Internet Engineering Task Force (IETF)
Request for Comments: 6714
Category: Standards Track
ISSN: 2070-1721

C. Holmberg
S. Blau
Ericsson
E. Burger
Georgetown University

August 2012

Connection Establishment for Media Anchoring (CEMA) for the Message Session Relay Protocol (MSRP)

Abstract

This document defines a Message Session Relay Protocol (MSRP) extension, Connection Establishment for Media Anchoring (CEMA). Support of this extension is OPTIONAL. The extension allows middleboxes to anchor the MSRP connection, without the need for middleboxes to modify the MSRP messages; thus, it also enables secure end-to-end MSRP communication in networks where such middleboxes are deployed. This document also defines a Session Description Protocol (SDP) attribute, 'msrp-cema', that MSRP endpoints use to indicate support of the CEMA extension.

Status of This Memo

This is an Internet Standards Track document.

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc6714.

Table of Contents

1. Introduction

The Message Session Relay Protocol (MSRP) [RFC4975] expects to use MSRP relays [RFC4976] as a means for Network Address Translation (NAT) traversal and policy enforcement. However, many Session Initiation Protocol (SIP) [RFC3261] networks, which deploy MSRP, contain middleboxes. These middleboxes anchor and control media; perform tasks such as NAT traversal, performance monitoring, and address domain bridging; interconnect Service Level Agreement (SLA) policy enforcement; and so on. One example is the Interconnection

Border Control Function (IBCF) [GPP23228], defined by the 3rd Generation Partnership Project (3GPP). The IBCF controls a media relay that handles all types of SIP session media, such as voice, video, MSRP, etc.

MSRP, as defined in RFC 4975 [RFC4975] and RFC 4976 [RFC4976], cannot anchor through middleboxes. The reason is that MSRP messages have routing information embedded in the message. Without an extension such as CEMA, middleboxes must read the message to change the routing information. This occurs because middleboxes modify the address:port information in the Session Description Protocol (SDP) [RFC4566] c/m-line in order to anchor media. An "active" [RFC6135] MSRP User Agent (UA) establishes the MSRP TCP or Transport Layer Security (TLS) connection based on the MSRP URI of the SDP 'path' attribute. This means that the MSRP connection will not be routed through the middlebox unless the middlebox also modifies the MSRP URI of the topmost SDP 'path' attribute. In many scenarios, this will prevent the MSRP connection from being established. In addition, if the middlebox modifies the MSRP URI of the SDP 'path' attribute, then the MSRP URI comparison procedure [RFC4975], which requires consistency between the address information in the MSRP messages and the address information carried in the MSRP URI of the SDP 'path' attribute, will fail.

The only way to achieve interoperability in this situation is for the middlebox to act as an MSRP back-to-back User Agent (B2BUA). Here, the MSRP B2BUA acts as the endpoint for the MSRP signaling and media, performs the corresponding modification in the associated MSRP messages, and originates a new MSRP session toward the actual remote endpoint. However, the enabling of MSRP B2BUA functionality requires substantially more resource usage in the middlebox, which normally results in a negative impact on performance. In addition, the MSRP message needs to be exposed in cleartext to the MSRP B2BUA, which violates the end-to-end principle [RFC3724].

This specification defines an MSRP extension, Connection Establishment for Media Anchoring (CEMA). In most cases, CEMA allows MSRP endpoints to communicate through middleboxes as defined in Section 2, without a need for the middleboxes to be MSRP B2BUAs. In such cases, middleboxes that want to anchor the MSRP connection simply modify the SDP c/m-line address information, similar to what the middleboxes do for non-MSRP media types. MSRP endpoints that support the CEMA extension will use the SDP c/m-line address information for establishing the TCP or TLS connection for sending and receiving MSRP messages.

The CEMA extension is backward compatible, meaning that CEMA-enabled MSRP endpoints can communicate with non-CEMA-enabled endpoints. In scenarios where MSRP endpoints do not support the CEMA extension, an MSRP endpoint that supports the CEMA extension behaves in the same way as an MSRP endpoint that does not support it. The CEMA extension only provides an alternative mechanism for negotiating and providing address information for the MSRP TCP connection. After the creation of the MSRP connection, an MSRP endpoint that supports the CEMA extension acts according to the procedures for creating MSRP messages, performing checks when receiving MSRP messages defined in RFC 4975 and, when it is using a relay for MSRP communications, RFC 4976.

2. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, RFC 2119 [RFC2119].

Definitions:

Fingerprint-Based TLS Authentication: An MSRP endpoint that uses a self-signed certificate and sends a fingerprint (i.e., a hash of the self-signed certificate) in SDP to the other MSRP endpoint. This fingerprint binds the TLS key exchange to the signaling plane and authenticates the other endpoint based on trust in the signaling plane.

Name-Based TLS Authentication: An MSRP endpoint that uses a certificate that is bound to the endpoint’s hostname or SIP address of record. In the TLS session setup, the other MSRP endpoint verifies that the identity associated with the certificate corresponds to that of the peer (as indicated in SIP/SDP) and that the binding of the identity to the public key was done by a party that the endpoint trusts. This definition includes traditional certificates issued by a well-known certification authority as well as self-signed certificates published via the SIP Certificate Management Service [RFC6072] and other similar mechanisms.

B2BUA: This is an abbreviation for back-to-back user agent.

MSRP B2BUA: A network element that terminates an MSRP connection from one MSRP endpoint and reoriginates that connection toward another MSRP endpoint. Note that the MSRP B2BUA is distinct from a SIP B2BUA. A SIP B2BUA terminates a SIP session and reoriginates that session toward another SIP endpoint. In the

context of MSRP, a SIP endpoint initiates a SIP session toward another SIP endpoint. However, that INVITE may go through, for example, an outbound proxy or inbound proxy to route to the remote SIP endpoint. As part of that SIP session, an MSRP session that may follow the SIP session path is negotiated. However, there is no requirement to co-locate the SIP network elements with the MSRP network elements.

TLS B2BUA: A network element that terminates security associations (SAs) from endpoints and establishes separate SAs between itself and each endpoint.

Middlebox: A SIP network device that modifies SDP media address:port information in order to steer or anchor media flows described in the SDP, including TCP and TLS connections used for MSRP communication, through a media proxy function controlled by the SIP endpoint. In most cases, the media proxy function relays the MSRP messages without modification, while in some circumstances it acts as an MSRP B2BUA. Other SIP-related functions -- such as those related to routing, modification of SIP information, etc. -- performed by the Middlebox, and whether it acts as a SIP B2BUA or not, are outside the scope of this document. Section 6 describes additional assumptions regarding how the Middlebox handles MSRP in order to support the extension defined in this document.

Media anchor: An entity that performs media anchoring inserts itself in the media path of a media communication session between two entities. The media anchor will receive, and forward, the media sent between the entities.

This document reuses the terms "answer", "answerer", "offer", and "offerer" as defined in [RFC3264].

3. Applicability Statement

This document defines a Message Session Relay Protocol (MSRP) extension, Connection Establishment for Media Anchoring (CEMA). Support of this extension is OPTIONAL. The extension allows Middleboxes to anchor the MSRP connection, without the need for Middleboxes to modify the MSRP messages; thus, it also enables secure end-to-end MSRP communication in networks where such Middleboxes are deployed. The document also defines a Session Description Protocol (SDP) attribute, 'msrp-cema', that MSRP endpoints use to indicate support of the CEMA extension.

The CEMA extension is primarily intended for MSRP endpoints that operate in networks in which Middleboxes that want to anchor media connections are deployed, without the need for the Middleboxes to

enable MSRP B2BUA functionality. An example of such a network is the IP Multimedia Subsystem (IMS), defined by the 3rd Generation Partnership Project (3GPP), which also has the capability for all endpoints to use name-based TLS authentication. The extension is also useful for other MSRP endpoints that operate in other networks but that communicate with MSRP endpoints in networks with such Middleboxes, unless there is a gateway between these networks that by default always enables MSRP B2BUA functionality.

This document assumes certain behaviors on the part of Middleboxes, as described in Section 6. These behaviors are not standardized. If Middleboxes do not behave as assumed, then the CEMA extension does not add any value over base MSRP behavior. MSRP endpoints that support CEMA are required to use RFC 4975 behavior in cases where they detect that the CEMA extension cannot be enabled.

4. Connection Establishment for Media Anchoring Mechanism

4.1. General

This section defines how an MSRP endpoint that supports the CEMA extension generates SDP offers and answers for MSRP, and which SDP information elements the MSRP endpoint uses when creating the TCP or TLS connection for sending and receiving MSRP messages.

Based on the procedures described in Sections 4.2 and 4.3, in the following cases the CEMA extension will not be enabled, and there will be a fallback to the MSRP connection establishment procedures defined in RFC 4975 and RFC 4976:

- A non-CEMA-enabled MSRP endpoint becomes "active" [RFC6135] (no matter whether it uses a relay for its MSRP communication or not), as it will always establish the MSRP connection using the SDP 'path' attribute, which contains the address information of the remote MSRP endpoint, instead of using the SDP c/m-line, which contains the address information of the Middlebox.

- A non-CEMA-enabled MSRP endpoint that uses a relay for its MSRP communication becomes "passive" [RFC6135], as it cannot be assumed that the MSRP endpoint inserts the address information of the relay in the SDP c/m-line.

- A CEMA-enabled MSRP endpoint that uses a relay for its MSRP communication becomes "active", since if it adds the received SDP c/m-line address information to the ToPath header field of the MSRP message (in order for the relay to establish the MSRP connection toward the Middlebox), the session matching [RFC4975] performed by the remote MSRP endpoint will fail.

4.2. MSRP SDP Offerer Procedures

When a CEMA-enabled offerer sends an SDP offer for MSRP, it generates the SDP offer according to the procedures in RFC 4975. In addition, the offerer follows RFC 4976 if it is using a relay for MSRP communication. The offerer also performs the following additions and modifications:

1. The offerer MUST include an SDP 'msrp-cema' attribute in the MSRP media description of the SDP offer.

2. If the offerer is not using a relay for MSRP communication, it MUST include an SDP 'setup' attribute in the MSRP media description of the SDP offer, according to the procedures in RFC 6135 [RFC6135].

3. If the offerer is using a relay for MSRP communication, it MUST, in addition to including the address information of the relay in the topmost SDP 'path' attribute, also include the address information of the relay, rather than its own address information, in the SDP c/m-line associated with the MSRP media description. In addition, it MUST include an SDP 'setup:actpass' attribute in the MSRP media description of the SDP offer.

When the offerer receives an SDP answer, if the MSRP media description of the SDP answer does not contain an SDP 'msrp-cema' attribute, and if any one or more of the criteria below are met, the offerer MUST fall back to RFC 4975 behavior by sending a new SDP offer according to the procedures in RFC 4975 and RFC 4976. The new offer MUST NOT contain an SDP 'msrp-cema' attribute.

1. The SDP c/m-line address information associated with the MSRP media description does not match (see Section 4.4) the information in the MSRP URI of the 'path' attribute(s) (in which case it is assumed that the SDP c/m-line contains the address of a Middlebox), and the MSRP endpoint will become "passive" (if the MSRP media description of the SDP answer contains an SDP 'setup: active' attribute).

NOTE: If an MSRP URI contains a domain name, it needs to be resolved into an IP address and port before it is checked against the SDP c/m-line address information, in order to determine whether the address information matches.

2. The offerer uses a relay for its MSRP communication, the SDP c/m-line address information associated with the MSRP media description does not match the information in the MSRP URI of the SDP 'path' attribute(s) (in which case it is assumed that the SDP

c/m-line contains the address of a Middlebox), and the offerer will become "active" (either by default or if the MSRP media description of the SDP answer contains an SDP 'setup:passive' attribute).

3. The remote MSRP endpoint, acting as an answerer, uses a relay for its MSRP communication, the SDP c/m-line address information associated with the MSRP media description does not match the information in the MSRP URI of the SDP 'path' attributes (in which case it is assumed that the SDP c/m-line contains the address of a Middlebox), and the MSRP offerer will become "active" (either by default or if the MSRP media description of the SDP answer contains an SDP 'setup:passive' attribute).

NOTE: As described in Section 6, in the absence of the SDP 'msrp-cema' attribute in the new offer, it is assumed that a Middlebox will act as an MSRP B2BUA in order to anchor MSRP media.

The offerer can send the new offer within the existing early dialog [RFC3261], or it can terminate the early dialog and establish a new dialog by sending the new offer in a new initial INVITE request.

The offerer MAY choose to terminate the session establishment if it can detect that a Middlebox acting as an MSRP B2BUA is not the desired remote MSRP endpoint.

If the answerer uses a relay for its MSRP communication, and the SDP c/m-line address information associated with the MSRP media description matches one of the SDP 'path' attributes, it is assumed that there is no Middlebox in the network. In that case, the offerer MUST fall back to RFC 4975 behavior, but it does not need to send a new SDP offer.

In other cases, where none of the criteria above are met, and where the MSRP offerer becomes "active", it MUST use the SDP c/m-line for establishing the MSRP TCP connection. If the offerer becomes "passive", it will wait for the answerer to establish the TCP connection, according to the procedures in RFC 4975.

4.3. MSRP SDP Answerer Procedures

If the MSRP media description of the SDP offer does not contain an SDP 'msrp-cema' attribute, and the SDP c/m-line address information associated with the MSRP media description does not match the information in the MSRP URI of the SDP 'path' attribute(s), the answerer MUST either reject the offered MSRP connection (by using a

zero port value number in the generated SDP answer) or reject the whole SIP request that carries the SDP offer with a 488 Not Acceptable Here [RFC3261] response.

NOTE: The reason for the rejection is that the answerer assumes that a middlebox that does not support the CEMA extension has modified the c/m-line address information of the SDP offer without enabling MSRP B2BUA functionality.

NOTE: If an MSRP URI contains a domain name, it needs to be resolved into an IP address and port before it is checked against the SDP c/m-line address information, in order to determine whether the address information matches.

If any one or more of the criteria below are met, the answerer MUST fall back to RFC 4975 behavior and generate the associated SDP answer according to the procedures in RFC 4975 and RFC 4976. The answerer MUST NOT insert an SDP 'msrp-cema' attribute in the MSRP media description of the SDP answer.

1. Both MSRP endpoints are using relays for their MSRP communication. The answerer can detect if the remote MSRP endpoint, acting as an offerer, is using a relay for its MSRP communication if the MSRP media description of the SDP offer contains multiple SDP 'path' attributes.

2. The offerer uses a relay for its MSRP communication and will become "active" (either by default or if the MSRP media description of the SDP offer contains an SDP 'setup:active' attribute). Note that a CEMA-enabled offerer would include an SDP 'setup:actpass' attribute in the SDP offer, as described in Section 4.2.

3. The answerer uses a relay for MSRP communication and is not able to become "passive" (if the MSRP media description of the offer contains an SDP 'setup:passive' attribute). Note that an offerer is not allowed to include an SDP 'setup:passive' attribute in an SDP offer, as described in RFC 6135.

In all other cases, the answerer generates the associated SDP answer according to the procedures in RFC 4975 and RFC 4976, with the following additions and modifications:

1. The answerer MUST include an SDP 'msrp-cema' attribute in the MSRP media description of the SDP answer.

2. If the answerer is not using a relay for MSRP communication, it MUST include an SDP 'setup' attribute in the MSRP media description of the answer, according to the procedures in RFC 6135.

3. If the answerer is using a relay for MSRP communication, it MUST, in addition to including the address information of the relay in the topmost SDP 'path' attribute, also include the address information of the relay, rather than its own address information, in the SDP c/m-line associated with the MSRP media description. In addition, the answerer MUST include an SDP 'setup:passive' attribute in the MSRP media description of the SDP answer.

If the answerer included an SDP 'msrp-cema' attribute in the MSRP media description of the SDP answer, and if the answerer becomes "active", it MUST use the received SDP c/m-line for establishing the MSRP TCP or TLS connection. If the answerer becomes "passive", it will wait for the offerer to establish the MSRP TCP or TLS connection, according to the procedures in RFC 4975.

4.4. Address Information Matching

When comparing address information in the SDP c/m-line and an MSRP URI, for address and port equivalence, the address and port values are retrieved in the following ways:

- SDP c/m-line address information: The IP address is retrieved from the SDP c-line, and the port from the associated SDP m-line for MSRP.

- In case the SDP c-line contains a Fully Qualified Domain Name (FQDN), the IP address is retrieved using DNS.

- MSRP URI address information: The IP address and port are retrieved from the authority part of the MSRP URI.

- In case the authority part of the MSRP URI contains an FQDN, the IP address is retrieved using DNS, according to the procedures in Section 6.2 of RFC 4975.

NOTE: According to RFC 4975, the authority part of the MSRP URI must always contain a port.

Before IPv6 addresses are compared for equivalence, they need to be converted into the same representation, using the mechanism defined in RFC 5952 [RFC5952].

NOTE: In case the DNS returns multiple records, each needs to be compared against the SDP c/m-line address information, in order to find at least one match.

NOTE: If the authority part of the MSRP URI contains special characters, they are handled according to the procedures in Section 6.1 of RFC 4975.

4.5. Usage with the Alternative Connection Model

An MSRP endpoint that supports the CEMA extension MUST support the mechanism defined in RFC 6135, as it extends the number of scenarios where one can use the CEMA extension. An example is where an MSRP endpoint is using a relay for MSRP communication, and it needs to be "passive" in order to use the CEMA extension, instead of doing a fallback to RFC 4975 behavior.

5. The SDP 'msrp-cema' Attribute

5.1. General

The SDP 'msrp-cema' attribute is used by MSRP entities to indicate support of the CEMA extension, according to the procedures in Sections 4.2 and 4.3.

5.2. Syntax

This section describes the syntax extensions to the ABNF syntax defined in RFC 4566 required for the SDP 'msrp-cema' attribute. The ABNF defined in this specification is conformant to RFC 5234 [RFC5234].

```
attribute          =/ msrp-cema-attr
;attribute defined in RFC 4566
msrp-cema-attr     = "msrp-cema"
```

6. Middlebox Assumptions

6.1. General

This document does not specify explicit Middlebox behavior, even though Middleboxes enable some of the procedures described here. However, as MSRP endpoints are expected to operate in networks where Middleboxes that want to anchor media are present, this document makes certain assumptions regarding how such Middleboxes behave.

6.2. MSRP Awareness

In order to support interoperability between UAs that support the CEMA extension and UAs that do not support the extension, the Middlebox is MSRP aware. This means that it implements MSRP B2BUA functionality. The Middlebox enables that functionality in cases where the offerer does not support the CEMA extension. In cases where the SDP offer indicates support of the CEMA extension, the Middlebox can simply modify the SDP c/m-line address information for the MSRP connection.

In cases where the Middlebox enables MSRP B2BUA functionality, it acts as an MSRP endpoint. If it does not use the CEMA procedures, it will never forward the SDP 'msrp-cema' attribute in SDP offers and answers.

If the Middlebox does not implement MSRP B2BUA functionality, or does not enable it when the SDP 'msrp-cema' attribute is not present in the SDP offer, CEMA-enabled MSRP endpoints will in some cases be unable to interoperate with non-CEMA-enabled endpoints across the Middlebox.

6.3. TCP Connection Reuse

Middleboxes do not need to parse and modify the MSRP payload when endpoints use the CEMA extension. A Middlebox that does not parse the MSRP payload probably will not be able to reuse TCP connections for multiple MSRP sessions. Instead, in order to associate an MSRP message with a specific session, the Middlebox often assigns a unique local address:port combination for each MSRP session. Due to this, between two Middleboxes there might be a separate connection for each MSRP session.

If the Middlebox does not assign a unique address:port combination for each MSRP session, and does not parse MSRP messages, it might end up forwarding MSRP messages toward the wrong destination.

6.4. SDP Integrity

This document assumes that Middleboxes are able to modify the SDP address information associated with the MSRP media.

NOTE: Even though the CEMA extension as such works with end-to-end SDP protection, the main advantage of the extension is in networks where Middleboxes are deployed.

If the Middlebox is unable to modify SDP payloads due to end-to-end integrity protection, it will be unable to anchor MSRP media, as the SIP signaling would fail due to integrity violations.

6.5. TLS

When UAs use the CEMA extension, this document assumes that Middleboxes relay MSRP media packets at the transport layer. The TLS handshake and resulting security association (SA) can be established peer-to-peer between the MSRP endpoints. The Middlebox will see encrypted MSRP media packets but is unable to inspect the cleartext content.

When UAs fall back to RFC 4975 behavior, Middleboxes act as TLS B2BUAs. The Middlebox decrypts MSRP media packets received from one MSRP endpoint and then re-encrypts them before sending them toward the other MSRP endpoint. Middleboxes can inspect and modify the MSRP message content.

7. Security Considerations

7.1. General

Unless otherwise stated, the security considerations in RFC 4975 and RFC 4976 still apply. This section only describes additions and changes introduced by the CEMA extension.

The purpose of CEMA is to enable MSRP communication over Middleboxes. These Middleboxes are commonly deployed by SIP network operators, who also commonly deploy firewall and routing policies that prevent media sessions from working unless they traverse the Middleboxes.

CEMA makes it possible for Middleboxes to tunnel TLS to allow end-to-end SAs between endpoints. This is an improvement over the status quo, since without CEMA, the Middleboxes would be forced to both read and modify the cleartext MSRP messages, which would make end-to-end confidentiality and integrity protection of the MSRP transport channel impossible.

RFC 4975 suggests two ways for MSRP endpoints to verify that the TLS connection is established end to end. The first option is to use certificates from a well-known certification authority and verify that the SubjectAltName matches the MSRP URI of the other side. The second option is to use self-signed certificates and include a fingerprint of the certificate in the SDP offer/answer. Provided the signaling is integrity protected, both endpoints can verify that the TLS SA is established with the correct host by matching the received certificate against the received fingerprint.

Fingerprint-based authentication is expected to be common for end clients. In order to ensure the integrity of the fingerprint, RFC 4975 recommends using the SIP Identity mechanism [RFC4474]. However, this mechanism may not be compatible with CEMA, which operates under the assumption that Middleboxes will modify the contents of SDP offers and answers. Until a mechanism is available that enables a subset of the SDP to be signed, end clients that support CEMA and use fingerprint-based authentication are forced to trust the entire signaling path. In other words, end clients must accept the fact that every signaling proxy could potentially replace the fingerprints and insert a Middlebox that acts as a TLS B2BUA.

An alternative solution that only requires a limited trust in the signaling plane is to use self-signed certificates together with the SIP Certificate Management Service [RFC6072]. The security provided by this solution is roughly equivalent to SIP Identity and fingerprint-based authentication (in fact, RFC 6072 is based on RFC 4474). Section 7.5 discusses this approach further.

In the remainder of this section, we will assume that fingerprint-based authentication is used without SIP Identity or similar mechanisms that protect the SDP across several hops.

7.2. Man-in-the-Middle (MITM) Attacks

If TLS is not used to protect MSRP, the CEMA extension might make it easier for a MITM to transparently insert itself in the communication between MSRP endpoints in order to monitor or record unprotected MSRP communication. This can be mitigated by the use of TLS. It is therefore RECOMMENDED that TLS [RFC5246] be used. It is also recommended that TLS be used end to end, which CEMA enables even in the case of Middleboxes. According to RFC 4975, MSRP endpoints are required to support TLS. This also applies to CEMA-enabled endpoints.

7.3. TLS Usage without Middleboxes

If TLS is used without Middleboxes, the security considerations in RFC 4975 and RFC 4976 still apply unchanged. Note that this is not the main use case for the CEMA extension.

7.4. TLS Usage with Middleboxes

This is the main use case for the CEMA extension; the endpoints expect one or more Middleboxes.

The CEMA extension supports the usage of both name-based authentication and fingerprint-based authentication for TLS in the presence of Middleboxes. The use of fingerprint-based authentication requires signaling integrity protection. This can, for example, be hop-by-hop cryptographic protection or cryptographic access protection combined with a suitably protected core network. As stated in Section 6.4, this document assumes that Middleboxes are able to modify the SDP address information associated with the MSRP media.

If a Middlebox acts as a TLS B2BUA, the security considerations are the same as those without the CEMA extension. In such a case, the Middlebox acts as a TLS endpoint.

If a Middlebox does not act as a TLS B2BUA, TLS is end to end and the Middlebox just forwards the TLS packets. This requires that both peers support the CEMA extension.

If fingerprint-based authentication is used, the MSRP endpoints might not be able to decide whether or not the Middlebox acts as a TLS B2BUA. But this is not an issue, as the signaling network is considered trusted by the endpoint (a requirement to use fingerprint-based authentication).

7.5. Authentication, Credentials, and Key Management

One issue with the usage of TLS (not specific to CEMA) is the availability of a PKI. Endpoints can always provide self-signed certificates and include fingerprints in the SDP offer and answer. However, this relies on SDP signaling being integrity protected, which may not always be the case.

Therefore, in addition to the authentication mechanisms defined in RFC 4975, it is RECOMMENDED that a CEMA-enabled MSRP endpoint also support self-signed certificates together with the Certificate Management Service [RFC6072], to which it publishes its self-signed certificate and from which it fetches on demand the self-signed certificates of other endpoints.

Alternate key distribution mechanisms, such as DNS-Based Authentication of Named Entities (DANE) [DANE], Pretty Good Privacy (PGP) [RFC6091], Ticket-Based Modes of Key Distribution in Multimedia Internet KEYing (MIKEY-TICKET) [RFC6043], or some other technology, might become ubiquitous enough to solve the key distribution problem in the future.

One of the target deployments for CEMA is the 3GPP IMS SIP network. In this environment, authentication and credential management are less of a problem, as the SDP signaling is mostly considered trusted, service providers provision signed certificates or manage signed certificates on behalf of their subscribers, and MIKEY-TICKET is available. Some of these options require trusting the service provider, but those issues are beyond the scope of this document.

7.6. Endpoint Procedures for TLS Negotiation

The CEMA extension does not change the endpoint procedures for TLS negotiation. As in RFC 4975, the MSRP endpoint uses the negotiation mechanisms in SDP and then the TLS handshake to agree on mechanisms and algorithms that both support. The mechanisms can be divided into three different security levels:

1. MSRPS: Security mechanisms that do not rely on trusted signaling, such as name-based authentication

2. MSRPS: Mechanisms that do rely on trusted signaling, such as fingerprint-based authentication

3. MSRP: Unprotected

If the endpoint uses security mechanisms that do not rely on trusted signaling, the endpoint can detect if a Middlebox that acts as a B2BUA is inserted. It is therefore RECOMMENDED that such a mechanism be used.

If the endpoint uses security mechanisms that rely on trusted signaling, the endpoint may not be able to detect if a Middlebox that acts as a B2BUA is inserted (by the trusted network operator). To be able to eavesdrop, a Middlebox must do an active "attack" on the setup signaling. A Middlebox cannot insert itself at a later point.

If unprotected MSRP is used, the endpoint cannot detect if a Middlebox that acts as a B2BUA is inserted and Middleboxes may be inserted at any time during the session.

The mechanism in RFC 6072 [RFC6072] provides end-to-end security without relying on trust in the signaling plane and eases the use and deployment of name-based authentication.

The procedures for choosing and offering name-based authentication, fingerprint-based authentication, and unprotected MSRP as described in RFC 4975 still apply.

7.7. Fingerprint-Based Authentication

If the endpoint cannot use a key management protocol that does not rely on trust in the signaling plane, such as name-based authentication, the only alternative is fingerprint-based authentication.

The use of fingerprint-based authentication requires integrity protection of the signaling plane. This can, for example, be hop-by-hop cryptographic protection or cryptographic access protection combined with a suitably protected core network. Unless cryptographic end-to-end SDP integrity protection or encryption is used, this may be hard for the endpoint to decide. In the end, it is up to the endpoint to decide whether the signaling path is trusted or not.

How this decision is done is implementation specific, but normally, signaling over the Internet SHOULD NOT be trusted. Signaling over a local or closed network might be trusted. Such networks can, for example, be a closed enterprise network or a network operated by an operator that the end user trusts. In IMS, for example, the signaling traffic in the access network is integrity protected and the traffic is routed over a closed network separated from the Internet. If the network is not trusted, the endpoints SHOULD NOT use fingerprint-based authentication.

When an endpoint receives a fingerprint, that fingerprint represents a binding between the identity as established by TLS and that established via SDP. As previously noted, the fingerprint is vulnerable to an active MITM attack from any on-path proxy. Endpoints SHOULD therefore locally store fingerprints associated with the relevant identities when first seen and SHOULD provide a warning when a new fingerprint is seen for what otherwise appears to be the same peer identity. While there are valid reasons for keys to change from time to time, that ought to be the exception -- hence the suggested warning.

It should, however, be noted that using fingerprint-based authentication over an insecure network increases the security compared to unencrypted MSRP. In order to intercept the plaintext media when fingerprint-based authentication is used, the attacker is required to be present on both the signaling and media paths and actively modify the traffic. It is very hard for the endpoints to detect when such an attack is taking place, though. A client using DTLS-SRTP (a Secure Real-time Transport Protocol (SRTP) extension for Datagram Transport Layer Security (DTLS)) [RFC5764] for Voice over IP (VoIP) media security might wish to use fingerprint-based authentication also for MSRP media security.

MSRPS with fingerprint-based authentication is vulnerable to attacks due to vulnerabilities in the SIP signaling. If there are weaknesses in the integrity protections on the SIP signaling, an attacker may insert malicious Middleboxes to alter, record, or otherwise harm the media. With insecure signaling, it can be difficult for an endpoint to even be aware that the remote endpoint has any relationship to the expected endpoint. Securing the SIP signaling does not solve all problems. For example, in a SIP Secure (SIPS) environment, the endpoints have no cryptographic way of validating that one or more SIP proxies in the proxy chain are not, in fact, malicious.

8. IANA Considerations

8.1. IANA Registration of the SDP 'msrp-cema' Attribute

IANA has added an attribute to the 'att-field (media level only)' registry of the Session Description Protocol (SDP) Parameters registry, according to the information provided in this section.

This section registers a new SDP attribute, 'msrp-cema'. The required information for this registration, as specified in RFC 4566, is as follows:

Contact name: Christer Holmberg

Contact email: christer.holmberg@ericsson.com

Attribute name: msrp-cema

Type of attribute: media level

Purpose: This attribute is used to indicate support of the MSRP Connection Establishment for Media Anchoring (CEMA) extension defined in RFC 6714. When present in an MSRP media description of an SDP body, it indicates that the creator of the SDP supports the CEMA mechanism.

Values: The attribute does not carry a value.

Charset dependency: none

9. Acknowledgements

Thanks to Ben Campbell, Remi Denis-Courmont, Nancy Greene, Hadriel Kaplan, Adam Roach, Robert Sparks, Salvatore Loreto, Shida Schubert, Ted Hardie, Richard L. Barnes, Inaki Baz Castillo, Saul Ibarra Corretge, Cullen Jennings, Adrian Georgescu, Miguel Garcia, and Paul Kyzivat for their guidance and input in order to produce this document.

Thanks to John Mattsson, Oscar Ohlsson, Ben Campbell, and Stephen Farrell for their help in restructuring the Security Considerations section, based on feedback from the IESG.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[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.

[RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.

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

[RFC4975] Campbell, B., Ed., Mahy, R., Ed., and C. Jennings, Ed., "The Message Session Relay Protocol (MSRP)", RFC 4975, September 2007.

[RFC4976] Jennings, C., Mahy, R., and A. Roach, "Relay Extensions for the Message Sessions Relay Protocol (MSRP)", RFC 4976, September 2007.

[RFC5234] Crocker, D., Ed., and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[RFC6072] Jennings, C. and J. Fischl, Ed., "Certificate Management Service for the Session Initiation Protocol (SIP)", RFC 6072, February 2011.

[RFC6135] Holmberg, C. and S. Blau, "An Alternative Connection Model for the Message Session Relay Protocol (MSRP)", RFC 6135, February 2011.

10.2. Informative References

[RFC3724] Kempf, J., Ed., Austein, R., Ed., and IAB, "The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture", RFC 3724, March 2004.

[RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.

[RFC5764] McGrew, D. and E. Rescorla, "Datagram Transport Layer Security (DTLS) Extension to Establish Keys for the Secure Real-time Transport Protocol (SRTP)", RFC 5764, May 2010.

[RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010.

[RFC6043] Mattsson, J. and T. Tian, "MIKEY-TICKET: Ticket-Based Modes of Key Distribution in Multimedia Internet KEYing (MIKEY)", RFC 6043, March 2011.

[RFC6091] Mavrogiannopoulos, N. and D. Gillmor, "Using OpenPGP Keys for Transport Layer Security (TLS) Authentication", RFC 6091, February 2011.

[GPP23228] 3GPP, "IP Multimedia Subsystem (IMS); Stage 2", 3GPP TS 23.228 11.5.0, June 2012, <http://www.3gpp.org/ftp/Specs/html-info/23228.htm>.

[DANE] "DNS-Based Authentication of Named Entities (DANE) Working Group", <https://datatracker.ietf.org/wg/dane/charter/>.

Authors' Addresses

Christer Holmberg
Ericsson
Hirsalantie 11
Jorvas 02420
Finland

EMail: christer.holmberg@ericsson.com

Staffan Blau
Ericsson
Stockholm 12637
Sweden

EMail: staffan.blau@ericsson.com

Eric Burger
Georgetown University
Department of Computer Science
37th and O Streets, NW
Washington, DC 20057-1232
United States of America

Fax: +1 530 267 7447
EMail: eburger@standardstrack.com
URI: http://www.standardstrack.com