

Linux Kernel Network Security - Transport Layer Security (TLS)

Deep Hacking

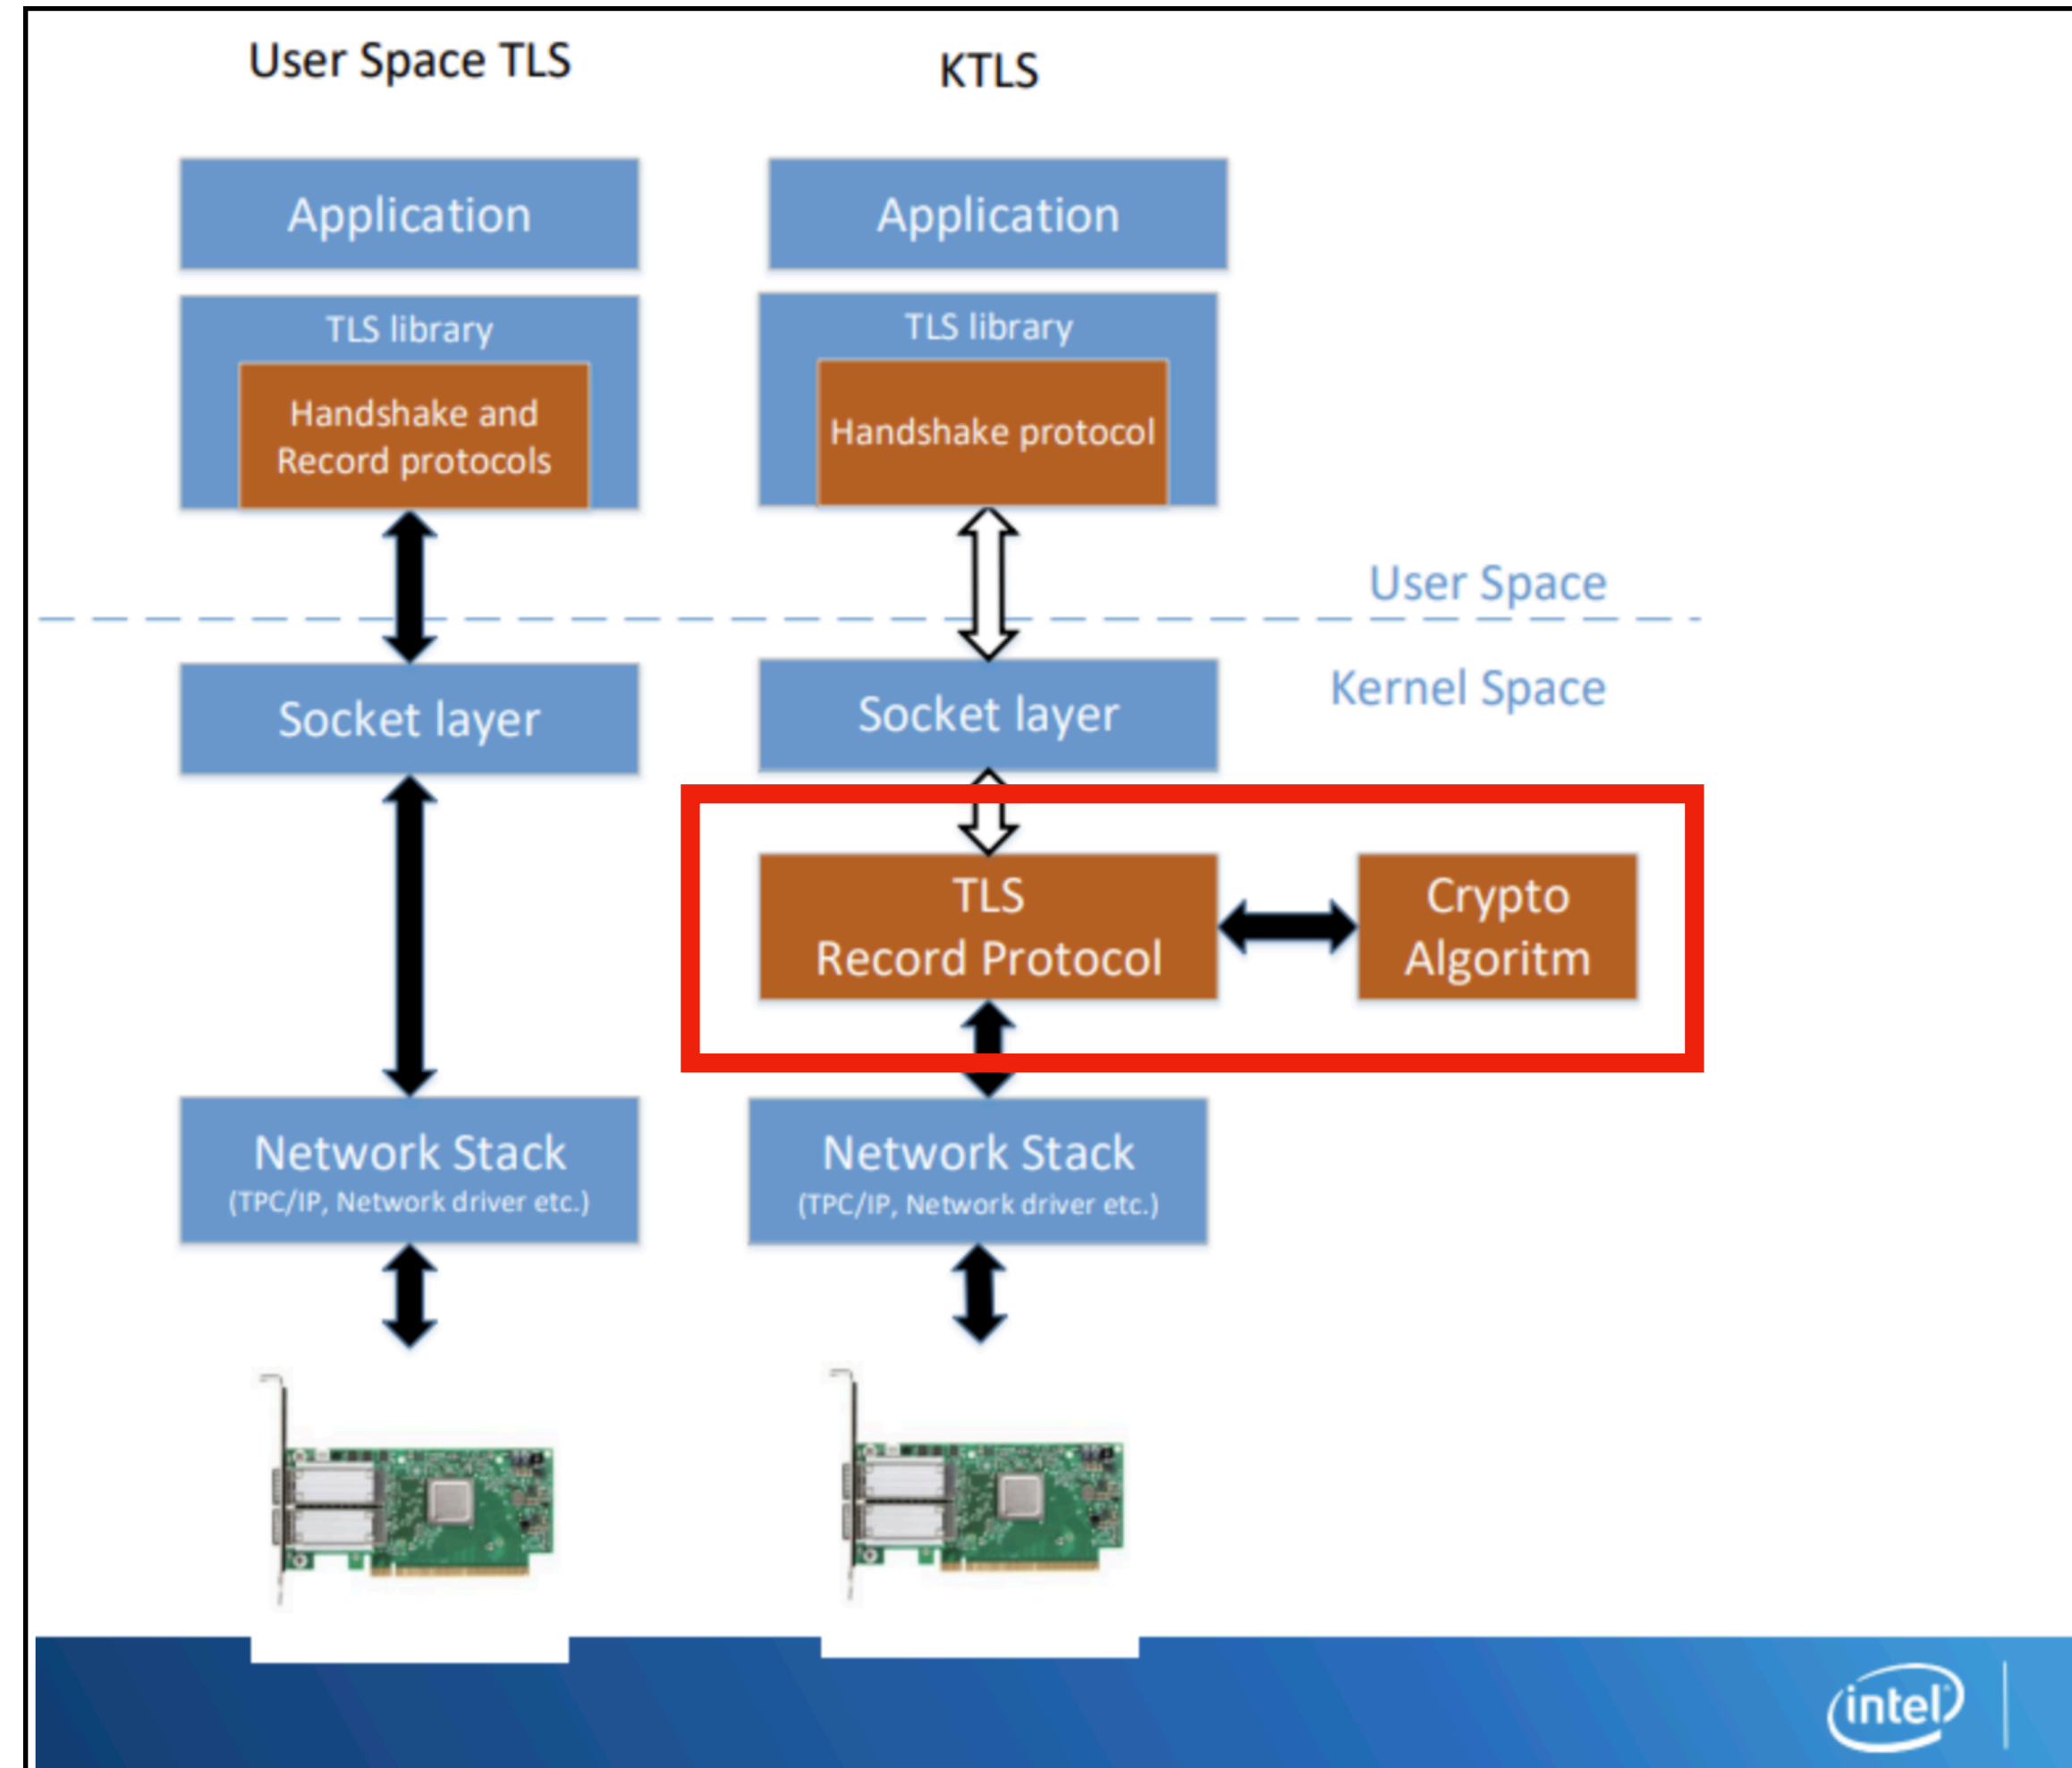
2025/01/19 Pumpkin 🎃

Outline

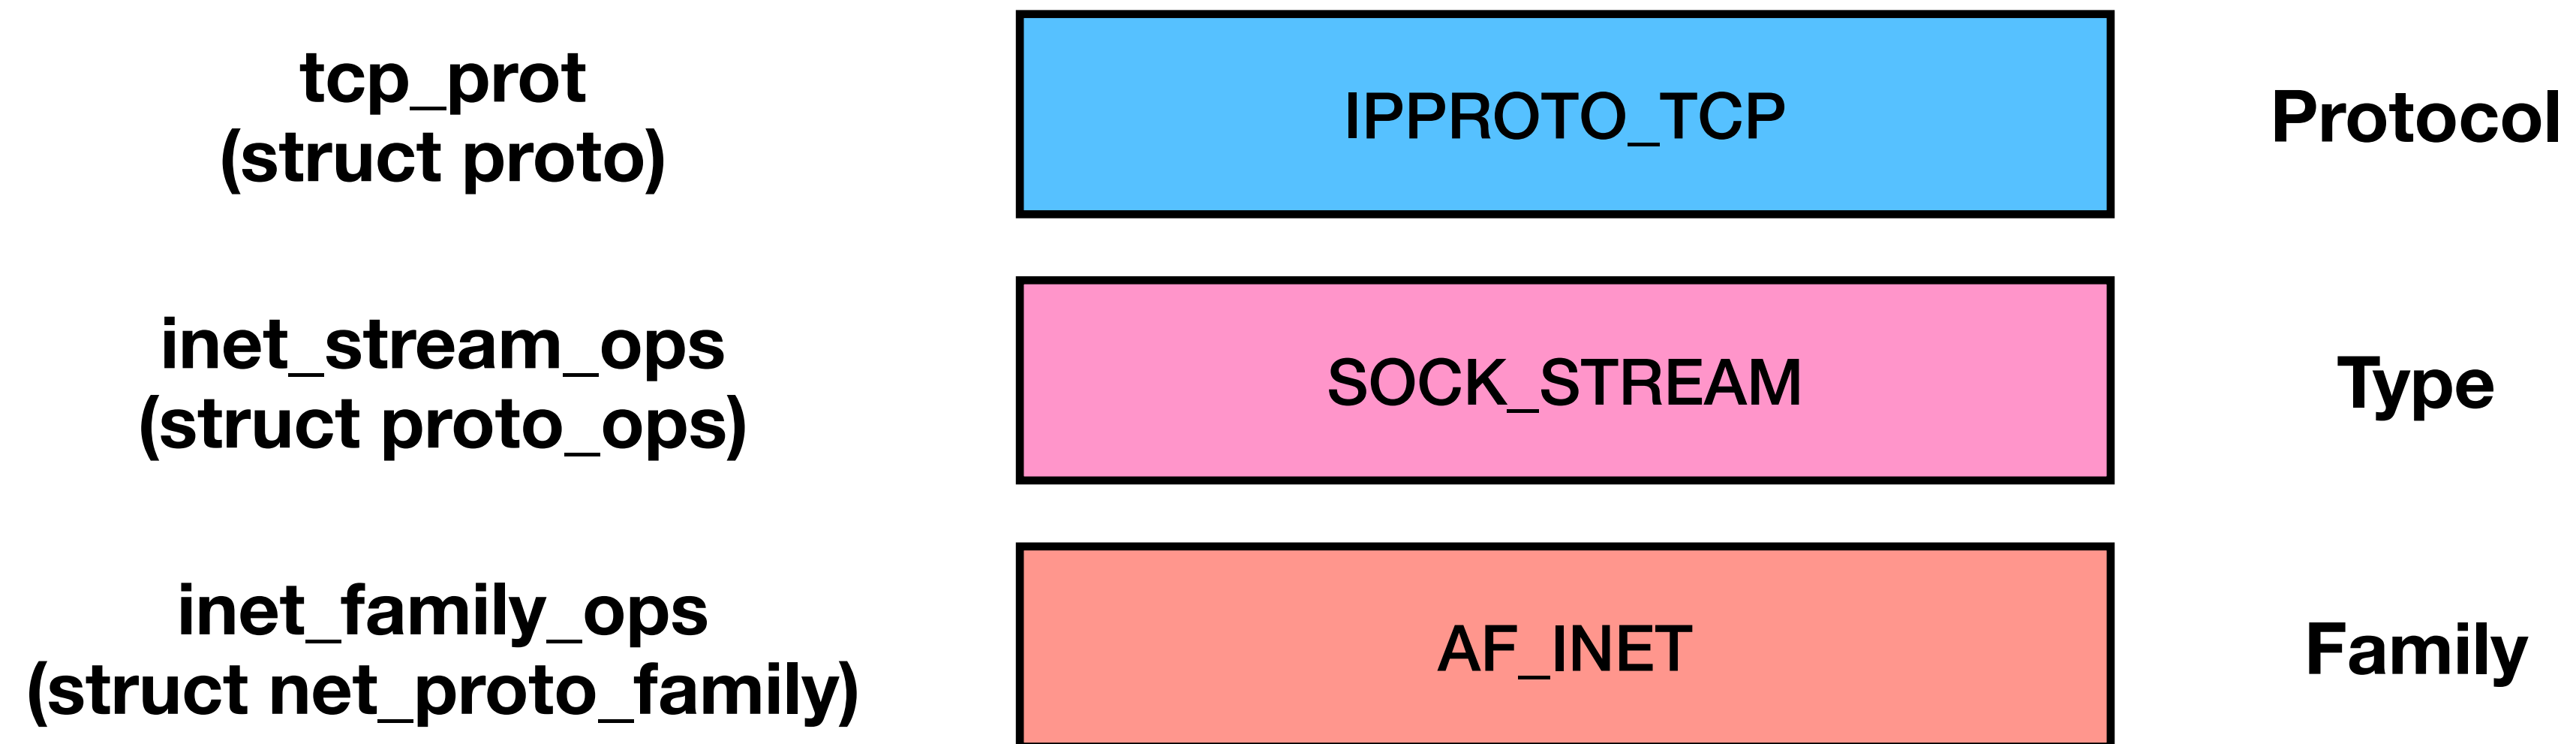
- Overview
- Vulnerability

Overview

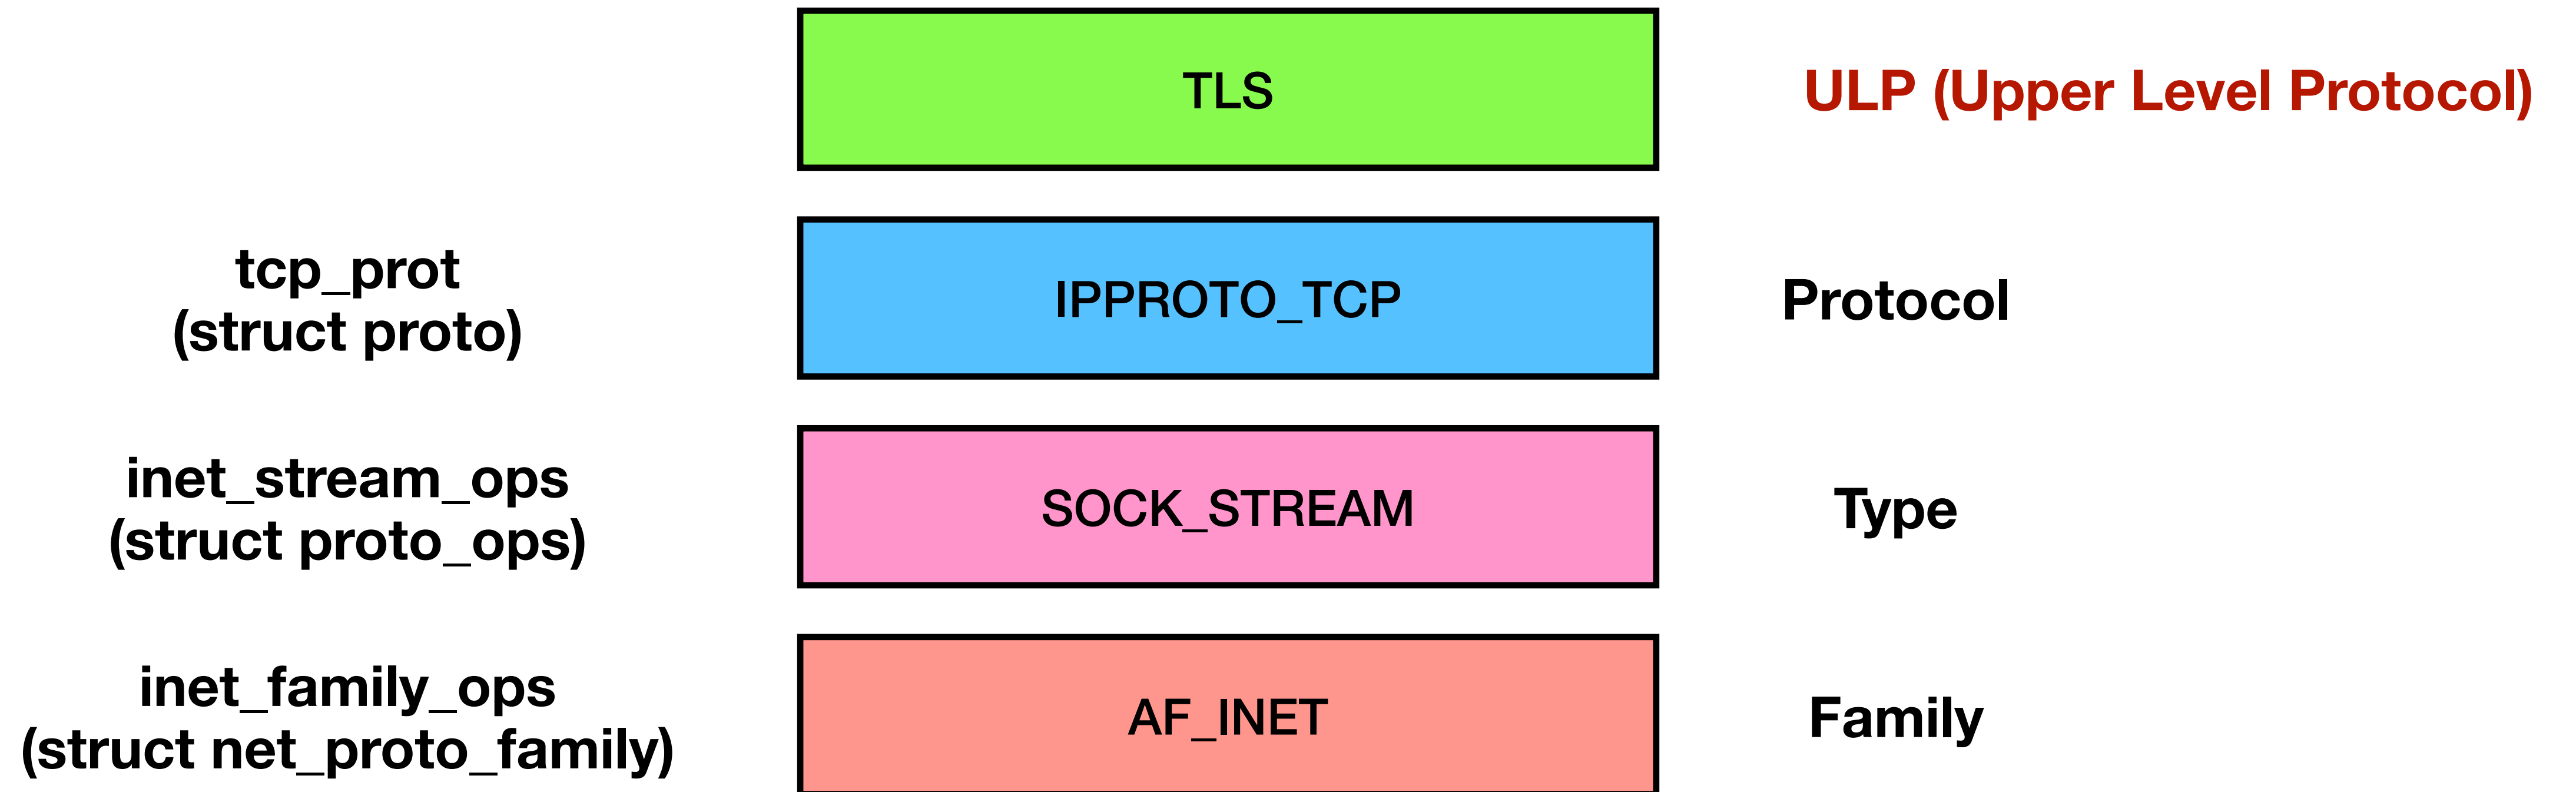
Overview



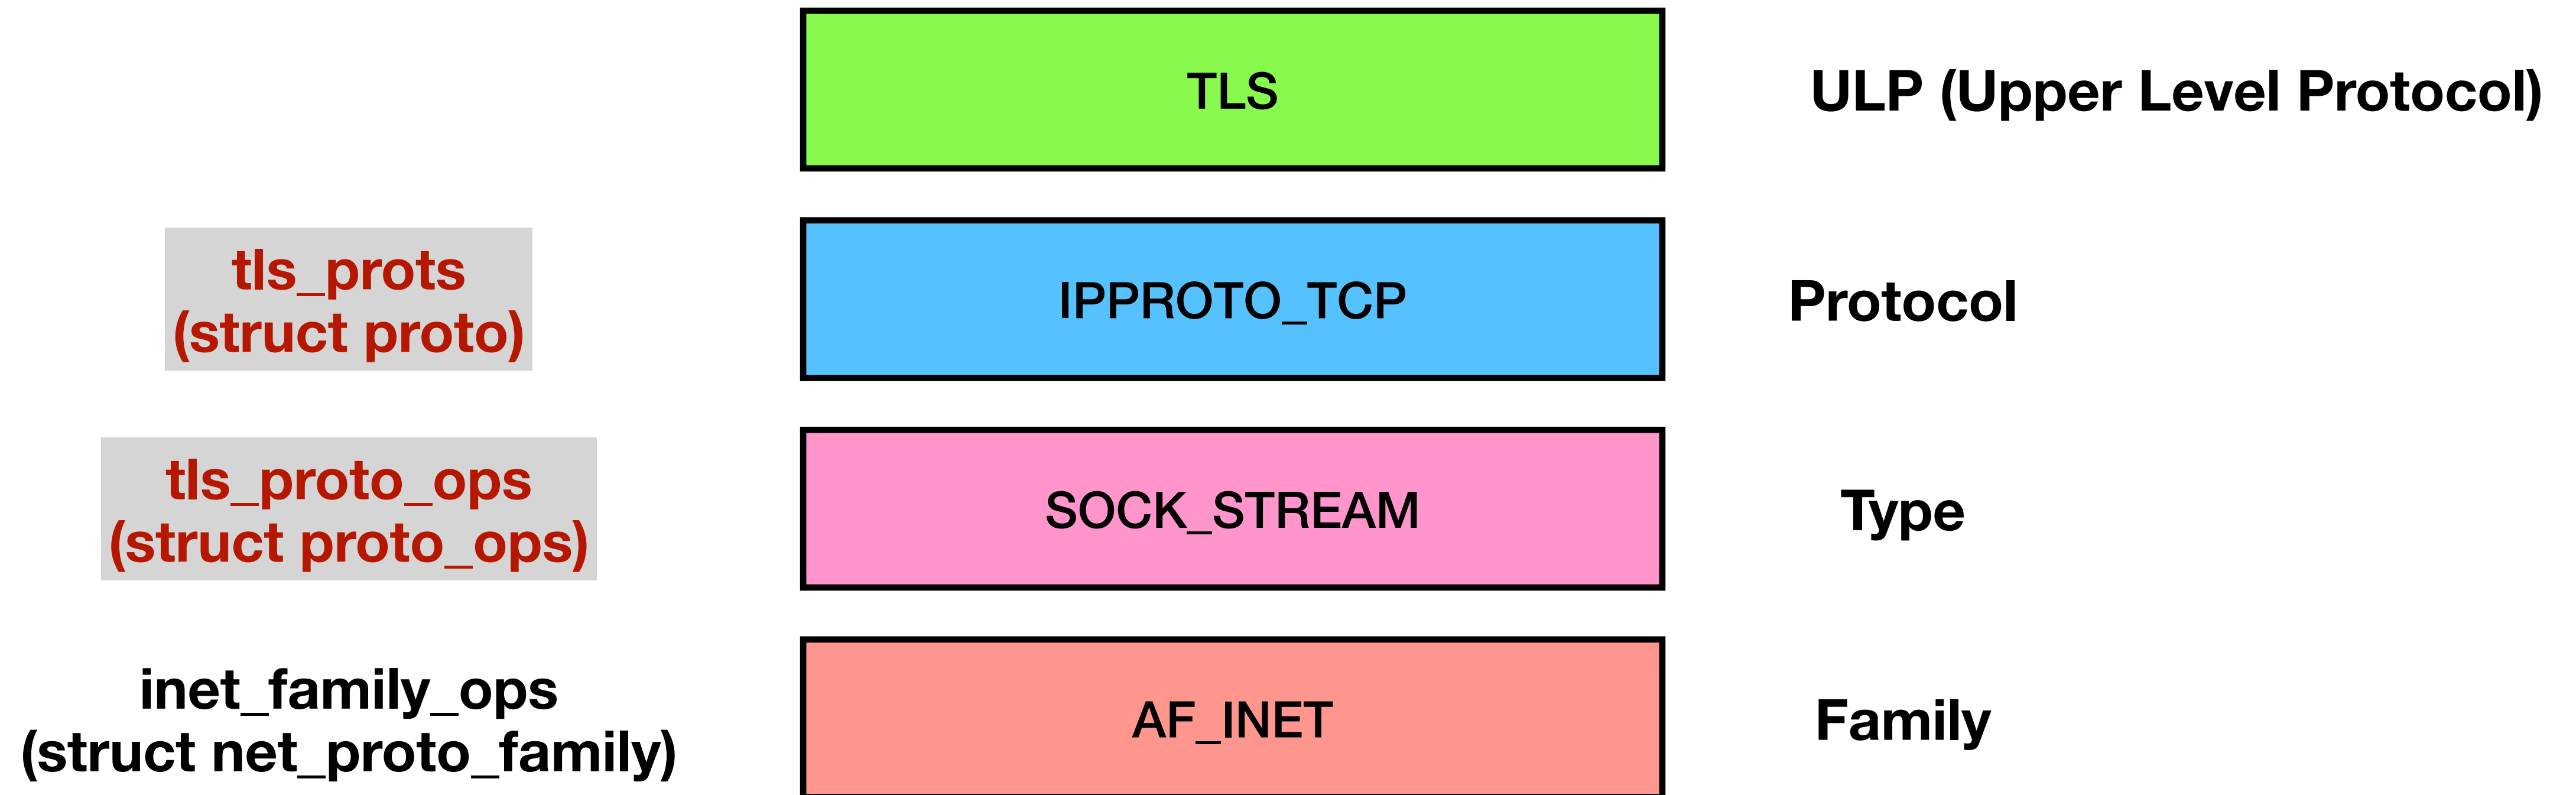
Overview

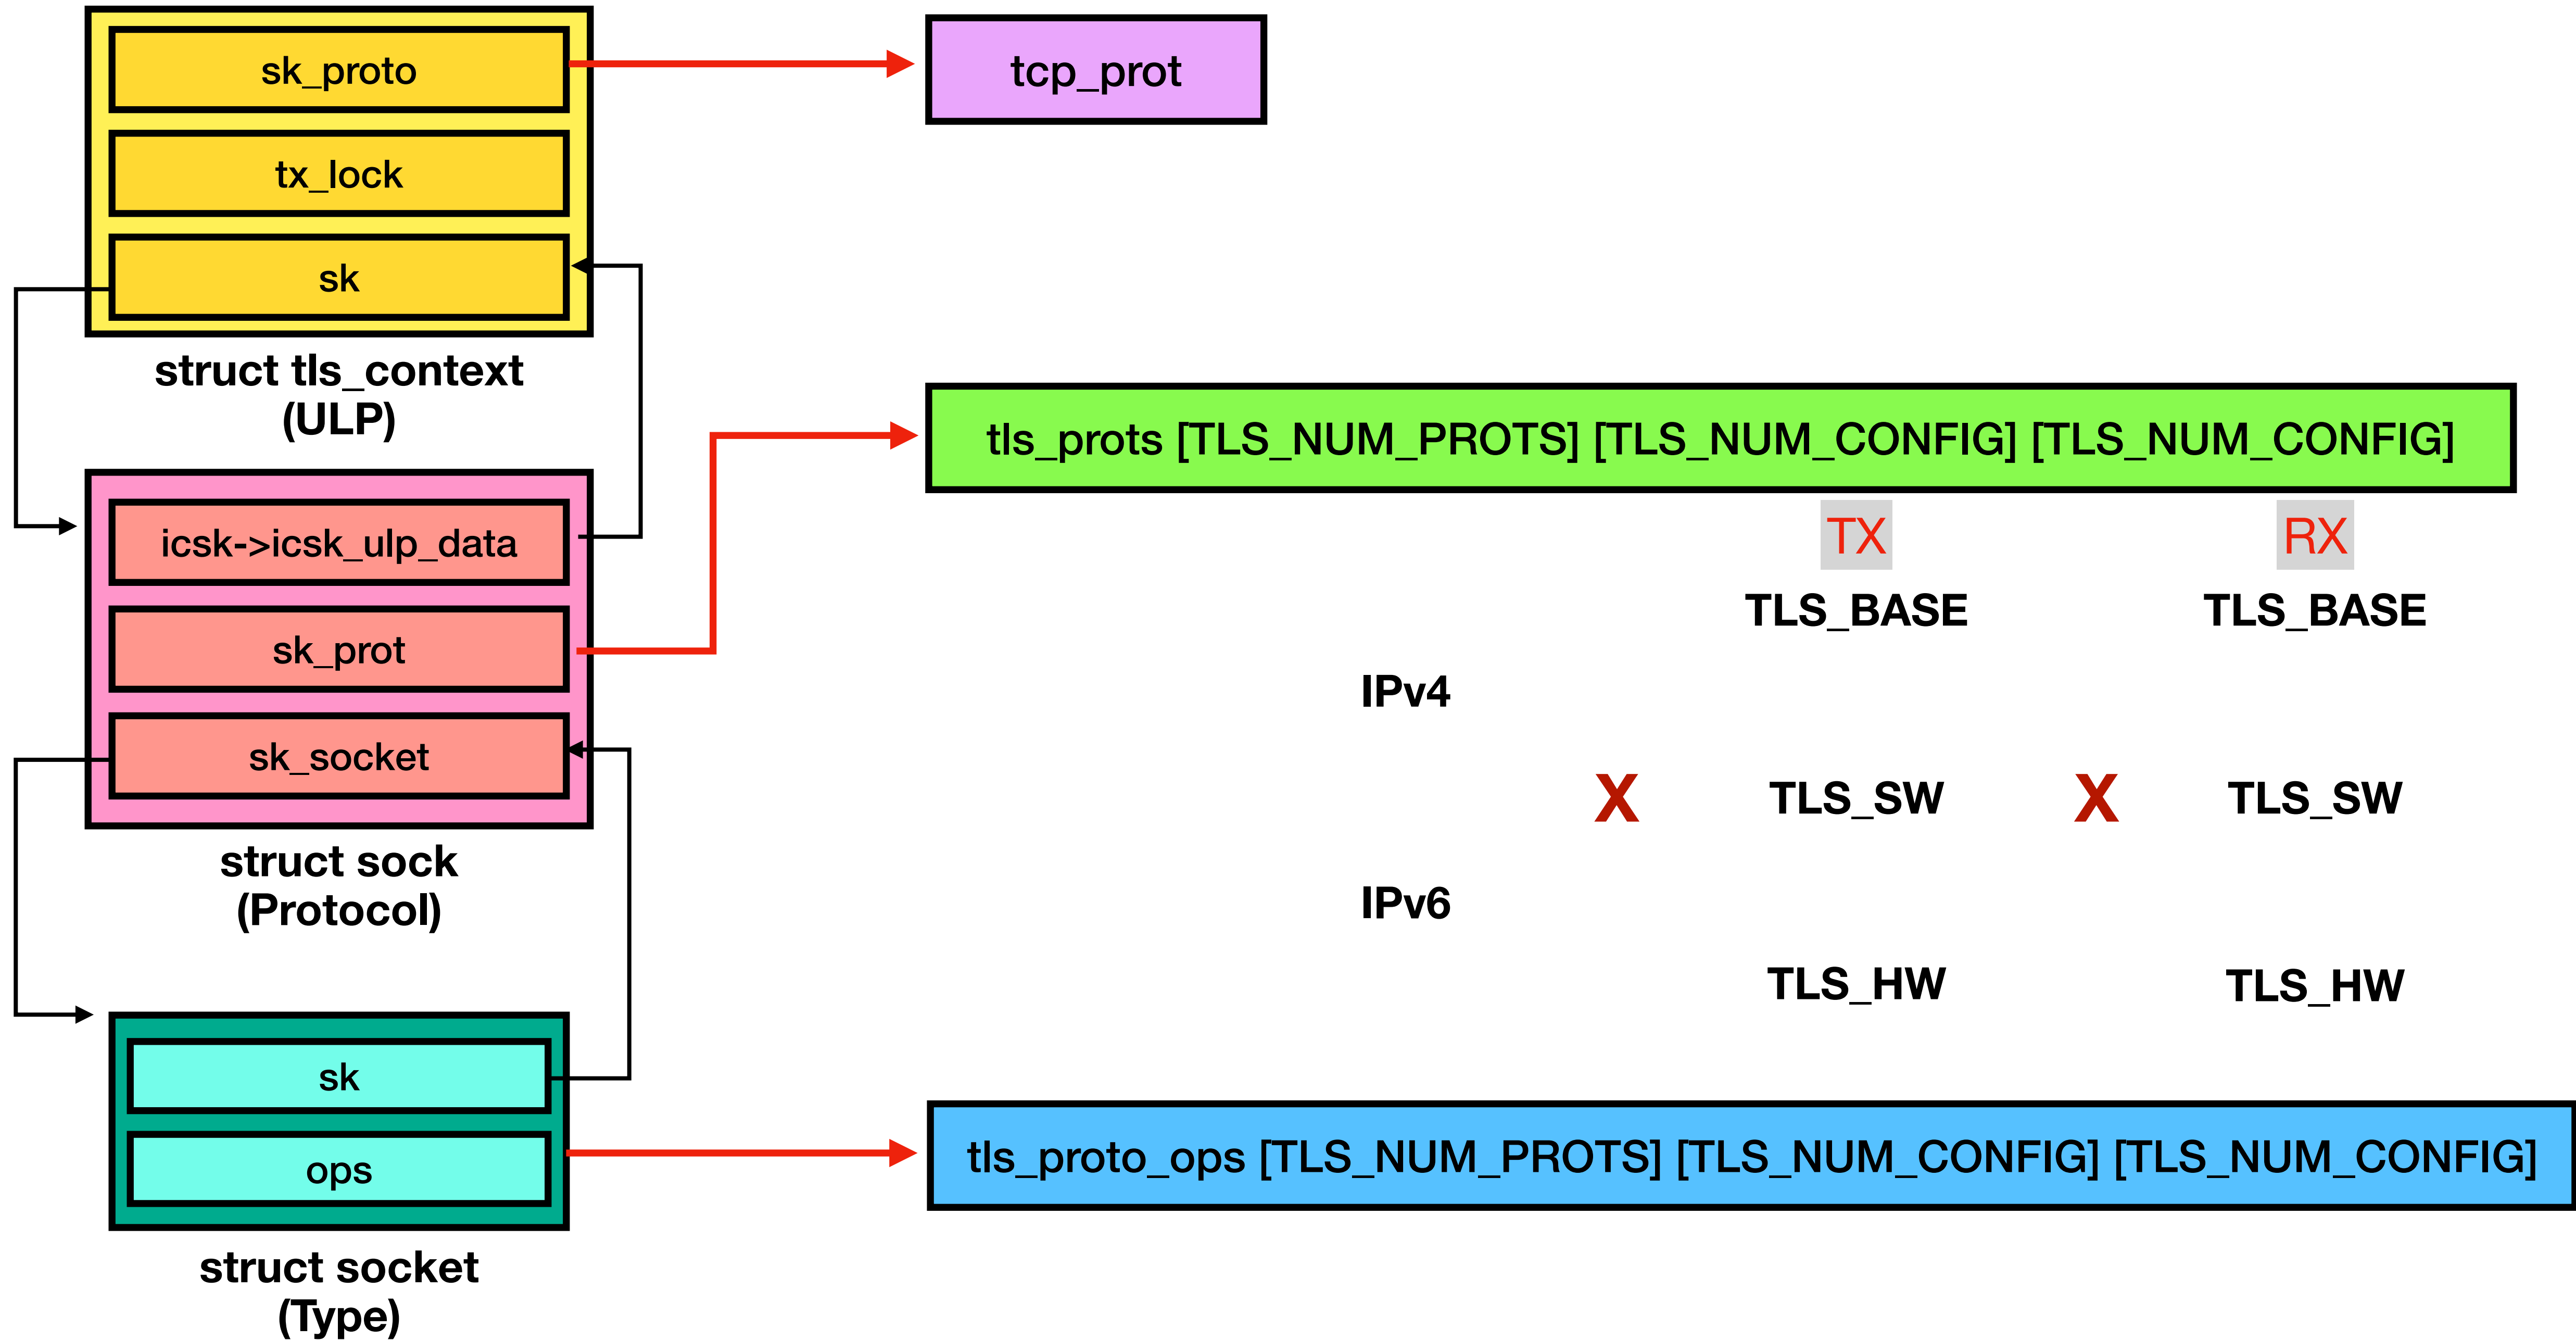


Overview

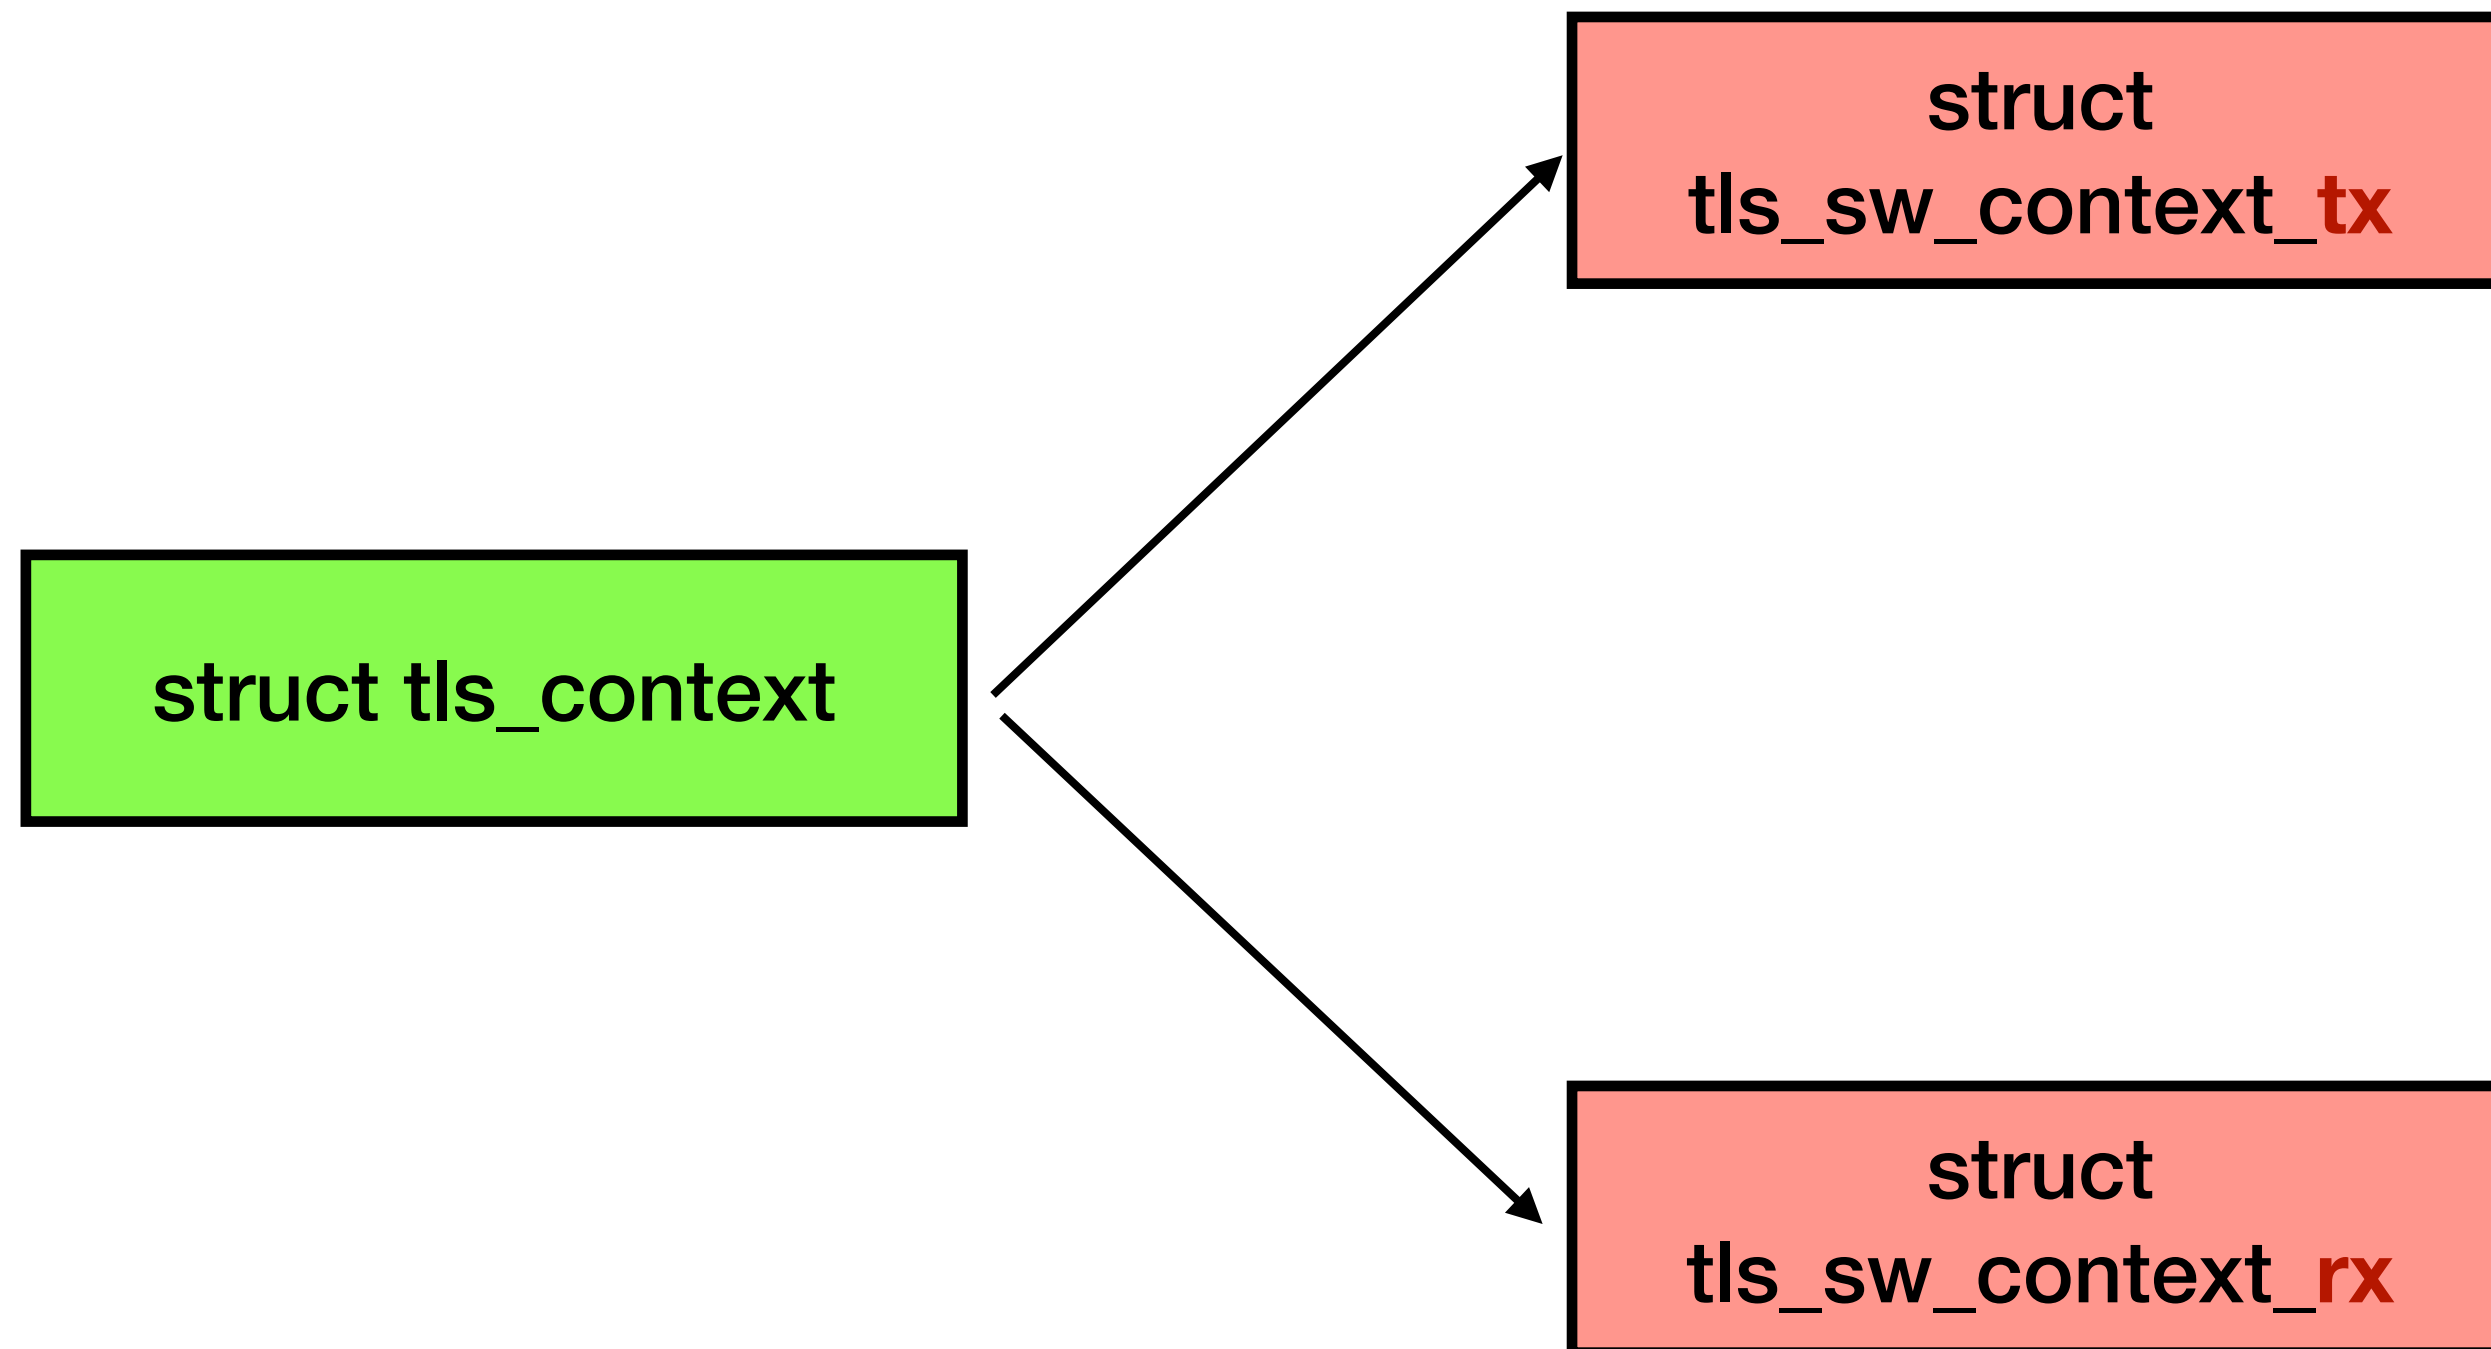


Overview





Overview



cipher type	AES_GCM_128
TLS vers.	TLS_1_2
KEY	0123...DEF
IV	12345678
SALT	SALT
...	...

cipher type	...
TLS vers.	...
KEY	...
IV	...
SALT	...
...	...

Overview

```
// 1. config a TCP socket to TLS
setsockopt(sockfd, SOL_TCP, TCP_ULP, "tls", sizeof("tls"));

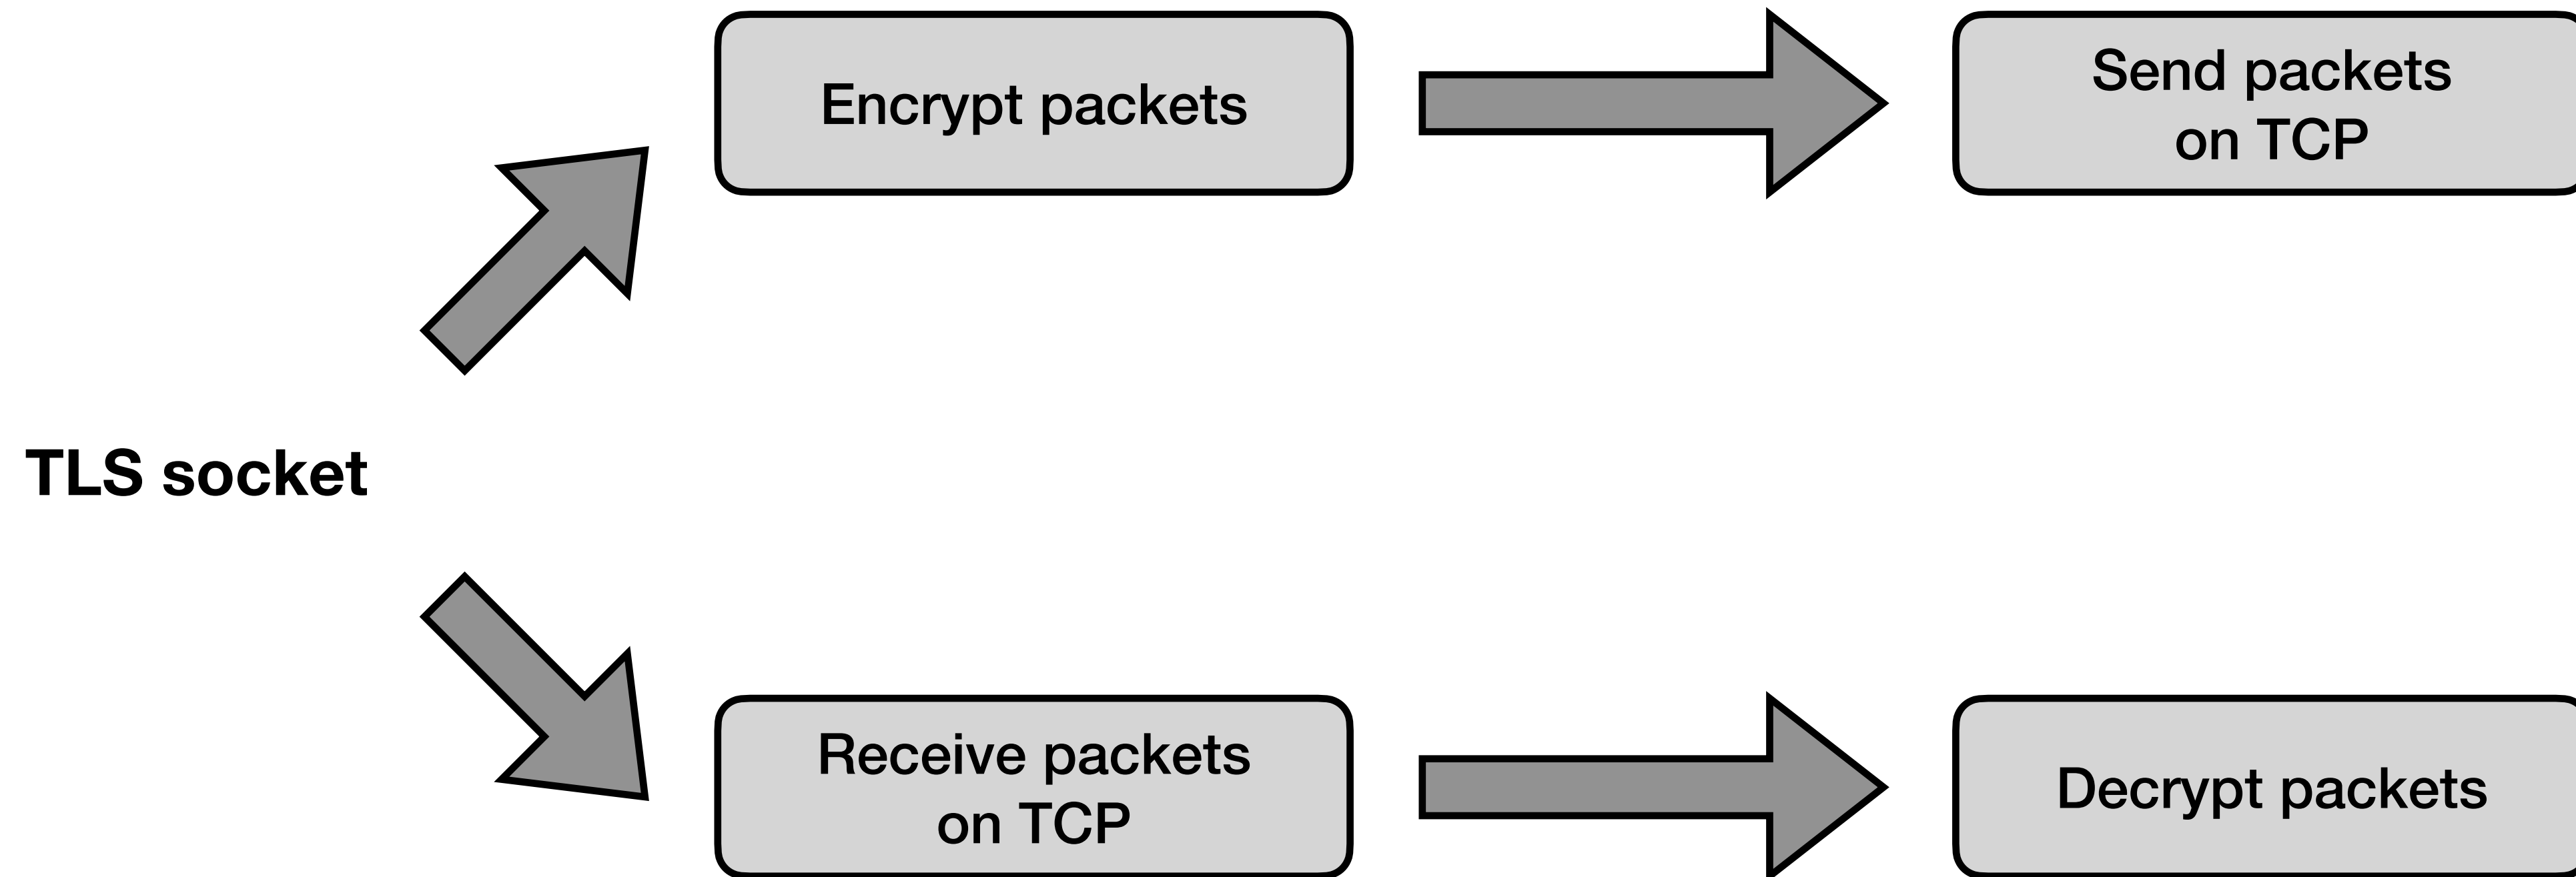
// 2. config TX / RX of the TLS socket
struct tls12_crypto_info_aes_gcm_128 crypto_info = {};

crypto_info.info.version = TLS_1_2_VERSION;
crypto_info.info.cipher_type = TLS_CIPHER_AES_GCM_128;

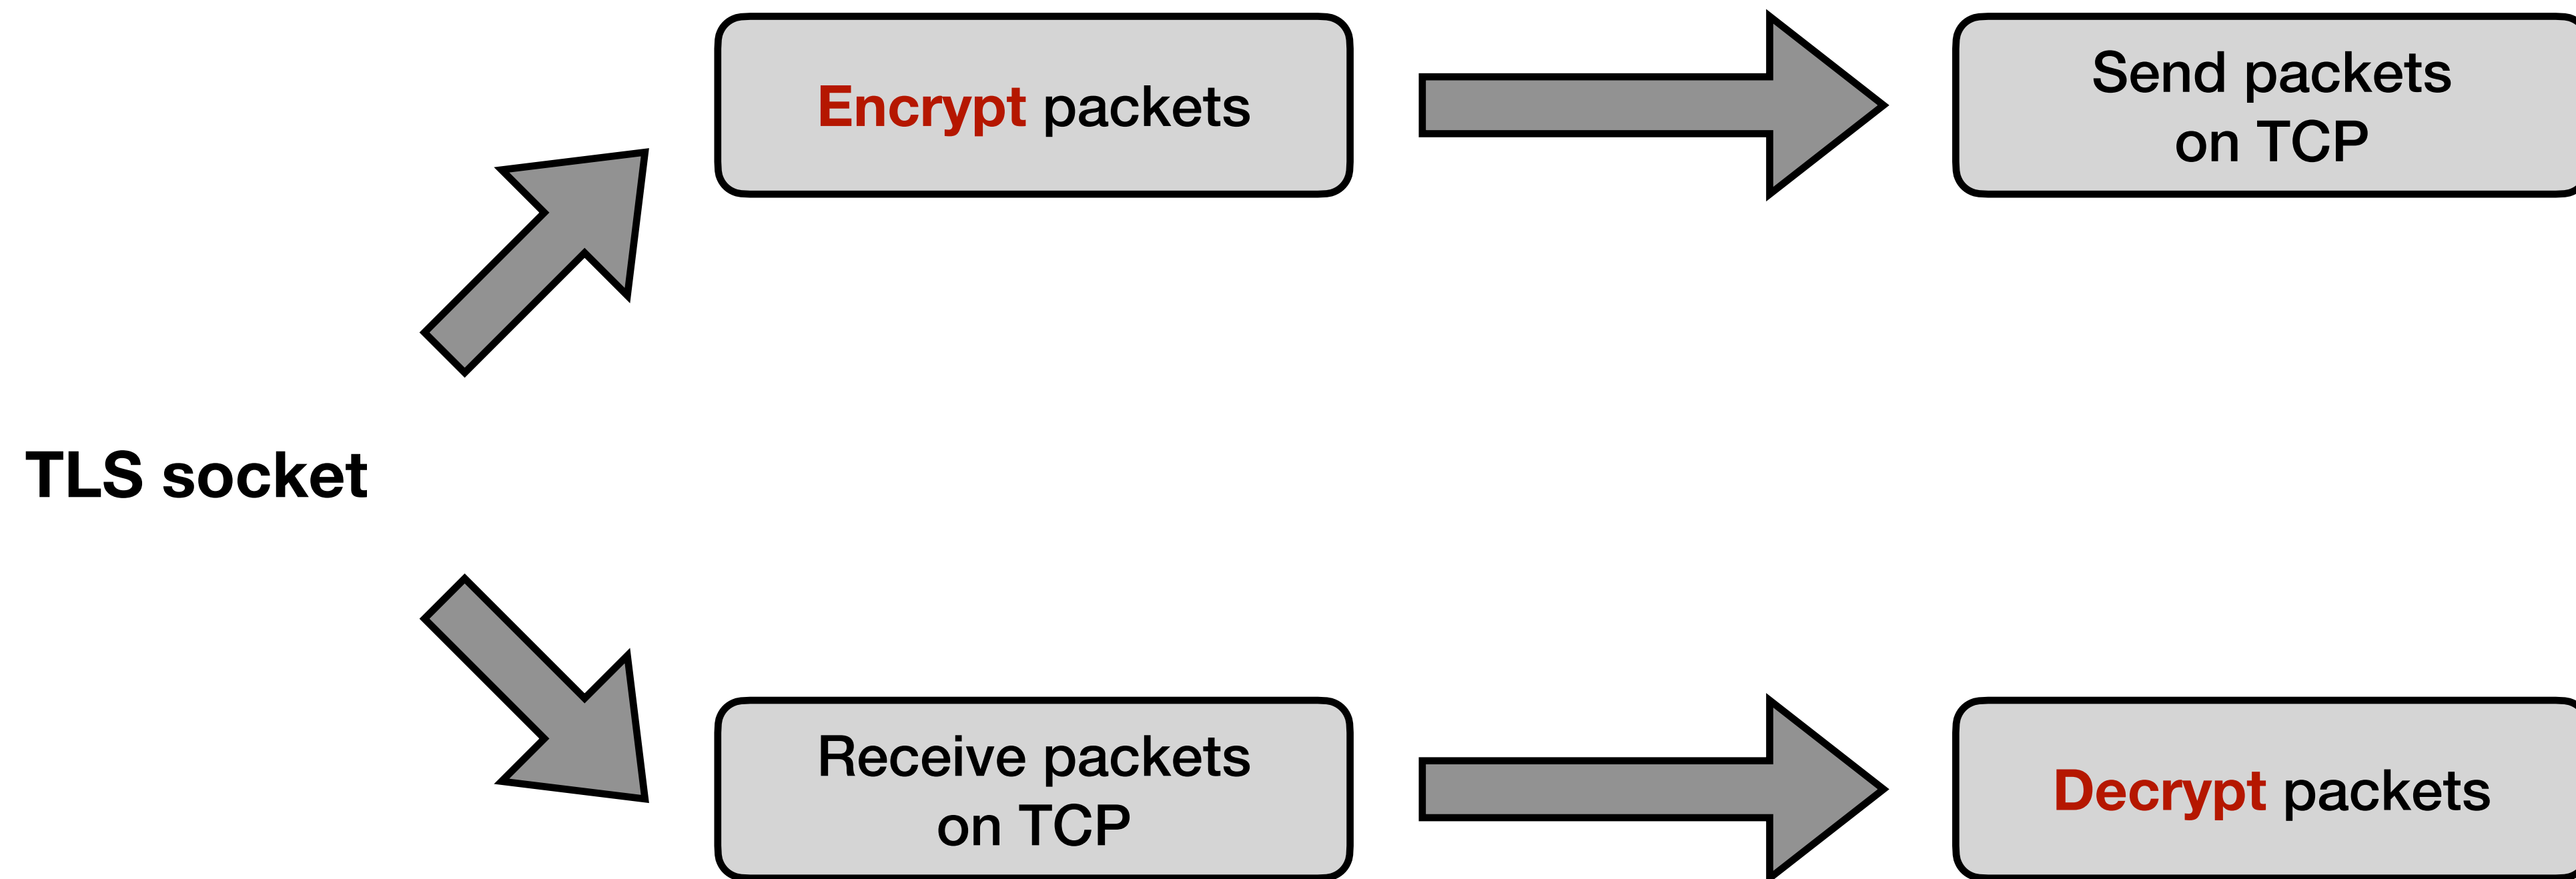
memcpy(crypto_info.key, "0123456789ABCDEF", TLS_CIPHER_AES_GCM_128_KEY_SIZE); // 16
memcpy(crypto_info.iv, "12345678", TLS_CIPHER_AES_GCM_128_IV_SIZE); // 8
memcpy(crypto_info.salt, "SALT", TLS_CIPHER_AES_GCM_128_SALT_SIZE); // 4

setsockopt(sockfd, SOL_TLS, TLS_TX, &crypto_info, sizeof(crypto_info));
setsockopt(sockfd, SOL_TLS, TLS_RX, &crypto_info, sizeof(crypto_info));
```

Overview



Overview



Overview

- Supported TLS algorithms
 - gcm(aes)
 - ccm(aes)
 - gcm(sm4)
 - ...

Overview

- Supported TLS algorithms
 - **gcm**(aes)
 - **ccm**(aes)
 - **gcm**(sm4)
 - ... **Template name**

Overview

- Supported TLS algorithms
 - gcm(**aes**)
 - ccm(**aes**)
 - gcm(**sm4**)
 - ... **Cipher name**

Overview

- **Algorithm**

- Implementation of a specific **cryptographic operation**, such as AES, SHA-256, or HMAC

- **Template**

- Constructing more complex cryptographic transformations by **combining** or **layering** simpler algorithms

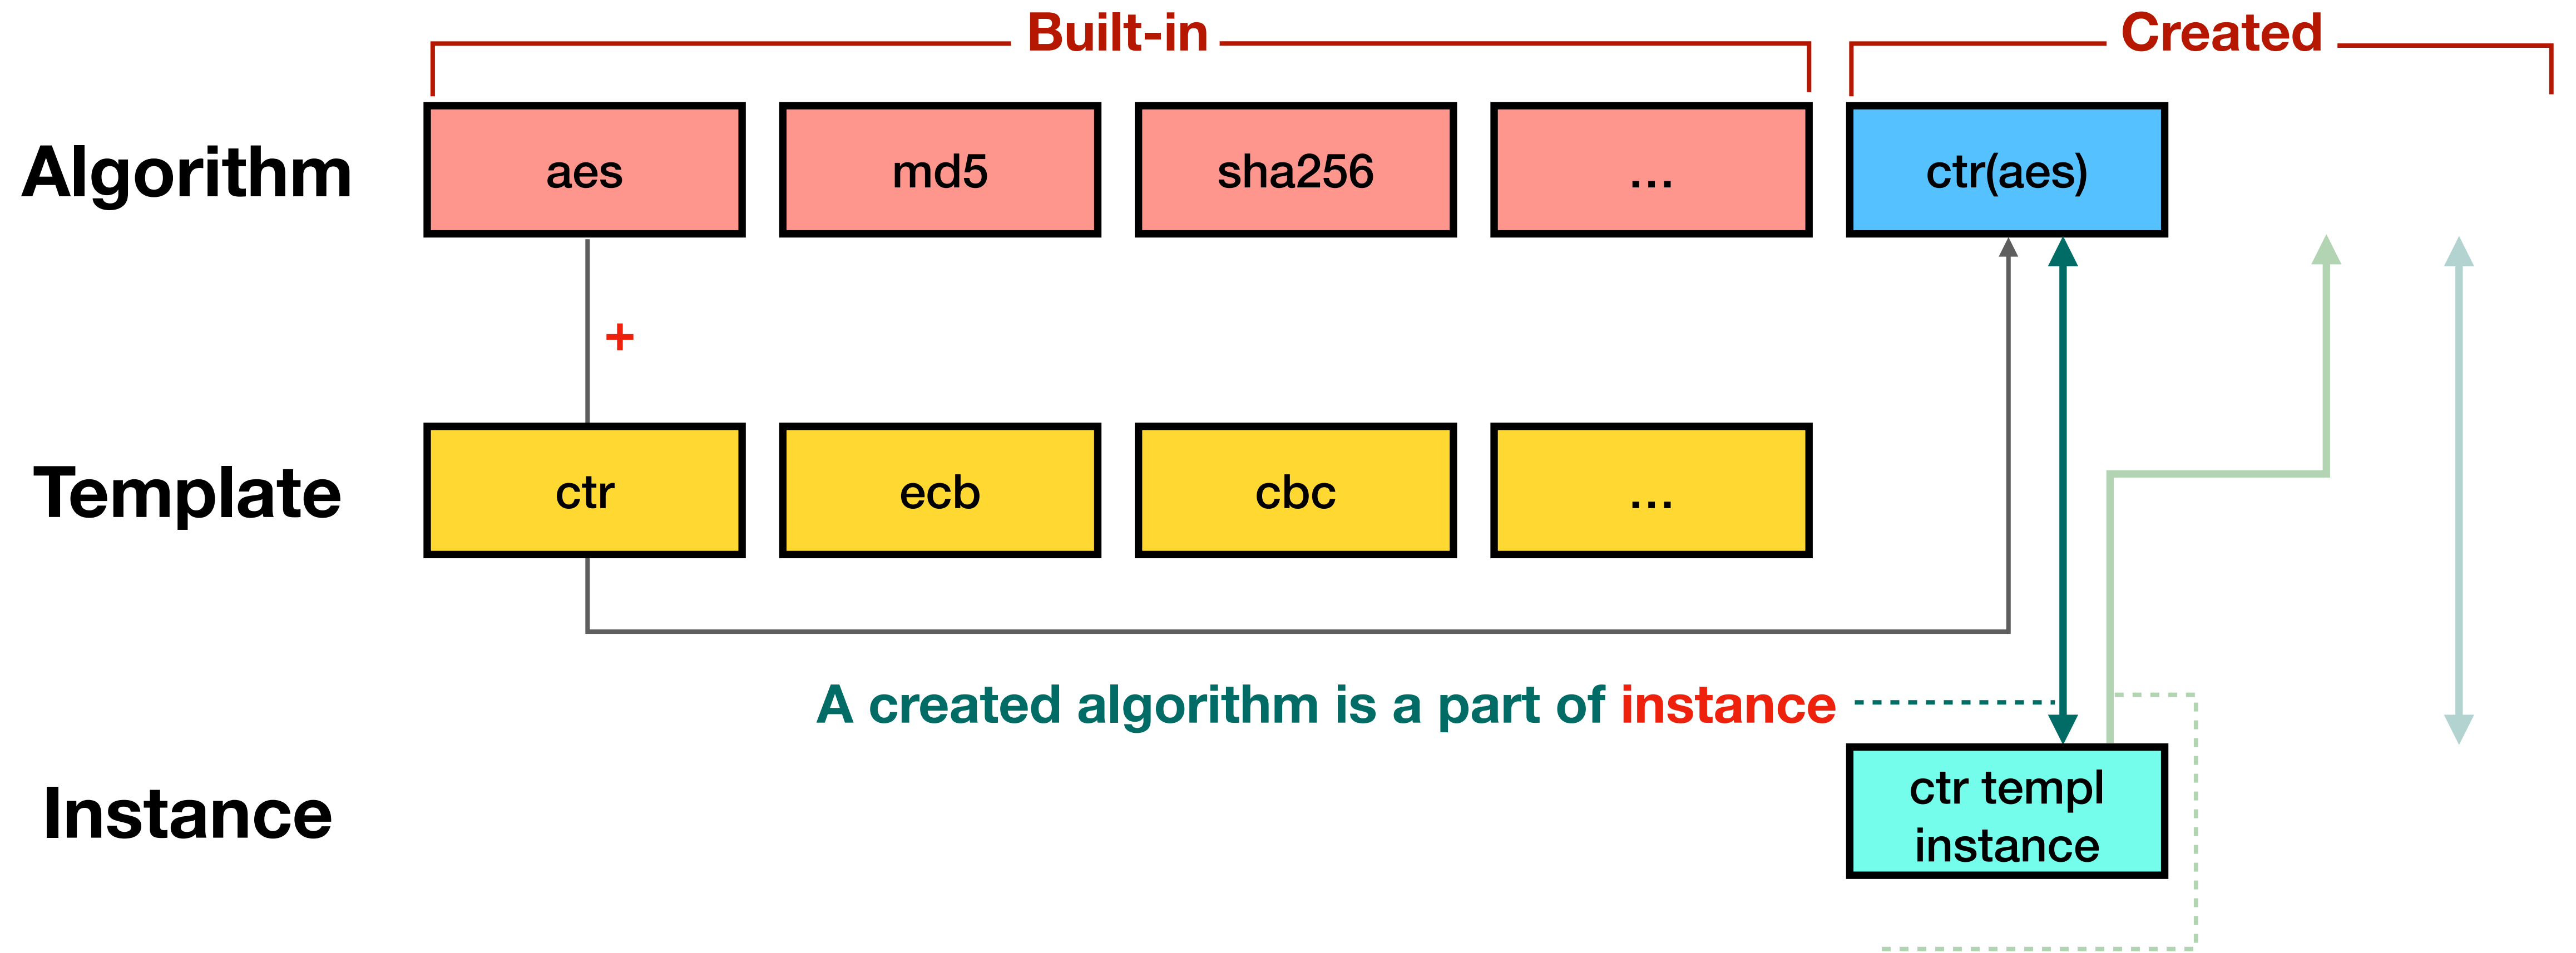
Overview

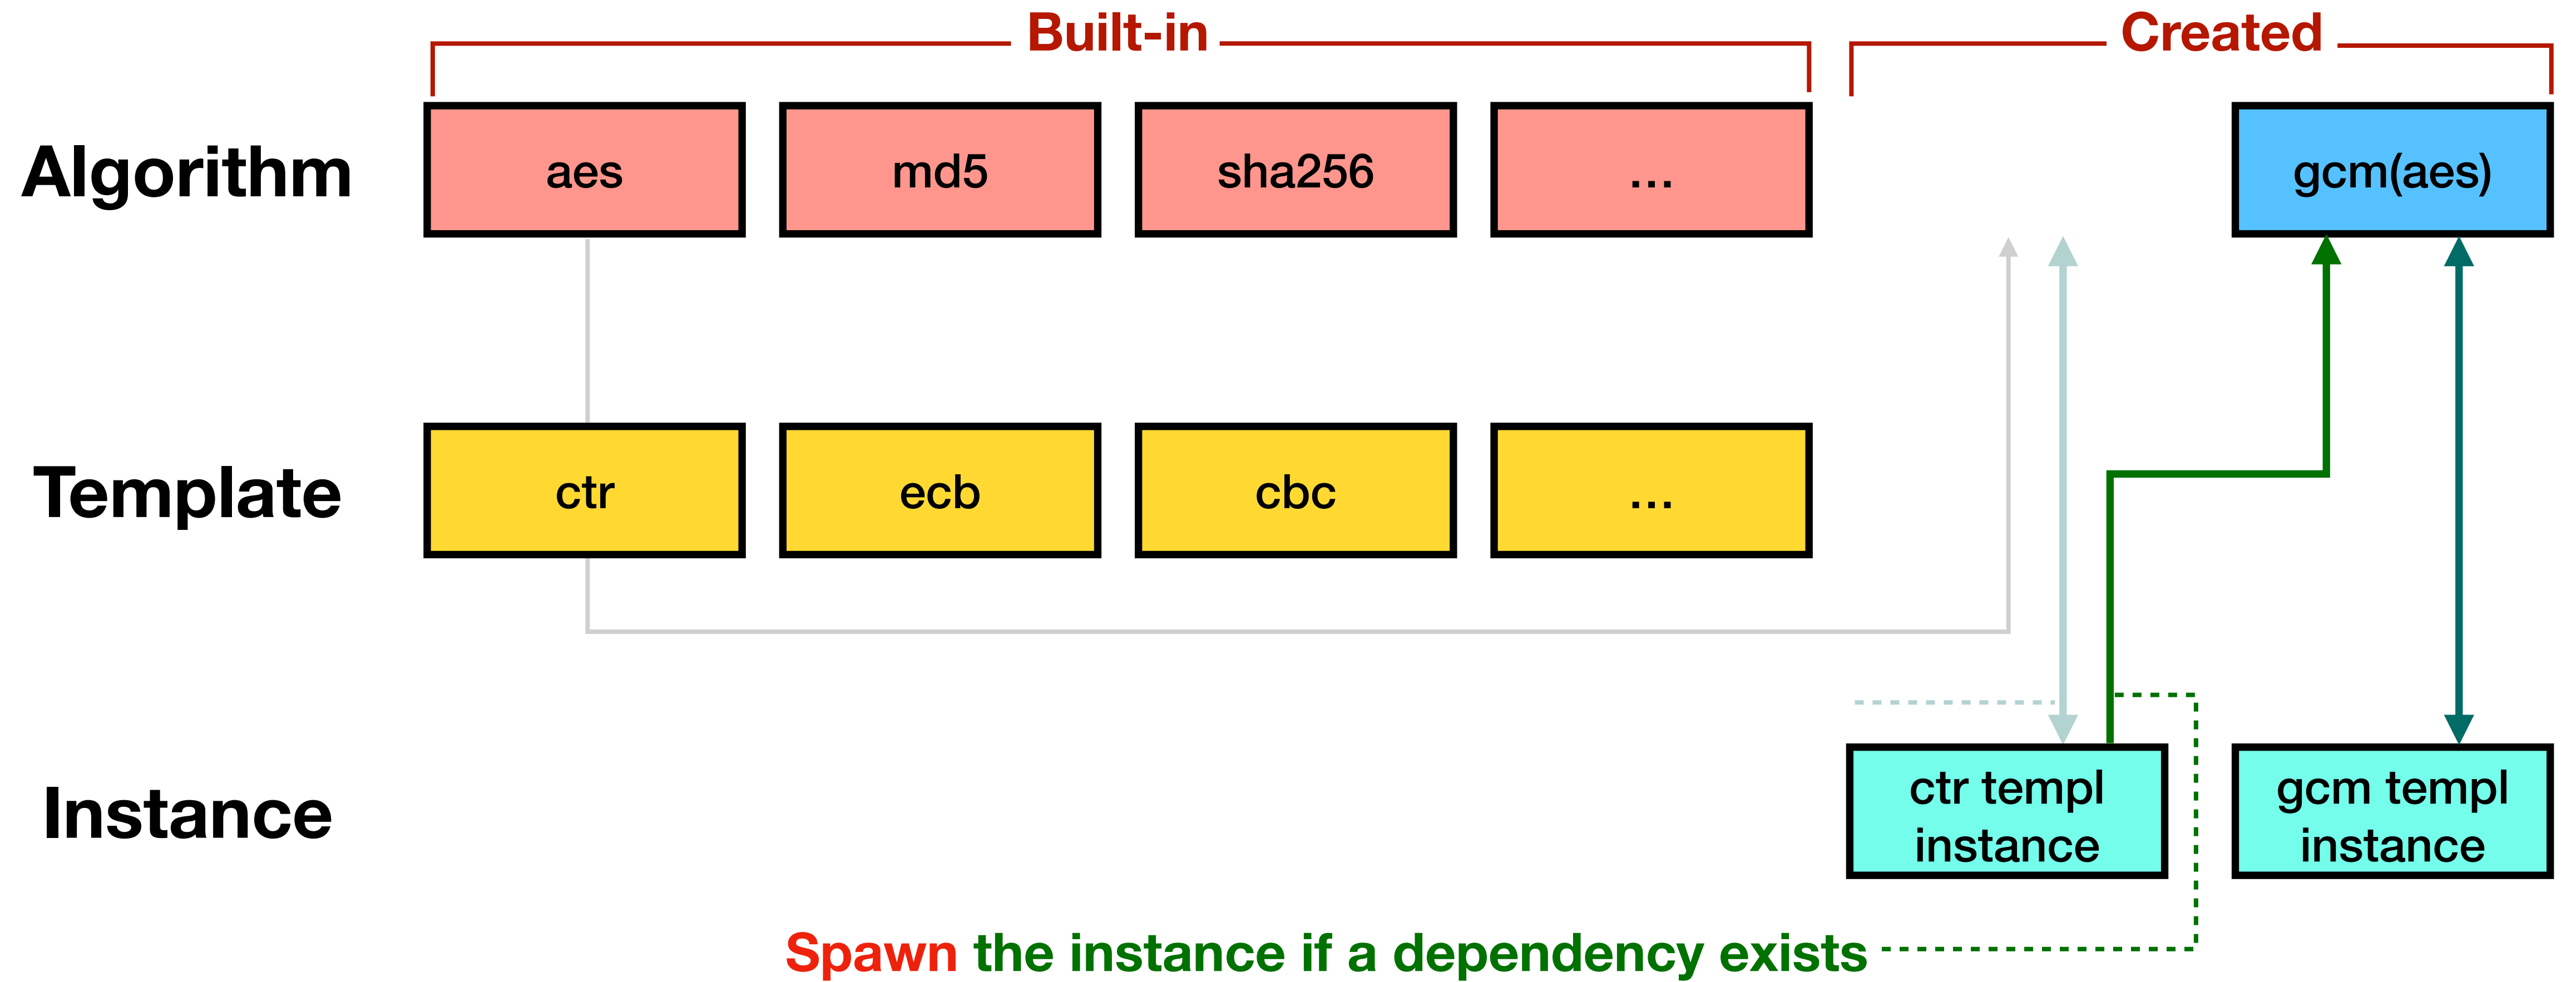
- **Instance**

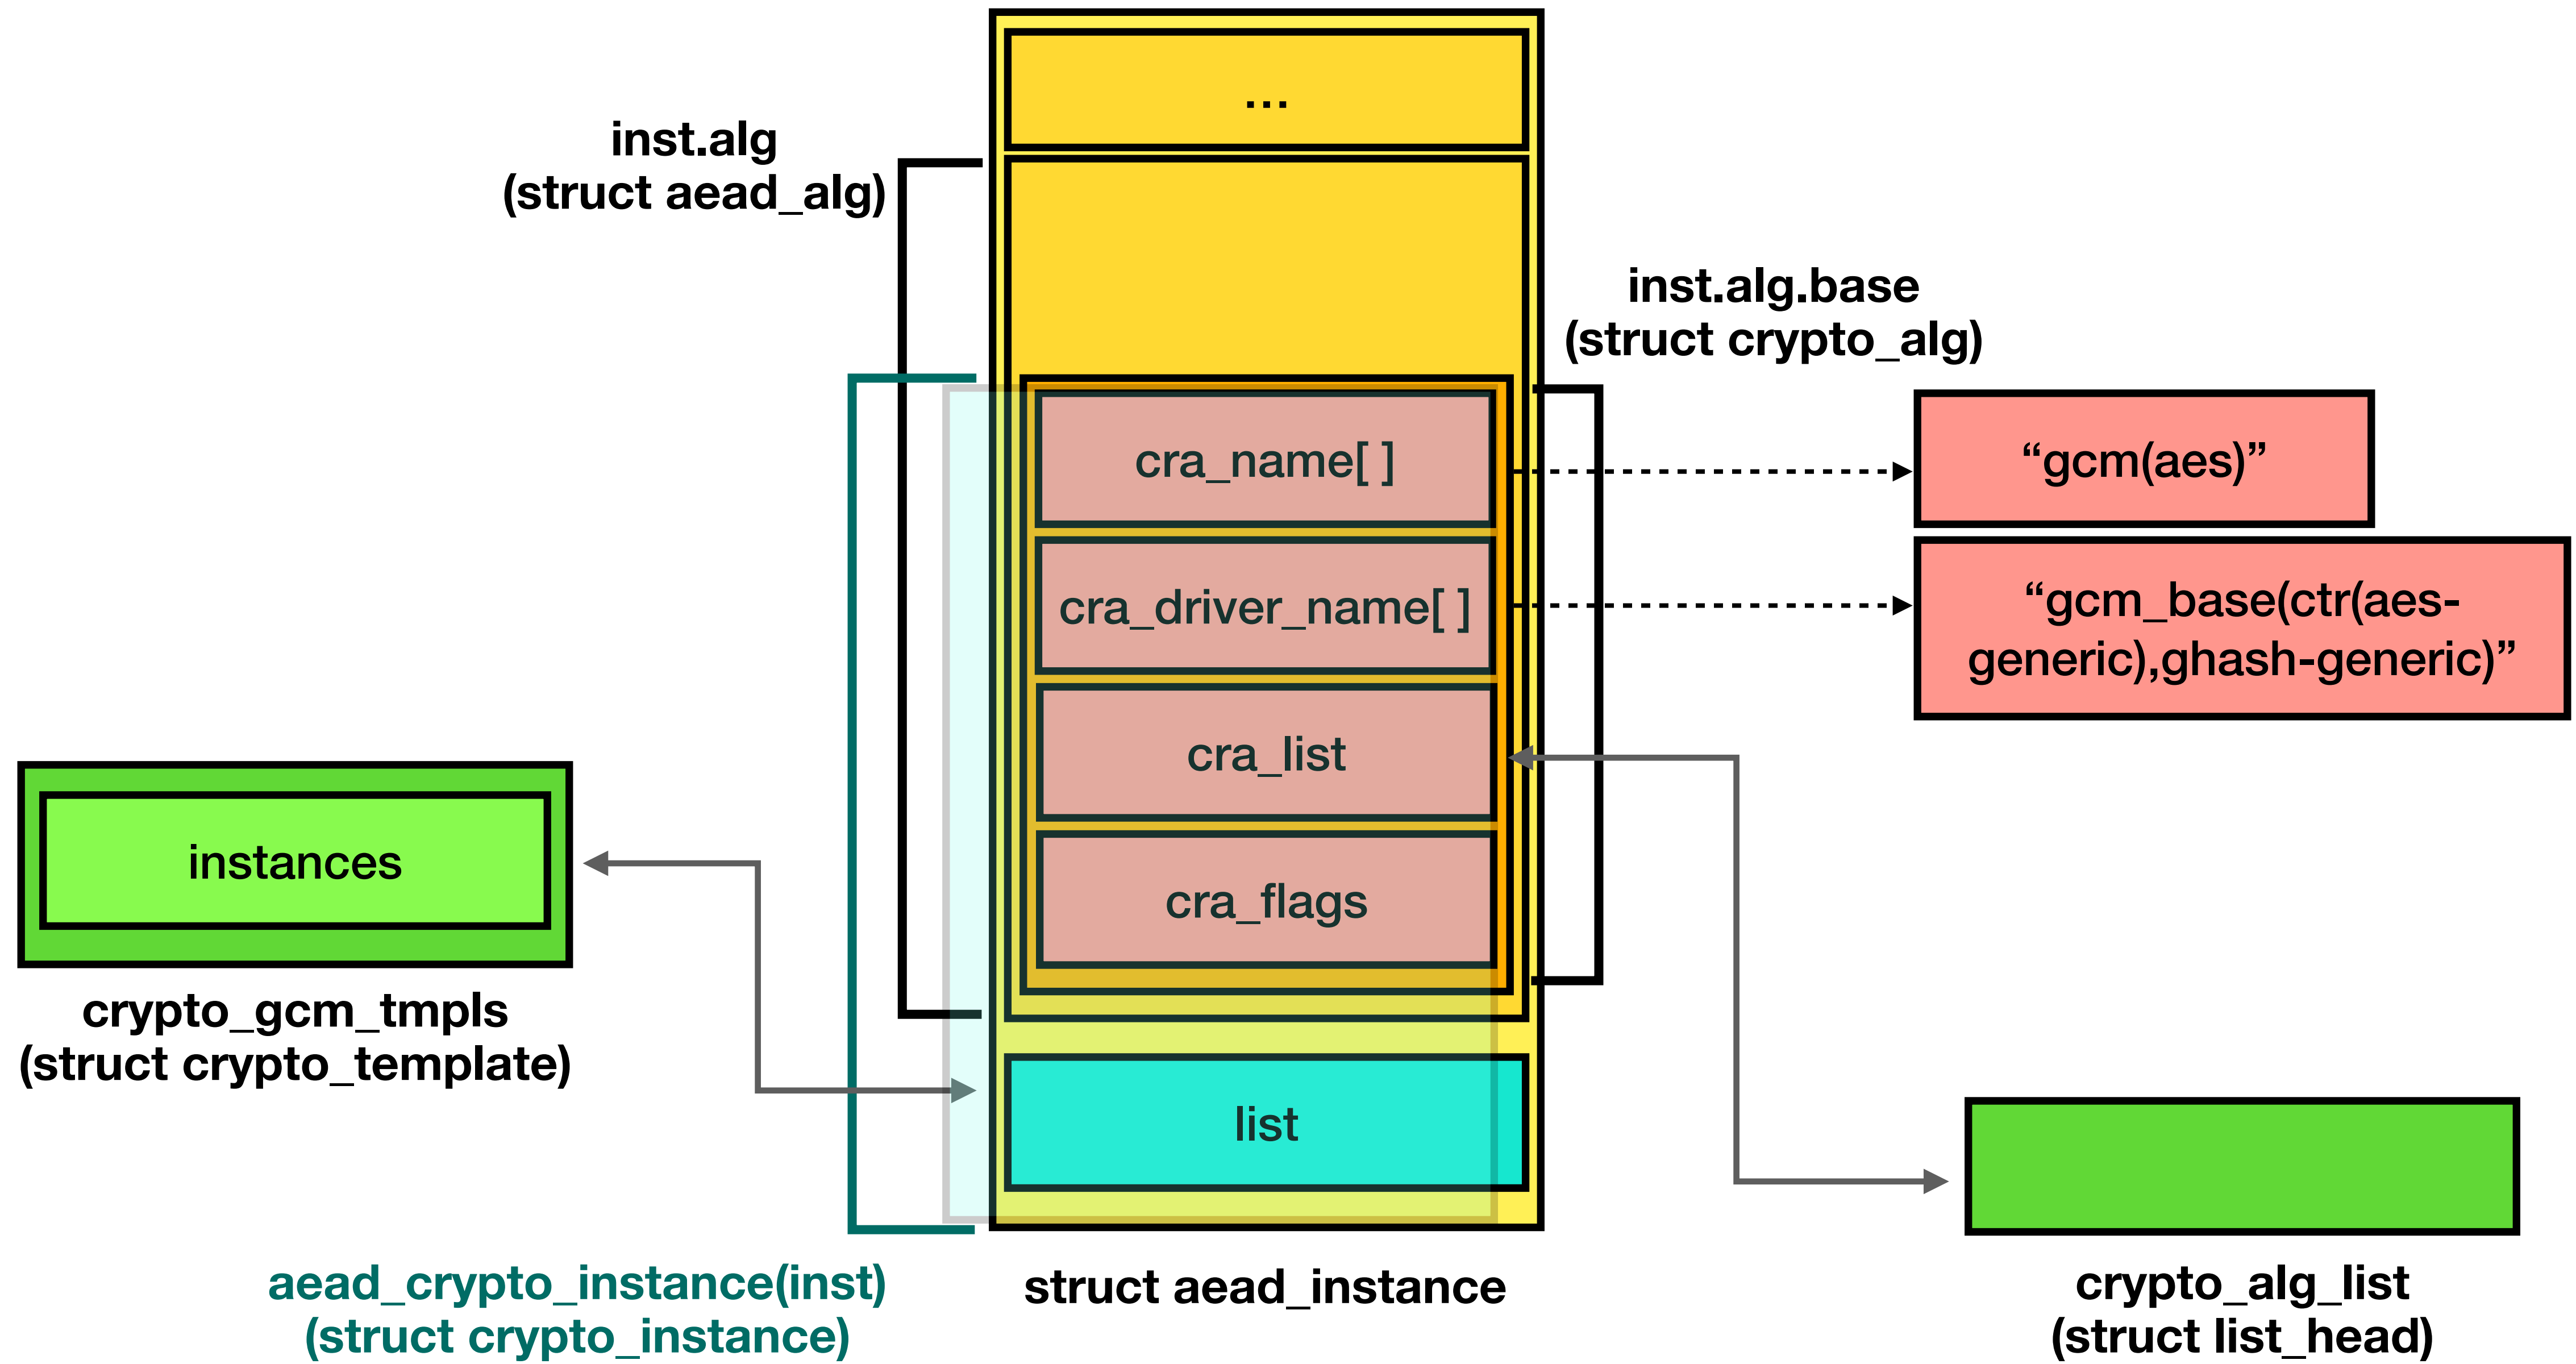
- Instantiation of a cryptographic template, where specific algorithms and parameters **have been configured**

- **Spawn**

- Create a linkage or **dependency** between cryptographic instances and algorithms





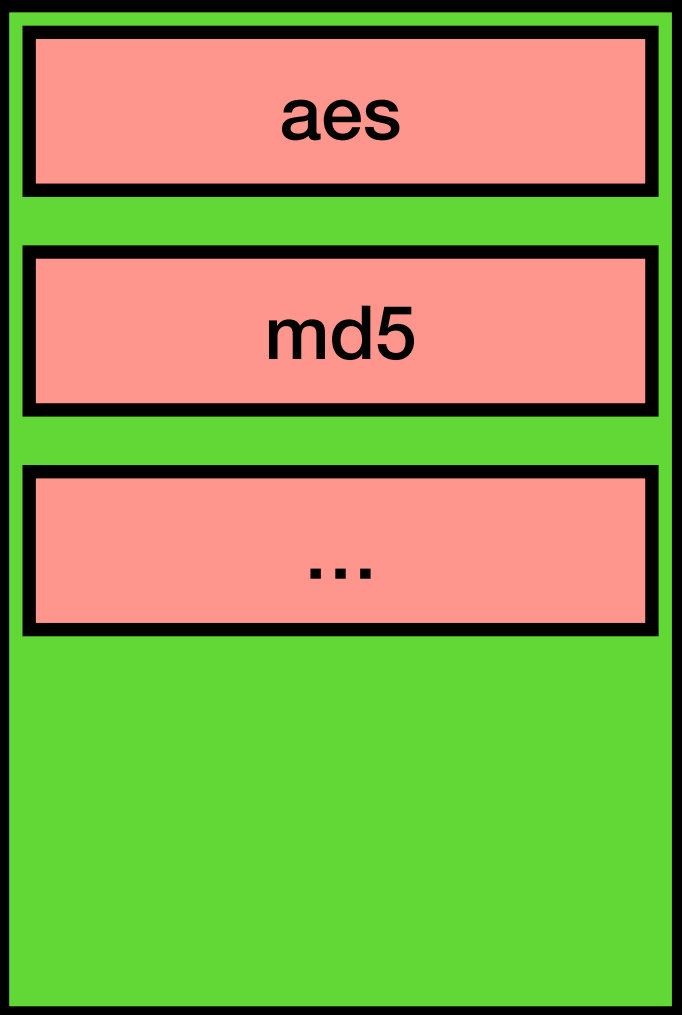


Overview

- For example, if we configure “gcm(aes)” as the crypto algorithm of TX...

Thread-A

Find "gcm(aes)"



Global variable

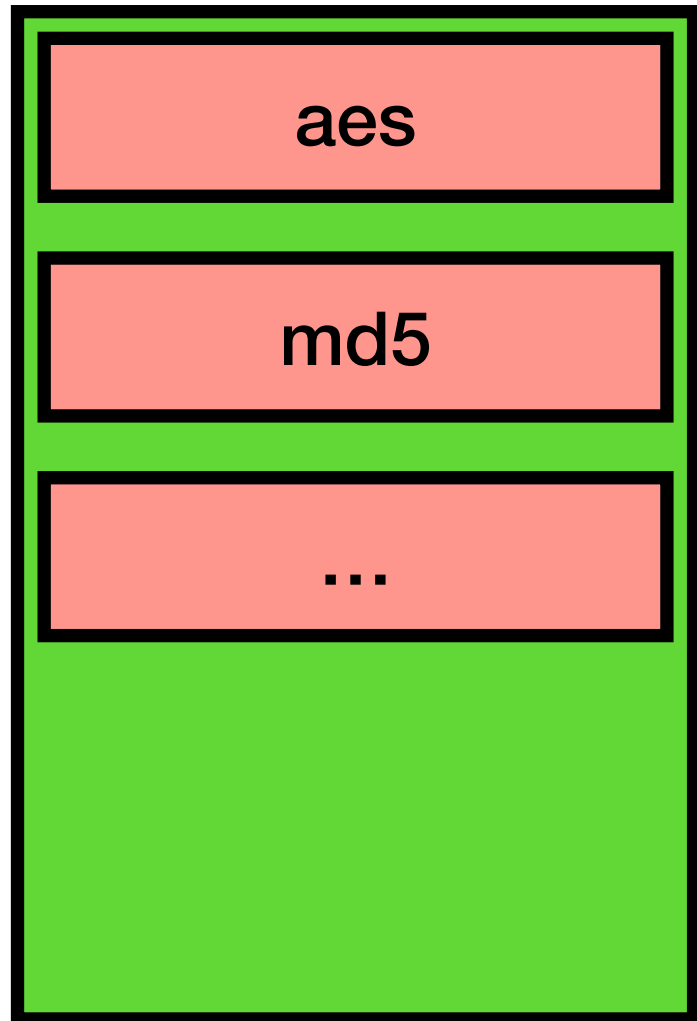
Thread-B

Thread-A

Thread-B

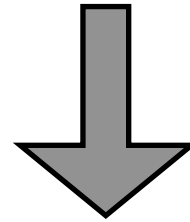
Not found

Find "gcm(aes)"



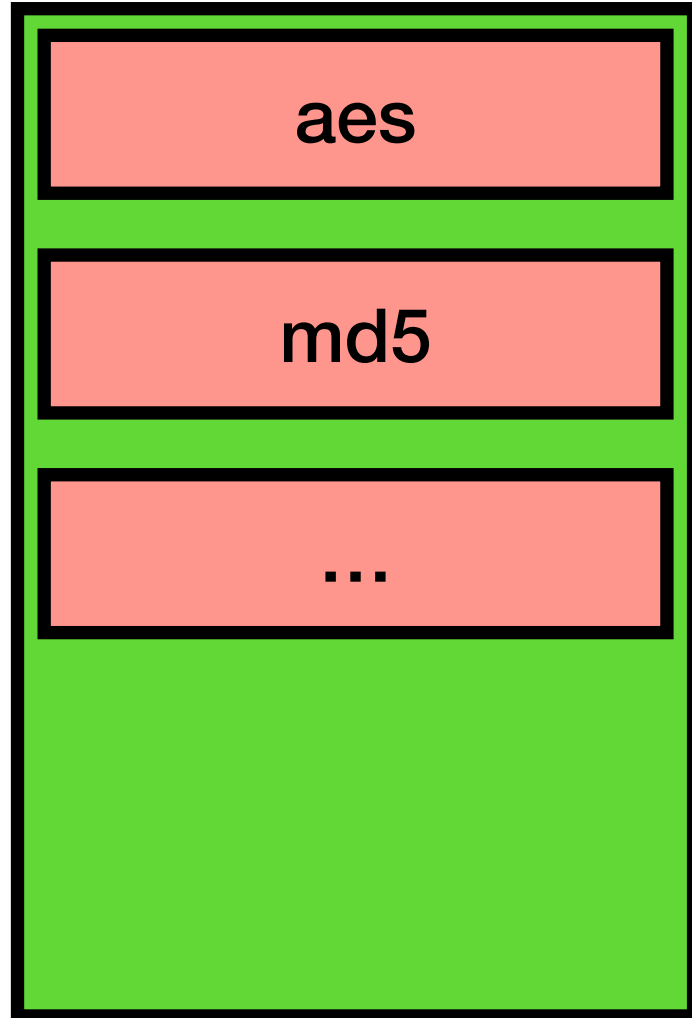
Thread-A

Find "gcm(aes)"



Setup probe

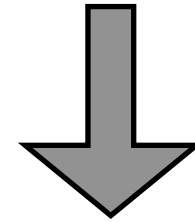
Thread-B



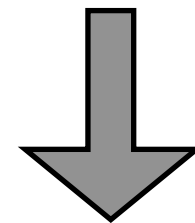
Global variable

Thread-A

Find "gcm(aes)"

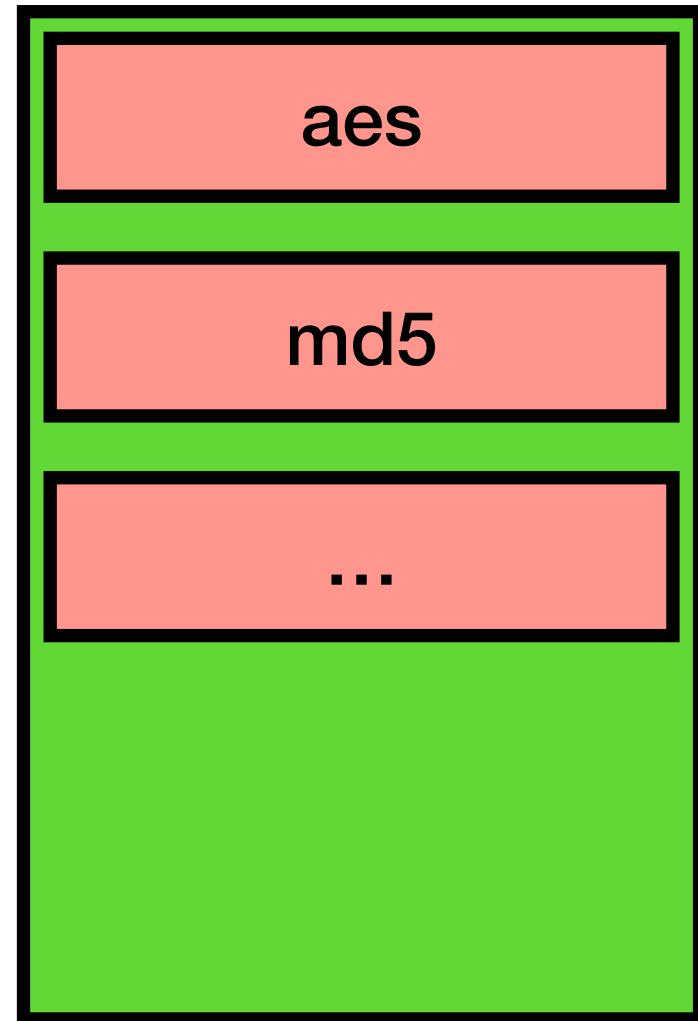


Setup probe



Dispatch probing

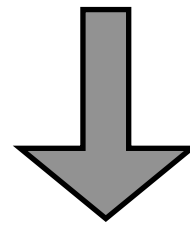
Thread-B



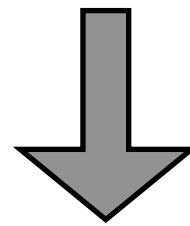
Global variable

Thread-A

Find "gcm(aes)"



Setup probe



Dispatch probing

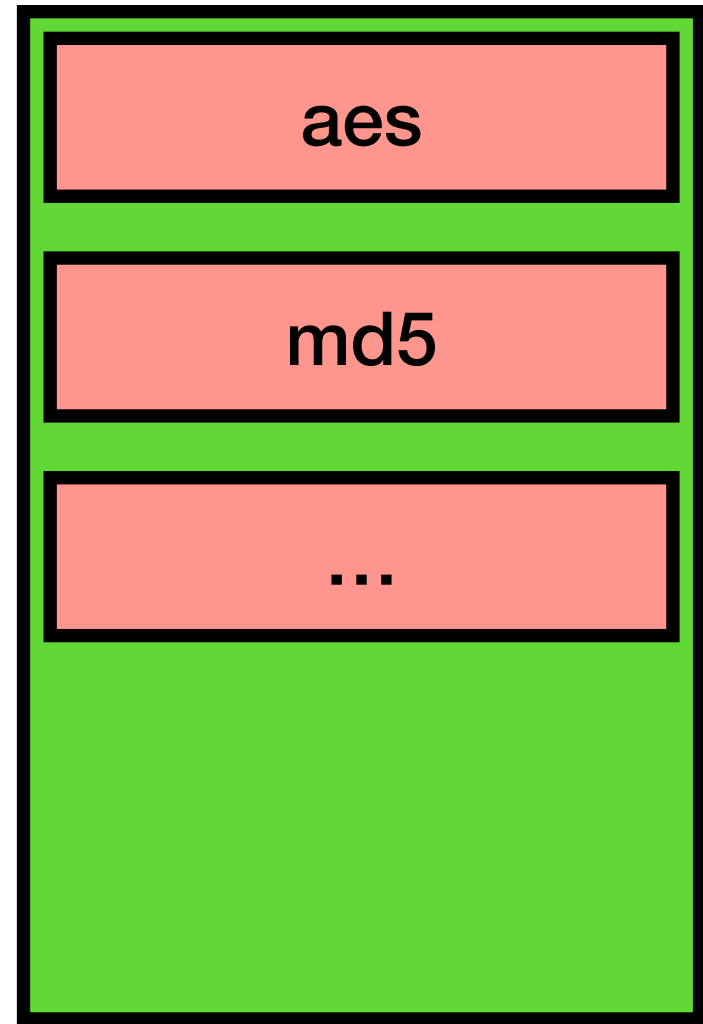
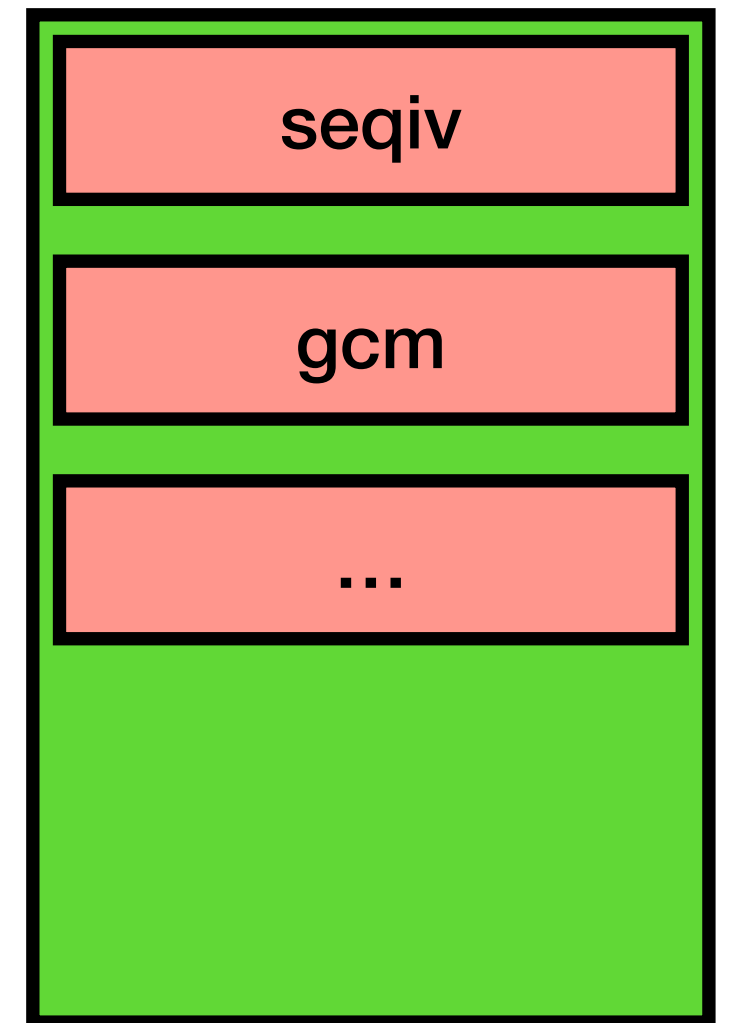
- 1. Template name = "gcm"
- 2. Cipher name = "aes"



Thread-B

"cryptomgr_probe"

Find template

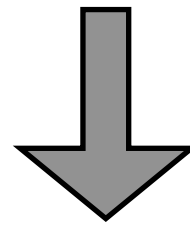


Global variable

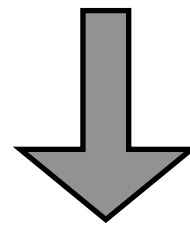
Global variable

Thread-A

Find "gcm(aes)"



Setup probe



Dispatch probing

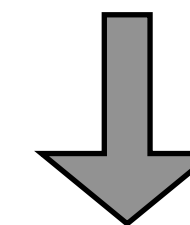
- 1. Template name = "gcm"
- 2. Cipher name = "aes"



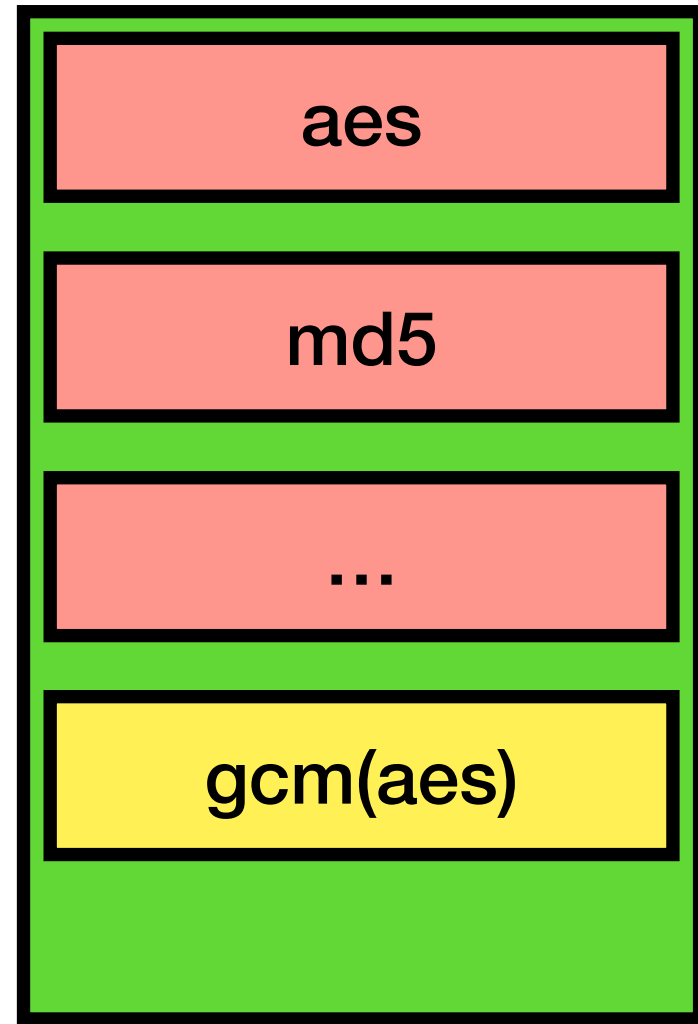
Thread-B

"cryptomgr_probe"

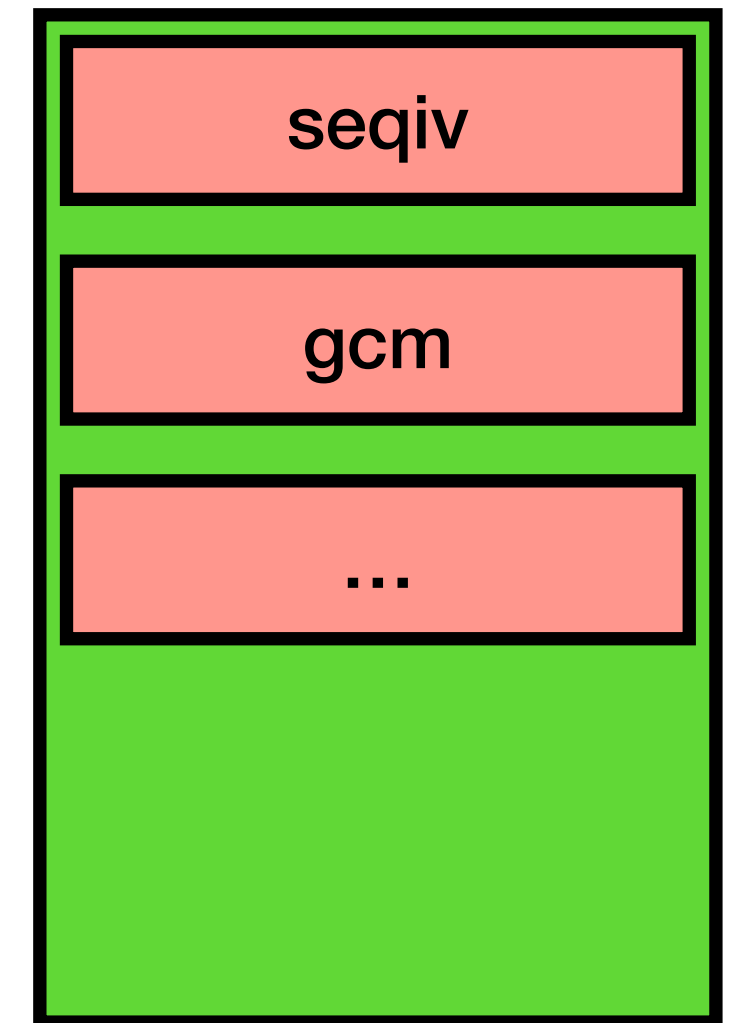
Find template



Create & initialize instance



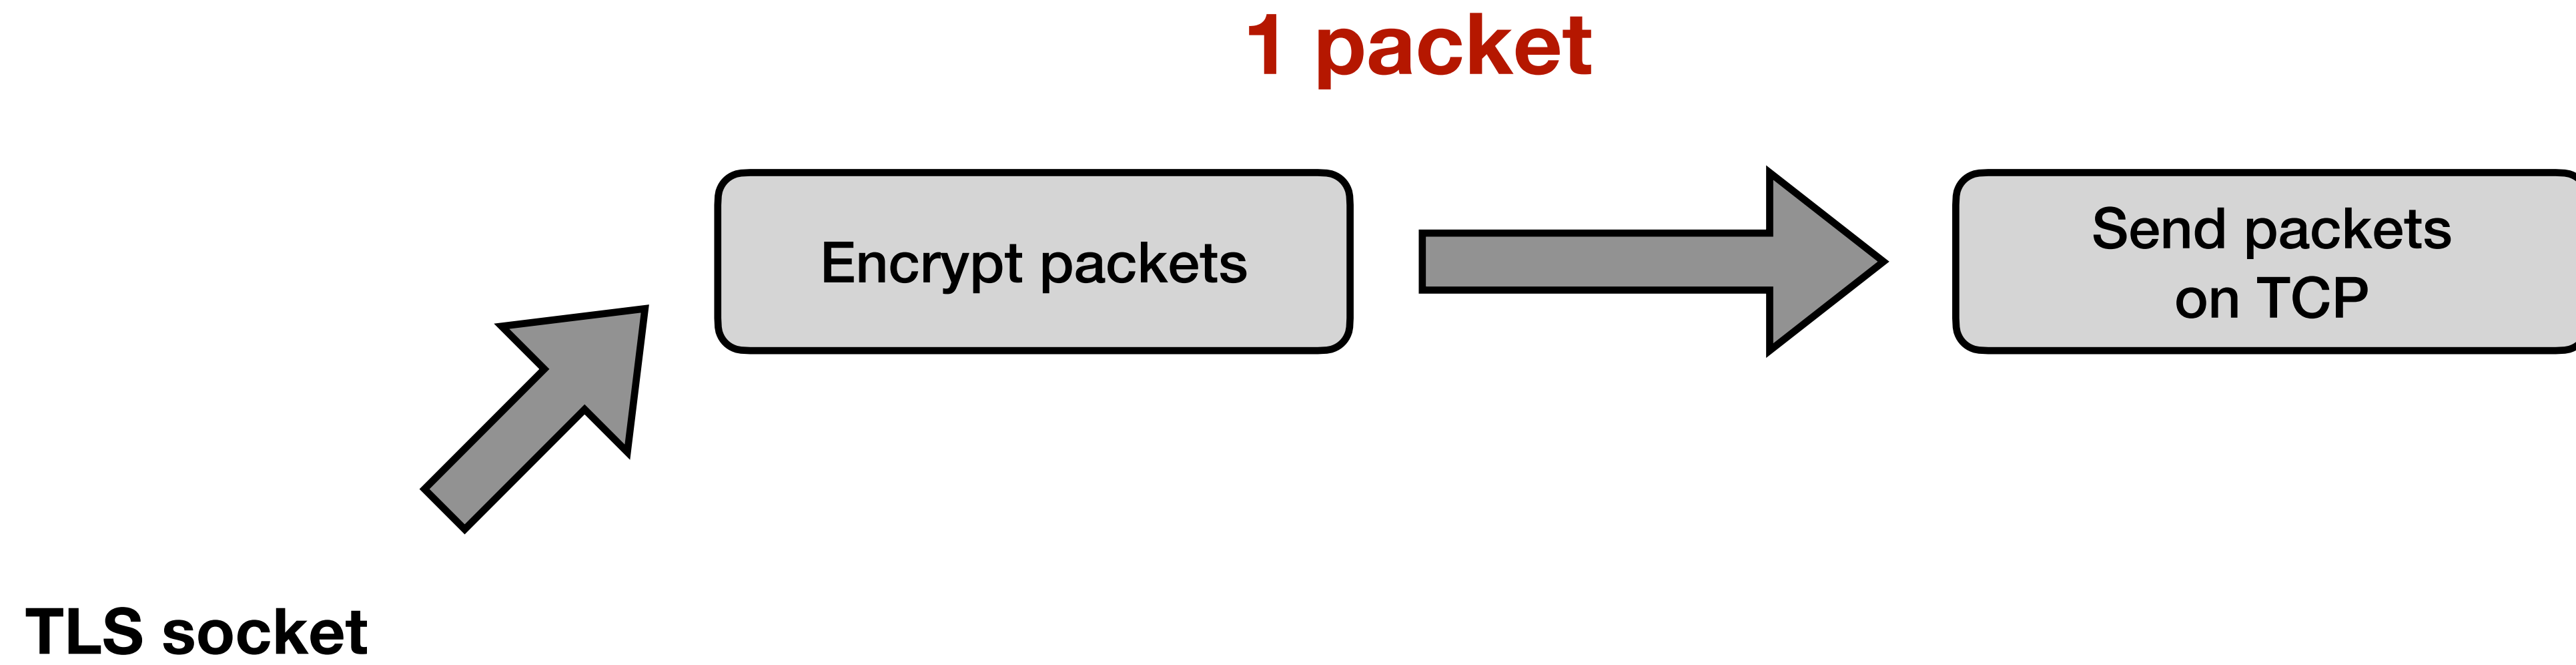
Global variable



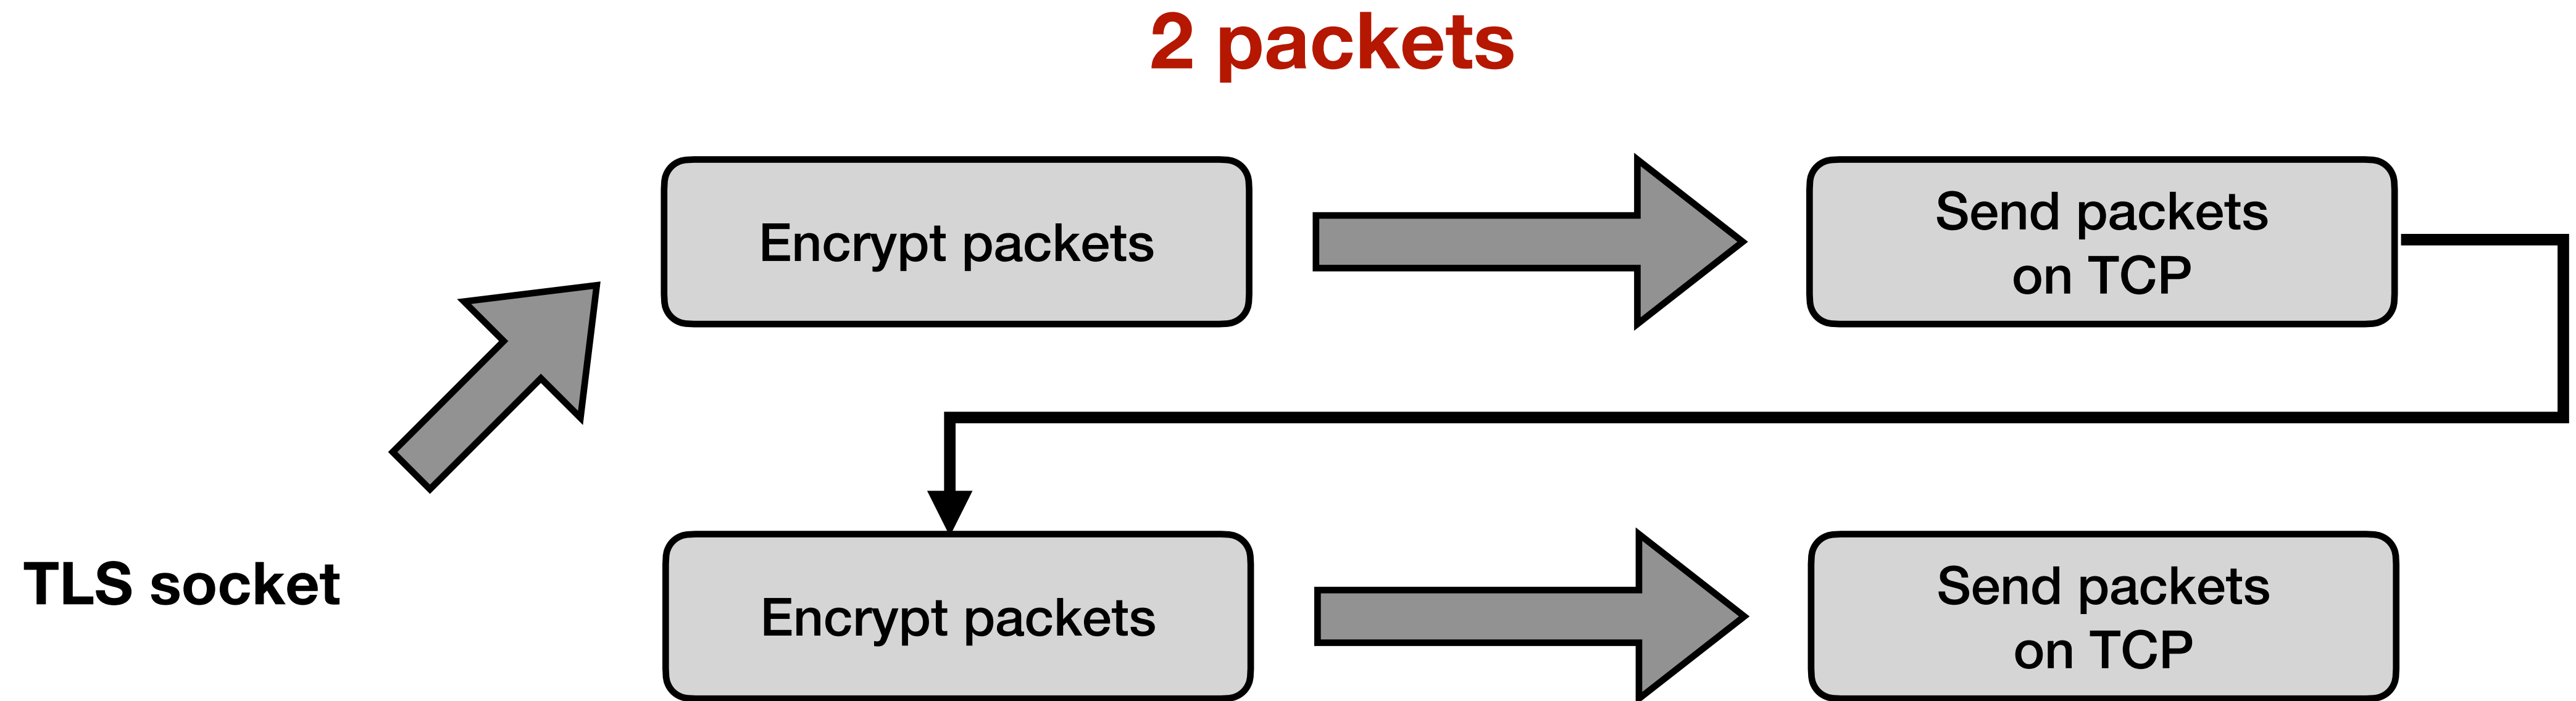
Global variable



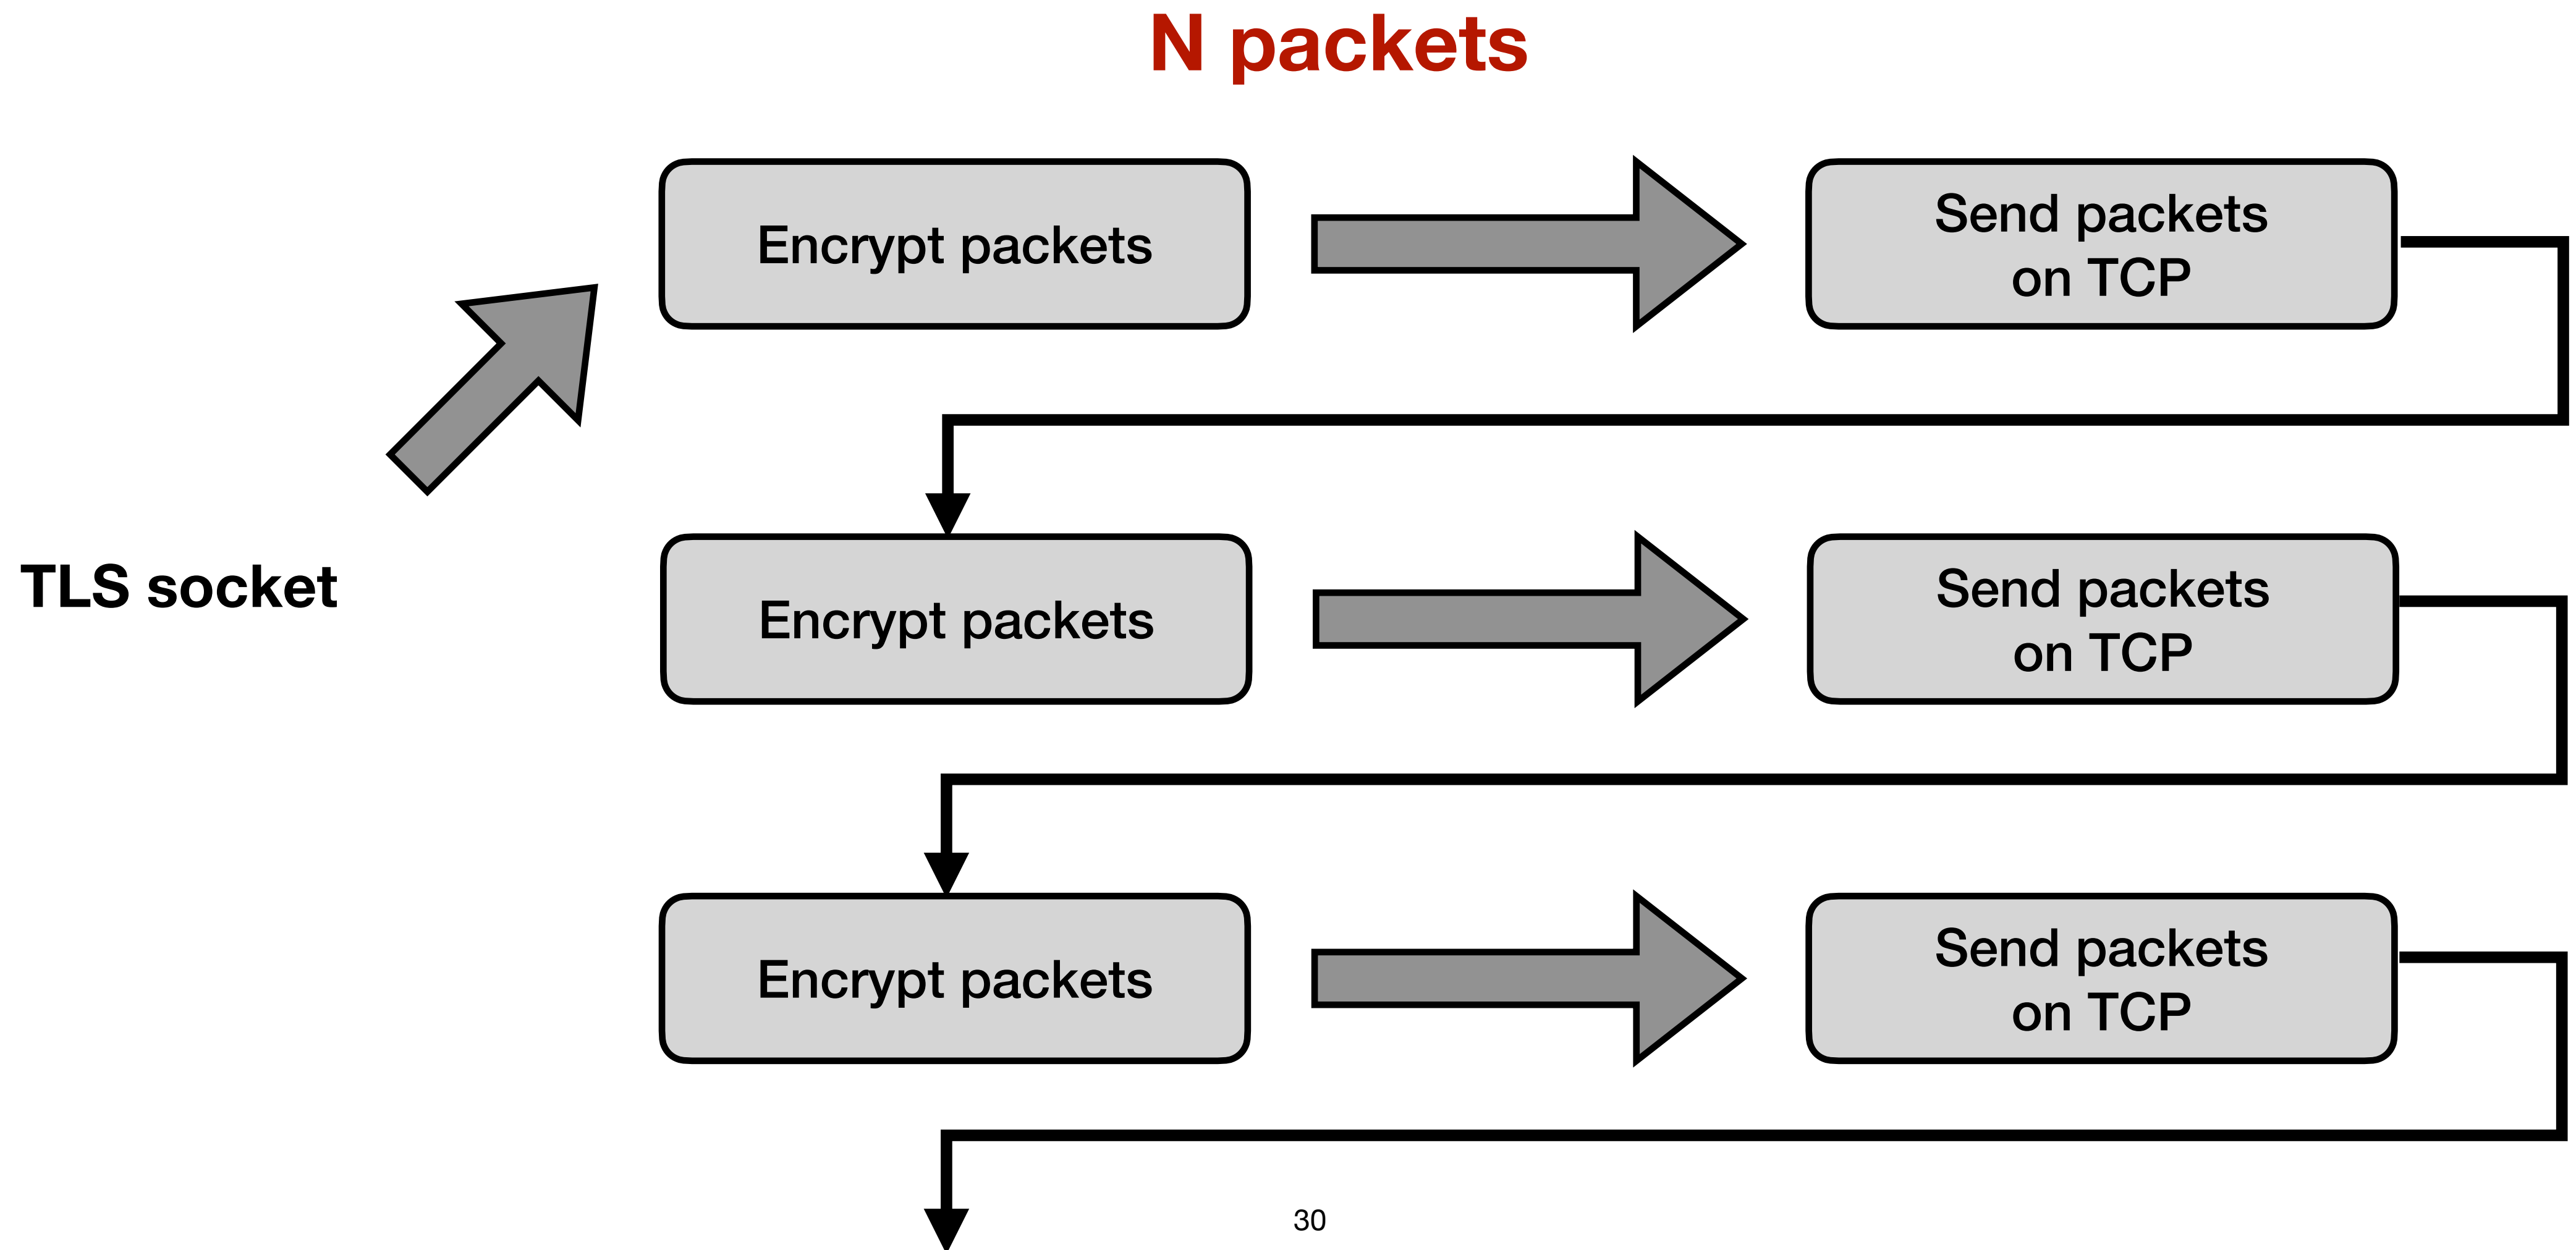
Overview



Overview



Overview

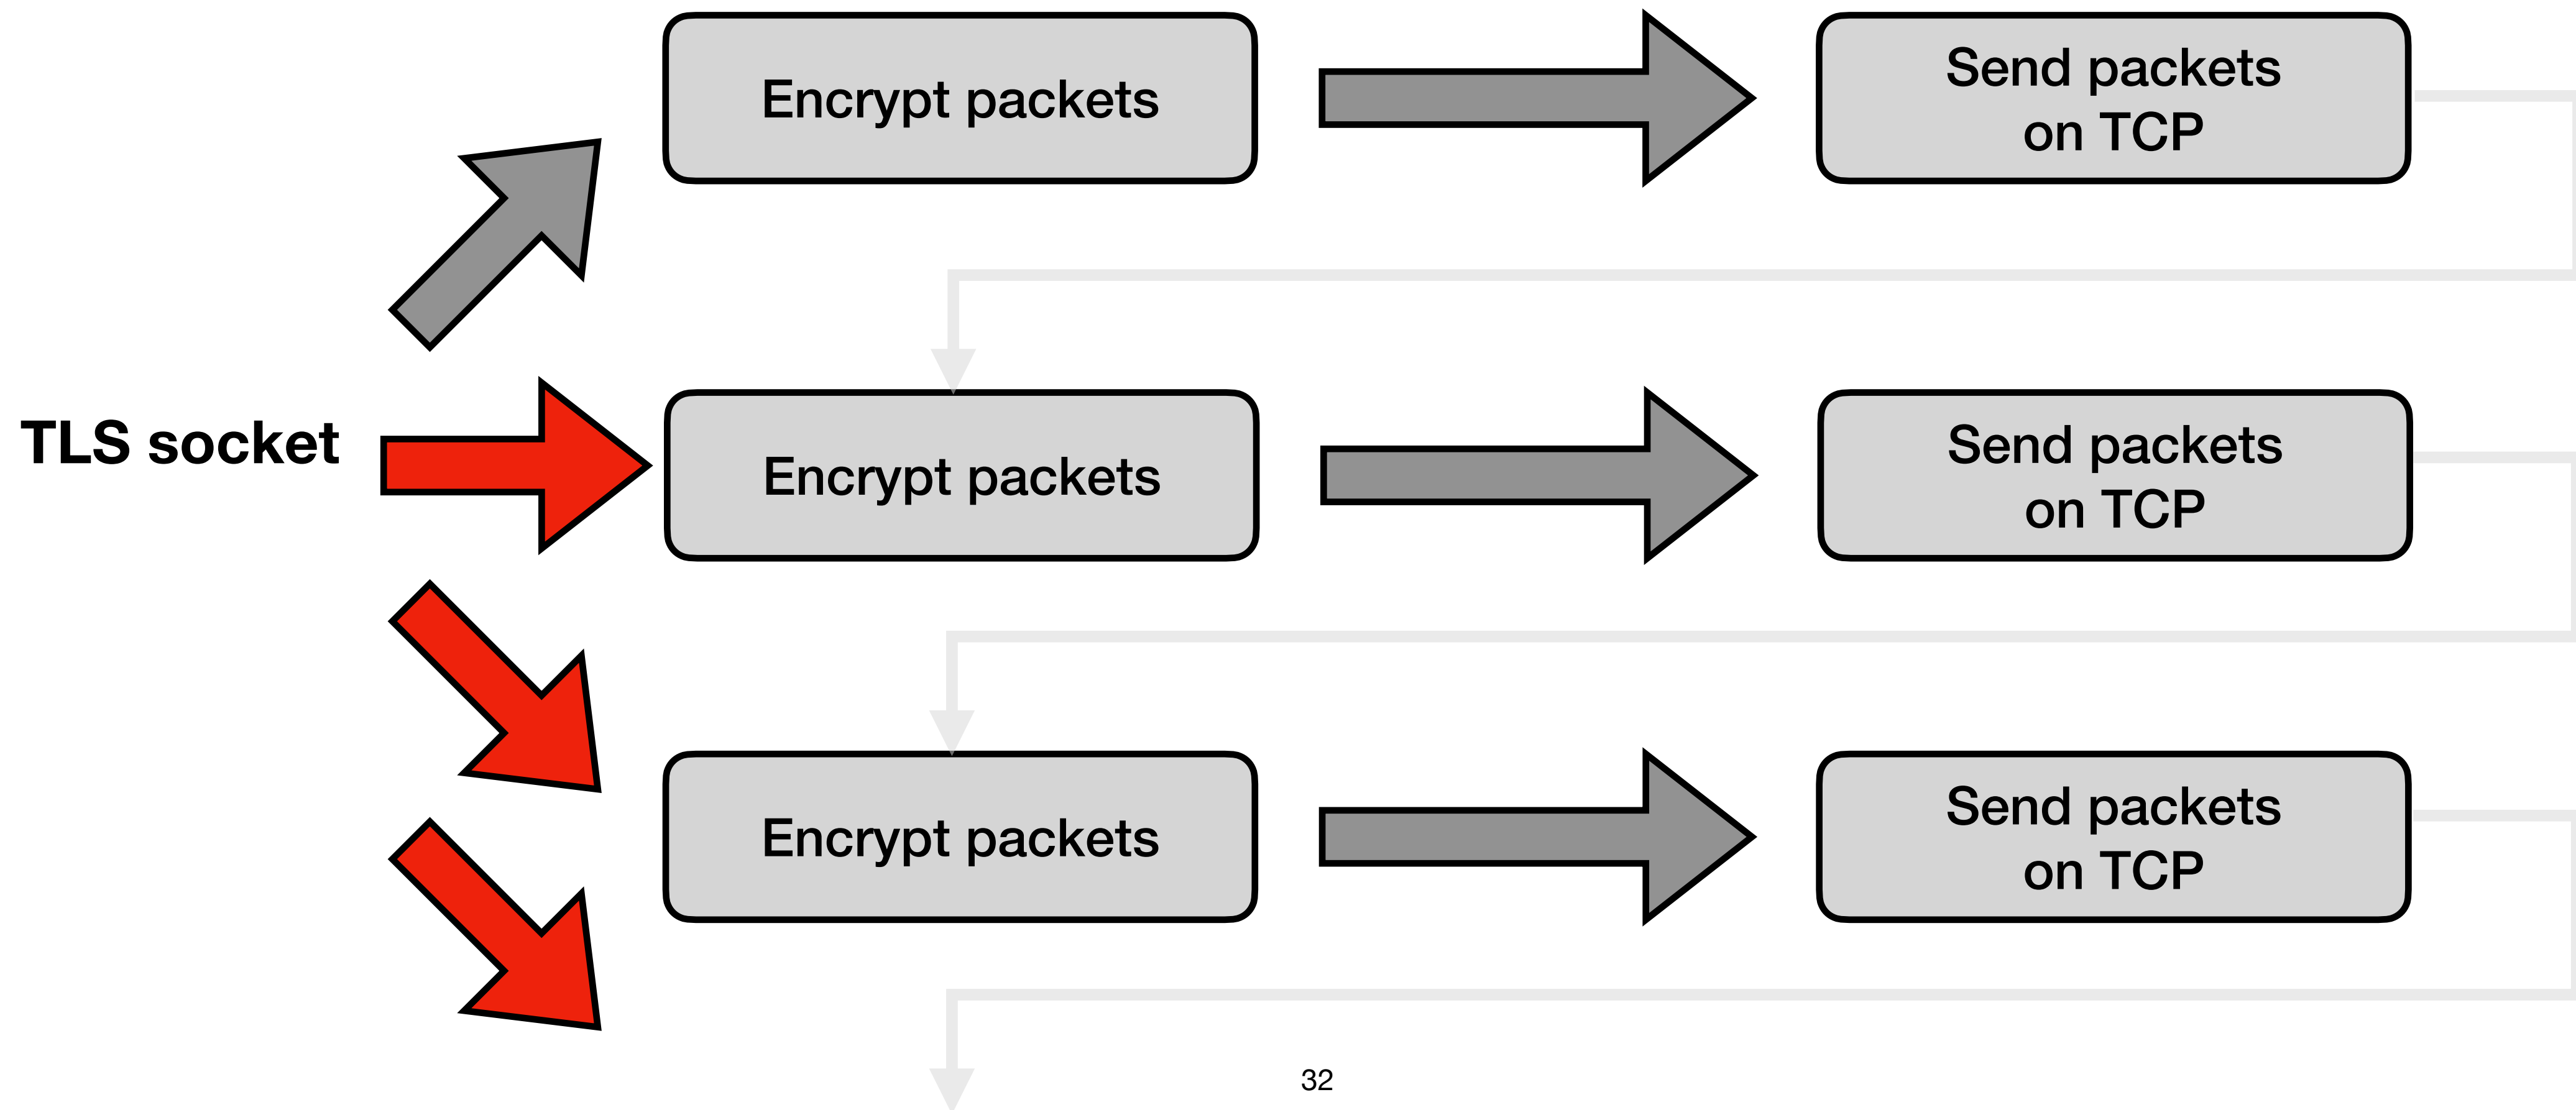


Overview



Overview

N packets (asynchronous mode)



Overview

Vendor specific drivers

```
✓ C aes_cbc.c drivers/crypto/vmx 1  
  CRYPTO_ALG_ASYNC);  
✓ C aes_ctr.c drivers/crypto/vmx 1  
  CRYPTO_ALG_ASYNC);  
✓ C aes_xts.c drivers/crypto/vmx 1  
  CRYPTO_ALG_ASYNC);  
✓ C zynqmp-aes-gcm.c drivers/crypto/xilinx 1  
  CRYPTO_ALG_ASYNC |  
✓ C dm-verity-target.c drivers/md 1  
  v->use_tasklet ? CRYPTO_ALG_ASYNC : 0);  
✓ C ppp_mppe.c drivers/net/ppp 1  
  ...crypto_has_ahash("sha1", 0, CRYPTO_ALG_ASYNC))  
✓ C tcp.c drivers/nvme/host 1  
  ...crypto_alloc_ahash("crc32c", 0, CRYPTO_ALG_ASYNC);  
✓ C tcp.c drivers/nvme/target 1  
  ...crypto_alloc_ahash("crc32c", 0, CRYPTO_ALG_ASYNC);  
✓ C iscsi_tcp.c drivers/scsi 1  
  ...crypto_alloc_ahash("crc32c", 0, CRYPTO_ALG_ASYNC);
```



Overview

- **Cryptd**
 - Enabled when the CONFIG_CRYPTD_CRYPTD compile option is set
 - A **crypto daemon** which converts an arbitrary synchronous crypto algorithm into an **asynchronous algorithm** that runs in a **kthread**



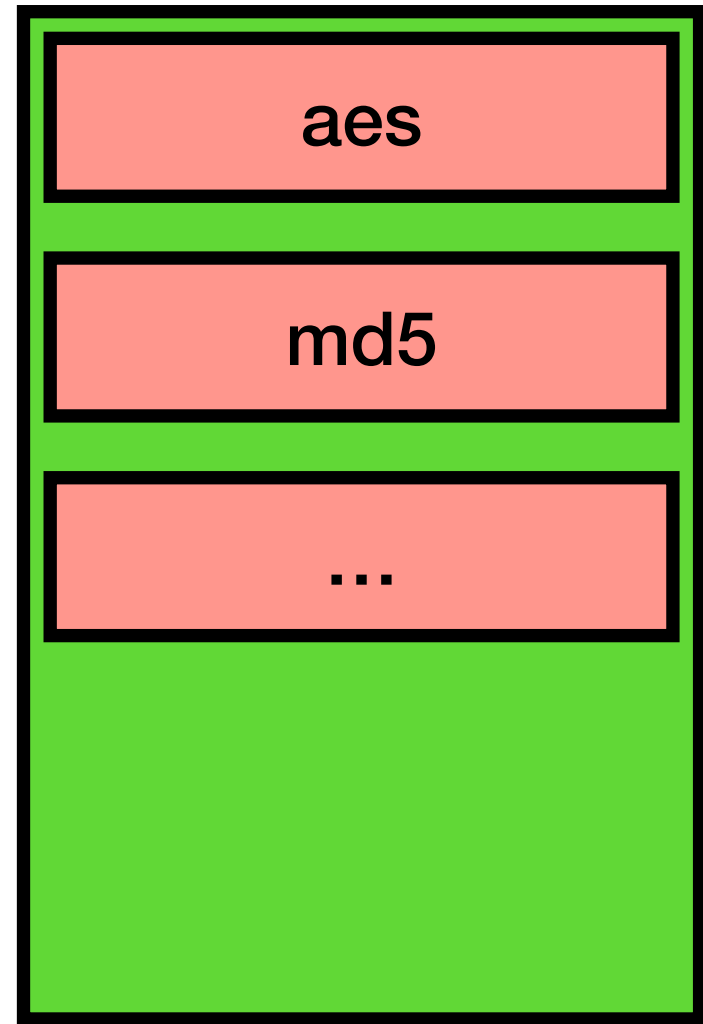
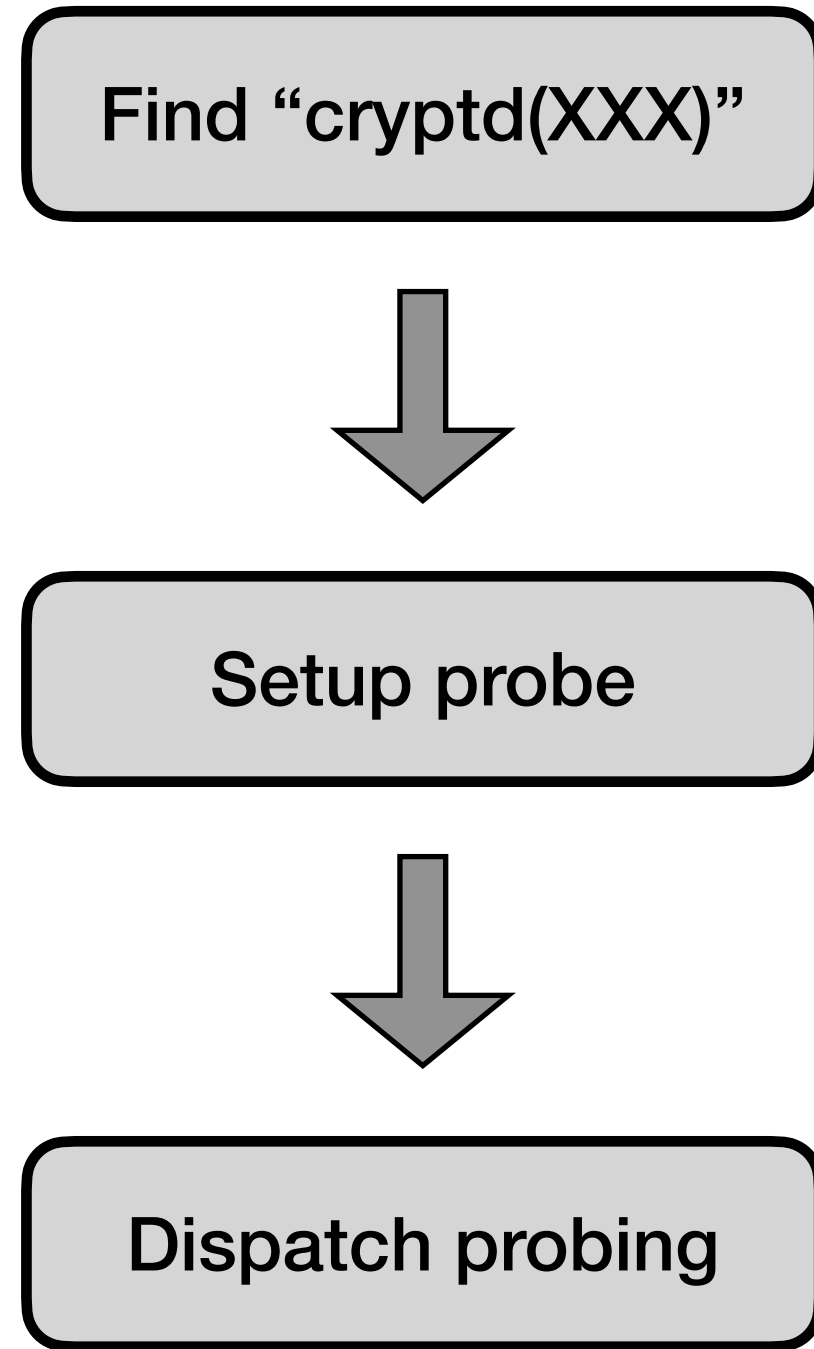
Overview

- **Cryptd**

- Enabled when the `CONFIG_CRYPTD_CRYPTD` compile option is set
- A **crypto daemon** which converts an arbitrary synchronous crypto algorithm into an **asynchronous algorithm** that runs in a **kthread**
- Used as a **template**



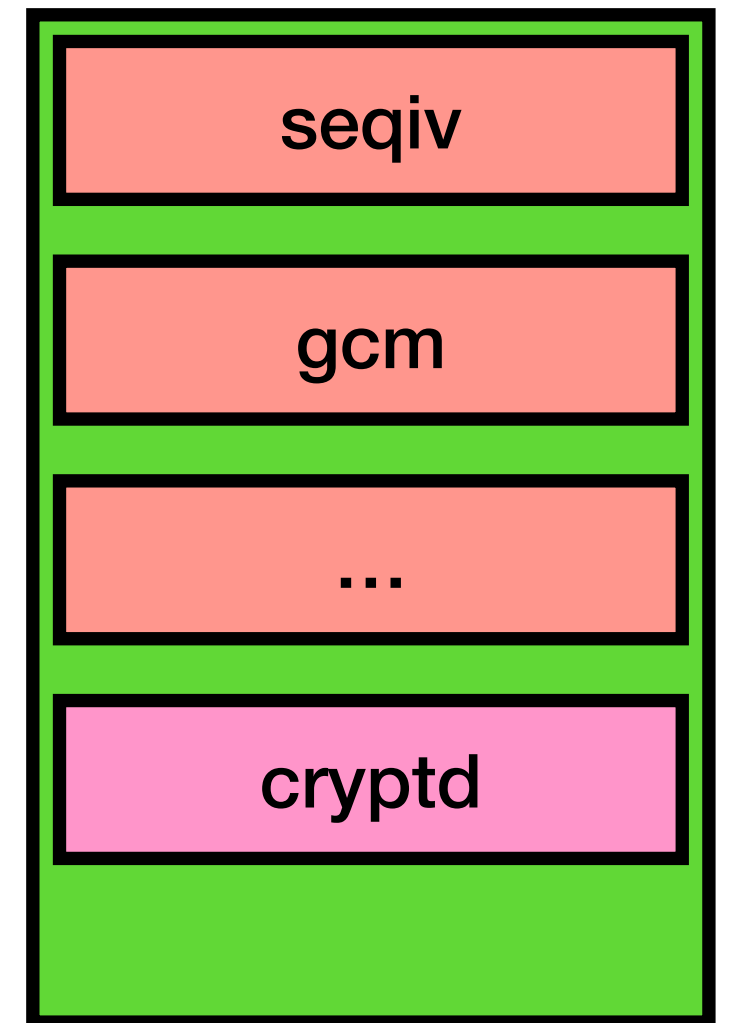
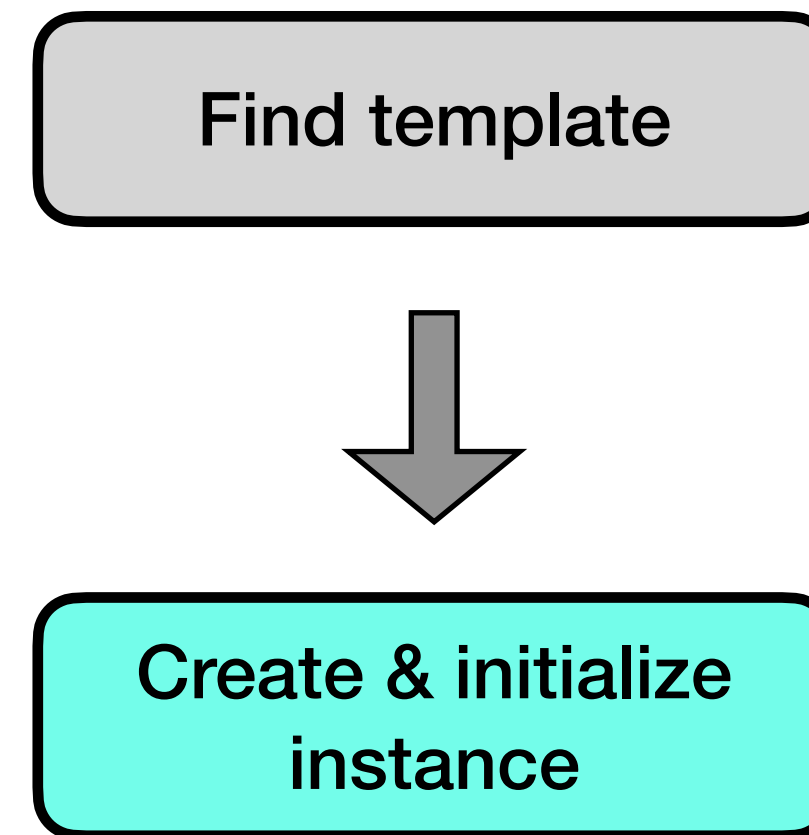
Thread-A



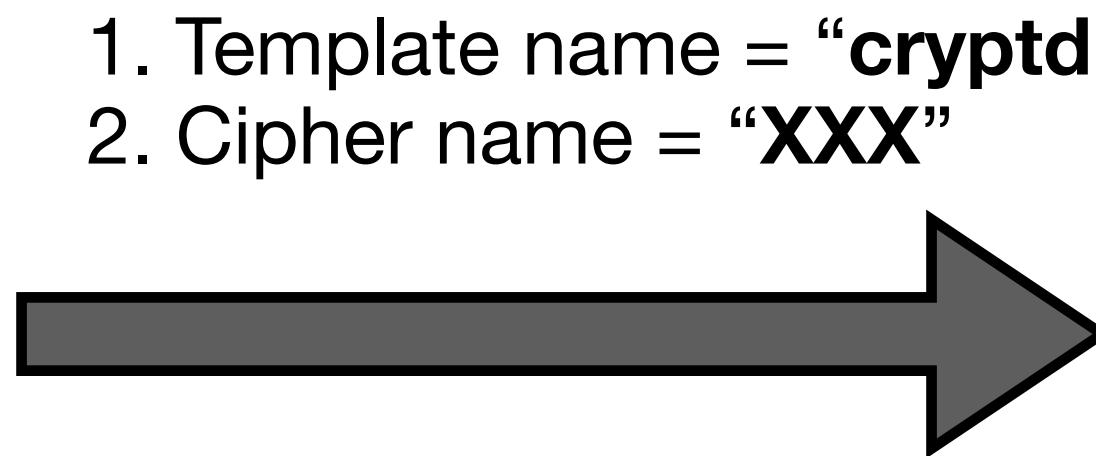
Global variable

Thread-B

"cryptomgr_probe"



Global variable



Thread-A

```

static int cryptd_create_aead(struct crypto_template *tmpl,
                             struct rtattr **tb,
                             struct crypto_attr_type *algt,
                             struct cryptd_queue *queue)
{
    struct aead_instance_ctx *ctx;
    struct aead_instance *inst;

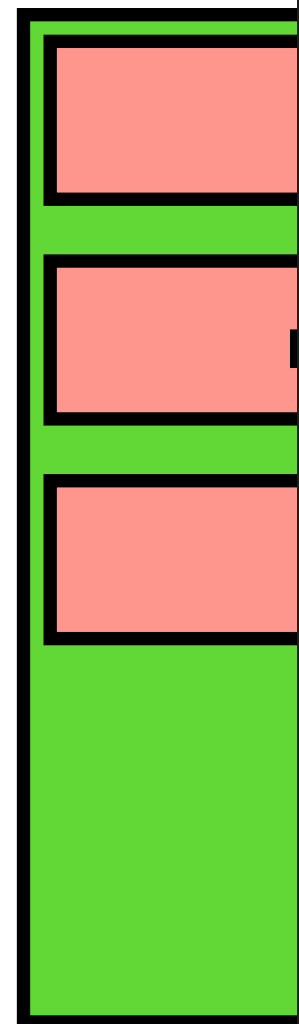
    // [...]
    inst = kzalloc(sizeof(*inst) + sizeof(*ctx), GFP_KERNEL);
    ctx = aead_instance_ctx(inst);

    // [...]
    inst->alg.base.cra_flags |= CRYPTO_ALG_ASYNC
        (alg->base.cra_flags & CRYPTO_ALG_INTERNAL);

    // [...]
    err = aead_register_instance(tmpl, inst);

    // [...]
}

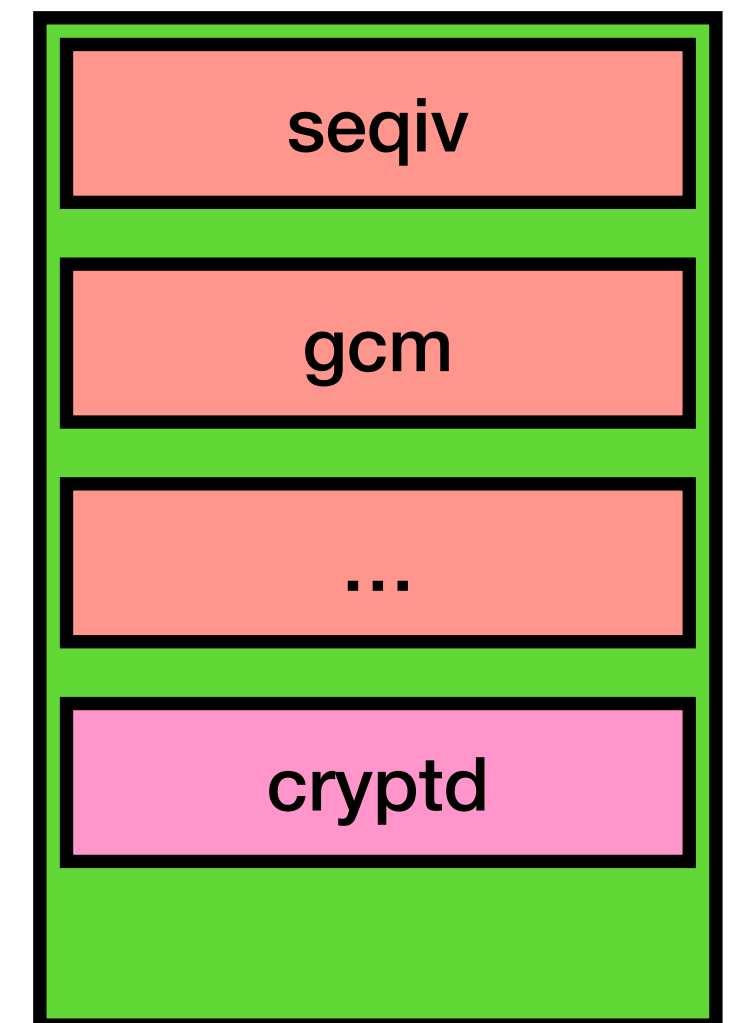
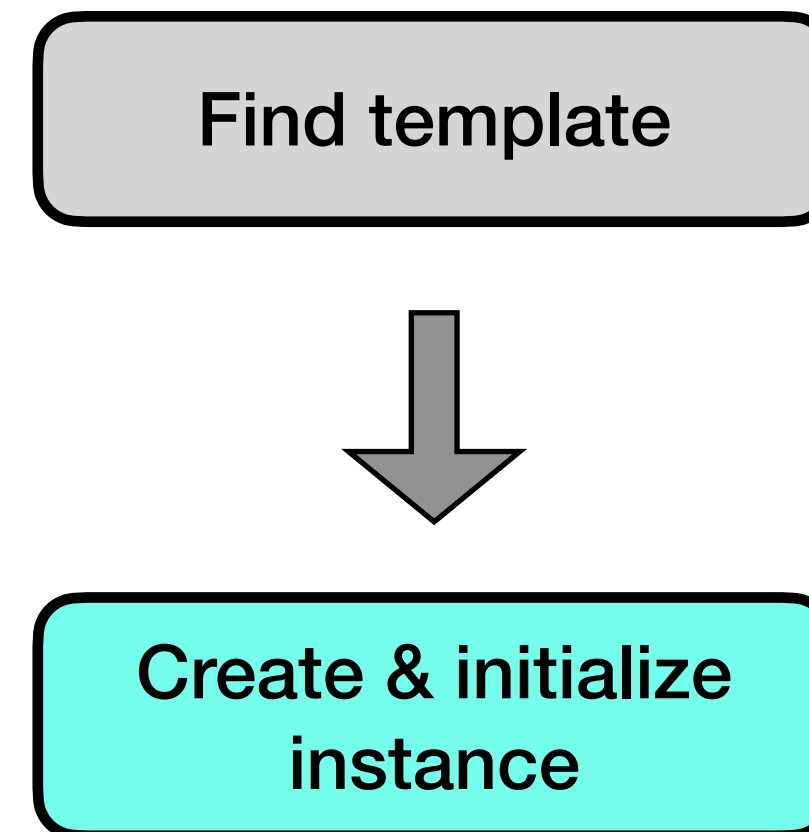
```



Global

Thread-B

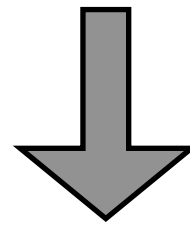
“cryptomgr_probe”



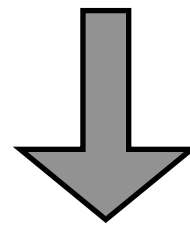
Global variable

Thread-A

Find "cryptd(XXX)"



Setup probe



Dispatch probing

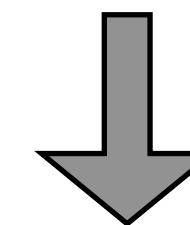
- 1. Template name = "cryptd"
- 2. Cipher name = "XXX"



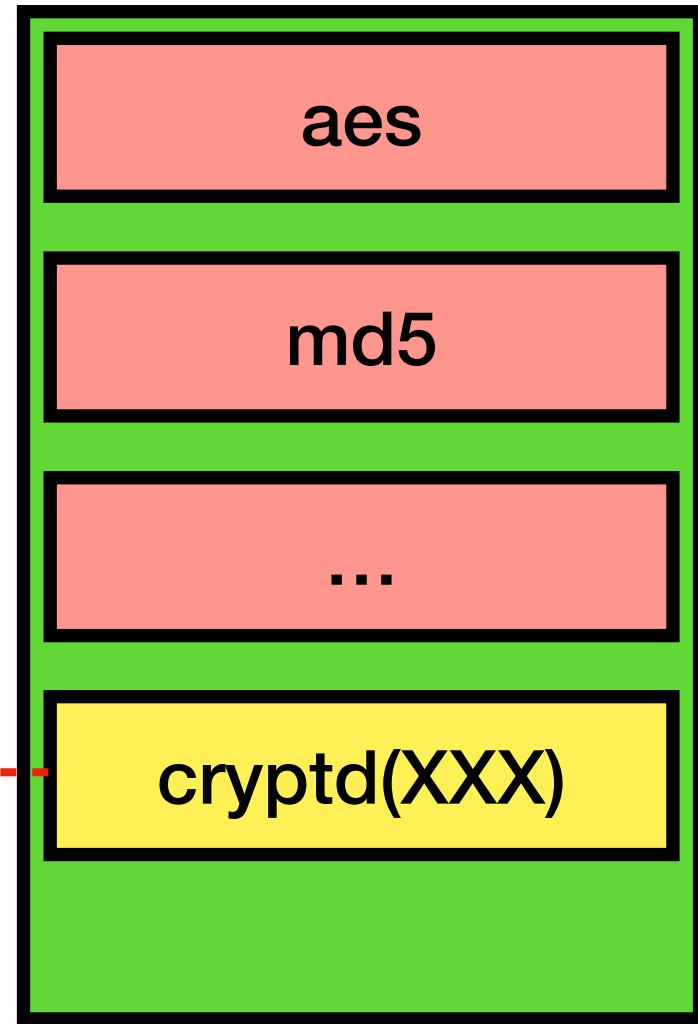
Thread-B

"cryptomgr_probe"

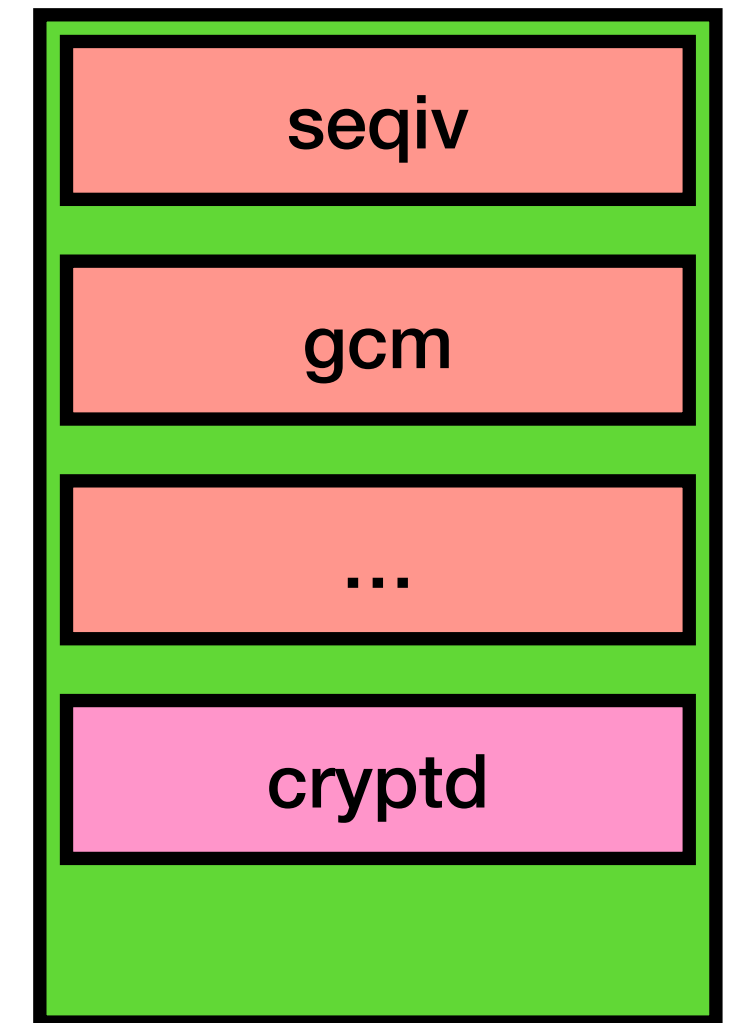
Find template



Create & initialize instance



Global variable



Global variable

Same as the original algorithm (XXX in here) but with asynchronous mode enabled

Overview

- But how?

```
const struct tls_cipher_desc tls_cipher_desc[TLS_CIPHER_MAX + 1 - TLS_CIPHER_MIN] = {  
    TLS_CIPHER_AES_GCM_128, ..., "gcm(aes)"  
    TLS_CIPHER_AES_GCM_256, ..., "gcm(aes)"  
    TLS_CIPHER_AES_CCM_128, ..., "ccm(aes)"  
    TLS_CIPHER_CHACHA20_POLY1305, ..., "rfc7539(chacha20,poly1305)"  
    TLS_CIPHER_SM4_GCM, ..., "gcm(sm4)"  
    TLS_CIPHER_SM4_CCM, ..., "ccm(sm4)"  
    TLS_CIPHER_ARIA_GCM_128, ..., "gcm(aria)"  
    TLS_CIPHER_ARIA_GCM_256, ..., "gcm(aria)"  
};
```

Overview

- **AF_ALG**
 - Interface to kernel crypto API
 - Algorithm probing with user-provided algorithm name

```
#include <linux/if_alg.h>

int sock = socket(AF_ALG, SOCK_SEQPACKET, 0);
struct sockaddr_alg sa = {
    .sa_family = AF_ALG,
    .sa_type = "aead",
    .sa_name = "cryptd(gcm(aes))",
};
bind(sock, (struct sockaddr *)&sa, sizeof(sa));
```


Thread-A

Encrypt a packet

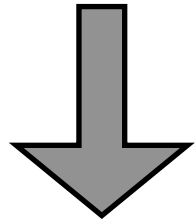
Thread-B
(`cryptd_queue_worker`)

Pending queue



Thread-A

Encrypt a packet



Enqueue request
(Cryptd)



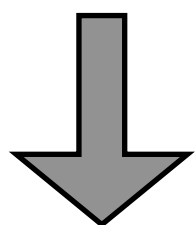
Pending queue



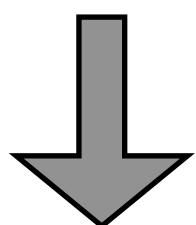
Thread-B
(cryptd_queue_worker)

Thread-A

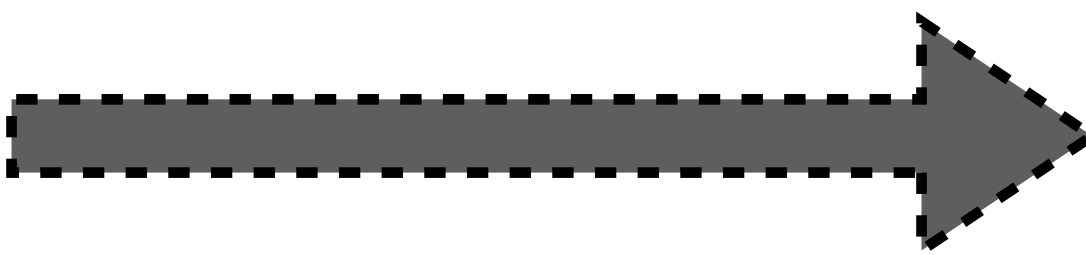
Encrypt a packet



Encqueue request
(Cryptd)



Wakeup worker



Thread-B (cryptd_queue_worker)

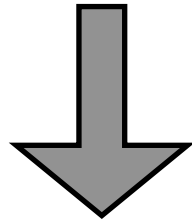


Pending queue

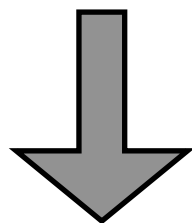


Thread-A

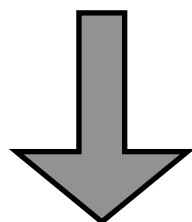
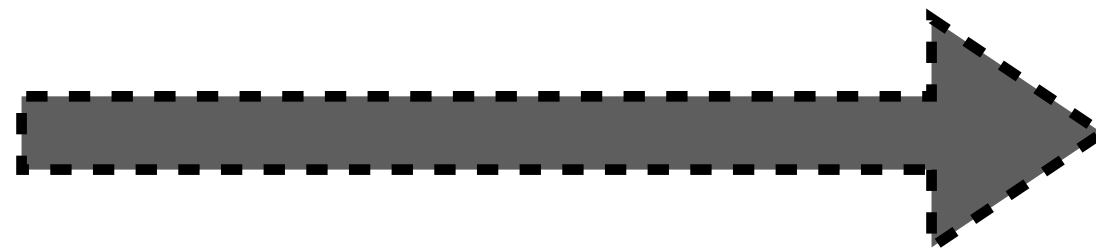
Encrypt a packet



Encqueue request
(Cryptd)



Wakeup worker



Return

Thread-B (cryptd_queue_worker)

Deque request

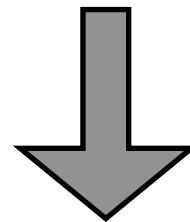


Pending queue

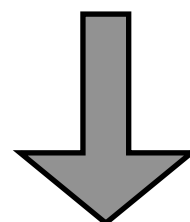


Thread-A

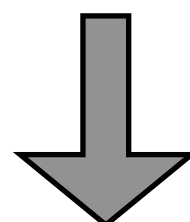
Encrypt a packet



Encqueue request
(Cryptd)



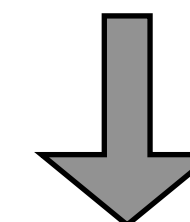
Wakeup worker



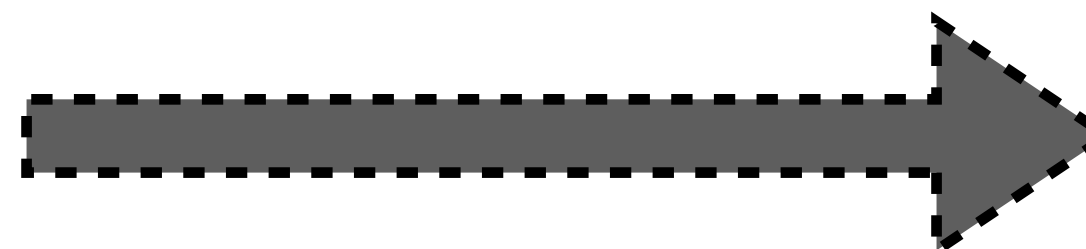
Return

Thread-B (cryptd_queue_worker)

Dequeue request



Handle request




Pending queue



Vulnerability

Vulnerability



CVE-2024-26800	2024-02-29	tls: fix use-after-free on failed backlog decryption
	2024-02-10	Merge branch 'tls-fixes'
	2024-02-10	net: tls: fix returned read length with async decrypt
	2024-02-10	selftests: tls: use exact comparison in recv_partial
CVE-2024-26582	2024-02-10	net: tls: fix use-after-free with partial reads and async decrypt
CVE-2024-26584	2024-02-10	net: tls: handle backlogging of crypto requests
CVE-2024-26585	2024-02-10	tls: fix race between tx work scheduling and socket close
CVE-2024-26583	2024-02-10	tls: fix race between async notify and socket close
	2024-02-10	net: tls: factor out tls_*crypt_async_wait()

Vulnerability

2024-02-29 tls: fix use-after-free on failed backlog decryption

2024-02-10 Merge branch 'tls-fixes'

2024-02-10 net: tls: fix returned read length with async decrypt

2024-02-10 selftests: tls: use exact comparison in recv_partial

2024-02-10 net: tls: fix use-after-free with partial reads and async decrypt

2024-02-10 net: tls: handle backlogging of crypto requests

2024-02-10 tls: fix race between tx work scheduling and socket close

2024-02-10 tls: fix race between async notify and socket close

2024-02-10 net: tls: factor out tls_*crypt_async_wait()

CVE-2024-26583

Vulnerability

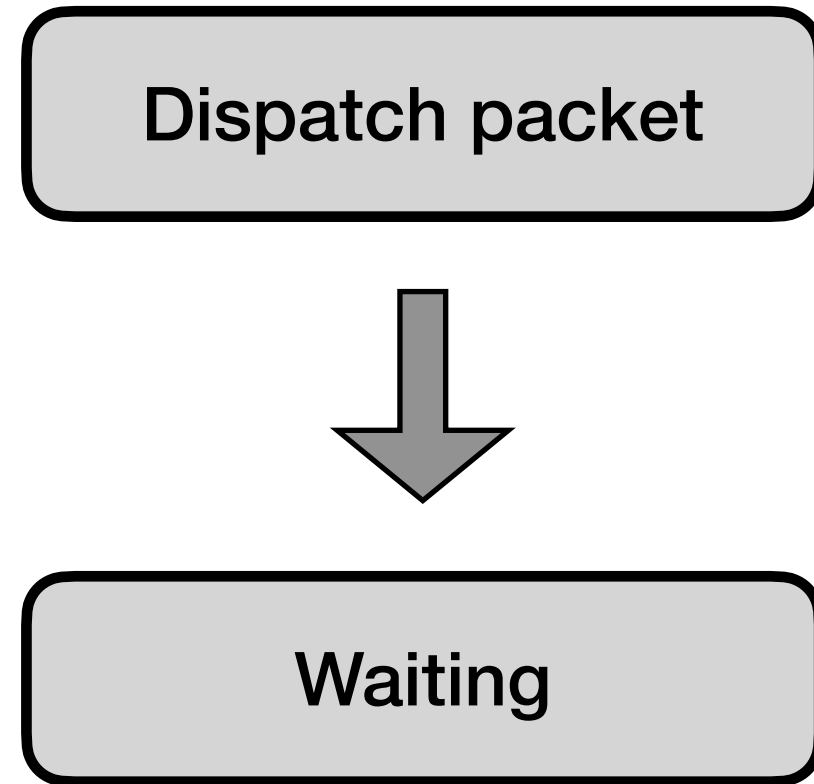
CVE-2024-26583

tls: fix race between async notify and socket close

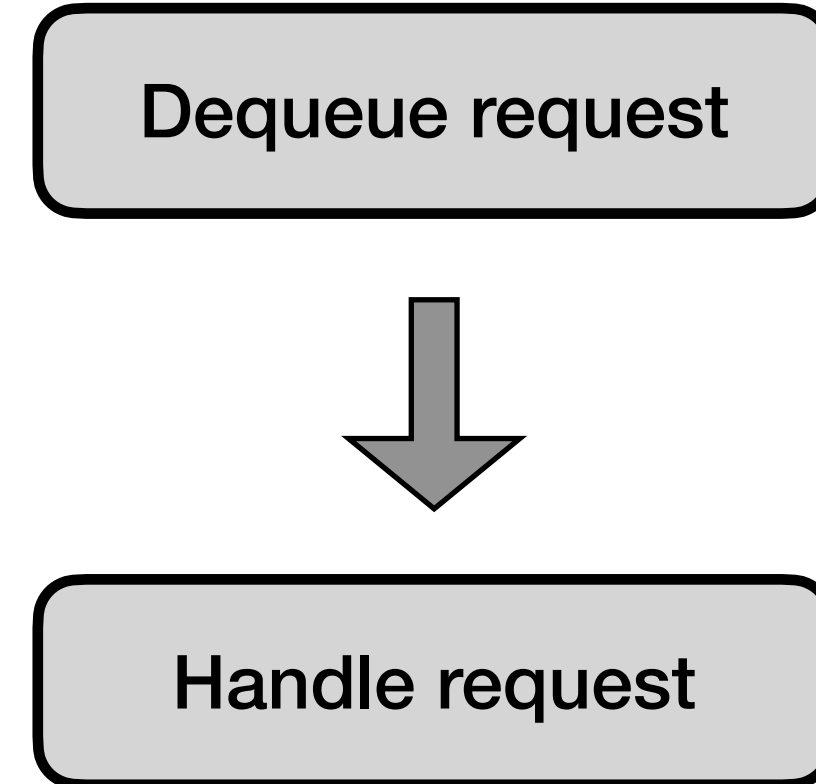
The submitting thread (one which called `recvmsg/sendmsg`) may exit as soon as the async crypto handler calls `complete()` so any code past that point risks touching already freed data.

Try to avoid the locking and extra flags altogether. Have the main thread hold an extra reference, this way we can depend solely on the atomic ref counter for synchronization.

Thread-A

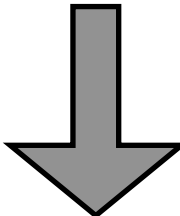


Thread-B (cryptd_queue_worker)



Thread-A

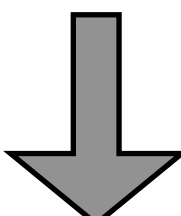
Dispatch packet



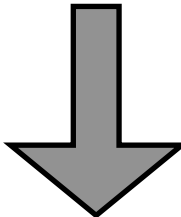
Waiting

**Thread-B
(cryptd_queue_worker)**

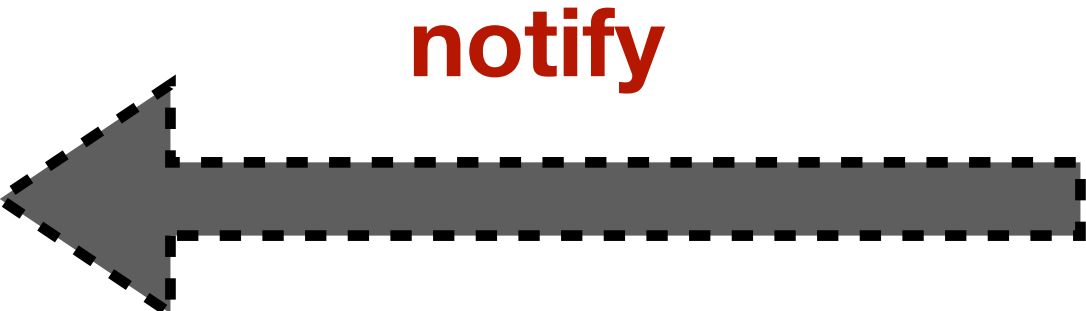
Dequeue request



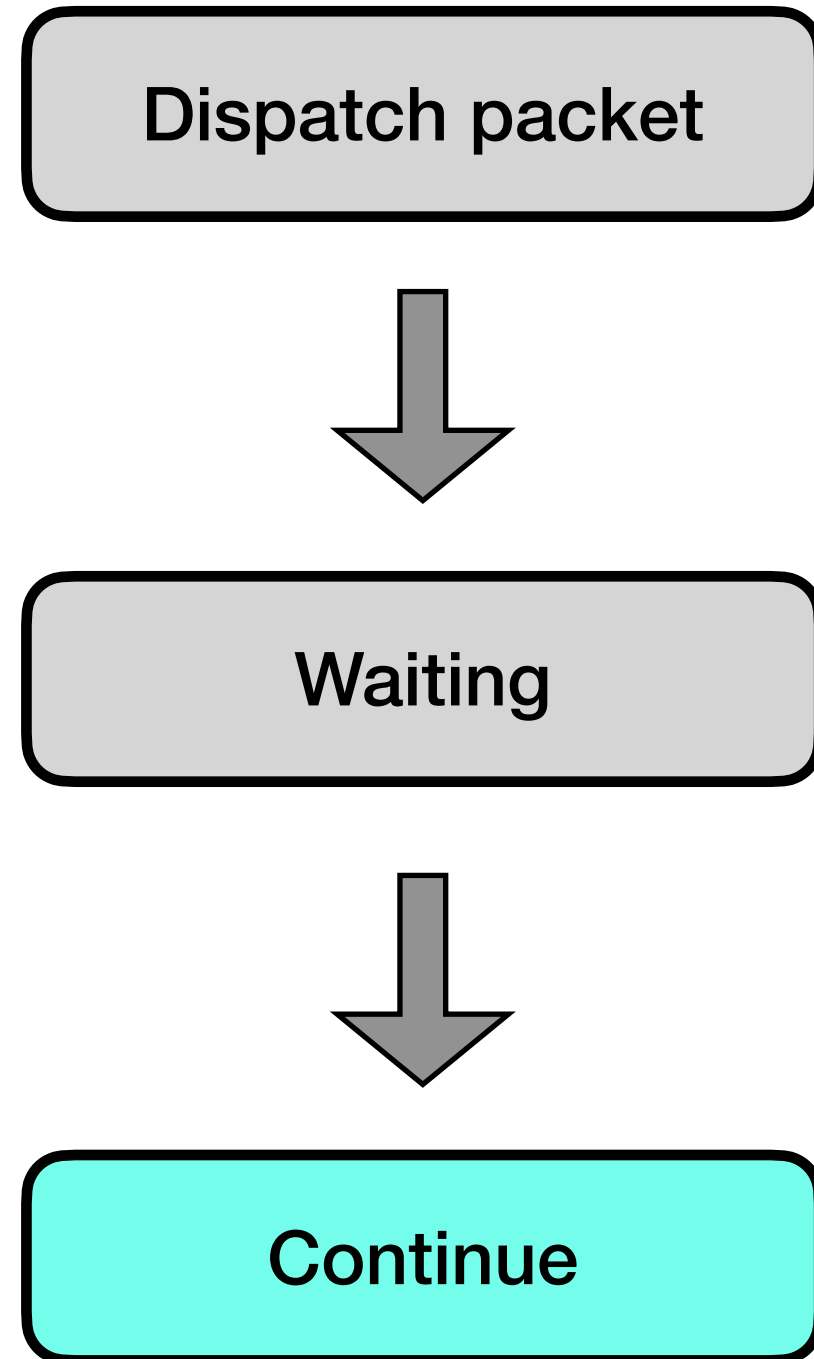
Handle request



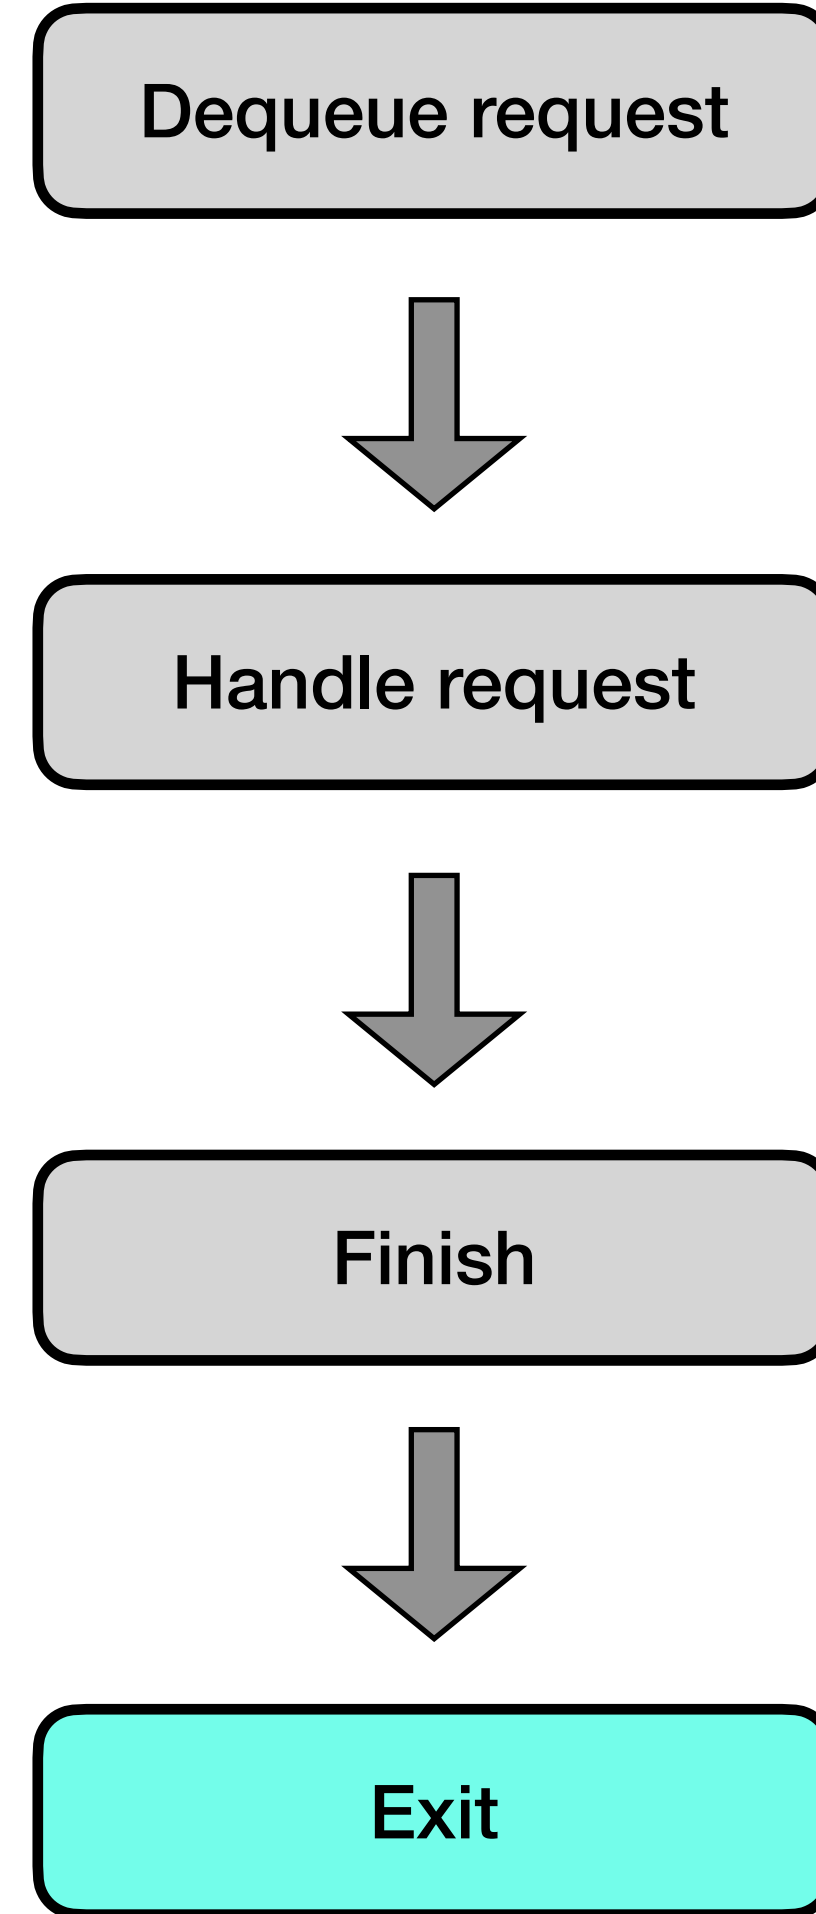
Finish



Thread-A

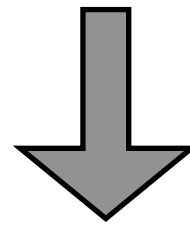


Thread-B (cryptd_queue_worker)

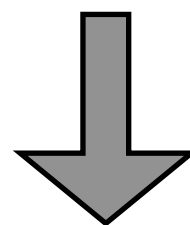


Thread-A

Dispatch packet



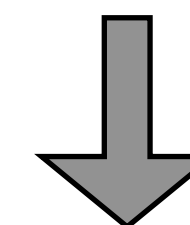
Waiting



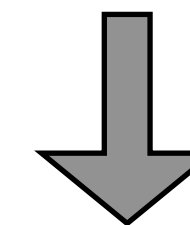
wait_for_completion()

Thread-B (cryptd_queue_worker)

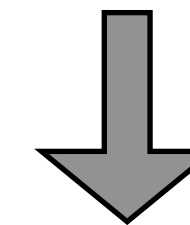
Dequeue request



Handle request



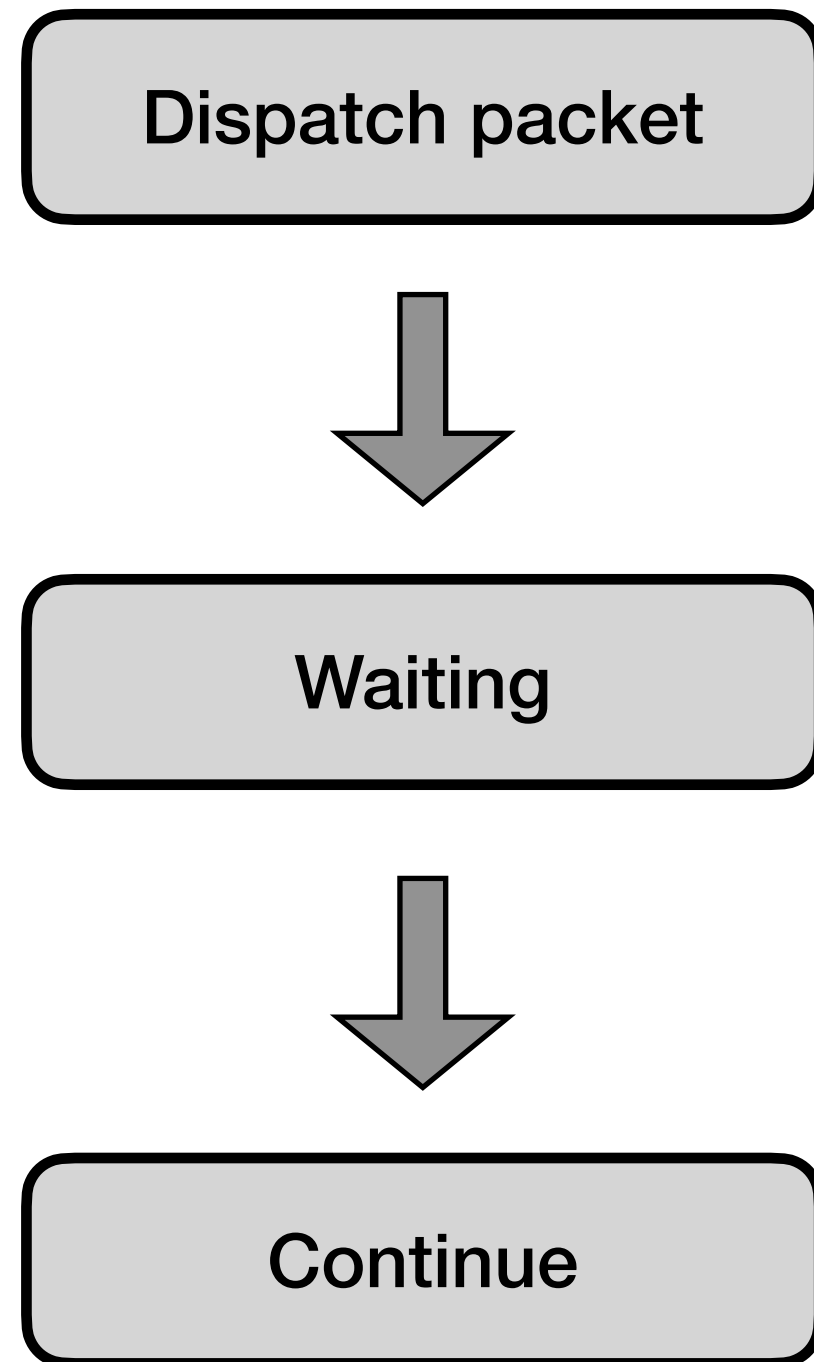
complete()



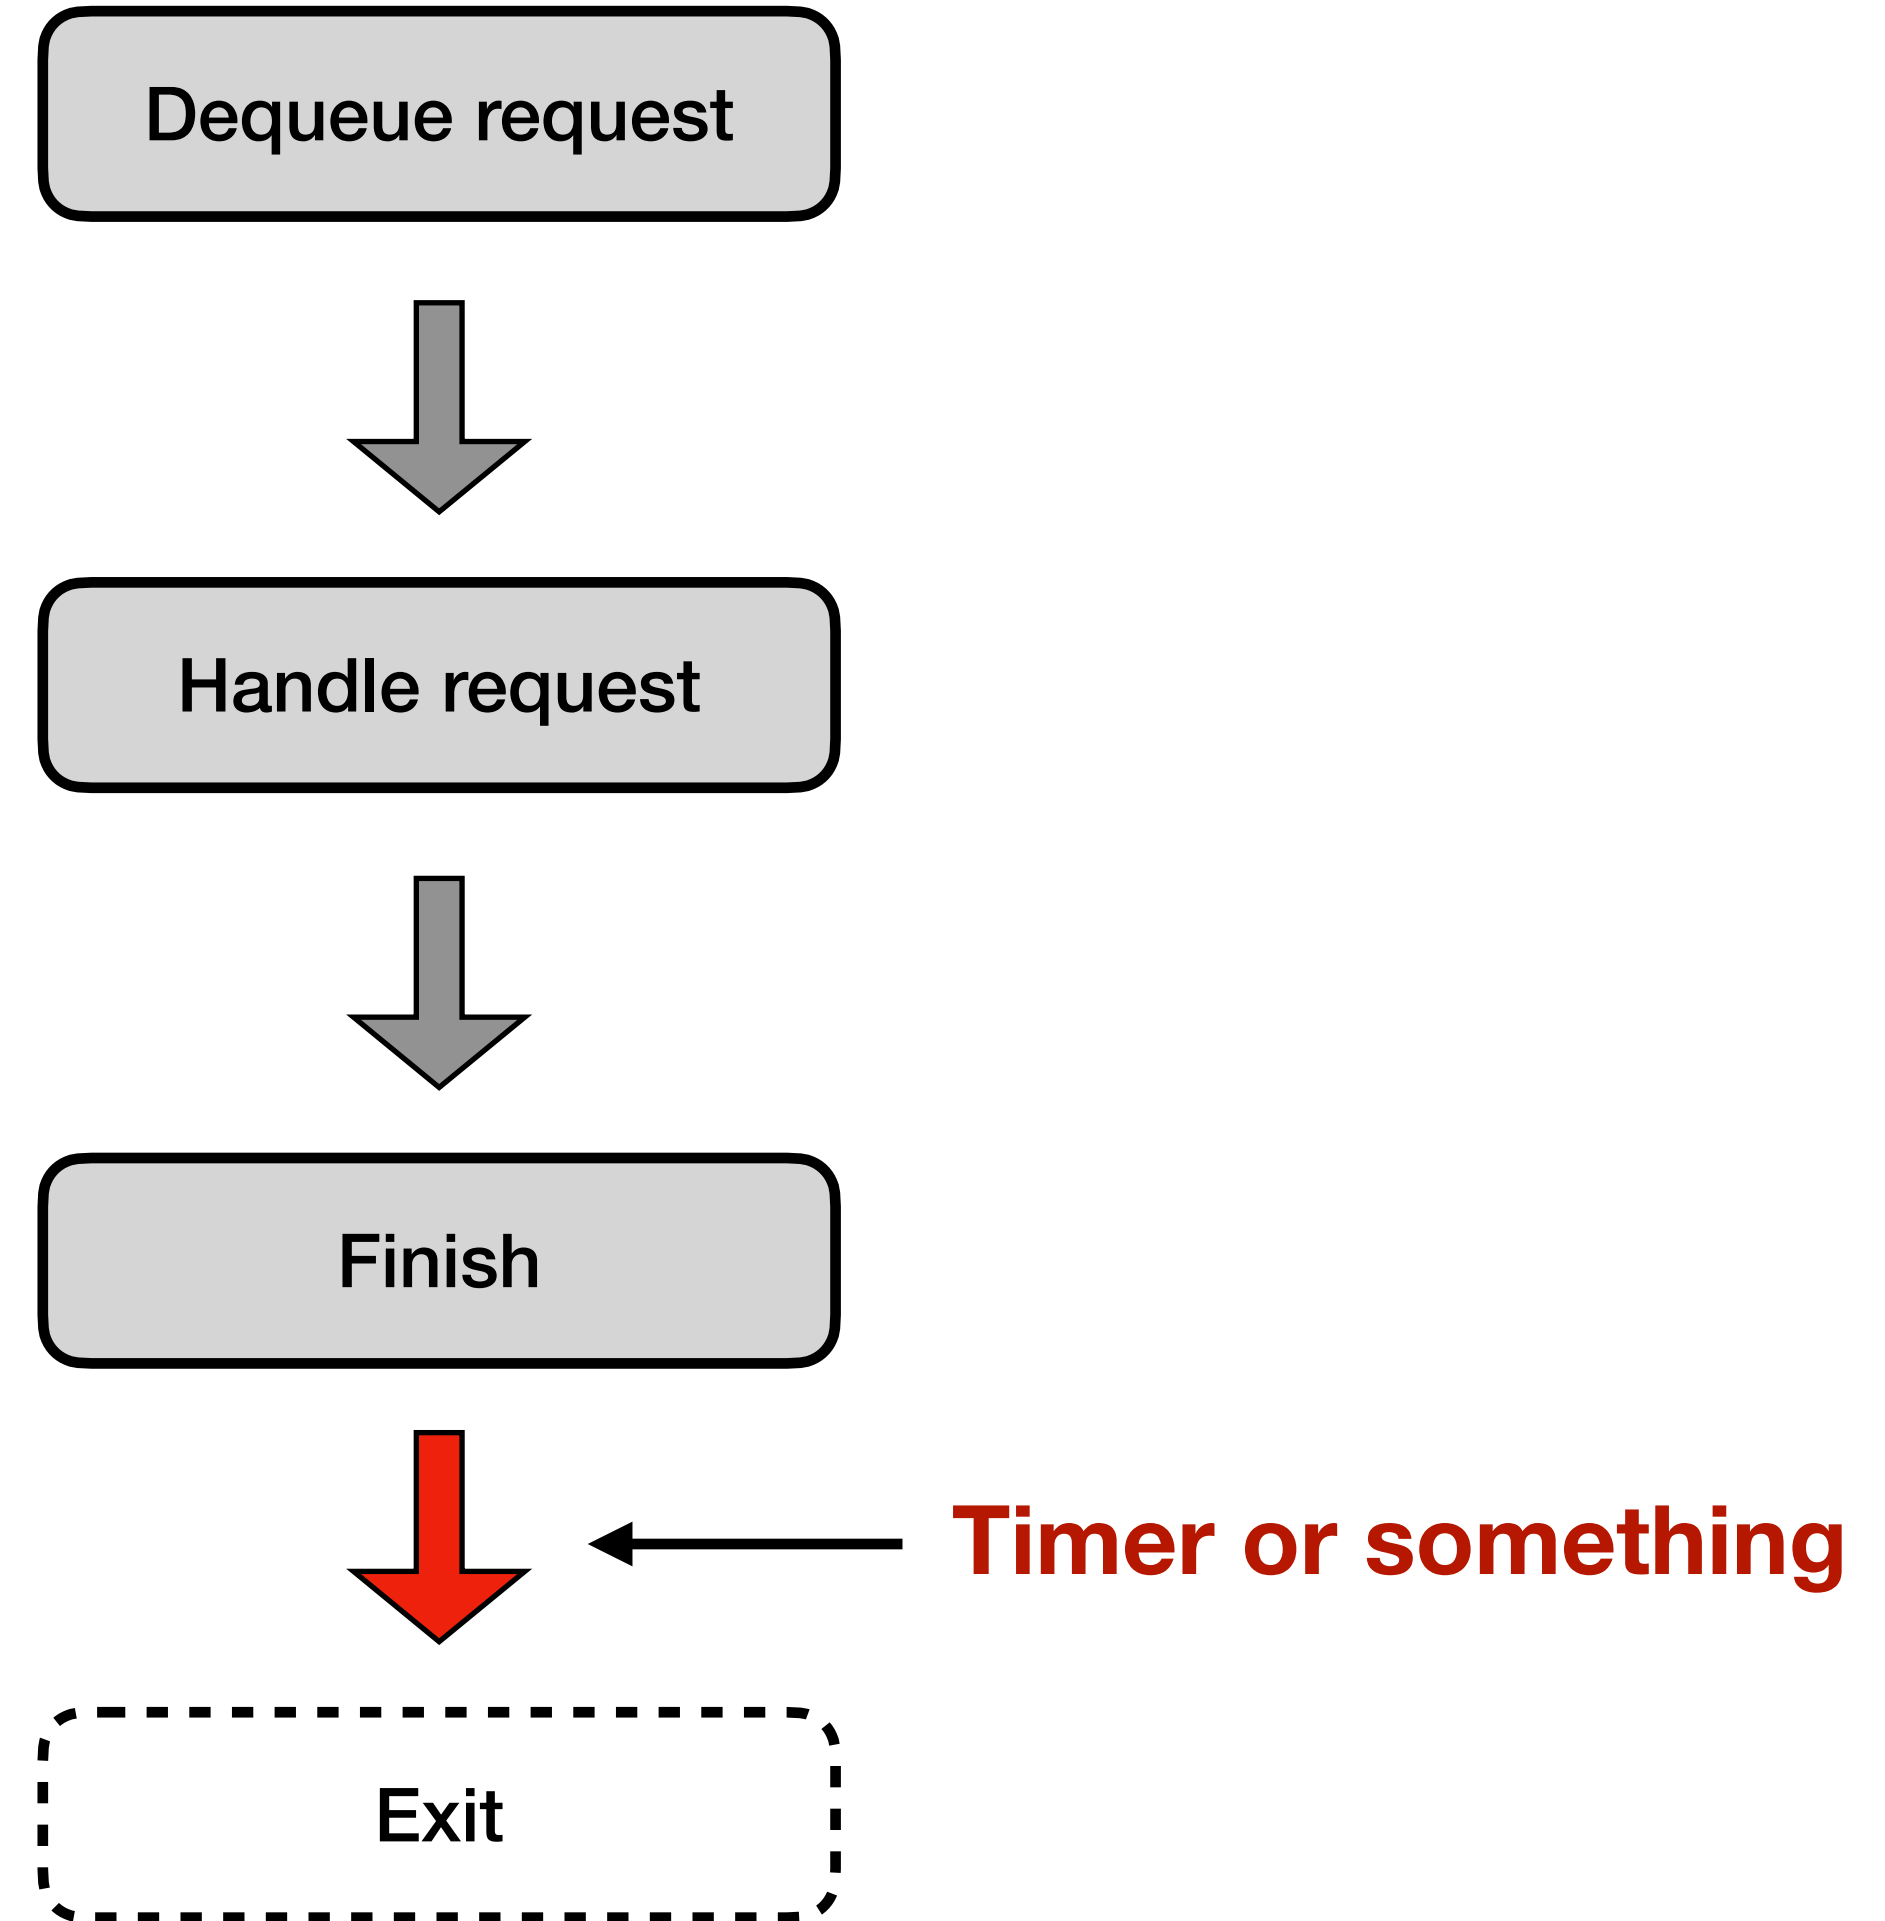
Exit



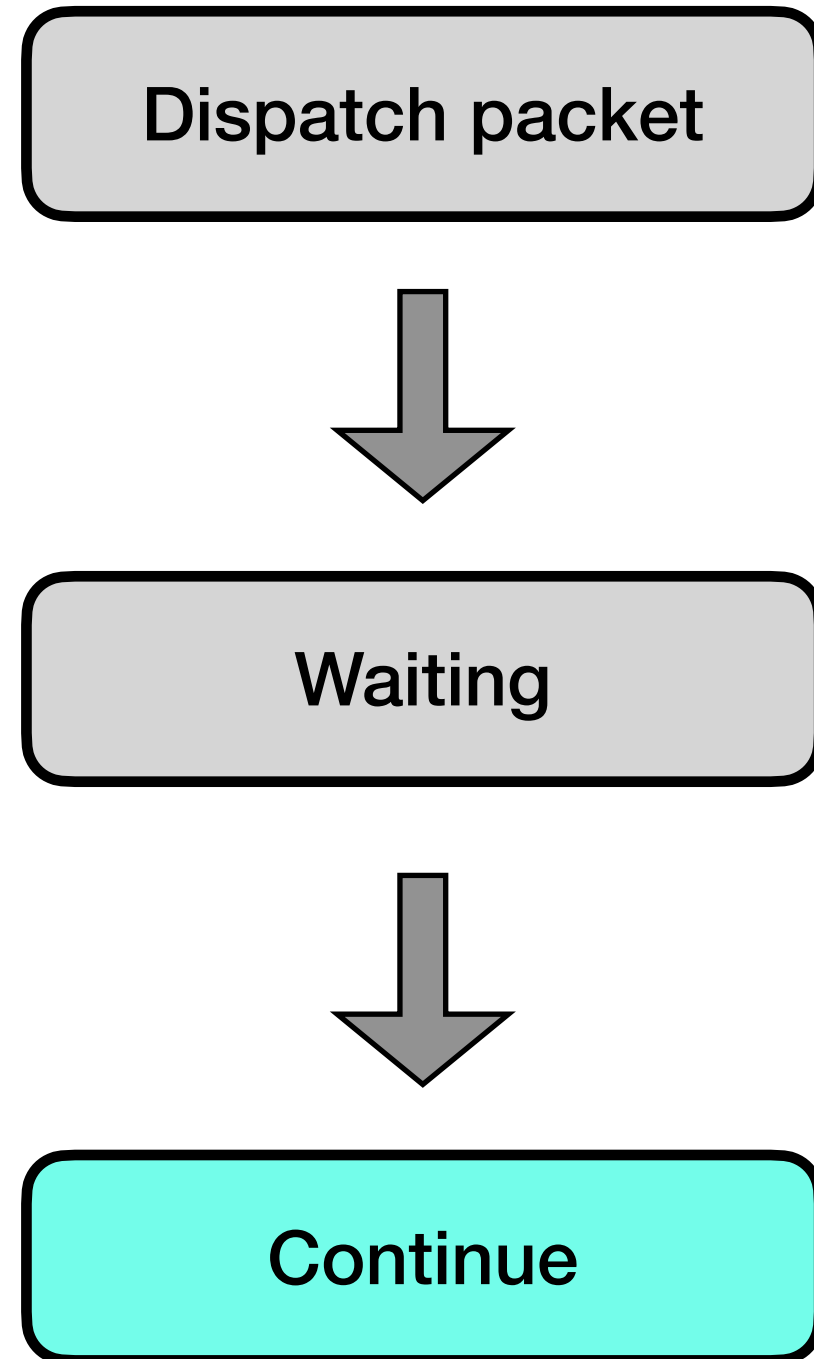
Thread-A



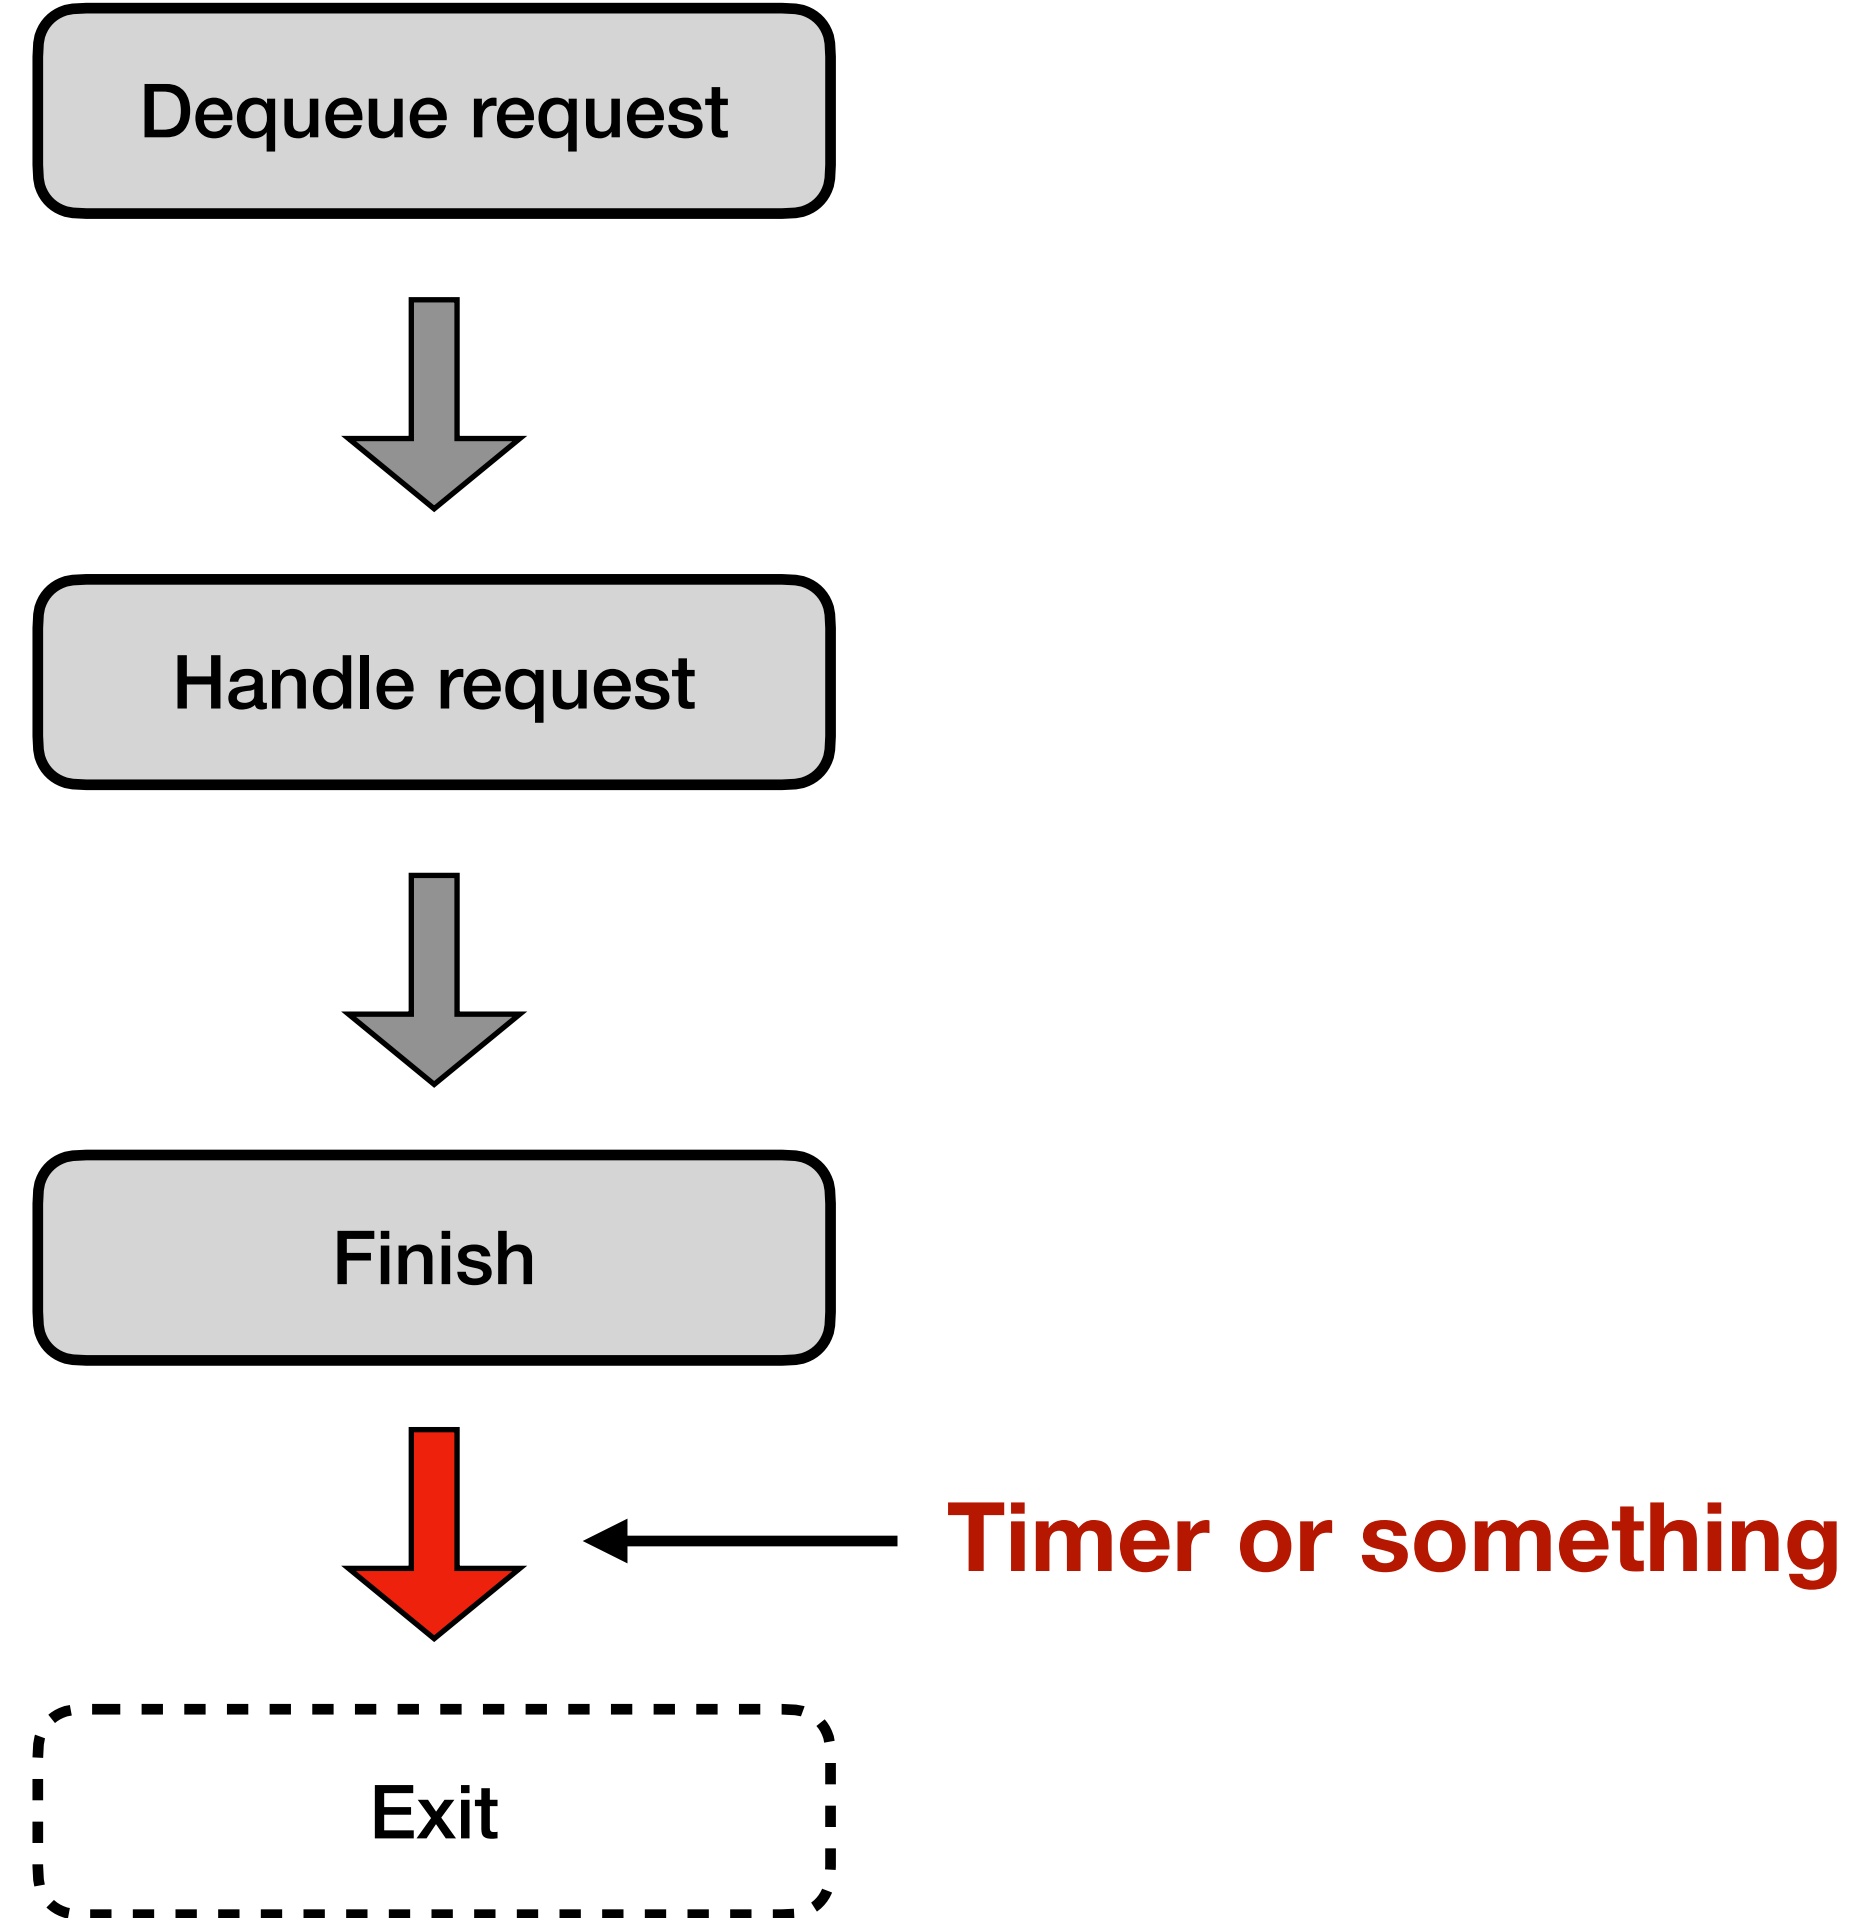
Thread-B (cryptd_queue_worker)



Thread-A

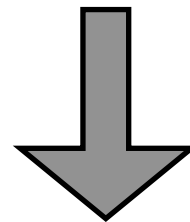


Thread-B (cryptd_queue_worker)

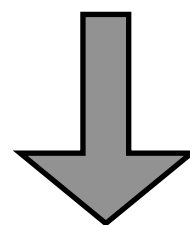


Thread-A

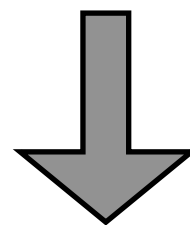
Dispatch packet



Waiting



Continue

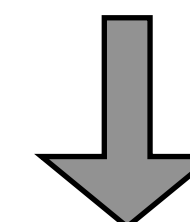


Exit

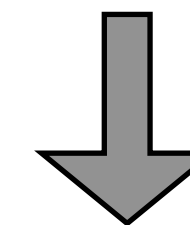
Free TX/RX context ...

Thread-B (cryptd_queue_worker)

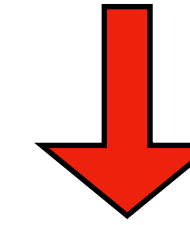
Dequeue request



Handle request



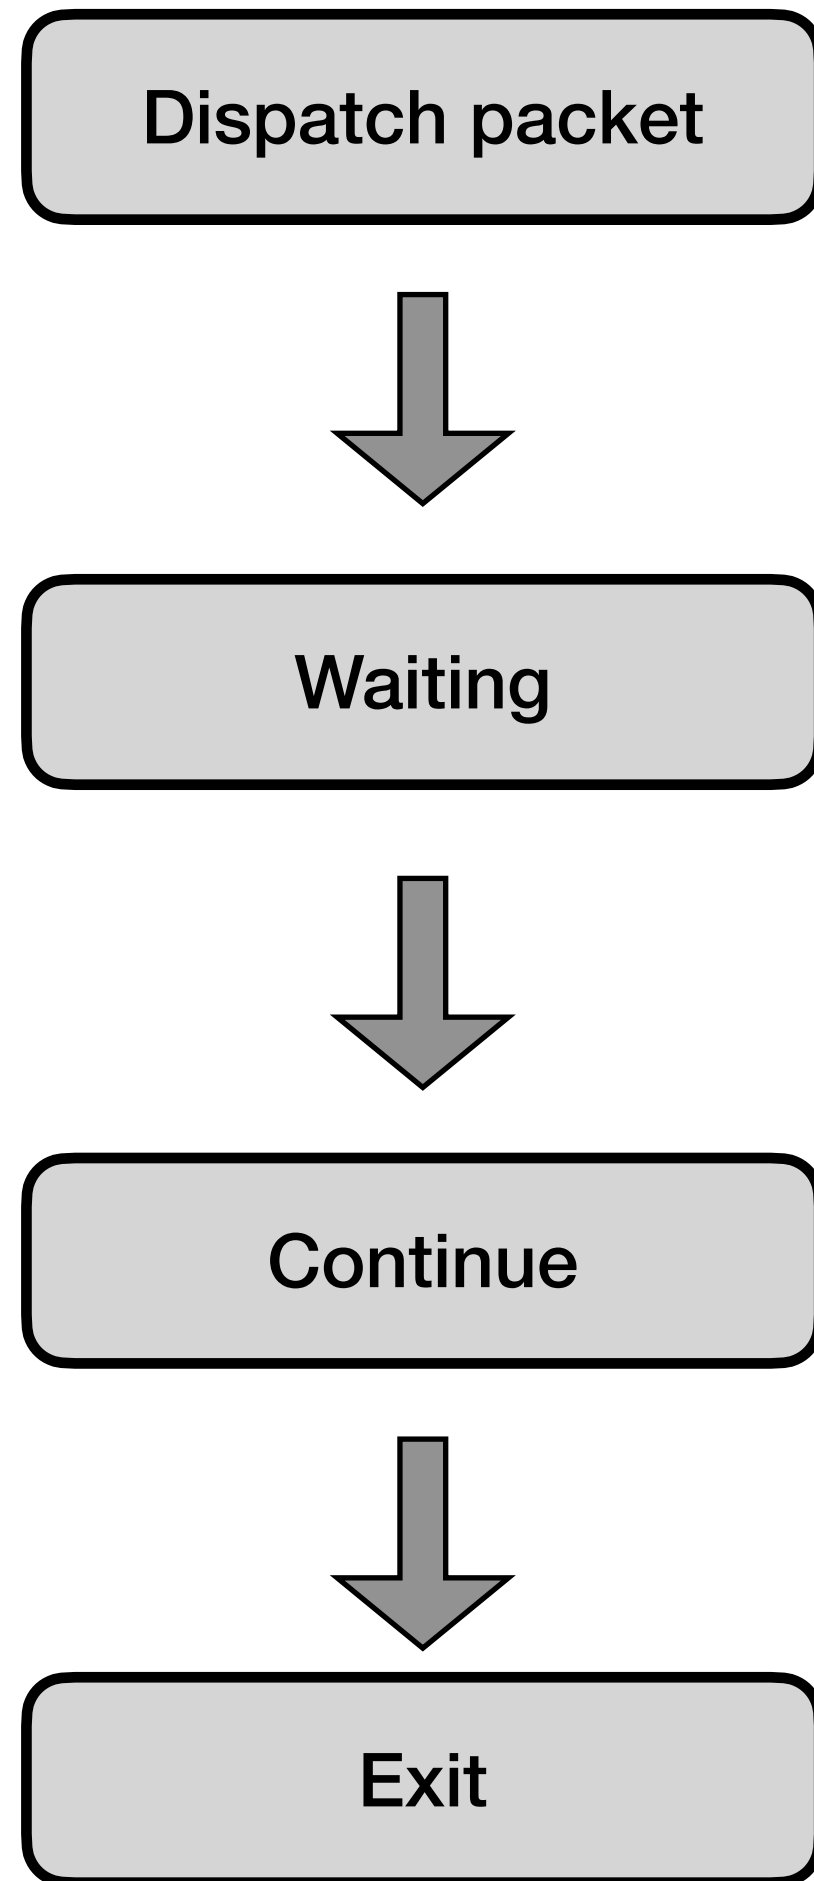
Finish



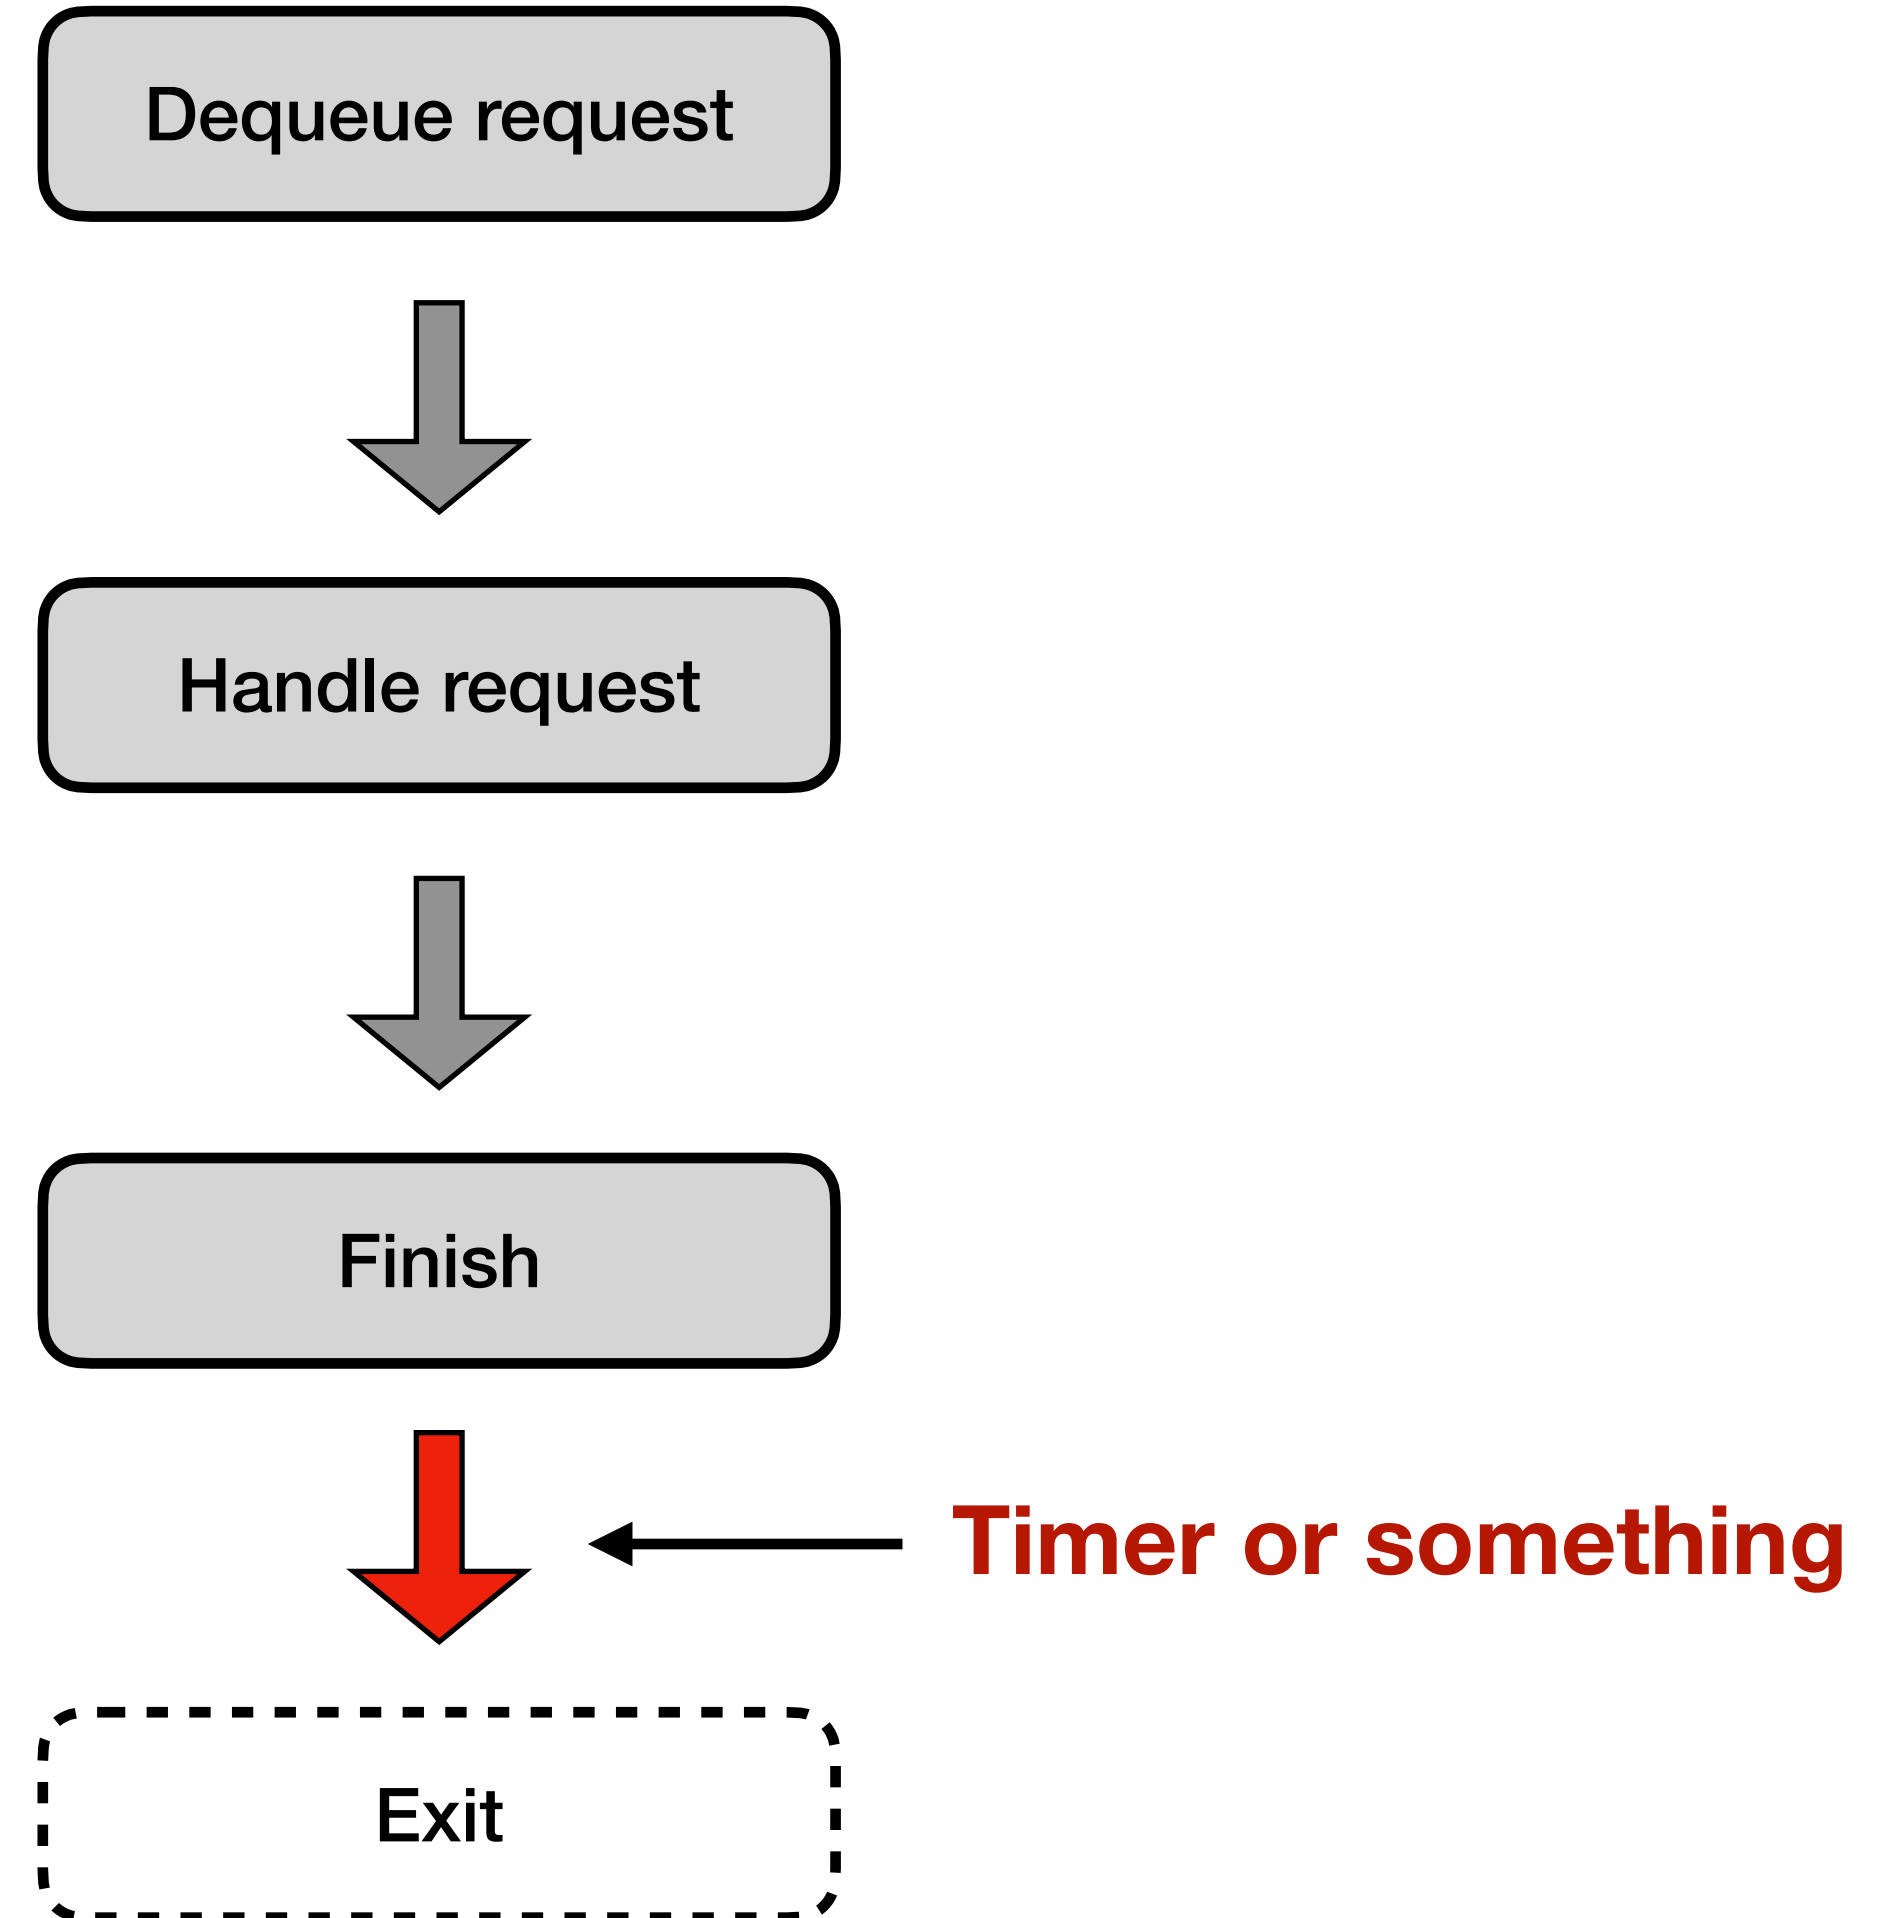
← Timer or something

Exit

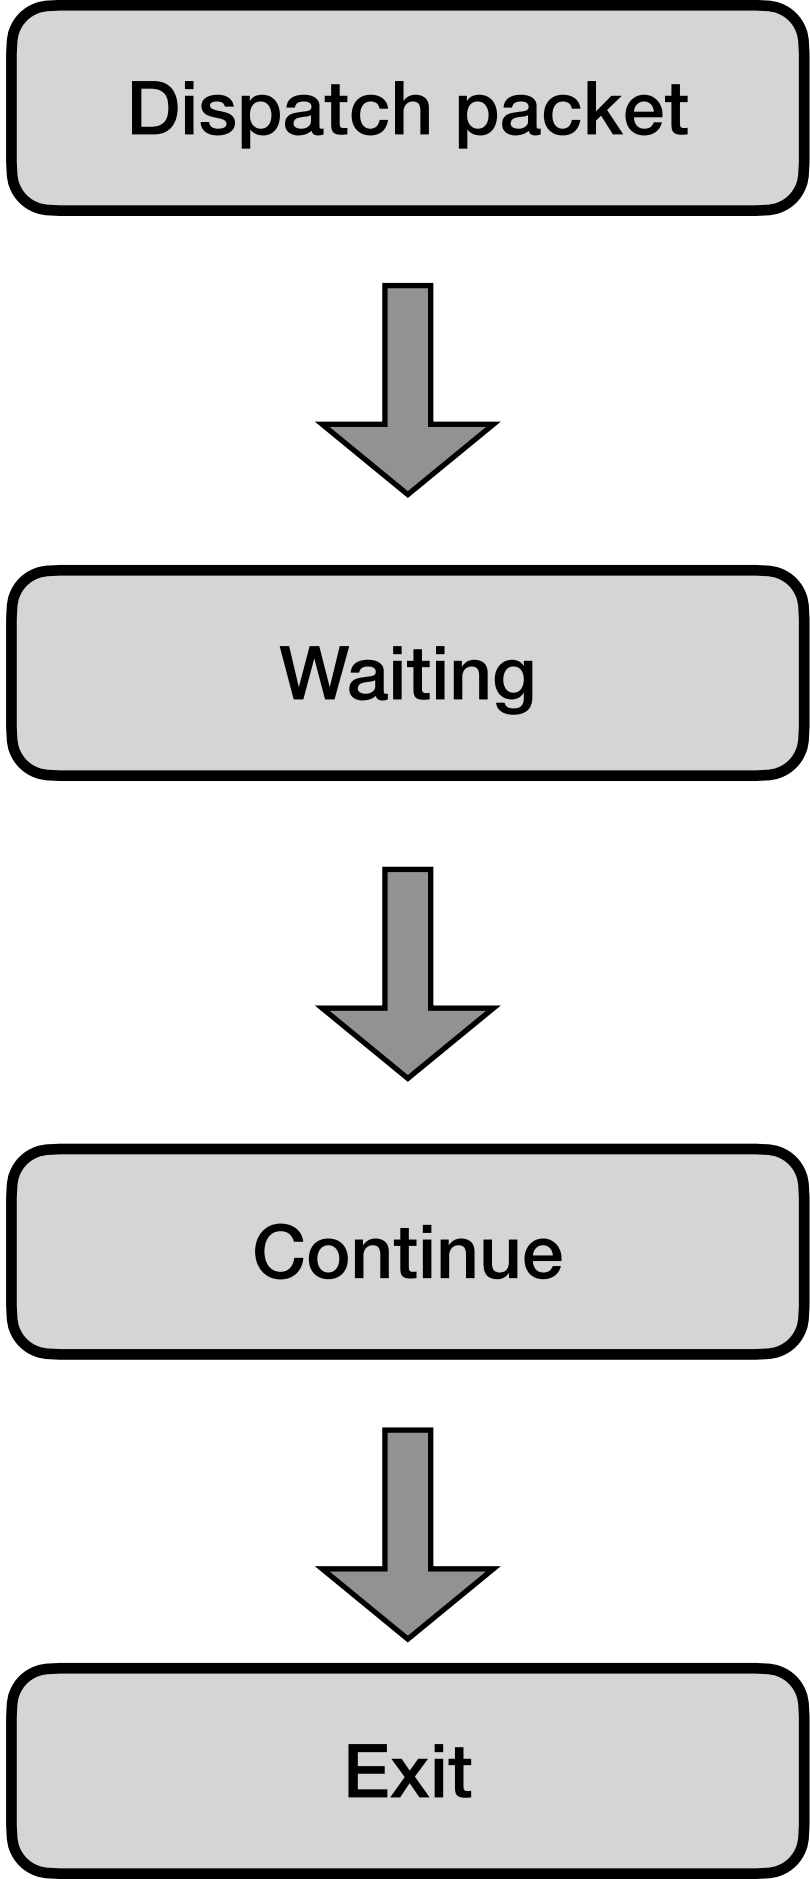
Thread-A



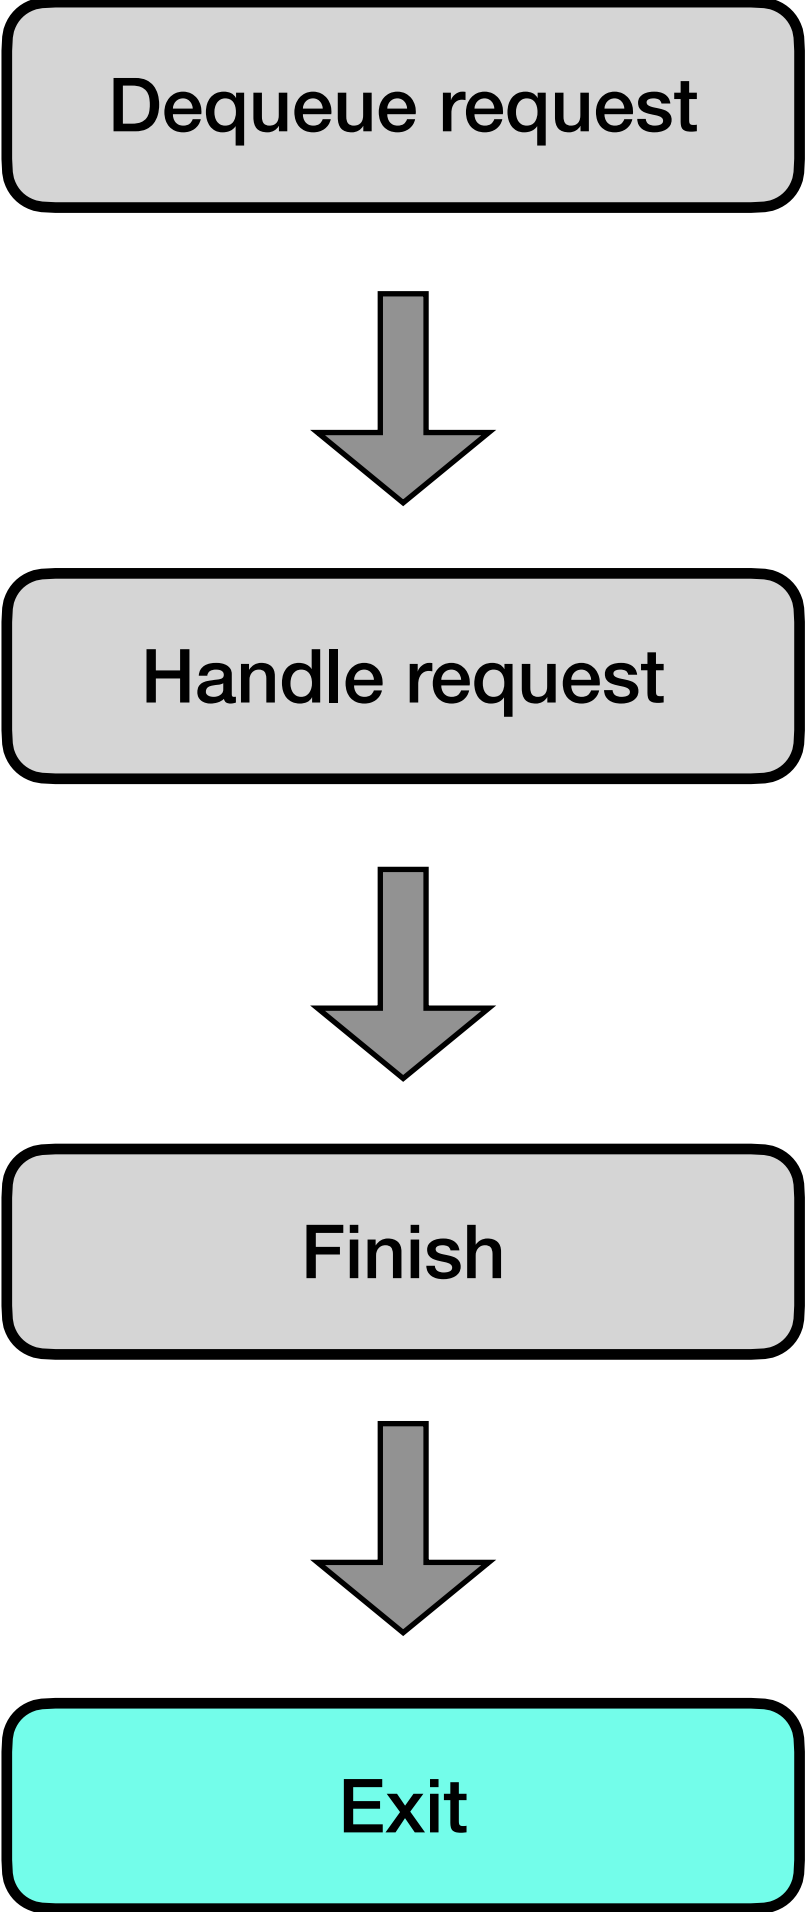
Thread-B (cryptd_queue_worker)

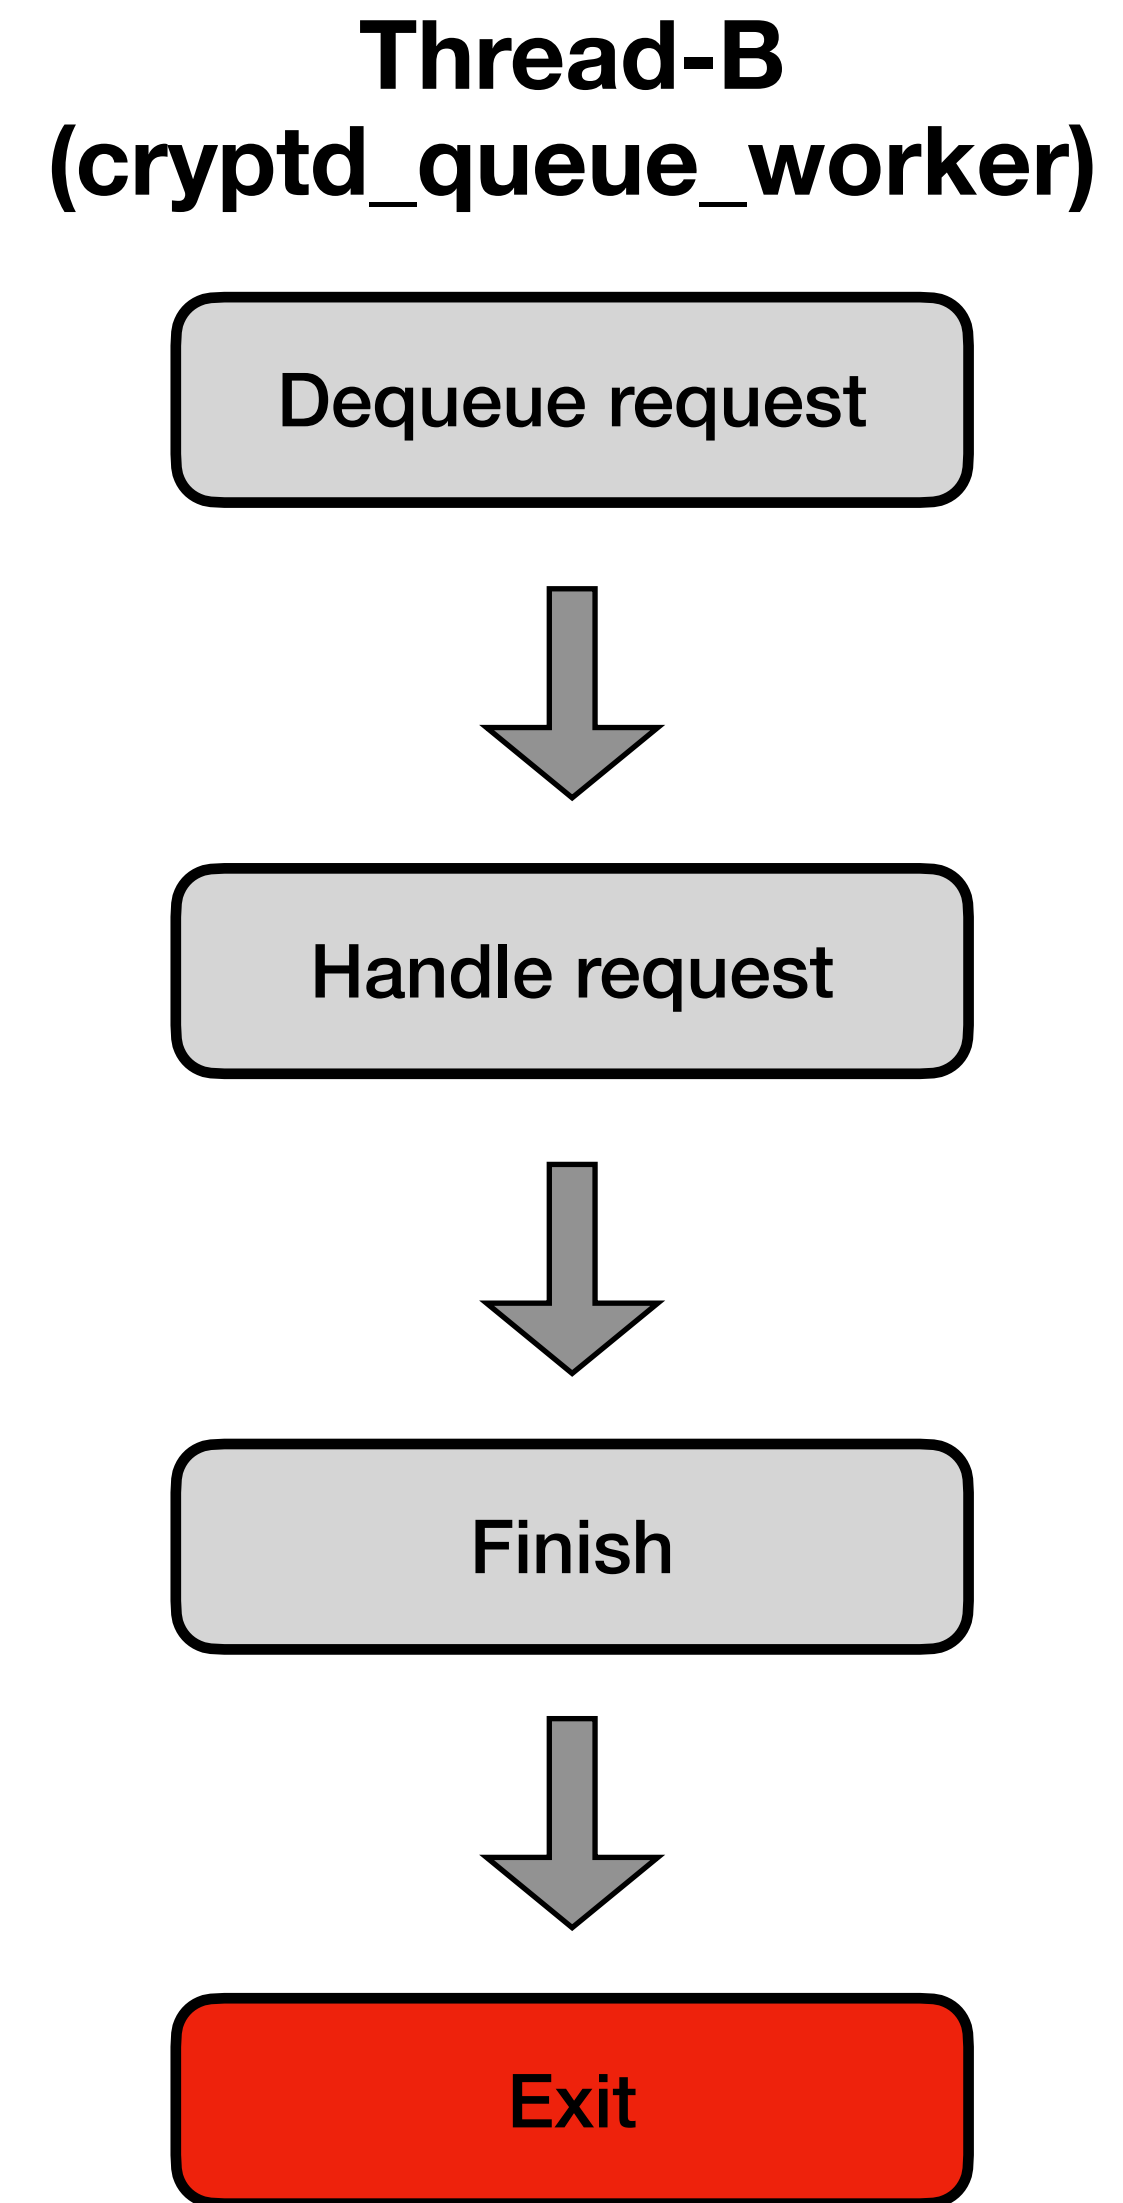
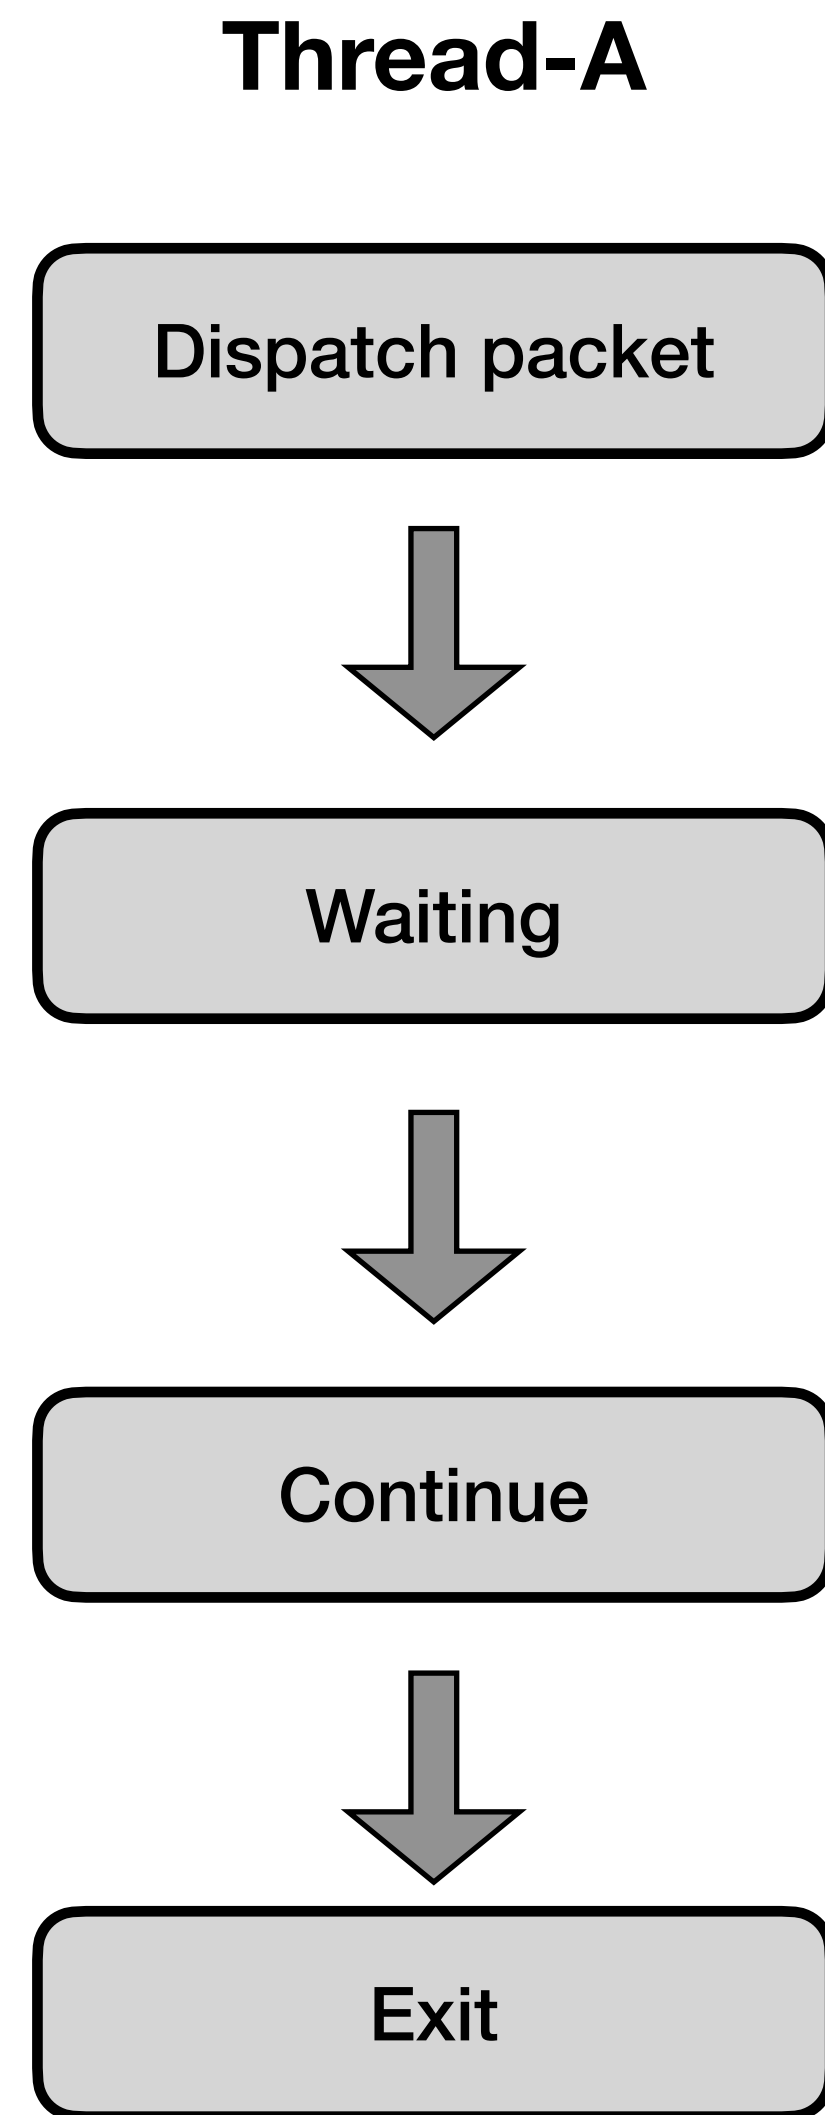


Thread-A



Thread-B (cryptd_queue_worker)





UAF when accessing TX/RX context object