CS 294-28: Network Security

Prof. Vern Paxson

http://inst.eecs.berkeley.edu/~cs294-28/
http://www.icir.org/vern/cs294-28/

vern@cs

May 12, 2008

Lecture Outline

- Research themes
  - General
  - Particular to network security
- Topics we didn't cover
  - Feedback here helpful
- Topics we did cover
  - "take-away" points/themes, not tech.
  - Chime in with feedback/questions
- Course survey ~ 2:15PM

What Is This Class?

- Brand new graduate course on network security
  - Brand new = it will be bumpy at times
  - Graduate = focus on reading papers, participatory discussion, major project
  - Network security = how do we keep our computer networks functioning as intended & free of abuse
  - Network = heavy emphasis on global Internet
    - And not much emphasis on host-side issues

General Research Themes

- All papers have shortcomings
  - Doesn't mean you can't extract value
- For your own work:
  - Frame limitations
  - Be thorough & generous towards prior work
  - Provide insight into tradeoffs
- Methodological issues
  - Gauging data quality
  - Bootstrapping (perhaps) ground truth
  - Partition development vs. assessment data

General Research Themes

- Replication/criticism of prior work is unfortunately very rare
  - Corollary: little research upside to publishing data
- Research does not proceed as presented in a well-written paper
- Topics can heat up excessively
  - Multicast, QoS; Traceback, worm models
  - Crucial task for successful research is problem selection

Network Security Research Themes

- Evasion-proof is not a realistic goal
  - Research progresses in often-pretty-modest steps (*building blocks*)
  - "Raising the bar" has definite utility
  - Today's evasion problem looks different tomorrow
  - But: *do* frame evasion picture
- Field changes very fast
  - Including serendipity
  - You need to figure out how to be nimble

Research Themes, con't

- Beware the problem of Crud
  - Surprising diversity of benign activity
  - Great utility in obtaining real data
- We're constantly trading off
  - Especially false positives vs. negatives
- Beware funding ecosystems (and popular press)
  - E.g., DARPA's need for metrics
- Historically, publishing attacks has been worthwhile
  - But not guaranteed

Some General Techniques

- Bayesian decision-making
  - $P(A|B) = P(B|A) \cdot P(A)/P(B)$
  - Underlies broad class of ways to compare the past versus what's happening now
- Non-actionable analysis can still yield high-quality input into additional analysis
- Related trick: cheap pre-filtering to winnow down resources spent on more expensive processing
- Offload state (securely) to untrusted party
- Leverage observations of trustworthy "neutrals"
  - E.g., resisting evasion; probing caches

Topics We Didn't Cover

- Infrastructure protection (DNSSEC, SBGP)
- Peer-to-peer
- IPSec/VPNs
- Phishing, spyware (and not much on spam)
- Underground economy
- Group Security
- Sensornets
- Vehicular networks
- Security of e-commerce
- Attacker infrastructure (scam sites)

Authentication / Identity

- Granularity of entity: person/service/system
- Low-level mechanisms all long worked out
  - Problem is cost: computational & management
  - Practical revocation is especially unsolved
- Attribution & filtering vs spoofing, laundering
- Leveraging more limited notions of identity
  - *Personas* via consistent-signing
  - "*Duckling*" model for imprinting

Denial-of-Service

- Via logic/algorithm errors
- Via flooding
- Amplification & reflection
- Attackers resisting filtering ⇒ spoofing
- Spoofing ⇒ backscatter ⇒ telescopes

Telescopes

- Global insight
- But:
  - Non-homogeneous
  - For global events, exhibits *lag*
- Passive systems readily scale
- Active/responsive systems much more challenging
  - Including problems of filtering, containment
- *Attractors*: how to bring in traffic (e.g., spam trap)

Traceback

- Spoofing ⇒ packet source localization
  - To provide relief; prevent future use; deterrence
- Marking: key notion of introducing header state reflecting packet's communication properties
- Hash-based: key notion of Bloom Filters
  - V. efficient (probabilistically correct)
  - Privacy-preserving ("provenance")
  - Fine-grained/single packet

Marking Evolves

- PI: *deterministic* marking provides filtering handle
  - Focus on relief, not traceback/attribution
  - Issues of collateral damage
- Capabilities: simple mechanism for network to enforce receiver consent
  - Design tradeoffs: header space vs. attacker work factor
  - Raises denial-of-capability problem
  - Can integrate with *puzzles*:
    - Source given partial capability, needs to compute the rest
- Fully end-system quasi-capabilities: SYN cookies

DDoS Defense Space

- Filtering: TTL, PI handles, Pushback
- ✓Spoof prevention: ingress filtering, SYN cookies
- Level-the-playing-field: Puzzles, Defense-by-Offense, CAPTCHAs
- Hiding: Overlays, lightweight authenticators
- ✓Spreading: CDNs, Anycast, load-balancing
- Incentives: Re-Feedback
  - Address *externalities* via bi-lateral agreements
- ✓Overprovisioning

Intrusion Detection

- Field developed reactively, from multiple directions
- Key issues:
  - False positives vs. negatives (Base Rate Fallacy)
  - Threat model (local; reflects policy)
  - Evasion
  - Actionable decisions (intrusion prevention)
- Network vs. host-based tradeoffs
  - Performance (HIDS), management (NIDS), disambiguation (HIDS), extent of trust (NIDS), visibility (HIDS), breadth of context (NIDS)

Styles of Intrusion Detection

- Known signatures (*syntax*)
- Known misuse (can include semantics)
- Anomaly detection
  - Must ground in the *domain*
- Specification-based
  - Define what's allowed, prohibit all else
- Behavioral (contextual evidence)
  - E.g., "unset HISTFILE"

High-Performance NIDS

- Bro's layered architecture
  - Initial packet filtering (no longer effective)
  - Judgment-neutral distillation of activity
    - In semantic (usually app-level) terms
  - Sharp separation of mechanism vs. policy
    - Event engine does **not** generate "alerts"
- Utility of extensive logging (*Time Machine*)
- Thorny problem: state management
- Increasing need for parallelized execution

NIDS Evasion

- Deep problem due to *ambiguity* + *crud*
  - Presence of ambiguity often not actionable
- Occurs at multiple layers
  - App-layer extremely problematic
- Pushes towards active network elements
  - "*normalizer*"
- Brings out additional issues:
  - State management abetted by in-line operation
  - Analysis-friendly protocols (e.g., reliable RSTs)

IDS Evaluation

- Deep Problem #1: rich "normal" behavior
- Deep Problem #2: desire for black-box analysis/ranking
- Goal should be: Illumination
  - *Why* do FPs/FNs occur (or TPs/TNs, even)?
  - *How* do system's different elements contribute?
  - *Which* parameters are relevant in what ways?

  ⇒How will it work in other contexts??

Worms

- Relevance:
  - Latent threat (cyberwarfare)
  - Groundwork for botnets (tech transfer)
  - Large-scale analysis
- Morris (1988): highly innovative; global
- Code Red, CR2, Nimda (2001): dynamics due to bugs, competition, programmed die-off
- Slammer (2003): speed from fire-and-forget
- Blaster/Nachia (2003): 10Ms of Windows boxes

Reasoning about Worms

- Amenable to math (SI model, logistic growth)
- Targeting: random scanning, localized, hit-lists, permutation, meta-server/topological, contagion, flash
- White worms: bad-idea magnet
- Self-stopping: don't rely on hearing chatter

Worm Detection

- Signature distillation:
  - Given pool of benign/malicious flows, find discriminating substrings
  - Polymorphism offers many degrees of freedom
- Network-based behavioral: *contact graphs*
  - Only works after infection has spread somewhat
- Host-based behavioral: taint-checking
  - Self-Certifying Alerts ameliorate trust issues

Honeyfarms

- Low vs. high-fidelity honeypots
- High-quality detection signal: network propagation
  - But no good for bots :-(
- *Toxicology spread*, signature distillation
- Issues:
  - Filtering (remove scans; lightweight replay)
  - How to detect VM is “done”?
  - Malware employing anti-VM technology
  - Containment (liability, fingerprinting)

Forensics

- Goals: assess damage / fix holes / attribution / legal pursuit
- Perhaps for large-scale events via analysis of contact graph
- Witty analysis: surprising power of structure

Scanning

- Difficult to crisply define
- "Background radiation" / evolution over time
- TRW:
  - Parameterized in terms of problem domain basics
  - Principled FP/FN/detection speed tradeoffs
- DNS-based:
  - Lookup precedes connection attempt
  - DNS provides remote visibility
    - As with cache probing; DNSBL counter-intelligence; sinkholes

Traffic Analysis

- Side channels: power/threat of information leakage
- SSH keystroke inference
  - Hidden Markov Model to reduce search space
  - Entropy as means of assessing opportunity
- Stepping stones
  - Structural model driven off of empirical invariant
  - Not actionable, but high-level feature detection

Legal / Ethical Issues

- Distinctions:
  - Contents vs. addressing vs. storage
- First two require consent or "provider exemption"
- Storage = after recipient has read it
  - Contents: only disclosed w/ consent or court order
  - Headers: yes, except "public" service providers cannot disclose to government entities

Detecting Web Server Attacks

- Exemplar of apt anomaly detection
  - However, detection not necessarily actionable
- Handy set of statistical approaches
  - Distribution outliers, mismatches, inferring/generalizing structure
- Should this topic have been structured differently/deeper?

Detecting Web Client Attacks

- Crawler-based browser honeypots
- Detection based on system state changes
- "Landing" pages that redirect to "malware distribution" pages
  - Drive-by download: automated infection
- Some content more likely problematic
  - But also plenty of ≈ innocuous content
  - Weird prevalence of Chinese sites

VOIP Security

- PSTN trust model: barrier is SS7 network
- Complexity of VOIP space: naming, name resolution, rendezvous, middlebox traversal, retargeting
- #1 operator goal: no free calls
- Setup decoupled from media path
  - Leads to trust oriented around proxies, not end users
- Skype: engineered to resist analysis
- VOIP spam: need real-time filtering, pre-content

Anonymity

- It *does* matter
  - Journalists, human rights, whistle-blowers, law enforcement
- Tor model: download topology from K server, build circuits
  - Also supports hidden services
- Core tensions: ease of deployment, acceptably low latency, liability burden
- Attacks leverage “corner-cutting”
  - External traffic (e.g., DNS resolution)
  - *Option distinguishability* (fingerprint via quirks)

Wireless

= zillions of limited devices coming our way

- Space is rapidly evolving/innovating

- Lots of information exposed
  - Device inference ("inventorying")
  - Persistent naming (linkable across changes)
  - Exposed resource discovery
- Medical devices: software radio attacks, zero-power defenses
- Cellular attacks via: cost of instantiating communication, shared control channels

Trace Anonymization

- Tension: utility of released information vs. threat model
  - Goal: managing risk, not preventing it
- Core problem: adversary possesses external information
  - E.g. unmasking of address anonymization via presence of sequential scanners

Botnets

- Key technique: infiltration
- Graybox testing to extract behavior profile
- Power of DNS monitoring:
  - Global meas.; C&C sinkholes; counter-intel
- Mark-and-recapture for population est.
- P2P-based C&C: lots of room to improve :-(
- But are they making Big Bucks?