Network Working Group
Internet-Draft
Intended status: BCP
Expires: June 6, 2014

S. Farrell
Trinity College Dublin
H. Tschofenig

December 3, 2013

Pervasive Monitoring is an Attack
draft-farrell-perpass-attack-02.txt

Abstract

The IETF has consensus that pervasive monitoring is a technical attack that should be mitigated in the design of IETF protocols, where possible.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 6, 2014.

1. It's an Attack

[[Note (to be removed before publication): This draft is written as if IETF consensus has been established for the text.]]

The technical plenary of IETF 88 [IETF88Plenary] discussed pervasive monitoring. Such pervasive surveillance requires the monitoring party to take actions that are indistinguishable from an attack on Internet communications. Participants at that meeting therefore expressed strong agreement that this was an attack that should be mitigated where possible via the design of protocols that make pervasive monitoring significantly more expensive or infeasible. This Best Current Practice (BCP) formally documents that consensus, having been through an IETF last call.

For the purposes of this BCP "pervasive monitoring" means very widespread privacy-invasive gathering of protocol artefacts including application content, protocol meta-data (such as headers) or keys used to secure protocols. Other forms of traffic analysis, for example, correlation, timing or measuring packet sizes can also be used for pervasive monitoring.

The term "attack" is used here in a technical sense that differs somewhat from common English usage. In common English usage, an "attack" is an aggressive action perpetrated by an opponent, intended to enforce the opponent's will on the attacked party. In the Internet, the term is used to refer to a behavior that subverts the intent of a communicator without the agreement of the parties to the communication. It may change the content of the communication, record the content of the communication, or through correlation with other communication events or attempts, reveal information the communicator did not intend to be revealed. It may also have other effects that similarly subvert the intent of a communicator. [RFC4949] contains a more complete definition for the term "attack" as used here. We also use the term in the singular here, even though pervasive monitoring in reality may require a multi-faceted set of coordinated attacks.

In particular, the term "attack", when used technically, implies nothing about the motivation of the actor mounting the attack. The motivation behind pervasive monitoring is not relevant for this document, but can range from non-targeted nation-state surveillance, to legal but privacy-unfriendly purposes by commercial enterprises, to illegal purposes by criminals. The same techniques can be used regardless of motivation and we cannot defend against the most nefarious actors while allowing monitoring by other actors no matter how benevolent some might consider them to be. As technology advances, techniques that were once only available to extremely well

funded actors become more widely accessible. Mitigating this attack is therefore a protection against wider usage of pervasive monitoring.

2. And We Will Continue to Mitigate the Attack

The IETF also has consensus to, where possible, work to mitigate the technical parts of the pervasive monitoring attack, in just the same way as we continually do for these and any other protocol vulnerability.

There are various ways in which IETF protocols can be designed in order to mitigate pervasive monitoring, but those will change over time as mitigation and attack techniques develop and so are not described here. This BCP simply records the consensus to design protocols so as to mitigate the attack, where possible.

More limited-scope monitoring to assist with network management that is required in order to operate the network or an application is not considered pervasive monitoring. There is though a clear potential for such limited monitoring mechanisms to be abused as part of pervasive monitoring, so this tension needs careful consideration in protocol design. Making networks unmanageable in order to mitigate pervasive monitoring would not be an acceptable outcome. But equally, ignoring pervasive monitoring in designing network management mechanisms would go against the consensus documented in this BCP. An appropriate balance will likely emerge over time as real instances of this tension are considered.

It is also important to note that the term "mitigation" is a technical term that does not necessarily imply an ability to completely prevent or thwart an attack. As in common English usage, this term is used here in the sense of "make less severe, serious, or painful." [OED] In this case, designing IETF protocols to mitigate pervasive monitoring will almost certainly not completely prevent such from happening, but can significantly increase the cost of such monitoring or force what was covert monitoring to be more overt or more likely to be detected (possibly later) via other means. And even where the IETF has done this work well and that has been fully deployed, there will still be some privacy-relevant information that will inevitably be disclosed by protocols.

While RFC 4949 does not contain a definition for the term mitigation, we prefer it here to the term countermeasure which is defined in RFC 4949 since the latter term is more often understood to mean putting in place a more fully effective mitigation of an attack.

Finally, we note that the IETF is not equipped to tackle the non-technical aspects of mitigating pervasive surveillance. Others need to step forward to tackle those if pervasive monitoring is to be fully addressed.

3. Process Note

In the past, architectural statements of this sort, e.g., [RFC1984] and [RFC2804] have been published as joint products of the IESG and IAB. However, since those documents were published, the IETF and IAB have separated their publication "streams" as described in [RFC4844] and [RFC5741]. This document was initiated by both the IESG and IAB, but it is published as an IETF-stream consensus document, having garnered the consensus of the IETF as approved by the IESG.

4. Security Considerations

This BCP is entirely about privacy. More information about the relationship between security and privacy threats can be found in [RFC6973]. Section 5.1.1 of [RFC6973] specifically addresses surveillance as a combined security-privacy threat.

5. IANA Considerations

There are none. We hope the RFC editor deletes this section before publication.

6. Acknowledgements

We would like to thank the participants of the IETF 88 technical plenary for their feedback. Thanks in particular to the following for useful suggestions that resulted in changes to this text: Jari Arkko, Fred Baker, Marc Blanchet, Brian Carpenter, Benoit Claise, Spencer Dawkins, Adrian Farrel, Russ Housley, Joel Jaeggli, Eliot Lear, Barry Leiba, Ted Lemon, Erik Nordmark, Pete Resnick, Peter Saint-Andre, and Sean Turner. Additionally, we would like to thank all those who contributed suggestions on how to improve Internet security and privacy or who commented on this on various IETF mailing lists, such as the ietf@ietf.org and the perpass@ietf.org lists.

7. Informative References

[IETF88Plenary]

IETF, "IETF 88 Plenary Meeting Materials", URL: https://datatracker.ietf.org/meeting/88/materials.html, Nov 2013.

[OED] Stevenson, Angus, "Oxford Dictionary of English", Oxford University Press http://www.oxforddictionaries.com/definition/english/mitigate, 2010.

[RFC1984] IAB, IESG, Carpenter, B., and F. Baker, "IAB and IESG Statement on Cryptographic Technology and the Internet", RFC 1984, August 1996.

[RFC2804] IAB and IESG, "IETF Policy on Wiretapping", RFC 2804, May 2000.

[RFC4844] Daigle, L. and Internet Architecture Board, "The RFC Series and RFC Editor", RFC 4844, July 2007.

[RFC4949] Shirey, R., "Internet Security Glossary, Version 2", RFC 4949, August 2007.

[RFC5741] Daigle, L., Kolkman, O., and IAB, "RFC Streams, Headers, and Boilerplates", RFC 5741, December 2009.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, July 2013.

Authors' Addresses

Stephen Farrell
Trinity College Dublin
Dublin, 2
Ireland

Phone: +353-1-896-2354
Email: stephen.farrell@cs.tcd.ie

Hannes Tschofenig
Brussels,
Belgium

Phone:
Email: hannes.tschofenig@gmx.net