

Writing Protocol Models

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

The IETF process depends on peer review. However, IETF documents are generally written to be useful for implementors, not reviewers. In particular, while great care is generally taken to provide a complete description of the state machines and bits on the wire, this level of detail tends to get in the way of initial understanding. This document describes an approach for providing protocol "models" that allow reviewers to quickly grasp the essence of a system.

1. Introduction

The IETF process depends on peer review. However, in many cases, the documents submitted for publication are extremely difficult to review. Because reviewers have only limited amounts of time, this leads to extremely long review times, inadequate reviews, or both. In our view, a large part of the problem is that most documents fail to present an architectural model for how the protocol operates, opting instead to simply describe the protocol and let the reviewer figure it out.

This is acceptable when documenting a protocol for implementors, because they need to understand the protocol in any case; but it dramatically increases the strain on reviewers. Reviewers need to get the big picture of the system and then focus on particular points. They simply do not have time to give the entire document the attention an implementor would.

One way to reduce this load is to present the reviewer with a MODEL -- a short description of the system in overview form. This provides the reviewer with the context to identify the important or difficult pieces of the system and focus on them for review. As a side benefit, if the model is done first, it can be serve as an aid to the detailed protocol design and a focus for early review, prior to protocol completion. The intention is that the model would either be the first section of the protocol document or be a separate document provided with the protocol.

2. The Purpose of a Protocol Model

A protocol model needs to answer three basic questions:

1. What problem is the protocol trying to achieve?
2. What messages are being transmitted and what do they mean?
3. What are the important, but unobvious, features of the protocol?

The basic idea is to provide enough information that the reader could design a protocol which was roughly isomorphic to the protocol being described. Of course, this doesn't mean that the protocol would be identical, but merely that it would share most important features. For instance, the decision to use a KDC-based authentication model is an essential feature of Kerberos [KERBEROS]. By contrast, the use of ASN.1 is a simple implementation decision. S-expressions -- or XML, had it existed at the time -- would have served equally well.

The purpose of a protocol model is explicitly not to provide a complete or alternate description of the protocol being discussed. Instead, it is to provide a big picture overview of the protocol so that readers can quickly understand the essential elements of how it works.

3. Basic Principles

In this section we discuss basic principles that should guide your presentation.

3.1. Less is more

Humans are only capable of keeping a very small number of pieces of information in their head at once. Because we're interested in ensuring that people get the big picture, we have to dispense with a lot of detail. That's good, not bad. The simpler you can make things the better.

3.2. Abstraction is good

A key technique for representing complex systems is to try to abstract away pieces. For instance, maps are better than photographs for finding out where you want to go because they provide an abstract, stylized, view of the information you're interested in. Don't be afraid to compress multiple protocol elements into a single abstract piece for pedagogical purposes.

3.3. A few well-chosen details sometimes help

The converse of the previous principle is that sometimes details help to bring a description into focus. Many people work better when given examples. Thus, it's often a good approach to talk about the material in the abstract and then provide a concrete description of one specific piece to bring it into focus. Authors should focus on the normal path. Error cases and corner cases should only be discussed where they help illustrate an important point.

4. Writing Protocol Models

Our experience indicates that it is easiest to grasp protocol models when they are presented in visual form. We recommend a presentation format centered around a few key diagrams, with explanatory text for each. These diagrams should be simple and typically consist of "boxes and arrows" -- boxes representing the major components, arrows representing their relationships, and labels indicating important features.

We recommend a presentation structured in three parts to match the three questions mentioned in the previous sections. Each part should contain 1-3 diagrams intended to illustrate the relevant points.

4.1. Describe the problem you're trying to solve

The most critical task that a protocol model must perform is to explain what the protocol is trying to achieve. This provides crucial context for understanding how the protocol works, and whether it meets its goals. Given the desired goals, an experienced reviewer will usually have an idea of how they would approach the problem and, thus, be able to compare that approach with the approach taken by the protocol under review.

The "Problem" section of the model should start with a short statement of the environments in which the protocol is expected to be used. This section should describe the relevant entities and the likely scenarios under which they would participate in the protocol. The Problem section should feature a diagram of the major

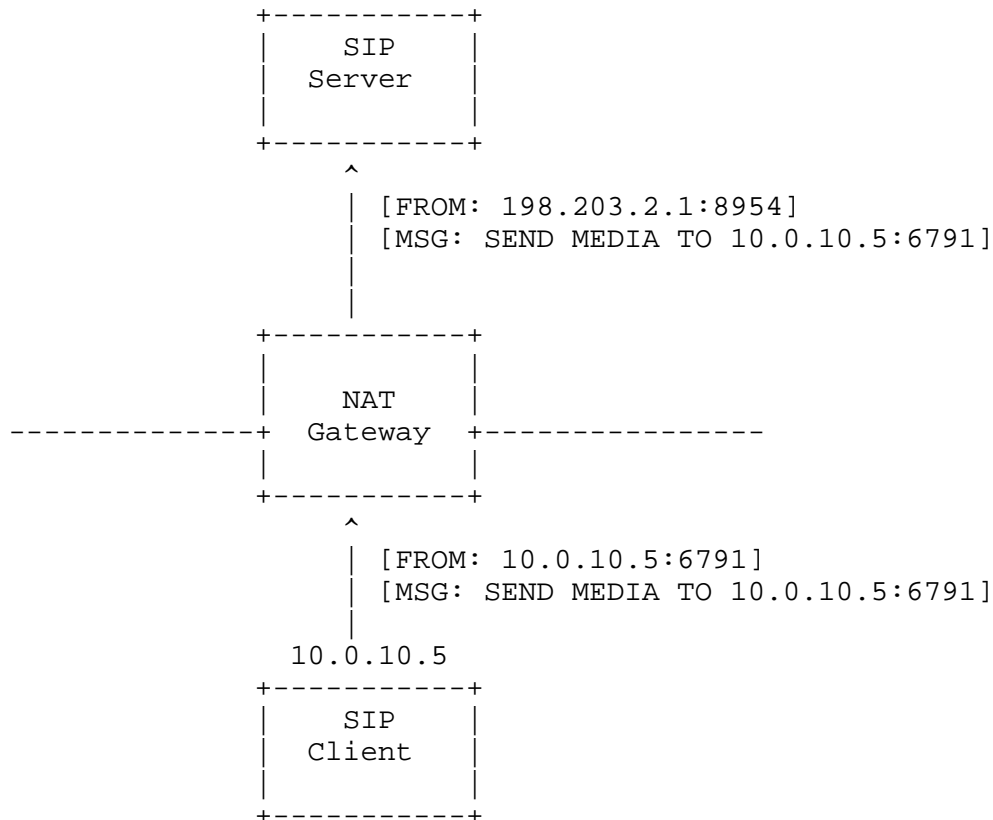
communicating parties and their inter-relationships. It is particularly important to lay out the trust relationships between the various parties, as these are often unobvious.

4.1.1. Example: STUN (RFC 3489)


STUN [STUN] is a UNilateral Self-Address Fixing (UNSAF) [UNSAF] protocol that allows a machine located behind a NAT to determine what its external apparent IP address is. Although STUN provides a complete and thorough description of the operation of the protocol, it does not provide a brief, up-front overview suitable for a quick understanding of its operation. The rest of this section shows what a suitable overview might look like.

Network Address Translation (NAT) makes it difficult to run a number of classes of service from behind the NAT gateway. This is particularly a problem when protocols need to advertise address/port pairs as part of the application layer protocol. Although the NAT can be configured to accept data destined for that port, address translation means the address that the application knows about is not the same as the one on which it is reachable.

Consider the scenario represented in the figure below. A SIP client is initiating a session with a SIP server in which it wants the SIP server to send it some media. In its Session Description Protocol (SDP) [SDP] request it provides the IP address and port on which it is listening. However, unbeknownst to the client, a NAT is in the way. The NAT translates the IP address in the header, but unless it is SIP aware, it doesn't change the address in the request. The result is that the media goes into a black hole.



The purpose of STUN is to allow clients to detect this situation and determine the address mapping. They can then place the appropriate address in their application-level messages. This is done by using an external STUN server. That server is able to determine the translated address and tell the STUN client, as shown below.



4.2. Describe the protocol in broad overview

Once the problem has been described, the next task is to give a broad overview of the protocol. This means showing, either in "ladder diagram" or "boxes and arrows" form, the protocol messages that flow between the various networking agents. This diagram should be accompanied with explanatory text that describes the purpose of each message and the MAJOR data elements.

This section SHOULD NOT contain detailed descriptions of the protocol messages or of each data element. In particular, bit diagrams, ASN.1 modules, and XML schema SHOULD NOT be shown. The purpose of this section is not to provide a complete description of the protocol, but to provide enough of a map that a person reading the full protocol document can see where each specific piece fits.

In certain cases, it may be helpful to provide a state machine description of the behavior of network elements. However, such state machines should be kept as minimal as possible. Remember that the purpose is to promote high-level comprehension, not complete understanding.

4.2.1. Example: DCCP

Datagram Congestion Control Protocol [DCCP] is a protocol for providing datagram transport with network-friendly congestion avoidance behavior. The DCCP base protocol document is over 100 pages long and the congestion control mechanisms themselves are separate. Therefore, it is very helpful to have an architectural overview of DCCP that abstracts away the details. The remainder of this section is an attempt to do so.

NOTE: The author of this document was on the DCCP review team and his experience with that document was one of the motivating factors for this document. Since the review, the DCCP authors have added some overview material, some of which derives from earlier versions of this document.

Although DCCP is datagram-oriented like UDP, it is stateful like TCP. Connections go through the following phases:

1. Initiation
2. Feature negotiation
3. Data transfer
4. Termination

4.2.1.1. Initiation

As with TCP, the initiation phase of DCCP involves a three-way handshake, shown below.

Client		Server
-----		-----
DCCP-Request	->	
[Ports, Service,		
Features]		
	<-	DCCP-Response
		[Features,
		Cookie]
DCCP-Ack	->	
[Features,		
Cookie]		

DCCP 3-way handshake

In the DCCP-Request message, the client tells the server the name of the service it wants to talk to and the ports it wants to communicate on. Note that ports are not tightly bound to services, as they are in TCP or UDP common practice. It also starts feature negotiation. For pedagogical reasons, we will present feature negotiation

separately in the next section. However, realize that the early phases of feature negotiation happen concurrently with initiation.

In the DCCP-Response message, the server tells the client that it is willing to accept the connection and continues feature negotiation. In order to prevent SYN flood-style DOS attacks, DCCP incorporates an IKE-style cookie exchange. The server can provide the client with a cookie that contains all of the negotiation state. This cookie must be echoed by the client in the DCCP-Ack, thus removing the need for the server to keep state.

In the DCCP-Ack message, the client acknowledges the DCCP-Response and returns the cookie to permit the server to complete its side of the connection. As indicated above, this message may also include feature negotiation messages.

4.2.1.2. Feature Negotiation

In DCCP, feature negotiation is performed by attaching options to other DCCP packets. Thus, feature negotiation can be piggybacked on any other DCCP message. This allows feature negotiation during connection initiation as well as during data flow.

Somewhat unusually, DCCP features are one-sided. Thus, it's possible to have a different congestion control regime for data sent from client to server than from server to client.

Feature negotiation is done with the Change and Confirm options. There are four feature negotiation options in all: Change L, Confirm L, Change R, and Confirm R. The "L" options are sent by the feature location, where the feature is maintained, and the "R" options are sent by the feature remote.

A Change R message says to the peer "change this option setting on your side". The peer can respond with a Confirm L, meaning "I've changed it". Some features allow Change R options to contain multiple values, sorted in preference order. For example:

Client	Server
-----	-----
Change R(CCID, 2) -->	
	<-- Confirm L(CCID, 2)
	* agreement that CCID/Server = 2 *
Change R(CCID, 3 4) -->	
	<-- Confirm L(CCID, 4, 4 2)
	* agreement that CCID/Server = 4 *

In the second exchange, the client requests that the server use either CCID 3 or CCID 4, with 3 preferred. The server chooses 4 and supplies its preference list, "4 2".

The Change L and Confirm R options are used for feature negotiations that are initiated by the feature location. In the following example, the server requests that CCID/Server be set to 3 or 2 (with 3 being preferred), and the client agrees.

Client	Server
-----	-----
	<-- Change L(CCID, 3 2)
Confirm R(CCID, 3, 3 2) -->	
* agreement that CCID/Server = 3 *	

4.2.1.3. Data Transfer

Rather than have a single congestion control regime, as in TCP, DCCP offers a variety of negotiable congestion control regimes. The DCCP documents describe two congestion control regimes: additive increase, multiplicative decrease (CCID-2 [CCID2]), and TCP-friendly rate control (CCID-3 [CCID3]). CCID-2 is intended for applications that want maximum throughput. CCID-3 is intended for real-time applications that want smooth response to congestion.

4.2.1.3.1. CCID-2

CCID-2's congestion control is extremely similar to that of TCP. The sender maintains a congestion window and sends packets until that window is full. Packets are Aced by the receiver. Dropped packets and ECN [ECN] are used to indicate congestion. The response to congestion is to halve the congestion window. One subtle difference between DCCP and TCP is that the Acks in DCCP must contain the sequence numbers of all received packets (within a given window), not just the highest sequence number, as in TCP.

4.2.1.3.2. CCID-3

CCID-3 is an equation-based form of rate control, intended to provide smoother response to congestion than CCID-2. The sender maintains a "transmit rate". The receiver sends Ack packets that contain information about the receiver's estimate of packet loss. The sender uses this information to update its transmit rate. Although CCID-3 behaves somewhat differently than TCP in its short-term congestion response, it is designed to operate fairly with TCP over the long term.

4.2.1.4. Termination

Connection termination in DCCP is initiated by sending a Close message. Either side can send a Close message. The peer then responds with a Reset message, at which point the connection is closed. The side that sent the Close message must quietly preserve the socket in TIMEWAIT state for 2MSL.

Client		Server
-----		-----
Close	->	
	<-	Reset
[Remains in TIMEWAIT]		

Note that the server may wish to close the connection but not remain in TIMEWAIT (e.g., due to a desire to minimize server-side state). In order to accomplish this, the server can elicit a Close from the client by sending a CloseReq message and, thus, keep the TIMEWAIT state on the client.

4.3. Describe any important protocol features

The final section (if there is one) should contain an explanation of any important protocol features that are not obvious from the previous sections. In the best case, all the important features of the protocol would be obvious from the message flow. However, this isn't always the case. This section is an opportunity for the author to explain those features. Authors should think carefully before writing this section. If there are no important points to be made, they should not populate this section.

Examples of the kind of feature that belongs in this section include: high-level security considerations, congestion control information, and overviews of the algorithms that the network elements are intended to follow. For instance, if you have a routing protocol, you might use this section to sketch out the algorithm that the router uses to determine the appropriate routes from protocol messages.

4.3.1. Example: WebDAV COPY and MOVE

The WebDAV standard [WEBDAV] is fairly terse, preferring to define the required behaviors and let the reader work out the implications. In some situations, explanatory material that details those implications can help the reader understand the overall model. The rest of this section describes one such case.

WebDAV [WEBDAV] includes both a COPY method and a MOVE method. While a MOVE can be thought of as a COPY followed by DELETE, COPY+DELETE and MOVE aren't entirely equivalent.

The use of COPY+DELETE as a substitute for MOVE is problematic because of the creation of the intermediate file. Consider the case where the user is approaching a quota boundary. A COPY+DELETE should be forbidden because it would temporarily exceed the quota. However, a simple rename should work in this situation.

The second issue is permissions. The WebDAV permissions model allows the server to grant users permission to rename files, but not to create new ones. This is unusual in ordinary filesystems, but nothing prevents it in WebDAV. This is clearly not possible if a client uses COPY+DELETE to do a MOVE.

Finally, a COPY+DELETE does not produce the same logical result as would be expected with a MOVE. Because COPY creates a new resource, it is permitted (but not required) to use the time of new file creation as the creation date property. By contrast, the expectation for MOVE is that the renamed file will have the same properties as the original.

5. Formatting Issues


The requirement that Internet-Drafts and RFCs be renderable in ASCII is a significant obstacle when writing the sort of graphics-heavy document being described here. Authors may find it more convenient to do a separate protocol model document in Postscript or PDF and simply make it available at review time -- though an archival version would certainly be handy.

6. A Complete Example: Internet Key Exchange (IKE)

Internet Key Exchange (IKE) [IKE] is one of the most complicated security protocols ever designed by the IETF. Although the basic IKE core is a fairly straightforward Diffie-Hellman-based handshake, this can often be difficult for new readers to understand abstractly, apart from the protocol details. The remainder of this section provides overview of IKE suitable for those new readers.

6.1. Operating Environment

Internet key Exchange (IKE) [IKE] is a key establishment and parameter negotiation protocol for Internet protocols. Its primary application is for establishing security associations (SAs) [IPSEC] for IPsec AH [AH] and ESP [ESP].



The general deployment model for IKE is shown above. The IPsec engines and IKE engines typically are separate modules. When no security association exists for a packet that needs to be processed (either sent or received), the IPsec engine contacts the IKE engine and asks it to establish an appropriate SA. The IKE engine contacts the appropriate peer and uses IKE to establish the SA. Once the IKE handshake is finished it registers the SA with the IPsec engine.

In addition, IKE traffic between the peers can be used to refresh keying material or adjust operating parameters, such as algorithms.

6.1.1. Initiator and Responder

Although IPsec is basically symmetrical, IKE is not. The party who sends the first message is called the INITIATOR. The other party is called the RESPONDER. In the case of TCP connections, the INITIATOR will typically be the peer doing the active open (i.e., the client).

6.1.2. Perfect Forward Secrecy

One of the major concerns in IKE design was that traffic be protected even if the keying material of the nodes was later compromised, provided that the session in question had terminated and so the session-specific keying material was gone. This property is often called Perfect Forward Secrecy (PFS) or back traffic protection.

6.1.3. Denial of Service Resistance

Because IKE allows arbitrary peers to initiate computationally-expensive cryptographic operations, it potentially allows resource consumption denial of service (DoS) attacks to be mounted against the IKE engine. IKE includes countermeasures designed to minimize this risk.

6.1.4. Keying Assumptions

Because Security Associations are essentially symmetric, both sides must, in general, be authenticated. Because IKE needs to be able to establish SAs between a broad range of peers with various kinds of prior relationships, IKE supports a very flexible keying model. Peers can authenticate via shared keys, digital signatures (typically from keys vouched for by certificates), or encryption keys.

6.1.5. Identity Protection

Although IKE requires the peers to authenticate to each other, it was considered desirable by the working group to provide some identity protection for the communicating peers. In particular, the peers should be able to hide their identity from passive observers and one peer should be able to require the author to authenticate before they self-identify. In this case, the designers chose to make the party who speaks first (the INITIATOR) identify first.

6.2. Protocol Overview

At a very high level, there are two kinds of IKE handshake:

- (1) Those that establish an IKE security association.
- (2) Those that establish an AH or ESP security association.

When two peers that have never communicated before need to establish an AH/ESH SA, they must first establish an IKE SA. This allows them to exchange an arbitrary amount of protected IKE traffic. They can then use that SA to do a second handshake to establish SAs for AH and ESP. This process is shown in schematic form below. The notation E(SA,XXXX) is used to indicate that traffic is encrypted under a given SA.

Initiator -----		Responder -----
Handshake MSG	->	
	<-	Handshake MSG
	[...]	
		\ Stage 1: \ Establish IKE / SA (IKesa) /
E(IKesa, Handshake MSG)	->	
	<-	E(IKesa, Handshake MSG)
		\ Stage 2: \ Establish AH/ESP / SA

The two kinds of IKE handshake

IKE terminology is somewhat confusing, referring under different circumstances to "phases" and "modes". For maximal clarity we will refer to the Establishment of the IKE SA as "Stage 1" and the Establishment of AH/ESP SAs as "Stage 2". Note that it's quite possible for there to be more than one Stage 2 handshake, once Stage 1 has been finished. This might be useful for establishing multiple AH/ESP SAs with different cryptographic properties.

The Stage 1 and Stage 2 handshakes are actually rather different, because the Stage 2 handshake can, of course, assume that its traffic is being protected with an IKE SA. Accordingly, we will first discuss Stage 1 and then Stage 2.

6.2.1. Stage 1

There are a large number of variants of the IKE Stage 1 handshake, necessitated by use of different authentication mechanisms. However, broadly speaking Stage 1 handshakes fall into one of two basic categories: MAIN MODE, which provides identity protection and DoS resistance, and AGGRESSIVE MODE, which does not. We will cover MAIN MODE first.

6.2.1.1. Main Mode

Main Mode is a six message (3 round trip) handshake, which offers identity protection and DoS resistance. An overview of the handshake is below.

Initiator -----		Responder -----
CookieI, Algorithms	->	\ Parameter
	<-	CookieR, Algorithms / Establishment
CookieR, Nonce, Key Exchange	->	
	<-	Nonce, Key Exchange \ Establish / Shared key
E(IKEsa, Auth Data)	->	
	<-	E(IKEsa, Auth data) \ Authenticate / Peers

IKE Main Mode handshake (Stage 1)

In the first round trip, the Initiator offers a set of algorithms and parameters. The Responder picks out the single set that it likes and responds with that set. It also provides CookieR, which will be used to prevent DoS attacks. At this point, there is no secure association but the peers have tentatively agreed upon parameters. These parameters include a Diffie-Hellman (DH) group, which will be used in the second round trip.

In the second round trip, the Initiator sends the key exchange information. This generally consists of the Initiator's Diffie-Hellman public share (Yi). He also supplies CookieR, which was provided by the responder. The Responder replies with his own DH share (Yr). At this point, both Initiator and Responder can compute the shared DH key (ZZ). However, there has been no authentication and, therefore, they don't know with any certainty that the connection hasn't been attacked. Note that as long as the peers generate fresh DH shares for each handshake, PFS will be provided.

Before we move on, let's take a look at the cookie exchange. The basic anti-DoS measure used by IKE is to force the peer to demonstrate that it can receive traffic from you. This foils blind attacks like SYN floods [SYNFLOOD] and also makes it somewhat easier to track down attackers. The cookie exchange serves this role in IKE. The Responder can verify that the Initiator supplied a valid CookieR before doing the expensive DH key agreement. This does not totally eliminate DoS attacks, because an attacker who was willing to reveal his location could still consume server resources; but it does protect against a certain class of blind attack.

In the final round trip, the peers establish their identities. Because they share an (unauthenticated) key, they can send their identities encrypted, thus providing identity protection from eavesdroppers. The exact method of proving identity depends on what form of credential is being used (signing key, encryption key, shared secret, etc.), but in general you can think of it as a signature over some subset of the handshake messages. So, each side would supply its certificate and then sign using the key associated with that certificate. If shared keys are used, the authentication data would be a key ID and a MAC. Authentication using public key encryption follows similar principles, but is more complicated. Refer to the IKE document for more details.

At the end of the Main Mode handshake, the peers share:

- (1) A set of algorithms for encryption of further IKE traffic.
- (2) Traffic encryption and authentication keys.
- (3) Mutual knowledge of the peer's identity.

6.2.1.2. Aggressive Mode

Although IKE Main Mode provides the required services, there was concern that the large number of round trips required added, excessive latency. Accordingly, an Aggressive Mode was defined. Aggressive mode packs more data into fewer messages, and thus reduces latency. However, it does not provide identity protection or protection against DoS.

Initiator	Responder
-----	-----
Algorithms, Nonce, Key Exchange,	
->	
	<-
	Algorithms, Nonce, Key Exchange, Auth Data
Auth Data	->

IKE Aggressive Mode Handshake (Stage 1)

After the first round trip, the peers have all the required properties, but the Initiator has not authenticated to the Responder. The third message closes the loop by authenticating the Initiator. Note that since the authentication data is sent in the clear, no identity protection is provided; and because the Responder does the DH key agreement without a round trip to the Initiator, there is no DoS protection

6.2.2. Stage 2

Stage 1 on its own isn't very useful. The purpose of IKE, after all, is to establish associations to be used to protect other traffic, not merely to establish IKE SAs. Stage 2 (what IKE calls "Quick Mode") is used for this purpose. The basic Stage 2 handshake is shown below.

Initiator		Responder
-----		-----
AH/ESP parameters, Algorithms, Nonce, Handshake Hash	->	
	<-	AH/ESP parameters, Algorithms, Nonce, Handshake Hash
Handshake Hash	->	

The Basic IKE Quick Mode (Stage 2)

As with quick mode, the first two messages establish the algorithms and parameters while the final message is a check over the previous messages. In this case, the parameters also include the transforms to be applied to the traffic (AH or ESP) and the kinds of traffic that are to be protected. Note that there is no key exchange information shown in these messages.

In this version of Quick Mode, the peers use the preexisting Stage 1 keying material to derive fresh keying material for traffic protection (with the nonces to ensure freshness). Quick mode also allows for a new Diffie-Hellman handshake for per-traffic key PFS. In that case, the first two messages shown above would also include Key Exchange payloads, as shown below.

Initiator		Responder
-----		-----
AH/ESP parameters, Algorithms, Nonce, Key Exchange, Handshake Hash	->	
	<-	AH/ESP parameters, Algorithms, Nonce, Key Exchange, Handshake Hash
Handshake Hash	->	

A Variant of Quick Mode with PFS (Stage 2)

6.3. Other Considerations

There are a number of features of IKE that deserve special consideration. They are discussed here.

6.3.1. Cookie Generation

As mentioned previously, IKE uses cookies as a partial defense against DoS attacks. When the responder receives Main Mode message 3 containing the Key Exchange data and the cookie, it verifies that the cookie is correct. However, this verification must not involve having a list of valid cookies. Otherwise, an attacker could potentially consume arbitrary amounts of memory by repeatedly requesting cookies from a responder. The recommended way to generate a cookie, as suggested by Phil Karn, is to have a single master key and compute a hash of the secret and the initiator's address information. This cookie can be verified by recomputing the cookie value based on information in the third message, and seeing if it matches.

6.3.2. Endpoint Identities

So far we have been rather vague about what kinds of endpoint identities are used. In principle, there are three ways a peer might be identified: by a shared key, a pre-configured public key, or a certificate.

6.3.2.1. Shared Key

In a shared key scheme, the peers share a symmetric key. This key is associated with a key identifier, which is known to both parties. It is assumed that the party verifying that identity also has a table that indicates for which traffic (i.e., what addresses) that identity is allowed to negotiate SAs.

6.3.2.2. Pre-Configured Public Key

A pre-configured public key scheme is the same as a shared key scheme except that the verifying party has the authenticating party's public key instead of a shared key.

6.3.2.3. Certificate

In a certificate scheme, the authenticating party presents a certificate containing their public key. It is straightforward to establish that this certificate matches the authentication data provided by the peer. What is less straightforward is to determine whether a given peer is entitled to negotiate for a given class of traffic. In theory, one might be able to determine this from the name in the certificate (e.g., the subject name contains an IP address that matches the ostensible IP address). In practice, this is not clearly specified in IKE and, therefore, is not really interoperable. Currently, it is likely that a configuration table maps certificates to policies, as in the other two authentication schemes.

7. Security Considerations

This document does not define any protocols and therefore has no security issues.

A. Appendix: IAB Members at the Time of This Writing

Bernard Aboba
Harald Alvestrand
Rob Austein
Leslie Daigle
Patrik Falstrom
Sally Floyd
Jun-ichiro Itojun Hagino
Mark Handley
Bob Hinden
Geoff Huston
Eric Rescorla
Pete Resnick
Jonathan Rosenberg

Normative References

There are no normative references for this document.

Informative References

- [AH] Kent, S., and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.
- [CCID2] Floyd, S. and E. Kohler, "Profile for DCCP Congestion Control ID 2: TCP-like Congestion Control", Work in Progress, October 2003.
- [CCID3] Floyd, S., Kohler, E., and J. Padhye, "Profile for DCCP Congestion Control ID 3: TFRC Congestion Control", Work in Progress, February 2004.
- [DCCP] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", Work in Progress, November 2004.
- [ECN] Ramakrishnan, K. Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, September 2001.
- [ESP] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.
- [IKE] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.

- [IPSEC] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [KERBEROS] Kohl, J. and C. Neuman, "The Kerberos Network Authentication Service (V5)", RFC 1510, September 1993.
- [SDP] Handley, M. and V. Jacobson, "SDP: Session Description Protocol" RFC 2327, April 1998.
- [STUN] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP)", RFC 3489, March 2003.
- [SYNFLOOD] CERT Advisory CA-1996-21 TCP SYN Flooding and IP Spoofing Attacks <<http://www.cert.org/advisories/CA-1996-21.html>>, September 19, 1996.
- [UNSAF] Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", RFC 3424, November 2002.
- [WEBDAV] Goland, Y., Whitehead, E., Faizi, A., Carter, S., and D. Jensen, "HTTP Extensions for Distributed Authoring -- WEBDAV", RFC 2518, February 1999.

Authors' Addresses

Eric Rescorla
RTFM, Inc.
2064 Edgewood Drive
Palo Alto, CA 94303

Phone: (650)-320-8549
EMail: ekr@rtfm.com

Internet Architecture Board
IAB

EMail: iab@iab.org

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

