Linux Kernel Crypto API

Stephan Mueller <smueller@chronox.de>
Marek Vasut <marek@denx.de>

Linux Kernel Crypto API

by Stephan Mueller and Marek Vasut
Copyright © 2014 Stephan Mueller

This documentation is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

For more details see the file COPYING in the source distribution of Linux.

Table of Contents

1. Kernel Crypto API Interface Specification 1
    Introduction 1
    Terminology 1
2. Kernel Crypto API Architecture 3
    Cipher algorithm types 3
    Ciphers And Templates 3
    Synchronous And Asynchronous Operation 3
    Crypto API Cipher References And Priority 4
    Key Sizes 5
    Cipher Allocation Type And Masks 5
    Internal Structure of Kernel Crypto API 6
        Generic AEAD Cipher Structure 6
        Generic Block Cipher Structure 8
        Generic Keyed Message Digest Structure 8
3. Developing Cipher Algorithms 10
    Registering And Unregistering Transformation 10
    Single-Block Symmetric Ciphers [CIPHER] 10
        Registration specifics 10
        Cipher Definition With struct cipher_alg 11
    Multi-Block Ciphers [BLKCIPHER] [ABLKCIPHER] 11
        Registration Specifics 11
        Cipher Definition With struct blkcipher_alg and ablkcipher_alg 11
        Specifics Of Asynchronous Multi-Block Cipher 12
    Hashing [HASH] 12
        Registering And Unregistering The Transformation 12
        Cipher Definition With struct shash_alg and ahash_alg 12
        Specifics Of Asynchronous HASH Transformation 13
4. User Space Interface 14
    Introduction 14
    User Space API General Remarks 14
    In-place Cipher operation 15
    Message Digest API 15
    Symmetric Cipher API 16
    AEAD Cipher API 16
        AEAD Memory Structure 17
    Random Number Generator API 18
    Zero-Copy Interface 18
    Setsockopt Interface 19
    User space API example 19
5. Programming Interface 20
    Block Cipher Context Data Structures 20
    Block Cipher Algorithm Definitions 21
    Asynchronous Block Cipher API 30
    Asynchronous Cipher Request Handle 40
    Authenticated Encryption With Associated Data (AEAD) Cipher API 46
    Asynchronous AEAD Request Handle 56
    Synchronous Block Cipher API 64
    Single Block Cipher API 77
    Synchronous Message Digest API 84
    Message Digest Algorithm Definitions 94
    Asynchronous Message Digest API 98

Asynchronous Hash Request Handle 111
Synchronous Message Digest API 116
Crypto API Random Number API 129
6. Code Examples 136
Code Example For Asynchronous Block Cipher Operation 136
Code Example For Synchronous Block Cipher Operation 138
Code Example For Use of Operational State Memory With SHASH 140
Code Example For Random Number Generator Usage 140

Chapter 1. Kernel Crypto API Interface Specification

Introduction

The kernel crypto API offers a rich set of cryptographic ciphers as well as other data transformation mechanisms and methods to invoke these. This document contains a description of the API and provides example code.

To understand and properly use the kernel crypto API a brief explanation of its structure is given. Based on the architecture, the API can be separated into different components. Following the architecture specification, hints to developers of ciphers are provided. Pointers to the API function call documentation are given at the end.

The kernel crypto API refers to all algorithms as "transformations". Therefore, a cipher handle variable usually has the name "tfm". Besides cryptographic operations, the kernel crypto API also knows compression transformations and handles them the same way as ciphers.

The kernel crypto API serves the following entity types:

- consumers requesting cryptographic services
- data transformation implementations (typically ciphers) that can be called by consumers using the kernel crypto API

This specification is intended for consumers of the kernel crypto API as well as for developers implementing ciphers. This API specification, however, does not discuss all API calls available to data transformation implementations (i.e. implementations of ciphers and other transformations (such as CRC or even compression algorithms) that can register with the kernel crypto API).

Note: The terms "transformation" and cipher algorithm are used interchangeably.

Terminology

The transformation implementation is an actual code or interface to hardware which implements a certain transformation with precisely defined behavior.

The transformation object (TFM) is an instance of a transformation implementation. There can be multiple transformation objects associated with a single transformation implementation. Each of those transformation objects is held by a crypto API consumer or another transformation. Transformation object is allocated when a crypto API consumer requests a transformation implementation. The consumer is then provided with a structure, which contains a transformation object (TFM).

The structure that contains transformation objects may also be referred to as a "cipher handle". Such a cipher handle is always subject to the following phases that are reflected in the API calls applicable to such a cipher handle:

1. Initialization of a cipher handle.
2. Execution of all intended cipher operations applicable for the handle where the cipher handle must be furnished to every API call.

3. Destruction of a cipher handle.

When using the initialization API calls, a cipher handle is created and returned to the consumer. Therefore, please refer to all initialization API calls that refer to the data structure type a consumer is expected to receive and subsequently to use. The initialization API calls have all the same naming conventions of crypto_alloc_*.

The transformation context is private data associated with the transformation object.

Chapter 2. Kernel Crypto API Architecture

Cipher algorithm types

The kernel crypto API provides different API calls for the following cipher types:

- Symmetric ciphers
- AEAD ciphers
- Message digest, including keyed message digest
- Random number generation
- User space interface

Ciphers And Templates

The kernel crypto API provides implementations of single block ciphers and message digests. In addition, the kernel crypto API provides numerous "templates" that can be used in conjunction with the single block ciphers and message digests. Templates include all types of block chaining mode, the HMAC mechanism, etc.

Single block ciphers and message digests can either be directly used by a caller or invoked together with a template to form multi-block ciphers or keyed message digests.

A single block cipher may even be called with multiple templates. However, templates cannot be used without a single cipher.

See /proc/crypto and search for "name". For example:

- aes
- ecb(aes)
- cmac(aes)
- ccm(aes)
- rfc4106(gcm(aes))
- sha1
- hmac(sha1)
- authenc(hmac(sha1),cbc(aes))

In these examples, "aes" and "sha1" are the ciphers and all others are the templates.

Synchronous And Asynchronous Operation

The kernel crypto API provides synchronous and asynchronous API operations.

When using the synchronous API operation, the caller invokes a cipher operation which is performed synchronously by the kernel crypto API. That means, the caller waits until the cipher operation completes. Therefore, the kernel crypto API calls work like regular function calls. For synchronous operation, the set of API calls is small and conceptually similar to any other crypto library.

Asynchronous operation is provided by the kernel crypto API which implies that the invocation of a cipher operation will complete almost instantly. That invocation triggers the cipher operation but it does not signal its completion. Before invoking a cipher operation, the caller must provide a callback function the kernel crypto API can invoke to signal the completion of the cipher operation. Furthermore, the caller must ensure it can handle such asynchronous events by applying appropriate locking around its data. The kernel crypto API does not perform any special serialization operation to protect the caller's data integrity.

Crypto API Cipher References And Priority

A cipher is referenced by the caller with a string. That string has the following semantics:

```
template(single block cipher)
```

where "template" and "single block cipher" is the aforementioned template and single block cipher, respectively. If applicable, additional templates may enclose other templates, such as

```
template1(template2(single block cipher)))
```

The kernel crypto API may provide multiple implementations of a template or a single block cipher. For example, AES on newer Intel hardware has the following implementations: AES-NI, assembler implementation, or straight C. Now, when using the string "aes" with the kernel crypto API, which cipher implementation is used? The answer to that question is the priority number assigned to each cipher implementation by the kernel crypto API. When a caller uses the string to refer to a cipher during initialization of a cipher handle, the kernel crypto API looks up all implementations providing an implementation with that name and selects the implementation with the highest priority.

Now, a caller may have the need to refer to a specific cipher implementation and thus does not want to rely on the priority-based selection. To accommodate this scenario, the kernel crypto API allows the cipher implementation to register a unique name in addition to common names. When using that unique name, a caller is therefore always sure to refer to the intended cipher implementation.

The list of available ciphers is given in /proc/crypto. However, that list does not specify all possible permutations of templates and ciphers. Each block listed in /proc/crypto may contain the following information -- if one of the components listed as follows are not applicable to a cipher, it is not displayed:

- name: the generic name of the cipher that is subject to the priority-based selection -- this name can be used by the cipher allocation API calls (all names listed above are examples for such generic names)
- driver: the unique name of the cipher -- this name can be used by the cipher allocation API calls
- module: the kernel module providing the cipher implementation (or "kernel" for statically linked ciphers)
- priority: the priority value of the cipher implementation
- refcnt: the reference count of the respective cipher (i.e. the number of current consumers of this cipher)
- selftest: specification whether the self test for the cipher passed

- type:
  - blkcipher for synchronous block ciphers
  - ablkcipher for asynchronous block ciphers
  - cipher for single block ciphers that may be used with an additional template
  - shash for synchronous message digest
  - ahash for asynchronous message digest
  - aead for AEAD cipher type
  - compression for compression type transformations
  - rng for random number generator
  - givcipher for cipher with associated IV generator (see the geniv entry below for the specification of the IV generator type used by the cipher implementation)
- blocksize: blocksize of cipher in bytes
- keysize: key size in bytes
- ivsize: IV size in bytes
- seedsize: required size of seed data for random number generator
- digestsize: output size of the message digest
- geniv: IV generation type:
  - eseqiv for encrypted sequence number based IV generation
  - seqiv for sequence number based IV generation
  - chainiv for chain iv generation
  - <builtin> is a marker that the cipher implements IV generation and handling as it is specific to the given cipher

Key Sizes

When allocating a cipher handle, the caller only specifies the cipher type. Symmetric ciphers, however, typically support multiple key sizes (e.g. AES-128 vs. AES-192 vs. AES-256). These key sizes are determined with the length of the provided key. Thus, the kernel crypto API does not provide a separate way to select the particular symmetric cipher key size.

Cipher Allocation Type And Masks

The different cipher handle allocation functions allow the specification of a type and mask flag. Both parameters have the following meaning (and are therefore not covered in the subsequent sections).

The type flag specifies the type of the cipher algorithm. The caller usually provides a 0 when the caller wants the default handling. Otherwise, the caller may provide the following selections which match the the aforementioned cipher types:

- CRYPTO_ALG_TYPE_CIPHER Single block cipher
- CRYPTO_ALG_TYPE_COMPRESS Compression
- CRYPTO_ALG_TYPE_AEAD Authenticated Encryption with Associated Data (MAC)
- CRYPTO_ALG_TYPE_BLKCIPHER Synchronous multi-block cipher
- CRYPTO_ALG_TYPE_ABLKCIPHER Asynchronous multi-block cipher
- CRYPTO_ALG_TYPE_GIVCIPHER Asynchronous multi-block cipher packed together with an IV generator (see geniv field in the /proc/crypto listing for the known IV generators)
- CRYPTO_ALG_TYPE_DIGEST Raw message digest
- CRYPTO_ALG_TYPE_HASH Alias for CRYPTO_ALG_TYPE_DIGEST
- CRYPTO_ALG_TYPE_SHASH Synchronous multi-block hash
- CRYPTO_ALG_TYPE_AHASH Asynchronous multi-block hash
- CRYPTO_ALG_TYPE_RNG Random Number Generation
- CRYPTO_ALG_TYPE_PCOMPRESS Enhanced version of CRYPTO_ALG_TYPE_COMPRESS allowing for segmented compression / decompression instead of performing the operation on one segment only. CRYPTO_ALG_TYPE_PCOMPRESS is intended to replace CRYPTO_ALG_TYPE_COMPRESS once existing consumers are converted.

The mask flag restricts the type of cipher. The only allowed flag is CRYPTO_ALG_ASYNC to restrict the cipher lookup function to asynchronous ciphers. Usually, a caller provides a 0 for the mask flag.

When the caller provides a mask and type specification, the caller limits the search the kernel crypto API can perform for a suitable cipher implementation for the given cipher name. That means, even when a caller uses a cipher name that exists during its initialization call, the kernel crypto API may not select it due to the used type and mask field.

Internal Structure of Kernel Crypto API

The kernel crypto API has an internal structure where a cipher implementation may use many layers and indirections. This section shall help to clarify how the kernel crypto API uses various components to implement the complete cipher.

The following subsections explain the internal structure based on existing cipher implementations. The first section addresses the most complex scenario where all other scenarios form a logical subset.

Generic AEAD Cipher Structure

The following ASCII art decomposes the kernel crypto API layers when using the AEAD cipher with the automated IV generation. The shown example is used by the IPSEC layer.

For other use cases of AEAD ciphers, the ASCII art applies as well, but the caller may not use the AEAD cipher with a separate IV generator. In this case, the caller must generate the IV.

The depicted example decomposes the AEAD cipher of GCM(AES) based on the generic C implementations (gcm.c, aes-generic.c, ctr.c, ghash-generic.c, seqiv.c). The generic implementation serves as an example showing the complete logic of the kernel crypto API.

It is possible that some streamlined cipher implementations (like AES-NI) provide implementations merging aspects which in the view of the kernel crypto API cannot be decomposed into layers any more. In case of the AES-NI implementation, the CTR mode, the GHASH implementation and the AES cipher are all merged into one cipher implementation registered with the kernel crypto API. In this case, the concept described by the following ASCII art applies too. However, the decomposition of GCM into the individual sub-components by the kernel crypto API is not done any more.

Each block in the following ASCII art is an independent cipher instance obtained from the kernel crypto API. Each block is accessed by the caller or by other blocks using the API functions defined by the kernel crypto API for the cipher implementation type.

The blocks below indicate the cipher type as well as the specific logic implemented in the cipher.

The ASCII art picture also indicates the call structure, i.e. who calls which component. The arrows point to the invoked block where the caller uses the API applicable to the cipher type specified for the block.

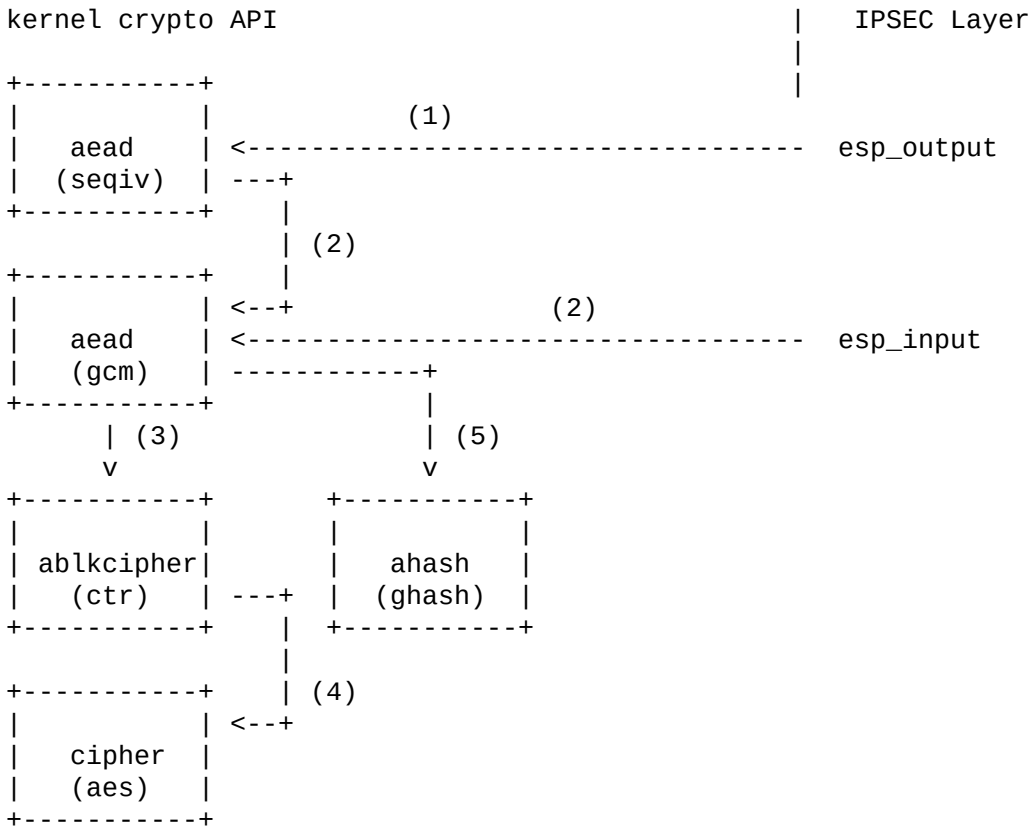

The following call sequence is applicable when the IPSEC layer triggers an encryption operation with the esp_output function. During configuration, the administrator set up the use of rfc4106(gcm(aes)) as the cipher for ESP. The following call sequence is now depicted in the ASCII art above:

1. esp_output() invokes crypto_aead_encrypt() to trigger an encryption operation of the AEAD cipher with IV generator.

   In case of GCM, the SEQIV implementation is registered as GIVCIPHER in crypto_rfc4106_alloc().

The SEQIV performs its operation to generate an IV where the core function is seqiv_geniv().

2. Now, SEQIV uses the AEAD API function calls to invoke the associated AEAD cipher. In our case, during the instantiation of SEQIV, the cipher handle for GCM is provided to SEQIV. This means that SEQIV invokes AEAD cipher operations with the GCM cipher handle.

   During instantiation of the GCM handle, the CTR(AES) and GHASH ciphers are instantiated. The cipher handles for CTR(AES) and GHASH are retained for later use.

   The GCM implementation is responsible to invoke the CTR mode AES and the GHASH cipher in the right manner to implement the GCM specification.

3. The GCM AEAD cipher type implementation now invokes the ABLKCIPHER API with the instantiated CTR(AES) cipher handle.

   During instantiation of the CTR(AES) cipher, the CIPHER type implementation of AES is instantiated. The cipher handle for AES is retained.

   That means that the ABLKCIPHER implementation of CTR(AES) only implements the CTR block chaining mode. After performing the block chaining operation, the CIPHER implementation of AES is invoked.

4. The ABLKCIPHER of CTR(AES) now invokes the CIPHER API with the AES cipher handle to encrypt one block.

5. The GCM AEAD implementation also invokes the GHASH cipher implementation via the AHASH API.

When the IPSEC layer triggers the esp_input() function, the same call sequence is followed with the only difference that the operation starts with step (2).

Generic Block Cipher Structure

Generic block ciphers follow the same concept as depicted with the ASCII art picture above.

For example, CBC(AES) is implemented with cbc.c, and aes-generic.c. The ASCII art picture above applies as well with the difference that only step (4) is used and the ABLKCIPHER block chaining mode is CBC.

Generic Keyed Message Digest Structure

Keyed message digest implementations again follow the same concept as depicted in the ASCII art picture above.

For example, HMAC(SHA256) is implemented with hmac.c and sha256_generic.c. The following ASCII art illustrates the implementation:

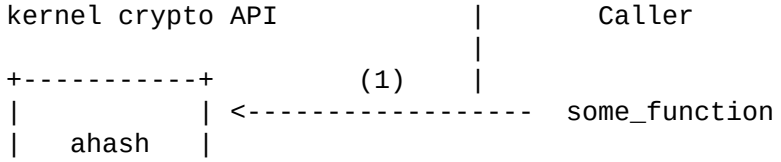

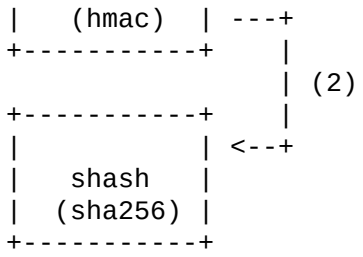

The following call sequence is applicable when a caller triggers an HMAC operation:

1. The AHASH API functions are invoked by the caller. The HMAC implementation performs its operation as needed.

   During initialization of the HMAC cipher, the SHASH cipher type of SHA256 is instantiated. The cipher handle for the SHA256 instance is retained.

   At one time, the HMAC implementation requires a SHA256 operation where the SHA256 cipher handle is used.

2. The HMAC instance now invokes the SHASH API with the SHA256 cipher handle to calculate the message digest.

Chapter 3. Developing Cipher Algorithms

Registering And Unregistering Transformation

There are three distinct types of registration functions in the Crypto API. One is used to register a generic cryptographic transformation, while the other two are specific to HASH transformations and COMPRESSion. We will discuss the latter two in a separate chapter, here we will only look at the generic ones.

Before discussing the register functions, the data structure to be filled with each, struct crypto_alg, must be considered -- see below for a description of this data structure.

The generic registration functions can be found in include/linux/crypto.h and their definition can be seen below. The former function registers a single transformation, while the latter works on an array of transformation descriptions. The latter is useful when registering transformations in bulk.

```
int crypto_register_alg(struct crypto_alg *alg);
int crypto_register_algs(struct crypto_alg *algs, int count);
```

The counterparts to those functions are listed below.

```
int crypto_unregister_alg(struct crypto_alg *alg);
int crypto_unregister_algs(struct crypto_alg *algs, int count);
```

Notice that both registration and unregistration functions do return a value, so make sure to handle errors. A return code of zero implies success. Any return code < 0 implies an error.

The bulk registration / unregistration functions require that struct crypto_alg is an array of count size. These functions simply loop over that array and register / unregister each individual algorithm. If an error occurs, the loop is terminated at the offending algorithm definition. That means, the algorithms prior to the offending algorithm are successfully registered. Note, the caller has no way of knowing which cipher implementations have successfully registered. If this is important to know, the caller should loop through the different implementations using the single instance *_alg functions for each individual implementation.

Single-Block Symmetric Ciphers [CIPHER]

Example of transformations: aes, arc4, ...

This section describes the simplest of all transformation implementations, that being the CIPHER type used for symmetric ciphers. The CIPHER type is used for transformations which operate on exactly one block at a time and there are no dependencies between blocks at all.

Registration specifics

The registration of [CIPHER] algorithm is specific in that struct crypto_alg field .cra_type is empty. The .cra_u.cipher has to be filled in with proper callbacks to implement this transformation.

See struct cipher_alg below.

Cipher Definition With struct cipher_alg

Struct cipher_alg defines a single block cipher.

Here are schematics of how these functions are called when operated from other part of the kernel. Note that the .cia_setkey() call might happen before or after any of these schematics happen, but must not happen during any of these are in-flight.

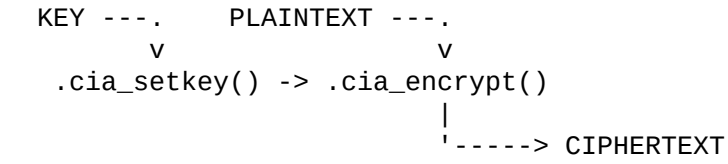

Please note that a pattern where .cia_setkey() is called multiple times is also valid:

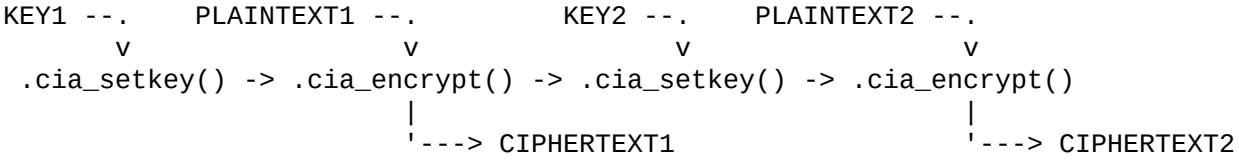

Multi-Block Ciphers [BLKCIPHER] [ABLKCIPHER]

Example of transformations: cbc(aes), ecb(arc4), ...

This section describes the multi-block cipher transformation implementations for both synchronous [BLKCIPHER] and asynchronous [ABLKCIPHER] case. The multi-block ciphers are used for transformations which operate on scatterlists of data supplied to the transformation functions. They output the result into a scatterlist of data as well.

Registration Specifics

The registration of [BLKCIPHER] or [ABLKCIPHER] algorithms is one of the most standard procedures throughout the crypto API.

Note, if a cipher implementation requires a proper alignment of data, the caller should use the functions of crypto_blkcipher_alignmask() or crypto_ablkcipher_alignmask() respectively to identify a memory alignment mask. The kernel crypto API is able to process requests that are unaligned. This implies, however, additional overhead as the kernel crypto API needs to perform the realignment of the data which may imply moving of data.

Cipher Definition With struct blkcipher_alg and ablkcipher_alg

Struct blkcipher_alg defines a synchronous block cipher whereas struct ablkcipher_alg defines an asynchronous block cipher.

Please refer to the single block cipher description for schematics of the block cipher usage. The usage patterns are exactly the same for [ABLKCIPHER] and [BLKCIPHER] as they are for plain [CIPHER].

Specifics Of Asynchronous Multi-Block Cipher

There are a couple of specifics to the [ABLKCIPHER] interface.

First of all, some of the drivers will want to use the Generic ScatterWalk in case the hardware needs to be fed separate chunks of the scatterlist which contains the plaintext and will contain the ciphertext. Please refer to the ScatterWalk interface offered by the Linux kernel scatter / gather list implementation.

Hashing [HASH]

Example of transformations: crc32, md5, sha1, sha256,...

Registering And Unregistering The Transformation

There are multiple ways to register a HASH transformation, depending on whether the transformation is synchronous [SHASH] or asynchronous [AHASH] and the amount of HASH transformations we are registering. You can find the prototypes defined in include/crypto/internal/hash.h:

```
int crypto_register_ahash(struct ahash_alg *alg);

int crypto_register_shash(struct shash_alg *alg);
int crypto_register_shashes(struct shash_alg *algs, int count);
```

The respective counterparts for unregistering the HASH transformation are as follows:

```
int crypto_unregister_ahash(struct ahash_alg *alg);

int crypto_unregister_shash(struct shash_alg *alg);
int crypto_unregister_shashes(struct shash_alg *algs, int count);
```

Cipher Definition With struct shash_alg and ahash_alg

Here are schematics of how these functions are called when operated from other part of the kernel. Note that the .setkey() call might happen before or after any of these schematics happen, but must not happen during any of these are in-flight. Please note that calling .init() followed immediately by .finish() is also a perfectly valid transformation.

```
I)   DATA -----------.
                     v
      .init() -> .update() -> .final()      ! .update() might not be called
                  ^    |         |            at all in this scenario.
                  '----'         '---> HASH

II)  DATA -----------.-----------.
                     v           v
```

```
      .init() -> .update() -> .finup()      ! .update() may not be called
                  ^    |         |            at all in this scenario.
                  '----'         '---> HASH

III) DATA -----------.
                     v
                  .digest()                 ! The entire process is handled
                     |                        by the .digest() call.
                     '---------------> HASH
```

Here is a schematic of how the .export()/.import() functions are called when used from another part of the kernel.

```
KEY--.                 DATA--.
     v                       v                 ! .update() may not be called
 .setkey() -> .init() -> .update() -> .export()   at all in this scenario.
                          ^     |         |
                          '-----'         '--> PARTIAL_HASH

----------- other transformations happen here -----------

PARTIAL_HASH--.   DATA1--.
              v          v
          .import -> .update() -> .final()     ! .update() may not be called
                      ^    |         |           at all in this scenario.
                      '----'         '--> HASH1

PARTIAL_HASH--.   DATA2-.
              v         v
          .import -> .finup()
                       |
                       '---------------> HASH2
```

Specifics Of Asynchronous HASH Transformation

Some of the drivers will want to use the Generic ScatterWalk in case the implementation needs to be fed separate chunks of the scatterlist which contains the input data. The buffer containing the resulting hash will always be properly aligned to .cra_alignmask so there is no need to worry about this.

Chapter 4. User Space Interface

Introduction

The concepts of the kernel crypto API visible to kernel space is fully applicable to the user space interface as well. Therefore, the kernel crypto API high level discussion for the in-kernel use cases applies here as well.

The major difference, however, is that user space can only act as a consumer and never as a provider of a transformation or cipher algorithm.

The following covers the user space interface exported by the kernel crypto API. A working example of this description is libkcapi that can be obtained from [1]. That library can be used by user space applications that require cryptographic services from the kernel.

Some details of the in-kernel kernel crypto API aspects do not apply to user space, however. This includes the difference between synchronous and asynchronous invocations. The user space API call is fully synchronous.

[1] http://www.chronox.de/libkcapi.html

User Space API General Remarks

The kernel crypto API is accessible from user space. Currently, the following ciphers are accessible:

- Message digest including keyed message digest (HMAC, CMAC)
- Symmetric ciphers
- AEAD ciphers
- Random Number Generators

The interface is provided via socket type using the type AF_ALG. In addition, the setsockopt option type is SOL_ALG. In case the user space header files do not export these flags yet, use the following macros:

```
#ifndef AF_ALG
#define AF_ALG 38
#endif
#ifndef SOL_ALG
#define SOL_ALG 279
#endif
```

A cipher is accessed with the same name as done for the in-kernel API calls. This includes the generic vs. unique naming schema for ciphers as well as the enforcement of priorities for generic names.

To interact with the kernel crypto API, a socket must be created by the user space application. User space invokes the cipher operation with the send()/write() system call family. The result of the cipher operation is obtained with the read()/recv() system call family.

The following API calls assume that the socket descriptor is already opened by the user space application and discusses only the kernel crypto API specific invocations.

To initialize the socket interface, the following sequence has to be performed by the consumer:

1. Create a socket of type AF_ALG with the struct sockaddr_alg parameter specified below for the different cipher types.

2. Invoke bind with the socket descriptor

3. Invoke accept with the socket descriptor. The accept system call returns a new file descriptor that is to be used to interact with the particular cipher instance. When invoking send/write or recv/read system calls to send data to the kernel or obtain data from the kernel, the file descriptor returned by accept must be used.

In-place Cipher operation

Just like the in-kernel operation of the kernel crypto API, the user space interface allows the cipher operation in-place. That means that the input buffer used for the send/write system call and the output buffer used by the read/recv system call may be one and the same. This is of particular interest for symmetric cipher operations where a copying of the output data to its final destination can be avoided.

If a consumer on the other hand wants to maintain the plaintext and the ciphertext in different memory locations, all a consumer needs to do is to provide different memory pointers for the encryption and decryption operation.

Message Digest API

The message digest type to be used for the cipher operation is selected when invoking the bind syscall. bind requires the caller to provide a filled struct sockaddr data structure. This data structure must be filled as follows:

```
struct sockaddr_alg sa = {
 .salg_family = AF_ALG,
 .salg_type = "hash", /* this selects the hash logic in the kernel */
 .salg_name = "sha1" /* this is the cipher name */
};
```

The salg_type value "hash" applies to message digests and keyed message digests. Though, a keyed message digest is referenced by the appropriate salg_name. Please see below for the setsockopt interface that explains how the key can be set for a keyed message digest.

Using the send() system call, the application provides the data that should be processed with the message digest. The send system call allows the following flags to be specified:

- MSG_MORE: If this flag is set, the send system call acts like a message digest update function where the final hash is not yet calculated. If the flag is not set, the send system call calculates the final message digest immediately.

With the recv() system call, the application can read the message digest from the kernel crypto API. If the buffer is too small for the message digest, the flag MSG_TRUNC is set by the kernel.

In order to set a message digest key, the calling application must use the setsockopt() option of ALG_SET_KEY. If the key is not set the HMAC operation is performed without the initial HMAC state change caused by the key.

Symmetric Cipher API

The operation is very similar to the message digest discussion. During initialization, the struct sockaddr data structure must be filled as follows:

```
struct sockaddr_alg sa = {
 .salg_family = AF_ALG,
 .salg_type = "skcipher", /* this selects the symmetric cipher */
 .salg_name = "cbc(aes)" /* this is the cipher name */
};
```

Before data can be sent to the kernel using the write/send system call family, the consumer must set the key. The key setting is described with the setsockopt invocation below.

Using the sendmsg() system call, the application provides the data that should be processed for encryption or decryption. In addition, the IV is specified with the data structure provided by the sendmsg() system call.

The sendmsg system call parameter of struct msghdr is embedded into the struct cmsghdr data structure. See recv(2) and cmsg(3) for more information on how the cmsghdr data structure is used together with the send/recv system call family. That cmsghdr data structure holds the following information specified with a separate header instances:

- specification of the cipher operation type with one of these flags:
  - ALG_OP_ENCRYPT - encryption of data
  - ALG_OP_DECRYPT - decryption of data
- specification of the IV information marked with the flag ALG_SET_IV

The send system call family allows the following flag to be specified:

- MSG_MORE: If this flag is set, the send system call acts like a cipher update function where more input data is expected with a subsequent invocation of the send system call.

Note: The kernel reports -EINVAL for any unexpected data. The caller must make sure that all data matches the constraints given in /proc/crypto for the selected cipher.

With the recv() system call, the application can read the result of the cipher operation from the kernel crypto API. The output buffer must be at least as large as to hold all blocks of the encrypted or decrypted data. If the output data size is smaller, only as many blocks are returned that fit into that output buffer size.

AEAD Cipher API

The operation is very similar to the symmetric cipher discussion. During initialization, the struct sockaddr data structure must be filled as follows:

```
struct sockaddr_alg sa = {
 .salg_family = AF_ALG,
 .salg_type = "aead", /* this selects the symmetric cipher */
 .salg_name = "gcm(aes)" /* this is the cipher name */
};
```

Before data can be sent to the kernel using the write/send system call family, the consumer must set the key. The key setting is described with the setsockopt invocation below.

In addition, before data can be sent to the kernel using the write/send system call family, the consumer must set the authentication tag size. To set the authentication tag size, the caller must use the setsockopt invocation described below.

Using the sendmsg() system call, the application provides the data that should be processed for encryption or decryption. In addition, the IV is specified with the data structure provided by the sendmsg() system call.

The sendmsg system call parameter of struct msghdr is embedded into the struct cmsghdr data structure. See recv(2) and cmsg(3) for more information on how the cmsghdr data structure is used together with the send/recv system call family. That cmsghdr data structure holds the following information specified with a separate header instances:

- specification of the cipher operation type with one of these flags:
  - ALG_OP_ENCRYPT - encryption of data
  - ALG_OP_DECRYPT - decryption of data
- specification of the IV information marked with the flag ALG_SET_IV
- specification of the associated authentication data (AAD) with the flag ALG_SET_AEAD_ASSOCLEN. The AAD is sent to the kernel together with the plaintext / ciphertext. See below for the memory structure.

The send system call family allows the following flag to be specified:

- MSG_MORE: If this flag is set, the send system call acts like a cipher update function where more input data is expected with a subsequent invocation of the send system call.

Note: The kernel reports -EINVAL for any unexpected data. The caller must make sure that all data matches the constraints given in /proc/crypto for the selected cipher.

With the recv() system call, the application can read the result of the cipher operation from the kernel crypto API. The output buffer must be at least as large as defined with the memory structure below. If the output data size is smaller, the cipher operation is not performed.

The authenticated decryption operation may indicate an integrity error. Such breach in integrity is marked with the -EBADMSG error code.

AEAD Memory Structure

The AEAD cipher operates with the following information that is communicated between user and kernel space as one data stream:

- plaintext or ciphertext
- associated authentication data (AAD)
- authentication tag

The sizes of the AAD and the authentication tag are provided with the sendmsg and setsockopt calls (see there). As the kernel knows the size of the entire data stream, the kernel is now able to calculate the right offsets of the data components in the data stream.

The user space caller must arrange the aforementioned information in the following order:

- AEAD encryption input: AAD || plaintext
- AEAD decryption input: AAD || ciphertext || authentication tag

The output buffer the user space caller provides must be at least as large to hold the following data:

- AEAD encryption output: ciphertext || authentication tag
- AEAD decryption output: plaintext

Random Number Generator API

Again, the operation is very similar to the other APIs. During initialization, the struct sockaddr data structure must be filled as follows:

```
struct sockaddr_alg sa = {
 .salg_family = AF_ALG,
 .salg_type = "rng", /* this selects the symmetric cipher */
 .salg_name = "drbg_nopr_sha256" /* this is the cipher name */
};
```

Depending on the RNG type, the RNG must be seeded. The seed is provided using the setsockopt interface to set the key. For example, the ansi_cprng requires a seed. The DRBGs do not require a seed, but may be seeded.

Using the read()/recvmsg() system calls, random numbers can be obtained. The kernel generates at most 128 bytes in one call. If user space requires more data, multiple calls to read()/recvmsg() must be made.

WARNING: The user space caller may invoke the initially mentioned accept system call multiple times. In this case, the returned file descriptors have the same state.

Zero-Copy Interface

In addition to the send/write/read/recv system call family, the AF_ALG interface can be accessed with the zero-copy interface of splice/vmsplice. As the name indicates, the kernel tries to avoid a copy operation into kernel space.

The zero-copy operation requires data to be aligned at the page boundary. Non-aligned data can be used as well, but may require more operations of the kernel which would defeat the speed gains obtained from the zero-copy interface.

The system-interent limit for the size of one zero-copy operation is 16 pages. If more data is to be sent to AF_ALG, user space must slice the input into segments with a maximum size of 16 pages.

Zero-copy can be used with the following code example (a complete working example is provided with libkcapi):

```
int pipes[2];

pipe(pipes);
```

```
/* input data in iov */
vmsplice(pipes[1], iov, iovlen, SPLICE_F_GIFT);
/* opfd is the file descriptor returned from accept() system call */
splice(pipes[0], NULL, opfd, NULL, ret, 0);
read(opfd, out, outlen);
```

Setsockopt Interface

In addition to the read/recv and send/write system call handling to send and retrieve data subject to the cipher operation, a consumer also needs to set the additional information for the cipher operation. This additional information is set using the setsockopt system call that must be invoked with the file descriptor of the open cipher (i.e. the file descriptor returned by the accept system call).

Each setsockopt invocation must use the level SOL_ALG.

The setsockopt interface allows setting the following data using the mentioned optname:

- ALG_SET_KEY -- Setting the key. Key setting is applicable to:
  - the skcipher cipher type (symmetric ciphers)
  - the hash cipher type (keyed message digests)
  - the AEAD cipher type
  - the RNG cipher type to provide the seed
- ALG_SET_AEAD_AUTHSIZE -- Setting the authentication tag size for AEAD ciphers. For a encryption operation, the authentication tag of the given size will be generated. For a decryption operation, the provided ciphertext is assumed to contain an authentication tag of the given size (see section about AEAD memory layout below).

User space API example

Please see [1] for libkcapi which provides an easy-to-use wrapper around the aforementioned Netlink kernel interface. [1] also contains a test application that invokes all libkcapi API calls.

[1] http://www.chronox.de/libkcapi.html

Chapter 5. Programming Interface

Please note that the kernel crypto API contains the AEAD givcrypt API (crypto_aead_giv* and aead_givcrypt_* function calls in include/crypto/aead.h). This API is obsolete and will be removed in the future. To obtain the functionality of an AEAD cipher with internal IV generation, use the IV generator as a regular cipher. For example, rfc4106(gcm(aes)) is the AEAD cipher with external IV generation and seqniv(rfc4106(gcm(aes))) implies that the kernel crypto API generates the IV. Different IV generators are available.

Block Cipher Context Data Structures

These data structures define the operating context for each block cipher type.

Name

struct aead_request — AEAD request

Synopsis

```
struct aead_request {
  struct crypto_async_request base;
  unsigned int assoclen;
  unsigned int cryptlen;
  u8 * iv;
  struct scatterlist * src;
  struct scatterlist * dst;
  void * __ctx[];
};
```

Members

base	Common attributes for async crypto requests
assoclen	Length in bytes of associated data for authentication
cryptlen	Length of data to be encrypted or decrypted
iv	Initialisation vector
src	Source data
dst	Destination data
__ctx[]	Start of private context data

Block Cipher Algorithm Definitions

These data structures define modular crypto algorithm implementations, managed via `crypto_register_alg` and `crypto_unregister_alg`.

Name

struct crypto_alg — definition of a cryptograpic cipher algorithm

Synopsis

```
struct crypto_alg {
  struct list_head cra_list;
  struct list_head cra_users;
  u32 cra_flags;
  unsigned int cra_blocksize;
  unsigned int cra_ctxsize;
  unsigned int cra_alignmask;
  int cra_priority;
  atomic_t cra_refcnt;
  char cra_name[CRYPTO_MAX_ALG_NAME];
  char cra_driver_name[CRYPTO_MAX_ALG_NAME];
  const struct crypto_type * cra_type;
  union cra_u;
  int (* cra_init) (struct crypto_tfm *tfm);
  void (* cra_exit) (struct crypto_tfm *tfm);
  void (* cra_destroy) (struct crypto_alg *alg);
  struct module * cra_module;
};
```

Members

cra_list — internally used

cra_users — internally used

cra_flags — Flags describing this transformation. See include/linux/crypto.h CRYPTO_ALG_* flags for the flags which go in here. Those are used for fine-tuning the description of the transformation algorithm.

cra_blocksize — Minimum block size of this transformation. The size in bytes of the smallest possible unit which can be transformed with this algorithm. The users must respect this value. In case of HASH transformation, it is possible for a smaller block than *cra_blocksize* to be passed to the crypto API for transformation, in case of any other transformation type, an error will be returned upon any attempt to transform smaller than *cra_blocksize* chunks.

cra_ctxsize — Size of the operational context of the transformation. This value informs the kernel crypto API about the memory size needed to be allocated for the transformation context.

cra_alignmask — Alignment mask for the input and output data buffer. The data buffer containing the input data for the algorithm must be aligned to this alignment mask. The data buffer for the output data must be aligned to this alignment mask. Note that the Crypto API will do the re-alignment in software, but only under special conditions and there is a performance hit. The re-alignment happens at these occasions for different

cra_priority	Priority of this transformation implementation. In case multiple transformations with same *cra_name* are available to the Crypto API, the kernel will use the one with highest *cra_priority*.
cra_refcnt	internally used
cra_name[CRYPTO_MAX_ALG_NAME]	Generic name (usable by multiple implementations) of the transformation algorithm. This is the name of the transformation itself. This field is used by the kernel when looking up the providers of particular transformation.
cra_driver_name[CRYPTO_MAX_ALG_NAME]	Unique name of the transformation provider. This is the name of the provider of the transformation. This can be any arbitrary value, but in the usual case, this contains the name of the chip or provider and the name of the transformation algorithm.
cra_type	Type of the cryptographic transformation. This is a pointer to struct crypto_type, which implements callbacks common for all transformation types. There are multiple options: crypto_blkcipher_type, crypto_ablkcipher_type, crypto_ahash_type, crypto_rng_type. This field might be empty. In that case, there are no common callbacks. This is the case for: cipher, compress, shash.
cra_u	Callbacks implementing the transformation. This is a union of multiple structures. Depending on the type of transformation selected by *cra_type* and *cra_flags* above, the associated structure must be filled with callbacks. This field might be empty. This is the case for ahash, shash.
cra_init	Initialize the cryptographic transformation object. This function is used to initialize the cryptographic transformation object. This function is called only once at the instantiation time, right after the transformation context was allocated. In case the cryptographic hardware has some special requirements which need to be handled by software, this function shall check for the precise requirement of the transformation and put any software fallbacks in place.
cra_exit	Deinitialize the cryptographic transformation object. This is a counterpart to *cra_init*, used to remove various changes set in *cra_init*.
cra_destroy	internally used
cra_module	Owner of this transformation implementation. Set to THIS_MODULE

Description

The struct crypto_alg describes a generic Crypto API algorithm and is common for all of the transformations. Any variable not documented here shall not be used by a cipher implementation as it is internal to the Crypto API.

Name

struct ablkcipher_alg — asynchronous block cipher definition

Synopsis

```
struct ablkcipher_alg {
  int (* setkey) (struct crypto_ablkcipher *tfm, const u8 *key,unsigned int keylen
  int (* encrypt) (struct ablkcipher_request *req);
  int (* decrypt) (struct ablkcipher_request *req);
  int (* givencrypt) (struct skcipher_givcrypt_request *req);
  int (* givdecrypt) (struct skcipher_givcrypt_request *req);
  const char * geniv;
  unsigned int min_keysize;
  unsigned int max_keysize;
  unsigned int ivsize;
};
```

Members

setkey — Set key for the transformation. This function is used to either program a supplied key into the hardware or store the key in the transformation context for programming it later. Note that this function does modify the transformation context. This function can be called multiple times during the existence of the transformation object, so one must make sure the key is properly reprogrammed into the hardware. This function is also responsible for checking the key length for validity. In case a software fallback was put in place in the *cra_init* call, this function might need to use the fallback if the algorithm doesn't support all of the key sizes.

encrypt — Encrypt a scatterlist of blocks. This function is used to encrypt the supplied scatterlist containing the blocks of data. The crypto API consumer is responsible for aligning the entries of the scatterlist properly and making sure the chunks are correctly sized. In case a software fallback was put in place in the *cra_init* call, this function might need to use the fallback if the algorithm doesn't support all of the key sizes. In case the key was stored in transformation context, the key might need to be re-programmed into the hardware in this function. This function shall not modify the transformation context, as this function may be called in parallel with the same transformation object.

decrypt — Decrypt a single block. This is a reverse counterpart to *encrypt* and the conditions are exactly the same.

givencrypt — Update the IV for encryption. With this function, a cipher implementation may provide the function on how to update the IV for encryption.

givdecrypt — Update the IV for decryption. This is the reverse of *givencrypt* .

geniv — The transformation implementation may use an “IV generator” provided by the kernel crypto API. Several use cases have a predefined approach how IVs are to be updated. For such use cases, the kernel crypto API provides ready-to-use implementations that can be referenced with this variable.

min_keysize — Minimum key size supported by the transformation. This is the smallest key length supported by this transformation algorithm. This must be set to one of the pre-defined values as this is not hardware specific. Possible values for this field can be found via git grep “_MIN_KEY_SIZE” include/crypto/

max_keysize	Maximum key size supported by the transformation. This is the largest key length supported by this transformation algorithm. This must be set to one of the pre-defined values as this is not hardware specific. Possible values for this field can be found via git grep “_MAX_KEY_SIZE” include/crypto/
ivsize	IV size applicable for transformation. The consumer must provide an IV of exactly that size to perform the encrypt or decrypt operation.

Description

All fields except *`givencrypt`* , *`givdecrypt`* , *`geniv`* and *`ivsize`* are mandatory and must be filled.

Name

struct aead_alg — AEAD cipher definition

Synopsis

```
struct aead_alg {
  int (* setkey) (struct crypto_aead *tfm, const u8 *key,unsigned int keylen);
  int (* setauthsize) (struct crypto_aead *tfm, unsigned int authsize);
  int (* encrypt) (struct aead_request *req);
  int (* decrypt) (struct aead_request *req);
  int (* init) (struct crypto_aead *tfm);
  void (* exit) (struct crypto_aead *tfm);
  const char * geniv;
  unsigned int ivsize;
  unsigned int maxauthsize;
};
```

Members

setkey — see struct ablkcipher_alg

setauthsize — Set authentication size for the AEAD transformation. This function is used to specify the consumer requested size of the authentication tag to be either generated by the transformation during encryption or the size of the authentication tag to be supplied during the decryption operation. This function is also responsible for checking the authentication tag size for validity.

encrypt — see struct ablkcipher_alg

decrypt — see struct ablkcipher_alg

init — Initialize the cryptographic transformation object. This function is used to initialize the cryptographic transformation object. This function is called only once at the instantiation time, right after the transformation context was allocated. In case the cryptographic hardware has some special requirements which need to be handled by software, this function shall check for the precise requirement of the transformation and put any software fallbacks in place.

exit — Deinitialize the cryptographic transformation object. This is a counterpart to *`init`*, used to remove various changes set in *`init`*.

geniv — see struct ablkcipher_alg

ivsize — see struct ablkcipher_alg

maxauthsize — Set the maximum authentication tag size supported by the transformation. A transformation may support smaller tag sizes. As the authentication tag is a message digest to ensure the integrity of the encrypted data, a consumer typically wants the largest authentication tag possible as defined by this variable.

Description

All fields except *`ivsize`* is mandatory and must be filled.

Name

struct blkcipher_alg — synchronous block cipher definition

Synopsis

```
struct blkcipher_alg {
  int (* setkey) (struct crypto_tfm *tfm, const u8 *key,unsigned int keylen);
  int (* encrypt) (struct blkcipher_desc *desc,struct scatterlist *dst, struct sca
  int (* decrypt) (struct blkcipher_desc *desc,struct scatterlist *dst, struct sca
  const char * geniv;
  unsigned int min_keysize;
  unsigned int max_keysize;
  unsigned int ivsize;
};
```

Members

setkey	see struct ablkcipher_alg
encrypt	see struct ablkcipher_alg
decrypt	see struct ablkcipher_alg
geniv	see struct ablkcipher_alg
min_keysize	see struct ablkcipher_alg
max_keysize	see struct ablkcipher_alg
ivsize	see struct ablkcipher_alg

Description

All fields except *`geniv`* and *`ivsize`* are mandatory and must be filled.

Name

struct cipher_alg — single-block symmetric ciphers definition

Synopsis

```
struct cipher_alg {
  unsigned int cia_min_keysize;
  unsigned int cia_max_keysize;
  int (* cia_setkey) (struct crypto_tfm *tfm, const u8 *key,unsigned int keylen);
  void (* cia_encrypt) (struct crypto_tfm *tfm, u8 *dst, const u8 *src);
  void (* cia_decrypt) (struct crypto_tfm *tfm, u8 *dst, const u8 *src);
};
```

Members

cia_min_keysize — Minimum key size supported by the transformation. This is the smallest key length supported by this transformation algorithm. This must be set to one of the pre-defined values as this is not hardware specific. Possible values for this field can be found via git grep "_MIN_KEY_SIZE" include/crypto/

cia_max_keysize — Maximum key size supported by the transformation. This is the largest key length supported by this transformation algorithm. This must be set to one of the pre-defined values as this is not hardware specific. Possible values for this field can be found via git grep "_MAX_KEY_SIZE" include/crypto/

cia_setkey — Set key for the transformation. This function is used to either program a supplied key into the hardware or store the key in the transformation context for programming it later. Note that this function does modify the transformation context. This function can be called multiple times during the existence of the transformation object, so one must make sure the key is properly reprogrammed into the hardware. This function is also responsible for checking the key length for validity.

cia_encrypt — Encrypt a single block. This function is used to encrypt a single block of data, which must be *cra_blocksize* big. This always operates on a full *cra_blocksize* and it is not possible to encrypt a block of smaller size. The supplied buffers must therefore also be at least of *cra_blocksize* size. Both the input and output buffers are always aligned to *cra_alignmask*. In case either of the input or output buffer supplied by user of the crypto API is not aligned to *cra_alignmask*, the crypto API will re-align the buffers. The re-alignment means that a new buffer will be allocated, the data will be copied into the new buffer, then the processing will happen on the new buffer, then the data will be copied back into the original buffer and finally the new buffer will be freed. In case a software fallback was put in place in the *cra_init* call, this function might need to use the fallback if the algorithm doesn't support all of the key sizes. In case the key was stored in transformation context, the key might need to be re-programmed into the hardware in this function. This function shall not modify the transformation context, as this function may be called in parallel with the same transformation object.

cia_decrypt — Decrypt a single block. This is a reverse counterpart to *cia_encrypt*, and the conditions are exactly the same.

Description

All fields are mandatory and must be filled.

Name

struct rng_alg — random number generator definition

Synopsis

```
struct rng_alg {
  int (* generate) (struct crypto_rng *tfm,const u8 *src, unsigned int slen,u8 *ds
  int (* seed) (struct crypto_rng *tfm, const u8 *seed, unsigned int slen);
  void (* set_ent) (struct crypto_rng *tfm, const u8 *data,unsigned int len);
  unsigned int seedsize;
  struct crypto_alg base;
};
```

Members

generate — The function defined by this variable obtains a random number. The random number generator transform must generate the random number out of the context provided with this call, plus any additional data if provided to the call.

seed — Seed or reseed the random number generator. With the invocation of this function call, the random number generator shall become ready for generation. If the random number generator requires a seed for setting up a new state, the seed must be provided by the consumer while invoking this function. The required size of the seed is defined with *seedsize* .

set_ent — Set entropy that would otherwise be obtained from entropy source. Internal use only.

seedsize — The seed size required for a random number generator initialization defined with this variable. Some random number generators does not require a seed as the seeding is implemented internally without the need of support by the consumer. In this case, the seed size is set to zero.

base — Common crypto API algorithm data structure.

Asynchronous Block Cipher API

Asynchronous block cipher API is used with the ciphers of type CRYPTO_ALG_TYPE_ABLKCIPHER (listed as type “ablkcipher” in /proc/crypto).

Asynchronous cipher operations imply that the function invocation for a cipher request returns immediately before the completion of the operation. The cipher request is scheduled as a separate kernel thread and therefore load-balanced on the different CPUs via the process scheduler. To allow the kernel crypto API to inform the caller about the completion of a cipher request, the caller must provide a callback function. That function is invoked with the cipher handle when the request completes.

To support the asynchronous operation, additional information than just the cipher handle must be supplied to the kernel crypto API. That additional information is given by filling in the ablkcipher_request data structure.

For the asynchronous block cipher API, the state is maintained with the tfm cipher handle. A single tfm can be used across multiple calls and in parallel. For asynchronous block cipher calls, context data supplied and only used by the caller can be referenced the request data structure in addition to the IV used for the cipher request. The maintenance of such state information would be important for a crypto driver implementer to

have, because when calling the callback function upon completion of the cipher operation, that callback function may need some information about which operation just finished if it invoked multiple in parallel. This state information is unused by the kernel crypto API.

Name

crypto_alloc_ablkcipher — allocate asynchronous block cipher handle

Synopsis

```
struct crypto_ablkcipher * crypto_alloc_ablkcipher (const char *
alg_name, u32 type, u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the ablkcipher cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Description

Allocate a cipher handle for an ablkcipher. The returned struct crypto_ablkcipher is the cipher handle that is required for any subsequent API invocation for that ablkcipher.

Return

allocated cipher handle in case of success; IS_ERR is true in case of an error, PTR_ERR returns the error code.

Name

crypto_free_ablkcipher — zeroize and free cipher handle

Synopsis

```
void crypto_free_ablkcipher (struct crypto_ablkcipher * tfm);
```

Arguments

tfm cipher handle to be freed

Name

crypto_has_ablkcipher — Search for the availability of an ablkcipher.

Synopsis

```
int crypto_has_ablkcipher (const char * alg_name, u32 type, u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the ablkcipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Return

true when the ablkcipher is known to the kernel crypto API; false otherwise

Name

crypto_ablkcipher_ivsize — obtain IV size

Synopsis

```
unsigned int crypto_ablkcipher_ivsize (struct crypto_ablkcipher * tfm);
```

Arguments

tfm cipher handle

Description

The size of the IV for the ablkcipher referenced by the cipher handle is returned. This IV size may be zero if the cipher does not need an IV.

Return

IV size in bytes

Name

crypto_ablkcipher_blocksize — obtain block size of cipher

Synopsis

```
unsigned int crypto_ablkcipher_blocksize (struct crypto_ablkcipher *
tfm);
```

Arguments

tfm cipher handle

Description

The block size for the ablkcipher referenced with the cipher handle is returned. The caller may use that information to allocate appropriate memory for the data returned by the encryption or decryption operation

Return

block size of cipher

Name

crypto_ablkcipher_setkey — set key for cipher

Synopsis

```
int crypto_ablkcipher_setkey (struct crypto_ablkcipher * tfm, const u8
* key, unsigned int keylen);
```

Arguments

tfm cipher handle

key buffer holding the key

keylen length of the key in bytes

Description

The caller provided key is set for the ablkcipher referenced by the cipher handle.

Note, the key length determines the cipher type. Many block ciphers implement different cipher modes depending on the key size, such as AES-128 vs AES-192 vs. AES-256. When providing a 16 byte key for an AES cipher handle, AES-128 is performed.

Return

0 if the setting of the key was successful; < 0 if an error occurred

Name

crypto_ablkcipher_reqtfm — obtain cipher handle from request

Synopsis

```
struct crypto_ablkcipher * crypto_ablkcipher_reqtfm (struct
ablkcipher_request * req);
```

Arguments

req ablkcipher_request out of which the cipher handle is to be obtained

Description

Return the crypto_ablkcipher handle when furnishing an ablkcipher_request data structure.

Return

crypto_ablkcipher handle

Name

crypto_ablkcipher_encrypt — encrypt plaintext

Synopsis

```
int crypto_ablkcipher_encrypt (struct ablkcipher_request * req);
```

Arguments

req reference to the ablkcipher_request handle that holds all information needed to perform the cipher operation

Description

Encrypt plaintext data using the ablkcipher_request handle. That data structure and how it is filled with data is discussed with the ablkcipher_request_* functions.

Return

0 if the cipher operation was successful; < 0 if an error occurred

Name

crypto_ablkcipher_decrypt — decrypt ciphertext

Synopsis

```
int crypto_ablkcipher_decrypt (struct ablkcipher_request * req);
```

Arguments

req reference to the ablkcipher_request handle that holds all information needed to perform the cipher operation

Description

Decrypt ciphertext data using the ablkcipher_request handle. That data structure and how it is filled with data is discussed with the ablkcipher_request_* functions.

Return

0 if the cipher operation was successful; < 0 if an error occurred

Asynchronous Cipher Request Handle

The ablkcipher_request data structure contains all pointers to data required for the asynchronous cipher operation. This includes the cipher handle (which can be used by multiple ablkcipher_request instances), pointer to plaintext and ciphertext, asynchronous callback function, etc. It acts as a handle to the ablkcipher_request_* API calls in a similar way as ablkcipher handle to the crypto_ablkcipher_* API calls.

Name

crypto_ablkcipher_reqsize — obtain size of the request data structure

Synopsis

```
unsigned int crypto_ablkcipher_reqsize (struct crypto_ablkcipher * tfm);
```

Arguments

tfm cipher handle

Return

number of bytes

Name

ablkcipher_request_set_tfm — update cipher handle reference in request

Synopsis

```
void ablkcipher_request_set_tfm (struct ablkcipher_request * req, struct
crypto_ablkcipher * tfm);
```

Arguments

req request handle to be modified

tfm cipher handle that shall be added to the request handle

Description

Allow the caller to replace the existing ablkcipher handle in the request data structure with a different one.

Name

ablkcipher_request_alloc — allocate request data structure

Synopsis

```
struct ablkcipher_request * ablkcipher_request_alloc (struct
crypto_ablkcipher * tfm, gfp_t gfp);
```

Arguments

tfm cipher handle to be registered with the request

gfp memory allocation flag that is handed to kmalloc by the API call.

Description

Allocate the request data structure that must be used with the ablkcipher encrypt and decrypt API calls. During the allocation, the provided ablkcipher handle is registered in the request data structure.

Return

allocated request handle in case of success; IS_ERR is true in case of an error, PTR_ERR returns the error code.

Name

ablkcipher_request_free — zeroize and free request data structure

Synopsis

```
void ablkcipher_request_free (struct ablkcipher_request * req);
```

Arguments

req request data structure cipher handle to be freed

Name

ablkcipher_request_set_callback — set asynchronous callback function

Synopsis

```
void ablkcipher_request_set_callback (struct ablkcipher_request * req,
u32 flags, crypto_completion_t compl, void * data);
```

Arguments

req request handle

flags specify zero or an ORing of the flags CRYPTO_TFM_REQ_MAY_BACKLOG the request queue may back log and increase the wait queue beyond the initial maximum size; CRYPTO_TFM_REQ_MAY_SLEEP the request processing may sleep

compl callback function pointer to be registered with the request handle

data The data pointer refers to memory that is not used by the kernel crypto API, but provided to the callback function for it to use. Here, the caller can provide a reference to memory the callback function can operate on. As the callback function is invoked asynchronously to the related functionality, it may need to access data structures of the related functionality which can be referenced using this pointer. The callback function can access the memory via the “data” field in the crypto_async_request data structure provided to the callback function.

Description

This function allows setting the callback function that is triggered once the cipher operation completes.

The callback function is registered with the ablkcipher_request handle and must comply with the following template

void callback_function(struct crypto_async_request *req, int error)

Name

ablkcipher_request_set_crypt — set data buffers

Synopsis

```
void ablkcipher_request_set_crypt (struct ablkcipher_request * req,
struct scatterlist * src, struct scatterlist * dst, unsigned int nbytes,
void * iv);
```

Arguments

req	request handle
src	source scatter / gather list
dst	destination scatter / gather list
nbytes	number of bytes to process from *src*
iv	IV for the cipher operation which must comply with the IV size defined by crypto_ablkcipher_ivsize

Description

This function allows setting of the source data and destination data scatter / gather lists.

For encryption, the source is treated as the plaintext and the destination is the ciphertext. For a decryption operation, the use is reversed - the source is the ciphertext and the destination is the plaintext.

Authenticated Encryption With Associated Data (AEAD) Cipher API

The AEAD cipher API is used with the ciphers of type CRYPTO_ALG_TYPE_AEAD (listed as type “aead” in /proc/crypto)

The most prominent examples for this type of encryption is GCM and CCM. However, the kernel supports other types of AEAD ciphers which are defined with the following cipher string:

authenc(keyed message digest, block cipher)

For example: authenc(hmac(sha256), cbc(aes))

The example code provided for the asynchronous block cipher operation applies here as well. Naturally all *ablkcipher* symbols must be exchanged the *aead* pendants discussed in the following. In addition, for the AEAD operation, the aead_request_set_assoc function must be used to set the pointer to the associated data memory location before performing the encryption or decryption operation. In case of an encryption, the associated data memory is filled during the encryption operation. For decryption, the associated data memory must contain data that is used to verify the integrity of the decrypted data. Another deviation from the asynchronous block cipher operation is that the caller should explicitly check for -EBADMSG of the crypto_aead_decrypt. That error indicates an authentication error, i.e. a breach in the integrity of the message. In essence, that -EBADMSG error code is the key bonus an AEAD cipher has over “standard” block chaining modes.

Memory Structure:

To support the needs of the most prominent user of AEAD ciphers, namely IPSEC, the AEAD ciphers have a special memory layout the caller must adhere to.

The scatter list pointing to the input data must contain:

* for RFC4106 ciphers, the concatenation of associated authentication data || IV || plaintext or ciphertext. Note, the same IV (buffer) is also set with the aead_request_set_crypt call. Note, the API call of aead_request_set_ad must provide the length of the AAD and the IV. The API call of aead_request_set_crypt only points to the size of the input plaintext or ciphertext.

* for “normal” AEAD ciphers, the concatenation of associated authentication data || plaintext or ciphertext.

It is important to note that if multiple scatter gather list entries form the input data mentioned above, the first entry must not point to a NULL buffer. If there is any potential where the AAD buffer can be NULL, the calling code must contain a precaution to ensure that this does not result in the first scatter gather list entry pointing to a NULL buffer.

Name

crypto_alloc_aead — allocate AEAD cipher handle

Synopsis

```
struct crypto_aead * crypto_alloc_aead (const char * alg_name, u32 type,
u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the AEAD cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Description

Allocate a cipher handle for an AEAD. The returned struct crypto_aead is the cipher handle that is required for any subsequent API invocation for that AEAD.

Return

allocated cipher handle in case of success; IS_ERR is true in case of an error, PTR_ERR returns the error code.

Name

crypto_free_aead — zeroize and free aead handle

Synopsis

```
void crypto_free_aead (struct crypto_aead * tfm);
```

Arguments

tfm cipher handle to be freed

Name

crypto_aead_ivsize — obtain IV size

Synopsis

```
unsigned int crypto_aead_ivsize (struct crypto_aead * tfm);
```

Arguments

tfm cipher handle

Description

The size of the IV for the aead referenced by the cipher handle is returned. This IV size may be zero if the cipher does not need an IV.

Return

IV size in bytes

Name

crypto_aead_authsize — obtain maximum authentication data size

Synopsis

```
unsigned int crypto_aead_authsize (struct crypto_aead * tfm);
```

Arguments

tfm cipher handle

Description

The maximum size of the authentication data for the AEAD cipher referenced by the AEAD cipher handle is returned. The authentication data size may be zero if the cipher implements a hard-coded maximum.

The authentication data may also be known as “tag value”.

Return

authentication data size / tag size in bytes

Name

crypto_aead_blocksize — obtain block size of cipher

Synopsis

```
unsigned int crypto_aead_blocksize (struct crypto_aead * tfm);
```

Arguments

tfm cipher handle

Description

The block size for the AEAD referenced with the cipher handle is returned. The caller may use that information to allocate appropriate memory for the data returned by the encryption or decryption operation

Return

block size of cipher

Name

crypto_aead_setkey — set key for cipher

Synopsis

```
int crypto_aead_setkey (struct crypto_aead * tfm, const u8 * key, unsigned int keylen);
```

Arguments

tfm cipher handle

key buffer holding the key

keylen length of the key in bytes

Description

The caller provided key is set for the AEAD referenced by the cipher handle.

Note, the key length determines the cipher type. Many block ciphers implement different cipher modes depending on the key size, such as AES-128 vs AES-192 vs. AES-256. When providing a 16 byte key for an AES cipher handle, AES-128 is performed.

Return

0 if the setting of the key was successful; < 0 if an error occurred

Name

crypto_aead_setauthsize — set authentication data size

Synopsis

```
int crypto_aead_setauthsize (struct crypto_aead * tfm, unsigned int authsize);
```

Arguments

tfm cipher handle

authsize size of the authentication data / tag in bytes

Description

Set the authentication data size / tag size. AEAD requires an authentication tag (or MAC) in addition to the associated data.

Return

0 if the setting of the key was successful; < 0 if an error occurred

Name

crypto_aead_encrypt — encrypt plaintext

Synopsis

```
int crypto_aead_encrypt (struct aead_request * req);
```

Arguments

req reference to the aead_request handle that holds all information needed to perform the cipher operation

Description

Encrypt plaintext data using the aead_request handle. That data structure and how it is filled with data is discussed with the aead_request_* functions.

IMPORTANT NOTE The encryption operation creates the authentication data / tag. That data is concatenated with the created ciphertext. The ciphertext memory size is therefore the given number of block cipher blocks + the size defined by the crypto_aead_setauthsize invocation. The caller must ensure that sufficient memory is available for the ciphertext and the authentication tag.

Return

0 if the cipher operation was successful; < 0 if an error occurred

Name

crypto_aead_decrypt — decrypt ciphertext

Synopsis

```
int crypto_aead_decrypt (struct aead_request * req);
```

Arguments

req reference to the ablkcipher_request handle that holds all information needed to perform the cipher operation

Description

Decrypt ciphertext data using the aead_request handle. That data structure and how it is filled with data is discussed with the aead_request_* functions.

IMPORTANT NOTE The caller must concatenate the ciphertext followed by the authentication data / tag. That authentication data / tag must have the size defined by the crypto_aead_setauthsize invocation.

Return

0 if the cipher operation was successful; -EBADMSG: The AEAD cipher operation performs the authentication of the data during the decryption operation. Therefore, the function returns this error if the authentication of the ciphertext was unsuccessful (i.e. the integrity of the ciphertext or the associated data was violated); < 0 if an error occurred.

Asynchronous AEAD Request Handle

The aead_request data structure contains all pointers to data required for the AEAD cipher operation. This includes the cipher handle (which can be used by multiple aead_request instances), pointer to plaintext and ciphertext, asynchronous callback function, etc. It acts as a handle to the aead_request_* API calls in a similar way as AEAD handle to the crypto_aead_* API calls.

Name

crypto_aead_reqsize — obtain size of the request data structure

Synopsis

```
unsigned int crypto_aead_reqsize (struct crypto_aead * tfm);
```

Arguments

tfm cipher handle

Return

number of bytes

Name

aead_request_set_tfm — update cipher handle reference in request

Synopsis

```
void aead_request_set_tfm (struct aead_request * req, struct crypto_aead
* tfm);
```

Arguments

req request handle to be modified

tfm cipher handle that shall be added to the request handle

Description

Allow the caller to replace the existing aead handle in the request data structure with a different one.

Name

aead_request_alloc — allocate request data structure

Synopsis

```
struct aead_request * aead_request_alloc (struct crypto_aead * tfm,
gfp_t gfp);
```

Arguments

tfm cipher handle to be registered with the request

gfp memory allocation flag that is handed to kmalloc by the API call.

Description

Allocate the request data structure that must be used with the AEAD encrypt and decrypt API calls. During the allocation, the provided aead handle is registered in the request data structure.

Return

allocated request handle in case of success; IS_ERR is true in case of an error, PTR_ERR returns the error code.

Name

aead_request_free — zeroize and free request data structure

Synopsis

```
void aead_request_free (struct aead_request * req);
```

Arguments

req request data structure cipher handle to be freed

Name

aead_request_set_callback — set asynchronous callback function

Synopsis

```
void aead_request_set_callback (struct aead_request * req, u32 flags,
crypto_completion_t compl, void * data);
```

Arguments

req request handle

flags specify zero or an ORing of the flags CRYPTO_TFM_REQ_MAY_BACKLOG the request queue may back log and increase the wait queue beyond the initial maximum size; CRYPTO_TFM_REQ_MAY_SLEEP the request processing may sleep

compl callback function pointer to be registered with the request handle

data The data pointer refers to memory that is not used by the kernel crypto API, but provided to the callback function for it to use. Here, the caller can provide a reference to memory the callback function can operate on. As the callback function is invoked asynchronously to the related functionality, it may need to access data structures of the related functionality which can be referenced using this pointer. The callback function can access the memory via the “data” field in the crypto_async_request data structure provided to the callback function.

Description

Setting the callback function that is triggered once the cipher operation completes

The callback function is registered with the aead_request handle and must comply with the following template

void callback_function(struct crypto_async_request *req, int error)

Name

aead_request_set_crypt — set data buffers

Synopsis

```
void aead_request_set_crypt (struct aead_request * req, struct scat-
terlist * src, struct scatterlist * dst, unsigned int cryptlen, u8 * iv);
```

Arguments

req	request handle
src	source scatter / gather list
dst	destination scatter / gather list
cryptlen	number of bytes to process from *src*
iv	IV for the cipher operation which must comply with the IV size defined by `crypto_aead_ivsize`

Description

Setting the source data and destination data scatter / gather lists which hold the associated data concatenated with the plaintext or ciphertext. See below for the authentication tag.

For encryption, the source is treated as the plaintext and the destination is the ciphertext. For a decryption operation, the use is reversed - the source is the ciphertext and the destination is the plaintext.

For both src/dst the layout is associated data, plain/cipher text, authentication tag.

The content of the AD in the destination buffer after processing will either be untouched, or it will contain a copy of the AD from the source buffer. In order to ensure that it always has a copy of the AD, the user must copy the AD over either before or after processing. Of course this is not relevant if the user is doing in-place processing where src == dst.

IMPORTANT NOTE AEAD requires an authentication tag (MAC). For decryption, the caller must concatenate the ciphertext followed by the authentication tag and provide the entire data stream to the decryption operation (i.e. the data length used for the initialization of the scatterlist and the data length for the decryption operation is identical). For encryption, however, the authentication tag is created while encrypting the data. The destination buffer must hold sufficient space for the ciphertext and the authentication tag while the encryption invocation must only point to the plaintext data size. The following code snippet illustrates the memory usage buffer = kmalloc(ptbuflen + (enc ? authsize : 0)); sg_init_one(sg, buffer, ptbuflen + (enc ? authsize : 0)); aead_request_set_crypt(req, sg, sg, ptbuflen, iv);

Name

include/crypto/aead.h — Document generation inconsistency

Oops

Warning

The template for this document tried to insert the structured comment from the file `include/crypto/aead.h` at this point, but none was found. This dummy section is inserted to allow generation to continue.

Name

aead_request_set_ad — set associated data information

Synopsis

```
void aead_request_set_ad (struct aead_request * req, unsigned int assoclen);
```

Arguments

req request handle

assoclen number of bytes in associated data

Description

Setting the AD information. This function sets the length of the associated data.

Synchronous Block Cipher API

The synchronous block cipher API is used with the ciphers of type CRYPTO_ALG_TYPE_BLKCIPHER (listed as type “blkcipher” in /proc/crypto)

Synchronous calls, have a context in the tfm. But since a single tfm can be used in multiple calls and in parallel, this info should not be changeable (unless a lock is used). This applies, for example, to the symmetric key. However, the IV is changeable, so there is an iv field in blkcipher_tfm structure for synchronous blkcipher api. So, its the only state info that can be kept for synchronous calls without using a big lock across a tfm.

The block cipher API allows the use of a complete cipher, i.e. a cipher consisting of a template (a block chaining mode) and a single block cipher primitive (e.g. AES).

The plaintext data buffer and the ciphertext data buffer are pointed to by using scatter/gather lists. The cipher operation is performed on all segments of the provided scatter/gather lists.

The kernel crypto API supports a cipher operation “in-place” which means that the caller may provide the same scatter/gather list for the plaintext and cipher text. After the completion of the cipher operation, the plaintext data is replaced with the ciphertext data in case of an encryption and vice versa for a decryption. The caller must ensure that the scatter/gather lists for the output data point to sufficiently large buffers, i.e. multiples of the block size of the cipher.

Name

crypto_alloc_blkcipher — allocate synchronous block cipher handle

Synopsis

```
struct crypto_blkcipher * crypto_alloc_blkcipher (const char * alg_name,
u32 type, u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the blkcipher cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Description

Allocate a cipher handle for a block cipher. The returned struct crypto_blkcipher is the cipher handle that is required for any subsequent API invocation for that block cipher.

Return

allocated cipher handle in case of success; IS_ERR is true in case of an error, PTR_ERR returns the error code.

Name

crypto_free_blkcipher — zeroize and free the block cipher handle

Synopsis

```
void crypto_free_blkcipher (struct crypto_blkcipher * tfm);
```

Arguments

tfm cipher handle to be freed

Name

crypto_has_blkcipher — Search for the availability of a block cipher

Synopsis

```
int crypto_has_blkcipher (const char * alg_name, u32 type, u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the block cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Return

true when the block cipher is known to the kernel crypto API; false otherwise

Name

crypto_blkcipher_name — return the name / cra_name from the cipher handle

Synopsis

```
const char * crypto_blkcipher_name (struct crypto_blkcipher * tfm);
```

Arguments

tfm cipher handle

Return

The character string holding the name of the cipher

Name

crypto_blkcipher_ivsize — obtain IV size

Synopsis

```
unsigned int crypto_blkcipher_ivsize (struct crypto_blkcipher * tfm);
```

Arguments

tfm cipher handle

Description

The size of the IV for the block cipher referenced by the cipher handle is returned. This IV size may be zero if the cipher does not need an IV.

Return

IV size in bytes

Name

crypto_blkcipher_blocksize — obtain block size of cipher

Synopsis

```
unsigned int crypto_blkcipher_blocksize (struct crypto_blkcipher * tfm);
```

Arguments

tfm cipher handle

Description

The block size for the block cipher referenced with the cipher handle is returned. The caller may use that information to allocate appropriate memory for the data returned by the encryption or decryption operation.

Return

block size of cipher

Name

crypto_blkcipher_setkey — set key for cipher

Synopsis

```
int crypto_blkcipher_setkey (struct crypto_blkcipher * tfm, const u8 *
key, unsigned int keylen);
```

Arguments

tfm cipher handle

key buffer holding the key

keylen length of the key in bytes

Description

The caller provided key is set for the block cipher referenced by the cipher handle.

Note, the key length determines the cipher type. Many block ciphers implement different cipher modes depending on the key size, such as AES-128 vs AES-192 vs. AES-256. When providing a 16 byte key for an AES cipher handle, AES-128 is performed.

Return

0 if the setting of the key was successful; < 0 if an error occurred

Name

crypto_blkcipher_encrypt — encrypt plaintext

Synopsis

```
int crypto_blkcipher_encrypt (struct blkcipher_desc * desc, struct scat-
terlist * dst, struct scatterlist * src, unsigned int nbytes);
```

Arguments

desc reference to the block cipher handle with meta data

dst scatter/gather list that is filled by the cipher operation with the ciphertext

src scatter/gather list that holds the plaintext

nbytes number of bytes of the plaintext to encrypt.

Description

Encrypt plaintext data using the IV set by the caller with a preceding call of crypto_blkcipher_set_iv.

The blkcipher_desc data structure must be filled by the caller and can reside on the stack. The caller must fill desc as follows: desc.tfm is filled with the block cipher handle; desc.flags is filled with either CRYPTO_TFM_REQ_MAY_SLEEP or 0.

Return

0 if the cipher operation was successful; < 0 if an error occurred

Name

crypto_blkcipher_encrypt_iv — encrypt plaintext with dedicated IV

Synopsis

```
int crypto_blkcipher_encrypt_iv (struct blkcipher_desc * desc, struct
scatterlist * dst, struct scatterlist * src, unsigned int nbytes);
```

Arguments

desc reference to the block cipher handle with meta data

dst scatter/gather list that is filled by the cipher operation with the ciphertext

src scatter/gather list that holds the plaintext

nbytes number of bytes of the plaintext to encrypt.

Description

Encrypt plaintext data with the use of an IV that is solely used for this cipher operation. Any previously set IV is not used.

The blkcipher_desc data structure must be filled by the caller and can reside on the stack. The caller must fill desc as follows: desc.tfm is filled with the block cipher handle; desc.info is filled with the IV to be used for the current operation; desc.flags is filled with either CRYPTO_TFM_REQ_MAY_SLEEP or 0.

Return

0 if the cipher operation was successful; < 0 if an error occurred

Name

crypto_blkcipher_decrypt — decrypt ciphertext

Synopsis

```
int crypto_blkcipher_decrypt (struct blkcipher_desc * desc, struct scat-
terlist * dst, struct scatterlist * src, unsigned int nbytes);
```

Arguments

desc reference to the block cipher handle with meta data

dst scatter/gather list that is filled by the cipher operation with the plaintext

src scatter/gather list that holds the ciphertext

nbytes number of bytes of the ciphertext to decrypt.

Description

Decrypt ciphertext data using the IV set by the caller with a preceding call of crypto_blkcipher_set_iv.

The blkcipher_desc data structure must be filled by the caller as documented for the crypto_blkcipher_encrypt call above.

Return

0 if the cipher operation was successful; < 0 if an error occurred

Name

crypto_blkcipher_decrypt_iv — decrypt ciphertext with dedicated IV

Synopsis

```
int crypto_blkcipher_decrypt_iv (struct blkcipher_desc * desc, struct
scatterlist * dst, struct scatterlist * src, unsigned int nbytes);
```

Arguments

desc reference to the block cipher handle with meta data

dst scatter/gather list that is filled by the cipher operation with the plaintext

src scatter/gather list that holds the ciphertext

nbytes number of bytes of the ciphertext to decrypt.

Description

Decrypt ciphertext data with the use of an IV that is solely used for this cipher operation. Any previously set IV is not used.

The blkcipher_desc data structure must be filled by the caller as documented for the crypto_blkcipher_encrypt_iv call above.

Return

0 if the cipher operation was successful; < 0 if an error occurred

Name

crypto_blkcipher_set_iv — set IV for cipher

Synopsis

```
void crypto_blkcipher_set_iv (struct crypto_blkcipher * tfm, const u8
* src, unsigned int len);
```

Arguments

tfm cipher handle

src buffer holding the IV

len length of the IV in bytes

Description

The caller provided IV is set for the block cipher referenced by the cipher handle.

Name

crypto_blkcipher_get_iv — obtain IV from cipher

Synopsis

```
void crypto_blkcipher_get_iv (struct crypto_blkcipher * tfm, u8 * dst,
unsigned int len);
```

Arguments

tfm cipher handle

dst buffer filled with the IV

len length of the buffer dst

Description

The caller can obtain the IV set for the block cipher referenced by the cipher handle and store it into the user-provided buffer. If the buffer has an insufficient space, the IV is truncated to fit the buffer.

Single Block Cipher API

The single block cipher API is used with the ciphers of type CRYPTO_ALG_TYPE_CIPHER (listed as type “cipher” in /proc/crypto).

Using the single block cipher API calls, operations with the basic cipher primitive can be implemented. These cipher primitives exclude any block chaining operations including IV handling.

The purpose of this single block cipher API is to support the implementation of templates or other concepts that only need to perform the cipher operation on one block at a time. Templates invoke the underlying cipher primitive block-wise and process either the input or the output data of these cipher operations.

Name

crypto_alloc_cipher — allocate single block cipher handle

Synopsis

```
struct crypto_cipher * crypto_alloc_cipher (const char * alg_name, u32 type, u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the single block cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Description

Allocate a cipher handle for a single block cipher. The returned struct crypto_cipher is the cipher handle that is required for any subsequent API invocation for that single block cipher.

Return

allocated cipher handle in case of success; IS_ERR is true in case of an error, PTR_ERR returns the error code.

Name

crypto_free_cipher — zeroize and free the single block cipher handle

Synopsis

```
void crypto_free_cipher (struct crypto_cipher * tfm);
```

Arguments

tfm cipher handle to be freed

Name

crypto_has_cipher — Search for the availability of a single block cipher

Synopsis

```
int crypto_has_cipher (const char * alg_name, u32 type, u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the single block cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Return

true when the single block cipher is known to the kernel crypto API; false otherwise

Name

crypto_cipher_blocksize — obtain block size for cipher

Synopsis

```
unsigned int crypto_cipher_blocksize (struct crypto_cipher * tfm);
```

Arguments

tfm cipher handle

Description

The block size for the single block cipher referenced with the cipher handle tfm is returned. The caller may use that information to allocate appropriate memory for the data returned by the encryption or decryption operation

Return

block size of cipher

Name

crypto_cipher_setkey — set key for cipher

Synopsis

```
int crypto_cipher_setkey (struct crypto_cipher * tfm, const u8 * key,
unsigned int keylen);
```

Arguments

tfm cipher handle

key buffer holding the key

keylen length of the key in bytes

Description

The caller provided key is set for the single block cipher referenced by the cipher handle.

Note, the key length determines the cipher type. Many block ciphers implement different cipher modes depending on the key size, such as AES-128 vs AES-192 vs. AES-256. When providing a 16 byte key for an AES cipher handle, AES-128 is performed.

Return

0 if the setting of the key was successful; < 0 if an error occurred

Name

crypto_cipher_encrypt_one — encrypt one block of plaintext

Synopsis

```
void crypto_cipher_encrypt_one (struct crypto_cipher * tfm, u8 * dst,
const u8 * src);
```

Arguments

tfm cipher handle

dst points to the buffer that will be filled with the ciphertext

src buffer holding the plaintext to be encrypted

Description

Invoke the encryption operation of one block. The caller must ensure that the plaintext and ciphertext buffers are at least one block in size.

Name

crypto_cipher_decrypt_one — decrypt one block of ciphertext

Synopsis

```
void crypto_cipher_decrypt_one (struct crypto_cipher * tfm, u8 * dst,
const u8 * src);
```

Arguments

tfm cipher handle

dst points to the buffer that will be filled with the plaintext

src buffer holding the ciphertext to be decrypted

Description

Invoke the decryption operation of one block. The caller must ensure that the plaintext and ciphertext buffers are at least one block in size.

Synchronous Message Digest API

The synchronous message digest API is used with the ciphers of type CRYPTO_ALG_TYPE_HASH (listed as type “hash” in /proc/crypto)

Name

crypto_alloc_hash — allocate synchronous message digest handle

Synopsis

```
struct crypto_hash * crypto_alloc_hash (const char * alg_name, u32 type,
u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the message digest cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Description

Allocate a cipher handle for a message digest. The returned struct crypto_hash is the cipher handle that is required for any subsequent API invocation for that message digest.

Return

allocated cipher handle in case of success; IS_ERR is true in case of an error, PTR_ERR returns the error code.

Name

crypto_free_hash — zeroize and free message digest handle

Synopsis

```
void crypto_free_hash (struct crypto_hash * tfm);
```

Arguments

tfm cipher handle to be freed

Name

crypto_has_hash — Search for the availability of a message digest

Synopsis

```
int crypto_has_hash (const char * alg_name, u32 type, u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the message digest cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Return

true when the message digest cipher is known to the kernel crypto API; false otherwise

Name

crypto_hash_blocksize — obtain block size for message digest

Synopsis

```
unsigned int crypto_hash_blocksize (struct crypto_hash * tfm);
```

Arguments

tfm cipher handle

Description

The block size for the message digest cipher referenced with the cipher handle is returned.

Return

block size of cipher

Name

crypto_hash_digestsize — obtain message digest size

Synopsis

```
unsigned int crypto_hash_digestsize (struct crypto_hash * tfm);
```

Arguments

tfm cipher handle

Description

The size for the message digest created by the message digest cipher referenced with the cipher handle is returned.

Return

message digest size

Name

crypto_hash_init — (re)initialize message digest handle

Synopsis

```
int crypto_hash_init (struct hash_desc * desc);
```

Arguments

desc cipher request handle that to be filled by caller -- desc.tfm is filled with the hash cipher handle; desc.flags is filled with either CRYPTO_TFM_REQ_MAY_SLEEP or 0.

Description

The call (re-)initializes the message digest referenced by the hash cipher request handle. Any potentially existing state created by previous operations is discarded.

Return

0 if the message digest initialization was successful; < 0 if an error occurred

Name

crypto_hash_update — add data to message digest for processing

Synopsis

```
int crypto_hash_update (struct hash_desc * desc, struct scatterlist *
sg, unsigned int nbytes);
```

Arguments

desc cipher request handle

sg scatter / gather list pointing to the data to be added to the message digest

nbytes number of bytes to be processed from *sg*

Description

Updates the message digest state of the cipher handle pointed to by the hash cipher request handle with the input data pointed to by the scatter/gather list.

Return

0 if the message digest update was successful; < 0 if an error occurred

Name

crypto_hash_final — calculate message digest

Synopsis

```
int crypto_hash_final (struct hash_desc * desc, u8 * out);
```

Arguments

desc cipher request handle

out message digest output buffer -- The caller must ensure that the out buffer has a sufficient size (e.g. by using the crypto_hash_digestsize function).

Description

Finalize the message digest operation and create the message digest based on all data added to the cipher handle. The message digest is placed into the output buffer.

Return

0 if the message digest creation was successful; < 0 if an error occurred

Name

crypto_hash_digest — calculate message digest for a buffer

Synopsis

```
int crypto_hash_digest (struct hash_desc * desc, struct scatterlist *
sg, unsigned int nbytes, u8 * out);
```

Arguments

desc see `crypto_hash_final`

sg see `crypto_hash_update`

nbytes see `crypto_hash_update`

out see `crypto_hash_final`

Description

This function is a “short-hand” for the function calls of crypto_hash_init, crypto_hash_update and crypto_hash_final. The parameters have the same meaning as discussed for those separate three functions.

Return

0 if the message digest creation was successful; < 0 if an error occurred

Name

crypto_hash_setkey — set key for message digest

Synopsis

```
int crypto_hash_setkey (struct crypto_hash * hash, const u8 * key,
unsigned int keylen);
```

Arguments

hash cipher handle

key buffer holding the key

keylen length of the key in bytes

Description

The caller provided key is set for the message digest cipher. The cipher handle must point to a keyed hash in order for this function to succeed.

Return

0 if the setting of the key was successful; < 0 if an error occurred

Message Digest Algorithm Definitions

These data structures define modular message digest algorithm implementations, managed via `crypto_register_ahash`, `crypto_register_shash`, `crypto_unregister_ahash` and `crypto_unregister_shash`.

Name

struct hash_alg_common — define properties of message digest

Synopsis

```
struct hash_alg_common {
  unsigned int digestsize;
  unsigned int statesize;
  struct crypto_alg base;
};
```

Members

digestsize — Size of the result of the transformation. A buffer of this size must be available to the *`final`* and *`finup`* calls, so they can store the resulting hash into it. For various predefined sizes, search include/crypto/ using git grep _DIGEST_SIZE include/crypto.

statesize — Size of the block for partial state of the transformation. A buffer of this size must be passed to the *`export`* function as it will save the partial state of the transformation into it. On the other side, the *`import`* function will load the state from a buffer of this size as well.

base — Start of data structure of cipher algorithm. The common data structure of crypto_alg contains information common to all ciphers. The hash_alg_common data structure now adds the hash-specific information.

Name

struct ahash_alg — asynchronous message digest definition

Synopsis

```
struct ahash_alg {
  int (* init) (struct ahash_request *req);
  int (* update) (struct ahash_request *req);
  int (* final) (struct ahash_request *req);
  int (* finup) (struct ahash_request *req);
  int (* digest) (struct ahash_request *req);
  int (* export) (struct ahash_request *req, void *out);
  int (* import) (struct ahash_request *req, const void *in);
  int (* setkey) (struct crypto_ahash *tfm, const u8 *key,unsigned int keylen);
  struct hash_alg_common halg;
};
```

Members

init — Initialize the transformation context. Intended only to initialize the state of the HASH transformation at the beginning. This shall fill in the internal structures used during the entire duration of the whole transformation. No data processing happens at this point.

update — Push a chunk of data into the driver for transformation. This function actually pushes blocks of data from upper layers into the driver, which then passes those to the hardware as seen fit. This function must not finalize the HASH transformation by calculating the final message digest as this only adds more data into the transformation. This function shall not modify the transformation context, as this function may be called in parallel with the same transformation object. Data processing can happen synchronously [SHASH] or asynchronously [AHASH] at this point.

final — Retrieve result from the driver. This function finalizes the transformation and retrieves the resulting hash from the driver and pushes it back to upper layers. No data processing happens at this point.

finup — Combination of *update* and *final*. This function is effectively a combination of *update* and *final* calls issued in sequence. As some hardware cannot do *update* and *final* separately, this callback was added to allow such hardware to be used at least by IPsec. Data processing can happen synchronously [SHASH] or asynchronously [AHASH] at this point.

digest — Combination of *init* and *update* and *final*. This function effectively behaves as the entire chain of operations, *init*, *update* and *final* issued in sequence. Just like *finup*, this was added for hardware which cannot do even the *finup*, but can only do the whole transformation in one run. Data processing can happen synchronously [SHASH] or asynchronously [AHASH] at this point.

export — Export partial state of the transformation. This function dumps the entire state of the ongoing transformation into a provided block of data so it can be *import* 'ed back later on. This is useful in case you want to save partial result of the transformation after processing certain amount of data and reload this partial result multiple times later on for multiple re-use. No data processing happens at this point.

import Import partial state of the transformation. This function loads the entire state of the ongoing transformation from a provided block of data so the transformation can continue from this point onward. No data processing happens at this point.

setkey Set optional key used by the hashing algorithm. Intended to push optional key used by the hashing algorithm from upper layers into the driver. This function can store the key in the transformation context or can outright program it into the hardware. In the former case, one must be careful to program the key into the hardware at appropriate time and one must be careful that `.setkey` can be called multiple times during the existence of the transformation object. Not all hashing algorithms do implement this function as it is only needed for keyed message digests. SHAx/MDx/CRCx do NOT implement this function. HMAC(MDx)/HMAC(SHAx)/CMAC(AES) do implement this function. This function must be called before any other of the *`init`*, *`update`*, *`final`*, *`finup`*, *`digest`* is called. No data processing happens at this point.

halg see struct hash_alg_common

Name

struct shash_alg — synchronous message digest definition

Synopsis

```
struct shash_alg {
  int (* init) (struct shash_desc *desc);
  int (* update) (struct shash_desc *desc, const u8 *data,unsigned int len);
  int (* final) (struct shash_desc *desc, u8 *out);
  int (* finup) (struct shash_desc *desc, const u8 *data,unsigned int len, u8 *out
  int (* digest) (struct shash_desc *desc, const u8 *data,unsigned int len, u8 *ou
  int (* export) (struct shash_desc *desc, void *out);
  int (* import) (struct shash_desc *desc, const void *in);
  int (* setkey) (struct crypto_shash *tfm, const u8 *key,unsigned int keylen);
  unsigned int descsize;
  unsigned int digestsize;
  unsigned int statesize;
  struct crypto_alg base;
};
```

Members

init	see struct ahash_alg
update	see struct ahash_alg
final	see struct ahash_alg
finup	see struct ahash_alg
digest	see struct ahash_alg
export	see struct ahash_alg
import	see struct ahash_alg
setkey	see struct ahash_alg
descsize	Size of the operational state for the message digest. This state size is the memory size that needs to be allocated for shash_desc.__ctx
digestsize	see struct ahash_alg
statesize	see struct ahash_alg
base	internally used

Asynchronous Message Digest API

The asynchronous message digest API is used with the ciphers of type CRYPTO_ALG_TYPE_AHASH (listed as type “ahash” in /proc/crypto)

The asynchronous cipher operation discussion provided for the CRYPTO_ALG_TYPE_ABLKCIPHER API applies here as well.

Name

crypto_alloc_ahash — allocate ahash cipher handle

Synopsis

```
struct crypto_ahash * crypto_alloc_ahash (const char * alg_name, u32 type, u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the ahash cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Description

Allocate a cipher handle for an ahash. The returned struct crypto_ahash is the cipher handle that is required for any subsequent API invocation for that ahash.

Return

allocated cipher handle in case of success; IS_ERR is true in case of an error, PTR_ERR returns the error code.

Name

crypto_free_ahash — zeroize and free the ahash handle

Synopsis

```
void crypto_free_ahash (struct crypto_ahash * tfm);
```

Arguments

tfm cipher handle to be freed

Name

crypto_ahash_init — (re)initialize message digest handle

Synopsis

```
int crypto_ahash_init (struct ahash_request * req);
```

Arguments

req ahash_request handle that already is initialized with all necessary data using the ahash_request_* API functions

Description

The call (re-)initializes the message digest referenced by the ahash_request handle. Any potentially existing state created by previous operations is discarded.

Return

0 if the message digest initialization was successful; < 0 if an error occurred

Name

crypto_ahash_digestsize — obtain message digest size

Synopsis

```
unsigned int crypto_ahash_digestsize (struct crypto_ahash * tfm);
```

Arguments

tfm cipher handle

Description

The size for the message digest created by the message digest cipher referenced with the cipher handle is returned.

Return

message digest size of cipher

Name

crypto_ahash_reqtfm — obtain cipher handle from request

Synopsis

```
struct crypto_ahash * crypto_ahash_reqtfm (struct ahash_request * req);
```

Arguments

req asynchronous request handle that contains the reference to the ahash cipher handle

Description

Return the ahash cipher handle that is registered with the asynchronous request handle ahash_request.

Return

ahash cipher handle

Name

crypto_ahash_reqsize — obtain size of the request data structure

Synopsis

```
unsigned int crypto_ahash_reqsize (struct crypto_ahash * tfm);
```

Arguments

tfm cipher handle

Description

Return the size of the ahash state size. With the crypto_ahash_export function, the caller can export the state into a buffer whose size is defined with this function.

Return

size of the ahash state

Name

crypto_ahash_setkey — set key for cipher handle

Synopsis

```
int crypto_ahash_setkey (struct crypto_ahash * tfm, const u8 * key,
unsigned int keylen);
```

Arguments

tfm cipher handle

key buffer holding the key

keylen length of the key in bytes

Description

The caller provided key is set for the ahash cipher. The cipher handle must point to a keyed hash in order for this function to succeed.

Return

0 if the setting of the key was successful; < 0 if an error occurred

Name

crypto_ahash_finup — update and finalize message digest

Synopsis

```
int crypto_ahash_finup (struct ahash_request * req);
```

Arguments

req reference to the ahash_request handle that holds all information needed to perform the cipher operation

Description

This function is a “short-hand” for the function calls of crypto_ahash_update and crypto_shash_final. The parameters have the same meaning as discussed for those separate functions.

Return

0 if the message digest creation was successful; < 0 if an error occurred

Name

crypto_ahash_final — calculate message digest

Synopsis

```
int crypto_ahash_final (struct ahash_request * req);
```

Arguments

req reference to the ahash_request handle that holds all information needed to perform the cipher operation

Description

Finalize the message digest operation and create the message digest based on all data added to the cipher handle. The message digest is placed into the output buffer registered with the ahash_request handle.

Return

0 if the message digest creation was successful; < 0 if an error occurred

Name

crypto_ahash_digest — calculate message digest for a buffer

Synopsis

```
int crypto_ahash_digest (struct ahash_request * req);
```

Arguments

req reference to the ahash_request handle that holds all information needed to perform the cipher operation

Description

This function is a “short-hand” for the function calls of crypto_ahash_init, crypto_ahash_update and crypto_ahash_final. The parameters have the same meaning as discussed for those separate three functions.

Return

0 if the message digest creation was successful; < 0 if an error occurred

Name

crypto_ahash_export — extract current message digest state

Synopsis

```
int crypto_ahash_export (struct ahash_request * req, void * out);
```

Arguments

req reference to the ahash_request handle whose state is exported

out output buffer of sufficient size that can hold the hash state

Description

This function exports the hash state of the ahash_request handle into the caller-allocated output buffer out which must have sufficient size (e.g. by calling crypto_ahash_reqsize).

Return

0 if the export was successful; < 0 if an error occurred

Name

crypto_ahash_import — import message digest state

Synopsis

```
int crypto_ahash_import (struct ahash_request * req, const void * in);
```

Arguments

req reference to ahash_request handle the state is imported into

in buffer holding the state

Description

This function imports the hash state into the ahash_request handle from the input buffer. That buffer should have been generated with the crypto_ahash_export function.

Return

0 if the import was successful; < 0 if an error occurred

Asynchronous Hash Request Handle

The ahash_request data structure contains all pointers to data required for the asynchronous cipher operation. This includes the cipher handle (which can be used by multiple ahash_request instances), pointer to plaintext and the message digest output buffer, asynchronous callback function, etc. It acts as a handle to the ahash_request_* API calls in a similar way as ahash handle to the crypto_ahash_* API calls.

Name

ahash_request_set_tfm — update cipher handle reference in request

Synopsis

```
void ahash_request_set_tfm (struct ahash_request * req, struct
crypto_ahash * tfm);
```

Arguments

req request handle to be modified

tfm cipher handle that shall be added to the request handle

Description

Allow the caller to replace the existing ahash handle in the request data structure with a different one.

Name

ahash_request_alloc — allocate request data structure

Synopsis

```
struct ahash_request * ahash_request_alloc (struct crypto_ahash * tfm,
gfp_t gfp);
```

Arguments

tfm cipher handle to be registered with the request

gfp memory allocation flag that is handed to kmalloc by the API call.

Description

Allocate the request data structure that must be used with the ahash message digest API calls. During the allocation, the provided ahash handle is registered in the request data structure.

Return

allocated request handle in case of success; IS_ERR is true in case of an error, PTR_ERR returns the error code.

Name

ahash_request_free — zeroize and free the request data structure

Synopsis

```
void ahash_request_free (struct ahash_request * req);
```

Arguments

req request data structure cipher handle to be freed

Name

ahash_request_set_callback — set asynchronous callback function

Synopsis

```
void ahash_request_set_callback (struct ahash_request * req, u32 flags,
crypto_completion_t compl, void * data);
```

Arguments

req request handle

flags specify zero or an ORing of the flags CRYPTO_TFM_REQ_MAY_BACKLOG the request queue may back log and increase the wait queue beyond the initial maximum size; CRYPTO_TFM_REQ_MAY_SLEEP the request processing may sleep

compl callback function pointer to be registered with the request handle

data The data pointer refers to memory that is not used by the kernel crypto API, but provided to the callback function for it to use. Here, the caller can provide a reference to memory the callback function can operate on. As the callback function is invoked asynchronously to the related functionality, it may need to access data structures of the related functionality which can be referenced using this pointer. The callback function can access the memory via the “data” field in the crypto_async_request data structure provided to the callback function.

Description

This function allows setting the callback function that is triggered once the cipher operation completes.

The callback function is registered with the ahash_request handle and must comply with the following template

void callback_function(struct crypto_async_request *req, int error)

Name

ahash_request_set_crypt — set data buffers

Synopsis

```
void ahash_request_set_crypt (struct ahash_request * req, struct scat-
terlist * src, u8 * result, unsigned int nbytes);
```

Arguments

req ahash_request handle to be updated

src source scatter/gather list

result buffer that is filled with the message digest -- the caller must ensure that the buffer has sufficient space by, for example, calling `crypto_ahash_digestsize`

nbytes number of bytes to process from the source scatter/gather list

Description

By using this call, the caller references the source scatter/gather list. The source scatter/gather list points to the data the message digest is to be calculated for.

Synchronous Message Digest API

The synchronous message digest API is used with the ciphers of type CRYPTO_ALG_TYPE_SHASH (listed as type “shash” in /proc/crypto)

The message digest API is able to maintain state information for the caller.

The synchronous message digest API can store user-related context in in its shash_desc request data structure.

Name

crypto_alloc_shash — allocate message digest handle

Synopsis

```
struct crypto_shash * crypto_alloc_shash (const char * alg_name, u32 type, u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the message digest cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Description

Allocate a cipher handle for a message digest. The returned struct crypto_shash is the cipher handle that is required for any subsequent API invocation for that message digest.

Return

allocated cipher handle in case of success; IS_ERR is true in case of an error, PTR_ERR returns the error code.

Name

crypto_free_shash — zeroize and free the message digest handle

Synopsis

```
void crypto_free_shash (struct crypto_shash * tfm);
```

Arguments

tfm cipher handle to be freed

Name

crypto_shash_blocksize — obtain block size for cipher

Synopsis

```
unsigned int crypto_shash_blocksize (struct crypto_shash * tfm);
```

Arguments

tfm cipher handle

Description

The block size for the message digest cipher referenced with the cipher handle is returned.

Return

block size of cipher

Name

crypto_shash_digestsize — obtain message digest size

Synopsis

```
unsigned int crypto_shash_digestsize (struct crypto_shash * tfm);
```

Arguments

tfm cipher handle

Description

The size for the message digest created by the message digest cipher referenced with the cipher handle is returned.

Return

digest size of cipher

Name

crypto_shash_descsize — obtain the operational state size

Synopsis

```
unsigned int crypto_shash_descsize (struct crypto_shash * tfm);
```

Arguments

tfm cipher handle

Description

The size of the operational state the cipher needs during operation is returned for the hash referenced with the cipher handle. This size is required to calculate the memory requirements to allow the caller allocating sufficient memory for operational state.

The operational state is defined with struct shash_desc where the size of that data structure is to be calculated as sizeof(struct shash_desc) + crypto_shash_descsize(alg)

Return

size of the operational state

Name

crypto_shash_setkey — set key for message digest

Synopsis

```
int crypto_shash_setkey (struct crypto_shash * tfm, const u8 * key,
unsigned int keylen);
```

Arguments

tfm cipher handle

key buffer holding the key

keylen length of the key in bytes

Description

The caller provided key is set for the keyed message digest cipher. The cipher handle must point to a keyed message digest cipher in order for this function to succeed.

Return

0 if the setting of the key was successful; < 0 if an error occurred

Name

crypto_shash_digest — calculate message digest for buffer

Synopsis

```
int crypto_shash_digest (struct shash_desc * desc, const u8 * data,
unsigned int len, u8 * out);
```

Arguments

desc see `crypto_shash_final`

data see `crypto_shash_update`

len see `crypto_shash_update`

out see `crypto_shash_final`

Description

This function is a “short-hand” for the function calls of crypto_shash_init, crypto_shash_update and crypto_shash_final. The parameters have the same meaning as discussed for those separate three functions.

Return

0 if the message digest creation was successful; < 0 if an error occurred

Name

crypto_shash_export — extract operational state for message digest

Synopsis

```
int crypto_shash_export (struct shash_desc * desc, void * out);
```

Arguments

desc reference to the operational state handle whose state is exported

out output buffer of sufficient size that can hold the hash state

Description

This function exports the hash state of the operational state handle into the caller-allocated output buffer out which must have sufficient size (e.g. by calling crypto_shash_descsize).

Return

0 if the export creation was successful; < 0 if an error occurred

Name

crypto_shash_import — import operational state

Synopsis

```
int crypto_shash_import (struct shash_desc * desc, const void * in);
```

Arguments

desc reference to the operational state handle the state imported into

in buffer holding the state

Description

This function imports the hash state into the operational state handle from the input buffer. That buffer should have been generated with the crypto_ahash_export function.

Return

0 if the import was successful; < 0 if an error occurred

Name

crypto_shash_init — (re)initialize message digest

Synopsis

```
int crypto_shash_init (struct shash_desc * desc);
```

Arguments

desc operational state handle that is already filled

Description

The call (re-)initializes the message digest referenced by the operational state handle. Any potentially existing state created by previous operations is discarded.

Return

0 if the message digest initialization was successful; < 0 if an error occurred

Name

crypto_shash_update — add data to message digest for processing

Synopsis

```
int crypto_shash_update (struct shash_desc * desc, const u8 * data,
unsigned int len);
```

Arguments

desc operational state handle that is already initialized

data input data to be added to the message digest

len length of the input data

Description

Updates the message digest state of the operational state handle.

Return

0 if the message digest update was successful; < 0 if an error occurred

Name

crypto_shash_final — calculate message digest

Synopsis

```
int crypto_shash_final (struct shash_desc * desc, u8 * out);
```

Arguments

desc operational state handle that is already filled with data

out output buffer filled with the message digest

Description

Finalize the message digest operation and create the message digest based on all data added to the cipher handle. The message digest is placed into the output buffer. The caller must ensure that the output buffer is large enough by using crypto_shash_digestsize.

Return

0 if the message digest creation was successful; < 0 if an error occurred

Name

crypto_shash_finup — calculate message digest of buffer

Synopsis

```
int crypto_shash_finup (struct shash_desc * desc, const u8 * data,
unsigned int len, u8 * out);
```

Arguments

desc see `crypto_shash_final`

data see `crypto_shash_update`

len see `crypto_shash_update`

out see `crypto_shash_final`

Description

This function is a “short-hand” for the function calls of crypto_shash_update and crypto_shash_final. The parameters have the same meaning as discussed for those separate functions.

Return

0 if the message digest creation was successful; < 0 if an error occurred

Crypto API Random Number API

The random number generator API is used with the ciphers of type CRYPTO_ALG_TYPE_RNG (listed as type “rng” in /proc/crypto)

Name

crypto_alloc_rng — - allocate RNG handle

Synopsis

```
struct crypto_rng * crypto_alloc_rng (const char * alg_name, u32 type,
u32 mask);
```

Arguments

alg_name is the cra_name / name or cra_driver_name / driver name of the message digest cipher

type specifies the type of the cipher

mask specifies the mask for the cipher

Description

Allocate a cipher handle for a random number generator. The returned struct crypto_rng is the cipher handle that is required for any subsequent API invocation for that random number generator.

For all random number generators, this call creates a new private copy of the random number generator that does not share a state with other instances. The only exception is the “krng” random number generator which is a kernel crypto API use case for the `get_random_bytes` function of the /dev/random driver.

Return

allocated cipher handle in case of success; `IS_ERR` is true in case of an error, `PTR_ERR` returns the error code.

Name

crypto_rng_alg — obtain name of RNG

Synopsis

```
struct rng_alg * crypto_rng_alg (struct crypto_rng * tfm);
```

Arguments

tfm cipher handle

Description

Return the generic name (cra_name) of the initialized random number generator

Return

generic name string

Name

crypto_free_rng — zeroize and free RNG handle

Synopsis

```
void crypto_free_rng (struct crypto_rng * tfm);
```

Arguments

tfm cipher handle to be freed

Name

crypto_rng_get_bytes — get random number

Synopsis

```
int crypto_rng_get_bytes (struct crypto_rng * tfm, u8 * rdata, unsigned
int dlen);
```

Arguments

tfm cipher handle

rdata output buffer holding the random numbers

dlen length of the output buffer

Description

This function fills the caller-allocated buffer with random numbers using the random number generator referenced by the cipher handle.

Return

0 function was successful; < 0 if an error occurred

Name

crypto_rng_reset — re-initialize the RNG

Synopsis

```
int crypto_rng_reset (struct crypto_rng * tfm, const u8 * seed, unsigned
int slen);
```

Arguments

tfm cipher handle

seed seed input data

slen length of the seed input data

Description

The reset function completely re-initializes the random number generator referenced by the cipher handle by clearing the current state. The new state is initialized with the caller provided seed or automatically, depending on the random number generator type (the ANSI X9.31 RNG requires caller-provided seed, the SP800-90A DRBGs perform an automatic seeding). The seed is provided as a parameter to this function call. The provided seed should have the length of the seed size defined for the random number generator as defined by crypto_rng_seedsize.

Return

0 if the setting of the key was successful; < 0 if an error occurred

Name

crypto_rng_seedsize — obtain seed size of RNG

Synopsis

```
int crypto_rng_seedsize (struct crypto_rng * tfm);
```

Arguments

tfm cipher handle

Description

The function returns the seed size for the random number generator referenced by the cipher handle. This value may be zero if the random number generator does not implement or require a reseeding. For example, the SP800-90A DRBGs implement an automated reseeding after reaching a pre-defined threshold.

Return

seed size for the random number generator

Chapter 6. Code Examples

Code Example For Asynchronous Block Cipher Operation

```
struct tcrypt_result {
 struct completion completion;
 int err;
};

/* tie all data structures together */
struct ablkcipher_def {
 struct scatterlist sg;
 struct crypto_ablkcipher *tfm;
 struct ablkcipher_request *req;
 struct tcrypt_result result;
};

/* Callback function */
static void test_ablkcipher_cb(struct crypto_async_request *req, int error)
{
 struct tcrypt_result *result = req->data;

 if (error == -EINPROGRESS)
  return;
 result->err = error;
 complete(&result->completion);
 pr_info("Encryption finished successfully\n");
}

/* Perform cipher operation */
static unsigned int test_ablkcipher_encdec(struct ablkcipher_def *ablk,
        int enc)
{
 int rc = 0;

 if (enc)
  rc = crypto_ablkcipher_encrypt(ablk->req);
 else
  rc = crypto_ablkcipher_decrypt(ablk->req);

 switch (rc) {
 case 0:
  break;
 case -EINPROGRESS:
 case -EBUSY:
  rc = wait_for_completion_interruptible(
   &ablk->result.completion);
  if (!rc && !ablk->result.err) {
```

```
   reinit_completion(&ablk->result.completion);
   break;
  }
 default:
  pr_info("ablkcipher encrypt returned with %d result %d\n",
         rc, ablk->result.err);
  break;
 }
 init_completion(&ablk->result.completion);

 return rc;
}

/* Initialize and trigger cipher operation */
static int test_ablkcipher(void)
{
 struct ablkcipher_def ablk;
 struct crypto_ablkcipher *ablkcipher = NULL;
 struct ablkcipher_request *req = NULL;
 char *scratchpad = NULL;
 char *ivdata = NULL;
 unsigned char key[32];
 int ret = -EFAULT;

 ablkcipher = crypto_alloc_ablkcipher("cbc-aes-aesni", 0, 0);
 if (IS_ERR(ablkcipher)) {
  pr_info("could not allocate ablkcipher handle\n");
  return PTR_ERR(ablkcipher);
 }

 req = ablkcipher_request_alloc(ablkcipher, GFP_KERNEL);
 if (IS_ERR(req)) {
  pr_info("could not allocate request queue\n");
  ret = PTR_ERR(req);
  goto out;
 }

 ablkcipher_request_set_callback(req, CRYPTO_TFM_REQ_MAY_BACKLOG,
     test_ablkcipher_cb,
     &ablk.result);

 /* AES 256 with random key */
 get_random_bytes(&key, 32);
 if (crypto_ablkcipher_setkey(ablkcipher, key, 32)) {
  pr_info("key could not be set\n");
  ret = -EAGAIN;
  goto out;
 }

 /* IV will be random */
 ivdata = kmalloc(16, GFP_KERNEL);
 if (!ivdata) {
  pr_info("could not allocate ivdata\n");
  goto out;
```

```
 }
 get_random_bytes(ivdata, 16);

 /* Input data will be random */
 scratchpad = kmalloc(16, GFP_KERNEL);
 if (!scratchpad) {
  pr_info("could not allocate scratchpad\n");
  goto out;
 }
 get_random_bytes(scratchpad, 16);

 ablk.tfm = ablkcipher;
 ablk.req = req;

 /* We encrypt one block */
 sg_init_one(&ablk.sg, scratchpad, 16);
 ablkcipher_request_set_crypt(req, &ablk.sg, &ablk.sg, 16, ivdata);
 init_completion(&ablk.result.completion);

 /* encrypt data */
 ret = test_ablkcipher_encdec(&ablk, 1);
 if (ret)
  goto out;

 pr_info("Encryption triggered successfully\n");

out:
 if (ablkcipher)
  crypto_free_ablkcipher(ablkcipher);
 if (req)
  ablkcipher_request_free(req);
 if (ivdata)
  kfree(ivdata);
 if (scratchpad)
  kfree(scratchpad);
 return ret;
}
```

Code Example For Synchronous Block Cipher Operation

```
static int test_blkcipher(void)
{
 struct crypto_blkcipher *blkcipher = NULL;
 char *cipher = "cbc(aes)";
 // AES 128
 charkey =
"\x12\x34\x56\x78\x90\xab\xcd\xef\x12\x34\x56\x78\x90\xab\xcd\xef";
 chariv =
```

```
"\x12\x34\x56\x78\x90\xab\xcd\xef\x12\x34\x56\x78\x90\xab\xcd\xef";
 unsigned int ivsize = 0;
 char *scratchpad = NULL; // holds plaintext and ciphertext
 struct scatterlist sg;
 struct blkcipher_desc desc;
 int ret = -EFAULT;

 blkcipher = crypto_alloc_blkcipher(cipher, 0, 0);
 if (IS_ERR(blkcipher)) {
  printk("could not allocate blkcipher handle for %s\n", cipher);
  return -PTR_ERR(blkcipher);
 }

 if (crypto_blkcipher_setkey(blkcipher, key, strlen(key))) {
  printk("key could not be set\n");
  ret = -EAGAIN;
  goto out;
 }

 ivsize = crypto_blkcipher_ivsize(blkcipher);
 if (ivsize) {
  if (ivsize != strlen(iv))
   printk("IV length differs from expected length\n");
  crypto_blkcipher_set_iv(blkcipher, iv, ivsize);
 }

 scratchpad = kmalloc(crypto_blkcipher_blocksize(blkcipher), GFP_KERNEL);
 if (!scratchpad) {
  printk("could not allocate scratchpad for %s\n", cipher);
  goto out;
 }
 /* get some random data that we want to encrypt */
 get_random_bytes(scratchpad, crypto_blkcipher_blocksize(blkcipher));

 desc.flags = 0;
 desc.tfm = blkcipher;
 sg_init_one(&sg, scratchpad, crypto_blkcipher_blocksize(blkcipher));

 /* encrypt data in place */
 crypto_blkcipher_encrypt(&desc, &sg, &sg,
     crypto_blkcipher_blocksize(blkcipher));

 /* decrypt data in place
  * crypto_blkcipher_decrypt(&desc, &sg, &sg,
  */    crypto_blkcipher_blocksize(blkcipher));


 printk("Cipher operation completed\n");
 return 0;

out:
 if (blkcipher)
  crypto_free_blkcipher(blkcipher);
 if (scratchpad)
```

```
  kzfree(scratchpad);
 return ret;
}
```

Code Example For Use of Operational State Memory With SHASH

```
struct sdesc {
 struct shash_desc shash;
 char ctx[];
};

static struct sdescinit_sdesc(struct crypto_shash *alg)
{
 struct sdescsdesc;
 int size;

 size = sizeof(struct shash_desc) + crypto_shash_descsize(alg);
 sdesc = kmalloc(size, GFP_KERNEL);
 if (!sdesc)
  return ERR_PTR(-ENOMEM);
 sdesc->shash.tfm = alg;
 sdesc->shash.flags = 0x0;
 return sdesc;
}

static int calc_hash(struct crypto_shashalg,
       const unsigned chardata, unsigned int datalen,
       unsigned chardigest) {
 struct sdescsdesc;
 int ret;

 sdesc = init_sdesc(alg);
 if (IS_ERR(sdesc)) {
  pr_info("trusted_key: can't alloc %s\n", hash_alg);
  return PTR_ERR(sdesc);
 }

 ret = crypto_shash_digest(&sdesc->shash, data, datalen, digest);
 kfree(sdesc);
 return ret;
}
```

Code Example For Random Number Generator Usage

```
static int get_random_numbers(u8 *buf, unsigned int len)
{
 struct crypto_rngrng = NULL;
 chardrbg = "drbg_nopr_sha256"; /* Hash DRBG with SHA-256, no PR */
 int ret;

 if (!buf || !len) {
  pr_debug("No output buffer provided\n");
  return -EINVAL;
 }

 rng = crypto_alloc_rng(drbg, 0, 0);
 if (IS_ERR(rng)) {
  pr_debug("could not allocate RNG handle for %s\n", drbg);
  return -PTR_ERR(rng);
 }

 ret = crypto_rng_get_bytes(rng, buf, len);
 if (ret < 0)
  pr_debug("generation of random numbers failed\n");
 else if (ret == 0)
  pr_debug("RNG returned no data");
 else
  pr_debug("RNG returned %d bytes of data\n", ret);

out:
 crypto_free_rng(rng);
 return ret;
}
```