Pervasive Monitoring

stephen.farrell@cs.tcd.ie
June 2014

It's an attack

- The actions of NSA and their partners (nation-state or corporate, coerced or not) are a multi-faceted form of attack, or are indistinguishable from that
- Not unique, others are likely doing the same... or will
- The scale arguably makes this an example of a new pervasive monitoring threat model that is neither purely passive nor a classic Man-in-the-Middle and that we have not normally considered in protocol design, implementation or deployment
- A purely technical response will not “solve the problem” but we should treat an attack as we usually do and try mitigate it

A Definition

From RFC7258/BCP188 “Pervasive Monitoring is an Attack”

> Pervasive Monitoring (PM) is widespread (and often covert) surveillance through intrusive gathering of protocol artefacts, including application content, or protocol meta-data such as headers. Active or passive wiretaps and traffic analysis, (e.g., correlation, timing or measuring packet sizes), or subverting the cryptographic keys used to secure protocols can also be used as part of pervasive monitoring. PM is distinguished by being indiscriminate and very large-scale, rather than by introducing new types of technical compromise.

IETF (Re)Action

- Overall: snowdonia has re-energised folks to do better on security and privacy in general (and not solely in response to PM)
  - Side meeting in Berlin @ IETF-87
  - Tech plenary, major discussion @ IETF-88
  - STRINT workshop before IETF-89
    - https://tools.ietf.org/html/draft-iab-strint-report
  - Topic at many meetings/BoFs @ IETF-89
  - Wanting to see results from IETF-90 onwards...
- Unsurprisingly this is similar to the more broad technical community reaction

New IETF work related to PM

- UTA WG formed, update BCPs on how to use TLS in applications
  - WG has to do work now of course
- RFC7258/BCP188 published after major IETF LC debate – sets the basis for further actions
- Proposals for new work discussed around IETF-89:
  - DNS Privacy - unthinkable before snowdonia
  - TCP encryption: was proposed two years ago but mistakenly rejected
    - Including by me, as ack'd at mic @ IETF-88, bummer
- Old-RFC privacy/PM review team formed
  - Please help! Mail me.
- IAB re-factoring their security and privacy programmes.

Other relevant IETF Things

- TLS 1.3 being developed aiming for better handshake encryption properties (and learning from previous TLS problems)
- HTTPBIS WG developing HTTP/2.0, the major deployment model for which seems to be to run much much more HTTP traffic over TLS
- And since all this is IETF stuff, you can (and please do) join in and help if you're willing and able

DNS Privacy

- IETF-89 BoF materials
  - https://www.ietf.org/proceedings/89/dnse.html
    - I stole slides from there mostly:-)
- Mailing list:
  - https://www.ietf.org/mailman/listinfo/dns-privacy
- Drafts – All “unofficial” remember, nothing here has consensus yet
  - Problem statement: draft-bortzmeyer-dnsop-dns-privacy
  - Some requirements: draft-hallambaker-dnse

DNS Privacy Problem

- Query timing and content is “meta-data” that can help a pervasive monitor
  - Could correlating DNS queries (e.g. via timing) be a fingerprint for which web page you're on?
- QNAME itself can be sensitive:
  - <political-party>.<cctld> or <ailment>.org
- Full QNAME sent too often, too far in queries
  - Can go to root, not needed there, at least in principle

Problems with this Problem (1)

- DNS names likely to be exposed elsewhere anyway, primarily in TLS ServerNameIndication (SNI) which is not easy to protect
  - SNI protection is being considered in TLS1.3, but load-balancers probably need something
- QNAME is used by some CDNs and others for valid networking purposes
  - Maybe. Could have interaction with “solving” public-suffix list via DNS
- DNS privacy was never a requirement and we have DNSSEC; don't make deploying DNSSEC harder!
  - Yep, but times change. Privacy solutions can almost certainly be independent of, and complementary to, DNSSEC

Problems with this Problem (2)

- Stub<->Recursive and Recursive<->Authoritative patterns of interaction are hugely different
  - Yep. May need different approaches to privacy for those, but that might well be ok, since the privacy issues are fairly different
- If you get your DNS from DHCP, there's no point since the resolver could anyone and could be spying on you
  - Yep. But that'd mean a more active attack.
- There's no way to get this deployable via UDP
  - Maybe. But maybe there is!

Solutions

- Basic ideas for solutions are “easy”
  - QNAME minimisation: just don't! No protocol change needed
  - Crypto moving parts are fairly obvious
    - See the BoF materials
  - How to arrange those parts so that something might be deployed and useful is not at all obvious
- State of play:
  - Mailing list are discussing.
  - If interested, sign up and go

Crypto + Data Minimisation

- Mitigation = Crypto + Minimisation
  - For DNS protocol and other cases
  - Though undoubtedly we will learn more/better as we go
- That includes registries too presumably
  - whois coudn't be controversial could it?
- Are there activities feeding into policy development that are already considering PM?
  - If not should/could there be?

What to do? (1)

- Turn on crypto
  - For applications and between data-centres
  - Current tools: TLS, IPsec, IEEE MAC-sec, DNSSEC
  - Future tools?: DNS-priv, TCPInc (tcpcrypt), MPLS-OE
    - Discussions ongoing
  - Measure/gamify what is being used
- Data minimisation
  - E.g. DNS QNAME minimisation
  - More uncertain, more to learn here

What to do? (2)

- Better implementations
  - https://cryptech.is/ and similar
  - Update/check/audit crypto support
  - Make security/privacy admin easier
- Deployments
  - Turn on stuff that helps privacy
  - Significant issues with business models and deployed base of services
- Users
  - Target diversity - Don't all use the same services all the time

What to do? (3)

- Discuss the issue openly
  - In whatever fora are relevant for you
- Agitate (if that's your kind of thing:-)
- Go and be responsible engineers/computer scientists/whatever and take the broader implications of your work/research into account before, while and after doing it

Conclusions

- IETF has consensus PM is an attack (RFC7258) and is working that problem
- We all should consider how we can work to make PM harder, since those doing it will not just stop
- When/if societies do decide that PM is as bad as it is, then the technical community should have in place the tools to effect that decision