

Network Working Group	S. Barber
Internet-Draft	Cox Communications
Intended status: Informational	O. DeLong
Expires: March 23, 2012	Hurricane Electric
	C. Grundemann
	CableLabs
	V. Kuarsingh
	Rogers Communications
	B. Schliesser
	Cisco Systems
	September 20, 2011

ARIN Draft Policy 2011-5: Shared Transition Space  
draft-bdgs-arin-shared-transition-space-03

## [Abstract](#)

This memo discusses the applicability of a Shared Transition Space, an IPv4 prefix designated for local use within service provider networks during the period of IPv6 transition. This address space has been proposed at various times in the IETF, and more recently come to consensus within the ARIN policy development community where it was recommended for adoption as Draft Policy 2011-5.

## [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 March 23, 2012.

## [Copyright Notice](#)

Copyright (c) 2011 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.

## [Table of Contents](#)

- \*1. [Introduction](#)
- \*2. [Applicability](#)
  - \*2.1. [Intended Use of Shared Transition Space](#)
    - \*2.1.1. [CGN](#)
    - \*2.1.2. [SP Services & Infrastructure](#)
    - \*2.1.3. [Note of Caution](#)
  - \*2.2. [Alternatives](#)
    - \*2.2.1. [Global Unicast Addresses](#)
    - \*2.2.2. [Private](#)
    - \*2.2.3. [Class E](#)
    - \*2.2.4. [Prefix Squatting](#)
    - \*2.2.5. [Regional Re-use of Allocated Prefix](#)
    - \*2.2.6. [Consortium](#)
- \*3. [Analysis of Benefits](#)
  - \*3.1. [Continued Operation Post-exhaustion](#)
  - \*3.2. [Delayed Need for CGN Deployment](#)
  - \*3.3. [Recovery of Existing Addresses](#)
    - \*3.3.1. [Re-deployment Where Needed](#)
    - \*3.3.2. [Return or Transfer](#)
  - \*3.4. [Impact on Allocations of RIR Inventory](#)
  - \*3.5. [Benefit of Standardization](#)
  - \*3.6. [IPv6 Deployments](#)
- \*4. [Analysis of Detractors' Arguments](#)
  - \*4.1. [It Breaks](#)
    - \*4.1.1. [NAT is Bad](#)

- \*4.1.2. [Breaks Assumptions about Address Scope](#)
- \*4.1.2.1. [6to4](#)
- \*4.1.3. [Potential Misuse as Private Space](#)
- \*4.2. [It's Not Needed](#)
- \*4.2.1. [Nobody Will Use It](#)
- \*4.2.2. [ISPs Are Not Actually Growing](#)
- \*4.2.3. [RIR IPv4 Inventory is Not Actually Exhausted](#)
- \*4.2.4. [ISP IPv4 Inventory is Not Actually Exhausted](#)
- \*4.3. [Address Inventory](#)
- \*4.3.1. [Shared Transition Space Uses Up Address Inventory](#)
- \*4.3.2. [/10 is not Enough](#)
- \*4.4. [IPv6 Arguments](#)
- \*4.4.1. [Use IPv6 Instead](#)
- \*4.4.2. [Delay of IPv6 Deployment](#)
- \*5. [ARIN Draft Policy 2011-5](#)
- \*5.1. [History](#)
- \*5.1.1. [Shared Address Space](#)
- \*5.1.2. [Proposal](#)
- \*5.2. [Policy Text](#)
- \*6. [Acknowledgements](#)
- \*7. [IANA Considerations](#)
- \*8. [Security Considerations](#)
- \*9. [References](#)
- \*[Authors' Addresses](#)

## **1. Introduction**

As the Internet community approaches exhaustion of unallocated IPv4 numbers, the value of globally unique addresses is becoming manifest. More than ever network operators recognize the need to transition to the IPv6 address family. However, the immediate necessity of continued IPv4 connectivity poses a near-term challenge - without adequate IPv4 resources, most network operators must deploy more efficient addressing architectures and many must deploy address-sharing technologies.

In order to facilitate these operators' need for near-term IPv4 connectivity, [\[I-D.weil-shared-transition-space-request\]](#) proposes the reservation of a /10 IPv4 prefix for use in Service Provider (SP) networks. Referred to as Shared Transition Space, this address block would facilitate SP deployment of non-unique address plans that do not conflict with traditional Private [\[RFC1918\]](#) address space. By using the Shared Transition Space operators may deploy CGN [\[I-D.ietf-behave-lsn-requirements\]](#) internal networks, extranet [\[RFC4364\]](#) communities, and/or SP-local services without consuming Global Unicast Addresses.

However, given the Feb 2011 depletion of the IANA Free Pool inventory [\[NRO-IANA-exhaust\]](#) it is not currently possible for the IANA to reserve an IPv4 /10 prefix as recommended in [\[I-D.weil-shared-transition-space-request\]](#). Thus the ARIN community has proposed in Draft Policy [\[ARIN-2011-5\]](#) the reservation of a Shared Transition Space from the ARIN inventory of unallocated IPv4 numbers. After much discussion by the ARIN community, [\[ARIN-2011-5\]](#) reached consensus and was recommended by the ARIN Advisory Council for approval by the ARIN Board of Trustees.

Following the community's recommendation of [\[ARIN-2011-5\]](#) the ARIN Board requested clarification from the IAB with regard to responsibilities outlined in [\[RFC2860\]](#). The ARIN Board received a response in [\[IAB-response\]](#) indicating that the IETF holds responsibility for the reservation of specialized address blocks. Thus, the ARIN Board believes that it is not within ARIN's authority to unilaterally make specialized allocations of the sort proposed in Draft Policy 2011-5. [\[PPML-022778\]](#)

This memo explains the intended use and discusses the merits and drawbacks of using Shared Transition Space.

## **2. Applicability**

### **2.1. Intended Use of Shared Transition Space**

The Shared Transition Space is intended for use by service providers and should not be thought of as additional RFC1918 space. There are a number of specific use-cases for the Shared Transition Space. This section discusses the primary scenarios envisioned at the time of this writing. Equipment vendors and non-ISP network operators should be aware that using the Shared Transition Space outside of its intended scope may result in unpredictable behavior.

### **2.1.1. CGN**

The primary use-case for the Shared Transition Space will be deployment in CGN [\[I-D.ietf-behave-lsn-requirements\]](#) internal networks. A key benefit of CGN is the ability to share a smaller number of Global Unicast Addresses (GUA) amongst a larger number of end-sites.

In one CGN deployment scenario sometimes referred to as NAT444 [\[I-D.shirasaki-nat444\]](#), the CGN internal network is numbered with IPv4 addresses that are not globally routed while the end-sites are numbered with Private [\[RFC1918\]](#) addresses. In this scenario the Shared Transition Space will be used to provide contextually unique IPv4 addresses to end-site CPE devices and intermediate infrastructure. [\[I-D.shirasaki-nat444-isp-shared-addr\]](#)

### **2.1.2. SP Services & Infrastructure**

In networks that contain local services (such as nameservers, content repositories or caches, etc) the Shared Transition Space will offer an alternative to GUA. For instance, video content servers that are available only to customers directly connected to the SP network might be addressed from the Shared Transition Space, preserving GUA for services that require global connectivity. Where these services are accessed by customers who have their own IPv4-only equipment, use of the Shared Transition Space will reduce or eliminate the need for NAT. Similarly, those infrastructure elements which touch IPv4-only customer-managed equipment could also be numbered from the Shared Transition Space. In cases where the provider manages both endpoints, IPv6 should be used.

### **2.1.3. Note of Caution**

In any case, care must be taken to ensure the Shared Transition Space is not used in scenarios where routing may be ambiguous. For instance, when multiple provider networks may be simultaneously reachable the use of Shared Transition Space might result in address conflicts etc. Conversely, operators may choose to allow (not filter) ICMP messages from the Shared Transition Space in order to enable Path MTU Discovery etc. This topic requires further investigation so that best practices may be developed.

## **2.2. Alternatives**

A number of possible alternatives to Shared Transition Space have been proposed and/or discussed by the Internet community. See, for instance, [\[RFC6319\]](#) for a discussion of alternatives and potential issues. This section outlines these possible alternatives and briefly discusses their applicability.

### **2.2.1. Global Unicast Addresses**

Every discussion of the Shared Transition Space begins with an assumption that Global Unicast Addresses (GUA) are a preferable choice for numbering. This is almost always technically true. However, given the fundamental driver of IPv4 address exhaustion, GUA is not a pragmatic alternative to the Shared Transition Space. Additionally, if various organizations use various GUA ranges to number CGN zones, it will be difficult for other networks and/or systems to deterministically know if the endpoints are using true Internet reachable IPs, or if the source network may be using them as CGN zone space. This situation would likely lead to additional technical issues during various leakage conditions, filter rule issues (routing) and for CDN or other third party providers who may be present within the source network, to name a few.

### **2.2.2. Private**

In each of the use-cases for Shared Transition Space, it may be possible to instead use Private [\[RFC1918\]](#) address space. In situations where all endpoints in the network are managed by a single organization, this may be a viable option. However when end-sites are administered by different organizations and/or individuals, the possibility of address conflict becomes a significant risk to operations. Private [\[RFC1918\]](#) address space is not generally intended to be used for purposes which cross administrative domains. Further, these recommendations involve use of the Shared Transition Space to provide services in one administrative domain to leaf networks which are generally single-homed to the serving administrative domain. This is also a significant difference from the intent of Private [\[RFC1918\]](#) address space.

A study of DNS traffic [\[v6ops-msg06187\]](#) has shown that effectively all of the existing Private [\[RFC1918\]](#) address space is currently being used by end-sites attached to the Internet. While individual network environments may vary in this regard, most SP operators face the risk that their use of Private address space will conflict with their customer end-sites. defined private space is not generally intended to be used for purposes which cross administrative domains.

In the event of conflict, it is possible that the end-site CPE will fail and/or not function correctly. Some CPE implementations are known to support overlapping addresses on the "inside" and "outside" interfaces, however many others are known to fail under such circumstances. For SP operators, the Shared Transition Space offers a less risky alternative to GUA that retains the benefit of non-conflict. Also, the use of Private [\[RFC1918\]](#) address space on interfaces and hosts often causes default behaviors on such hosts which may not be desirable when the endpoint is actually connected to the Internet. There are often behavioral expectations for Internet connected endpoints, regardless of them being subject to a NAT.

Incorrect affiliation of the WAN side interface being in a "protected" zone and/or on a trusted network may not be desirable. With NAT444 deployments, it is important that the endpoint (i.e. CPE) behave like any other Internet node. One example of this from our testing was observed behaviors where some CPEs did not filter and/or firewall correctly when Private [\[RFC1918\]](#) address space was used on both WAN and LAN interfaces.

### **2.2.3. Class E**

One proposed alternative to Shared Transition Space is the re-classification and use of the 240.0.0.0/4 "Class E" address space as unicast. This has been proposed, for instance, by [\[I-D.fuller-240space\]](#) and [\[I-D.wilson-class-e\]](#). While this alternative might be possible in tightly constrained environments, where all of the network elements are known to support Class E address space, it is not generally useful in the use-cases described above. At this time, a significant number of IPv4 stack implementations treat the Class E address space as reserved and will not route, forward, and/or originate traffic for that range. For example, [\[CISCO\]](#) states that: "No addresses are allowed with the highest-order bits set to 1111." For the scenarios described herein, it should be noted that this alternative would create additional SP dependencies on customer selected CPE support for Class E addressing.

### **2.2.4. Prefix Squatting**

An unfortunate alternative to the Shared Transition Space is "prefix squatting", in which the operator re-uses another organization's IPv4 allocation for their own numbering needs. When this approach results in the other organization's prefix being announced globally by the "squatting" operator, it is often referred to as "prefix hijacking". However, this discussion is focused on scenarios in which the prefix is not announced globally but is, rather, used for internal numbering only.

In this scenario, the allocation may not be routed globally by the legitimate address holder, making it attractive for such purposes. Or it may be routed but "uninteresting" to the SP network's endpoints. In either case there is a potential for conflict in the event that any end-site actually wishes to communicate with the legitimate address holder. Indeed, various RIRs attempt to discover and "recycle" abandoned or unused IPv4 address space, making it more likely that such conflicts will be experienced in time. As such, this alternative is to be discouraged with prejudice.

It is important to note that there are no behavioral advantages to using "squat space" over using assigned "shared space". Both options subject the CPE to the same general behaviors (GUA space, but not globally reachable). The only real difference is the negative impacts of squatting (as noted above) and the advantages of a community coordinated and standardized prefix.

The primary reason that any network would be likely to adopt "prefix squatting" is if they are faced with the operational realities of CGN before/without the allocation of a shared transition space.

#### **2.2.5. Regional Re-use of Allocated Prefix**

Similar to "Prefix Squatting" but significantly less dangerous, this alternative involves the reuse by an operator of their own address allocations. In this scenario, a network operator might use the same prefix for multiple "regions" and/or extranet communities. For instance, in CGN deployments the operator might reuse the same GUA prefix across multiple geographic regions (e.g. without announcing it globally).

Here again, it is important to note that there are no behavioral advantages gained over a "shared space" but there is the added community cost of each network having to dedicate a unique block of addresses to this purpose, consuming far more resources than a single block of "shared space".

#### **2.2.6. Consortium**

In the event that the Internet community doesn't set aside an IPv4 prefix for Shared Transition Space, it is possible that a number of SP operators can come together and designate an address block to be "shared" amongst them for an identical purpose. This would have the same technical merits as an IETF and/or RIR sponsored Shared Transition Space, however it would lack the efficiency of a community coordinated and standardized prefix for such purposes, gain no behavioral advantages, remove the deterministic nature of managing a single range and also subjects the Internet (users of the space) to additional risk since any member of the consortium who has contributed space could later pull out and potentially cause disruptions in multiple networks.

### **3. Analysis of Benefits**

#### **3.1. Continued Operation Post-exhaustion**

Availability of a Shared Transition Space helps SPs continue to meet the demands of IPv4 addressing and/or connectivity post exhaustion. For environments where CGN in a NAT444 scenario is necessary, addresses from this space can be used to provide addressing for the network between the CGN device(s) and CPE which will enable IPv4 flow continuity for customers using these services. In other circumstances, the shared transition space allows SPs to number devices in the network which do not require global reachability without the need for fulfillment thorough an RIR.



### **3.2. Delayed Need for CGN Deployment**

If operators are required to use their individually allocated GUA where "shared space" would have applied, e.g. for internal services, they will face exhaustion sooner and thus be forced to deploy CGN sooner as well. Operators may be able to postpone the deployment of CGN by using "shared space" for internal uses, because that allows more efficient use of their remaining GUA in places where global uniqueness is truly mandatory.

Further, without this shared transition space, some service providers may be forced to reclaim GUA from existing customers in order to deploy CGN and address the required infrastructure. Having this transition space will enable deployment of CGN where it is required, in a manner that is less disruptive and with impact to fewer customers.

### **3.3. Recovery of Existing Addresses**

The shared transition space can also be used to number and reclaim IPv4 addresses within provider networks which do not require global reachability. This option can be used by many networks worldwide, it provides an option for using currently assigned space much more efficiently.

#### **3.3.1. Re-deployment Where Needed**

Operators can re-deploy recovered addresses for customers that need them (including new / static / GUA customers), hosted servers, etc. or to facilitate other efforts that might provide even more efficient use of GUA space within the network. The freed addresses can be assigned to endpoints which require IPv4 global reachability and thus help delay and/or remove the need for CGN.

#### **3.3.2. Return or Transfer**

In cases where the operator is not deploying CGN and doesn't need the recovered addresses, they can be made available to others that do need them for connectivity to the public IPv4 Internet. This may be through voluntary return to the RIR, or through transfer to another network operator. For example, in the ARIN region, there are transfer mechanisms defined in the ARIN NRPM 8.3 [\[ARIN-NRPM-8.3\]](#).

### **3.4. Impact on Allocations of RIR Inventory**

While making Shared Transition Space available to the community may or may not lessen the demand on the RIRs for allocations, it will help ensure that the address resources which remain in inventory are used most efficiently, maximizing the use of that inventory for services that require Global Unicast Addresses.

### **3.5. Benefit of Standardization**

Standardizing on a single block will help the community develop standard ways of selecting, routing, filtering and managing shared space. This task would be much more difficult or impractical for any of the alternative options.

Standard internal routing policy and filtering can be applied uniformly inside network environments. Additionally, exchange points between networks can have standard policies applied allowing operators to protect each other from CGN zone IPs leaking between networks. This may not be possible with squat space since many operators will not divulge what space may be used and with Private [\[RFC1918\]](#) address space where each operator may only be able to free up certain portions of the space which are not likely to be consistent between networks.

### **3.6. IPv6 Deployments**

Operators will need to grapple with the need to provide IPv4 based flow continuity to customers post exhaustion. By removing the burden of operators needing to find adequate IPv4 address space to meet the needs that a Shared Transition Space can fulfill, they can concentrate on the real task at hand: Deploying IPv6.

## **4. Analysis of Detractors' Arguments**

### **4.1. It Breaks**

#### **4.1.1. NAT is Bad**

NAT is understood to be less than optimal [\[RFC6269\]](#), especially when implemented as CGN [\[I-D.donley-nat444-impacts\]](#). That said, it is a necessary technology for many networks and cannot be completely avoided. Since the number of IPv4 Internet endpoints will exceed the number of IPv4 addresses which are available for Internet connectivity, NATs are needed.

While the authors agree that "NAT is bad", it must also be understood that shared transition space does not change the fundamental motivations or issues with NAT and so those problems will not be discussed at length here.

#### **4.1.2. Breaks Assumptions about Address Scope**

Some host or CPE functions incorrectly assume global reachability based on the type of address that is configured, potentially causing issues when deployed in a NAT444 scenario. Whether an operator uses this proposed Shared Transition Space or some other GUA space (e.g. through squatting or reuse), the net effect on hosts and/or CPE making such assumptions about reachability is identical. Conversely, with an identified Shared Transition Space hosts that make these mistaken assumptions can be modified to treat the identified block as having

restricted reachability semantics. This would not be possible (or at least not nearly as easy) with the other solutions.

#### **4.1.2.1. 6to4**

Although 6to4 can break in CGN scenarios using the Shared Transition Space, recent guidance suggests that it should be turned off by default. [\[RFC6343\]](#) [\[I-D.ietf-v6ops-6to4-to-historic\]](#) Indeed, recent versions of operating systems de-preference 6to4 addresses as described in [\[I-D.ietf-6man-rfc3484-revise\]](#), mitigating effects from incorrect 6to4 instantiation behind a firewall that obstructs its function. Since the volume of impacted endpoints will be low, operators can likely manage the disabling of 6to4 when needed. More fundamentally, broken 6to4 should not be an issue if service providers deploy (and user equipment supports) native IPv6 connectivity.

#### **4.1.3. Potential Misuse as Private Space**

Shared Transition Space is intended to be used solely by Service Providers for IPv4 to IPv6 transition purposes. [\[I-D.weil-shared-transition-space-request\]](#) The value of a Shared Transition Space may be diminished if commonly misused by end-sites as generic Private addresses. Thus, the reservation must be clearly designated for use by SPs that are providing infrastructure as described herein.

### **4.2. It's Not Needed**

#### **4.2.1. Nobody Will Use It**

This argument is simply incorrect. Post IPv4-exhaustion, any SP that wishes to continue providing IPv4 connectivity will necessarily deploy network architectures and technologies that require such an address space. Thus, in absence of a designated Shared Transition Space, operators will use GUA space in essentially the same ways described in this memo, with or without IETF or RIR acknowledgement.

#### **4.2.2. ISPs Are Not Actually Growing**

While customer growth for some ISPs has slowed, for many service providers new services are growing at a faster rate than has been anticipated. Wireline voice customers for example require two-way communication paths to allow them to function properly. IP enabled televisions is another example of devices that support video and voice services and require IP addresses. The only way to maintain these services, which in many cases are considered lifeline, is to provide them with an IP address that is unique within the service provider network.

Likewise, growth continues to exist in some geographical regions. While some areas have slower growth, as a result of significant penetration

of Internet access, there are still many areas with unmet needs, growing populations, or both.

#### **4.2.3. RIR IPv4 Inventory is Not Actually Exhausted**

With the IANA inventory essentially exhausted [\[NRO-IANA-exhaust\]](#) it is only a matter of time before each of the RIRs are unable to satisfy requests for IPv4 addresses. [\[GIH-When\]](#) In fact, the APNIC has already allocated all but their final /8 of inventory [\[APNIC-final-slash8\]](#) and is no longer making allocations larger than a /22 prefix. Each of the other RIRs is on a trajectory toward exhaustion in the near future.

#### **4.2.4. ISP IPv4 Inventory is Not Actually Exhausted**

While some SPs have existing inventory that will outlast the RIR inventories, this is not universally true. In fact, the distribution of IPv4 number resources amongst operators is highly variable (based on size, history, etc) and in the worst cases is already becoming problematic.

### **4.3. Address Inventory**

#### **4.3.1. Shared Transition Space Uses Up Address Inventory**

While true that this Shared Transition Space will remove a block of global unicast IPv4 addresses from the free pool, it must also be noted that the use of the same "shared space" repeatedly across multiple networks will very likely increase the available pool of unique IPv4 addresses through operational efficiency. For example, if just two operators use their own GUA /10, the Internet community effectively loses a /9 of unique space while if both operators use the same "shared" /10, the Internet community loses that single /10. This benefit becomes more significant as more operators use the Shared Transition Space.

It remains to be seen whether the reservation of a Shared Transition Space will actually delay the impending exhaustion of RIRs' IPv4 inventory. Certainly, the availability of this Shared Transition Space will satisfy a number of demands that would otherwise become requests for GUA resources. However, whether this translates to an actual reduction in requests is up to the RIRs and requesting organizations. Regardless of the allocation of Shared Transition Space, RIR IPv4 exhaustion may happen at roughly the same time. However, as noted above, Shared Transition Space does provide the opportunity for more efficient use of the remaining RIR IPv4 addresses. Additionally, the reservation of a Shared Transition Space will enable continued deployment of IPv4 connectivity by SP networks beyond the free pool depletion horizon; another clear benefit.

#### **4.3.2. /10 is not Enough**

Although previous requests for Shared Transition Space asked for a full /8, it has been determined by many operators that a /10 will in fact be sufficient. A /10 provides for roughly 4 million hosts and although many of the largest SPs have subscriber counts in the tens of millions, none will be placing all of their subscribers behind a single CGN. In the event that a /10 does not provide enough addresses for an operators entire CGN deployment, it could be re-used multiple times in distinct "NAT zones" or regions.

#### **4.4. IPv6 Arguments**

##### **4.4.1. Use IPv6 Instead**

Although IPv6 is the strategic long term answer for IPv4 address exhaustion, it does not immediately solve IPv4 connectivity requirements. There is an entire eco-system which exists on the Internet today and is not IPv6 ready at this time [[I-D.arkko-ipv6-only-experience](#)]. IPv4 flow continuity will be required for at least several years.

Many businesses have long procurement and fulfillment cycles which will need to be used to upgrade networks to support IPv6. Also, the consumer (home) space is years away from being all IPv6 capable. Many homes are filled with IPv4 only consumer electronics, computers, TVs, accessories and other systems.

There are still a number of products that are either not IPv6 compliant, or for which the necessary criteria for being "IPv6 compliant" is unclear or undefined. Some examples include security products, a large number of software applications, and there are still production systems (both inside companies and as products) being rolled out that are not IPv6 aware.

##### **4.4.2. Delay of IPv6 Deployment**

The whole Internet needs to get to IPv6 more or less at the same time in order to avoid significant deployment of transition technologies. This proposal may help delay some transition technology deployment while IPv6 deployments move ahead. More IPv6 should mean less transition technology.

#### **5. ARIN Draft Policy 2011-5**

##### **5.1. History**

###### **5.1.1. Shared Address Space**

Proposals for additional Private space date back at least to [[I-D.hain-1918bis](#)] in April 2004. Some of these proposals and surrounding discussion may have considered similar use-cases as described herein.

However, they were fundamentally intended to increase the size of the [\[RFC1918\]](#) pool, whereas a defining characteristic of Shared Transition Space is that it is distinctly not part of the Private [\[RFC1918\]](#) address pool.

Rather, the Shared Transition Space is reserved for more narrow deployment scenarios, specifically for use by SPs to meet the needs of ongoing IPv4 connectivity during the period of IPv6 transition. This key feature emerged in more recent proposals such as [\[I-D.shirasaki-isp-shared-addr\]](#) in June 2008 where a proposal to set aside "ISP Shared Space" was made. During discussion of these more recent proposals, many of the salient points were commented upon, including challenges with RFC1918 in the ISP NAT Zone, Avoidance of IP Address Conflicts, and challenges with 240/4.

This effort was followed up by [\[I-D.weil-opsawg-provider-address-space\]](#) in July 2010 which was re-worked in November 2010 as [\[I-D.weil-shared-transition-space-request\]](#). These latter efforts continued to point out the operators' need for Shared Transition Space, with full acknowledgement that challenges may arise with NAT444 as per [\[I-D.donley-nat444-impacts\]](#) and that such space does not reduce the need for IPv6 deployment.

Most recently, following exhaustion of the IANA's /8 pool [\[NRO-IANA-exhaust\]](#), this proposal has been discussed by various RIR policy development participants. As described herein, the body of ARIN policy development participants has reached consensus and recommended a Shared Address Space policy for adoption [\[ARIN-2011-5\]](#).

This history shows that operators and other industry contributors have consistently identified the need for a Shared Transition Space assignment, based on pragmatic operational needs. The previous work has also described the awareness of the challenges of this space, but points to this option as the most manageable and workable solution for IPv6 transition space.

### **[5.1.2. Proposal](#)**

The following is a brief history of the proposal for Shared Address Space within ARIN, ultimately resulting in the recommendation of ARIN Draft Policy 2011-5 [\[ARIN-2011-5\]](#).

In January 2011, a policy was proposed to the ARIN policy development community called ARIN-prop-127: Shared Transition Space for IPv4 Address Extension [\[ARIN-prop-127\]](#). This policy proposal would reserve an IPv4 /10 prefix by ARIN, to be shared by any Service Providers who wish to use it with no further registration actions required.

After generating much discussion (over 100 posts) on the ARIN Public Policy Mailing List (PPML), the ARIN Advisory Council (AC) accepted the proposal as Draft Policy 2011-5 [\[ARIN-AC-28Jan2011\]](#), formally announced on PPML 3 February 2011 [\[ARIN-2011-5-AC\]](#).

On 14 February 2011, ARIN staff made the following statement with regard to 2011-5: "In keeping with the spirit of RFC 2860 with respect to the assignment of specialized address blocks, ARIN Staff will

consult with the IANA and the IAB regarding implementation of this draft policy." [\[ARIN-2011-5-Staff\]](#)

In the ensuing PPML discussion there was a roughly two to one ratio of those clearly in support of the policy versus those clearly against. ARIN Draft Policy 2011-5 was then discussed at the ARIN XXVII public policy meeting on 12 April 2011. Following the discussion, there was a straw poll of the room. With a total number of people in the meeting room and remote of 116; in favor of it were 30 and against it were 15. [\[ARIN27.2011-5\]](#)

Seeing an obvious need in the service provider community, the AC voted to send the Draft Policy to last call [\[ARIN-AC-13Apr2011\]](#) for final comments 18 April through 2 May 2011. [\[ARIN-2011-5-LC\]](#) Following a strong show of support from the operator community during last call, the AC voted [\[ARIN-AC-19May2011\]](#) to recommend adoption of 2011-5 to the ARIN Board of Trustees with a vote of 10 in favor and 2 abstentions. [\[ARIN-2011-5-Rec\]](#)

Following this recommendation, ARIN staff consulted with the IAB and IANA as committed. The IAB response [\[IAB-response\]](#) stated, in short, that they believed the adoption of [\[ARIN-2011-5\]](#) was in conflict with the provisions in [\[RFC2860\]](#) and requested that the community re-review the operational and technical merits of shared transition space in the IETF. That process is now underway, with this draft an attempt at more fully analyzing said operational and technical merits.

## **[5.2. Policy Text](#)**

Draft Policy ARIN-2011-5

Shared Transition Space for IPv4 Address Extension

Date: 20 January 2011

Policy statement:

Updates 4.10 of the NRPM:

A second contiguous /10 IPv4 block will be reserved to facilitate IPv4 address extension. This block will not be allocated or assigned to any single organization, but is to be shared by Service Providers for internal use for IPv4 address extension deployments until connected networks fully support IPv6. Examples of such needs include: IPv4 addresses between home gateways and NAT444 translators.

Rationale:

The Internet community is rapidly consuming the remaining supply of unallocated IPv4 addresses. During the transition period to IPv6, it is imperative that Service Providers maintain IPv4 service for devices and networks that are currently incapable of upgrading to IPv6. Consumers must be able to reach the largely IPv4 Internet after exhaustion.

Without a means to share addresses, people or organizations who gain Internet access for the first time, or those who switch providers, or move to another area, will be unable to reach the IPv4 Internet.

Further, many CPE router devices used to provide residential or small-medium business services have been optimized for IPv4 operation, and typically require replacement in order to fully support the transition



to IPv6 (either natively or via one of many transition technologies). In addition, various consumer devices including IP-enabled televisions, gaming consoles, medical and family monitoring devices, etc. are IPv4-only, and cannot be upgraded. While these will eventually be replaced with dual-stack or IPv6 capable devices, this transition will take many years. As these are typically consumer-owned devices, service providers do not have control over the speed of their replacement cycle. However, consumers have an expectation that they will continue to receive IPv4 service, and that such devices will continue to have IPv4 Internet connectivity after the IPv4 pool is exhausted, even if the customer contracts for new service with a new provider.

Until such customers replace their Home Gateways and all IPv4-only devices with IPv6-capable devices, Service Providers will be required to continue to offer IPv4 services through the use of an IPv4 address sharing technology such as NAT444. A recent study showed that there is no part of RFC1918 space which would not overlap with some IPv4 gateways, and therefore to prevent address conflicts, new address space is needed.

Service providers are currently presented with three options for obtaining sufficient IPv4 address space for NAT444/IPv4 extension deployments: (1) Request allocations under the NRPM; (2) share address space with other providers (this proposal); or (3) use address space allocated to another entity (i.e. 'squat'). Of the three options, option 2 (this proposal) is preferable, as it will minimize the number of addresses used for IPv4 extension deployments while preserving the authority of IANA and RIRs.

Timetable for implementation: immediately

## 6. Acknowledgements

The authors would like to thank the following individuals for their contributions: John Curran, David Farmer, Jeffrey Finkelstein, William Herrin, and Dan Wing.

The authors would also like to thank the following people for their review, comments, and support: Gary Buhrmaster, Chris Donley, Wes George, Chris Metz, Richard Von Scherr, and Lane Wigley.

## 7. IANA Considerations

Upon notification by the IAB that that an address reservation should be made, ARIN is willing to proceed with the implementation of its Draft Policy 2011-5 which would result in ARIN reserving IPv4 /10 block for shared transition. The IANA is to record the allocation of the IPv4 address block for this purpose. Alternatively, the IAB may direct the IANA to request return of sufficient address space from ARIN's available IPv4 number resource pool to allow the IANA to perform this reservation directly.



## 8. Security Considerations

This memo makes reference to a number of deployment scenarios that have unique security considerations, and the reader is advised to investigate these independently.

While this memo does not introduce any specific technical issues that may be subject to detailed security considerations, it does recommend the reservation of a new IPv4 address space that might have unique properties when deployed. As such, all implementors of this Shared Transition Space are encouraged to consider carefully the best practices associated with the use of this space, including considerations relating to filtering, routing, etc.

## 9. References

, "

[RFC1918]	<a href="#">Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G. and E. Lear, "Address Allocation for Private Internets"</a> , BCP 5, RFC 1918, February 1996.
[RFC4364]	Rosen, E. and Y. Rekhter, " <a href="#">BGP/MPLS IP Virtual Private Networks (VPNs)</a> ", RFC 4364, February 2006.
[RFC2860]	Carpenter, B., Baker, F. and M. Roberts, " <a href="#">Memorandum of Understanding Concerning the Technical Work of the Internet Assigned Numbers Authority</a> ", RFC 2860, June 2000.
[I-D.ietf-behave-lsn-requirements]	Perreault, S, Yamagata, I, Miyakawa, S, Nakagawa, A and H Ashida, " <a href="#">Common requirements for Carrier Grade NAT (CGN)</a> ", Internet-Draft draft-ietf-behave-lsn-requirements-03, August 2011.
[I-D.donley-nat444-impacts]	Donley, C, Howard, L, Kuarsingh, V, Chandrasekaran, A and V Ganti, " <a href="#">Assessing the Impact of NAT444 on Network Applications</a> ", Internet-Draft draft-donley-nat444-impacts-01, October 2010.
[I-D.weil-shared-transition-space-request]	Weil, J, Kuarsingh, V, Donley, C, Liljenstolpe, C and M Azinger, " <a href="#">IANA Reserved IPv4 Prefix for Shared Transition Space</a> ", Internet-Draft draft-weil-shared-transition-space-request-03, August 2011.
[I-D.weil-opsawg-provider-address-space]	Weil, J, Kuarsingh, V and C Donley, " <a href="#">IANA Reserved IPv4 Prefix for IPv6 Transition</a> ", Internet-Draft draft-weil-opsawg-provider-address-space-02, September 2010.
[I-D.shirasaki-isp-shared-addr]	Yamagata, I, Miyakawa, S, Nakagawa, A, Yamaguchi, J and H Ashida, " <a href="#">ISP Shared</a>

	<a href="#">Address</a> ", Internet-Draft draft-shirasaki-isp-shared-addr-06, July 2011.
[I-D.shirasaki-nat444-isp-shared-addr]	Yamaguchi, J, Shirasaki, Y, Miyakawa, S, Nakagawa, A and H Ashida, " <a href="#">NAT444 addressing models</a> ", Internet-Draft draft-shirasaki-nat444-isp-shared-addr-06, July 2011.
[I-D.shirasaki-nat444]	Yamagata, I, Shirasaki, Y, Nakagawa, A, Yamaguchi, J and H Ashida, " <a href="#">NAT444</a> ", Internet-Draft draft-shirasaki-nat444-04, July 2011.
[I-D.ietf-6man-rfc3484-revise]	Matsumoto, A, Kato, J, Fujisaki, T and T Chown, " <a href="#">Update to RFC 3484 Default Address Selection for IPv6</a> ", Internet-Draft draft-ietf-6man-rfc3484-revise-04, July 2011.
[RFC6269]	Ford, M., Boucadair, M., Durand, A., Levis, P. and P. Roberts, " <a href="#">Issues with IP Address Sharing</a> ", RFC 6269, June 2011.
[RFC6319]	Azinger, M. and L. Vegoda, " <a href="#">Issues Associated with Designating Additional Private IPv4 Address Space</a> ", RFC 6319, July 2011.
[I-D.fuller-240space]	Fuller, V, " <a href="#">Reclassifying 240/4 as usable unicast address space</a> ", Internet-Draft draft-fuller-240space-02, March 2008.
[I-D.wilson-class-e]	Wilson, P, Michaelson, G and G Huston, " <a href="#">Redesignation of 240/4 from "Future Use" to "Private Use"</a> ", Internet-Draft draft-wilson-class-e-02, September 2008.
[I-D.ietf-v6ops-6to4-to-historic]	Troan, O, " <a href="#">Request to move Connection of IPv6 Domains via IPv4 Clouds (6to4) to Historic status</a> ", Internet-Draft draft-ietf-v6ops-6to4-to-historic-05, June 2011.
[RFC6343]	Carpenter, B., " <a href="#">Advisory Guidelines for 6to4 Deployment</a> ", RFC 6343, August 2011.
[I-D.arkko-ipv6-only-experience]	Arkko, J and A Keranen, " <a href="#">Experiences from an IPv6-Only Network</a> ", Internet-Draft draft-arkko-ipv6-only-experience-03, April 2011.
[ARIN-prop-127]	Donley, C., "ARIN-prop-127: Shared Transition Space for IPv4 Address Extension", Jan 2011.
[ARIN27.2011-5]	ARIN, "ARIN XXVII Meeting - Participant Vote on 2011-5", Apr 2011.
[ARIN-2011-5]	ARIN, "Draft Policy ARIN-2011-5: Shared Transition Space for IPv4 Address Extension", 2011.
[ARIN-2011-5-AC]	

	ARIN, "Message to ARIN-PPML, announcing selection of ARIN-prop-127 for Discussion as Draft Policy 2011-5", Feb 2011.
[ARIN-2011-5-Staff]	ARIN, "Message to ARIN-PPML, providing additional ARIN Staff Assessment of Draft Policy 2011-5", Feb 2011.
[ARIN-2011-5-LC]	ARIN, "Message to ARIN-PPML, announcing Last Call for Draft Policy 2011-5", Apr 2011.
[ARIN-2011-5-Rec]	ARIN, "Message to ARIN-PPML, announcing Advisory Council meeting results Recommending 2011-5 for Board Approval", May 2011.
[ARIN-AC-28Jan2011]	ARIN, "Minutes: Meeting of the ARIN Advisory Committee - 28 Jan 2011", Jan 2011.
[ARIN-AC-13Apr2011]	ARIN, "Minutes: Meeting of the ARIN Advisory Committee - 13 Apr 2011", Apr 2011.
[ARIN-AC-19May2011]	ARIN, "Minutes: Meeting of the ARIN Advisory Committee - 19 May 2011", May 2011.
[ARIN-NRPM-8.3]	ARIN, "ARIN Number Resource Policy Manual, section 8.3 - Transfers to Specified Recipients", Jul 2011.
[IAB-response]	IAB, "IAB responds to ARIN request for guidance regarding Draft Policy ARIN-2011-5", Jun 2011.
[NRO-IANA-exhaust]	NRO, "Free Pool of IPv4 Address Space Depleted", Feb 2011.
[v6ops-msg06187]	WIDE, "Re: [v6ops] IETF 79 Meeting minutes - Draft", Nov 2010.
[GIH-When]	Huston, G, "When?", Sep 2010.
[APNIC-final-slash8]	APNIC, "APNIC IPv4 Address Pool Reaches Final /8", Apr 2011.
[I-D.hain-1918bis]	Hain, T, " <a href="#">Expanded Address Allocation for Private Internets</a> ", Internet-Draft draft-hain-1918bis-01, January 2005.
[PPML-022778]	Message to ARIN-PPML, indicating the Board's disposition toward 2011-5", July 2011.
[CISCO]	, , "TCP/IP Overview: Class E Addresses", .

### Authors' Addresses

Stan Barber Barber Cox Communications EMail: [stan.barber2@cox.com](mailto:stan.barber2@cox.com)

Owen Delong Delong Hurricane Electric EMail: [owen@delong.com](mailto:owen@delong.com)

Chris Grundemann Grundemann CableLabs EMail:

[c.grundemann@cablelabs.com](mailto:c.grundemann@cablelabs.com)

Victor Kuarsingh Kuarsingh Rogers Communications EMail:

[victor.kuarsingh@rci.rogers.com](mailto:victor.kuarsingh@rci.rogers.com)

Benson Schliesser Schliesser Cisco Systems EMail: [bschlies@cisco.com](mailto:bschlies@cisco.com)