

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 28, 2009

P. Zimmermann
Zfone Project
A. Johnston, Ed.
Avaya
J. Callas
PGP Corporation
October 25, 2008

ZRTP: Media Path Key Agreement for Secure RTP
draft-zimmermann-avt-zrtp-10

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on April 28, 2009.

Abstract

This document defines ZRTP, a protocol for media path Diffie-Hellman exchange to agree on a session key and parameters for establishing Secure Real-time Transport Protocol (SRTP) sessions. The ZRTP protocol is media path keying because it is multiplexed on the same port as RTP and does not require support in the signaling protocol. ZRTP does not assume a Public Key Infrastructure (PKI) or require the complexity of certificates in end devices. For the media session, ZRTP provides confidentiality, protection against man-in-the-middle (MiTM) attacks, and, in cases where the signaling protocol provides

Internet-Draft

ZRTP

October 2008

end-to-end integrity protection, authentication. ZRTP can utilize a Session Description Protocol (SDP) attribute to provide discovery and authentication through the signaling channel. To provide best effort SRTP, ZRTP utilizes normal RTP/AVP profiles.

Table of Contents

1.	Introduction	5
2.	Terminology	6
3.	Overview	6
3.1.	Key Agreement Modes	7
3.1.1.	Diffie-Hellman Mode Overview	7
3.1.2.	Multistream Mode Overview	9
3.1.3.	Preshared Mode Overview	9
4.	Protocol Description	10
4.1.	Discovery	10
4.1.1.	Protocol Version Negotiation	11
4.2.	Commit Contention	13
4.3.	Matching Shared Secret Determination	13
4.3.1.	Responder Behavior	15
4.3.2.	Initiator Behavior	16
4.3.3.	Handling a Shared Secret Cache Mismatch	16
4.4.	DH and non-DH key agreements	18
4.4.1.	Diffie-Hellman Mode	18
4.4.1.1.	Hash Commitment in Diffie-Hellman Mode	18
4.4.1.2.	Responder Behavior in Diffie-Hellman Mode	19
4.4.1.3.	Initiator Behavior in Diffie-Hellman Mode	20
4.4.1.4.	Shared Secret Calculation for DH Mode	20
4.4.2.	Multistream Mode	22
4.4.2.1.	Commitment in Multistream Mode	22
4.4.2.2.	Shared Secret Calculation for Multistream Mode	23
4.4.3.	Preshared Mode	24
4.4.3.1.	Commitment in Preshared Mode	25
4.4.3.2.	Initiator Behavior in Preshared Mode	25
4.4.3.3.	Responder Behavior in Preshared Mode	25
4.4.3.4.	Shared Secret Calculation for Preshared Mode	26
4.5.	Key Generation	27
4.6.	Confirmation	28
4.6.1.	Updating the Cache of Shared Secrets	29
4.7.	Termination	30
4.7.1.	Termination via Error message	30
4.7.2.	Termination via GoClear message	30

4.7.2.1.	Key Destruction for GoClear message	32
4.7.3.	Key Destruction at Termination	32
4.8.	Random Number Generation	33
4.9.	ZID and Cache Operation	33
4.9.1.	Cacheless implementations	34

5.	ZRTP Messages	35
5.1.	ZRTP Message Formats	36
5.1.1.	Message Type Block	36
5.1.2.	Hash Type Block	37
5.1.2.1.	Implicit Hash and HMAC algorithm	38
5.1.3.	Cipher Type Block	38
5.1.4.	Auth Tag Block	39
5.1.5.	Key Agreement Type Block	39
5.1.6.	SAS Type Block	41
5.1.7.	Signature Type Block	41
5.2.	Hello message	41
5.3.	HelloACK message	43
5.4.	Commit message	44
5.5.	DHPart1 message	47
5.6.	DHPart2 message	49
5.7.	Confirm1 and Confirm2 messages	51
5.8.	Conf2ACK message	52
5.9.	Error message	53
5.10.	ErrorACK message	54
5.11.	GoClear message	55
5.12.	ClearACK message	55
5.13.	SASrelay message	56
5.14.	RelayACK message	58
6.	Retransmissions	59
7.	Short Authentication String	60
7.1.	SAS Verified Flag	61
7.2.	Signing the SAS	63
7.3.	Relaying the SAS through a PBX	63
7.3.1.	PBX Enrollment and the PBX Enrollment Flag	65
8.	Signaling Interactions	66
8.1.	Binding the media stream to the signaling layer via the Hello Hash	67
8.1.1.	Integrity-protected signaling enables integrity-protected DH exchange	69
8.2.	Deriving the SRTP secret (srtps) from the signaling layer	70

8.3.	Codec Selection for Secure Media	71
9.	False ZRTP Packet Rejection	71
10.	Intermediary ZRTP Devices	73
11.	The ZRTP Disclosure flag	75
11.1.	Guidelines on Proper Implementation of the Disclosure Flag	76
12.	RTP Header Extension Flag for ZRTP	77
13.	IANA Considerations	78
14.	Appendix – Media Security Requirements	78
15.	Security Considerations	80
15.1.	Self-healing Key Continuity Feature	83
16.	Acknowledgments	84

17.	References	85
17.1.	Normative References	85
17.2.	Informative References	86
	Authors' Addresses	89
	Intellectual Property and Copyright Statements	90

1. Introduction

ZRTP is a key agreement protocol which performs Diffie-Hellman key exchange during call setup in the media path, and is transported over the same port as the Real-time Transport Protocol (RTP) [[RFC3550](#)] media stream which has been established using a signaling protocol such as Session Initiation Protocol (SIP) [[RFC3261](#)]. This generates a shared secret which is then used to generate keys and salt for a Secure RTP (SRTP) [[RFC3711](#)] session. ZRTP borrows ideas from PGPfone [[pgpfone](#)]. A reference implementation of ZRTP is available as Zfone [[zfone](#)].

The ZRTP protocol has some nice cryptographic features lacking in many other approaches to media session encryption. Although it uses a public key algorithm, it does not rely on a public key infrastructure (PKI). In fact, it does not use persistent public keys at all. It uses ephemeral Diffie-Hellman (DH) with hash commitment, and allows the detection of man-in-the-middle (MiTM) attacks by displaying a short authentication string (SAS) for the users to read and verbally compare over the phone. It has Perfect Forward Secrecy, meaning the keys are destroyed at the end of the call, which precludes retroactively compromising the call by future

disclosures of key material. But even if the users are too lazy to bother with short authentication strings, we still get reasonable authentication against a MiTM attack, based on a form of key continuity. It does this by caching some key material to use in the next call, to be mixed in with the next call's DH shared secret, giving it key continuity properties analogous to SSH. All this is done without reliance on a PKI, key certification, trust models, certificate authorities, or key management complexity that bedevils the email encryption world. It also does not rely on SIP signaling for the key management, and in fact does not rely on any servers at all. It performs its key agreements and key management in a purely peer-to-peer manner over the RTP packet stream.

In cases where the short authentication string (SAS) cannot be verbally compared by two human users, the SAS can be authenticated by exchanging an optional signature over the SAS (described in [Section 7.2](#)).

ZRTP can be used and discovered without being declared or indicated in the signaling path. This provides a best effort SRTP capability. Also, this reduces the complexity of implementations and minimizes interdependency between the signaling and media layers. However, when ZRTP is indicated in the signaling via the `zrtp-hash` SDP attribute, ZRTP has additional useful properties. By sending a hash of the ZRTP Hello message in the signaling, ZRTP provides a useful binding between the signaling and media paths, which is explained in

[Section 8.1](#). When this is done through a signaling path that has end-to-end integrity protection, the DH exchange is automatically protected from a MiTM attack, which is explained in [Section 8.1.1](#).

[2.](#) Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [RFC 2119](#) and indicate requirement levels for compliant implementations [[RFC2119](#)].

[3.](#) Overview

This section provides a description of how ZRTP works. This description is non-normative in nature but is included to build understanding of the protocol.

ZRTP is negotiated the same way a conventional RTP session is negotiated in an offer/answer exchange using the standard AVP/RTP profile. The ZRTP protocol begins after two endpoints have utilized a signaling protocol such as SIP and are ready to exchange media. If ICE [[I-D.ietf-mmusic-ice](#)] is being used, ZRTP begins after ICE has completed its connectivity checks.

ZRTP is multiplexed on the same ports as RTP. It uses a unique header that makes it clearly differentiable from RTP or STUN.

In environments in which sending ZRTP packets to non-ZRTP endpoints might cause problems and signaling path discovery is not an option, ZRTP endpoints can include the RTP header extension flag for ZRTP in normal RTP packets sent at the start of a session as a probe to discover if the other endpoint supports ZRTP. If the flag is received from the other endpoint, ZRTP messages can then be exchanged.

A ZRTP endpoint initiates the exchange by sending a ZRTP Hello message to the other endpoint. The purpose of the Hello message is to confirm the endpoint supports the protocol and to see what algorithms the two ZRTP endpoints have in common.

The Hello message contains the SRTP configuration options, and the ZID. Each instance of ZRTP has a unique 96-bit random ZRTP ID or ZID that is generated once at installation time. ZIDs are discovered during the Hello message exchange. The received ZID is used to look up retained shared secrets from previous ZRTP sessions with the endpoint.

A response to a ZRTP Hello message is a ZRTP HelloACK message. The HelloACK message simply acknowledges receipt of the Hello. Since RTP commonly uses best effort UDP transport, ZRTP has retransmission timers in case of lost datagrams. There are two timers, both with exponential backoff mechanisms. One timer is used for retransmissions of Hello messages and the other is used for retransmissions of all other messages after receipt of a HelloACK.

If an integrity protected signaling channel is available, a hash of the Hello message can be sent. This allows rejection of false injected ZRTP Hello messages by an attacker.

Hello and other ZRTP messages also contain a hash image that is used to link the messages together. This allows rejection of false injected ZRTP messages during an exchange.

[3.1.](#) Key Agreement Modes

After both endpoints exchange Hello and HelloACK messages, the key agreement exchange can begin with the ZRTP Commit message. ZRTP supports a number of key agreement modes including both Diffie-Hellman and non-Diffie-Hellman modes as described in the following sections.

The Commit message may be sent immediately after both endpoints have completed the Hello/HelloAck discovery handshake. Or it may be deferred until later in the call, after the participants engage in some unencrypted conversation. The Commit message may be manually activated by a user interface element, such as a GO SECURE button, which becomes enabled after the Hello/HelloAck discovery phase. This emulates the user experience of a number of secure phones in the PSTN world [[comsec](#)]. However, it is expected that most simple ZRTP user agents will omit such buttons and proceed directly to secure mode by sending a Commit message immediately after the Hello/HelloAck handshake.

[3.1.1.](#) Diffie-Hellman Mode Overview

An example ZRTP call flow is shown in Figure 1 below. Note that the order of the Hello/HelloACK exchanges in F1/F2 and F3/F4 may be reversed. That is, either Alice or Bob might send the first Hello message. Note that the endpoint which sends the Commit message is considered the initiator of the ZRTP session and drives the key agreement exchange. The Diffie-Hellman public values are exchanged in the DHPart1 and DHPart2 messages. SRTP keys and salts are then calculated.

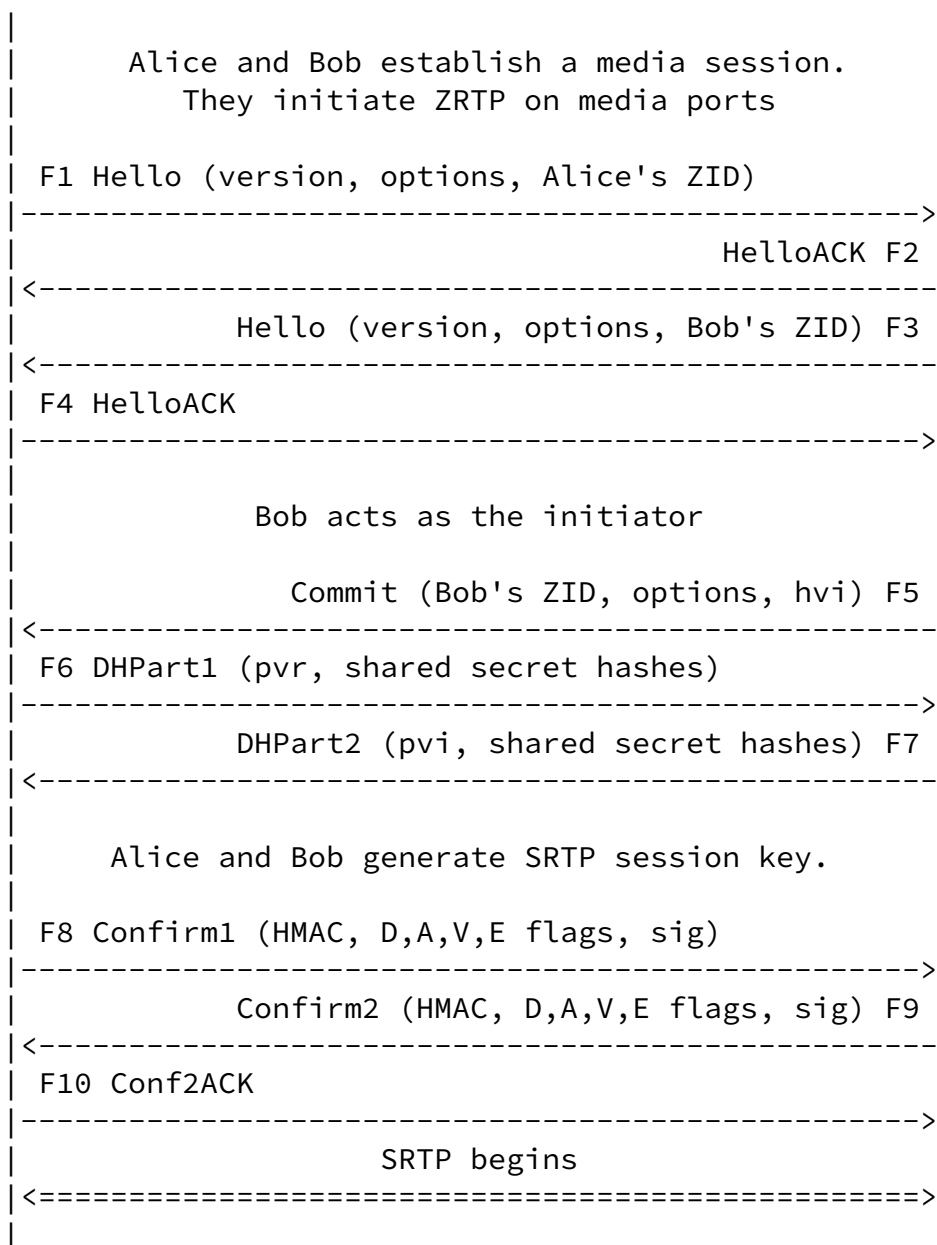


Figure 1: Establishment of an SRTP session using ZRTP

ZRTP authentication uses a Short Authentication String (SAS) which is ideally displayed for the human user. Alternatively, the SAS can be authenticated by exchanging an OPTIONAL digital signature (sig) over the short authentication string in the Confirm1 or Confirm2 messages (described in [Section 7.2](#)).

The ZRTP Confirm1 and Confirm2 messages are sent for a number of reasons, not the least of which is they confirm that all the key agreement calculations were successful and thus the encryption will work. They also carry other information such as the Disclosure flag

(D), the Allow Clear flag (A), the SAS Verified flag (V), and the PBX Enrollment flag (E). All flags are encrypted to shield them from a passive observer.

[3.1.2.](#) Multistream Mode Overview

Multistream mode is an alternative key agreement method when two endpoints have an established SRTP media stream between them and hence an active ZRTP Session key. ZRTP can derive multiple SRTP keys from a single DH exchange. For example, an established secure voice call that adds a video stream must use Multistream mode to quickly initiate the video stream without a second DH exchange.

When Multistream mode is indicated in the Commit message, a call flow similar to Figure 1 is used, but no DH calculation is performed by either endpoint and the DHPart1 and DHPart2 messages are omitted. The Confirm1, Confirm2, and Conf2ACK messages are still sent. Since the cache is not affected during this mode, multiple Multistream ZRTP exchanges can be performed in parallel between two endpoints.

When adding additional media streams to an existing call, only Multistream mode is used. Only one DH operation is performed, just for the first media stream.

[3.1.3.](#) Preshared Mode Overview

In the Preshared Mode, endpoints can skip the DH calculation if they have a shared secret from a previous ZRTP session. Preshared mode is indicated in the Commit message and results in the same call flow as Multistream mode. The principal difference between Multistream mode and Preshared mode is that Preshared mode uses a previously cached shared secret, *rs1*, instead of an active ZRTP Session key as the initial keying material.

This mode could be useful for slow processor endpoints so that a DH calculation does not need to be performed every session. Or, this mode could be used to rapidly re-establish an earlier session that was recently torn down or interrupted without the need to perform another DH calculation.

Preshared mode has forward secrecy properties. If a phone's cache is captured by an opponent, the cached shared secrets cannot be used to recover earlier encrypted calls, because the shared secrets are replaced with new ones in each new call, as in DH mode. However, the captured secrets can be used by a passive wiretapper in the media path to decrypt the next call, if the next call is in Preshared mode.

This differs from DH mode, which requires an active MiTM wiretapper to exploit captured secrets in the next call. However, if the next

call is missed by the wiretapper, he cannot wiretap any further calls. It thus preserves most of the self-healing properties ([Section 15.1](#)) of key continuity enjoyed by DH mode.

[4.](#) Protocol Description

This section begins the normative description of the protocol.

ZRTP MUST be multiplexed on the same ports as the RTP media packets.

To support best effort encryption from the Media Security Requirements [[I-D.ietf-sip-media-security-requirements](#)], ZRTP uses normal RTP/AVP profile (AVP) media lines in the initial offer/answer exchange. The ZRTP SDP attribute `a=zrtp-hash` defined in [Section 8](#) SHOULD be used in all offers and answers to indicate support for the ZRTP protocol. The Secure RTP/AVP (SAVP) profile MAY be used in subsequent offer/answer exchanges after a successful ZRTP exchange has resulted in an SRTP session, or if it is known the other endpoint supports this profile.

The use of the RTP/SAVP profile has caused failures in negotiating best effort SRTP due to the limitations on negotiating profiles using SDP. This is why ZRTP supports the RTP/AVP profile and includes its own discovery mechanisms.

In all key agreement modes, the initiator SHOULD NOT send RTP media after sending the Commit message, and MUST NOT send SRTP media before receiving either the Conf2ACK or the first SRTP media (with a valid SRTP auth tag) from the responder. The responder SHOULD NOT send RTP media after receiving the Commit message, and MUST NOT send SRTP media before receiving the Confirm2 message.

[4.1.](#) Discovery

During the ZRTP discovery phase, a ZRTP endpoint discovers if the other endpoint supports ZRTP and the supported algorithms and options. This information is transported in a Hello message, described in [Section 5.2](#).

ZRTP endpoints SHOULD include the SDP attribute `a=zrtp-hash` in offers and answers, as defined in [Section 8](#). ZRTP MAY use an RTP [[RFC3550](#)] extension field as a flag to indicate support for the ZRTP protocol in RTP packets as described in [Section 12](#).

The Hello message includes the ZRTP version, hash type, cipher type, authentication method and tag length, key agreement type, and Short Authentication String (SAS) algorithms that are supported. The Hello

message also includes a hash image as described in [Section 9](#). In addition, each endpoint sends and discovers ZIDs. The received ZID is used later in the protocol as an index into a cache of shared secrets that were previously negotiated and retained between the two parties.

A Hello message can be sent at any time, but is usually sent at the start of an RTP session to determine if the other endpoint supports ZRTP, and also if the SRTP implementations are compatible. A Hello message is retransmitted using timer T1 and an exponential backoff mechanism detailed in [Section 6](#) until the receipt of a HelloACK message or a Commit message.

The use of the `a=zrtp-hash` SDP attribute to authenticate the Hello message is described in [Section 8.1](#).

[4.1.1](#). Protocol Version Negotiation

This specification defines ZRTP version 1.00. Since new versions of ZRTP may be developed in the future, this specification defines a protocol version negotiation in this section.

Each party declares what version of the ZRTP protocol they support via the version field in the Hello message ([Section 5.2](#)). If both parties have the same version number in their Hello messages, they can proceed with the rest of the protocol. To facilitate both parties reaching this state of protocol version agreement in their Hello messages, ZRTP should use information provided in the signaling layer, if available. If a ZRTP endpoint supports more than one version of the protocol, it SHOULD declare them all in a list of SIP SDP `a=zrtp-hash` attributes (defined in [Section 8](#)), listing separate hashes, with separate ZRTP version numbers in each item in the list.

Both parties should inspect the list of ZRTP version numbers supplied by the other party in the SIP SDP a=zrtp-hash attributes. Both parties should choose the highest version number that appear in both parties' list of a=zrtp-hash version numbers, and use that version for their Hello messages. If both parties use the SIP signaling in this manner, their initial Hello messages will have the same ZRTP version number, provided they both have at least one supported protocol version in common. Before the ZRTP key agreement can proceed, an endpoint MUST have sent and received Hellos with the same protocol version.

It is best if the signaling layer is used to negotiate the protocol version number. However, the a=zrtp-hash SDP attribute is not always present in the SIP packet, as explained in [Section 8.1](#). In the absence of any guidance from the signaling layer, an endpoint MUST

send the highest supported version in initial Hello messages. If the two parties send different protocol version numbers in their Hello messages, they can reach agreement to use a common version, if one exists. They iteratively apply the following rules until they both have matching version fields in their Hello messages and the key agreement can proceed:

- o If an endpoint receives a Hello message with an unsupported version number that is higher than the endpoint's current Hello message version, the received Hello message MUST be ignored. The endpoint continues to retransmit Hello messages on the standard retry schedule ([Section 6](#)).
- o If an endpoint receives a Hello message with a version number that is lower than the endpoint's current Hello message, and the endpoint supports a version that is less than or equal to the received version number, the endpoint MUST stop retransmitting the old version number and MUST start sending a new Hello message with the highest supported version number that is less than or equal to the received version number.
- o If an endpoint receives a Hello message with an unsupported version number that is lower than the endpoint's current Hello message, the endpoint MUST send an Error message ([Section 5.9](#)) indicating failure to support this ZRTP version.

The above comparisons are iterated until the version numbers match,

or until it exits on a failure to match.

For example, assume that Alice supports protocol version 1.00 and 2.00, and Bob supports version 1.00 and 1.10. Alice initially sends a Hello with version 2.00, and Bob initially sends a Hello with version 1.10. Bob ignores Alice's 2.00 Hello and continues to send his 1.10 Hello. Alice detects that Bob does not support 2.00 and she stops sending her 2.00 Hellos and starts sending a stream of 1.00 Hellos. Bob sees the 1.00 Hello from Alice and stops sending his 1.10 Hellos and switches to sending 1.00 Hellos. At that point, they have converged on using version 1.00 and the protocol proceeds on that basis.

When comparing protocol versions, a ZRTP endpoint MUST include only the first three octets of the version field in the comparison. The final octet is ignored, because it is not significant for interoperability. For example, "1.0 ", "1.00", "1.01", or "1.0a" are all regarded as a version match, because they would all be interoperable versions.

Changes in protocol version numbers are expected to be infrequent after version 1.00. Supporting multiple versions adds code complexity and may introduce security weaknesses in the implementation. The old

adage about keeping it simple applies especially to implementing security protocols. Endpoints SHOULD NOT support protocol versions earlier than version 1.00.

[4.2.](#) Commit Contention

After both parties have received compatible Hello messages, a Commit message ([Section 5.4](#)) can be sent to begin the ZRTP key exchange. The endpoint that sends the Commit is known as the initiator, while the receiver of the Commit is known as the responder.

If both sides send Commit messages initiating a secure session at the same time the following rules are used to break the tie:

- o If one Commit is for a DH mode while the other is for Preshared mode, then the Preshared Commit MUST be discarded and the DH Commit proceeds.
- o If the two Commits are both Preshared mode, and one party has set

the MiTM (M) flag in the Hello message and the other has not, the Commit message from the party who set the (M) flag MUST be discarded, and the one who has not set the (M) flag becomes the initiator, regardless of the nonce values. In other words, for Preshared mode, the phone is the initiator and the PBX is the responder.

- o If the two Commits are either both DH modes or both non-DH modes, then the Commit message with the lowest hvi value (for DH Commits), or lowest nonce value (for non-DH Commits), MUST be discarded and the other side is the initiator, and the protocol proceeds with the initiator's Commit. The two hvi or nonce values are compared as large unsigned integers in network byte order.

If one Commit is for Multistream mode while the other is for non-Multistream (DH or Preshared) mode, a software error has occurred and the ZRTP negotiation should be terminated. This should never occur because of the constraints on Multistream mode described in [Section 4.4.2](#).

In the event that Commit messages are sent by both ZRTP endpoints at the same time, but are received in different media streams, the same resolution rules apply as if they were received on the same stream. The media stream in which the Commit will proceed through the ZRTP exchange while the media stream with the discarded Commit must wait for the completion of the other ZRTP exchange.

[4.3](#). Matching Shared Secret Determination

The following sections describe how ZRTP endpoints generate and/or use the set of shared secrets s1, auxsecret, and pbxsecret through

the exchange of the DHPart1 and DHPart2 messages. This doesn't cover the Diffie-Hellman calculations. It only covers the method whereby the two parties determine if they already have shared secrets in common in their caches.

Each ZRTP endpoint maintains a long-term cache of shared secrets that it has previously negotiated with the other party. The ZID of the other party, received in the other party's Hello message, is used as an index into this cache to find the set of shared secrets, if any exist. This cache entry may contain previously retained shared secrets, rs1 and rs2, which give ZRTP its key continuity features.

If the other party is a PBX, the cache may also contain a trusted MiTM PBX shared secret, called pbxsecret, defined in [Section 7.3.1](#).

The DHPart1 and DHPart2 messages contain a list of hashes of these shared secrets to allow the two endpoints to compare the hashes with what they have in their caches to detect whether the two sides share any secrets that can be used in the calculation of the session key. The use of this shared secret cache is described in [Section 4.9](#).

If no secret of a given type is available, a random value is generated and used for that secret to ensure a mismatch in the hash comparisons in the DHPart1 and DHPart2 messages. This prevents an eavesdropper from knowing which types of shared secrets are available between the endpoints.

[Section 4.3.1](#) and [Section 4.3.2](#) both refer to the auxiliary shared secret auxsecret. The auxsecret shared secret may be defined by the VoIP user agent out-of-band from the ZRTP protocol. In some cases it may be provided by the signaling layer as srtps, which is defined in [Section 8.2](#). If it is not provided by the signaling layer, the auxsecret shared secret may be manually provisioned in other application-specific ways that are out-of-band, such as computed from a hashed pass phrase by prior agreement between the two parties. Or it may be a family key used by an institution that the two parties both belong to. It is a generalized mechanism for providing a shared secret that is agreed to between the two parties out of scope of the ZRTP protocol. It is expected that most typical ZRTP endpoints will rarely use auxsecret.

For both the initiator and the responder, the shared secrets s1, s2, and s3 will be calculated so that they can all be used later to calculate s0 in [Section 4.4.1.4](#). Here is how s1, s2, and s3 are calculated by both parties:

The shared secret s1 will be either the initiator's rs1 or the initiator's rs2, depending on which of them can be found in the responder's cache. If the initiator's rs1 matches the responder's

rs1 or rs2, then s1 MUST be set to the initiator's rs1. If and only if that match fails, then if the initiator's rs2 matches the responder's rs1 or rs2, then s1 MUST be set to the initiator's rs2. If that match also fails, then s1 MUST be set to null. The

complexity of the s1 calculation is to recover from any loss of cache sync from an earlier aborted session, due to the Byzantine Generals' Problem [[Byzantine](#)].

The shared secret s2 MUST be set to the value of auxsecret if and only if both parties have matching values for auxsecret, as determined by comparing the hashes of auxsecret sent in the DH messages. If they don't match, s2 MUST be set to null.

The shared secret s3 MUST be set to the value of pbxsecret if and only if both parties have matching values for pbxsecret, as determined by comparing the hashes of pbxsecret sent in the DH messages. If they don't match, s3 MUST be set to null.

If s1, s2, or s3 have null values, they are assumed to have a zero length for the purposes of hashing them later during the s0 calculation in [Section 4.4.1.4](#).

The comparison of hashes of rs1, rs2, auxsecret, and pbxsecret is described in the next sections.

[4.3.1](#). Responder Behavior

The responder calculates an HMAC keyed hash using the first retained shared secret, rs1, as the key on the string "Responder" which generates a retained secret ID, rs1IDr, which is truncated to the leftmost 64 bits. HMACs are calculated in a similar way for additional shared secrets:

```
rs1IDr = HMAC(rs1, "Responder")
rs2IDr = HMAC(rs2, "Responder")
auxsecretIDr = HMAC(auxsecret, "Responder")
pbxsecretIDr = HMAC(pbxsecret, "Responder")
```

The set of keyed hashes (HMACs) of shared secrets are included by the responder in the DHPart1 message.

The HMACs of the possible shared secrets received in the DHPart2 can be compared against the HMACs of the local set of possible shared secrets. From these comparisons, s1, s2, and s3 are calculated per the methods described above in [Section 4.3](#). The expected HMAC values of the shared secrets are calculated (using the string "Initiator" instead of "Responder") as in [Section 4.3.2](#) and compared to the HMACs received in the DHPart2 message. The secrets corresponding to

matching HMACs are kept while the secrets corresponding to the non-matching ones are replaced with a null, which is assumed to have a zero length for the purposes of hashing them later. The resulting s1, s2, and s3 values are used later to calculate s0 in [Section 4.4.1.4](#).

[4.3.2](#). Initiator Behavior

The initiator calculates an HMAC keyed hash using the first retained shared secret, rs1, as the key on the string "Initiator" which generates a retained secret ID, rs1IDi, which is truncated to the leftmost 64 bits. HMACs are calculated in a similar way for additional shared secrets:

```
rs1IDi = HMAC(rs1, "Initiator")
rs2IDi = HMAC(rs2, "Initiator")
auxsecretIDi = HMAC(auxsecret, "Initiator")
pbxsecretIDi = HMAC(pbxsecret, "Initiator")
```

These HMACs of shared secrets are included by the initiator in the DHPart2 message.

The initiator then calculates the set of secret IDs that are expected to be received from the responder in the DHPart1 message by substituting the string "Responder" instead of "Initiator" as in [Section 4.3.1](#).

The HMACs of the possible shared secrets received are compared against the HMACs of the local set of possible shared secrets. From these comparisons, s1, s2, and s3 are calculated per the methods described above in [Section 4.3](#). The secrets corresponding to matching HMACs are kept while the secrets corresponding to the non-matching ones are replaced with a null, which is assumed to have a zero length for the purposes of hashing them later. The resulting s1, s2, and s3 values are used later to calculate s0 in [Section 4.4.1.4](#).

For example, consider two ZRTP endpoints who share secrets rs1 and pbxsecret (defined in [Section 7.3.1](#)). During the comparison, rs1ID and pbxsecretID will match but auxsecretID will not. As a result, s1 = rs1, s2 will be null, and s3 = pbxsecret.

[4.3.3](#). Handling a Shared Secret Cache Mismatch

A shared secret cache mismatch is defined to mean that we expected a cache match because rs1 exists in our local cache, but we computed a null value for s1 (per the method described in [Section 4.3](#)).

Internet-Draft

ZRTP

October 2008

If one party has a cached shared secret and the other party does not, this indicates one of two possible situations. Either there is a man-in-the-middle (MiTM) attack, or one of the legitimate parties has lost their cached shared secret by some mishap. Perhaps they inadvertently deleted their cache, or their cache was lost or disrupted due to restoring their disk from an earlier backup copy. The party that has the surviving cache entry can easily detect that a cache mismatch has occurred, because they expect their own cached secret to match the other party's cached secret, but it does not match. It is possible for both parties to detect this condition if both parties have surviving cached secrets that have fallen out of sync, due perhaps to one party restoring from a disk backup.

If either party discovers a cache mismatch, the user agent who makes this discovery must treat this as a possible security event and **MUST** alert their own user that there is a heightened risk of a MiTM attack, and that the user should verbally compare the SAS with the other party to ascertain that no MiTM attack has occurred. If a cache mismatch is detected and it is not possible to compare the SAS, either because the user interface does not support it or because one or both endpoints are unmanned devices, and no other SAS comparison mechanism is available, the session **MAY** be terminated.

The session need not be terminated on a cache mismatch event if the mechanism described in [Section 8.1.1](#) is available, which allows authentication of the DH exchange without human assistance. Or if any mechanism is available to determine if the SAS matches. This would require either circumstances that allow human verbal comparisons of the SAS, or by using the **OPTIONAL** digital signature feature on the SAS hash, as described in [Section 7.2](#). Even if the user interface does not permit an SAS comparison, the human user **MUST** be warned, and may elect to proceed with the call at their own risk.

Here is a non-normative example of a cache-mismatch alert message from a ZRTP user agent (specifically, Zfone [[zfone](#)]), designed for a desktop PC graphical user interface environment. It is by no means required that the alert be this detailed:

"We expected the other party to have a shared secret cached from a previous call, but they don't have it. This may mean your partner

simply lost his cache of shared secrets, but it could also mean someone is trying to wiretap you. To resolve this question you must check the authentication string with your partner. If it doesn't match, it indicates the presence of a wiretapper." If the alert is rendered by a robot voice instead of a GUI, brevity may be more important: "Something's wrong. You must check the authentication string with your partner. If it doesn't match, it indicates the presence of a wiretapper."

[4.4.](#) DH and non-DH key agreements

The next step is the generation of a secret for deriving SRTP keying material. ZRTP uses Diffie-Hellman and two non-Diffie-Hellman modes, described in the following sections.

[4.4.1.](#) Diffie-Hellman Mode

The purpose of the Diffie-Hellman (either Finite Field Diffie-Hellman or Elliptic Curve Diffie-Hellman) exchange is for the two ZRTP endpoints to generate a new shared secret, s_0 . In addition, the endpoints discover if they have any cached or previously stored shared secrets in common, and uses them as part of the calculation of the session keys.

Because the DH exchange affects the state of the retained shared secret cache, only one in-process ZRTP DH exchange may occur at a time between two ZRTP endpoints. Otherwise, race conditions and cache integrity problems will result. When multiple media streams are established in parallel between the same pair of ZRTP endpoints (determined by the ZIDs in the Hello Messages), only one can be processed. Once that exchange completes with Confirm2 and Conf2ACK messages, another ZRTP DH exchange can begin. This constraint does not apply when Multistream mode key agreement is used since the cached shared secrets are not affected.

[4.4.1.1.](#) Hash Commitment in Diffie-Hellman Mode

From the intersection of the algorithms in the sent and received Hello messages, the initiator chooses a hash, cipher, auth tag, key agreement type, and SAS type to be used.

A Diffie-Hellman mode is selected by setting the Key Agreement Type

to one of the DH or ECDH values in Table 5 in the Commit. In this mode, the key agreement begins with the initiator choosing a fresh random Diffie-Hellman (DH) secret value (svi) based on the chosen key agreement type value, and computing the public value. (Note that to speed up processing, this computation can be done in advance.) For guidance on generating random numbers, see [Section 4.8](#). The value for the DH generator g, the DH prime p, and the length of the DH secret value, svi, are defined in [Section 5.1.5](#).

$$pvi = g^{svi} \bmod p$$

where g and p are determined by the key agreement type value. The pvi value is formatted as a big-endian octet string, fixed to the width of the DH prime, and leading zeros MUST NOT be truncated.

The hash commitment is performed by the initiator of the ZRTP exchange. The hash value of the initiator, hvi, includes a hash of the entire DHPart2 message as shown in Figure 9 (which includes the Diffie-Hellman public value, pvi), and the responder's Hello message:

$$hvi = \text{hash}(\text{initiator's DHPart2 message} \mid \text{responder's Hello message})$$

Note that the Hello message includes the fields shown in Figure 3.

The information from the responder's Hello message is included in the hash calculation to prevent a bid-down attack by modification of the responder's Hello message.

The initiator sends hvi in the Commit message.

The use of hash commitment in the DH exchange constrains the attacker to only one guess to generate the correct short authentication string (SAS) ([Section 7](#)) in his attack, which means the SAS can be quite short. A 16-bit SAS, for example, provides the attacker only one chance out of 65536 of not being detected.

[4.4.1.2](#). Responder Behavior in Diffie-Hellman Mode

Upon receipt of the Commit message, the responder generates its own fresh random DH secret value, svr, and computes the public value.

(Note that to speed up processing, this computation can be done in advance.) For guidance on random number generation, see [Section 4.8](#). The value for the DH generator g , the DH prime p , and the length of the DH secret value, svr , are defined in [Section 5.1.5](#).

$$pvr = g^{svr} \bmod p$$

The pvr value is formatted as a big-endian octet string, fixed to the width of the DH prime, and leading zeros MUST NOT be truncated.

Upon receipt of the DHPart2 message, the responder checks that the initiator's public DH value is not equal to 1 or $p-1$. An attacker might inject a false DHPart2 packet with a value of 1 or $p-1$ for $g^{svi} \bmod p$, which would cause a disastrously weak final DH result to be computed. If pvi is 1 or $p-1$, the user should be alerted of the attack and the protocol exchange MUST be terminated. Otherwise, the responder computes its own value for the hash commitment using the public DH value (pvi) received in the DHPart2 packet and its Hello packet and compares the result with the hvi received in the Commit packet. If they are different, a MiTM attack is taking place and the user is alerted and the protocol exchange terminated.

The responder then calculates the Diffie-Hellman result:

$$DHResult = pvi^{svr} \bmod p$$

[4.4.1.3](#). Initiator Behavior in Diffie-Hellman Mode

Upon receipt of the DHPart1 message, the initiator checks that the responder's public DH value is not equal to 1 or $p-1$. An attacker might inject a false DHPart1 packet with a value of 1 or $p-1$ for $g^{svr} \bmod p$, which would cause a disastrously weak final DH result to be computed. If pvr is 1 or $p-1$, the user should be alerted of the attack and the protocol exchange MUST be terminated.

The initiator then sends a DHPart2 message containing the initiator's public DH value and the set of calculated shared secret IDs as defined in [Section 4.3.2](#).

The initiator calculates the same Diffie-Hellman result using:

$DHResult = pvr^{svi} \bmod p$

[4.4.1.4](#). Shared Secret Calculation for DH Mode

A hash of the received and sent ZRTP messages in the current ZRTP exchange in the following order is calculated by both parties:

$total_hash = \text{hash}(\text{Hello of responder} \mid \text{Commit} \mid \text{DHPart1} \mid \text{DHPart2})$

Note that only the ZRTP messages (Figure 3, Figure 5, Figure 8, and Figure 9), not the entire ZRTP packets, are included in the `total_hash`.

For both the initiator and responder, the `DHResult` is formatted as a big-endian octet string, fixed to the width of the DH prime, and leading zeros MUST NOT be truncated. For example, for a 3072-bit `p`, `DHResult` would be a 384 octet value, with the first octet the most significant.

The calculation of the final shared secret, `s0`, is in compliance with the recommendations in sections [5.8.1](#) and [6.1.2.1](#) of NIST SP 800-56A [SP800-56A]. This is done by hashing a concatenation of a number of items, including the `DHResult`, the ZID's of the initiator (`ZIDi`) and the responder (`ZIDr`), the `total_hash`, and the set of non-null shared secrets as described in Section 4.3.

In [section 5.8.1](#) of NIST SP 800-56A [SP800-56A], NIST requires certain parameters to be hashed together in a particular order, which NIST refers to as: Z, AlgorithmID, PartyUInfo, PartyVInfo,

`SuppPubInfo`, and `SuppPrivInfo`. In our implementation, our `DHResult` corresponds to Z, "ZRTP-HMAC-KDF" corresponds to AlgorithmID, our `ZIDi` and `ZIDr` correspond to PartyUInfo and PartyVInfo, our `total_hash` corresponds to `SuppPubInfo`, and the set of three shared secrets `s1`, `s2`, and `s3` corresponds to `SuppPrivInfo`. NIST also requires a 32-bit big-endian integer counter to be included in the hash each time the hash is computed, which we have set to the fixed value of 1, because we only compute the hash once.

$s0 = \text{hash}(\text{counter} \mid \text{DHResult} \mid \text{"ZRTP-HMAC-KDF"} \mid \text{ZIDi} \mid \text{ZIDr} \mid \text{total_hash} \mid \text{len}(s1) \mid s1 \mid \text{len}(s2) \mid s2 \mid \text{len}(s3) \mid s3)$

Note that temporary values s1, s2, and s3 were calculated per the methods described above in [Section 4.3](#), and they are erased from memory immediately after they are used to calculate s0.

The length of the DHResult field was implicitly agreed to by the negotiated DH prime size. The length of total_hash is implicitly determined by the negotiated hash algorithm. All of the explicit length fields, len(), in the above hash are 32-bit big-endian integers, giving the length in octets of the field that follows. Some members of the set of shared secrets (s1, s2, and s3) may have lengths of zero if they are null (not shared), and are each preceded by a 4-octet length field. For example, if s2 is null, len(s2) is 0x00000000, and s2 itself would be absent from the hash calculation, which means len(s3) would immediately follow len(s2). While inclusion of ZIDi and ZIDr may be redundant, because they are implicitly included in the total_hash, we explicitly include them here to follow NIST SP800-56A. The string "ZRTP-HMAC-KDF" (not null-terminated) identifies what purpose the resulting s0 will be used for, which is to serve as the master key for the ZRTP HMAC-based key derivation function defined in [Section 4.5](#).

A ZRTP Session Key is generated which then allows the ZRTP Multistream mode to be used to generate SRTP key and salt pairs for additional concurrent media streams between this pair of ZRTP endpoints. If a ZRTP Session Key has already been generated between this pair of endpoints and is available, no new ZRTP Session Key is calculated.

ZRTPSess = HMAC(s0,"ZRTP Session Key")

The ZRTPSess key is kept for the duration of the call signaling session between the two ZRTP endpoints. That is, if there are two separate calls between the endpoints (in SIP terms, separate SIP dialogs), then a ZRTP Session Key MUST NOT be used across the two call signaling sessions. ZRTPSess MUST be destroyed no later than the end of the call signaling session.

The two endpoints proceed with key generation as described in [Section 4.5](#), now that there is a defined s0 and ZRTPSess key.

[4.4.2](#). Multistream Mode

The Multistream key agreement mode can be used to generate SRTP keys and salts for additional media streams established between a pair of endpoints. Multistream mode cannot be used unless there is an active SRTP session established between the endpoints which means a ZRTP Session key is active. This ZRTP Session key can be used to generate keys and salts without performing another DH calculation. In this mode, the retained shared secret cache is not used or updated. As a result, multiple ZRTP Multistream mode exchanges can be processed in parallel between two endpoints.

Multistream mode is also used to resume a secure call that has gone clear using a GoClear message as described in [Section 4.7.2.1](#).

When adding additional media streams to an existing call, Multistream mode MUST be used. The first media stream MUST use either DH mode or Preshared mode. Only one DH exchange or Preshared exchange is performed, just for the first media stream. The DH exchange or Preshared exchange MUST be completed for the first media stream before Multistream mode is used to add any other media streams.

[4.4.2.1](#). Commitment in Multistream Mode

Multistream mode is selected by the initiator setting the Key Agreement Type to "Mult" in the Commit message (Figure 6). The Cipher Type, Auth Tag Length, and Hash in Multistream mode SHOULD be set by the initiator to the same as the values as in the initial DH Mode Commit. The SAS Type is ignored as there is no SAS authentication in this mode.

Note: This requirement is needed since some endpoints cannot support different SRTP algorithms for different media streams. However, in the case of Multistream mode being used to go secure after a GoClear, the requirement to use the same SRTP algorithms is relaxed if there are no other active SRTP sessions.

In place of hvi in the Commit, a random nonce of length 4-words (16 octets) is chosen. Its value MUST be unique for all nonce values chosen for active ZRTP sessions between a pair of endpoints. If a Commit is received with a reused nonce value, the ZRTP exchange MUST be immediately terminated.

Note: Since the nonce is used to calculate different SRTP key and salt pairs for each media stream, a duplication will result in the same key and salt being generated for the two media streams, which would have disastrous security consequences.

If a Commit is received selecting Multistream mode, but the responder does not have a ZRTP Session Key available, the exchange MUST be terminated. Otherwise, the responder proceeds to the next section on Shared Secret Calculation, [Section 4.4.2.2](#).

If both sides send Multistream Commit messages at the same time, the contention is resolved and the initiator/responder roles are settled according to [Section 4.2](#), and the protocol proceeds.

In Multistream mode, both the DHPart1 and DHPart2 messages are skipped. After receiving the Commit message from the initiator, the responder sends the Confirm1 message after calculating this stream's SRTP keys, as described below.

[4.4.2.2](#). Shared Secret Calculation for Multistream Mode

A hash of the received and sent ZRTP messages in the current ZRTP exchange for the current media stream is calculated:

```
total_hash = hash(Hello of responder | Commit )
```

This refers to the Hello and Commit messages for the current media stream which is using Multistream mode, not the original media stream that included a full DH key agreement. Note that only the ZRTP messages (Figure 3 and Figure 6), not the entire ZRTP packets, are included in the hash.

The SRTP keys and salts for the initiator and responder are calculated using the ZRTP Session Key ZRTPSess and the nonce from the Commit message. The nonce from the Commit message is implicitly included in the total_hash, which hashed the entire Commit message and the other party's Hello message. For the nth media stream:

```
s0n = HMAC(ZRTPSess, total_hash)
```

Note that the responder's Hello message, included in the total_hash, includes some unique nonce-derived material of its own (the H3 hash image), thereby ensuring that each of the two parties can unilaterally force the resulting s0n shared secret to be unique for each media stream, even if one party by some error fails to produce a unique nonce. Note also that the ZRTPSess key is derived from material that also includes a different and more inclusive total_hash from the entire packet sequence that performed the original DH

Internet-Draft

ZRTP

October 2008

exchange for the first media stream in this ZRTP session.

At this point in Multistream mode, the two endpoints begin key generation as described in [Section 4.5](#) using s_0n in place of s_0 in the key generation formulas for this media stream.

[4.4.3](#). Preshared Mode

The Preshared key agreement mode can be used to generate SRTP keys and salts without a DH calculation, instead relying on a shared secret from previous DH calculations between the endpoints.

This key agreement mode is useful to rapidly re-establish a secure session between two parties who have recently started and ended a secure session that has already performed a DH key agreement, without performing another lengthy DH calculation, which may be desirable on slow processors in resource-limited environments. Preshared mode **MUST NOT** be used for adding additional media streams to an existing call. Multistream mode **MUST** be used for this purpose.

In the most severe resource-limited environments, Preshared mode may be useful with processors that cannot perform a DH calculation in an ergonomically acceptable time limit. Shared key material may be manually provisioned between two such endpoints in advance and still allow a limited subset of functionality. Such a "better than nothing" implementation would have to be regarded as non-compliant with the ZRTP specification, but it could interoperate in Preshared (and if applicable, Multistream) mode with a compliant ZRTP endpoint.

Because Preshared mode affects the state of the retained shared secret cache, only one in-process ZRTP Preshared exchange may occur at a time between two ZRTP endpoints. This rule is explained in more detail in [Section 4.4.1](#), and applies for the same reasons as in DH mode.

Preshared mode **MUST NOT** be used for establishing a second media stream. Multistream mode is designed for that.

Preshared mode is only included in this specification to meet the R-REUSE requirement in the Media Security Requirements [\[I-D.ietf-sip-media-security-requirements\]](#) document. A series of preshared-keyed calls between two ZRTP endpoints should use a DH key

exchange periodically. Preshared mode is only used if a cached shared secret has been established in an earlier session by a DH exchange, as discussed in [Section 4.9](#).

[4.4.3.1](#). Commitment in Preshared Mode

Preshared mode is selected by setting the Key Agreement Type to Preshared in the Commit message. This results in the same call flow as Multistream mode. The principal difference between Multistream mode and Preshared mode is that Preshared mode uses a previously cached shared secret, `rs1`, instead of an active ZRTP Session key, `ZRTPSess`, as the initial keying material.

Because Preshared mode depends on having a reliable shared secret in its cache, it is RECOMMENDED that Preshared mode only be used when the SAS Verified flag has been previously set.

[4.4.3.2](#). Initiator Behavior in Preshared Mode

The Commit message (Figure 7) is sent by the initiator of the ZRTP exchange. From the intersection of the algorithms in the sent and received Hello messages, the initiator chooses a hash, cipher, auth tag, key agreement type, and SAS type to be used.

To assemble a Preshared commit, we must first construct a temporary `preshared_key`, which is constructed from one of several possible combinations of cached key material, depending on what is available in the shared secret cache. If `rs1` is not available in the initiator's cache, then Preshared mode MUST NOT be used.

```
preshared_key = hash( len(rs1) | rs1 | len(auxsecret) | auxsecret
| len(pbxsecret) | pbxsecret )
```

All of the explicit length fields, `len()`, in the above hash are 32-bit big-endian integers, giving the length in octets of the field that follows. Some members of the set of shared secrets (`rs1`, `auxsecret`, and `pbxsecret`) may have lengths of zero if they are null (not available), and are each preceded by a 4-octet length field. For example, if `auxsecret` is null, `len(auxsecret)` is `0x00000000`, and

auxsecret itself would be absent from the hash calculation, which means len(pbxsecret) would immediately follow len(auxsecret).

In place of hvi in the Commit message, two smaller fields are inserted by the initiator:

- A random nonce of length 4-words (16 octets).
- A keyID = HMAC(preshared_key, "Prsh") truncated to 64 bits.

[4.4.3.3](#). Responder Behavior in Preshared Mode

The responder uses the received keyID to search for matching key material in its cache. It does this by computing a preshared_key

value and keyID value using the same formula as the initiator, depending on what is available in the responder's local cache. If the locally computed keyID does not match the received keyID in the Commit, the responder recomputes a new preshared_key and keyID from a different subset of shared keys from the cache, dropping auxsecret or pbxsecret or both from the hash calculation, until a matching preshared_key is found or it runs out of possibilities. Note that rs2 is not included in the process.

If it finds the appropriate matching shared key material, it is used to derive s0 and a new ZRTPSess key, as described in the next section on Shared Secret Calculation, [Section 4.4.3.4](#).

If the responder determines that it does not have a cached shared secret from a previous DH exchange, or it fails to match the keyID hash from the initiator with any combination of its shared keys, it SHOULD respond with its own DH Commit message. This would reverse the roles and the responder would become the initiator, because the DH Commit must always "trump" the Preshared Commit message as described in [Section 4.2](#). The key exchange would then proceed using DH mode. However, if a severely resource-limited responder lacks the computing resources to respond in a reasonable time with a DH Commit, it MAY respond with a ZRTP Error message ([Section 5.9](#)) indicating that no shared secret is available.

If both sides send Preshared Commit messages initiating a secure session at the same time, the contention is resolved and the initiator/responder roles are settled according to [Section 4.2](#), and

the protocol proceeds.

In Preshared mode, both the DHPart1 and DHPart2 messages are skipped. After receiving the Commit message from the initiator, the responder sends the Confirm1 message after calculating this stream's SRTP keys, as described below.

[4.4.3.4](#). Shared Secret Calculation for Preshared Mode

A hash of the received and sent ZRTP messages in the current ZRTP exchange for the current media stream is calculated:

```
total_hash = hash(Hello of responder | Commit )
```

Note that only the ZRTP messages (Figure 3 and Figure 7), not the entire ZRTP packets, are included in the hash. The nonce from the Commit message is implicitly included in the total_hash, which hashed the entire Commit message and the other party's Hello message. Next, the preshared_key is used to derive s0 and ZRTPSess:

```
s0 = HMAC(preshared_key, total_hash)
ZRTPSess = HMAC(s0,"ZRTP Session Key")
```

The preshared_key is erased as soon as it has been used to calculate s0 and ZRTPSess. The ZRTPSess key allows the later use of Multistream mode for adding additional media streams to this session.

Note that the responder's Hello message, included in the total_hash, includes some unique nonce-derived material of its own (the H3 hash image), thereby ensuring that each of the two parties can unilaterally force the resulting s0 shared secret to be unique for each media stream, even if one party by some error fails to produce a unique nonce.

Note: Since the nonce is used to calculate different SRTP key and salt pairs for each media stream, a duplication will result in the same key and salt being generated for the two media streams, which would have disastrous security consequences.

At this point in Preshared mode, the two endpoints begin key generation as described in [Section 4.5](#), now that there is a defined

s0 and ZRTPSess key.

4.5. Key Generation

The following calculations derive a set of keys from s0. For the original media stream that calculated s0 from the DH exchange, s0 means the original s0. For any additional media streams that were activated in Multistream mode, s0 means s0n, for the n-th media stream. It is also assumed that the ZRTPSess key has been defined.

Various keys, such as those used by SRTP, must be derived from the shared secret s0. To do this, ZRTP uses an HMAC-based key derivation function, keyed by s0, instead of simply drawing subkey material directly from s0, as defined in NIST SP800-56A. The possibly greater noninvertability of HMAC may add an extra measure of isolation for the derived keys.

The SRTP master key and master salt are derived from s0. Separate SRTP keys and salts are used in each direction for each media stream. Unless otherwise specified, ZRTP uses SRTP with no MKI, 32 bit authentication using HMAC-SHA1, AES-CM 128 or 256 bit key length, 112 bit session salt key length, 2^48 key derivation rate, and SRTP prefix length 0.

The ZRTP initiator encrypts and the ZRTP responder decrypts packets by using srtpkeyi and srtpsalti, while the ZRTP responder encrypts and the ZRTP initiator decrypts packets by using srtpkeyr and

srtpsaltr. These are generated by:

```
srtpkeyi = HMAC(s0,"Initiator SRTP master key")
srtpsalti = HMAC(s0,"Initiator SRTP master salt")
srtpkeyr = HMAC(s0,"Responder SRTP master key")
srtpsaltr = HMAC(s0,"Responder SRTP master salt")
```

The SRTP key and salt values are truncated (taking the leftmost bits) to the length determined by the chosen SRTP algorithm.

The HMAC keys are the same length as the output of the underlying hash function, and are thus generated without truncation by:

```
hmackeyi = HMAC(s0,"Initiator HMAC key")
```

```
hmackeyr = HMAC(s0,"Responder HMAC key")
```

Note that these HMAC keys are used only by ZRTP and not by SRTP.

Note: Different HMAC keys are needed for the initiator and the responder to ensure that GoClear messages in each direction are unique and can not be cached by an attacker and reflected back to the endpoint.

ZRTP keys are generated for the initiator and responder to use to encrypt the Confirm1 and Confirm2 messages. They are truncated to the same size as the negotiated SRTP key size.

```
zrtkeyi = HMAC(s0,"Initiator ZRTP key")
```

```
zrtkeyr = HMAC(s0,"Responder ZRTP key")
```

All key material is destroyed as soon as it is no longer needed, no later than the end of the call. *s0* is erased in [Section 4.6.1](#), and the rest of the session key material is erased in [Section 4.7.2.1](#) and [Section 4.7.3](#).

The Short Authentication String (SAS) value is calculated from the HMAC of a fixed string, keyed with the ZRTPSess key derived from the DH key agreement. This means the same SAS is used for all media streams which are derived from a single DH key agreement in a ZRTP session.

```
sashash = HMAC(ZRTPSess,"SAS")
```

```
sasvalue = sashash [truncated to leftmost 32 bits]
```

[4.6](#). Confirmation

The Confirm1 and Confirm2 messages (Figure 10) contain the cache expiration interval (defined in [Section 4.9](#)) for the newly generated retained shared secret. The flagoctet is an 8 bit unsigned integer made up of these flags: the PBX Enrollment flag (E) defined in

[Section 7.3.1](#), SAS Verified flag (V) defined in [Section 7.1](#), Allow Clear flag (A) defined in [Section 4.7.2](#), and Disclosure flag (D) defined in [Section 11](#).

```
flagoctet = (E * 2^3) + (V * 2^2) + (A * 2^1) + (D * 2^0)
```

Part of the Confirm1 and Confirm2 messages are encrypted using full-

block Cipher Feedback Mode, and contain a 128-bit random CFB Initialization Vector (IV). The Confirm1 and Confirm2 messages also contain an HMAC covering the encrypted part of the Confirm1 or Confirm2 message which includes a string of zeros, the signature length, flag octet, cache expiration interval, signature type block (if present) and signature block ([Section 7.2](#)) (if present). For the responder:

```
hmac = HMAC(hmackeyr, encrypted part of Confirm1)
```

For the initiator:

```
hmac = HMAC(hmackeyi, encrypted part of Confirm2)
```

The hmackeyi and hmackeyr keys are computed in [Section 4.5](#).

The exchange is completed when the responder sends either the Conf2ACK message or the responder's first SRTP media packet (with a valid SRTP auth tag). The initiator MUST treat the first valid SRTP media from the responder as equivalent to receiving a Conf2ACK. The responder may respond to Confirm2 with either SRTP media or Conf2ACK, or both, in whichever order the responder chooses (or whichever order the "cloud" chooses to deliver them).

[4.6.1](#). Updating the Cache of Shared Secrets

After receiving the Confirm messages, both parties must now update their retained shared secret rs1 in their respective caches, provided the following conditions hold:

- 1) This key exchange is either DH or Preshared mode, not Multistream mode, which does not update the cache.
- 2) Depending on the values of the cache expiration intervals that are received in the two Confirm messages, there are some scenarios that do not update the cache, as explained in [Section 4.9](#).
- 3) The responder MUST receive the initiator's Confirm2 message before updating the responder's cache.
- 4) The initiator MUST receive either the responder's Conf2ACK message or the responder's SRTP media (with a valid SRTP auth tag) before updating the initiator's cache.

For DH mode only, before updating the retained shared secret rs1 in the cache, each party first discards their old rs2 and copies their old rs1 to rs2. The old rs1 is saved to rs2 because of the risk of session interruption after one party has updated his own rs1 but before the other party has enough information to update her own rs1. If that happens, they may regain cache sync in the next session by using rs2 (per [Section 4.3](#)). This mitigates the well-known Byzantine Generals' Problem [[Byzantine](#)]. The old rs1 value is not saved in Preshared mode.

For DH mode and Preshared mode, both parties compute a new rs1 value from s0 this way:

```
rs1 = HMAC(s0,"retained secret")
```

After s0 is used to derive the new rs1, it MUST be erased. Even if rs1 is not updated (in the case of Multistream mode), s0 MUST still be destroyed.

[4.7.](#) Termination

A ZRTP session is normally terminated at the end of a call, but it may be terminated early by either the Error message or the GoClear message.

[4.7.1.](#) Termination via Error message

The Error message ([Section 5.9](#)) is used to terminate an in-progress ZRTP exchange due to an error. The Error message contains an integer Error Code for debugging purposes. The termination of a ZRTP key agreement exchange results in no updates to the cached shared secrets and deletion of all crypto context.

The ZRTP Session key, ZRTPSess, is only deleted if the ZRTP session in which it was generated and all ZRTP sessions which are using it are terminated.

[4.7.2.](#) Termination via GoClear message

The GoClear message ([Section 5.11](#)) is used to switch from SRTP to RTP, usually because the user has chosen to do that by pressing a button. The GoClear uses an HMAC of the Message Type Block sent in the GoClear Message computed with the hmackey derived from the shared secret. This HMAC is truncated to the leftmost 64 bits. When sent by the initiator:

Internet-Draft

ZRTP

October 2008

```
clear_hmac = HMAC(hmackeyi, "GoClear ")
```

When sent by the responder:

```
clear_hmac = HMAC(hmackeyr, "GoClear ")
```

A GoClear message which does not receive a ClearACK response must be resent. If a GoClear message is received with a bad HMAC, it must be ignored, and no ClearACK is sent.

A ZRTP endpoint MAY choose to accept GoClear messages after the session has switched to SRTP, allowing the session to revert to RTP. This is indicated in the Confirm1 or Confirm2 messages (Figure 10) by setting the Allow Clear flag (A). If an endpoint sets the Allow Clear (A) flag in their Confirm message, it indicates that they support receiving GoClear messages.

A ZRTP endpoint that receives a GoClear MUST authenticate the message by checking the clear_hmac. If the message authenticates, the endpoint stops sending SRTP packets, and generates a ClearACK in response. It MUST also delete all the crypto key material for all the SRTP media streams, as defined in [Section 4.7.2.1](#).

Until confirmation from the user is received (e.g. clicking a button, pressing a DTMF key, etc.), the ZRTP endpoint MUST NOT resume sending RTP packets. The endpoint then renders to the user an indication that the media session has switched to clear mode, and waits for confirmation from the user. This blocks the flow of sensitive discourse until the user is forced to take notice that he's no longer protected by encryption. To prevent pinholes from closing or NAT bindings from expiring, the ClearACK message MAY be resent at regular intervals (e.g. every 5 seconds) while waiting for confirmation from the user. After confirmation of the notification is received from the user, the sending of RTP packets may begin.

After sending a GoClear message, the ZRTP endpoint stops sending SRTP packets. When a ClearACK is received, the ZRTP endpoint deletes the crypto context for the SRTP session, as defined in [Section 4.7.2.1](#), and may then resume sending RTP packets.

In the event a ClearACK is not received before the retransmissions of GoClear are exhausted, the key material is deleted, as defined in [Section 4.7.2.1](#).

After the users have transitioned from SRTP media back to RTP media (clear mode), they may decide later to return to secure mode by manual activation, usually by pressing a GO SECURE button. In that case, a new secure session is initiated by the party that presses the

button, by sending a new Commit packet, leading to a new session key negotiation. It is not necessary to send another Hello packet, as the two parties have already done that at the start of the call and thus have already discovered each other's ZRTP capabilities. It is possible for users to toggle back and forth between clear and secure modes multiple times in the same call, just as they could in the old days of secure PSTN phones.

[4.7.2.1](#). Key Destruction for GoClear message

All SRTP session key material **MUST** be erased by the receiver of the GoClear message upon receiving a properly authenticated GoClear. The same key destruction **MUST** be done by the sender of GoClear message, upon receiving the ClearACK.

In particular, the destroyed key material includes the SRTP session keys and salts, SRTP master keys and salts, and all material sufficient to reconstruct the SRTP keys and salts, including ZRTPSess (s0 should have been destroyed earlier, in [Section 4.6.1](#)). All key material that would have been erased at the end of the SIP session **MUST** be erased. However, ZRTPSess is destroyed in a manner different from the other key material. Both parties replace ZRTPSess with a hash of itself, without truncation:

```
ZRTPSess = hash(ZRTPSess)
```

This meets the requirements of Perfect Forward Secrecy (PFS), but preserves a new version of ZRTPSess, so that the user can later re-initiate secure mode during the same call without performing another Diffie-Hellman calculation using Multistream mode which requires and assumes the existence of ZRTPSess with the same value at both ZRTP endpoints. A new key negotiation after a GoClear **SHOULD** use a Multistream Commit message.

Note: Multistream mode is preferred over a Diffie-Hellman mode since this does not require the generation of a new hash chain and

a new signaling exchange to exchange new hash values.

Later, at the end of the entire call, ZRTPSess is finally destroyed along with the other key material, as described in [Section 4.7.3](#).

[4.7.3](#). Key Destruction at Termination

All SRTP session key material MUST be erased by both parties at the end of the call. In particular, the destroyed key material includes the SRTP session keys and salts, SRTP master keys and salts, and all material sufficient to reconstruct the SRTP keys and salts, including ZRTPSess and s0 (although s0 should have been destroyed earlier, in

[Section 4.6.1](#)). The only exceptions are the cached shared secrets needed for future calls, including rs1, rs2, and pbxsecret.

[4.8](#). Random Number Generation

The ZRTP protocol uses random numbers for cryptographic key material, notably for the DH secret exponents and nonces, which must be freshly generated with each session. Whenever a random number is needed, all of the following criteria must be satisfied:

Random numbers MUST be freshly generated, meaning that it must not have been used in a previous calculation.

When generating a random number k of L bits in length, k MUST be chosen with equal probability from the range of $[1 < k < 2^L]$.

It MUST be derived from a physical entropy source, such as RF noise, acoustic noise, thermal noise, high resolution timings of environmental events, or other unpredictable physical sources of entropy. For a detailed explanation of cryptographic grade random numbers and guidance for collecting suitable entropy, see [RFC 4086](#) [[RFC4086](#)] and Chapter 10 of Practical Cryptography [[Ferguson](#)]. The raw entropy must be distilled and processed through a deterministic random bit generator (DRBG). Examples of DRBGs may be found in NIST SP 800-90 [[SP800-90](#)], and in [[Ferguson](#)]. Failure to use true entropy from the physical environment as a basis for generating random cryptographic key material would lead to a disastrous loss of security.

[4.9.](#) ZID and Cache Operation

Each instance of ZRTP has a unique 96-bit random ZRTP ID or ZID that is generated once at installation time. It is used to look up retained shared secrets in a local cache. A single global ZID for a single installation is the simplest way to implement ZIDs. However, it is specifically not precluded for an implementation to use multiple ZIDs, up to the limit of a separate one per callee. This then turns it into a long-lived "association ID" that does not apply to any other associations between a different pair of parties. It is a goal of this protocol to permit both options to interoperate freely.

Each time a new s_0 is calculated, a new retained shared secret rs_1 is generated and stored in the cache, indexed by the ZID of the other endpoint. This cache updating is described in [Section 4.6.1](#). For the new retained shared secret, each endpoint chooses a cache expiration value which is an unsigned 32 bit integer of the number of seconds that this secret should be retained in the cache. The time

interval is relative to when the Confirm1 message is sent or received.

The cache intervals are exchanged in the Confirm1 and Confirm2 messages (Figure 10). The actual cache interval used by both endpoints is the minimum of the values from the Confirm1 and Confirm2 messages. A value of 0 seconds means the newly-computed shared secret SHOULD NOT be stored in the cache, and if a cache entry already exists from an earlier call, the stored cache interval should be set to 0. This means if either Confirm message contains a null cache expiration interval, and there is no cache entry already defined, no new cache entry is created. A value of 0xffffffff means the secret should be cached indefinitely and is the recommended value. If the ZRTP exchange is Multistream Mode, the field in the Confirm1 and Confirm2 is set to 0xffffffff and ignored, and the cache is not updated.

The expiration interval need not be used to force the deletion of a shared secret from the cache when the interval has expired. It just means the shared secret MAY be deleted from that cache at any point after the interval has expired without causing the other party to note it as an unexpected security event when the next key negotiation

occurs between the same two parties. This means there need not be perfectly synchronized deletion of expired secrets from the two caches, and makes it easy to avoid a race condition that might otherwise be caused by clock skew.

If the expiration interval is not properly agreed to by both endpoints, it may later result in false alarms of MiTM attacks, due to apparent cache mismatches ([Section 4.3.3](#)).

[4.9.1](#). Cacheless implementations

It is possible to implement a simplified but nonetheless useful profile of the ZRTP protocol that does not support any caching of shared secrets. In this case the cache expiration interval should always be set to zero, and the SAS Verified (V) flag ([Section 7.1](#)) should always be set to false. The users would have to rely exclusively on the verbal SAS comparison for every call. That is, unless MiTM protection is provided by the mechanisms in [Section 8.1.1](#) or [Section 7.2](#), which introduce their own forms of complexity.

If caching of shared secrets is not supported, it would sacrifice the key continuity features, as well as Preshared mode ([Section 4.4.3](#)). There would also be no PBX trusted MiTM ([Section 7.3](#)) features, including the PBX security enrollment ([Section 7.3.1](#)) mechanism.

[5](#). ZRTP Messages

All ZRTP messages use the message format defined in Figure 2. All word lengths referenced in this specification are 32 bits or 4 octets. All integer fields are carried in network byte order, that is, most significant byte (octet) first, commonly known as big-endian.

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0 0 0 1|Not Used (set to zero) |           Sequence Number           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           ZRTP Magic Cookie (0x5a525450)           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Source Identifier           |

```


can not be overstated. The probability of such a false alarm hinges on a mere 16-bit checksum that usually protects UDP packets, so more error detection is needed. For these reasons, this belt-and-suspenders approach is used to minimize the chance of a transmission error affecting the ZRTP key agreement.

The CRC is calculated across the entire ZRTP packet shown in Figure 2, including the ZRTP Header and the ZRTP Message, but not including the CRC field. If a ZRTP message fails the CRC check, it is silently discarded.

[5.1.](#) ZRTP Message Formats

ZRTP messages are designed to simplify endpoint parsing requirements and to reduce the opportunities for buffer overflow attacks (a good goal of any security extension should be to not introduce new attack vectors).

ZRTP uses 8 octets (2 words) blocks to encode Message Type. 4 octets (1 word) blocks are used to encode Hash Type, Cipher Type, and Key Agreement Type, and Authentication Tag. The values in the blocks are ASCII strings which are extended with spaces (0x20) to make them the desired length. Currently defined block values are listed in Tables 1-6 below.

Additional block values may be defined and used.

ZRTP uses this ASCII encoding to simplify debugging and make it "Wireshark (Ethereal) friendly".

[5.1.1.](#) Message Type Block

Currently 14 Message Type Blocks are defined - they represent the set of ZRTP message primitives. ZRTP endpoints MUST support the Hello, HelloACK, Commit, DHPart1, DHPart2, Confirm1, Confirm2, Conf2ACK, SASrelay, RelayACK, Error and ErrorACK message types. ZRTP endpoints MAY support the GoClear and ClearACK messages. Additional messages may be defined in extensions to ZRTP.

"Hello "		Hello Message
"HelloACK"		HelloACK Message
"Commit "		Commit Message
"DHPart1 "		DHPart1 Message
"DHPart2 "		DHPart2 Message
"Confirm1"		Confirm1 Message
"Confirm2"		Confirm2 Message
"Conf2ACK"		Conf2ACK Message
"Error "		Error Message
"ErrorACK"		ErrorACK Message
"GoClear "		GoClear Message
"ClearACK"		ClearACK Message
"SASrelay"		SASrelay Message
"RelayACK"		RelayACK Message

Table 1. Message Type Block Values

[5.1.2.](#) Hash Type Block

Only one Hash Type is currently defined, SHA-256 [[FIPS-180-2](#)], and all ZRTP endpoints MUST support this hash. Additional Hash Types can be registered and used, such as the NIST SHA-3 hash [[SHA-3](#)] when it becomes available. Note that the Hash Type refers to the hash algorithm that will be used throughout the ZRTP key exchange, not the hash algorithm to be used in the SRTP Authentication Tag.

ZRTP makes use of HMAC message authentication codes based on the negotiated Hash Type. The HMAC function is defined in [[FIPS-198-1](#)]. Test vectors for HMAC-SHA-256 may be found in [[RFC4231](#)].

Hash Type Block	Meaning
"S256"	SHA-256 Hash defined in FIPS 180-2

Table 2. Hash Type Block Values

All hashes and HMACs used throughout the ZRTP protocol will use the negotiated Hash Type, except for the special cases noted in [Section 5.1.2.1](#).

[5.1.2.1](#). Implicit Hash and HMAC algorithm

While most of the HMACs used in ZRTP are defined by the negotiated Hash Type ([Section 5.1.2](#)), some hashes and HMACs must be precomputed prior to negotiations, and thus cannot have their algorithms negotiated during the ZRTP exchange. They are implicitly predetermined to use SHA-256 [[FIPS-180-2](#)] and HMAC-SHA-256.

These are the hashes and HMACs that MUST use the Implicit hash and HMAC algorithm:

The hash chain H0-H3 defined in [Section 9](#).

The HMACs that are keyed by this hash chain, as defined in [Section 8.1.1](#).

The Hello Hash in the a=zrtp-hash attribute defined in [Section 8.1](#).

ZRTP defines a method for negotiating different ZRTP protocol versions ([Section 4.1.1](#)). SHA-256 is the Implicit Hash for ZRTP protocol version 1.00. Future ZRTP protocol versions may, if appropriate, use another hash algorithm as the Implicit Hash, such as the NIST SHA-3 hash [[SHA-3](#)] when it becomes available. For example, a future SIP packet may list two a=zrtp-hash SDP attributes, one based on SHA-256 for ZRTP version 1.00, and another based on SHA-3 for ZRTP version 2.00.

[5.1.3](#). Cipher Type Block

All ZRTP endpoints MUST support AES-128 (AES1) and MAY support AES-256 (AES3) or other Cipher Types. The choice of the AES key length is coupled to the Key Agreement type, as explained in [Section 5.1.5](#).

The use of AES-128 in SRTP is defined by [[RFC3711](#)]. The use of AES-256 in SRTP is defined by [[I-D.ietf-avt-srtp-big-aes](#)].

Internet-Draft

ZRTP

October 2008

Cipher Type Block	Meaning
"AES1"	AES-CM with 128 bit keys as defined in RFC 3711
"AES3"	AES-CM with 256 bit keys

Table 3. Cipher Type Block Values

[5.1.4.](#) Auth Tag Block

All ZRTP endpoints MUST support HMAC-SHA1 authentication, 32 bit and 80 bit length tags as defined in [[RFC3711](#)].

Auth Tag Block	Meaning
"HS32"	HMAC-SHA1 32 bit authentication tag as defined in RFC 3711
"HS80"	HMAC-SHA1 80 bit authentication tag as defined in RFC 3711

Table 4. Auth Tag Values

[5.1.5.](#) Key Agreement Type Block

All ZRTP endpoints MUST support DH3k, SHOULD support Preshared, and MAY support EC25, EC38, and EC52.

If a ZRTP endpoint supports multiple concurrent media streams, such as audio and video, it MUST support Multistream ([Section 4.4.2](#)) mode. Also, if a ZRTP endpoint supports the GoClear message ([Section 4.7.2](#)), it SHOULD support Multistream, to be used if the two parties choose to return to the secure state after going Clear (as explained in [Section 4.7.2.1](#)).

For Finite Field Diffie-Hellman, ZRTP endpoints MUST use the DH parameters defined in [RFC 3526](#) [[RFC3526](#)], as follows. DH3k uses the 3072-bit MODP group. The DH generator g is 2. The random Diffie-Hellman secret exponent SHOULD be twice as long as the AES key length. If AES-128 is used, the DH secret value SHOULD be 256 bits long. If AES-256 is used, the secret value SHOULD be 512 bits long.

If Elliptic Curve DH is used, the ECDH algorithm and key generation is from NIST SP 800-56A [[SP800-56A](#)]. The curves used are from NSA

Suite B [[NSA-Suite-B](#)], which uses the same curves as ECDSA defined by FIPS 186-3 [[FIPS-186-3](#)], and can also be found in [RFC 4753](#) [[RFC4753](#)], sections 3.1 through 3.3. The validation procedures are from NIST SP 800-56A [[SP800-56A](#)] [section 5.6.2.6](#), method 3, ECC Partial Validation. Both the X and Y coordinates of the point on the curve are sent, in the first and second half of the ECDH public value, respectively.

The choice of AES key length is coupled to the choice of key agreement type. If either EC38 or EC52 is chosen as the key agreement, AES-256 (AES3) SHOULD be used. If DH3K or EC25 is chosen, either AES-128 (AES1) or AES-256 (AES3) MAY be used.

ZRTP also defines two non-DH modes, Multistream and Preshared, in which the SRTP key is derived from a shared secret and some nonce material.

Table 5 lists the pv length in words and DHPart1 and DHPart2 message length in words for each Key Agreement Type Block.

Key Agreement Type Block	pv words	message words	Meaning
"DH3k"	96	117	DH mode with $p=3072$ bit prime as defined in RFC 3526
"Prsh"	-	-	Preshared Non-DH mode
"Mult"	-	-	Multistream Non-DH mode

"EC25"		16		37		Elliptic Curve DH, P-256 per RFC 4753, section 3.1

"EC38"		24		45		Elliptic Curve DH, P-384 per RFC 4753, section 3.2

"EC52"		33		54		Elliptic Curve DH, P-521 per RFC 4753, section 3.3

Table 5. Key Agreement Type Block Values

[5.1.6.](#) SAS Type Block

The SAS Type determines how the SAS is rendered to the user so that the user may verbally compare it with his partner over the voice channel. This allows detection of a man-in-the-middle (MitM) attack.

All ZRTP endpoints MUST support the base32 and MAY support the base256 rendering schemes for the Short Authentication String, and other SAS rendering schemes. The ZRTP SAS rendering schemes are described in [Section 7](#).

SAS Type Block		Meaning

"B32 "		Short Authentication String using base32 encoding

"B256"		Short Authentication String using base256 encoding (PGP Word List)

Table 6. SAS Type Block Values

[5.1.7.](#) Signature Type Block

The signature type block is a 4 octet (1 word) block used to

represent the signature algorithm discussed in [Section 7.2](#). Suggested signature algorithms and key lengths are a future subject of standardization.

[5.2](#). Hello message

The Hello message has the format shown in Figure 3. The Hello ZRTP message begins with the preamble value 0x505a then a 16 bit length in 32 bit words. This length includes only the ZRTP message (including the preamble and the length) but not the ZRTP header or CRC.

Next is the Message Type Block and a 4 character string containing the version (ver) of the ZRTP protocol which is "1.00" for this specification. Next is the Client Identifier string (cid) which is 4 words long and identifies the vendor and release of the ZRTP software. The 256-bit hash image H3 is defined in [Section 9](#). The next parameter is the ZID, the 96 bit long unique identifier for the ZRTP endpoint.

The next four bits contains flag bits. The MiTM flag (M) is a Boolean that is set to true if and only if this Hello message is sent from a device, usually a PBX, that has the capability to send an SASrelay message ([Section 5.13](#)). The Passive flag (P) is a Boolean

normally set to False. A ZRTP endpoint which is configured to never initiate secure sessions is regarded as passive, and would set the P bit to True. The next 8 bits are unused and SHOULD be set to zero when sent and MUST be ignored on receipt.

Next is a list of supported Hash algorithms, Cipher algorithms, SRTP Auth Tag types, Key Agreement types, and SAS types. The number of listed algorithms are listed for each type: hc=hash count, cc=cipher count, ac=auth tag count, kc=key agreement count, and sc=sas count. The values for these algorithms are defined in Tables 2, 3, 4, 5, and 6. A count of zero means that only the mandatory to implement algorithms are supported. Mandatory algorithms MAY be included in the list. The order of the list indicates the preferences of the endpoint. If a mandatory algorithm is not included in the list, it is added to the end of the list for preference.

Note: Implementers are encouraged to keep these algorithm lists small - the list does not need to include every cipher and hash

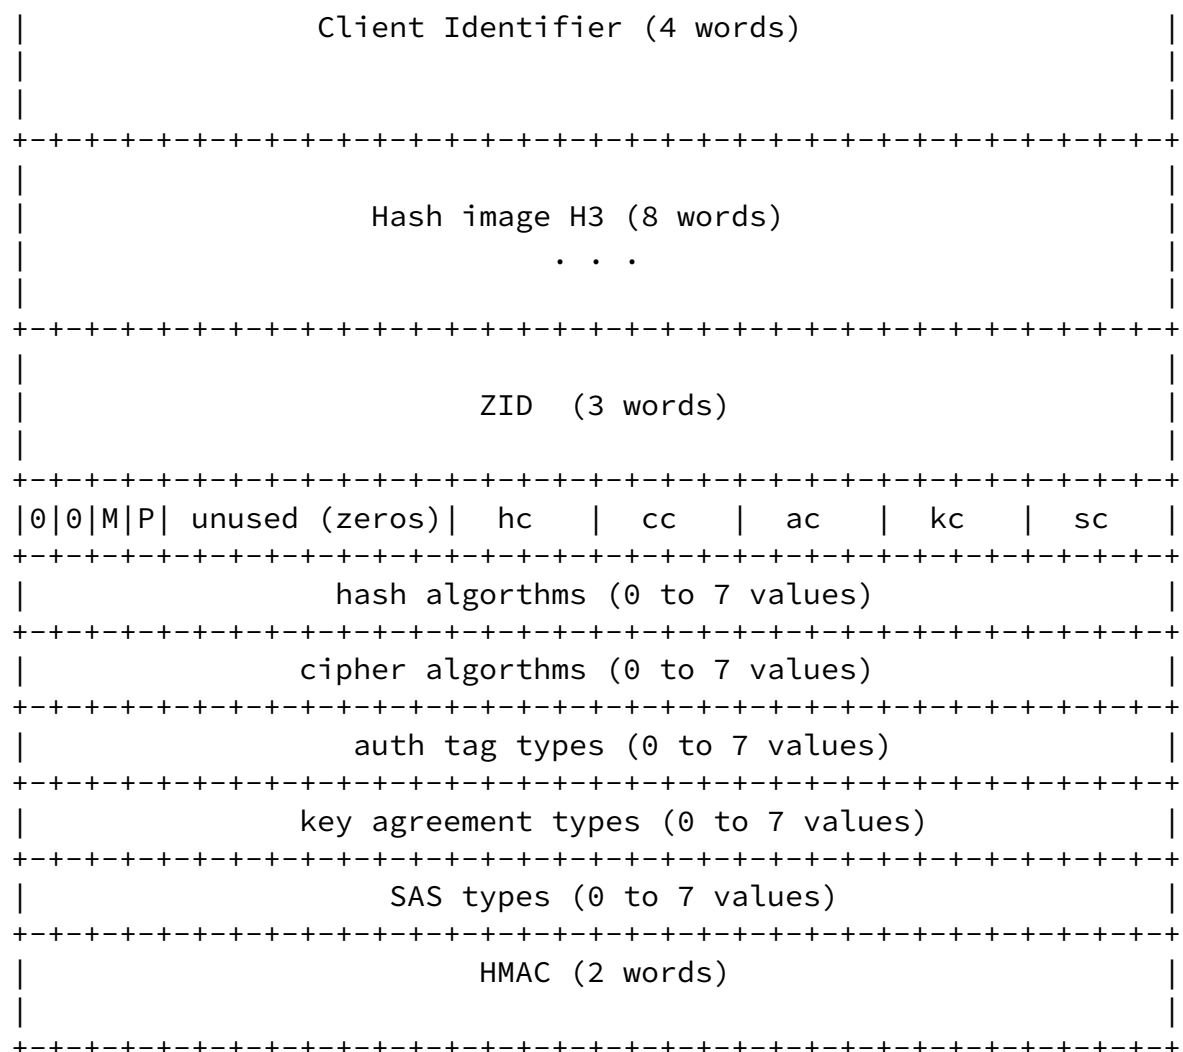


Figure 3: Hello message format

5.3. HelloACK message

The HelloACK message is used to stop retransmissions of a Hello message. A HelloACK is sent regardless if the version number in the Hello is supported or the algorithm list supported. The receipt of a HelloACK stops retransmission of the Hello message. The format is

shown in the Figure below. Note that a Commit message can be sent in place of a HelloACK by an Initiator.

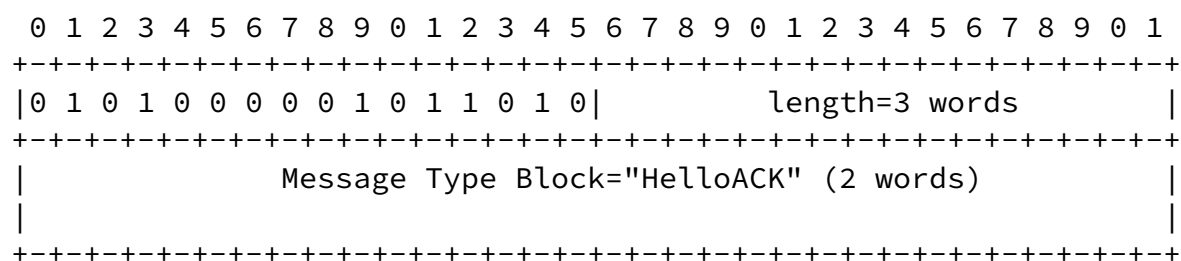


Figure 4: HelloACK message format

5.4. Commit message

The Commit message is sent to initiate the key agreement process after both sides have received a Hello message, which means it can only be sent after receiving both a Hello message and a HelloACK message. There are three subtypes of Commit messages, whose formats are shown in Figure 5, Figure 6, and Figure 7.

The Commit message contains the Message Type Block, then the 256-bit hash image H2 which is defined in [Section 9](#). The next parameter is the initiator's ZID, the 96 bit long unique identifier for the ZRTP endpoint.

Next is a list of algorithms selected by the initiator (hash, cipher, auth tag type, key agreement, sas type). For a DH Commit, the hash value hvi is a hash of the DHPart2 of the Initiator and the Responder's Hello message, as explained in [Section 4.4.1.1](#).

The 64-bit HMAC at the end of the message is computed across the whole message, not including the HMAC. The HMAC key is the sender's H1 (defined in [Section 9](#)), and thus the HMAC cannot be checked by the receiving party until the sender's H1 value is known to the receiving party later in the protocol.

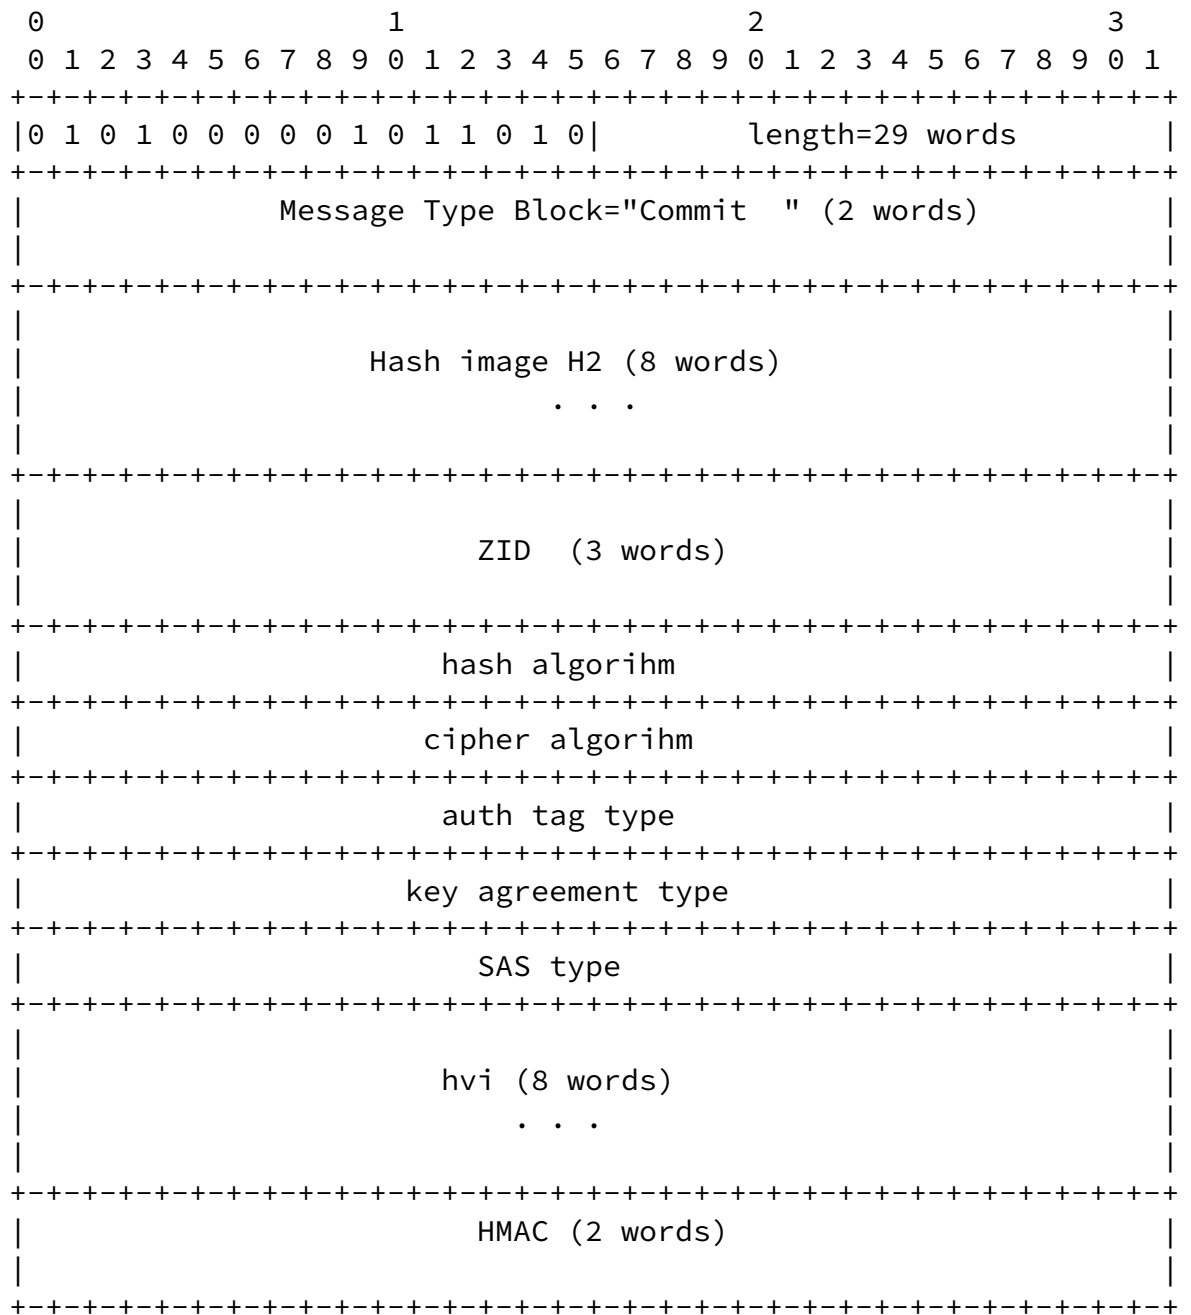


Figure 5: DH Commit message format

Internet-Draft

ZRTP

October 2008

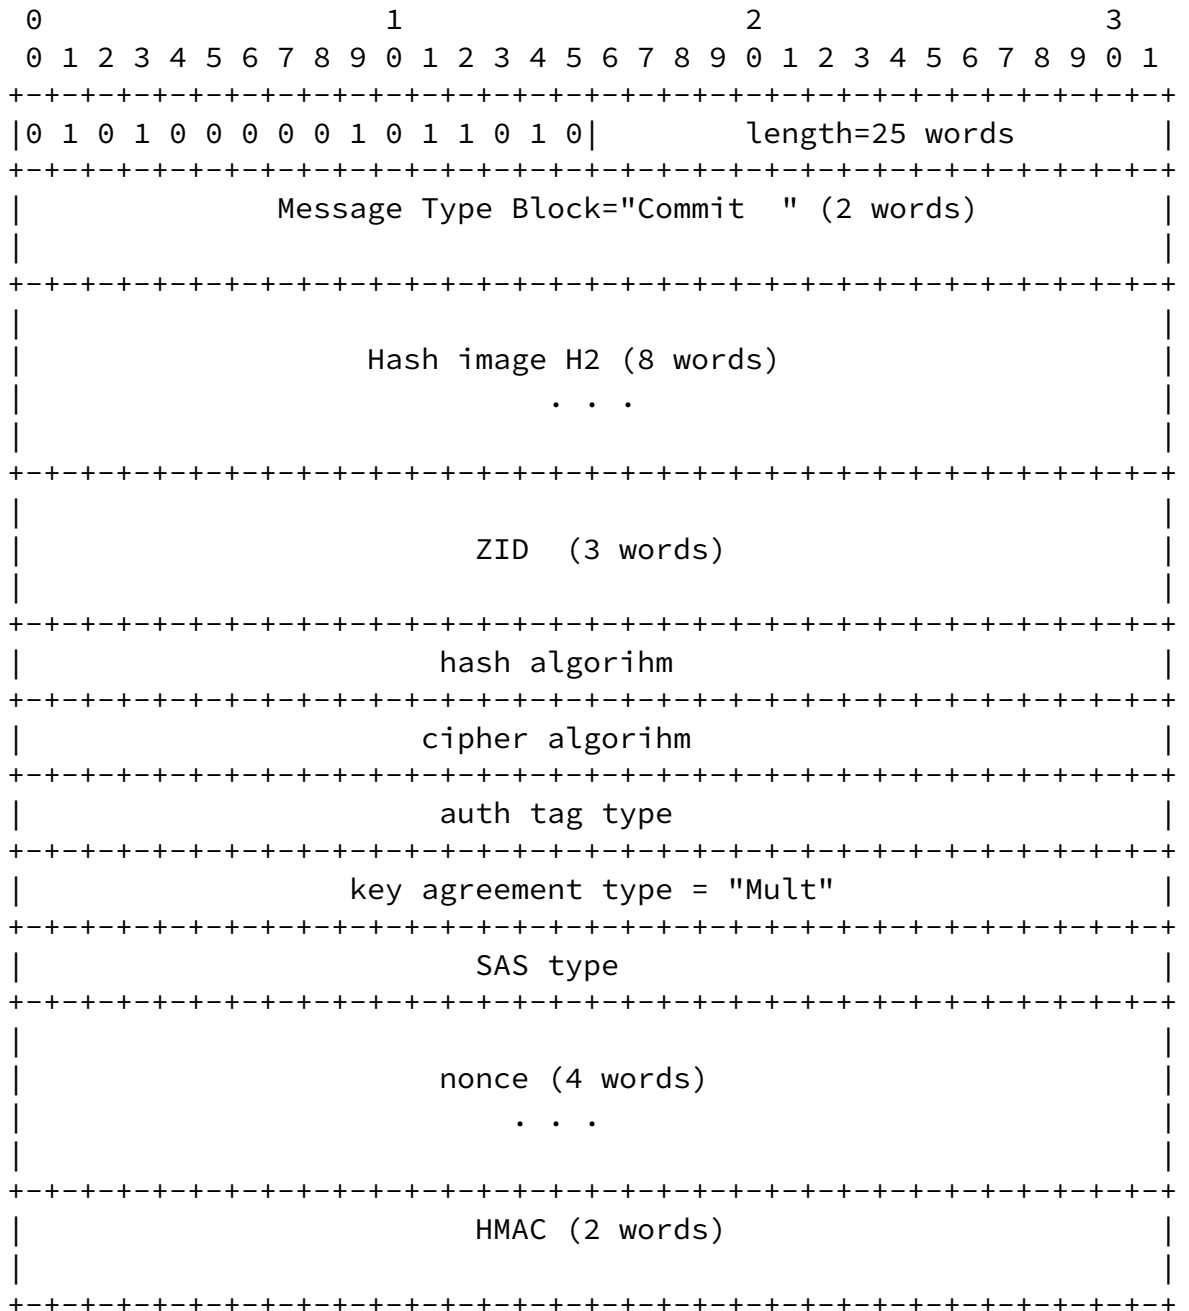


Figure 6: Multistream Commit message format

Internet-Draft

ZRTP

October 2008

```

      0              1              2              3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0 1 0 1 0 0 0 0 0 1 0 1 1 0 1 0|          length=27 words      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                Message Type Block="Commit  " (2 words)                                |
|                                                                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                Hash image H2 (8 words)                                |
|                                . . .                                |
|                                                                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                ZID (3 words)                                |
|                                                                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                hash algorihm                                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                cipher algorihm                                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                auth tag type                                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                key agreement type = "Prsh"                                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                SAS type                                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                nonce (4 words)                                |
|                                . . .                                |
|                                                                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                keyID (2 words)                                |

```




Figure 7: Preshared Commit message format

5.5. DHPart1 message

The DHPart1 message begins the DH exchange. The format is shown in Figure 8 below. The DHPart1 message is sent by the Responder if a valid Commit message is received from the Initiator. The length of the pvr value and the length of the DHPart1 message depends on the Key Agreement Type chosen. This information is contained in Table 5.

Note that for both Multistream and Preshared modes, no DHPart1 or DHPart2 message will be sent.

The 256-bit hash image H1 is defined in [Section 9](#).

The next four parameters are HMACs of potential shared secrets used in generating the ZRTP secret. The first two, rs1IDr and rs2IDr, are the HMACs of the responder's two retained shared secrets, truncated to 64 bits. Next is auxsecretIDr, the HMAC of the responder's auxsecret (defined in [Section 4.3](#)), truncated to 64 bits. The last parameter is the HMAC of the trusted MiTM PBX shared secret pbxsecret, defined in [Section 7.3.1](#). The Message format for the DHPart1 message is shown in Figure 8.

The 64-bit HMAC at the end of the message is computed across the whole message, not including the HMAC. The HMAC key is the sender's H0 (defined in [Section 9](#)), and thus the HMAC cannot be checked by the receiving party until the sender's H0 value is known to the receiving party later in the protocol.

[illegible]

party later in the protocol.

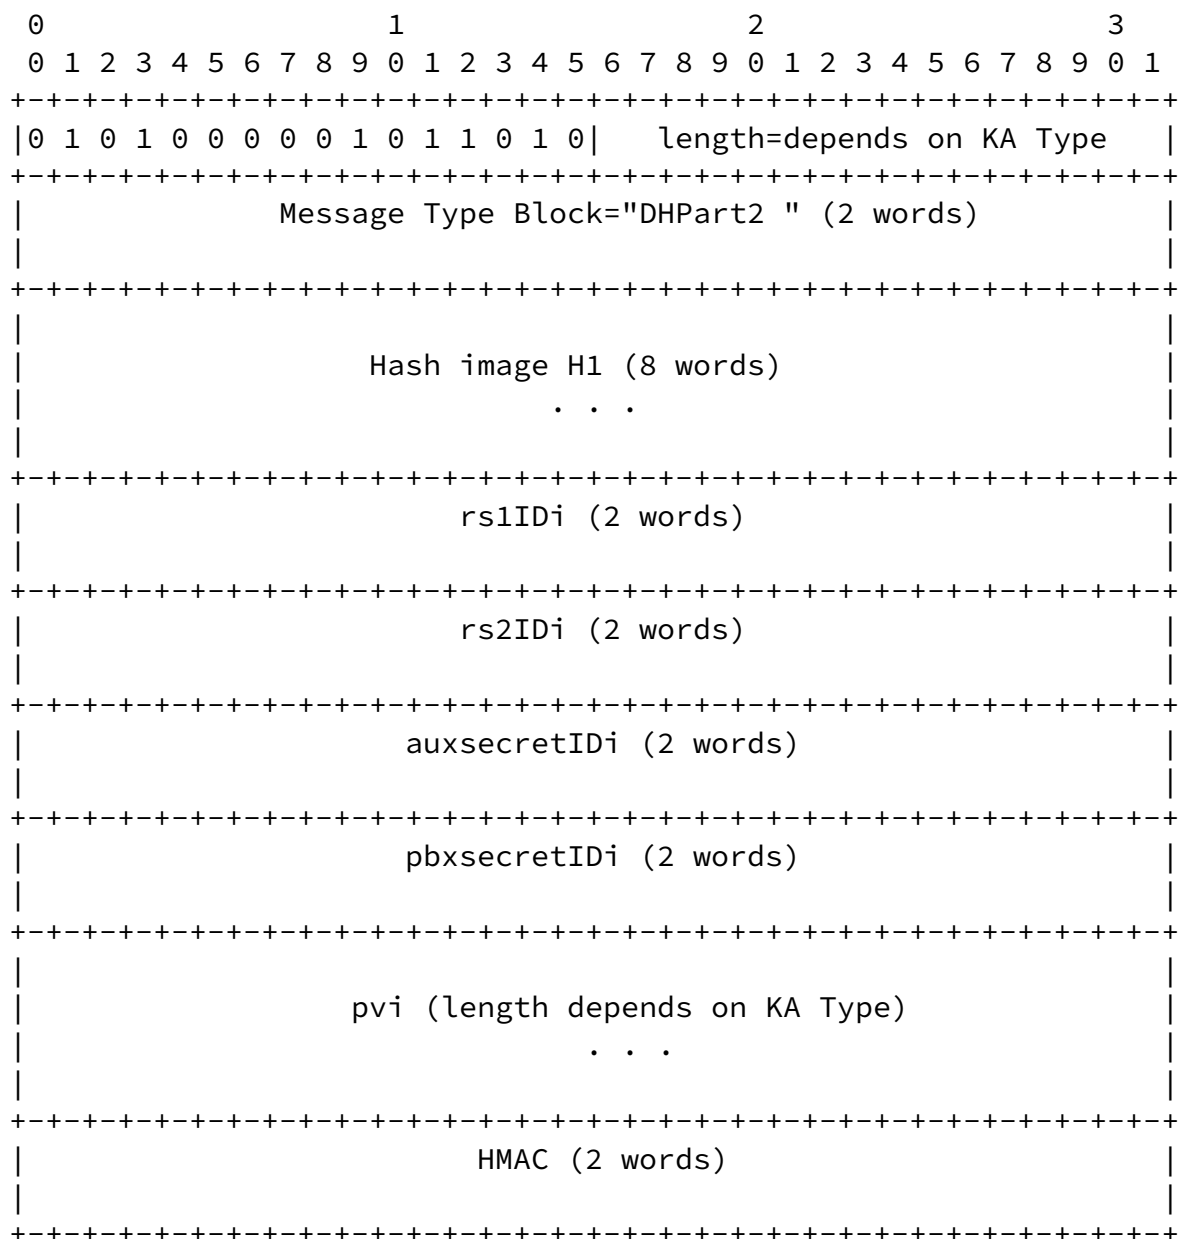


Figure 9: DHPart2 message format

5.7. Confirm1 and Confirm2 messages

The Confirm1 message is sent by the Responder in response to a valid

DHPart2 message after the SRTP session key and parameters have been negotiated. The Confirm2 message is sent by the Initiator in response to a Confirm1 message. The format is shown in Figure 10 below. The message contains the Message Type Block "Confirm1" or "Confirm2". Next is the HMAC, a keyed hash over encrypted part of the message (shown enclosed by "==" in Figure 10). This HMAC is keyed and computed according to [Section 4.6](#). The next 16 octets contain the CFB Initialization Vector. The rest of the message is encrypted using CFB and protected by the HMAC.

The first field inside the encrypted region is the hash pre-image H0, which is defined in detail in [Section 9](#).

The next 15 bits are not used and SHOULD be set to zero when sent and MUST be ignored when received in Confirm1 or Confirm2 messages.

The next 9 bits contain the signature length. If no SAS signature (described in [Section 7.2](#)) is present, all bits are set to zero. The signature length is in words and includes the signature type block. If the calculated signature octet count is not a multiple of 4, zeros are added to pad it out to a word boundary. If no signature block is present, the overall length of the Confirm1 or Confirm2 Message will be set to 19 words.

The next 8 bits are used for flags. Undefined flags are set to zero and ignored. Four flags are currently defined. The PBX Enrollment flag (E) is a Boolean bit defined in [Section 7.3.1](#). The SAS Verified flag (V) is a Boolean bit defined in [Section 7.1](#). The Allow Clear flag (A) is a Boolean bit defined in [Section 4.7.2](#). The Disclosure Flag (D) is a Boolean bit defined in [Section 11](#). The cache expiration interval is defined in [Section 4.9](#).

If the signature length (in words) is non-zero, a signature type block will be present along with a signature block. Next is the signature block. The signature block includes the key used to generate the signature ([Section 7.2](#)).

CFB [[SP800-38A](#)] mode is applied with a feedback length of 128-bits, a full cipher block, and the final block is truncated to match the exact length of the encrypted data. The CFB Initialization Vector is a 128 bit random nonce. The block cipher algorithm and the key size is the same as what was negotiated for the media encryption. CFB is used to encrypt the part of the Confirm1 message beginning after the

CFB IV to the end of the message (the encrypted region is enclosed by "=====" in Figure 10).

The responder uses the `zrtpkeyr` to encrypt the Confirm1 message. The initiator uses the `zrtpkeyi` to encrypt the Confirm2 message.

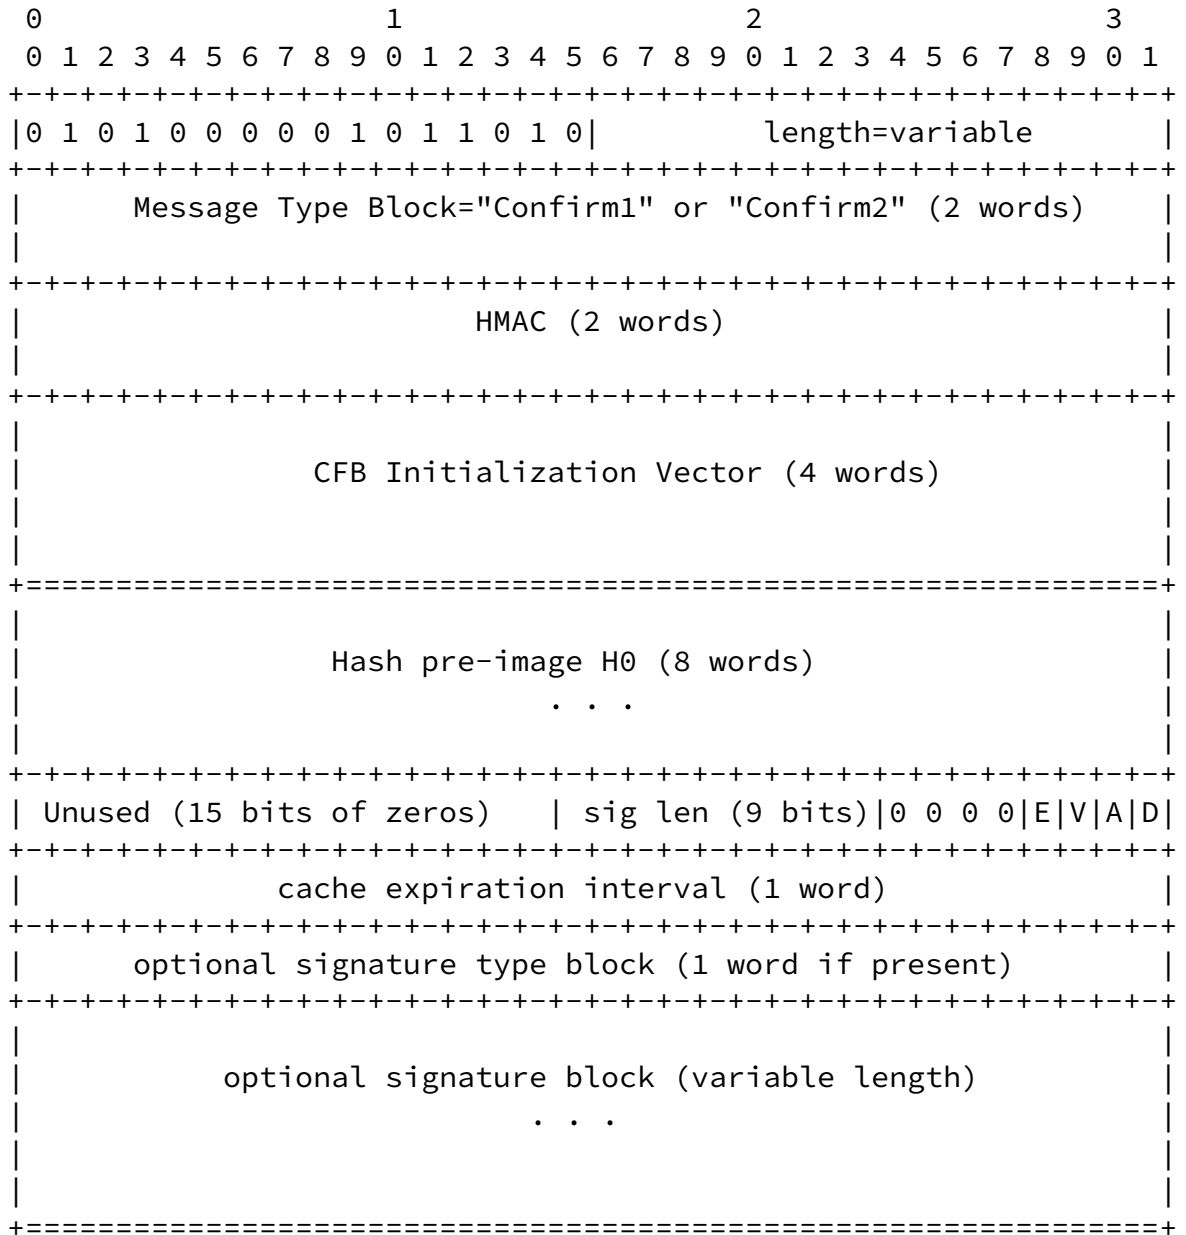


Figure 10: Confirm1 and Confirm2 message format

5.8. Conf2ACK message

The Conf2ACK message is sent by the Responder in response to a valid Confirm2 message. The message format for the Conf2ACK is shown in the Figure below. The receipt of a Conf2ACK stops retransmission of the Confirm2 message. Note that the first SRTP media (with a valid

[illegible]

5.9. Error message

[illegible]

Defined hexadecimal values for the Error Code are listed in Table 7.

Internet-Draft

ZRTP

October 2008

Error Code	Meaning
0x10	Malformed packet (CRC OK, but wrong structure)
0x20	Critical software error
0x30	Unsupported ZRTP version
0x40	Hello components mismatch
0x51	Hash type not supported
0x52	Cipher type not supported
0x53	Public key exchange not supported
0x54	SRTP auth. tag not supported
0x55	SAS scheme not supported
0x56	No shared secret available, DH mode required
0x61	DH Error: bad pvi or pvr (== 1, 0, or p-1)
0x62	DH Error: hvi != hashed data
0x63	Received relayed SAS from untrusted MiTM
0x70	Auth. Error: Bad Confirm pkt HMAC
0x80	Nonce reuse
0x90	Equal ZIDs in Hello
0xA0	Service unavailable

```
0x100    | GoClear packet received, but not allowed
```

Table 7. ZRTP Error Codes

5.10. ErrorACK message

The ErrorACK message is sent in response to an Error message. The receipt of an ErrorACK stops retransmission of the Error message. The format is shown in the Figure below.

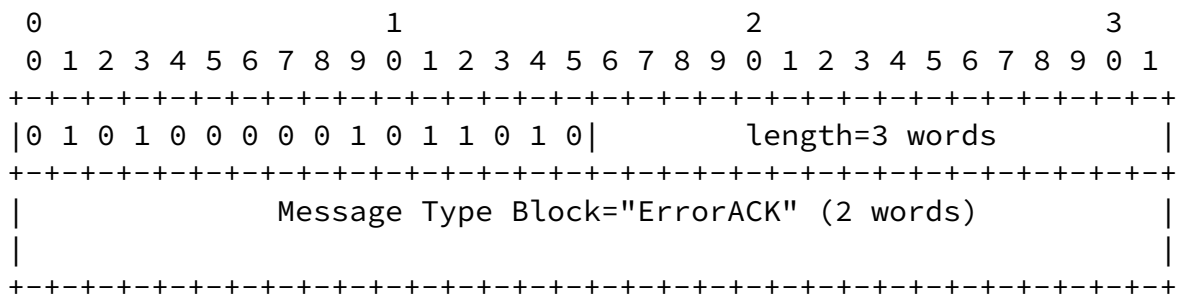
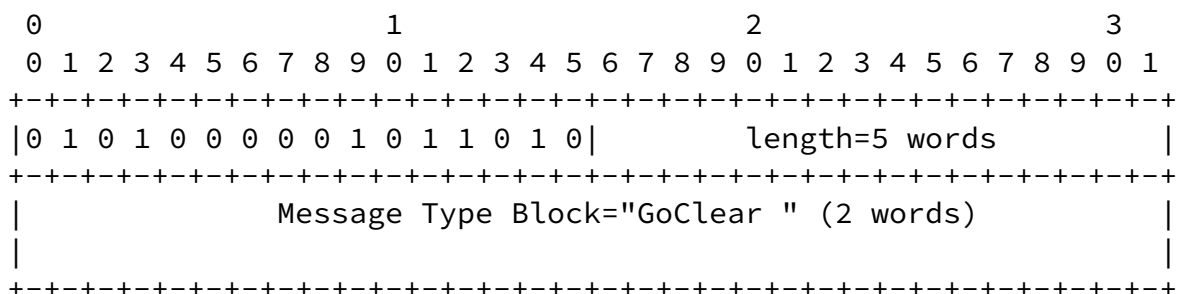


Figure 13: ErrorAck message format

5.11. GoClear message

Support for the GoClear message is OPTIONAL in the protocol, and it is sent to switch from SRTP to RTP. The format is shown in the Figure below. The clear_hmac is used to authenticate the GoClear message so that bogus GoClear messages introduced by an attacker can be detected and discarded. The use of GoClear is described in [Section 4.7.2](#).






Figure 14: GoClear message format

5.12. ClearACK message

Support for the ClearACK message is OPTIONAL in the protocol, and it is sent to acknowledge receipt of a GoClear. A ClearACK is only sent if the clear_hmac from the GoClear message is authenticated. Otherwise, no response is returned. The format is shown in the Figure below.

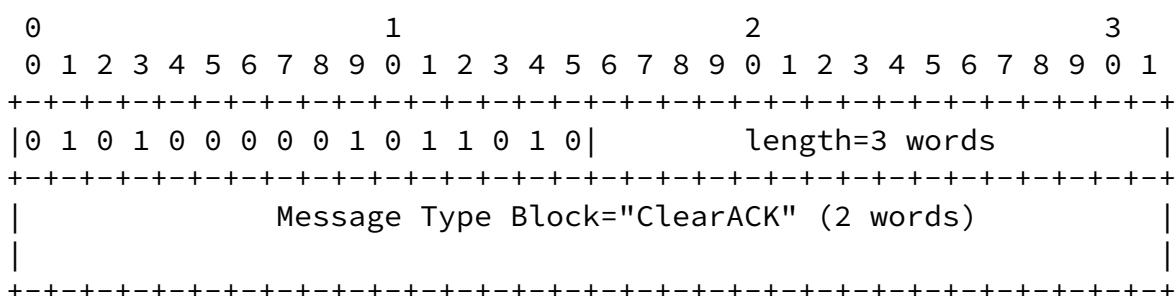


Figure 15: ClearAck message format

5.13. SASrelay message

The SASrelay message is sent by a trusted Man in The Middle (MiTM), most often a PBX. It is not sent as a response to a packet, but is sent as a self-initiated packet by the trusted MiTM. It can only be sent after the rest of the ZRTP key negotiations have completed, after the Confirm packets and their ACKs. It can only be sent after the trusted MiTM has finished key negotiations with the other party, because it is the other party's SAS that is being relayed. It is sent with retry logic until a RelayACK message ([Section 5.14](#)) is

received or the retry schedule has been exhausted.

If a device, usually a PBX, sends an SASrelay message, it MUST have previously declared itself as a MiTM device by setting the MiTM (M) flag in the Hello message ([Section 5.2](#)). If the receiver of the SASrelay message did not previously receive a Hello message with the MiTM (M) flag set, the Relayed SAS SHOULD NOT be rendered. A RelayACK is still sent, but no Error message is sent.

The SASrelay message format is shown in Figure 16 below. The message contains the Message Type Block "SASrelay". Next is the HMAC, a keyed hash over encrypted part of the message (shown enclosed by "==" in Figure 16). This HMAC is keyed the same way as the HMAC in the Confirm messages (see [Section 4.6](#)). The next 16 octets contain the CFB Initialization Vector. The rest of the message is encrypted using CFB and protected by the HMAC.

The next 15 bits are not used and SHOULD be set to zero when sent and MUST be ignored when received in SASrelay messages.

The next 9 bits contain the signature length. The trusted MiTM MAY compute a digital signature on the SAS hash, as described in [Section 7.2](#), using a persistent signing key owned by the trusted MiTM. If no SAS signature is present, all bits are set to zero. The signature length is in words and includes the signature type block. If the calculated signature octet count is not a multiple of 4, zeros are added to pad it out to a word boundary. If no signature block is

present, the overall length of the SASrelay Message will be set to 12 words.

The next 8 bits are used for flags. Undefined flags are set to zero and ignored. Three flags are currently defined. The Disclosure Flag (D) is a Boolean bit defined in [Section 11](#). The Allow Clear flag (A) is a Boolean bit defined in [Section 4.7.2](#). The SAS Verified flag (V) is a Boolean bit defined in [Section 7.1](#). These flags are updated values to the same flags provided earlier in the Confirm packet, but they are updated to reflect the new flag information relayed by the PBX from the other party.

The next 32 bit word contains the rendering scheme for the relayed sasvalue, which will be the same rendering scheme used by the other

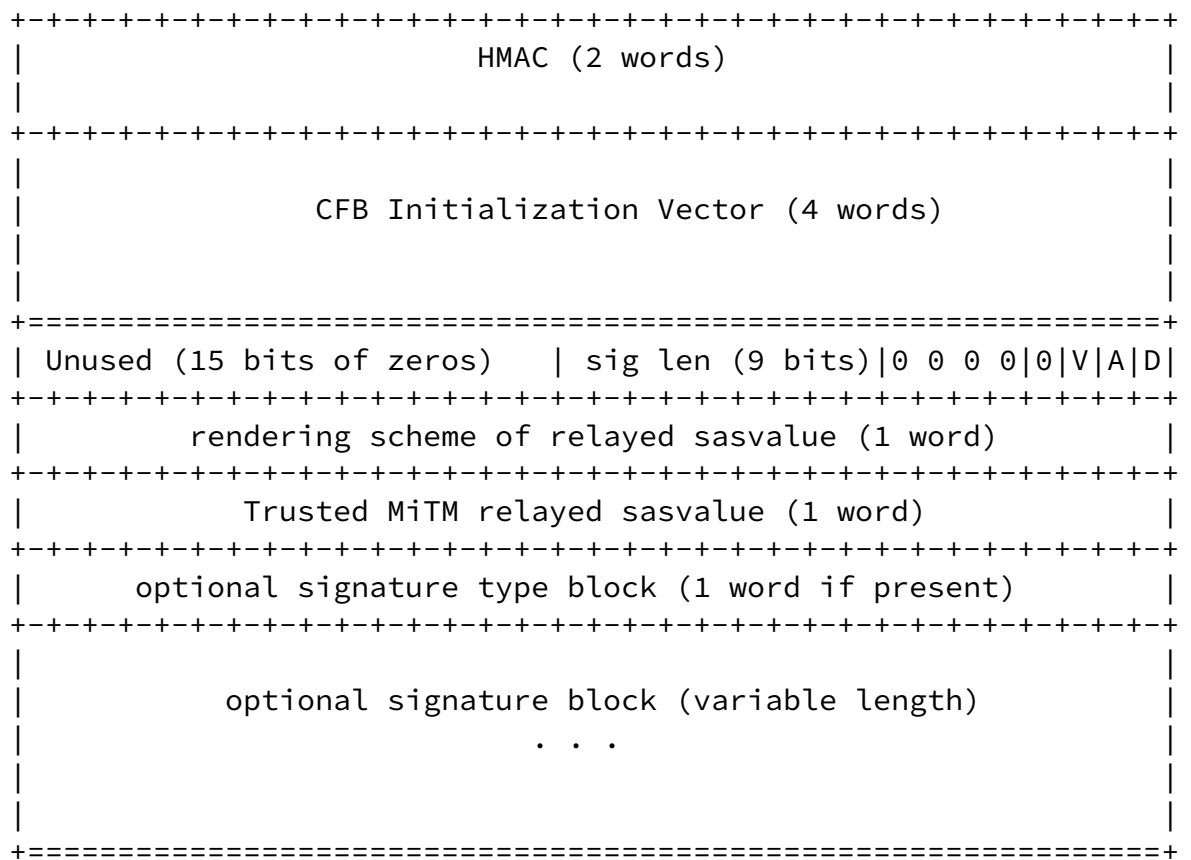


Figure 16: SASrelay message format

5.14. RelayACK message

The RelayACK message is sent in response to a valid SASrelay message. The message format for the RelayACK is shown in the Figure below. The receipt of a RelayACK stops retransmission of the SASrelay message.




Figure 17: RelayACK message format

6. Retransmissions

ZRTP uses two retransmission timers T1 and T2. T1 is used for retransmission of Hello messages, when the support of ZRTP by the other endpoint may not be known. T2 is used in retransmissions of all the other ZRTP messages.

All message retransmissions MUST be identical to the initial message including nonces, public values, etc; otherwise, hashes of the message sequences may not agree.

Practical experience has shown that RTP packet loss at the start of an RTP session can be extremely high. Since the entire ZRTP message exchange occurs during this period, the defined retransmission scheme is defined to be aggressive. Since ZRTP packets with the exception of the DHPart1 and DHPart2 messages are small, this should have minimal effect on overall bandwidth utilization of the media session.

ZRTP endpoints MUST NOT exceed the bandwidth of the resulting media session as determined by the offer/answer exchange in the signaling layer.

Hello ZRTP messages are retransmitted at an interval that starts at T1 seconds and doubles after every retransmission, capping at 200ms. T1 has a recommended initial value of 50 ms. A Hello message is retransmitted 20 times before giving up, which means the entire retry schedule for Hello messages is exhausted after 3.75 seconds (50 + 100 + 18*200 ms). Retransmission of a Hello ends upon receipt of a HelloACK or Commit message.

The post-Hello ZRTP messages are retransmitted only by the session initiator - that is, only Commit, DHPart2, and Confirm2 are retransmitted if the corresponding message from the responder, DHPart1, Confirm1, and Conf2ACK, are not received. Note that the Confirm2 message retransmission can also be stopped by receiving the first SRTP media (with a valid SRTP auth tag) from the responder.

Internet-Draft

ZRTP

October 2008

The GoClear, Error, and SASrelay messages may be initiated and retransmitted by either party, and responded to by the other party, regardless of which party is the overall session initiator. They are retransmitted if the corresponding response message ClearACK, ErrorACK, and RelayACK, are not received.

Non-Hello ZRTP messages are retransmitted at an interval that starts at T2 seconds and doubles after every retransmission, capping at 600ms. T2 has a recommended initial value of 150 ms. Each non-Hello message is retransmitted 10 times before giving up, which means the entire retry schedule is exhausted after 5.25 seconds ($150 + 300 + 8 \times 600$ ms). Only the initiator performs retransmissions. Each message has a response message that stops retransmissions, as shown below in Table 8. The higher values of T2 means that retransmissions will likely only occur with packet loss.

These recommended retransmission intervals are designed for a typical broadband Internet connection. In some high latency communication channels, such as those provided by some mobile phone environments or geostationary satellites, the initial value for the T1 or T2 retransmission timer should be increased to be no less than the round trip time provided by the communications channel. It should take into account the time required to transmit the entire message and the entire reply.

Message	Acknowledgement Message
-----	-----
Hello	HelloACK or Commit
Commit	DHPart1 or Confirm1
DHPart2	Confirm1
Confirm2	Conf2ACK or SRTP media
GoClear	ClearACK
Error	ErrorACK
SASrelay	RelayACK

Table 8. Retransmitted ZRTP Messages and Responses

7. Short Authentication String

This section will discuss the implementation of the Short Authentication String, or SAS in ZRTP. The SAS can be verbally verified by the human users reading the string aloud, or by

validating an OPTIONAL digital signature (described in [Section 7.2](#)) exchanged in the Confirm1 or Confirm2 messages.

The use of hash commitment in the DH exchange ([Section 4.4.1.1](#)) constrains the attacker to only one guess to generate the correct SAS

in his attack, which means the SAS can be quite short. A 16-bit SAS, for example, provides the attacker only one chance out of 65536 of not being detected.

The rendering of the SAS value to the user depends on the SAS Type agreed upon in the Commit message. For the SAS Type of base32, the leftmost 20 bits of the 32-bit sasvalue are rendered as a form of base32 encoding known as z-base-32 [[z-base-32](#)]. The purpose of z-base-32 is to represent arbitrary sequences of octets in a form that is as convenient as possible for human users to manipulate. As a result, the choice of characters is slightly different from base32 as defined in [RFC 3548](#). The leftmost 20 bits of the sasvalue results in four base32 characters which are rendered to both ZRTP endpoints. For the SAS Type of base256, the leftmost 16 bits of the 32-bit sasvalue are rendered using the PGP Wordlist [[pgpwordlist](#)] [[Juola1](#)][[Juola2](#)]. Other SAS Types may be defined to render the SAS value in other ways.

The SAS SHOULD be rendered to the user for authentication.

The SAS is not treated as a secret value, but it must be compared to see if it matches at both ends of the communications channel. The two users read it aloud to their partners to see if it matches. This allows detection of a man-in-the-middle (MiTM) attack.

There is only one SAS value computed per call. That is the SAS value for the first media stream established, which computes the ZRTPSess key, using DH mode. The ZRTPSess key is used to compute the SAS, as well as the SRTP session keys for each additional media stream in Multistream mode. This SAS applies to all media streams for the same call.

[7.1](#). SAS Verified Flag

The SAS Verified flag (V) is set based on the user indicating that SAS comparison has been successfully performed. The SAS Verified

flag is exchanged securely in the Confirm1 and Confirm2 messages (Figure 10) of the next session. In other words, each party sends the SAS Verified flag from the previous session in the Confirm message of the current session. It is perfectly reasonable to have a ZRTP endpoint that never sets the SAS Verified flag, because it would require adding complexity to the user interface to allow the user to set it. The SAS Verified flag is not required to be set, but if it is available to the client software, it allows for the possibility that the client software could render to the user that the SAS verify procedure was carried out in a previous session.

Regardless of whether there is a user interface element to allow the

user to set the SAS Verified flag, it is worth caching a shared secret, because doing so reduces opportunities for an attacker in the next call.

If at any time the users carry out the SAS comparison procedure, and it actually fails to match, then this means there is a very resourceful man-in-the-middle. If this is the first call, the MiTM was there on the first call, which is impressive enough. If it happens in a later call, it also means the MiTM must also know the cached shared secret, because you could not have carried out any voice traffic at all unless the session key was correctly computed and is also known to the attacker. This implies the MiTM must have been present in all the previous sessions, since the initial establishment of the first shared secret. This is indeed a resourceful attacker. It also means that if at any time he ceases his participation as a MiTM on one of your calls, the protocol will detect that the cached shared secret is no longer valid -- because it was really two different shared secrets all along, one of them between Alice and the attacker, and the other between the attacker and Bob. The continuity of the cached shared secrets make it possible for us to detect the MiTM when he inserts himself into the ongoing relationship, as well as when he leaves. Also, if the attacker tries to stay with a long lineage of calls, but fails to execute a DH MiTM attack for even one missed call, he is permanently excluded. He can no longer resynchronize with the chain of cached shared secrets.

Some sort of user interface element (maybe a checkbox) is needed to allow the user to tell the software the SAS verify was successful, causing the software to set the SAS Verified flag (V), which

(together with our cached shared secret) obviates the need to perform the SAS procedure in the next call. An additional user interface element can be provided to let the user tell the software he detected an actual SAS mismatch, which indicates a MiTM attack. The software can then take appropriate action, clearing the SAS Verified flag, and erase the cached shared secret from this session. It is up to the implementer to decide if this added user interface complexity is warranted.

If the SAS matches, it means there is no MiTM, which also implies it is now safe to trust a cached shared secret for later calls. If inattentive users don't bother to check the SAS, it means we don't know whether there is or is not a MiTM, so even if we do establish a new cached shared secret, there is a risk that our potential attacker may have a subsequent opportunity to continue inserting himself in the call, until we finally get around to checking the SAS. If the SAS matches, it means no attacker was present for any previous session since we started propagating cached shared secrets, because this session and all the previous sessions were also authenticated

with a continuous lineage of shared secrets.

[7.2.](#) Signing the SAS

In some applications, it may be hard to arrange for two human users to verbally compare the SAS. To handle these cases, ZRTP allows for an OPTIONAL signature feature, which allows the SAS to be checked without human participation. The SAS MAY be signed and the signature sent inside the Confirm1, Confirm2 (Figure 10), or SASrelay (Figure 16) messages. The signature algorithm, length of the signature and the key used to create the signature are all sent along with the signature. The key types and signature algorithms are for future study. The signature is calculated over the entire SAS hash result (sashash) that was truncated down to derive the sasvalue. The signatures exchanged in the encrypted Confirm1, Confirm2, or SASrelay messages MAY be used to authenticate the ZRTP exchange.

[7.3.](#) Relaying the SAS through a PBX

ZRTP is designed to use end-to-end encryption. The two parties' verbal comparison of the short authentication string (SAS) depends on this assumption. But in some PBX environments, such as Asterisk,

there are usage scenarios that have the PBX acting as a trusted man-in-the-middle (MiTM), which means there are two back-to-back ZRTP connections with separate session keys and separate SAS's.

For example, imagine that Bob has a ZRTP-enabled VoIP phone that has been registered with his company's PBX, so that it is regarded as an extension of the PBX. Alice, whose phone is not associated with the PBX, might dial the PBX from the outside, and a ZRTP connection is negotiated between her phone and the PBX. She then selects Bob's extension from the company directory in the PBX. The PBX makes a call to Bob's phone (which might be offsite, many miles away from the PBX through the Internet) and a separate ZRTP connection is negotiated between the PBX and Bob's phone. The two ZRTP sessions have different session keys and different SAS's, which would render the SAS useless for verbal comparison between Alice and Bob. They might even mistakenly believe that a wiretapper is present because of the SAS mismatch, causing undue alarm.

ZRTP has a mechanism for solving this problem by having the PBX relay the Alice/PBX SAS to Bob, sending it through to Bob in a special SASrelay packet as defined in [Section 5.13](#), which is sent after the PBX/Bob ZRTP negotiation is complete, after the Confirm packets. Only the PBX, acting as a special trusted MiTM (trusted by the recipient of the SAS relay packet), will relay the SAS. The SASrelay packet protects the relayed SAS from tampering via an included HMAC, similar to how the Confirm packet is protected. Bob's ZRTP-enabled

phone accepts the relayed SAS for rendering only because Bob's phone had previously been configured to trust the PBX. This special trusted relationship with the PBX can be established through a special security enrollment procedure. After that enrollment procedure, the PBX is treated by Bob as a special trusted MiTM. This results in Alice's SAS being rendered to Bob, so that Alice and Bob may verbally compare them and thus prevent a MiTM attack by any other untrusted MiTM.

A real bad-guy MiTM cannot exploit this protocol feature to mount a MiTM attack and relay Alice's SAS to Bob, because Bob has not previously carried out a special registration ritual with the bad guy. The relayed SAS would not be rendered by Bob's phone, because it did not come from a trusted PBX. The recognition of the special trust relationship is achieved with the prior establishment of a

special shared secret between Bob and his PBX, which is called pbxsecret (defined in [Section 7.3.1](#)), also known as the trusted MiTM key.

The trusted MiTM key can be stored in a special cache at the time of the initial enrollment (which is carried out only once for Bob's phone), and Bob's phone associates this key with the ZID of the PBX, while the PBX associates it with the ZID of Bob's phone. After the enrollment has established and stored this trusted MiTM key, it can be detected during subsequent ZRTP call negotiations between the PBX and Bob's phone, because the PBX and the phone MUST pass the hash of the trusted MiTM key in the DH packet. It is then used as part of the key agreement to calculate s_0 .

During a key agreement with two other ZRTP endpoints, the PBX may have a shared trusted MiTM key with both endpoints, only one endpoint, or neither endpoint. If the PBX has a shared trusted MiTM key with neither endpoint, the PBX SHOULD NOT relay the SAS. If the PBX has a shared trusted MiTM key with only one endpoint, the PBX SHOULD relay the SAS from one party the other by sending an SASrelay message to the endpoint that it shares a trusted MiTM key. If the PBX has a shared trusted MiTM key with both endpoints, the PBX SHOULD relay the SAS from one party the other by sending an SASrelay message to only one of the endpoints.

Note: In the case of sharing trusted MiTM key with both endpoints, it does not matter which endpoint receives the relayed SAS as long as only one endpoint receives it.

The PBX can determine whether it is trusted by the ZRTP user agent of the caller or callee. The presence of a shared trusted MiTM key in the key negotiation sequence indicates that the phone has been enrolled with this PBX and therefore trusts it to act as a trusted

MiTM. The PBX SHOULD relay the SAS from the other party in this case.

The relayed SAS fields contain the SAS rendering type and the binary 32-bit sasvalue. The receiver absolutely MUST NOT render the relayed SAS if it does not come from a specially trusted ZRTP endpoint. The security of the ZRTP protocol depends on not rendering a relayed SAS from an untrusted MiTM, because it may be relayed by a MiTM attacker.

See the SASrelay packet definition (Figure 16) for further details.

To ensure that both Alice and Bob will use the same SAS rendering scheme after the keys are negotiated, the PBX also sends the SASrelay message to the unenrolled party (which does not regard this PBX as a trusted MiTM), conveying the SAS rendering scheme, but not the SAS value, which it sets to zero. The unenrolled party will ignore the relayed SAS field, but will use the specified SAS rendering scheme.

The next section describes the initial enrollment procedure that establishes a special shared secret between the PBX and Bob's phone, a trusted MiTM key, so that the phone will learn to recognize the PBX as a trusted MiTM.

[7.3.1.](#) PBX Enrollment and the PBX Enrollment Flag

Both the PBX and the endpoint need to know when enrollment is taking place. One way of doing this is to setup an enrollment extension on the PBX which a newly configured endpoint would call and establish a ZRTP session. The PBX would then play audio media that offers the user an opportunity to configure his phone to trust this PBX as a trusted MiTM. The PBX calculates and stores the trusted MiTM shared secret in its cache and associates it with this phone, indexed by the phone's ZID. The trusted MiTM PBX shared secret is calculated this way:

```
pbxsecret = HMAC(ZRTPSess,"Trusted MiTM key")
```

The PBX signals the enrollment process by setting the PBX Enrollment flag (E) in the Confirm message (Figure 10). This flag is used to trigger the ZRTP endpoint's user interface to prompt the user if they want to trust this PBX and calculate and store the pbxsecret in the cache. If the user decides to respond by activating the appropriate user interface element (a menu item, checkbox, or button), his ZRTP user agent calculates pbxsecret using the same formula and saves it in a special cache entry associated with this PBX.

If the user elects not to enroll, perhaps because he dialed a wrong number or does not yet feel comfortable with this PBX, he can simply hang up and not save the pbxsecret in his cache. The PBX will have

it saved in the PBX cache, but that will do no harm. The SASrelay

scheme does not depend on the PBX trusting the phone. It only depends on the phone trusting the PBX. It is the phone (the user) who is at risk if the PBX abuses its MiTM privileges.

An endpoint MUST NOT store the pbxsecret in the cache without explicit user authorization.

After this enrollment process, the PBX and the ZRTP-enabled phone both share a secret that enables the phone to recognize the PBX as a trusted MiTM in future calls. This means that when a future call from an outside ZRTP-enabled caller is relayed through the PBX to this phone, the phone will render a relayed SAS from the PBX. If the SASrelay packet comes from a MiTM which does not know the pbxsecret, the phone treats it as a "bad guy" MiTM, and refuses to render the relayed SAS. Regardless of which party initiates any future phone calls through the PBX, the enrolled phone or the outside phone, the PBX will relay the SAS to the enrolled phone.

There are other ways that ZRTP user agents can be configured to trust a PBX. Perhaps the pbxsecret can be configured into the phone by some automated provisioning process in large IT environments. This specification does not require that products be configured solely by this enrollment process. Any process that results in a pbxsecret to be computed and shared between the PBX and the phone will suffice. This is one such method that has been shown to work.

[8.](#) Signaling Interactions

This section discusses how ZRTP, SIP, and SDP work together.

Note that ZRTP may be implemented without coupling with the SIP signaling. For example, ZRTP can be implemented as a "bump in the wire" or as a "bump in the stack" in which RTP sent by the SIP UA is converted to ZRTP. In these cases, the SIP UA will have no knowledge of ZRTP. As a result, the signaling path discovery mechanisms introduced in this section should not be definitive – they are a hint. Despite the absence of an indication of ZRTP support in an offer or answer, a ZRTP endpoint SHOULD still send Hello messages.

ZRTP endpoints which have control over the signaling path include a ZRTP SDP attributes in their SDP offers and answers. The ZRTP attribute, a=zrtp-hash is used to indicate support for ZRTP and to convey a hash of the Hello message. The hash is computed according to [Section 8.1](#).

Aside from the advantages described in [Section 8.1](#), there are a

number of potential uses for this attribute. It is useful when signaling elements would like to know when ZRTP may be utilized by endpoints. It is also useful if endpoints support multiple methods of SRTP key management. The ZRTP attribute can be used to ensure that these key management approaches work together instead of against each other. For example, if only one endpoint supports ZRTP but both support another method to key SRTP, then the other method will be used instead. When used in parallel, an SRTP secret carried in an `a=keymgmt` [RFC4567] or `a=crypto` [RFC4568] attribute can be used as a shared secret for the srtps computation defined in [Section 8.2](#). The ZRTP attribute is also used to signal to an intermediary ZRTP device not to act as a ZRTP endpoint, as discussed in [Section 10](#).

The `a=zrtp-hash` attribute can only be included in the SDP at the media level since Hello messages sent in different media streams will have unique hashes.

The ABNF for the ZRTP attribute is as follows:

```
zrtp-attribute    = "a=zrtp-hash:" zrtp-version zrtp-hash-value
zrtp-version      = token
zrtp-hash-value   = 1*(HEXDIG)
```

Example of the ZRTP attribute in an initial SDP offer or answer used at the session level:

```
v=0
o=bob 2890844527 2890844527 IN IP4 client.biloxi.example.com
s=
c=IN IP4 client.biloxi.example.com
t=0 0
m=audio 3456 RTP/AVP 97 33
a=rtpmap:97 iLBC/8000
a=rtpmap:33 no-op/8000
a=zrtp-hash:1.00 fe30efd02423cb054e50efd0248742ac7a52c8f91bc2df881ae642c371b
```

[8.1](#). Binding the media stream to the signaling layer via the Hello Hash

It is desirable to tie the media stream to the signaling channel to prevent a third party from inserting false media packets. If the signaling layer contains information that ties it to the media stream, false media streams can be rejected.

To accomplish this, a 256-bit hash (using the hash algorithm defined

in [Section 5.1.2.1](#)) is computed across the entire ZRTP Hello message (as shown in Figure 3). This hash image is made available to the

signaling layer, where it is transmitted as a hexadecimal value in the SIP channel using the SDP attribute, `a=zrtp-hash` defined in this specification. Each media stream (audio or video) will have a separate Hello packet, and thus will require a separate `a=zrtp-hash` in an SDP attribute. The recipient of the SIP/SDP message can then use this hash image to detect and reject false Hello packets in the media channel, as well as identify which media stream is associated with this SIP call. Each Hello packet hashes uniquely, because it contains the H3 field derived from a random nonce, defined in [Section 9](#).

The Hello Hash as an SDP attribute is an OPTIONAL feature, because some ZRTP endpoints do not have the ability to add SDP attributes to the signaling. For example, if ZRTP is implemented in a hardware bump-in-the-wire device, it might only have the ability to modify the media packets, not the SIP packets, especially if the SIP packets are integrity protected and thus cannot be modified on the wire. If the SDP has no hash image of the ZRTP Hello message, the recipient's ZRTP user agent cannot check it, and thus will not be able to reject Hello messages based on this hash.

After the Hello Hash is used to properly identify the ZRTP Hello message as belonging to this particular SIP call, the rest of the ZRTP message sequence is protected from false packet injection by other protection mechanisms. For example, the use of the `total_hash` in the shared secret calculation, and also the hash chaining mechanism defined in [Section 9](#).

An attacker who controls only the signaling layer, such as an uncooperative VoIP service provider, may be able to deny service by corrupting the hash of the Hello message in the SDP attribute, which would force ZRTP to reject perfectly good Hello messages. If there is reason to believe this is happening, the ZRTP endpoint MAY allow Hello messages to be accepted that do not match the hash image in the SDP attribute.

Even in the absence of SIP integrity protection, the inclusion of the `a=zrtp-hash` SDP attribute, when coupled with the hash chaining mechanism defined in [Section 9](#), meets the R-ASSOC requirement in the

"...a mechanism for associating key management messages with both the signaling traffic that initiated the session and with protected media traffic. Allowing such an association also allows the SDP offerer to avoid performing CPU-consuming operations (e.g., Diffie-Hellman or public key operations) with attackers that have not seen the signaling messages."

The a=zrtp-hash SDP attribute becomes especially useful if the SDP is integrity-protected end-to-end by SIP Identity ([RFC 4474](#)) [[RFC4474](#)] or better still, Dan Wing's SIP Identity using Media Path [[I-D.wing-sip-identity-media](#)]. This leads to an ability to stop MiTM attacks independent of ZRTP's SAS mechanism, as explained in [Section 8.1.1](#) below.

[8.1.1](#). Integrity-protected signaling enables integrity-protected DH exchange

If and only if the signaling path and the SDP is protected by some form of end-to-end integrity protection, such as one of the abovementioned mechanisms, so that it can guarantee delivery of the a=zrtp-hash attribute without any tampering by a third party, and if there is good reason to trust the signaling layer to protect the interests of the end user, it is possible to authenticate the key exchange and prevent a MiTM attack. This can be done without requiring the users to verbally compare the SAS, by using the hash chaining mechanism defined in [Section 9](#) to provide a series of HMAC keys that protect the entire ZRTP key exchange. Thus, an end-to-end integrity-protected signaling layer automatically enables an integrity-protected Diffie-Hellman exchange in ZRTP, which in turn means immunity from a MiTM attack. Here's how it works.

The integrity-protected SIP SDP contains a hash commitment to the entire Hello message. The Hello message contains H3, which provides a hash commitment for the rest of the hash chain H0-H2 ([Section 9](#)). The Hello message is protected by a 64-bit HMAC, keyed by H2. The Commit message is protected by a 64-bit HMAC keyed by H1. The DHPart1 or DHPart2 messages are protected by a 64-bit HMAC keyed by H0. The HMAC protecting the Confirm messages are computed by a different HMAC key derived from the resulting key agreement. Each

message's HMAC is checked when the HMAC key is received in the next message. If a bad HMAC is discovered, it MUST be treated as a security exception indicating a MiTM attack, perhaps by logging or alerting the user, and MUST NOT be treated as a random error. Random errors are already discovered and quietly rejected by bad CRCs (Figure 2).

The Hello message must be assembled before any hash algorithms are negotiated, so an implicit predetermined hash algorithm and HMAC algorithm (both defined in [Section 5.1.2.1](#)) must be used. All of the aforementioned HMACs keyed by the hashes in the aforementioned hash chain MUST be computed with the HMAC algorithm defined in [Section 5.1.2.1](#), with the HMAC truncated to 64 bits.

The Media Security Requirements

[[I-D.ietf-sip-media-security-requirements](#)] R-EXISTING requirement can

be fully met by leveraging a certificate-backed PKI in the signaling layer to integrity-protect the delivery of the a=zrtp-hash SDP attribute. This would thereby protect ZRTP against a MiTM attack, without requiring the user to check the SAS, without adding any explicit signatures or signature keys to the ZRTP key exchange, and without any extra public key operations or extra packets.

Without an end-to-end integrity protection mechanism in the signaling layer to guarantee delivery of the a=zrtp-hash SDP attribute without modification by a third party, these HMACs alone will not prevent a MiTM attack. In that case, ZRTP's built-in SAS mechanism will still have to be used to authenticate the key exchange. At the time of this writing, very few deployed VoIP clients offer a fully implemented SIP stack that provides end-to-end integrity protection for the delivery of SDP attributes. Also, end-to-end signaling integrity becomes more problematic if E.164 numbers [[RFC3824](#)] are used in SIP. Thus, real-world implementations of ZRTP endpoints will continue to depend on SAS authentication for quite some time. Even after there is widespread availability of SIP user agents that offer integrity protected delivery of SDP attributes, many users will still be faced with the fact that the signaling path may be controlled by institutions that do not have the best interests of the end user in mind. In those cases, SAS authentication will remain the gold standard for the prudent user.

Even without SIP integrity protection, the Media Security Requirements [[I-D.ietf-sip-media-security-requirements](#)] R-ACT-ACT requirement can be met by ZRTP's SAS mechanism. Although ZRTP may benefit from an integrity-protected SIP layer, it is fortunate that ZRTP's self-contained MiTM defenses do not actually require an integrity-protected SIP layer. ZRTP can bypass the delays and problems that SIP integrity faces, such as E.164 number usage, and the complexity of building and maintaining a PKI.

In contrast, DTLS-SRTP [[I-D.ietf-avt-dtls-srtp](#)] appears to depend heavily on end-to-end integrity protection in the SIP layer. Further, DTLS-SRTP must bear the additional cost of a signature calculation of its own, in addition to the signature calculation the SIP layer uses to achieve its integrity protection. ZRTP needs no signature calculation of its own to leverage the signature calculation carried out in the SIP layer.

[8.2.](#) Deriving the SRTP secret (srtps) from the signaling layer

The shared secret calculations defined in [Section 4.3](#) make use of the SRTP secret (srtps), if it is provided by the signaling layer.

It is desirable for only one SRTP key negotiation protocol to be

used, and that protocol should be ZRTP. But in the event the signaling layer negotiates its own SRTP master key and salt, using the SDES [[RFC4568](#)] or [[RFC4567](#)], it can be passed from the signaling to the ZRTP layer and mixed into ZRTP's own shared secret calculations, without compromising security by creating a dependency on the signaling for media encryption.

ZRTP computes srtps from the SRTP master key and salt parameters provided by the signaling layer in this manner:

```
srtps = hash(SRTP master key | SRTP master salt)
```

It is expected that the srtps parameter will be rarely computed or used in typical ZRTP endpoints, because it is likely and desirable that ZRTP will be the sole means of negotiating SRTP keys, needing no help from SDES [[RFC4568](#)] or [[RFC4567](#)]. If srtps is computed, it will be stored in the auxiliary shared secret auxsecret, defined in [Section 4.3](#), and used in [Section 4.3.1](#) and [Section 4.3.2](#).

[8.3.](#) Codec Selection for Secure Media

Codec selection is negotiated in the signaling layer. If the signaling layer determines that ZRTP is supported by both endpoints, this should provide guidance in codec selection to avoid variable bit-rate (VBR) codecs that leak information.

When voice is compressed with a VBR codec, the packet lengths vary depending on the types of sounds being compressed. This leaks a lot of information about the content even if the packets are encrypted, regardless of what encryption protocol is used [[Wright1](#)]. It is RECOMMENDED that VBR codecs be avoided in encrypted calls. It is not a problem if the codec adapts the bit rate to the available channel bandwidth. The vulnerable codecs are the ones that change their bit rate depending on the type of sound being compressed.

It also appears that voice activity detection (VAD) leaks information about the content of the conversation, but to a lesser extent than VBR. This effect can be ameliorated by lengthening the VAD hangover time by about 1 to 2 seconds, if this is feasible in your application. This is a topic that requires further study.

[9.](#) False ZRTP Packet Rejection

An attacker who is not in the media path may attempt to inject false ZRTP protocol packets, possibly to effect a denial of service attack, or to inject his own media stream into the call. VoIP by its nature invites various forms of denial of service attacks and requires

protocol features to reject such attacks. While bogus SRTP packets may be easily rejected via the SRTP auth tag field, that can only be applied after a key agreement is completed. During the ZRTP key negotiation phase, other false packet rejection mechanisms are needed. One such mechanism is the use of the total_hash in the final shared secret calculation, but that can only detect false packets after performing the computationally expensive Diffie-Hellman calculation.

The VoIP developer community expects to see a lot of denial of service attacks, especially from attackers who are not in the media

path. Such an attacker might inject false ZRTP packets to force a ZRTP endpoint to engage in an endless series of pointless and expensive DH calculations. To detect and reject false packets cheaply and rapidly as soon as they are received, ZRTP uses a hash chain, which is a series of successive hash images. Before each session, the following values are computed:

- H0 = 256-bit random nonce (different for each party)
- H1 = hash (H0)
- H2 = hash (H1)
- H3 = hash (H2)

The hash chain MUST use the hash algorithm defined in [Section 5.1.2.1](#). Each 256-bit hash image is the pre-image of the next, and the sequence of images is sent in reverse order in the ZRTP packet sequence. The hash image H3 is sent in the Hello packet, H2 is sent in the Commit packet, H1 is sent in the DHPart1 or DHPart2 packets, and H0 is sent in the Confirm1 or Confirm2 packets. The initial random H0 nonces that each party generates MUST be unpredictable to an attacker and unique within a ZRTP call, which thereby forces the derived hash images H1-H3 to also be unique and unpredictable.

The recipient checks if the packet has the correct hash pre-image, by hashing it and comparing the result with the hash image for the preceding packet. Packets which contain an incorrect hash pre-image MUST NOT be used by the recipient, but MAY be processed as security exceptions, perhaps by logging or alerting the user. As long as these bogus packets are not used, and correct packets are still being received, the protocol SHOULD be allowed to run to completion, thereby rendering ineffective this denial of service attack.

Because these hash images alone do not protect the rest of the contents of the packet they reside in, this scheme assumes the attacker cannot modify the packet contents from a legitimate party, which is a reasonable assumption for an attacker who is not in the media path. This covers an important range of denial-of-service

attacks. For dealing with the remaining set of attacks that involve packet modification, other mechanisms are used, such as the total_hash in the final shared secret calculation, and the hash commitment in the Commit packet.

False Hello packets may be detected and rejected by the mechanism defined in [Section 8.1](#). This mechanism requires that each Hello packet be unique, and the inclusion of the H3 hash image meets that requirement.

If and only if an integrity-protected signaling channel is available, this hash chaining scheme can be used to key HMACs to authenticate the entire ZRTP key exchange, and thereby prevent a MiTM attack, without relying on the users verbally comparing the SAS. See [Section 8.1.1](#) for details.

Some ZRTP user agents allow the user to manually switch to clear mode (via the GoClear packet) in the middle of a secure call, and then later initiate secure mode again. Many consumer client products will omit this feature, but those that allow it may return to secure mode again in the same media stream. Although the same chain of hash images will be re-used and thus rendered ineffective the second time, no real harm is done because the new SRTP session keys will be derived in part from a cached shared secret, which was safely protected from the MiTM in the previous DH exchange earlier in the same call.

[10](#). Intermediary ZRTP Devices

This section discusses the operation of a ZRTP endpoint which is actually an intermediary. For example, consider a device which proxies both signaling and media between endpoints. There are three possible ways in which such a device could support ZRTP.

An intermediary device can act transparently to the ZRTP protocol. To do this, a device **MUST** pass RTP header extensions and payloads (to allow the ZRTP Flag) and non-RTP protocols multiplexed on the same port as RTP (to allow ZRTP and STUN). This is the **RECOMMENDED** behavior for intermediaries as ZRTP and SRTP are best when done end-to-end.

An intermediary device could implement the ZRTP protocol and act as a ZRTP endpoint on behalf of non-ZRTP endpoints behind the intermediary device. The intermediary could determine on a call-by-call basis whether the endpoint behind it supports ZRTP based on the presence or absence of the ZRTP SDP attribute flag (`a=zrtp-hash`). For non-ZRTP endpoints, the intermediary device could act as the ZRTP endpoint

using its own ZID and cache. This approach SHOULD only be used when there is some other security method protecting the confidentiality of the media between the intermediary and the inside endpoint, such as IPSec or physical security.

The third mode, which is NOT RECOMMENDED, is for the intermediary device to attempt to back-to-back the ZRTP protocol. The only exception to this case is where the intermediary device is a trusted element providing services to one of the endpoints - e.g. a Private Branch Exchange or PBX. In this mode, the intermediary would attempt to act as a ZRTP endpoint towards both endpoints of the media session. This approach MUST NOT be used except as described in [Section 7.3](#) as it will always result in a detected man-in-the-middle attack and will generate alarms on both endpoints and likely result in the immediate termination of the session.

In cases where centralized media mixing is taking place, the SAS will not match when compared by the humans. However, this situation is known in the SIP signaling by the presence of the isfocus feature tag [[RFC4579](#)]. As a result, when the isfocus feature tag is present, the DH exchange can be authenticated by the mechanism defined in [Section 8.1.1](#) or by validating signatures ([Section 7.2](#)) in the Confirm or SASrelay messages. For example, consider a audio conference call with three participants Alice, Bob, and Carol hosted on a conference bridge in Dallas. There will be three ZRTP encrypted media streams, one encrypted stream between each participant and Dallas. Each will have a different SAS. Each participant will be able to validate their SAS with the conference bridge by using signatures optionally present in the Confirm messages (described in [Section 7.2](#)). Or, if the signaling path has end-to-end integrity protection, each DH exchange will have automatic MiTM protection by using the mechanism in [Section 8.1.1](#).

SIP feature tags can also be used to detect if a session is established with an automaton such as an IVR, voicemail system, or speech recognition system. The display of SAS strings to users should be disabled in these cases.

It is possible that an intermediary device acting as a ZRTP endpoint might still receive ZRTP Hello and other messages from the inside endpoint. This could occur if there is another inline ZRTP device which does not include the ZRTP SDP attribute flag. An intermediary acting as a ZRTP endpoint receiving ZRTP Hello and other messages from the inside endpoint MUST NOT pass these ZRTP messages.

Internet-Draft

ZRTP

October 2008

[11.](#) The ZRTP Disclosure flag

There are no back doors defined in the ZRTP protocol specification. The designers of ZRTP would like to discourage back doors in ZRTP-enabled products. However, despite the lack of back doors in the actual ZRTP protocol, it must be recognized that a ZRTP implementer might still deliberately create a rogue ZRTP-enabled product that implements a back door outside the scope of the ZRTP protocol. For example, they could create a product that discloses the SRTP session key generated using ZRTP out-of-band to a third party. They may even have a legitimate business reason to do this for some customers.

For example, some environments have a need to monitor or record calls, such as stock brokerage houses who want to discourage insider trading, or special high security environments with special needs to monitor their own phone calls. We've all experienced automated messages telling us that "This call may be monitored for quality assurance". A ZRTP endpoint in such an environment might unilaterally disclose the session key to someone monitoring the call. ZRTP-enabled products that perform such out-of-band disclosures of the session key can undermine public confidence in the ZRTP protocol, unless we do everything we can in the protocol to alert the other user that this is happening.

If one of the parties is using a product that is designed to disclose their session key, ZRTP requires them to confess this fact to the other party through a protocol message to the other party's ZRTP client, which can properly alert that user, perhaps by rendering it in a graphical user interface. The disclosing party does this by sending a Disclosure flag (D) in Confirm1 and Confirm2 messages as described in [Section 5.7](#).

Note that the intention here is to have the Disclosure flag identify products that are designed to disclose their session keys, not to identify which particular calls are compromised on a call-by-call basis. This is an important legal distinction, because most government sanctioned wiretap regulations require a VoIP service provider to not reveal which particular calls are wiretapped. But there is nothing illegal about revealing that a product is designed to be wiretap-friendly. The ZRTP protocol mandates that such a

product "out" itself.

You might be using a ZRTP-enabled product with no back doors, but if your own graphical user interface tells you the call is (mostly) secure, except that the other party is using a product that is designed in such a way that it may have disclosed the session key for monitoring purposes, you might ask him what brand of secure telephone he is using, and make a mental note not to purchase that brand

yourself. If we create a protocol environment that requires such back-doored phones to confess their nature, word will spread quickly, and the "invisible hand" of the free market will act. The free market has effectively dealt with this in the past.

Of course, a ZRTP implementer can lie about his product having a back door, but the ZRTP standard mandates that ZRTP-compliant products MUST adhere to the requirement that a back door be confessed by sending the Disclosure flag to the other party.

There will be inevitable comparisons to Steve Bellovin's 2003 April fool's joke, when he submitted [RFC 3514](#) [[RFC3514](#)] which defined the "Evil bit" in the IPV4 header, for packets with "evil intent". But we submit that a similar idea can actually have some merit for securing VoIP. Sure, one can always imagine that some implementer will not be fazed by the rules and will lie, but they would have lied anyway even without the Disclosure flag. There are good reasons to believe that it will improve the overall percentage of implementations that at least tell us if they put a back door in their products, and may even get some of them to decide not to put in a back door at all. From a civic hygiene perspective, we are better off with having the Disclosure flag in the protocol.

If an endpoint stores or logs SRTP keys or information that can be used to reconstruct or recover SRTP keys after they are no longer in use (i.e. the session is active), or otherwise discloses or passes SRTP keys or information that can be used to reconstruct or recover SRTP keys to another application or device, the Disclosure flag D MUST be set in the Confirm1 or Confirm2 message.

11.1. Guidelines on Proper Implementation of the Disclosure Flag

Some implementers have asked for guidance on implementing the

Disclosure Flag. Some people have incorrectly thought that a connection secured with ZRTP cannot be used in a call center, with voluntary voice recording, or even with a voicemail system. Similarly, some potential users of ZRTP have over considered the protection that ZRTP can give them. These guidelines clarify both concerns.

The ZRTP Disclosure Flag only governs the ZRTP/SRTP stream itself. It does not govern the underlying RTP media stream, nor the actual media itself. Consequently, a PBX that uses ZRTP may provide conference calls, call monitoring, call recording, voicemail, or other PBX features and still say that it does not disclose the ZRTP key material. A video system may provide DVR features and still say that it does not disclose the ZRTP key material. The ZRTP Disclosure Flag, when not set, means only that the ZRTP cryptographic key

material stays within the bounds of the ZRTP subsystem.

If an application has a need to disclose the ZRTP cryptographic key material, the easiest way to comply with the protocol is to set the flag to the proper value. The next easiest way is to overestimate disclosure. For example, a call center that commonly records calls might choose to set the disclosure flag even though all recording is an analog recording of a call (and thus outside the ZRTP scope) because it sets an expectation with clients that their calls might be recorded.

Note also that the ZRTP Disclosure Flag does not require an implementation to preclude hacking or malware. Malware that leaks ZRTP cryptographic key material does not create a liability for the implementor from non-compliance with the ZRTP specification.

A user of ZRTP should note that ZRTP is not a panacea against unauthorized recording. ZRTP does not and cannot protect against an untrustworthy partner who holds a microphone up to the speaker. It does not protect against someone else being in the room. It does not protect against analog wiretaps in the phone or in the room. It does not mean your partner has not been hacked with spyware. It does not mean that the software has no flaws. It means that the ZRTP subsystem is not knowingly leaking ZRTP cryptographic key material.

12. RTP Header Extension Flag for ZRTP

This specification defines a new RTP header extension used only for discovery of support for ZRTP. No ZRTP data is transported in the extension. When used, the X bit is set in the RTP header to indicate the presence of the RTP header extension.

[Section 5.3.1 in RFC 3550](#) [[RFC3550](#)] defines the format of an RTP Header extension. The Header extension is appended to the RTP header. The first 16 bits are an identifier for the header extension, and the following 16 bits are length of the extension header in 32 bit words. The ZRTP flag RTP header extension has the value of 0x505A and a length of 0. The format of the header extension is as shown in the Figure below.




Figure 18: RTP Extension header format for ZRTP Flag

ZRTP endpoints MAY include the ZRTP Flag in RTP packets sent at the start of a session. For example, an endpoint may decide to include the flag in the first 2 seconds of RTP packets sent. The inclusion of the flag MAY be ended if a ZRTP message (such as Hello) is received.

13. IANA Considerations

This specification defines a new SDP [[RFC4566](#)] attribute in [Section 8](#).

Contact name: Philip Zimmermann <prz@mit.edu>

Attribute name: "zrtp-hash".

Type of attribute: Media level.

Subject to charset: Not.

Purpose of attribute: The 'zrtp-hash' indicates that a UA supports the ZRTP protocol and provides a hash of the ZRTP Hello message. The ZRTP protocol version number is also specified.

Allowed attribute values: Hex.

14. Appendix - Media Security Requirements

This section discusses how ZRTP meets all RTP security requirements discussed in the Media Security Requirements [[I-D.ietf-sip-media-security-requirements](#)] document without any dependencies on other protocols or extensions, unlike DTLS-SRTP [[I-D.ietf-avt-dtls-srtp](#)] which requires additional protocols and mechanisms.

R-FORK-RETARGET is met since ZRTP is a media path key agreement protocol.

R-DISTINCT is met since ZRTP uses ZIDs and allows multiple independent ZRTP exchanges to proceed.

R-REUSE is met using the Multistream and Preshared modes.

R-AVOID-CLIPPING is met since ZRTP is a media path key agreement protocol

R-RTP-VALID is met since the ZRTP packet format does not pass the RTP validity check

R-ASSOC is met using the a=zrtp-hash SDP attribute in INVITEs and responses.

R-NEGOTIATE is met using the Commit message.

R-PSTN is met since ZRTP can be implemented in Gateways.

R-PFS is met using ZRTP Diffie-Hellman key agreement methods.

R-COMPUTE is met using the Hello/Commit ZRTP exchange.

R-CERTS is met using the optional signature field in ZRTP Confirm messages.

R-FIPS is met since ZRTP uses algorithms that allow FIPS certification.

R-DOS is met since ZRTP does not introduce any new denial of service attacks.

R-EXISTING is met since ZRTP can support the use of certificates or keys.

R-AGILITY is met since the set of hash, cipher, authentication tag length, key agreement method, SAS type, and signature type can all be extended and negotiated.

R-DOWNGRADE is met since ZRTP has protection against downgrade attacks.

R-PASS-MEDIA is met since ZRTP prevents a passive adversary with access to the media path from gaining access to keying material used to protect SRTP media packets.

R-PASS-SIG is met since ZRTP prevents a passive adversary with access to the signaling path from gaining access to keying material used to protect SRTP media packets.

R-SIG-MEDIA is met using the a=zrtp-hash SDP attribute in INVITES and responses.

R-ID-BINDING is met using the a=zrtp-hash SDP attribute.

R-ACT-ACT is met using the a=zrtp-hash SDP attribute in INVITES and responses.

R-BEST-SECURE is met since ZRTP utilizes the RTP/AVP profile and

hence best effort SRTP in every case.

R-OTHER-SIGNALING is met since ZRTP can utilize modes in which there is no dependency on the signaling path.

R-RECORDING is met using the ZRTP Disclosure flag.

R-TRANSCODER is met if the transcoder operates as a trusted MitM (i.e. a PBX).

R-ALLOW-RTP is met due to ZRTP's best effort encryption.

15. Security Considerations

This document is all about securely keying SRTP sessions. As such, security is discussed in every section.

Most secure phones rely on a Diffie-Hellman exchange to agree on a common session key. But since DH is susceptible to a man-in-the-middle (MitM) attack, it is common practice to provide a way to authenticate the DH exchange. In some military systems, this is done by depending on digital signatures backed by a centrally-managed PKI. A decade of industry experience has shown that deploying centrally managed PKIs can be a painful and often futile experience. PKIs are just too messy, and require too much activation energy to get them started. Setting up a PKI requires somebody to run it, which is not practical for an equipment provider. A service provider like a carrier might venture down this path, but even then you have to deal with cross-carrier authentication, certificate revocation lists, and other complexities. It is much simpler to avoid PKIs altogether, especially when developing secure commercial products. It is therefore more common for commercial secure phones in the PSTN world to augment the DH exchange with a Short Authentication String (SAS) combined with a hash commitment at the start of the key exchange, to shorten the length of SAS material that must be read aloud. No PKI is required for this approach to authenticating the DH exchange. The AT&T TSD 3600, Eric Blossom's COMSEC secure phones [[comsec](#)], PGPfone [[pgpfone](#)], and CryptoPhone [[cryptophone](#)] are all examples of products that took this simpler lightweight approach.

The main problem with this approach is inattentive users who may not execute the voice authentication procedure, or unattended secure

phone calls to answering machines that cannot execute it.

Additionally, some people worry about voice spoofing. But it is a mistake to think this is simply an exercise in voice impersonation (perhaps this could be called the "Rich Little" attack). Although there are digital signal processing techniques for changing a person's voice, that does not mean a man-in-the-middle attacker can safely break into a phone conversation and inject his own short authentication string (SAS) at just the right moment. He doesn't know exactly when or in what manner the users will choose to read aloud the SAS, or in what context they will bring it up or say it, or even which of the two speakers will say it, or if indeed they both will say it. In addition, some methods of rendering the SAS involve using a list of words such as the PGP word list [Juola2], in a manner analogous to how pilots use the NATO phonetic alphabet to convey information. This can make it even more complicated for the attacker, because these words can be worked into the conversation in unpredictable ways. Remember that the attacker places a very high value on not being detected, and if he makes a mistake, he doesn't get to do it over. Some people have raised the question that even if the attacker lacks voice impersonation capabilities, it may be unsafe for people who don't know each other's voices to depend on the SAS procedure. This is not as much of a problem as it seems, because it isn't necessary that they recognize each other by their voice, it is only necessary that they detect that the voice used for the SAS procedure matches the voice in the rest of the phone conversation.

A popular and field-proven approach is used by SSH (Secure Shell) [RFC4251], which Peter Gutmann likes to call the "baby duck" security model. SSH establishes a relationship by exchanging public keys in the initial session, when we assume no attacker is present, and this makes it possible to authenticate all subsequent sessions. A successful MiTM attacker has to have been present in all sessions all the way back to the first one, which is assumed to be difficult for the attacker. ZRTP's key continuity features are actually better than SSH, at least for VoIP, for reasons described in [Section 15.1](#). All this is accomplished without resorting to a centrally-managed PKI.

We use an analogous baby duck security model to authenticate the DH exchange in ZRTP. We don't need to exchange persistent public keys, we can simply cache a shared secret and re-use it to authenticate a long series of DH exchanges for secure phone calls over a long period of time. If we read aloud just one SAS, and then cache a shared secret for later calls to use for authentication, no new voice authentication rituals need to be executed. We just have to remember we did one already.

Internet-Draft

ZRTP

October 2008

If one party ever loses this cached shared secret, it is no longer available for authentication of DH exchanges. This cache mismatch situation is easy to detect by the party that still has a surviving shared secret cache entry. If it fails to match, either there is a MiTM attack or one side has lost their shared secret cache entry. The user agent that discovers the cache mismatch must alert the user that a cache mismatch has been detected, and that he must do a verbal comparison of the SAS to distinguish if the mismatch is because of a MiTM attack or because of the other party losing her cache. From that point on, the two parties start over with a new cached shared secret. Then they can go back to omitting the voice authentication on later calls.

A particularly compelling reason why this approach is attractive is that SAS is easiest to implement when a graphical user interface or some sort of display is available, which raises the question of what to do when a display is less conveniently available. For example, some devices that implement ZRTP might have a graphical user interface that is only visible through a web browser, such as a PBX or some other nearby device that implements ZRTP as a "bump-in-the-wire". If we take an approach that greatly reduces the need for a SAS in each and every call, we can operate in products without a graphical user interface with greater ease. Then the SAS can be compared less frequently through a web browser, or it might even be presented as needed to the local user through a locally generated voice prompt, which the local user hears and verbally repeats and compares with the remote party. Using a voice prompt in this way is purely for the local ZRTP user agent to render the SAS to the local user, and is not to be confused with the verbal comparison of the SAS between two human users.

It is a good idea to force your opponent to have to solve multiple problems in order to mount a successful attack. Some examples of widely differing problems we might like to present him with are: Stealing a shared secret from one of the parties, being present on the very first session and every subsequent session to carry out an active MiTM attack, and solving the discrete log problem. We want to force the opponent to solve more than one of these problems to succeed.

ZRTP can use different kinds of shared secrets. Each type of shared secret is determined by a different method. All of the shared

secrets are hashed together to form a session key to encrypt the call. An attacker must defeat all of the methods in order to determine the session key.

First, there is the shared secret determined entirely by a Diffie-Hellman key agreement. It changes with every call, based on random

numbers. An attacker may attempt a classic DH MiTM attack on this secret, but we can protect against this by displaying and reading aloud an SAS, combined with adding a hash commitment at the beginning of the DH exchange.

Second, there is an evolving shared secret, or ongoing shared secret that is automatically changed and refreshed and cached with every new session. We will call this the cached shared secret, or sometimes the retained shared secret. Each new image of this ongoing secret is a non-invertable function of its previous value and the new secret derived by the new DH agreement. It is possible that no cached shared secret is available, because there were no previous sessions to inherit this value from, or because one side loses its cache.

There are other approaches for key agreement for SRTP that compute a shared secret using information in the signaling. For example, [\[RFC4567\]](#) describes how to carry a MIKEY (Multimedia Internet KEYing) [\[RFC3830\]](#) payload in SDP [\[RFC4566\]](#). Or [RFC 4568](#) (SDS) [\[RFC4568\]](#) describes directly carrying SRTP keying and configuration information in SDP. ZRTP does not rely on the signaling to compute a shared secret, but if a client does produce a shared secret via the signaling, and makes it available to the ZRTP protocol, ZRTP can make use of this shared secret to augment the list of shared secrets that will be hashed together to form a session key. This way, any security weaknesses that might compromise the shared secret contributed by the signaling will not harm the final resulting session key.

The shared secret provided by the signaling (if available), the shared secret computed by DH, and the cached shared secret are all hashed together to compute the session key for a call. If the cached shared secret is not available, it is omitted from the hash computation. If the signaling provides no shared secret, it is also omitted from the hash computation.

No DH MiTM attack can succeed if the ongoing shared secret is available to the two parties, but not to the attacker. This is because the attacker cannot compute a common session key with either party without knowing the cached secret component, even if he correctly executes a classic DH MiTM attack.

[15.1.](#) Self-healing Key Continuity Feature

The key continuity features of ZRTP are analogous to those provided by SSH (Secure Shell) [[RFC4251](#)], but they differ in one respect. SSH caches public signature keys that never change, and uses a permanent private signature key that must be guarded from disclosure. If someone steals your SSH private signature key, they can impersonate

you in all future sessions and mount a successful MiTM attack any time they want.

ZRTP caches symmetric key material used to compute secret session keys, and these values change with each session. If someone steals your ZRTP shared secret cache, they only get one chance to mount a MiTM attack, in the very next session. If they miss that chance, the retained shared secret is refreshed with a new value, and the window of vulnerability heals itself, which means they are locked out of any future opportunities to mount a MiTM attack. This gives ZRTP a "self-healing" feature if any cached key material is compromised.

A MiTM attacker must always be in the media path. This presents a significant operational burden for the attacker in many VoIP usage scenarios, because being in the media path for every call is often harder than being in the signaling path. This will likely create coverage gaps in the attacker's opportunities to mount a MiTM attack. ZRTP's self-healing key continuity features are better than SSH at exploiting any temporary gaps in MiTM attack coverage. Thus, ZRTP quickly recovers from any disclosure of cached key material.

The infamous Debian OpenSSL weak key vulnerability [[dsa-1571](#)] (discovered and patched in May 2008) offers a real-world example of why ZRTP's self-healing scheme is a good way to do key continuity. The Debian bug resulted in the production of a lot of weak SSH (and TLS/SSL) keys, which continued to compromise security even after the bug had been patched. In contrast, ZRTP's key continuity scheme adds new entropy to the cached key material with every call, so old

deficiencies in entropy are washed away with each new session.

It should be noted that the addition of shared secret entropy from previous sessions can extend the strength of the new session key to AES-256 levels, even if the new session uses Diffie-Hellman keys no larger than DH-3072 or ECDH-256, provided the cached shared secrets were initially established when the wiretapper was not present. This is why AES-256 MAY be used with the smaller DH key sizes in [Section 5.1.5](#).

Caching shared symmetric key material is also less CPU intensive compared with using digital signatures, which may be important for low-power mobile platforms.

[16.](#) Acknowledgments

The authors would like to thank Bryce Wilcox-O'Hearn and Colin Plumb for their contributions to the design of this protocol, and to thank Hal Finney, Viktor Krikun, Werner Dittmann, Jon Peterson, Dan Wing,

Zimmermann, et al.

Expires April 28, 2009

[Page 84]

Internet-Draft

ZRTP

October 2008

Sagar Pai, Colin Perkins, David McGrew, and Roni Even for their helpful comments and suggestions.

The use of hash chains to key HMACs in ZRTP is similar to Adrian Perrig's TESLA protocol [[TESLA](#)].

[17.](#) References

[17.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", [RFC 3711](#), March 2004.

[RFC3526] Kivinen, T. and M. Kojo, "More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)", [RFC 3526](#), May 2003.

[RFC3309] Stone, J., Stewart, R., and D. Otis, "Stream Control Transmission Protocol (SCTP) Checksum Change", [RFC 3309](#), September 2002.

[RFC4231] Nystrom, M., "Identifiers and Test Vectors for HMAC-SHA-224, HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512", [RFC 4231](#), December 2005.

[SP800-90]
Barker, E. and J. Kelsey, "Recommendation for Random Number Generation Using Deterministic Random Bit Generators", NIST Special Publication 800-90 (Revised) March 2007.

[SP800-56A]
Barker, E., Johnson, D., and M. Smid, "Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography", NIST Special Publication 800-56A Revision 1, March 2007.

[FIPS-180-2]
"Secure Hash Signature Standard (SHS)", NIST FIPS PUB 180-

2 August 2002.

[FIPS-198-1]
"The Keyed-Hash Message Authentication Code (HMAC)", NIST FIPS PUB 198-1 July 2008.

[NSA-Suite-B]
"Fact Sheet NSA Suite B Cryptography", NSA Information Assurance Directorate Fact Sheet NSA Suite B.

[RFC4753] Fu, D. and J. Solinas, "ECP Groups For IKE and IKEv2", [RFC 4753](#), January 2007.

[FIPS-186-3]

"Digital Signature Standard (DSS)", NIST FIPS PUB 186-3 Draft, March 2006.

[SP800-38A]

Dworkin, M., "Recommendation for Block Cipher: Methods and Techniques", NIST Special Publication 800-38A 2001 Edition.

[z-base-32]

Wilcox, B., "Human-oriented base-32 encoding", <http://zooko.com/repos/z-base-32/base32/DESIGN> .

[pgpwordlist]

"PGP Words", http://en.wikipedia.org/wiki/PGP_Words .

[RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", [RFC 4566](#), July 2006.

17.2. Informative References

[I-D.ietf-sip-media-security-requirements]

Wing, D., Fries, S., Tschofenig, H., and F. Audet, "Requirements and Analysis of Media Security Management Protocols", [draft-ietf-sip-media-security-requirements-07](#) (work in progress), June 2008.

[Ferguson]

Ferguson, N. and B. Schneier, "Practical Cryptography", Wiley Publishing 2003.

[RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", [BCP 106](#), [RFC 4086](#), June 2005.

[Juola1] Juola, P. and P. Zimmermann, "Whole-Word Phonetic

Zimmermann, et al.

Expires April 28, 2009

[Page 86]

Internet-Draft

ZRTP

October 2008

Distances and the PGPfone Alphabet", Proceedings of the International Conference of Spoken Language Processing (ICSLP-96) 1996.

[Juola2] Juola, P., "Isolated Word Confusion Metrics and the PGPfone Alphabet", Proceedings of New Methods in Language Processing 1996.

- [pgpfone] Zimmermann, P., "PGPfone",
<http://philzimmermann.com/docs/pgpfone10b7.pdf> .
- [zfone] Zimmermann, P., "Zfone",
<http://www.philzimmermann.com/zfone> .
- [Byzantine]
"The Two Generals' Problem",
http://en.wikipedia.org/wiki/Two_Generals%27_Problem .
- [TESLA] Perrig, A., Canetti, R., Tygar, J., and D. Song, "The TESLA Broadcast Authentication Protocol", <http://www.ece.cmu.edu/~adrian/projects/tesla-cryptobytes/tesla-cryptobytes.pdf> .
- [SHA-3] "Cryptographic Hash Algorithm Competition", NIST Computer Security Resource Center Cryptographic Hash Project.
- [comsec] Blossom, E., "The VP1 Protocol for Voice Privacy Devices Version 1.2", <http://www.comsec.com/vp1-protocol.pdf> .
- [cryptophone]
"CryptoPhone", <http://www.cryptophone.de/> .
- [Wright1] Wright, C., Ballard, L., Coull, S., Monroe, F., and G. Masson, "Spot me if you can: Uncovering spoken phrases in encrypted VoIP conversations", Proceedings of the 2008 IEEE Symposium on Security and Privacy 2008.
- [dsa-1571]
"Debian Security Advisory - OpenSSL predictable random number generator",
<http://www.debian.org/security/2008/dsa-1571> .
- [I-D.ietf-avt-srtp-big-aes]
McGrew, D., "The use of AES-192 and AES-256 in Secure RTP", <http://www1.tools.ietf.org/html/draft-ietf-avt-srtp-big-aes> .
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston,

- A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", [RFC 3261](#), June 2002.
- [RFC4251] Ylonen, T. and C. Lonvick, "The Secure Shell (SSH) Protocol Architecture", [RFC 4251](#), January 2006.
- [RFC4568] Andreasen, F., Baugher, M., and D. Wing, "Session Description Protocol (SDP) Security Descriptions for Media Streams", [RFC 4568](#), July 2006.
- [RFC4567] Arkko, J., Lindholm, F., Naslund, M., Norrman, K., and E. Carrara, "Key Management Extensions for Session Description Protocol (SDP) and Real Time Streaming Protocol (RTSP)", [RFC 4567](#), July 2006.
- [RFC3830] Arkko, J., Carrara, E., Lindholm, F., Naslund, M., and K. Norrman, "MIKEY: Multimedia Internet KEYing", [RFC 3830](#), August 2004.
- [RFC3514] Bellovin, S., "The Security Flag in the IPv4 Header", [RFC 3514](#), April 1 2003.
- [RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", [RFC 4474](#), August 2006.
- [I-D.ietf-mmusic-ice]
Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols",
[draft-ietf-mmusic-ice-19](#) (work in progress), October 2007.
- [RFC4579] Johnston, A. and O. Levin, "Session Initiation Protocol (SIP) Call Control - Conferencing for User Agents",
[BCP 119](#), [RFC 4579](#), August 2006.
- [I-D.wing-sip-identity-media]
Wing, D. and H. Kaplan, "SIP Identity using Media Path",
[draft-wing-sip-identity-media-02](#) (work in progress),
February 2008.
- [RFC3824] Peterson, J., Liu, H., Yu, J., and B. Campbell, "Using E.164 numbers with the Session Initiation Protocol (SIP)",
[RFC 3824](#), June 2004.
- [I-D.ietf-avt-dtls-srtp]
McGrew, D. and E. Rescorla, "Datagram Transport Layer

Internet-Draft

ZRTP

October 2008

Security (DTLS) Extension to Establish Keys for Secure
Real-time Transport Protocol (SRTP)",
[draft-ietf-avt-dtls-srtp-05](#) (work in progress),
September 2008.

Authors' Addresses

Philip Zimmermann
Zfone Project

Email: prz@mit.edu

Alan Johnston (editor)
Avaya
St. Louis, MO 63124

Email: alan@sipstation.com

Jon Callas
PGP Corporation

Email: jon@pgp.com

Internet-Draft

Z RTP

October 2008

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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, THE IETF TRUST 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.