Internet Traffic Trends

A View from 67 ISPs

Craig Labovitz (labovit@arbor.net)

Danny McPherson (danny@arbor.net)
Scott Iekel-Johnson (scottij@arbor.net)
Mike Hollyman (mhollyman@arbor.net)

Internet Statistics

- Golden Age
  - NSFNet 1986-1995
  - Monthly ARTs reports
    - protocol, ports
  - Filtered prefix reports

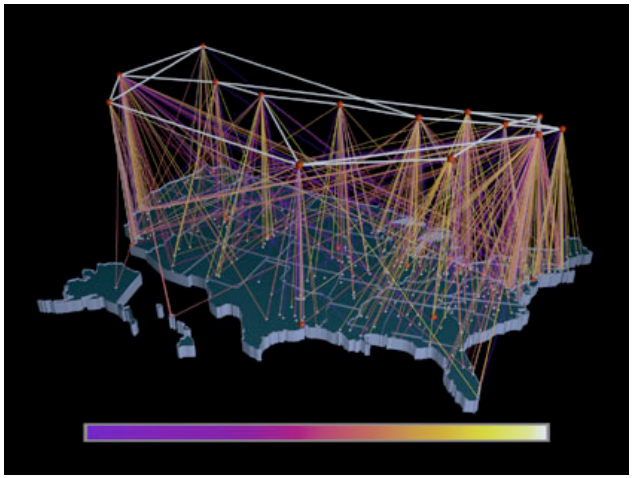

- Today
  - Some ISP specific traffic research and commercial datasets
    - e.g. Akamai, Google, etc.
  - Lots of BGP data
  - Many, many analyst reports
  - Mostly people make stuff up

Internet Traffic Project

Global view of Internet traffic and attack trends

- Leverage commercial probe deployments
  - Pool of 2,500+ active Flow / DPI collectors
  - Across 250 ISPs / Content Provider / Higher Ed
  - Deployed adjacent interesting bits of infrastructure
- Internet scale data collection
  - Traffic, DPI, Mitigation and Security datasets
  - Geographically and topologically diverse

Internet Traffic Project

- Outgrowth Fingerprint Sharing Initiative
  - 45 publicly disclosed participants
  - and Annual ISP Security Survey


THE WALL STREET JOURNAL.

Firms Join Forces Against Hackers

BY RIVA RICHMOND DOW JONES NEWSWIRES

- All data voluntary anonymous data sharing agreement with ISPs
- Still more research project than commercial
  - Arbor, University of Michigan, Princeton (Intern)
  - And 78 ISPs (and growing)

Internet Traffic Deployment

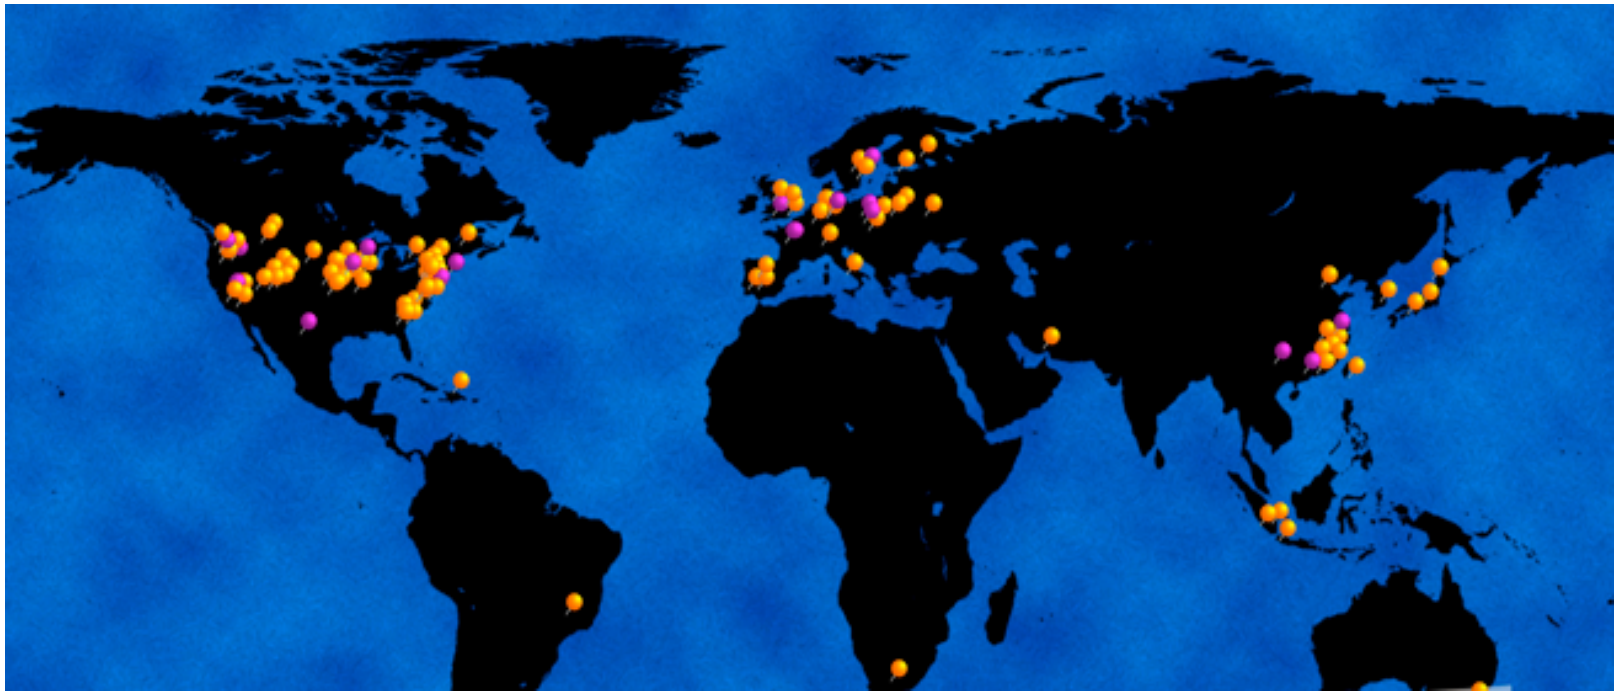

- 67 long-term participants (2 years)
- 17 unique countries
- 27 in US, though many have global footprint

Current Traffic Project Deployment

- 67 long-term ISPs (now 78)
  - 5 MSO, 4 Tier1, 15 Tier2, 4 Content, 1 R&E
  - Remainder not self-catagorized

- 1,270 routers
- 141,629 interfaces
- > 1.8 Tbps of average inter-domain traffic
- 485 days and counting (began September 2006)

Typical ISP Deployment

- Majority deployments are Flow from all peering edge routers
  - NetFlow / JFlow/ Sflow / IPFIX / etc
- Growing DPI from gigabit inline / portspan in front of customers or server clouds
- Exported to commercial probes
  - Usually 1/100 - 1/000 sampling
  - Regexp or BGP based classification of border interfaces to avoid double counting
  - Data validated against interface SNMP counters

Anonymous XML Data

- Five minute traffic samples
  - Traffic In/out of network (subset of backbone traffic)
  - Cross-products based on top N protocols, ASNs, ports, applications, etc.
- Traffic anomaly data
  - Combination protocol signatures, behavior and statistical variance from baselines
  - Distinguish Attack versus Flash Crowd
    - Annotations and mitigation status
- Self-Categorization
  - Tier1/2/3, Content, High Ed, etc
  - Predominant geographic coverage area

Internet Traffic Project

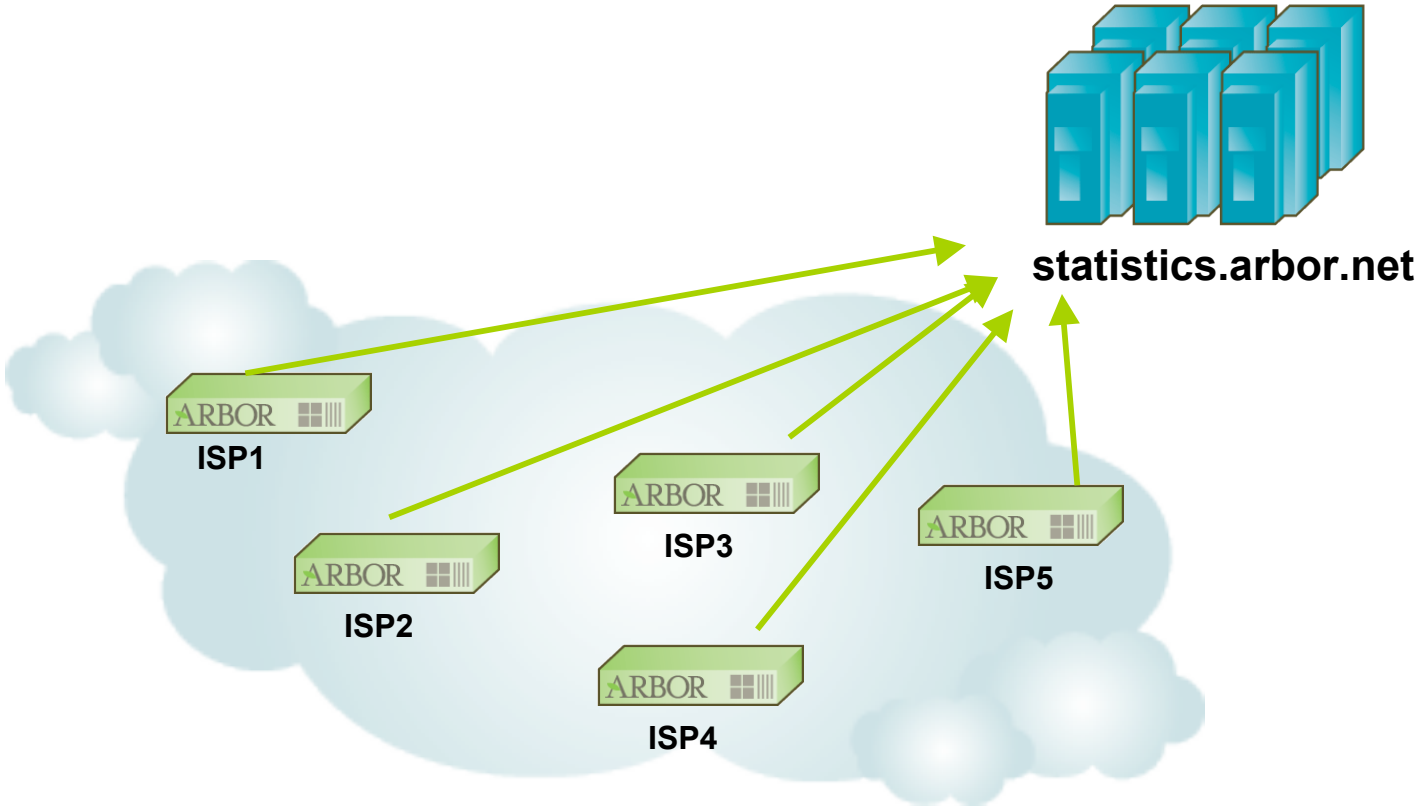

- Each participating ISP deployment submits XML
- Anonymous XML over SSL every hour
- Arbor managed servers collect/process

90 Day Protocol Distribution Trends

- No real surprises: TCP dominates followed by UDP
- Possible North America / Europe bias to dataset given diurnal patterns
- Wither IPv6?

60 Day TCP Port Trends

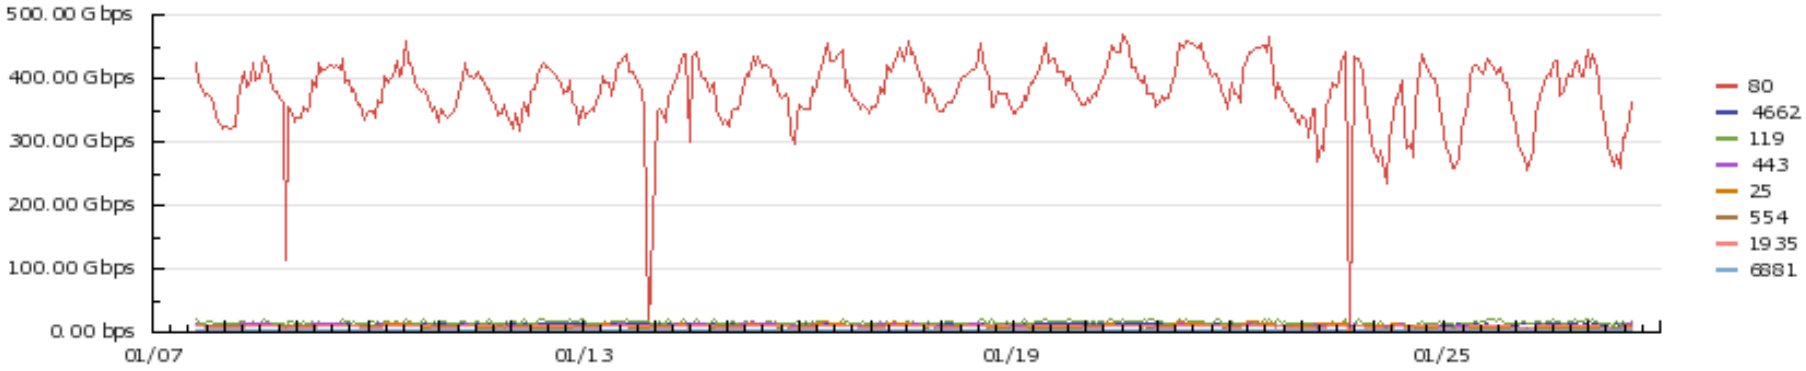

- Again, no surprises: http/80 by far most prominent TCP port
- In second place, TCP/4662 (edonkey) most prominent inter-domain peer-2-peer file sharing protocol
- Rises of NNTP (ranks 3rd) as file sharing alternative (alt.binaries!)

Internet Anomaly / Attack Summary

- 485 days
- Anomalies
  - 616,631 total anomalies reported
  - 1,271 average anomalies/day
- Attacks
  - 353,588 classified attacks (57.34%)
  - 729 average attacks/day

Internet Attack Propagation

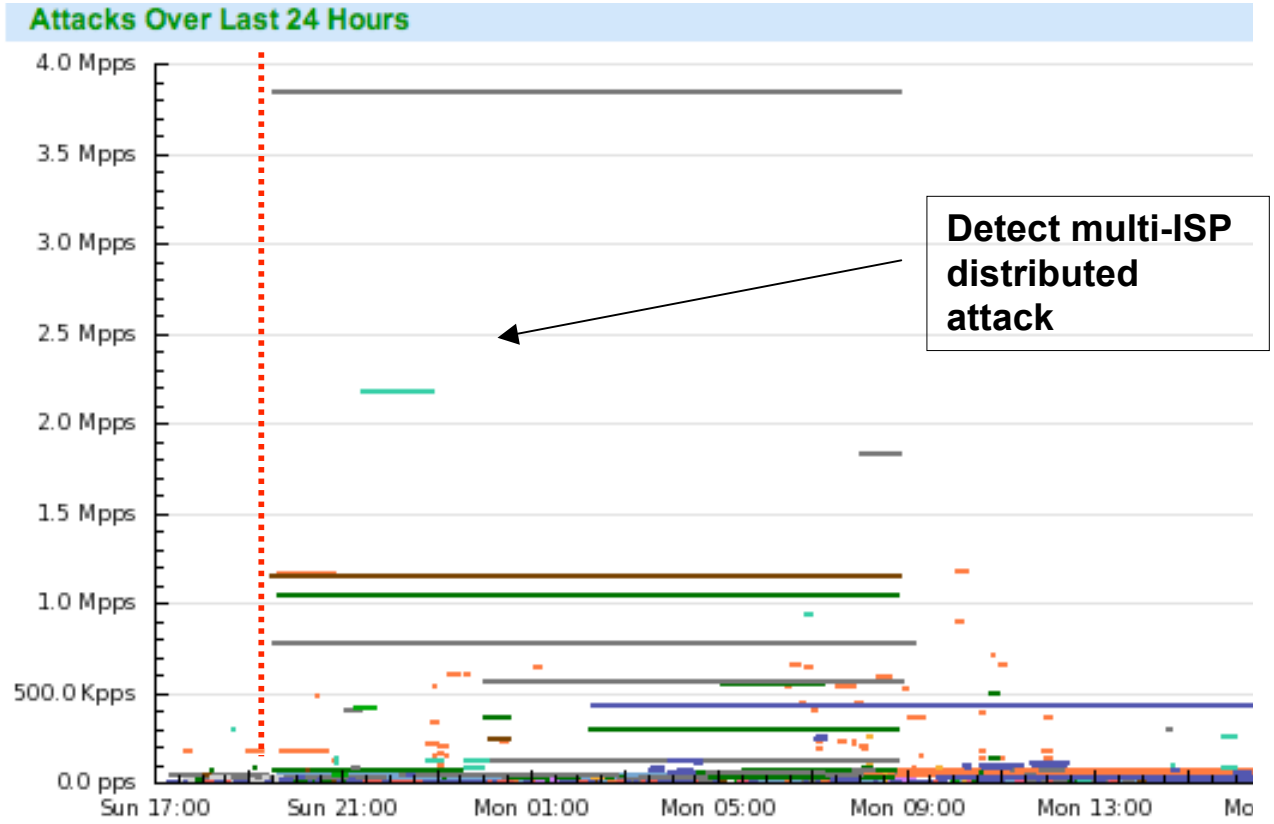

- Each color represents different anonymous ISP (30 represented)
- Each line represents different attack
- 7 Outbound ISPs, 10 attack streams (7 tcpsyn, 3 icmp) generating 6.312 Mpps, one Russian AV Vendor

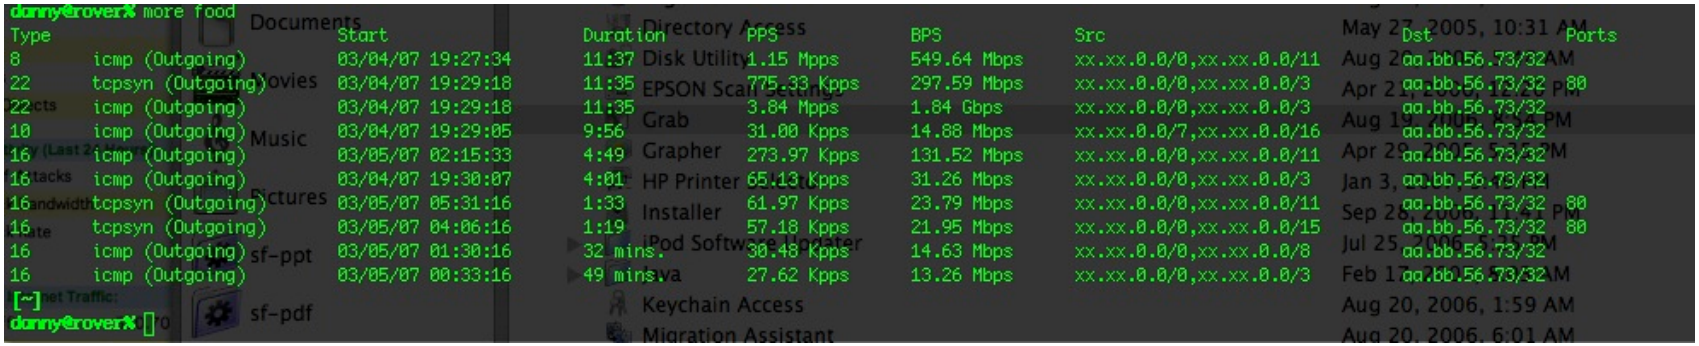

Internet Attack Packet Size Distribution

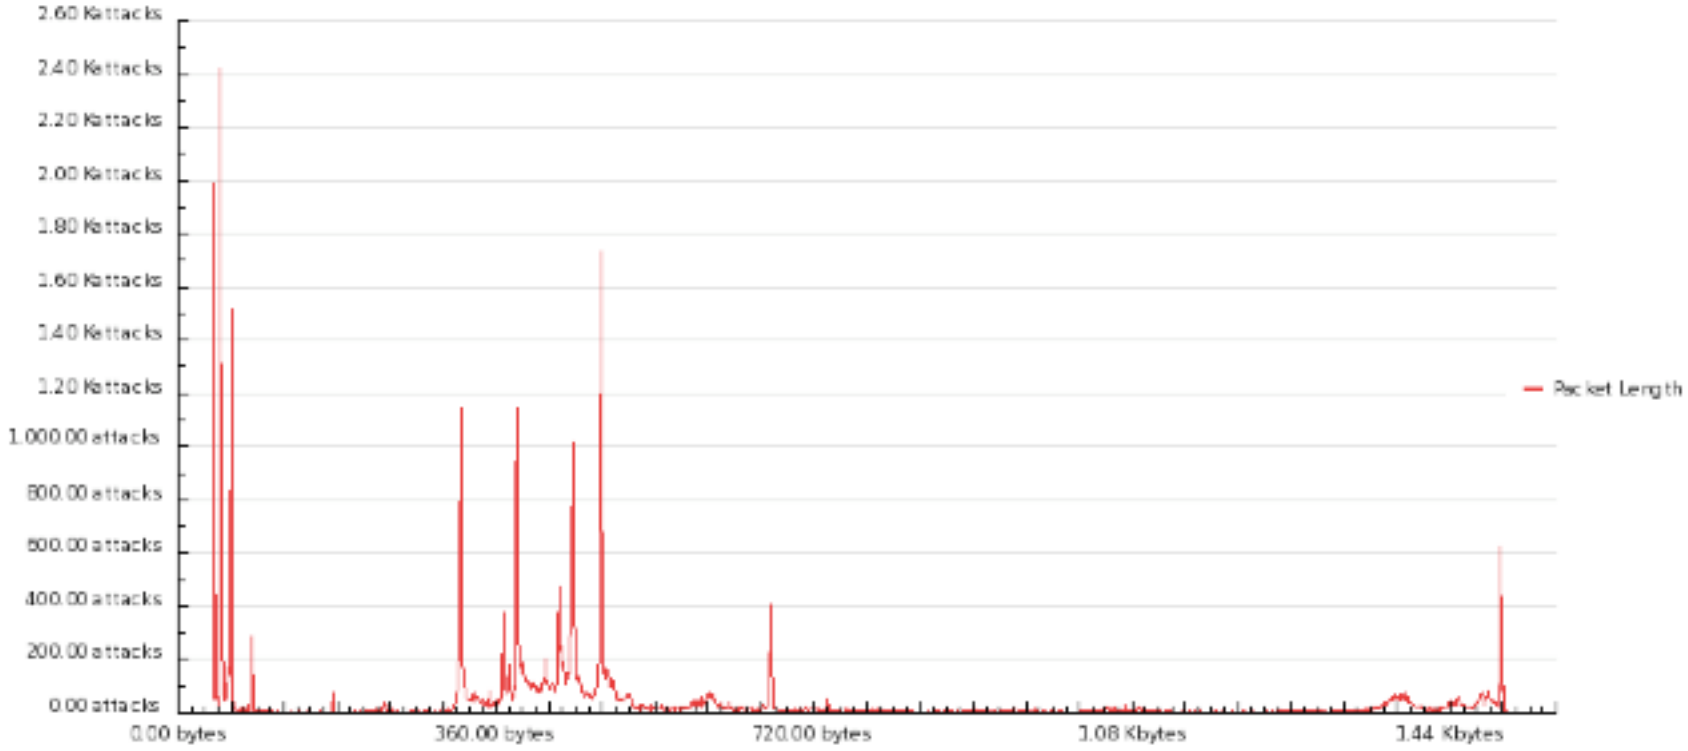

- Small packets predominate (pps attacks)
- Spectral analysis-like fingerprints of other attack types and tools
- Some issues with data collection methodology…

Internet Attack Types

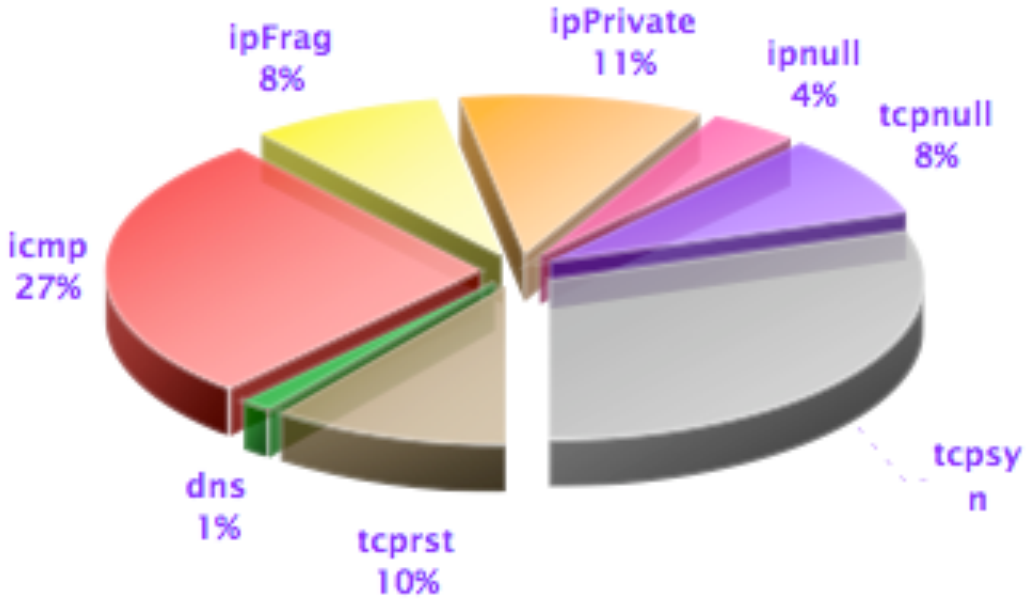

- Data excluding floods
- After nine years, TCP SYN (31%) still dominates DDoS
- ICMP (27%) also prominent

Most Frequently Attacked Ports

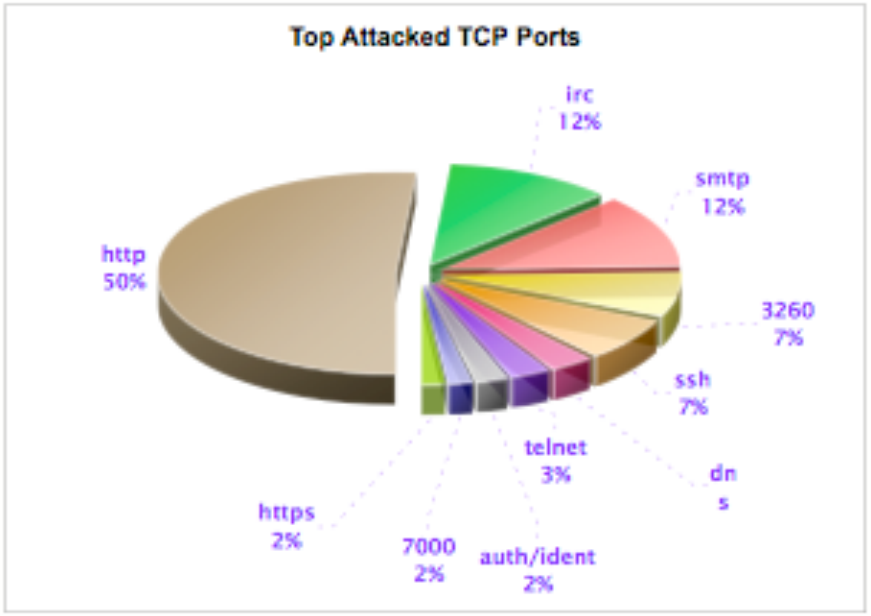

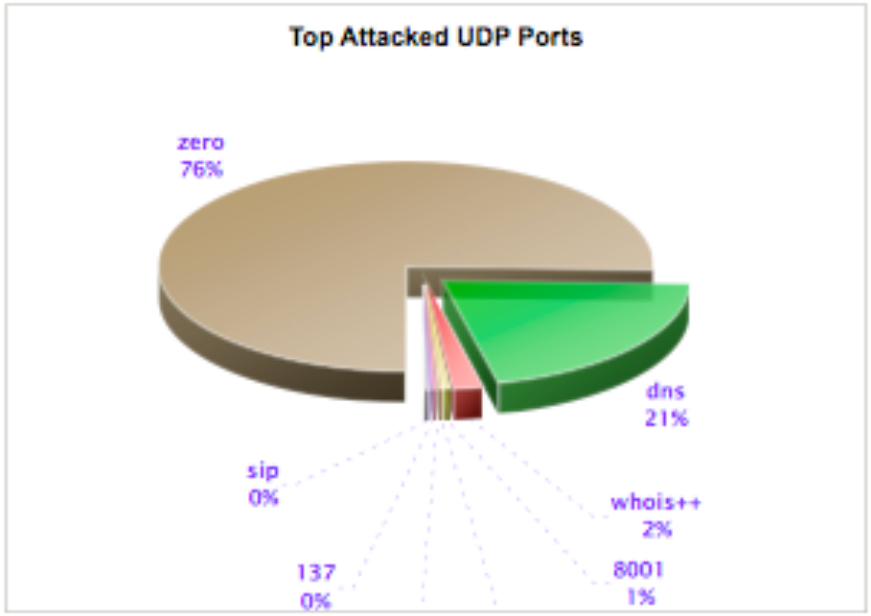

- HTTP ports account for bulk of TCP-based attacks
- Fragmentation attacks lead the pack on the UDP front

Internet Attack Scale

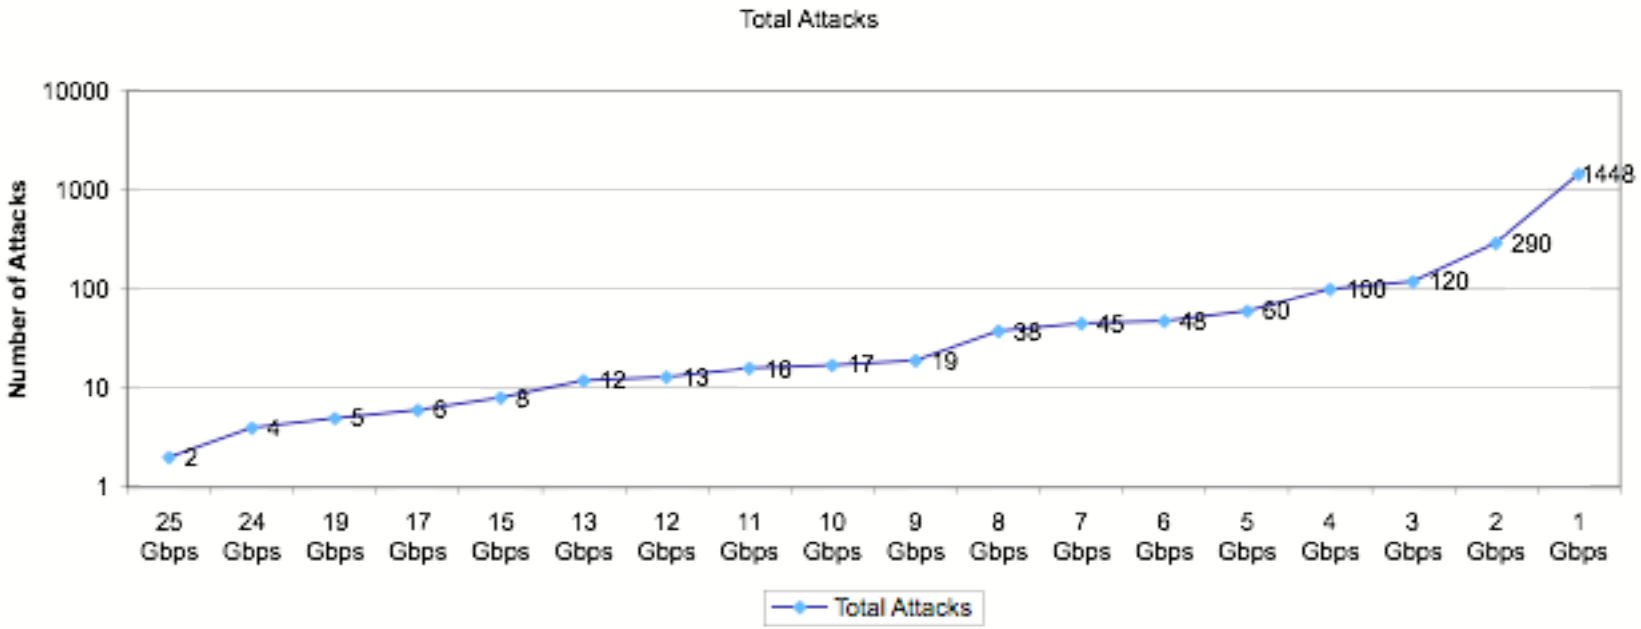

- Unique attacks exceeding indicated BPS threshold for single ISP
- Average of three 1-Gbps or larger attacks per day over 485 days of collection
- Two ~25 Gbps attacks reported by a single ISP (on same day, about one hour apart, duration of ~35 minutes)

21 Days Y/Y

- Significant Y/Y growth
- Identify additional trends: Holiday Season typically slow time for attackers

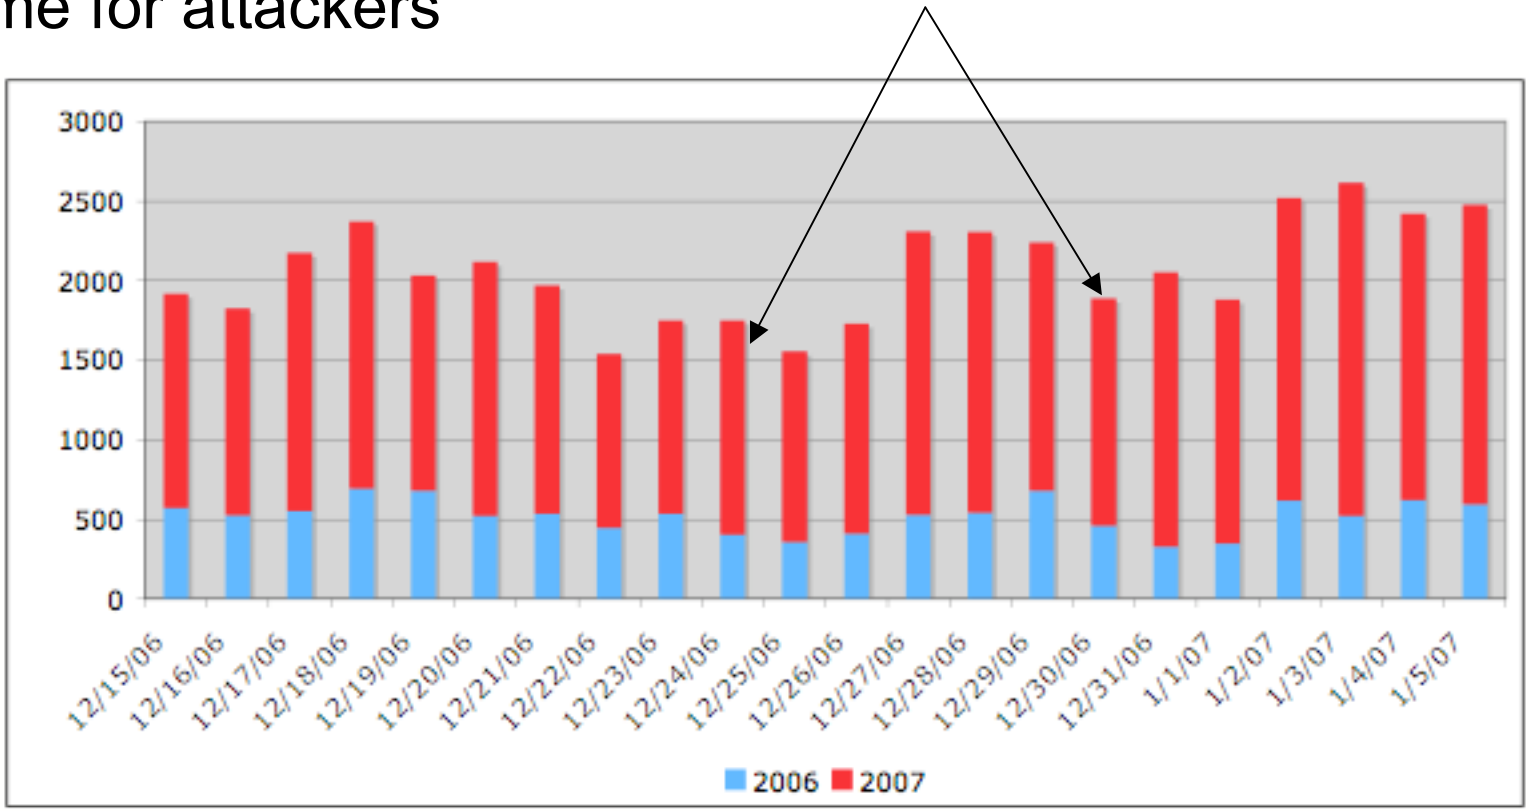

Challenges

- Balance commercial privacy with research and business interests
- Data normalization / extrapolation
  - Differing notions tier1
  - Many business units within an ISP
- Data availability to other researchers

Questions?

Craig Labovitz (labovit@arbor.net),
Danny McPherson (danny@arbor.net)
Scott Iekel-Johnson (scottij@arbor.net)
Mike Hollyman (mhollyman@arbor.net)