

Internet Engineering Task Force  
Internet-Draft  
Intended status: Informational  
Expires: August 28, 2017

Shi, Ed.  
Yang  
Wang  
Wu  
Yin  
Tsinghua Univ.  
February 24, 2017

**Classless AS-based Addressing**  
**draft-shi-v6ops-caba-01**

Abstract

Addressing is the foundation of routing and it impacts the routing scalability. The IPv6 routing system will face severe scalability issue if the IPv6 addresses are recklessly used without careful aggregation strategy. This draft describes an addressing scheme and the corresponding allocation scheme which can facilitate the aggregation of IPv6 addresses and keep the IPv6 inter-domain routing system from being overloaded.

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 August 28, 2017.

Copyright Notice

Copyright (c) 2017 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 (<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="#">2</a>
<a href="#">2.</a>	Address format . . . . .	<a href="#">2</a>
<a href="#">3.</a>	ASN allocation . . . . .	<a href="#">3</a>
<a href="#">4.</a>	effects on the size of IPv6 inter-domain routing tables . . .	<a href="#">4</a>
<a href="#">5.</a>	IANA Considerations . . . . .	<a href="#">5</a>
<a href="#">6.</a>	Security Considerations . . . . .	<a href="#">5</a>
<a href="#">7.</a>	Conclusions . . . . .	<a href="#">5</a>
<a href="#">8.</a>	Normative References . . . . .	<a href="#">5</a>
	Authors' Addresses . . . . .	<a href="#">5</a>

## [1.](#) Introduction

IP address fragment is one of the most contributing factors that cause scalability problems to the IPv4 inter-domain routing system. The IPv6 routing system will face severe scalability issue if the IPv6 addresses are recklessly used without careful aggregation strategy. This document propose Classless AS-Based Addressing (CABA) scheme, which embeds 32 bits AS number (ASN) into the IPv6 address. If ASNs are properly allocated, CABA can facilitate aggregation of IP addresses and reduce the size of inter-domain routing tables.

## [2.](#) Address format

The Classless AS-Based address has five fields of fixed lengths, as shown in Figure 1.

```
+-----+-----+-----+-----+-----/ /-----+
| IANA /8 | ASN Low /16 | ASN High /16 | Subnet /24 | Interface /64 |
+-----+-----+-----+-----+-----/ /-----+
```

Figure 1: Classless AS-Based address format

IANA /8 field should be allocated by IANA.

ASN Low /16 is the lower 16 bits of the AS number.

ASN High /16 is the higher 16 bits of the AS number.

Subnet /24 should be allocated by individual ASes according to their intra-domain routing policies.



Interface /64 is interface ID.

All ASNs are assumed as 32 bits long. For a legacy ASN of 16 bits, the higher 16 bits are filled with 0. The document uses AS x.y to denote an ASN whose higher part is x and lower part is y. For example, AS 199081 (0x309a9 in hex) is denoted as AS 3.2473, and the ASN Low /16 field and the ASN High /16 field are 0x09a9 and 0x0003 respectively, as shown in Figure 2.

ASN in decimal	ASN in hexadecimal	IPv6 Block
AS 0.2473	0x9a9	[xx]09:a900::/40
AS 3.2473	0x309a9	[xx]09:a900:300::/40

Figure 2: ASN and IPv6 Block examples

In order to prevent routing scalability issues, any prefixes that are more specific than the allocated /40 should not be advertised into global routing system. Meanwhile, IPv6 operator must filter BGP updates that contain information of more specific prefixes.

### 3. ASN allocation

To take advantage of this AS-based addressing scheme, this document suggests to allocate ASNs in a way similar to provider-aggregatable (PA) addresses, where each allocation is made by ASes instead of RIRs.

Initially, a legacy 16-bit ASN is regarded as holding 65536 32-bit ASNs. For example, AS 333 holds ASNs from 0.333 to 65535.333. Existing 32-bit ASNs are regarded as holding only one 32-bit ASN and it is not hold by any other AS. For example, if 2.333 exists, AS 333 holds ASNs 0.333, 1.333 and from 3.333 to 65535.333. And each AS can use only one ASN for routing purpose and this ASN is the smallest among all ASNs it holds.

When a new AS wants to obtain an ASN, it should choose a provider who holds several ASNs and ask for a unrouted ASN. The provider should response by allocating an ASN or a range of ASNs from its ASN pool according to its policy. Each AS can have its own allocation policy but must conform to the following guidelines:

1. The higher part of allocated ASNs should be successive.
2. The higher part of remained ASNs should be as successive as possible (There might be holes due to historical reasons).



For example, if AS 0.333 allocates 4 ASNs to its customer, it should allocate ASNs from 65532.333 to 65535.333.

Therefore, an ASN can only decide the number of ASNs in each allocation. As examples, there might be two kind of policies:

1. Allocating a fixed ratio of all ASNs currently hold by the AS.
2. Allocating a fixed number of ASNs.

For example, an AS who has ASNs from 0.100 to 65535.100 may choose to allocate 1/4 of its ASNs (49152.100~65535.100) or 256 ASNs (65280.100~65535.100). An AS can define its own allocation policy provided the policy conforms to the above guidelines.

According to the allocation strategy above, we encode the ASN's lower bits first. Thus ASNs who have the same lower part share a common prefix. As Figure 2 shows, AS 0.2473 and AS 3.2473 share the same prefix [xx]09:a900::/32 and their IPv6 blocks are more likely to be aggregated by ASN allocator AS 0.2473. (xx denote the bits provided by IANA). Therefore, any prefixes that are more specific than the corresponding AS-Based prefix should not be advertised into global routing system. As a result, IPv6 operator must filter the anomalous BGP updates which contain information of more specific prefixes.

#### **4. effects on the size of IPv6 inter-domain routing tables**

For core routers in the Internet, their inter-domain routing tables are usually large and complex. routing tables list the routes to particular network prefixes and some metrics associated with those routes. As there are so many prefixes spreading around the current Internet, core routers' inter-domain routing tables become larger, causing routing system overloaded.

In the current Internet, an AS usually announces several prefixes to its neighbors. CABA scheme embeds 32-bit ASN into its IPv6 address. If CABA is adopted in the Internet, the number of spreading prefixes will reduce obviously. One AS deploying CABA can only announce one prefix to other networks. Thus, the size of inter-domain routing tables will never be more than the number of ASes. Figure 3 shows the result of our simulation. The routes from inter-domain routing tables are collected by Routeviews in August, 2016. We choose different proportions of ASes in the Internet randomly to adopt CABA scheme and count the size of inter-domain routing tables.



+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+
adoption rate	0	20	40	60	80	100
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+
RIB size	634050	515201	393595	280183	173500	53483
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+

Figure 3: CABA deployment influences sizes of inter-domain routing tables

In addition, since the allocation of ASNs facilitates aggregation of IP addresses (providers can aggregate classless AS-Based addresses which they allocate to customers), the size of routing tables will reduce further.

## 5. IANA Considerations

IANA should allocate a /8 IPv6 address block for this purpose, which is similar to 2002::/8 [RFC3068] used by 6to4. No IPv4 allocation required.

## 6. Security Considerations

This document includes no request to security.

## 7. Conclusions

This draft describes an addressing scheme CABA (Classless AS-Based Addressing) and an address allocation scheme which can facilitate aggregation of addresses. CABA can solve the scalability problem of IP address and reduce the size of core RIBs obviously.

## 8. Normative References

[RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", [RFC 3068](https://www.rfc-editor.org/rfc/rfc3068), DOI 10.17487/RFC3068, June 2001, <<http://www.rfc-editor.org/info/rfc3068>>.

### Authors' Addresses

Xingang Shi (editor)  
Tsinghua Univ.  
Beijing  
CN

Email: shixg@cernet.edu.cn



Yan Yang  
Tsinghua Univ.  
Beijing  
CN

Email: yangyan15@mails.tsinghua.edu.cn

Zhiliang Wang  
Tsinghua Univ.  
Beijing  
CN

Email: wzl@csnet1.cs.tsinghua.edu.cn

Jianping Wu  
Tsinghua Univ.  
Beijing  
CN

Email: jianping@csnet1.cs.tsinghua.edu.cn

Xia Yin  
Tsinghua Univ.  
Beijing  
CN

Email: yxia@csnet1.cs.tsinghua.edu.cn