Towards Best Practices for Active Network Measurement

Christos Papadopoulos (Colorado State University)

and

John Heidemann (USC/ISI)

Feb 13, 2009

AIMS Workshop, San Diego CA

USC Viterbi School of Engineering ISI Information Sciences Institute

Colorado State University

If You Run Active Experiments then You've Seen This:

To Colorado State University Team:
Your activity is not acceptable and will not be tolerated.
Please keep your activities from attempting to access or scanning our networks.

Additional attacks or scans will be reported to the authorities.

What was the "attack"?

We pinged their network once every 11mins..

USC Viterbi School of Engineering ISI Information Sciences Institute

Colorado State University

The Aegis of Research is not Enough..

…..

If our IP is not removed from your list by the deadline.....that would be a very bad thing.

A very bad thing indeed.

I will leave it at that. I have about hit my limit with spam.....the "important research".....and all the other crap that is going on on the internet....You bastards have directed a probe directed to our static IP address....and I am flipping ticked off about it....big time.

…

This person appears to know that our pings were part of a research project.

We still got no sympathy.

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

..nor is Being Very Careful!

NANOG list, Jan 12 this year:

I'm not entirely certain what is going on but has anyone noticed some strange announcements for 174.128.31.0/24?

....

Interestingly enough, ARIN indicates this is a part of range they have assigned for reachability testing.

http://ws.arin.net/whois/?queryinput=174.128.31.0

Follow-ups:

-> I would think that it would only be polite to notify people about what is going on so that other people do not waste their time looking for phantom issues.

-> ... having the courtesy to notify next time would be very much appreciated. I was headed into a family member's funeral when I received the hijack notification. I took the 15 minutes to do some quick investigation, fire off a few emails informing my colleagues of the issue and "arrived" at the funeral a bit late. Perhaps in the future it would be better not to play with my toys without asking my permission first?

-> The thread generated **96 messages!**

Pretty high, even for Noisy NANOG™

USC Viterbi School of Engineering ISI Information Sciences Institute

Colorado State University

Solving the Meeting Problem

- Network operators dealing with our experiments **lose** money sorting experiments from attacks
- Academics running such experiments **gain** knowledge and publications
  - .. but lose goodwill when mistaken for attackers
- The level of clue among network operators varies a lot
  - .. some network operators are really small-business people
- The level of clue among academics running experiments **should** be high
- Therefore, we believe it is **OUR** job to design experiments carefully and inform the network operators in a timely manner

USC Viterbi School of Engineering ISI Information Sciences Institute

Colorado State University

Some Current Community Efforts

- RFC 1262: Guidelines for Internet Measurement Activities (1991) (Vint Cerf)
- Issues and etiquette in use and sharing measured data (Allman 2007)
- Planetlab documents
- CAIDA, Wisconsin, ISI measurements and web documentation
- Legal issues (Paul Ohm 2007, KC's pamphlet)
- Predict MOA's
- etc..

USC Viterbi School of Engineering ISI Information Sciences Institute

Colorado State University

Learning from Other Communities

- the medical community's has decades of best practices in human-subjects research
  - The Belmont Report (1979, now a historical document) and HHS regulations
  - Institutional Review Boards (IRBs)
- some principles from Belmont
  - risk-benefit criteria
  - subject selection
  - informed consent of subjects
- but we need to figure out how we're *alike* and *different*
  - and *our* best practices and standards for IRBs

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

Current Practices not Enough

- RFC 1262 is too brief, too generic and too old – little beyond calling to "Do the Right Thing"
- There is no easy/standard way for operators to lean about our experiments
  - We currently rely on researcher expertise, good will and coziness with network operators
  - ..and fail all too frequently
- There are virtually no guidelines for newcomers into Internet-wide active experiments
- There is no way to find out if anyone else is doing the same experiment to avoid duplication
- We generally do not involve nor provide guidance to IRBs

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

Important Questions

- Is my experiment appropriate?
- Is my experiment necessary?
  - Can I do it without active measurement (in a lab)?
- Is my experiment harmful?
  - How disruptive will it be to others?
  - How should I notify others?
- Where is the balance?
  - How can I manage complaints?
  - Is the balance positive?
- much more…
- *We don't have answers to these and many other such questions.*

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

Secondary Issues: Anonymization

- What is the appropriate anonymization strategy for my experiment?
- Are there tools and guidance available to help me?
- Has anyone developed such tools and are they willing to share?
- Do I need to talk to a lawyer?

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

More Secondary Issues

- What is the appropriate etiquette for passive measurements?
- How should I secure my collectors?
- What should I expect the provider to disclose about the data I am capturing?
- What are the guidelines for reporting results?

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

Proposal: Two Prong Approach

- **First: Need updated Best Practices document**
  - Should be a community effort
- **Second: Centralized Database** of current and planned academic experiments
  - Make it trivial for interested parties to get information about current and planned experiments

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

Best Practices Document

- What should be included?
  - What sections should it cover?
  - What should be in these sections?
- What existing documents should we tap?
- What legal issues should it address?
- How do we get network operators to contribute?
- What is the right process to publicize the draft and who should approve it?
  - Should we push it through the RFC process?

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

Blog

- Need something more than a passive document
- A Blog can be a centralized location for quick and hopefully painless answers
- So with apologies to John Stuart, let us propose..

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

PeskyAcademics.com

- Blog with search facilities to find academic experiments quickly
- Academics (and others?) register all current and planned experiments
- Entries vetted so hackers cannot register bogus experiments
- Users search based on IP address and possibly other fields
- Result has a link to the experiment with explanations and contact information

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

Other Possible Features

- Mailing list announcing new experiments with information and opt-out links
- "Do not call list" that can be proactively distributed to researchers
- Announcements/updates of experiments
- Help operators with technical issues, e.g., how to quiet alerts
- Track top queries to determine peskiness
- What else?

Downsides?

- Would public disclosure hurt any experiments?
- How do we avoid an inrush to the "do not call" list?
- Are there any security concerns? Any unintentional benefits to hackers?
- Other downsides?

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

Publicity

- Publicity will be hard, but proportional to usefulness
- Publicize the blog (and BP document) with operators and academics so advertise everywhere (mailing lists, NANOG, conferences, etc.)
- Provide links on the blog that showcase research results to educate and encourage operator tolerance

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

Good for the Community Too

- Easy way for academics to find each other
  - Informs others of what type of data is being collected and how to get it
  - (Segway into Predict?)
  - Avoids duplicate efforts, encourages collaboration
- Showcases academic activities to funding agencies

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University

In Summary

- We believe that a Best Practices document is well overdue
  - RFC?
- Document alone is not enough: need an active entity to track experiments
  - Blog?
- We are calling on the community to do both

USC Viterbi School of Engineering ISI Information Sciences Institute Colorado State University