

Standard ECMA-355

3rd Edition / June 2008

**Corporate
Telecommunication
Networks - Tunnelling
of QSIG over SIP**

Standard

Standard
ECMA-355
3rd Edition / June 2008

**Corporate
Telecommunication
Networks - Tunnelling of
QSIG over SIP**

Introduction

This Standard is one of a series of Ecma Standards defining the interworking of services and signalling protocols deployed in corporate telecommunication networks (CNs) (also known as enterprise networks). The series uses telecommunication concepts as developed by ITU-T and conforms to the framework of International Standards on Open Systems Interconnection as defined by ISO/IEC.

This particular Standard specifies tunnelling of QSIG over the Session Initiation Protocol (SIP). This enables calls between "islands" of circuit switched networks that use QSIG signalling to be interconnected by an IP network that uses SIP signalling without loss of QSIG functionality. This standard facilitates the introduction of enhanced SIP and SDP functionality that was specified after publication of the early editions of this standard. These enhancements include payload encryption and mechanisms to negotiate SDP capabilities.

The changes in this Standard comprise a mandatory payload renegotiation with reversed direction of the offer/answer exchange compared with early editions. In order to achieve backward compatibility with early editions an indicator for the changed signalling procedures is introduced. This indicator is used to dynamically detect if fallback to signalling procedures compliant to early editions is necessary.

This Standard is based upon the practical experience of Ecma member companies and the results of their active and continuous participation in the work of ISO/IEC JTC1, ITU-T, IETF, ETSI and other international and national standardization bodies. It represents a pragmatic and widely based consensus.

This Ecma Standard has been adopted by the General Assembly of June 2008.

Table of contents

1	Scope	1
2	Normative references	1
3	Terms and definitions	2
3.1	External definitions	2
3.2	Other definitions	2
3.2.1	Corporate telecommunication Network (CN)	2
3.2.2	Egress gateway	2
3.2.3	Gateway	2
3.2.4	Ingress gateway	2
3.2.5	IP network	2
3.2.6	Media stream	2
3.2.7	Private Integrated Services Network (PISN)	2
3.2.8	Private Integrated services Network eXchange (PINX)	3
4	Abbreviations and acronyms	3
5	Background and architecture	3
5.1	Architecture	3
5.2	Basic operation	4
5.3	QSIG connectionless transport	5
5.4	Late availability of SDP parameters at the egress gateway	5
6	Procedures	5
6.1	General	5
6.2	Encapsulation of QSIG messages in SIP messages	5
6.3	QSIG SETUP message handling at an ingress gateway	6
6.3.1	Sending a SIP INVITE request	6
6.3.2	Receipt of responses to the INVITE request	6
6.4	QSIG SETUP message handling at an egress gateway	7
6.4.1	Receiving a SIP INVITE request	7
6.4.2	Rejecting a QSIG message in an INVITE request	8
6.5	Subsequent QSIG messages	8
6.6	Terminating the SIP dialog	8
6.7	QSIG connectionless message handling at an ingress gateway	9
6.7.1	Sending a SIP INVITE request	9
6.7.2	Receipt of responses to the INVITE request	9

6.8	QSIG connectionless message handling at an egress gateway	10
7	Example message sequences	10
7.1	Call establishment	10
7.2	Call clearing	11
7.3	QSIG connectionless message	12
7.4	Call establishment with port=0 in first SDP answer	13
7.5	Backwards compatibility with early editions	14
8	Security considerations	16
Annex A	Changes from early editions (informative)	17

1 Scope

This Standard specifies tunnelling of "QSIG" over the Session Initiation Protocol (SIP) within a corporate telecommunication network (CN).

"QSIG" is a signalling protocol that operates between Private Integrated services Network eXchanges (PINX) within a Private Integrated Services Network (PISN). A PISN provides circuit-switched basic services and supplementary services to its users. QSIG is specified in Standards, in particular [1] (call control in support of basic services), [2] (generic functional protocol for the support of supplementary services) and a number of Standards specifying individual supplementary services.

NOTE

The name QSIG was derived from the fact that it is used for signalling at the Q reference point. The Q reference point is a point of demarcation between two PINXs [1].

SIP is an application layer protocol for establishing, terminating and modifying multimedia sessions. It is typically carried over IP [4], [6]. Telephone calls are considered as a type of multimedia session where just audio is exchanged. SIP is defined in [9].

Often a CN comprises both PISNs employing QSIG and IP networks employing SIP. A call or call independent signalling can originate at a user connected to a PISN and terminate at a user connected to an IP network or vice versa. In either case, a gateway provides interworking between QSIG and SIP at the boundary between the PISN and the IP network. Basic call interworking at a gateway is specified in [3]. Another case is where a call or call independent signalling originates at a user connected to a PISN, traverses an IP network using SIP, and terminates at a user connected to another (or another part of the same) PISN. This document addresses this last case in a way that preserves all QSIG capabilities across the IP network. It achieves this by tunnelling QSIG messages within SIP requests and responses in the context of a SIP dialog.

The tunnelling of QSIG through a public IP network employing SIP is outside the scope of this specification. However, the functionality specified in this specification is in principle applicable to such a scenario when deployed in conjunction with other relevant functionality (e.g., address translation, security functions, etc.).

This specification is applicable to any interworking unit that can act as a gateway between a PISN employing QSIG and a corporate IP network employing SIP, with QSIG tunnelled within SIP requests and responses.

2 Normative references

- [1] International Standard ISO/IEC 11572 "Information technology -- Telecommunications and information exchange between systems -- Private Integrated Services Network -- Circuit mode bearer services -- Inter-exchange signalling procedures and protocol" (also published by Ecma as Standard ECMA-143).
- [2] International Standard ISO/IEC 11582 "Information technology -- Telecommunications and information exchange between systems -- Private Integrated Services Network -- Generic functional protocol for the support of supplementary services -- Inter-exchange signalling procedures and protocol " (also published by Ecma as Standard ECMA-165).
- [3] International Standard ISO/IEC 17343 "Information technology -- Telecommunications and information exchange between systems -- Corporate telecommunication networks -- Signalling interworking between QSIG and SIP -- Basic services" (also published by Ecma as Standard ECMA-339).
- [4] J. Postel, "Internet Protocol", RFC 791.
- [5] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119.
- [6] S. Deering, R. Hinden, "Internet Protocol, Version 6 (IPv6)", RFC 2460.
- [7] Donovan, S., "The SIP INFO Method", RFC 2976.

- [8] Zimmerman, E., Peterson, J., Vemuri, A., Ong, L., Audet, F., Watson, M. and M. Zonoun, "MIME media types for ISUP and QSIG objects", RFC 3204.
- [9] J. Rosenberg, H. Schulzrinne, et al., "SIP: Session initiation protocol", RFC 3261.
- [10] J. Rosenberg, H. Schulzrinne, et al., "An Offer/Answer Model with the Session Description Protocol (SDP)", RFC 3264.
- [11] Rosenberg, J., "The Session Initiation Protocol (SIP) UPDATE message", RFC 3311.
- [12] J. Rosenberg, H. Schulzrinne, P. Kyzivat, "Indicating User Agent Capabilities in the Session Initiation Protocol (SIP)", RFC 3840.

3 Terms and definitions

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [5] and indicate requirement levels for compliant SIP implementations.

For the purposes of this document, the following definitions apply.

3.1 External definitions

The definitions in [1] and [9] apply as appropriate.

3.2 Other definitions

3.2.1 Corporate telecommunication Network (CN)

Sets of privately-owned or carrier-provided equipment that are located at geographically dispersed locations and are interconnected to provide telecommunication services to a defined group of users.

NOTE

A CN can comprise a PISN, a private IP network (intranet) or a combination of the two.

3.2.2 Egress gateway

A gateway handling a QSIG call or call-independent signalling connection established in the direction IP network to PISN.

3.2.3 Gateway

An entity that behaves as a QSIG Transit PINX with QSIG carried over a circuit-switched link within a PISN on one side and QSIG tunnelled over SIP within an IP network on the other side.

3.2.4 Ingress gateway

A gateway handling a QSIG call or call-independent signalling connection established in the direction PISN to IP network.

3.2.5 IP network

A network, unless otherwise stated a corporate network, offering connectionless packet-mode services based on the Internet Protocol (IP) as the network layer protocol.

3.2.6 Media stream

Audio or other user information transmitted in UDP packets, typically containing RTP, in a single direction between the gateway and a peer entity participating in a session established using SIP.

NOTE

Normally a SIP session establishes a pair of media streams, one in each direction.

3.2.7 Private Integrated Services Network (PISN)

A CN or part of a CN that employs circuit-switched technology and QSIG signalling.

3.2.8 Private Integrated services Network eXchange (PINX)

A PISN nodal entity comprising switching and call handling functions and supporting QSIG signalling in accordance with [1].

4 Abbreviations and acronyms

CN	corporate telecommunication network
IP	Internet Protocol
PINX	Private Integrated services Network eXchange
PISN	Private Integrated Services Network
QSIG	Signalling system for the Q reference point
RTP	Real-time Transport Protocol
SDP	Session Description Protocol
SIP	Session Initiation Protocol
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UA	User Agent
UAC	User Agent Client
UAS	User Agent Server
UDP	User Datagram Protocol
URI	Universal Resource Identifier

5 Background and architecture

5.1 Architecture

This document concerns the case of a call or call independent signalling that originates at a user connected to a PISN employing QSIG, traverses an IP network employing SIP, and terminates at a user connected to another (or another part of the same) PISN. This can be achieved by employing a gateway at each boundary between a PISN employing QSIG and an IP network employing SIP, as shown in Figure 1.




Figure 1 — Call from QSIG via SIP to QSIG

Each gateway can provide interworking as specified in [3]. This provides a basic call capability. However, [3] only specifies interworking for QSIG basic call, as specified in [1]. Many of the other capabilities of QSIG (support for supplementary services and additional network features) as specified in other standards and in vendor-specific specifications are not covered. Some of these additional capabilities of QSIG are suitable for interworking with SIP and might be the subject of future Standards or other specifications. Other capabilities of QSIG are unsuitable for interworking with SIP because corresponding capabilities do not exist in SIP or are achieved in ways that are

incompatible with QSIG. Therefore interworking at a gateway between QSIG and SIP will be limited to those QSIG capabilities that have sufficiently compatible equivalents in SIP. Each capability requires special implementation in the gateway, and therefore a typical gateway might provide interworking for only a subset of capabilities for which interworking is feasible. The result of this is that there will be a loss of capability on a call or call independent signalling from QSIG to SIP or vice versa. For a case similar to that shown in Figure 1 there will likewise be a loss of capability. This can be compounded if the two gateways are of different types, since only those capabilities common to both gateways will survive end-to-end. The solution is to tunnel QSIG messages through the IP network within SIP messages so that no end-to-end QSIG capabilities are lost. One of the two gateways originates a SIP dialog to the other gateway. SIP messages within the dialog are used to tunnel QSIG messages. Through the use of SDP [10], the dialog also establishes a session in which media streams carry user information (e.g., speech) between the two QSIG gateways, if required. The two gateways act as QSIG Transit PINXs, which relay QSIG messages with little or no modification.

In a conventional PISN employing QSIG, two PINXs are connected by means of an inter-PINX link, which comprises a signalling channel (carrying QSIG messages) and one or more user information channels carrying speech, modem information or data. With the tunnelling solution, the IP network provides the inter-PINX link between the two gateways acting as Transit PINXs. The tunnel provided by SIP for QSIG messages acts as the signalling channel and the media streams act as the user information channels.

This document covers the case where a single dialog between two gateways is used for a single QSIG call or call independent signalling connection, as specified in [2]. This means that the dialog is established when the QSIG call or call independent signalling connection is established and cleared down when the QSIG call or call independent signalling connection is cleared down.

An enhanced scenario in which a single SIP dialog is maintained long term and used to tunnel a multiplicity of QSIG calls or call independent signalling connections, with the possibility of multiple QSIG calls or call independent signalling connections being in progress at any one time, is outside the scope of this document.

This document also covers call-independent connectionless transport.

5.2 Basic operation

When a gateway (the ingress gateway) receives a QSIG call or call-independent signalling connection establishment request (QSIG SETUP message) from the PISN, it needs to generate a SIP INVITE request using a Request-URI that will route the request to an appropriate egress gateway. The Request-URI must be derived in some way from the required destination of the QSIG call or call-independent signalling connection (as indicated in the Called party number information element of the QSIG SETUP message). The Request-URI can explicitly identify the egress gateway or it can simply identify the required destination. The first case is likely to require some sort of look-up capability in the ingress gateway, the configuration of which is outside the scope of this document. For the latter case algorithmic mapping of the called party number to a Request-URI might be sufficient, but this delegates the task of selecting an appropriate egress gateway to SIP proxies.

An ingress gateway may determine from the required destination of a call or call-independent signalling connection that the destination is not reachable via QSIG tunnelling. In this case the QSIG gateway can either route the call or call-independent signalling connection onwards within the PISN or can route the call or call-independent signalling connection into the IP network using interworking as specified in [3]. How an ingress gateway determines that the destination is not reachable via QSIG tunnelling is outside the scope of this document.

If an ingress gateway maps the QSIG called party number to a Request URI that does not explicitly identify a particular egress gateway, routing of the INVITE request is left to SIP proxies. A proxy might route the request to a UAS that is not an egress gateway to QSIG, in which case QSIG tunnelling will not be possible. Allowing the call or call-independent signalling connection to proceed in this situation is likely to be undesirable, since the ingress gateway expects to carry out QSIG tunnelling whereas interworking with SIP, as specified in [3], would be more appropriate. To

cater for this situation, a mechanism is defined that causes an INVITE request containing tunnelled QSIG to be rejected by an egress gateway that does not support this capability.

NOTE

Allowing the INVITE request to be routed by proxies either to an egress gateway to QSIG or to some other UAS without the ability for the ingress gateway to choose in advance is undesirable. It implies that the ingress gateway maps the QSIG SETUP message to a SIP INVITE request in accordance with both this document and [3] simultaneously. Although this may seem feasible superficially, architecturally it is dangerous because with QSIG tunnelling the ingress gateway should act as a QSIG Transit PINX whereas with interworking in accordance with [3] it should act as a QSIG Outgoing Gateway PINX. The ingress gateway will not know for certain which behaviour to adopt until a 200 OK arrives, and therefore in the meantime it will not know how to handle information relating to certain QSIG capabilities (supplementary services and additional network features) in the QSIG SETUP message. It is not clear whether this can be handled safely for all possible QSIG capabilities (including vendor-specific capabilities). For this reason, this specification and [3] require the ingress gateway to make a decision between tunnelling and interworking respectively.

5.3 QSIG connectionless transport

When a gateway (the ingress gateway) receives a QSIG connectionless FACILITY message from the PISN, it needs to generate a SIP INVITE request using a Request-URI that will route the request to an appropriate egress gateway. The Request-URI must be derived in some way from the required destination of the QSIG FACILITY message (as indicated in the Called party number information element of the QSIG SETUP message). Techniques similar to those used for QSIG SETUP messages can be used.

An ingress gateway may determine from the required destination of a QSIG connectionless FACILITY message that the destination is not reachable via QSIG tunnelling. In this case the QSIG gateway can either route the message onwards within the PISN or discard the message.

5.4 Late availability of SDP parameters at the egress gateway

In conjunction with certain gateway architectures (e.g. decomposed gateways comprising a separate signalling gateway and media gateway), valid SDP parameters might not be available to the egress gateway (e.g., the signalling gateway part of a decomposed gateway) at the time when an SDP answer to the initial SDP offer needs to be sent. In this case the egress gateway will have to send a dummy SDP answer. The final SDP parameters need to be agreed upon during a later phase of the QSIG call establishment.

In order to achieve identical signalling procedures between ingress gateway and egress gateway irrespective of the gateway architecture, the ingress gateway starts SDP re-negotiation immediately after confirmation of the SIP dialog provided that both gateways support the changed procedures. This mandatory re-negotiation during QSIG call establishment is redundant in the case of an integrated gateway that can provide valid parameters for an SDP answer at the time of the first SDP negotiation, but is still employed in the interests of a single solution that works for all gateway architectures. In the case of decomposed gateways, the solution allows the SDP offer-answer cycle to be conducted end-to-end between the ingress and egress media gateways, with only minimal intervention by the intermediate signalling gateways.

6 Procedures

6.1 General

A gateway SHALL behave as a QSIG Transit PINX as specified in [1] and modified as specified below.

6.2 Encapsulation of QSIG messages in SIP messages

When encapsulating a QSIG message inside a SIP message, a gateway SHALL include the QSIG message in a MIME body of the SIP request or response in accordance with [8] using media type application/QSIG. QSIG segmentation SHALL NOT apply.

If any other MIME body is to be included (e.g., SDP), the gateway SHALL use multi-part MIME.

The gateway SHALL include a Content-Disposition header field indicating "signal" and "handling=required" as a SIP header field (in the case of single-part MIME) or as a MIME header field in the body containing the QSIG message (in the case of multi-part MIME).

6.3 QSIG SETUP message handling at an ingress gateway

6.3.1 Sending a SIP INVITE request

The ingress gateway, on receipt of a QSIG SETUP message eligible for tunnelling over SIP to an egress gateway, SHALL build a SIP INVITE request message containing a Request-URI suitable for routing towards a suitable egress gateway.

NOTE

The Request-URI should be derived in some way from the Called party number information in the QSIG SETUP message. The Request-URI can explicitly identify a particular egress gateway. Alternatively it can identify the final destination in a way that leaves selection of a suitable egress gateway to SIP proxies.

The From: header field SHOULD contain a URI identifying either the ingress gateway or the calling party (derived from the QSIG Calling party number information element).

The Contact: header field SHALL include a user agent capability in accordance with [12] in which URI parameter +u.ecma-international.org/ecma355/new_sdp_by_ingress is present, indicating support for the procedures specified by this Standard. This URI parameter provides a means to achieve backward compatibility to early editions of this standard.

The ingress gateway SHALL encapsulate the QSIG SETUP message in the SIP INVITE request.

The encapsulated QSIG SETUP message MAY differ from the received SETUP message in accordance with acceptable modification at a QSIG Transit PINX. For example, the Channel identification information element MAY change. Because in the encapsulated QSIG SETUP message the contents of the Channel identification information element have no significance, the channel number field SHOULD contain value 1 and the preferred/exclusive field SHOULD contain value 1 (exclusive).

For call establishment the INVITE request SHALL contain an SDP offer proposing a pair of media streams, one in each direction (a=sendrecv), that the gateway can map to the user information channel indicated in the Channel identification information element in the received QSIG message. The media streams SHALL be suitable for use in accordance with the Bearer capability information element in the received QSIG SETUP message. For call-independent signalling connection establishment the INVITE request SHALL contain an SDP offer [10] containing zero "m=" lines.

After sending the SIP INVITE request, the ingress gateway SHALL NOT encapsulate any further message from QSIG until a SIP 200 OK response has been received.

6.3.2 Receipt of responses to the INVITE request

The action specified below is in addition to normal UA handling of a SIP response.

On receipt of a SIP 4xx, 5xx or 6xx final response, the ingress gateway SHALL either take alternative action to route the call or call-independent signalling connection (outside the scope of this document) or clear the call or call-independent signalling connection using an appropriate cause value in the QSIG Cause information element of the QSIG clearing message concerned (DISCONNECT, RELEASE or RELEASE COMPLETE). If the SIP response contains an encapsulated QSIG RELEASE COMPLETE message, the ingress gateway SHOULD use the cause value in that message to determine the cause value when clearing. The ingress gateway SHOULD choose a cause that reflects the fact that the next PINX cannot be reached (e.g., Cause value 3 "no route to destination").

NOTE

A SIP 415 Unsupported Media Type final response can be expected if the UAS does not support encapsulated QSIG.

On receipt of a SIP 200 OK response, the ingress gateway SHALL carry out normal SIP processing, including transmission of an ACK request, and SHALL act upon any encapsulated QSIG message. The ingress gateway SHALL check for presence of a URI parameter +u.ecma-

international.org/ecma355/new_sdp_by_ingress in the Contact: header field of the SIP 200 OK response.

If present, the ingress gateway SHALL immediately issue a SIP re-INVITE request re-using the SDP data from the initial SIP INVITE request as specified in 6.3.1. Afterwards the ingress gateway SHALL act in accordance with [9]. The ingress gateway SHALL accept incoming SIP INFO requests that tunnel QSIG information independently of whether a response to the pending SIP re-INVITE request has been received.

If not present, the ingress gateway SHALL NOT immediately issue the SIP re-INVITE request. If the received SDP answer indicates port=0, the ingress gateway SHALL wait for receipt of a SIP re-INVITE request including a new SDP offer from the egress gateway and act in accordance with [9] afterwards. The ingress gateway SHALL accept incoming SIP INFO requests including a QSIG message.

The ingress gateway SHALL also connect the QSIG user information channel to the media streams indicated in the SDP answer if able to do so.

NOTE

If an "m=" line indicates port 0 the ingress gateway will be unable to start the media stream in the direction ingress to egress at this stage. 2-way connection will be achieved after a successful re-INVITE or UPDATE transaction.

If the ingress gateway has not received a non-zero port for a media stream by the time the QSIG CONNECT message has been received, the ingress gateway SHALL not act upon the QSIG CONNECT message and instead SHALL act as if QSIG timer T301 has expired.

6.4 QSIG SETUP message handling at an egress gateway

6.4.1 Receiving a SIP INVITE request

On receipt of a SIP INVITE request containing a QSIG message in the body of the request, the egress gateway SHALL check for presence of the URI parameter +u.ecma-international.org/ecma355/new_sdp_by_ingress in the Contact: header field and send a SIP 200 OK response. If present in the SIP INVITE request, the SIP 200 OK response SHALL contain the URI parameter +u.ecma-international.org/ecma355/new_sdp_by_ingress in the Contact: header field together with an SDP answer. The egress gateway SHALL omit the URI parameter from the SIP 200 OK response if it was absent in the SIP INVITE request. The SDP answer SHOULD establish symmetrical media streams, unless the SDP offer contained zero "m=" lines (for a call-independent signalling connection).

If the egress gateway cannot determine appropriate SDP parameters at the time the SDP answer is to be sent, the egress gateway SHALL send an SDP answer in which the "m=" lines indicate port 0. In this case the IP address in the "c=" line is unimportant but may, for example, be that of the SIP signalling part of the egress gateway.

If the QSIG message is a SETUP message acceptable for routing onwards into the PISN, the egress gateway SHALL select a user information channel on the PISN side and shall forward the QSIG SETUP message. The forwarded QSIG SETUP message MAY differ from the received SETUP message in accordance with acceptable modification at a QSIG Transit PINX. In particular, the Channel identification information element SHALL reflect the selected user information channel. For call establishment the egress gateway SHALL also connect the QSIG user information channel to the established media streams.

The egress gateway MAY include in the SIP 200 OK response to the initial SIP INVITE request an encapsulated QSIG SETUP ACKNOWLEDGE message or CALL PROCEEDING message. Otherwise the gateway SHALL transmit this first responding QSIG message later in a SIP INFO message in accordance with 6.5.

Further procedures depend on the presence or absence of the URI parameter +u.ecma-international.org/ecma355/new_sdp_by_ingress in the Contact: header field of the initial SIP INVITE request.

If the URI parameter +u.ecma-international.org/ecma355/new_sdp_by_ingress was present, the egress gateway SHALL wait for receipt of a SIP re-INVITE request from the ingress gateway

independently of whether valid SDP parameters could be determined previously. The egress gateway SHALL acknowledge the receipt of the SIP re-INVITE request by sending a 100 Trying provisional response to the ingress gateway in case that valid SDP parameter are not yet available. When valid SDP parameters are already available or become available at a later stage of QSIG call establishment, the egress gateway SHALL use the SDP parameters as SDP answer in a 200 OK response to the received SIP re-INVITE request.

If the URI parameter `+u.ecma-international.org/ecma355/new_sdp_by_ingress` was present, the egress gateway MAY include in the SIP 200 response to the SIP re-INVITE request an encapsulated QSIG message.

If the URI parameter `+u.ecma-international.org/ecma355/new_sdp_by_ingress` was absent, the egress gateway will not receive a SIP re-INVITE request. When the SDP parameters are available at a later stage of QSIG call establishment, the egress gateway SHALL re-negotiate media streams by sending a SIP re-INVITE request or SIP UPDATE request with an SDP offer reflecting these parameters.

NOTE 1

This fall-back to the procedures of early editions may cause problems with specified behaviour of certain SDP extensions such as key management and negotiation of SDP capabilities.

NOTE 2

The egress gateway may reject a SIP INVITE request in accordance with [9]. For example, if the SIP UAS does not support encapsulated QSIG and therefore is not capable of being an egress gateway, SIP response code 415 Unsupported Media Type will apply. If the SIP UAS is unable to accept the SDP offer, SIP response code 488 Not Acceptable Here will apply. However, an egress gateway that is unable to validate the SDP offer in the initial SIP INVITE request prior to sending a SIP 200 response will accept the offer and use the SDP offer received in the SIP re-INVITE request for the media stream negotiation. An egress gateway compliant to an early edition of this standard would send a SIP re-INVITE request or SIP UPDATE request for media stream re-negotiation.

6.4.2 Rejecting a QSIG message in an INVITE request

If the egress gateway receives an INVITE request containing a QSIG message that is not acceptable (e.g., a SETUP message not suitable for routing onwards, an inappropriate QSIG message, a SETUP message for a call for which suitable media streams cannot not be established) the egress gateway SHALL send back a responding QSIG message in accordance with [1] or [2], e.g., a RELEASE COMPLETE message containing an appropriate value in the QSIG Cause information element. The responding QSIG message SHALL be sent either in an INFO message in accordance with 6.5 or, in the case of a RELEASE COMPLETE message, in a BYE message in accordance with 6.6.

If there is a pending SIP re-INVITE request at the time of call rejection, the egress gateway SHALL terminate the pending transaction by a 487 Request Terminated response.

6.5 Subsequent QSIG messages

After transmitting a 200 OK response to the initial SIP INVITE request (egress gateway) or transmitting an ACK following receipt of a 200 OK response (ingress gateway), a gateway SHALL encapsulate any further QSIG messages for transmission to the peer gateway in the body of a SIP INFO request [7] and SHALL be able to receive further QSIG messages from the peer gateway encapsulated in the body of SIP INFO request. The exceptions are:

- a SIP BYE request, which MAY encapsulate a QSIG RELEASE COMPLETE message, in accordance with 6.6; and
- a SIP re-INVITE request or SIP UPDATE request [11], which MAY encapsulate a waiting QSIG message during QSIG call establishment or during payload renegotiation for an active call.

6.6 Terminating the SIP dialog

When a gateway determines that a QSIG call or call-independent signalling connection has terminated, it SHALL terminate the SIP session by transmitting a BYE request. If a gateway transmits the final QSIG message of the call or call-independent signalling connection (RELEASE COMPLETE), the gateway MAY encapsulate that QSIG message in the BYE request. Otherwise

the gateway SHALL transmit the BYE request after the final QSIG message has been sent or received and SHALL NOT encapsulate a QSIG message in the BYE request.

If there is a pending SIP re-INVITE request at the time of dialog termination, the egress gateway SHALL terminate the pending transaction by a 487 Request Terminated response.

6.7 QSIG connectionless message handling at an ingress gateway

The QSIG FACILITY message (including addressing information elements) is used to convey call independent connectionless signalling information, as specified in [2].

For tunnelling of a QSIG FACILITY message a SIP INVITE request is used, but the session does not need to be completely established. The SIP INVITE request is rejected after the receipt of the QSIG FACILITY message.

NOTE

By rejecting the SIP INVITE request the overhead of having to release the established dialog is avoided.

6.7.1 Sending a SIP INVITE request

An ingress gateway, on receipt of a QSIG connectionless FACILITY message eligible for tunnelling over SIP to an egress gateway, SHALL build a SIP INVITE request message containing a Request-URI suitable for routing towards a suitable egress gateway.

NOTE

The Request-URI should be derived in some way from the Called party number information element in the QSIG FACILITY message. The Request-URI can explicitly identify a particular egress gateway. Alternatively it can identify the final destination in a way that leaves selection of a suitable egress gateway to SIP proxies.

The From: header field SHOULD contain a URI identifying either the ingress gateway or the calling party (derived from the QSIG Calling party number information element).

The ingress gateway SHALL encapsulate the QSIG FACILITY message in the SIP INVITE request. The encapsulated QSIG FACILITY message MAY differ from the received FACILITY message in accordance with acceptable modification at a QSIG Transit PINX.

The INVITE request SHALL contain an SDP offer [10] containing zero "m=" lines.

After transmission of an INVITE request containing a QSIG connectionless FACILITY message, the ingress gateway SHALL NOT use any dialog created (e.g., on receipt of a 1xx response) for transmitting further encapsulated QSIG messages.

The ingress gateway MAY generate further INVITE requests containing encapsulated QSIG FACILITY messages to the same egress gateway and/or to other egress gateways without waiting for a response to previous INVITE requests.

6.7.2 Receipt of responses to the INVITE request

The action specified below is in addition to normal UA handling of a SIP response.

On receipt of a SIP 4xx, 5xx or 6xx final response, the ingress gateway SHOULD consider the handling of the QSIG connectionless FACILITY message complete.

NOTE

A 603 response normally indicates successful delivery as far as the egress gateway (but not necessarily to the final destination), whereas any other final response will normally indicate failure to deliver. There is no obligation on the ingress gateway to try to recover from failure.

On receipt of a SIP 2xx final response, the ingress gateway SHOULD consider the handling of the QSIG connectionless FACILITY message complete and terminate the dialog by transmitting a BYE request.

NOTE

A 2xx final response should not normally be received.

6.8 QSIG connectionless message handling at an egress gateway

On receipt of a SIP INVITE request containing a QSIG FACILITY message in the body of the request, the egress gateway SHALL treat the QSIG message as a connectionless message and send a SIP 603 Decline response.

NOTE

The reason for the negative response is because there is no value in establishing a dialog, since delivery of the QSIG connectionless message is complete.

If the QSIG FACILITY message is suitable for routing onwards into the PISN, the egress gateway SHALL forward it accordingly. Otherwise the egress gateway SHALL discard the message. The forwarded QSIG FACILITY message MAY differ from the received FACILITY message in accordance with acceptable modification at a QSIG Transit PINX.

NOTE

The egress gateway may reject a SIP INVITE request in accordance with [9]. For example, if the SIP UAS does not support encapsulated QSIG and therefore is not capable of being an egress gateway, SIP response code 415 Unsupported Media Type will apply.

7 Example message sequences

7.1 Call establishment

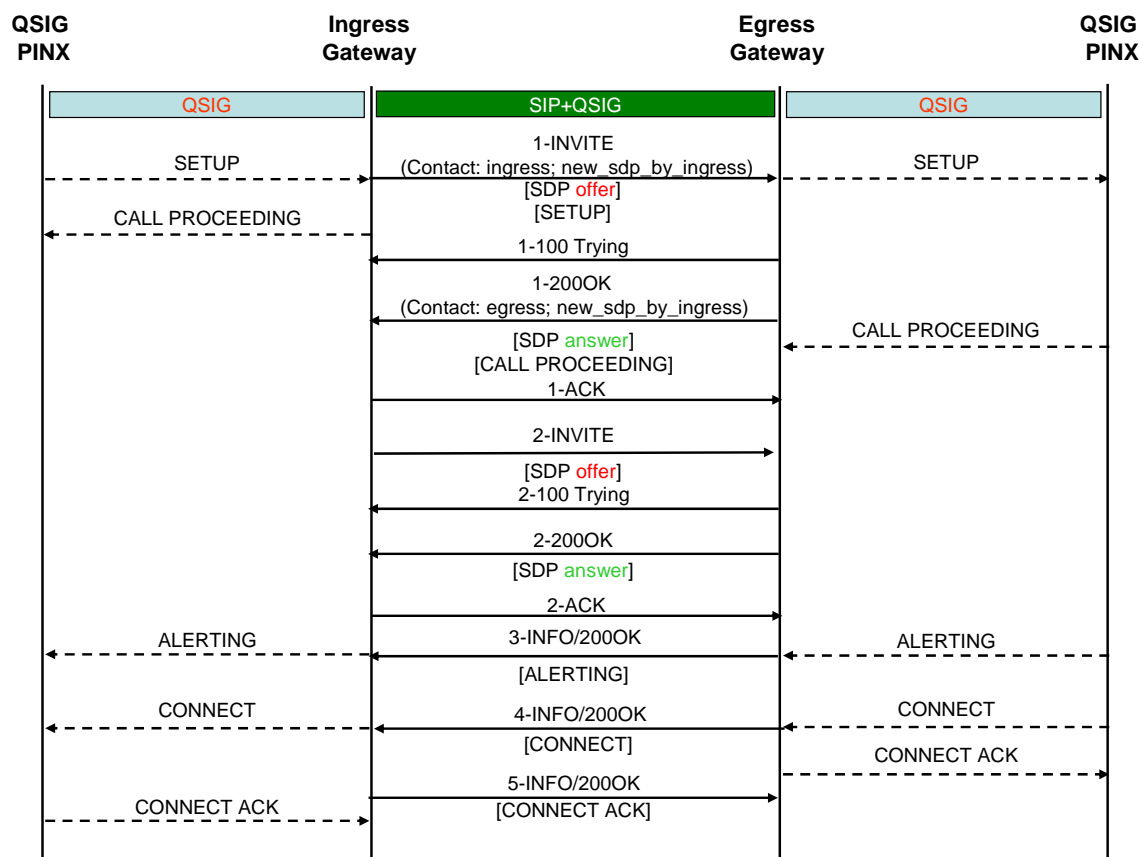


Figure 2 — Call establishment QSIG-SIP-QSIG

7.2 Call clearing

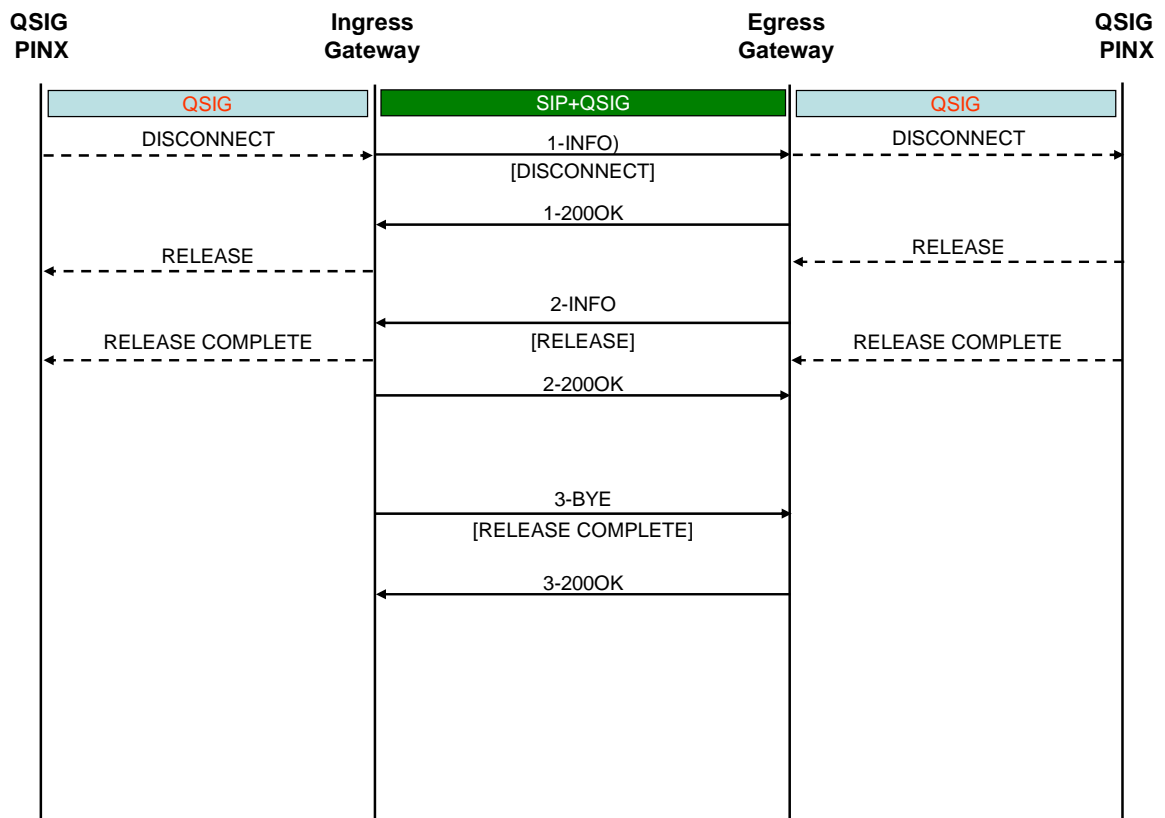


Figure 3 — Call clearing QSIG-SIP-QSIG

7.3 QSIG connectionless message

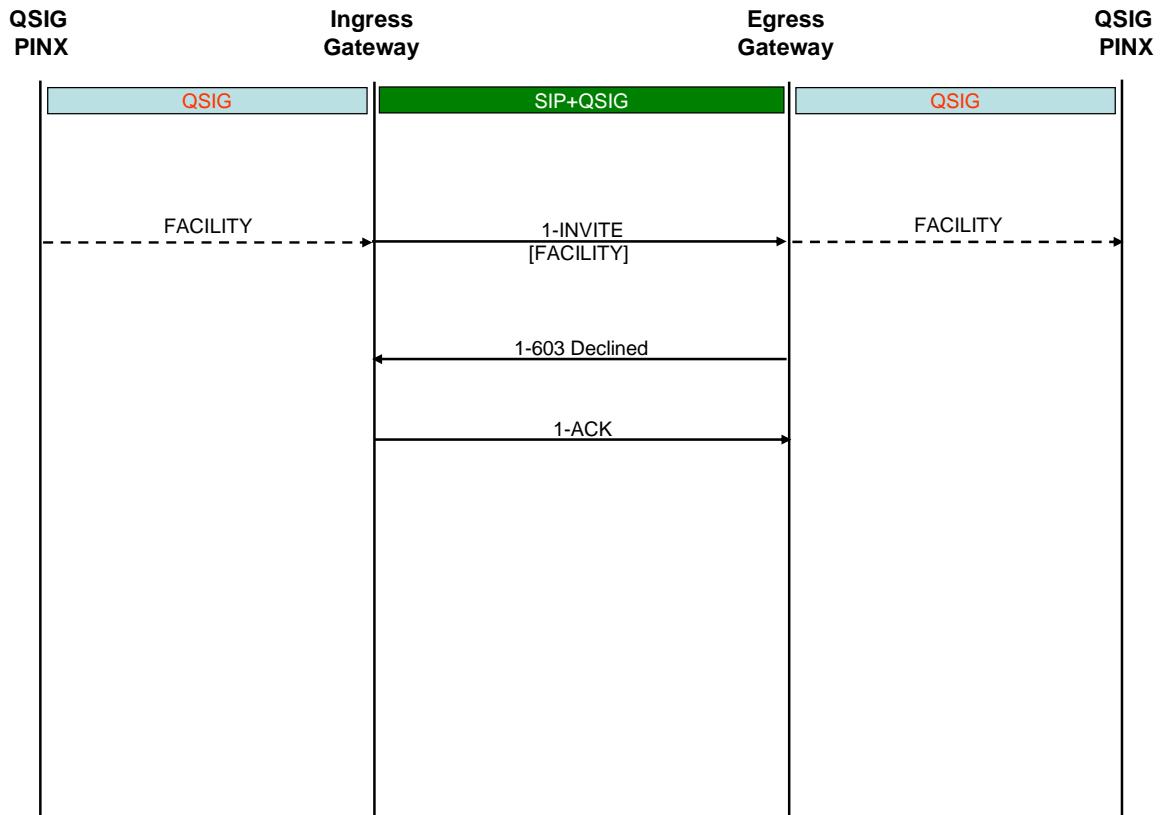


Figure 4 — QSIG call independent connectionless signalling

7.4 Call establishment with port=0 in first SDP answer

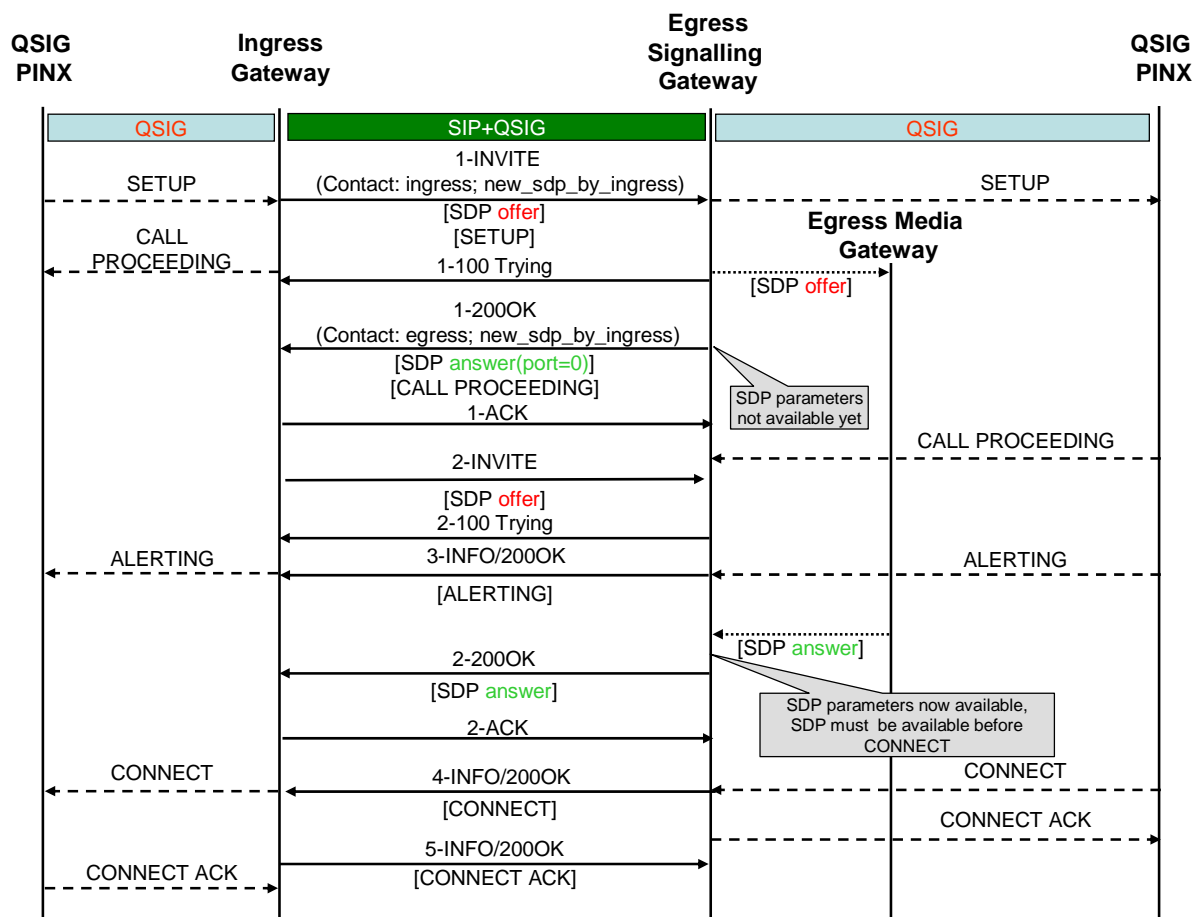


Figure 5 — Call establishment QSIG-SIP-QSIG with port=0 in SDP answer in SIP 200 response to INVITE request

7.5 Backwards compatibility with early editions

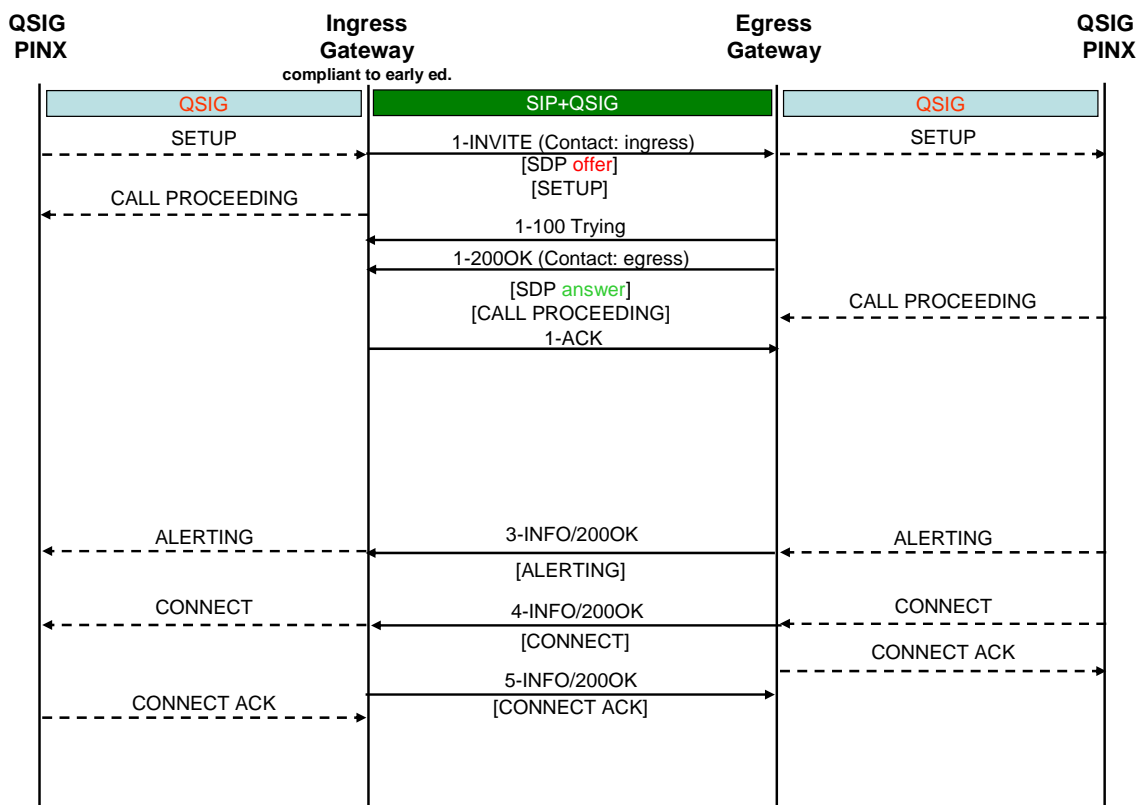


Figure 6 — Call establishment QSIG-SIP-QSIG with ingress gateway compliant to an early edition

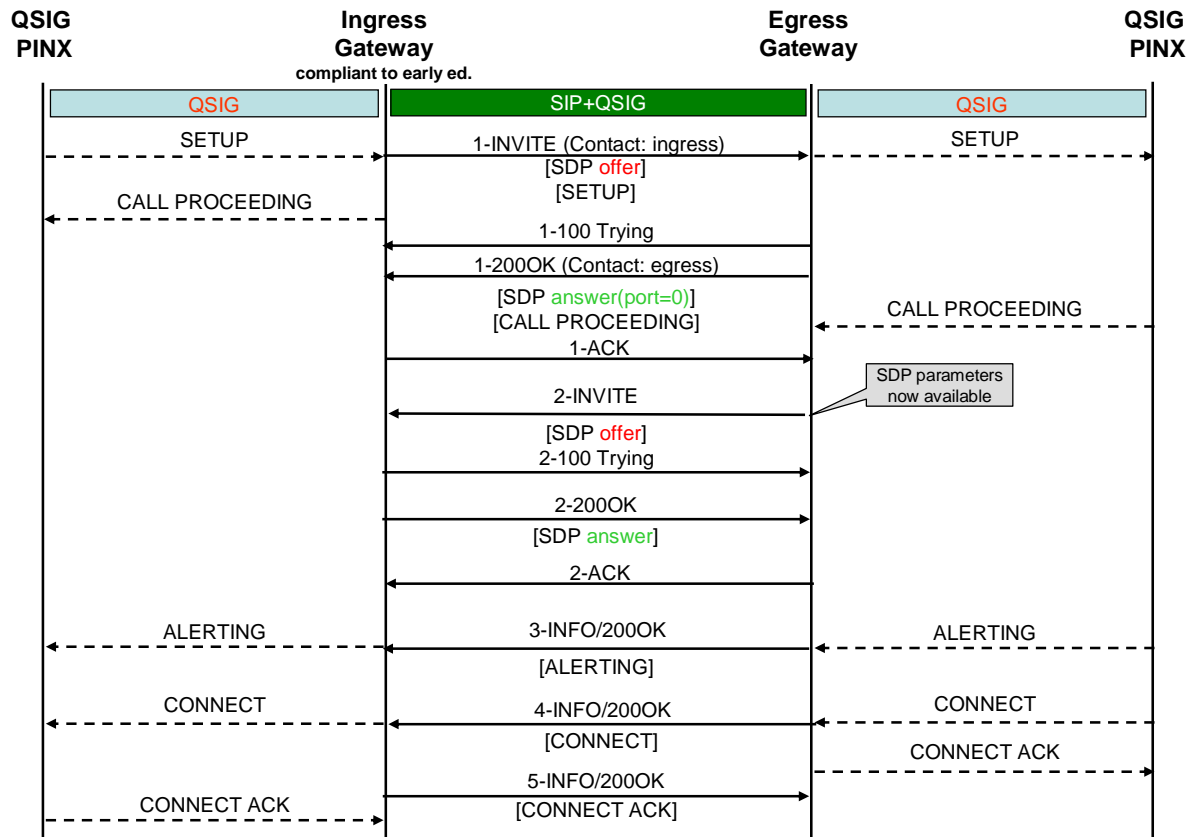


Figure 7 — Call establishment QSIG-SIP-QSIG with ingress gateway compliant to an early edition and port=0 in SDP answer

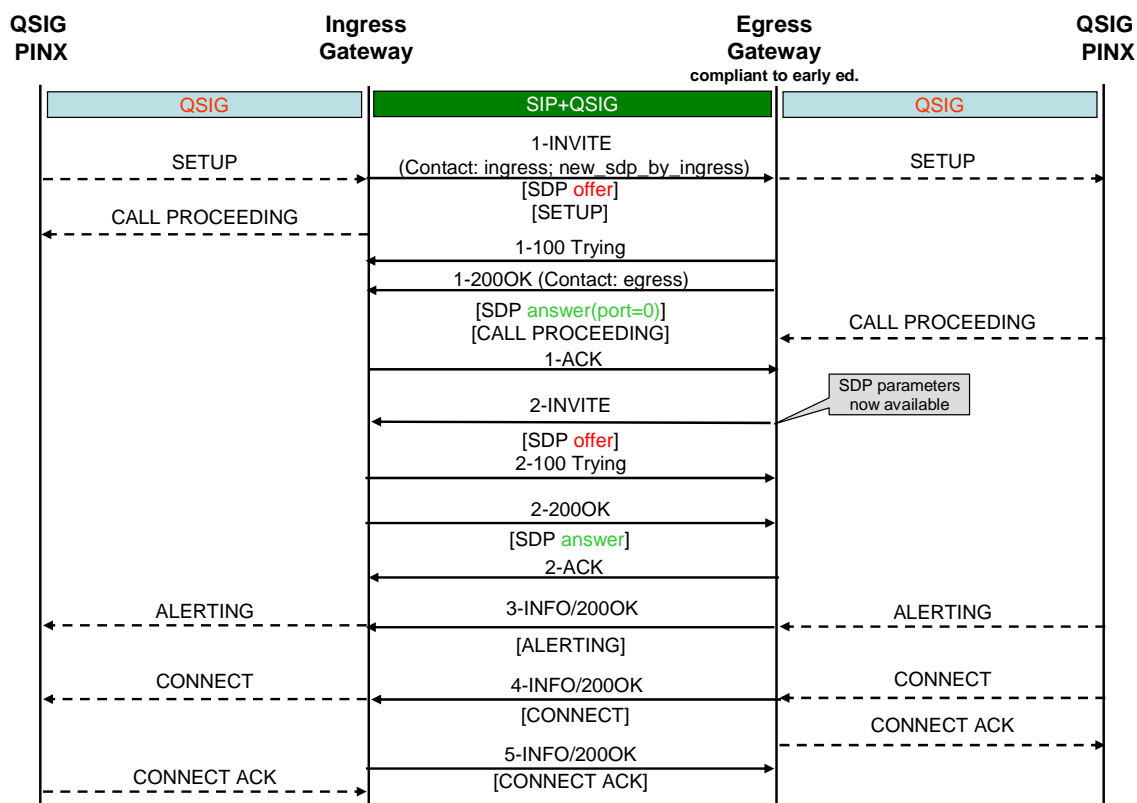


Figure 8 — Call establishment QSIG-SIP-QSIG with egress gateway compliant to an early edition and port=0 in SDP answer

8 Security considerations

QSIG can contain potentially sensitive information, i.e., numbers and names of call participants. Therefore a gateway needs to take care that any QSIG information it transmits is sent only to trusted QSIG gateways and cannot be accessed by third parties. Furthermore a gateway needs to be sure of the source and integrity of any QSIG information it receives, to avoid harming or sending misleading information to the QSIG network.

A gateway SHALL support the transport of SIP over Transport Layer Security (TLS) and the use of the SIPS URI as defined in [9] for identifying the gateway and for establishing a call or call-independent signalling connection to a peer gateway. In addition a gateway SHALL either be able to act as a TLS server, which requires it to have its own private key and certificate, or support the retention of TLS connections for use by incoming SIP calls or call-independent signalling connections.

NOTE

Support of TLS and SIPS meet the security requirements to the extent that each link (between gateway and proxy and between proxies) is secured. This is sufficient in a typical enterprise environment, where proxies can be trusted.

In addition, a gateway MAY support the use of S/MIME for securing a QSIG body in accordance with [9].

NOTE

This avoids the need for trusted proxies, but requires each gateway to have a private key and certificate.

Annex A (informative)

Changes from early editions

A.1 Motivation for changes in payload renegotiation procedures

Early editions of this standard introduced a procedure for payload renegotiation that can be applied in case of a decomposed egress gateway. This procedure deals with a situation where the egress signalling gateway needs to send an SDP answer when appropriate SDP parameters are not yet available from the media gateway. In order to comply with SIP procedures, the SIP 200 OK response to the initial INVITE request needs to contain an SDP answer, and since that response needs to be sent in order to convey the QSIG CALL PROCEEDING message in a timely manner, SDP parameters may not be available from the media gateway. Waiting for the correct SDP parameters to be available may lead to expiry of QSIG timer T303.

Using a straight-forward implementation concept, the egress signalling gateway would send a dummy SDP answer in the 200 OK response to the initial INVITE request. It also would send SDP parameters to the media gateway as received in the SDP offer from the ingress gateway and wait for an answer with SDP parameters from the media gateway. Once received the egress gateway would issue a new SDP offer towards the ingress gateway using the newly-acquired parameters, which were intended as an answer to the original offer. Based on the assumption that SDP offers and SDP answers can be used interchangeably such a practice would work adequately.

However, since the publication of the early editions of this standard various extensions to SDP were specified. Among these SDP extensions are key-management as needed for payload encryption or mechanisms for negotiation of SDP capabilities that must be exchanged via end-to-end offers and answers. Both mentioned extensions introduce asymmetries between SDP offers and SDP answers, such that SDP offers and SDP answers cannot be used interchangeably any longer. Thus, based on a straight-forward implementation, the payload renegotiation procedures between egress gateway and ingress gateway break when such asymmetries are present. It is important that this standard be tolerant of all present and future enhancements to SDP, including, in particular, various key management methods for media security and SDP capability negotiation. This means that it must be possible to conduct the SDP offer-answer cycle end-to-end between the ingress and egress media gateways and/or devices e.g. behind an IP-IP gateway, with only minimal intervention by the intermediate signalling gateways.

NOTE

Such an IP-IP gateway is meant to be an entity that behaves as a QSIG Transit PINX with QSIG tunnelled over SIP within an IP network on one side and QSIG tunnelled over another protocol within an IP network on the other side or as a QSIG End PINX with QSIG tunnelled over SIP within an IP network on one side and providing interworking to another IP-based signalling protocol on the other side. Detailed architecture and procedures for such a device are outside the scope of this standard.

A.2 Late availability of SDP parameters at the egress gateway

Early editions of this standard specified that the egress gateway caters for late availability of SDP parameters as needed. Newer editions specify signalling procedures that are identical for both decomposed and integrated gateways and are tolerant of all present and future enhancements to SDP, including, in particular, any that prevent the offer and answer being used interchangeably.

This is achieved when the ingress gateway, as opposed to the egress gateway, starts payload renegotiation immediately after confirmation of the SIP dialog when both gateways support the changed procedures. Compared to the early editions, this leads to an extra SIP round-trip during QSIG call establishment in the case of integrated gateways, but achieves identical signalling procedures between ingress gateway and egress gateway irrespective of the gateway architecture.

Furthermore, it allows the SDP offer-answer cycle to be conducted end-to-end between the ingress and egress media gateways, with only minimal intervention by the intermediate signalling gateways.

A.3 Compatibility with gateways that support early editions of this Standard

Reversing the direction of the re-INVITE request during call establishment may lead to collisions or glare conditions if the egress gateway is conformant to an early edition of this standard and the ingress gateway is conformant to a newer edition. On the other hand, if the ingress gateway conforms to an early edition of this standard and the egress gateway conforms to this new approach, neither side will send a re-INVITE request.

For these reasons an indicator for the new procedures is necessary for the ingress and the egress gateway. As the proposed change affects the SIP signalling procedures and not the QSIG procedures it is appropriate to use an indication at the SIP level. SIP User agent capabilities [12] can provide such an indicator by using IANA registered or vendor-specific unregistered media feature tags as a Contact: header field parameter. The latter option is used in this standard.

Both ingress and egress gateway are required to include such an indicator during the initial SIP INVITE transaction. Afterwards both gateways use either signalling procedures conformant to this standard (if both gateways support this standard) or fall back to procedures conformant to early editions of this standard. Such a fallback will of course result in potential loss of payload encryption, SDP capability negotiation or other extensions that introduce asymmetries into SDP.