Individual Submission
draft-haberman-ipv6-anycast-rr-00.txt
October 2002
Expires April 2003

B. Haberman
Caspian Networks
E. Nordmark
Sun Microsystems

IPv6 Anycast Binding using Return Routability

Status of this Memo

Abstract

Today, the use of IPv6 anycast is limited. This document proposes a mechanism by which TCP/SCTP and stateful protocols using UDP can securely discover the mapping from an anycast address to a unicast address that can be used until a failure is detected. The mechanism reuses the Mobile IPv6 Return Routability and Binding Update mechanism.

Introduction

Several limitations in IPv6 anycast have been identified[ANALYSIS]. This document proposes a solution for securely discovering the mapping from an anycast address to a unicast address so that protocols which require a sequence of packets to go to the same anycast receiver can take advantage of anycast for the initial discovery of the unicast address of an anycast member.

This document proposes an approach to IPv6 anycast communication utilizing the Return Routability Procedure defined in [MOBILEIP].

This approach will allow client nodes to use IPv6 anycast for initial packet exchanges and determine the unicast address of the anycast member. After which, the anycast member will provide sufficient information for the client node to "bind" a unicast address assigned to the member.

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 [RFC 2119].

Other terms in this document are defined in [MOBILEIP].

Overview

The conceptual model of anycast communication contains many similarities to the Mobile IPv6 model. One similarity is that the node listening on an anycast address is reachable at two addresses (the anycast address and a unicast address) much like a MIP Mobile Node is reachable at a Home Address and a Care-of Address. As with MIPv6, there has to be a secure mechanism to establish a binding between the anycast and unicast address. The node originating the communication takes on the role of Mobile IPv6 correspondent node. The correspondent node can utilize the binding information to make communication efficient.

The primary difference between the anycast model and the Mobile IPv6 model is that once a binding is known, it is permanent for the duration of the communication session between the nodes or until a failure is detected. That is point, the originating node will want to establish a new binding.

What this memo defines is a mechanism for applying the return routability procedure from Mobile IPv6 to anycast communication.

Anycast Return Routability Procedure

The return routability signaling for anycast proceeds as follows:

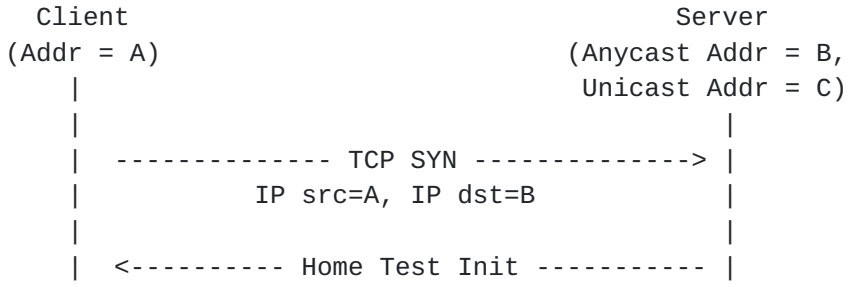

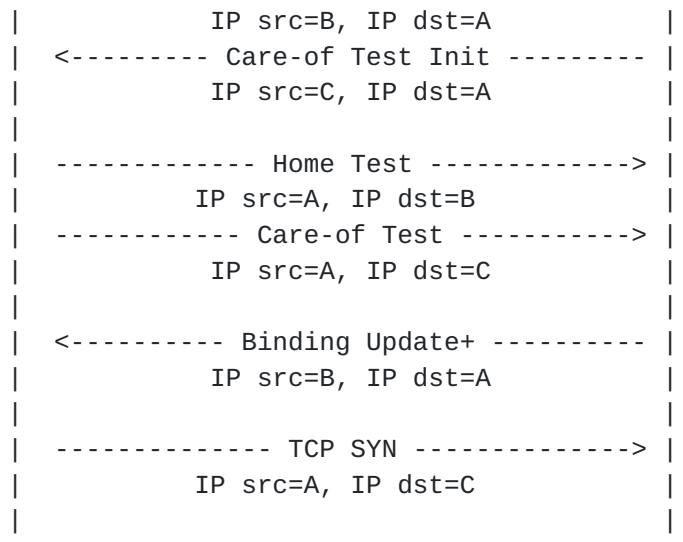

The Home Test Init (HoTI) and the Care-of Test Init (CoTI) are transmitted at the same time by the server. The client responds with the Home Test (HoT) and the Care-of Test (CoT) messages as quickly as possible.

The contents of each message, their formats, and the basic processing rules are defined in [MOBILEIP].

Binding Update+

The Binding Update message returned by the server requires an additional indication that the binding is for an anycast address. This allows the client stack to drop any connection state associated with the anycast address and recreate it using the server's unicast address.

This indication can be accomplished by reserving an additional flag within the Mobile IPv6 Binding Update message. The new Binding Update message format would be as follows:

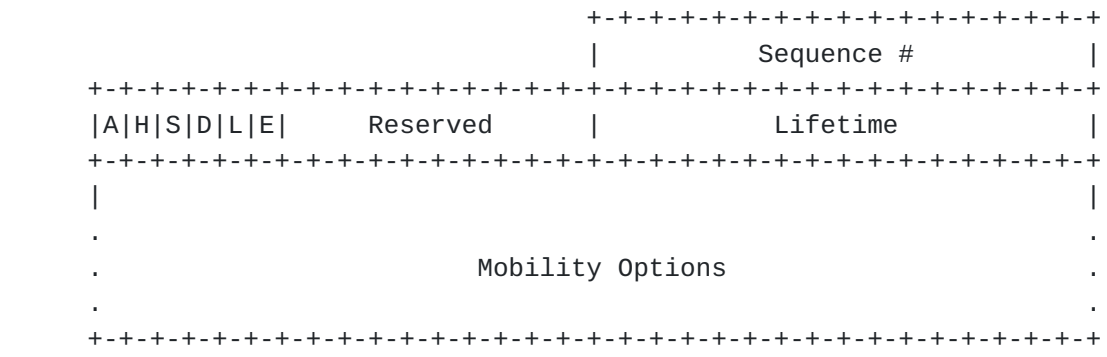

The E (for equivalent) bit indicates a switchover from an anycast

address to a unicast address. It MUST be set when sending a Binding Update for an anycast address. It should be noted that a Binding Update with E=1 indicates that a Mobile IPv6 RHtype2 extension header is not included in the packet and that transport protocols should be made aware of the change in the destination address.

The Lifetime field of the Binding Update MUST be set to a value of one (16 seconds). This allows the binding to be used for immediate communication. Longer lifetimes will cause the anycast receivers to exert control over which member of the anycast group receives packets addressed to it. By having a short lifetime, the control over delivery of packets to the anycast group is maintained by the routing system.

Open Issues

- This mechanism requires the use of an anycast address as a source address for the CoTI message. The issues with this needs to be carefully understood with respect to [ANALYSIS].
- Additional text needed on Mobile IPv6 message exchange?
- Additional security threats?
- Should an RHtype2 be used to deliver the packets for better consistency with MIPv6?

Security Considerations

Security considerations are discussed in [MOBILEIP].

References

[RFC 2373] Hinden, R., and Deering, S., "IP Version 6 Addressing Architecture", RFC 2373, July 1998.

[ANALYSIS] Hagino, J., and Ettikan, K., "An Analysis of IPv6 Anycast", work in progress.

[MOBILEIP] Johnson, D., Perkins, C., and Arkko, J., "Mobility Support in IPv6", work in progress.

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

Authors' Address

Brian Haberman
Caspian Networks
One Park Drive
Suite 400
Research Triangle Park, NC 27709
Tel: +1-919-949-4828
EMail: bkhabs@nc.rr.com

Erik Nordmark
Sun Microsystems Laboratories
180, avenue de l'Europe
38334 SAINT ISMIER Cedex, France
Tel: +33 (0)4 76 18 88 03
Fax: +33 (0)4 76 18 88 88
EMail: erik.nordmark@sun.com

Full Copyright Statement