

Proposed TLA and NLA Assignment Rules

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

1.0 Introduction

This document proposes rules for Top-Level Aggregation Identifiers (TLA ID) and Next-Level Aggregation Identifiers (NLA ID) as defined in [AGGR]. These proposed rules apply to registries allocating TLA ID's and to organizations receiving TLA ID's.

This proposal is intended as input from the IPng working group to the IANA and Registries. It is not intended for any official IETF status. Its content represents the result of extensive discussion between the IPng working group, IANA, and Registries.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119].

2.0 Scope



The proposed TLA and NLA assignment rules described in this document are intended for the first two years of IPv6 TLA address assignments. As routing technology evolves and we gain additional experience with allocating IPv6 addresses the procedures proposed in this document may change.

3.0 IPv6 Aggregatable Global Unicast Address Format

This document proposes assignment rules for the TLA ID and NLA ID fields in the IPv6 Aggregatable Global Unicast Address Format. This address format is designed to support both the current provider-based aggregation and a new type of exchange-based aggregation. The combination will allow efficient routing aggregation for sites that connect directly to providers and for sites that connect to exchanges. Sites will have the choice to connect to either type of aggregation entity.

While this address format is designed to support exchange-based aggregation (in addition to current provider-based aggregation) it is not dependent on exchanges for its overall route aggregation properties. It will provide efficient route aggregation with only provider-based aggregation.

The aggregatable global unicast address format as defined in [AGGR] is as follows:



Where

- FP Format Prefix (001)
- TLA ID Top-Level Aggregation Identifier
- RES Reserved for future use
- NLA ID Next-Level Aggregation Identifier
- SLA ID Site-Level Aggregation Identifier
- INTERFACE ID Interface Identifier

4.0 Technical Motivation

The design choices for the size of the fields in the aggregatable address format were based on the need to meet a number of technical requirements that are described in [AGGR]. An extract of the technical requirements from [AGGR] is as follows:

The size of the Top-Level Aggregation Identifier is 13 bits. This allows for 8,192 TLA ID's. This size was chosen to insure that the default-free routing table in top level routers in the Internet is kept within the limits, with a reasonable margin, of the current routing technology. The margin is important because default-free routers will also carry a significant number of longer (i.e., more-specific) prefixes for optimizing paths internal to a TLA and between TLAs.

The important issue is not only the size of the default-free routing table, but the complexity of the topology that determines the number of copies of the default-free routes that a router must examine while computing a forwarding table. In current practice with IPv4, it is common to see a prefix announced fifteen times via different paths. The complexity of Internet topology is very likely to increase in the future. It is important that IPv6 default-free routing support additional complexity as well as a considerably larger internet.

It should be noted for comparison that the current IPv4 default-free routing table is approximately 50,000 prefixes. While this shows that it is possible to support more routes than 8,192 it is matter of debate if the number of prefixes supported today in IPv4 is already too high for current routing technology. There are serious issues of route stability as well as cases of providers not supporting all top level prefixes. The technical requirement was to pick a TLA ID size that was below, with a reasonable margin, what was being done with IPv4.

The choice of 13 bits for the TLA field was an engineering compromise. Fewer bits would have been too small by not supporting enough top level organizations. More bits would have exceeded what can be reasonably accommodated, with a reasonable margin, with current routing technology in order to deal with the issues described in the previous paragraphs.

If in the future, routing technology improves to support a larger number of top level routes in the default-free routing tables there are two choices on how to increase the number TLA identifiers. The first is to expand the TLA ID field into the reserved field. This would increase the number of TLA ID's to approximately 2 million. The second approach is to allocate another format prefix (FP) for use with this address format. Either or a combination of these approaches allows the number of TLA ID's to increase significantly.

The size of the Reserved field is 8 bits. This size was chosen to allow significant growth of either the TLA ID and/or the NLA ID fields.

The size of the Next-Level Aggregation Identifier field is 24 bits. This allows for approximately sixteen million NLA ID's if used in a flat manner. Used hierarchically it allows for a complexity roughly equivalent to the IPv4 address space (assuming an average network size of 254 interfaces). If in the future additional room for complexity is needed in the NLA ID, this may be accommodated by extending the NLA ID into the Reserved field.

The size of the Site-Level Aggregation Identifier field is 16 bits. This supports 65,535 individual subnets per site. The design goal for the size of this field was to be sufficient for all but the largest of organizations. Organizations which need additional subnets can arrange with the organization they are obtaining Internet service from to obtain additional site identifiers and use this to create additional subnets.

The Site-Level Aggregation Identifier field was given a fixed size in order to force the length of all prefixes identifying a particular site to be the same length (i.e., 48 bits). This facilitates movement of sites in the topology (e.g., changing service providers and multi-homing to multiple service providers).

The Interface ID Interface Identifier field is 64 bits. This size was chosen to meet the requirement specified in [ARCH] to support EUI-64 based Interface Identifiers.

The proposed TLA/NLA assignment rules described in this document are consistent with these technical requirements.

The specific technical motivation for the proposed TLA/NLA assignment rules described in this document is as follows:

- Limit the number of top level prefixes in the Internet to a manageable size. This is important to insure that the default-free routing table in the top level routers in the Internet is kept within the limits, with a reasonable margin, of current routing technology.
- Only assign top level prefixes to transit providers, not to leaf sites even if they are multiply homed. The aggregation address format is designed to have a clear separation between transit providers and leaf sites. Sites which wish to be multihomed to multiple transit providers have in IPv6 a number of alternatives to having a top level prefix.

- Only assign top level prefixes to organizations who are capable and intend to provide operational IPv6 transit services within three months of assignment. The goal is to not assign top level prefixes to organizations who only want a prefix in case they might provide service sometime in the future. The assignment of prefixes is intended to closely match the operational IPv6 Internet and to be consistent with the current practice of registries making assignments when addresses are actually used.
- Organizations assigned TLA ID's are required to make all the assignments publically available. This is necessary in order for the registries to have accurate information on assignments and to enable trouble shooting Internet problems.
- Allocation of prefixes that are consistent with the address format in [AGGR]. Specifically the allocation prefixes that are not longer than 48 bits as to not infringe into the SLA and Interface Identifier fields. This is to facilitate movement of sites in the topology (e.g., changing service providers and multi-homing to multiple service providers).

5.0 Proposed Rules for Assignment of Top-Level Aggregation ID's

TLA ID's are assigned to organizations providing transit topology. They are specifically not assigned to organizations only providing leaf topology. TLA ID assignment does not imply ownership. It does imply stewardship over a valuable Internet resource.

The IAB and IESG have authorized the Internet Assigned Numbers Authority (IANA) as the appropriate entity to have the responsibility for the management of the IPv6 address space as defined in [ALLOC].

The IANA will assign small blocks (e.g., few hundred) of TLA ID's to registries. The registries will assign the TLA ID's to organizations meeting the requirements for TLA ID assignment. When the registries have assigned all of their TLA ID's they can request that the IANA give them another block. The blocks do not have to be contiguous. The IANA may also assign TLA ID's to organizations directly. This includes the temporary TLA assignment for testing and experimental usage for activities such as the 6bone or new approaches like exchanges.

5.1 Proposed TLA Allocation Stages

TLA allocations will be done in two stages. The first stage is to allocate a Sub-TLA ID. When the recipient has demonstrated that they have assigned more than 90% of the NLA ID for their Sub-TLA ID, they will be allocated a TLA ID. The Sub-TLA ID does not have to be returned.

Sub-TLA ID's are assigned out of TLA ID 0x0001 as follows. Note that use of the Reserved field to create the Sub-TLA field is specific to TLA ID 0x0001. It does not affect any other TLA.

3	13	13	19
FP	TLA ID	Sub-TLA	NLA ID

where:

FP = 001 = Format Prefix

This is the Format Prefix used to identify aggregatable global unicast addresses.

TLA ID = 0x0001 = Top-Level Aggregation Identifier

This is the TLA ID assigned by the IANA for Sub-TLA allocation.

Sub-TLA ID = Sub-TLA Aggregation Identifier

The Sub-TLA ID field is used by the registries for initial allocations to organizations meeting the requirements in Section 5.2 of this document. The IANA will assign small blocks (e.g., few hundred) of Sub-TLA ID's to registries. The registries will assign the Sub-TLA ID's to organizations meeting the requirements specified in Section 5.2. When the registries have assigned all of their Sub-TLA ID's they can request that the IANA give them another block. The blocks do not have to be contiguous. The

IANA may also assign Sub-TLA ID's to organizations directly. This includes the temporary TLA assignment for testing and experimental usage for activities such as the 6bone or new approaches like exchanges.

NLA ID = Next-Level Aggregation Identifier

Next-Level Aggregation ID's are used by organizations assigned a TLA ID to create an addressing hierarchy and to identify sites. The organization can assign the top part of the NLA ID in a manner to create an addressing hierarchy appropriate to its network. See Section 6.0 for more detail.

Sub-TLA allocations are interim until the organization receiving the Sub-TLA can show evidence of IPv6 Internet transit service. If transit service can not be demonstrated by three months from the date of allocation the Sub-TLA allocation will be revoked.

As part of assigning a TLA ID to an organization, the IANA or Registries may initially only assign a fraction of the NLA ID space for a particular TLA ID to the organization receiving the TLA ID assignment. When the organization has assigned more than 90% of the NLA ID space it may request additional NLA ID space in its TLA ID.

5.2 Proposed Assignment Requirements

The proposed assignment requirements are intended as input from the IPng working group to the IANA and Registries. It is not intended for any official IETF status.

Registries enforce the following requirements for organizations assigned Sub-TLA and TLA ID's:

- 1) Must have a plan to offer native IPv6 service within 3 months from assignment. The plan must include NLA ID allocation and registration procedures. NLA ID allocation and registration may be subcontracted to other organizations such as a registry.

Native IPv6 service is defined as providing IPv6 service as defined in the appropriate "IPv6 over <link>" specification such as "IPv6 over Ethernet" [ETHER], "IPv6 over FDDI" [FDDI], etc., for the link at the boundary of the organization. This should include running Neighbor Discovery (as appropriate) and exchanging IPv6 routing information. The method the organization uses to carry IPv6 traffic across its network is independent of this definition and is a local issue for the organization.

- 2) Must have a verifiable track record of providing Internet transit to other organizations. Sub-TLA and/or TLA ID's must not be assigned to organizations that are only providing leaf service even if multihomed.

Verification of an organization's track record in providing Internet transit service must be verified by techniques such as traceroute, BGP advertisements, etc.

- 3) Payment of a registration fee to the Internet Assigned Numbers Authority (IANA). Registries may also charge some fee for services rendered, generally in relation to the cost of providing those services. All payment of registration and service fees must be made prior to the actual Sub-TLA ID and/or TLA ID assignment.
- 4) Must provide registry services for the NLA ID address space it is responsible for under its Sub-TLA ID and/or TLA ID. This must include both sites and next level providers. The database of NLA assignments must be public and made available to the registries.
- 5) Periodically (interval set by registry) provide to registry utilization statistics of the Sub-TLA ID and/or TLA ID it has custody of. The organization must also show evidence of carrying TLA routing and transit traffic. This can be in the form of traffic statistics, traceroutes, routing table dumps, or similar means.
- 6) Organizations requesting another Sub-TLA and/or TLA ID must show evidence to the registries that they have assigned more than 90% of the NLA ID space in their previous allocations.

Organizations which are given custody of a Sub-TLA ID and/or TLA ID, and fail to continue to meet all the above requirements may have the Sub-TLA ID and/or TLA ID custody revoked.

6.0 Proposed Rules Assignment of Next-Level Aggregation ID's

Next-Level Aggregation ID's are used by organizations assigned a Sub-TLA ID and/or TLA ID to create an addressing hierarchy and to identify sites. The organization can assign the top part of the NLA ID in a manner to create an addressing hierarchy appropriate to its network.

Registries may initially only assign a fraction of the NLA ID space for a particular Sub-TLA ID and/or TLA ID to the organization receiving the Sub-TLA ID and/or TLA ID assignment. When the organization has assigned more than 90% of the NLA ID space it may request additional NLA ID space in its Sub-TLA ID and/or TLA ID.

Organizations assigned Sub-TLA ID and/or TLA ID's are required to assume (directly or indirectly) registry duties for the NLA ID's they assign. Each organization assigned a NLA ID is required to assume registry duties for the next level NLA ID's it assigns and follow

Registry guidelines. This responsibility includes passing this information back to the registry that assigned the TLA and/or Sub-TLA. The TLA ID and/or Sub-TLA ID holder collects this information from the next level, the next level holder collects this information from the level below, etc.

The design of the bit layout of the NLA ID space for a specific Sub-TLA ID and/or TLA ID is left to the organization responsible for that Sub-TLA ID and/or TLA ID. Likewise the design of the bit layout of the next level NLA ID is the responsibility of the organization assigned the previous level NLA ID. It is recommended that organizations assigning NLA address space use "slow start" allocation procedures as is currently done with IPv4 CIDR blocks [CIDR].

The design of an NLA ID allocation plan is a tradeoff between routing aggregation efficiency and flexibility. Creating hierarchies allows for greater amount of aggregation and results in smaller routing tables. Flat NLA ID assignment provides for easier allocation and attachment flexibility, but results in larger routing tables.

7.0 Acknowledgments

The author would like to express his thanks to Thomas Narten, Steve Deering, Bob Fink, Matt Crawford, Rebecca Nitzan, Allison Mankin, Jim Bound, Christian Huitema, Scott Bradner, Brian Carpenter, John Stewart, Eric Hoffman, Jon Postel, Daniel Karrenberg, Kim Hubbard, Mirjam Kuehne, Paula Caslav, David Conrad, and David Kessens for their review and constructive comments.

8.0 Security Considerations

IPv6 addressing documents do not have any direct impact on Internet infrastructure security. Authentication of IPv6 packets is defined in [AUTH]. Authentication of the ownership of prefixes to avoid "prefix stealing" is a related security issue but is beyond the scope of this document.

9.0 References

- [AGGR] Hinden, R., Deering, S. and M. O'Dell, "An Aggregatable Global Unicast Address Format", RFC 2374, July 1998.
- [ALLOC] IAB and IESG, "IPv6 Address Allocation Management", RFC 1881, December 1995.
- [ARCH] Hinden, R., "IP Version 6 Addressing Architecture", RFC 2373, July 1998.

- [AUTH] Atkinson, R. and S. Kent, "IP Authentication Header", RFC 2402, November 1998.
- [CIDR] Fuller, V., Li, T., Varadhan, K. and J. Yu, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.
- [ETHER] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [FDDI] Crawford, M., "Transmission of IPv6 Packets over FDDI Networks", RFC 2467, December 1998.
- [IPV6] Deering, S. and R. Hinden, Editors, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

10.0 Author's Address

Robert M. Hinden
Nokia
232 Java Drive
Sunnyvale, CA 94089
USA

Phone: +1 408 990-2004
EMail: hinden@iprg.nokia.com

11.0 Full Copyright Statement

Copyright (C) The Internet Society (1998). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

