

Network Working Group
Internet-Draft
Obsoletes: [3268](#), [4346](#), [4366](#), [5246](#) (if
approved)
Updates: [4492](#) (if approved)
Intended status: Standards Track
Expires: September 10, 2015

E. Rescorla
RTFM, Inc.
March 09, 2015

The Transport Layer Security (TLS) Protocol Version 1.3
draft-ietf-tls-tls13-05

Abstract

This document specifies Version 1.3 of the Transport Layer Security (TLS) protocol. The TLS protocol provides communications security over the Internet. The protocol allows client/server applications to communicate in a way that is designed to prevent eavesdropping, tampering, or message forgery.

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 September 10, 2015.

Copyright Notice

Copyright (c) 2015 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

Internet-Draft

TLS

March 2015

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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	4
1.1.	Requirements Terminology	5
1.2.	Major Differences from TLS 1.2	5
2.	Goals	6
3.	Goals of This Document	7
4.	Presentation Language	7
4.1.	Basic Block Size	7
4.2.	Miscellaneous	8
4.3.	Vectors	8
4.4.	Numbers	9
4.5.	Enumerateds	9
4.6.	Constructed Types	10
4.6.1.	Variants	11
4.7.	Cryptographic Attributes	12
4.8.	Constants	14
5.	The Pseudorandom Function	15
6.	The TLS Record Protocol	16
6.1.	Connection States	16
6.2.	Record Layer	19
6.2.1.	Fragmentation	19
6.2.2.	Record Payload Protection	20
6.3.	Key Calculation	22
7.	The TLS Handshaking Protocols	23
7.1.	Alert Protocol	24
7.1.1.	Closure Alerts	25

7.1.2.	Error Alerts	25
7.2.	Handshake Protocol Overview	29
7.3.	Handshake Protocol	33
7.3.1.	Hello Messages	34
7.3.2.	Client Key Share Message	37

Rescorla

Expires September 10, 2015

[Page 2]

Internet-Draft

TLS

March 2015

7.3.3.	Server Key Share Message	48
7.3.4.	Encrypted Extensions	49
7.3.5.	Server Certificate	49
7.3.6.	Certificate Request	52
7.3.7.	Server Certificate Verify	53
7.3.8.	Server Finished	55
7.3.9.	Client Certificate	56
7.3.10.	Client Certificate Verify	58
8.	Cryptographic Computations	58
8.1.	Computing the Master Secret	59
8.1.1.	The Session Hash	60
8.1.2.	Diffie-Hellman	61
8.1.3.	Elliptic Curve Diffie-Hellman	61
9.	Mandatory Cipher Suites	61
10.	Application Data Protocol	61
11.	Security Considerations	62
12.	IANA Considerations	62
13.	References	63
13.1.	Normative References	63
13.2.	Informative References	65
13.3.	URIs	68
Appendix A.	Protocol Data Structures and Constant Values	69
A.1.	Record Layer	69
A.2.	Alert Messages	69
A.3.	Handshake Protocol	70
A.3.1.	Hello Messages	71
A.3.2.	Key Exchange Messages	74
A.3.3.	Authentication Messages	74
A.3.4.	Handshake Finalization Messages	75
A.4.	The Cipher Suite	75
A.5.	The Security Parameters	77
A.6.	Changes to RFC 4492	77
Appendix B.	Glossary	78
Appendix C.	Cipher Suite Definitions	81
Appendix D.	Implementation Notes	81
D.1.	Random Number Generation and Seeding	81

D.2.	Certificates and Authentication	82
D.3.	Cipher Suites	82
D.4.	Implementation Pitfalls	82
Appendix E.	Backward Compatibility	83
E.1.	Compatibility with prior versions	83
E.2.	Compatibility with SSL	85
Appendix F.	Security Analysis	85
F.1.	Handshake Protocol	85
F.1.1.	Authentication and Key Exchange	86
F.1.2.	Version Rollback Attacks	87
F.1.3.	Detecting Attacks Against the Handshake Protocol	87
F.1.4.	Resuming Sessions	87

F.2.	Protecting Application Data	88
F.3.	Denial of Service	88
F.4.	Final Notes	89
Appendix G.	Working Group Information	89
Appendix H.	Contributors	89

[1.](#) Introduction

DISCLAIMER: This is a WIP draft of TLS 1.3 and has not yet seen significant security analysis.

RFC EDITOR: PLEASE REMOVE THE FOLLOWING PARAGRAPH The source for this draft is maintained in GitHub. Suggested changes should be submitted as pull requests at <https://github.com/tlswg/tls13-spec>.

Instructions are on that page as well. Editorial changes can be managed in GitHub, but any substantive change should be discussed on the TLS mailing list.

The primary goal of the TLS protocol is to provide privacy and data integrity between two communicating applications. The protocol is composed of two layers: the TLS Record Protocol and the TLS Handshake Protocol. At the lowest level, layered on top of some reliable transport protocol (e.g., TCP [[RFC0793](#)]), is the TLS Record Protocol. The TLS Record Protocol provides connection security that has two basic properties:

- The connection is private. Symmetric cryptography is used for data encryption (e.g., AES [[AES](#)], etc.). The keys for this symmetric encryption are generated uniquely for each connection

and are based on a secret negotiated by another protocol (such as the TLS Handshake Protocol). The Record Protocol can also be used without encryption, i.e., in integrity-only modes.

- The connection is reliable. Messages include an authentication tag which protects them against modification.
- The Record Protocol can operate in an insecure mode but is generally only used in this mode while another protocol is using the Record Protocol as a transport for negotiating security parameters.

The TLS Record Protocol is used for encapsulation of various higher-level protocols. One such encapsulated protocol, the TLS Handshake Protocol, allows the server and client to authenticate each other and to negotiate an encryption algorithm and cryptographic keys before the application protocol transmits or receives its first byte of data. The TLS Handshake Protocol provides connection security that has three basic properties:

- The peer's identity can be authenticated using asymmetric, or public key, cryptography (e.g., RSA [[RSA](#)], DSA [[DSS](#)], etc.). This authentication can be made optional, but is generally required for at least one of the peers.
- The negotiation of a shared secret is secure: the negotiated secret is unavailable to eavesdroppers, and for any authenticated connection the secret cannot be obtained, even by an attacker who can place himself in the middle of the connection.
- The negotiation is reliable: no attacker can modify the negotiation communication without being detected by the parties to the communication.

One advantage of TLS is that it is application protocol independent. Higher-level protocols can layer on top of the TLS protocol transparently. The TLS standard, however, does not specify how protocols add security with TLS; the decisions on how to initiate TLS handshaking and how to interpret the authentication certificates exchanged are left to the judgment of the designers and implementors of protocols that run on top of TLS.

1.1. Requirements Terminology

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

1.2. Major Differences from TLS 1.2

draft-05

- Prohibit SSL negotiation for backwards compatibility.
- Fix which MS is used for exporters.

draft-04

- Modify key computations to include session hash.
- Remove ChangeCipherSpec
- Renumber the new handshake messages to be somewhat more consistent with existing convention and to remove a duplicate registration.
- Remove renegotiation.

- Update format of signatures with context.
- Remove point format negotiation.

draft-03

- Remove GMT time.
- Merge in support for ECC from [RFC 4492](#) but without explicit curves.
- Remove the unnecessary length field from the AD input to AEAD ciphers.
- Rename {Client,Server}KeyExchange to {Client,Server}KeyShare

- Add an explicit HelloRetryRequest to reject the client's

[draft-02](#)

- Increment version number.
- Reworked handshake to provide 1-RTT mode.
- Remove custom DHE groups.
- Removed support for compression.
- Removed support for static RSA and DH key exchange.
- Removed support for non-AEAD ciphers

[2.](#) Goals

The goals of the TLS protocol, in order of priority, are as follows:

1. Cryptographic security: TLS should be used to establish a secure connection between two parties.
2. Interoperability: Independent programmers should be able to develop applications utilizing TLS that can successfully exchange cryptographic parameters without knowledge of one another's code.
3. Extensibility: TLS seeks to provide a framework into which new public key and record protection methods can be incorporated as necessary. This will also accomplish two sub-goals: preventing the need to create a new protocol (and risking the introduction

of possible new weaknesses) and avoiding the need to implement an entire new security library.

4. Relative efficiency: Cryptographic operations tend to be highly CPU intensive, particularly public key operations. For this reason, the TLS protocol has incorporated an optional session caching scheme to reduce the number of connections that need to be established from scratch. Additionally, care has been taken

to reduce network activity.

3. Goals of This Document

This document and the TLS protocol itself are based on the SSL 3.0 Protocol Specification as published by Netscape. The differences between this protocol and SSL 3.0 are not dramatic, but they are significant enough that the various versions of TLS and SSL 3.0 do not interoperate (although each protocol incorporates a mechanism by which an implementation can back down to prior versions). This document is intended primarily for readers who will be implementing the protocol and for those doing cryptographic analysis of it. The specification has been written with this in mind, and it is intended to reflect the needs of those two groups. For that reason, many of the algorithm-dependent data structures and rules are included in the body of the text (as opposed to in an appendix), providing easier access to them.

This document is not intended to supply any details of service definition or of interface definition, although it does cover select areas of policy as they are required for the maintenance of solid security.

4. Presentation Language

This document deals with the formatting of data in an external representation. The following very basic and somewhat casually defined presentation syntax will be used. The syntax draws from several sources in its structure. Although it resembles the programming language "C" in its syntax and XDR [[RFC4506](#)] in both its syntax and intent, it would be risky to draw too many parallels. The purpose of this presentation language is to document TLS only; it has no general application beyond that particular goal.

4.1. Basic Block Size

The representation of all data items is explicitly specified. The basic data block size is one byte (i.e., 8 bits). Multiple byte data items are concatenations of bytes, from left to right, from top to

bottom. From the byte stream, a multi-byte item (a numeric in the

example) is formed (using C notation) by:

```
value = (byte[0] << 8*(n-1)) | (byte[1] << 8*(n-2)) |  
        ... | byte[n-1];
```

This byte ordering for multi-byte values is the commonplace network byte order or big-endian format.

[4.2.](#) Miscellaneous

Comments begin with `/*` and end with `*/`.

Optional components are denoted by enclosing them in `"[[]]"` double brackets.

Single-byte entities containing uninterpreted data are of type `opaque`.

[4.3.](#) Vectors

A vector (single-dimensioned array) is a stream of homogeneous data elements. The size of the vector may be specified at documentation time or left unspecified until runtime. In either case, the length declares the number of bytes, not the number of elements, in the vector. The syntax for specifying a new type, `T'`, that is a fixed-length vector of type `T` is

```
T T'[n];
```

Here, `T'` occupies `n` bytes in the data stream, where `n` is a multiple of the size of `T`. The length of the vector is not included in the encoded stream.

In the following example, `Datum` is defined to be three consecutive bytes that the protocol does not interpret, while `Data` is three consecutive `Datum`, consuming a total of nine bytes.

```
opaque Datum[3];      /* three uninterpreted bytes */  
Datum Data[9];        /* 3 consecutive 3 byte vectors */
```

Variable-length vectors are defined by specifying a subrange of legal lengths, inclusively, using the notation `<floor..ceiling>`. When these are encoded, the actual length precedes the vector's contents in the byte stream. The length will be in the form of a number consuming as many bytes as required to hold the vector's specified maximum (ceiling) length. A variable-length vector with an actual length field of zero is referred to as an empty vector.

```
T T'<floor..ceiling>;
```

In the following example, `mandatory` is a vector that must contain between 300 and 400 bytes of type `opaque`. It can never be empty. The actual length field consumes two bytes, a `uint16`, which is sufficient to represent the value 400 (see [Section 4.4](#)). On the other hand, `longer` can represent up to 800 bytes of data, or 400 `uint16` elements, and it may be empty. Its encoding will include a two-byte actual length field prepended to the vector. The length of an encoded vector must be an even multiple of the length of a single element (for example, a 17-byte vector of `uint16` would be illegal).

```
opaque mandatory<300..400>;
    /* length field is 2 bytes, cannot be empty */
uint16 longer<0..800>;
    /* zero to 400 16-bit unsigned integers */
```

[4.4](#). Numbers

The basic numeric data type is an unsigned byte (`uint8`). All larger numeric data types are formed from fixed-length series of bytes concatenated as described in [Section 4.1](#) and are also unsigned. The following numeric types are predefined.

```
uint8 uint16[2];
uint8 uint24[3];
uint8 uint32[4];
uint8 uint64[8];
```

All values, here and elsewhere in the specification, are stored in network byte (big-endian) order; the `uint32` represented by the hex bytes 01 02 03 04 is equivalent to the decimal value 16909060.

Note that in some cases (e.g., DH parameters) it is necessary to represent integers as opaque vectors. In such cases, they are represented as unsigned integers (i.e., leading zero octets are not required even if the most significant bit is set).

[4.5](#). Enumerateds

An additional sparse data type is available called `enum`. A field of type `enum` can only assume the values declared in the definition. Each definition is a different type. Only enumerateds of the same type may be assigned or compared. Every element of an enumerated must be assigned a value, as demonstrated in the following example. Since the elements of the enumerated are not ordered, they can be

assigned any unique value, in any order.

```
enum { e1(v1), e2(v2), ... , en(vn) [[, (n)]] } Te;
```

An enumerated occupies as much space in the byte stream as would its maximal defined ordinal value. The following definition would cause one byte to be used to carry fields of type Color.

```
enum { red(3), blue(5), white(7) } Color;
```

One may optionally specify a value without its associated tag to force the width definition without defining a superfluous element.

In the following example, Taste will consume two bytes in the data stream but can only assume the values 1, 2, or 4.

```
enum { sweet(1), sour(2), bitter(4), (32000) } Taste;
```

The names of the elements of an enumeration are scoped within the defined type. In the first example, a fully qualified reference to the second element of the enumeration would be Color.blue. Such qualification is not required if the target of the assignment is well specified.

```
Color color = Color.blue;    /* overspecified, legal */  
Color color = blue;          /* correct, type implicit */
```

For enumerations that are never converted to external representation, the numerical information may be omitted.

```
enum { low, medium, high } Amount;
```

[4.6.](#) Constructed Types

Structure types may be constructed from primitive types for convenience. Each specification declares a new, unique type. The syntax for definition is much like that of C.

```
struct {  
    T1 f1;  
    T2 f2;
```

```

    ...
    Tn fn;
} [[T]];

```

The fields within a structure may be qualified using the type's name, with a syntax much like that available for enumerations. For example, `T.f2` refers to the second field of the previous declaration. Structure definitions may be embedded.

[4.6.1.](#) Variants

Defined structures may have variants based on some knowledge that is available within the environment. The selector must be an enumerated type that defines the possible variants the structure defines. There must be a case arm for every element of the enumeration declared in the select. Case arms have limited fall-through: if two case arms follow in immediate succession with no fields in between, then they both contain the same fields. Thus, in the example below, "orange" and "banana" both contain `V2`. Note that this is a new piece of syntax in TLS 1.2.

The body of the variant structure may be given a label for reference. The mechanism by which the variant is selected at runtime is not prescribed by the presentation language.

```

struct {
    T1 f1;
    T2 f2;
    ...
    Tn fn;
    select (E) {
        case e1: Te1;
        case e2: Te2;
        case e3: case e4: Te3;
        ...
        case en: Ten;
    } [[fv]];
} [[Tv]];

```

For example:

```
enum { apple, orange, banana } VariantTag;

struct {
    uint16 number;
    opaque string<0..10>; /* variable length */
} V1;

struct {
    uint32 number;
    opaque string[10];    /* fixed length */
} V2;

struct {
    select (VariantTag) { /* value of selector is implicit */
        case apple:
            V1; /* VariantBody, tag = apple */
        case orange:
        case banana:
            V2; /* VariantBody, tag = orange or banana */
    } variant_body; /* optional label on variant */
} VariantRecord;
```

[4.7.](#) Cryptographic Attributes

The two cryptographic operations -- digital signing, and

authenticated encryption with additional data (AEAD) -- are designated digitally-signed, and aead-ciphered, respectively. A field's cryptographic processing is specified by prepending an appropriate key word designation before the field's type specification. Cryptographic keys are implied by the current session state (see [Section 6.1](#)).

A digitally-signed element is encoded as a struct DigitallySigned:

```
struct {  
    SignatureAndHashAlgorithm algorithm;  
    opaque signature<0..2^16-1>;  
} DigitallySigned;
```

The algorithm field specifies the algorithm used (see [Section 7.3.2.5.1](#) for the definition of this field). Note that the algorithm field was introduced in TLS 1.2, and is not in earlier versions. The signature is a digital signature using those algorithms over the contents of the element. The contents themselves do not appear on the wire but are simply calculated. The length of the signature is specified by the signing algorithm and key.

In previous versions of TLS, the ServerKeyExchange format meant that attackers can obtain a signature of a message with a chosen, 32-byte prefix. Because TLS 1.3 servers are likely to also implement prior versions, the contents of the element always start with 64 bytes of octet 32 in order to clear that chosen-prefix.

Following that padding is a NUL-terminated context string in order to disambiguate signatures for different purposes. The context string will be specified whenever a digitally-signed element is used.

Finally, the specified contents of the digitally-signed structure follow the NUL at the end of the context string. (See the example at the end of this section.)

In RSA signing, the opaque vector contains the signature generated using the RSASSA-PKCS1-v1_5 signature scheme defined in [\[RFC3447\]](#). As discussed in [\[RFC3447\]](#), the DigestInfo MUST be DER-encoded [\[X680\]](#) [\[X690\]](#). For hash algorithms without parameters (which includes SHA-

1), the `DigestInfo.AlgorithmIdentifier.parameters` field MUST be NULL, but implementations MUST accept both without parameters and with NULL parameters. Note that earlier versions of TLS used a different RSA signature scheme that did not include a `DigestInfo` encoding.

In DSA, the 20 bytes of the SHA-1 hash are run directly through the Digital Signing Algorithm with no additional hashing. This produces two values, `r` and `s`. The DSA signature is an opaque vector, as above, the contents of which are the DER encoding of:

```
Dss-Sig-Value ::= SEQUENCE {  
    r INTEGER,  
    s INTEGER  
}
```

Note: In current terminology, DSA refers to the Digital Signature Algorithm and DSS refers to the NIST standard. In the original SSL and TLS specs, "DSS" was used universally. This document uses "DSA" to refer to the algorithm, "DSS" to refer to the standard, and it uses "DSS" in the code point definitions for historical continuity.

All ECDSA computations MUST be performed according to ANSI X9.62 [X962] or its successors. Data to be signed/verified is hashed, and the result run directly through the ECDSA algorithm with no additional hashing. The default hash function is SHA-1 [SHS]. However, an alternative hash function, such as one of the new SHA hash functions specified in FIPS 180-2 may be used instead if the certificate containing the EC public key explicitly requires use of another hash function. (The mechanism for specifying the required hash function has not been standardized, but this provision

anticipates such standardization and obviates the need to update this document in response. Future PKIX RFCs may choose, for example, to specify the hash function to be used with a public key in the `parameters` field of `subjectPublicKeyInfo`.) [[OPEN ISSUE: This needs updating per 4492-bis <https://github.com/tlswg/tls13-spec/issues/59>]]

In AEAD encryption, the plaintext is simultaneously encrypted and integrity protected. The input may be of any length, and `aead-ciphered` output is generally larger than the input in order to accommodate the integrity check value.

In the following example

```
struct {
    uint8 field1;
    uint8 field2;
    digitally-signed opaque {
        uint8 field3<0..255>;
        uint8 field4;
    };
} UserType;
```

Assume that the context string for the signature was specified as "Example". The input for the signature/hash algorithm would be:

[illegible]

followed by the encoding of the inner struct (field3 and field4).

The length of the structure, in bytes, would be equal to two bytes for field1 and field2, plus two bytes for the signature and hash algorithm, plus two bytes for the length of the signature, plus the length of the output of the signing algorithm. The length of the signature is known because the algorithm and key used for the signing are known prior to encoding or decoding this structure.

4.8. Constants

Typed constants can be defined for purposes of specification by declaring a symbol of the desired type and assigning values to it.

Under-specified types (opaque, variable-length vectors, and structures that contain opaque) cannot be assigned values. No fields of a multi-element structure or vector may be elided.

For example:

```
struct {
    uint8 f1;
    uint8 f2;
} Example1;
```



```
Example1 ex1 = {1, 4}; /* assigns f1 = 1, f2 = 4 */
```

5. The Pseudorandom Function

A construction is required to do expansion of secrets into blocks of data for the purposes of key generation or validation. This pseudorandom function (PRF) takes as input a secret, a seed, and an identifying label and produces an output of arbitrary length.

In this section, we define one PRF, based on HMAC [[RFC2104](#)]. This PRF with the SHA-256 hash function is used for all cipher suites defined in this document and in TLS documents published prior to this document when TLS 1.2 is negotiated. New cipher suites **MUST** explicitly specify a PRF and, in general, **SHOULD** use the TLS PRF with SHA-256 or a stronger standard hash function.

First, we define a data expansion function, $P_hash(secret, data)$, that uses a single hash function to expand a secret and seed into an arbitrary quantity of output:

$$P_hash(secret, seed) = \text{HMAC_hash}(secret, A(1) + seed) + \\ \text{HMAC_hash}(secret, A(2) + seed) + \\ \text{HMAC_hash}(secret, A(3) + seed) + \dots$$

where $+$ indicates concatenation.

$A()$ is defined as:

$$A(0) = seed \\ A(i) = \text{HMAC_hash}(secret, A(i-1))$$

P_hash can be iterated as many times as necessary to produce the required quantity of data. For example, if P_SHA256 is being used to create 80 bytes of data, it will have to be iterated three times (through $A(3)$), creating 96 bytes of output data; the last 16 bytes of the final iteration will then be discarded, leaving 80 bytes of output data.

TLS's PRF is created by applying P_hash to the secret as:

$$\text{PRF}(secret, label, seed) = P_hash(secret, label + seed)$$

The label is an ASCII string. It should be included in the exact form it is given without a length byte or trailing null character. For example, the label "slithy toves" would be processed by hashing the following bytes:

73 6C 69 74 68 79 20 74 6F 76 65 73

[6.](#) The TLS Record Protocol

The TLS Record Protocol is a layered protocol. At each layer, messages may include fields for length, description, and content. The Record Protocol takes messages to be transmitted, fragments the data into manageable blocks, protects the records, and transmits the result. Received data is decrypted and verified, reassembled, and then delivered to higher-level clients.

Three protocols that use the record protocol are described in this document: the handshake protocol, the alert protocol, and the application data protocol. In order to allow extension of the TLS protocol, additional record content types can be supported by the record protocol. New record content type values are assigned by IANA in the TLS Content Type Registry as described in [Section 12](#).

Implementations **MUST NOT** send record types not defined in this document unless negotiated by some extension. If a TLS implementation receives an unexpected record type, it **MUST** send an `unexpected_message` alert.

Any protocol designed for use over TLS must be carefully designed to deal with all possible attacks against it. As a practical matter, this means that the protocol designer must be aware of what security properties TLS does and does not provide and cannot safely rely on the latter.

Note in particular that type and length of a record are not protected by encryption. If this information is itself sensitive, application designers may wish to take steps (padding, cover traffic) to minimize information leakage.

[6.1.](#) Connection States

A TLS connection state is the operating environment of the TLS Record Protocol. It specifies a record protection algorithm and its parameters as well as the record protection keys and IVs for the connection in both the read and the write directions. The security parameters are set by the TLS Handshake Protocol, which also determines when new cryptographic keys are installed and used for

Internet-Draft

TLS

March 2015

record protection. The initial current state always specifies that records are not protected.

The security parameters for a TLS Connection read and write state are set by providing the following values:

connection end

Whether this entity is considered the "client" or the "server" in this connection.

PRF algorithm

An algorithm used to generate keys from the master secret (see [Section 5](#) and [Section 6.3](#)).

record protection algorithm

The algorithm to be used for record protection. This algorithm must be of the AEAD type and thus provides integrity and confidentiality as a single primitive. It is possible to have AEAD algorithms which do not provide any confidentiality and [Section 6.2.2](#) defines a special NULL_NULL AEAD algorithm for use in the initial handshake). This specification includes the key size of this algorithm and the lengths of explicit and implicit initialization vectors (or nonces).

handshake master secret

A 48-byte secret shared between the two peers in the connection and used to generate keys for protecting the handshake.

master secret

A 48-byte secret shared between the two peers in the connection and used to generate keys for protecting application data.

client random

A 32-byte value provided by the client.

server random

A 32-byte value provided by the server.

These parameters are defined in the presentation language as:

Internet-Draft

TLS

March 2015

```
enum { server, client } ConnectionEnd;

enum { tls_prf_sha256 } PRFAlgorithm;

enum { aes_gcm } RecordProtAlgorithm;

/* The algorithms specified in PRFAlgorithm and
   RecordProtAlgorithm may be added to. */

struct {
    ConnectionEnd          entity;
    PRFAlgorithm           prf_algorithm;
    RecordProtAlgorithm    record_prot_algorithm;
    uint8                  enc_key_length;
    uint8                  block_length;
    uint8                  fixed_iv_length;
    uint8                  record_iv_length;
    opaque                 hs_master_secret[48];
    opaque                 master_secret[48];
    opaque                 client_random[32];
    opaque                 server_random[32];
} SecurityParameters;
```

The record layer will use the security parameters to generate the following four items (some of which are not required by all ciphers, and are thus empty):

```
client write key
server write key
client write IV
server write IV
```

The client write parameters are used by the server when receiving and

processing records and vice versa. The algorithm used for generating these items from the security parameters is described in [Section 6.3](#)

Once the security parameters have been set and the keys have been generated, the connection states can be instantiated by making them the current states. These current states **MUST** be updated for each record processed. Each connection state includes the following elements:

cipher state

The current state of the encryption algorithm. This will consist of the scheduled key for that connection.

sequence number

Each connection state contains a sequence number, which is maintained separately for read and write states. The sequence number **MUST** be set to zero whenever a connection state is made the active state. Sequence numbers are of type uint64 and may not exceed $2^{64}-1$. Sequence numbers do not wrap. If a TLS implementation would need to wrap a sequence number, it must terminate the connection. A sequence number is incremented after each record: specifically, the first record transmitted under a particular connection state **MUST** use sequence number 0.

[6.2.](#) Record Layer

The TLS record layer receives uninterpreted data from higher layers in non-empty blocks of arbitrary size.

[6.2.1.](#) Fragmentation

The record layer fragments information blocks into TLSPlaintext records carrying data in chunks of 2^{14} bytes or less. Client message boundaries are not preserved in the record layer (i.e., multiple client messages of the same ContentType **MAY** be coalesced into a single TLSPlaintext record, or a single message **MAY** be fragmented across several records).

```
struct {  
    uint8 major;
```

```

    uint8 minor;
} ProtocolVersion;

ProtocolVersion version = { 3, 4 };    /* TLS v1.3*/

enum {
    reserved(20), alert(21), handshake(22),
    application_data(23), (255)
} ContentType;

struct {
    ContentType type;
    ProtocolVersion version;
    uint16 length;
    opaque fragment[TLSPlaintext.length];
} TLSPlaintext;

```

type

The higher-level protocol used to process the enclosed fragment.

version

The version of the protocol being employed. This document describes TLS Version 1.3, which uses the version { 3, 4 }. The version value 3.4 is historical, deriving from the use of {3, 1} for TLS 1.0. (See [Appendix A.1](#).) Note that a client that supports multiple versions of TLS may not know what version will be employed before it receives the ServerHello. See [Appendix E](#) for discussion about what record layer version number should be employed for ClientHello.

length

The length (in bytes) of the following TLSPlaintext.fragment. The length MUST NOT exceed 2^{14} .

fragment

The application data. This data is transparent and treated as an independent block to be dealt with by the higher-level protocol specified by the type field.

Implementations MUST NOT send zero-length fragments of Handshake or

Alert types. Zero-length fragments of Application data MAY be sent as they are potentially useful as a traffic analysis countermeasure.

[6.2.2.](#) Record Payload Protection

The record protection functions translate a TLSPlaintext structure into a TLSCiphertext. The deprotection functions reverse the process. In TLS 1.3 as opposed to previous versions of TLS, all ciphers are modelled as "Authenticated Encryption with Additional Data" (AEAD) [[RFC5116](#)]. AEAD functions provide a unified encryption and authentication operation which turns plaintext into authenticated ciphertext and back again.

AEAD ciphers take as input a single key, a nonce, a plaintext, and "additional data" to be included in the authentication check, as described in [Section 2.1 of \[RFC5116\]](#). The key is either the client_write_key or the server_write_key.

```
struct {
    ContentType type;
    ProtocolVersion version;
    uint16 length;
    opaque nonce_explicit[SecurityParameters.record_iv_length];
    aead-ciphered struct {
        opaque content[TLSPlaintext.length];
    } fragment;
} TLSCiphertext;
```

type

Rescorla	Expires September 10, 2015	[Page 20]
----------	----------------------------	-----------

Internet-Draft	TLS	March 2015
----------------	-----	------------

The type field is identical to TLSPlaintext.type.

version

The version field is identical to TLSPlaintext.version.

length

The length (in bytes) of the following TLSCiphertext.fragment.
The length MUST NOT exceed $2^{14} + 2048$.

fragment

The AEAD encrypted form of TLSPlaintext.fragment.

Each AEAD cipher suite MUST specify how the nonce supplied to the AEAD operation is constructed, and what is the length of the `TLSCiphertext.nonce_explicit` part. In many cases, it is appropriate to use the partially implicit nonce technique described in [Section 3.2.1 of \[RFC5116\]](#); with `record_iv_length` being the length of the explicit part. In this case, the implicit part SHOULD be derived from `key_block` as `client_write_iv` and `server_write_iv` (as described in [Section 6.3](#)), and the explicit part is included in `GenericAEADCipher.nonce_explicit`.

The plaintext is the `TLSPlainText.fragment`.

The additional authenticated data, which we denote as `additional_data`, is defined as follows:

$$\text{additional_data} = \text{seq_num} + \text{TLSPlainText.type} + \text{TLSPlainText.version}$$

where "+" denotes concatenation.

Note: In versions of TLS prior to 1.3, the `additional_data` included a length field. This presents a problem for cipher constructions with data-dependent padding (such as CBC). TLS 1.3 removes the length field and relies on the AEAD cipher to provide integrity for the length of the data.

The AEAD output consists of the ciphertext output by the AEAD encryption operation. The length will generally be larger than `TLSPlainText.length`, but by an amount that varies with the AEAD cipher. Since the ciphers might incorporate padding, the amount of overhead could vary with different `TLSPlainText.length` values. Each AEAD cipher MUST NOT produce an expansion of greater than 1024 bytes. Symbolically,

$$\text{AEADEncrypted} = \text{AEAD-Encrypt}(\text{write_key}, \text{nonce}, \text{plaintext}, \text{additional_data})$$

[[OPEN ISSUE: Reduce these values? <https://github.com/tlswg/tls13-spec/issues/55>]]

In order to decrypt and verify, the cipher takes as input the key, nonce, the "`additional_data`", and the `AEADEncrypted` value. The

output is either the plaintext or an error indicating that the decryption failed. There is no separate integrity check. That is:

```
    TLSPlaintext.fragment = AEAD-Decrypt(write_key, nonce,  
                                         AEADEncrypted,  
                                         additional_data)
```

If the decryption fails, a fatal bad_record_mac alert MUST be generated.

As a special case, we define the NULL_NULL AEAD cipher which is simply the identity operation and thus provides no security. This cipher MUST ONLY be used with the initial TLS_NULL_WITH_NULL_NULL cipher suite.

[6.3.](#) Key Calculation

[[OPEN ISSUE: This needs to be revised. See <https://github.com/tlswg/tls13-spec/issues/5>]] The Record Protocol requires an algorithm to generate keys required by the current connection state (see [Appendix A.5](#)) from the security parameters provided by the handshake protocol.

The master secret is expanded into a sequence of secure bytes, which is then split to a client write encryption key and a server write encryption key. Each of these is generated from the byte sequence in that order. Unused values are empty. Some ciphers may additionally require a client write IV and a server write IV.

When keys are generated, the current master secret (MS) is used as an entropy source. For handshake records, this means the hs_master_secret. For application data records, this means the regular master_secret.

To generate the key material, compute

```
    key_block = PRF(MS,  
                   "key expansion",  
                   SecurityParameters.server_random +  
                   SecurityParameters.client_random);
```

where MS is the relevant master secret. The PRF is computed enough times to generate the necessary amount of data for the key_block, which is then partitioned as follows:

```
client_write_key[SecurityParameters.enc_key_length]
server_write_key[SecurityParameters.enc_key_length]
client_write_IV[SecurityParameters.fixed_iv_length]
server_write_IV[SecurityParameters.fixed_iv_length]
```

Currently, the client_write_IV and server_write_IV are only generated for implicit nonce techniques as described in [Section 3.2.1 of \[RFC5116\]](#).

7. The TLS Handshaking Protocols

TLS has three subprotocols that are used to allow peers to agree upon security parameters for the record layer, to authenticate themselves, to instantiate negotiated security parameters, and to report error conditions to each other.

The Handshake Protocol is responsible for negotiating a session, which consists of the following items:

session identifier

An arbitrary byte sequence chosen by the server to identify an active or resumable session state.

peer certificate

X509v3 [\[RFC3280\]](#) certificate of the peer. This element of the state may be null.

cipher spec

Specifies the authentication and key establishment algorithms, the pseudorandom function (PRF) used to generate keying material, and the record protection algorithm (See [Appendix A.5](#) for formal definition.)

resumption premaster secret

48-byte secret shared between the client and server.

is resumable

A flag indicating whether the session can be used to initiate new connections.

These items are then used to create security parameters for use by the record layer when protecting application data. Many connections can be instantiated using the same session through the resumption feature of the TLS Handshake Protocol.

Internet-Draft

TLS

March 2015

[7.1.](#) Alert Protocol

One of the content types supported by the TLS record layer is the alert type. Alert messages convey the severity of the message (warning or fatal) and a description of the alert. Alert messages with a level of fatal result in the immediate termination of the connection. In this case, other connections corresponding to the session may continue, but the session identifier **MUST** be invalidated, preventing the failed session from being used to establish new connections. Like other messages, alert messages are encrypted as specified by the current connection state.

```
enum { warning(1), fatal(2), (255) } AlertLevel;
```

```
enum {  
    close_notify(0),  
    unexpected_message(10),  
    bad_record_mac(20),  
    decryption_failed_RESERVED(21),  
    record_overflow(22),  
    decompression_failure_RESERVED(30),  
    handshake_failure(40),  
    no_certificate_RESERVED(41),  
    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_RESERVED(60),  
    protocol_version(70),  
    insufficient_security(71),  
    internal_error(80),  
    user_canceled(90),  
    no_renegotiation(100),  
    unsupported_extension(110),  
    (255)  
} AlertDescription;
```

```
struct {  
    AlertLevel level;  
    AlertDescription description;  
} Alert;
```

[7.1.1.](#) Closure Alerts

The client and the server must share knowledge that the connection is ending in order to avoid a truncation attack. Either party may initiate the exchange of closing messages.

close_notify

This message notifies the recipient that the sender will not send any more messages on this connection. Note that as of TLS 1.1, failure to properly close a connection no longer requires that a session not be resumed. This is a change from TLS 1.0 to conform with widespread implementation practice.

Either party MAY initiate a close by sending a close_notify alert. Any data received after a closure alert is ignored.

Unless some other fatal alert has been transmitted, each party is required to send a close_notify alert before closing the write side of the connection. The other party MUST respond with a close_notify alert of its own and close down the connection immediately, discarding any pending writes. It is not required for the initiator of the close to wait for the responding close_notify alert before closing the read side of the connection.

If the application protocol using TLS provides that any data may be carried over the underlying transport after the TLS connection is closed, the TLS implementation must receive the responding close_notify alert before indicating to the application layer that the TLS connection has ended. If the application protocol will not transfer any additional data, but will only close the underlying transport connection, then the implementation MAY choose to close the transport without waiting for the responding close_notify. No part of this standard should be taken to dictate the manner in which a usage profile for TLS manages its data transport, including when connections are opened or closed.

Note: It is assumed that closing a connection reliably delivers pending data before destroying the transport.

[7.1.2.](#) Error Alerts

Error handling in the TLS Handshake protocol is very simple. When an error is detected, the detecting party sends a message to the other party. Upon transmission or receipt of a fatal alert message, both parties immediately close the connection. Servers and clients **MUST** forget any session-identifiers, keys, and secrets associated with a failed connection. Thus, any connection terminated with a fatal alert **MUST NOT** be resumed.

Whenever an implementation encounters a condition which is defined as a fatal alert, it **MUST** send the appropriate alert prior to closing the connection. For all errors where an alert level is not explicitly specified, the sending party **MAY** determine at its discretion whether to treat this as a fatal error or not. If the implementation chooses to send an alert but intends to close the connection immediately afterwards, it **MUST** send that alert at the fatal alert level.

If an alert with a level of warning is sent and received, generally the connection can continue normally. If the receiving party decides not to proceed with the connection (e.g., after having received a `no_renegotiation` alert that it is not willing to accept), it **SHOULD** send a fatal alert to terminate the connection. Given this, the sending party cannot, in general, know how the receiving party will behave. Therefore, warning alerts are not very useful when the sending party wants to continue the connection, and thus are sometimes omitted. For example, if a peer decides to accept an expired certificate (perhaps after confirming this with the user) and wants to continue the connection, it would not generally send a `certificate_expired` alert.

The following error alerts are defined:

`unexpected_message`

An inappropriate message was received. This alert is always fatal and should never be observed in communication between proper implementations.

bad_record_mac

This alert is returned if a record is received which cannot be deprotected. Because AEAD algorithms combine decryption and verification, this message is used for all deprotection failures. This message is always fatal and should never be observed in communication between proper implementations (except when messages were corrupted in the network).

decryption_failed_RESERVED

This alert was used in some earlier versions of TLS, and may have permitted certain attacks against the CBC mode [[CBCATT](#)]. It MUST NOT be sent by compliant implementations.

record_overflow

A TLSCiphertext record was received that had a length more than $2^{14}+2048$ bytes, or a record decrypted to a TLSPlaintext record with more than 2^{14} bytes. This message is always fatal and should never be observed in communication between proper

implementations (except when messages were corrupted in the network).

decompression_failure

This alert was used in previous versions of TLS. TLS 1.3 does not include compression and TLS 1.3 implementations MUST NOT send this alert when in TLS 1.3 mode.

handshake_failure

Reception of a handshake_failure alert message indicates that the sender was unable to negotiate an acceptable set of security parameters given the options available. This is a fatal error.

no_certificate_RESERVED

This alert was used in SSLv3 but not any version of TLS. It MUST NOT be sent by compliant implementations.

bad_certificate

A certificate was corrupt, contained signatures that did not verify correctly, etc.

unsupported_certificate

A certificate was of an unsupported type.

certificate_revoked

A certificate was revoked by its signer.

certificate_expired

A certificate has expired or is not currently valid.

certificate_unknown

Some other (unspecified) issue arose in processing the certificate, rendering it unacceptable.

illegal_parameter

A field in the handshake was out of range or inconsistent with other fields. This message is always fatal.

unknown_ca

A valid certificate chain or partial chain was received, but the certificate was not accepted because the CA certificate could not be located or couldn't be matched with a known, trusted CA. This message is always fatal.

access_denied

A valid certificate was received, but when access control was applied, the sender decided not to proceed with negotiation. This message is always fatal.

decode_error

A message could not be decoded because some field was out of the specified range or the length of the message was incorrect. This message is always fatal and should never be observed in communication between proper implementations (except when messages were corrupted in the network).

decrypt_error

A handshake cryptographic operation failed, including being unable to correctly verify a signature or validate a Finished message. This message is always fatal.

export_restriction_RESERVED

This alert was used in some earlier versions of TLS. It MUST NOT

be sent by compliant implementations.

`protocol_version`

The protocol version the client has attempted to negotiate is recognized but not supported. (For example, old protocol versions might be avoided for security reasons.) This message is always fatal.

`insufficient_security`

Returned instead of `handshake_failure` when a negotiation has failed specifically because the server requires ciphers more secure than those supported by the client. This message is always fatal.

`internal_error`

An internal error unrelated to the peer or the correctness of the protocol (such as a memory allocation failure) makes it impossible to continue. This message is always fatal.

`user_canceled`

This handshake is being canceled for some reason unrelated to a protocol failure. If the user cancels an operation after the handshake is complete, just closing the connection by sending a `close_notify` is more appropriate. This alert should be followed by a `close_notify`. This message is generally a warning.

`no_renegotiation`

Sent by the client in response to a hello request or by the server in response to a client hello after initial handshaking. Versions of TLS prior to TLS 1.3 supported renegotiation of a previously established connection; TLS 1.3 removes this feature. This message is always fatal.

`unsupported_extension`

sent by clients that receive an extended server hello containing an extension that they did not put in the corresponding client hello. This message is always fatal.

New Alert values are assigned by IANA as described in [Section 12](#).

[7.2](#). Handshake Protocol Overview

The cryptographic parameters of the session state are produced by the TLS Handshake Protocol, which operates on top of the TLS record layer. When a TLS client and server first start communicating, they agree on a protocol version, select cryptographic algorithms, optionally authenticate each other, and use public-key encryption techniques to generate shared secrets.

The TLS Handshake Protocol involves the following steps:

- Exchange hello messages to agree on a protocol version, algorithms, exchange random values, and check for session resumption.
- Exchange the necessary cryptographic parameters to allow the client and server to agree on a premaster secret.
- Exchange certificates and cryptographic information to allow the client and server to authenticate themselves.
- Generate a master secret from the premaster secret and exchanged random values.
- Provide security parameters to the record layer.
- Allow the client and server to verify that their peer has calculated the same security parameters and that the handshake occurred without tampering by an attacker.

Note that higher layers should not be overly reliant on whether TLS always negotiates the strongest possible connection between two peers. There are a number of ways in which a man-in-the-middle attacker can attempt to make two entities drop down to the least secure method they support. The protocol has been designed to minimize this risk, but there are still attacks available: for example, an attacker could block access to the port a secure service runs on, or attempt to get the peers to negotiate an unauthenticated connection. The fundamental rule is that higher levels must be cognizant of what their security requirements are and never transmit information over a channel less secure than what they require. The TLS protocol is secure in that any cipher suite offers its promised

level of security: if you negotiate AES-GCM [GCM] with a 1024-bit DHE key exchange with a host whose certificate you have verified, you can expect to be that secure.

These goals are achieved by the handshake protocol, which can be summarized as follows: The client sends a ClientHello message which contains a random nonce (ClientHello.random), its preferences for Protocol Version, Cipher Suite, and a variety of extensions. In the same flight, it sends a ClientKeyShare message which contains its share of the parameters for key agreement for some set of expected server parameters (DHE/ECDHE groups, etc.).

If the client has provided a ClientKeyShare with an appropriate set of keying material, the server responds to the ClientHello with a ServerHello message. The ServerHello contains the server's nonce (ServerHello.random), the server's choice of the Protocol Version, Session ID and Cipher Suite, and the server's response to the extensions the client offered.

The server can then generate its own keying material share and send a ServerKeyShare message which contains its share of the parameters for the key agreement. The server can now compute the shared secret (the premaster secret). At this point, the server starts encrypting all remaining handshake traffic with the negotiated cipher suite using a key derived from the premaster secret (via the "handshake master secret"). The remainder of the server's handshake messages will be encrypted using that key.

Following these messages, the server will send an EncryptedExtensions message which contains a response to any client's extensions which are not necessary to establish the Cipher Suite. The server will then send its certificate in a Certificate message if it is to be authenticated. The server may optionally request a certificate from the client by sending a CertificateRequest message at this point. Finally, if the server is authenticated, it will send a CertificateVerify message which provides a signature over the entire handshake up to this point. This serves both to authenticate the server and to establish the integrity of the negotiation. Finally, the server sends a Finished message which includes an integrity check over the handshake keyed by the shared secret and demonstrates that the server and client have agreed upon the same keys. [[TODO: If the server is not requesting client authentication, it MAY start sending application data following the Finished, though the server has no way of knowing who will be receiving the data. Add this.]]

Once the client receives the ServerKeyShare, it can also compute the premaster secret and decrypt the server's remaining handshake messages. The client generates its own sending keys based on the

Internet-Draft

TLS

March 2015

premaster secret and will encrypt the remainder of its handshake messages using those keys and the newly established cipher suite. If the server has sent a CertificateRequest message, the client MUST send the Certificate message, though it may contain zero certificates. If the client has sent a certificate, a digitally-signed CertificateVerify message is sent to explicitly verify possession of the private key in the certificate. Finally, the client sends the Finished message.

At this point, the handshake is complete, and the client and server may exchange application layer data, which is protected using a new set of keys derived from both the premaster secret and the handshake transcript (see [[I-D.ietf-tls-session-hash](#)] for the security rationale for this.)

Application data MUST NOT be sent prior to the Finished message. [[TODO: can we make this clearer and more clearly match the text above about server-side False Start.]] Client Server

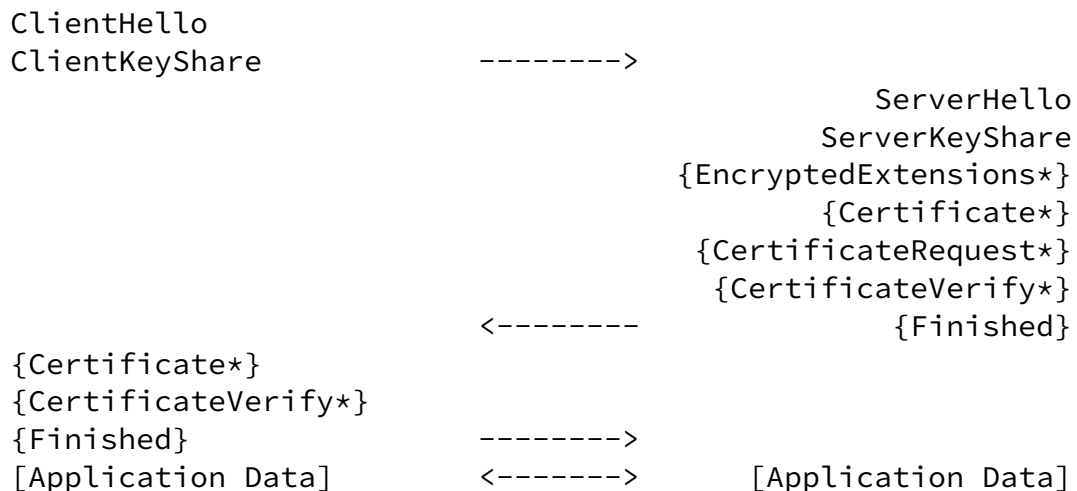


Figure 1. Message flow for a full handshake

* Indicates optional or situation-dependent messages that are not always sent.

{ } Indicates messages protected using keys derived from the handshake master secret.

[] Indicates messages protected using keys derived from the master

secret.

If the client has not provided an appropriate ClientKeyShare (e.g. it includes only DHE or ECDHE groups unacceptable or unsupported by the server), the server corrects the mismatch with a HelloRetryRequest

and the client will need to restart the handshake with an appropriate ClientKeyShare, as shown in Figure 2:

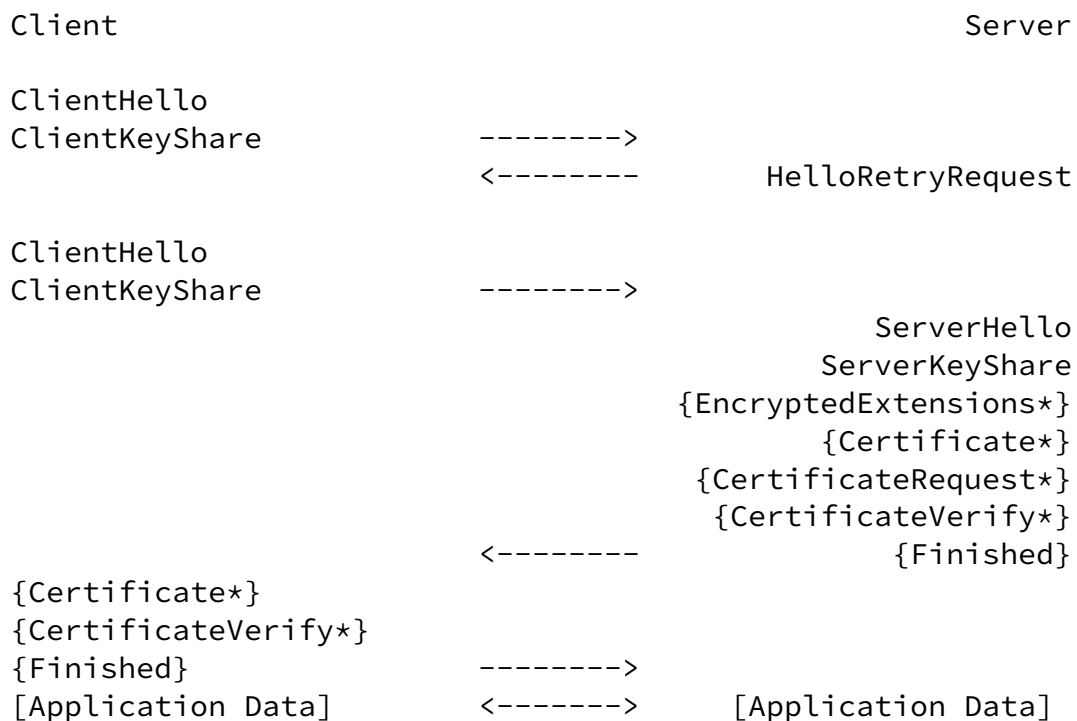


Figure 2. Message flow for a full handshake with mismatched parameters

[[OPEN ISSUE: Should we restart the handshake hash?
<https://github.com/tlswg/tls13-spec/issues/104>.]] [[OPEN ISSUE: We need to make sure that this flow doesn't introduce downgrade issues. Potential options include continuing the handshake hashes (as long as clients don't change their opinion of the server's capabilities with aborted handshakes) and requiring the client to send the same ClientHello (as is currently done) and then checking you get the same negotiated parameters.]]

If no common cryptographic parameters can be negotiated, the server

will send a fatal alert.

When the client and server decide to resume a previous session or duplicate an existing session (instead of negotiating new security parameters), the message flow is as follows:

The client sends a ClientHello using the Session ID of the session to be resumed. The server then checks its session cache for a match. If a match is found, and the server is willing to re-establish the connection under the specified session state, it will send a ServerHello with the same Session ID value. At this point, both client and server MUST proceed directly to sending Finished messages,

which are protected using handshake keys as described above, computed using resumption premaster secret created in the first handshake as the premaster secret. Once the re-establishment is complete, the client and server MAY begin to exchange application layer data, which is protected using the application secrets (See flow chart below.) If a Session ID match is not found, the server generates a new session ID, and the TLS client and server perform a full handshake.

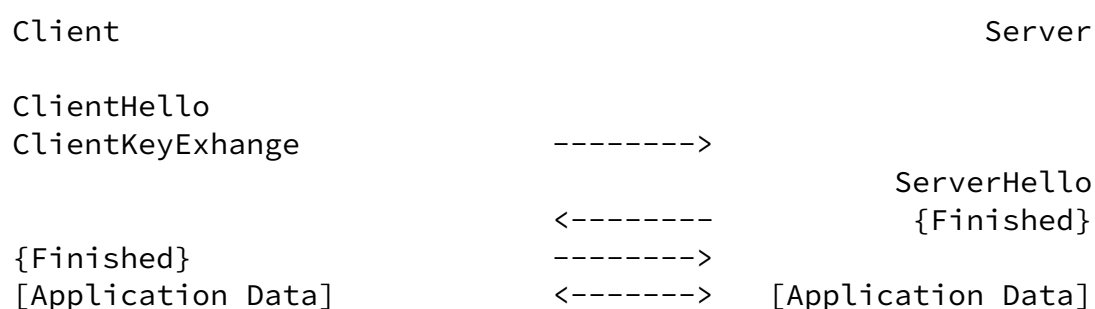


Figure 3. Message flow for an abbreviated handshake

The contents and significance of each message will be presented in detail in the following sections.

[7.3.](#) Handshake Protocol

The TLS Handshake Protocol is one of the defined higher-level clients of the TLS Record Protocol. This protocol is used to negotiate the secure attributes of a session. Handshake messages are supplied to the TLS record layer, where they are encapsulated within one or more TLSPlaintext structures, which are processed and transmitted as

specified by the current active session state.

```
enum {
    reserved(0), client_hello(1), server_hello(2),
    client_key_share(5), hello_retry_request(6),
    server_key_share(7), certificate(11), reserved(12),
    certificate_request(13), certificate_verify(15),
    reserved(16), finished(20), (255)
} HandshakeType;

struct {
    HandshakeType msg_type;      /* handshake type */
    uint24 length;              /* bytes in message */
    select (HandshakeType) {
        case client_hello:      ClientHello;
        case client_key_share:  ClientKeyShare;
        case server_hello:      ServerHello;
        case hello_retry_request: HelloRetryRequest;
        case server_key_share:  ServerKeyShare;
        case certificate:       Certificate;
        case certificate_request: CertificateRequest;
        case certificate_verify: CertificateVerify;
        case finished:          Finished;
    } body;
}
```

```
} Handshake;
```

The handshake protocol messages are presented below in the order they MUST be sent; sending handshake messages in an unexpected order results in a fatal error. Unneeded handshake messages can be omitted, however.

New handshake message types are assigned by IANA as described in [Section 12](#).

[7.3.1](#). Hello Messages

The hello phase messages are used to exchange security enhancement capabilities between the client and server. When a new session begins, the record layer's connection state AEAD algorithm is initialized to NULL_NULL.

[7.3.1.1](#). Client Hello

When this message will be sent:

When a client first connects to a server, it is required to send the ClientHello as its first message. The client will also send a ClientHello when the server has responded to its ClientHello with a ServerHello that selects cryptographic parameters that don't match the client's ClientKeyShare. In that case, the client MUST

send the same ClientHello (without modification) along with the new ClientKeyShare. If a server receives a ClientHello at any other time, it MUST send a fatal no_renegotiation alert.

Structure of this message:

The ClientHello message includes a random structure, which is used later in the protocol.

```
struct {  
    opaque random_bytes[32];  
} Random;
```

random_bytes

32 bytes generated by a secure random number generator.

Note: Versions of TLS prior to TLS 1.3 used the top 32 bits of the Random value to encode the time since the UNIX epoch.

Note: The ClientHello message includes a variable-length session identifier. If not empty, the value identifies a session between the same client and server whose security parameters the client wishes to reuse. The session identifier MAY be from an earlier connection, this connection, or from another currently active connection. The second option is useful if the client only wishes to update the random structures and derived values of a connection, and the third option makes it possible to establish several independent secure connections without repeating the full handshake protocol. These independent connections may occur sequentially or simultaneously; a SessionID becomes valid when the handshake negotiating it completes with the exchange of Finished messages and persists until it is removed due to aging or because a fatal error was encountered on a connection associated with the session. The actual contents of the SessionID are defined by the server.

opaque SessionID<0..32>;

Warning: Because the SessionID is transmitted without confidentiality or integrity protection, servers MUST NOT place confidential information in session identifiers or let the contents of fake session identifiers cause any breach of security. (Note that the content of the handshake as a whole, including the SessionID, is protected by the Finished messages exchanged at the end of the handshake.)

The cipher suite list, passed from the client to the server in the ClientHello message, contains the combinations of cryptographic algorithms supported by the client in order of the client's

preference (favorite choice first). Each cipher suite defines a key exchange algorithm, a record protection algorithm (including secret key length) and a PRF. The server will select a cipher suite or, if no acceptable choices are presented, return a handshake failure alert and close the connection. If the list contains cipher suites the server does not recognize, support, or wish to use, the server MUST ignore those cipher suites, and process the remaining ones as usual.


```

uint8 CipherSuite[2];    /* Cryptographic suite selector */

enum { null(0), (255) } CompressionMethod;

struct {
    ProtocolVersion client_version;
    Random random;
    SessionID session_id;
    CipherSuite cipher_suites<2..2^16-2>;
    CompressionMethod compression_methods<1..2^8-1>;
    select (extensions_present) {
        case false:
            struct {};
        case true:
            Extension extensions<0..2^16-1>;
    };
} ClientHello;

```

TLS allows extensions to follow the `compression_methods` field in an extensions block. The presence of extensions can be detected by determining whether there are bytes following the `compression_methods` at the end of the `ClientHello`. Note that this method of detecting optional data differs from the normal TLS method of having a variable-length field, but it is used for compatibility with TLS before extensions were defined.

`client_version`

The version of the TLS protocol by which the client wishes to communicate during this session. This SHOULD be the latest (highest valued) version supported by the client. For this version of the specification, the version will be 3.4 (see [Appendix E](#) for details about backward compatibility).

`random`

A client-generated random structure.

`session_id`

The ID of a session the client wishes to use for this connection. This field is empty if no `session_id` is available, or if the client wishes to generate new security parameters.

This is a list of the cryptographic options supported by the client, with the client's first preference first. If the session_id field is not empty (implying a session resumption request), this vector MUST include at least the cipher_suite from that session. Values are defined in [Appendix A.4](#).

compression_methods

Versions of TLS before 1.3 supported compression and the list of compression methods was supplied in this field. For any TLS 1.3 ClientHello, this field MUST contain only the "null" compression method with the code point of 0. If a TLS 1.3 ClientHello is received with any other value in this field, the server MUST generate a fatal "illegal_parameter" alert. Note that TLS 1.3 servers may receive TLS 1.2 or prior ClientHellos which contain other compression methods and MUST follow the procedures for the appropriate prior version of TLS.

extensions

Clients MAY request extended functionality from servers by sending data in the extensions field. The actual "Extension" format is defined in [Section 7.3.2.5](#).

In the event that a client requests additional functionality using extensions, and this functionality is not supplied by the server, the client MAY abort the handshake. A server MUST accept ClientHello messages both with and without the extensions field, and (as for all other messages) it MUST check that the amount of data in the message precisely matches one of these formats; if not, then it MUST send a fatal "decode_error" alert.

After sending the ClientHello message, the client waits for a ServerHello or HelloRetryRequest message.

[7.3.2](#). Client Key Share Message

When this message will be sent:

This message is always sent by the client. It MUST immediately follow the ClientHello message. In backward compatibility mode (see Section XXX) it will be included in the EarlyData extension ([Section 7.3.2.5.3](#)) in the ClientHello.

Meaning of this message:

This message contains the client's cryptographic parameters for zero or more key establishment methods.

Structure of this message:

```
struct {  
    NamedGroup group;  
    opaque key_exchange<1..2^16-1>;  
} ClientKeyShareOffer;
```

group The named group for the key share offer. This identifies the specific key exchange method that the ClientKeyShareOffer describes. Finite Field Diffie-Hellman parameters are described in [Section 7.3.2.1](#); Elliptic Curve Diffie-Hellman parameters are described in [Section 7.3.2.2](#).

key_exchange Key exchange information. The contents of this field are determined by the value of NamedGroup entry and its corresponding definition.

```
struct {  
    ClientKeyShareOffer offers<0..2^16-1>;  
} ClientKeyShare;
```

offers

A list of ClientKeyShareOffer values.

Clients may offer an arbitrary number of ClientKeyShareOffer values, each representing a single set of key agreement parameters; for instance a client might offer shares for several elliptic curves or multiple integer DH groups. The shares for each ClientKeyShareOffer MUST be generated independently. Clients MUST NOT offer multiple ClientKeyShareOffers for the same parameters. It is explicitly permitted to send an empty ClientKeyShare message, as this is used to elicit the server's parameters if the client has no useful information. [TODO: Recommendation about what the client offers. Presumably which integer DH groups and which curves.] [TODO: Work out how this interacts with PSK and SRP.]

[7.3.2.1](#). Diffie-Hellman Parameters

Diffie-Hellman parameters for both clients and servers are encoded in the opaque key_exchange field of the ClientKeyShareOffer or ServerKeyShare structures. The opaque value contains the Diffie-Hellman public value ($dh_Y = g^X \bmod p$), encoded as a big-endian integer.

```
opaque dh_Y<1..2^16-1>;
```

Internet-Draft

TLS

March 2015

[7.3.2.2](#). ECHDE Parameters

ECDHE parameters for both clients and servers are encoded in the opaque `key_exchange` field of the `ClientKeyShareOffer` or `ServerKeyShare` structures. The opaque value conveys the Elliptic Curve Diffie-Hellman public value (`ecdh_Y`) represented as a byte string `ECPPoint.point`.

opaque point <1..2⁸-1>;

point

This is the byte string representation of an elliptic curve point following the conversion routine in [Section 4.3.6](#) of ANSI X9.62 {{X962}}.

Although X9.62 supports multiple point formats, any given curve MUST specify only a single point format. All curves currently specified in this document MUST only be used with the uncompressed point format.

Note: Versions of TLS prior to 1.3 permitted point negotiation; TLS 1.3 removes this feature in favor of a single point format for each curve.

[[OPEN ISSUE: We will need to adjust the compressed/uncompressed point issue if we have new curves that don't need point compression. This depends on the CFRG's recommendations. The expectation is that future curves will come with defined point formats and that existing curves conform to X9.62.]]

[7.3.2.3](#). Server Hello

When this message will be sent:

The server will send this message in response to a `ClientHello` message when it was able to find an acceptable set of algorithms and the client's `ClientKeyShare` message was acceptable. If the client proposed groups are not acceptable by the server, it will respond with an `insufficient_security` fatal alert.

Structure of this message:

Rescorla

Expires September 10, 2015

[Page 39]

Internet-Draft

TLS

March 2015

```
struct {
    ProtocolVersion server_version;
    Random random;
    SessionID session_id;
    CipherSuite cipher_suite;
    select (extensions_present) {
        case false:
            struct {};
        case true:
            Extension extensions<0..2^16-1>;
    };
} ServerHello;
```

The presence of extensions can be detected by determining whether there are bytes following the cipher_suite field at the end of the ServerHello.

server_version

This field will contain the lower of that suggested by the client in the client hello and the highest supported by the server. For this version of the specification, the version is 3.4. (See [Appendix E](#) for details about backward compatibility.)

random

This structure is generated by the server and MUST be generated independently of the ClientHello.random.

session_id

This is the identity of the session corresponding to this connection. If the ClientHello.session_id was non-empty, the server will look in its session cache for a match. If a match is found and the server is willing to establish the new connection

using the specified session state, the server will respond with the same value as was supplied by the client. This indicates a resumed session and dictates that the parties must proceed directly to the Finished messages. Otherwise, this field will contain a different value identifying the new session. The server may return an empty session_id to indicate that the session will not be cached and therefore cannot be resumed. If a session is resumed, it must be resumed using the same cipher suite it was originally negotiated with. Note that there is no requirement that the server resume any session even if it had formerly provided a session_id. Clients MUST be prepared to do a full negotiation -- including negotiating new cipher suites -- during any handshake.

cipher_suite

Rescorla

Expires September 10, 2015

[Page 40]

Internet-Draft

TLS

March 2015

The single cipher suite selected by the server from the list in ClientHello.cipher_suites. For resumed sessions, this field is the value from the state of the session being resumed.

extensions

A list of extensions. Note that only extensions offered by the client can appear in the server's list. In TLS 1.3 as opposed to previous versions of TLS, the server's extensions are split between the ServerHello and the EncryptedExtensions [Section 7.3.4](#) message. The ServerHello MUST only include extensions which are required to establish the cryptographic context.

[7.3.2.4](#). HelloRetryRequest

When this message will be sent:

The server will send this message in response to a ClientHello message when it was able to find an acceptable set of algorithms but the client's ClientKeyShare message did not contain an acceptable offer. If it cannot find such a match, it will respond with a handshake failure alert.

Structure of this message:

```
struct {
```

```

    ProtocolVersion server_version;
    CipherSuite cipher_suite;
    NamedGroup selected_group;
    Extension extensions<0..2^16-1>;
} HelloRetryRequest;

```

[[OPEN ISSUE: Merge in DTLS Cookies?]]

selected_group

The group which the client MUST use for its new ClientHello.

The "server_version", "cipher_suite" and "extensions" fields have the same meanings as their corresponding values in the ServerHello. The server SHOULD send only the extensions necessary for the client to generate a correct ClientHello/ClientKeyShare pair.

Upon receipt of a HelloRetryRequest, the client MUST send a new ClientHello/ClientKeyShare pair to the server. The ClientKeyShare MUST contain both the groups in the original ClientKeyShare as well as a ClientKeyShareOffer consistent with the "selected_group" field. I.e., it MUST be a superset of the previous ClientKeyShareOffer.

Upon re-sending the ClientHello/ClientKeyShare and receiving the server's ServerHello/ServerKeyShare, the client MUST verify that the selected ciphersuite and NamedGroup match that supplied in the HelloRetryRequest.

[7.3.2.5.](#) Hello Extensions

The extension format is:

```

struct {
    ExtensionType extension_type;
    opaque extension_data<0..2^16-1>;
} Extension;

enum {
    signature_algorithms(13), early_data(TBD), (65535)
} ExtensionType;

```

Here:

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

The initial set of extensions is defined in a companion document [[TLSEXT](#)]. The list of extension types is maintained by IANA as described in [Section 12](#).

An extension type MUST NOT appear in the ServerHello unless the same extension type appeared in the corresponding ClientHello. If a client receives an extension type in ServerHello that it did not request in the associated ClientHello, it MUST abort the handshake with an `unsupported_extension` fatal alert.

Nonetheless, "server-oriented" extensions may be provided in the future within this framework. Such an extension (say, of type x) would require the client to first send an extension of type x in a ClientHello with empty `extension_data` to indicate that it supports the extension type. In this case, the client is offering the capability to understand the extension type, and the server is taking the client up on its offer.

When multiple extensions of different types are present in the ClientHello or ServerHello messages, the extensions MAY appear in any order. There MUST NOT be more than one extension of the same type.

Finally, note that extensions can be sent both when starting a new session and when requesting session resumption. Indeed, a client that requests session resumption does not in general know whether the server will accept this request, and therefore it SHOULD send the same extensions as it would send if it were not attempting resumption.

In general, the specification of each extension type needs to describe the effect of the extension both during full handshake and session resumption. Most current TLS extensions are relevant only when a session is initiated: when an older session is resumed, the

server does not process these extensions in Client Hello, and does not include them in Server Hello. However, some extensions may specify different behavior during session resumption.

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

- Some cases where a server does not agree to an extension are error conditions, and some are simply refusals to support particular features. 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 cipher suite 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.

[7.3.2.5.1](#). Signature Algorithms

The client uses the "signature_algorithms" extension to indicate to the server which signature/hash algorithm pairs may be used in

digital signatures. The "extension_data" field of this extension contains a "supported_signature_algorithms" value.

```
enum {
    none(0), md5(1), sha1(2), sha224(3), sha256(4), sha384(5),
    sha512(6), (255)
} HashAlgorithm;

enum { anonymous(0), rsa(1), dsa(2), ecdsa(3), (255) }
    SignatureAlgorithm;

struct {
    HashAlgorithm hash;
    SignatureAlgorithm signature;
} SignatureAndHashAlgorithm;

SignatureAndHashAlgorithm
    supported_signature_algorithms<2..2^16-2>;
```

Each SignatureAndHashAlgorithm value lists a single hash/signature pair that the client is willing to verify. The values are indicated in descending order of preference.

Note: Because not all signature algorithms and hash algorithms may be accepted by an implementation (e.g., DSA with SHA-1, but not SHA-256), algorithms here are listed in pairs.

hash

This field indicates the hash algorithm which may be used. The values indicate support for unhashed data, MD5 [[RFC1321](#)], SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512 [[SHS](#)], respectively. The "none" value is provided for future extensibility, in case of a signature algorithm which does not require hashing before signing.

signature

This field indicates the signature algorithm that may be used. The values indicate anonymous signatures, RSASSA-PKCS1-v1_5 [[RFC3447](#)] and DSA [[DSS](#)], and ECDSA [[ECDSA](#)], respectively. The "anonymous" value is meaningless in this context but used in [Section 7.3.3](#). It MUST NOT appear in this extension.

The semantics of this extension are somewhat complicated because the cipher suite indicates permissible signature algorithms but not hash

algorithms. [Section 7.3.5](#) and [Section 7.3.3](#) describe the appropriate rules.

If the client supports only the default hash and signature algorithms (listed in this section), it MAY omit the `signature_algorithms` extension. If the client does not support the default algorithms, or supports other hash and signature algorithms (and it is willing to use them for verifying messages sent by the server, i.e., server certificates and server key share), it MUST send the `signature_algorithms` extension, listing the algorithms it is willing to accept.

If the client does not send the `signature_algorithms` extension, the server MUST do the following:

- If the negotiated key exchange algorithm is one of (DHE_RSA, ECDHE_RSA), behave as if client had sent the value {sha1,rsa}.
- If the negotiated key exchange algorithm is DHE_DSS, behave as if the client had sent the value {sha1,dsa}.
- If the negotiated key exchange algorithm is ECDHE_ECDSA, behave as if the client had sent value {sha1,ecdsa}.

Note: this is a change from TLS 1.1 where there are no explicit rules, but as a practical matter one can assume that the peer supports MD5 and SHA-1.

Note: this extension is not meaningful for TLS versions prior to 1.2. Clients MUST NOT offer it if they are offering prior versions. However, even if clients do offer it, the rules specified in [\[TLSEXT\]](#) require servers to ignore extensions they do not understand.

Servers MUST NOT send this extension. TLS servers MUST support receiving this extension.

When performing session resumption, this extension is not included in Server Hello, and the server ignores the extension in Client Hello (if present).

[7.3.2.5.2](#). Negotiated Groups

When sent by the client, the "supported_groups" extension indicates the named groups which the client supports, ordered from most preferred to least preferred.

Internet-Draft

TLS

March 2015

Note: In versions of TLS prior to TLS 1.3, this extension was named "elliptic curves" and only contained elliptic curve groups. See [\[RFC4492\]](#) and [\[I-D.ietf-tls-negotiated-ff-dhe\]](#).

The "extension_data" field of this extension SHALL contain a "NamedGroupList" value:

```
enum {  
    // Elliptic Curve Groups.  
    sect163k1 (1), sect163r1 (2), sect163r2 (3),  
    sect193r1 (4), sect193r2 (5), sect233k1 (6),  
    sect233r1 (7), sect239k1 (8), sect283k1 (9),  
    sect283r1 (10), sect409k1 (11), sect409r1 (12),  
    sect571k1 (13), sect571r1 (14), secp160k1 (15),  
    secp160r1 (16), secp160r2 (17), secp192k1 (18),  
    secp192r1 (19), secp224k1 (20), secp224r1 (21),  
    secp256k1 (22), secp256r1 (23), secp384r1 (24),  
    secp521r1 (25),  
  
    // Finite Field Groups.  
    ffdhe2048(256), ffdhe3072(257), ffdhe4096(258),  
    ffdhe8192(259),  
  
    // Reserved Code Points.  
    reserved (0xFE00..0xFEFF),  
    reserved(0xFF01),  
    reserved(0xFF02),  
    (0xFFFF)  
} NamedGroup;  
  
struct {  
    NamedGroup named_group_list<1..2^16-1>  
} NamedGroupList;
```

sect163k1, etc

Indicates support of the corresponding named curve. The named curves defined here are those specified in SEC 2 [13]. Note that many of these curves are also recommended in ANSI X9.62 [\[X962\]](#) and FIPS 186-2 [\[DSS\]](#). Values 0xFE00 through 0xFEFF are reserved for private use. Values 0xFF01 and 0xFF02 were used in previous

versions of TLS but MUST NOT be offered by TLS 1.3 implementations. [[OPEN ISSUE: Triage curve list.]]

ffdhe2432, etc

Indicates support of the corresponding finite field group, defined in [[I-D.ietf-tls-negotiated-ff-dhe](#)]

Rescorla

Expires September 10, 2015

[Page 46]

Internet-Draft

TLS

March 2015

Items in `named_curve_list` are ordered according to the client's preferences (favorite choice first).

As an example, a client that only supports `secp192r1` (aka NIST P-192; value 19 = 0x0013) and `secp224r1` (aka NIST P-224; value 21 = 0x0015) and prefers to use `secp192r1` would include a TLS extension consisting of the following octets. Note that the first two octets indicate the extension type (Supported Group Extension):

```
00 0A 00 06 00 04 00 13 00 15
```

The client MUST supply a "named_groups" extension containing at least one group for each key exchange algorithm (currently DHE and ECDHE) for which it offers a cipher suite. If the client does not supply a "named_groups" extension with a compatible group, the server MUST NOT negotiate a cipher suite of the relevant type. For instance, if a client supplies only ECDHE groups, the server MUST NOT negotiate finite field Diffie-Hellman. If no acceptable group can be selected across all cipher suites, then the server MUST generate a fatal "handshake_failure" alert.

NOTE: A server participating in an ECDHE-ECDSA key exchange may use different curves for (i) the ECDSA key in its certificate, and (ii) the ephemeral ECDH key in the ServerKeyExchange message. The server must consider the supported groups in both cases.

[[TODO: IANA Considerations.]]

[7.3.2.5.3](#). Early Data Extension

TLS versions before 1.3 have a strict message ordering and do not permit additional messages to follow the ClientHello. The EarlyData extension allows TLS messages which would otherwise be sent as

separate records to be instead inserted in the ClientHello. The extension simply contains the TLS records which would otherwise have been included in the client's first flight.

```
struct {  
    TLSCipherText messages<5 .. 2^24-1>;  
} EarlyDataExtension;
```

Extra messages for the client's first flight MAY either be transmitted standalone or sent as EarlyData. However, when a client does not know whether TLS 1.3 can be negotiated - e.g., because the server may support a prior version of TLS or because of network intermediaries - it SHOULD use the EarlyData extension. If the EarlyData extension is used, then clients MUST NOT send any messages other than the ClientHello in their initial flight.

Any data included in EarlyData is not integrated into the handshake hashes directly. E.g., if the ClientKeyShare is included in EarlyData, then the handshake hashes consist of ClientHello + ServerHello, etc. However, because the ClientKeyShare is in a ClientHello extension, it is still hashed transitively. This procedure guarantees that the Finished message covers these messages even if they are ultimately ignored by the server (e.g., because it is sent to a TLS 1.2 server). TLS 1.3 servers MUST understand messages sent in EarlyData, and aside from hashing them differently, MUST treat them as if they had been sent immediately after the ClientHello.

Servers MUST NOT send the EarlyData extension. Negotiating TLS 1.3 serves as acknowledgement that it was processed as described above.

[[OPEN ISSUE: This is a fairly general mechanism which is possibly overkill in the 1-RTT case, where it would potentially be more attractive to just have a "ClientKeyShare" extension. However, for the 0-RTT case we will want to send the Certificate, CertificateVerify, and application data, so a more general extension seems appropriate at least until we have determined we don't need it for 0-RTT.]]

[7.3.3.](#) Server Key Share Message

When this message will be sent:

This message will be sent immediately after the ServerHello message if the client has provided a ClientKeyShare message which is compatible with the selected cipher suite and group parameters.

Meaning of this message:

This message conveys cryptographic information to allow the client to compute the premaster secret: a Diffie-Hellman public key with which the client can complete a key exchange (with the result being the premaster secret) or a public key for some other algorithm.

Structure of this message:

```
struct {  
    NamedGroup group;  
    opaque key_exchange<1..216-1>;  
} ServerKeyShare;
```

group The named group for the key share offer. This identifies the selected key exchange method from the ClientKeyShare message

([Section 7.3.2](#)), identifying which value from the ClientKeyShareOffer the server has accepted as is responding to.

key_exchange Key exchange information. The contents of this field are determined by the value of NamedGroup entry and its corresponding definition.

[7.3.4](#). Encrypted Extensions

When this message will be sent:

If this message is sent, it MUST be sent immediately after the server's ServerKeyShare.

Meaning of this message:

The EncryptedExtensions message simply contains any extensions which should be protected, i.e., any which are not needed to establish the cryptographic context. The same extension types

MUST NOT appear in both the ServerHello and EncryptedExtensions. If the same extension appears in both locations, the client MUST rely only on the value in the EncryptedExtensions block. [[OPEN ISSUE: Should we just produce a canonical list of what goes where and have it be an error to have it in the wrong place? That seems simpler. Perhaps have a whitelist of which extensions can be unencrypted and everything else MUST be encrypted.]]

Structure of this message:

```
struct {  
    Extension extensions<0..2^16-1>;  
} EncryptedExtensions;
```

extensions
A list of extensions.

[7.3.5.](#) Server Certificate

When this message will be sent:

The server MUST send a Certificate message whenever the agreed-upon key exchange method uses certificates for authentication (this includes all key exchange methods defined in this document except DH_anon). This message will always immediately follow either the EncryptedExtensions message if one is sent or the ServerKeyShare message.

Meaning of this message:

Rescorla	Expires September 10, 2015	[Page 49]
----------	----------------------------	-----------

Internet-Draft	TLS	March 2015
----------------	-----	------------

This message conveys the server's certificate chain to the client.

The certificate MUST be appropriate for the negotiated cipher suite's key exchange algorithm and any negotiated extensions.

Structure of this message:

```
opaque ASN1Cert<1..2^24-1>;  
  
struct {  
    ASN1Cert certificate_list<0..2^24-1>;  
} Certificate;
```


certificate_list

This is a sequence (chain) of certificates. The sender's certificate MUST come first in the list. Each following certificate MUST directly certify the one preceding it. Because certificate validation requires that root keys be distributed independently, the self-signed certificate that specifies the root certificate authority MAY be omitted from the chain, under the assumption that the remote end must already possess it in order to validate it in any case.

The same message type and structure will be used for the client's response to a certificate request message. Note that a client MAY send no certificates if it does not have an appropriate certificate to send in response to the server's authentication request.

Note: PKCS #7 [[PKCS7](#)] is not used as the format for the certificate vector because PKCS #6 [[PKCS6](#)] extended certificates are not used. Also, PKCS #7 defines a SET rather than a SEQUENCE, making the task of parsing the list more difficult.

The following rules apply to the certificates sent by the server:

- The certificate type MUST be X.509v3, unless explicitly negotiated otherwise (e.g., [[RFC5081](#)]).
- The end entity certificate's public key (and associated restrictions) MUST be compatible with the selected key exchange algorithm.

Key Exchange Alg. Certificate Key Type

DHE_RSA	RSA public key; the certificate MUST allow the
ECDHE_RSA	key to be used for signing (the
	digitalSignature bit MUST be set if the key

usage extension is present) with the signature scheme and hash algorithm that will be employed in the server key exchange message.

Note: ECDHE_RSA is defined in [[RFC4492](#)].

DHE_DSS DSA public key; the certificate MUST allow the key to be used for signing with the hash algorithm that will be employed in the server key exchange message.

ECDHE_ECDSA ECDSA-capable public key; the certificate MUST allow the key to be used for signing with the hash algorithm that will be employed in the server key exchange message. The public key MUST use a curve and point format supported by the client, as described in [[RFC4492](#)].

- The "server_name" and "trusted_ca_keys" extensions [[TLSEXT](#)] are used to guide certificate selection.

If the client provided a "signature_algorithms" extension, then all certificates provided by the server MUST be signed by a hash/signature algorithm pair that appears in that extension. Note that this implies that a certificate containing a key for one signature algorithm MAY be signed using a different signature algorithm (for instance, an RSA key signed with a DSA key). This is a departure from TLS 1.1, which required that the algorithms be the same.

If the server has multiple certificates, it chooses one of them based on the above-mentioned criteria (in addition to other criteria, such as transport layer endpoint, local configuration and preferences, etc.). If the server has a single certificate, it SHOULD attempt to validate that it meets these criteria.

Note that there are certificates that use algorithms and/or algorithm combinations that cannot be currently used with TLS. For example, a certificate with RSASSA-PSS signature key (id-RSASSA-PSS OID in SubjectPublicKeyInfo) cannot be used because TLS defines no corresponding signature algorithm.

As cipher suites that specify new key exchange methods are specified for the TLS protocol, they will imply the certificate format and the required encoded keying information.

7.3.6. Certificate Request

When this message will be sent:

A non-anonymous server can optionally request a certificate from the client, if appropriate for the selected cipher suite. This message, if sent, will immediately follow the server's Certificate message).

Structure of this message:

```
enum {  
    rsa_sign(1), dss_sign(2), rsa_fixed_dh(3), dss_fixed_dh(4),  
    rsa_ephemeral_dh_RESERVED(5), dss_ephemeral_dh_RESERVED(6),  
    fortezza_dms_RESERVED(20), (255)  
} ClientCertificateType;
```

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

```
struct {  
    ClientCertificateType certificate_types<1..28-1>;  
    SignatureAndHashAlgorithm  
        supported_signature_algorithms<2..216-2>;  
    DistinguishedName certificate_authorities<0..216-1>;  
} CertificateRequest;
```

certificate_types

A list of the types of certificate types that the client may offer.

rsa_sign	a certificate containing an RSA key
dss_sign	a certificate containing a DSA key
rsa_fixed_dh	a certificate containing a static DH key.
dss_fixed_dh	a certificate containing a static DH key

supported_signature_algorithms

A list of the hash/signature algorithm pairs that the server is able to verify, listed in descending order of preference.

certificate_authorities

A list of the distinguished names [\[X501\]](#) of acceptable certificate_authorities, represented in DER-encoded format. These distinguished names may specify a desired distinguished name for a root CA or for a subordinate CA; thus, this message can be used to describe known roots as well as a desired authorization space. If the certificate_authorities list is empty, then the client MAY send any certificate of the appropriate ClientCertificateType, unless there is some external arrangement to the contrary.

Internet-Draft

TLS

March 2015

The interaction of the `certificate_types` and `supported_signature_algorithms` fields is somewhat complicated. `certificate_types` has been present in TLS since SSLv3, but was somewhat underspecified. Much of its functionality is superseded by `supported_signature_algorithms`. The following rules apply:

- Any certificates provided by the client MUST be signed using a hash/signature algorithm pair found in `supported_signature_algorithms`.
- The end-entity certificate provided by the client MUST contain a key that is compatible with `certificate_types`. If the key is a signature key, it MUST be usable with some hash/signature algorithm pair in `supported_signature_algorithms`.
- For historical reasons, the names of some client certificate types include the algorithm used to sign the certificate. For example, in earlier versions of TLS, `rsa_fixed_dh` meant a certificate signed with RSA and containing a static DH key. In TLS 1.2, this functionality has been obsoleted by the `supported_signature_algorithms`, and the certificate type no longer restricts the algorithm used to sign the certificate. For example, if the server sends `dss_fixed_dh` certificate type and `{{sha1, dsa}, {sha1, rsa}}` signature types, the client MAY reply with a certificate containing a static DH key, signed with RSA-SHA1.

New `ClientCertificateType` values are assigned by IANA as described in [Section 12](#).

Note: Values listed as `RESERVED` may not be used. They were used in SSLv3.

Note: It is a fatal `handshake_failure` alert for an anonymous server to request client authentication.

[7.3.7](#). Server Certificate Verify

When this message will be sent:

This message is used to provide explicit proof that the server possesses the private key corresponding to its certificate.

certificate and also provides integrity for the handshake up to this point. This message is only sent when the server is authenticated via a certificate. When sent, it MUST be the last server handshake message prior to the Finished.

Structure of this message:

```
struct {  
    digitally-signed struct {  
        opaque handshake_messages_hash[hash_length];  
    }  
} CertificateVerify;
```

Here `handshake_messages_hash` is a digest of all handshake messages sent or received, starting at `ClientHello` and up to, but not including, this message, including the type and length fields of the handshake messages. This is a digest of the concatenation of all the Handshake structures (as defined in [Section 7.3](#)) exchanged thus far. For the PRF defined in [Section 5](#), the digest MUST be the Hash used as the basis for the PRF. Any cipher suite which defines a different PRF MUST also define the Hash to use in this computation. Note that this is the same running hash that is used in the Finished message [Section 7.3.8](#).

The context string for the signature is "TLS 1.3, server CertificateVerify". A hash of the handshake messages is signed rather than the messages themselves because the digitally-signed format requires padding and context bytes at the beginning of the input. Thus, by signing a digest of the messages, an implementation need only maintain one running hash per hash type for CertificateVerify, Finished and other messages.

If the client has offered the "signature_algorithms" extension, the signature algorithm and hash algorithm MUST be a pair listed in that extension. Note that there is a possibility for inconsistencies here. For instance, the client might offer DHE_DSS key exchange but omit any DSA pairs from its "signature_algorithms" extension. In order to negotiate correctly, the server MUST check any candidate cipher suites against the "signature_algorithms" extension before selecting them. This is somewhat inelegant but is a compromise designed to minimize changes to the original cipher suite design.

In addition, the hash and signature algorithms MUST be compatible with the key in the server's end-entity certificate. RSA keys MAY be used with any permitted hash algorithm, subject to restrictions in the certificate, if any.

Because DSA signatures do not contain any secure indication of hash algorithm, there is a risk of hash substitution if multiple hashes may be used with any key. Currently, DSA [DSS] may only be used with SHA-1. Future revisions of DSS [DSS-3] are expected to allow the use of other digest algorithms with DSA, as well as guidance as to which digest algorithms should be used with each key size. In addition, future revisions of [RFC3280] may specify

mechanisms for certificates to indicate which digest algorithms are to be used with DSA. [[TODO: Update this to deal with DSS-3 and DSS-4. <https://github.com/tlswg/tls13-spec/issues/59>]]

[7.3.8.](#) Server Finished

When this message will be sent:

The Server's Finished message is the final message sent by the server and indicates that the key exchange and authentication processes were successful.

Meaning of this message:

Recipients of Finished messages MUST verify that the contents are correct. Once a side has sent its Finished message and received and validated the Finished message from its peer, it may begin to send and receive application data over the connection. This data will be protected under keys derived from the `hs_master_secret` (see [Section 8](#)).

Structure of this message:

```
struct {  
    opaque verify_data[verify_data_length];  
} Finished;
```

verify_data

```
PRF(hs_master_secret, finished_label, Hash(handshake_messages))
[0..verify_data_length-1];
```

finished_label

For Finished messages sent by the client, the string "client finished". For Finished messages sent by the server, the string "server finished".

Hash denotes a Hash of the handshake messages. For the PRF defined in [Section 5](#), the Hash MUST be the Hash used as the basis for the PRF. Any cipher suite which defines a different PRF MUST also define the Hash to use in the Finished computation.

In previous versions of TLS, the verify_data was always 12 octets long. In the current version of TLS, it depends on the cipher suite. Any cipher suite which does not explicitly specify verify_data_length has a verify_data_length equal to 12. This includes all existing cipher suites. Note that this representation has the same encoding as with previous versions.

Future cipher suites MAY specify other lengths but such length MUST be at least 12 bytes.

handshake_messages

All of the data from all messages in this handshake (not including any HelloRequest messages) up to, but not including, this message. This is only data visible at the handshake layer and does not include record layer headers. This is the concatenation of all the Handshake structures as defined in [Section 7.3](#), exchanged thus far.

The value handshake_messages includes all handshake messages starting at ClientHello up to, but not including, this Finished message. This may be different from handshake_messages in [Section 7.3.7](#) or [Section 7.3.10](#). Also, the handshake_messages for the Finished message sent by the client will be different from that for the Finished message sent by the server, because the one that is sent second will include the prior one.

Note: Alerts and any other record types are not handshake messages and are not included in the hash computations. Also, HelloRequest

messages are omitted from handshake hashes.

[7.3.9](#). Client Certificate

When this message will be sent:

This message is the first handshake message the client can send after receiving the server's Finished. This message is only sent if the server requests a certificate. If no suitable certificate is available, the client **MUST** send a certificate message containing no certificates. That is, the `certificate_list` structure has a length of zero. If the client does not send any certificates, the server **MAY** at its discretion either continue the handshake without client authentication, or respond with a fatal `handshake_failure` alert. Also, if some aspect of the certificate chain was unacceptable (e.g., it was not signed by a known, trusted CA), the server **MAY** at its discretion either continue the handshake (considering the client unauthenticated) or send a fatal alert.

Client certificates are sent using the `Certificate` structure defined in [Section 7.3.5](#).

Meaning of this message:

This message conveys the client's certificate chain to the server; the server will use it when verifying the `CertificateVerify`

message (when the client authentication is based on signing) or calculating the premaster secret (for non-ephemeral Diffie-Hellman). The certificate **MUST** be appropriate for the negotiated cipher suite's key exchange algorithm, and any negotiated extensions.

In particular:

- The certificate type **MUST** be X.509v3, unless explicitly negotiated otherwise (e.g., [[RFC5081](#)]).
- The end-entity certificate's public key (and associated restrictions) has to be compatible with the certificate types listed in `CertificateRequest`:

Client Cert. Type Certificate Key Type

rsa_sign RSA public key; the certificate MUST allow the key to be used for signing with the signature scheme and hash algorithm that will be employed in the certificate verify message.

dss_sign DSA public key; the certificate MUST allow the key to be used for signing with the hash algorithm that will be employed in the certificate verify message.

ecdsa_sign ECDSA-capable public key; the certificate MUST allow the key to be used for signing with the hash algorithm that will be employed in the certificate verify message; the public key MUST use a curve and point format supported by the server.

rsa_fixed_dh
dss_fixed_dh Diffie-Hellman public key; MUST use the same parameters as server's key.

rsa_fixed_ecdh
ecdsa_fixed_ecdh ECDH-capable public key; MUST use the same curve as the server's key, and MUST use a point format supported by the server.

- If the certificate_authorities list in the certificate request message was non-empty, one of the certificates in the certificate chain SHOULD be issued by one of the listed CAs.

- The certificates MUST be signed using an acceptable hash/signature algorithm pair, as described in [Section 7.3.6](#). Note that this relaxes the constraints on certificate-signing algorithms found in prior versions of TLS.

Note that, as with the server certificate, there are certificates that use algorithms/algorithm combinations that cannot be currently

used with TLS.

[7.3.10](#). Client Certificate Verify

When this message will be sent:

This message is used to provide explicit verification of a client certificate. This message is only sent following a client certificate that has signing capability (i.e., all certificates except those containing fixed Diffie-Hellman parameters). When sent, it MUST immediately follow the client's Certificate message. The contents of the message are computed as described in [Section 7.3.7](#), except that the context string is "TLS 1.3, client CertificateVerify".

The hash and signature algorithms used in the signature MUST be one of those present in the supported_signature_algorithms field of the CertificateRequest message. In addition, the hash and signature algorithms MUST be compatible with the key in the client's end-entity certificate. RSA keys MAY be used with any permitted hash algorithm, subject to restrictions in the certificate, if any.

Because DSA signatures do not contain any secure indication of hash algorithm, there is a risk of hash substitution if multiple hashes may be used with any key. Currently, DSA [[DSS](#)] may only be used with SHA-1. Future revisions of DSS [[DSS-3](#)] are expected to allow the use of other digest algorithms with DSA, as well as guidance as to which digest algorithms should be used with each key size. In addition, future revisions of [[RFC3280](#)] may specify mechanisms for certificates to indicate which digest algorithms are to be used with DSA.

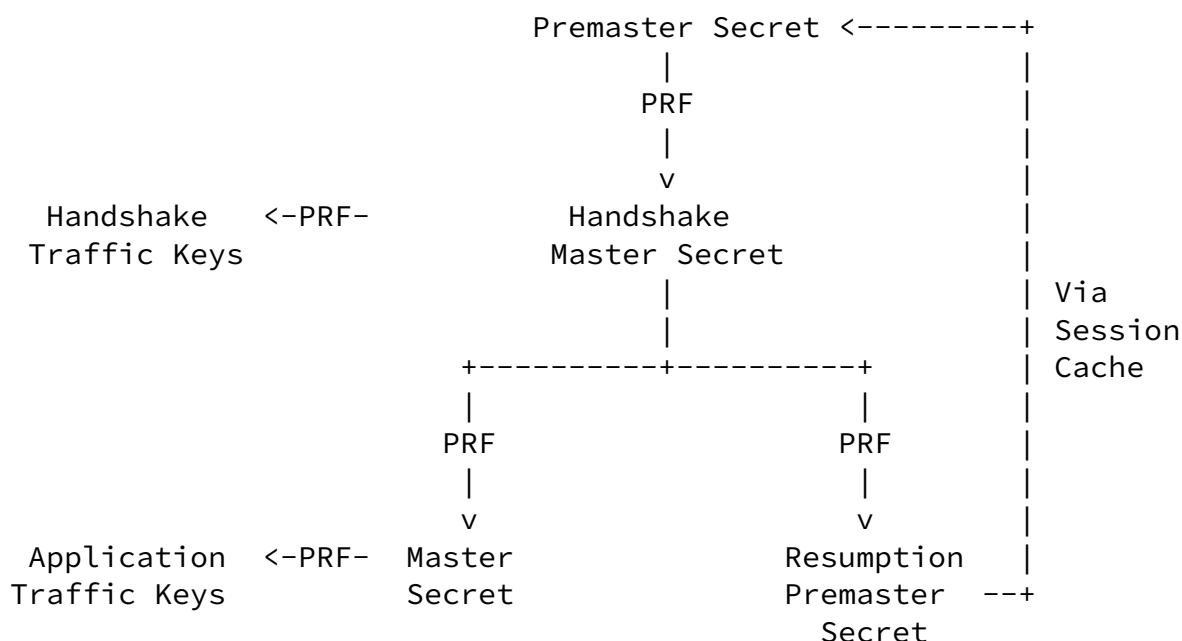
[8](#). Cryptographic Computations

In order to begin connection protection, the TLS Record Protocol requires specification of a suite of algorithms, a master secret, and the client and server random values. The authentication, key agreement, and record protection algorithms are determined by the cipher_suite selected by the server and revealed in the ServerHello

message. The random values are exchanged in the hello messages. All that remains is to calculate the master secret.

8.1. Computing the Master Secret

The `pre_master_secret` is used to generate a series of master secret values, as shown in the following diagram and described below.



First, as soon as the `ClientKeyShare` and `ServerKeyShare` messages have been exchanged, the client and server each use the unauthenticated key shares to generate a master secret which is used for the protection of the remaining handshake records. Specifically, they generate:

```

hs_master_secret = PRF(pre_master_secret, "handshake master secret",
                        session_hash
                        [0..47]);
  
```

During resumption, the premaster secret is initialized with the "resumption premaster secret", rather than using the values from the `ClientKeyShare/ServerKeyShare` exchange.

This master secret value is used to generate the record protection keys used for the handshake, as described in [Section 6.3](#).

Once the `hs_master_secret` has been computed, the premaster secret **SHOULD** be deleted from memory.

Once the last non-Finished message has been sent, the client and server then compute the master secret which will be used for the

Internet-Draft

TLS

March 2015

remainder of the session. It is also used with TLS Exporters [\[RFC5705\]](#).

```
master_secret = PRF(hs_master_secret, "extended master secret",  
                    session_hash  
                    [0..47]);
```

If the server does not request client authentication, the master secret can be computed at the time that the server sends its Finished, thus allowing the server to send traffic on its first flight (see [TODO] for security considerations on this practice.) If the server requests client authentication, this secret can be computed after the client's Certificate and CertificateVerify have been sent, or, if the client refuses client authentication, after the client's empty Certificate message has been sent.

For full handshakes, each side also derives a new secret which will be used as the premaster_secret for future resumptions of the newly established session. This is computed as:

```
resumption_premaster_secret = PRF(hs_master_secret,  
                                   "resumption premaster secret",  
                                   session_hash  
                                   [0..47]);
```

The session_hash value is a running hash of the handshake as defined in [Section 8.1.1](#). Thus, the hs_master_secret is generated using a different session_hash from the other two secrets.

All master secrets are always exactly 48 bytes in length. The length of the premaster secret will vary depending on key exchange method.

[8.1.1](#). The Session Hash

When a handshake takes place, we define

```
session_hash = Hash(handshake_messages)
```

where "handshake_messages" refers to all handshake messages sent or received, starting at client hello up to the present time, with the exception of the Finished message, including the type and length fields of the handshake messages. This is the concatenation of all the exchanged Handshake structures.

For concreteness, at the point where the handshake master secret is derived, the session hash includes the ClientHello, ClientKeyShare, ServerHello, and ServerKeyShare, and HelloRetryRequest (if any) (though see [<https://github.com/tlswg/tls13-spec/issues/104>]). At

the point where the master secret is derived, it includes every handshake message, with the exception of the Finished messages. Note that if client authentication is not used, then the session hash is complete at the point when the server has sent its first flight. Otherwise, it is only complete when the client has sent its first flight, as it covers the client's Certificate and CertificateVerify.

[8.1.2.](#) Diffie-Hellman

A conventional Diffie-Hellman computation is performed. The negotiated key (Z) is used as the `pre_master_secret`, and is converted into the `master_secret`, as specified above. Leading bytes of Z that contain all zero bits are stripped before it is used as the `pre_master_secret`.

[8.1.3.](#) Elliptic Curve Diffie-Hellman

All ECDH calculations (including parameter and key generation as well as the shared secret calculation) are performed according to [6] using the ECKAS-DH1 scheme with the identity map as key derivation function (KDF), so that the premaster secret is the x-coordinate of the ECDH shared secret elliptic curve point represented as an octet string. Note that this octet string (Z in IEEE 1363 terminology) as output by FE2OSP, the Field Element to Octet String Conversion Primitive, has constant length for any given field; leading zeros found in this octet string MUST NOT be truncated.

(Note that this use of the identity KDF is a technicality. The complete picture is that ECDH is employed with a non-trivial KDF because TLS does not directly use the premaster secret for anything other than for computing the master secret.)

[9.](#) Mandatory Cipher Suites

In the absence of an application profile standard specifying otherwise, a TLS-compliant application MUST implement the cipher

suite TODO:Needs to be selected [1]. (See [Appendix A.4](#) for the definition).

[10.](#) Application Data Protocol

Application data messages are carried by the record layer and are fragmented and encrypted based on the current connection state. The messages are treated as transparent data to the record layer.

Rescorla

Expires September 10, 2015

[Page 61]

Internet-Draft

TLS

March 2015

[11.](#) Security Considerations

Security issues are discussed throughout this memo, especially in Appendices D, E, and F.

[12.](#) IANA Considerations

[[TODO: Update <https://github.com/tlswg/tls13-spec/issues/62>]]

This document uses several registries that were originally created in [\[RFC4346\]](#). IANA has updated these to reference this document. The registries and their allocation policies (unchanged from [\[RFC4346\]](#)) are listed below.

- TLS ClientCertificateType Identifiers Registry: Future values in the range 0-63 (decimal) inclusive are assigned via Standards Action [\[RFC2434\]](#). Values in the range 64-223 (decimal) inclusive are assigned via Specification Required [\[RFC2434\]](#). Values from 224-255 (decimal) inclusive are reserved for Private Use [\[RFC2434\]](#).
- TLS Cipher Suite Registry: Future values with the first byte in the range 0-191 (decimal) inclusive are assigned via Standards Action [\[RFC2434\]](#). Values with the first byte in the range 192-254 (decimal) are assigned via Specification Required [\[RFC2434\]](#). Values with the first byte 255 (decimal) are reserved for Private Use [\[RFC2434\]](#).
- TLS ContentType Registry: Future values are allocated via

Standards Action [[RFC2434](#)].

- TLS Alert Registry: Future values are allocated via Standards Action [[RFC2434](#)].
- TLS HandshakeType Registry: Future values are allocated via Standards Action [[RFC2434](#)].

This document also uses a registry originally created in [[RFC4366](#)]. IANA has updated it to reference this document. The registry and its allocation policy (unchanged from [[RFC4366](#)]) is listed below:

- TLS ExtensionType Registry: Future values are allocated via IETF Consensus [[RFC2434](#)]. IANA has updated this registry to include the signature_algorithms extension and its corresponding value (see [Section 7.3.2.5](#)).

This document also uses two registries originally created in [[RFC4492](#)]. IANA [should update/has updated] it to reference this

document. The registries and their allocation policies are listed below.

- TLS NamedCurve registry: Future values are allocated via IETF Consensus [[RFC2434](#)].
- TLS ECPointFormat Registry: Future values are allocated via IETF Consensus [[RFC2434](#)].

In addition, this document defines two new registries to be maintained by IANA:

- TLS SignatureAlgorithm Registry: The registry has been initially populated with the values described in [Section 7.3.2.5.1](#). Future values in the range 0-63 (decimal) inclusive are assigned via Standards Action [[RFC2434](#)]. Values in the range 64-223 (decimal) inclusive are assigned via Specification Required [[RFC2434](#)]. Values from 224-255 (decimal) inclusive are reserved for Private Use [[RFC2434](#)].
- TLS HashAlgorithm Registry: The registry has been initially populated with the values described in [Section 7.3.2.5.1](#). Future

values in the range 0-63 (decimal) inclusive are assigned via Standards Action [[RFC2434](#)]. Values in the range 64-223 (decimal) inclusive are assigned via Specification Required [[RFC2434](#)]. Values from 224-255 (decimal) inclusive are reserved for Private Use [[RFC2434](#)].

[13.](#) References

[13.1.](#) Normative References

- [AES] National Institute of Standards and Technology, "Specification for the Advanced Encryption Standard (AES)", NIST FIPS 197, November 2001.
- [DSS] National Institute of Standards and Technology, U.S. Department of Commerce, "Digital Signature Standard", NIST FIPS PUB 186-2, 2000.
- [RFC1321] Rivest, R., "The MD5 Message-Digest Algorithm", [RFC 1321](#), April 1992.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), February 1997.

Rescorla Expires September 10, 2015 [Page 63]

Internet-Draft TLS March 2015

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 2434](#), October 1998.
- [RFC3280] Housley, R., Polk, W., Ford, W., and D. Solo, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", [RFC 3280](#), April 2002.
- [RFC3447] Jonsson, J. and B. Kaliski, "Public-Key Cryptography Standards (PKCS) #1: RSA Cryptography Specifications Version 2.1", [RFC 3447](#), February 2003.

- [RFC5288] Salowey, J., Choudhury, A., and D. McGrew, "AES Galois Counter Mode (GCM) Cipher Suites for TLS", [RFC 5288](#), August 2008.
- [SCH] Schneier, B., "Applied Cryptography: Protocols, Algorithms, and Source Code in C, 2nd ed.", 1996.
- [SHS] National Institute of Standards and Technology, U.S. Department of Commerce, "Secure Hash Standard", NIST FIPS PUB 180-2, August 2002.
- [TRIPLEDES] National Institute of Standards and Technology, "Recommendation for the Triple Data Encryption Algorithm (TDEA) Block Cipher", NIST Special Publication 800-67, May 2004.
- [X680] ITU-T, "Information technology - Abstract Syntax Notation One (ASN.1): Specification of basic notation", ISO/IEC 8824-1:2002, 2002.
- [X690] ITU-T, "Information technology - ASN.1 encoding Rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", ISO/IEC 8825-1:2002, 2002.
- [X962] ANSI, "Public Key Cryptography For The Financial Services Industry: The Elliptic Curve Digital Signature Algorithm (ECDSA)", ANSI X9.62, 1998.

Rescorla Expires September 10, 2015 [Page 64]

Internet-Draft TLS March 2015

[13.2.](#) Informative References

- [CBCATT] Moeller, B., "Security of CBC Ciphersuites in SSL/TLS: Problems and Countermeasures", May 2004, <http://www.openssl.org/~bodo/tls-cbc.txt>.
- [CCM] "NIST Special Publication 800-38C: The CCM Mode for Authentication and Confidentiality", May 2004,

<<http://csrc.nist.gov/publications/nistpubs/800-38C/SP800-38C.pdf>>.

- [DES] "Data Encryption Standard (DES)", NIST FIPS PUB 46-3, October 1999.
- [DSS-3] National Institute of Standards and Technology, U.S., "Digital Signature Standard", NIST FIPS PUB 186-3 Draft, 2006.
- [ECDSA] American National Standards Institute, "Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital Signature Algorithm (ECDSA)", ANSI ANS X9.62-2005, November 2005.
- [ENCAUTH] Krawczyk, H., "The Order of Encryption and Authentication for Protecting Communications (Or: How Secure is SSL?)", 2001.
- [FI06] "Bleichenbacher's RSA signature forgery based on implementation error", August 2006, <<http://www.imc.org/ietf-openpgp/mail-archive/msg14307.html>>.
- [GCM] Dworkin, M., "Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC", NIST Special Publication 800-38D, November 2007.
- [I-D.ietf-tls-negotiated-ff-dhe]
Gillmor, D., "Negotiated Finite Field Diffie-Hellman Ephemeral Parameters for TLS", [draft-ietf-tls-negotiated-ff-dhe-07](#) (work in progress), March 2015.
- [I-D.ietf-tls-session-hash]
Bhargavan, K., Delignat-Lavaud, A., Pironti, A., Langley, A., and M. Ray, "Transport Layer Security (TLS) Session Hash and Extended Master Secret Extension", [draft-ietf-tls-session-hash-03](#) (work in progress), November 2014.

Barnes, R., Thomson, M., Pironti, A., and A. Langley, "Deprecating Secure Sockets Layer Version 3.0", [draft-ietf-tls-sslv3-diediedie-02](#) (work in progress), March 2015.

- [PKCS6] RSA Laboratories, "PKCS #6: RSA Extended Certificate Syntax Standard, version 1.5", November 1993.
- [PKCS7] RSA Laboratories, "PKCS #7: RSA Cryptographic Message Syntax Standard, version 1.5", November 1993.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [RFC1948] Bellare, S., "Defending Against Sequence Number Attacks", [RFC 1948](#), May 1996.
- [RFC2246] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", [RFC 2246](#), January 1999.
- [RFC2785] Zuccherato, R., "Methods for Avoiding the "Small-Subgroup" Attacks on the Diffie-Hellman Key Agreement Method for S/MIME", [RFC 2785](#), March 2000.
- [RFC3268] Chown, P., "Advanced Encryption Standard (AES) Ciphersuites for Transport Layer Security (TLS)", [RFC 3268](#), June 2002.
- [RFC3526] Kivinen, T. and M. Kojo, "More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)", [RFC 3526](#), May 2003.
- [RFC3766] Orman, H. and P. Hoffman, "Determining Strengths For Public Keys Used For Exchanging Symmetric Keys", [BCP 86](#), [RFC 3766](#), April 2004.
- [RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", [BCP 106](#), [RFC 4086](#), June 2005.
- [RFC4302] Kent, S., "IP Authentication Header", [RFC 4302](#), December 2005.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), December 2005.

-
- [RFC4307] Schiller, J., "Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)", [RFC 4307](#), December 2005.
- [RFC4346] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.1", [RFC 4346](#), April 2006.
- [RFC4366] Blake-Wilson, S., Nystrom, M., Hopwood, D., Mikkelsen, J., and T. Wright, "Transport Layer Security (TLS) Extensions", [RFC 4366](#), April 2006.
- [RFC4492] Blake-Wilson, S., Bolyard, N., Gupta, V., Hawk, C., and B. Moeller, "Elliptic Curve Cryptography (ECC) Cipher Suites for Transport Layer Security (TLS)", [RFC 4492](#), May 2006.
- [RFC4506] Eisler, M., "XDR: External Data Representation Standard", STD 67, [RFC 4506](#), May 2006.
- [RFC5081] Mavrogiannopoulos, N., "Using OpenPGP Keys for Transport Layer Security (TLS) Authentication", [RFC 5081](#), November 2007.
- [RFC5116] McGrew, D., "An Interface and Algorithms for Authenticated Encryption", [RFC 5116](#), January 2008.
- [RFC5705] Rescorla, E., "Keying Material Exporters for Transport Layer Security (TLS)", [RFC 5705](#), March 2010.
- [RFC6176] Turner, S. and T. Polk, "Prohibiting Secure Sockets Layer (SSL) Version 2.0", [RFC 6176](#), March 2011.
- [RSA] Rivest, R., Shamir, A., and L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems", Communications of the ACM v. 21, n. 2, pp. 120-126., February 1978.
- [SSL2] Netscape Communications Corp., "The SSL Protocol", February 1995.
- [SSL3] Freier, A., Karlton, P., and P. Kocher, "The SSL 3.0 Protocol", November 1996.
- [TIMING] Boneh, D. and D. Brumley, "Remote timing attacks are practical", USENIX Security Symposium, 2003.
- [TLSEXT] Eastlake 3rd, D., "Transport Layer Security (TLS)

Internet-Draft TLS March 2015

[X501] "Information Technology - Open Systems Interconnection -
The Directory: Models", ITU-T X.501, 1993.

13.3. URIs

[1] <https://github.com/tlswg/tls13-spec/issues/32>

[2] <mailto:tls@ietf.org>

[Appendix A](#). Protocol Data Structures and Constant Values

This section describes protocol types and constants.

[A.1](#). Record Layer

```
struct {
    uint8 major;
    uint8 minor;
} ProtocolVersion;

ProtocolVersion version = { 3, 4 };    /* TLS v1.3*/

enum {
    reserved(20), alert(21), handshake(22),
    application_data(23), (255)
} ContentType;

struct {
    ContentType type;
    ProtocolVersion version;
    uint16 length;
    opaque fragment[TLSPlainText.length];
} TLSPlainText;

struct {
    ContentType type;
    ProtocolVersion version;
    uint16 length;
    opaque nonce_explicit[SecurityParameters.record_iv_length];
    aead-ciphered struct {
        opaque content[TLSPlainText.length];
    }
}
```

```
    } fragment;  
} TLSCiphertext;
```

[A.2.](#) Alert Messages

```
enum { warning(1), fatal(2), (255) } AlertLevel;  
  
enum {  
    close_notify(0),  
    unexpected_message(10),  
    bad_record_mac(20),  
    decryption_failed_RESERVED(21),  
    record_overflow(22),  
    decompression_failure_RESERVED(30),  
    handshake_failure(40),  
    no_certificate_RESERVED(41),  
    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_RESERVED(60),  
    protocol_version(70),  
    insufficient_security(71),  
    internal_error(80),
```

```

        user_canceled(90),
        no_renegotiation(100),
        unsupported_extension(110),
        (255)
    } AlertDescription;

    struct {
        AlertLevel level;
        AlertDescription description;
    } Alert;

```

[A.3.](#) Handshake Protocol

```

enum {
    reserved(0), client_hello(1), server_hello(2),
    client_key_share(5), hello_retry_request(6),
    server_key_share(7), certificate(11), reserved(12),
    certificate_request(13), certificate_verify(15),
    reserved(16), finished(20), (255)
} HandshakeType;

struct {
    HandshakeType msg_type;      /* handshake type */
    uint24 length;              /* bytes in message */
    select (HandshakeType) {
        case client_hello:      ClientHello;
        case client_key_share:  ClientKeyShare;
        case server_hello:      ServerHello;
        case hello_retry_request: HelloRetryRequest;
        case server_key_share:  ServerKeyShare;
        case certificate:       Certificate;
    }

```



```

        case certificate_request: CertificateRequest;
        case certificate_verify:  CertificateVerify;
        case finished:           Finished;
    } body;
} Handshake;

```

[A.3.1.](#) Hello Messages

```

opaque SessionID<0..32>;

uint8 CipherSuite[2];    /* Cryptographic suite selector */

enum { null(0), (255) } CompressionMethod;

struct {
    ProtocolVersion client_version;
    Random random;
    SessionID session_id;
    CipherSuite cipher_suites<2..2^16-2>;
    CompressionMethod compression_methods<1..2^8-1>;
    select (extensions_present) {
        case false:
            struct {};
        case true:
            Extension extensions<0..2^16-1>;
    };
} ClientHello;

struct {
    ProtocolVersion server_version;

```

```

    Random random;
    SessionID session_id;
    CipherSuite cipher_suite;
    select (extensions_present) {
        case false:
            struct {};
        case true:
            Extension extensions<0..2^16-1>;
    };
} ServerHello;

```

```

struct {
    ProtocolVersion server_version;
    CipherSuite cipher_suite;
    NamedGroup selected_group;
    Extension extensions<0..2^16-1>;
} HelloRetryRequest;

struct {
    ExtensionType extension_type;
    opaque extension_data<0..2^16-1>;
} Extension;

enum {
    signature_algorithms(13), early_data(TBD), (65535)
} ExtensionType;

    struct {
        TLSCipherText messages<5 .. 2^24-1>;
    } EarlyDataExtension;

struct {
    Extension extensions<0..2^16-1>;
} EncryptedExtensions;

```

[A.3.1.1.](#) Signature Algorithm Extension

```

enum {
    none(0), md5(1), sha1(2), sha224(3), sha256(4), sha384(5),
    sha512(6), (255)
} HashAlgorithm;

```

```

enum { anonymous(0), rsa(1), dsa(2), ecdsa(3), (255) }
    SignatureAlgorithm;

struct {
    HashAlgorithm hash;
    SignatureAlgorithm signature;
} SignatureAndHashAlgorithm;

SignatureAndHashAlgorithm
    supported_signature_algorithms<2..2^16-2>;

```

[A.3.1.2](#). Named Group Extension

```

enum {
    // Elliptic Curve Groups.
    sect163k1 (1), sect163r1 (2), sect163r2 (3),
    sect193r1 (4), sect193r2 (5), sect233k1 (6),
    sect233r1 (7), sect239k1 (8), sect283k1 (9),
    sect283r1 (10), sect409k1 (11), sect409r1 (12),
    sect571k1 (13), sect571r1 (14), secp160k1 (15),
    secp160r1 (16), secp160r2 (17), secp192k1 (18),
    secp192r1 (19), secp224k1 (20), secp224r1 (21),
    secp256k1 (22), secp256r1 (23), secp384r1 (24),
    secp521r1 (25),

    // Finite Field Groups.
    ffdhe2048(256), ffdhe3072(257), ffdhe4096(258),
    ffdhe8192(259),

    // Reserved Code Points.
    reserved (0xFE00..0xFEFF),
    reserved(0xFF01),
    reserved(0xFF02),
    (0xFFFF)
} NamedGroup;

struct {
    NamedGroup named_group_list<1..2^16-1>
} NamedGroupList;

```

[A.3.2.](#) Key Exchange Messages

```
struct {
    NamedGroup group;
    opaque key_exchange<1..216-1>;
} ClientKeyShareOffer;

struct {
    ClientKeyShareOffer offers<0..216-1>;
} ClientKeyShare;

opaque dh_Y<1..216-1>;

opaque point <1..28-1>;

struct {
    NamedGroup group;
    opaque key_exchange<1..216-1>;
} ServerKeyShare;
```

[A.3.3.](#) Authentication Messages

```
opaque ASN1Cert<1..224-1>;

struct {
    ASN1Cert certificate_list<0..224-1>;
} Certificate;

enum {
    rsa_sign(1), dss_sign(2), rsa_fixed_dh(3), dss_fixed_dh(4),
    rsa_ephemeral_dh_RESERVED(5), dss_ephemeral_dh_RESERVED(6),
    fortezza_dms_RESERVED(20), (255)
} ClientCertificateType;

opaque DistinguishedName<1..216-1>;

struct {
    ClientCertificateType certificate_types<1..28-1>;
    SignatureAndHashAlgorithm
        supported_signature_algorithms<2..216-2>;
    DistinguishedName certificate_authorities<0..216-1>;
} CertificateRequest;

struct {
    digitally-signed struct {
        opaque handshake_messages_hash[hash_length];
    }
}
```

```
} CertificateVerify;
```

[A.3.4.](#) Handshake Finalization Messages

```
struct {  
    opaque verify_data[verify_data_length];  
} Finished;
```

[A.4.](#) The Cipher Suite

The following values define the cipher suite codes used in the ClientHello and ServerHello messages.

A cipher suite defines a cipher specification supported in TLS Version 1.2.

TLS_NULL_WITH_NULL_NULL is specified and is the initial state of a TLS connection during the first handshake on that channel, but MUST NOT be negotiated, as it provides no more protection than an unsecured connection.

```
CipherSuite TLS_NULL_WITH_NULL_NULL = { 0x00,0x00 };
```

The following cipher suite definitions, defined in [{{RFC5288}}](#), are used for server-authenticated (and optionally client-authenticated) Diffie-Hellman. DHE denotes ephemeral Diffie-Hellman, where the Diffie-Hellman parameters are signed by a signature-capable certificate, which has been signed by the CA. The signing algorithm used by the server is specified after the DHE component of the CipherSuite name. The server can request any signature-capable certificate from the client for client authentication.

```
CipherSuite TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 = {0x00,0x9E}  
CipherSuite TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 = {0x00,0x9F}  
CipherSuite TLS_DHE_DSS_WITH_AES_128_GCM_SHA256 = {0x00,0xA2}  
CipherSuite TLS_DHE_DSS_WITH_AES_256_GCM_SHA384 = {0x00,0xA3}
```

The following cipher suite definitions, defined in [{{RFC5289}}](#), are used for server-authenticated (and optionally client-authenticated) Elliptic Curve Diffie-Hellman. ECDHE denotes ephemeral Diffie-Hellman, where the Diffie-Hellman parameters are signed by a signature-capable certificate, which has been signed by the CA. The

signing algorithm used by the server is specified after the DHE component of the CipherSuite name. The server can request any signature-capable certificate from the client for client authentication.

```
CipherSuite TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 = {0xC0,0x2B};
CipherSuite TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 = {0xC0,0x2C};
CipherSuite TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256   = {0xC0,0x2F};
CipherSuite TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384   = {0xC0,0x30};
```

The following ciphers, defined in [RFC5288], are used for completely anonymous Diffie-Hellman communications in which neither party is authenticated. Note that this mode is vulnerable to man-in-the-middle attacks. Using this mode therefore is of limited use: These cipher suites MUST NOT be used by TLS 1.2 implementations unless the application layer has specifically requested to allow anonymous key exchange. (Anonymous key exchange may sometimes be acceptable, for example, to support opportunistic encryption when no set-up for authentication is in place, or when TLS is used as part of more complex security protocols that have other means to ensure authentication.)

```
CipherSuite TLS_DH_anon_WITH_AES_128_GCM_SHA256 = {0x00,0xA6}
CipherSuite TLS_DH_anon_WITH_AES_256_GCM_SHA384 = {0x00,0xA7}
```

[[TODO: Add all the defined AEAD ciphers. This currently only lists GCM. <https://github.com/tlswg/tls13-spec/issues/53>]] Note that using non-anonymous key exchange without actually verifying the key exchange is essentially equivalent to anonymous key exchange, and the same precautions apply. While non-anonymous key exchange will generally involve a higher computational and communicational cost than anonymous key exchange, it may be in the interest of interoperability not to disable non-anonymous key exchange when the application layer is allowing anonymous key exchange.

The PRFs SHALL be as follows:

- o For cipher suites ending with _SHA256, the PRF is the TLS PRF with

SHA-256 as the hash function.

- o For cipher suites ending with _SHA384, the PRF is the TLS PRF with SHA-384 as the hash function.

New cipher suite values are been assigned by IANA as described in [Section 12](#).

Note: The cipher suite values { 0x00, 0x1C } and { 0x00, 0x1D } are reserved to avoid collision with Fortezza-based cipher suites in SSL 3.

[A.5](#). The Security Parameters

These security parameters are determined by the TLS Handshake Protocol and provided as parameters to the TLS record layer in order to initialize a connection state. SecurityParameters includes:

```
enum { server, client } ConnectionEnd;

enum { tls_prf_sha256 } PRFAlgorithm;

enum { aes_gcm } RecordProtAlgorithm;

/* The algorithms specified in PRFAlgorithm and
   RecordProtAlgorithm may be added to. */

struct {
    ConnectionEnd          entity;
    PRFAlgorithm           prf_algorithm;
    RecordProtAlgorithm    record_prot_algorithm;
    uint8                  enc_key_length;
    uint8                  block_length;
    uint8                  fixed_iv_length;
    uint8                  record_iv_length;
    opaque                 hs_master_secret[48];
    opaque                 master_secret[48];
    opaque                 client_random[32];
```

```
        opaque                server_random[32];  
    } SecurityParameters;
```

[A.6.](#) Changes to [RFC 4492](#)

[RFC 4492](#) [[RFC4492](#)] adds Elliptic Curve cipher suites to TLS. This document changes some of the structures used in that document. This section details the required changes for implementors of both [RFC 4492](#) and TLS 1.2. Implementors of TLS 1.2 who are not implementing [RFC 4492](#) do not need to read this section.

This document adds a "signature_algorithm" field to the digitally-signed element in order to identify the signature and digest algorithms used to create a signature. This change applies to digital signatures formed using ECDSA as well, thus allowing ECDSA signatures to be used with digest algorithms other than SHA-1, provided such use is compatible with the certificate and any restrictions imposed by future revisions of [[RFC3280](#)].

As described in [Section 7.3.5](#) and [Section 7.3.9](#), the restrictions on the signature algorithms used to sign certificates are no longer tied to the cipher suite (when used by the server) or the

ClientCertificateType (when used by the client). Thus, the restrictions on the algorithm used to sign certificates specified in Sections [2](#) and [3](#) of [RFC 4492](#) are also relaxed. As in this document, the restrictions on the keys in the end-entity certificate remain.

[Appendix B.](#) Glossary

Advanced Encryption Standard (AES)

AES [[AES](#)] is a widely used symmetric encryption algorithm. AES is a block cipher with a 128-, 192-, or 256-bit keys and a 16-byte block size. TLS currently only supports the 128- and 256-bit key sizes.

application protocol

An application protocol is a protocol that normally layers directly on top of the transport layer (e.g., TCP/IP). Examples include HTTP, TELNET, FTP, and SMTP.

asymmetric cipher

See public key cryptography.

authenticated encryption with additional data (AEAD)

A symmetric encryption algorithm that simultaneously provides confidentiality and message integrity.

authentication

Authentication is the ability of one entity to determine the identity of another entity.

certificate

As part of the X.509 protocol (a.k.a. ISO Authentication framework), certificates are assigned by a trusted Certificate Authority and provide a strong binding between a party's identity or some other attributes and its public key.

client

The application entity that initiates a TLS connection to a server. This may or may not imply that the client initiated the underlying transport connection. The primary operational difference between the server and client is that the server is generally authenticated, while the client is only optionally authenticated.

client write key

The key used to protect data written by the client.

connection

A connection is a transport (in the OSI layering model definition) that provides a suitable type of service. For TLS, such connections are peer-to-peer relationships. The connections are transient. Every connection is associated with one session.

Digital Signature Standard (DSS)

A standard for digital signing, including the Digital Signing Algorithm, approved by the National Institute of Standards and Technology, defined in NIST FIPS PUB 186-2, "Digital Signature Standard", published January 2000 by the U.S. Department of Commerce [[DSS](#)]. A significant update [[DSS-3](#)] has been drafted and was published in March 2006.

digital signatures

Digital signatures utilize public key cryptography and one-way hash functions to produce a signature of the data that can be authenticated, and is difficult to forge or repudiate.

handshake

An initial negotiation between client and server that establishes the parameters of their transactions.

Initialization Vector (IV)

Some AEAD ciphers require an initialization vector to allow the cipher to safely protect multiple chunks of data with the same keying material. The size of the IV depends on the cipher suite.

Message Authentication Code (MAC)

A Message Authentication Code is a one-way hash computed from a message and some secret data. It is difficult to forge without knowing the secret data. Its purpose is to detect if the message has been altered.

master secret

Secure secret data used for generating keys and IVs.

MD5

MD5 [[RFC1321](#)] is a hashing function that converts an arbitrarily long data stream into a hash of fixed size (16 bytes). Due to significant progress in cryptanalysis, at the time of publication of this document, MD5 no longer can be considered a 'secure' hashing function.

public key cryptography

A class of cryptographic techniques employing two-key ciphers. Messages encrypted with the public key can only be decrypted with the associated private key. Conversely, messages signed with the private key can be verified with the public key.

one-way hash function

A one-way transformation that converts an arbitrary amount of data into a fixed-length hash. It is computationally hard to reverse the transformation or to find collisions. MD5 and SHA are examples of one-way hash functions.

RSA

A very widely used public key algorithm that can be used for either encryption or digital signing. [[RSA](#)]

server

The server is the application entity that responds to requests for connections from clients. See also "client".

session

A TLS session is an association between a client and a server. Sessions are created by the handshake protocol. Sessions define a set of cryptographic security parameters that can be shared among multiple connections. Sessions are used to avoid the expensive negotiation of new security parameters for each connection.

session identifier

A session identifier is a value generated by a server that identifies a particular session.

server write key

The key used to protect data written by the server.

SHA

The Secure Hash Algorithm [[SHS](#)] is defined in FIPS PUB 180-2. It produces a 20-byte output. Note that all references to SHA (without a numerical suffix) actually use the modified SHA-1 algorithm.

SHA-256

The 256-bit Secure Hash Algorithm is defined in FIPS PUB 180-2. It produces a 32-byte output.

SSL

Netscape's Secure Socket Layer protocol [[SSL3](#)]. TLS is based on SSL Version 3.0.

Transport Layer Security (TLS)

This protocol; also, the Transport Layer Security working group of the Internet Engineering Task Force (IETF). See "Working Group Information" at the end of this document (see page 99).

[Appendix C](#). Cipher Suite Definitions

Cipher Suite	Key Exchange	Record Protection	PRF
TLS_NULL_WITH_NULL_NULL	NULL	NULL_NULL	N/A
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	DHE_RSA	AES_128_GCM	SHA256
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	DHE_RSA	AES_256_GCM	SHA384
TLS_DHE_DSS_WITH_AES_128_GCM_SHA256	DHE_DSS	AES_128_GCM	SHA256
TLS_DHE_DSS_WITH_AES_256_GCM_SHA384	DHE_DSS	AES_256_GCM	SHA384
TLS_DH_anon_WITH_AES_128_GCM_SHA256	DH_anon	AES_128_GCM	SHA256
TLS_DH_anon_WITH_AES_256_GCM_SHA384	DH_anon	AES_128_GCM	SHA384

Cipher	Key Material	Implicit IV Size	Explicit IV Size
-----	-----	-----	-----
NULL	0	0	0
AES_128_GCM	16	4	8
AES_256_GCM	32	4	8

Key Material

The number of bytes from the key_block that are used for generating the write keys.

Implicit IV Size

The amount of data to be generated for the per-connection part of the initialization vector. This is equal to `SecurityParameters.fixed_iv_length`).

Explicit IV Size

The amount of data needed to be generated for the per-record part of the initialization vector. This is equal to `SecurityParameters.record_iv_length`).

[Appendix D](#). Implementation Notes

The TLS protocol cannot prevent many common security mistakes. This section provides several recommendations to assist implementors.

[D.1](#). Random Number Generation and Seeding

TLS requires a cryptographically secure pseudorandom number generator (PRNG). Care must be taken in designing and seeding PRNGs. PRNGs based on secure hash operations, most notably SHA-1, are acceptable, but cannot provide more security than the size of the random number generator state.

Internet-Draft

TLS

March 2015

To estimate the amount of seed material being produced, add the number of bits of unpredictable information in each seed byte. For example, keystroke timing values taken from a PC compatible's 18.2 Hz timer provide 1 or 2 secure bits each, even though the total size of the counter value is 16 bits or more. Seeding a 128-bit PRNG would thus require approximately 100 such timer values.

[RFC4086] provides guidance on the generation of random values.

[D.2.](#) Certificates and Authentication

Implementations are responsible for verifying the integrity of certificates and should generally support certificate revocation messages. Certificates should always be verified to ensure proper signing by a trusted Certificate Authority (CA). The selection and addition of trusted CAs should be done very carefully. Users should be able to view information about the certificate and root CA.

[D.3.](#) Cipher Suites

TLS supports a range of key sizes and security levels, including some that provide no or minimal security. A proper implementation will probably not support many cipher suites. For instance, anonymous Diffie-Hellman is strongly discouraged because it cannot prevent man-in-the-middle attacks. Applications should also enforce minimum and maximum key sizes. For example, certificate chains containing 512-bit RSA keys or signatures are not appropriate for high-security applications.

[D.4.](#) Implementation Pitfalls

Implementation experience has shown that certain parts of earlier TLS specifications are not easy to understand, and have been a source of interoperability and security problems. Many of these areas have been clarified in this document, but this appendix contains a short list of the most important things that require special attention from implementors.

TLS protocol issues:

- Do you correctly handle handshake messages that are fragmented to multiple TLS records (see [Section 6.2.1](#))? Including corner cases

like a ClientHello that is split to several small fragments? Do you fragment handshake messages that exceed the maximum fragment size? In particular, the certificate and certificate request handshake messages can be large enough to require fragmentation.

- Do you ignore the TLS record layer version number in all TLS records before ServerHello (see [Appendix E.1](#))?
- Do you handle TLS extensions in ClientHello correctly, including omitting the extensions field completely?
- When the server has requested a client certificate, but no suitable certificate is available, do you correctly send an empty Certificate message, instead of omitting the whole message (see [Section 7.3.9](#))?

Cryptographic details:

- What countermeasures do you use to prevent timing attacks against RSA signing operations [[TIMING](#)].
- When verifying RSA signatures, do you accept both NULL and missing parameters (see [Section 4.7](#))? Do you verify that the RSA padding doesn't have additional data after the hash value? [[FI06](#)]
- When using Diffie-Hellman key exchange, do you correctly strip leading zero bytes from the negotiated key (see [Section 8.1.2](#))?
- Does your TLS client check that the Diffie-Hellman parameters sent by the server are acceptable (see [Appendix F.1.1.2](#))?
- Do you use a strong and, most importantly, properly seeded random number generator (see [Appendix D.1](#)) Diffie-Hellman private values, the DSA "k" parameter, and other security-critical values?

[Appendix E](#). Backward Compatibility

[E.1](#). Compatibility with prior versions

[[TODO: Revise backward compatibility section for TLS 1.3.

<https://github.com/tlswg/tls13-spec/issues/54>]] Since there are various versions of TLS (1.0, 1.1, 1.2, 1.3, and any future versions) and SSL (2.0 and 3.0), means are needed to negotiate the specific protocol version to use. The TLS protocol provides a built-in mechanism for version negotiation so as not to bother other protocol components with the complexities of version selection.

TLS versions 1.0, 1.1, and 1.2, and SSL 3.0 are very similar, and use compatible ClientHello messages; thus, supporting all of them is relatively easy. Similarly, servers can easily handle clients trying to use future versions of TLS as long as the ClientHello format remains compatible, and the client supports the highest protocol version available in the server.

A TLS 1.3 client who wishes to negotiate with such older servers will send a normal TLS 1.3 ClientHello, containing { 3, 4 } (TLS 1.3) in ClientHello.client_version. If the server does not support this version, it will respond with a ServerHello containing an older version number. If the client agrees to use this version, the negotiation will proceed as appropriate for the negotiated protocol.

If the version chosen by the server is not supported by the client (or not acceptable), the client MUST send a "protocol_version" alert message and close the connection.

If a TLS server receives a ClientHello containing a version number greater than the highest version supported by the server, it MUST reply according to the highest version supported by the server.

A TLS server can also receive a ClientHello containing a version number smaller than the highest supported version. If the server wishes to negotiate with old clients, it will proceed as appropriate for the highest version supported by the server that is not greater than ClientHello.client_version. For example, if the server supports TLS 1.0, 1.1, and 1.2, and client_version is TLS 1.0, the server will proceed with a TLS 1.0 ServerHello. If server supports (or is willing to use) only versions greater than client_version, it MUST send a "protocol_version" alert message and close the connection.

Whenever a client already knows the highest protocol version known to a server (for example, when resuming a session), it SHOULD initiate the connection in that native protocol.

Note: some server implementations are known to implement version negotiation incorrectly. For example, there are buggy TLS 1.0 servers that simply close the connection when the client offers a version newer than TLS 1.0. Also, it is known that some servers will refuse the connection if any TLS extensions are included in ClientHello. Interoperability with such buggy servers is a complex topic beyond the scope of this document, and may require multiple connection attempts by the client.

Earlier versions of the TLS specification were not fully clear on what the record layer version number (TLSPlaintext.version) should contain when sending ClientHello (i.e., before it is known which version of the protocol will be employed). Thus, TLS servers compliant with this specification MUST accept any value {03,XX} as the record layer version number for ClientHello.

TLS clients that wish to negotiate with older servers MAY send any value {03,XX} as the record layer version number. Typical values would be {03,00}, the lowest version number supported by the client,

and the value of ClientHello.client_version. No single value will guarantee interoperability with all old servers, but this is a complex topic beyond the scope of this document.

[E.2.](#) Compatibility with SSL

The security of SSL 2.0 [[SSL2](#)] is considered insufficient for the reasons enumerated in [[RFC6176](#)], and MUST NOT be negotiated for any reason.

Implementations MUST NOT send an SSL version 2.0 compatible CLIENT-HELLO. Implementations MUST NOT negotiate TLS 1.3 or later using an SSL version 2.0 compatible CLIENT-HELLO. Implementations are NOT RECOMMENDED to accept an SSL version 2.0 compatible CLIENT-HELLO in order to negotiate older versions of TLS.

Implementations MUST NOT send or accept any records with a version less than { 3, 0 }.

The security of SSL 3.0 [[SSL3](#)] is considered insufficient for the reasons enumerated in [[I-D.ietf-tls-ssl3-diediedie](#)], and MUST NOT be

negotiated for any reason.

Implementations MUST NOT send a `ClientHello.version` or `ServerHello.version` set to `{ 3, 0 }` or less. Any endpoint receiving a Hello message with `ClientHello.version` or `ServerHello.version` set to `{ 3, 0 }` MUST respond with a "protocol_version" alert message and close the connection.

[Appendix F](#). Security Analysis

The TLS protocol is designed to establish a secure connection between a client and a server communicating over an insecure channel. This document makes several traditional assumptions, including that attackers have substantial computational resources and cannot obtain secret information from sources outside the protocol. Attackers are assumed to have the ability to capture, modify, delete, replay, and otherwise tamper with messages sent over the communication channel. This appendix outlines how TLS has been designed to resist a variety of attacks.

[F.1](#). Handshake Protocol

The handshake protocol is responsible for selecting a cipher spec and generating a master secret, which together comprise the primary cryptographic parameters associated with a secure session. The handshake protocol can also optionally authenticate parties who have certificates signed by a trusted certificate authority.

[F.1.1](#). Authentication and Key Exchange

TLS supports three authentication modes: authentication of both parties, server authentication with an unauthenticated client, and total anonymity. Whenever the server is authenticated, the channel is secure against man-in-the-middle attacks, but completely anonymous sessions are inherently vulnerable to such attacks. Anonymous servers cannot authenticate clients. If the server is authenticated, its certificate message must provide a valid certificate chain leading to an acceptable certificate authority. Similarly, authenticated clients must supply an acceptable certificate to the server. Each party is responsible for verifying that the other's certificate is valid and has not expired or been revoked.

The general goal of the key exchange process is to create a `pre_master_secret` known to the communicating parties and not to attackers. The `pre_master_secret` will be used to generate the `master_secret` (see [Section 8.1](#)). The `master_secret` is required to generate the Finished messages and record protection keys (see [Section 7.3.8](#) and [Section 6.3](#)). By sending a correct Finished message, parties thus prove that they know the correct `pre_master_secret`.

[F.1.1.1](#). Anonymous Key Exchange

Completely anonymous sessions can be established using Diffie-Hellman for key exchange. The server's public parameters are contained in the server key share message, and the client's are sent in the client key share message. Eavesdroppers who do not know the private values should not be able to find the Diffie-Hellman result (i.e., the `pre_master_secret`).

Warning: Completely anonymous connections only provide protection against passive eavesdropping. Unless an independent tamper-proof channel is used to verify that the Finished messages were not replaced by an attacker, server authentication is required in environments where active man-in-the-middle attacks are a concern.

[F.1.1.2](#). Diffie-Hellman Key Exchange with Authentication

When Diffie-Hellman key exchange is used, the client and server use the client key exchange and server key exchange messages to send temporary Diffie-Hellman parameters. The signature in the certificate verify message (if present) covers the entire handshake up to that point and thus attests the certificate holder's desire to use the the ephemeral DHE keys.

Peers SHOULD validate each other's public key Y (dh_Ys offered by the server or DH_Yc offered by the client) by ensuring that $1 < Y < p-1$. This simple check ensures that the remote peer is properly behaved and isn't forcing the local system into a small subgroup.

Additionally, using a fresh key for each handshake provides Perfect Forward Secrecy. Implementations SHOULD generate a new X for each

handshake when using DHE cipher suites.

[F.1.2.](#) Version Rollback Attacks

Because TLS includes substantial improvements over SSL Version 2.0, attackers may try to make TLS-capable clients and servers fall back to Version 2.0. This attack can occur if (and only if) two TLS-capable parties use an SSL 2.0 handshake.

Although the solution using non-random PKCS #1 block type 2 message padding is inelegant, it provides a reasonably secure way for Version 3.0 servers to detect the attack. This solution is not secure against attackers who can brute-force the key and substitute a new ENCRYPTED-KEY-DATA message containing the same key (but with normal padding) before the application-specified wait threshold has expired. Altering the padding of the least-significant 8 bytes of the PKCS padding does not impact security for the size of the signed hashes and RSA key lengths used in the protocol, since this is essentially equivalent to increasing the input block size by 8 bytes.

[F.1.3.](#) Detecting Attacks Against the Handshake Protocol

An attacker might try to influence the handshake exchange to make the parties select different encryption algorithms than they would normally choose.

For this attack, an attacker must actively change one or more handshake messages. If this occurs, the client and server will compute different values for the handshake message hashes. As a result, the parties will not accept each others' Finished messages. Without the master_secret, the attacker cannot repair the Finished messages, so the attack will be discovered.

[F.1.4.](#) Resuming Sessions

When a connection is established by resuming a session, new ClientHello.random and ServerHello.random values are hashed with the session's master_secret. Provided that the master_secret has not been compromised and that the secure hash operations used to produce the record protection keys are secure, the connection should be secure and effectively independent from previous connections.

Attackers cannot use known keys to compromise the master_secret without breaking the secure hash operations.

Sessions cannot be resumed unless both the client and server agree. If either party suspects that the session may have been compromised, or that certificates may have expired or been revoked, it should force a full handshake. An upper limit of 24 hours is suggested for session ID lifetimes, since an attacker who obtains a master_secret may be able to impersonate the compromised party until the corresponding session ID is retired. Applications that may be run in relatively insecure environments should not write session IDs to stable storage.

[F.2.](#) Protecting Application Data

The master_secret is hashed with the ClientHello.random and ServerHello.random to produce unique record protection secrets for each connection.

Outgoing data is protected using an AEAD algorithm before transmission. The authentication data includes the sequence number, message type, message length, and the message contents. The message type field is necessary to ensure that messages intended for one TLS record layer client are not redirected to another. The sequence number ensures that attempts to delete or reorder messages will be detected. Since sequence numbers are 64 bits long, they should never overflow. Messages from one party cannot be inserted into the other's output, since they use independent keys.

[F.3.](#) Denial of Service

TLS is susceptible to a number of denial-of-service (DoS) attacks. In particular, an attacker who initiates a large number of TCP connections can cause a server to consume large amounts of CPU doing asymmetric crypto operations. However, because TLS is generally used over TCP, it is difficult for the attacker to hide his point of origin if proper TCP SYN randomization is used [[RFC1948](#)] by the TCP stack.

Because TLS runs over TCP, it is also susceptible to a number of DoS attacks on individual connections. In particular, attackers can forge RSTs, thereby terminating connections, or forge partial TLS records, thereby causing the connection to stall. These attacks cannot in general be defended against by a TCP-using protocol. Implementors or users who are concerned with this class of attack should use IPsec AH [[RFC4302](#)] or ESP [[RFC4303](#)].

Internet-Draft

TLS

March 2015

[F.4.](#) Final Notes

For TLS to be able to provide a secure connection, both the client and server systems, keys, and applications must be secure. In addition, the implementation must be free of security errors.

The system is only as strong as the weakest key exchange and authentication algorithm supported, and only trustworthy cryptographic functions should be used. Short public keys and anonymous servers should be used with great caution. Implementations and users must be careful when deciding which certificates and certificate authorities are acceptable; a dishonest certificate authority can do tremendous damage.

[Appendix G.](#) Working Group Information

The discussion list for the IETF TLS working group is located at the e-mail address tls@ietf.org [2]. Information on the group and information on how to subscribe to the list is at <https://www1.ietf.org/mailman/listinfo/tls>

Archives of the list can be found at: <http://www.ietf.org/mail-archive/web/tls/current/index.html>

[Appendix H.](#) Contributors

Christopher Allen (co-editor of TLS 1.0)
Alacrity Ventures
ChristopherA@AlacrityManagement.com

Martin Abadi
University of California, Santa Cruz
abadi@cs.ucsc.edu

Karthikeyan Bhargavan (co-author of [[I-D.ietf-tls-session-hash](#)])
INRIA
karthikeyan.bhargavan@inria.fr

Steven M. Bellovin
Columbia University
smb@cs.columbia.edu

Simon Blake-Wilson (co-author of [RFC4492](#))

BCI
sblakewilson@bcisse.com

Nelson Bolyard
Sun Microsystems, Inc.

Rescorla

Expires September 10, 2015

[Page 89]

Internet-Draft

TLS

March 2015

nelson@bolyard.com (co-author of [RFC4492](#))

Ran Canetti
IBM
canetti@watson.ibm.com

Pete Chown
Skygate Technology Ltd
pc@skygate.co.uk

Antoine Delignat-Lavaud (co-author of [[I-D.ietf-tls-session-hash](#)])
INRIA
antoine.delignat-lavaud@inria.fr

Tim Dierks (co-editor of TLS 1.0, 1.1, and 1.2)
Independent
tim@dierks.org

Taher Elgamal
taher@securify.com
Securify

Pasi Eronen
pasi.eronen@nokia.com
Nokia

Anil Gangolli
anil@busybuddha.org

Vipul Gupta (co-author of [RFC4492](#))
Sun Microsystems Laboratories
vipul.gupta@sun.com

Kipp Hickman

Chris Hawk (co-author of [RFC4492](#))

Corriente Networks LLC
chris@corriente.net

Alfred Hoenes

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

Daniel Kahn Gillmor
ACLU
dkg@fifthhorseman.net

Rescorla

Expires September 10, 2015

[Page 90]

Internet-Draft

TLS

March 2015

Phil Karlton (co-author of SSLv3)

Paul Kocher (co-author of SSLv3)
Cryptography Research
paul@cryptography.com

Hugo Krawczyk
IBM
hugo@ee.technion.ac.il

Adam Langley (co-author of [[I-D.ietf-tls-session-hash](#)])
Google
agl@google.com

Ilari Liusvaara
ilari.liusvaara@elisanet.fi

Jan Mikkelsen
Transactionware
janm@transactionware.com

Bodo Moeller (co-author of [RFC4492](#))
Google
bodo@openssl.org

Magnus Nystrom
RSA Security
magnus@rsasecurity.com

Alfredo Pironti (co-author of [[I-D.ietf-tls-session-hash](#)])
INRIA
alfredo.pironti@inria.fr

Marsh Ray (co-author of [[I-D.ietf-tls-session-hash](#)])
Microsoft
maray@microsoft.com

Robert Relyea
Netscape Communications
relyea@netscape.com

Jim Roskind
Netscape Communications
jar@netscape.com

Michael Sabin

Dan Simon

Rescorla

Expires September 10, 2015

[Page 91]

Internet-Draft

TLS

March 2015

Microsoft, Inc.
dansimon@microsoft.com

Martin Thomson
Mozilla
mt@mozilla.com

Tom Weinstein

Tim Wright
Vodafone
timothy.wright@vodafone.com

Author's Address

Eric Rescorla
RTFM, Inc.

EMail: ekr@rtfm.com

