



Implementing the
Global Policy for
Post Exhaustion
IPv4 Allocation
Mechanisms by the
IANA

Table of Contents

Initial consultation	1
Registry users	1
Registry features.....	1
Phase one: Create the “Recovered IPv4 Pool”	2
Phase two: Allocating from the “Recovered IPv4 Pool”	4
Input sought	4

Initial consultation

On 6 May 2012, the [Global Policy for Post Exhaustion IPv4 Allocation Mechanisms by the IANA \(Global Policy\)](#) was ratified by the ICANN Board. The IESG and NRO are key policy-making bodies for IPv4 allocations and IP address registries. In June 2012, ICANN consulted them about a plan to implement the Global Policy.

The policy requires ICANN to “make appropriate modifications to the ‘Internet Protocol V4 Address Space’ page of the IANA website.” Feedback given to ICANN during the initial consultation included a request to see the complete Recovered IPv4 Pool in a single view, separate from the main IPv4 Registry. This feedback leads to the question of what “appropriate modifications” means in the context of this policy and what is desirable.

Registry users

There are at least four distinct groups of registry users. It is important that changes to the registry either represent an improvement or do not affect them. The registry group users are:

- Registry operators (RIRs, NIRs & others);
- Network operators;
- Other technical registry users, such as software developers and Computer Security Incident Response Teams (CSIRTs); and
- Ordinary Internet users, such as people who have received spam e-mail.

Each of these groups has slightly different needs and has a different approach to the way they use the data published in the registry. For instance, some want a registry that is simple to parse programmatically; while others want to search for a particular address; and yet others are looking for the high level IPv4 address space distribution.

Registry features

Some of the key options that should be considered when implementing the “appropriate modifications” to the IPv4 Registry are listed below.

- Allow users to sort and filter using the values for each of the key data elements. For instance, showing all allocations made in a specific year or all prefixes with a particular status. The ability to sort is already present in the current registry stylesheet but the ability to filter the registry is not.
- Make allocations to legacy registrants with /8s from the former Class A space and RIRs available as a single file in a structured format that can be programmatically parsed. This is what is currently provided but if the Recovered IPv4 Pool is published as a separate registry some kind of composite view would also be needed to provide this functionality.
- Allow users to easily search for the registration containing the address they want information about. This is a particularly important feature for users who are querying the registry but do not have an understanding of what an IP address is and are nor familiar with IP address ranges expressed as prefixes.

- Allow users to see graphical representations of particular elements in the registry. For instance, charts showing the proportion of address space with each status; the proportion of address space allocated to each RIR; or a map of the whole address space.

Phase one: Create the “Recovered IPv4 Pool”

The initial consultation considered two implementations for the Recovered IPv4 Pool: 1) through increasing the granularity of the IPv4 Registry and 2) in a separate registry, with the current IPv4 Registry remaining unchanged. There are advantages and disadvantages to both approaches.

	Pro	Con
Single registry	Gives a complete overview of the IPv4 unicast address space	Could grow to several thousand entries
	Provides an exact implementation of the policy	
	Can be implemented more quickly	
	Makes it easier for ordinary Internet users to identify special-use assignments made in RFCs	
Multiple registries	Provides an easy-to-read snapshot of the current Recovered IPv4 Pool	Provides an incomplete overview of the IPv4 unicast address space
		Is incompatible with the the policy requirement for the Recovered IPv4 Pool to be shown as part of the "Internet Protocol V4 Address Space" page on the IANA web site

One way to mitigate the disadvantages arising from a multiple-registries approach is to have multiple ‘views’ of an enhanced IPv4 unicast registry. However, this approach would require some software development work and so might not be possible in the immediate future, necessitating a short-term multi-registry model as an interim solution.

If there is a consensus in favour of a multi-registry model, the design of the Recovered IPv4 Pool registry needs to be considered. One possible approach is to use a composite registry which shows the status of prefixes using a status field. Another possible approach is to use two sub-registries: one for the current Recovered IPv4 Pool and the other as a record of allocations made from the Recovered IPv4 Pool. There are advantages and disadvantages to both approaches.

	Pro	Con
Single Recovered IPv4 Pool registry	Requires less software to generate a synthesized	More complicated for occasional registry users to

	complete overview of the IPv4 unicast address space	parse and understand
Two Recovered IPv4 Pool sub-registries	Creating a new registry structure means not requiring dates for all entries, making record keeping simpler	

The example below shows what an integrated Recovered IPv4 Pool registry might look like. The design could be varied to provide a date for entries with a RECOVERED status. However, registry users who are not aware of the policy might not understand that 172.16.0.0/16 is allocated to AFRINIC in this example, while 192.0.2.0/24 was returned by AFRINIC. Experience gained while there were unallocated /8s suggests that this is likely to lead to misdirected abuse reports when IPv4 address space in the Recovered IPv4 Pool is used illicitly.

Prefix	RIR	Date	Whois Server	Status
172.16.0.0/16	AFRINIC	2013-09-01	whois.afrinic.net	ALLOCATED
172.17.0/16	APNIC	2013-09-01	whois.apnic.net	ALLOCATED
172.18.0.0/16	ARIN	2013-09-01	whois.arin.net	ALLOCATED
172.19.0.0/16	LACNIC	2013-09-01	whois.lacnic.net	ALLOCATED
172.20.0.0/16	RIPE NCC	2013-09-01	whois.ripe.net	ALLOCATED
192.0.2.0/24	AFRINIC			RECOVERED
198.51.100.0/24	APNIC			RECOVERED
203.0.113.0/24	LACNIC			RECOVERED

The examples below show what a registry composed of two sub-registries might look like. It makes it clearer that the donors to the Recovered IPv4 Pool no longer have control of the address space they have returned.

(Example) Current Recovered IPv4 Pool

Prefix	Returning RIR	Status
192.0.2.0/24	AFRINIC	RECOVERED
198.51.100.0/24	APNIC	RECOVERED
203.0.113.0/24	LACNIC	RECOVERED

(Example) Allocations made from the Recovered IPv4 Pool

Prefix	RIR	Date	Whois Server	Status
172.16.0.0/16	AFRINIC	2013-09-01	whois.afrinic.net	ALLOCATED
172.17.0/16	APNIC	2013-09-01	whois.apnic.net	ALLOCATED
172.18.0.0/16	ARIN	2013-09-01	whois.arin.net	ALLOCATED
172.19.0.0/16	LACNIC	2013-09-01	whois.lacnic.net	ALLOCATED
172.20.0.0/16	RIPE NCC	2013-09-01	whois.ripe.net	ALLOCATED

Additional discussion is required with all key stakeholders before finalizing an approach to the registry changes.

Phase two: Allocating from the “Recovered IPv4 Pool”

During the initial consultation, a preference for allocations from the Recovered IPv4 Pool to take account of two criteria in addition to those specified in the policy, was expressed. These are:

- Existing reverse delegation boundaries; and
- Aggregation maximization.

There are at least two possible approaches to the implementation of this part of the policy. The first is to develop and publish software to make selections from the Recovered IPv4 Pool. The alternative approach is similar to the “perform actions” and “confirm actions” steps in the current [Internet-Draft Approval Process, as described in ICANN’s Proposal to Perform IANA Functions](#). Staff would “calculate the size of the ‘IPv4 allocation unit’ [...] at the beginning of each IPv4 allocation period” and manually calculate the distribution from the pool based on the policy requirements and the RIRs’ preferred criteria. The proposed distribution would then be published, either to the RIRs alone or more widely, for a review period. Changes could be made based on feedback, with the actual allocations taking place a month after the start of each allocation period.

	Pro	Con
Software selection process	Reduces staff operational workload	Complicated to design, code and review
	High degree of transparency as anyone can use published software to get identical results	Changes to the rule-set need software update following the initial implementation & release
Manual selection process	No software development costs	Susceptible to human error
		Increases staff operational workload
		Degree of transparency depends on who is invited to review proposed distribution
	Allows stakeholders to influence prefix selection in each IPv4 allocation period, potentially increasing customer satisfaction	Allows stakeholders to influence prefix selection in each IPv4 allocation period, potentially reducing transparency
	Allows 3 rd -party review prior to allocation, providing transparency	Slows the allocation process in each IPv4 allocation period

The second approach requires less up-front investment in software development and should provide customer satisfaction with all allocations at the cost of a manual process.

Input sought

Input is sought on implementation options.

1. Which approach should be taken to the registration of the Recovered IPv4 Pool?

2. Which approach should be taken to selecting the prefixes to be allocated from the Recovered IPv4 Pool?

This public comment period is intended to gather input from registry users.