



Revoke and Let Live: A Secure Key Revocation API for Cryptographic Devices

Véronique Cortier, Graham Steel, Cyrille Wiedling

**RESEARCH
REPORT**

N° 7949

July 2012

Project-Team Cassis



Revoke and Let Live: A Secure Key Revocation API for Cryptographic Devices*

Véronique Cortier[†], Graham Steel[‡], Cyrille Wiedling[†]

Project-Team Cassis

Research Report n° 7949 — July 2012 — 38 pages

Abstract: While extensive research addresses the problem of establishing session keys through cryptographic protocols, relatively little work has appeared addressing the problem of revocation and update of long term keys. We present an API for symmetric key management on embedded devices that supports revocation and prove security properties design in the symbolic model of cryptography. Our API supports two modes of revocation: a passive mode where keys have an expiration time, and an active mode where revocation messages are sent to devices. For the first we show that once enough time has elapsed after the compromise of a key, the system returns to a secure state, i.e. the API is robust against attempts by the attacker to use a compromised key to compromise other keys or keep the compromised key alive past its validity time. For the second we show that once revocation messages have been received the system is immediately in a secure state. Notable features of our designs are that all secret values on the device are revocable, and the device returns to a functionally equivalent state after revocation is complete.

Key-words: api, formal methods, revocation

* The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 258865, project ProSecure, and the Direction Générale de l'Armement under contract no 11810242, Secure Interfaces.

[†] LORIA, CNRS & INRIA project Cassis

[‡] INRIA, PARIS, France

RESEARCH CENTRE
NANCY – GRAND EST

615 rue du Jardin Botanique
CS20101
54603 Villers-lès-Nancy Cedex

Révoquer et laisser vivre : Une API sécurisée avec révocation de clefs.

Résumé : De nombreuses recherches traitent du problème de l'établissement de clefs de session à l'aide de protocoles cryptographiques. Inversement, il y a eu relativement peu de travaux concernant le problème de la révocation et la mise à jour de clefs à longue durée de vie. Nous présentons une API, pour la gestion de clefs symétriques sur des systèmes embarqués, qui supporte la révocation et nous démontrons ses propriétés de sécurité dans le modèle symbolique. Notre API permet deux modes de révocation : le premier, passif, se base sur une durée de vie limitée attribuée aux clefs et le second, actif, envoie des instructions de révocation de clef aux différents dispositifs.

Notre premier résultat démontre que, lorsqu'une clef est compromise, le système retourne, après un certain temps, dans un état sécurisé, c.à.d. un état où les API's respectent la politique de sécurité. Le second résultat montre, qu'une fois les messages de révocation réceptionnés, les appareils retournent immédiatement dans un état sécurisé. Les particularités principales de notre schéma sont que tous les secrets contenus dans un appareil sont révocables et, qu'après révocation, l'appareil retourne dans un état fonctionnel.

Mots-clés : api, méthodes formelles, révocation

1 Introduction

Embedded systems deployed in hostile environments often employ some dedicated tamper-resistant secure hardware to handle cryptographic operations and keep keys secure. Examples include mobile phones (which contain SIM cards), smartphones (recent models include ‘Secure Elements’), public transport ticketing systems (such as the Calypso system which employs smartcards and ‘SAM’ modules [10]), smart utility meters (that include a smartcard-like chip for cryptography), on-vehicle cryptographic devices to support vehicle-to-vehicle networking [16] *et cetera*. In such systems, it is often necessary to support the possibility of remotely revoking and updating the long-term keys on the device. However, while extensive research addresses the problem of establishing a new session key or determining what security properties can be guaranteed in the event of long-term key corruption, relatively little work has appeared addressing the problem of revocation and update of long-term keys.

Most existing solutions for key revocation follow one of two approaches: either key revocation actually relies on some ‘longer term’ key that cannot be itself revoked, or key revocation is simply achieved by disabling, resetting or isolating the compromised device. For many applications, both these approaches are unsatisfactory. We propose that a key revocation API should ideally satisfy the following properties:

1. *The device should remain functional* - specifically, whenever possible, the device should return to an equivalent functional state after revocation (but with fresh keys in place).
2. *Any key should be revocable* - side-channel attacks may compromise (perhaps with significant effort) any of the keys stored on a cryptographic device, and the more sensitive a key is, the more likely an attacker is to dedicate effort to breaking it. Hence it is not prudent to decide in advance which keys may or may not be compromised.

The first main contribution of this paper is the design of an API with update and revocation functionalities that satisfy these properties. We assume a very general scenario in which a population of user devices (which could be anything from employee smartcards to enterprise HSMs and can include clients and servers) wish to establish keys and communicate securely between themselves, and a small number of administrator devices are used only to setup new devices and to revoke compromised keys. We start from an existing API for symmetric key management [4] capable of supporting a variety of protocols and show how to extend it with commands for revocation and update of long-term keys. Our design allows the right to update long-term keys to be shared or delegated in a secure way, making it suitable for the modern embedded systems environment where many actors (manufacturer, infrastructure provider, service provider, application developer) may need different access rights. We also discuss some limitations of revocation when symmetric key cryptography is the only operation available (as is in fact the case in many real-world embedded systems).

Our second main contribution is a formal proof of security of our API. We show that it ensures two main properties. First, when a key is lost, the system is able to repair itself after some time. Indeed, our first technique for revoking

a key is time-based: each key expires after some delay. Second, if the need to revoke a key is urgent, an explicit revocation command allows us to immediately secure the device, by prohibiting keys of certain levels for a certain period of time. We show in an example how with careful labelling of key levels this can be used to maintain full functionality. Our API makes use of special management keys, which may of course be revoked as well. Our security proofs are carried out in a symbolic model, based on terms for representing messages. The detailed proofs are included in an appendix. A short version of this report has been accepted to CCS'12 [5].

Related work Most deployed key revocation mechanisms are quite simple. For example, the Trusted Platform Module [15] supports a `TPM_OwnerClear` command that resets the 'ownership' status of the device and erases most keys such as the root key for the storage hierarchy. It does not support, for example, the revocation of individual keys in the hierarchy. In multicast group key management schemes, hierarchies are also standard, and here revocation of a key corresponds either to permanent removal of an entity from the group, or removal until the entity resubscribes, in which case long term secrets are assumed to be still intact. Our work will not make this assumption. One lesson from this domain is that many proposals published without proofs of security have turned out to have attacks [14].

Richer key revocation schemes have received some previous attention in the academic literature, e.g. in the key management schemes of wireless sensor networks. Here nodes in the network are expected to be deployed in hostile environments where key compromise may occur. Eschenauer and Glgor propose a scheme under which every node i shares a long-term symmetric key K_{Ci} with a control server [7]. If a node n is compromised, the server sends a fresh signature key S_i to each nodes i sharing keys with n encrypted under K_{Ci} , and then signs a list of keys to be revoked and sends this list to each i . The nodes then delete the keys, isolating the compromised node. This scheme has two disadvantages for the general case: first it assumes that a central authority knows the key identifiers of all the keys in the network and between what parties they are shared, and second it has no way to recover a device which has lost its K_{Ci} to the attacker. An alternative scheme is KeyRev [17], which proposes the use of a secret sharing scheme to distribute session keys only to unrevoked nodes, thereby isolating compromised nodes without explicitly revoking their long term keys. Again this has the disadvantage that the loss of a long term key means the unrecoverable loss of a device: acceptable in a wireless sensor network perhaps but not in the general case. In this paper we aim to be able to securely update all keys on the device, provided not too many keys have been compromised.

The Sevecom API [9] is a proposal for an on-board tamper-resistant device to handle cryptography in next-generation vehicles supporting VANETs (vehicle to vehicle ad-hoc networks). It includes two root public keys which are used to check authenticity of messages coming from a central server. The signed messages from the central server are used to trigger updates of working keys. The Sevecom API is interesting because it allows update of the long term root keys using a simple two step protocol. We will examine this example in more detail in the next section.

There are other proposals for key management APIs with security proofs in

the literature, but none of these address the question of revocation [3, 6, 8].

2 Revocation API Design Issues

We set the following requirements for our design: (1) the device should return to full functionality whenever possible, and (2) all secret values should be revocable. Resources within the tamper-resistant boundary of an embedded crypto device are usually tightly restricted. We therefore design our API to use a minimum of memory (instead of storing lists of blacklisted keys, we blacklist sets of keys through their attributes) and cryptographic functionality (just symmetric encryption using an authenticated encryption scheme). We do however require that the devices contain loosely synchronized real-time clocks. This assumption is standard in some domains, e.g. on-vehicle cryptographic units [9], Smart-meters, but is not suitable for every application - we will comment briefly in section 6 on how to adapt our design to devices without clocks.

Our starting point is a generalisation of the symmetric key management API of Cortier and Steel [4], which defines a simple hierarchy on keys under management whereby keys higher in the hierarchy are allowed to encrypt (wrap) keys lower down, for storage outside the device or as part of a key exchange protocol. We will describe the API fully in the next section. Here we note some design constraints that are imposed on us by our key revocation requirements, since these may be of independent interest to future designers of revocation APIs.

Firstly, since we can have no single upper bound to the key hierarchy in order to satisfy requirement (2), our design must incorporate a set of keys of the same level *Max* which can revoke each other. The Sevecom API mentioned in section 1 tackles this by having two root keys, where knowledge of the private halves of both is required in order to revoke and update any one of them. The update protocol is given in figure 1, where the terms to the left of the semi-colon represent protocol messages, and the terms to the right are predicates on the state of the device. Initially, the device is in a ‘two key state’, with root keys k_1 and k_2 available for use. Then a revoke message is given for key k_1 , which consists of k_1 signed under its own private half $inv(k_1)$. The device receives this and deletes k_1 , moving into a ‘one key state’ where no further revocation messages will be processed. Then an update message is given, using k_2 to sign the new root key k_3 .

RevokeRootKey1 :

$$\{k_1\}_{inv(k_1)}; keys(k_1, k_2) \rightarrow ; keys(_, k_2)$$

UpdateRootKey1 :

$$\{k_3\}_{inv(k_2)}; keys(_, k_2) \rightarrow ; keys(k_2, k_3)$$

Figure 1: Sevecom Revocation Protocol (simplified)

At first it may seem that this protocol achieves its security goal: unless the intruder manages to obtain access to both $inv(k_1)$ and $inv(k_2)$, he cannot replace a key with one of his own. However, analysis by Mödersheim and

Modesti showed that there is another attack scenario [12]: suppose an attacker had corrupted only $inv(k_2)$, and then waited until the server sent a revocation message for k_1 . He could then intervene and inject his own root key k_i using $inv(k_2)$ to sign it. In a footnote, Mödersheim and Modesti propose that both signature keys be used to sign both messages, and the current public keys be included in the message to prevent replays of old updates. They did not verify this solution however. In our scheme we generalise the idea to a scheme where some number N_{Max} of the K level Max keys are required in order to revoke and replace one. We similarly require the revocation message to demonstrate knowledge of the current key in order to avoid replays, but we do this by encrypting under the old key rather than adding it to the plaintext, to prevent possible key cycles. Since the old key may be corrupted, it is the innermost encryption. A revocation message, e.g. ‘replace k_i with k_j ’, will therefore have the following form:

$$\{\text{update}, k_j\}_{k_i k_1 \dots k_n}$$

In our setting, we only have symmetric key cryptography available, so unlike the public keys in Sevecom, our root keys must remain confidential. This imposes an upper limit to the security we can achieve: if at any point in the future, it is possible for the attacker to obtain all the keys k_1, \dots, k_n , then he will obtain also k_j , and by repeating this for all the subsequent key updates be able to obtain N_{Max} of the current level Max keys. To prevent this, for our security proofs, we demand that ‘honest’ updates to the level Max keys, i.e. updates generated by the server, remain unknown to the attacker. In practice this could mean that the level Max keys are only updated when the device is connected to a trusted host. If this is too cumbersome for the application then a quantitative risk assessment would have to be undertaken to set the N and K high enough to achieve the required degree of security. A full solution would require asymmetric cryptography, as we will discuss in section 6. Note that we will still prove that the device will resist attacks against fake update messages constructed by the intruder, even when he has corrupted individual level Max keys.

For keys of lower level than Max, we assume that the attacker can see all (encrypted) update messages. We need some way to ensure freshness of messages containing keys encrypted by other keys, otherwise revocation will be ineffective. To see this, consider the following example: Assume $\{\dots, k_3, \dots\}_{k_5}$ and $\{\dots, k_3, \dots\}_{k_4}$ have been sent out through normal key encryption commands, with k_i at level $i(< \text{Max})$ in the hierarchy. Assume that at some point k_4 is lost to the attacker. Then he learns k_3 as well (by decrypting the second message). Our revocation API would allow us to remove k_4 , and if we are smart we will also remove k_3 since it must be assumed to be lost as well. However, this will not suffice: the attacker could replay the key distribution message $\{\dots, k_3, \dots\}_{k_5}$, and therefore re-inject the corrupted k_3 into the device. To avoid this one could blacklist all keys below a corrupted key, which is problematic (this blacklist will take up device memory and will have to be kept indefinitely), or revoke all the keys in the hierarchy all the way to the top (impractical), or assure some kind of freshness of key distribution messages. We choose the latter option, requiring that every key is given a validity time when it is issued. After the expiry date passes, the device will refuse to use the key. Without this prop-

erty, we would need some other way to ensure freshness of all messages, such as challenge-response exchanges for every message.

Finally, note that though we design our API to use memory efficiently, we do not explicitly analyse its resistance to denial of service attacks.

3 Description of the API

We describe the design and commands for our key management and key revocation API.

3.1 Setting

We assume that keys are structured in a hierarchy, such that a key k_1 may encrypt a key k_2 only if k_1 is greater than k_2 . More precisely, we assume a set of levels \mathcal{L} equipped with a (partial) order $<$, a maximal element \mathbf{Max} and minimal element 0 . Management keys used for revocation or updates will be keys of maximum level \mathbf{Max} , and knowledge of at least $N_{\mathbf{Max}}$ of them will be required for many operations. We also assume that a level can only be compared to finite number of levels. More precisely, we assume that for any $l \in \mathcal{L}$, the set $\{l' \mid l' < l\}$ is finite. This will ensure that we have no infinite sequence of the form $\{\dots k_1 \dots\}_{k_2}, \{\dots k_2 \dots\}_{k_3}, \dots$

We assume that each tamper resistant device (TRD) a has:

- a clock, whose current time is given by $t_a \in \mathcal{T}$ where \mathcal{T} is an infinite ordered set of time events. For example, \mathcal{T} may be \mathbb{R}^+ , the set of non negative real numbers. For simplicity, we will assume that all clocks are synchronized with a global clock, referred to as the time of the global system. Our security properties could be adapted to take account of clock drift, provided some limit on the drift is assured.
- a table Θ_a of keys. Each entry in the table is indexed by a *handle* h and the corresponding entry $\Theta_a(h)$ is (k, l, v, m) where k is the actual key stored on the TRD, l is its corresponding level, v its *validity date* of the key, and $m \in \mathcal{M}$ is a *miscellaneous field* that may describe some other attributes of the key (e.g. describing the purpose of the key).
- a blacklist \mathfrak{B}_a of elements of the form (l, t) where l is a level and t is an expiration time. Intuitively, whenever (l, t) occurs in a blacklist, it means that the TRD will never accept a key of level l (or below), unless time t has now passed.

Intuitively, the design of our API will ensure that a key may only encrypt other keys lower in the hierarchy whose validity dates have not expired.

We also assume that the TRDs share a function $\delta : \mathcal{L} \rightarrow \mathcal{T}$, that associates lifetime to keys depending on their levels. We may sometimes abbreviate $\delta(l)$ by δ_l . Since we have assumed that a level can only be compared to finite number of levels, we can compute the maximum lifetime of a chain of levels, smaller than a given level. Formally, we consider the function

$$\Delta : l \mapsto \max_{l_1 < \dots < l_n < l, n \in \mathbb{N}} \sum_{i=1}^n \delta(l_i)$$

Intuitively, $\Delta(l)$ is the time where compromising a key k of level l may compromise keys of lower levels, even if the validity time of k has expired, due to chains of encryptions (see Section 5).

For initialisation, we assume that each device contains at least N_{Max} keys of level Max which are known to an administrator. From this, using the commands of the API, the administrator can bootstrap the system and, if necessary, update all the level Max keys.

3.2 Commands

We first give the standard commands which do not concern revocation or keys of level Max , but suffice for normal key management operations. They are mostly the same as in the original API, [4], where they were shown to be sufficient to implement a number of key establishment protocols while always keeping sensitive keys in the secure memory of the TRD, never exposing them in the clear on the host machine. We have generalised the ordering on the hierarchy and introduced checking of the validity time of keys and a blacklist \mathfrak{B}_a of key levels together with expiration times. We will write $l \in \mathfrak{B}_a$ if there is $(l', t) \in \mathfrak{B}_a$ with $l \leq l'$ and t a valid time. We also assume a test $\text{Distinct}(h_1, \dots, h_n)$ which checks that the h_i are pairwise distinct.

Public generation

The generation command for public (level 0) data (e.g a fresh public nonce), which possibly takes an argument $m \in \mathcal{M}$, is defined as follows. It returns both the value and a handle to the value as stored on the device.

```
generatePublic( $m$ )
  let  $h = \text{Fresh}(H_a)$  in
   $H_a := H_a \cup \{h\}$ 
  let  $n = \text{Fresh}(N)$  in
   $N := N \cup \{n\}$ 
  let  $v = t + \delta(0)$  in
   $\Theta_a := \Theta_a \cup \{h \rightarrow (n, 0, v, m)\}$ 
  return  $h, n$ 
```

where $\text{Fresh}(E)$ (respectively with $E \in \{H_a, N, K\}$) returns an element of the set $\mathcal{E} \setminus E$ (respectively with $\mathcal{E} \in \{\mathcal{H}_a, \mathcal{N}, \mathcal{K}\}$).

Secret generation

The generation command for a secret key, with level of security l and, possibly, an argument $m \in \mathcal{M}$:

```
generateSecret( $l, m$ )
  if  $0 < l < \text{Max}$ 
  let  $h = \text{Fresh}(H_a)$  in
   $H_a := H_a \cup \{h\}$ 
  let  $k = \text{Fresh}(K)$  in
   $K := K \cup \{k\}$ 
  let  $v = t_a + \delta(l)$  in
```

```

 $\Theta_a := \Theta_a \cup \{h \rightarrow (k, l, v, m)\}$ 
return  $h$ 

```

Encryption

The encryption command takes as input a list of data that are meant to be encrypted (which may include handles pointing to values stored on the device) and a handle for the key k that will be used for encryption. We check that k has not expired and is not below a level which has been blacklisted. For each term to be encrypted, we check the level is lower than that of k , the expiration date is valid, and that the level is not blacklisted. Note that **break** aborts the entire command.

```

encrypt( $[X_1, \dots, X_n], h$ )
  let  $(k, l, v, m) = \Theta_a(h)$  in
  if  $(l = \text{Max}) \vee (v \leq t_a) \vee (l \in \mathfrak{B}_a)$ 
    break
  for  $i = 1..n$ 
    if  $X_i = M_i, m_i$  /* message */
       $Y_i := (M_i, 0, t_a + \delta(0), m_i)$ 
    if  $X_i = h_i$  /* handle */
      let  $(k_i, l_i, v_i, m_i) = \Theta_a(h_i)$  in
      if  $(l_i < l) \wedge (v_i > t_a) \wedge (l_i \notin \mathfrak{B}_a)$ 
         $Y_i := (k_i, l_i, v_i, m_i)$ 
      else break
  return  $\{Y_1, \dots, Y_n\}_k$ 

```

Decryption

The decryption command takes as inputs a handle for the key that will be used for decryption and a cipher-text. We also assume that decrypt throws **break** on failure of authentication (i.e. we assume an authenticated encryption scheme). Note that the checks on the levels etc. performed during encryption are repeated. This is important for security in the presence of corrupted keys.

```

decrypt( $C, h$ )
  let  $(k, l, v, m) = \Theta_a(h)$  in
  if  $(l = \text{Max}) \vee (v \leq t_a) \vee (l \in \mathfrak{B}_a)$ 
    break
  let  $X_1, \dots, X_n = \text{dec}(k, C)$  in
  for  $i = 1..n$ 
    let  $(k_i, l_i, v_i, m_i) = X_i$  in
    if  $(l_i = 0) \wedge (t_a < v_i \leq t_a + \delta(l_i)) \wedge (l_i \notin \mathfrak{B}_a)$ 
       $Y_i := k_i, m_i$  /* message */
    elseif  $(l_i < l) \wedge (t_a < v_i \leq t_a + \delta(l_i)) \wedge (l_i \notin \mathfrak{B}_a)$ 
      let  $h_i = \text{Fresh}(H_a)$  in
       $H_a := H_a \cup \{h_i\}$ 
       $Y_i := h_i$  /* handle */
       $\Theta_a := \Theta_a \cup \{h_i \rightarrow (k_i, l_i, v_i, m_i)\}$ 
    else break

```

```
return  $Y_1, \dots, Y_n$ 
```

3.3 Management of Revocation keys

We now introduce the commands for managing keys of level **Max**, which will only need to be used if a compromise to one of these keys is suspected. These keys, called *revocation keys*, can be used either to revoke and update lower level keys (keys with a smaller level than **Max**), or to revoke and update the revocation keys. At least N_{Max} revocation keys must be given to revoke or update a revocation key.

Update Max

If we cannot trust a certain revocation key anymore, the administrator can update it using the `updateMax` function. It takes as inputs a cipher-text of the key to be updated and its update encrypted with n revocation keys k_1, \dots, k_n such that $n \geq N_{\text{Max}}$ and takes also the corresponding handles where these keys are stored, h_1, \dots, h_n . `updateMax` commands will be assumed to be sent on secure channels to avoid that, if an attacker breaks very old **Max** keys, he could immediately deduce the current active **Max** keys. Actually, we do not need all `updateMax` commands to be runned under secure channels but simply that this occurs sufficiently regularly.

```
updateMax( $C, h_1, \dots, h_n$ )
  for  $i = 1..n$ 
    let  $(k_i, \text{Max}, v_i, m_i) = \Theta_a(h_i)$  in
    if  $\exists j \in \llbracket 1, n \rrbracket$  s.t.  $v_j \leq t_a \vee \neg \text{Distinct}(h_1, \dots, h_n)$ 
      break
    let  $(\text{updateMax}, k', v'_k, m'_k) = \text{dec}(k_1, \dots, \text{dec}(k_n, C))$  in
    if  $(v'_k \leq t_a) \vee (v'_k > t_a + \delta_{\text{Max}})$ 
      break
     $\Theta_a(h_1) := (k', \text{Max}, v'_k, m'_k)$ 
```

Note that the `updateMax` command does not introduce key cycles and avoid replays.

3.4 Management of Working keys

We call any key of level $l : l < \text{Max}$ a *working key*. We include separate commands for creating working keys on the device and for updating or revoking them, each of which require a N_{Max} level **Max** keys. These commands are in addition to the usual operational key management functions arising from the encryption and decryption commands given in section 3.2: they are intended to be used for bootstrapping the system or for removing and updating possibly compromised keys. In the descriptions below, $n \geq N_{\text{Max}}$.

Create

The create function takes as inputs a cipher-text containing the “order” to create and the different data to create, the whole encrypted with n revocation keys

k_1, \dots, k_n and takes also the corresponding handles where these keys are stored, h_1, \dots, h_n .

```

create( $C, h_1, \dots, h_n$ )
  for  $i = 1..n$ 
    let  $(k_i, \text{Max}, v'_i, m'_i) = \Theta_a(h_i)$  in
    if  $\exists j \in \llbracket 1, n \rrbracket$  s.t.  $v'_j \leq t_a \vee \neg \text{Distinct}(h_1, \dots, h_n)$ 
      break
    let  $C' = \text{dec}(k_1, \dots, \text{dec}(k_n, C))$  in
    let  $(\text{create}, x_1, l_1, v_1, m_1, \dots, x_p, l_p, v_p, m_p) = C'$  in
    if  $\exists j$  s.t.  $(l_j \geq \text{Max}) \vee (v_j \leq t_a) \vee$ 
       $(v_j > t_a + \delta(l_j)) \vee (l_j \in \mathfrak{B}_a)$ 
      break
    for  $j = 1..p$ 
      let  $h'_i = \text{Fresh}(H_a)$  in
       $H_a := H_a \cup \{h'_i\}$ 
       $\Theta(h'_i) := (x_i, l_i, v_i, m_i)$ 
    return  $h'_1, \dots, h'_p$ 

```

Update

The update function takes as inputs a cipher-text containing the “order” to update, the different values to change and their updates, the whole encrypted with n revocation keys k_1, \dots, k_n and takes also the corresponding handles where these keys are stored, h_1, \dots, h_n .

```

update( $C, h_1, \dots, h_n$ )
  for  $i = 1..n$ 
    let  $(k_i, \text{Max}, v_i, m_i) = \Theta_a(h_i)$  in
    if  $\exists j \in \llbracket 1, n \rrbracket$  s.t.  $v_j \leq t_a \vee \neg \text{Distinct}(h_1, \dots, h_n)$ 
      break
    let  $C' = \text{dec}(k_1, \dots, \text{dec}(k_n, C))$  in
    let  $(\text{update}, x_1, x'_1, l'_1, v'_1, m'_1, \dots, x_p, x'_p, \dots, m'_p) = C'$  in
    for  $j = 1..p$ 
      for  $h \in H_a$  s.t.  $\Theta(h) = (x_j, l, v, m)$ 
        if  $(l < \text{Max}) \wedge (l'_j = l) \wedge (t_a < v'_j \leq t_a + \delta(l_j))$ 
           $\wedge (l'_j \notin \mathfrak{B}_a)$ 
             $\Theta(h) := (x'_j, l'_j, v'_j, m'_j)$ 

```

Revoke

We first define the revoke command in its most general form. The administrator may wish to revoke keys, for example if he suspects them to be compromised. He may wish to revoke a precise set of keys but he may also wish to revoke them on the basis of their attributes, e.g. validity time. To be as flexible as possible, we consider a revoke command that is parametrized by a function F such that :

$$F : \mathcal{L} \times \mathcal{T} \times \mathcal{M} \rightarrow \{\perp, \top\}$$

which defines, according to some criteria (e.g. level, validity, etc) what is going to be kept or deleted.

The revoke function takes as input such a function F , encrypted with n revocation keys, and the corresponding handles h_1, \dots, h_n .

```

revoke( $C, h_1, \dots, h_n$ )
  for  $i = 1..n$ 
    let  $(k_i, \text{Max}, v_i, m_i) = \Theta_a(h_i)$  in
    if  $\exists j \in \llbracket 1, n \rrbracket$  s.t.  $v_j \leq t_a \vee \neg \text{Distinct}(h_1, \dots, h_n)$ 
      break
    let  $(\text{revoke}, F) = \text{dec}(k_1, \dots, \text{dec}(k_n, C))$  in
    for  $h \in H_a$  s.t.  $\Theta_a(h) = (x, i, v, m)$ 
      if  $F(i, v, m) = \top$ 
         $\Theta_a := \Theta_a \setminus \{h \mapsto (x, i, v, m)\}$ 

```

Note that, in practice an implementation should offer some specialized revoke functions for the particular application. Indeed, sending an arbitrary Boolean function as a parameter would raise both implementation and security issues. However, by defining F in this general way we ensure that our security results hold for any choice of F .

Blacklist

The Blacklist function takes as input a cipher-text containing the levels to blacklist together with expiration times, encrypted with n revocation keys k_1, \dots, k_n . The effect of the command is to add the levels to the blacklist and erase all the keys of the corresponding levels, including smaller ones. Blacklisting the level instead of all the keys saves memory. It also means the administrator does not need to retain the actual values of all the working keys.

```

blacklist( $C, h_1, \dots, h_n$ )
  for  $i = 1..n$ 
    let  $(k_i, \text{Max}, v_i, m_i) = \Theta_a(h_i)$  in
    if  $\exists j \in \llbracket 1, n \rrbracket$  s.t.  $v_j \leq t_a \vee \neg \text{Distinct}(h_1, \dots, h_n)$ 
      break
    let  $(\text{blacklist}, (l_1, t_1), \dots, (l_p, t_p)) =$ 
       $\text{dec}(k_1, \dots, \text{dec}(k_n, C))$  in
    for  $i = 1..p$ 
      /* add to blacklist */
       $\mathfrak{B}_a := \mathfrak{B}_a \cup \{(l_i, t_i)\}$ 
      /* erase affected keys */
      for  $h \in H_a$  s.t.  $\Theta_a(h) = (x', l', v', m')$ 
        if  $l' \leq l_i$ 
           $\Theta_a := \Theta_a \setminus \{h \mapsto (x', l', v', m')\}$ 

```

3.5 Example

Let us consider two TRDs a and b (see Figure 2) initialized with handles containing keys of level Max k and k' and sharing key k_1 of level l_1 . In a first step, the user of device a generates, using the `generateSecret` command, a secret key k_2 with a level l_2 (such that $l_2 < l_1$) and a validity v_2 , this new secret appears in TRD a stored under a new handle. To share the key with b , a encrypts k_2 under k_1 using `encrypt`. When the message is received by the user of b he uses the `decrypt` command which will store the new key k_2 under a new handle.

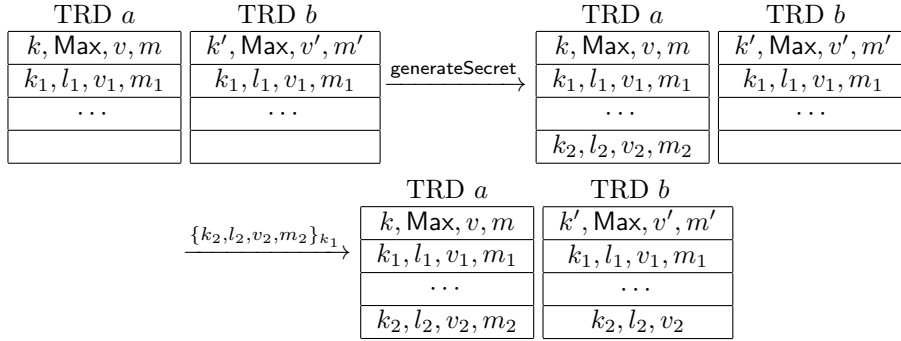


Figure 2: A simple execution. A key is created on device a and shared with device b .

3.6 Threat Scenario

We consider a scenario where the attacker:

- controls the network (he may read and send messages on the network),
- controls the host machines on which the TRD are connected; therefore he can execute API commands on any TRD
- may break some arbitrary keys of his choice, typically by brute-forcing some keys or employing side-channel attacks.

In the remainder of the paper we show that, provided the attacker does not break too many keys of level Max (an attacker should not break more than $N_{\text{Max}} - 1$ keys simultaneously stored on a TRD), then our system can self-repair given time or one can explicitly fix a TRD. More precisely, we show the two following results:

- If some key of level $l \neq \text{Max}$ is lost then after some time (actually, at time $v + \delta_l$ where v is the validity date of the lost key) the keys of level l are secure again.
- Moreover, if some key of level $l \neq \text{Max}$ is lost and a TRD receives a command blacklisting the level l , then all the keys on the TRD are secure again.

Note that our analysis will not deal with denial of service attacks.

4 Formal Model

We study the security offered by our model in a symbolic model, where messages are represented by terms, following a now standard approach used for many protocols (see e.g. [1, 2, 11]).

4.1 Syntax

We assume a finite set of names \mathcal{A} and two infinite sets \mathcal{N} and \mathcal{K} respectively for nonces and keys. We also assume a finite set \mathcal{H} representing handles and an infinite set \mathcal{M} representing the miscellaneous fields, with $\epsilon \in \mathcal{M}$ representing the empty element. We recall that \mathcal{L} denotes the set of levels. Messages are represented using a term algebra **Terms** defined by the following grammar:

$$T, T_1, T_2, \dots := a \mid n \mid k \mid l \mid t \mid m \mid \{T\}_k \mid \langle T_1, T_2 \rangle$$

where $a \in \mathcal{A}$, $n \in \mathcal{N}$, $k \in \mathcal{K}$, $l \in \mathcal{L}$, $t \in \mathcal{T}$, $m \in \mathcal{M}$. The term $\{t\}_k$ represents the message t (symmetrically) encrypted by the key k while the term $\langle t_1, t_2 \rangle$ represents the concatenation (or more precisely, the pairing) of the two messages t_1 and t_2 . For simplicity, we will often write t_1, t_2, \dots, t_n instead of $\langle t_1, \langle t_2, \dots \langle t_{n-1}, t_n \rangle \dots \rangle$. The notion of *subterm* is defined as usual: t' is a subterm of t if t' occurs at some position p in t , that is $t|_p = t'$. We denote by $\text{St}(m)$ the set of subterms of m and by extension $\text{St}(S)$ is the set of subterms of terms in S .

A *global state* of our system is described by a tuple $(\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ where $\mathcal{P} \subseteq \mathcal{A}$ is the set of TRDs in the system, $t \in \mathcal{T}$ represents the current time, $\mathfrak{M} \subseteq \text{Terms}$ represents the set of messages sent so far over the network, N and K are respectively sets of currently used nonces and keys¹ by all the APIs and \mathcal{I} is a function:

$$\mathcal{I} : a \mapsto (\Theta_a, H_a, \mathfrak{B}_a, t_a, N_a, K_a)$$

that represents the current local state of the TRD $a \in \mathcal{P}$. More precisely, $H_a \subseteq \mathcal{H}$ represents the finite set of handles currently used in the API a and $\mathfrak{B}_a \subseteq \mathcal{L} \times \mathcal{T}$ represents the set of blacklisted levels for which the TRD does not accept keys anymore. $N_a \subseteq N$ and $K_a \subseteq K$ are respectively the nonces and keys that have been generated or stored on the API. Θ_a represents the key table of the TRD. Formally, it is a function

$$\Theta_a : H_a \rightarrow (\mathcal{K} \cup \mathcal{N}) \times \mathcal{L} \times \mathcal{T} \times \mathcal{M}$$

Indeed, as seen in Section 3.1, each handle $h \in H_a$ points to an entry (x, l, v, m) corresponding a nonce or a key x and its attributes: the level l , its validity v , and other miscellaneous information m .

As indicated in the definition of a global state, all keys come with a level.

Definition 1. Let $(\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ be a global state and let k be a key in K . The level of k is defined as follows :

$$\text{Level}(k) = \{l \mid \exists a \in \mathcal{P}, h \in H_a \text{ s.t. } \Theta_a(h) = (k, l, v, m)\}$$

It is *a priori* possible for a key k to have several levels (i.e. $\text{Level}(k)$ is not a singleton) but we will show in our proofs that this never happens for uncompromised keys. We will therefore say that a key k is of level l if $l \in \text{Level}(k)$.

¹ N and K are here as artifacts of the model. Since an API gets a random key (or nonce) when it creates it, there is very little chance that it generates a key (or nonce) which has already been used. To capture this, we model a global knowledge of all nonces and keys used by all APIs to be sure that freshly generated keys (or nonces) are new.

$$\begin{array}{c}
\frac{t \in \mathfrak{M}}{\mathfrak{M} \vdash t} \quad \frac{\mathfrak{M} \vdash x, \mathfrak{M} \vdash y}{\mathfrak{M} \vdash \langle x, y \rangle} \quad \frac{\mathfrak{M} \vdash \langle x, y \rangle}{\mathfrak{M} \vdash x} \\
\\
\frac{\mathfrak{M} \vdash \langle x, y \rangle}{\mathfrak{M} \vdash y} \quad \frac{\mathfrak{M} \vdash x, \mathfrak{M} \vdash y}{\mathfrak{M} \vdash \{x\}_y} \quad \frac{\mathfrak{M} \vdash \{x\}_t, \mathfrak{M} \vdash y}{\mathfrak{M} \vdash x}
\end{array}$$

Figure 3: Deduction Rules (Intruder)

4.2 Semantics

The behavior of our API is modeled by the transitions of our formal system. We define four kinds of transitions: transition of time, forgery by the attacker, silent transitions, and key management commands.

Transition of time We assume that the execution of API commands is instantaneous (this could be adapted to take into account worst case execution time). The effect of the time is modelled by a separate and explicit TIM transition:

$$(TIM) \quad (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t) \xrightarrow{(\text{Time passes})} (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t')$$

with $t' > t$.

Forgery by the attacker As is usually the case in formal models, the ability of the attacker to construct, or *deduce*, is modeled by a relation \vdash . We write $S \vdash m$ if m can be deduced from S . The deduction relation is formally defined in Figure 3. Intuitively, the attacker can deduce any term that can be obtained by pairing, encrypting, projecting, and decrypting whenever it has the encryption key.

Example 1. Let $S = \{\{k_1\}_{k_2}, \langle k_2, k_3 \rangle, \{k_4\}_{k_5}\}$. Then $S \vdash k_1$, $S \vdash k_2$, $S \vdash k_3$, but $S \not\vdash k_4$, $S \not\vdash k_5$. We also have $S \vdash \langle k_2, k_1 \rangle$, but $S \not\vdash \langle k_2, k_4 \rangle$.

The attacker may send any deducible message over the network. This is reflected by the following transition.

$$(DED) \quad (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t) \xrightarrow{(\text{Deduction})} (\mathcal{P}, \mathcal{I}, \mathfrak{M} \cup \{m\}, N, K, t)$$

provided $\mathfrak{M} \vdash m$.

Silent transitions Each API command executed on a TRD a with input m defines a transition $\Theta_a, H_a, \mathfrak{B}_a \rightarrow \Theta'_a, H'_a, \mathfrak{B}'_a$. Moreover, the TRD outputs some message m' . Given a global state $(\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ such that $m \in \mathfrak{M}$, the corresponding global transition is

$$(\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t) \rightarrow (\mathcal{P}, \mathcal{I}', \mathfrak{M} \cup \{m'\}, N \cup N', K \cup K', t)$$

with N' and K' being respectively the sets of nonces and keys generated by a during the execution of the command. The function \mathcal{I}' is defined by $\mathcal{I}' : a \mapsto (\Theta'_a, H'_a, \mathfrak{B}'_a, t_a)$ for the corresponding API a and $\mathcal{I}'(a') = \mathcal{I}(a')$ otherwise.

$$\begin{aligned}
(\text{UPM}) \quad & (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t) \xrightarrow{\text{updateMax}} (\mathcal{P}, \mathcal{I}', \mathfrak{M}, N', K', t) \\
(\text{NEW}) \quad & (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t) \xrightarrow{\text{create}} (\mathcal{P}, \mathcal{I}', \mathfrak{M} \cup \{m\}, N', K', t) \\
& \text{with } m = \{\text{create}, N, N', x_1, l_1, v_1, m_1 \dots, m_p\}_{k_1 \dots k_n} \\
(\text{UPD}) \quad & (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t) \xrightarrow{\text{update}} (\mathcal{P}, \mathcal{I}', \mathfrak{M} \cup \{m\}, N', K', t) \\
& \text{with } m' = \{\text{update}, x_1, x'_1, l'_1, v'_1, m'_1 \dots, m'_p\}_{k_1 \dots k_n} \\
(\text{RVK}) \quad & (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t) \xrightarrow{\text{revoke}} (\mathcal{P}, \mathcal{I}', \mathfrak{M} \cup \{m\}, N, K, t) \\
& \text{with } m = \{\text{revoke}, F\}_{k_1 \dots k_n} \\
(\text{BLK}) \quad & (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t) \xrightarrow{\text{blacklist}} (\mathcal{P}, \mathcal{I}', \mathfrak{M} \cup \{m\}, N, K, t) \\
& \text{with } m = \{\text{blacklist}, (l_1, t_1), \dots, (l_n, t_n)\}_{k_1 \dots k_n}
\end{aligned}$$

For each transition, the function \mathcal{I}' and the sets N' and K' are updated as specified by the corresponding command, as for silent transitions.

Figure 4: Key management command transitions

Key management commands As explained in Section 2, at least some revocation commands need to be kept hidden from the attacker. Otherwise, assuming that the attacker controls and memorizes all the traffic on the network and assuming he can break keys (which corresponds to our threat scenario), then breaking N_{Max} old keys may compromise the entire system, even if the compromised keys had not been in use for some time.

We therefore assume that the key management commands for keys of level **Max** are sent over a private channel. In practice, this could be achieved by several means. For example, we may assume that the key administrator has a physical access to the TRD that needs to be updated. Or we may also assume that the user would connect his/her TRD to a trusted machine, on which a secure channel (e.g. via TLS) is established with the key administrator. Note that the key management commands for keys of level **Max** are executed only when a key of level **Max** is lost (or suspected to be lost), or when keys of level **Max** are updated (e.g. when their validity has expired). We expect these events to occur infrequently.

Note that our assumption does not prevent an adversary from trying to run the UPM command. In particular, in case it has sufficiently many keys of level **Max**, it may well build a well formed Update Max command and send it to a TRD. This is reflected in the silent transitions defined above.

All the other key management commands can be sent over an insecure network. The corresponding transition system is presented in Figure 4. The fact that we assume the UPM command is sent over a secure channel is reflected in the fact that the set \mathfrak{M} remains unchanged for this transition.

In the remainder of the paper we assume that managers in charge of generating key management commands behave consistently. More precisely, we assume that they only use fresh keys when updating and creating keys and that they never give different attributes to the same key value.

Key compromise We model the fact that the attacker may compromise a key by adding a transition lost, which allows an attacker to obtain a key of his choice.

$$(\text{LST}) \quad (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t) \xrightarrow{\text{Lost}(k)} (\mathcal{P}, \mathcal{I}, \mathfrak{M} \cup \{k\}, N, K, t)$$

where k is a key that appears on at least one TRD, that is k occurs in the image of \mathcal{I} .

4.3 Example

Let us reconsider the example of Figure 2. Let E be the initial global state, we have $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$. To simplify, let us consider that $\mathcal{P} = \{a, b\}$, $K = \{k, k', k_1\}$, t and \mathfrak{M} are arbitrary. Then $\mathcal{I} = \{a \mapsto (\Theta_a, H_a, \mathfrak{B}_a, t_a, N_a, K_a), b \mapsto (\Theta_b, H_b, \mathfrak{B}_b, t_b, N_b, K_b)\}$ where $t_a = t_b = t$, $H_a = \{h_1, h_2\}$, $H_b = \{h'_1, h'_2\}$, $K_a = K_b = K$. We have $\Theta_a(h_1) = (k, \text{Max}, v, m)$, $\Theta_b(h'_1) = (k', \text{Max}, v', m')$ and $\Theta_a(h_2) = \Theta_b(h'_2) = (k_1, l_1, v_1, m_1)$.

Using the `generationSecret` command with TRD a will lead to a global state $E' = (\mathcal{P}, \mathcal{I}', \mathfrak{M}, N, K', t)$ where $\mathcal{I}' = \{a \mapsto (\Theta'_a, H'_a, \mathfrak{B}_a, t_a, N_a, K'_a), b \mapsto (\Theta_b, H_b, \mathfrak{B}_b, t_b, N_b, K_b)\}$ with $K'_a = K' = K \cup \{k_2\}$, $H'_a = H_a \cup \{h_2\}$ and $\Theta'_a = \Theta_a \cup \{h_3 \mapsto (k_2, l_2, v_2, m_2)\}$.

Then the `encrypt` command leads to a global state $E'' = (\mathcal{P}, \mathcal{I}, \mathfrak{M}', N, K, t)$ where $\mathfrak{M}' = \mathfrak{M} \cup \{(k_2, l_2, v_2)_{k_1}\}$. Finally the `decrypt` command is used by TRD b on the message $\{k_2, l_2, v_2\}_{k_1}$. That leads to $E''' = (\mathcal{P}, \mathcal{I}''', \mathfrak{M}, N, K, t)$ with $\mathcal{I}''' = \{a \mapsto (\Theta'_a, H'_a, \mathfrak{B}_a, t_a, N_a, K'_a), b \mapsto (\Theta'''_b, H'''_b, \mathfrak{B}_b, t_b, N_b, K'''_b)\}$ with $K'''_b = K_b \cup \{k_2\}$, $H'''_b = H_b \cup \{h'_3\}$ and $\Theta'''_b = \Theta_b \cup \{h'_3 \mapsto (k_2, l_2, v_2, m_2)\}$.

5 Security

We are now ready to formally state the security properties discussed in section 3.6. Intuitively, we would like to show that apart from explicitly lost keys, any key that has not expired is secure, that is, cannot be deduced by the attacker. However, this is not strictly the case: several commands such as `encrypt` or the administrator commands produce messages of the form $\{\dots k' \dots\}_k$ that are sent over the network and possibly stored by the attacker. It could be the case that k expires before k' . Clearly, if k becomes known to the attacker then k' will be immediately known as well, even if k has expired and k' has not.

Therefore, we distinguish three kinds of key: *valid* keys, *latent* keys, and *dead* keys. Informally, valid keys are those stored on a device whose validity time has not passed, latent keys are those which are not valid but on which security of some encryption of a valid lower level key might still depend, and dead keys are everything else.

Definition 2. Let $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ be a global state and $k \in K$. Given $a \in \mathcal{P}$, we denote by Θ_a the key table stored on a and defined by $\mathcal{I}(a) = (\Theta_a, \dots)$.

- k is a **valid** key if $\exists a \in \mathcal{P}, \exists h \in H_a$ s.t. $\Theta_a(h) = (k, l, v, m)$ with $v > t_a$
or if there exists $\{\dots, k, l, v, m, \dots\} \in St(\mathfrak{M})$ with $v > t$.
We denote by \mathcal{V}_E the set of valid keys in E .

- k is a **latent** key if $\exists a \in \mathcal{P}, \exists h \in H_a$ s.t. $\Theta_a(h) = (k, l, v, m)$ with $v + \Delta(l) > t$ or if there exists $\{\dots, k, l, v, m, \dots\} \in St(\mathfrak{M})$ with $v + \Delta(l) > t$.
We denote by \mathcal{U}_E the set of latent keys in E . The notation \mathcal{U}_E comes from the fact that these keys may also be viewed as “undead”.
- Otherwise, k is a **dead** key.

A key is latent if it expired but less than $\Delta(l)$ ago. The reason why latent keys still influence the security of the whole system can be illustrated with the following scenario. Let k be a key of level l and validity time v and consider $l_1 < \dots < l_n < l$ such that $\Delta(l) = \sum_{i=1}^n \delta(l_i)$. Then it may be the case that k was used to encrypt a key k_n of level l_n just before its expiration time, yielding the message $\{k_n, l_n, v_n, m_n\}_k$ with $v_n = v + \delta(l_n)$. Similarly, each key k_i may have been used to encrypt a key k_{i-1} of level l_{i-1} , yielding a message $\{k_{i-1}, l_{i-1}, v_{i-1}, m_{i-1}\}_{k_i}$ with $v_{i-1} = v_i + \delta(l_{i-1})$. Then it is easy to see that $v_1 = v + \Delta(l)$ and clearly if k is lost at anytime after its expiration date v but before $v + \Delta(l)$ then the security of k_1 (valid) would be immediately compromised as well.

5.1 Initial states

As discussed in section 3.2, before being deployed, a TRD is loaded with some level Max keys so that the administrator can then remotely equip the device with an initial set of working keys.

Definition 3. A global state E_0 is said to be *originating* if $E_0 = (\mathcal{P}, \mathcal{I}_0, \emptyset, 0, \emptyset, \emptyset)$ with $\mathcal{I}_0 : a \mapsto (\Theta_a, H_a, 0, \emptyset, \emptyset)$ and for all $a \in \mathcal{P}$, we have $H_a = \{h_1^a, \dots, h_n^a\}$ with $\Theta_a(h_i^a) = (k_i^a, \text{Max}, v_i^a, m_i^a)$. The k_i^a keys are *revocation keys* initially placed in the memory of each API and should be all distinct: $k_i^a = k_j^a$ implies $i = j$.

5.2 Keys of level Max

The keys of level Max are crucial since they can be used to update and revoke all the keys, including keys of level Max themselves. We first show that security of these keys is preserved, assuming that not too many of the keys of level Max stored on an API are simultaneously lost.

Hypothesis 1. Let $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ be a global state such that $E_0 \rightarrow^* E$ and let \mathcal{L} the set of keys that have been lost so far (i.e. the set of keys k such that the transition $\text{Lost}(k)$ has been executed). We assume that $\forall a \in \mathcal{P}$, we have

$$\#\{k \mid k \in \mathcal{U}_E \cap \mathcal{L} \cap K_a \text{ and } \text{Max} \in \text{Level}(k)\} \leq N_{\text{Max}} - 1$$

where K_a is the set of keys associated to a by the function \mathcal{I} (that is $\mathcal{I}(a) = (\dots, K_a)$).

We assume this hypothesis to hold in all the global states we consider in the remainder of the paper. From a practical point of view, it means that the lifetime of a key of level Max should be set such that during the time when a Max key is latent, it is sufficiently unlikely that an attacker will break more than $N_{\text{Max}} - 1$ such latent keys. More precisely, our assumption is slightly weaker

since the attacker may break more keys provided that no more than $N_{\text{Max}} - 1$ of them have occurred on the same TRD.

The design of our API ensures that no Max keys can be learned by an attacker, except for the explicitly lost ones.

Theorem 1. *Let E be a global state such that $E_0 \rightarrow^* E$, then :*

$$\forall k \in \mathcal{V}_E^{\text{Max}} \setminus \mathcal{L}, E \not\models k.$$

The proof relies on the fact that keys of level Max only appear in key position in the messages sent over the network and the fact that for any Max key, $\text{Level}(k)$ is a singleton. These two properties are shown to be preserved by application of the rules and rely on Hypothesis 1 that ensures that the attacker never knows sufficiently many Max keys to forge an administrator command.

5.3 Lower level keys

The case of keys of lower level is more difficult than the case of Max keys. First, as soon as a key k of level l is lost, then any key of lower level $l' < l$ is compromised as well. Indeed, for any key k' of level $l' < l$ stored on the same TRD as the key k , the attacker may use the `encrypt` command to get k' encrypted by k and therefore deduce k' . This attack does not only compromise the keys stored on TRDs that contain the compromised key k . Indeed, assume that the key k is stored on some TRD a_1 and does not appear in some other TRD a_2 . Then as soon as a_1 and a_2 share some key k'' of level $l'' > l$, the attacker may use the `encrypt` command on a_1 to get k encrypted by k'' . Sending this encryption to a_2 and executing the `decrypt` command on a_2 , the attacker registers k on a_2 and can now gain any key of level lower than l .

Therefore, as soon as a key of level l is lost, any key of level $l' < l$ should be considered as lost too. We introduce the notion of keys *compatible* with a set of levels, which are latent keys that have an uncompromised level.

Definition 4. *Let $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ be a global state, \mathcal{L} a set of keys (typically lost keys), and $k \in K$. Let L_v be a set of levels (that typically corresponds to the levels of keys lost so far). Then k is compatible with L_v if $k \in \mathcal{U}_E \setminus \mathcal{L}$, $\text{Max} \notin \text{Level}(k)$, and there exists $l \in \text{Level}(k)$ such that $l \not\prec L_v$ (that is $l \not\prec l'$ for any $l' \in L_v$).*

We note $\mathcal{G}_{(E, \mathcal{L})}^{L_v}$ the set of keys compatible with L_v in the global state E . We may write $\mathcal{G}_E^{L_v}$ instead of $\mathcal{G}_{(E, \mathcal{L})}^{L_v}$ when \mathcal{L} is clear from the context.

The first property of our API is that keys of levels not marked as lost are uncompromised. We also show that for such keys, the level and all other attributes are unique (i.e. consistent across all TRDs). We actually need to enforce a stronger property, called *robustness*.

Definition 5. *A global state $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ is robust up to a set of levels L_v and a set \mathcal{L} of (lost) keys if the following properties are satisfied:*

- (1) *We consider that (l, v, m) are attributes of k if any of these three cases holds :*
 - (a) $\exists a \in \mathcal{P}, h \in H_a$ s.t. $\Theta_a(h) = (k, l, v, m)$,

- (b) $\exists q \in K, m' \in \mathfrak{M}$ s.t. $\text{Max} \in \text{Level}(q)$ and $\{\text{cmd}, \dots, k, l, v, m, \dots\}_{q_1 \dots q_j q} \in \text{St}(m')$ with $\text{cmd} \neq \text{blacklist}$ and $\text{Max} \in \text{Level}(q_i)$,
- (c) $\exists q \in K, m' \in \mathfrak{M}$ s.t. $\{\dots, k, l, v, m, \dots\}_q \in \text{St}(m')$.

If (l_1, v_1, m_1) and (l_2, v_2, m_2) are attributes of a key $k \in \mathcal{G}_E^{\text{Lv}}$, then $(l_1, v_1, m_1) = (l_2, v_2, m_2)$.

- (2) $\forall k \in \mathcal{G}_E^{\text{Lv}}, \forall m \in \mathfrak{M}$ s.t. $k \in \text{St}(m)$, then, any occurrences of k is of one of these three forms :

- (i) $\{m'\}_k$,
- (ii) $\{\dots, k, \dots\}_{k'}$ with $k' \in K, v_k \leq v_{k'} + \delta_k$ and $\text{Level}(k) < \text{Level}(k') \neq \text{Max}$,
- (iii) $\{\dots, k, \dots\}_{q_1 \dots q_n k'}$ with $q_1, \dots, q_n, k' \in K$ s.t., for $i \in [1, n]$, $\text{Max} \in \text{Level}(q_i)$, $\text{Max} \in \text{Level}(k')$, $k' \notin \mathcal{L}$ and $v_k \leq v_{k'} + \delta_k$,
- (3) $\text{Lv} \geq_s \{l \mid \exists k \in \mathcal{L} \cap \mathcal{U}_E \text{ s.t. } l \in \text{Level}(k)\}$ where $S_1 \geq_s S_2$ if for any $l_1 \in S_1$, there exists $l_2 \in S_2$ such that $l_1 \geq l_2$.

Robustness is meant to precisely control how compatible keys occur in a global state E . It ensures two main properties. First, compatible keys have a unique attribute in the system, that is, levels and validity times of compatible keys are consistently propagated through the system. Moreover, robustness ensures that compatible keys occur only in key position, except if they are encrypted by a key of greater level or if they occur in an administrator command (thus encrypted by a key of level Max). The third item is a technical condition that ensures that levels corresponding to lost (and not yet dead) keys should always be considered as compromised.

Whenever a state is robust, all its compatible keys remain secret from the attacker.

Proposition 1. Let $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ be a global state, \mathcal{L} a set of keys (typically lost keys), and $k \in K$. Let Lv be a set of levels. Assume that E is robust up to Lv and \mathcal{L} then:

$$\forall k \in \mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}} \quad \mathfrak{M} \not\models k.$$

Our three main results are:

- A global state is always robust up to the levels of lost keys (Theorem 2). This means that losing a key does not compromise any other key apart from those at lower levels than the lost key, i.e. those which are trivially compromised.
- Once a lost key is dead, the corresponding level is protected again (Theorem 3), meaning that even when a key is lost and no special action is taken, the system repairs itself after a certain time.
- When a TRD receives a blacklist command, the corresponding levels are immediately protected again on this TRD (Theorem 4).

Theorem 2. *Let E be a global state such that $E_0 \rightarrow^* E$ and let \mathcal{L} be the set of keys lost along the execution $E_0 \rightarrow^* E$. Then E is robust up to $L_v = \{l \mid \exists \text{Lost}(k) \in \mathcal{L} \text{ and } \text{Level}(k) = l\}$ and \mathcal{L} .*

Theorem 2 is the key theorem of our paper and its proof consists in proving that robustness is invariant under the rules of our API. Theorem 1 ensures that it is impossible for an attacker to break more than $N_{\text{Max}} - 1$ keys of level Max. This, in turn, ensures that all accepted key management commands come from the administrator. Then it remains to show that the robustness property is invariant under application of the other rules of the API. For the deduction rule, the first step is to show that compatible (latent) keys are all non-deducible. For decryption, the robustness property in the global state E ensures that encrypted messages have the right form. Therefore when decrypting a message, the newly stored keys also satisfy the invariant. In case the attacker forges a message, the checks on the levels and the validity times ensure that the newly stored keys are of an incompatible level and therefore their security does not need to be assessed. The arguments for encryption and the other commands are along the same lines.

Theorem 3. *Let E, E' be two global states such that $E_0 \rightarrow^* E \rightarrow^* E'$. Assume that E is robust up to L_v and \mathcal{L} and let $l \in L_v$ be a compromised level. Let \mathcal{L}' be the set of keys lost along the execution $E \rightarrow^* E'$. We further assume that there has been no ‘lost’ event for keys of level greater than l , that is, there is no $k \in \mathcal{L}'$ such that $\text{Level}(k) \geq l$. Let $v = \sup(\{v_k \mid k \in \mathcal{L} \text{ and } \text{Level}(k) = l\})$, where v_k is the validity time associated to k (Theorem 2 ensures v_k is unique). Then:*

$$t' \geq v + \Delta(l) \implies E'' \text{ is robust up to } L_v'' \text{ for } \mathcal{L}'',$$

where t' is the global time of E' and E'' is such that $E' \xrightarrow{\text{Lost}(A)} E''$ with $A = \{k \mid \text{Level}(k) = l' \text{ with } l' < l\}$ and $L_v'' = (L_v \cup \bigcup_{k \in \mathcal{L}''} \text{Level}(k)) \setminus \{l\}$ where $\mathcal{L}'' = \mathcal{L}' \cup A$.

Intuitively, Theorem 3 ensures that whenever a key k has been lost then once k is not latent any more (*i.e.* once k is dead), the system is now robust up to $L_v \setminus \{l\}$ plus the levels $\{l' \mid l' < l\}$ which can still be compromised and the newly lost keys \mathcal{L}' , that is, the keys that might have been lost during the execution between E and E' . Theorem 3 is actually a consequence of Theorem 2, which can be seen by carefully analysing the robustness property and noticing that once the lost key k is dead, the set of keys compatible with L_v'' and L_v coincide (if the keys of lower levels are explicitly lost).

We now state the security gained with the **blacklist** command. This command only fixes the system locally. We therefore need to define robustness locally to a TRD.

Definition 6. *Let E be a global state, \mathcal{L} be a set of lost keys, and L_v be a set of levels. Let $a \in \mathcal{P}$ be a TRD. Then a is locally robust up to L_v and \mathcal{L} in E if $E_a = (\{a\}, \mathcal{I}_a, \mathfrak{M}, N, K, t)$ is robust up to L_v where \mathcal{I}_a is the restriction of \mathcal{I} on $\{a\}$ and where the conditions of robustness are considered for keys that are stored on a (that is, appear in the image of Θ_a).*

Proposition 2 is the analog of Proposition 1 for locally robustness and states that compatible keys on locally robust TRDs are secure.

Proposition 2. *Let $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ be a global state, \mathcal{L} a set of keys (typically lost keys), and $k \in K$. Let L_v be a set of levels. Assume that a is locally robust up to L_v and \mathcal{L} in E then:*

$$\forall k \in \mathcal{G}_{(E, \mathcal{L})}^{L_v} \cap \{k' \mid \exists h \text{ s.t. } \Theta_a(h) = (k', l, v, m)\} \quad \mathfrak{M} \not\models k.$$

Finally, we can state our final theorem that ensures that blacklisting a level protects the corresponding keys on a TRD.

Theorem 4. *Let $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ be a global state s.t. $E_0 \longrightarrow^* E$. Let \mathcal{L} be a set of (lost) keys. Assume that $a \in \mathcal{P}$ is locally robust up to L_v and \mathcal{L} in E . Assume that $E \xrightarrow{L} \longrightarrow^* E'$ with L a blacklist command for TRD a , level l , and validity time t_h (with t_h greater or equal to $v_h = \sup(\{v_k \mid k \in \mathcal{L} \text{ and } l \in \text{Level}(k)\})$). Let \mathcal{L}' be the set of keys lost along the execution $E \xrightarrow{L} \longrightarrow^* E'$. We further assume that there has been no lost event for keys of level greater than l , that is, there is no $k \in \mathcal{L}'$ such that $\text{Level}(k) \geq l$.*

Then a is locally robust up to $L_v'' = (L_v \cup \bigcup_{k \in \mathcal{L}''} \text{Level}(k)) \setminus \{l\}$ in E'' where E'' is such that $E' \xrightarrow{\text{Lost}(A)} E''$ with $A = \{k \mid \text{Level}(k) = l' \text{ with } l' < l\}$ and $\mathcal{L}'' = \mathcal{L}' \cup A$.

Similarly to Theorem 3, Theorem 4 ensures that once a level is blacklisted, the corresponding level is secure again. More precisely, it states that the system is then robust up to $L_v \setminus \{l\}$ plus the levels $\{l' \mid l' < l\}$ which can still be compromised and the newly lost keys \mathcal{L}' , that is, the keys that might have been lost during the execution between E and E' . The proof follows easily from Theorem 3.

5.4 Application

Suppose we are using our API to implement a protocol similar to Kerberos [13] for granting access to resources (see Figure 5). We consider key distribution centre (*KDC*), a ticket granting server (*TGS*) and principals Alice (a) and Bob (b).

We will order the levels using an inclusion ordering, where K_a has level (a) , $K_{a, TGS}$ has level (a, TGS, KDC) and $K_{a, b}$ has level (a, b, TGS, KDC) (i.e. the levels correspond to the parties assumed to have access to the key). Other terms are considered public (level 0). For two levels l, l' , $l > l' \Leftrightarrow l \subset l'$. The protocol is now implementable using the commands of the API.

In a deployment of the system, we assume that devices for the users a, b, \dots and for the servers TGS, KDC are loaded with appropriate keys of level Max as part of some personalisation process. The administrator can then load long term keys $K_a, K_{a, TGS}$ etc. on to the user devices and the servers. They can now play out the protocol using just the working key commands **encrypt**, **decrypt**, **secretGenerate**, **publicGenerate**. Note that when implemented by our API the protocol messages will be tagged following our scheme, e.g. message 2 will appear as $\{tgt, 0, v', TGT\}_{K_{TGS}}$, $\{K_{a, TGS}, (a, TGS, KDC), v, K_{A, TGS}\}_{K_a}$ where TGT and $K_{A, TGS}$ are elements in the miscellaneous field and v, v' are some validity times.

1. $a \rightarrow KDC : request, KDC$
2. $KDC \rightarrow a : \{tgt\}_{K_{TGS}}, \{K_{a,TGS}\}_{K_a}$
3. $a \rightarrow TGS : a, b, \{tgt\}_{K_{TGS}}, \{Auth1\}_{K_{a,TGS}}$
4. $TGS \rightarrow a : \{K_{a,b}\}_{K_b}, \{K_{a,b}\}_{K_{a,TGS}}$
5. $a \rightarrow b : b, \{K_{a,b}\}_{K_b}, \{Auth2\}_{K_{a,b}}$

Figure 5: A Kerberos-like protocol

Suppose key $K_{a,TGS}$ is compromised. Then the administrator has two choices. One is to do nothing. Then he knows that after time $v + \Delta((a, TGS, KDC))$, where v is the expiry time of $K_{a,TGS}$, the system returns to a state where keys at level (a, TGS, KDC) are secure. Lower level keys like session keys $K_{a,b}$ might still be compromised, but new session keys will be safe. Alternatively he can signal to TGS and a to blacklist the key. In this case this level (a, TGS, KDC) becomes unusable. However, to preserve functionality, and administrator can add extra tags t_i to all levels, so level (a) becomes (a, t_1) , level (a, TGS, KDC) becomes (a, TGS, KDC, t_1, t_2) , and the session keys $K_{a,b}$ will have level $(a, b, TGS, KDC, t_1, t_2, t_3)$. If $K_{a,TGS}$ is compromised the administrator only needs to issue a new key with a new level $a, b, TGS, KDC, t_1, t'_2$ and the device can continue to issue tickets securely while level (a, TGS, KDC, t_1, t_2) is blacklisted.

6 Conclusions

We have presented our API for key management with revocation and proved its security properties. While the underlying assumptions for our system to be secure are not unrealistic, in future work we intend to weaken them further. To relax the assumption on level Max messages remaining confidential seems to require more cryptographic primitives. In particular, in order to ensure that the loss of long term keys does not lead to the loss of update messages we require perfect forward secrecy, which would seem to imply the use of asymmetric cryptography. Extending our API to asymmetric cryptography is ongoing work. We also assume here that each device has a real time clock. This could be weakened by adding some sort of nonce based freshness test for each new key accepted, as required in the restricted version of the API on which our revocation scheme is based [4]. All our security proofs are made in the symbolic model, but it should be possible to extend them to the standard model of cryptography under the usual assumptions, i.e. IND-CCA2 for the authenticated encryption scheme, since we avoid problematic constructions such as key cycles.

References

- [1] M. Abadi, B. Blanchet, and C. Fournet. Just fast keying in the pi calculus. *ACM Transactions on Information and System Security (TISSEC)*, 10(3):1–59, July 2007.
- [2] B. Blanchet and A. Chaudhuri. Automated formal analysis of a protocol for secure file sharing on untrusted storage. In *Proceedings of the 29th IEEE*

- Symposium on Security and Privacy (S&P'08)*, pages 417–431, Oakland, CA, 2008. IEEE.
- [3] C. Cachin and N. Chandran. A secure cryptographic token interface. In *Computer Security Foundations (CSF-22)*, pages 141–153, Long Island, New York, 2009. IEEE Computer Society Press.
 - [4] V. Cortier and G. Steel. A generic security API for symmetric key management on cryptographic devices. In *Proceedings of the 14th European Symposium on Research in Computer Security (ESORICS'09)*, volume 5789 of *Lecture Notes in Computer Science*, pages 605–620, Saint Malo, France, September 2009. Springer.
 - [5] V. Cortier, G. Steel, and C. Wiedling. Revoke and let live : A secure key revocation api for cryptographic devices. In *19th ACM Conference on Computer and Communications Security (CCS'12)*, Raleigh, USA, October 2012. ACM.
 - [6] J. Courant and J.-F. Monin. Defending the bank with a proof assistant. In *Proceedings of the 6th International Workshop on Issues in the Theory of Security (WITS'06)*, pages 87 – 98, Vienna, Austria, March 2006.
 - [7] L. Eschenauer and V. D. Gligor. A key management scheme for distributed sensor networks. In *Proceedings of the 9th ACM Conference on Computer and Communications Security (CCS)*, pages 41–47, 2002.
 - [8] S. Fröschle and G. Steel. Analysing PKCS#11 key management APIs with unbounded fresh data. In *Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security (ARSPA-WITS'09)*, volume 5511 of *Lecture Notes in Computer Science*, pages 92–106, York, UK, 2009. Springer.
 - [9] F. E. Kargl. Sevecom baseline architecture. Deliverable D2.1-App.A for EU Project Sevecom, 2009.
 - [10] F. Levy. SAM and key management functional presentation. Available from <http://www.calypsostandard.net/>, December 2010.
 - [11] G. Lowe. Breaking and fixing the Needham-Schroeder public-key protocol using FDR. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS'96)*, volume 1055 of *LNCS*, pages 147–166. Springer-Verlag, march 1996.
 - [12] S. Mödersheim and P. Modesti. Verifying sevecom using set-based abstraction. In *IWCMC*, pages 1164–1169. IEEE, 2011.
 - [13] C. Neuman, T. Yu, S. Hartman, and K. Raeburn. The kerberos network authentication service - rfc 4120. Available at <http://tools.ietf.org/html/rfc4120>.
 - [14] G. Steel and A. Bundy. Attacking group protocols by refuting incorrect inductive conjectures. *Journal of Automated Reasoning*, 46(1–2):149–176, January 2006. Special Issue on Automated Reasoning for Security Protocol Analysis.

- [15] Trusted Computing Group. *TPM Main Part 3 Commands.*, level 2 revision 116 edition, March 2011. Specification Version 1.2.
- [16] B. Weyl. Secure on-board architecture specification. Deliverable D3.2 for EU Project EVITA, <http://evita-project.org/Deliverables/EVITAD3.2.pdf>, August 2011.
- [17] X. Z. Yong Wan, Byrav Ramamurthy. Keyrev: An efficient key revocation scheme for wireless sensor networks. In *IEEE International Conference on Communications (ICC)*, pages 1260 – 1265, 2007.

This appendix contains the proofs of our main theorems.

A Proving Theorem 1

We first prove an easy and rather straightforward lemma that will be often used in the proofs: if a key appears only in key position in a set of messages, then it is not deducible.

Lemma 1. *We suppose that keys are atomic. Let \mathfrak{M} be a set of terms and k a key such that k only appears in key position in m , $\forall m \in \mathfrak{M}$. Then, $\forall t$ such that $\mathfrak{M} \vdash t$, k only appears in key position in t , if k appears in it. In particular, $\mathfrak{M} \not\vdash k$.*

Proof. Let k be a key and the set \mathfrak{M} be such as defined. Let us prove the lemma by induction on the number of steps needed to deduce t .

Base case : All terms in \mathfrak{M} satisfy the property by hypothesis.

Induction Hypothesis : $\forall t$ such that $\mathfrak{M} \vdash t$ in n or less steps, then t satisfies the property.

Now suppose that $\mathfrak{M} \vdash t$ in $n + 1$ steps :

- If the last step is a decomposition rule (all rules but encryption and pairing). We that t is a subterm of a term t' with $t' = \{t\}_v$ or $\langle t, v \rangle$. In the first one, $t' = \{t\}_v$, if k appears not in a key position in t , it appears not in a key position in t' too, and we have a contradiction with the induction hypothesis. The conclusion is the same in the case of the pairing. Thus t , verifies the property.
- If the last step is the encryption rule, then t is of the form $\{t'\}_y$ with $\mathfrak{M} \vdash t'$ in n steps. If $y = k$, then $\exists \Pi$ a proof, in n or less steps, of $\mathfrak{M} \vdash k$, which is a contradiction, since k is not in key position in the term k . Then, we must have $y \neq k$ and, in that case, the induction hypothesis gives us the fact that k only appears in key position in t' and, due to the form of t , it is clear that t also verifies this property.
- If the last step is the pairing rule, then t is of the form $\langle t', t'' \rangle$ with $\mathfrak{M} \vdash t', t''$ in, at most, n steps. Using the induction hypothesis on t' and t'' it is clear that t also satisfy the property.

Moreover, suppose that $\mathfrak{M} \vdash k$. This contradicts the fact that k only appears in key position. Thus $\mathfrak{M} \not\vdash k$. \square

Now let us consider the following hypothesis :

Hypothesis 2. We suppose that the administrator does not make any mistakes when performing administration commands, especially, keys are distributed consistently, with the same attributes and the administrator should not create equalities between keys. It means that if $\exists m_1, m_2 \in \mathfrak{M}$ such that $m_1 = \{\text{cmd}_1, \dots, (k_1), k'_1, \text{att}_1, \dots\}_{q_1 \dots q_n}$ and $m_2 = \{\text{cmd}_2, \dots, (k_2), k'_2, \text{att}_2, \dots\}_{q'_1 \dots q'_n}$, where att_i are the attributes associated to k'_i in the command $\text{cmd}_i \in \{\text{blacklist}, \text{create}, \text{update}\}$ (and k_1, k_2 may not appear), then we have that $k'_1 = k'_2$ (resp. $k_1 = k_2$) implies $\text{cmd}_1 = \text{cmd}_2$, $\text{att}_1 = \text{att}_2$ and $k_1 = k_2$ (resp. $k'_1 = k'_2$).

It also means that if there are two `updateMax` commands m_1 and m_2 such that $m_1 = \{\text{updateMax}, k, v, m\}_{k_1 \dots k_n}$ and $m_2 = \{\text{updateMax}, k', v', m'\}_{k'_1 \dots k'_n}$, then if $k = k'$ (or $k_1 = k'_1$) then we have $m_1 = m_2$.

We also assume that `updateMax`, `update` and `create` commands are made with fresh keys. Moreover, the administrator performs operations according to the specification.

Theorem 1 is proved by induction, showing that revocation keys satisfy a stronger invariant, which is reflected in the following proposition.

Proposition 3. Let us consider global states $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ and $E' = (\mathcal{P}', \mathcal{I}', \mathfrak{M}', N', K', t')$ such that $E_\emptyset \xrightarrow{*} E \xrightarrow{L} E'$. Let \mathcal{L} be a set of keys. Let us consider the following properties :

- (1) $\forall a \in \mathcal{P}, \forall h \in H_a$ s.t. $\Theta_a(h) = (k_1, \text{Max}, v_1, m_1)$ and $\forall b \in \mathcal{P}, \forall h' \in H_b$ s.t. $\Theta_b(h') = (k_2, \text{Max}, v_2, m_2)$, then, we have:

$$k_1 = k_2 \text{ and } k_1, k_2 \notin \mathcal{L} \implies (v_1, m_1) = (v_2, m_2).$$

Moreover, $\forall a \in \mathcal{P}$, all revocation keys stored in a given API are different, i.e. $\forall h, h' \in H_a$ s.t. $\Theta_a(h) = (k, \text{Max}, v, m)$ and $\Theta_a(h') = (k', \text{Max}, v', m')$, then $k = k' \implies h = h'$.

We consider the set $\mathcal{U}_E^{\text{Max}}$ of latent revocation keys of global state E .

- (2) $\forall k \in \mathcal{U}_E^{\text{Max}} \setminus \mathcal{L}, \forall m \in \mathfrak{M}$ we have, for all occurrences of k in m is of the form $\{m'\}_k$.

Then, if E satisfies (1) and (2) for \mathcal{L} , then E' also satisfies them for \mathcal{L}' with $\mathcal{L}' = \mathcal{L} \cup \{k\}$ if $E \xrightarrow{\text{Lost}(k)} E'$ and $\mathcal{L}' = \mathcal{L}$ otherwise.

Proof. To prove the proposition, we will show that properties (1) and (2) are invariant under application of transition rules :

- $E \xrightarrow{(\text{DED})} E'$. This transition does not affect handles at all, then it is obvious that (1) is invariant. We also have $\mathfrak{M}' = \mathfrak{M} \cup \{u\}$ with $\mathfrak{M} \vdash u$. Let $k \in \mathcal{U}_E^{\text{Max}} \setminus \mathcal{L}$ such that $\text{Level}(k) = \text{Max}$, and $m \in \mathfrak{M}'$ such that $k \in \text{St}(m)$. If $m \in \mathfrak{M}$, then we conclude using the induction hypothesis that E verifies (2). We now consider the case where $m = u$. Using Lemma 1 and induction hypothesis, we know that u is a term where k can only appear in key position and we conclude that (2) holds.

Considering that and using Lemma 1, we easily obtain following lemma:

Lemma 2. *Let E be a global state satisfying to properties (1) and (2) for \mathcal{L} , then:*

$$\forall k \in \mathcal{U}_E^{\text{Max}} \setminus \mathcal{L}, \mathfrak{M} \not\models k.$$

Now let us continue our proof of Proposition 3.

- $E \xrightarrow{(\text{TIM})} E'$. In this case, we have $\mathcal{I}' = \mathcal{I}$ since there is no modification of the memory of all TRD's, thus, (1) holds obviously. Moreover, as we also have $\mathfrak{M}' = \mathfrak{M}$, using the fact that there are no keys which are lost during the transition and since latent revocation keys may only become dead with time, i.e. $\mathcal{U}_{E'}^{\text{Max}} \setminus \mathcal{L}' \subseteq \mathcal{U}_E^{\text{Max}} \setminus \mathcal{L}$, (2) clearly holds for E' .
- $E \xrightarrow{(\text{LST})} E'$. As in the previous case, we have $\mathcal{I}' = \mathcal{I}$ (no modification of handles on TRD's), thus, (1) holds. In addition, we have that $\mathfrak{M}' = \mathfrak{M} \cup \{k\}$, $\mathcal{L}' = \mathcal{L} \cup \{k\}$, with k the lost key, and $\mathcal{U}_{E'}^{\text{Max}} = \mathcal{U}_E^{\text{Max}}$ since $t' = t$. Let $p \in \mathcal{U}_{E'}^{\text{Max}} \setminus \mathcal{L}'$ and let $m \in \mathfrak{M}'$ such that $p \in \text{St}(m)$. Then, since $p \neq k$ and keys are atomic, we must have $p \notin \text{St}(k)$, thus $m \in \mathfrak{M}$. Using the fact that $p \in \mathcal{U}_{E'}^{\text{Max}} \setminus \mathcal{L}' \subseteq \mathcal{U}_E^{\text{Max}} \setminus \mathcal{L}$ and using the fact that E satisfies (2), we can conclude easily that all occurrence of p in m are of the form $\{m'\}_k$ and conclude that the property (2) is invariant.
- $E \xrightarrow{(\text{UPM})} E'$.
 - (1) Let a be the TRD modified by the transition. Consider $a_1, a_2 \in \mathcal{P}$ and $h_1 \in H_{a_1}$, $h_2 \in H_{a_2}$ s.t. $h_1 \neq h_2$, $\Theta_{a_1}(h_1) = (k_1, \text{Max}, v_1, m_1)$ and $\Theta_{a_2}(h_2) = (k_2, \text{Max}, v_2, m_2)$ with $k_1, k_2 \notin \mathcal{L}'$ and $k_1 = k_2$. We note that $\mathcal{L}' = \mathcal{L}$. We will show that $(v_1, m_1) = (v_2, m_2)$:
 - If neither h_1 nor h_2 has been updated by the transition. Then, we conclude using the induction hypothesis.
 - The case where h_1 and h_2 where both updated by the transition is impossible according to the specification of the API.
 - Now assume that h_1 has been updated (and thus $a_1 = a$) and h_2 has not. We must have that k_1 appears in the command of the transition of the form $\{\text{updateMax}, k_1, v_1, m_1\}_{q_1 \dots q_n}$, and we have two subcases :
 - * Either k_1 has never occurred in the system before, then $k_1 = k_2$ is impossible.
 - * Or k_1 was already existing in the system, that is, there is a TRD c and $h \in H_c$ s.t. $\Theta_c(h) = (k_1, \text{Max}, v, m)$. Using induction hypothesis over h and h_2 , we have that $(v, m) = (v_2, m_2)$. Then, considering that, by case inspection, it is only possible to put a revocation key in a TRD using an `updateMax` command, thus, either k_1 was updated in c using such command or was in c since E_\emptyset . Since it is impossible, according to Hypothesis 1 and Lemma 2, for an intruder to forge an `updateMax` command, this must be an action of the administrator. Then, according to Hypothesis 2, we must have $(v_1, m_1) = (v, m)$ otherwise it implies that the administrator does not act properly. Since $(v, m) = (v_2, m_2)$, we

conclude.

Now, let us show that $\forall a \in \mathcal{P}, \forall h, h' \in H_a$ s.t. $\Theta_a(h) = (k, \text{Max}, v, m)$ and $\Theta_a(h') = (k', \text{Max}, v', m')$, then $k = k' \Rightarrow h = h'$. Let us consider $b \in \mathcal{P}$ and $h, h' \in H_b$ s.t. $\Theta_b(h) = (k, \text{Max}, v, m)$ and $\Theta_b(h') = (k', \text{Max}, v', m')$ with $k = k'$. If $b \neq a$, the modified TRD, then, using induction hypothesis, we have that $h = h'$. If $b = a$ with k a revocation key not added in the TRD during the transition, then, again, the induction hypothesis gives us that $h = h'$. If k is added in the TRD during the transition, let us suppose that we have $h \neq h'$. Then, we must have that $\exists m_1, m_2$ two **updateMax** commands coming from the administrator (we showed previously that an intruder cannot forge such commands) s.t. $m_1 = \{\text{updateMax}, k, v, m\}_{q_1 \dots q_n}$ and $m_2 = \{\text{updateMax}, k, v, m\}_{q'_1 \dots q'_n}$ with $q_1 \neq q'_1$ according to the specification since q_1 and q'_1 determine the handle where k will be stored. Such a case is in contradiction with Hypothesis 2. Thus $h' = h$. Finally, we conclude that (1) holds.

- (2) In this transition we have $\mathfrak{M}' = \mathfrak{M}$ and even if the set of latent but not lost revocation key is changed by L , there is no trace of these new keys in \mathfrak{M}' , thus we can conclude.

- $E \xrightarrow{L} E'$ with $L \in \{(\text{BLK}), (\text{NEW}), (\text{UPD}), (\text{RVK})\}$. Since these transitions do not affect handles containing revocation keys, it is clear that (1) still holds. Concerning (2), we have $\mathfrak{M}' = \mathfrak{M} \cup \{m'\}_{k_1 \dots k_n}$, m' may only contain keys with a level strictly lesser than **Max**, and k_1, \dots, k_n are revocation keys. Let $k \in \mathcal{U}_E^{\text{Max}} \setminus \mathcal{L}'$, i.e. $k \in \mathcal{U}_E^{\text{Max}} \setminus \mathcal{L}$, with $\text{Level}(k) = \text{Max}$ and $m \in \mathfrak{M}'$ such that $k \in \text{St}(m)$. If $m \in \mathfrak{M}$, we use the fact that E verifies (2) and conclude, if $m = \{m'\}_{k_1 \dots k_n}$, then, according to the structure of m' it must be the case where $k = k_i$ for $i \in \{1, \dots, n\}$ and it also satisfies (2).

- If $E \longrightarrow E'$. Let us consider three subcases :

- The transition corresponds to the **generationPublic**, **generationSecret** or **decrypt** command. Then, handles containing revocation keys are not modified and (1) is invariant. Moreover, there is no output containing revocation keys, so (2) is also invariant.
- The transition corresponds to the **encrypt** command. Such a transition does not affect handles containing revocation keys and (1) is trivially verified. Moreover, we have $\mathfrak{M}' = \mathfrak{M} \cup \{m\}$ with $m = \{U_1, \dots, U_N\}_k$ and U_i may only contain a key of a level strictly less than **Max**. Moreover k can not be a revocation key according to the specification and then, (2) holds.
- The transition corresponds to the **blacklist**, **create**, **update**, **revoke** or **updateMax** command. Such a transition would implies that the attacker is capable to perform an admin command, i.e. is able to forge a message with N_{Max} revocation keys for a given TRD which is in contradiction with Hypothesis 1 and Lemma 2.

□

We can now prove the Theorem 1.

Theorem 1. *Let E be a global state such that $E_\emptyset \longrightarrow^* E$, then :*

$$\forall k \in \mathcal{V}_E^{\text{Max}} \setminus \mathcal{L}, E \not\vdash k.$$

Proof. According to the Definition 3, we have that E_\emptyset clearly satisfies properties (1) and (2) of Proposition 3. Using the result of that proposition, we know that E verifies these properties too. Let us consider $k \in \mathcal{V}_E^{\text{Max}} \setminus \mathcal{L}$, then we have that $k \in \mathcal{U}_E^{\text{Max}} \setminus \mathcal{L}$. The Lemma 2 allows us to conclude that $\mathfrak{M} \not\vdash k$, that is $E \not\vdash k$. \square

B Proving Proposition 1

We show that good keys are never deducible by an attacker. More precisely, we show that the property of robustness is preserved under deduction.

Lemma 3. *The properties defined in Definition 5 are invariant under deduction.*

Proof. Let $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ be a global state robust up to L_v and \mathcal{L} . We consider E' another global state such that $E \xrightarrow{(\text{DED})} E'$ and we prove that E' verifies properties (1), (2) and (3). To do so, we show that if E verifies (1), (2) and (3) for L_v and \mathcal{L} , then $E_p = (\mathcal{P}, \mathcal{I}, \mathfrak{M}_p, N, K, t)$ where $\mathfrak{M}_p = \mathfrak{M} \cup \{p\}$ with $\mathfrak{M} \vdash p$ in one step also verifies (1), (2) and (3) for L_v and \mathcal{L} . We can then conclude by induction.

(1) Let (l, v, m) be attributes of k in E_p and let us show that they must be attributes of k in E already.

- If (l, v, m) are attributes of k by case (a) in E_p , then, since the transition does not affect handles, we must have that (l, v, m) were attributes of k by case (a) already in E .
- If (l, v, m) are attributes of k by case (b) in E_p , we know that $\exists q \in K$, $m \in \mathfrak{M}_p$ s.t. $\text{Max} \in \text{Level}(q)$ and $\{\text{cmd}, \dots, k, l_2, v_2, m_2, \dots\}_{q_1 \dots q_j q} \in \text{St}(m)$ with $\text{Max} \in \text{Level}(q_i)$. We consider two subcases:

- If $m \in \mathfrak{M}$, then (l, v, m) were already attributes of k in E .
- If $m = p$. Then we have two sub-subcases:

* The deduction rule is a decomposition rule or the rule $\frac{m \in \mathfrak{M}}{\mathfrak{M} \vdash m}$. Then, we have that $m \in \text{St}(\mathfrak{M})$ and we can conclude as in the previous case.

* The deduction rule is a composition rule. If the rule $\frac{x \in \mathfrak{M} \ y \in \mathfrak{M}}{\mathfrak{M} \vdash \langle x, y \rangle}$ is applied, then $\{\text{cmd}, \dots, k, l_2, v_2, m_2, \dots\}_{q_1 \dots q_j q}$ must be subterm of either x or y and we can conclude. If the rule is $\frac{x \in \mathfrak{M} \ y \in \mathfrak{M}}{\mathfrak{M} \vdash \{x\}_y}$, then we have two possibilities :

* $\{\text{cmd}, \dots, k, l, v, m, \dots\}_{q_1 \dots q_j q}$ is subterm of $x \in \mathfrak{M}$ and we conclude.

- Inria

satisfies (2).

- If the rule is $\frac{\{x\}_y \in \mathfrak{M} \ y \in \mathfrak{M}}{\mathfrak{M} \vdash x}$, then k occurring in $m = x$ implies that k occurs in $\{x\}_y$. Since $\{x\}_y \in \mathfrak{M}$, using the induction hypothesis, we know that it verifies (2). Moreover, since $y \in \mathfrak{M}$ we have that $y = k$ leads to a contradiction with (2) in E . We consider three subcases:
 - In the first subcase, k appears in key position in $\{x\}_y$ and as we have just seen it, it can not be the case $y = k$, then $\exists m'$ s.t. $\{m'\}_k$ occurs in x . In that case, this occurrence of k in x still satisfies (2.i).
 - In the second subcase, we have $\{\dots, k, \dots\}_{k'} \in \text{St}(\{x\}_y)$ with $k' \in K$ s.t. $v_k \leq v_{k'} + \delta(l)$ with $l = \text{Level}(k) < \text{Level}(k')$ and $\text{Max} \notin \text{Level}(k')$ (k appears in $\{x\}_y$ according to (2.ii)). If $\{\dots, k, \dots\}_{k'} \in \text{St}(x)$ we can conclude easily, the most complicated case is when $\{\dots, k, \dots\}_{k'} = \{x\}_y$. In that case, since $v_k \leq v_{k'} + \delta(l)$ and $k \in \mathcal{U}_{E_p}$, we have that $k' \in \mathcal{U}_{E_p}$, indeed, if we consider $l' = \max(\text{Level}(k'))$, we have:

$$\begin{aligned} v_{k'} + \Delta(l') &= v_{k'} + \delta(l) - \delta(l) + \Delta(l'), \\ &\geq v_k + \Delta(l') - \delta(l), \\ &\geq v_k + \Delta(l), \text{ since } l < l' \\ &> t. \end{aligned}$$

Now, with $k \in \mathcal{G}_{(E_p, \mathcal{L})}^{\text{Lv}}$ and that $k' \in \mathcal{U}_{E_p}$, we have that $k' \notin \mathcal{L}$ otherwise, using property (3), there would be some level l' in Lv corresponding to k' and since $\text{Level}(k) < \text{Level}(k')$, k would not be in $\mathcal{G}_{(E_p, \mathcal{L})}^{\text{Lv}}$. Using this fact and that $\text{Level}(k) < \text{Level}(k')$ we also have that $\nexists i, j$ s.t. $l_i < l_j$ with $l_i \in \text{Level}(k')$ and $l_j \in \text{Lv}$. Thus, we have that $k' \in \mathcal{G}_{(E_p, \mathcal{L})}^{\text{Lv}}$, especially $k' \in \mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$ (the transition does not affect time or the set of lost keys). Then, $y = k'$ would implies that $k' \in \mathfrak{M}$ which is a contradiction with (2).

- In the third subcase, we have $\{\dots, k, \dots\}_{q_1 \dots q_n k'} \in \text{St}(\{x\}_y)$ with $q_1, \dots, q_n, k' \in K$ s.t., for $i \in [1, n]$, $\text{Max} \in \text{Level}(q_i)$, $k' \notin \mathcal{L}$, $\text{Max} \in \text{Level}(k')$ and $v_k \leq v_{k'} + \delta(l)$ with $l = \text{Level}(k)$ (The occurrence of k is of the form (2.iii)). As always if the whole occurrence is included in x , the conclusion is easy. We consider the worst case where there is the equality. Using the argument above on validities, we know that $k' \in \mathcal{U}_{E_p}$, i.e. in \mathcal{U}_E , and since it is not lost and that $\text{Max} \in \text{Level}(k')$, we have that $k' \in \mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$. Then, using Lemma 2, we know that $E \not\models k'$ and then that $y = k'$ leads to a contradiction.

(3) Since $\mathcal{L}' = \mathcal{L}$ and $\mathcal{U}_E = \mathcal{U}_{E'}$, we have that $\{l \mid \exists k \in \mathcal{L} \cap \mathcal{U}_E \text{ s.t. } l \in$

$\text{Level}(k) = \{l \mid \exists k \in \mathcal{L}' \cap \mathcal{U}_{E'} \text{ s.t. } l \in \text{Level}(k)\}$ and (3) holds obviously with the same L_v .

□

We can now easily prove Proposition 1.

Proposition 1. *Let $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ be a global state, \mathcal{L} a set of keys (typically lost keys), and $k \in K$. Let L_v be a set of levels. Assume that E is robust up to L_v and \mathcal{L} then:*

$$\forall k \in \mathcal{G}_{(E, \mathcal{L})}^{L_v} \quad \mathfrak{M} \not\vdash k.$$

Proof. Let us consider that $\exists k \in \mathcal{G}_{(E, \mathcal{L})}^{L_v}$ such that $E \vdash k$, i.e. $\mathfrak{M} \vdash k$. According to this, $\exists E'$ a global state such that $E \xrightarrow{(\text{DED})} E'$ where $\mathfrak{M}' = \mathfrak{M} \cup \{k\}$. Since E is robust up to L_v and \mathcal{L} , using Lemma 3, we know that E' is robust up to L_v and \mathcal{L} . But $k \in \mathcal{G}_{(E, \mathcal{L})}^{L_v}$ implies $k \in \mathcal{G}_{(E', \mathcal{L})}^{L_v}$ and since $k \in \mathfrak{M}'$, then E' cannot satisfies (2). Contradiction.

□

C Proving Theorem 2

Theorem 2. *Let E be a global state such that $E_0 \longrightarrow^* E$ and let \mathcal{L} be the set of keys lost along the execution $E_0 \longrightarrow^* E$. Then E is robust up to $L_v = \{l \mid \exists \text{Lost}(k) \in \mathcal{L} \text{ and } \text{Level}(k) = l\}$ and \mathcal{L} .*

The proof of this theorem by induction. The initial case is obvious and the general case is proved by the following proposition.

Proposition 4. *Let us consider global states $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ and $E' = (\mathcal{P}', \mathcal{I}', \mathfrak{M}', N', K', t')$ such that $E_0 \longrightarrow^* E \xrightarrow{L} E'$. We suppose that E is robust up to $L_v = \{l \mid \exists \text{Lost}(k) \in \mathcal{L} \text{ and } \text{Level}(k) = l\}$ and \mathcal{L} . If $L \neq \text{Lost}(k)$, then E' is robust up to L_v and \mathcal{L} , otherwise E' is robust up to $L'_v = L_v \cup \text{Level}(k)$.*

Proof. We will show that (1), (2) and (3) are invariant under application of transition rules :

- If $E \xrightarrow{(\text{DED})} E'$. This step is already proven in Lemma 3.
- If $E \xrightarrow{(\text{TIM})} E'$. During this transition, only some latent keys may become dead since $t_{E'} > t_E$. In that case, we have $\mathcal{G}_{(E', \mathcal{L})}^{L_v} \subseteq \mathcal{G}_{(E, \mathcal{L})}^{L_v}$, moreover, by inspection, we have that existential quantifications over keys of properties (1) and (2) are stable by time passing. (Time does not turn non lost keys into lost keys.) Thus, we conclude using the induction hypothesis on E . Finally, since we may only have $\{l \mid \exists k \in \mathcal{L} \cap \mathcal{U}_E \text{ s.t. } l \in \text{Level}(k)\} \supseteq \{l \mid \exists k \in \mathcal{L} \cap \mathcal{U}_{E'} \text{ s.t. } l \in \text{Level}(k)\}$, (3) holds obviously.

- If $E \xrightarrow{(\text{LST})} E'$. Then, we consider $L'_v = L_v \cup \text{Level}(k)$ and $\mathcal{L}' = \mathcal{L} \cup \{k\}$ with k the key lost during the transition. Considering this, we have that $\mathcal{G}_{(E', \mathcal{L}')}^{L'_v} \subseteq \mathcal{G}_{(E, \mathcal{L})}^{L_v}$ and, since, by inspection, existential quantifications over keys in property (1) are stable under the lost transition, using the induction hypothesis on E , (1) holds easily.

For property (2), There are two issues. The first one concerns the case where $k_1 \in \mathcal{G}_{(E, \mathcal{L})}^{L_v}$ s.t. we have $\{\dots, k_1, \dots\}_k \in \text{St}(\mathfrak{M})$ with $k \in K$ and $\text{Level}(k_1) = l_1 < \text{Level}(k)$, $\text{Max} \notin \text{Level}(k)$ and $v_{k_1} \leq v_k + \delta(l_1)$. But since we lost k during the transition, then k_1 is no more in $\mathcal{G}_{(E', \mathcal{L}')}^{L'_v}$ and then there is nothing to verify. The second problem appears when we have $k_1 \in \mathcal{G}_{(E, \mathcal{L})}^{L_v}$ s.t. $\{\dots, k_1, \dots\}_{q_1 \dots q_n k} \in \text{St}(\mathfrak{M})$ with $q_1, \dots, q_n, k \in K$ s.t., for $i \in [1, n]$, $\text{Max} \in \text{Level}(q_i)$, $k \notin \mathcal{L}$, $\text{Max} \in \text{Level}(k')$ and $v_{k_1} \leq v_k + \delta_{l_1}$ with $l_1 = \text{Level}(k_1)$. Even if we lose k during the transition, we may have that k_1 remains in $\mathcal{G}_{(E', \mathcal{L}')}^{L'_v}$. Let us suppose that, then, using the last inequality, we have :

$$\begin{aligned} v_{k_1} + \Delta(l_1) &\leq v_k + \delta(l_1) + \Delta(l_1), \\ &\leq v_k + \Delta(\text{Max}), \text{ since } l_1 < \text{Max} \in \text{Level}(k'). \end{aligned}$$

This shows that $k \in \mathcal{U}_{E'}$ (i.e. in \mathcal{U}_E) since, if it is not, then $v_k + \Delta(\text{Max}) \leq t$ which implies $v_{k_1} + \Delta(\text{Max}) \leq t$ and, especially, that k_1 is dead implying a contradiction with the fact that $k_1 \in \mathcal{G}_{(E', \mathcal{L}')}^{L'_v}$. Moreover, since $k \notin \mathcal{L}$ in E , we know that k was in $\mathcal{G}_{(E, \mathcal{L})}^{L_v}$ and, using Lemma 2, we have that $E \not\models k$. Let $m \in \mathfrak{M}$ containing the occurrence of k_1 we focus on. Then, we have two possibilities : m is an admin message or m was deduced by the intruder. In the second case, we have that $\exists m' \in \mathfrak{M}$ an admin message such that $\{m', t_1, \dots, t_n\} \vdash m$ with $t_1, \dots, t_n \in \mathfrak{M}$ and $\{\dots, k_1, \dots\}_{q_1 \dots q_n k} \in \text{St}(m')$. Otherwise, it would implies that the intruder forged m containing an undeducible revocation key without any admin message, which is impossible. Then, in both cases, we have that $\exists m \in \mathfrak{M}$ an admin message s.t. $\{\dots, k_1, \dots\}_{q_1 \dots q_n k} \in \text{St}(m)$. We distinguish three subcases :

- If we have $n > N_{\text{Max}} - 1$, we have a term with more than N_{Max} revocation keys subterm of a admin message which is in contradiction with Hypothesis 2 since the specification requires N_{Max} revocation keys, no more.
- If we have $n = N_{\text{Max}} - 1$, then $m = \{\dots, k, \dots\}_{q_1 \dots q_n k}$. Suppose that $\forall i \in [1, n]$, $q_i \in \mathcal{L}$. Then, if no keys are dead, we have N_{Max} latent revocation keys lost, which is a contradiction to Hypothesis 1, since all these revocation keys belong to one TRD. Then, one key is dead. If k is dead, we already saw that k_1 is dead too and we have a

contradiction. If $\exists j \in [1, n]$ such that q_j is dead, we have :

$$\begin{aligned} v_{k_1} + \Delta(l_1) &= v_{k_1} - v_{q_j} + v_{q_j} + \Delta(l_1), \\ &\leq \delta(l_1) + v_{q_j} + \Delta(l_1), \\ &\leq v_{q_j} + \Delta(\text{Max}), \\ &\leq t. \end{aligned}$$

Indeed, we have $v_{k_1} - v_{q_j} \leq \delta(l_1)$ since at the time of emission of the message, time t_e , we know that $v_{k_1} \leq t_e + \delta(l_1)$ and that $t_e < v_{q_j}$ since q_j was necessarily valid. Thus, in that case, we also have a contradiction on k_1 , then, we must have that $\exists i \in [1, n]$ such that $q_i \notin \mathcal{L}$. Moreover, we have that $v_{k_1} \leq v_{q_i} + \delta(l_1)$ using the two previous inequalities. Then, we show that the occurrence of k_1 still verifying (2.iii) after the loss of k .

- If $n < N_{\text{Max}} - 1$. We have $m = \{\dots, k, \dots\}_{q_1 \dots q_n k' q'_1 \dots q'_p}$ with $n + p = N_{\text{Max}} - 1$ and $\forall i \in [1, p]$, $\text{Level}(q'_i) = \text{Max}$ and $E \vdash q'_i$ for $i \in [1, p]$. ($\{\dots, k, \dots\}_{q_1 \dots q_n k'}$ is deduced from m , then it implies that the keys 'protecting' this term were deducible otherwise there will be a contradiction.) This implies, for each $i \in [1, p]$ that either $q'_i \in \mathcal{L}$ or q'_i is dead. If q'_i is dead, we prove, using the argument with the inequalities above, that k_1 is dead, which is a contradiction with the fact that $k_1 \in \mathcal{G}_{(E', \mathcal{L}')}^{\text{L}'_v}$, thus, we have that $\forall i \in [1, p]$, $q'_i \in \mathcal{L}$. Now, we can apply the previous case ($n = N_{\text{Max}} - 1$) to m in order to prove that $\exists i \in [1, n]$ s.t. $q_i \notin \mathcal{L}$ and $v_{k_1} \leq v_{q_i} + \delta(l_1)$. Thus, we show that the occurrence of k_1 still verifying (2.iii).

For property (3), $\{l \mid \exists k \in \mathcal{L}' \cap \mathcal{U}_{E'} \text{ s.t. } l \in \text{Level}(k)\} = \{l \mid \exists k \in \mathcal{L} \cap \mathcal{U}_E \text{ s.t. } l \in \text{Level}(k)\} \cup \text{Level}(k)$ and since $\text{L}'_v = \text{L}_v \cup \text{Level}(k)$ and $\text{L}_v \geq_s \{l \mid \exists k \in \mathcal{L} \cap \mathcal{U}_E \text{ s.t. } l \in \text{Level}(k)\}$, (3) holds.

- If $E \xrightarrow{(\text{UPM})} E'$. This transition only affects handles containing revocation keys, implying that we have $\mathcal{G}_{(E', \mathcal{L}')}^{\text{L}'_v} = \mathcal{G}_{(E, \mathcal{L})}^{\text{L}_v}$ and, according to the semantic, we also have $\mathfrak{M}' = \mathfrak{M}$. Moreover, by inspection and using hypothesis 2 on admin commands, we have that existential quantifications over keys of properties (1) and (2) are stable by these transitions. Then, it is clear, that, using the induction hypothesis on E , we have that (1) and (2) holds for E' . Finally, since $\mathcal{L} = \mathcal{L}'$, (3) holds obviously.
- If $E \xrightarrow{(\text{RVK})} E'$. In this case, we remove a certain number of handles (keys) from a TRD and we have that $\mathfrak{M}' = \mathfrak{M} \cup \{m\}$ where m does not contain key values. Then, using the induction hypothesis, we have that (1) and (2) holds. Finally, since $\mathcal{L} = \mathcal{L}'$, (3) holds obviously.
- If $E \xrightarrow{(\text{BLK})} E'$. In this case, we remove a certain number of handles (keys) from a TRD and we have that $\mathfrak{M}' = \mathfrak{M} \cup \{m\}$ where m contains no key values, i.e. $\forall k \in \mathcal{G}_{(E', \mathcal{L}')}^{\text{L}'_v}$ s.t. $k \notin \text{St}(m)$. Considering this, and the induction hypothesis, it is obvious that (1) and (2) holds for E' . Finally, since $\mathcal{L} = \mathcal{L}'$, (3) holds obviously.

- If $E \xrightarrow{L} E'$ with $L \in \{(\text{NEW}), (\text{UPD})\}$. In this case, we have $\mathcal{G}_{(E', \mathcal{L}')}^{\text{Lv}} \supseteq \mathcal{G}_E^{\text{Lv}}$ and $\mathfrak{M}' = \mathfrak{M} \cup \{m\}$ where $m = \{m'\}_{q_1 \dots q_r}$ with q_1, \dots, q_r valid revocation keys for the concerned TRD, and then, using Hypothesis 1, $\exists j \in [1, r]$ s.t. $q_r \notin \mathcal{L}'$.
 - (1) Let $k \in \mathcal{G}_{(E', \mathcal{L}')}^{\text{Lv}}$. If k appears in m with no attributes (in the case where k will be updated by the command), then, using the induction hypothesis is enough to conclude. If k appears with attributes (then it is the new key created or updated by the command), then let us consider two cases :
 - Either k has never occurred before in \mathfrak{M} (in the commands that have been previously executed), then (1) holds obviously.
 - Or k has occurred before in \mathfrak{M} . Then, during the evolution $E_\emptyset \longrightarrow^* E$, we consider the increasing sequence of history $(\mathfrak{M}_i)_i$ and consider the first one such that k does not appear in \mathfrak{M}_{i-1} and k appears in \mathfrak{M}_i . To set a fixed lower level key in a TRD, there are three ways : a **create**, **update** or **decrypt** command. Since the two first commands require N_{Max} valid revocation keys, if k appears with such a command, this one must have been done by the administrator according to Hypothesis 1 and Lemma 2. The use of a **decrypt** command implies that the command comes from the intruder (an administrator will not use this way to put k in a TRD) but this would be in contradiction with the fact that $k \in \mathcal{G}_{(E', \mathcal{L}')}^{\text{Lv}}$ since k would be lost from the start. Thus k was set up by an admin command m' , and, using the Hypothesis 2, we know that the attributes associated to k in m are the same to those which are associated to k in m' . Now, using the induction hypothesis in E , we know that all previous attributes associated to k are the same as those in m' , which are the same as those in the new handle and m . Thus, (1) holds.
 - (2) Let $k \in \mathcal{G}_{(E', \mathcal{L}')}^{\text{Lv}}$ s.t. $\exists m' \in \mathfrak{M}'$ s.t. $k \in \text{St}(m')$. We have the same discussion as is the previous transition : there is at least one revocation key q which is not lost in m and we must have $v_k \leq v_q + \delta(l)$ where $l = \text{Level}(k)$. Then, (2) holds.
 - (3) Finally, since $\mathcal{L} = \mathcal{L}'$, (3) holds obviously.
- If $E \longrightarrow E'$. Let us consider four subcases :
 - The transition corresponds to the **generatePublic** or **generateSecret** command. For the first one, only information of level 0 is created and then properties hold obviously. The **generateSecret** command creates fresh keys, thus (1) holds and since $\mathfrak{M}' = \mathfrak{M}$ we have also that (2) holds. Finally, since $\mathcal{L} = \mathcal{L}'$, (3) holds obviously.
 - The transition corresponds to the **blacklist**, **create**, **update**, **revoke** or **updateMax** command. This implies that the intruder may forge an

admin message, and, therefore, can know N_{Max} different and valid revocation keys for the API, which is not possible according to Hypothesis 1 and the Theorem 1. Finally, since $\mathcal{L} = \mathcal{L}'$, (3) holds obviously.

- The transition corresponds to the **encrypt** command. For property (1), the command only uses handles, or information of level 0, already existing, then, using induction hypothesis on E , we can conclude that the property holds. For property (2), we have that $\mathfrak{M}' = \mathfrak{M} \cup \{\{U_1, \dots, U_n\}_k\}$ with U_i a level 0 message or a (k_i, l_i, v_i, m_i) set with $1 \leq l_i < \text{Max}$. Suppose that $\exists i \in [1, n]$ s.t. $k_i \in \mathcal{G}_{(E', \mathcal{L}')}^{\text{Lv}}$, then, according to the specification, we have that $\text{Level}(k_i) < \text{Level}(k)$, $\text{Max} \notin \text{Level}(k)$ and $v_k > t$, $v_{k_i} > t$. Thus, we must have $v_{k_i} \leq v_k + \delta(l_i)$ otherwise, we have $v_{k_i} > v_k + \delta(l_i) > t + \delta(l_i)$ which is impossible. Indeed, let us consider the first apparition, as in a previous case, of k_i in the system. We have that k_i appeared from a **create**, **update**, **decrypt** or **generateSecret** command. But, whatever the applied command, the TRD does a test on validity (or sets it itself) such that $v_{k_i} \leq t_c + \delta(l_i)$ where t_c is the time of the state where k_i appears for the first time. Since $t + \delta(l_i) > t_c + \delta(l_i)$, having $v_{k_i} > v_k + \delta(l_i)$ leads to a contradiction. Then (2) holds. Finally, since we still have $\mathcal{L} = \mathcal{L}'$, (3) holds obviously.
- The transition corresponds to the **decrypt** command. We have that \mathfrak{M}' does not contain information of a level greater or equal to 1. Moreover, we add a key in the API concerned by the **decrypt** command and $\mathcal{G}_{(E', \mathcal{L}')}^{\text{Lv}}$ may be larger than $\mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$. Let us consider $k, q \in K'$ s.t. $(\text{decrypt}, \{\dots, k, l, v, m, \dots\}_q, h_q)$, where h_q is the handle pointing to q in the TRD, is the message implying the considered **decrypt** command. Suppose that $k \in \mathcal{G}_{(E', \mathcal{L}')}^{\text{Lv}} \setminus \mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$. Then, the message is coming from the intruder, otherwise, k should already appear in a handle of another TRD and we should have $k \in \mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$. Since the message comes from the intruder, it implies that $E \vdash q$. Using Proposition 1, we must have that $q \notin \mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$. But we know that the message was accepted by the TRD, then, q must be valid (and then latent), then, if $q \notin \mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$, we must have that $q \in \mathcal{L}$ or $\exists q' \in K \cap \mathcal{L}$ and $\text{Level}(q) \leq \text{Level}(q')$. Adding the fact that we must have $\text{Level}(k) < \text{Level}(q)$, we must have $k \notin \mathcal{G}_{(E', \mathcal{L}')}^{\text{Lv}}$ which is a contradiction. Thus, k was already in $\mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$ and $\mathcal{G}_{(E', \mathcal{L}')}^{\text{Lv}} = \mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$ which, using the induction hypothesis, implies that E' verifies (2).

For property (1), we consider two cases :

- * If k , the key used to decrypt, is in $\mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$. Then, we know, using Proposition 1, that $E \not\vdash k$. Thus, the message was already existing in \mathfrak{M} and verifies property (1) by induction hypothesis.
- * If k , the key used to decrypt, is not in $\mathcal{G}_{(E, \mathcal{L})}^{\text{Lv}}$. Then, it must be the case that $k \in \mathcal{L}$ or k is under a lost key (and may be

considered as lost.) since it must be valid (then, a fortiori, latent). Since all keys q in the encrypted message must verify $\text{Level}(q) < \text{Level}(k)$, then, $q \notin \mathcal{G}_{(E, \mathcal{L})}^{L_v}$, i.e. $q \notin \mathcal{G}_{(E', \mathcal{L}')}^{L_v}$ and then does not interfere with property (1).

Finally, since $\mathcal{L} = \mathcal{L}'$, (3) holds obviously. \square

D Retrieving higher security levels

We show that higher security levels are guaranteed either when some time elapses (Theorem 3) or when a compromised key is blacklisted (Theorem 4).

Theorem 3. *Let E, E' be two global states such that $E_\emptyset \rightarrow^* E \rightarrow^* E'$. Assume that E is robust up to L_v and \mathcal{L} and let $l \in L_v$ be a compromised level. Let \mathcal{L}' be the set of keys lost along the execution $E \rightarrow^* E'$. We further assume that there has been no ‘lost’ event for keys of level greater than l , that is, there is no $k \in \mathcal{L}'$ such that $\text{Level}(k) \geq l$. Let $v = \sup(\{v_k \mid k \in \mathcal{L} \text{ and } \text{Level}(k) = l\})$, where v_k is the validity time associated to k (Theorem 2 ensures v_k is unique). Then:*

$$t' \geq v + \Delta(l) \implies E'' \text{ is robust up to } L_v'' \text{ for } \mathcal{L}'',$$

where t' is the global time of E' and E'' is such that $E' \xrightarrow{\text{Lost}(A)} E''$ with $A = \{k \mid \text{Level}(k) = l' \text{ with } l' < l\}$ and $L_v'' = (L_v \cup \bigcup_{k \in \mathcal{L}''} \text{Level}(k)) \setminus \{l\}$ where $\mathcal{L}'' = \mathcal{L}' \cup A$.

Proof. We know, using Theorem 2, that E' is robust up to $L_v' = L_v \cup \bigcup_{k \in \mathcal{L}'} \text{Level}(k)$ and \mathcal{L}' . Using Theorem 2 again, we have that E'' is robust up to $L = L_v' \cup \{l' \mid l' < l\} = L_v \cup \bigcup_{k \in \mathcal{L}''} \text{Level}(k)$ and $\mathcal{L}'' = \mathcal{L}' \cup A$.

Now, if we consider $L_v'' = (L_v \cup \bigcup_{k \in \mathcal{L}''} \text{Level}(k)) \setminus \{l\} = L \setminus \{l\}$, we have that $\mathcal{G}_{(E'', \mathcal{L}'')}^L = \mathcal{G}_{(E'', \mathcal{L}'')}^{L_v''}$. Indeed, we have:

$$\begin{aligned} \mathcal{G}_{(E'', \mathcal{L}'')}^L &= \{k \mid \text{Max} \notin \text{Level}(k), k \in \mathcal{U}_{E''} \setminus \mathcal{L}'', \\ &\quad \exists l' \in \text{Level}(k) \text{ s.t. } l' \not\prec L\} \\ &= \{k \mid \text{Max} \notin \text{Level}(k), k \in \mathcal{U}_{E''} \setminus \mathcal{L}'', \\ &\quad \exists l' \in \text{Level}(k) \text{ s.t. } l' \not\prec L_v''\} \\ &\quad \setminus \{k \mid \text{Max} \notin \text{Level}(k), k \in \mathcal{U}_{E''} \setminus \mathcal{L}'', \\ &\quad \exists l' \in \text{Level}(k) \text{ s.t. } l' < l\} \\ &= \mathcal{G}_{(E'', \mathcal{L}'')}^{L_v''} \setminus G, \end{aligned}$$

where $G = \{k \mid \text{Max} \notin \text{Level}(k), k \in \mathcal{U}_{E''} \setminus \mathcal{L}'', \exists l' \in \text{Level}(k) \text{ s.t. } l' < l\}$.

But, in E'' , $\forall k$ s.t. $l' \in \text{Level}(k)$ with $l' < l$, then $k \in \mathcal{L}''$. Then, $G = \emptyset$. Thus, we have that E'' verifies (1) and (2) for L_v'' .

Moreover, we know that $\forall k \in \mathcal{L}''$ s.t. $\text{Level}(k) = l$, then $k \notin \mathcal{U}_{E''}$, according to this, we have $L_v'' \geq_s \{l_i \mid k_i \in \mathcal{L}'' \cap \mathcal{U}_{E''} \text{ where } l_i = \text{Level}(k_i)\}$ using the fact above and the construction of L_v'' . Thus E'' also ensures (3) for L_v'' and E'' is robust up to L_v'' for \mathcal{L}'' . \square

Theorem 4. *Let $E = (\mathcal{P}, \mathcal{I}, \mathfrak{M}, N, K, t)$ a be a global state s.t. $E_\emptyset \rightarrow^* E$. Let \mathcal{L} be a set of (lost) keys. Assume that $a \in \mathcal{P}$ is locally robust up to L_v and \mathcal{L} in E . Assume that $E \xrightarrow{L} \rightarrow^* E'$ with L a blacklist command for TRD a , level l , and validity time t_h (with t_h greater or equal to $v_h = \sup(\{v_k \mid k \in \mathcal{L} \text{ and } l \in \text{Level}(k)\})$). Let \mathcal{L}' be the set of keys lost along the execution $E \xrightarrow{L} \rightarrow^* E'$. We further assume that there has been no lost event for keys of level greater than l , that is, there is no $k \in \mathcal{L}'$ such that $\text{Level}(k) \geq l$.*

Then a is locally robust up to $L_v' = (L_v \cup \bigcup_{k \in \mathcal{L}''} \text{Level}(k)) \setminus \{l\}$ in E'' where E'' is such that $E' \xrightarrow{\text{Lost}(A)} E''$ with $A = \{k \mid \text{Level}(k) = l' \text{ with } l' < l\}$ and $\mathcal{L}'' = \mathcal{L}' \cup A$.

Proof. When the blacklist command is performed, we know, according to the specification, that a is robust up to $(L_v \cup \bigcup_{k \in \mathcal{L}} \text{Level}(k)) \setminus \{l' \mid l' \leq l\}$ and \mathcal{L} (in the state E_L such that $E \xrightarrow{L} E_L \rightarrow^* E'$). Then, if t' , the current time of global state E' is smaller than t_h , then, since the TRD a is blocked under level l , it is clear that it is robust up to $(L_v \cup \bigcup_{k \in \mathcal{L}'} \text{Level}(k)) \setminus \{l' \mid l' \leq l\}$ and \mathcal{L}' in E' , and, robust up to $(L_v \cup \bigcup_{k \in \mathcal{L}''} \text{Level}(k)) \setminus \{l\}$ and \mathcal{L}'' in E'' . (In fact it is even robust to a greater set of levels before the end of the effect of the blacklist.)

If $t' > t_h$, then, according to the hypothesis on t_h , we can use the Theorem 3 to prove that E'' itself is robust up to $L_v' = (L_v \cup \bigcup_{k \in \mathcal{L}''} \text{Level}(k)) \setminus \{l\}$ and \mathcal{L}'' , implying immediately that a is robust up to L_v' and \mathcal{L}'' in E'' . \square



**RESEARCH CENTRE
NANCY – GRAND EST**

615 rue du Jardin Botanique
CS20101
54603 Villers-lès-Nancy Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399