Security Implications for Network Architecture

Avi Freedman
Chief Network Scientist, Akamai
Chief Network Architect, FastNet

Main Question

- Do we spend enough energy worrying about security, vs. scalability and reliability?
  - Proposing some thought topics; lab and real-world testing needed
  - Mainly, the goal is to get a dialog going about what medium-sized and large networks can do about potentially serious attacks

Why Worry?

- We have had huge outages caused by OSPF or IS-IS redistribution into BGP over the last 5-6 years. Often, networks have many tens, or hundreds, of routers in their core IGP areas, and it only takes physical access to one and 5 minutes to melt a network, large or small. Also, networks often don't have the most secure backend authentication and monitoring systems.

Why Else?

- The attacks that we could see in the relatively short-term could be devastating in terms of ability to knock out hundreds or core routers – there have really been no serious attacks on the infrastructure *itself*.
- Of course, BGP instability etc could then cause even larger cascading failures. Since few have budget to buy new routers, and top-of-the-line deployed routers still have little sophistication on CPU protection, how can we get the beginnings of some protection?

Vulnerabilities: CPU Protection

- CPU Protection
  - We know that there are affects on router CPUs and BGP churn due to attacks on applications on the Internet (past worms)
  - Imagine 100k-1m hosts infected, with most hosts attacking routers with forged-source packets
  - 100mbits/sec is about it to the CPU on most deployed high-end routers, and filtering tricks are useful – but many boxes can’t run much in the way of filtering.

Vulnerabilities: Redistribution

- There's no evidence that route redistribution has ever been an attack, but when BGP -> IGP redistribution happens it generally means a multi-hour outage (avg. one big network per year for the last 5-6 years is affected, though).
- There are some things that can be done to help prevent redistribution from happening or from being as severe.
- Main question for avoiding attacks: Segmenting areas and ensuring limited access to routers in critical (backbone) area(s)

Vulnerabilities: Interconnection

- Just a note. Many of the networks that I work with are gravitating towards 4-5 buildings for interconnection, and are moving or decommissioning old interconnects.
- Not sure what the solution is, but as a community we need to think hard about putting all of the eggs in one basket.

Vulnerabilities: DNS Infrastructure

- Two pieces: Resolving and authoritative. In some sense, authoritative may be a larger concern for minute-by-minute outages.
- Another piece: Software diversity. Nsd, ?nominum work? – on the authoritative end.
- May want to consider anycast and filtering for both resolving and authoritative.

Vulnerabilities: Router Control

- Some networks are running on fairly old (un-maintained) systems doing AAA, logging, control, etc.
- Worth a watch against internal attack on these systems. Also, may want to consider ensuring that enable access != root access on these systems.

Architecture Implications: Common Sense

- DNS: Diverse DNS software; anycast and filtering for both resolving and authoritative NSs. Also remember that most deployed NS software is limited at the box CPU/software, not at the ethernet.
- Router hardening: OOB Access – make sure there's out of band console and power cycle (if possible) access so that IGPs can be shut or routers rebooted if endless churn ensues.
- Interconnection: Hard fiscal question, but worth considering "insurance" value of not putting all the eggs in 4-5 baskets.

Architecture Implications: R&D (1)

- Redistribution
  - Does segmenting IGP areas help limit the effect of churn with 100k+ route redistribution?
  - In any case, physical security needs to be ensured for core IGP routers
  - Also worth real look at incremental SPF or IGP CPU throttling

Architecture Implications: R&D (2)

- CPU Protection: Isolate control and monitoring from customer packets
  - Please, no screams…
  - Suggest investigation of ways of logically partitioning the network that routing, control, AAA (CPU packets) pass on from the "user packet population"
  - One possibility: MPLS and separate routing tables
  - 2nd possibility: Frame encap on SONET circuits/use of VLANs on gige/faste; filters to prevent input to "internal" prefix(es) on "dirty" sub-interfaces.
  - Still need some filtering or vendor work to limit the cases where packets from other interfaces can touch the CPU (unreachables, ICMP echo, etc)