Flocon2010 – January 12th, 2010

Introduction to SIE (condensed)

Eric Ziegast
<info@sie.isc.org>

Security Information Exchange

Raison d'être

Providing common legal and privacy framework for sharing sensitive data

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

Enabling cross-analysis between disparate data sets

Creating network effect between participants (stone soup)

Decentralized - bi-lateral

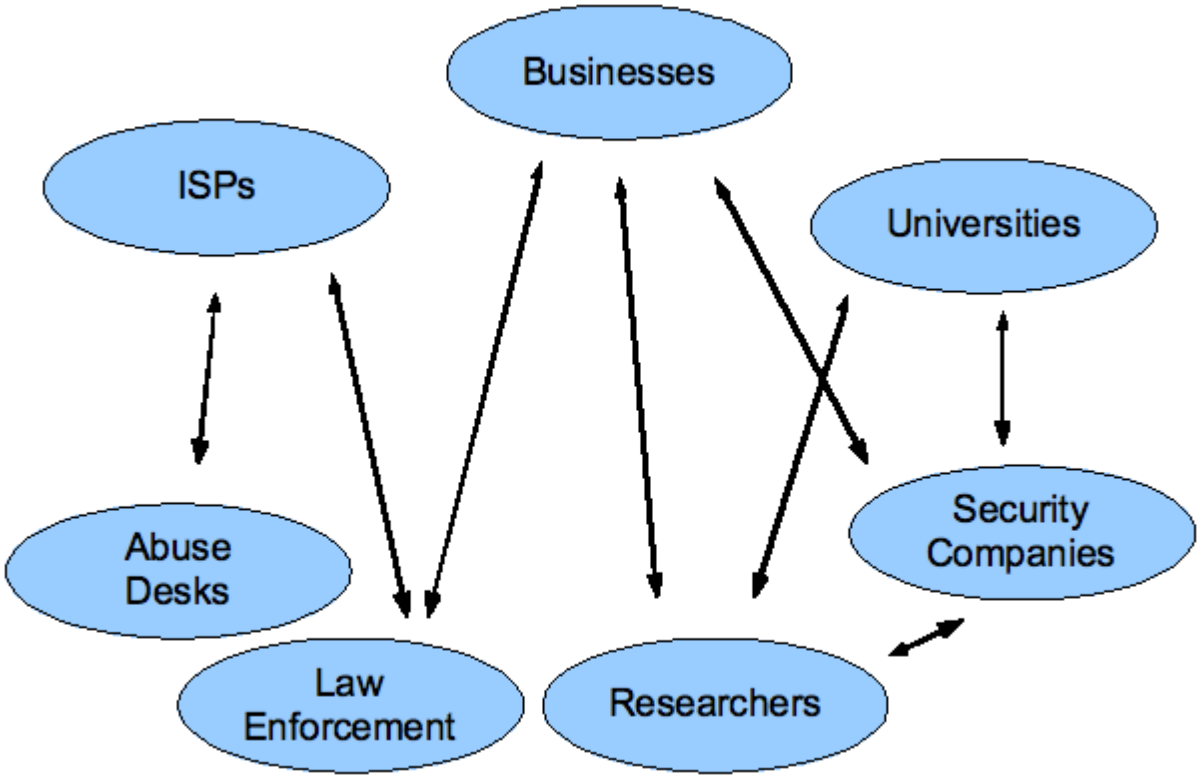

Centralized - multi-lateral

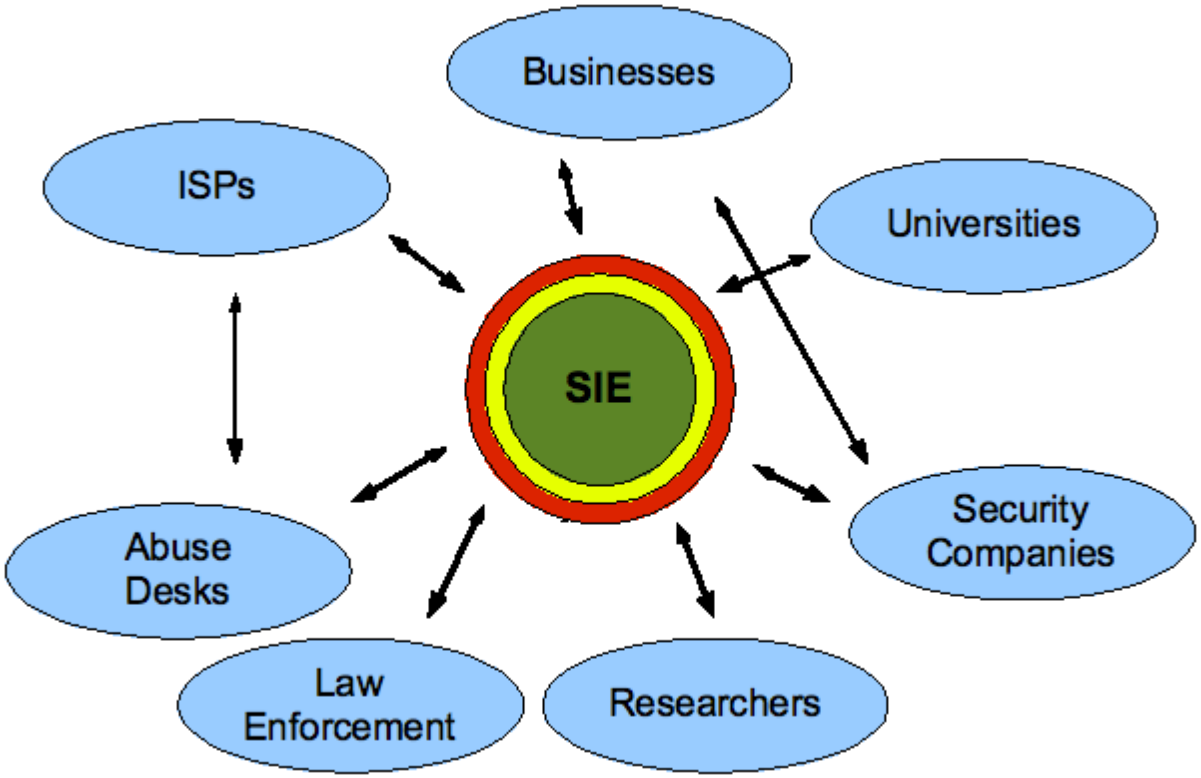

Efficient sharing within common legal/privacy framework

Data distribution model - today

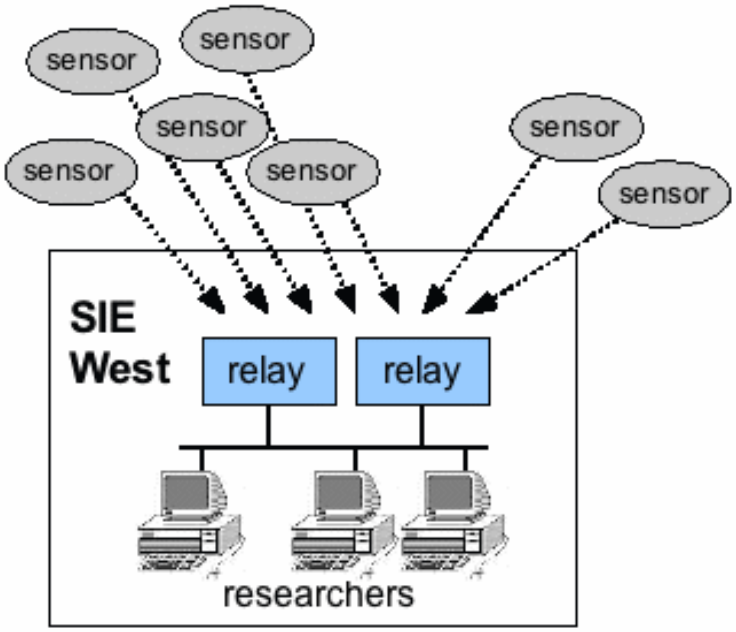

- SF Bay Area, US (PAIX)
- Main sensor relays
- Some researchers getting feeds off switches

Data distribution model - east

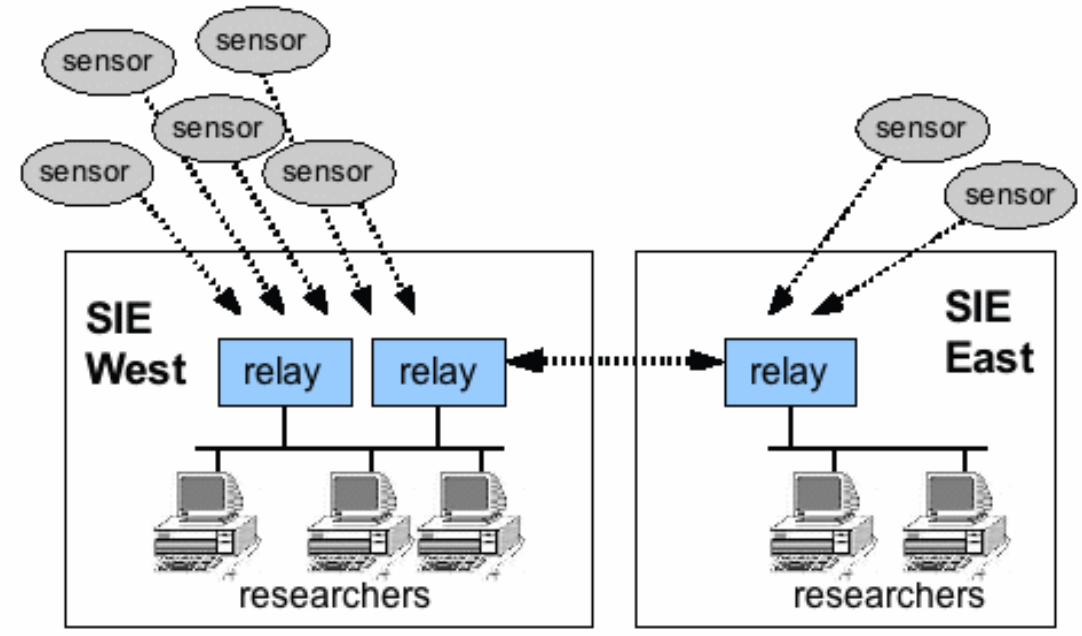

- NYC, US
- Redundant facilities
- More researchers

SRV load balancing

Data distribution model - relay

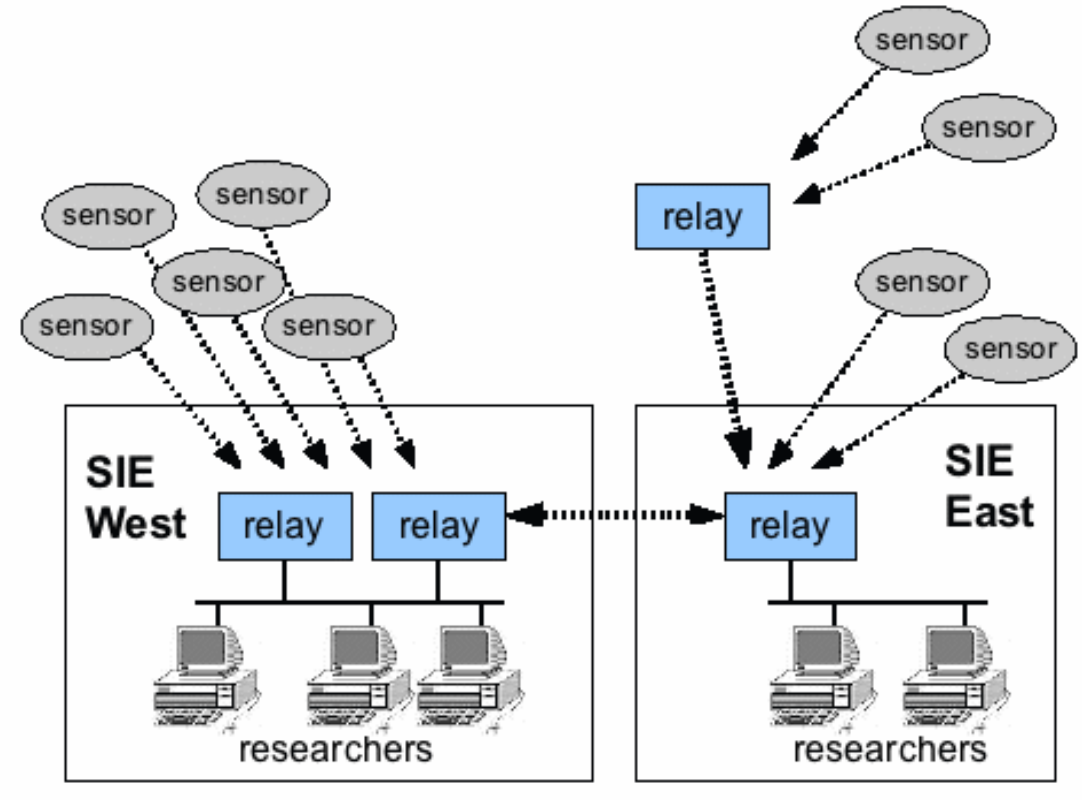

- Add relays at exchanges in different countries
- Add local sensors
- Local sharing or tools possible within relay

Data distribution model - promote

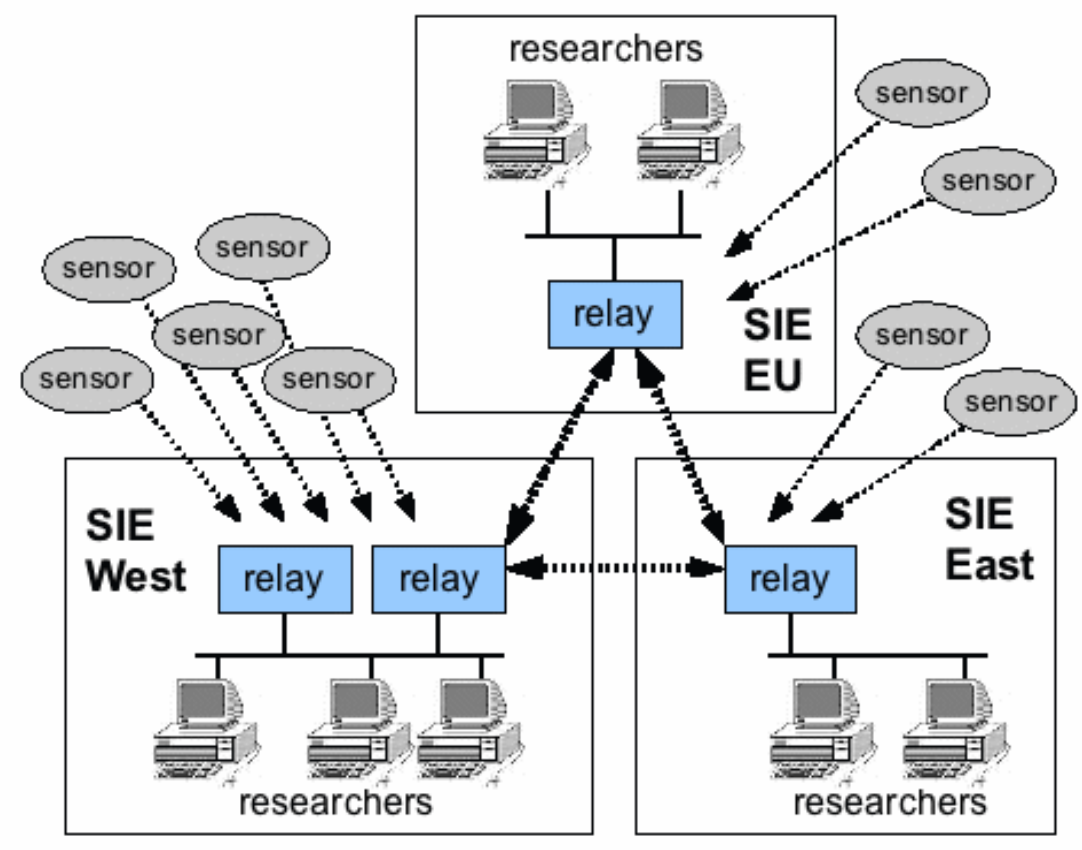

- Promote node when number of researchers is significant
- Scaling issues

Disparate data

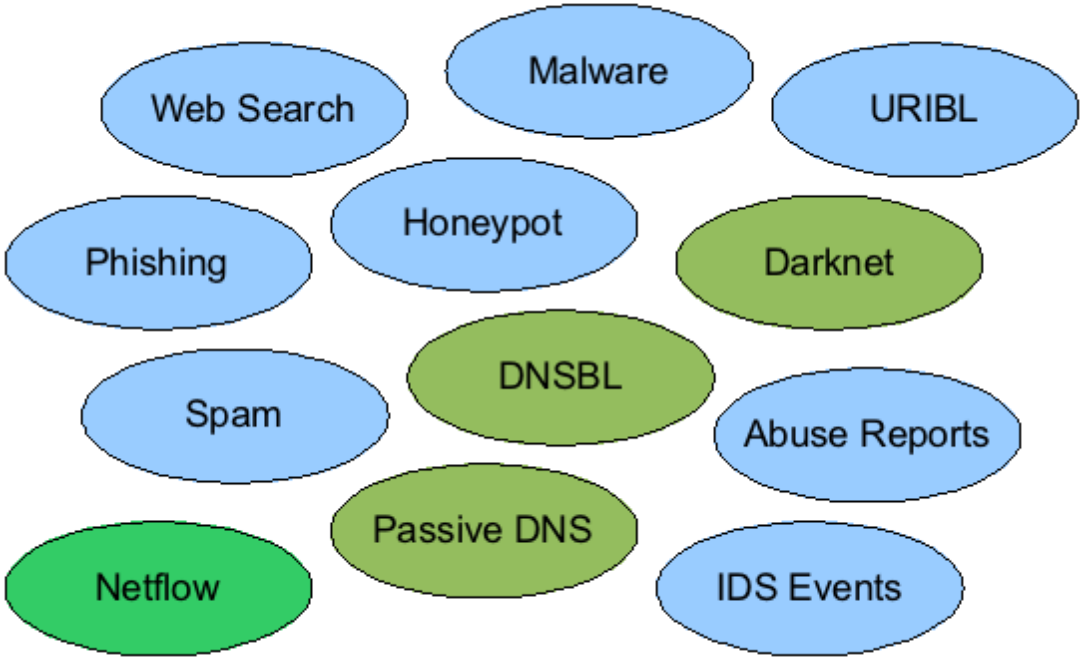

pDNS

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

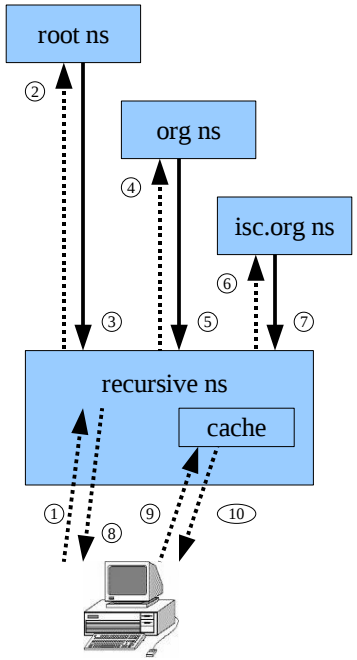

1. client queries for www . isc . org to recursive nameserver
2. recursive nameserver queries root server for “www.isc.org”
3. root nameserver responds with information for “org” nameservers
4. recursive nameserver queries org nameserver for “www.isc.org”
5. org nameserver responds with information for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

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

Fire hose

150 kpps (not-including sub-packets)

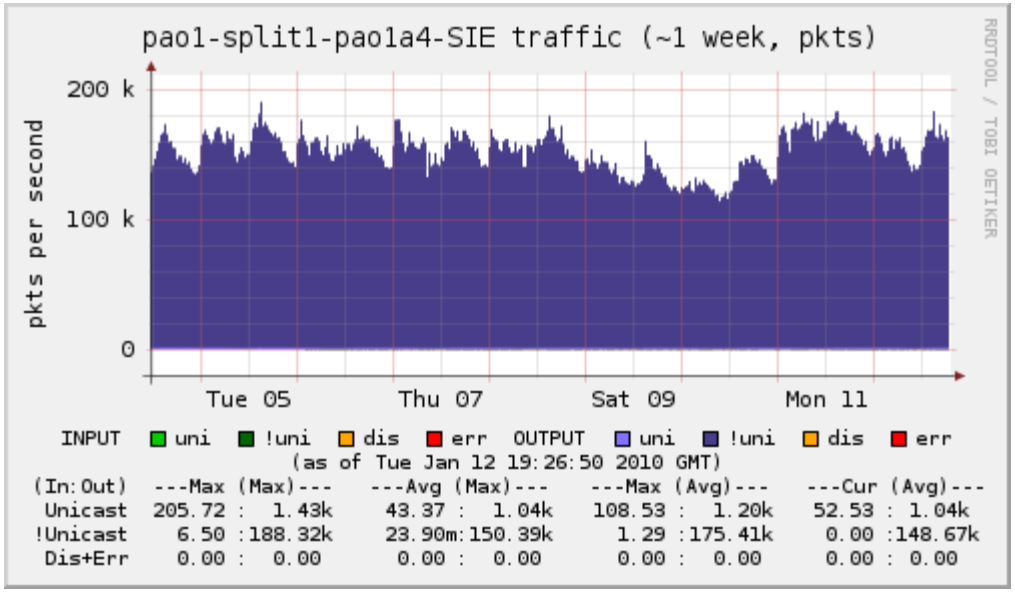

Tools

ncap – used primarily for DNS data

plugins – filter data for rebroadcast

nmsg – used to describe any data context

multi-site VPN – services and lookup tols

hardware – high packet rates

fast switch – line rate GigE, jumbo frames

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

storage – large disk for arrears, SSD

ncap

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

Evolution from pcap/dnscap

Defragmentation

Drop link layer info

Normalized network format

Nanosecond timestamps

User-defined flags

New features key to SIE

I/O – File, BPF, Unicast, broadcast, multicast

Plug-in modules

dedupe, pattern matching, table lookups

Broadcast architecture

Not a database – flow freely

Not a just a packet dump

Real time, not arrears

Loosely coupled multi-processor

Ethernet switch

Partitions via VLANs or “channels”

ncap

plug-in filters in action

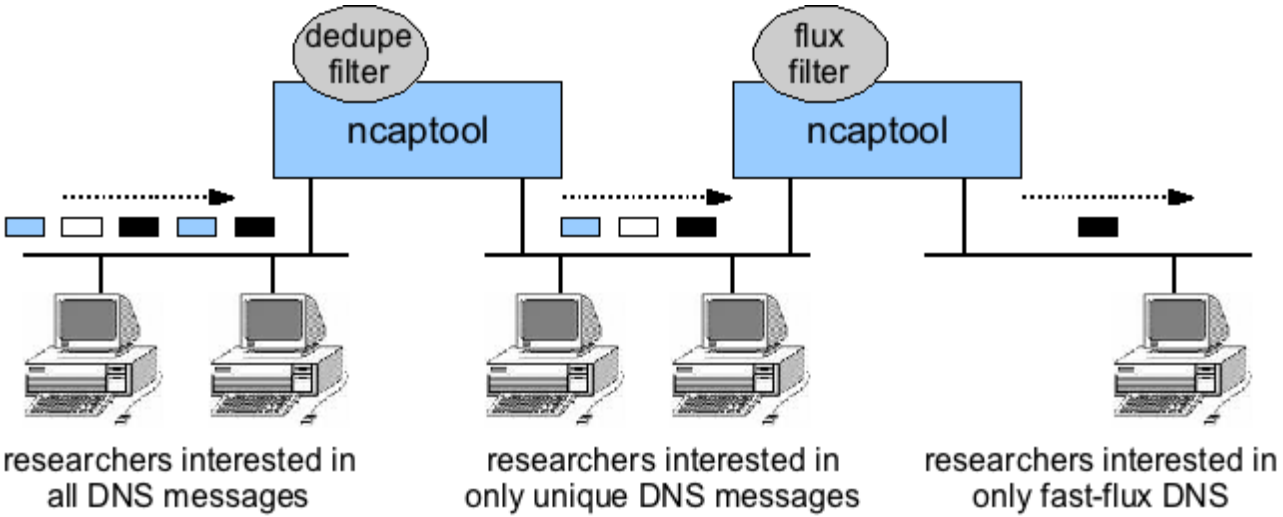

filters – raw passive dns

max at 40000 pps – too much

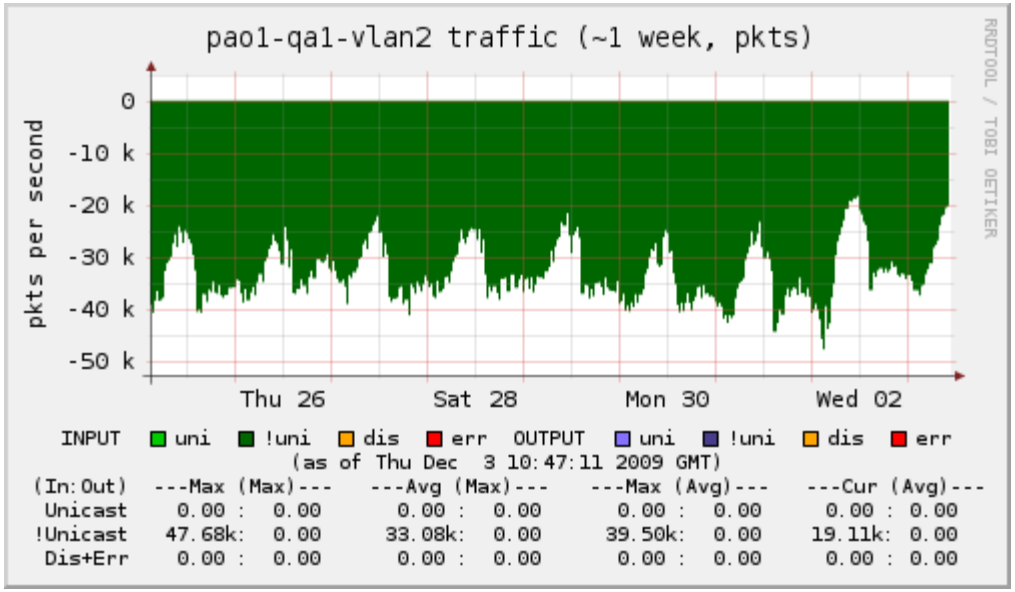

filters – deduplicated dns

max at 5000 pps - better

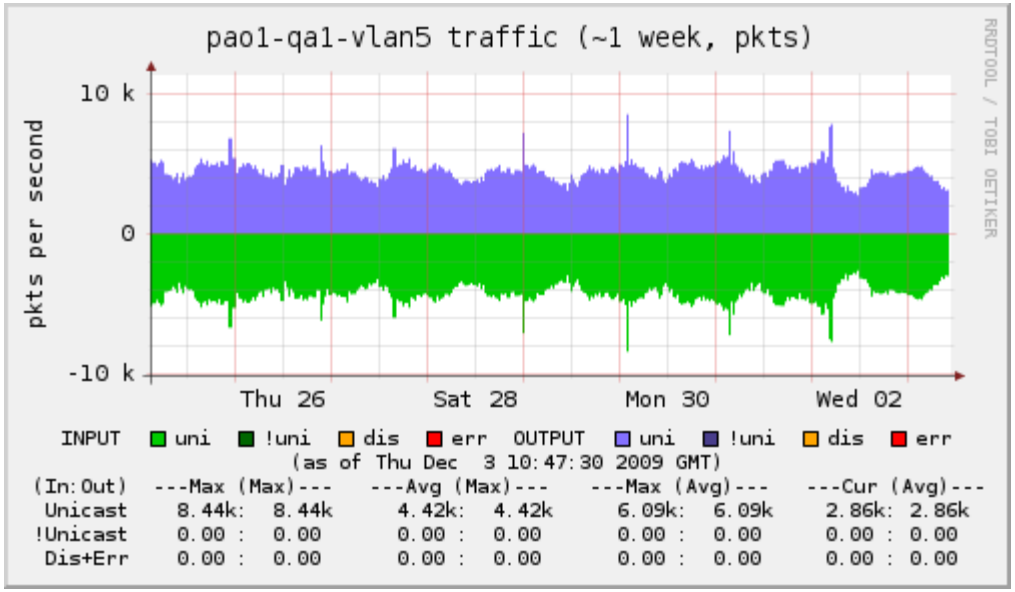

filters – fast-flux dns

rate is 2-4 pps

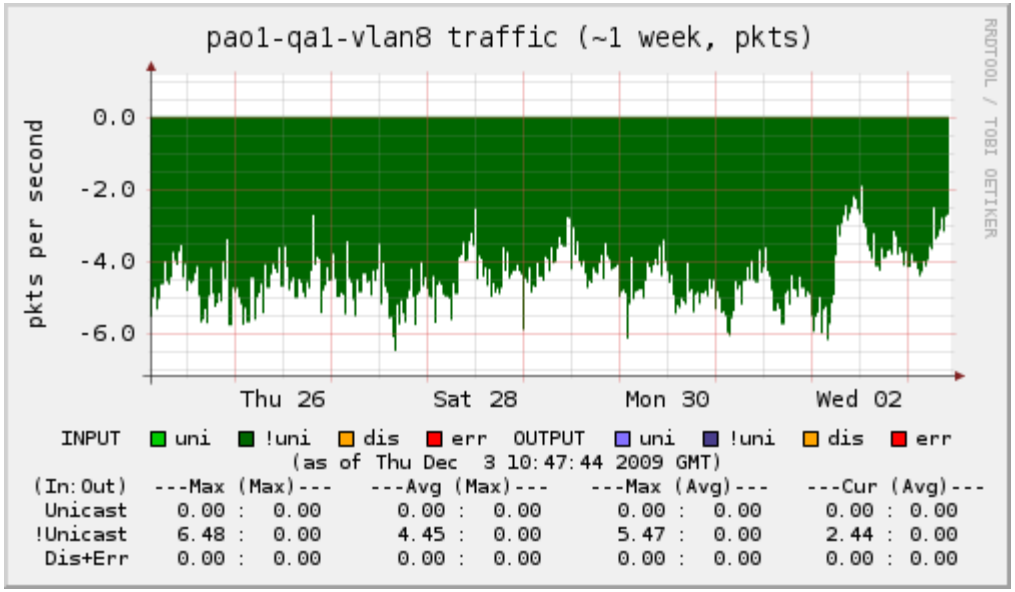

nmsg

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

Needs:

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

Channels (nmsg, et.al.)

Spam – yes, we want spam

URL Link Pairs – search engines

Conficker

sinkhole URL, DNS, P2P

https://conficker.sie.isc.org

Mitigation efforts – other botnets

Upcoming

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

Malware – not just hash

Channel example - pDNS

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

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

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

dns QUERY,NOERROR,3324,qr

1 radio.wareznet.net,IN,A

0

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

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

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

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

.,CLASS512,TYPE41,32768,[0]

Channel example – web sinkhole

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

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

type: sinkhole

srcip: YYY.YY.YYY.YY

srcport: 64707

dstip: 149.XX.XX.XX

dstport: 80

request:

GET /search?q=0 HTTP/1.0

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

Host: 149.XX.XX.XX

Pragma: no-cache

p0f data (kudos to Chris Lee)

Channel example - spam

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

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

type: spamtrap

srcip: 189.15.60.161

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

from: REDACTED@spamtrapdomain.net

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

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

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

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

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

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

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

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

Expand with pDNS

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

faa76.gupicfd.cn

fb6886169.yepekfy.cn

...etc...

ns3.j8w.ru

www.vapagnj.cn

xidisqs.cn

xivivxt.cn

yepekfy.cn

zuvidtn.cn

zuwohxc.cn

zuyimqg.cn

zuzewnp.cn

zuzovgw.cn

Combining data

Jose Nazario / Thorsten Holz - Malware08

- Heuristics, Point system, Fast-flux

David Dagon / Wenke Lee

- pDNS + string matching -> FakeAV

Richard Clayton – UKNOF13

- pDNS hosts + active scans => blocking policies

Andrew Fried – Blackhat DC 2010

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

Ed Stoner – Flocon 2010

- pDNS/ncap + netflow/silk

How organizations can help

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

Bring servers to SIE and create value-added services

lots of data yet to be analyzed

get familiar with tools

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

pDNS, spam, netflow, darknet blocks, etc.

Questions?

Email: info@sie.isc.org

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

Eric Ziegast – SIE Programme Manager

+1.650.423.1363 (PST8PDT)

Nmsg:

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

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

Ncap:

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