

IPsecME Working Group	Y. Nir, Ed.
Internet-Draft	Check Point
Intended status: Standards Track	D. Wierbowski
Expires: September 11, 2011	IBM
	F. Detienne
	P. Sethi
	Cisco
	March 10, 2011

A Quick Crash Detection Method for IKE  
draft-ietf-ipsecme-failure-detection-06

## Abstract

This document describes an extension to the IKEv2 protocol that allows for faster detection of 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 Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet- Drafts is at <http://datatracker.ietf.org/drafts/current/>. Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress." This Internet-Draft will expire on September 11, 2011.

## 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](#)
- \*1.1. [Conventions Used in This Document](#)
- \*2. [RFC 5996 Crash Recovery](#)
- \*3. [Protocol Outline](#)
- \*4. [Formats and Exchanges](#)
- \*4.1. [Notification Format](#)
- \*4.2. [Passing a Token in the AUTH Exchange](#)
- \*4.3. [Replacing Tokens After Rekey or Resumption](#)
- \*4.4. [Replacing the Token for an Existing SA](#)
- \*4.5. [Presenting the Token in an Unprotected Message](#)
- \*5. [Token Generation and Verification](#)
- \*5.1. [A Stateless Method of Token Generation](#)
- \*5.2. [A Stateless Method with IP addresses](#)
- \*5.3. [Token Lifetime](#)
- \*6. [Backup Gateways](#)
- \*7. [Interaction with Session Resumption](#)
- \*8. [Operational Considerations](#)
- \*8.1. [Who should implement this specification](#)
- \*8.2. [Response to unknown child SPI](#)
- \*9. [Security Considerations](#)
- \*9.1. [QCD Token Generation and Handling](#)
- \*9.2. [QCD Token Transmission](#)
- \*9.3. [QCD Token Enumeration](#)
- \*10. [IANA Considerations](#)
- \*11. [Acknowledgements](#)

- \*12. [Change Log](#)
- \*12.1. [Changes from draft-ietf-ipsecme-failure-detection-05](#)
- \*12.2. [Changes from draft-ietf-ipsecme-failure-detection-04](#)
- \*12.3. [Changes from draft-ietf-ipsecme-failure-detection-03](#)
- \*12.4. [Changes from draft-ietf-ipsecme-failure-detection-02](#)
- \*12.5. [Changes from draft-ietf-ipsecme-failure-detection-01](#)
- \*12.6. [Changes from draft-ietf-ipsecme-failure-detection-00](#)
- \*12.7. [Changes from draft-nir-ike-qcd-07](#)
- \*12.8. [Changes from draft-nir-ike-qcd-03 and -04](#)
- \*12.9. [Changes from draft-nir-ike-qcd-02](#)
- \*12.10. [Changes from draft-nir-ike-qcd-01](#)
- \*12.11. [Changes from draft-nir-ike-qcd-00](#)
- \*12.12. [Changes from draft-nir-qcr-00](#)
- \*13. [References](#)
- \*13.1. [Normative References](#)
- \*13.2. [Informative References](#)
- \*Appendix A. [The Path Not Taken](#)
- \*Appendix A.1. [Initiating a new IKE SA](#)
- \*Appendix A.2. [SIR](#)
- \*Appendix A.3. [Birth Certificates](#)
- \*Appendix A.4. [Reducing Liveness Check Length](#)
- \*[Authors' Addresses](#)

## **[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 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., different environment and situation requiring different rules).

Some existing IKEv2 implementations already do that (i.e., both shorten timeouts or limit number of retries) based on these kind of hints and also start liveness checks quickly after the other end goes silent. However, see [Appendix 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 extensions'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:



```

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+]

                    / --> EAP
repeat 1..N times |
                    \ <-- 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 which 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 ESP or 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:

- \*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.
- \*It should not be possible for an external attacker to guess the QCD token generated by an implementation. Cryptographic mechanisms such as PRNG and hash functions are RECOMMENDED.
- \*The token maker MUST be able to re-generate or retrieve the token based on the IKE SPIs even after it reboots.
- \*The method of token generation MUST be such that a collision of QCD tokens between different pairs of IKE SPI will be highly unlikely.

### 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.

- \*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.
- \*Those 32 bytes, called the "QCD\_SECRET", are stored in non-volatile storage on the machine, and kept indefinitely.
- \*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.
- \*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.

## [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 stand-by gateway ready to take over as soon as the primary gateway fails for any reason. [\[cluster\]](#) 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 stand-by 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 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 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 the this discussion:

\*For remote-access clients it makes sense to implement the token taker role.

\*For remote-access gateways it makes sense to implement the token maker role.

\*For inter-domain VPN gateways it makes sense to implement both roles, because it can't be known in advance where the traffic originates.

\*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 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.

## 8.2. Response to unknown child SPI

After a reboot, it is more likely that an implementation receives 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 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, she 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 which 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 balancer 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 it 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:

- \*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.
- \*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.
- \*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 is requested to assign a notify message type from the status types range (16406-40959) of the "IKEv2 Notify Message Types" registry with 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. Change Log](#)

This section lists all changes in this document

NOTE TO RFC EDITOR : Please remove this section in the final RFC

### [12.1. Changes from draft-ietf-ipsecme-failure-detection-05](#)

\*Some clarifications suggested by Magnus Nystrom.

### [12.2. Changes from draft-ietf-ipsecme-failure-detection-04](#)

\*Some more rephrasing of section 9.2 based on suggestions by Tero Kivinen and Dave Wierbowski.

### [12.3. Changes from draft-ietf-ipsecme-failure-detection-03](#)

\*Merged section 9.4 into section 9.2.

\*Multiple typos discovered by Scott Moonen, Keith Welter and Yaron.

### [12.4. Changes from draft-ietf-ipsecme-failure-detection-02](#)

\*Moved section 7 to Appendix A. Also changed some wording.

\*Fixed some language in the "interaction with session resumption" section to say that although liveness check MUST be done, there are no time limits to how long an implementation takes before starting liveness check, or ending it.

### [12.5. Changes from draft-ietf-ipsecme-failure-detection-01](#)

\*Fixed the language requiring random IKE SPIs.

\*Some better explanation of the reasons to choose the methods in [Section 5.2](#) and the method in [Section 5.1](#), to close issue #193.

\*Added text to the beginning of [Section 9](#) to accomodate issue #194.

## [12.6. Changes from draft-ietf-ipsecme-failure-detection-00](#)

\*Nits pointed out by Scott and Yaron.

\*Pratima and Frederic are back on board.

\*Changed IKEv2bis draft reference to RFC 5996.

\*Resolved issues #189, #190, #191, and #192:

- Renamed section 4.5 and removed the requirement to send an acknowledgement for the unprotected message.

- Moved the QCD token from the last to the first IKE\_AUTH request.

- Added a MUST to [Section 9.3](#) to require that IKE SPIs be randomly generated.

- Changed the language in [Section 8.1](#), to not use RFC 2119 terminology.

- Moved the section describing why one would want the method dependant on IP addresses (in [Section 5.2](#) from operational considerations to security considerations.

## [12.7. Changes from draft-nir-ike-qcd-07](#)

\*First WG version.

\*Addressed Scott C Moonen's concern about collisions of QCD tokens.

\*Updated references to point to IKEv2bis instead of RFC 4306 and 4718. Also converted draft reference for resumption to RFC 5723.

\*Added Dave Wiebrowski as author, and removed Pratima and Frederic.

## [12.8. Changes from draft-nir-ike-qcd-03 and -04](#)

Mostly editorial changes and cleaning up.

## [12.9. Changes from draft-nir-ike-qcd-02](#)

\*Described QCD token enumeration, following a question by Lakshminath Dondeti.

\*Added the ability to replace the QCD token for an existing IKE SA.

\*Added tokens dependent on peer IP address and their interaction with MOBIKE.

#### 12.10. Changes from draft-nir-ike-qcd-01

\*Removed stateless method.

\*Added discussion of rekeying and resumption.

\*Added discussion of non-synchronized load-balanced clusters of gateways in the security considerations.

\*Other wording fixes.

#### 12.11. Changes from draft-nir-ike-qcd-00

\*Merged proposal with draft-detienne-ikev2-recovery

\*Changed the protocol so that the rebooted peer generates the token. This has the effect, that the need for persistent storage is eliminated.

\*Added discussion of birth certificates.

#### 12.12. Changes from draft-nir-qcr-00

\*Changed name to reflect that this relates to IKE. Also changed from quick crash recovery to quick crash detection to avoid confusion with IFARE.

\*Added more operational considerations.

\*Added interaction with IFARE.

\*Added discussion of backup gateways.

### 13. References

#### 13.1. Normative References

[1]	<a href="#">Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels"</a> , BCP 14, RFC 2119, March 1997.
[2]	Kaufman, C, Hoffman, P, Nir, Y and P Eronen, " <a href="#">Internet Key Exchange Protocol: IKEv2</a> ", RFC 5996, September 2010.
[3]	Eronen, P., " <a href="#">IKEv2 Mobility and Multihoming Protocol (MOBIKE)</a> ", RFC 4555, June 2006.

#### 13.2. Informative References

[1]	
-----	--

	Sheffer, Y. and H. Tschofenig, " <a href="#">IKEv2 Session Resumption</a> ", RFC 5723, January 2010.
[2]	Detienne, F., Sethi, P. and Y. Nir, " <a href="#">Safe IKE Recovery</a> ", Internet-Draft draft-detienne-ikev2-recovery, January 2010.
[3]	Nir, Y., " <a href="#">IPsec Cluster Problem Statement</a> ", Internet-Draft draft-ietf-ipsecme-ipsec-ha, July 2010.

## [Appendix A. The Path Not Taken](#)

### [Appendix 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 piggy-backed 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.

### [Appendix 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.

### [Appendix 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.

#### [Appendix 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 Nir Check Point Software Technologies Ltd. 5  
Hasolelim st. Tel Aviv, 67897 Israel EMail: [ynir@checkpoint.com](mailto:ynir@checkpoint.com)

David Wierbowski Wierbowski International Business Machines 1701  
North Street Endicott, New York 13760 United States EMail:  
[wierbows@us.ibm.com](mailto:wierbows@us.ibm.com)

Frederic Detienne Detienne Cisco Systems, Inc. De Kleetlaan, 7  
Diegem, B-1831 Belgium Phone: +32 2 704 5681 EMail: [fd@cisco.com](mailto:fd@cisco.com)

Pratima Sethi Sethi Cisco Systems, Inc. O'Shaughnessy Road, 11  
Bangalore, Karnataka 560027 India Phone: +91 80 4154 1654 EMail:  
[psethi@cisco.com](mailto:psethi@cisco.com)