

Internals of Nut/TLS

Daniel Otte
(daniel.otte@rub.de)

September 2012

Abstract

This document will introduce you to the inner workings of Nut/TLS, the TLS 1.2 server implementation for the Nut/OS embedded operating system.

Contents

1	About	3
2	Overview	3
2.1	Features	3
2.2	A few words about TLS	3
2.3	Basic operation of Nut/TLS	4
2.3.1	Driver Initialization	4
2.3.2	Opening A Nut/TLS stream	5
2.3.3	Assigning A Stream To A Nut/TLS stream	5
2.3.4	Handshake	5
2.3.5	Send And Receive Data	5
2.3.6	Destruction Of A Nut/TLS stream	6
2.4	Drivers Point-Of-View	6
2.4.1	Reading Data	6
2.4.2	Sending Data	6
2.5	Applications Point-Of-View	6
2.6	How The Parts Fit Together	7
2.7	Implementation content	7
2.7.1	Header Files	7
2.7.2	Implementation Files	8

3	Basic Data-structures	8
3.1	tls_connection_t	9
3.1.1	name	9
3.1.2	next	9
3.1.3	connection	9
3.1.4	write_buffer	9
3.1.5	write_buffer_size	9
3.1.6	write_buffer_fill	10
3.1.7	read_buffer	10
3.1.8	write_buffer_protocol	10
3.1.9	closing	10
3.1.10	write_lock	10
3.2	tls_record_connection_t	10
3.2.1	base_stream	11
3.2.2	session_id	11
3.2.3	version	11
3.2.4	secure	11
3.2.5	post_clienthello	11
3.2.6	tx_pre_pattern	11
3.2.7	rx_pre_pattern	12
3.2.8	sec_parameters	12
3.2.9	sec_states_local	12
3.2.10	sec_states_remote	12
3.2.11	pending	12
3.2.12	pending_status	12
3.2.13	error_state	12
4	Buffering	12
4.1	Write Buffer	13
4.1.1	Allocation	13
4.2	Read Buffer	13
5	Packet Reception	13
6	Packet Transmission	15
7	The Handshake Process	15
8	Alert Handling	16

1 About

This section will tell you a little bit about the background of Nut/TLS and its development.

*** will be written later ***

2 Overview

2.1 Features

Supported Cipher-Suites:

- TLS_RSA_WITH_AES_128_CBC_SHA256 (0x00,0x3C)
- TLS_RSA_WITH_AES_128_CBC_SHA (0x00,0x2F)
- TLS_RSA_WITH_AES_256_CBC_SHA256 (0x00,0x3D)
- TLS_RSA_WITH_AES_256_CBC_SHA (0x00,0x35)
- TLS_RSA_WITH_3DES_EDE_CBC_SHA (0x00,0x0A)
- TLS_RSA_WITH_RC4_128_SHA (0x00,0x05)
- TLS_RSA_WITH_RC4_128_MD5 (0x00,0x04)
- TLS_RSA_WITH_NULL_SHA256 (0x00,0x3B)
- TLS_RSA_WITH_NULL_SHA (0x00,0x02)
- TLS_RSA_WITH_NULL_MD5 (0x00,0x01)

2.2 A few words about TLS

TLS means Transport-Layer-Security and is available in several versions, which are all available as RFCs. We are here concerned only with version 1.2 which is specified in RFC-5246. TLS is also the successor of SSL, as the new name did not propagate very well most people still say "SSL" when they mean TLS.

TLS forms a tunnel for sensitive data over insecure channels. It neither defines what data is transported nor over which channel the secure channel is established. A common combination is transferring HTTP-protocol data via TLS over a TCP-connection. This is widely known as HTTPS.

TLS is basically a record protocol operating over a reliable stream (like a TCP-connection), with two specified additional protocols operating on top of it. A third protocol would be the application stream which also has a specified protocol id.

Each record transmitted over a TLS connection may be encrypted and/or signed by the transmitting party.

The two important protocols which are part of TLS and are operating over the record connection are the handshake-protocol and the alert-protocol.

The handshake-protocol allows establishing a secure connection using asymmetric cryptography while the alert-protocol handles error- and warning-messages which are also used to gracefully close the connection.

It is important to know that the standard explicitly forbids application-layer communication before the first successful handshake is done.

2.3 Basic operation of Nut/TLS

Nut/OS offers you to use TCP connections like normal streams in C. Nut/TLS extends this idea to use TLS connections also like normal streams. It even allows you to operate your connection over any bidirectional stream (not only TCP connections but also serial/uart lines).

Nut/TLSs core is implemented as Nut/OS driver. It offers the standard interface of Nut/OS drivers but also special functions which are necessary for it to work.

The usage schema is quite simple:

1. initialization of the driver
2. opening an Nut/TLS stream
3. assigning a stream to the Nut/TLS stream
4. doing a handshake
5. send and receive data
6. close the Nut/TLS stream

2.3.1 Driver Initialization

The initialization of the driver is started by the following code:

```
NutRegisterDevice(&devTLS, 0, 0);
```

This causes the creation of a dedicated thread which flushes the write buffers in fixed intervals and the allocation of memory for the resumption table.

2.3.2 Opening A Nut/TLS stream

A Nut/TLS stream is opened like a regular file with `fopen()`. The name must start with `"tls:"` to indicate that the Nut/TLS driver should be used to open this file. The `"tls:"`-prefix is followed by an identifier which is unique for each distinct TLS-connection running at the same time. The identifier may consist of any combination of the letters "A" to "Z", "a" to "z", the numerical characters "0" to "9" and the character "_". Opening an existing connection will result in a stream which is an alias to the already opened stream.

2.3.3 Assigning A Stream To A Nut/TLS stream

Due to the flexibility of Nut/TLS it is necessary to explicitly assign a stream to operate on to the Nut/TLS stream. This stream may be of any kind, but must be bidirectional. In most cases this will be a stream obtained from a socket-TCP connection. The stream to operate on is also called the "base stream".

2.3.4 Handshake

After the base stream is assigned a basic record layer connection is established. This connection does neither provide confidentiality nor integrity. To secure the connection a "handshake" has to take place. Within this handshake the cryptographic parameters are negotiated and cryptographic keys are exchanged which will protect following data.

The handshake is normally initiated by the client after establishing the record connection. To proceed the handshake the `make_handshake()` function is called. It just takes a pointer to the Nut/TLS stream as parameter and returns zero on success.

2.3.5 Send And Receive Data

When the handshake is done data exchange can begin. Exchanging data is as simple as using `fread()`, `fwrite()`, `fprintf()`. All functions which can handle streams will also work with Nut/TLS streams.

2.3.6 Destruction Of A Nut/TLS stream

The connection is destroyed by calling `fclose()` with a pointer to the Nut/TLS stream as parameter.

2.4 Drivers Point-Of-View

2.4.1 Reading Data

When data is requested by the application, the driver first tries to fulfill the request by using already buffered data. If the drivers buffer does not contain enough data, a packet is read from the base stream, in-place decrypted and then chained into the buffer. If the request can still not be fulfilled it reads the next packet, decrypts it and chains it into the buffer. This is repeated until the request can be fulfilled or an error occurs. If an error occurs the data which could be read is returned to the application. The eof-flag and error indicator are set according to the reason of the error.

2.4.2 Sending Data

If the data to send fits into the write buffer it is simply copied into the write buffer. If the write buffer is full or the data does not fit into the write buffer the content of the write buffer is encrypted and send to the base stream. The remaining data is send in packet with a size which is equal to the size of the write buffer. If there is a last block, which is smaller than the write buffer, it is copied into the write buffer.

Encrypting and sending the packet to the base stream is, opposed to the reading behavior, handled by the lower layer functions. This avoids additional copying of the data.

2.5 Applications Point-Of-View

Nut/TLS appears as filter to the application. It behaves like a normal C-style stream (like `stdin`, `stdout` and `stderr`). The stream is bidirectional and binary only, All the tools which are able to operate on this kind of streams are adequate for use with Nut/TLS streams (like `fread()`, `fwrite()`, `fprintf()`, ...).

But before normal operation can take place, it is necessary to properly start the connection. Therefore a stream has to be assigned to the Nut/TLS stream and a handshake has to be done. Since Nut/TLS is quite flexible it is possible to run TLS connections over nearly any other bidirectional stream. So a stream to operate on has to be assigned to the

Nut/TLS stream. This is done by `int tls_ctl_set_stream(FILE* tls_file, FILE* base_stream)`.

2.6 How The Parts Fit Together

The driver-style implementation allows using TLS connections like normal streams. But this abstraction covers only the TLS-record-level and the TLS-alert-protocol. The implementation of handshake itself uses this level of abstraction and can therefore not be integrated into the driver. Therefore it is necessary to initiate a handshake manually after opening the connection. Also renegotiation is not implemented for this and security reasons.

2.7 Implementation content

2.7.1 Header Files

<code>hexdump.h</code>	hexdump routines for debugging purposes
<code>sec_memcmp.h</code>	secure version of <code>memcmp()</code>
<code>server.key.h</code>	declaration of TLS-private-key structures
<code>tls_alert_protocol.h</code>	declaration of alert functions
<code>tls_alert_types.h</code>	declaration of alert related types and values
<code>tls_certificate.h</code>	declaration of TLS-certificate structures
<code>tls_cipher_suites.h</code>	declaration of TLS-cipher-suites related types and values
<code>tls_crypto.h</code>	declaration of cryptographic core functions
<code>tls_driver.h</code>	declaration of driver related functions and types
<code>tls_handshake.h</code>	declaration of handshake function
<code>tls_handshake_protocol.h</code>	
<code>tls_random.h</code>	declaration of PRNG related functions
<code>tls_record_layer.h</code>	declaration of record-level TLS functions
<code>tls_resumption.h</code>	declaration of functions and types related to TLS-session-resumption

2.7.2 Implementation Files

<code>hexdump.c</code>	hexdump routines for debugging purposes
<code>nuttls_test_system.cert.tls.elf</code>	certificate wrapped into an elf-object
<code>sec_memcmp.c</code>	secure version of <code>memcmp()</code>
<code>server.key.c</code>	structures containing the private key material
<code>tls_alert_protocol.c</code>	alert protocol related functions
<code>tls_cipher_suites.c</code>	structures and functions related to TLS-cipher-suites
<code>tls_crypto.c</code>	functions handling low level symmetric crypto
<code>tls_driver.c</code>	driver related functions and structures
<code>tls_handshake.c</code>	functions for the TLS-handshake-protocol
<code>tls_random.c</code>	PRNG wrapper around entropy ¹
<code>tls_record_layer_common.c</code>	functions shared between TLS-record-layer transmit and receive functions
<code>tls_record_layer_rx.c</code>	functions handling the reception of TLS-record-layer packets
<code>tls_record_layer_tx.c</code>	functions handling the transmission of TLS-record-layer packets
<code>tls_resumption.c</code>	functions for managing the resumption table
<code>tls_test.c</code>	Demo application implementing a HTTPS server

3 Basic Data-structures

The two most important structures are `tls_connection_t` and `tls_record_connection_t`. `tls_connection_t` contains the driver related data, while `tls_record_connection_t` contains the TLS related data.

¹entropy is part of the ARM-Crypto-Lib

3.1 `tls_connection_t`

```
1 typedef struct _tls_connection_t tls_connection_t;
2
3 struct _tls_connection_t {
4     char *name;
5     tls_connection_t *next;
6     tls_record_connection_t connection;
7     void *write_buffer;
8     size_t write_buffer_size;
9     size_t write_buffer_fill;
10    tls_read_buffer_block_header_t *read_buffer;
11    uint8_t write_buffer_protocol;
12    uint8_t closing;
13    uint8_t write_lock;
14 };
```

3.1.1 `name`

The `name` field contains a pointer to a string of characters representing the name given to the connection during `fopen()`. The name is actual the part after the initial `"tls:"`-prefix of the path.

3.1.2 `next`

The `next` pointer points to the next connection in memory, so that all connections form a linked list.

3.1.3 `connection`

The `connection` field is a `tls_record_connection` (not a pointer), holding the parameters of the TLS connection.

3.1.4 `write_buffer`

The `write_buffer` pointer points to some memory holding the write buffer.

3.1.5 `write_buffer_size`

`write_buffer_size` is the size of the buffer. The size is considered to be bytes.

3.1.6 write_buffer_fill

`write_buffer_fill` is the amount of actual data in the write buffer. Its unit is bytes.

3.1.7 read_buffer

`read_buffer` points to the first block header for the read buffer.

3.1.8 write_buffer_protocol

`write_buffer_protocol` is the numeric protocol id to use for data transmission.

3.1.9 closing

`closing` is a flag which indicates if the connection is currently in the process of closing. If it is closing the field reads one else it reads zero.

3.1.10 write_lock

`write_lock` is a flag indicating if a write is in process. Write requests will be blocked while this field is non-zero.

3.2 tls_record_connection_t

```
1 typedef struct _tls_record_connection_t tls_record_connection_t;
2
3 struct _tls_record_connection_t {
4     FILE* base_stream;
5     uint32_t session_id;
6     union __attribute__((packed)){
7         struct {
8             uint8_t major;
9             uint8_t minor;
10        } names;
11        uint16_t id16;
12    } version;
13    uint8_t secure;
14    uint8_t post_clienthello;
15    tls_record_structure_pattern_t tx_pre_pattern;
16    tls_record_ciphertext_struct_t rx_pre_pattern;
17    tls_sec_parameters_t sec_parameters;
```

```
18     tls_record_states_t  sec_states_local;
19     tls_record_states_t  sec_states_remote;
20     tls_record_pending_state_t *pending;
21     uint8_t pending_status;
22     tls_error_state_t error_state;
23 };
```

3.2.1 base_stream

`base_stream` is a pointer to the base stream.

3.2.2 session_id

`session_id` is a 32-bit session id assigned to the connection. It is automatically generated during handshake and can be used for session resumption.

3.2.3 version

The `version` field is a union holding the TLS protocol version used for communication. It can be used as 16-bit id (`id16`) or as major (`names.major`) and minor (`names.minor`) part.

3.2.4 secure

The `secure` flag is used to signalize if a successful handshake has already took place. It reads zero before the initial handshake and one afterwards.

3.2.5 post_clienthello

The `post_clienthello` flag is used to signalize if checks for packet version should be relaxed. This is necessary since the initial `clienthello` may use a older record layer version for compatibility.

3.2.6 tx_pre_pattern

The `tx_pre_pattern` field holds precomputed offset information for sending packets. These data are computed after the crypto parameters are negotiated.

3.2.7 rx_pre_pattern

The `rx_pre_pattern` field holds precomputed offset information for receiving packets. These data are computed after the crypto parameters are negotiated.

3.2.8 sec_parameters

The `sec_parameters` field holds the negotiated crypto parameters (incl. `master_secret`, `client_random` and `server_random`).

3.2.9 sec_states_local

The `sec_states_local` field holds the contexts for cryptographic algorithms for the sending side.

3.2.10 sec_states_remote

The `sec_states_remote` field holds the contexts for cryptographic algorithms for the receiving side.

3.2.11 pending

`pending` points to a structure which contains the cryptographic parameters just negotiated but not yet active. The will be used after receiving a `change_cipher_spec` message.

3.2.12 pending_status

`pending_status` signalizes for which direction there are currently pending states.

3.2.13 error_state

`error_state` contains information about the last error which occurred.

4 Buffering

There are two buffers per connection, one for the receiving side and one for the transmitting side.

4.1 Write Buffer

The write buffer implementation is quite simple. The buffer consists of a memory area allocated by `create_new_connection()` (in `tls_driver.h`). All data which should be send out normally are just copied into the buffer and the whole buffer is transmitted (by `tls_connection_send_buffer` (in `tls_driver.h`)) when it is full. An exception are messages which are to large for the buffer. In this case the buffer is send on the wire and then blocks of the message which all have the size of the buffer. If data remains (which will not fill a whole buffer) it is copied into the buffer.

4.1.1 Allocation

The whole allocation of the buffer happens in `create_new_connection()` (in `tls_driver.h`). The two macros `INITIAL_WRITE_BUFFER_SIZE` and `MIN_WRITE_BUFFER_SIZE` are used to configure the allocation process. The function first tries to allocate `INITIAL_WRITE_BUFFER_SIZE` bytes of memory, if that fails it successively reduces the memory request to the half of the previous request and tries again, until the requested is fulfilled or the requested size is below `MIN_WRITE_BUFFER_SIZE`. If the requested size is lower than `MIN_WRITE_BUFFER_SIZE` an error is returned.

4.2 Read Buffer

Since TLS is a packet based protocol, always whole packets are received and buffered as such. The buffered packets are managed by a linked list of header blocks which contain the management information for those packets.

```
1 struct tls_read_buffer_block_header_st {
2     uint8_t *head;
3     uint8_t *buffer;
4     tls_read_buffer_block_header_t *next;
5     size_t remaining_bytes;
6     uint8_t protocol;
7 };
```

5 Packet Reception

The process of packet reception starts in `read_packet_to_buffer()` (in `tls_driver.h`) which is called if the read buffer could not fulfill a read request (see 4.1.1).

First it allocates three buffers:

- **buffer** for holding the packet content
- **packet_info** for holding information generated by packet decryption
- **header** for holding information about the memory block (for inclusion in the buffers linked list)

Then it reads the 5 byte packet header from the base stream into the just allocated buffer **buffer**.

offset	size	meaning
0 (byte)	1 (byte)	sub-protocol
1 (byte)	2 (byte)	protocol version
3 (byte)	2 (byte)	remaining data bytes in packet

The standard limits the maximum amount of remaining bytes to $2^{14} + 2048$ (rfc5246 page 21). This limit is checked, then optionally a compile time specified limit (the macro `MAX_RX_PACKET_LENGTH`). If the length checks are passed the size of the header buffer is expanded so it can hold the whole packet.

Then it reads the remaining part of the packet from the base stream into **buffer**.

Now the decryption routine `tls_record_decode()` gets called which decrypts the packet in place and returns data in an information structure (`tls_record_packet_info_t`).

Here the packet version is checked (for all packets except an initial `client_hello`) and the length is checked again and decryption is delegated to `tls_crypto_decipher_packet()`

In `tls_crypto_decipher_packet()` the decryption work is delegated to the CBC or stream functions of the ARM-Crypto-Lib.

Back in `tls_record_decode()` the next thing is trying to compute the structure of the packet. this computation is accelerated by precomputed tables based on the negotiated parameters (`compute_pattern()`). After the structure is computed the Message Authentication Code (MAC) is checked (`check_mac()`) and then the message padding (`check_padding()`).

If any check goes wrong a proper alert message is generated and send to the base stream.

The driver, after `tls_record_decode()` returns, copies necessary information from `packet_info` into `header` and frees `packet_info`. The next step s checking the protocol if it is set to `alert_proto` (21) and takes action depending on what kind of alert message was received (see 8 on page 16).

If the message is no fatal alert message or a `close_notify` it is chained into the linked list forming the read buffer.

6 Packet Transmission

The process of packet transmission starts in `tls_connection_send_buffer()` (in `driver.h`) which is called when the write buffer is flushed (which occurs if the write buffer is full or an external flush is issued also see 4.1 on page 13

`tls_connection_send_buffer()` itself first checks the validity of the given protocol number (it must be 20, 21, 22 or 23). After this basic check memory is allocated for the `pinfo` structure of type `tls_record_packet_info_t`. This structure is filled with the available data.

If the packet needs padding a valid padding length is randomly selected by `random_pad_length()`.

Then a buffer for the packet is allocated. All further processing of the packet takes place in this buffer.

After the address of the buffer is fixed, the structure of the packet is computed with absolute values by `compute_pattern()`.

Now the buffer is filled with copied data and a header is prepended.

7 The Handshake Process

The handshake initializes the cryptographic parameters of a TLS connection. It is run at the beginning of a session and is initiated by the client. Multiple messages are exchanged of which some are optional. A normal handshake (as shown by figure 1 on the next page) involves asymmetric cryptographic computations and so is very computing intense. The standard allows re-handshaking, which means doing a handshake via an already initialized session. This is not supported since it might be a security problem². An alternative to do a "normal" full handshake is to do an abbreviated handshake to resume a session using common secrets that already have been computed between the two parties. This abbreviated handshake is shown in figure ?? on page ?. It is very fast since only symmetric primitives are used to compute the new keys (derived from the common secret and the random values exchanged with the hello messages).

The handshake implementation (`tls_handshake.c`) works on top of the record layer. It is currently restricted to RSA cipher suites but can easily

²See RFC 5746 for a description of the attack

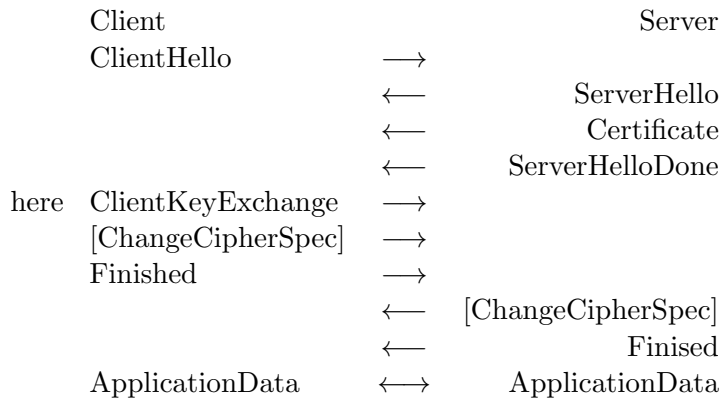


Figure 1: TLS handshake as done by Nut/TLS

be replaced. The handshake handling routine (`make_handshake()`) is called directly by the application after opening the Nut/TLS stream file. `make_handshake()` waits for an incoming `client_hello` message and then starts the negotiation. It returns zero on success and non-zero if an error occurred.

Since all handshake messages start with a three byte length field, a generic function is used to read handshake messages and allocate memory for them (`read_handshake_packet`).

The certificate message send to the client is stored completely (with handshake header) in FLASH. It is created by specific tools (`cert-merge`) before compilation and is linked in during the linking stage. Since usage of the data positioned by the linker requires knowledge of the symbols referring to the binary blob, those are defined as macros (`CERT_BLOB_START`, `CERT_BLOB_END` and `CERT_BLOB_SIZE`).

The private RSA key is stored in `server.key.c` which defines multiple biginteger objects needed to do the RSA computations.

Users **should** replace the certificate and the key by self created ones where the key should be kept strictly secret.

8 Alert Handling

Alert handling is directly done by the driver and the record layer. The main purpose of the alert protocol is to signal warnings, errors and regular termination of the communication. The last is especially important since abnormal termination by a fatal error causes the session to be not resumable. Handling of incoming alert messages is done directly after the packet

is decrypted and before it would be stored in the read buffer. Every alert is send to the `tls_alert_receive()` function (in `alert_protocol.c`) and further checked and processed. This would be the place to handle specific errors if the application needs such handling. The return value of `tls_alert_receive()` signal if the connection should be terminated or if it can further operate. If a fatal error or a `close_notify` occurs the connection will be terminated, which is handled by the function in the driver.

Sending alert messages (due to problems) happens by a call to `tls_alert_send()`. This function can also handle an additional comment which will also be logged (if logging is enabled).

References