

AusNOG 2019

Getting

IPv6 Private Addressing

Right

Mark Smith

markzzzsmith@gmail.com
@ipv6tao
@markzzzsmith

IPv6 Private Addressing:

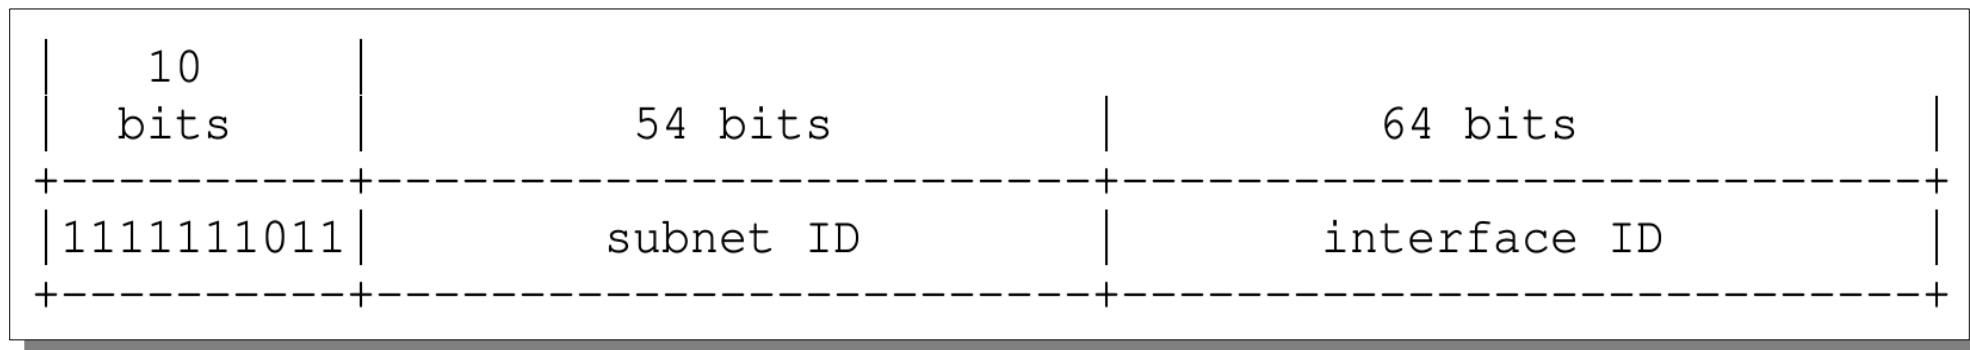
First Attempt

IPv6 Site-Local Addresses

RFC 4291

IPv6 Addressing Architecture

February 2006



IANA Allocation: **fec0::/10**

IETF IPv6 WG Circa 2002

What is a "site"?


Merge sites with no NAT or
renumbering?

"Site"?



Buildings **or** Campus?

Buildings **and** Campus?

"Site"?



"Site"?



"Site"?


site

"1. The **place** where anything is **fixed**; situation; local position"

"2. A **place** fitted or chosen for any certain **permanent** use or occupation"

<https://en.wiktionary.org/wiki/site>

"Site"?



`fec0::/10` prefix/addresses on the link?


“Site”

Really too geographical.




Site-Locals

Sometimes
unsuitable.



Site-Local

Networking Issues



Distinguish instances?

Node internal
Site
Identifier?




VRFs in all
IPv6 nodes?

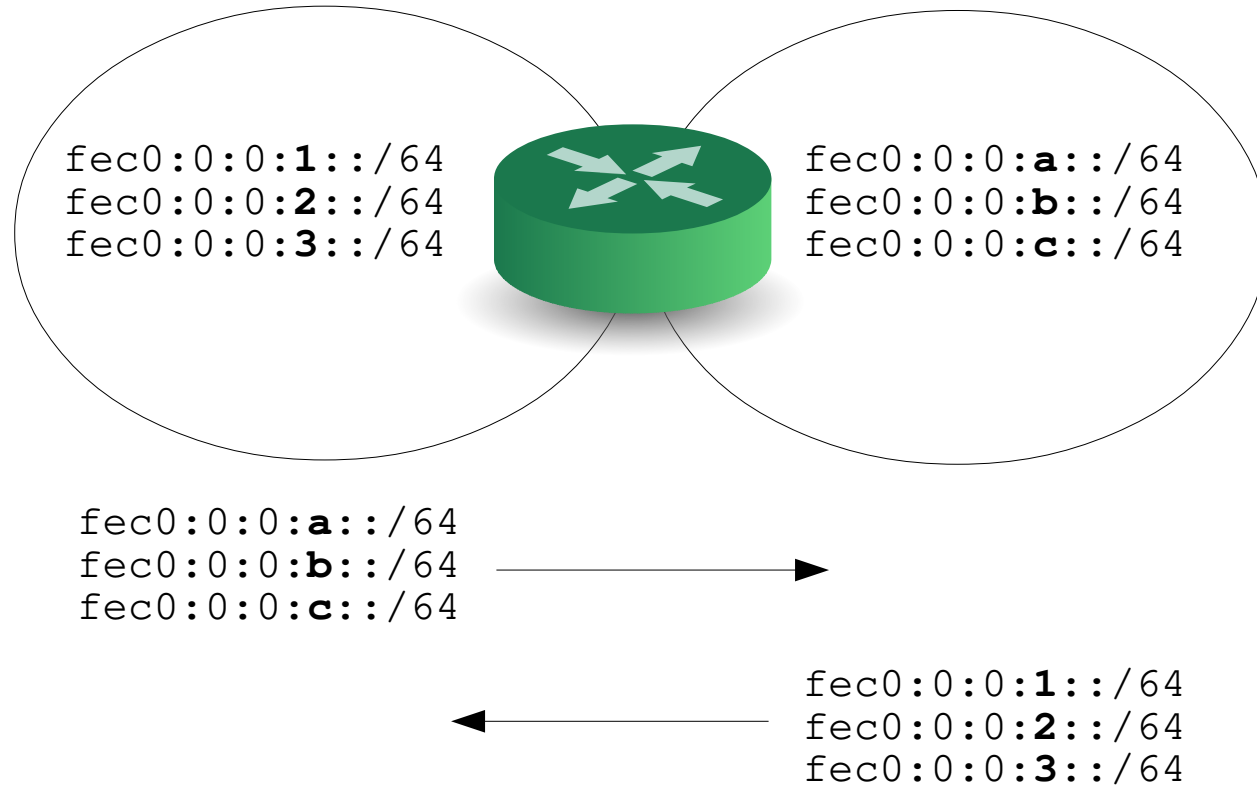
COMPLEXITY!

Merging Site-Local Sites

Overlapping or Duplicated Address Spaces



Lucky? Unique Subnets?



Unlucky. Duplicate Subnets.




NAT? NO.

The Trouble with NAT (Or why I care about IPv6)


AusNOG Conference
2016

Mark Smith
markzzzsmith@gmail.com
@markzzzsmith

Renumber Site Subnets?



Fundamental problem



Trying to send **1:1** (**unicast**) to a **non-unique** destination.

—→ Want
—→ Get

Uniquifying Addresses

Process -

Renumbering

Function -

NAT

Adding Context -

Site ID
VRF

"Prevention is better than cure"

Desiderius Erasmus

UNIQUE' em

in the


FIRST PLACE.

ALL CONTEXTS!

48-bit Absolute Internet and Ethernet Host Numbers

by Yogen K. Dalal and Robert S. Printis

OPD-T8101 July 1981



Abstract: Xerox internets and Ethernet local computer networks use 48-bit absolute host numbers. This is a radical departure from practices currently in use in internetwork systems and local networks. This paper describes how the host numbering scheme was designed in the context of an overall internetwork and distributed systems architecture.

<https://ethernethistory.typepad.com/papers/HostNumbers.pdf>

~~"Site"~~ too.

Network Working Group
Request for Comments: 3879
Category: Standards Track

C. Huitema
Microsoft
B. Carpenter
IBM
September 2004

Deprecating Site Local Addresses

Abstract

This document describes the issues surrounding the use of IPv6 site-local unicast addresses in their original form, and formally deprecates them. This deprecation does not prevent their continued use until a replacement has been standardized and implemented.



IPv6 Private Addressing:

Unique Local Unicast
Addressing


Network Working Group
Request for Comments: 4193
Category: Standards Track

R. Hinden
Nokia
B. Haberman
JHU-APL
October 2005

Unique Local IPv6 Unicast Addresses

Abstract

This document defines an IPv6 unicast address format that is globally unique and is intended for local communications, usually inside of a site. These addresses are not expected to be routable on the global Internet.



Prefix FC00::/7 prefix to identify Local IPv6 unicast addresses.

L Set to 1 if the prefix is locally assigned.
Set to 0 may be defined in the future. See
Section 3.2 for additional information.

Global ID 40-bit global identifier used to create a
globally unique prefix. See Section 3.2 for
additional information.

Subnet ID 16-bit Subnet ID is an identifier of a subnet
within the site.

Interface ID 64-bit Interface ID as defined in [ADDARCH].



ULA-C or ULA-L

Type	Purpose	L Bit	Prefix	Status
ULA-C	Central ULA Prefix Registry	L=0	fc 00::/8	Never took off
ULA-L	Local Network Generated	L=1	fd 00::/8	All current ULAs



All currently valid ULAs start with **fd**.

40 bit Global ID

RFC 4193

Unique Local IPv6 Unicast Addresses

October 2005

3.2. Global ID

The allocation of Global IDs is pseudo-random [RANDOM]. They MUST NOT be assigned sequentially or with well-known numbers. This is to ensure that there is not any relationship between allocations and to help clarify that these prefixes are not intended to be routed globally. Specifically, these prefixes are not designed to aggregate.



[RANDOM] Eastlake, D., 3rd, Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.

40 bit Global ID

RFC 4193

Unique Local IPv6 Unicast Addresses

October 2005

3.2. Global ID

The allocation of Global IDs is pseudo-random [RANDOM]. **They MUST NOT be assigned sequentially or with well-known numbers.** This is to ensure that there is not any relationship between allocations and to help clarify that these prefixes are not intended to be routed globally. Specifically, these prefixes are not designed to aggregate.



40 bit Global ID

“Their values SHOULD NOT be
generated by a human.” –
Me!

Humans may not [RANDOM] .



This presentation?



40 bit Global ID

AusNOG 2011

Residential IPv6 CPE What Not to Do and Other Observations

Mark Smith
Nextgen Networks
mark.smith@nn.com.au
September 2011


40 bit Global ID

AusNOG 2011

ULAs with random part of all zeros

- One implementation would announce ULAs with an all zero random part, when IPv6 on the WAN interface went down i.e. would attempt to “swap” ULAs for globals, rather than make ULAs constant and independent of WAN IPv6 delegated prefix.
- Effectively makes ULAs the deprecated IPv6 site-locals.

40 bit Global ID

[Main](#) | [About](#) | [Contact](#) | [News](#) | [User Home](#) | [PoPs](#) | [Presentations](#)

[GRH](#) | [DFP](#) | [ULA](#) | [Compare](#) | [Looking Glass](#) | [Status](#)

IPv6 ULA (Unique Local Address) RFC4193 Registration List

The ULA register currently has the following 5625 prefixes registered.

Informal registry

fd36:ddb1:26da::/48	Administrator	Resources Holding
fd36:ddb1:26db::/48	Administrator	Resources Holding
fd36:ddb1:26dc::/48	Administrator	Resources Holding
fd36:ddb1:26dd::/48	Administrator	Resources Holding
fd36:ddb1:26de::/48	Administrator	Resources Holding
fd36:ddb1:26df::/48	Administrator	Resources Holding

40 bit Global ID

kube-v6

Instructions on how to instantiate a multi-node, IPv6-only Kubernetes cluster using the CNI bridge plugin for developing or exploring IPv6 on Kubernetes.

No advice on
correct ULAs!

Set up node IP addresses

For the example topology show above, the eth2 interface requires a router external to the Kubernetes cluster. The router would be statically configured with IPv6 Unique L

Node	IP Address
-----	-----
NAT64/DNS64	fd00::64
Kube Master	fd00::100
Kube Node 1	fd00::101
Kube Node 1	fd00::102

Add static routes between nodes, po

In the list of static routes below, the subnets/addresses used

Subnet/Address	Description
-----	-----
64:ff9b::/96	Prefix used inside the cluster
fd00::101	Kube Node 1
fd00::102	Kube Node 2
fd00:101::/64	Kube Node 1's pod subnet
fd00:102::/64	Kube Node 2's pod subnet
fd00:1234::/64	Cluster's Service subnet

Site-Locals and their
problems reproduced!

ULA-L Global ID Algorithm

RFC 4193

Unique Local IPv6 Unicast Addresses

October 2005

3.2.2. Sample Code for Pseudo-Random Global ID Algorithm

Summary:

1. 64 bit NTP current time of day
2. System EUI-64 (IEEE), or system serial number
3. Concatenate those two values
4. SHA-1 to get 160 bit hash
5. Least significant 40 bits of SHA-1 hash will be Global ID
6. Append to 8 bits of 0xfd to produce ULA /48 prefix (i.e. L=1)

Example

My home ULA:

fd[REDACTED]:[REDACTED]:[REDACTED]:[REDACTED]: : / 48


ULA Generators

RFC 7084,
"Basic Requirements for
IPv6 Customer Edge
Routers"

Android Apps

Online - search "IPv6
ULA generator"

iPhone Apps (?)



Subnet ID 16-bit Subnet ID is an identifier of a subnet within the site.

Interface ID 64-bit Interface ID as defined in [ADDARCH].

Routing

RFC 4193

Unique Local IPv6 Unicast Addresses

October 2005

fc00::/7 and longer prefixes

– by default, don't send or accept in EGP

fd[REDACTED]:[REDACTED]:[REDACTED]:[REDACTED]::/48



– network operator override to allow inter-ULA domain routing

Other Operational Things

RFC 4193

Unique Local IPv6 Unicast Addresses

October 2005

4. Operational Guidelines	7
4.1. Routing	7
4.2. Renumbering and Site Merging	7
4.3. Site Border Router and Firewall Packet Filtering	8
4.4. DNS Issues	8
4.5. Application and Higher Level Protocol Issues	9
4.6. Use of Local IPv6 Addresses for Local Communication	9
4.7. Use of Local IPv6 Addresses with VPNs	10
5. Global Routing Considerations	11
5.1. From the Standpoint of the Internet	11
5.2. From the Standpoint of a Site	11
6. Advantages and Disadvantages	12
6.1. Advantages	12
6.2. Disadvantages	13



Would GUAs do?

Network Working Group
Request for Comments: 3587
Obsoletes: 2374
Category: Informational

R. Hinden
Nokia
S. Deering
Cisco
E. Nordmark
Sun
August 2003

IPv6 Global Unicast Address Format





Also known as **public IPv6 Internet addresses**

Not really

	GUA	ULA
Globally Unique (Assured or Likely)	✓	✓
Designed for Internet connectivity/reachability	✓	✗
Assigned by	RIR/LIR/ISP	Local administrator or local CPE
\$\$\$	Per Month/Annum	\$0 forever



ULA Use Cases




ULA Use Cases


RFC 2473

Generic Packet Tunneling in IPv6

December 1998



ULA Addressing



IPv6-in-IPv6
IPv4-in-IPv6
GRE
L2TP
SRv6
etc.


GUA OR ULA



WHY NOT BOTH?




Why Both?



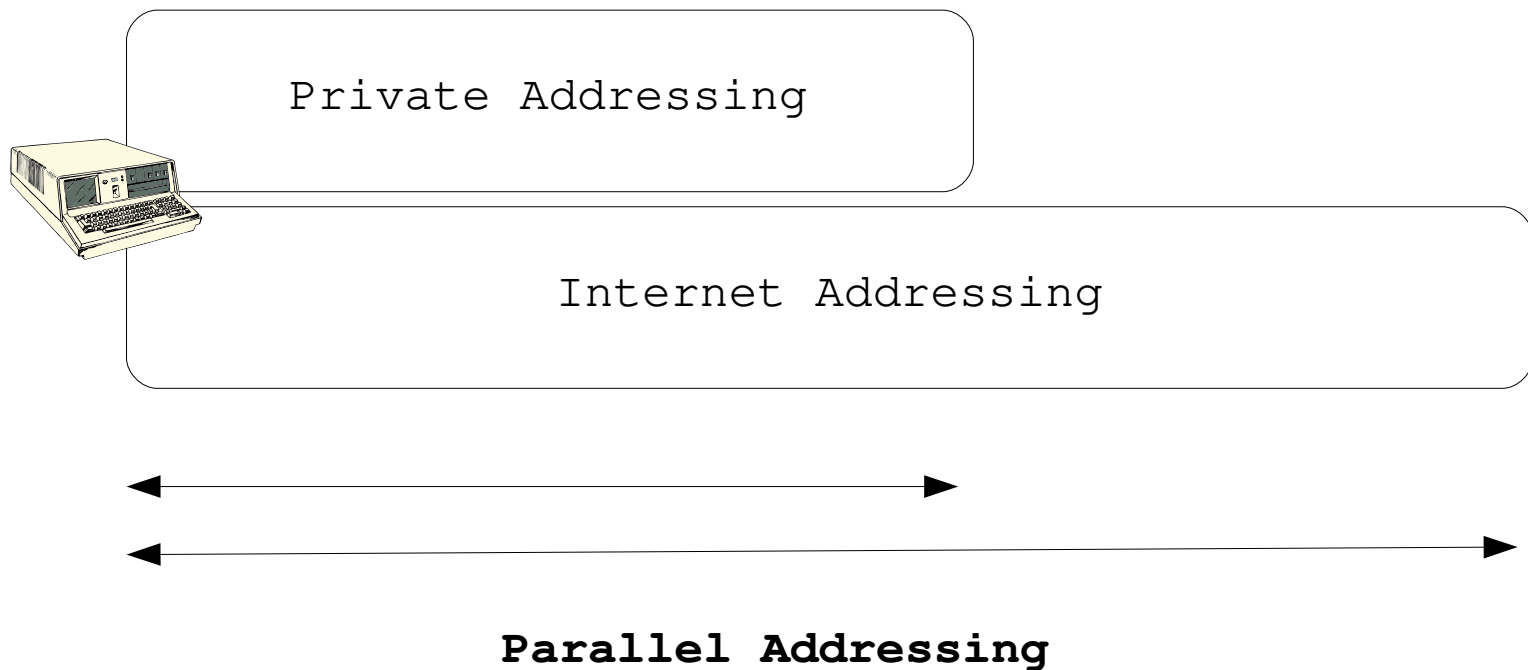
Internal communication
independent of
external addressing

IPv6 Multi-Addressing


IPv4 Private and Internet Addressing





IPv6 Private and Internet Addressing



IPv6 Node Address Mix





Private
Addressing



Internet Addressing

Local Network

IPv6 Node Interface Multi-Addressing




Link Locals

ULAs

GUAs

Link-Local addresses are for use on a single link. Link-Local addresses have the following format:



IANA Allocation: **fe80::/10**

Destination Address? Source Address?

RFC 6724

Default Address Selection for IPv6

September 2012

Very basically:

1. Pick smallest scope

Loopback < Link-Local < ULA, GUA

2. Pick ULA over GUA

3. Pick addresses with the most bits in common

Internet Engineering Task Force (IETF)

Request for Comments: 7934

BCP: 204

Category: Best Current Practice

ISSN: 2070-1721

L. Colitti

V. Cerf

Google

S. Cheshire

D. Schinazi

Apple Inc.

July 2016

Host Address Availability Recommendations

Abstract

This document recommends that networks provide general-purpose end hosts with multiple global IPv6 addresses when they attach, and it describes the benefits of and the options for doing so.

Origins?

Network Working Group
Request for Comments: 1681
Category: Informational

S. Bellovin
AT&T Bell Laboratories
August 1994

On Many Addresses per Host

Overview and Rational


Currently, most hosts have only one address. With comparatively rare exceptions, hosts as hosts -- as opposed to hosts acting as routers or PPP servers -- are single-homed. Our address space calculations reflect this; we are assuming that we can estimate the size of the address space by counting hosts. But this may be a serious error. I suggest that that model may -- and should -- change.

Bad for Security?

Firewalls and Internet Security

Repelling the Wily Hacker

William R. Cheswick
Steven M. Bellovin



ADDISON-WESLEY PROFESSIONAL COMPUTING SERIES

Copyright **1994**
AT&T and Lumeta
Corporation.

Innovative Uses

Transient Addressing for Related Processes: Improved Firewalling by Using IPV6 and Multiple Addresses per Host

Peter M. Gleitz*

Steven M. Bellovin

smb@research.att.com

AT&T Labs Research

Abstract

Traditionally, hosts have tended to assign relatively few network addresses to an interface for extended periods. Encouraged by the new abundance of addressing possibilities provided by IPv6, we propose a new method, called Transient Addressing for Related Processes (TARP), whereby hosts temporarily employ and subsequently discard IPv6 addresses in servicing a client host's network requests. The method provides certain security advantages and neatly finesses some well-known firewall problems caused by dynamic port negotiation used in a variety of application protocols. A pro-

[https://
www.cs.columbia.edu/~smb/
papers/tarp.pdf](https://www.cs.columbia.edu/~smb/papers/tarp.pdf)

Proceedings of the Eleventh Usenix Security Conference, August 2001, Washington, D.C.

Internet Engineering Task Force (IETF)
Request for Comments: 8273
Category: Informational
ISSN: 2070-1721

J. Brzozowski
Comcast Cable
G. Van de Velde
Nokia
December 2017

Unique IPv6 Prefix per Host

Abstract

This document outlines an approach utilizing existing IPv6 protocols to allow hosts to be assigned a unique IPv6 prefix (instead of a unique IPv6 address from a shared IPv6 prefix). Benefits of using a unique IPv6 prefix over a unique service-provider IPv6 address include improved host isolation and enhanced subscriber management on shared network segments.


SSID:ABCD_Free

2001:db8:0:**1234**::/64

2001:db8:0:**5678**::/64

2001:db8:0:**abcd**::/64

..... Client isolation



Summary

Site-Local
Addressing

Unique Local
Addressing

IPv6 Multi-
Addressing

IPv6 Addressing
Innovations



UNIQUE your **ULAs** !



Questions?

CC image courtesy of Kiwithing
[http://www.flickr.com/photos/kiwisaotome/
8261132558/sizes/c/
in/photostream/](http://www.flickr.com/photos/kiwisaotome/8261132558/sizes/c/in/photostream/)