Resilience of the Internet Interconnection Ecosystem

Chris Hall
Ross Anderson
Richard Clayton
Evangelos Ouzounis
Panagiotis Trimintzios

WEIS
14th June 2011

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

NPL
National Physical Laboratory

enisa
European Network and Information Security Agency

ENISA report

- European Network and Information Security Agency: ENISA
- Formal study written for them, accompanied by questionnaire responded to by many domain experts and a report giving a detailed analysis of the results
- Written (mainly) by Chris Hall, one time peering coordinator for a large UK ISP – documents the reality of how and why ISPs interconnect and the resilience issues that arise
- Original report 240 pages, has executive summary that has been reworked for an academic audience as our WEIS paper
- Read the original, you're guaranteed to learn dozens of things that you never knew before.

What's "peering"

- ISPs have customers who want access to "the Internet"
- ISP purchases "transit" ie: a contracted service to swap packets with any other address on the Internet
- ISP may reduce their costs by "peering" (usually for free) with others nearby (to reduce costs of link) ISPs. Saves the both having to pay for transit; so win-win
- IXPs (Internet Exchange Points) provide many potential peers at a single place (usually a shared "peering LAN")
- One of things the report draws attention to is the rise of "content networks" who will peer with anyone (often at IXPs)
  - they are now so important that transit providers probably could not cope if content provider network failed.

Reachability and performance

- BGP (Border Gateway Protocol) distributes reachability info
  - it's insecure (and can be slow to converge in the face of change)
- Customers care about congestion (and latency and jitter)
  - BGP cannot signal information about capacity
- BGP has very few mechanisms for "traffic engineering"
  - in the face of congestion engineers have little info & little to tweak
- Disasters have been dealt with by ad hoc routing and by neighbourly assistance
- But that assumes that it's routes that are lost, not capacity
  - no provisions for traffic prioritisation in a disaster
  - and probably not a decision that society would wish ISPs to make

Economics of transit

- Marginal cost of providing transit to a new ISP is almost zero
- Hence prices have been falling rapidly as networks compete
- Partial transit (regional routes only) undercuts full transit

- Effect is that all the transit providers are losing money
- #1 and #2 have recently merged (to have 55-60% of market)

- Risk of misuse of “significant market power”... our recommendation that regulators start to get up to speed predated this merger, but is given impetus by it

Measurement difficulties

- ISPs may have a limited understanding of where traffic is flowing on their networks – they know next to nothing about their neighbours' networks.
- Can probe but
  - mainly establishes reachability, not capacity
  - tells you nothing about backup routes (if any)
- Most of what we know comes from "experiments"
  - catastrophes (Katrina, 9/11 etc)
  - cock-ups (PK blocking of YouTube, route leaks etc)
  - side-effects of academic research (big BGP packet incident)

Recommendations

1. Incident investigation (by independent body?)
2. Network performance measurement
3. Research into network performance & resilience
4. Develop & deploy secure inter-domain routing
5. Research into AS (ie ISP) incentives
6. Sponsor Best Practice
7. Independently test equipment & protocols
8. Regular disaster recovery exercises
9. Contingency plans for possible transit market failure
10. Traffic prioritisation may be needed in disasters, preplan
11. Greater transparency on security (maybe educating purchasers)

Resilience of the Internet Interconnection Ecosystem

http://www.enisa.europa.eu/act/res/other-areas/inter-x

http://www.lightbluetouchpaper.org

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

NPL
National Physical Laboratory