

GAPI: A Geographically Aggregatable Provider Independent
Address Space to Support Multihoming in IPv6

1 Mandatory statements

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>

2 Abstract

This document establishes an address allocation framework for Geographically Aggregatable Provider Independent (GAPI) IPv6 addresses for the purpose of multihoming. A /16 is divided in a hierarchical manner over geographical entities such as parts of continents, countries, states, provinces and metropolitan areas, with each receiving one or more /32 allocations from which end-user assignments can be made. The number of /32 allocations for a geographical entity depends on the current population. This address allocation framework in itself does not present a scalable multihoming solution, it merely provides a necessary building block for such solutions.

3 Introduction

At present, there are at least two ways to support multihoming in IPv6 under development that are based in whole or in part on geographical address aggregation to support the necessary scalability:

provider internal aggregation and [1] and Multi Homing Aliasing Protocol (MHAP) [2]. Since a common address allocation framework provides for easy migration from [1] to [2], the authors thought it prudent to cooperate in developing such a framework. The new address allocation policies outlined in this document are intended to be used

for multihomed end-users only; regular provider aggregatable address space should be used for internet service providers and their single homed customers.

The basic presumption behind all kinds of geographical aggregation is that global flat routing, where all routers have full routing information for all possible multihomed destinations connected to the Internet world wide, isn't feasible in the long run. The number of multihomed networks is likely to grow faster than the memory in routers that is available to store this information. This relationship is linear.

However, processing updates for a larger number of routes scales (slightly) worse than linear. To get around this lack of scalability, a geographical aggregation scheme splits the global routing domain into a number of smaller regional ones, where flat routing happens in each region. Ideally, outside the region only aggregates are visible. For simplicity and to allow efficient implementation, the framework presented here requires "areas" where flat routing takes place to have a fixed size: a /32 holding up to 65536 (2^{16}) fixed sized end-user /48 assignments. The maximum number of these /32 areas is also 65536. Areas are grouped in CIDR-like fashion if a geographic region has a population that warrants allocating more than a single /32. The highest level of aggregation is the subcontinent or "zone" level. There are 13 entities at this level, in order of population:

- [1.](#) China
- [2.](#) Continental Asia
- [3.](#) India
- [4.](#) Northern Africa
- [5.](#) Asian Islands
- [6.](#) Western Europe
- [7.](#) North America
- [8.](#) South America
- [9.](#) Eastern Europe
- [10.](#) Middle East
- [11.](#) Southern Africa
- [12.](#) Central America
- [13.](#) Oceania

The next level is the country level. Every country is assigned a range of /32 blocks, depending on population. Countries that are medium-sized or larger may be subdivided according to existing administrative boundaries, such as by state or province. The allocation size per state or province must match the population relative to the country and other states or provinces.

The lowest level of aggregation is the metropolitan level. Cities of sufficient size are allocated one or more "metro areas". Assignments to end-users in, or very close to, a city are drawn from one of the metro area /32s allocated to the city. Addresses for end-users in small cities or rural areas are drawn from one of the /32 areas

allocated to the country (if not subdivided), state or province (a country/state/province or "CSP" area).

4 Rationale for population base

If all geographic entities at each aggregation level were to receive the same amount of address space, a GAPI allocation scheme would either waste enormous amounts of address space, or it would only allow a limited number of multihomed networks in densely populated areas. Neither of these options ensures the necessary scalability. In theory, it would be possible to allocate a completely new address range when the previously allocated one depletes. This is done in current IPv4 provider-based aggregation. However, this new range wouldn't fit in the existing hierarchy. In addition, having a limited number of entities at each hierarchical level improves scalability and manageability. Thus, allocating completely new GAPI address ranges should be done as infrequently as possible.

Using current multihoming deployment as a base for GAPI address allocation to regions would make sense for the short term, but in the long run it seems reasonable to expect regions with little multihoming to catch up, while in regions where multihoming deployment is currently higher, the growth in multihoming is likely to level off at some point. Predicting future multihoming needs is hard, especially over a longer period of time. Also, such predictions might be very politically charged.

This leaves two possibilities to base the initial GAPI allocation sizes on: geography and population. Geography has the advantage it doesn't change much over time, but the main disadvantage is that it uses the address space very inefficiently. This makes current population the best variable to base initial allocations on. Even

using very optimistic assumptions, there are no geographic areas of significant size where current levels of multihoming per capita are higher than 1:25000. The framework presented here should provide up to 1:10 at the bottom of the aggregation hierarchy using an initial /16 allocation, based on current population. This growth path should provide enough time to develop new criteria to base allocation of a subsequent /16 or larger block on.

[5 Allocation policy](#)

The goals of the allocation policy are:

- [1.](#) Be completely neutral, fair and unbiased, in order to minimize the potential for political complications
- [2.](#) Good aggregation at all levels
- [3.](#) Reasonable flexibility
- [4.](#) Ease of implementation

Py & Van Beijnum
[draft-py-multi6-gapi-00.txt](#)

Page 3
October 2002

[5.1 Country allocations](#)

Each independent country is allocated at least one /32 area. The allocation size depends on the country's population figure for the year 2001. This is divided by the number D1, which equals 131072. The result of the division is rounded up to the next power of two.

This is the number of /32s constituting the country's allocation.

[5.2 Zone allocations](#)

The subdivision of the globe in 13 zones is relatively arbitrary. However, this division fits current and expected future Regional Internet Regions well, and limits the population per zone somewhat over a strict by-continent subdivision. Zone allocations are chosen such that they are large enough to hold the country allocations for all countries located within the geographic bounds of the zone. If for any of the zones that encompass more than a single country, the number of /32s not allocated to countries is less than 25%, the zone allocation size is doubled.

[5.3 Subdivision of large countries](#)

When a country has an allocation of 32 or more /32s, this address

space may be distributed over the country by allocating blocks of /32s to existing sub-entities such as provinces or states. The exact geographic bounds of these sub-entities must be clear to the general public and not subject to any controversy. The size of each allocation is determined by dividing the population of the sub-entity by the number D2, which is twice D1.

At least 40% of the country allocation must remain unallocated. If necessary, a higher value than D2 may be used as a divisor in this country to reach this objective. The average number of /32 areas per state or province must be at least 4.

[5.4](#) Metro allocations

All cities with a population of at least D2 within the city limits are allocated a block of /32s. The population for small cities or municipalities that do not qualify for an address block of their own is added to the closest city that qualifies, if there is one within 40 kilometers. (Distance measured center to center.) The size of the address block for a city and its surroundings is determined by taking the total populace, dividing it by D2 and rounding down to the next power of two.

[5.5](#) Reserved for future use

The first 1/64th of each allocation at the country/state/province level is reserved for future special uses and must not be allocated to lower aggregation levels and not be assigned to end-users.

[5.6](#) Subsequent allocations

Whenever allocated address space gets close to running out, the IANA, Regional Internet Registry or other organization managing (part of) the address space should draw new allocations from the next higher level. New blocks of address space may be allocated in a way that is different from what is outlined here, if analysis of the coordinates for current assignments warrants this.

[6](#) End-user address assignments

Per country or state/province, only one /32 block is used initially. A new block is used when the first is exhausted, and so on. The /32s allocated to a metropolitan area may be put into use concurrently, if there is a reason to do so.

When requesting IPv6 GAPI addresses, an organization should provide justification for the use of GAPI space, and information that makes it possible to assign addresses from the right geographic area, in addition to the information required by current assignment policies. Multihoming is justification for the use of GAPI space. Geographic information should consist of the longitude and latitude of the primary location where the addresses will be used. This information should be accurate within 2 kilometers, as long as any inaccuracies don't make the organization appear to be at the other side of an administrative border or natural barrier (such as a river). Preferably, the requesting organization should also include the longitude and latitude of the ISP locations they connect to. However, this information may be omitted.

The minimum assignment size is always /48. Future multihoming solutions may not support the longest match first rule.

[7](#) IANA considerations

IANA considerations are discussed throughout this memo. More specifically, the IANA/ICANN is requested to allocate /16 worth of IPv6 address space for GAPI, and the Regional Internet Registries are asked to further assign this address space to end-user organizations.

[8](#) Security considerations

Having addresses that are closely tied to an organization's location may be undesirable in certain situations. Organizations requesting address space should consider the consequences of using GAPI address space, and are encouraged to use provider aggregatable address space if and when they want to avoid disclosing their location.

Some organizations may be uncomfortable with providing very accurate longitude and latitude information when requesting address space. They may introduce a 2 kilometer inaccuracy to avoid exact pinpointing, as described in [section 6](#). In addition, the Regional Internet Registry

or other organization responsible for assigning address space must not make location information public. Specifically, this information should not appear as a result of whois queries. Registries are encouraged to provide aggregated location information for policy development purposes, but only as long as this information is anonymized and can't be tied to a single organization or small group of organizations.

9 Document and author information

The latest version of this document is always available at:

<http://arneill-py.sacramento.ca.us/ipv6mh/>

<http://www.muada.com/drafts/>

The authors can be reached at:

Michel Py

[2825](#) Marshall Way

Sacramento, CA 95818-3525

USA

Email: michel@arneill-py.sacramento.ca.us

Iljitsch van Beijnum

Karel Roosstraat 95

[2571](#) BG Den Haag

Netherlands

Email: iljitsch@muada.com

References

- [1] "Provider-Internal Aggregation based on Geography to Support Multihoming in IPv6", work in progress
- [2] "Multi Homing Aliasing Protocol (MHAP)", work in progress