

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: July 28, 2013

M. Levy  
Hurricane Electric  
M. Pounsett  
Afiliias  
January 24, 2013

A mechanism to allocate IPv6 blocks for BGP networks based on the  
networks AS Number  
draft-mlevy-v6ops-auto-v6-allocation-per-asn-00.txt

## Abstract

This document provides a methodology for automatically allocating IPv6 [[RFC2460](#)] address blocks for networks that run BGP [[RFC4271](#)] and are either single-homed or multi-homed [[BARBER2011](#)]. The automatic allocation is taken from a specific /16 block assigned by IANA for this purpose.

Networks that require more than this single /48 can still request additional allocations via the existing RIR process. Networks are not forced to use this allocation and can ignore this completely. Availability of the /48 assignment via this mechanism does not change existing mechanisms for obtaining IPv6 assignments through the existing RIR (Regional Internet Registry) or LIR (Local Internet Registry) mechanisms.

There is an implicit assumption that it's a good thing to promote networks to enable IPv6 with a near-zero effort.

## Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on July 28, 2013.

## Copyright Notice

---

Internet-Draft   Auto allocation mechanism for IPv6 blocks   January 2013

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Defining the /48 address . . . . .	<a href="#">3</a>
<a href="#">3.</a>	IANA allocated /16 . . . . .	<a href="#">4</a>
<a href="#">4.</a>	ASN bit length . . . . .	<a href="#">4</a>
<a href="#">5.</a>	ASN allocation . . . . .	<a href="#">5</a>
<a href="#">6.</a>	BGP Filtering and Validation . . . . .	<a href="#">5</a>
<a href="#">7.</a>	Whois . . . . .	<a href="#">5</a>
<a href="#">8.</a>	RIR support . . . . .	<a href="#">5</a>
<a href="#">9.</a>	Reverse DNS . . . . .	<a href="#">5</a>
<a href="#">10.</a>	Routing Table Impact . . . . .	<a href="#">6</a>
<a href="#">11.</a>	IANA Considerations . . . . .	<a href="#">6</a>
<a href="#">12.</a>	Security Considerations . . . . .	<a href="#">6</a>
<a href="#">13.</a>	Acknowledgements . . . . .	<a href="#">6</a>
<a href="#">14.</a>	References . . . . .	<a href="#">6</a>
<a href="#">14.1.</a>	Normative References . . . . .	<a href="#">6</a>
<a href="#">14.2.</a>	Informative References . . . . .	<a href="#">7</a>
	Authors' Addresses . . . . .	<a href="#">7</a>

---

Internet-Draft   Auto allocation mechanism for IPv6 blocks   January 2013

## [1.](#) Introduction

IPv6's massive address space provides a wonderful opportunity to radically change the process of IP address allocation. Presently IPv6 allocation is the same process as used by the IPv4 world. The RIRs allocate IPv6 address blocks based on member requests and their space requirements. A minimum allocation of a /48 is believed to be sufficient to start any enterprise in the world of IPv6.

Described below is a method to automatically define an IPv6 /48 block that can be used in the global routing tables for any ASN.

## [2.](#) Defining the /48 address

The /48 address is allocated using a fixed pattern.

ASN-based address assignment

```
+-----+-----+-----+-----+-----+-----+-----+-----+
| IANA /16 |      32 bit ASN      |      80 bits for IPv6 /48 space      |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

The sum of the bits in the IANA allocated /16 prefix plus the 32 bits of the ASN plus the /48 of user defined usage adds up to 128 bits.

Bits	Usage
16	IANA provided prefix
32	ASN
80	User defined
128	TOTAL

Table 1: Address space math

The end network can allocate out of the assigned /48 as needed. It is assumed that the end network will use this allocation for global routing; however the network may choose to not announce this allocation.

If any route is announced from this allocation, any prefixes more specific than the allocated /48 must not be propagated in to the global IPv6 routing table. This is to prevent the IPv6 routing table from becoming too large. Therefore, a site which uses this allocation MUST NOT advertise a more specific than the allocated /48 routing prefix. All native IPv6 network operators MUST filter out

and discard any routing prefix advertisements longer than /48 from within this /16 allocation.

### [3.](#) IANA allocated /16

IANA is to allocate a /16 in a similar manner to the 2002::/16 allocation for 6to4 [[RFC3068](#)] which is used by the 6to4 protocol[RFC3056].. No IPv4 allocation by IANA is required.

### [4.](#) ASN bit length

ASNs are now defined as 32 bits[RFC4893] long. If a network operates a smaller 16 bit ASN (for example an earlier allocation) then there are 16 bits of zeros prepending the ASN number. No provision is provided within this allocation methodology to provide 16 bit ASNs with an advantage as this could cause an overlapping allocation.

No special treatment is provided an operator of a 16 bit ASN; All ASNs are considered to use 32 bits.

16 bit ASNs (shown at binary)

```
+-----+-----+-----+-----+-----+-----+-----+-----+
| 0000 | 0000 | 0000 | 0000 | ##### | ##### | ##### | ##### |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

32 bit ASNs (shown at binary)

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| ##### | ##### | ##### | ##### | ##### | ##### | ##### | ##### |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

ASNs are normally expressed as human-readable decimal numbers; yet for this allocation the number should be converted into a hexadecimal notation. All IPv6 addresses are written in hexadecimal. (NNNN represents the /16 allocation by IANA)

ASN in decemal	ASN in hexadecimal	Sample IP block
AS1	1	NNNN:0000:0001::/48
AS6939	1B1B	NNNN:0000:1B1B::/48
AS29001	7149	NNNN:0000:0001::/48
AS393220	60004	NNNN:0006:0004::/48

Table 2: ASN examples

A network is assigned an ASN via the existing RIR process. Only valid ASN holders can use this ASN-based prefix.

## 5. ASN allocation

ASNs are allocated by RIRs and this RFC does not handle that arena.

ASNs defined as private ASNs MUST NOT use this scheme. The special 16-bit ASN 23456 MUST NOT use this scheme.

## 6. BGP Filtering and Validation

Filtering would be a simple case of mapping the final ASN in a path to the prefix in an exact bit-order match.

For example; the prefix NNNN:0000:1B1B::/48 should only be seen as announced from AS6939 (6939 equals 0x1B1B in hex). Networks would have their upstream transit providers add this /48 prefix to their existing inbound BGP route filters.

## 7. Whois

In order for the IPv6 /48 to have valid whois information RIRs will have to map it from their existing ASN whois data. The whois data for the /48 should be the same as the ASN whois data.

## 8. RIR support

It is suggested that RIRs, via their Internet Registry services, support these automatic assignments (including but not limited to whois and reverse DNS). The RIRs should provide these services in addition to the services already provided for the associated AS number.

## 9. Reverse DNS

Reverse DNS is vital to the operation of the global Internet. Any block allocated via this mechanism should include the ability to have the network operator provide reverse DNS entries.

It is proposed that the first /64 from the allocated /48 have the rDNS services hard coded into the ::1 and ::2 addresses. This allows reverse DNS to operate without intervention or RIR involvement.

Nameserver 1: NNNN:####:####:0000::1

Naneserver 2: NNNN:####:####:0000::2

## 10. Routing Table Impact

This mechanism is not expected to have any impact to the global Internet routing table since existing policies in the RIR system already readily provide for the allocation of provider-independent IPv6 prefixes. Additionally, AS number holders are likely to be multihomed entities which were going to be independently routed in any case. Service Providers are, as always, not obligated to route these IPv6 assignments and/or may establish conditions of service which offset any additional routing cost.

## 11. IANA Considerations

This memo includes a request to IANA to allocate a /16 from the available IPv6 address space.

## 12. Security Considerations

There is clearly a need for RPKI ROAs for all allocated ASNs within an RIR. The mapping would be 1:1 from ASN to /48 prefix. Hence an RIR allocating an ASN should also allow the holder of that ASN to manage this IPv6 prefix via their RIR portal.

There is nothing different in with a prefix from this space than from a prefix allocated by an RIR.

## 13. Acknowledgements

We would like to thank the following people. ????

We would also like to also thank the contributions from people in the industry that have promoted IPv6 and multihoming: Bill Manning (as himself) and Randy Bush (IIJ).

## 14. References

### 14.1. Normative References

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.

- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", [RFC 3056](#), February 2001.
- [RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", [RFC 3068](#), June 2001.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.

[RFC4893] Vohra, Q. and E. Chen, "BGP Support for Four-octet AS Number Space", [RFC 4893](#), May 2007.

#### [14.2](#). Informative References

[BARBER2011]

Barber, S., "IPv6 in the Real World: Practical Multihoming", February 2011, <<http://www.stanbarber.com/rpsl/ipv6-in-the-real-world-practical-multihoming>>.

#### Authors' Addresses

Martin J. Levy  
Hurricane Electric  
760 Mission Court  
Fremont, CA 94359  
US

Phone: +1 510 580-4100  
Email: [martin@he.net](mailto:martin@he.net)  
URI: <http://he.net/>

Matthew Pounsett  
Afilias  
4141 Yonge St., Suite 204  
Toronto, Ontario M2P 2A8  
CA

Phone: +1 416 646-3304  
Email: [mpounsett@afilias.info](mailto:mpounsett@afilias.info)  
URI: <http://www.afilias.info/>