Regional Address Registries, Governance and Internet Freedom

26 November, 2008

Drafters:

Milton Mueller

Professor, Syracuse University School of Information Studies

XS4All Professor, Delft University of Technology

Brenden Kuerbis

Ph.D. candidate, Syracuse University School of Information Studies

Operations Director, Internet Governance Project

Concurring:

Michel van Eeten

Assoc. Professor, Delft University of Technology, School of Technology, Policy and Management

Abstract: *Regional Internet Address Registries (RIRs) are private, nonprofit and transnational governance entities that evolved organically with the growth of the Internet to manage and coordinate Internet Protocol addresses. The RIR's management of Internet address resources is becoming more contentious and more central to global debates over Internet governance. This is happening because of two transformational problems: 1) the depletion of the IPv4 address space; and 2) the attempt to introduce more security into the Internet routing system. We call these problems "transformational" because they raise the stakes of the RIR's policy decisions, make RIR processes more formal and institutionalized, and have the potential to create new, more centralized control mechanisms over Internet service providers and users. A danger in this transition is that the higher stakes and centralized control mechanisms become magnets for political contention, just as ICANN's control of the DNS root did. In order to avoid a repeat of the problems of ICANN, we need to think carefully about the relationship between RIRs, governments, and Internet freedom. In particular, we need to shield RIRs from interference by national governments, and strengthen and institutionalize their status as neutral technical coordinators with limited influence over other areas of Internet governance.*

From 1995 to the present, the domain name system was the center of controversy and institutional change for global Internet governance. As this occurred, the major Regional Address Registries, ARIN, RIPE and APNIC, quietly evolved into highly technical, self-governing spaces that were off most people's radar screen. As some academic observers noted at the time, the contrast between the politically charged ICANN space and the relatively obscure and quiet address governance space was puzzling. Address resources have always been as economically valuable as domain names, if not more so. Technically speaking, IP addresses are far more essential to the functioning of the Internet than domain names. Why did one resource management regime become the center of a global governance controversy while the other did not? If we can understand the answer to this question, we should also be able to explain what is now making IP addresses almost as controversial as domain names. And if we are really lucky, by comparing our prior experience with ICANN and domain names to the new predicament of the RIRs, we might be able to learn some useful lessons and anticipate and avoid some known problems.

Two transformational problems are converging to make the management of

Internet address resources more contentious, and the RIRs more prominent in the debates over global internet governance.

The most important is the impending exhaustion of the IPv4 address space. The RIRs were created to manage an Internet identifier resource, namely IPv4 addresses. Some time in the last two years, it became evident that we will run out of IPv4 addresses soon. Near the end of 2008, the last remaining stash of unallocated address blocks – the so-called "free pool" – had dwindled to only 36 blocks. In recent times, these blocks of addresses have been distributed to regional address management entities at a rate of about 12 per year, which means we have only about a three year supply left. Demand may accelerate as exhaustion of this pool approaches, but even if it doesn't the end of the free pool is within sight.[1] The depletion of that resource pool, as we shall see, is a seismic shift of the very ground on which the RIRs stand, and will produce major transformations of their practices.

Secondly, the growing demand for a more secure Internet is another critical arena of change. RIRs are affected by efforts to make the Internet more secure, especially in the area of routing. There has always been an intimate interaction between address management and routing. Up to now, however, the RIRs' function of registering and rationing address resource utilization has been only loosely related to the routing practices of Internet service providers. This could change. Some of the secure routing proposals might give the RIRs the power to exert direct, operational control over Internet service providers.

We call these problems "transformational" because they force the RIRs to become more formally institutionalized, make their role in internet governance more contentious and political, and raise the stakes of their policy decisions. As we shall see, these impending changes in the RIRs institutional capacity have both negative and

---

[1] A 2005 study by Cisco's Tony Hain projected depletion of the IANA pool by 2010. Tony Hain, "A Pragmatic Report on IPv4 Address Space Consumption," *The Internet Protocol Journal*, Volume 8, Number 3. http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_8-3/ipv4.html Another respected computer scientist who, in 2003, projected that the IPv4 address space would last "another three decades or so" now projects the depletion date of the IANA pool at January 2011, and the depletion of the regional subpools at November 2011. See Geoff Huston, "IPv4 Address Report" http://www.potaroo.net/tools/ipv4/ .

positive possibilities. RIRs could become a point of more centralized control over suppliers and users, and that could attract all the political problems that we now associate with ICANN. But these changes might also make address management more efficient and help fix some of well-known security problems. This means that we will need to think more broadly and comprehensively about the role of RIRs in Internet governance than we have before. It is no longer sufficient, or even helpful, for the RIRs simply to proclaim that they make policy in an open and bottom up manner. Open participation by itself offers no solution to the problems identified in this paper, and could actually make things worse. Instead, we need to think more deeply about what RIRs should and *should not* do, regardless of who participates in them. We need to carefully assess how their activities relate to governmental power, and about what rights individual Internet users and Internet service providers have within their governance regime.

This paper begins with an analysis of what the RIR actually do. It then examines the two transformational problems facing the RIRs. The concluding section focuses on how the RIRs might respond to their impending politicization. It discusses the parallel with the ICANN experience and discusses the need for public policy principles applicable to the RIRs that can help safeguard the freedom of the Internet.

The IP Address Governance Regime

The Internet we know is based upon a data communication protocol known as Internet Protocol version 4 (IPv4). IPv4 is a software procedure that moves data packets from one unique numerical address to the other.[2] The 1981 standard that defined Internet Protocol created a fixed address field of 32 bits, which creates a mathematical possibility of about 4 billion unique addresses. Address blocks are assigned to private users in a hierarchical fashion. (See Figure 1) At the top of the hierarchy is ICANN, whose IANA function distributes large blocks of 16,777,216 addresses (known as /8's) to one of five regional Internet address registries (RIRs).

---

[2] RFC 791, Internet Protocol, DARPA INTERNET PROGRAM PROTOCOL SPECIFICATION, September 1981. See also RFC 790, Assigned Numbers, Jon Postel, September 1981.

The RIRs then accept applications from organizations with networks that need addresses within their territory. Some larger blocks may be assigned directly to end user organizations, but most will go to Internet service providers who will then re-assign them to their customers.[3]

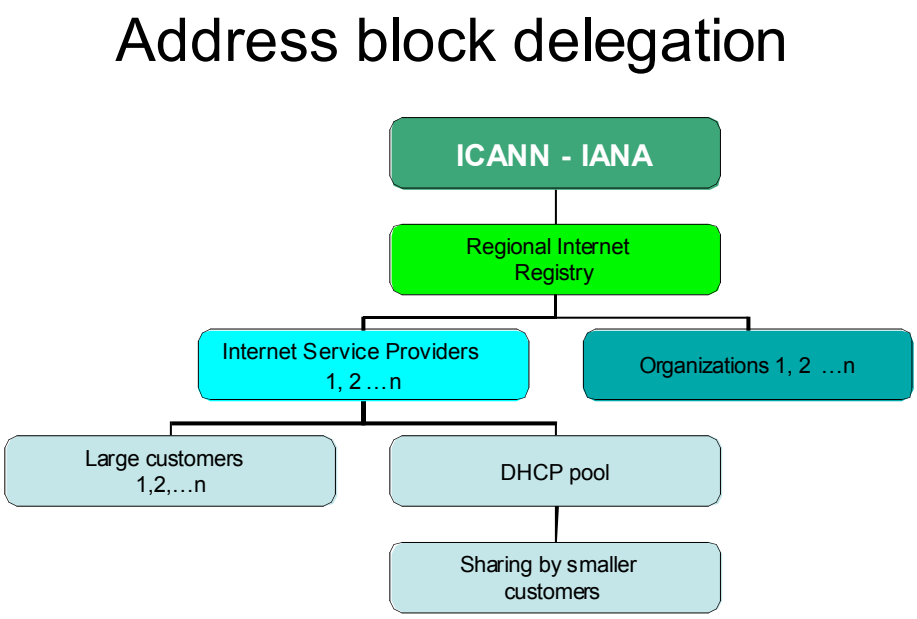

Like all social institutions, the RIRs have developed an ideology about what they do and why they do it, which is codified in their own policy documents. Address resources are considered a "shared public resource" and the RIRs are considered their "stewards." Addresses are said to be "loaned" to private users, not sold, and users are not supposed to gain any property rights in an address block they are granted. One can understand the political and economic pressures the RIRs face better, however, if one uses the tools of institutional economics to examine the historical evolution of the

[3] An *allocation* is defined as an address block given to an Internet service provider for intermediate use in selling internet service to other users. An address *assignment* is a block given to end users for their own use; e.g., corporations or universities with private networks.

institutional arrangements, the way they define rights to address resources, and the gaps and shortcomings of the regime.[4]

The initial distribution of IPv4 addresses began very informally in the early 1980s, when the Internet was basically a research project. A computer scientist at the University of Southern California handed out big swaths of the available space to members of the American military-industrial-university complex: MIT, Stanford, Hewlett Packard, Motorola, General Electric, Halliburton, Defense Department agencies, etc. These are called "legacy allocations." The creation of the RIRs led to tighter and more formalized address allocation and assignment policies, and more careful registration and tracking policies. Before they were created, however, a large portion of the IPv4 address space – one estimate says one third (OECD, 2008), another says 45%[5] -- had already been assigned or allocated. A significant amount of these legacy allocations are still held without any contractual obligations.

RIRs were first created in the 1990s as the Internet protocols began to be widely adopted.[6] The first Regional Internet Registry, RIPE-NCC, was created in 1991 to serve the European region. In 1995 APNIC was created to serve the Asia Pacific region. Both were incorporated as private sector nonprofits. In 1997, parallel to the creation of ICANN, the address administration functions performed by several U.S. government contractors were privatized and placed in the hands of a new nonprofit entity known as the American Registry for Internet Numbers (ARIN). (Later, two new regional registries were created for the Latin America and African regions: LACNIC and AfriNIC, respectively.)

From a resource management standpoint, the RIRs perform three functions: 1) they serve as a *registry* that coordinates the uniqueness of IP addresses; 2) they act as

---

[4] For a sampling of the institutional economics literature, see Elinor Ostrom, 1990. *Governing the Commons: The evolution of institutions for collective action* Oxford University Press; Elinor Ostrom, Roy Gardner, James Walker, 1994. *Rules, Games, and Common-pool Resources*. University of Michigan Press; Anderson, Terry L., Grewell, J. Bishop (2000) "Property Rights Solutions for the Global Commons: Bottom-Up or Top-Down?" In: Duke Environmental Law & Policy Forum, Vol. X, No. 2, Spring 2000

[5] Author interview with Paul Wilson, APNIC Director, July 21, 2004.

[6] See RFC 1174 (1990) and RFC 1466 (1993) for the earliest documentation of the rationale for creating Regional Internet Registries.

gatekeepers that *conserve* address resources; and 3) they distribute addresses in a way that *aggregates routes*, thereby conserving scarce routing table space.

The registry function

The most basic function of the RIRs is to maintain a registry that keeps track of which organization has been allocated or assigned which IP address block(s). Maintaining a registry allows service providers and others to see who has been assigned a specific address block, and so to coordinate their selections so that there are no conflicts or overlaps. Coordinating the uniqueness of IP addresses is an essential technical function; two different machines on the Internet attempt cannot use the same address at the same time without causing conflicts or malfunctions that can disable Internet connectivity. Registry information about address assignments can be retrieved from the RIRs' "Whois databases." This information is also useful for Internet service providers in determining their routing policies.

The conservation function

The RIRs also fulfill a conservation function by rationing access to address resources. Internet service providers and other organizations that want addresses must fill out a detailed request form and submit it to the RIR in their region. The form is an attempt to "justify" their "need" for the addresses. The RIR bases its needs assessment on engineering studies of the applicant's plans, among other things. Thus, the definition of "need" underlying this regime is entirely technical; it does not take into account the relative economic value of alternative uses of addresses, nor does it reward technical configurations that conserve address space. In effect, the RIRs have instituted a system of central planning in which a bureaucracy rations access to address resources based on pre-formulated rules and criteria.[7] In the course of

[7] See the general APNIC FAQ, http://www.apnic.net/info/faq/apnic_faq/obtaining.html; the *ARIN Number Resources Policy Manual* http://www.arin.net/policy/nrpm.html; and the RIPE *IPv4 Address Assignment and Allocation Policies for the RIPE NCC Service Region* http://www.ripe.net/ripe/docs/ipv4-policies.html

receiving address resources, the organizations and Internet service providers sign contracts with the RIRs, and these contracts are used to enforce the applicable policies.

The route aggregation function

One of the most important policy guidelines of the RIR regime is the need to aggregate routes at the top of the routing hierarchy. Hierarchical route aggregation is a response to a scaling problem caused by the growth of the Internet. Whenever possible, Internet service providers or organizations applying for addresses are given large blocks of contiguous addresses. By preventing end users or assignees from breaking up these contiguous blocks into smaller units and fragmenting them across different Internet service providers, the RIRs economize on the number of routes announced in the core of the Internet. If there were no aggregation of routes, the number of routes announced, it is feared, would expand beyond the capacity of routers to handle them. Thus in the mid-1990s the Internet community rejected what it called the "address ownership" model and adopted instead what it called an "address lending" policy. RFC 2008 was the seminal document making the case for a lending policy as only the model consistent with the need for hierarchical route aggregation.[8] Although RIRs finance themselves via address-related fees and membership charges, they insist that members are not "buying" addresses but are merely paying the RIR for services associated with administering the address space and its registry. The RIRs formally prohibit assignees from reselling or transferring the addresses directly to other private users.[9]

The line between permitted and not-permitted transfers is not always clear, however.[10] The prohibition on decentralized transfers is justified by the RIRs on two

---

[8] RFC 2008, "Implications of Various Address Allocation Policies for Internet Routing" Y. Rekhter, T. Li. (1996).

[9] Route aggregation serves an important technical efficiency function but it also helps lock customers in to their Internet service providers. Internet service providers, not their customers, control the addresses, so when one switches from one vendor to another an organization may have to completely renumber its network, which is costly.

[10] Internet service providers who hold address allocations sell services commercially to their customers, and among these services are fixed IP addresses, with specific charges associated with addresses. When companies with IP address allocations or assignments are merged or acquired, RIRs allow the address resources to be transferred along with ownership of the company.

grounds. First, it is perceived as necessary to maintain their own role as central planners, or "need assessors." If organizations could obtain addresses on an open market then willingness to pay, rather than a technical justification, would determine how addresses were allocated. Second, some control of how blocks are allocated is perceived as needed to maintain route aggregation.

From Loose Governance to Centralized Governance?

Overall, the RIRs have evolved into an effective but fairly loose self-governance regime. Within this relatively informal governance regime, there is a great deal of slack. When considering the economic and political forces that are changing this regime, it is important to understand: where do they get their authority or influence, and when they are or are not able to enforce compliance; what is not controlled or imperfectly controlled by their methods; and how economic incentives put pressure on network operators to break or bend the regime's rules.

The RIRs' primary source of governance leverage is their maintenance of an authoritative IP address registry, with an associated Whois service. As long as the world's Internet service providers collectively recognize RIR's registries as authoritative, an entry in the registry constitutes a legitimate and globally recognized claim to the exclusive use of IP address resources. The public Whois service allows Internet service providers (Internet service providers) to look up the identity of organizations that hold address blocks. Internet service providers will usually not issue route announcements for address blocks unless they first check the IP address Whois to see that the organization holding the block is actually who it claims to be and is their customer, vendor or peering partner. Thus while it is technically possible for any organization to appropriate any IP address block they please, appropriations of IP addresses that bypass the RIRs are more likely to encounter trouble in their attempt to be recognized and used by Internet service providers. In short, the RIRs derive most of their legitimacy and impact from the fact that Internet service providers pay attention to them and choose to use them as a basis for coordination. Bear in mind that Internet service providers and similar organizations with IP networks constitute the

membership of RIRs, and thus play a major role in determining their officers and policies.

The same could be said, however, of the DNS root managed by ICANN and the U.S. Commerce Department. Theoretically, it is possible for the world's Internet service providers to abandon the ICANN root and point to another. And yet it is obvious that strong network effects have made the world completely dependent on the DNS managed by ICANN and the U.S. Commerce Department, and any attempts to defect would be both costly and likely to cause global compatibility problems.[11] Thus, over the past decade the seemingly distributed and voluntary system of domain name governance has become more centralized, more regulatory and increasingly political.

In the analysis that follows, we show how the same could happen to the RIRs. We examine two forces that are creating pressures to change the informality of this governance regime: address scarcity and the demand for security.

Transformational Problem #1: IPv4 Address Scarcity

The most important change in the RIR environment is the depletion of the IPv4 address space. Running out of IPv4 addresses poses serious challenges for the RIR's governance model. Unfortunately, many people within that regime are still in various stages of denial about this.

The RIR's approach to resource management is based on two key assumptions: 1) that there are free, unallocated resources available; and 2) that the task of the resource manager is to distribute unused addresses to organizations that actually want to use them in functioning networks (as opposed to hoarders, speculators or hogs). But once the IPv4 free pool runs out, the established methods of allocating address resources according to "justified need" loses its relevance. When all the IPv4 addresses have been allocated to users, the concept of "need" becomes *relative need* rather than some engineering-based calculation of absolute need. In other words, it is possible to imagine five, ten or a hundred organizations that have a meritorious technical claim on some IPv4 addresses, but there may not be enough available

[11] Mueller 2002

addresses to satisfy all their requests. In the post-free pool world, address management does not mean judging whether an engineering plan or utilization levels justify a specific number of address blocks. In means, instead, the following things:

- Making decisions about which of two, equally justified *competing* applications should get available addresses;
- *Transferring* resources from less important, lower-valued uses to more important, higher-valued uses;
- *Reclaiming* address resources that have been allocated but remain unused;
- *Trimming back* overly large allocations to organizations that could make do with fewer addresses.

The problem here is that the existing procedures of the RIRs cannot perform any of these functions well. To respond to these new demands, the RIRs must undertake a massive transformation of their policies and practices.

Reclamation of unused address space, for example, is one of the biggest weaknesses of the existing regime. A large part of the allocated IPv4 address space seems to be unused, especially legacy allocations in the North American region.[12] Address blocks that are lying fallow can be surreptitiously taken over by spammers, illegal pornographers, or other Internet malefactors with a need to operate under cover. Spammers hijacked an entire /8 originally allocated to Halliburton in the 1980s.[13] Two /16 address blocks, containing tens of thousands of IPv4 addresses, were hijacked from NASA and a small software company and used to facilitate spamming.[14] In these two cases, the address blocks were essentially abandoned; their delegated users had completely lost track of their status and were not even aware of their appropriation by a third party. Routes for these misappropriated addresses can be announced even if they are not officially sanctioned by the RIRs. For example, the CIDR report for

---

[12] An OECD report cited a claim by Geoff Huston, Chief Scientist at APNIC, that 90% of RIR-allocated space is actually used (routed) while only 40% of legacy space is used. OECD (2008) p. 26-27. The report cites surveys that examine the population of visible IPv4 Internet hosts, and find that "only a low percentage of advertised addresses respond, which could mean that even among routed address space, significant address space is unused."

[13] Google cache of The Complete Whois web site, retrieved 10 June 2008, http://completewhois.org/hijacked/hijackers.htm

[14] http://www.47-usc-230c2.org/ 47-usc-230c2, A web site maintained by Ronald F. Guilmette.

March 2008 lists 317 potentially bogus Autonomous System numbers and nearly 400 possibly bogus route announcements.

The appropriation of address blocks by spammers illustrates the weakness of the current RIR policies in dealing with unused or underutilized address space. When IP addresses are not used by those to whom they have been allocated, they do not automatically return into the common pool for use by others.[15] Organizations that have been given IP address allocations retain them until they choose to give them up, regardless of whether they are actually being used. And they have very weak incentives to return addresses to RIRs.[16] If they don't give them back, nothing bad happens. If they do give them back, they incur both administrative costs (the cost of altering their records and interacting with the RIR) and opportunity costs (the cost of foregoing future use of the addresses). RIRs' ability to monitor the actual usage of assignments is limited. Even if they did have perfect information about actual usage and "needs" of applicants, their enforcement powers are weak.

The problem of address block reclamation is difficult enough when the RIRs are dealing with organizations with which they have established contractual agreements. But what about all the holders of legacy address blocks, which were allocated before the RIRs existed? Because they have no contracts with legacy holders, RIRs lack the authority to recover their resources or regulate their behavior, unless the legacy holders choose to give it to them. How big is this problem? Consider that one estimate says that 45% of the IPv4 address space was allocated prior to 1997. In the ARIN region, the Whois records associated with 2,357 Autonomous Systems and 22,718 Organizations have not changed since the end of 1997, when ARIN was created. This is a good indication of the large number of entities that are formally outside the governance regime in North America, even though they may cooperate

---

[15] In a true common pool model, the IP address space would work like a gigantic DHCP address pool. (DHCP stands for Dynamic Host Configuration Protocol, and defines mechanisms through which clients can be assigned a network address for a finite lease, allowing for serial reassignment of network addresses to different clients when one client ceases using an address. See RFC 2131.) Organizations would grab addresses (like catching fish) only when they were actually using them, and as soon as they were not using them the addresses would be released back into the common pool for use by others.

[16] In principle APNIC allocates addresses on an "annually renewable" basis. In practice it seems to rely primarily on the organization's initiative to reclaim addresses.

with or even participate in it. Since 2007, RIRs have made concerted efforts to bring the legacy holders into the contractual governance fold by developing a "Legacy Registry Service Agreement."[17] Take up by legacy holders has been modest, at best.[18]

If reclamation will be difficult, what about transferring address resources from lower to higher-valued uses? Here again, the RIRs also do not appear to be well-positioned to facilitate transfers. There are basically two ways for an agency to make judgments about the relative value of resources in alternate uses. One is to institute trading and competitive bidding for them; i.e., to institute a market allocation system.[19] Transfer markets would allow organizations with address resources they think they can do without, to sell address blocks to an organization that needs/wants them.[20] RIPE, APNIC and ARIN have all proposed and debated transfer policies. The proposals are extremely controversial within the RIRs because they undermine the RIR's centralized control of the resources and due to anti-market sentiment. The alternative is to hold ongoing beauty contests based on expert judgment and/or community input; something akin to what was called "comparative hearings" in broadcast licensing, where administrative processes are used to assess the merit of competing claimants. These procedures are time and labor intensive and would not be welcomed by many in the industry.

A robust policy debate about which of these options is preferable is now going on inside the RIRs.[21] This paper, however, is concerned not with the normative debate

---

[17] ARIN's new Legacy RSA contracts, http://www.arin.net/registration/legacy/index.html, bring uncontracted legacy address resource holders into the contractual system but promise not to reclaim unused resources.

[18] 129 Organizations have signed the ARIN Legacy RSA as of August 2008. Of those, 66 of them already received new address resources and thus had signed a separate standard RSA with ARIN. So only about 63 new organizations were brought in by the program, out of tens of thousands.

[19] Each of the three largest RIRs is considering proposals to permit market-based address transfers. They are: Asia-Pacific region: prop-050-v002: IPv4 address transfers (Huston); European region: RIPE 2007-08, "Enabling Methods for Reallocation of IPv4 Resources." (Titley and van Mook); North America region: ARIN: Policy Proposal 2008-2 IPv4 Transfer Policy Proposal.

[20] Milton Mueller, Scarcity in IP addresses: IPv4 Address Transfer Markets and the Regional Internet Address Registries (July 20, 2008). *Internet Governance Project*. Paper IGP08-002. Available at http://internetgovernance.org/pdf/IPAddress_TransferMarkets.pdf

[21] For an assessment of the transfer proposals, see Mueller, M. "Scarcity in IP addresses: IPv4 Address Transfer Markets and the Regional Internet Address Registries," (July 20, 2008). *Internet Governance*

over markets vs. no markets, but with a scientific assessment of the impact of IPv4 scarcity on RIRs as governance institutions. From that perspective, it doesn't really matter which of the two directions they choose. Either way, *major changes will be required in the RIRs' processes*. RIRs will no longer be rationing access to a free pool of unoccupied addresses; they will be facilitating transfers among users and uses and reclaiming unused or underutilized address blocks. Both available methods of doing this (markets or administrative procedures) point in the same general direction: toward much more extensive record keeping, surveillance and enforcement capabilities. In short, an environment of resource scarcity will lead to much greater formalization and institutionalization of the RIRs' administration of the address space, and a reconfiguration of their relationship to Internet service providers and other consumers of address resources.

The market path

One proposal for dealing with the problem of IPv4 address scarcity is a more liberal transfer market. As others have remarked, a market transfer policy requires RIRs to act as effective and authoritative "title agencies" for address resources. Before addresses can be exchanged between two parties, the buying party must feel confident that the selling party is the legitimate owner of the resource being exchanged, and that the resources it acquires will be recognized as its own on a global basis. Thus, the Whois record showing that an organization has been assigned a specific block of addresses will have much greater economic consequences once those addresses become, in effect, transferable use rights. The boundaries, limits and regulations associated with those rights will have to be more precisely defined. In authorizing and recording address transactions, the RIRs will have to be vigilant against fraud and misappropriation of address resources, otherwise its attempt to facilitate transactions will enmesh them in litigation and conflict when the transferred property rights are found to be defective or contested. The RIRs will have to monitor those who engage in these transactions carefully to make sure that they conform to regulations designed to

---

Project. Paper IGP08-002. Available at http://internetgovernance.org/pdf/IPAddress_TransferMarkets.pdf

prevent speculation and hoarding. Transfer may require instituting "Resource Public Key Infrastructure" (RPKI) certificates to automatically authenticate resource claims.[22] If transfers become common, the RIRs will also have to develop and enforce a whole new set of regulations designed to maintain route aggregation under the new market transfer regime.

The nonmarket path

Assume that the RIRs forsake market-oriented transfer policies and instead rely on more systematic reclamation and administrative procedures to respond to scarcity in IPv4. If this path is taken, the RIRs would have to make *even more sweeping changes* in their practices. First, they would have to invest in the institutional capacity to promulgate new standards regarding what would be considered "unused" or "insufficiently used" address resources, and then (somehow) acquire the legal authority to take those resources away from their nominal holders so that they could be transferred to other applicants. Second, they would have to apply those standards to all organizations in their region and identify and reassign the identified resources. Third, in re-assigning reclaimed resources, RIRs could no longer rely on their established engineering criteria; instead they would have to set up a competing claims process. In an environment with no free pool, multiple organizations will be able to present the RIRs with viable engineering plans for using the addresses. Engineering plans that technically justify the use of a certain number of addresses do not necessarily justify taking address allocations away from someone else, nor do they unambiguously tell you which of a dozen competing proposals is more worthy. So an RIR would have to decide which plan was more important or more valuable. Once RIRs start taking away resources from current holders, and making decisions among competing claimants about which plans are more worthy, they will need to be very formal and careful in their procedures, and hire lawyers, because mediating any such distributional conflict is bound to be contentious. These distributional conflicts in turn

---

[22] See the presentation of Randy Bush, "IPv4 runout, Trading, and the RPKI," presentation at the Internet Institute of Japan, April 15, 2008. http://www.menog.net/meetings/menog3/presentations/bush-080415.menog-v4-trad-rpki.pdf

will produce additional procedural formality and complexity in RIR processes, to protect both the RIR and the parties seeking resources. Fourth and finally, there is a danger that the absence of a legitimate private transfer market will fuel a black or gray market, which could lead to a breakdown in the accuracy and universality of the RIRs' databases. If in an environment of scarcity the RIRs institute procedures for acquiring IPv4 addresses that are perceived as too costly or time-consuming, organizations may simply bypass them and begin trading address resources among themselves, leaving the RIRs out of the loop. Although trading address resources is supposed to be forbidden, by creating shell companies, or retaining organizational IDs associated with merged companies, it is possible to skirt these rules.

So the RIR's will have to invest significant resources in surveillance of address utilization and find stronger enforcement mechanisms to prevent transfers and appropriation outside the regime's rules. Thus, even if address resources are not considered transferable property rights, in an environment of scarcity the RIRs will have to function as an effective title agency.

Transformational Problem #2: Routing Security

Communication over the Internet is dependent on two things: the ability to *identify hosts*, and the ability to *identify routes that exist between hosts*. Routing is currently a kind of self-governed commons, relying mainly on the implementation of the Border Gateway Protocol (BGP) and decentralized routing registries by Internet service providers, with a bow to the RIRs' IP addressing allocation and aggregation policies. Though it has worked remarkably well for 15 years, the Internet's routing system has no systematic security or authentication, which makes it susceptible to problems.

The problem of insecure routing, a long known vulnerability, was illustrated dramatically last year when YouTube disappeared from the Internet for about an hour. The disappearance occurred because an ISP in Pakistan, acting under the orders of government censors to block YouTube nationally, accidentally propagated a route

announcement globally through a transnational Internet connectivity provider.[23] As the route announcement was picked up by other Internet service providers, it had the effect of blocking YouTube across the entire global Internet. The same openness and flexibility, however, made it possible for Internet service providers to quickly correct the problem.

Authenticating IP address block assignments

Currently, there are efforts within the IETF and the RIRs to provide a method to authenticate the allocation of IP address block prefixes to Autonomous System Numbers (ASNs). An Autonomous System (AS) is a technical name for a network operator,[24] and AS numbers are unique integers assigned to network operators that are used to control routing. Since beginning work in 2005, the Secure Inter-Domain Routing (SIDR) Working Group has published several documents, including the Routing Public Key Infrastructure (RPKI) protocol,[25] which proposes to create a system containing digital certificates (as defined in RFC 3779) that bind organizational identity (i.e., ASNs) to assigned IP address block prefix(s).[26] According to the protocol, authenticating these certificates could rely on one or more trust anchors within the IP address space and AS number allocation hierarchy. However, the RPKI protocol does not specify which institutional entity(s) (e.g., the United Nations, the International Telecommunication Union, the U.S. Department of Commerce, ICANN, the RIRs) might maintain them. Alongside the standardization efforts at the IETF, the RIRs are developing the software necessary to support RPKI systems.

---

[23] Hong Kong-based telecommunication firm PCCW International.
[24] RFC 1930 defines an autonomous system as "a connected group of one or more IP prefixes run by one or more network operators which has a single and clearly defined routing policy.", *Guidelines for creation, selection, and registration of an Autonomous System (AS)*, available at http://tools.ietf.org/html/rfc1930#section-3
[25] See http://tools.ietf.org/html/draft-ietf-sidr-arch-03
[26] See http://www.ietf.org/mail-archive/web/sidr/current/msg00000.html, and SIDR Working Group charter and documents at http://tools.ietf.org/wg/sidr/

Authenticating route information

While being able to authenticate the allocation of IP address block prefixes to ASNs is important, it is only a piece of securing Internet routing. One must also have accurate and verified information about the interconnections or "routes" an autonomous system maintains and announces to the rest of the Internet.[27] Sometimes "route objects," which associate Internet routes with organizational ASNs, are stored in an Internet Route Registry (IRR). In theory, information in an IRR can be used to verify route announcements made by Internet service providers and to filter fake or erroneous ones. However, there is currently no way to globally authenticate route object data contained in IRRs. The IRR system is decentralized and voluntary. Several types of organizations (e.g., Internet service providers, RIRs) operate their own route registries, often mirroring route object data found in other registries. A meta-registry of all operating IRRs is maintained by Merit Network's Routing Assets Database (RADB). While it could be considered the *de facto* authoritative list for the IRR system, it does not have any contractual arrangements with other IRRs to maintain that information.

Creating property rights in address assignments and route announcements

In the past year, proposals have been submitted to ARIN,[28] APNIC,[29] and RIPE[30] to develop RIR-based routing registries that combine global RPKI authentication of prefix assignments with route object authorization information. This would provide the ability to authenticate not only what AS was using a particular prefix, but also what routes it announced to the Internet. But as Arbor Network's chief security officer and Internet Architecture Board (IAB) member Danny McPherson pointed out, implementing this kind of system also raises major governance issues. It could fundamentally change the role of IANA and the RIRs because it could be used

---

[27] Route objects are exchanged between organizations over the Internet using the Border Gateway Protocol, which is another fundamental, insecure Internet protocol. Efforts to make BGP more secure include sBGP, soBGP.
[28] See http://lists.arin.net/pipermail/arin-ppml/2008-May/010788.html
[29] See http://www.apnic.net/policy/proposals/prop-059-v001.html
[30] See http://www.ripe.net/ripe/policies/proposals/2008-04.html

to link the control of IP number resources to control over what is routed on the Internet. To quote McPherson,

> "Upon full employment of such a system, …the IP resources allocation hierarchy that exists today, which is sort of an out of band function that has no direct consequence on what's actually routed, now could have direct control over what's actually routed on the Internet, and perhaps most importantly, what's not. So, if you don't pay your RIR membership fees, your address allocations could actually be revoked, and this could trickle its way into the routing system, where filters might be augmented to discard your route announcements, or into a protocol like SBGP where it's actually automated."[31]

The basic point McPherson makes is that to "secure" a route, someone must be assigned an exclusive property right over the addresses to which a routing announcement refers, and also over the routing announcement space. Depending on how it is implemented, RPKI schemes could be used to make the enforcement of this exclusivity relatively automatic. Because the creation of this exclusivity function will likely rely on a hierarchical chain of certificate authentication, whoever controls the trust anchor(s) at the top of the hierarchy would be in a position to disconnect from the Internet anyone immediately below them in the hierarchy. This kind of power is, obviously, analogous to the control of the DNS root that became so contentious during and after the creation of ICANN. It creates a form of exclusivity and central control over the use of the addresses and routing that simply didn't exist before. It would constitute a fundamental change in the governance status of the RIR regime, as the RIRs could gain a direct impact on operations that they have never had before.

But it is not just the RIRs' role that might change. The newly created centralized point of control might also attract the attention of litigants and governments. Think of indirect liability for peer to peer file sharing, or efforts to track down spammers or phishers. ICANN's control of the DNS root gave it unavoidable forms of operational leverage over domain name registries and registrars, which served as a magnet for trademark/copyright interests and national governments

---

[31] D. McPherson, "IPv4 Exhaustion: Trading Routing Autonomy for Security." Arbor Networks blog http://asert.arbornetworks.com/2008/03/ipv4-exhaustion-trading-routing-autonomy-for-security/

seeking to assert forms of control over the internet. So could SIDR and RPKI give whoever controls the address assignment and routing authentication hierarchy the ability to exert significant forms of policy leverage over Internet service providers and their users. This trend could make the RIRs very much like ICANN indeed.

Or, worse, it could expand the concentration of authority within ICANN even more. In June 2008, ICANN's Security and Stability Advisory Committee (SSAC) asked the Board to budget a specific line item for "Management of certificates for the addressing system (RPKI).[32] Like DNSSEC, there are policy questions to be explored around the hierarchical chain of authentication and trust anchor control. These issues are only beginning to be discussed outside of the technical community.

The Need for a Policy Framework

More effective systems of keeping track of who is using which IP addresses, and more secure and exclusive address assignments can bring important benefits to Internet administration: more efficient resource utilization, better technical compatibility, avoidance of hijacking. Likewise, more secure routing and better authentication of the interactions among Internet service providers can prevent major abuses and failures and help to deter criminal activity on the Internet.

But any system of surveillance and administrative control can be politicized and abused. The ability to manage IP addresses is no exception. IP addressing has become the mechanism of choice both for tracking down Internet users and for enforcing actions against them.[33] Routing announcements could also become regulated for similar purposes. The RIRs currently lack any mechanisms or principles to regulate or limit efforts to utilize their leverage over address assignments and routing management for political, policy or legal purposes. There is nothing in their procedures that encourage them to consider, much less protect, individual rights to

---

[32] B. Kuerbis, "Will ICANN move to control routing security? IGP blog, 25 June 2008. http://blog.internetgovernance.org/blog/_archives/2008/6/25/3762527.html

[33] See the Electronic Frontier Foundation's case archive on *RIAA v. Verizon Internet Services*, in which an appeals court upheld Verizon's right to protect the privacy of its customers when challenged by copyright interests seeking to link IP addresses to specific customer identities. http://www.eff.org/cases/riaa-v-verizon-case-archive

privacy, freedom of expression or due process of law. Yet, the RIRs are slowly and, apparently, unconsciously backing themselves into a position in which they may exercise forms of power over Internet service providers and users that bear on these broader questions of public policy. This has to change. The stakes are too high.

The ICANN example: something to avoid

The RIRs' predicament creates a sense of *déjà vu*. As with ICANN, we see a global technical coordination system with a hierarchical delegation structure raising questions about the relationship between technical coordination and public policy for the Internet. It is our contention that that relationship as it first emerged around the formation of ICANN has since been handled badly – by the U.S. government, by WSIS, and by ICANN itself. Two aspects of the ICANN experience are worthy of notice. First, its unbalanced approach to rights protection, and second, the progressively expanding, camel's-nose- under-the-tent role played by governments.

During the creation of ICANN the U.S. Commerce Dept persistently rebuffed efforts to include freedom of expression as a principle of the ICANN regime, claiming that it was only "technical management." The U.S. also has persistently opposed efforts to incorporate privacy rights protections into ICANN's Whois. In both cases, it claimed to be focused on technical coordination, but in fact it was creating an Internet governance regime oriented around protecting the rights of trademark holders in the DNS. There was nothing inherently wrong with the effort to protect trademark rights, of course; what was wrong was the decision to elevate trademark protection above all other rights. The ultimate result was a biased regime that simply refused to recognize some kinds of rights claims.

Another disturbing aspect of the ICANN experience has been the long term evolution of governments' role in ICANN. The initial plan of ICANN was to *privatize* DNS administration and to keep governments out entirely. However, the U.S. first asserted for itself a temporary supervisory role. Once in that position, the U.S. refused to give it up its oversight powers – and the trademark interests and national economic interests (such as VeriSign) who benefited from the regime encouraged that stance.

The special U.S. role provoked other governments, heightening their demands for influence and control over ICANN's day to day decision making processes. This power competition found a voice during WSIS, but seems to have only strengthened the U.S. resolve to hold on to its privileged position.

ICANN originally dealt with the intersection of its technical management and public policy functions by creating a Governmental Advisory Committee (GAC). GAC was initially conceived as a way of allowing governments a purely advisory and informational role while keeping them at arm's length. But this did not work. GAC has regularly expanded its role in ICANN, and there are strong pressures to keep expanding it.[34] It began by asserting authority over ccTLD delegations and moved on to assert for itself a privileged right to dictate anything that might be called a "public policy matter." The GAC's "advisory" role has now become a parallel policy making process that is equal if not superior in status to ICANN's other Supporting Organizations. But GAC is not a treaty organization. When it proclaims something as its official "policy" it is not required to negotiate binding instruments and its constituent national governments are not required to go back to their legislatures to get approval. In effect, ICANN enables an informal consensus among a handful of governments to have binding force of international law. This isn't right.

ICANN's new gTLD process provides a case in point. Five years ago it was inconceivable that ICANN would attempt to regulate the content of Internet web sites. And if governments attempted to negotiate a treaty that regulated speech internationally, both the U.S. and Europe would be bound by legal or constitutional guarantees of free expression and legislative and electoral checks and balances. But GAC policy advice is not bound by these constraints. Thus, due to a lethal combination of GAC policy principles and trademark/copyright owners, ICANN will censor a wide class of top level domain names irrespective of any free expression

[34] See the speech of ITU Secretary-General Hamadoun Toure, 6 November 2008 at the ICANN Cairo meeting. https://cai.icann.org/files/meetings/cairo2008/toure-speech-06nov08.txt

claims, not because of anything harmful in the names themselves, but because of the content that might be found under such domain names if they are assigned.[35]

Thus, ICANN is rapidly evolving into a "worst of both worlds" situation: it eliminates many state-granted legal rights and due process protections by giving a private actor extensive regulatory and taxing power over the Internet, but it also facilitates unbalanced and extra-legal governmental interventions.

Multistakeholderism and "bottom up" are not the answers

The new fashion, when confronted with these problems, is to sing the praises of "multistakeholder governance," and so-called "bottom up decision making." The ideology of multistakeholderism celebrates the indiscriminate mixing of states, private business and technical governance. While multistakeholderism has served a good purpose in opening up intergovernmental organizations to participation by broader segments of society, it does not solve the problems identified in this paper, and could easily make them worse.

Up to this point, RIRs have presented themselves to the public as bottom-up, open and effective governance structures. In many respects this is true. They are far more transparent and open to public participation than your typical intergovernmental policy making organization. Board members and Advisory Council members are directly elected by members, not through complicated and indirect representational and nominating committee processes as in ICANN. The policy making processes they run seem much less manipulated by management, and there is no equivalent to the U.S. Commerce Department political oversight hovering in the background to manipulate the process.

The RIRs are currently based on a membership model, and joining an RIR presumes that the party joining the policy making process has a material stake in management of the resources. This is, on the whole, a good thing, because it closely aligns the primary users of the resources with the policy making process. RIRs are

---

[35] See the ICANN memo on "Morality and Public Order Objection Considerations in new gTLDs," 29 October 2008. http://www.icann.org/en/topics/new-gtlds/morality-public-order-draft-29oct08-en.pdf

similar to the IETF in that they are composed primarily of technical experts with a common background and shared norms. To invite the general public and government policy makers into this process actually contributes to the problem we have identified, namely the growing danger that technical management will be used to exert policy leverage without any well-defined rules or constraints to protect basic human rights. Broader, indiscriminate participation in RIR processes simply invites anyone with a political agenda to meddle in technical governance processes. While this may seem more democratic, in reality it only opens the door to organized special interests. Mr. and Ms. Ordinary Internet User are not going to spend weeks of their time reading email messages about route aggregation, provider-independent address blocks, and IPv6 migration strategies. The only people who will have the commitment and energy to enter into these processes will be organized interests with a political axe to grind. That means, e.g., copyright and trademark interests seeking leverage over Internet users,' censorship advocates looking for new ways to get a grip on content providers, and so on.

Toward a neutral core

In order to avoid a repeat of the problems of ICANN, we need to think carefully about the relationship between RIRs, governments, and Internet freedom. As the RIRs migrate into a more formal and potentially politicized global governance regime, we need to find ways to institutionalize and preserve the RIRs' mandate to remain minimalist technical coordinators whose sole mandate is maintaining efficient and effective Internet connectivity. Efforts to exploit their technical leverage to achieve policy objectives unrelated to address management must be resisted. We need to develop a legal and institutional firewall that separates the address management policies from the public policy concerns that need to be carried out at the national level by governments. This does not mean that governments, law and public policy considerations have no role to play in Internet governance. It simply means that those considerations have to be developed separately, in other institutions. The following principles might serve as guidelines:

- IP address management is best maintained through voluntary participation in transnational, private sector nonprofit organizations that do not have centralized and coercive authority over Internet service providers.
- The RIR's transnational mechanisms for Internet address resource allocation policies cannot and should not attempt to implement national-level public policies.
- Insofar as possible, technical standards and protocols designed to solve security problems should avoid creating global bottleneck facilities that concentrate a dependency on a single point in the network.

Current policy trends threaten to drag the RIRs into the more political Internet governance debates. The impact of their decisions remains centered on the technical infrastructure, but is expanding to include more economic, legal and political dimensions. The small community congregated around RIRs would probably be the first to admit that it is not in a position to make authoritative decisions about the proper tradeoffs between free expression and legal responsibility on the Internet, between individual privacy rights and public safety, between technical efficiency and competition, for the entire internet. Thus it is important for RIRs to restrict themselves to a narrowly technical mission and to resist attempts to load policy functions onto their address management mechanisms.