

TLS Working Group

INTERNET-DRAFT

June 20, 2001

Expires December 20, 2001

Intended Category: Standards track

Simon Blake-Wilson, Certicom

Magnus Nystrom, RSA Security

David Hopwood, Independent Consultant

Jan Mikkelsen, Transactionware

Tim Wright, Vodafone

TLS Extensions

<[draft-ietf-tls-extensions-00.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#). Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or made obsolete 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.

The list of current Internet-Drafts may be found at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories may be found at <http://www.ietf.org/shadow.html>.

Abstract

This document describes extensions that may be used to add functionality to TLS. It provides both generic extension mechanisms for the TLS handshake client and server hellos, and specific extensions using these generic mechanisms.

The extensions may be used by TLS clients and servers. The extensions are backwards compatible - communication is possible between TLS 1.0 clients that support the extensions and TLS 1.0 servers that do not support the extensions, and vice versa.

This document is based on discussions within the TLS working group and within the WAP security group.

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](#) [[KEYWORDS](#)].

Please send comments on this document to the TLS mailing list.

Table of Contents

1. Introduction	2
2. General Extension Mechanisms	4
2.1. Extended Client Hello	4
2.2. Extended Server Hello	5
2.3. Hello Extensions	5
2.4. Extensions to the handshake protocol	6
3. Specific Extensions	7
3.1. DNS Name Indication	7
3.2. Maximum Record Size Negotiation	8
3.3. Client Certificate URLs	9
3.4. Trusted CA Indication	11
3.5. Truncated HMAC	12
3.6. OCSP Status Request	13
4. Error alerts	14
5. Procedure for Defining New Extensions	16
6. Security Considerations	17
6.1. Security of dns_name	17
6.2. Security of max_record_size	17
6.3. Security of client_certificate_url	18
6.4. Security of trusted_ca_keys	18
6.5. Security of truncated_hmac	19
6.6. Security of status_request	19
7. Internationalisation Considerations	19
8. Intellectual Property Rights	20
9. Acknowledgments	20
10. References	20
11. Authors' Addresses	21

[1. Introduction](#)

This document describes extensions that may be used to add functionality to TLS. It provides both generic extension mechanisms for the TLS handshake client and server hellos, and specific extensions using these generic mechanisms.

TLS is now used in an increasing variety of operational environments - many of which were not envisioned when the original design criteria for TLS were determined. The extensions introduced in this document are designed to enable TLS to operate as effectively as possible in new environments like wireless networks.

Wireless environments often suffer from a number of constraints not commonly present in wired environments - these constraints may include bandwidth limitations, computational power limitations,

memory limitations, and battery life limitations.

The extensions described here focus on extending the functionality provided by the TLS protocol message formats. Other issues, such as the addition of new cipher suites, are deferred.

Specifically, the extensions described in this document are designed to:

- Allow TLS clients to provide to the TLS server the DNS name of the server they are contacting. This functionality is desirable to facilitate secure connections to servers which host multiple 'virtual' servers at a single underlying network address.
- Allow TLS clients and servers to negotiate the maximum record size to be sent. This functionality is desirable as a result of memory constraints among some clients, and bandwidth constraints among some access networks.
- Allow TLS clients and servers to negotiate the use of client certificate URLs. This functionality is desirable in order to conserve memory on constrained clients.
- Allow TLS clients to indicate to TLS servers which CA root keys they possess. This functionality is desirable in order to prevent multiple handshake failures involving TLS clients which are only able to store a small number of CA root keys due to memory limitations.
- Allow TLS clients and servers to negotiate the use of truncated MACs. This functionality is desirable in order to conserve bandwidth in constrained access networks.
- Allow TLS clients and servers to negotiate that the server sends the client an OCSP [[OCSP](#)] response during a TLS handshake. This functionality is desirable in order to avoid sending a CRL over a constrained access network and therefore save bandwidth.

In order to support the extensions above, general extension mechanisms for the client hello message and the server hello message are introduced.

The extensions described in this document may be used by TLS 1.0 clients and TLS 1.0 servers. The extensions are designed to be backwards compatible - meaning that TLS 1.0 clients that support the extensions can talk to TLS 1.0 servers that do not support the extensions, and vice versa.

Backwards compatibility is primarily achieved via two considerations:

- Clients typically request the use of extensions via the extended client hello message described in [Section 2.1](#). TLS 1.0 [[TLS](#)] requires servers to "accept" extended client hello messages, even if the server does not "understand" the extension.
- For the specific extensions described here, no mandatory server response is required when clients request extended functionality.

Note however, that although backwards compatibility is supported, some constrained clients may be forced to reject communications with servers that do not support the extensions as a result of the limited capabilities of such clients.

The remainder of this document is organized as follows. [Section 2](#) describes general extension mechanisms for the client hello and server hello handshake messages. [Section 3](#) describes specific extensions to TLS 1.0. [Section 4](#) describes new error alerts for use with the TLS extensions. The final sections of the document address IPR, security considerations, acknowledgements, and references.

[2. General Extension Mechanisms](#)

This section presents general extension mechanisms for the TLS handshake client hello and server hello messages.

These general extension mechanisms are necessary in order to enable clients and servers to negotiate whether to use specific extensions, and how to use specific extensions. The extension formats described are based on [MAILING LIST].

[Section 2.1](#) specifies the extended client hello message format, [Section 2.2](#) specifies the extended server hello message format, and [Section 2.3](#) describes the actual extension format used with the extended client and server hellos.

2.1. Extended Client Hello

The extended client hello message format MAY be sent in place of the client hello message format when clients wish to request extended functionality from servers. The extended client hello message format is:

```
struct {  
    ProtocolVersion client_version;  
    Random random;  
    SessionID session_id;
```



```
CipherSuite cipher_suites<2..2^16-1>;
CompressionMethod compression_methods<1..2^8-1>;
Extension client_hello_extension_list<0..2^16-1>;
} ClientHello;
```

Here the new "client_hello_extension_list" field contains a list of extensions. The actual "Extension" format is defined in [Section 2.3](#).

In the event that clients request additional functionality using the extended client hello, and this functionality is not supplied by the server, clients MAY abort the handshake.

Note that TLS, [Section 7.4.1.2](#), allows additional information to be added to the client hello message. Thus the use of the extended client hello defined above should not "break" existing TLS 1.0 servers.

2.2. Extended Server Hello

The extended server hello message format MAY be sent in place of the server hello message when the client has requested extended functionality via the extended client hello message specified in [Section 2.1](#). The extended server hello message format is:

```
struct {
    ProtocolVersion server_version;
    Random random;
    SessionID session_id;
    CipherSuite cipher_suite;
    CompressionMethod compression_method;
    Extension server_hello_extension_list<0..2^16-1>;
} ServerHello;
```

Here the new "server_hello_extension_list" field contains a list of extensions. The actual "Extension" format is defined in [Section 2.3](#).

Note that the extended server hello message is only sent in response to an extended client hello message. This prevents the possibility that the extended server hello message could "break" existing TLS 1.0 clients.

2.3. Hello Extensions

The extension format for extended client hellos and extended server hellos is:

```
struct {
    ExtensionType extensionType;
```



```
    opaque extension_data<0..2^16-1>;  
  } Extension;
```

Here:

- "extensionType" identifies the particular extension type.
- "extension_data" contains information specific to the particular extension type.

The extension types defined in this document are:

```
enum {  
    dns_name(0), max_record_size(1), client_certificate_url(2),  
    trusted_ca_keys(3), truncated_hmac(4), status_request(5),  
    (65535)  
} ExtensionType;
```

Note that for all extension types (including those defined in future), the extension type should appear in the extended server hello only if the same extension type appeared in the corresponding client hello. Thus clients **MUST** abort the handshake if they receive an extension type in the extended server hello that they did not request in the associated (extended) client hello.

Also note that when multiple extensions are present in the extended client hello or the extended server hello, the extensions must appear in the order identified in "ExtensionType". Thus clients and servers **MUST** abort the handshake if they receive an extended hello message in which the extensions are not in the correct order.

Finally note that all the extensions defined in this document are relevant only when a session is initiated. Extensions appearing in client and server hellos sent during session resumption **MUST** be ignored, and the extension functionality negotiated during session initiation applied to the resumed session.

2.4. Extensions to the handshake protocol

This document suggests the use of two new handshake messages, "CertificateURL" and "CertificateAndOCSPResponse". These messages are described in [Section 3.3](#) and [Section 3.6](#), respectively. The new handshake message structure therefore becomes:

```
enum {  
    hello_request(0), client_hello(1), server_hello(2),  
    certificate(11), server_key_exchange (12),  
    certificate_request(13), server_hello_done(14),
```



```
    certificate_verify(15), client_key_exchange(16),
    finished(20), certificate_url(21), ocsp_response(22), (255)
} HandshakeType;

struct {
    HandshakeType msg_type;      /* handshake type */
    uint24 length;              /* bytes in message */
    select (HandshakeType) {
        case hello_request:      HelloRequest;
        case client_hello:       ClientHello;
        case server_hello:       ServerHello;
        case certificate:         Certificate;
        case server_key_exchange: ServerKeyExchange;
        case certificate_request: CertificateRequest;
        case server_hello_done:   ServerHelloDone;
        case certificate_verify:  CertificateVerify;
        case client_key_exchange: ClientKeyExchange;
        case finished:            Finished;
        case certificate_url:      CertificateURL;
        case ocsp_response:        CertificateAndOCSPResponse;
    } body;
} Handshake;
```

3. Specific Extensions

This section describes the specific TLS extensions specified in this document.

Note that any messages associated with these extensions that are sent during the TLS handshake MUST be included in the hash calculations involved in "Finished" messages.

[Section 3.1](#) describes the extension of TLS to allow clients to indicate which server they are contacting. [Section 3.2](#) describes the extension to provide maximum record size negotiation. [Section 3.3](#) describes the extension to allow client certificate URLs. [Section 3.4](#) describes the extension to allow clients to indicate which CA root keys they possess. [Section 3.5](#) describes the extension to allow the use of truncated HMAC. [Section 3.6](#) describes the extension to support integration of OCSP into TLS handshakes.

3.1. DNS Name Indication

TLS does not provide a mechanism for clients to tell servers the DNS name of the server they are contacting. It may be desirable for clients to provide this information to facilitate secure connections to servers which host multiple 'virtual' servers at a single underlying network address.

In order to provide the server DNS name, clients MAY include an extension of type "dns_name" in the (extended) client hello. The "extension_data" field of this extension shall contain:

```
opaque DNSName<1..216-1>;
```

"DNSName" contains the fully qualified domain name of the server, as understood by the client. The domain name is represented as a byte string using UTF-8 encoding [UTF8], without a trailing dot. (Note that the use of UTF-8 here for encoding internationalized domain names is independent of the choice of encoding for these names in the DNS protocol. The latter has yet to be decided by the IETF International Domain Name Working Group.)

Literal IPv4 and IPv6 addresses are not permitted.

It is RECOMMENDED that clients include an extension of type "DNSName" in the client hello whenever they locate a server by its domain name.

Servers that receive a client hello containing the "dns_name" extension, MAY use the information contained in the extension to guide their selection of an appropriate certificate to return to the client. In this event, the server shall include an extension of type "dns_name" in the (extended) server hello. The "extension_data" field of this extension shall be empty.

If the server understood the client hello extension but does not recognize the DNS name as belonging to a domain it is responsible for, it should send an unrecognized_domain alert (which may or may not be fatal).

3.2. Maximum Record Size Negotiation

TLS specifies a fixed maximum record size of 2¹⁴ bytes. It may be desirable for constrained clients to negotiate a smaller maximum record size due to memory limitations or bandwidth limitations.

In order to negotiate smaller maximum record sizes, clients MAY include an extension of type "max_record_size" in the (extended) client hello. The "extension_data" field of this extension shall contain:

```
enum{
    29(1), 210(2), 211(3), 212(4), (255)
} MaxRecordSize;
```

whose value is the desired maximum record size. The allowed values for this field are: 2⁹, 2¹⁰, 2¹¹, and 2¹².

Servers that receive an extended client hello containing a "max_record_size" extension, MAY accept the requested maximum record size by including an extension of type "max_record_size" in the (extended) server hello. The "extension_data" field of this extension shall contain "MaxRecordSize" whose value is the same as the requested maximum record size.

Servers receiving maximum record size negotiation requests for values other than the allowed values MUST abort the handshake with an "illegal_parameter" alert. Similarly, clients receiving maximum record size negotiation responses that differ from the size they requested MUST also abort the handshake with an "illegal_parameter" alert.

Once a maximum record size other than 2^{14} has been successfully negotiated, the client and server MUST immediately begin fragmenting messages (including handshake messages), to ensure that no message larger than the negotiated size is sent. Note that TLS already requires clients and servers to support fragmentation of handshake messages.

The negotiated size applies for the duration of the session including session resumptions.

The negotiated size limits the input that the record layer may process without fragmentation. Note that the output of the record layer may be larger. For example, if the negotiated size is $2^9=512$, then for currently defined cipher suites (those defined in [\[TLS\]](#), [\[KERB\]](#), and planned AES ciphersuites), the record layer output can be at most 793 bytes: 5 bytes of headers, 512 bytes of application data, 256 bytes of padding, and 20 bytes of MAC. This means that in this event a TLS record layer peer receiving a TLS record layer message larger than 793 bytes may discard the message and output an error without decrypting the message. The exact error message sent will depend on the size of the received message - either "record_overflow" if the message is longer than $2^{14}+2048$ bytes, or "decryption_failed" otherwise.

3.3. Client Certificate URLs

TLS specifies that when client authentication is performed, client certificates are sent by clients to servers during the TLS handshake. It may be desirable for constrained clients to send certificate URLs in place of certificates so that they do not need to store their certificates and can therefore save memory.

In order to negotiate to send certificate URLs to a server, clients MAY include an extension of type "client_certificate_url" in the

(extended) client hello. The "extension_data" field of this extension shall be empty.

(Note that it is necessary to negotiate use of client certificate URLs in order to avoid "breaking" existing TLS 1.0 servers.)

Servers that receive an extended client hello containing a "client_certificate_url" extension, MAY indicate that they are willing to accept certificate URLs by including an extension of type "client_certificate_url" in the (extended) server hello. The "extension_data" field of this extension shall be empty.

After negotiation of the use of client certificate URLs has been successfully completed (by exchanging hellos including "client_certificate_url" extensions), clients MAY send a "CertificateURL" message in place of a "Certificate" message:

```
struct {  
    URLAndHash url_and_hash_list<1..2^16-1>;  
} CertificateURL;  
  
struct {  
    opaque URL<1..2^16-1>;  
    CertHash certificate_hash;  
} URLAndHash;  
  
opaque CertHash[20];
```

Here "url_and_hash_list" contains a sequence of URLs and SHA-1 hashes, each URL referring to a single, DER-encoded X.509v3 certificate, and the SHA-1 hash of that certificate. The URLs should occur in the list in the same order that the corresponding certificates appear in the certificate chain.

Servers receiving "CertificateURL" shall attempt to retrieve the client's certificate chain from the URLs, and then process the certificate chain as usual. HTTP SHOULD be used to retrieve the certificate chain from the URLs, and MUST be supported by servers supporting this extension.

Servers MUST check that the SHA-1 hash of any certificates retrieved from a CertificateURL matches the given hash, and then process the certificate chain as usual.

If any retrieved certificate does not have the correct SHA-1 hash, the server MUST abort the handshake with a bad_certificate alert.

Note that clients may choose to send either "Certificate" or

"CertificateURL" after successfully negotiating the option to send certificate URLs. The option to send a certificate is included to provide flexibility to clients possessing multiple certificates.

If a server encounters an unreasonable delay in obtaining certificates in a given CertificateURL, it SHOULD time out and signal a "certificate_unobtainable" error alert.

3.4. Trusted CA Indication

Constrained clients which, due to memory limitations, possess only a small number of CA root keys, may wish to indicate to servers which root keys they possess, in order to avoid repeated handshake failures.

In order to indicate which CA root keys they possess, clients MAY include an extension of type "trusted_ca_keys" in the (extended) client hello. The "extension_data" field of this extension shall contain "TrustedAuthorities" where:

```
struct {
    TrustedAuthority trusted_authorities_list<0..2^16-1>;
} TrustedAuthorities;

struct {
    IdentifierType identifier_type;
    select (identifier_type) {
        case null: struct {};
        case key_hash_sha: KeyHash;
        case x509_name: DistinguishedName;
        case cert_hash: CertHash;
    } Identifier;
} TrustedAuthority;

enum { null(0), key_hash_sha(1), x509_name(2), cert_hash(3),
      (255)}
    IdentifierType;

opaque DistinguishedName<1..2^16-1>;

opaque KeyHash[20];
```

Here "TrustedAuthorities" provides a list of CA root key identifiers that the client possesses. Each CA root key is identified via either:

- "null" - no CA root key identity supplied.
- "key_hash_sha" - contains the SHA-1 hash of the CA root key. For

DSA and ECDSA keys, this is the hash of the "subjectPublicKey" value. For RSA keys, this is the hash of the byte string representation of the modulus (without any initial 0-valued bytes). (This copies the key hash formats deployed in other environments.)

- "cert_hash" - contains the SHA-1 hash of a certificate containing the CA root key.
- "x509_name" - contains the X.509 distinguished name of the CA.

Note that clients may include none, some, or all of the CA root keys they possess in this extension.

Note also that it is possible that a key hash or a distinguished name alone may not uniquely identify a certificate issuer - for example if a particular CA has multiple key pairs - however here we assume this is the case following the use of distinguish names to identify certificate issuers in TLS.

The option to include no CA root keys is included to allow the client to indicate possession of some pre-defined set of CA root keys.

Servers that receive a client hello containing the "trusted_ca_keys" extension, MAY use the information contained in the extension to guide their selection of an appropriate certificate chain to return to the client.

3.5. Truncated HMAC

Currently defined TLS ciphersuites use the MAC construction HMAC with either MD5 or SHA-1 [[HMAC](#)] to authenticate record layer communications. In TLS the entire output of the hash function is used as the MAC tag. However it may be desirable in constrained environments to save bandwidth by truncating the output of the hash function to 80 bits when forming MAC tags.

In order to negotiate the use of 80-bit truncated HMAC, clients MAY include an extension of type "truncated_hmac" in the extended client hello. The "extension_data" field of this extension shall be empty.

Servers that receive an extended hello containing a "truncated_hmac" extension, MAY agree to use a truncated HMAC by including an extension of type "truncated_hmac" in the extended server hello.

Note that if new ciphersuites are added that do not use HMAC, and the session negotiates one of these ciphersuites, this extension will have no effect. It is strongly recommended that any new cipher suites

using other MACs consider the MAC size as an integral part of the cipher suite definition, taking into account both security and bandwidth considerations.

If HMAC truncation has been successfully negotiated during a TLS handshake, and the negotiated ciphersuite uses HMAC, both the client and the server pass this fact to the TLS record layer along with the other negotiated security parameters. Subsequently during the session, clients and servers **MUST** use truncated HMACs, calculated as specified in [\[HMAC\]](#). Note that this extension does not affect the calculation of the PRF as part of handshaking or key derivation.

The negotiated HMAC truncation size applies for the duration of the session including session resumptions.

3.6. OCSP Status Request

Constrained clients may wish to use OCSP [\[OCSP\]](#) to check the validity of server certificates, in order to avoid transmission of CRLs and therefore save bandwidth on constrained networks.

In order to indicate their desire to use OCSP, clients **MAY** include an extension of type "status_request" in the (extended) client hello. The "extension_data" field of this extension shall contain "StatusRequest" where:

```
struct {  
    ResponderID responder_id_list<0..2^16-1>;  
    Extensions request_extensions;  
} StatusRequest;  
  
opaque ResponderID<1..2^16-1>;  
opaque Extensions<0..2^16-1>;
```

Here "ResponderIDs" provides a list of OCSP responders that the client trusts. A zero-length "responder_id_list" sequence has the special meaning that the responders are implicitly known to the server - e.g. by prior arrangement. "Extensions" is a DER encoding of the OCSP request extensions.

Both "ResponderID" and "Extensions" are DER-encoded ASN.1 types as defined in [\[OCSP\]](#).

Servers that receive a client hello containing the "status_request" extension, **MAY** return an OCSP response to the client along with their certificate. If so, they **SHOULD** use the information contained in the extension when selecting an OCSP responder, and **SHOULD** include request_extensions in the OCSP request.

Servers return an OCSP response along with their certificate by sending "CertificateAndOCSPResponse" in place of the "Certificate" message. If a server returns an OCSP response, then the server MUST have included an extension of type "status_request" with empty "extension_data" in the extended server hello.

```
struct {  
    ASN.1Cert certificate_list<0..2^24-1>;  
    OCSPResponse ocp_response;  
} CertificateAndOCSPResponse;  
  
opaque ASN.1Cert<1..2^24-1>;  
  
opaque OCSPResponse<1..2^24-1>;
```

Here "ocsp_response" contains a complete, DER-encoded OCSP response (using the ASN.1 type OCSPResponse defined in [[OCSP](#)]). Note that only one OCSP response may be sent.

The "CertificateAndOCSPResponse" message is conveyed using the handshake message type "ocsp_response".

Note that a server MAY also choose not to send the "CertificateAndOCSPResponse" message, and instead send the "Certificate" message, even if it receives a "status_request" extension in the client hello message.

Note in addition that servers MUST NOT send the "CertificateAndOCSPResponse" message unless it received a "status_request" extension in the client hello message.

Clients requesting an OCSP response, and receiving an OCSP response in a "CertificateAndOCSPResponse" field:

- MUST process the certificate as if it was received in a "Certificate" message, and;
- SHOULD check the OCSP response and abort the handshake if the response is not satisfactory.

4. Error Alerts

This section defines new error alerts for use with the TLS extensions defined in this document.

The following new error alerts are defined. To avoid "breaking" existing clients and servers, these alerts MUST NOT be sent unless the sending party has received an extended hello message from the

party they are communicating with.

- "unsupported_extension" - this alert is sent by clients that receive an extended server hello containing an extension that they did not put in the corresponding client hello (see [Section 2.3](#)). This message is always fatal.
- "bad_extension_order" - this alert is sent by clients or servers that receive an extended hello with the extensions in the wrong order (see [Section 2.3](#)). This message is always fatal.
- "unrecognised_domain" - this alert is sent by servers that receive a dns_name extension request, but do not recognize the DNS name as belonging to a domain they are responsible for. This message MAY be fatal.
- "certificate_unobtainable" - this alert is sent by servers who are unable to retrieve a certificate chain from the URL supplied by the client (see [Section 3.3](#)). This message is always fatal.
- "bad_ocsp_response" - this alert is sent by clients that receive an invalid OCSP response (see [Section 3.6](#)). This message is always fatal.

These error alerts are conveyed using the following syntax:

```
enum {  
    close_notify(0),  
    unexpected_message(10),  
    bad_record_mac(20),  
    decryption_failed(21),  
    record_overflow(22),  
    decompression_failure(30),  
    handshake_failure(40),  
    certificate_unobtainable(41), /* new */  
    bad_certificate(42),  
    unsupported_certificate(43),  
    certificate_revoked(44),  
    certificate_expired(45),  
    certificate_unknown(46),  
    illegal_parameter(47),  
    unknown_ca(48),  
    access_denied(49),  
    decode_error(50),  
    decrypt_error(51),  
    export_restriction(60),  
    protocol_version(70),  
    insufficient_security(71),
```



```
        internal_error(80),
        user_canceled(90),
        no_renegotiation(100),
        unsupported_extension(110),    /* new */
        bad_extension_order(111),      /* new */
        unrecognised_domain(112),      /* new */
        bad_ocsp_response(113),        /* new */
        (255)
    } AlertDescription;
```

5. Procedure for Defining New Extensions

Traditionally for Internet protocols, the Internet Assigned Numbers Authority (IANA) handles the allocation of new values for future expansion, and RFCs usually define the procedure to be used by the IANA. However, there are subtle (and not so subtle) interactions that may occur in this protocol between new features and existing features which may result in a significant reduction in overall security.

Therefore, requests to define new extensions (including assigning extension and error alert numbers) should be forwarded to the IETF TLS Working Group for discussion.

The following considerations should be taken into account when designing new extensions:

- All of the extensions defined in this document follow the convention that for each extension that a client requests and that the server understands, the server replies with an extension of the same type.
- Some cases where a server does not agree to an extension are error conditions, and some simply a refusal to support a particular feature. In general error alerts should be used for the former, and a field in the server extension response for the latter.
- Extensions should as far as possible be designed to prevent any attack that forces use (or non-use) of a particular feature by manipulation of handshake messages. This principle should be followed regardless of whether the feature is believed to cause a security problem.

Often the fact that the extension fields are included in the inputs to the Finished message hashes will be sufficient, but extreme care is needed when the extension changes the meaning of messages sent in the handshake phase.

Designers and implementors should be aware of the fact that until the handshake has been authenticated, active attackers can modify messages and insert, remove, or replace extensions.

- It would be technically possible to use extensions to change major aspects of the design of TLS; for example the design of ciphersuite negotiation. This is not recommended; it would be more appropriate to define a new version of TLS - particularly since the TLS handshake algorithms have specific protection against version rollback attacks based on the version number, and the possibility of version rollback should be a significant consideration in any major design change.

6. Security Considerations

Security considerations for the extension mechanism in general, and the design of new extensions, are described in the previous section. A security analysis of each of the extensions defined in this document is given below.

In general, implementers should continue to monitor the state of the art, and address any weaknesses identified.

Additional security considerations are described in the TLS 1.0 RFC [[TLS](#)].

6.1. Security of dns_name

If a single server hosts several domains, then clearly it is necessary for the owners of each domain to ensure that this satisfies their security needs. Apart from this, dns_name does not appear to introduce significant security issues.

The length of the domain name should be checked for buffer overflow (note that [RFC 1035](#) restricts domain names to 255 bytes).

6.2. Security of max_record_size

The maximum record size takes effect immediately, including for handshake messages. However, that does not introduce any security complications that are not already present in TLS, since [[TLS](#)] requires implementations to be able to handle fragmented handshake messages.

Note that as described in [section 3.2](#), once a non-null ciphersuite has been activated, the effective maximum record size depends on the ciphersuite, as well as on the negotiated max_record_size. This must

be taken into account when sizing buffers, and checking for buffer overflow.

6.3. Security of client_certificate_url

When client authentication is used **without** the client_certificate_url extension, the client certificate chain is covered by the Finished message hashes. The purpose of checking a SHA-1 hash of the certificate chain contents, is to ensure that the same property holds when this extension is used - i.e. that all of the information in the certificate chain retrieved by the server is as the client intended.

(The attack that this protects against is admittedly fairly unrealistic: the attacker would have to get a valid certificate on the client's key which is different from the client's certificate. Nevertheless, including this hash guarantees that there can be no unexpected problems due to the certificate chain not being directly covered by the Finished hashes.)

A similar approach should also be taken if possible for any future extensions that involve fetching information from external sources.

Note that although TLS uses both MD5 and SHA-1 hashes in several other places, this was not believed to be necessary here. The property required of SHA-1 is second pre-image resistance.

Support for client_certificate_url involves the server acting as a client in another protocol (usually HTTP, but other URL schemes are not prohibited). It is therefore subject to many of the same security considerations that apply to a publically accessible HTTP proxy server. This includes the possibility that an attacker might use the server to indirectly attack another host that is vulnerable to some security flaw. It also includes potentially increased exposure to denial of service attacks: an attacker can make many connections, each of which results in the server making an HTTP request.

It is recommended that the client_certificate_url extension should have to be specifically enabled by a server administrator, rather than being enabled by default.

As discussed in [[URI](#)], URLs that specify ports other than the default may cause problems, as may very long URLs (which are more likely to be useful in exploiting buffer overflow bugs).

6.4. Security of trusted_ca_keys

It is possible that which CA root keys a client possesses could be

regarded as confidential information. As a result, the CA root key indication extension should be used with care.

The use of the SHA-1 certificate hash alternative ensures that each certificate is specified unambiguously. As for the previous extension, it was not believed necessary to use both MD5 and SHA-1 hashes.

6.5. Security of truncated_hmac

It is possible that truncated MACs are weaker than "un-truncated" MACs. However, no significant weaknesses are currently known or expected to exist for HMAC with MD5 or SHA-1, truncated to 80 bits. Note that the output length of a MAC need not be as long as the length of a symmetric cipher key, since forging of MAC values cannot be done off-line: in TLS, a single failed MAC guess will cause the immediate termination of the TLS session.

Since the MAC algorithm only takes effect after the handshake messages have been authenticated by the hashes in the Finished messages, it is not possible for an active attacker to force negotiation of the truncated HMAC extension where it would not otherwise be used (to the extent that the handshake authentication is secure). Therefore, in the event that any security problem were found with truncated HMAC in future, if either the client or the server for a given session have been updated to take into account the problem, they would be able to veto use of this extension.

6.6. Security of status_request

If a client requests an OCSP response, it must take into account that an attacker's server using a compromised key could (and probably would) pretend not to support the extension. A client that requires OCSP validation of certificates SHOULD be prepared to contact the OCSP server directly in this case.

Use of the OCSP nonce request extension (id-pkix-ocsp-nonce) may improve security against attacks that attempt to replay OCSP responses; see section 4.4.1 of [[OCSP](#)] for further details.

7. Internationalisation Considerations

None of the extensions defined here directly use strings subject to localisation. Domain names are encoded using UTF-8. If future extensions use text strings, then internationalisation should be considered in their design.

8. Intellectual Property Rights

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this document. Please address the information to the IETF Executive Director.

9. Acknowledgments

The authors wish to thank the TLS Working Group and the WAP Security Group. This document is based on discussion within these groups.

10. References

[HMAC] Krawczyk, H., Bellare, M., and Canetti, R. - HMAC: Keyed-hashing for message authentication. IETF [RFC 2104](#), February 1997.

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

[KERB] A. Medvinsky, M. Hur, "Addition of Kerberos Cipher Suites to Transport Layer Security (TLS)," IETF [RFC 2712](#), October 1999.

[KEYWORDS] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels," IETF [RFC 2119](#), March 1997.

[MAILING LIST] Mikkelsen, J. Eberhard, R., and J. Kistler, "General ClientHello extension mechanism and virtual hosting," Ietf-tls mailing list posting, August 14, 2000.

[OCSP] Myers, M., Ankney, R., Malpani, A., Galperin, S., and C. Adams, "Internet X.509 Public Key Infrastructure: Online Certificate Status Protocol - OCSP," IETF [RFC 2560](#), June 1999.

[TLS] Dierks, T., and C. Allen, "The TLS Protocol - Version 1.0," IETF [RFC 2246](#), January 1999.

[URI] T. Berners-Lee, R. Fielding, L. Masinter, "Uniform Resource Identifiers (URI): Generic Syntax," IETF [RFC 2396](#), August 1998.

[UTF8] F. Yergeau, "UTF-8, a transformation format of ISO 10646," IETF [RFC 2279](#), January 1998.

11. Authors' Addresses

Simon Blake-Wilson
Certicom Corp.
sblake-wilson@certicom.com

Magnus Nystrom
RSA Security
magnus@rsasecurity.com

David Hopwood
Independent Consultant
david.hopwood@zetnet.co.uk

Jan Mikkelsen
Transactionware
jam@transactionware.com

Tim Wright
Vodafone
timothy.wright@vf.vodafone.co.uk

