RIPE 58 – European Internet Exchange Working Group

Introduction to SIE and 2009 Update

Keith Mitchell - ISC

<keith_mitchell@isc.org>

Exchange concepts

- Operator is a neutral trusted intermediary
- Participants are often direct competitors
- Examples
  - Internet traffic (PAIX, Equinix, LINX)
  - Equity/futures (NY/London SE, NASDAQ)
  - Telco/Meet Me Room (Telehouse, CRG West)

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

Decentralized - bi-lateral

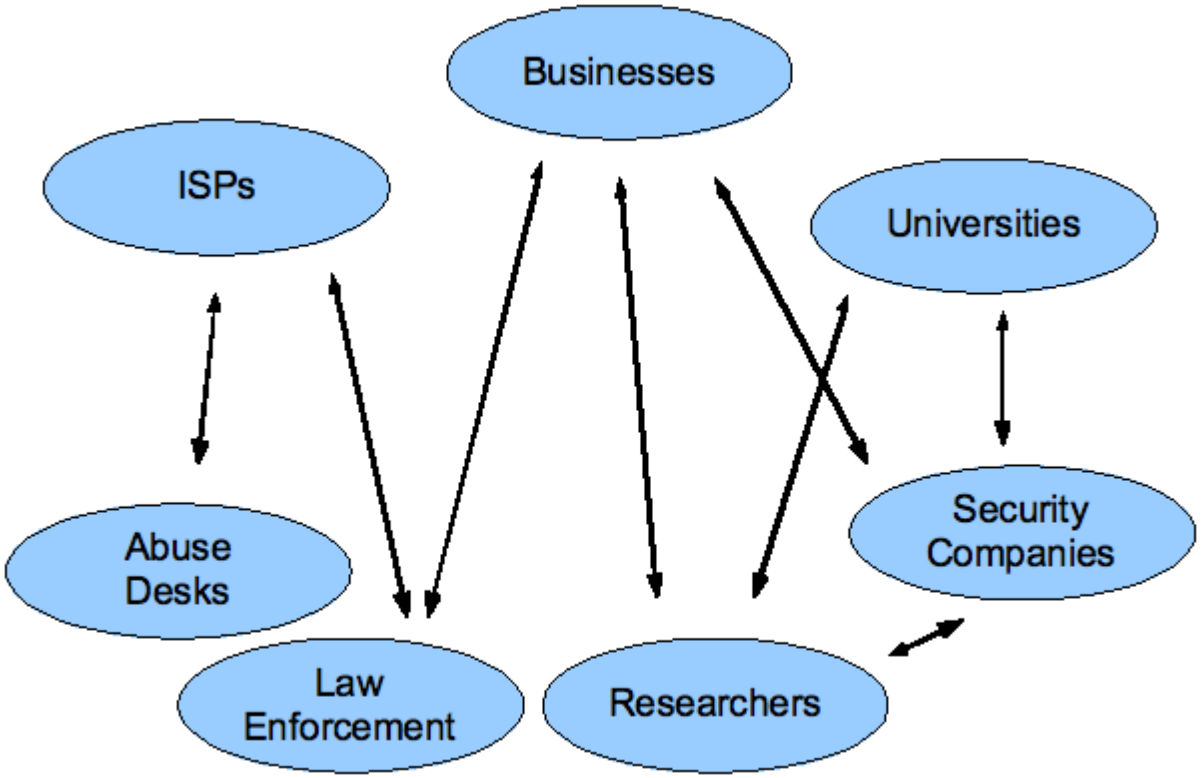

Centralized - multi-lateral

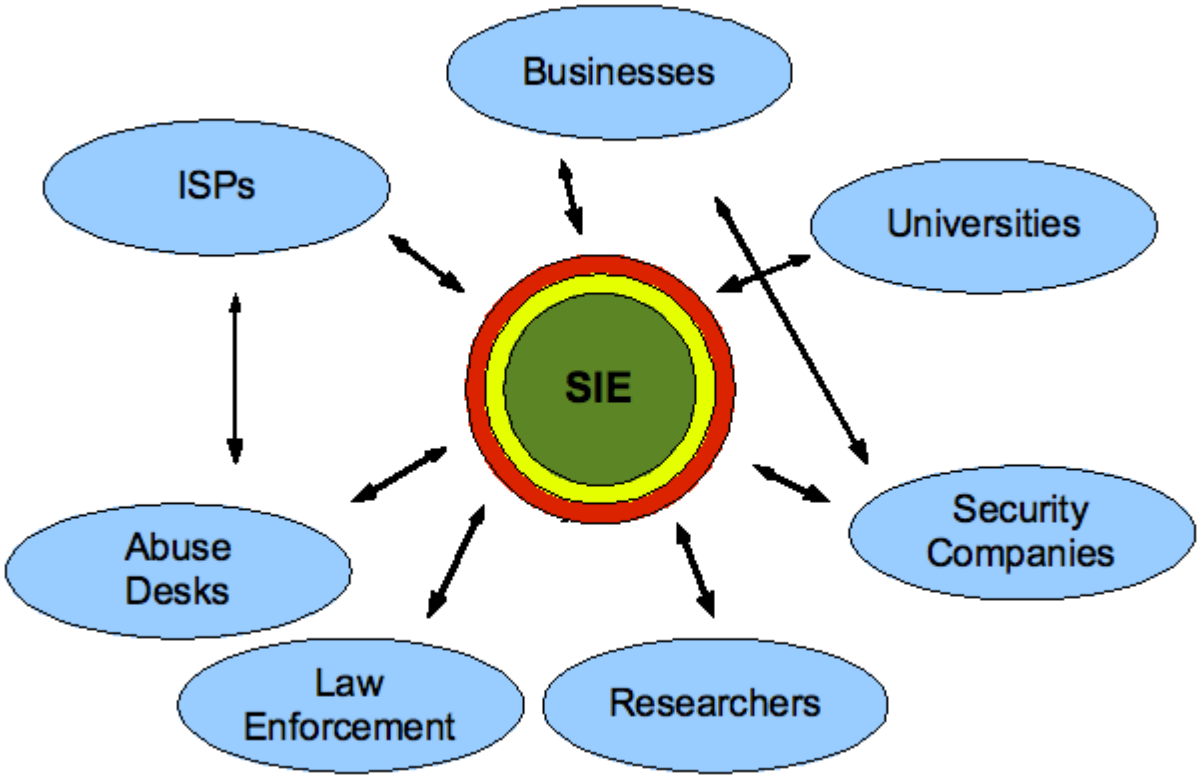

Efficient sharing within common legal/privacy framework

Data distribution model - today

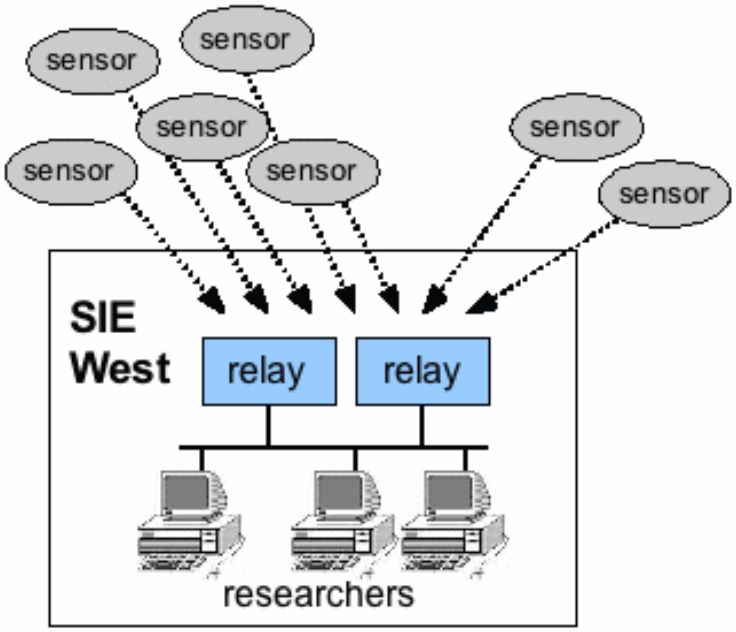

- SF Bay Area, US (PAIX)
- Main sensor relays
- Some researchers getting feeds off switches

Data distribution model - east

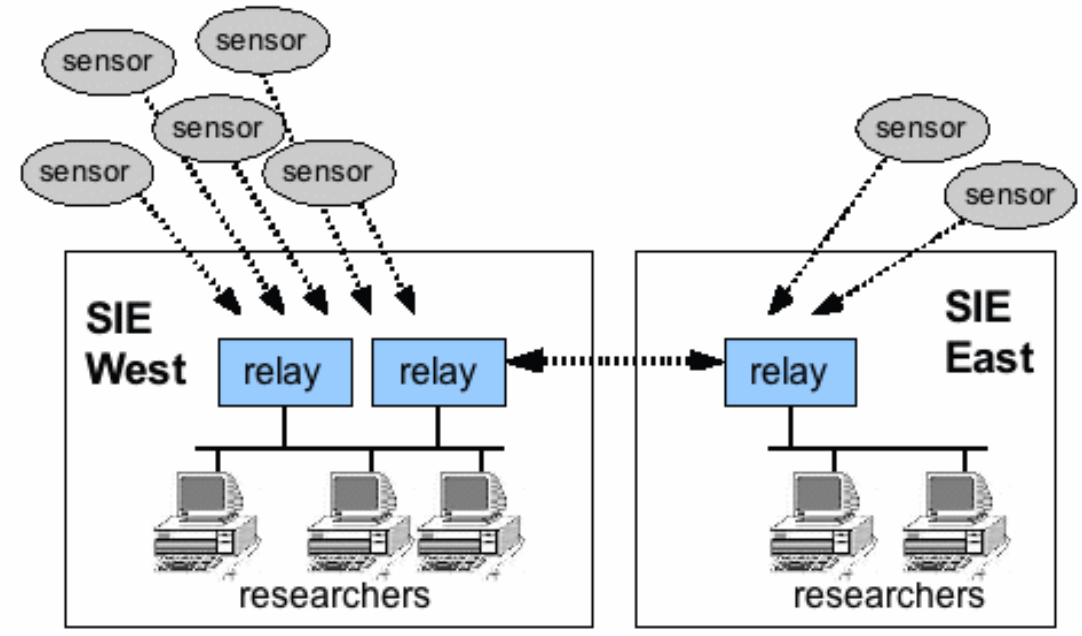

- DC or NY, US
- Redundant facilities
- More researchers

Data distribution model - relay

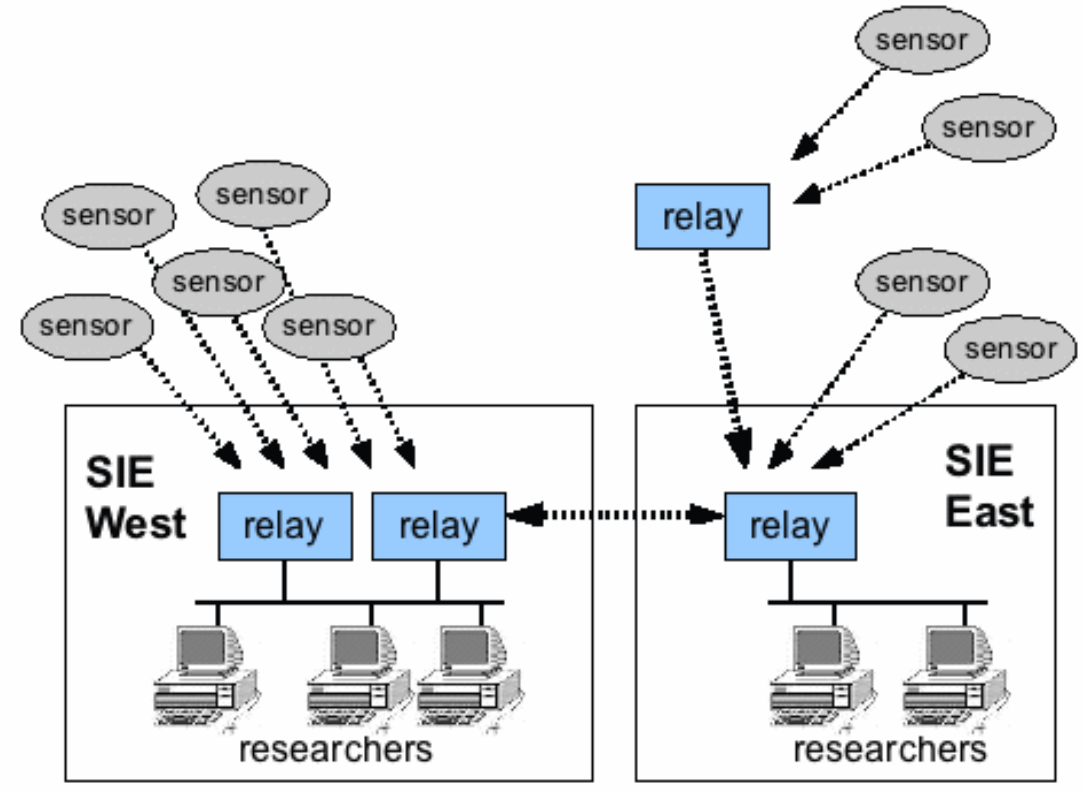

- Add relays at exchanges in different countries
- Add local sensors
- Local sharing or tools possible within relay

Data distribution model - promote

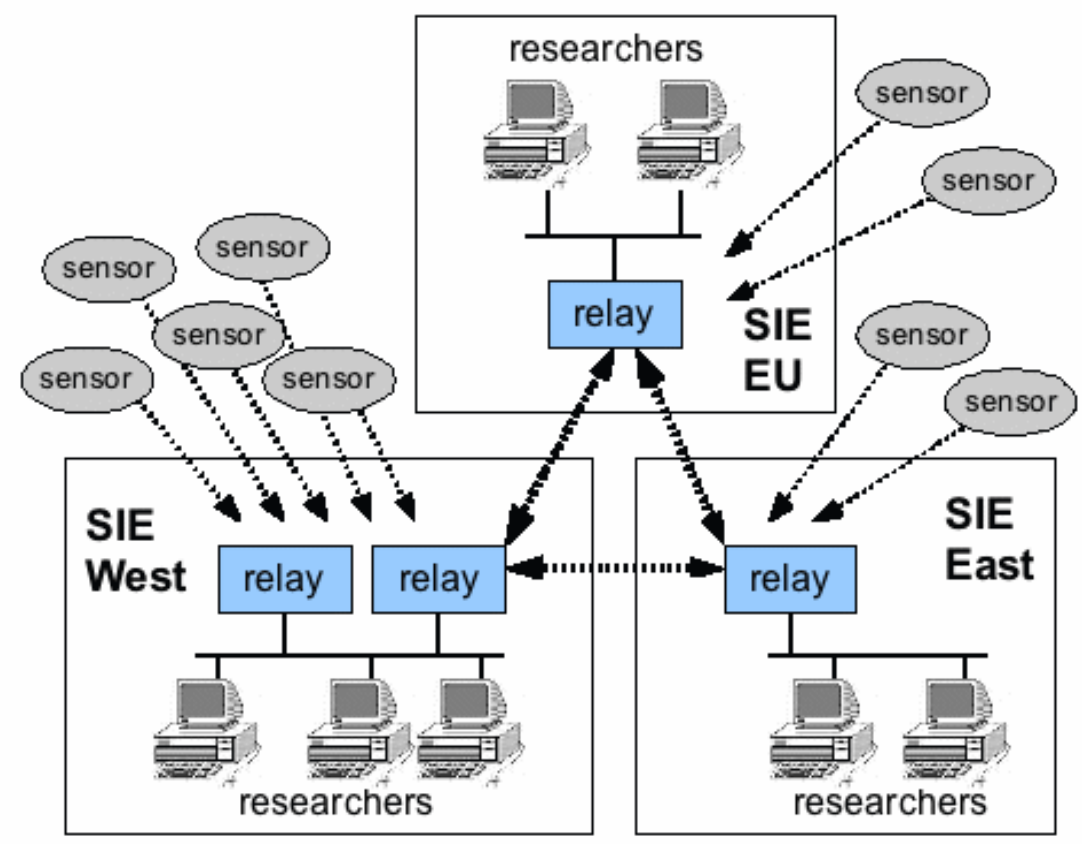

- Promote node when number of researchers is significant
- Scaling issues

Disparate data

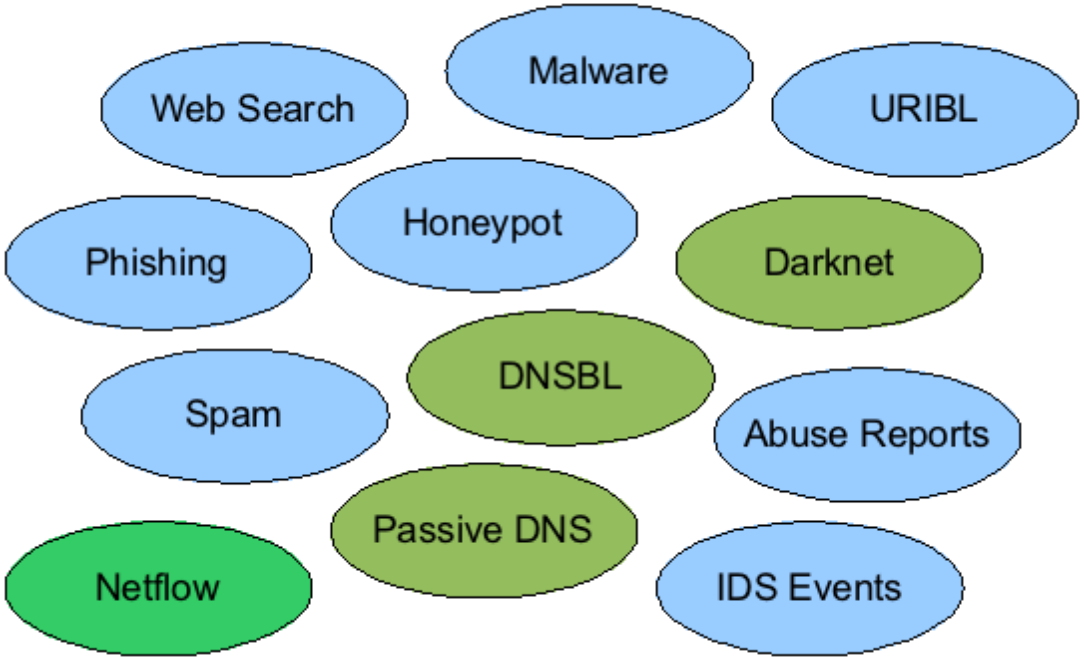

Tools

- ncap – used primarily for DNS data
- plugins – filter data for rebroadcast
- nmsg – used to describe any data context
- hardware – high packet rates
  - fast switch – line rate GigE, jumbo frames, no packet loss
  - servers – Linux/FreeBSD, 64-bit, multi core, 1333+MHz FSB, DDR2-667+ memory
  - storage – large disk for arrears, SSD

ncap

- ftp://ftp.isc.org/isc/ncap
- Enhancements to 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, database lookups

ncap

- plug-in filters in action

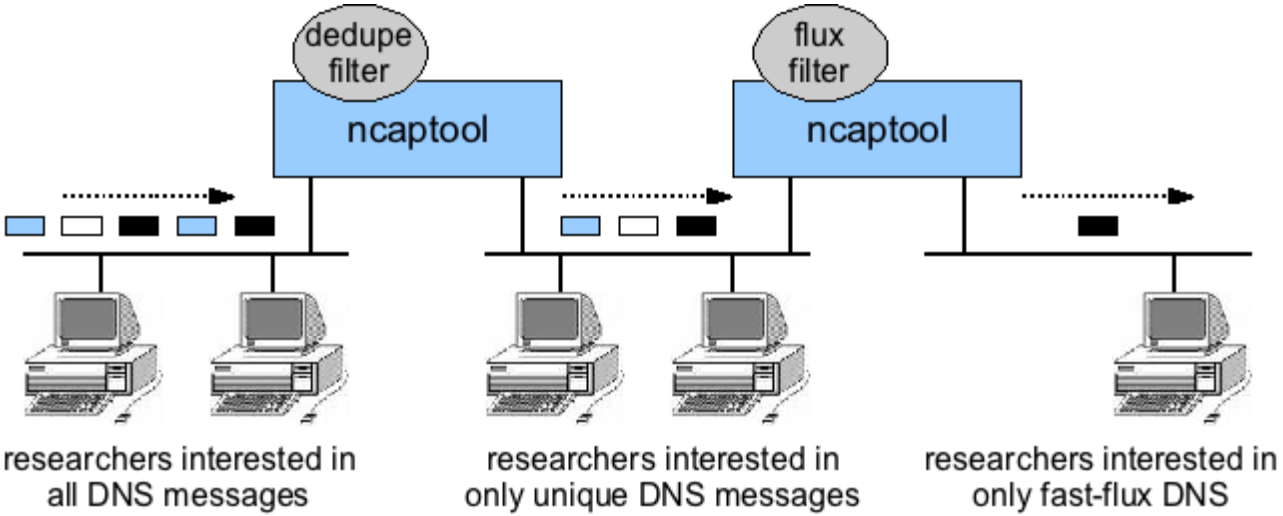

filters – raw passive dns

- max at 20000 pps – too much

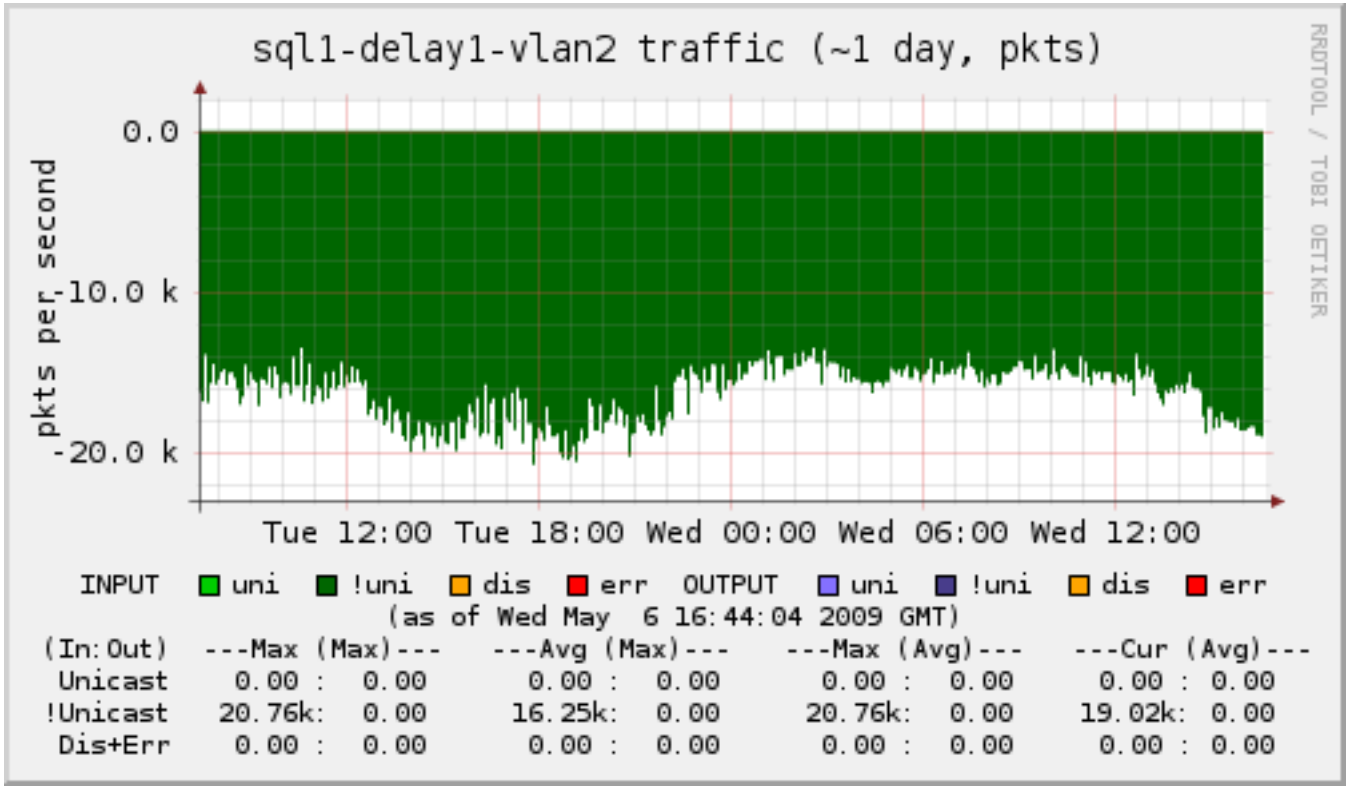

filters – deduplicated dns

- max at 3000 pps - better

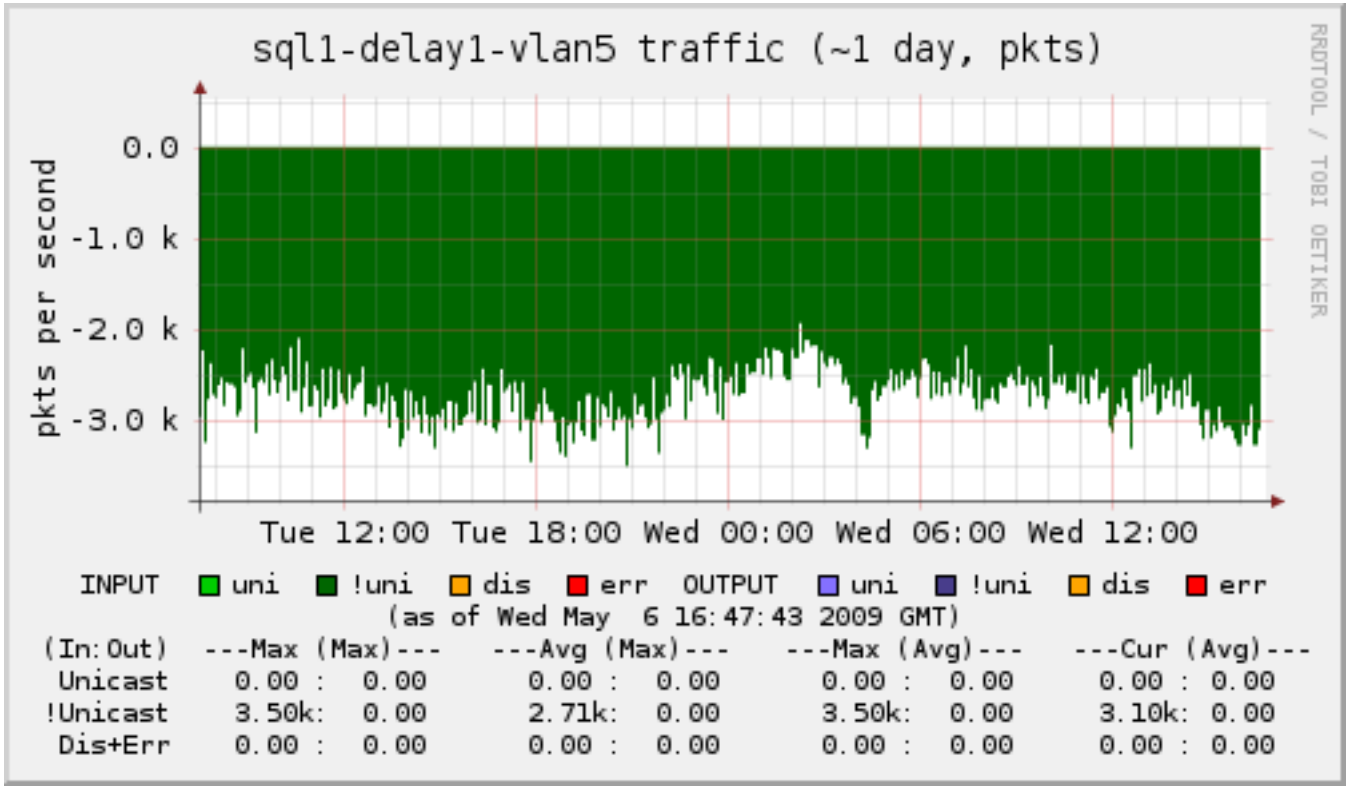

filters – fast-flux dns

- rate is <1 pps

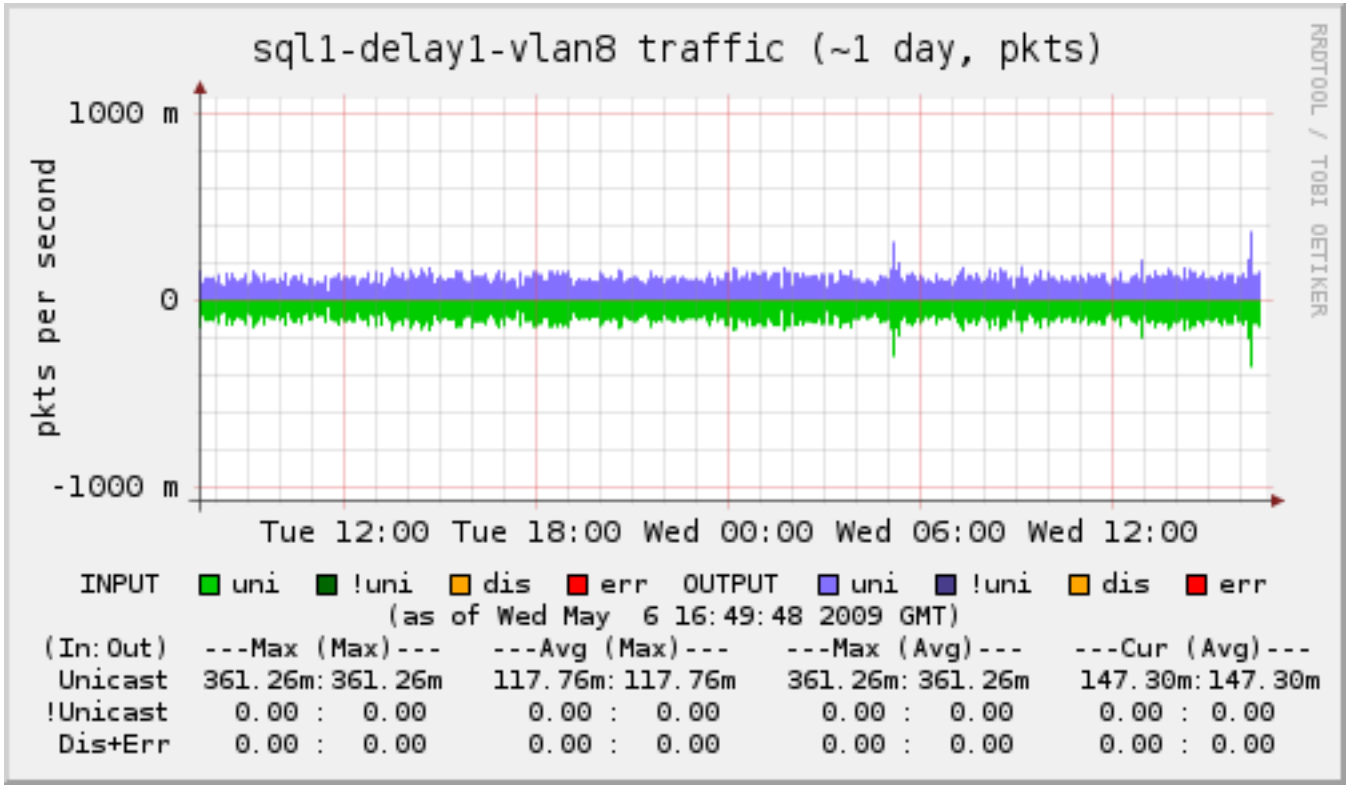

nmsg

- ftp://ftp.isc.org/isc/nmsg
- Any structured data
- Format:

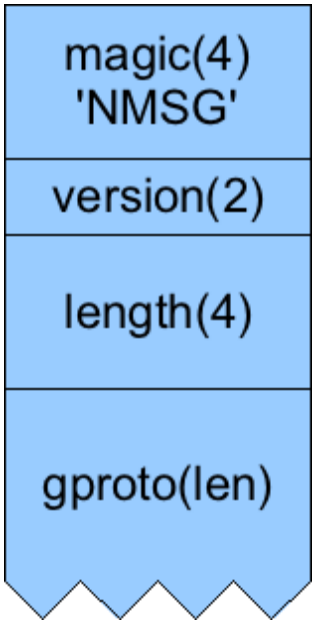

Google protocol buffers

- http://code.google.com/apis/protocolbuffers
- APIs for C, C++, Python, Java, Perl
- Arguably better than XML
  - 3x-10x smaller, 20x-100x faster
  - simpler, easy to read/filter raw off wire
- Open source
- Extensible:
  - ISC-defined (vendor id=1 for ISC)
  - user-defined (vendor id=__)

nmsg.proto

```
package nmsg;

message Nmsg {
    repeated NmsgPayload    payloads = 1;
}

message NmsgFragment {
    required uint32         id = 1;
    required uint32         current = 2;
    required uint32         last = 3;
    required bytes          fragment = 4;
}
```

```
message NmsgPayload {
    required int32          vid = 1;
    required int32          msgtype = 2;
    required int64          time_sec = 3;
    required fixed32        time_nsec = 4;
    optional bytes          payload = 5;
    repeated uint32         user = 6;
}
```

Channels (ncap)

- Raw Passive DNS
  - Deduplicated Passive DNS
    - Fast-flux DNS
    - IP blacklist DNS – CBL, SORBS
- DNS Queries
  - RBL / DNSBL
  - Dynamic DNS providers
  - Popular sites / some TLDs
  - AS112 RFC1918 PTR lookups
  - Legacy root queries

Channels (nmsg, et.al.)

- Spam – send us more!
- URL Link Pairs – search engines
- Netflow
- Conficker
  - sinkhole URL, DNS, P2P
  - https://conficker.sie.isc.org
- Coming next
  - malware
  - darknet

Channel example - spam

- Preprocessing scripts (python) packetizes into {envelope, headers, URLs}
- Spam types:
  - spam traps
  - "this is spam" reports/submissions
  - spamassassin-scored email
- Good starting point for analysis
  - Malware, phishing, botnets

isc/email.proto

```
package nmsg.isc;

enum EmailType {
    unknown = 0;
    spamtrap = 1;       // email sent to a spamtrap
    rej_network = 2;    // rejected by network or SMTP (pre-DATA) checks
    rej_content = 3;    // rejected by content filter (including domain blacklists)
    rej_user = 4;       // classified by user as spam
}

message Email {
    optional EmailType  type = 8;
    optional bytes      headers = 2;    // SMTP headers
    optional bytes      srcip = 3;      // remote client IP
    optional bytes      srchost = 4;    // remote client PTR, if known
    optional bytes      helo = 5;       // HELO/EHLO parameter
    optional bytes      from = 6;       // MAIL FROM parameter (brackets stripped)
    repeated bytes      rcpt = 7;       // RCPT TO parameter(s) (brackets stripped)
    repeated bytes      bodyurl = 9;    // URL(s) found in decoded body
}
```

Channel in progress – malware

- Malware submissions
  - L1: Upload, provide internal URL link, hashes
  - L2: automated unpackaging and sandboxing tools
  - L3: behavioral descriptions
- Value: Data sharing vs. Analysis sharing

Future examples

- Domain/IP/AS reputation and other spam/malware/DNS/URL analysis byproducts
- DNS scans
  - ISC hosts some scans – maintains do-not-scan list
  - Real-time ISC domain survey?
- Automated abuse or DDoS reporting
- Flash-mob URLs
- BGP updates/changes/flapping

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
  - DNS, spam, netflow, darknet, etc.

Exchange operators

- Organize within your sphere of influence – ISPs, Hosting, local CERTs, law enforcement
- Actively manage security relationships with customers and ISPs (beyond abuse@)
- Get involved with local/national security organizations
- Investigate privacy laws and policies

Questions?

- Email: info@sie.isc.org
- Web: https://sie.isc.org/
- Keith Mitchell – here at RIPE
- Eric Ziegast - +1.650.423.1363 (GMT-0700)