Experimentation with network-based security mechanisms

GENI Security Workshop
January 22-23, 2009
UC Davis

G. Kesidis
EE and CSE Depts
Penn State
kesidis@engr.psu.edu

.

Outline

- GENI experimental context.
- Experimental progression as part of a generic engineering design cycle.
- Theoretical/formal phase.
- Simulation/emulation phase.
- Prototypical deployment.
- Specific examples and component problems, *e.g.*,
  - Experimental scale-down, and
  - Traffic generation.

GENI experimental context

- In the following, we are considering:
  - a network-based security mechanism under test in the context of
  - a “clean slate” network architecture.
- A security experiment would therefore need to specify:
  - a network topology (open or closed) including peripheral end-systems,
  - background and attack traffic,
  - a network architecture spanning:
    - addressing/packet-format,
    - name resolution,
    - routing/forwarding,
    - and possibly layer-3 protocols for connection establishment, authentication, *etc*.,
  - and the security mechanism under test (possibly implicitly part of the network architecture).

Generic engineering design cycle

1. Device conception/design.
2. Formal/theoretical evaluation based on models of the designed device/system and its operating conditions.
3. Testing with increasingly greater realism and cost:
    1. Simulation
    2. Emulation
    3. Prototypical deployments

- Redesigns possible after each test phase.
- Each test phase should be conducted and documented so as to be "repeatable" by a third party, to within assessed statistical confidences in the performance and complexity metrics.
- Each test phase may involve consideration of:
    - Presence of and interoperation with competitive devices/systems,
    - Incremental deployment strategies to improve rate of adoption,
    - Assessment of a "control" device/system and comparison against the competition.

Different perspectives on theoretical/formal study

- "Unfortunately, understanding network performance is more of an art than a science. There is little underlying theory that is actually of any use in practice. The best we can do is give rules of thumb gained from hard experience and present examples taken from the real world." from A.S. Tanenbaum. Computer Networks, 3rd Ed. Prentice Hall, 1996, p. 555,556.
- V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling", *ACM/IEEE ToN*, 1995.
- J. Cao, W.S. Cleveland, D. Lin and D.X. Sun, "Internet traffic tends *toward* Poisson and independent as the load increases", in *Nonlinear Estimation and Classification*, Springer, 2002.
- My own experience is that theoretical/formal performance evaluation, consciously conducted in highly idealized and simplified network settings, are pursued and valued by industry.

Trace-based traffic replay, modelling

- Attack traffic recreation.
  - In a network setting, might not need to recreate the host exploit and thereby avoid containment issues.
  - How to develop “variations” of known attacks to test the robustness of a defense, while avoiding the perception of developing new attacks.
- Background traffic recreation is obviously important for assessment of false positives.
  - Far more activity in the research community on forensics than on employment of network trace traffic traces for testing purposes.
  - Dissemination of anonymized traces greatly improved through the activities of, *e.g.*, CAIDA and PREDICT.
- Methods of realistically “light” salting of traces by, *e.g.*, cover low-intensity attack activity (b/g traffic rarely captured with interesting attacks *in situ*).
- Need to characterize session-level “demand” from traces motivated by the need to experiment with:
  - high volume attacks, and
  - new network architectures (even just different layer 4).

Scale-down for theoretical, simulation or emulation testing phases

- Given limited resources for simulation and emulation, may need to reduce the scale of the experiment by using a much smaller “open” or “closed” network.
- Scale-down of an open topology as in Thevenin-Norton equivalent circuits.
- Clearly, also need metrics to assess fidelity of scaled-down model to the original.
- Some preliminary results by the DHS/NSF EMIST team for scale-down of
  - attack traffic (scanning worms), and
  - inter-AS network topologies (attacks targeting BGP).

Example: 128:1 scaled-down SQL Slammer worm

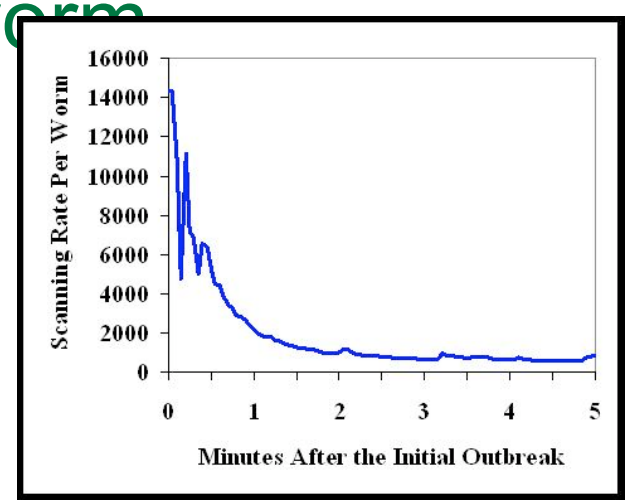

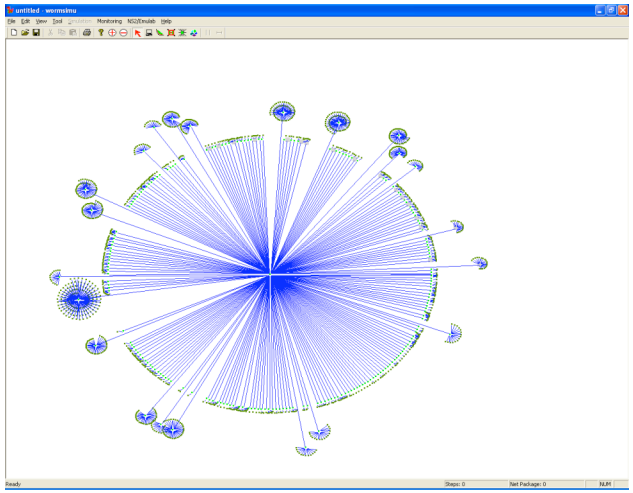

- Characteristic outbound scan-rate saturation, but total attack traffic negligible compared to core background traffic.
- A SIR model of worm spread recreated this characteristic saturation [TOMACS'08].
- 128:1 scale-down experiment on DETER [WORM'04]:
  - Idealized Internet core connecting access (stub) links
  - 600 susceptible SQLs residing behind <1000 stubs
  - Stub links to core are distinct access (bandwidth) bottlenecks
- Need not emulate actual method of host infection, but can vary scanning strategy, see EMIST's scanning worm tool on DETER experimenter's dashboard.
- EMIST's development of tools for experiment specification, visualization, scale-down, traffic generation, etc., was potent outreach activity.

Incentives and security

- Network trace data can inform “utilities” modeling user behavior – important for games studying effects of economic incentives.
- As a result of the network neutrality ruling by the FCC, renewed interest in incentives to deter excessive consumption (BitTorrent) under flat rate plans.
- Comcast residential broadband access recently migrating from flat rate *F* toward pricing formula involving usage-based charges above a threshold (*i.e.,* overages, as commonly used at NNI): $F + R(U-\Theta)^+$
- Such a pricing formula naturally leads to an authentication problem, renewed interest in differentiated services, *etc*.
- Examples many such mechanisms have already been standardized and are already deployed in the commodity Internet, *e.g.*, CMTS DOCSIS, AT&T’s DSL U-verse, MIDCOM, and a lot of “traffic engineering” (TE) technology including diffserv, scheduling, Ethernet (802.1p).

Prototypical deployment - GENI

- I understood GENI's original story as a testbed intended to be able to:
  - simultaneously mount different network architectures under test,
  - somehow assess them through engaging
    - actual data/service providers (& p2p networks) which could shop data or services they want to mount among the architectures under test, and
    - actual end-users that would access the "GENI" data/services through interfaces with the existing Internet.
- Security experimentation:
  - Need for attack isolation my preclude "deliberate" attack experimentation.
  - Need to facilitate defense deployment and assessment tools.
  - Need to facilitate deployment of other types of security mechanisms to manage different types of authentications, reputation/referral systems, *etc*.
- Generally, the testbed thus conceived faces problems, *e.g.*,
  - reconfigurable routers, associated experimental artifacts, and
  - Management of experimental resource allocation and associated fairness issues.
- Again, interesting research problem of *incremental deployment strategies* to
  - maximize performance in a "hybrid" environment and, thereby,
  - maximize likelihood of growth in deployment/adoption, *i.e.*, survival of the fittest.

GENI defaults

- Availability of a “default” network architecture, or one chosen from a library, which can be modified by the experimenter.
- In particular, availability of default:
  - connection-oriented network architectures,
  - incentive systems, and
  - associated billing/book-keeping and authentication mechanisms.
- Note that modifications may need oversight so that commodity Internet does not experience unexpected problems through the GENI interface.

Acknowledgements

- EMIST(-DETER) ‘03-’07 project team, particularly
  - S.F. Wu, J. Rowe of UC Davis
  - V. Paxson and N. Weaver of ICSI/UC Berkeley
  - P. Liu and D.J. Miller of Penn State
  - S. Fahmy of Purdue
  - P. Porras of SRI
  - S. Schwab of SPARTA
  - J. Evans (NSF) and D. Maughan (DHS)
  - Tools: DETER’s experimenter’s dashboard at www.isi.edu/deter
- Talks on security experimentation at the NSF `08 “Science of Security” workshop by
  - Roy Maxion
  - John Mitchell