

speermint
Internet-Draft
Intended status: Standards Track
Expires: May 17, 2007

B. Rosen
NeuStar
November 13, 2006

Best Current Practices for Session Peering on the Internet by Carriers
through Federationsr
draft-rosen-speermint-peeringbcp-v1-00.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of 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/ietf/lid-abstracts.txt>.

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

This Internet-Draft will expire on May 17, 2007.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This memo defines a first version Best Current Practice for peering among a federation of voice or other multimedia service providers

Internet-Draft

BCP for Peering Through Federations

November 2006

Table of Contents

1.	Terminology	3
2.	Overview	3
2.1.	Addressing and Routing	4
2.2.	Connectivity	5
2.3.	Security (Accountability)	5
2.4.	Quality	5
2.5.	Protocol Mediation	6
2.6.	Model	6
3.	Responsibilities of Federations	6
3.1.	Federation Static Policies	6
3.1.1.	Membership	7
3.1.2.	Identity	7
3.1.3.	Media Exchange	7
3.1.4.	Capacity Controls	7
3.1.5.	Protocol Specification	7
3.1.6.	CODEC choices	8
3.1.7.	Billing	8
3.2.	Layer 3 interconnection	8
3.3.	Layer 5 interconnection	8
3.4.	Routing	8
3.5.	NAT Traversal	9
3.6.	Transcode	9
3.7.	Capacity Controls	9
3.8.	Protocol Mediation	9
4.	Responsibilities of peers	9
4.1.	Conformance to Policies	9
4.2.	Identity	10
4.3.	Transcode	10
5.	Security Considerations	10
6.	Normative References	10
	Author's Address	11
	Intellectual Property and Copyright Statements	12

1. Terminology

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 [[RFC2119](#)].

It is assumed that the reader is familiar with the terminology and acronyms defined in [[RFC3261](#)]

2. Overview

The SIP standard [[RFC3261](#)] models much of its protocol operations on familiar Internet applications such as email and the web. For example, SIP clients contact remote domains by resolving SIP URIs which are, like email addresses, composed of a username and domain name portion. As is the case with email, a client need not have any pre-association with a remote domain in order to initiate a session with a user in that domain. The security of SIP is designed to counter the sorts of threats that arise on the Internet - threats based on eavesdropping, packet spoofing, impersonation of identity, and the like. Endpoints must assume responsibility for most of these security functions. Like email and the web, SIP assumes that the Internet end-to-end model applies: that is, that entities using SIP have unimpeded connectivity to one another.

There are however operational models in which the assumptions of traditional Internet applications do not hold up. The end-to-end reachability of user agents is commonly obstructed by network-layer impediments like network address translators (NATs) or firewalls. Many SIP user agents, and SIP deployments, utilize telephone numbers rather than email-like SIP URIs, which introduces requirements for a new resolution process in the routing of requests. Some SIP service providers are also uncomfortable leaving the management of security to user agents, for a variety of reasons.

To a voice service provider that leverages SIP in a commercial offering, these concerns are all basic issues of usability. Most SIP user agents today have dial pads, and users are accustomed to the use of telephone numbers of voice applications. Without a transparent and automatic NAT traversal function such as ICE (which in turn relies on STUN and TURN services), SIP calls on the public Internet may face serious problems establishing media paths. Users similarly cannot be burdened with understanding and supporting security functions.

Solving these concerns also motivates voice service providers to establish formal peering relationships with one another. Peering

represents an agreement between parties to permit the exchange of traffic, generally in accordance with some pre-established policy. Without agreement between two VSP domains, for example, on how telephone numbers are resolved, it is impractical for users in different domains to call one another. Security, and in particular authorization, is perhaps the most fundamental reason for peering. The email model, while quite successful, has a very widespread problem with undesirable traffic (namely spam), and a comparable problem for SIP could be quite harmful to commercial offerings. Peering agreements would allow providers to trace accountable sources of undesired traffic and to make appropriate authorization decisions based on traffic sources.

Peering at the session layer can occur on a bilateral basis or a multilateral basis, where the latter generally takes the form of a federation (typically in an "assisted" peering configuration). At lower layers of the Internet architecture, there are also various forms of bilateral and multilateral connections which are established between Internet service providers; the more service providers require interconnection, however, the less attractive bilateral connections become, if only because the cost of constructing physical links between networks becomes unwieldy. While it might seem that there is no similar difficulty with the establishment of bilateral peering at the session layer, there are a number of reasons why voice service providers might want to minimize the number of connections they establish to peer networks: for example, to reduce load on gateways (SBCs and IPSec gateways), or to simplify authorization or routing decisions by delegating that responsibility to network elements operated by the federation.

Moreover, if there are media-based applications which need to be made available to the federation as a whole, a point of lower layer interconnection, such as a traditional layer 3 interexchange point, is an ideal place to stage them. Any such applications like transcoders and media relays would best be situated in a layer 3 point of interconnection.

The following sections detail some of the functions that might be performed at a peering point and briefly explain how they benefit from being deployed in a federation environment.

[2.1.](#) Addressing and Routing

Several forms of private directories are useful in a peering context. Aside from the widely-attested need to translate telephone numbers into an identifier that can be routed on the Internet (typically via ENUM or an ENUM-like mechanisms), there is a further need in a peering environment to manage points of egress and ingress on the

networks of peers. While this can take the form of a conventional DNS lookup, such a lookup could return IP addresses that are only routable within a peering point, thereby restricting their access to the federation.

[2.2.](#) Connectivity

Most NAT traversal schemes, such as TURN, require the availability of a relay that is reachable by both parties in a call. A peering point is a natural place to stage such a relay precisely because its position in the network is unlikely to introduce additional latency in the media by lengthening the call path.

[2.3.](#) Security (Accountability)

SIP can be used in constrained environments, effectively closed IP networks, where the threats that are quite plausible on the public Internet become very unlikely - especially environments that are based on traditional telephone networks. In those closed environments the use of the baseline SIP security mechanisms may seem very unattractive. Some form of transitive trust is typically viewed as sufficient in this sort of environment. The use of network-layer

security gateways that connect individual networks to a closed peering point VLAN would be one example of how this might operate.

Certain application-layer functions can assist with the establishment of transitive trust and the management of service provider authorization based on that trust. For example, mechanisms like [RFC3325](#) or [RFC4474](#), which provide assurance of the identity of the originator of a SIP request, can be performed by a SIP proxy server resident in the peering point. Note that especially when telephone number translations are centrally managed by the federation, providing identity functions for Caller-ID typical must also be managed by the federation.

[2.4.](#) Quality

The use of protocols to establish quality of service across a traffic path in an IP network is quite controversial, especially when tied to a real-time application like voice over IP. Traffic engineering, management of quality across a particular link, is more common and generally less complex than resource reservation on a per-call basis. Moreover, the extremely large deployments of certain VoIP applications on the public Internet which lack any per-call resource reservation have created a great deal of skepticism about the need to incur any significant expense to assure quality.

Layer 3 peering points are general points of optimal quality in the

Rosen

Expires May 17, 2007

[Page 5]

Internet-Draft

BCP for Peering Through Federations

November 2006

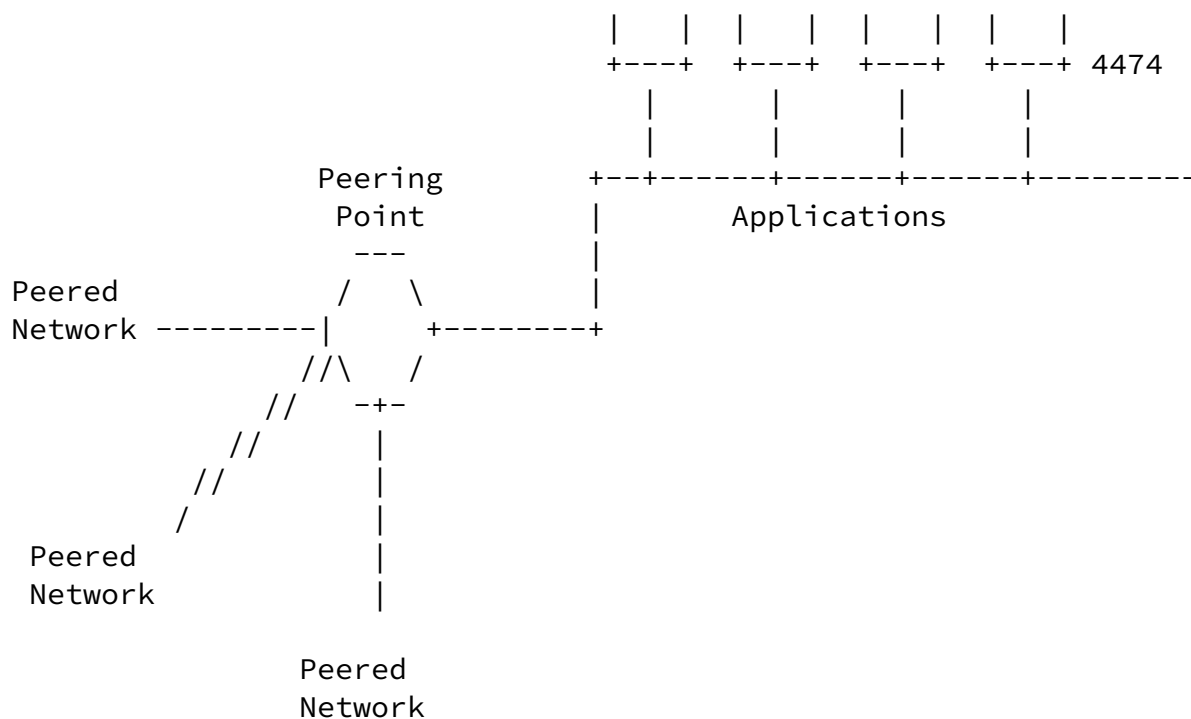
network from a latency and bandwidth perspective, and accordingly they are likely to be the best place to stage any network-based mechanism which would help to assure call quality.

[2.5.](#) Protocol Mediation

[[unsure if we'd want to talk about different SIP "variants" and the normalization of SIP signaling, this is just a placeholder entry]]

[2.6.](#) Model

Directories	NAT	SIP
+---+	+---+	+---+ 3325



[3. Responsibilities of Federations](#)

[3.1. Federation Static Policies](#)

Federations must establish explicit policies on at least the following matters. Such policies may be in the form of a contract or other agreement between the federation and peers, a web page, or other prominent mechanism. In some cases, the federation MAY explicitly permit or prohibit parts of the policy described to be a matter for bilateral agreements within the federation. In this version of the BCP, we provide no standardized way to express this form of policy.

[3.1.1. Membership](#)

The Federation MUST specify who is allowed to peer at the federation, and how the peers are made known to one another. The policy MUST include a statement of whether indirect (i.e. transit) peers are permitted.

[3.1.2. Identity](#)

The Federation MUST specify the requirements on peers to identify users. The Federation MAY permit [[RFC3325](#)] asserted identity. The Federation SHOULD permit [[RFC4474](#)] Identity. Federations supporting [RFC4474](#) MUST specify the CA(s) permitted to issue certificates of the authentication service (which MAY be operated by the Federation). The Federation policy MUST specify the date maximum discrepancy period, The policy MUST specify what is permitted in the display name of the From: header, and what mechanisms peers must have to control such content.

[3.1.3.](#) Media Exchange

The Federation MUST specify mechanisms for the interchange of media among the peers. This MUST include mechanisms for NAT traversal.

[3.1.4.](#) Capacity Controls

The Federation MUST specify policy to control the traffic sent to and received by peers. The policy SHOULD include limits on the maximum number of active calls, maximum number of calls/messages per specified unit time and the aggregate media bandwidth. Specifications for both aggregate traffic to/from the federation as well as limits between two peers MAY be specified.

The Federation MAY specify policy for individual ingress/egress elements as well as total traffic to/from a peer.

Federations MAY permit peers to specify and/or form bilateral agreements on the limits. This version of the BCP does not specify mechanisms for dynamic discovery or modification of such policies.

[3.1.5.](#) Protocol Specification

The Federation MUST specify details of the (SIP) signaling messages that peers must conform to. Such specification SHOULD include minimal extensions that MUST be supported, and what options MUST be supported, MUST NOT be supported or MAY be supported. This specification SHOULD include a description of the services that are expected to be supported across the Federation, and the signaling

[3.1.6.](#) CODEC choices

The federation defines a codec policy to which all peers must adhere which would include designation of one or more mandatory to deploy codec and/or local transcode capability for each supported media type (audio, video, text) so that all calls can successfully complete offer/answer exchange. Consideration should be given in specifying mandatory-to-deploy codecs to include at least one that has minimal degradation of signal fidelity when two transcodes are required to achieve actual end to end compatibility.

[3.1.7.](#) Billing

The Federation MUST specify what charging mechanisms for the exchange of traffic it permits, and any support for such practices (e.g. CDR production) it provides. Where the Federation provides explicit billing arrangements, such arrangements must be described, including currency choices.

[3.2.](#) Layer 3 interconnection

The federation MUST specify and/or provide the mechanism by which peers exchange packets at the TCP/IP layer. This may involve addressing issues (if not using public IP addresses), and VPN or other tunneling mechanisms. The federation MUST detail the processes by which peers establish, test and maintain their TCP/IP connections.

[3.3.](#) Layer 5 interconnection

The Federation MUST provide a mechanism for discovery and addressability of multiple ingress elements (proxy servers, SBCs or B2BUAs) from multiple egress elements to allow exchange of signaling between the peers. Where multiple ingress elements are permitted, the Federation must specify how origination peers select one of the ingress elements, and how termination peers may control such selection mechanisms. This version of the BCP does not define automatic load sharing or overload recovery mechanisms.

[3.4.](#) Routing

The Federation MUST specify the mechanism by which peers discover routing information for the exchange of traffic. Routing mechanisms MUST permit any peer to discover how to route to any other peer's subscribers (or, in the case of a transit peer, the indirect peer's subscribers) based on the Address of Record. Where transit peers are permitted, the Federation MUST either prohibit two or more transit

peers from providing access from the same indirect peer (requiring the indirect peer to choose which transit peer represents it at the Federation), or provide mechanisms allow an origination network to choose from more than one transit peer who provides transit to the indirect peer.

For interoperability reason's, each Federation MUST support at least a Federation supplied ENUM query interface where the Federation supports TN based addressing. If the ENUM data is not in the public DNS tree, the Federation MUST support a provisioning mechanism for a peer to supply it's TNS for peering

The Federation SHOULD provide mechanisms for peers to change their routing information dynamically. The change mechanism SHOULD have reasonable ways to bound the time from the initiation of the change to it's effectivity for all peers in the Federation

[3.5.](#) NAT Traversal

The Federation MAY provide facilities to assist NAT traversal, including STUN and TURN servers.

[3.6.](#) Transcode

A Federation MAY provide transcode capability. If it does, it MUST specify the mechanism by which peers engage it's transcoder service.

[3.7.](#) Capacity Controls

A Federation MAY provide mechanisms to monitor and/or limit capacity. This may take the form of mechanisms to determine and report traffic (active calls, calls/messages per unit time, media bandwidth) as well as mechanisms to limit traffic to federation/peer policy.

[3.8.](#) Protocol Mediation

A Federation MAY provide protocol mediation services to peers which would ameliorate protocol specification limits described in [Section 3.1.5](#)

[4.](#) Responsibilities of peers

[4.1.](#) Conformance to Policies

Peers MUST adhere to the policies of the Federation

[4.2.](#) Identity

Each peer must make certain the identity of the originating and terminating endpoints are reliably marked. If [[RFC3325](#)] is permitted by the Federation, the peer MUST restrict access to its services to its subscribers, provide a reliable authentication mechanism to identify them, and assert P-A-I with the actual TN for that endpoint. The peer MUST NOT allow endpoints to assert their own P-A-I unless the peer checks the validity of the assertion. If the Federation permits [[RFC4474](#)] identity, the peer MUST operate an authorization service or make a 3rd party service available to its subscribers that meet the policy of the Federation. The certificate of the authorization service MUST be signed by a CA authorized by the Federation. The peer SHOULD operate a verification service to validate Identity assertions in traffic recieved from the Federation.

[4.3.](#) Transcode

Where the peer permits endpoints to offer a codec list that does not contain a codec on the federation mandatory-to-deploy list, the peer must provide transcode capability to at least one of the codecs on the federations list for each type of media. The transcoder should be transparent to another federation peer; the offer/answer from the peer should include a codec on the federation's list, with no action required by the other side to engage the transcoder (unless it has its own, equivalent transcode issue),

[5.](#) Security Considerations

[[RFC3261](#)].

[6.](#) Normative References

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

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston,

A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", [RFC 3261](#), June 2002.

- [RFC3325] 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.

Rosen

Expires May 17, 2007

[Page 10]

Internet-Draft

BCP for Peering Through Federations

November 2006

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

Author's Address

Brian Rosen
NeuStar
470 Conrad Dr
Mars, PA 16046
US

Phone: +1 724 382 1051
Email: brian.rosen@neustar.biz

Internet-Draft BCP for Peering Through Federations November 2006

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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 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.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).