

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 28 December 2023

S. Jiang
BUPT
B. Liu
Huawei Technologies
N. Buraglio
ForwardingPlane, LLC
26 June 2023

Considerations For Using Unique Local Addresses
draft-ietf-v6ops-ula-usage-considerations-03

Abstract

This document provides considerations for using IPv6 Unique Local Addresses (ULAs). Based on an analysis of different ULA usage scenarios, this document identifies use cases where ULA addresses are helpful as well as potential problems caused by using them.

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 <https://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 28 December 2023.

Copyright Notice

Copyright (c) 2023 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Requirements Language	3
3. General Considerations For Using ULAs	3
3.1. Do Not Treat ULA Equal to RFC1918	3
3.2. Using ULAs in a Limited Scope	4
4. Analysis and Operational Considerations for Scenarios Using ULAs	4
4.1. ULA-only in Isolated Networks	4
4.2. ULA+PA in Connected Networks	5
4.3. ULA-Only in Connected Networks	7
4.4. Some Specific Use Cases	8
4.4.1. Special Routing	8
4.4.2. Used as Identifier	8
4.5. IPv4 Co-existence Considerations	9
5. Security Considerations	9
6. IANA Considerations	10
7. Acknowledgements	10
8. References	10
8.1. Normative References	10
8.2. Informative References	10
Authors' Addresses	12

1. Introduction

Unique Local Addresses (ULA) is defined in [RFC4193], and it is an alternative to site-local address (deprecated in [RFC3879]). ULAs have the following features:

- Automatically Generated

ULA prefixes can be automatically generated using the algorithms described in [RFC4193]. This feature allows automatic prefix allocation. Thus one can get a network working immediately without applying for prefix(es) from an RIR/LIR (Regional Internet Registry/Local Internet Registry).

- Globally Unique

ULAs are defined as a global scope address space. However, they are not intended to be used globally on the public Internet; in contrast, they are mostly used locally, for example, in isolated networks, internal networks, or VPNs.

ULAs are intended to have an extremely low probability of collision. The randomization of 40 bits in a ULA prefix is considered sufficient enough to ensure a high degree of uniqueness (refer to [RFC4193] Section 3.2.3 for details) and simplifies merging of networks by avoiding the need to renumber overlapping IP address space.

- Provider Independent Address Space

ULAs can be used for internal communications even without Internet connectivity. They need no registration, so they can support on-demand usage and do not carry any RIR/LIR burden of documentation or fees.

- Well Known Prefix

The prefixes of ULAs are well known thus they are easily identified and filtered.

This document aims to introduce the usage of ULAs in various scenarios, provide some operational considerations, and clarify the advantages and disadvantages of the usage in each scenario. Thus, the administrators could choose to use ULAs in a certain way that considered beneficial for them.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. General Considerations For Using ULAs

3.1. Do Not Treat ULA Equal to RFC1918

ULA and [RFC1918] are similar in some aspects. The most obvious one is as described in Section 3.1.3 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 this does not mean that ULA is equal to an IPv6 version of [RFC1918] deployment. [RFC1918] usually combines with NAT/NAPT for global connectivity. But it is not necessary to combine ULAs with any kind of NAT. Operators can use ULA for local communications along with global addresses for global communications (see Section 4.2). This is a big advantage brought by default support of multiple-addresses-per-interface feature in IPv6. (People may still have a requirement for NAT with ULA, this is discussed in Section 4.3. But people also need to 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].

3.2. 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 need to 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 is enabled, the administrators might also need to use internal-only DNS names for ULAs and might need to split the DNS so that the internal DNS server includes records that are not presented in the external DNS server.

4. Analysis and Operational Considerations for Scenarios Using ULAs

4.1. ULA-only in Isolated Networks

IP is used ubiquitously. Some networks like industrial control bus (e.g. [RS-485], [SCADA], or even non-networked digital interfaces like [MIL-STD-1397] have begun to use IP. In these kinds of networks, the system may lack the ability to communicate with the public networks.

As another example, there may be some networks in which the equipment has the technical capability to connect to the Internet, but is prohibited by administration. These networks may include data center networks, separate financial networks, lab networks. machine-to-machine (e.g. vehicle networks), sensor networks, or even normal LANs, and can include very large numbers of addresses.

ULA is a straightforward way to assign the IP addresses in the kinds of networks just described, 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 assigning vehicles with ULAs, it is then possible to separate in-vehicle embedded networks into different subnets depending on real-time situation.

However, each isolated network has the possibility to be connected in the future. Administrators need to consider the following before deciding whether to use ULAs:

- * If the network eventually connects to another isolated or private network, the potential for address collision arises. However, if the ULAs were generated in the standard way, this will not be a big problem.
- * If the network eventually connects to the global Internet, then the operator will need to add a new global prefix and ensure that the address selection policy is properly set up on all interfaces.

Operational considerations:

- * Prefix generation: randomly generated according to the algorithms defined in [RFC4193] or manually assigned. Normally, automatic generation of the prefixes is recommended, following [RFC4193]. If there are some specific reasons that call for manual assignment, administrators have to plan the prefixes carefully to avoid collision.
- * Prefix announcement: in some cases, networks might need to announce prefixes to each other. For example, in vehicle networks with infrastructure-less settings such as Vehicle-to-Vehicle (V2V) communication, prior knowledge of the respective prefixes is unlikely. Hence, a prefix announcement mechanism is needed to enable inter-vehicle communications based on IP. As one possibility, such announcements could rely on extensions to the Router Advertisement message of the Neighbor Discovery Protocol [RFC4861].

4.2. ULA+PA in Connected Networks

Two classes of network might need to use ULA with PA (Provider Aggregated) addresses:

- * Home network. Home networks are normally assigned with one or more globally routed PA prefixes to connect to the uplink of an ISP. In addition, they may need internal routed networking even when the ISP link is down. Then ULA is a proper tool to fit the

requirement. [RFC7084] requires the CPE to support ULA. Note: ULAs provide more benefit for multiple-segment home networks; for home networks containing only one segment, link-local addresses are better alternatives.

- * 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 can be used to improve internal connectivity and make it more resilient, or to isolate certain functions like OAM for servers.

Benefits of Using ULAs in this scenario:

- * Separated local communication plane: for either home networks or enterprise networks, the main purpose of using ULAs along with PA addresses is to provide a logically local routing plane separated from the global 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 for 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 rest of the network because they are changing ISPs, because the ISP has restructured its address allocations, or for some other reason. In these situations, ULA is an effective tool for addressing internal-only nodes. Even public nodes can benefit from ULA for renumbering, on their internal interfaces. When renumbering, as [RFC4192] suggests, old prefixes continue to be valid until 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 local communication disentangled from global routing plane change. If we use ULAs for local communication, the separated local routing plane can isolate the effects of global routing change.

Drawbacks:

- * Operational Complexity: there are some arguments that in practice the use of ULA+PA creates additional operational complexity. This is not a ULA-specific problem; the multiple-addresses-per-interface is an important feature of IPv6 protocol. Nevertheless, running multiple prefixes needs more operational consideration than running a single one.

Operational considerations:

- * **Default Routing:** connectivity may be broken if ULAs are used as default route. When using RIO (Route Information Option) in [RFC4191], specific routes can be added without a default route, thus avoiding bad user experience due to timeouts on ICMPv6 redirects. This behavior was well documented in [RFC7084] as rule ULA-5 "An IPv6 CE router MUST NOT advertise itself as a default router with a Router Lifetime greater than zero whenever all of its configured and delegated prefixes are ULA prefixes." and along with rule L-3 "An IPv6 CE router MUST advertise itself as a router for the delegated prefix(es) (and ULA prefix if configured to provide ULA addressing) using the "Route Information Option" specified in Section 2.3 of [RFC4191]. This advertisement is independent of having or not having IPv6 connectivity on the WAN interface.". However, it needs to be noticed that current OSes don't all support [RFC4191].
- * **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 PA prefixes in accordance with the two mechanisms. The administrators need to know the current issue of the SLAAC/DHCPv6 interaction.
- * **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. [RFC6724] claims 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 choose not to do reverse DNS delegation inside of their local control of ULA prefixes, a significant amount of information about the ULA population may leak to the outside world. Because reverse queries will be made and naturally routed to the global reverse tree, so external parties will be exposed to the existence of a population of ULA addresses. [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.

4.3. ULA-Only in Connected Networks

In theory, a site numbered with ULAs only can get connected via a NPTv6[RFC6296] (which is an experimental specification that provides a stateless one-to-one mapping between internal addresses and external addresses) or application-layer proxy. This approach could get provider independent addresses or get connected from the isolated stage without applying to any RIRs/LIRs. This might make small

organizations saving time and address fee.

However, this approach breaks the end-to-end transparency. People have suffered from the NAT/Proxy middle boxes so much in the IPv4 era, there is no reason to continue the suffering when IPv6 is available. This document does not consider ULA+NPTv6/Proxy as a good choice for normal cases. Rather, this document considers ULA+PA (Provider Aggregated) as a better approach to connect to the global network when ULAs are expected to be retained.

4.4. Some Specific Use Cases

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

4.4.1. Special Routing

For various reasons the administrators may want to have private routing be controlled and separated from other routing. For example, in the business-to-business case, 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 prefixes that would be used in accordance with an agreement between the parties.

4.4.2. 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 upper 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:

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

It needs to be noticed again that in theory ULA has the possibility of collision. However, the probability is desirably small enough and can be ignored in most cases when ULAs are used as identifiers.

4.5. IPv4 Co-existence Considerations

Generally, this document does not consider IPv4 to be in scope. But regarding ULA, there is a special case needs to be recognized, which is described in Section 3.2.2 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 policy table defined in the old address selection standard [RFC3484]. This will clearly result in a connection failure. The new address selection standard [RFC6724] has corrected this behavior by preferring IPv4 than ULAs in the default policy table. However, there are still lots of hosts using the old standard [RFC3484], thus this could be an issue in real networks.

Happy Eyeballs [RFC8305] solves this connection failure problem, but unwanted timeouts will obviously lower the user experience. One possible approach to eliminating the timeouts is to deprecate the IPv6 default route and simply configure a scoped route on hosts (in the context of this document, only configure the ULA prefix routes). Another alternative is to configure IPv4 preference on the hosts, and not include 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 and can be connected through IPv4 as default and internal nodes can always connect through IPv6. But since IPv6 preference is default, changing the default in all nodes is not suitable at scale.

5. 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 4.2, when using NPTv6, the administrators need to know where the firewall is located to set proper filtering rules.

Also as mentioned in Section 4.2, if administrators choose not to do reverse DNS delegation inside their local control of ULA prefixes, a significant amount of information about the ULA population may leak to the outside world.

6. IANA Considerations

This memo has no actions for IANA.

7. Acknowledgements

Many valuable comments were received in the IETF v6ops WG mail list, especially from Cameron Byrne, Fred Baker, Brian Carpenter, Lee Howard, Victor Kuarsingh, Alexandru Petrescu, Mikael Abrahamsson, Tim Chown, Jen Linkova, Christopher Palmer Jong-Hyouk Lee, Mark Andrews, Lorenzo Colitti, Ted Lemon, Joel Jaeggli, David Farmer, Doug Barton, Owen Delong, Gert Doering, Bill Jouris, Bill Cervený, Dave Thaler, Nick Hilliard, Jan Zorz, Randy Bush, Anders Brandt, , Sofiane Imadali 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.

Tom Taylor did a language review and revision throughout the whole document. The authors appreciate a lot for his help.

This document was produced using the xml2rfc tool [RFC7991] (initially prepared using 2-Word-v2.0.template.dot.).

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<https://www.rfc-editor.org/info/rfc4193>>.
- [RFC7991] Hoffman, P., "The "xml2rfc" Version 3 Vocabulary", RFC 7991, DOI 10.17487/RFC7991, December 2016, <<https://www.rfc-editor.org/info/rfc7991>>.

8.2. Informative References

- [MIL-STD-1397] "Military Standard, Input/Output Interfaces, Standard Digital Data, Navy Systems (MIL-STD-1397B), 3 March 1989".

- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, DOI 10.17487/RFC3484, February 2003, <<https://www.rfc-editor.org/info/rfc3484>>.
- [RFC3879] Huitema, C. and B. Carpenter, "Deprecating Site Local Addresses", RFC 3879, DOI 10.17487/RFC3879, September 2004, <<https://www.rfc-editor.org/info/rfc3879>>.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <<https://www.rfc-editor.org/info/rfc4191>>.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, DOI 10.17487/RFC4192, September 2005, <<https://www.rfc-editor.org/info/rfc4192>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, DOI 10.17487/RFC4864, May 2007, <<https://www.rfc-editor.org/info/rfc4864>>.
- [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, DOI 10.17487/RFC5220, July 2008, <<https://www.rfc-editor.org/info/rfc5220>>.
- [RFC6281] Cheshire, S., Zhu, Z., Wakikawa, R., and L. Zhang, "Understanding Apple's Back to My Mac (BTMM) Service", RFC 6281, DOI 10.17487/RFC6281, June 2011, <<https://www.rfc-editor.org/info/rfc6281>>.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, DOI 10.17487/RFC6296, June 2011, <<https://www.rfc-editor.org/info/rfc6296>>.

- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<https://www.rfc-editor.org/info/rfc6724>>.
- [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, DOI 10.17487/RFC7084, November 2013, <<https://www.rfc-editor.org/info/rfc7084>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8305] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", RFC 8305, DOI 10.17487/RFC8305, December 2017, <<https://www.rfc-editor.org/info/rfc8305>>.
- [RS-485] "Electronic Industries Association (1983). Electrical Characteristics of Generators and Receivers for Use in Balanced Multipoint Systems. EIA Standard RS-485.".
- [SCADA] "Boyer, Stuart A. (2010). SCADA Supervisory Control and Data Acquisition. USA: ISA - International Society of Automation.".
- [ULA-IN-WILD]
"G. Michaelson, "conference.apnic.net/data/36/apnic-36-ula_1377495768.pdf"."

Authors' Addresses

Sheng Jiang
Beijing University of Posts and Telecommunications
No. 10 Xitucheng Road
Haidian District
Beijing
100083
China
Email: shengjiang@bupt.edu.cn

Bing Liu
Huawei Technologies
Q14, Huawei Campus, No.156 Beiqing Road
Hai-Dian District, Beijing, 100095
P.R. China

Email: leo.liubing@huawei.com

Nick Buraglio
ForwardingPlane, LLC
Email: buraglio@forwardingplane.net