Can cryptography secure the Internet ?

Richard Clayton

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Santa's Crypto
Get-Together 2008
4th December 2008

Outline

- Threat scenarios
- Type of attackers
- The Domain Name System (DNS)
- Secure DNS
- The Border Gateway Protocol (BGP)
- Secure BGP ?
- Securing email

CAUTION

This talk describes possible attacks on Internet infrastructure, especially DNS & BGP. But, not all of these attacks work everywhere, and people may be reluctant to discuss whether they work or not in their part of the real world.

So don't assume it's all entirely true!

However, it isn't entirely false either!

Any mention made of particular networks, ISPs or countries is merely to make abstract ideas concrete, not an analysis of actual flaws.

NB: Do not try any of this at home (OR at work)

Threat scenarios

- “I wish to stop the Internet working”
  - not very lucrative – bad guys need a working Internet too!
  - but may be relevant if you’re fighting a war!
- “If you don’t pay, I’ll stop the Internet working”
  - however, getting paid is not a trivial problem (c.f. kidnaps)
- “I wish to capture a significant amount of incoming email to a major ISP mail server” (and the attack should “scale”, so 0-day exploit of **`sendmail`** not considered here)
  - much more realistic way to improve your lifestyle!
  - email may contain passwords etc
  - email can be made to contain passwords etc
  - answering email often “proves” identity
  - also provides further opportunities to blackmail the ISP, or trash their reputation as being secure

Types of attackers

- Back bedroom attackers
  - can nowadays have control of a reasonable size botnet
- Criminal entrepreneurs
  - may 0wn (or own!) a smallish ISP in Ruritania
- Organised crime ??
  - simpler for them just to bribe an employee!
- Nation States
  - danger! can afford an Internet scale(-ish) test-bed
- For none of these attackers can it be assumed that BGP or DNS are too obscure to be attacked effectively!
  - DNS is what translates human names into host addresses
  - BGP is how routers learn where packets are to be sent
  - both underpin every other Internet system we have

Underlying strategies

- Cannot just steal every packet – people notice
  - c.f. YouTube outage in February 2008 (Pakistan Telecom)
  - in fact, most of what we know about BGP vulnerabilities comes directly from an extensive history of configuration errors!
- Accept email, resend to the correct ISP
  - top 50 senders in logs is a give-away, so use a botnet
- Reject email end of data with a 4xx response
  - email generally re-delivered after a delay, so suitable for intermittent attacks
- Tunnel stolen SMTP packets to correct place
  - either a peer of target or customer within target
  - complex to get this right, but is unlikely to be noticed
- Arrange for packets to take a detour past your monitoring

Domain Name System (DNS)

- Request sent by UDP : **`dig www.highwayman.com A`**
- Response sent by UDP

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 35955
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 3, ADDITIONAL: 3
;; ANSWER SECTION:
www.highwayman.com.     7200    IN      A       193.111.200.151
;; AUTHORITY SECTION:
highwayman.com.         7200    IN      NS      ns0.gradwell.com.
highwayman.com.         7200    IN      NS      ns1.gradwell.net.
highwayman.com.         7200    IN      NS      ns2.gradwell.net.
;; ADDITIONAL SECTION:
ns0.gradwell.com.       36603   IN      A       193.111.200.7
ns1.gradwell.net.       66051   IN      A       193.84.87.10
ns2.gradwell.net.       65985   IN      A       216.218.244.58
```

DNS (I): Active attacks

- DNS server asks for data (using UDP)
  - attacker supplies incorrect answer first
    - 16 bit identifier is not long enough!
    - but, modern software randomises request port
    - new idea is to add entropy by case of characters in the name
- Older software was often flawed
  - predictable random numbers!
    - or even accepts non-authorised data!
- No-one monitors for attacks
  - however this scales badly, so of limited interest
  - BUT WAIT!

DNS (II): Dan Kaminsky 2008

- Ask for multiple sub-domains (sub1, sub2 etc.)
  - neat way to ensure that resolver always has to ask
- Attacker tries to get their answer in first
  - BUT of course only poisons some obscure sub-domain
- Kaminsky realised could supply NS data as well
  - “in-bailiwick” data (extra info from authoritative server)
  - relied upon for some purposes! So devastating attack!
- Mitigate (only) with lots of entropy (as before)
  - and what of clever servers behind rather dumb firewalls?
- Mitigate (expensively) using TCP requests
  - can be a useful technique if realise that under attack
- Only real fix is DNSSEC (of which more in a moment)

DNS (III): Phishing

- “Rock-phish” gang spoofed GoDaddy Aug07 (and this year Network Solutions and others have also been attacked)
  - probably just wanted some cheap domains
  - BUT control of a registrar account permits changes to name server identities
  - So significant amounts of infrastructure protected by eight character passwords passed over the web !
- Registrars for grown-ups will check validity of changes out-of-band, $10 hosting services will not
  - a significant number of US banks were vulnerable
- Attack vector (delivering password sniffing capability onto end-user machines) might also be drive-by malware…

DNS (IV): Subvert a “Root of Trust”

- 13 top level name servers (A-M)
  - maximum that will fit in a DNS response
- Included with BIND (etc) as a text file
  - you have to start bootstrapping somewhere!
- L moved from 198.32.64.12 to 199.7.83.42
  - moved 1 Nov 2007 (warnings sent 24 Oct 2007)
  - AS20144 (ICANN) announced route until 2 May
- BUT other AS’s also announced route
  - Dec 15 (AS42909), Mar 18 (AS 4555), Apr 1 (AS9584)
  - all serving the right thing (through May)
    - we think!

DNSSEC (cryptographers to the rescue!)

- Provides cryptographic signing of DNS results
- First attempt (RFC2535: get your parent to sign) appears unworkable at Internet scale
- DNSSEC-bis (RFC4033, 4034, 4035) signs at domain level, and gets that key signed by the next level up the tree
- Eventually the root (“.”) will be signed, until then you can start from “trust anchors” and work down from there
  - so far, Brazil (.br), Bulgaria (.bg), Czech Republic (.cz), Puerto Rico (.pr) and Sweden (.se) have signed their zones
- To avoid excess work by parent two keys are used, a zone signing key ZSK and a key signing key KSK
  - parent signs the KSK, the KSK signs today’s ZSK

DNSSEC – the zone walking problem

- DNS needs a proof-of-non-existence facility
  - wildcard values (MX etc) only apply to names that don't exist
  - so need to respond "qwerty.example.com doesn't exist"
  - scope for a Denial of Service attack if create denials on the fly
  - so pre-sign "nothing before smtp.example.com"
- Can now enumerate the zone in linear time
  - geeks took view that DNS was public… but lawyers educated them
- Fix is to use hashes (hurrah for the cryptographers)
  - hash the "qwerty" name and respond with next higher hash
  - incorporate "salt" to prevent building of pre-image dictionaries
  - also permit <n> applications of the hash (to heat the planet)
  - now defined (and being deployed) as NSEC3 records

DNSSEC – what goes wrong?

- Hard to tell yet!
- Needs more resources in caches, servers, resolvers…
- EDNS (larger packets) has given us new DDoS attacks!
- Keys time out (30 day default for ZSK – for “security”)
- Keys have to be passed securely to the parent
- Applications don’t understand fancy responses
- Users don’t know whether to trust unsigned responses

“If you think that cryptography solves your problem then you don’t understand cryptography and you don’t understand your problem!” (Butler Lampson & Roger Needham)

DNSSEC – what does it mean?

- There's a lot of politics involved in "signing the root"
  - should it be ICANN? the US Government? Verisign?
- Currently you can make up your own mind who to trust
  - ISC "lookaside" registry
  - or collect Trust anchors from websites (using https!)
- Perhaps everyone should sign?
  - Ruritania signs the root and local ISPs use that trust anchor
  - people worry about fragmentation (c.f. Chinese root servers)
    - you don't fix social problems by having one root
- Perhaps no-one should want to sign?
  - if an incorrect record is signed in error then who do I sue?
  - how much money are they good for when I'm defrauded?

Border Gateway Protocol (BGP)

- Routers announce their AS number and “routes” or “prefixes” that they will accept traffic for (e.g. 158.152.0.0/16 in AS2529)
- They pass on appropriate information prepending their AS
  - 158.152.0.0/16 3549 1239 1239 1239 2529
- Routing is over the “most specific” route with lowest AS count
  - NB: ignoring communities, local preferences etc
- Hence prefer “shortest path” to AS2529 from “here”
- BUT if 158.152.17.0/24 is announced by someone else (perhaps this is the network of a dual-homed customer) then that route will be preferred for 158.152.17.42, even if it is a longer path!

Using BGP to attack mail systems

- Basic idea: announce a /32 for mailserver
  - BGP will prefers the “more specific” announcement
- Traffic then flows to announcing AS in Ruritania
  - email contents are available for inspection
- /32 may not propagate, so /24 may be better
  - /24 leads to complexity if other hosts or services
  - hence tunnelling packets back to ISP may be a sensible strategy (and just sniff them as they pass)
- Sniffing possible anyway at other ISPs
  - difference here is scale and remoteness

Cryptographers to the rescue ?

- Secure BGP (S-BGP, from BBN)
  - PKI to authenticate ownership/linkage of prefix, AS and router
  - UPDATE records carry a signed “attestation”
  - IPSEC used to secure links between routers
- Secure Origin BGP (soBGP, from Cisco)
  - certificates statically bind AS and prefix list
  - policy certificates state that more specifics are invalid
  - routers must attest to a path to originating AS
  - envisaged as a web of trust (local decisions to accept signatures)
- Neither approach is going anywhere fast

Attacks on router-router links

- In 2004 Paul A. Watson pointed out that it was relatively easy to reset TCP links by sending a single RST packet.
  - many stacks only checked that the RST was within the “window”
  - some only checked the offset was positive !
  - need to get the port numbers right, but not a lot of entropy, especially for long-lived BGP connections!
- There was already a cryptographic defence!
  - TCP MD5 option (RFC2385)
  - Every packet contains MD5(packet header & contents | secret key)
  - Great enthusiasm for this (despite it being MD5) – until people started worrying about denial of service attacks!
  - sysadmin mailboxes now full of MD5 “secrets”
- Neater approach is to check that “hop count” is 255
  - prevents random attacks from elsewhere on the Internet

Today's reality

- Regional Internet Registries (RIPE, ARIN &c) hold signed data
  - route objects can only be changed by cryptographically signed requests. The objects can indicate with "mnt-lower" that more specifics do not exist (which wouldn't have saved YouTube)
  - this data doesn't yet affect minute-to-minute routing decisions
- More specific routes may be spotted by monitoring
  - sFlow/Netflow can check if traffic coming from correct peer
  - MyASN @ RIPE, Renesys etc
  - note that "bogon filtering" hides route from owner! and so Best Practice can prevent give-away failures
- Almost impossible to detect unauthorised announcements (multi-homed customers are commonplace), but monitoring will indicate changes. Many day-to-day failures are customer announcements that are not correctly filtered "up-stream".

What do Route Object signatures mean?

- Suppose people started to care about the signatures on route objects then how do we know that they're correct?
  - viz: when they aren't correct, then who do we sue?
  - RIRs are currently trusted to hold valid information, so some imaginative prefix hijacking (City of Los Angeles) commenced with forged documents being sent to ARIN
- Conversely, whose signature should be present on objects, and can they be forced to withhold it.
  - There was significant interest in de-peering the "Russian Business Network", and recently Atrivo and McColo have had their connectivity removed. Would it be a good idea to permit this at the registry level (avoiding all those tedious phone calls to Russia)?
  - centralisation has problems, decentralisation may not work
- Cryptography isn't good at "meaning"

Stopping spam (using cryptography)

- All spam is the same – so collect MD5 digests
  - trivially avoided by trivial changes!
- All spam is forged so sign the headers
  - latest scheme is DKIM (DomainKeys Identified Mail)
  - SHA-256 (body & a selection of header fields (& a list thereof))
  - signs the digest and places result into email header
  - certificate stored in the DNS (so that must be trustworthy)
  - receiver can spot forgeries (eg phishing)
  - needs a reputation system to deal with unknown senders
  - so just one small part of the puzzle…
- Encrypted email (S/MIME & PGP)
  - easiest way through spam filters!
  - spammers don't use it, because no-one else (non-geek) does!

Securing SMTP (more cryptography)

- Opportunistic encryption (RFC3207)
  - uses STARTTLS capability & command
  - negotiate mutually acceptable algorithm
    - ✓ works out of the box for major MTAs
    - ✓ only end-points can decrypt the traffic
    - × increases processing load (may not matter)
    - × no “man-in-the-middle” protection
- Check certificates before sending email
  - prevents man-in-the-middle
    - ✓ works out of the box for major MTAs
    - × increases processing load (may not matter)
    - × needs a PKI (or a lot of bilateral arrangements)

Negligence

- The failure to use reasonable care
- Current test for “duty of care”:
  - harm must

    (1) be reasonably foreseeable

    (2) there must be a relationship of proximity between the plaintiff and defendant and

    (3) it must be “fair, just and reasonable” to impose liability
- If an attack is effective on a mailserver, because an ISP has failed to deploy secure DNS or BGP is it their fault?
- Short term specific: if a firewall makes DNS IP-IDs predictable, is there negligence?
- Doe$ anyone under$tand what a $ignature mean$?

...and if you'd like to know more

- DNSSEC

  Web is full of tutorials!

- RIPE

  MyASN & lots of other initiatives

- Experimental alerting systems

  `http://iar.cs.unm.edu/alerts.php`

  `http://phas.netsec.colostate.edu`

- Anirudh Ramachandran and Nick Feamster (SIGCOMM 2006)

  Understanding the Network-Level Behavior of Spammers

Where are we ?

- Secure DNS almost here
  - some TLDs already signed, more to come
  - unlikely that will be fully deployed for years
  - BUT Kaminsky exploit has given it a huge boost
- Secure BGP(s) entirely experimental at present
  - concerns about performance (cf MD5)
  - concerns about key distribution
  - when will it be stable and inter-working?
  - State-of-the-art is monitoring for wickedness
- Email
  - Crypto mechanisms work “out of the box”, but not widely used
- Economics
  - No-one knows what a signature is worth!

Can cryptography secure the Internet ?

http://www.lightbluetouchpaper.org/

UNIVERSITY OF CAMBRIDGE
Computer Laboratory