



Pervasive Monitoring

stephen.farrell@cs.tcd.ie
May 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

Contents

- What is pervasive monitoring?
- Who's been a naughty signals-intelligence agency then?
- What to do?

- Bored and prefer video?
 - <https://www.youtube.com/watch?v=oV71hhEpQ20>
 - 2.5 hours though:-)
- Disclaimer: This talk is what I (sf) think, not what the IETF thinks

A Definition

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.

RFC7258/BCP188

was: <http://tools.ietf.org/html/draft-farrell-perpass-attack-06>



Aspects of PM

- Non-targeted, attempts to catch “all” traffic/users, perhaps within some application/protocol/geographic scope
 - Targeted monitoring is different and is sometimes justifiable according to most people, though there are vast differences in opinion as to how to do that acceptably (see also RFC2804)
- PM is neither a classic active nor passive attack
 - Active attack on DNS/HTTP might be used to determine target for passive attack (QUANTUMINSERT)
 - Passive attack on SSH might be used to determine target for active attack (“We target sysadmins”)
- We talk about PM as “an” attack but as we'll see it really involves a whole plethora of different, mostly well known techniques used in a co-ordinated manner

Attack Summary (Briefly)

- Attacks don't respect network layers but those are still useful, so let's look at things roughly from bottom (h/w) up the layers
 - We DO NOT KNOW what exactly has been done
 - Most of the media stories are not specific enough for our purposes
 - Much of this rests on speculation
- Hard to keep an up to date timeline...
 - https://en.wikipedia.org/wiki/Timeline_of_mass_surveillance_disclosures

Low Level Stuff

- Doping h/w from various places in production
 - <http://people.umass.edu/gbecker/BeckerChes13.pdf>
- Dual_EC_DRBG
 - http://csrc.nist.gov/publications/nistbul/itlbul2013_09_supplemental.pdf
 - <http://arstechnica.com/security/2013/09/stop-using-nsa-influence-code-in-our-product-rsa-tells-customers/>
 - <http://dual-ec.org/>
- Other (P)RNG code in kernels etc. often unsafe
 - <https://freedom-to-tinker.com/blog/nadiah/new-research-theres-no-need-panic-over-factorable-keys-just-mind-your-ps-and-qs/>
- Stuxnet/Flame: different motive but similar modus-operandi
 - [http://www.h4ckr.us/library/Documents/ICS_Events/Stuxnet%20Dossier%20\(Symantec\)%20v1.4.pdf](http://www.h4ckr.us/library/Documents/ICS_Events/Stuxnet%20Dossier%20(Symantec)%20v1.4.pdf)
- *Even more paranoid: compromised crypto APIs might leak key bits through use of IVs that appear random but are actually related to application data encryption key*
 - <http://www.metzdowd.com/pipermail/cryptography/2013-September/017571.html>

Mid-Level Stuff

- Tempora
 - <http://www.theguardian.com/uk/2013/jun/21/gchq-cables-secret-world-communications-nsa>
- PRISM
 - <http://www.theguardian.com/world/2013/jun/06/us-tech-giants-nsa-data>
- XKEYSCORE
 - <http://www.theguardian.com/world/2013/jul/31/nsa-top-secret-program-online-data>
- Phone records
 - <http://www.theguardian.com/world/2013/sep/26/nsa-surveillance-intelligence-chiefs-testify-before-senate-live-updates>
- Financial records
 - <http://www.spiegel.de/international/world/spiegel-exclusive-nsa-spies-on-international-bank-transactions-a-922276.html>
- Email records
 - <http://www.theguardian.com/world/2013/jun/27/nsa-data-mining-authorized-obama>

High-ish Level Stuff

- MitM attacks on popular web sites (collaborating PKI?)
 - <https://www.net-security.org/secworld.php?id=15579>
- Directed breaches of ISP/provider n/w in Belgacom
 - <http://www.spiegel.de/international/europe/british-spy-agency-gchq-hacked-belgian-telecoms-firm-a-923406.html>
- Meta-data vs. data
 - <http://www.theguardian.com/technology/interactive/2013/jun/12/what-is-metadata-nsa-surveillance>
- Databases of passwords, secret and private keys, incl. Kerberos!
 - <http://g1.globo.com/fantastico/noticia/2013/09/nsa-documents-show-united-states-spied-brazilian-oil-giant.html> reports on programme, from which some screenshots were posted...
 - <http://leaksource.files.wordpress.com/2013/09/nsa-brazil-4.png> (last bullet says “Results can frequently be verified using Kerberos etc. data”)

High-ish Level Stuff (2, sigh;-)

- “Tor stinks”
 - <http://www.theguardian.com/world/2013/oct/04/nsa-gchq-attack-tor-network-encryption>
- Optic Nerve – GCHQ watching Yahoo video
 - <http://www.theguardian.com/world/2014/feb/27/gchq-nsa-webcam-images-internet-yahoo>
 - “a surprising number of people use webcam conversations to show intimate parts of their body”
- Bahamas “full take” ~30-days of telephone (GSM?) audio
 - <https://firstlook.org/theintercept/article/2014/05/19/data-pirates-caribbean-nsa-recording-every-cell-phone-call-bahamas/>
 - Reportedly involves (ab)use of LI interfaces (cf RFC2804!)

Higher Level Stuff

- “Legal” compulsion
 - <http://www.theguardian.com/world/2013/jun/12/microsoft-twitter-rivals-nsa-requests>
- Corporate collusion
 - (non-objective:-)
 - <http://www.phibetaiota.net/2013/06/rickard-falkvinge-nsa-as-poster-child-for-government-corporate-corruption-collusion-treason/>

Highest Level Stuff

- Damage to reputation of Internet
- Damage to reputation of US & partners and for any government capable of this kind of attack
- The attack is highly unlikely to stop and highly likely to get worse as others join in
- Damage to IETF processes – how many IETFers could you get for a chunk of \$250M per year? (BULLRUN)
 - There is damage to people sponsored by attackers, even if all of those people are entirely good-actors (which afaik they are)
- Nation-state level reactions could damage themselves and the Internet
- All the usual players (or wannabe players) in Internet governance will try benefit from this situation

Silly, silly agencies

- Getting caught is really really bad
 - If you are going to do PM, then not getting caught is an absolute requirement
 - With 1,000,000 people cleared to see something, how secret can it ever be?
- Most country's signals intelligence agencies ostensibly have a dual purpose – to spy/snoop and to help protect their nation, e.g. with “cybersecurity” expertise
 - The former is damaging to the latter it seems

Why is PM (said to be) “good”?

- If targetted surveillance is agreed as a good, then how do you find out who to target?
- Asymmetric threat model, “terrorism”
- The public want it or don't really care (cf. Facebook, gmail)
- Telecommunications has always been monitored to the maximum extent feasible and this is just the same
- We're not looking, and copying/storing is not “collecting”
- Network administrators need to see plaintext to manage their networks according to the legitimate policies for that

The public doesn't care?

- Tapping a government leader's phone is not PM but the common reaction from such people when its known to have happened seems to be outrage (real or feigned)
- How would you react if you knew for sure that recordings of your calls were actually being listened to without permission?
- Related artistic “experiment”:
 - <http://www.theguardian.com/technology/2014/may/16/conver-snitch-covert-surveillance-art-form-twitter>

But PM is, in fact, “bad” (tech)

- Its a technical attack, same as any other, handle it via risk analysis as always – This is the only one with IETF consensus
- Requires bad actor to compromise running systems (can't be only collusive)
- Creates and pumps a market for zero-day attacks
- Zero-day attack “weapons” are wonderful foot-guns because using them often hands their design to your “enemy”
- Probability of collateral damage is high if I secretly inject complex new functions into network
- Data will leak out eventually (deliberately or accidentally)
- Even the end-to-end argument says its bad: can't innovate @ edges if n/w has to support PM

But PM is, in fact, “bad” (non-tech)

- Chilling effects – leads to self-censorship
- Potential for holder of DB to exercise undue influence, think DDR (former East Germany)
- Privacy is a human right (details are of course TBD)
- Work on or for PM is 100% certain to be re-used for repression somewhere in the world
- Everyone has something to hide, some of the time
- Assumption behind PM is guilt-by-association fallacy
- In the end, I didn't give them permission

Guilt-by-association Fallacy

- PM is to an extent motivated by g-b-a
 - State actors consider indirect contact with potential suspect creates a new suspect
- Clearly wrong – while terrorists do contact one another, 99%+ of people contacted even by real terrorists are not terrorists, never mind 2nd or 3rd level effects (note: 99% is guess)
- BUT - We need to not fall into the same trap
 - Some e.g. NSA employees do try to subvert the security of the Internet, but that does not imply anyone who was ever funded by US DoD is now suspect
 - Some of us who think BULLRUN is a gigantic disgrace have fallen for this fallacy
- We need to treat technical proposals and people on face value, and rigorously and openly examine proposals
 - Same and best defence as for handling commercial bias
 - The new information we have will be considered in such evaluation

A tiny bit of IETF history

- 1990's: Crypto wars
 - RFC 1984
- 2000's: wiretap/lawful interception
 - RFC 2804
- Today's equivalent is PM?
 - RFC 7258 (published May 13 2014)

IETF as Part of Problem/Solution

- IETF and broader (re)actions need to note that we are part of the problem as well as part of the solution
 - Whether due to BULLRUN or not, we have made it too hard to use secure protocol variants and we could have done better
- Personally (i.e. not consensus IETF position):
 - My guess is BULLRUN was hugely ineffective, but may have had some minor impact
 - Much more likely: Large complex organisations setting complex requirements and lack of deployment experience meant we made non-optimal engineering choices (over complex mainly, “gold-plating”)
 - We can and will do better though, with confidence for that being largely based on significantly large real deployments – that will force us to do the right thing
- Reminder: BULLRUN == US\$250M/year to make Internet security worse

IETF just one part of the story

- Getting deployments to turn things on is the point and a lot of that does not need any IETF work, since the protocols often just work (or can be forced to “just work”;-)
 - XMPP “Manifesto” from Nov 2013, May 19th 2014 deadline
 - <http://blog.prosody.im/mandatory-encryption-on-xmpp-starts-today/>
 - MTA-MTA STARTTLS
 - <https://www.facebook.com/notes/1453015901605223/>
 - Others...

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
- See also: IETF breakout, for a lot more detail

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
- BoFs at IETF-89: DNS Privacy and TCP encryption
 - Former unthinkable before snowdonia
 - Latter: 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

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