Pervasive Monitoring 3 years in

stephen.farrell@cs.tcd.ie

NORDUnet Conference September 22nd 2016

Talk Outline

- Me: Trinity College Dublin, topics: security/privacy/DTN;
  - IETF security area director (but not speaking for IETF)
  - Expecting wisdom? Vision? Apologies:-)
- A quick reminder of how we got here...
- Good stuff people are doing or have done
- Some harder questions arising
- What should we all be doing about this?
- These slides https://down.dsg.cs.tcd.ie/nordunet/
  - More refs at end

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

Nov 2013 IETF Technical Plenary
(W) https://www.ietf.org/proceedings/88/technical-plenary.html

What have we learned?

- We've mostly learned the unexpectedly broad scope and scale of Intelligence agency snooping
- If there is a way to get at data or meta-data, then they are trying or doing it, including “offensively”
- “Offensive” weapon foot-gun offends common sense
  - Send the bytes of your “offensive” weapon out and they will be used against your customers
- If there is >1 way, they'll try/do them all
- “Collection” fallacy (and others) happily trotted out

What else have we learned?

- Partial timelines:
  - https://en.wikipedia.org/wiki/Global_surveillance_disclosure
  - https://www.theguardian.com/us-news/nsa
- My favourite:
  - https://www.theguardian.com/world/2014/feb/27/gchq-nsa-webcam-images-internet-yahoo
- My most interesting (politically):
  - http://www.scmagazineuk.com/gchq-faces-new-belgacom-hack-allegations/article/388531/
- My most interesting (technically):
  - The short-range radar thing
  - https://en.wikipedia.org/wiki/NSA_ANT_catalog

And more recently...

- Juniper use of EC_DBRG and code changes (Jan 2016)
  - “unauthorized code in the ScreenOS software in our products”
  - https://forums.juniper.net/t5/Security-Incident-Response/Advancing-the-Security-of-Juniper-Products/ba-p/286383
- IKEv1 vulnerabilities in Cisco products courtesy of “shadow brokers” (Sep 2016)
  - “allow an unauthenticated, remote attacker to retrieve memory contents”
  - https://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20160916-ikev1
- UK Parliament Report on “Bulk Powers” (Aug 2016)
  - Claims to justify bulk interception based on mostly secret case studies
  - Does at least express some concerns that it may not be ok for govt to bulk hack lots and lots of devices
  - Ignores the effects of bulk collection for all those of us who are not guilty
  - https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/546925/56730_Cm9326_WEB.PDF

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.

PM is not everything

- PM is far from the only security or privacy issue on which we need to work
  - Spam, malware, DDoS, …
  - But mitigations for PM can also help a lot with other problems
- Hypothesis: If we work to address PM, and prioritise services and mechanisms that mitigate PM and that are also effective against other attacks then we will be doing the “right thing”

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 (July 2013)
  - Tech plenary, major discussion @ IETF-88 (Nov 2013)
  - STRINT workshop before IETF-89 (Feb 2014) [RFC7687]
  - Topic at many meetings/BoFs @ IETF-89 (July 2014)
  - Seeing results from IETF-90 (Nov 2014) onwards...
- Unsurprisingly this is similar to the more broad technical community reaction

IETF work related to PM

- RFC 7258/BCP188 published after major IETF LC debate – sets the basis for further actions
- RFC 7435 defines “Opportunistic Security” - less gold-plating, more deployment
- IAB Statement on Internet Confidentiality: basically: encrypt everything!
- New working groups established:
  - UTA: update BCPs on how to “Use TLS in Applications” - RFC7525
  - DPRIVE: “DNS Privacy”- unthinkable before snowdonia RFC7626
  - TCPINC: “TCP INCreased security”: tcpcrypt proposed two years earlier but rejected
    - Mistakenly, including by me, as ack'd at mic @ IETF-88, bummer
- IAB re-factored security and privacy programme
  - Developed PM threat model document (RFC7624)
- Stuff not going so well
  - Old-RFC privacy/PM review team – go back and see what needs fixing: moribund
  - Endymail email list for discussion of ways IETF can help those working on new e2e interpersonal messaging solutions: hard problem

PM is an Attack RFC 7258

- RFC7258/BCP188 says that all IETF work will consider PM as an attack to be mitigated as part of our normal design processes for all protocol development
  - Note: this does not mean PM is always relevant nor that it's always practical to mitigate PM via protocol mechanisms, but if you can't, you need to be able to say why
- Took ~1000 emails to get rough consensus on that since countering PM is not free
    - Impacts on network management
    - Some folks scared of unreasonable security/privacy nerd dominance

Opportunistic Security RFC 7435

- IETF modus operandi has (in practice) been to define mandatory to implement security that works for higher security environments

  => often hard/expensive to deploy => often not used => cleartext often sent even when better options exist

- Opportunistic Security (OS) aims to evaluate these trade-offs on a connection-by-connection basis, explicitly allowing for e.g. unauthenticated endpoints for confidentiality (open-channel key exchange) as an option that is better than cleartext
- I (personally) hope that this concept is followed very often and is fleshed out to the point where we end up with a new security development approach that is based around OS
  - Not there yet: TLS deprecation of RC4 was interesting because of differing perspectives from web and mail folks about what conclusion to draw when following the OS approach

OS example: Deprecating RC4

- RC4 past sell-by date: agreed by all
- For the web ~15% of https sites were using TLS/RC4 (FF 2014 measurement)
  - When RC4 zapped 99% of those just picked a better option (AES, 3DES)
- SMTP+STARTTLS between MTAs
  - There is a widely deployed MTA that only does RC4, 3DES is buggy and won't work (so I'm told)
  - Zapping RC4 means emails will be sent in clear between MTAs when one is the buggy one
- So – which is better: deprecate RC4 entirely or add this and possibly other caveats?
  - IETF rough consensus was to deprecate entirely, but some mail folks were in the rough
- Interesting example implying conclusion from following OS protocol design pattern will depend on scope
  - OS requires us each to figure out some kind of utility or objective function and where those differ enough, different well meaning folks will reach different conclusions
- It is OK that it is harder to figure out what to do when following the OS approach

IAB Statement

"We recommend that encryption be deployed throughout the protocol stack since there is not a single place within the stack where all kinds of communication can be protected.

The IAB urges protocol designers to design for confidential operation by default. We strongly encourage developers to include encryption in their implementations, and to make them encrypted by default. We similarly encourage network and service operators to deploy encryption where it is not yet deployed, and we urge firewall policy administrators to permit encrypted traffic."

https://www.iab.org/2014/11/14/iab-statement-on-internet-confidentiality/

DNS Privacy

- DNSSEC provides integrity and origin authentication but confidentiality/privacy was never considered a requirement
- Since 2013 that has changed, IETF DPRIVE working group was formed to tackle this issue
- Problem statement set out in RFC 7626
- QNAME minimisation RFC 7816
- TLS/TCP on port 853 between stub and recursive almost done, DTLS/UDP equivalent in the works
- Work on recursive to authoritative starting
- Discussion on DNS/HTTP at IETF97 (Nov2016) will certainly include DNS/HTTPS (way too early to say where that might go)

QUIC

- QUIC is a proposed new transport protocol that runs over UDP and that encrypts a lot
  - https://datatracker.ietf.org/doc/charter-ietf-quic/
  - Goal is the same security properties as TLS1.3/TCP
- QUIC is already deployed to some extent
- Privacy is not the only reason things like QUIC use encryption
  - Cleartext allows middleboxes to see and mess with traffic, which has good and bad aspects
- Will likely provide examples of the tension between privacy and the ability to manage a network mentioned in RFC 7258

Other relevant IETF Things

- TLS 1.3 aiming for better handshake encryption properties and learning from previous TLS problems in various ways
- HTTP/2.0, [RFC7540] the major deployment model for which seems to be to run much much more HTTP traffic over TLS
- Extension to HTTP/2.0 defining opportunistic security way of sending http URI schemed content over TLS
- Negatively: deprecate RC4 [RFC7465] in TLS, SSL3 [RFC7568]
- And since all this is IETF stuff, you can (and please do) join in and help if you're willing and able – that's how to make it better!
  - Even a small amount of good researcher input is hugely valuable (but you need to be able to deal with a noisy environment;-)
- New Curves and deprecating old crypto (CURDLE WG)
- OPENPGP WG updating crypto

Non-IETF Things Relevant to IETF

- Crypto Forum Research Group (CFRG) in Internet Research Task Force (IRTF)
  - Goal: provide venue bridging academic crypto and Internet technical community
  - Curve25519, Curve448 [RFC7748]
  - Chacha20/Poliy1305 AEAD construct [RFC7664]
- IEEE 802 have started work on privacy and are considering e.g. MAC address randomisation
  - Collaborating with IETF
- W3C TAG statement on “Securing the Web”
  - Builds on RFC7258 and IAB statement
  - https://www.w3.org/2001/tag/doc/web-https
- … there are loads more

Longer Term Factors at Play

- Spooks will be spooks (whether govt. or private sector)
- Privacy invasive commerce (legitimate and not)
- Legal accountability mechanisms (courts of various kinds)
- Small+good things can transition to (big+bad, dead or living-dead)
- Badly-informed decision makers/commentators/twits
- Government regulation of business (e.g. Data Protection Agencies)
- Commercial reaction to user privacy requirements (even evil corporate behemoths have many good folks working for 'em)
- NGOs working to enhance privacy (and get attention)
- Constantly refreshed naivety of yet another generation of clean-slaters (producing occasional good ideas)
- Guilt-by-association is a fallacy no matter who makes the error
- Technical privacy enhancing/enforcement mechanisms (when those work)

So what else?

- We've outlined the problem
- We've seen there's work ongoing
  - Most addressing relatively low-hanging fruit in a sense
- Still hard to get agreed/done/finished/deployed
  - Esp. **deployed**, which is REQUIRED for this to be at all useful – Fantasy is of no use here
- But almost are fairly obvious things to do
  - Encrypt more, do more/better security, yay!
- So how about a hard problem or two?

Hard Problem #1
Targeted Monitoring or so-called “LI”

- Many people against PM, but ok with targeted monitoring (e.g. wiretap under warrant), sometimes called “lawful interception” (LI)
- Requirements for LI from 19th century, (IMO)
- LI deployments can be, and have been, used for PM
- 21st century Internet is not amenable to those requirements

  => tension, tussle, cases like Apple/FBI
- The “right thing”™ : openly evaluate requirements taking equal account of Internet-scale interoperability, security and privacy alongside law enforcement needs
  - I doubt it'll happen, for various reasons
- More on this: “Requirements Analysis Required – otherwise Targeted Monitoring Enables Pervasive Monitoring,” March 2016 IEEE Computer Magazine
  - http://www.qmags.com/R/?i=3110a8&e=3193789&doi=60830608&uk=2FE1161B1640F7E013144D46111630BBC52FF14A391.htm
  - https://down.dsg.cs.tcd.ie/cpus/ (preprint/extended version of above)

Hard Problem #2
Collaborative Corporate Collection

- Advertising-driven or captured-user walled-gardens are inimical to privacy; always will be; ensnare even educated users; almost ubiquitous; lead corporations to “collect it all,” even more than TLAs
- Non-solutions include: user education, FOSS, legislation, cleverer crypto
  - Someone has to pay to keep the lights on
- We, the technical community, are bad actors here (along with advertisers)
- How to balance toxicity of data vs. data-mining benefits?
- How to convince browsers/web-sites to not enable ad-trackers?
- I have no idea how to get to substantive improvement
- But if we don't, PM will continue and accelerate
  - TLAs can convince or coerce companies into collaborating

What to do? (1)

- Discuss the issue openly in whatever fora are relevant for you, (it's not a secret:-)
- Consider privacy issues in your networks and the data you make available
  - Avoid logging potentially sensitive data if you can
  - Find and delete old crap you no longer need
  - That means more work! But you should do it

What to do? (2)

- Turn on crypto – ciphertext should be base assumption for new things
  - Consider the OS approach to make that easier
- Don't use new stuff without considering privacy implications
  - Data minimisation will save you some later leaks
- Help with better implementations
  - https://cryptech.is/ and similar
- Encourage target diversity - Don't all use the same services all the time
  - Even if you're not a huge population, you can start trends

What to do? (3)

- Don't demand the impossible (and do nothing in the meantime)!
  - Encourage clean-slate work, but don't imagine it can all be deployed now – and only deployed things help
- Agitate (if that's your kind of thing:-)
- Consider privacy trade-offs when deploying e.g. IDS, anti-spam or malware detection technologies
- Go and be responsible network operators and take the broader implications of your work into account before, while and after doing it

Summary

- IETF has consensus PM is an attack (RFC7258) and is working that problem, as are others
- 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

Thanks

Offline questions welcome too
stephen.farrell@cs.tcd.ie

These slides:
https://down.dsg.cs.tcd.ie/nordunet/

More References (1)

- General IETF stuff:
- https://www.ietf.org/
- https://www.ietf.org/newcomers.html
- Working group details for WG <foo>:
  - https://tools.ietf.org/wg/<foo> - links to charter, docs, mail archive etc
  - Suggested <foo> values:
    - tls, dprive, tcpinc, httpbis, uta

References (2)

- Relevant IETF non-wg lists:
  - All of them (loads): https://www.ietf.org/list/nonwg.html
  - Perpass – triage list for PM related stuff:
    - https://www.ietf.org/mailman/listinfo/perpass
  - Security area list (saag)
    - https://www.ietf.org/mailman/listinfo/perpass
  - Possible e2e interpersonal messaging discussion
    - https://www.ietf.org/mailman/listinfo/endymail
  - General privacy discussion
    - https://www.ietf.org/mailman/listinfo/ietf-privacy
- IRTF:
  - https://www.irtf.org/
  - IRTF Crypto Research Forum Goup: https://irtf.org/cfrg

References (3)

- Videos (ISOC hint:-)
  - IETF youtube stuff in general
  - https://www.youtube.com/user/ietf/videos?sort=p&view=0&flow=grid
- Nov'13 IETF technical plenary video
  - https://www.youtube.com/watch?v=oV71hhEpQ20
- Dan york videos 5 minute summaries of IETF meetings
  - There are loads but these are about PM
  - https://www.youtube.com/watch?v=HG54EsHYKr0
  - https://www.youtube.com/watch?v=fbjs_6Mz-6s
- STRINT workshop (RFC7687)
  - Has all 66 position papers
  - https://www.w3.org/2014/strint/

References (4)

- IEEE Internet Computing “soapbox” column on why PM is bad:
  - http://www.computer.org/csdl/mags/ic/2014/04/mic2014040004.pdf
- Some Internet drafts not referenced above:
  - PM Threat model
    - https://tools.ietf.org/html/draft-iab-privsec-confidentiality-threat
  - DNS Privacy problem statement`
    - RFC7626
  - “Modern” TLS best current practices
    - RFC7525