                 Opportunistic Security: some protection most of the time
                      draft-dukhovni-opportunistic-security-00

Abstract

   This memo defines the term "opportunistic security". In contrast to
   the established approach of delivering strong protection some of the
   time, opportunistic security strives to deliver at least some
   protection most of the time. The primary goal is therefore broad
   interoperability, with security policy tailored to the capabilities
   of peer systems.

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 January 3, 2015.

Copyright Notice

Table of Contents

1. Introduction

   Historically, Internet security protocols have prioritized strong
   protection for peers capable and motivated to absorb the associated
   costs. Since strong protection is not universally applicable, while
   communications traffic was sometimes strongly secured, more typically
   it was not protected at all. The fact that most traffic is
   unprotected facilitates nation-state pervasive monitoring (PM
   [I-D.farrell-perpass-attack]) by making it cost-effective (or at
   least not cost-prohibitive). Indiscriminate collection of
   communications traffic would be substantially less attractive if
   security protocols were designed to operate at a range of protection
   levels with encrypted transmission accessible to most if not all
   peers, and stronger security still available where required by policy
   or opportunistically negotiated.

   Encryption is easy, but key management is difficult. Key management
   at Internet scale remains an incompletely solved problem. The PKIX
   ([RFC5280]) key management model introduces costs that not all peers
   are willing to bear and is also not sufficient to secure
   communications when the peer reference identity is obtained
   indirectly over an insecure channel or communicating parties don't
   agree on a mutually trusted certification authority (CA). DNSSEC is
   not at this time sufficiently widely adopted to make DANE a viable
   alternative at scale. Trust on first use (TOFU) key management
   models (as with saved SSH fingerprints and various certificate
   pinning approaches) don't protect initial contact and require user
   intervention when key continuity fails.

   Without Internet-scale key management, authentication is often not
   possible. When protocols only offer the options of strongly-
   authenticated secure channels or else no security, most traffic gets
   no security protection. Therefore, in order to make encryption more
   ubiquitous, authentication needs to be optional. When strongly
   authenticated communication is not possible, unauthenticated
   encryption is still substantially stronger than cleartext.
   Opportunistic security encourages peers to employ as much security as
   possible, without falling back to unnecessarily weak options. In

   particular, opportunistic security encourages unauthenticated
   encryption when authentication is not an option.

[2](#). **Opportunistic Security Design Philosophy**

   Interoperate to maximize deployment: The primary goal of designs
      that feature opportunistic security is to be able to communicate
      with any reasonably configured peer. If many peers are only
      capable of cleartext, then it is acceptable to fall back to
      cleartext when encryption is not possible. If authentication is
      only possible for some peers, then it is acceptable to
      authenticate only those peers and not the rest. Interoperability
      must be possible without bilateral coordination. Applications
      employing opportunistic security need to be deployable at Internet
      scale, with each peer independently configured to meet its own
      security needs (within the practical bounds of the application
      protocol). Opportunistic security must not get in the way of the
      peers communicating if neither end is misconfigured.

   Maximize security peer by peer: Subject to the above large-scale
      interoperability goal, opportunistic security strives to maximize
      security based on the capabilities of the peer (or peers). For
      some opportunistic security protocols the maximal protection
      possible may be just unauthenticated encryption. For others,
      greater security may be an option, and opportunistic security may
      at times (in partial conflict with the interoperability goal)
      refuse to continue with peers where higher security is expected,
      but for some reason not achieved. The conditions under which
      connections fail should generally be limited to operational errors
      at one or the other peer or an active attack, so that well-
      maintained systems rarely encounter problems in normal use of
      opportunistic security.

   Encrypt by default: An opportunistic security protocol MUST
      interoperably achieve at least unauthenticated encryption between
      peer systems that don't explicitly disable this capability. Over
      time, as peer software is updated to support opportunistic
      security, only legacy systems or a minority of systems where
      encryption is disabled should be communicating in cleartext.
      Whenever possible, opportunistic security should employ Perfect
      Forward Secrecy (PFS) to make recovery of previously sent keys and
      plaintext computationally expensive even after disclosure of long-
      term keys.

   No misrepresentation of security: Unauthenticated communication or
      use of authentication that is vulnerable to MiTM attacks is not
      represented as strong security. Where strong security is

required, opportunistic security is not a substitute, though the
underlying mechanisms may in some cases be very similar.

3. Terminology

The following definitions are derived from the Internet Security
Glossary [RFC4949], where applicable.

Perfect Forward Secrecy (PFS): For a key management protocol, the
   property that compromise of long-term keys does not compromise
   session/traffic/content keys that are derived from or distributed
   using the long-term keys.

Man-in-the-Middle attack (MiTM): A form of active wiretapping attack
   in which an attacker intercepts and may selectively modify
   communicated data to masquerade as one of the entities involved in
   a communication. Masquerading enables the MiTM to violate the
   confidentiality and/or the integrity of communicated data passing
   through it.

Trust on First Use (TOFU): In a protocol, TOFU typically consists of
   accepting an asserted identity, without authenticating that
   assertion, and caching a key or credential associated with the
   identity. Subsequent communication using the cached key/
   credential is secure against a MiTM attack, if such an attack did
   not succeed during the (vulnerable) initial communication or if
   the MiTM is not present for all subsequent communications. The
   SSH protocol makes use of TOFU. The phrase "leap of faith" (LoF)
   is sometimes used as a synonym.

Unauthenticated Encryption: Encryption using a key management
   technique that enables unauthenticated communication between
   parties. The communication may be 1-way or 2-way unauthenticated.
   If 1-way, the initiator (client) or the target (server) may be
   anonymous.

4. Security Considerations

Though opportunistic security potentially supports transmission in
cleartext, unauthenticated encryption, or other protection levels
short of the strongest potentially applicable, the effective security
for users is increased, not reduced. Provided strong security is not
required by policy or securely negotiated, nothing is lost by
allowing weaker protection levels, indeed opportunistic security is
strictly stronger than the alternative of providing no security
services when maximal security is not applicable.

5. Acknowledgements

   I would like to thank Steve Kent. Some of the text in this document
   is based on his earlier draft.

6. References

   [I-D.farrell-perpass-attack]
             Farrell, S. and H. Tschofenig, "Pervasive Monitoring is an
             Attack", draft-farrell-perpass-attack-06 (work in
             progress), February 2014.

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

   [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S.,
             Housley, R., and W. Polk, "Internet X.509 Public Key
             Infrastructure Certificate and Certificate Revocation List
             (CRL) Profile", RFC 5280, May 2008.

Author's Address

   Viktor Dukhovni
   Two Sigma

   Email: ietf-dane@dukhovni.org