

Multihomed ISPs and Policy Control

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [Section 10 of RFC2026](#).

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/lid-abstracts.html>

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

Abstract

Policy control of next level ISPs, delegated address spaces from top level ISPs, is discussed.

While global policy coordination requires top level aggregators, local policy can be controlled with next level aggregators.

1. Introduction

Considering that some people has been arguing to have 4 byte AS numbers, the number of ISPs will grow indefinitely

On the other hand, having small number of TLIs will make full routing table small that it can be expected that most hosts has full routing table, reducing the problems with destination and source address selection.

An obvious solution is to have layers of ISPs, as was specified in IPv6 to have TLA (Top Level Aggregator) and NLA (Next Level Aggregator) that they can be allocated to TLIs (Top Level ISPs) and

INTERNET DRAFT

Multihomed ISPs and Policy Control

June 2003

NLIs (Next Level ISPs, correspondingly.

The problem, however, is whether and how the number of TLIs can be controlled.

[2. Robustness](#)

An essential property of ISPs is robustness of its service that it is almost mandatory that NLIs are multihomed to multiple TLIs.

It is, then, expected that some sites are multihomed to multiple TLIs and/or NLIs.

It is expected that NLIs have multiple prefixes each belonging to multiple TLAs, all of which is delegated to sites.

[3. BGP Policy Control](#)

BGP Policy is controlled by identifying ISPs not by address prefix but by AS numbers.

Thus, a next level ISP not having its own TLA can still fully control its policy.

Moreover, neighbour ISPs can adjust their policy using the full prefix of the ISP.

However, to limit the size of global routing table, an AS and prefixes of the next level ISP at the distance must be discarded or merged with its TLA.

But, it is better than a current multihoming practice that prefix of a multihomed site is propagated locally to give robustness against local failures, because multiple TLAs give robustness against global failures.

Thus, it is not essential that ISPs have their own TLAs.

[4. Limiting the Number of TLAs.](#)

There should be hard upper bound on the number of TLAs in the Internet.

For example, some TLA may be supplied from RIRs with bidding.

Some TLA may be allocated to each country (proportional to the population of the country) and delivered with the country's policy.

M. Ohta

Expires on December 23, 2003

[Page 2]

INTERNET DRAFT

Multihomed ISPs and Policy Control

June 2003

The proper number of TLAs, it seems to the author, should be somewhere between 1024 and 8192.

[5.](#) Author's Address

Masataka Ohta
Graduate School of Information Science and Engineering
Tokyo Institute of Technology
2-12-1, O-okayama, Meguro-ku, Tokyo 152-8552, JAPAN

Phone: +81-3-5734-3299

Fax: +81-3-5734-3299

EMail: mohta@necom830.hpcl.titech.ac.jp

[M.](#) Ohta

Expires on December 23, 2003

[Page 3]