Implications of SIE

Paul Vixie, ISC

ISC/CAIDA Workshop @M^{3}AAWG

October 2012

SIE Characteristics

- History: conceived in 2007, piloted in 2008 (NCAP), formally launched in ~2009 (NMSG)
- General purpose, scalable, distributed data collection; shared real-time analysis
- Multiple channels, multiple schemas
- Channels: passive DNS, honeypot results, spamtrap results, network telescope packets
- Growth (traffic, sensors, data types) continues

Known SIE Applications

- First blood goes to Andrew Fried, who cross-correlated passive DNS and spam trap results
- ISC DNS channel package (dedup, filtering)
- ISC DNSDB is a general purpose passive DNS database – ~2 years so far, fully indexed
- Some security companies are feeding SIE data into their pre-existing analysis systems
- Many SIE credits in published research

An Impedence Mismatch?

- Noting:
  - Science requires objective repeatability
  - SIE is like Heraclitus’ river: never the same
- Open questions:
  - Does academic rigor require known data sets?
  - Would that hold back innovation of real-time analysis methods?

ISC Passive DNS Channel Package

- Raw sensor data is mightily self-similar
  - ISC’s dedup processes reduce by ~15:1
- Lots of chaff among the wheat
  - ISC’s filter separates the RBL, PTR, netflix, and DNS tunnel traffic into a “DNS chaff channel”
  - The remaining wheat gets its own channel
- Note: NXDOMAIN, REFUSED, FORMERR, other errors, are only present in the raw data
  - ISC will shortly add a “DNS Error channel”

ISC’s Attitude Toward SIE

- The SIE port and all raw sensor data donated to ISC is a service of ISC’s non-profit parent company
    - This is our deal with our sensor operators
- Private raw sensor data (“proprietary spam”?) is available by negotiation with channel manager
    - SIE is a convenient and trusted place to interchange
- ISC’s commercial value-added activities include:
    - “DNS channel package” (dedup, chaff filter, [errors])
    - “DNSDB” (passive DNS database: full indexing forever)

ISC's Corporate Structure

- ISC's heart and soul is a non-profit tax-exempt public charity – "the parent/holding company"
  - Operates Security Information Exchange (ISC SIE)
- ISC deploys commercial subsidiaries for non-charitable activities – for fund-raising and to ensure relevance in the I.T. market
  - Value added security products (like DNSDB)
  - Also BIND/DHCP support, training, consulting, software enhancement; open source routing; etc.

Some Service Ideas

- Anyone with an SIE port can build services
  - For themselves, or as a [commercial] service
- ISC has a commercial subsidiary which is now in the process of building these examples
  - Commercial, so, not open-source
  - As with DNS channel package and DNSDB, price is reduced for sensor operators and poor non-profits
- These are straightforward applications of SIE, presented here to stimulate some discussion

Service Idea #1: Real Time Monitoring

- A network owner ("the customer") registers:
  - Their global identifiers (IP addresses, DNS names)
  - Notification preferences (RSS, e-mail, SNMP trap)
- The real time monitoring system watches for:
  - Spam or darknet from, or new passive DNS results in, customer IP address blocks
  - Spam mentioning, or passive DNS results about, customer domain or subdomain names

Service Idea #2: DNS Poisoning Detection

- Customer registers:
  - The names of their DNS primary zones
  - Notification preferences (RSS, e-mail, SNMP trap)
- Service operation:
  - Run a stealth DNS slave for all customer zones
  - Passive DNS results using any customer domain name are resolved in parallel inside stealth slave
  - Trigger if observed response is “wrong”

Discussion

- We need new ideas in at least these areas:
  - Channels, sensors, data types
  - Gateways, translators, tools
  - [Commercial] services
  - SIE-enabled research

- Discuss!