

Network Working Group
Request for Comments: 2002
Category: Standards Track

C. Perkins, Editor
IBM
October 1996

IP Mobility Support

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document specifies protocol enhancements that allow transparent routing of IP datagrams to mobile nodes in the Internet. Each mobile node is always identified by its home address, regardless of its current point of attachment to the Internet. While situated away from its home, a mobile node is also associated with a care-of address, which provides information about its current point of attachment to the Internet. The protocol provides for registering the care-of address with a home agent. The home agent sends datagrams destined for the mobile node through a tunnel to the care-of address. After arriving at the end of the tunnel, each datagram is then delivered to the mobile node.

Table of Contents

1. Introduction	3
1.1. Protocol Requirements	3
1.2. Goals	4
1.3. Assumptions	4
1.4. Applicability	4
1.5. New Architectural Entities	5
1.6. Terminology	6
1.7. Protocol Overview	8
1.8. Specification Language	11
1.9. Message Format and Protocol Extensibility	12
2. Agent Discovery	14
2.1. Agent Advertisement	14
2.1.1. Mobility Agent Advertisement Extension	16
2.1.2. Prefix-Lengths Extension	18
2.1.3. One-byte Padding Extension	19
2.2. Agent Solicitation	19
2.3. Foreign Agent and Home Agent Considerations	19
2.3.1. Advertised Router Addresses	20

C.2. TCP Congestion Management	73
D. Example Scenarios	74
D.1. Registering with a Foreign Agent Care-of Address	74
D.2. Registering with a Co-Located Care-of Address	75
D.3. Deregistration	76
E. Applicability of Prefix Lengths Extension	76
Editor's Address	79

1. Introduction

IP version 4 assumes that a node's IP address uniquely identifies the node's point of attachment to the Internet. Therefore, a node must be located on the network indicated by its IP address in order to receive datagrams destined to it; otherwise, datagrams destined to the node would be undeliverable. For a node to change its point of attachment without losing its ability to communicate, currently one of the two following mechanisms must typically be employed:

- a) the node must change its IP address whenever it changes its point of attachment, or
- b) host-specific routes must be propagated throughout much of the Internet routing fabric.

Both of these alternatives are often unacceptable. The first makes it impossible for a node to maintain transport and higher-layer connections when the node changes location. The second has obvious and severe scaling problems, especially relevant considering the explosive growth in sales of notebook (mobile) computers.

A new, scalable, mechanism is required for accommodating node mobility within the Internet. This document defines such a mechanism, which enables nodes to change their point of attachment to the Internet without changing their IP address.

1.1. Protocol Requirements

A mobile node must be able to communicate with other nodes after changing its link-layer point of attachment to the Internet, yet without changing its IP address.

A mobile node must be able to communicate with other nodes that do not implement these mobility functions. No protocol enhancements are required in hosts or routers that are not acting as any of the new architectural entities introduced in Section 1.5.

All messages used to update another node as to the location of a mobile node must be authenticated in order to protect against remote redirection attacks.

1.2. Goals

The link by which a mobile node is directly attached to the Internet may often be a wireless link. This link may thus have a

powered, and minimizing power consumption is important. Therefore, the number of administrative messages sent over the link by which a mobile node is directly attached to the Internet should be minimized, and the size of these messages should be kept as small as is reasonably possible.

1.3. Assumptions

The protocols defined in this document place no additional constraints on the assignment of IP addresses. That is, a mobile node can be assigned an IP address by the organization that owns the machine.

This protocol assumes that mobile nodes will generally not change their point of attachment to the Internet more frequently than once per second.

This protocol assumes that IP unicast datagrams are routed based on the destination address in the datagram header (and not, for example, by source address).

1.4. Applicability

Mobile IP is intended to enable nodes to move from one IP subnet to another. It is just as suitable for mobility across homogeneous media as it is for mobility across heterogeneous media. That is, Mobile IP facilitates node movement from one Ethernet segment to another as well as it accommodates node movement from an Ethernet segment to a wireless LAN, as long as the mobile node's IP address remains the same after such a movement.

One can think of Mobile IP as solving the "macro" mobility management problem. It is less well suited for more "micro" mobility management

Perkins	Standards Track	[Page 4]
RFC 2002	IP Mobility Support	October 1996

applications -- for example, handoff amongst wireless transceivers, each of which covers only a very small geographic area. As long as node movement does not occur between points of attachment on different IP subnets, link-layer mechanisms for mobility (i.e., link-layer handoff) may offer faster convergence and far less overhead than Mobile IP.

1.5. New Architectural Entities

Mobile IP introduces the following new functional entities:

Mobile Node

A host or router that changes its point of attachment from one network or subnetwork to another. A mobile node may change its location without changing its IP address; it may continue to communicate with other Internet nodes at any location using its (constant) IP address, assuming link-layer connectivity to a point of attachment is available.

Home Agent

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.