

Network Working Group  
Internet-Draft  
Intended status: BCP  
Expires: June 23, 2014

S. Farrell  
Trinity College Dublin  
H. Tschofenig  
December 20, 2013

**Pervasive Monitoring is an Attack  
draft-farrell-perpass-attack-03.txt**

Abstract

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 23, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## **1. Pervasive Monitoring is Indistinguishable from an Attack**

The technical plenary of the November 2013 IETF meeting [[IETF88Plenary](#)] discussed pervasive monitoring (or surveillance) which 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, see [[RFC2026](#)] [Section 5](#)) formally documents that consensus.

For the purposes of this document "pervasive monitoring" means often covert and very widespread intrusive gathering of protocol artefacts including application content, protocol meta-data such as headers, or cryptographic keys used to secure protocols. Active or passive wiretaps, traffic analysis, correlation, timing or measuring packet sizes can also be used as part of 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. Here, 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, 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." 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. The IETF will work to Mitigate Pervasive Monitoring

"Mitigation" is a technical term that does not imply an ability to completely prevent or thwart an attack. Protocols that mitigate pervasive monitoring will not prevent the attack, but can significantly change the threat. (See the diagram on page 24 of [RFC 4949](#) for how the terms attack and threat are related.) This can significantly increase the cost of attacking, force what was covert to be overt, or make the attack more likely to be detected, possibly later.

IETF standards already provide mechanisms to protect Internet communications and there are guidelines [[RFC3552](#)] for applying these in protocol design. But those generally do not consider pervasive monitoring, the confidentiality of protocol meta-data, countering traffic analysis nor data minimisation. [[RFC6973](#)] And in all cases, there will remain some privacy-relevant information that is inevitably disclosed by protocols.

It is nonetheless timely to revisit the security and privacy properties of our standards. The IETF will work to mitigate the technical parts of the pervasive monitoring threat, just as we do for other protocol vulnerabilities. The ways in which IETF protocols mitigate pervasive monitoring will change over time as mitigation and attack techniques evolve and so are not described here.

Those developing IETF specifications need to be able to describe how they have considered pervasive monitoring, and, if the attack is relevant to the work to be published, be able to justify related design decisions. This does not mean a new "pervasive monitoring considerations" section is needed in IETF documentation. It means that, if asked, there needs to be a good answer to the question "is pervasive monitoring relevant to this work and if so how has it been addressed?"

While pervasive monitoring is an attack, other forms of monitoring can be beneficial and not part of any attack, e.g. network management functions monitor packets or flows, anti-spam mechanisms see mail message content and monitoring can even be a mitigation for pervasive monitoring in the case of Certificate Transparency. [[RFC6962](#)] There is though a clear potential for monitoring mechanisms to be abused for pervasive monitoring, so this tension needs careful consideration in protocol design. Making networks unmanageable to mitigate pervasive monitoring is not an acceptable outcome, but ignoring pervasive monitoring 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.

Finally, the IETF, as a standards development organisation, does not control the implementation or deployment of our specifications (though IETF participants do develop many implementations), nor does the IETF specify all layers of the protocol stack. And the non-technical (e.g. legal and political) aspects of mitigating pervasive monitoring are outside of the scope of the IETF. The broader Internet community will need to step forward to tackle pervasive monitoring, if it 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 Internet Engineering Steering Group (IESG) and the Internet Architecture Board (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 is published as an IETF-stream consensus document, in order to ensure that it properly reflects the consensus of the IETF community as a whole.

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

### **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 or comments: Jari Arkko, Fred Baker, Marc Blanchet, Tim Bray, Scott Brim, Randy Bush, Brian Carpenter, Benoit Claise, Alissa Cooper, Dave Crocker, Spencer Dawkins, Avri Doria, Wesley Eddy, Adrian Farrel, Joseph Lorenzo Hall, Ted Hardie, Sam

Hartmann, Bjoern Hoehrmann, Phillip Hallam-Baker, Russ Housley, Joel Jaeggli, Stephen Kent, Eliot Lear, Barry Leiba, Ted Lemon, Subrahmanian Moonesamy, Erik Nordmark, Pete Resnick, Peter Saint-Andre, Andrew Sullivan, Sean Turner, and Stefan Winter. 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](mailto:ietf@ietf.org) and the [perpass@ietf.org](mailto: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.
- [RFC1984] IAB, IESG, Carpenter, B., and F. Baker, "IAB and IESG Statement on Cryptographic Technology and the Internet", [RFC 1984](#), August 1996.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [RFC2804] IAB and IESG, "IETF Policy on Wiretapping", [RFC 2804](#), May 2000.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), July 2003.
- [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.
- [RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", [RFC 6962](#), June 2013.
- [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](mailto:stephen.farrell@cs.tcd.ie)

Hannes Tschofenig  
Brussels,  
Belgium

Phone:  
Email: [hannes.tschofenig@gmx.net](mailto:hannes.tschofenig@gmx.net)