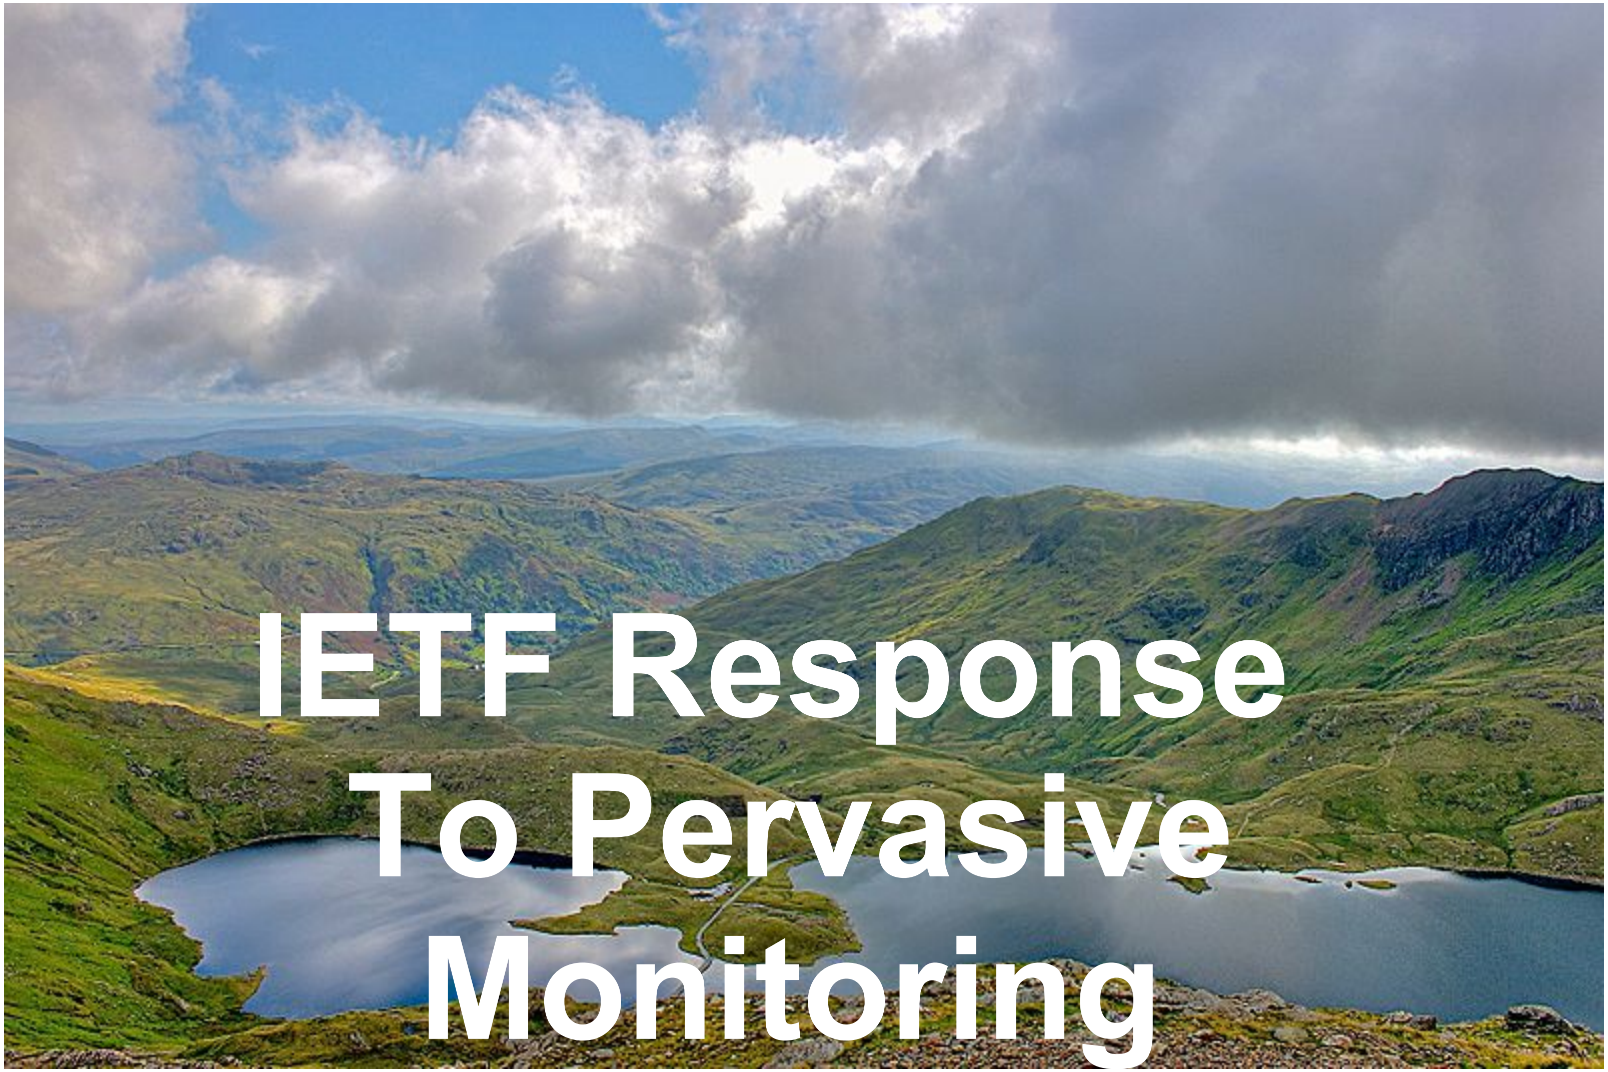

IETF Response To Pervasive Monitoring

stephen.farrell@cs.tcd.ie
November 7th 2013

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
- A purely technical response will not “solve the problem” but we should treat an attack as we usually do and try mitigate it
- Will we have rough consensus on the above?
  - Be good to know, personally I *think* we will.

There are things we can do

- There are technical things we can do that might significantly affect the cost of pervasive monitoring and that can improve security and privacy generally
- Some of those are short-term “point” changes (or BCPs), others may take time to be agreed, mature and get deployed
- If we're serious about tackling the problem, some changes may affect IETF processes, long-held positions, deployments or business models
  - Mantatory-to-Implement (MTI) vs. more-than-MTI
  - Confidentiality vs packet inspection
  - Anonymity/pseudonymity vs authent/law enforcement/ advertising

So let's do them

- There is a time element to some of this – it could be that we can get some changes made or started more easily while the news is fresh
- Equally, being seen not to act in this situation could inflict more damage
- We should do and be seen to be doing as much as we can to counter this attack, and now is the time – publicity counts and the attackers haven't just crossed a line, they've moved it
- NOTE: “we” in all the above means the IETF ***and*** each of us outside the IETF

Trusted Computing Base →Dodgy Computing Base

- Crypto-implementation-worries++
  - Some but nowhere near all paranoia justified: RNGs, side-channels, …
  - Affects IETF protocol design, implementation & deployment
  - Some significantly: DNSSEC, RPKI, others...
  - Turns assumptions about crypto APIs on their head a bit
- Idea: a set of IETFers and others help organise & fund a team of developers to make **a high-assurance open-source h/w and s/w crypto engine platform**
  - Only limited knowledge and funds needed to make small numbers of devices from COTS components
  - Meet the crypto requirements of a set of interesting IETF protocols and applications that use those
  - Think PKCS#11 + key-handling-ceremonies
- **Not** an IETF activity, but...
  - IETF & others generating use-cases and requirements
  - Core development team not an IETF WG nor DT
  - High risk, (high-assurance open-h/w?) but pretty cool if it works
- Interested in helping with use-cases, reqs?
  - Thursday 1145-1300 in Plaza B, bring your own lunch; A bit more info: https://cryptech.is/

Dramatis Personae for ACT-I: "get started"

- Bart Preneel
- Jari Arkko
- **Leif Johansson**
- **Linus Nordberg**
- **Lucy Lynch**
- Lynn StAmour
- Olaf Kolkman
- **Randy Bush**
- Russ Housley
- Sean Turner
- Stephen Farrell
- Steve Bellovin
- Tero Kivinen

IETF Actions ("easy")

- First, and most important: Discuss the situation and what to do openly
  - perpass list mainly for triage of issues, not intended as a WG
  - Discussion at various IETF-88 sessions: Appsarea, HTTPbis, Perpass BoF,...
  - IAB workshop on Internet hardening just before IETF-89 (London)
    - Call for participation/position-papers in a couple of weeks
    - With some help from EU FP7 STREWS project
    - Maybe spin up IRTF RG around then?
- Second, work the problem, some obvious bits:
  - Threat analyses, draft-trammell-perpass-ppa
  - Deployable PFS ciphersuite BCPs for TLS and for foo-over-TLS (foo=smtp, imap, xmpp, …)
  - Encourage operational changes, e.g. more local IXPs, more direct fibre...
  - Good problem statement text from various folks

IETF Actions (trickier)

- For a couple, a start has been made:
  - Privacy BCP, draft-cooper-ietf-privacy-requirements
  - More-than-MTI – get closer to “secure by default” discussed, but no clear outcome yet
- Some relevant issues from hard → very hard:
  - The impact of turning on TLS everywhere for the web
    - And/or tcpcrypt for TCP. And/or IPsec.
  - The practicality of end-to-end security for, e.g. email, IM, VoIP
  - Could WebRTC and IoT make it all worse? Or better?
  - Fingerprinting and traffic analysis from RF->Application layer
    - IP addresses as personally identifying information? Location traces?
  - Corporate cloudy privacy-busting will be affected if we succeed

Conclusions

- It is an attack.
- It is a new scale of attack
- The right response is for the IETF is to develop technical mitigations, as before and as usual
  - Goal: make it significantly more expensive for a bad actor
- There are things we can and should do
  - Do them! And openly, starting now.
- For things where we're not sure: work the problem
  - What are **you** doing about this?

Backup Stuff

References

- Perpass list info:
  - https://www.ietf.org/mailman/listinfo/perpass
- List of relevant sessions at IETF-88:
  - http://down.dsg.cs.tcd.ie/misc/perpass-sessions.txt
- Rough list of useful material from perpass list:
  - http://down.dsg.cs.tcd.ie/misc/perpass.txt
- Technical overview of attacks
  - Overview from perpass BoF session (next session)
  - See meeting materials: https://datatracker.ietf.org/meeting/88/materials.html#perpass
- The above are more lists of lists and not direct references, but having you doing cut'n'paste is easier than me typing:-)

Significantly More Expensive

- “Significantly more expensive” means something like at least 2^80+ more work compared to simply recording plaintext, which is quite doable with current crypto protocols and deployments in many cases
  - That is significant even for these bad actors
  - Yes, 2^128 is the target, but there may be corner cases that take a while to go away
- That is also likely to force them towards more active attacks, which are riskier for the attacker and detectable or preventable when we have good key management
  - E.g. using Certificate Transparency (RFC6962) or some other “big DB of public keys” approach

More-than-MTI

- MTI has gotten us some very good things but still too many RFC 6919 cases and/or we mess up security because we don't really mean it in v1.0 of a protocol
  - Interop events that just don't even try the “secure” version
- More-than-MTI aims to get security turned-on/used by default
  - Likely less than Mandatory-To-Use
  - Perhaps: “MUST offer/use security by default. MAY allow a way to turn off security via local configuration.”
  - But more work on that is definitely needed
- Arguments:
  - For: More-than-MTI could get usable security in v1.0
  - Against: That's policy and just won't work for enough protocols