

Internet Engineering Task Force (IETF)
Request for Comments: 6290
Category: Standards Track
ISSN: 2070-1721

Y. Nir, Ed.
Check Point
D. Wierbowski
IBM
F. Detienne
P. Sethi
Cisco
June 2011

A Quick Crash Detection Method for the Internet Key Exchange Protocol (IKE)

Abstract

This document describes an extension to the Internet Key Exchange Protocol version 2 (IKEv2) that allows for faster detection of Security Association (SA) desynchronization using a saved token.

When an IPsec tunnel between two IKEv2 peers is disconnected due to a restart of one peer, it can take as much as several minutes for the other peer to discover that the reboot has occurred, thus delaying recovery. In this text, we propose an extension to the protocol that allows for recovery immediately following the restart.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6290>.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Conventions Used in This Document	3
2. RFC 5996 Crash Recovery	4
3. Protocol Outline	5
4. Formats and Exchanges	6
4.1. Notification Format	6
4.2. Passing a Token in the AUTH Exchange	7
4.3. Replacing Tokens after Rekey or Resumption	8
4.4. Replacing the Token for an Existing SA	9
4.5. Presenting the Token in an Unprotected Message	9
5. Token Generation and Verification	10
5.1. A Stateless Method of Token Generation	11
5.2. A Stateless Method with IP Addresses	11
5.3. Token Lifetime	12
6. Backup Gateways	12
7. Interaction with Session Resumption	13
8. Operational Considerations	14
8.1. Who Should Implement This Specification	14
8.2. Response to Unknown Child SPI	15
9. Security Considerations	16
9.1. QCD Token Generation and Handling	16
9.2. QCD Token Transmission	17
9.3. QCD Token Enumeration	18
10. IANA Considerations	18
11. Acknowledgements	18
12. References	19
12.1. Normative References	19
12.2. Informative References	19
Appendix A. The Path Not Taken	20
A.1. Initiating a New IKE SA	20
A.2. SIR	20
A.3. Birth Certificates	20
A.4. Reducing Liveness Check Length	21

1. Introduction

IKEv2, as described in [RFC5996] and its predecessor RFC 4306, has a method for recovering from a reboot of one peer. As long as traffic flows in both directions, the rebooted peer should re-establish the tunnels immediately. However, in many cases, the rebooted peer is a VPN gateway that protects only servers, so all traffic is inbound. In other cases, the non-rebooted peer has a dynamic IP address, so the rebooted peer cannot initiate IKE because its current IP address is unknown. In such cases, the rebooted peer will not be able to re-establish the tunnels. Section 2 describes how recovery works under RFC 5996, and explains why it may take several minutes.

The method proposed here is to send an octet string, called a "QCD token", in the IKE_AUTH exchange that establishes the tunnel. That token can be stored on the peer as part of the IKE SA. After a reboot, the rebooted implementation can re-generate the token and send it to the peer, so as to delete the IKE SA. Deleting the IKE SA results in a quick establishment of new IPsec tunnels. This is described in Section 3.

1.1. Conventions Used in This Document

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

The term "token" refers to an octet string that an implementation can generate using only the properties of a protected IKE message (such as IKE Security Parameter Indexes (SPIs)) as input. A conforming implementation MUST be able to generate the same token from the same input even after rebooting.

The term "token maker" refers to an implementation that generates a token and sends it to the peer as specified in this document.

The term "token taker" refers to an implementation that stores such a token or a digest thereof, in order to verify that a new token it receives is identical to the old token it has stored.

The term "non-volatile storage" in this document refers to a data storage module that persists across restarts of the token maker. Examples of such a storage module include an internal disk, an internal flash memory module, an external disk, and an external database. A small non-volatile storage module is required for a token maker, but a larger one can be used to enhance performance, as described in Section 8.2.

2. RFC 5996 Crash Recovery

When one peer loses state or reboots, the other peer does not get any notification, so unidirectional IPsec traffic can still flow. The rebooted peer will not be able to decrypt it, however, and the only remedy is to send an unprotected INVALID_SPI notification as described in Section 3.10.1 of [RFC5996]. That section also describes the processing of such a notification:

If this Informational Message is sent outside the context of an IKE_SA, it should be used by the recipient only as a "hint" that something might be wrong (because it could easily be forged).

Since the INVALID_SPI can only be used as a hint, the non-rebooted peer has to determine whether the IPsec SA and indeed the parent IKE SA are still valid. The method of doing this is described in Section 2.4 of [RFC5996]. This method, called "liveness check", involves sending a protected empty INFORMATIONAL message, and awaiting a response. This procedure is sometimes referred to as "Dead Peer Detection" or DPD.

Section 2.4 does not mandate how many times the liveness check message should be retransmitted, or for how long, but does recommend the following:

It is suggested that messages be retransmitted at least a dozen times over a period of at least several minutes before giving up on an SA...

Those "at least several minutes" are a time during part of which both peers are active, but IPsec cannot be used.

Especially in the case of a reboot (rather than fail-over or administrative clearing of state), the peer does not recover immediately. Reboot, depending on the system, may take from a few seconds to a few minutes. This means that at first the peer just goes silent, i.e., does not send or respond to any messages. IKEv2 implementations can detect this situation and follow the rules given in Section 2.4:

If there has only been outgoing traffic on all of the SAs associated with an IKE SA, it is essential to confirm liveness of the other endpoint to avoid black holes. If no cryptographically protected messages have been received on an IKE SA or any of its Child SAs recently, the system needs to perform a liveness check in order to prevent sending messages to a dead peer.

[RFC5996] does not mandate any time limits, but it is possible that the peer will start liveness checks even before the other end is sending INVALID_SPI notification, as it detected that the other end is not sending any packets anymore while it is still rebooting or recovering from the situation.

This means that the several minutes recovery period is overlapping the actual recover time of the other peer; i.e., if the security gateway requires several minutes to boot up from the crash, then the other peers have already finished their liveness checks before the crashing peer even has a chance to send INVALID_SPI notifications.

There are cases where the peer loses state and is able to recover immediately; in those cases it might take several minutes to recreate the IPsec SAs.

Note that the IKEv2 specification specifically gives no guidance for the number of retries or the length of timeouts, as these do not affect interoperability. This means that implementations are allowed to use the hints provided by the INVALID_SPI messages to shorten those timeouts (i.e., a different environment and situation requiring different rules).

Some existing IKEv2 implementations already do that (i.e., shorten timeouts or limit number of retries) based on these kinds of hints and also start liveness checks quickly after the other end goes silent. However, see Appendix A.4 for a discussion of why this may not be enough.

3. Protocol Outline

Supporting implementations will send a notification, called a "QCD token", as described in Section 4.1 in the first IKE_AUTH exchange messages. These are the first IKE_AUTH request and final IKE_AUTH response that contain the AUTH payloads. The generation of these tokens is a local matter for implementations, but considerations are described in Section 5. Implementations that send such a token will be called "token makers".

A supporting implementation receiving such a token MUST store it (or a digest thereof) along with the IKE SA. Implementations that support this part of the protocol will be called "token takers". Section 8.1 has considerations for which implementations need to be token takers, and which should be token makers. Implementations that are not token takers will silently ignore QCD tokens.

When a token maker receives a protected IKE request message with unknown IKE SPIs, it SHOULD generate a new token that is identical to the previous token, and send it to the requesting peer in an unprotected IKE message as described in Section 4.5.

When a token taker receives the QCD token in an unprotected notification, it MUST verify that the TOKEN_SECRET_DATA matches the token stored with the matching IKE SA. If the verification fails, or if the IKE SPIs in the message do not match any existing IKE SA, it SHOULD log the event. If it succeeds, it MUST silently delete the IKE SA associated with the IKE_SPI fields and all dependent child SAs. This event MAY also be logged. The token taker MUST accept such tokens from any IP address and port combination, so as to allow different kinds of high-availability configurations of the token maker.

A supporting token taker MAY immediately create new SAs using an Initial exchange, or it may wait for subsequent traffic to trigger the creation of new SAs.

See Section 7 for a short discussion about this extension's interaction with IKEv2 Session Resumption ([RFC5723]).

4. Formats and Exchanges

4.1. Notification Format

The notification payload called "QCD token" is formatted as follows:

```

          1               2               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
! Next Payload !C!  RESERVED   !               Payload Length   !
+-----+-----+-----+-----+-----+-----+-----+-----+
! Protocol ID  !  SPI Size    ! QCD Token Notify Message Type !
+-----+-----+-----+-----+-----+-----+-----+-----+
!
~                               TOKEN_SECRET_DATA                               ~
!
+-----+-----+-----+-----+-----+-----+-----+-----+

```

- o Protocol ID (1 octet) MUST be 1, as this message is related to an IKE SA.
- o SPI Size (1 octet) MUST be zero, in conformance with Section 3.10 of [RFC5996].

- o QCD Token Notify Message Type (2 octets) - MUST be 16419, the value assigned for QCD token notifications.
- o TOKEN_SECRET_DATA (variable) contains a generated token as described in Section 5.

4.2. Passing a Token in the AUTH Exchange

For brevity, only the Extensible Authentication Protocol (EAP) version of an AUTH exchange will be presented here. The non-EAP version is very similar. The figures below are based on Appendix C.3 of [RFC5996].

```

first request      --> IDi,
                    [N(INITIAL_CONTACT)],
                    [[N(HTTP_CERT_LOOKUP_SUPPORTED)], CERTREQ+],
                    [IDr],
                    [N(QCD_TOKEN)]
                    [CP(CFG_REQUEST)],
                    [N(IPCOMP_SUPPORTED)+],
                    [N(USE_TRANSPORT_MODE)],
                    [N(ESP_TFC_PADDING_NOT_SUPPORTED)],
                    [N(NON_FIRST_FRAGMENTS_ALSO)],
                    SA, TSi, TSr,
                    [V+]

first response     <-- IDr, [CERT+], AUTH,
                    EAP,
                    [V+]

repeat 1..N times  / --> EAP
                   |
                   \ <-- EAP

last request       --> AUTH

last response      <-- AUTH,
                    [N(QCD_TOKEN)]
                    [CP(CFG_REPLY)],
                    [N(IPCOMP_SUPPORTED)],
                    [N(USE_TRANSPORT_MODE)],
                    [N(ESP_TFC_PADDING_NOT_SUPPORTED)],
                    [N(NON_FIRST_FRAGMENTS_ALSO)],
                    SA, TSi, TSr,
                    [N(ADDITIONAL_TS_POSSIBLE)],
                    [V+]

```

Note that the QCD_TOKEN notification is marked as optional because it is not required by this specification that every implementation be both token maker and token taker. If only one peer sends the QCD token, then a reboot of the other peer will not be recoverable by this method. This may be acceptable if traffic typically originates from the other peer.

In any case, the lack of a QCD_TOKEN notification MUST NOT be taken as an indication that the peer does not support this standard. Conversely, if a peer does not understand this notification, it will simply ignore it. Therefore, a peer may send this notification freely, even if it does not know whether the other side supports it.

The QCD_TOKEN notification is related to the IKE SA and should follow the AUTH payload and precede the Configuration payload and all payloads related to the child SA.

4.3. Replacing Tokens after Rekey or Resumption

After rekeying an IKE SA, the IKE SPIs are replaced, so the new SA also needs to have a token. If only the responder in the rekey exchange is the token maker, this can be done within the CREATE_CHILD_SA exchange. If the initiator is a token maker, then we need an extra informational exchange.

The following figure shows the CREATE_CHILD_SA exchange for rekeying the IKE SA. Only the responder sends a QCD token.

```
request          --> SA, Ni, [KEi]

response         <-- SA, Nr, [KEr], N(QCD_TOKEN)
```

If the initiator is also a token maker, it SHOULD initiate an INFORMATIONAL exchange immediately after the CREATE_CHILD_SA exchange as follows:

```
request          --> N(QCD_TOKEN)

response         <--
```

For session resumption, as specified in [RFC5723], the situation is similar. The responder, which is necessarily the peer that has crashed, SHOULD send a new ticket within the protected payload of the IKE_SESSION_RESUME exchange. If the Initiator is also a token maker, it needs to send a QCD_TOKEN in a separate INFORMATIONAL exchange.

The INFORMATIONAL exchange described in this section can also be used if QCD tokens need to be replaced due to a key rollover. However, since token takers are required to verify at least 4 QCD tokens, this is only necessary if secret QCD keys are rolled over more than four times as often as IKE SAs are rekeyed. See Section 5.1 for an example method that uses secret keys that may require rollover.

4.4. Replacing the Token for an Existing SA

With some token generation methods, such as that described in Section 5.2, a QCD token may sometimes become invalid, although the IKE SA is still perfectly valid.

In such a case, the token maker MUST send the new token in a protected message under that IKE SA. That exchange could be a simple INFORMATIONAL, such as in the last figure in the previous section, or else it can be part of a MOBIKE INFORMATIONAL exchange such as in the following figure taken from Section 2.2 of [RFC4555] and modified by adding a QCD_TOKEN notification:

```
(IP_I2:4500 -> IP_R1:4500)
HDR, SK { N(UPDATE_SA_ADDRESSES),
          N(NAT_DETECTION_SOURCE_IP),
          N(NAT_DETECTION_DESTINATION_IP) } -->

<-- (IP_R1:4500 -> IP_I2:4500)
      HDR, SK { N(NAT_DETECTION_SOURCE_IP),
                N(NAT_DETECTION_DESTINATION_IP) }

<-- (IP_R1:4500 -> IP_I2:4500)
      HDR, SK { N(COOKIE2), [N(QCD_TOKEN)] }

(IP_I2:4500 -> IP_R1:4500)
HDR, SK { N(COOKIE2), [N(QCD_TOKEN)] } -->
```

A token taker MUST accept such gratuitous QCD_TOKEN notifications as long as they are carried in protected exchanges. A token maker SHOULD NOT generate them unless it is no longer able to generate the old QCD_TOKEN.

4.5. Presenting the Token in an Unprotected Message

This QCD_TOKEN notification is unprotected, and is sent as a response to a protected IKE request, which uses an IKE SA that is unknown.

```
message --> N(INVALID_IKE_SPI), N(QCD_TOKEN)+
```

If child SPIs are persistently mapped to IKE SPIs as described in Section 8.2, a token taker may get the following unprotected message in response to an Encapsulating Security Payload (ESP) or Authentication Header (AH) packet.

message --> N(INVALID_SPI), N(QCD_TOKEN)+

The QCD_TOKEN and INVALID_IKE_SPI notifications are sent together to support both implementations that conform to this specification and implementations that don't. Similar to the description in Section 2.21 of [RFC5996], the IKE SPI and message ID fields in the packet headers are taken from the protected IKE request.

To support a periodic rollover of the secret used for token generation, the token taker MUST support at least four QCD_TOKEN notifications in a single packet. The token is considered verified if any of the QCD_TOKEN notifications matches. The token maker MAY generate up to four QCD_TOKEN notifications, based on several generations of keys.

If the QCD_TOKEN verifies OK, the receiver MUST silently discard the IKE SA and all associated child SAs. If the QCD_TOKEN cannot be validated, a response MUST NOT be sent, and the event may be logged. Section 5 defines token verification.

5. Token Generation and Verification

No token generation method is mandated by this document. Two methods are documented in the following sub-sections, but they only serve as examples.

The following lists the requirements for a token generation mechanism:

- o Tokens MUST be at least 16 octets long, and no more than 128 octets long, to facilitate storage and transmission. Tokens SHOULD be indistinguishable from random data.
- o It should not be possible for an external attacker to guess the QCD token generated by an implementation. Cryptographic mechanisms such as a pseudo-random number generator (PRNG) and hash functions are RECOMMENDED.
- o The token maker MUST be able to re-generate or retrieve the token based on the IKE SPIs even after it reboots.

- o The method of token generation MUST be such that a collision of QCD tokens between different pairs of IKE SPI will be highly unlikely.

For verification, the token taker makes a bitwise comparison of the token stored along with the IKE SA with the token sent in the unprotected message. Multihomed takers might flip back-and-forth between several addresses, and have their tokens replaced as described in Section 4.4. To help avoid the case where the latest stored token does not match the address used after the maker lost state, the token taker MAY store several earlier tokens associated with the IKE SA, and silently discard the SA if any of them matches.

5.1. A Stateless Method of Token Generation

The following describes a stateless method of generating a token. In this case, 'stateless' means not maintaining any per-tunnel state, although there is a small amount of non-volatile storage required.

- o At installation or immediately after the first boot of the token maker, 32 random octets are generated using a secure random number generator or a PRNG.
- o Those 32 bytes, called the "QCD_SECRET", are stored in non-volatile storage on the machine, and kept indefinitely.
- o If key rollover is required by policy, the implementation MAY periodically generate a new QCD_SECRET and keep up to 3 previous generations. When sending an unprotected QCD_TOKEN, as many as 4 notification payloads may be sent, each from a different QCD_SECRET.
- o The TOKEN_SECRET_DATA is calculated as follows:

$$\text{TOKEN_SECRET_DATA} = \text{HASH}(\text{QCD_SECRET} \mid \text{SPI-I} \mid \text{SPI-R})$$

5.2. A Stateless Method with IP Addresses

This method is similar to the one in the previous section, except that the IP address of the token taker is also added to the block being hashed. This has the disadvantage that the token needs to be replaced (as described in Section 4.4) whenever the token taker changes its address.

See Section 9.2 for a discussion of a use-case for this method. When using this method, the `TOKEN_SECRET_DATA` field is calculated as follows:

$$\text{TOKEN_SECRET_DATA} = \text{HASH}(\text{QCD_SECRET} \mid \text{SPI-I} \mid \text{SPI-R} \mid \text{IPaddr-T})$$

The `IPaddr-T` field specifies the IP address of the token taker. Secret rollover considerations are similar to those in the previous section.

Note that with a multihomed token taker, the QCD token matches just one of the token taker IP addresses. Usually this is not a problem, as packets sent to the token maker come out the same IP address. If for some reason this changes, then the token maker can replace the token as described in Section 4.4. If IKEv2 Mobility and Multihoming (MOBIKE) is used, replacing the tokens *SHOULD* be piggybacked on the INFORMATIONAL exchange with the `UPDATE_SA_ADDRESSES` notifications.

There is a corner case where the token taker begins using a new IP address (because of multihoming, roaming, or normal network operations) and the token maker loses state before replacing the token. In that case, it will send a correct QCD token, but the token taker will still have the old token. In that case, the extension will not work, and the peers will revert to RFC 5996 recovery.

5.3. Token Lifetime

The token is associated with a single IKE SA and *SHOULD* be deleted by the token taker when the SA is deleted or expires. More formally, the token is associated with the pair (SPI-I, SPI-R).

6. Backup Gateways

Making crash detection and recovery quick is a worthy goal, but since rebooting a gateway takes a non-zero amount of time, many implementations choose to have a standby gateway ready to take over as soon as the primary gateway fails for any reason. [RFC6027] describes considerations for such clusters of gateways with synchronized state, but the rest of this section is relevant even when there is no synchronized state.

If such a configuration is available, it is *RECOMMENDED* that the standby gateway be able to generate the same token as the active gateway. If the method described in Section 5.1 is used, this means that the `QCD_SECRET` field is identical in both gateways. This has the effect of having the crash recovery available immediately.

Note that this refers to "high-availability" configurations, where only one gateway is active at any given moment. This is different from "load sharing" configurations where more than one gateway is active at the same time. For load sharing configurations, please see Section 9.2 for security considerations.

7. Interaction with Session Resumption

Session resumption, specified in [RFC5723], allows the setting up of a new IKE SA to consume less computing resources. This is particularly useful in the case of a remote access gateway that has many tunnels. A failure of such a gateway requires all these many remote access clients to establish an IKE SA either with the rebooted gateway or with a backup. This tunnel re-establishment occurs within a short period of time, creating a burden on the remote access gateway. Session resumption addresses this problem by having the clients store an encrypted derivative of the IKE SA for quick re-establishment.

What Session Resumption does not help is the problem of detecting that the peer gateway has failed. A failed gateway may go undetected for an arbitrarily long time, because IPsec does not have packet acknowledgement, and applications cannot signal the IPsec layer that the tunnel "does not work". Section 2.4 of RFC 5996 does not specify how long an implementation needs to wait before beginning a liveness check, and only says "not recently" (see full quote in Section 2). In practice, some mobile devices wait a very long time before beginning a liveness check, in order to extend battery life by allowing parts of the device to remain in low-power modes.

QCD tokens provide a way to detect the failure of the peer in the case where a liveness check has not yet ended (or begun).

A remote access client conforming to both specifications will store QCD tokens, as well as the Session Resumption ticket, if provided by the gateway. A remote access gateway conforming to both specifications will generate a QCD token for the client. When the gateway reboots, the client will discover this in either of two ways:

1. The client does regular liveness checks, or else the time for some other IKE exchange has come. Since the gateway is still down, the IKE exchange times out after several minutes. In this case, QCD does not help.

2. Either the primary gateway or a backup gateway (see Section 6) is ready and sends a QCD token to the client. In that case, the client will quickly re-establish the IPsec tunnel, either with the rebooted primary gateway or the backup gateway as described in this document.

The full combined protocol looks like this:

```

Initiator                      Responder
-----
HDR, SAi1, KEi, Ni  -->

                                <-- HDR, SAR1, KEr, Nr, [CERTREQ]

HDR, SK {IDi, [CERT,]
[CERTREQ,] [IDr,]
AUTH, N(QCD_TOKEN)
SAi2, TSi, TSr,
N(TICKET_REQUEST)}  -->
                                <-- HDR, SK {IDr, [CERT,] AUTH,
                                N(QCD_TOKEN), SAR2, TSi, TSr,
                                N(TICKET_LT_OPAQUE) }

      ---- Reboot ----

HDR, {}  -->
                                <-- HDR, N(QCD_TOKEN)

HDR, [N(COOKIE),]
Ni, N(TICKET_OPAQUE)
[,N+]  -->
                                <-- HDR, Nr [,N+]

```

8. Operational Considerations

8.1. Who Should Implement This Specification

Throughout this document, we have referred to reboot time alternately as the time that the implementation crashes and the time when it is ready to process IPsec packets and IKE exchanges. Depending on the hardware and software platforms and the cause of the reboot, rebooting may take anywhere from a few seconds to several minutes. If the implementation is down for a long time, the benefit of this protocol extension is reduced. For this reason, critical systems should implement backup gateways as described in Section 6.

Implementing the "token maker" side of QCD makes sense for IKE implementation where protected connections originate from the peer, such as inter-domain VPNs and remote access gateways. Implementing the "token taker" side of QCD makes sense for IKE implementations where protected connections originate, such as inter-domain VPNs and remote access clients.

To clarify this discussion:

- o For remote-access clients it makes sense to implement the token taker role.
- o For remote-access gateways it makes sense to implement the token maker role.
- o For inter-domain VPN gateways it makes sense to implement both roles, because it can't be known in advance where the traffic originates.
- o It is perfectly valid to implement both roles in any case, for example, when using a single library or a single gateway to perform several roles.

In order to limit the effects of Denial-of-Service (DoS) attacks, a token taker SHOULD limit the rate of QCD_TOKENs verified from a particular source.

If excessive amounts of IKE requests protected with unknown IKE SPIs arrive at a token maker, the IKE module SHOULD revert to the behavior described in Section 2.21 of [RFC5996] and either send an INVALID_IKE_SPI notification or ignore it entirely.

Section 9.2 requires that token makers never send a QCD token in the clear for a valid IKE SA and describes some configurations where this could occur. Implementations that may be installed in such configurations SHOULD automatically detect this and disable this extension in unsafe configurations and MUST allow the user to control whether the extension is enabled or disabled.

8.2. Response to Unknown Child SPI

After a reboot, it is more likely that an implementation will receive IPsec packets than IKE packets. In that case, the rebooted implementation will send an INVALID_SPI notification, triggering a liveness check. The token will only be sent in a response to the liveness check, thus requiring an extra round trip.

To avoid this, an implementation that has access to enough non-volatile storage MAY store a mapping of child SPIs to owning IKE SPIs, or to generated tokens. If such a mapping is available and persistent across reboots, the rebooted implementation SHOULD respond to the IPsec packet with an INVALID_SPI notification, along with the appropriate QCD_TOKEN notifications. A token taker SHOULD verify the QCD token that arrives with an INVALID_SPI notification the same as if it arrived with the IKE SPIs of the parent IKE SA.

However, a persistent storage module might not be updated in a timely manner and could be populated with tokens relating to IKE SPIs that have already been rekeyed. A token taker MUST NOT take an invalid QCD token sent along with an INVALID_SPI notification as evidence that the peer is either malfunctioning or attacking, but it SHOULD limit the rate at which such notifications are processed.

9. Security Considerations

The extension described in this document must not reduce the security of IKEv2 or IPsec. Specifically, an eavesdropper must not learn any non-public information about the peers.

The proposed mechanism should be secure against attacks by a passive man in the middle (MITM) (eavesdropper). Such an attacker must not be able to disrupt an existing IKE session, either by resetting the session or by introducing significant delays. This requirement is especially significant, because this document introduces a new way to reset an IKE SA.

The mechanism need not be similarly secure against an active MITM, since this type of attacker is already able to disrupt IKE sessions.

9.1. QCD Token Generation and Handling

Tokens MUST be hard to guess. This is critical, because if an attacker can guess the token associated with an IKE SA, they can tear down the IKE SA and associated tunnels at will. When the token is delivered in the IKE_AUTH exchange, it is encrypted. When it is sent again in an unprotected notification, it is not, but that is the last time this token is ever used.

An aggregation of some tokens generated by one maker together with the related IKE SPIs MUST NOT give an attacker the ability to guess other tokens. Specifically, if one taker does not properly secure the QCD tokens and an attacker gains access to them, this attacker MUST NOT be able to guess other tokens generated by the same maker. This is the reason that the QCD_SECRET in Section 5.1 needs to be sufficiently long.

The token taker **MUST** store the token in a secure manner. No attacker should be able to gain access to a stored token.

The QCD_SECRET **MUST** be protected from access by other parties. Anyone gaining access to this value will be able to delete all the IKE SAs for this token maker.

The QCD token is sent by the rebooted peer in an unprotected message. A message like that is subject to modification, deletion, and replay by an attacker. However, these attacks will not compromise the security of either side. Modification is meaningless because a modified token is simply an invalid token. Deletion will only cause the protocol not to work, resulting in a delay in tunnel re-establishment as described in Section 2. Replay is also meaningless, because the IKE SA has been deleted after the first transmission.

9.2. QCD Token Transmission

A token maker **MUST NOT** send a valid QCD token in an unprotected message for an existing IKE SA.

This requirement is obvious and easy in the case of a single gateway. However, some implementations use a load balancer to divide the load between several physical gateways. It **MUST NOT** be possible even in such a configuration to trick one gateway into sending a valid QCD token for an IKE SA that is valid on another gateway. This is true whether the attempt to trick the gateway uses the token taker's IP address or a different IP address.

IPsec failure detection is not applicable to deployments where the QCD secret is shared by multiple gateways and the gateways cannot assess whether the token can be legitimately sent in the clear while another gateway may actually still own the SA's. Load balancing configurations typically fall in this category. In order for a load balancing configuration of IPsec gateways to support this specification, all members **MUST** be able to tell whether a particular IKE SA is active anywhere in the cluster. One way to do this is to synchronize a list of active IKE SPIs among all the cluster members.

Because it includes the token taker's IP address in the token generation, the method in Section 5.2 can (under certain conditions) prevent revealing the QCD token for an existing pair of IKE SPIs to an attacker who is using a different IP address, even in a load-sharing cluster without state synchronization. That method does not prevent revealing the QCD token to an active attacker who is spoofing the token taker's IP address. Such an attacker may attempt to direct messages to a cluster member other than the member responsible for

the IKE SA in an attempt to trick that gateway into sending a QCD token for a valid IKE SA. That method should not be used unless the load balancer guarantees that IKE packets from the same source IP address always go to the same cluster member.

9.3. QCD Token Enumeration

An attacker may try to attack QCD if the generation algorithm described in Section 5.1 is used. The attacker will send several fake IKE requests to the gateway under attack, receiving and recording the QCD tokens in the responses. This will allow the attacker to create a dictionary of IKE SPIs to QCD tokens, which can later be used to tear down any IKE SA.

Three factors mitigate this threat:

- o The space of all possible IKE SPI pairs is huge: 2^{128} , so making such a dictionary is impractical. Even if we assume that one implementation always generates predictable IKE SPIs, the space is still at least 2^{64} entries, so making the dictionary is extremely hard. To ensure this, token makers MUST generate unpredictable IKE SPIs by using a cryptographically strong pseudo-random number generator.
- o Throttling the amount of QCD_TOKEN notifications sent out, as discussed in Section 8.1, especially when not soon after a crash will limit the attacker's ability to construct a dictionary.
- o The methods in Section 5.1 and Section 5.2 allow for a periodic change of the QCD_SECRET. Any such change invalidates the entire dictionary.

10. IANA Considerations

IANA has assigned a notify message type (16419) from the status types range (16406-40959) of the "IKEv2 Notify Message Types" registry with the name "QUICK_CRASH_DETECTION".

11. Acknowledgements

We would like to thank Hannes Tschofenig and Yaron Sheffer for their comments about Session Resumption.

Others who have contributed valuable comments are, in alphabetical order, Lakshminath Dondeti, Paul Hoffman, Tero Kivinen, Scott C Moonen, Magnus Nystrom, and Keith Welter.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4555] Eronen, P., "IKEv2 Mobility and Multihoming Protocol (MOBIKE)", RFC 4555, June 2006.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.

12.2. Informative References

- [RFC5723] Sheffer, Y. and H. Tschofenig, "Internet Key Exchange Protocol Version 2 (IKEv2) Session Resumption", RFC 5723, January 2010.
- [RFC6027] Nir, Y., "IPsec Cluster Problem Statement", RFC 6027, October 2010.
- [recovery] Detienne, F., Sethi, P., and Y. Nir, "Safe IKE Recovery", Work in Progress, July 2009.

Appendix A. The Path Not Taken

A.1. Initiating a New IKE SA

Instead of sending a QCD token, we could have the rebooted implementation start an Initial exchange with the peer, including the INITIAL_CONTACT notification. This would have the same effect, instructing the peer to erase the old IKE SA, as well as establishing a new IKE SA with fewer rounds.

The disadvantage here is that in IKEv2, an authentication exchange MUST have a piggybacked Child SA set up. Since our use-case is such that the rebooted implementation does not have traffic flowing to the peer, there are no good selectors for such a Child SA.

Additionally, when authentication is asymmetric, such as when EAP is used, it is not possible for the rebooted implementation to initiate IKE.

A.2. SIR

Another proposal that was considered for this work item is the SIR extension, which is described in [recovery]. Under that proposal, the non-rebooted peer sends a non-protected query to the possibly rebooted peer, asking whether the IKE SA exists. The peer replies with either a positive or negative response, and the absence of a positive response, along with the existence of a negative response, is taken as proof that the IKE SA has really been lost.

The working group preferred the QCD proposal to this one.

A.3. Birth Certificates

Birth Certificates is a method of crash detection that has never been formally defined. Bill Sommerfeld suggested this idea in a mail to the IPsec mailing list on August 7, 2000, in a thread discussing methods of crash detection:

If we have the system sign a "birth certificate" when it reboots (including a reboot time or boot sequence number), we could include that with a "bad spi" ICMP error and in the negotiation of the IKE SA.

We believe that this method would have some problems. First, it requires Alice to store the certificate, so as to be able to compare the public keys. That requires more storage than does a QCD token. Additionally, the public key operations needed to verify the self-signed certificates are more expensive for Alice.

We believe that a symmetric-key operation such as proposed here is more light-weight and simple than that implied by the Birth Certificate idea.

A.4. Reducing Liveness Check Length

Some implementations require fewer retransmissions over a shorter period of time for cases of liveness check started because of an `INVALID_SPI` or `INVALID_IKE_SPI` notification.

We believe that the default retransmission policy should represent a good balance between the need for a timely discovery of a dead peer, and a low probability of false detection. We expect the policy to be set to take the shortest time such that this probability achieves a certain target. Therefore, we believe that reducing the elapsed time and retransmission count may create an unacceptably high probability of false detection, and this can be triggered by a single `INVALID_IKE_SPI` notification.

Additionally, even if the retransmission policy is reduced to, say, one minute, it is still a very noticeable delay from a human perspective, from the time that the gateway has come up (i.e., is able to respond with an `INVALID_SPI` or `INVALID_IKE_SPI` notification) and until the tunnels are active, or from the time the backup gateway has taken over until the tunnels are active. The use of QCD tokens can reduce this delay.

Authors' Addresses

Yoav Nir (editor)
Check Point Software Technologies, Ltd.
5 Hasolelim st.
Tel Aviv 67897
Israel

EMail: ynir@checkpoint.com

David Wierbowski
International Business Machines
1701 North Street
Endicott, New York 13760
United States

EMail: wierbows@us.ibm.com

Frederic Detienne
Cisco Systems, Inc.
De Kleetlaan, 7
Diegem B-1831
Belgium

Phone: +32 2 704 5681
EMail: fd@cisco.com

Pratima Sethi
Cisco Systems, Inc.
O'Shaugnessy Road, 11
Bangalore, Karnataka 560027
India

Phone: +91 80 4154 1654
EMail: psethi@cisco.com

