Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

SIP-PBX / Service Provider Interoperability

"SIPconnect 2.0 Technical Recommendation"

SIP Forum Document Number: TWG-11

Abstract

The SIPconnect 2.0 Technical Recommendation is a profile of the Session Initiation Protocol (SIP) and related media aspects that enables direct connectivity between a SIP-enabled Service Provider Network and a SIP-enabled Enterprise Network. It specifies the minimal set of IETF and ITU-T standards that must be supported, provides precise guidance in the areas where the standards leave multiple implementation options, and specifies a minimal set of capabilities that should be supported by the Service Provider and Enterprise Networks.

SIPconnect 2.0 effectively extends SIPconnect 1.1. Where SIPconnect 1.0, and 1.1, focused primarily on basic network registration, identity/privacy management, call originations, call terminations, and advanced services, this version adds additional guidance on Security, Emergency Calling, and IPv6.

Where appropriate, recommendations from SIPconnect 1.1 have been left unchanged, although some modifications to prior recommendations have been made based on experience and feedback gathered through adoption of SIPconnect 1.1 in the industry.

Status of this Memo

SIPconnect 2.0 FINAL (v.18).

Disclaimer

The SIP Forum takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the SIP Forum's procedures with respect to rights in SIP Forum Technical Recommendations, both drafts and final versions, or other similar documentation can be found in the SIP Forum's current adopted intellectual property right Recommendation. Copies of claims of rights made available for publication 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 Technical Recommendation can be obtained from the SIP Forum.

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

SIPconnect Certified and SIPconnect Compliant

SIPconnect, SIPconnect Compliant, and SIPconnect Certified are certification marks of the SIP Forum. Implementers who wish to certify their products and services as SIPconnect Compliant and SIPconnect Certified may do so under the SIPconnect Certification Testing program of the SIP Forum. To learn more about this opportunity and obtain other useful information about SIPconnect Certification, please visit http://www.sipforum.org/content/view/289/307/.

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

Table of Contents

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

List of Figures

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

1. Introduction

The Session Initiation Protocol (SIP) is the dominant industry standard for signaling in support of VoIP and other services. The deployment of Session Initiation Protocol (SIP)-enabled PBXs (SIP-PBXs) among Enterprises of all sizes is increasing rapidly. Deployment of SIP infrastructure by Service Providers is also increasing, driven by the demand for commercial VoIP offerings. Many new SIP-PBXs support SIP phones and SIP-based communication with other SIP-PBXs. The result of these parallel deployments is a present need for direct IP peering between SIP-enabled SIP-PBXs and Service Providers.

Currently published ITU-T Recommendations and IETF RFCs offer a comprehensive set of building blocks that can be used to achieve direct IP peering between SIP-enabled SIP-PBX systems and a Service Provider's SIP-enabled network. However, due to the sheer number of these standards documents, Service Providers and equipment manufacturers have no clear "master reference" that outlines which standards they must specifically support in order to ensure success. This has led to a number of interoperability problems and has unnecessarily slowed the migration to SIP as replacement for traditional TDM (Time Division Multiplexed) connections.

This SIP Forum document aims to address this issue. In short, this document defines the protocol support, implementation rules, and features required for predictable interoperability between SIP-enabled Enterprise Networks and SIP-enabled Service Providers. Note that this document does not preclude or discourage the negotiation of additional functionality.

SIPconnect 2.0 restates, updates, and extends the areas of implementation guidance found in SIPconnect 1.1, including:

- Specification of a reference architecture that describes the common network elements necessary for Service Provider-to-SIP-PBX peering for the primary purpose of call origination and termination.
- Specification of the basic protocols (and protocol extensions) that must be supported by each element of the reference architecture.
- Specification of the exact standards associated with these protocols that must or should be supported by each element of the reference architecture.
- Specification of two modes of operation – Registration mode and Static mode - whereby a Service Provider can locate a SIP-PBX.
- Specification of standard forms of Enterprise Public Identities.
- Specification of signaling messages for Basic 2-Way Calls, Call Forwarding, and Call Transfer.
- Specification of minimum requirements for codec support, packetization intervals, and capability negotiation.
- Specification of minimum requirements for handling fax and modem transmissions.
- Specification of minimum requirements for handling echo cancellation.
- Specification of minimum requirements for transporting DTMF tones.
- Specification of security mechanisms for both signaling and media security.
- Specification of minimum requirements for supporting IPv6.
- Specification of minimum requirements for emergency calling.

2. Conventions and Terminology

The keywords "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in [RFC 2119]

3. Reference Architecture

The reference architecture diagram in Figure 1 shows the functional elements that may be deployed to support the interface described in this Technical Recommendation. The diagram shows two reference points between the Enterprise Network and the Service Provider Network; a signaling reference point (1) and a media reference point (2).

The signaling reference point carries SIP signaling messages to support voice services between the Enterprise Network and the Service Provider network.
The media reference point carries the RTP and RTCP packets between the Service Provider and Enterprise Media Endpoints. An Enterprise Media Endpoint could be contained within an SBC, SIP-PBX, an IP-based user device (e.g., SIP phone) in the Enterprise, or a media-relay device in the Enterprise Network. The Service Provider Media Endpoint could be a SBC, PSTN Gateway, an IP-based user endpoint device, a media server, or any other IP-based media-capable entity.

The signaling reference point and the media reference point together comprise the SIPconnect 2.0 interface.

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

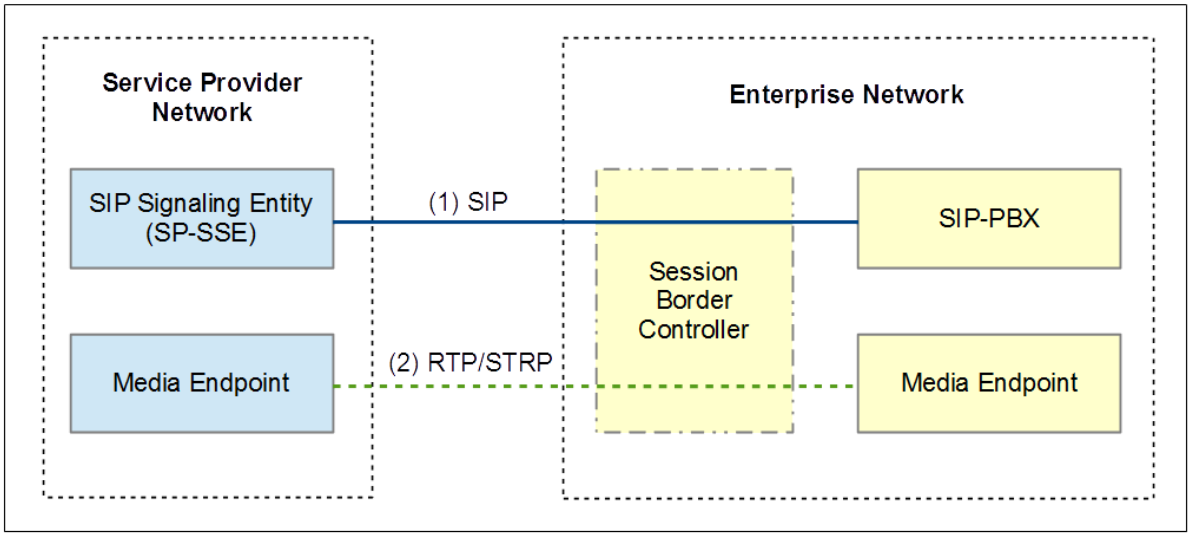

Figure 1: Reference Architecture

It is important to note that this Technical Recommendation presents these functional elements as separate physical components for the purposes of illustration only. It is perfectly acceptable for an equipment manufacturer to combine these entities. For example, a manufacturer may choose to integrate the SIP-PBX and Media Endpoint functions. Both integrated and non-integrated implementations are equally conformant as long as they fully adhere to the individual rules governing each of the defined functions.

Additionally, just as multiple logical functions can be collapsed into one physical entity, a single logical function in this Technical Recommendation can be decomposed into multiple physical entities. For example, the SP-SSE can be decomposed into the functional nodes of an IMS core network. The internal interfaces of the SP-SSE is however not covered by this Technical Recommendation.

[RFC 7092] provides a Taxonomy of SIP Back-to-Back User Agents (B2BUA's) including examples of how the entities described in this recommendation may be combined in different ways.

Note that many deployments will include a Network Address Translator (NAT) between the Service Provider Network and the Enterprise Network. This document does not describe NATs as part of the SIPconnect 2.0 interface. This document describes functionality as observed at the reference points. The requirements at the reference points are unaffected by the presence of a NAT.

Note that a single SIP-PBX may serve Media Endpoints in a number of geographically-distributed locations.

4. Definitions

Service Provider SIP-Signaling Entity (SP-SSE) – the Service Provider's point of SIP signaling interconnection with the Enterprise.

SIP-PBX – The Enterprise's point of SIP signaling interconnection with the Service Provider.

SIP Endpoint – The term used in this specification to refer to both SP-SSEs and SIP-PBXes.

Enterprise Public Identity - An Address of Record (AOR) represented as a SIP URI, used to identify a user or group of users served by the SIP-PBX. Enterprise Public Identities are used in conjunction with delivering incoming and outgoing calls.

Registration AOR – An AOR represented as a SIP URI, used solely to identify the SIP-PBX during registration.

Media Endpoint – Any entity that terminates an SRTP/RTP/RTCP stream.

Back-to-Back User Agent (B2BUA) –A logical entity that receives a request and processes it as a user agent server (UAS). In order to determine how the request should be answered, it acts as a user agent client (UAC) and generates a request to another SIP user agent server (UAS).

5. Key Assumptions and Limitations of Scope

This Technical Recommendation lists a number of IETF and ITU-T specifications needed to meet the requirements for interconnection between a Service Provider and an Enterprise Network.

The following key assumptions have been made:

- The primary service to be delivered over this interface is audio-based call origination and/or termination between the Enterprise and Service Provider Networks, including emergency services. The delivery of any other service (e.g. instant messaging, etc.) is out of scope.
- All reference architecture elements specified for the Service Provider and Enterprise Networks are in place and operational.
- Signaling considerations between the SP-SSE and other Service Provider devices (e.g. Trunking Gateway) are outside the scope of this document.
- Signaling considerations between the SIP-PBX and other Enterprise devices (e.g. IP phones) are outside the scope of this document.
- Layer 3 network design and QoS considerations are outside of the scope of this document

- Element management, network management, network security, and other operational considerations are outside the scope of this document.

SIPconnect assumes a peering model in which the both the Service Provider and the Enterprise deploy advanced call processing platforms, which communicate via a SIP trunk. The Service Provider enables communication between users in the Enterprise network, who are served by the SIP-PBX, and users outside the Enterprise network. The SIP-PBX in the Enterprise typically provides a variety of call services (Voice Mail, Call Forwarding, Hunt Group, and so on); the SP-SSE in the Service Provider's network may provide additional call services. The Service Provider is assumed to have knowledge of the E.164 numbers associated with the SIP-PBX, which, together with a domain name, form the Enterprise Public Identities of the Enterprise users.

6. Basic SIP Support

SIP-PBXs and SP-SSEs **MUST** support SIP in accordance with [RFC 3261] and offer-answer in accordance with [RFC 3264], as qualified by statements in later sections of this document. Requirements for support of other IETF RFCs and other standards are as stated in later sections of this document.

This document specifies a profile of SIP, as well as specifying some media aspects. Implementations of this Technical Recommendation **MUST NOT** simply assume that a particular feature or option listed as mandatory in this document is supported by a peer SIP-PBX or SP-SSE. Instead, a SIP-PBX or SP-SSE **MUST** use mechanisms specified for SIP (e.g., Supported, Require and Allow header fields) and SDP (e.g., attributes, payload formats) for ascertaining support of a given SIP or SDP extension at a peer SP-SSE or SIP-PBX. Failure to do this can lead to interoperability problems.

7. Modes of Operation

This document describes two modes of operation for SIPconnect 2.0; the Registration mode (specified in Annex A) and the Static mode (specified in Annex B). These modes differ primarily in the way the Service Provider Network discovers the SIP signaling address of the SIP-PBX.

In the Registration mode, the SIP-PBX conveys its SIP signaling address to the Service Provider Network using the SIP registration procedure defined in [RFC 6140] In effect, the SIP-PBX registers with the Service Provider Network, using a REGISTER request with a specially-formatted Contact URI. After the SIP-PBX is authenticated, the registrar updates its location service with a unique AOR-to-Contact mapping for each of the AORs associated with the SIP-PBX. The primary advantage of the Registration mode is that it enables the SIP-PBX to be easily deployed in a "plug-and-play" fashion; i.e., with only a minimum of configuration data the SIP-PBX can initiate the registration procedure to automatically establish connectivity with the Service Provider Network.

In Registration mode:

- The SIP-PBX uses SIP registration procedures to advertise the SIP-PBX's SIP signaling address to the SP-SSE, and

- The SP-SSE authenticates the SIP-PBX using SIP Digest.

In the Static mode, the Service Provider Network views the SIP-PBX as a peer SIP-based network that is responsible for the Enterprise Public Identities that it serves. In this mode the Service Provider Network is either configured with the SIP-PBX signaling address, or it discovers the address using the Domain Name Service (DNS). The Service Provider Network procedures for routing out-of-dialog requests to the SIP-PBX align closely with the SIP routing procedures defined in [RFC 3261] (and [RFC 3263] if DNS is used).

In Static mode:

- The Enterprise Network can use DNS to advertize its publicly-reachable SIP-PBX SIP signaling address to the SP-SSE.

Advantages of Registration mode over Static mode include:

- It enables the Service Provider Network to discover the signaling address of the SIP-PBX that is assigned a dynamic IP address (so that the SIP-PBX is not required to have a static signaling address publicly viewable in DNS),
- It provides a mechanism for a SIP-PBX located behind a NAT to automatically establish connectivity with the Service Provider Network,
- It provides a mechanism for a failed SIP-PBX to automatically inform the network when it is back online, and
- It enables the Service Provider to tap into streamlined and scalable subscriber provisioning and management processes (e.g., a Service Provider Network that is designed to support the heavy registration traffic generated by millions of users is well suited to support registration traffic generated by large numbers of SIP-PBXs operating in the Registration mode).

Advantages of Static mode over Registration mode include:

- Since Static-mode SIP-PBXes do not send REGISTER requests when they initialize, Static mode operation is less susceptible to "avalanche restart" issues, when a large geographic area restores power, and
- The SP-SSE is not dependent on the SIP-PBX to re-establish any broken registration before the SP-SSE can deliver inbound requests to the SIP-PBX.

The Static mode is often used for larger Enterprises, where the size of the Enterprise warrants more explicit provisioning of connection and service information by the Service Provider. For example, large Enterprise trunks often have unique requirements for SLAs (Service Level Agreements), call routing, load balancing, codec support, etc., which make explicit provisioning necessary.

SIP-PBXs **MUST** support either Registration mode, as specified in Annex A, or Static mode, as described in Annex B. SIP-PBXs **MAY** support both modes.

SP-SSEs **MUST** support either Registration mode, as specified in Annex A, or Static mode, as described in Annex B. SP-SSEs **MAY** support both modes.

Note that an SP-SSE supporting only Annex A and a SIP-PBX supporting only Annex B, or vice versa, will not interoperate. Both sides must support the same Annex in order to communicate.

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

8. Supported Signaling Transport Protocols

SIP-PBXs and SP-SSEs **MUST** implement TCP. TCP does not have to be **used** for a SIPconnect 2.0 signaling connection, if both sides agree not to, but it **MUST** be available in order to comply with this Technical Recommendation.

UDP support is allowed in order to accommodate legacy devices. TCP support is mandated in order to accommodate large and growing SIP requests and responses (see Section 18.1.1 for more background), and for use with TLS.

8.1 TLS

While SIPconnect 2.0 continues to require TLS support at **MUST** strength, we should note that using TLS for signaling as described in Sections 16.2 and 17.2 does not require the use of the SIPS URI scheme.

[RFC 3261] Section 26.2.2 deprecates the "transport=TLS" URI parameter. SIP-PBXes and SP-SSEs **MUST** ignore this parameter.

When presenting a certificate, a SIP-PBX or SP-SSE **SHOULD** identify itself by means of a SIP URI using type uniformResourceIdentifier in the subjectAltName field, in accordance with [RFC 5280].

[RFC 3261] Section 26.3.1 states:

> Proxy servers, redirect servers, and registrars **SHOULD** possess a site certificate whose subject corresponds to their canonical hostname.

When receiving a certificate, SIP-PBX and SP-SSE implementations **MUST** support extraction of the canonical hostname from the subjectCommonName (CN) if (and only if) there are no subjectAltName extension fields, following the rules documented in Section 7.1 of [RFC 5922]. SIP-PBX and SP-SSE implementations **MUST** comply with guidelines relating to usage of the Subject field, specified in [RFC 5280] Section 4.1.2.6, and the SubjectAltName field as specified in [RFC 5280] Section 4.2.1.6. Compliance with [RFC 5280] Section 4.1.2.6 is necessary to support existing certificate signer implementations that use the CN field instead of the subjectAltName field.

Furthermore, SIP-PBX and SP-SSE implementations **MUST** be able to accept a DNS name as an identity (e.g. proxy1.example.com), instead of a SIP URI as defined in [RFC 3261] (e.g., sip:proxy.example.com). This is to allow for supporting SP-SSE or SIP-PBX implementations that commonly use certificates that were created for HTTP instead of for SIP. It is also **RECOMMENDED** that SIP-PBX and SP-SSE implementations be able to provide a certificate with either a URI or DNS name for backward compatibility.

8.1.1 SP-SSE TLS Requirements

The SP-SSE **MUST** support TLS version 1.2, higher versions **MAY** be supported when available. The SP-SSE **MAY** be configured to support TLS version 1.0 in order to enable interworking with SIP-PBX which does not support higher versions. The SP-SSE **MUST** avoid TLS protocol version intolerance. I.e., even if only TLS 1.2 is supported, TLS handshakes with peers that try to negotiate higher - yet unknown - versions (e.g. TLS 1.3 or TLS 2.98) **MUST** succeed (ending up in TLS 1.2 negotiation).

An SP-SSE **MUST** support the following cipher suite:

- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256.

The SP-SSE **MAY** support the following cipher suites for backwards compatibility:

- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_128_CBC_SHA

The SP-SSE when acting as the TLS server **MUST** determine the cipher to be used based on its own preference order (i.e. TLS_ECDH**E**_RSA_WITH_AES_128_GCM_SHA256, and then optional ciphers, then the TLS 1.0 fallback cipher(s)) and use the first in its own list, that is also available in the list sent by the TLS client.

The SIP-SSE acting as the TLS client (Annex B: Static Mode) **MUST** send the list of supported ciphers in the order of preference as above (i.e. TLS_ECDH**E**_RSA_WITH_AES_128_GCM_SHA256, and then optional ciphers, then the TLS 1.0 fallback cipher(s)).

8.1.2 SIP-PBX TLS Requirements

The SIP-PBX **MUST** support TLS version 1.2, higher versions **MAY** be supported when available. The SIP-PBX **MAY** be configured to support TLS version 1.0 in order to enable interworking with SIP-SSE which does not support higher versions. The SIP-PBX **MUST** avoid TLS protocol version intolerance. I.e., even if only TLS 1.2 is supported, TLS handshakes with peers that try to negotiate higher - yet unknown - versions (e.g. TLS 1.3 or TLS 2.98) **MUST** succeed (ending up in TLS 1.2 negotiation).

An SIP-PBX **MUST** support the following cipher suite:

- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256.

The SP-SSE **MAY** support the following cipher suites for backwards compatibility:

- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_128_CBC_SHA

The SIP-PBX when acting as the TLS server (Annex B: Static Mode) **MUST** determine the cipher to be used based on its own preference order (i.e. TLS_ECDH**E**_RSA_WITH_AES_128_GCM_SHA256, and then optional ciphers, then the TLS 1.0 fallback cipher(s)) and use the first in its own list, that is also available in the list sent by the TLS client.

The SIP-PBX acting as the TLS client **MUST** send the list of supported ciphers in the order of preference as above (i.e. TLS_ECDH**E**_RSA_WITH_AES_128_GCM_SHA256, and then optional ciphers, then the TLS 1.0 fallback cipher(s)).

9. Enterprise Public Identities

SIP-PBXs and SP-SSEs **MUST** be able to support Enterprise Public Identities in the form of a SIP URI containing a global E.164 [ITU-T E.164] number and the "user=phone" parameter.

For example:

sip:+16132581234@example.com;user=phone

The global E.164 number **MUST** begin with a leading "+", **MUST NOT** contain a phone-context parameter and **MUST NOT** include visual separators.

For a given SIPconnect 2.0 interface, the choice of value for the host part of Enterprise Public Identities is a contractual matter between the enterprise and the Service Provider. For Registration mode, the value of the host part of Enterprise Public Identities will be the domain name or sub-domain name of the Service Provider. For Static mode, the value of the host part of Enterprise Public Identities can be in the form of a

sub-domain of the Service Provider domain assigned to the SIP-PBX (e.g. "pbx1.operator.net"), or the SIP-PBX IP address, or the domain of the Enterprise (e.g. "enterprise.com").

Support for other forms of Enterprise Public Identity (including identities based on telephone numbers that are not global E.164 numbers (e.g., sip:7042;phone-context=enterprise.com@example.com;user=phone) and identities not based on telephone numbers (e.g., sip:alice@example.com) is out of scope of this Technical Recommendation.

9.1 Routing SIP Requests to Enterprise Public Identities

The SP-SSE is responsible for routing SIP requests to the appropriate SIP-PBX; i.e. on receiving a SIP request and translating the destination address to an Enterprise Public Identity, the SP-SSE **MUST** use that Enterprise Public Identity to discover the SIP signaling address of the SIP-PBX. The mechanism to perform this discovery depends on whether the SIP-PBX is deployed using Registration or Static mode:

- In Registration mode, the SP-SSE determines the SIP-PBX signaling address using the address binding that was established when the SIP-PBX registered, as described in Section 16.
- In Static mode the SP-SSE determines the SIP-PBX signaling address using either statically configured data or DNS, as described in Section 17.

10. Establishing Basic 2-Way Calls

This section describes the procedures for establishing basic 2-way calls between the Enterprise and the Service Provider Network.

10.1 Incoming Calls from the Service Provider to the Enterprise

Calls to Enterprise Public Identities are routed by the SP-SSE to the SIP-PBX and are usually routed by the SIP-PBX directly to a specific user station – bypassing the attendant or operator. This is commonly referred to as "Directed Inward Dial" (DID) service.

This section describes guidelines for populating the Request-URI, and the "P-Asserted-Identity" [RFC 3325] and [RFC 5876], "To" and "From" header fields for new-dialog INVITE requests sent from the SP-SSE to the SIP-PBX. The SP-SSE **MUST** ensure that all other header fields in the INVITE request comply with [RFC 3261].

10.1.1 Request-URI

The SP-SSE **MUST** populate the Request-URI of the INVITE request in accordance with Section 16.7 for Registration mode and in accordance with Section 17.6 for Static mode.

On receiving an INVITE request from the SP-SSE, the SIP-PBX **MUST** identify the called user based on the contents of the Request-URI.

10.1.2 "To" Header Field

The "To" header field URI of a SIP request generated by the SP-SSE is frequently populated with the Enterprise Public Identity to which the Request-URI relates. However, there may be cases, such as a prior redirection, where the "To" header field URI does not contain the desired destination. As such, the SIP-PBX **MUST NOT** rely on the contents of "To" header field for routing decisions, but **MUST** use the Request-URI instead.

10.1.3 "From" Header Field

For IP-based originations, there are no special restrictions on the contents of the "From" header field URI, beyond the requirements specified in [RFC 3261]. For example, the "From" header field URI could contain either a SIP or Tel URI. Typically the "From" header field URI is set by the originating UAC, and either carried transparently through to the terminating UAS, or modified en-route. For example, a network-based "anonymizing" service could update the "From" header field URI to obscure the identity of the caller and originating Service Provider. In cases where the SP-SSE needs to generate an anonymous URI (e.g., for a call incoming to the Service Provider Network from the PSTN for which calling number privacy is requested), the SP-SSE **MUST** send a URI as shown here.

sip:anonymous@anonymous.invalid

Note: Where a display-name is included, no semantic meaning should be attributed to the display name. This has resulted in reported interoperability problems, because the display name could be in any language.

If the originating SIP entity supplied an E.164 calling number, and the caller did not request calling number privacy, then the SP-SSE **MUST** populate the "From" header field with a SIP URI containing the E.164 calling number, the Service Provider domain name, and the "user=phone" parameter as shown below. If any display name information is available and has not been restricted for delivery, it **SHOULD** also be provided.

sip:+15616261234@example.com;user=phone

where "example.com" is the domain name of the Service Provider Network.

If no caller identity is available and privacy has not been requested, the SP-SSE **SHOULD** send a URI containing a host portion with a top level domain of ".invalid", as shown below.

sip:unavailable@unknown.invalid

There are no special requirements placed on the SIP-PBX in processing the "From" header field, beyond the requirements specified in [RFC 3261].

10.1.4 "P-Asserted-Identity" and "Privacy" Header Fields

If the caller requested privacy, and the Service Provider Network does not trust the Enterprise Network, then the SP-SSE **MUST** remove all "P-Asserted-Identity" header fields in the INVITE request before sending the request to the SIP-PBX.

If the caller requested privacy, and the SP-SSE is able to assert an identity, and the Service Provider Network trusts the Enterprise Network, then the SP-SSE **MUST** include a "P-Asserted-Identity" header field and a "Privacy" header field with value 'id' in the INVITE request, in addition to providing an anonymous "From" header field URI as specified in Section 10.1.3, before sending the request to the SIP-PBX. When privacy was not requested by the remote UE, the SIP-PBX **MUST** anyway support to both receive and to not receive the PAI header (the SP-SSE may have policies to for not sending the PAI header).

If the caller did not request privacy, and the SP-SSE is able to assert an identity, then the SP-SSE **SHOULD** include a "P-Asserted-Identity" header field containing a URI identifying the calling user in the INVITE request before sending the request to the SIP-PBX.

In general, there are no restrictions on the contents of the "P-Asserted-Identity" header field, beyond the requirements specified in [RFC 3325] and [RFC 5876]. This is due to the fact that when the SP-SSE receives a "P-Asserted-Identity" header field, in a SIP request or response, from a trusted entity that conforms to [RFC 3325] and [RFC 5876], it transparently passes the header field to the SIP-PBX without modification. This means that the SIP-PBX **MUST** support receiving a "P-Asserted-Identity" header field containing any form of URI permissible according to [RFC 3325] and [RFC 5876].

The "domain-name" identifies the domain of the originating network; e.g. "domain-name" could be domain of the Service Provider Network, domain of a peer to the Service Provider Network, or domain of another Enterprise Network. As described in [RFC 3325], the SIP-PBX **MUST** accept up to two "P-Asserted-Identity" header field values, one in the form of a Tel URI, and one in the form of a SIP URI. The SIP-PBX **MUST** prefer the SIP URI when two are present.

If the "P-Asserted-Identity" header field is to be included, then the SP-SSE **SHOULD** also include display name information along with the SIP or Tel URI in the "P-Asserted-Identity" header field, if the display name is available and has not been restricted for delivery.

For example:

P-Asserted-Identity: "John Smith" <sip:+15616261234@example.com;user=phone>

The SIP-PBX **MUST** support receiving a "Privacy" header field from the SP-SSE that contains a priv-value of either ‘id’ or ‘none’, as per [RFC 3325], [RFC 5876] and [RFC 3323].

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

10.2 Outgoing Calls from the Enterprise to the Service Provider

This section describes SIP-PBX and SP-SSE requirements for populating and receiving the Request-URI and "To" and "From" header fields for new dialog INVITE requests sent from the SIP-PBX to the SP-SSE. It also specifies how the "P-Asserted-Identity" header field can be used by the Enterprise Network to assert the identity of the caller, and usage of the "Privacy" header field to suppress the delivery of caller identity, as described in [RFC 3325] and [RFC 5876]. The SIP-PBX **MUST** ensure that all other header fields in the INVITE request comply with [RFC 3261].

This section covers the case where the call is initiated by an Enterprise user served by the SIP-PBX. The case where the SIP-PBX sends an INVITE request to the SP-SSE to establish the forward-to leg of a call forwarded by an Enterprise user is covered in Section 11.

10.2.1 Request-URI

If the SIP-PBX has an E.164 number identifying the called user (e.g., derived from a Tel URI or a dial string), the SIP-PBX **MUST** populate the Request-URI of the INVITE request with a SIP URI of the following form, using the domain name of the Service Provider in the host part:

sip:+12128901234@sp.example.com;user=phone

If the SIP-PBX has a dial string identifying the called user and is unable to convert it to a SIP URI of the form "user=phone", the SIP-PBX **MUST** populate the Request-URI of the INVITE request with a SIP URI in the following form:

sip:92125551212@sp.example.com

10.2.2 "To" Header Field

The "To" header field URI in a SIP request generated by the SIP-PBX is normally populated with the same URI as the Request-URI. However, there may be cases, such as a prior redirection, where the "To" header field URI does not contain the desired destination. As such, the SP-SSE **MUST NOT** rely on the "To" header field URI for routing decisions, but use the Request-URI instead.

10.2.3 "P-Asserted-Identity" Header Field

The SIP-PBX **MUST** include a "P-Asserted-Identity" header field in the INVITE request in accordance with the rules of [RFC 3325] and [RFC 5876]. If the URI is an Enterprise Public Identity, then it must be formed in accordance with section 9 of this specification.

10.2.4 "From" Header Field

The SIP-PBX **MUST** populate the "From" header field URI with a URI that the SIP PBX wishes to be used for caller identification. This may be an Enterprise Public Identity, an anonymous URI, or a SIP or Tel URI that the SIP-PBX has received from an entity behind the SIP-PBX.

If the "From" URI is not an Enterprise Public Identity, the Service Provider's ability to deliver this information as caller identification will depend on policy.

In cases where the Enterprise Network needs to generate an anonymous URI on behalf of a caller (as opposed to passing on a received anonymous URI), the SIP-PBX **MUST** send a URI of the form

sip:anonymous@anonymous.invalid

10.2.5 "Privacy" Header Field

If the SIP-PBX requires privacy for a call by suppressing delivery of caller identity to downstream entities, it **MUST** include a "Privacy" header field with value 'id' in the INVITE request, in addition to providing an anonymous "From" header field URI as specified in Section 10.2.4. If the SP-SSE provides privacy by default and the SIP-PBX requires privacy to be overridden for a call, the SIP-PBX **MUST** include a "Privacy" header field with value 'none' in the INVITE request.

The SP-SSE **MUST** support receiving a "Privacy" header, from the SIP-PBX that contains a priv-value of either ‘id’ or ‘none’, as per [RFC 3325], [RFC 5876] and [RFC 3323].

10.2.6 “P-Preferred-Identity” Header Field

The SIP-PBX **MAY** include a "P-Preferred-Identity" header field in the INVITE request in accordance with the rules of [RFC 3325]. The handling of the "P-Preferred-Identity" header field by the SP-SSE is out of scope here.

11. Call Forwarding

The ability for the Enterprise to forward calls through the SIP Connect interface is considered a basic requirement. In order to forward a call the SIP-PBX can do either of the two alternatives:

- Forwarding by initial INVITE
- Forwarding by Call deflection

An SP-SSE **MUST** be able to accept forwarded calls from a SIP-PBX. Note that an SP-SSE may enforce policies that include a variety of restrictions on calls forwarded from an untrusted SIP-PBX (e. g.,

mandating the inclusion of a "History-Info" header field [RFC 7044] with a "From" header field that does not correspond to an Enterprise Public Identity assigned to the SIP-PBX). These policies are outside the scope of the SIPconnect Technical Recommendation.

11.1 Forwarding by New INVITE

To forward with a new INVITE, the SIP-PBX **MUST** send an initial INVITE request to the SP-SSE, populated as specified in Section 10.2 and with:

- The request-URI identifying the forwarded-to target destination.
- A History-Info header containing the Enterprise Public Identity of the forwarding user, formatted according to section 9, and in accordance with [RFC 7044].
- A From header containing the original calling party identity.
- A P-Asserted-Identity header containing a valid identity of the forwarding SIP-PBX.

A simplified example call flow for Call Forwarding is shown in Figure 2. The initial call leg is on dialog [1] and the forwarded leg is on dialog [2]. Note: For a SIP-PBX that acts as a B2BUA, dialog [1] and dialog [2] are generally different dialogs. For a SIP-PBX that acts as a proxy server, dialog [1] and dialog [2] are the same dialog.

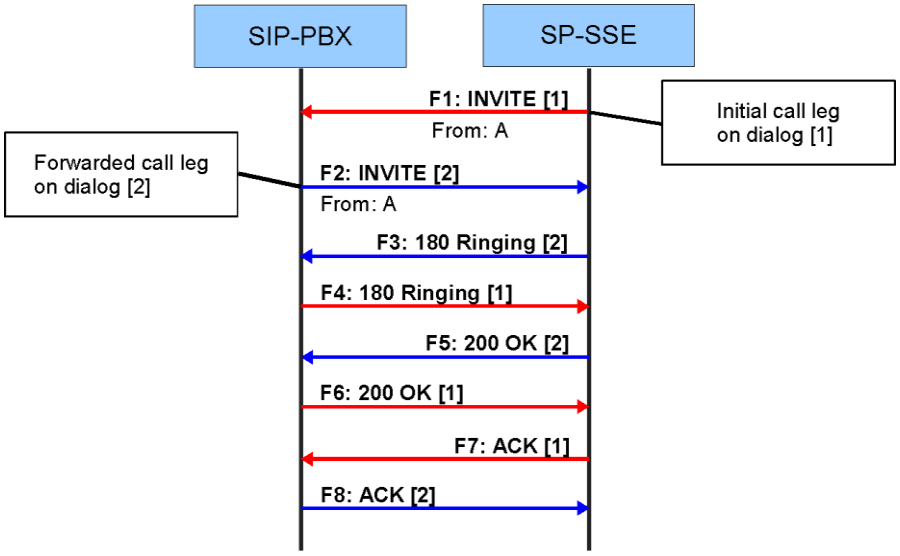

Figure 2: Call Forward by New INVITE

11.2 Forwarding by Call Deflection (302)

To deflect a call from the SIP-PBX, the SIP-PBX responds to the INVITE by a deflection response, 302, which includes the target in the Contact header. The SP-SSE **MUST** execute network based forwarding as a result of receiving such 302 response from the SIP-PBX (as opposed to relaying the response to the remote UE).

12. Call Transfer

The ability for the SIP-PBX or the SP-SSE to transfer calls that cross the SIPconnect 2.0 interface is considered a basic requirement in this Technical Recommendation. This section specifies a set of SIP primitives that can be used to support the transfer of calls that cross a SIPconnect 2.0 interface.

12.1 Overview

Call transfer can be accomplished by the use of REFER requests (a "proxy model") in accordance with [RFC 5589], or by the use of one or more INVITE/re-INVITE requests (a "third party call control model"). The SP-SSE and SIP-PBX **MUST** support the use of INVITE/re-INVITE for initiating and responding to call transfers.

Support for initiating and responding to call transfers using the REFER method is outside the scope of SIPconnect. SIPconnect has selected the use of INVITE/re-INVITE for call transfer because that is what is commonly deployed at the time of writing and because of Enterprise or Service Provider policies that might require rejection of received REFER requests (e.g., because of charging considerations).

12.1.1 Blind Transfer

Blind transfer, known as basic transfer in [RFC 5589], is where a new call is established from the transferee to the transfer target and the transferor drops out immediately, without waiting for the transfer target to answer.

A SIP-PBX acting as a B2BUA can accomplish blind transfer using INVITE/re-INVITE as follows. Assuming that the call with the transferee crosses the SIPconnect interface and the transfer target is reachable across the SIPconnect interface, the SIP-PBX sends a new dialog INVITE request to the SP-SSE targeted at the transfer target and sends a re-INVITE request to the SP-SSE on the existing dialog with the transferee, changing the SDP for this dialog, so media goes between the transferee and transfer target.

The SP-SSE can accomplish blind transfer in a similar manner using INVITE/re-INVITE. The INVITE and re-INVITE transactions are used to achieve an offer-answer exchange between the transferee and transfer target.

For example, the SIP-PBX can send an offerless INVITE request towards the transfer target. In response, the transfer target supplies an SDP offer, which the SIP-PBX includes in a re-INVITE request towards the transferee. The SIP-PBX then forwards the SDP answer from the transferee in an ACK request towards the transfer target. If the transferee is within the SIP-PBX, only the INVITE transaction towards the transfer target will cross the SIPconnect interface. If the transfer target is within the SIP-PBX, only the re-INVITE request towards the transferee will cross the SIPconnect interface.

A simplified example call flow for Blind Transfer is shown in Figure 3. Note that the initial call leg is on dialog [1] and the transferred leg is on dialog [2]. It should be noted that this call flow is illustrative only, and does not mandate a specific implementation. More complex call flows may be required to support feature interactions encountered in real-world deployments; for example when the transfer target has a terminating feature that sends early media toward the transferee.

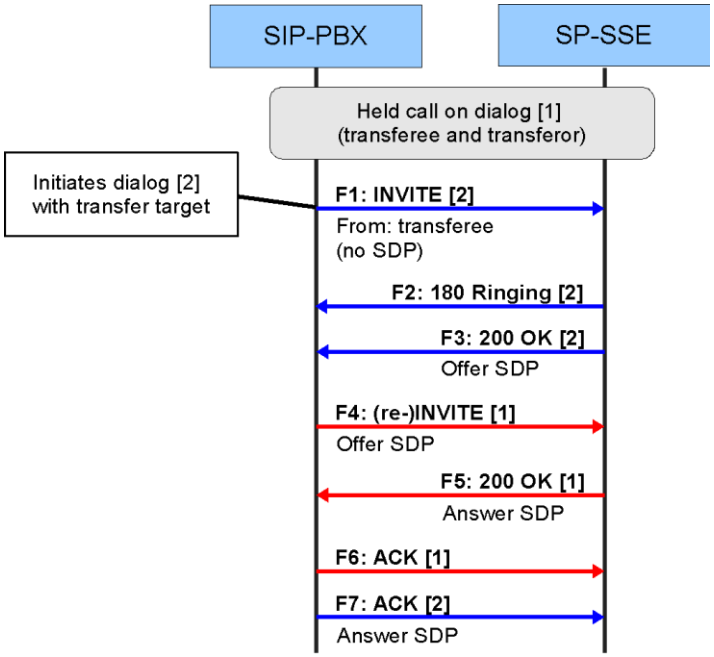

Figure 3: Blind Transfer

Requirements for the support of re-INVITE are given in Section 12.2.

12.1.2 Attended Transfer

Attended transfer is where the transferor has already established a new call to the transfer target and the transfer target has answered. Transfer then involves replacing the two existing calls (with the transferee and with the transfer target) by a single call.

The SIP-PBX can accomplish attended transfer using re-INVITE as follows. Assuming that each call crosses the SIPconnect 2.0 interface, the SIP-PBX sends a re-INVITE request to the SP-SSE on each of the existing dialogs. The two re-INVITE transactions are used to achieve an offer-answer exchange between the transferee and transfer target.

For example, the SIP-PBX can send an offer-less re-INVITE request towards the transfer target. In response, the transfer target supplies an SDP offer, which the SIP-PBX includes in a re-INVITE request towards the transferee. The SIP-PBX then forwards the SDP answer from the transferee in an ACK request towards the transfer target. If the transferee is within the SIP-PBX, only the re-INVITE transaction towards the transfer target will cross the SIPconnect 2.0 interface. If the transfer target is within the SIP-PBX, only the re-INVITE transaction towards the transferee will cross the SIPconnect 2.0 interface. A simplified example call flow for Attended Transfer is shown in Figure 4. Note that the initial call leg is on dialog [1] and the transferred leg is on dialog [2].

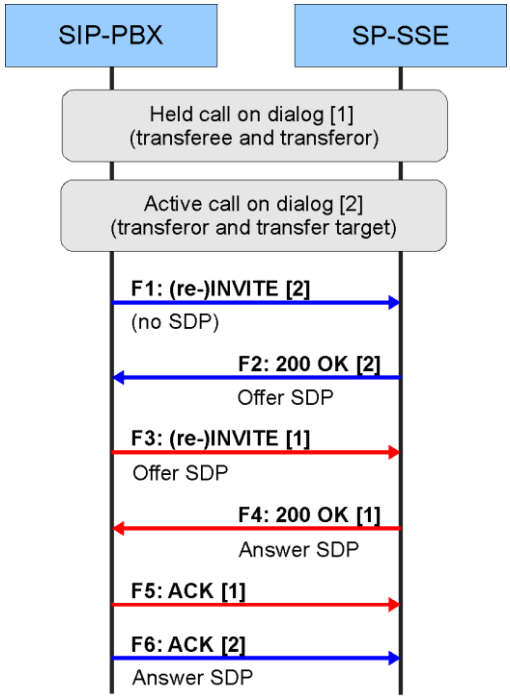

Figure 4: Attended Transfer

The SP-SSE can accomplish attended transfer in a similar manner using re-INVITE.

Requirements for the support of re-INVITE are given in Section 12.2.

12.2 Requirements for Use of the re-INVITE Method in the Context of Call Transfer

The SIP-PBX and the SP-SSE **MUST** support both sending and receiving a re-INVITE request with an SDP offer, and sending and receiving a re-INVITE request without an SDP offer.

13. Emergency Services

The SIP-PBX **MUST** have a dial plan that recognizes emergency calls.

When a SIP-PBX routes a call recognized as an emergency call to the SP-SSE, it **MUST** populate the Request-URI using a dial string URI, as specified in Section 10.2.1, that contains the national emergency services number.

The SIP PBX **MUST** include the identity of the caller in the "P-Asserted-Identity" header field, as described in Section 10.2.3, and in the "From" header field, as described in Section 10.2.4, except in territories where the SIP-PBX is required to include other information (such as a Location Identification Number) in one of these header fields. The SIP PBX **MUST NOT** withhold the "P-Asserted-Identity" header field for privacy reasons and **MUST NOT** anonymize the "From" header field.

The SP-SSE **MUST** be able to recognize emergency calls based on the presence of the agreed emergency services number in the Request-URI.

If an originating session is an emergency session, then SIP session limits do not apply. The SP-SSE **MUST NOT** apply SIP session limits to emergency calls originated by the SIP-PBX. Note that this does not preclude the SP-SSE rejecting the emergency call for other reasons including local congestion or exceeding limits explicitly applicable for emergency calls.

13.1 Location Conveyance

Information relating to the location of a SIP-PBX user or device **MAY** be provided depending on local regulatory requirements. The SIP-PBX when providing location **SHOULD** do so using the SIP Geolocation Header field as specified in [RFC 6442], location **MAY** be provided by value or by reference.

When location is provided by value it **MUST** be structured in accordance with the formats and rules defined in [RFC 5491] and transported in a PIDF-LO as defined in [RFC 4119] . SIP-PBX and SIP-SSE implementations which add additional MIME bodies, including PIDF-LO, should note that [RFC 5621] , which specifies how message bodies are handled in SIP, states that the default value for the content-disposition 'handling' parameter is "required". Therefore to prevent calls being rejected by a SIP entity that does not support a specific MIME body the SIP-PBX MUST set the content-disposition 'handling' parameter to "optional". This is especially important for emergency calls.

The SIP-PBX **SHOULD** insert a Geolocation-Routing Header field with a value of "yes" if and only if it wants the call to be routed based on the location information it inserted. The SP-SSE **MAY** use the location information sent by the SIP-PBX together with a Geolocation-Routing Header field set to "yes" to route an emergency call to the local PSAP.

To ensure the privacy of location when conveying location information over the SIPconnect interface by value it is **RECOMMENDED** that TLS is used as the SIP transport according to sections 16.2 (Registration Mode) or 17.2 (Static Mode) meaning that all calls are **RECOMMENDED** to use TLS, not just those containing the Geolocation Header field.

13.2 Additional Data

The use of additional data as described in [RFC 7852] is for further study.

14. Media and Session Interactions

14.1 SDP Offer/Answer

A SP-SSE/SIP-PBX acting on behalf of a Media Endpoint that originates and/or terminates RTP traffic **MUST** utilize the Session Description Protocol (SDP) as described in [RFC 4566] in conjunction with the offer/answer model described in [RFC 3264] to exchange media capabilities (IP address, port number, media type, send/receive mode, codec, DTMF mode, etc).

SIP-PBXs and SP-SSEs **MUST** be capable of receiving INVITE requests without an SDP offer and supplying an SDP offer in an appropriate response, in accordance with [RFC 3261].

During a call, media capability negotiation **MAY** be initiated by either end, for the purpose of verifying dialog state or for other reasons, and experience has shown that some SIP implementations don't handle offers with unchanged SDP correctly.

A SP-SSE/SIP-PBX that participates in SDP offer/answer negotiation **MUST** be prepared to accept additional offers containing SDP with a version that has not changed, and **MUST** generate a valid answer (which could be the same SDP sent previously, or could be different).

A SP-SSE/SIP-PBX that sends additional SDP offers with the same version **MUST** be prepared to accept answers with SDP which may be the same as the previously received SDP, or may be different.

A SP-SSE/SIP-PBX that sends SDP with a change compared to the previously sent SDP **MUST** increase the version number in the o-line, in accordance with [RFC 4566].

SIP-PBX and SP-SSE implementations sending changes to negotiated media capabilities via SIP re-INVITE **MUST** support [RFC 3261], Section 14 "Modifying an Existing Session". SIP UPDATE **MAY** be used for this purpose when both endpoints advertise support for [RFC 3311].

14.2 Media Transport

A Media Endpoint **MUST** send and receive voice samples using the real-time transport protocol (RTP) as described in [RFC 3550] and **SHOULD** support SRTP [RFC 3711] using SDP security descriptions [RFC 4568] for the key exchange, as specified in Section 14.4.

RTP itself comprises two parts: the RTP data transfer protocol, and the RTP control protocol (RTCP). RTCP is a fundamental and integral part of RTP, and MUST be implemented.

Any Media Endpoint that originates and/or terminates RTP or SRTP traffic over UDP **MUST** use the same UDP port for sending and receiving session media (i.e. symmetric RTP).

Any Media Endpoint that originates and/or terminates RTP traffic **MUST** be capable of processing RTP packets with a different packetization rate than the rate used for sending.

14.3 Audio Profile

Any Media Endpoint that originates and/or terminates voice traffic **MUST** support the [ITU-T G.711] µ-Law and A-Law PCM codecs with a packetization rate of 20 ms. Any device intended for low-bandwidth operation **SHOULD** support [ITU-T G.729] codecs with a packetization rate of 20 ms.

In the absence of a specific indication that receiving G.711 discontinuously using the Comfort Noise (CN) payload type defined in [RFC 3389] is supported, the SIP-PBX or SP-SSE **MUST** assume that the far end Media Endpoint does not support receiving G.711 discontinuously. In order to indicate in SDP that receiving G.711 discontinuously is supported by the local Media Endpoint, the SIP-PBX/SP-SSE **MUST** include payload type 13 in the "m=audio" line as described in [RFC 3389].

It is possible that the Media Endpoint associated with the Offerer or Answerer supports receiving CN packets but not sending them. In that case, it would be perfectly legal to send SDP with Audio Video Profile (AVP) 13 in the "m=audio" line. The Offerer or Answerer in this case is expressing its Media Endpoint's willingness to receive CN packets even if its Media Endpoint never sends any itself.

In the absence of a specific indication that receiving G.729 discontinuously (i.e., [ITU-T G.729] Annex B) is not supported, the SP-SSE/SIP-PBX **MUST** assume that the far end Media Endpoint supports receiving G729 discontinuously. In order to indicate in SDP that receiving G,729 discontinuously is not supported by the local Media Endpoint, the "a=fmtp:18 annexb=no" attribute **MUST** be included. See Section 2.1.9 in [RFC 4856].

It is possible that the Media Endpoint associated with the Offerer or Answerer supports receiving [ITU-T G.729] Annex B but not sending it. In that case, it would be perfectly legal to send SDP with "annexb=yes" (or without any parameter since that means the same thing). The Offerer or Answerer in this case is expressing its Media Endpoint's willingness to receive [ITU-T G.729] Annex B packets, even if the local Media Endpoint never sends any itself.

14.4 Media Security Using Secure RTP (SRTP)

Secure RTP [RFC 3711] is a RTP profile which provides confidentiality, authentication and replay protection for both RTP and RTCP.

SIPconnect 2.0 Media Endpoints **SHOULD** secure the media using SRTP [RFC 3711] and when doing so **MUST** use SDP Security Descriptions [RFC 4568] for the necessary key exchange.

SDP Security Descriptions allows for negotiation of various crypto-suites and SRTP parameters in the a=crypto: attribute as defined in [RFC 4568]. As a least common denominator that allows for successful interoperability, the Offerer **MUST** include at least one a=crypto: attribute that uses the following values:

- crypto-suite: AES_CM_128_HMAC_SHA1_80
- key||salt: dynamically and randomly calculated for each new offer, unique to the entire SDP message and unique per direction. This means that in case of a new SDP offer/answer exchange, the Offerer **SHOULD** include a new master key and master salt that is unique and generated independently from the key and salt provided during the previous SDP offer/answer exchange.

The Offerer **MUST NOT** include the following elements in the above a=crypto: attribute:

- lifetime:
- MKI:length
- any session parameters, e.g. KDR, UNENCRYPTED_SRTP, UNENCRYPTED_SRTCP, UNAUTHENTICATED_SRTP, FEC_ORDER, FEC_KEY and WSH.

Since the a=crypto: attribute carries the key material in cleartext, the call signaling MUST be protected by TLS as described in section 8.1.

Media Security is a configuration option that is agreed between the SIP-PBX administrator and the Service Provider therefore SIPconnect 2.0 does not specify any mechanism for negotiating media security. Negotiation and fallback mechanisms are for further study and may be included in future versions of this specification.

Media Endpoints **SHOULD** use the confidentiality mechanisms in SRTP and SRTCP to ensure media confidentiality as described in [RFC 3711].

Media Endpoints **SHOULD** use the integrity mechanisms in SRTP and SRTCP to ensure media integrity as described in [RFC 3711].

Media Endpoints **SHOULD** use the replay protection mechanism for protecting both SRTP and SRTCP as described in [RFC 3711].

Subject to the above recommendations, the SDP offer **MAY** include further a=crypto: attributes allowing for other crypto-suites or carrying any valid combination of optional elements that were disallowed for the mandatory a=crypto: attribute from above. The recommendation for using new key material in subsequent SDP offer/answer exchanges remains valid also when one of these further a=crypto: attributes is negotiated. Usage of new key material is motivated due to SIP forking and due to Transfer, in which case the offerer's key is distributed to several peers.

14.5 Transport of DTMF Tones

A SP-SSE/SIP-PBX **MUST** advertise support for telephone-events [RFC 4733] in its SDP on behalf of any Media Endpoint that supports receiving DTMF digits using [RFC 4733] procedures.

Any Media Endpoint that supports receiving DTMF **MUST** support [RFC 4733] procedures.

Any Media Endpoint that supports sending DTMF **MUST** use the [RFC 4733] procedures to transmit DTMF tones using the RTP telephone-event payload format, provided that the other side has advertized support for receiving [RFC 4733] in the offer/answer exchange.

For any local Media Endpoint that supports receiving telephone-event packets, the SIP-PBX or SP-SSE **MUST** include the supported events in an "a=fmtp:" line as is described as mandatory in [RFC 4733].

To provide backward compatibility with [RFC 2833] implementations, any Media Endpoint **MUST** be prepared to receive telephone-event packets for all events in the range 0-15 and a SIP-PBX or SP-SSE **MUST** be prepared to accept SDP with a payload type mapped to telephone-event, even if it does not have an associated "a=fmtp" line.

14.6 Echo Cancellation

Any Media Endpoint that can introduce echo **MUST** provide [ITU-T G.168]-compliant echo cancellation.

14.7 FAX Calls

In-band fax transmissions are especially problematic over packet networks, especially for calls that traverse the public Internet or other network that doesn't offer adequate QoS.

Media Endpoints that support fax (e.g., a SIP media server that can originate/terminate faxes) and Media Endpoints that can act as a T.30 gateway (e.g., a Media Endpoint that supports an RJ11 analog telephone interface) **SHOULD** support the [ITU-T T.38] Recommendation.

Media Endpoints that support [ITU-T T.38] **MUST** support User Datagram Protocol Transport Layer (UDPTL) transport.

14.8 Call Progress Tones

Media Endpoints **SHOULD** locally generate call progress tones or announcements, or other suitable indications, when the response to an INVITE request indicates call failure. Selection of the particular tone or announcement for a given response code might depend on local practices and regulation, but otherwise is left to the equipment manufacturer's discretion.

14.9 Ringback Tone, In-band Tones, and Early Media

The delivery of in-band announcements and call progress tones from the Service Provider, or from a SIP-PBX, to a caller before a call is answered is achieved through early media.

According to the reference architecture, section 3, it is actually the Media Endpoint which handles the early media and therefore the function split between the SIP-PBX and Media Endpoint is implementation dependent. Therefore the Media Endpoint requirements specified in this section could be performed by the SIP-PBX, for example when the Media Endpoint is not a SIP endpoint.

When acting as a call originator the Media Endpoint **MAY** indicate support for the P-Early-Media header [RFC 5009] by including a P-Early-media header field set to "supported" in the INVITE request.

When acting as a call originator, the SIP-PBX, upon receipt of a 180 provisional response message without SDP (whether reliable [RFC 3262] or unreliable) **MUST** instruct the Media Endpoint to play local ringback tone to the user. Upon receipt of SDP in any 18x provisional response message (reliable [RFC 3262] or unreliable), the SIP-PBX **MUST** forward this information to the Media Endpoint. If the SIP-PBX indicated support for P-Early-Media header it **MUST** also forward any P-Early-Media header received to the Media Endpoint.

When acting as a call terminator and expecting the originating end to provide local ringback tone, the Media Endpoint **MUST NOT** send RTP packets to the originator if a 180 provisional response message was sent. In addition, if the received INVITE request contained the P-Early-Media header indicating "supported" the Media Endpoint **MAY** include a P-Early-Media header in the 180 response.

When acting as a call terminator and wanting to provide tones and announcements during an early dialog to the caller, the Media Endpoint **MAY** include the P-Early-Media header indicating "sendonly" or "sendrecv" in any response containing SDP.

A Media Endpoint that does not support the P-Early-Media header [RFC 5009] or does not receive a P-

Early-Media header in the 18x response, on receipt of an instruction to play local ringback tone, **MUST** do so until it receives valid RTP packets or is instructed by the SIP-PBX that the call has been answered. On receipt of valid RTP packets, a Media Endpoint **MUST** disable any local ringback tone and play the received media. A Media Endpoint, on receipt of information concerning received SDP, **MAY** use the information to determine whether RTP packets received are valid and **MAY** discard RTP packets arriving before that time.

A Media Endpoint **MUST** play any received early media when:

- The Media Endpoint supports the P-Early-Media header [RFC 5009], and receives a 18x response containing a P-Early-Media header field with "sendonly" or "sendrecv", and
- SDP is present in the 18x or was present in an earlier 18x response.

14.10 Putting a Session on Hold

A 2-way session can be put on hold by using an offer-answer exchange (Section 14.1) and the directionality attributes as described below.

When the hold initiator (which may be the SIP-PBX or SP-SSE acting transparently as Media Endpoint) provides music-on-hold (MOH) treatment:

- The MOH source in the SP-SSE/SIP-PBX is based on local policy.

- The hold initiator **MUST** set the SDP directionality attribute to "a=sendonly".

If the hold initiator does not provide MOH, it **MUST** set the SDP directionality attribute to "a=inactive" or "a=sendonly". The attribute "a=inactive" is **RECOMMENDED** because it provides an indication to the held entity that MOH is not being provided by the hold initiator.

A SP-SSE/SIP-PBX **MAY** support the ability to receive SDP session descriptions that have the ‘c=’ field set to all zeros (0.0.0.0), when the addrtype field is IPv4, for support of non-compliant remote SIP signaling entities that use this deprecated syntax from [RFC 2543], rather than the "a=sendonly" or "a=inactive" syntax specified in [RFC 3264]. However, note that the deprecated syntax is not supported in specifications of other Standard Bodies, such as 3GPP IMS, used by some incumbent carriers for their SIP deployments.

Note - If the SP-SSE operates an IMS Announcement/MOH Application Server (acting as described in [3GPP TS 24.610] the SP-SSE MUST ensure that this Application Server is not triggered for the purpose of providing tones and announcements, including MOH, if the sender of the offer/answer with a=sendonly is connected via a SIP-Trunk complying to this specification.

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

15. IPv6

The SIP Forum expects that all networks will eventually migrate from IPv4 to IPv6. Therefore, this recommendation defines a migration path, which starts with the Service Provider providing connectivity over both IPv4, for those Enterprises that have yet to migrate, and IPv6, for those Enterprises that have migrated.

For the sake of simplicity and to avoid interoperability issues, neither the Service Provider nor the Enterprise on the SIPconnect interface is **REQUIRED** to support a dual stack implementation. In particular, media negotiations via ICE ([RFC 5245]), ALTC ([RFC 6947]), or similar mechanisms are out of scope.

The work in the SIP Forum IPv6 working group has proven that SIP dual stack operation has some unresolved issues which require further work in the standards bodies to resolve. Therefore the SIP Forum strongly recommends to use of different service names for IPv4 and IPv6 address families such that the SRV record for the SIPconnect 2.0 interface points to hosts with either A or AAAA records, not both.

The Service Provider **MUST** support connectivity over the SIPconnect interface via IPv4 or IPv6. The Service Provider **SHOULD** support connectivity over both IPv4 and IPv6.

The Enterprise **MUST** support connectivity via either IPv4 or IPv6. An Enterprise **MAY** split its subscribers between an IPv4-connected network and an IPv6-connected network; however, this split must be considered as two separate instances of the SIPconnect interface.

A Service Provider or Enterprise that supports connectivity over IPv4 **MUST** use IPv4 for both signaling and media.

A Service Provider or Enterprise that supports connectivity over IPv6 **MUST** use IPv6 for both signaling and media.

The SP-SSE and the SIP-PBX **MUST** support IPv6 address syntax in SIP requests and responses, even if the SIPconnect interface uses IPv4. This includes IPv6 addresses in SIP URIs as well as the Via header field.

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

16. Annex A: Registration Mode

As stated in Section 7, in Registration mode, the SIP-PBX conveys its SIP signaling address to the Service Provider Network using the SIP registration procedure. In effect, the SIP-PBX registers with the Service Provider Network, just as a directly hosted SIP endpoint would register. However, because a SIP-PBX has multiple Enterprise Public Identities, it needs to register a contact address on behalf of each of these. Rather than performing a separate registration procedure for each Enterprise Public Identity, Registration mode makes use of the mechanism in [RFC 6140] to achieve multiple registrations using a single REGISTER transaction.

According to this mechanism, the SIP-PBX delivers to the SP-SSE in the "Contact" header field of a REGISTER request a template from which the SP-SSE can construct contact URIs for each of the AORs (Enterprise Public Identities) assigned to the SIP-PBX, and thus can register these contact URIs within its location service. These registered contact URIs can then be used to deliver to the SIP-PBX inbound requests targeted at the AORs concerned. The mechanism can be used with AORs comprising SIP URIs based on global E.164 numbers and the Service Provider's domain name or sub-domain name. This is consistent with requirements for Enterprise Public Identities for Registration mode in Section 9.

As a pre-requisite, the SIP-PBX and the SP-SSE need to be provisioned with the set of E.164 numbers (and hence the set of Enterprise Public Identities) assigned to the SIP-PBX and with a Registration AOR for use in the "To" header field of the REGISTER request. The SIP-PBX **MUST** be capable of provisioning any format of SIP-URI as the Registration AOR, in order to accommodate SP-SSE requirements (i.e., the Registration AOR is not subject to the same constraints as Enterprise Public Identities and could, for example, be an "email-style" SIP URI).

The requirements of this section apply only to SIP-PBXs and SP-SSEs that support Registration mode.

16.1 Locating SIP Servers

16.1.1 Enterprise Requirements

The SIP-PBX **MUST** provide its SIP signaling address(es) and port(s) to the SP-SSE using the SIP registration procedure described in Section 16.4.

The SIP-PBX **MUST** be capable of obtaining information about the SP-SSE, using the procedure described in Section 17.1.1.2.

16.1.2 Service Provider Network Requirements

The SP-SSE **MUST** make its SIP signaling address(es) and port(s) available to the Enterprise Network as specified in Section 17.1.2.1.

The SP-SSE **MUST** obtain the SIP-PBX signaling address/port using SIP registration, as described in Section 16.4.

16.2 Signaling Security

In Registration mode, the following rules for using TLS apply:

- Both SIP-PBX and SP-SSE **MUST** support the TLS Server Authentication model, whereby the SP-SSE (acting as TLS server), provides its certificate to the SIP-PBX (acting as TLS client) as part of the TLS establishment phase. Note that this is essentially the same model as secure TLS/SSL connections on the Public Internet for HTTP. This avoids the need for the SIP-PBX to have a certificate. However, a consequence is that the SIP-PBX **MUST** initiate the TLS session (in order to act as the TLS client).
- The SIP-PBX **MUST** be capable of initiating the establishment of a TLS session.
- The SIP-PBX **MUST** be capable of being provisioned with either a certification authority certificate or with a copy of the certificate the SP-SSE plans to use (or a fingerprint thereof). However, the SIP-PBX does not need to be provisioned with a certificate.
- The SIP-PBX **MUST** validate the certificate received during TLS establishment using the path validation procedure described in [RFC 5280].
- The SIP-PBX **SHOULD** verify the status of the certificate received during TLS establishment. For TLS certificate revocation checks, Online Certificate Status Protocol (OCSP) **SHOULD** be used. It is **RECOMMENDED** to use OCSP Stapling ([RFC 6066] and [RFC 6961]) to avoid delays in call setup.
- The SIP-PBX **MUST** be capable of being configured to require use of TLS to initiate a session.

In Registration mode, when the SIP-PBX is configured to require use of TLS with an SP-SSE, the following requirements apply:

- The SIP-PBX **MUST** initiate the establishment of the TLS session.
- The SIP-PBX **MUST NOT** utilize other transports (UDP or TCP), even if the SP-SSE indicates that these are available via configuration of DNS NAPTR and/or SRV resource records.

When the SP-SSE is configured to accept TLS connections, the following requirements apply:

- When configuring DNS NAPTR and/or SRV resource records in accordance with Section 15.1.2, the SP-SSE **SHOULD** indicate support for TLS.
- The SP-SSE **MUST** be configured with a verifiable digital certificate to secure a TLS session.
- The SP-SSE **MUST** use certificates that are signed by a third party certification authority unless the certificates can be validated through some other means, such as being pre-installed at the SIP-PBX or signed by the SP-SSE itself.

When using TLS (as a result of being configured to require use of TLS, or as a result of discovering the availability of TLS from DNS), the SIP-PBX **MUST** establish a TLS connection (if not already established) prior to registration and **MUST** use that connection to deliver the REGISTER request and all subsequent SIP messages to the SP-SSE. The SP-SSE **MUST** authenticate the SIP-PBX using SIP digest authentication, as specified in Section 16.4, and reject the REGISTER request if authentication fails. Following successful registration, the SP-SSE **MUST** use a TLS connection that is authenticated as a connection to this SIP-PBX to deliver all SIP requests to the SIP-PBX.

The SIP-PBX and the SP-SSE **MUST** avoid closing down the TLS connection, other than in exceptional circumstances (e.g., for maintenance). The SIP-PBX is responsible for attempting to keep the connection alive, and if the TLS connection fails, the SIP-PBX is responsible for re-establishing the TLS connection at the earliest opportunity and registering again, in order that the SP-SSE can deliver SIP requests to the SIP-PBX at any time (e.g., in support of incoming calls).

16.2.1 The Use of transport=tls Parameter

When a SIP-PBX registers, the SP-SSE **MUST** ignore the transport=tls parameter in the "Contact" header field URI.

The reachability through TLS is indirectly determined by the SP-SSE because the registration itself is using TLS.

16.3 Firewall and NAT Traversal

Any IP addresses contained within the header fields and message body parts (e.g. SDP) of SIP messages exchanged between the Service Provider and Enterprise Networks **MUST** be publicly routable addresses, unless the Service Provider Network is providing an implicit NAT traversal function or the two are using a private VPN-style address space.

16.4 Registration

The SIP-PBX and SP-SSE **MUST** support multiple AOR registration in accordance with [RFC 6140], using the provisioned Registration AOR and the set of provisioned Enterprise Public Identities, even if there is only a single provisioned Enterprise Public Identity.

In the REGISTER request, the SIP-PBX **MUST** include a Contact URI in accordance with [RFC 6140] using a suitable domain part, e.g., the SIP-PBX's IP address. The SIP-PBX **MUST** insert the Registration AOR in the "From"and "To" header fields of the REGISTER request.

The SIP-PBX and SP-SSE **MUST** support the authentication mechanisms outlined in Section 16.6 for digest authentication for the REGISTER requests, using a user name and password agreed to by both parties.

16.4.1 Registration Failures

This section details the behavior requirements for the SP-SSE and SIP-PBX for Registration failure scenarios.

16.4.1.1 Failure of SIP-PBX to Reach the SP-SSE

If the SIP-PBX fails to receive any response to a REGISTER request in Timer_F time (typically 32 seconds) or encounters a transport error when sending a REGISTER request, the SIP-PBX **MUST** consider the SP-SSE unreachable and try to register with an alternate SP-SSE address if it has one. If the SIP-PBX has an established connection-based transport (e.g., TCP) to the SP-SSE, and Timer_F expires or a transport error is encountered as above, it **MUST** try to re-establish a connection to the same SP-SSE before considering it unreachable, by resetting Timer_F and sending a new REGISTER request. The SIP-PBX **MUST NOT** attempt to re-establish the connection to the same SP-SSE more than once before considering the SP-SSE unreachable. This allows for cases where the SP-SSE lost previous transport connection state but is otherwise reachable, such that the SIP-PBX will try a second time and only consider the SP-SSE unreachable if that second attempt fails.

If no SP-SSE is reachable, or no alternates are available, the SIP-PBX **MUST** delay reattempting Registration for 30 seconds, and increasing this delay value by doubling it for each successive delivery failure until delivery succeeds, up to a maximum value of 960 seconds.

Note that receiving an explicit non-2xx final response from the SP-SSE does not constitute a delivery failure. Instead, behaviors for such final responses are noted in the following sections.

16.4.1.2 Redirection of SIP-PBX from SP-SSE

The SP-SSE **MUST NOT** issue a 302 Moved Temporarily redirect response to a REGISTER request, to get the SIP-PBX to Register with an alternate SP-SSE address identified by the Contact URI in the response.

16.4.1.3 Unknown SIP-PBX Identity

The SP-SSE **MUST** authenticate all REGISTER requests. If the authentication fails, the SP-SSE **SHOULD** issue a new authentication challenge. The SP-SSE should not validate any data in the request before a successful authentication.

A SIP PBX that fails authentication **MUST** consider the Registration attempt to have failed, and notify the SIP-PBX administrator if possible through some means. The SIP-PBX **SHOULD** follow the backoff procedures defined previously in Section 16.4.1.1.

16.4.1.4 Incorrect SIP-PBX Password

If the digest challenge response of the SIP-PBX in its REGISTER request is stale or invalid, the SP-SSE **MUST** issue one of the following response codes, unless the SP-SSE is configured to silently discard these requests based on policy:

- a 401 Unauthorized,
- a 407 Proxy Authentication Required, or

- a 403 Forbidden

If a SIP-PBX receives more than three responses of 401, 407 or 403 in aggregate, without a different response other than one of those in between, then the SIP-PBX **MUST** consider the Registration attempt to have failed, and notify the SIP-PBX administrator if possible through some means. The SIP-PBX **SHOULD** follow the backoff procedures defined previously in Section 16.4.1.1.

16.4.1.5 Other Servers Unreachable from SP-SSE

If an SP-SSE is unable to complete registration, it **MAY** issue a 480 Temporarily Unavailable response code for a REGISTER request. An SIP-PBX receiving such a response to a REGISTER request **MUST** act exactly as if delivery to the SP-SSE had failed per Section 16.4.1.1, and **MUST** follow the backoff procedures defined previously in Section 16.4.1.1.

16.4.1.6 SP-SSE Administratively Disabled or Overloaded

An overloaded SP-SSE **MUST** generate a 503 Service Unavailable or 500 Internal Error response code to a REGISTER request, unless it is silently discarding requests due to overload, and **SHOULD** include a "Retry-After" header field value indicating how long the SIP-PBX **SHOULD** wait before re-attempting a REGISTER request to the same SP-SSE.

This "Retry-After" header field value **SHOULD** include an element of randomness so that all served SIP-PBXes don't become synchronized and repeatedly attempt to register en mass.

A SIP-PBX receiving such a response **MUST** support the "Retry-After" header field, and **MUST** honor the value as follows: if the value is 32 seconds or less, it **MUST** wait the requested time and retry the request to the same SP-SSE; if the value is larger, it **MUST** remember the value for that SP-SSE address instance, and try any alternate SP-SSE addresses it can. If an alternate SP-SSE can be successfully reached and Registration succeeds through the alternate, the SIP-PBX **MAY** discard the "Retry-After" value of the original. Otherwise, it **MUST** wait to reattempt registration to the original SP-SSE for the "Retry-After" interval.

16.4.1.7 Other 4xx/5xx/6xx Responses

Any 4xx, 5xx or 6xx-class response to a REGISTER request not explicitly identified above **SHOULD** be treated in a similar manner as Section 16.4.1.1 unless it can automatically be resolved by the SIP-PBX internally - i.e., unless it is part of an explicit negotiation mechanism or procedure. It **SHOULD** be treated as a delivery failure with a maximum retry interval of 960 seconds (16 minutes), unless a longer "Retry-After" header field is specified.

16.4.2 Registration-related Failures for Other Requests

If a SIP-PBX encounters a transport error when attempting to contact the SP-SSE, encounters Timer F expiry (non-INVITE requests) or Timer B expiry (INVITE requests), or receives a 403 response for any non-REGISTER request, the SIP-PBX **MUST:**

- consider the request attempt to have failed,
- assume that the SIP-PBX's registration is no longer active at the SP-SSE, and
- notify the SIP-PBX administrator if possible through some means.

The SIP-PBX **SHOULD** attempt re-registration using the procedures defined previously in Section 16.4.1.1.

16.5 Maintaining Registration

It is important that registrations are maintained and, in the event of failure, are re-established quickly, since the SP-SSE depends on the SIP-PBX being registered in order to deliver inbound requests to the SIP-PBX. Where TCP (with or without TLS) is used, the TCP connection needs to be maintained as the means for delivering inbound requests.

Because NATs and firewalls may drop a TCP connection through lack of use, measures need to be taken to keep the connection alive and detect whether it has been dropped. Similarly, where UDP is used, it is necessary to keep the path through NATs and firewalls alive. Therefore the SIP-PBX **MUST** honor the REGISTER expiry time provided by the SP-SSE, and **MAY** send REGISTER requests more frequently if NAT and firewall policies require this.

If failure is detected a SIP-PBX **MUST** attempt reconnection, and if that fails **MUST** try an alternative SP-SSE if available, in accordance with Section 16.4.1.1.

16.6 Authentication

16.6.1 Authentication of the Enterprise by the Service Provider

The SP-SSE authenticates the SIP-PBX using SIP Digest authentication mechanism.

The SIP-PBX and SP-SSE **MUST** support the digest authentication scheme as described in Section 22.4 of [RFC 3261]. The Service Provider assigns the SIP-PBX a username and associated password that are valid within the Service Provider's domain (realm).

The following rules apply:

1. The SP-SSE **MAY** challenge any SIP request. The SIP-PBX **MUST** support receiving 401 Unauthorized and 407 Proxy Authentication Required from the SP-SSE. When so challenged by the SP-SSE, the SIP-PBX **MUST** respond with authentication credentials that are valid within the

Service Provider's realm (i.e. based on the username and password supplied by the Service Provider).

2. In order to avoid unnecessary challenges, the SIP-PBX **SHOULD** include its authentication credentials using the current nonce in each subsequent request that allows authentication credentials to be sent to the SP-SSE.

When Digest Authentication is used over a path that is not protected by TLS, the credentials used are subject to offline "dictionary attacks", and successful attackers can then make calls that are billed to the SIP-PBX. Credentials provided to the SIP-PBX **SHOULD** be selected with this threat in mind. For example, passwords that appear in dictionaries would be poor choices. The credentials used for Digest Authentication **SHOULD** be machine-generated to have at least 64 bits of cryptographic randomness and then delivered via an automated provisioning mechanism. Human-memorable passwords are not the best choices. Since no end user has to enter one of these passwords, it is practical to use strong credentials.

16.6.2 Authentication of the Service Provider by the Enterprise

Authentication of the Service Provider by the Enterprise is supported using TLS server authentication. If TLS is required (based on local configuration data), then the SIP-PBX **MUST** perform TLS server authentication as described in Section 16.2.

16.6.3 Accounting

Accounting places no special requirements on the SIPconnect 2.0 interface. The SP-SSE may generate billing records for calls originating from the SIP-PBX, based on the local policy of the Service Provider. The SIP-PBX is not **REQUIRED** to signal a billing number to the SP-SSE (i.e., the SP-SSE will be configured with the billing number associated with billable incoming calls from the SIP-PBX).

16.7 Routing Inbound Requests to the SIP-PBX

The SP-SSE **MUST** route inbound out-of-dialog requests targeted at Enterprise Public Identities to the registered SIP-PBX in accordance with [RFC 6140]. This means that the Request-URI will comprise a SIP-URI containing the user part of the target Enterprise Public Identity and the domain part of the registered contact for that AOR.

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

17. Annex B: Static Mode

In the Static mode, the Service Provider and Enterprise Networks view each other as peer networks. The SP-SSE is configured with the domain name of the Enterprise and is either configured with the static IP address of the SIP-PBX or obtains the IP address of the SIP-PBX via DNS.

17.1 Locating SIP Servers

17.1.1 Enterprise Requirements

17.1.1.1 Providing Enterprise Address to SP-SSE

The SIP-PBX **MUST** provide its SIP signaling address and port to the SP-SSE using one of the following mechanisms:

- DNS: The Enterprise Network ensures the existence of a publicly-accessible DNS server that is authoritative for its domain (or a sub-domain delegated by the Service Provider for use by the Enterprise). This DNS server **SHOULD** provide a DNS interface that supports NAPTR resource records and **MUST** provide a DNS interface that supports SRV resource records [RFC 2782].

- Configuration: The Enterprise Network provides information to allow the Service Provider to configure mapping of the Enterprise Fully Qualified Domain Name (FQDN) to the SIP-PBX signaling address/port and transport at the SP-SSE.

17.1.1.2 Obtaining SP-SSE Address

Except when a TLS connection already exists, the SIP-PBX **MUST** use one of the following mechanisms to obtain the address and port of the SP-SSE and the transport protocol (UDP, TCP or TLS) to be used:

- [RFC 3263] "Locating SIP Servers": SIP-PBX utilizes DNS NAPTR and SRV queries as described in [RFC 3263] to determine the IP address(es), transport protocol(s), and port number(s) of the SP-SSE(s) associated with the Service Provider's domain name. This option assumes that the SIP-PBX has been pre-configured with the domain name of the Service Provider Network.

- Configuration: One or more transport protocols and SIP signaling address(es)/port(s) of the SP-SSE are configured in the SIP-PBX. A configured SP-SSE signaling address **SHOULD** be in the form of a hostname that can be resolved through DNS A/AAAA resource records, rather than an IP address (see additional guidance in Section 15).

When a TLS connection already exists, the SIP-PBX **MUST** reuse that TLS connection for all SIP messages.

17.1.2 Service Provider Network Requirements

17.1.2.1 Providing SP-SSE Address to Enterprise

The SP-SSE **MUST** be reachable through a publicly-accessible DNS server. The DNS server **SHOULD** provide a DNS interface that supports NAPTR resource records and **MUST** provide a DNS interface that supports SRV resource records.

17.1.2.2 Obtaining the Enterprise Network Address

The SP-SSE **MUST** support both of the following mechanisms to obtain the address and port of the SIP-PBX and the transport protocol (UDP, TCP or TLS) to be used and, except when a TLS connection already exists, **MUST** use one of these mechanisms:

- DNS: SP-SSE utilizes DNS NAPTR and SRV queries for the pre-configured domain name of the Enterprise Network, as described in [RFC 3263], to determine the IP address, transport protocol, and port number of the SIP-PBX(s) associated with the Enterprise Network's domain name.

- Configuration: The mapping of the Enterprise FQDN to the SIP-PBX signaling address/port and transport protocol is statically configured in the SP-SSE. A configured SIP-PBX signaling address **SHOULD** be in the form of a hostname that can be resolved through DNS A/AAAA resource records, rather than an IP address (see additional guidance in Section 15).

When a TLS connection already exists, the SP-SSE **MUST** reuse that TLS connection for all SIP messages.

17.2 Signaling Security

When using TLS in SIP-PBX and SP-SSE in Static mode, the following general requirements apply:

- Both SIP-PBX and SP-SSE **MUST** support the TLS Mutual Authentication model, whereby both the SP-SSE and the SIP-PBX provide their respective certificate as part of the TLS establishment phase.
- Both SIP-PBX and SP-SSE **MUST** be able to initiate the establishment of a TLS session.
- Both SIP-PBX and SP-SSE **MUST** be capable of being provisioned with either a certification authority certificate or with a copy of the certificate the peer SIP endpoint plans to use (or a fingerprint thereof).
- Both SIP-PBX and SP-SSE **MUST** validate a certificate received during TLS establishment using the path validation procedure described in [RFC 5280].
- Both SIP-PBX and SP-SSE **SHOULD** verify the status of a certificate received during TLS establishment. For TLS certificate revocation checks, OCSP **SHOULD** be used. It is **RECOMMENDED** to use OCSP Stapling ([RFC 6066] and [RFC 6961]) therefore both sides

need to provide OCSP staples as well as understand OCSP staples along with the TLS certificates.

- Both SIP-PBX and SP-SSE **MUST** be capable of being configured to require use of TLS to initiate a session to a particular peer. When TLS is configured to be required for session initiation to a peer, a SIP-PBX or SP-SSE **MUST NOT** initiate sessions with other transports (UDP or TCP), even if the peer indicates that these are available via configuration of DNS NAPTR and/or SRV resource records.
- Both SIP-PBX and SP-SSE **MUST** be capable of being configured to require use of TLS to accept sessions initiated to it by a peer. When TLS is configured to be required to accept sessions initiated from all peers, a SIP-PBX **MUST NOT** advertise support for other transports (UDP or TCP), via configuration of DNS NAPTR and/or SRV resource records.
- Existing TLS connections **SHOULD** be reused by both the SP-SEE and the SIP-PBX.
- The SP-SSE **MAY,** by policy, refuse connections without SIP-PBX client certificates.

When a SIP-PBX is configured to accept TLS connections, the following requirements apply:

- When configuring DNS NAPTR and/or SRV resource records in accordance with Section 17.1.1.1, the SIP-PBX **SHOULD** indicate support for TLS.
- The SIP-PBX **MUST** be configured with a verifiable digital certificate to secure a TLS session.
- The SIP-PBX **MUST** be configured with a certificate signed by a third party certification authority unless the configured certificate can be validated through some other means, such as being pre-installed on the SP-SSE or signed by the SIP-PBX itself.

When an SP-SSE is configured to accept TLS connections, the following requirements apply:

- When configuring DNS NAPTR and/or SRV resource records in accordance with Section 17.1.2.1, the SP-SSE **SHOULD** indicate support for TLS.
- The SP-SSE **MUST** be configured with a verifiable digital certificate to secure a TLS session.
- The SP-SSE **MUST** be configured with a certificate signed by a third party certification authority unless the configured certificate can be validated through some other means, such as being pre-installed on the SIP-PBX or signed by the SP-SSE itself.

When a SIP-PBX is configured to use TLS without a local SIP-PBX certificate, the following requirements apply:

- The SIP-PBX **MUST** always be the initiator of the TLS connections.
- The SIP-PBX **MUST**, whilst operational, ensure that a TLS connection exists and is kept alive with the SP-SSE (i.e. the SIP-PBX cannot only rely on outbound SIP signaling to trigger establishment of TLS connections or it may not be able to receive calls from the SP-SSE).
- If all TLS connections are lost, the SIP-PBX **MUST** immediately establish new TLS connection(s) with the SP-SSE in order to maintain full operational state and reachability.
- After accepting the TLS connection, the SP-SSE **MUST** still authenticate the SIP-PBX. The SP-SSE **MAY** do so using SIP digest authentication by which it **MUST** authenticate the first request and **MAY** authenticate all subsequent requests. When using digest the SIP-PBX and SP-SSE

MUST support the authentication mechanisms defined in Section 16.6 and **MUST** follow the procedure define in section 16.4.1.4 for failure case handling.

17.3 Firewall and NAT Traversal

The same considerations described for Registration mode in Section 16.3 apply here.

In addition, Static mode requires that both the SIP-PBX and the SP-SSE be directly reachable, which may require configuration of a static binding if NATs or firewalls are present between those elements.

17.4 Failover and Recovery

SIP-PBXes that require timely detection of SIP peer failure **MAY** use any of these mechanisms as keep-alives:

- Sending an OPTIONS request periodically, or
- Sending a carriage return/line feed periodically (TCP only – Note: this is a unidirectional CR/LF with no application layer acknowledgement. This can generate TCP resets if the SIP peer fails).

SIP-PBXes that support one of these mechanisms **MUST** also support a mechanism that allows the keep-alive interval to be configured.

17.5 Authentication

The SP-SSE and SIP-PBX authenticate each other using TLS mutual authentication. If TLS is required (based on local configuration data), then the SP-SSE and SIP-PBX **MUST** perform TLS mutual authentication as described in Section 17.2.

17.6 Routing Inbound Requests to the SIP-PBX

The SP-SSE **MUST** populate the Request-URI of the INVITE request with the Enterprise Public Identity of the called Enterprise user in the valid form defined in Section 9, or with a Contact URI provided by the SIP PBX in a previous request or response.

18. References

ITU-T E.164 International Telecommunications Union, "Recommendation E.164: The international public telecommunication numbering plan", May 1997, <http://www.itu.int>.

ITU-T G.168 International Telecommunications Union, "Recommendation G.168:Digital network echo cancellers", January 2007, <http://www.itu.int>.

ITU-T G.711 International Telecommunications Union, "Recommendation G.711: Pulse code modulation (PCM) of voice frequencies ", November 1988, <http://www.itu.int>.

ITU-T G.729 International Telecommunications Union, "Recommendation ITU-T G.729: Coding of speech at 8 kbit/s using conjugate-structure algebraic-code-excited linear prediction (CS-ACELP)", January 2007, <http://www.itu.int>.

ITU-T T.38 International Telecommunications Union, "Recommendation T.38: Procedures for real-time Group 3 facsimile communication over IP networks ", April 2007, <http://www.itu.int/rec/T-REC-T.38/e>.

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

RFC 2543 Handley, M., Schulzrinne, H., Schooler, E., and J. Rosenberg, "SIP: Session Initiation Protocol", RFC 2543, DOI 10.17487/RFC2543, March 1999, <http://www.rfc-editor.org/info/rfc2543>.

RFC 2782 A. Gulbrandsen, P. Vixie, L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.

RFC 2833 H. Schulzrinne, S. Petrack, "RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals", RFC 2833, May 2000.

RFC 3261 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.

RFC 3262 J. Rosenberg, H. Schulzrinne, "Reliability of Provisional Responses in Session Initiation Protocol (SIP), RFC 3262, June 2002.

RFC 3263 J. Rosenberg, H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers", RFC 3263, June 2002.

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

RFC 3311 J. Rosenberg, "The Session Initiation Protocol (SIP) UPDATE Method", RFC 3311, September 2002.

RFC 3323 J. Peterson, "A Privacy Mechanism for the Session Initiation Protocol (SIP)", RFC 3323, November 2002.

RFC 3325 C. Jennings, J. Peterson, M. Watson, "Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks", RFC 3325, November 2002.
RFC 3389 R. Zopf, "Real-time Transport Protocol (RTP) Payload for Comfort Noise (CN)", RFC 3389, September 2002.

RFC 3550 H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", RFC 3550, July 2003.

RFC 3711 M Baugher, D McGrew, M. Naslund, E. Carrara, K. Norrman "The Secure Real-Time Transport Protocol (SRTP)" RFC 3711, March 2004.

RFC 4119 Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, DOI 10.17487/RFC4119, December 2005, <http://www.rfc-editor.org/info/rfc4119>.

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

RFC 4568 F.Andreasen, M. Baugher, D.Wing,, "SDP: Security Descriptions For Media Streams", RFC 4568, July 2006.

RFC 4733 H. Schulzrinne, T. Taylor, "RTP Payload for DTMF Digits, Telephony Tones, and Telephony Signals", RFC 4733 (Obsoletes RFC 2833), December 2006.

RFC 4856 S. Casner, "Media Type Registration of Payload Formats in the RTP Profile for Audio and Video Conferences", RFC 4856, March 2007.

RFC 5009 Ejzak, R., "Private Header (P-Header) Extension to the Session Initiation Protocol (SIP) for Authorization of Early Media", RFC 5079, September 2007.

RFC 5031 H. Schulzrinne, "A Uniform Resource Name (URN) for Emergency and Other Well-Known Services", RFC 5031, January 2008.

RFC 5245 Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.

RFC 5280 D. Cooper et. al., "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2009.

RFC 5491 Winterbottom, J., Thomson, M., and H. Tschofenig, "GEOPRIV Presence Information Data Format Location Object (PIDF-LO) Usage Clarification, Considerations, and Recommendations", RFC 5491, DOI 10.17487/RFC5491, March 2009, <http://www.rfc-editor.org/info/rfc5491>.

RFC 5589 R, Sparks, A. Johnston, D. Petrie, "Session Initiation Protocol Call Control – Transfer", RFC 5589, March 2009.

RFC 5621 Camarillo, G., "Message Body Handling in the Session Initiation Protocol (SIP)", RFC 5621, DOI 10.17487/RFC5621, September 2009, <http://www.rfc-editor.org/info/rfc5621>.

RFC 5876 J. Elwell, "Updates to Asserted Identity in the Session Initiation Protocol (SIP)", RFC 5876, April 2010.

RFC 5922 Gurbani, V., Lawrence, S., and B. Laboratories, "Domain Certificates in the Session Initiation Protocol (SIP)", RFC 5922, June 2010.

RFC 6066 Eastlake, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, January 2011.

RFC 6140 A. B. Roach, "Registration for Multiple Phone Numbers in the Session Initiation Protocol (SIP)", RFC 6140, March 2011.

RFC 6442 Polk, J., Rosen, B., and J. Peterson, "Location Conveyance for the Session Initiation Protocol", RFC 6442, DOI 10.17487/RFC6442, December 2011, <http://www.rfc-editor.org/info/rfc6442>.

RFC 6947 Boucadair, M., Kaplan, H., Gilman, R., and S. Veikkolainen, "The Session Description Protocol (SDP) Alternate Connectivity (ALTC) Attribute", RFC 6947, DOI 10.17487/RFC6947, May 2013, <http://www.rfc-editor.org/info/rfc6947>.

RFC 6961 Pettersen, Y., "The Transport Layer Security (TLS) Multiple Certificate Status Request Extension", RFC 6961, June 2013.

RFC 7044 Barnes, M., Audet, F., Schubert, S., van Elburg, J., and C. Holmberg, "An Extension to the Session Initiation Protocol (SIP) for Request History Information", RFC 7044, DOI 10.17487/RFC7044, February 2014, <http://www.rfc-editor.org/info/rfc7044>.

RFC 7092 Kaplan, H. and V. Pascual, "A Taxonomy of Session Initiation Protocol (SIP) Back-to-Back User Agents", RFC 7092, December 2013.

RFC7852 Gellens, R., Rosen, B., Tschofenig, H., Marshall, R., and J. Winterbottom, "Additional Data Related to an Emergency Call", RFC 7852, DOI 10.17487/RFC7852, July 2016, <http://www.rfc-editor.org/info/rfc7852>.

TS 24.610 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Communication HOLD using IP Multimedia Core Network subsystem; Protocol Specification, TS 24.610, V13,1,9, 2015-12.

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

19. Contributors to SIPconnect 2.0 and Contact Information

The SIP Forum would like to thank those individuals and organizations which contributed to the development of SIPconnect 2.0.

19.1 Individual Contributors

Tolga Asveren
Sonus
E-mail: tasveren@sonusnet.com

Mathias Barkow
Vodafone GmbH
Phone: +49 211 533 5441
E-mail: mathias.barkow@vodafone.com

Spencer Dawkins
Huawei Technologies (USA) and Wonder Hamster Internetworking LLC
E-mail: spencerdawkins.ietf@gmail.com

Charles Eckel
Cisco
E-mail: eckelcu@cisco.com

Andrew Hutton
Unify
Technology Drive
Nottingham, UK
E-mail: andrew.hutton@unify.com
Twitter: @huttonandy

Roland Jesske
Deutsche Telekom
Heinrich-Hertz-Strasse 3-7
Darmstadt 64307
Germany
Phone: +4961515812766
E-mail: r.jesske@telekom.de

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

Olle E. Johansson
Edvina AB
Runbovaegen 10
Sollentuna SE-192 48
Sweden
E-mail: oej@edvina.net

Dr. Alan Johnston
E-mail: alan.b.johnston@gmail.com

Laura Liess
Deutsche Telekom
E-mail: l.liess@telekom.de
Phone: +49 6151 5812761

Ingemar Lindblad
Ericsson
E-mail: Ingemar.lindblad@ericsson.com

Neel Bala Neelakantan
Sonus
E-mail: nneelakantan@sonusnet.com

Gonzalo Salgueiro
Cisco
7200-12 Kit Creek Road
RTP, NC 27709
Phone: +1-919-392-3266
E-mail: gsalguei@cisco.com

Doug Sauder
BroadSoft
E-mail: doug@broadsoft.com

Thomas Stach
Independent Consultant, Vienna
E-mail: thomass.stach@gmail.com

James Winterbottom
Winterb Consulting
E-mail: winterb.consulting@outlook.com

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

Rifaat Shekh-Yusef
Avaya
250 Sidney Street
Belleville, Ontario
Canada
Phone: +1-613-967-5267
E-mail: rifaat.ietf@gmail.com

19.2 Organizational Contributors

The following organizations contributed to the development of SIPconnect 2.0:

bitkom
Federal Association for Information Technology, Telecommunications and New Media (Bitkom)
Albrechtstraße 10
10117 Berlin
Germany
www.bitkom.org

NICC Standards Ltd – UK Interoperability Standards
http://www.niccstandards.org.uk/
E-mail: secretary@niccstandards.org.uk

19.3 Acknowledgements to Contributors to Previous Versions

SIPconnect 2.0 is a revision of SIPconnect 1.1 so it's appropriate to recognize the contributors to SIPconnect 1.1 who were Bernard Aboba, François Audet, Shaun Bharrat, Eric Burger, Spencer Dawkins, John Elwell, Alan Johnston, David Hancock, Cullen Jennings, Hadriel Kaplan, Brian Lindsay, Richard Shockey, Mark Stewart, Theo Zourzouvillys.

SIPconnect 1.1 was itself a revision of SIPconnect 1.0, so it's also appropriate to thank the contributors to SIPconnect 1.0 which formed the basis for this work.

Editors: Andrew Hutton (Unify),
Gonzalo Salgueiro (Cisco)

20. Full Copyright Statement