

[MS-IKEE]:

Internet Key Exchange Protocol Extensions

Intellectual Property Rights Notice for Open Specifications Documentation

- **Technical Documentation.** Microsoft publishes Open Specifications documentation for protocols, file formats, languages, standards as well as overviews of the interaction among each of these technologies.
- **Copyrights.** This documentation is covered by Microsoft copyrights. Regardless of any other terms that are contained in the terms of use for the Microsoft website that hosts this documentation, you may make copies of it in order to develop implementations of the technologies described in the Open Specifications and may distribute portions of it in your implementations using these technologies or your documentation as necessary to properly document the implementation. You may also distribute in your implementation, with or without modification, any schema, IDL's, or code samples that are included in the documentation. This permission also applies to any documents that are referenced in the Open Specifications.
- **No Trade Secrets.** Microsoft does not claim any trade secret rights in this documentation.
- **Patents.** Microsoft has patents that may cover your implementations of the technologies described in the Open Specifications. Neither this notice nor Microsoft's delivery of the documentation grants any licenses under those or any other Microsoft patents. However, a given Open Specification may be covered by Microsoft [Open Specification Promise](#) or the [Community Promise](#). If you would prefer a written license, or if the technologies described in the Open Specifications are not covered by the Open Specifications Promise or Community Promise, as applicable, patent licenses are available by contacting iplg@microsoft.com.
- **Trademarks.** The names of companies and products contained in this documentation may be covered by trademarks or similar intellectual property rights. This notice does not grant any licenses under those rights. For a list of Microsoft trademarks, visit www.microsoft.com/trademarks.
- **Fictitious Names.** The example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted in this documentation are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

Reservation of Rights. All other rights are reserved, and this notice does not grant any rights other than specifically described above, whether by implication, estoppel, or otherwise.

Tools. The Open Specifications do not require the use of Microsoft programming tools or programming environments in order for you to develop an implementation. If you have access to Microsoft programming tools and environments you are free to take advantage of them. Certain Open Specifications are intended for use in conjunction with publicly available standard specifications and network programming art, and assumes that the reader either is familiar with the aforementioned material or has immediate access to it.

Revision Summary

Date	Revision History	Revision Class	Comments
10/22/2006	0.01		Version 0.01 release
1/19/2007	1.0		Version 1.0 release
3/2/2007	1.1		Version 1.1 release
4/3/2007	1.2		Version 1.2 release
5/11/2007	1.3		Version 1.3 release
6/1/2007	1.3.1	Editorial	Changed language and formatting in the technical content.
7/3/2007	2.0	Major	Updated and revised the technical content.
7/20/2007	2.0.1	Editorial	Changed language and formatting in the technical content.
8/10/2007	3.0	Major	Updated and revised the technical content.
9/28/2007	3.0.1	Editorial	Changed language and formatting in the technical content.
10/23/2007	3.0.2	Editorial	Changed language and formatting in the technical content.
11/30/2007	3.0.3	Editorial	Changed language and formatting in the technical content.
1/25/2008	4.0	Major	Updated and revised the technical content.
3/14/2008	4.0.1	Editorial	Changed language and formatting in the technical content.
5/16/2008	4.0.2	Editorial	Changed language and formatting in the technical content.
6/20/2008	5.0	Major	Updated and revised the technical content.
7/25/2008	6.0	Major	Updated and revised the technical content.
8/29/2008	6.1	Minor	Clarified the meaning of the technical content.
10/24/2008	6.2	Minor	Clarified the meaning of the technical content.
12/5/2008	7.0	Major	Updated and revised the technical content.
1/16/2009	8.0	Major	Updated and revised the technical content.
2/27/2009	9.0	Major	Updated and revised the technical content.
4/10/2009	10.0	Major	Updated and revised the technical content.
5/22/2009	11.0	Major	Updated and revised the technical content.
7/2/2009	12.0	Major	Updated and revised the technical content.
8/14/2009	12.1	Minor	Clarified the meaning of the technical content.
9/25/2009	12.2	Minor	Clarified the meaning of the technical content.
11/6/2009	13.0	Major	Updated and revised the technical content.
12/18/2009	13.1	Minor	Clarified the meaning of the technical content.
1/29/2010	14.0	Major	Updated and revised the technical content.

Date	Revision History	Revision Class	Comments
3/12/2010	15.0	Major	Updated and revised the technical content.
4/23/2010	16.0	Major	Updated and revised the technical content.
6/4/2010	17.0	Major	Updated and revised the technical content.
7/16/2010	18.0	Major	Updated and revised the technical content.
8/27/2010	18.0	None	No changes to the meaning, language, or formatting of the technical content.
10/8/2010	18.0	None	No changes to the meaning, language, or formatting of the technical content.
11/19/2010	18.0	None	No changes to the meaning, language, or formatting of the technical content.
1/7/2011	18.1	Minor	Clarified the meaning of the technical content.
2/11/2011	18.1	None	No changes to the meaning, language, or formatting of the technical content.
3/25/2011	18.1	None	No changes to the meaning, language, or formatting of the technical content.
5/6/2011	18.1	None	No changes to the meaning, language, or formatting of the technical content.
6/17/2011	18.2	Minor	Clarified the meaning of the technical content.
9/23/2011	18.2	None	No changes to the meaning, language, or formatting of the technical content.
12/16/2011	19.0	Major	Updated and revised the technical content.
3/30/2012	19.0	None	No changes to the meaning, language, or formatting of the technical content.
7/12/2012	19.1	Minor	Clarified the meaning of the technical content.
10/25/2012	19.1	None	No changes to the meaning, language, or formatting of the technical content.
1/31/2013	20.0	Major	Updated and revised the technical content.
8/8/2013	21.0	Major	Updated and revised the technical content.
11/14/2013	21.0	None	No changes to the meaning, language, or formatting of the technical content.
2/13/2014	22.0	Major	Updated and revised the technical content.
5/15/2014	22.0	None	No changes to the meaning, language, or formatting of the technical content.
6/30/2015	23.0	Major	Significantly changed the technical content.

Table of Contents

1	Introduction	8
1.1	Glossary	8
1.2	References	11
1.2.1	Normative References	11
1.2.2	Informative References	12
1.3	Overview	13
1.3.1	Network Address Translation Traversal (NAT-T).....	13
1.3.2	IKE Fragmentation.....	14
1.3.3	Authentication Using a Cryptographically Generated Address	14
1.3.4	Fast Failover	15
1.3.5	Negotiation Discovery	15
1.3.6	Reliable Delete	15
1.3.7	Denial of Service Protection	16
1.3.8	IKE/AuthIP Co-Existence	16
1.3.9	IKE SA Correlation (IKEv2)	16
1.3.10	IKE Server Internal Addresses Configuration Attributes (IKEv2)	16
1.3.11	Xbox Multiplayer Gaming (IKEv2)	16
1.3.12	IPsec Security Realm (IKEv2 transport mode)	16
1.3.13	Extension to RFC Cross Reference.....	17
1.4	Relationship to Other Protocols	18
1.5	Prerequisites/Preconditions	18
1.5.1	General Prerequisites/Preconditions	18
1.5.2	CGA Authentication Prerequisites/Preconditions	18
1.6	Applicability Statement	18
1.7	Versioning and Capability Negotiation	19
1.8	Vendor-Extensible Fields	19
1.9	Standards Assignments.....	19
2	Messages.....	20
2.1	Transport.....	20
2.2	Message Syntax.....	20
2.2.1	NAT-T Payload Types	20
2.2.2	NAT-T UDP Encapsulation Modes	20
2.2.3	IKE Message Fragment.....	21
2.2.3.1	Fragment Payload Packet	21
2.2.4	AUTH_CGA Authentication Method Packet	22
2.2.5	ID_IPV6_CGA Identification Type Packet	22
2.2.6	Notify Payload Packet.....	23
2.2.7	Notify Payload (IKEv2) Packet	25
2.2.8	Configuration Attribute (IKEv2) Packet	26
2.2.9	Correlation Payload (IKEv2) Packet	26
2.2.10	Security Realm Vendor ID Payload (IKEv2)	27
3	Protocol Details	28
3.1	Common Details	28
3.1.1	Abstract Data Model.....	28
3.1.2	Timers	29
3.1.3	Initialization.....	29
3.1.4	Higher-Layer Triggered Events	30
3.1.5	Message Processing Events and Sequencing Rules	30
3.1.6	Timer Events.....	31
3.1.7	Other Local Events.....	31
3.2	NAT Traversal Details	31
3.2.1	Abstract Data Model.....	32
3.2.2	Timers	32

3.2.3	Initialization.....	32
3.2.4	Higher-Layer Triggered Events	33
3.2.4.1	Start of an IKE MM SA Negotiation	33
3.2.5	Message Processing Events and Sequencing Rules	33
3.2.5.1	Receiving Message #1	33
3.2.5.2	Receiving Message #2	33
3.2.5.3	Receiving Other Messages.....	33
3.2.6	Timer Events.....	33
3.2.7	Other Local Events.....	34
3.3	IKE Fragmentation Details.....	34
3.3.1	Abstract Data Model.....	34
3.3.2	Timers	35
3.3.3	Initialization.....	35
3.3.4	Higher-Layer Triggered Events	36
3.3.4.1	Start of an IKE MM SA Negotiation	36
3.3.5	Message Processing Events and Sequencing Rules	36
3.3.5.1	Receiving Message #1	36
3.3.5.2	Receiving Message #2	36
3.3.5.3	Receiving Other IKE Messages.....	36
3.3.6	Timer Events.....	37
3.3.6.1	Expiration of Fragmentation Timer	37
3.3.6.2	Expiration of the Fragment Reassembly Timer	37
3.3.7	Other Local Events.....	38
3.4	CGA Authentication Details.....	38
3.4.1	Abstract Data Model.....	39
3.4.2	Timers	39
3.4.3	Initialization.....	39
3.4.4	Higher-Layer Triggered Events	39
3.4.4.1	Start of an IKE MM SA Negotiation	39
3.4.5	Message Processing Events and Sequencing Rules	40
3.4.5.1	Receiving Message #1	40
3.4.5.2	Receiving Message #2	40
3.4.5.3	Receiving Message #3	40
3.4.5.4	Receiving Message #4	40
3.4.5.5	Receiving Message #5	40
3.4.5.6	Receiving Message #6	41
3.4.6	Timer Events.....	41
3.4.7	Other Local Events.....	41
3.5	Fast Failover Client Details	41
3.5.1	Abstract Data Model.....	42
3.5.2	Timers	42
3.5.3	Initialization.....	42
3.5.4	Higher-Layer Triggered Events	42
3.5.4.1	Start of an IKE MM SA Negotiation	42
3.5.5	Message Processing Events and Sequencing Rules	42
3.5.5.1	Receiving Message #1	42
3.5.5.2	Receiving Message #2	42
3.5.6	Timer Events.....	43
3.5.6.1	Expiration of the QM SA Idle Timer.....	43
3.5.7	Other Local Events.....	43
3.5.7.1	Successful Negotiation of a QM SA	43
3.6	Fast Failover Server Details	43
3.6.1	Abstract Data Model.....	43
3.6.2	Timers	43
3.6.3	Initialization.....	43
3.6.4	Higher-Layer Triggered Events	44
3.6.4.1	Start of an IKE MM SA Negotiation	44
3.6.5	Message Processing Events and Sequencing Rules	44

3.6.5.1	Receiving Message #1	44
3.6.5.2	Receiving Message #2	44
3.6.6	Timer Events.....	44
3.6.7	Other Local Events.....	44
3.7	Negotiation Discovery Details	44
3.7.1	Abstract Data Model.....	48
3.7.2	Timers	49
3.7.3	Initialization.....	49
3.7.4	Higher-Layer Triggered Events	49
3.7.4.1	Outbound Packet	49
3.7.4.2	Inbound Packet.....	50
3.7.5	Message Processing Events and Sequencing Rules	51
3.7.5.1	Receiving Message #1	51
3.7.5.2	Receiving Message #2	51
3.7.5.3	Receiving Message #5.....	51
3.7.5.4	Receiving Message #6	52
3.7.6	Timer Events.....	52
3.7.7	Other Local Events.....	52
3.8	Reliable Delete Details	52
3.8.1	Abstract Data Model.....	52
3.8.2	Timers	53
3.8.3	Initialization.....	53
3.8.4	Higher-Layer Triggered Events	53
3.8.4.1	SA Deletion/Invalidation	53
3.8.5	Message Processing Events and Sequencing Rules	53
3.8.5.1	Receiving Message #1	53
3.8.5.2	Receiving Message #2	54
3.8.6	Timer Events.....	54
3.8.6.1	Expiration of the Delete Retransmission Timer	54
3.8.7	Other Local Events.....	54
3.8.7.1	Shutdown	55
3.8.7.2	MM SA Exhaustion.....	55
3.9	Denial of Service Protection Details	55
3.9.1	Abstract Data Model.....	56
3.9.2	Timers	56
3.9.3	Initialization.....	56
3.9.4	Higher-Layer Triggered Events	56
3.9.5	Message Processing Events and Sequencing Rules	56
3.9.5.1	Receiving Message #1	56
3.9.5.2	Receiving Message #2	56
3.9.5.3	Receiving Message #3	57
3.9.6	Timer Events.....	57
3.9.7	Other Local Events.....	57
3.10	IKE SA Correlation (IKEV2) Details	57
3.10.1	Abstract Data Model.....	58
3.10.2	Timers	58
3.10.3	Initialization.....	58
3.10.4	Higher-Layer Triggered Events	58
3.10.5	Message Processing Events and Sequencing Rules	58
3.10.5.1	Receiving Message #1	59
3.10.5.2	Receiving Subsequent Messages	59
3.10.5.3	Receiving the Error Notify	60
3.10.6	Timer Events.....	60
3.10.7	Other Local Events.....	60
3.11	IKE Server Internal Addresses Configuration Attributes (IKEv2) Details	60
3.11.1	Abstract Data Model.....	60
3.11.2	Timers	60
3.11.3	Initialization.....	60

3.11.4	Higher-Layer Triggered Events	60
3.11.5	Message Processing Events and Sequencing Rules	61
3.11.5.1	Receiving Message #1	61
3.11.5.2	Receiving Message #2	62
3.11.6	Timer Events.....	62
3.11.7	Other Local Events.....	62
3.12	Dead Peer Detection Details	62
3.12.1	Abstract Data Model.....	62
3.12.2	Timers	62
3.12.3	Initialization.....	62
3.12.4	Higher-Layer Triggered Events	63
3.12.4.1	TCP Dead Peer Detection	63
3.12.4.2	UDP Dead Peer Detection.....	63
3.12.5	Message Processing Events and Sequencing Rules	63
3.12.5.1	Receiving a UDP Packet	63
3.12.6	Timer Events.....	63
3.12.6.1	Expiration of the QM SA Idle Timer.....	63
3.12.7	Other Local Events.....	63
3.12.7.1	Successful Negotiation of a QM SA and MM SA.....	63
3.13	Xbox Multiplayer Gaming (IKEv2) Vendor IDs Details	64
3.13.1	Abstract Data Model.....	64
3.13.2	Timers	64
3.13.3	Initialization.....	64
3.13.4	Higher-Layer Triggered Events	64
3.13.5	Message Processing Events and Sequencing Rules	64
3.13.5.1	Microsoft Xbox One 2013 Vendor ID	65
3.13.5.2	Xbox IKEv2 Negotiation Vendor ID	65
3.13.6	Timer Events.....	65
3.13.7	Other Local Events.....	65
3.14	Security Realm ID (IKEv2) Vendor IDs Details	65
3.14.1	Abstract Data Model.....	65
3.14.2	Timers	65
3.14.3	Initialization.....	65
3.14.4	Higher-Layer Triggered Events	65
3.14.5	Message Processing Events and Sequencing Rules	66
3.14.5.1	IKE_SA_INIT Messages.....	67
3.14.5.2	IKE_SA_AUTH and CREATE_CHILD_SA Messages.....	67
3.14.6	Timer Events.....	67
3.14.7	Other Local Events.....	67
4	Protocol Examples	68
4.1	Negotiation Discovery Examples.....	68
5	Security	70
5.1	Security Considerations for Implementers	70
5.1.1	Negotiation Discovery	70
5.2	Index of Security Parameters	70
6	Appendix A: Product Behavior	71
7	Change Tracking.....	91
8	Index.....	94

1 Introduction

Internet Key Exchange (IKE) Protocol Extensions apply to the IKE Protocol versions 1 and 2, as specified in [\[RFC2407\]](#), [\[RFC2408\]](#), [\[RFC2409\]](#), [\[RFC3947\]](#), and [\[RFC4306\]](#). These extensions provide additional capabilities to **IKE**, including interoperation between different revisions of the **network address translation** traversal (NAT-Traversal or NAT-T) specification, fragmentation of large IKE version 1 messages, authentication by using **cryptographically generated addresses (CGAs)**, fast failover when communicating with a **cluster** of hosts, easier interoperation with non-**Internet Protocol security (IPsec)**-capable peers, acknowledgment of **security association (SA)** deletion messages, denial of service protection, IKE security association correlation (IKEv2), and IKE server internal addresses configuration attributes (IKEv2).

Sections 1.8, 2, and 3 of this specification are normative and can contain the terms MAY, SHOULD, MUST, MUST NOT, and SHOULD NOT as defined in [\[RFC2119\]](#). Sections 1.5 and 1.9 are also normative but do not contain those terms. All other sections and examples in this specification are informative.

1.1 Glossary

The following terms are specific to this document:

Authenticated IP (AuthIP): An **Internet Key Exchange (IKE)** protocol extension, as specified in [\[MS-AIPS\]](#).

authentication header (AH): An **Internet Protocol Security (IPsec)** encapsulation mode that provides authentication and message integrity. For more information, see [\[RFC4302\]](#) section 1.

certificate: A certificate is a collection of attributes (1) and extensions that can be stored persistently. The set of attributes in a certificate can vary depending on the intended usage of the certificate. A certificate securely binds a public key to the entity that holds the corresponding private key. A certificate is commonly used for authentication (2) and secure exchange of information on open networks, such as the Internet, extranets, and intranets. Certificates are digitally signed by the issuing certification authority (CA) and can be issued for a user, a computer, or a service. The most widely accepted format for certificates is defined by the ITU-T X.509 version 3 international standards. For more information about attributes and extensions, see [\[RFC3280\]](#) and [\[X509\]](#) sections 7 and 8.

certificate chain: A sequence of **certificates**, where each certificate in the sequence is signed by the subsequent certificate. The last certificate in the chain is normally a self-signed certificate.

cluster: A group of computers that are able to dynamically assign resource tasks among nodes in a group. The group can be accessed as though they are a single host. A cluster is generally accessed by using a virtual IP address. For more information, see [\[MSFT-WLBS\]](#).

cryptographic hash function: A function that maps an input of any length to a short output bit string of fixed length, such that finding an input that maps to a particular bit string of the correct output length, or even finding two inputs that map to the same output bit string, is computationally infeasible. For more information, see [\[SCHNEIER\]](#) chapters 2 and 18.

cryptographically generated address (CGA): An IPv6 address for which the interface identifiers (the low-order 64 bits) are generated by computing a **cryptographic hash function** on a public key. The corresponding private key can be used to sign messages sent from this IPv6 address. **CGA** is specified in [\[RFC3972\]](#).

domain of interpretation (DOI): A domain that defines the manner in which a group of protocols uses the **ISAKMP** (as specified in [\[RFC2408\]](#)) framework to negotiate **security associations (SAs)** (for example, identifiers for cryptographic algorithms, interpretation of payload contents, and so on). For example, the Internet Protocol security (IPsec) **DOI** (as specified in [\[RFC2407\]](#))

defines the use of the **ISAKMP** framework for protocols that negotiate **main mode (MM)** and **quick mode security associations (SAs)**. Both **Internet Key Exchange (IKE)** and **AuthIP** fall under the IPsec **DOI**.

Encapsulating Security Payload (ESP): An **Internet Protocol security (IPsec)** encapsulation mode that provides authentication, data confidentiality, and message integrity. For more information, see [\[RFC4303\]](#) section 1.

exchange: A pair of messages, consisting of a request and a response.

flow: A TCP session or User Datagram Protocol (UDP) pseudo session, identified by a 5-tuple (source and destination IP and ports, and protocol). By extension, a request/response Internet Control Message Protocol (ICMP) exchange (for example, ICMP echo) is also a **flow**.

Generic Security Services (GSS): An Internet standard, as described in [\[RFC2743\]](#), for providing security services to applications. It consists of an application programming interface (GSS-API) set, as well as standards that describe the structure of the security data.

initiator: The party that sends the first message of an **Internet Key Exchange (IKE)**.

Internet Key Exchange (IKE): The protocol that is used to negotiate and provide authenticated keying material for **security associations (SAs)** in a protected manner. For more information, see [\[RFC2409\]](#).

Internet Protocol security (IPsec): A framework of open standards for ensuring private, secure communications over Internet Protocol (IP) networks through the use of cryptographic security services. IPsec supports network-level peer authentication, data origin authentication, data integrity, data confidentiality (encryption), and replay protection. The Microsoft implementation of IPsec is based on standards developed by the Internet Engineering Task Force (IETF) IPsec working group.

Internet Security Association and Key Management Protocol (ISAKMP): A cryptographic protocol specified in [\[RFC2408\]](#) that defines procedures and packet formats to establish, negotiate, modify and delete **security associations (SAs)**. It forms the basis of the **Internet Key Exchange (IKE)** protocol, as specified in [\[RFC2409\]](#).

ISAKMP payload: A modular building block for constructing **ISAKMP** messages. A payload is used to transfer information such as **security association (SA)** data, or key generation and authentication data. The presence and order of payloads in a packet is defined by and dependent upon the type of exchange specified in the **ISAKMP** header of the **ISAKMP** message. For more information, see [\[RFC2408\]](#) section 4.1.

main mode (MM): The first phase of an **Internet Key Exchange (IKE)** negotiation that performs authentication and negotiates a **main mode security association (MM SA)** between the peers. For more information, see [\[RFC2409\]](#) section 5.

main mode security association (MM SA): A security association that is used to protect **Internet Key Exchange (IKE)** traffic between two peers. For more information, see [\[RFC2408\]](#) section 2.

main mode security association database (MMSAD): A database that contains operational state for each **main mode (MM) security association (SA)**. For more information, see [\[MS-AIPS\]](#) section 3.1.1 and [\[MS-IKEE\]](#) section 3.1.1.

maximum transmission unit (MTU): The size, in bytes, of the largest packet that a given layer of a communications protocol can pass onward.

negotiation: A series of exchanges. The successful outcome of a **negotiation** is the establishment of one or more **security associations (SAs)**. For more information, see [\[RFC2408\]](#) section 2.

negotiation discovery: An **Internet Key Exchange (IKE)** extension that improves interoperability between **Internet Protocol security (IPsec)** and non-IPsec-aware hosts. Detecting that the peer host is not capable of **IPsec** usually involves waiting for the **IKE** negotiation to time out, then sending traffic in the clear. With **negotiation discovery**, the host starts the **IKE** negotiation and sends clear text traffic in parallel. If the **IKE** negotiation succeeds and **security associations (SAs)** are established, further traffic is secured.

network address translation (NAT): The process of converting between IP addresses used within an intranet, or other private network, and Internet IP addresses.

nonce: A number that is used only once. This is typically implemented as a random number large enough that the probability of number reuse is extremely small. A nonce is used in authentication protocols to prevent replay attacks. For more information, see [\[RFC2617\]](#).

phase: A series of exchanges that provide a particular set of security services (for example, authentication or creation of **security associations (SAs)**).

quick mode: The second phase of an **Internet Key Exchange (IKE)** negotiation, during which the peers negotiate **quick mode security associations (QM SAs)**. For more information, see [\[RFC2409\]](#) section 5.5.

quick mode security association (QM SA): A **security association (SA)** that is used to protect IP packets between peers (the **Internet Key Exchange (IKE)** traffic is protected by the **main mode security association (MM SA)**). For more information, see [\[RFC2409\]](#) section 5.5.

responder: The computer that responds to request messages.

root certificate: A self-signed **certificate** that identifies the public key of a root certification authority (CA) and has been trusted to terminate a **certificate chain**.

security association (SA): A simplex "connection" that provides security services to the traffic carried by it. See [\[RFC4301\]](#) for more information.

security association database (SAD): A database that contains parameters that are associated with each established (keyed) **security association**.

security policy database (SPD): A database that specifies the policies that determine the disposition of all IP traffic inbound or outbound from a host or security gateway.

self-signed certificate: A **certificate** that is signed by its creator and verified using the public key contained in it. Such certificates are also termed **root certificates**.

transport mode: An IP encapsulation mechanism, as specified in [\[RFC4301\]](#), that provides **Internet Protocol security (IPsec)** security for host-to-host communication.

tunnel mode: An IP encapsulation mechanism, as specified in [\[RFC4301\]](#), that provides Internet Protocol security (IPsec) security to tunneled IP packets. IPsec processing is performed by the tunnel endpoints, which can be (but are typically not) the end hosts.

vendor ID payload: A particular type of **ISAKMP payload** that contains a vendor-defined constant. The constant is used by vendors to identify and recognize remote instances of their implementations. This mechanism allows a vendor to experiment with new features while maintaining backward compatibility. For more information, see [\[RFC2408\]](#) section 3.16.

MAY, SHOULD, MUST, SHOULD NOT, MUST NOT: These terms (in all caps) are used as defined in [\[RFC2119\]](#). All statements of optional behavior use either MAY, SHOULD, or SHOULD NOT.

1.2 References

Links to a document in the Microsoft Open Specifications library point to the correct section in the most recently published version of the referenced document. However, because individual documents in the library are not updated at the same time, the section numbers in the documents may not match. You can confirm the correct section numbering by checking the [Errata](#).

1.2.1 Normative References

We conduct frequent surveys of the normative references to assure their continued availability. If you have any issue with finding a normative reference, please contact dochelp@microsoft.com. We will assist you in finding the relevant information.

[ECP] Fu, D. and Solinas, J., "ECP Groups For IKE and IKEv2", September 2005, <http://tools.ietf.org/id/draft-ietf-ipsec-ike-ecp-groups-02.txt>

[GSS] Piper, D., and Swander, B., "A GSS-API Authentication Method for IKE", Internet Draft, July 2001, <http://tools.ietf.org/html/draft-ietf-ipsec-isakmp-gss-auth-07>

[IANAIPSEC] IANA, "Internet Key Exchange (IKE) Attributes", November 2006, <http://www.iana.org/assignments/ipsec-registry>

[IANAISAKMP] IANA, "'Magic Numbers' for ISAKMP Protocol", October 2006, <http://www.iana.org/assignments/isakmp-registry>

[MS-AIPS] Microsoft Corporation, "[Authenticated Internet Protocol](#)".

[MS-ERREF] Microsoft Corporation, "[Windows Error Codes](#)".

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <http://www.rfc-editor.org/rfc/rfc2119.txt>

[RFC2403] Madson, C. and Glenn, R., "The Use of HMAC-MD5-96 Within ESP and AH", RFC 2403, November 1998, <http://www.ietf.org/rfc/rfc2403.txt>

[RFC2407] Piper, D., "The Internet IP Security Domain of Interpretation for ISAKMP", RFC 2407, November 1998, <http://www.ietf.org/rfc/rfc2407.txt>

[RFC2408] Maughan, D., Schertler, M., Schneider, M., and Turner, J., "Internet Security Association and Key Management Protocol (ISAKMP)", RFC 2408, November 1998, <http://www.ietf.org/rfc/rfc2408.txt>

[RFC2409] Harkins, D. and Carrel, D., "The Internet Key Exchange (IKE)", RFC 2409, November 1998, <http://www.ietf.org/rfc/rfc2409.txt>

[RFC2451] Pereira, R. and Adams, R., "The ESP CBC-Mode Cipher Algorithms", RFC 2451, November 1998, <http://www.ietf.org/rfc/rfc2451.txt>

[RFC3447] Jonsson, J. and Kaliski, B., "Public-Key Cryptography Standards (PKCS) #1: RSA Cryptography Specifications Version 2.1", RFC 3447, February 2003, <http://www.ietf.org/rfc/rfc3447.txt>

[RFC3526] Kivinen, T. and Kojo, M., "More Modular Exponential (MODP) Diffie-Hellman Groups for Internet Key Exchange (IKE)", RFC 3526, May 2003, <http://www.ietf.org/rfc/rfc3526.txt>

[RFC3947] Kivinen, T., Swander, B., Huttunen, A., and Volpe, V., "Negotiation of NAT-Traversal in the IKE", RFC 3947, January 2005, <http://www.ietf.org/rfc/rfc3947.txt>

[RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005, <http://www.ietf.org/rfc/rfc3972.txt>

[RFC4301] Kent, S. and Seo, K., "Security Architecture for the Internet Protocol", RFC 4301, December 2005, <http://www.ietf.org/rfc/rfc4301.txt>

[RFC4306] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005, <http://www.ietf.org/rfc/rfc4306.txt>

[RFC4555] P. Eronen, Ed., "IKEv2 Mobility and Multihoming Protocol (MOBIKE)", RFC 4555, June 2006, <http://www.ietf.org/rfc/rfc4555.txt>

[RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and Eronen, P., "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010, <http://tools.ietf.org/html/rfc5996>

[RFC768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980, <http://www.ietf.org/rfc/rfc768.txt>

[RFC792] Postel, J., "Internet Control Message Protocol", RFC 792, September 1981, <http://www.ietf.org/rfc/rfc792.txt>

1.2.2 Informative References

[DRAFT-NATT] Kivinen, T., Huttunen, A., Swander, B., and Volpe, V., "Negotiation of NAT-Traversal in the IKE", June 2002, <http://tools.ietf.org/id/draft-ietf-ipsec-nat-t-ike-03.txt>

[FIPS140] FIPS PUBS, "Security Requirements for Cryptographic Modules", FIPS PUB 140, December 2002, <http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf>

[MSFT-WLBS] Microsoft Corporation, "MS Windows NT Load Balancing Service (WLBS)", January 1999, <http://www.microsoft.com/technet/archive/winntas/deploy/depovg/wlbsdepl.msp?mfr=true>

[RFC2404] Madson, C. and Glenn, R., "The Use of HMAC-SHA-1-96 Within ESP and AH", RFC 2404, November 1998, <http://www.ietf.org/rfc/rfc2404.txt>

[RFC2405] Madson, C. and Doraswamy, N., "The ESP DES-CBC Cipher Algorithm With Explicit IV", RFC 2405, November 1998, <http://www.ietf.org/rfc/rfc2405.txt>

[RFC2410] Glenn, R. and Kent, S., "The NULL Encryption Algorithm and Its Use With IPsec", RFC 2410, November 1998, <http://www.ietf.org/rfc/rfc2410.txt>

[RFC3602] Frankel, S., Glenn, R., and Kelly, S., "The AES-CBC Cipher Algorithm and Its Use with IPsec", RFC 3602, September 2003, <http://www.ietf.org/rfc/rfc3602.txt>

[RFC3715] Aboba, B. and Dixon, W., "IPsec-Network Address Translation (NAT) Compatibility Requirements", RFC 3715, March 2004, <http://www.ietf.org/rfc/rfc3715.txt>

[RFC3948] Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and Stenberg, M., "UDP Encapsulation of IPsec ESP Packets", RFC 3948, January 2005, <http://www.ietf.org/rfc/rfc3948.txt>

[RFC4106] Viega, J. and McGrew, D., "The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP)", RFC 4106, June 2005, <http://www.ietf.org/rfc/rfc4106.txt>

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005, <http://www.ietf.org/rfc/rfc4302.txt>

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005, <http://www.ietf.org/rfc/rfc4303.txt>

[RFC4543] McGrew, D., and Viega, J., "The Use of Galois Message Authentication Code (GMAC) in IPsec ESP and AH", RFC 4543, May 2006, <http://www.ietf.org/rfc/rfc4543.txt>

[RFC4621] Kivinen, T., and Tschofenig, H., "Design of the IKEv2 Mobility and Multihoming (MOBIKE) Protocol", RFC 4621, August 2006, <http://www.ietf.org/rfc/rfc4621.txt>

[RFC791] Postel, J., Ed., "Internet Protocol: DARPA Internet Program Protocol Specification", RFC 791, September 1981, <http://www.rfc-editor.org/rfc/rfc791.txt>

[SHA256] National Institute of Standards and Technology, "FIPS 180-2, Secure Hash Standard (SHS)", August 2002, <http://csrc.nist.gov/publications/fips/fips180-2/fips180-2withchangenotice.pdf>

1.3 Overview

The Internet Key Exchange (IKE) Protocol version 1 is used to negotiate security associations (SAs), as specified in [RFC2409], for the purpose of keying **authentication header (AH)** and **Encapsulating Security Payload (ESP)** packet transformations. For more information, see [RFC4302] and [RFC4303], respectively. For the general security architecture of IPsec, see [RFC4301].

The IKE Protocol version 1 is specified in [RFC2409] and is closely tied to [RFC2407] and [RFC2408]. In addition, IKE is clearly the most commonly implemented protocol that uses [RFC2407] and [RFC2408]. Also, version 2 of the IKE protocol is specified by a single Request for Comments [RFC4306]. For these reasons, industry practice supports use of the term IKE to collectively refer to [RFC2407], [RFC2408], [RFC2409], and more recently, [RFC4306].

In the remainder of this document, the term IKE collectively applies to [RFC2407], [RFC2408], [RFC2409], and [RFC4306]. Where applicable, the appropriate section of each RFC is referenced in the document. <1>

This document specifies the extensions to IKE. Each of these IKE extensions is independent and can be implemented in isolation. There is no sequencing between the individual extensions. An implementation of this protocol can support any combination of these IKE extensions. <2>

1.3.1 Network Address Translation Traversal (NAT-T)

In the original IPsec specifications, the interposition of network address translation (NAT) devices between IPsec peers prevents correct IPsec operation. For more information about the incompatibilities, see [RFC3715] section 2.

Two specifications have been defined to address these incompatibilities. For more information about the User Datagram Protocol (UDP) encapsulation of ESP packets, see [RFC3948]. UDP-encapsulated ESP packets are correctly translated by NAT devices. [RFC3947] specifies an IKE extension to detect the presence of NAT devices between two IPsec peers and to negotiate the use of a UDP-encapsulated ESP.

Network address translation traversal (NAT-T) **negotiation** for IKE was first published as an Internet draft before becoming [RFC3947]. In [DRAFT-NATT], the IKE parameter numbers for NAT-T negotiation are chosen from the appropriate private use ranges, as specified in [IANASAKMP]. In specification [RFC3947], different IKE parameter numbers were assigned by the Internet Assigned Numbers Authority (IANA). As a result, a [DRAFT-NATT]-compliant implementation is incompatible with an [RFC3947]-compliant implementation. For more information, see [DRAFT-NATT].

The NAT-T extension specified in this document enables IKE implementations supporting NAT-T to negotiate the use of either the [DRAFT-NATT] or the [RFC3947] parameters. This specification does not extend the NAT-T protocol itself. It negotiates only the interpretation of the NAT-T IKE parameter numbers. Also, this document specifies the support of NAT-T IKE for IPsec **transport mode** only.

The extension negotiates the use of the [DRAFT-NATT] or [RFC3947] parameters as follows:

1. The host signals which revisions of the specification it supports (that is, [DRAFT-NATT], [RFC3947], or both) by sending **vendor ID payloads** ("RFC 3947" or "draft-ietf-ipsec-nat-t-ike-02\n") with its first IKE message. See section [1.7](#), Capability Negotiation.
2. On receipt of the first IKE message from the peer, the host looks up the vendor ID payloads to determine which revision of the NAT-T protocol to use. If both revisions are supported by both hosts, preference is given to [RFC3947] over [DRAFT-NATT].

For details, see section [3.2](#).

1.3.2 IKE Fragmentation

IKE uses UDP as a transport. IKE messages can be sufficiently large; so the underlying IP layer may fragment them, as described in [\[RFC791\]](#) section 2.3. This fragmentation typically happens with IKE messages that contain **certificate chains**. To avoid fragmentation-based attacks, fragmented UDP packets are commonly blocked by firewalls and routers. Blocking the fragmented UDP packets can lead to IKE failures that are especially difficult to diagnose. The IKE fragmentation extension that is specified in this document avoids fragmentation at the IP level by fragmenting IKE packets into smaller UDP packets that the underlying IP layer is guaranteed not to fragment.

Hosts that support IKE fragmentation advertise this capability through a "FRAGMENTATION" vendor ID payload; for more information, see section [1.7](#). If both peers support fragmentation, a fragmentation timer is started whenever a message is sent. If the timer expires, it is assumed that the message that is associated with the timer did not reach its destination because it was too large to traverse the intervening network. In this case, the message is split into several small fragments, and all these small fragments are sent.

So that the destination host can correctly reassemble the fragmented message, each fragment carries a fragment ID that is unique to the original message and a fragment number that is unique to the particular fragment. Fragment numbers range from 1 to N, where N is the number of fragments for a message.

Upon receipt of a fragment, the receiving host verifies whether it has already received other fragments for that fragment ID. If not, the receiving host starts a reassembly timer. It then verifies whether it has received all N fragments for the message, where the Nth fragment is indicated by a particular bit in the fragment. If the fragment reassembly timer expires before all fragments are correctly received, the receiving host must discard all fragments.

For details, see section [3.3](#).

1.3.3 Authentication Using a Cryptographically Generated Address

This extension specifies a new authentication method for IKE based on cryptographically generated addresses (CGAs), as specified in [\[RFC3972\]](#). A CGA is an IPv6 address for which the interface identifier (that is, the low-order 64 bits) is generated by computing a **cryptographic hash function** of a public key.

Hosts that support CGA authentication advertise their capability through an "IKE CGA version 1" vendor ID payload. CGA authentication is negotiated as a regular IKE authentication method; see section [1.7](#), Capability Negotiation. The CGA verification that occurs during this authentication ensures that the remote peer has access to the private key that was used to generate the CGA. This CGA verification uses the corresponding public key and a parameters structure that contains information originally used to generate the CGA. The public key and parameters structure must, therefore, be sent to the host that verifies the CGA. The public key is transmitted within an IKE **certificate** payload, and the parameters structure is transmitted by using a new CGA identification payload as part of the IKE

main mode (MM) negotiation. Successful validation of the CGA completes the IKE main mode negotiation.

For details, see section [3.4](#).

1.3.4 Fast Failover

This extension reduces the time required for a client to restore an IPsec security association (SA) to the virtual IP address for a cluster of hosts after a failure on one of the hosts that is sharing the virtual IP address.

The client uses a "Vid-Initial-Contact" vendor ID payload (see section [1.7](#), Capability Negotiation) to signal to the cluster that it does not have any **main mode security association (MM SA)** or **quick mode security association (QM SA)** established with the cluster so that the IKE session may be reallocated to a different node within the cluster. The server uses an "NLBS_PRESENT" vendor ID payload (see section [1.7](#), Capability Negotiation) to indicate to the client that the client should use a shorter **quick mode** idle timer. In this way, a new QM SA is renegotiated faster if a failover occurs.

For more information about clusters based on virtual IP addresses, see [\[MSFT-WLBS\]](#). For specifications, see sections [3.5](#) and [3.6](#).

1.3.5 Negotiation Discovery

IKE Protocol Extensions enable a client to determine whether a remote peer supports IPsec-protected communications.

Negotiation discovery introduces new IPsec policy options. In the case of outbound traffic, if the traffic matches a negotiation discovery policy, the host sends the packet in Cleartext and starts an IKE negotiation in parallel. If the remote peer is not IPsec-capable, the IKE negotiation eventually times out, and the connection stays in Cleartext. If the peer is IPsec-capable and the IKE negotiation eventually succeeds, the connection starts using the negotiated SA. To enforce that a once-secured **flow** can never downgrade back to Cleartext, this extension maintains a per-flow state table that is looked up for every packet.

In the case of inbound traffic, negotiation discovery supports a policy-specified boundary mode in which the host can accept both Cleartext and secured connections to allow inbound traffic from non-IPsec-capable hosts in addition to secure connections from IPsec-capable hosts. The flow state table determines if an incoming Cleartext packet should be accepted.

For details, see section [3.7](#).

1.3.6 Reliable Delete

This extension enables a peer to reliably confirm the deletion of a security association that is established with another peer. The original IKE specification does not require the acknowledgment of Delete payloads.

This capability is advertised through additional **ISAKMP payloads**. The standard IKE Delete message is sent with an additional **ISAKMP** Nonce payload (as specified in [\[RFC2408\]](#) section 3.13) appended. The host starts a retransmission timer when sending the Delete message. On receipt of the Delete message, the host constructs an acknowledgment message that contains an ISAKMP Nonce payload, an ISAKMP Delete payload, and the Message ID from the received Delete message in the ISAKMP header. On receipt of the acknowledgment message, the host verifies that the Message ID matches the Message ID that was sent with the Delete message. On expiration of the retransmission timer, the Delete message is retransmitted.

For details, see section [3.8](#).

1.3.7 Denial of Service Protection

A **responder** that implements the IKE protocol must create states for all correctly formed initial requests, even if the **initiator** is flooding the responder with packets from multiple incorrect IP addresses. The vulnerability to denial-of-service (DOS) attacks is mitigated if responders do not create any state until the peer can prove that it exists at a routable address.

This extension enables a responder to delay creating state until it has verified the following:

1. That the source of a message is not a spoofed IP address.
2. When a threshold of incoming requests has been reached.

For details, see section [3.9](#).

1.3.8 IKE/AuthIP Co-Existence

This extension allows two peers that are both IKEv1 and authenticated IP (**AuthIP**)-capable to negotiate the use of AuthIP over IKEv1. This extension is specified in [\[MS-AIPS\]](#) section 1.7 and also applies to IKE. [<3>](#)

1.3.9 IKE SA Correlation (IKEv2)

This extension allows two different IKEv2 IKE_SA to be correlated together. Assume that an IKE_SA has been established. This is called SAoriginal. At a later time, to ensure that the client credentials are still valid, but without tearing down the existing SA, a new IKE_SA (called SAcurrent) can be built to embed a new payload in this **exchange** that securely correlates this SA with the original SA.

1.3.10 IKE Server Internal Addresses Configuration Attributes (IKEv2)

This extension allows the IKEv2 client endpoint of an IPsec remote access client (IRAC), as specified in [\[RFC4306\]](#) section 2.19, to determine the internal IPv4 and IPv6 addresses of the IPsec remote access server (IRAS), as also specified in [\[RFC4306\]](#) section 2.19.

1.3.11 Xbox Multiplayer Gaming (IKEv2)

This extension is used by two IKEv2 peers negotiating SAs for Xbox multiplayer gaming scenarios. There are two vendor ID payloads used for this extension. The first vendor ID payload, "Microsoft Xbox One 2013", is used by an IKEv2 initiator endpoint to show that this SA negotiation is for Xbox multiplayer gaming. The second vendor ID payload, "Xbox IKEv2 Negotiation", and an associated identifier are used by negotiating peers to distinguish between various types of multiplayer gaming secure connections and to do some throttling based on the type. Details of these extensions are specified in section [3.13](#).

1.3.12 IPsec Security Realm (IKEv2 transport mode)

An IPsec "Security Realm" defines per-application IPsec policies and the set of "related" applications whose network traffic is secured by these policies. The security realm refers to the common set of crypto settings used for IPsec SA negotiation, and the credentials used for authentication. Details of this extension are specified in section [3.14](#).

This extension is used by two IKEv2 peers negotiating transport mode SAs for scenarios involving per-application IPsec policies. This extension uses a vendor ID payload called "MSFT IPsec Security Realm Id". The vendor ID payload is associated with a 16-byte identifier. This identifier is used as an optional selector to choose an appropriate IPsec policy for negotiation.

If the message from the initiator for negotiating the child SA does not have a "MSFT IPsec Security Realm Id" vendor ID, but the parent IKE SA is associated with a security realm policy, then this message will be discarded by the responder and the child SA negotiation will be failed.

1.3.13 Extension to RFC Cross Reference

The following table summarizes how each IKE extension extends each of the applicable RFCs.

IKE extension	Extends [RFC2407]	Extends [RFC2408]	Extends [RFC2409]	Extends [RFC3947]	Extends [RFC4306]	IKE version
NAT-T transport mode only	(1)	(2) (3)		(7)		IKEv1
IKE fragmentation		(3)	(8)			IKEv1
CGA authentication	(4) (5)	(3)	(9)			IKEv1
Fast failover		(3)	(10)			IKEv1
Negotiation discovery		(3) (6)	(10)			IKEv1
Reliable delete			(11)			IKEv1
Denial of Service protection		(6)	(12)			IKEv1
IKE SA Correlation					(13)	IKEv2
Configuration Attribute					(14)	IKEv2

1. Adjunction of an encapsulation mode in the private range. Encapsulation mode is specified in [\[RFC2407\]](#) section 4.5.
2. Adjunction of a vendor ID. Vendor ID is as specified in [\[RFC2408\]](#) section 3.16.
3. Adjunction of Payload Types in the private range. Payload Types are specified in [RFC2408] section 3.1.
4. Adjunction of an authentication method within an ISAKMP SA payload, as specified in [RFC2407] section 4.6.1.
5. Adjunction of an identification type for an ISAKMP Identification payload from the private Identification Type range, as specified in [RFC2407] section 4.6.2.
6. Adjunction of a notify message type from the private range. The notify message types are specified in [RFC2408] section 3.14.1.
7. Negotiation of the interpretation of Payload Types and Encapsulation Modes.
8. Fragmentation and reassembly. Packet construction and decoding for IKE are specified in [\[RFC2409\]](#) section 5.
9. Extends the IKE phase 1 exchange using certificates. For more information, see [RFC2409] section 5.1.

10. Extends the IKE phase 1 exchange. For more information, see [RFC2409] section 5. Extends the QM SAs negotiation. For more information, see [RFC2409] section 5.5.
11. Extends the Notify exchange. For more information, see [RFC2409] section 5.7.
12. Extends the IKE phase 1 exchange. For more information, see [RFC2409] section 5.1.
13. This extension allows two different IKEv2 IKE_SA to be correlated together for the purpose of ensuring that the client credentials are still valid but without tearing down the existing SA. When validation is required, a new IKE_SA (called SA_{current}) can be built to embed a new payload in this exchange that securely correlates this SA with the original SA.
14. This extension allows the IKEv2 client endpoint of an IPsec remote access client (IRAC), as specified in [RFC4306], to determine the internal IPv4 and IPv6 addresses of the IPsec remote access server (IRAS), also as specified in [RFC4306].

1.4 Relationship to Other Protocols

IKE is used for the authentication and keying of IPsec SAs, as specified in [RFC4301] section 3. IKE relies on UDP as a transport, as specified in [RFC768].

1.5 Prerequisites/Preconditions

The following sections describe the prerequisites and preconditions for using IKE protocol extensions:

- [General Prerequisites/Preconditions \(section 1.5.1\)](#)
- [CGA Authentication Prerequisites/Preconditions \(section 1.5.2\)](#)

1.5.1 General Prerequisites/Preconditions

IKE assumes that both the initiator and the responder have an IP address and have UDP connectivity. IKE also assumes that the initiator knows the responder's IP address (for example, through manual configuration or through a policy lookup in the case of **tunnel mode**).

Successful establishment of a QM SA using IKEv1 requires that the initiator and the responder have at least one common authentication method and a common set of cryptographic parameters for the MM and the QM SAs. For authentication using certificates, each peer validates the remote peer certificate chain to a locally trusted **root certificate**, as specified in [RFC2409] section 5.1. For pre-shared key authentication, both peers are required to share the same pre-shared secret, as specified in [RFC2409] section 5.4.

1.5.2 CGA Authentication Prerequisites/Preconditions

For CGA authentication, as specified in [RFC3972] section 1, the peers must possess a CGA and the associated **self-signed certificate**.

1.6 Applicability Statement

- NAT-T applies when NAT devices between the IPsec peers can otherwise prevent the establishment of IPsec SAs.
- IKE fragmentation applies when intermediary devices in the path between the IPsec peers can drop fragmented UDP datagrams, that can prevent the establishment of IPsec SAs.
- Authentication using CGA applies when the IPsec peers do not share a common credential distribution infrastructure. CGA authentication allows such peers to verify that the remote peer

has access to the public-private key pair used to generate the CGA. CGA authentication only applies to IPv6 addresses.

- Fast failover applies when IPsec clients connect to a cluster of hosts using IPsec, and it is necessary to minimize the amount of time required for a client to failover from one host in the cluster to another.
- Negotiation discovery applies when hosts communicate with both IPsec-aware and non-IPsec-aware devices, and it is necessary to minimize the amount of time required to detect IPsec-awareness on each peer.
- Reliable delete applies when a peer needs to reliably confirm the deletion of a security association (SA) established with another peer.
- IKEv2 SA Correlation applies when two different IKEv2 SAs need to be correlated.
- IKEv2 Server Internal Addresses Configuration Attributes apply when the client endpoint of an IPsec remote access client needs to determine the internal IPv4 and IPv6 addresses of the IPsec remote access server.

1.7 Versioning and Capability Negotiation

This section covers versioning issues in the following areas:

- **Protocol Versions:** The protocol version is part of the ISAKMP header. IKEv1 uses protocol version 1.0, as specified in [\[RFC2408\]](#) section 3.1. IKEv2 uses protocol version 2.0, as specified in [\[RFC4306\]](#) section 3.1.
- **Security and Authentication Methods:** IKE supports multiple authentication and encryption algorithms for both the MM SAs and QM SAs, as specified in [\[RFC2408\]](#) section 5.6. IKE supports the negotiation of the authentication method, the Diffie-Hellman group, and the hashing and authentication algorithm using [\[RFC2409\]](#), [\[GSS\]](#), or [\[RFC3972\]](#).<4>
- **Cryptographic Parameters:** Cryptographic parameters are negotiated in different **phases** of the protocol (that is, initial exchange, MM, and quick mode, as specified in [\[RFC2409\]](#) section 5). Details about algorithm and parameter numbers are specified in [\[IANAIPSEC\]](#) and [\[IANAISAKMP\]](#).<5>
- **Capability Negotiation:** IKE can advertise specific capabilities through vendor ID payloads, as specified in [\[RFC2408\]](#) section 3.16.<6>

1.8 Vendor-Extensible Fields

The IKE extensions specified in this document do not introduce any new vendor-extensible fields. These extensions inherit the extensibility features of ISAKMP (as specified in [\[RFC2408\]](#)) and IKE (as specified in [\[RFC2409\]](#)).

1.9 Standards Assignments

No standards assignments have been received for the IKE extensions described in this document. All values used in these extensions are in private ranges, as specified in [\[IANAIPSEC\]](#) and [\[IANAISAKMP\]](#).

2 Messages

2.1 Transport

IKE messages MUST be transported over ISAKMP, as specified in [\[RFC2408\]](#), which uses UDP port 500 by default. IKE MUST run over ports 500 and 4500 if a NAT has been detected, as specified in [\[RFC3947\]](#) section 3.2; otherwise, it MAY be run over a different port. <7>

All fields are sent and encoded in network order unless otherwise specified.

2.2 Message Syntax

2.2.1 NAT-T Payload Types

Each ISAKMP message consists of a header and a variable number of payloads, each identified by a 1-octet payload type value in its Payload Type field, as specified in [\[RFC2408\]](#) section 3.1. NAT-T adds two new payload types: NAT Discovery (NAT-D) and NAT Original Address (NAT-OA). The payload type values for these payload types are specified in [\[RFC3947\]](#). For more information about an alternative set of payload type values, see [\[DRAFT-NATT\]](#). <8>

The following table describes the NAT-D payload type.

NAT Discovery (NAT-D) payload type value	Revision
0x82	[DRAFT-NATT]
0x14	[RFC3947]

The following table describes the supported NAT-OA payload type.

Supported NAT Original Address (NAT-OA) payload type	Revision
0x83	[DRAFT-NATT]
0x15	[RFC3947]

2.2.2 NAT-T UDP Encapsulation Modes

The Encapsulation Mode field is located in the SA payload, as specified in [\[RFC2407\]](#) section 4.5. Specification [\[RFC3947\]](#) introduces new encapsulation mode values for this field. For more information about an alternative set of these values, see [\[DRAFT-NATT\]](#). <9>

The following table describes the UDP-Encapsulated-Tunnel values.

UDP-Encapsulated-Tunnel	Revision
0xF003	[DRAFT-NATT]
0x0003	[RFC3947]

The following table describes the UDP-Encapsulated-Transport values.

UDP-Encapsulated-Transport	Revision
0xF004	[DRAFT-NATT]
0x0004	[RFC3947]

2.2.3 IKE Message Fragment

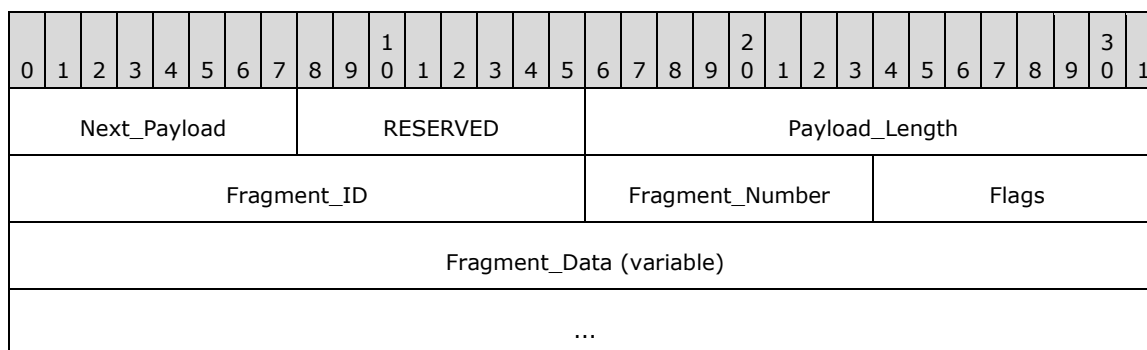
An IKE message fragment contains:

- An ISAKMP header, as specified in [\[RFC2408\]](#) section 3.1.
- A single, non-encrypted, Fragment payload.

2.2.3.1 Fragment Payload Packet

The Fragment Payload is an ISAKMP payload, as specified in [\[RFC2408\]](#) section 3.1. The Payload Type value for a Fragment payload is 0x84 from the private payload type range, as specified in [\[RFC2408\]](#) section 3.1. A Fragment payload MUST be preceded by an ISAKMP header that has this payload type.

The following illustration describes the Fragment Payload packet.



Next_Payload (1 byte): Identifier for the payload type, which MUST specify the next payload in the message. For a Fragment payload, this field MUST be set to 0.

RESERVED (1 byte): This field MUST be set to zero. The responder MUST ignore this field on receipt. This behavior is identical to IKE.

Payload_Length (2 bytes): This field MUST be the length, in bytes, of the payload, including the generic payload header. This is identical to IKE.

Fragment_ID (2 bytes): The **Fragment ID** field is 2 bytes that MUST specify the same value for every fragment that is generated from a particular IKE message.

Fragment_Number (1 byte): The **Fragment Number** field MUST indicate the order in which the fragments are sent. The first fragment MUST have a fragment number of 1, and each subsequent fragment MUST have a fragment number that is one greater than that of the previous fragment. Because the maximum size of an IKE message is limited to 64 KB by UDP and fragments are aligned on the minimum **MTU** for IPv4 and IPv6, the fragment number cannot wrap.

Flags (1 byte): The flag field MUST have the following value.

Value	Meaning
LAST_FRAGMENT 0x01	This flag indicates the last fragment in the message.

All other bits of the Flags field MUST be set to zero on the initiator and ignored on the responder. For more details on flag semantics, see section [3.1](#).

Fragment_Data (variable): The **Fragment Data** field MUST contain the fragment. The size of the **Fragment Data** field MUST be computed by subtracting the size of the Fragment Payload header (8 bytes) from the value of the **Payload Length** field.

2.2.4 AUTH_CGA Authentication Method Packet

AUTH_CGA is an authentication method within an ISAKMP SA payload, as specified in [RFC2407](#) section 4.6.1. The format of the SA payload is the following, as specified in [RFC2408](#) section 3.4.

- A number of Proposal payloads, as specified in [RFC2408](#) section 3.5.
- Within each Proposal payload, there is a number of Transform payloads, as specified in [RFC2408](#) section 3.6.
- Within each Proposal payload, there is a number of Data Attributes payloads, as specified in [RFC2408](#) section 3.3. In a Data Attribute payload, an authentication method is indicated by the value 0x0003 in the Attribute Type field of the Data Attribute payload, as specified in [RFC2409](#) Appendix A. The particular authentication method is determined by the value of the Attribute Value field, as specified in [RFC2409](#) Appendix A.

The Data Attribute payload for the AUTH_CGA Authentication method has the format seen in the following AUTH_CGA packet.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
A	Attribute_Type															Attribute_Value															

A - One (1 bit): This field MUST be set to 1.

Attribute_Type (15 bits): For the AUTH_CGA authentication method, this field MUST be set to the value 0x0003. This value corresponds to the authentication method, as specified in [RFC2409](#) Appendix A.

Attribute_Value (2 bytes): For the AUTH_CGA authentication method, this field MUST be set to the value 0xFDED in network order. This value is from the private authentication method range, as specified in [RFC2409](#) Appendix A.

2.2.5 ID_IPV6_CGA Identification Type Packet

ID_IPV6_CGA is an identification type for an ISAKMP Identification payload, as specified in [RFC2407](#) section 4.6.2. The ID_IPV6_CGA Identification Type is 0xFA from the private Identification Type range, as specified in [IANAISAKMP](#).

The format of the Identification payload for an ID_IPV6_CGA identification type is seen in the following packet.

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										RESERVED										Payload_Length											
Identification_Type										Protocol_ID										Port											
Modifier (16 bytes)																															
...																															
...																															
Collision_Count										Extension_fields (variable)																					
...																															

Next_Payload (1 byte): This field is the identifier for the payload type of the next payload in the message. This field MUST be identical to the corresponding IKE field.

RESERVED (1 byte): This field MUST be set to zero. The responder MUST ignore this field on receipt. This behavior is identical to IKE.

Payload_Length (2 bytes): This field MUST be the length in bytes of the payload, including the Generic Payload header. This is identical to IKE.

Identification_Type (1 byte): This field is the value describing how the fields after the Port field should be interpreted. The ID_IPV6_CGA identification type MUST be 0xFA, from the private Identification Type range, as specified in [IANAISAKMP].

Protocol_ID (1 byte): This field MUST be set to zero. The responder MUST ignore this field on receipt. This is identical to IKE.

Port (2 bytes): This field MUST be set to zero. The responder MUST ignore this field on receipt. This is identical to IKE.

Modifier (16 bytes): This field MUST be as specified in [RFC3972] section 3.

Collision_Count (1 byte): This field MUST be as specified in [RFC3972] section 3.

Extension_fields (variable): This field MUST be as specified in [RFC3972] section 3.

2.2.6 Notify Payload Packet

The Notify Payload packet is specified in [RFC2408] section 3.14. The format is as follows.

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										RESERVED										Payload_Length											
Domain_of_Interpretation																															
Protocol-ID										SPI_size										Notify_Message_Type											

Security_Parameter_Index (variable)
...
Notification_Data (variable)
...

Next_Payload (1 byte): This field MUST be as specified in [RFC2408] section 3.14.

RESERVED (1 byte): This field MUST be as specified in [RFC2408] section 3.14.

Payload_Length (2 bytes): This field MUST be as specified in [RFC2408] section 3.14.

Domain_of_Interpretation (4 bytes): The **domain of interpretation (DOI)** field MUST be set to 1 (IPSEC_DOI) as specified in [RFC2408] section A.2.

Protocol-ID (1 byte): This field MUST be as specified in [RFC2408] section 3.14.

SPI_size (1 byte): This field MUST be as specified in [RFC2408] section 3.14. The **SPI_size** is updated to a value of 8 when the Message ID is appended to the notification data as described in this section under **Notification_Data**.

Notify_Message_Type (2 bytes): This MUST identify the type of notification being sent with this message, in network byte order. The notify message types MUST be one of the following values, which are from the private range, as specified in [RFC2408] section 3.14.1.

Value	Meaning
0x9C43	NOTIFY_STATUS (check) This notify message type is a status code indicating the failure to establish a security association (SA) with a peer.
0x9C44	NOTIFY_DOS_COOKIE (check) This notify message type is used by the DOS protection extension.
0x9C45	EXCHANGE_INFO This notify message type is used by the negotiation discovery extension.

Security_Parameter_Index (variable): This is the Security Parameter Index (SPI) of size SPI_size. This field MUST be as specified in [RFC2408] section 3.14.

Notification_Data (variable): The content of this field depends on the **Notify_Message_Type** field. The following list describes field content for various notify message types. If the peer has previously sent the Vendor ID "MS NT5 ISAKMPOAKLEY" as specified in the footnote regarding Capability Negotiation in section 1.7, and the notify corresponds to the quick mode exchange, then the Message ID (in network order) of the quick mode is appended as the first 4 bytes of the notification data. In particular, the NOTIFY_DOS_COOKIE will never have the Message ID in the notification data because that is always a main mode operation. The EXCHANGE_INFO notify will always have the Message ID appended if the peer sends the above vendor ID. The NOTIFY_STATUS will only have the Message ID appended if the failure is a quick mode failure.

Field content MUST correspond to the **Notify_Message_Type** as follows:

- NOTIFY_STATUS (4 Bytes): MUST be a status code indicating failure. The values transmitted as status codes are implementation-specific. <10>
- NOTIFY_DOS_COOKIE (8 Bytes): MUST be the responder cookie value.
- EXCHANGE_INFO (4 Bytes): The flag values MUST be one of the following values.

Value	Meaning
0x00000001	IKE_EXCHANGE_INFO_ND_BOUNDARY This flag is used by the negotiation discovery extension.
0x00000002	IKE_EXCHANGE_INFO_GUARANTEE_ENCRYPTION This flag is used by the negotiation discovery extension.

2.2.7 Notify Payload (IKEv2) Packet

The Notify Payload packet is specified in [RFC4306](#) section 3.10. The format is as follows.

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		
Protocol-ID										SPI_size										Notify_Message_Type													
SPI																																	
Notification_Data (variable)																																	
...																																	

Protocol-ID (1 byte): This field MUST be as specified in [RFC4306] section 3.10.

SPI_size (1 byte): This field MUST be as specified in [RFC4306] section 3.10.

Notify_Message_Type (2 bytes): This MUST identify the type of notification being sent with this message, in network byte order. The notify message types MUST be one of the following values, which are from the private error range, as specified in [RFC4306] section 3.10.1.

Value	Meaning
0x3039	Notify status. This notify message type is used to tell the peer of a private failure reason.

SPI (4 bytes): The Security Parameter Index (SPI) field MUST be as specified in [RFC4306] section 3.10.

Notification_Data (variable): The content of this field depends on the **Notify_Message_Type** field. The following list describes field content for various notify message types. Field content MUST correspond to the notify message type as follows:

- NOTIFY_STATUS (4 bytes): MUST be a status code indicating failure. The values transmitted as status codes are implementation specific. <11>

2.2.8 Configuration Attribute (IKEv2) Packet

The Configuration Attribute packet is specified in [\[RFC4306\]](#) section 3.15.1. The format is as follows.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
R	Attribute Type															Length															
Value (variable)																															
...																															

R (1 bit): This reserved field MUST be as specified in [\[RFC4306\]](#) section 3.15.1.

Attribute Type (15 bits): This field MUST be as specified in [\[RFC4306\]](#) section 3.15.1.

Length (2 bytes): The length of the data in the value field.

Value (variable): The internal IPv4 or IPv6 address of the server.

Two additional **Attribute Types** from the private-use range are defined as follows.

Attribute type	Length (bytes)	Value
INTERNAL_IP4_SERVER 0x5BA0	4	The internal IPv4 address of the server.
INTERNAL_IP6_SERVER 0x5BA1	16	The internal IPv6 address of the server.

2.2.9 Correlation Payload (IKEv2) Packet

The Correlation Payload (IKEv2) packet format is as follows. There are two IKE_SAs here, SA_{current} and SA_{original}. This payload is sent under the protection of SA_{current}. The Payload Type value for a Correlation payload is 0xc8 from the private payload type range, as specified in [\[RFC4306\]](#) section 3.2.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Next_Payload										RESERVED										Payload_Length											
IKE_SA_Initiator_SPI																															
...																															
IKE_SA_Responder_SPI																															
...																															
Correlation_Hash (variable)																															

...

Next_Payload (1 byte): This field MUST be as specified in [\[RFC2408\]](#) section 3.2.

RESERVED (1 byte): This field MUST be as specified in [\[RFC2408\]](#) section 3.2.

Payload_Length (2 bytes): This field MUST be as specified in [\[RFC2408\]](#) section 3.2.

IKE_SA_Initiator_SPI (8 bytes): This MUST be set to the initiator's SPI from the IKE_SA being correlated, SAoriginal. This value is taken from the IKEv2 header of the prior IKE_SA, as specified in [\[RFC4306\]](#) section 3.1.

IKE_SA_Responder_SPI (8 bytes): This MUST be set to the responder's SPI from the IKE_SA being correlated, SAoriginal. This value is taken from the IKEv2 header of the prior IKE_SA, as specified in [\[RFC4306\]](#) section 3.1.

Correlation_Hash (variable): This computes a keyed hash using the SAcurrent's negotiated PRF function. The key used is the SK_ai on the initiator and the SK_ar for the responder from SAoriginal. See [\[RFC4306\]](#) section 2.14. The correlation hash is as follows.

```
prf(SK_a(i or r),  
SAcurrent.InitiatorSpi|SAcurrent.ResponderSpi|SAoriginal.InitiatorSpi|SAoriginal.responde  
rSpi)
```

2.2.10 Security Realm Vendor ID Payload (IKEv2)

The "MSFT IPsec Security Realm Id" vendor ID payload is constructed as specified in [\[RFC5996\]](#) section 3.12. The vendor ID payload has a variable length field called Vendor ID or VID. In the case of this extension, the first 16 bytes is an MD5 hash of the string "MSFT IPsec Security Realm Id". The subsequent bytes contain the actual Security Realm ID. [<12>](#)

3 Protocol Details

The following sections specify protocol details, including abstract data models and message processing rules, that are common and that are specific to NAT-T, IKE fragmentation, CGAs, the fast-failover client, the fast-failover server, negotiation discovery, reliable delete, IKE SA correlation (IKEv2), IKE Server Internal Addresses Configuration Attributes (IKEv2), and DOS protection.

3.1 Common Details

This section documents deviations from "The Internet IP Security Domain of Interpretation for ISAKMP", as specified in [\[RFC2407\]](#); "Internet Security Association and Key Management Protocol (ISAKMP)", as specified in [\[RFC2408\]](#); "The Internet Key Exchange (IKE)", as specified in [\[RFC2409\]](#); "Internet Key Exchange (IKEv2) Protocol", as specified in [\[RFC4306\]](#); and "Negotiation of NAT-Traversal in the IKE", as specified in [\[RFC3947\]](#). These deviations affect each of these RFC standards as described in the table in section [1.3.13](#).

The flags bit semantics used by this document are as follows: for a flag, its "value" signifies a mask which, when its bitwise logical AND with the flags field is computed, yields either a zero value (all zero bits) if the flag is unset (set to FALSE), and a nonzero value otherwise. For example, a flag mask/value of 0x01 signifies that the bitwise logical AND of a single-byte flag field with 0x01 is zero if and only if the flag is set to FALSE. Assuming no other flag masks/values for this field, then, both 0x00 and 0x01 are valid values for this single-byte flag field: the former corresponding to the flag being unset, and the latter to the flag being set.

3.1.1 Abstract Data Model

This section describes a conceptual model of possible data organization that an implementation maintains to participate in this protocol in addition to what is specified in [\[RFC2407\]](#), [\[RFC2408\]](#), [\[RFC2409\]](#), [\[RFC3947\]](#), and [\[RFC4301\]](#) for IKEv1, or [\[RFC4306\]](#) for IKEv2. The described organization is provided to explain how the protocol behaves. This document does not mandate that implementations adhere to this model as long as their external behavior is consistent with the behavior described in this document.

The following main data elements are required by any implementation:

- **Main mode security association database (MMSAD):** A database that contains the operational state for each MM SA. The entry for each MM SA contains the following data elements.

For each IKE MM SA, the following information MUST be maintained:

- All states that are necessary for managing a standard IKE MM SA as defined in [\[RFC2409\]](#) appendix A for IKEv1 and [\[RFC4306\]](#) section 3.3.2 for IKEv2.
- All states that are necessary for management of other IKE extensions for the SA, as specified in this section and in sections [3.2.1](#), [3.3.1](#), [3.4.1](#), [3.5.1](#), [3.6.1](#), [3.7.1](#), [3.8.1](#) for IKEv1 only, and [3.10.1](#) for IKEv2 only.

The MMSAD MUST be indexed by the local and peer IP addresses and the initiator and responder cookies found in the ISAKMP header, as specified in [\[RFC2408\]](#).

- Peer authorization database (PAD): The PAD and its management operations are specified in [\[RFC4301\]](#) section 4.4.3. This specification does not extend that definition. The PAD that is referred to in this specification contains rules that describe if and how IKE should negotiate SAs with a remote peer, as specified in [\[RFC4301\]](#).

All states that are necessary for the management of IKE extensions are described in section [3.4.1](#) for IKEv1 only.

The PAD MUST be looked up by using tuples that are composed of local and remote IP addresses.

- **Security policy database (SPD):** The SPD and its management operations are specified in [RFC4301] section 4.4.1. The SPD that is referred to in this specification contains rules that describe if and how IPsec protection is applied to inbound or outbound IP traffic. The SPD MUST be looked up by using tuples that are composed of flow information (that is, source and destination IP addresses, port numbers, and protocol) for the packet.

All states that are necessary for management of IKE extensions are described in section 3.7.1 for IKEv1 only.

- **Security association database (SAD):** The SAD contains the parameters of each QM SA. The SAD and its management operations are specified in [RFC4301] section 4.4.2.

All states that are necessary for management of IKE extensions are described in section 3.7.1 for IKEv1 only.

- **Connection state table:** Stores a set of connection entries. These connection entries correspond to active TCP/UDP/ICMP or protocol-only connections.

The possible connection entries are:

- V4 TCP/UDP state entry: {IPv4 source address {DWORD}, IPv4 destination address {DWORD}, IP protocol {DWORD}, source port {DWORD}, destination port {DWORD}}.
- V6 TCP/UDP state entry: {IPv6 source address {16 bytes}, IPv6 destination address {16 bytes}, IP protocol {DWORD}, source port {DWORD}, destination port {DWORD}}.
- V4 ICMP state entry: {IPv4 source address {DWORD}, IPv4 destination address {DWORD}, IP protocol {DWORD}, ICMP type {DWORD}, ICMP code {DWORD}}, as defined in [RFC792].
- V6 ICMP state entry: {IPv6 source address {16 bytes}, IPv6 destination address {16 bytes}, IP protocol {DWORD}, ICMP type {DWORD}, ICMP code {DWORD}}, as defined in [RFC792].
- V4 protocol-only state entry: {IPv4 source address {DWORD}, IPv4 destination address {DWORD}, IP protocol {DWORD}}.
- V6 protocol-only state entry: {IPv6 source address {16 bytes}, IPv6 destination address {16 bytes}, IP protocol {DWORD}}.

All states that are necessary for management of IKE extensions are described in section 3.7.1 for IKEv1 only.

- **Other states:** Additional states are defined in section 3.9.1 and section 3.11.1.

Note The preceding conceptual data can be implemented by using a variety of techniques. Any data structure that stores the preceding conceptual data may be used in the implementation.

3.1.2 Timers

None beyond what is specified in [RFC2407], [RFC2408], [RFC2409], [RFC3947], or [RFC4306].

3.1.3 Initialization

None beyond what is specified in [RFC2407], [RFC2408], [RFC2409], [RFC3947], or [RFC4306].

3.1.4 Higher-Layer Triggered Events

None except what is specified in [\[RFC2407\]](#), [\[RFC2408\]](#), [\[RFC2409\]](#), [\[RFC3947\]](#), or [\[RFC4306\]](#).

3.1.5 Message Processing Events and Sequencing Rules

[\[RFC2407\]](#): Message processing MUST be as specified in [RFC2407] with the following exceptions:

- [RFC2407] section 4.5.2: "If conflicting attributes are detected, an ATTRIBUTES-NOT-SUPPORTED Notification Payload SHOULD be returned and the security association setup MUST be aborted."

The IKE variant specified by this document MUST NOT terminate the SA setup when it encounters an unknown attribute.

- [RFC2407] section 4.5.3: "If an implementation receives a defined IPSEC DOI attribute (or attribute value) that it does not support, an ATTRIBUTES-NOT-SUPPORTED SHOULD be sent and the security association setup MUST be aborted, unless the attribute value is in the reserved range."

The IKE variant specified by this document MUST NOT terminate the SA setup when it encounters an unknown attribute.

- [RFC2407] section 4.5.3: "Notification Status Messages MUST be sent under the protection of an ISAKMP SA, either as a payload in the last main mode exchange; in a separate informational exchange after main mode or aggressive mode processing is complete; or as a payload in any quick mode exchange."

The IKE variant specified by this document SHOULD send notifications unprotected by an SA, without the hash payload, as specified in [\[RFC2409\]](#) section 5.7, if the notify occurs during the first two round trips of main mode. If the notify occurs in the last round trip of main mode, then this notify SHOULD be protected by the SA. [<13>](#)

[\[RFC2408\]](#): Message processing MUST be as specified in [RFC2408] with the following exceptions:

- [RFC2408] section 3.9: "The certificate payload MUST be accepted at any point during an exchange."

The IKE variant specified by this document MUST NOT accept certificate payloads at any time; a certificate payload MUST be in a message that contains an ID payload.

- [RFC2408] section 5.1: "When transmitting an ISAKMP message, the transmitting entity (initiator or responder) MUST do the following: 1. Set a timer and initialize a retry counter."

The IKE variant timer specified by this document does not set a retransmission timer in the following cases:

- The responder never sets a retransmission timer.
- A notify message is sent to a peer.
- A delete message is sent to a peer that does not support reliable deletes, that is, a peer that has not sent the Microsoft Implementation Vendor ID.

[\[RFC2409\]](#): Message processing MUST be as specified in [RFC2409].

[\[RFC3947\]](#): Message processing MUST be as specified in [RFC3947] with the following exceptions:

- [RFC3947] section 5.2: "In the case of transport mode, both ends MUST send both original initiator and responder addresses to the other end" and "The initiator MUST send the payloads if it

proposes any UDP-Encapsulated-Transport mode, and the responder MUST send the payload only if it selected UDP-Encapsulated-Transport mode."

The IKE variant specified by this document MUST send the NAT-OA if the host is behind a NAT.

[RFC4306]: Message processing MUST be as specified in [RFC4306] with the following exceptions:

- [RFC4306] section 2.7: "This hierarchical structure was designed to efficiently encode proposals for cryptographic suites when the number of supported suites is large because multiple values are acceptable for multiple transforms. The responder MUST choose a single suite, which MAY be any subset of the SA proposal following the rules below:"

The responder MUST consult its SPD and loop through the SPD entries, comparing each SPD entry in turn with all the proposal suites from the peer. If a match is found from the list of proposal suites, the responder MUST accept that proposal suite. This MUST repeat until a match is found, or policy comparison, and the negotiation fails.

- [RFC4306] section 3.12: "Writers of Internet-Drafts who wish to extend this protocol MUST define a Vendor ID payload to announce the ability to implement the extension in the Internet-Draft."

The IKE variant specified by this document does not define a Vendor ID to announce the implementation of CFG attributes described in section [3.11](#).

3.1.6 Timer Events

None beyond what is specified in [\[RFC2407\]](#), [\[RFC2408\]](#), [\[RFC2409\]](#), [\[RFC3947\]](#), or [\[RFC4306\]](#).

3.1.7 Other Local Events

None beyond what is specified in [\[RFC2407\]](#), [\[RFC2408\]](#), [\[RFC2409\]](#), [\[RFC3947\]](#), or [\[RFC4306\]](#).

3.2 NAT Traversal Details

Using the notation specified in [\[RFC2409\]](#) section 3.2, the generalized form of an IKE phase 1 exchange that uses NAT-T is as shown in the following figure and as specified in [\[RFC3947\]](#) section 3.2.




Figure 1: IKE phase 1 exchange using NAT-T

The description in this section uses the message numbers from the protocol sequence diagram.

The IKE NAT Traversal Protocol extension exists in two revisions. The [RFC3947] revision is specified in [RFC3947]. The [DRAFT-NATT] revision is identical to the [RFC3947] revision, except that the values used for the types defined in sections 2.2.1 and 2.2.2 are those that are specified in [DRAFT-NATT], instead of those that are specified in [RFC3947]. Both revisions include the negotiation of a choice of revision supported by both peers. <14> For more information, see [DRAFT-NATT].

3.2.1 Abstract Data Model

When this extension is implemented, the following additional state should be maintained. This is an extension to IKE Protocol version 1 as specified in [RFC2409].

Main mode security association database (MMSAD): The entry for each MM SA contains the following specific data element for NAT-T:

- Selected Revision: A flag that MUST specify what revision of the NAT-T protocol extension (as specified in [RFC3947]) has been selected for this MM SA. For more information, see [DRAFT-NATT].

3.2.2 Timers

The NAT-T keep-alive timer (per MM SA) is as specified in [RFC3948] section 4. <15>

3.2.3 Initialization

None.

3.2.4 Higher-Layer Triggered Events

3.2.4.1 Start of an IKE MM SA Negotiation

As part of the construction of message #1 for a new MM SA negotiation (as specified in [\[RFC2409\]](#) section 5), a NAT-T supporting host MUST include with its first IKE message extra vendor ID payloads (as specified in [\[RFC2408\]](#) section 3.16) to advertise its NAT-T revision support (as specified in [\[RFC3947\]](#) section 3.1). If the host supports only [\[DRAFT-NATT\]](#), it MUST include only the vendor ID "draft-ietf-ipsec-nat-t-ike-02\n" within message #1. If it supports only [\[RFC3947\]](#), it MUST include only the vendor ID "RFC 3947" within message #1. If it supports both [\[DRAFT-NATT\]](#) and [\[RFC3947\]](#), it MUST include both vendor IDs "draft-ietf-ipsec-nat-t-ike-02\n" and "RFC 3947" within message #1. [<16>](#)

3.2.5 Message Processing Events and Sequencing Rules

3.2.5.1 Receiving Message #1

On receipt of message #1, a NAT-T supporting host MUST check for the presence of the NAT-T vendor ID payloads that are specified in section [3.2.4.1](#). If NAT-T vendor ID payloads are present in the message, the host MUST set the Selected Revision for the corresponding MMSAD entry according to the following rules:

- If both hosts support [\[RFC3947\]](#) and [\[DRAFT-NATT\]](#), the host MUST set the Selected Revision to [\[RFC3947\]](#). For more information, see [\[DRAFT-NATT\]](#).
- If both hosts share only one common revision, the host MUST set the Selected Revision to the common revision.
- If the hosts do not share a common revision, the host MUST ignore the payload.

Then, the host MUST construct message #2 (as specified in [\[RFC2409\]](#) section 5) and add vendor ID payloads that advertise its NAT-T capabilities, setting the values of those payloads exactly as it would if it were constructing IKE message #1. For details, see section [3.2.4](#).

3.2.5.2 Receiving Message #2

On receipt of message #2, the host MUST check for the presence of NAT-T vendor ID payloads and set the Selected Revision as specified in section [3.2.5.1](#).

3.2.5.3 Receiving Other Messages

As specified in [\[RFC3947\]](#) section 5.2, NAT-OA payloads can be sent within the first two quick mode messages. On receipt of the first or second quick mode message, the host MUST use the Selected Revision flag of the SA's corresponding entry in the MMSAD to interpret the payload type, as defined in section [2.2.1](#).

A UDP Encapsulation type can be negotiated through the SA payload, as specified in [\[RFC3947\]](#) section 5.1. On receipt of an IKE message that may contain an SA payload, the host MUST use the Selected Revision flag of the SA's corresponding entry in the MMSAD to interpret the Encapsulation Type, as defined in section [2.2.2](#).

3.2.6 Timer Events

None.

3.2.7 Other Local Events

None.

3.3 IKE Fragmentation Details

Using the notation as specified in [\[RFC2409\]](#) section 3.2, the generalized form of an IKE phase 1 exchange that is authenticated with signatures is as shown in the following figure, as a fragmentation example. For more information, see [\[RFC2409\]](#) section 5.

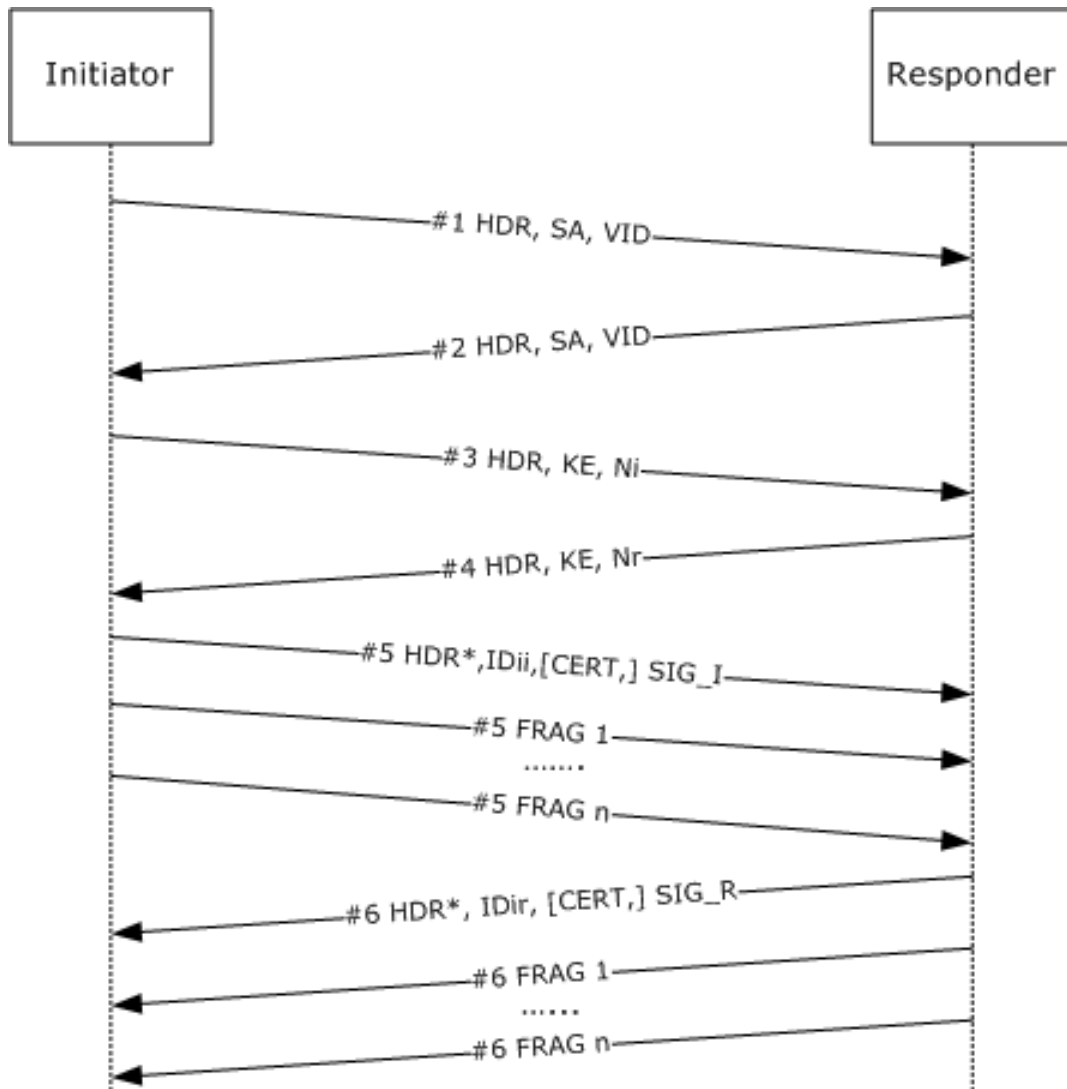


Figure 2: IKE phase 1 exchange

The description in this section uses the message numbers from the protocol sequence diagram.

3.3.1 Abstract Data Model

When this extension is implemented, the following additional state should be maintained. This is an extension to IKE Protocol version 1 as specified in [\[RFC2409\]](#).

Main mode security association database (MMSAD): The entry for each MM SA contains the following IKE fragmentation-specific data elements.

- Fragmentation supported: A flag that MUST be set if the peer supports receiving fragmented messages.
- Fragmentation active: A flag that MUST be set if the IKE messages MUST be fragmented.
- Fragmentation determination: The fragmentation need is determined by the firing of the fragmentation timer. See section [3.3.2](#) and the associated endnotes for more details. After determining that fragmentation is needed, the chosen MTU MUST be the minimum MTU for the protocol, which is 576 bytes for IPv4 and 1280 bytes for IPv6.
- Fragment queue: A queue holding the fragments that correspond to incomplete IKE messages, indexed by the Fragment ID. Each entry in the queue MUST contain:
 - The Fragment ID.
 - The Fragment Number.
 - A Flag that indicates if this fragment is the last one (that is, the Last Fragment bit is set in the Fragment payload).
 - The Fragment Data.

For definitions of the previous values, see section [2.2.3.1](#).

Flow state table: The following information MUST be maintained.

- Fragment ID counter: MUST be maintained and MUST be a 16 bit number. A Fragment ID counter SHOULD be implemented as a global counter.

3.3.2 Timers

IKE fragmentation uses the following timers:

- Fragmentation timer (for each IKE message): This timer MUST trigger fragmentation. The fragmentation timer MUST be started after sending each IKE message. The expiration of the fragmentation timer MUST indicate that the message should be fragmented the next time it is retransmitted. There MUST be one fragmentation timer per MM SA. The fragmentation timer must fire within the retransmission duration of the IKE negotiation and MUST be between 1 and 5 seconds. [<17>](#)
- Fragment reassembly timer (for each Fragment ID value): This timer MUST trigger the discarding of all the fragments received for this message. The fragment reassembly timer MUST be started when a fragment payload is received and the timer has not been started for the corresponding Fragment ID value. When the fragmentation reassembly timer fires, the delay MUST NOT exceed 90 seconds. [<18>](#)

3.3.3 Initialization

The Fragment ID counter MUST be set to zero.

3.3.4 Higher-Layer Triggered Events

3.3.4.1 Start of an IKE MM SA Negotiation

As part of the construction of message #1 for a new MM SA negotiation (as specified in [\[RFC2409\]](#) section 5), an IKE fragmentation-supporting host MUST include a "FRAGMENTATION" vendor ID payload (that is, a vendor ID payload that is generated by using the Vendor ID string "FRAGMENTATION", as specified in [\[RFC2408\]](#) section 3.16) to advertise its fragmentation capability.

3.3.5 Message Processing Events and Sequencing Rules

3.3.5.1 Receiving Message #1

On receipt of message #1, the host MUST check for the presence of a "FRAGMENTATION" vendor ID payload. If a "FRAGMENTATION" vendor ID payload is present in the message, the host MUST set the Fragmentation supported flag for the corresponding MMSAD entry.

Then, the host MUST construct message #2 (as specified in [\[RFC2409\]](#) section 5) and add the "FRAGMENTATION" vendor ID payload to advertise its fragmentation capability.

3.3.5.2 Receiving Message #2

On receipt of message #2, the host MUST check for the presence of a "FRAGMENTATION" vendor ID payload and set the Fragmentation supported flag, as specified in section [3.3.5.1](#).

3.3.5.3 Receiving Other IKE Messages

On receipt of an IKE message, the host MUST check if the message contains a Fragment payload. If a Fragment payload is present, this payload MUST be the only payload in the message. If not, the host MUST silently discard the message.

On receipt of a Fragment payload, the host MUST:

- Retrieve the Fragment ID from the Fragment payload.
- Start a fragmentation reassembly timer for this Fragment ID if no fragments are currently queued for this Fragment ID.
- If the queue for this Fragment ID already contains a fragment with the same Fragment number, the host MUST silently discard the message. If not, the host MUST queue the Fragment payload's fields in the corresponding entry of the MMSAD, indexed by the Fragment ID.

In addition, the host SHOULD set the Fragmentation active flag in the corresponding MMSAD entry. [<19>](#)

The host MUST then check whether all Fragment payloads for this Fragment ID have been received (that is, whether Fragment payloads that have a Fragment number from 1 to n have been received, and fragment n has the Last Fragment flag set).

The host MUST silently discard all Fragment payloads for this Fragment ID if any of the following error conditions occur:

- More than one Fragment payload has the Last Fragment flag set.
- A Fragment payload has been received with a Fragment number greater than the Fragment number of the fragment with the Last Fragment flag set.

If all Fragment payloads for a Fragment ID have been received, the host MUST construct the reassembled message by concatenating the following:

- The ISAKMP header from the first fragment.
- Fragment payloads (without the Fragment payload header) in the order of their Fragment number.

The host MUST then stop the fragment reassembly timer and process the reassembled IKE message as a typical message.

If the received message is a response to a previously sent message, the host MUST clear the fragmentation timer for the previously sent message.

If the processing of the IKE message results in the host sending a message, and the Fragmentation active flag is set for the corresponding MM SA, the host SHOULD fragment this message following the steps specified in section [3.3.6.1](#). If the Fragmentation active flag is not set, the host MUST start the fragmentation timer for the message it is about to send. [<20>](#)

3.3.6 Timer Events

3.3.6.1 Expiration of Fragmentation Timer

When the fragmentation timer expires, the host SHOULD start fragmenting the message that caused the timer to start. Note that the host does not need to buffer every message for fragmentation purposes because the IKE protocol has provisions for regenerating lost messages. [<21>](#)

The fragments MUST be constructed as follows:

- The Fragment ID counter is incremented.
- The IKE message is split into "n" fragments that are numbered 1 to n; the size of each fragment (after adding IP, UDP, and ISAKMP headers) is 576 bytes for IPv4 and 1,280 bytes for IPv6; however, the last fragment, which contains the remainder of the message, may be smaller.
- IKE does not adjust packet size based on router MTU advertisement; it continues to send packets for IPv4 (576 bytes) and IPv6 (1,280 bytes). Therefore, IP-level fragmentation is possible in this case.
- For each fragment, a message MUST be constructed as follows:
 - The ISAKMP header of the original IKE message has the Next Payload field set to the Fragment payload and the Encrypted flag cleared (as specified in [\[RFC2408\]](#) section 3.1).
 - The Fragment payload header has the Fragment ID set to the current value of the Fragment ID counter, the Fragment number set to the current Fragment number, and the Last Fragment flag set to Fragment number n.

The fragments MUST be sent back-to-back to the peer.

The only messages that IKE fragments are those that contain the Identification payload, as specified in [\[RFC2408\]](#) section 3.8.

3.3.6.2 Expiration of the Fragment Reassembly Timer

When the fragment reassembly timer expires, the host MUST silently discard all the fragments currently queued under the Fragment ID of the Fragment payload whose receipt caused the timer to start.

3.3.7 Other Local Events

None.

3.4 CGA Authentication Details

Using the notation as specified in [\[RFC2409\]](#) section 3.2, the generalized form of an IKE phase 1 exchange using certificates is as shown in the following figure. For more information, see [\[RFC2409\]](#) section 5.1.

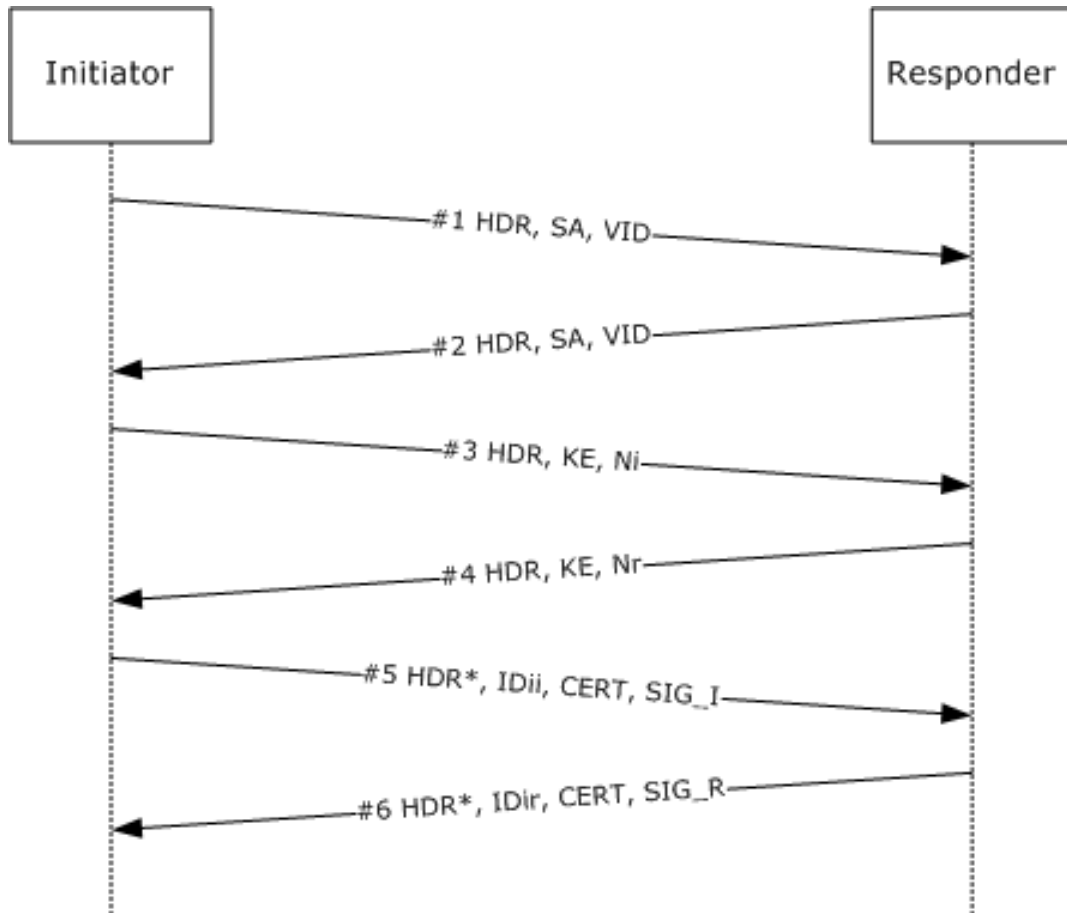


Figure 3: IKE phase 1 exchange using certificates

The CGA Authentication Protocol extension uses the same exchanges as an IKE phase 1 certificate exchange. The description in this section uses the message numbers from the protocol sequence diagram above.

The ID_IPV6_CGA identification type packet (section [2.2.5](#)) does not contain the subnet. The subnet is determined by using the following algorithm.

1. Compare the first 4 bytes of the CGA address to a well-known prefix—0x3f, 0xfe, 0x83, 0x1e—to get the prefix length. If the values match, the prefix length is equal to 88 bits; otherwise, the prefix length is 64 bits.
2. Using the prefix length, the subnet is determined by taking the leftmost number of bits equal to the prefix length from the CGA address in the packet from the peer.

3.4.1 Abstract Data Model

When this extension is implemented, the following additional state should be maintained. This is an extension to IKE Protocol version 1 as specified in [\[RFC2409\]](#).

Main mode security association database (MMSAD): The entry for each MM SA contains the following CGA authentication-specific data elements:

- CGA_CAPABLE: A flag that indicates if the authentication type 0xFDED MUST be interpreted as the AUTH_CGA authentication method.

Peer authorization database (PAD): The following information MUST be maintained:

- A new valid value AUTH_CGA that identifies the CGA authentication method, added to the locally-configurable list of acceptable authentication methods.
- A new CGA ID data structure to hold the following parameters:
 - Modifier: size: 16 octets, type: unsigned integer. See [\[RFC3972\]](#) section 3.
 - Subnet Prefix: size: 8 octets, type: IPv6 subnet. See [\[RFC3972\]](#) section 3.
 - Collision Count: size: 1 octet, type: unsigned integer. See [\[RFC3972\]](#) section 3.
 - Public Key: size: variable, type: cryptographic key. See [\[RFC3972\]](#) section 3.
- A self-signed certificate (type X.509) compatible with the IKE exchange. See [\[RFC2409\]](#) section 5.1.

This data structure is used during:

- Generation of a CGA and its associated self-signed certificate (see section [3.4.3](#)).
- Construction of an identity payload (see section [3.4.5.4](#)).
- Verification of its association with a public key (see section [3.4.5.5](#)).

3.4.2 Timers

None.

3.4.3 Initialization

Each host configured to use CGA authentication MUST generate an RSA public/private key pair (see [\[RFC3447\]](#) section 3 and [\[RFC3972\]](#) section 3). The host MUST then generate a X.509 self-signed certificate that uses this key pair and is compatible with IKE (see [\[RFC2409\]](#) section 5.1).

The CGA itself MUST be created as described in [\[RFC3972\]](#) section 4. This IP address is used to send and receive the IKE packets described in section [3.4.5](#).

3.4.4 Higher-Layer Triggered Events

3.4.4.1 Start of an IKE MM SA Negotiation

As part of the construction of message #1, a CGA authentication-supporting host MUST include an "IKE CGA version 1" vendor ID payload (that is, a vendor ID payload generated by using the vendor ID string "IKE CGA version 1", as specified in [\[RFC2408\]](#) section 3.16) to advertise its CGA authentication capability.

If the PAD requires CGA authentication, the host MUST include the AUTH_CGA Authentication method in its SA payload, as specified in section [2.2.4](#).

The host MUST use its CGA to communicate with the peer for this negotiation.

3.4.5 Message Processing Events and Sequencing Rules

3.4.5.1 Receiving Message #1

On receipt of message #1, a CGA authentication-supporting host MUST check for the presence of the "IKE CGA version 1" vendor ID payload. If an "IKE CGA version 1" vendor ID payload is present in message #1, the host MUST set the CGA_CAPABLE flag for the corresponding MMSAD entry.

The host MUST then look up its PAD to select one of the transforms that the peer proposes, as specified in [\[RFC2408\]](#) section 5.4.

If the host selects the proposed AUTH_CGA authentication method defined in section [3.4.1](#), the host MUST construct message #2, as specified in [\[RFC2409\]](#) section 5.1, and add an "IKE CGA version 1" vendor ID payload to advertise its CGA authentication capability.

The host MUST also use its CGA to communicate with the peer for this negotiation.

3.4.5.2 Receiving Message #2

On receipt of message #2, the host MUST check whether the proposal that the peer selected contains the AUTH_CGA authentication method defined in section [3.4.1](#). The host then MUST construct message #3, as specified in [\[RFC2409\]](#) section 5.1.

3.4.5.3 Receiving Message #3

Processing MUST be identical to that specified in [\[RFC2409\]](#) section 5.1.

3.4.5.4 Receiving Message #4

Processing MUST be identical to that specified in [\[RFC2409\]](#) section 5.1.

The host MUST then construct message #5, as specified in [\[RFC2409\]](#) section 5.1, with the following differences:

- The Identity payload MUST have the Identification type [ID_IPV6_CGA](#) and contain the identification data that corresponds to the host CGA (for details, see section 2.2.5). The ID_IPV6_CGA fields are read from the CGA ID (see section [3.4.1](#)).
- The CERT payload MUST contain the self-signed certificate that corresponds to the CGA.

3.4.5.5 Receiving Message #5

On receipt of message #5, the host MUST validate the message in the following ways:

- Use the SIG_I payload to verify the signature, as specified in [\[RFC2409\]](#) section 5.1. A successful verification proves that the peer has access to the private key that corresponds to the self-signed certificate passed in the CERT payload of message #5.
- Retrieve the CGA parameter structure (that is, Modifier, Collision Count, and Extension Fields) from the [ID_IPV6_CGA](#) Identity payload (for details, see section [2.2.4](#)).

- Verify that the public key contained in the self-signed certificate and the parameter structure were used to generate the peer CGA, as specified in [\[RFC3972\]](#) section 5.

If an error is encountered during payload processing, or the CGA cannot be validated, the host MUST fail the negotiation, as specified in [\[RFC2408\]](#) section 5.

Then, the host MUST construct message #6 by using the procedure for constructing message #5, as specified in section [3.4.5.4](#).

3.4.5.6 Receiving Message #6

On receipt of message #6, the host MUST validate the message using the procedure specified for validating message #5 in section [3.4.5.5](#).

3.4.6 Timer Events

None.

3.4.7 Other Local Events

None.

3.5 Fast Failover Client Details

Using the notation as specified in [\[RFC2409\]](#) section 3.2, the generalized form of an IKE phase 1 exchange is as shown in the following figure. For more information, see [\[RFC2409\]](#) section 5.

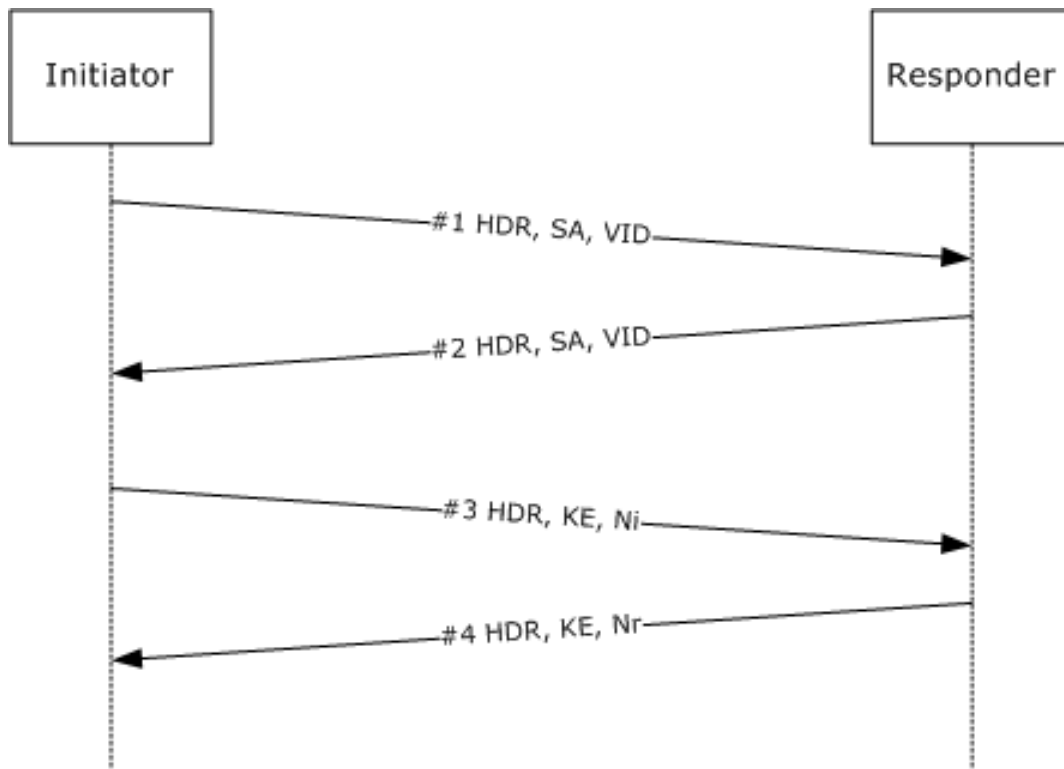


Figure 4: IKE phase 1 exchange

The description in this section uses the message numbers from the protocol sequence diagram.

3.5.1 Abstract Data Model

When this extension is implemented, the following additional state should be maintained. This is an extension to IKE Protocol version 1 as specified in [\[RFC2409\]](#).

Main mode security association database (MMSAD): The entry for each MM SA contains the following fast-failover client-specific data elements:

- Fast Failover: A flag that indicates that the "NLBS_PRESENT" vendor ID was received from the peer for this MM SA. For more details, see section [3.6.4.1](#).

3.5.2 Timers

QM SA idle timer (for each QM SA): This timer controls the inactivity time before the QM SA can be deleted (as specified in section [3.5.7.1](#)). This timer MUST be set when the QM SA has been negotiated. The QM SA idle timer is 1 minute if the peer has sent an "NLBS_PRESENT" vendor ID payload during the negotiation of the MM SA under which this QM SA was negotiated (as specified in section [3.6.4.1](#)). Otherwise, the QM SA idle timer is 5 minutes.

3.5.3 Initialization

None.

3.5.4 Higher-Layer Triggered Events

3.5.4.1 Start of an IKE MM SA Negotiation

As part of the construction of message #1 for a new MM SA negotiation (as specified in [\[RFC2409\]](#) section 5), a fast failover-supporting host MUST include a "Vid-Initial-Contact" vendor ID payload (that is, a vendor ID payload that is generated using the vendor ID string "Vid-Initial-Contact", as specified in [\[RFC2408\]](#) section 3.16) if the host does not have any active MM SAs to the peer. This is determined by looking up the MMSAD using the peer IP address.

In addition, the host MAY also add the "Vid-Initial-Contact" vendor ID payload to message #1 if it has no open TCP connections to the peer and if new connection attempts cause the retransmission of SYN packets. <22>

3.5.5 Message Processing Events and Sequencing Rules

3.5.5.1 Receiving Message #1

On receipt of message #1, a fast failover-supporting host MUST check for the presence of the "NLBS_PRESENT" vendor ID (as specified in section [3.6.4.1](#)). If the "NLBS_PRESENT" vendor ID payload is present in the message, the host MUST set the Fast Failover flag for the corresponding MMSAD entry.

If no errors are found, the host MUST construct message #2 in response. The host MUST add the "Vid-Initial-Contact" vendor ID payload to message #2 under the conditions that are specified in section [3.5.4.1](#). Otherwise, the host MUST silently ignore the packet.

3.5.5.2 Receiving Message #2

On receipt of message #2, the host MUST check for the presence of the "NLBS_PRESENT" vendor ID (for details, see section [3.6.4.1](#)). If the "NLBS_PRESENT" vendor ID payload is present in the message, the host MUST set the Fast Failover flag for the corresponding MMSAD entry.

3.5.6 Timer Events

3.5.6.1 Expiration of the QM SA Idle Timer

Upon expiration of the QM SA idle timer, the host MUST delete all states for the corresponding QM SA in the SAD.

3.5.7 Other Local Events

3.5.7.1 Successful Negotiation of a QM SA

QM SAs MUST be negotiated as specified in [\[RFC2409\]](#) section 5.5. Upon successful negotiation of a QM SA, the host MAY set the QM SA idle timer to a lower value than the default value if the Fast Failover flag is set on the corresponding MM SA. [<23>](#)

3.6 Fast Failover Server Details

The description in this section uses the message numbers from the protocol sequence diagram in section [3.5](#).

3.6.1 Abstract Data Model

This section describes a conceptual model of possible data organization that an implementation maintains to participate in this protocol. The described organization is provided to explain how the protocol behaves. This document does not mandate that implementations adhere to this model as long as their external behaviors are consistent with what is described in this document. This is an extension to IKE Protocol version 1 as specified in [\[RFC2409\]](#).

The data elements any implementation requires include the following:

- Main mode security association database (MMSAD):

For each MM SA (as specified in [\[RFC2409\]](#)), the following information MUST be maintained:

- All IKE states necessary for managing an IKE MM SA, without extensions.
- All states necessary for managing other IKE extensions for the SA, as specified in sections [3.1.1](#) and 3.6.1.
- Initial Contact: A flag indicating if the "Vid-Initial-Contact" vendor ID payload (see section [3.5.4.1](#)) has been received for the MM SA.

The MMSAD MUST be indexed by the local and peer IP addresses and the initiator and responder cookies found in the ISAKMP header (as specified in [\[RFC2408\]](#)).

Note The preceding conceptual data can be implemented by using a variety of techniques. An implementation is at liberty to implement such data in any way it pleases.

3.6.2 Timers

None.

3.6.3 Initialization

None.

3.6.4 Higher-Layer Triggered Events

3.6.4.1 Start of an IKE MM SA Negotiation

As part of the construction of message #1, a fast failover-supporting host MUST include an "NLBS_PRESENT" vendor ID payload (that is, a vendor ID payload generated by using the vendor ID string "NLBS_PRESENT", as specified in [\[RFC2408\]](#) section 3.16).

3.6.5 Message Processing Events and Sequencing Rules

3.6.5.1 Receiving Message #1

On receipt of message #1, the host MUST check for the presence of the "Vid-Initial-Contact" vendor ID (as specified in section [3.5.4.1](#)). If the "Vid-Initial-Contact" vendor ID payload is present in the message, the host MUST set the Initial Contact flag for the corresponding MMSAD entry.

If the host is part of a cluster, it MAY use this information to rebalance the MM SA to a different host within the cluster. [<24>](#)

3.6.5.2 Receiving Message #2

Message #2 has the same processing as message #1.

3.6.6 Timer Events

None.

3.6.7 Other Local Events

None.

3.7 Negotiation Discovery Details

Using the notation as specified in [\[RFC2409\]](#) section 3.2, the generalized form of an IKE phase 1 (MM) exchange is as shown in the following figure. For more information, see [\[RFC2409\]](#) section 5.




Figure 5: IKE phase 1 (MM) exchange

The description in this section uses the MM message numbers from the protocol sequence diagram.

Using the notation as specified in [RFC2409] section 3.2, the generalized form of an IKE phase 2 (quick mode) exchange is as shown in the following figure. For more information, see [RFC2409] section 5.5.

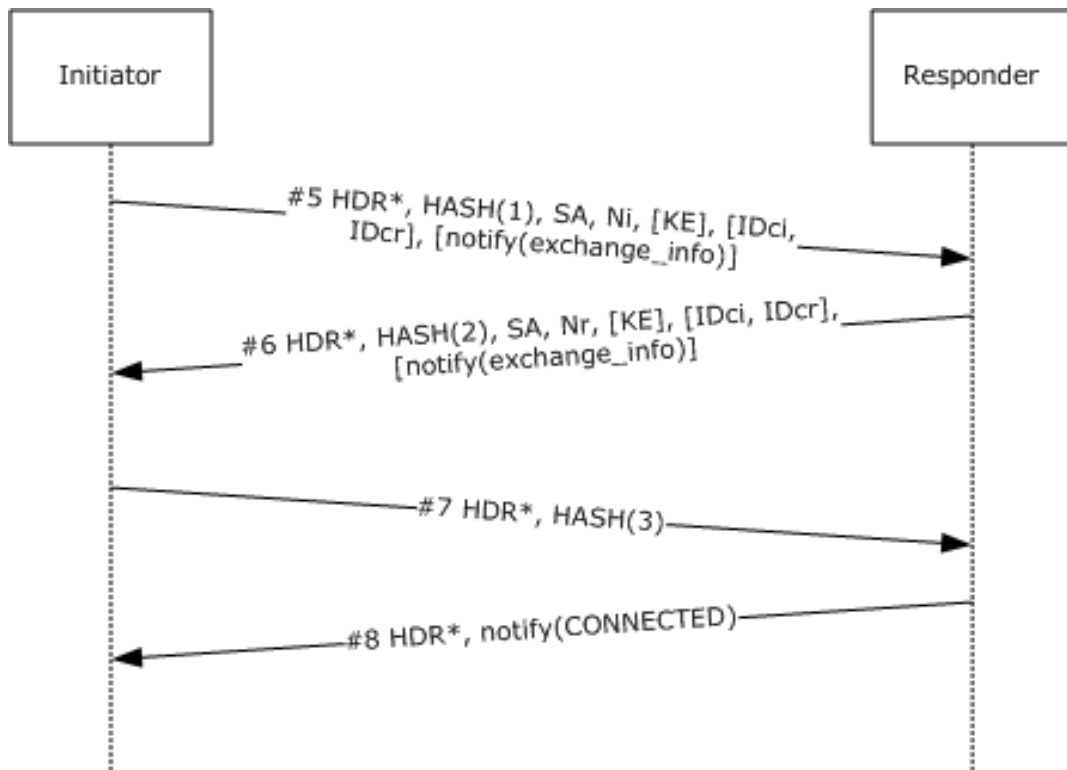


Figure 6: IKE phase 2 (QM) exchange

The description in this section uses the quick mode message numbers from the protocol sequence diagram.

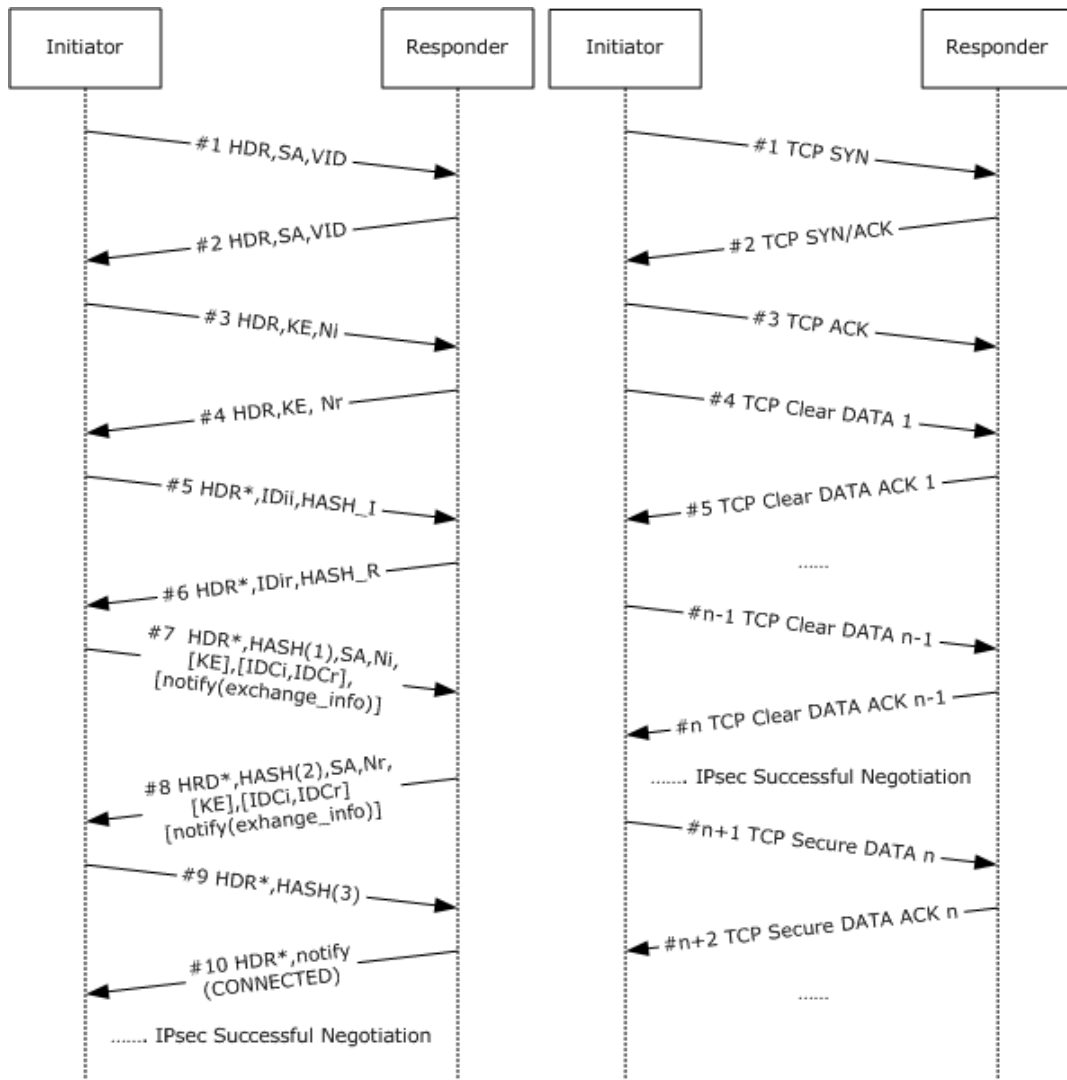


Figure 7: Negotiation discovery of a TCP connection between two IPsec-capable peers.

The TCP packet exchanges happen in parallel with the IKE exchanges that are described in figures 5 and 6. The preceding figure (figure 7) illustrates one of many ways in which the packets might interleave. When the IKE exchange completes (figure 6) the successful IPsec negotiation, the TCP connection is secured.

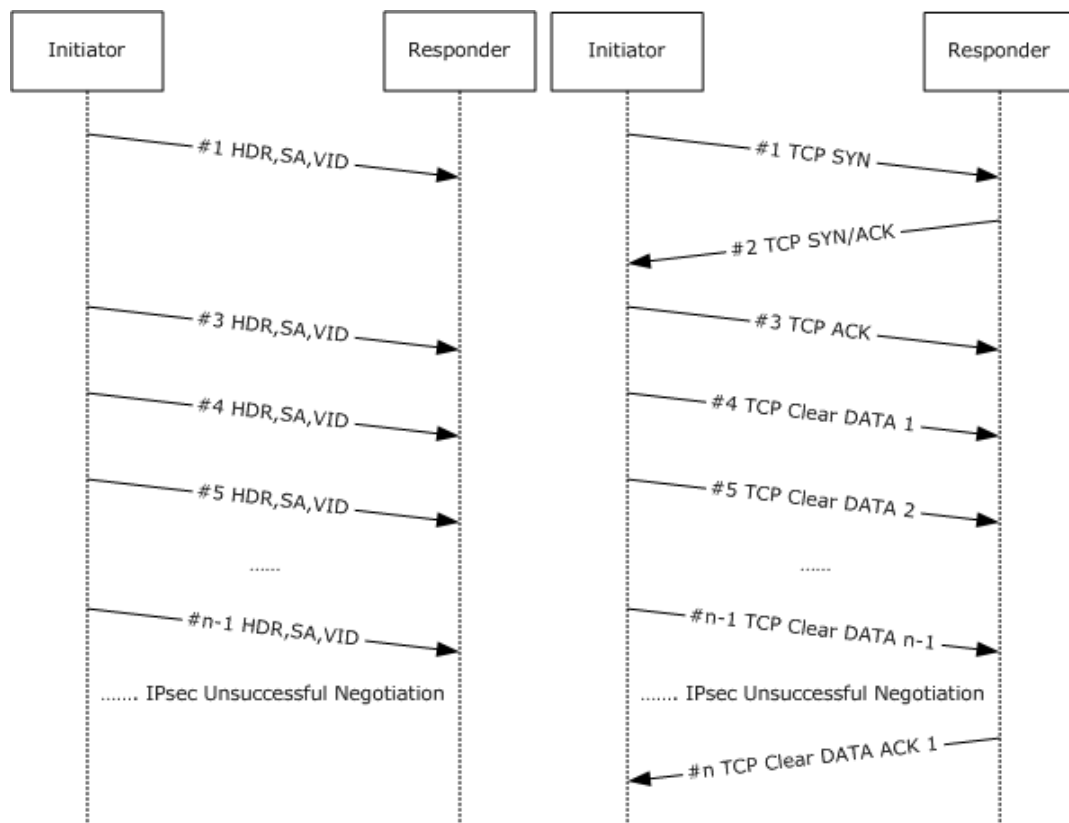


Figure 8: Negotiation discovery of a TCP connection between an IPsec-capable peer and a non-IPsec-capable peer.

The TCP packet exchanges happen in parallel with the IKE exchanges that are described in figures 5 and 6. The preceding figure (figure 8) illustrates one of many ways in which the packets might interleave. The responder does not respond to the IKE negotiation (an unsuccessful IPsec negotiation), and the TCP connection continues in the clear.

If the responder responds to the IKE negotiation, IKE fails because the responder does not have, by definition, a valid credential (it is non-IPsec-capable). However, the IKE failure does not affect the TCP stream, and the TCP connection continues in the clear.

3.7.1 Abstract Data Model

When this extension is implemented, the following additional states should be maintained. This is an extension to IKE Protocol version 1 as specified in [\[RFC2409\]](#).

Main mode security association database (MMSAD): The entry for each MM SA contains the following specific data element for negotiation discovery:

- Negotiation Discovery Supported: A flag that MUST be set if the peer supports negotiation discovery.

Security policy database (SPD): The following information MUST be maintained:

- A policy flag indicating that negotiation discovery MUST be applied to inbound and/or outbound traffic.
- A Boundary policy flag for negotiation discovery inbound rules that MUST be set if plaintext is accepted for this rule.

- A policy flag that MUST be set if encryption is guaranteed for this traffic.

Security association database (SAD): The following information MUST be maintained:

- Boundary flag: A flag that MUST be set if the QM SA matches an inbound negotiation discovery rule on the remote host.
- Guaranteed Encryption flag: A flag that MUST be set if the QM SA is an encryption SA and can be used for flows that have the Guaranteed Encryption flag set.

Flow state table: The following information MUST be maintained:

- Secure flag: A flag that MUST be set if one or more packets for this flow have been sent over a QM SA.
- Guaranteed Encryption flag: A flag that MUST be set if encryption is guaranteed for this flow.
- Acquire flag: A flag that MUST be set if a QM SA negotiation has already been triggered for this flow. This flag prevents triggering of an Acquire for each packet over a connection that stays in plaintext.

3.7.2 Timers

None.

3.7.3 Initialization

None.

3.7.4 Higher-Layer Triggered Events

3.7.4.1 Outbound Packet

An outbound packet MUST be matched against the SPD to determine if and how it needs to be protected, as specified in [\[RFC4301\]](#) section 5.

- If the packet matches a negotiation discovery rule in the SPD, and no QM SA matches the packet, one of the following MUST occur:
 - If the Secure flag is not set for the corresponding flow:

The IPsec implementation MUST send the packet and MUST trigger IKE to negotiate the corresponding QM SA if the Acquire flag is not set on the corresponding flow. Otherwise, the IPsec implementation MUST send the packet and MUST NOT trigger IKE. The first quick mode negotiation message is message #5. Message #5 MUST be constructed as follows:

- The header and payloads MUST be constructed as specified in [\[RFC2409\]](#) section 5.5.
- If the SPD rule matching the traffic has the Boundary flag set, or if the Guarantee Encryption flag is set for the flow, the host MUST include a notification payload with the following fields and values:

Notify Message Type (2 bytes): 0x9C45 (EXCHANGE_INFO).

The Notification Data field is interpreted as a flags field.

- Flag 0x00000001 (IKE_EXCHANGE_INFO_ND_BOUNDARY) MUST be set if the corresponding rule in the SPD has the Boundary flag set.

- Flag 0x00000002 (IKE_EXCHANGE_INFO_GUARANTEE_ENCRYPTION) MUST be set if the Guarantee Encryption flag is set on the corresponding flow.
- This notification payload MUST be constructed as specified in section [2.2.6](#).

The host MUST then set the Acquire flag on the corresponding flow.

- If the Secure flag is set for the corresponding flow:

The IPsec implementation MUST NOT send the packet (it MAY queue or silently discard the packet) and MUST trigger IKE to negotiate the corresponding QM SA. Message #5 MUST be constructed as previously specified.

If a QM SA needs to be negotiated, and no corresponding MM SA exists (as determined by using the outbound packet destination IP address to look up the MMSAD), an MM SA MUST be negotiated. The host MUST construct and send packet #1 as specified in [RFC2409] section 5. The host MUST include in it an "MS-Negotiation Discovery Capable" vendor ID payload (a vendor ID payload generated by using the vendor ID string "MS-Negotiation Discovery Capable", as specified in [\[RFC2408\]](#) section 3.16).

- If the packet matches a negotiation discovery rule in the SPD, and a QM SA matches the packet, the following MUST occur:

If the matching QM SA and the corresponding flow do not have the same value for the Guaranteed Encryption flag, the host MUST trigger IKE to negotiate the corresponding QM SA, as previously described in the case where there is no matching QM SA for the packet.

Otherwise, one of the following MUST occur:

- If the matching QM SA is a UDP-ESP SA ([\[RFC3947\]](#) section 5) with the Boundary flag (defined in section [3.7.1](#)) set, the host MUST send the packet in Cleartext.
- Otherwise, the IPsec implementation MUST send the packet encapsulated by using the matching QM SA, and it MUST set the Secure flag for this flow.
- If the packet does not match a negotiation discovery rule, packet processing MUST be performed as specified in [RFC4301] section 5.

If the packet matches a Guaranteed Encryption rule in the SPD, the host MUST set the Guaranteed Encryption flag on the corresponding flow. This rule MUST apply regardless of whether a matching QM SA is found or not.

3.7.4.2 Inbound Packet

An inbound packet is matched against the SPD after IPsec decapsulation to determine if and how it needs to be treated, as specified in [\[RFC4301\]](#) section 5. The following rules MUST be applied to the packet:

- If the packet is in Cleartext:
 - If the packet is the first packet for a new flow (for example, an inbound TCP SYN packet):

If the packet matches an inbound negotiation discovery rule in the SPD, the host MUST accept the packet. Otherwise, the host MUST silently discard the packet.
 - If the packet belongs to an already existing flow:

If the Secure flag is not set on the flow, the host MUST accept the packet. Otherwise, the host MUST silently discard the packet.
- If the packet was encapsulated using ESP or authentication header (AH):

The host MUST set the Secure flag on the flow and process the packet as specified in [RFC4301] section 5.

Regardless of whether the packet is in plaintext, if there is an SA that matches the packet, and its Guaranteed Encryption flag is set, the host MUST set the Guaranteed Encryption flag on the corresponding flow.

3.7.5 Message Processing Events and Sequencing Rules

3.7.5.1 Receiving Message #1

On receipt of message #1, the host MUST check for the presence of the "MS-Negotiation Discovery Capable" vendor ID payload (as specified in section [3.7.4.1](#)). If the "MS-Negotiation Discovery Capable" vendor ID payload is present in the message, the host MUST set the Negotiation Discovery Supported flag for the corresponding MMSAD entry.

Then, the host MUST construct message #2, as specified in [\[RFC2409\]](#) section 5, and add the "MS-Negotiation Discovery Capable" vendor ID payload to advertise its negotiation discovery capability.

3.7.5.2 Receiving Message #2

On receipt of message #2, the host MUST check for the presence of the "MS-Negotiation Discovery Capable" vendor ID payload (for details, see section [3.7.4.1](#)) and set the Negotiation Discovery Supported flag for the corresponding MMSAD entry.

Messages #3 and #4 MUST be constructed and processed as specified in [\[RFC2409\]](#) section 5.

3.7.5.3 Receiving Message #5

On receipt of message #5, the host MUST check for the presence of flags within a notification payload of type EXCHANGE_INFO.

- IKE_EXCHANGE_INFO_ND_BOUNDARY: If this flag is set, the host MUST set the Boundary flag for the corresponding QM SA.
- IKE_EXCHANGE_INFO_GUARANTEE_ENCRYPTION: If this flag is set, the host MUST set the Guaranteed Encryption flag for the corresponding QM SA.

Message #6 MUST be constructed in response as follows:

The IPsec implementation MUST send the packet and MUST trigger IKE to negotiate the corresponding QM SA. The first quick mode negotiation message is message #5. Message #6 MUST be constructed as follows:

- The header and payloads MUST be constructed as specified in [\[RFC2409\]](#) section 5.5.
- If the SPD rule matching the traffic for which the QM SA is negotiated has the Boundary flag set, the host MUST add a notification payload with the following fields:

Notify Message Type (2 bytes): 0x9C45 (EXCHANGE_INFO).

The Notification Data field is interpreted as a flags field.

Flag 0x00000001 (IKE_EXCHANGE_INFO_ND_BOUNDARY) MUST be set if the corresponding rule in the SPD has the Boundary flag set.

This notification payload MUST be constructed as specified in section [2.2.6](#).

3.7.5.4 Receiving Message #6

On receipt of message #6, the host MUST check for the presence of flags within a notification payload of type EXCHANGE_INFO:

- IKE_EXCHANGE_INFO_ND_BOUNDARY: If this flag is set, the host MUST set the Boundary flag for the QM SA. For more details see section [2.2.6](#).

Messages #7 and #8 are constructed and processed as specified in [\[RFC2408\]](#) section 3.1.

3.7.6 Timer Events

None.

3.7.7 Other Local Events

None.

3.8 Reliable Delete Details

Using the notation as specified in [\[RFC2408\]](#) section 4.1.1, the generalized form of an IKE Delete exchange using the Reliable Delete extension is as shown in the following figure. For more information, see [\[RFC2409\]](#) section 5.




Figure 9: IKE Delete exchange

The description in this section uses the message numbers from the protocol sequence diagram.

3.8.1 Abstract Data Model

When this extension is implemented, the following additional state should be maintained. This is an extension to IKE Protocol version 1 as specified in [\[RFC2409\]](#).

Flow state table: The following information MUST be maintained:

- Ni payload: The exact Ni payload that is sent with the delete message#1 should be preserved as part of the IKE MM SA state in order to validate the acknowledgment response. The Ni payload is a Nonce payload and MUST be constructed as specified in [\[RFC2408\]](#) section 3.13.

3.8.2 Timers

The delete retransmission timer (for each MM and QM SA): This triggers a Delete payload retransmission. The start and duration of the timer MUST be as specified in sections [3.8.4.1](#), [3.8.6.1](#), and [3.8.7.1](#).

3.8.3 Initialization

None.

3.8.4 Higher-Layer Triggered Events

3.8.4.1 SA Deletion/Invalidation

The higher layer application may cause SAs to be deleted by changing the underlying security policy, or by triggering a local state cleanup (see section [3.8.7](#)). In such cases, the host SHOULD delete the SAs, as specified in [\[RFC2408\]](#) section 5.15.

After a delete has been triggered, a delete notify MUST be sent immediately, but the MM SA MUST NOT be deleted until quick mode delete processing has been completed. Moreover, the QM SAs associated with the MM SA MUST NOT be deleted until deletion is triggered by other protocol events, as specified in [\[RFC2409\]](#) section 5.5. These protocol events are quick mode lifetime expiry as specified in [\[RFC2409\]](#) Section 5.5, policy changes (see section 3.8.7) or the peer sending a quick mode delete (See section [3.8.5](#)). Once all the QM SAs associated with the MM SA have been deleted the MM SA MUST be deleted.

The host MUST then construct message #1 as follows:

- Message #1 MUST consist only of an ISAKMP header, a Hash payload, a Nonce payload, and a Delete payload, as specified in [\[RFC2408\]](#) section 3.15. [<25>](#)
- The ISAKMP header MUST be constructed as specified in [\[RFC2409\]](#) section 5.7.
- The Hash payload MUST be constructed in the following manner:

```
HASH(1) = prf(SKEYID_a, M-ID | Ni | Delete)
```

as specified in [\[RFC2409\]](#) section 5.7.

- The Ni payload is a Nonce payload and MUST be constructed as specified in [\[RFC2408\]](#) section 3.13.
- The Delete payload MUST be constructed as specified in [\[RFC2408\]](#) section 3.15.

If the "MS NT5 ISAKMPOAKLEY" vendor ID payload (see section [1.7](#)) has been received from the peer for the corresponding MM SA, the host MUST then start the delete retransmission timer and set it to expire in 1 second. Otherwise, the host MUST NOT start the delete retransmission timer.

3.8.5 Message Processing Events and Sequencing Rules

3.8.5.1 Receiving Message #1

On receipt of message #1, the host MUST validate the message, as specified in [\[RFC2408\]](#) section 5. If message #1 is correctly validated, the host MUST delete the corresponding SA and MUST construct message #2 in response.

- The message MUST consist only of an ISAKMP header as specified in [RFC2408] section 3.1, a Hash payload as specified in [RFC2408] section 3.11, a Delete payload as specified in [RFC2408] section 3.15, and a Nonce payload structured as specified in [RFC2408] section 3.13.
- The ISAKMP header MUST be constructed as specified in as specified in [RFC2408] section 3.1. The Message ID field MUST be copied from message #1.
- The Hash payload MUST be constructed in the following manner:

$$\text{HASH}(2) = \text{prf}(\text{SKEYID}_a, \text{Ni} \mid \text{M-ID} \mid \text{Nr} \mid \text{Delete})$$

Once computed as above, this hash value MUST be sent on the wire format specified in section 3.11 of [RFC2408].

- The Ni payload is the Nonce payload without a generic payload header.
- The Delete payload MUST be copied from message #1.
- The Nr payload is a Nonce payload and MUST be constructed as specified in [RFC2408] section 3.13.

Otherwise, the host MUST silently discard message #1.

3.8.5.2 Receiving Message #2

On receipt of message #2, the host MUST validate the message as follows:

- Validate the ISAKMP header, as specified in [RFC2408] section 5.2.
- Verify that the message ID in the ISAKMP payload is identical to the message ID from message #1.

If this verification succeeds, the host MUST stop the delete retransmission timer. Otherwise, the host MUST silently discard message #2.

3.8.6 Timer Events

3.8.6.1 Expiration of the Delete Retransmission Timer

When this timer expires, the initiator MUST retransmit message #1, as specified in section 3.8.4.1, and it SHOULD reset the timer to double the previous duration unless a total of four retransmissions has already occurred. If four retransmissions have occurred, the host MUST remove the corresponding MM SA or QM SA from the MMSAD or the SAD without retransmitting message #1 or resetting the timer. <26>

When each timer expires, if a message #2 has not been received and verified for that SA, as specified in section 3.8.5.2, it SHOULD retransmit the notification message for that SA without resetting the timer.

3.8.7 Other Local Events

An administrator can trigger local SA state deletion via a local-only interface to delete all active SAs.

The abstract interface for security policy configuration changes is specified in [RFC4301] section 4.4.1. The administrator MUST be able to specify a new local security policy as defined in [RFC4301] section 4.4.1. Any MM SAs established with a policy invalidated by the new policy are deleted as specified in section 3.8.4.1.

3.8.7.1 Shutdown

IKE protocol shutdown: IKE MUST send Delete notification messages for all SAs, as specified in section [3.8.4.1](#), and then SHOULD set the delete retransmission timer to 1 second for each SA. [<27>](#)

3.8.7.2 MM SA Exhaustion

Establishment of a successful QM SA may exhaust the limits for the number of QM SAs allowed for a given MM. This quick mode limit is a local policy setting in the PAD. [<28>](#) In this case, the host MUST NOT explicitly delete the SA. Instead, the SA MUST be invalidated, and not used for establishing any new QM SAs.

3.9 Denial of Service Protection Details

IKE goes into DOS protection under the condition described in section [3.9.7](#).

Using the notation, as specified in [\[RFC2408\]](#) section 4.1.1, the generalized form of an IKE exchange using the DOS Protection extension is as shown in the following figure. For more information, see [\[RFC2409\]](#) section 5.

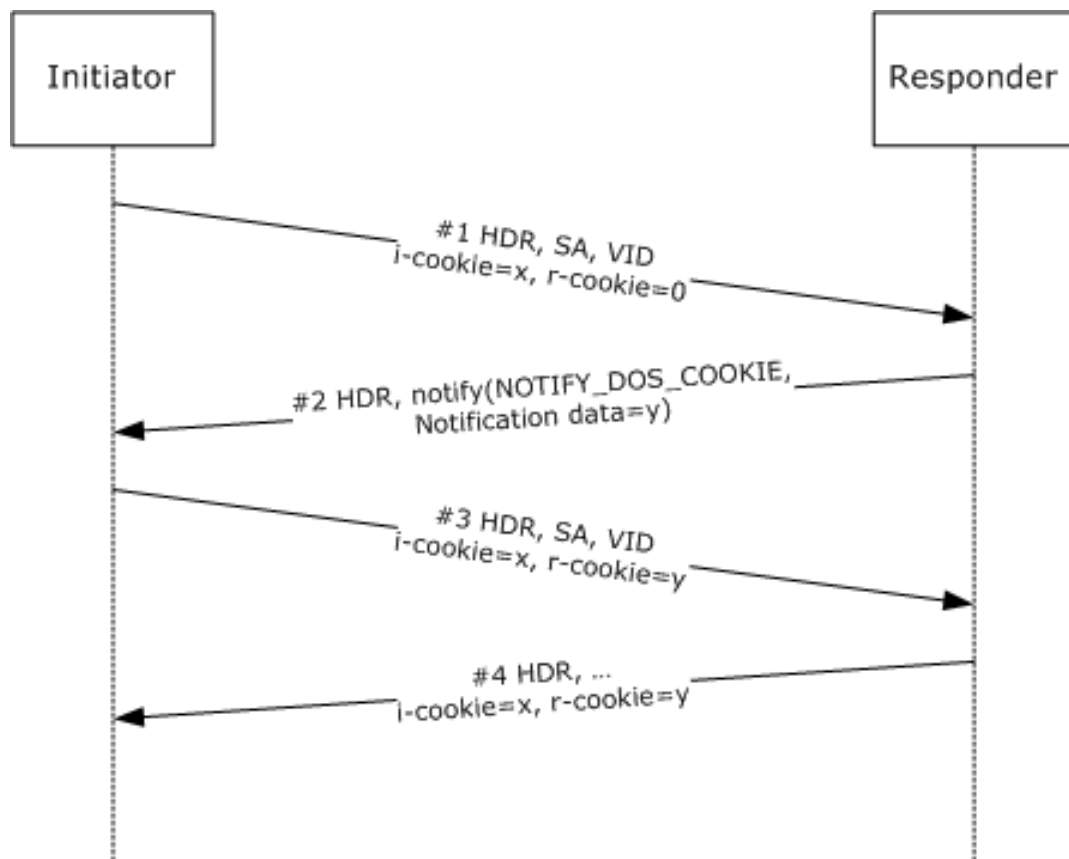


Figure 10: IKE using the DOS Protection extension

The description in this section uses the message numbers from the protocol sequence diagram.

3.9.1 Abstract Data Model

When this extension is implemented, the following additional state should be maintained. This is an extension to IKE Protocol version 1 as specified in [\[RFC2409\]](#).

Flow state table: The following information MUST be maintained:

- A flag indicating that DOS protection is active.

DOS Protection mode state: responder MUST maintain the following state to implement Denial of Service Protection mode.

- A cookie field consisting of random data.
- A cookie timeout period, initialized to 150 secs.

This state is used by the cookie generation algorithm that is described in section [3.9.5.1](#).

3.9.2 Timers

None.

3.9.3 Initialization

None.

3.9.4 Higher-Layer Triggered Events

None.

3.9.5 Message Processing Events and Sequencing Rules

3.9.5.1 Receiving Message #1

On receipt of message #1, the host MUST validate the message, as specified in [\[RFC2408\]](#) section 5. If message #1 is correctly validated, the host MUST construct message #2 in response, as follows:

- The message MUST consist of only an ISAKMP header and a Notify payload structure, as specified in [\[RFC2408\]](#) section 3.14.
- The ISAKMP header MUST be constructed as specified in [\[RFC2409\]](#) section 5.7. The message ID field is unique to this exchange, as specified in [\[RFC2409\]](#) section 5.7.
- The notify message type MUST be set to NOTIFY_DOS_COOKIE, and the notification data MUST contain an 8-byte cookie value. The cookie generation mechanism is implementation-dependent but SHOULD be stateless to provide good DOS protection. [<29>](#)

The host MUST then silently discard message #1, even if the message is correctly validated.

3.9.5.2 Receiving Message #2

On receipt of message #2, the host MUST validate the message, as specified in [\[RFC2408\]](#) section 5. In addition, the host MUST:

- Verify that the message contains a single Notify payload, that the notify message type is set to NOTIFY_DOS_COOKIE, and that the notification data contains an 8-byte cookie value. No checks on the actual value are performed at this stage.

If this verification succeeds, the host MUST construct message #3 as follows:

- Message #3 is the same as message #1, except that the **Responder Cookie** field of the ISAKMP header ([RFC2408] section 3.1) is the cookie from the notify NOTIFY_DOS_COOKIE payload in message #2.

Otherwise the host MUST process message #2 as a normal ISAKMP message.

3.9.5.3 Receiving Message #3

On receipt of message #3, the host MUST validate the message, as specified in [\[RFC2408\]](#) section 5. In addition, the host MUST:

- Verify that the **Responder Cookie** field in the ISAKMP header is not zero.
- Verify that the **Responder Cookie** field in the ISAKMP header is the same as the cookie sent in the Notify payload of message #2. The actual verification mechanism is implementation-dependent. [<30>](#)

If this verification succeeds, the host MUST process message #3 as a normal ISAKMP message. Otherwise, the host MUST process message #3 in the same way as message #1.

Subsequent messages received for this SA on the host in DOS Protection mode MUST be processed the same as message #3.

Subsequent messages received for SAs for which no state exists in the SAD MUST be processed in the same way as message #1.

3.9.6 Timer Events

None.

3.9.7 Other Local Events

DOS Protection threshold: If the number of negotiations for which only one message has been received from any initiator is above a predefined threshold, IKE MUST go into DOS Protection mode (see section [3.1](#) for details). The threshold may be implemented in a number of ways. [<31>](#)

3.10 IKE SA Correlation (IKEV2) Details

See [\[RFC4306\]](#) section 1.2. If SA Correlation is used, during the IKE_SA exchange the Correlation payload MUST be inserted immediately prior to the SA payload.

On initiator:

HDR, SK {IDi, [CERT,] [CERTREQ,] [IDr,] NOTIFY, AUTH, CORRELATION, SAI2, TSi, TSr}

This is similar to the behavior for the Extensible Authentication Protocol (EAP) exchange, as defined in [\[RFC4306\]](#) section 2.16.

NOTIFY is related to the Mobility and Multihoming Protocol (MOBIKE). See [\[RFC4555\]](#) section 4 for information about the Notify message type. See [\[RFC4306\]](#) section 3.10 for the general Notify header format.

The correlation exchange MUST use the same authentication as the original exchange. If the original exchange did EAP authentication, then the correlation exchange MUST use EAP authentication. Similarly, if the original exchange used certificate authentication (and not EAP authentication), then the correlation exchange MUST use certificate authentication, and MUST NOT use EAP authentication.

3.10.1 Abstract Data Model

When this extension is implemented, the following additional state should be maintained. This is an extension to IKE Protocol version 2 as specified in [RFC4306](#).

Main mode security association database (MMSAD): The entry for each MM SA contains the following specific data elements for IKE SA Correlation.

For IKE_SA correlation (IKEv2), the following information MUST be maintained:

- The index of the entry in the MMSAD for the other SA to which this SA has been correlated, if it exists (see section [3.10.5.1](#)).

3.10.2 Timers

None.

3.10.3 Initialization

None.

3.10.4 Higher-Layer Triggered Events

None.

3.10.5 Message Processing Events and Sequencing Rules

The following figures show the standard and EAP exchange sequences, as specified in [RFC4306](#) sections 1.2 and 2.16, respectively.

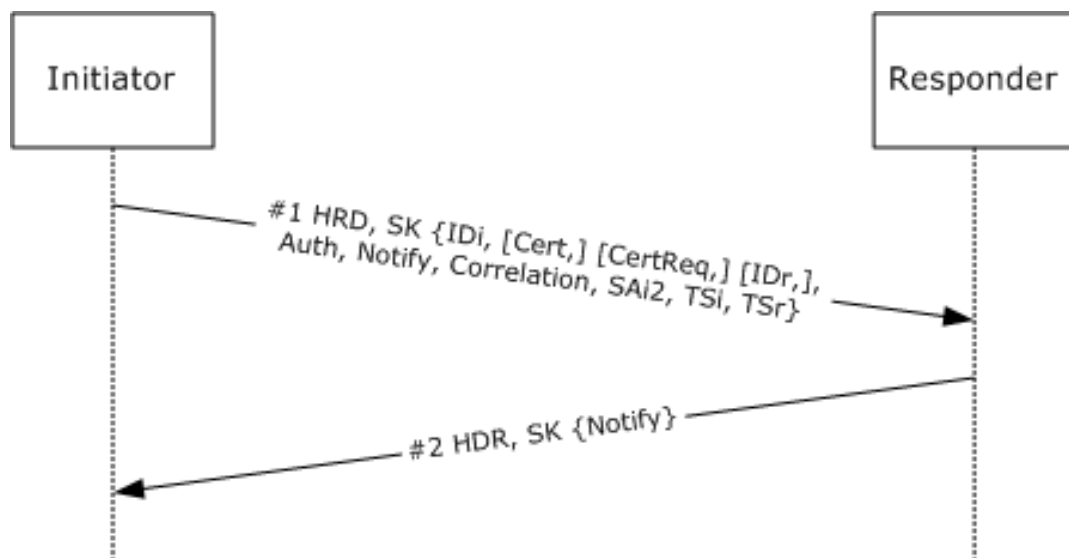


Figure 11: Standard IKEv2 exchange

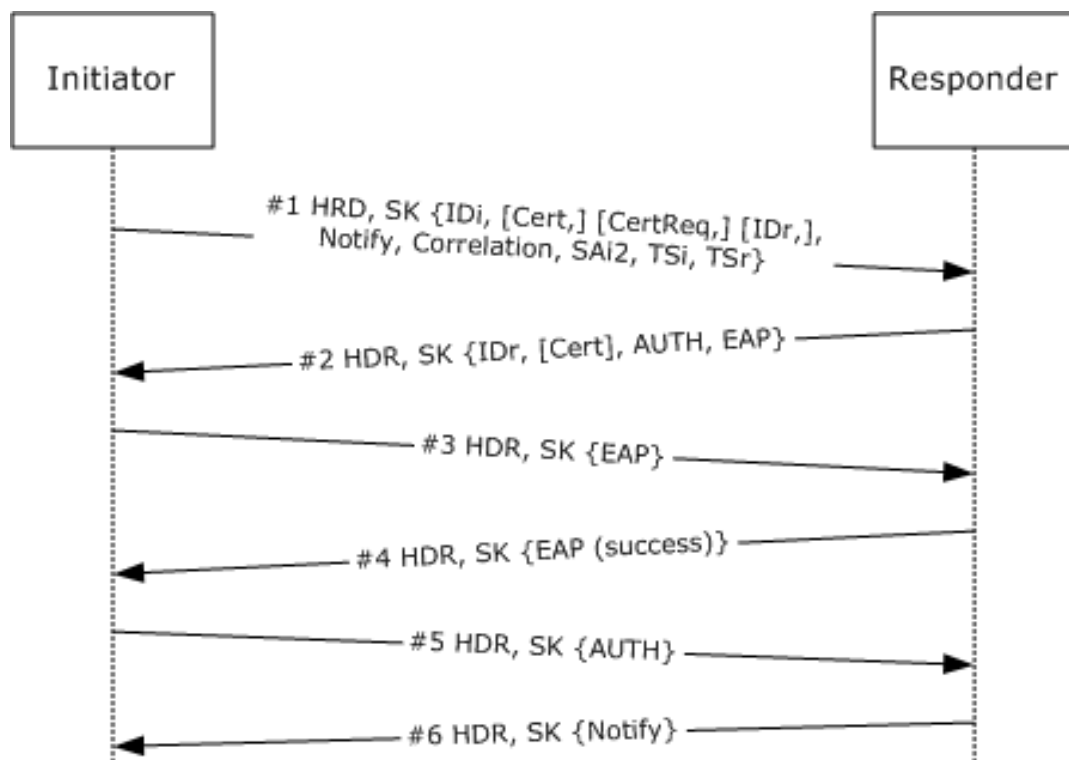


Figure 12: IKEv2 EAP exchange

3.10.5.1 Receiving Message #1

The responder processes all payloads prior to the correlation payload as per [\[RFC4306\]](#), [\[RFC4555\]](#), and [\[RFC4621\]](#). Note that message #1 corresponds to the third packet in the IKEv2 exchange. See [\[RFC4306\]](#) section 1.2.

When the host receives the correlation payload, it MUST validate its generic header as specified in [\[RFC4306\]](#) section 3.2. Additionally, the host MUST:

1. See whether an existing IKE_SA in its SADB table matches the initiator and responder SPIs from the correlation payload.
2. If there is an existing SA, the host MUST validate the correlation hash by computing its own value given its local SA state, and comparing it with the value of the correlation hash in the payload. If they are equal, the host flags these SAs as correlated.

Any failures in this exchange MUST NOT affect the state of the correlated IKE_SA.

3.10.5.2 Receiving Subsequent Messages

All subsequent messages in the exchange—except the final message—are processed as usual. At the end of the exchange, when the responder has successfully finished processing the final message, the responder tears down this exchange and sends back an IKEV2 error notify via the notification mechanism in [\[RFC4306\]](#) section 1.4.

For the standard exchange, there are no subsequent messages. For the EAP exchange, the subsequent messages 2–5 are constructed and processed identically to [\[RFC4306\]](#).

3.10.5.3 Receiving the Error Notify

The error notify MUST be processed as specified in [\[RFC4306\]](#) section 1.4 and MUST delete the SA as specified in [\[RFC4306\]](#) section 3.10.1.

The initiator, who is receiving the error notify, SHOULD process the extended error information as defined in [2.2.7](#).

3.10.6 Timer Events

None.

3.10.7 Other Local Events

None.

3.11 IKE Server Internal Addresses Configuration Attributes (IKEv2) Details

See [\[RFC4306\]](#) section 2.19. During the IKE_AUTH exchange, the IPsec remote access client (IRAC) MUST request the IPsec remote access server (IRAS)-controlled address. [<32>](#)

On initiator:

HDR, SK {IDi, [CERT,] [CERTREQ,] [IDr,] AUTH, CP(CFG_REQUEST),SAi2, TSi, TSr}

The server (IRAS) replies with:

HDR, SK {IDr, [CERT,] AUTH, CP(CFG_REPLY), SAR2, TSi, TSr}

3.11.1 Abstract Data Model

When this extension is implemented, [<33>](#) the following additional state should be maintained. This is an extension to IKE Protocol version 2 as specified in [\[RFC4306\]](#).

Flow state table: The following information MUST be maintained:

- The internal IPv4 address of the server.
- The internal IPv6 address of the server.

The initiator SHOULD request this attribute for each IP version it supports.

3.11.2 Timers

None.

3.11.3 Initialization

None.

3.11.4 Higher-Layer Triggered Events

None.

3.11.5 Message Processing Events and Sequencing Rules

The following figure shows the exchange sequence for IKEv2 Non-EAP embedded quick mode negotiation with Configuration payloads.

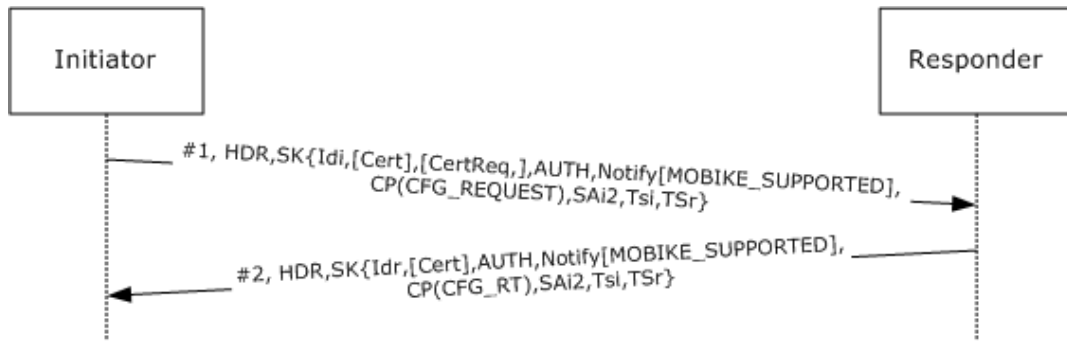


Figure 13: IKEv2 Non-EAP embedded quick mode negotiation with Configuration payload exchange

The following figure shows the Configuration payload exchange sequence with EAP, as specified in [\[RFC4306\]](#) section 3.15.




Figure 14: IKEv2 Configuration payload exchange with EAP

3.11.5.1 Receiving Message #1

When the host receives the CFG_REQUEST (as specified in [\[RFC4306\]](#) section 3.15) for the INTERNAL_IP4_SERVER or INTERNAL_IP6_SERVER attribute, it MUST validate the message as also specified in [\[RFC4306\]](#) section 3.15. Additionally, the host MAY [<34>](#):

- See whether the server has an internal IPv4 address or an internal IPv6 address.
- If either or both are present, add these attributes in CFG_REPLY.

Any failures in this exchange MUST NOT affect the state of the IKE_SA.

3.11.5.2 Receiving Message #2

When the host receives the CFG_REPLY (as specified in [\[RFC4306\]](#) section 3.15) for the INTERNAL_IP4_SERVER or INTERNAL_IP6_SERVER attribute, it MUST validate the message as also specified in [\[RFC4306\]](#) section 3. Additionally, the host MAY: [<35>](#)

- See whether the server has sent an internal IPv4 address or an internal IPv6 address.
- If either or both are present, store these values in its local data structures and use these addresses to send packets to the internal address of IRAS.

Any failures in this exchange MUST NOT affect the state of the IKE_SA.

3.11.6 Timer Events

None.

3.11.7 Other Local Events

None.

3.12 Dead Peer Detection Details

3.12.1 Abstract Data Model

When this extension is implemented [<36>](#), the following additional state should be maintained. This is an extension to IKE Protocol version 1 as specified in [\[RFC2409\]](#).

Main mode security association database (MMSAD): The entry for each MM SA contains the following fast-failover client-specific data elements:

- InboundPacketTimeStamp: 1 octet, type: unsigned integer. A time stamp field present if the SA has the NLB present flag described in section [3.5.1](#)
- A DeadPeerDetection flag: A flag that indicates whether the current SA is in dead peer detection mode.

3.12.2 Timers

QM SA idle timer (for each QM SA): This timer controls the inactivity time before the QM SA can be deleted (as specified in section [3.5.7.1](#)). This timer MUST be set when the QM SA has been negotiated as described in section [3.5.2](#).

3.12.3 Initialization

None.

3.12.4 Higher-Layer Triggered Events

3.12.4.1 TCP Dead Peer Detection

The stack sends a TCP packet and makes a lookup of the corresponding connection state in the state table defined in section [3.1.1](#). It determines whether the packet is a TCP retransmission. If it is a retransmission, the flag `DeadPeerDetection` defined in section [3.12.1](#) is set to TRUE and the dead peer detection is executed as follows:

- The host implementing this feature MUST attempt to rekey the QM SA (as described in [\[RFC2409\]](#) section 5.5) when a new connection is attempted to the peer.
- On failure of a quick mode rekey, the host implementing this extension MUST attempt to rekey MM SA (as described in [\[RFC2409\]](#) section 5.4) with a maximum of two retransmissions.
- If MM rekey fails, the peer is deemed dead and a new MM SA negotiation ([\[RFC2409\]](#) section 5.4) MAY be attempted.

3.12.4.2 UDP Dead Peer Detection

The stack sends a UDP packet and makes a lookup of the corresponding connection state in the state table defined in section [3.1.1](#). It determines whether the corresponding SA has seen a packet in the other direction by checking the `InboundPacketTimeStamp` field. If the difference is more than 20 seconds, the flag `DeadPeerDetection` defined in section [3.12.1](#) is set to TRUE and the dead peer detection is executed as follows:

- The host implementing this feature MUST attempt to rekey the QM SA (as described in [\[RFC2409\]](#) section 5.5).
- On failure of a quick mode rekey, the host implementing this extension MUST attempt to rekey MM SA (as described in [\[RFC2409\]](#) section 5.4) with a maximum of two retransmissions.
- If the MM rekey fails, the peer is deemed dead and a new MM SA negotiation ([\[RFC2409\]](#) section 5.4) MAY be attempted.

3.12.5 Message Processing Events and Sequencing Rules

3.12.5.1 Receiving a UDP Packet

The stack receives an inbound UDP packet and determines the corresponding connection state in the state table defined in section [3.1.1](#), and then it sets the `InboundPacketTimeStamp` to the current time.

3.12.6 Timer Events

3.12.6.1 Expiration of the QM SA Idle Timer

Upon expiration of the QM SA idle timer, the host MUST delete all states for the corresponding QM SA in the SAD.

3.12.7 Other Local Events

3.12.7.1 Successful Negotiation of a QM SA and MM SA

QM SAs MUST be negotiated as specified in [\[RFC2409\]](#) section 5.5. Upon successful negotiation of a QM SA, the host MUST set the `DeadPeerDetection` to FALSE, and the host MAY set the QM SA idle

timer to a lower value than the default value if the Fast Failover flag is set on the corresponding MM SA.<37>

MM SAs MUST be negotiated as specified in [RFC2409] section 5.4. Upon successful negotiation of a MM SA, the host MUST set the DeadPeerDetection to FALSE.

3.13 Xbox Multiplayer Gaming (IKEv2) Vendor IDs Details

3.13.1 Abstract Data Model

When this extension is implemented<38>, the following additional state should be maintained. This is an extension to IKE Protocol version 2 as specified in [RFC4306].

Main mode security association database (MMSAD): The entry for each MM SA contains the following Xbox multiplayer gaming-specific data element:

- Xbox IKEv2 Negotiation Type: 4 octets, type: unsigned integer. An integer representing the type of Xbox multiplayer identifier associated with the "Xbox IKEv2 Negotiation" vendor ID payload.<39>

3.13.2 Timers

None.

3.13.3 Initialization

For Xbox multiplayer gaming, secure connections can be of various types. This type information is stored in the XBox IKEv2 Negotiation Type ADM element discussed in section 3.13.1. The significance of the different types of secure connections for Xbox multiplayer gaming is out of scope for this document. However, a limit can be imposed on the number of simultaneous IKE negotiations that are available for each type of Xbox multiplayer gaming secure connection. Absence of such a configuration would mean that there is no limit to the number of simultaneous ongoing negotiations.

3.13.4 Higher-Layer Triggered Events

None.

3.13.5 Message Processing Events and Sequencing Rules




Figure 15: IKE_SA_INIT message exchange for Xbox multiplayer gaming secure-connection establishment

IKE initiators that are participating in Xbox multiplayer gaming scenarios and establishing a secure connection with a remote peer can send the "Microsoft Xbox One 2013" vendor ID and the "Xbox IKEv2 Negotiation" vendor ID payloads in the IKE_SA_INIT message.

3.13.5.1 Microsoft Xbox One 2013 Vendor ID

The "Microsoft Xbox One 2013" vendor ID simply indicates that the IKEv2 message exchange is for negotiation of an IKE SA for Xbox multiplayer gaming secure connections.

3.13.5.2 Xbox IKEv2 Negotiation Vendor ID

The "Xbox IKEv2 Negotiation" vendor ID can be looked up by the responder and stored in the Xbox IKEv2 Negotiation Type ADM element discussed in section [3.13.1](#). For the associated negotiation type, the host MUST increment the number of ongoing IKE negotiations. If the number of such IKE negotiations exceeds the configured limit for the given Xbox secure connection, the negotiation is failed.

3.13.6 Timer Events

If an IKE SA is associated with an Xbox negotiation type, then IKE_SA_INIT messages for those SAs are not retransmitted if no response is received from the peer after the first timeout period ([\[RFC5996\]](#) section 2.1).

3.13.7 Other Local Events

None.

3.14 Security Realm ID (IKEv2) Vendor IDs Details

3.14.1 Abstract Data Model

When this extension is implemented, [<40>](#) the following additional state should be maintained. This is an extension to IKE Protocol version 2 as specified in [\[RFC5996\]](#).

Security policy database (SPD): The following information MUST be maintained for a security realm IPsec policy:

- **Security Realm ID:** A variable length array of bytes stored as an HMAC-MD5 hash of the string that identifies the security realm IPsec policy. For more information, see section [1.3.12.<41>](#)

3.14.2 Timers

None.

3.14.3 Initialization

None.

3.14.4 Higher-Layer Triggered Events

None.

3.14.5 Message Processing Events and Sequencing Rules




Figure 16: Sending Security Realm ID Vendor ID in IKE_SA_INIT and IKE_SA_AUTH messages

IKE initiators can send the Security Realm ID vendor ID in the IKE_SA_INIT and IKE_SA_AUTH messages if the policy used to negotiate the IKE and IPsec SAs are security realm-based IPsec policies.

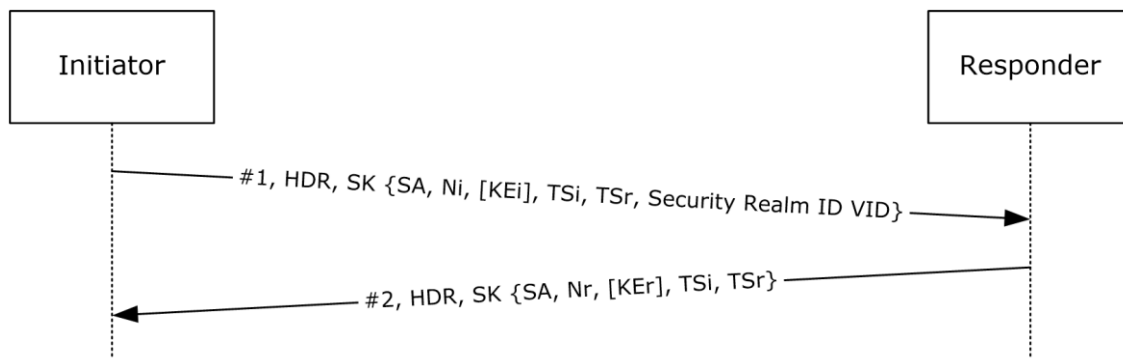


Figure 17: Sending Security Realm ID Vendor ID payload in CREATE_CHILD_SA messages

The security realm vendor ID payload (section [2.2.10](#)) is sent in CREATE_CHILD_SA messages if the parent SA is associated with a security realm-based policy.

3.14.5.1 IKE_SA_INIT Messages

Initiator: If the initiator chooses a security realm-based IPsec policy to trigger an SA negotiation, it reads the **Security Realm ID** ADM element defined in section [3.14.1](#), and includes it in the "MSFT IPsec Security Realm Id" vendor ID payload in the IKE_SA_INIT message.

Responder: If the responder receives an IKE_SA_INIT message that contains a "MSFT IPsec Security Realm Id" vendor ID, it reads the last 16 bytes of the payload, and uses that data to look up a matching IPsec policy. Note that there might be implicit priorities associated with IPsec policies. A higher priority IPsec policy that is not associated with any security realm can be selected over a lower priority IPsec policy that might be associated with the said security realm ID. However, if a security realm-based IPsec policy were to be chosen, then the security realm ID associated to the policy MUST exactly match the security realm ID as received in the vendor ID. On the other hand, if the IKE_SA_INIT message does not have a "MSFT IPsec Security Realm Id" vendor ID, then the responder MAY [42](#) skip any security realm-based IPsec policies while selecting an IKE policy.

For additional product-specific behavior, see the following citation: [43](#)

3.14.5.2 IKE_SA_AUTH and CREATE_CHILD_SA Messages

Initiator: If the initiator chooses a security realm-based IPsec policy to trigger an SA negotiation, it takes the security realm ID in the policy and includes it in the "MSFT IPsec Security Realm Id" vendor ID payload in an IKE_SA_AUTH message (for embedded child IPsec SA negotiation) or a CREATE_CHILD_SA message (for standalone child IPsec SA negotiation). Note that in these messages the vendor ID payload is part of the encrypted payloads. Also, the initiator MUST select the same security realm ID in both the IKE_SA_INIT message and the IKE_SA_AUTH/CREATE_CHILD_SA messages.

Responder: If the responder receives a "MSFT IPsec Security Realm Id" vendor ID in the IKE_SA_AUTH or CREATE_CHILD_SA messages, it looks up an IPsec (QM) policy in the same way as for IKE (MM) policy. However, if a security realm ID-based IPsec policy is chosen, the responder MUST ensure that the corresponding IKE (MM) policy is associated with the same security realm ID. If the message from the initiator for negotiating the child SA does not have a "MSFT IPsec Security Realm Id" vendor ID, but the parent IKE SA is associated to a security realm policy, then this message will be discarded by the responder and the child SA negotiation will fail.

Note that for rekeying IKE and child IPsec SAs, CREATE_CHILD_SA messages are used, and the security realm vendor ID is used in a manner that is similar to that in the preceding paragraphs.

3.14.6 Timer Events

None.

3.14.7 Other Local Events

None.

4 Protocol Examples

4.1 Negotiation Discovery Examples

The following protocol sequence diagram depicts communication between a client with a negotiation discovery policy and a server with negotiation discovery in boundary mode.




Figure 18: Negotiation discovery between client and server

In this example, the client initiates a TCP connection to the server. At the same time that it sends the TCP SYN packet, the client initiates the IKE to the server. TCP traffic flows in the clear until the IKE negotiation completes with IKE message #6. Then, the traffic for this connection is protected.

In the second example, the server requires all inbound traffic to be protected.

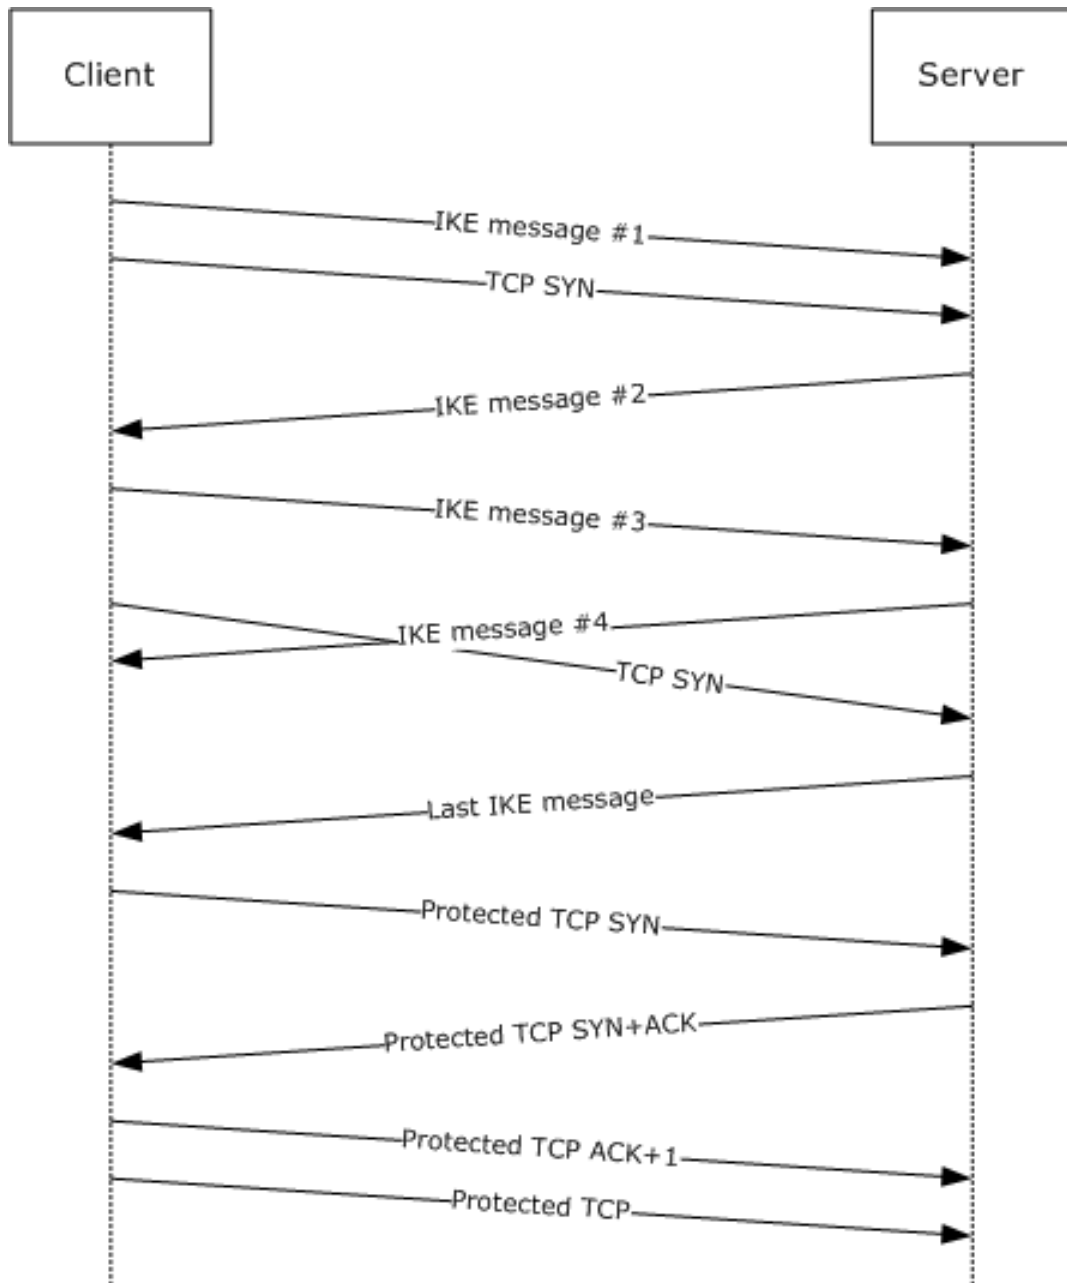


Figure 19: Negotiation discovery between client and server, all inbound traffic protected

In this example, the client initiates a TCP connection to the server. At the same time that it sends the TCP SYN packet, the client initiates the IKE to the server. The Cleartext TCP SYN packets are dropped by the server and retransmitted by the client until the IKE negotiation completes with IKE message #6. The server then accepts the protected traffic.

5 Security

5.1 Security Considerations for Implementers

5.1.1 Negotiation Discovery

Negotiation discovery allows Cleartext outbound and inbound connections if the peer is not IPsec-capable. Connections that are Cleartext should be considered when designing the policy.

5.2 Index of Security Parameters

Security parameter	Section
Authentication method	1.7
Encryption/authentication algorithms	1.7
Diffie-Hellman	1.7

6 Appendix A: Product Behavior

The information in this specification is applicable to the following Microsoft products or supplemental software. References to product versions include released service packs.

Note: Some of the information in this section is subject to change because it applies to an unreleased, preliminary version of the Windows Server operating system, and thus may differ from the final version of the server software when released. All behavior notes that pertain to the unreleased, preliminary version of the Windows Server operating system contain specific references to Windows Server 2016 Technical Preview as an aid to the reader.

- Windows 2000 operating system
- Windows XP operating system
- Windows Server 2003 operating system
- Windows Vista operating system
- Windows Server 2008 operating system
- Windows 7 operating system
- Windows Server 2008 R2 operating system
- Windows 8 operating system
- Windows Server 2012 operating system
- Windows 8.1 operating system
- Windows Server 2012 R2 operating system
- Windows 10 operating system
- Windows Server 2016 Technical Preview operating system

Exceptions, if any, are noted below. If a service pack or Quick Fix Engineering (QFE) number appears with the product version, behavior changed in that service pack or QFE. The new behavior also applies to subsequent service packs of the product unless otherwise specified. If a product edition appears with the product version, behavior is different in that product edition.

Unless otherwise specified, any statement of optional behavior in this specification that is prescribed using the terms SHOULD or SHOULD NOT implies product behavior in accordance with the SHOULD or SHOULD NOT prescription. Unless otherwise specified, the term MAY implies that the product does not follow the prescription.

[<1> Section 1.3: IKEv2 Protocol Implementation Notes](#)

[\[RFC4306\]](#) IKEv2 MUST / MUST NOT implementation notes

RFC Requirement	RFC 4306 Section	Compliance statement
"If a node receives a delete request for SAs for which it has already issued a delete request, it MUST delete the outgoing SAs while processing the request and the incoming SAs while processing the response."	Section 1.4	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016

RFC Requirement	RFC 4306 Section	Compliance statement
		Technical Preview.
"Note that Message IDs are cryptographically protected and provide protection against message replays. In the unlikely event that Message IDs grow too large to fit in 32 bits, the IKE_SA MUST be closed."	Section 2.2	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"The management interface by which the Shared Secret is provided MUST accept ASCII strings of at least 64 octets and MUST NOT add a null terminator before using them as shared secrets. It MUST also accept a HEX encoding of the Shared Secret."	Section 2.15	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"IKEv2 simplifies this situation by requiring that ECN be usable in the outer IP headers of all tunnel-mode IPsec SAs created by IKEv2. Specifically, tunnel encapsulators and decapsulators for all tunnel-mode SAs created by IKEv2 MUST support the ECN full-functionality option for tunnels specified in [RFC3168] and MUST implement the tunnel encapsulation and decapsulation processing specified in [RFC4301] to prevent discarding of ECN congestion indications."	Section 3.6	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"MUST be capable of being configured to send and accept the first two Hash and URL formats (with HTTP URLs)."	Section 3.6	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.

[RFC4306] IKEv2 SHOULD / SHOULD NOT implementation notes

RFC Requirement	RFC 4306 Section	Compliance statement
"The initiator SHOULD repeat the request, but now with a KEi payload from the group the responder selected."	Section 1.3	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"...regard half-closed connections as anomalous and audit their existence should they persist."	Section 1.4	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"IKEv2 implementations SHOULD be aware of the maximum UDP message size supported."	Section 2	Not supported in Windows 7, Windows Server 2008 R2,

RFC Requirement	RFC 4306 Section	Compliance statement
		Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"An IKE endpoint supporting a window size greater than one SHOULD be capable of processing incoming requests out of order to maximize performance in the event of network failures or packet reordering."	Section 2.4	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"An endpoint SHOULD suspect that the other endpoint has failed based on routing information and initiate a request to see whether the other endpoint is alive."	Section 2.4	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"...implementations SHOULD reject as invalid a message with those payloads in any other order."	Section 2.5	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"Implementations SHOULD support in-place rekeying of SAs."	Section 2.8	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"If redundant SAs are created though such a collision, the SA created with the lowest of the four nonces used in the two exchanges SHOULD be closed by the endpoint that created it."	Section 2.8	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"If an initiator receives a message on an SA for which it has not received a response to its CREATE_CHILD_SA request, it SHOULD interpret that as a likely packet loss and retransmit the CREATE_CHILD_SA request."	Section 2.8	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"If an error occurs outside the context of an IKE request (e.g., the node is getting ESP messages on a nonexistent SPI), the node SHOULD initiate an INFORMATIONAL exchange with a Notify payload describing the problem."	Section 2.21	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.

RFC Requirement	RFC 4306 Section	Compliance statement
"A node SHOULD treat such a message (and also a network message like ICMP destination unreachable) as a hint that there might be problems with SAs to that IP address and SHOULD initiate a liveness test for any such IKE_SA. An implementation SHOULD limit the frequency of such tests."	Section 2.21	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"There are cases where a NAT box decides to remove mappings that are still alive (for example, the keepalive interval is too long, or the NAT box is rebooted). To recover in these cases, hosts that are not behind a NAT SHOULD send all packets (including retransmission packets) to the IP address and port from the last valid authenticated packet from the other end."	Section 2.23	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"To assure maximum interoperability, implementations MUST be configurable to send at least one of ID_IPV4_ADDR, ID_FQDN, ID_RFC822_ADDR, or ID_KEY_ID, and MUST be configurable to accept all of these types. Implementations SHOULD be capable of generating and accepting all of these types."	Section 3.5	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"Implementations SHOULD be capable of being configured to send and accept Raw RSA keys."	Section 3.6	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"Note that since IKE passes an indication of initiator identity in message 3 of the protocol, the responder SHOULD NOT send EAP Identity requests. The initiator SHOULD, however, respond to such requests if it receives them."	Section 3.16	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.

[RFC4306] IKEv2 MAY implementation notes

RFC Requirement	RFC 4306 Section	Compliance statement
"The traffic selectors for traffic to be sent on that SA are specified in the TS payloads, which may be a subset of what the initiator of the CHILD_SA proposed."	Section 1.3	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"If the receiving node has an active IKE_SA to the IP address from whence the packet came, it MAY send a notification of the wayward packet over that IKE_SA in an INFORMATIONAL exchange."	Section 1.5	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2,

RFC Requirement	RFC 4306 Section	Compliance statement
		Windows 10, and Windows Server 2016 Technical Preview.
"IKEv2 implementations SHOULD be aware of the maximum UDP message size supported and MAY shorten messages by leaving out some certificates or cryptographic suite proposals if that will keep messages below the maximum."	Section 2	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"In order to maximize IKE throughput, an IKE endpoint MAY issue multiple requests before getting a response to any of them if the other endpoint has indicated its ability to handle such requests. For simplicity, an IKE implementation MAY choose to process requests strictly in order and/or wait for a response to one request before issuing another."	Section 2.3	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"To prevent this, the initiator MAY be willing to accept multiple responses to its first message, treat each as potentially legitimate, respond to it, and then discard all the invalid half-open connections when it receives a valid cryptographically protected response to any one of its requests."	Section 2.4	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"The responder in that case MAY reject the message by sending another response with a new cookie or it MAY keep the old value of <secret> around for a short time and accept cookies computed from either one."	Section 2.6	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"An implementation MAY refuse all CREATE_CHILD_SA requests within an IKE_SA."	Section 2.8	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"An initiator MAY send a dummy message on a newly created SA if it has no messages queued in order to assure the responder that the initiator is ready to receive messages."	Section 2.8	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"If more than one subset is acceptable but their union is not, the responder MUST accept some subset and MAY include a Notify payload of type ADDITIONAL_TS_POSSIBLE to indicate that the	Section 2.9	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server

RFC Requirement	RFC 4306 Section	Compliance statement
initiator might want to try again."		2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"CP(CFG_REQUEST) MUST contain at least an INTERNAL_ADDRESS attribute (either IPv4 or IPv6) but MAY contain any number of additional attributes the initiator wants returned in the response."	Section 2.19	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"An IKE peer wishing to inquire about the other peer's IKE software version information MAY use the method below."	Section 2.20	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"A node receiving a suspicious message from an IP address with which it has an IKE_SA MAY send an IKE Notify payload in an IKE INFORMATIONAL exchange over that SA."	Section 2.21	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"A node requesting a CHILD_SA MAY advertise its support for one or more compression algorithms through one or more Notify payloads of type IPCOMP_SUPPORTED. The response MAY indicate acceptance of a single compression algorithm with a Notify payload of type IPCOMP_SUPPORTED."	Section 2.22	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"IPv6-only implementations MAY be configurable to send only ID_IPV6_ADDR."	Section 3.5	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"INVALID_SPI MAY be sent in an IKE INFORMATIONAL exchange when a node receives an ESP or AH packet with an invalid SPI."	Section 3.10.11	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"INVALID_SELECTORS MAY be sent in an IKE INFORMATIONAL	Section	Not supported in Windows 7,

RFC Requirement	RFC 4306 Section	Compliance statement
exchange when a node receives an ESP or AH packet whose selectors do not match those of the SA on which it was delivered (and that caused the packet to be dropped)."	3.10.11	Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"INITIAL_CONTACT: It MAY be sent when an IKE_SA is established after a crash,..."	Section 3.10.11	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"NAT_DETECTION_SOURCE_IP: There MAY be multiple Notify payloads of this type in a message if the sender does not know which of several network attachments will be used to send the packet."	Section 3.10.11	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"NAT_DETECTION_DESTINATION_IP: Alternately, it MAY reject the connection attempt if NAT traversal is not supported."	Section 3.10.11	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"HTTP_CERT_LOOKUP_SUPPORTED: This notification MAY be included in any message that can include a CERTREQ payload and indicates that the sender is capable of looking up certificates based on an HTTP-based URL."	Section 3.10.11	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"The CFG_REPLY Configuration Payload MAY return that value, or a new one. It MAY also add new attributes and not include some requested ones. Requestors MUST ignore returned attributes that they do not recognize. Some attributes MAY be multi-valued, in which case multiple attribute values of the same type are sent and/or returned."	Section 3.15	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.
"INTERNAL_IP4_ADDRESS, INTERNAL_IP6_ADDRESS: With IPv6, a requestor MAY supply the low-order address bytes it wants to use. Multiple internal addresses MAY be requested by requesting multiple internal address attributes."	Section 3.15.1	Not supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.

<2> [Section 1.3](#): IKE extensions by operating system cross-reference. Implemented in Windows NT 4.0 operating system, Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The following tables describe the extensions that each release supports.

The IKE proposal for Encapsulating Security Payload (ESP) and Authentication Headers (AH) is deprecated in the Windows 7 implementation of IKE version 2 but supported by IKE version 1 in Windows NT 4.0, Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.

IKE extension	Windows NT 4.0 (with additional download)	Windows 2000	Windows 2000 SP4 post-SP4 rollup
NAT-T	X		X
IKE fragmentation			X
CGA authentication			
Fast failover			
Negotiation discovery			
Reliable delete			X

IKE extension	Windows XP	Windows XP SP2	Windows Server 2003	Windows Vista and Windows Server 2008	Windows 7 and Windows Server 2008 R2	Windows 8, Windows Server 2012, Windows 8.1, and Windows Server 2012 R2	Windows 10, and Windows Server 2016 Technical Preview
NAT-T		X	X	X	X	X	X
IKE fragmentation		X	X	X	X	X	X
CGA authentication				X	X	X	X
Fast failover		X	X	X	X	X	X
Negotiation discovery				X	X	X	X
Reliable delete	X	X	X	X	X	X	X
IKEv2 SA Correlation					X	X	X
IKEv2 Configuration					X	X	X

IKE extension	Windows XP	Windows XP SP2	Windows Server 2003	Windows Vista and Windows Server 2008	Windows 7 and Windows Server 2008 R2	Windows 8, Windows Server 2012, Windows 8.1, and Windows Server 2012 R2	Windows 10, and Windows Server 2016 Technical Preview
n Attributes							
Denial of Service protection	X	X	X	X	X	X	X
Dead Peer Detection						X *	X
Xbox Multiplayer Gaming (IKEv2) Vendor IDs							X
Security Realm (IKEv2) Vendor IDs							X

* Dead Peer Detection is implemented only for IKEv2-based site-to-site IPsec tunnels on Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016 Technical Preview. Dead Peer Detection is not implemented on Windows 8, Windows 8.1, or Windows 10.

<3> [Section 1.3.8](#): IKE extensions by operating system cross reference. Implemented in Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The following table describes the extensions that are supported by IKE in each release.

IKE extension	Windows XP SP2	Windows Server 2003	Windows Vista and Windows Server 2008	Windows 7 and Windows Server 2008 R2	Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
IKE/AuthIP co-existence			X	X	X
Exchange information notification payload	X	X	X	X	X

<4> [Section 1.7](#): Algorithms implemented by operating system cross reference. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The following table lists the algorithms that are implemented.

Authentication method	Windows 2000	Windows XP	Windows Server 2003	Windows Vista and Windows Server 2008	Windows 7 and Windows Server 2008 R2	Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
Pre-shared key (as specified in [RFC2409])	X	X	X	X	X	X
RSA signature (as specified in [RFC2409])	X	X	X	X	X	X
Kerberos using GSS-API (as specified in [GSS])	X	X	X	X	X	X
CGA (as specified in [RFC3972])				X	X	X

<5> [Section 1.7](#): Cryptographic parameters implemented by operating system cross reference. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The following tables list the parameters that are implemented.

Diffie-Hellman group	Windows 2000	Windows XP	Windows Server 2003	Windows Vista and Windows Server 2008	Windows 7 and Windows Server 2008 R2	Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
Default 768-bit MODP group (as specified in [RFC2409])	X	X	X	X	X	X
Alternate 1,024-bit MODP group (as specified in [RFC2409])	X	X	X	X	X	X
2,048-bit MODP group (as specified in [RFC3526])			X	X	X	X
ECP256 (as specified in [ECP])				X	X	X
ECP384 (as specified in [ECP])				X	X	X

Authentication algorithm	Windows 2000	Windows XP	Windows Server 2003	Windows Vista and Windows Server 2008	Windows 7 and Windows Server 2008 R2	Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
NULL (as specified in [RFC2410])	X	X	X	X	X	X
HMAC-SHA1-96 (for more information, see [RFC2404])	X	X	X	X	X	X
HMAC-MD5-96 (as specified in [RFC2403])	X	X	X	X	X	X
AES-MAC (for more information, see [RFC4543])					X	X
SHA-256 (for more information, see [SHA256])					X	X

Encryption algorithm	Windows 2000	Windows XP	Windows Server 2003	Windows Vista and Windows Server 2008	Windows 7 and Windows Server 2008 R2	Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
NULL (as specified in [RFC2410])	X	X	X	X	X	X
DES-CBC (for more information, see [RFC2405])	X	X	X	X	X	X
3DES-CBC (as specified in [RFC2451])	X	X	X	X	X	X
AES-CBC with 128, 192, and 256 Bit Keys (for more information, see [RFC3602])				X	X	X
AES-GCM with 128, 192, and 256 Bit Keys (for more information, see [RFC5288])					X	X

see [RFC4106])						
---------------------------------	--	--	--	--	--	--

<6> [Section 1.7](#): Vendor ID payloads. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The Microsoft implementation of IKE supports the following vendor IDs.

The Microsoft implementation vendor ID (the first rows of the second table that follows, where the common name starts with "Microsoft implementation") is constructed by appending a 32-bit (4-byte) version number in network order to the 128-bit (16-byte) MD5 hash of the "MS NT5 ISAKMPOAKLEY" string. The version number is the additional 4 bytes that denote the Windows operating system version as detailed in the first table that follows.

Operating system version	4-byte version number
Windows 2000	00 00 00 02
Windows XP	00 00 00 03
Windows Server 2003	00 00 00 04
Windows Vista	00 00 00 05
Windows Server 2008	00 00 00 06
Windows 7	00 00 00 07
Windows Server 2008 R2	00 00 00 08
Windows 8	00 00 00 09
Windows Server 2012	00 00 00 09
Windows 8.1	00 00 00 09
Windows Server 2012 R2	00 00 00 09
Windows 10	00 00 00 09
Windows Server 2016 Technical Preview	00 00 00 09

Common name	String representation	Wire representation (MD5 hash of string)	Version
Microsoft implementation Windows 2000	"MS NT5 ISAKMPOAKLEY" + version number 2	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 02	Windows 2000
Microsoft implementation Windows XP	"MS NT5 ISAKMPOAKLEY" + version number 3	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 03	Windows XP

Microsoft implementation Windows Server 2003	"MS NT5 ISAKMPOAKLEY" + version number 4	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 04	Windows Server 2003
Microsoft implementation Windows Vista	"MS NT5 ISAKMPOAKLEY" + version number 5	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 05	Windows Vista
Microsoft implementation Windows Server 2008	"MS NT5 ISAKMPOAKLEY" + version number 6	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 06	Windows Server 2008
Microsoft implementation Windows 7	"MS NT5 ISAKMPOAKLEY" + version number 7	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 07	Windows 7
Microsoft implementation Windows Server 2008 R2	"MS NT5 ISAKMPOAKLEY" + version number 8	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 08	Windows Server 2008 R2
Microsoft implementation Windows 8	"MS NT5 ISAKMPOAKLEY" + version number 9	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 09	Windows 8
Microsoft implementation Windows Server 2012	"MS NT5 ISAKMPOAKLEY" + version number 9	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 09	Windows Server 2012
Microsoft implementation Windows 8.1	"MS NT5 ISAKMPOAKLEY" + version number 9	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 09	Windows 8.1
Microsoft implementation Windows 10	"MS NT5 ISAKMPOAKLEY" + version number 9	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 09	Windows 10
Microsoft implementation Windows Server 2012 R2	"MS NT5 ISAKMPOAKLEY" + version number 9	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 09	Windows Server 2012 R2
Microsoft implementation Windows Server 2016 Technical Preview	"MS NT5 ISAKMPOAKLEY" + version number 9	1E 2B 51 69 05 99 1C 7D 7C 96 FC BF B5 87 E4 61 00 00 00 09	Windows Server 2016 Technical Preview

Kerberos authentication supported (as specified in [GSS])	"GSSAPI"	62 1B 04 BB 09 88 2A C1 E1 59 35 FE FA 24 AE EE	Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
NLB/MSCS fast failover supported	"Vid-Initial-Contact"	26 24 4D 38 ED DB 61 B3 17 2A 36 E3 D0 CF B8 19	Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
NLB/MSCS fast failover supported	"NLBS_PRESENT"	72 87 2B 95 FC DA 2E B7 08 EF E3 22 11 9B 49 71	Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
Fragmentation avoidance supported	"FRAGMENTATION"	40 48 B7 D5 6E BC E8 85 25 E7 DE 7F 00 D6 C2 D3	Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
NAT-T supported	"draft-ietf-ipsec-nat-t-ike-02\n"	90 CB 80 91 3E BB 69 6E 08 63 81 B5 EC 42 7B 1F	Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
NAT-T supported	"RFC 3947"	4A 13 1C 81 07 03 58 45 5C 57 28 F2 0E 95 45 2F	Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
AuthIP supported	"MS-MamieExists"	21 4C A4 FA FF A7 F3 2D 67 48 E5 30 33 95 AE 83	Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
CGA supported	"IKE CGA version 1"	E3 A5 96 6A 76 37 9F E7 07 22 82 31 E5 CE 86 52	Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview
Negotiation discovery supported	"MS-Negotiation Discovery Capable"	FB 1D E3 CD F3 41 B7 EA 16 B7 E5 BE 08 55 F1	Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016

		20	Technical Preview
Microsoft Xbox One 2013	"Microsoft Xbox One 2013"	8A A3 94 CF 8A 55 77 DC 31 10 C1 13 B0 27 A4 F2	Windows 10 and Windows Server 2016 Technical Preview
Xbox IKEv2 Negotiation	"Xbox IKEv2 Negotiation"	66 08 22 B3 A7 3A 24 41 49 57 8D 62 E0 EB 46 A0	Windows 10 and Windows Server 2016 Technical Preview
Security Realm ID	"MSFT IPsec Security Realm Id"	68 6A 8C BD FE 63 4B 40 51 46 FB 2B AF 33 E9 E8	Windows 10 and Windows Server 2016 Technical Preview

<7> [Section 2.1](#): IKE transport port assignments. Implemented in Windows NT 4.0, Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. These IKE extensions run on UDP ports 500 and 4500 only.

<8> [Section 2.2.1](#): NAT-T payload types. Implemented in Windows NT 4.0, Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The NAT-T payload type values that are used by each Windows version are specified in section [3.2.4.1](#).

<9> [Section 2.2.2](#): NAT-T UDP encapsulation modes. Implemented in Windows NT 4.0, Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The UDP encapsulation mode values used by each Windows version are specified in section [3.2.4.1](#).

<10> [Section 2.2.6](#): Error codes. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. This field may contain any Windows error code value. For more information about these codes, see [\[MS-ERREF\]](#).

<11> [Section 2.2.7](#): Error codes. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. This field may take on any Windows error code value. For more information about these codes, see [\[MS-ERREF\]](#).

<12> [Section 2.2.10](#): For Windows implementations, the "MSFT IPsec Security Realm Id" payload is 32 bytes long. The first 16 bytes is the MD5 hash of the vendor ID string. The remaining 16 bytes is an HMAC-MD5 encrypted string that identifies a particular security realm (as discussed in section [1.3.12](#)). The key that is used for this purpose is the string "SecurityRealmPolicyHmacKey".

<13> [Section 3.1.5](#): Initialization vectors (IV) choice for encrypted notifications sent prior to MM SA establishment:

If the peer sent the MS_NT5_ISAKMPOAKLEY notify vendor ID and the 4-byte version number is 0x00000002, 0x00000003, 0x00000004, or 0x00000005, (denoting Windows 2000, Windows XP, Windows Server 2003 and Windows Vista, respectively), the IV used in encrypting the notify is the last cipher block of the last sent packet. Otherwise, the IV will be the last cipher block of the last decrypted packet.

<14> [Section 3.2](#): [\[RFC3947\]](#) and [\[DRAFT-NATT\]](#) implementation. Both revisions are implemented in Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. Windows 2000 operating system Service Pack 4 (SP4), Windows XP operating system Service Pack 2 (SP2), and Windows Server 2003 implement the [\[DRAFT-NATT\]](#) revision. For more information, see [\[DRAFT-NATT\]](#).

<15> [Section 3.2.2](#): NAT-T keep-alive timer: A keep-alive message is sent every 20 seconds. Implemented in Windows NT 4.0, Windows 2000, Windows XP, Windows Server 2003, Windows Vista operating system with Service Pack 2 (SP2), Windows Server 2008 SP1, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. Windows Vista, Windows Vista operating system with Service Pack 1 (SP1), and Windows Server 2008 do not send keep-alives.

<16> [Section 3.2.4.1](#): [\[NAT-T IKE\]](#) message construction. Implemented in Windows 2000 Server operating system Service Pack 4 (SP4), Windows XP SP2, Windows Server 2003, Windows Vista, Windows Server 2008 operating system, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.

NAT-T revision support	Version
[DRAFT-NATT] and [RFC3947]	Windows Vista, Windows Server 2008, Windows 7, and Windows Server 2008 R2
[DRAFT-NATT]	Windows 2000 Server SP4, Windows XP SP2, and Windows Server 2003

Windows does not support NAT-T for IPv6 and therefore, does not send the NAT-T vendor IDs for IPv6 negotiations.

<17> [Section 3.3.2](#): Fragmentation timer. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The fragmentation timer is variable.

This timer interval is computed as the sum of the first 2 packet retransmission times. In Windows 2000, Windows XP, and Windows Server 2003, this is started from the IKE exchange (the second round trip in main mode). For Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview, this is started from the first main mode packet of the exchange. This should sum to 3 seconds. However, there is variance in the timer implementation up to ½ second per retransmission. This is an implementation artifact in that the underlying timer implementation is only this accurate. Hence the observed timer will be within the range of 2 to 4 seconds.

- In Windows 2000, Windows XP, and Windows Server 2003, both the initiator and the responder implement a fragmentation timer.
- In Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview, only the initiator implements a fragmentation timer.

[<18> Section 3.3.2](#): Fragment reassembly timer: Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The fragment reassembly timer is set to 70 seconds.

[<19> Section 3.3.5.3](#): Fragmentation active flag. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The Fragmentation active flag is set on receipt of a Fragment payload.

[<20> Section 3.3.5.3](#): IKE Message Fragmentation active flag behavior. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, and Windows Server 2008 R2. IKE messages are fragmented if the Fragmentation active flag is set, as per the conditions specified in section [3.3.6.1](#).

[<21> Section 3.3.6.1](#): Expiration of the fragmentation timer. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.

[<22> Section 3.5.4.1](#): Vendor ID payload. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. In Windows Vista and Windows Server 2008, the host sends the "Vid-Initial-Contact" vendor ID payload if it has no open TCP connections to the peer and new connection attempts cause the retransmission of SYN packets.

[<23> Section 3.5.7.1](#): QM SA idle timer. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The QM SA idle timer is set to 1 minute if the Fast Failover flag is set on the parent MM SA, and it is set to 5 minutes if the Fast Failover flag is not set.

[<24> Section 3.6.5.1](#): Vendor ID processing. Implemented in Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. This information is used to evaluate whether the MM SA should be allocated to a different host within the cluster. For more information, see [\[MSFT-WLBS\]](#).

[<25> Section 3.8.4.1](#): Nonce. Implemented in Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. Nonces are 32-byte random numbers that are generated from a FIPS-140-compliant random-number generator. For more information, see [\[FIPS140\]](#).

[<26> Section 3.8.6.1](#): Delete Retransmission timer. Implemented in Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The first retransmission occurs after 1 second. The time-out is doubled for each subsequent retransmission up to a maximum of six retransmissions. The maximum retransmission interval is capped at 16 seconds; so if the doubling of the previous interval exceeds 16 seconds, 16 seconds is used. The timer is started only if the remote host is a Windows peer, as identified by the "MS NT5 ISAKMPOAKLEY" vendor ID payload.

[<27> Section 3.8.7.1](#): Shutdown behavior. On shutdown for Windows 2000, Windows XP, and Windows Server 2003, IKE runs as specified in the footnote regarding the delete transmission timer in section [3.8.6.1](#). Note that the machine may shut down before the maximum number of retransmissions has actually been sent.

<28> [Section 3.8.7.2](#): After a delete has been triggered, Windows immediately sends the delete notify, and delays deleting the MM state internally to handle quick mode delete processing. Also, Windows does not immediately delete the quick mode(s) associated with the MM on receiving the MM delete, but waits for them to be deleted as a result of other protocol events.

<29> [Section 3.9.5.1](#): DOS Protection mode cookie generation. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The Windows implementation uses the following algorithm to generate the cookie (*prevTimeSlice* is a Boolean input parameter to the algorithm). iCookie is the Initiator Cookie as defined in [\[RFC2408\]](#) section 3.1.

```
Set Curtime to the 32 bits number of seconds
  elapsed since midnight, January 1, 1970
Set LocalIPAddr to the local IP address in
  network order
Set Localport to the 16 bits local listening UDP
  port (500 or 4500) in network order /* This port is the local port that the packet was
  received on. */
Set Peerport to the 16 bits remote port in
  network order
Set PeerIPAddr to the peer IP address in network order
Set cookieKey to a 50-byte random number
Set COOKIE KEY TIME to 150 seconds
If LocalIPAddr and PeerIPAddr are IPv4 addresses then
  Compute localAddr as 01 00 02 00 concatenated with LocalPort concatenated with LocalIPAddr
  concatenated with 26 bytes of 0
  Compute peerAddr as 01 00 02 00 concatenated with peerPort concatenated with peerIPAddr
  concatenated with 26 bytes of 0
end if
If LocalIPAddr and PeerIPAddr are IPv6 addresses then
  Compute localAddr as 0x01 0x00 0x02 0x00 concatenated with LocalPort
  concatenated with LocalIPAddr concatenated with 14 bytes of 0
  Compute peerAddr as 0x05 0x00 0x17 0x00 concatenated with peerPort
  concatenated with peerIPAddr concatenated with 14 bytes of 0
end if
Compute Curtime as ((Curtime + COOKIE KEY TIME) / COOKIE KEY TIME) * COOKIE KEY TIME
If prevTimeSlice is true then
  Compute Curtime as Curtime - COOKIE_KEY_TIME
End if
Compute tempCookie as SHA1(cookieKey concatenated with iCookie concatenated with peerAddr
  concatenated with localAddr concatenated with curTime)
Compute cookie as the first 8 bytes of tempCookie
```

<30> [Section 3.9.5.3](#): DOS Protection mode cookie validation. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The Windows implementation checks the validity of the **Responder Cookie** field by regenerating the cookie using the algorithm specified in section [3.9.5.1](#). The algorithm is as follows.

```
Set RCookie to the cookie field from message #2
Set prevTimeslice to FALSE
Compute cookie as described in <ref2>
If RCookie=cookie then
  RCookie is valid
Else
  Set prevTimeslice to TRUE
  Compute cookie as described in <ref2>
  If RCookie=cookie then
    RCookie is valid
  Else
    RCookie is invalid
  End if
End if
```


End if

[<31> Section 3.9.7](#): DOS Protection Threshold. Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. In Windows Vista and Windows Server 2008, Windows goes into DOS Protection mode if the number of negotiations for which only one message has been received from any initiator is more than 500. This is detected when the number of MM SAs in the MMSAD (see section [3.1.1](#)) is more than 500, and these SAs have only received one message. For a given IP address, if the number of negotiations for which only one message has been received is above 35, Windows drops new incoming negotiations from this IP address. For this reason, incoming messages have to come from multiple IP addresses in order to trigger the Denial of Service Protection mode. In Windows 2000, Windows XP, or Windows Server 2003, Windows goes into DOS protection mode immediately after setting the registry key and restarting the service.

Windows goes out of DOS Protection mode if the number of MM SAs in the MMSAD for which only one message has been received from any initiator is less than 100.

To enable the Windows DOS Protection mode in Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, or Windows Server 2016 Technical Preview, set the following Windows registry DWORD to 1.

SYSTEM\\CurrentControlSet\\Services\\IKEEXT\\Parameters\\EnableDOSProtect (DWORD)

To enable Windows DOS Protection mode in Windows 2000, Windows XP, or Windows Server 2003, set the following Windows registry DWORD to 1.

SYSTEM\\CurrentControlSet\\Services\\PolicyAgent\\Oakley\\EnableDOSProtect (DWORD). Stop and restart the PolicyAgent service for this setting to take effect.

[<32> Section 3.11](#): This feature is supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.

[<33> Section 3.11.1](#): This feature is supported in Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview.

[<34> Section 3.11.5.1](#): The Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview responder adds this attribute.

[<35> Section 3.11.5.2](#): The Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview initiator processes this attribute.

[<36> Section 3.12.1](#): Dead Peer Detection is implemented only for server-to-server site-to-site-tunnel mode IPsec tunnels on Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016 Technical Preview. Dead Peer Detection is not implemented on Windows 8, Windows 8.1, or Windows 10 for IKEv2-based VPN (that is, VPN Reconnect).

[<37> Section 3.12.7.1](#): Implemented in Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, Windows Server 2012, Windows 8.1, Windows Server 2012 R2, Windows 10, and Windows Server 2016 Technical Preview. The QM SA idle timer is set to 1 minute if the Fast Failover flag is set on the parent MM SA, and it is set to 5 minutes if the Fast Failover flag is not set.

<38> [Section 3.13.1](#): Implemented for Xbox multiplayer gaming scenarios.

<39> [Section 3.13.1](#): The integer value associated with the "Xbox IKEv2 Negotiation" vendor ID can be 0 or 1. These values denote different types of secure connections for Xbox multiplayer gaming. Their significance is beyond the scope of this document.

<40> [Section 3.14.1](#): Implemented for Xbox multiplayer gaming scenarios.

<41> [Section 3.14.1](#): This data is used to negotiate various IPsec SA proposals and different authentication methods for securing traffic for different game titles.

<42> [Section 3.14.5.1](#): If the IKE_SA_INIT message does not have a "MSFT IPsec Security Realm Id" vendor ID, Windows skips all security realm-based IPsec policies.

<43> [Section 3.14.5.1](#): If a IKEv2 responder receives an IKE_SA_INIT message with "MSFT IPsec Security Realm Id" vendor payload, then the Windows implementation does not send the optional CERTREQ payload ([RFC5996](#) section 1.2) in the IKE_SA_INIT response message.

7 Change Tracking

This section identifies changes that were made to this document since the last release. Changes are classified as New, Major, Minor, Editorial, or No change.

The revision class **New** means that a new document is being released.

The revision class **Major** means that the technical content in the document was significantly revised. Major changes affect protocol interoperability or implementation. Examples of major changes are:

- A document revision that incorporates changes to interoperability requirements or functionality.
- The removal of a document from the documentation set.

The revision class **Minor** means that the meaning of the technical content was clarified. Minor changes do not affect protocol interoperability or implementation. Examples of minor changes are updates to clarify ambiguity at the sentence, paragraph, or table level.

The revision class **Editorial** means that the formatting in the technical content was changed. Editorial changes apply to grammatical, formatting, and style issues.

The revision class **No change** means that no new technical changes were introduced. Minor editorial and formatting changes may have been made, but the technical content of the document is identical to the last released version.

Major and minor changes can be described further using the following change types:

- New content added.
- Content updated.
- Content removed.
- New product behavior note added.
- Product behavior note updated.
- Product behavior note removed.
- New protocol syntax added.
- Protocol syntax updated.
- Protocol syntax removed.
- New content added due to protocol revision.
- Content updated due to protocol revision.
- Content removed due to protocol revision.
- New protocol syntax added due to protocol revision.
- Protocol syntax updated due to protocol revision.
- Protocol syntax removed due to protocol revision.
- Obsolete document removed.

Editorial changes are always classified with the change type **Editorially updated**.

Some important terms used in the change type descriptions are defined as follows:

- **Protocol syntax** refers to data elements (such as packets, structures, enumerations, and methods) as well as interfaces.
- **Protocol revision** refers to changes made to a protocol that affect the bits that are sent over the wire.

The changes made to this document are listed in the following table. For more information, please contact dochelp@microsoft.com.

Section	Tracking number (if applicable) and description	Major change (Y or N)	Change type
1.3.11 Xbox Multiplayer Gaming (IKEv2)	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
1.3.12 IPsec Security Realm (IKEv2 transport mode)	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
2.2.10 Security Realm Vendor ID Payload (IKEv2)	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.13 Xbox Multiplayer Gaming (IKEv2) Vendor IDs Details	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.13.1 Abstract Data Model	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.13.2 Timers	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.13.3 Initialization	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.13.4 Higher-Layer Triggered Events	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.13.5 Message Processing Events and Sequencing Rules	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.13.5.1 Microsoft Xbox One 2013 Vendor ID	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.13.5.2 Xbox IKEv2 Negotiation Vendor ID	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.13.6 Timer Events	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.13.7 Other Local Events	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.

Section	Tracking number (if applicable) and description	Major change (Y or N)	Change type
3.14 Security Realm ID (IKEv2) Vendor IDs Details	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.14.1 Abstract Data Model	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.14.2 Timers	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.14.3 Initialization	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.14.4 Higher-Layer Triggered Events	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.14.5 Message Processing Events and Sequencing Rules	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.14.5.1 IKE_SA_INIT Messages	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.14.5.2 IKE_SA_AUTH and CREATE_CHILD_SA Messages	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.14.6 Timer Events	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
3.14.7 Other Local Events	Added section with content for Windows 10 and Windows Server 2016 Technical Preview.	Y	New content added.
6 Appendix A: Product Behavior	Updated the product applicability list and product behavior notes to include Windows 10 operating system.	Y	Content update.
6 Appendix A: Product Behavior	Updated the product behavior notes to include Windows Server 2016 Technical Preview operating system.	Y	Product behavior note updated.

8 Index

A

Abstract data model
CGA authentication ([section 3.1.1](#) 28, [section 3.4.1](#) 39)
[client](#) 42
[denial of service](#) 56
[Denial of Service \(DOS\)](#) 28
fast failover client ([section 3.1.1](#) 28, [section 3.5.1](#) 42)
fast failover server ([section 3.1.1](#) 28, [section 3.6.1](#) 43)
IKE fragmentation ([section 3.1.1](#) 28, [section 3.3.1](#) 34)
NAT traversal ([section 3.1.1](#) 28, [section 3.2.1](#) 32)
negotiation discovery ([section 3.1.1](#) 28, [section 3.7.1](#) 48)
reliable delete ([section 3.1.1](#) 28, [section 3.8.1](#) 52)
[server](#) 43
[Applicability](#) 18
[AUTH CGA Authentication Method Packet message](#) 22
[AUTH CGA Authentication Method packet](#) 22
[Authentication - cryptographically generated address](#) 14

C

[Capability negotiation](#) 19
CGA authentication
abstract data model ([section 3.1.1](#) 28, [section 3.4.1](#) 39)
[higher-layer triggered events](#) 39
[initialization](#) 39
[local events](#) 41
[message processing](#) 40
[overview](#) 38
[preconditions](#) 18
[prerequisites](#) 18
[receiving message #1](#) 40
[receiving message #2](#) 40
[receiving message #3](#) 40
[receiving message #4](#) 40
[receiving message #5](#) 40
[receiving message #6](#) 41
[sequencing rules](#) 40
[timer events](#) 41
[timers](#) 39
[Change tracking](#) 91
Client
[abstract data model](#) 42
[initialization](#) 42
overview ([section 3.1](#) 28, [section 3.5](#) 41)
[timers](#) 42
[Configuration Attribute \(IKEv2\) Packet message](#) 26
[Configuration Attribute packet](#) 26
[Correlation Payload \(IKEv2\) Packet message](#) 26
[Correlation Payload IKEv2 packet](#) 26

D

Data model - abstract
CGA authentication ([section 3.1.1](#) 28, [section 3.4.1](#) 39)
[client](#) 42
[denial of service](#) 56
[Denial of Service \(DOS\)](#) 28
fast failover client ([section 3.1.1](#) 28, [section 3.5.1](#) 42)
fast failover server ([section 3.1.1](#) 28, [section 3.6.1](#) 43)
IKE fragmentation ([section 3.1.1](#) 28, [section 3.3.1](#) 34)
NAT traversal ([section 3.1.1](#) 28, [section 3.2.1](#) 32)
negotiation discovery ([section 3.1.1](#) 28, [section 3.7.1](#) 48)
reliable delete ([section 3.1.1](#) 28, [section 3.8.1](#) 52)
[server](#) 43
[Delete retransmission timer expiration](#) 54
[Denial of service](#) 16
abstract data model ([section 3.1.1](#) 28, [section 3.9.1](#) 56)
[higher-layer triggered events](#) 56
[initialization](#) 56
[local events](#) 57
[message processing](#) 56
[overview](#) 55
[receiving message #1](#) 56
[receiving message #2](#) 56
[receiving message #3](#) 57
[sequencing rules](#) 56
[timer events](#) 57
[timers](#) 56

E

[Encapsulation modes - NAT-T syntax](#) 20
[Examples - negotiation discovery](#) 68

F

[Fast failover](#) 15
Fast failover client
abstract data model ([section 3.1.1](#) 28, [section 3.5.1](#) 42)
[expiration QM SA idle timer](#) 43
[higher-layer triggered events](#) 42
[initialization](#) 42
[local events](#) 43
[message processing](#) 42
[overview](#) 41
receiving message #1 ([section 3.5.5.1](#) 42, [section 3.5.5.2](#) 42)
[sequencing rules](#) 42

[timer events](#) 43
[timers](#) 42
Fast failover server
 abstract data model ([section 3.1.1](#) 28, [section 3.6.1](#) 43)
 [higher-layer triggered events](#) 44
 [initialization](#) 43
 [local events](#) 44
 [message processing](#) 44
 [overview](#) 43
 [receiving message #1](#) 44
 [receiving message #2](#) 44
 [sequencing rules](#) 44
 [timer events](#) 44
 [timers](#) 43
[Fields - vendor-extensible](#) 19
[Fragment Payload packet](#) 21
[Fragmentation](#) 14

G

[Glossary](#) 8

H

Higher-layer triggered events
 [CGA authentication](#) 39
 [denial of service](#) 56
 [fast failover client](#) 42
 [fast failover server](#) 44
 [IKE fragmentation](#) 36
 [NAT traversal](#) 33
 [negotiation discovery](#) 49
 [protocol](#) 30
 [reliable delete](#) 53

I

[ID_IPV6_CGA Identification Type Packet message](#) 22
[ID_IPV6_CGA packet](#) 22
[IKE fragmentation](#) 14
 abstract data model ([section 3.1.1](#) 28, [section 3.3.1](#) 34)
 [fragmentation reassembly timer expiration](#) 37
 [fragmentation timer expiration](#) 37
 [higher-layer triggered events](#) 36
 [initialization](#) 35
 [local events](#) 38
 [message processing](#) 36
 [overview](#) 34
 [receiving message #1](#) 36
 [receiving message #2](#) 36
 [receiving other messages](#) 36
 [sequencing rules](#) 36
 [timer events](#) 37
 [timers](#) 35
[IKE Message Fragment message](#) 21
[IKE message fragment syntax](#) 21

IKE MM SA negotiation ([section 3.2.4.1](#) 33, [section 3.3.4.1](#) 36, [section 3.4.4.1](#) 39, [section 3.5.4.1](#) 42, [section 3.6.4.1](#) 44)
[IKE/AuthIP coexistence](#) 16
[Implementer - security considerations](#) 70
[Inbound packets](#) 50
[Index of security parameters](#) 70
[Informative references](#) 12
Initialization
 [CGA authentication](#) 39
 [client](#) 42
 [denial of service](#) 56
 [fast failover client](#) 42
 [fast failover server](#) 43
 [IKE fragmentation](#) 35
 [negotiation discovery](#) 49
 [protocol](#) 29
 [reliable delete](#) 53
 [server](#) 43
[Initialization - NAT traversal](#) 32
[Introduction](#) 8

L

Local events
 [CGA authentication](#) 41
 [denial of service](#) 57
 [fast failover client](#) 43
 [fast failover server](#) 44
 [IKE fragmentation](#) 38
 [NAT traversal](#) 34
 [negotiation discovery](#) 52
 [protocol](#) 31
 [reliable delete](#) 54

M

Message processing
 [CGA authentication](#) 40
 [denial of service](#) 56
 [fast failover client](#) 42
 [fast failover server](#) 44
 [IKE fragmentation](#) 36
 [NAT traversal](#) 33
 [negotiation discovery](#) 51
 [protocol](#) 30
 receiving message #1 ([section 3.2.5.1](#) 33, [section 3.3.5.1](#) 36, [section 3.4.5.1](#) 40, [section 3.5.5.1](#) 42, [section 3.5.5.2](#) 42, [section 3.6.5.1](#) 44, [section 3.7.5.1](#) 51, [section 3.8.5.1](#) 53, [section 3.8.5.2](#) 54, [section 3.9.5.1](#) 56)
 receiving message #2 ([section 3.2.5.2](#) 33, [section 3.3.5.2](#) 36, [section 3.4.5.2](#) 40, [section 3.6.5.2](#) 44, [section 3.7.5.2](#) 51, [section 3.9.5.2](#) 56)
 receiving message #3 ([section 3.4.5.3](#) 40, [section 3.9.5.3](#) 57)
 [receiving message #4](#) 40
 receiving message #5 ([section 3.4.5.5](#) 40, [section 3.7.5.3](#) 51)
 receiving message #6 ([section 3.4.5.6](#) 41, [section 3.7.5.4](#) 52)

receiving other messages ([section 3.2.5.3](#) 33,
[section 3.3.5.3](#) 36)
[reliable delete](#) 53
Messages
[AUTH CGA Authentication Method Packet](#) 22
[Configuration Attribute \(IKEv2\) Packet](#) 26
[Correlation Payload \(IKEv2\) Packet](#) 26
[ID_IPV6 CGA Identification Type Packet](#) 22
[IKE Message Fragment](#) 21
[NAT-T Payload Types](#) 20
[NAT-T UDP Encapsulation Modes](#) 20
[Notify Payload \(IKEv2\) Packet](#) 25
[Notify Payload Packet](#) 23
[Security Realm Vendor ID Payload \(IKEv2\)](#) 27
[syntax](#) 20
[transport](#) 20

N

NAT traversal
abstract data model ([section 3.1.1](#) 28, [section 3.2.1](#) 32)
[higher-layer triggered events](#) 33
[initialization](#) 32
[local events](#) 34
[message processing](#) 33
overview ([section 1.3.1](#) 13, [section 3.2](#) 31)
[payload types syntax](#) 20
[receiving message #1](#) 33
[receiving message #2](#) 33
[receiving other messages](#) 33
[sequencing rules](#) 33
[timer events](#) 33
[timers](#) 32
[UDP encapsulation modes syntax](#) 20
[NAT-T Payload Types message](#) 20
[NAT-T UDP Encapsulation Modes message](#) 20
[Negotiation discovery](#) 15
abstract data model ([section 3.1.1](#) 28, [section 3.7.1](#) 48)
[higher-layer triggered events](#) 49
[initialization](#) 49
[local events](#) 52
[message processing](#) 51
[overview](#) 44
[receiving message #1](#) 51
[receiving message #2](#) 51
[receiving message #5](#) 51
[receiving message #6](#) 52
[sequencing rules](#) 51
[timer events](#) 52
[timers](#) 49
[Negotiation discovery example](#) 68
[Negotiation discovery security](#) 70
[Normative references](#) 11
[Notify Payload \(IKEv2\) Packet message](#) 25
[Notify Payload Packet message](#) 23
[Notify Payload packet](#) 23
[Notify Payload IKEV2 packet](#) 25

O

Other local events
[server](#) 44
[Outbound packets](#) 49
[Overview](#) 13
[Overview \(synopsis\)](#) 13

P

Packets
[inbound](#) 50
[outbound](#) 49
[Parameters - security index](#) 70
[Preconditions](#) 18
[CGA authentication](#) 18
[general](#) 18
[Prerequisites](#) 18
[CGA authentication](#) 18
[general](#) 18
[Product behavior](#) 71
Protocol
[higher-layer triggered events](#) 30
[initialization](#) 29
[local events](#) 31
[message processing](#) 30
[sequencing rules](#) 30
[timer events](#) 31
[timers](#) 29
Protocol Details
[overview](#) 28

Q

[QM SA idle timer expiration](#) 43
[QM SA negotiation](#) 43

R

[References](#) 11
[informative](#) 12
[normative](#) 11
[Relationship to other protocols](#) 18
[Reliable delete](#) 15
abstract data model ([section 3.1.1](#) 28, [section 3.8.1](#) 52)
[delete retransmission timer expiration](#) 54
[higher-layer triggered events](#) 53
[initialization](#) 53
[local events](#) 54
[message processing](#) 53
[overview](#) 52
receiving message #1 ([section 3.8.5.1](#) 53, [section 3.8.5.2](#) 54)
[sequencing rules](#) 53
[shutdown](#) 55
[timer events](#) 54
[timers](#) 53
[RFC cross-reference extension](#) 17

S

[SA deletion](#) 53

Security

- [implementer considerations](#) 70
- [negotiation discovery security](#) 70
- [parameter index](#) 70

[Security Realm Vendor ID Payload \(IKEv2\) message](#) 27

Sequencing rules

- [CGA authentication](#) 40
- [denial of service](#) 56
- [fast failover client](#) 42
- [fast failover server](#) 44
- [IKE fragmentation](#) 36
- [NAT traversal](#) 33
- [negotiation discovery](#) 51
- [protocol](#) 30
- [reliable delete](#) 53

Server

- [abstract data model](#) 43
- [initialization](#) 43
- [other local events](#) 44
- overview ([section 3.1](#) 28, [section 3.6](#) 43)
- [timer events](#) 44
- [timers](#) 43

[Shutdown](#) 55

[Standards assignments](#) 19

Syntax

- [IKE message fragment](#) 21
- [messages](#) 20
- [NAT-T payload types](#) 20
- [NAT-T UDP encapsulation modes](#) 20

T

Time events

- [expiration QM SA idle timer](#) 43
- [fast failover client](#) 43

Timer events

- [CGA authentication](#) 41
- [denial of service](#) 57
- [fast failover server](#) 44
- [IKE fragmentation](#) 37
- [NAT traversal](#) 33
- [negotiation discovery](#) 52
- [protocol](#) 31
- [reliable delete](#) 54
- [server](#) 44

Timers

- [CGA authentication](#) 39
- [client](#) 42
- [delete retransmission timer expiration](#) 54
- [denial of service](#) 56
- [fast failover client](#) 42
- [fast failover server](#) 43
- [fragmentation reassembly timer expiration](#) 37
- [fragmentation timer expiration](#) 37
- [IKE fragmentation](#) 35
- [NAT traversal](#) 32
- [negotiation discovery](#) 49
- [protocol](#) 29

[reliable delete](#) 53

[server](#) 43

[Tracking changes](#) 91

[Transport](#) 20

Triggered events - higher-layer

- [CGA authentication](#) 39
- [denial of service](#) 56
- [fast failover client](#) 42
- [fast failover server](#) 44
- [IKE fragmentation](#) 36
- [NAT traversal](#) 33
- [negotiation discovery](#) 49
- [protocol](#) 30
- [reliable delete](#) 53

V

[Vendor-extensible fields](#) 19

[Versioning](#) 19