

IPv6 Operations Working Group
Internet Draft
Intended status: Experimental
Expires: September 2011

Z. Kanizsai
BME
L. Bokor
BME
G. Jeney
BME
G. Panza
CEFRIEL
March 8, 2011

IPv6 anycast based feedback data aggregation
draft-kanizsai-v6ops-anycast-data-aggregation-00.txt

Status of this Memo

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

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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 August 8, 2011.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

This document describes how to use anycast addresses for collecting feedback information on the reverse link in case of multicast forward transmission. The application for anycast addressing in the case of multicast transmission is the novelty. The draft describes the fundamentals and requirements about how to collect and aggregate feedback information if anycasting is applied.

Table of Contents

1. Introduction.....	3
2. IPv6 anycast based feedback data aggregation.....	3
2.1. Application and usage scenarios.....	3
2.2. Terminology.....	4
2.3. Protocol Operation Overview and Addressing.....	5
3. Benefits of using anycast based aggregation.....	7
4. Security Considerations.....	7
5. IANA Considerations.....	7
6. References.....	7
6.1. Normative References.....	7
6.2. Informative References.....	7
7. Acknowledgments.....	8

1. Introduction

Cross-optimized communication architectures deeply rely on collection and evaluation of different feedback information provided by network nodes periodically or on-demand. However, information is often represented in a redundant way (e.g., series of measurement data with same values can be shortly represented by a single value together with zero standard deviation). As a technique to remove redundancy and achieve efficient communication, data aggregation has been introduced and widely investigated in the literature. The aim of data aggregation is to cut back the amount of data to be transmitted while still distributing the required information about events of interest. An adequate aggregation scheme can reduce the usage of bandwidth/energy/computational power of all architectural components and nodes in the network. Data aggregation considers two main aspects. On one hand, data-centric aggregation schemes are designed to address the encoding, calculation, and compression of aggregatable data coming from multiple sources (using aggregation functions such as MAX, MIN, AVERAGE, or the probabilistic aggregation). On the other hand, routing-centric aggregation mechanisms are supposed to cover routing problems: how (e.g., when and where) information pieces (i.e., datagrams) can meet each other in order to be aggregated. In the sections below we provide a general solution for the latter issue by introducing a feedback data aggregation architecture deploying designated entities inside the network (called aggregation servers) that collect individual feedback information pieces and relay the newly composed aggregated data towards further processing in an optimal way, all based on IPv6 anycasting.

2. IPv6 anycast based feedback data aggregation

In this section anycast based feedback aggregation is introduced according to different points of view, such as the typical scenario where this technique can be used, the required new entities introduced in the network and the newly proposed addressing architecture for efficient operation.

2.1. Application and usage scenarios

The typical scenario where the application of IPv6 anycast based feedback aggregation can be beneficial is depicted in Figure 1.

This scenario includes one single Server or Source of a general service (S), i.e. content(s), which requires feedback information from the subscribers called User Entities (UE). The UEs' connection type can be fixed or mobile as well. The feedback information helps the server to provide the best service given the current network

conditions by adaptively modifying the server working parameters as required.

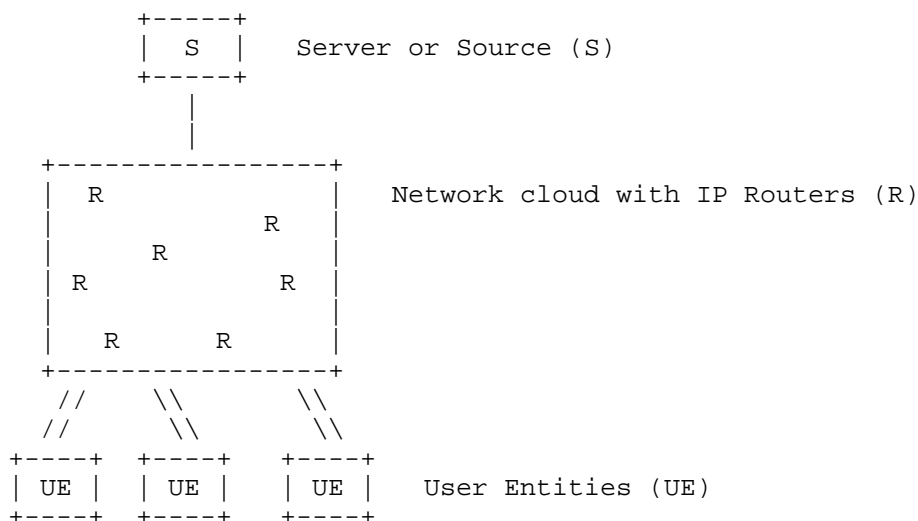


Figure 1 A typical scenario for anycast based aggregation

A feedback message is usually composed of several small numeric values which provide information about the actual quality of the service (QoS), network parameters, etc. These values are more often generated with high frequency and are only a few bits or bytes long, so sending each one to the server separately in an IP packet is a critical waste of bandwidth [FeedAgg]. To avoid this, a good solution is using IPv6 anycasting [RFC1546] and feedback aggregation in combination.

2.2. Terminology

- o Server or Source (S): A node in the network which provides for example adaptive multimedia streaming to the User Entities. This service is identified with a service anycast address which is practically identical and indistinguishable from the IPv6 unicast address of the server [RFC4291].
- o User Entity (UE): A user terminal which is able to subscribe to the service provided by the Server or Source. It measures some parameters of the received service data (e.g. multimedia stream) continuously and sends this information back to the adaptive Server to keep the QoS of the service as high as possible despite the constantly changing network conditions.

- o Anycast Capable Router (ACR): An IP packet router which is capable of handling anycast addresses and services in the network. The anycast service providers (i.e. the Feedback Aggregation Servers) are registering themselves in the closest ACR. They send their unicast address and the ID of the anycast service they are intended to participate in. During the operation of the ACR the packets that are addressed to the service's anycast address are forwarded to "at least one and preferably only one" service provider according to a parameter like hop count, the load of the servers, etc [RFC1546] [RFC4786].
- o Anycast routing protocol: A routing protocol running on the ACRs besides the normal routing process. This protocol maintains the anycast group information which is updated by the service providers periodically. A packet addressed to an anycast address is routed according to the current group state information to "at least one and preferably only one" service provider [RFC1546].
- o Feedback Aggregation Server (FAS): A server node which processes the incoming IP messages addressed to the anycast address of the service. The individual feedback messages are decapsulated and the FAS, which is aware of the feedback types of the given service, stores them in separate queues. Every feedback type has a lifetime, so the various types of feedback messages must be sent within different time constraints. When the timer expires for the first message in a queue, the complete content of this queue is placed in a new IP packet and sent to the server of the service.
- o Feedback Aggregation Address (FAA): This is the address the IPv6 packets, containing individual feedback messages, are addressed to. The FAA identifies the anycast group of the FASs and the Source. In practice, this address should be one of the Source's unicast addresses. This provides feedback delivery also in the cases when no ACR is present in the path from the UE back to the Source.

2.3. Protocol Operation Overview and Addressing

The service Source and the Feedback Aggregation Servers are in the same anycast group, addressable with the same anycast address, which should be one of the unicast addresses of the server [RFC4291]. The Source and the Feedback Aggregation Server are marked in the same way in Figure 2 according to the above reason. The Server should have assigned at least two unicast addresses, one used as the anycast group address and the other used by the aggregated IP packets sent by the FASs. The unicast packet forwarding between the FASs and the Source prevents packet looping between ACRs and FASs (Figure 2).

3. Benefits of using anycast based aggregation

In accordance with the literature, the aggregation ratio at network-level is determined by the length of the tracking history and the MTU size on the aggregation server's uplink. On average an aggregation ratio between 2 and 10 can be achieved. By applying this solution the overhead in the core network can be significantly reduced allowing also for an increased number of servable UEs for a given uplink transmission capacity of the Source.

4. Security Considerations

The above introduced solution does not raise new security issues or requirements, thus the considerations from [RFC1546] and [RFC4786] apply as well to this document.

5. IANA Considerations

This document has no new IANA considerations.

6. References

6.1. Normative References

- [RFC1546] Partridge, C., Mendez, T. and W. Milliken, "Host Anycasting Service", RFC 1546, November 1993.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC4291] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 4291, February 2006.
- [RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", RFC 4786, December 2006.

6.2. Informative References

- [FeedAgg] Kanizsai, Z., Bokor, L. and G. Jeney, "An Anycast based Feedback Aggregation Scheme for Efficient Network Transparency in Cross-layer Design", Submitted to Periodica Polytechnica Special Issue, Under review process
- [AnyTerm] Hashimoto, M., Ata, S., Kitamura, H. and M. Murata, "IPv6 Anycast Terminolgy Definition", IETF Internet Draft, draft-doi-ipv6-anycast-func-term-05.txt, January 2006, work in progress

[AnyApp] Matsunaga, S., Ata, S., Kitamura, H. and M. Murata,
"Applications of IPv6 Anycasting", IETF Internet Draft,
draft-ata-ipv6-anycast-app-01.txt, February 2005, work in
progress

7. Acknowledgments

This proposal results from the work carried out within the framework of OPTIMIX project (www.ict-optimix.eu) which is partly funded by the 7th Framework Programme (FP7) of the European Union's Information and Communication Technologies (ICT) under the contract FP7 No. INFSO-ICT-214625. The authors would like to thank all participants and contributors who take part in the studies.

The support of the Hungarian Government through the TAMOP-4.2.1/B-09/1/KMR-2010-0002 project at the Budapest University of Technology and Economics is also acknowledged.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Zoltan Kanizsai
Department of Telecommunications
Budapest University of Technology and Economics (BME)
Magyar Tudosok krt. 2. IB121
H-1117, Budapest
Hungary

Email: kanizsai@hit.bme.hu

Laszlo Bokor
Mobile Innovation Centre
Budapest University of Technology and Economics (BME)
Bertalan Lajos u. 2. Z301
H-1111, Budapest
Hungary

Email: goodzi@mcl.hu

Gabor Jeney
Department of Telecommunications
Budapest University of Technology and Economics (BME)
Magyar Tudosok krt. 2. IE450
H-1117, Budapest
Hungary

Email: jeneyg@hit.bme.hu

Gianmarco Panza
Digital Platform and Pervasive ICT Division
CEFRIEL - Politecnico di Milano
via Fucini 2, 20133 Milan
Italy

Email: Gianmarco.panza@cefriel.com

