

Network Working Group
Internet Draft
Intended status: Informational
Expires: August 18, 2014

B. Liu
S. Jiang
Huawei Technologies
February 14, 2014

Recommendations of Using Unique Local Addresses
draft-ietf-v6ops-ula-usage-recommendations-02.txt

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 18, 2014.

Copyright Notice

Copyright (c) 2014 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.

Abstract

This document provides guidance of how to use ULAs. It analyzes ULA usage scenarios and recommends use cases where ULA addresses might be beneficially used.

Table of Contents

1.	Introduction	3
2.	The analysis of ULA features	3
2.1.	Automatically Generated	3
2.2.	Globally unique.....	3
2.3.	Independent address space	4
2.4.	Well known prefix	4
2.5.	Stable or Temporary Prefix	4
3.	Enumeration of Scenarios Using ULAs	4
3.1.	Isolated network	4
3.2.	Connected network	5
3.2.1.	ULA-only Deployment	5
3.2.2.	ULA along with GUA	7
3.3.	IPv4 Co-existence consideration	9
4.	General Guidelines of using ULA	9
4.1.	Do not treat ULA equal to RFC1918	9
4.2.	Using ULAs in a limited scope	10
5.	ULA usage recommendations	10
5.1.	Used in Isolated Networks	10
5.2.	ULA along with GUA	10
5.3.	Recommended Specific Use Cases	10
5.3.1.	Special routing	11
5.3.2.	Used as NAT64 prefix	11
5.3.3.	Used as identifier	12
6.	Security Considerations	12
7.	IANA Considerations	13
8.	Conclusions	13
9.	References	13
9.1.	Normative References	13
9.2.	Informative References	13
10.	Acknowledgments	14

1. Introduction

Unique Local Addresses (ULAs) are defined in [[RFC4193](#)] as provider-independent prefixes that can be used on isolated networks, internal networks, and VPNs. Although ULAs may be treated like global scope by applications, normally they should not be used on the publicly routable internet.

However, the ULAs haven't been widely used since IPv6 hasn't been widely deployed yet. The use of ULA addresses in various types of networks has been confusing to network operators. This document aims to clarify the advantages and disadvantages of ULAs and how they can be most appropriately used.

2. The analysis of ULA features

2.1. Automatically Generated

ULA prefixes could be automatically generated according to the algorithms described in [[RFC4193](#)]. This feature allows automatic address allocation, which is beneficial for some lightweight systems and can leverage minimal human management.

2.2. Globally unique

ULA is intended to have an extremely low probability of collision. Since the hosts assigned with ULA may occasionally be merged into one network, this uniqueness is necessary. The prefix uniqueness is based on randomization of 40 bits and is considered random enough to ensure a high degree of uniqueness (refer to [[RFC4193](#)] [section 3.2.3](#) for details) and make merging of networks simple and without the need to renumbering overlapping IP address space. Overlapping is cited as a deficiency with how [[RFC1918](#)] addresses were deployed, and ULA was designed to overcome this deficiency.

Notice that, as described in [[RFC4864](#)], in practice, applications may treat ULAs like global-scope addresses, but address selection algorithms may need to distinguish between ULAs and ordinary GUA (Global-scope Unicast Address) to ensure bidirectional communications. (Note: the new address selection standard has supported this in the default policy table. [[RFC6724](#)])

2.3. Independent address space

ULA provides an internal address independence capability in IPv6ULA can be used for internal communications without having any permanent or only intermittent Internet connectivity. And it needs no registration so that it can support on-demand usage and does not carry any RIR documentation burden or disclosures.

2.4. Well known prefix

The prefixes of ULAs are well known and they are easy to be identified and easy to be filtered.

This feature may be convenient to management of security policies and troubleshooting. For example, the administrators can decide what parameters have to be assembled or transmitted globally, by a separate function, through an appropriate gateway/firewall, to the Internet or to the telecom network.

2.5. Stable or Temporary Prefix

A ULA prefix can be generated once, at installation time or "factory reset", and then never change unless the network manager wants to change. Alternatively, it could be regenerated regularly, if desired for some reason.

3. Enumeration of Scenarios Using ULAs

In this section, we cover possible ULA use cases. Some of them might have been discussed in other documents and are briefly reviewed in this document as well as other potential valid usage is discussed.

3.1. Isolated network

IP is used ubiquitously. Some networks like industrial control bus (e.g. [[RS-485](#)], [[SCADA](#)], or even non-networked digital interface like [[MIL-STD-1397](#)] began to use IP protocol. In such kind of networks, the system might lack the ability/requirement of connecting to the Internet. or explicitly designed not to connect.

Besides, there might be some networks which could connect to the Internet, but prohibited by administration or just temporally not connected. These networks may include machine-to-machine (e.g. vehicle networks), sensor networks, or even normal LANs, which may include very large numbers of addresses.

Since the serious disadvantages and impact on applications by ambiguous address space have been well documented in [[RFC1918](#)], ULA is a straightforward way to assign the IP addresses in these kinds of networks with minimal administrative cost or burden. Also, ULAs fit in multiple subnet scenarios, in which each subnet has its own ULA prefix. For example, when we assign vehicles with ULA addresses, it is then possible to separate in-vehicle embedded networks into different subnets depending on real-time requirements, devices types, services and more.

- o Benefits of Using ULAs in Isolated Networks:
 - No cost of RIR/LIR fees or operational burden
 - Could be used instantly on-demand
 - Scalable to support multiple subnets
 - Extremely low probability of collision when isolated networks merge
- o Drawbacks:
 - The uniqueness of the global ULA prefix is not guaranteed, however, the probability of collision is extremely low
- o Operational considerations
 - Prefix generation: Randomly generated according to the algorithms defined in [[RFC4193](#)] or literally assigned by human. Normally, it is recommended to following the standard way to automatically generate the prefixes; if there are some specific reasons that need to be assigned by human, the administrators must carefully plan the prefixes to avoid collision.
 - Prefix announcement: In some cases, the networks might need to announce prefixes to each other, for example in vehicle networks with infrastructure-less settings such as Vehicle-to-Vehicle (V2V) settings, prior knowledge of respective prefixes is unlikely. Hence, a prefix announcement mechanism is needed to enable inter-vehicles communications based on IP. For instance, such announcement could rely on an extension of the Router Advertisement message of Neighbor Discovery Protocol (e.g. [[I-D.petrescu-autoconf-ra-based-routing](#)] and [[I-D.jhlee-mext-mnpp](#)]).

3.2. Connected network

3.2.1. ULA-only Deployment

In some situations, hosts/interfaces are assigned with ULA-only, but the networks need to communicate with the outside. It mainly includes the following two models.

- o Using Network Prefix Translation

Network Prefix Translation (NPTv6) [[RFC6296](#)] is an experimental specification that provides a stateless one to one mapping between internal addresses and external addresses.

In some very constrained situations(for example, in the sensors), the network needs ULA as the on-demand and stable addressing which doesn't need much code to support address assignment mechanisms like DHCP or full ND (Note: surely it needs SLAAC). If the network also needs to connect to the outside, then there can be an NPTv6 gateway which is not subject to extreme resource constraints. Especially when a lightweight isolated network needs to add Internet connectivity, this is quite a straightforward and efficient way.

This document does not intend to encourage the ULA binding with NPTv6 model, since in [[RFC5902](#)] the IAB had already gave opinions on IPv6 NAT; but this document considers it as an effective approach in some specific situations as described above.

- o Using application-layer proxies

The proxies terminate the network-layer connectivity of the hosts and delegate the outgoing/incoming connections.

In some environments (e.g. information security sensitive enterprise or government), the endpoints are default disconnected to the Internet, and need the proxies to connect for central control. In IPv4, using private address space with proxies is an effective and common practice for this purpose, and it is natural to pick ULA in IPv6.

- o Benefits of Using ULAs in This Scenario:

- Allowing minimal management burden on address assignment for some specific environments.

- o Drawbacks:

- The serious disadvantages and impact on applications by NAT have been well documented in [[RFC2993](#)] and [[RFC3027](#)]. However, it should be noted that, For NPTv6, as described in [[RFC6296](#)], it is "a mechanism that has fewer architectural problems than merely implementing a traditional stateful Network Address Translator in an IPv6 environment."

- o Operational considerations
 - Firewall Deployment: As described in [[RFC6296](#)], firewall is recommended to be implemented with NPTv6 translator. Then the administrators need to know about where the firewall is located. If the firewall is located outside the NPTv6 translator, then the filtering is based on the translated GUA prefixes, and when the internal ULA prefixes are renumbered, the filtering rules don't need to be changed (however, when GUA prefixes of the NPTv6 are renumbered, the filtering rules need to be updated accordingly.). If the firewall is located inside the NPTv6 translator, the filtering is then based on the ULA prefixes, and the rules need to be updated according to the ULA prefixes renumbering (but don't need to update when NPTv6 GUA prefixes are renumbered.).

3.2.2. ULA along with GUA

There are two classes of network probably to use ULA with GUA addresses:

- Home network. Home networks are normally assigned with one or more globally routed PA prefixes to connect to the uplink of some an ISP. And besides, they may need internal routed networking even when the ISP link is down. Then ULA is a proper tool to fit the requirement. And in [[RFC6204](#)], it requires the CPE to support ULA. Note, ULAs provide more benefit for multiple-segment home networks; for home networks only containing one segment, link-local addresses are better alternative.
 - Enterprise network. An enterprise network is usually a managed network with one or more PA prefixes or with a PI prefix, all of which are globally routed. The ULA could be used for internal connectivity redundancy and better internal connectivity or isolation of certain functions like OAM of servers.
- o Benefits of Using ULAs in This Scenario:
 - Separated Local Communication Plane: For either home networks or enterprise networks, the main purpose of using ULA along with GUA is to provide a logically local routing plane separated from the globally routing plane. The benefit is to ensure stable and specific local communication regardless of the ISP uplink failure. This benefit is especially meaningful for the home network or private OAM function in an enterprise.
 - Renumbering: In some special cases such as renumbering, enterprise administrators may want to avoid the need to renumber their internal-only, private nodes when they have to renumber the PA addresses of the whole network because of changing ISPs, ISPs

restructure their address allocations, or whatever reasons. In these situations, ULA is an effective tool for the internal-only nodes.

Besides the internal-only nodes, the public nodes can also benefit from ULA for renumbering. When renumbering, as [RFC4192](#) suggested, it has a period to keep using the old prefix(es) before the new prefix(es) is(are) stable. In the process of adding new prefix(es) and deprecating old prefix(es), it is not easy to keep the local communication immune of global routing plane change. If we use ULA for the local communication, the separated local routing plane can isolate the affecting by global routing change.

o Drawbacks:

- Operational Complexity: There are some arguments that in practice the ULA+PA makes additional operational complexity. It is not a ULA-specific problem; the multiple-addresses-per-interface is an important feature of IPv6 protocol. Never the less, running multiple prefixes needs more operational considerations than running a single one.

o Operational considerations

- Default Routing: connectivity might be broken if ULAs are used as default route.

- SLAAC/DHCPv6 co-existing: Since SLAAC and DHCPv6 might be enabled in one network simultaneously; the administrators need to carefully plan how to assign ULA and GUA prefixes in accordance with the two mechanisms. The administrators need to know the current issue of the SLAAC/DHCPv6 interaction (please refer to [I-D.liu-bonica-dhcpv6-slaac-problems] for details).

- Address Selection: As mentioned in [[RFC5220](#)], there is a possibility that the longest matching rule will not be able to choose the correct address between ULAs and global unicast addresses for correct intra-site and extra-site communication. In [[RFC6724](#)], it claimed that a site-specific policy entry can be used to cause ULAs within a site to be preferred over global addresses.

- DNS Relevant: if administrators chose not to do reverse DNS delegation inside of their local control of ULA prefixes, a significant amount of information about the ULA population might leak to the outside world. Because reverse queries will be made and naturally routed to the global reverse tree, so people will be exposed to the existence of population the ULA uses, not in traffic terms, but in knowledge terms. [[ULA-IN-WILD](#)] provides more detailed situations on this issue.

Administrators may need a split DNS to separate the queries from internal and external for ULA entries and GUA entries.

3.3. IPv4 Co-existence consideration

Generally, this document does not consider IPv4 relevant as in the scope. But regarding ULA, there is a special case that needs to be noticed, which is described in [section 2.2.2 of \[RFC5220\]](#). When an enterprise has IPv4 Internet connectivity but does not yet have IPv6 Internet connectivity, and the enterprise wants to provide site-local IPv6 connectivity, a ULA is the best choice for site-local IPv6 connectivity. Each employee host will have both an IPv4 global or private address and a ULA. Here, when this host tries to connect to an outside node that has registered both A and AAAA records in the DNS, the host will choose AAAA as the destination address and the ULA for the source address according to the IPv6 preference of the default address selection policy. This will clearly result in a connection failure.

Although with Happy Eyeballs [\[RFC6555\]](#) this connection failure problem could be solved, but unwanted timeouts would obviously lower the user experience. One possible approach of eliminating the timeouts is deprecating IPv6 default route and simply configuring a scoped route on hosts (in the context of this document, only configure the ULA prefix routes).

Another alternative is configuring IPv4 preference on the hosts, and not including DNS A records but only AAAA records for the internal nodes in the internal DNS server, then outside nodes have both A and AAAA records could be connected through IPv4 as default and internal nodes could be always connected through IPv6. But since IPv6 preference is default, changing the default in all nodes is not suitable at scale.

4. General Guidelines of using ULA

4.1. Do not treat ULA equal to [RFC1918](#)

ULA and [\[RFC1918\]](#) are similar in some aspects. The most obvious one is as described in [section 2.1.3](#) that ULA provides an internal address independence capability in IPv6 that is similar to how [\[RFC1918\]](#) is commonly used. ULA allows administrators to configure the internal network of each platform the same way it is configured in IPv4. Many organizations have security policies and architectures based around the local-only routing of [\[RFC1918\]](#) addresses and those policies may directly map to ULA [\[RFC4864\]](#).

But it doesn't mean ULA is equal to an IPv6 version of [\[RFC1918\]](#) deployment. [\[RFC1918\]](#) usually combines with NAT/NAPT for global connectivity. But it is not necessarily to combine ULAs with any kind of NAT. People could use ULA for local communications along with global addresses for global communications (see [section 5.2](#)). This is a big advantage brought by default support of multiple-addresses-per-interface feature in IPv6. (People might still have requirement of ULA with NAT, this is discussed in [section 3.2.1](#). But people also should keep in mind that ULA is not intentionally designed for this kind of use case.)

Another important difference is the ability to merge two ULA networks without renumbering (because of the uniqueness), which is a big advantage over [\[RFC1918\]](#).

[4.2](#). Using ULAs in a limited scope

A ULA is by definition a prefix that is never advertised outside a given domain, and is used within that domain by agreement of those networked by the domain.

So when using ULAs in a network, the administrators should clearly set the scope of the ULAs and configure ACLs on relevant border routers to block them out of the scope. And if internal DNS are enabled, the administrators might also need to use internal-only DNS names for ULAs and should split the DNS so that the internal DNS server includes records that are not presented in the external DNS server.

[5](#). ULA usage recommendations

[5.1](#). Used in Isolated Networks

As analyzed in [section 3.1](#), ULA is very suitable for isolated networks. Especially when you have subnets in the isolated networks, ULA is the best choice.

[5.2](#). ULA along with GUA

As the benefits described in [Section 3.2.2](#), using ULA along with GUA to provide a logically separated local plane could benefit to OAM functions and renumbering.

[5.3](#). Recommended Specific Use Cases

Along with the general scenarios, this section provides some specific use cases that could benefit from using ULA.

5.3.1. Special routing

For various reasons the administrators might want to have private routing be controlled and separated from other routing. For example, in the b2b case described in [\[I-D.baker-v6ops-b2b-private-routing\]](#), two companies might want to use direct connectivity that only connects stated machines, such as a silicon foundry with client engineers that use it. A ULA provides a simple way to assign such prefixes that would be used in accordance with an agreement between the parties.

5.3.2. Used as NAT64 prefix

Since the NAT64 pref64 is just a group of local fake addresses for the DNS64 to point traffic to a NAT64, When using a ULA prefix as the pref64, it could easily ensures that only local systems can use the translation resources of the NAT64 system since the ULA is not intended to be globally routable and helps clearly identify traffic that is locally contained and destined to a NAT64. Using ULA for Pref64 is deployed and it is an operational model.

But there's an issue should be noticed. The NAT64 standard [RFC6146] mentioned the pref64 should align with [\[RFC6052\]](#), in which the IPv4-Embedded IPv6 Address format was specified. If we pick a /48 for NAT64, it happened to be a standard 48/ part of ULA (7bit ULA famous prefix+ 1 "L" bit + 40bit Global ID). Then the 40bit of ULA is not violated to be filled with part of the 32bit IPv4 address. This is important, because the 40bit assures the uniqueness of ULA, if the prefix is shorter than /48, the 40bit would be violated, and this may cause conformance issue. But it is considered that the most common use case will be a /96 PREF64, or even /64 will be used. So it seems this issue is not common in current practice.

It is most common that ULA Pref64 will be deployed on a single internal network, where the clients and the NAT64 share a common internal network. ULA will not be effective as Pref64 when the access network must use an Internet transit to receive the translation service of a NAT64 since the ULA will not route across the internet.

According to the default address selection table specified in [\[RFC6724\]](#), the host would always prefer IPv4 over ULA. This could be a problem in NAT64-CGN scenario as analyzed in Section 8 of [\[I-D.ietf-v6ops-nat64-experience\]](#). So administrators need to add additional site-specific address selection rules to the default table to steer traffic flows going through NAT64-CGN. However, updating the default policy tables in all hosts involves significant management cost. This may be possible in an enterprise (using a group policy

object, or other configuration mechanisms), but it is not suitable at scale for home networks.

5.3.3. Used as identifier

ULAs could be self-generated and easily grabbed from the standard IPv6 stack. And ULAs don't need to be changed as the GUA prefixes do. So they are very suitable to be used as identifiers by the up layer applications. And since ULA is not intended to be globally routed, it is not harmful to the routing system.

Such kind of benefit has been utilized in real implementations. For example, in [[RFC6281](#)], the protocol BTMM (Back To My Mac) needs to assign a topology-independent identifier to each client host according to the following considerations:

- o TCP connections between two end hosts wish to survive in network changes.
- o Sometimes one needs a constant identifier to be associated with a key so that the Security Association can survive the location changes.

It should be noticed again that in theory ULA has the possibility of collision. However, the probability is desirable small enough and could be ignored by most of the cases when used as identifiers.

6. Security Considerations

Security considerations regarding ULAs, in general, please refer to the ULA specification [[RFC4193](#)].

Also refer to [[RFC4864](#)], which shows how ULAs help with local network protection.

As mentioned in [Section 3.2.1](#), when use NPTv6, the administrators need to know about where the firewall is located to set proper filtering rules.

As mentioned in [Section 3.2.2](#), if administrators chose not to do reverse DNS delegation inside their local control of ULA prefixes, a significant amount of information about the ULA population might leak to the outside world.

7. IANA Considerations

IANA considerations should be updated to point to [RFC4193](#) in a similar manner to [section 4](#).

8. Conclusions

ULA is a useful tool; it could be successfully deployed in a diverse set of circumstances including large private machine-to-machine type networks, enterprise networks with private systems, and within service providers to limit Internet communication with non-public services such as caching DNS servers and NAT64 translation resources. Some of the deployment is already in real production networks.

We should eliminate the misunderstanding that ULA is just an IPv6 version of [\[RFC1918\]](#). The features of ULA could be beneficial for various use cases.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), [BCP14](#), March 1997.
- [RFC4193] Hinden, R., B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), October 2005.

9.2. Informative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", [BCP 5](#), [RFC 1918](#), February 1996.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", [RFC 4864](#), May 2007.
- [RFC5220] Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Problem Statement for Default Address Selection in Multi-Prefix Environments: Operational Issues of [RFC 3484](#) Default Rules", [RFC 5220](#), July 2008.
- [RFC6281] Cheshire, S., Zhu, Z., Wakikawa, R., and L. Zhang, "Understanding Apple's Back to My Mac (BTMM) Service", [RFC 6281](#), June 2011.

- [RFC6296] Wasserman, M., and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", [RFC 6296](#), June 2011.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", [RFC 6555](#), April 2012.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", [RFC 6724](#),
- [I-D.ietf-v6ops-nat64-experience]
Chen G., Cao Z., Xie C., and D. Binet, "NAT64 Operational Experiences", Working in Progress, October, 2013
- [I-D.baker-v6ops-b2b-private-routing]
F. Baker, "Business to Business Private Routing", Expired
- [I-D.petrescu-autoconf-ra-based-routing]
Petrescu, A., Janneteau, C., Demailly, N. and S. Imadali, "Router Advertisements for Routing between Moving Networks", Working in Progress, October, 2013
- [I-D.jhlee-mext-mnpp]
Lee, J.-H., and T. Ernst, "Mobile Network Prefix Provisioning", Expired
- [RS-485] <http://en.wikipedia.org/wiki/RS-485>
- [MIL-STD-1397]
<http://en.wikipedia.org/wiki/MIL-STD-1397>
- [SCADA] <http://en.wikipedia.org/wiki/SCADA>
- [ULA-IN-WILD]
G. Michaelson, "www.ietf.org/proceedings/87/slides/slides-87-v6ops-0.pdf"

10. Acknowledgments

Many valuable comments were received in the IETF v6ops WG mail list, especially from Fred Baker, Brian Carpenter, Lee Howard, Victor Kuarsingh, Alexandru Petrescu, Mikael Abrahamsson, Jong-Hyouk Lee, Doug Barton, Owen DeLong, Anders Brandt, Tim Chown, Jen Linkova, Christopher Palmer and Wesley George.

Some test of using ULA in the lab was done by our research partner BNRC-BUPT (Broad Network Research Centre in Beijing University of

Posts and Telecommunications). Thanks for the work of Prof. Xiangyang Gong and student Dengjia Xu.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Bing Liu
Huawei Technologies Co., Ltd
Huawei Q14 Building, No.156 Beiqing Rd.,
Zhong-Guan-Cun Environmental Protection Park, Beijing
P.R. China

EMail: leo.liubing@huawei.com

Sheng Jiang
Huawei Technologies Co., Ltd
Huawei Q14 Building, No.156 Beiqing Rd.,
Zhong-Guan-Cun Environmental Protection Park, Beijing
P.R. China

EMail: jiangsheng@huawei.com