Identifying Compromised Hosts by Analyzing Real-Time Blacklists

NANOG 35
October 2005, Los Angeles
Rick Wesson
Alice's Registry, Inc.

Problem Statement

- Without actively port-scanning your network, it's difficult to tell when malicious or compromised hosts appear and begin causing problems for the rest of the world.

Background

- DNS-based Real-Time Blacklists contain information about potentially compromised systems, including spam senders, open SOCKS and HTTP proxies, botnet members, and open SMTP relays.
- We're using nine of the largest DNS RBLs as data sources.
- Combined with a real-time feed of a BGP transit routing table.

Methodology

- Real-time BGP feed allows a baseline pairing of addresses to the AS numbers which **normally** advertise them, and allows one to see **anomalous** advertisements.
- The maintainer for the normally-originating AS is assumed to be the responsible party for the addresses.

Methodology

- Hourly import of DNS RBL data into a MySQL database
- Each record is tagged with a timestamp, source RBL, type of issue (open proxy, spam sender, vulnerable web server, et cetera), and the ASN currently originating the most specific advertised enclosing prefix

Methodology

- Reports are generated daily on a per-ASN basis, and include a rolling fourteen-day window of all new, ongoing, and resolved issues on hosts within that ASN's originated prefixes.
- Reports are currently limited to the most recent 1,000 issues, and are emailed as CSV ASCII. (Although some networks have more than 450,000 active issues within a fourteen-day window.)

Findings

- 9,200 of the 21,000 ASNs currently visible to us contain at least one issue.
- China, the U.S., the E.U., and Korea are the leading overall sources of issues.
- The highest ratio of issues to addresses are on Chinese networks.

Announced Hosts by Country

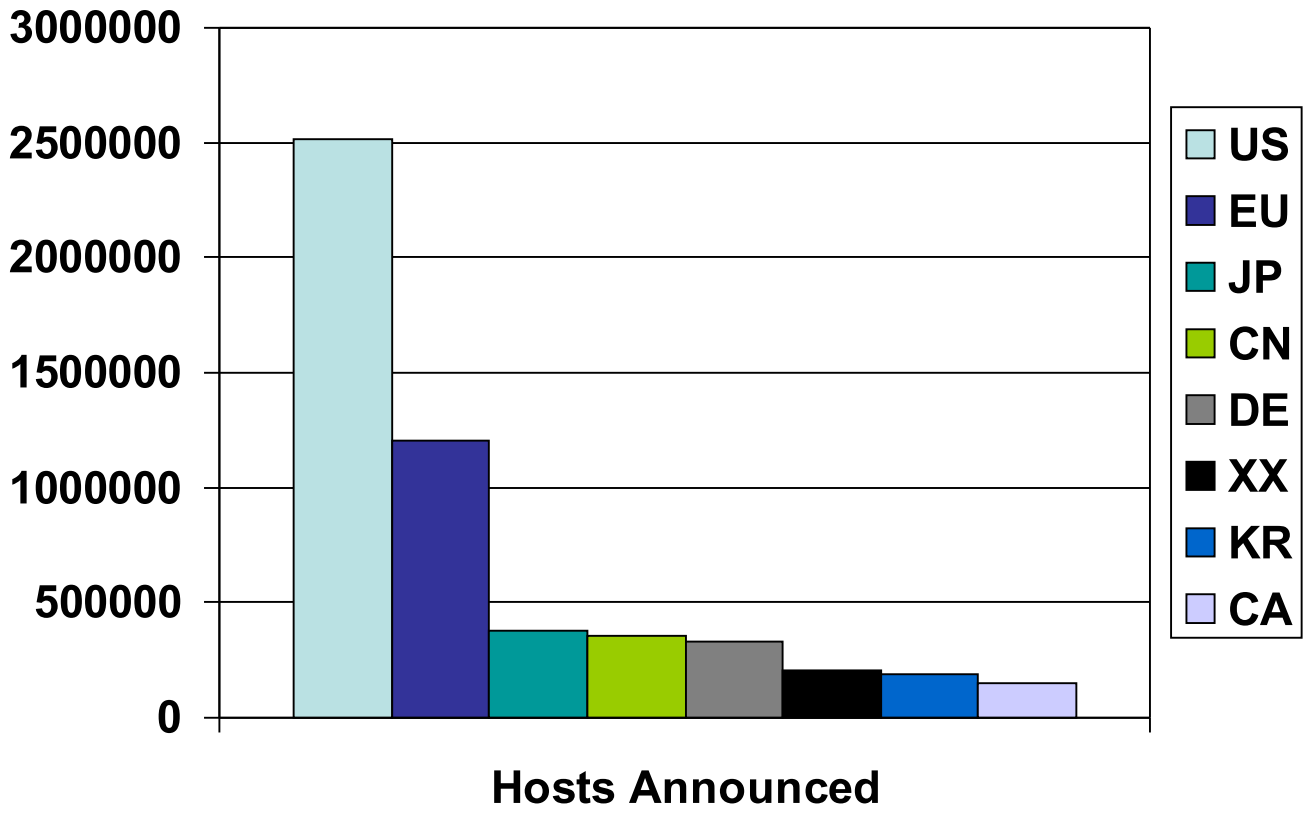

Total Issues Globally

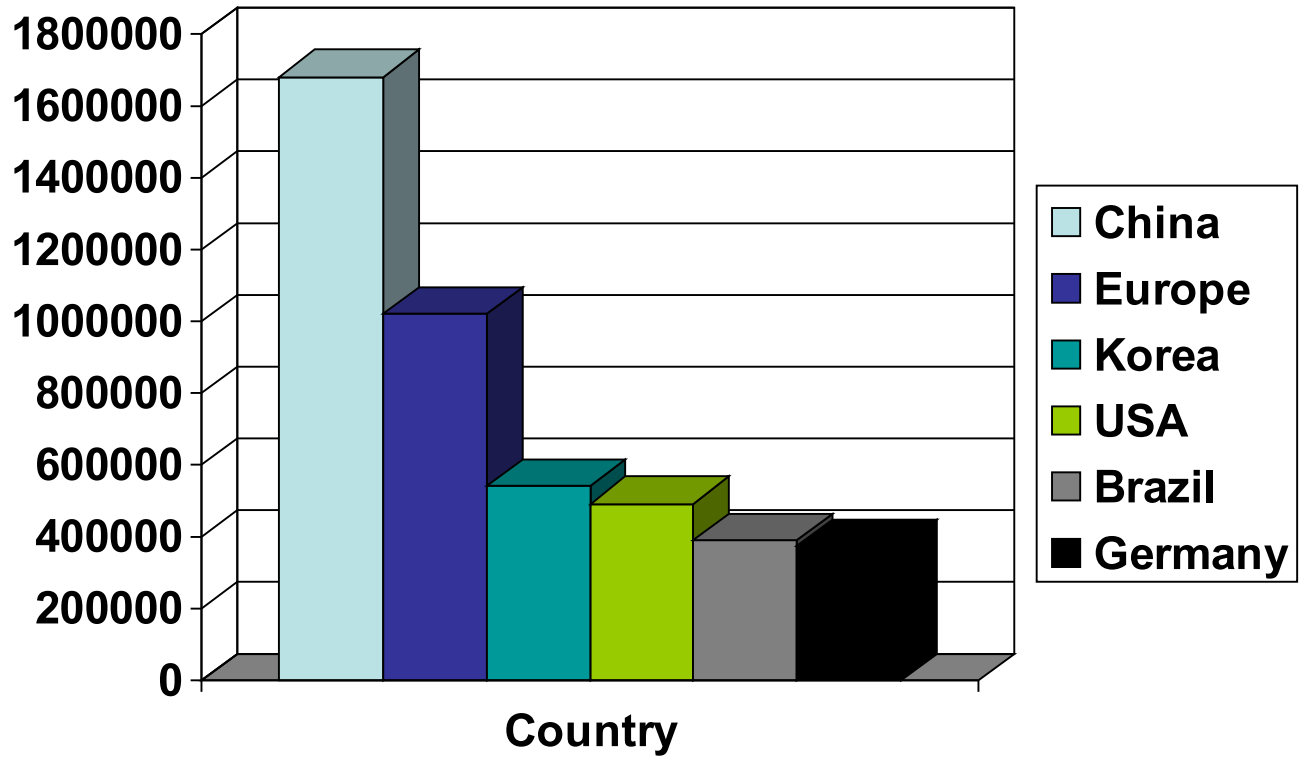

I:A Ratio Globally

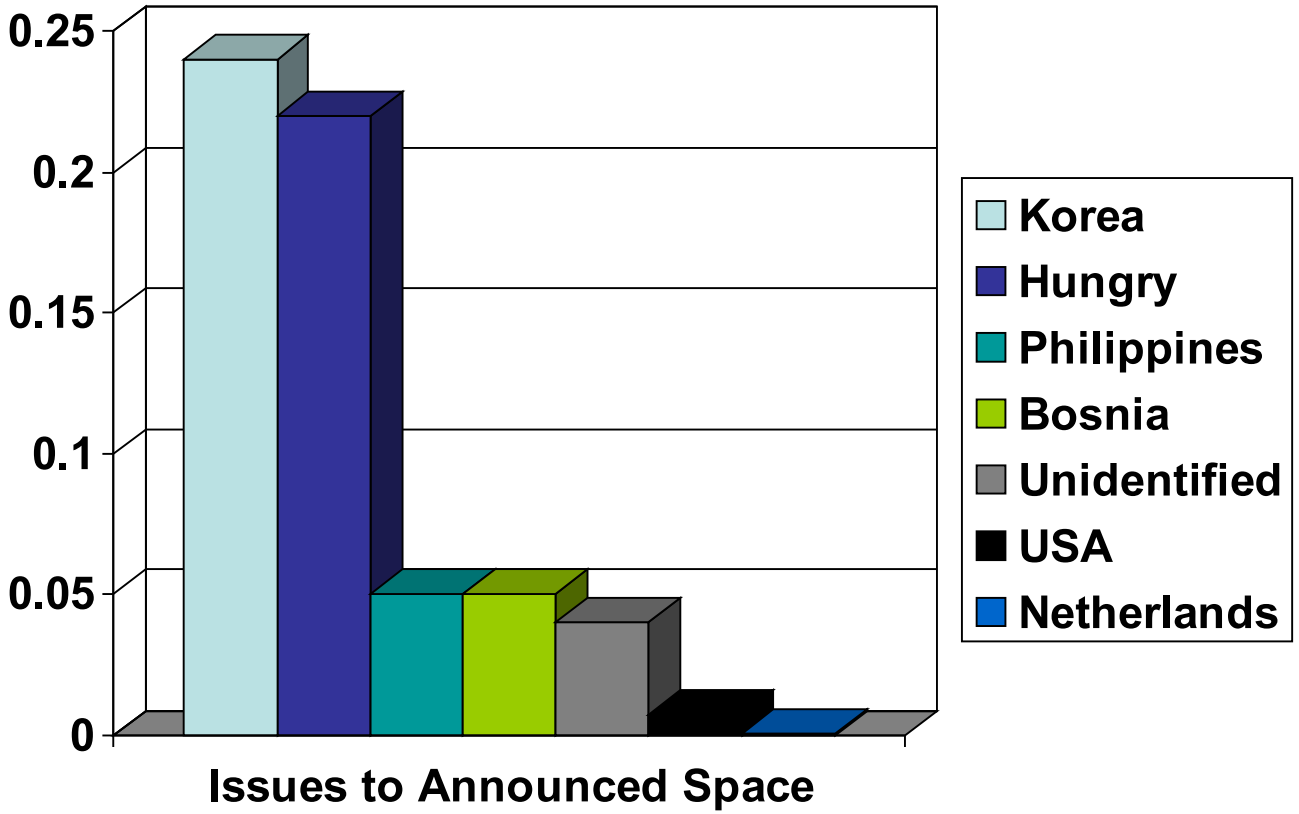

Total Issues in the U.S.

- 30 day window as of Oct 7, 2005
- Verizon (49,591:6,065,664) ~ 0.08175
- Quest 13,981:14,639,616 ~0.000955
- Charter Communications 13,209:2,128,128 ~ 0.006206
- EarthLink 9,51:1233664~ 0.007709

Worst I:A Ratio in the U.S.

- New Liberty Hospital District of Clay County, Missouri (1:2 for 1 /24)
- Foxworth-Galbraith Lumber Company (1:3 for 1 /24)
- American Central Gas Technologies, Inc (1:3 for 1/24)

Future Directions

- For large ASNs, we are beginning to provide reports as real-time XMPP / Jabber feeds
- Authenticated web display for NOC staff
- Negotiating more data sources

Acknowledgments

- Spamhaus, DSBL, CBL
- PCH (bgp feed)
- ServePath (San Francisco ISP)