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.”

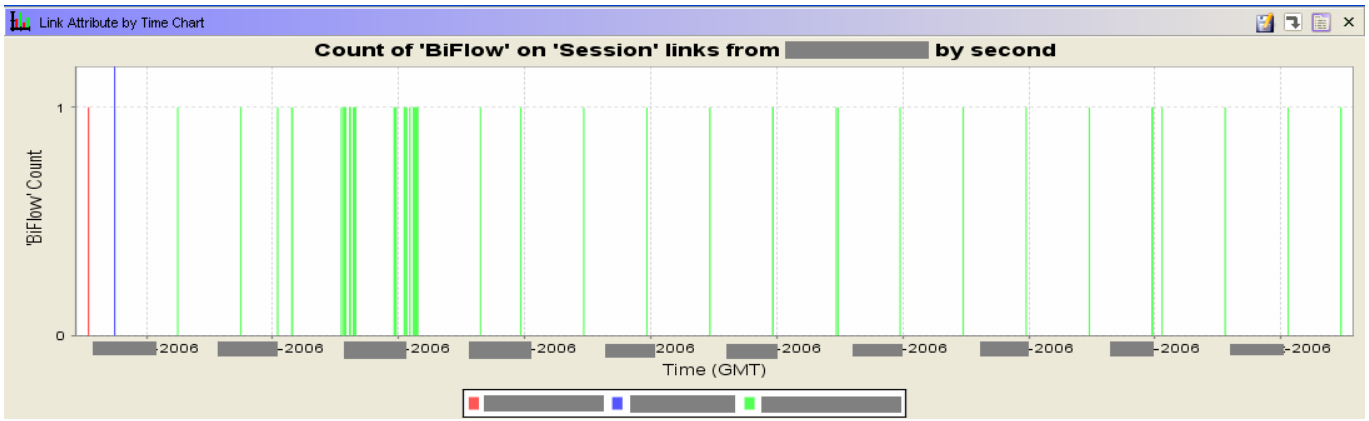

Network Beaconing Detection Strategy

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

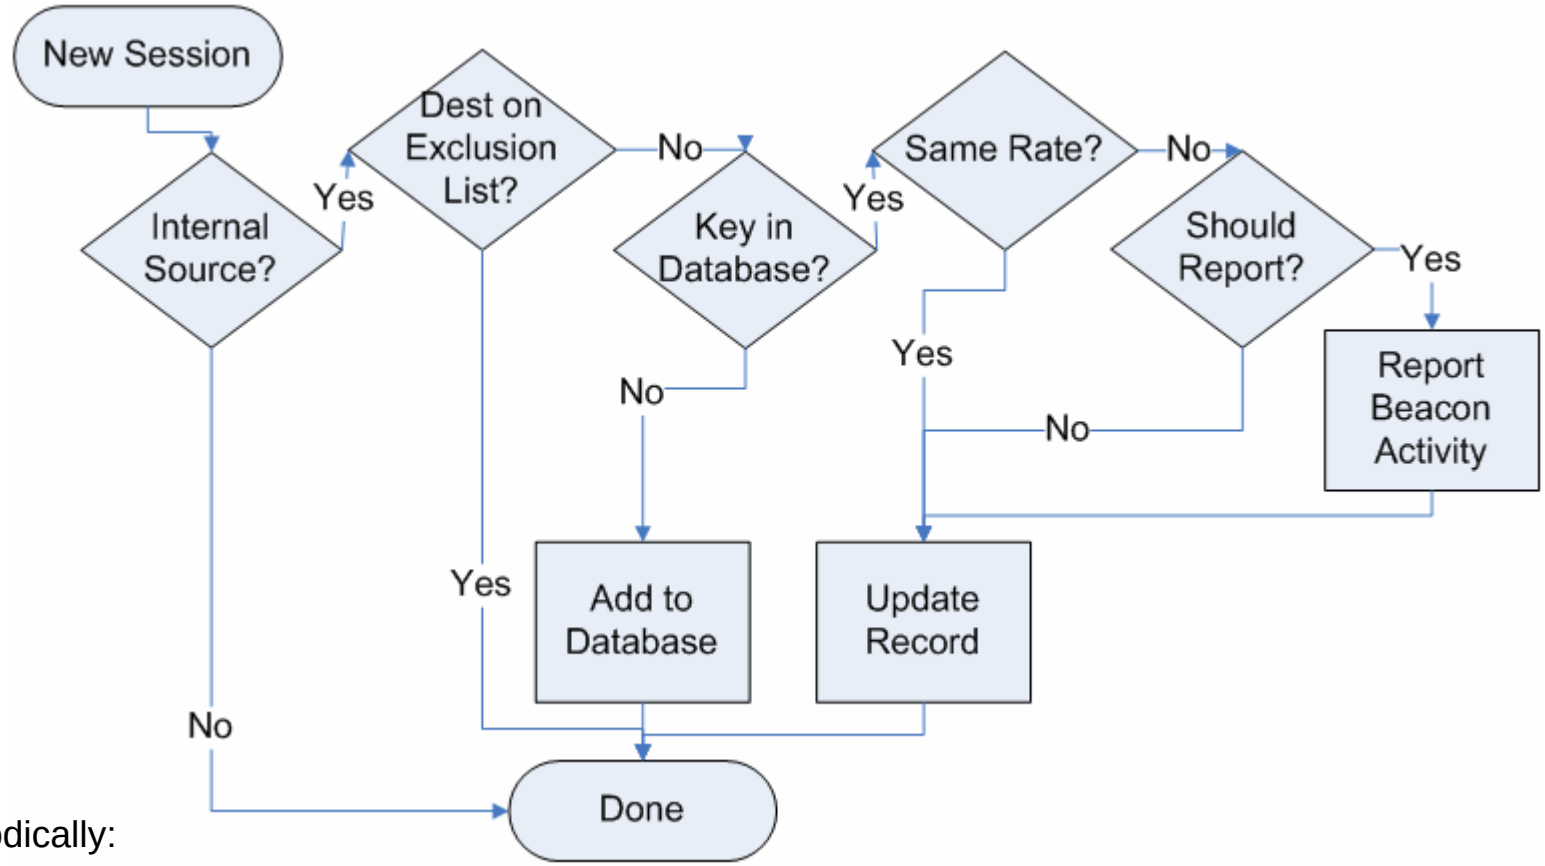

Periodically:

- 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

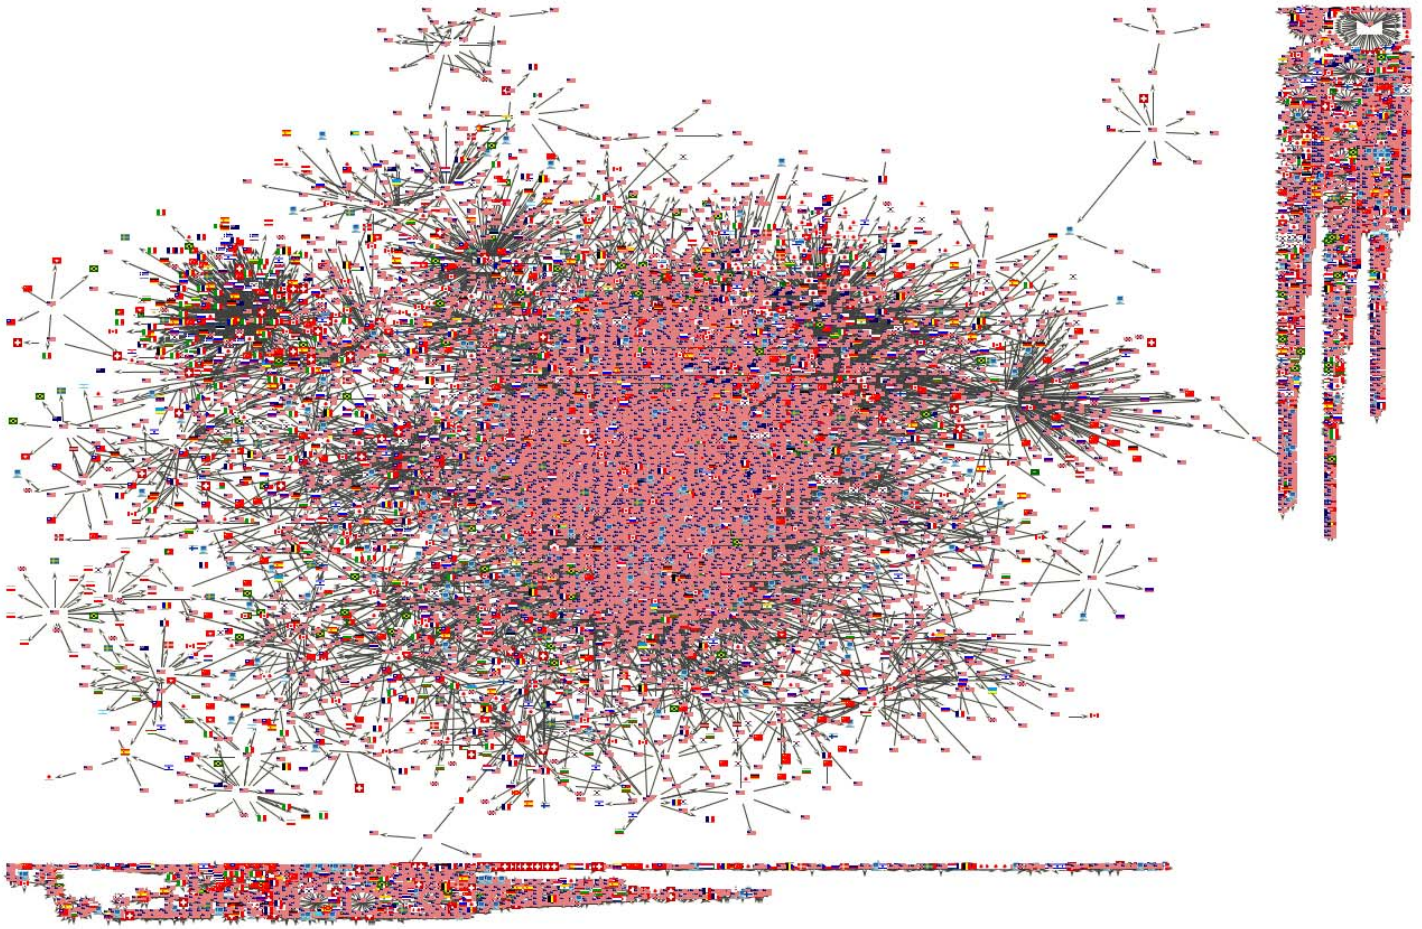

Beacons identified one day of November, 2007

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.

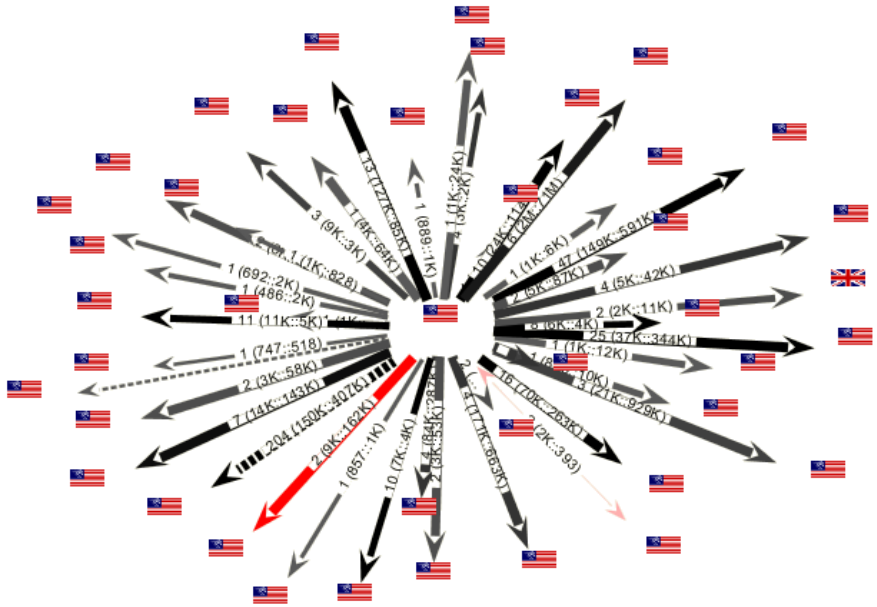

Graph view of Netflow (black), intrusion detection (red), and beaconing (dotted) records from a host.

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.

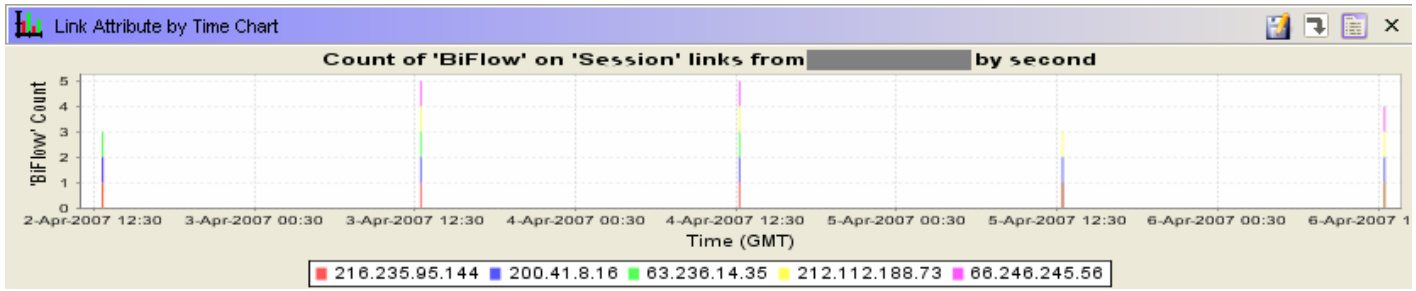

Interesting beaconing to 5 hosts worldwide. Later explained by a popular media player refreshing ads.

Incident Detection : What's Going On?

Two Hosts beaconing to 262 hosts (TCP 2170) over several hours with large response bytes.
[globus]

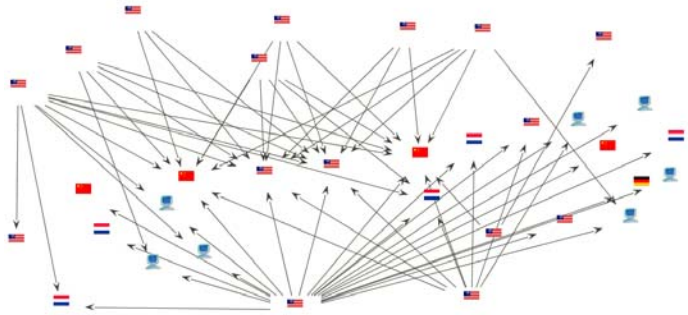

Several hosts beaconing to multiple destinations on TCP and UDP; some beacons never respond
[peer to peer download manager]

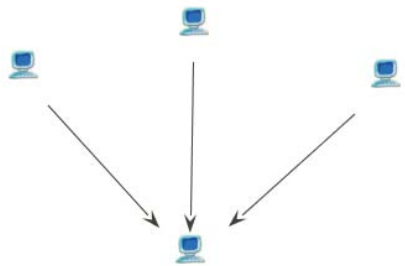

Three Hosts beaconing to a host (TCP 80) every 3 hours.
[i********** spyware phoning home]

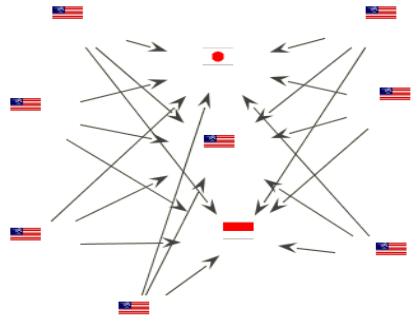

Seven Hosts beaconing to 3 hosts (TCP 30000) over several hours with no response.
["canadapost" shipping module ???]

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