

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 22, 2010

D. Jen
L. Zhang
UCLA
October 19, 2009

Understand Mapping
draft-jen-mapping-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/ietf/1id-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 April 22, 2010.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

This draft discusses the different requirements that mapping to support mobility has versus mapping to support routing scalability. In mobility solutions, packets are forwarded to the location where

mapping information is stored so that they can be encapsulated to the final destination. In routing scalability solutions, mapping information needs to be available at packet entry points to the network. Attempts to use one mapping system for both purposes can lead to high overhead in either mapping information distribution or otherwise mapping information discovery, as well as stale mapping information being used for data forwarding.

Table of Contents

1.	Introduction	3
2.	Mapping for Routing Scalability	3
3.	Mapping for Mobility Support	4
4.	Differences in Mapping Requirements: Mobility Support vs Scalable Routing	5
5.	Concluding Remarks: Separating Mobility Support from Scalable Routing	6
6.	References	7
	Authors' Addresses	7

1. Introduction

Since early 2007, the IRTF has recharted the Routing Research Group (RRG) to develop a scalable routing architecture. Over the last two years a number of proposed solutions have been presented to the RRG. Although the final recommendation is still being worked on, the general direction taken by most RRG proposals is based on the map & encap approach [[RFC1955](#)]. This has led many people to ask the same recurring question: since mapping is needed to solve the routing scalability problem, and since mobility solutions typically involve a binding/mapping function between an identifier and a mobile's current IP address, can the same mapping solution for routing scalability be used to support mobility?

We believe that some basic differences exist between the requirements for the mapping needed for mobility support and that for routing scalability. This draft is an attempt to articulate the basic differences.

2. Mapping for Routing Scalability

Mapping for routing scalability is meant to reduce the information storage requirements at routers. Scaling the routing system requires an effective way to remove some number of entries in today's routing system from some routers, without losing reachability to any destinations. Once this is done, not all routers have routing information for all destinations anymore. Thus packets are mapped and then either encapsulated or translated before being forwarded over a network of routers that do not have the specific routing information for the packet destinations.

Since part of the current routing information is replaced by mapping information, the mapping information must be stored in some locations. The number and placement of these locations differ vastly from one RRG proposal to another. They range from the distribution of a full mapping database to all entry points to the network (thus only reducing the routing table size of internal routers), to keeping the mapping information at distributed locations and building a search tree to find individual pieces.

NERD represents one extreme point of the mapping system design by distributing the full mapping table to all network entry points. NERD assumes a central database to collect all the mapping information and updates, from which all the ITRs (ingress tunnel router) fetch the full mapping table and updates periodically. This design enables ITRs to encapsulate incoming packets directly to ETRs. A few other schemes such as IVIP and APT can be viewed as variations

of NERD in terms of mapping information distribution. APT assumes that a small number of nodes in every AS (collectively) hold the full mapping table, while ITRs query for mapping entries and cache the results (which prevents suboptimal paths and overload of the mapping holders). Although such variations reduce the number of nodes that hold the full mapping table, all mapping entry changes still must be pushed out to many nodes.

LISP ALT may be viewed as representing the other extreme point on the design spectrum. Instead of distributing mapping information out, ALT builds a hierarchical search tree for non-globally-routable IP addresses (the "EID space" in LISP terminology), so that each ITR can find the mapping for a given destination address by walking up and down the tree. ALT avoids the cost of any nodes holding a full table, by paying the cost of delay in finding the mapping information. In addition, because the hierarchical search tree is shared by all ITRs for all destinations, the tree itself, especially nodes at the top, can become overloaded. ALT lets ITRs cache the mapping information to avoid the lookup delay and reduce the load on the search tree.

3. Mapping for Mobility Support

The basic question for mobility support is how to reach a moving receiver. Whoever sends packets to a moving receiver must be able to identify the receiver via a piece of stable information (stable in the sense that it does not change as the receiver moves). However, if the sender's knowledge about the receiver does not change while the receiver moves, some means must exist to bind that unchanging identifier of the receiver with its dynamically changing location. Locations in the Internet are represented by IP addresses.

The above leads to the following intuitive observations: mobility support essentially involves three basic concepts: a stable identifier for a moving receiver, an IP address of the receiver's current location, and a mapping in between. Different mobility support designs are simply different ways of choosing receiver identifiers and different approaches to providing mappings between the identifiers and the receivers' current IP addresses.

One example of a mobility support scheme is Mobile IP [[RFC3344](#)]. It uses a home agent IP address as the moving host's identifier. This address plays a dual role of being both an identifier and also an IP address to send packets to. MIP requires that a home agent be able to keep track of the current location of the moving host and to intercept packets sent to the mobile's home address. Mobility means potentially constant movement of mobile nodes, hence constant changes

to the mapping between mobile identifiers and their current locations. Individual home agents keep track of dynamic mapping changes for the mobile nodes under their care. All packets to mobile nodes are sent to corresponding home agents, thus home agents do not need to propagate the dynamic mapping changes anywhere. All senders wishing to communicate with the same mobile node obtains its home agent address via DNS lookup, and they send data directly to that IP address. If a single node wants to serve as a home agent for a large number of mobiles, it will receive all the traffic for all the mobile nodes.

4. Differences in Mapping Requirements: Mobility Support vs Scalable Routing

A fundamental difference between the two mapping systems is as follows. In scalable routing solutions, generally speaking some information is deliberately removed from nodes to relieve their routing storage and associated update processing requirements. Destinations are subtracted from routing tables. Subsequently, nodes do not know how to route to all IP addresses. Mapping is required to map the many unroutable destinations to routable addresses for encapsulation (or translation) and delivery. In mobility support solutions such as MIP, however, mapping and encapsulation serve an entirely different purpose. Unlike in scalable routing schemes, mobility solutions assume that all IP addresses are reachable in the global Internet (it is up to the routing system to assure that). Rather, new reachability information is to the tables of home agents in the form of mappings. Home agents map mobile node identifiers to their most recent IP addresses. Since mobility is made transparent to a correspondent, it always sends packets to the mobile node's home address, the packets are then intercepted by the home agent and encapsulated to reach the mobile node. The home agent keeps the mapping information, but does not need to distribute it anywhere, because all packets going to the mobile are sent to the home subnet that the home agent sits on. The encapsulation simply preserves the original packet as sent by the correspondent.

From the observations above, one can see that a number of issues arise if one tries to add mobility support to the same mapping system used for routing scalability. Recall that movement leads to mapping changes. As more and more portable devices are becoming Internet-capable, the number of mobile nodes will likely grow larger and larger over time, which would lead to more and more overall mapping changes (even if the movement frequency of a single node remains steady). If one tries to push the changes out to all the nodes that hold full mapping tables as needed in NERD, IVIP, and APT, this simply does not scale well enough to support networks with large

numbers of mobile nodes and subnets. On the other hand, if one uses the distributed mapping nodes themselves as "home agent" equivalents, so that mapping changes can simply be kept at "home agents" without being propagated anywhere, such as the case in ALT, this also comes with its own issues. In mobility schemes such as MIP, packets to the mobile node go directly to the home address and get redirected by the home agent. However, in the scalable routing context, ITR routers do not know the ETRs (=home agent) which hold the mapping information; packets have to walk up and down the search tree to reach ETRs. This would overload the search tree with packets. This problem is exactly why caching was introduced into scalability schemes in the first place. However, caching here conflicts with mobility support. Once a entry router caches the current address of a mobile nodes, it has no way to know when the mobile moves again, resulting in stale cache entries being used for packet forwarding. So far there has been no good solution to this problem. Using short TTLs for caching simply lead us back to an overload of the mapping system.

Having each mobile node keep track of its current correspondents and directly informing them of movement can help with the stale cache problem for existing communications, but does not help new senders who rely on ITRs having the correct mapping information to contact mobiles. Even for the ongoing communication, when both ends are mobiles and move at the same time, one falls back to the mapping system to get reconnected again. A fundamental problem with stale mapping information at ITRs is a failure in packet delivery that end hosts have no way to get around. Caching is mandatory for mapping in scalable routing solutions, yet caching makes the mapping system no longer support mobility very well anymore.

5. Concluding Remarks: Separating Mobility Support from Scalable Routing

It is very tempting to believe that the same mapping service can provide both mobility and scalability, since so many mobility and scalability solutions all involve some type of mapping. However, attempts to inject mobility mappings into the scalability mapping system have revealed that the two do not fit together very well. We hope this draft clarifies the reasons that mobility mappings and scalability mappings do not mix. It is still very possible that a mobility solution can co-exist with a scalability solution, but one must be careful not to try to bundle both solutions into one mapping system.

6. References

[RFC1955] Hinden, R., "New Scheme for Internet Routing and Addressing (ENCAPS) for IPNG", [RFC 1955](#), 1996.

Authors' Addresses

Dan Jen
UCLA
4805 Boelter Hall, UCLA
Los Angeles, CA 90095
US

Email: jenster@cs.ucla.edu

Lixia Zhang
UCLA
3713 Boelter Hall, UCLA
Los Angeles, CA 90095
US

Email: lixia@cs.ucla.edu

