

Network Working Group
Request for Comments: 5009
Category: Informational

R. Ejzak
Alcatel-Lucent
September 2007

Private Header (P-Header) Extension to
the Session Initiation Protocol (SIP) for Authorization of Early Media

Status of This Memo

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

Abstract

This document describes a private Session Initiation Protocol (SIP) header field (P-header) to be used by the European Telecommunications Standards Institute (ETSI) Telecommunications and Internet-converged Services and Protocols for Advanced Networks (TISPAN) for the purpose of authorizing early media flows in Third Generation Partnership Project (3GPP) IP Multimedia Subsystems (IMS). This header field is useful in any SIP network that is interconnected with other SIP networks and needs to control the flow of media in the early dialog state.

Table of Contents

1. Introduction	2
2. Applicability Statement	3
3. Conventions and Acronyms	3
4. Background on Early Media Authorization	4
4.1. Backward Early Media	5
4.2. Forward Early Media	5
5. Applicability of RFC 3959 and RFC 3960	6
6. Overview of Operation	6
7. Limitations of the P-Early-Media Header Field	8
8. The P-Early-Media Header Field	8
8.1. Procedures at the User Agent Client	10
8.2. Procedures at the User Agent Server	10
8.3. Procedures at the Proxy	11
9. Formal Syntax	11
10. Security Considerations	11
11. IANA Considerations	12
11.1. Registration of the "P-Early-Media" SIP Header Field	12
12. Acknowledgements	12
13. References	12
13.1. Normative References	12
13.2. Informative References	13

1. Introduction

This document defines the use of the P-Early-Media header field for use within SIP [1] messages in certain SIP networks to authorize the cut-through of backward and/or forward early media when permitted by the early media policies of the networks involved. The P-Early-Media header field is intended for use in a SIP network, such as a 3GPP IMS [13][14] that has the following characteristics: its early media policy prohibits the exchange of early media between end users; it is interconnected with other SIP networks that have unknown, untrusted, or different policies regarding early media; and it has the capability to "gate" (enable/disable) the flow of early media to/from user equipment.

Within an isolated SIP network, it is possible to gate early media associated with all endpoints within the network to enforce a desired early media policy among network endpoints. However, when a SIP network is interconnected with other SIP networks, only the boundary node connected to the external network can determine which early media policy to apply to a session established between endpoints on different sides of the boundary. The P-Early-Media header field provides a means for this boundary node to communicate this early media policy decision to other nodes within the network.

2. Applicability Statement

The use of this extension is only applicable inside a "Trust Domain" as defined in RFC 3325 [6]. Nodes in such a Trust Domain are explicitly trusted by its users and end-systems to authorize early media requests only when allowed by early media policy within the Trust Domain.

This document does NOT offer a general early media authorization model suitable for inter-domain use or use in the Internet at large. Furthermore, since the early media requests are not cryptographically certified, they are subject to forgery, replay, and falsification in any architecture that does not meet the requirements of the Trust Domain.

An early media request also lacks an indication of who specifically is making or modifying the request, and so it must be assumed that the Trust Domain is making the request. Therefore, the information is only meaningful when securely received from a node known to be a member of the Trust Domain.

Although this extension can be used with parallel forking, it does not improve on the known problems with early media and parallel forking, as described in RFC 3960 [4], unless one can assume the use of symmetric RTP.

Despite these limitations, there are sufficiently useful specialized deployments that meet the assumptions described above, and can accept the limitations that result, to warrant publication of this mechanism. An example deployment would be a closed network that emulates a traditional circuit switched telephone network.

3. Conventions and Acronyms

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [2].

The following acronyms are used in this document:

3GPP	- the Third Generation Partnership Project
ABNF	- Augmented Backus-Naur Form [5]
DTMF	- Dual Tone Multi-Frequency
ETSI	- European Telecommunications Standards Institute
IMS	- Internet Protocol Multimedia Subsystem [13][14]
MIME	- Multipurpose Internet Mail Extensions
NAT	- Network Address Translation
PSTN	- Public Switched Telephone Network

SDP - Session Description Protocol [7]
SIP - Session Initiation Protocol [1]
TISPAN - Telecommunications and Internet-converged Services and
Protocols for Advanced Networks
UA - User Agent [1]
UAC - User Agent Client [1]
UAS - User Agent Server [1]

4. Background on Early Media Authorization

PSTN networks typically provide call progress information as backward early media from the terminating switch towards the calling party. PSTN networks also use forward early media from the calling party towards the terminating switch under some circumstances for applications, such as digit collection for secondary dialing. PSTN networks typically allow backward and/or forward early media since they are used for the purpose of progressing the call to the answer state and do not involve the exchange of data between endpoints.

In a SIP network, backward early media flows from the User Agent Server (UAS) towards the User Agent Client (UAC). Forward early media flows from the UAC towards the UAS. SIP networks by default allow both forms of early media, which may carry user data, once the media path is established. Early media is typically desirable with a PSTN gateway as UAS, but not with SIP user equipment as UAS.

To prevent the exchange of user data within early media while allowing early media via PSTN gateways, a SIP network may have a policy to prohibit backward early media from SIP user equipment and to prohibit forward media towards SIP user equipment, either of which may contain user data. A SIP network containing both PSTN gateways and SIP end devices, for example, can maintain such an early media policy by gating "off" any early media with a SIP end device acting as UAS, gating "on" early media with a SIP end device acting as UAC, and gating "on" early media at each PSTN gateway.

Unfortunately, a SIP network interconnected with another SIP network may have no means of assuring that the interconnected network is implementing a compatible early media policy, thus allowing the exchange of user data within early media under some circumstances. For example, if a network "A" allows all early media with user equipment as UAC and an interconnected network "B" allows all early media with user equipment as UAS, any session established between user equipment as UAC in "A" and user equipment as UAS in "B" will allow bidirectional user data exchange as early media. Other combinations of early media policies may also produce similar undesirable results.

The purpose of the extension is to allow a SIP network interconnected to other SIP networks with different early media policies to correctly identify and enable authorized early media according to its policies.

4.1. Backward Early Media

Backward early media in the PSTN typically comprises call progress information, such as ringing feedback ("ringback"), or announcements regarding special handling such as forwarding. It may also include requests for further information, such as a credit card number to be entered as forward early media in the form of Dual Tone Multi-Frequency (DTMF) tones or speech. Backward early media of this type provides information to the calling party strictly for the purpose of progressing the call and involves no exchange of data between end users. The usual PSTN charging policy assumes that no data is exchanged between users until the call has been answered.

A terminating SIP User Agent (UA) outside of the SIP network, on the other hand, may provide any user data in a backward early media stream. Thus, if the network implements the usual early media policy, the network equipment gating the backward early media flow for the originating UA must distinguish between authorized early media from a terminating SIP endpoint and unauthorized early media from another SIP device outside of the network. Given the assumption of a transitive trust relationship between SIP servers in the network, this can be accomplished by including some information in a backward SIP message that identifies the presence of authorized backward early media. Since it is necessary to verify that this indication comes from a trusted source, it is necessary for each server on the path back to the originating UA to be able to verify the trust relationship with the previous server and to remove such an indication when it cannot do so. A server on the boundary to an untrusted SIP network can assure that no indication of authorized backward early media passes from an external UAS to a UAC within the network. Thus, the use of a private header field that can be modified by SIP proxies is to be preferred over the use of a Multipurpose Internet Mail Extensions (MIME) attachment that cannot be modified in this way.

4.2. Forward Early Media

Forward early media is less common than backward early media in the PSTN. It is typically used to collect secondary dialed digits, to collect credit card numbers, or to collect other DTMF or speech responses for the purpose of further directing the call. Forward early media in the PSTN is always directed toward a network server

for the purpose of progressing a call and involves no exchange of data between end users.

A terminating SIP UA outside of the SIP network, on the other hand, may receive any user data in a forward early media stream. Thus, if the network implements the usual early media policy, the network equipment gating the forward early media flow for the originating UA must distinguish between a terminating endpoint that is authorized to receive forward early media, and another SIP device outside of the network that is not authorized to receive forward early media containing user data. This authorization can be accomplished in the same manner as for backward early media by including some information in a backward SIP message that identifies that the terminating side is authorized to receive forward early media.

5. Applicability of RFC 3959 and RFC 3960

The private header extension defined in this document is applicable to the gateway model defined in RFC 3960 [4], since the PSTN gateway is the primary requestor of early media in an IMS. For the same reason, neither the application server model of RFC 3960, nor the early-session disposition type defined in RFC 3959 [3] is applicable.

The gateway model of RFC 3960 [4] allows for individual networks to create local policy with respect to the handling of early media, but does not address the case where a network is interconnected with other networks with unknown, untrusted, or different early media policies. Without the kind of information in the P-Early-Media header field, it is not possible for the network to determine whether cut-through of early media could lead to the transfer of data between end-users during session establishment.

Thus, the private header extension in this document is a natural extension of the gateway model of RFC 3960 [4] that is applicable within a transitive trust domain.

6. Overview of Operation

This document defines a new P-Early-Media header field for the purpose of requesting and authorizing requests for backward and/or forward early media. A UAC capable of recognizing the P-Early-Media header field may include the header field in an INVITE request. The P-Early-Media header field in an INVITE request contains the "supported" parameter.

As members of the Trust Domain, each proxy receiving an INVITE request must decide whether to insert or delete the P-Early-Media header field before forwarding.

A UAS receiving an INVITE request can use the presence of the P-Early-Media header field in the request to decide whether to request early media authorization in subsequent messages towards the UAC. After receiving an incoming INVITE request, the UAS requesting backward and/or forward early media will include the P-Early-Media header field in a message towards the UAC within the dialog, including direction parameter(s) that identify for each media line in the session whether the early media request is for backward media, forward media, both, or neither. The UAS can change its request for early media by including a modified P-Early-Media header field in a subsequent message towards the UAC within the dialog.

Each proxy in the network receiving the P-Early-Media header field in a message towards the UAC has the responsibility for assuring that the early media request comes from an authorized source. If a P-Early-Media header field arrives from either an untrusted source, a source not allowed to send backward early media, or a source not allowed to receive forward early media, then the proxy may remove the P-Early-Media header field or alter the direction parameter(s) of the P-Early-Media header field before forwarding the message, based on local policy.

A proxy in the network not receiving the P-Early-Media header field in a message towards the UAC may insert one based on local policy.

If the proxy also performs gating of early media, then it uses the parameter(s) of the P-Early-Media header field to decide whether to open or close the gates for backward and forward early media flow(s) between the UAs. The proxy performing gating of early media may also add a "gated" parameter to the P-Early-Media header field before forwarding the message so that other gating proxies in the path can choose to leave open their gates.

If the UAC is a trusted server within the network (e.g., a PSTN gateway), then the UAC may use the parameter(s) of the P-Early-Media header field in messages received from the UAS to decide whether to perform early media gating or cut-through and to decide whether or not to render backward early media in preference to generating ringback based on the receipt of a 180 Ringing response.

If the UAC is associated with user equipment, then the network will have assigned a proxy the task of performing early media gating, so that the parameter(s) of the P-Early-Media header field received at such a UAC do not require that the UAC police the early media flow(s), but they do provide additional information that the UAC may use to render media.

The UAC and proxies in the network may also insert, delete, or modify the P-Early-Media header field in messages towards the UAS within the dialog according to local policy, but the interpretation of the header field when used in this way is a matter of local policy and not defined herein. The use of direction parameter(s) in this header field could be used to inform the UAS of the final early media authorization status.

7. Limitations of the P-Early-Media Header Field

The P-Early-Media header field does not apply to any SDP with Content-Disposition: early-session [3].

When parallel forking occurs, there is no reliable way to correlate early media authorization in a dialog with the media from the corresponding endpoint unless one can assume the use of symmetric RTP, since the SDP messages do not identify the RTP source address of any media stream. When a UAC or proxy receives multiple early dialogs and cannot accurately identify the source of each media stream, it SHOULD use the most restrictive early media authorization it receives on any of the dialogs to decide the policy to apply towards all received media. When early media usage is desired for any reason and one cannot assume the use of symmetric RTP, it is advisable to disable parallel forking using callerprefs [9].

Although the implementation of media gating is outside the scope of this extension, note that media gating must be implemented carefully in the presence of NATs and protocols that aid in NAT traversal. Media gating may also introduce a potential for media clipping that is similar to that created during parallel forking or any other feature that may disable early media, such as custom ringback.

8. The P-Early-Media Header Field

The P-Early-Media header field with the "supported" parameter MAY be included in an INVITE request to indicate that the UAC or a proxy on the path recognizes the header field.

A network entity MAY request the authorization of early media or change a request for authorization of early media by including the P-Early-Media header field in any message allowed by Table 1 within the dialog towards the sender of the INVITE request. The P-Early-Media header field includes one or more direction parameters where each has one of the values: "sendrecv", "sendonly", "recvonly", or "inactive", following the convention used for Session Description Protocol (SDP) [7][8] stream directionality. Each parameter applies, in order, to the media lines in the corresponding SDP messages establishing session media. Unrecognized parameters SHALL be

silently discarded. Non-direction parameters are ignored for purposes of early media authorization. If there are more direction parameters than media lines, the excess SHALL be silently discarded. If there are fewer direction parameters than media lines, the value of the last direction parameter SHALL apply to all remaining media lines. A message directed towards the UAC containing a P-Early-Media header field with no recognized direction parameters SHALL NOT be interpreted as an early media authorization request.

The parameter value "sendrecv" indicates a request for authorization of early media associated with the corresponding media line, both from the UAS towards the UAC and from the UAC towards the UAS (both backward and forward early media). The value "sendonly" indicates a request for authorization of early media from the UAS towards the UAC (backward early media), and not in the other direction. The value "recvonly" indicates a request for authorization of early media from the UAC towards the UAS (forward early media), and not in the other direction. The value "inactive" indicates either a request that no early media associated with the corresponding media line be authorized, or a request for revocation of authorization of previously authorized early media.

The P-Early-Media header field in any message within a dialog towards the sender of the INVITE request MAY also include the non-direction parameter "gated" to indicate that a network entity on the path towards the UAS is already gating the early media, according to the direction parameter(s). When included in the P-Early-Media header field, the "gated" parameter SHALL come after all direction parameters in the parameter list.

When receiving a message directed toward the UAC without the P-Early-Media header field and no previous early media authorization request has been received within the dialog, the default early media authorization depends on local policy and may depend on whether the header field was included in the INVITE request. After an early media authorization request has been received within a dialog, and a subsequent message is received without the P-Early-Media header field, the previous early media authorization remains unchanged.

The P-Early-Media header field in any message within a dialog towards the UAS MAY be ignored or interpreted according to local policy.

The P-Early-Media header field does not interact with SDP offer/answer procedures in any way. Early media authorization is not influenced by the state of the SDP offer/answer procedures (including preconditions and directionality) and does not influence the state of the SDP offer/answer procedures. The P-Early-Media header field may or may not be present in messages containing SDP. The most recently

received early media authorization applies to the corresponding media line in the session established for the dialog until receipt of the 200 OK response to the INVITE request, at which point all media lines in the session are implicitly authorized. Early media flow in a particular direction requires that early media in that direction is authorized, that media flow in that direction is enabled by the SDP direction attribute for the stream, and that any applicable preconditions [11] are met. Early media authorization does not override the SDP direction attribute or preconditions state, and the SDP direction attribute does not override early media authorization.

Table 1 is an extension of Tables 2 and 3 in RFC 3261 [1] for the P-Early-Media header field. The column "PRA" is for the PRACK method [12]. The column "UPD" is for the UPDATE method [10].

Header field	where	proxy	ACK	BYE	CAN	INV	OPT	REG	PRA	UPD
P-Early-Media	R	amr	-	-	-	o	-	-	o	o
P-Early-Media	18x	amr	-	-	-	o	-	-	-	-
P-Early-Media	2xx	amr	-	-	-	-	-	-	o	o

Table 1: P-Early-Media Header Field

8.1. Procedures at the User Agent Client

A User Agent Client MAY include the P-Early-Media header field with the "supported" parameter in an INVITE request to indicate that it recognizes the header field.

A User Agent Client receiving a P-Early-Media header field MAY use the parameter(s) of the header field to gate or cut-through early media, and to decide whether to render early media from the UAS to the UAC in preference to any locally generated ringback triggered by a 180 Ringing response. If a proxy is providing the early media gating function for the User Agent Client, then the gateway model of RFC 3960 [4] for rendering of early media is applicable. A User Agent Client without a proxy in the network performing early media gating that receives a P-Early-Media header field SHOULD perform gating or cut-through of early media according to the parameter(s) of the header field.

8.2. Procedures at the User Agent Server

A User Agent Server that is requesting authorization to send or receive early media MAY insert a P-Early-Media header field with appropriate parameters(s) in any message allowed in table 1 towards the UAC within the dialog. A User Agent Server MAY request changes in early media authorization by inserting a P-Early-Media header

field with appropriate parameter(s) in any subsequent message allowed in table 1 towards the UAC within the dialog.

If the P-Early-Media header field is not present in the INVITE request, the User Agent Server MAY choose to suppress early media authorization requests and MAY choose to execute alternate early media procedures.

8.3. Procedures at the Proxy

When forwarding an INVITE request, a proxy MAY add, retain, or delete the P-Early-Media header field, depending on local policy and the trust relationship with the sender and/or receiver of the request.

When forwarding a message allowed in Table 1 towards the UAC, a proxy MAY add, modify, or delete a P-Early-Media header field, depending on local policy and the trust relationship with the sender and/or receiver of the message. In addition, if the proxy controls the gating of early media for the User Agent Client, it SHOULD use the contents of the P-Early-Media header field to gate the early media, according to the definitions of the header field parameters defined in clause 8.

9. Formal Syntax

The syntax of the P-Early-Media header field is described below in ABNF, according to RFC 4234 [5], as an extension to the ABNF for SIP in RFC 3261 [1]. Note that not all combinations of em-param elements are semantically valid.

```
P-Early-Media = "P-Early-Media" HCOLON
                [ em-param *(COMMA em-param) ]
em-param      = "sendrecv" / "sendonly" / "recvonly"
                / "inactive" / "gated" / "supported" / token
```

10. Security Considerations

The use of this extension is only applicable inside a "Trust Domain", as defined in RFC 3325 [6]. This document does NOT offer a general early media authorization model suitable for inter-domain use or use in the Internet at large.

There are no confidentiality concerns associated with the P-Early-Media header field. It is desirable to maintain the integrity of the direction parameters in the header field across each hop between servers to avoid the potential for unauthorized use of early media. It is assumed that the P-Early-Media header field is used within the context of the 3GPP IMS trust domain or a similar trust domain,

consisting of a collection of SIP servers maintaining pair wise security associations.

Within the trust domain of a network it is only necessary to police the use of the P-Early-Media header field at the boundary to user equipment served by the network and at the boundary to peer networks. It is assumed that boundary servers in the trust domain of a network will have local policy for the treatment of the P-Early-Media header field as it is sent to or received from any possible server external to the network. Since boundary servers are free to modify or remove any P-Early-Media header field in SIP messages forwarded across the boundary, the integrity of the P-Early-Media header field can be verified to the extent that the connections to external servers are secured. The authenticity of the P-Early-Media header field can only be assured to the extent that the external servers are trusted to police the authenticity of the header field.

11. IANA Considerations

11.1. Registration of the "P-Early-Media" SIP Header Field

Name of Header field: P-Early-Media

Short form: none

Registrant: Richard Ejzak
ejzak@alcatel-lucent.com

Normative description: Section 8 of this document

12. Acknowledgements

The author would like to thank Miguel Garcia-Martin, Jan Holm, Sebastien Garcin, Akira Kurokawa, Erick Sasaki, James Calme, Greg Tevonian, Aki Niemi, Paul Kyzivat, Gonzalo Camarillo, Brett Tate, Jon Peterson, Alfred Hoenes, and David Black for their significant contributions made throughout the writing and reviewing of this document.

13. References

13.1. Normative References

- [1] 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.

- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [3] Camarillo, G., "The Early Session Disposition Type for the Session Initiation Protocol (SIP)", RFC 3959, December 2004.
- [4] Camarillo, G. and H. Schulzrinne, "Early Media and Ringing Tone Generation in the Session Initiation Protocol (SIP)", RFC 3960, December 2004.
- [5] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 4234, October 2005.
- [6] Jennings, C., Peterson, J., and M. Watson, "Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks", RFC 3325, November 2002.
- [7] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [8] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.
- [9] Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Caller Preferences for the Session Initiation Protocol (SIP)", RFC 3841, August 2004.
- [10] Rosenberg, J., "The Session Initiation Protocol (SIP) UPDATE Method", RFC 3311, October 2002.
- [11] Camarillo, G., Marshall, W., and J. Rosenberg, "Integration of Resource Management and Session Initiation Protocol (SIP)", RFC 3312, October 2002.
- [12] Rosenberg, J. and H. Schulzrinne, "Reliability of Provisional Responses in Session Initiation Protocol (SIP)", RFC 3262, June 2002.

13.2. Informative References

- [13] 3GPP "TS 23.228: IP Multimedia Subsystem (IMS); Stage 2 (Release 7)", 3GPP 23.228, March 2007,
ftp://ftp.3gpp.org/specs/archive/23_series/23.228/.
- [14] 3GPP "TS 24.229: IP Multimedia Call Control Protocol based on SIP and SDP; Stage 3 (Release 7)", 3GPP 24.229, March 2007,
ftp://ftp.3gpp.org/specs/archive/24_series/24.229/.

ETSI documents can be downloaded from the ETSI Web server, "<http://www.etsi.org/>". Any 3GPP document can be downloaded from the 3GPP Web server, "<http://www.3gpp.org/>". See specifications.

Authors Address

Richard Ejzak
Alcatel-Lucent
1960 Lucent Lane
Naperville, IL 60566
USA

Phone: +1 630 979 7036
EMail: ejzak@alcatel-lucent.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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

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

Intellectual Property

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

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

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

