

- Help I need more IPv6 addresses!
  - Lets turn no into yes
    - Time crunch

# GOAL

- Get the community to decide on one of the two proposals
- Get the community to decide on any modifications to the preferred proposal
  - What you want removed
  - What you want added
- Get ARIN AC to write the preferred modifications and move it forward for ratification so the current roadblock to IPv6 is removed now and not 6 to 12 months down the road.

# Policy 2010-9 IPv6 for 6rd

*Policy statement:*

*6rd is an incremental method for Service Providers to deploy IPv6, defined in the IETF Standards Track RFC 5969.*

*If you have IPv4 addresses then you automatically qualify for IPv6 space for 6rd. Upon receipt of a 6rd request, an appropriate additional IPv6 allocation will be made that supports 6rd to be counted as a separate & parallel deployment to IPv4 and native IPv6. There is no requirement to segregate address space requested under this policy from regular IPv6 Allocation Supernets. From a management perspective this address space is to be treated as regular IPv6 address space.*

*While it is possible for an operator to transition to native IPv6 within the same address space used by 6rd, some operators may wish to keep native IPv6 users separate from 6rd users to permit optimization of aggregation. If an operator chooses to renumber users to an address space outside the 6rd aggregate when transitioning them to native IPv6, the 6rd allocation may be returned to ARIN when it is no longer in use.*

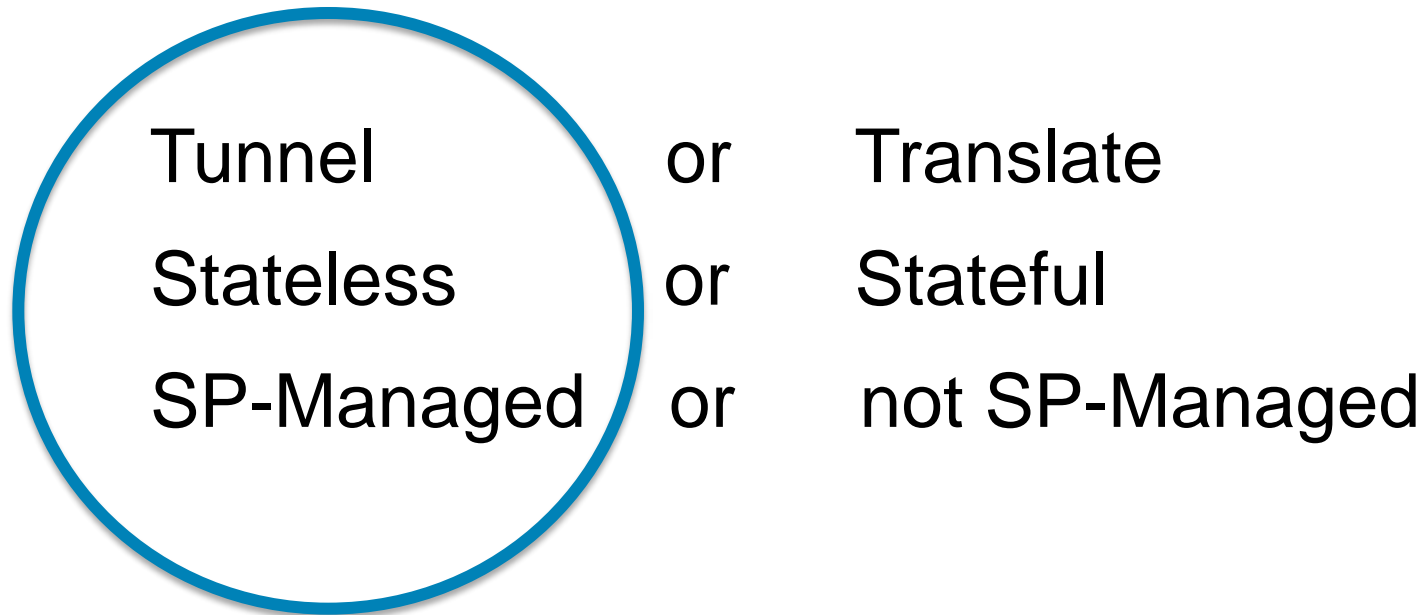
*It is suggested that contiguous allocations be made to any prior existing ones in the event justification for more IPv6 address space exists when the organization transitions 6rd out of their network.*

*Justification for use of IPv6 for 6rd will be reviewed after the first 3 years and reclaimed if it is not in use. After the first 3 years, any additional reviews will follow regular IPv6 policy. Requester will be exempt from returning all or a portion of the address space when 6rd is no longer used if they can show justification for need of the 6rd address space for other existing IPv6 addressing requirements be it native IPv6 or some other IPv6 network technology.*

# Instructive example (from DRAFT Draft Policy 2010-9)

- I. *Transition mechanisms which embed IPv4 addresses within ISP allocated prefixes allow for and SP-Managed, stateless operation. One such mechanism is “6rd” defined in RFC 5969.*
- II. 6rd is intended to be an incremental method for deploying IPv6 and bridge the gap for End Users to the IPv6 Internet. The method provides a native dual-stack service to a subscriber site by leveraging existing infrastructure. If an entity already has a /32 of IPv6 they can not use the same /32 for native IPv6 as they do for the 6rd routing and a separate minimum size of a /32 is required while a larger subnet like a /28 may be needed based on a non-contiguous IPv4 addressing plan.
- III. The 6rd prefix is an RIR delegated IPv6 prefix. It must encapsulate an IPv4 address and must be short enough so that a /56 or /60 can be given to subscribers. This example shows how the 6rd prefix is created based on a /32 IPv6 prefix using RFC1918 address space from 10.0.0.0/8:
- IV. SP IPv6 prefix: 2001:0DB8::/32 v4suffix-length: 24 (from 10/8, first octet (10) is excluded from the encoding) 6rd CE router IPv4 address: 10.100.100.1 6rd site IPv6 prefix: 2001:0DB8:6464:0100::/56
- V. This example shows how the 6rd prefix is created based on a /28 IPv6 prefix using one of several non-contiguous global address ranges:
- VI. SP IPv6 prefix: 2001:0DB0::/28 v4suffix-length: 32 (unable to exclude common bits due to non-contiguous IPv4 allocations) 6rd CE router IPv4 address: 192.0.2.1 6rd site IPv6 prefix: 2001:0DBC:0000:2010::/60

# Aspects of IPv6 Transition Mechanisms




**6rd is a Stateless, SP-Managed, Tunneling Protocol**

# IPv6 Prefix from an IPv4 Address


ISP's 6rd IPv6 Prefix

Subscriber IPv6 Prefix

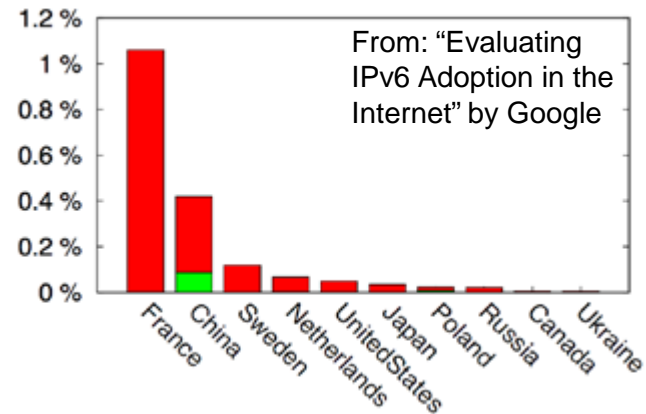


- The construction is what allows 6rd to be **stateless** and **SP-managed**

# 6rd - Packet Flow and Encapsulation




“...externally 6rd looks, feels and smells like native IPv6 ” –RIPE Labs



<http://labs.ripe.net/Members/emileaben/content-measuring-ipv6-web-clients-and-caching-resolvers-part-2-country-level-and-other-statistics>


"Evaluating IPv6 Adoption in the Internet" Lorenzo Colitti, Steinar H. Gunderson, Erik Kline, Tiziana Refice

# What should /n and /m be?



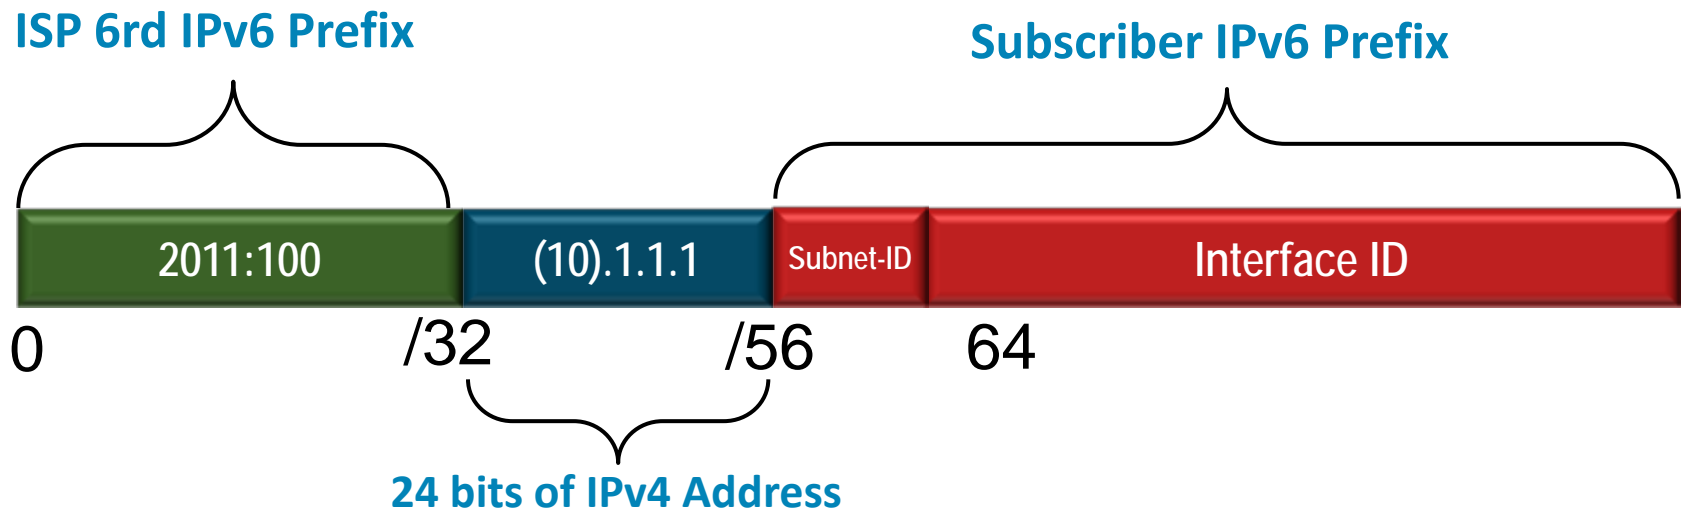


**n = /32, m = /64**



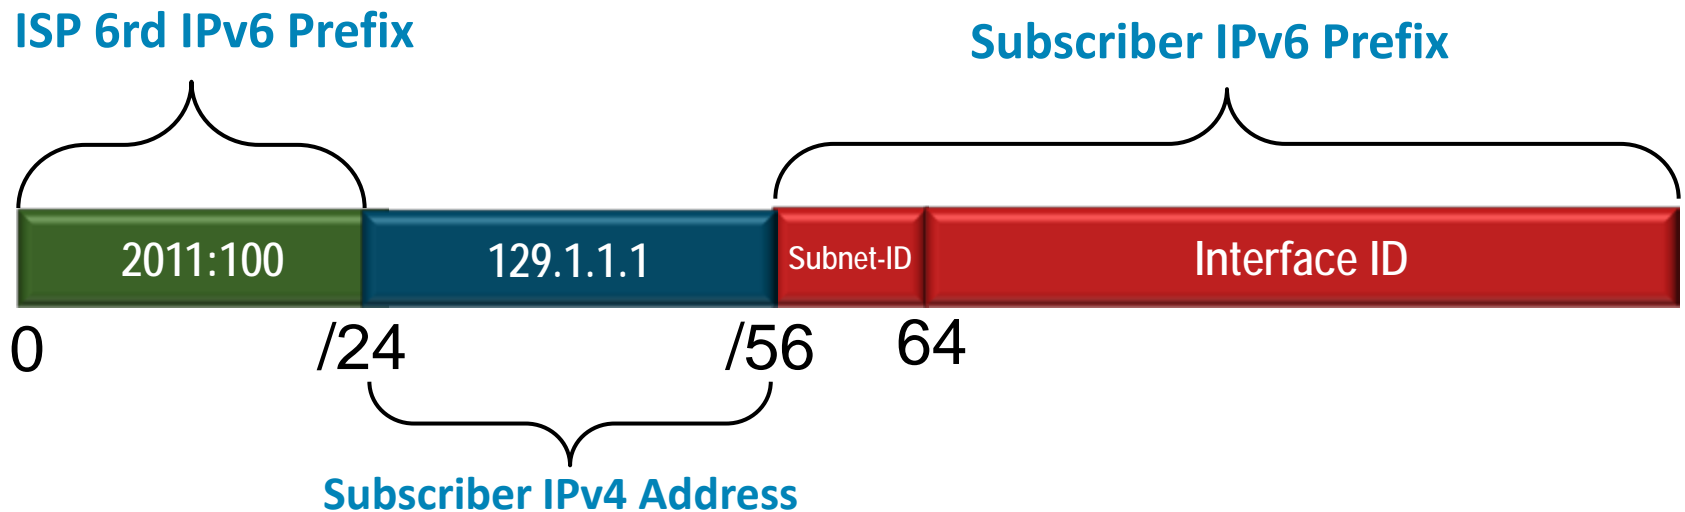
- /64 does not work with multiple subnets
  - Breaks existing residential IPv6 equipment
  - Does not allow IP routing, forces bridging
  - Leads to IPv6 NAT

**n = /32, m = /56**



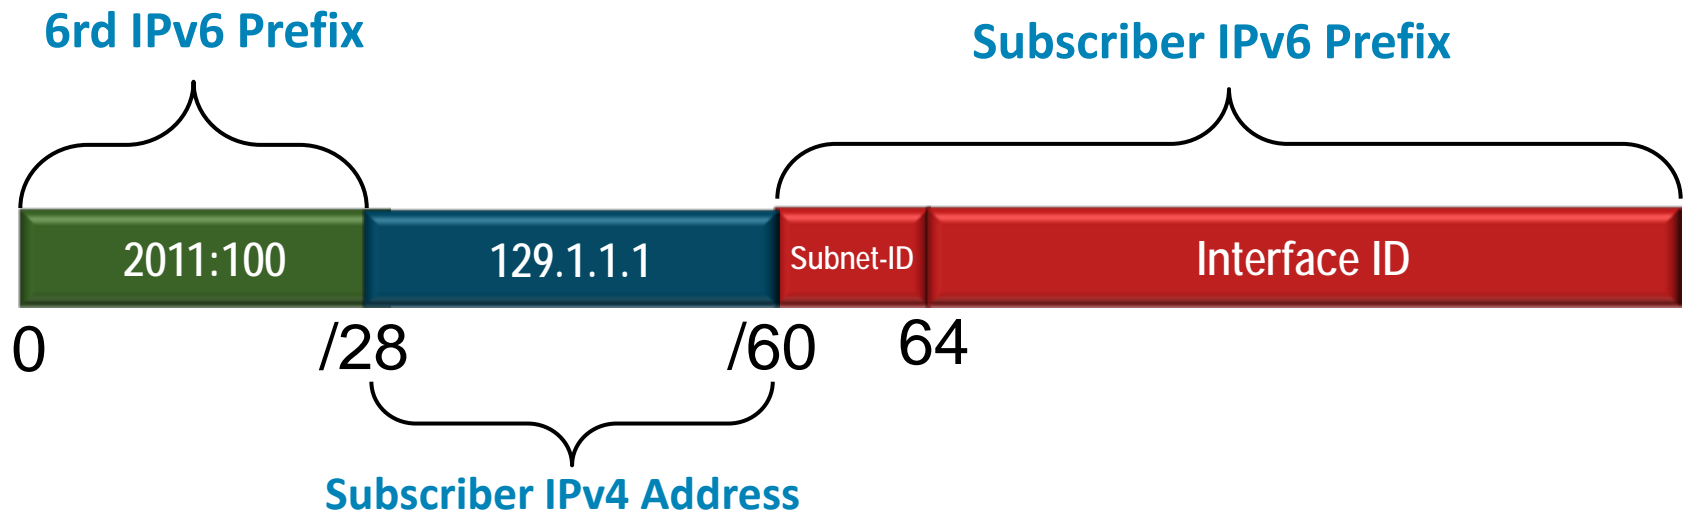
- Allows 256 subnets
- Works well with a CGN (10/8)
- Does not work well with a set of non-aggregated IPv4 blocks

**n = /24, m = /56**



- Allows up to 256 subnets in the home
- But /24 seems like overkill

# n = /28, m = /60 the right balance?



- Allows up to 16 subnets in the home, which is far better than zero
- Matches most 6rd deployments today

# PPML: Possible ways to encode a size?

- No organization can justify holding total IPv6 allocations that exceed /24 under this policy.

Rationale: If you think you need more than that, you haven't thought it through well enough.

- No organization can justify more than two disaggregate allocations under this policy irrespective of individual or total size.

Rationale: You get a couple tries but then you have to clear out and return one of your earlier tries before you can make attempt number three.

- Unless an organization is mapping more than “n” (5?) number of disaggregate IPv4 allocations with the transition mechanism the the largest additional allocation they can justify is a /32.

# IPv6 Subsequent Allocation

- Alternative proposal thanks to community input
- A cleaner proposal is born

# Draft Policy 2010-12: Subsequent for Transitional Technology

- Modify 6.5.2.1 Subsequent allocation criteria. ADD the following sentence:

**Subsequent allocations will also be considered for transitional technologies that cannot be accommodated by, nor were accounted for, under the initial allocation**

- Justification for the subsequent subnet size will be based on the plan and technology provided.

**Justification for these allocations will be reviewed every 3 years and reclaimed if it is not in use. Requester will be exempt from returning all or a portion of the address space if they can show justification for need of this allocation for other existing IPv6 addressing requirements be it Native V6 or some other V6 network technology.**

# Instructive example (from DRAFT Draft Policy 2010-9)

- I. *Transition mechanisms which embed IPv4 addresses within ISP allocated prefixes allow for and SP-Managed, stateless operation. One such mechanism is “6rd” defined in RFC 5969.*
- II. 6rd is intended to be an incremental method for deploying IPv6 and bridge the gap for End Users to the IPv6 Internet. The method provides a native dual-stack service to a subscriber site by leveraging existing infrastructure. If an entity already has a /32 of IPv6 they can not use the same /32 for native IPv6 as they do for the 6rd routing and a separate minimum size of a /32 is required while a larger subnet like a /28 may be needed based on a non-contiguous IPv4 addressing plan.
- III. The 6rd prefix is an RIR delegated IPv6 prefix. It must encapsulate an IPv4 address and must be short enough so that a /56 or /60 can be given to subscribers. This example shows how the 6rd prefix is created based on a /32 IPv6 prefix using RFC1918 address space from 10.0.0.0/8:
- IV. SP IPv6 prefix: 2001:0DB8::/32 v4suffix-length: 24 (from 10/8, first octet (10) is excluded from the encoding) 6rd CE router IPv4 address: 10.100.100.1 6rd site IPv6 prefix: 2001:0DB8:6464:0100::/56
- V. This example shows how the 6rd prefix is created based on a /28 IPv6 prefix using one of several non-contiguous global address ranges:
- VI. SP IPv6 prefix: 2001:0DB0::/28 v4suffix-length: 32 (unable to exclude common bits due to non-contiguous IPv4 allocations) 6rd CE router IPv4 address: 192.0.2.1 6rd site IPv6 prefix: 2001:0DBC:0000:2010::/60



# Two Draft Policy Proposals

## **2010-9: “IPv6 for 6rd” (latest rev: 24 Sept 2010)**

Con: Policy language not crisp and succinct

Con: Too specific to 6rd

## **2010-12: “IPv6 Subsequent Allocation” (20 July 2010)**

Pro: Crisp and succinct policy language

Pro: 6rd mentioned as an justification example only

Con: Mailing List feedback that it is missing stronger policy language on allocation size and limits

# How to proceed

# Which Draft?

## **2010-9: “IPv6 for 6rd” (latest rev: 24 Sept 2010)**

Con: Policy language not crisp and succinct

Con: Too specific to 6rd

## **2010-12: “IPv6 Subsequent Allocation” (20 July 2010)**

Pro: Crisp and succinct policy language

Pro: 6rd mentioned as an justification example only

Con: Mailing List feedback that it is missing stronger policy language on allocation size and limits

# Size

- Specify a size limit?
- Specify a standard one size fits all size?
- What size fits?
  - /32 where  $n = /32, m = /64$
  - /24 where  $n = /24, m = /56$
  - /28 where  $n = /28, m = /60$  the right balance?

# Review Terms

- Every 3 yrs?
- One review after 3 yrs then BAU?

# Backup

# 6rd - Packet Flow and Encapsulation

