Network Working Group
INTERNET-DRAFT
Intended status: Informational
Expires: June 17, 2009

R. Ejzak
Alcatel-Lucent

December 17, 2008

Extension to the Session Description Protocol (SDP) for Bypass of Border Gateways <draft-ejzak-mmusic-bg-bypass-00.txt>

Status of this memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

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

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

The list of current Internet-Drafts can be accessed at http://www.ietf.org/1id-abstracts.txt

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

Abstract

This document describes an extension to the Session Description Protocol (SDP) that can be used by systems of cooperating networks using Application Level Gateways (ALG) to insert border gateways performing as Network Address Port Translators (NAPT) between their IP realms to identify when border gateways can be bypassed for more efficient media flow. This extension can be used by networks based on a protocol using the SDP offer/answer model, such as the IP Multimedia Subsystem (IMS) of the Third Generation Partnership Project (3GPP), which is based on the Session Initiation Protocol (SIP). ALGs using this extension can determine within a single SDP offer/answer transaction when the insertion of a new border gateway would cause the media path to re-enter an IP realm visited elsewhere

within the media path, and to bypass one or more border gateways that would otherwise be included in the media path. This extension also works with hosted NAPT traversal schemes to establish a direct media path between endpoints within the same IP realm. Optional procedures provide additional means to improve media flow.

Table of Contents

1. Introduction

The IP Multimedia Subsystem (IMS) [20] [21] and other SIP networks have the option to deploy border gateways between the IP realms defined by each network. Within an IP realm every endpoint is reachable from any other endpoint using a common address space. Each border gateway typically provides a firewall or Network Address Port Translator (NAPT) [13] to limit access to endpoints within a realm. An Application Layer Gateway (ALG) controls each border gateway to allocate new IP addresses and ports as necessary for each SDP media line and updates the SDP connection and port information in each forwarded SDP offer and answer to effectively insert the border gateway into each end-to-end multimedia stream.

The media path associated with a multimedia stream may traverse an arbitrary number of IP realms between endpoints. As long as each border gateway in the media path has no connection to IP realms on the media path other than its two directly connected IP realms, there is no option to optimize the media path using the allocated border gateway resources. But if either endpoint or any border gateway on the path has direct access to one of the other IP realms on the path, then a shorter media path exists. A sequence of ALGs implementing the procedures herein, where each ALG can determine the IP address and port information for entities on the media path in its interconnected IP realms, will be able to establish a media path with the minimum number of border gateways without compromising any of the access controls associated with the border gateways on the path. If one or more ALGs on the signaling path do not implement the procedures then border gateway bypass can still occur but some potentially bypassable border gateways may remain in the media path.

The procedures described herein also include an "active-bypass" option to attempt to find a shorter media path segment between existing border gateways associated with the path. This option requires additional SIP signaling to establish a SIP dialog for each alternate media path segment candidate, whereas the base algorithm works by adding information to existing SDP offer/answer messages. Due to this additional signaling overhead, this option should only be used when it can be determined that dramatic improvement is possible for a media path segment.

This extension also works with hosted NAPT traversal schemes to establish a direct media path between endpoints within the same IP realm. If the endpoints are in different IP realms, this extension cannot bypass an ALG/BG that is coordinating the traversal of a Residential Gateway (RG) NAPT, although it is possible that a combination of NAPT traversal techniques can achieve this. This

document does not analyze combination methods to address this limitation. Since networks using ALG/BGs typically perform other

media path functions at the ALG/BG configured to traverse the RG/NAPT, this is not a significant limitation.

RFC 3264 [3] describes the SDP offer/answer model, which enables SIP networks to establish end-to-end media paths for the multimedia streams in each session. This document describes two SDP extension attributes and some extensions to ALG procedures for forwarding SDP offers and answers. ALGs on the path manipulate the SDP as necessary within a single end-to-end SDP offer/answer transaction to enable establishment of an end-to-end media path with the minimum of border gateways. The SDP extension attributes describe media connection and port information for each IP realm on the path that is a candidate to bypass one or more border gateways on the path.

This document describes an extension and optimization of the ALG approach to NAPT traversal. Other options for NAPT traversal include the Middlebox Control Protocol [14], Session Traversal Utilities for NAT (STUN) [18], the STUN Relay Usage [19], and Realm Specific IP [11] [12]. The most recent and comprehensive approach to NAPT traversal is Interactive Connectivity Establishment (ICE) [17], which uses STUN to identify candidate addresses for NAPT traversal for media streams established by the offer/answer model.

While an ALG approach may require the insertion of a SIP back to back user agent (B2BUA) to modify SDP whenever a border gateway is inserted in the media path, ICE also has several disadvantages. ICE requires the deployment of STUN servers in each IP realm, a means of advertising the location of available STUN servers to SIP endpoints, extra signaling to discover candidate addresses for inclusion in SDP offers and answers, extra signaling to communicate the selected connection information, and implementation of the ICE procedures in the endpoints. With ICE, border gateways must be configured to allow signaling between endpoints and STUN servers, and do not receive definitive information on which ones are actually used and which remote addresses will be used in the RTP [15] stream. This makes it difficult for border gateways to limit access to known IP source addresses and to predict bandwidth usage, which are two important reasons for deploying border gateways.

The border gateway bypass procedures in this document, while requiring the use of ALGs, avoid the requirement to deploy STUN servers, require no additional signaling beyond what is needed for a single end-to-end SDP offer/answer transaction (although an optional procedure does generate additional signaling), require no new procedures to be supported by endpoints, allow border gateways to limit access to known IP source addresses, and allow border gateways to predictably manage aggregate bandwidth usage for all sessions.

Since this extension does not incorporate end-to-end connectivity checks of the media path, it requires accurate provisioning of the IP realms.

2. Applicability Statement

The use of this extension is only applicable inside a "Trust Domain" as defined in RFC 3325 [4]. Nodes in such a Trust Domain are explicitly trusted by its users and end-systems to inspect and manipulate SDP messages as necessary to traverse and/or bypass firewalls and NATS while limiting access from unauthorized sources to endpoints in IP realms associated with the Trust Domain.

Since the procedures in this document include an option to cryptographically certify the candidate connection and port information from each IP realm, they can be used under some circumstances when the signaling traverses non-trusted networks or the Internet at large.

This extension requires that ALGs on the signaling path have the ability to access and manipulate SDP messages, which is inconsistent with the general recommendation that these messages be encrypted and integrity protected end-to-end.

In the interest of algorithmic simplicity, this extension finds improved media paths in most cases according to the available information, but not under all circumstances.

This document does NOT offer a general model for optimal configuration of border gateways in the Internet at large.

This extension assumes that there is at most a single set of connection and port information for each SDP media line, consistent with existing RFCs. Possible future SDP extensions that allow description of alternative connection or port capabilities may not be compatible.

This extension makes some assumptions about the behavior of ALGs not implementing the extension that may not always be valid. See section 6.5 for a discussion of the compatibility issues and work-arounds. The extension also has some limitations when handling an unspecified address as connection information from an endpoint. See section 6.4.

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 an IMS network using

border gateways to interconnect multimedia sessions with other networks.

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 [1].

The following acronyms are used in this document:

- 3GPP - the Third Generation Partnership Project
- 3pcc - Third Party Call Control [16]
- ABNF - Augmented Backus-Naur Form [6]
- ALG - Application Layer Gateway [13]
- B2BUA - Back to Back User Agent [2]
- BG - Border Gateway
- FQDN - Fully Qualified Domain Name
- GRUU - Globally Reachable UA URI [8]
- ICE - Interactive Connectivity Establishment [17]
- IMS - Internet Protocol Multimedia Subsystem [20] [21]
- IP - Internet Protocol
- IPSEC - IP Security
- IPv4 - IP Version 4
- IPv6 - IP Version 6
- LAN - Local Area Network
- MD5 - Message-Digest 5 Algorithm [9]
- NAT - Network Address Translation [13]
- NAPT - Network Address Port Translation [13]
- RG - Residential Gateway
- RTCP - RTP Control Protocol [15]
- RTP - Real-time Transport Protocol [15]
- SDP - Session Description Protocol [7]
- SIP - Session Initiation Protocol [2]
- SP - Space
- STUN - Session Traversal Utilities for NAT [18]
- TCP - Transport Control Protocol
- UA - User Agent [2]
- UDP - User Datagram Protocol
- URI - Uniform Resource Identifier
- WGS - World Geodetic System [24]

4. Overview of Operation

4.1.

Figure 1 shows a typical call configuration between endpoints UA1 and UA2, where the SIP signaling goes between the UAs via at least one ALG (four are shown) and other SIP servers not shown, and one RTP multimedia stream goes between the UAs via the BGs and possibly an RG associated with each UA (only one RG is shown associated with UA2). Each BG is controlled by its corresponding ALG. R1, R2, etc., in the figure represent the IP realms associated with each segment of the media path.

The media path for each multimedia stream between the UAs is established via an end-to-end SDP offer/answer exchange where each ALG may choose to modify the connection and port information associated with each media line in the SDP to insert its BG in the media path according to normal ALG procedures. Each ALG may also perform the base algorithm procedures to identify when one or more BGs and/or RGs can be bypassed and to modify the forwarded SDP messages to implement the corresponding changes in the media path to bypass the BGs.

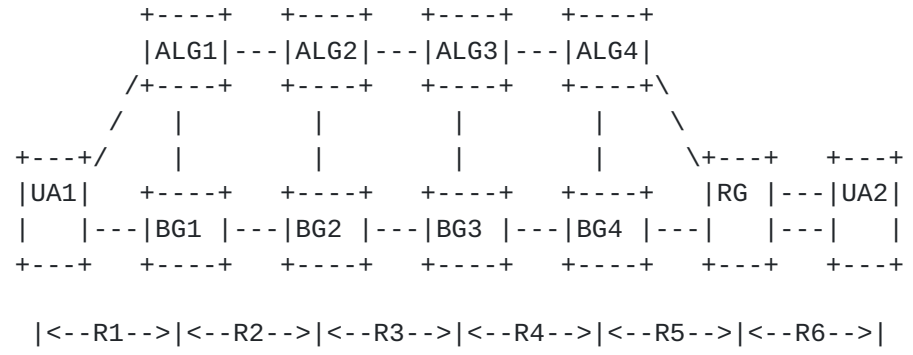

Figure 1: Example Call Configuration

Figure 2 shows another example call configuration where secondary BGs are used to establish a media path with fewer BGs. ALG1 through ALG5 initially allocate BG1a, BG2, BG4, BG4 and BG5a as ALGs forward the initial SDP offer towards UA2 from UA1. These BGs enable traversal of unique IP realms R1 through R6 (not labeled in the figure). Since these BGs do not create any loop in the media path, there is no possibility to bypass any of them if the algorithm is limited to finding loops in a fixed media path.

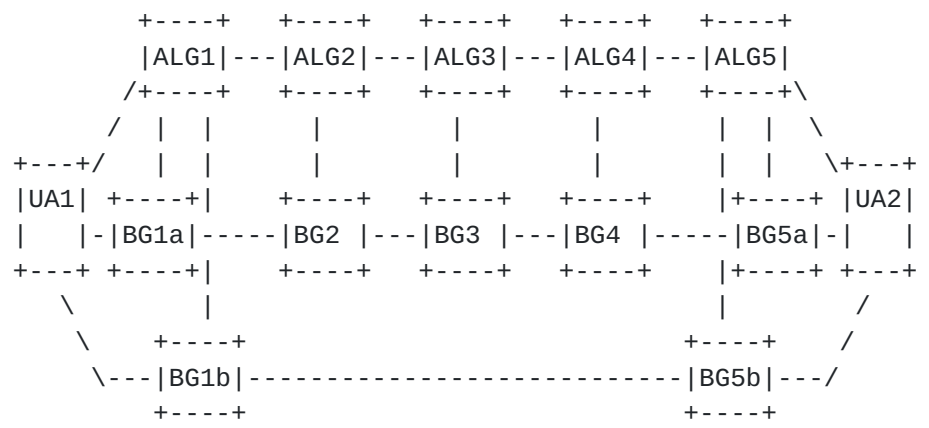

Figure 2: Example Configuration using Secondary BGs

While forwarding the initial SDP, if an ALG along the way, such as ALG1, controls BG(s) that have access to IP realm(s) other than those IP realms that it controls on the default media path (i.e., not R1 or R2), then the ALG can advertise its ability to access additional IP realm(s) by including information about them in the forwarded SDP.

If a subsequent ALG (e.g., ALG5) determines that it controls a BG (e.g., BG5b) that has a direct connection to an IP realm accessible from a BG controlled by a previous ALG in the path (e.g., ALG1 and BG1b), then the ALG may choose to use this alternative media path if it appears to be an improvement over the initial path. In this example, the algorithm establishes an alternative media path from UA1 to UA2 via BG1b and BG5b while significantly reducing the number of BGs traversed. Note that the IP realm between BG1b and BG5b in the example (R7) will not match any of the IP realms R1 through R6. If the connections exist, the algorithm may also generate alternative paths either via BG1a and BG5b, via BG1b and BG5a, or via BG1a and BG5a, for example (not shown).

The border gateway bypass base algorithm and active-bypass option (described in the next section) assume ICE is not used by any entity in the architecture. Although hybrid procedures are possible, they are beyond the scope of this document.

It is assumed that the UAs participate in standard SDP offer/answer negotiation by presenting standard connection and port information for each media line according to RFC 4566 [7], RFC 3264 [3] and possibly other extensions. If necessary, the ALGs may use the rtcp attribute defined in RFC 3605 [5] to identify an RTCP port not using the expected default value.

The border gateway bypass base algorithm and the active-bypass option are may be implemented only within the ALGs. The procedures have no impact on any aspect of SDP offer/answer negotiation other than the connection and port information associated with each media line.

This document defines an SDP extension attribute 'visited-realm' that provides connection and port information for a prior IP realm visited on the signaling path. Each instance of visited-realm has an instance number, realm identifier, connection/port data, and optional cryptographic signature computed using an algorithm private to each IP realm so as to ensure the integrity of the visited-realm data.

This document also defines an SDP extension attribute 'secondary-realm' that provides connection and port information for secondary IP realms associated with the signaling path. The secondary-realm attribute includes the same types of information as the visited-realm attribute.

Note that the connection and port information in each SDP offer/answer transaction within a SIP dialog must be handled the same way, as described in this document, re-allocating and de-allocating BGs as necessary with each SDP offer/answer transaction to accommodate any potential changes in the IP realms associated with the session endpoints.

4.2. Overview of Operation of the Active-Bypass Option

Figure 3 shows an example of the use of the base algorithm with the active-bypass option. If the initial BG allocations traversing IP realms R1 through R6 do not offer an opportunity to bypass any BGs (as in figure 2), and if no connections exist to offer any of the alternative options available in the base algorithm, then the active-bypass option can discover additional alternative(s). Note that in this case BG1b and BG5b do not share a common IP realm (in fact, all of the IP realms are different in this example), so the active-bypass option creates a new signaling path via ALG6 to establish a new media path segment via BG6.

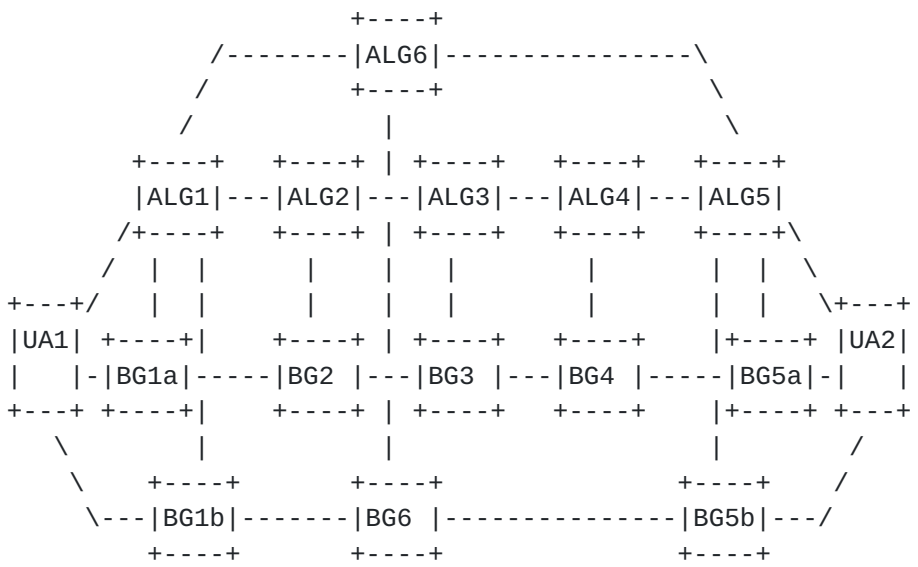

Figure 3: Example Configuration with Active-Bypass Option

When implementing the active-bypass option, the following additional information may be included in each visited-realm and secondary-realm attribute generated by the base algorithm for an SDP offer, if available: the approximate geo-location of the corresponding BG; the approximate delay of IP packets on the previous media path segment between this BG and the immediately preceding BG or endpoint; the approximate packet loss rate on the same media path segment; and if the ALG is reachable via a globally unique host name, then a globally reachable address of the ALG with a unique instance id for the corresponding SIP dialog and media line, in the form of a temporary GRUU [8].

Each ALG should include the geo-location, delay and loss information in the first visited-realm attribute generated for an SDP offer, and may include them for other visited-realm or secondary-realm attributes if the information differs significantly from the first. Each ALG may include the GRUU in the first visited-realm attribute generated for a media line in an SDP offer. There is no need to repeat the GRUU in subsequent visited-realm or secondary-realm attributes for the same media line.

When processing the SDP answer in the second phase of the base algorithm, after determining which BGs (if any) are to be bypassed as a result of the base algorithm, each ALG that still controls a BG determines if there is the possibility that a significantly shorter media path segment can be established via another ALG reachable via a GRUU. Each ALG makes this determination based on the available geo-location, delay and packet loss information associated with each

BG and media path segment.

If an ALG determines that it may be able to establish a shorter media path segment, the ALG (e.g., ALG5) sends a SIP INVITE request to the "best" ALG reachable via a GRUU (e.g., ALG1) to establish a separate dialog and corresponding alternate media path segment (e.g., via ALG6 and BG6). If the ALG is successful in establishing the alternate media path segment and it appears to be significantly better than the corresponding one determined by the base algorithm, then the ALGs instruct the BGs to insert the shorter path segment into the overall media path.

Figure 4 shows a call flow that corresponds to the configuration in figure 3.

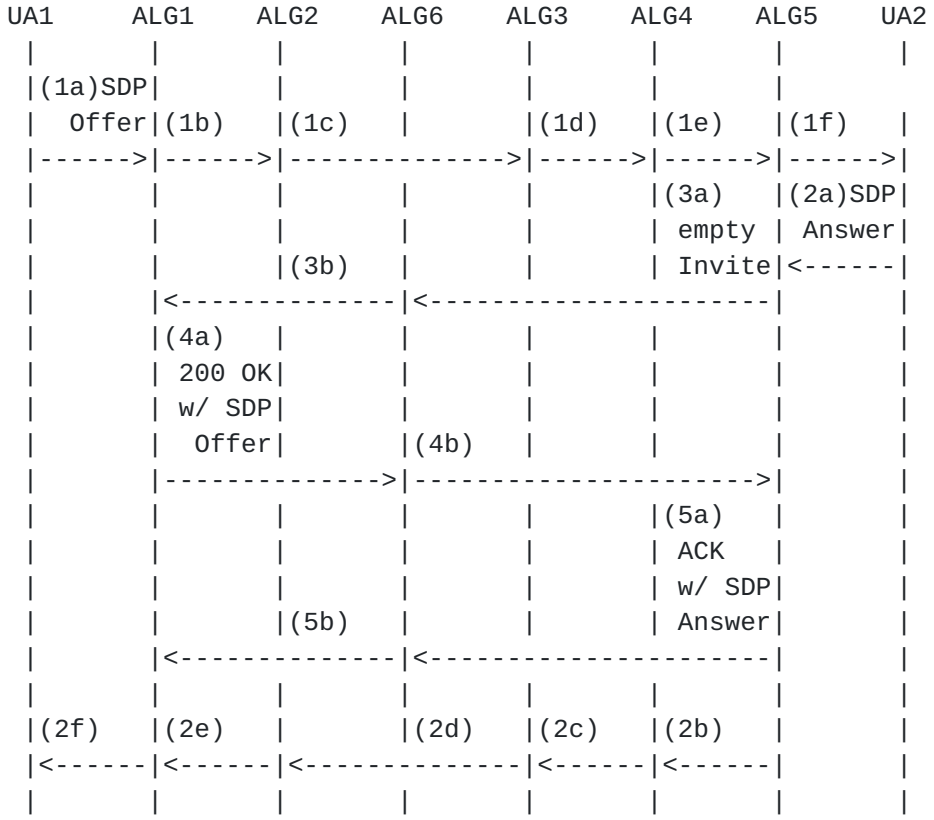

Figure 4: Example Flow with Active-Bypass Option

Steps 1a to 1f describe the progression of SDP offers via the ALGs from UA1 to UA2 and steps 2a to 2f describe the corresponding progression of SDP answers according to the base algorithm. After step 2a, ALG5 determines that it may be able to establish a shorter media path segment via ALG1 and sends an empty SIP INVITE request to ALG1 via ALG6 in steps 3a and 3b. Steps 4a, 4b, 5a and 5b describe a new SDP offer/answer transaction between ALG1 and ALG5 via ALG6 which attempts to establish an alternate media path segment. If an

alternate media path segment is successfully established and is a

significant improvement, ALG5 signals the selection of the alternate media path segment to ALG1 in steps 2b through 2e. ALG1 incorporates the alternate media path segment into the media path for the primary dialog before forwarding the final SDP answer to UA1 in step 2f.

5. IP realm considerations

For the procedures in this specification, the term "IP realm" has a specific meaning beyond the use of the term "realm" for digest authentication [10]. An IP realm has two purposes: 1) to identify a private means by which network entities sharing private information can verify that data communicated via intermediaries remains unchanged; and 2) to identify when one network entity is reachable from another via a fully interconnected common IP address space.

The syntax for the visited-realm and secondary-realm extension attributes in section 7 clearly describes means of accomplishing purpose 1) using security credentials.

There are many network configurations for which 2) is applicable, as described below.

For example, all hosts in a residence on a private LAN behind an RG/NAPT can be considered to be in their own IP realm. An operator providing hosted NAPT traversal from an ALG in the network can identify a separate IP realm for each such residence and provide the security framework to ensure, for example, that it is possible to provide a media path directly between hosts in the same residence when they are involved in an end-to-end session established via SIP servers in an external network, thus bypassing a potentially significant number of BGs that would otherwise have been allocated using normal ALG procedures.

A very similar example is when there is a private enterprise network using a private IP address space with one or more NAPTs to external networks. The same principles apply as in the residential case. An ALG providing hosted NAPT traversal creates an IP realm for the enterprise, associates the appropriate IP addresses from the enterprise IP realm with a selected identifier and looks for opportunities to bypass BGs in the network.

Session endpoints not associated with NAPTs may also be directly connected to an ALG in the network. Those mutually reachable endpoints connected to an ALG may be assigned an IP realm.

Once a media path enters a network isolated with ALGs from access

and peer networks, all addresses associated with media connections

to BGs that are mutually reachable within the network can be considered part of another IP realm. Whenever an ALG forwards an SDP offer back into such an IP realm after visiting it on a prior hop, there is an opportunity to bypass all BGs visited on the "loop" back into the IP realm.

Two interconnected networks may have ALG/BGs directly connected via IPSEC associations over the Internet. There may be one or more IP realms created just to identify these limited connectivity options. Since there will be limited opportunities to bypass BGs via these IP realms, a network MAY choose to leave these IP realms unidentified and MAY choose not to forward visited-realm or secondary-realm information for these IP realms.

IP addresses reachable from the open internet are associated with the pre-defined IP realm "IN".

These are just a few examples of IP realms. Since no connectivity checks are used to verify reachability, IP realms MUST be provisioned to correctly identify mutually reachable IP addresses. It is RECOMMENDED that networks provide other means to verify reachability between endpoints in their defined IP realms.

6. ALG procedures

The ALG procedures apply in this section SHALL apply separately to each media line with non-zero port value in each SDP message, and SHALL apply separately to each SDP offer/answer transaction.

6.1. ALG handling of SDP offer

When an ALG receives an SDP offer from a UA or another ALG, it first determines the IP realm for the outgoing segment of the media path associated with the outgoing signaling. For example, in Figure 1, if UA1 initiates an SDP offer towards UA2, then the outgoing IP realm for ALG1 is R2, the outgoing IP realm for ALG2 is R3, and the outgoing IP realm for ALG4 is R6 (rather than R5). Since ALG4 is managing the traversal of the RG to R6, BG4 and IP realm R5 are not eligible for bypass, unless both media path IP endpoints are in the same IP realm R6, so that all BGs and RGs in the media path are bypassed.

The ALG examines all previously visited IP realms represented by the visited-realm and secondary-realm instances for the media line in the received SDP offer. If the outgoing IP realm matches any of the visited-realm or secondary-realm instances, then the ALG can bypass

one or more BGs, including the one it controls. The ALG SHOULD select the earliest matching IP realm and determine the number of

BGs that can be bypassed by substituting the connection and port information from this earliest IP realm into the forwarded SDP offer.

The ALG then determines if a BG under its control has access both to the outgoing IP realm and to an IP realm associated with a prior visited-realm or secondary-realm instance in the received SDP offer. In this case the ALG may be able to bypass one or more BGs, but not the one it controls. The ALG SHOULD select the earliest IP realm accessible from a BG under its control and determine the number of BGs that can be bypassed by connecting the prior IP realm directly to the BG. Note that in this case use of a visited-realm instance associated with the immediately prior ALG is pointless since no BGs are bypassed. Also note that in this case use of a secondary-realm instance associated with the immediately prior ALG will not reduce the number of BGs in the path, but may still result in a superior media path if, for example, it can be determined that there is less IP layer congestion using this path.

The ALG SHALL then select one of the following four cases depending on applicability and local policy.

1. Bypass the controlled BG and one or more prior BGs.
2. Bypass the controlled BG.
3. Bypass prior BGs.
4. Bypass no BGs.

The most common local policy will be to select the case that bypasses the largest number of BGs. In cases 3 and 4, the ALG MAY signal that it is not to be bypassed by removing all visited-realm and secondary-realm instances associated with incoming and prior IP realms from the forwarded SDP offer. The ALG SHOULD signal that it is not to be bypassed if it performs any necessary media function other than address translation, e.g., transcoding.

6.1.1. SDP offer case 1: bypass controlled BG and prior BGs

In case 1, the ALG determines that there exists a visited-realm or secondary-realm instance for the media line in the received SDP offer that does not match the incoming IP realm for that media line but does match the IP realm to be used for the media line in the forwarded SDP offer.

The ALG

1. SHALL replace the connection and port information for the media line in the SDP offer with the connection and port information from the earliest visited-realm or secondary-realm instance

associated with the outgoing IP realm;

2. SHALL delete every visited-realm or secondary-realm instance with realm-number value higher than the one used to populate the outgoing connection and port data, and
3. SHALL forward the modified SDP offer.

An example of case 1, using Figure 1 as reference, is that upon receiving an SDP offer from the direction of UA1, ALG3 determines that R4 and R1 are instances of the same IP realm. ALG3 substitutes the connection and port information from UA1 into the outgoing SDP offer and deletes the visited-realm instances for R2 and R3 from the SDP before forwarding. After the end-to-end SDP offer/answer transaction is completed, the media path will bypass BG1, BG2 and BG3.

6.1.2. SDP offer case 2: bypass controlled BG

In case 2 (bypass only the controlled BG), the ALG determines that the outgoing IP realm is accessible from the incoming IP realm represented by the IP connection and port information for the media line in the received SDP offer. If there is a visited-realm or secondary-realm instance for the incoming IP realm that matches the media line in the received SDP offer (not necessarily matching the incoming connection information), the ALG SHALL forward the received SDP offer without change. Otherwise the ALG SHALL construct a new visited-realm instance from the connection and port information for the media line in the incoming SDP offer and SHALL add this visited-realm instance to the SDP offer before forwarding.

For case 2, the received SDP offer will normally include a visited-realm or secondary-realm instance that matches the incoming IP realm unless the previous ALG does not support the BG bypass procedures. Adding this missing information provides for more opportunities to perform BG bypass.

6.1.3. SDP offer case 3: bypass prior BGs

In case 3, the ALG determines that a BG under its control has access both to the outgoing IP realm and to an IP realm other than the incoming IP realm that matches a prior visited-realm or secondary-realm instance for the media line in the received SDP offer.

The ALG:

1. SHALL use the connection and port information from the earliest visited-realm or secondary-realm instance accessible from the BG as the remote connection and port information for the side of the BG directed towards the offerer;

2. SHALL replace the connection and port information for the media line in the SDP offer with the connection and port information from the side of its BG directed toward the answerer;

3. SHALL delete from the SDP answer every visited-realm and secondary-realm instance with realm-number higher than the realm-number for the earliest visited-realm or secondary-realm instance accessible from the BG;
4. MAY, if the ALG requires that its BG remain in the media path, remove all visited-realm and secondary-realm instances from the SDP offer;
5. SHOULD, if the outgoing IP realm does not match any of the visited-realm or secondary-realm instances in the SDP offer, add to the SDP offer a visited-realm instance for the IP realm associated with the connection and port information for the media line in the modified SDP offer;
6. MAY add to the SDP offer a secondary-realm instance for each IP realm that does not match any other visited-realm or secondary-realm instance for the media line but for which there is a BG controlled by the ALG that has access both to this IP realm and to the incoming IP realm associated with the BG previously allocated by this ALG and
7. SHALL forward the modified SDP offer.

An example of case 3, using Figure 1 as reference, is that upon receiving an SDP offer from the direction of UA1, ALG4 determines that BG4 has access to R2. ALG4 substitutes its BG connection and port information into the SDP offer, uses the connection and port information from the visited-realm instance for R2 as the remote connection and port information for the UA1 side of BG4, deletes the visited-realm instances for R3 and R4 from the SDP offer, and adds the visited-realm instance for R5 before forwarding. After the end-to-end SDP offer/answer transaction is completed, the media path will bypass BG2 and BG3.

6.1.4. SDP offer case 4: bypass no BGs

In case 4, the ALG bypasses no BGs.

The ALG:

1. SHOULD, if there is no visited-realm or secondary-realm instance that matches the IP realm associated with the media line in the received SDP offer and the ALG allows bypass of its BG, construct a new visited-realm instance from the connection and port information for the media line in the incoming SDP offer and add this visited-realm instance to the SDP offer to be forwarded;
2. SHALL replace the connection and port information for the media line in the SDP offer with the connection and port information from the side of its BG directed toward the answerer;
3. MAY, if the ALG requires that its BG remain in the media path,

remove all visited-realm and secondary-realm instances from the SDP offer;

4. SHOULD, if the outgoing IP realm does not match any of the visited-realm or secondary-realm instances in the SDP offer, add a visited-realm instance for the IP realm associated with the connection and port information for the media line in the forwarded SDP offer;
5. MAY add to the SDP offer a secondary-realm instance for each IP realm that does not match any other visited-realm or secondary-realm instance for the media line but for which there is a BG controlled by the ALG that has access both to this IP realm and to the IP realm associated with the received SDP offer and
6. SHALL forward the modified SDP offer.

If the ALG is not performing hosted NAPT traversal on the side towards the SDP offerer, the ALG SHALL use the connection and port information from the incoming SDP offer as the remote connection and port information for the side of the BG directed towards the offerer. If the ALG is performing hosted NAPT traversal on the side towards the SDP offerer, the ALG/BG MUST discover the address of the RG via latching or other unspecified technique. Except for the insertion of the visited-realm and secondary-realm instance(s) in the outgoing SDP offer, case 4 corresponds to standard ALG behavior.

6.2. ALG handling of SDP answer in base algorithm

After forwarding an SDP offer, the ALG SHALL keep information about which of the four cases it selected for handling of BG bypass and which visited-realm and secondary-realm instances it received and added to the forwarded SDP offer. The ALG uses this information in the processing of the corresponding SDP answer, but there are additional sub-cases to be considered since downstream ALGs can also bypass BGs already visited, and other ALGs in the path may or may not support the BG bypass procedures. Note that there is at most one identified instance of each IP realm (as represented by a visited-realm or secondary-realm instance) in the SDP offer that reaches its final destination. The ALG uses this fact to correctly process the SDP answer. Unidentified IP realms represent lost opportunities for BG bypass.

To help distinguish the additional sub-cases when processing the SDP answer, the ALG SHALL insert into the connection information for the media line in the forwarded SDP answer either: 1) a valid IP address for the corresponding IP realm or 2) an unspecified address. For this purpose the unspecified address for IPv4 is '0.0.0.0' and for IPv6 is a domain name within the .invalid DNS top level domain (rather than the IPv6 unspecified address '0::0'). When signaling the unspecified address for the connection information, the port information MUST have a non-zero value.

The ALG must consider the following sub-cases when receiving an SDP answer:

a. The connection and port information for the media line in the SDP answer received by the ALG is *valid* for its IP realm. This IP realm matches the IP realm associated with the connection and port information for the corresponding media line in the SDP offer forwarded by the ALG.
b. The connection information for the media line in the SDP answer received by the ALG is the *unspecified address*. The visited-realm instance in the SDP answer matches a visited-realm or secondary-realm instance previously *received* in the SDP offer.
c. The connection information for the media line in the SDP answer received by the ALG is the *unspecified address*. The visited-realm instance in the SDP answer matches the IP realm associated with the connection and port information for the corresponding media line in the SDP offer *forwarded* by the ALG, and sub-case b does not apply.
d. The connection information for the media line in the SDP answer received by the ALG is the *unspecified address*. The visited-realm instance in the SDP answer matches the IP realm associated with the connection and port information for the corresponding media line in the SDP offer *received* by the ALG, and sub-cases b and c do not apply.
e. The connection information for the media line in the SDP answer received by the ALG is the *unspecified address*. The visited-realm instance in the SDP answer matches the IP realm associated with a secondary-realm instance previously inserted by the ALG in the forwarded SDP offer, and sub-cases b, c and d do not apply.
f. The connection information for the media line in the SDP answer received by the ALG is the *unspecified address*. Sub-cases b, c, d and e do not apply.

Note that after completing the processing for the appropriate sub-case, the ALG MAY release any BG resources no longer used by the resulting media path.

6.2.1. SDP answer sub-case a: valid connection information

In sub-case a, the ALG receives connection information for the media line in the SDP answer that corresponds to a valid IP address in its IP realm. The ALG behavior depends on which SDP offer case it selected when forwarding the SDP offer:

. In case 1, since the ALG bypassed its BG and at least one prior BG when forwarding the SDP offer, the ALG must forward an SDP

answer containing the unspecified address to signal that the ALG receiving the forwarded SDP answer controls a BG that is to be bypassed. The ALG SHALL construct a new visited-realm instance from the connection and port information for the media line in the incoming SDP answer, SHALL add this visited-realm instance to the SDP answer, replacing any other visited-realm instances that may appear in the SDP answer, SHALL replace the connection information for the media line in the SDP answer with the unspecified address, and SHALL forward the modified SDP answer.

- In case 2, since the ALG already bypassed its BG and no others in the SDP offer, it SHALL forward the received SDP answer with no changes.
- In case 3, since the ALG already bypassed at least one prior BG in the SDP offer, but did not bypass its own BG, the forwarded SDP answer must contain the unspecified address to signal that the ALG receiving the forwarded SDP answer controls a BG that is to be bypassed. The ALG SHALL construct a new visited-realm instance from the local connection and port information for the side of the BG directed towards the offerer, SHALL add this visited-realm instance to the SDP answer, SHALL replace the connection information for the media line in the SDP answer with the unspecified address, and SHALL forward the modified SDP answer.
- In case 4, since the ALG does not bypass any BGs, the ALG SHALL replace the connection and port information for the media line in the SDP answer with the local connection and port information for the side of its BG directed toward the offerer, and SHALL forward the modified SDP answer.

In addition, when the controlled BG remains allocated, as in cases 3 and 4 with sub-case a, if the ALG is not performing hosted NAPT traversal on the side towards the SDP answerer, the ALG SHALL use the connection and port information from the incoming SDP answer as the remote connection and port information for the side of the BG directed towards the answerer. If the ALG is performing hosted NAPT traversal on the side towards the SDP answerer, the ALG/BG MUST discover the IP address of the RG via latching or other unspecified technique.

6.2.2. SDP answer sub-case b: match on other IP realm

In sub-case b, the ALG receives an unspecified address in the connection information for the media line in the SDP answer. The visited-realm instance in the SDP answer matches a visited-realm or secondary-realm instance previously *received* by the ALG in the SDP offer. Regardless which case 1-4 the ALG previously applied to the

SDP offer, the ALG is not required to provide a BG for the media path. The ALG SHALL forward the SDP answer with no changes.

6.2.3.

SDP answer sub-case c: match on forwarded SDP offer

In sub-case c, the ALG receives an unspecified address in the connection information for the media line in the SDP answer. The visited-realm instance in the SDP answer matches the IP realm associated with the connection and port information for the corresponding media line in the SDP offer *forwarded* by the ALG, and sub-case b does not apply. The ALG behavior depends on which SDP offer case it selected when forwarding the SDP offer:

- Sub-case b applies exclusively to case 1.
- In case 2, since the ALG already bypassed its BG and no others in the SDP offer, the visited-realm instance in the received SDP answer also matches the IP realm associated with the connection and port information for the corresponding media line in the SDP offer *received* by the ALG. The ALG SHALL replace the connection and port information for the media line in the SDP answer with the connection and port information from the visited-realm instance in the received SDP answer, SHALL delete the visited-realm instance from the SDP answer, and SHALL forward the modified SDP answer.
- In case 3, since the ALG already bypassed at least one prior BG in the SDP offer, but did not bypass its own BG, the forwarded SDP answer must contain the unspecified address to signal that the ALG receiving the forwarded SDP answer controls a BG that is to be bypassed. The ALG SHALL replace the visited-realm instance for the media line in the SDP answer with a new visited-realm instance constructed from the local connection and port information for the side of the BG directed towards the offerer, SHALL retain the unspecified address in the connection information for the media line in the SDP answer, and SHALL forward the modified SDP answer.
- In case 4, since the ALG does not bypass any BGs, the ALG SHALL replace the connection and port information for the media line in the SDP answer with the local connection and port information for the side of its BG directed toward the offerer, SHALL delete the visited-realm instance from the SDP answer, and SHALL forward the modified SDP answer.

In addition, when the controlled BG remains allocated, as in cases 3 and 4 with sub-case c, the ALG SHALL use the connection and port information from the visited-realm instance in the received SDP answer as the remote connection and port information for the side of the BG directed towards the answerer.

6.2.4.

SDP answer sub-case d: match on received SDP offer

In sub-case d, the ALG receives an unspecified address in the connection information for the media line in the SDP answer. The visited-realm instance in the SDP answer matches the IP realm associated with the connection and port information for the corresponding media line in the SDP offer *received* by the ALG, and sub-cases b and c do not apply. The ALG bypasses its BG in all cases. The ALG behavior depends on which SDP offer case it selected when forwarding the SDP offer:

- Sub-case b applies exclusively to case 1.
- Either sub-case b or c applies to case 2.
- Sub-case b applies exclusively to case 3.
- In case 4, since the ALG did not bypass any BGs when processing the SDP offer, it must now signal the forwarded SDP answer to bypass its own BG. The ALG SHALL replace the connection and port information for the media line in the SDP answer with the connection and port information from the visited-realm instance for the media line in the received SDP answer, SHALL delete the visited-realm instance from the SDP answer, and SHALL forward the modified SDP answer.

6.2.5. SDP answer sub-case e: match on own secondary-realm

In sub-case e, the ALG receives the unspecified address in the connection information for the media line in the SDP answer. The visited-realm instance in the SDP answer matches a secondary-realm instance previously inserted by the ALG in the SDP offer, and sub-cases b, c and d do not apply. The ALG behavior depends on which SDP offer case it selected when forwarding the SDP offer:

- SDP offer cases 1 and 2 do not apply since the ALG does not insert secondary-realm instances into the SDP offer in these cases.
- In case 3, since the ALG already bypassed at least one prior BG in the SDP offer, but did not bypass its own BG, the forwarded SDP answer must contain the unspecified address to signal that the ALG receiving the forwarded SDP answer controls a BG that is to be bypassed. The ALG uses the BG associated with the secondary-realm instance rather than the original BG allocated for the forwarded SDP offer. The ALG SHALL construct a new visited-realm instance from the local connection and port information for the side of the secondary BG directed towards the offerer, SHALL add this visited-realm instance to the SDP answer, SHALL replace the connection information for the media line in the SDP answer with the unspecified address, and SHALL forward the modified SDP answer.
- In case 4, since the ALG does not bypass any BGs, the ALG SHALL

replace the connection and port information for the media line
in the SDP answer with the local connection and port

information for the side of its secondary BG directed toward the offerer, and SHALL forward the modified SDP answer.

In addition, since the secondary BG remains allocated for this sub-case, if the ALG is not performing hosted NAPT traversal on the side towards the SDP answerer, the ALG SHALL use the connection and port information from the incoming SDP answer as the remote connection and port information for the side of the BG directed towards the answerer. If the ALG is performing hosted NAPT traversal on the side towards the SDP answerer, the ALG/BG MUST discover the address of the RG via latching or other unspecified technique.

6.2.6. SDP answer sub-case f: no match

In sub-case f, the ALG receives an unspecified address in the connection information for the media line in the SDP answer, and sub-cases b, c, d and e do not apply. Since either there is no visited-realm instance or the instance does not match any of the listed cases, then either the unspecified address comes from the SDP answerer or the active-bypass option has been invoked by another ALG. In all cases 1-4, the ALG SHALL forward the SDP answer with no changes.

6.3. ALG procedures for Active-Bypass Option

During the processing of the SDP answer in the base algorithm, any ALG that still retains a BG in the media path (i.e., SDP answer sub-cases a, c or e with SDP offer cases 3 or 4) MAY choose to perform the active-bypass option as a candidate anchor ALG for an alternate media path segment. The candidate anchor ALG contacts the best candidate target ALG to mutually determine if a superior media path segment is available.

6.3.1. Anchor ALG sends an alternate path request

Each ALG handling one of SDP answer sub-cases a, c or e with SDP offer case 3 or 4 MAY examine the information within visited-realm and secondary-realm instances previously received in the SDP offer to determine if there is a possibility that a significantly "better" remaining path can be constructed than the one already determined by the base algorithm. In particular, the ALG examines the geo-location, delay and loss data from its BG back to the earliest ALG reachable via a GRUU to make this determination. The method of

using the information to identify better paths and the threshold of improvement required (given the extra signaling needed for the active-bypass option) is a matter of local policy.

For example, if the earliest ALG reachable via a GRUU controls a BG that is geographically close to the BG controlled by the determining ALG, yet there are other visited-realm or secondary-realm instances on the path between them that are geographically distant from them, then there is good reason to expect that a better media path segment exists.

If a possible "better" path exists for one or more SDP media lines to the same earlier ALG, the determining ALG (now called the anchor ALG) SHALL send a SIP INVITE request without SDP to the earlier ALG (now called the target ALG). This INVITE request is called an alternate path request. This alternate path request will, if successful, result in an alternate path dialog and one or more alternate media path segments, if they have not already been established by earlier alternate path requests. This is in contrast to the original dialog, for which the anchor ALG is still processing the SDP answer.

If an alternate path dialog associated with the original dialog already exists between the anchor and target ALGs, the alternate path request SHALL comprise a re-INVITE request within the existing alternate path dialog. This may occur, for example, if a previous SDP offer/answer transaction has already completed within the original dialog. Otherwise the alternate path request SHALL comprise a new INVITE request, placing the GRUU of the target ALG in the Request-URI and the GRUU of the anchor ALG in the From and P-Asserted-Identity headers.

According to normal IMS routing procedures, the alternate path request may traverse one or more ALGs on its path to the target ALG. If the alternate path request fails prematurely with any non-success final response, the anchor ALG SHOULD abort the active-bypass option and continue handling of the SDP answer within the original dialog according to the base algorithm.

6.3.2. Target ALG processing of alternate path request

Upon receipt of an alternate path request in a new INVITE request, the target ALG SHALL identify the corresponding original dialog via the unique value of the GRUU in the Request-URI. Upon receipt of an alternate path request in a re-INVITE request, the target ALG SHALL identify the associated alternate path dialog and its corresponding original dialog. The target ALG uniquely identifies either request as an alternate path request associated with the original dialog since the assigned GRUU is the only address for which the target ALG will establish a corresponding alternate path dialog.

For each SDP media line in the previously forwarded SDP offer within the original dialog for which SDP offer case 3 or 4 has been applied (i.e., the target ALG has allocated a BG for the media line), the target ALG SHALL determine the IP realm associated with the alternate path request. Then for each applicable media line, the target ALG SHALL determine whether the BG resource(s) allocated during the processing of the SDP offer for the original dialog has access to the IP realm associated with the alternate path request. If so, then the BG resource can be re-used, else the target ALG MUST allocate a new BG resource.

Then the target ALG SHALL construct a new SDP offer from the SDP offer forwarded within the original dialog by:

1. copying the original SDP offer;
2. modifying the o line as appropriate;
3. deleting all visited-realm and secondary-realm instances;
4. constructing the visited-realm information for each applicable media line;
5. inserting the corresponding connection and visited-realm instance information for each applicable media line and
6. setting port value to zero for all other media lines.

For each applicable media line in the new SDP offer, if BG resources are available with access to additional IP realms as well as access to the IP realm previously selected for the portion of the bearer path towards the original SDP offerer, the target ALG MAY construct the corresponding secondary-realm instances and add them to the media line.

Then the target ALG SHALL send the constructed SDP offer to the anchor ALG in the SIP 200 OK response message according to normal SIP procedures. If the alternate path request received by the target ALG traversed one or more ALGs on its path from the anchor ALG, this new SDP offer will also traverse the same ALGs, which will recursively apply the base algorithm and optionally the active-bypass option to the SDP offer.

If an error such as any of the following occurs during the processing of the alternate path request, the target ALG responds with an appropriate SIP final error response:

- The target ALG does not recognize the GRUU.
- There are no BG resources allocated for any media line in the original SDP offer.
- The INVITE request included SDP.

6.3.3.

Anchor ALG processing of SDP offer from Target ALG

When the anchor ALG receives the SDP offer from the target ALG in the 200 OK response, the anchor ALG SHALL apply the following procedure independently to each media line in the received SDP offer before returning the corresponding SDP offer in the ACK request towards the target ALG.

If the port value is set to zero in the media line, the anchor ALG SHALL set the port value to zero in the corresponding media line in the SDP answer to be sent towards the target ALG and SHALL proceed with the base algorithm (i.e., the active-bypass option has no impact on the base algorithm for this media line).

If the media line has a non-zero port value, then the anchor ALG SHALL attempt to identify the corresponding media line in the original SDP answer. There is a possibility that the order of the media lines in the received SDP offer is different from the order of the media lines in the original SDP answer due to intermediate applications performing 3rd party call control procedures to split/merge SDP media lines. If there is a visited-realm or secondary-realm instance in the received SDP offer with a GRUU for the target ALG, then this can be matched against the GRUU received for the target ALG in the original SDP offer to identify the corresponding media line. If no GRUU is present to assist in matching media lines, the anchor ALG may be able to uniquely match the media lines based on other information, e.g., only one applicable media line is common to both the original and alternate path dialogs.

If the anchor ALG cannot identify the corresponding original media line for a received media line with a non-zero port value, the anchor ALG SHALL set the port value to zero in the corresponding media line in the SDP answer to be sent towards the target ALG.

If the anchor ALG can identify the corresponding original media line for a received media line with a non-zero port value, the anchor ALG SHOULD use available visited-realm and secondary-realm instance information in the received SDP offer and MAY use other unspecified data to determine if the alternate media path segment is significantly "better" than the corresponding portion of the original media path. The algorithm used to assess the quality of each media path segment and to determine the minimum threshold of significance is a matter of local policy.

If the anchor ALG determines that the alternate media path segment is not significantly better than the corresponding portion of the original media path, the anchor ALG SHALL set the port value to zero in the corresponding media line in the SDP answer to be sent towards the target ALG and SHALL proceed with the base algorithm.

If the anchor ALG determines that the alternate media path segment is significantly better than the corresponding portion of the original media path, the anchor ALG:

1. SHALL allocate BG resources for the IP realm associated with the alternate media path segment, if not already available;
2. SHALL set the connection information and/or visited-realm attribute for the corresponding media line in the SDP answer in the alternate path dialog according to the recursive application of the base algorithm by choosing SDP offer case 3 or 4 according to the processing of the received media line from the alternate path dialog and by applying SDP answer sub-case a, c or e from the processing of the original SDP answer; and
3. SHALL modify the processing of the original SDP answer in the base algorithm as follows.

For the corresponding media line of the SDP answer received during the course of the base algorithm, the anchor ALG:

1. SHALL select the remote connection and port information for the side of the BG directed towards the answerer according to the SDP offer case applied to the media line in the alternate path dialog and the applicable original SDP answer sub-case;
2. SHALL delete any visited-realm instance for the media line in the SDP answer;
3. SHALL construct a new visited-realm instance for the special IP realm "NOMATCH" including the GRUU of the media line received from the target ALG, if available;
4. SHALL add this visited-realm instance to the SDP answer;
5. SHALL replace the connection information for the media line in the SDP answer with the unspecified address; and
6. SHALL forward the modified SDP answer within the original dialog.

6.3.4. Other ALG processing of SDP answer in original dialog

After the anchor ALG forwards the original SDP answer, every other conformant ALG on the signaling path prior to the target ALG will forward the SDP answer without change according to SDP answer sub-case f of the base algorithm.

6.3.5. Target ALG processing of SDP answers

Upon receipt of the SDP answer within the original dialog, recognizing that it has recently received and responded to an

alternate path request for this media line (and possibly others), the target ALG:

1. SHALL determine if SDP answer sub-case f applies with special IP realm "NOMATCH" in the corresponding visited-realm attribute (if one is present);
2. SHALL verify that the corresponding media line for the alternate path dialog is to be associated with this original media line, using either the GRUU in the received visited-realm attribute or other unspecified means;
3. SHALL determine if the SDP answer for the alternate path dialog is received (in the ACK request) in a reasonable amount of time;
4. SHALL determine if the port for the corresponding media line for the alternate path dialog has non-zero value and
5. SHALL determine that SDP answer sub-case a, c or e applies to the corresponding media line for the alternate path dialog.

If any of the above conditions do not apply, then the target ALG SHOULD continue with the normal processing of the base algorithm and mark the media line for the alternate path request as "unused". Note that some combinations of conditions (representing error cases) will fail to establish an end-to-end media path. If this occurs, the target ALG SHOULD reject subsequent alternate path requests within the original dialog and MAY apply other unspecified recovery actions.

If all of the above conditions apply, the target ALG SHALL apply the applicable SDP offer case 3 or 4 and the applicable SDP answer sub-case a, c or e for the corresponding media line for the alternate path dialog to configure the BG and modify the received SDP answer for the original dialog before forwarding the SDP answer.

The net result of the successful application of the procedure in sections 6.3.1 through 6.3.5 is to replace the portion of the end-to-end media path generated by the base algorithm between the target and anchor ALGs with the alternate media path segment generated by the alternate path request.

6.3.6. Release of alternate path dialog

The target ALG and anchor ALG SHOULD release the alternate path dialog and associated resources not otherwise needed using standard SIP procedures when either the original dialog is released or when all of the media lines for the alternate path dialog either have port value zero or are marked "unused".

If the alternate path dialog is released while in use to maintain an alternate media path segment, the anchor ALG and target ALG MAY release the corresponding original dialog or perform other

unspecified recovery actions.

6.4. Special handling of unspecified address from endpoints

If the UA initiating an SDP offer includes an unspecified address in the connection information, the unspecified address SHALL be associated with the IP realm of the UA. The ALG SHALL follow case 1 when forwarding an SDP offer with an unspecified address, where it is understood that the SDP offer contains an implicit visited-realm instance with the unspecified address for every IP realm. The net result of this procedure is that if there is an unspecified address in the initial SDP offer, every ALG will forward an unspecified address. If the received SDP answer includes a valid IP address, it will be transformed into an unspecified address by the first ALG using sub-case a, and subsequent ALGs will include the unspecified address in the forwarded SDP answer using a sub-case b through f. Since this procedure does not support the use of a "black hole" address [16] to discover the connection information for the answering UA, there are some limitations to the applicability of these procedures, although none of the recommended 3pcc procedures [16] depend on the use of the "black hole" address.

If the UA initiating an SDP answer includes an unspecified address in the connection information, the ALG procedures for handling of SDP answers remain unchanged, with the result that if any BGs were allocated when forwarding SDP offers, they will all be released. Each ALG SHALL treat an SDP answer with an unspecified address but without an explicit visited-realm instance as if it contains a single implicit visited-realm instance for an unknown IP realm. Thus sub-case f always applies.

Note that if the initial SDP offer or initial SDP answer includes an unspecified address in the connection information, there can be no media flow until a subsequent SDP offer/answer transaction is performed using actual IP addresses from the endpoint IP realms.

6.5. Assumptions about non-compliant ALGs

A non-compliant ALG will usually delete unknown SDP attributes before forwarding SDP offers or answers. Such an ALG will delete any visited-realm or secondary-realm instances from the SDP offer before allocating a BG and forwarding the SDP offer, making it impossible for subsequent ALGs to bypass the allocated BG. Optimizations can still be applied independently to the portions of the end-to-end media path before and after the non-compliant ALG to successfully establish the end-to-end media path via the BG allocated by the non-compliant ALG.

If a non-compliant ALG in a session signaling path does forward visited-realm and secondary-realm attributes after BG allocation,

compliant ALGs retain most opportunities for BG bypass while establishing the end-to-end media path if the non-compliant ALG exhibits the following behaviors:

. When receiving an SDP message with an unspecified address in the connection information, the non-compliant ALG retains the unspecified address in the forwarded SDP message. If the ALG both converts an unspecified address into a valid address and forwards visited-realm attributes, then the procedures may fail to establish a media path. The ALGs bordering a non-compliant ALG known to do this MAY implement a work-around by manipulating the signaling to keep the non-compliant ALG in the media path, although this forfeits significant opportunities for BG bypass.

  To keep a neighbor ALG in the path, a compliant ALG selects an applicable case or sub-case from the detailed procedures that ensures that real connection information is provided in all SDP messages destined to the neighbor ALG and to delete all visited-realm attributes in SDP messages destined to or coming from the neighbor ALG.

. A non-compliant ALG will not terminate a session for which there is no media flow in its BG. The ALG must implicitly accept that its BG may be bypassed.

  The ALGs bordering a non-compliant ALG that is known to violate this assumption MAY implement a work-around by manipulating the signaling to keep the non-compliant ALG in the media path, although this forfeits significant opportunities for BG bypass.

6.6. Operation in the presence of forking

Forking has no impact on the processing of SDP offers according to either the base algorithm or the active-bypass option.

When an ALG forwards an INVITE request including an SDP offer that is subsequently forked, an ALG may receive multiple SDP answers associated with the SDP offer in the early dialog state, where each SDP answer is within a separate early dialog. The ALG SHALL apply the base algorithm and the active-bypass option separately to each SDP answer associated with a forked branch. The ALG MUST retain any resources reserved during the handling of the SDP offer until a dialog is fully established or until it can receive no other forked SDP answers. If the ALG allocates a BG resource that is shared by multiple media paths created for parallel-forked dialogs, the ALG MAY apply local policy to selectively filter the media streams associated with the forked endpoints according to the gateway model of RFC 3960 [22] or RFC 5009 [23].

7. The visited-realm and secondary-realm attributes

The visited-realm and secondary-realm SDP attributes are media-level attributes only.

The visited-realm attribute contains an IP realm identifier and transport address for a previously visited realm that can potentially be used to bypass allocated BGs.

The secondary-realm attribute contains an IP realm identifier and transport address for a secondary realm that can potentially be used to bypass allocated BGs.

The syntax of these attributes is defined using Augmented BNF as defined in RFC 4234 [6]:

```
   visited-realm         = "visited-realm" ":" realm-number SP
                           realm SP
                           nettype SP                ;from RFC 4566
                           addrtype SP               ;from RFC 4566
                           connection-address SP     ;from RFC 4566
                           port                      ;from RFC 4566
                           [SP rtcp-port [SP rtcp-address]]
                           [SP coordinates]
                           [SP delay]
                           [SP loss]
                           [SP temp-gruu]
                           [SP credentials]
                           *(SP extension-att-name SP
                                extension-att-value)

   secondary-realm       = "secondary-realm" ":" realm-number SP
                           realm SP
                           nettype SP                ;from RFC 4566
                           addrtype SP               ;from RFC 4566
                           connection-address SP     ;from RFC 4566
                           port                      ;from RFC 4566
                           [SP rtcp-port [SP rtcp-address]]
                           [SP coordinates]
                           [SP delay]
                           [SP loss]
                           [SP temp-gruu]
                           [SP credentials]
                           *(SP extension-att-name SP
                                extension-att-value)

   realm-number          = 1*DIGIT
   realm                 = non-ws-string           ;from RFC 4566
   rtcp-port             = "rtcp-port" SP port
   rtcp-address          = "rtcp-address" SP connection-address
   coordinates           = "coordinates" SP latitude "," longitude
   latitude              = [ "-" ] 1*2DIGIT [ "." *DIGIT ]
   longitude             = [ "-" ] 1*3DIGIT [ "." *DIGIT ]
   delay                 = "delay" SP delay-value
   delay-value           = 1*DIGIT
   loss                  = "loss" SP loss-value
   loss-value            = "-" 1*DIGIT ["." 1*DIGIT]
   temp-gruu             = "temp-gruu" SP SIP-URI    ;from RFC 3261
   credentials           = "credentials" SP credentials-value
   credentials-value     = non-ws-string
   extension-att-name    = token
   extension-att-value   = non-ws-string
```

This grammar encodes the primary information about each visited-realm and secondary-realm instance: the sequence in which the realm was visited, the realm identity, its IP address and port, and optional geo-location, IP packet delay, IP packet loss, temporary-GRUU and security credentials:

<realm-number>: For a visited-realm instance, realm-number is a positive decimal integer between 1 and 256 which identifies the sequence in which this visited-realm instance was visited during the forwarding of an SDP offer, compared to other visited-realm instances for the media line in the same SDP offer. It MUST start at 1 and MUST increment by 1 compared to the highest existing realm-number for the media line when inserting a new visited-realm instance into an SDP offer. The realm-number can be ignored in an SDP answer since there should only be one visited-realm instance and no secondary-realm instance in an SDP answer. It is RECOMMENDED that the realm-number have value 1 in an SDP answer. For a secondary-realm instance in a forwarded SDP offer, realm-number MUST have the same value as the realm-number for the visited-realm instance created for the same media line by the same ALG for the connection information in the forwarded SDP offer.

<realm>: identifies a set of mutually reachable IP endpoints that share a common IP addressing scheme. Each realm also defines a protection domain for all hosts using visited-realm or secondary-realm attribute instances for the realm, to help ensure the integrity of the remaining information in each attribute instance. A public address reachable from the open internet MAY be associated with the special realm "IN", for which no credentials are required. The special realm "NOMATCH" is used to signify a realm only reachable via an alternate media path segment created by the active-bypass option. Operators of ALGs that wish to ensure the integrity of the visited-realm instance information for their realm(s) MUST adhere to the following guidelines for creation of a realm string for their servers: 1) Realm strings MUST be globally unique. It is RECOMMENDED that a realm string contain a hostname or domain name, following the recommendation in Section 3.2.1 of RFC 2617 [10]. 2) Realm strings SHOULD present a human-readable identifier that can be rendered to a user.

<nettype>, <addrtype> and <connection-address>: are taken from the connection-field (c= line) of RFC 4566 [7]. They describe the IP address associated with the visited-realm instance, allowing for IPv4 addresses, IPv6 addresses and

FQDNs. An IP address SHOULD be used, but an FQDN MAY be

used in place of an IP address. When receiving an offer or answer containing an FQDN in an a=visited-realm attribute, if there is a match on the realm according to the procedures herein, the FQDN is looked up in the DNS using an A or AAAA record, and the resulting IP address is used for the remainder of the procedure.

<port>: is also taken from RFC 4566 [7]. It is the port associated with the visited-realm instance. Its meaning depends on the network being used for the connection-address, and on the transport protocol selected for the corresponding media line, e.g., UDP or TCP.

<rtcp-port> and <rtcp-address>: taken together are semantically equivalent to the rtcp attribute defined in RFC 3605 [5]. They optionally encode the RTCP port and address information when the visited-realm instance is for an RTP stream and the RTCP port number is not exactly one greater than the port for the RTP stream at the same address.

<coordinates>: provides the approximate geographic coordinates (geo-location) of the BG or endpoint associated with the connection information in the visited-realm or secondary-realm attribute. The "latitude" component MUST contain the decimal latitude of the identified location in the reference system WGS 84 [24]. The "longitude" component MUST contain the decimal longitude of the identified location in the reference system WGS 84 [24]. The number of decimal places indicates the precision of the value. The coordinates need only be accurate enough to estimate the minimum IP packet propogation delay between successive BGs/endpoints based on distance. The ALG SHOULD include known coordinates for each visited-realm or secondary-realm attribute in a forwarded SDP offer. The procedures in this document do not require the use of coordinates in SDP answers.

<delay-value>: is an estimate of the delay in transporting IP packets between the controlled BG and the next BG or endpoint towards the SDP offerer (through the previous IP realm). delay-value is a positive decimal integer representing the delay in milliseconds. The ALG SHOULD include delay-value for each visited-realm or secondary-realm attribute in a forwarded SDP offer if the information is available and is significantly different from an estimated minimum value based on the coordinates of the respective BGs/endpoints. The procedures in this document do not require the use of delay-value in SDP answers.

<loss-value>: is an estimate of the rate of IP packet loss on the link between the controlled BG and the next BG or endpoint towards the SDP offerer. loss-value is equal to log(packet-loss-rate) in negative decimal format, where packet-loss-rate is the average ratio of lost IP packets to all IP packets sent on the link. The packet-loss-rate can be reconstructed as 10**(loss-value). The ALG SHOULD include loss-value for each visited-realm or secondary-realm attribute in a forwarded SDP offer if the information is available. The procedures in this document do not require the use of loss-value in SDP answers.

<temp-gruu>: is a temporary GRUU assigned uniquely by each ALG for a specific dialog and media line. Draft-ietf-sip-gruu [8] defines the format of the temporary GRUU. For each media line in a forwarded SDP offer, if the ALG supports the target ALG procedures of the active-bypass option, is reachable via a globally unique host name, and controls the BG associated with the connection information for the media line in the forwarded SDP offer, the ALG SHOULD include a temp-gruu in the corresponding visited-realm attribute generated by the ALG. See the active-bypass option procedures for use of the temp-gruu in an SDP answer. The procedures in this document do not require the use of temp-gruu in the secondary-realm attribute.

<credentials-value>: is a digital signature computed on the other contents of the attribute and other secret data. The authority for the protection domain associated with the realm MAY choose MD5 [9] or other algorithm to compute the credentials. For additional security, extension attributes (such as nonce and opaque used for digest [10]) MAY be used to link the credentials calculated on the attribute in one SDP message to prior SDP offers or answers used within a SIP dialog. Only servers within the protection domain need to verify the integrity of the attribute contents.

The candidate attribute can itself be extended. The grammar allows for new name/value pairs to be added at the end of the attribute. An implementation MUST ignore any name/value pairs it does not understand.

Since the connection and port information in an instance of the visited-realm attribute can only be used by a trusted node within the corresponding IP realm, the realm MAY choose to put encrypted versions of the connection-address and port information into the extension parameters while putting dummy values into the connection-

address and port fields.

8. Security Considerations

The use of this extension is only applicable inside a "Trust Domain" as defined in RFC 3325 [4]. Nodes in such a Trust Domain are explicitly trusted by its users and end-systems to inspect and manipulate SDP messages as necessary to traverse and/or bypass firewalls and NATS while limiting access from unauthorized sources to endpoints in IP realms associated with the Trust Domain.

Since the procedures in this document include an option to cryptographically certify the candidate connection and port information from each IP realm, they can be used under some circumstances when the signaling traverses non-trusted networks or the Internet at large.

Since the base algorithm in this extension requires no additional signaling outside of an end-to-end SDP offer/answer exchange, it is likely to be impacted by any attack that can modify or disrupt an SDP offer/answer exchange. Such an attack could direct media to a target of a DoS attack, insert a third party into the media stream, and so on. These are similar to the general security considerations for offer/answer exchanges, and the security considerations in RFC 3264 [3] apply. These require techniques for message integrity and encryption for offers and answers, which can be satisfied by the SIPS mechanism [2] or IMS security mechanisms when SIP is used. As such, the usage of hop-by-hop message integrity and encryption with this extension is RECOMMENDED.

In addition to the above considerations, the active-bypass option in this extension establishes alternate path dialogs and alternate media path segments using GRUUs with values that cannot always be certified. Thus the active-bypass option is NOT RECOMMENDED for signaling that traverses non-trusted networks or the Internet at large.

This extension is not consistent with end-to-end security procedures that are otherwise recommended for SDP messages.

9. IANA Considerations

This specification registers two new SDP attributes per the procedures of Section 8.2.4 of [7]. The required information for the registration is included here.

9.1. visited-realm Attribute

Contact Name: Richard Ejzak, ejzak@alcatel-lucent.com

Attribute Name: visited-realm

Long Form: visited-realm

Type of Attribute: media level

Charset Considerations: The attribute is not subject to the charset attribute.

Purpose: This attribute is used in private networks employing border gateways to identify configurations in which IP realms are re-entered when establishing an end-to-end multimedia session, so that border gateways can be bypassed without compromising their role in securing access to the networks. The attribute provides a means to identify connection information for visited IP realms to help select the most optimal available path.

Appropriate Values: See Section 7 of RFC XXXX [Note to RFC-ed: please replace XXXX with the RFC number of this specification].

9.2. secondary-realm Attribute

Contact Name: Richard Ejzak, ejzak@alcatel-lucent.com

Attribute Name: secondary-realm

Long Form: secondary-realm

Type of Attribute: media level

Charset Considerations: The attribute is not subject to the charset attribute.

Purpose: This attribute is used in private networks employing border gateways to identify configurations in which secondary IP realms are available to establish an end-to-end multimedia session, so that border gateways can be bypassed without compromising their role in securing access to the networks. The attribute provides a means to identify connection information for secondary IP realms to help

select the most optimal available path.

Appropriate Values: See Section 7 of RFC XXXX [Note to RFC-ed: please replace XXXX with the RFC number of this specification].

10. References

10.1. Normative References

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

[2] 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.

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

[4] Jennings, C., Peterson, J. and Watson, M., "Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Network", RFC 3325, November 2002.

[5] Huitema, C., "Real Time Control Protocol (RTCP) attribute in Session Description Protocol (SDP)", RFC 3605, October 2003.

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

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

[8] Rosenberg, J., "Obtaining and Using Globally Routable User Agent (UA) URIs (GRUU) in the Session Initiation Protocol (SIP)", draft-ietf-sip-gruu-15 (RFC editor's queue), October 2007.

10.2. Informative References

[9] Rivest, R., "The MD5 Message-Digest Algorithm", RFC 1321, April 1992.

[10] Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A. and L. Stewart, "HTTP authentication: Basic and Digest Access Authentication", RFC 2617, June 1999.

[11] Borella, M., Lo, J., Grabelsky, D., and G. Montenegro, "Realm Specific IP: Framework", RFC 3102, October 2001.

[12] Borella, M., Grabelsky, D., Lo, J., and K. Taniguchi, "Realm Specific IP: Protocol Specification", RFC 3103, October 2001.

[13] Senie, D., "Network Address Translator (NAT)-Friendly Application Design Guidelines", RFC 3235, January 2002.

[14] Srisuresh, P., Kuthan, J., Rosenberg, J., Molitor, A., and A. Rayhan, "Middlebox communication architecture and framework",

RFC 3303, August 2002.

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

[16] Rosenberg, J., Peterson, J., Schulzrinne, H., and G. Camarillo, "Best Current Practices for Third Party Call Control (3pcc) in the Session Initiation Protocol (SIP)", BCP 85, RFC 3725, April 2004.
[17] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", draft-ietf-mmusic-ice-19 (RFC editor's queue), October 2007.
[18] Rosenberg, J., Mahy, R., Matthews, P. and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.
[19] Rosenberg, J., Mahy, R., and P. Matthews, "Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)", draft-ietf-behave-turn-12 (work in progress), November 2008.
[20] 3GPP "TS 23.228: IP Multimedia Subsystem (IMS); Stage 2 (Release 8)", 3GPP 23.228, September 2008, ftp://ftp.3gpp.org/specs/archive/23_series/23.228/.
[21] 3GPP "TS 24.229: IP Multimedia Call Control Protocol based on SIP and SDP; Stage 3 (Release 8)", 3GPP 24.229, September 2008, ftp://ftp.3gpp.org/specs/archive/24_series/24.229/.
[22] Camarillo, G. and H. Schulzrinne, "Early Media and Ringing Tone Generation in the Session Initiation Protocol (SIP)", RFC 3960, December 2004.
[23] Ejzak, R., "Private Header (P-Header) Extension to the Session Initiation Protocol (SIP) for Authorization of Early Media", RFC 5009, September 2007.
[24] National Imagery and Mapping Agency, "Department of Defense World Geodetic System 1984, Third Edition", NIMA TR8350.2, January 2000.

Any 3GPP document can be downloaded from the 3GPP webserver, http://www.3gpp.org/. See specifications.

Author's Address

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

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

Copyright Notice

This Internet-Draft expires June 17, 2009.