

TLS
Internet-Draft
Intended status: Informational
Expires: January 17, 2013

D. McGrew
D. Wing
Cisco
Y. Nir
Checkpoint
P. Gladstone
Independent
July 16, 2012

TLS Proxy Server Extension
draft-mcgrew-tls-proxy-server-01

Abstract

Transport Layer Security (TLS) is commonly used to protect HTTP and other protocols; it provides encrypted and authenticated conversations between a client and a server. In some scenarios, two TLS sessions are used, so that a third device can participate in the protected communication. In these cases, separate TLS sessions are run between the client and the middle device, on one side, and the middle device and the server on the other side. This provides the needed security, as long as the client, server, and middle device use appropriate and consistent security policies. However, this last part is problematic; how can the middle device know if a client trusts a server? At present, TLS provides no mechanism to coordinate policies, and there is no convenient way to do so.

This note defines a TLS extension that allows a TLS server to provide a TLS client with all of information about the other TLS server (or servers) that are participating in the application layer traffic that the client needs to make a well-informed access control decision. This empowers the client to reject TLS sessions that include servers that it does not trust.

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 17, 2013.

Copyright Notice

Copyright (c) 2012 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.

Table of Contents

1.	Introduction	4
1.1.	Requirements Language	5
2.	Motivation	5
3.	Operation	7
3.1.	ProxyInfoExtension	9
3.2.	Client Certificates	11
4.	Discussion	14
5.	IANA Considerations	14
6.	Security Considerations	15
7.	References	17
7.1.	Normative References	17
7.2.	Informative References	17
	Authors' Addresses	18

1. Introduction

Transport Layer Security (TLS) [RFC 5246](#) [[RFC5246](#)] is commonly used to protect HTTP [[RFC2616](#)] as described in [[RFC2818](#)]. In some scenarios an HTTP proxy is used, for instance, to allow caching, to provide anonymity to a client, or to provide security by using an application-layer firewall to inspect the HTTP traffic on behalf of the client (e.g. to protect it against cross-site scripting attacks). A TLS session cannot protect traffic between the client and server when an HTTP proxy is present. It is possible to have separate TLS sessions between the client and the proxy, on one side, and the proxy and the server on the other side, as show in Figure 1 . This technique provides the appropriate cryptographic security (see below for a discussion of why some other alternatives are less attractive). But there is a problem: the presence of the proxy removes the client's knowledge about the server. Without this knowledge, the client has no way to decide what trust, if any, it should have in the server. This is most problematic when the client trusts multiple different servers for different applications, or trusts servers from different domains.



A proxied HTTPS session, with two independent TLS sessions.

Figure 1

A further issue is that the client cannot determine the security level of the TLS session between the proxy and the server. For instance, a client can negotiate a high security ciphersuite between itself and the proxy, but it will have no way of knowing what ciphersuite is in use between the client and the server, which could be using the obsolete 56-bit Data Encryption Standard (DES) cipher.

Another point of difficulty is the fact that there can be multiple proxies on a particular path. To solve the security issues introduced by TLS proxies in a way that is generally applicable, it is necessary to accommodate scenarios involving multiple proxies.

A separate issue is the provisioning of the proxy with information about what servers (or rather, which certificates) should be trusted. If the laptop has installed certificates that are specific to its

organization or to a particular domain, how can the proxy know to trust these certificates on behalf of the laptop?

We propose a solution in this note, by describing a TLS extension that can be used by a proxy to provide information to a TLS client about the TLS server. When this extension is used, the client is well informed about the proxy as well as the server, and can make a knowledgeable access control decision about the server, using the same processes that it uses when the proxy is not present. The data in the extension are signed by the proxy in order to bind the information about the server to a particular session between the client and the proxy. When there are multiple proxies, the client is informed about all of them. This extension also works for DTLS.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

2. Motivation

One motivation for the proxy extension is the benefit that it provides a client with a clear and explicit indication whenever a device is attempting to act in the role of a TLS proxy server. This allows TLS clients to reject sessions that include proxy servers that it does not trust. As an extreme example, a client could be configured to reject all sessions that involve proxy servers, in order to enforce a conservative security policy.

The following motivating example describes a typical situation with a TLS proxy, as in Figure 1. A laptop trusts the server A for a particular banking application, and trusts server B for a social media application, and can authenticate both servers by using standard PKIX certificate checking [[RFC5280](#)] and locally stored root certificates. Or rather, the client trusts a set of root certificates, and uses them to authenticate the TLS servers that it connects with. The laptop also trusts the proxy, and has a certificate by which it can authenticate the proxy. When making a connection directly with B, the laptop can authenticate the server as being trusted (that is, the server's public key appears in a certificate that has been signed by the appropriate trusted certificate authority), and it can also check the authorizations of that server (that is, B is authorized to provide the social media service, but not any other services such as banking). If the web traffic from the laptop goes through an HTTP proxy, then the proxy will need to know that it should trust both A and B to act as TLS

servers. Assuming that it does have this knowledge, it will proxy TLS connections from both A and B. However, when the client attempts to establish an HTTPS connection to A through the proxy, it has no way of knowing what security checks the proxy has applied to the connection between the proxy and A. The client cannot tell whether the trusted certificate that it associates with A was used on the connection between the proxy and A. The inability of the client to be confident of the identity of the actual server forces the client to trust all TLS servers indiscriminately.

This obstacle could be overcome by pushing the client's policy (that is, information about what servers it trusts for what applications) onto the proxy, so that the proxy can make well-informed decisions on behalf of the client. However, this alternative has significant drawbacks: it requires that the proxy obtain and store a significant amount of information about each client, and it requires the construction of a syntax by which the client's policy can be expressed and understood. In contrast, our solution moves the information about the server to the client, which does not require the communication or storage of any security policy between the client and server.

TLS proxies without this extension also defeat some recent security mechanisms that other groups have added to TLS:

- o Extended Validation certificates ([\[CAB.EV\]](#)) are certificates that contain a special indication of the actual organization for which the certificate had been issued. Only a minority of public CAs are authorized to issue such certificates. By hiding the actual server certificate, proxies do not allow the browser to determine the EV status of the server certificate. This loses the visual indication that browsers typically show when EV certificates are present.
- o The DANE protocol ([\[I-D.ietf-dane-protocol\]](#)) stores hashes of server certificates in the DNS. Clients are expected to verify that the certificate that the server uses is the one specified in the DNS record. This fails if the certificate is one generated by the proxy.
- o Key Pinning ([\[I-D.ietf-websec-key-pinning\]](#)) is an alternative way to make sure the client is connecting to the correct server. Unlike DANE, that stores the certificate hash in the DNS, key pinning sends that hash in an HTTP header. Still, a client that moves behind a proxy will have stored the correct hash, but will get in TLS a certificate that does not match that hash, causing the connection to fail, unless that feature is disabled behind a proxy.

All these cases are solved if the client has access to the actual certificate sent by the server. This is provided by this extension.

3. Operation

In this note, a TLS proxy is a device that acts as a TLS server in one session and acts as a TLS client in another session, and passes all of the data from one session to the other, possibly modifying it in the process. That is, it is a non-transparent proxy, in the terms of [[RFC2616](#)].



A TLS session with a single proxy (top) and a TLS session with two proxies (bottom).

Figure 2

The essential idea is as follows. When a TLS proxy is contacted by a client, it does not respond to the client until it completes a TLS session with the server. It then sends the client an assertion about the server and the session, signed with the same private key that it uses in its role as the TLS proxy server. When the client receives this assertion, it checks the data in the assertion to determine whether or not it trusts the server. The assertion is carried in a ProxyInfoExtension, which is defined below.

This extension carries all of the information that is available to a TLS client about a TLS server; thus the client can use existing authorization checking processes. The client will need to verify the hostname and/or address, and check to see if the certificate has been revoked. The client authenticates the proxy server as usual during the TLS session. This ensures that the client trusts the proxy, and because of the signature on the assertion, it should trust the server certificate carried in the assertion. The proxy need not perform any checking on the server certificate, because this check is done by the client. Of course, by completing a TLS exchange with the server, the proxy verifies that the server holds the private key associated with that certificate.

If the client attempts session resumption, and the proxy can resume the session, then it must attempt to resume the session with the

server. If the server resumes the session, then the proxy must resume the session with the client. If the proxy cannot resume the session with the client, then it MUST begin a fresh session with the server, and not resume the session with the client. If the server does not resume, the proxy MUST not resume the session with the client.

Because there may be more than one proxy in any path, the TLS extension carries a list of assertions.

On receiving a ClientHello from the client, the proxy:

1. Checks for a ProxyInfoExtension in the ClientHello; if there is no such extension, then the following steps cannot be performed and are omitted,
2. Establishes a TLS session with the server (session #2 in Figure 1); a ProxyInfoExtension is included in that session,
3. Constructs a ProxyInfo structure by populating it with information about the server and the current session with that server; if the server sends back a ProxyInfoExtension, then the ProxyInfo structure is included as the next_proxy_info,
4. Signs the ProxyInfo structure with the private key corresponding to the server certificate it uses in session #1,
5. Completes the session with the client (session #1 in Figure 1) and provides the ProxyInfoExtension in that session,

The proxy MAY

Perform revocation checking on the certificate chain of the server in session #2, and indicate that it has done this in the extension by setting performed_revocation_checking to "true".

Note that the entity acting in the role of the server in session #2 could be a proxy, but in the above it is referred to as a server because that is the role that it performs in that TLS session.

When TLS is used in HTTPS, the proxy MUST perform the Server Identity checks described in [Section 3.1 of \[RFC2818\]](#).

The normal operation of the proxy is to accept the (extended) ClientHello from the client and then send a ClientHello to the server. It is recommended that the TLS Proxy support commonly deployed TLS extensions (as defined in [\[RFC4366\]](#) et al). Any TLS extensions present on the original ClientHello MUST be examined and

either ignored, processed or forwarded (possibly after modification) to the TLS server as part of the new ClientHello.

The client:

1. Includes a ProxyInfoExtension in the ClientHello message,
2. Checks for ProxyInfoExtension in the ServerHello message; if there is no such extension, then the TLS processing continues as usual; otherwise,
3. Processes the ProxyInfo extension by checking the validity of the digitally-signed struct, then performing the usual server authentication and authorization checking on the server_certificate_list in the ProxyInfo,
4. Checks the revocation_checking_performed flag in the ProxyInfo; if it is "false", then the client SHOULD perform revocation checking on the server_certificate_list,
5. Checks the ProxyInfoFlag in the next_proxy_info field; if it is not_empty, then the client returns to step 3 and performs that processing on the next_proxy_info.

In order to maintain backwards compatibility for existing TLS clients, the TLS proxies MUST (by default) perform certificate validation for the certificates that they receive from the server. The use of the ProxyInfoExtension in the extended ClientHello is an indication by the client to request the alternate processing defined by this note. In particular, if this extension is present in the extended ClientHello, then the TLS proxy should not use its own private key to dynamically generate a certificate.

The proxy will relay the data between the client and peer data connections. End-to-end flow control is maintained by the relay process: if the relay process is no longer able to write data to the destination of the relayed data, the relay process stops reading data from the source.

3.1. ProxyInfoExtension

The syntax of the ProxyInfo extension is as follows:


```
struct {
    ProtocolVersion    tls_version;
    CipherSuite        cipher_suite;
    CompresseionMethod compression_method;
} ConnectionSecurityParameters;

enum { client_to_proxy, proxy_to_client,
        proxy_to_server, server_to_proxy } ProxyInfoFlag;

struct {
    select (ProxyInfoFlag) {
        case client_to_proxy:
            /* zero length body */
        case proxy_to_client:
            digitally-signed struct {
                ConnectionSecurityParameters connection_parameters;
                ASN.1Cert server_certificate_list<0..2^8-1>;
                Random server_random; /* server-side server random */
                Boolean revocation_checking_performed;
                ProxyInfo next_proxy_info;
            } SignedProxyInfo;
        case proxy_to_server:
            digitally-signed struct {
                ASN.1Cert proxy_certificate_list<0..2^8-1>;
                ProxyInfo next_proxy_info;
            } SignedProxyInfo2;
        case server_to_proxy:
            /* zero length body */
    }
} ProxyInfo;

struct {
    ProxyInfo proxy_info;
} ProxyInfoExtension;
```

The ProxyInfo structure is defined recursively, so that the signature of each proxy authenticates the information provided by the proxies that follow it on the path. The ProxyInfo contains the ProxyInfoFlag, which indicates whether or not the ProxyInfo is empty (in which case it contains no other fields) or not (in which case it contains a SignedProxyInfo structure). The SignedProxyInfo structure is signed with the public key that the proxy uses in its role as the TLS server (in session #1). That structure contains the connection_parameters that describe the security of session #2, and the certificate chain of the server from session #2 in the server_certificate_list. If the proxy has performed revocation checking on that certificate chain, it indicates this by setting the

Boolean `revocation_checking_performed`. If the server in session #2 was actually a proxy itself, and it provides a `ProxyInfo` struct, then that struct is included in the `next_proxy_info` field. Otherwise, the `next_proxy_info` field contains an empty `ProxyInfo`.

```
enum {  
    /* ... */  
    proxy_info(TBD1), (65535)  
} ExtensionType;
```

[3.2.](#) Client Certificates

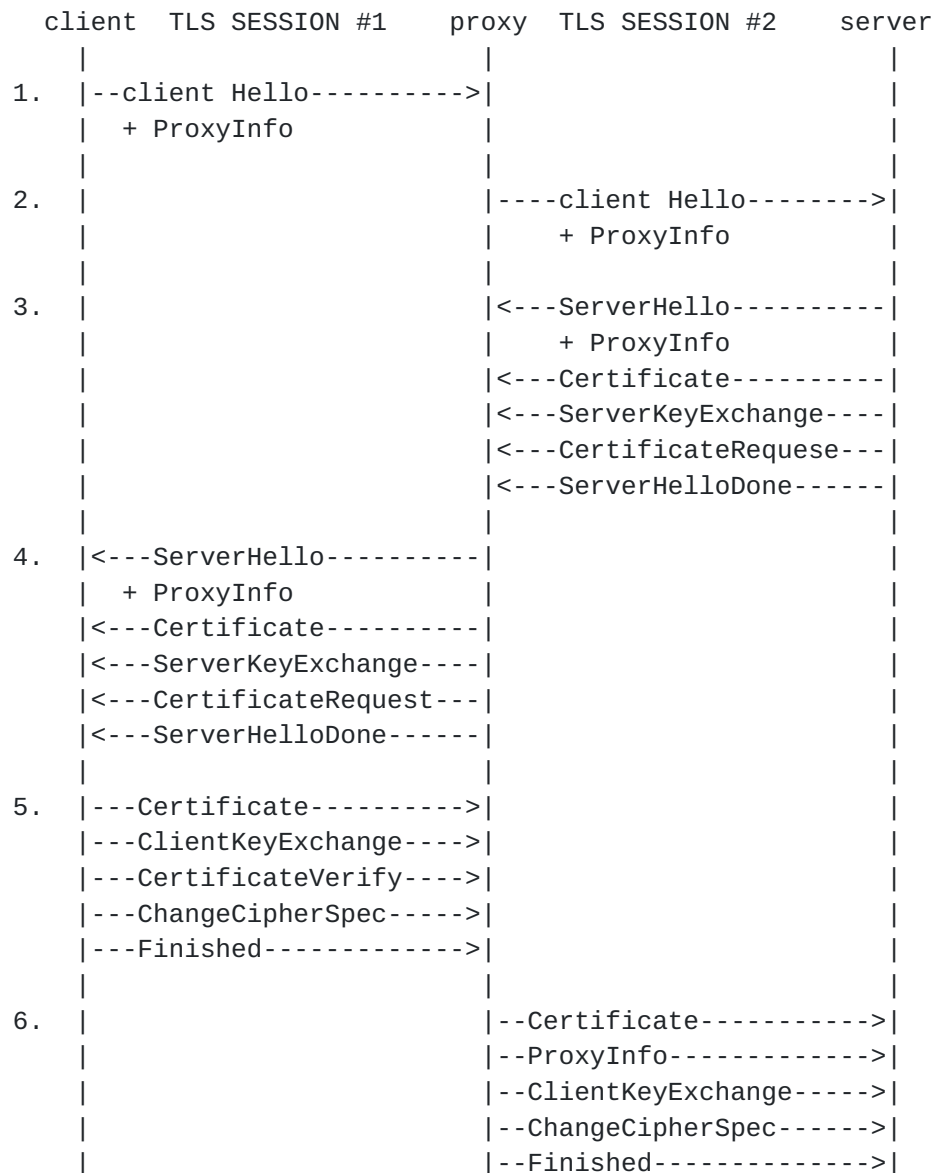
The mechanism described above supports server authentication and requires updates to the TLS client (e.g., the web browser) to support TLS proxying. Some TLS connections use client certificates (sometimes called mutual authentication). When TLS client certificates are used, the TLS server must be updated to support TLS proxy. This section describes the handshake changes beyond those described in the previous section.

On the first TLS connection to a server, the TLS proxy does not know if the TLS server will request a client certificate (that is, if the server will send a `CertificateRequest` in its TLS handshake). So, the TLS proxy first establishes a TCP and TLS connection to the server (as described in previous section) and when the proxy sees the TLS `CertificateRequest` from the server, it starts a new TCP connection. Then the following steps occur:

1. The client sends `ClientHello` with `ProxyInfo`. If the `ClientHello` does not include `ProxyInfo`, processing stops and the following steps cannot be performed.
2. The proxy, after determining the server is going to request the client's certificate (see proceeding paragraph), initiates a TLS session (TLS Session #2) to the server. The `ProxyInfo` in TLS Session #2 contains two `ProxyInfo` datastructures -- the `ProxyInfo` from Session #1 and the second containing the proxy's own certificate.
3. The server acknowledges it supports `ProxyInfo`, by including `ProxyInfo` in its `ServerHello` with a nonce it wants signed by the client, along with its `CertificateRequest`.
4. On Session #1, the proxy sends a `ServerHello`, and includes the `ProxyInfo` from the previous step.

5. On Session #1, the client sends its Certificate, its CertificateVerify, ChangeCipherSpec, and Finished.
6. The proxy validates the CertificateVerify message. If it fails validation, both of the TLS sessions are abandoned.
7. On Session #2, the proxy sends the client's certificate (obtained in the previous step) in ProxyInfo, the client's CertificateVerify (obtained in the previous step), its own ClientKeyExchange, ChangeCipherSpec, and Finished.
8. The TLS server verifies the TLS client initiated the TLS communication by using the data in ProxyInfo, the nonce it sent in Step 3, and CertificateVerify.

A message flow diagram of the TLS proxy, after the proxy has determined the TLS server requests a client certificate



The ProxyCertVerify is carried from the client to the server (that is, it is un-modified by the proxy), which allows the TLS server to check that certificate. Because there are two separate TLS, Session #1 and Session #2, ProxyCertVerify cannot utilize a signature over the TLS handshake messages (as with the classic CertificateVerify). Instead, the new ProxyCertVerify message contains two signatures which provide a similar (but not identical) function. One signature is over TLS SESSION #2's ServerHello (which is conveyed into Session #1's ProxyInfo), signed using the client's private key. The second signature is over TLS SESSION #2's entire handshake, signed using the private key of the TLS proxy.

4. Discussion

The ProxyInfo extension could contain information about the checking that the proxy performed on the server and its certificate. For example, if the DNS name of the server matched the subjectAltName, this fact could be indicated. It may be desirable to enumerate the ways in which the server can match its certificate, to allow the proxy to indicate to the client which of those ways was positive for a particular server.

A potential issue with the ProxyInfo extension is that it can be large, because the certificate chains that it carries can be large. Roughly speaking, the amount of certificate data presented to the client is proportional to the number of proxies on the path. It is undesirable to require that so much data be sent, but on the other hand, the client does need all of the data in order to make a well-informed access control decision. It appears that the data is the minimum required, in the sense that removing any of the data would make it impossible for the client to assess the security of the entire path.

The proxy is required to do the authentication checking on the signatures created by the server, but not the authorization checking or revocation checking. The responsibility for authorization checking is not put onto the proxy because it does not know the security policy of the client; in particular, the proxy does not know which servers the client trusts for which applications. The responsibility for revocation checking is not put onto the proxy because that process is better left to the client. The client can perform revocation checking on all of the certificate lists for all of the proxies and the server in parallel, whereas if each proxy performed the revocation checking, those processes would necessarily be serial. Since revocation checking can take a significant amount of time, the serial approach could add a significant amount of latency to the TLS session, and potentially trigger retransmissions. The parallel approach not only reduces the overall latency, but it moves it outside of the client's retransmission timer for the ClientHello message.

The ProxyInfo extension could convey the IP address of the server, or other network layer information such as the DNS name. However, it is not clear that this information is needed, so it was not included.

5. IANA Considerations

This document requests IANA to update its registry of TLS extension types to assign an entry, referred herein as proxy_info, with the

number TBD1.

6. Security Considerations

In a situation with a client, server, and a middle device that all need to participate in an encrypted and authenticated session, the appropriate security goals are to

- preserve the security properties of the cryptographic protocols in use,

- make the client aware of both the middle device and the server, able to authenticate the both of those devices, and able to check that both of the devices are trusted/authorized to act in their roles,

- allow the client to make access control decisions that are as well-informed as when only the client and server are present.

The idea in this note meets these goals.

We briefly describe some alternative approaches that do not meet these security goals.

First, we consider the proliferation of private keys. In order to allow one device to act as a proxy for a server, the private key of the server could be shared with the proxy. This practice may be workable when there is a one-to-one correspondence between proxies and servers, but it substantially increases the security risk. If a proxy contains multiple private keys, it becomes an attractive target for an attacker.

Second, we consider the session-key proliferation approach in which there is only a single TLS session, negotiated between the client and server, and the proxy participates in the session because either the client or the server has passed the secret session keys to the proxy (using some secure channel). One attempt at this approach is in the now abandoned [[I-D.nir-tls-keyshare](#)]. If the proxy is completely passive, and it only decrypts traffic from the TLS session and never modifies the data in that session, then this method can be secure. However, if the proxy rewrites the data inside the session, or originates messages, then the security of the TLS protocol will be undermined. Message authentication can be subverted because an attacker can intercept a message sent by the server, and forward it on to the client, bypassing the proxy. By interleaving messages sent by the proxy with ones sent by the server, an attacker can potentially confuse a client, and can certainly cause a denial of

service. Confidentiality may be undermined as well; if RC4, AES-GCM, or AES-CCM is in use, information about the plaintext will be leaked due to keystream reuse. Session-key proliferation is not secure when the proxy needs to edit the session. Most proxies do need to edit the session, and we regard it as potentially hazardous to construct a TLS proxy along these lines. Suppose that such a proxy were implemented because it was anticipated that the application proxy would be read-only, but then a future revision to the application protocol or the goals of the application proxy made it necessary to have the proxy edit the application session. If the session-key proliferation approach had been used, the implementer would be in the awkward position of having to choose between the costly path of implementing a completely new approach that preserved security, and the quick and inexpensive path of allowing the proxy to edit the session to the detriment of the security of the application.

With the ProxyInfo extension, there is no protection against the proxy lying about the security characteristics of the onward connection, unless client certificates are used. However, in any proxying scenario, it is necessary to trust the proxy, just as a client must trust the server. For instance, any proxy (not just one using the ProxyInfo extension) could choose to forward the plaintext from the session to untrusted third parties, and violate the trust of the client. It is the responsibility of the client to decide whether or not a particular device should be trusted to act in the role of proxy. The ProxyInfo proposal has the benefit of making the presence of the proxy obvious, and allows the client to refuse to deal with untrusted proxies.

Many clients use password-based authentication within a TLS tunnel. When a proxy is present, it can learn plaintext passwords, and it can gain the information needed to perform offline dictionary attacks against authentication systems that use challenge-response methods. This is a highly undesirable aspect of TLS proxying. The ProxyInfo extension does nothing to directly help this issue. However, it does indirectly improve the situation, because it empowers the client with information that enables it to reject proxies and servers that it should not trust. Since the TLS authentication (including both sever and proxy authentication) takes place before the password-based authentication, the client can protect itself by rejecting sessions with inappropriate proxies, or inappropriate servers on the path beyond the proxy.

In theory, the cryptographic proxying scenario could be considered as multiparty security negotiation and key establishment. It may be interesting to investigate such ideas because they can allow for more equitable negotiation of session parameters, and additional security properties. This note focuses on compatibility with existing

specifications and implementations, so these considerations are beyond its scope.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2818] Rescorla, E., "HTTP Over TLS", [RFC 2818](#), May 2000.
- [RFC4366] Blake-Wilson, S., Nystrom, M., Hopwood, D., Mikkelsen, J., and T. Wright, "Transport Layer Security (TLS) Extensions", [RFC 4366](#), April 2006.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), August 2008.
- [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.

7.2. Informative References

- [CAB.EV] CA/Browser Forum, "GUIDELINES for the PROCESSING of EXTENDED VALIDATION CERTIFICATES", CAB GUIDELINES for the PROCESSING of EXTENDED VALIDATION CERTIFICATES, January 2009.
- [I-D.ietf-dane-protocol]
Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", [draft-ietf-dane-protocol-23](#) (work in progress), June 2012.
- [I-D.ietf-websec-key-pinning]
Evans, C. and C. Palmer, "Public Key Pinning Extension for HTTP", [draft-ietf-websec-key-pinning-02](#) (work in progress), June 2012.
- [I-D.nir-tls-keyshare]
Nir, Y., "A Method for Sharing Record Protocol Keys with a Middlebox in TLS", [draft-nir-tls-keyshare-02](#) (work in progress), March 2012.

[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H.,
Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext
Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.

Authors' Addresses

David A. McGrew
Cisco Systems, Inc.
510 McCarthy Blvd.
Milpitas, CA 95035
US

Phone: (408) 525 8651
Email: mcgrew@cisco.com
URI: <http://www.mindspring.com/~dmcgrew/dam.htm>

Dan Wing
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134
USA

Email: dwing@cisco.com

Yoav
Check Point Software Technologies Ltd.
5 Ha'Solelim Street
Tel Aviv, 67897
Israel

Email: ynir@checkpoint.com

Philip Gladstone
Independent

