IP Version 6 Working Group
Internet-Draft
Intended status: Informational
Expires: 20 July 2026

W. O. Ideafarm
IDEAFARM.COM

16 January 2026

IPv6 Sparse Random Addressing
draft-ideafarm-ipv6-sparse-random-addressing-00

Abstract

The IPv6 address space is huge. This document discusses how "sparse random addressing" can be used to defeat DDOS attacks, and also to create paid-access websites. The essential idea is to use the host ID portion of an IPv6 address as a time-based password that only paid subscribers know. (They know it by running a program on the device that they use to access the website. That program uses the current time and a secret shared with the web server to calculate the current host ID.) Each subscriber has a unique secret so, at any point in time, uses a unique IPv6 address to access the website. Sparseness prevents attack packets from reaching the web server. Uniqueness creates accountability. To protect routers from flood attacks, the globally routable /48 prefix is also randomized, in a way that exploits BGP route propagation delay to partially or completely quash the flooding at the source, thereby protecting all upstream routers.

This early draft only contains a conceptual overview.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://ideafarm.github.io/RFC-IPv6-Sparse-Random-Addressing/draft-ideafarm-ipv6-sparse-random-addressing.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-ideafarm-ipv6-sparse-random-addressing/.

Discussion of this document takes place on the IP Version 6 Working Group Working Group mailing list (mailto:ipv6@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/ipv6. Subscribe at https://www.ietf.org/mailman/listinfo/ipv6/.

Source for this draft and an issue tracker can be found at https://github.com/ideafarm/RFC-IPv6-Sparse-Random-Addressing.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction 2
2. Conventions and Definitions 5
3. Security Considerations 5
4. IANA Considerations 5
5. Normative References 5
Acknowledgments 6
Author's Address 6

1. Introduction

The IPv6 address space is huge. For example, it contains 68 billion prefixes that are 36 bits long, so every human being alive today could be given a /36 prefix for his or her exclusive use for life, and this policy could continue indefinitely without ever consuming more than a small fraction of the available address space. Yet all of the Tier-1 networking providers and all of the Regional Internet Registries still operate as if addresses are scarce and that each device should only have a single, unchanging, published, address.

For example, ARIN’s IPv6 block assignment policy requires that an applicant must have multiple physical locations to obtain a prefix shorter than the minimum prefix length (48 bits) that is globally routable, and that the length of the prefix granted will depend upon the number of physical locations such that each location is only given a single routable prefix. The policy neither contemplates nor provides for using the huge IPv6 address space in any manner other than the way that IPv4 addresses have been used for a half century, one unchanging public address per device.

This document proposes that portions of an IPv6 address be used as time-based passwords, describes a particular experimental implementation and preliminary findings from studying it, and discusses the costs and benefits, both to individual organizations and to the Internet community, of using the huge IPv6 address space in this new way.

Time-based passwords have become widely used to provide "two factor authentication" when a user logs into a website. In that application, brute force password attacks are prevented by requiring that the user also enter a time-based password that is calculated using the current time and a secret that only he and the web server knows. Current practice allocates 64 bits for the host ID portion of an IPv6 address, which is huge relative to the number of distinct values needed to identify each host on a physical link of any conceivable size. By using some (or all) of those bits as a time-based password, attack packets can be prevented, by sparseness, from reaching that server.

For such a scheme to work, a would-be attacker must not be able to discover the secret. A simple way to prevent this is to give a unique secret to each paid subscriber, which requires that the server accepts packets from many IPv6 addresses, one for each subscriber. An attacker can purchase a subscription and use a tool like Wireshark to see the current IPv6 address that gives him access to the server, but if he uses that information to attack the server, the IPv6 addresses that the attack uses will reveal his identity.

A similar approach can be used to secure upstream routers from a flooding attack. A well funded and motivated attacker might use a botnet to flood the /48 prefix with randomly addressed packets, not to reach the web server, but to get the upstream ISP to nullroute the entire prefix. Randomizing a portion of the /48 prefix can mitigate and, in some scenarios, entirely quash such an attack.

Using the huge IPv6 address space in this way consumes orders of magnitude more addresses, so such applications revive address exhaustion concerns. Fortunately, these concerns are mitigated by

two facts. First, the space is so huge that, even with current routing practice, every person on the planet can be given a /36 prefix for his exclusive use, and each such block contains 4096 routable blocks (blocks that use a 48 bit prefix), since 12 bits are available for randomizing the prefix. Second, the RIR's can reserve a portion of the address space for sparse random addressing in a way that permits routers to enforce a requirement that routes in that space are sparse. This would enable routers to treat much longer prefixes as globally routable in that space. (In that space, the number of routes in the global routing table would be constrained by enforced sparseness rather than prefix length.) It might take several years for such prefixes to become reliably routable, but after that transition, sparse random addressing would no longer require prefixes shorter than /48. So sparse random addressing would not, long term, create an address exhaustion concern.

Explosive growth of the global routing table, not address exhaustion, will likely be what constrains use of the huge IPv6 address space. Even if prefixes longer than 48 bits are dropped, that leaves up to 281 trillion prefixes requiring global routability. Sparse random addressing "theoretically" does not increase the number of routes advertised, but in practice it will, because multiple routes must be advertised at the same time. As discussed below, the objective in randomizing the prefix is to exploit asymmetries that put a DDOS attacker and his botnet at a disadvantage. In the experimental implementation, this is done by withdrawing the eldest prefix and advertising a new prefix approximately once each minute, with each prefix being advertised for about 15 minutes. Deploying sparse random addressing in this way, with each server cluster emitting a BGP UPDATE (add+withdraw) message once each minute, would place a new processing burden on the global BGP system.

This document is a working draft that presents the current status of an experimental implementation of sparse random addressing using a single router and a single web server at a single datacenter to serve a simple paid-access website. The initial objective is to get it all to work, and then to assess robustness. (Can it effectively withstand a DDOS attack?) This document will specify the experimental system in enough detail to enable others to replicate the experiment, using the same proprietary open source software. In this initial phase, the focus will be to find out whether it is possible for a well funded and motivated DDOS attacker to bring down the experimental website.

The next research objective will be to quantify the resource impact that the deployment of such systems will have on the global BGP and routing system. If the benefits of defeating DDOS attacks and of facilitating paid-access websites are significant, the Internet

community can then discuss what level of incremental processing cost it is willing to bear, in the form of more expensive routers and higher energy consumption, to obtain those benefits.

At this early stage, this approach appears to promise to give the Internet community the "upper hand" in its battle against bad actors who use botnets to attack servers, and sparse random addressing, which can only be done with IPv6, just might be the "killer app" that could finally drive universal migration to IPv6.

2. Conventions and Definitions

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. Security Considerations

[TODO Security]

4. IANA Considerations

This document proposes that a /16 block be reserved for sparse random addressing:

Block: 2003::/16 Name: Sparse Random Addressing RFC: [This document] Allocation Date: [TBD] Termination Date: N/A Source: TRUE Destination: TRUE Forwardable: TRUE Globally Reachable: TRUE Reserved-by-Protocol: TRUE

Required IP compliant processing: External BGP peers SHALL NOT drop routes for prefixes within this block based upon prefix lengths from /0 through /128. To limit the number of global routing table entries for destinations in this block, routers SHALL enforce sparseness, using any of the algorithms specified in this document.

5. 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/rfc/rfc2119>.

[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/rfc/rfc8174>.

Acknowledgments

The author thanks the ARIN Regional Internet Registry for temporarily providing an IPv6 /36 block to accommodate his study of, and experimentation with, sparse random addressing.

Author’s Address

Wo Of Ideafarm
IDEAFARM.COM
Email: wo.ideafarm.publication.delayed.4.yrs@ideafarm.com