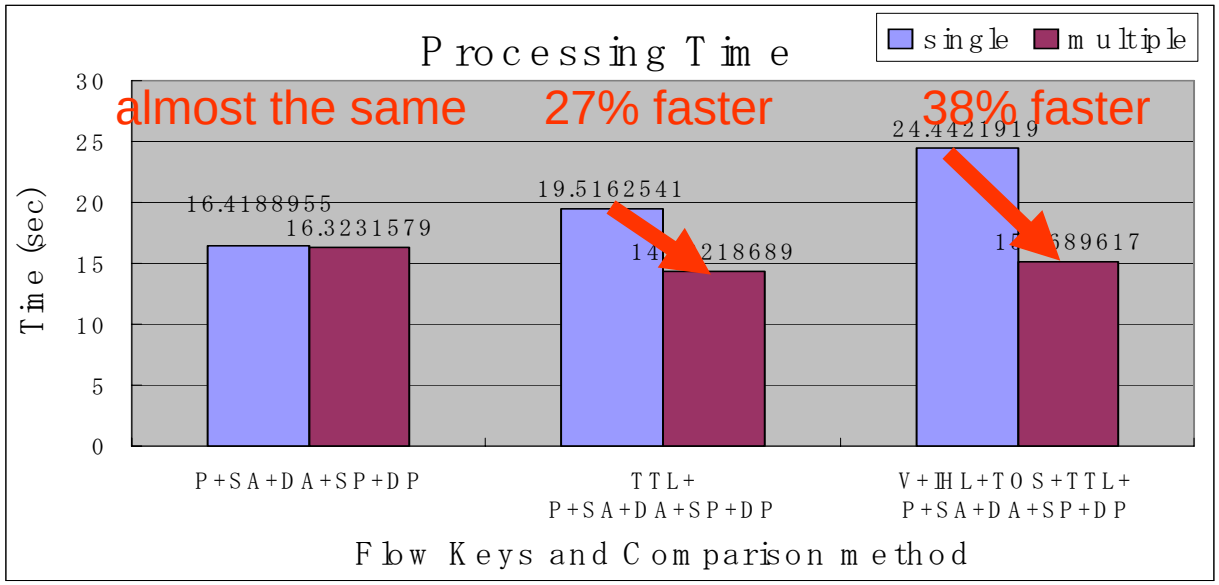
Evaluation & Conclusion

NTT Network Service System Laboratories, NTT Corporation

This material contains an evaluation about only comparison method.

If you want to see an evaluation about copy method, please see a material I talked in past IETF, http://www3.ietf.org/proceedings/07jul/slides/ipfix-10.pdf.

Evaluation of comparison method for multiple fields



v	IHL	TOS	Total Length
ID		F	Offset
TTL	Protoco		IP Checksum
Source IPv4 Address			
Destination IPv4 Address			
Source Port		Dst Port	

P+SA+DA+SP+DP

v	IHL	TOS	Total Length
ID		F	Offset
TTL	Protoco		IP Checksum
Source IPv4 Address			
Destination IPv4 Address			
Source Port		Dst Port	

TTL+
P+SA+DA+SP+DP

v	IHL	TOS	Total Length
ID		F	Offset
TTL	Protoco		IP Checksum
Source IPv4 Address			
Destination IPv4 Address			
Source Port		Dst Port	

V+IHL+TOS+TTL+
P+SA+DA+SP+DP

- **When the density of Flow Key fields is higher, this method works faster.**

Computing environment for the evaluation

- Software Exporter program
  - runs on Intel Xeon 3.06 GHz HT architecture
  - runs on Linux (debian/gnu Linux 4.0)
  - compiled by gcc4
    - optimized option: -O3
- Data used as observed packets:
  - PCAP data published by WIDE project.
  - contains 6,906,333 packets.
  - ftp://mawi.nezu.wide.ad.jp/pub/mawi/samplepoint-B/20060303/200603030100.dump.gz

Conclusion

- Introduced ideas to improve performances of IPFIX processes
  - **Comparison method for multiple fields in MPs**
  - **Copy method for multiple fields in EPs, and CPs**
- These ideas are based on **defining the order rule of IEs/fields**
  - Our Recommendation: IEs/f**ields are placed in the order referring to the packet header fields.**
- The order rule is published as an individual Internet Draft
  - http://tools.ietf.org/id/draft-irino-ipfix-ie-order-03.txt
  - If you agree with these ideas, work with us.

High Level Flow Correlation

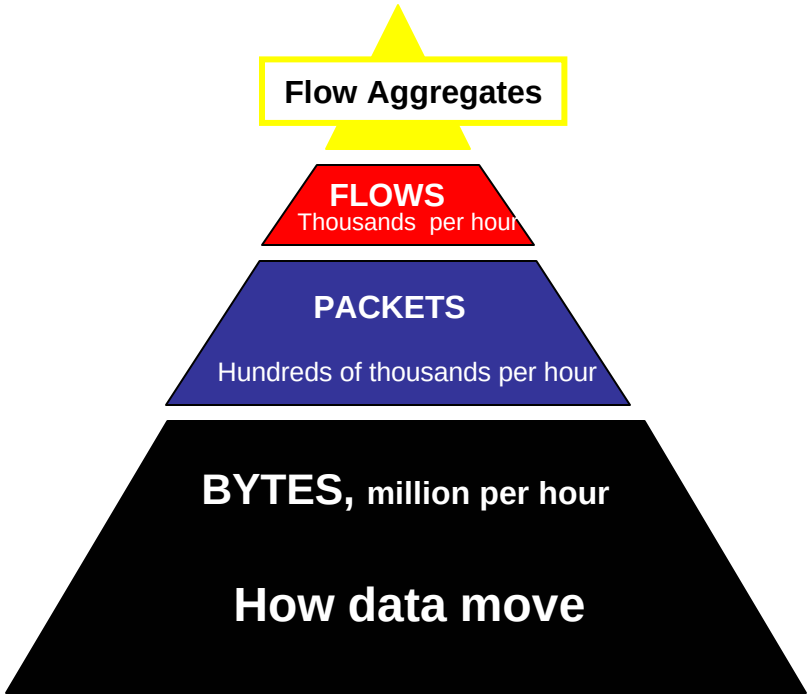
Valentino Crespi, California State Los Angeles, CA
Annarita Giani, UC Berkeley, CA
Rajiv Raghunarayan, Cisco Systems, Inc.

FloCon 2008, Savannah GA, January 7-10, 2008.

Outline

1. Extension of previous work on Flow Aggregation, (Flocon 2006).
2. Embedding of network traffic in a Euclidian Space.
3. Complex modeling.
4. Planned work.

Behind Flow Aggregation



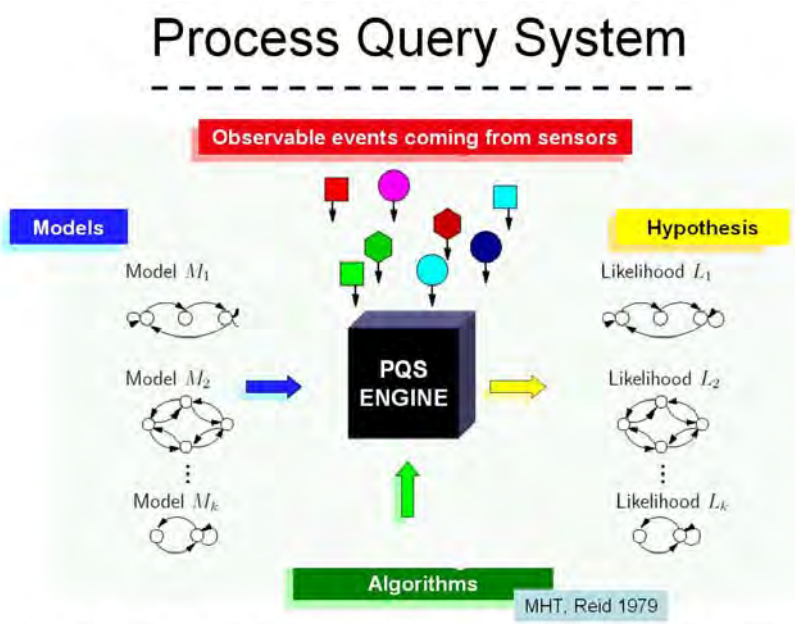
- Monitoring
- Anomaly detection
- Security analysis
- Traffic profiling
- Debugging
- Traffic engineering
- Usage-based profiling
- Network planning
- Pricing, peering

Data Reduction = Fewer events to be analyzed

Process based analysis of flows

We believe that **automated correlation at the raw flow level** is complicated and susceptible to false positives. The world consists of **processes** so our approach to correlation is process-based..

Implementation of a **PQS based process detection for Cyber Situational Awareness.**



The idea is to formulate hypotheses by associating new observations to an existing pool of rated hypotheses in all the possible ways and then calculate the new rating recursively.

Giani, De Souza, Berk, Cybenko, " Attribution and Aggregation of Network Flows for Security Analysis ," *in Proc. Flocon 2006,* Portland, OR.

Flow + Snort Alerts

Scenario: several packets in a flow triggered IDS alerts

Snort rule **1560** generates an alert when an attempt is made to exploit a known vulnerability in a web server or a web application.

Snort rule **1852** generates an alert when an attempt is made to access the 'robots.txt' file directly.

Timestamp	Sensor	src IP	dst IP	Proto
Jul 09 16:28:32	S1852	65.54.188.140	208.253.154.195	TCP
Jul 09 16:29:35	S1852	65.54.188.140	208.253.154.195	TCP
Jul 09 16:44:44	S1560	65.54.188.140	208.253.154.195	TCP
Jul 09 18:26:08	S1560	65.54.188.140	208.253.154.195	TCP
Jul 09 21:05:03	S1852	65.54.188.140	208.253.154.195	TCP
Jul 09 22:31:08	S1852	65.54.188.140	208.253.154.195	TCP
Jul 09 22:31:08	S1560	65.54.188.140	208.253.154.195	TCP
Jul 10 02:45:19	S1852	65.54.188.140	208.253.154.195	TCP
Jul 10 02:45:23	S1852	65.54.188.140	208.253.154.195	TCP
Jul 10 09:21:15	S1852	65.54.188.140	208.253.154.195	TCP
Jul 10 14:33:43	S1852	65.54.188.140	208.253.154.195	TCP
Jul 10 17:54:54	S1852	65.54.188.140	208.253.154.195	TCP
Jul 10 22:07:02	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 01:38:09	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 04:05:54	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 04:20:00	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 04:20:00	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 11:07:12	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 11:56:12	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 17:16:59	S1852	65.54.188.140	208.253.154.195	TCP
S Jul 10 02:30:27	F	65.54.188.140	208.253.154.195	TCP
E Jul 10 23:55:56				

SNORT ALERTS

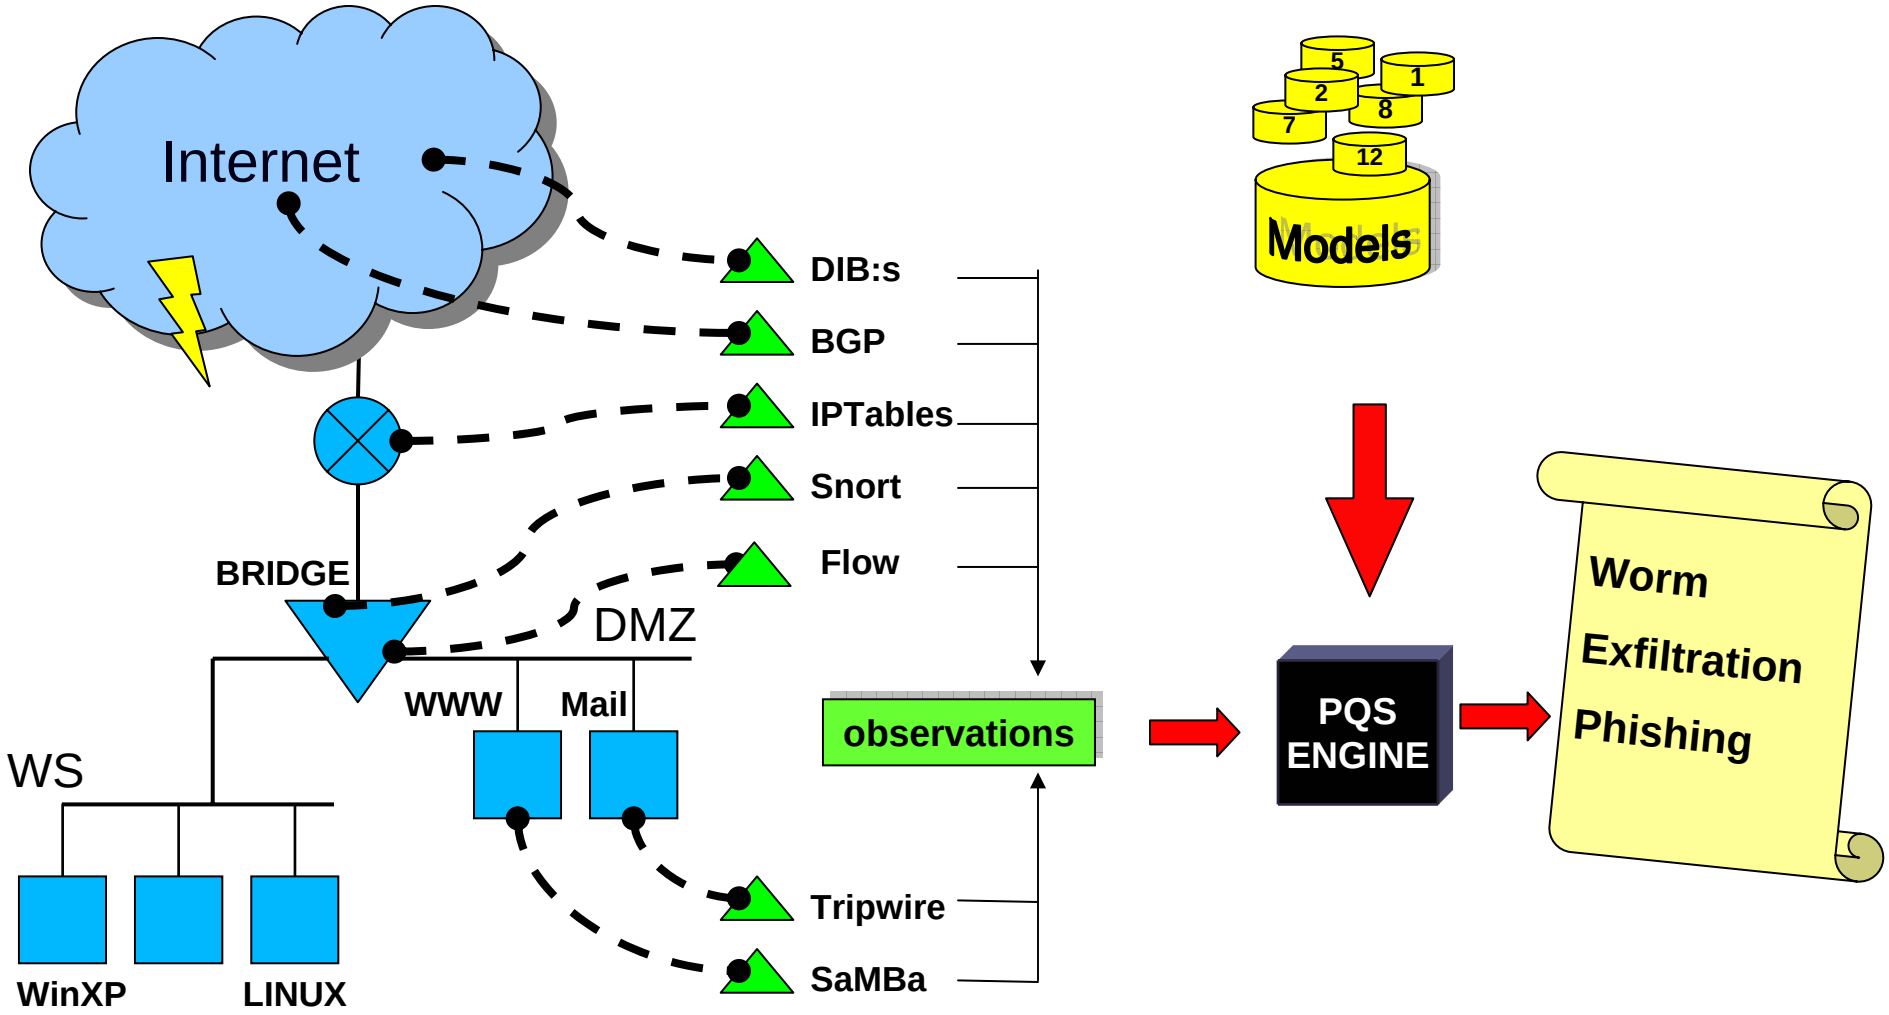
FLOW

Table 2: A sample track of correlated IDS and Flow events

The flow can be characterized as malicious and further investigation must be done.

Flow aggregation and correlations between flow data with security events

Correlation Engine for Computer Security



Sensors and Models

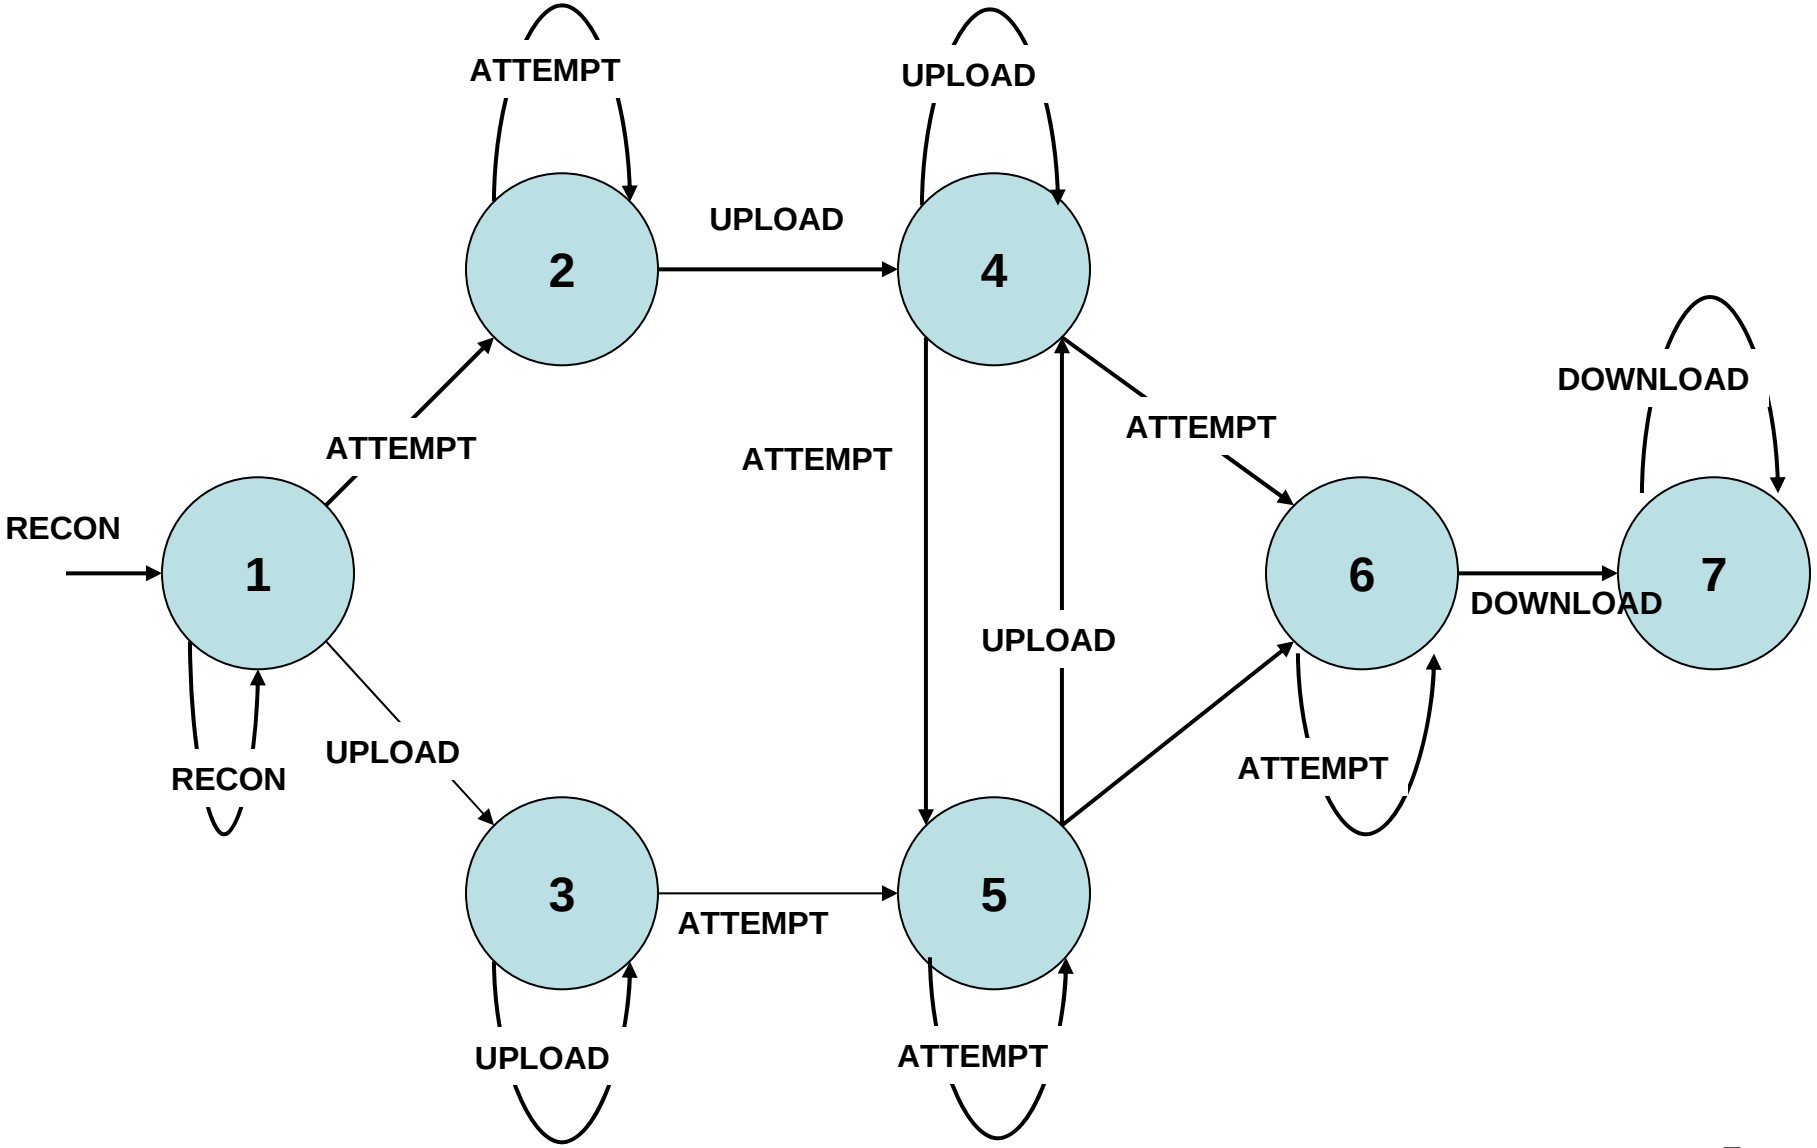
Sensors

1	**DIB:s**	**Dartmouth ICMP-T3 Bcc: System**
2	**Snort, Dragon**	**Signature Matching IDS**
3	**IPtables**	**Linux Netfilter firewall, log based**
4	**Samba**	**SMB server - file access reporting**
5	**Flow sensor**	**Network analysis**
6	**ClamAV**	**Virus scanner**
7	**Tripwire**	**Host filesystem integrity checker**

Models

1	**Noisy Internet Worm Propagation – fast scanning**
2	**Email Virus Propagation – hosts aggressively send emails**
3	**Low&Slow Stealthy Scans – of our entire network**
4	**Unauthorized Insider Document Access – insider information theft**
5	**Multistage Attack – several penetrations, inside our network**
6	**DATA movement**
7	**TIER 2 models**

Example - Phishing Attack Model



Current aggregators and analyzers

- **POWERFUL TOOLS** to understand the behavior of the network according to certain parameters, e.g. the amount of resources consumed, the variance on the various characteristics of the communication (source ip, destination ip), port.
- **PROBLEM**: They do not provide an analysis and a description of the dynamic evolution of network traffic.
- **NEED** for a structure that summarizes the behavior of the network.

OUR IDEA

Combine flow aggregation techniques with our previous process-based approach:

Use aggregators and flow analyzers to translate traffic into a process to be modeled and estimated.

Build circuits of Aggregating gates

1. Place observing nodes in multiple locations of the network (e.g. on each local router).
2. Each observing nodes dumps traffic flows to a Macro Aggregator.
3. Macro Aggregator: *circuit*. Each gate is a flow aggregator
   - First layer consists of classical aggregators that output flow aggregates. Successive layers process aggregates of flow aggregates
   - Final output: a vector function of the dumped traffic ranging in $\mathbf{R}^n$:

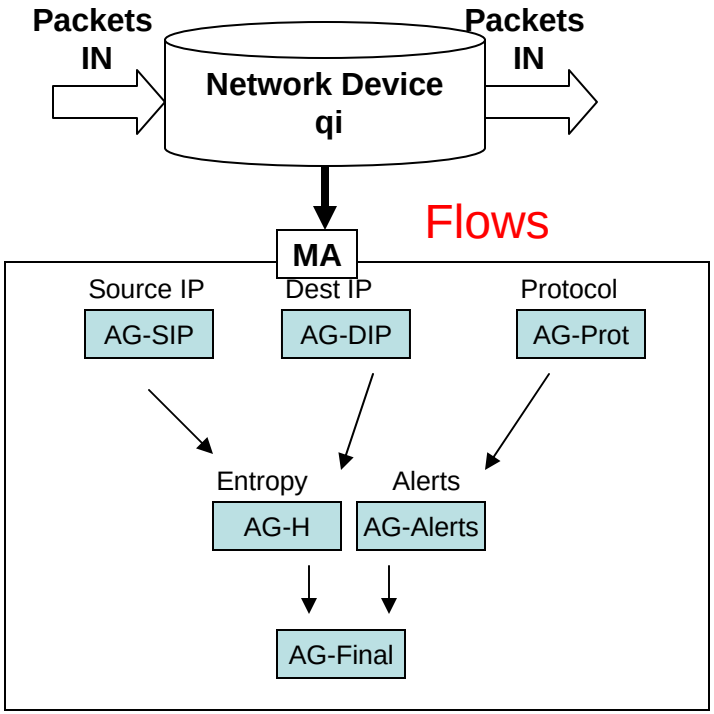
$$\mathbf{X}(t) = (x_1(t), x_2(t), \cdots, x_n(t))$$

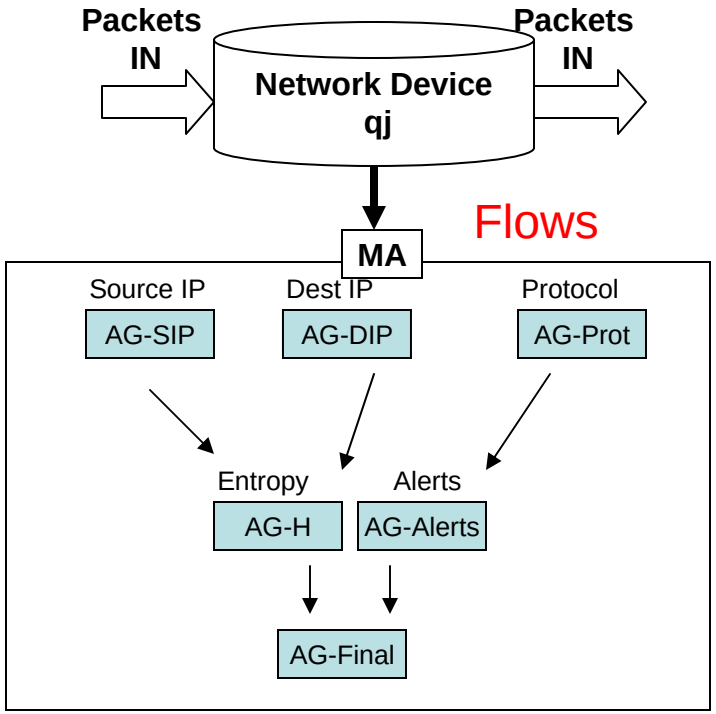
At each time the observing nodes produce a set of vectors:

$$S(t) = \{X_1(t), X_2(t), \cdots, X_n(t)\}$$

1. Identify and Analyze properties of S(t) over time to characterize/detect anomalies.

Embed Traffic in Euclidean Space





$$X_i(t) = (x_{i,1}(t), x_{i,2}(t), \cdots, x_{i,n}(t))$$

$$X_j(t) = (x_{j,1}(t), x_{j,2}(t), \cdots, x_{j,n}(t))$$

(Entropy S-IP,Entropy D-IP, Average Size,…,%TCP Traffic,%UDP Traffic)

$$S(t) = \{..., X_i(t), ..., X_j(t), ...\}$$

Aggregating flows (1)

Flows of security attacks usually have common patterns and form conspicuous traffic clusters. Identifies clusters of attack flows in real time and aggregates those large number of short attack flows into few metaflows.

Same sourceIP ~ worm propagation
Same destIP ~ Denial of Service Attack
Same destIP and SourceIP ~ most portscan

Purpose is mostly security.

Example:

Keys: srcIP, dstIP, srcPort (Proto), dstPort (Proto)

Properties of clusters containing attack traffic:

- Fixed value in one of the keys.
- Number of flows in cluster large.
- Size of flows in cluster small.

Random Dimensions

(srcIP, * , * , dstPort (Proto)

Fixed Keys

Aggregation Priority Parameter

X = dstIP

$$H(X) = -\sum_i P[X = ip_i] \cdot \log P[X = ip_i]$$

$$APP = \min_{X \in \{dstIP, srcPort\}} \{H(X)\}$$

$$x_i = APP$$

Entropy Based Flow Aggregation (2006)
Yan Hu, Dah-Ming Chiu, and John C.S. Lui
The Chinese University of Hong Kong

Aggregating flows (2)

A small percentage of flows consume most of the network bandwidth.

Study of heavy flows in 4 orthogonal dimensions:

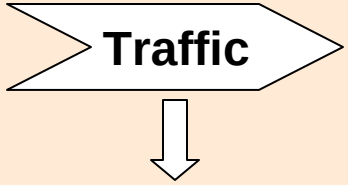
- Size
- Duration
- Rate
- Burstiness

Flow is divided in bins bi of duration T.
Burstiness is the standard deviation of bi

and examine their correlations.

The flow can be:

- Elephant - large size flows
- Tortoise - high duration flows
- Cheetah - large rate flows
- Porcupine – high burstiness



% of Elephant = E
% of Tortoise = T
% of Cheetah = C
% of Porcupine = P

$$x_i = (E, T, C, P)$$

On the correlation of Internet flow characteristics (2003)
Kun-Chan Lan, John Heidemann
Information Science Institute, University of Southern California

Analyze *S(t)* over time

Our Interest: Study the evolution of correlations of flow aggregate attributes over time in order to detect/estimate anomalies and attacks.

$$S(t) = \{X_1(t), X_2(t), \cdots, X_n(t)\}$$

- Build models (continuous, discrete event, probabilistic, etc.) to describe the observation process S(t).
- Track the unknown states of the built model given S(t).
- Apply Learning Techniques to learn models.
- Use Tracking Machines to estimate the hidden state sequence given the observables and the models.

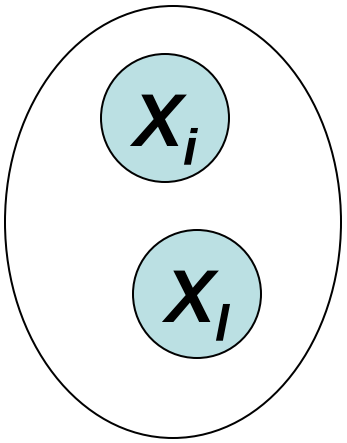
Metric of traffic features

$$S(t) = \{X_1(t), X_2(t), \cdots, X_n(t)\}$$

Use clustering techniques (e.g., spectral clustering, k-means based algorithms, etc.) to clusterize the observing nodes and detect “shifts” of traffic characteristics between different observation nodes:

1. Study how clusters change over time and characterize/detect anomalies.
2. Use clusters to produce a graphic representation of the traffic.
3. Define discrete models to describe the evolution of clusters in relation to specific events: coordinated computer attacks, presence of covert channels, bugs in the network software, hardware breakdowns, etc.

Spectral Clustering

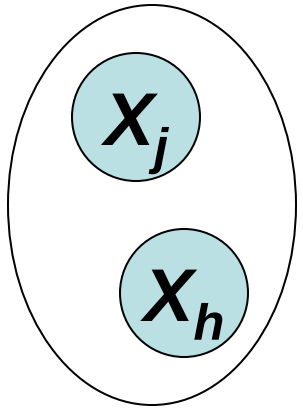


Input: Similarity Matrix M=[aij], , number k>0

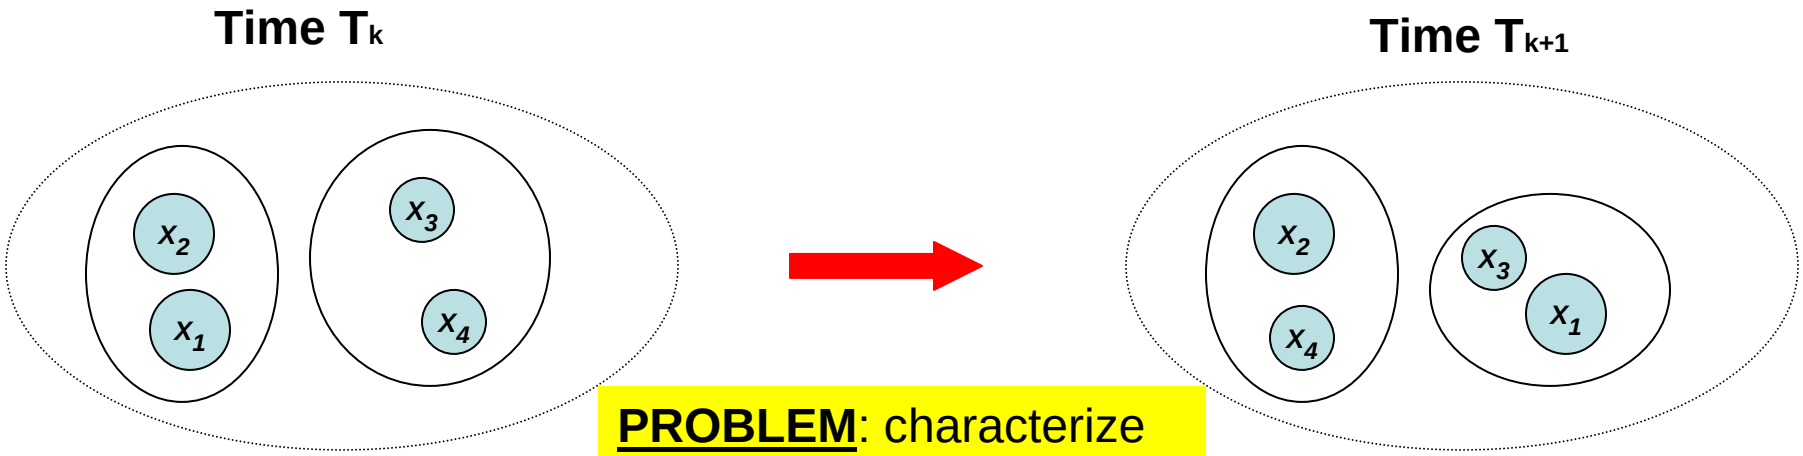
$a_{ij} = s(X_i, X_j)$ e.g $a_{ij} = \exp(-\|X_i - X_j\| / 2\sigma^2)$

- Build similarity graph. For example the Graph whose adjacency matrix AG = M.
- L = Laplacian(AG)
- Compute the k eigenvectors of L associated with the k smallest eigenvalues: v1, v2,…,vk
- V = [v1 v2 … vk], nxk matrix
- Pick the rows of V: y1, y2,…,yn
- Cluster yi's using k-means algorithm into C1,C2,…Ck

Output: clusters C1,C2,…,Ck



Discrete Models of Cluster Evolution



No attacks

PROBLEM: characterize anomalous unobservable events associated with transitions between clusters.

Attack

Idea: Build Deterministic Finite State Automata models to identify transitions. In this case we identify anomalies by studying the current clustering in relation to the previous "snapshot" of traffic

Challenges

- Parameter estimation: in our example of clustering k was fixed.
- Define and learn models of the system’s dynamics.
- Identify relevant attributes of flow aggregators to obtain significant vectors.
- Define appropriate similarity function.
- Use a realistic Data Set to verify approaches.

Planned Work

- Implement clustering method.
- Develop discrete models.
- Build a software monitor to analyze traffic through clusters and vector representation.
- Experimental analysis of the efficaciousness of our approach.

References

[1] I. S. Dhillon, Y. Guan, and B. Kulis. Kernel k-means, Spectral Clustering and Normalized Cuts. In *Proceedings of the KDD'04 Workshop*, Seattle, Washington, August 2004.

[2] C. Estan, S. Savage, and G. Varghese. Automatically Inferring Patterns of Resource Consumption in Network Traffic. In *Proceedings of the 2004 SIGCOMM*.

[3] A. Giani, I. G. D. Souza, V. Berk, and G. Cybenko. Attribution and Aggregation of Network Flows for Security Analysis. In *Proceedings of FloCon 2006*.

[4] Y. Hu, D.-M. Chiu, and J. C. Lui. Adaptive Flow Aggregation - A New Solution for Robust Flow Monitoring under Security Attacks. In *Proceedings of 2006 IEEE/IFIP Network Operations and Management Symposium (IEEE/IFIP NOMS)*.

[5] Y. Hu, D. Chui, and J. C. Lui. Adaptive Flow Aggregation - a New Solution for Robust Flow Monitoring under Security Attacks. In *Proceedings of the 2006 Network Operations and Management Symposium*.

[6] K. Keys, D. Moore, and C. Esten. A Robust System for Accurate Real-time Summaries of Internet Traffic. In *Proceedings of the 2005 SIGMETRICS*, June 2005.

[7] L. Rodrigues and P. R. Guardieiro. A Spatial and Temporal Analysis of Internet Aggregate Traffic at the Flow Level. In *Proceedings of the 2004 Global Telecommunications Conference (GLOBECOM)*, volume 2.

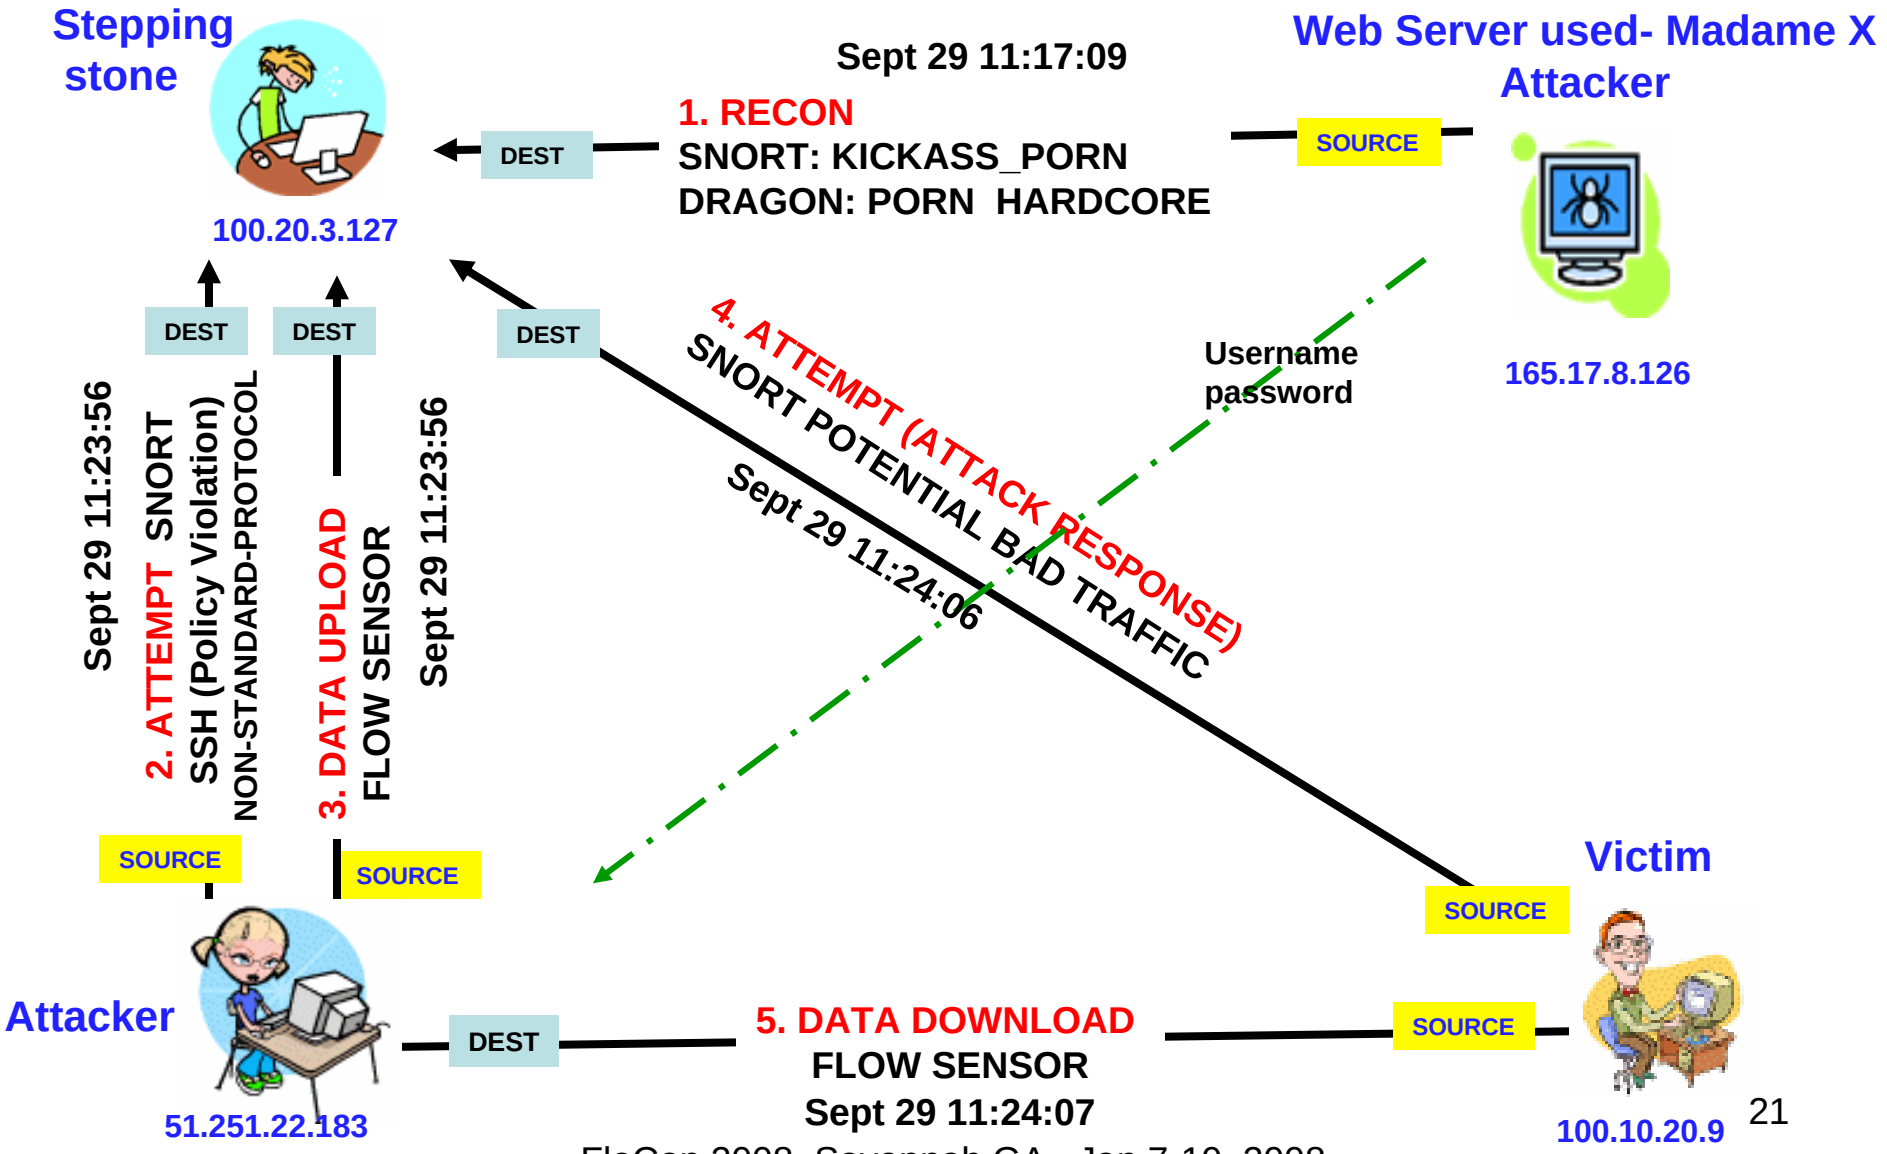
[8] B. Trammell and C. Gates. Naf: The NetSA Aggregated Flow Tool Suite. In *Proceedings of the Large Installation System Administration Conference (LISA 2006)*, 2006.

[9] U. von Luxburg. A Tutorial on Spectral Clustering. Technical Report TR-149, Max-Planck-Institut für biologische Kybernetik, 2006.

Thanks

Annarita Giani <agiani@eecs.berkeley.edu>
Valentino Crespi <vcrespi@calstatela.edu>
Rajiv Raghunarayan <raraghun@cisco.com>

Complex Phishing Attack Observables



Identifying Anomalous Traffic Using Delta Traffic

Tsuyoshi KONDOH and Keisuke ISHIBASHI
Information Sharing Platform Labs.
NTT

Flocon2008, January 7–10, 2008, Savannah GA

Outline

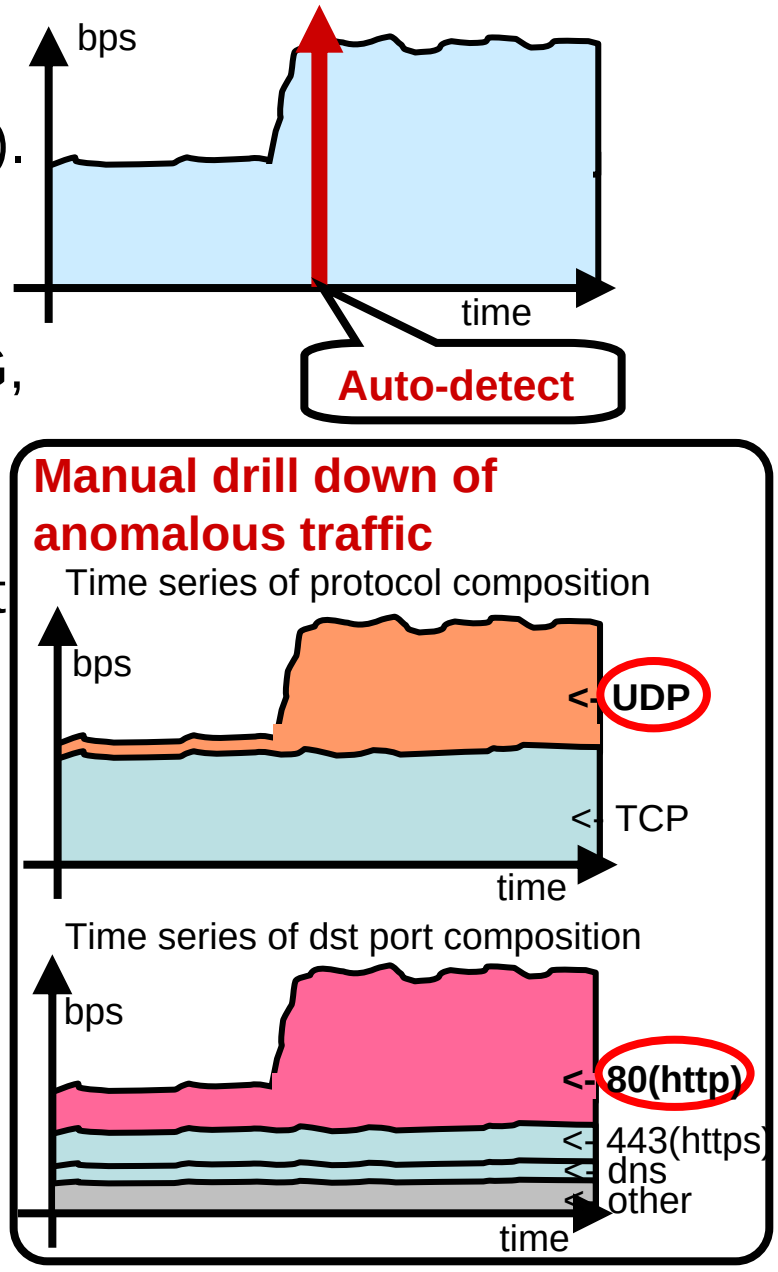
- Background and Motivation
  - Identifying anomalous traffic is the missing piece.
- Our Technique: DELTAA
  - Concepts
    1. Extract anomalous traffic as the delta of normal and anomalous time periods.
    2. Auto-aggregate extracted anomalous traffic.
  - Operation of our technique
    - Show the step by step operation of our technique.
- Evaluation
  - Evaluation using synthesized DDoS traffic.
- Summary

Copyright NTT Information Sharing Platform Laboratories, 2008

Background and Motivation

- Monitoring of traffic volumes is widely used for network operation (e.g. MRTG).
- Many techniques for detecting anomalous volume change have been proposed (NBAD, Holt-winters in MRTG, … etc.).
- Some tools to mitigate damage from anomalous traffic. (e.g. drop/rate limit at router, detour to Cisco Guard, etc.)

- However, accurate mitigation needs accurate ACLs (ACL set).
- But, Generating accurate ACL set requires manual drill down by operator.
  - **It's too costly.**



Our Technique: DELTAA

- **DELTAA** outputs ACL set for filtering or rate limiting to mitigate the damage from anomalous traffic.
  - DELTAA: Delta Traffic Automatic Aggregator

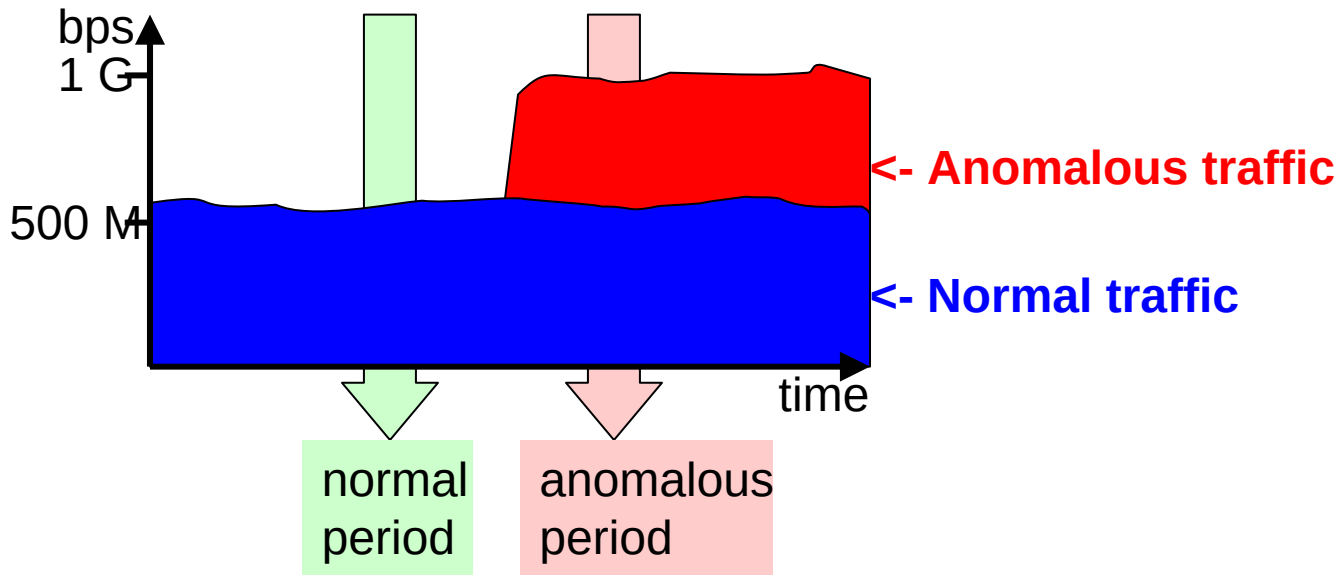
- Three concepts of DELTAA:
  1. Reveal anomalous traffic using delta traffic, between normal and anomalous periods.
  2. Aggregate delta traffic and generate optimized ACL set on single dimensions.
     - Dimension means source IP address, destination IP address, protocol or port numbers.
  3. Generate multi-dimensional ACL set by integrating each single dimensional ACL set.

Copyright NTT Information Sharing Platform Laboratories, 2008

Concept #1:
(1) Definition of "Normal" and "Anomalous" Traffic

Throughout this presentation, I use the following definitions.

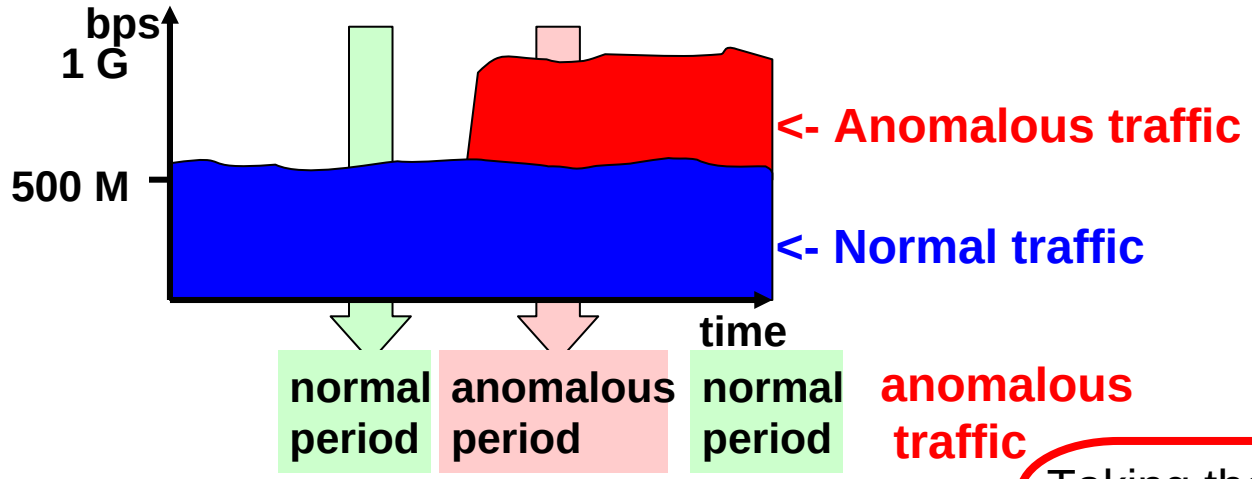
1. Anomalous traffic: Traffic that causes a change in traffic volume (bps/pps/fps).
   - BitTorrent and server intrusion are out of scope because they always exist or do not cause a change in traffic volume.
2. Normal period: Period when traffic volume is normal.
3. Anomalous period: Period when traffic volume is anomalous.



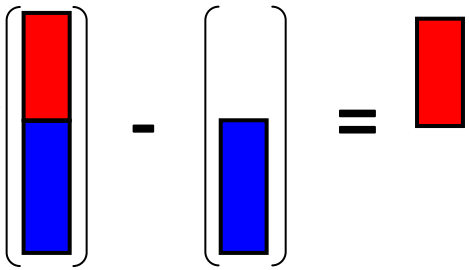
Concept #1 :
(2) Reveal Anomalous Traffic

- Make two assumptions
  1. traffic of normal period = normal traffic
  2. traffic of anomalous period = normal traffic + anomalous traffic

- anomalous traffic = traffic of anomalous period – traffic of normal period



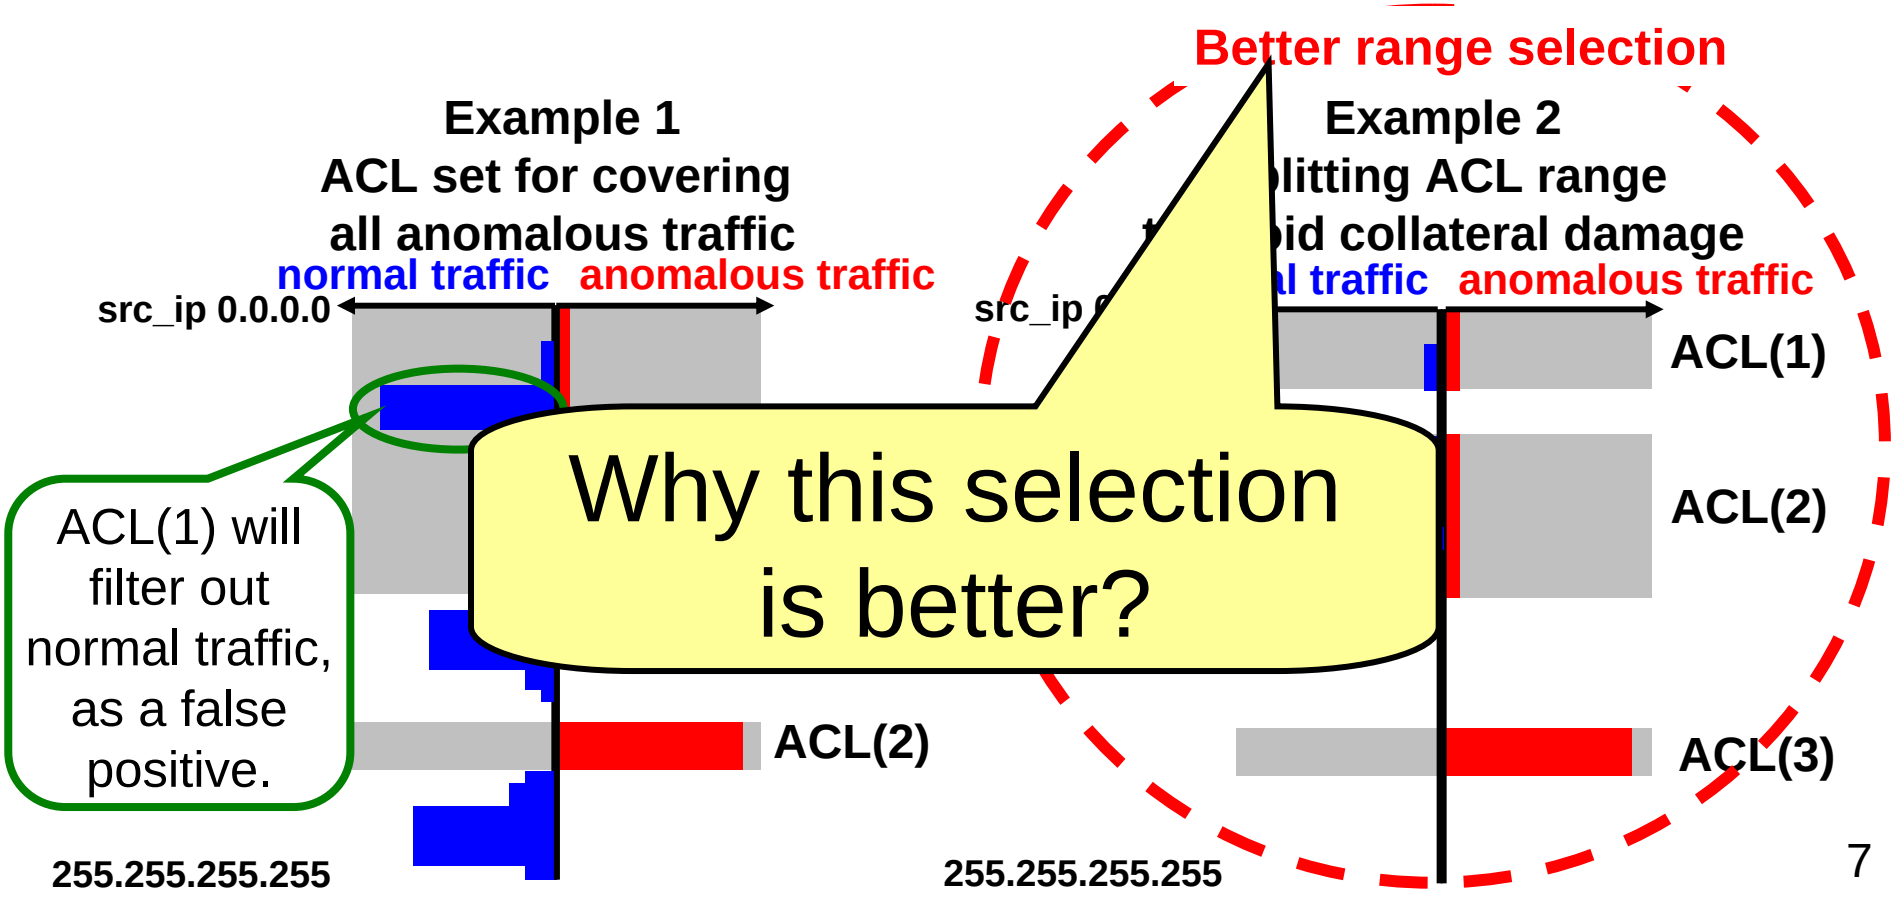
Extracting anomalous traffic from "traffic of anomalous period" is difficult because it is a mixture of normal and anomalous traffic.



Taking the delta between "traffic of normal period" and that of anomalous period, we can effectively extract anomalous traffic.

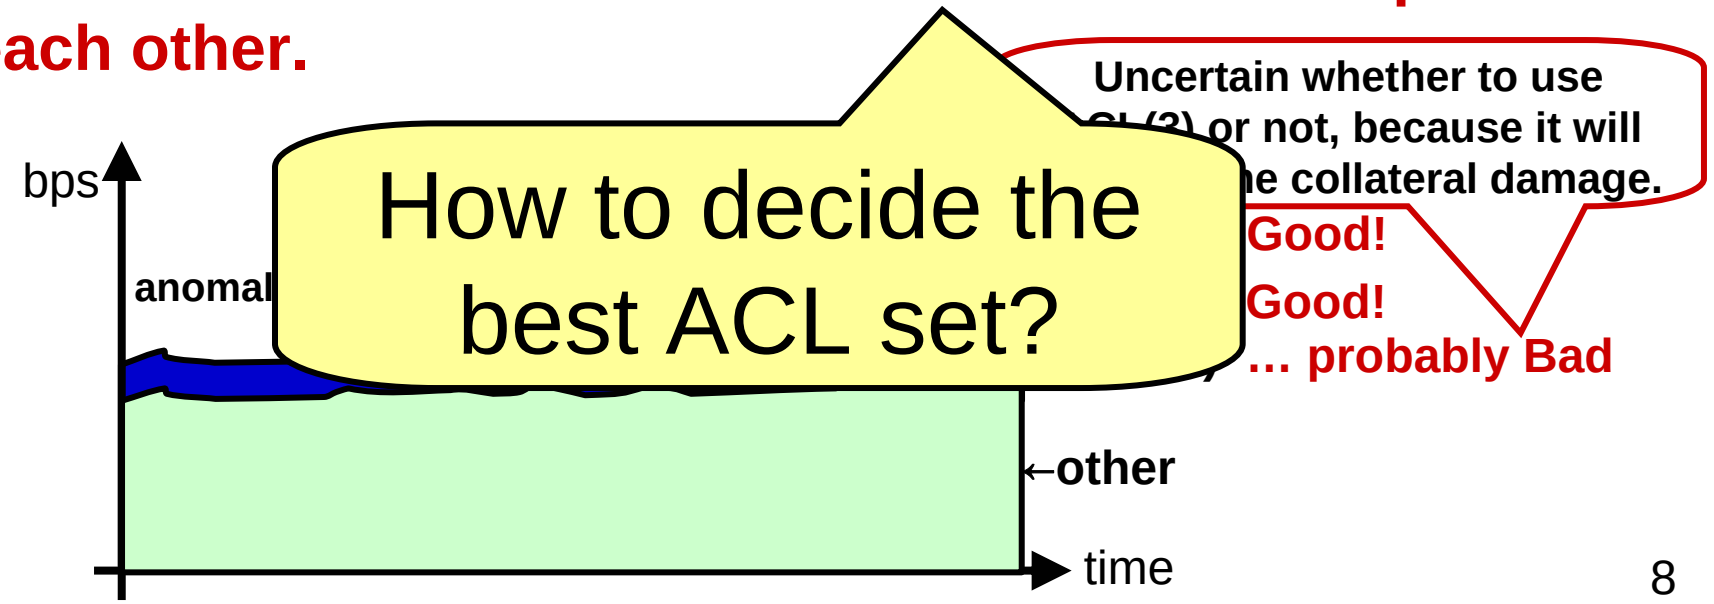
Concept #2:
Auto-aggregate Delta Traffic

- Our technique expresses anomalous traffic with some number of ranges.
  - For example source IP address ranges.
  - The ranges should be optimal for filtering.



Criteria of "Goodness"

- We introduce three criteria of identification.
  1. Coverage ratio: (1 – FNR)
     Maximize filtered anomalous traffic
  2. Collateral (damage) ratio: (FPR)
     Minimize filtered (normal) legitimate traffic
  3. **Number of ACLs**:
     ACL entry budget is limited, so fewer ACLs is better.
- **These three criteria have a trade-off relationship with each other.**



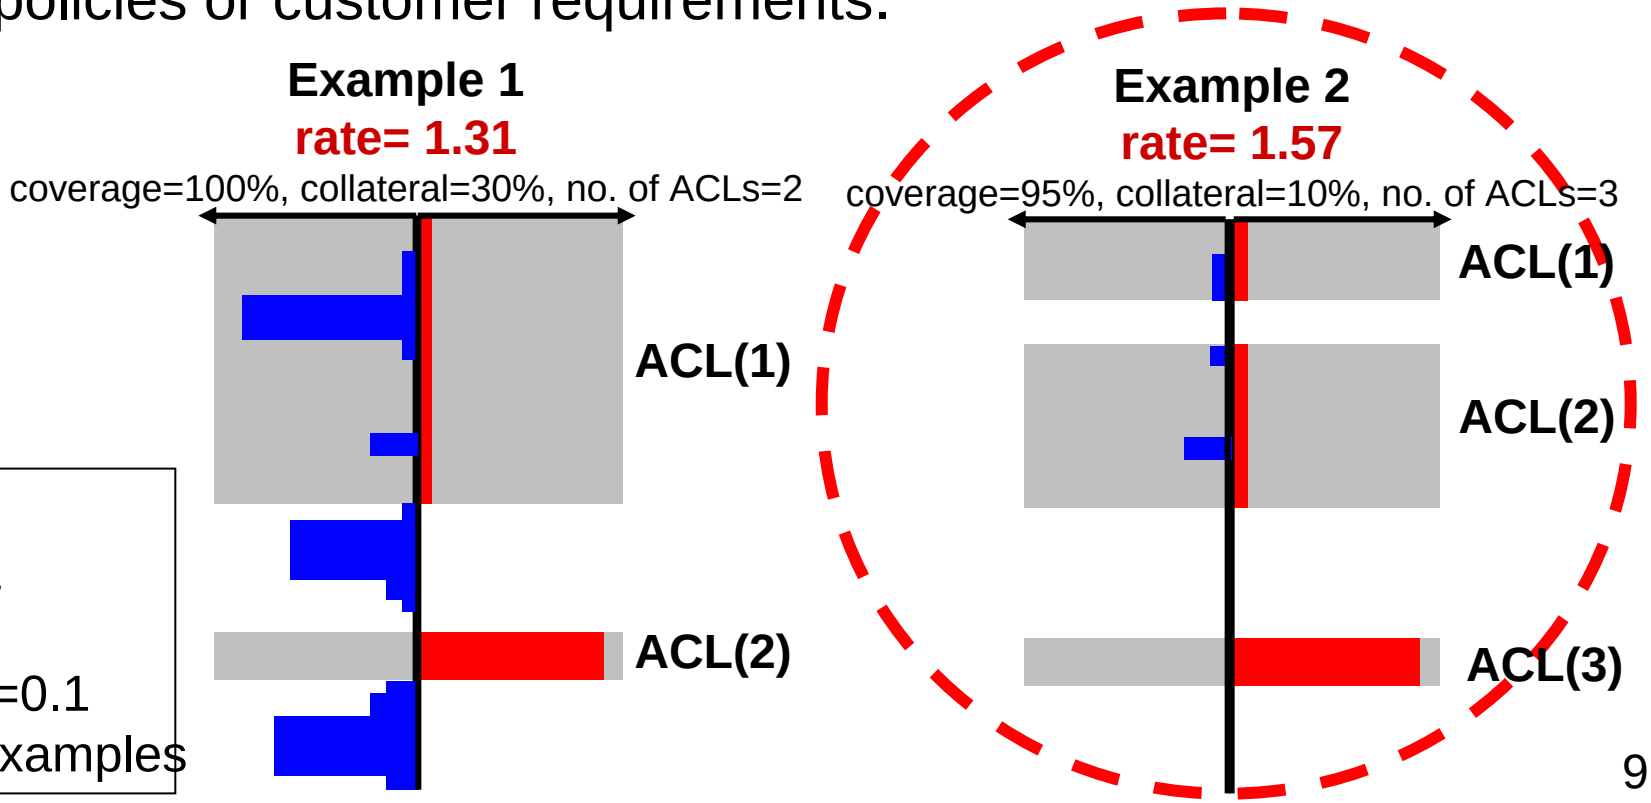
Evaluation Formula for Goodness

To decide the best ACL set, we introduce this formula:

$$\text{coverage}: cov,\ \text{collateral ratio}: coll,\ \text{no. of ACLs}: n$$

$$\text{rate} = \frac{(\beta - \alpha) + \alpha \cdot cov - \beta \cdot coll}{n^{\gamma}} \quad (\alpha, \beta, \gamma : \text{weighting coefficients})$$

– By tuning the weighting coefficients, we can reflect network policies or customer requirements.



Use
-alpha=1
-beta=2
-gamma=0.1
for this examples

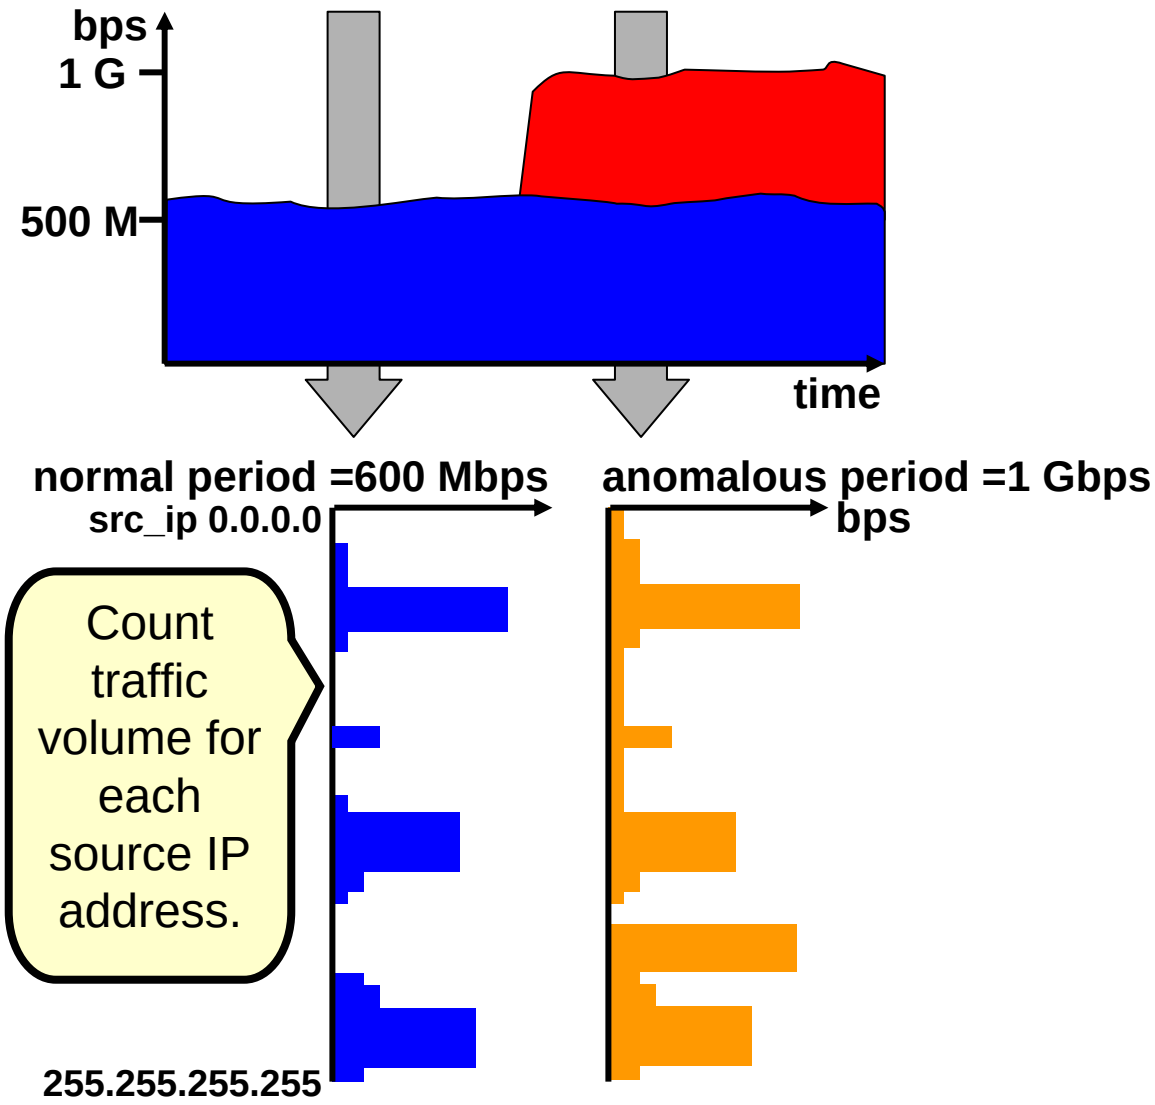
Step by Step Explanation of Our Technique

- Following seven pages show the step by step operation of our technique including above two concepts and concept #3.

Copyright NTT Information Sharing Platform Laboratories, 2008

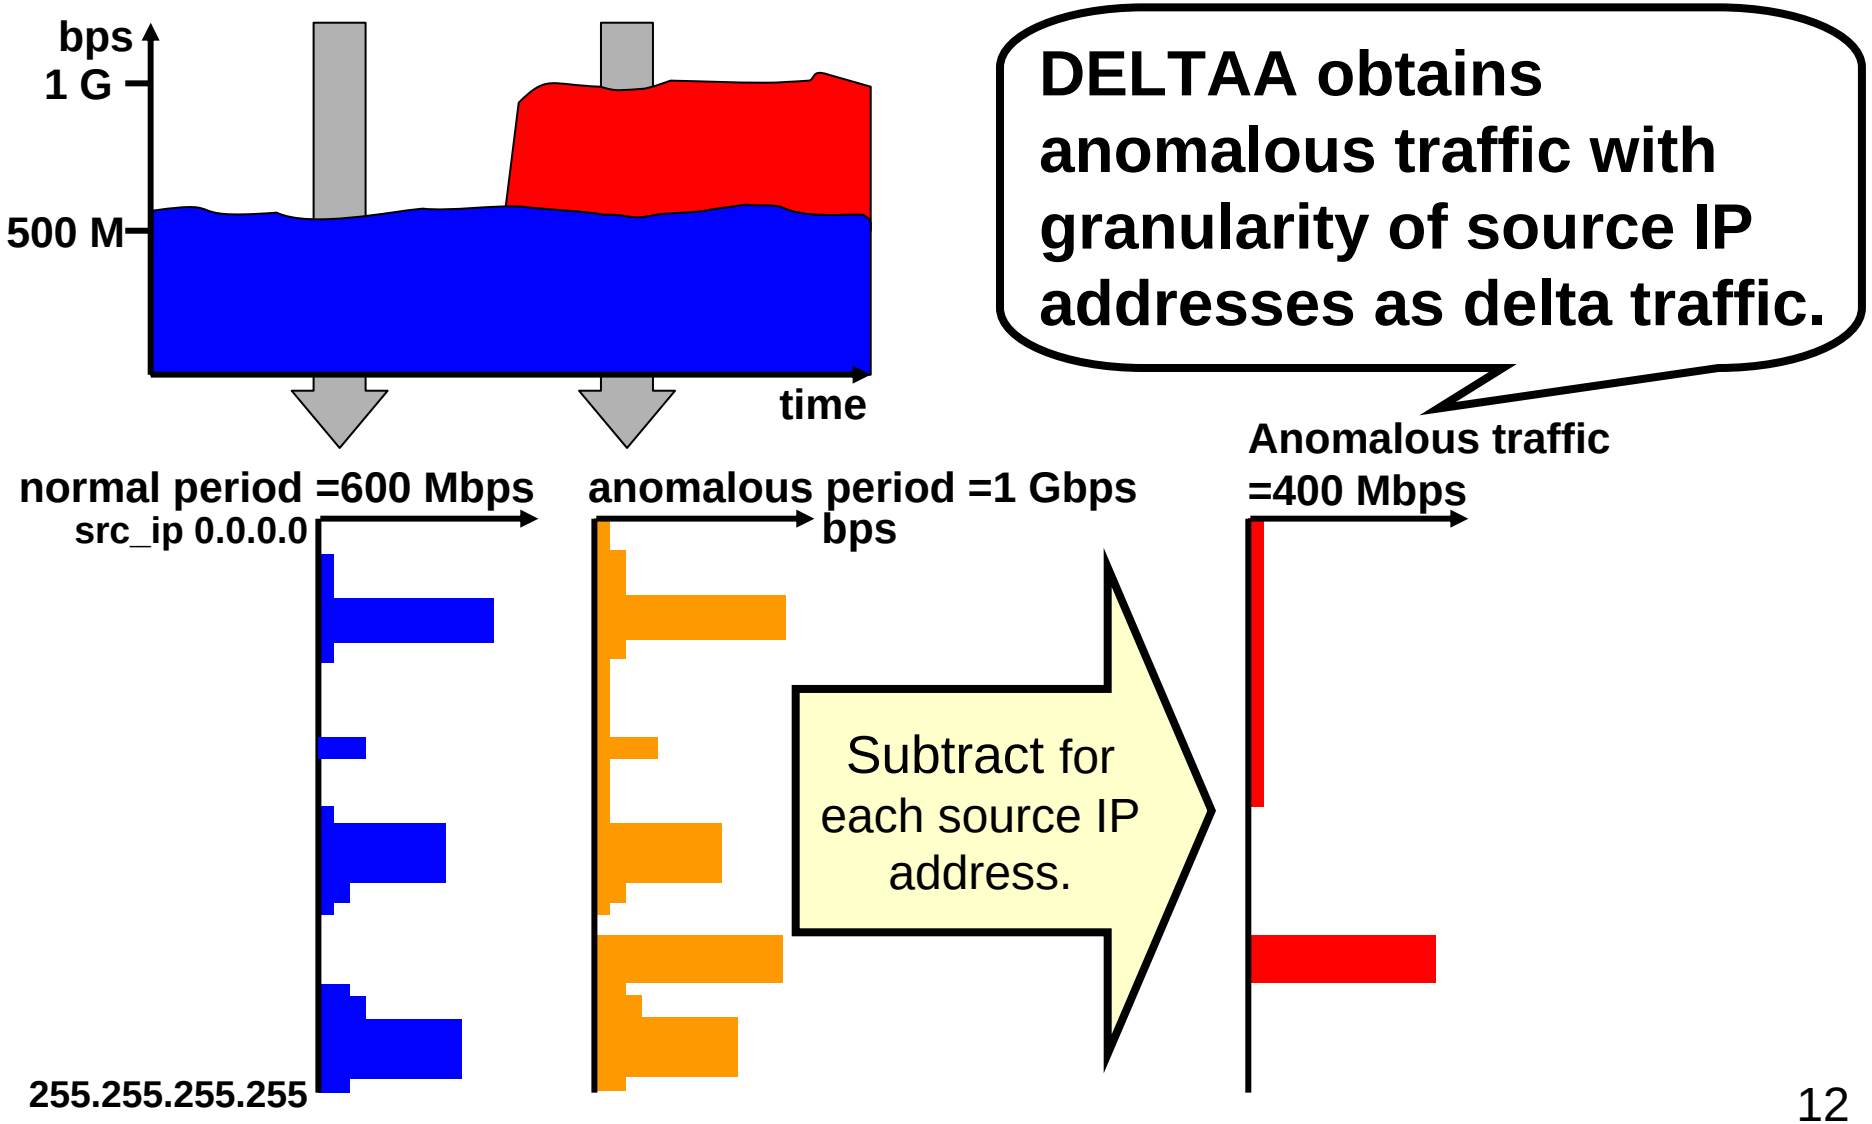
Step 1: (1) Counting Up

Count traffic of both normal and anomalous periods for each source IP address.



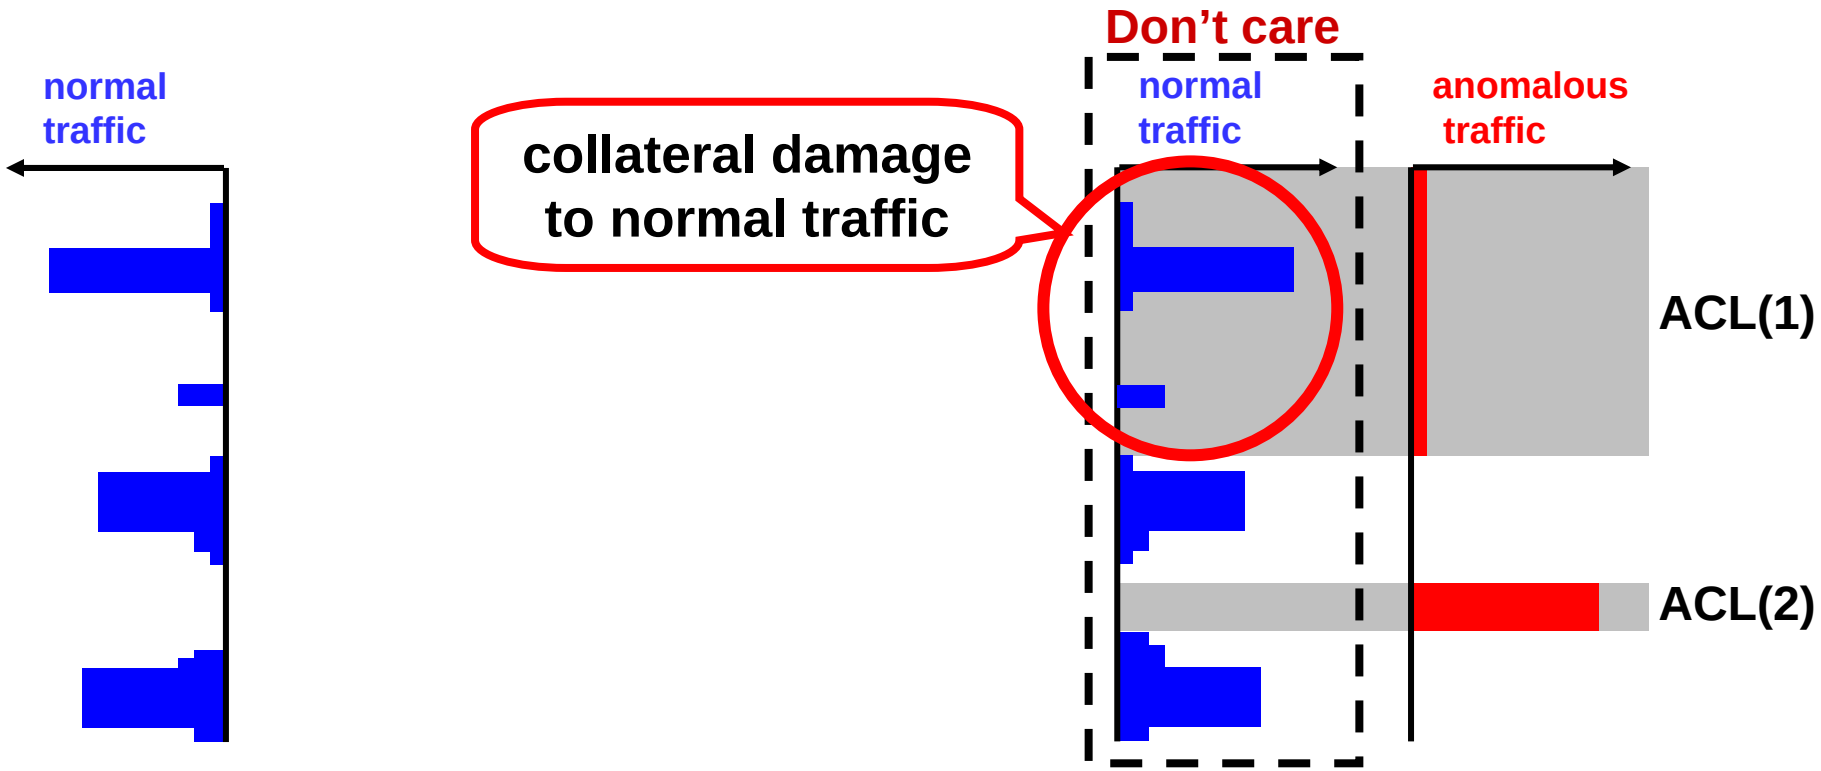
Step 1: (2) Making Delta Traffic

Make delta traffic by subtracting traffic of normal period from that of anomalous period.



Step 2: (1) Deciding ACL Set as IP Address Ranges

- When using anomalous traffic information only, collateral damage cannot be avoided.
  - Causes miss-filtering of normal traffic.
- So, we need to use information on both normal and anomalous traffic.

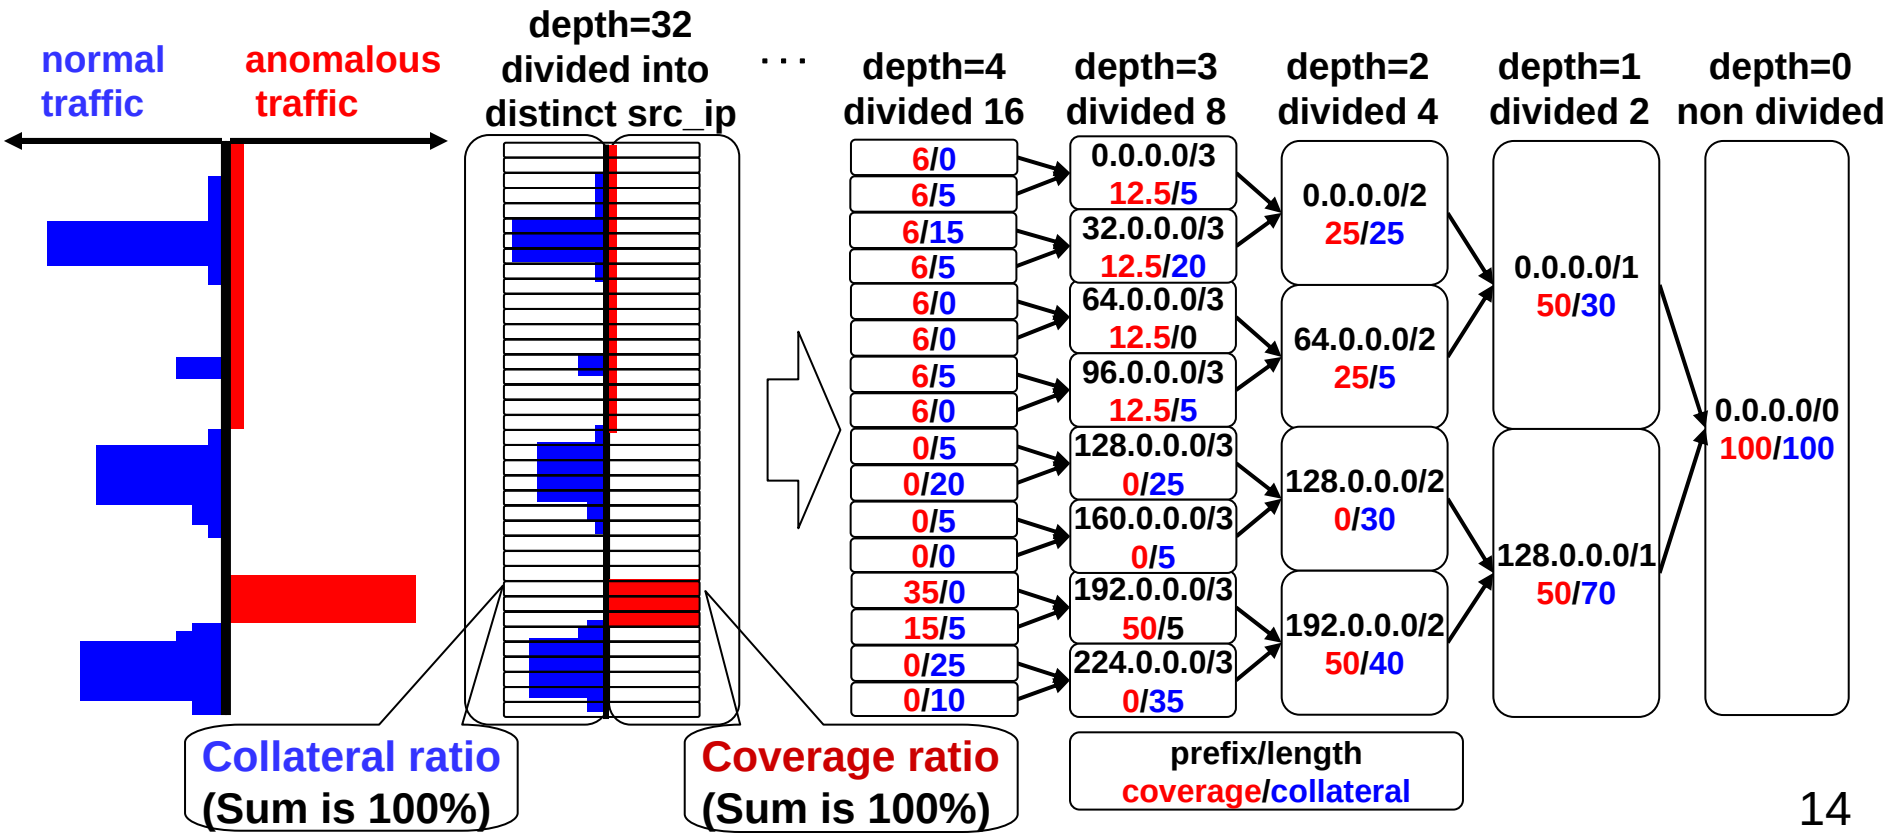


Copyright NTT Information Sharing Platform Laboratories, 2008

Step 2: (2) Building Tree of Normal and Anomalous Traffic

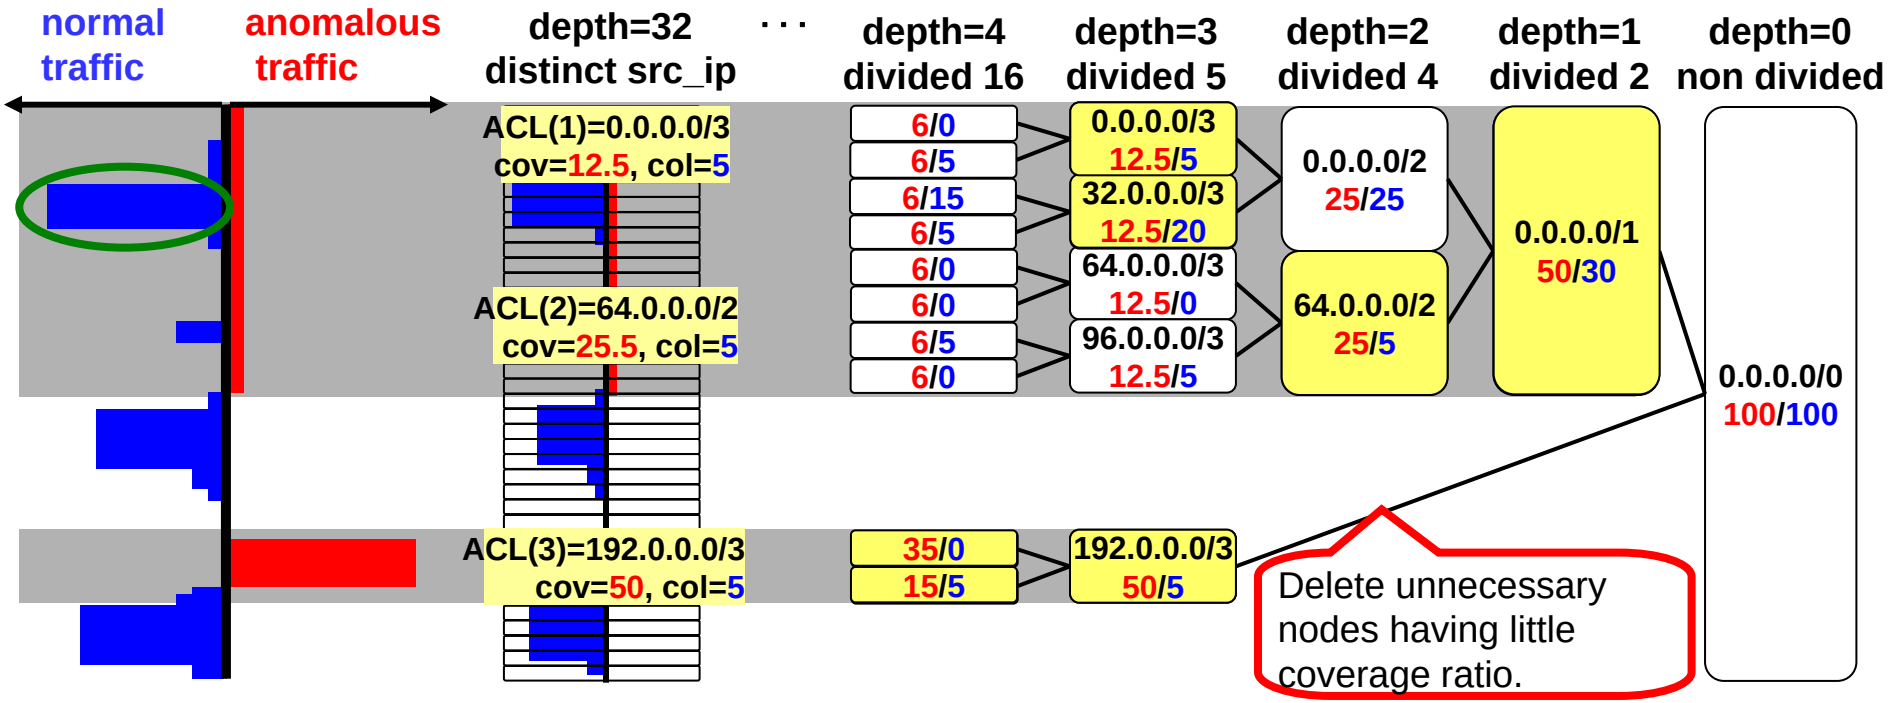
Making Traffic Tree

- Build up from individual source IP addresses (depth=32).
- Each node has information about coverage and collateral ratio.
  - Collateral ratio: normal traffic of the node ÷ total normal traffic
  - Coverage ratio: anomalous traffic of the node ÷ total anomalous traffic
- Make parent nodes by merging child node information.



Step 3: Selecting Best Node Set from Traffic Tree

1. To reduce search space, delete unnecessary nodes.
   - Unnecessary node: node which having little coverage ratio (little anomalous traffic) or little difference from its descendant nodes.
2. Search best node combination by applying the formula for all non-overlap node combinations in a brute force way.
   - Best node combination = **Best ACL set for source IP dimension**



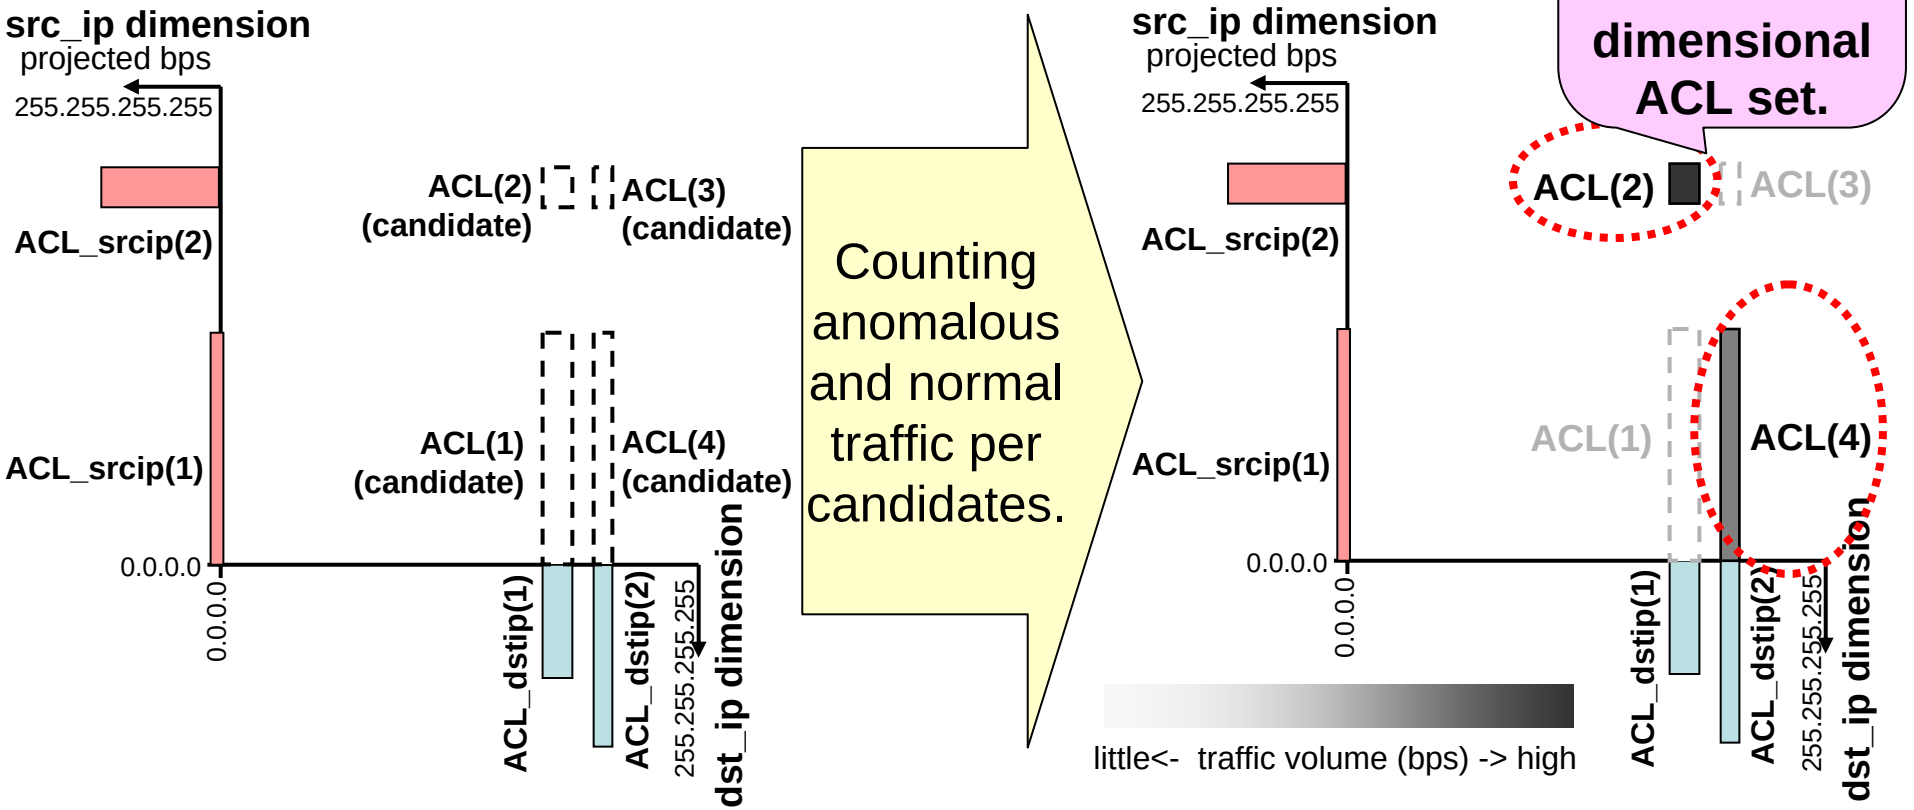
Example 2. rate= 0.71 : Can cover all the anomalous traffic, but too much collateral damage

rate= 1.44 : Can cover all the anomalous traffic with little collateral damage

rate= 0.71 : Can cover all the anomalous traffic, but too much collateral damage and must be aggregated to /3 (depth=3)

Concept #3 Generate Multi-Dimensional ACL Set

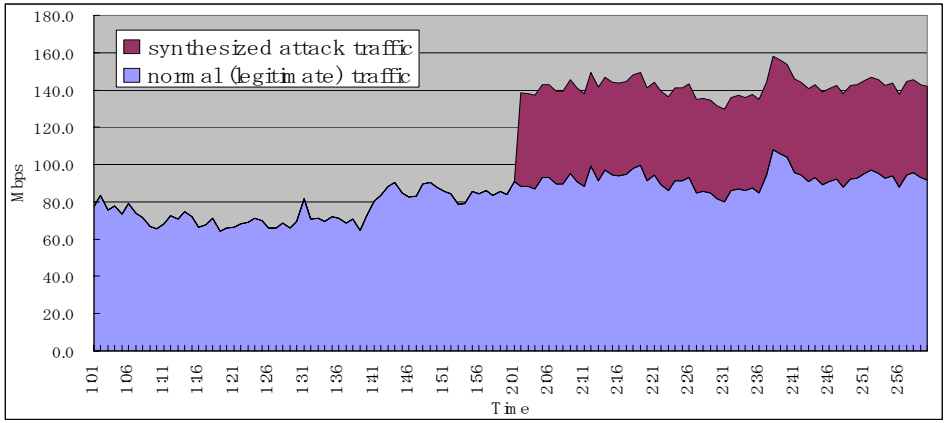
- Generate single dimensional ACL set in parallel.
  - ‘source IP’, ‘destination IP’ , ‘protocol’, ‘source port’ and ‘destination port’
- Make candidates of multi-dimensional ACL sets as a product sets of each dimension.
- Count anomalous/normal traffic for every candidates.
- Select best combination of candidates in terms of goodness score.



Evaluation and Results: Test Data Set

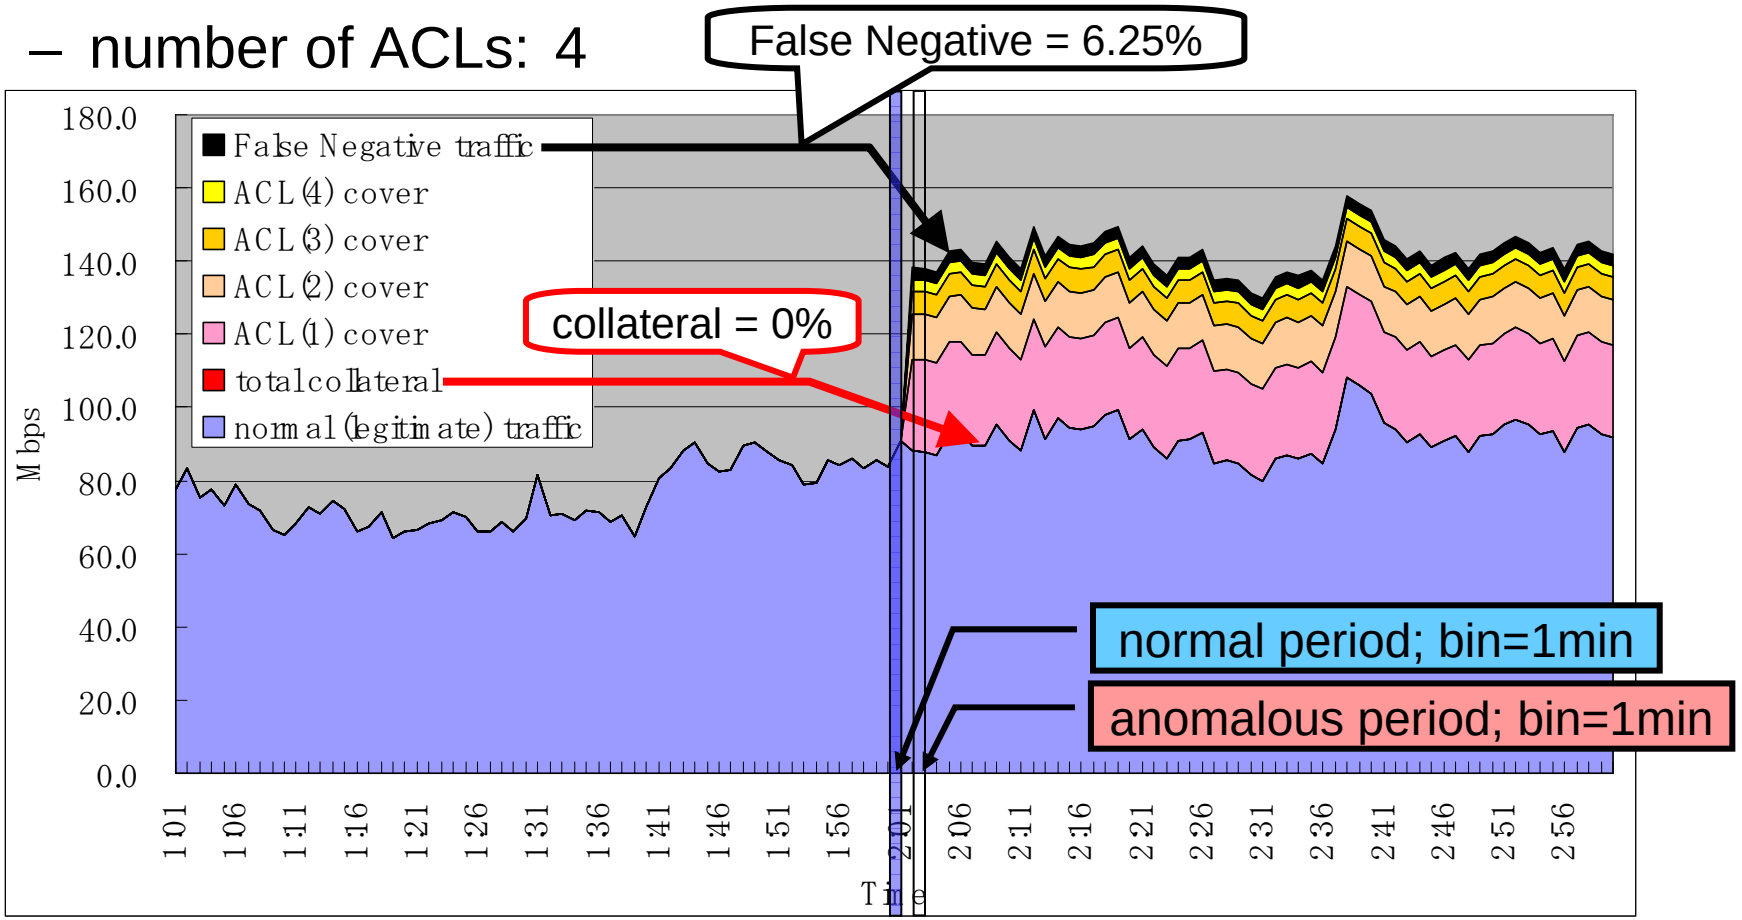
- Normal traffic: publicly available traffic data, captured on transpacific line (100 Mbps)
- Anomalous traffic: injected synthesized DDoS attack traffic
  - Mimic large DDoS attack
    - We choose source/destination IP addresses that have large normal traffic, because simple identification would cause collateral.
    - Destination: Popular server appeared in normal traffic
    - Source: Choose IP address block (/16) from which volume of normal traffic to the destination is largest.
- Test how well our technique can extract the injected anomalous traffic.

Use the "weighting coefficients"
-alpha=1 (weight for coverage)
-beta=10,000 (weight for collateral)
-gamma=0.0001 (weight for no. of ACLs)
to avoid collateral damage



Evaluation and Results: Results (1)

- Results: We get four ACLs (Four ACLs are one set.)
  - coverage: 93.75%
  - collateral: 0.00%
  - number of ACLs: 4



Time series of traffic with output ACLs displayed in separate colors

Evaluation and Results: Results (2) OUTPUT

basetime_len= 60.0 (sec) : (1168362060.0 - 1168362120.0) **basic information**
anomtime_len= 60.0 (sec) : (1168362180.0 - 1168362240.0)
base_total_bps= 89,121,539.5
anom_total_bps= 137,729,812.7
diff_total_bps= 48,608,273.2 +54.5 %

1-D_OUTPUT: PROTOCOL= 6 coverage= 100.42 collateral= 95.52 **single dimensional**
1-D_OUTPUT: SRC_PORT= high coverage= 108.27 collateral= 33.42 **identification**
1-D_OUTPUT: DST_PORT= high coverage= 100.09 collateral= 96.40 **results**
1-D_OUTPUT: **SRC_IP** coverage= 96.43 collateral= 0.00
119.170.0.0/17 coverage= 51.43 collateral= 0.00
119.170.128.0/18 coverage= 25.72 collateral= 0.00
119.170.192.0/19 coverage= 12.86 collateral= 0.00
119.170.240.0/20 coverage= 6.43 collateral= 0.00
1-D_OUTPUT: **DST_IP** coverage= 102.93 collateral= 2.17
134.45.182.70/32 coverage= 102.93 collateral= 2.17

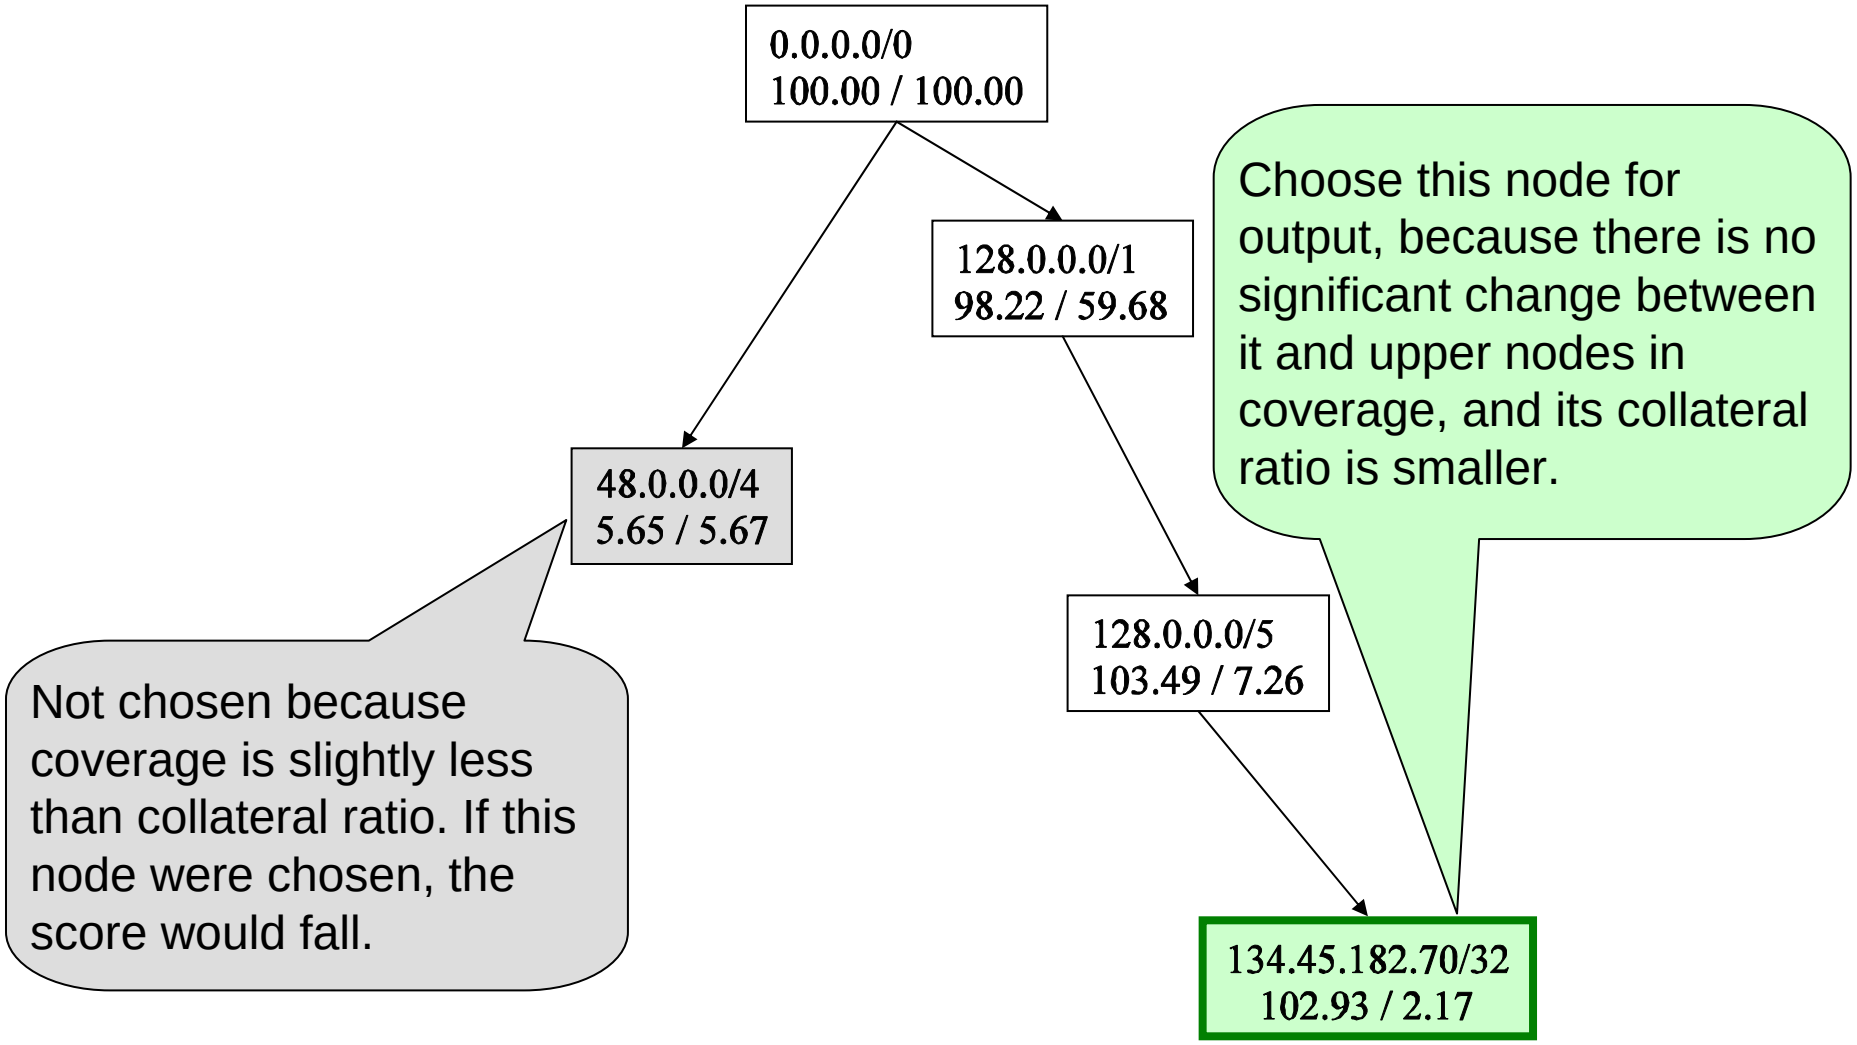
MULTI-DIMENSION_FLOW_OUTPUT coverage= 96.43 collateral= 0.00

flowID_0: cov= 51.43	col= 0.00:	**119.170.0.0/17**	**134.45.182.70/32**	6	high	high
flowID_1: cov= 25.72	col= 0.00:	**119.170.128.0/18**	**134.45.182.70/32**	6	high	high
flowID_2: cov= 12.86	col= 0.00:	**119.170.192.0/19**	**134.45.182.70/32**	6	high	high
flowID_3: cov= 6.43	col= 0.00:	**119.170.240.0/20**	**134.45.182.70/32**	6	high	high

↑ **coverage** ↑ **collateral**: ↑ src_ip ↑ dst_ip protocl scr_port dst_port

Evaluation and Results (3): Destination IP Tree

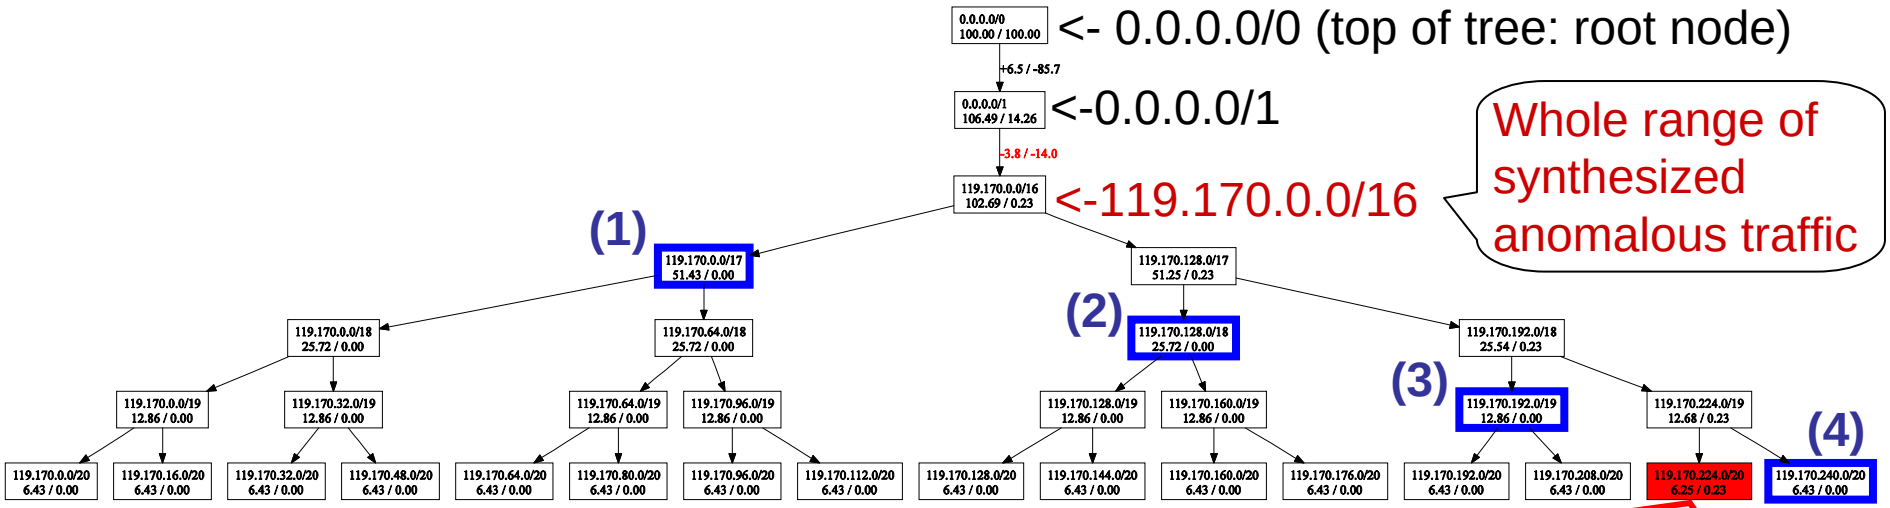
1-D_OUTPUT: **DST_IP** coverage= 102.93 collateral= 2.17
134.45.182.70/32 coverage= 102.93 collateral= 2.17



Copyright NTT Information Sharing Platform Laboratories, 2008

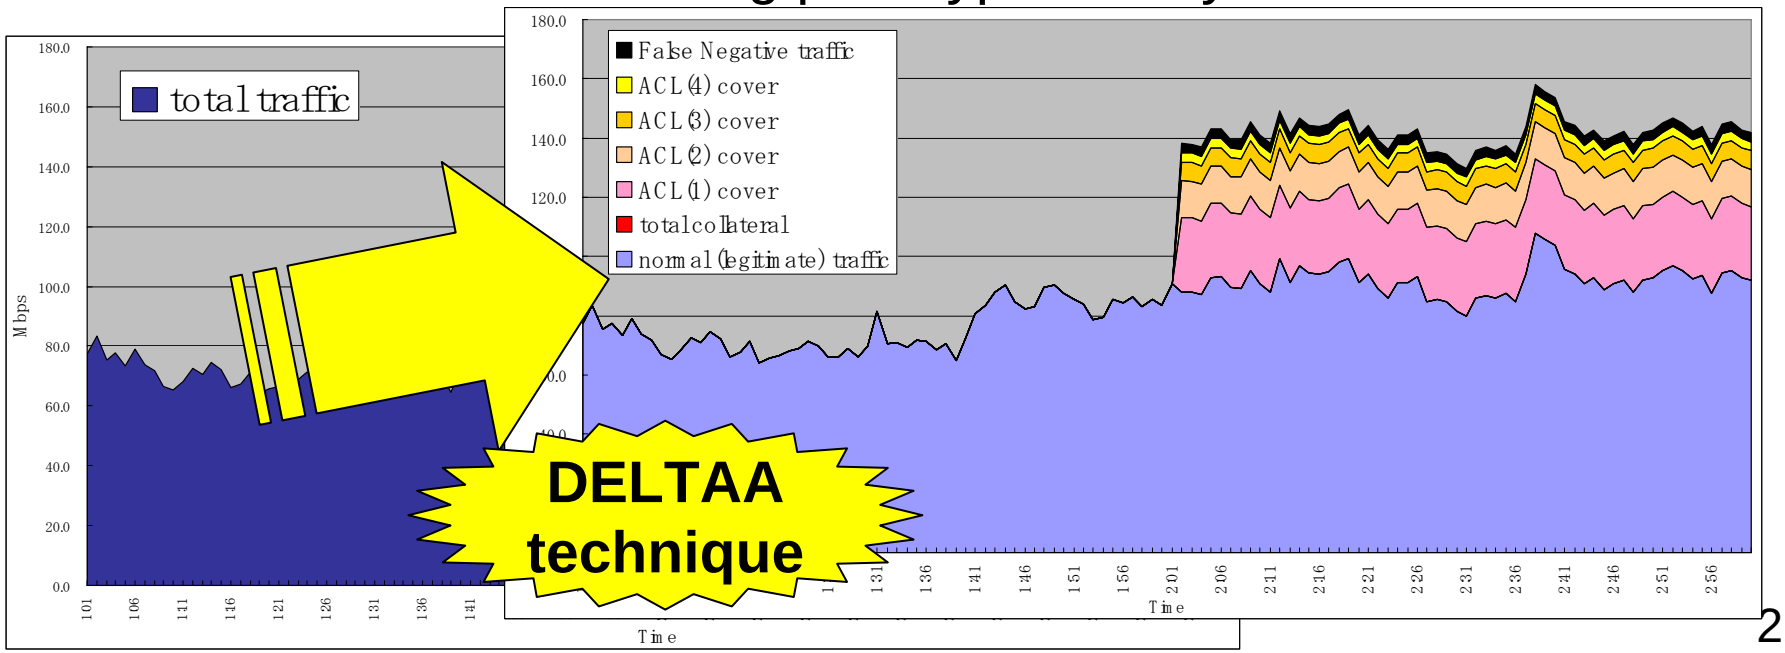
Evaluation and Results (4): Source IP Tree

1-D_OUTPUT: **SRC_IP**		coverage= 96.43	collateral= 0.00
(1)	**119.170.0.0/17**	coverage= 51.43	collateral= 0.00
(2)	**119.170.128.0/18**	coverage= 25.72	collateral= 0.00
(3)	**119.170.192.0/19**	coverage= 12.86	collateral= 0.00
(4)	**119.170.240.0/20**	coverage= 6.43	collateral= 0.00



Summary

- Introduced three criteria of optimal ACL set.
  - for mitigating DDoS attacks on router
- Proposed DELTAA technique: Optimizes trade-off among the these criteria, using normal and anomalous traffic.
- Presented an example of applying DELTAA to extract injected anomalous traffic.
  - Evaluation results using prototype and synthesized data set.



Thank you.

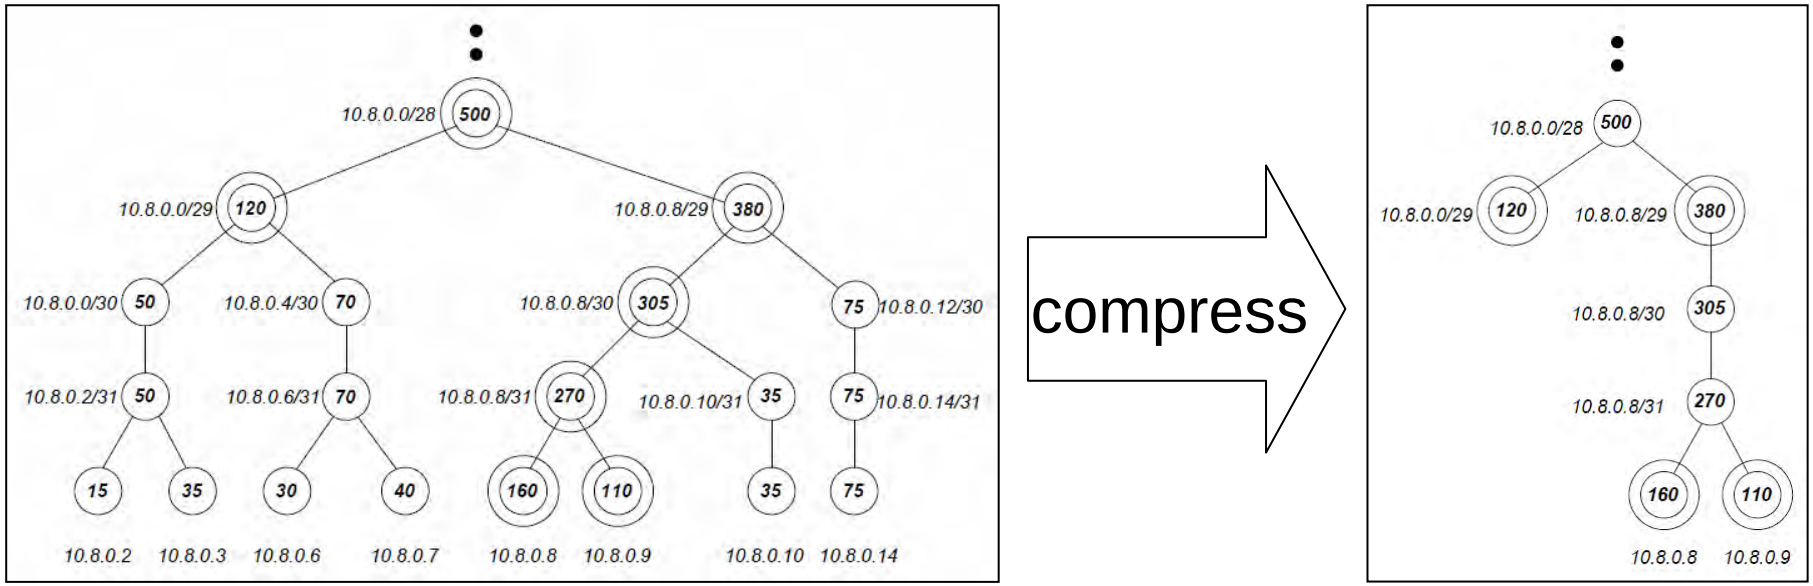
Any questions are welcome.

tsuyoshi.kondoh [at] lab.ntt.co.jp
ishibashi.keisuke [at] lab.ntt.co.jp

This study was supported by
the Ministry of Internal Affairs and Communications of Japan.

Q: Will Calculation Complexity Be Explosion?

- The way of making single dimensional tree and compressing way is similar to Estan's way in [Automatically].
- So, number of nodes on compressed tree is limited,
- We can Search all non-overlap node combinations in a brute force way within realistic time and resource.



[Automatically]: C. Estan, S. Savage and G. Varghese, "Automatically Inferring Patterns of Resource Consumption in Network Traffic," *SIGCOMM, August 2003.*

Copyright NTT Information Sharing Platform Laboratories, 2008



YAF

A Case Study in Flow Meter Design

presented at
FloCon 2008 - Savannah, Georgia

Brian Trammell
Technical Lead, Engineering
CERT Network Situational Awareness

© 2008 Carnegie Mellon University

YAF

Open-source, IPFIX-compliant bidirectional flow meter

- Available from http://tools.netsa.cert.org

Processes packets from multiple inputs

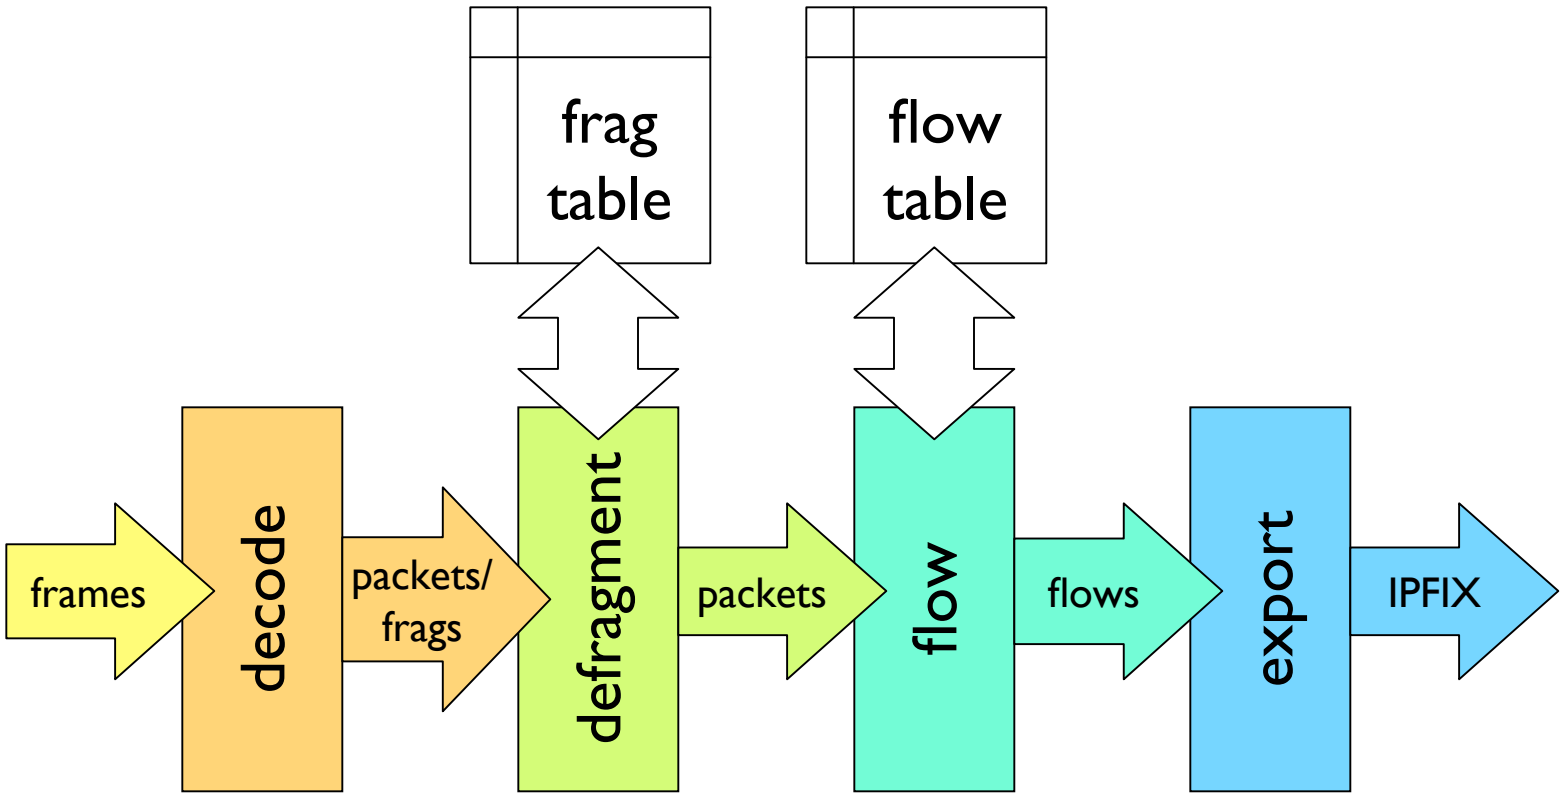
- libpcap dumpfiles (ad-hoc packet analysis)
- libpcap live capture (including proprietary pcap interfaces, e.g. Bivio)
- Endace DAG live capture

Performance is network hardware and I/O bound…

- …easily handles OC3, OC12, GigE at line speed, but
- 10GigE requires proprietary hardware at saturation.

© 2008 Carnegie Mellon University

Flow Meter Design



© 2008 Carnegie Mellon University

Flow Meter Effects on Flow Data

Fragmentation

End Conditions

Timeouts

Delta Counters

Biflows

The Packet Clock

© 2008 Carnegie Mellon University

Fragmentation

Three approaches for flowing fragmented traffic:

- pretend there's no such thing as fragmentation,
- drop all fragmented packets, or
- full or partial fragment reassembly

Each approach has tradeoffs, and is applicable in certain situations.

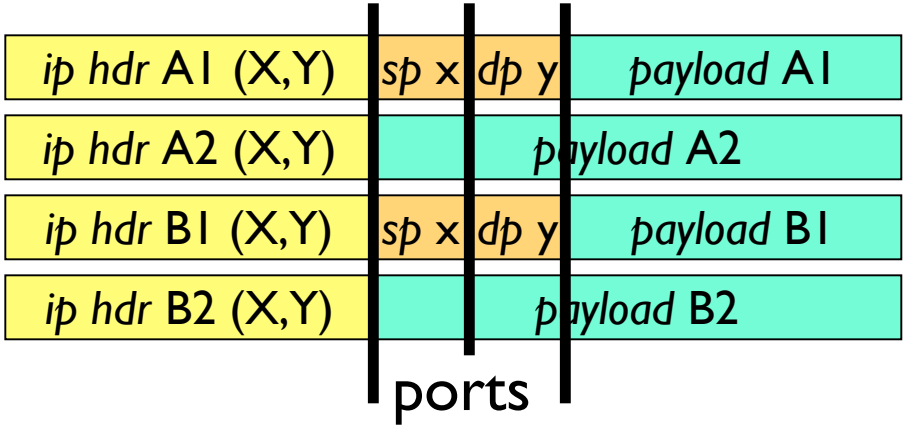
YAF supports partial reassembly.

© 2008 Carnegie Mellon University

Fragmentation?

Easiest way to handle fragmentation: don't.

Leads to inaccurate flow data as subsequent fragment port numbers are incorrectly decoded:



sip	*dip*	*proto*	*sp*	*dp*	*pkts*
X.X.X.X	Y.Y.Y.Y	6	x	y	2
X.X.X.X	Y.Y.Y.Y	6	$A2_0$	$A2_2$	1
X.X.X.X	Y.Y.Y.Y	6	$B2_0$	$B2_2$	1

© 2008 Carnegie Mellon University

Fragmentation? (2)

Often used in resource-restricted environments (e.g., routers).

- Much faster: no requirement even to recognize fragmented packets.
- Much less memory consumption: no fragment table.
- Less susceptible to resource exhaustion attacks.

Trivially easy to implement.

Difficult or impossible to recover actual flows from random fragment offset port data.

© 2008 Carnegie Mellon University

Dropping fragmented packets

Requires minimal resources at flow meter:

- need to recognize fragments, but not store them.

Leads to meter blindness:

- all an attacker must do to hide from the measurement infrastructure is fragment all packets.

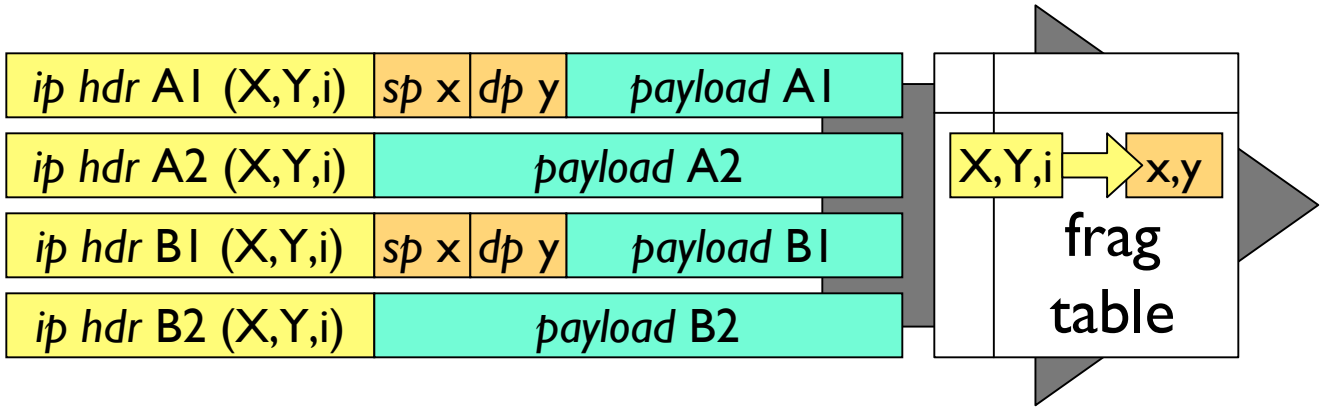
Only applicable behind perimeter devices which also drop all fragmented packets.

```
sip             dip             proto sp    dp    pkts
                        [no flows]
```


© 2008 Carnegie Mellon University

Partial fragment reassembly

Associate each fragmented packet with its actual transport ports:



sip	*dip*	*proto*	*sp*	*dp*	*pkts*
X.X.X.X	Y.Y.Y.Y	6	x	y	4

© 2008 Carnegie Mellon University

Partial fragment reassembly (2)

Accurately assigns fragments to respective flows.

Requires additional resources at flow meter:

- need to recognize, look up, and store every fragment.

More difficult to implement and maintain.

Requires care to avoid vulnerability to resource exhaustion attacks.

© 2008 Carnegie Mellon University

Flow End Conditions

Flow meter must recognize actual connection shutdown…

- …through varying degrees of modeling the host TCP state machine.

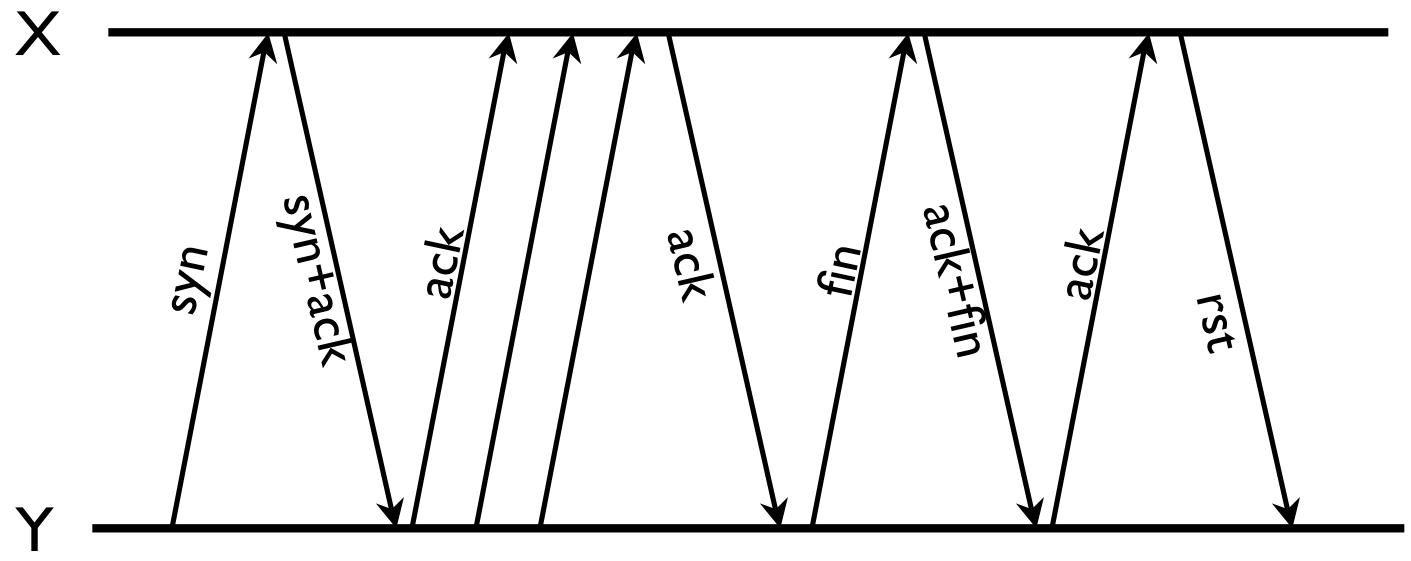
Flows on the wire are not always so well-behaved.

Example: multiple-RST teardown.

© 2008 Carnegie Mellon University

Multiple RST teardown

How many flows here?



sip	*dip*	*flags*	*sp*	*dp*	*pkts*
Y.Y.Y.Y	X.X.X.X	SAF	x	y	6
Y.Y.Y.Y	X.X.X.X	SAF	y	x	3
Y.Y.Y.Y	X.X.X.X	R	y	x	1

© 2008 Carnegie Mellon University

Multiple RST teardown (2)

Tempting to group RSTs on teardown into original flow…

- …how long to keep closed flow state?
- …how far to take this RST grouping?
- …how to communicate new configuration parameters to analysts?

YAF stays predictable, at the expense of generating multiple flow records for this behavior.

© 2008 Carnegie Mellon University

Passive Timeouts

Flows which have no packets over $TO_{passive}$ seconds are closed.

Necessary to terminate flows for all non-connection-oriented transports,

- i.e., anything but TCP.

Longer passive timeouts consolidate low-frequency periodic activity into fewer flows.

Shorter passive timeouts reduce flow table resource consumption for such activity.

© 2008 Carnegie Mellon University

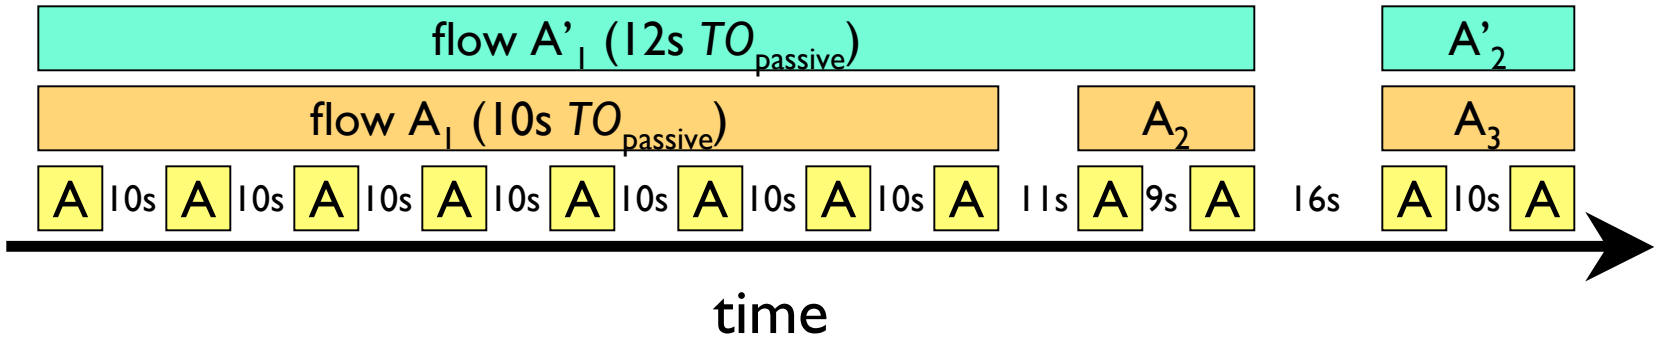
Passive timeouts (2)

Generally chosen to match common protocol timeouts…

- … which are generally round numbers, e.g., 10, 30, 60 sec.

May be chosen to avoid flow closure ambiguity due to minor variations:

- e.g., 12, 33, 64 sec.



© 2008 Carnegie Mellon University

Active Timeouts

Flows which have been open for TO_{active} seconds are closed.

- Maximum flow duration is TO_{active} seconds.

Necessary to ensure long-lived flows are eventually flushed from the flow table.

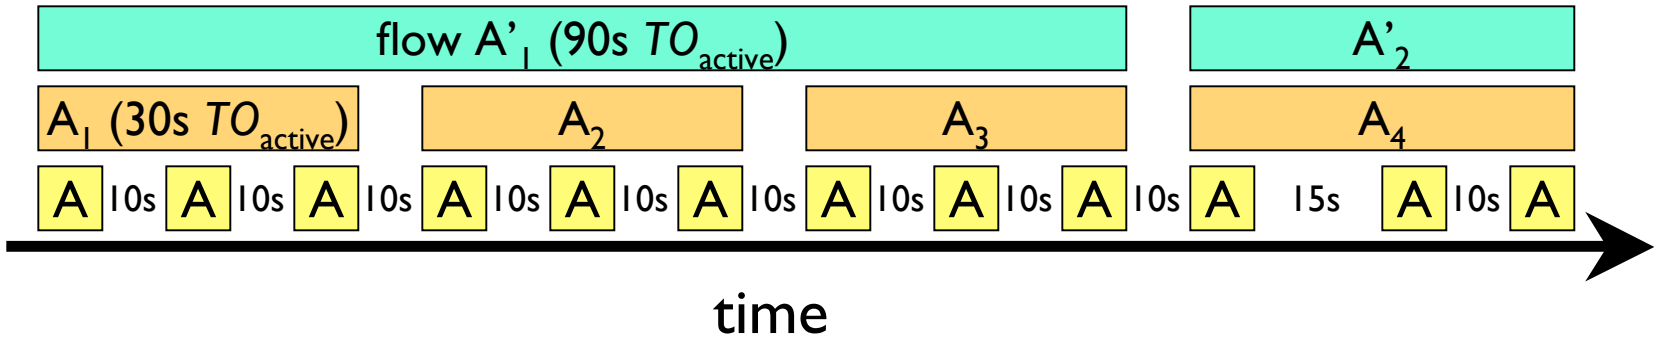
Active timeout determines reporting delay.

© 2008 Carnegie Mellon University

Active Timeouts (2)

Shorter active timeouts used for more rapid reporting.

Longer active timeouts used for better data reduction.



© 2008 Carnegie Mellon University

Delta Counters

Flow meters which periodically emit multiple flow records per flow (for rapid reporting) may use total or delta counters.

Total counters replace values in previous flow records.

Delta counters add to values in previous flow records…

- …thereby reducing state requirements on meter and increasing them on collector.

YAF uses total counters, but doesn't emit multiple records per flow…

- …uses active timeout instead.

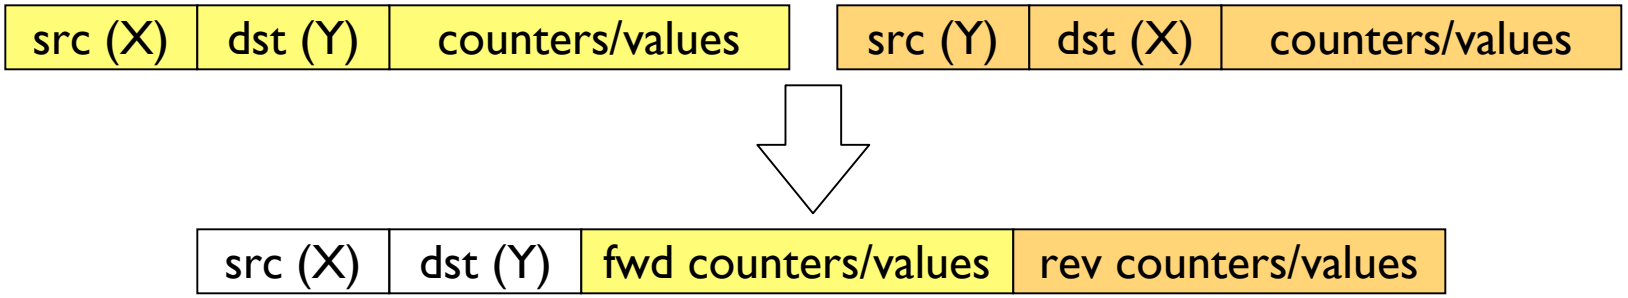
© 2008 Carnegie Mellon University

Biflows

Representation of two sides of a connection with a single flow record:

- Allows additional data reduction
- Enables easier connection analysis
- Improves flow state modeling at flow meter

YAF is a biflow meter, but SiLK stores uniflows.



© 2008 Carnegie Mellon University

The Packet Clock

Important to drive all processes within a flow meter with a single clock

- fragment timeouts, flow timeouts, time stamping, etc.

When building a flow meter, gettimeofday(2) is not your friend.

- often a problem with porting host-based software into a network-based monitoring environment

Use the timestamp from the packet instead!

- ensures that the resulting flow stream identical whether captured live or generated from dumpfile.

© 2008 Carnegie Mellon University

Getting YAF

http://tools.netsa.cert.org

Builds on Mac OS X, Linux, BSD, Solaris

- Bug reports from these or other Unices welcome!

Some prerequisites

- glib-2.0 (C modernization layer)
- libairframe (application utility library from NetSA)
- libfixbuf (IPFIX protocol implementation from NetSA)
- libpcap (generally available on most modern Unices)
- libdag (only required for Endace DAG capture)

© 2008 Carnegie Mellon University

Questions?

Ask now…

…or later:

- Brian Trammell <bht@cert.org>
- Chris Inacio <inacio@cert.org>

© 2008 Carnegie Mellon University

Network Analysis of Point of Sale System Compromises

Operation Terminal Guidance

Chicago Electronic & Financial Crimes Task Force

U.S. Secret Service



Outline

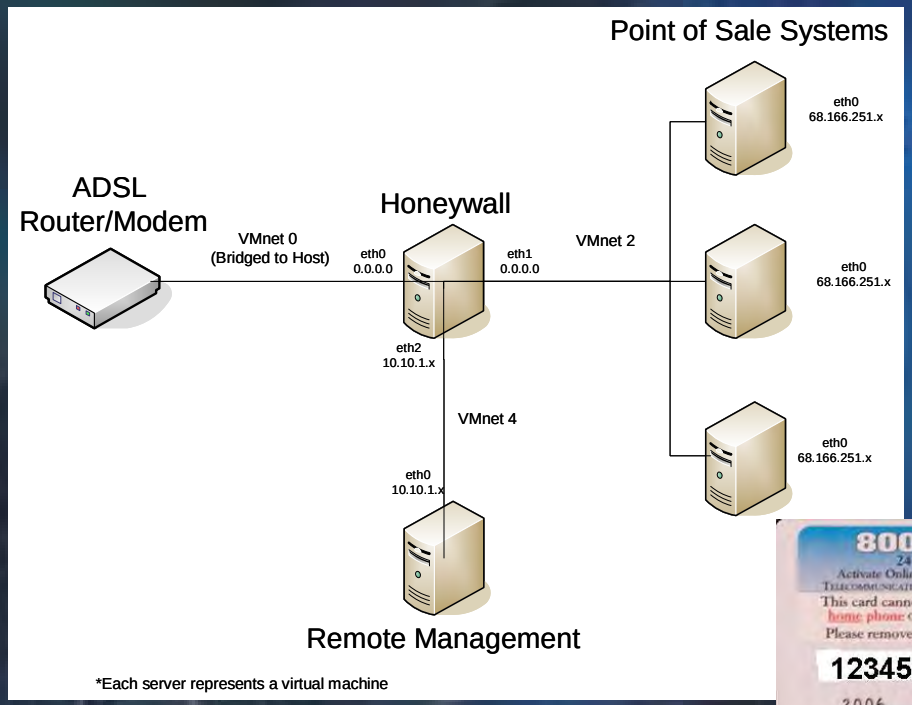
- Background
- Hypothesis
- Deployment Methodology
- Data Analysis
- Findings
- Discussion

Investigative Goals

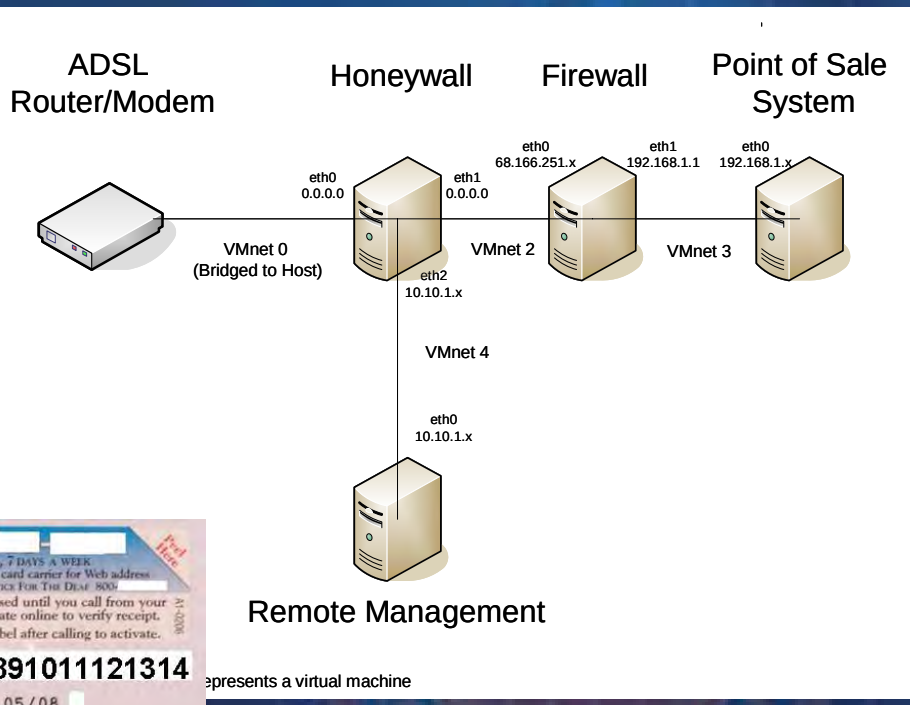
- Hypothesis: Remote attackers were not targeting point of sale (POS) system software, rather POS system compromises are a result of insecure deployment of the underlying operating system by automated scanning and vulnerability exploitation

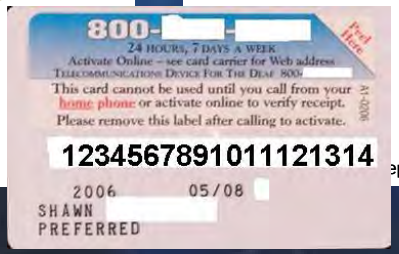
Deployment Methodology

Test Group Honeynet



Control Group Honeynet



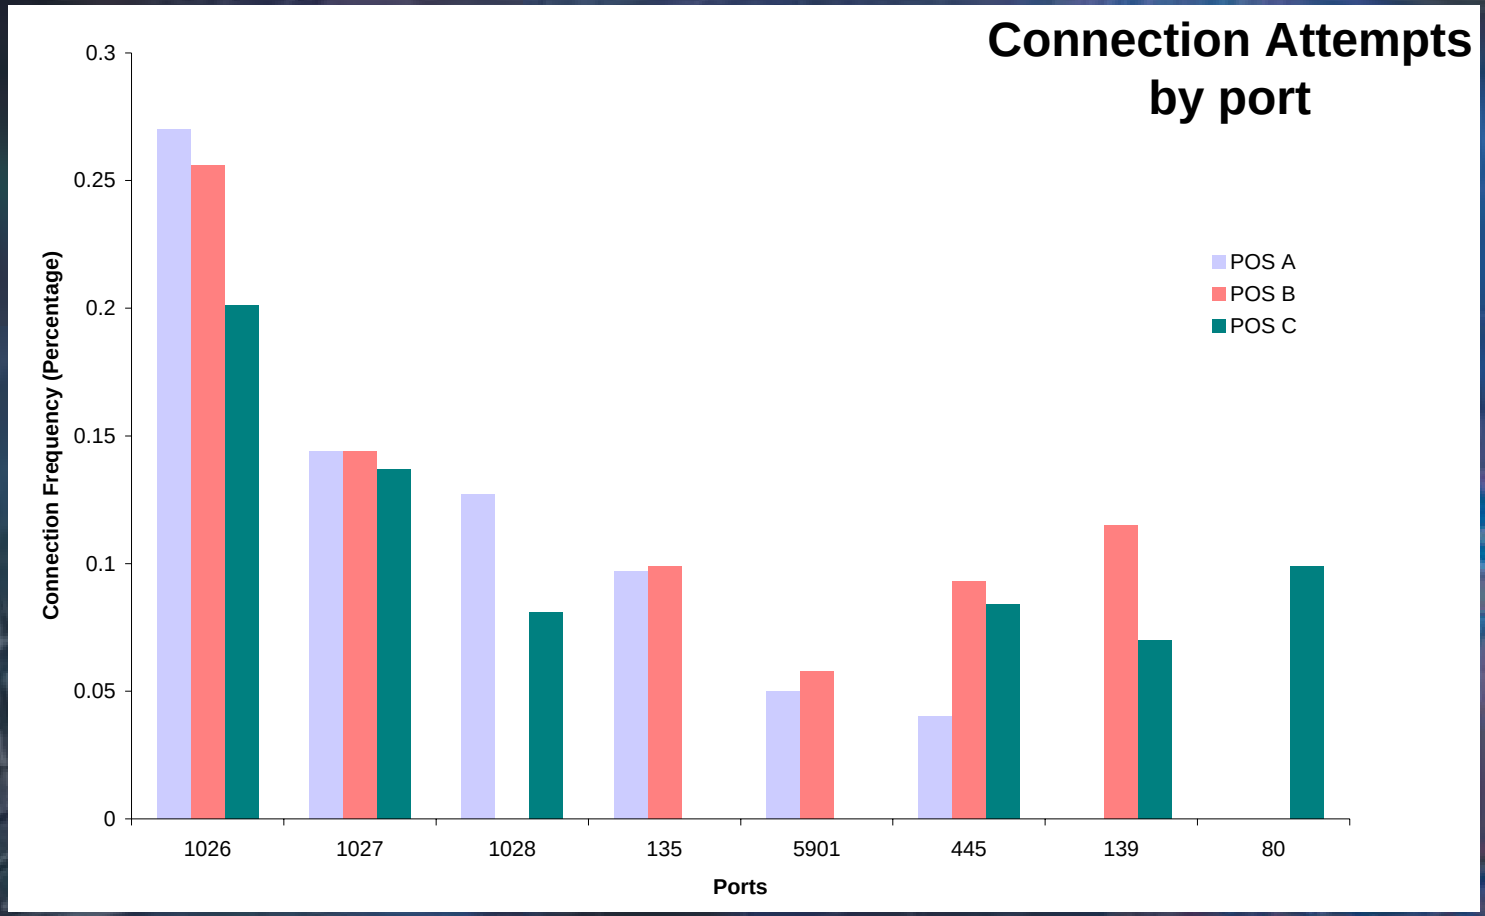


Honeytoken

U.S. Secret Service

Data Analysis

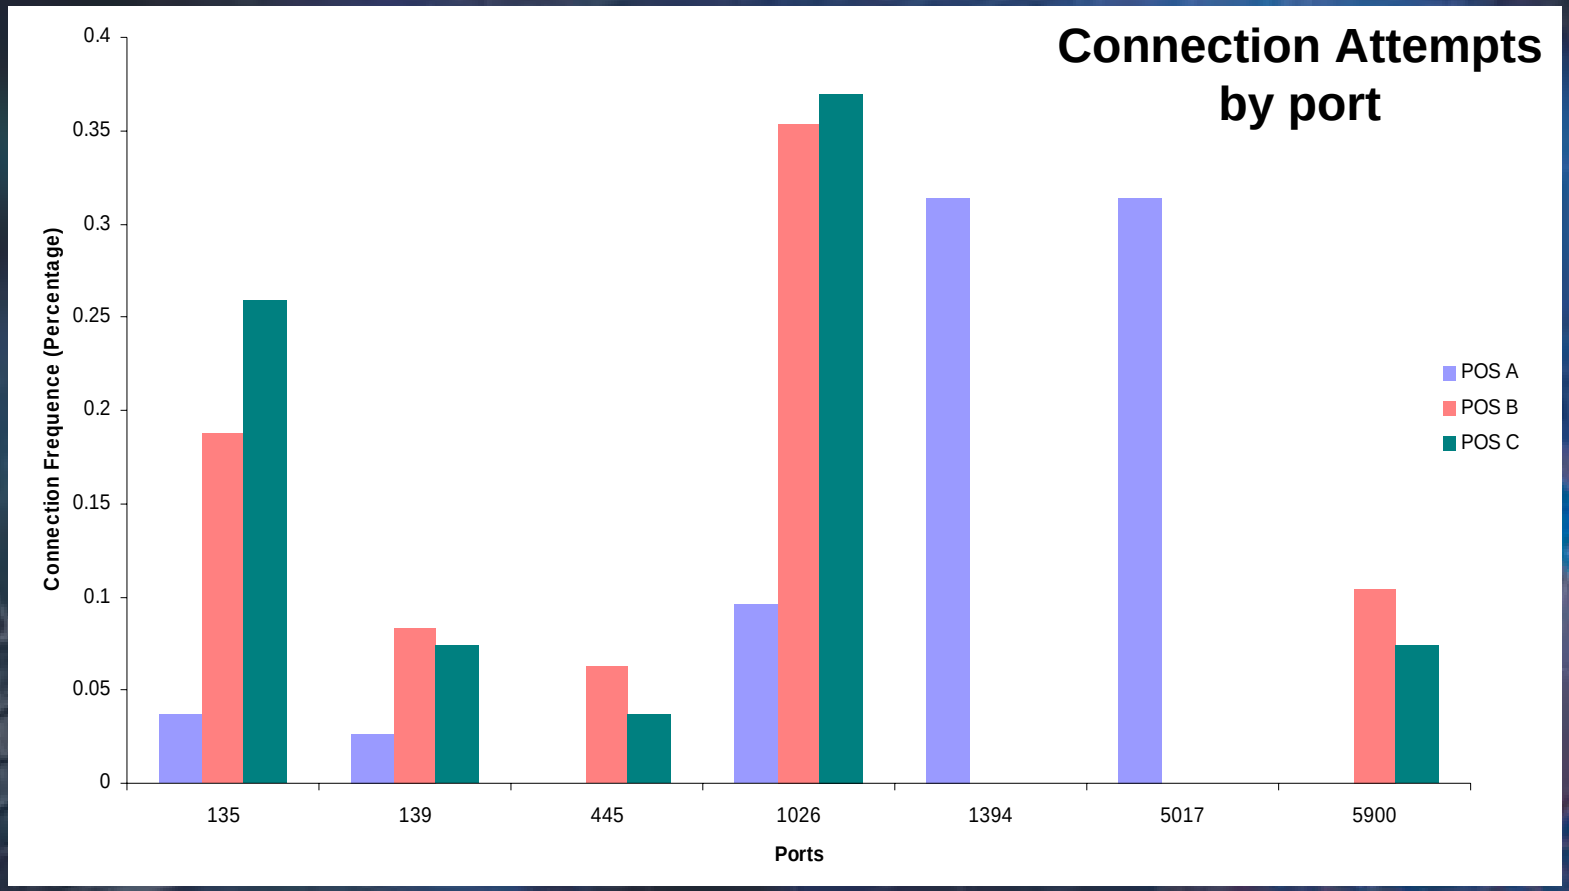
Control Group



U.S. Secret Service

Data Analysis

Test Group



U.S. Secret Service

Data Analysis

- Association rules
  - Clustering
    - T: Number of virtual POS systems with connection attempts from a single source
    - n_i: Number of packets from a source to a virtual POS system
    - N: Total number of packets from a source to all three POS systems
    - $N=\sum n_i$

$$\text{Support}(R) = \frac{\#\text{ connections (POS system A, B, and C)}}{\#\text{connections}}$$

Data analysis methodology from F. Pouget and M. Dacier. "Honeypot Based Forensics."

U.S. Secret Service

Data Analysis

Control Group Clusters

Port	Item Sets	Support %	Support % > 1%
80	Cluster 1: T=1, N=3 Cluster 2: T=1, N=1 Cluster 3: T=2, N=8 (n=5, n=3)	43.5% 10.9% 4.3%	1
135	Cluster 4: T=1, N=1 Cluster 5: T=1, N=2	54.5% 22%	2
139	Cluster 6: T=1, N=2 Cluster 7: T=1, N=3	75% 10.1%	1
445	Cluster 8: T=1, N=1 Cluster 9: T=1, N=2 Cluster 10: T=1, N=3	20% 70% 7.1%	2
1026	Cluster 11: T=1, N=1	53.5%	1
1027	Cluster 12: T=1, N=1	98%	1
1028	Cluster 13: T=1, N=1	83%	1
5901	Cluster 14: T=1, N=2	90.9%	1

Data Analysis

Test Group Clusters

Port	Item Sets	Support %	Support % > 1%
445	Cluster 1: T=2, N=34	22.2%	0
1026	Cluster 2: T=2, N=3 Cluster 3: T=3, N=3 (n=1,n=1, n=1) Cluster 4: T=1, N=1	1.8% 20% 50.9%	2
1394	Cluster 5: T=1, N=12 Cluster 6: T=1, N=15 Cluster 7: T=1, N=6 Cluster 8: T=1, N=9	20% 16.7% 1.7% 16.7%	3
2967	Cluster 9: T=3, N=8 (n=2, n=3, n=3) Cluster 10: T=3, N=30 (n=10, n=10, n=10)	10% 10%	0
5900	Cluster 11: T=3, N=3	20%	0

U.S. Secret Service

Data Analysis

- Edit Distance Analysis
  - Extract TCP payloads from previous identified cluster members
  - Compare packets from each IP address against all others identified through clustering

Source A	Source B
<mss E..0..@.o.A.;W\. D..s.]..... p...^2..... <mss E..0..@.o.A.;W\. D..s.]..... p...^2.....	<mss E..0.{@.k.l\=.y. D..s.....jd..... p..... <mss E..0.{@.k.l\=.y. D..s.....jd..... p.....

Attack Phrases

U.S. Secret Service

Data Analysis

Control Group Phrase Distance

Cluster	Port	Phrase Distance (Lines)	Std Deviation
Cluster 6	139	2	9
Cluster 7	139	1	5
Cluster 8	445	3	10
Cluster 9	445	5	8
Cluster 10	445	4	18
Cluster 11	1026	86	169
Cluster 13	1028	12	65
Cluster 14	5901	32	12

***Clusters 1,2, 3,4,5, and 12 were discarded as not statistically significant

U.S. Secret Service

Data Analysis

Test Group Phrase Distance

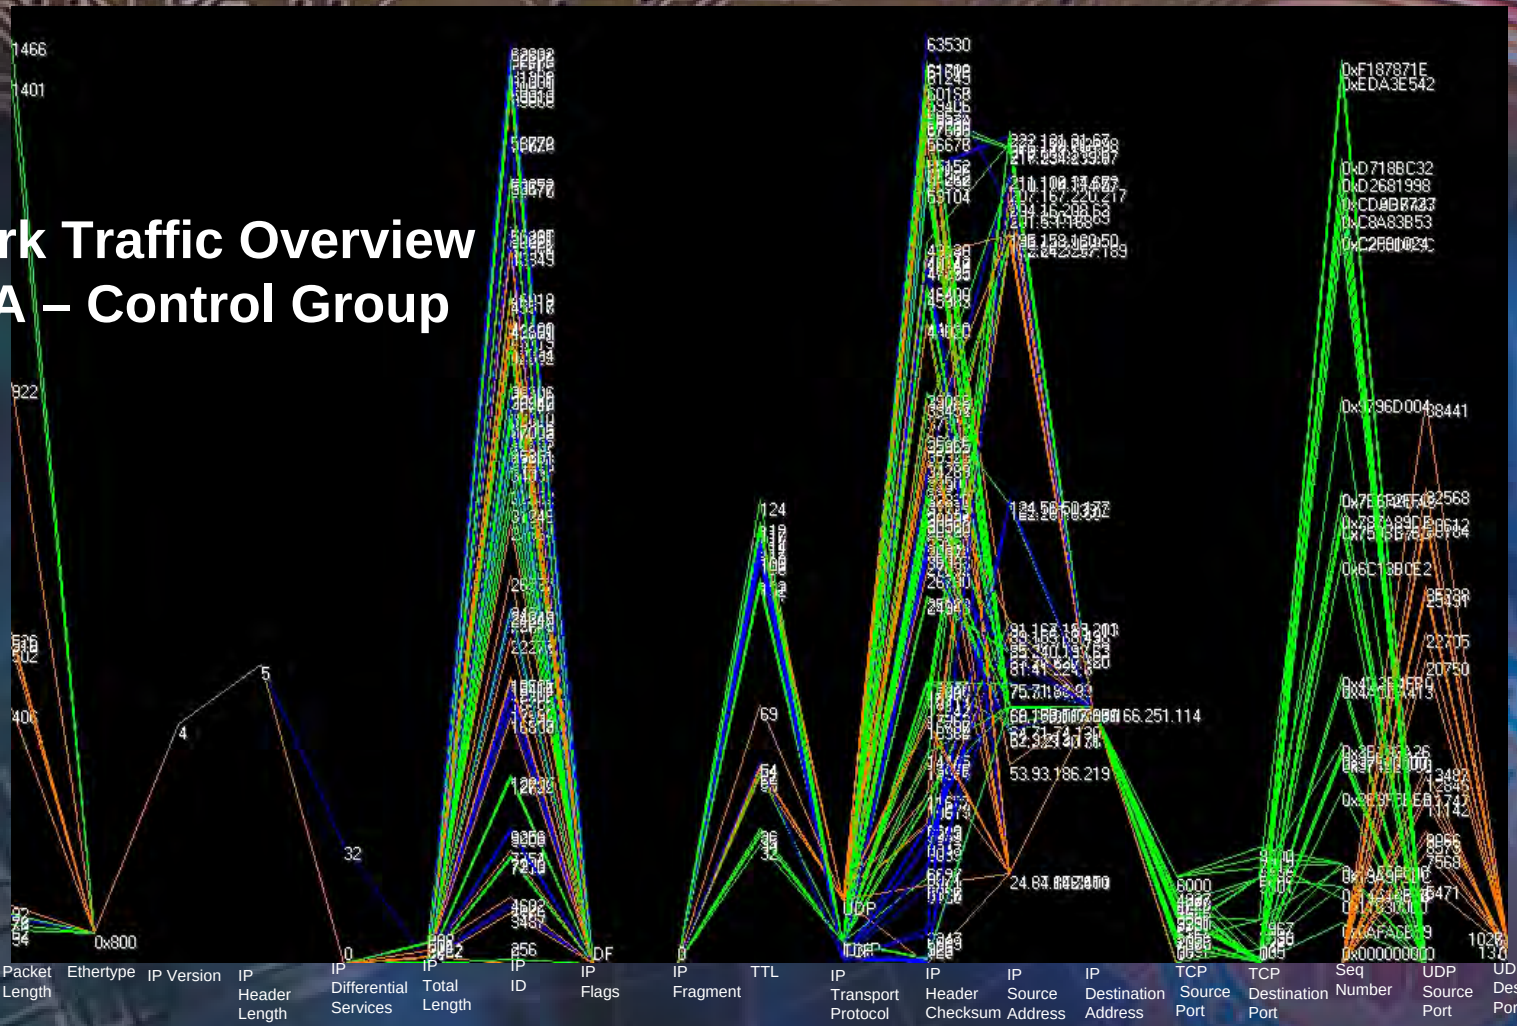
Cluster	Port	Phrase Distance (Lines)	Std Deviation
Cluster 2	1026	324	238
Cluster 5	1394	360	85
Cluster 6	1394	280	170
Cluster 7	1394	529	136
Cluster 8	1394	1422	1143
Cluster 11	5900	240	257

***Clusters 1,3,4,9,10 were discarded as not statistically significant

U.S. Secret Service

Data Analysis

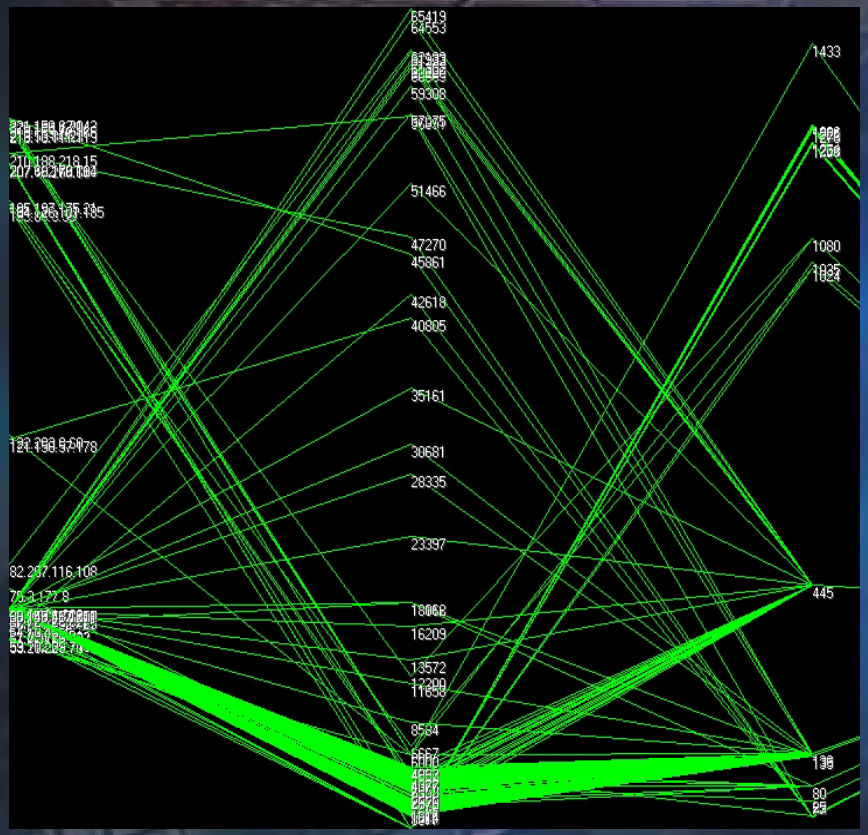
Network Traffic Overview
POS A – Control Group

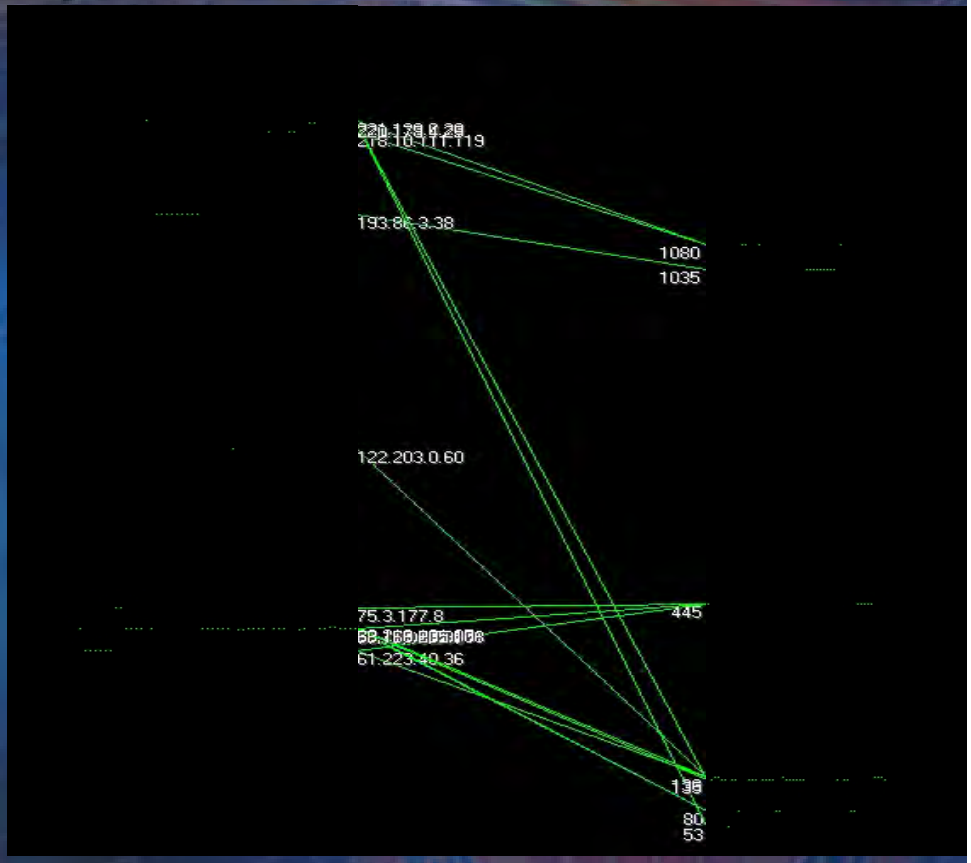


Visualization methodology from Greg Conti's. "Security Data Visualization."

U.S. Secret Service

Data Analysis

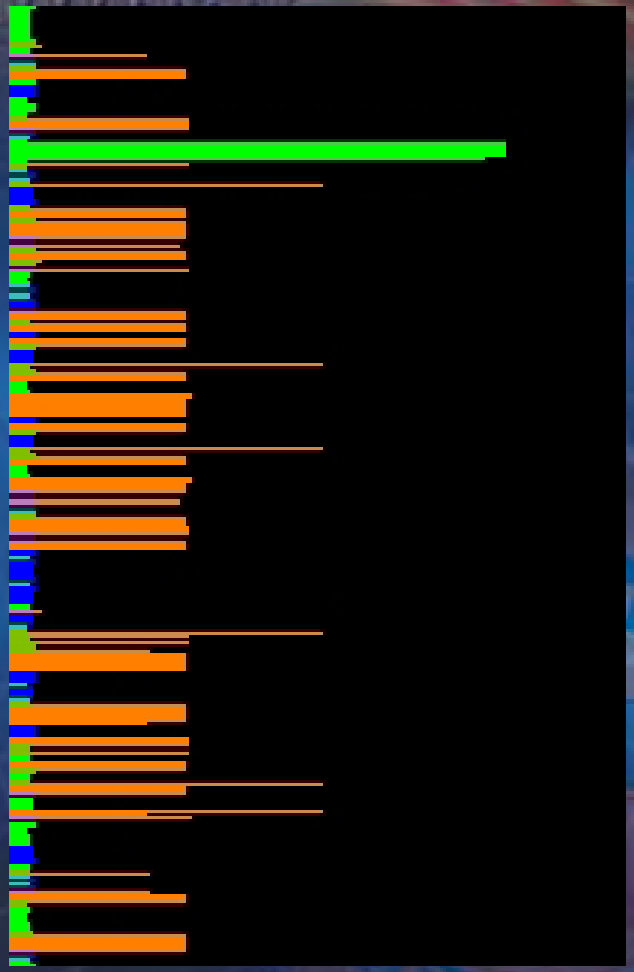




U.S. Secret Service

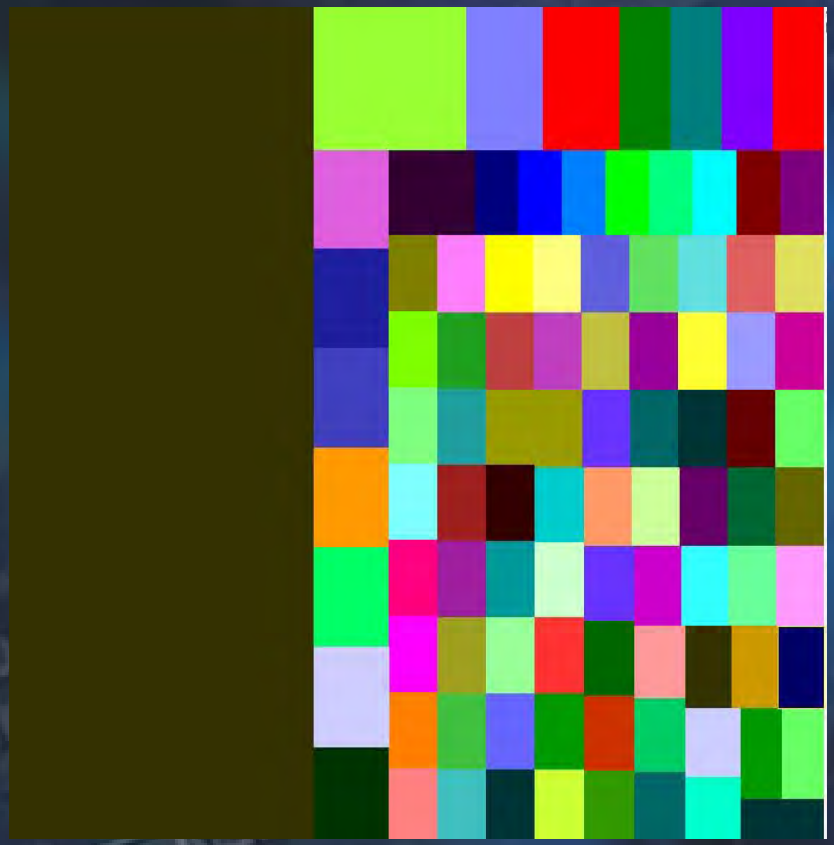
Data Analysis

- The TCP outlier is associated with browsing public web site to ensure connectivity
- Uniform length of packets

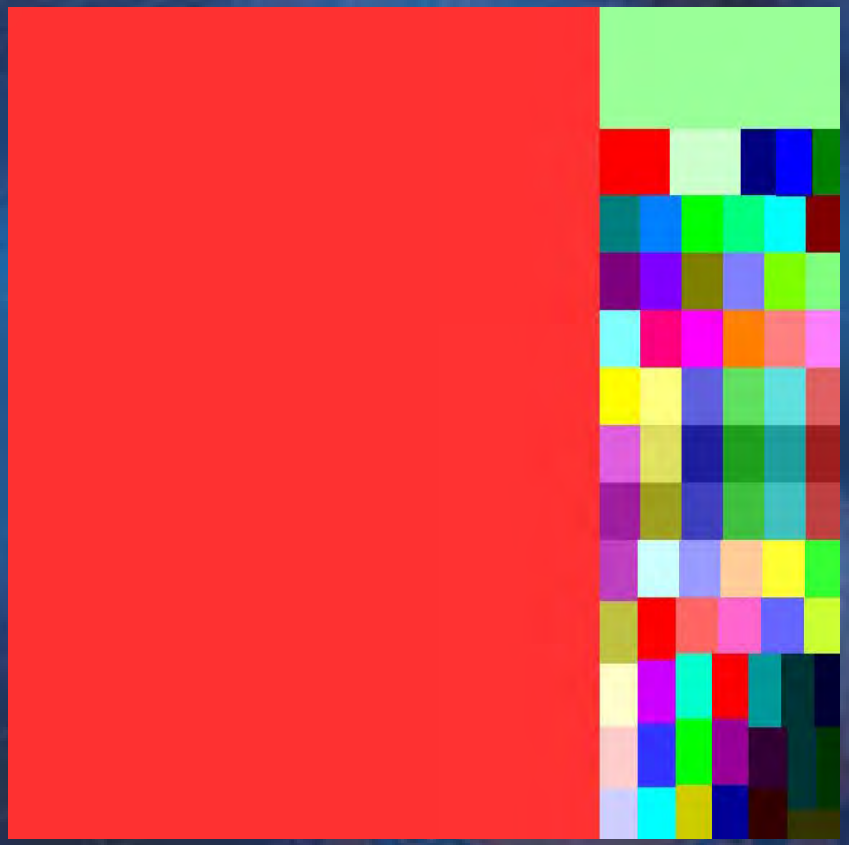


Data Analysis

TCP Packet Tree Map

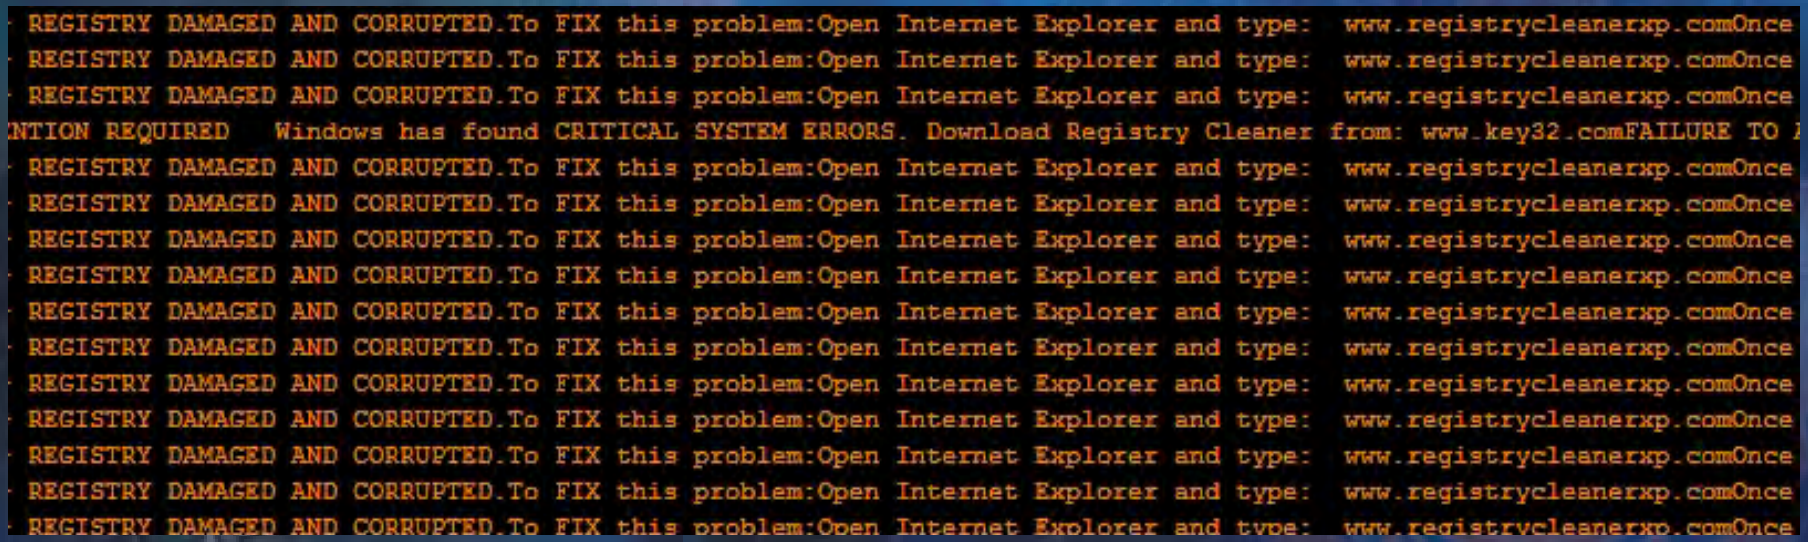


UDP Packet Tree Map



Data Analysis

- Examination of the UDP packets identified in the previous tree map revealed them to be spam targeting messenger applications



U.S. Secret Service

Findings

- Automated scanning of select set of ports
- Multiple exploits targeting multiple OS's from single source IP address
- Attackers not aware compromised system is a POS system until after compromise and exploit
- Insecure installation of operating system and applications lead to compromise

Discussion

All references available upon request

Ryan E. Moore
Special Agent
U.S. Secret Service
312-353-5431
ryan.moore@usss.dhs.gov

U.S. Secret Service

One Year of Peer to Peer

Ron McLeod, BCSc, MCSc.
Director - Corporate Development Telecom Applications Research Alliance
Doctoral Student, Faculty of Computer Science, Dalhousie University

Presentation Summary

This presentation will profile the result of the growth in peer-to-peer applications on a sample network and describe the resultant massive increase in the diversity of traffic. This diversity impacts the ability to profile baseline normative behaviour using Blind Flow Analysis.

I will also briefly discuss the application of SiLKtools, Neural Networks and Bioinformatic strategies to Blind Flow Analysis of real world security problems and how that analysis is affected by the growth in recreational/user driven applications.

What began as a basic design principal of end-to-end management with popular applications in recreational computing is quickly becoming a dominant evolutionary force in network traffic patterns.

Traffic patterns are becoming emergent properties influenced by the voluntary adoption of new systems by individuals without any collective intent.

The network is evolving at the edges.

“Peer-to-Peer is the basic design of the Internet” – Christian Huitema

Sample Network Description

- A Multi-tenant Commercial Network consisting of:
  - ~ 40 user assigned hosts, actual number subject to minor fluctuations over time.
  - ~40 special hosts not assigned to individual users. These hosts form parts of various temporary development and experimental environments.
  - Users were apprised that Network flow data was now being captured for experimental and management reasons.
  - Payload data was neither collected nor examined.
  - Analysts did not have access to the content of specific hosts for further investigation.
  - For confidentiality reasons the identity of the Network is not specified in this Presentation.

A Review of Blind Flow Analysis

The Need for Classification Based on Minimal Information (the extreme case in the world of tomorrow)

- Capturing and examining payload contents is widely viewed as a potential violation of privacy and placed in a category similar to listening in on a telephone call.
- Even attempts to use information derived from the payload (such as ngrams) do little to alleviate the fundamental concern of the user surrounding access to the payload.
- In multi-tenant commercial environments this user concern may be based in protection of commercial confidentiality.
- There is less (although not zero) concern among the user community with regard to the capture and investigation of packet header data (some concern for Source and Destination IP's and MAC's).
- Therefore, the network analyst may be limited to examining a severely reduced subset of the packet header information in an attempt to determine if the system under their management (or monitoring) is operating properly or experiencing anomalous behavior.
- The loss of access to the originating address information means that the analyst no longer has access to a unique field in the data that identifies the individual hosts in the traffic (i.e. they cannot tell one computer from another by looking at the remaining flow record traffic alone).
- In such an environment, what is required is a method of classification that relies on minimal information and the development of traffic flow behaviour models that use only this information.

One Strategy for Comparing A Suspicious Host to a Standard Workstation Using Blind Flow Analysis

Local Baseline Workstation Behaviour (BWB)	**Suspicious Host**
Bytes Transferred in one month < 20 million per month	45 billion per month
Internal DIPs < 10 per month	3 per month
External DIPs < 20 per month	1.74 million per month
Protocols: 1 < 2 %	1 1 %
6 > 70 %	6 9 %
17 < 30 %	17 90 %
Number of Protocols < 5	3

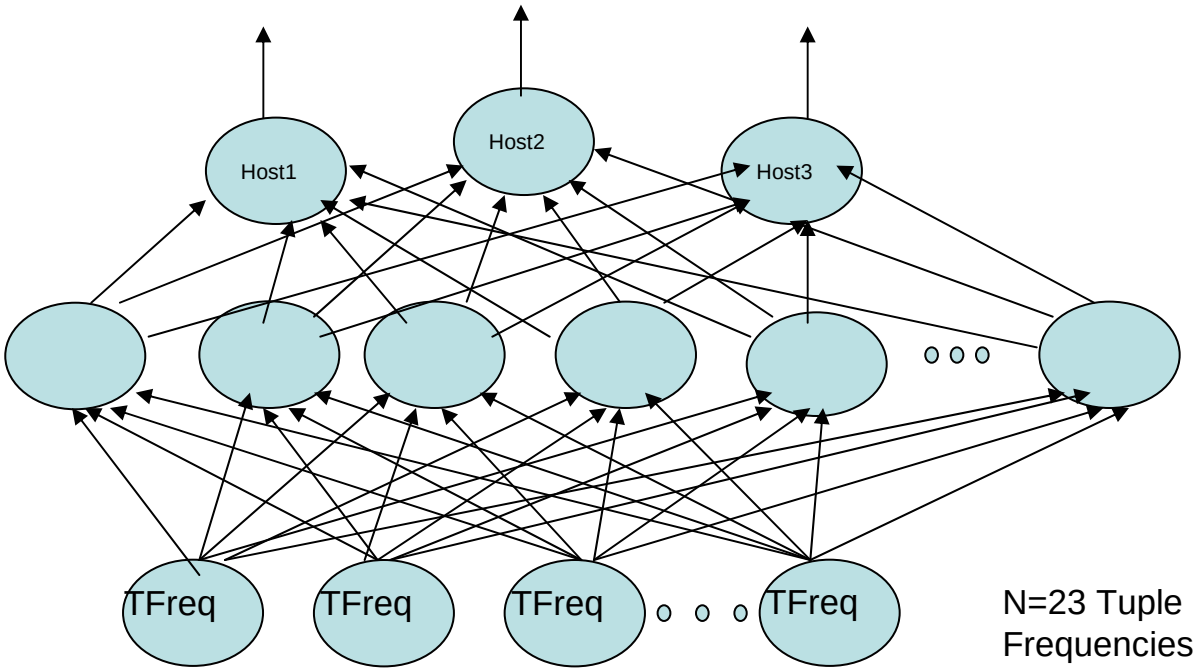
Port Number Range	# of Ports Accessed	%of Ports Accessed	%of Total Bytes Traffic	# of Ports Accessed	%of Ports Accessed	%of Total Bytes Traffic
<1024	< 7	20-50%	<1%	45	0.07%	
1024-5000	< 10	>30%	>90%	3,976	6%	1%
>5000	< 5	<20%	<9%	60,059	93%	99%

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information

- In early 2006 Neural Network was used to classify workstation traffic based on a localized “Workstation Genome”.
- It was found workstation behaviour could be fully described by a set of 23 unique 3-tuples formed by the combination of Protocol, Destination Port, and Byte Range ID – Where Byte Range ID was one of five levels given by:

Bytes	Range
0 – 100	1
100 – 999	2
1000 – 9,999	3
10,000 – 49,999	4
50,000 +	5

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information



Each input frequency vector contains an observed frequency for each 3-tuple for a 24 hour period.
Each 3-tuple is defined as Protocol, Destination Port, Byte Range.
All observed Workstations could be described by a 23 element Vector.

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information

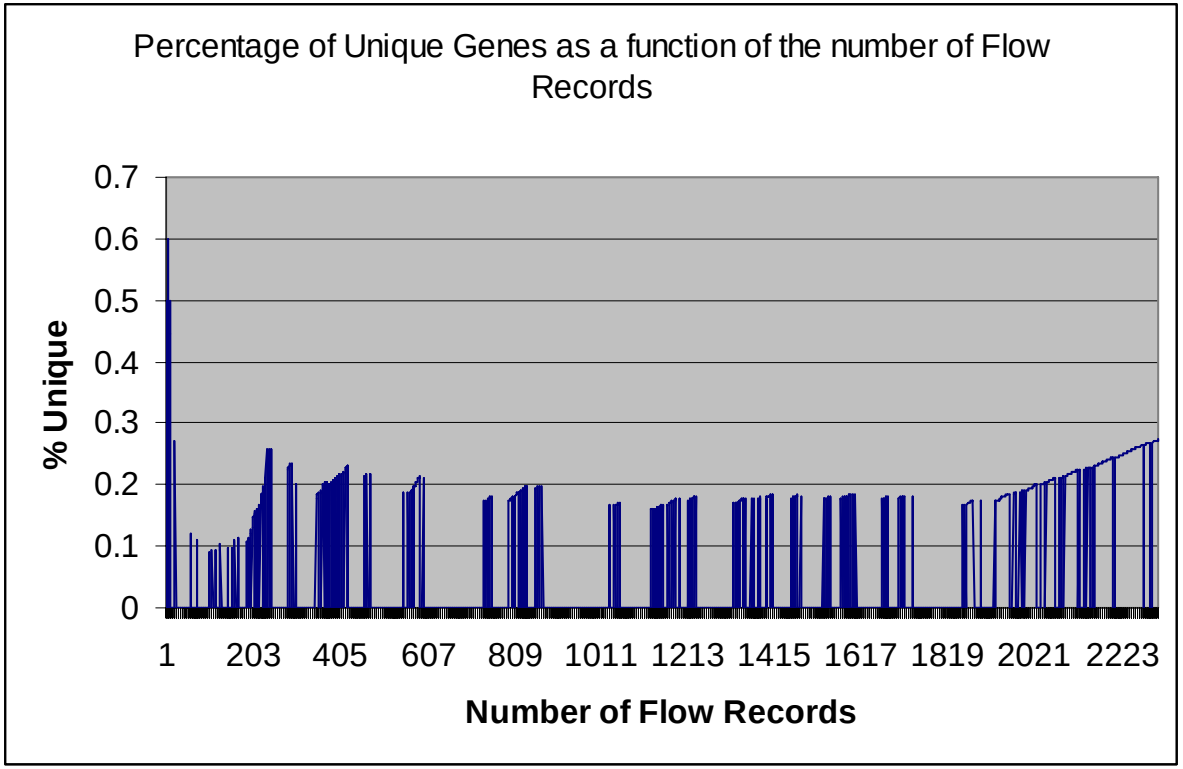
Host ID	Day	Output Vector	Classification (Hit/Miss/Unknown
1 [0 1 0]	1	[0.04 0.86 0.08]	HIT
	2	[0.17 0.97 0.00]	HIT
	3	[0.10 0.91 0.02]	HIT
	4	[0.09 0.95 0.01]	HIT
2 [1 0 0]	1	[0.95 0.06 0.00]	HIT
	2	[0.96 0.04 0.00]	HIT
	3	[0.95 0.06 0.00]	HIT
	4	[0.95 0.07 0.00]	HIT
3 [0 0 1]	1	[0.00 0.09 0.92]	HIT
	2	[0.00 0.00 0.99]	HIT
	3	[0.00 0.12 0.92]	HIT
	4	[0.00 0.00 0.99]	HIT

100% Success rate on uniquely classifying a small sample of the population

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information

- In early 2007 a similar population of workstations was chosen with the goal of testing a Support Vector Machine approach to classification.
- *To the great surprise of the author, the number of unique 3-tuples required to uniquely describe the Workstation Genome had risen from 23 to over 600 in 16 months.*
- Subsequent investigation showed that the diversity of the observed behaviour increased as a function of both population size as well as the length of the sampling period.

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information



By limiting the traffic to ICMP and TCP flow records, the number of unique tuples required to adequately describe the population reached a steady state of approximately 18% of the total number of all expressed tuples.

When UDP traffic was introduced into the sample, the percentage of unique tuples in the population did *not* reach a steady state in proportionality but rather the number of the unique tuples increased in linear proportion to the number of total tuples observed.

Impact of Peering Traffic on Blind Flow Analysis and the Uniqueness of Minimal Information

- What happened to the network traffic to create such diversity in such a short period of time?

- Expected monthly unique destination IPs =1200 (40 hosts * 30 external and internal DIP contacts).

  Actual values:

  Average monthly destination IPs = 140,000
  Average monthly number of flows = 2.8 million
  Average monthly byte volume of approximately 31 billion

- In addition to unusual volumes, two fundamental behaviours changed.
  - Protocol Ratio
    - From TCP 70% UDP 30%
    - To TCP 50% UDP 50%
  - Use of Unique Destination Ports by Workstations now parallels Server behaviour.

One Year of Peer-to-Peer

Much has been written lately of the growth and deployment of Peer-to-Peer Protocols

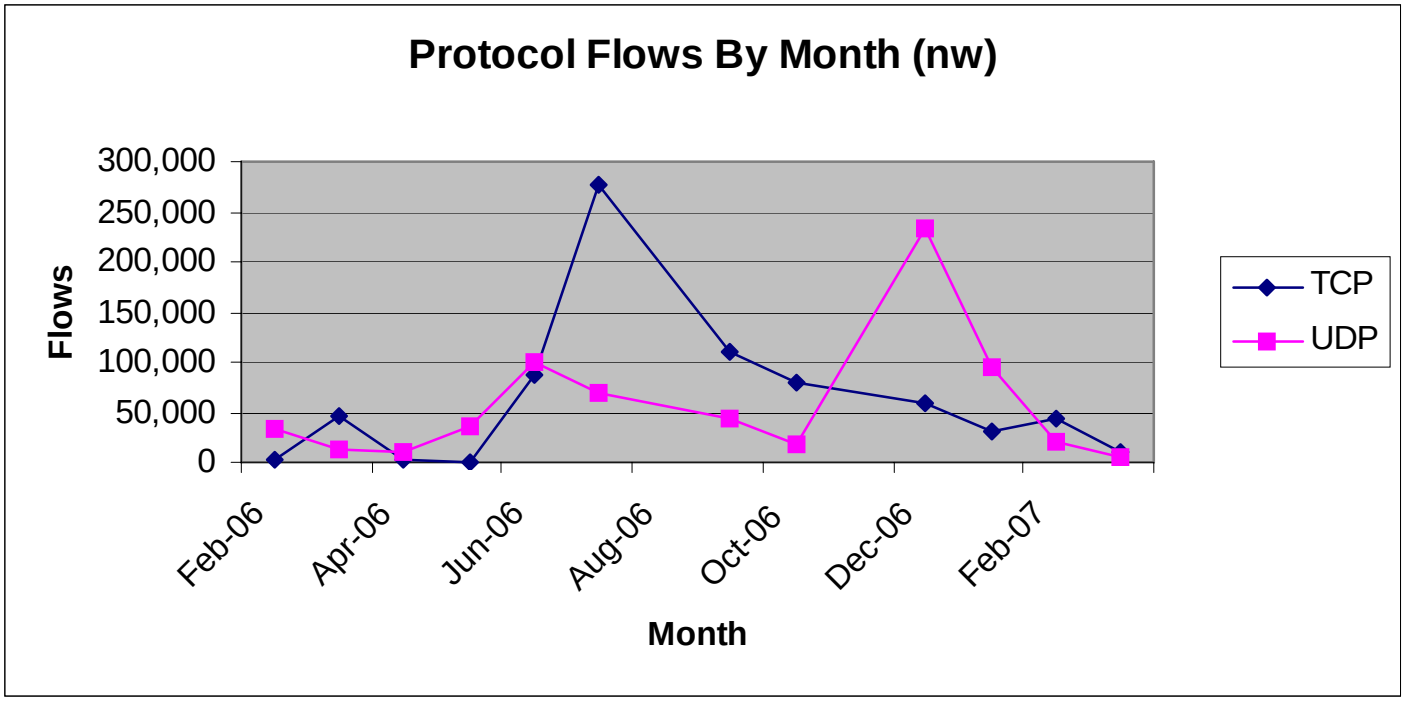
Recommended reading *“Transport Layer Identification of P2P Traffic”, Thomas Karagiannis, et al, IMC’ 04, 2004, Taorimina, Italy.*

Perhaps Peer-to-Peer is the culprit.

Decided to check for the presence of known P2P in the traffic

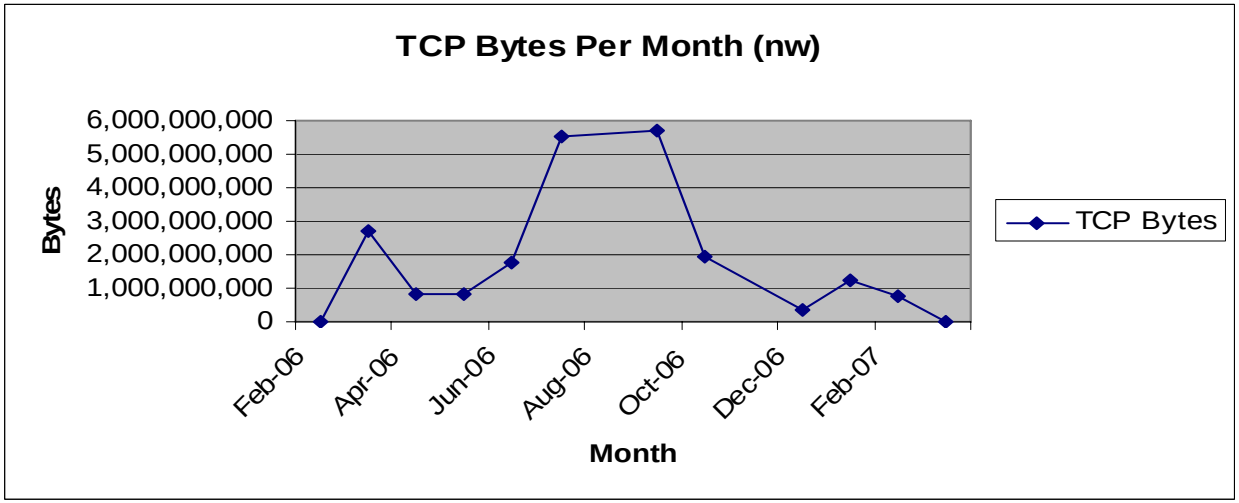
eDonkey2000
Fasttrack
Bittorent
Gnutella
MP2P

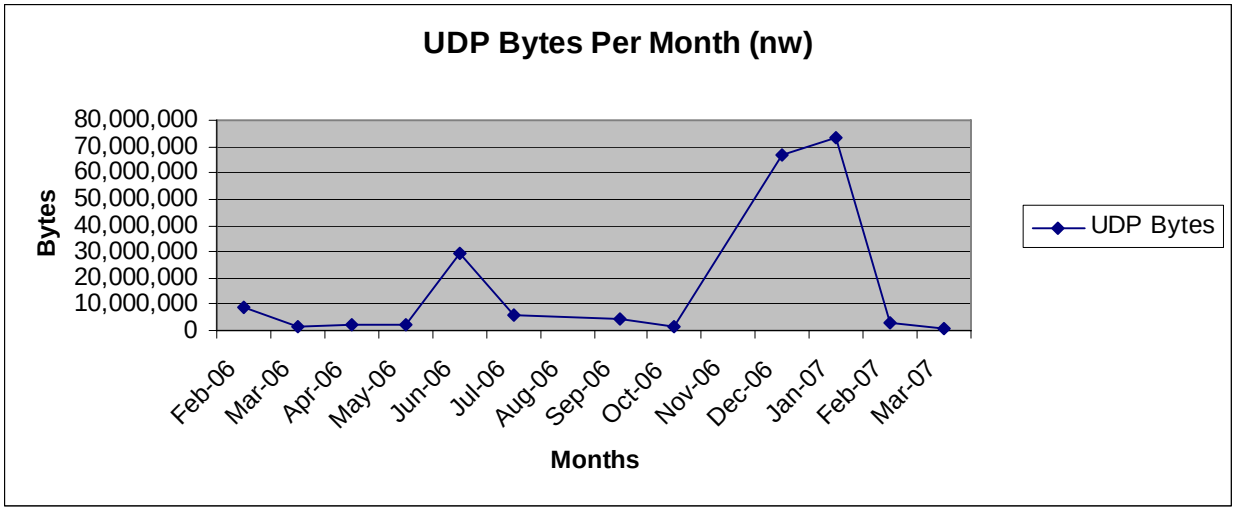
One Year of Peer-to-Peer



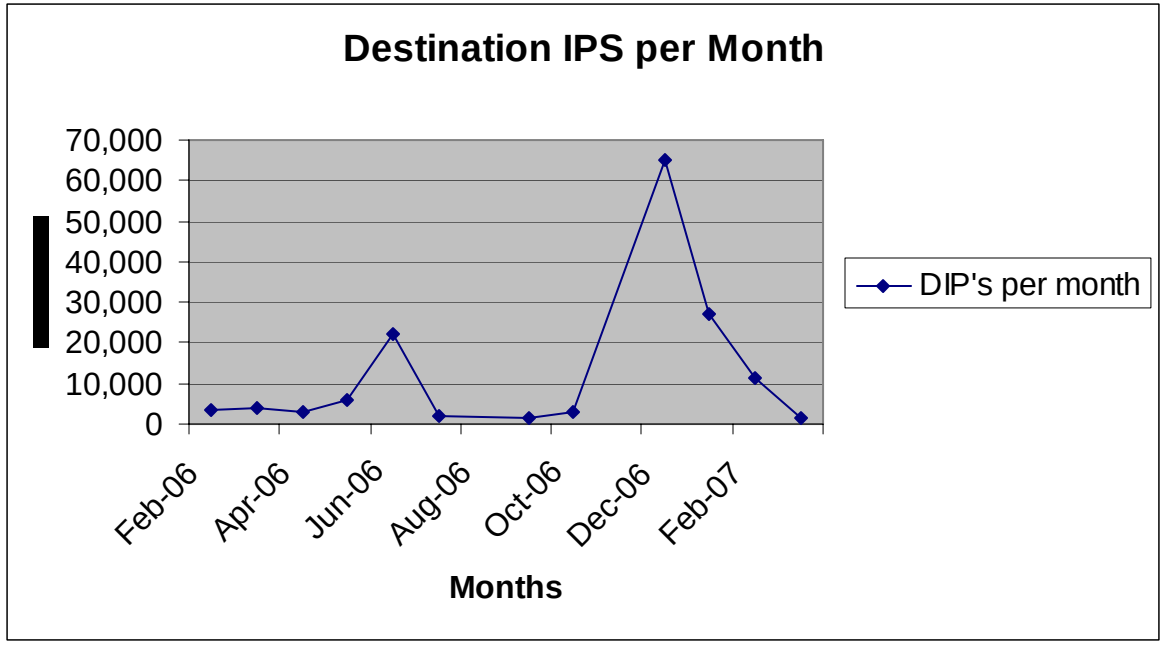
The graph above shows the pattern of flows by protocol for one year for the Target network.

One Year of Peer-to-Peer



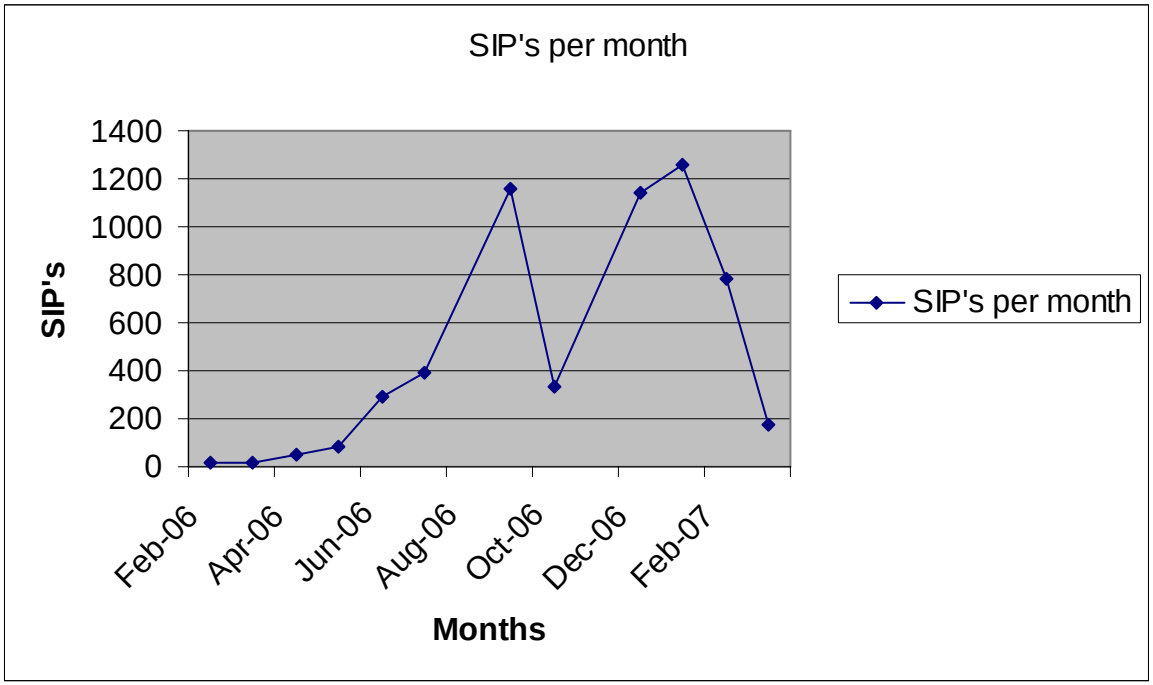


One Year of Peer-to-Peer



For a small network they talked to quite a few friends.

One Year of Peer-to-Peer

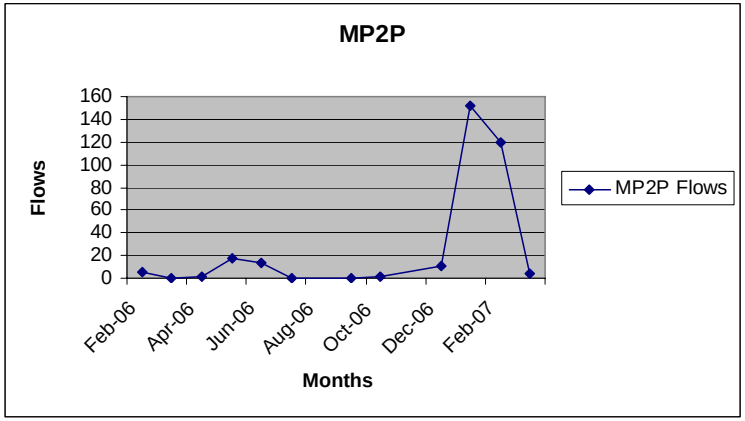


The feeling was mutual.

One Year of Peer-to-Peer

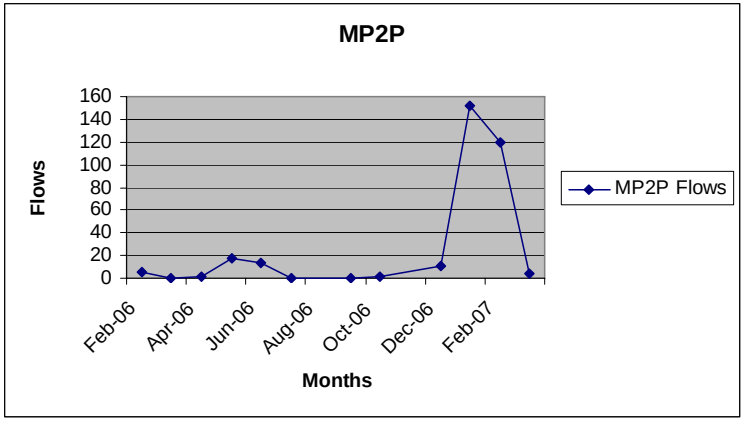
Let’s consider the traffic contribution for each P2P Application in the table.

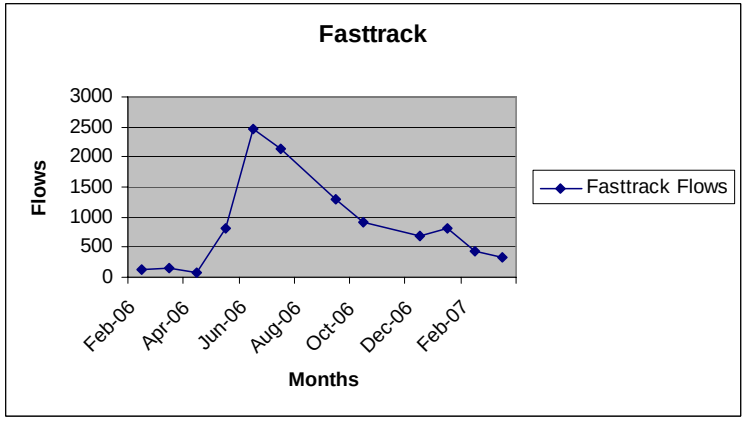
One Year of Peer-to-Peer



MP2P, or Manolito, is a P2P system primarily used to share music files. MP2P traffic was the least contributor to the overall network traffic among the observed systems. This traffic reached a peak flow count of just under 160 in January 2007.

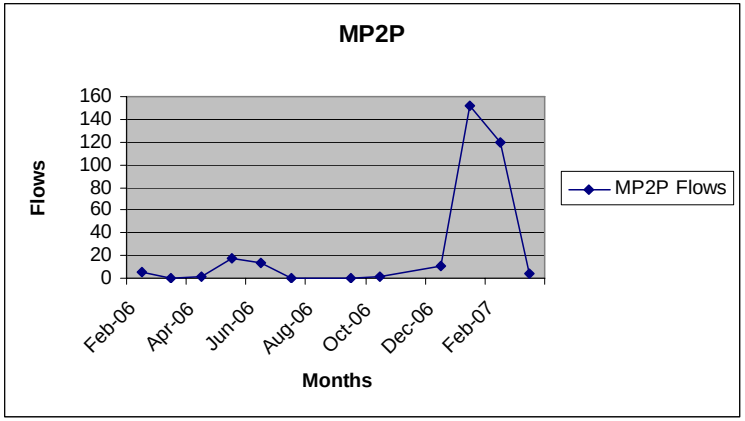
One Year of Peer-to-Peer

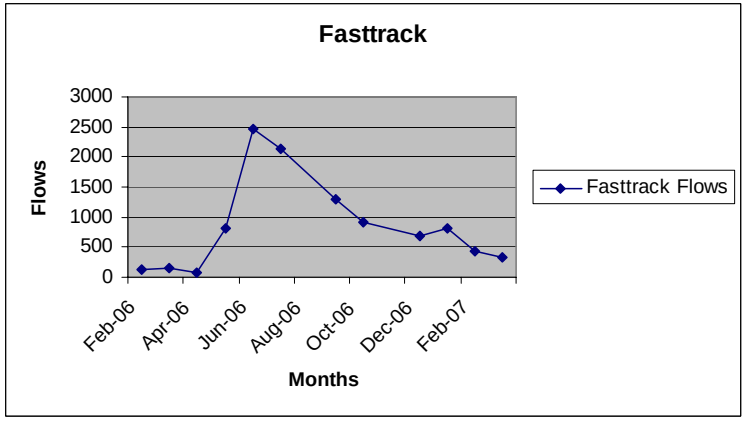


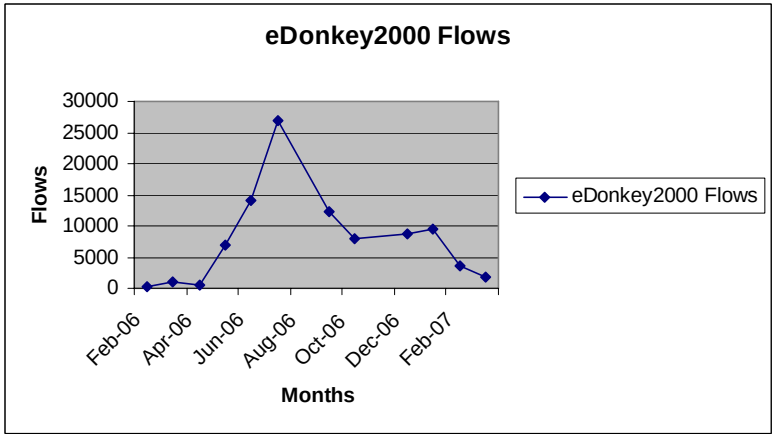


The Fasttrack P2P system is primarily used by Kazaa and its variants to exchange mp3 music files. Fasttrack traffic reached a peak flow count of 2,500 in July 2006.

One Year of Peer-to-Peer

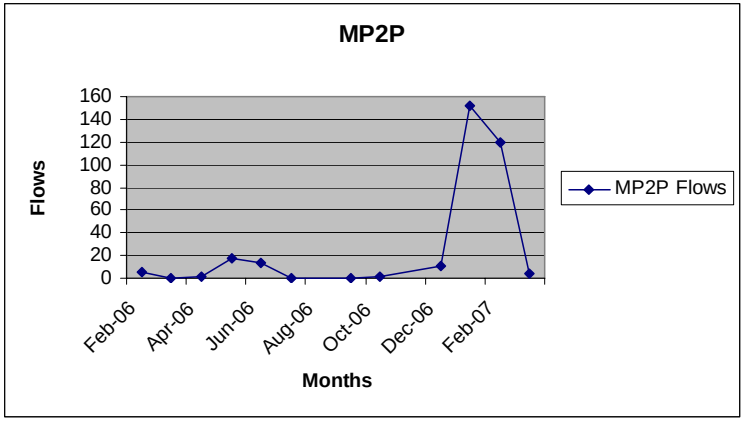


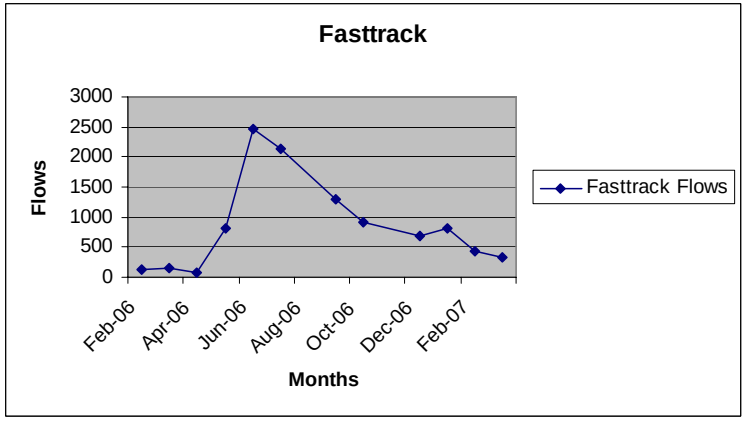




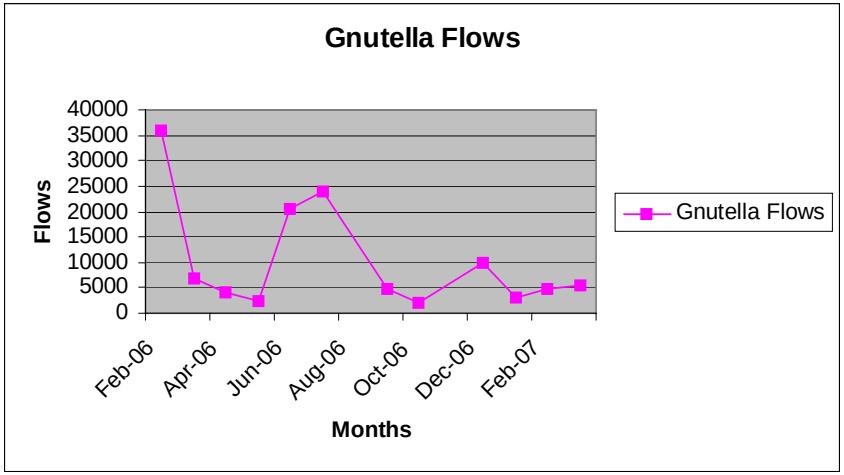
EDonkey2000 was a peer-to-peer system primarily used to distribute large images, video games and software. Although officially discontinued in September 2005 due to legal action brought by the Recording Industry Association of America (RIAA), we speculate, based on our profiling, that we observed eDonkey2000 communication during 2006. EDonkey traffic passed 25,000 flows in July 2006.

One Year of Peer-to-Peer

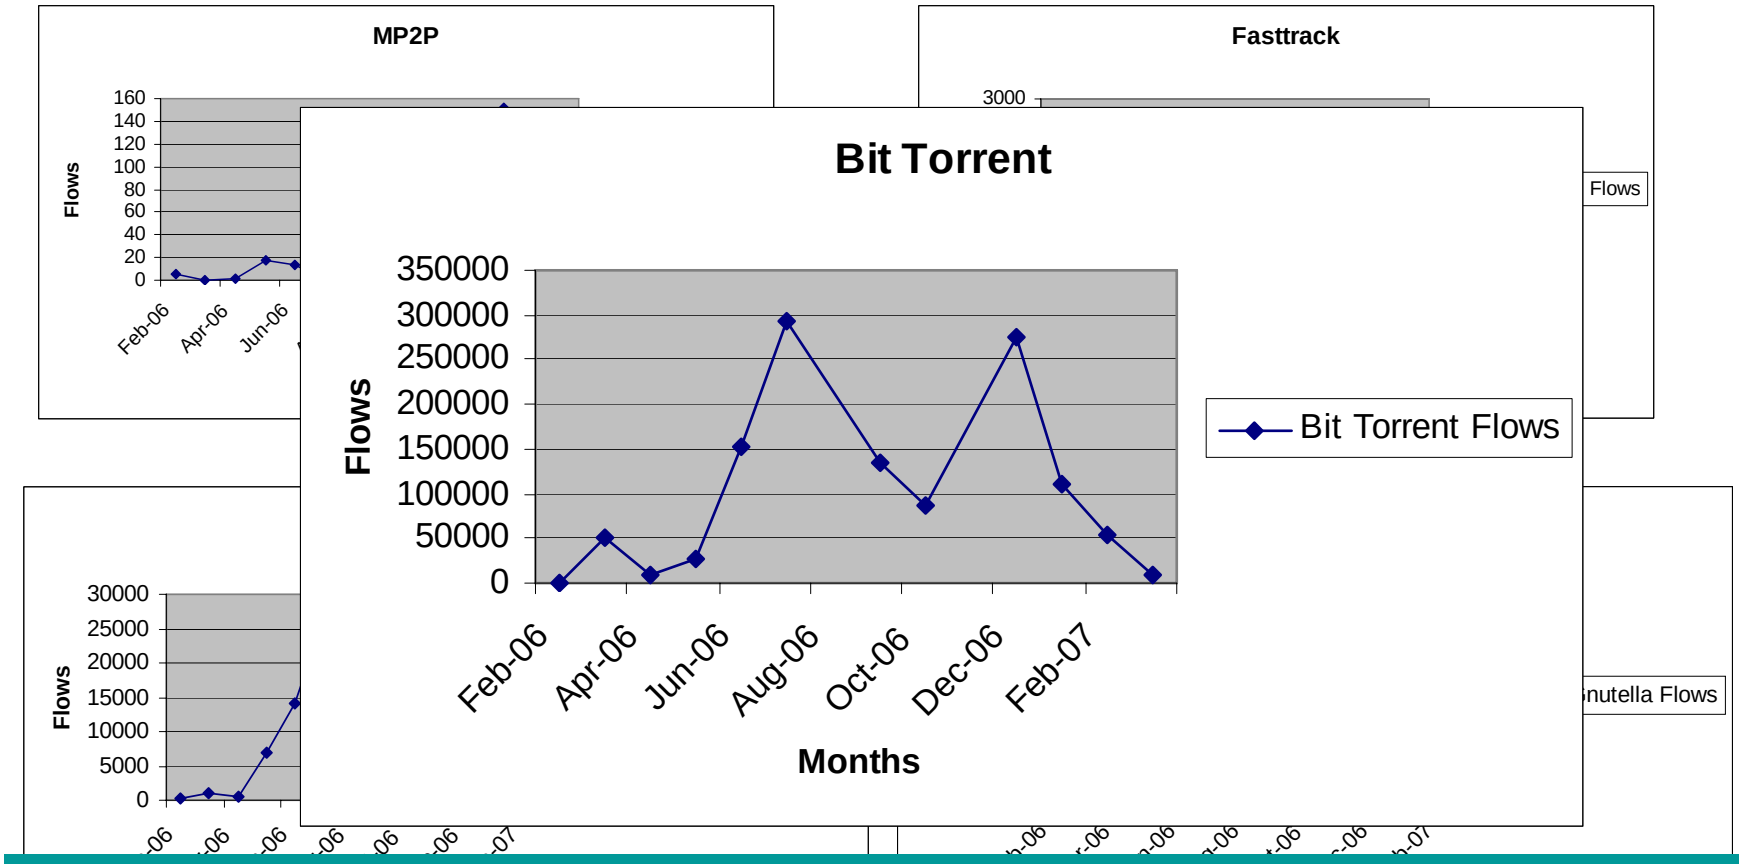




Gnutella is a multi-tier Peer based file exchange system. Traffic from Gnutella ranged from 5,000 to 35,000 flows per month.

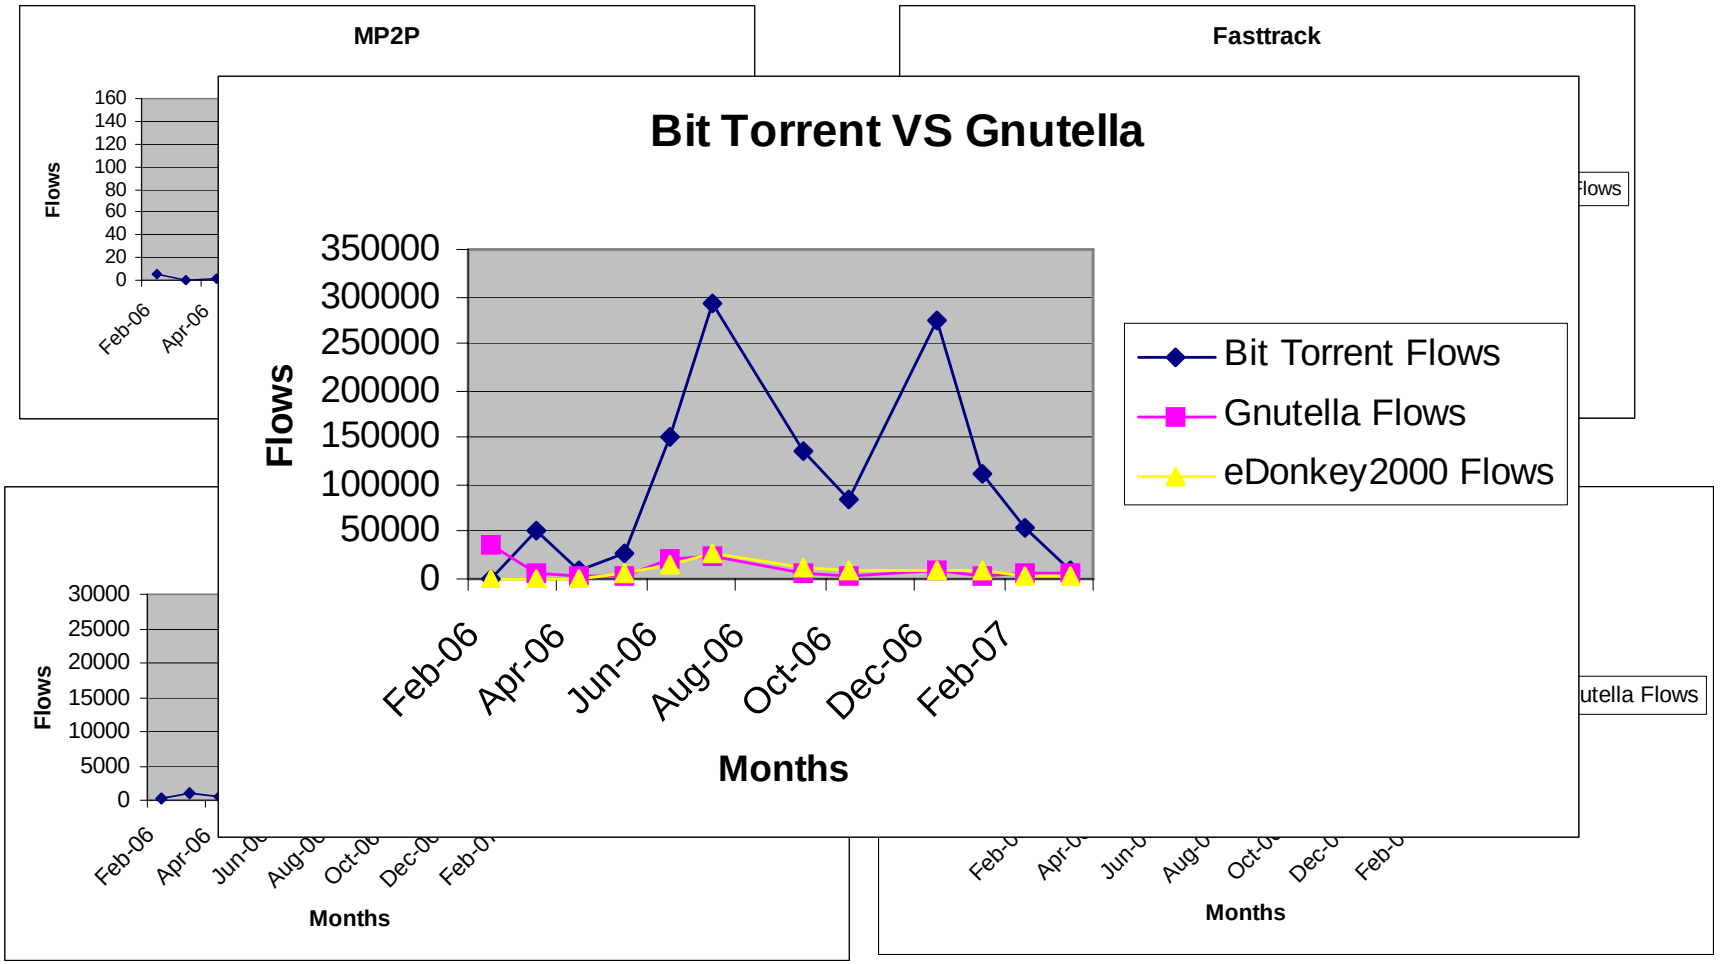


One Year of Peer-to-Peer



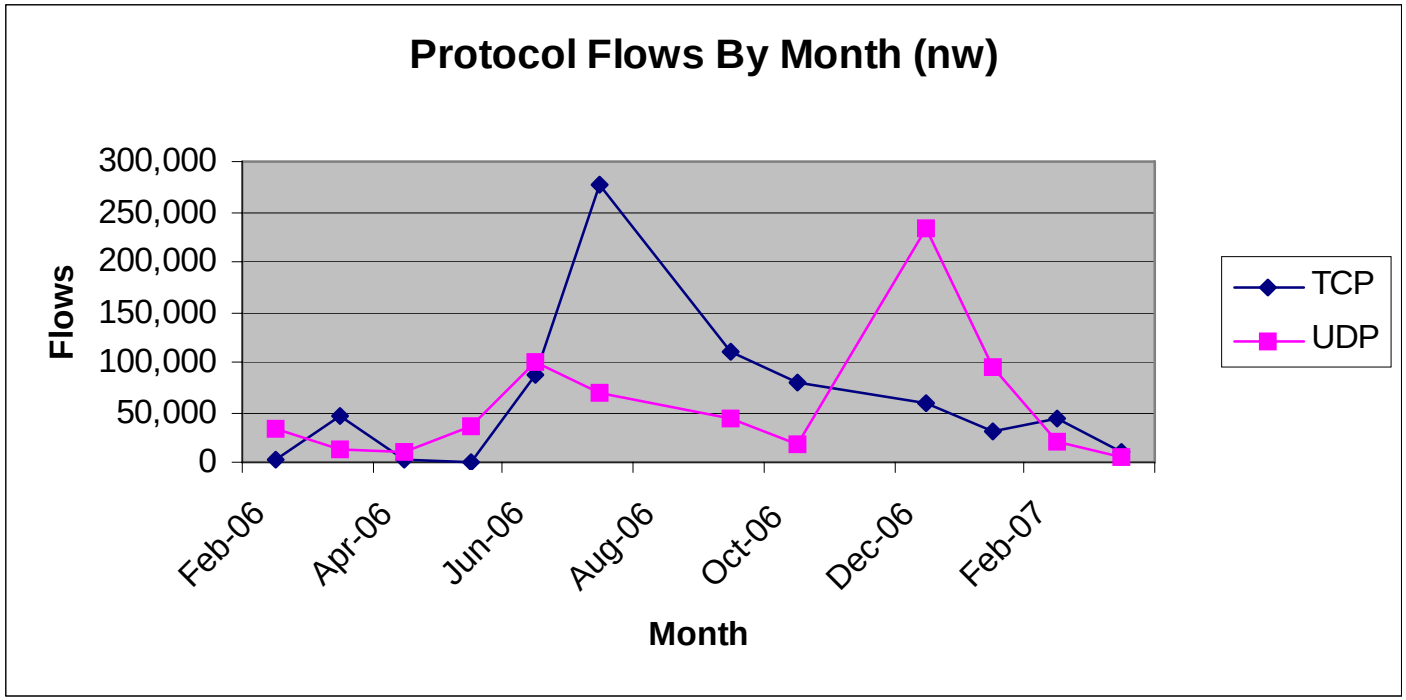
BitTorrent is an ever increasing popular P2P system used for exchanging large data files. Many open source software releases are distributed using BitTorrent. It is also used to distribute legal movie and music downloads. BitTorrent traffic eclipsed most P2P traffic at 300,000 flows.

One Year of Peer-to-Peer

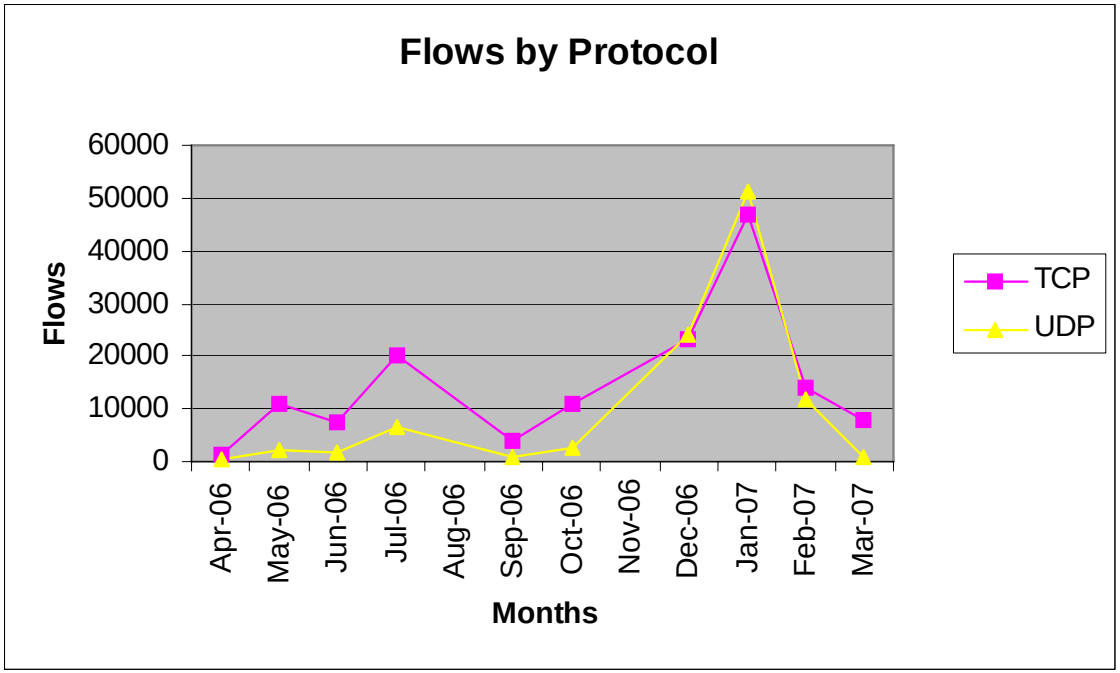


One Year of Peer-to-Peer

Unfortunately the overall Peer-to-Peer flow pattern did not match the pattern that we were seeking. That being a 50/50 ratio of TCP to UDP.

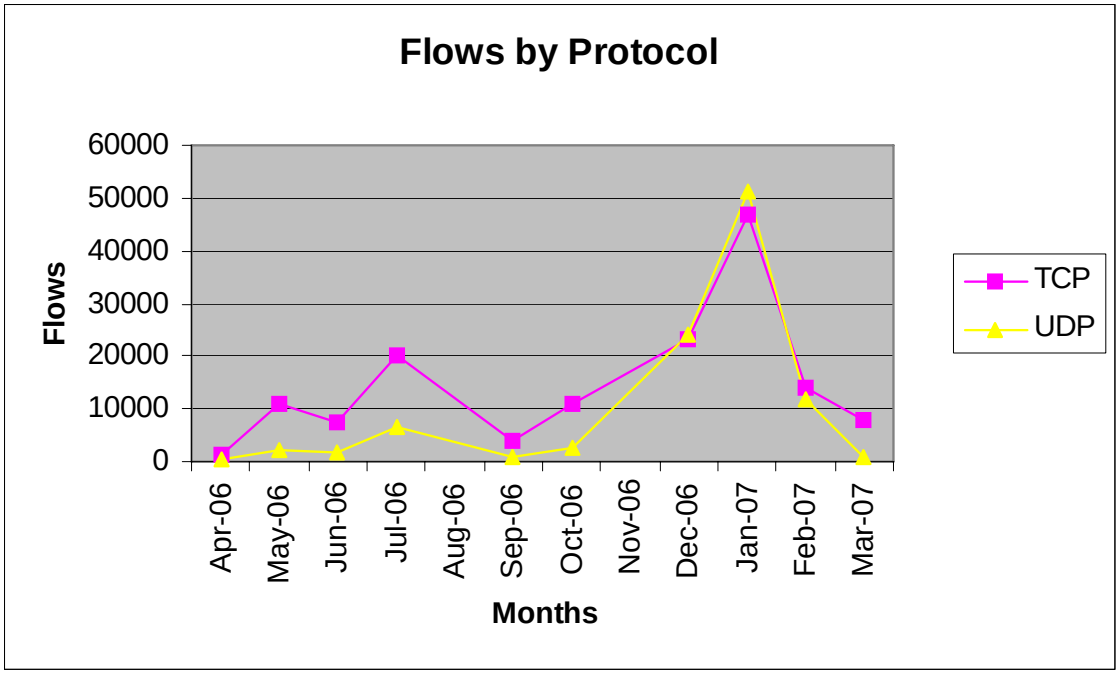


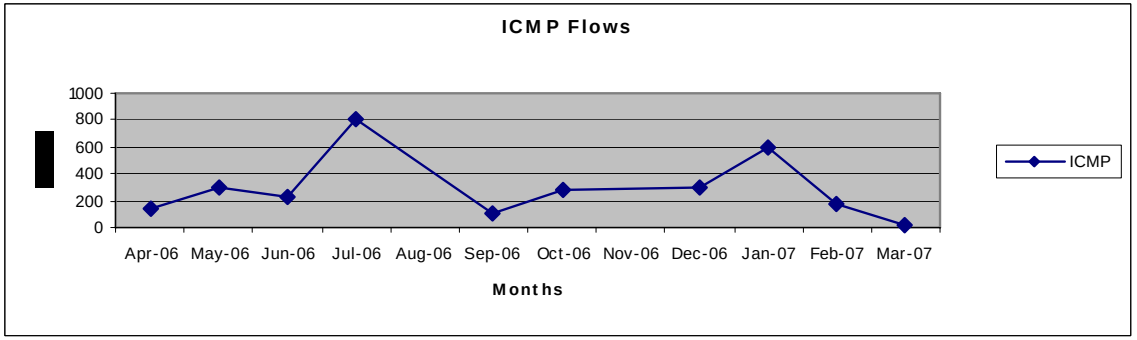
One Year of Peer-to-Peer



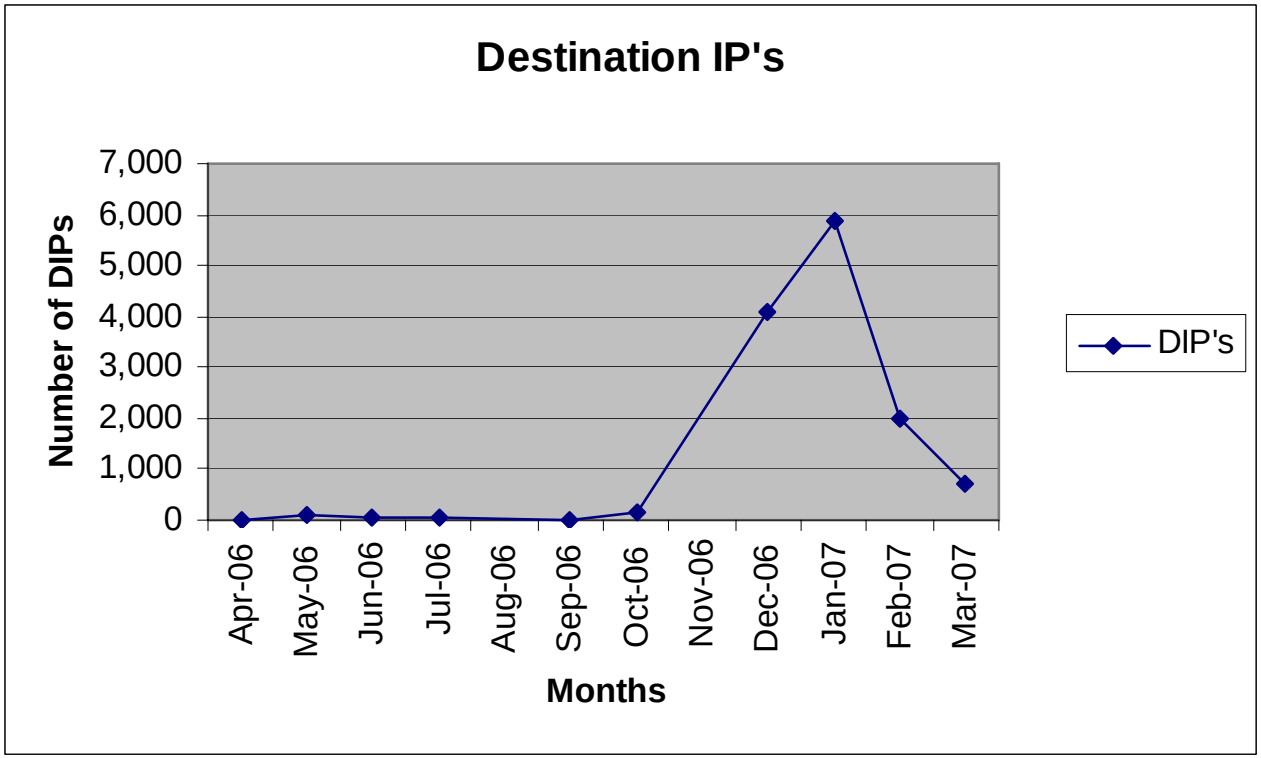
The graph above shows the pattern for which we were searching. This is the traffic from a single user workstation, with a peak flow count of 50,000 flows per month.

One Year of Peer-to-Peer



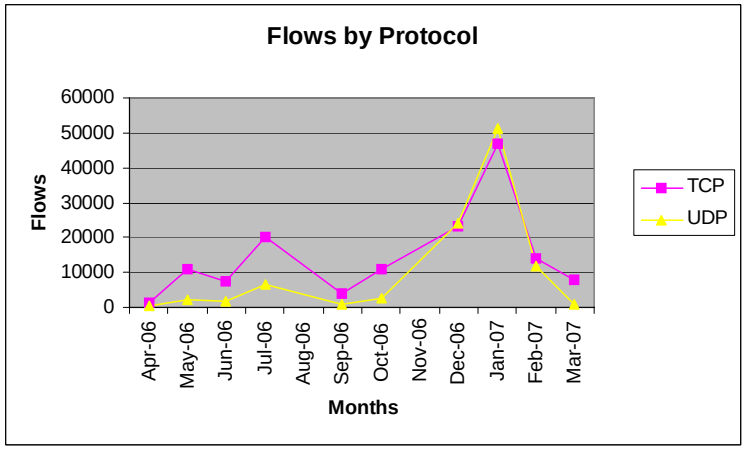


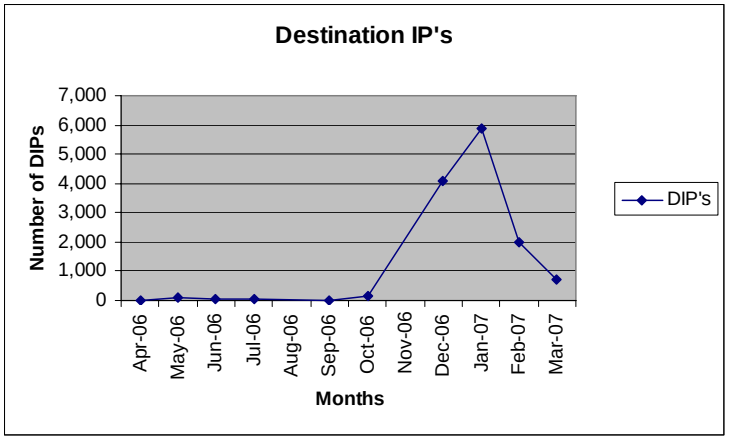
One Year of Peer-to-Peer



This workstation changed its behaviour in late fall 2006 from talking to less than 100 DIPs per month to 6,000 DIPs per month.

One Year of Peer-to-Peer





Who am I ?

One Year of Peer-to-Peer

SKYPE

This traffic pattern is driven by the adoption of Voip by a single user in the target network.

Disclaimer: It is important to point out that since the experimenter had no access to the actual machine or payload data this conclusion is simply conjecture based on known user Behaviour within the target network. (Skype is a wonderful App)

Observations on Traffic for Clients and Peers

- Consumes considerable Resources.
- Represents an Application Level WAN Network for Communication.
- Provides a channel to hide Malicious Activity.

"*McAfee suggested hackers were likely to create malicious software to target instant messaging services, Voice over Internet Protocol (VoIP) telephony services and online gaming sites.*" **Hackers will target social networking sites: security firms - Thursday, November29, 2007, CBC News http://www.cbc.ca**

Evidence that all is not as it Appears

- One day in February a conversation took place between a user host on the Network and a host compromised by an on-line game server.
- Two hours later the user host was attempting to contact a few friends….

Sequentially….

Destination IP	sPort	dPort	Proto	bytes
XXX.XXX.026.000	0	2048	1	56
XXX.XXX.026.000	0	2048	1	168
XXX.XXX.026.001	0	2048	1	56
XXX.XXX.026.001	0	2048	1	168
XXX.XXX.026.002	0	2048	1	56
XXX.XXX.026.002	0	2048	1	168
XXX.XXX.026.003	0	2048	1	56
XXX.XXX.026.003	0	2048	1	168
XXX.XXX.026.004	0	2048	1	56
XXX.XXX.026.004	0	2048	1	168
XXX.XXX.026.005	0	2048	1	56
XXX.XXX.026.005	0	2048	1	168
XXX.XXX.026.006	0	2048	1	56
XXX.XXX.026.006	0	2048	1	168
XXX.XXX.026.007	0	2048	1	56
XXX.XXX.026.007	0	2048	1	168
XXX.XXX.026.008	0	2048	1	56
XXX.XXX.026.008	0	2048	1	168
XXX.XXX.026.009	0	2048	1	56
XXX.XXX.026.009	0	2048	1	168
XXX.XXX.026.010	0	2048	1	56
XXX.XXX.026.010	0	2048	1	168
XXX.XXX.026.011	0	2048	1	56
XXX.XXX.026.011	0	2048	1	168
XXX.XXX.026.012	0	2048	1	56
XXX.XXX.026.012	0	2048	1	168
XXX.XXX.026.013	0	2048	1	56
XXX.XXX.026.013	0	2048	1	168
XXX.XXX.026.014	0	2048	1	56
XXX.XXX.026.014	0	2048	1	168
XXX.XXX.026.015	0	2048	1	56
XXX.XXX.026.015	0	2048	1	168
XXX.XXX.026.016	0	2048	1	56

We Need to Re-Consider our Willingness to be a Peer

- Users willingly download and install client/peer/server software.
- They even participate in strategies to avoid barriers and impediments (like Nat'ing).
- There is an implied trust that the communication is exclusively what it claims to be.
- "When they thought they were playing at war craft, they were actually playing at war craft."

Concluding Notes

- The network is evolving at the edges
- This means that network architectures, management and provisioning strategies are now more responsive then ever.
- Global communication resources are primarily influenced by the uncoordinated activities of individuals.
- Traffic patterns are emergent properties without intent.

Future Work

- Study the growth in diversity of patterns in traffic.
- Study the form and distribution of applications and participants.
- Track Unidentified Anomalies.

- February 2008, TARA will announce the InTARA project

Intelligent Network Traffic Analyzers for Reconstructive and Real Time Analysis

- InTARA will be a multi-million dollar, multi-year project to develop intelligent traffic analysis capabilities for the good guys.
- We are seeking global collaborative research and commercialization partners. Early stage interest from Australia, India, Switzerland, Canada.

Revisiting the Threshold Random Walk Scan Detector

Vagishwari Nagaonkar
Dr.John Mchugh
Faculty of Computer Science
Dalhousie University

Presented for FLOCON 2008

DALHOUSIE UNIVERSITY
Inspiring Minds

Introduction

- Initial Activity in many intrusions
  - Scanning
- Techniques to detect these initial scans
- One of the effective algorithms
  - Threshold Random Walk

DALHOUSIE UNIVERSITY
Inspiring Minds

Introduction (contd.)

- Challenges when using TRW
  - UDP and ICMP Traffic
  - Repetitive Scanning
  - Slow and Stealthy Scans
- Using Bloom filters
  - eliminate repetitive input to TRW
  - look for reverse matches in time ordered data

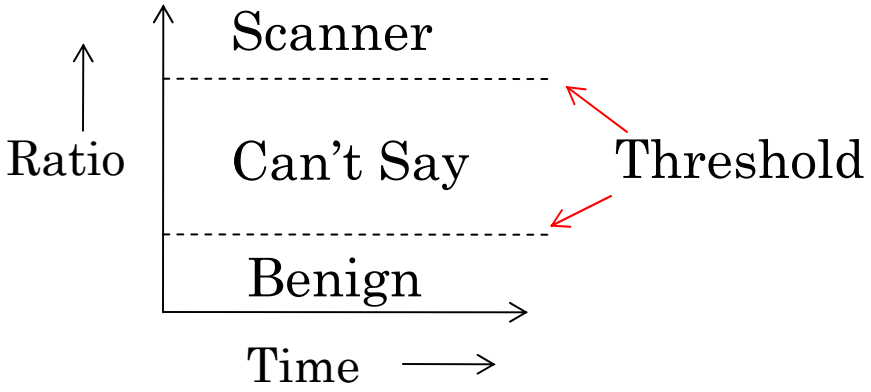
DALHOUSIE UNIVERSITY
Inspiring Minds

Threshold Random Walk

- Scan Detection Algorithm based on sequential hypothesis testing.

- Uses a positive reward based scan detection.
  - For a given host, records connection attempt made :

Connection	Ratio
Successful	Decreases
Failed	Increases



DALHOUSIE UNIVERSITY
Inspiring Minds

Threshold Random Walk

- The ratio is calculated as :

$$\Lambda(Y) \equiv \frac{\Pr[Y|H_1]}{\Pr[Y|H_0]} = \Pi_{i=1}^{n} \frac{\Pr[Y_i|H_1]}{\Pr[Y_i|H_0]}$$

- Where the probabilities are :

$$\Pr[Y_i = 0|H_0] = \theta_0, \quad \Pr[Y_i = 1|H_0] = 1 - \theta_0$$
$$\Pr[Y_i = 0|H_1] = \theta_1, \quad \Pr[Y_i = 1|H_1] = 1 - \theta_1$$

  - **Y = success (0) or failed (1) connection attempt**
  - **H0 = benign hypothesis**
  - **H1 = scanner hypothesis**
  - **Θ0 = probability that the source is benign, for a successful connection attempt**
  - **Θ1 = probability that the source is scanner for a successful connection attempt**

DALHOUSIE UNIVERSITY
Inspiring Minds

Threshold Random Walk

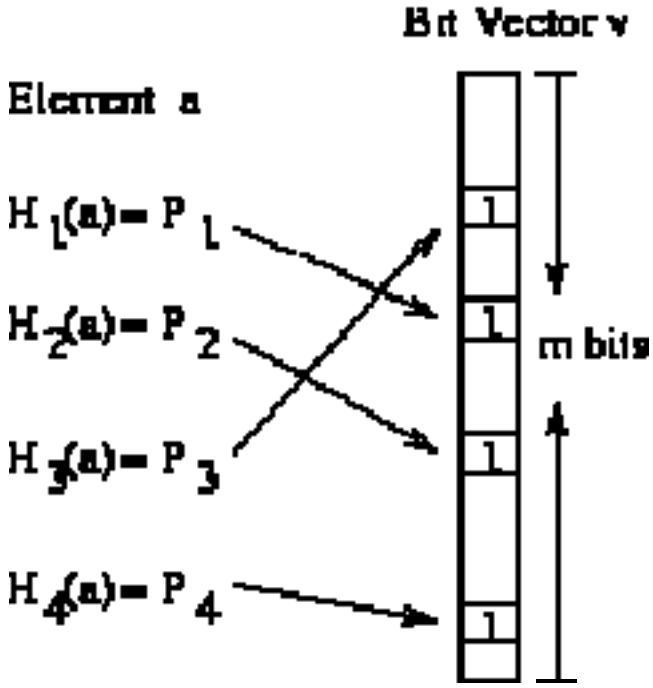
- The thresholds are calculated based on
  - desired true positive ($\beta = 0.99$)
  - desired false positive ($\alpha = 0.01$)

$$\eta_1 \leftarrow \frac{\beta}{\alpha} \qquad \eta_0 \leftarrow \frac{1-\beta}{1-\alpha}$$

DALHOUSIE UNIVERSITY
Inspiring Minds

Bloom Filter

- It's a Data Structure
  - test the membership of an element for a given set
- Definition of the Structure
  - bit array of m bits
  - k different hash functions
  - Hash functions maps a key value to one of the m array positions.



DALHOUSIE UNIVERSITY
Inspiring Minds

Bloom Filter

- Properties :
  - False positives possible
  - No false negatives
  - Elements can be added
  - No deletion possible
  - Greater the number of elements, higher the probability of false positives.
  - Space Efficient
  - Cannot determine the elements present in it.

DALHOUSIE UNIVERSITY
Inspiring Minds

Modified TRW with Bloom Filter

- TRW hit or miss definition
  - For a given pair in the flow record
    eg {sip, dip}
    - HIT = if a corresponding entry {dip, sip, sport, dport, proto} is found within a specified timeout period
    - MISS = if a corresponding entry {dip, sip, sport, dport, proto} is not found within a specified timeout period

DALHOUSIE UNIVERSITY
Inspiring Minds

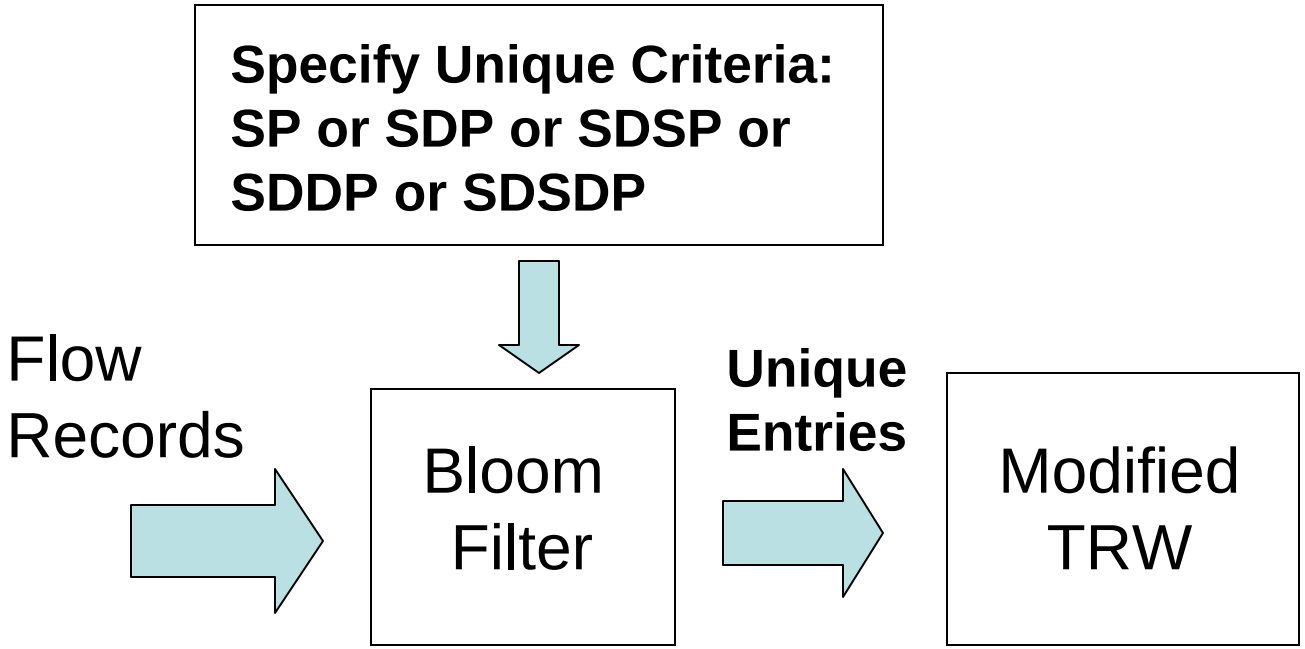
Modified TRW with Bloom Filter

- Bloom Filter uses 10 hash functions and a bit vector of size 2^32

- Experiment Set up :
  - Pass the flow records through the bloom filter.
  - Specify selection criteria: {sip, dip}, {sip, dip, proto}, {sip, dip, sport}, {sip, dip, dport}, {sip, dip, sport, dport, proto}
  - Use the TRW scanning algorithm.

DALHOUSIE UNIVERSITY
Inspiring Minds

Modified TRW with Bloom Filter



DALHOUSIE UNIVERSITY
Inspiring Minds

The Dataset

- A year long trace collected on a /22 enterprise network
- Using Silk Tools
- Internal Network Hosts
  - Total Address Space = 1024
  - #Active hosts in a given day = varies between 60-70
  - Active Address Space ~ 6%

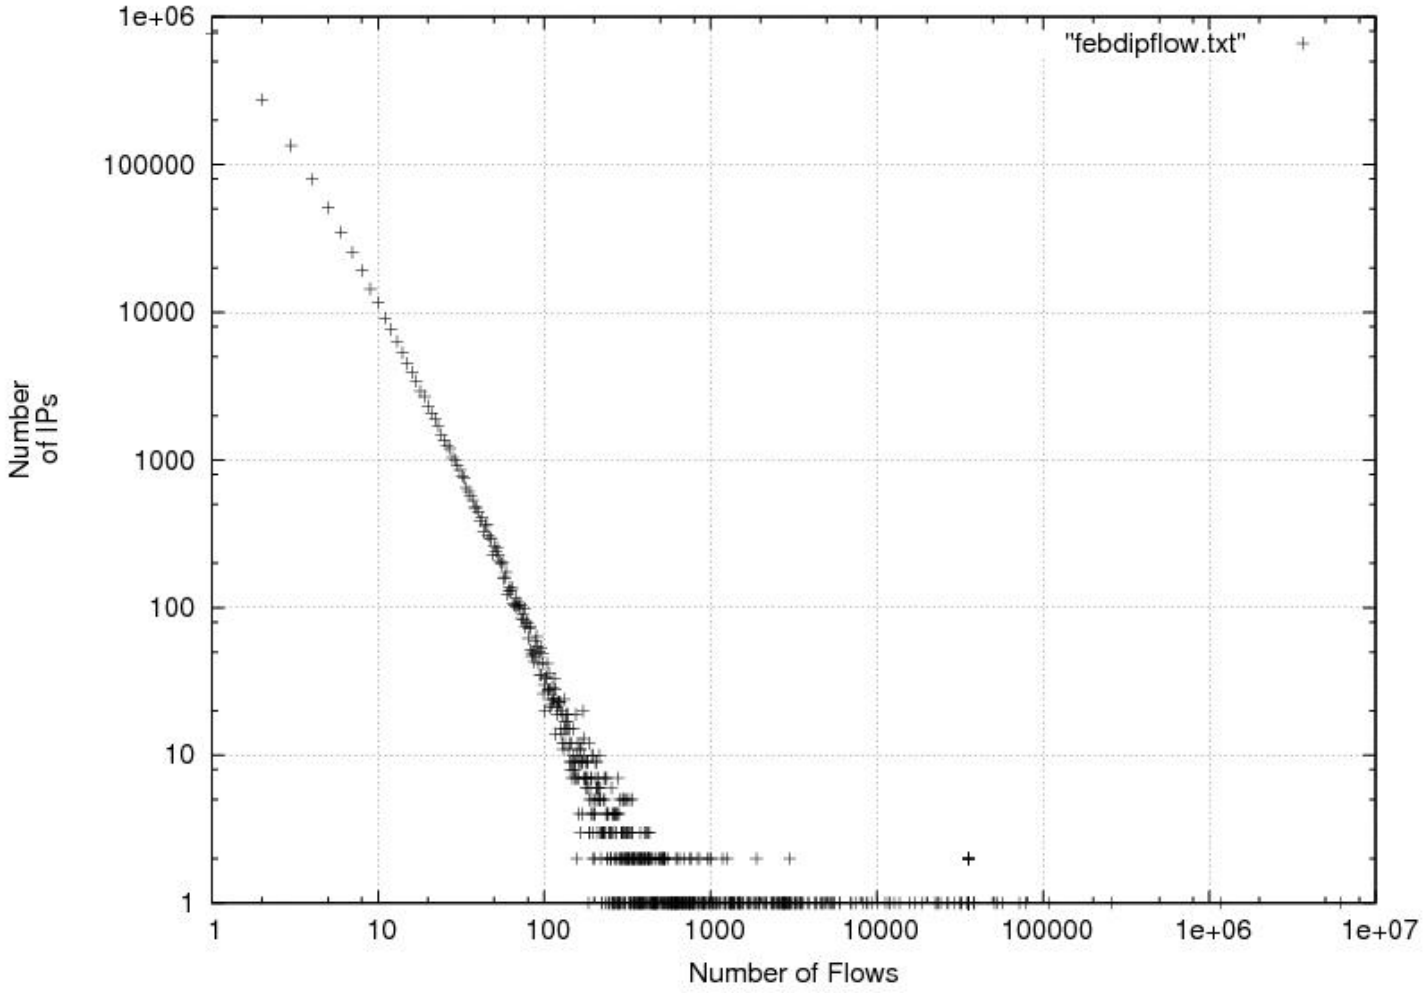
DALHOUSIE UNIVERSITY
Inspiring Minds

The Dataset

	OutIps Seen			
	EtoO	**OtoE**	**Non Responsive Out ips**	**% Non Responsive Out ips**
Feb	26680	7270	19410	72.75112444
Mar	30232	3866	26366	87.21222546
Apr	56126	14576	41550	74.02986138
May	2355612	106893	2248719	95.46219836
June	2847371	283270	2564101	90.05152472
July	2601834	246312	2355522	90.53313932
Aug	30181	29097	1084	3.591663629
Sept	126913	126549	364	0.28681065
Oct	330740	277438	53302	16.11598234
Nov	4050	2932	1118	27.60493827
Dec	2226535	254484	1972051	88.57040199
Total	**10636274**	**1352687**	**9283587**	**87.28232274**

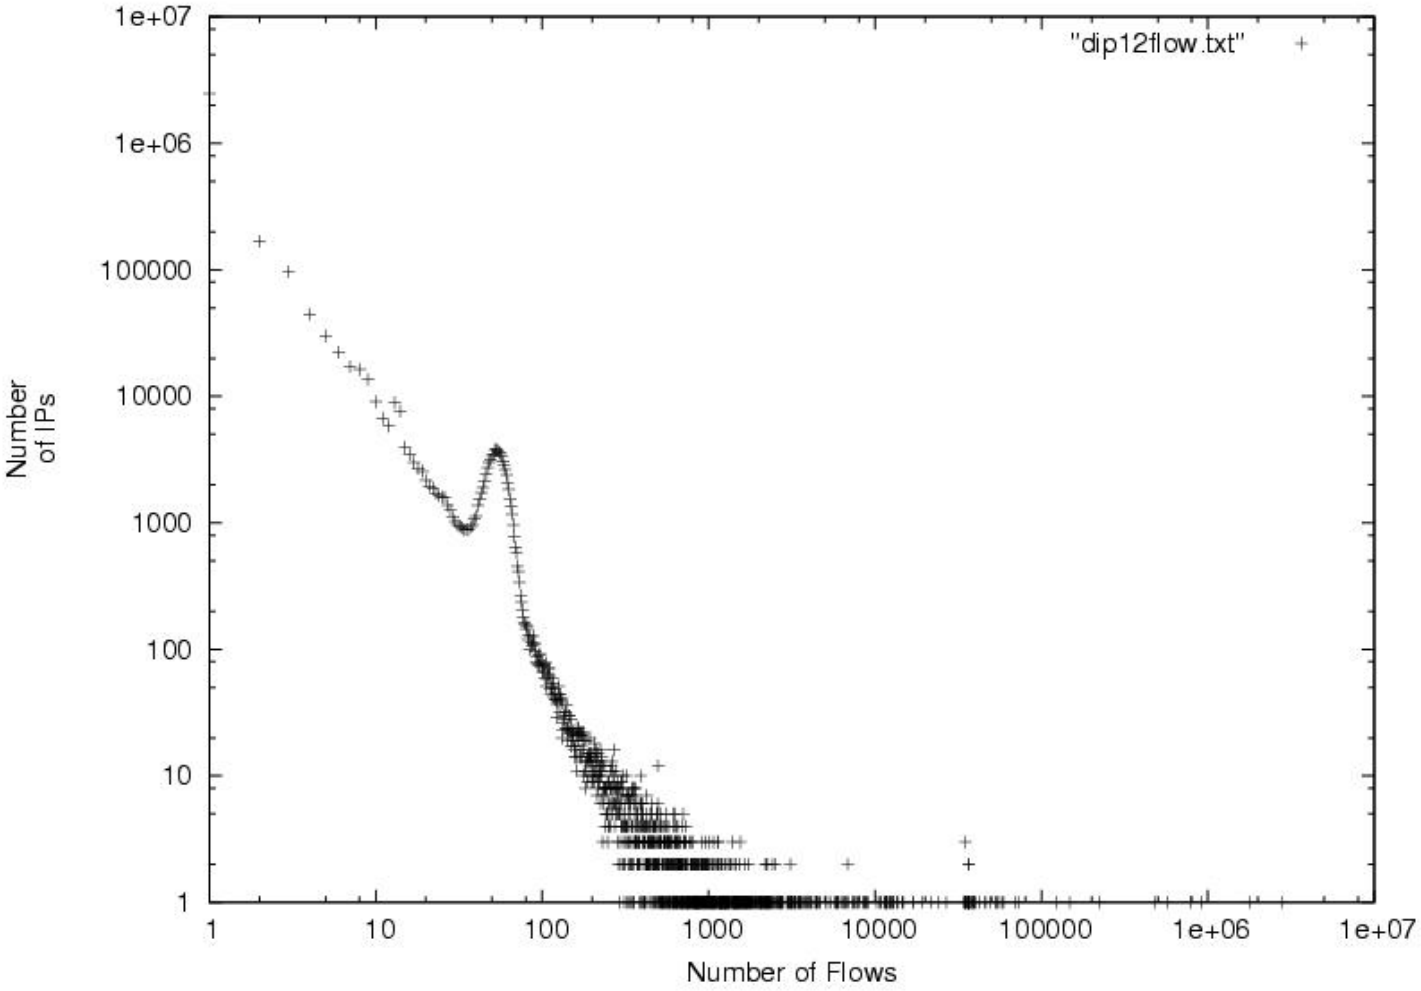
DALHOUSIE UNIVERSITY
Inspiring Minds

The Dataset



DALHOUSIE UNIVERSITY
Inspiring Minds

The Dataset



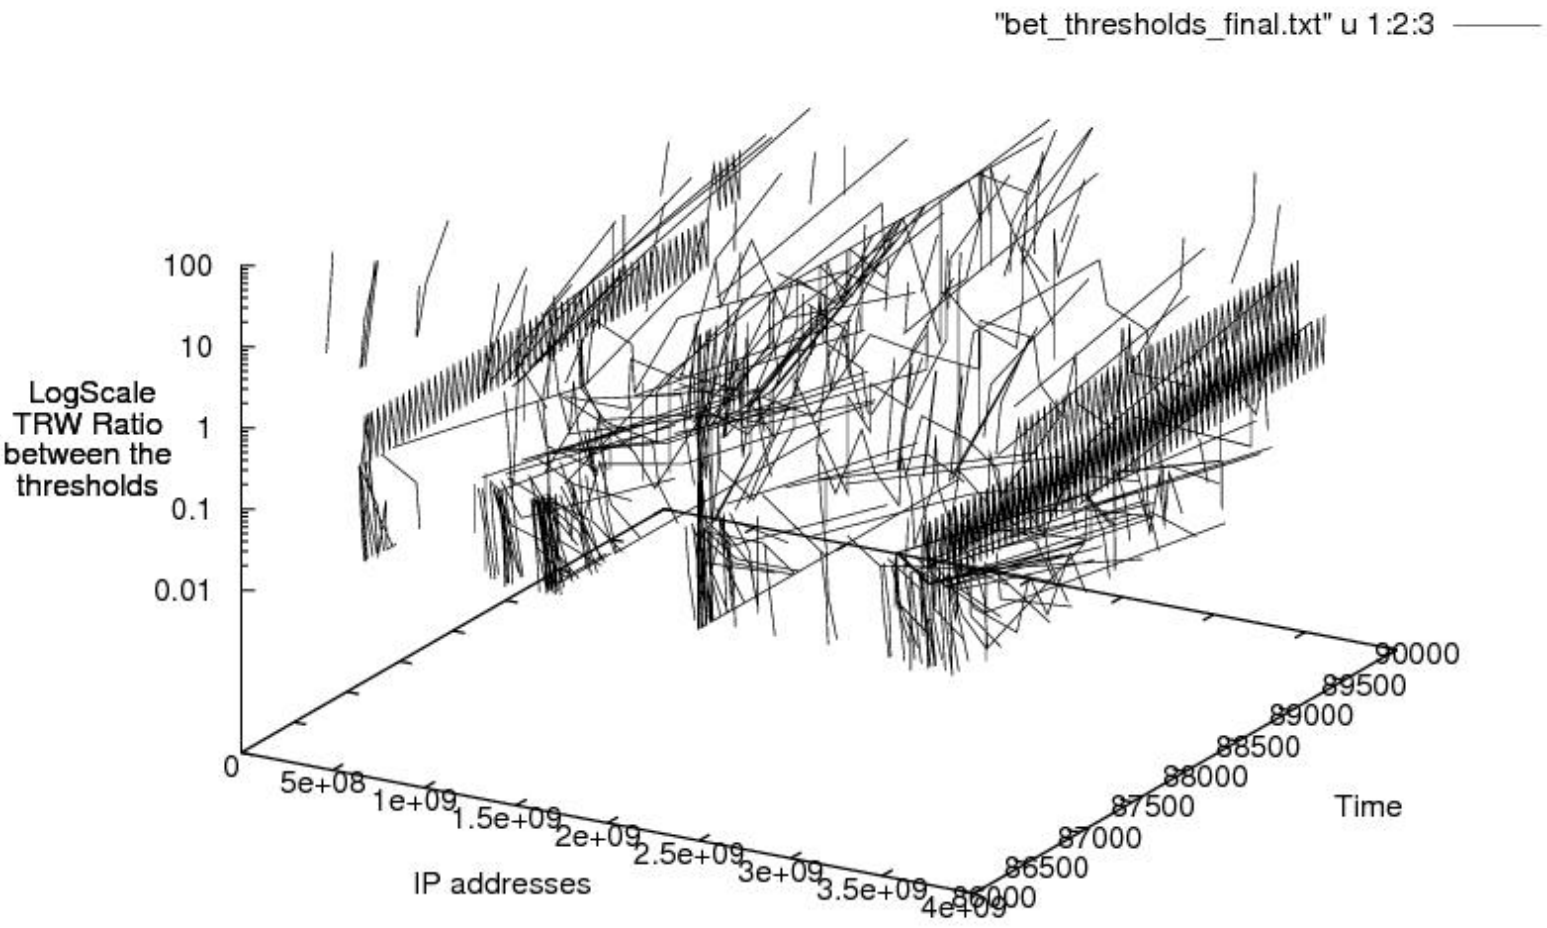
DALHOUSIE UNIVERSITY
Inspiring Minds

Problems faced during Analysis

- Time granularity
  - millisecond not available.
  - The order of flow records for the same second is the outside to inside put first.
- Background noise in the traffic.
- ICMP ping traffic causes false detection.

DALHOUSIE UNIVERSITY
Inspiring Minds

Problems faced during Analysis



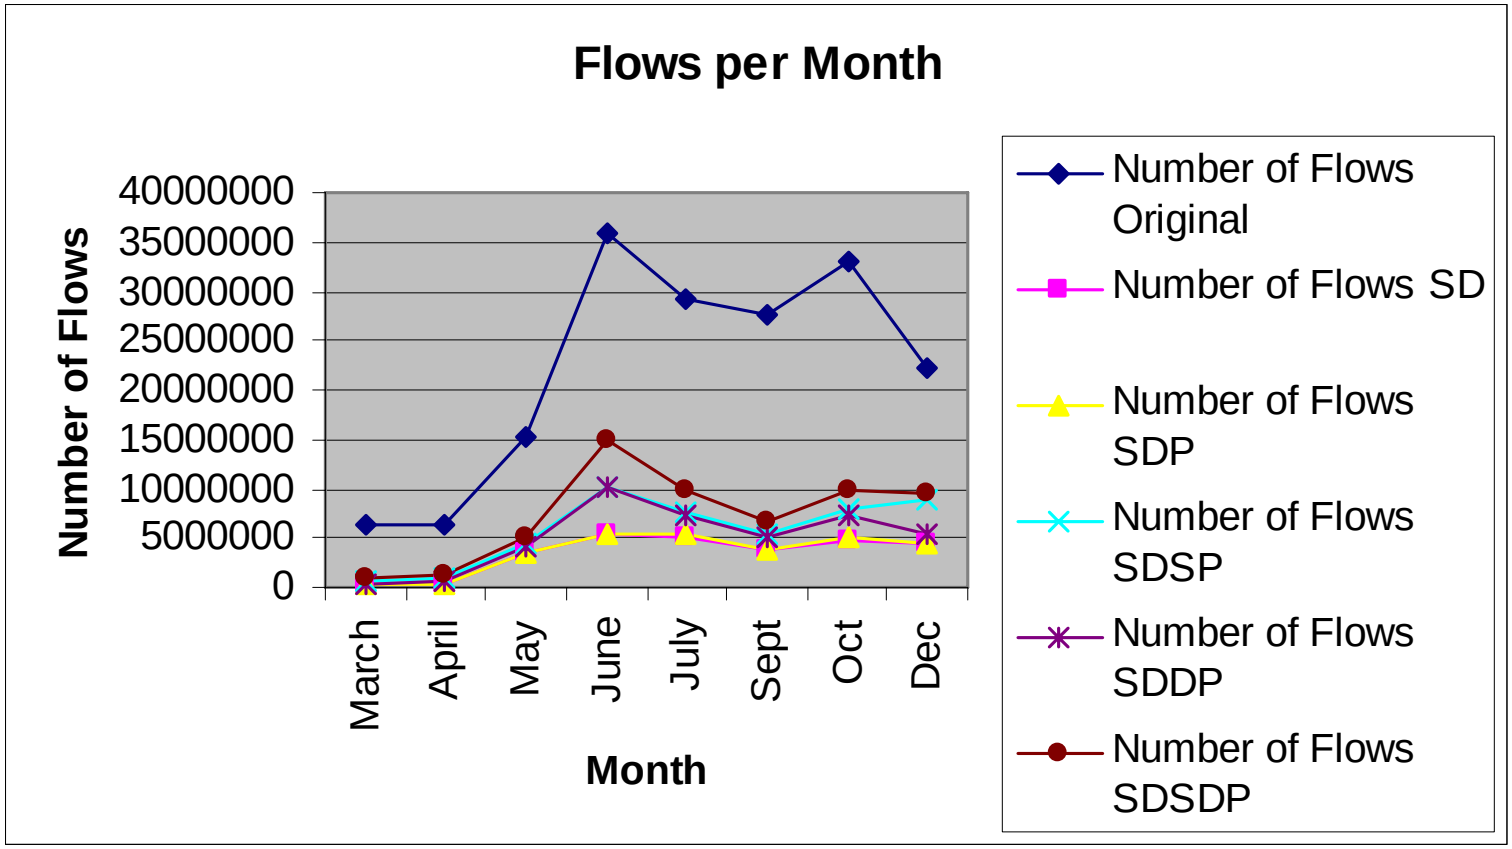
DALHOUSIE UNIVERSITY
Inspiring Minds

Preliminary Results

- TRW Parameters used:
  - Theta1 determined based on the %active internal hosts compared to the total address space ~ 0.0654
  - Theta0 ~ 0.8
    - Changed theta0 for benign hosts to hits / (hits + miss)
    - The value of new theta0 ranged from 0.45 to 1.00
    - All benign hosts still classified as benign
  - Alpha (desired false positive) = 0.01
  - Beta (desired true positive) = 0.99

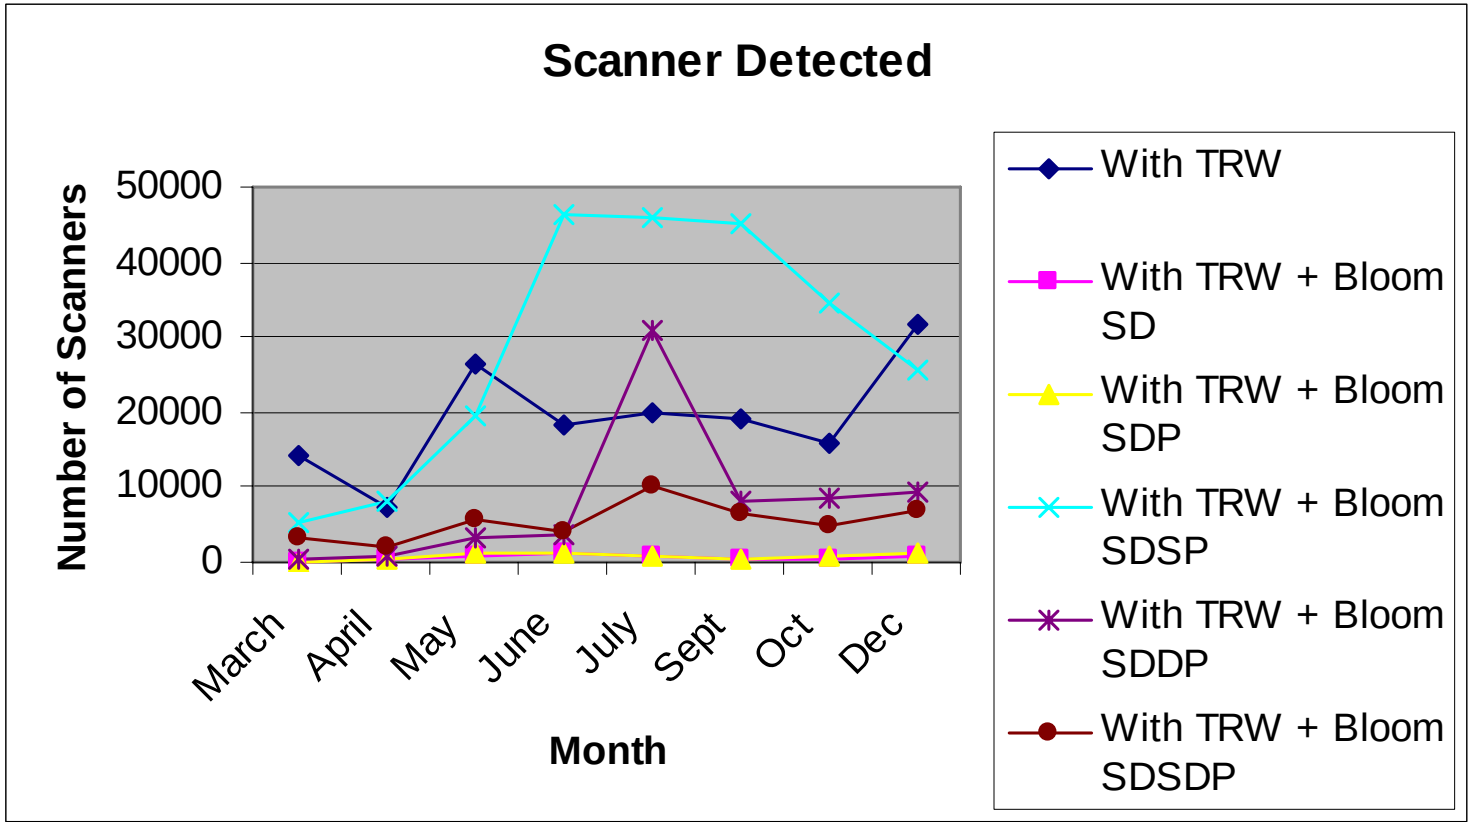
DALHOUSIE UNIVERSITY
Inspiring Minds

Preliminary Results



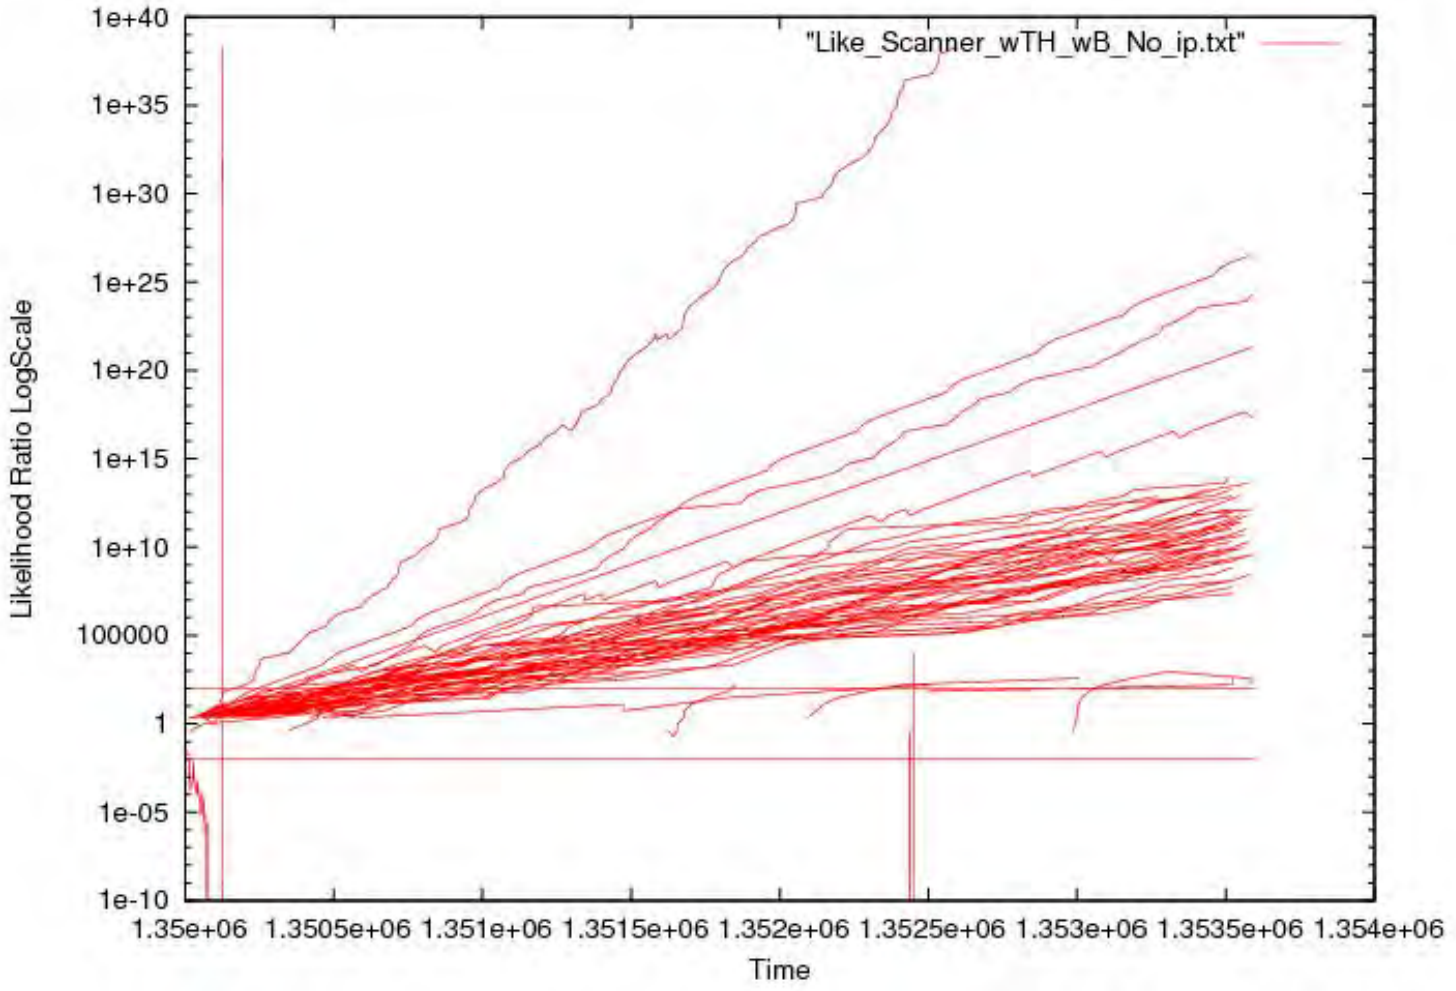
DALHOUSIE UNIVERSITY
Inspiring Minds

Preliminary Results



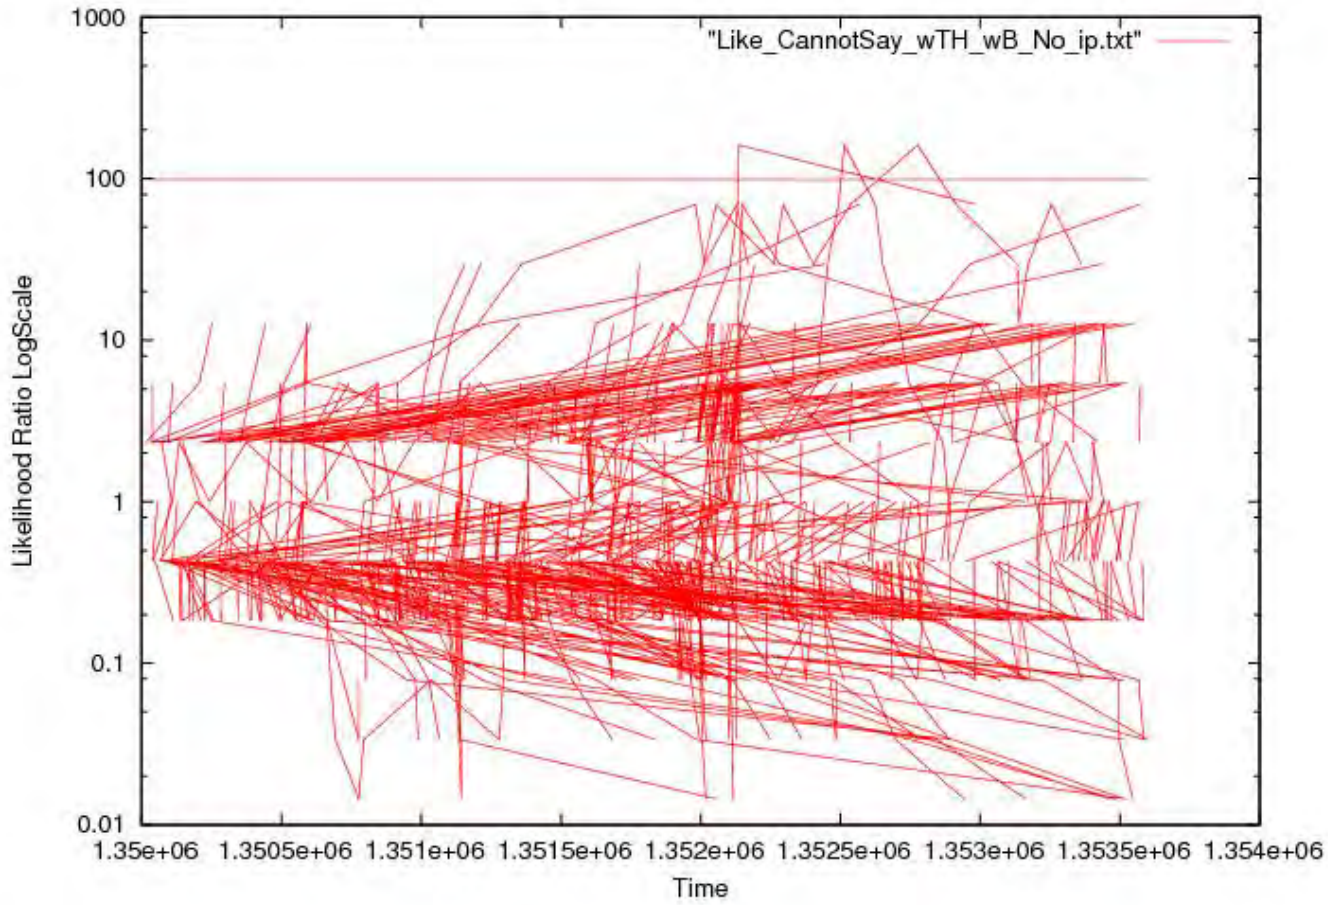
DALHOUSIE UNIVERSITY
Inspiring Minds

Preliminary Results
Plot of Likelihood ration for Scanners



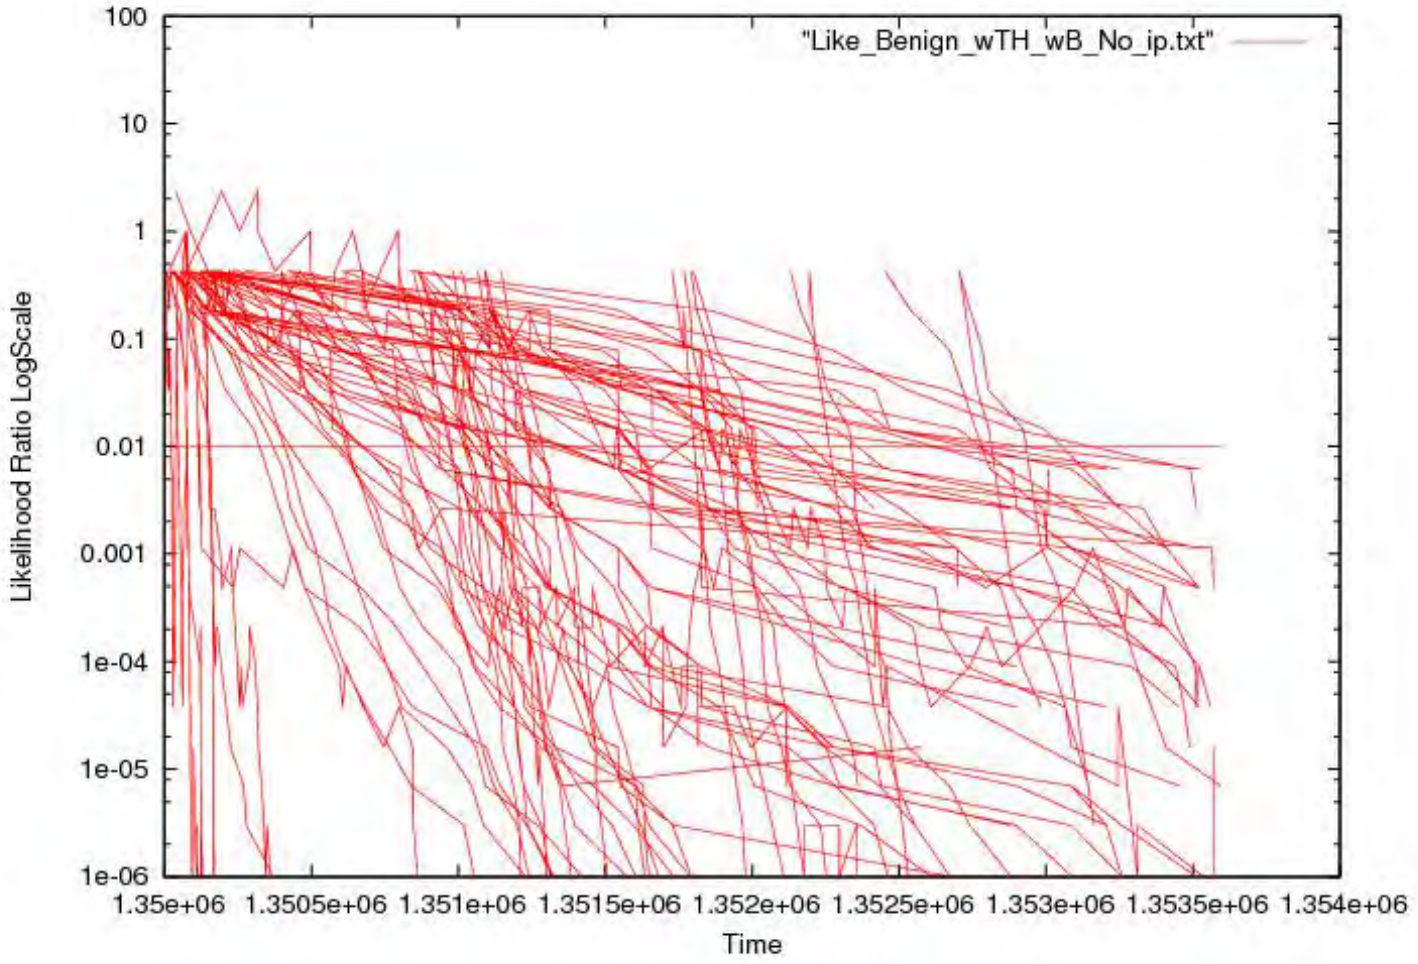
DALHOUSIE UNIVERSITY
Inspiring Minds

Preliminary Results
Plot of Likelihood ration for Can't Says



DALHOUSIE UNIVERSITY
Inspiring Minds

Preliminary Results
Plot of Likelihood ration for Benign



DALHOUSIE UNIVERSITY
Inspiring Minds

Initial Conclusions

- Using Bloom filter, reduces the false positives, (by how much ?)
  - unique entries considered for a given filter criteria
- Using specific filter criteria for the bloom filter
  - detects vertical scanning
  - detects horizontal scanning

DALHOUSIE UNIVERSITY
Inspiring Minds

Further Work In Progress

- Need to improve the technique by
  - Vary theta0 and theta1 values
  - Effect of timeout period
  - Real time scenario
- Long term analysis of IPs toggling between the three regions
  - Esp. from scanning to Can't say or benign

DALHOUSIE UNIVERSITY
Inspiring Minds

Acknowledgments

- Ron McLeod
- TARA
- Faculty of Computer Science, Dalhousie University

DALHOUSIE UNIVERSITY
Inspiring Minds

Thank you

Questions ?

DALHOUSIE UNIVERSITY
Inspiring Minds

Flow Analysis in a Wireless Environment with short DHCP Leases

Sanket Parikh

John McHugh

Dept. of Computer Science

Dalhousie University

DALHOUSIE UNIVERSITY
Inspiring Minds

Project Objectives

- Analysis of Wireless Network Data from University of Dartmouth (Crawdad Archive)
- Adding MAC Layer information in Net Flow tools for identification of nodes and Activities performed by a node.
- Return converted flow data to the Crawdad archive.

DALHOUSIE UNIVERSITY
Inspiring Minds

Project Rationale

- The main issue in analyzing wireless network data from many environments is the assignment of temporary IP Addresses using DHCP with short leases.
- The total user population often exceeds the available address space, and a given user may connect to the network for short sessions from a number of different locations making complicating per platform analyses.
- Work to date has concentrated on mobility rather than platform behaviour.

DALHOUSIE UNIVERSITY
Inspiring Minds

The Data

- 160 GB of compressed tcpdump packet headers.
- Collected continuously from 2 Nov 04 - 28 Feb 04
- 18 collection points academic, library, residence
- Nothing beyond IP Headers except TCP ports and flags, UDP ports.
- Anonymized with prefix preserving technique
  - Usage agreement precludes attacking anonymization to determine user identity.
  - Low order 24 bits of MAC also anonymized
    - List of known wireless MAC addresses provided

DALHOUSIE UNIVERSITY
Inspiring Minds

Technical Approach - 1

- Tried to use vlan tag fields to avoid altering YAF record format.
- Use the Forward and Reverse vlan tag fields to get source and destination MAC addresses into the yafscii
- Since these are 16 bits use perfect hash of MAC
- Problems:
  - vlan tag is in unidirectional extension of flow. Need both, even for unidirectional flows.
  - would like to use with real time and when MAC set not completely known

DALHOUSIE UNIVERSITY
Inspiring Minds

Technical Approach

- We added MAC to the bidirectional flow root in yaf, with both source and destination MAC addresses.
- There are a number of subtleties here, including the use of memcopy that introduces field order dependencies (an IPv4 optimization) and the assumption that MAC flag implies vlanid not zero.
- Once the MAC addresses are into the yafscii output, we started converting it into SiLK for further data analysis
- Shortly after we finished, CERT added MAC address support to YAF and we will use it in the future.

Technical Approach

- We created a module *yafscii2tuc.c*
  - Inserts minimal perfect hash index of MAC in in / out
  - Adds sensor id from command line to identify the sniffers.
- We split the output of the *yafscii2tuc* into separate hourly streams and use *popen* to send each one to a separate invocation of *rwtuc* so that the resulting files are in a proper date hierarchy.
- We also use *rwsort* on the *rwtuc* output to ensure time order and because *rwtuc* does not compress.

Minimal perfect hashes

- A Minimal Perfect Hash maps a set of N unique strings into integers in [0...N-1]
  - Packages available on internet designed for null terminated strings
  - Modified for counted strings
  - Extracted all MACS from Dartmouth packet data
  - Grouped to bring common usages together, e.g. known wireless, gateways, etc. then created MPH
  - 17000+ MACs, 11,000+ with IP packets.
- Lookup is constant time, collision free

DALHOUSIE UNIVERSITY
Inspiring Minds

MAC types

- There are 5 categories of MACS actively involved
  - Known Wireless MACs with IP traffic
  - Other MACs with IP packets
  - Multi cast MACs
  - Gateway MACs
  - Broadcast MACs
- A large number of MACs have no IP traffic
  - Some appear only at link layer, others in MAC list but not seen
- We used rwfilter to build sets for each type of MAC address based on the input and output field values

DALHOUSIE UNIVERSITY
Inspiring Minds

Project Outcomes

- We found some interesting information during analysis of the datasets. There are traces which shows some IP addresses appeared in two different sniffers located to different locations.
- The reason may be the physical location of sniffers for collecting data. Though sniffers were not located at proper distance from each other, there might be the chances for getting same IP traces in two different sniffers.
- This seems improbable and needs further study

DALHOUSIE UNIVERSITY
Inspiring Minds

Remaining problems

- yaf does not deal with decreasing time well
  - In live capture, packets are always in increasing time order no matter what the clock says
  - In playback the same holds unless the file has been reordered.
  - Several Dartmouth sensors exhibit decreasing time, probably due to ntp or other clock adjustments.
- Data from one of the sensors “breaks” the pipe
  - This may be related to the time problem above or may be due to another problem
  - Truncated packets may lead to other pathologies in yaf

DALHOUSIE UNIVERSITY
Inspiring Minds

Next steps

- We want to reassign the IPs currently used to a consistent IP that is related to the MAC index.
- First we need to determine if any wireless IPs are associated with gateway MACs.
  - This would occur if a wireless unit talked to another wireless unit via a routed connection, e.g. units connecting via separate sniffers.
  - Start by creating sets for each MAC type and looking for intersections
  - May have to explore DHCP strategy in more detail.
- This is currently underway.

DALHOUSIE UNIVERSITY
Inspiring Minds

Next Steps

- With the technique we used for this research should prove useful for similar data from wireless “hot spots”, airport, hotels and convention center networks and more.
- Same approach can be used to analyze data by using MAC layer information in Flow Analysis tools to identify the activities and movements of nodes in Wireless Networks.

DALHOUSIE UNIVERSITY
Inspiring Minds

???

DALHOUSIE UNIVERSITY
Inspiring Minds

Design for Large-Scale Collection System Using Flow Mediators

Atsushi Kobayashi, Tsuyoshi Kondoh, and
Keisuke Ishibashi
NTT Information Sharing Laboratories

Outline

- Introduction
  - Why do we need a large-scale collection system?
  - What is Flow Mediator?
- Requirements
  - I tried to explore the possibility of a large-scale collection system for large networks.
- Heuristic method of designing traffic collection system
  - Estimate number of flow records after aggregation or sampling
  - Adjust several parameters based on this result
- Summary

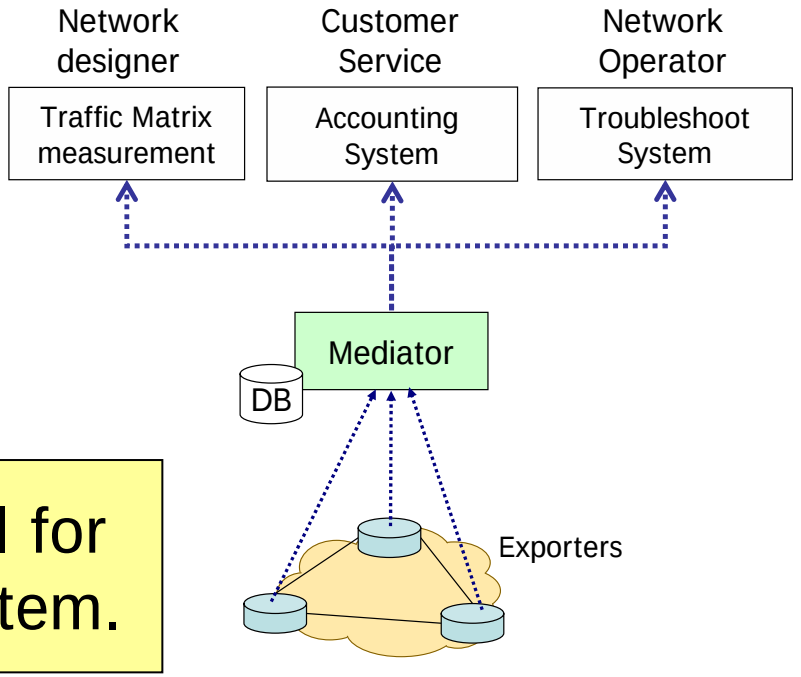
Introduction

- Traffic volumes in ISP networks are becoming huge in the last few years.
  - The number of exported flow records is becoming so huge that a single collector cannot handle them.
- A smaller sampling rate makes small flows invisible.
  - Even if traffic grows, network operators would like to maintain the same sampling rate as much as possible.
- Aggregated flow records from router make port number or IP address invisible.
  - Exporting 5-tuple flow records from router is better.

The demand for a large-scale traffic-collection system is growing.

What is Flow Mediator?

- Flow Mediator† is a device that “mediates” flow records and has the following functions:
  - collects Flow Records from various exporters
  - stores original flow records
  - aggregates flow records flexibly
  - distributes appropriate flow records for collectors/analyzers

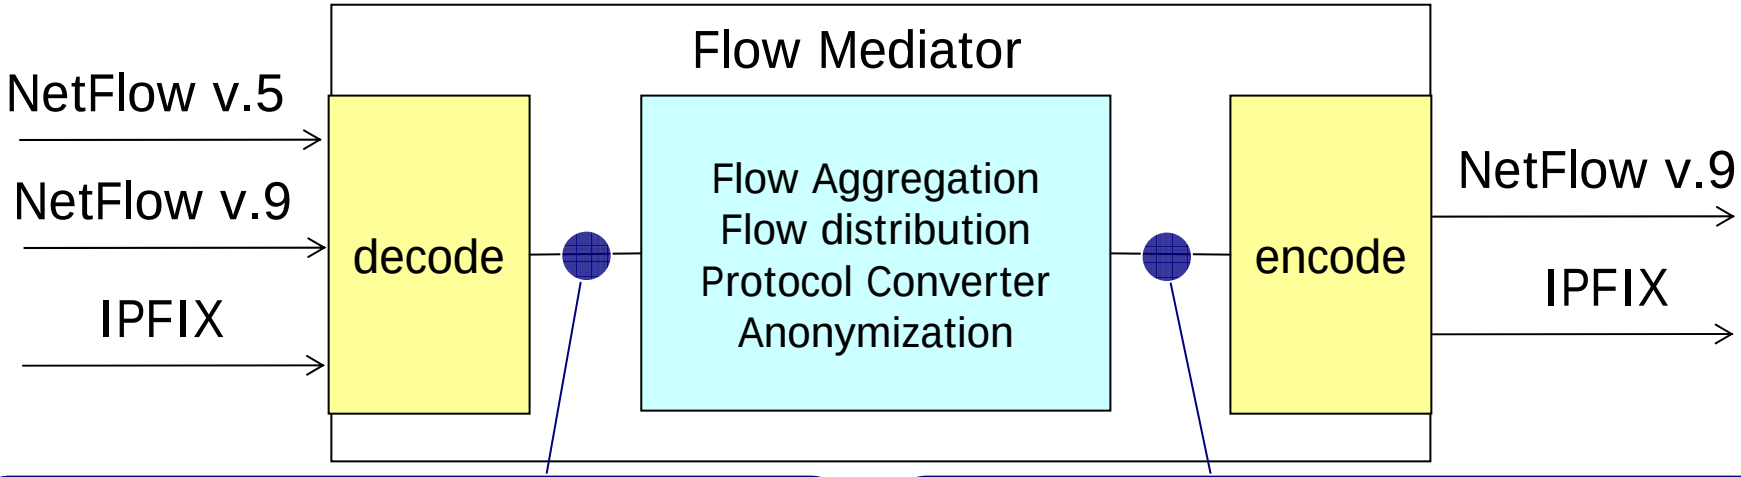


Flow mediator ought to be useful for making large-scale collection system.

† draft-kobayashi-ipfix-mediator-model-01.txt

You can easily make Flow Mediation code

- Net::Flow perl module is available on CPAN.
  - http://search.cpan.org/~akoba/Net-Flow-0.02/
  - The module can encode and decode NetFlow/IPFIX packets.
  - The encoding and decoding functions have a similar IF.




```
my ( $HeaderHashRef,
     $TemplateArrayRef,
   $FlowArrayRef,
   $ErrorsArrayRef) =
   Net::Flow::decode(
                  \$packet,
                  $TemplateArrayRef  ) ;
```

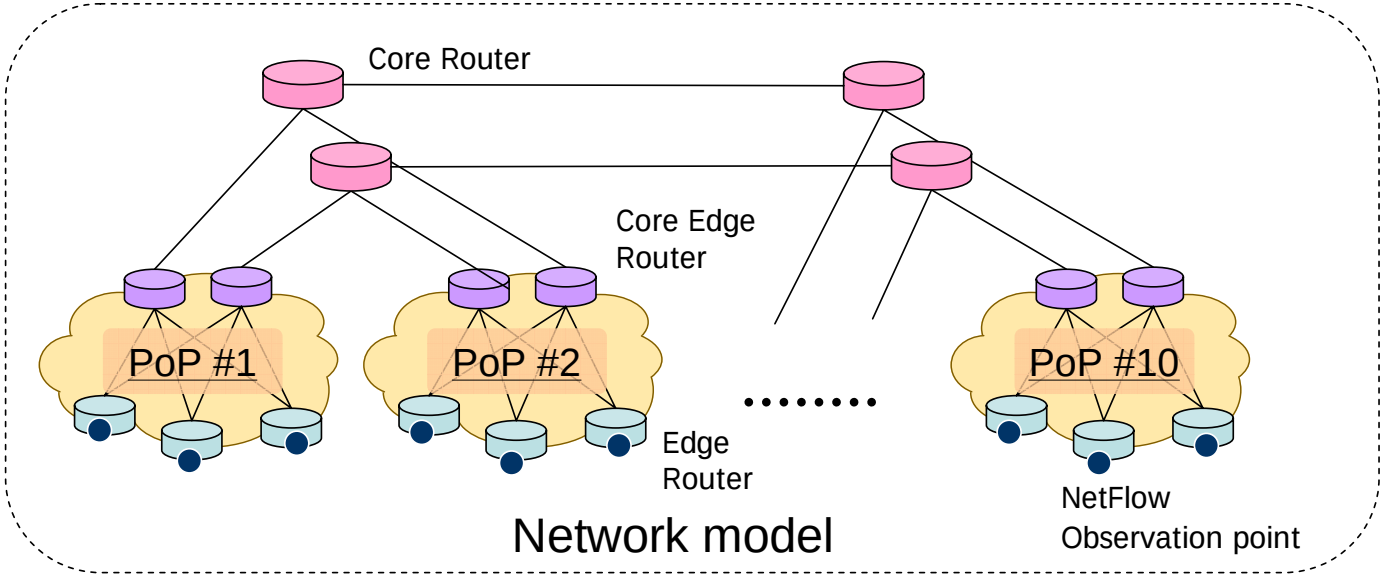
```
my ( $EncodeHeaderHashRef,
     $PktsArrayRef,
     $ErrorsArrayRef) =
     Net::Flow::encode(
                  $EncodeHeaderHashRef,
                  \@MyTemplates,
                  $FlowArrayRef,
                  1400 ) ;
```

Requirements

- Make traffic-collection system to meet following requirements
  - Requirement 1: measure traffic flow of entire networks
    - measure traffic matrices PoP by PoP and router by router
  - Requirement 2: store received 5-tuple flow records from router
    - When traffic incident happens, allow inspection of traffic.
  - Requirement 3: design scalable architecture to accommodate large ISP traffic volume

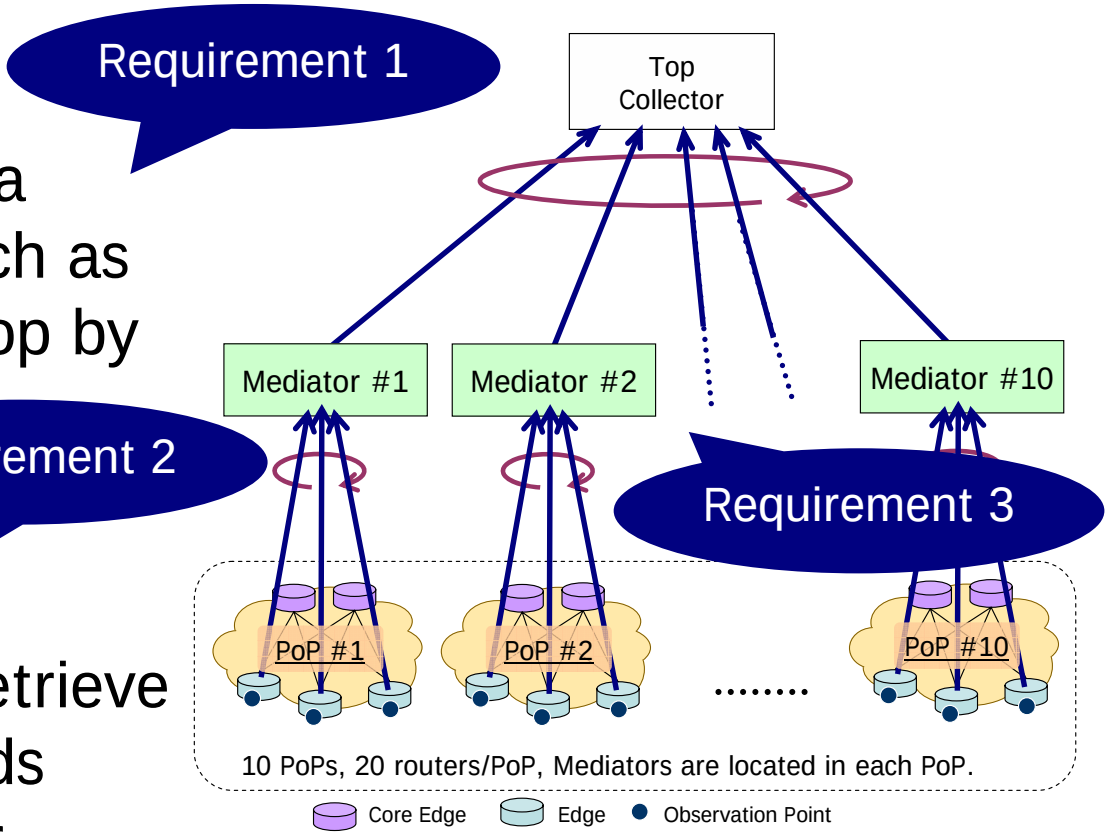
Goal

- Explore heuristic method of designing collection system for introduction into actual network.
- Proposed collection system needs to accommodate following network model.
  - Total traffic volume 500 Gb/s, 100 Mp/s
    - Edge Router 20/PoP×10 PoP = 200
    - NetFlow is enabled on IngressIF of Edge router.



Hierarchical Collection System

- Mediators are allocated in each PoP.
  - They store all flow records, aggregate them, and export them to next collector.
- Top Collector
  - measures wide-area traffic matrices, such as router by router, pop by pop.
- Inspection
  - If traffic incident happens, we can retrieve detailed flow records from Flow Mediator.

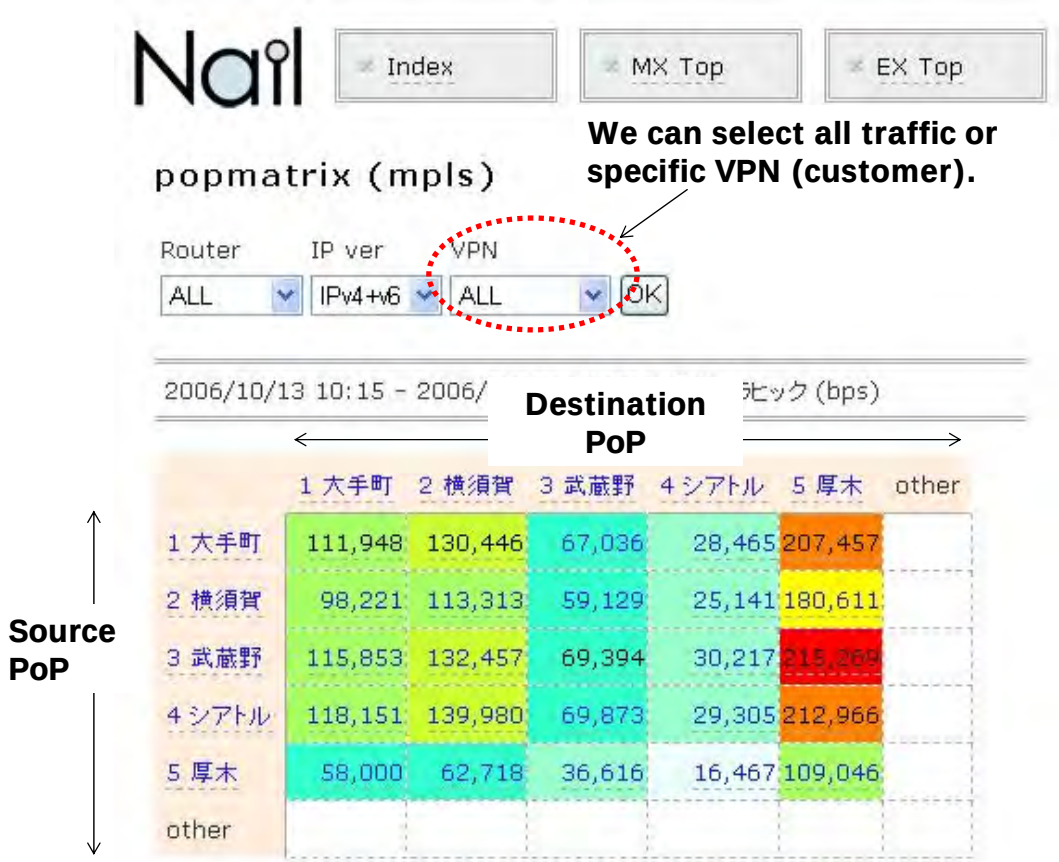


Visualize Traffic Matrices

- Top collector can visualize Router/PoP/AS Traffic Matrixes.

Nail is the name of our traffic matrix visualizer.

Color indicates traffic volume of Source/ Destination pair.

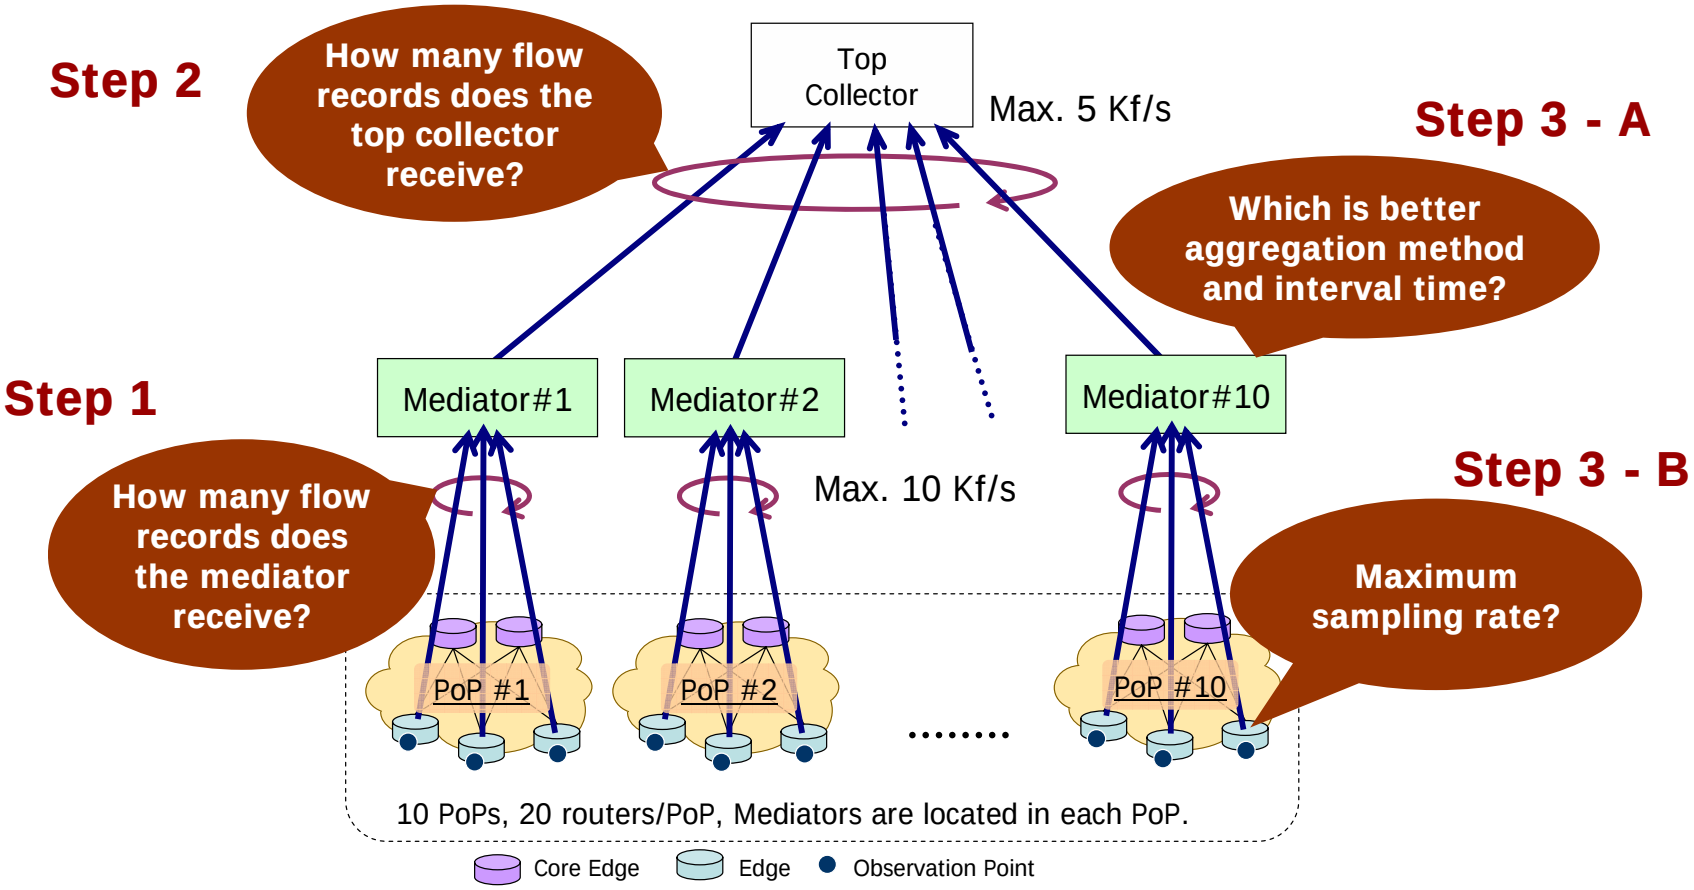


Heuristic Design Method

- Suitable values of several parameters are decided by the following steps.
  - Step 0: measure performance limit of flow mediator and top collector.
  - Step 1: reveal relation between number of flow records and packet sampling
  - Step 2: reveal relation between number of flow records and aggregation that depends on several factors.
    - Aggregation methods (BGP Next-Hop, Prefix, host)
    - Aggregation interval time (20 s, 60 s, 90 s...)
  - Step 3: select suitable value within performance limit.
    - Large sampling rate is preferable.
    - Small granularity of aggregation is preferable.

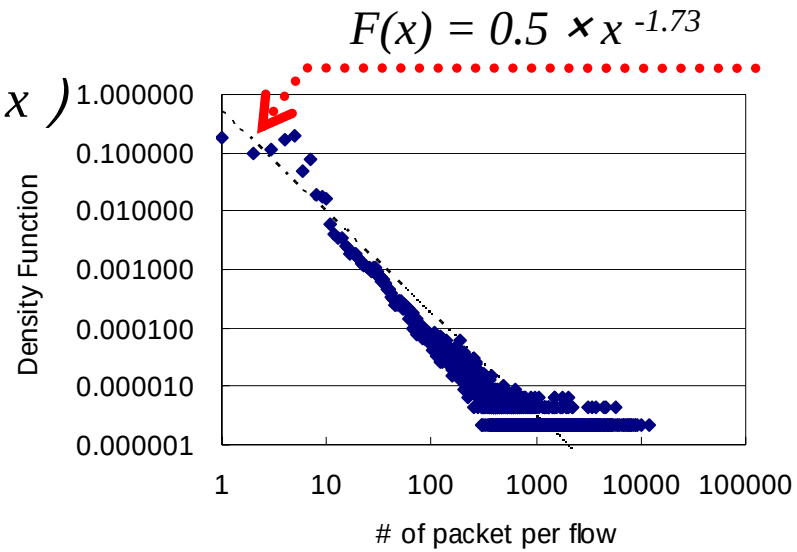
Consideration Points

- List several considerations, as follows.
  - Maximum performances of the top collector and mediators are 5 Kf/s and 10 Kf/s.



Step 1: estimate flow records after sampling

- Estimate number of flow records based on density function of packets per flow .
  - # of packets per flow: x
  - Packets per flow density function: $F(x)$
  - Sampling rate: $1/r$
  - Total number of unsampled flow: f_{all}



$$f_{sampled} = \sum_{x=1}^{\infty} \left(1-(1-1/r)^x\right) \times F(x) \times f_{all}$$

Extraction probability

$0.5x^{-1.73}$

Roughly estimate as follows.
100 Mpps ÷ 20 packets = 5 Mf/s

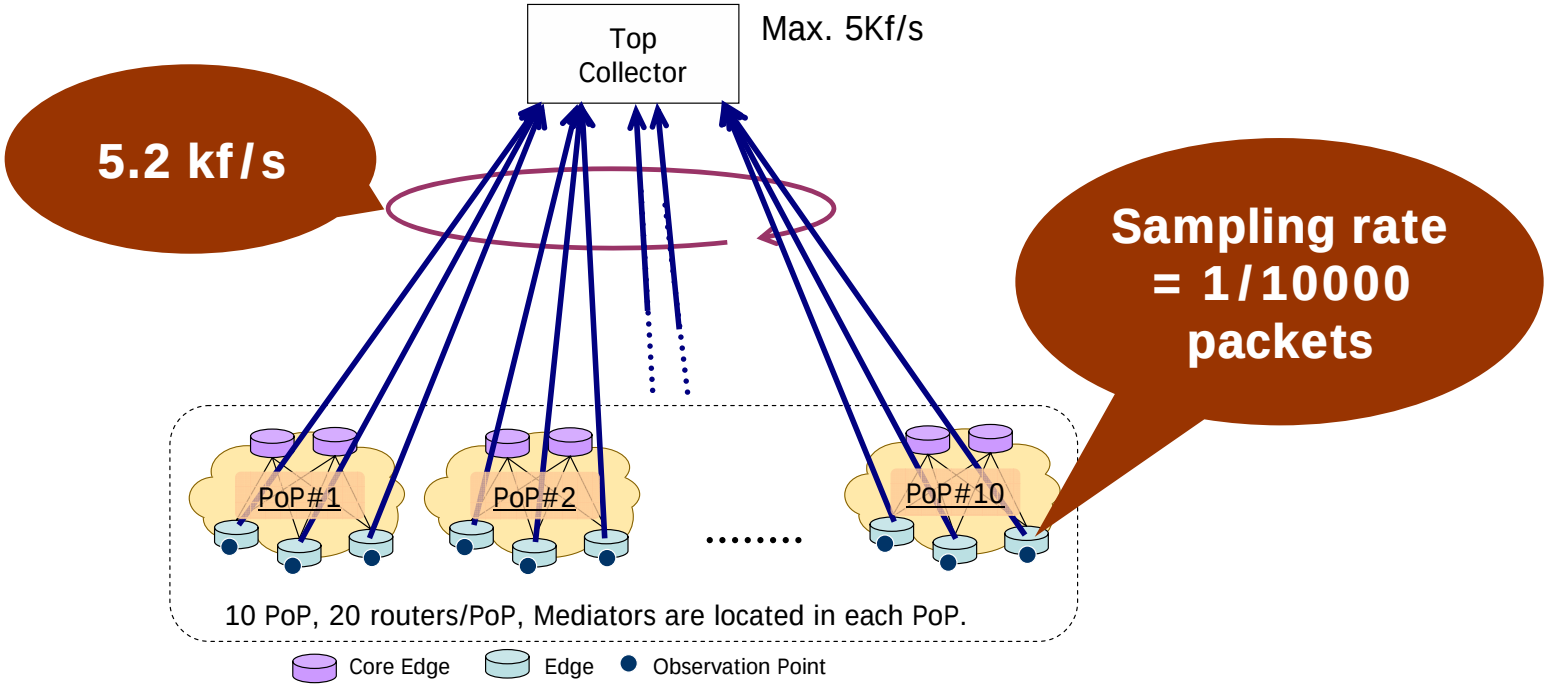
Approximate # of flows when total traffic volume is 500 Gb/s.

Sampling rate	1/100	1/1000	1/10000
$f_{sampled}$	*305 kf/s*	*43 kf/s*	*5.2 kf/s*

Too many flow records without mediator

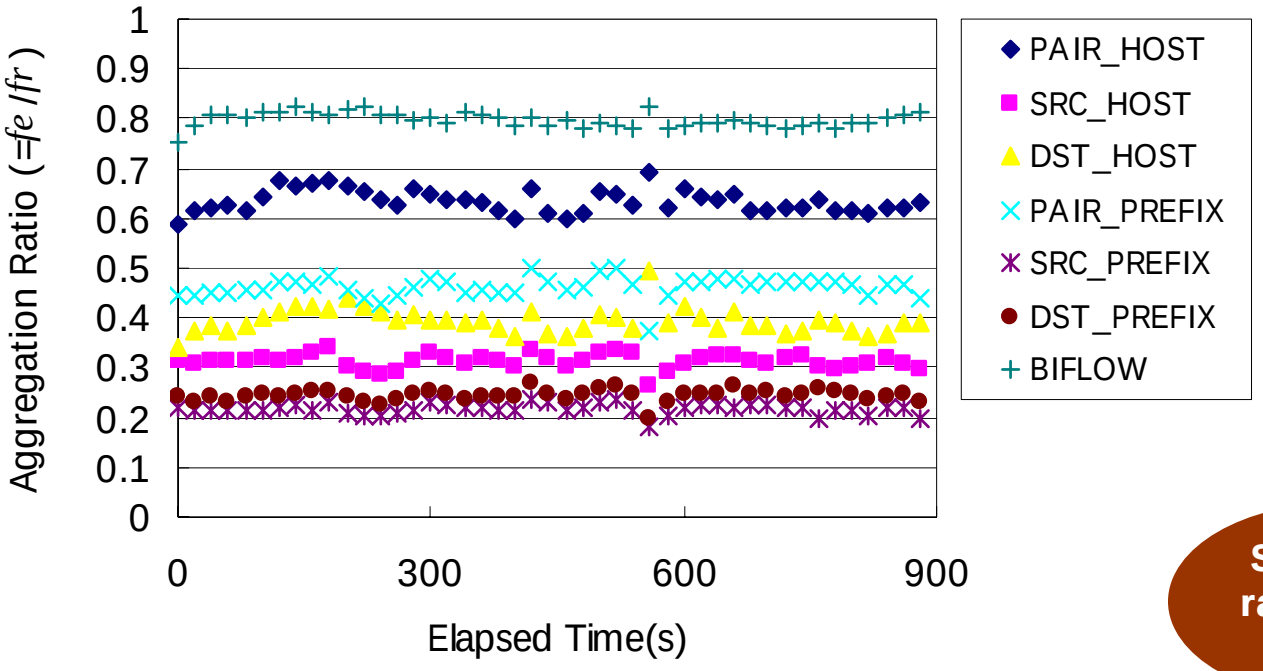
- Even if sampling rate is 1/10,000 packets, the number of flow records exceeds performance limit.

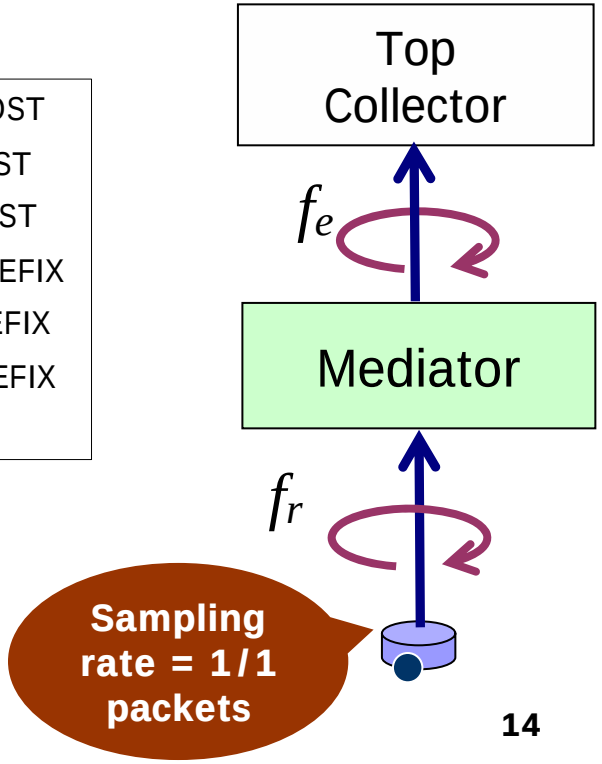
Sampling rate	1/100	1/1000	1/10000
$f_{sampled}$	*305 kf/s*	*43 kf/s*	*5.2 kf/s*



Step 2: flow records after aggregation

- What is the # of flow records after aggregation?
- Mediator aggregates unsampled flow records at 20-second interval.
  - Aggregation efficiency: Prefix > HOST > Pair Prefix > Pair HOST > Bi-Flow
    - The prefix length "/24" is uniformly applied to Prefix Aggregation.
    - Bi-flow is aggregated from two flow directions.

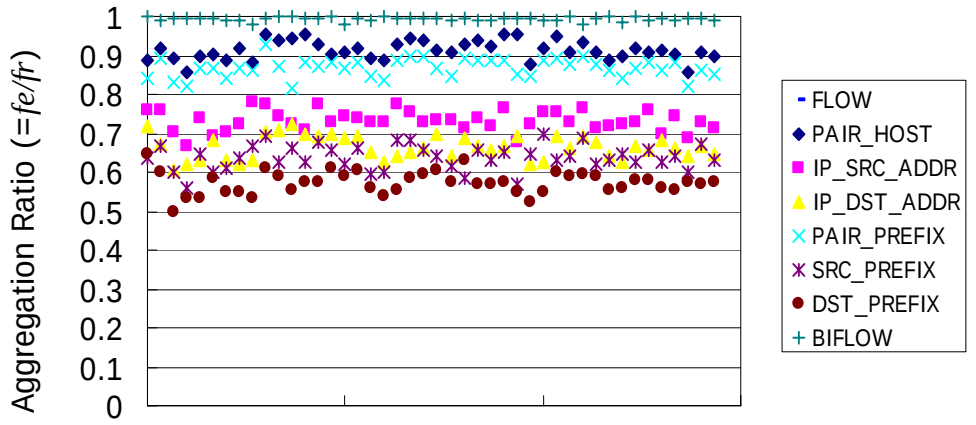




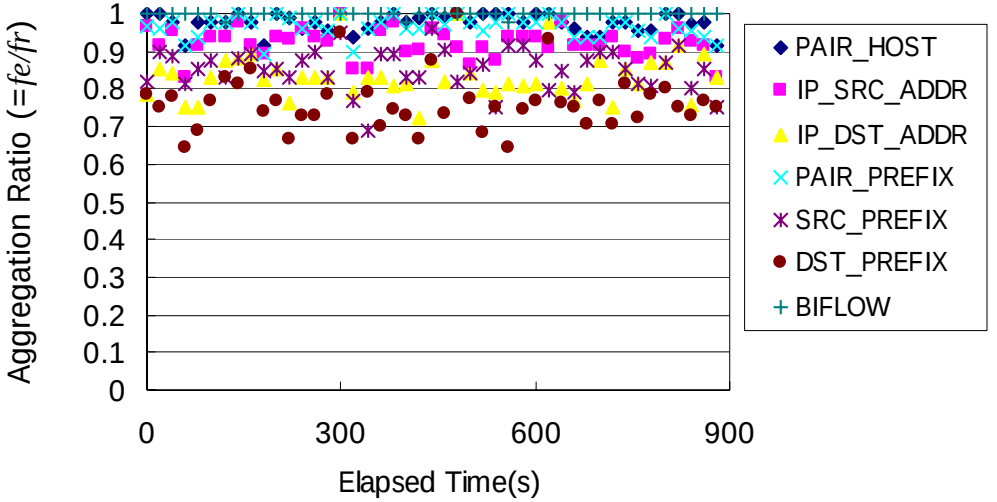
Step 2: Flow records after aggregation, sampling

- Each aggregation method becomes ineffective gradually.
- Bi-flow becomes ineffective immediately.
  - sensitive to sampling rate.

Sampling rate 1/128



Sampling rate 1/1024



Step 2: Which factor influences aggregation?

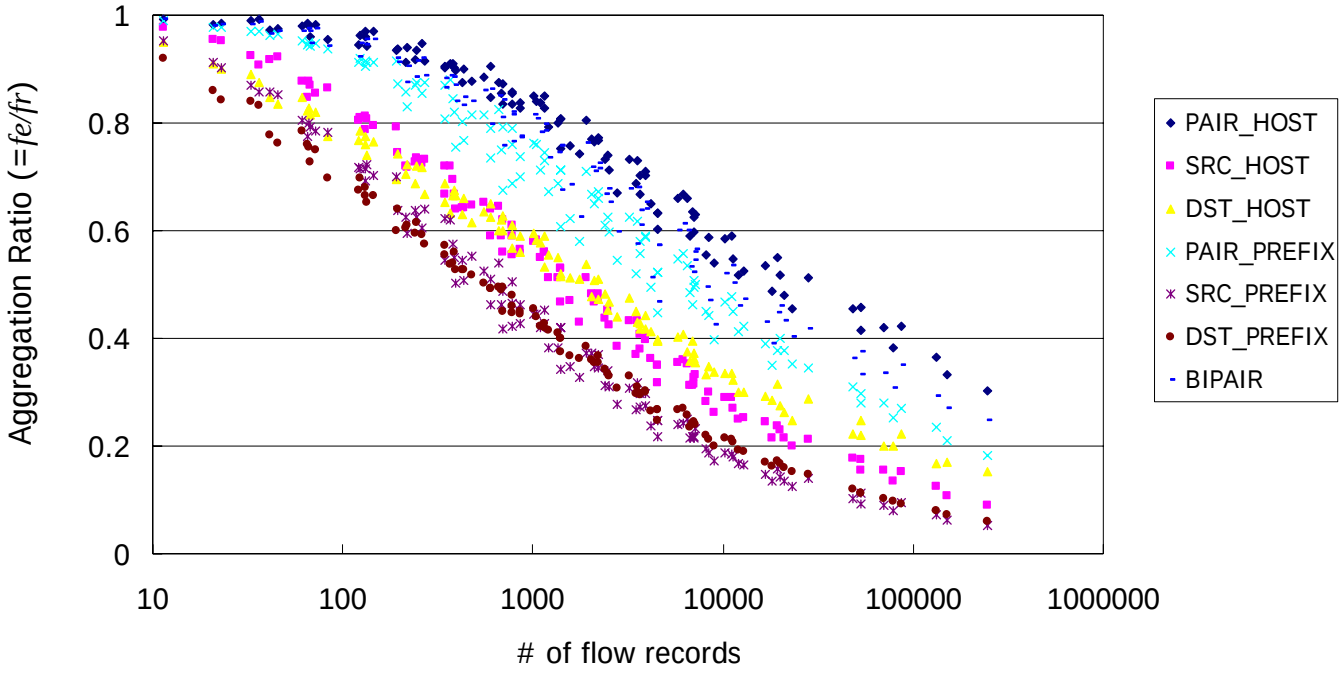
- Aggregation ratio depends on several factors.
  - Traffic Volume through observation point.
  - Sampling rate
  - Aggregation interval time

I guess that the aggregation ratio depends on the number of flow records received in interval time.

Received Flows	3450	3562
Aggregation Interval Time (s)	10	300
Sampling rate (1/r)	1	128
DST_HOST Aggregation ratio	45%	43%
DST_PREFIX Aggregation ratio	30%	32%

Step 2: Which factor influences aggregation?

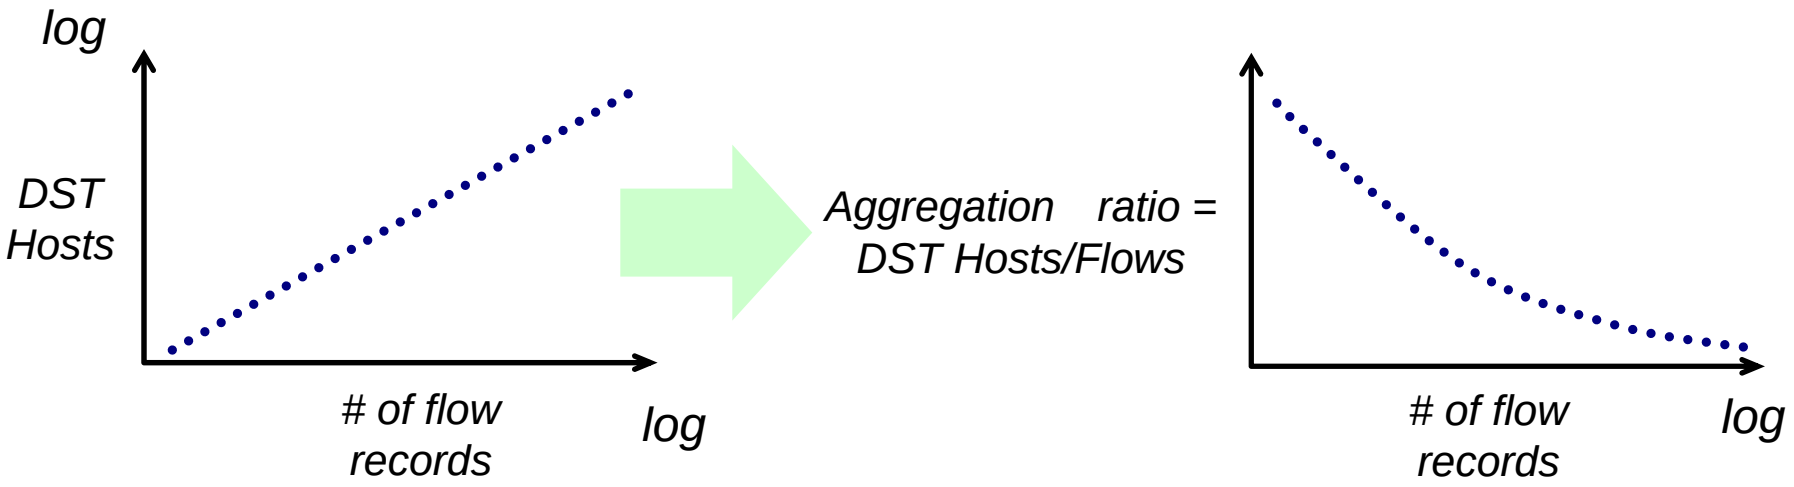
- I plotted all experimental data into one graph.
  - Three MAWI traffic data samples have different volumes.
  - Aggregation Interval time : 5 – 300s
  - Sampling rate : 1/1 – 1/1024



Aggregation ratio depends on number of received flow records.

Step 2: Formulation of Aggregation Ratio

- Aggregation ratio (R) can be estimated from number of flow records (f_r), as follows.
  - DST Host aggregation: $R_{dsthost} = 1.80 \times f_r^{-0.18}$
  - DST Prefix aggregation: $R_{dstprefix} = 2.34 \times f_r^{-0.26}$
- After all, the aggregation ratio depends on the # of unique hosts or prefixes versus # of flows.



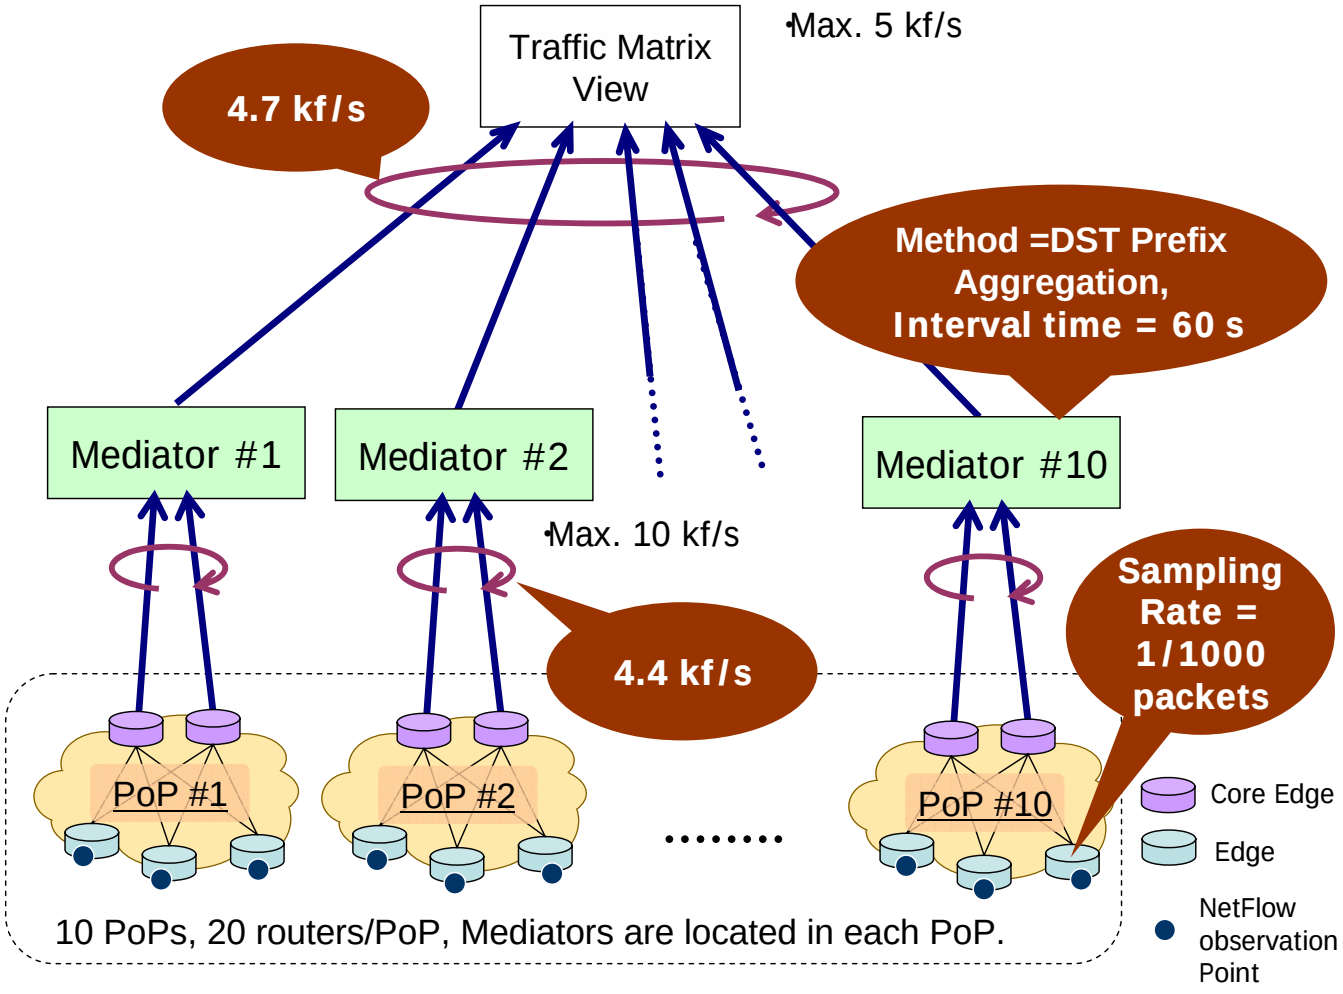
Step 3: Selection of Suitable Values

- I selected suitable value within performance limit.

		Sampling Rate	1/100	1/1000	1/10000
# of received flow records in top collector ($=\sum f_e$)	DST_HOST aggregation	Interval time = 60s	*45 kf/s*	9.0 kf/s	***1.6 kf/s***
	DST_Prefix aggregation	Interval time = 60s	*21 kf/s*	***4.7 kf/s***	***0.94 kf/s***
	DST_HOST aggregation	Interval time = 300s	*34 kf/s*	7.0 kf/s	***1.2 kf/s***
	DST_Prefix aggregation	Interval time = 300s	*12 kf/s*	***3.0 kf/s***	***0.62 kf/s***
# of received flow records in mediator (f_r)	—	—	*30 kf/s*	***4.4 kf/s***	***0.6 kf/s***

Example of collection system

- Sampling Rate: 1/1000
- Aggregation Interval time: 60 s



Conclusion

- To make large scale traffic collection system, flow mediator is efficient.
- Revealed relation between number of flow records and several factors:
  - Traffic volume
  - Sampling rate
  - Aggregation method
  - Aggregation interval time
- Demonstrated that traffic collection system using mediator can be introduced into actual large-scale networks.

Thank you for your attention.

This study was supported by the Ministry of Internal Affairs and Communications of Japan.



Abnormal traffic detection and alert

Yiming Gong
XO Communications

http://security.zz.ha.cn/flocon2008.pdf

Flocon 2008

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

The problem and request

- XO network
  - OC-192 IP backbone with OC-12 uplinks in our markets and data centers, AS 2828
- Backbone level abnormal traffic detection
  - netflow

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

The problem and request

- Commercial product not good enough
  - You get what GUI gives you
  - Very likely to miss low volume traffic attack
    - (storm worm, scans)
  - By default, alert based on thresholds
  - Lacking data mining ability
  - Cost

- Free flow-based tool
  - Powerful but you need tell them what to do

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

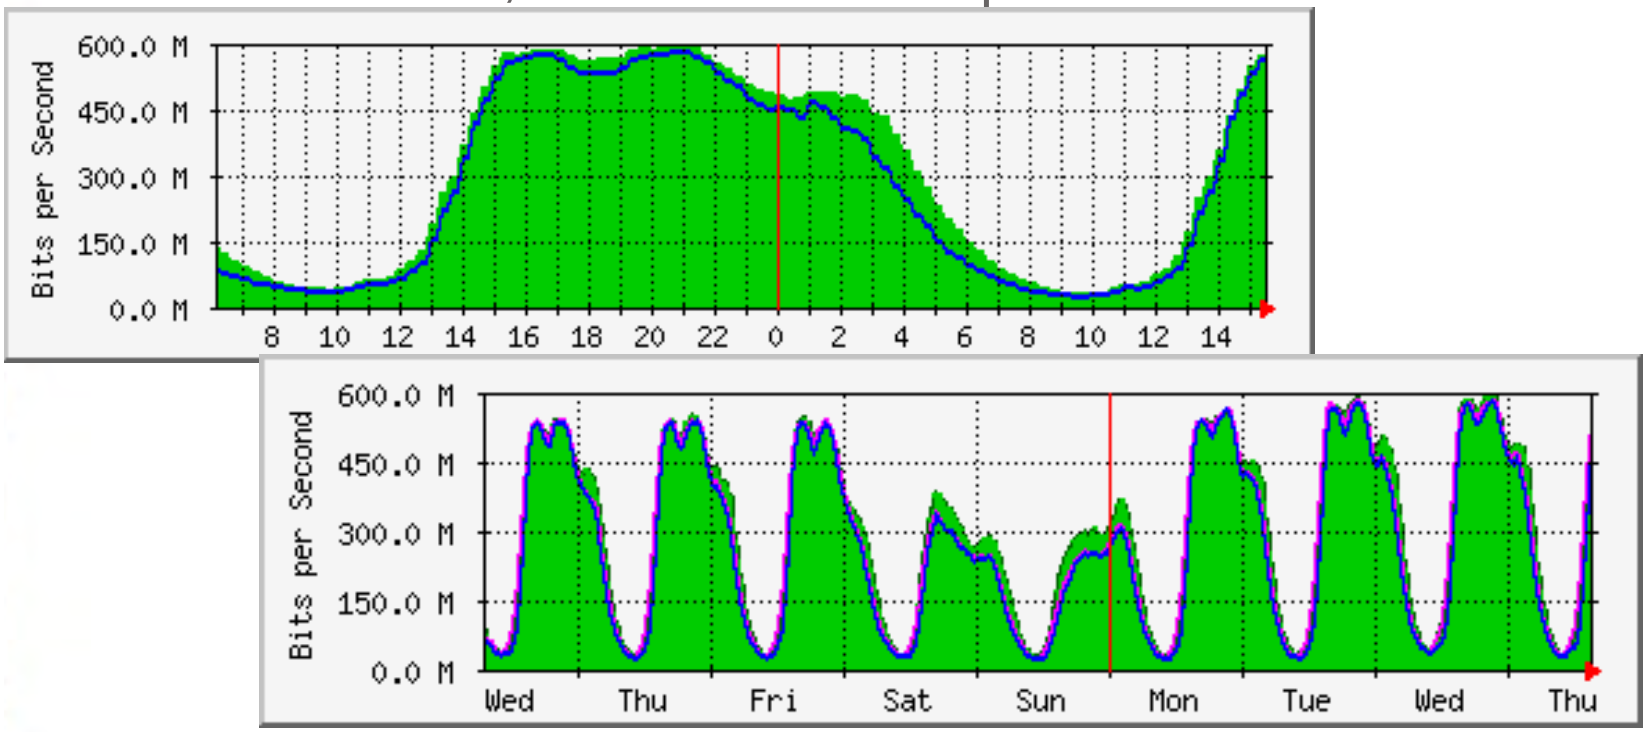
So what we want

- Detect network abnormal traffic
  - both low and high volume
- Non-threshold based
- Automatically
- Fully controlled and customized
- Data mining
- Better be free

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

In a perfect world

- Smooth curve, recurrent traffic pattern



- Spike means……?

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Our thought

- Break down raw netflow records to
  - TCP SYN, UDP total, ICMP type|code, protocol on each IFIndex of each edge router
    - Session
    - Traffic
- For each element
  - establish a weekly traffic profile
  - Profile is a band
- When
  - real data higher than the tolerant (upper) band
  - Match some other conditions

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Head results

- $ head -4 SYN_session_routerx_past_5_minutes
- 9 3005
- 75 2844
- 76 2121
- 8 2120

- $ head -4 SYN_traffic_routerx_past_5_minutes
- 9 137792
- 75 128952
- 8 101084
- 76 100092

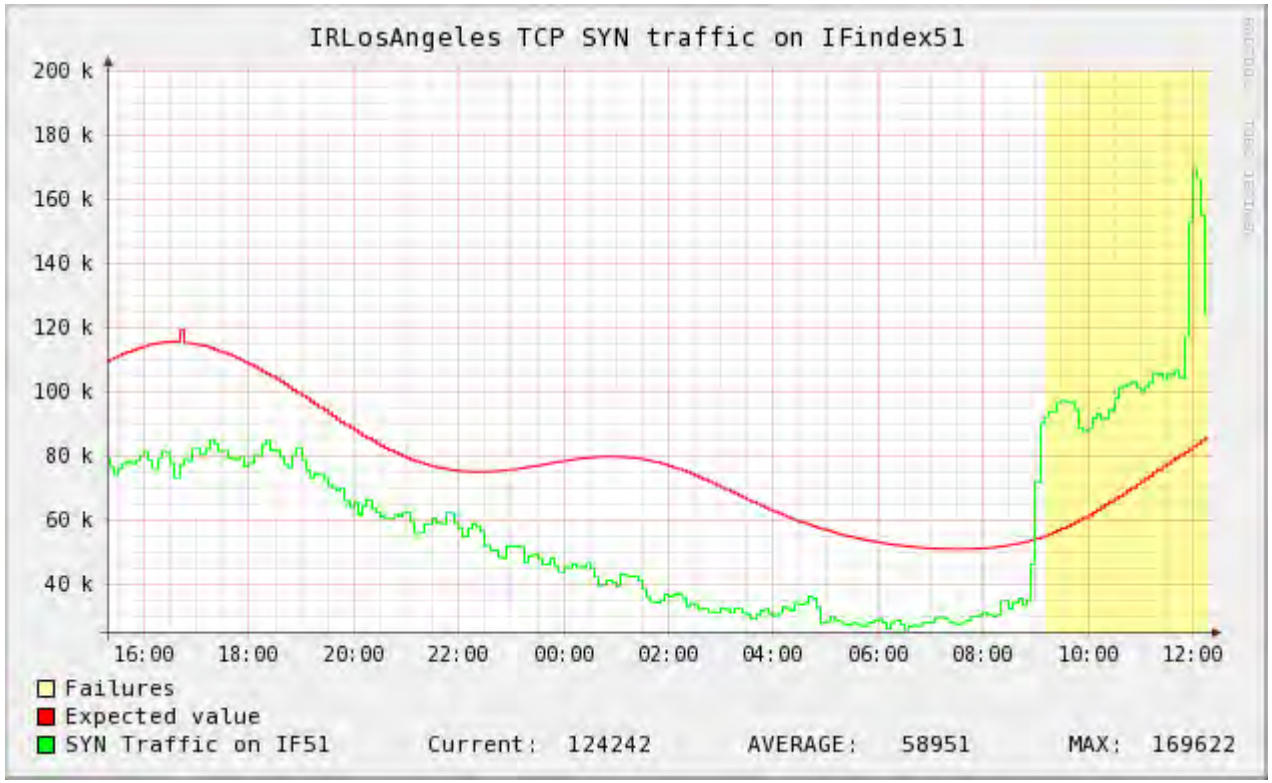
- $ head -4 PROTO_session_routerx_past_5_minutes
- 6 668344
- 17 104205
- 50 22725
- 1 4517

- $ head -4 ICMP_typecode_session_routerx_past_5_minutes
- 0 5431
- 2048 1953
- 2816 792
- 771 586

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Dynamic profile

- example



CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Dynamic profile

- Establishing a profile
  - Using NFDUMP receive, store and process netflow data
  - rrdtool with aberrant behavior module
  - rrdtool (http://oss.oetiker.ch/rrdtool/)
  - aberrant behavior module
    - Learns from past values and uses them to predict the future

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Dynamic profile

```
yiming> more IR-syn-Amsterdam
13 1864
9 144
21 85

rrdtool create IR-syn-Amsterdam.rrd -s 300
DS:13:GAUGE:1200:0:U       \
DS:9:GAUGE:1200:0:U       \
DS:21:GAUGE:1200:0:U       \
RRA:HWPREDICT:2016:0.001:0.0035:288

rrdtool tune IR-syn-Amsterdam.rrd  --deltapos 8
#deltapos set the scale parameter for the upper tolerant band
#different element should use different value
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Failure

- Only an entry
  - IR-syn-Amsterdam:[1196800800]RRA[FAILURES][1]DS[13] = 1.0000000000e+00
  - Need script do the trace back work
- Every 10 minutes, scans the rrd output for failures
- now rrdtool generates a failure alert, so what?

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

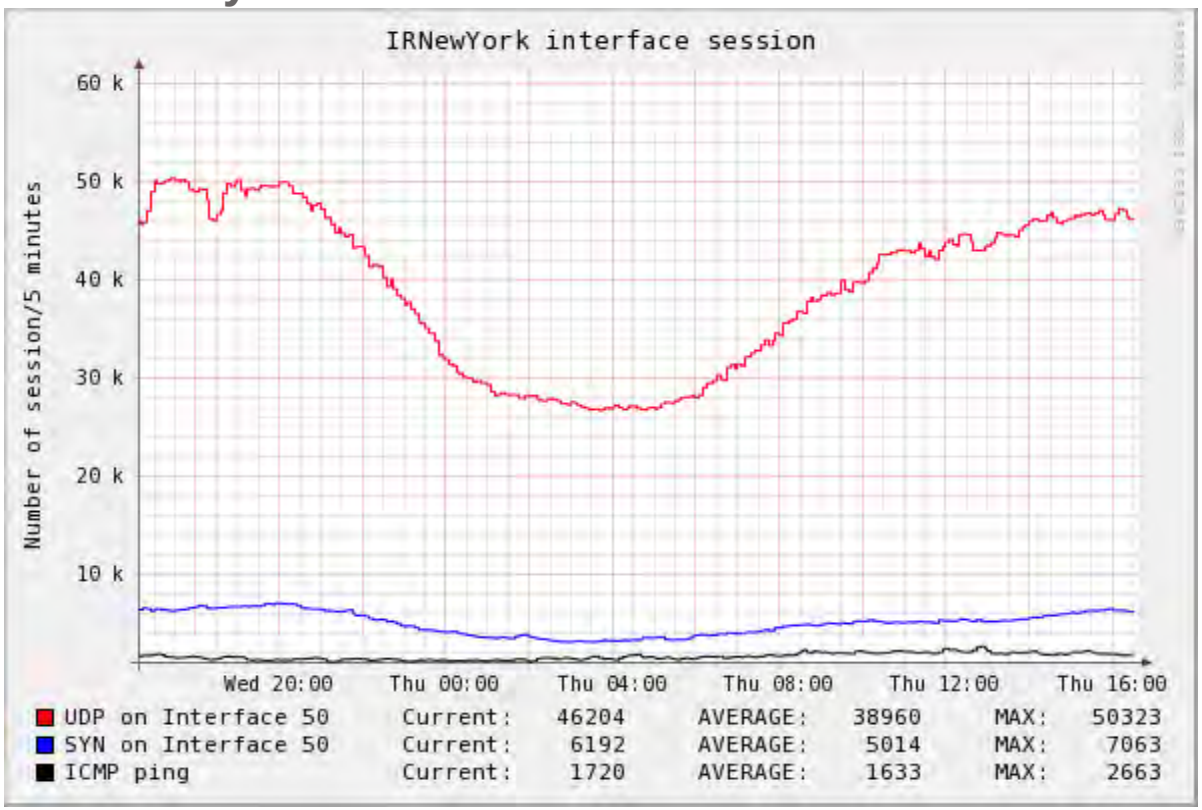
Failure

- Tracking down
  - Past_5|10 minutes_flow of 'TCP + SYN bit only + IFindex 13 + router Amsterdam'
  - "and"
  - Who is|are behind the spike?
    - **Spike should be caused by one or several hosts**
    - **these hosts can be either victims, attackers or normal hosts**
      - **Scan -> attacker**
      - **DoS|DDos –> victim or attacker**
      - **Email server and others**
    - **They have too many sessions or traffic**
      - **How many is too many?**

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Finding active host

- **For different protocol, different network, the definition of "too many" is different**



- **Alert! ICMP ping Used to be 1, now is 10!**

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Finding active host

- After rrdtool generates a failure

```
session-icmp*)
  alert-trigger-number="500";    #conditon a
  flowfilter="proto icmp and port 2048 and if $if";
  session-generated-by-single-host="280"; #conditon b
;;
  session-syn*)
  alert-trigger-number="2000";
  flowfilter="proto tcp and flags 2 and if $if";
  session-generated-by-single-host ="600“;
```

- A failure matches all the conditions can be regarded as a real failure and further actions will be needed

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Netflow records

- Pull out necessary data
- Generate alert
  - Picture, email

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert

- Scan alert

```
>IR LosAngeles has 5462 sessions on proto tcp and flags 2 and if 50 in 5 minutes

50 = STRING: [redacted]
50 = STRING: [redacted]

>Snapshot picture

http://[redacted]LosAngeles-50-abnormal.png

>One week|month picture

http://[redacted]LosAngeles-50-abnormal-week.png

http://[redacted]LosAngeles-50-abnormal-month.png

>Top IPs in 10 minutes

Date first seen          Duration Proto          IP Addr         Flows   Packets     Bytes       pps       bps   bpp
2007-12-05 08:51:02.520   289.718 any      218.233.1[redacted]    2114      2114     84560         7      2334    40
2007-12-05 08:51:20.493   130.413 any        218.234[redacted]     605       605     24200         4      1484    40

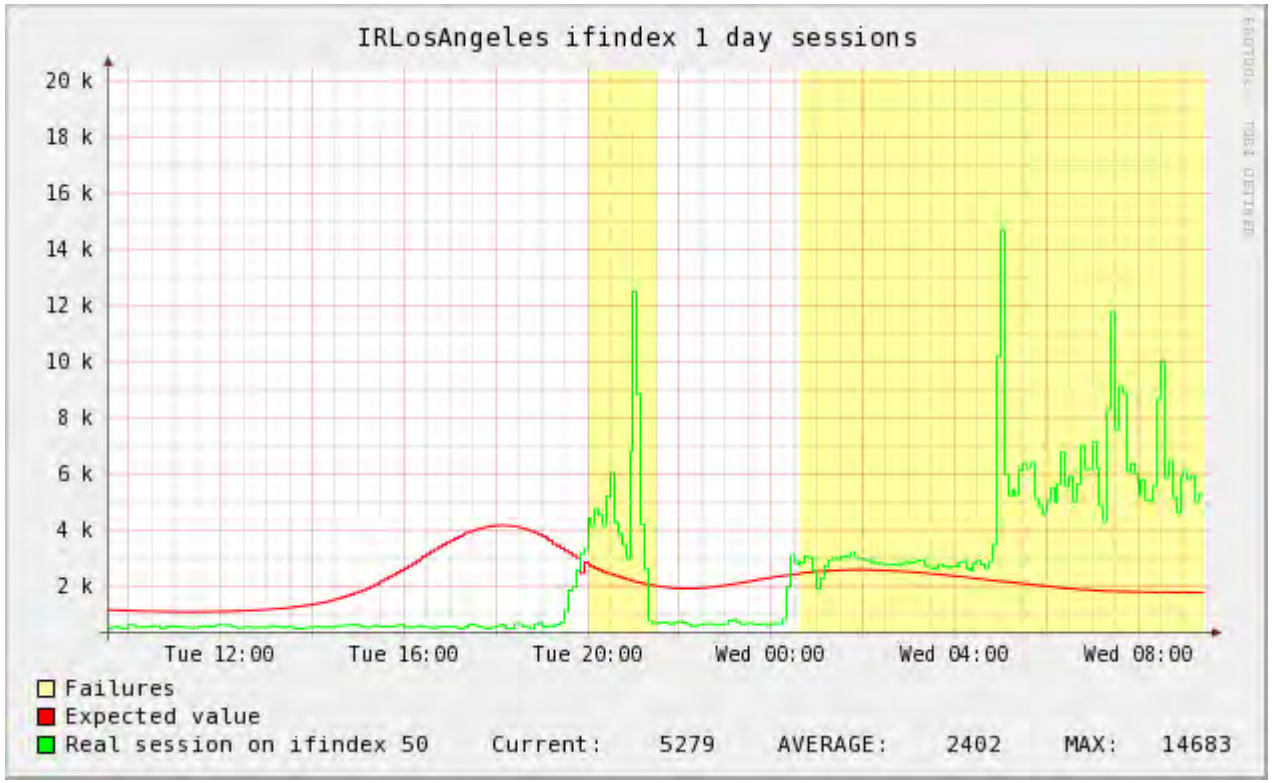
>Top IP info

* AS            IP              AS name                                    FQDN
[redacted] | 218.233.[redacted] | [redacted] Telecom Inc.   |
[redacted] | 218.234.[redacted] | [redacted] Telecom Inc.   |
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert

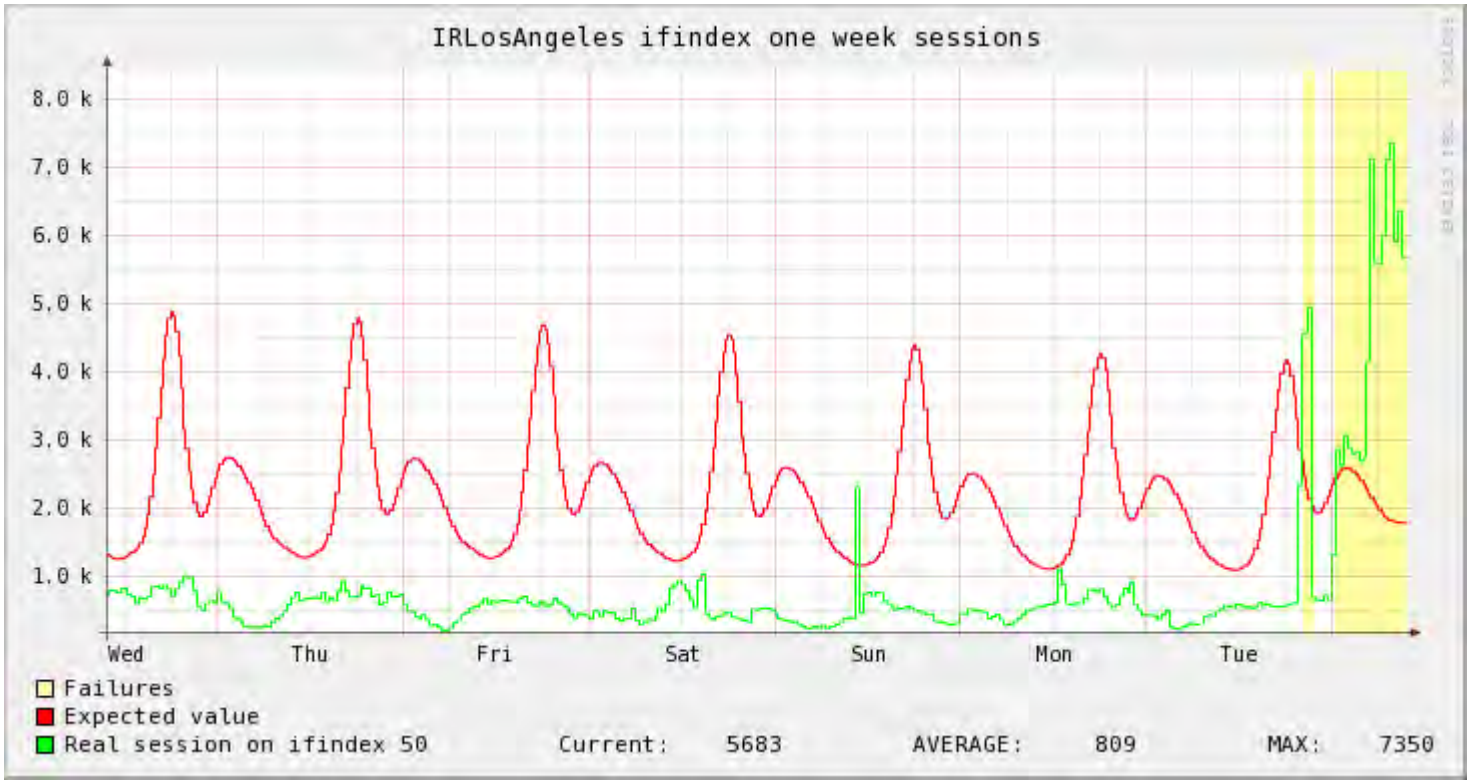
- Day



CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert

- Week



CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert

- Scan alert

```
>Top IP detail

/ ip 218.233.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted]65.106.6.[redacted]  6.621 ms
 6  [redacted]gin.ntt.net (206.223.[redacted]19.12)  7.089 ms
 7  [redacted]ntt.net ([redacted])  7.317 ms
 8  [redacted].gin.ntt.net ([redacted]5.20)  73.034 ms
 9  [redacted]ntt.net ([redacted])  73.922 ms

**Protocol summary for 218.233.198.25

Proto  Flows  Packets    Bytes      pps      bps   bpp
  6     2116     2116    84640        7     2337    40
 17        1        1      257        0        0   257

**Sampled netflow records

TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
TCP     218.233.[redacted]:6000      65.99[redacted]:7212  ....S.      40   0  [redacted]  50   73
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert

- Scan alert

```
/ ip 218.234.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted]106.6.[redacted] ms
 6  [redacted].ntt.[redacted]119.12)  7.135 ms
 7  [redacted].net ([redacted])  7.268 ms
 8  [redacted].ntt.[redacted].5.20)  79.209 ms
 9  [redacted].net ([redacted])  74.073 ms

**Protocol summary for 218.234.[redacted]

Proto  Flows  Packets    Bytes      pps      bps   bpp
 6       605      605    24200        4     1484    40

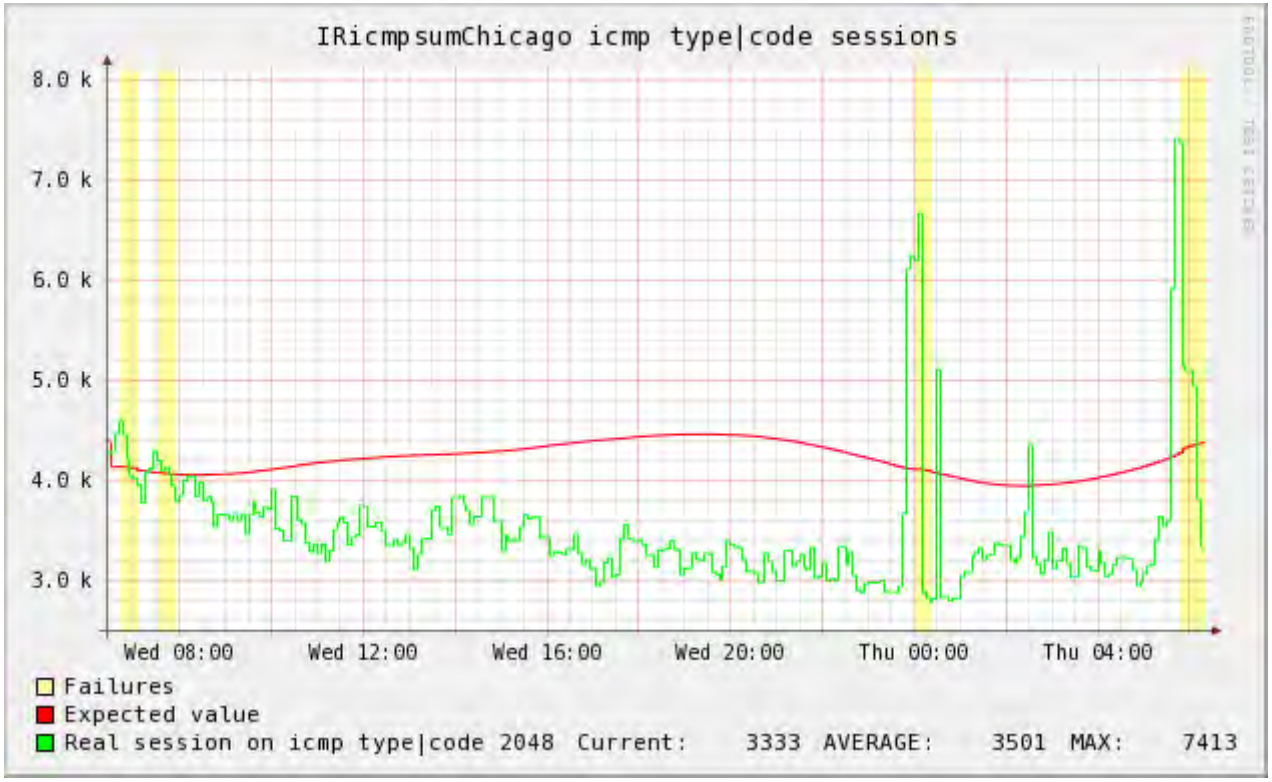
**Sampled netflow records

TCP       218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP       218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP       218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP       218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP       218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP       218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP       218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP       218.234.[redacted]:6000      71.60.[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP       218.234.[redacted]:6000      71.60.1[redacted]:6588  ....S.      40   0  [redacted]  50   73
TCP       218.234.[redacted]:6000      71.60.1[redacted]:6588  ....S.      40   0  [redacted]  50   73
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

alert

- Storm worm



CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert - one week later

- DDos

```
IR LosAngeles has 177002 sessions on proto tcp and flags 2 and if 50 in 5 minutes

50 = STRING:[redacted]
50 = STRING:[redacted]

>Snapshot picture

http://[redacted]LosAngeles-50-abnormal.png

>One week|month picture

http://[redacted]LosAngeles-50-abnormal-week.png

http://[redacted]LosAngeles-50-abnormal-month.png

>Top IPs in 10 minutes

Date first seen          Duration Proto          IP Addr         Flows   Packets     Bytes       pps       bps    bpp
2007-12-12 09:00:23.705   320.317 any      89.144.[redacted]      173554    176183   10.0 M       550    261507     59
2007-12-12 09:00:28.361   297.573 any      211.211[redacted]        1048      1056     50688         3      1362     48
2007-12-12 09:00:43.293   282.273 any      211.206[redacted]         692       708     33984         2       963     48
2007-12-12 09:00:43.269   291.093 any      211.44.[redacted]         658       667     42688         2      1173     64
2007-12-12 09:00:43.401   289.437 any      218.48.[redacted]         633       684     32832         2       907     48
2007-12-12 09:00:37.353   288.445 any      123.214[redacted]         627       640     30720         2       852     48
2007-12-12 09:00:23.705   311.869 any      58.127.[redacted]         603       618     39552         1      1014     64

>Top IP info

* AS          IP              AS name                                                    FQDN
[redacted]  | 89.144.[redacted]  | [redacted]Autonomus System number for [redacted]Net  |  ;; connection timed
[redacted]  | 211.211[redacted]  | [redacted]Telecom Inc.  |
[redacted]  | 211.206[redacted]  | [redacted]Telecom Inc.  |
[redacted]  | 211.44.[redacted]  | [redacted]Telecom Inc.  |
[redacted]  | 218.48.[redacted]  | [redacted]Telecom Inc.  |
[redacted]  | 123.214[redacted]  | [redacted]Telecom Inc.  |
[redacted]  | 58.127.[redacted]  | [redacted]Telecom Inc.  |
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert - one week later

- Traceroute returns nothing

```
>Top IP detail

/ ip 89.144.[redacted]

**Traceroute (from hop 5 to 9)  <----------  no traceroute info here

**Protocol summary for 89.144.[redacted]

Proto  Flows  Packets    Bytes      pps      bps   bpp
 6    173974   176610   10.0 M      551   261950    59

**Sampled netflow records

TCP      219.254.[redacted]:2391      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  11
TCP    123.212.[redacted]:2735      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP    219.233.[redacted]:3878      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP     58.124.[redacted]:4375      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP    219.251.[redacted]:4049      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP     58.123.[redacted]:3642      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  11
TCP      123.21[redacted]:2340      89.144.[redacted]:80    ....S.    48  0  [redacted]  50  10
TCP     211.200[redacted]:4256      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  11
TCP    221.143.[redacted]:4313      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP     218.234[redacted]:4353      89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert - one week later

```
/ ip 211.211.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted]L.us.xo[redacted].85)  7.113 ms
 6  [redacted]et (65.[redacted]  521 ms
 7  [redacted]net (20[redacted]6.505 ms
 8  [redacted]et (65.[redacted]  604 ms
 9  [redacted].net (6[redacted]6.511 ms

**Protocol summary for 211.211.[redacted]

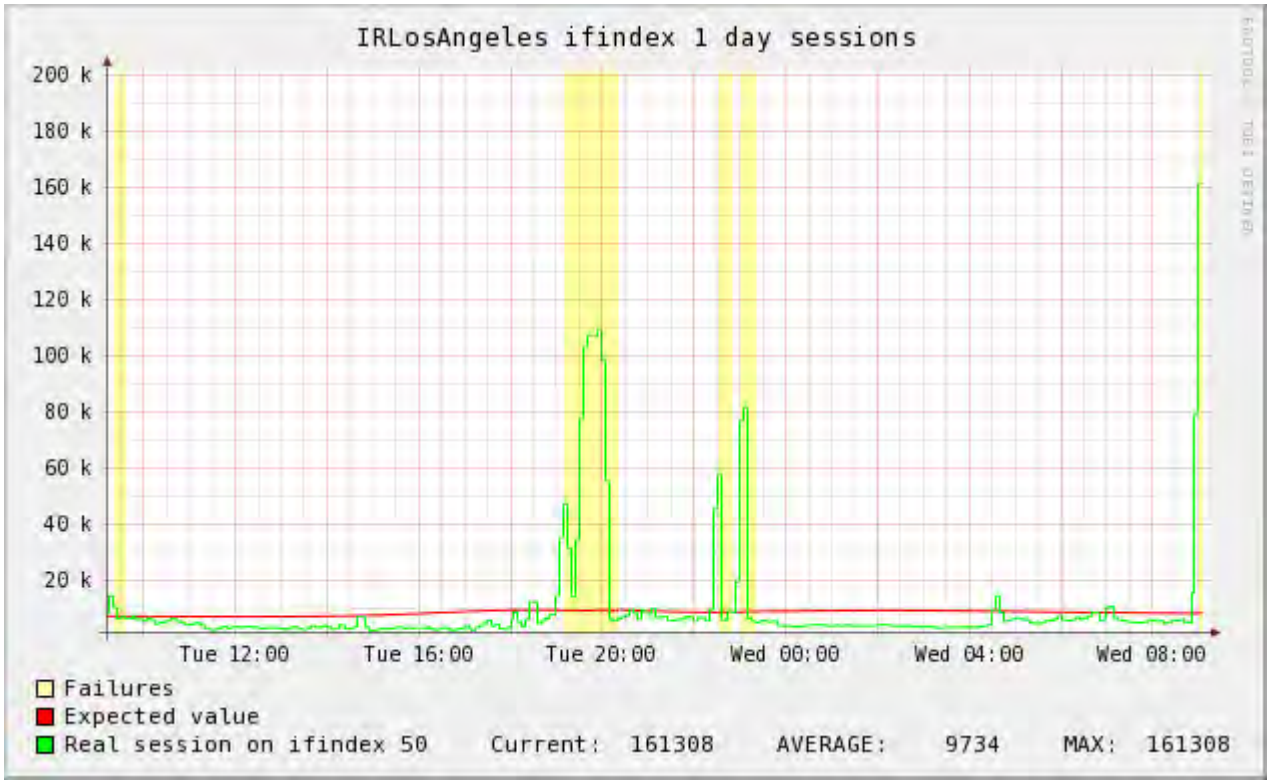
Proto  Flows  Packets    Bytes      pps      bps   bpp
 6     1057     1065    51120        3     1374    48

**Sampled netflow records

TCP    211.211.[redacted]:29937    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:32301    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:30596    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:35573    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:26497    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:31263    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:27378    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:34829    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:28267    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:59695    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert - one week later



CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert - whitelist

- Special customers

```
>Top IP info

* AS            IP              AS name                                    FQDN
[redacted]  | 64.39.[redacted]  | QUALYS-[redacted] Inc.   |  scanner[redacted].com.
[redacted]  | 63.245.[redacted] | MOZILLA-CORP - [redacted] Corporation  |  core2.[redacted].com.

>Top IP detail

/ ip 64.39.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted]
 6  [redacted].com ([redacted])  6.814 ms
 7  [redacted].com ([redacted])  66.692 ms
 8  [redacted]com ([redacted])  66.836 ms
 9  [redacted]com ([redacted])  66.824 ms

**Protocol summary for 64.39.[redacted]

Proto  Flows  Packets    Bytes      pps      bps   bpp
  6      182      200     8584        0      231    42
 17        1        1       58        0        0    58

**Sampled netflow records

TCP       64.39.[redacted]:2681    63.245.[redacted]:25     ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:37672   63.245.[redacted]:35459  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:38206   63.245.[redacted]:47123  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:38318   63.245.[redacted]:2870   ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:39700   63.245.[redacted]:34739  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:40293   63.245.[redacted]:26733  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:40210   63.245.[redacted]:53606  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:40603   63.245.[redacted]:63822  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:41626   63.245.[redacted]:55450  ....S.      40   0  [redacted]  31   8
TCP       64.39.[redacted]:41565   63.245.[redacted]:2361   ....S.      40   0  [redacted]  31   8
```


CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Alert – whitelist and misc

- Whitelist <cont>
  - Email servers
  - We don't want to miss real attack even if an IP is on whitelist

- Alert email
  - Suppression period
  - Subject
    - 12-05 abnormal sessions at LosAngeles proto tcp and flags 2 and if 50

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Data mining

- Database
  - 3 tables
    - IP,FQDN,AS
    - Summary
    - Raw netflow data
  - Data mining
    - Which peering neighbor sends out most attack traffic, who is the most attacked, which port is the most popular being scanned…etc.

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Data mining

- Database
  - 3rd party outside data
    - Dshield TOP 10000
    - Dshield AS
    - CBL data
    - Mynetwatchman
    - Our own darknet project output
    - Other private outside data
  - If XO host gets involved, these tables will be checked

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.



problem

- Problem
  - Peering neighbor
  - Alert correlation
    - But you can do it in database.

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

What you need

- Nfdump (or any other free flow software), rrdtool, mysql, net-snmp, dig, apache, some unix commands
- A box

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

- For more info
  - yiming.gong@xo.com
  - http://security.zz.ha.cn

- Thanks!

CONFIDENTIAL © 2007 XO. ALL RIGHTS RESERVED. XO, THE XO DESIGN LOGO, XOPTIONS AND ALL RELATED MARKS ARE TRADEMARKS OF XO. ALL OTHER TRADEMARKS ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Visual Representations of Flow Data

and the Value of Visual Language

Presented by Sunny Fugate
Space and Naval Warfare Systems Center, San Diego



Copyright 2007 Sunny Fugate

Human-Machine Efficiency

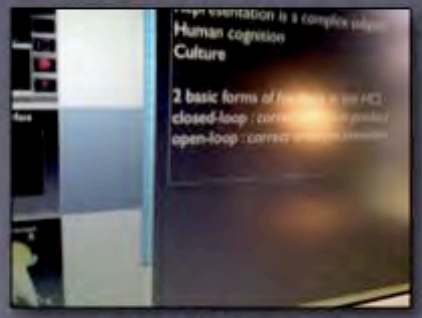
Over-Learned: **Feedback**

non-volitional feedback



haptic

volitional feedback



visual / aural

© 2007 Sunny Fugate

Human-Machine Efficiency

Over-Learned: **Feedback**

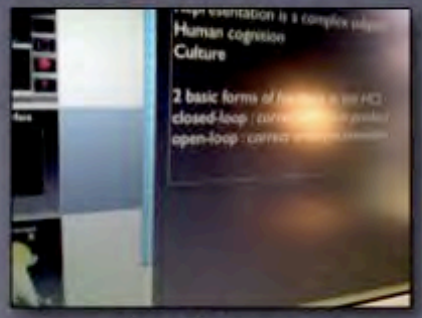
non-volitional feedback



correct errors in **production**

haptic

volitional feedback



visual / aural

© 2007 Sunny Fugate

Human-Machine Efficiency

Over-Learned: **Feedback**

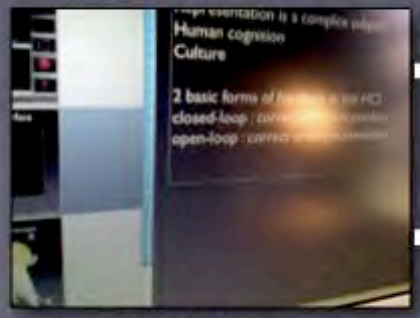
non-volitional feedback



haptic

correct errors in **production**

volitional feedback



visual / aural

correct errors in **semantics**

© 2007 Sunny Fugate

Human-Machine Efficiency

Over-Learned: **Feedback** - haptic vs visual/aural

Haptic Feedback

Visual / Aural Feedback

Sequential access



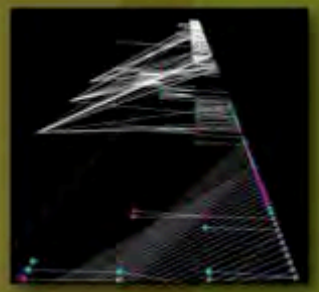
iFeel™



Sidewinder™
Force Feedback



Falcon™





joystick



mouse

Random access



keyboard



data/gesture glove





multi-touch



voice control

© 2007 Sunny Fugate

Human-Machine Efficiency

Over-Learned: **Feedback** - haptic vs visual/aural

Haptic Feedback

Visual / Aural Feedback

Sequential access



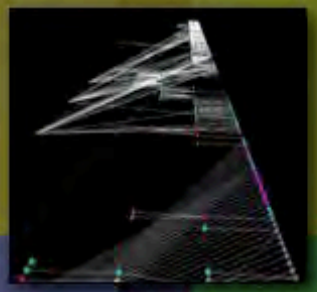
iFeel™



Sidewinder™
Force Feedback



Falcon™





joystick



mouse

Random access



keyboard



data/gesture glove





multi-touch



voice control

© 2007 Sunny Fugate

Human-Machine Efficiency

Over-Learned: **Feedback** - haptic vs visual/aural

Haptic Feedback

Visual / Aural Feedback

Sequential access



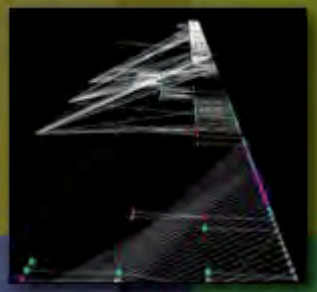
iFeel™



Sidewinder™
Force Feedback



Falcon™





joystick



mouse

Random access





keyboard



data/gesture glove



multi-touch



voice control

© 2007 Sunny Fugate

Human-Machine Efficiency

Over-Learned: **Feedback** - haptic vs visual/aural

Haptic Feedback

Visual / Aural Feedback

Sequential access



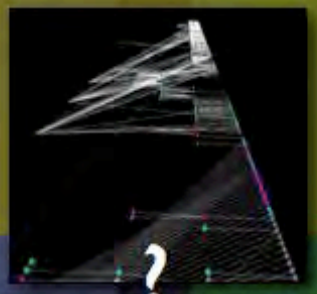
iFeel™



Sidewinder™
Force Feedback



Falcon™





joystick



mouse

Random access



keyboard



data/gesture glove





multi-touch

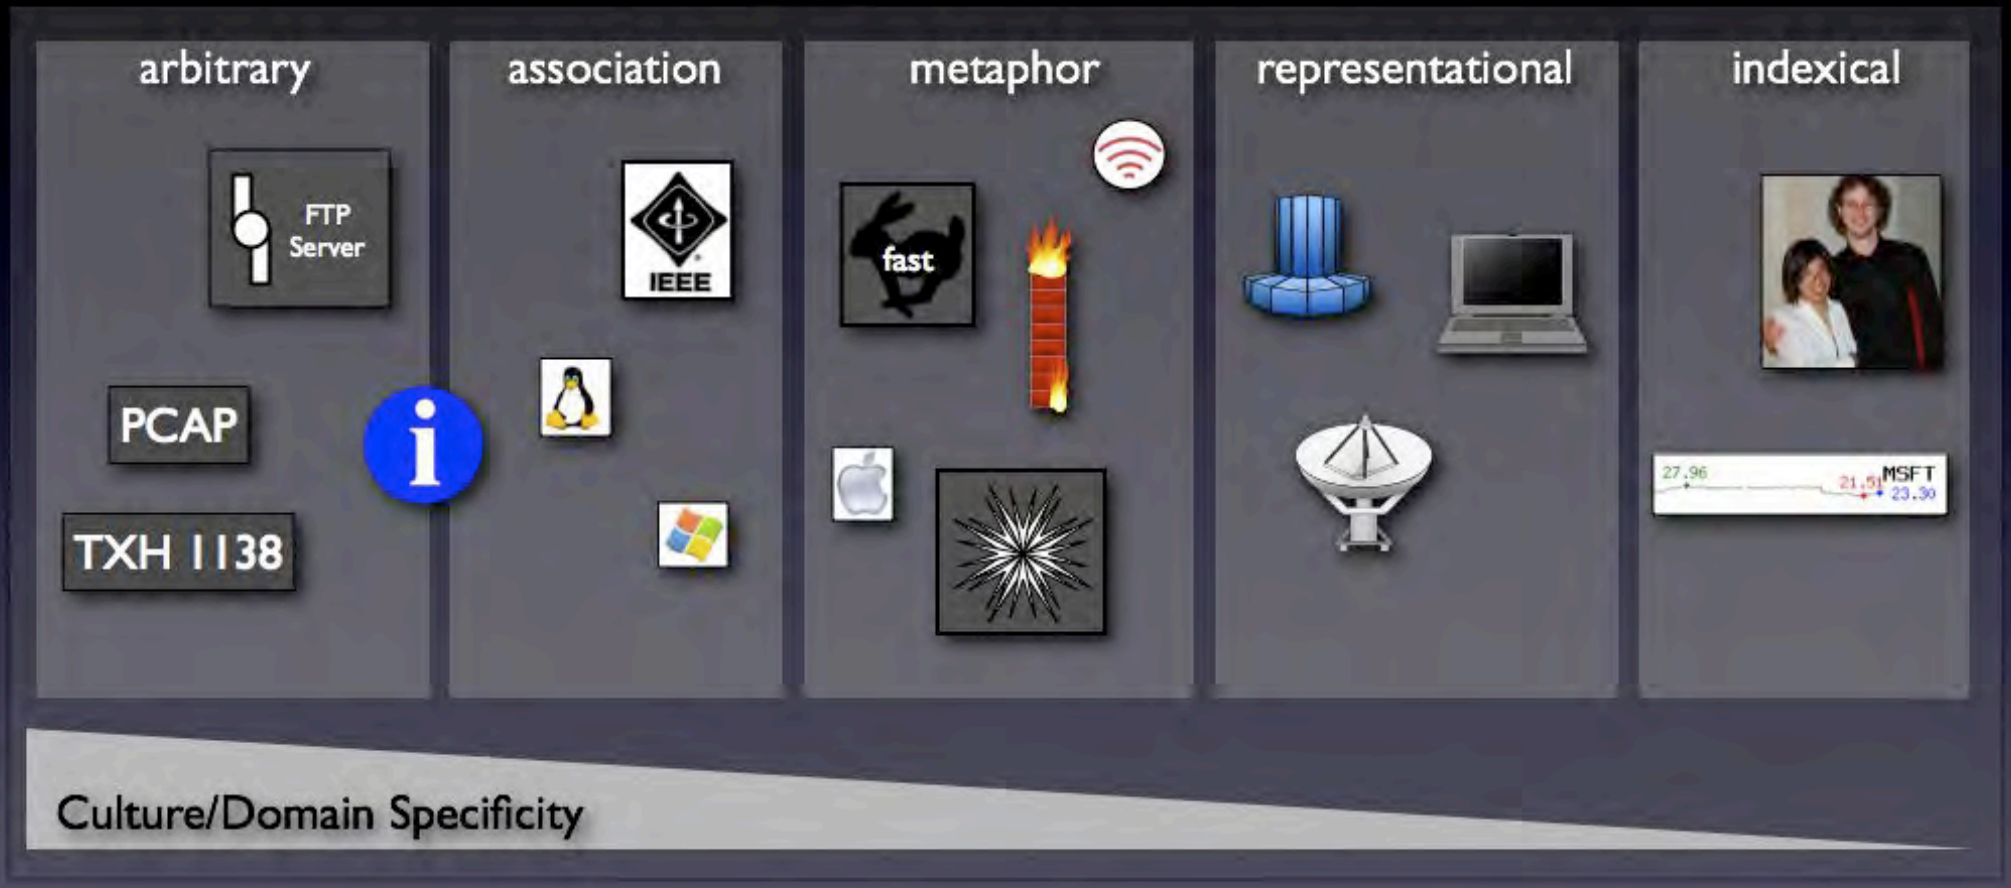


voice control

© 2007 Sunny Fugate

Human-Machine Efficiency

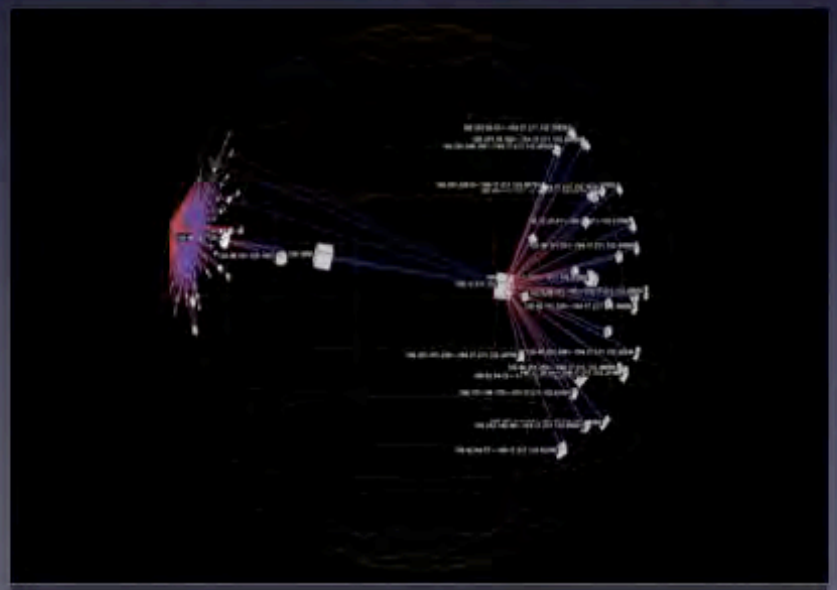
Under-Learned: **Representation**

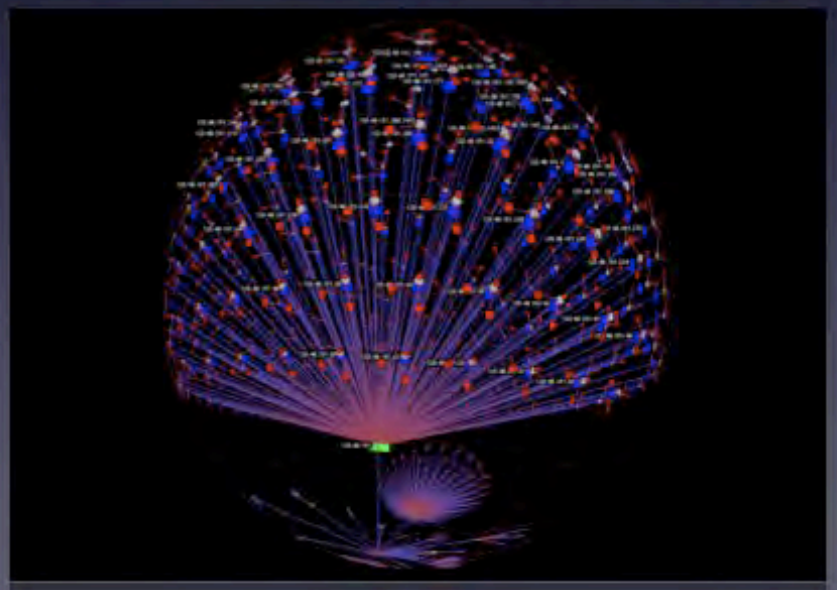


© 2007 Sunny Fugate

Flow in hyperbolic space

- 3 month SSC project in 2002
- discover and apply network visualization tools
- **Hyperviewer**: quasi-hierarchical hyperbolic space
- '**fish-eye**' 3-d
- Created by Stanford researcher **Tamara Munzner**













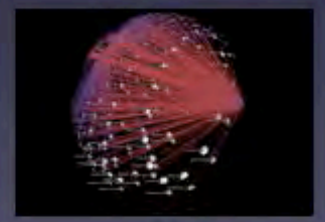
© 2007 Sunny Fugate

Flow in hyperbolic space

- Easily adapted to a forced-hierarchy view of flow
- **Opensource** C++ library and UI
- Experimented with visual methods

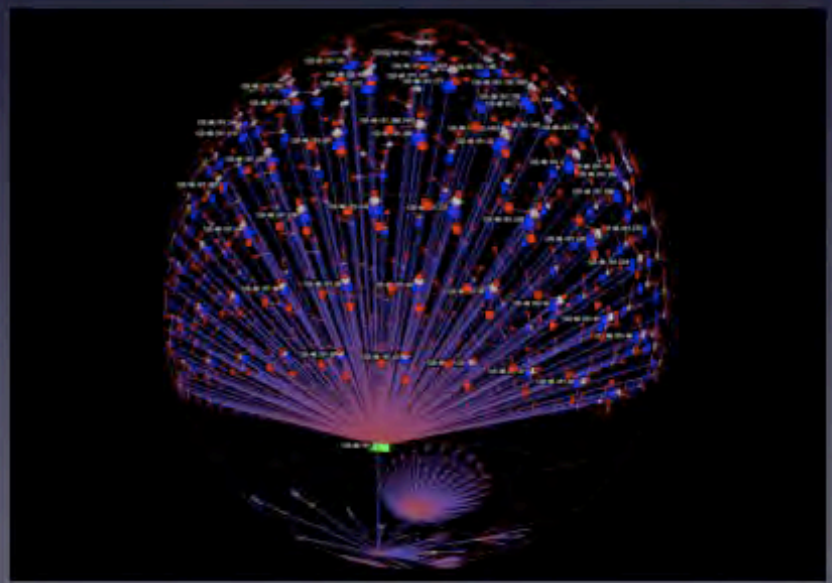




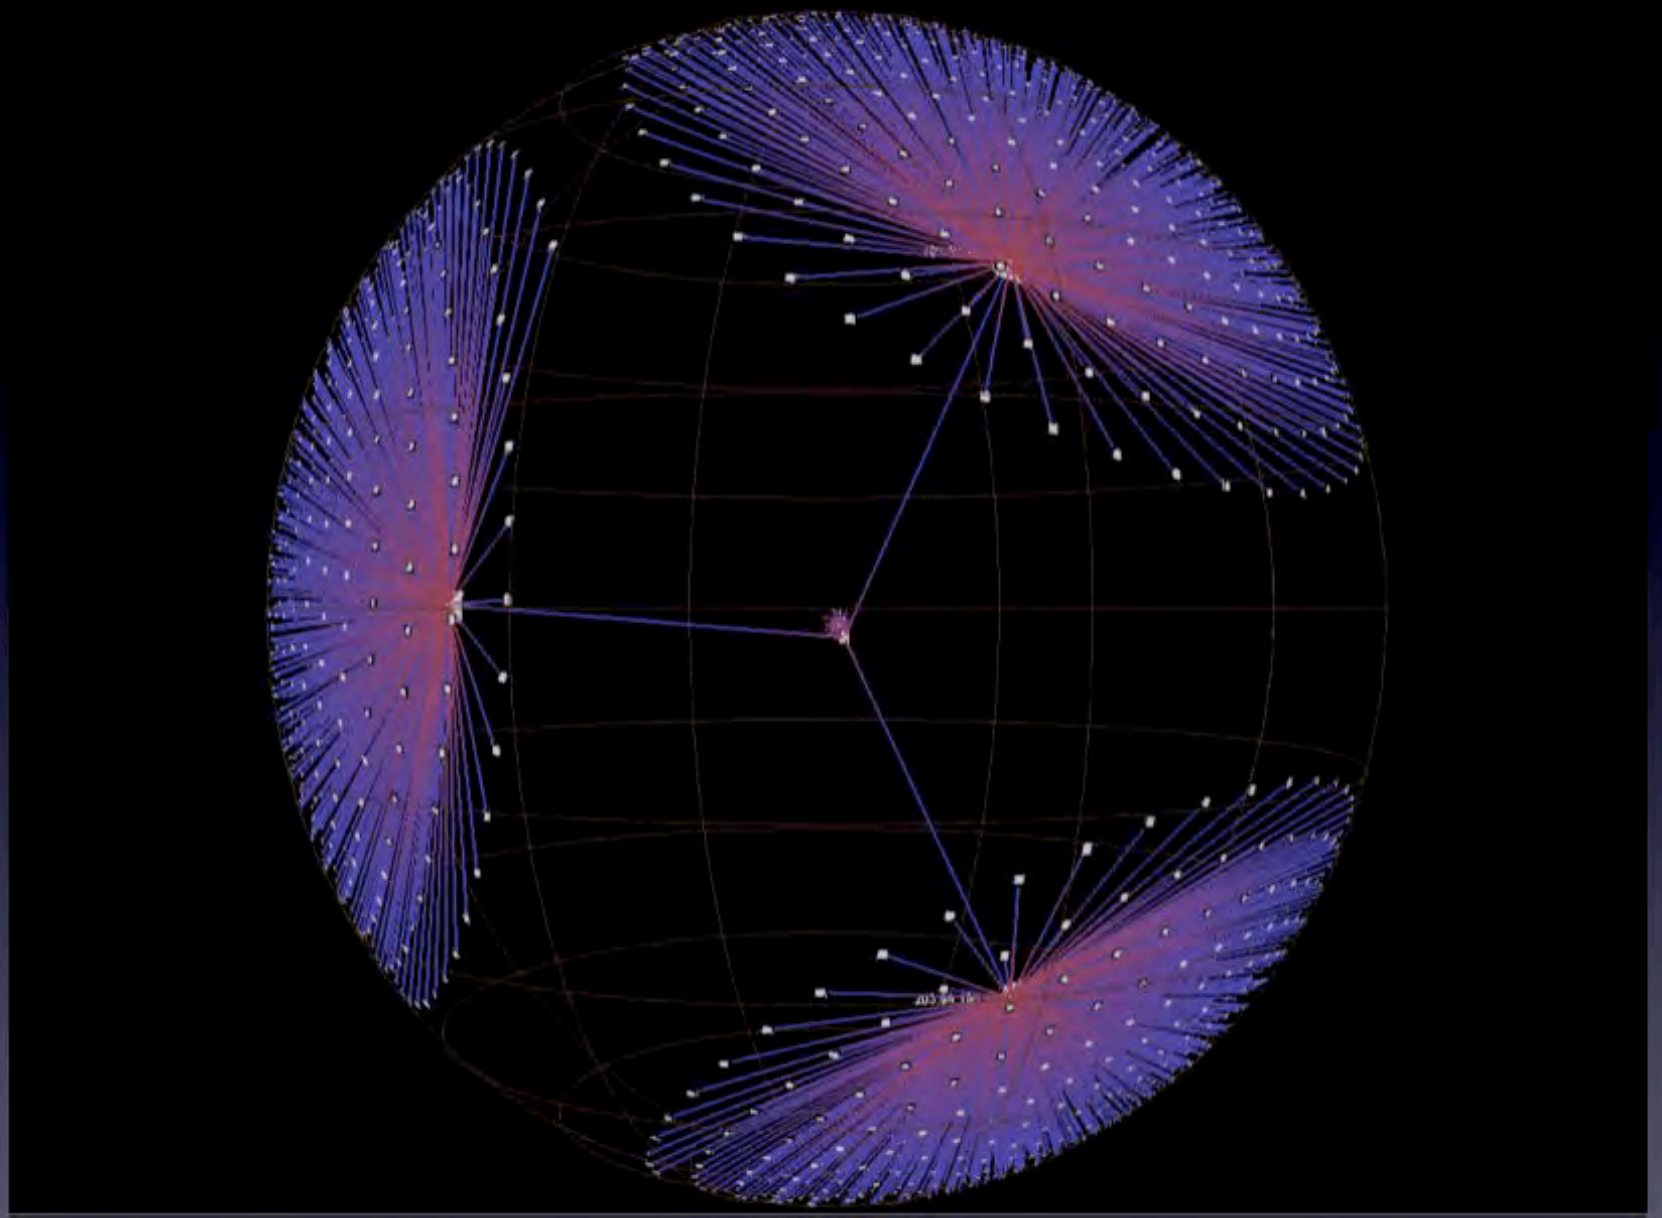




- colors
- graph cycles
- scaling
- text labels
- **graph size**
- search automation

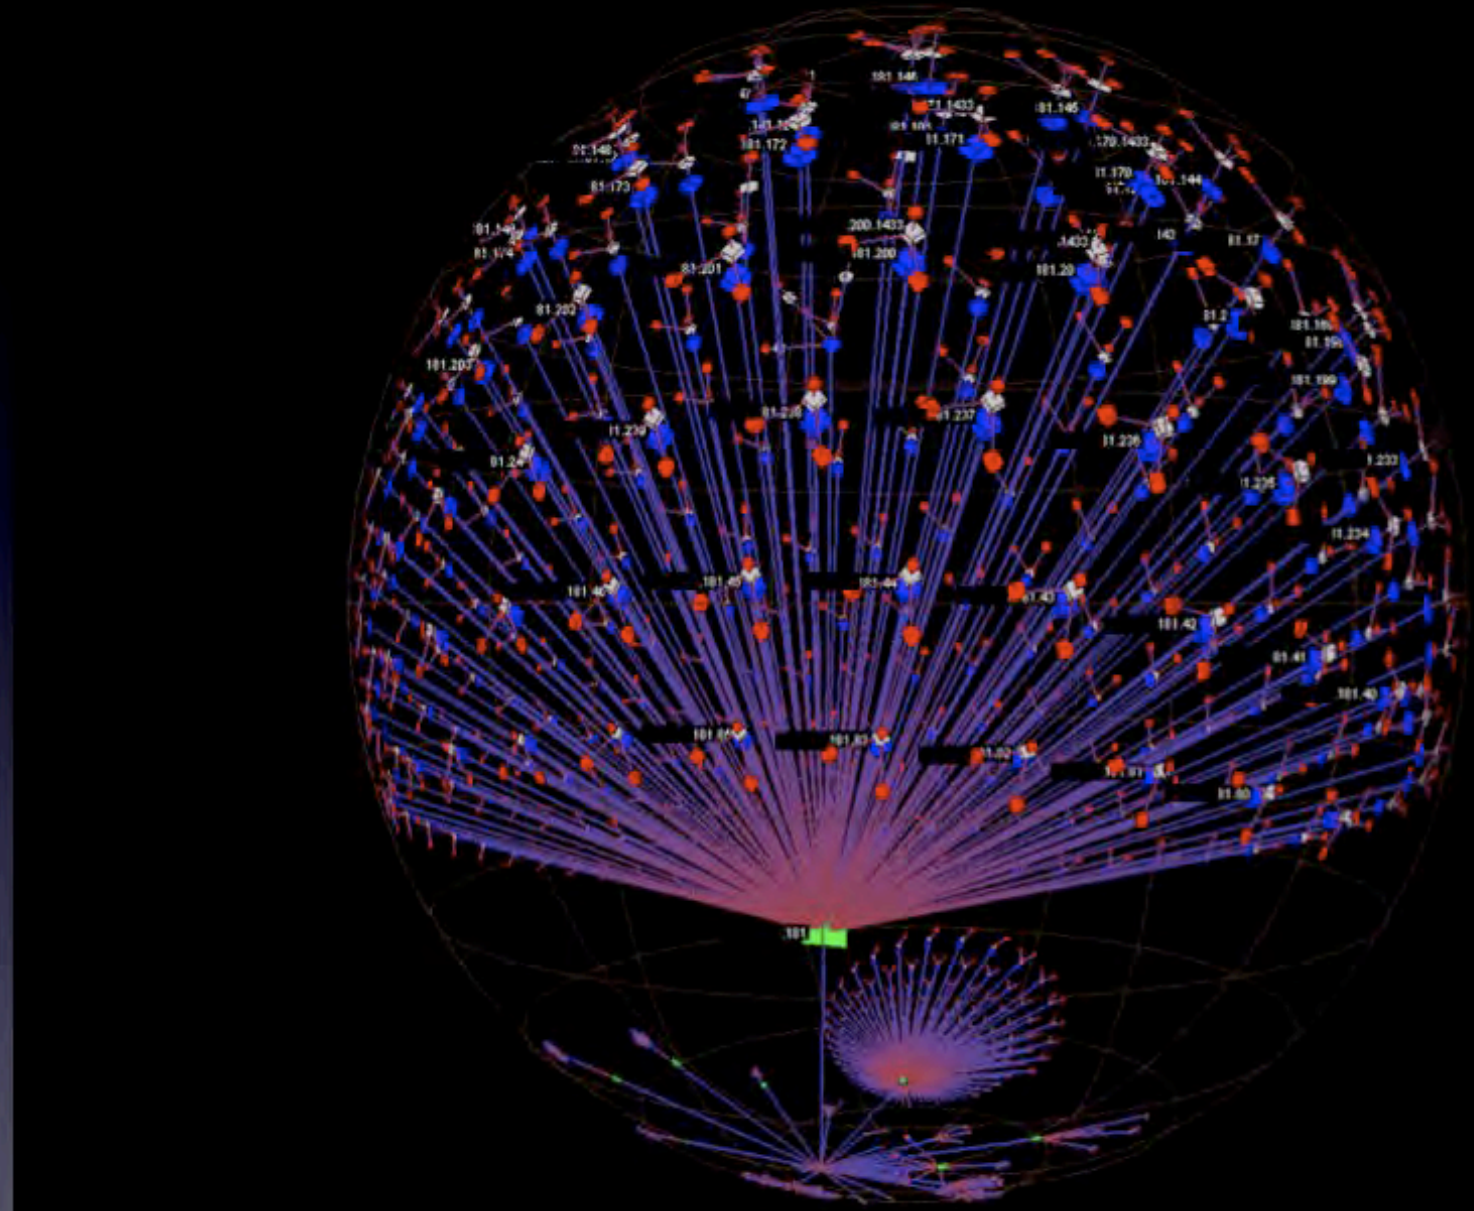


© 2007 Sunny Fugate



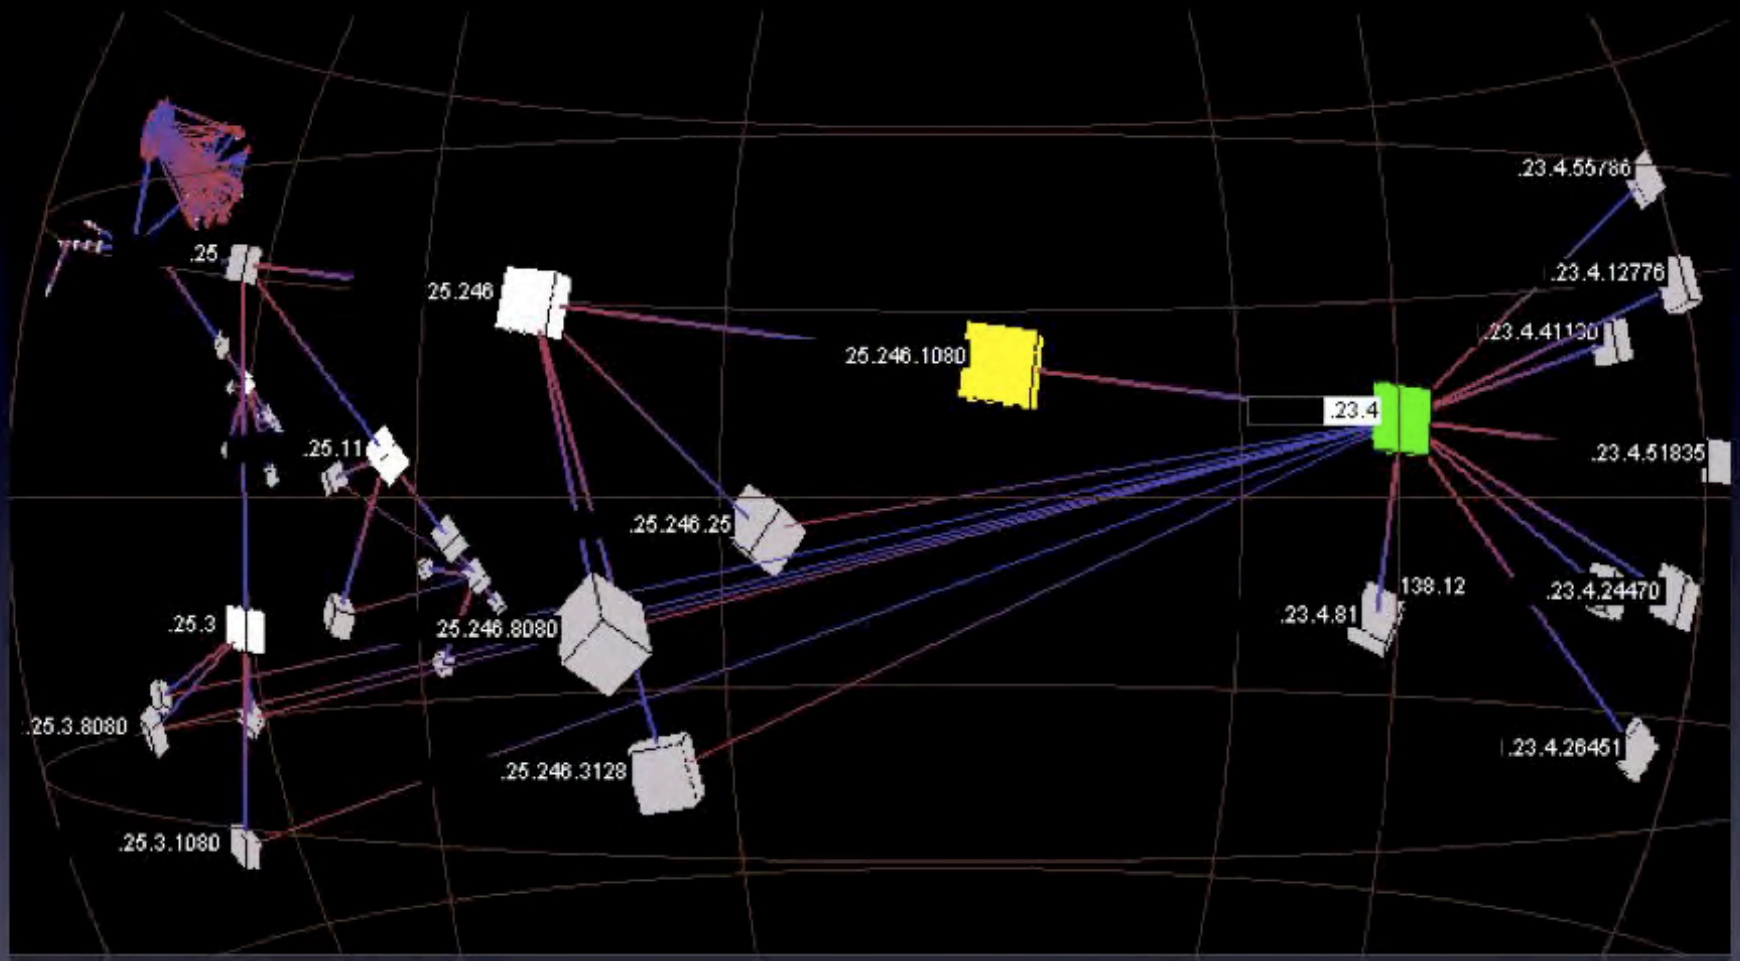
Symmetry in port access from 3 separate clients.

© 2007 Sunny Fugate



src/dst ports colored red/blue

© 2007 Sunny Fugate



Hierarchy showing client subnet and server ports

© 2007 Sunny Fugate

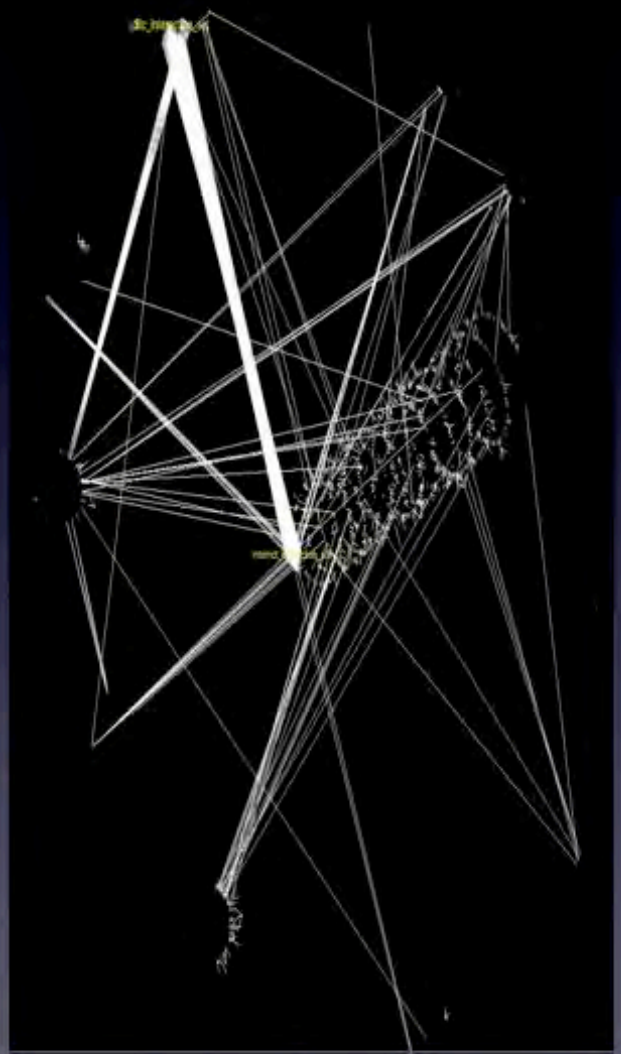
Shapes Vector

- Acquired by DARPA in 2002
- Developed by Australian DSTO

  (Defence Science Technology Organisation)
- JTF-GNO pilot program from 2003-2006

What is it?

- Intelligent Agents gather information and produce inferences
- Gathers information from multiple sources
  - pcap, flow, Snort, syslog, etc
- IAs performs automated data correlation & knowledge extraction
- Integrates visual and command-line analysis
- Integrated visualization makes use of human vision
- Supports visual analysis and decision-making



© 2007 Sunny Fugate

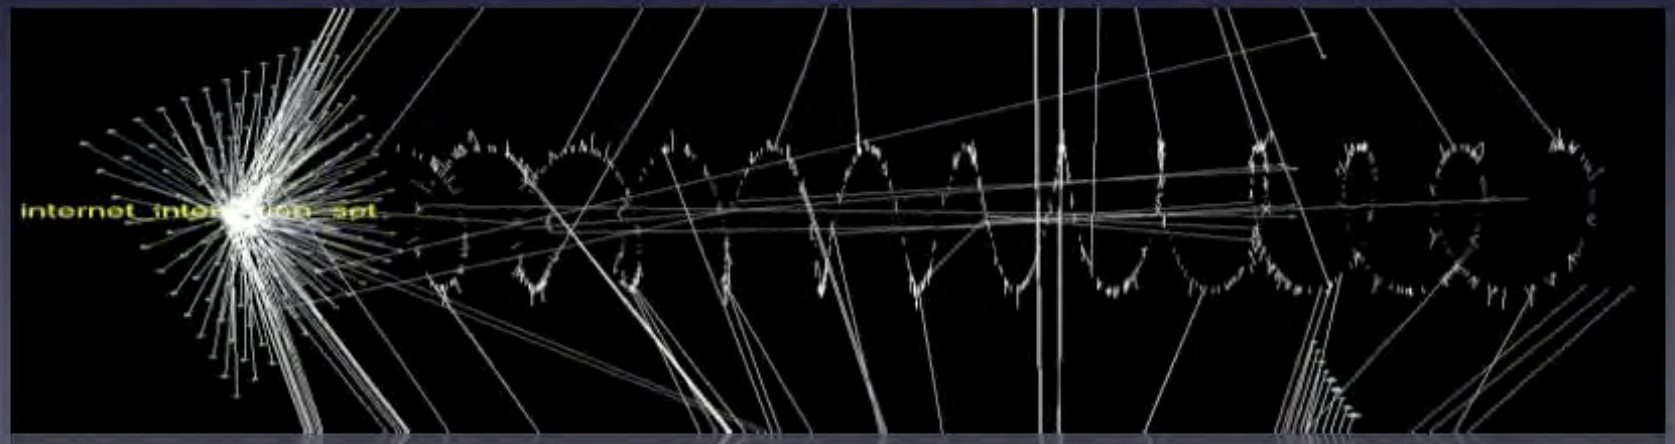
Shapes Vector

Contextual spatial, temporal, social, topological

Spatial physical geography or metaphor

Temporal sequences in time, correlated

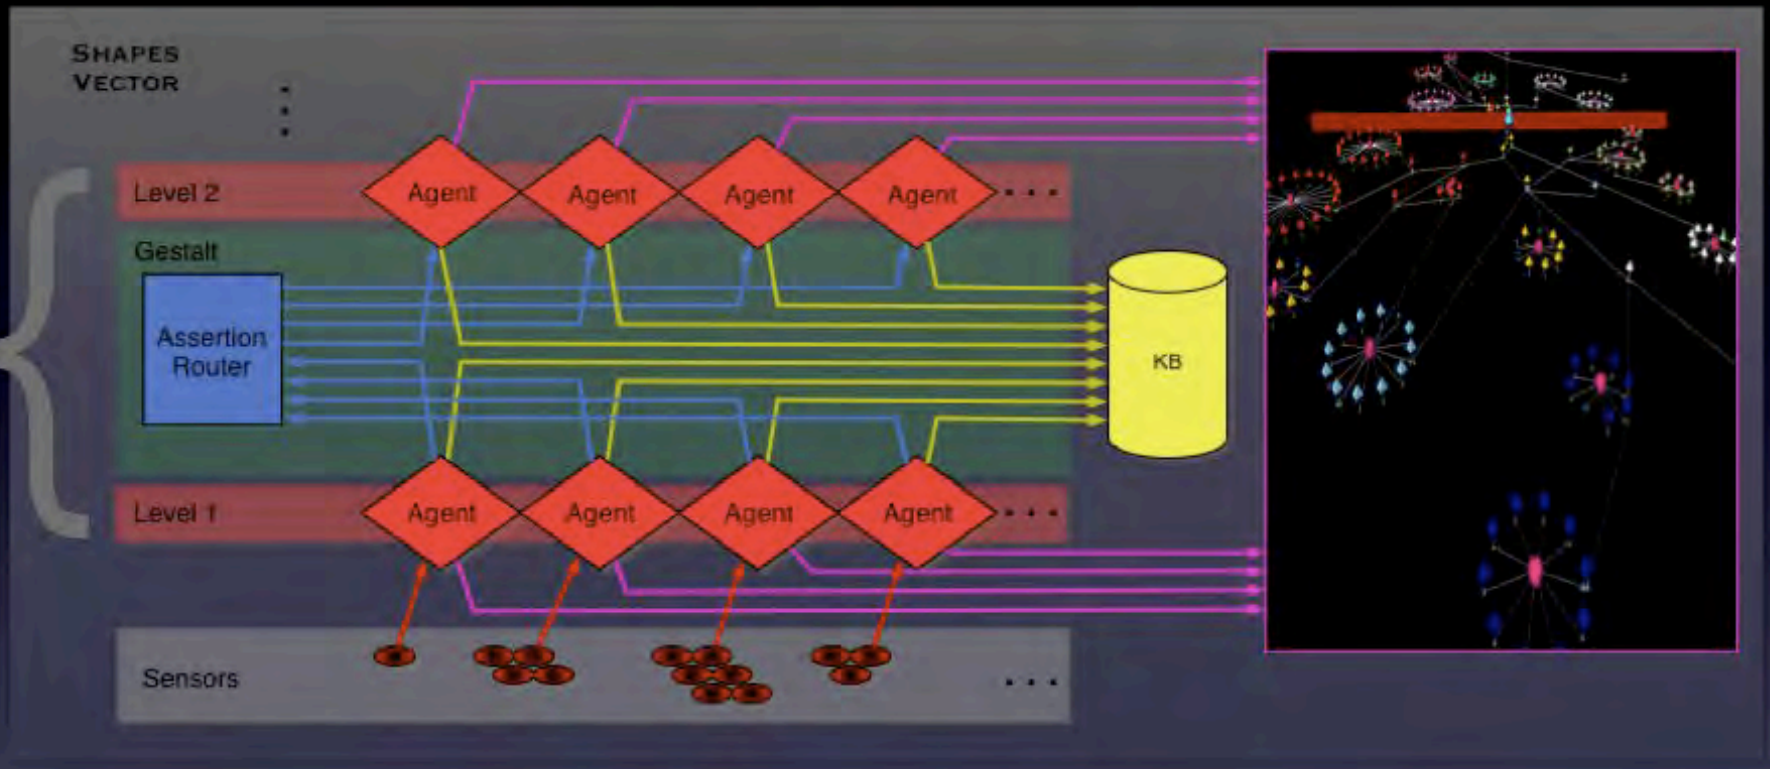
Visual use visual language to depict objects & events



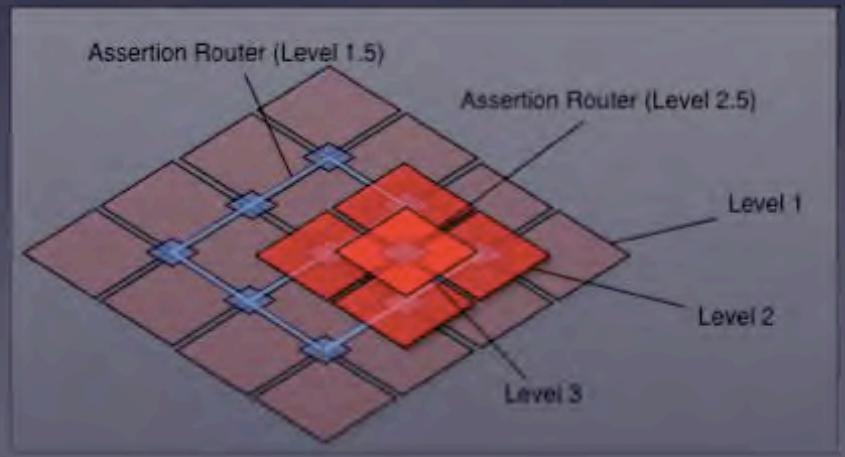
© 2007 Sunny Fugate

Architecture

SVKA



SVKA



- Agents can be written in many languages - must conform to the SV ontology and knowledge architecture (SVKA) specification
- Sensors can be built to wrap many information sources - must produce SV ontology
- SV ontology is a knowledge description language for network defense

© 2007 Sunny Fugate

Shapes Vector - Visual Language

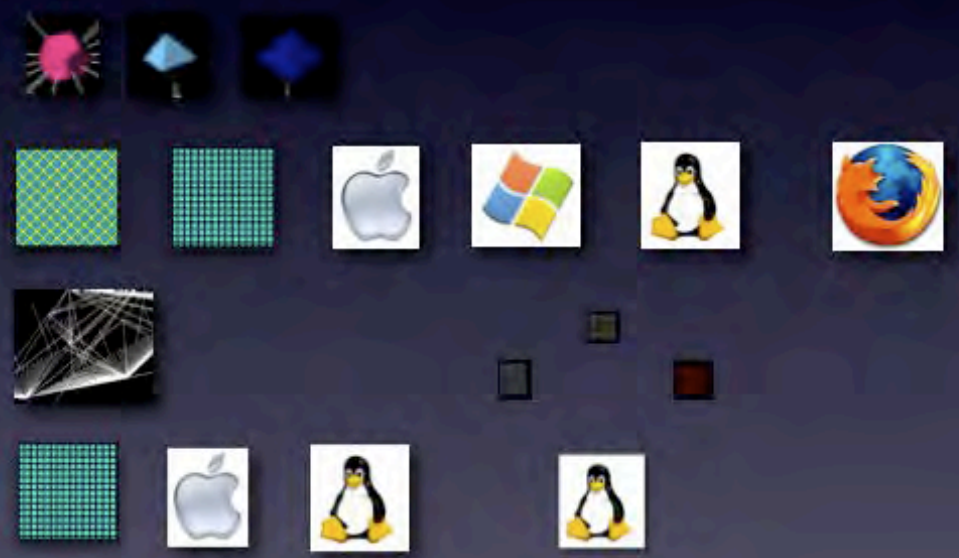
- Easily defined visual mappings
- No applied theory of visual language

shape/color/scale

texture/icon

connection / topology

movement



packet events, information exchange, attribute changes, attribute values, host id, software, processes, machine purpose, network topology, social topology, intrusion events, event type, event priority, client vs server, routing, ...

© 2007 Sunny Fugate

Shapes Vector - Visual Language

- Easily defined visual mappings
- No applied theory of visual language

shape/color/scale



texture/icon



connection / topology



movement



packet events, information exchange, attribute changes, attribute values, host id, software, processes, machine purpose, network topology, social topology, intrusion events, event type, event priority, client vs server, routing, ...

© 2007 Sunny Fugate

Shapes Vector - Visual Language

- Easily defined visual mappings
- No applied theory of visual language

shape/color/scale

texture/icon



connection / topology



movement



packet events, information exchange, attribute changes, attribute values, host id, software, processes, machine purpose, network topology, social topology, intrusion events, event type, event priority, client vs server, routing, ...

© 2007 Sunny Fugate

Shapes Vector - Visual Language

- Easily defined visual mappings
- No applied theory of visual language

shape/color/scale

texture/icon



connection / topology



movement



packet events, information exchange, attribute changes, attribute values, host id, software, processes, machine purpose, network topology, social topology, intrusion events, event type, event priority, client vs server, routing, ...

© 2007 Sunny Fugate

Shapes Vector - Visual Language

- Easily defined visual mappings
- No applied theory of visual language

shape/color/scale

texture/icon



connection / topology



movement



packet events, information exchange, attribute changes, attribute values, host id, software, processes, machine purpose, network topology, social topology, intrusion events, event type, event priority, client vs server, routing, ...

© 2007 Sunny Fugate

Shapes Vector - Visual Language

- Easily defined visual mappings
- No applied theory of visual language

shape/color/scale

texture/icon



connection / topology



movement



packet events, information exchange, attribute changes, attribute values, host id, software, processes, machine purpose, network topology, social topology, intrusion events, event type, event priority, client vs server, routing, ...

© 2007 Sunny Fugate

Shapes Vector - Visual Language

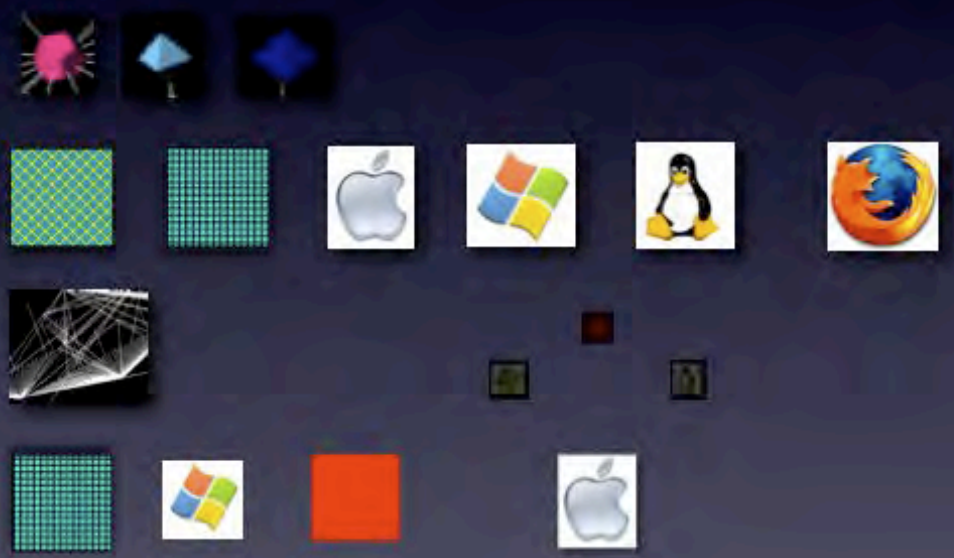
- Easily defined visual mappings
- No applied theory of visual language

shape/color/scale

texture/icon

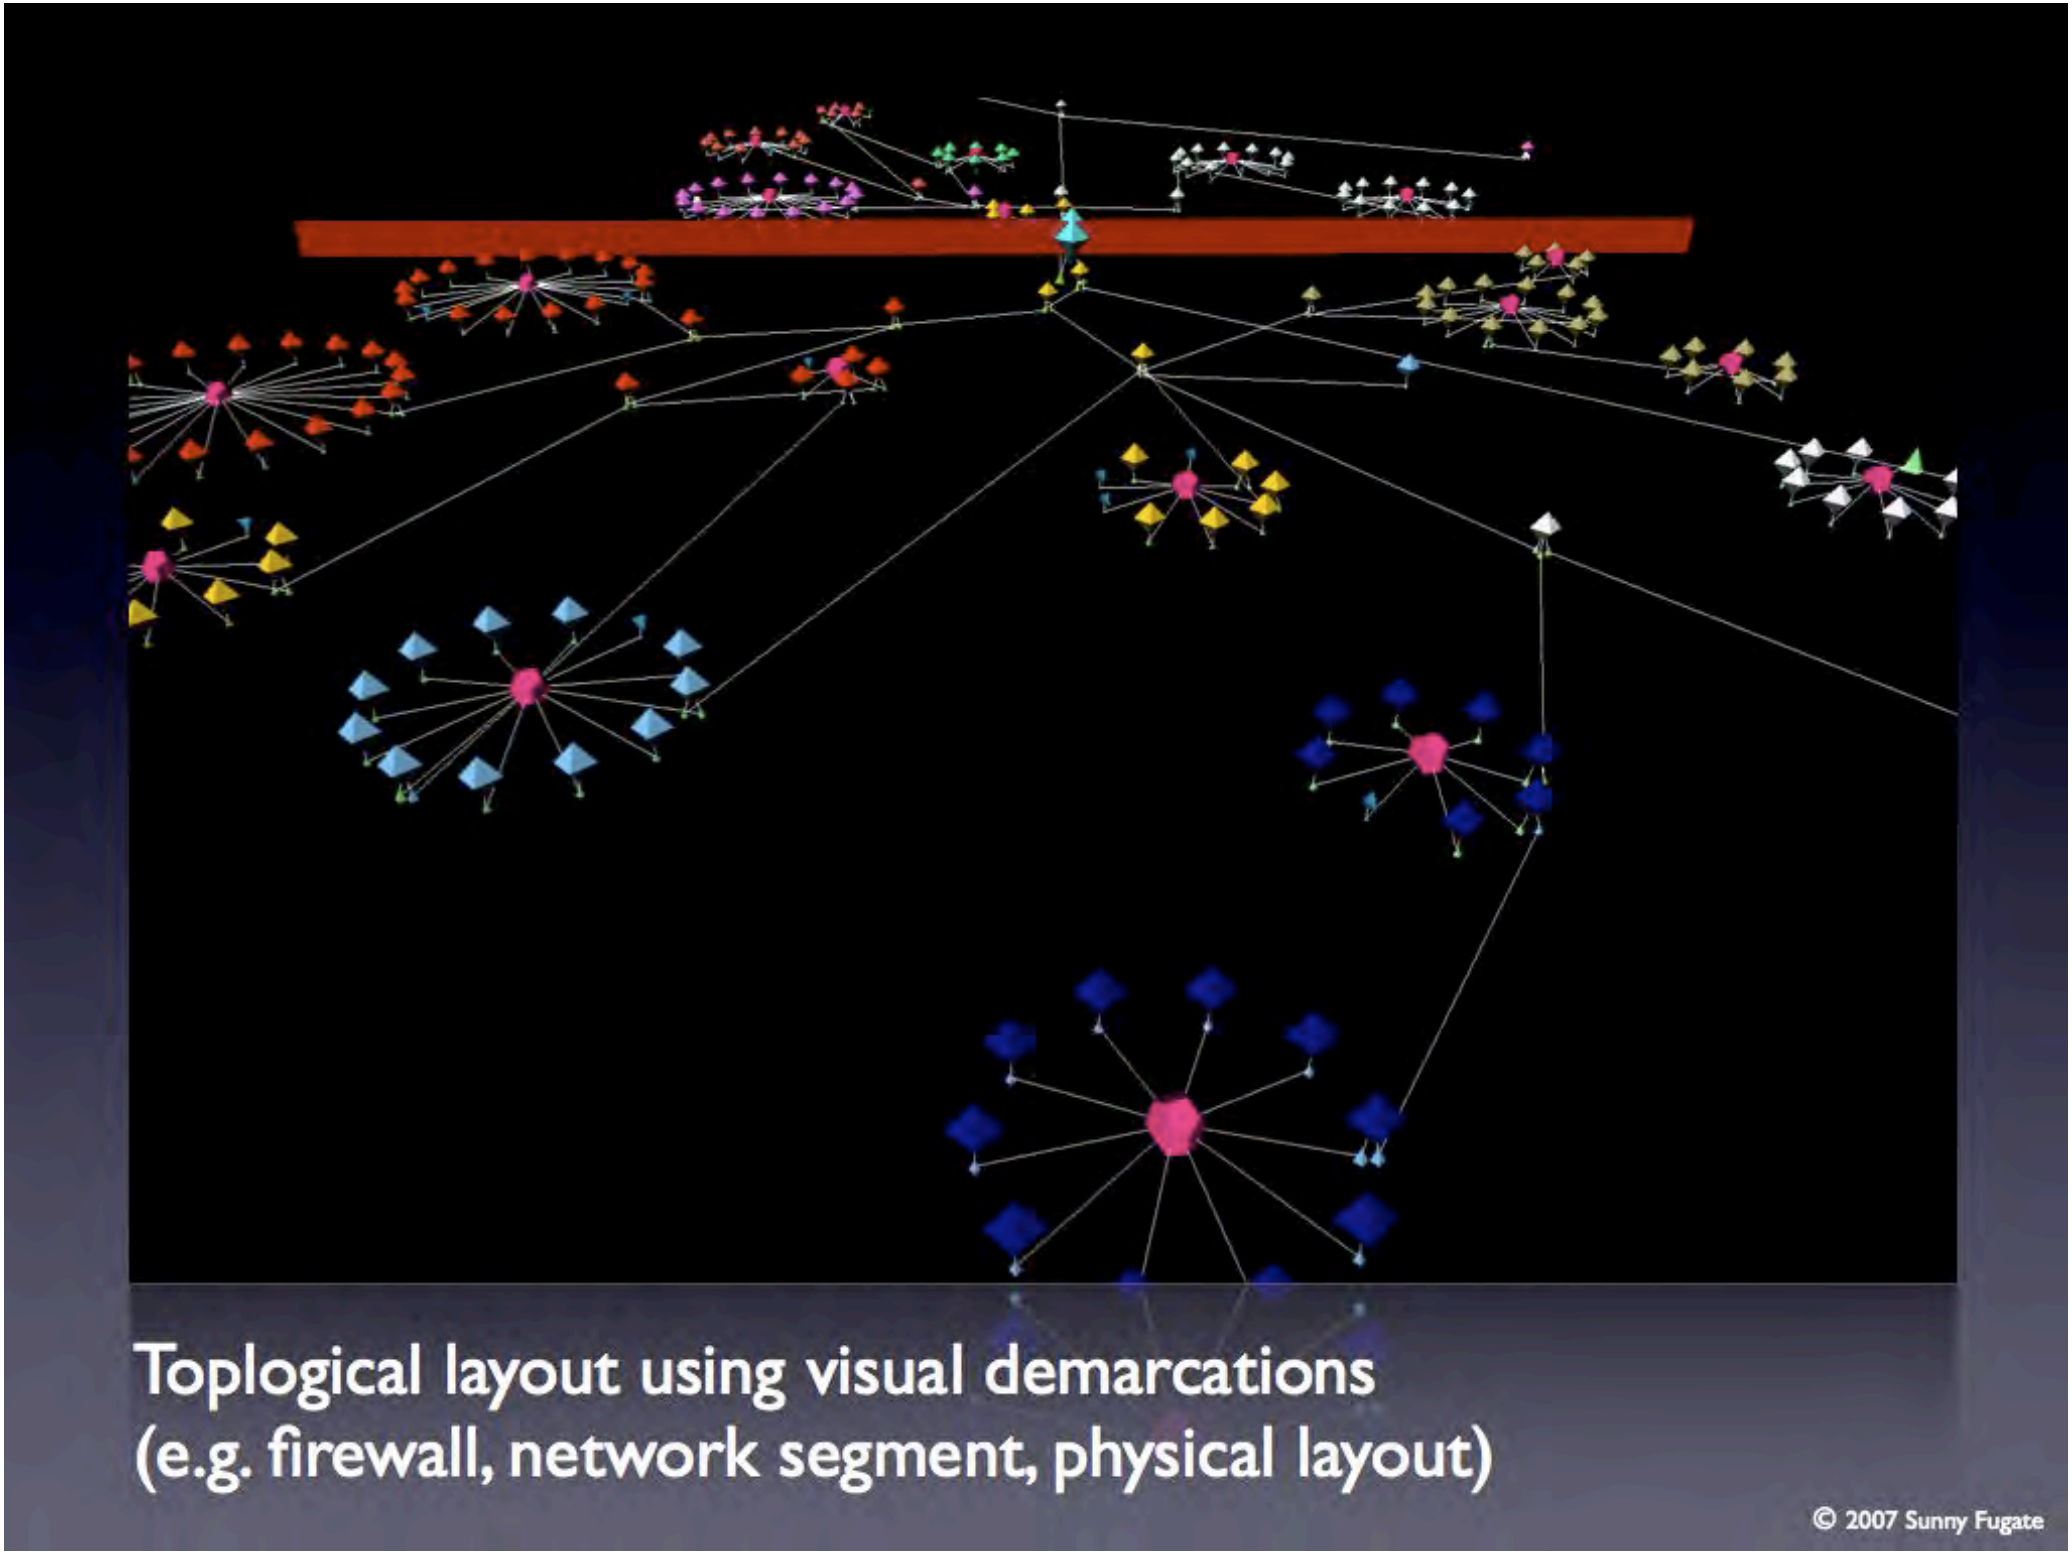
connection / topology

movement



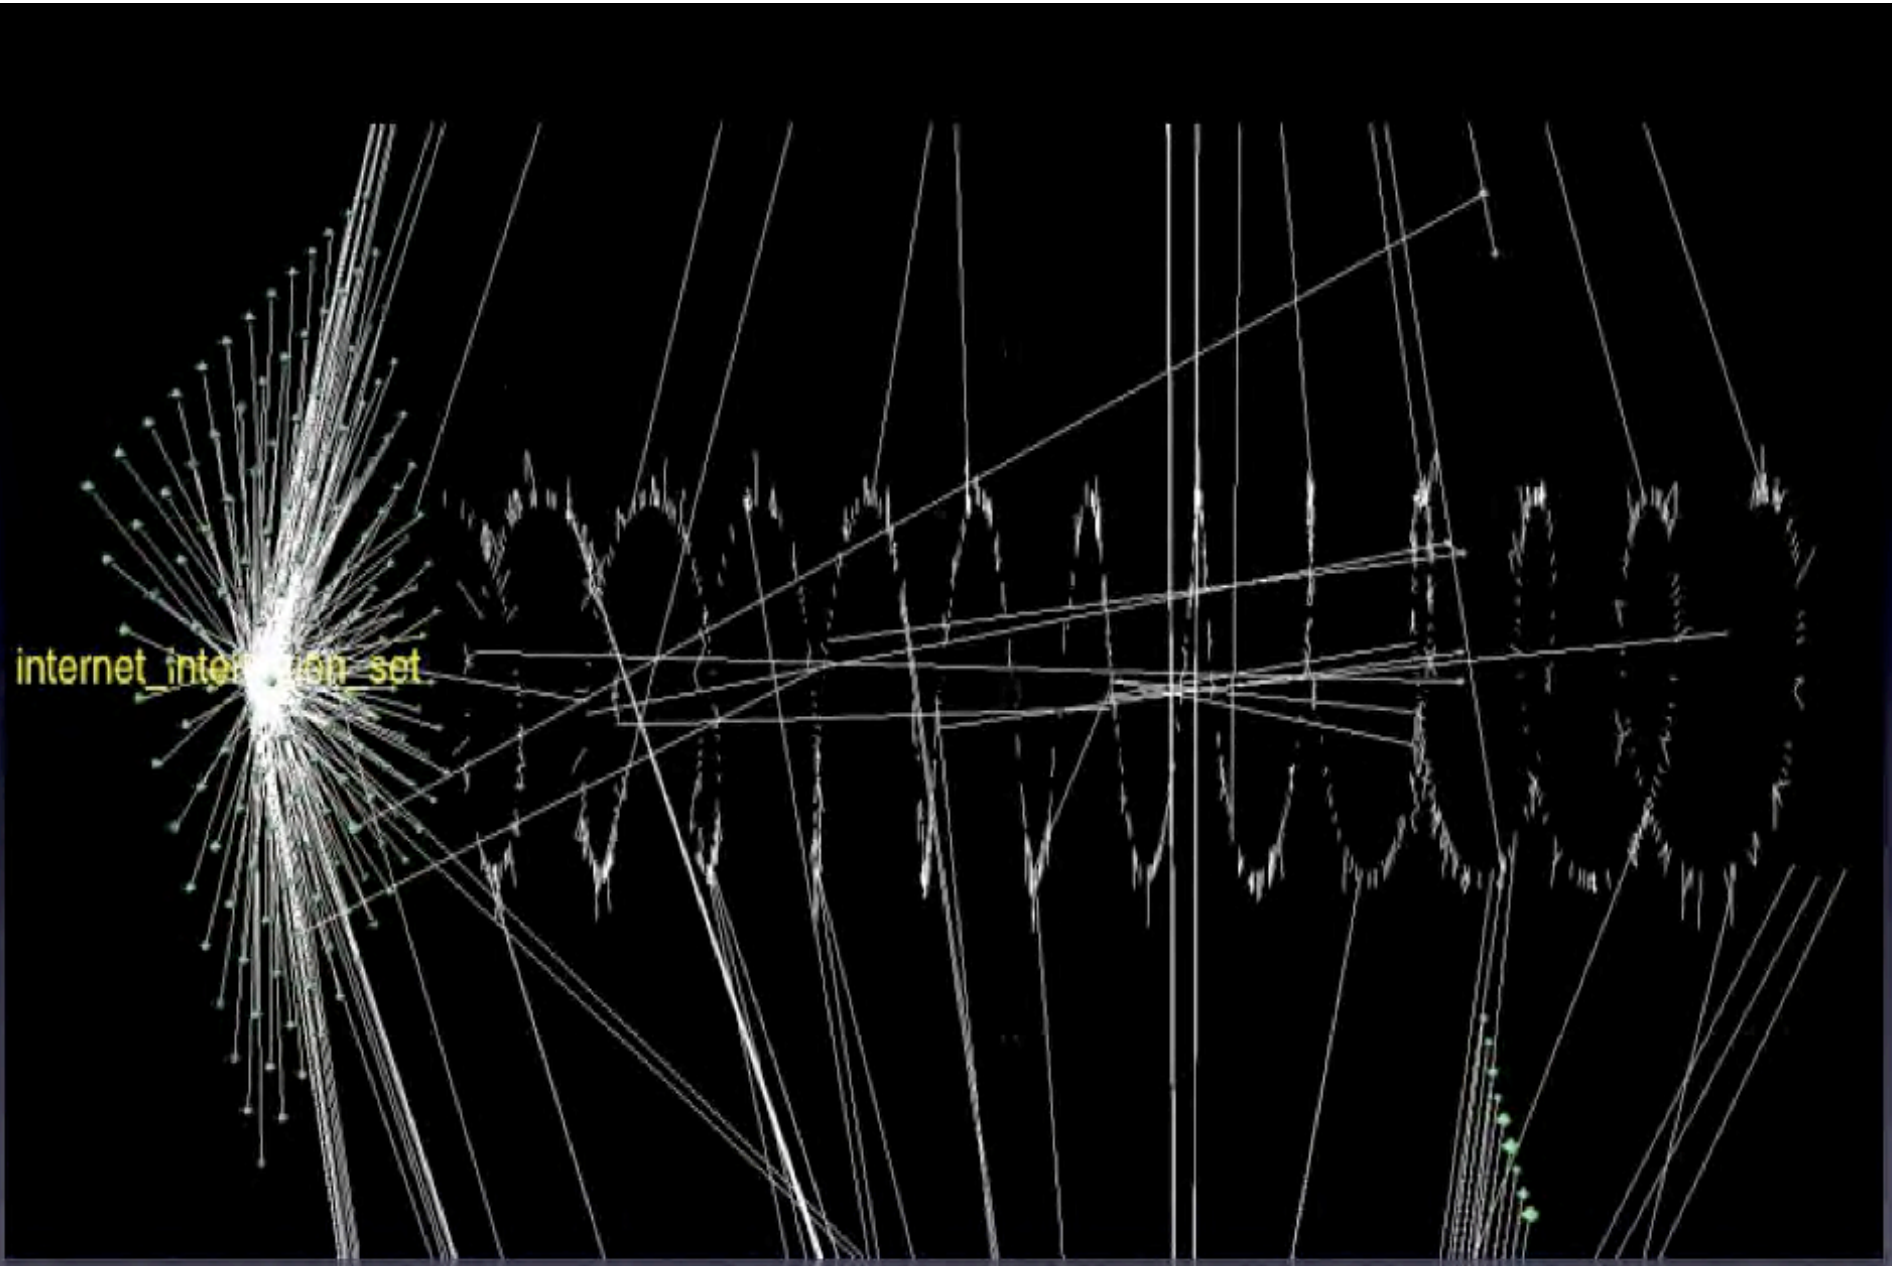
packet events, information exchange, attribute changes, attribute values, host id, software, processes, machine purpose, network topology, social topology, intrusion events, event type, event priority, client vs server, routing, ...

© 2007 Sunny Fugate



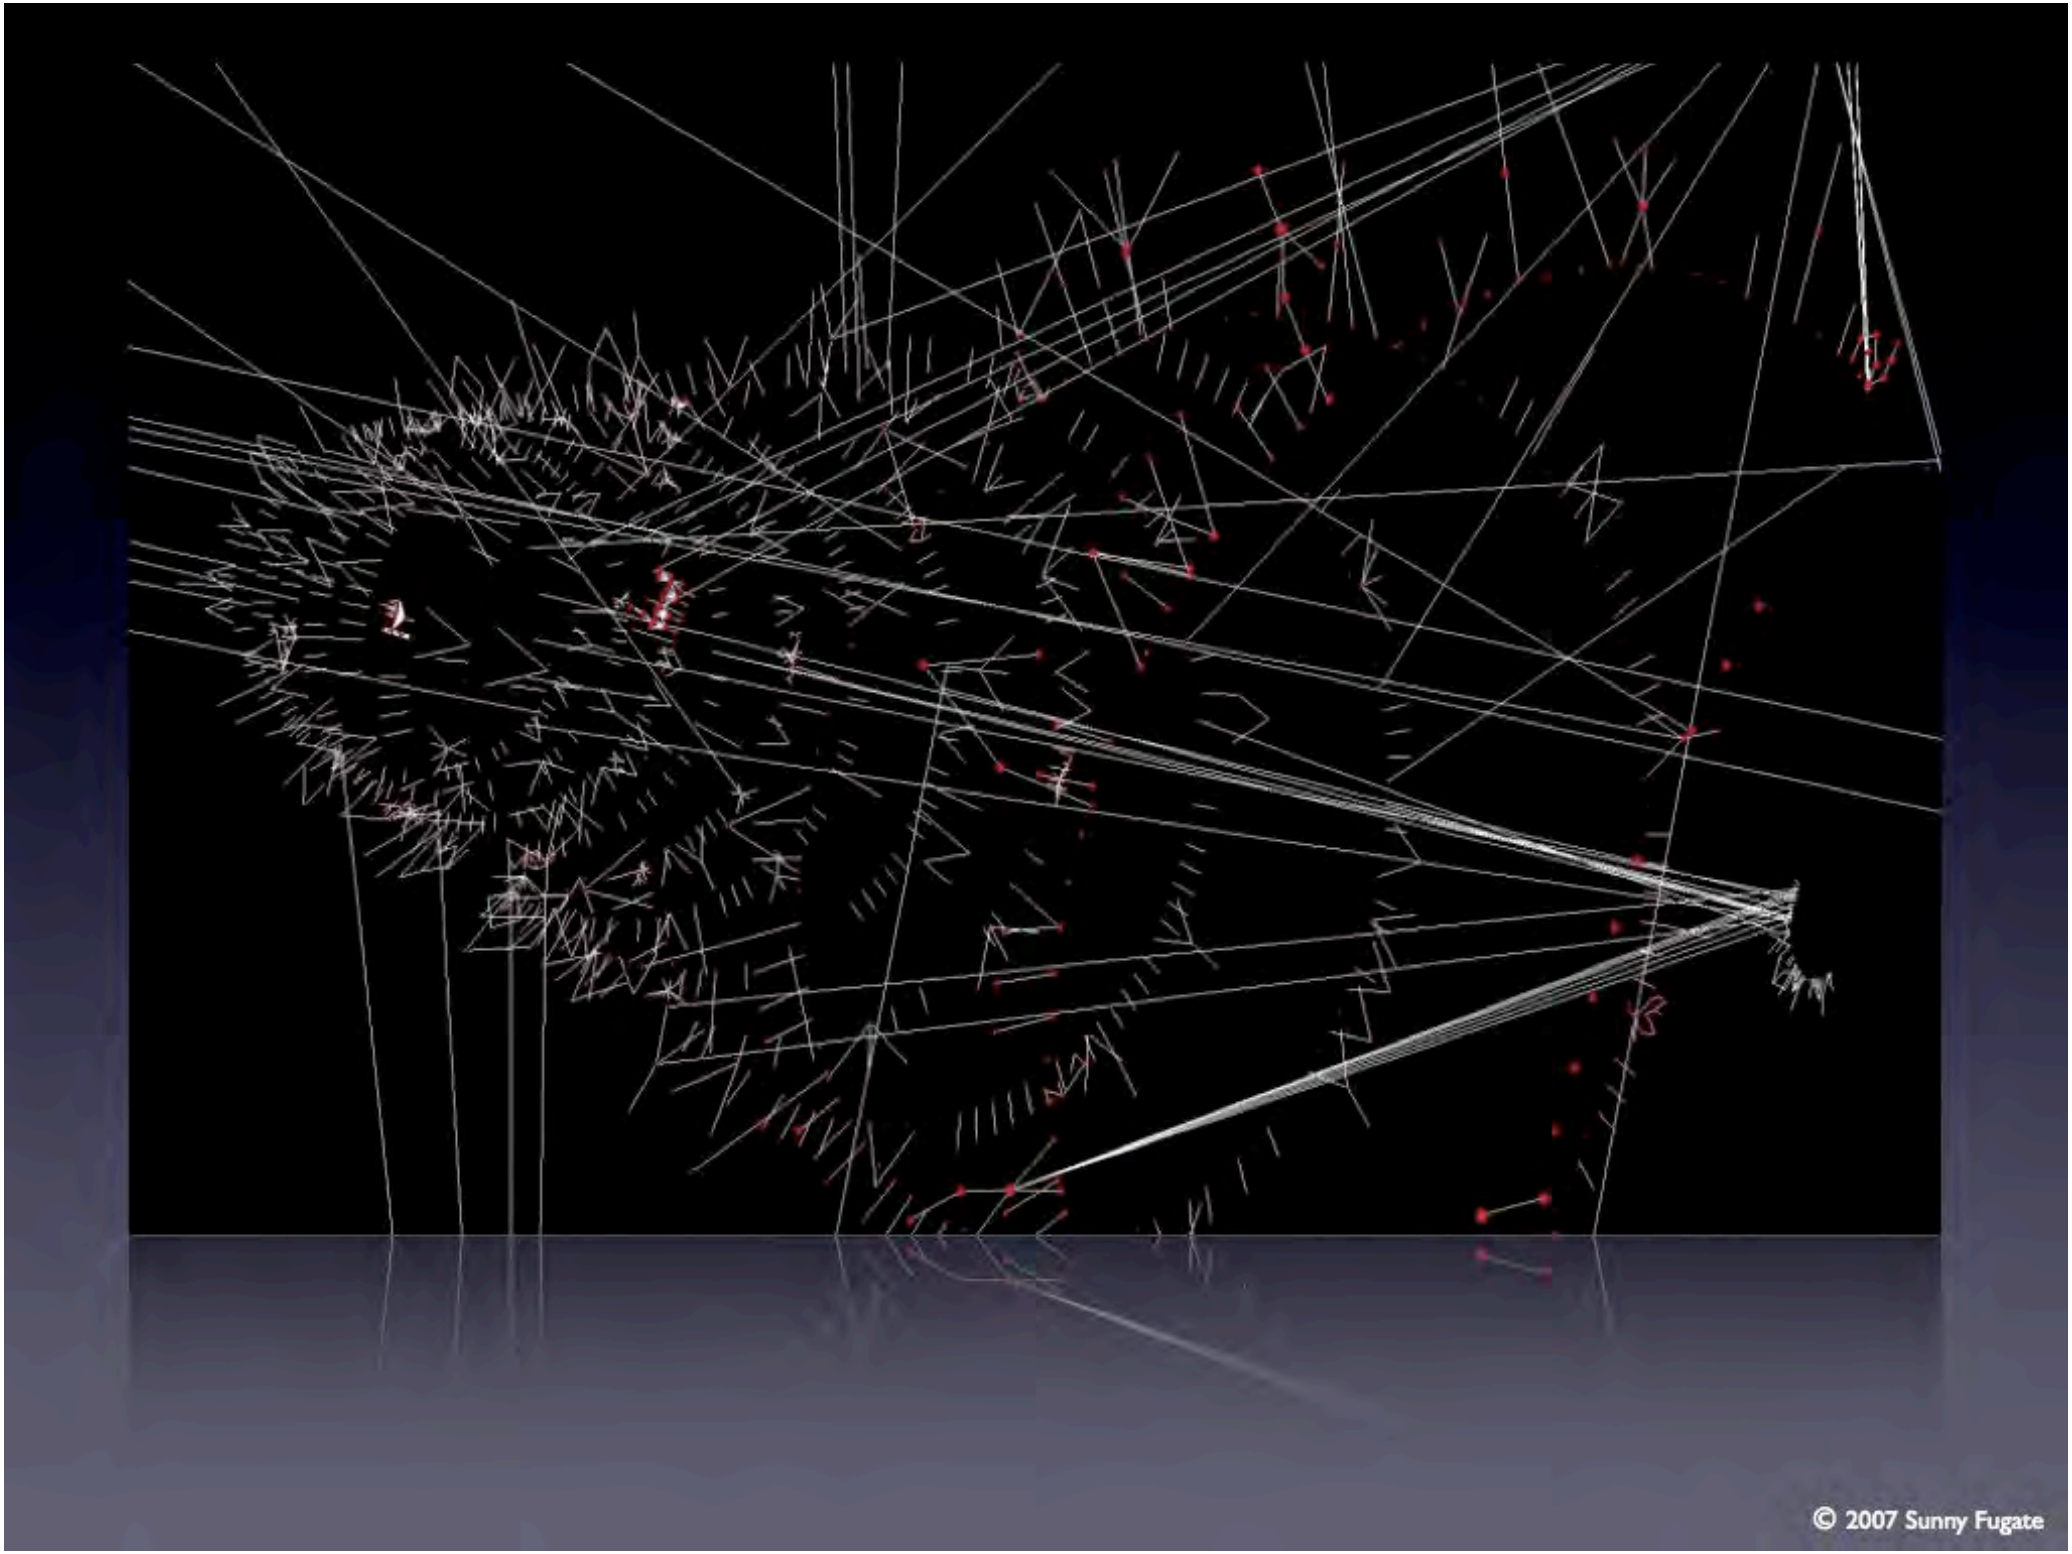
Toplogical layout using visual demarcations
(e.g. firewall, network segment, physical layout)

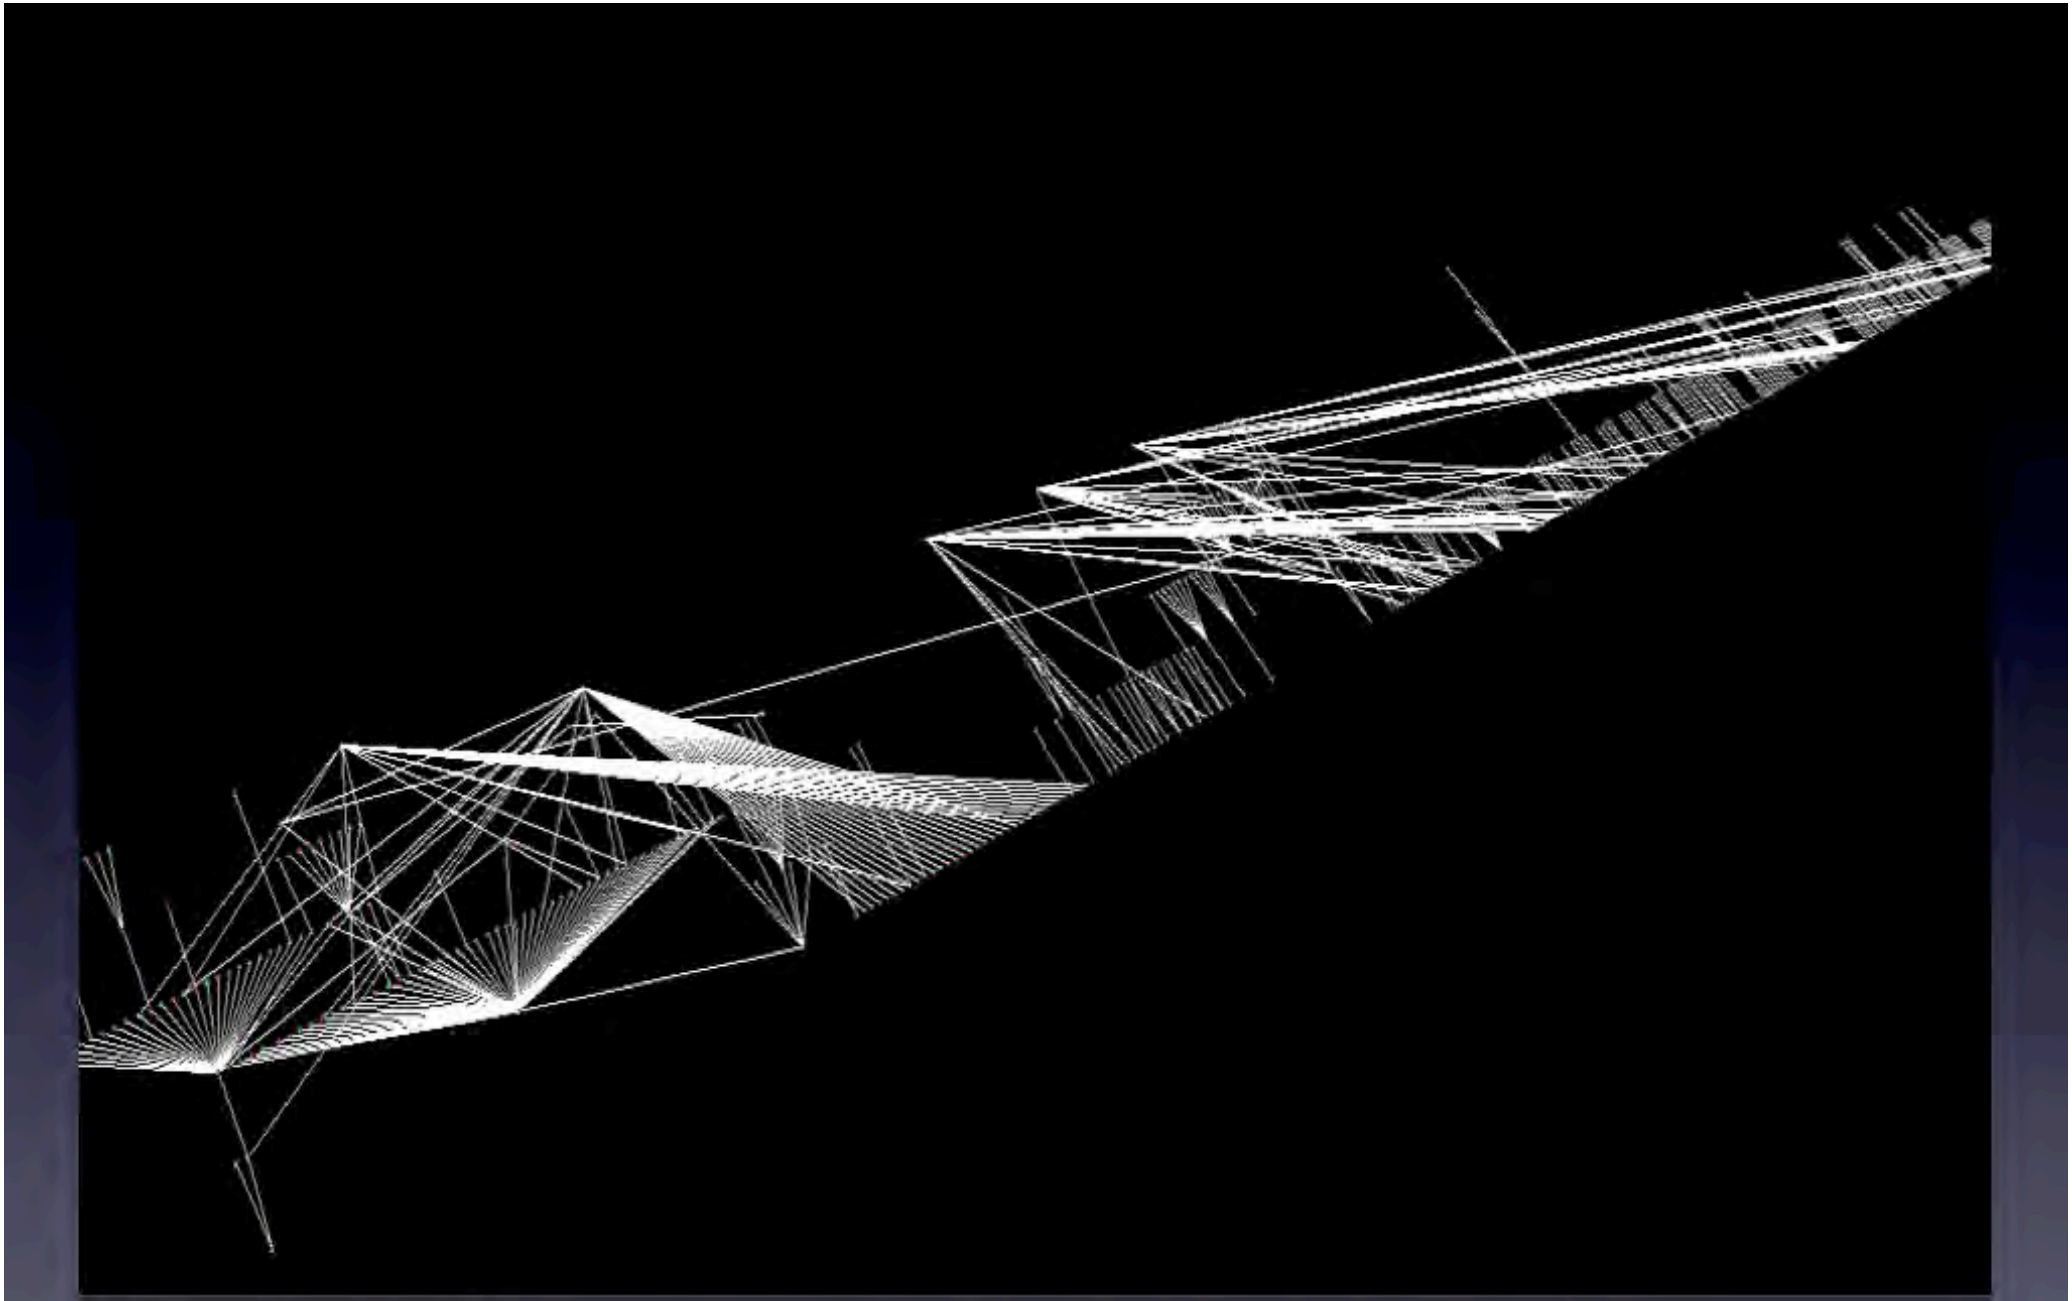
© 2007 Sunny Fugate



Automated layout to arrange hundreds of sub-graphs in a non-overlapping manner.

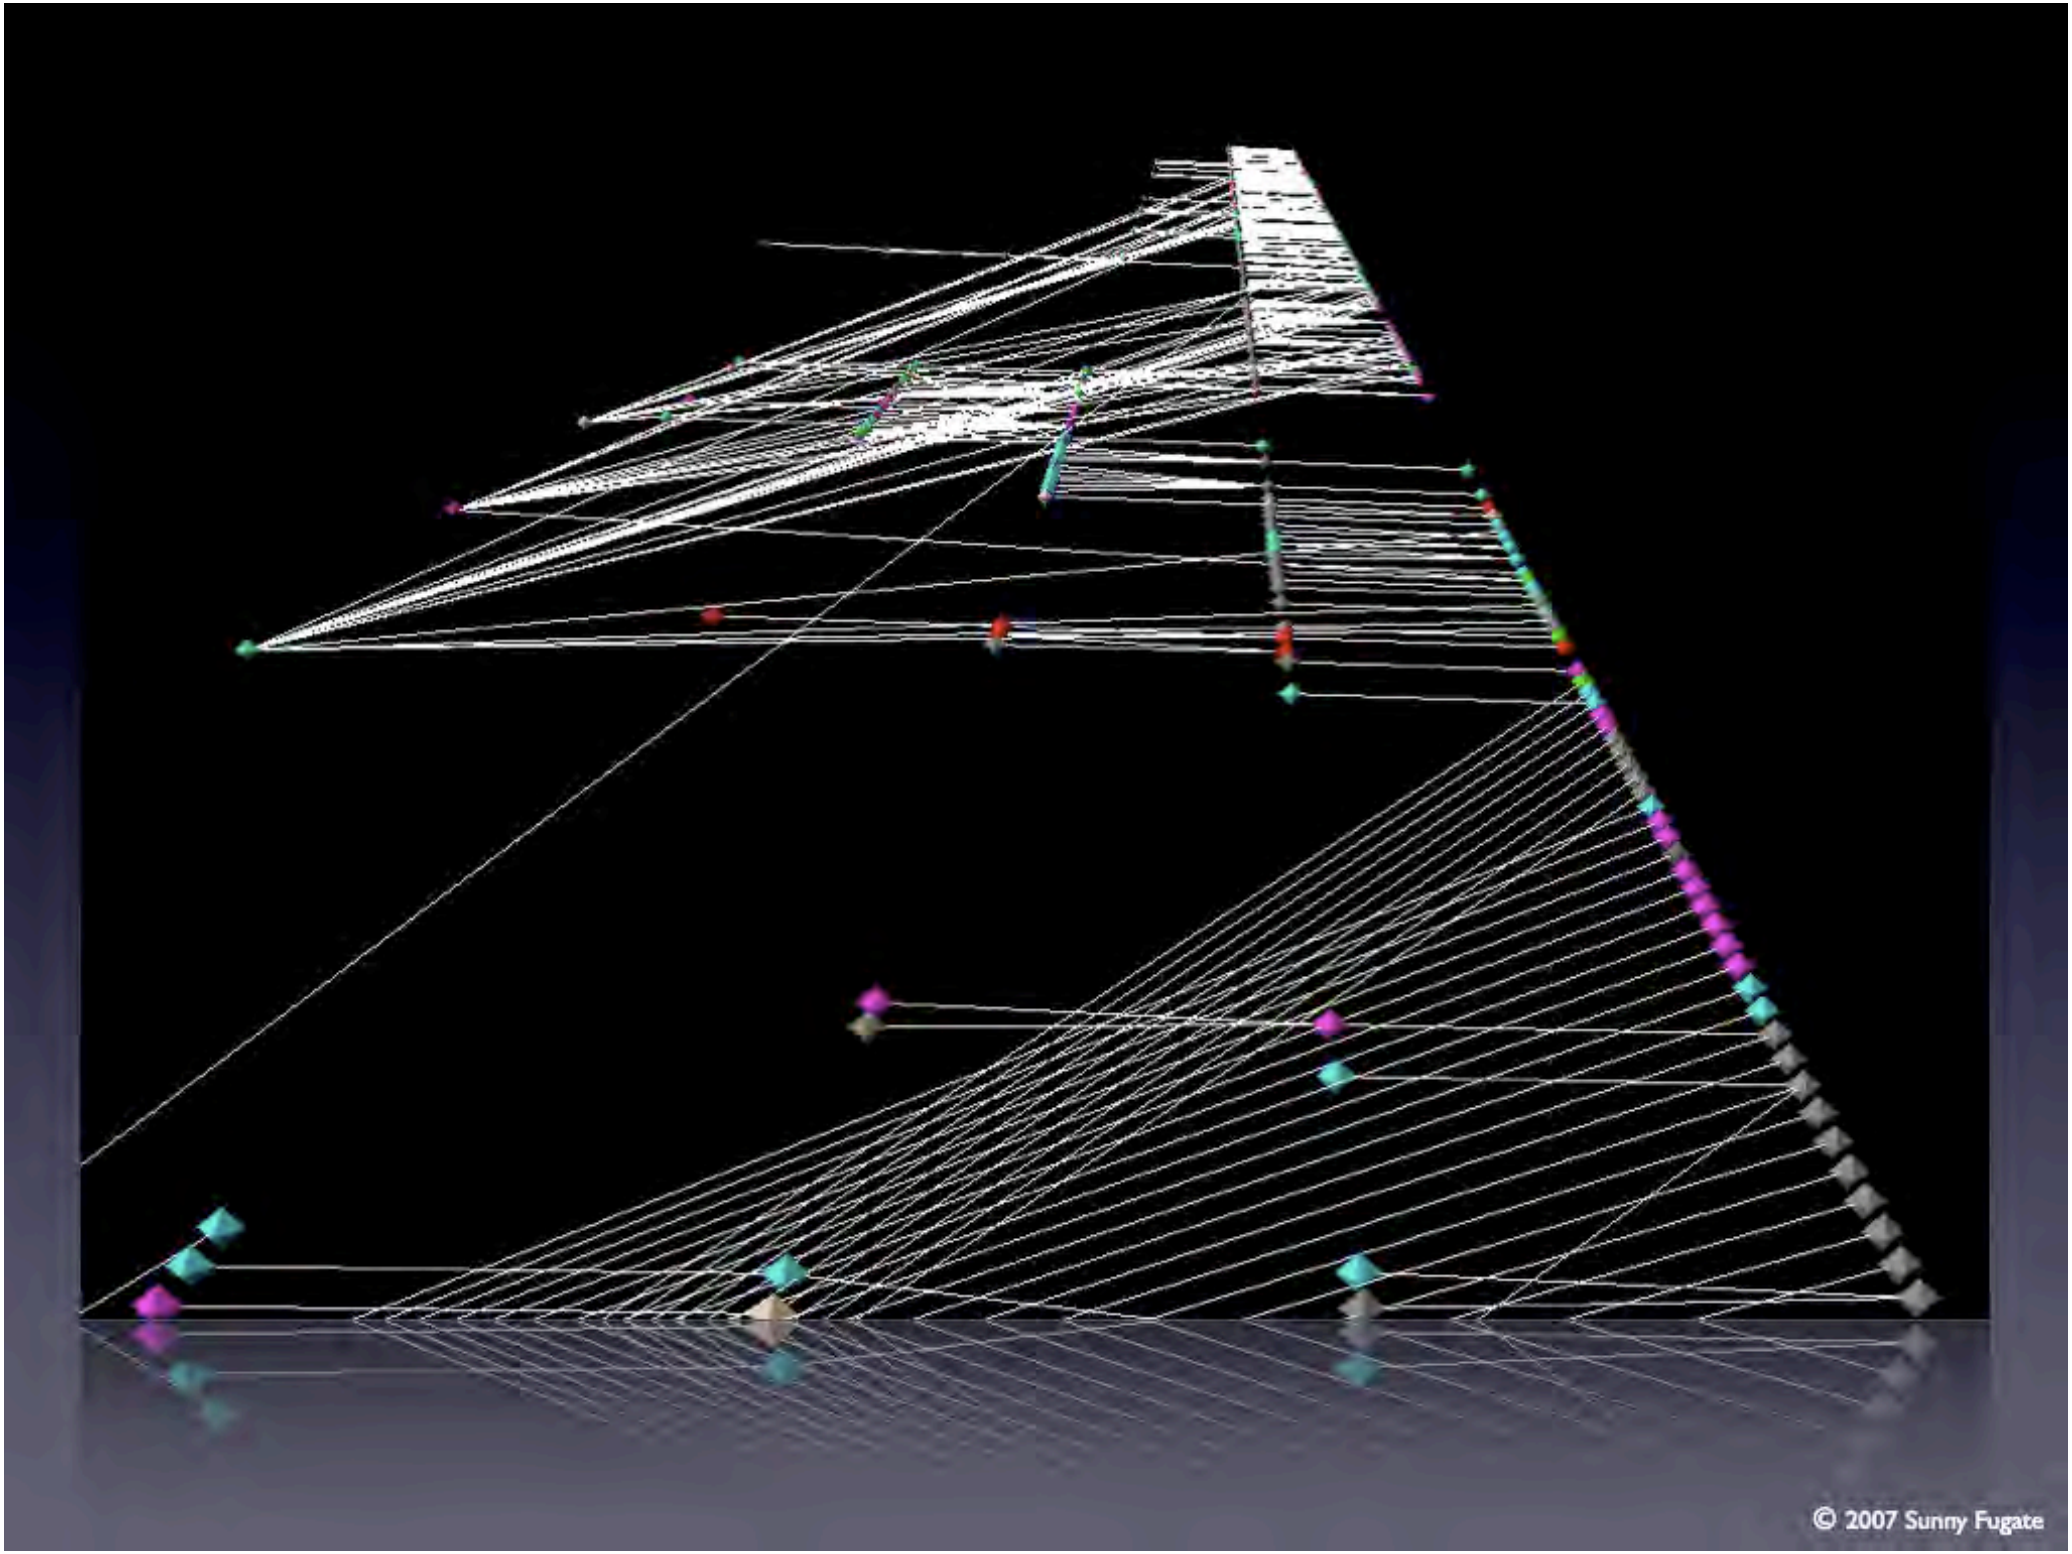
© 2007 Sunny Fugate

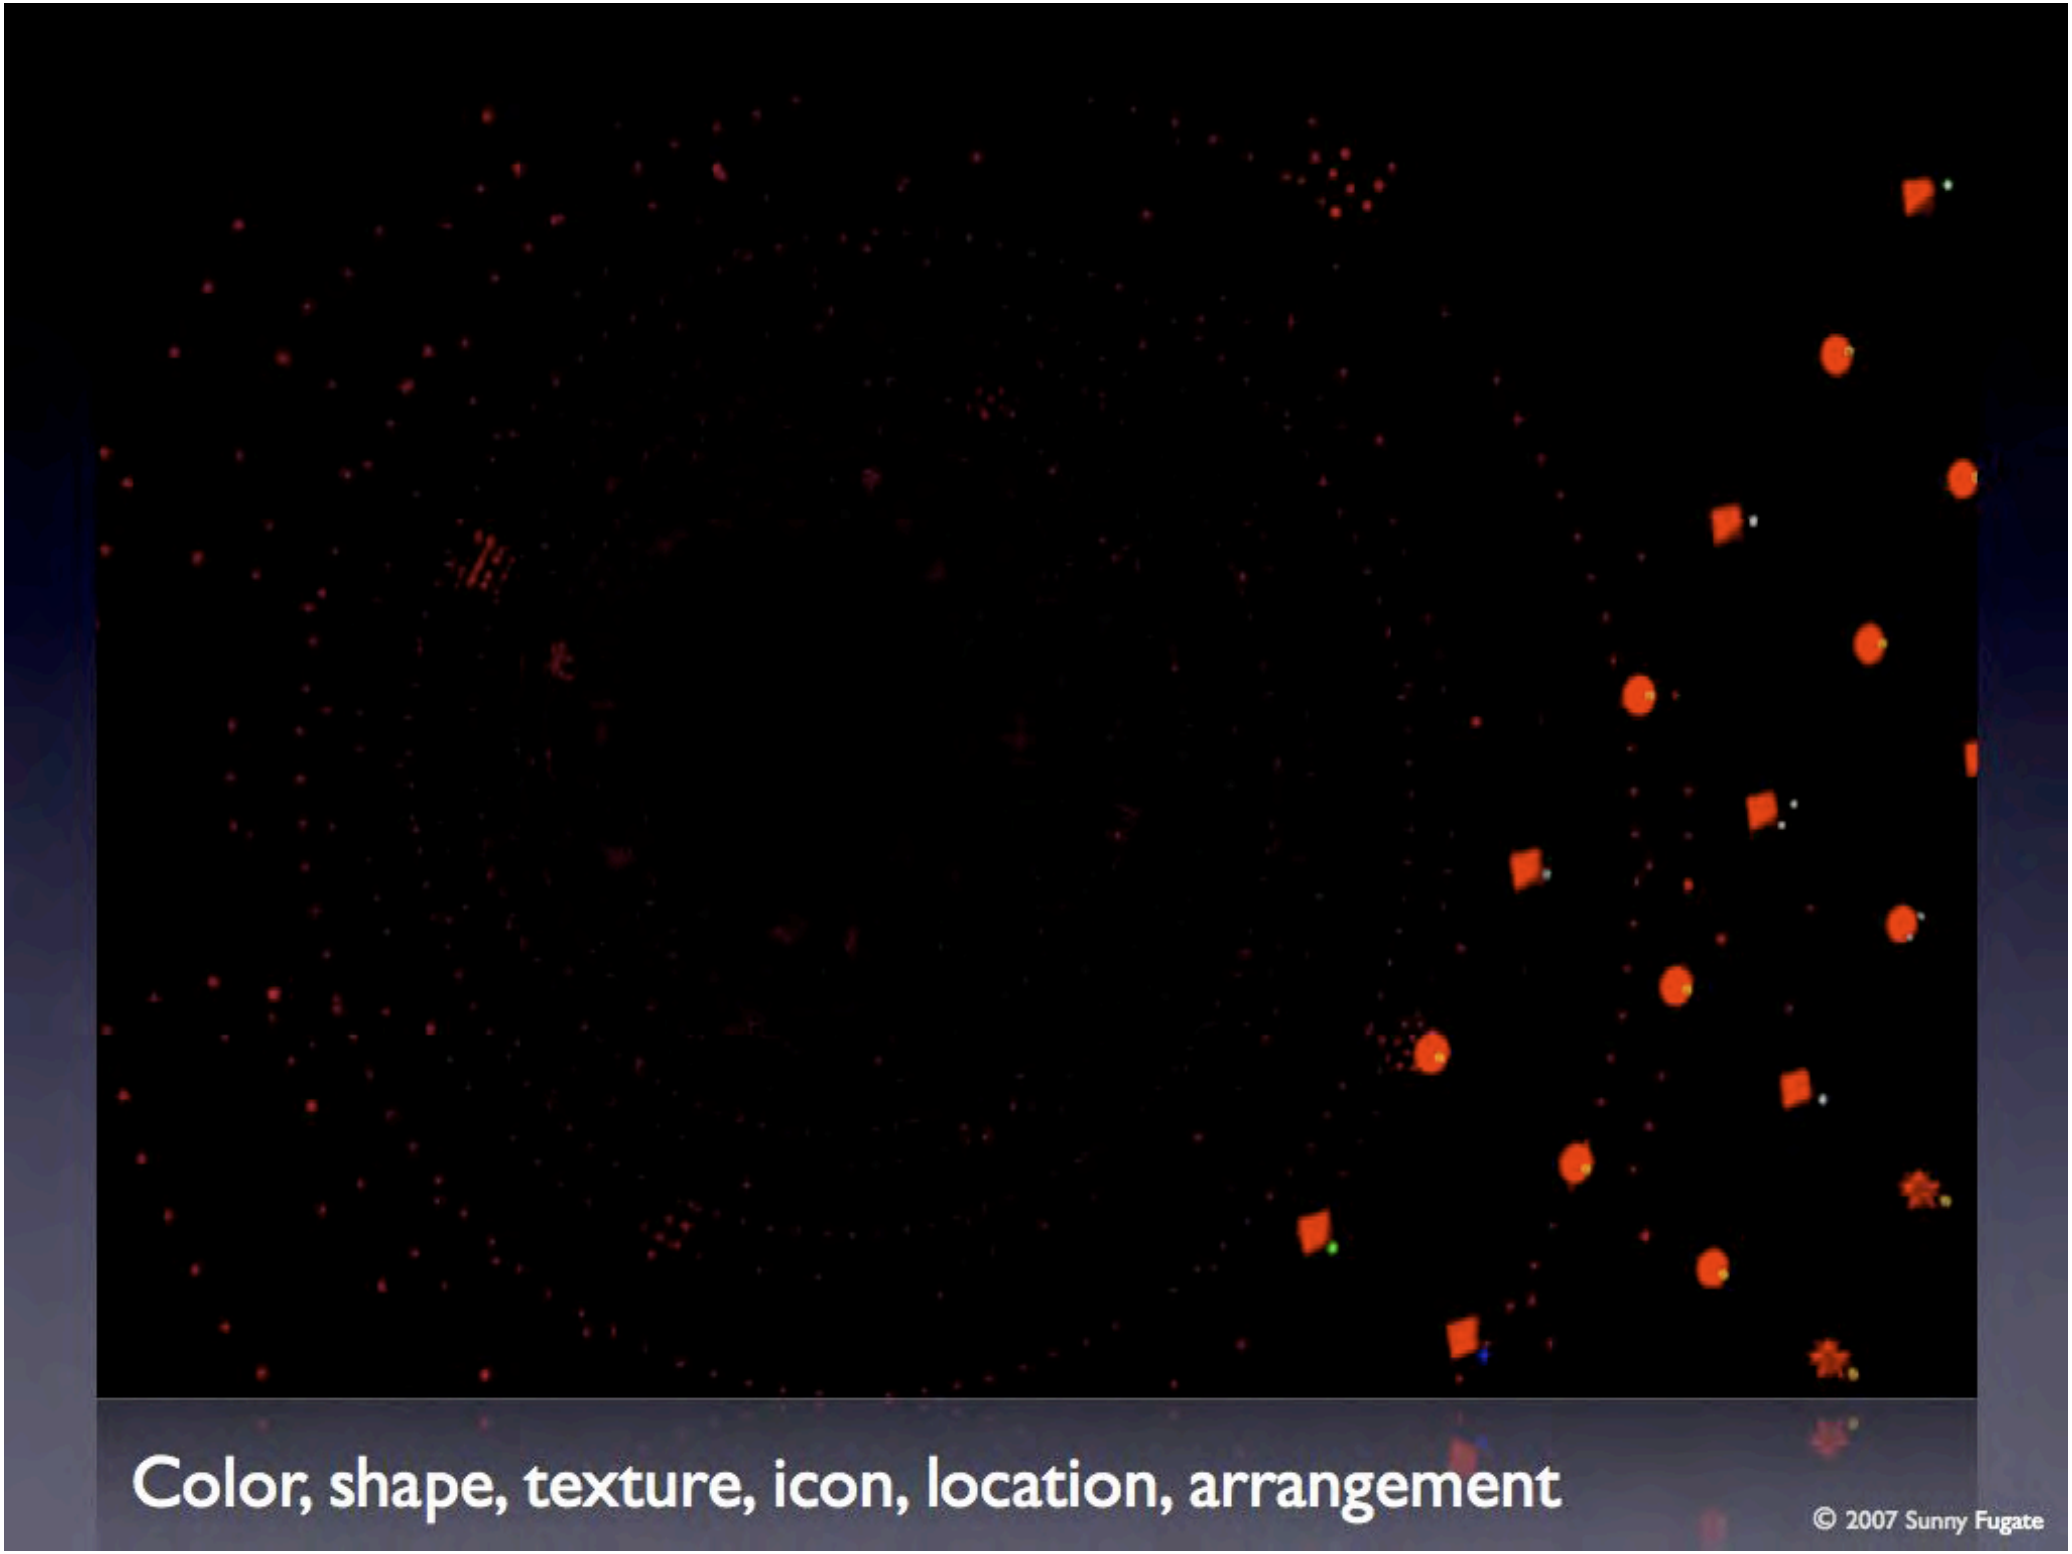

© 2007 Sunny Fugate



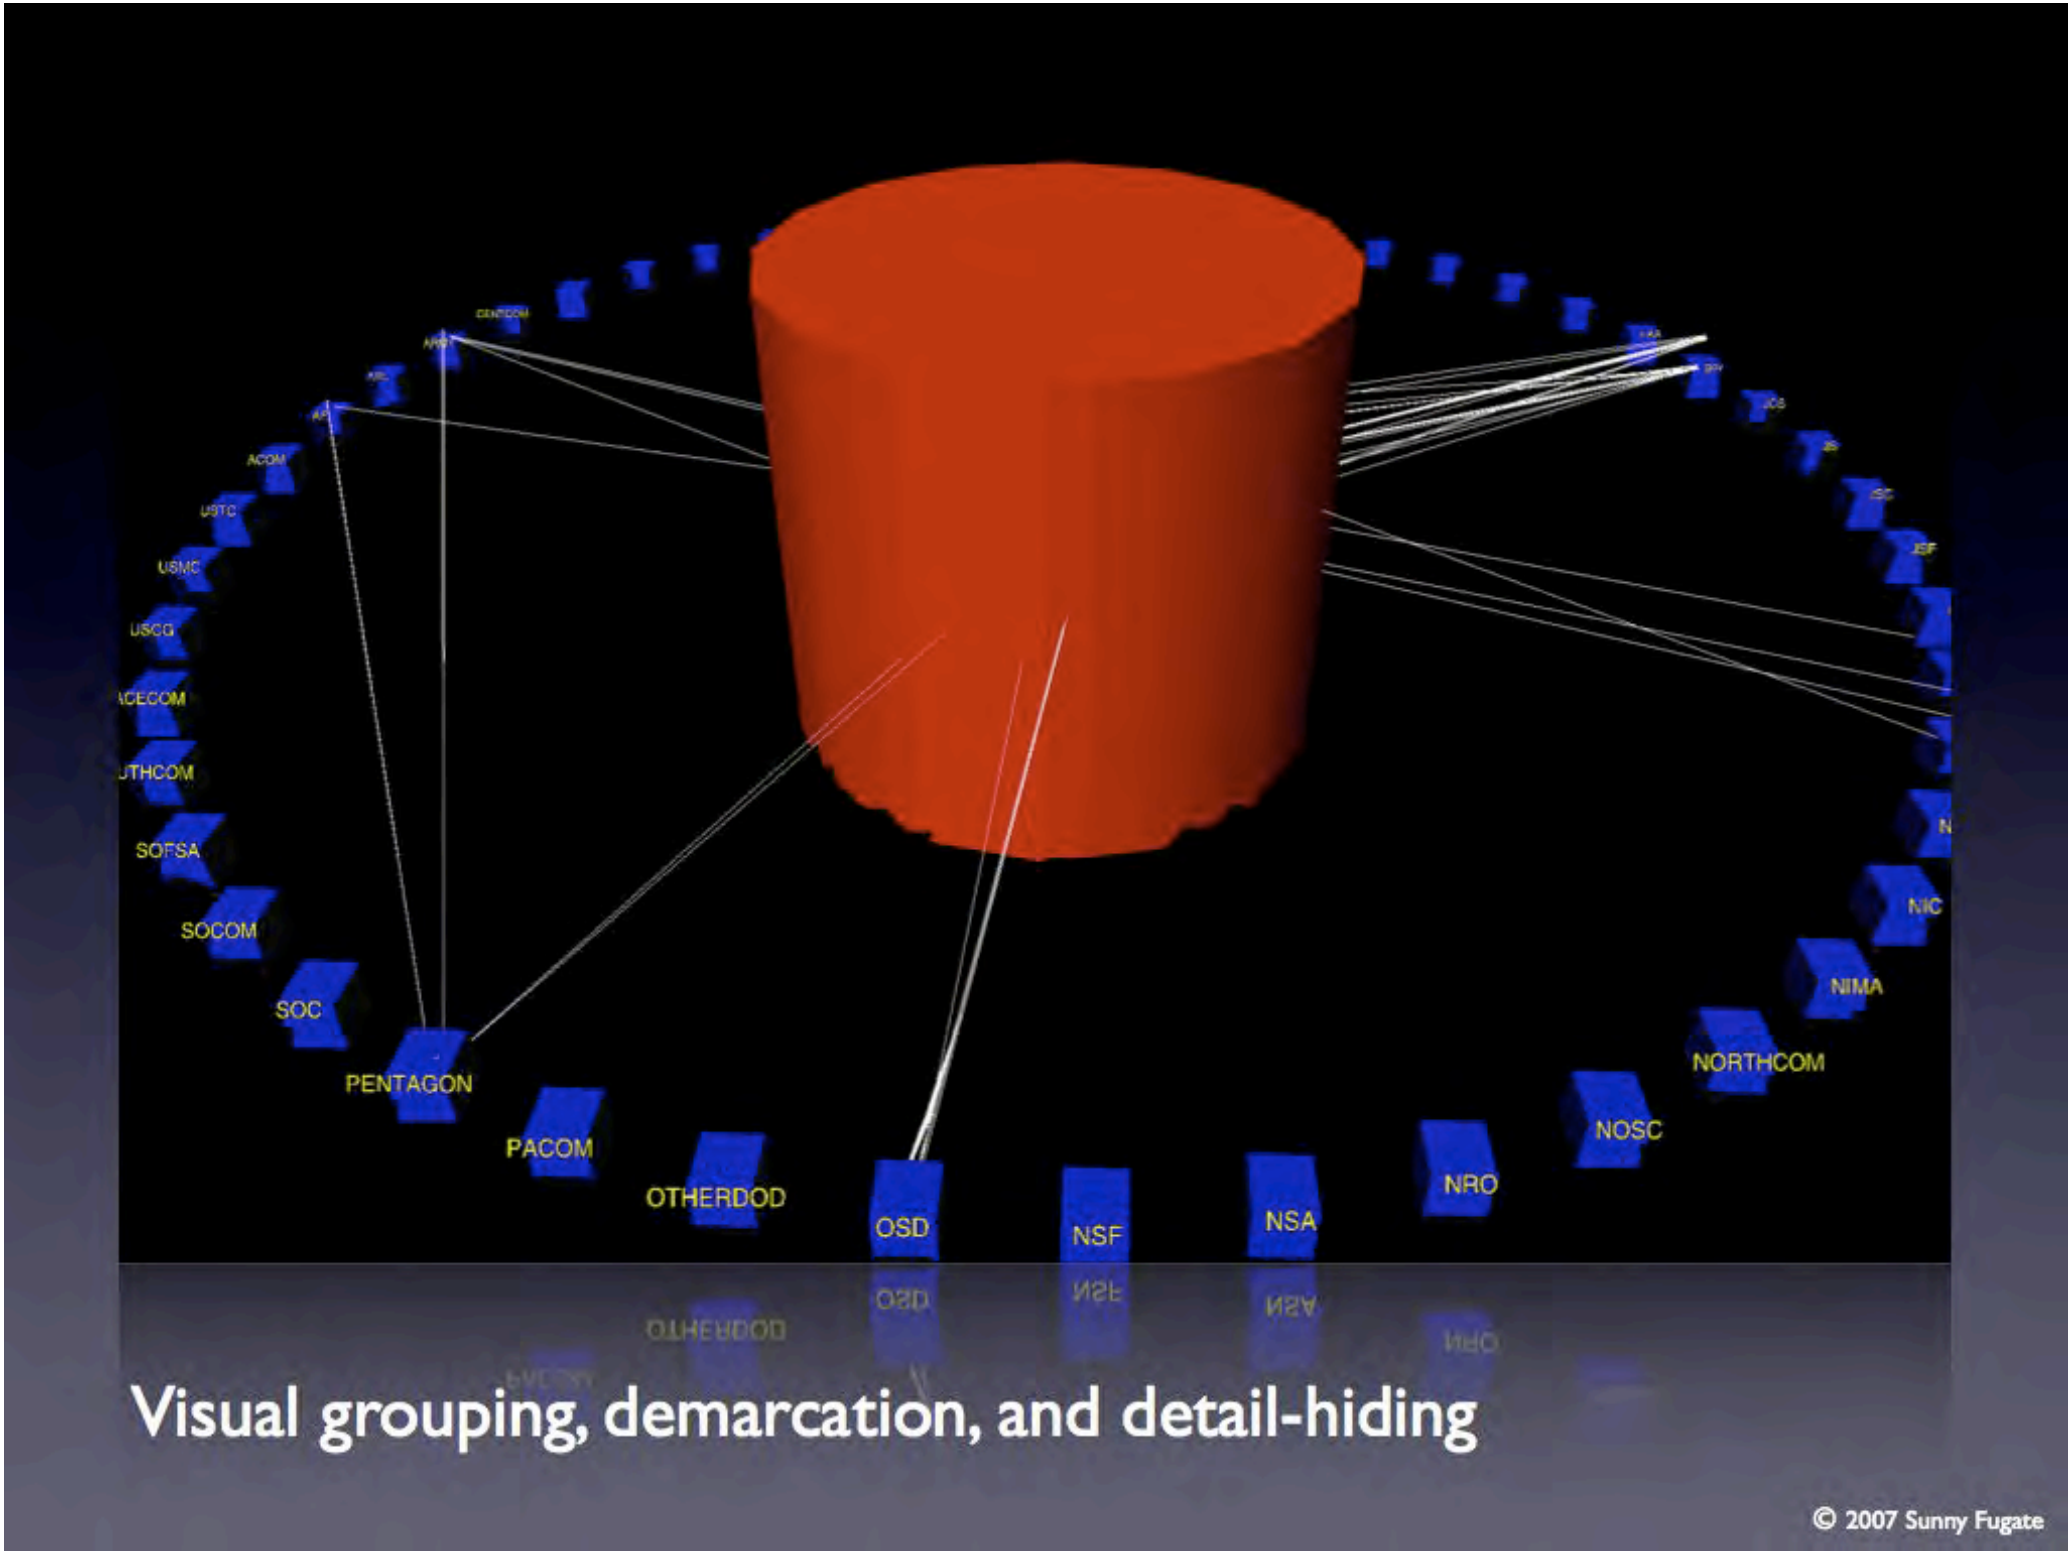
Topological layout discovered using hints in the data (e.g. TTL)

© 2007 Sunny Fugate


© 2007 Sunny Fugate

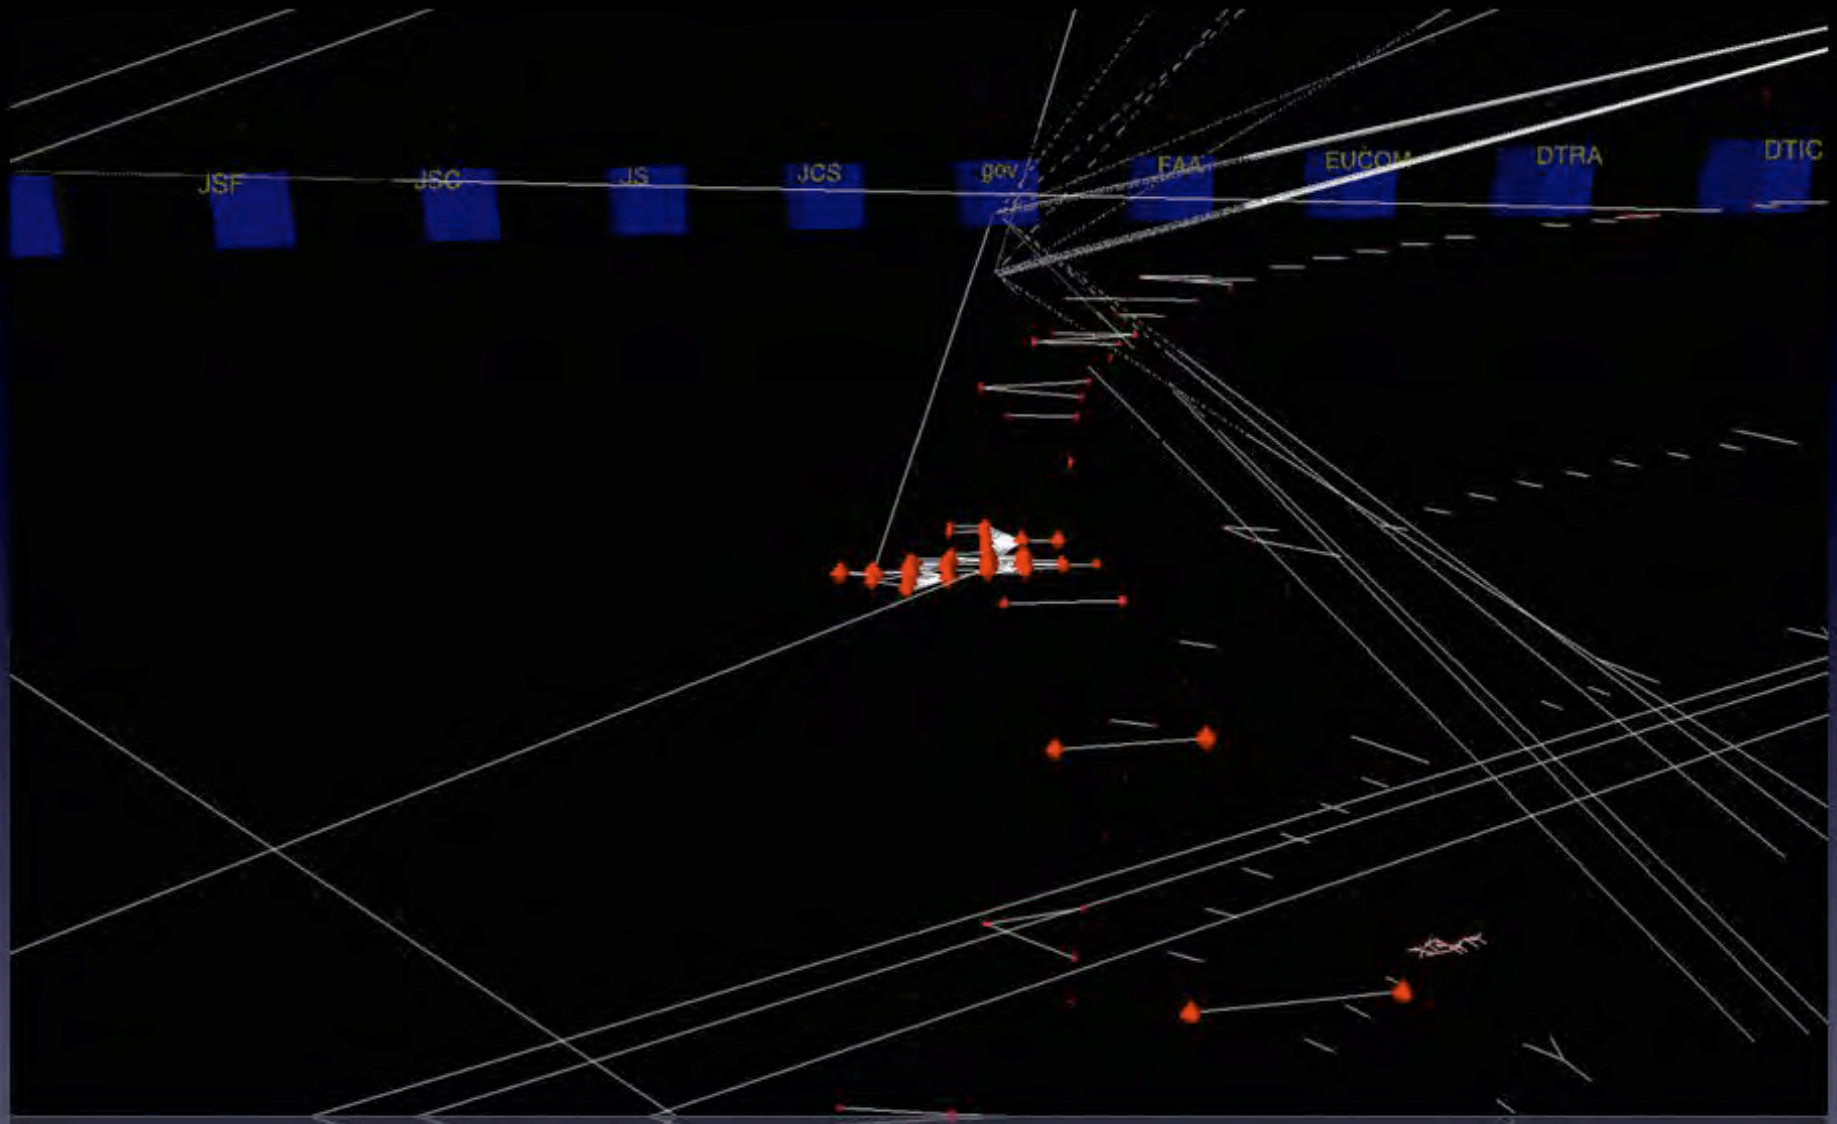


© 2007 Sunny Fugate



Visual grouping, demarcation, and detail-hiding

© 2007 Sunny Fugate



Expansive vantage points for network analysis

© 2007 Sunny Fugate

Shapes Vector Flow Viewer

- JTF-GNO funded effort to implement SV
  - Use SV architecture and components
  - DARPA demo system > operational system
  - New scripts, sensors, agents, and GUI
- Results
  - A visual augmentation of CLI
  - Produces a view of social topology
  - Intuitive view of gobs of data
    - static topology and event replay
  - Links statistical views and topology view

© 2007 Sunny Fugate

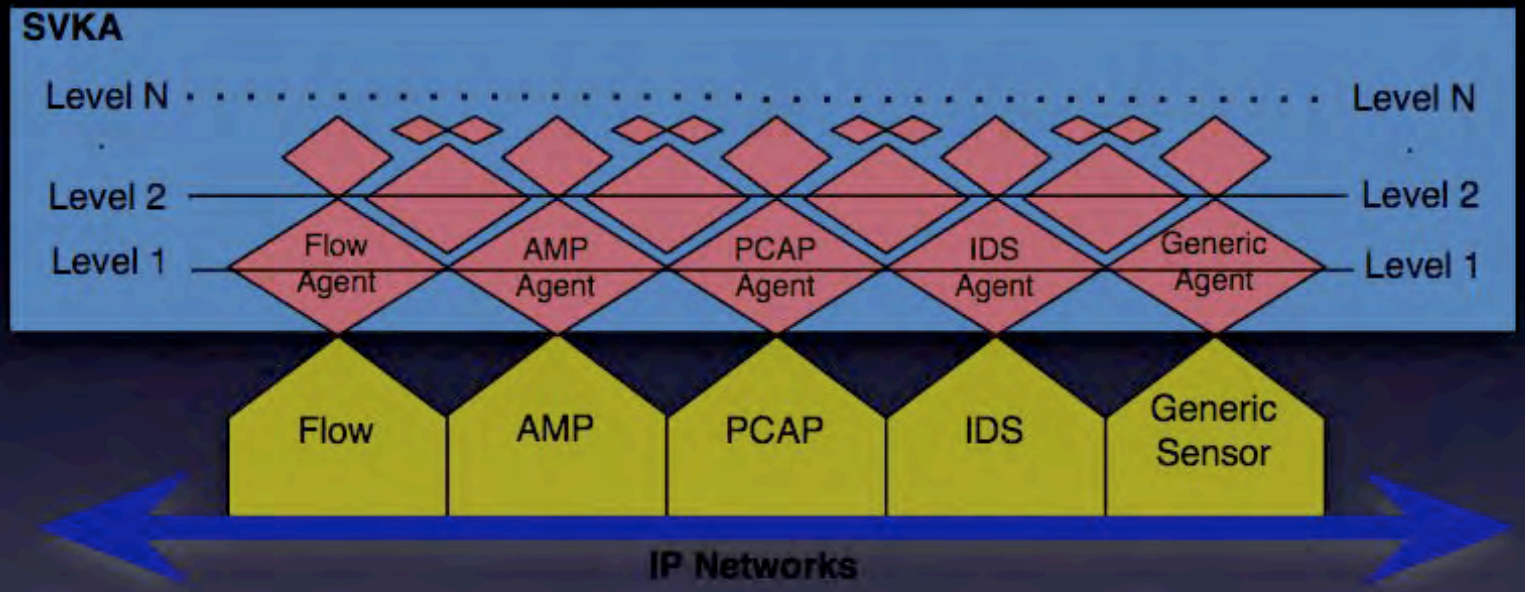
Flow Viewer GUI

- multiple stats views linked to visuals
- playback specific ranges & loop
- adjust replay velocity
- time-skip
- IP and attribute hotlists
- dynamic filtering controls
  - GUI managed rwfilter
  - filter using SV ontology
- integration between flow, AMP, IDS, & PCAP



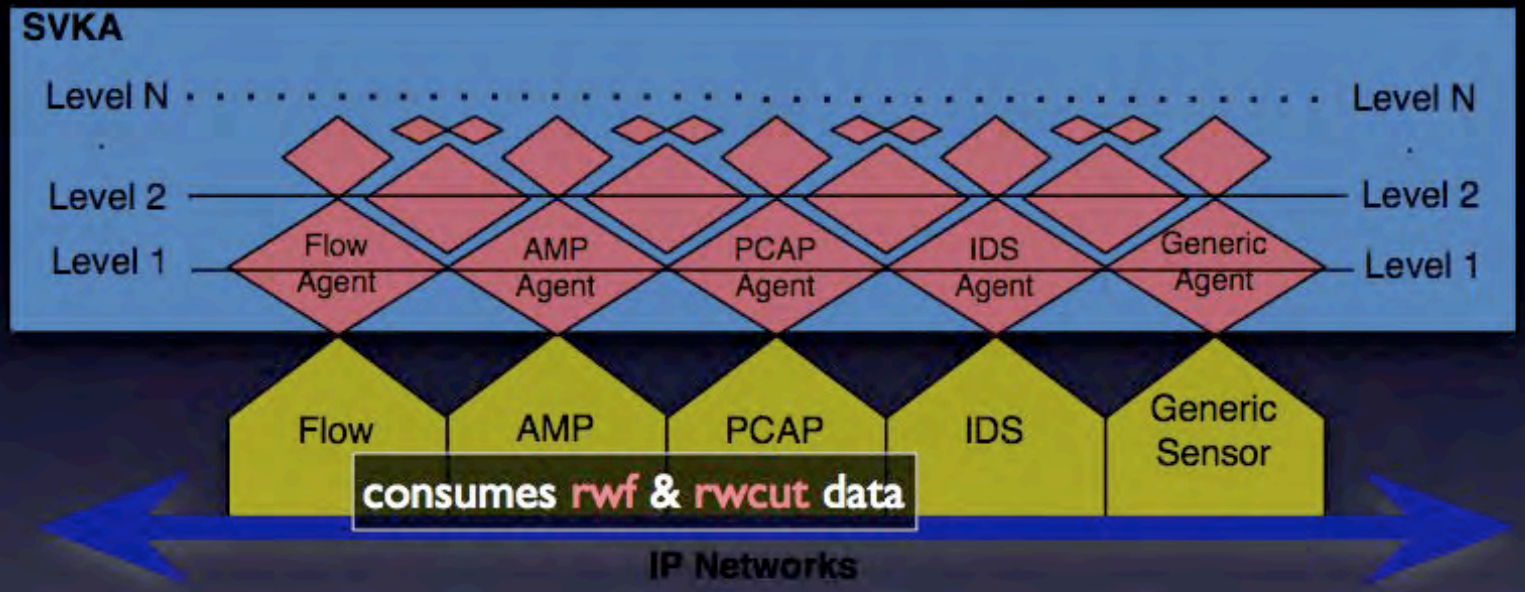
© 2007 Sunny Fugate

Flow Viewer Sensors



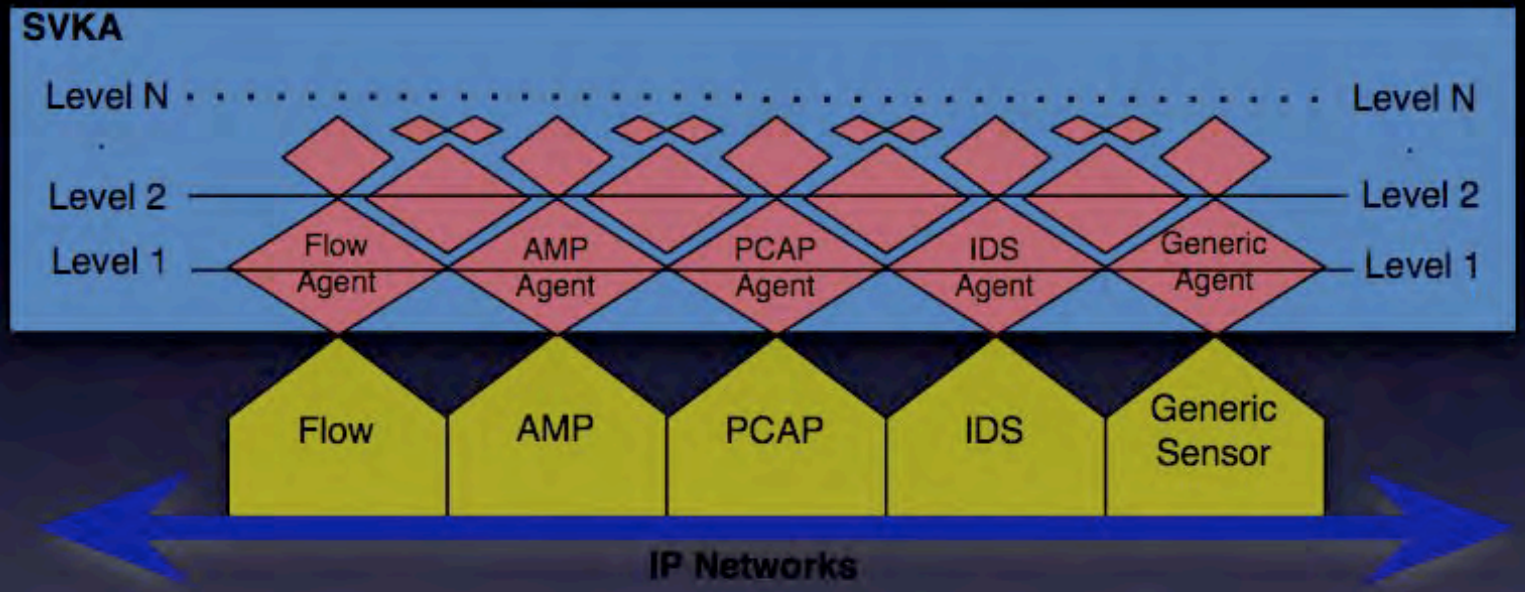
© 2007 Sunny Fugate

Flow Viewer Sensors



© 2007 Sunny Fugate

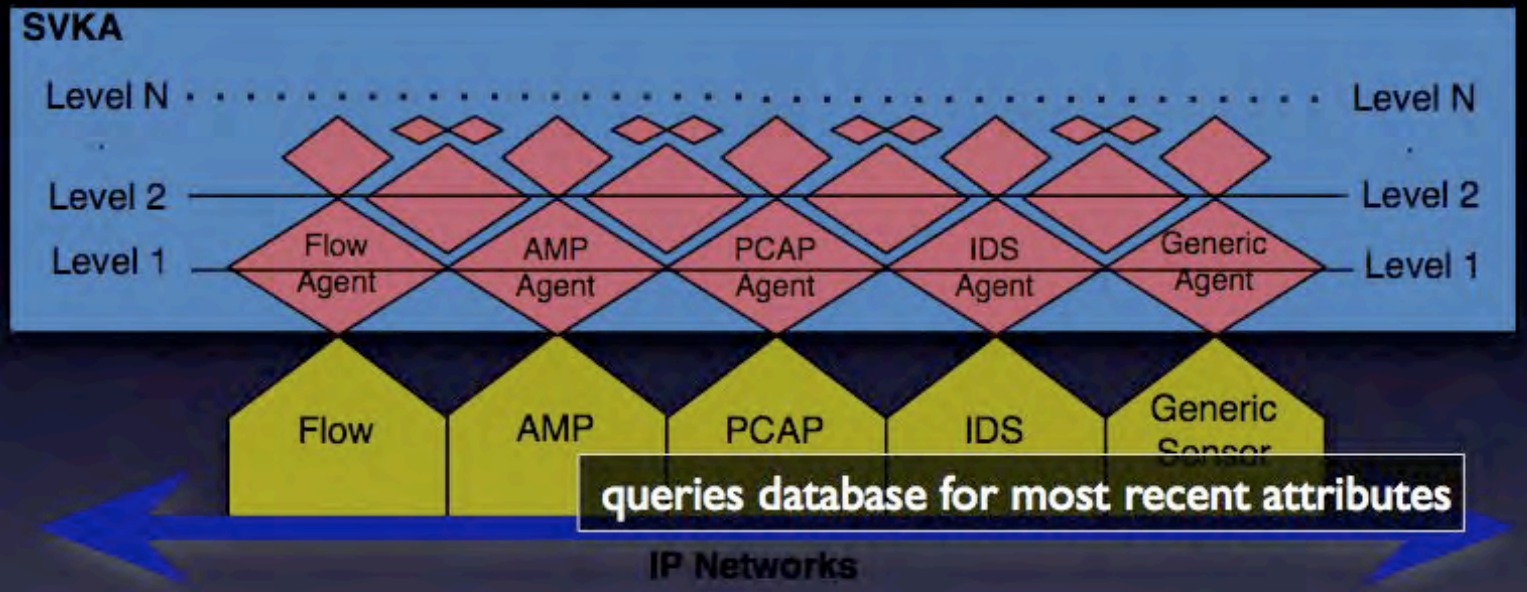
Flow Viewer Sensors



Flow Agent consumes rwf & rwcut data

© 2007 Sunny Fugate

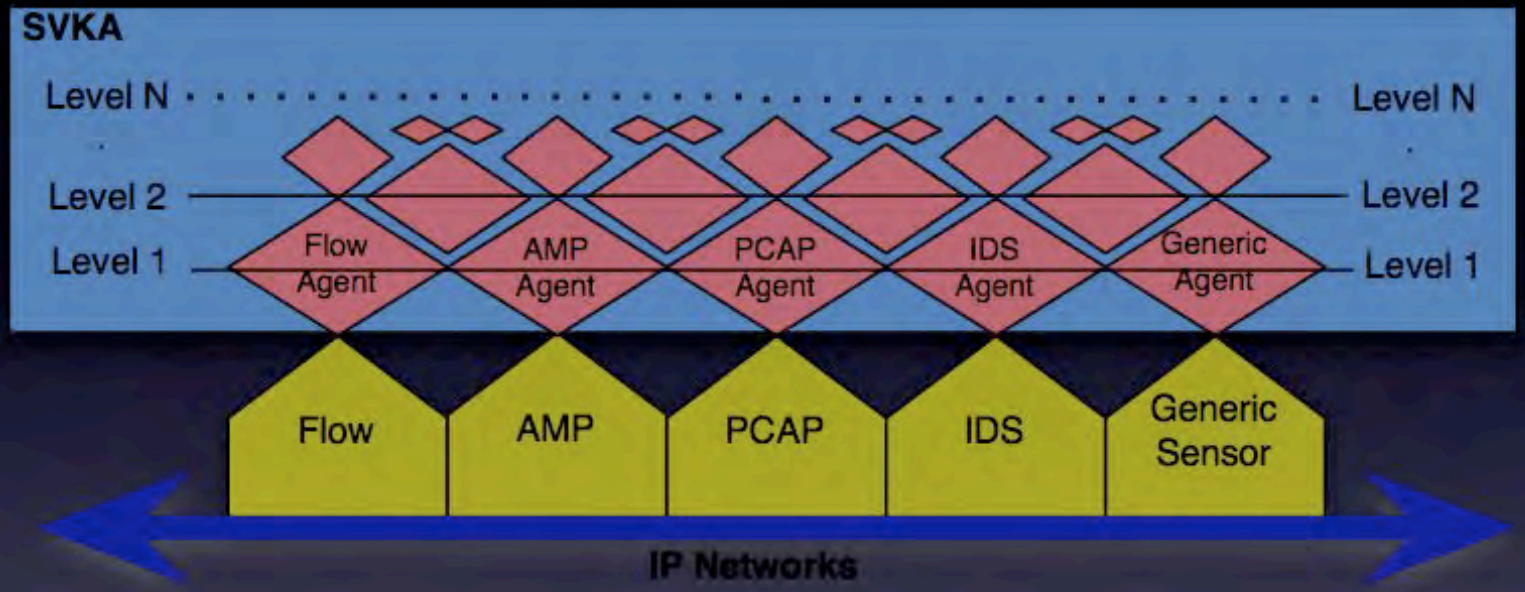
Flow Viewer Sensors



Flow Agent consumes rwf & rwcut data

© 2007 Sunny Fugate

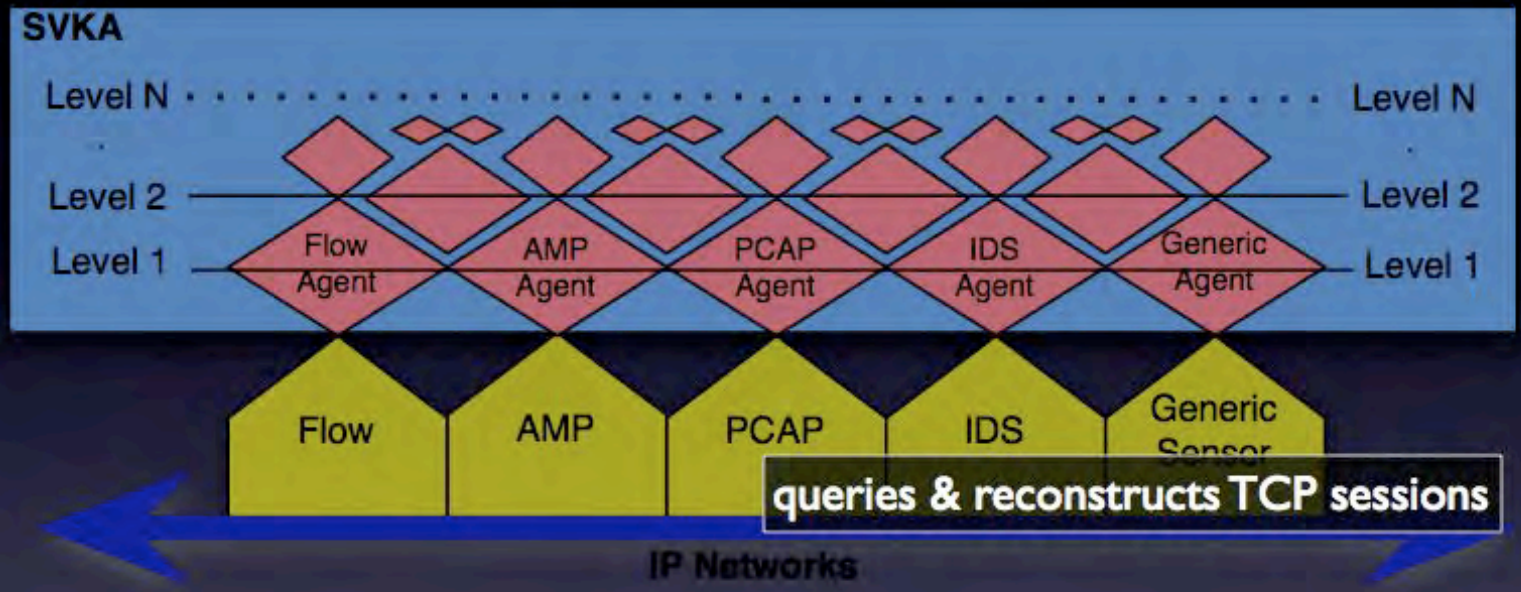
Flow Viewer Sensors



Flow Agent	consumes rwf & rwcut data
AMP Agent	queries database for most recent attributes

© 2007 Sunny Fugate

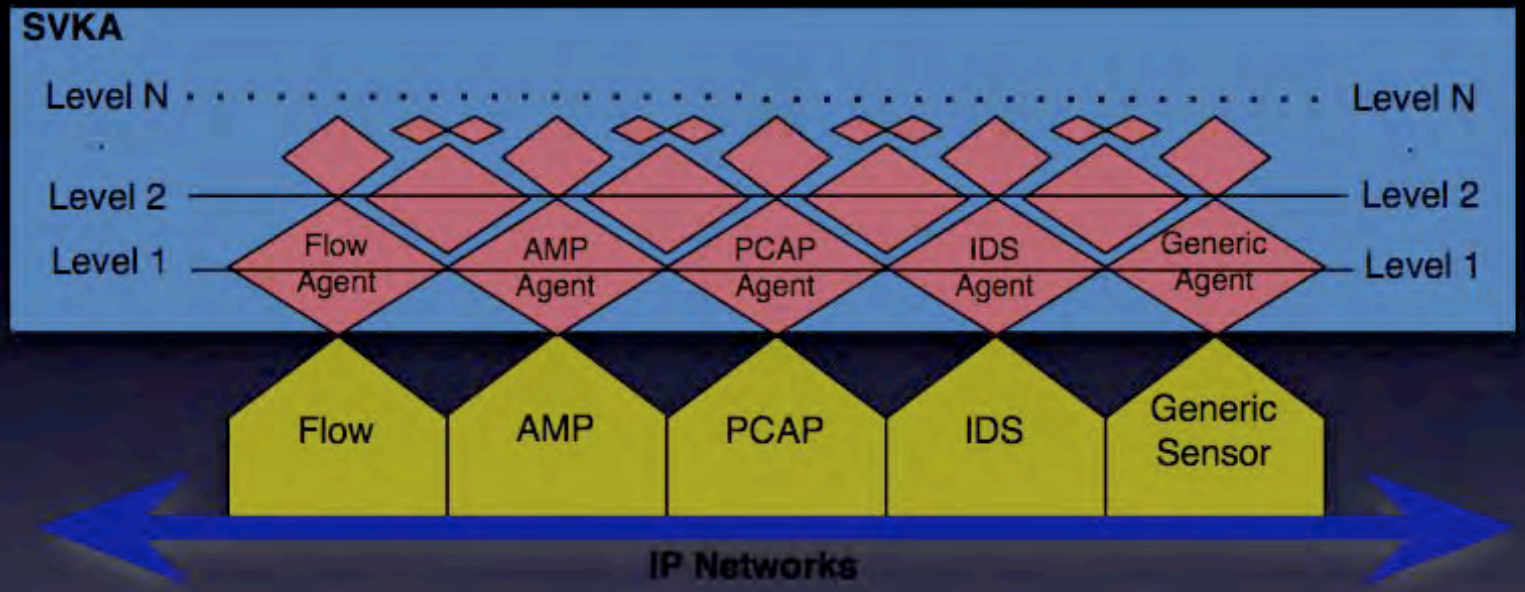
Flow Viewer Sensors



Flow Agent	consumes rwf & rwcut data
AMP Agent	queries database for most recent attributes

© 2007 Sunny Fugate

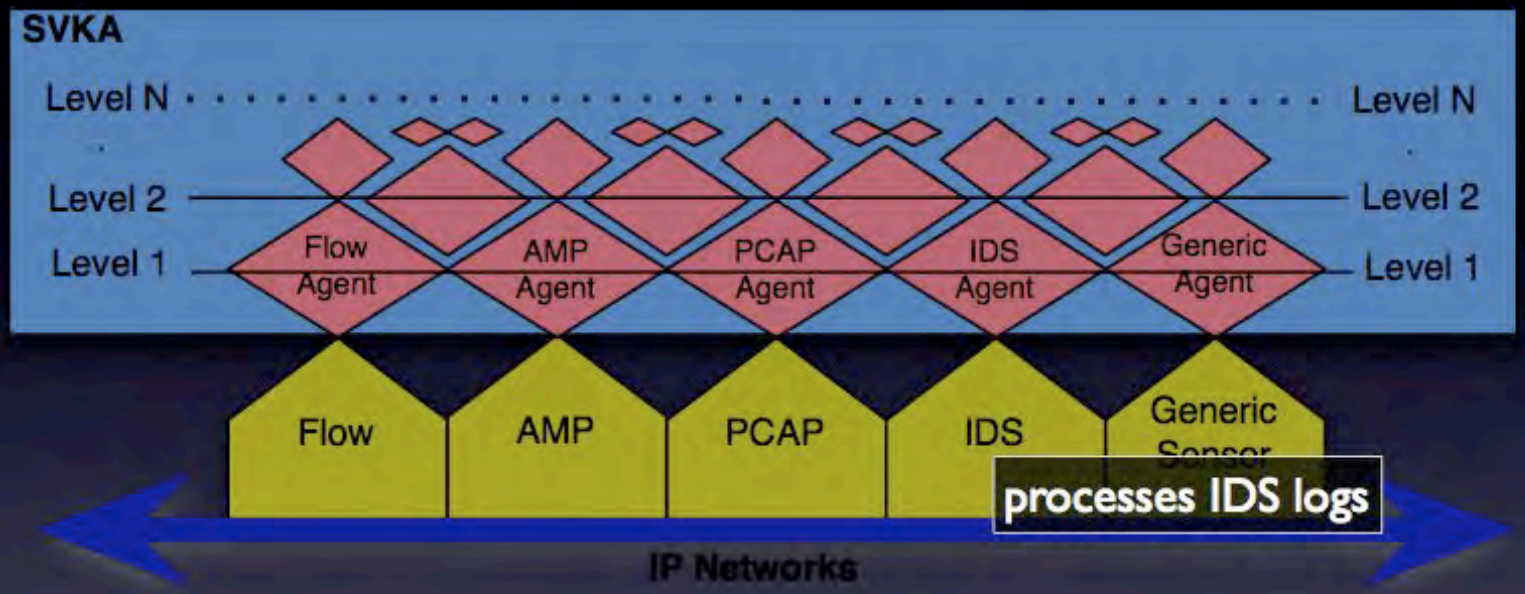
Flow Viewer Sensors



Flow Agent	consumes rwf & rwcut data
AMP Agent	queries database for most recent attributes
PCAP Agents	queries & reconstructs TCP sessions

© 2007 Sunny Fugate

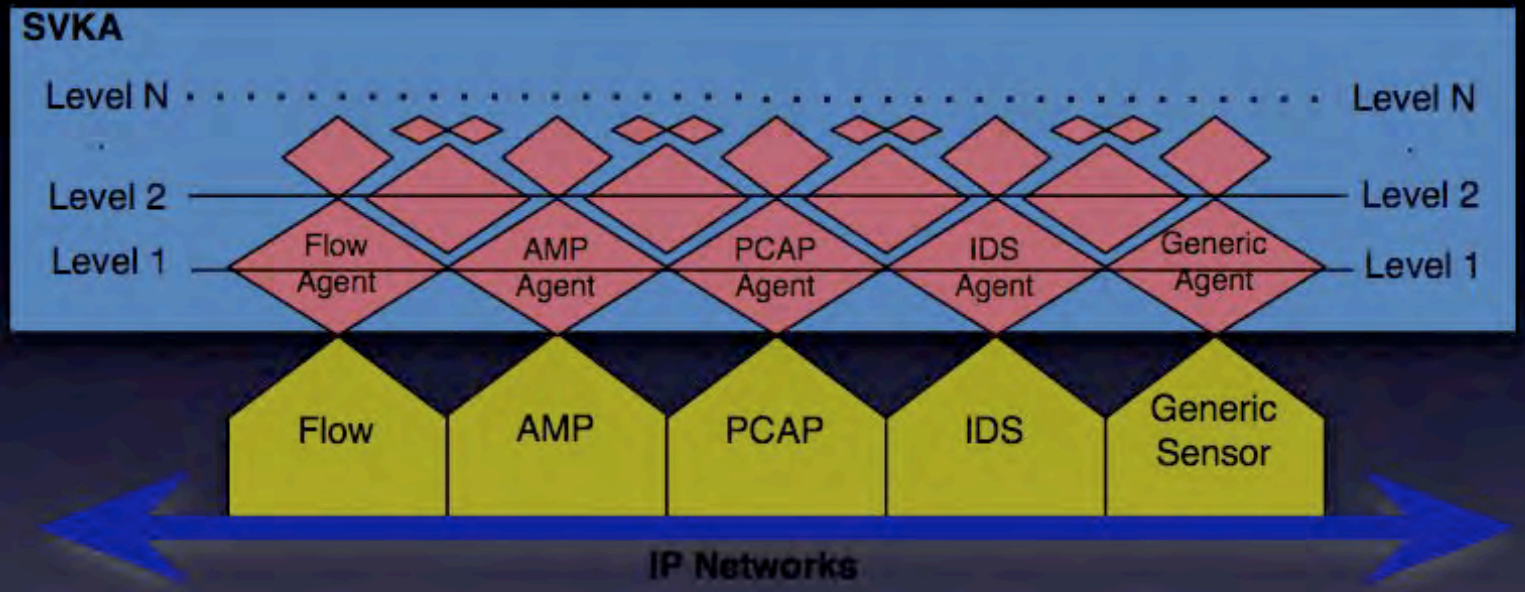
Flow Viewer Sensors



Flow Agent	consumes rwf & rwcut data
AMP Agent	queries database for most recent attributes
PCAP Agents	queries & reconstructs TCP sessions

© 2007 Sunny Fugate

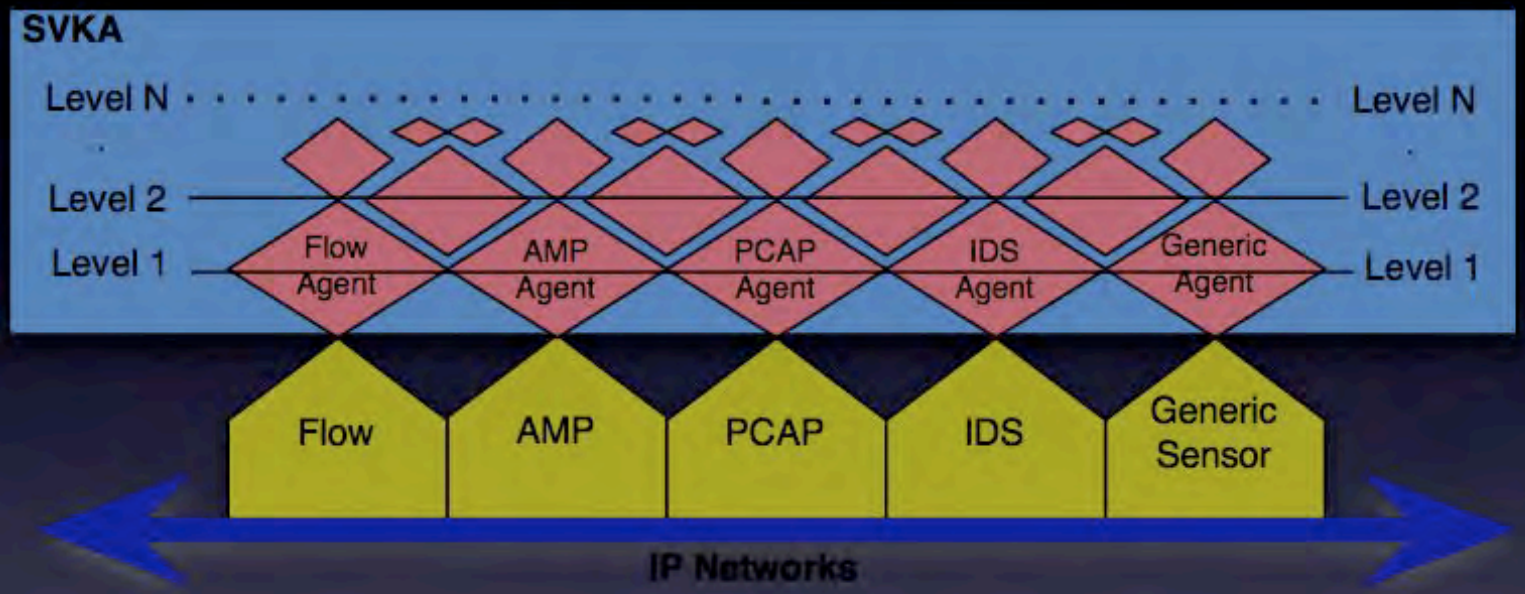
Flow Viewer Sensors



Flow Agent	consumes rwf & rwcut data
AMP Agent	queries database for most recent attributes
PCAP Agents	queries & reconstructs TCP sessions
IDS Agents	processes IDS logs

© 2007 Sunny Fugate

Flow Viewer Sensors



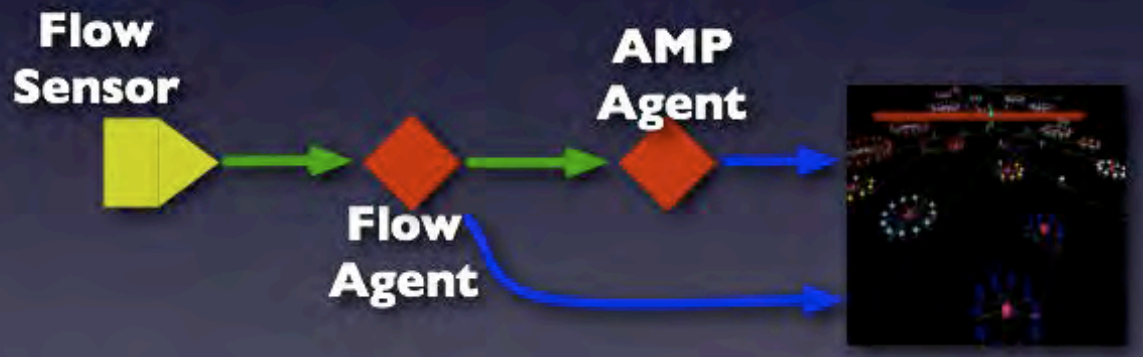
Flow Agent	consumes rwf & rwcut data ✓
AMP Agent	queries database for most recent attributes ✓
PCAP Agents	queries & reconstructs TCP sessions
IDS Agents	processes IDS logs

© 2007 Sunny Fugate

Flow Viewer
Intelligent Agents

Flow Sensor
- Converts flow into ontology
- produces facts

AMP Agent
- uses correlations from Flow Agent
- query made on every unique IP seen
- produces visual events

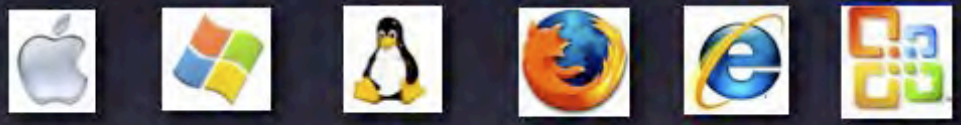


Flow Agent
- correlates records
- counts and corroborates
- produces inferences
- produces visual events

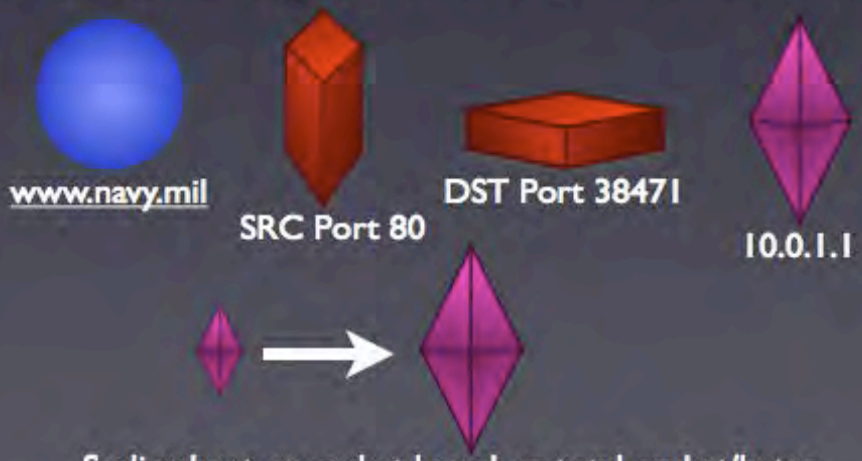
© 2007 Sunny Fugate

Flow Viewer Visual Language

Leverage cultural knowledge



Use metaphors for abstract

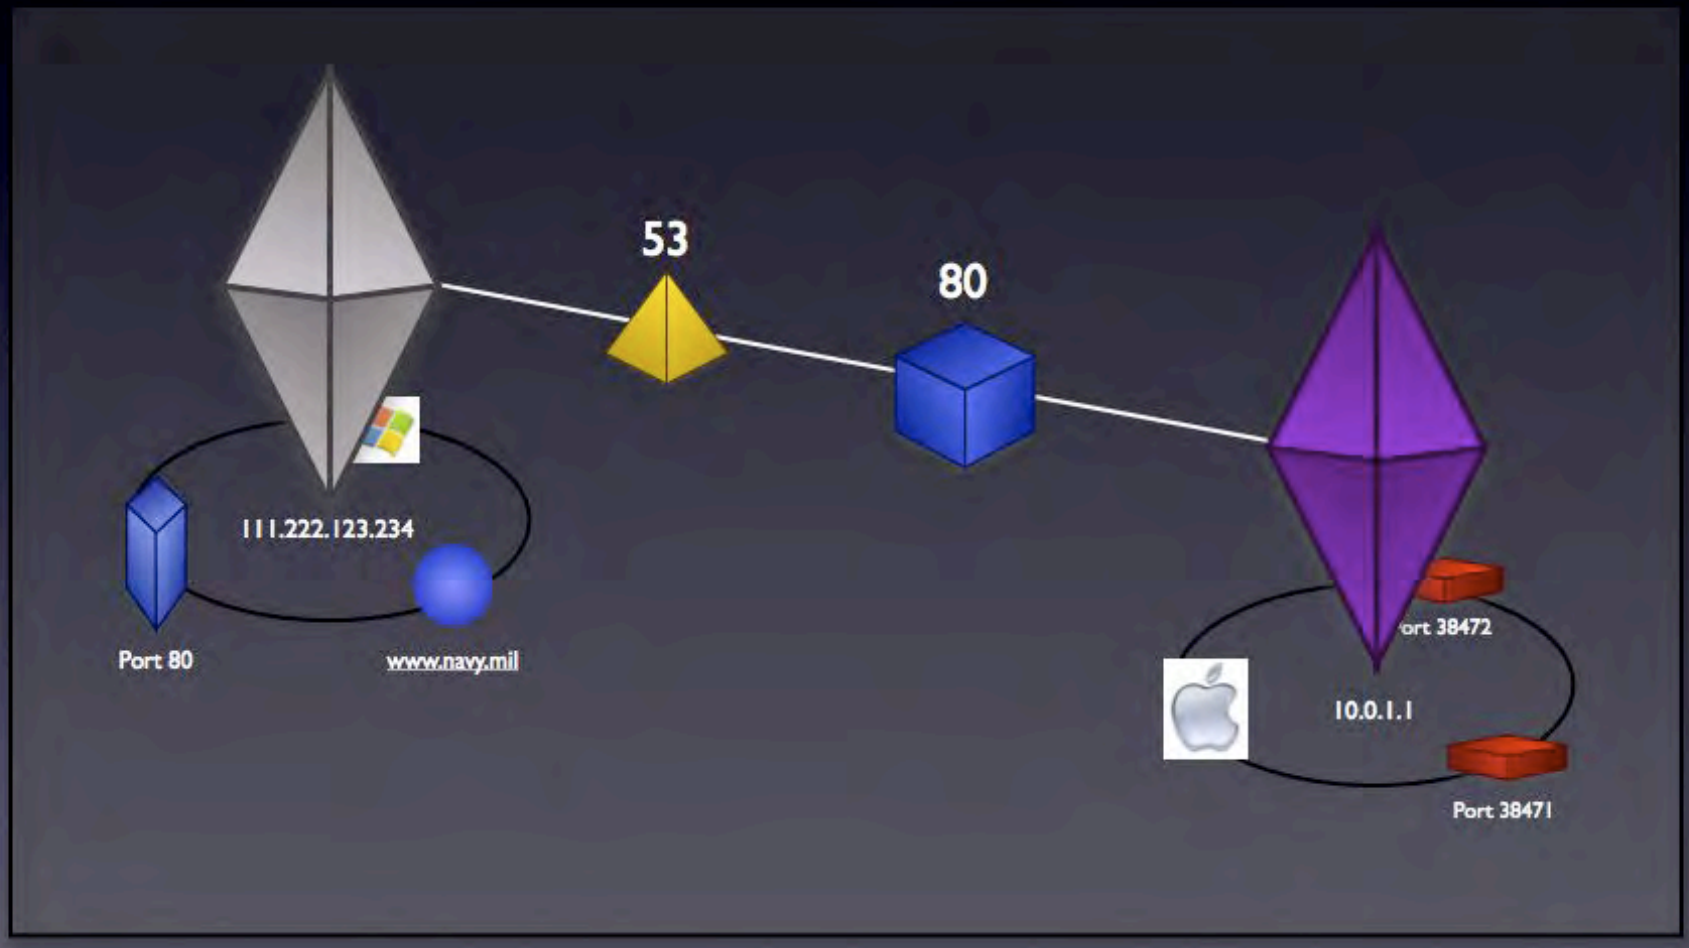


Color by ownership



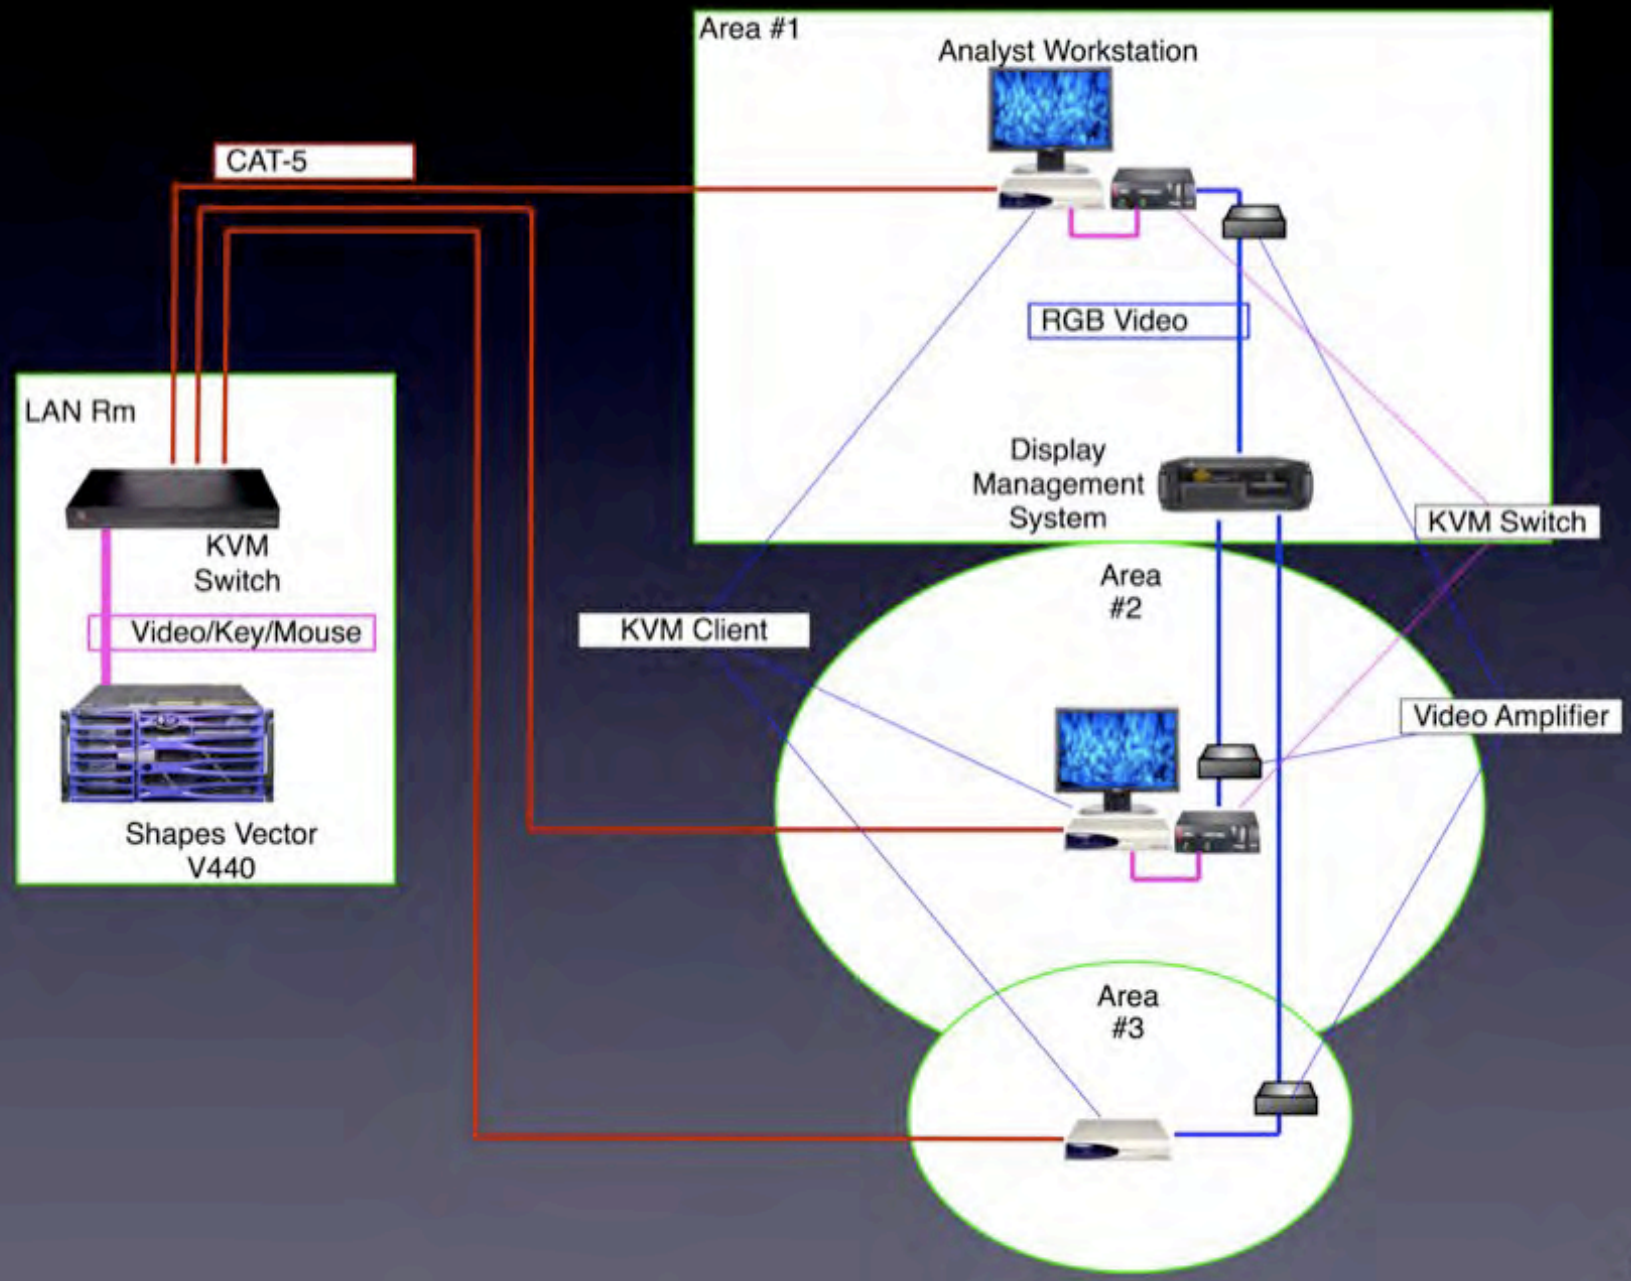
© 2007 Sunny Fugate

Flow Viewer Visual Language



© 2007 Sunny Fugate

Test installation



© 2007 Sunny Fugate

Flow Viewer

Visualization

- **Tested** using:
  - 100-5000 nodes
  - 1M-3M flows
  - 10K-300K flows per hour
- Integrated **filtering** (rwfilter, SVKA filtering, visual filter)
- Visual ID
- **Queries**
- **Grouping** (e.g. domain, netblock, vulnerability)
- **Replay**-mode or Real-time
- Historic **visual context**
  - Replay 'on top of' known incident

© 2007 Sunny Fugate

Flow Viewer data prep

Include

- Incoming & outgoing
- Hub & core-to-core traffic
- Wide port ranges
- Time-span wider than the activity (minutes to hours)
- Suspect IPs and ranges

Filter

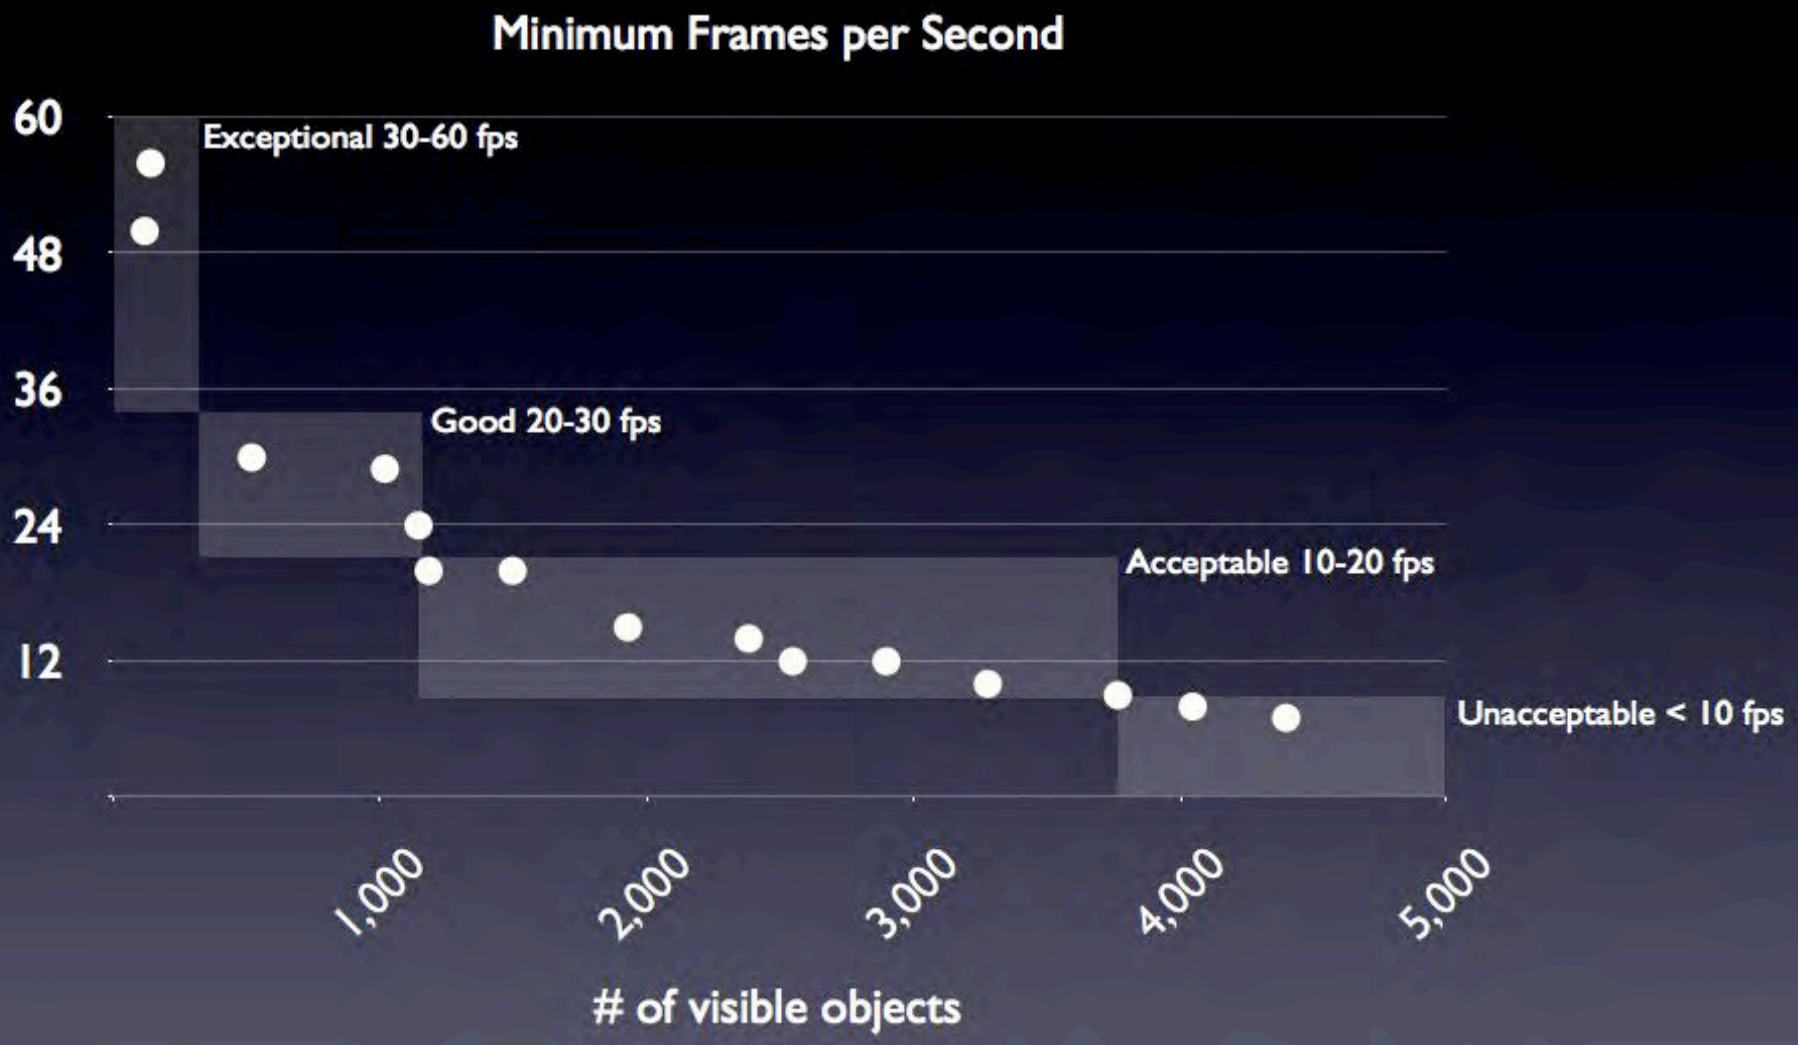
- Superfluous port traffic (e.g. 80, 53, 25)
- IPs that are unrelated to the incident

Sampling & Time

- Dense data
- Smear data across time resolution (~1 second)

© 2007 Sunny Fugate

Flow Viewer Performance



**Graphics performance on dual 1.5GHz SPARC SunFire v440 with Sun XVR 1200

© 2007 Sunny Fugate

Flow Viewer Performance

Real-time Performance	Real-time Records / Hour	Optimal playback rate
Optimal	10K-30K/hour	10X Real-time
Acceptable	40K-100K/hour	Real-time
Poor	100K-300K/hour	1/10 X Real-Time

Sparse data sets can be viewed quickly
e.g. months of data in minutes

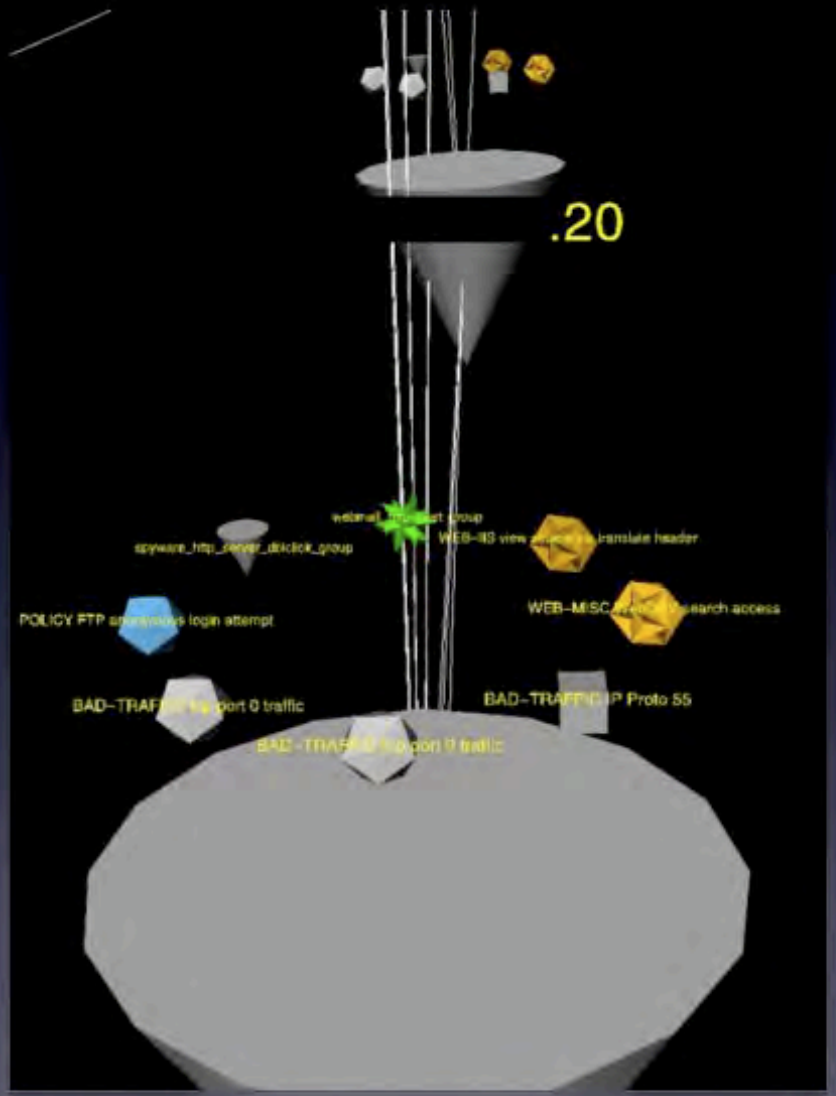
Dense data sets can be viewed slowly or filtered
e.g. seconds of data in minutes

© 2007 Sunny Fugate

Knowledge Depth vs Breadth

What trade-offs are we making?

- UI Feedback?
  - Haptic vs visual feedback
- Data access?
  - Random sequential access
- Training?
  - Under-learned vs over-learned
  - Tool complexity
- Meaning?
  - Visual semantic vs text
  - Intuitive/Iconic vs cryptic/coded



© 2007 Sunny Fugate

References

- [1] T. Abraham, Electronics, and Surveillance Research Laboratory (Australia). Information Technology Division. IDDM: Intrusion Detection Using Data Mining Techniques. DSTO Electronics and Surveillance Research Laboratory, 2001.
- [2] Mark Anderson, Dean Engelhardt, Damian Marriott, and Suneel Randhawa. Data processing and observation system, August 1 2006.
- [3] Mark Anderson, Dean Engelhardt, Damian Marriott, and Suneel Randhawa. Data view of a modelling system, April 11 2006.
- [4] H. H. Clark and W. G. Chase. On the process of comparing sentences against pictures. Cognitive Psychology, 3:472–517, 1972.
- [5] Herbert A. Colle and Gary B. Reid. The room effect: Metric spatial knowledge of local and separated regions. Presence: Teleoperators and Virtual Environments, 7(2):116–128, 1998.
- [6] Science Applications International Corporation. Intrusion Detection System System Protection Profile. National Security Agency, 9800 Savage Road, Fort Meade MD, 20755, version 1.4 edition, February 2002.
- [7] Stephen W. Draper and Donald A. Norman. User Centered System Design: New Perspectives on Human-computer Interaction. CRC, 1 edition, 1986.
- [8] D. Engelhardt and M. Anderson. A distributed multi-agent architecture for computer security situational awareness. Information Fusion, 2003. Proceedings of the Sixth International Conference of, 1, 2003.
- [9] Sunny Fugate. Visual language for tactical communication. In Proceedings of the First Annual Visual and Iconic Language Conference. SPAWAR Systems Center, San Diego, August 2007.
- [10] Sunny Fugate, Emily W. Medina, LorRaine Duffy, Dennis Magsombol, Omar Amezcua, Gary Rogers, and Marion Ceruti. Next-generation tactical-situation-assessment technology (tsat): Iconic language. In Sunny Fugate, editor, Proceedings of the First Annual Visual and Iconic Language Conference. SPAWAR Systems Center, August 2007.
- [11] David Gamon and Allen D. Bragdon. Brains That Work A Little Bit Differently: Recent Discoveries About Common Brain Diversities. Barnes and Noble, 2000.
- [12] James K. Hahn, Hesham Fouad, Larry Gritz, and Jong Won Lee. Integrating sounds and motions in virtual environments. Presence: Teleoperators and Virtual Environments, 7(1):67–77, 1998.
- [13] T. Munzner. INTERACTIVE VISUALIZATION OF LARGE GRAPHS AND NETWORKS. PhD thesis, STANFORD UNIVERSITY, 2000.
- [14] Jakob Nielsen. Usability Engineering (Interactive Technologies). Morgan Kaufmann, 1st edition, 1993.
- [15] CM Reed and NI Durlach. Short paper: Note on information transfer rates in human communication. Presence: Teleoperators and Virtual Environments, 7(5): 509–518, 1998.
- [16] Walter Shepherd. Shepherd's glossary of graphic signs and symbols;. Dent, London,, 1971.
- [17] Edward R. Tufte. Visual Explanations: Images and Quantities, Evidence and Narrative. Graphics Press, Cheshire, Conn., 1997.
- [18] TS TULLIS. An evaluation of alphanumeric, graphic, and color information displays. Human Factors, 23:541–550, 1981.
- [19] D.J. Ward, A.F. Blackwell, and D.J.C. MacKay. Dasher—a data entry interface using continuous gestures and language models. Proceedings of the 13th annual ACM symposium on User interface software and technology, pages 129–137, 2000.
- [20] G.J. Wills. Nicheworks-interactive visualization of very large graphs. Graph Drawing: 5th International Symposium, GD'97, Rome, Italy, September 18-20, 1997. Proceedings, 1997.

Images

- Jeff Han's Multi-Touch Screen Interface, Jeff Kubina, Flickr.com, license: http://creativecommons.org/licenses/by-sa/2.0/deed.en
- Atari joystick, duncan, Flickr.com, license: http://creativecommons.org/licenses/by-nc/2.0/deed.en
- Headphones, daxtoor, Flickr.com, license: http://creativecommons.org/licenses/by-sa/2.0/deed.en

© 2007 Sunny Fugate

SPAWAR
Systems Center
San Diego

Next Generation Tactical Situation Assessment Technology (NG-TSAT)

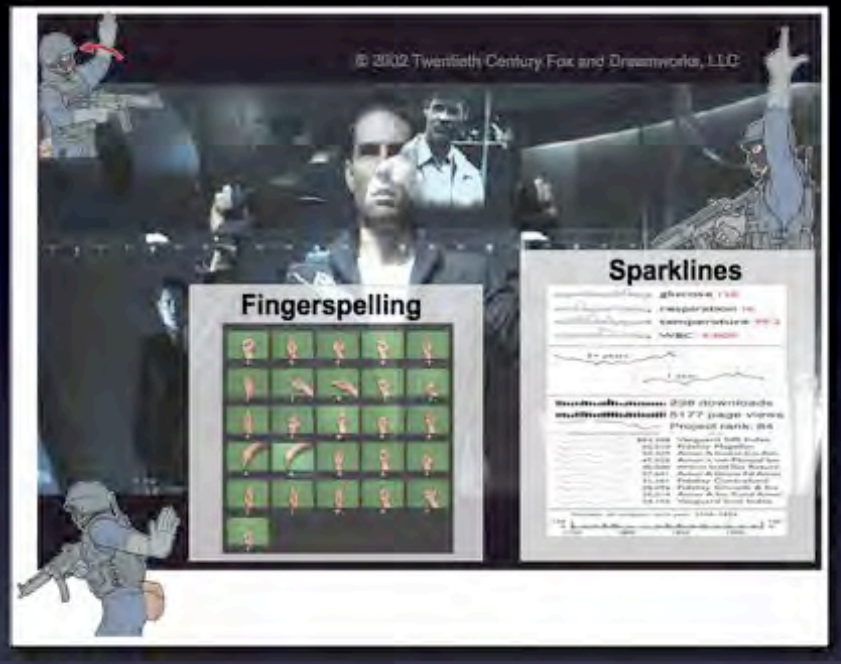
Objective: Next-generation Tactical Chat. Icon-based situation assessment (SA)language supported by wireless gesture-recognition gloves used in hostile or noisy (silence-mandated) environments

Description of Effort:

1.Linguistic Analysis: Analysis of current C² chat logs to determine speech patterns and repetitive SA concepts/themes

2. Iconic Language Development: Output of linguistic analysis determines candidate icons representing most prevalent SA "themes;" development of prototype C² iconic SA language

3. Wireless, Gesture-Recognition Gloves: Develop wireless gloves that recognize C² icons/gestures which can transmit across network to distributed warfighters (replacing keyboard input when in MOPP)



Benefits of TSAT:
Compressed Chat (25% ↓content; 50% ↓ reduction in production time) for rapid SA dissemination.
Gesture-recognition in very noisy, distributed ops, or in very austere environments (e.g., the moon)
Challenges:
1. No current method or theory for chat-meaning compression; currently done in prose; computer linguistic analysis of unstructured text still neoteric.
2. Wireless gesture recognition glove technology still in infant stages of development; focused on commercial animation support, not on disciplined language support
TRL: Chat: TRL 1-2; Gesture-recognition: TRL 1-4

Major Milestones FY06:
Linguistic analysis discovery of common C² SA themes
Development of icon/symbols for candidate SA themes
Development of proof-of-concept wireless gesture-recognition glove

Period of Performance: 2007-2012

PI contact info: Dr. LorRaine Duffy, (619) 553-9222, LorRaine.Duffy@navy.mil, SSC San Diego, CA

© SSC SD LorRaine Duffy 03042007

Synaesthesia

Synaesthesia: "a neurological condition in which two or more senses are coupled."

"loud color" "sharp laugh" "bitter wind"

grapheme color synesthesia - letters or numbers are perceived as inherently colored

© 2007 Sunny Fugate

How many numbers contain the digit 6?

9910 9972 3292 7602 82 9054
5636 2710 1944 6330 6560 8101
5177 1955 7029 4083 4643 5710
4935 2256 1495 1025 8375 8518
80 797 2610 3008 8784 1854 2383
9728 4523 573 5914 7975 281
6664 2682 7689 7753 273 5597
799 9960 1437 4534 8601 4563
6734 647 9409 6543 4827 2398
1532

© 2007 Sunny Fugate

Is this easier?



© 2007 Sunny Fugate

Emulating Synaesthesia

These methods can be used achieve **sequence disambiguation** and

9910	9972	3292	7602	82	9054
5636	2710	1944	6330	6560	8101
5177	1955	7029	4083	4643	5710
4935	2256	1495	1025	8375	8518
80	797	2610	3008	8784	1854	2383
9728	4523	573	5914	7975	281
6664	2682	7689	7753	273	5597
799	9960	1437	4534	8601	4563
6734	647	9409	6543	4827	2398
1532					



© 2007 Sunny Fugate

Emulating Synaesthesia

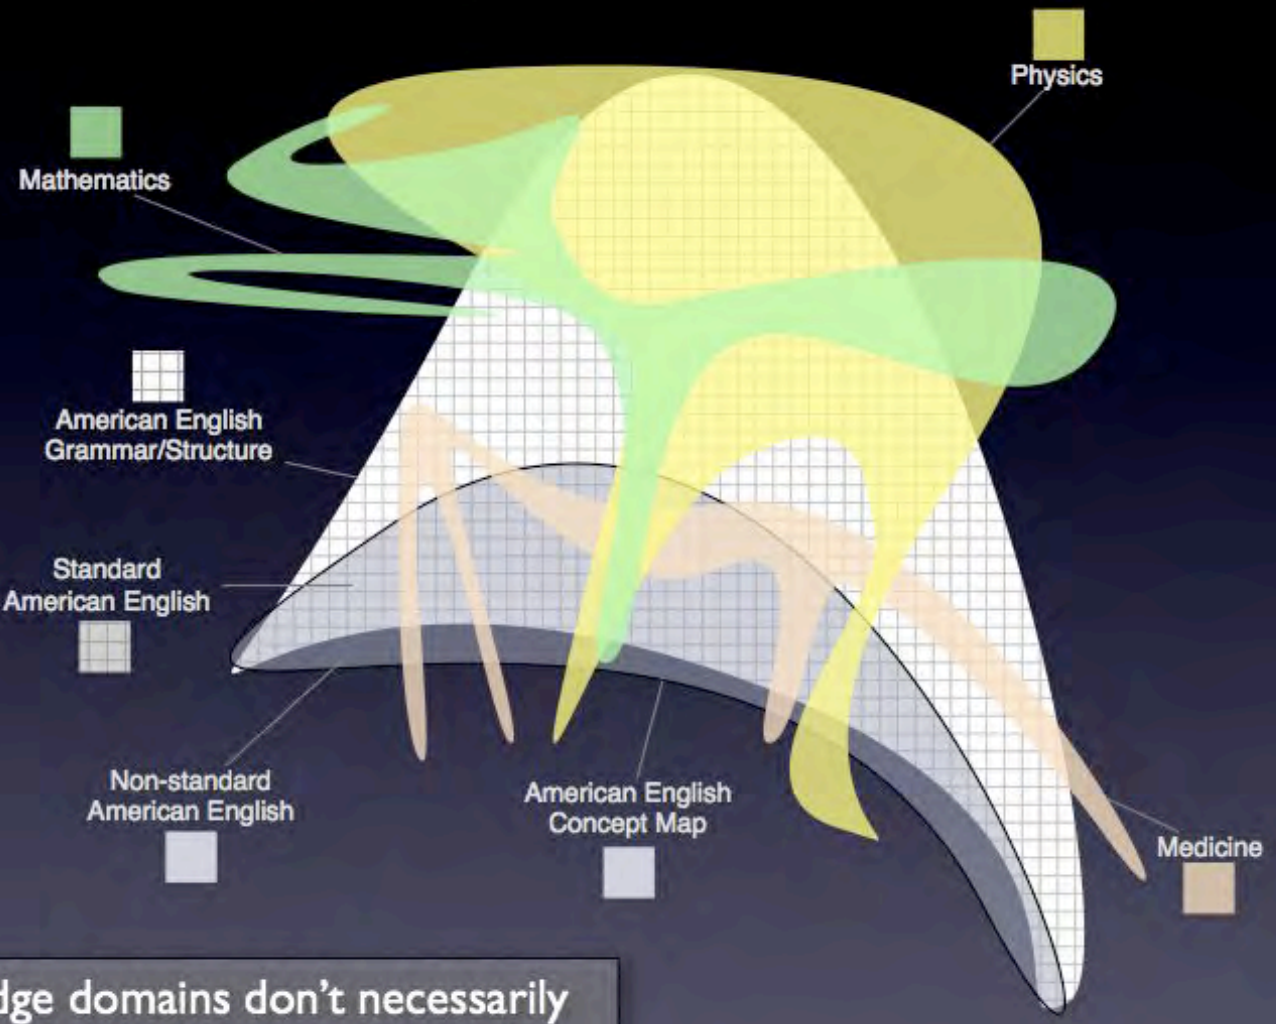
192.168.1.232
129.168.1.233

192.168.1.232
129.168.1.233

192.168.1.232
129.168.1.233

© 2007 Sunny Fugate

Language Domains



Cultures and knowledge domains don't necessarily use the same lexicon or even the same grammar!

How does the CND lexicon map to common language? Technical language? Military/tactical language?

© 2007 Sunny Fugate

IBM

Zurich Research Laboratory



Automating the configuration of flow monitoring probes

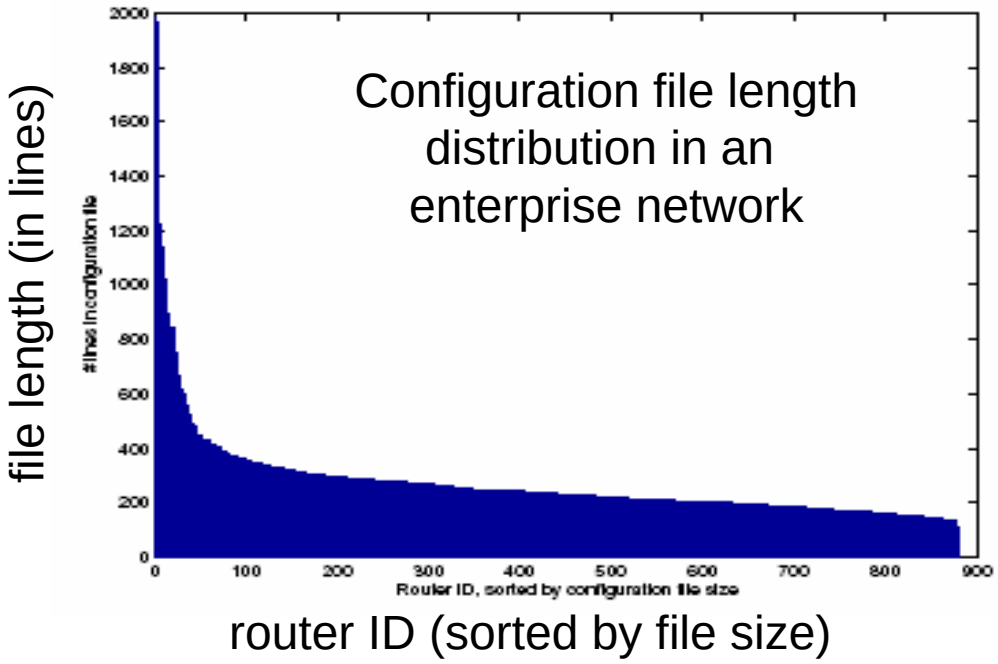
Xenofontas (Fontas) Dimitropoulos (xed@zurich.ibm.com)
Andreas Kind (ank@zurich.ibm.com)

Outline

- Background and motivation.
- Probe configuration architecture:
  - Requirements and goals.
  - Design.
  - Implementation.
- Future work and conclusions.

Network configuration

- Network elements are typically configured with low-level commands, e.g., Cisco IOS commands.
- Network administrators manage numerous network elements with lengthy configuration files.
- Network configuration is an error-prone and time-consuming process.
- Configuration errors can be costly, e.g.:
  - network outages
  - violations of SLAs

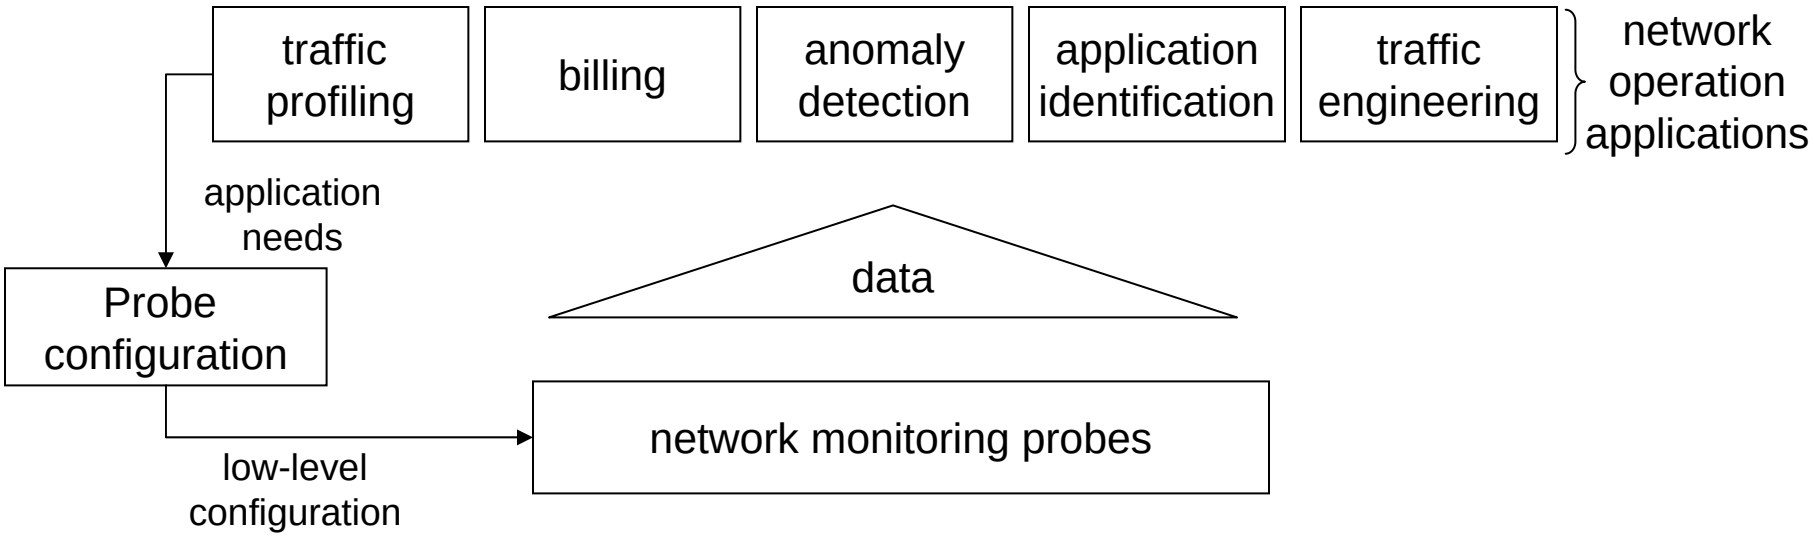


Source of figure: 100x100 project

Probe configuration

- The configuration of monitoring probes is part of the more general network configuration problem.
- Monitoring probes are gradually becoming more intelligent, for example, using advanced sampling and data aggregation techniques. Consequently, their configuration becomes more involved.
- Flexible Netflow (FNF) and IPFIX provide numerous configuration options that were not available earlier:
  - FNF has 58 different configuration commands.
  - FNF provides 65 different fields, arbitrary combinations of which can be used in the definition of flow key and non-key fields.
- Certain network operation applications need to dynamically change configuration to:
  - adapt to changing traffic conditions.
  - investigate on-going network anomalies.

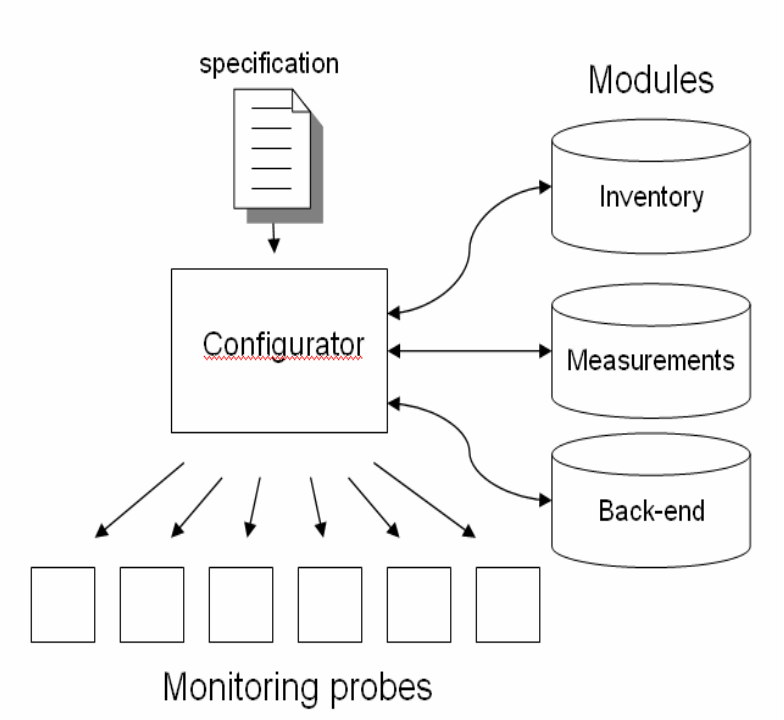
Configuration requirements



- Probe configuration should:
  1. take into account application needs.
  2. be aware of the available monitoring probes.
  3. generate low-level configuration commands.
  4. configure or update the configuration of probes.

Probe configuration architecture

- Three modules:
  - the measurements module describes different measurements, i.e., application needs.
  - the inventory module describes the monitoring probes of a network.
  - the back-end module provides necessary information for generating low-level commands.
- The specification identifies application needs.
- The configurator:
  - uses the modules and specification to generate low-level commands.
  - configures the probes

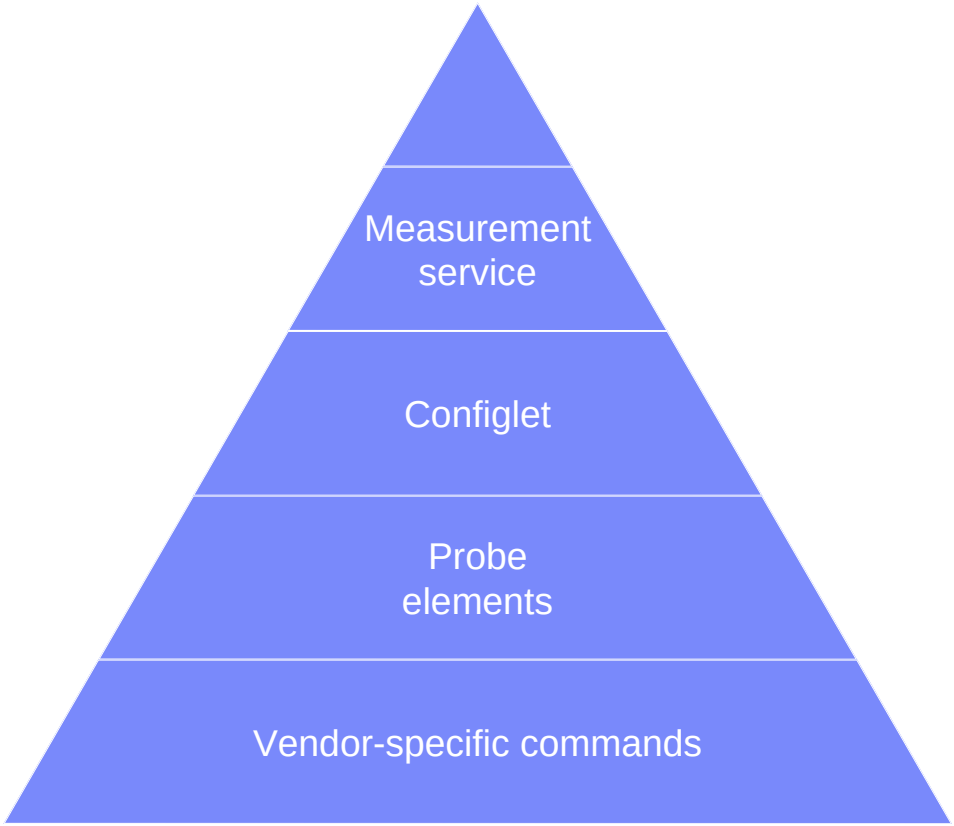


Design goals for simplifying configuration

1. Abstraction: hide low-level configuration commands.
2. Objective-oriented configuration expression:
   - express configuration in terms of measurement objectives.
   - focus on measurements instead of devices.
3. Network-wide configuration: configure a network instead of configuring individual devices.
4. Re-usability: make parts of configuration network-independent.
5. Extensibility: easily introduce support for new commands, measurements, etc.

Configuration abstraction hierarchy

- 1st level: vendor-specific configuration commands.
- 2nd level: probe elements (pe), i.e., logical components of a probe, like interface, flow cache, exporter.
- 3rd level: configlet, i.e., a set of specific probe elements that realizes a measurement.
- 4th level: measurement services, i.e., a configlet with certain probe selection rules.



Back-end module

- Specifies different probe elements.
- A probe element specification:
  - is written in XML.
  - has a unique id.
  - identifies parameters and parameter default values.
  - determines the low-level vendor-specific commands.

```
<!– Probe Element Exporter -->
<pe id=‘generic_exporter'>

 <params>
  <param id='port'>90</param>
  <param id='transport'>udp</param>
  <param id='destination'>192.0.0.1</param>
  <param id='label'>EXPORTER</param>
 </params>

 <template>
  <ios>
        flow exporter $label
        destination $destination
        transport $transport $port
  </ios>
  <yaf>
        --out $destination --ipfix $transport --ipfix-port $port
  </yaf>
  <junos>
  </junos>
 </template>

</pe>
```

Inventory module

- Specifies network probes, i.e., lists the characteristics that can be useful for their configuration.
- Besides describing location, system, and interface information, it declares tags that can be used for grouping probes and for probe selection.

```
<probe id='trabant.zurich.ibm.com'>
  <address>9.4.68.154</address>

  <location>
    <city>Zurich</city>
    <state>Central CH</state>
    <country>Switzerland</country>
  </location>

  <system>
    <os>ios</os>
    <version>12.4</version>
  </system>

  <interface id='FastEthernet0/0'>
    <capacity>100Mbits</capacity>
    <tag>internal</tag>
  </interface>

  <interface id='FastEthernet0/1'>
    <capacity>100Mbits</capacity>
    <tag>customer</tag>
  </interface>

  <tags>
    <tag>edge</tag>
  </tags>

 </probe>
```

Measurements module

```
<!--   Monitor how much traffic is send --
<!--   between IP blocks. -->
<msr id='traffic_matrix'>

  <params> <!-- Default pa
    <param id='collector_ad
    <param id='collector_po
    <param id='collector_tra
  </params>

  <!-- Probe element chain -
  <configlet>
  </configlet>

  <rules>
  </rules>

</msr>
```

```
<rules>
    <interface>
            if ( $interface.tag eq “external" and
                $probe.tag eq "edge" ) {
               return 1;
            } else {
               return 0;
            }
    </interface>
</rules>
```

```
<!-- Probe element chain -->
<configlet>
   <pe>
            <name>exporter</name>
            <params>
                                                  XPORTER</param>
                                                 >$collector_address</param>
                                                 ctor_port</param>
                                                 collector_transport</param>


                                                 me>

                                                 ACHE</param>
                                                 _DST_PREFIX_REC</param>
                                                 EXPORTER</param>
            </params>
   </pe>
   <pe>
            <name>interface</name>
            <params>
              <param id='monitor'>TM_CACHE</param>
              <param id='interface'>$interface->id</param>
              <param id='direction'>output</param>
            </params>
   </pe>
</configlet>
```

Input specification

- Lists the measurements and the probes in which to enable these measurements.
- Is the user interface and can be generated through a GUI.

```
<!-- Probes to apply measurements on -->
<probe id='wassen.zurich.ibm.com'></probe>
<probe id='trabant.zurich.ibm.com'></probe>


<!-- Measurements -->
<msr id='traffic_matrix'>
  <params> <!-- overwrite default values -->
    <param id='collector_address'>9.4.68.204</param>
    <param id='collector_port'>2055</param>
    <param id='collector_transport'>udp</param>
  </params>
</msr>

<msr id='app_monitoring'>
  <params> <!-- overwrite default values -->
    <param id='collector_address'>9.4.68.205</param>
    <param id='collector_port'>2055</param>
    <param id='collector_transport'>udp</param>
  </params>
</msr>
```

Design goals for simplifying configuration

1. Abstraction: hide low-level configuration commands.
2. Objective-oriented configuration expression:
   - express configuration in terms of measurement objectives.
   - focus on measurements instead of devices.
3. Network-wide configuration: configure a network instead of configuring individual devices.
4. Re-usability: make parts of configuration network-independent.
5. Extensibility: easily introduce support for new commands, measurements, etc.

Conclusions

- Described an architecture for simplifying the configuration of flow monitoring probes:
  - abstract configuration of probes and hide low-level details.
  - focus on measurement services that satisfy the objectives of applications.
  - generate and set configuration automatically.
- Future work:
  - Incorporate error-checking techniques.
  - Develop libraries for typical measurements.
  - Use NetConf.
  - Configuration optimization.

Privacy, Data Protection Law and Flow Data Anonymisation: requirements, issues, and challenges

Elisa Boschi, Hitachi Europe
Ralph Gramigna, KPMG

Acknowledgement: M. Bossardt (KPMG), D. Battisti (ETH Zurich)

Outline

- Review of law principles and requirements on data protection
  - European viewpoint
  - What is personal data?
  - Why is data protection law relevant for network monitoring?
  - Law principles overview
- The role of flow data anonymisation to support data protection
  - Discussion on its applicability and weaknesses
  - Suggestions for future steps

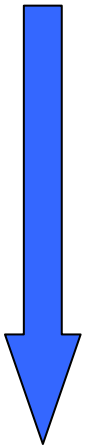
Data Protection Law: EU Directives

- Goal: protect the privacy of individuals
  - Not limited to information confidentiality
- EU Directives define the the minimum law requirements to be implemented by each EU member state
  - Applicable to international data transfers with EU
- Relevant to data protection:
  - Directive 1995/46/EC - on data protection
  - Directive 2002/58/EC - on privacy and electronic communications

Applicability and Personal Data

- Directive 95/46/EC applies to the

  „*processing* of *personal data*“





"*any information relating to an identified or identifiable natural person ('data subject'); an identifiable person is one who can be identified, directly* ***or indirectly****, in particular by reference to an identification number or to one or more factors specific to his ... identity*".

"any operation performed upon personal data, such as e.g. collection, storage, adaptation or alteration, consultation, disclosure by transmission, dissemination or otherwise making available, alignment or combination, erasure or destruction"

- Note: in some countries (e.g. Switzerland) this applies to „legal entities“ as well

Applicability to Network Monitoring

- *Indirect identification data comprise any information that may lead to identification of the data subject through association with other available information*
  - information available to the entity in charge of the data processing (ISP),
  - any information possessed by third parties
- IP addresses can identify someone “directly”
  - Esp. legal entities
- Many more attributes in a flow record can contribute to identifying someone “indirectly”

Principles: legitimation for processing

1. Consent
2. Data processing is „*necessary for the performance of a contract to which the data subject is a party*”
3. …

- Processing must be limited to specified purposes

- Further processing of data for historical, statistical or scientific purposes is possible provided that appropriate safeguards are provided
  - Left to national laws

Principles: Information of the Subject

The subject must be informed about:

1. Identity of the data controller
2. Purpose of the processing
3. Other information, e.g. the recipient of the data.

- It does not apply to scientific research, **IF** the provision of such information
  - proves impossible
  - would involve a disproportionate effort

- Appropriate safeguards must be provided
  - Their specification is let to national law

Border Crossing

- Transfer to third countries is generally possible if the third country ensures an adequate level of protection

http://ec.europa.eu/justice_home/fsj/privacy/thridcountries/index_en.htm

- E.g.
  - ✓ Switzerland, Canada, Argentina
  - ✗ USA (except Safe Harbor)

Traffic data and location data

- Introduced in Directive 2002/58/EC
  - *Traffic data: any data processed for the purpose of the conveyance of a communication or for the billing thereof*
  - *Location data: data indicating the geographic position of the terminal equipment of a user*

- Objectives:
  - Minimise the processing of personal data
  - Use anonymous or pseudonymous data where possible.

- „Anonymous“ = it is no longer possible to identify the data subject

Processing of Traffic and Location Data

- Traffic and location data relating to subscribers and users must be erased or made anonymous when no longer needed
- The processing of traffic data must be restricted
  - To persons acting under authority of providers
  - To certain activities (e.g. traffic management, fraud detection...)
- Location data can be processed only if
  - There is consent, or
  - Data is made anonymous

The Role of Flow Data Anonymisation to Support Data Protection

- The well known problem:
  - The more you anonymise the better privacy is protected...
  - ...but the less useful the data
- Anonymisation aims at removing sensitive information referring to an individual
- Attacks to anonymisation schemes have proved that those schemes could be broken allowing to "indirectly" identify people.
- Are known flow anonymisation techniques effective in protecting the privacy of individuals?

(4) Anonymization Techniques

Field to be anonymized:

IP address

IP	Truncation	Permutation	Black Marker	Prefix Preserving
135.98.111.17	135.98	141. 2. 32.37	10.1.1.1	22.131.88.67
135.98.111.128	135.98	41.12.96. 67	10.1.1.1	22.131.88.157
135.98.132.37	135.98	142.72.8.5	10.1.1.1	22.131.201.29
141.161.3.3	141.161	21.33.4.1	10.1.1.1	12.192.32.51
141.72.8.5	141.72	11.14.96.118	10.1.1.1	12.78.201.97
32.53.48.1	32.53	12.161.3.3	10.1.1.1	31.197.3.82

Some Anonymisation Attack Methods

- Data injection → injecting information to be logged with the purpose of later recognizing that data in the anonymized trace
- Fingerprinting → matching attributes of an anonymized object against those of a known object (e.g. web server) to discover a mapping between them
- Semantic attacks → system is exploited in a way that the victim thinks to do something, but is doing something different. The attacker may infer part of the unanonymized IP address by exploiting the semantics of prefix preserving.
- Structure recognition → recognizing structure between anonymized and unanonymized objects

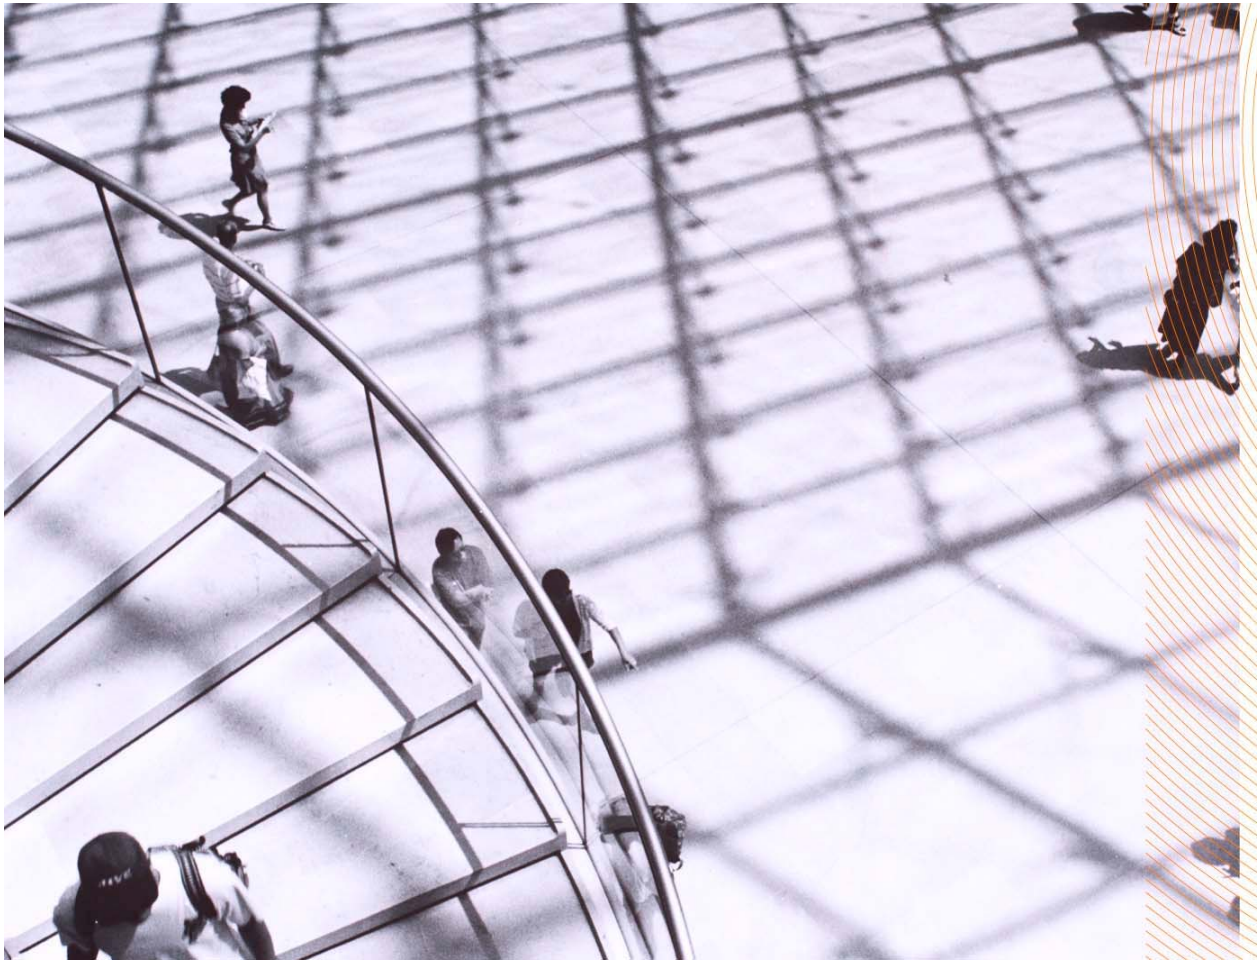
Attacks vs. Anonymisation Techniques

Attacks \ Anonymisation	Prefix-preserving	Cryptographic approach	Truncation	Permutation
Semantic attack	■	■		
Cryptographic attack		■		
Data Injection	■		■	■
Fingerprinting	■		■	■
Structure Recognition	■		■	■

- the attack can be used, (partial) results achieved

Conclusions

- We need to pay attention to data protection laws
- Anonymisation is part of the solution to protecting privacy, but
  - Research is still needed
  - This is not only a technical problem; a technical solution alone is not enough
- Legal solutions, policies, guidelines, interdisciplinary work are needed
- Anonymisation support is needed in standard flow data export protocols such as IPFIX



Automatic anomaly detection using NfSen

Wim Biemolt, SURFnet

Werner Schram, SURFnet



SURFnet network



Slightly less than twice the size of New Jersey



Automatic anomaly detection using NfSen

- SURFnet and netflow anomaly detection
  - NERD
  - NfSen
  - PeakFlow SP
- Currently used methods
  - DDos detection
  - Botnet detection
  - Holt-Winters aberrant behavior detection



SURFnet and netflow anomaly detection

- NERD v1
  - Developed by TNO
  - Based on cflowd
  - cflowd is no longer supported
- NERD v2
  - Initially developed by TNO
  - Has serious performance problems
  - NfSen can do the same but without the performance problems



NfSen

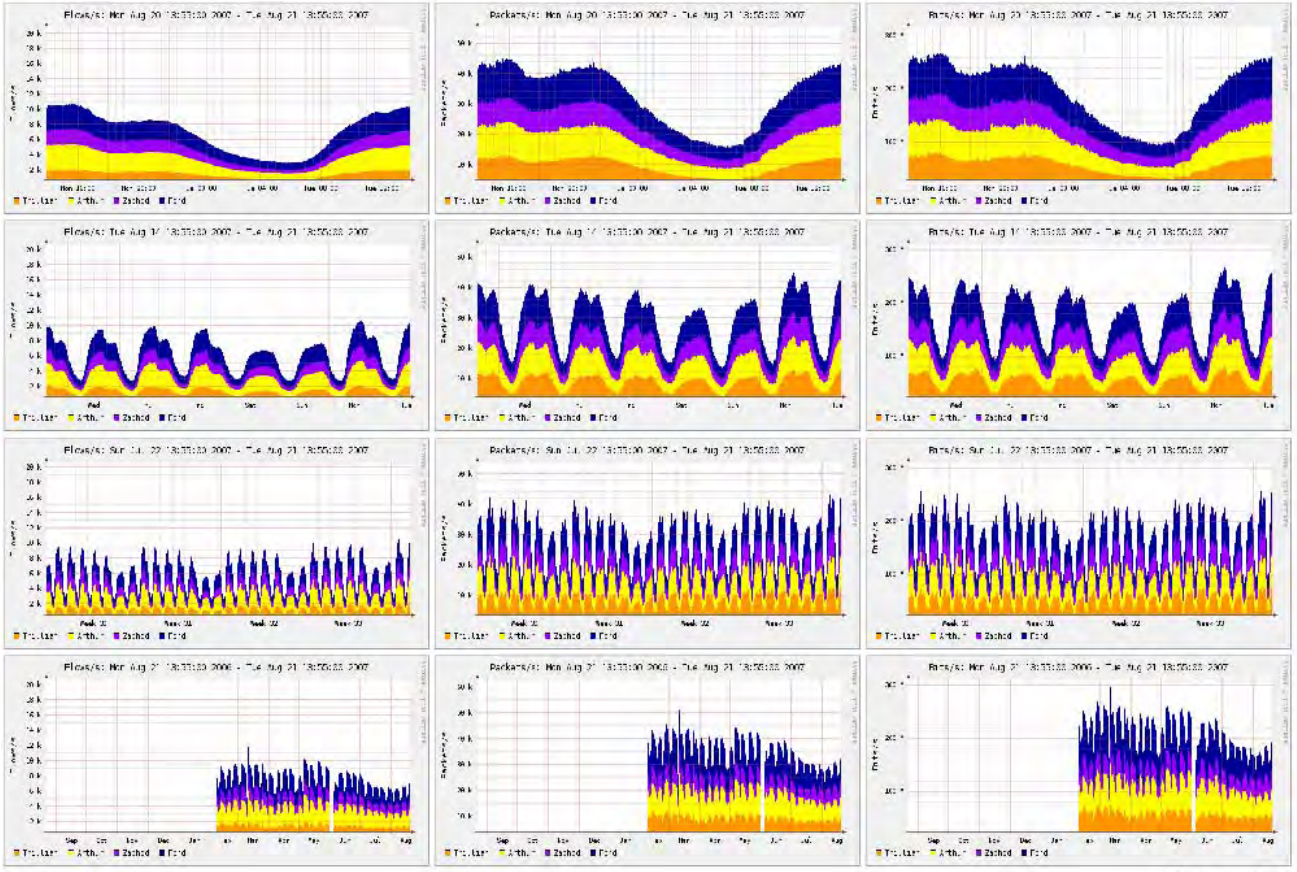
- Netflow Sensor (NfSen) is a network statistics tool
  - Developed by Peter Haag
  - Currently in active development
  - Alert plug-in system
  - Generic plug-in system
  - Some plug-ins already available



NfSen

Home | Graphs | Details | Alerts | Stats | Plugins | live | Bookmark URL | Profile: | live ▾

Overview Profile: live, Group: (nogroup)



nfsen snapshot-20070312



DDoS detection

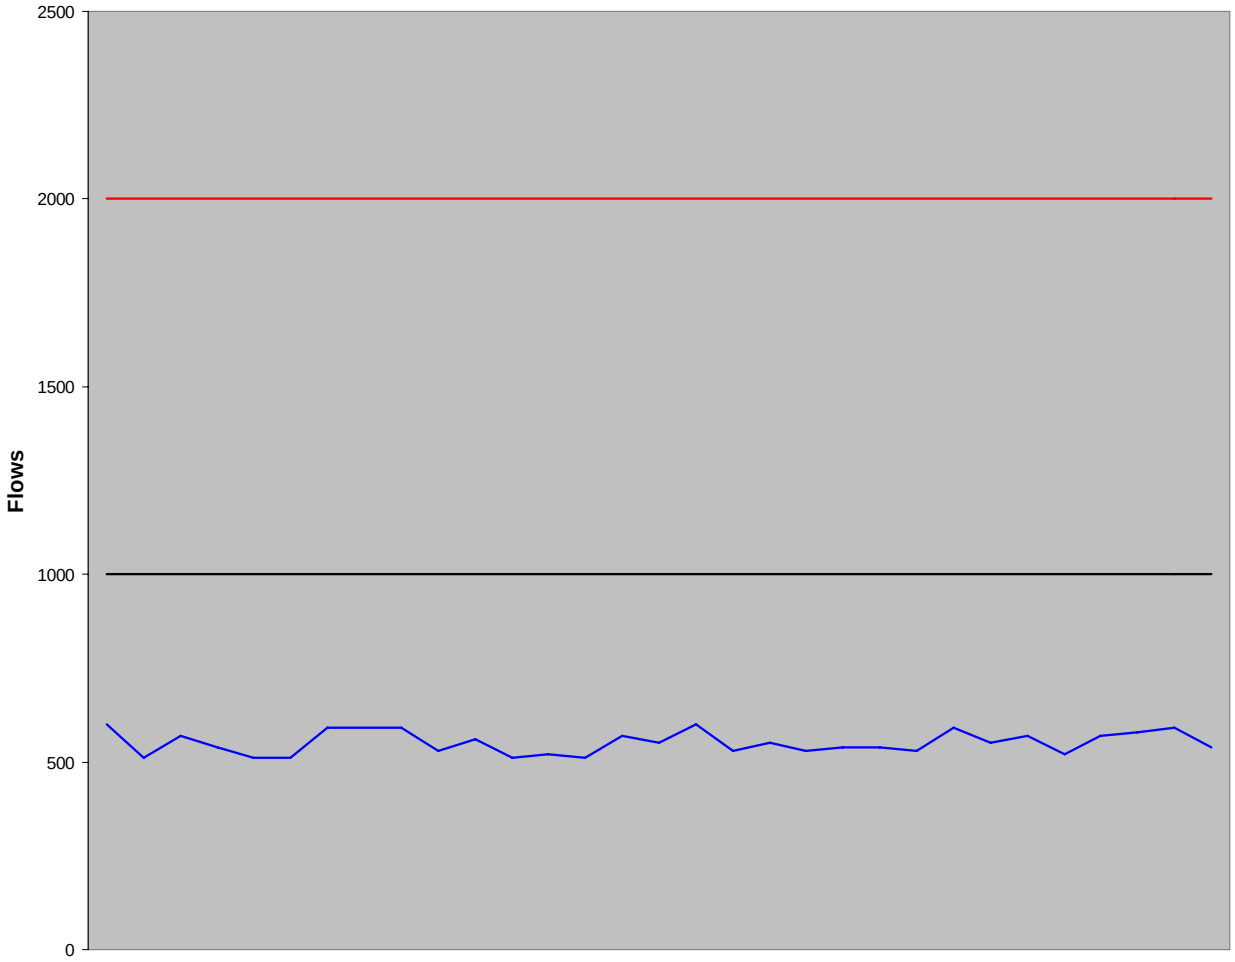


DDos detection

- Simple flow analysis based on NERD v1 DDos detection:
  - Low threshold
  - High threshold
  - Rules for traffic between those thresholds
  - Custom thresholds for high load services

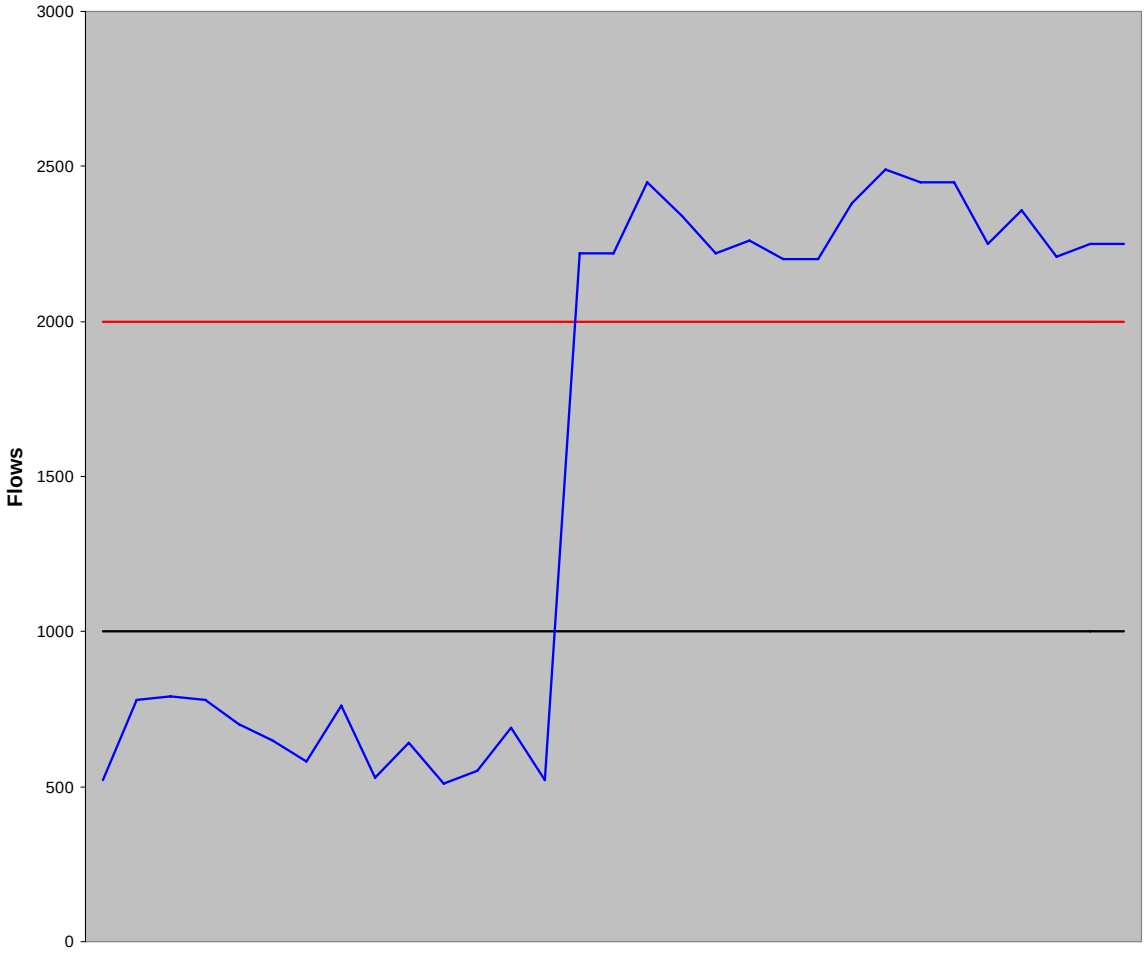


Expected traffic



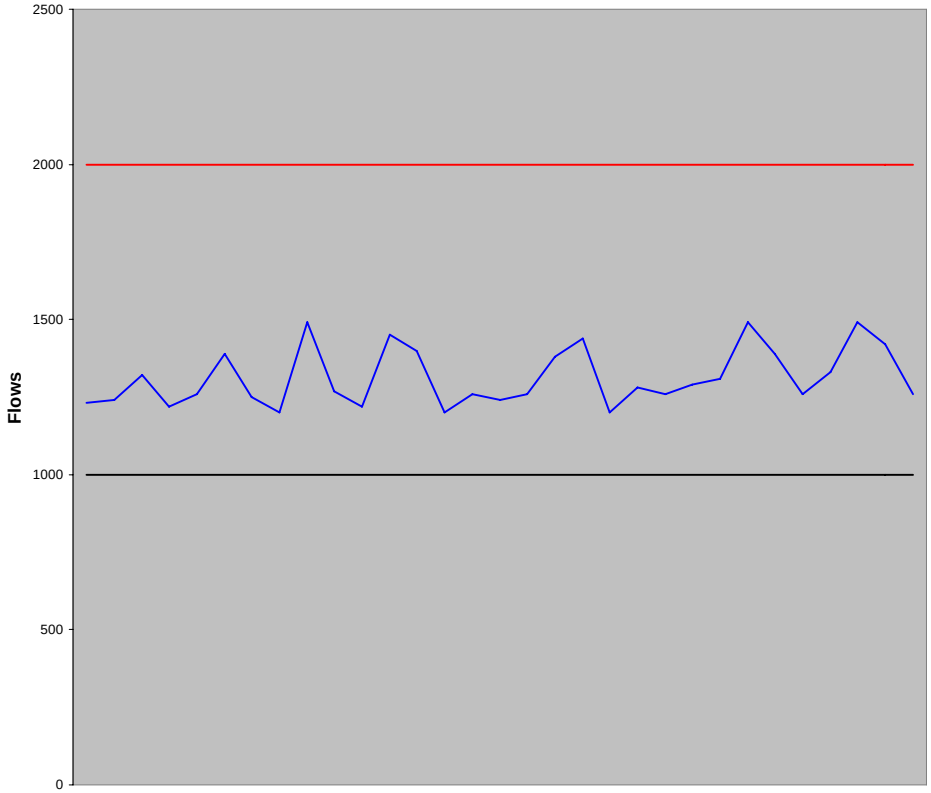


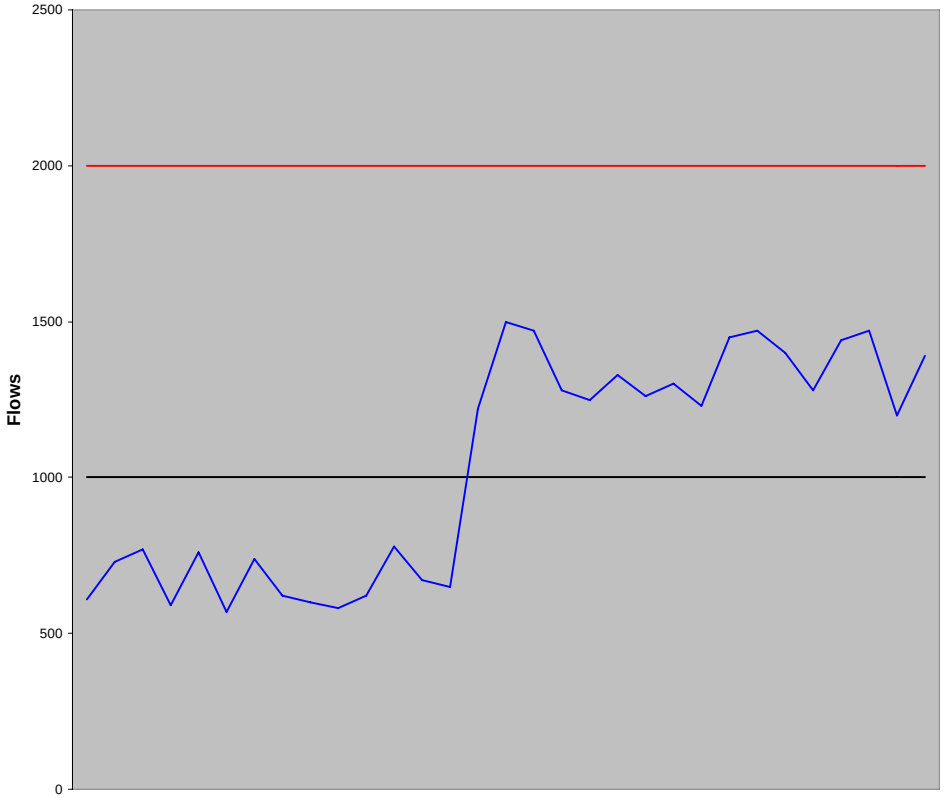
Definitively Suspicuous Traffic





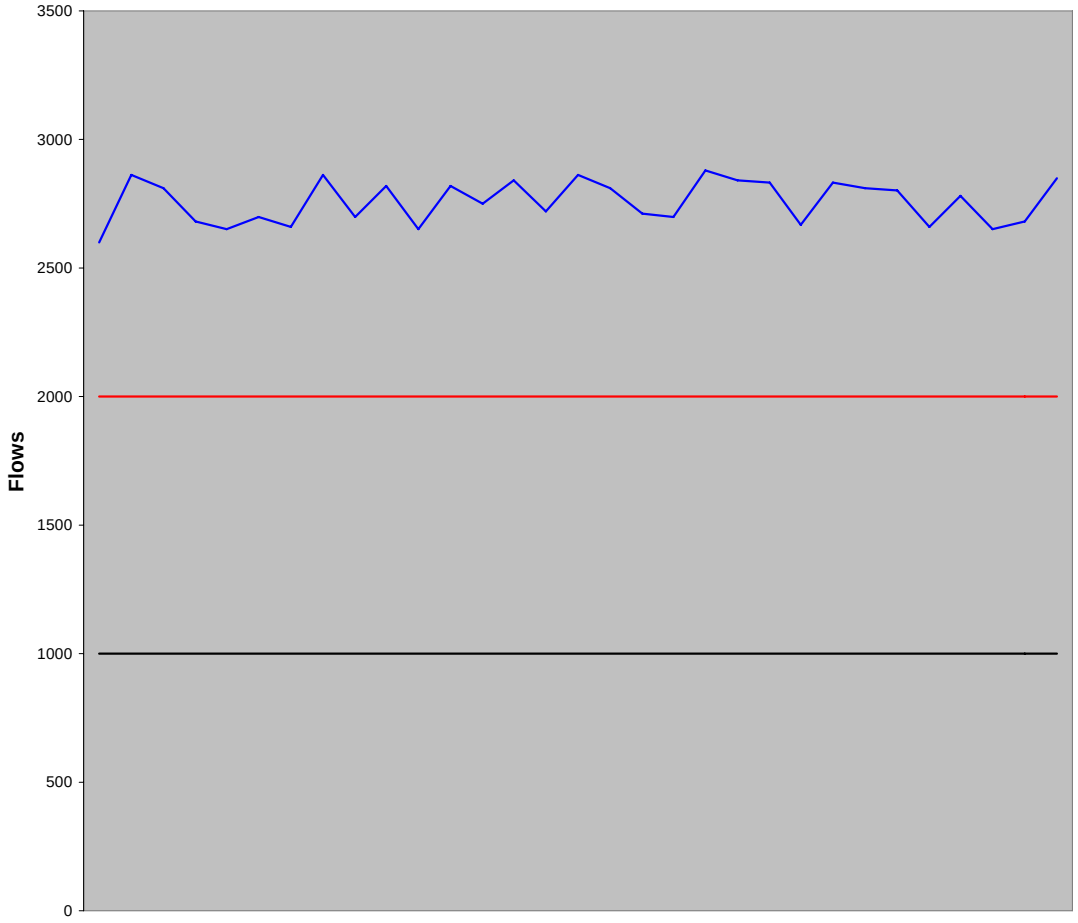
Border cases





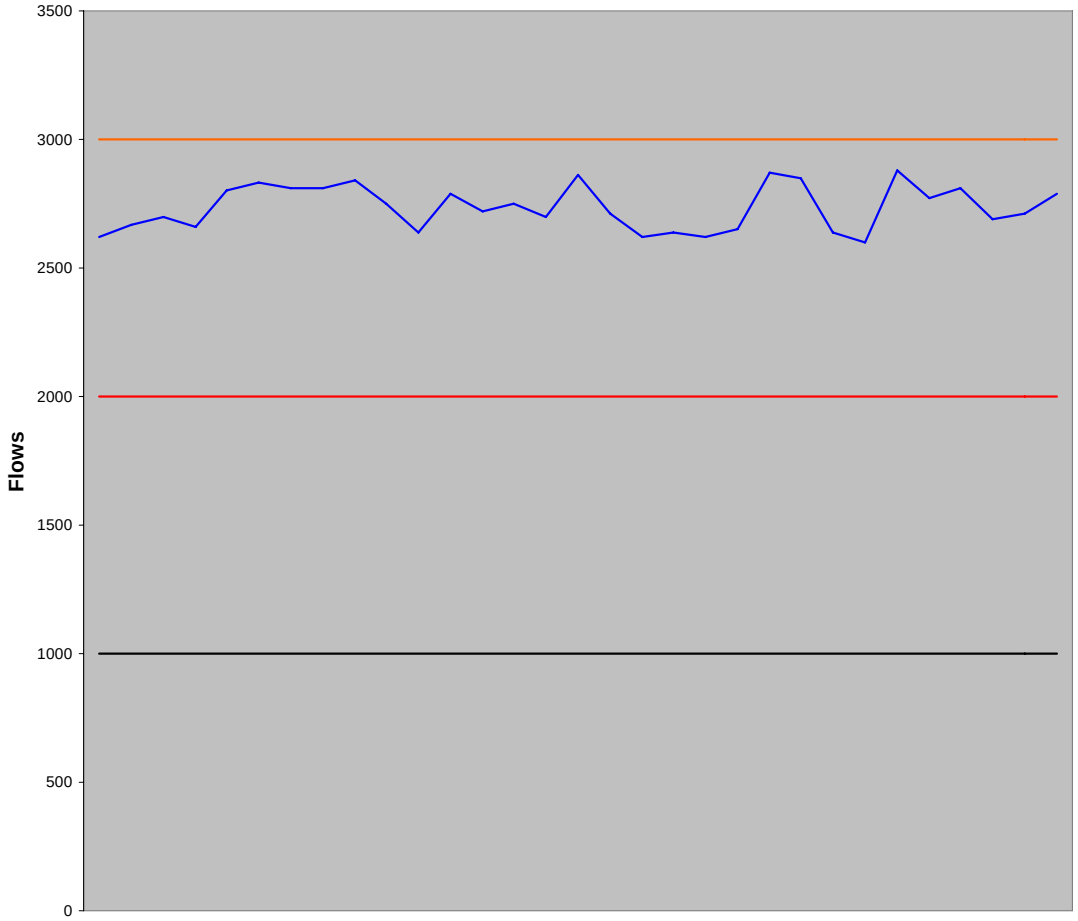


High load servers





Custom thresholds





DDos interface: report

Home | Graphs | Details | Alerts | Stats | Plugins | live | Bookmark URL | Profile: live ▼

alarm

report | setup | thresholds | botnets

number of alarms to show: 10 (0 for all)
from 7 days ago
up to 0 days ago
alarms: ddos
Ok

The ddos alarms between **2007-08-14** and **2007-08-22**

ID	Destination	Flows per 5 minutes	Average packets/flow	Average bytes/flow	Starttime	Stoptime	Active	
#33915	[redacted]	8125	3180	3	2007-08-21 08:44:59	2007-08-21 13:51:06	1	Delete
#33914	[redacted]	7674	4117	4	2007-08-21 08:39:59	2007-08-21 13:51:06	1	Delete
#33912	[redacted]	7725	117	1	2007-08-21 08:24:24	2007-08-21 13:51:06	1	Delete
#33911	[redacted]	7939	3194	3	2007-08-21 08:19:58	2007-08-21 13:51:06	1	Delete
#33910	[redacted]	10676	3647	3	2007-08-21 08:15:01	2007-08-21 13:51:06	1	Delete
#33909	[redacted]	6857	96	1	2007-08-21 08:10:01	2007-08-21 13:51:06	1	Delete
#33935	[redacted]	5399	126	1	2007-08-21 12:54:18	2007-08-21 13:46:05	1	Delete
#33926	[redacted]	5918	5259	4	2007-08-21 11:09:52	2007-08-21 13:46:05	1	Delete
#33922	[redacted]	5625	4977	4	2007-08-21 10:24:55	2007-08-21 13:46:05	1	Delete
#33934	[redacted]	5213	4173	3	2007-08-21 12:35:07	2007-08-21 13:31:04	1	Delete

Tuesday, 21-Aug-07 13:55:33 CEST
nfsen-alarm revision 86 (release 0.2.5b)
nfsen snapshot-20070312



DDos interface: Details

Home | Graphs | Details | Alerts | Stats | Plugins | live | Bookmark URL | Profile: live ▾

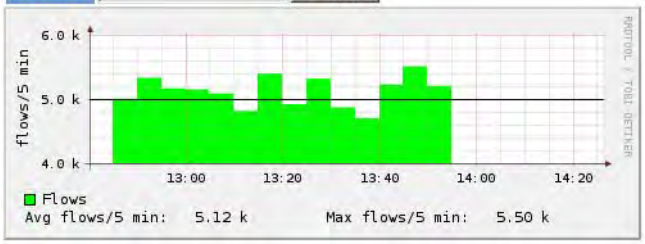
alarm

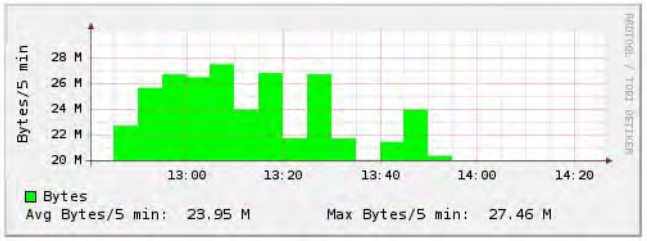
report | setup | thresholds | botnets | details: 33934

Remove tab

ip	[redacted]
hostname	[redacted]
first packet	2007-08-21 12:35
last packet	2007-08-21 13:56
threshold	Change

analyse





Top 10 flows per 5 minutes at 2007-08-21 13:56:05:

address	flows	bytes	port usage	last seen	actions
[redacted]	203	351310	min: 1042, max: 4990	2007-08-21 13:06:04	Report port scan / analyse
[redacted]	168	298160	min: 1064, max: 4647	2007-08-21 13:11:04	Report port scan / analyse
[redacted]	163	29627	min: 1103, max: 4990	2007-08-21 12:51:02	Report port scan / analyse
[redacted]	151	14326	min: 1033, max: 4925	2007-08-21 13:51:06	Report port scan / analyse
[redacted]	148	277170	min: 1032, max: 4992	2007-08-21 12:56:06	Report port scan / analyse



Botnet detection



Botnet detection

- Hosts infected by viruses connect to hosts known as botnet controllers
- List of botnet controllers are available, for example:
  http://www.bleedingthreats.net/rules/bleeding-botcc.rules
- Our plug-in logs all hosts that connect to known botnet controllers
- Automatically reports to incident report system using IODEF



Botnet IODEF reports

```
<?xml version="1.0" encoding="iso-8859-1"?>
<io:IODEF-Document xmlns:io="urn:ietf:params:xml:ns:iodef-1.0" lang="en">
  <io:Incident purpose="reporting">
    <io:IncidentID
    <io:StartTime>2
    <io:EndTime>200
    <io:ReportTime>
    <io:Assessment>
      <io:Impact ty
    </io:Assessment
    <io:Contact>
      <io:ContactNa
    </io:Contact>
    <io:EventData>
      <io:Method>
        <io:Referen
          <io:Refer
        </io:Refere
      </io:Method>
      <io:Flow>
        <io:System
          <io:Node>
            <io:Add
            <io:Cou
          </io:Node
        </io:System
        <io:System
          <io:Node>
            <io:Add
          </io:Node
          <io:Servi
            <io:Por
          </io:Serv
        </io:System
      </io:Flow>
    </io:EventData>
    <io:AdditionalD
NFSen</io:Additiona
  </io:Incident>
</io:IODEF-Document>
```

AIRT **IncidentdetailsSURFcert#019038**

Main menu | Import queue | Incidents | Search | Close current incident | Mail templates | Edit settings | Logout

(Bewerken) Externe identificatie:

(Bewerken) Ticket number(s):

Elementaire incidentgegevens

incidentsoort	infected
incidenttoestand	inspection requested
Incidentstatus	open
Datum van incident	20 aug 2007 17 02 03
Logboekinformatie	Source (ip) : 192.168.1.1 Target (ip:port) : 192.168.1.2 Packet (type:count): flow:23 Start time : 2007-08-13T15:07:47+02:00 End time : 2007-08-13T21:06:12+02:00

update

Beïnvloedde IP-adressen

IP adres	Machinenaam	Constituency	Rol in incident	Bewerken	Verwijder
192.168.1.1	infected.host	utwente.nl	Unknown	bewerken	verwijderen

IP adres Unknown Toevoegen

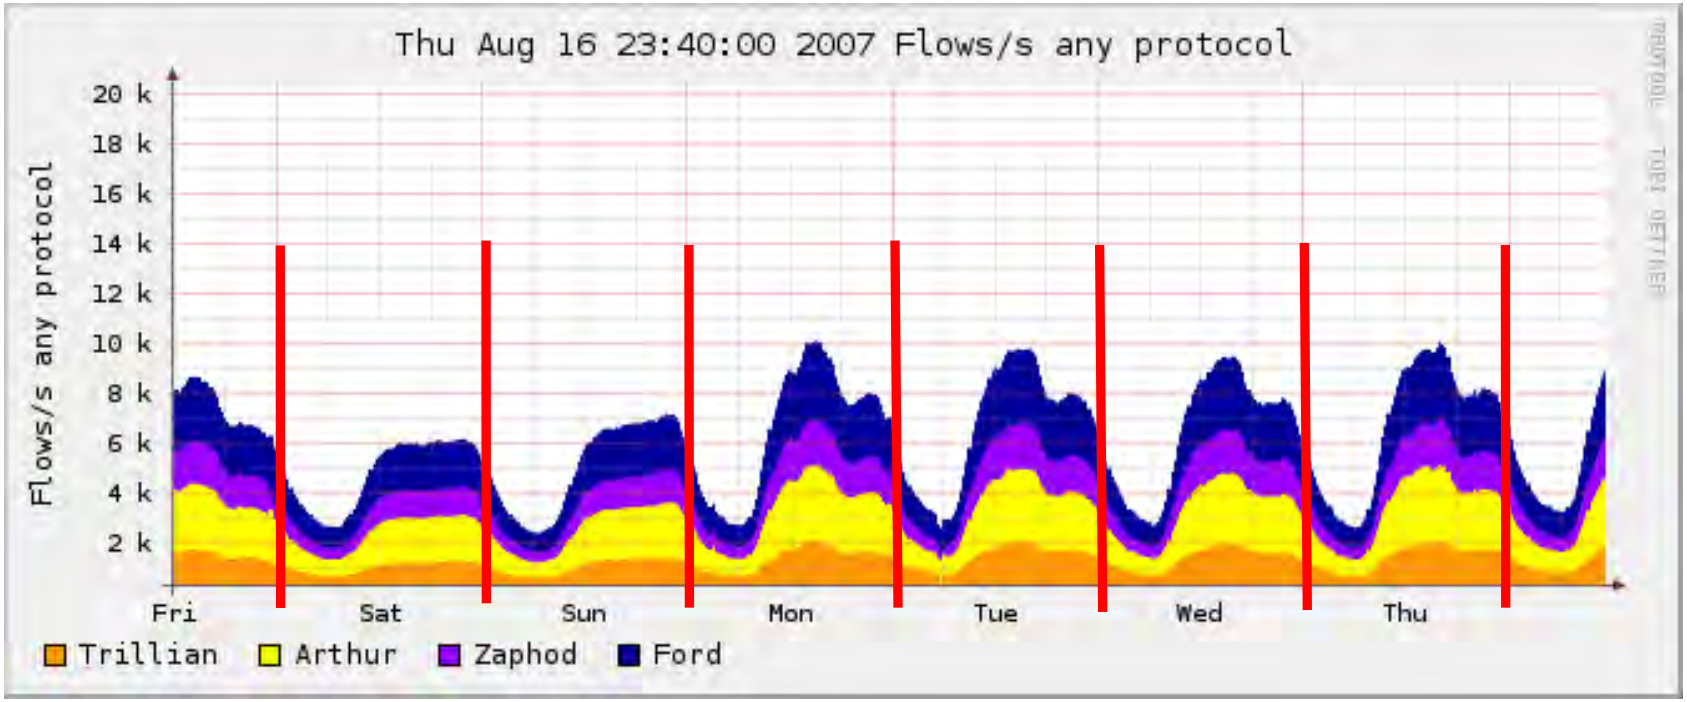


Holt-Winter abarrent behavior detection



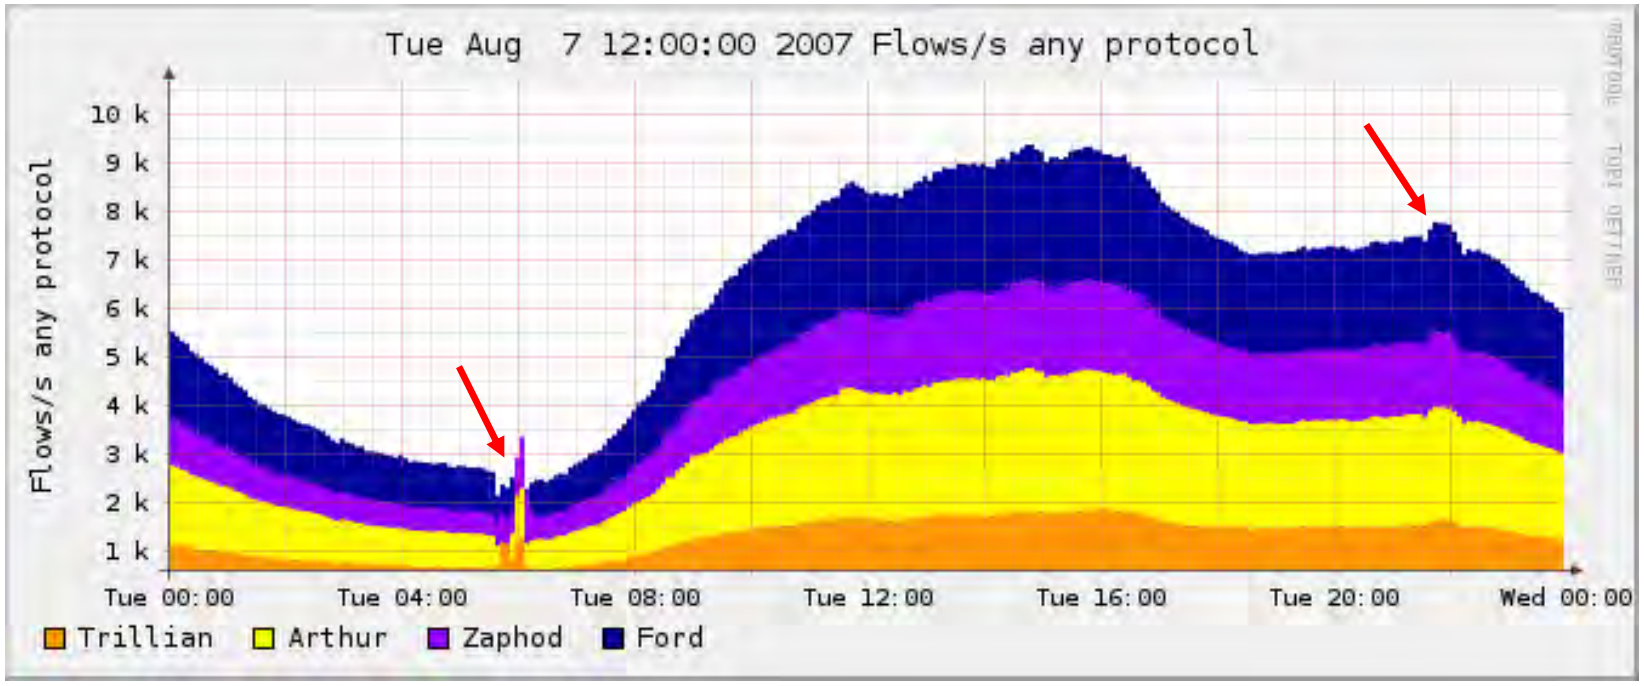
Holt-Winters aberrant behavior detection

- Uses information about periodic data to predict aberrant behavior.



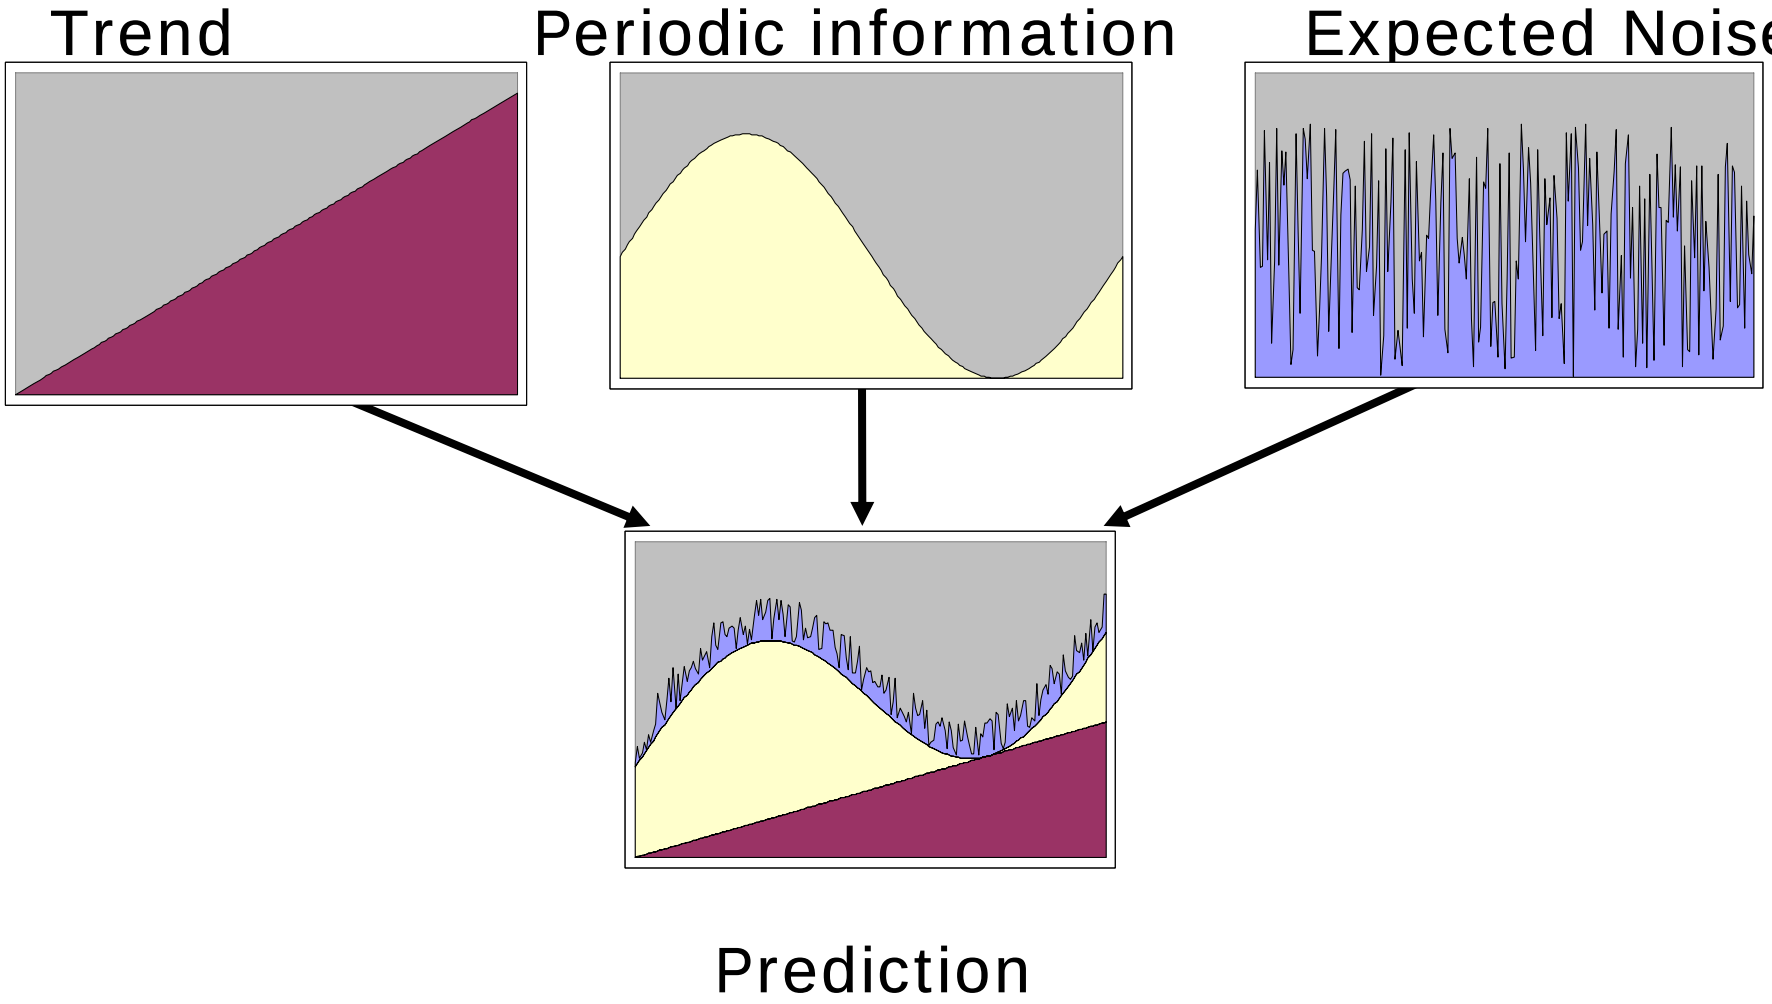


Holt-Winters: Example





Holt-Winters: Original implementation





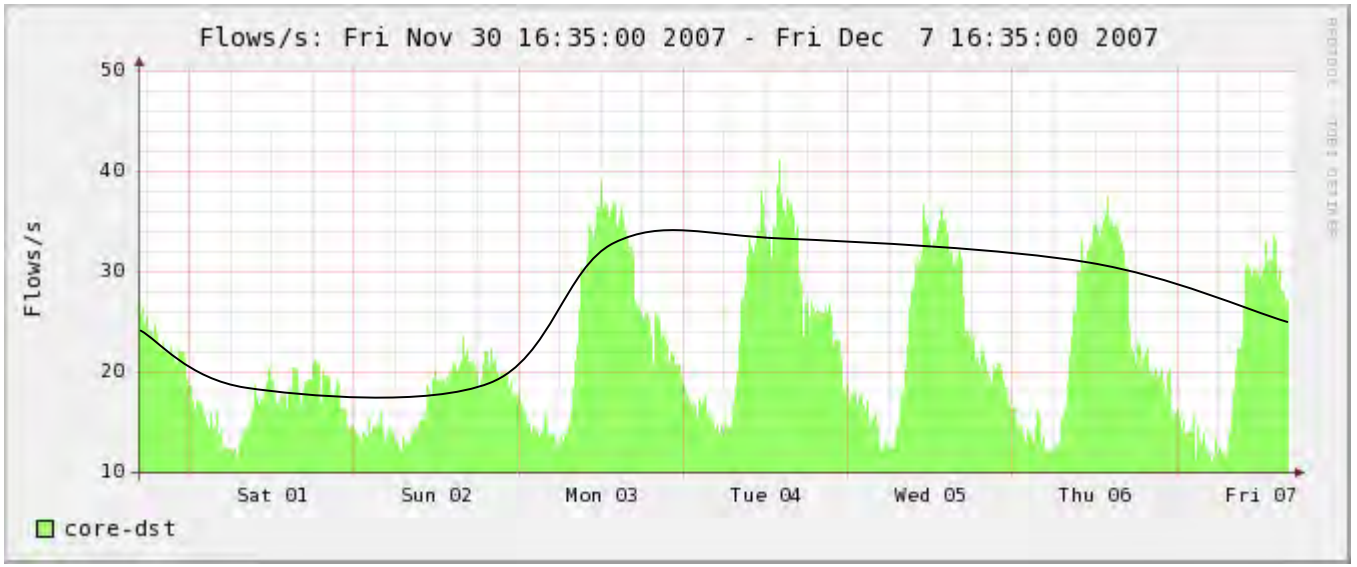
Limitations of the original implementation

- The original algorithm has three parameters:
  - One that defines the weight of historical data
  - One that defines the weight of the trend
  - One that defines the amount of expected noise
- The original algorithm has a constant learning rate
  - With a low learning rate, the selection of the initial values is critical. This will introduce false positives for a long time
  - With a high learning rate, the model will likely be overfitted. This will introduce false negatives
- The trend parameter has no significant influence with the resolution we are using



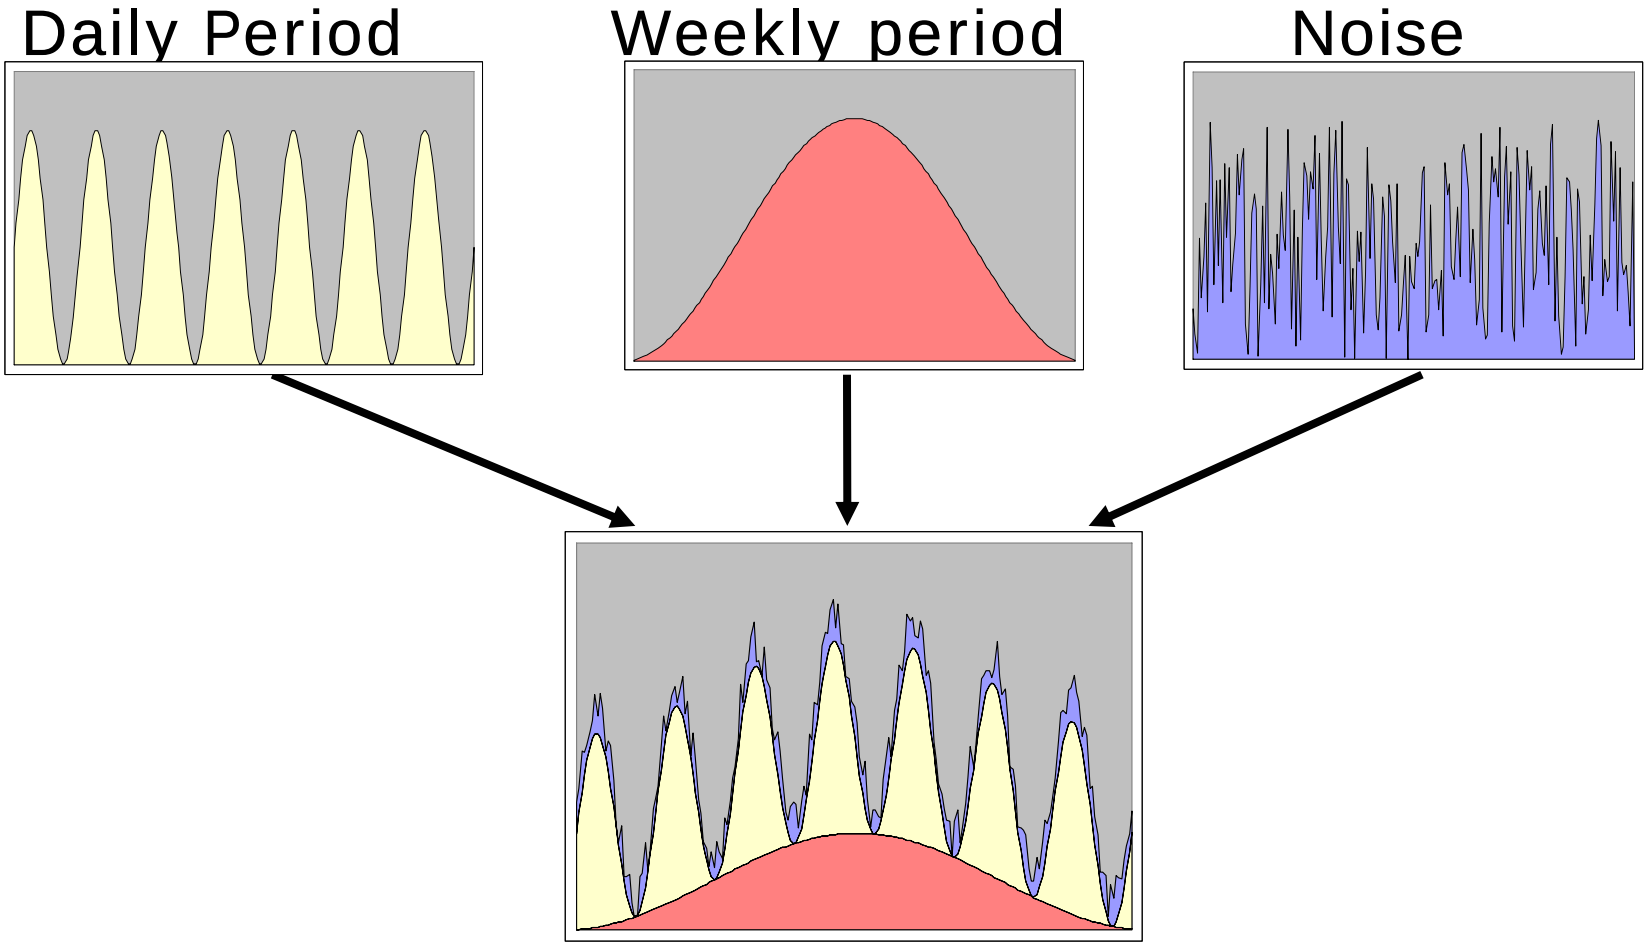
Holt-Winters: Multiple trends

Network traffic time series often show multiple recurring patterns



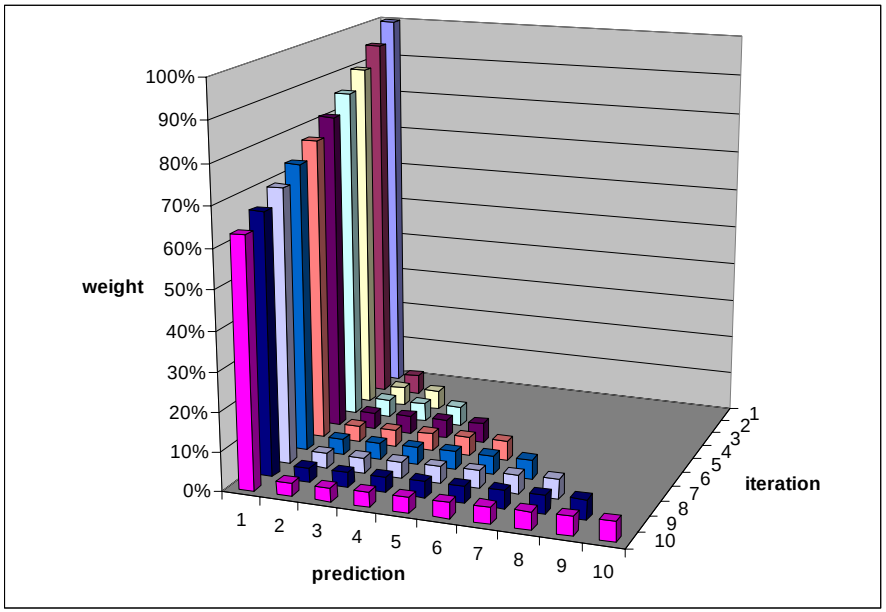


Holt-Winters: Multiple periods

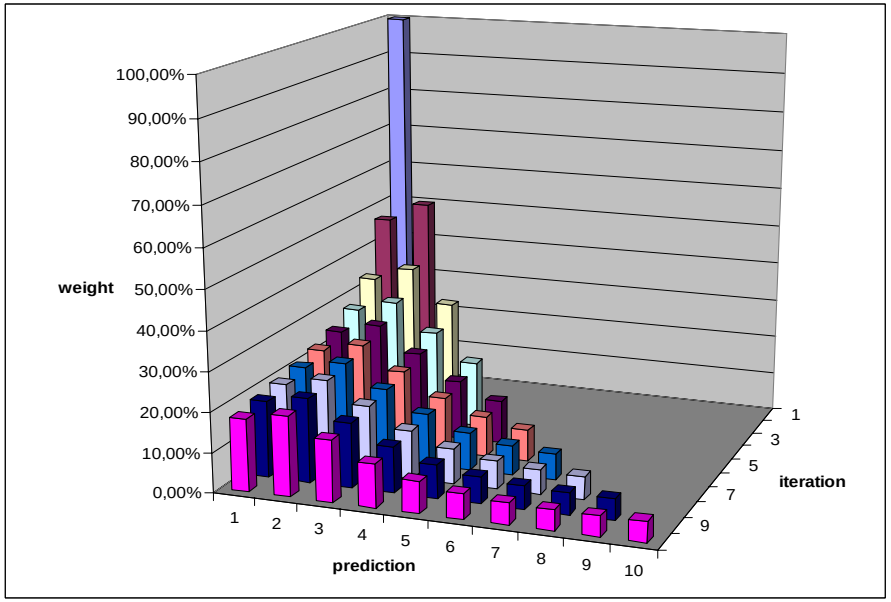




Learning rate



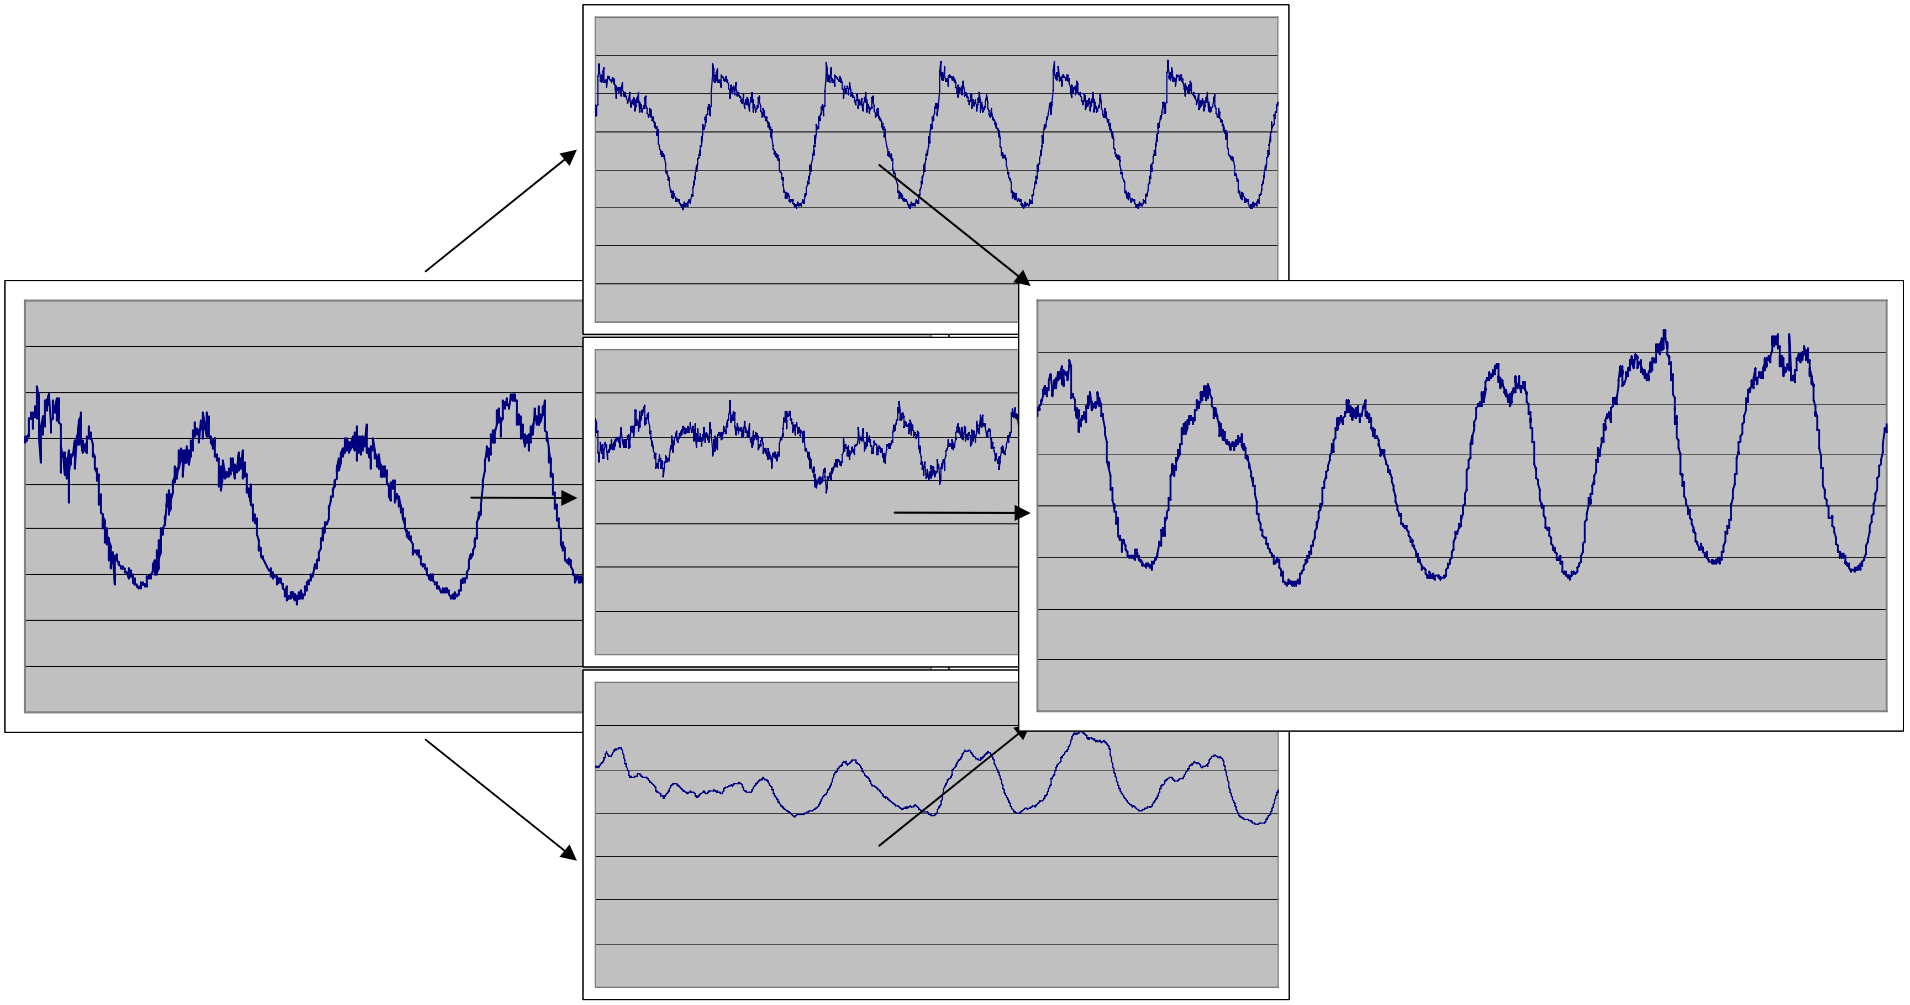
Fixed learning rate:
The first pattern is overweighted



Adaptive learning rate:
The weight of the first pattern is relative to the rest

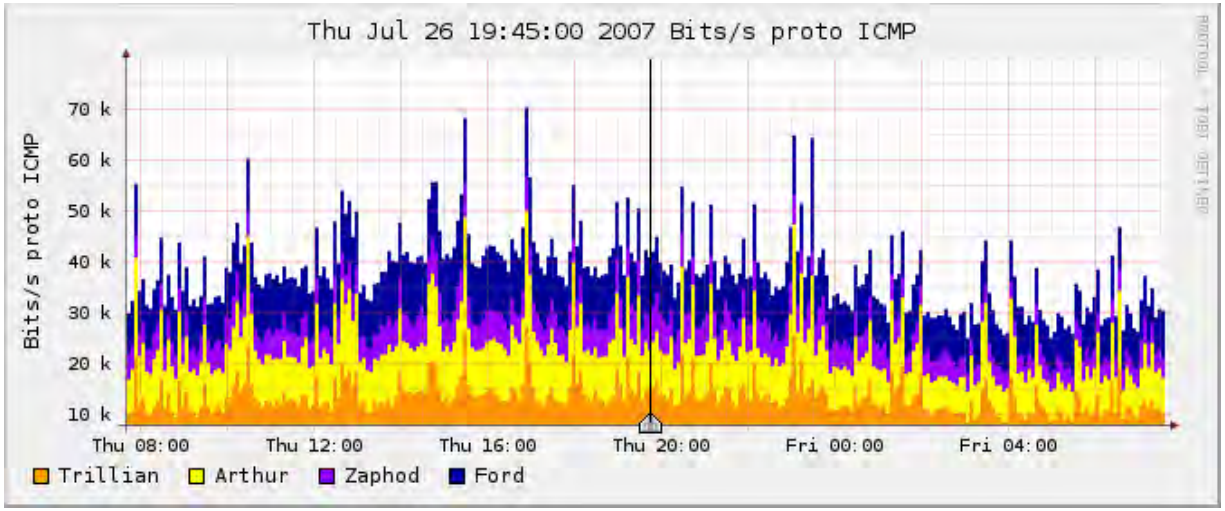


Real data example



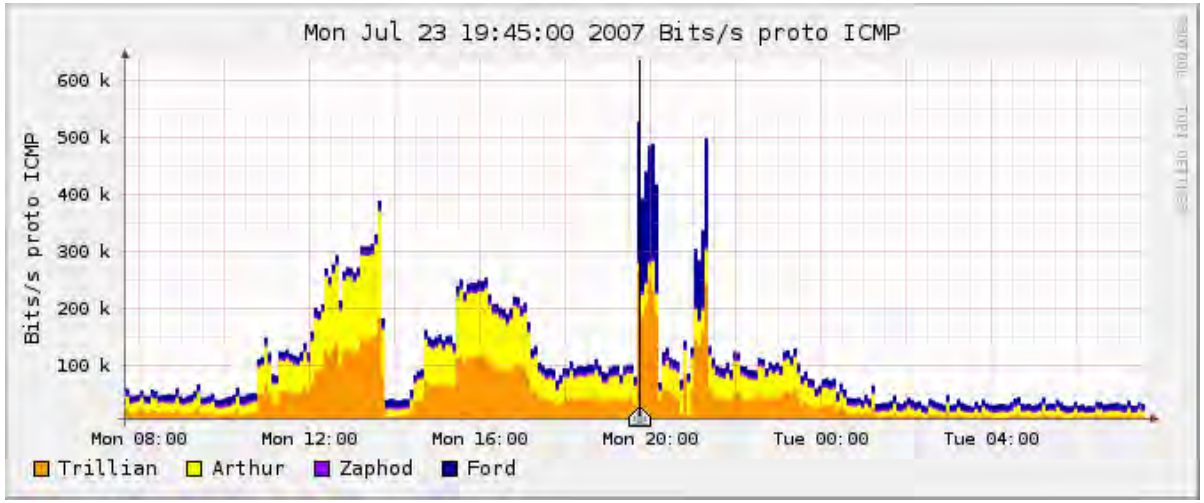


Holt Winters: Usage Example



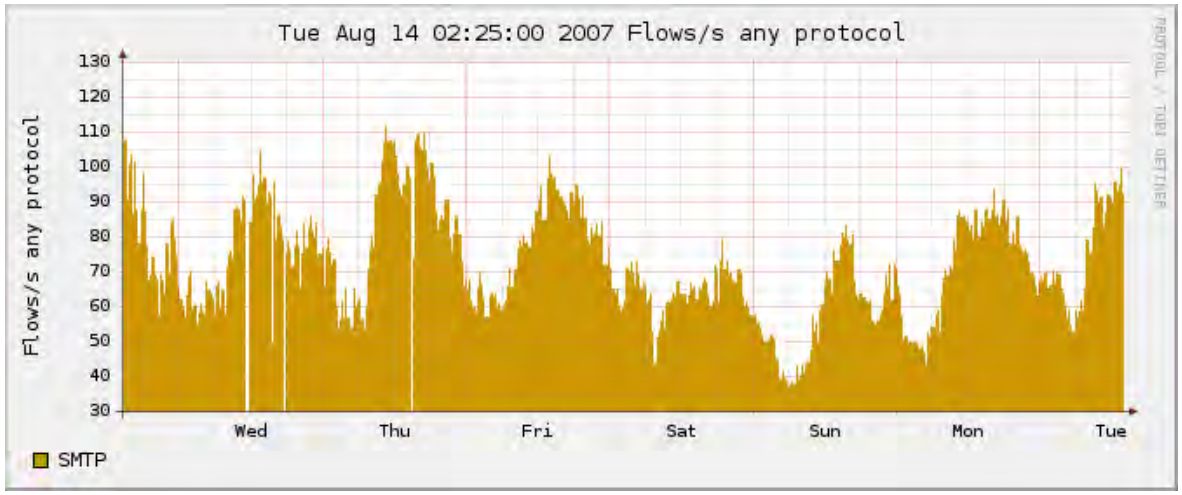
Normal ICMP Traffic

Aberrant ICMP Traffic:
Caused by DDos attack
by Stormworm
botnet



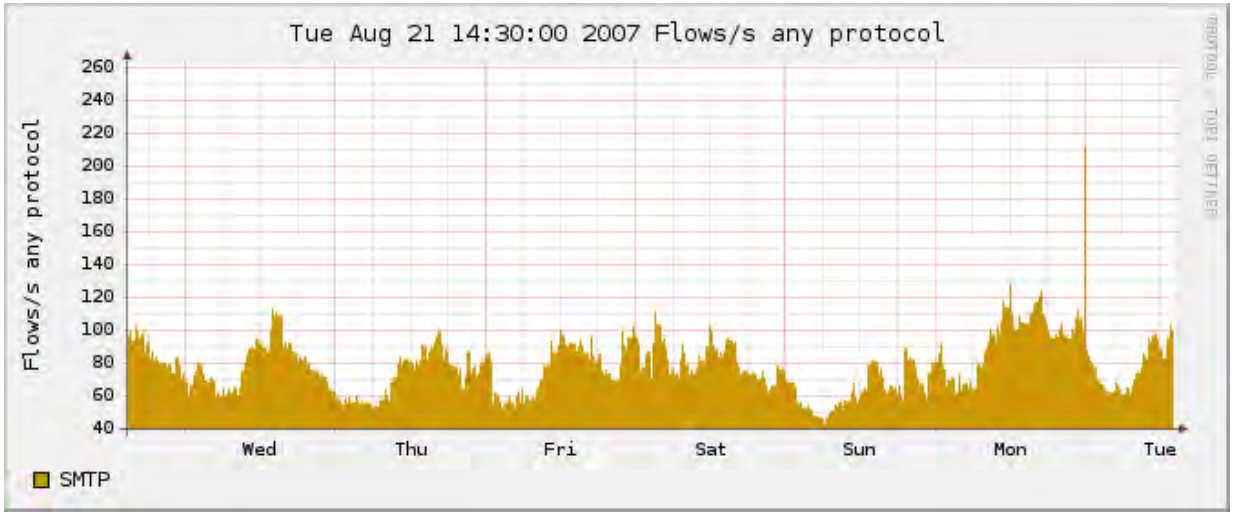


Holt Winters: Other possible uses



Common SMTP Traffic

Last week SMTP Traffic





Conclusions

- Simple netflow analysis can be a great added value to network security for ISP's
- Holt-Winters analysis has some quirks:
  - Parameters are hard to understand and hard to choose
  - The trend parameter (mostly) has no significant value in network analysis
  - Ignoring other patterns decreases accuracy
- But an addapted version can improve this:
  - A flexible learning rate can simplify the selection of the parameters
  - The trend parameter can be removed
  - We can look at more than 1 pattern



Availability

- DDoS plugin:
  - http://sourceforge.net/projects/nfsen-overflow/
- Botnet and Holt-Winters plugins:
  - In development, but contact me if you want to try it:
  - werner.schram@surfnet.nl

Wim Biemolt
Wim.Biemolt@surfnet.nl
www.surfnet.nl

Werner Schram
Werner.Schram@surfnet.nl
www.surfnet.nl

REDJACK

AMP-Based Flow Collection

Greg Virgin - RedJack

REDJACK

AMP- Based Flow Collection

- AMP - “Analytic Metadata Producer”: Patented US Government flow / metadata producer
- AMP generates data including
  - Flows
  - Host metadata (TCP stack information, software banners)
  - Metrics
- Purpose of this talk: To discuss the flow data collection implications of these additional data types for forensic analysis (not just correlation and alerting)
  - Additional data sources
  - Analysis scenarios
  - Collection schemes

REDJACK

Additional Data Sources

- Core data source: flow data
  - Netflow-like data with additional TCP flag information
- Flow-derived data sources: port details
  - Ports accepting connections
  - Bandwidth statistics
- Additional data sources (Not appropriate for flow records- aggregated data sources by IP, not communication)
  - TCP Stack information reflecting running O/S
  - Server Banners (as seen by the Internet)
  - Client Banners (as sent to the Internet)
  - DNS Names collected from both the DNS protocol and other protocols (NEVER trust DNS!)
  - Search strings from search engines (HTTP “referer” tags)

REDJACK

Scenario 1: Server “Importance”

- Server Profile
  - Configuration (“Windows 2000”)
  - List of listening ports (80, 443)
  - List of available services (“IIS/6”)
  - Domain name(s) (“www.golfcarts.com”)
  - Traffic Volume (X connections today, per week, per month)
  - Associated search strings (“golf carts”, “high performance golf carts”)
- Why?
  - Provides metrics to automatically partition servers by volume, type, vulnerability
  - Provides forensic value through server details often unavailable at time of analysis
- Flow analysis scenarios:
  - Which active servers were impacted by flow traffic / scans / attacks
  - Scrutinize payload-bearing traffic going to these servers
  - Make sure you’re not picking up potentially “normal” activity in other anomaly detection approaches (your concept of normal doesn’t necessarily have to be perfect)
  - Assign real world concepts to traffic activity and perform sanity checks through search strings

REDJACK

Scenario 2: DNS / Name Analysis

- Naming Information:
  - DNS Response packets
  - HTTP Get requests, mail protocol name announcements
- Why?
  - The current DNS implementation presents major risks because threats can masquerade as well known sites
  - The web protocol is dominated by virtual servers
  - We have found interesting discrepancies between DNS and naming in other protocols
  - Dealing with hosts as domain names is more natural (the purpose of the protocol)
- Flow analysis scenarios:
  - Name-based queries (possible with SiLK)
  - Names or name checksums incorporated into flow records for web traffic, followed by correlation with a name for the IP once the data is collected (helps with virtual servers)
  - Forensic analysis of traffic to or from bogus domain names to determine potential damage (but you have to do the above correlation first)

REDJACK

Scenario 3: Making IP Space Heterogeneous

- Required data:
  - Host Configuration
  - listening ports
  - running services
- Why?
  - Too often IP space is considered one big homogeneous blob - analysis is done on traffic between nodes without considering types of nodes
  - The diagnosis of activities such as worms can be made from hosts in a set running the same piece of software rather than signature
- Flow analysis scenarios:
  - What has been called a “similarity” analysis: take an IP set and run it against host profiles to provide statistics on what the hosts in the set have in common
  - Flow analysis broken down by host attributes isn’t very common, so there are a number of possibilities

REDJACK

Scenario 4: The “Alternate Use” Flag

- Marking flows for statistically significant attributes is marking flows based on signatures, not necessarily “new” data
- “Alternate Use” refers to the proper use of an Internet protocol without being used for the *purpose* of the protocol (this is not protocol analysis)
- Why?
  - This type of traffic can be a huge portion of the traffic
  - Of unique DNS names seen by your network, more than half of them may come from just a handful of sources
- Flow analysis scenarios:
  - Often port and protocol numbers are considered synonymous with legitimate use of protocols; this can be used to filter out alternate uses
  - Most of the “alternate” uses for DNS appear to be spam reporting, that information could be harvested

REDJACK

Scenario 5: IDS Verification

- Use host information or flow data to validate IDS records
  - If hosts aren't running the software that IDS signatures think they are…
- Not a new concept and done in practice

REDJACK

Summary of Scenarios

- New data sources can be used with flow data to:
  - Add contextual information and increase situational awareness
  - Create filters that could be useful for both queries and data collection
  - Partition data into bins or streams with more (or less) analytic meaning
- The best result is for these techniques to impact the data or be recorded as additional data
- This has an obvious impact on collection infrastructure
  - Data production software should be able to mark, reformat, or drop flow data based on this information
  - Data collection and storage software should be able to process or partition this information
  - Since most of these techniques don't amount to much more than a filter definition, a registry for these filters that different parts of the flow collection infrastructure can use is appropriate

REDJACK

New Sensor Attributes

- (This is in addition to flows with TCP options, host information, and DNS)
- Filters based on additional information
- Domain name value for the web protocol
- “Alternate Use” flag
- Not yet discussed:
  - Change ICMP to include third IP address in some instances

REDJACK

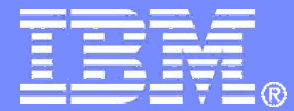
New Data Collection Attributes

- Marking or partitioning flows with domain names
- Metrics, filtering, and additional aggregation (flows for large servers can be compacted)

REDJACK

New Data Store Attributes

- Flow data closely tied to new data sources
- Registry for filtering techniques that can be leveraged by the sensor and collection
- Questions?
  - Greg Virgin, greg.virgin@redjack.com



Zurich Research Laboratory

Dynamic Adaptation of Flow Information Granularity for Incident Analysis

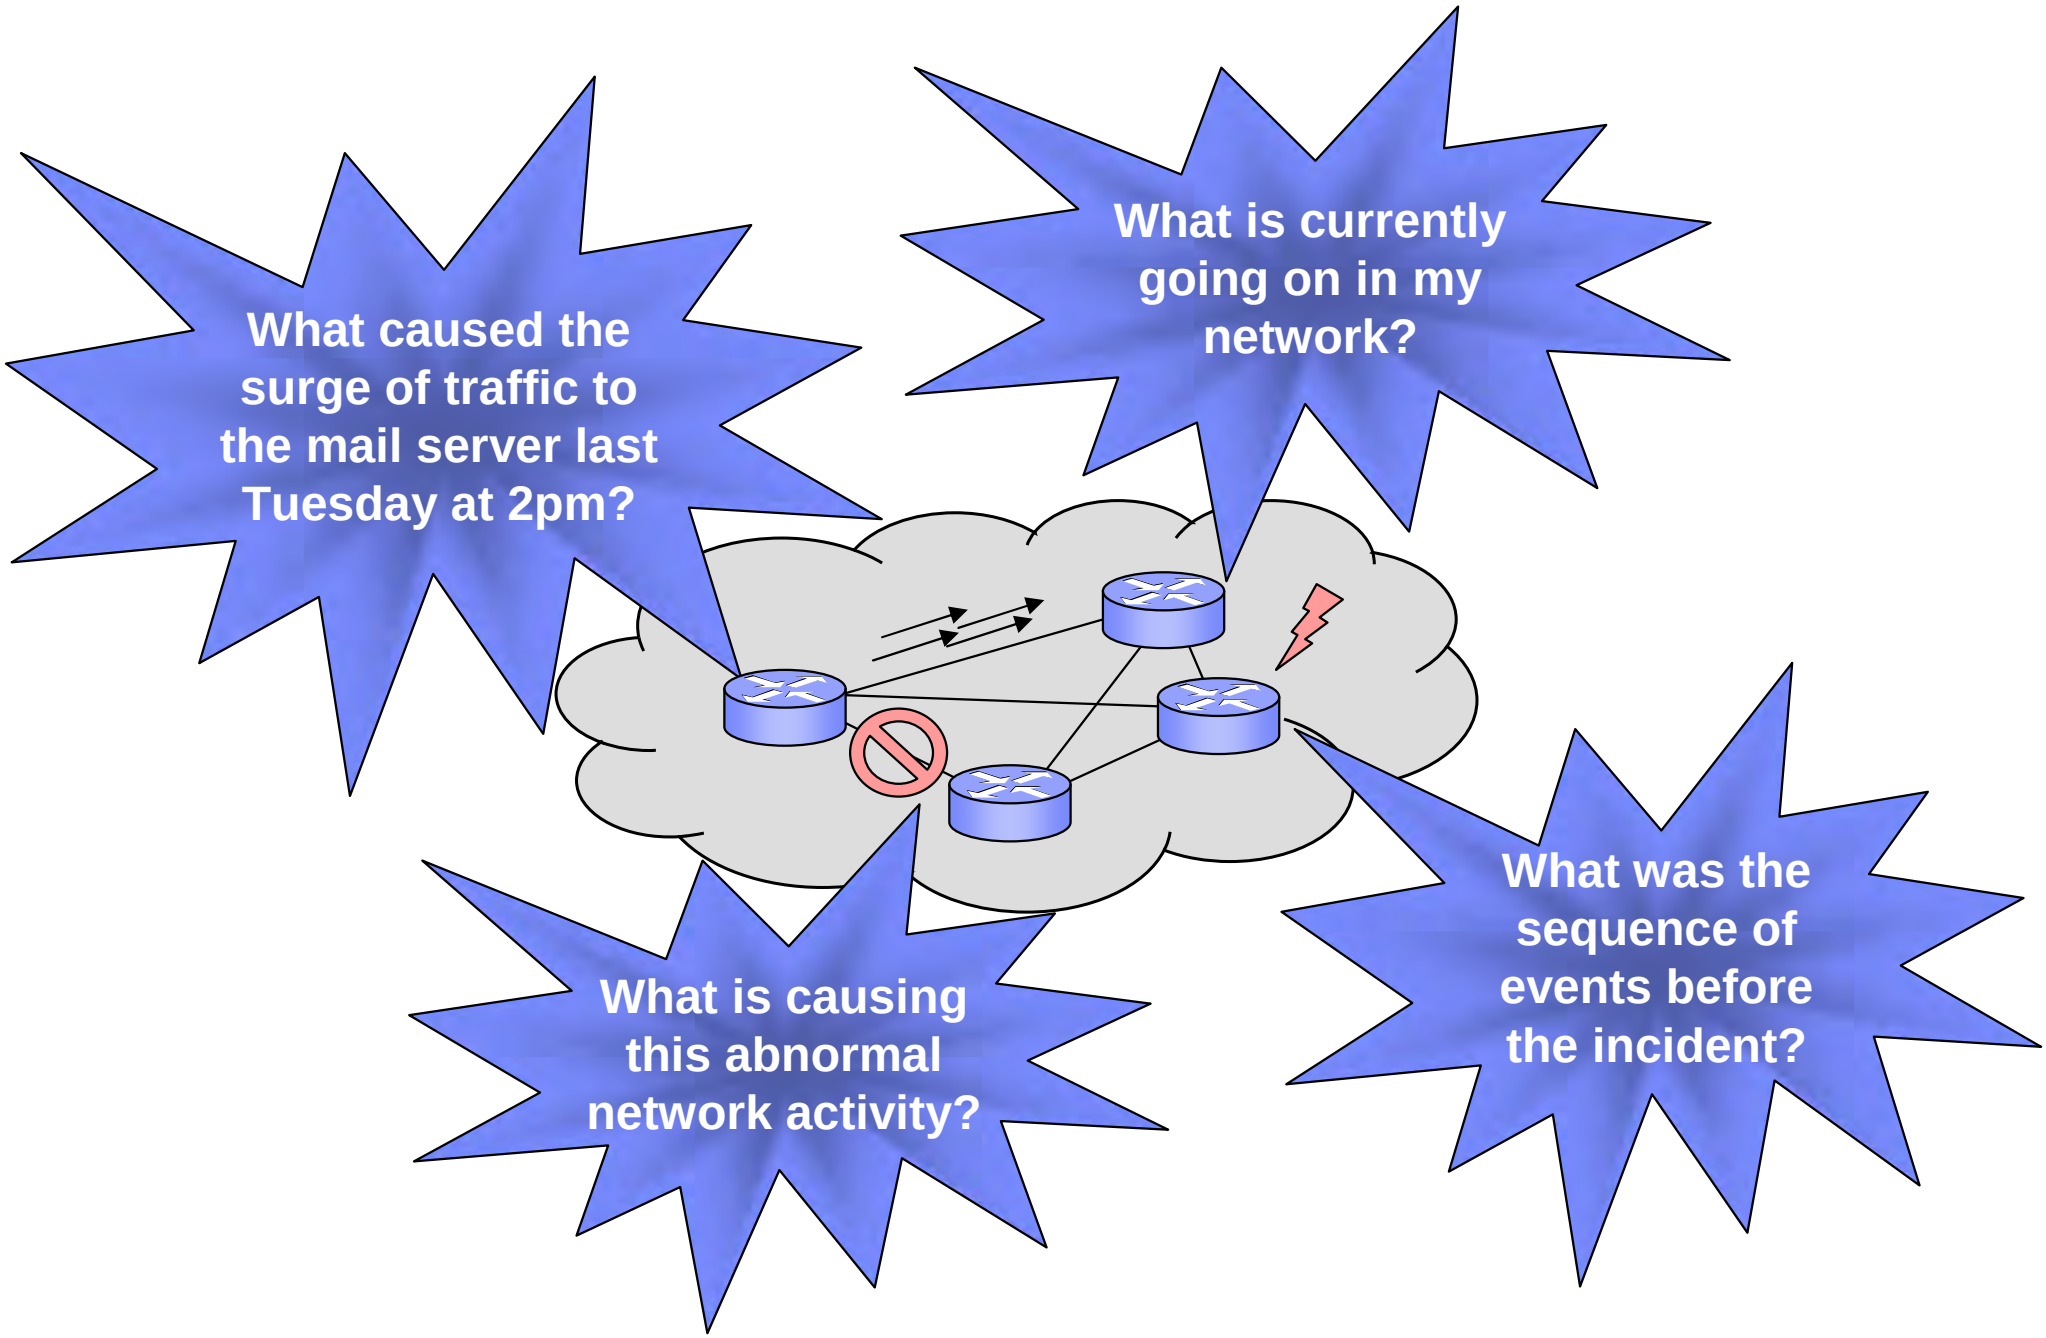
Marc Ph. Stoecklin <mtc@zurich.ibm.com>
Andreas Kind <ank@zurich.ibm.com>
Jean-Yves Le Boudec <jean-yves.leboudec@epfl.ch>

© 2008 IBM Corporation

Outline

- Problem statement and objectives
- Adapting flow information granularity
  - Increasing granularity with Zoom Monitors
  - Decreasing granularity with lossy compression
- Implementation
- Results
- Conclusion and outlook

© 2008 IBM Corporation



© 2008 IBM Corporation

Problem Statement

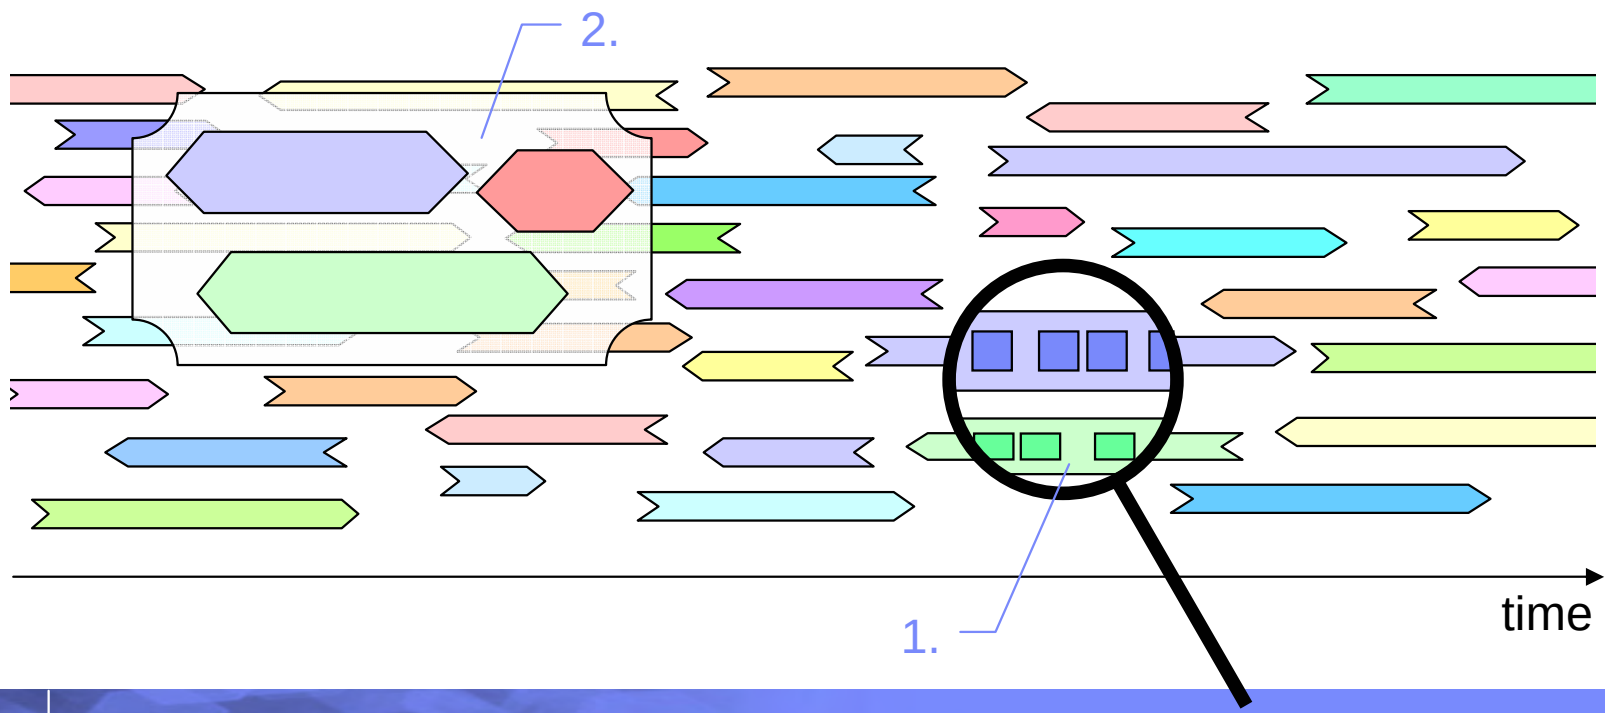
- Trade-off in network **traffic information collection** for **incident analysis**
  - **Raw packet traces**: finest level of detail but impractical to manage and search
  - **Flow traces**: high-level traffic abstraction but aggregated
- Traditional flow exports may **not provide traffic details required** to understand causes of incidents
  - Missing layer 3 and layer 4 header information
  - No packet content information
- Flow-level information is still a **considerable amount of data**
  - Flow record collections are still tedious to search, store, and analyze
  - Majority of this (raw) information is never accessed

© 2008 IBM Corporation

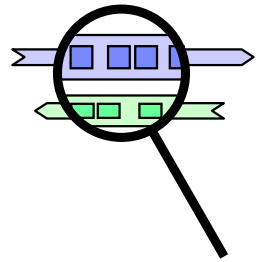
Objectives and Goals

- Extend a collector system to provide more accurate incident analysis
- Adapt information granularity depending on relevance of the traffic:

1. Focus in on particular traffic events to obtain more details
2. Compress known/less relevant traffic events (conserve a meaningful abstraction)



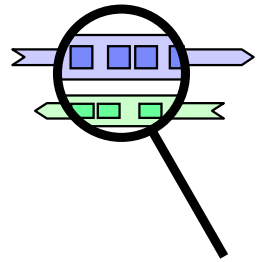
© 2008 IBM Corporation



Increasing Traffic Information Granularity

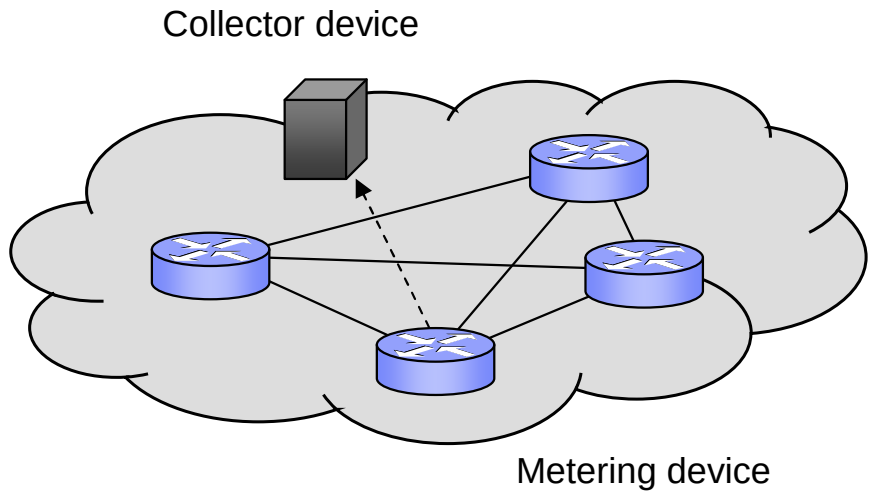
- **Problem**
  - Collecting detailed traffic information is cumbersome
  - Fixed and limited amount of information in default flow exports (e.g., NetFlow v5)
    - Valuable information may have been lost along with flow aggregation
- **Traditional approach (on-going anomaly)**
  - Physically attach a probe or packet dumping device at router (e.g., tcpdump with filtering)
  - Collection of rigid traffic information (e.g., entire packets): complex analysis
- **How to simplify data collection? Create Zoom Monitors!**
  - Dynamically controlled collection of traffic information at desired level of detail
  - Central management console for coordination
  - Make use of capabilities of network device inventory (routers, switches): reporting/dumping

© 2008 IBM Corporation



Zoom Monitors

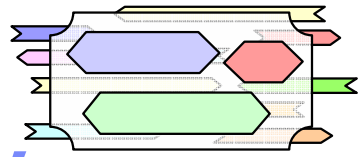
- **Specification**
  - Metering point and collector device
  - Zoom monitor lifespan
  - Filter criteria
  - Traffic aspects to be exported



- **Export collection and display**
  - Reconfigure metering device to create specific exports
  - Prepare collector device to store exported traffic information
  - Centralized management and display

- **Examples**
  - Show me the payload of all DNS requests of host 10.3.4.5 during the next 10 minutes
  - Look for all internal hosts scanning on TCP service port 9996 (e.g., candidate worm traffic)
  - Inspect GET/POST requests and virtual servers accessed on web server 10.4.5.6
  - Export unsampled flow measurements from subnet 10.9.3.1/24

© 2008 IBM Corporation

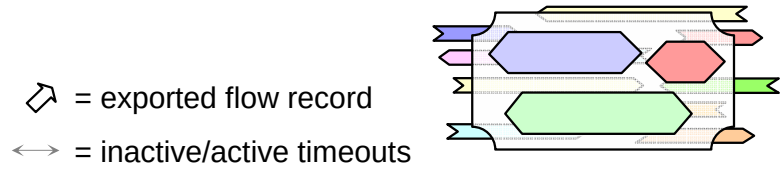


Decreasing Traffic Information Granularity

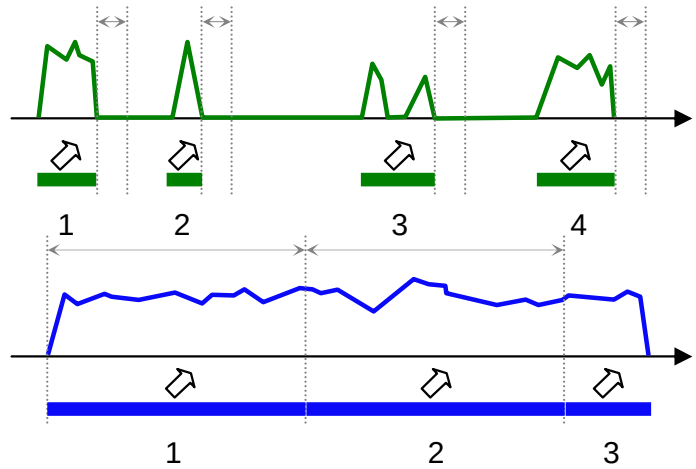
- **Problem**
  - Most stored traffic information is irrelevant for incident analysis (never accessed)
  - Redundancy (limited value): Increased storage overhead and search complexity
- **Traditional approaches**
  - Rolling database (FIFO): keep all records up to a limit (e.g., #records, age): information removal
  - Uniform summarization: adapt resolution of information (hourly, daily, weekly)
  - Keep top-k entries (according to some aspect)
- **How can we do better?**
  - Majority of network events is known or recurring
  - Gradually compress information of irrelevant traffic events in a lossy fashion
    - With minimal impact on incident analysis tasks
  - Summarize similar events (coarse-grained representation)

© 2008 IBM Corporation

Observations



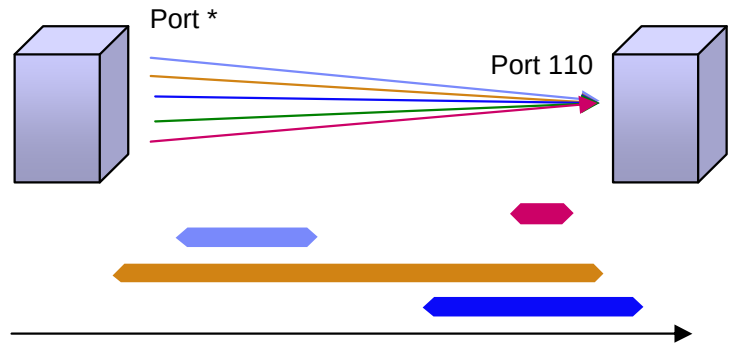
- **Flow exports**
  - Multiple exports for a single connection
  - Examples:
    - Long-lived connections (streams, remote sessions, etc.)
    - Timeouts on routers (inactive/active timeout)



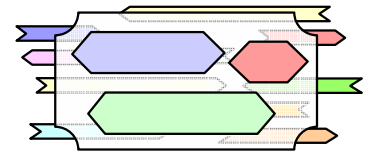
- **Bi-directionality**
  - Most flows have a reversed counterpart



- **Information similarity**
  - Sets of records with limited added value on the flow level
  - Groups of flows with similar properties (Web, mail, printer traffic, polling)
  - Uniqueness: ephemeral port, time stamps, byte and packet counters



© 2008 IBM Corporation



Compression Model[1]

	Abstraction models			
	Flow record	**Flow**	**Conversation**	**Session**
Raw exports	Yes	No	No	No
Flow definition	Yes	Yes	Yes	No (subset thereof)
Direction	Uni-directional	Uni-directional	Bi-directional	Bi-directional
# Flow records	1	≥ 1	≥ 1	≥ 1
# Flows	1	1	1 or 2	≥ 1 or ≥ 2
# Conversations	1	1	1	≥ 1

[1] without prior knowledge such as domain or application specific information

© 2008 IBM Corporation

Implementation

- **Metering device configuration for Zoom Monitors**
  - Reconfiguration of metering devices
  - Management console
- **Export collector**
  - Collection and storage
  - Traffic information compression
  - Data querying

© 2008 IBM Corporation

Metering Device Configuration

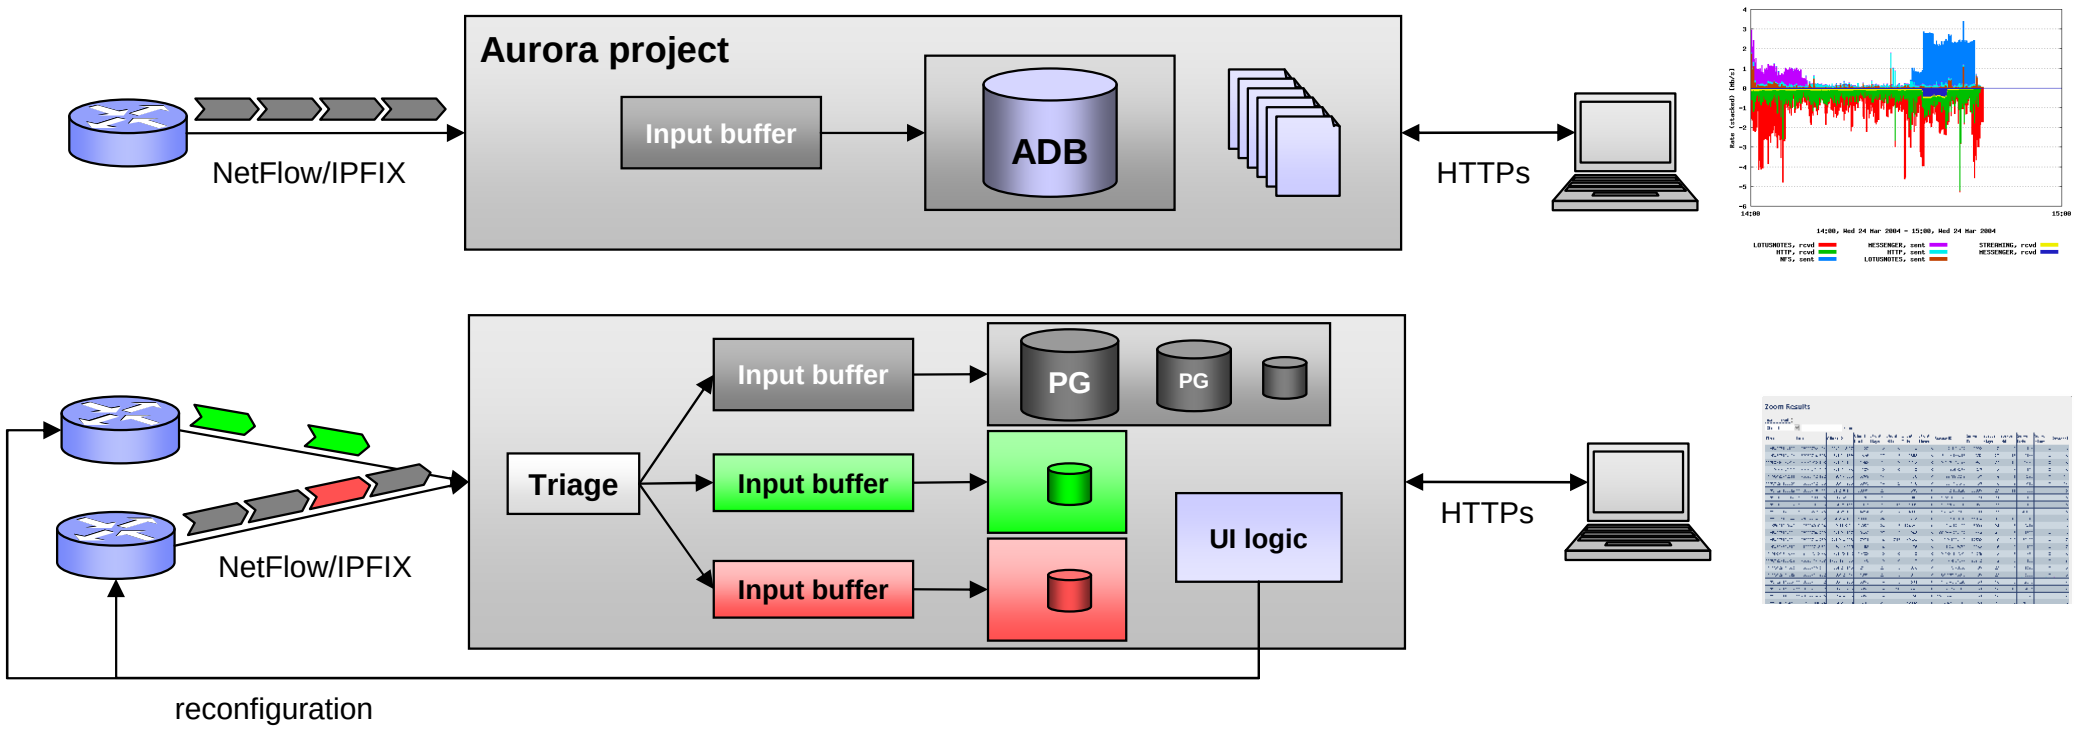
- **Technologies**
  - **Cisco IOS Flexible NetFlow (FNF)**
    - Configuration of multiple customized monitors
    - Currently: input filtering for FNF monitors not available (input filters needed at collector)
  - **Hespera Traffic Meter (IBM Research)**
    - Software-based flow monitor supporting NetFlow v5 and v9, IETF IPFIX exports
    - Customized flow exports (variable templates), CLI-based reconfiguration
    - Filtering with BPF filter syntax

- **User-based creation of dynamic zoom monitors**
  - Web-based specification of zoom monitors
  - Deployment on metering device (CLI-based) and management (e.g., lifespan)
    - Future: XML-based configuration (cf. [Dimitropoulos/Kind] or [NetConf])
  - Registering the zoom monitor at collector device (for disambiguation/triage)
  - Pre-defined zoom monitor templates from library

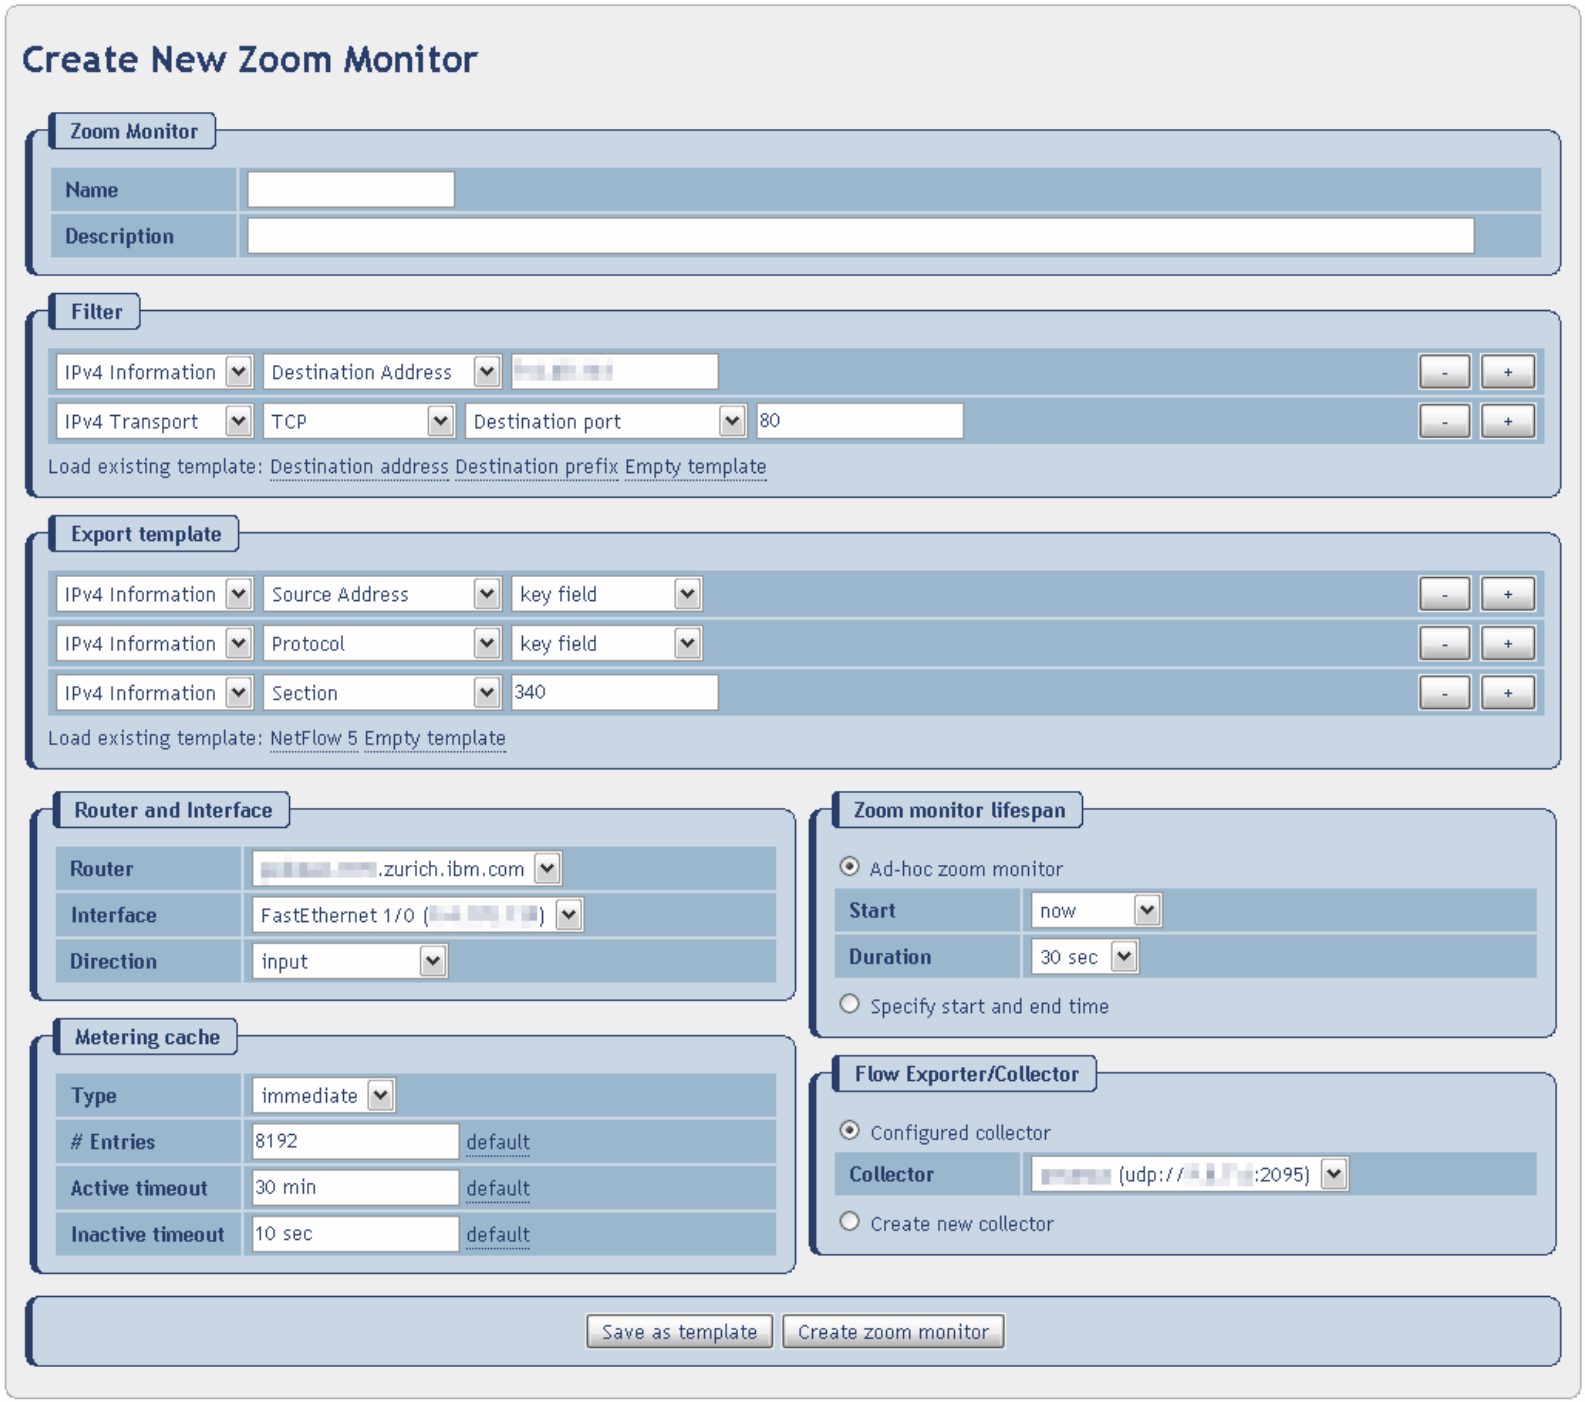
© 2008 IBM Corporation

Export Collector

- **Prototype based on the Aurora flow analyzing system (IBM Research)**
  - Replaced existing Aggregation Database (ADB) with PostgreSQL (PG) backend
  - Input triage according to zoom monitors
  - Incremental population/gradually remove detailed representation: keep “Session”



© 2008 IBM Corporation



Filter definition

Export information

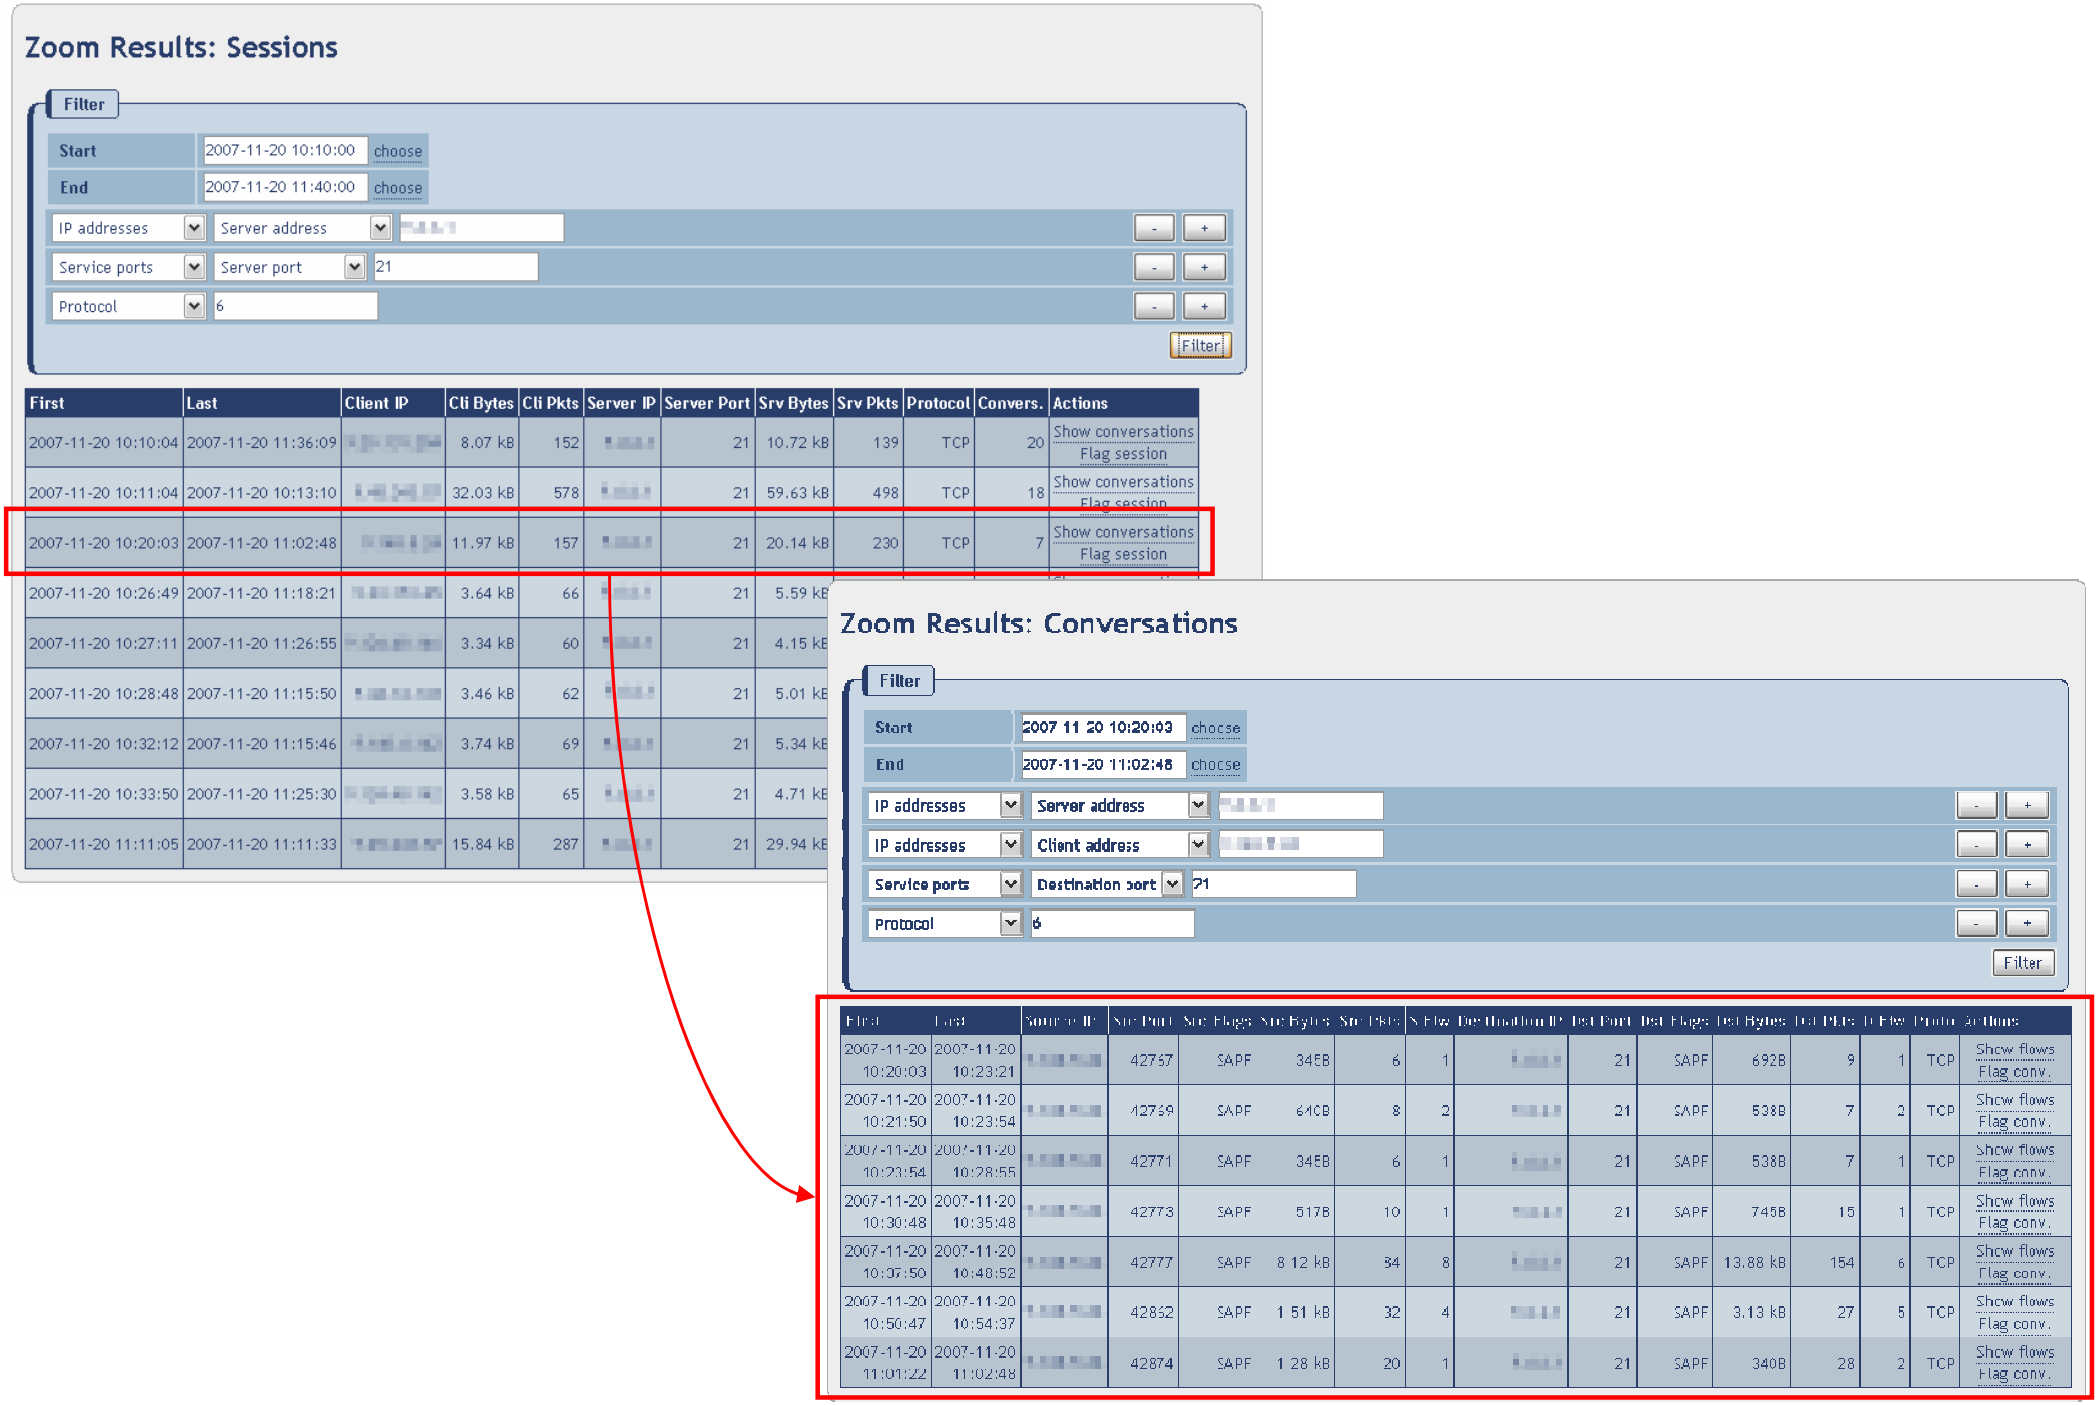
Router/Interface

Lifespan

Collector

Cache

© 2008 IBM Corporation



© 2008 IBM Corporation

Zoom Results: Zoom Monitor 'Payload Section'

Filter

Start [choose]

End [choose]

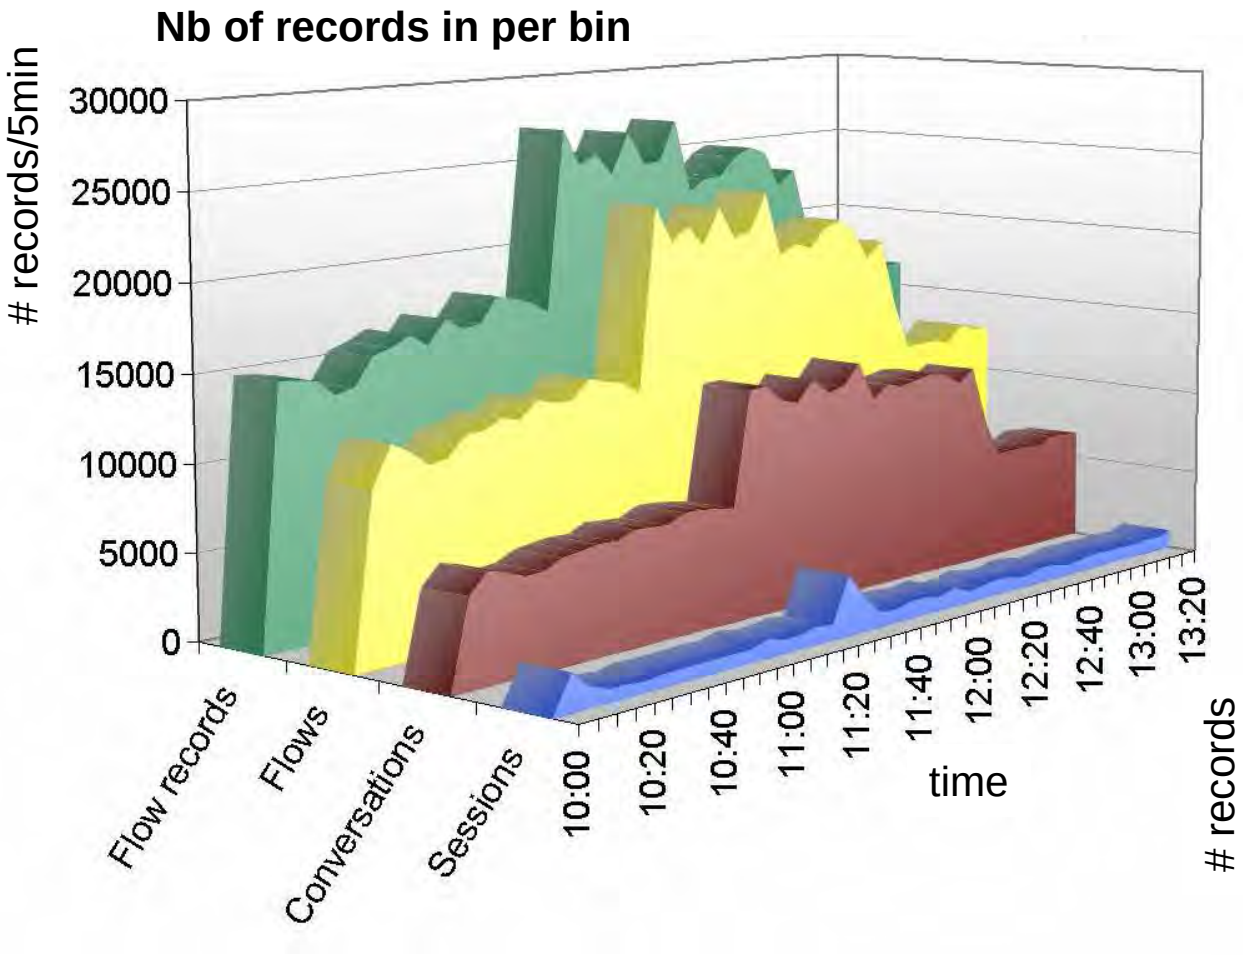
Please select ... [-] [+]

[Filter]

First	Src IP	Dst IP	Protocol	Src Port	Dst Port	Octets	Packets	Payload
2007-11-28 15:53:45.998			UDP	33859	53	57	1	0000 84 43 00 35 00 25 5e e0 3d a3 01 00 00 01 00 00 .....exa mple.com 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 63 6f 6d 0020 00 00 01 00 ....
2007-11-28 15:53:46.002			UDP	53	33859	73	1	0000 00 35 84 43 00 35 e6 83 3d a3 81 80 00 01 00 01 ....M. 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 63 6f 6d 0020 00 00 01 00 01 c0 0c 00 01 00 01 00 00 fe a0 00 0030 04 d0 4d bc
2007-11-28 15:53:47.568			UDP	33859	53	57	1	0000 84 43 00 35 00 25 5e e0 93 c1 01 00 00 01 00 00 .....exa mple.net 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 6e 65 74 0020 00 00 01 00 ....
2007-11-28 15:53:47.573			UDP	53	33859	73	1	0000 00 35 84 43 00 35 91 53 93 c1 81 80 00 01 00 01 ....M. 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 6e 65 74 0020 00 00 01 00 01 c0 0c 00 01 00 01 00 00 fe a9 00 0030 04 d0 4d bc
2007-11-28 15:53:51.698			UDP	33859	53	57	1	0000 84 43 00 35 00 25 5e e0 3c ea 01 00 00 01 00 00 .....exa mple.org 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 6f 72 67 0020 00 00 01 00 ....
2007-11-28 15:53:51.705			UDP	53	33859	73	1	0000 00 35 84 43 00 35 d0 36 3c ea 81 80 00 01 00 01 ....M. 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 6f 72 67 0020 00 00 01 00 01 c0 0c 00 01 00 01 00 00 fe b4 00 0030 04 d0 4d bc
2007-11-28 15:54:04.132			UDP	33859	53	56	1	0000 84 43 00 35 09 40 91 7b 78 ae 01 00 00 01 00 00 .....exa mple.fr. ... 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 02 66 72 00 0020 00 01 00
2007-11-28 15:54:04.143			UDP	53	33859	162	1	0000 00 35 84 43 00 8e fd fa b8 ae 81 80 00 01 00 01 ....Q.. 0010 00 02 00 02 07 65 78 61 6d 70 6c 65 02 66 72 00 0020 00 01 00 01 c0 0c 00 01 00 01 00 01 51 80 00 04

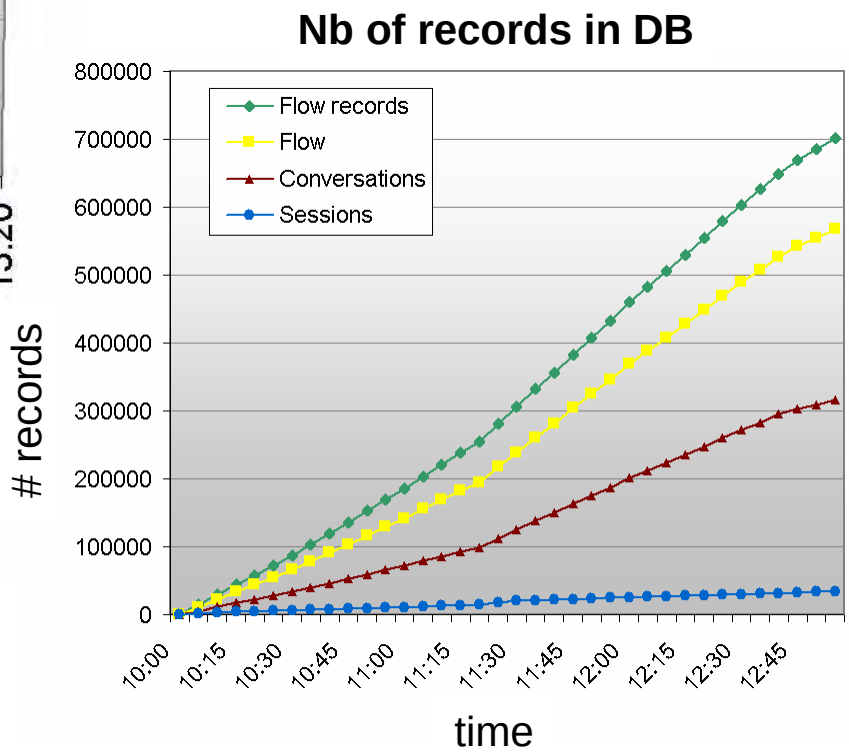
© 2008 IBM Corporation

Results: Compression (WAN traffic)



- **Average compression ratio**

#flow records : #flows	**1.26**	σ = 0.07
#flow records : #conversations	**2.34**	σ = 0.28
#flow records : #sessions	**22.80**	σ = 7.00

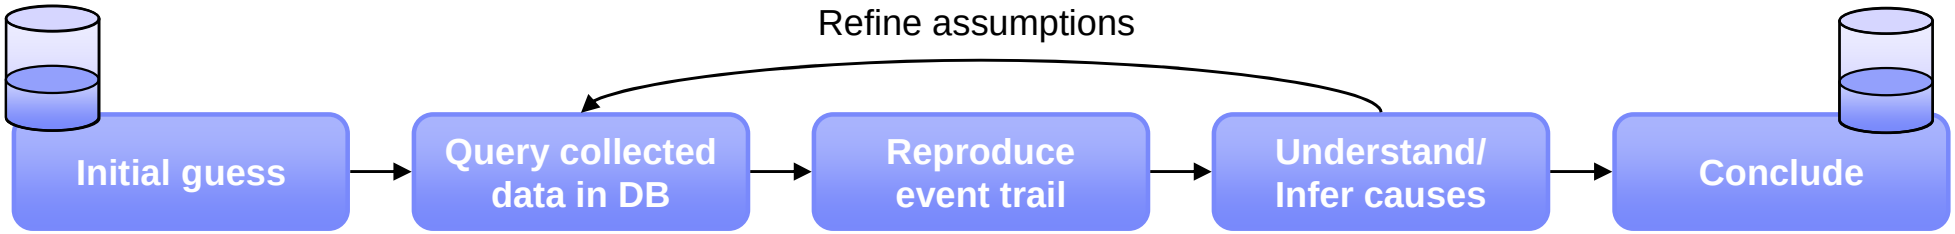


- Session inactive timeout: 20min

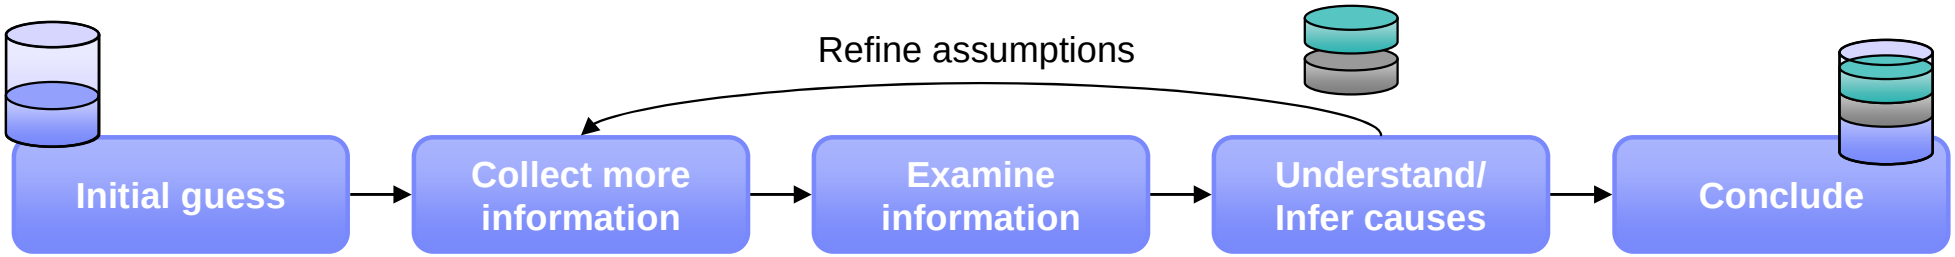
© 2008 IBM Corporation

Traffic Collection for Incident Analysis

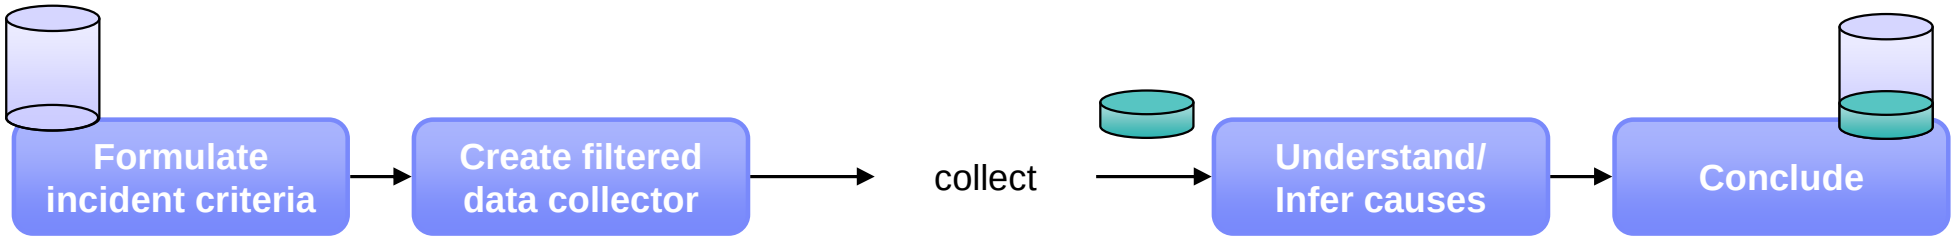
- **After-the-fact analysis**



- **Real-time analysis**



- **Future incident trap**



© 2008 IBM Corporation

Future Work and Vision

- **Automated zoom monitor creation**
  - Interface to a behavior-based network anomaly detection system
  - Proactive collection of evidence for off-line forensic analysis of abnormal events
- **Distributed collector infrastructure**
  - Distributed collectors, e.g., at multiple sites (scalability)
  - Transfer required information to central reporting system on demand
- **Cisco IOS Flexible NetFlow with input filters**
  - Perform filtering on routers to replace software-based metering (and filtering)

© 2008 IBM Corporation

Conclusion

- **Incident analysis tool adapting flow information granularity**
  - Increase level of detail of relevant/unknown traffic events
  - Decrease level of detail (lossy compression) of less relevant events
  - Keep a meaningful abstraction of all traffic events

- **Creation of customized zoom monitors**
  - Zoom in on specific traffic to gain additional information about its properties and behavior
  - Centralized management of metering devices for traffic detail collection

© 2008 IBM Corporation

References

- IBM Research. “Aurora – Network Traffic Analysis and Visualization”. http://www.zurich.ibm.com/aurora/
- Xenofontas Dimitropoulos and Andreas Kind. “Configuration of Monitors”. FloCon2008.
- NETCONF IETF Working Group. http://www.ops.ietf.org/netconf/
- Cisco Systems, Inc. “Cisco IOS Flexible Netflow”. Product website: http://www.cisco.com/en/US/products/ps6965/products_ios_protocol_option_home.html

© 2008 IBM Corporation



Dynamic Adaptation of Flow Information Granularity for Incident Analysis

Marc Ph. Stoecklin <mtc@zurich.ibm.com>
Andreas Kind <ank@zurich.ibm.com>
Jean-Yves Le Boudec <jean-yves.leboudec@epfl.ch>

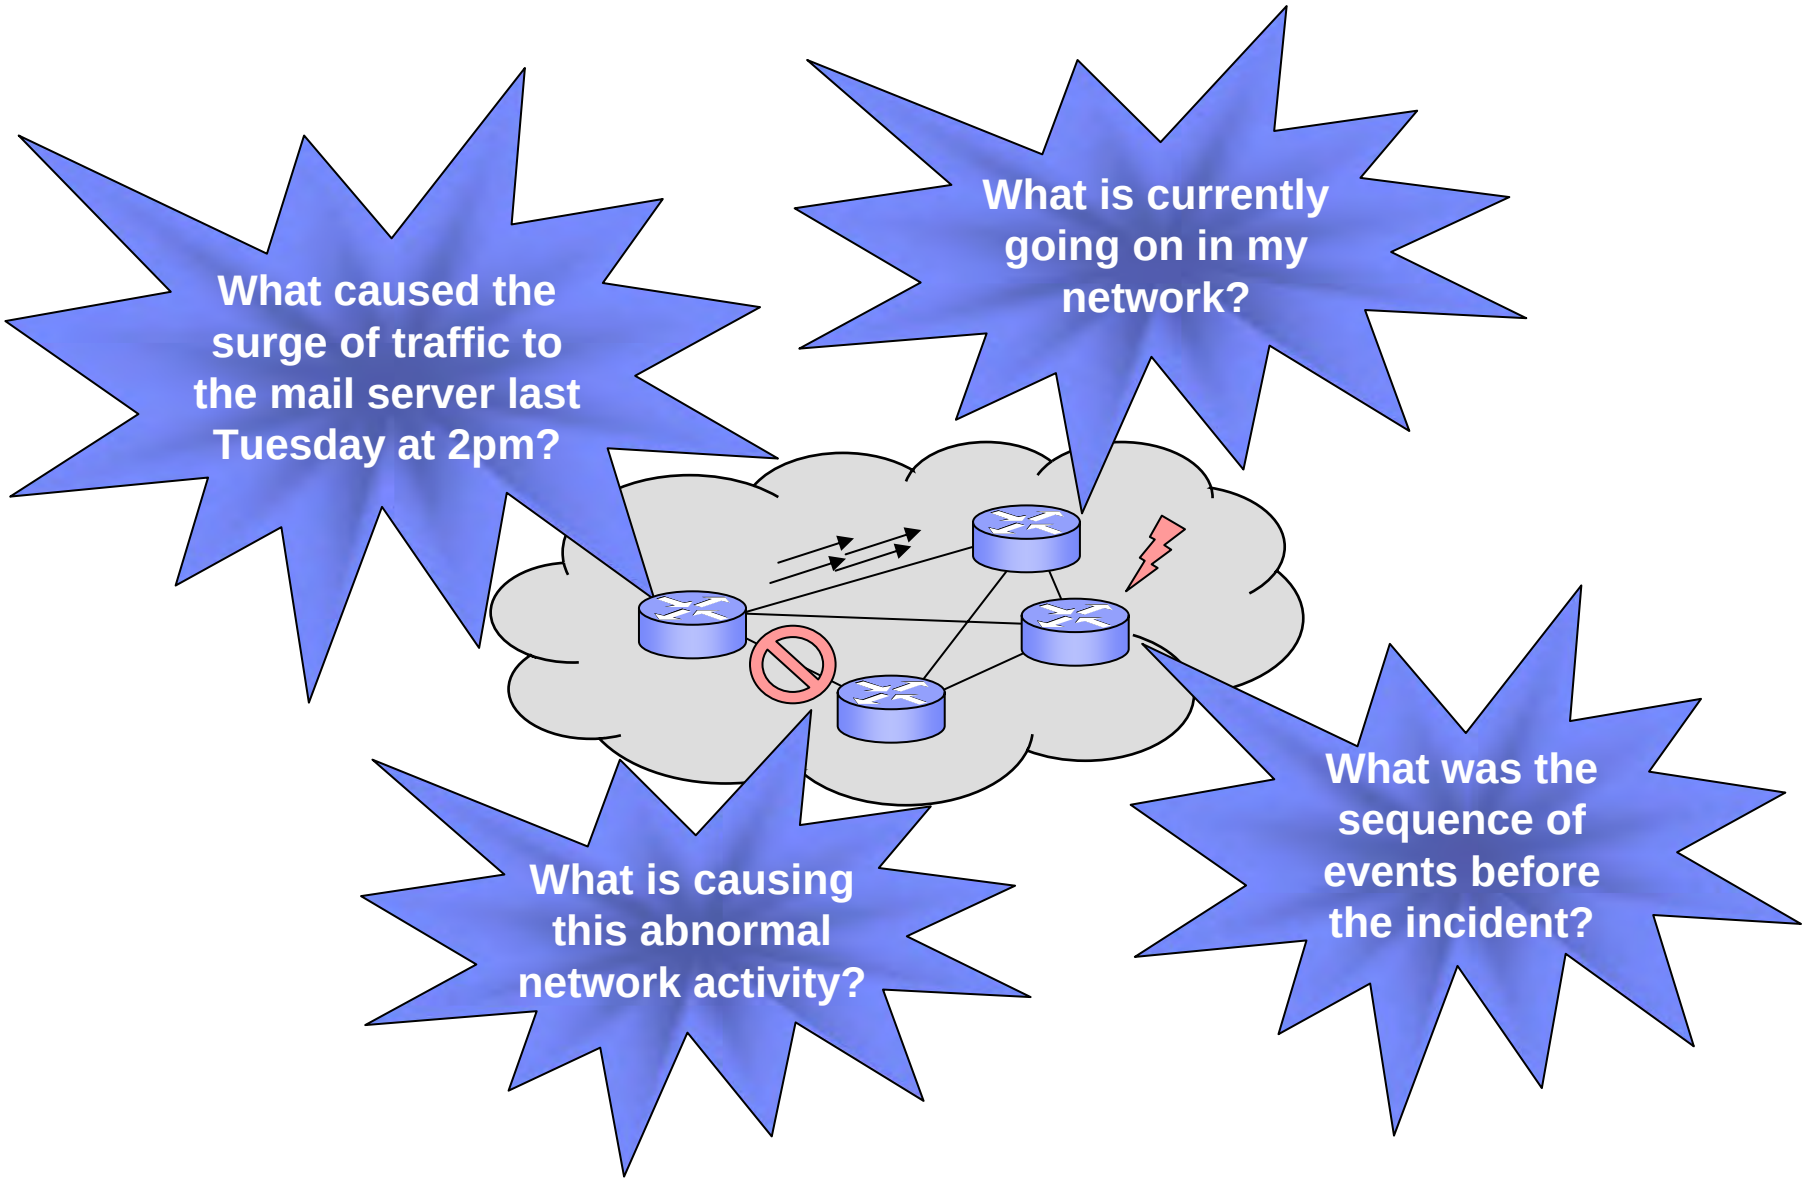


© 2008 IBM Corporation

Outline

- Problem statement and objectives
- Adapting flow information granularity
  - Increasing granularity with Zoom Monitors
  - Decreasing granularity with lossy compression
- Implementation
- Results
- Conclusion and outlook

© 2008 IBM Corporation



© 2008 IBM Corporation

Problem Statement

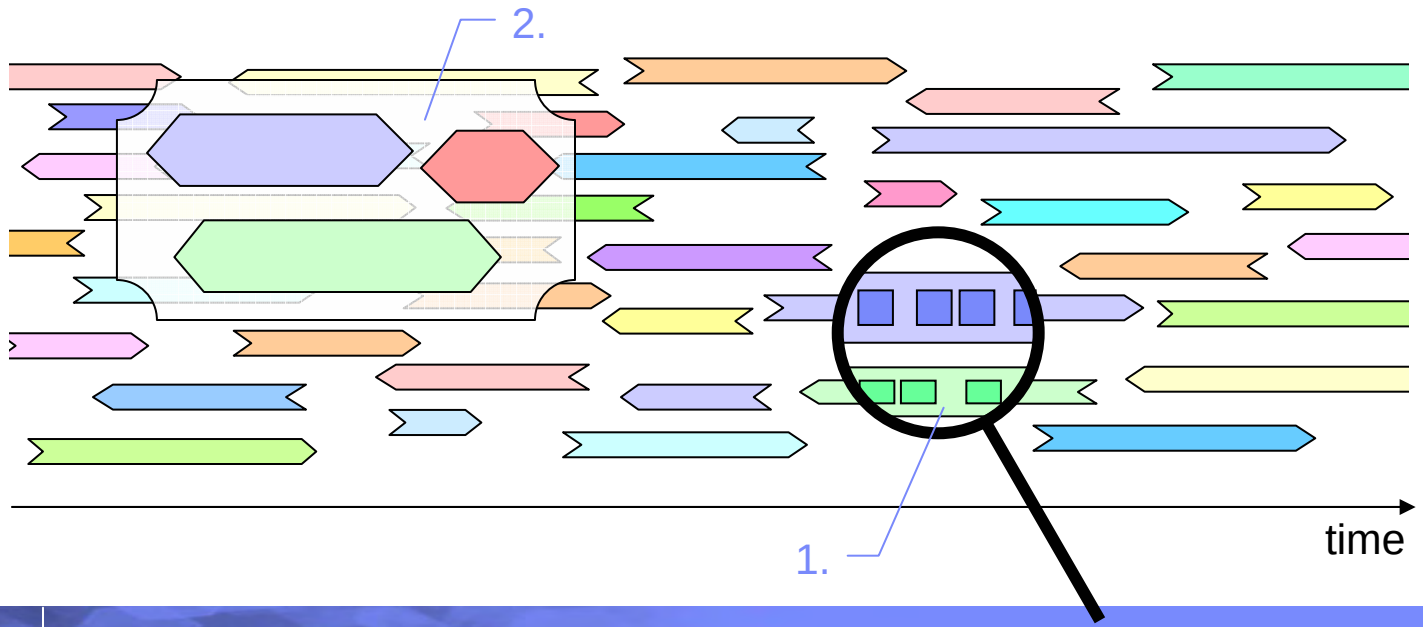
- Trade-off in network **traffic information collection** for **incident analysis**
  - **Raw packet traces**: finest level of detail but impractical to manage and search
  - **Flow traces**: high-level traffic abstraction but aggregated
- Traditional flow exports may **not provide traffic details required** to understand causes of incidents
  - Missing layer 3 and layer 4 header information
  - No packet content information
- Flow-level information is still a **considerable amount of data**
  - Flow record collections are still tedious to search, store, and analyze
  - Majority of this (raw) information is never accessed

© 2008 IBM Corporation

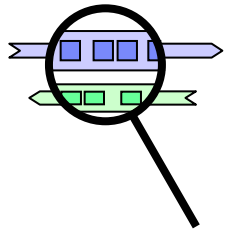
Objectives and Goals

- Extend a collector system to provide more accurate incident analysis
- Adapt information granularity depending on relevance of the traffic:

1. Focus in on particular traffic events to obtain more details
2. Compress known/less relevant traffic events (conserve a meaningful abstraction)



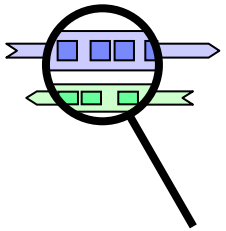
© 2008 IBM Corporation



Increasing Traffic Information Granularity

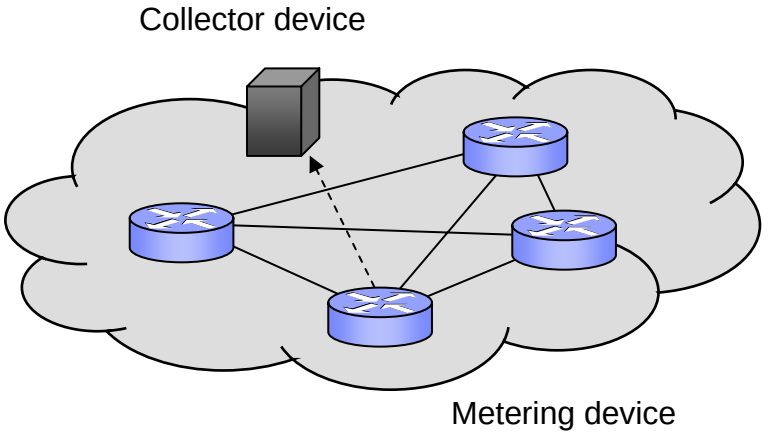
- **Problem**
  - Collecting detailed traffic information is cumbersome
  - Fixed and limited amount of information in default flow exports (e.g., NetFlow v5)
    - Valuable information may have been lost along with flow aggregation
- **Traditional approach (on-going anomaly)**
  - Physically attach a probe or packet dumping device at router (e.g., tcpdump with filtering)
  - Collection of rigid traffic information (e.g., entire packets): complex analysis
- **How to simplify data collection? Create Zoom Monitors!**
  - Dynamically controlled collection of traffic information at desired level of detail
  - Central management console for coordination
  - Make use of capabilities of network device inventory (routers, switches): reporting/dumping

© 2008 IBM Corporation



Zoom Monitors

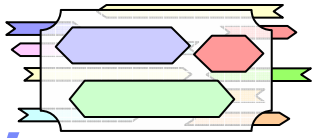
- **Specification**
  - Metering point and collector device
  - Zoom monitor lifespan
  - Filter criteria
  - Traffic aspects to be exported



- **Export collection and display**
  - Reconfigure metering device to create specific exports
  - Prepare collector device to store exported traffic information
  - Centralized management and display

- **Examples**
  - Show me the payload of all DNS requests of host 10.3.4.5 during the next 10 minutes
  - Look for all internal hosts scanning on TCP service port 9996 (e.g., candidate worm traffic)
  - Inspect GET/POST requests and virtual servers accessed on web server 10.4.5.6
  - Export unsampled flow measurements from subnet 10.9.3.1/24

© 2008 IBM Corporation

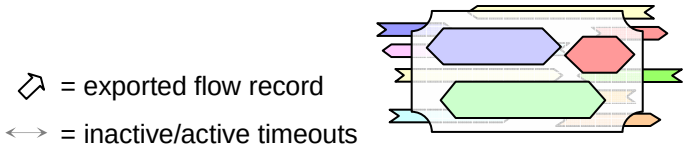


Decreasing Traffic Information Granularity

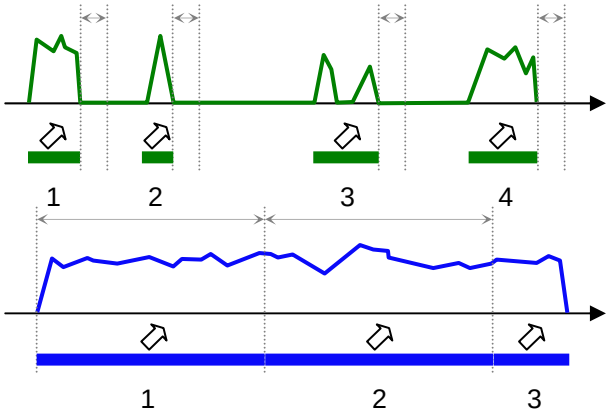
- **Problem**
  - Most stored traffic information is irrelevant for incident analysis (never accessed)
  - Redundancy (limited value): Increased storage overhead and search complexity
- **Traditional approaches**
  - Rolling database (FIFO): keep all records up to a limit (e.g., #records, age): information removal
  - Uniform summarization: adapt resolution of information (hourly, daily, weekly)
  - Keep top-k entries (according to some aspect)
- **How can we do better?**
  - Majority of network events is known or recurring
  - Gradually compress information of irrelevant traffic events in a lossy fashion
    - With minimal impact on incident analysis tasks
  - Summarize similar events (coarse-grained representation)

© 2008 IBM Corporation

Observations



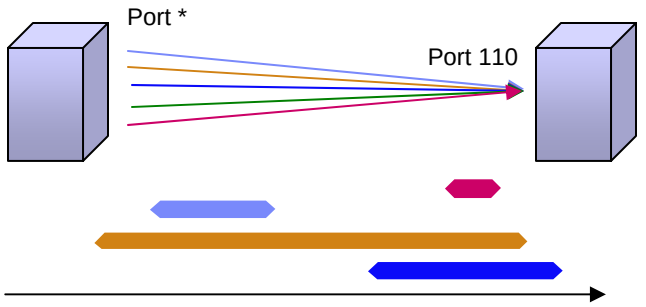
- **Flow exports**
  - Multiple exports for a single connection
  - Examples:
    - Long-lived connections (streams, remote sessions, etc.)
    - Timeouts on routers (inactive/active timeout)



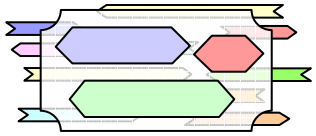
- **Bi-directionality**
  - Most flows have a reversed counterpart



- **Information similarity**
  - Sets of records with limited added value on the flow level
  - Groups of flows with similar properties (Web, mail, printer traffic, polling)
  - Uniqueness: ephemeral port, time stamps, byte and packet counters



© 2008 IBM Corporation



Compression Model[1]

	Abstraction models			
	Flow record	**Flow**	**Conversation**	**Session**
Raw exports	Yes	No	No	No
Flow definition	Yes	Yes	Yes	No (subset thereof)
Direction	Uni-directional	Uni-directional	Bi-directional	Bi-directional
# Flow records	1	≥ 1	≥ 1	≥ 1
# Flows	1	1	1 or 2	≥ 1 or ≥ 2
# Conversations	1	1	1	≥ 1

[1] without prior knowledge such as domain or application specific information

© 2008 IBM Corporation

Implementation

- **Metering device configuration for Zoom Monitors**
  - Reconfiguration of metering devices
  - Management console
- **Export collector**
  - Collection and storage
  - Traffic information compression
  - Data querying

© 2008 IBM Corporation

Metering Device Configuration

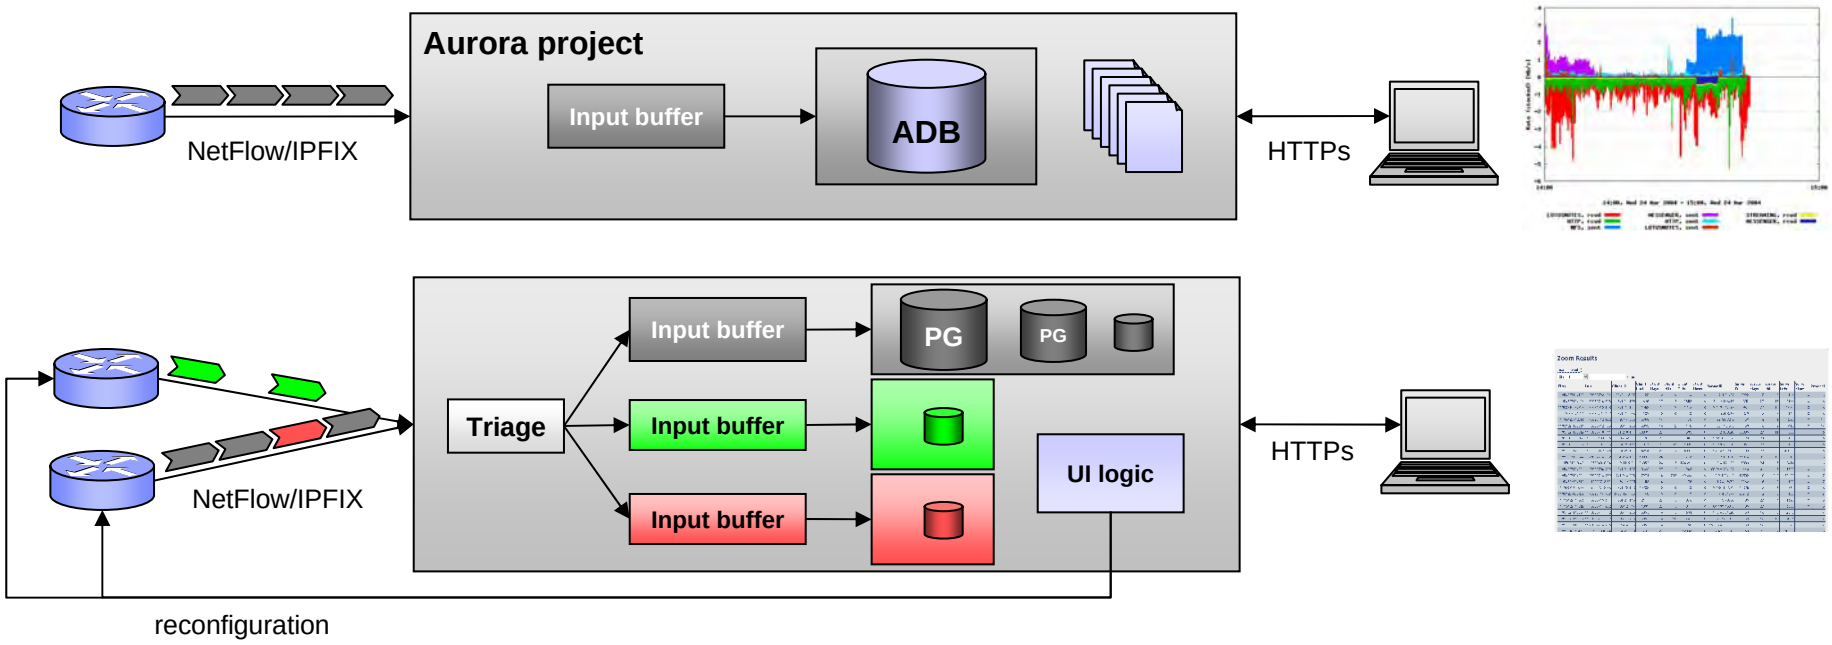
- **Technologies**
  - **Cisco IOS Flexible NetFlow (FNF)**
    - Configuration of multiple customized monitors
    - Currently: input filtering for FNF monitors not available (input filters needed at collector)
  - **Hespera Traffic Meter (IBM Research)**
    - Software-based flow monitor supporting NetFlow v5 and v9, IETF IPFIX exports
    - Customized flow exports (variable templates), CLI-based reconfiguration
    - Filtering with BPF filter syntax

- **User-based creation of dynamic zoom monitors**
  - Web-based specification of zoom monitors
  - Deployment on metering device (CLI-based) and management (e.g., lifespan)
    - Future: XML-based configuration (cf. [Dimitropoulos/Kind] or [NetConf])
  - Registering the zoom monitor at collector device (for disambiguation/triage)
  - Pre-defined zoom monitor templates from library

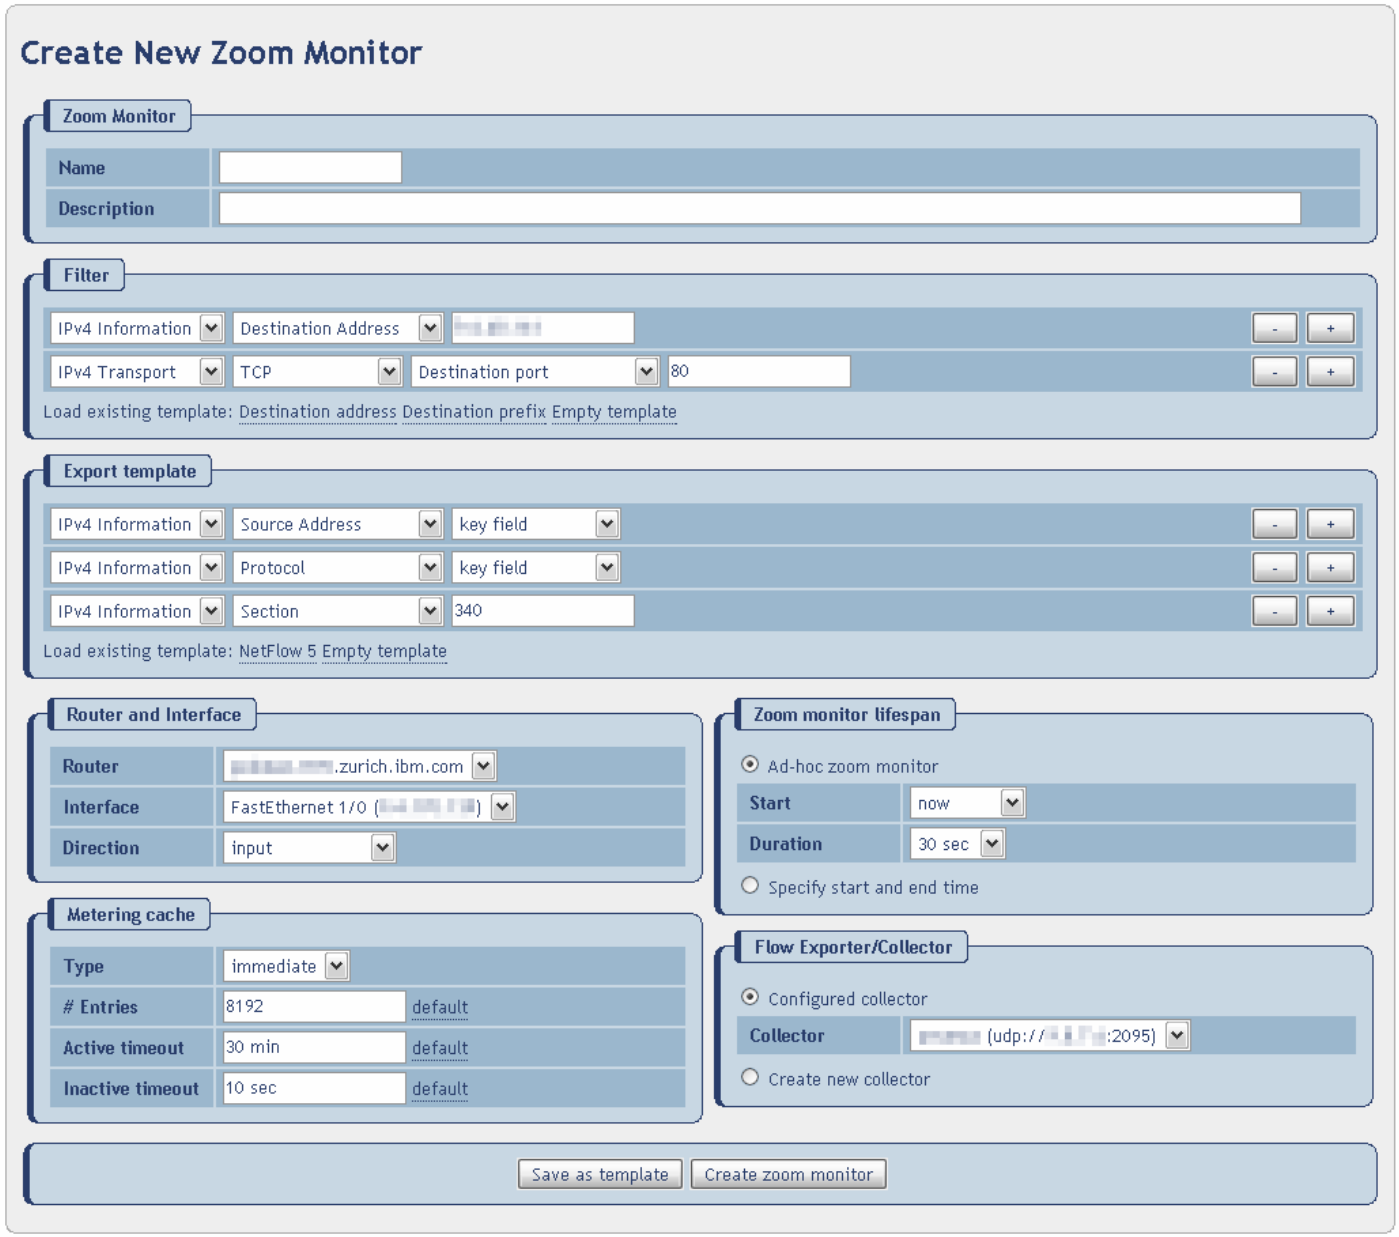
© 2008 IBM Corporation

Export Collector

- **Prototype based on the Aurora flow analyzing system (IBM Research)**
  - Replaced existing Aggregation Database (ADB) with PostgreSQL (PG) backend
  - Input triage according to zoom monitors
  - Incremental population/gradually remove detailed representation: keep “Session”



© 2008 IBM Corporation



Filter definition

Export information

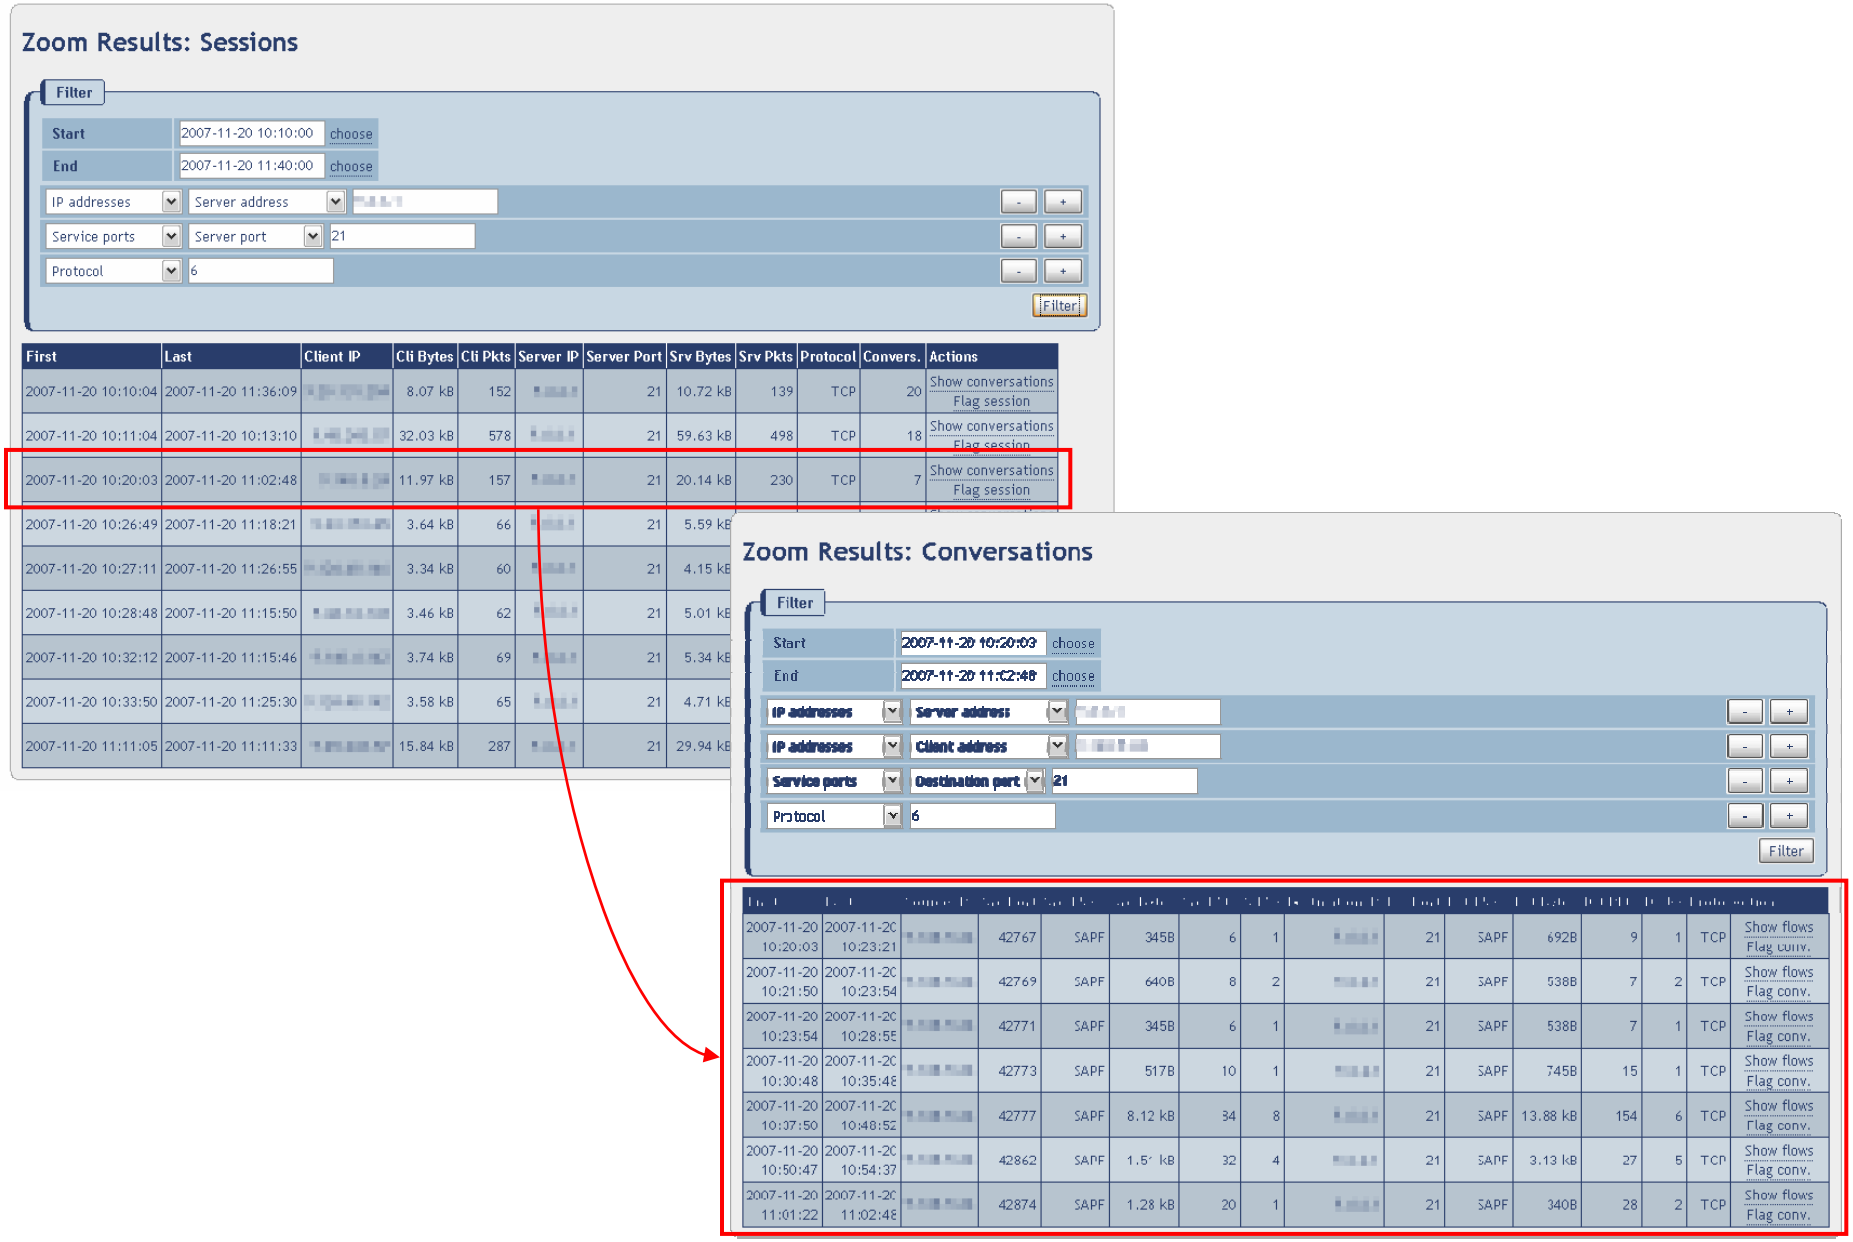
Router/Interface

Lifespan

Collector

Cache

© 2008 IBM Corporation



 © 2008 IBM Corporation

Zoom Results: Zoom Monitor 'Payload Section'

Filter

Start choose

End choose

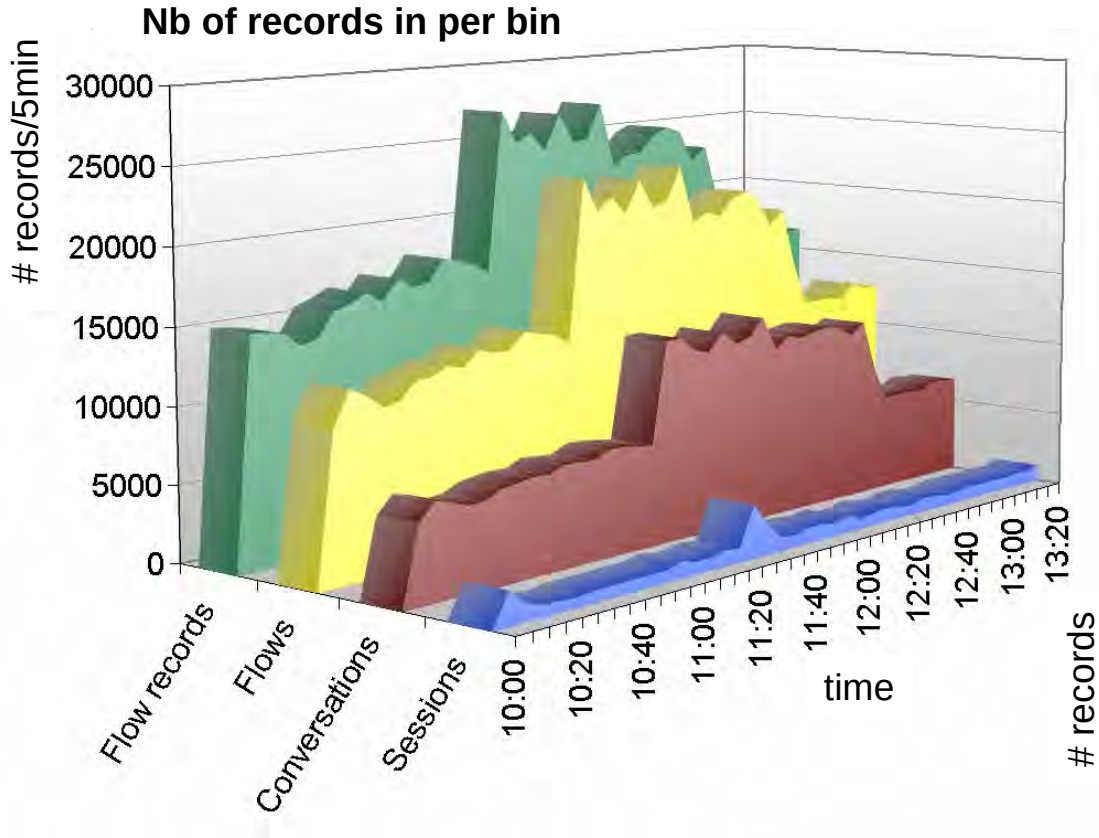
Please select ... - +

Filter

First	Src IP	Dst IP	Protocol	Src Port	Dst Port	Octets	Packets	Payload
2007-11-28 15:53:45.998			UDP	33859	53	57	1	0000 84 43 00 35 00 25 5e e0 3d a3 01 00 00 01 00 00 .....exa mple.com 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 63 6f 6d 0020 00 00 01 00
2007-11-28 15:53:46.002			UDP	53	33859	73	1	0000 00 35 84 43 00 35 e6 83 3d a3 81 80 00 01 00 01 .....exa mple.com 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 63 6f 6d .... 0020 00 00 01 00 01 c0 0c 00 01 00 01 00 00 fe a0 00 ..M. 0030 04 d0 4d bc
2007-11-28 15:53:47.568			UDP	33859	53	57	1	0000 84 43 00 35 00 25 5e e0 93 c1 01 00 00 01 00 00 .....exa mple.net 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 6e 65 74 .... 0020 00 00 01 00
2007-11-28 15:53:47.573			UDP	53	33859	73	1	0000 00 35 84 43 00 35 91 53 93 c1 81 80 00 01 00 01 .....exa mple.net 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 6e 65 74 .... 0020 00 00 01 00 01 c0 0c 00 01 00 01 00 00 fe a9 00 ..M. 0030 04 d0 4d bc
2007-11-28 15:53:51.698			UDP	33859	53	57	1	0000 84 43 00 35 00 25 5e e0 3c ea 01 00 00 01 00 00 .....exa mple.org 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 6f 72 67 .... 0020 00 00 01 00
2007-11-28 15:53:51.705			UDP	53	33859	73	1	0000 00 35 84 43 00 35 d0 36 3c ea 81 80 00 01 00 01 .....exa mple.org 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 03 6f 72 67 .... 0020 00 00 01 00 01 c0 0c 00 01 00 01 00 00 fe b4 00 ..M. 0030 04 d0 4d bc
2007-11-28 15:54:04.132			UDP	33859	53	56	1	0000 84 43 00 35 09 40 91 7b 78 ae 01 00 00 01 00 00 .....exa mple.fr. 0010 00 00 00 00 07 65 78 61 6d 70 6c 65 02 66 72 00 ... 0020 00 01 00
2007-11-28 15:54:04.143			UDP	53	33859	162	1	0000 00 35 84 43 00 8e fd fa b8 ae 81 80 00 01 00 01 .....exa mple.fr. 0010 00 02 00 02 07 65 78 61 6d 70 6c 65 02 66 72 00 ....Q.. 0020 00 01 00 01 c0 0c 00 01 00 01 00 01 51 80 00 04

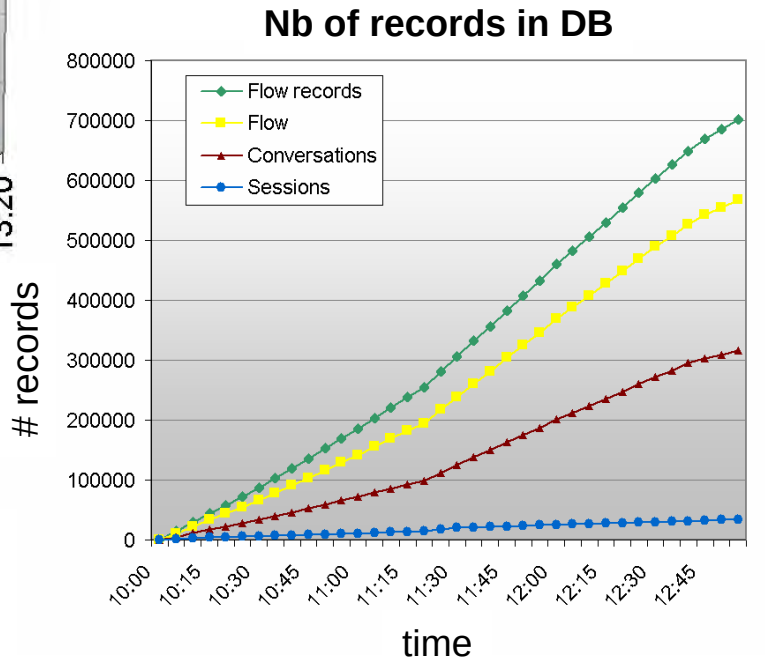
© 2008 IBM Corporation

Results: Compression (WAN traffic)



- **Average compression ratio**

#flow records : #flows	**1.26**	σ = 0.07
#flow records : #conversations	**2.34**	σ = 0.28
#flow records : #sessions	**22.80**	σ = 7.00

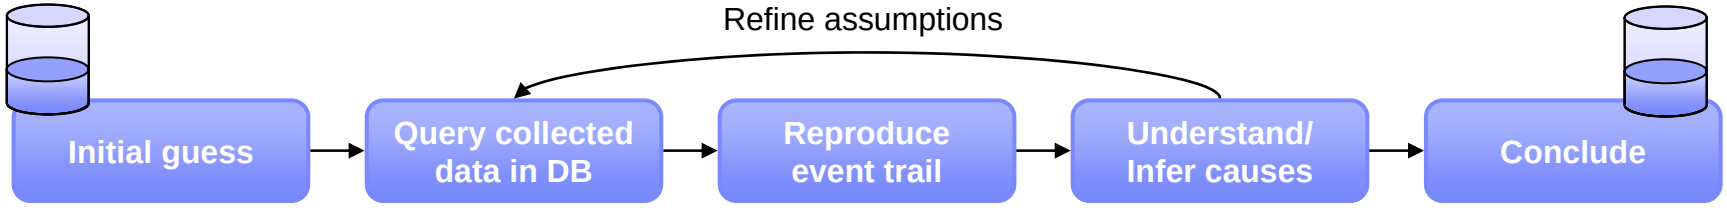


- Session inactive timeout: 20min

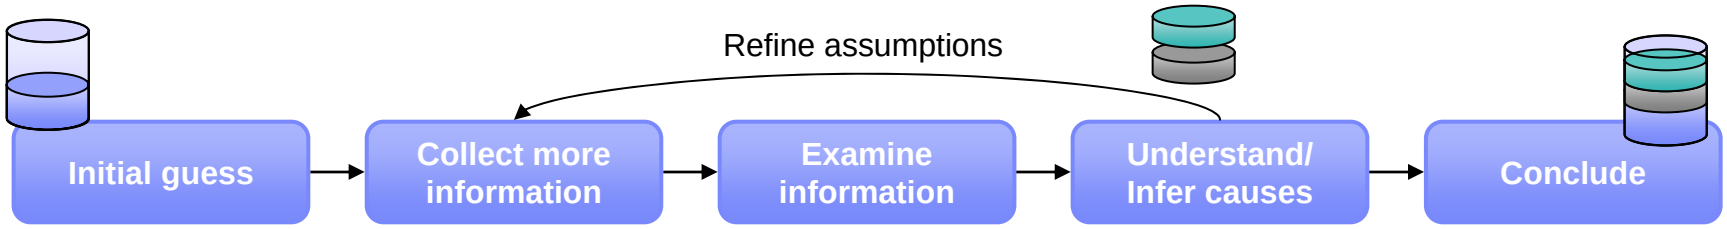
© 2008 IBM Corporation

Traffic Collection for Incident Analysis

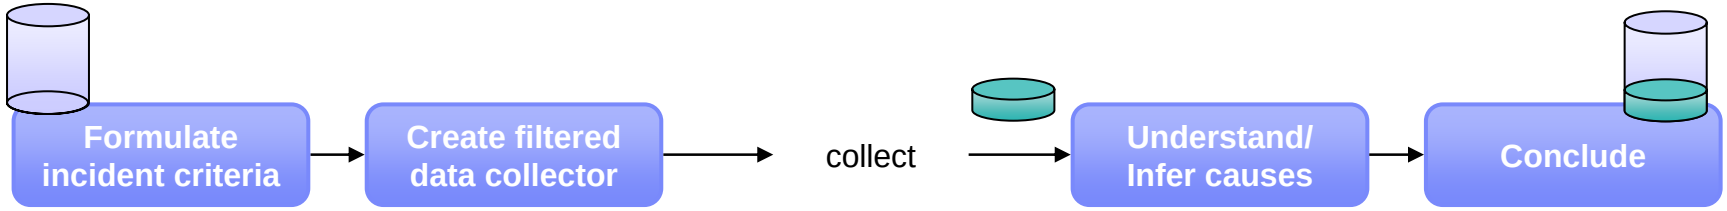
- **After-the-fact analysis**



- **Real-time analysis**



- **Future incident trap**



© 2008 IBM Corporation

Future Work and Vision

- **Automated zoom monitor creation**
  - Interface to a behavior-based network anomaly detection system
  - Proactive collection of evidence for off-line forensic analysis of abnormal events
- **Distributed collector infrastructure**
  - Distributed collectors, e.g., at multiple sites (scalability)
  - Transfer required information to central reporting system on demand
- **Cisco IOS Flexible NetFlow with input filters**
  - Perform filtering on routers to replace software-based metering (and filtering)

© 2008 IBM Corporation

Conclusion

- **Incident analysis tool adapting flow information granularity**
  - Increase level of detail of relevant/unknown traffic events
  - Decrease level of detail (lossy compression) of less relevant events
  - Keep a meaningful abstraction of all traffic events
- **Creation of customized zoom monitors**
  - Zoom in on specific traffic to gain additional information about its properties and behavior
  - Centralized management of metering devices for traffic detail collection

© 2008 IBM Corporation

References

- IBM Research. “Aurora – Network Traffic Analysis and Visualization”. http://www.zurich.ibm.com/aurora/
- Xenofontas Dimitropoulos and Andreas Kind. “Configuration of Monitors”. FloCon2008.
- NETCONF IETF Working Group. http://www.ops.ietf.org/netconf/
- Cisco Systems, Inc. “Cisco IOS Flexible Netflow”. Product website: http://www.cisco.com/en/US/products/ps6965/products_ios_protocol_option_home.html

© 2008 IBM Corporation



High performance. Delivered.

Assessing Disclosure Risk in Anonymized Datasets

Michele Bezzi (ATL) & Alexei Kounine (EPFL)

Copyright © 2008 Accenture All Rights Reserved. Accenture, its logo, and High Performance Delivered are trademarks of Accenture.

Outline

- Background & Motivation
- Anonymisation
- Disclosure Risk Estimation
  - Entropy measure
  - Properties
- Case Study: Flows
- Final remarks

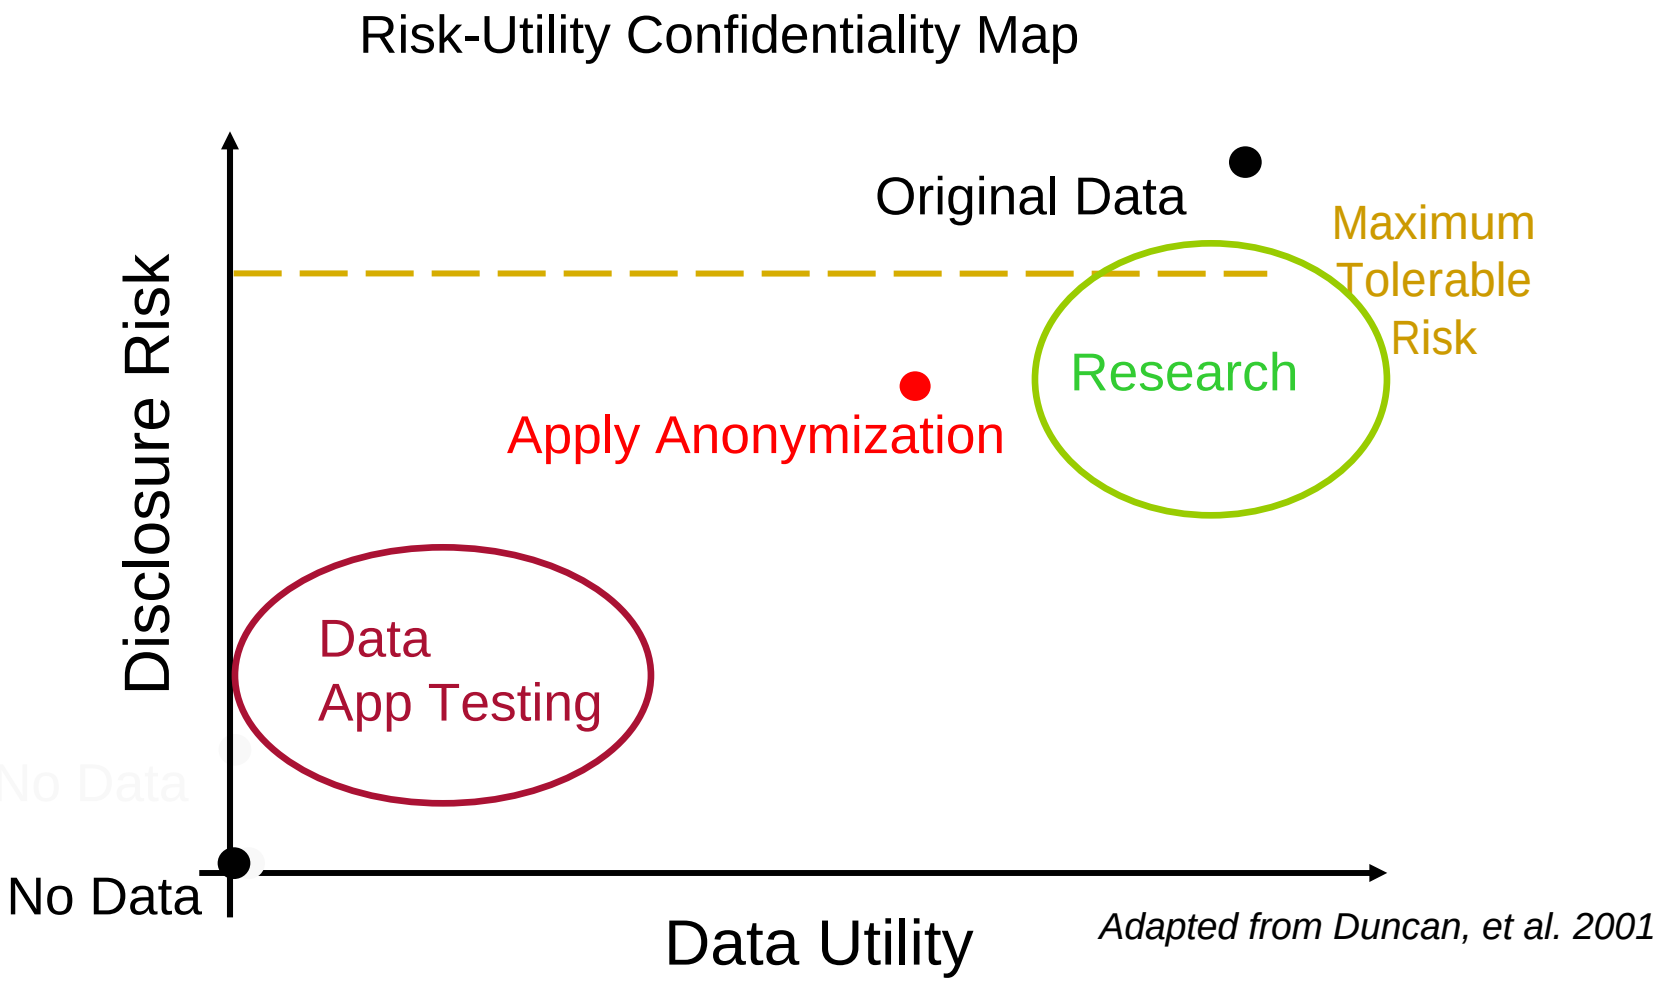
Copyright © 2008 Accenture All Rights Reserved.

Goal

- Problem:
  - The goal is to transform original data records so that no sensitive personal data are disclosed, whereas preserving the maximum amount of relevant information (*anonymity vs. utility* trade off), data integrity and consistency.
- Application
  - Creating datasets for application testing, whenever production DB contains sensitive data. (Our original goal)
  - Allowing researchers to share data and run analytical models on micro-data (e.g., log files), preserving privacy.

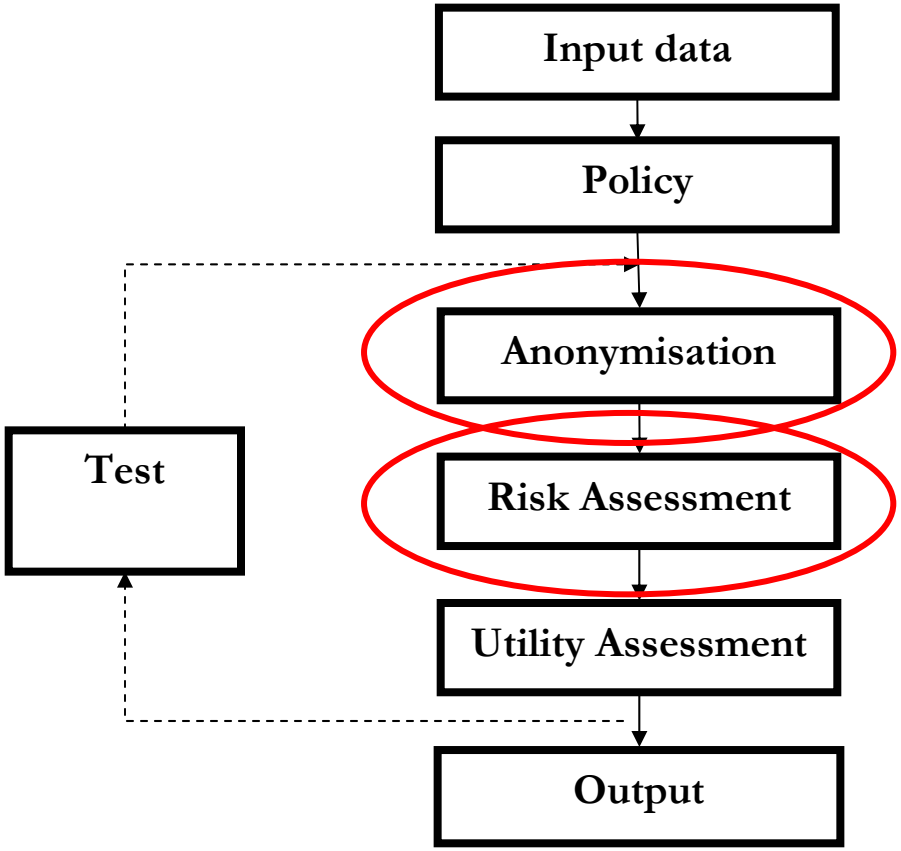
Copyright © 2008 Accenture All Rights Reserved.

Goal



Copyright © 2008 Accenture All Rights Reserved.

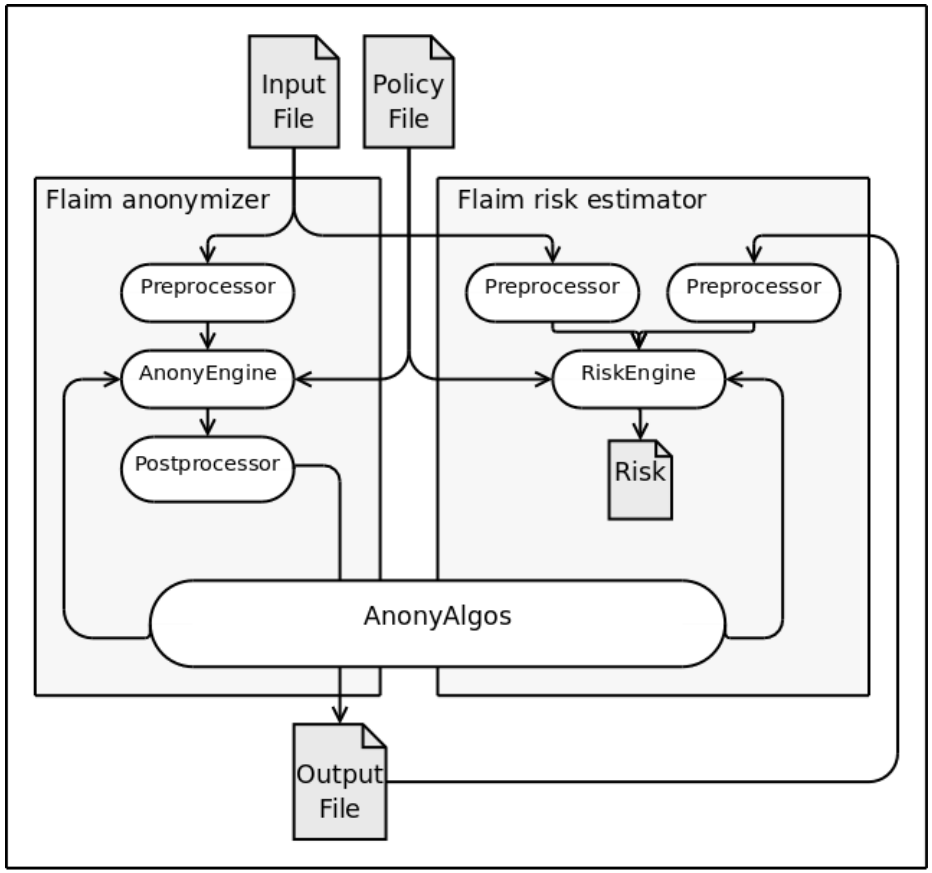
Anonymisation engine & risk estimation



Copyright © 2008 Accenture All Rights Reserved.

Implementation

- Using FLAIM (Framework for Log Anonymization and Information Management), developed by NCSA
- FLAIM anonymization engine (adapted) + risk module



Copyright © 2008 Accenture All Rights Reserved.

Anonymisation primitives

IPs

- Black Marker (16 bits):
- Random Permutation (one-to-one mapping)
- Prefix-preserving (random permutation, but preserving structure)

IP Address	Black Marker (16-bit)	Random Permutation	Prefix-preserving
168.125.96.167	*168.125.0.0*	*124.12.132.37*	*12.131.102.67*
168.125.96.18	*168.125.0.0*	*231.45.36.167*	*12.131.102.17*
168.125.132.37	*168.125.0.0*	*12.72.8.5*	*12.131.201.29*

Port number

- Bilateral Classification: Replace with 0 or 65535 (the port smaller or larger than 1024): E.g., *27 -> 0 , 2048->65535*

Number of packets/bytes

- Add random noise (zero-average)
- Classification

Copyright © 2008 Accenture All Rights Reserved.

Attack scenario

• The attacker aims at re-identifying released data by linking them with some background knowledge, which has some overlapping attributes with the released dataset.

• Estimating *P(r|s):* knowing data masking transformations, distance based similarity

• More uncertain mapping is - lower risk

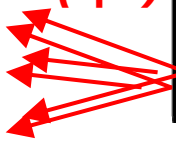
• Because the data holder does not know in advance which records and attributes might be available to the attacker, it must run the risk analysis on the whole released dataset and assume a set of key attributes the attacker might know and use for re-identification.

Original data *S*

SrcIP	SrcPort	DestIP	DestPort	Packets
168.125.253.2	80	147.81.124.1	3157	40
39.109.219.43	7310	142.68.22.108	59959	126
35.187.130.82	161	213.48.19.68	22	83

Anonymised data *R*

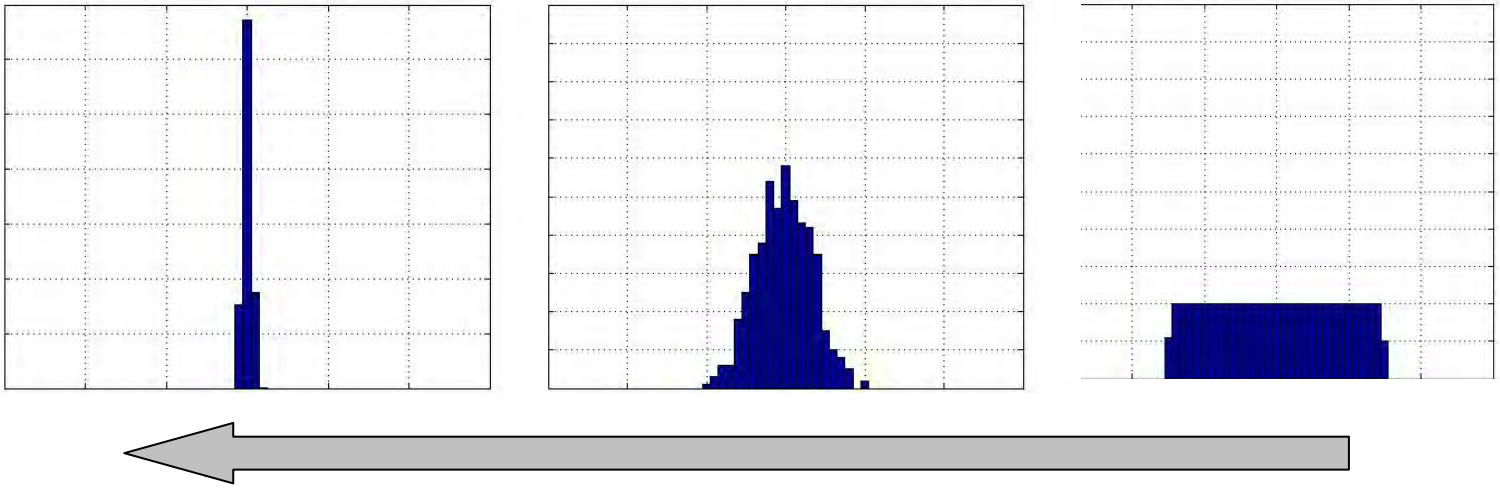
SrcIP	SrcPort	DestIP	DestPort	Packets
168.125.253.0	1023	10.1.1.1	65535	42
39.109.219.0	65535	10.1.1.1	65535	132
35.187.130.0	1023	10.1.1.1	0	81



Original data *S*
Background knowledge *S'*

SrcIP	SrcPort	DestIP	DestPort	Packets
168.125.253.2	80	147.81.124.1	3157	40
39.109.219.43	7310	142.68.22.108	59959	126
35.187.130.82	161	213.48.19.68	22	83

Estimating risk



High Risk
of re-identification

Low Risk
of re-identification

H=1.2 *H=3.7* *H=4.9*

Shannon entropy: Average # of binary questions to identify *s*

Small: risky

Large: safe

Copyright © 2008 Accenture All Rights Reserved.

Entropy as a risk measure

Shannon entropy: Average # of binary questions to identify a *single s*

$$H(\mathcal{R}|s) = -\sum_{r \in \mathcal{R}} P(r|s) \log_2 P(r|s)$$

Global risk:

Expected number of correct matches

$$E_{CM} = \sum_{s \in \mathcal{S}} \frac{1}{2^{H(\mathcal{R}|s)}}$$

k-anonymity condition

Copyright © 2008 Accenture All Rights Reserved.

Some properties

- Directly linked to information loss (utility):

$$I(\mathcal{S},\mathcal{R}) = H(\mathcal{R}) - \sum_{s\in\mathcal{S}} P(s)H(\mathcal{R}|s)$$

  - Minimal info loss:

$$\sum_{s\in\mathcal{S}} P(s)\log_2 H(R|s) \qquad \text{with constraint } H(R|s) \geq h_{min}$$

- Additivity

$$H(\mathcal{R}_1,\mathcal{R}_2|s) = H(\mathcal{R}_1|s) + H(\mathcal{R}_2|\mathcal{R}_1,s)$$

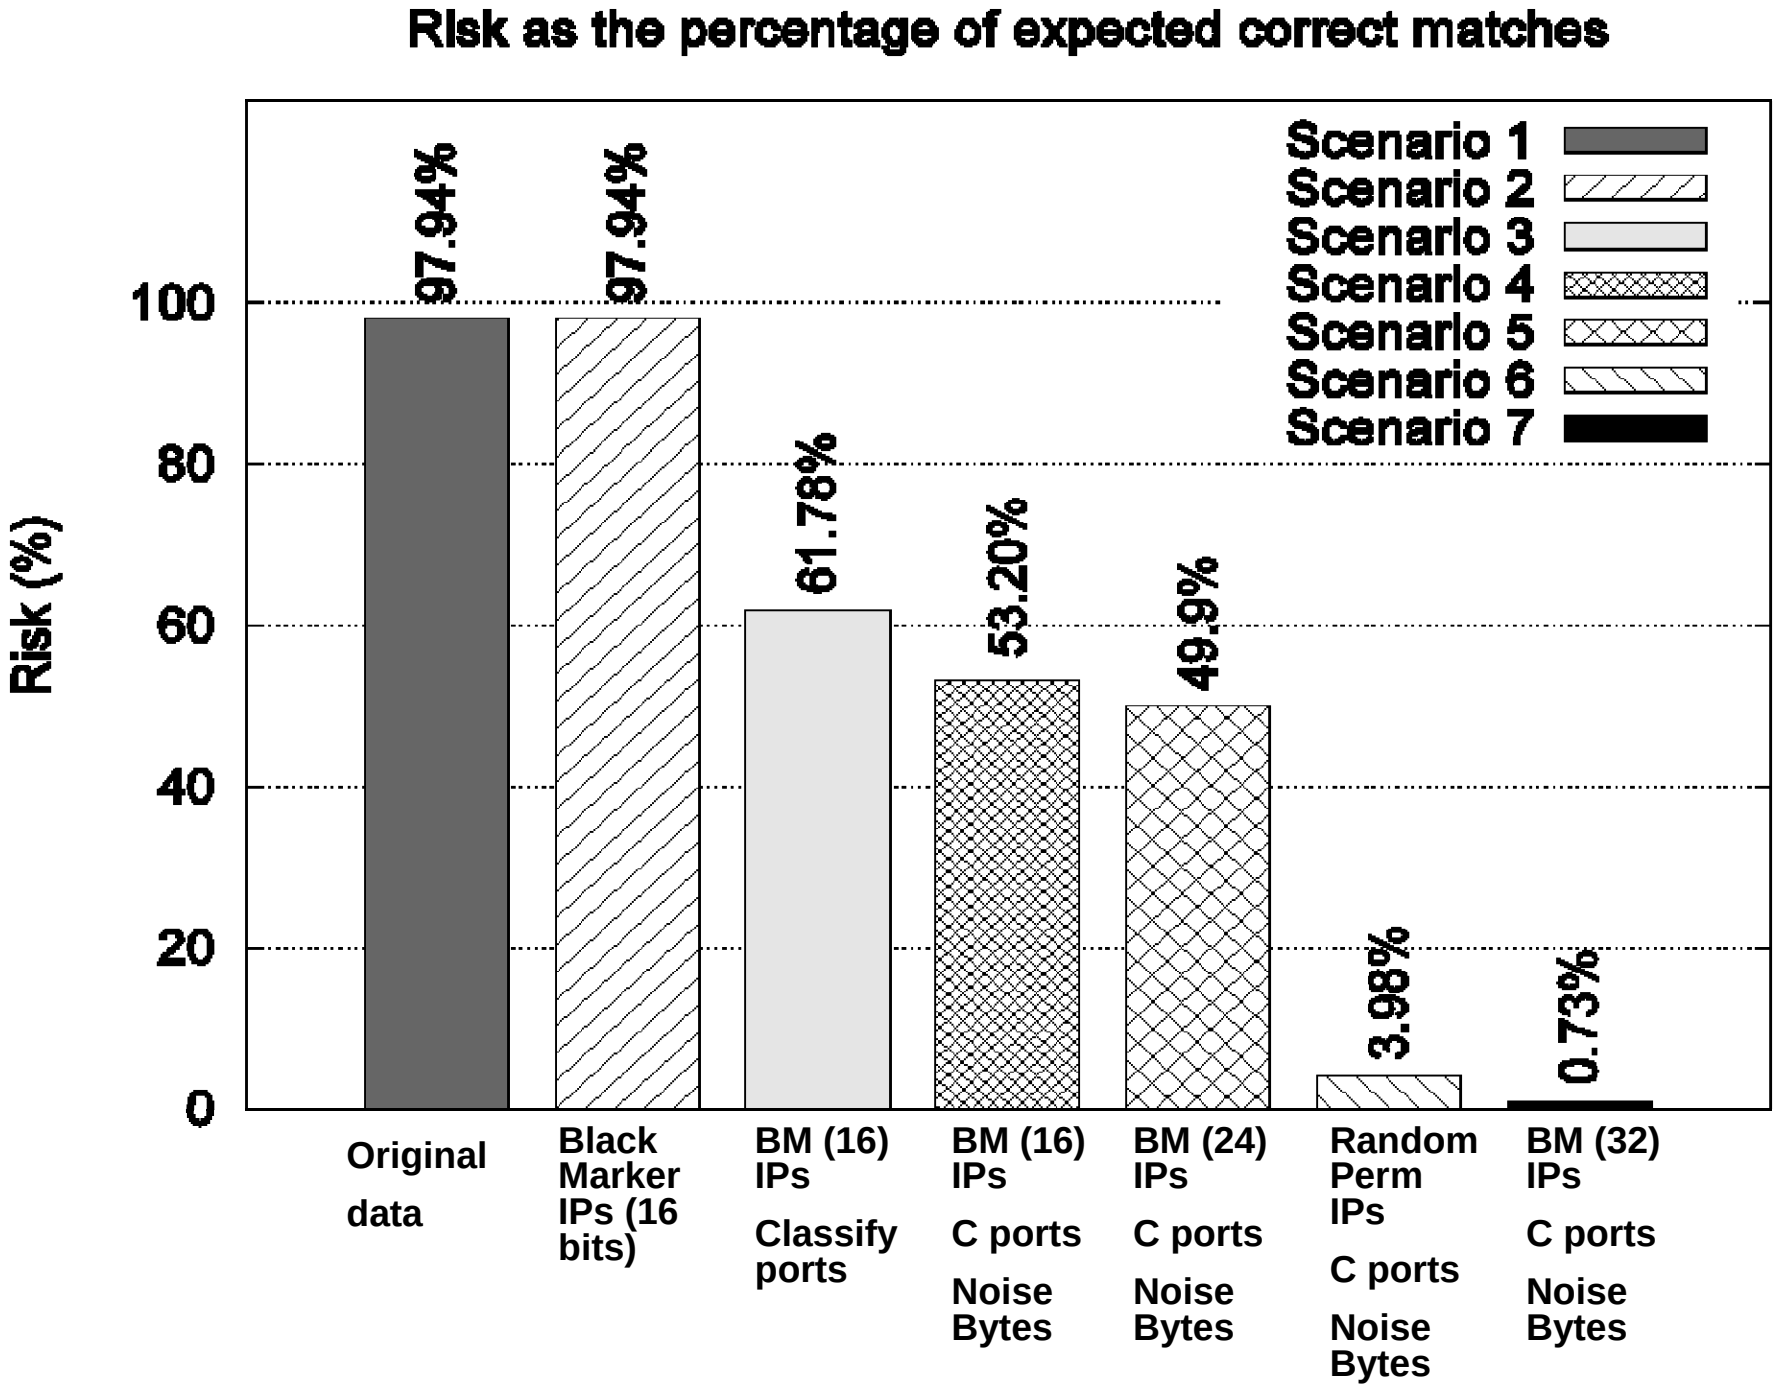
Copyright © 2008 Accenture All Rights Reserved.

Case study: flow

- nfdump testing dataset provided by FLAIM group
- 10000 records
- Src/Dst IPs, Src/Dst ports, Bytes used

Date flow start	Duration	Proto	Src IP Addr:Port		Dst IP Addr:Port	Packets	Bytes	Flows
2007-09-15 21:09:11.401	0.392	TCP	93.82.215.84:36073	->	215.177.13.213:80	8	672	1
2007-09-15 21:09:12.491	0.000	UDP	57.28.244.23:48549	->	204.23.1.67:33467	1	40	1
2007-09-15 21:09:12.431	0.000	UDP	57.28.244.23:48549	->	204.23.1.67:33465	1	40	1
2007-09-15 21:09:12.356	0.354	TCP	89.240.246.94:60717	->	190.0.95.202:3128	7	1253	1
2007-09-15 21:09:12.127	0.000	UDP	154.159.232.119:56395	->	204.23.1.67:33524	1	40	1
2007-09-15 21:09:11.617	0.000	UDP	72.252.1.23:53	->	191.69.116.86:4489	1	165	1
2007-09-15 21:19:20.043	4294216.796	UDP	151.117.100.51:111	->	106.243.186.60:967	5	280	1
2007-09-15 21:19:21.348	1430.067	UDP	111.96.210.161:61718	->	70.114.202.209:161	2	154	1
2007-09-15 21:19:22.694	0.000	UDP	169.53.207.33:53	->	247.215.39.74:3337	1	329	1
2007-09-15 21:19:20.074	0.000	TCP	141.245.94.187:39414	->	217.242.169.109:479	1	60	1
2007-09-15 21:19:21.323	4293905.249	UDP	111.96.210.161:51937	->	80.187.116.29:161	2	154	1
2007-09-15 21:19:21.314	1388.111	UDP	111.96.210.161:53427	->	80.187.116.29:161	3	231	1
2007-09-15 21:19:19.139	0.000	UDP	169.53.207.33:53	->	99.74.24.233:51878	1	284	1
2007-09-15 21:19:19.321	0.000	UDP	169.53.207.33:53	->	99.74.24.233:51879	1	284	1
2007-09-15 21:19:21.321	0.000	UDP	111.96.210.161:53877	->	80.187.116.29:161	1	77	1
2007-09-15 21:19:26.305	4294392.436	UDP	169.53.207.33:53	->	98.14.24.3:50999	2	348	1
2007-09-15 21:19:15.297	69.143	TCP	121.191.230.139:25	->	135.219.55.50:1674	4	291	1
2007-09-15 21:19:21.375	5.023	TCP	103.6.42.145:20144	->	88.118.84.209:51024	552	28712	1

Copyright © 2008 Accenture All Rights Reserved.


Risk as the percentage of expected correct matches
Risk (%)
100
80
60
40
20
0
97.94%
97.94%
61.78%
53.20%
49.9%
3.98%
0.73%
Scenario 1
Scenario 2
Scenario 3
Scenario 4
Scenario 5
Scenario 6
Scenario 7
Original data
Black Marker IPs (16 bits)
BM (16) IPs Classify ports
BM (16) IPs C ports Noise Bytes
BM (24) IPs C ports Noise Bytes
Random Perm IPs C ports Noise Bytes
BM (32) IPs C ports Noise Bytes

Final remarks

Quantifying disclosure risk is essential for finding the optimal trade-off between privacy and utility.

Measure disclosure risk using entropy:

- General: applicable to any anonymization algorithm (unlike k-anonymity)
- Stable: depends on shape of the distribution
- Linked to Information Theory

Future works (a lot…):

- More realistic testing (larger dataset, correlation across fields/records)
- Utility, Optimisation, **…**

Copyright © 2008 Accenture All Rights Reserved.

Thanks for the attention

Michele.bezzi@accenture.com

Copyright © 2008 Accenture All Rights Reserved.

IBM

Zurich Research Laboratory



Simplifying the configuration of flow monitoring probes

Xenofontas (Fontas) Dimitropoulos (xed@zurich.ibm.com)
Andreas Kind (ank@zurich.ibm.com)

Outline

- Background and motivation.
- Probe configuration architecture:
  - Requirements and goals.
  - Design.
  - Implementation.
- Future work and conclusions.

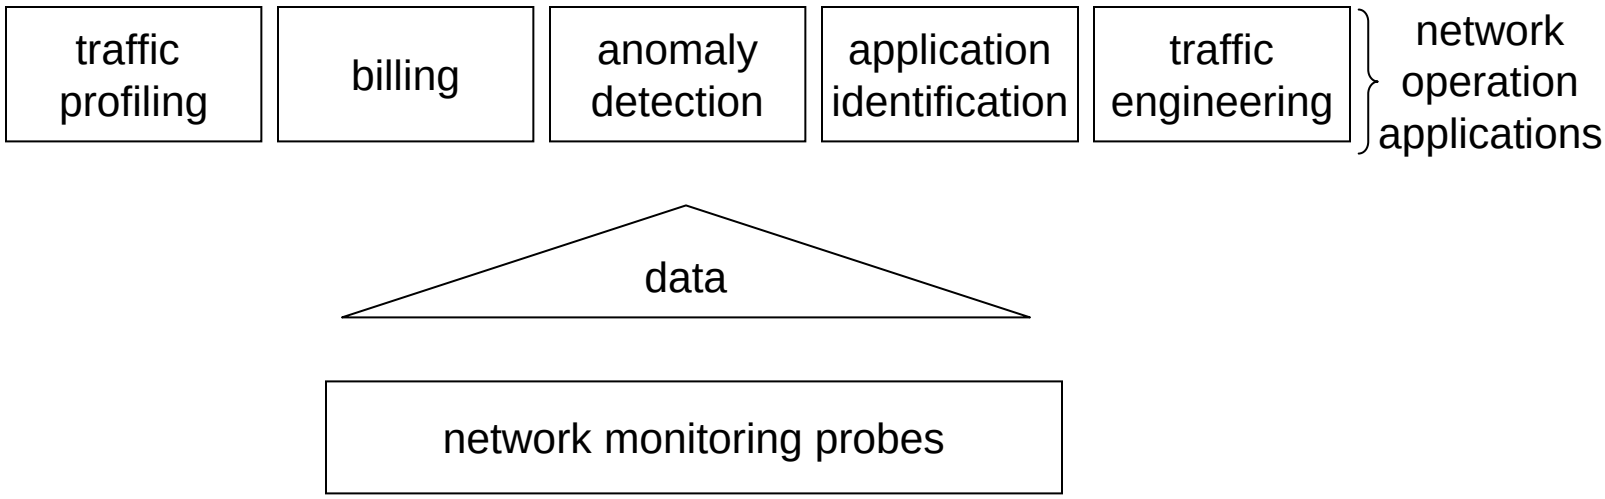
Network configuration

- Network elements are typically configured with low-level commands, e.g., Cisco IOS commands.
- Network administrators manage numerous network elements with lengthy configuration files.
- Network configuration is an error-prone and time-consuming process.
- Configuration errors can be costly, e.g.:
  - network outages
  - violations of SLAs

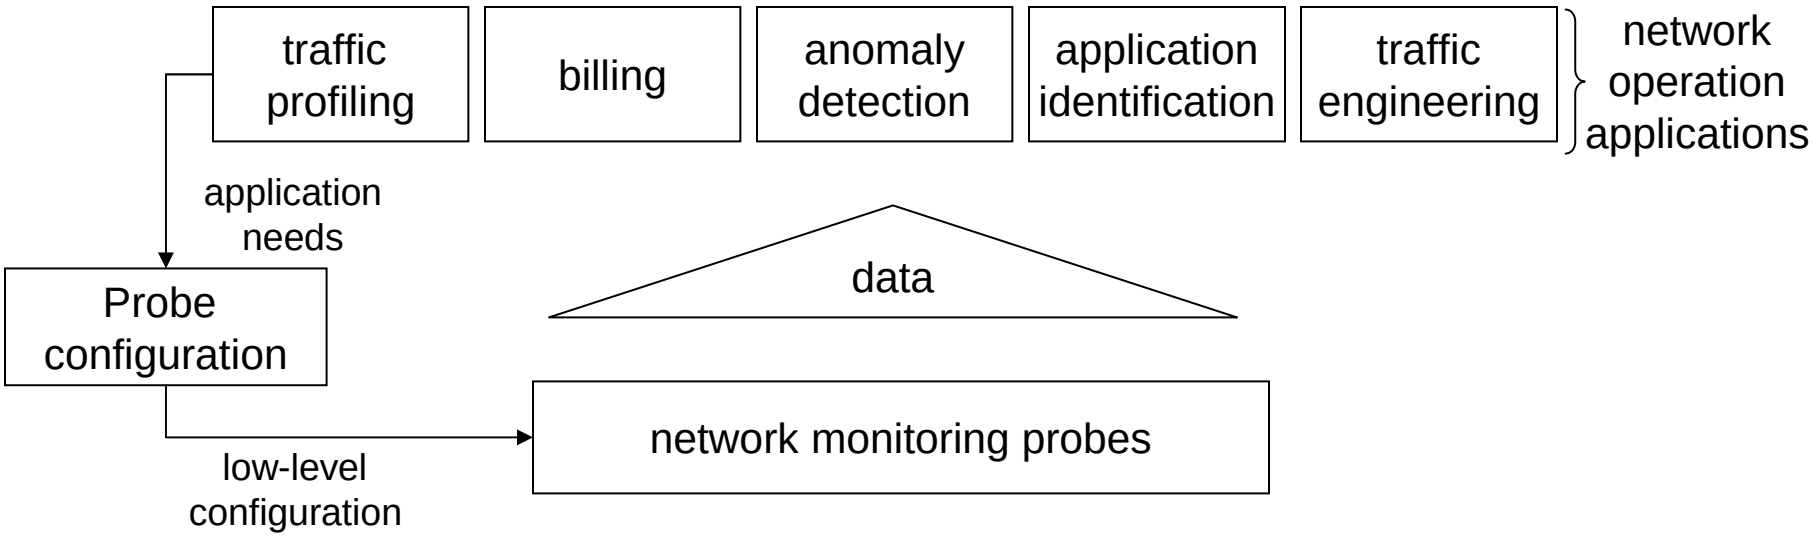
Probe configuration

- The configuration of monitoring probes is part of the more general network configuration problem.
- Monitoring probes are gradually becoming more intelligent, for example, using advanced sampling and data aggregation techniques. Consequently, their configuration becomes more involved.
- Flexible Netflow (FNF) and IPFIX provide numerous configuration options that were not available earlier:
  - FNF has 58 different configuration commands.
  - FNF provides 65 different fields, arbitrary combinations of which can be used in the definition of flow key and non-key fields.
- Certain network operation applications need to dynamically change configuration to:
  - adapt to changing traffic conditions.
  - investigate on-going network anomalies.

Configuration requirements



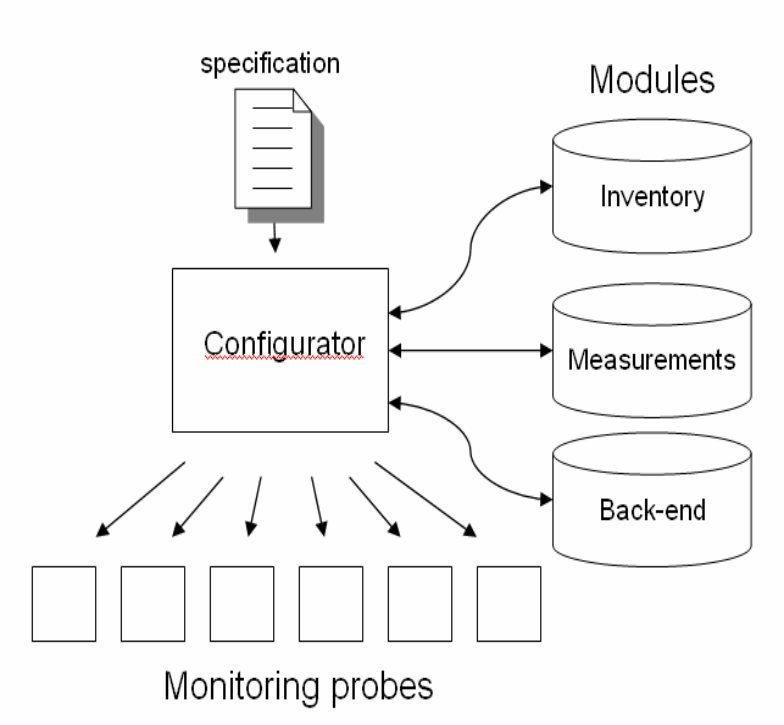
Configuration requirements



- Probe configuration should:
  1. take into account application needs.
  2. be aware of the available monitoring probes.
  3. generate low-level configuration commands.
  4. configure or update the configuration of probes.

Probe configuration architecture

- Three modules:
  - the measurements module describes different measurements, i.e., application needs.
  - the inventory module describes the monitoring probes of a network.
  - the back-end module provides necessary information for generating low-level commands.
- The specification identifies application needs.
- The configurator:
  - uses the modules and specification to generate low-level commands.
  - configures the probes

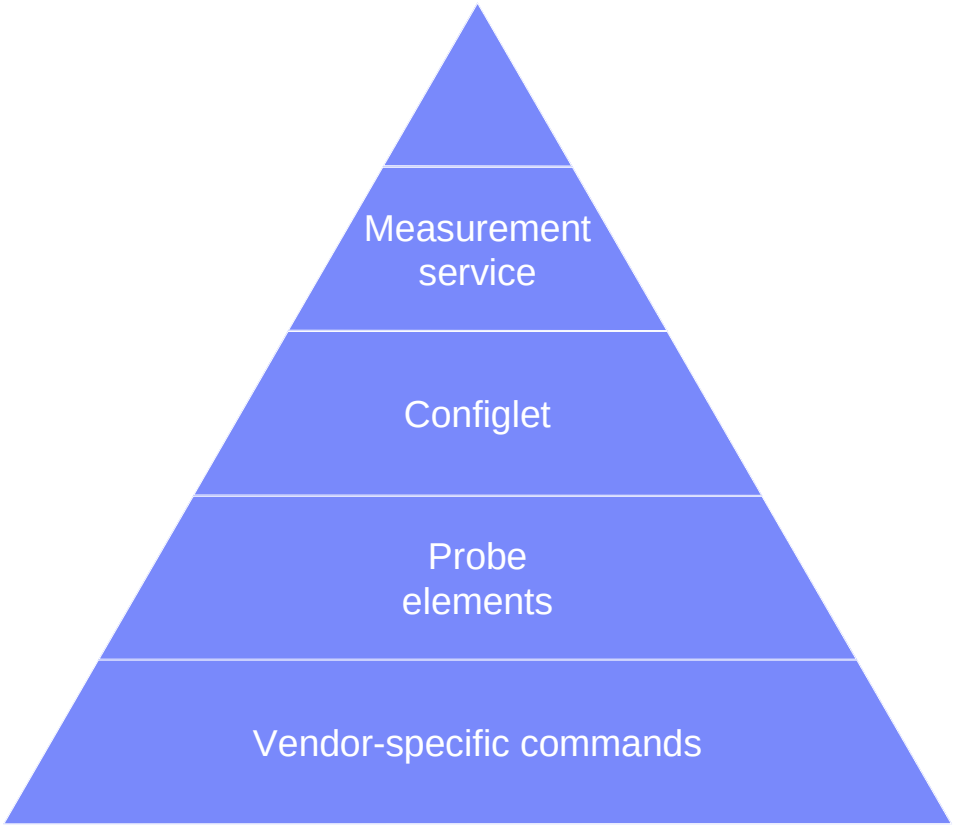


Design goals for simplifying configuration

1. Abstraction: hide low-level configuration commands.
2. Objective-oriented configuration expression:
   - express configuration in terms of measurement objectives.
   - focus on measurements instead of devices.
3. Network-wide configuration: configure a network instead of configuring individual devices.
4. Re-usability: make parts of configuration network-independent.
5. Extensibility: easily introduce support for new commands, measurements, etc.

Configuration abstraction hierarchy

- 1st level: vendor-specific configuration commands.
- 2nd level: probe elements (pe), i.e., logical components of a probe, like interface, flow cache, exporter.
- 3rd level: configlet, i.e., a set of specific probe elements that realizes a measurement.
- 4th level: measurement services, i.e., a configlet with certain probe selection rules.



Back-end module

- Specifies different probe elements.
- A probe element specification:
  - is written in XML.
  - has a unique id.
  - identifies parameters and parameter default values.
  - determines the low-level vendor-specific commands.

```
<!– Probe Element Exporter -->
<pe id='generic_exporter'>

 <params>
  <param id='port'>90</param>
  <param id='transport'>udp</param>
  <param id='destination'>192.0.0.1</param>
  <param id='label'>EXPORTER</param>
 </params>

 <template>
  <ios>
        flow exporter $label
        destination $destination
        transport $transport $port
  </ios>
  <yaf>
        --out $destination --ipfix $transport --ipfix-port $port
  </yaf>
  <junos>
  </junos>
 </template>

</pe>
```

Inventory module

- Specifies network probes, i.e., lists the characteristics that can be useful for their configuration.
- Besides describing location, system, and interface information, it declares tags that can be used for grouping probes and for probe selection.

```
<probe id='trabant.zurich.ibm.com'>
  <address>9.4.68.154</address>

  <location>
    <city>Zurich</city>
    <state>Central CH</state>
    <country>Switzerland</country>
  </location>

  <system>
    <os>ios</os>
    <version>12.4</version>
  </system>

  <interface id='FastEthernet0/0'>
    <capacity>100Mbits</capacity>
    <tag>internal</tag>
  </interface>

  <interface id='FastEthernet0/1'>
    <capacity>100Mbits</capacity>
    <tag>customer</tag>
  </interface>

  <tags>
    <tag>edge</tag>
  </tags>

 </probe>
```

Measurements module

```
<!--   Monitor how much traffic is send -->
<!--   between IP blocks. -->
<msr id='traffic_matrix'>

  <params> <!-- Default parameter values -->
    <param id='collector_address'>localhost</param>
    <param id='collector_port'>2055</param>
    <param id='collector_transport'>tcp</param>
  </params>

  <!-- Probe element chain -->
  <configlet>
  </configlet>

  <rules>
  </rules>

</msr>
```

Measurements module

```
<!-- Probe element chain -->
<configlet>
   <pe>
            <name>exporter</name>
            <params>
              <param id='label'>TM_EXPORTER</param>
              <param id='destination'>$collector_address</param>
              <param id='port'>$collector_port</param>
              <param id='transport'>$collector_transport</param>
            </params>
   </pe>
   <pe>
            <name>flow_cache</name>
            <params>
              <param id='label'>TM_CACHE</param>
              <param id='record'>SRC_DST_PREFIX_REC</param>
              <param id='export'>TM_EXPORTER</param>
            </params>
   </pe>
   <pe>
            <name>interface</name>
            <params>
              <param id='monitor'>TM_CACHE</param>
              <param id='interface'>$interface->id</param>
              <param id='direction'>output</param>
            </params>
   </pe>
</configlet>
```

Measurements module

```
<rules>
    <interface>
                if ( $interface.tag eq “external" and
                     $probe.tag eq "edge" ) {
                    return 1;
                } else {
                    return 0;
                }
    </interface>
</rules>
```

Input specification

- Lists the measurements and the probes in which to enable these measurements.
- Is the user interface and can be generated through a GUI.

```
<!-- Probes to apply measurements on -->
<probe id='wassen.zurich.ibm.com'></probe>
<probe id='trabant.zurich.ibm.com'></probe>


<!-- Measurements -->
<msr id='traffic_matrix'>
   <params> <!-- overwrite default values -->
     <param id='collector_address'>9.4.68.204</param>
     <param id='collector_port'>2055</param>
     <param id='collector_transport'>udp</param>
   </params>
</msr>

<msr id='app_monitoring'>
   <params> <!-- overwrite default values -->
     <param id='collector_address'>9.4.68.205</param>
     <param id='collector_port'>2055</param>
     <param id='collector_transport'>udp</param>
   </params>
</msr>
```

Design goals for simplifying configuration

1. Abstraction: hide low-level configuration commands.
2. Objective-oriented configuration expression:
   - express configuration in terms of measurement objectives.
   - focus on measurements instead of devices.
3. Network-wide configuration: configure a network instead of configuring individual devices.
4. Re-usability: make parts of configuration network-independent.
5. Extensibility: easily introduce support for new commands, measurements, etc.

Conclusions

- Described an architecture for automating the configuration of flow monitoring probes.
  - Configuration abstraction.
  - Reuse configuration.
  - Extensibility.
- Future/on-going work:
  - Incorporate error-checking techniques.
  - Develop libraries for typical measurements.
  - Configuration optimization.
  - Use NetConf.

High Level Flow Correlation

Valentino Crespi, California State Los Angeles, CA
Annarita Giani, UC Berkeley, CA
Rajiv Raghunarayan, Cisco Systems, Inc.

FloCon 2008, Savannah GA, January 7-10, 2008.

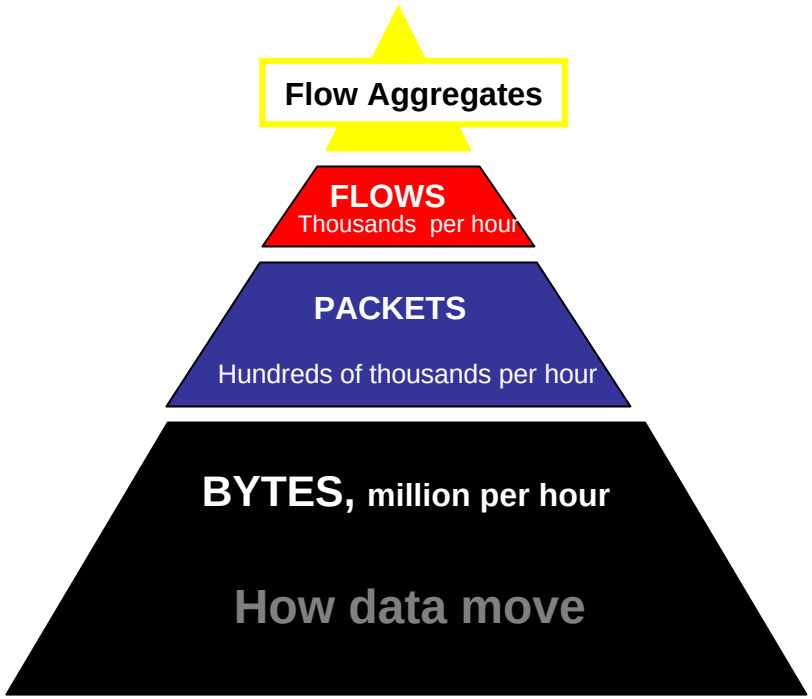
Outline

1. Extension of previous work on Flow Aggregation, (Flocon 2006).
2. Embedding of network traffic in an Euclidian Space.
3. Complex modeling through clustering.
4. Planned work.

Outline

1. Extension of previous work on Flow Aggregation, (Flocon 2006).
2. Embedding of network traffic in an Euclidian Space.
3. Complex modeling through clustering.
4. Planned work.

Behind Flow Aggregation



- Monitoring
- Anomaly detection
- Security analysis
- Traffic profiling
- Debugging
- Traffic engineering
- Usage-based profiling
- Network planning
- Pricing, peering

Data Reduction = Fewer events to be analyzed

Our Previous Work

A. Giani, I. De Souza, V. Berk, G. Cybenko, " Attribution and Aggregation of Network Flows for Security Analysis ," *in Proc. Flocon 2006,* Portland, OR.

We believe that **automated correlation at the raw flow level** is complicated and susceptible to false positives. The world consists of **processes** so our approach to correlation is process-based..

Flow aggregation and correlations between flow data with security events

Implementation of a **PQS based process detection for Cyber Situational Awareness.**

Flow + Snort Alerts

Scenario: several packets in a flow triggered IDS alerts

Snort rule **1560** generates an alert when an attempt is made to exploit a known vulnerability in a web server or a web application.

Snort rule **1852** generates an alert when an attempt is made to access the 'robots.txt' file directly.

Timestamp	Sensor	src IP	dst IP	Proto
Jul 09 16:28:32	S1852	65.54.188.140	208.253.154.195	TCP
Jul 09 16:29:35	S1852	65.54.188.140	208.253.154.195	TCP
Jul 09 16:44:44	S1560	65.54.188.140	208.253.154.195	TCP
Jul 09 18:26:08	S1560	65.54.188.140	208.253.154.195	TCP
Jul 09 21:05:03	S1852	65.54.188.140	208.253.154.195	TCP
Jul 09 22:31:08	S1852	65.54.188.140	208.253.154.195	TCP
Jul 09 22:31:08	S1560	65.54.188.140	208.253.154.195	TCP
Jul 10 02:45:19	S1852	65.54.188.140	208.253.154.195	TCP
Jul 10 02:45:23	S1852	65.54.188.140	208.253.154.195	TCP
Jul 10 09:21:15	S1852	65.54.188.140	208.253.154.195	TCP
Jul 10 14:33:43	S1852	65.54.188.140	208.253.154.195	TCP
Jul 10 17:54:54	S1852	65.54.188.140	208.253.154.195	TCP
Jul 10 22:07:02	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 01:38:09	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 04:05:54	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 04:20:00	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 04:20:00	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 11:07:12	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 11:56:12	S1852	65.54.188.140	208.253.154.195	TCP
Jul 11 17:16:59	S1852	65.54.188.140	208.253.154.195	TCP
S Jul 10 02:30:27	F	65.54.188.140	208.253.154.195	TCP
E Jul 10 23:55:56				

SNORT ALERTS

→ FLOW

Table 2: A sample track of correlated IDS and Flow events

The flow can be characterized as malicious and further investigation must be done.

Outline

1. Extension of previous work on Flow Aggregation, (Flocon 2006).
2. Embedding of network traffic in an Euclidian Space.
3. Complex modeling through clustering.
4. Planned work.

Current aggregators and analyzers

- **POWERFUL TOOLS** to understand the behavior of the network according to certain parameters, e.g. the amount of resources consumed, the variance on the various characteristics of the communication (source ip, destination ip), port.
- **PROBLEM**: They do not provide an analysis and a description of the dynamic evolution of network traffic.
- **NEED** for a structure that summarizes the behavior of the network.

OUR IDEA

Combine flow aggregation techniques with our previous process-based approach:

Use aggregators and flow analyzers to translate traffic into a process to be modeled and estimated.

Build circuits of Aggregating gates

1. Place observing nodes in multiple locations of the network (e.g. on each local router).
2. Each observing nodes dumps traffic flows to a Macro Aggregator (MA).
3. Macro Aggregator: *circuit*. Each gate is a flow aggregator
   - First layer consists of classical aggregators that output flow aggregates. Successive layers process aggregates of flow aggregates
   - Final output: a vector function of the dumped traffic ranging in $\mathbf{R}^n$:

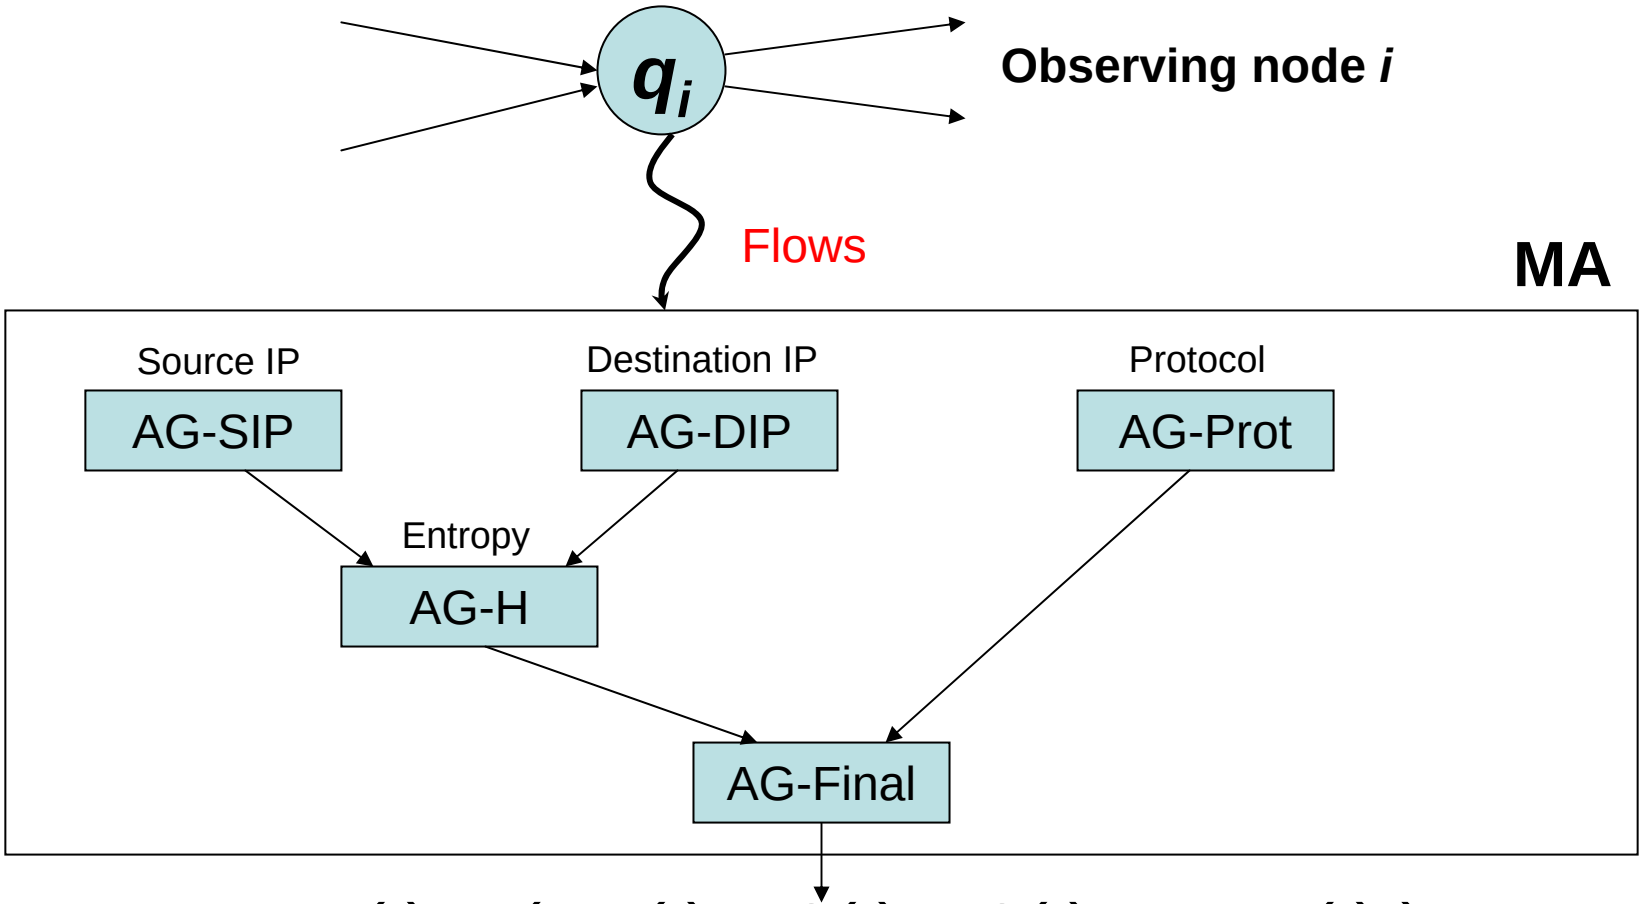
$$\mathbf{X}(t) = (x_1(t), x_2(t), \cdots, x_n(t))$$

   At each time the observing nodes produce a set of vectors:

$$\mathbf{S}(t) = (X_1(t), X_2(t), \cdots, X_n(t))$$

4. Identify and Analyze properties of S(t) over time to characterize/detect anomalies.

Embed Traffic in Euclidean Space



$$\mathbf{X}_i(t) = (x1(t), x2(t), x3(t),\ldots, xn(t))$$

(Entropy S-IP,Entropy D-IP, Average Size,…,%TCP Traffic,%UDP Traffic)

Entropy Based Flow Aggregation (2006)

Yan Hu, Dah-Ming Chiu, and John C.S. Lui
The Chinese University of Hong Kong

Based on Cisco's NetFlow – during flooding attacks the memory and network bandwidth consumed by flow records can increase beyond what is available.
A solution: Adapting sampling rate.

Flows of security attacks usually have common patterns and form conspicuous traffic clusters.
Identifies clusters of attacks flows in real time and aggregated those large number of short attack flows to a few meta flows.

Same sourceIP ~ worm propagation
Same destIP ~ Denial of Service Attack
Same destIP and SourceIP ~ most portscan

Purpose is mostly security.

On the correlation of Internet flow characteristics (2003)

Kun-Chan Lan, JOHN HEIDEMANN
Information Science Institute, University of Southern California

A small percentage of flows consume most of the network bandwidth.

Study of heavy flows in 4 orthogonal dimensions:

- Size
- Duration
- Rate
- Burstiness

and examine their correlations.

Strong correlation between size, rate, burstiness

Automatically Inferring Patterns of Resource Consumption in Network Traffic (2003)

Cristian Estan, Stefan Savage, George Varghese
University of California, San Diego

Method of traffic characterization that automatically groups traffic into minimal clusters of conspicuous consumption.

It is not a static analysis that captures flow characteristics but instead produces hybrid traffic definition that match the underline usage.

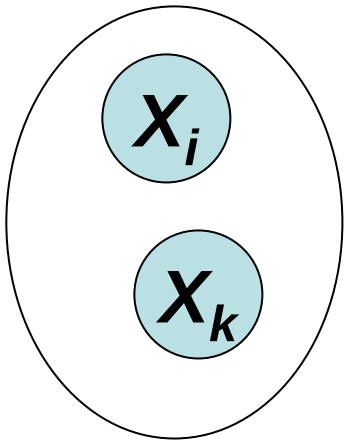
Purpose is mostly resource consumption.

Analyze *S(t)* over time

Approaches:

1. Use clustering techniques (e.g., spectral clustering, k-means based algorithms, etc.) to clusterize the observing nodes and infer correlations between observations and snapshots across the network.
    1. Study how clusters change over time and characterize/detect anomalies.
    2. Use clusters to produce a graphic representation of the traffic.
    3. Define discrete models to describe the evolution of clusters in relation to specific events: coordinated computer attacks, presence of covert channels, bugs in the network software, hardware breakdowns, etc.
2. Define State Space models.
3. Apply learning techniques to learn models.

Spectral Clustering

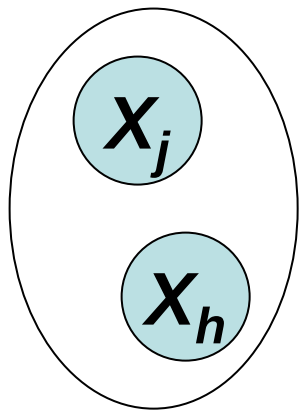


Input: Similarity Matrix M=[aij], , number k>0

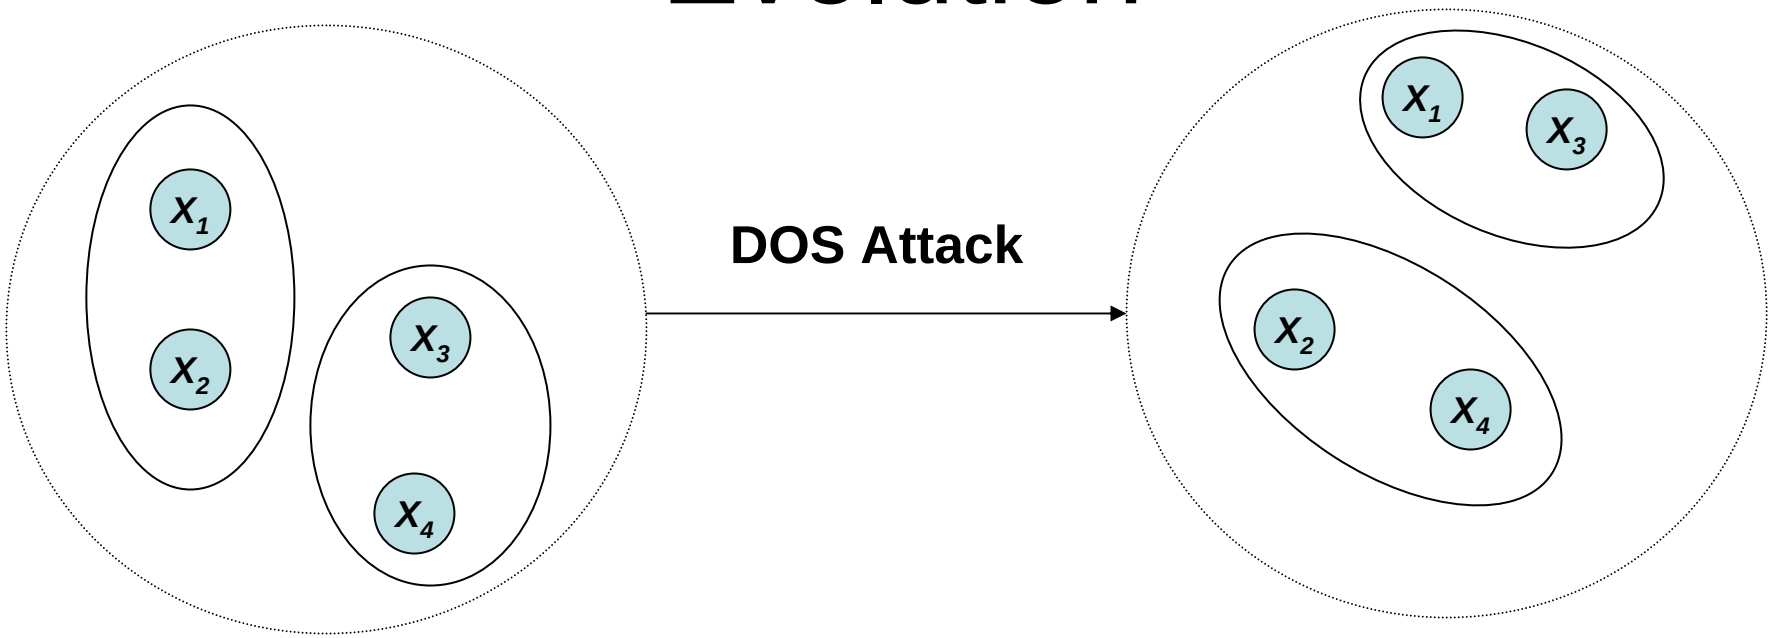
$a_{ij} = s(X_i, X_j)$ e.g $a_{ij} = \exp(-\|X_i - X_j\| / 2\sigma^2)$

- Build similarity graph. For example the Graph whose adjacency matrix AG = M.
- L = Laplacian(AG)
- Compute the k eigenvectors of L associated with the k smallest eigenvalues: v1, v2,…,vk
- V = [v1 v2 … vk], nxk matrix
- Pick the rows of V: y1, y2,…,yn
- Cluster yi's using k-means algorithm into C1,C2,…Ck

Output: clusters C1,C2,…,Ck



Discrete Models of Cluster Evolution



Idea: Build DFA models to identify transitions. In this case we identify anomalies by studying the current clustering in relation to the previous “snapshot” of traffic

Challenges

- Parameter estimation: in our example of clustering k was fixed.
- Apply Bayesian learning techniques to infer k.
- Apply *mixture models* technique to clustering
- Define and learn models of the system's dynamics.
- Identify relevant attributes of flow aggregators to obtain significant vectors.
- Define appropriate similarity function.

Outline

1. Extension of previous work on Flow Aggregation, (Flocon 2006).
2. Embedding of network traffic in an Euclidian Space.
3. Complex modeling through clustering.
4. Planned work.

Planned Work

- Implement clustering method.
- Develop discrete models.
- Build a software monitor to analyze traffic through clusters and vector representation.
- Experimental analysis of the efficaciousness of our approach.

References

[1] I. S. Dhillon, Y. Guan, and B. Kulis. Kernel k-means, Spectral Clustering and Normalized Cuts. In *Proceedings of the KDD'04 Workshop*, Seattle, Washington, August 2004.

[2] C. Estan, S. Savage, and G. Varghese. Automatically Inferring Patterns of Resource Consumption in Network Traffic. In *Proceedings of the 2004 SIGCOMM*.

[3] A. Giani, I. G. D. Souza, V. Berk, and G. Cybenko. Attribution and Aggregation of Network Flows for Security Analysis. In *Proceedings of FloCon 2006*.

[4] Y. Hu, D.-M. Chiu, and J. C. Lui. Adaptive Flow Aggregation - A New Solution for Robust Flow Monitoring under Security Attacks. In *Proceedings of 2006 IEEE/IFIP Network Operations and Management Symposium (IEEE/IFIP NOMS)*.

[5] Y. Hu, D. Chui, and J. C. Lui. Adaptive Flow Aggregation - a New Solution for Robust Flow Monitoring under Security Attacks. In *Proceedings of the 2006 Network Operations and Management Symposium*.

[6] K. Keys, D. Moore, and C. Esten. A Robust System for Accurate Real-time Summaries of Internet Traffic. In *Proceedings of the 2005 SIGMETRICS*, June 2005.

[7] L. Rodrigues and P. R. Guardieiro. A Spatial and Temporal Analysis of Internet Aggregate Traffic at the Flow Level. In *Proceedings of the 2004 Global Telecommunications Conference (GLOBECOM)*, volume 2.

[8] B. Trammell and C. Gates. Naf: The NetSA Aggregated Flow Tool Suite. In *Proceedings of the Large Installation System Administration Conference (LISA 2006)*, 2006.

[9] U. von Luxburg. A Tutorial on Spectral Clustering. Technical Report TR-149, Max-Planck-Institut für biologische Kybernetik, 2006.

Thanks

Annarita Giani <agiani@eecs.berkeley.edu>
Valentino Crespi <vcrespi@calstatela.edu>
Rajiv Raghunarayan <raraghun@cisco.com>



Attack Reducation and Anomaly Modeling in Popularly Targeted Protocols

Michael Collins, CERT/NetSA

© 2007 Carnegie Mellon University

Talk outline

The Problem

- Noise in traffic flows
- Impact on anomaly detection

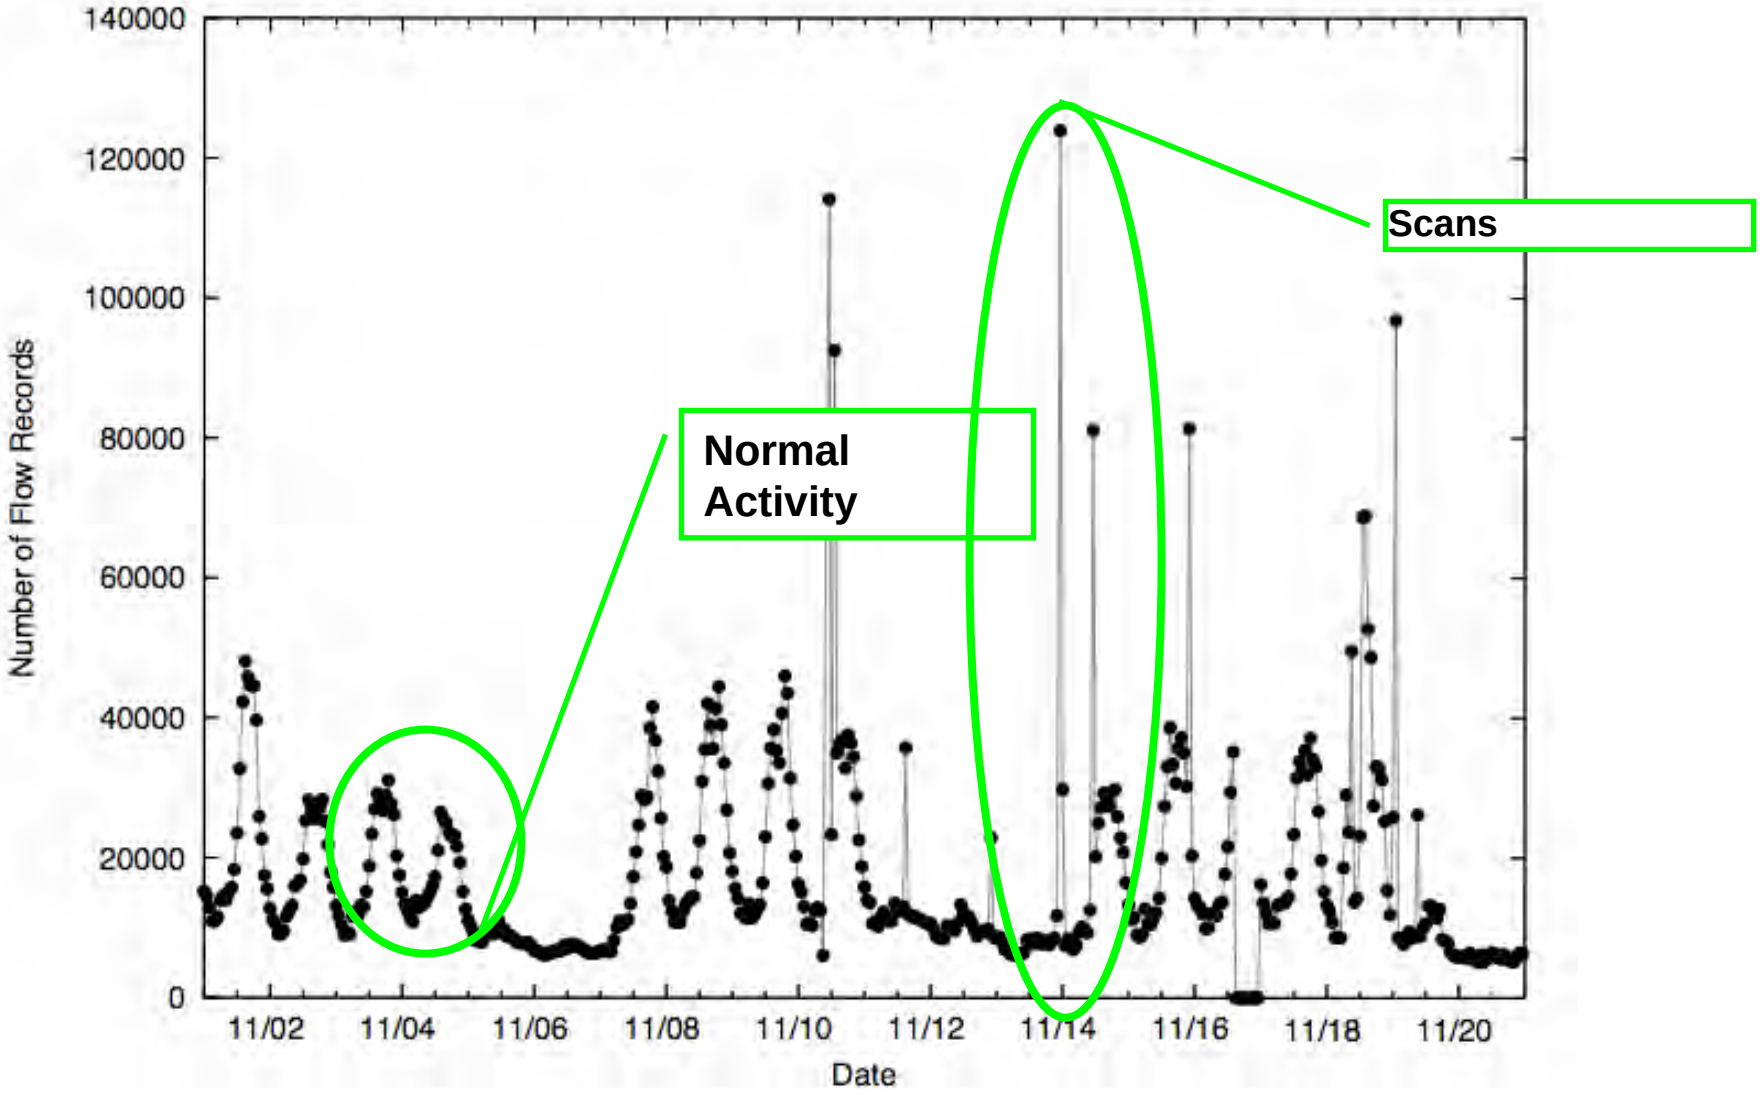
Two Stage Filtering

- Log Filtering
- State Filtering

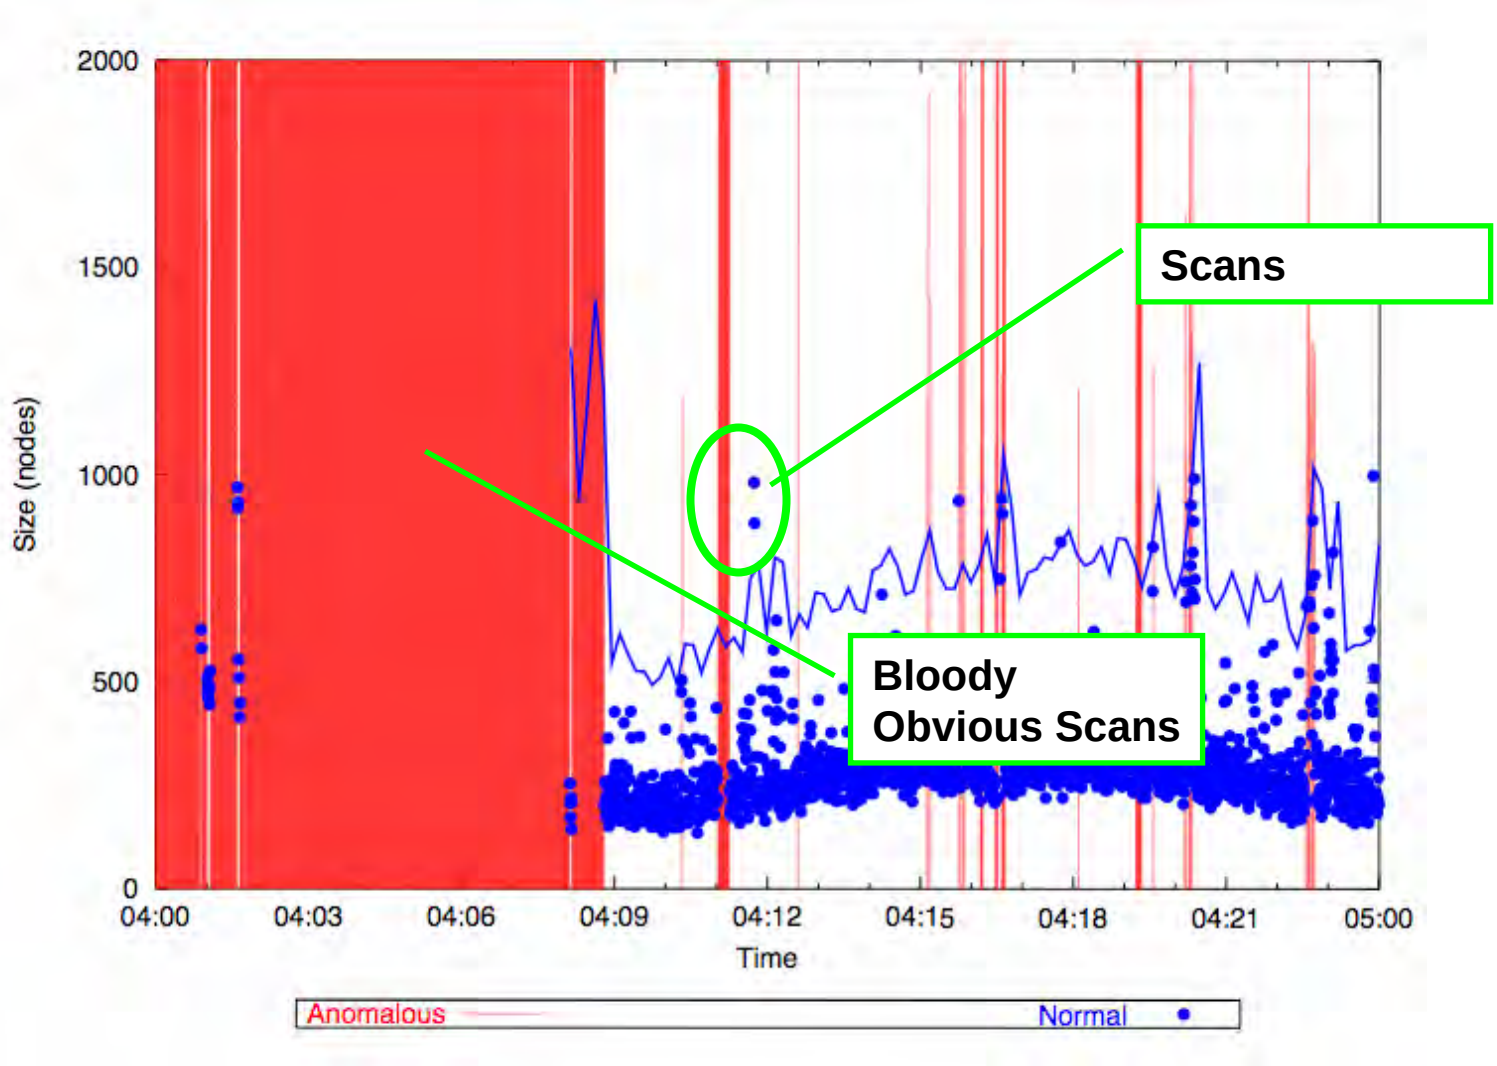
Attack reduction

- Core assumptions
- Method for data removal
- Impact

Innocuous Attacks



Normal SSH Activity



Raw SSH Data



A Hypothesis

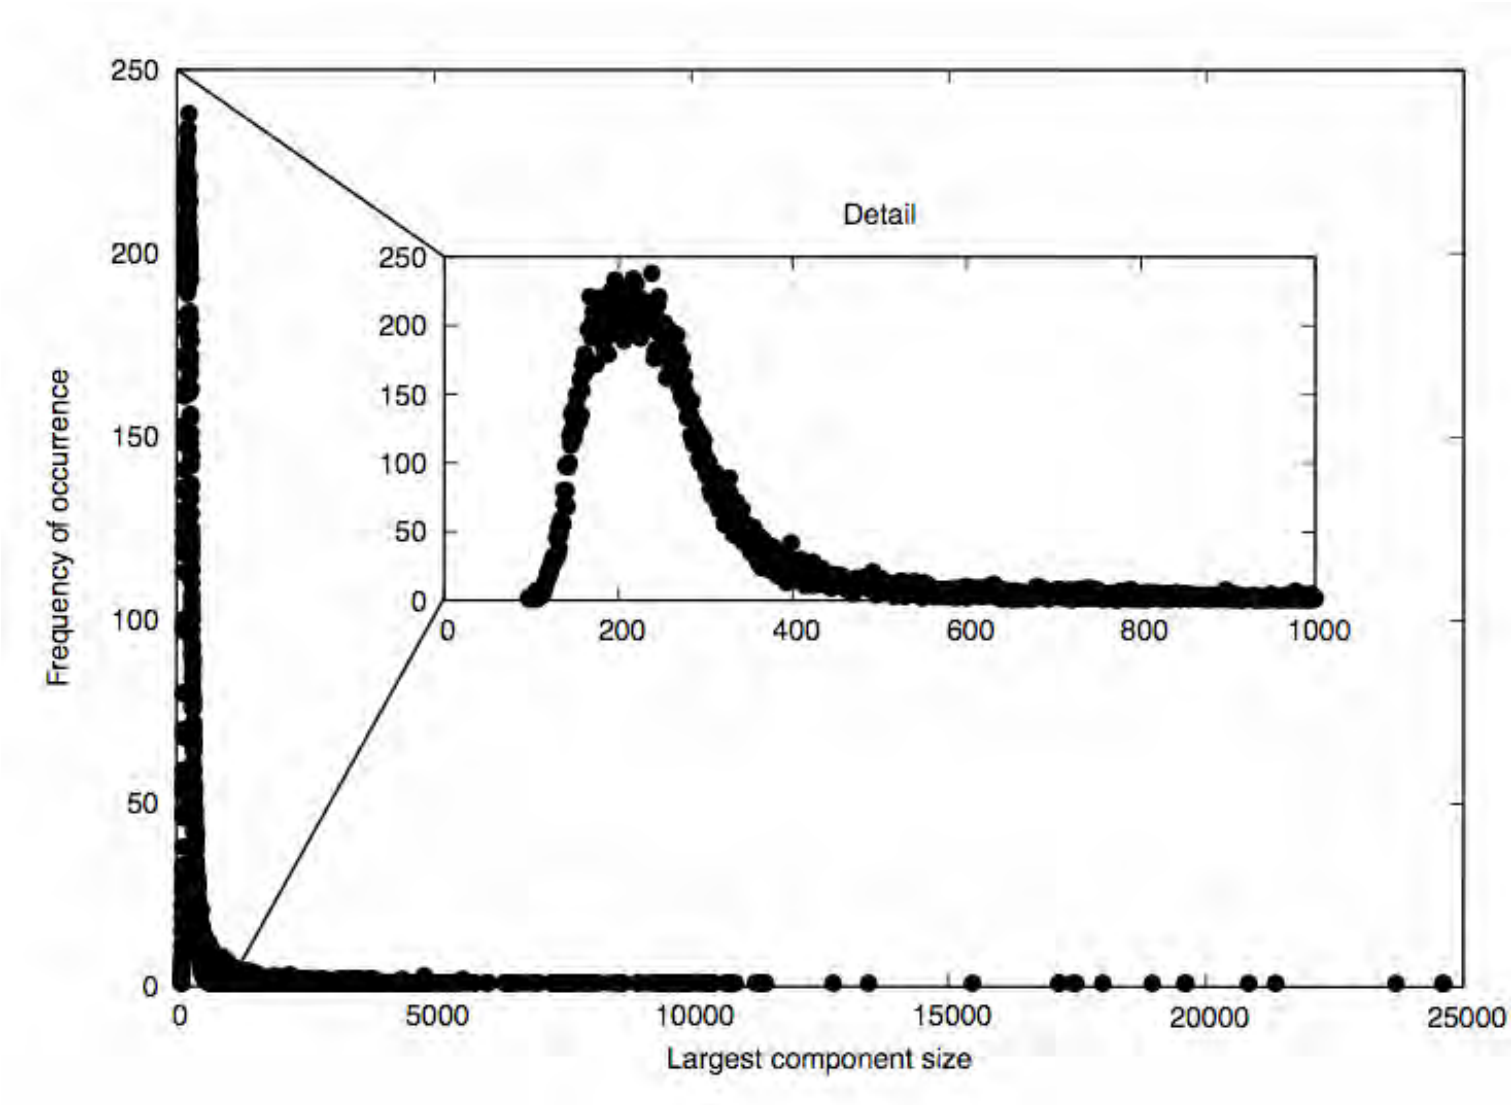
We see two populations:

- Normal users, who know where they're going
- Attackers, primarily scanners, who have no idea about the network's structure

The majority of attackers are clumsy

- Low success rates
- Picking targets effectively at random
- Pick many more targets than there are actual targets
  - >350,000 per 30s period, vs. ~ 10,000 real targets

Comparing the two populations



Impact on anomaly detection

Almost every anomaly detection system requires advance knowledge

- Mean, standard deviations
- Map of known servers

This information may not be easily acquired

- Inventory is nontrivial
- Going by the data can lead to false positives from attackers

We need to train the system while acknowledging the hostility

Filtering: Log Filtering

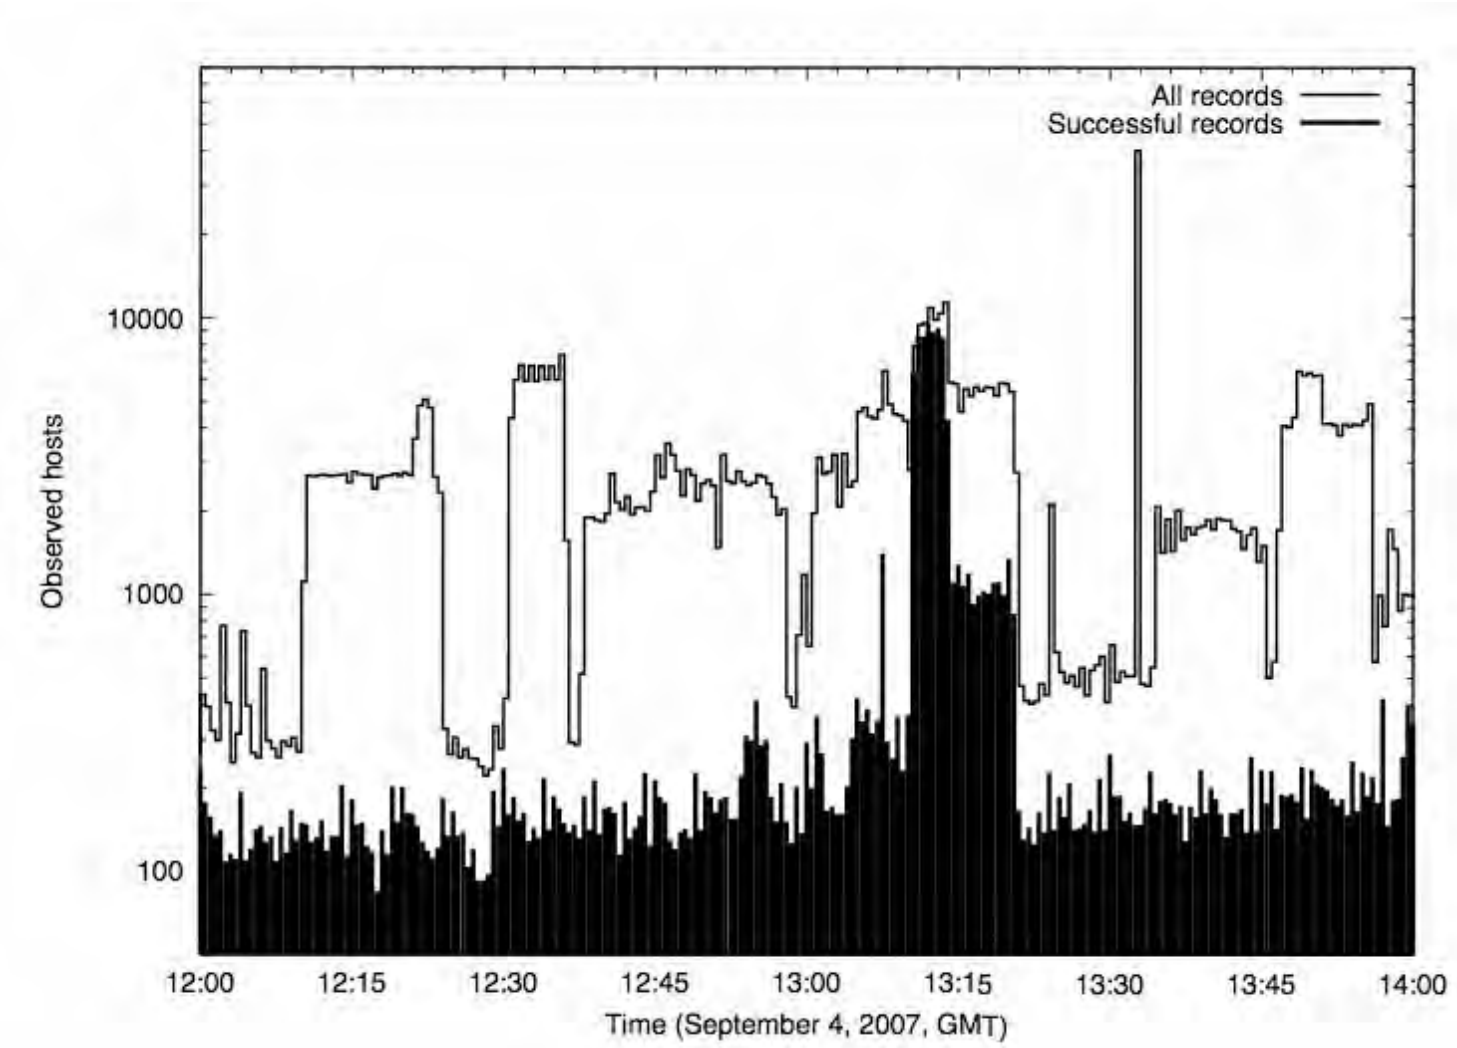
At least with TCP data, we can rely with the state machine

- ≤ 3 packets implies it is most likely a scan
- > 3 packets may be legitimate

In a two week ssh dataset:

- ≤ 3 packets make up 87% of the flows
- ≤3 packets make up 1% of total bandwidth

Log Filtering is Insufficient



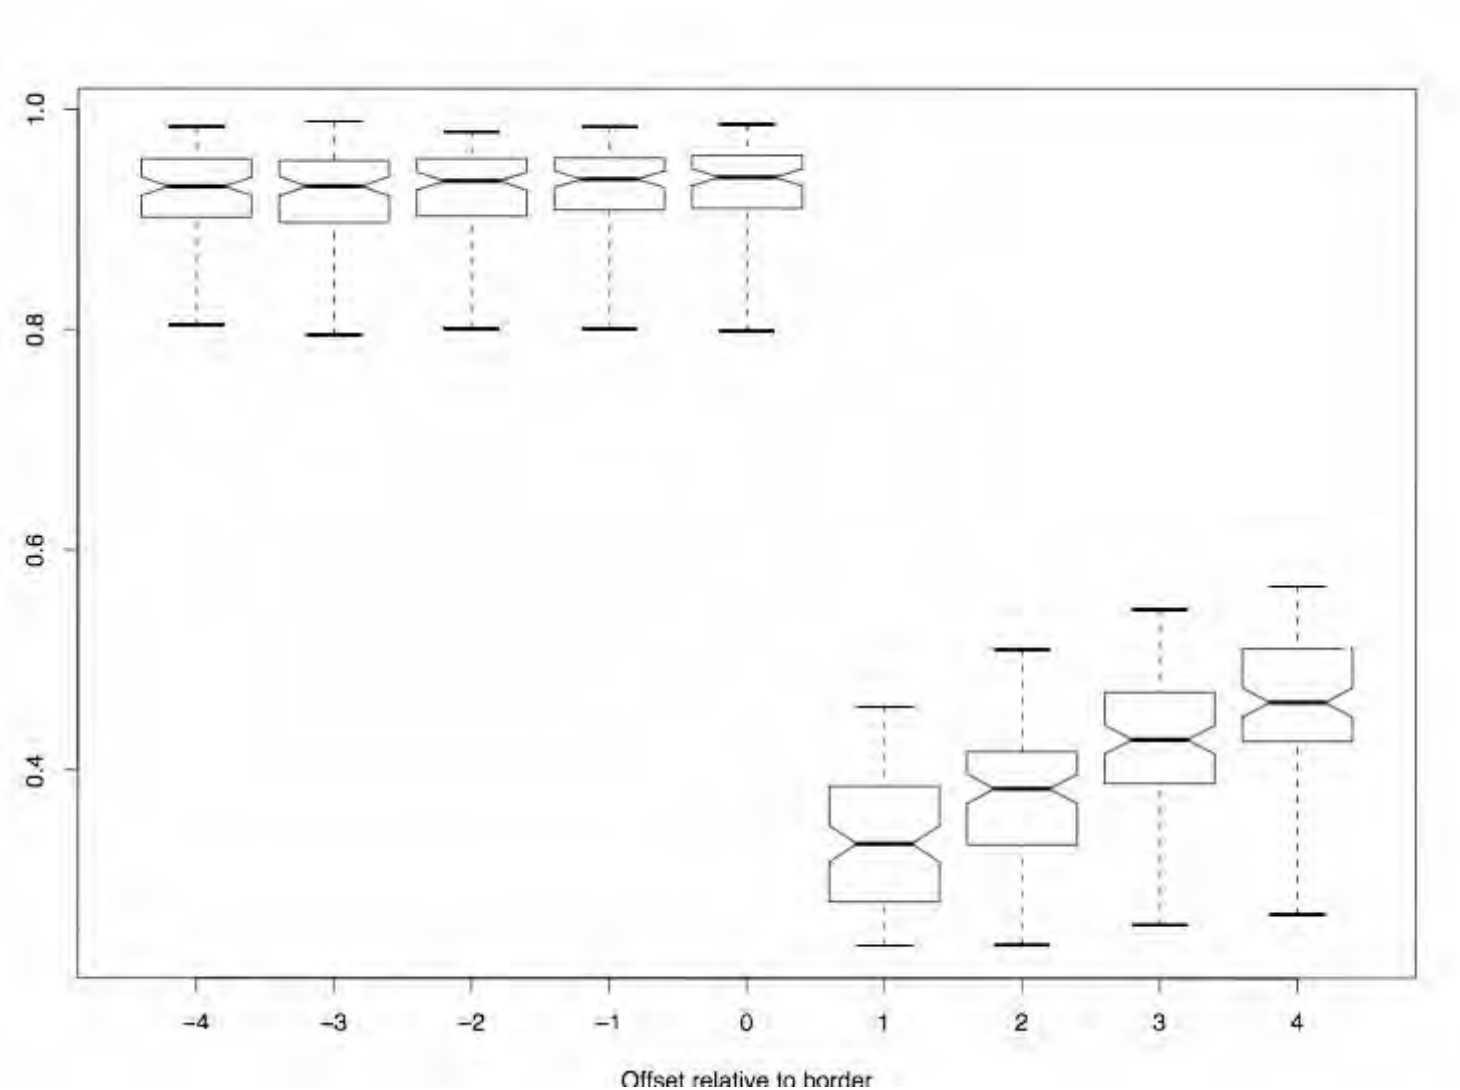
State Filtering

If we assume activity is Gaussian, then we can identify and eliminate outliers

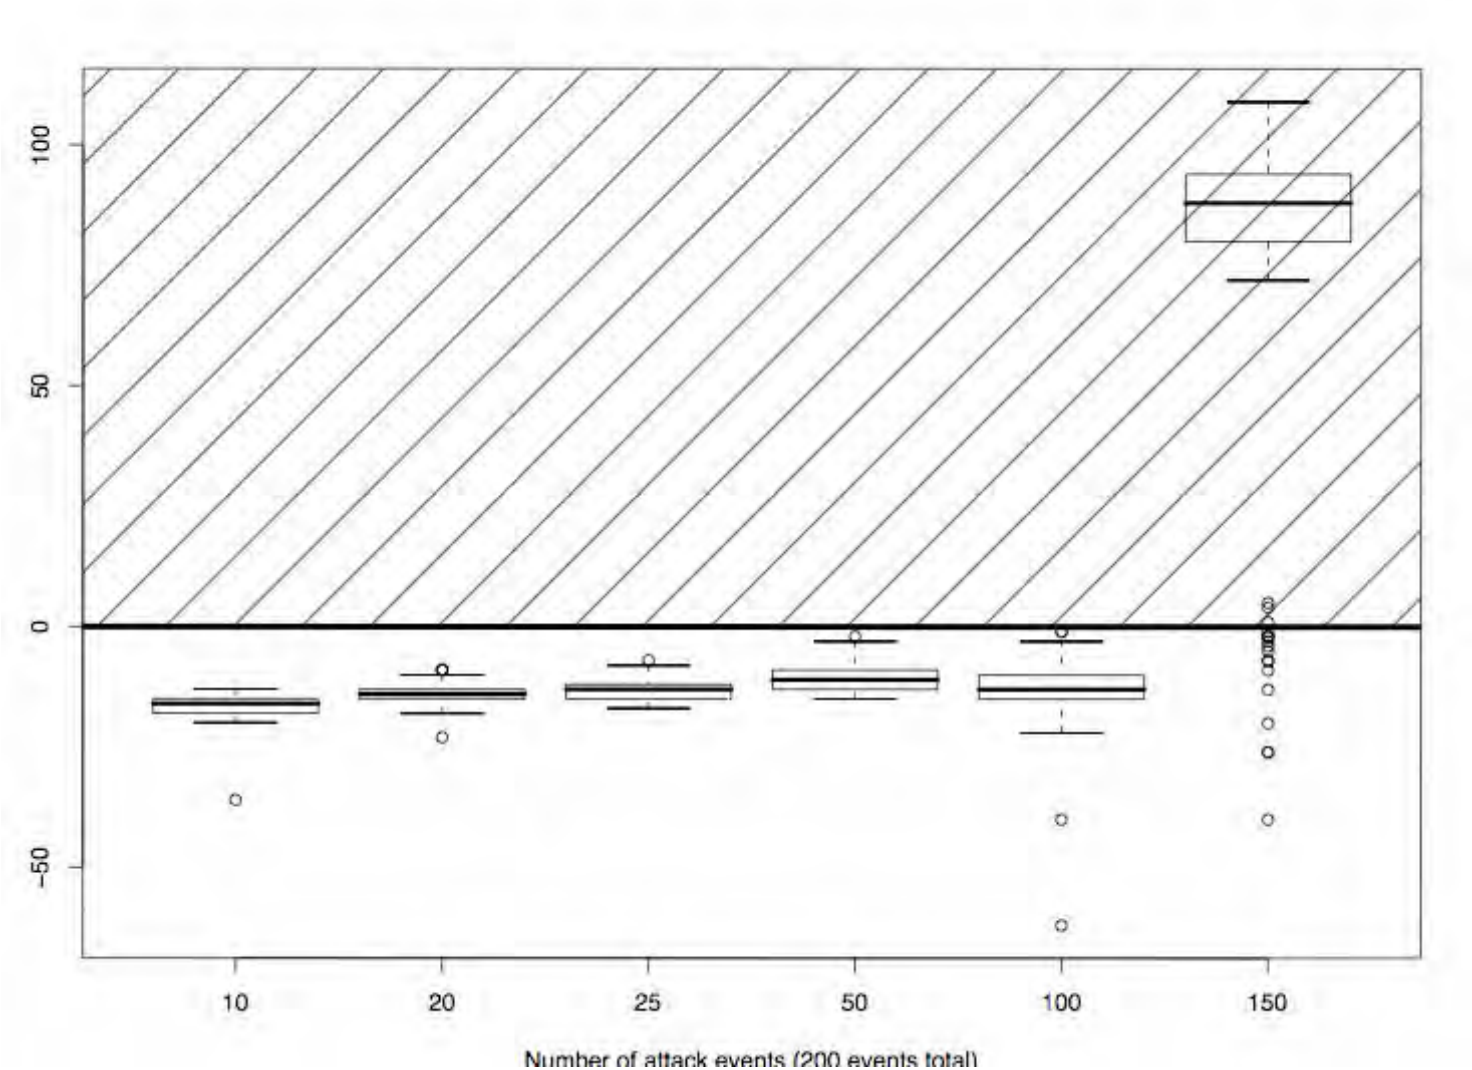
Simple test: Shapiro-Wilk test for normalcy

- God for 25-2000 samples
- Doesn’t require an estimate of mean or standard deviation

Very coarse…



How many attacks can we stand?



Conclusions

Constant noise is managable

- But it requires integrating multiple filtering mechanisms
- It also means assuming a certain mode of behavior
  - This method assumes gaussian, other tests are available

Open questions:

- What do we do with scans once we know they're there?

Privacy, Data Protection Law and Flow Data Anonymisation: requirements, issues, and challenges

Elisa Boschi, Hitachi Europe
Ralph Gramigna, KPMG

Acknowledgement: M. Bossardt (KPMG), D. Battisti (ETH)

Outline

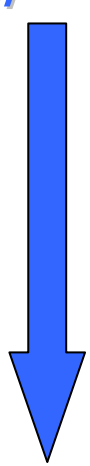
- Review of law principles and requirements on data protection
  - European viewpoint
  - What is personal data?
  - Why is data protection law relevant for network monitoring?
  - Law principles overview
- The role of flow data anonymisation to support data protection
  - Discussion on its applicability and weaknesses
  - Suggestions for future steps

Data Protection Law: EU Directives

- Goal: protect the privacy of individuals
  - Not limited to information confidentiality
- EU Directives define the the minimum law requirements to be implemented by each EU member state
  - Applicable to international data transfers with EU
- Relevant to data protection:
  - Directive 1995/46/EC - on data protection
  - Directive 2002/58/EC - on privacy and electronic communications

Applicability and Personal Data

- Directive 95/46/EC applies to the
  „*processing* of *personal data*“





"*any information relating to an identified or identifiable natural person ('data subject'); an identifiable person is one who can be identified, directly* ***or indirectly****, in particular by reference to an identification number or to one or more factors specific to his ... identity*".

"any operation performed upon personal data, such as e.g. collection, storage, adaptation or alteration, consultation, disclosure by transmission, dissemination or otherwise making available, alignment or combination, erasure or destruction"

- Note: in some countries (e.g. Switzerland) this applies to „legal entities“ as well

Applicability to Network Monitoring

- *Indirect identification data comprise any information that may lead to identification of the data subject through association with other available information*
  - information available to the entity in charge of the data processing (ISP),
  - any information possessed by third parties
- IP addresses can identify someone “directly”
  - Esp. legal entities
- Many more attributes in a flow record can contribute to identifying someone “indirectly”

Principles: legitimation for processing

1. Consent
2. Data processing is „*necessary for the performance of a contract to which the data subject is a party*"
3. ...

- Processing must be limited to specified purposes

- Further processing of data for historical, statistical or scientific purposes is possible provided that appropriate safeguards are provided
  - Left to national laws

Principles: Information of the Subject

The subject must be informed about:

1. Identity of the data controller
2. Purpose of the processing
3. Other information, e.g. the recipient of the data.

- It does not apply to scientific research, **IF** the provision of such information
  - proves impossible
  - would involve a disproportionate effort

- Appropriate safeguards must be provided
  - Their specification is let to national law

Border Crossing

- Transfer to third countries is generally possible if the third country ensures an adequate level of protection

http://ec.europa.eu/justice_home/fsj/privacy/thrid countries/index_en.htm

- E.g.
  - ✓ Switzerland, Canada, Argentina
  - ✗ USA (except Safe Harbor)

Traffic data and location data

- Introduced in Directive 2002/58/EC
  - *Traffic data: any data processed for the purpose of the conveyance of a communication or for the billing thereof*
  - *Location data: data indicating the geographic position of the terminal equipment of a user*

- Objectives:
  - Minimise the processing of personal data
  - Use anonymous or pseudonymous data where possible.

- „Anonymous“ = it is no longer possible to identify the data subject

Processing of Traffic and Location Data

- Traffic and location data relating to subscribers and users must be erased or made anonymous when no longer needed

- The processing of traffic data must be restricted
  - To persons acting under authority of providers
  - To certain activities (e.g. traffic management, fraud detection...)

- Location data can be processed only if
  - There is consent, or
  - Data is made anonymous

The Role of Flow Data Anonymisation to Support Data Protection

- The well known problem:
  - The more you anonymise the better privacy is protected...
  - ...but the less useful the data
- Anonymisation aims at removing sensitive information referring to an individual
- Attacks to anonymisation schemes have proved that those schemes could be broken allowing to "indirectly" identify people.
- Are known flow anonymisation techniques effective in protecting the privacy of individuals?

(4) Anonymization Techniques

Field to be anonymized:

IP address

IP	Truncation	Permutation	Black Marker	Prefix Preserving
135.98.111.17	135.98	141. 2. 32.37	10.1.1.1	22.131.88.67
135.98.111.128	135.98	41.12.96. 67	10.1.1.1	22.131.88.157
135.98.132.37	135.98	142.72.8.5	10.1.1.1	22.131.201.29
141.161.3.3	141.161	21.33.4.1	10.1.1.1	12.192.32.51
141.72.8.5	141.72	11.14.96.118	10.1.1.1	12.78.201.97
32.53.48.1	32.53	12.161.3.3	10.1.1.1	31.197.3.82

Some Anonymisation Attack Methods

- Data injection → injecting information to be logged with the purpose of later recognizing that data in the anonymized trace
- Fingerprinting → matching attributes of an anonymized object against those of a known object (e.g. web server) to discover a mapping between them
- Semantic attacks → system is exploited in a way that the victim thinks to do something, but he is doing something different. The attacker may infer part of the unanonymized IP address by exploiting the semantics of prefix preserving.
- Structure recognition → recognizing structure between anonymized and unanonymized objects

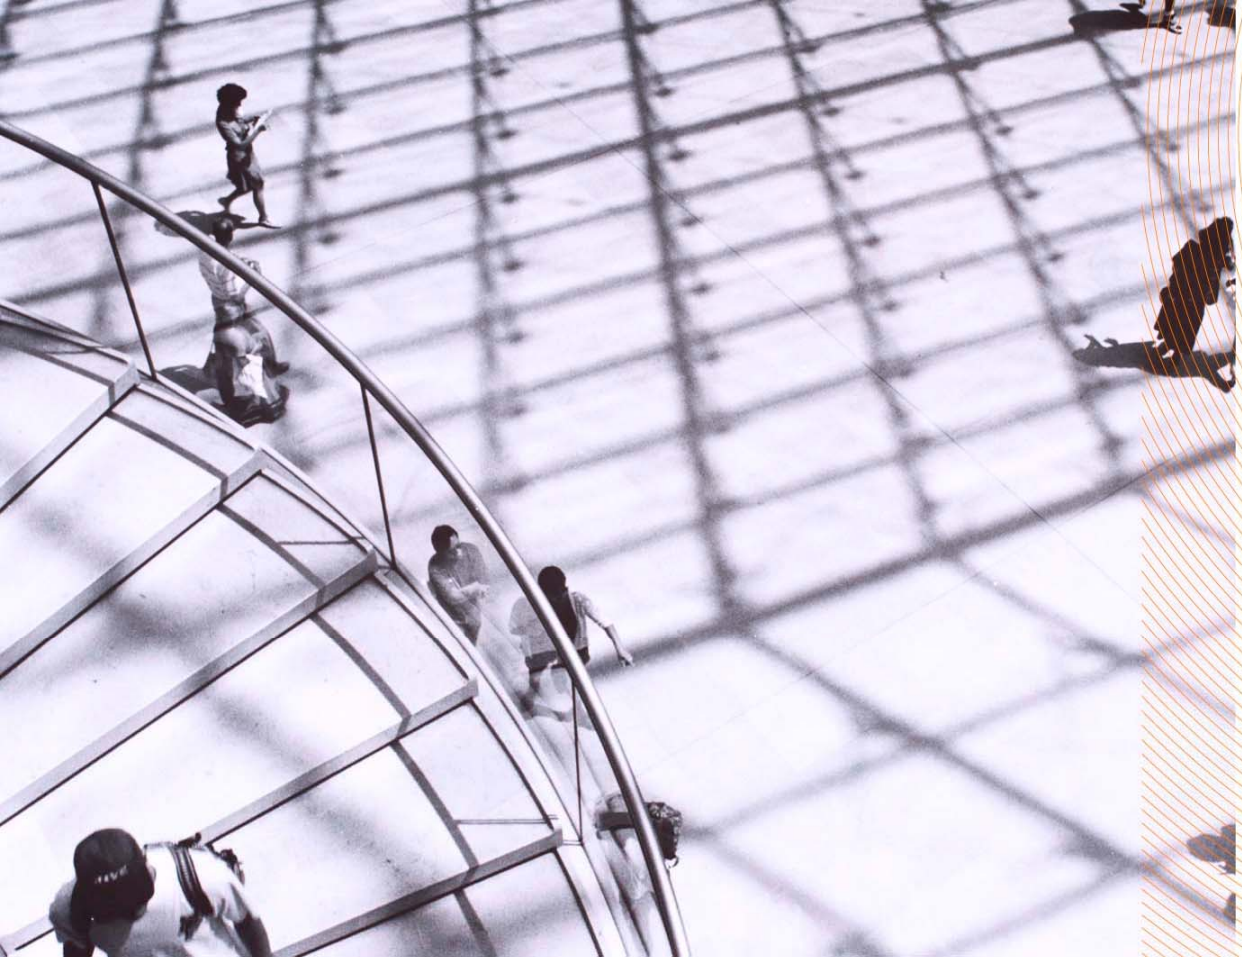
Attacks vs. Anonymisation Techniques

Anonymisation / Attacks	Prefix-preserving	Cryptographic approach	Truncation	Permutation
Semantic attack	■	■		
Cryptographic attack	■	■		
Data Injection	■		■	■
Fingerprinting	■		■	■
Structure Recognition	■		■	■

- ■ the attack can be used, (partial) results achieved

Conclusions

- We need to pay attention to data protection laws
- Anonymisation is part of the solution to protecting privacy, but
  - Research is still needed
  - This is not only a technical problem; a technical solution alone is not enough
- Legal solutions, policies, guidelines, interdisciplinary work are needed
- Anonymisation support is needed in standard flow data export protocols such as IPFIX



Automatic anomaly detection using NfSen

Wim Biemolt, SURFnet

Werner Schram, SURFnet



Automatic anomaly detection using NfSen

- SURFnet and netflow anomaly detection
  - NERD
  - NfSen
  - PeakFlow SP
- Currently used detection methods
  - DDos
  - Botnet
  - Holt-Winters aberrant behavior



SURFnet and netflow anomaly detection

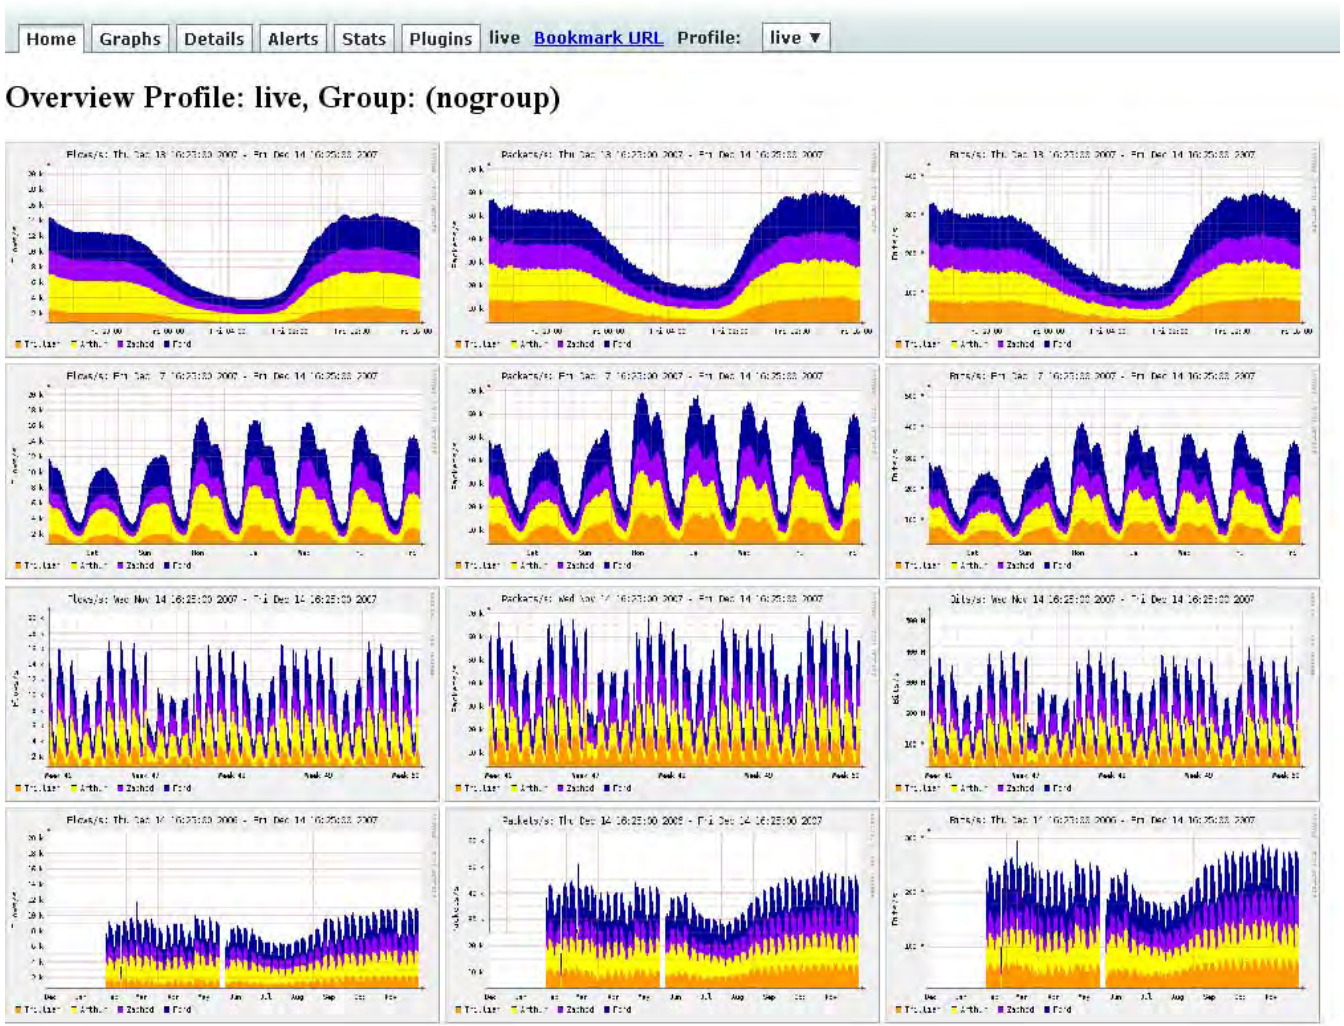
- NERD v1
  - Developed by TNO
  - Based on cflowd
  - cflowd is no longer supported
- NERD v2
  - Initially developed by TNO
  - Has serious performance problems
  - NfSen can do the same but without the performance problems



NfSen

- Netflow Sensor (NfSen) is a
  - network statistics tool
  - Developed by Peter Haag
  - Currently in active development
  - Alert plug-in system
  - Generic plug-in system
  - Some plug-ins already available

NfSen

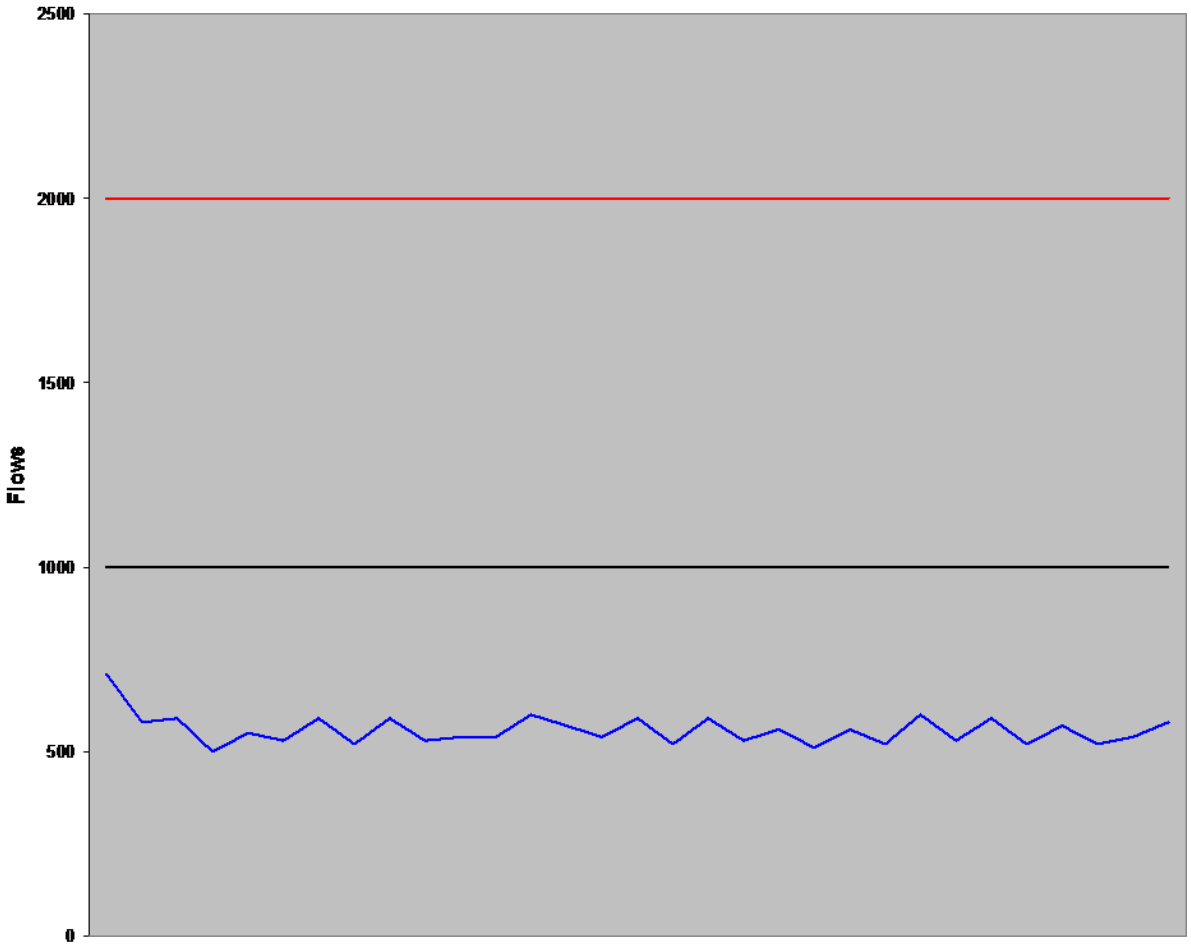

Home
Graphs
Details
Alerts
Stats
Plugins
live
Bookmark URL
Profile:
live
Overview Profile: live, Group: (nogroup)



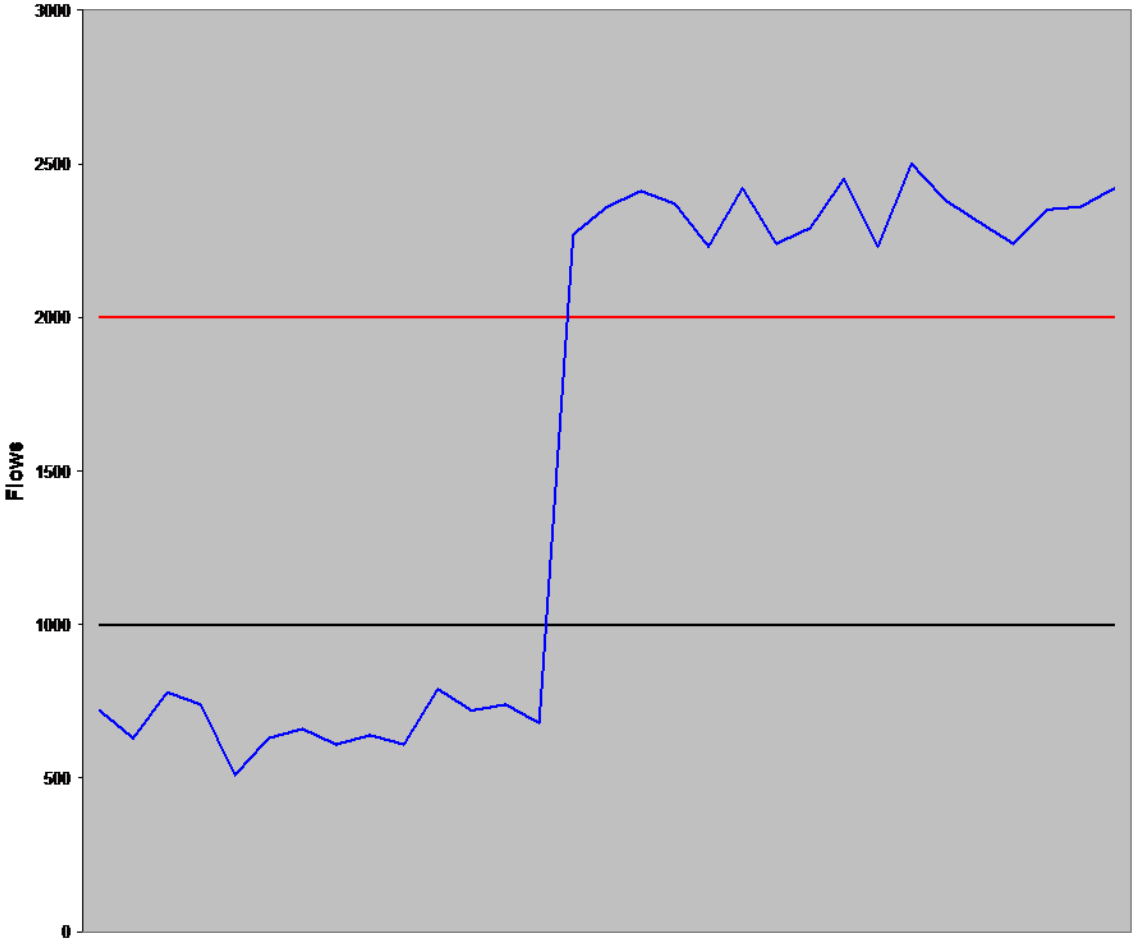
DDos detection

- Simple flow analysis
  - based on NERD v1 DDos detection
  - using a low threshold and a high threshold
  - Rules for traffic between those thresholds
  - Custom thresholds for high load services

Expected traffic

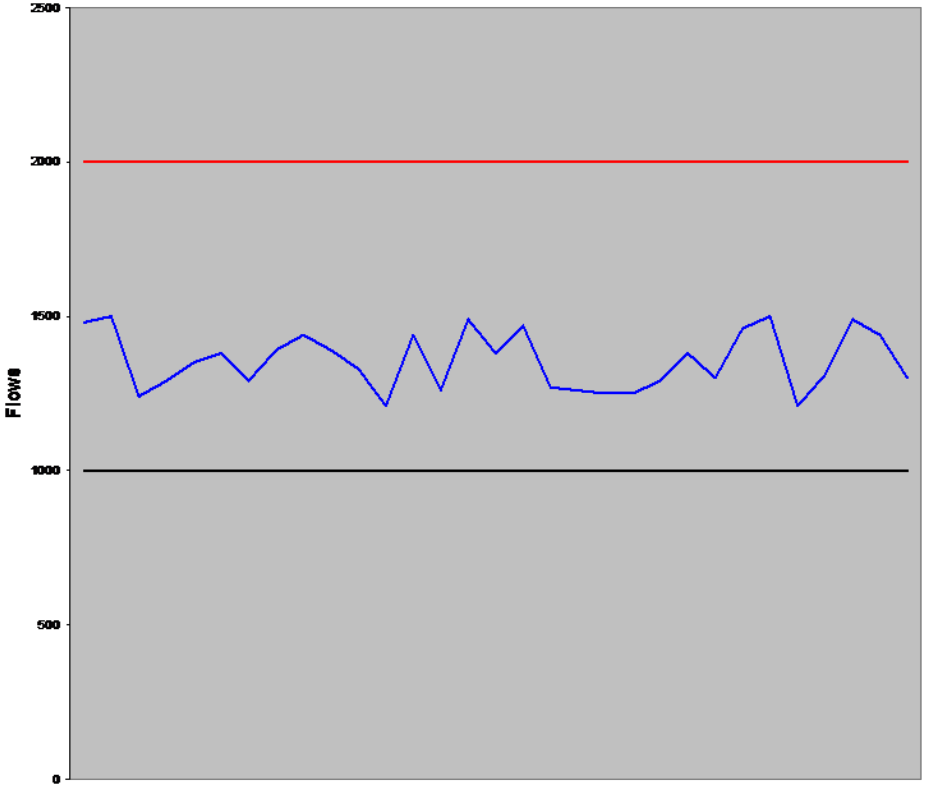


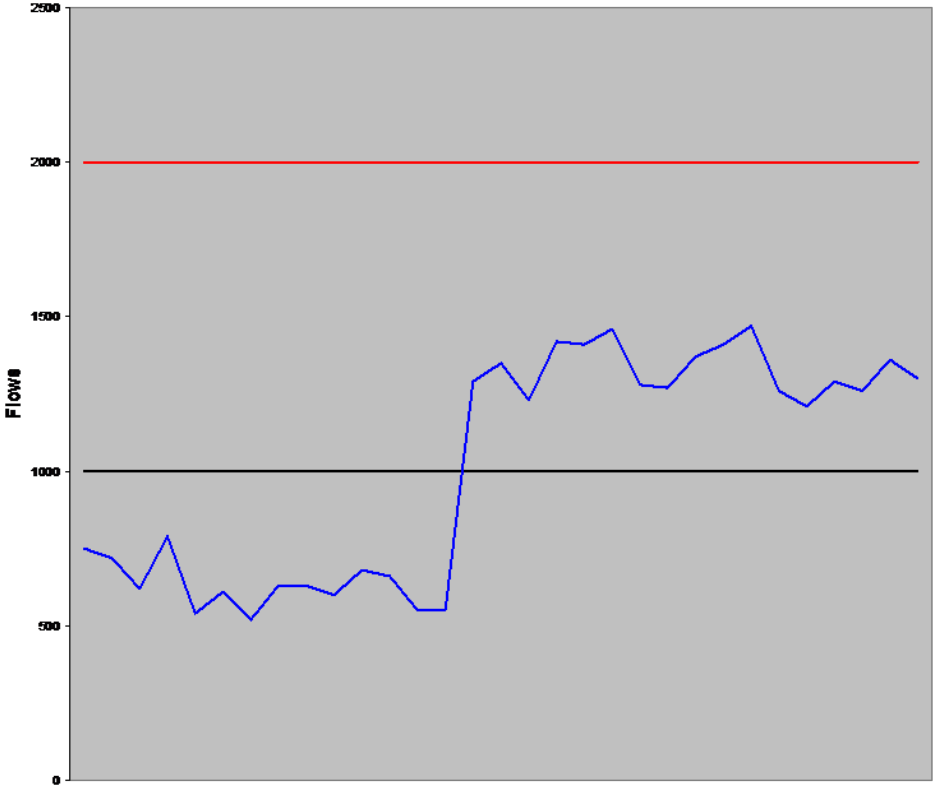
Definitively Conspicuous Traffic



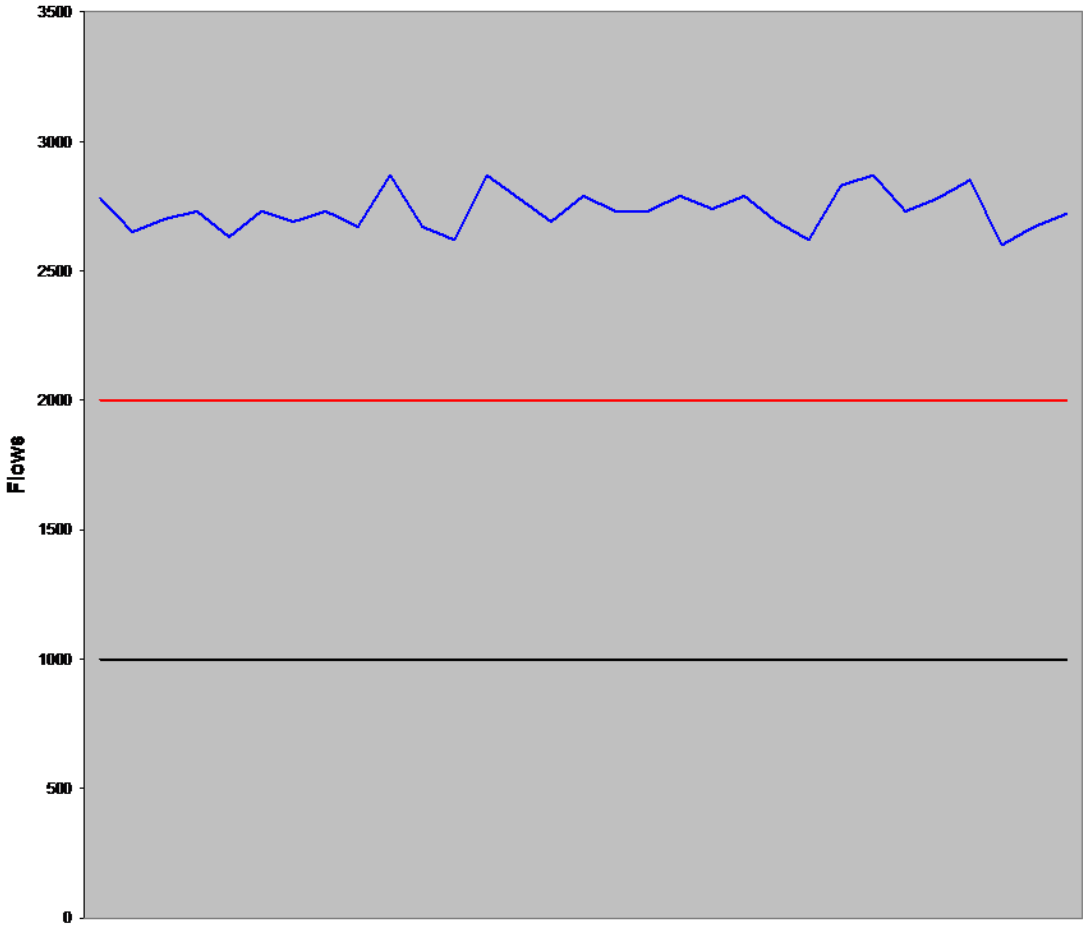


Border cases

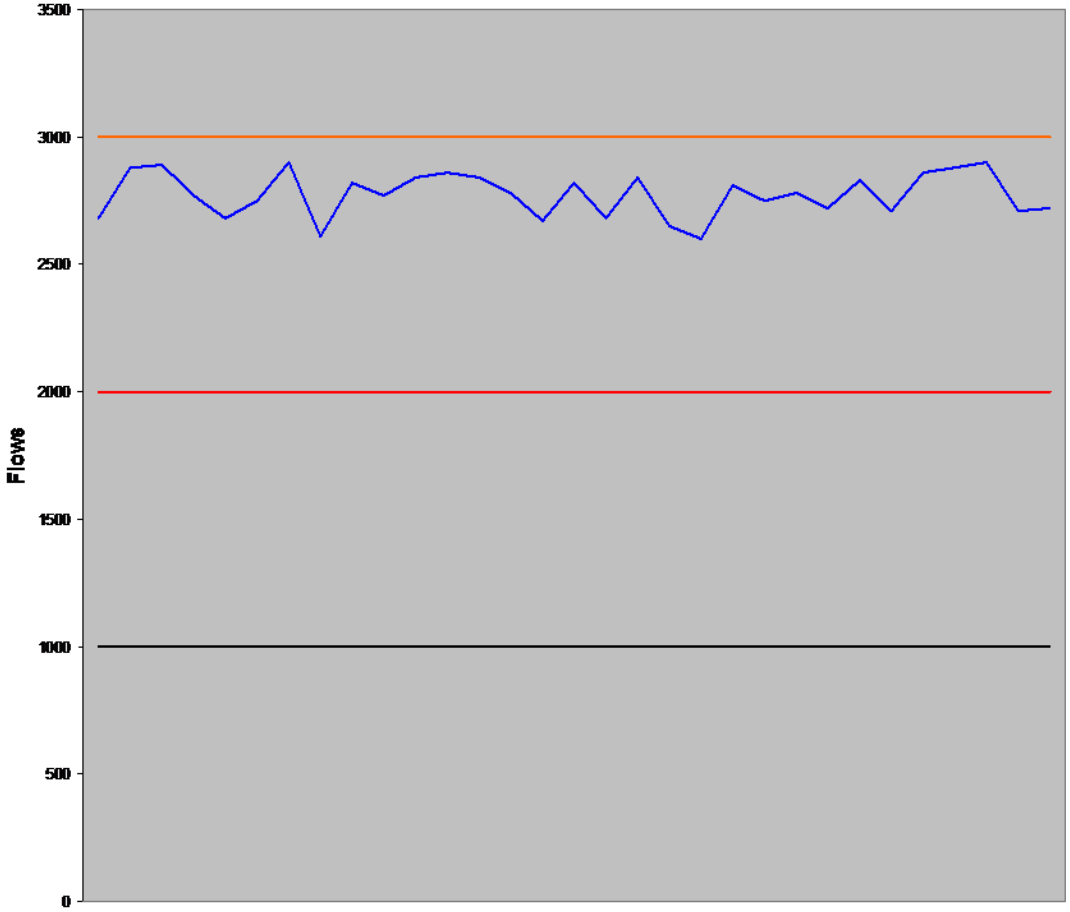




High load servers



Custom thresholds



DDos interface: report

Home | Graphs | Details | Alerts | Stats | Plugins | live | Bookmark URL | Profile: live ▼

alarm | Events

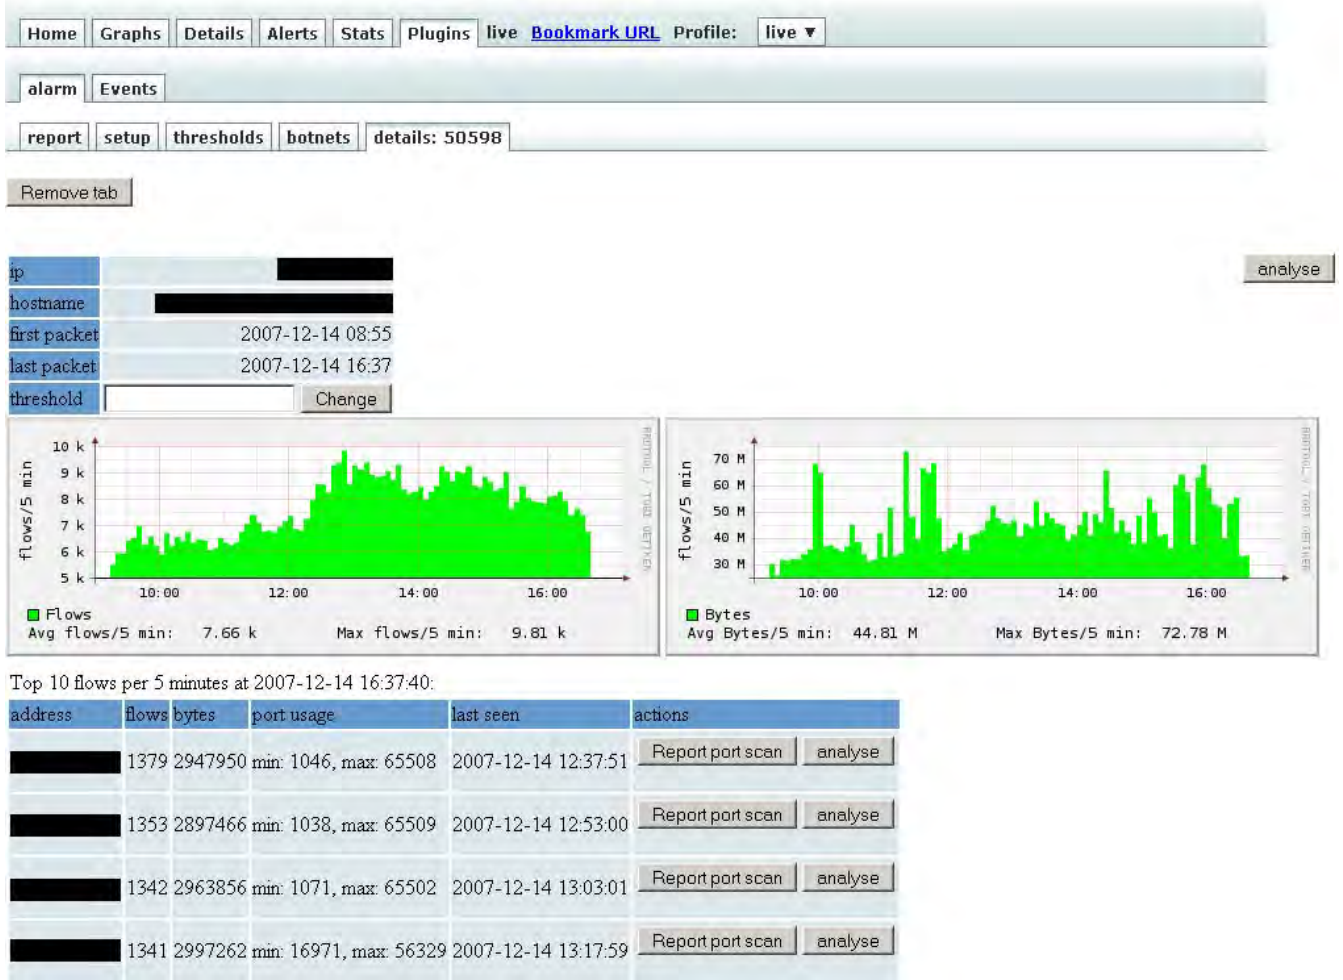
report | setup | thresholds | botnets

number of alarms to show: 10 (0 for all)
from 7 days ago
up to 0 days ago
alarms: ddos
Ok

The ddos alarms between **2007-12-07** and **2007-12-15**

ID	Destination	Flows per 5 minutes	Average packets/flow	Average bytes/flow	Starttime	Stoptime	Active	
#50598	[redacted]	7772	5054	4	2007-12-14 08:55:00	2007-12-14 16:32:50	1	Delete
#50596	[redacted]	10620	3859	4	2007-12-14 08:39:54	2007-12-14 16:32:50	1	Delete
#50594	[redacted]	9510	3147	3	2007-12-14 08:25:01	2007-12-14 16:32:50	1	Delete
#50593	[redacted]	12951	129	2	2007-12-14 08:24:58	2007-12-14 16:32:50	1	Delete
#50490	[redacted]	9517	73	1	2007-12-13 06:13:41	2007-12-14 16:32:50	1	Delete
#49820	[redacted]	281618	163	1	2007-12-04 14:47:47	2007-12-14 16:32:50	1	Delete
#49191	[redacted]	327975	125	1	2007-11-27 13:19:14	2007-12-14 16:32:50	1	Delete
#49074	[redacted]	22047	171	2	2007-11-26 13:32:20	2007-12-14 16:32:50	1	Delete
#50656	[redacted]	5222	2550	3	2007-12-14 16:20:07	2007-12-14 16:29:56	1	Delete
#50635	[redacted]	6031	1155	7	2007-12-14 11:44:53	2007-12-14 16:22:51	1	Delete

DDos interface: Details





Botnet detection

- Hosts infected by viruses connect to hosts known as botnet controllers
- List of botnet controllers are available, for example:
  http://www.bleedingthreats.net/rules/bleeding-botcc.rules
- Our plug-in logs all hosts that connect to known botnet controllers
- Automatically reports to incident report system using IODEF

Botnet IODEF reports

```
<?xml version="1.0" encoding="iso-8859-1"?>
<io:IODEF-Document xmlns:io="urn:ietf:params:xml:ns:iodef-1.0" lang="en">
  <io:Incident purpose="reporting">
    <io:IncidentID
    <io:StartTime>2
    <io:EndTime>200
    <io:ReportTime>
    <io:Assessment>
      <io:Impact ty
    </io:Assessment
    <io:Contact>
      <io:ContactNa
    </io:Contact>
    <io:EventData>
      <io:Method>
        <io:Referen
          <io:Refer
        </io:Refere
      </io:Method>
      <io:Flow>
        <io:System
          <io:Node>
            <io:Add
            <io:Cou
          </io:Node
        </io:System
        <io:System
          <io:Node>
            <io:Add
          </io:Node
          <io:Servi
            <io:Por
          </io:Serv
        </io:System
      </io:Flow>
    </io:EventData>
    <io:AdditionalD
NFSen</io:Additiona
  </io:Incident>
</io:IODEF-Document>
```

AIRT

IncidentdetailsSURFcert#019038

| Main menu | Import queue | Incidents | Search | Close current incident | Mail templates | Edit settings | Logout |

(Bewerken) Externe identificatie:

(Bewerken) Ticket number(s):

Elementaire incidentgegevens

incidentsoort	infected
incidenttoestand	inspection requested
Incidentstatus	open
Datum van incident	20 aug 2007 17 02 03
Logboekinformatie	Source (ip) : 192.168.1.1 Target (ip:port) : 192.168.1.2 Packet (type:count): flow:23 Start time : 2007-08-13T15:07:47+02:00 End time : 2007-08-13T21:06:12+02:00

update

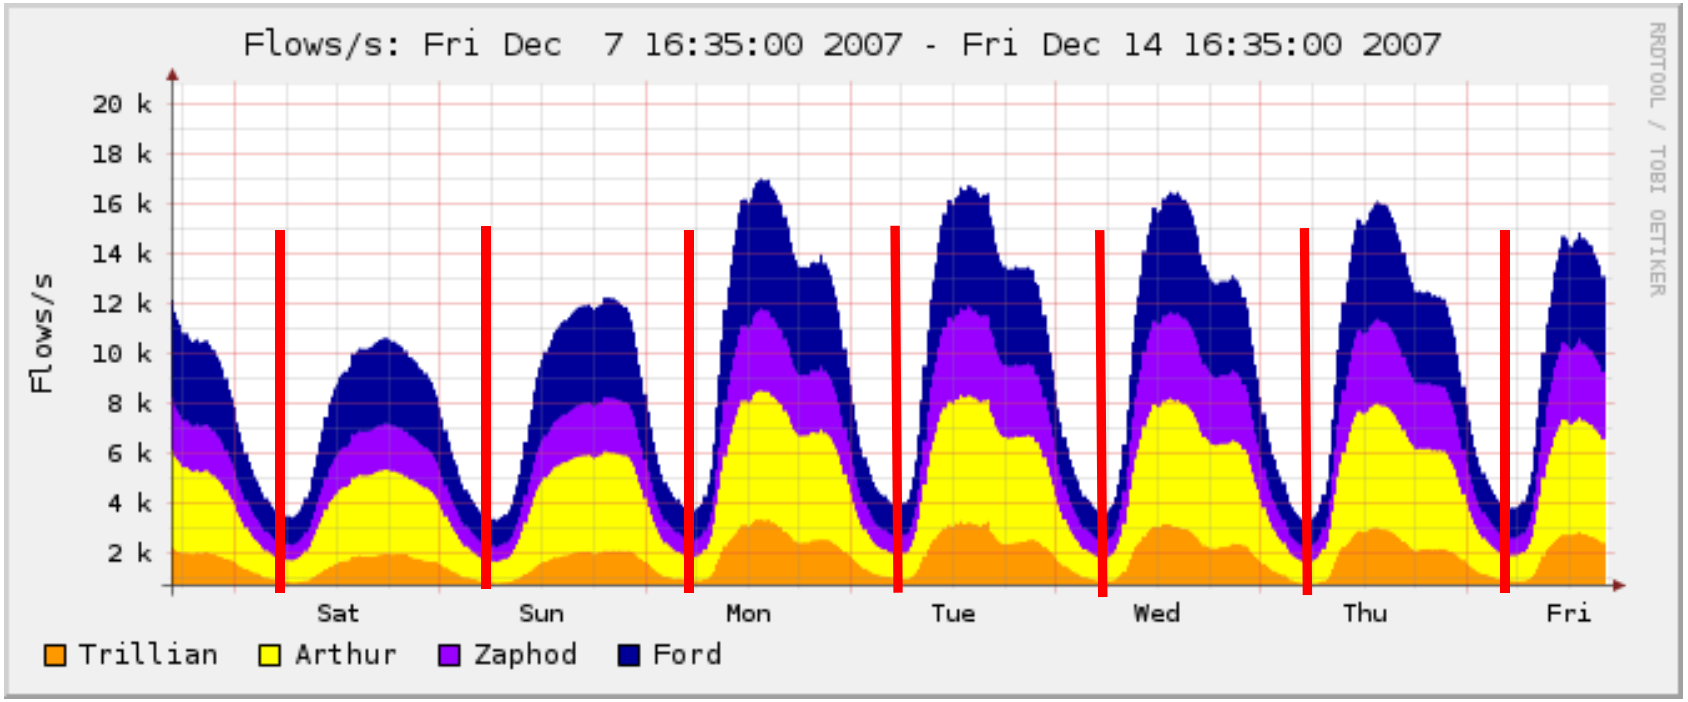
Beïnvloedde IP-adressen

IP adres	Machinenaam	Constituency	Rol in incident	Bewerken	Verwijder
192.168.1.1	infected.host	utwente.nl	Unknown	bewerken	verwijderen

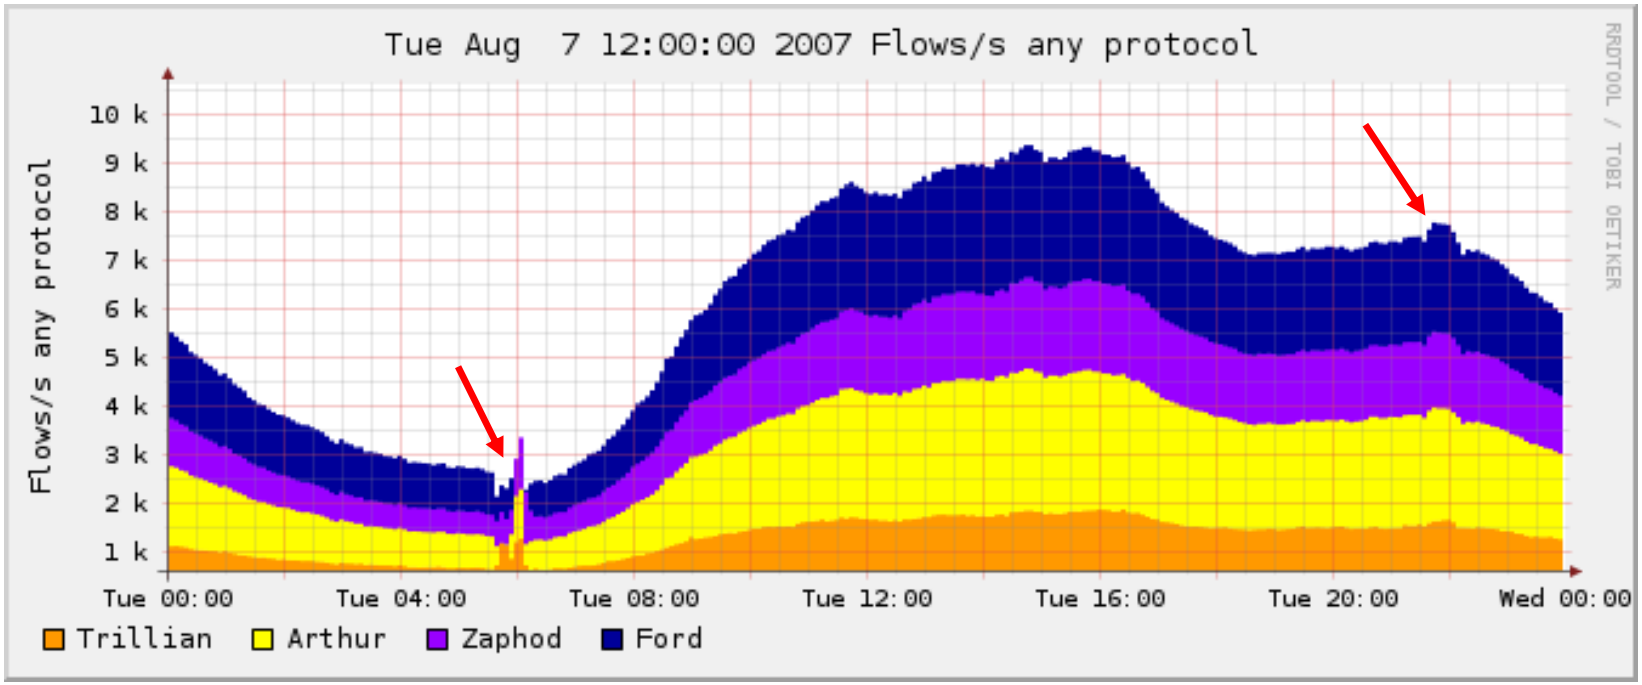
IP adres Unknown Toevoegen

Holt-Winters aberrant behavior detection

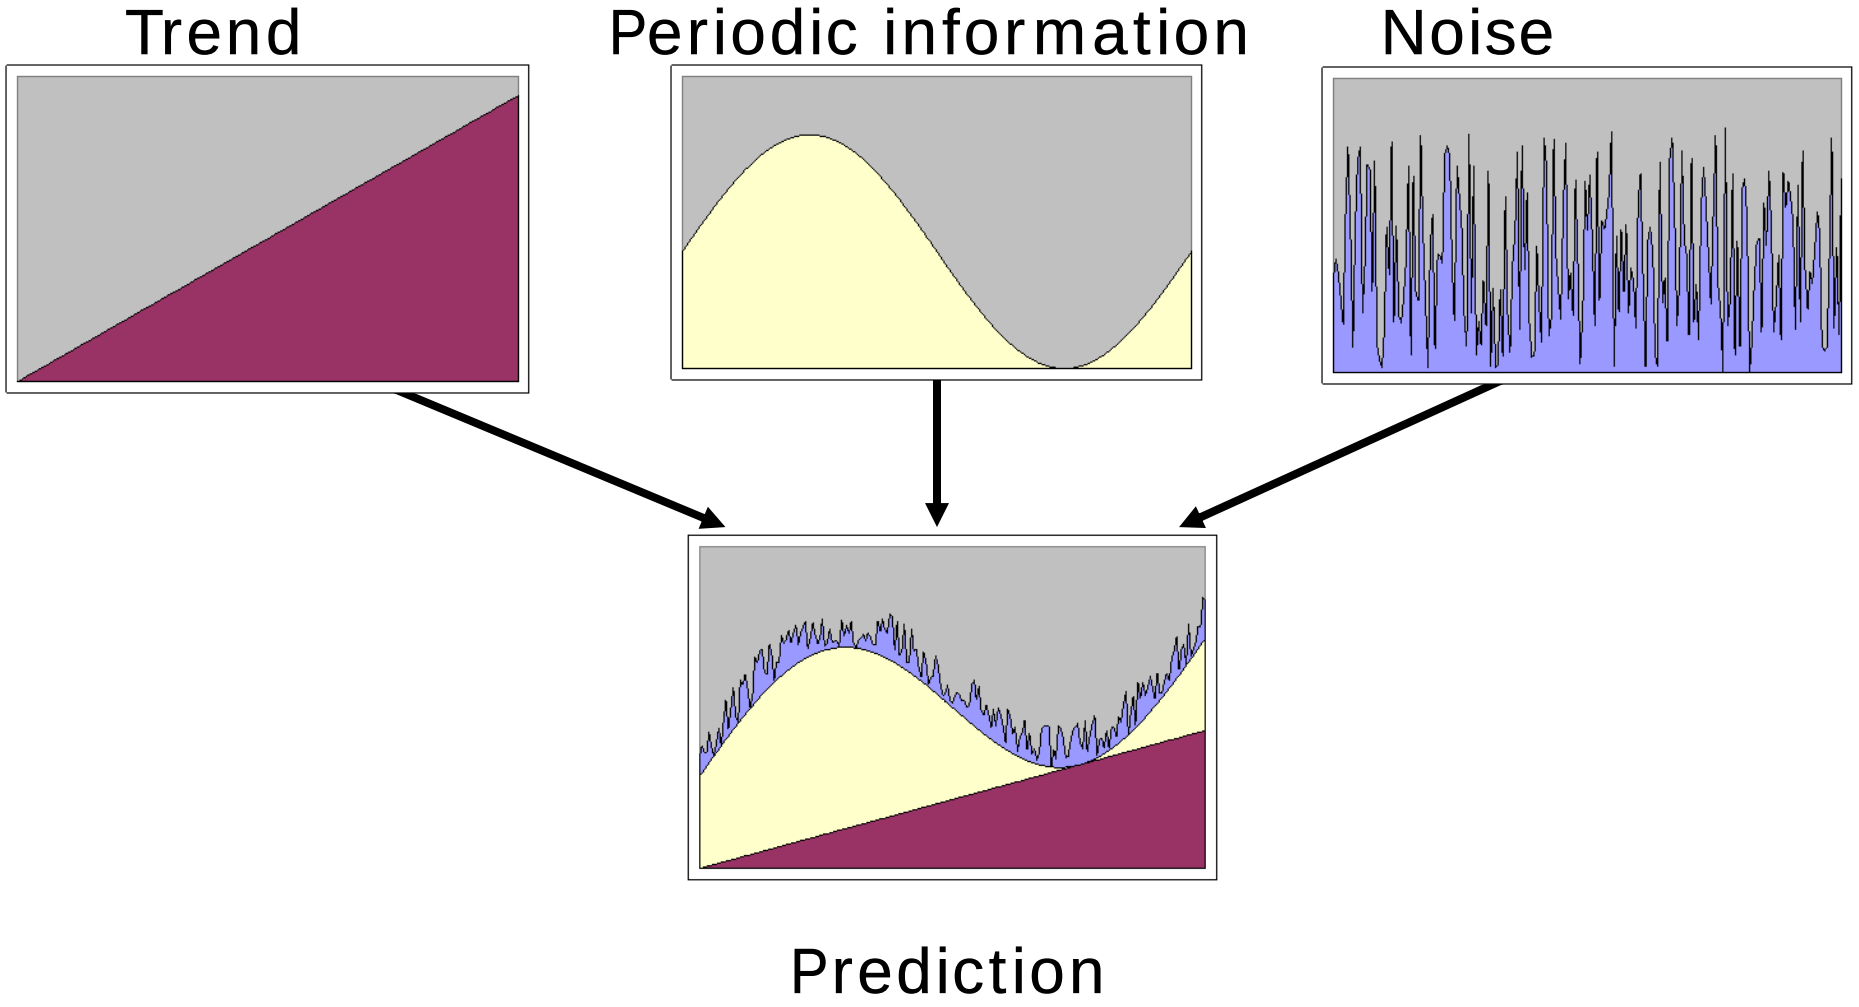
- Uses information about periodic data to predict aberrant behavior.



Holt-Winters: Example



Holt-Winters: Original implementation



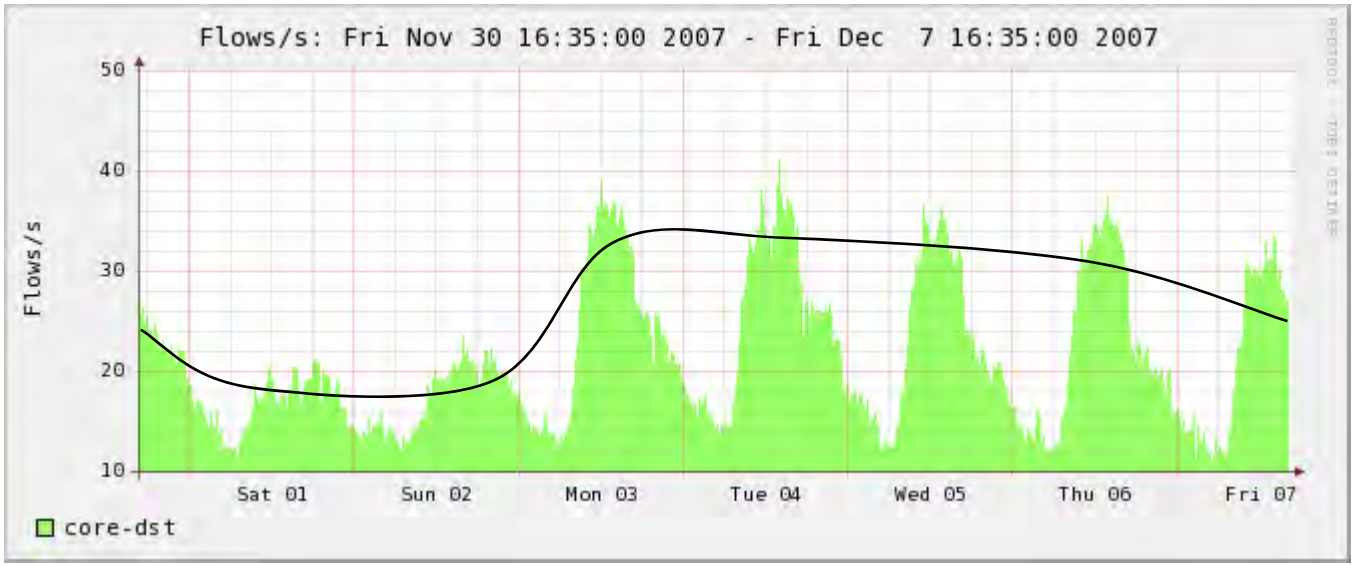


Limitations of the original implementation

- The original algorithm has three parameters which define:
  - the weight of historical data
  - the weight of the trend
  - the amount of expected noise
- The original algorithm has a constant learning rate
  - If a low learning rate is used, the selection of the initial values is critical. This will introduce false positives for a long time.
  - With a high learning rate, the model will likely be overfitted. This will introduce false negatives
- The trend parameter has no significant influence with the resolution we are using

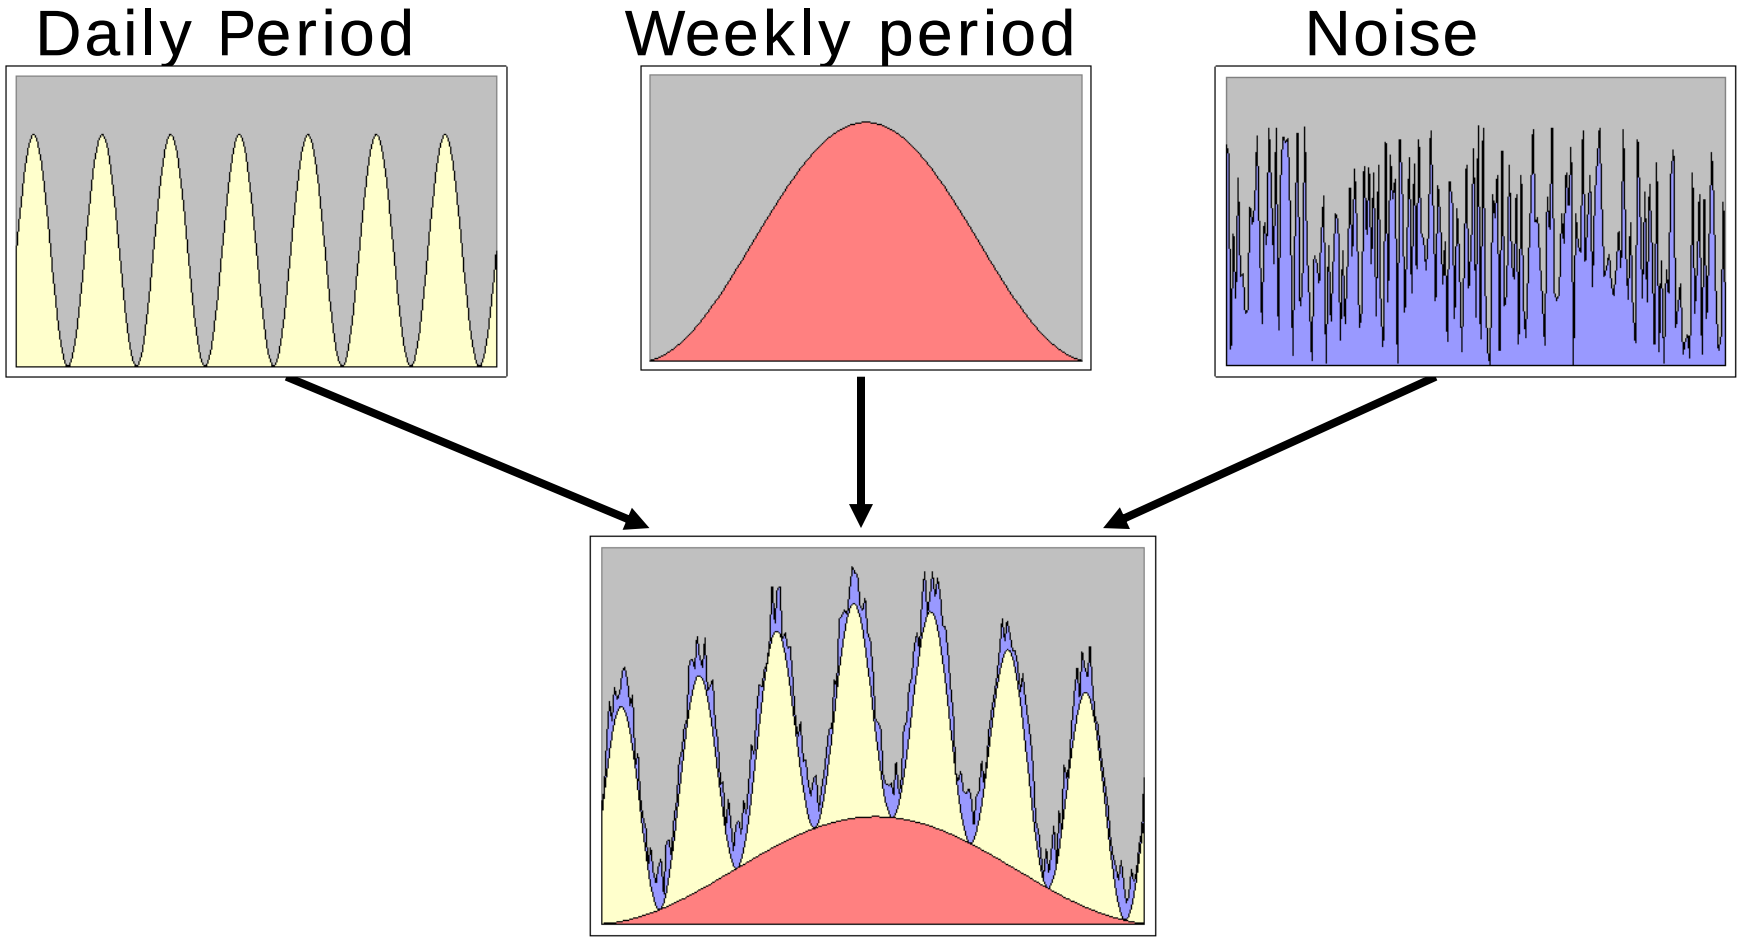
Holt-Winters: Multiple trends

Network traffic time series often show multiple recurring patterns, for example a weekly trend:

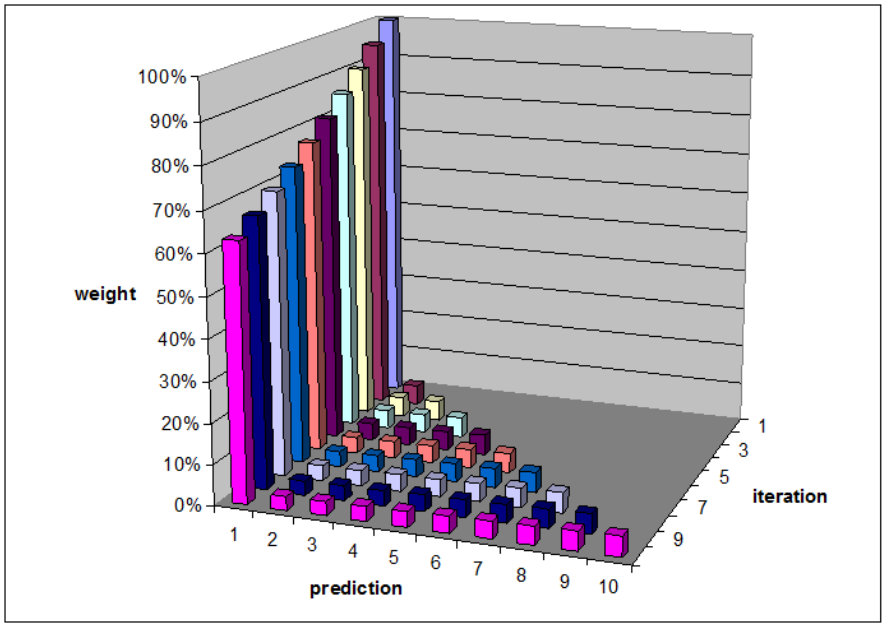




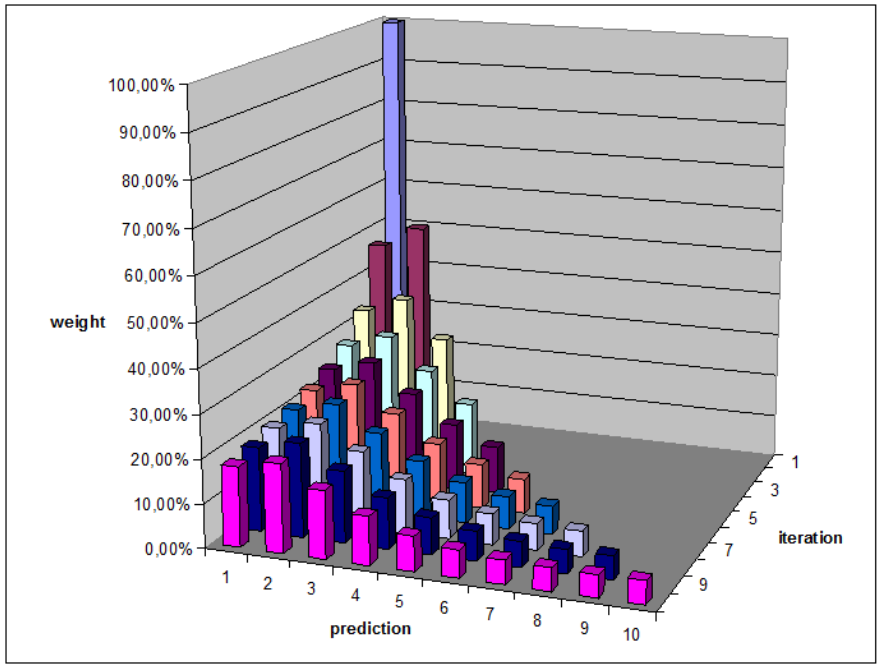
Holt-Winters: Multiple periods



Learning rate



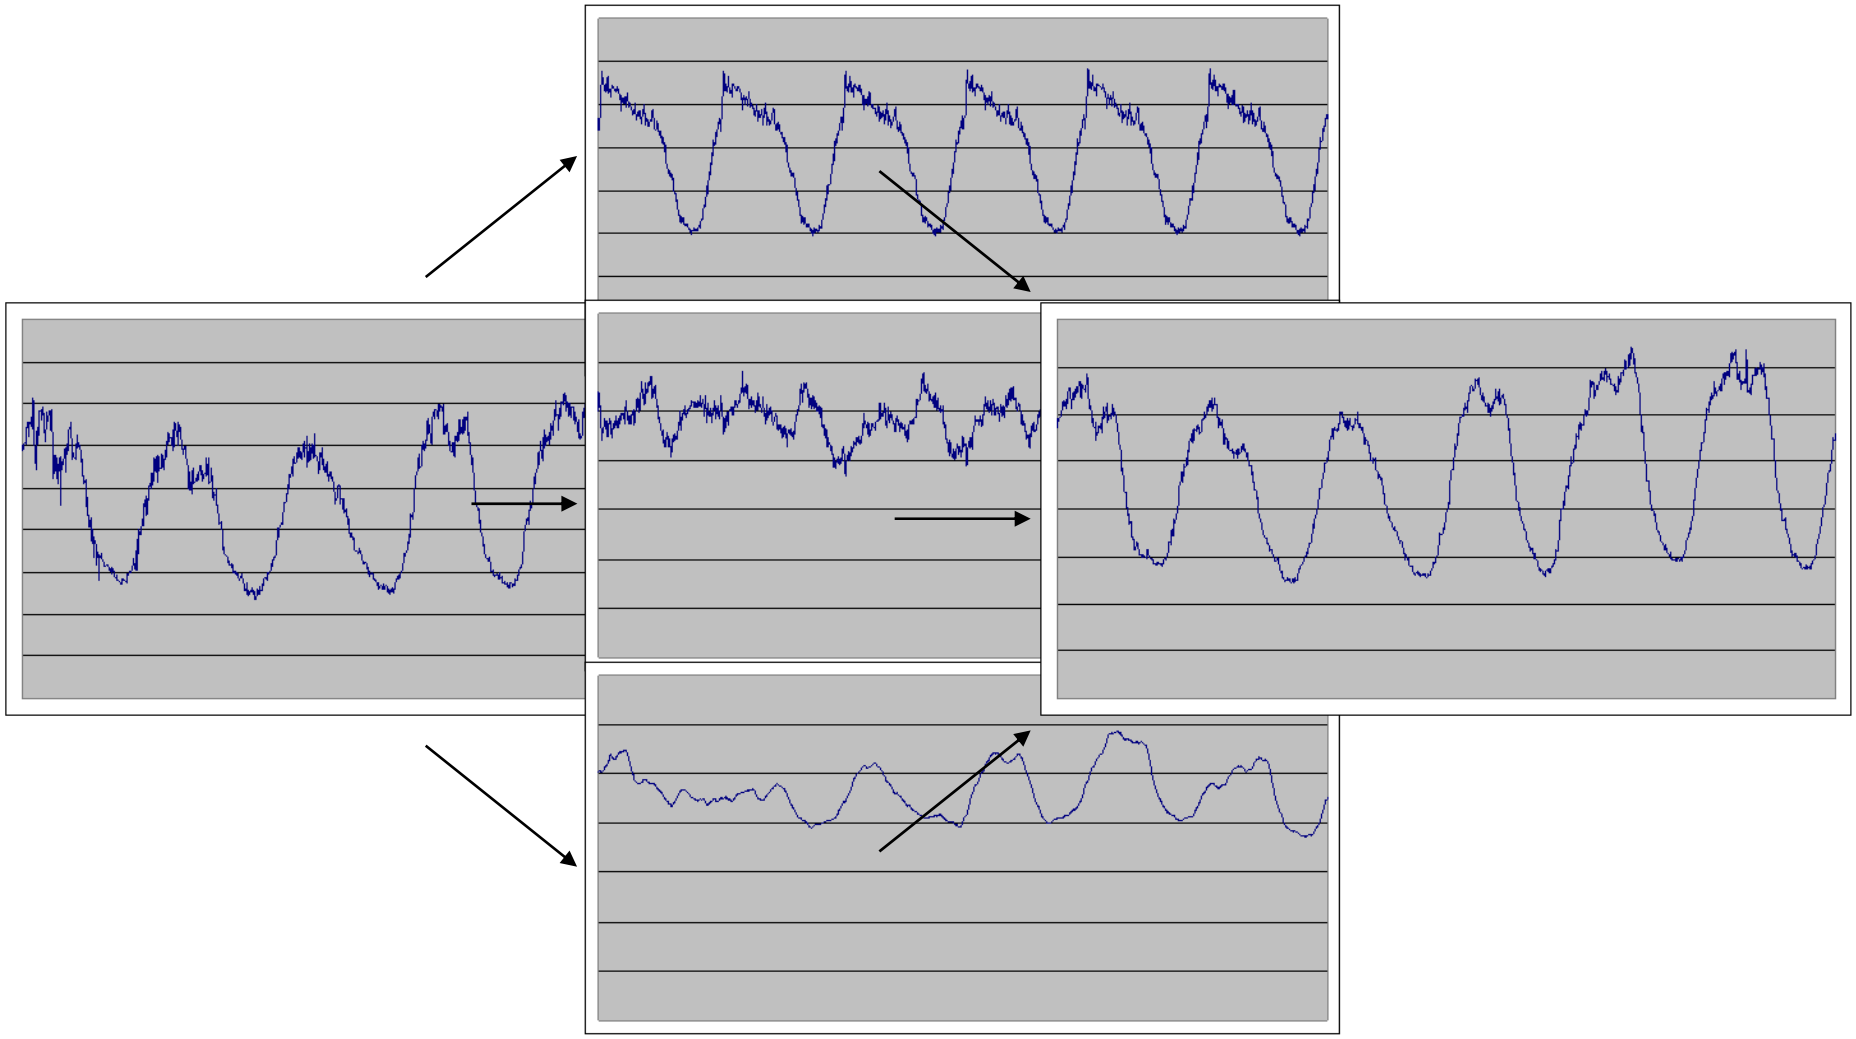
Fixed learning rate:
The first pattern is overweighted



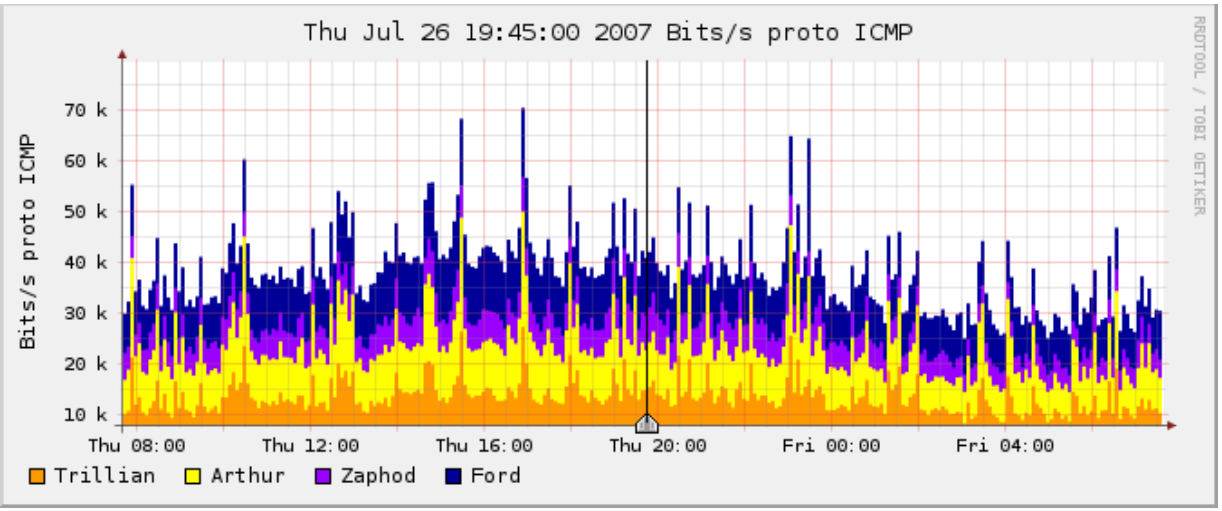
Adaptive learning rate:
The weight of the first pattern is relative to the rest



Real data example

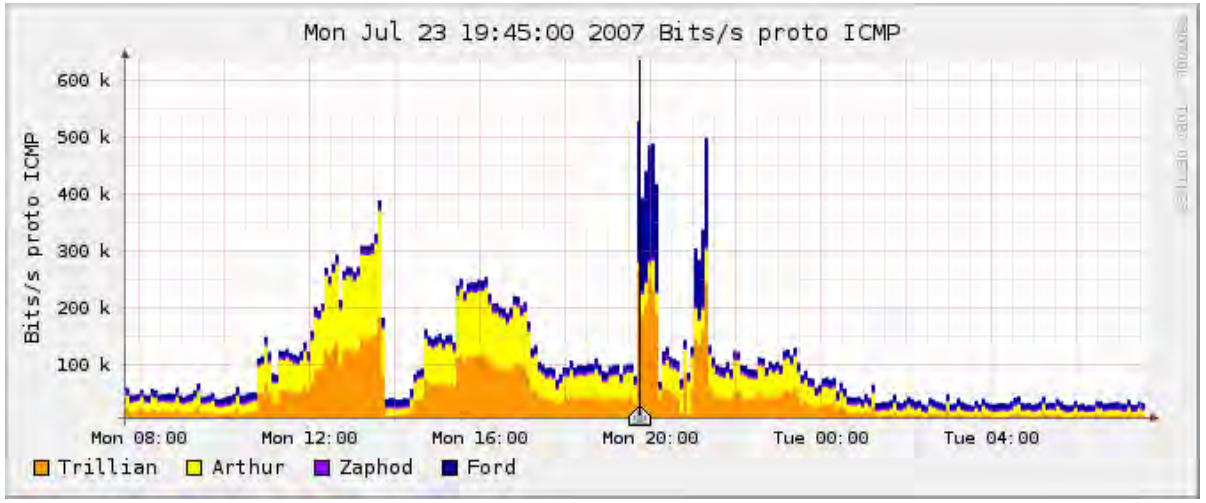


Holt Winters: Usage Example

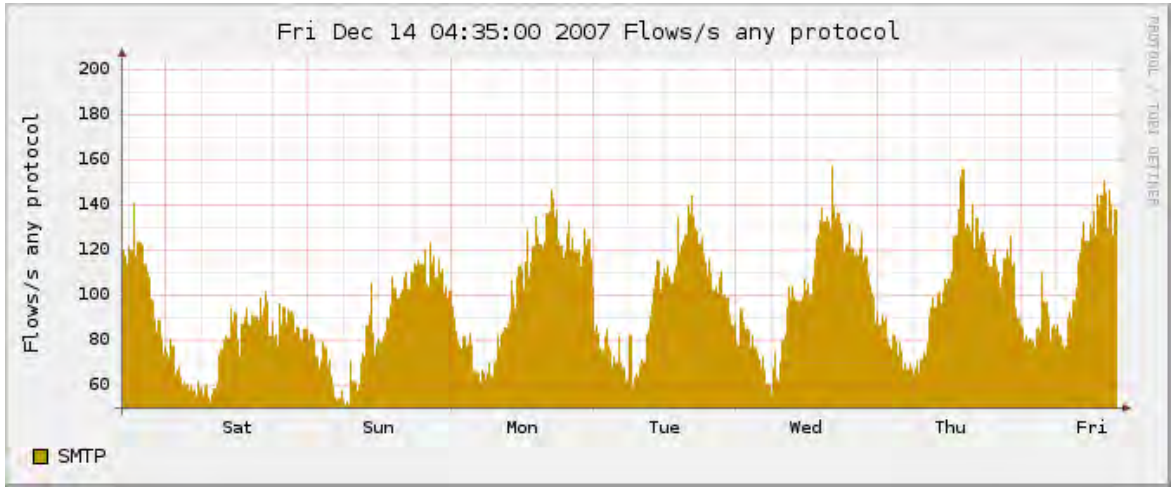


Normal ICMP Traffic

Aberrant ICMP Traffic:
Caused by DDos attack
by Stormworm
botnet

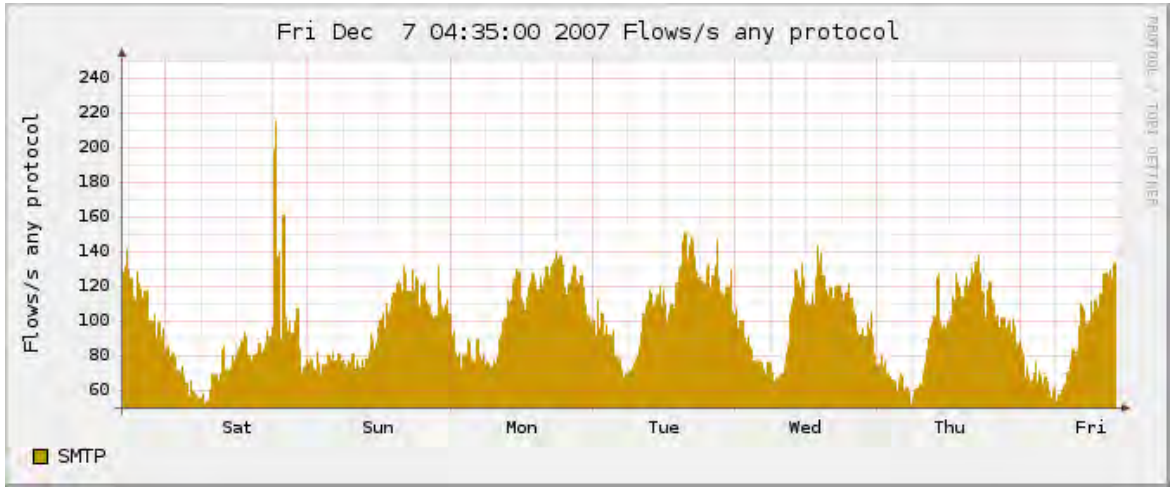


Holt Winters: Other possible uses



Common SMTP Traffic

Last week SMTP Traffic



Wim Biemolt
Wim.Biemolt@surfnet.nl
www.surfnet.nl

Werner Schram
Werner.Schram@surfnet.nl
www.surfnet.nl



YAF

A Case Study in Flow Meter Design

presented at
FloCon 2008 - Savannah, Georgia

Brian Trammell
Technical Lead, Engineering
CERT Network Situational Awareness

© 2008 Carnegie Mellon University

YAF

Open-source, IPFIX-compliant bidirectional flow meter

- Available from http://tools.netsa.cert.org

Processes packets from multiple inputs

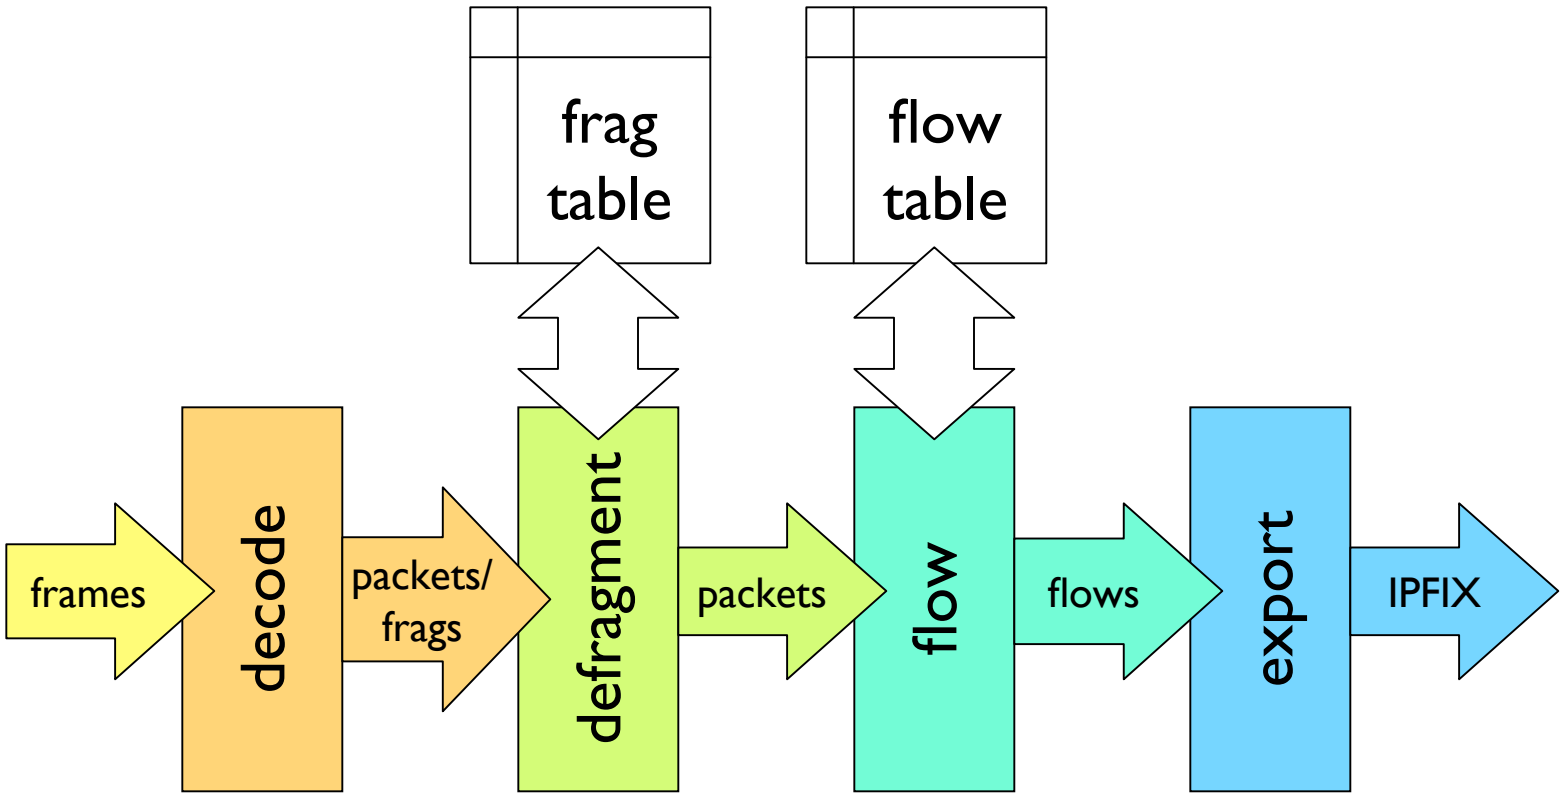
- libpcap dumpfiles (ad-hoc packet analysis)
- libpcap live capture (including proprietary pcap interfaces, e.g. Bivio)
- Endace DAG live capture

Performance is network hardware and I/O bound…

- …easily handles OC3, OC12, GigE at line speed, but
- 10GigE requires proprietary hardware at saturation.

© 2008 Carnegie Mellon University

Flow Meter Design



© 2008 Carnegie Mellon University

Flow Meter Effects on Flow Data

Fragmentation

End Conditions

Timeouts

Delta Counters

Biflows

The Packet Clock

© 2008 Carnegie Mellon University

Fragmentation

Three approaches for flowing fragmented traffic:

- pretend there's no such thing as fragmentation,
- drop all fragmented packets, or
- full or partial fragment reassembly

Each approach has tradeoffs, and is applicable in certain situations.

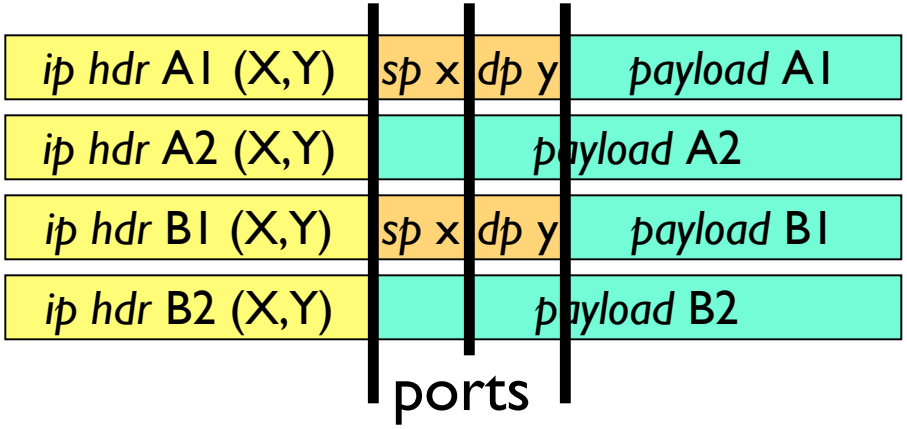
YAF supports partial reassembly.

© 2008 Carnegie Mellon University

Fragmentation?

Easiest way to handle fragmentation: don't.

Leads to inaccurate flow data as subsequent fragment port numbers are incorrectly decoded:



sip	*dip*	*proto*	*sp*	*dp*	*pkts*
X.X.X.X	Y.Y.Y.Y	6	x	y	2
X.X.X.X	Y.Y.Y.Y	6	$A2_0$	$A2_2$	1
X.X.X.X	Y.Y.Y.Y	6	$B2_0$	$B2_2$	1

© 2008 Carnegie Mellon University

Fragmentation? (2)

Often used in resource-restricted environments (e.g., routers).

- Much faster: no requirement even to recognize fragmented packets.
- Much less memory consumption: no fragment table.
- Less susceptible to resource exhaustion attacks.

Trivially easy to implement.

Difficult or impossible to recover actual flows from random fragment offset port data.

© 2008 Carnegie Mellon University

Dropping fragmented packets

Requires minimal resources at flow meter:

- need to recognize fragments, but not store them.

Leads to meter blindness:

- all an attacker must do to hide from the measurement infrastructure is fragment all packets.

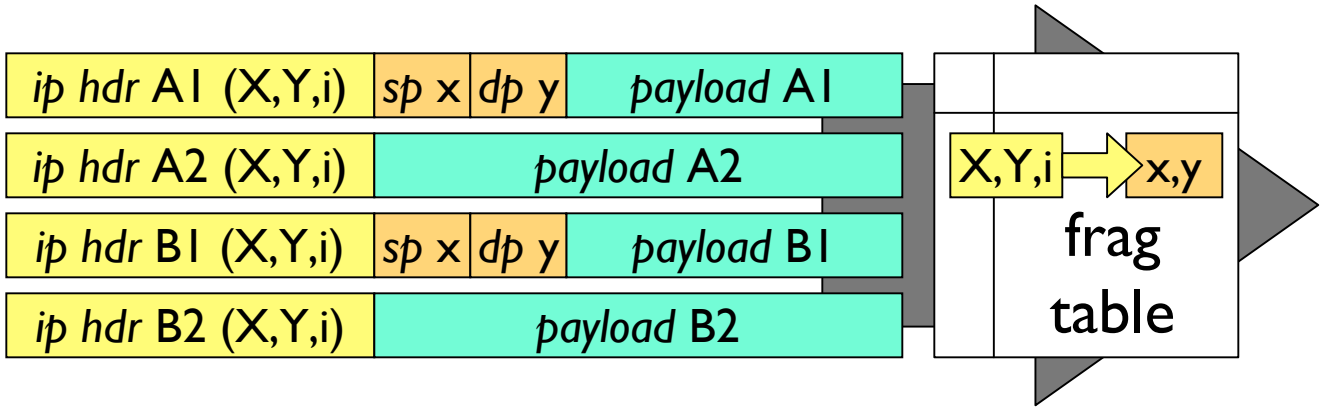
Only applicable behind perimeter devices which also drop all fragmented packets.

```
sip             dip             proto sp    dp    pkts
                        [no flows]
```


© 2008 Carnegie Mellon University

Partial fragment reassembly

Associate each fragmented packet with its actual transport ports:



sip	*dip*	*proto*	*sp*	*dp*	*pkts*
X.X.X.X	Y.Y.Y.Y	6	x	y	4

© 2008 Carnegie Mellon University

Partial fragment reassembly (2)

Accurately assigns fragments to respective flows.

Requires additional resources at flow meter:

- need to recognize, look up, and store every fragment.

More difficult to implement and maintain.

Requires care to avoid vulnerability to resource exhaustion attacks.

© 2008 Carnegie Mellon University

Flow End Conditions

Flow meter must recognize actual connection shutdown…

- …through varying degrees of modeling the host TCP state machine.

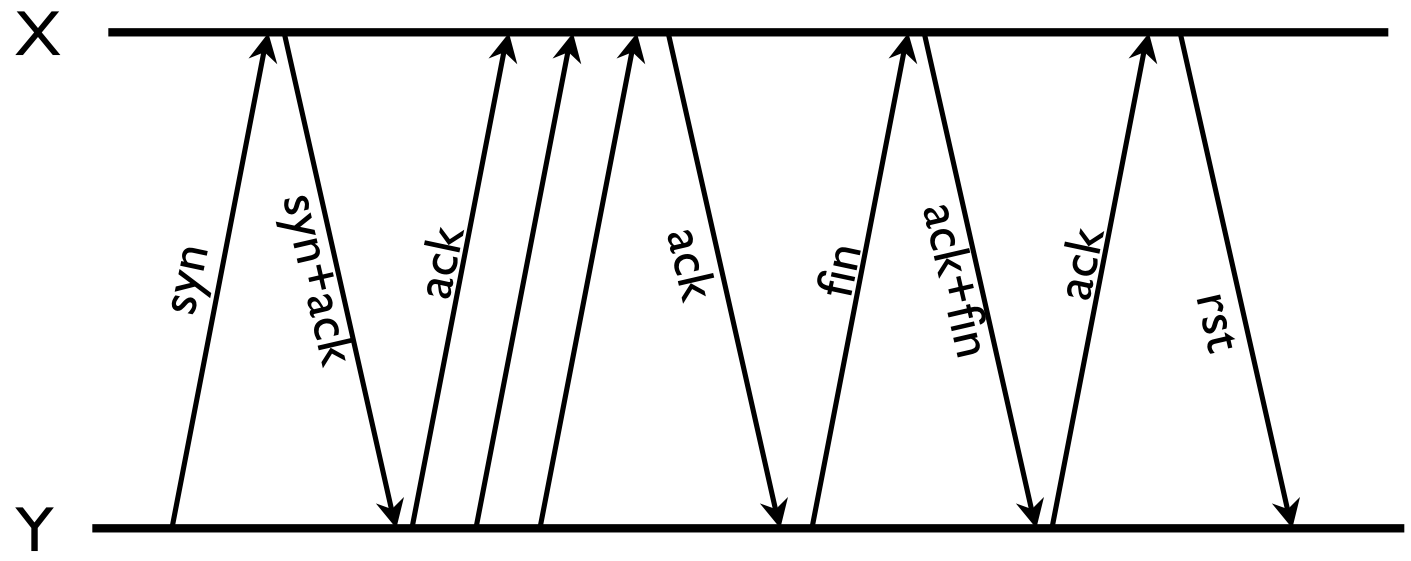
Flows on the wire are not always so well-behaved.

Example: multiple-RST teardown.

© 2008 Carnegie Mellon University

Multiple RST teardown

How many flows here?



sip	*dip*	*flags*	*sp*	*dp*	*pkts*
Y.Y.Y.Y	X.X.X.X	SAF	x	y	6
Y.Y.Y.Y	X.X.X.X	SAF	y	x	3
Y.Y.Y.Y	X.X.X.X	R	y	x	1

© 2008 Carnegie Mellon University

Multiple RST teardown (2)

Tempting to group RSTs on teardown into original flow…

- …how long to keep closed flow state?
- …how far to take this RST grouping?
- …how to communicate new configuration parameters to analysts?

YAF stays predictable, at the expense of generating multiple flow records for this behavior.

© 2008 Carnegie Mellon University

Passive Timeouts

Flows which have no packets over $TO_{passive}$ seconds are closed.

Necessary to terminate flows for all non-connection-oriented transports,

- i.e., anything but TCP.

Longer passive timeouts consolidate low-frequency periodic activity into fewer flows.

Shorter passive timeouts reduce flow table resource consumption for such activity.

© 2008 Carnegie Mellon University

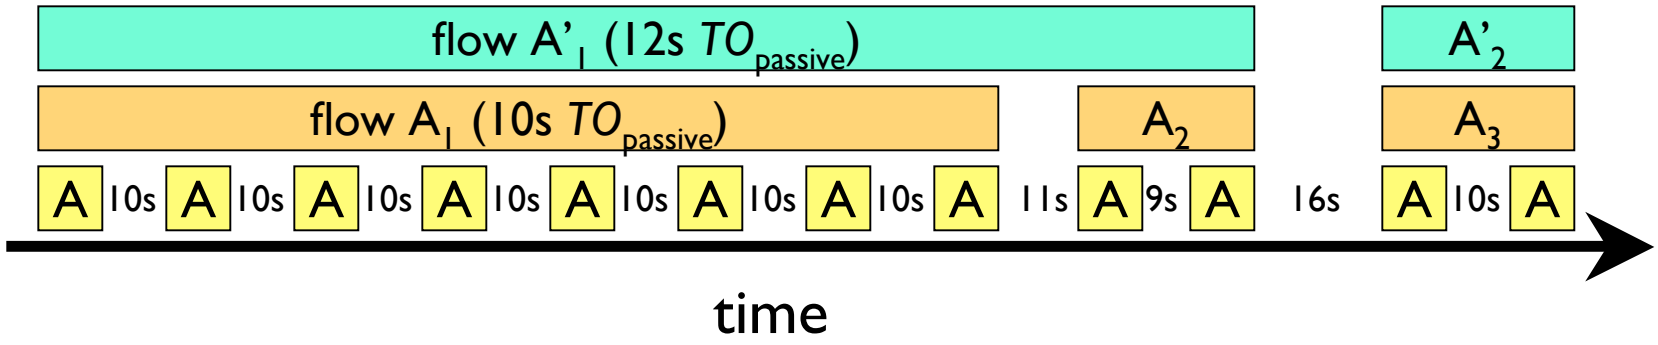
Passive timeouts (2)

Generally chosen to match common protocol timeouts…

- … which are generally round numbers, e.g., 10, 30, 60 sec.

May be chosen to avoid flow closure ambiguity due to minor variations:

- e.g., 12, 33, 64 sec.



© 2008 Carnegie Mellon University

Active Timeouts

Flows which have been open for TO_{active} seconds are closed.

- Maximum flow duration is TO_{active} seconds.

Necessary to ensure long-lived flows are eventually flushed from the flow table.

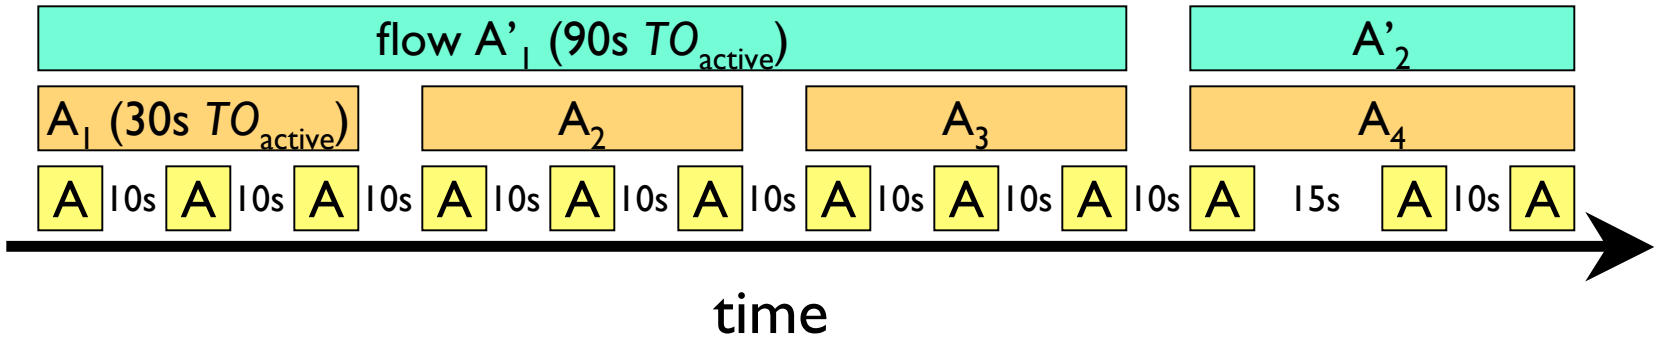
Active timeout determines reporting delay.

© 2008 Carnegie Mellon University

Active Timeouts (2)

Shorter active timeouts used for more rapid reporting.

Longer active timeouts used for better data reduction.



© 2008 Carnegie Mellon University

Delta Counters

Flow meters which periodically emit multiple flow records per flow (for rapid reporting) may use total or delta counters.

Total counters replace values in previous flow records.

Delta counters add to values in previous flow records…

- …thereby reducing state requirements on meter and increasing them on collector.

YAF uses total counters, but doesn't emit multiple records per flow…

- …uses active timeout instead.

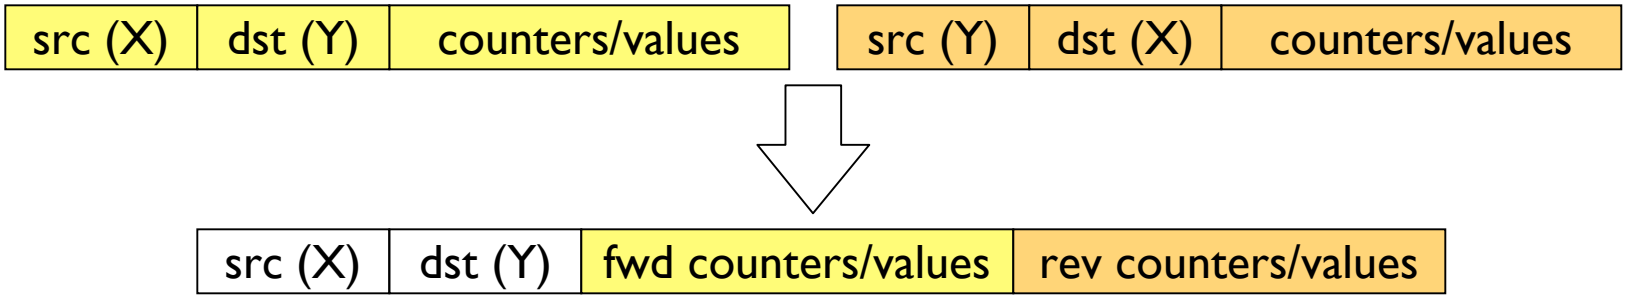
© 2008 Carnegie Mellon University

Biflows

Representation of two sides of a connection with a single flow record:

- Allows additional data reduction
- Enables easier connection analysis
- Improves flow state modeling at flow meter

YAF is a biflow meter, but SiLK stores uniflows.



© 2008 Carnegie Mellon University

The Packet Clock

Important to drive all processes within a flow meter with a single clock

- fragment timeouts, flow timeouts, time stamping, etc.

When building a flow meter, gettimeofday(2) is not your friend.

- often a problem with porting host-based software into a network-based monitoring environment

Use the timestamp from the packet instead!

- ensures that the resulting flow stream identical whether captured live or generated from dumpfile.

© 2008 Carnegie Mellon University

Getting YAF

http://tools.netsa.cert.org

Builds on Mac OS X, Linux, BSD, Solaris

- Bug reports from these or other Unices welcome!

Some prerequisites

- glib-2.0 (C modernization layer)
- libairframe (application utility library from NetSA)
- libfixbuf (IPFIX protocol implementation from NetSA)
- libpcap (generally available on most modern Unices)
- libdag (only required for Endace DAG capture)

© 2008 Carnegie Mellon University

Questions?

Ask now…

…or later:

- Brian Trammell <bht@cert.org>
- Chris Inacio <inacio@cert.org>

© 2008 Carnegie Mellon University

Assessing Disclosure Risk in Anonymized Datasets

Alexei Kounine* and Michele Bezzi†
*Ecole Polytechnique Fédérale de Lausanne
Lausanne, Switzerland
Email: alexei.kounine@epfl.ch

†Accenture Technology Labs
Sophia Antipolis, France
Email: michele.bezzi@accenture.com

***Abstract*—Sharing of log data is a valuable step towards the improvement of network security. However, logs often contain sensitive information and organizations are hesitant to share them. Anonymization methods are used for increasing protection, lowering the disclosure risk to a level considered safe. Accordingly, a metric for anonymity is necessary to quantitatively assess the risk before releasing log data. In this paper, we propose a general framework for estimating disclosure risk using conditional entropy between the original and the anonymized datasets. We demonstrate our approach using network log files.**

I. Introduction

Log data analysis is a powerful tool for improving network security. Typically, each organization uses only their own logs, but, with the increasing number of coordinated attacks, sharing log information between organizations is becoming essential [1]. The problem is that organizations are often reluctant to share their logs, since the information contained in them can be sensitive. Anonymization methods are used for limiting disclosure risk in releasing such sensitive datasets. Anonymizing data increases protection, lowering the disclosure risk, but, it also decreases the quality of the data and hence its utility [2]. Finding the optimal trade off between risk and utility is the main scope of the anonymization process. Both quantities are hard to define, and strongly depend on context variables, e.g., data usage, level of knowledge of the attacker, amount of data released. In this paper we focus on the evaluation of disclosure risk (Section II). The main contribution of this paper is to introduce a general measure of disclosure risk, which is applicable to any set of masking transformations. Unlike previous measures, we do not assume any specific masking algorithm. Moreover, our measure provides a robust estimation of risk both at single record level (local risk) and at global level, i.e. for the whole dataset (Section III). Our model is therefore general and can be applied for quantitatively comparing different anonymization policies. Furthermore, it is directly related to the measure of the information lost in the anonymization procedure. We implemented this risk estimator using the FLAIM framework [3] and tested on network log files (Section IV).

II. Privacy in public datasets

Data holders, such as national statistical institutes, often have to release data files containing information on individual people or firms (micro-data) for research purpose. At the same time they have to preserve the privacy of individuals. This problem also occurs for sharing log files, since they may contain personal information which cannot be released in its original form (IP addresses, port numbers, timestamps, quantities...). Consequently, these data holders need to anonymize their databases before release, using data masking algorithms such as: generalizing the data, i.e., recoding variables into broader classes (e.g., releasing only the first two digits of the zip code or removing the last octet of an IP address), suppressing part of or entire records (also known as black marker [3]), randomly swapping some fields among original data records, applying permutations (one-to-one mapping on a defined set) or perturbative masking, i.e., adding random noise to numerical data values.

When masking methods have been applied, data holders have to quantitatively assess the disclosure risk (or anonymity level), to verify whether it is below a defined threshold, in which case it is assumed to be acceptable. To this scope, various measures for estimating disclosure risk have been proposed so far [4], [5], [6]; their validity strongly depends on the application scenarios considered, but still, there is a consensus that the risk of disclosure cannot be reduced to zero (but removing all the information). Thus, in general, a threshold should be determined to decide whether to release a dataset or not. Broadly speaking, there are two different approaches for assessing disclosure risk: estimating the *rareness* in the sample or population, or estimating the probability of re-identifying a masked record using some external information.

Let us examine these two methods in detail. In a typical scenario an attacker has knowledge about some variables, which may identify a record in the dataset. Considering the example of a medical database, the attacker may know a few attributes (age, gender, marital status) from an external public register (census data) or some private source of information (e.g., knowing age and address of his neighbor). He then tries to

(a) Original log file $\mathcal{S}$

SrcIP	SrcPort	DestIP	DestPort	Packets
168.125.253.23	80	147.81.124.173	3157	40
39.109.219.43	7310	142.68.227.108	59959	126
35.187.130.82	161	213.48.191.68	55867	83

(b) Anonymized log file $\mathcal{R}$

SrcIP	SrcPort	DestIP	DestPort	Packets
168.125.253.0	1023	10.1.1.1	65535	42
39.109.219.0	65535	10.1.1.1	65535	132
35.187.130.0	1023	10.1.1.1	65535	81

(c) Background knowledge $\widehat{\mathcal{S}}$

SrcIP	SrcPort	DestIP	DestPort	Packets
39.109.219.43	7310	142.68.0.0	—	—

TABLE I: Example of original ($\mathcal{S}$) and anonymized log files ($\mathcal{R}$). In the anonymization process, the least-significant 8 bits of the SrcIP are blacked out, BM(8) (replaced with 0s). SrcPort and DestPort are partitioned in two classes (1023 and 65535), called binary classification (C). DestPrt is completely blacked out, BM(32). Packets are perturbed with random Gaussian noise.

match these variables (*keys*) with the partly altered records in the released database. In the case of log files, an attacker may inject some information (e.g., scanning some specific ports), with the goal of later recognizing them in the anonymized logs. When a unique record matches a combination of key variables, the intruder can re-identify the masked record, assuming he is certain that the record is in the dataset. In fact, even if there is more than a unique match, but the number of linked records characterized by that combination of keys is still low (say it does not exceed a threshold *k*), these records have a high risk of re-identification. This rule is known as k-anonymity [7]. This approach has some limitations: it does not consider intruder's knowledge explicitly, and, in case of continuous variables the number of population uniques could be extremely large, especially when these data are randomly perturbed during the masking process.

The second approach consists of estimating the probability of re-identification. As in the previous case, the attacker aims at linking pairs of records in the released database with his background information [8], [4], [9]). This method permits to assess the risk in both categorical and continuous data: a record is considered at risk if this probability exceeds a fixed threshold. The main issue with this approach is finding a reliable strategy to compute these probabilities, since in case there are many records with similar, close to threshold, probabilities of re-identification, the risk estimation can be strongly affected by random fluctuations.

III. Entropy based risk estimator

The protection model we propose here creates a measure of disclosure risk for micro-data release, which combines together the two approaches described above. This allows us to develop a measure applicable in general cases (i.e. for any kind of data transformation, as when using the probability of re-identification method) and, at the same time, it considers the whole distribution of original records (as in k-anonymization). The basic idea is to use Shannon entropy as a measure of disclosure risk for a single record. Entropy metrics have previously been proposed for computing information loss [10], and, more recently for estimating disclosure risk for tabular data [11] and in network communications [12], [13].

In this section, we briefly review the theoretical framework and analyze its mathematical features. We refer the reader to [14] for a more extended discussion on the topic.

Let us consider a dataset $\mathcal{S}$ containing some sensitive data, e.g., network log files (Table I(a)). Each entry $s \in \mathcal{S}$ of this dataset is transformed using a data masking procedure, for example one or more of the ones mentioned in the previous section. The final result is an anonymized version of $\mathcal{S}$ dataset, which we call $\mathcal{R}$ (Table I(b)).

The attacker aims at re-identifying released data by linking them with some external information or background knowledge $\widehat{\mathcal{S}}$ (Table I(c)), which has some overlapping attributes with the released dataset. If the attacker is able to reconstruct some attribute values of the original record, we have a privacy breach. Because the data holder does not know in advance which records and attributes might be available to the attacker, it must run the risk analysis on the whole released dataset $\widehat{\mathcal{S}} \equiv \mathcal{S}$ and assume a set of key attributes (called *quasi-identifiers* in the k-anonymity framework) the attacker might know and use for re-identification. These key attributes can coincide with the whole set of attributes. The re-identification procedure consists of estimating for each $\hat{s} \in \widehat{\mathcal{S}}$ the probability of linking it with a record $r \in \mathcal{R}$: $P(r|\hat{s})$. Because we are assuming $\widehat{\mathcal{S}} \equiv \mathcal{S}$, thereafter we will consider the $P(r|s)$ instead of $P(r|\hat{s})$.

We can estimate this probability assuming the attacker simulates the data masking transformations [15], uses the information released by data holders (such as the structure of the noise added) or defines a distance function between records [16]. Intuitively, the more uncertain the mapping $P(r|s)$, the lower the disclosure risk. Shannon's entropy can be used to estimate this uncertainty. By applying it to the conditional probability $P(r|s)$, the conditional entropy is obtained:

$$H(\mathcal{R}|s) = -\sum_{r\in\mathcal{R}} P(r|s)\log_2 P(r|s) \tag{1}$$

This quantity measures the risk at the level of single record s. It represents the average number of binary question we have to ask to identify the corresponding r given s. Low entropy values indicate an almost deterministic mapping, and high risk accordingly, whereas large entropy is associated to low disclosure risk.

For example, in the case where a selected record s can be linked to exactly k_s indistinguishable records in $\mathcal{R}$ (as in k-anonymity [7]), we have a uniform distribution over the k_s records: and the corresponding specific entropy, Eq. (1) is:

$$H(\mathcal{R}|s) = \log_2 k_s \tag{2}$$

The k-anonymity condition over the whole dataset can be written as:

$$k \geq min_{s\in\mathcal{S}} 2^{H(\mathcal{R}|s)} \tag{3}$$

Global identification risk, that is at the dataset level, can be derived from the local risk measures, Eq. (1). One possible choice (see [14] for other options) is to calculate the expected number of correct matches (E_{CM}, herein):

$$E_{CM} = \sum_{s \in \mathcal{S}} \frac{1}{2^{H(\mathcal{R}|s)}} \tag{4}$$

E_{CM} is the *average number* of correct matches considering the intruder is randomly guessing according to $P(r|s)$. In fact, the entropy $H(\mathcal{R}|s)$ represents the average number of binary questions required to determine r, given s [14].

E_{CM} differs from the *estimated* number of correct matches, called N_{TM} herein, typically used for global risk assessment (see [15], [9]). These two measures differ because N_{TM} is based on maximum likelihood, which implies verifying a posteriori whether a match is correct, whereas E_{CM} is the average number of correct matches considering a random guess according to $P(r|s)$. So, the latter lacks the *decoding* part (i.e., the maximum likelihood step) and relies on the shape of the distribution only. In practice, they coincide when $P(r|s)$ has a single sharp peak, that is an almost deterministic one to one mapping. In contrast, they may strongly deviate in presence of multiple peaks and/or a smooth distribution. In addition, because E_{CM} depends on the shape of the whole distribution (not only on its peak value), it is less sensitive to random fluctuations [14]. Lastly, note that conditional entropy is directly linked to the mutual information between $\mathcal{S}$ and $\mathcal{R}$, and it can be used as an estimation of the information lost in the anonymization transformation [14].

IV. Measuring risk on anonymized network logs

We tested the entropy-based risk estimator on a publicly available Netflow log file [1]. In this analysis, we only use a limited set of fields and records. In addition we do not consider the utility of the masked dataset. Consequently, results presented here should be viewed as a proof of concept and recommendations for selecting a specific anonymization policy are not provided.

To run these tests, we developed a risk-estimating module based on FLAIM (Framework for Log Anonymization and Information Management) [3]. FLAIM is a modular and scalable framework for anonymizing log files which includes an anonymization engine with various anonymization primitives (BlackMarker, Permutation, Enumeration, etc ...). We developed a component, RiskEngine (see Figure 1), capable of estimating disclosure risk (Eq. (4)) by comparing the original and anonymized log files. As the other FLAIM components, the risk estimator works on streamed data, allowing us to process very large datasets.

A. Results

To illustrate the previously described method and its implementation in FLAIM, seven different anonymization scenarios are presented. As testing dataset we used the sample Netflow file available on the FLAIM website. The nfdump module provided in FLAIM is used for parsing the log file. We considered a subset of the available fields: the source and destination IPs, the source and destination ports and the number of bytes in a flow. The seven scenarios are summarized in Table II.

[1] Available at http://flaim.ncsa.uiuc.edu/downloads/flaim/sample.nfdump.log

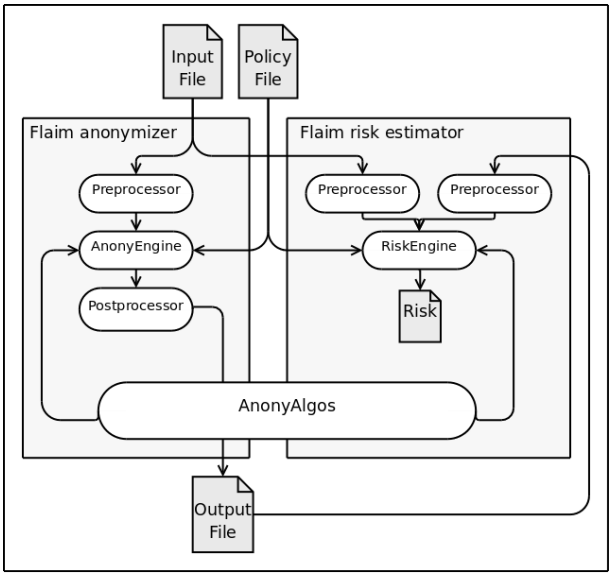


Fig. 1: The structure of the risk estimation component (RiskEngine). It is implemented as a subclass of AnonyEngine. The BasicPreprocessor and BasicPostprocessor classes were extended with interfaces to the RiskEngine. AnonyAlg and each of its subclasses now implement a method for estimating the probability $P(r|s)$ for each anonymization primitive

	SRC_IP	DST_IP	SRC_PRT	DST_PRT	BYTES
S1	None	None	None	None	None
S2	BM(16)	BM(16)	None	None	None
S3	BM(16)	BM(16)	C	C	None
S4	BM(16)	BM(16)	C	C	NA(10%)
S5	BM(24)	BM(24)	C	C	NA(10%)
S6	BRP	BRP	C	C	NA(10%)
S7	BM(32)	BM(32)	C	C	NA(10%)

TABLE II: List of the 7 anonymization scenarios discussed in the main text, in order of increasing anonymization *strength*. Legend: BM(16) (BM(24)): Black Marker applied on the 16 (24) least-significant bits. C: Classify: bins ports below 1024 in one bin and ports greater or equal to 1024 in another. NA(10%): Noise Addition: adds zero averaged Gaussian noise with a standard deviation equal to 10% of the value to anonymize. BRP: Binary Random permutation: maps each IP into a randomly generated IP in a consistent way (all IPs equal in the original log file are also equal in the anonymized log file). For more details about these transformations see Ref. [3].

Each anonymization primitive has its corresponding function for calculating the probability $P(r|s)$. For the sake of simplicity we assumed that the different fields are independent. Therefore, in the example above, $P(r|s)$ reads:

$$\begin{aligned} P(r|s) \quad = \quad & P(r|s)_{SRC_IP} \cdot P(r|s)_{DST_IP} \cdot \\ & \cdot P(r|s)_{SRC_PRT} \cdot P(r|s)_{DST_PRT} \cdot P(r|s)_{BYTES} \end{aligned}$$

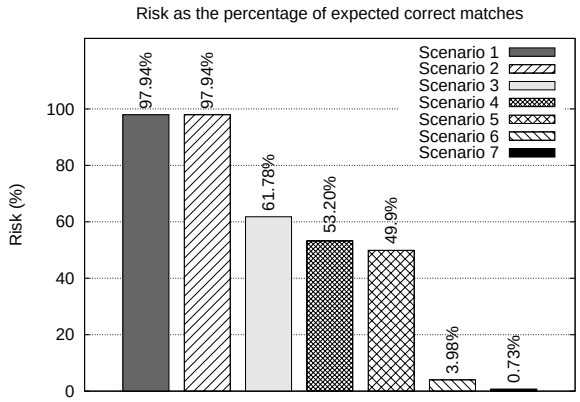


Fig. 2: Entropy-based risk for the 7 anonymization scenarios described in Table II

Figure 2 shows the expected number of correct matches, E_{CM}, as a percentage of the total number of records for the seven scenarios. Intuitively, increasing the number and strength of the anonymization methods leads to a reduced disclosure risk. In more details, we observed that removing the last 8 bits in the IP addresses have no impact on the estimated risk (Scenarios 1 and 2). Similarly suppressing the 16 or the 24 least-significant bits of the IP addresses lead to similar risk values (Scenarios 4 and 5). This indicates that, in this sample, most of the IP addresses sharing the same first octet are actually the same address (however, the corresponding port is not necessarily the same). In other words, due to a lack of diversity, most of the IPs can be identified by their first 8 most-significant bits. By generalizing the port number (Scenario 3), we observed a $\simeq 36\%$ decrease in the risk, suggesting that port re-coding could be a valuable anonymization strategy in this context. Adding random noise on the number of packets transmitted gives a further $\simeq 8\%$ decrease in the risk (Scenario 4). To obtain low risk values, we needed to remove most of the information contained in IP addresses by either using a one-to-one mapping into a predefined set (binary random permutation, scenario 6) or black marking all the 32 bits of the address (BM(32) in Scenario 7).

V. Summary

The advantage of using Shannon's entropy as a measure of disclosure risk for log file release is twofold: First, it can be applied to any general masking transformation, unlike k-anonymity measure, which is limited to non-perturbative masking transformations. Second, it only depends on the shape of the probability distribution; thus it is less sensitive to random fluctuations than measures where decoding of the masked record is needed.

The main technical issue is that computing the probability of re-identification can be hard for complex masking transformations [15]. Furthermore, these probabilities depend on the attack scenarios (attacker's knowledge, data sensitivity, etc ...), that are often difficult to model and application-specific. In the simple example we presented here, we could easily derive these probabilities under the assumptions of independence among fields and records. Both these hypotheses are unrealistic in many real world scenarios, such as in port scanning attack, where multiple ports are scanned in sequence on a single target host. Further analysis is needed to investigate the viability of this approach in realistic settings.

Acknowledgments

Alexei Kounine contributed to this study during his internship at Accenture Technology Labs in Sophia Antipolis. We thank Kiran Lakkaraju and Adam J. Slagall, for their help with FLAIM.

References

[1] A. Slagell and W. Yurcik, "Sharing computer network logs for security and privacy: A motivation for new methodologies of anonymization," 2005. [Online]. Available: citeseer.ist.psu.edu/slagell05sharing.html

[2] G. Duncan, S. Keller-McNulty, and S. Stokes, "Disclosure risk versus data utility: The RU confidentiality map," *Technical paper, Los Alamos National Laboratory, Los Alamos, NM*, 2001.

[3] A. J. Slagell, K. Lakkaraju, and K. Luo, "Flaim: A multi-level anonymization framework for computer and network logs," in *LISA*. USENIX, 2006, pp. 63–77.

[4] G. Duncan and D. Lambert, "The risk of disclosure for microdata," *Journal of Business & Economic Statistics*, vol. 7, p. 207, xx 1989, 10.2307/1391438. [Online]. Available: http://dx.doi.org/10.2307/1391438

[5] T. M. Truta, F. Fotouhi, and D. Barth-Jones, "Assessing global disclosure risk in masked microdata," in *WPES '04: Proceedings of the 2004 ACM workshop on Privacy in the electronic society*. New York, NY, USA: ACM Press, 2004, pp. 85–93.

[6] R. Benedetti and L. Franconi, "Statistical and technological solutions for controlled data dissemination," *Pre-proceedings of New Techniques and Technologies for Statistics*, vol. 1, pp. 225–232, 1998.

[7] L. Sweeney, "k-anonymity: a model for protecting privacy," *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, vol. 10, no. 5, pp. 557–570, 2002.

[8] W. E. Yancey, W. E. Winkler, and R. H. Creecy, "Disclosure risk assessment in perturbative microdata protection." in *Inference Control in Statistical Databases*, ser. Lecture Notes in Computer Science, J. Domingo-Ferrer, Ed., vol. 2316. Springer, 2002, pp. 135–152.

[9] C. J. Skinner and M. J. Elliot, "A measure of disclosure risk for microdata," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 64, no. 4, pp. 855–867, 2002. [Online]. Available: http://www.blackwell-synergy.com/doi/abs/10.1111/1467-9868.00365

[10] L. Willenborg and T. de Waal, *Elements of statistical disclosure control*. Springer New York, 2001.

[11] A. Oganian and J. Domingo-Ferrer, "A posteriori disclosure risk measure for tabular data based on conditional entropy," *SORT*, vol. 2, pp. 175–190, 2003.

[12] A. Serjantov and G. Danezis, "Towards an Information Theoretic Metric for Anonymity," *Privacy Enhancing Technologies: Second International Workshop, PET 2002, San Francisco, CA, USA, April 14-15, 2002*, 2003.

[13] C. Diaz, S. Seys, J. Claessens, B. Preneel, and K. ESAT-COSIC, "Towards Measuring Anonymity," *Privacy Enhancing Technologies: Second International Workshop, PET 2002, San Francisco, CA, USA, April 14-15, 2002: Revised Papers*, 2003.

[14] M. Bezzi, "An entropy-based method for measuring anonymity," in *Proceedings of the IEEE/CreateNet SECOVAL Workshop on the Value of Security through Collaboration*, Nice, France, September 2007.

[15] J. P. Reiter, "Estimating risks of identification disclosure in microdata," *Journal of the American Statistical Association*, vol. 100, pp. 1103–1112, December 2005, available at http://ideas.repec.org/a/bes/jnlasa/v100y2005p1103-1112.html.

[16] A. Narayanan and V. Shmatikov, "How to break anonymity of the netflix prize dataset," Oct 2006. [Online]. Available: http://arxiv.org/abs/cs/0610105

Integration of Context into Data Analysis and Visualization

Ashley Thomas, Uday Banerjee
SecureWorks

Existing approaches to Analysis

Existing workflow in a typical environment

- Mostly analyze data from separate sources (IDS/IPS/Firewall/Syslog/etc.) in a semi-integrated textual view (SIM).
  - Although the view may be integrated, typically the correlation is left up to the analyst. This is typically a complex task, demanding continuously high levels of cognition, and may lead to incomplete analyses.
- Analysts have to reference other tools (IDS signature details, packet captures, historical information, etc.) to make the proper determination)

Existing approaches to Analysis (contd.)

Most commercial environments are SLA driven, so no motivation to use ‘yet another tool’ (read 'visualization') to perform analysis.

- Millions of alerts per day
- High rate of false positives from alerts in the field
- Limited number of analysts
- Time spent on each alert is very limited
  - quality of analysis affected

Existing approaches to Analysis - Data Visualization

A well studied field:

- Several tools documented here: http://www.vizsec.org/applications

Visualization has faced problems with getting adapted into a typical analyst's workflow

- Tool is not purpose built for the environment
- Flexibility (is not always there to build your own visualizations)
- Performance (of viz tools is very important. A slow tool is going to be abandoned sooner or later)
- Gives the analyst a 'free flow' exploration of the data, but depends on him/her for finding the needle in the haystack. There is a need for some additional context to be provided to the analyst.
- Most systems just allow for exploration of data, but do not allow for inferences to be translated into 'work done'. In a typical commercial environment, SLAs dictate workflows, and the ROI on a given tool (investment = time spent as part of analysis, return=inference that other tools in the workflow did not give us) needs to be very high in order to become a standard part of the workflow.

Cross Platform Data Analysis

The ultimate goal is to have a unified data set that can be analyzed across different services, devices, applications, etc.

- Normalize data from different sources (IDS alerts, traffic flows, firewall and application logs, etc.)
- Extract context where applicable and present to the Analyst
- Visualize this data and present the ‘Big Picture’
- Allow the Analyst to resolve these events in the visualization GUI itself

A sample alert: analysis

[**] [1:648:7] SHELLCODE x86 NOOP [**]
[Classification: Executable code was detected] [Priority: 1]
08/09-15:46:51.632771 192.168.1.121:54835 -> 192.168.1.136:80
TCP TTL:64 TOS:0x0 ID:3403 IpLen:20 DgmLen:457 DF
AP Seq: 0x1E0C3C55 Ack: 0xB33C734D Win: 0x5C TcpLen: 32
TCP Options (3) => NOP NOP TS: 1221541188 17773996
[Xref => http://www.whitehats.com/info/IDS181]

- An example alert:
  - Server vulnerable?
  - False positive?
  - Attempt successful?

More context; Better analysis

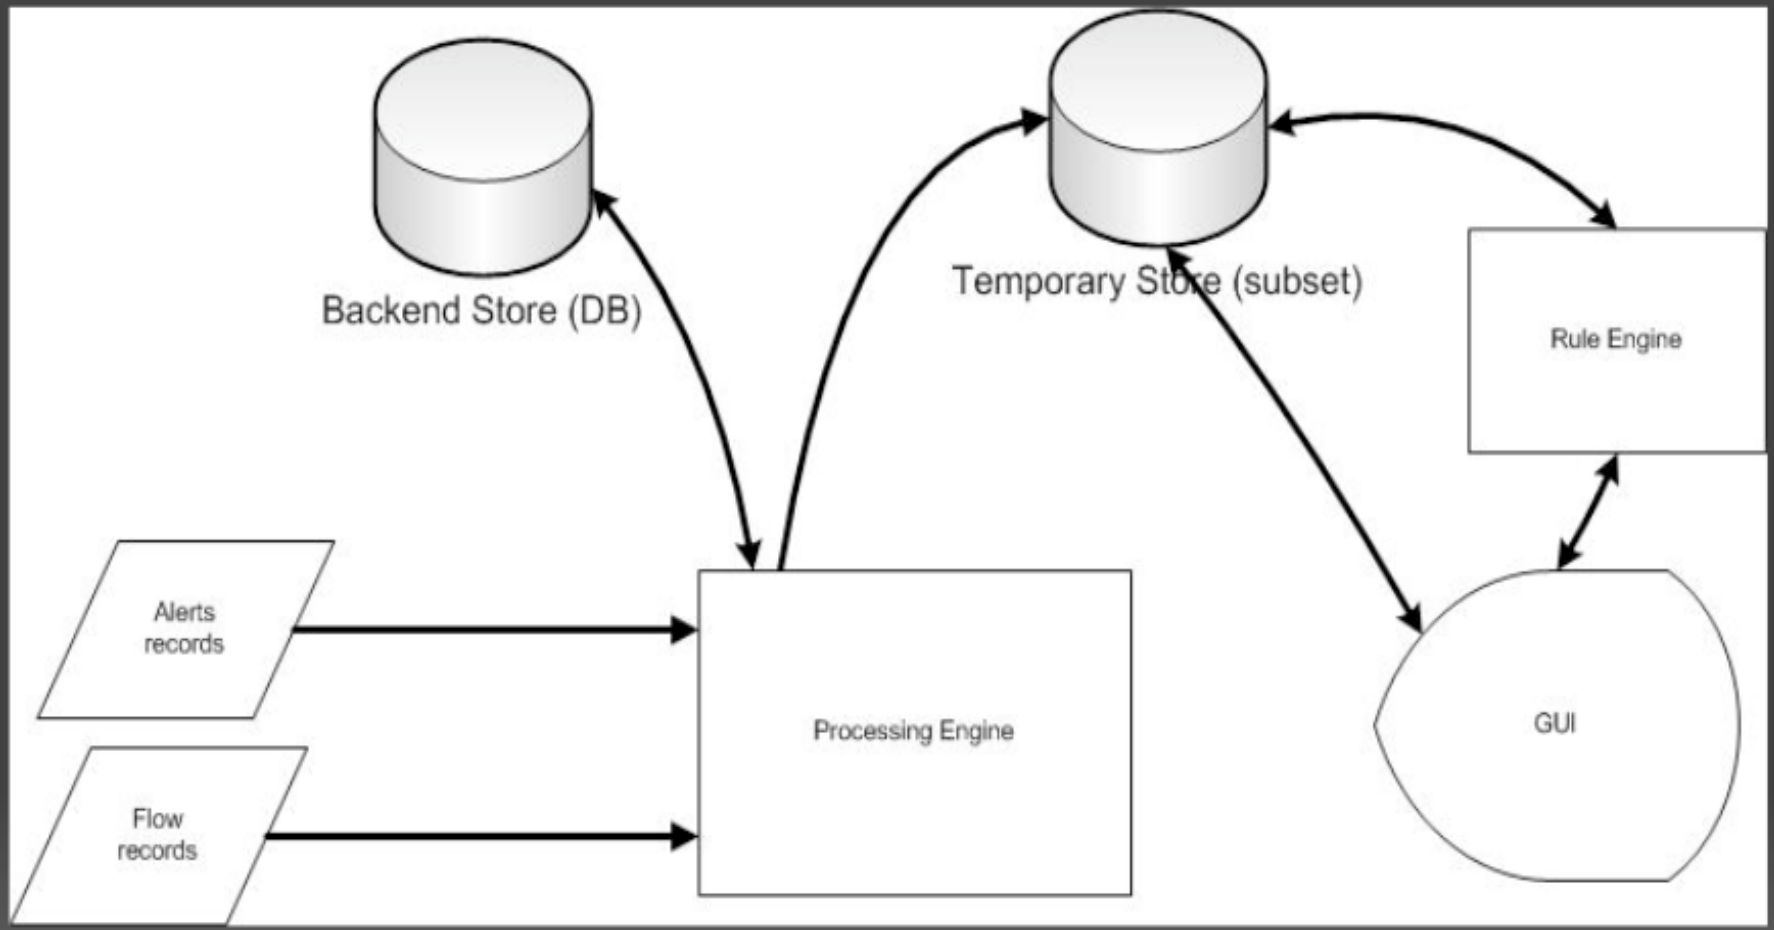
Flow record right after
08/09-15:47:12 192.168.1.136 209.185.243.135 TCP 2255 21 2666 15MB <snip>

FTP download by the server from an unknown site - suspicious.

An integrated platform

- Correlating alerts, flows, logs into the same platform
  - More context; better analysis
- Ability to visualize data flexibly (Analyst can override default visualizations and create new ones - e.g. Bar Chart over Pie Graph)
- Ability to drill-down/up based on time, ip address, other variables
- Provides guidance (via predefined rules)
- Integration of the analysis and taking action (ability for the Analyst to resolve events via the visualization interface)

Architecture



Architecture: Processing Engine

- **Ability to integrate and correlate.**
- **Ability to zoom-in**
- **Plug-in architecture:**

Each record type that is supported will be handled by an appropriate plug-in.

  - **IDS alert plug-in**
    - **isensor IPS**
    - **Snort IDS**
    - **cisco**
    - **mcafee**
  - **netflow record plug-in**
  - **Firewall plugin.**

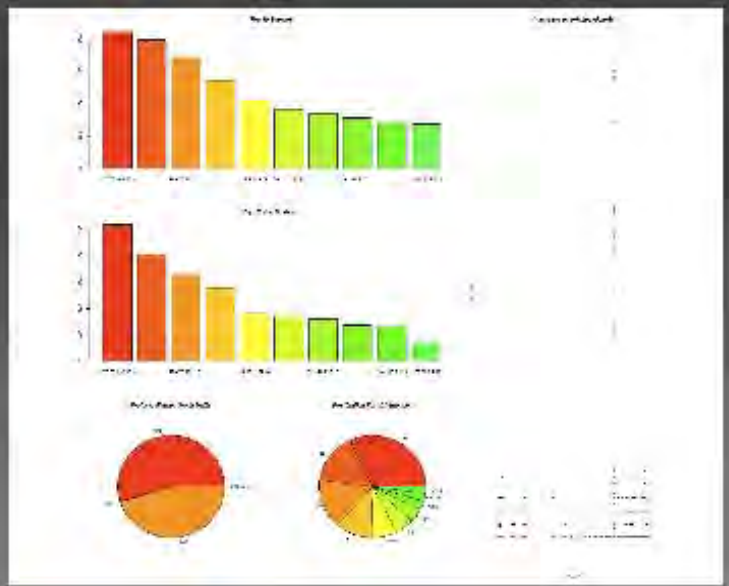
This plugin is aware of the formatting of each type of record. Finally when it is stored into the DB it is stored in a consistent fashion.

Architecture: Rule Engine

- Provides guidance
  - Simple predefined rules search the data for the existence of certain conditions, and highlight **certain records or flows in order to provide guidance** to the Analyst if applicable. Some examples below:
    - TCP Syn packets to external addresses on port 135/139/445
    - Change in threshold of flow activity (>50%) for a given host in a time window
    - Outbound activity to port 25 to Yahoo, AOL, Hotmail, etc. mail servers
    - Traffic directed to bogon IP addresses
- **Temporary Store:**
  - **Data subset for a certain window of time, e.g. (now – 2 hours ago). This may be the data the analyst will work on.**

Architecture: Visualization interface

- Flexible, fast interface that allows drill-down/up capability and the ability to assign a determination to the result set
- Consists of a 'parameter' section that allows the Analyst to shape the data set to be visualized (basically creating a SQL query)
- Once this query is submitted, the resultant data set is visualized using a set of default templates



Architecture: Visualization interface (contd.)

- The Analyst has the flexibility to change these default visualizations to something they feel could be more appropriate.
- R (www.r-project.org) was our first choice to display the graphics
  - Areas of investigation: Interactive images (Image Maps) that allow for 'click and drill down', better suited packages to display some relationships (Lattice for portscans, etc.)
  - Commercial tools exist that do a very good job of visualizing data (but external development can be an issue) (e.g. www.advizorsolutions.com, www.vizsec.org/applications/commercial-applications)

Architecture: Doing work

- The tool also enables the analyst to take action from the same GUI front end.
  - This may improve efficiency and speed of analysis
  - Allows the Analyst to resolve events in a larger scale
    - Mark all events from a source IP as benign (e.g. known scanner)
    - Escalate all events from a given source IP (established to be a known bad IP after analysis).

Discussion & Q/A

athomas@secureworks.com
ubanerjee@secureworks.com

Lawrence Livermore National Laboratory

Analysis of Network Beaconing Activity for Incident Response

FloCon2008



Peter Balland

DOE Computer Incident Advisory Capability (CIAC)

Lawrence Livermore National Laboratory, P. O. Box 808, Livermore, CA 94551

This work performed under the auspices of the U.S. Department of Energy by
Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344

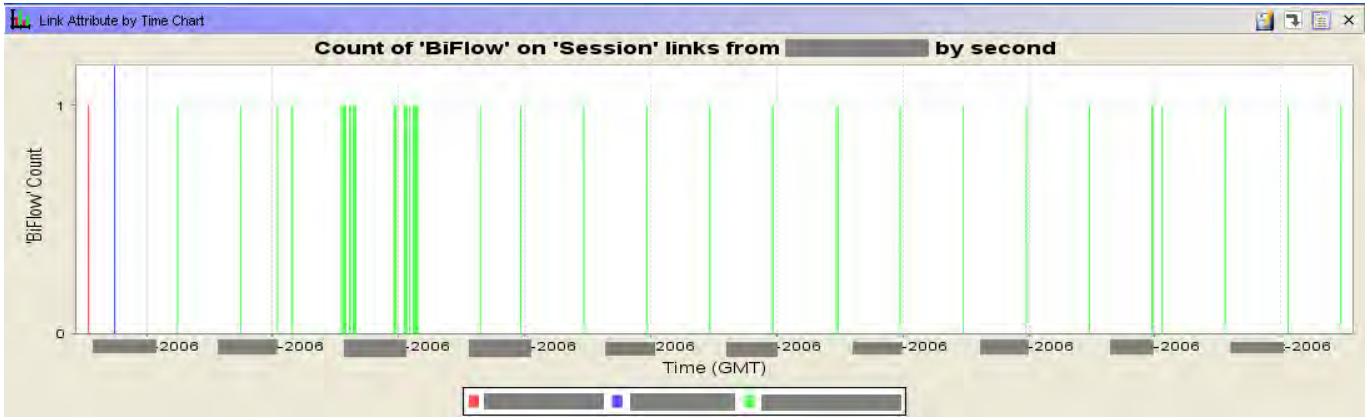
UCRL-PRES-236878

Background

- CIAC provides 24x7 “on-call” operational cyber security services to the Department of Energy (DOE)
- CIAC’s Mission:
  - *Prevent* cyber incidents whenever possible
  - Perform predictive analysis to *Watch and Warn* for any real or potential threats to DOE
  - Assist in the *Response* and restoration of operations should and incident occur
- CIAC collaborates with local site security personnel and other cyber security agencies

Motivation for Identifying Network Beaconing

- We seek additional indicators of malware infection to support proactive incident detection as well as to supplement incident response and forensics efforts.
- Analysis of previously identified incidents has uncovered network sessions sharing common characteristics that recur at regular intervals. We identify this as “network beaconing activity.”



Network Beaconing Detection Strategy

Our objective is to detect the following intrusion scenario:

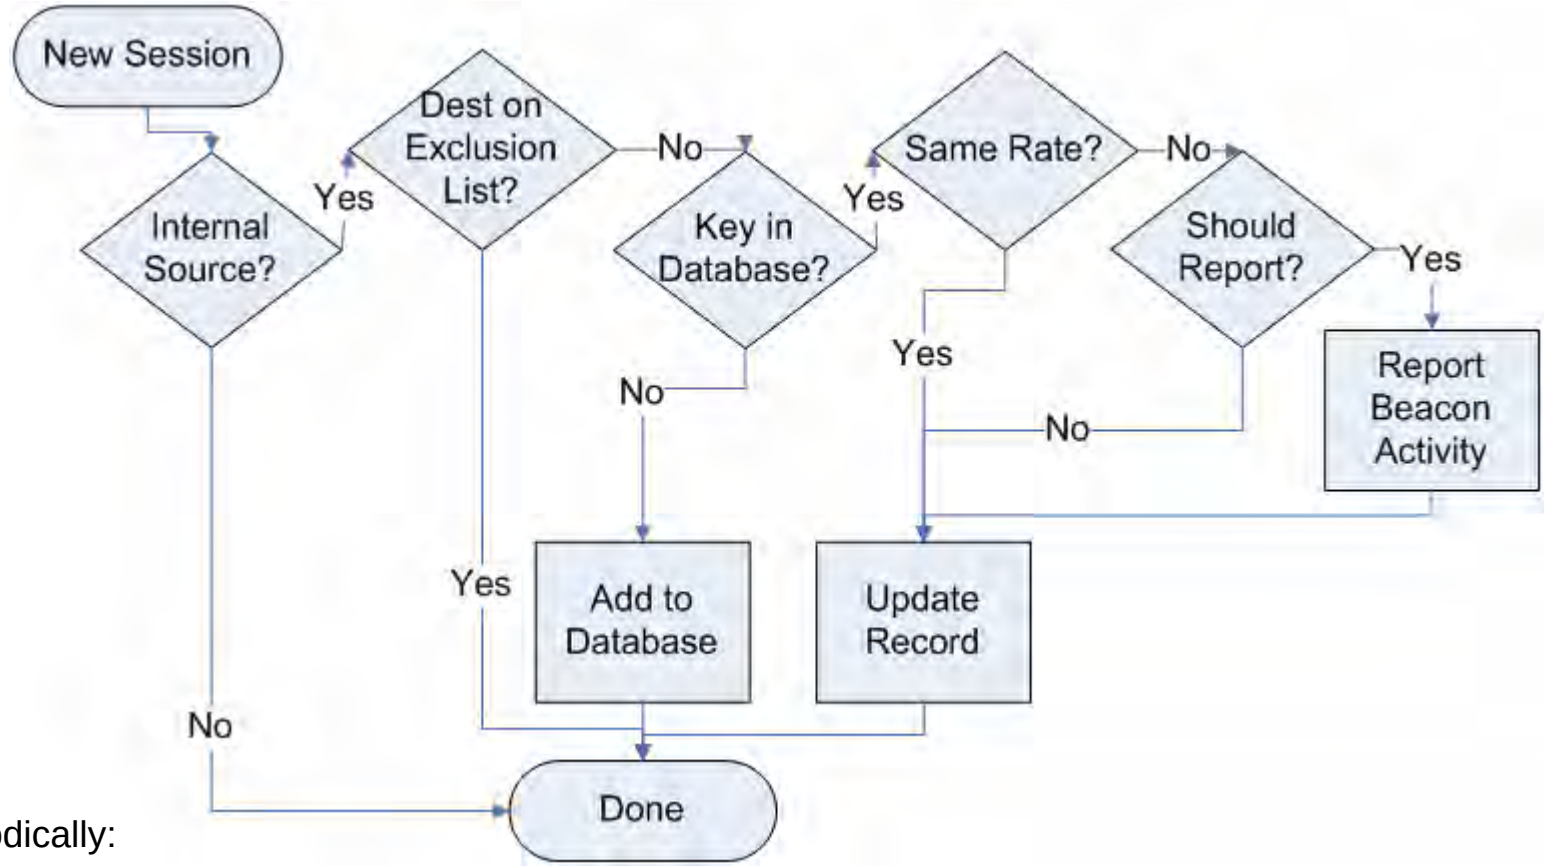
- Malware delivered via phishing email, drive-by-download, etc.
- Malware attempts connection to an unknown controller
  - If controller is not available, malware sleeps for a fixed duration and retries connection

We use this retry interval as an indicator of possible malware activity

Discovery Methodology : Overview

- Aggregate flow session summaries into bi-directional records and order by start time
- Check each session against whitelist criteria
- Maintain a database of inter-session times for each source and destination IP; update for each new session
- Report session groups that match a threshold of network beaconing activity

Discovery Methodology : Logical Flow



Periodically:

- Report and prune stale records
- Report ongoing records

Discovery Methodology : Aggregate Session Information

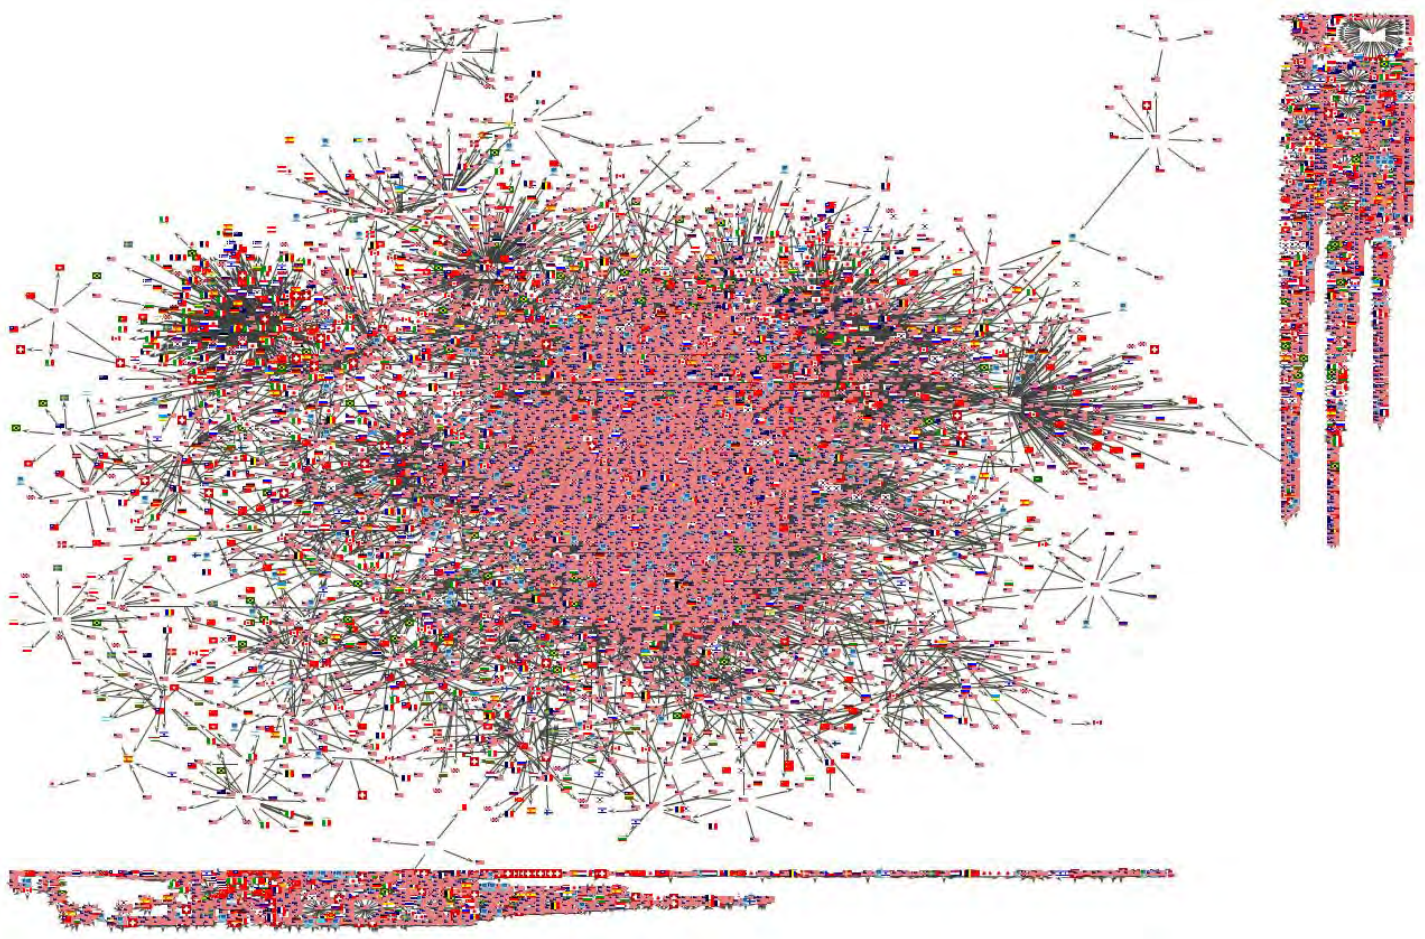
Flow Record

- Source IP
- Destination IP
- Protocol
- Source Port
- Destination Port
- Source Bytes
- Destination Bytes
- Source Packets
- Destination Packets
- Source Flags
- Destination Flags
- Flags of 1st Packet in Session

Database Record (61 Bytes)

- {Source, Destination} IP (Key)
- {Start, End} Timestamp
- Session Count
- First Seen Protocol
- Is Multiple Protocols
- First Seen {Source, Destination} Port
- Is Multiple {Source, Destination} Ports
- {Source, Destination} Bytes Mean
- {Source, Destination} Bytes Std Dev
- {Source, Destination} Packet Count Total
- {Source, Destination} Flags (Logical OR)
- Session Starting With SYN Count

Results : Qualitative



Beacons identified one day of November, 2007

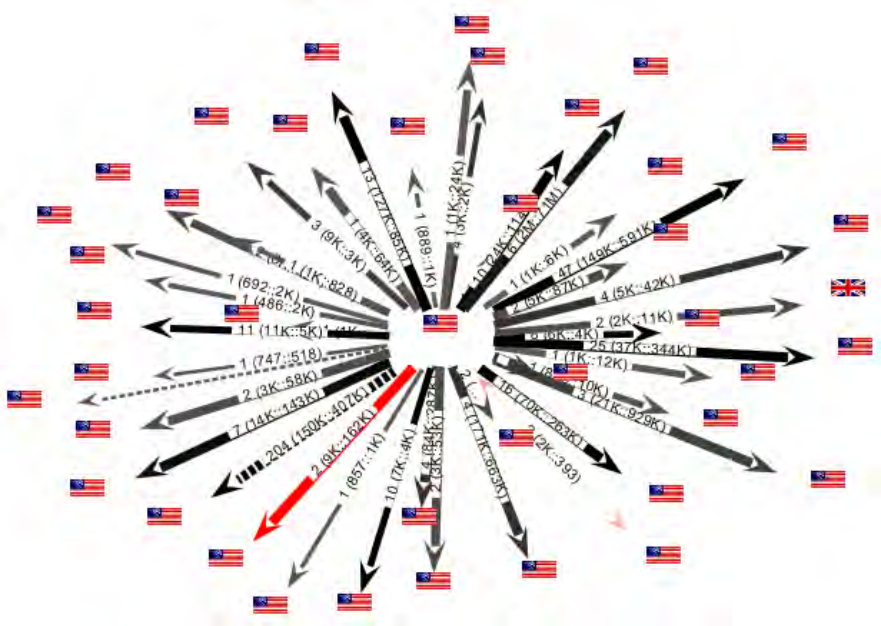
57,258 Beacon Records, 17,706 IPs, 21,224 Src-Dst IP Pairs

Results : Quantatative

- Prototype script using Perl + Berkeley DB on 2.8GHz Xeon Processor processes ~4800 sessions per second
- Midday on a work day in November 2007:
  - ~500,000 unique “active” internal IP addresses monitored
  - 2,351,565 unique src-dst pairs being tracked
  - ~1GB disk space for Berkeley DB database files (~140M raw data size)
- A week in November 2007:
  - 732,959 beacon records generated
    - 14,842 unique source IPs
    - 74,753 unique destination IPs

Analysis Methodology : Incident Response

- If compromised host is identified, past beaconing behavior of host may provide a toehold into the start of the intrusion
- If malicious IP is identified (watchlist, other intrusion, etc), beaconing activity to that IP may warrant additional concern.



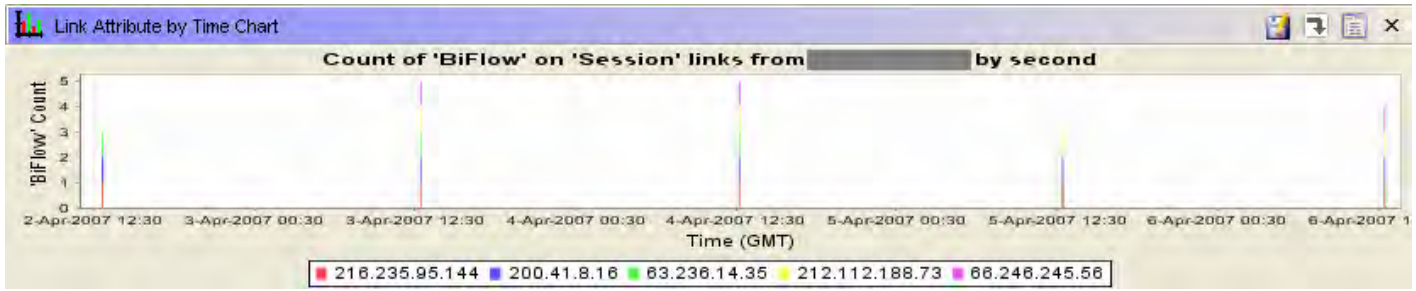
Graph view of Netflow (black), intrusion detection (red), and beaconing (dotted) records from a host.

Incident Detection : False Alarms

- Network beaconing activity is prevalent in many applications and protocols (NTP, RSS Feeds, automated software patching, etc)
  - Can be somewhat mitigated by whitelisting “trusted” IP addresses
- Keep-alive traffic in long lived sessions may appear as beacons
  - For TCP traffic, we can investigate the Flags field
- Does adware on a host constitute a false alarm? What about spyware?

Analysis Methodology : Incident Detection

- Rank identified beacons by how ‘interesting’ they are
  - Attempt to determine the cause of the beaconing
    - Significantly helped by domain knowledge of internal hosts, software configuration, security policy, and acceptable use policy
- In our experience of proactive investigations, fewer than 5% of beacons investigated were determined to be malicious. Several potential policy violations identified.

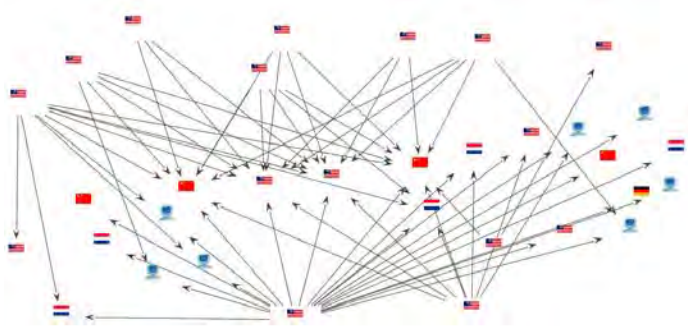


Interesting beaconing to 5 hosts worldwide. Later explained by a popular media player refreshing ads.

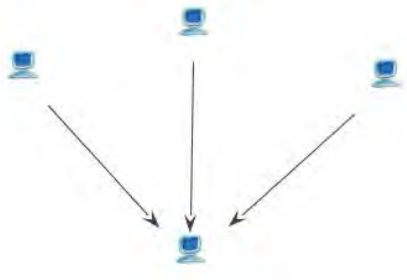
Incident Detection : What's Going On?



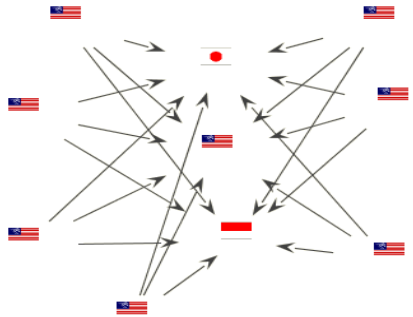
Two Hosts beaconing to 262 hosts (TCP 2170) over several hours with large response bytes.
[globus]



Several hosts beaconing to multiple destinations on TCP and UDP; some beacons never respond
[peer to peer download manager]



Three Hosts beaconing to a host (TCP 80) every 3 hours.
[i********** spyware phoning home]



Seven Hosts beaconing to 3 hosts (TCP 30000) over several hours with no response.
["canadapost" shipping module ???]

Conclusion

- Identification and analysis of network beaconing activity in flow data was readily achievable in our environment.
- Network beaconing logs have provided us with additional indicators that support incident detection and forensics.
- A high false positive rate hinders conclusive findings in the absence of additional evidence.
- When combined with other available security indicators, network beaconing activity has led to the discovery of network misconfigurations, policy violations, and compromises.

Useful resources

- Usual Internet Metadata (Whois, Search Engines, etc)
- Passive DNS Repositories
- Detailed host usage information (server, desktop, honeypot, etc)
- A really quick way to slice and dice lots of data

On the Anonymization and Deanonymization of NetFlow Traffic

Michalis Foukarakis*, Demetres Antoniades*, Spiros Antonatos*, Evangelos P. Markatos*
*Institute of Computer Science (ICS),
Foundation for Research and Technology Hellas (FORTH),
P.O BOX 1385, Heraklion Crete,
GR7-1110, Greece
{mfukar,danton,antonat,markatos}@ics.forth.gr

***Abstract*—Netflow is an efficient and flexible mechanism to collect network data and share them to security applications that require distributed knowledge. As information sharing breeds the danger of revealing user and network information, anonymization of Netflow data has to be applied before they are shared. To accomodate anonymization needs we have developed *anontool*. *Anontool* allows per-field anonymization up to the NetFlow layer offering a wide range of primitives to choose from. However, although we have the tools to perform anonymization, work needs to be done on the policy part. Some policies may be proven weak, if a third party can deanonymize the data and reveal user information. We demonstrate 2 possible attacks on anonymized traces. The first one is called active fingerprinting, where a malicious user injects packets that she can identify in an anonymized trace and thus reveal the mappings of IP addresses. The second one is the disclosure of web pages that users access based on the flow sizes recorded in a trace. We also present solutions to these attacks along with two anonymization primitives which we have implemented in *anontool* in order to defend against such attacks, to the extent possible.**

I. Introduction

As computer networks evolve and grow in size, the need for distributed network management and monitoring becomes more important than ever. Network activity log sharing has gained significant popularity recently, not only among computer security engineers and administrators, but also among researchers, developers and educators. To accommodate this increasingly popular need for information sharing as well as the fundamental lack of trust between members of different communities, several tools have emerged which enable their users to anonymize potentially sensitive information within those logs.

A popular format used by network activity logs is the Cisco NetFlow [1] format. The NetFlow format is based on the concept of a flow, which Cisco has defined as a set of packets that have the following five properties in common: source and destination IP address, source and destination port numbers, as well as the IP protocol field value. The most recent evolution of the Netflow format is version 9, which is currently the basis of the IETF [2] standard for information export. Given this fact, NetFlow is likely to gain even more in popularity.

Many diverse tools and techniques have been implemented for anonymization purposes. Most of them, however, are customized for specific purposes or provide limited functionality. Our approach, *anontool*, is a general-purpose tool that can anonymize live or stored traffic. Using Anontool, a user can choose from a large variety of functions to use on each and every application field, to implement her anonymization policy of choice. *Anontool* supports a number of protocols, but in this work we will focus on the NetFlow protocol.

Despite what tools one might use, the final result log of the anonymization process is likely to be publicly available. There is always the potential for an adversary to be able to infer a large amount of information from the log, if the anonymization policy is not chosen carefully. We will describe a few scenarios where an adversary can manipulate the log anonymization process in order to deduce useful information about the original trace. Given those attack scenarios, the need for techniques to defend against them becomes clear. We therefore describe two anonymization primitives which can aid in protecting against such attacks. A proof-of-concept implementation of them has already been incorporated into the latest version of *Anontool*.

II. Anontool Description

Anontool is a command line tool which enables users to anonymize both live and stored traffic. Its functionality is based upon the Anonymization API [3], in short AAPI. AAPI allows users to write their own anonymization applications. They can define which anonymization function be applied on any field, having complete freedom in deciding their policy. It provides a large set of anonymization primitives, from setting fields to constant values and performing basic mapping functions to prefix-preserving anonymization and several hash functions and block ciphers, as well as support for regular expression matching and replacement. AAPI can operate on a wide variety of protocols, ranging from Ethernet to HTTP and FTP in the application layer. All protocol fields are being made available to the user application.

AAPI has been implemented as a user-level library in the C language; it provides function calls for creating packet "streams", filtering using BPF filters, and of course applying anonymization functions. One of its main design goals was to accomodate extensibility, and potential developers are able to write their own protocol decoders similar to the ones already available for HTTP or NetFlow protocols, such as SMTP, or

their own anonymization primitives. It is also straightforward to write code that supports new input sources with few code additions.

Since the NetFlow format for packet export continuously gains on popularity [4], [5], [6] we have decided to extend AAPI, and subsequently Anontool, with support of the Cisco NetFlow packet export format. We took advantage of AAPI's extensibility and implemented decoding and anonymization functions for both version 5 and the newly defined version 9 of the NetFlow format. We take full advantage of the template-based nature of the NetFlow v9 format, to accurately provide the user with complete control of every field made available from information export nodes, such as Cisco routers or network monitoring applications that support the NetFlow export format, even in the event NetFlow templates change during a monitoring period.

Anontool is a fairly simple C application that makes use of the AAPI library to support anonymization of packet traces. It does not implement any anonymization functions in itself; it is much more transparent and less error prone for all the anonymization functionality to reside inside the AAPI implementation. It provides users the choice of protocols and functions to apply in order to create their anonymization policy. *Anontool* also implements some functionality such as preprocessing a trace to extract information which may be consequently used in the actual anonymization process; we will explain in further detail in Section IV.

III. Attacks Against NetFlow Anonymization Policies

In this section, we describe in detail two attacks against conventional anonymization policies and outline related work in both packet traces and flows.

We assume that the anonymization process is applied onto NetFlow traces, which are, in the general case, generated from a router at the border of a monitored network, and export information for traffic entering and exiting this particular network. The network could be of any size or topology; from a small home network to larger networks belonging to research institutes, universities, and so on.

In our threat model, we assume an active adversary that is able to direct traffic to the monitored network at will, has knowledge of the address space it occupies and can potentially compromise hosts inside it. We assume a rational attacker, for whom it is less costly, or more useful even, to "probe" and profile the monitored network before mounting attacks against it. The adversary may also have several external hosts under her command. She is also able to gain access to the anonymized traces, which will most likely be publicly released.

The first attack, which we call *"Active Fingerprinting"*, aims to break the mapping algorithm when used on IP addresses. Mapping takes the set of IP addresses in a trace and performs a simple mapping function onto another totally different set. The second set may be the output of a deterministic function seeded by a random quantity, such as the *drand48()* family of functions, or a very simple sequential assignment of unique IP address numbers which results in a one-to-one mapping.

The second attack aims at using the information about flow sizes contained inside NetFlow traces in order to deduce information about either hosts inside the monitored network or hosts that may be outside it, such as their IP address, network usage profiles, etc. We name this attack *"Statistical Signature Inference"*.

A. Active Fingerprinting

The idea that active fingerprinting exploits is that the mapping between real and anonymized IP addresses is one-to-one. Consequently, if the mapping on one flow is discovered, the mapping on the whole trace is compromised. This attack has been described on packet traces in [7], [8] and we demonstrate its applicability on NetFlow records below.

Using this idea, an adversary can establish flows from a host under her control, which resides outside the victim network, to one or more victim hosts inside it. These flows will appear in the anonymized trace. The challenge for the adversary is to construct those flows in such a way that they will be easily distinguished in the final trace. This can be accomplished in a variety of ways; she can craft a flow with specific attributes which are known not to be anonymized (the list is as large as the potential fields listed in a NetFlow record, and may usually include TCP flags, IP ToS, and so forth), or in the unlikely scenario where flows are fully anonymized, she can use temporal patterns which are easy to detect. Even a specific packet size can be used as an identifier for the packets involved in a dictionary attack. If the traces from the monitored network span a wide enough time period, the latter attack is very feasible as the trace contains a large number of attack packets.

A trivial way an adversary could create these flows is to perform a SYN scan on the victim network's address space. In this case, even if there is a clear temporal pattern which is easily detectable in the anonymized trace, it can be defeated in an easy way. Setting only the SYN bit in TCP flags and setting the number of bytes to a specific quantity makes the adversary unable to distinguish live hosts from unused address space.

We discuss a more general measure to defend against this kind of attack in Section IV.

B. Statistical Signature Inference

The idea behind this kind of attack is that each web page has a unique and complex enough structure which allows them to be identifiable despite our best efforts to anonymize their presence in NetFlow logs and preserve useful information in them as well.

A naive first effort would be the following. Consider the web sites **interesting.com** and **newssite.com**, and that a web session with each of them is **n** and **m** bytes long, respectively. The adversary can use one of the hosts under her command to initiate a web session to these sites and view the NetFlow records for source and destination IP addresses, port numbers,

and the total size of traffic exchanged. Assuming web page sizes do not radically differ from one session to another, and that NetFlow data records TCP traffic in its entirety, it is possible to filter out the set of web browsing sessions from an anonymized trace and construct a frequency histogram with the number of bytes transferred in each flow. According to our assumptions above, it is possible to see a great deal of flows around the values of **n** and **m**. The adversary can employ the same tactic to find those flows, and then gather further information about hosts inside the monitored network, which can then be used to answer questions such as: *"What web sites does host A visit?"*, *"Which hosts do frequently visit www.google.com?"*, and make user profiles.

Past work [8] has demonstrated this kind of attack on packet traces. Recent work [9] has extended and demonstrated this attack on NetFlow logs as well. We argue that the fundamental property of web sessions that allows this kind of analysis to be exploited is the fact that web sessions to different hosts produce flows with similar characteristics, especially in flow size. In Section IV we are going to view our proposal of a primitive which deals with this issue.

IV. Countermeasures

The previous section described two attacks for revealing sensitive information from an anonymized NetFlow trace. In this section, we will describe our proposed ways to deal with the aforementioned attacks, and evaluate their consistency.

A. Bidirectional Mapping

We propose a way to deal with the issue that does not iteratively consider all the combination of fields an adversary may use to craft her flows. Instead, we aim to eliminate the one-to-one mapping property without losing all of the information the trace can provide. To this goal, we propose a bidirectional mapping to be used, that is different mapping for each traffic direction. Let **A** be the IP address of a live host inside the monitored network, and **B** the IP address of a host outside the monitored network. Conventional mapping functions would map a flow *(A, B)* to *(A", B")* and a flow *(B, A)* to *(B", A")*. Bidirectional mapping maps a different address to **A** according to the direction of the flow that involves it. In the case of our example, the flow *(A, B)* would be mapped to *(A", B")* yet the flow *(B, A)* would receive a different mapping, say *(C, D)*.

Using this anonymization scheme we prevent the attacker from identifying her own network flows inside the anonymized trace. Thus, it is made impossible to correlate her data with the trace information and reveal any sensitive data from it. Also, most of the statistic information derived from the trace remains the same. We can still gather information about the incoming and outgoing traffic of the organization and identify the producers and the consumers of the network. Correlation of incoming and outgoing traffic for a specific IP address can not be done, but we argue that this is a general trade-off of the anonymization process and is up to the organization to decide whether to sacrifice sensitivity for usability in the data it makes public.

The implementation of such a primitive is quite easy, and it is already included in the stable version of *Anontool*.

B. Random Value Shifting

In order to diminish the viability of a statistical identification approach, but still be able to calculate some basic representative statistics about a NetFlow log, we propose a randomized shifting of values, which we will describe below.

Given a NetFlow data field with a given value range, our intent is to "scramble" its values across the NetFlow log to the point that we make an adversary unable to distinguish between two web sessions with the same web site and two web sessions with web sites that have similar web pages in size, but not as much as to destroy all the useful information a log may provide. More specifically, we intend to preserve metrics such as arithmetic averages and standard deviation, as well as other descriptive statistics. On the other hand, we wish to obfuscate inferential statistics, so that an adversary would be unable to reach conclusions that extend beyond the immediate data alone.

For clarity, we are going to use the flow size NetFlow field as an example from here on.

Our method is to add to the value of the flow size field a random value. This value is chosen uniformly at random from a fixed range *[-d, d]*. One of the basic properties of our choice is it allows us to directly preserve the arithmetic average and standard deviation of the original distribution of flow size values. The parameter d may be chosen arbitrarily, but we will demonstrate the importance of an educated choice with an example. Consider, as an elementary example, three flows with sizes of *15*, *17* and *25* bytes, which repeatedly occur within a NetFlow log. Choosing d to be equal to two, this is what happens in the anonymized trace: The flows with the initial size of *15* now occur with flow sizes from thirteen to seventeen. Flows with the size of *17* bytes now occur with sizes from fifteen to nineteen. These two groups of flows are now "mixed up"; what happens is that the confidence intervals for the random variables which represent the flow sizes of each flow are now different, and they overlap. On the other hand, that is not the case for the flow with size *25* bytes. It now occurs on the anonymized trace with values from 23 to 27. An adversary is still able to distinguish this flow from the other two with relative ease. Now we can easily conclude that a proper choice for the parameter d will have to take the entirety of the NetFlow log into consideration. This is an interesting topic for future work, which we will not further explore in the rest of the paper due to lack of space.

To verify our assumptions about the descriptive statistics of a NetFlow log being preserved after the application of random value shifting, we implemented it in *Anontool* and proceeded to process a NetFlow packet trace with it. Our choice for the parameter d was the minimum flow size observed, divided by 2. As we previously mentioned, this is most likely not a good choice for real world applications, but it is good

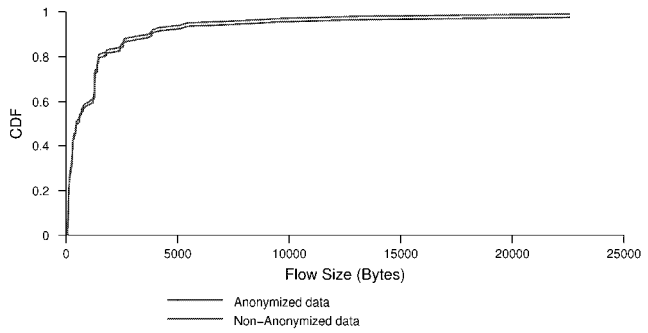


Fig. 1. Cumulative distribution function of flow sizes

enough for the experimental evaluation we describe. We then calculated the arithmetic average and the standard deviation for both the original and the anonymized trace, which we present here. The NetFlow trace spanned a time period of three minutes and a bit more than 150.000 bytes transferred. Table I presents the values calculated for both traces. We can see that the average and standard deviation do not largely differ. This supports our initial hypothesis, that we can preserve some amount of general information about NetFlows in the trace even after performing random value shifting. Figure 1 presents the cumulative distribution function of the distribution of flow sizes in the original, and the anonymized traces, for comparison and reference. As we can see the distribution remains, almost, identical after the anonymization process. This enforces our initial argument that the information that can be derived by the anonymization process are not affected by the value scrambling.

Trace	Average Flow Size (bytes)	Std. Deviation
Original	1843.69	7336
Anonymized	1845.02	7335.52

TABLE I
SOME BASIC DESCRIPTIVE STATISTICS REGARDING A NETFLOW TRACE BEFORE, AND AFTER ANONYMIZATION.

Currently, *Anontool* performs some basic trace preprocessing when random value shifting is going to be used. It processes all packets in a trace in order to extract the information it needs to calculate d. This information is dependent on the field that random value shifting is being applied on. After the value of d is calculated and chosen according to the user's method of choice, the actual anonymization process takes place. It is possible to estimate d during the anonymization process, however, as knowledge of the whole trace is impossible to have until the whole trace has been processed, we believe the estimated value will not yield as good a result as when a trace can be preprocessed and the value of d calculated on it.

V. CONCLUSIONS AND FUTURE WORK

We described two scenarios where an attacker can, by manipulating the anonymization process, deduce useful information about non-anonymized data in NetFlow traces. Those attacks have already been carried out in packet traces, and we have shown their applicability on NetFlow logs as well. In order to protect against these types of attacks, we have introduced two anonymization primitives and discussed their use and parameterization in order to make educated choices about anonymization policies. We also provided data which suggest their use still preserves useful data about NetFlow logs, without exposing inferential statistics to potential adversaries.

VI. AVAILABILITY

Anontool can be downloaded from http://dcs.ics.forth.gr/Activities/Projects/anontool.html. The application has been installed and tested on RedHat and Debian distributions of the Linux operating system.

REFERENCES

[1] Cisco Systems, Inc, "Netflow Specification." [Online]. Available: http://www.cisco.com/warp/public/732/Tech/nmp/netflow/index.shtml
[2] "Ip flow information export (ipfix)." [Online]. Available: http://www.ietf.org/html.charters/ipfix-charter.html
[3] D. Koukis, S. Antonatos, D. Antoniades, P. Trimintzios, and E. Markatos, "A generic anonymization framework for network traffic," in *Proceedings of the IEEE International Conference on Communications (ICC 2006)*, Jun. 2006.
[4] M. P. Collins and M. K. Reiter, "Finding peer-to-peer file-sharing using coarse network behaviors," in *ESORICS*, 2006, pp. 1–17.
[5] H.-J. Kang, M.-S. Kim, and J. W.-K. Hong, "A method on multimedia service traffic monitoring and analysis," in *DSOM*, 2003, pp. 93–105.
[6] B. Nickless, J.-P. Navarro, and L. Winkler, "Combining cisco netflow exports with relational database technology for usage statistics, intrusion detection, and network forensics." pp. 285–290.
[7] R. Pang, M. Allman, V. Paxson, and J. Lee, "The devil and packet trace anonymization," *ACM Computer Communication Review*, vol. 36, no. 1, pp. 29–38, Jan. 2006.
[8] D. Koukis, S. Antonatos, and K. G. Anagnostakis, "On the privacy risks of publishing anonymized ip network traces," in *Communications and Multimedia Security*, 2006, pp. 22–32.
[9] S. Coull, M. Collins, C. Wright, F. Monrose, and M. Reiter, "On web browsing privacy in anonymized netflows," in *16th USENIX Security Symposium*, 2007, pp. 339–352.