

Managing the Last eights: Three Ways forward for IPv6

L. Jean Camp[†], Rui Wang[†], Xiaoyong Zhou[†]
ljcamp@indiana.edu[†], wang63@indiana.edu[†], zhou@indiana.edu[†]
[†]School of Informatics, Indiana University at Bloomington

August 9, 2009

Abstract

IP addresses are critical resources for Internet development. The available IPv4 address pool is projected to fully allocated within 3-5 years, but the deployment of the alternative protocol, IPv6, is not accelerating. Recently, several proposals have introduced new approaches to managing IPv4 allocation. However, the market proposals in particular imply fundamental changes to the nature of IPv4 address allocation. This work analyzes the allocation history of ARIN, the organization which manages address allocation in the America region. Based on the historical data and projected trends, we propose three simple management policies which can be immediately implemented with limited resources and without either significant registrar cost nor drastic modifications on current allocation strategy. We verify the effectiveness of our policies using historical data. ¹

1 Introduction

Internet Protocol version 4 (IPv4) [1] is the fourth revision in the development of the Internet Protocol (IP) [10], and was the first version of the protocol to be widely deployed. IP address exhaustion refers to the decreasing supply of unallocated IPv4 addresses. IPv4 exhaustion has been a concern since the 1980s when the Internet began to experience dramatic growth. Address scarcity or perceived address scarcity was a driving factor in creating and adopting several new technologies, including classful networks, CIDR addressing, Network Address Translation (NAT) and a new version of the Internet Protocol, IPv6. The transition of the Internet to IPv6 is argued to be the only practical and readily available long-term solution to IPv4 address exhaustion. As the predicted IPv4 address exhaustion approaches its final stages, most ISPs, software vendors and service providers have neither adopted nor deployed IPv6.

Recent analysis shows that the unallocated IPv4 address pool will be exhausted within 3-5 years [6]. The barrier to entry on the Internet for new service providers and some class of services could become insurmountable. This is because there is the potential to lock out innovations that require routable blocks of IP addresses. To alleviate the issue, several proposals [7, 12, 13, 15] have discussed the possibility of enabling a market for transferring IPv4 addresses, with which

¹This paper is a continuation of the previous WEIS work. We acknowledge it may be out of scope.

organizations in need of more addresses could purchase them from organizations which have more allocated addresses than they require. However, building a market is far from trivial. Pricing, market clearing, dispute resolution and the prevention of speculation all must be correctly addressed. In addition, inappropriately allowing the trading of IPv4 addresses could have significant negative effects on the costs of Internet routing, retarding the migration to IPv6, and thereby adversely impacting Internet growth and architecture. [15]

We argue that a successful allocation strategy of IPv4 addresses should be simple and easy to deploy. It should require minimum structural changes to current allocation policy. Our proposals require minor revisions of current policies, and also avoid a market and a property rights regime. Based on these observations, we propose three simple policies to enhance the current allocation policies of ARIN for the purpose of extending the lifetime of ARIN's unallocated IPv4 addresses:

1. Only allocate to organizations with a small address space previous allocated,
2. Only allocate a given amount of the v4 space per year, or
3. Provide only a minimal routable allocations per organization.

Section 2 presents the related prior research. Section 3 describes the our analysis and introducing data collection. Section 4 gives a detailed description of our modeling, in which we propose three policies for ARIN, and evaluate the effectiveness of our approach with historical data. Section 5 concludes the paper and discusses the future research.

2 Related Work

Depletion of the IPv4 unallocated address pool was originally predicted to occur in 2037 [2]. However, the most recent estimates for RIRs is April 2012, less than 4 years from now [6, 4]. With the imminence of IPv4 depletion, many proposals and papers have introduced approaches to deal with the depletion issue. [14] applies several mathematical models to project IPv4 address depletion trends, and concludes that the recent consumption rates of IPv4 will not be sustainable from the central pool beyond this decade. [13] identifies valid measures of IPv6 diffusion and uses classic diffusion models to bound the uncertainty in those measures, and concludes that there is no reasonable case for diffusion of IPv6 before IPv4 full allocation under current policies. [15] introduces the challenge of managing Internet addresses, and discusses how market forces might be introduced to enhance the management of the Internet address space. [12] proposes that network incentives impede the transition from IPv4 to IPv6 - effectively requiring mechanisms to preserve the current IPv4 numbering system. The author argues for a a paid transfer system for IPv4 addresses to ameliorate the negative effects of IPv4 scarcity. The online report [6] gives a detailed analysis of IPv4 addresses: current allocation policies, the status of address consumption, and projection of address depletion.

Several RIRs, including ARIN, are considering modification of RIR policies to slow the depletion of the available addresses in their pools, or speed up the deployment of IPv6 protocols. One proposal [7] endeavors to enable "Simple Transfer of IPv4 Addresses" in which organizations are allowed to transfer IPv4 addresses from each other if they are subject to several restricted conditions. Another proposal [8] suggests setting a contiguous /10 IPv4 block which is dedicated

to facilitate IPv6 deployment. In order to receive an allocation of this block, applicants have to demonstrate that they are in immediate need of IPv4 addresses for the purpose of IPv6 deployment.

Our work is different from the above proposals in two aspects. Firstly, the prior work focuses on the depletion of the IANA unallocated address pool. Our work studies the allocation characteristics of ARIN, and tries to propose effective policies which could help ARIN extend the lifetime of its available addresses. Secondly, our approach is simple, and can be immediately implemented, as opposed to the prior-proposed approaches (e.g., paid transfer system) which requires both more processes and a fundamental change in the conception of IPv4 allocation for implementation.

3 Data Experiment Setup

In this section, we first introduce research method. Then we introduce the data collection approach with detailed explanation of how the data is compiled.

3.1 Framework

An overview of our research method is provided in Figure 1. Our data collection procedure consisted the tasks required to assemble historical allocation data provided by ARIN, and populate the data to the database for further analysis. The details of data collection are explained in Section 3.2. The data in mysql database has a simple format which is shown in Figure 1. There are three fields in the table: date, address block, and organization. Each entry of the table denotes an event in which ARIN allocated the designated address block to the organization in the denoted date. The database can provide the IPv4 address allocation distribution across different organizations. It also gives us the allocation history of how IPv4 address was allocated and its historical allocation rate to different organizations. With this dataset, we conducted data analysis and trend analysis as described in detail in Section 4.

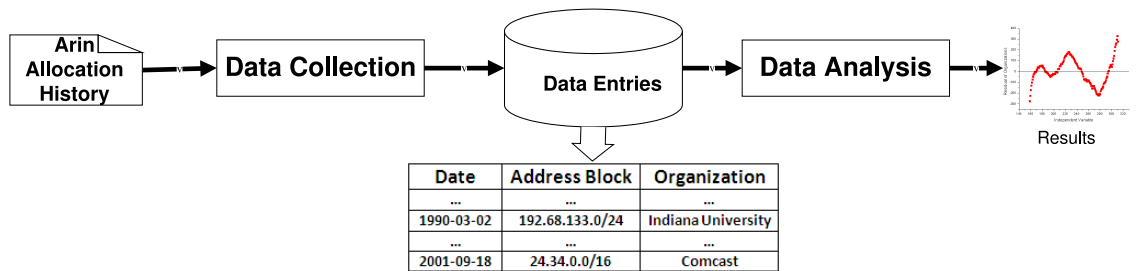


Figure 1: Data Compilation Framework.

3.2 Data Collection

We have two data sources. One data source is a Regional Internet Registry (RIR) statistics file, which represents all of the allocations and assignments made by RIR [9]. The RIR statistics file

Table 1: Data format.

Field name	Description	Values
registry	RIR	{apnic,arin,iana,lacnic,ripenc}
cc	country code	{A,PEU,UK}
type	type of Internet resource	{asn,ipv4,ipv6}
start	first address of a IP block	e.g., 156.56.0.0
value	the number of addresses in the block	e.g., 65536
date	date on the allocation	e.g., 1991-12-10
status	type of the allocation	{allocated, assigned}

summarizes the current state of allocations and assignments of Internet number resources. They are intended to provide a snapshot of the status of Internet number resources, without any transactional or historical details. The format for each record in the file is shown in Table 1. Each line in the file represents a single allocation (or assignment) of a specific range of Internet number resources (IPv4, IPv6 or ASN), made by the RIR identified in the record. In the case of IPv4, the records may represent non-CIDR ranges or CIDR blocks, and therefore the record format represents a beginning of range, and a count. This can be converted to prefix/length using simple algorithms. In the case of IPv6 the record format represents the prefix and the count of /128 instances under that prefix. We extracted the entries in which the RIR is ARIN, and the protocol is IPv4.

The first data source gives the IPv4 address allocation history of ARIN. However, there is no information which maps an assigned address block to the AS which received the IP block. Since there is no consistent data source of which IPv4 block went to which AS at a given date, we had to reconstruct that data set. We did this by querying two whois databases to get the information. One database [5] has incomplete data but can be queried continuously, ARINs [3] is complete but rate limited.

4 Data Analysis

Our analysis has four parts. The first part of our analysis projects the ending date of IPv4 address allocation for ARIN under current policies, and analyzes the distribution characteristics of ARIN’s address allocation. This not only provides a breakdown of the data, but provides a high level validity check on our constructed database. Then we introduce three policies designed not only to delay the exhaustion date of IPv4 but also to encourage IPv6 adoption by those organizations with the greatest technical resources and need. These policies are limiting allocation based on previous allocations, setting a maximum per organization annual allocation regardless of organization, and setting an annual allocation amount without regard to organization. These are described in Section 4.2, Section 4.3, and Section 4.4 respectively.

4.1 Address Distribution across Organizations

Estimate of available addresses for ARIN. Since not all addresses have been assigned from IANA to RIRs, we cannot know the exact pool of addresses that IANA will allocate to ARIN. Our current design uses a simple strategy to estimate the number. Since the current IANA address pool has 36 /8 unallocated blocks, one can imagine that if IANA evenly assigned these blocks to the

Table 2: Organization Statistics.

Organization Category	# of Organizations	Avg. of Allocated Addresses
IP \geq /8 block	34	395 /16s
IP in [/12 block,/8 block]	103	49 /16s
IP in [/16 block,/12 block]	1268	712 /24s
IP in [/20 block,/16 block]	3207	48 /24s
IP in [/24 block,/20 block]	4302	5 /24s
IP \leq /24 block	8202	1 /24s




Figure 2: Organizations.

five RIRs, each of the RIRs could get 7.2 /8 blocks. In addition, the current ARIN address pool has roughly 3 /8 blocks of unallocated addresses. We anticipate ARIN will have 10.2 /8 blocks (171,127,603 addresses) in total for future allocation. Thus we can argue that we are presenting optimistic but reasonable results. That is, our results are within the range of other work.

Table 3: Summary of Exhaustion Projections with Models, Year - Month

Linear Estimate		Polynomial Estimate	
2014-07		2012-11	
Lower Bound	Upper Bound	Lower Bound	Upper Bound
2013-08	2015-06	2012-04	2013-05

Exhaustion date of ARIN address pool. We used a linear projection based on the historical the addresses allocation by ARIN to predict the exhaustion date of its unallocated address pool, which is similar to prior research [6, 14]. However, we calculated a confidence interval(CI) [11] for the projected model, which describes the uncertainty of the projection under specified possibility. For example, a 95% CI means that the real data will fall into the predicted area of the model with a probability of 95%. The starting date for the projection is important. We estimated an appropriate starting date by analyzing Figure 2. According to the figure, the number of organizations applying for addresses from ARIN had a significant changes in 1996. After 1996, the number of organizations increased monotonically without much nonlinear change. So we use a linear projection starting from 1996. Our projection is described in Figure 3. Taking into consideration of the 95% CI, the figure shows that available addresses in ARIN pool will be exhausted with a certainty of 95%, between 2012 and 2014.

Organizational distribution of allocated addresses. In Table 2, we categorize organizations based on number of their allocated addresses into 6 classes, and compute several attributes for




Figure 3: Exhaustion Projection with Current Trend.

these categories. The category in which organizations at least owned a /8 block of addresses has 34 members, and each member has 395 /16s on average. A similar situation appears in the second category of 103 members, in which organizations have been allocated addresses between a /12 block and a /8 block. The average number of addresses in this category is 49 /16s. The other four categories have 16979 organizations in total, but received only a small portion of addresses for each. Based on the above observation, Those in the last four categories may have both more urgent needs for IPv4 addresses than members in the first two categories and lower total organizational network expertise. Policies which benefit the four categories could resolve the urgent needs of most organizations. There is also a category that cannot be predicted of future innovators who now have zero.

4.2 Analysis of Policy 1: Addresses Reserved for those with Smaller Past Allocations




Figure 4: Addresses Allocated to Organizations Based on a Threshold.

Table 4: Projected exhaustion date for policy 1.

Projection Type	Threshold /12			Threshold /16			Threshold /20		
	Projected	-% 95 PI	+% 95 PI	Projected	-% 95 PI	+% 95 PI	Projected	-% 95 PI	+% 95 PI
Linear	2024-01	2023-11	2024-04	2041-12	2040-07	2043-08	2060-12	2056-08	2066-04
Polynomial	2019-03	2019-03	2019-03	2029-06	2029-05	2029-07	2045-12	2045-11	2046-02




Figure 5: Organizations with more than /12 addresses no longer receive addresses.

Taking into consideration of Table 2, we introduce an allocation policy for ARIN which allows allocations only to organizations that have previously been allocated addresses below a given threshold. This policy prohibits the organizations with large blocks of addresses from being allocated more in the future. This may entail agreements upon allocation that transfers to these large organizations would result in revocation of previous allocations to prevent bundling and the use of subsidiaries. To the extent that cost drives IPv6 adoption, the cost of organizational subterfuge would create a price on IPv4 allocations that applies only to those with current generous allocations.

To validate our policy, we modeled three thresholds for allocation: /12 block, /16 block, and /20 block. We applied a linear projection model to the data to predict an exhaustion date for each tested threshold. Figure 4 shows the "residual" addresses allocated to organizations since 1996. The "residual" addresses are the portion of addresses that have been allocated to organizations above the specified threshold. This figure shows that at least 250 million of "residual" addresses have been allocated. To project the exhaustion date under the policies, we subtract the "residual" addresses from total addresses, and apply projection models on the generated data.

That is, each projection is a four step process.

1. Subtract residual addresses from historical data.




Figure 6: Organizations with more than /16 addresses no longer receive addresses.

2. Use this data to form a projection, calculating linear and polynomial fits.
3. Add the residual addresses back into projection for a starting or initial point for the projection.
4. Use the curve generated to project exhaustion.

Figure 5 shows the projection result for the policy with threshold /12. The linear model gives an exhaustion date of 2024, and the polynomial model gives 2019, extending the lifetime of IPv4 at least 5 years. Figure 6 and Figure 7 show the projection graphs with threshold of /16 and /20 respectively, which give more promising results of at least a 15 years gap. The summary results are listed in Table 4. This corresponds to previous findings, which illustrate that few large players have the largest influence on exhaustion.

The substantive argument for this policy is threefold: capacity, motivation, and impact. First, those organizations with the largest networks have the greatest aggregate experience. These organizations are likely to employ address network management techniques and employ highly skilled network engineers. Therefore these organizations are most likely to have the capacity to adopt next generation addressing. Second, these organizations are most likely to profit from the exclusion of innovations that might result from a lack of available IPv4 addresses. As the largest, the limitation of competition is structurally advantageous. To the extent that IPv4 scarcity increases the value of their asset (i.e., a large allocation) then that scarcity is advantageous to these organizations. Limiting allocation can alter that motivation. Third, the larger the organization and network that adopts next generation addressing the larger the impact on the network as a whole.




Figure 7: Organizations with more than /20 addresses no longer receive addresses.

4.3 Analysis of Policy 2: Fixed Address Allocation Per Organization

Our analysis of the first policy demonstrates that it has the potential to achieve the effective management of address depletion. However, that policy may be unfair if it ignores the needs of organizations which have already been allocated a large block or blocks of addresses. This section introduces second allocation policy, in which any applicants can receive a block of addresses no more than a threshold. Since many applicants only seek allocations beneath that threshold, the needs of these can be met under this policy.

To verify the effectiveness of the proposed non-discrimination policy, we estimate an upper bound of address consumption per month under the policy, which can be computed by simply multiplying the number of applicants each month with the specified threshold. The number of applicants in the future is predicted by projecting number of recent applicants since 1996, which is shown in Figure 8. Interestingly, a polynomial model was found to be the best fit for the projection. With the data shown in this figure, we computed the upper bound of address consumption per month, and tried to identify the date in which the upper bound reaches the amount of available addresses. The date denotes that, in worst case, when the available addresses will be exhausted. Figure 9 shows the result with the threshold of /22 addresses, which promisingly gives an exhaustion date of around 2060. We also studied the result with a threshold of /20, shown in Figure 10, and observed an exhaustion date of around 2030. Our last result is shown in Figure 11 with a threshold of /18 and an exhaustion date of 2017.

The argument for this policy is that it maintains a standard of non-discrimination. Any organization which needs either a new allocation or marginally more IPv4 addresses may be unready for a move to next-generation addressing, regardless of size and previous allocations. Again, it is possible that organizations will result to subterfuge. In this case, the subterfuge would be to obtain

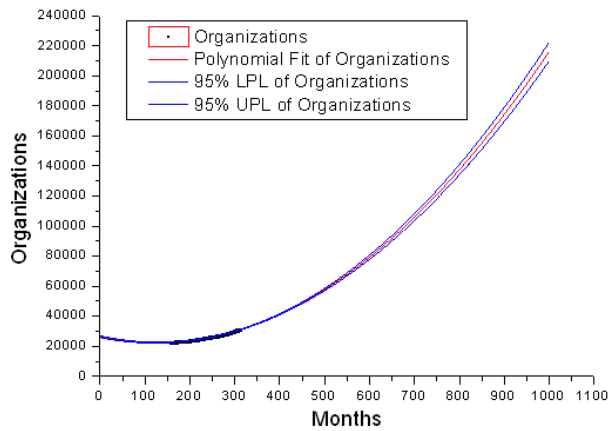


Figure 8: Projection for organizations.




Figure 9: Each organization receives no more than /22 addresses per allocation.

larger blocks. Should organizational subterfuge prove possible the result would be non-contiguous blocks of IPv4 addresses. Should organizational subterfuge be expensive, then this adds to the price of staying with IPv4.

4.4 Analysis of Policy 3: Fixed Address Allocation Pool Per Year

We studied the distribution of allocations in 2008. The distribution is shown in Figure 12. In this figure, the X-axis refers to the allocation events, and Y-axis refers to number of addresses allocated in the corresponding event. It shows that most of the events only allocated a small portion of IPv4 addresses (well below 1,000,000), but a few events allocated a large portion. This observation serves as the basis of our third policy. This policy enables the registrar to flexibly control the lifetime of available addresses without discouraging the needs of most organizations. Specifically, the enforcer first designates an anticipated IPv4 lifetime (e.g., 20 years), and then an upper bound of addresses is calculated by dividing the anticipated available address pool with the lifetime. The upper bound serves as the hard threshold of addresses available for allocation per year, so the




Figure 10: Each organization receives no more than /20 addresses per allocation.




Figure 11: Each organization receives no more than /18 addresses per allocation.

exhaustion date is under control of the RIR.

We consider the outcome of this policy in Table 5. The first column of the table gives different lifetimes for IPv4 allocation, the second column gives the calculated upper bound of addresses to allocate, and the last column shows the percentage of organizations which could receive addresses requested had the policy been applied in 2008. The result is that the policy could extend the lifetime of IPv4 to 30 years while still meeting the expressed needs of more than 95% applicants. Even if the lifetime were extended to 50 years, the requests of more than 75% of organizations could in theory, be met.

Again, an advantage of this policy is organizational nondiscrimination. The strictly first-come first-serve nature of the policy could provide certainty, limit speculation, and allow an orderly transition to next-generation addressing.

5 Conclusion

From historical data we classified allocations by organization, size, and date. Based on these data, we proposed the three simple policies above to extend the lifetime of the unallocated pool




Figure 12: Address allocation events in 2008.

Table 5: How many organizations' request would have been met under policy 3.

Lifetime of IPv4(year)	Upper bound of addresses to allocate	Percent Requests Fulfilled 2008
10	261 /16s	99.4%
20	130 /16s	97.9%
30	87 /16s	95.0%
40	65 /16s	87.9%
50	52 /16s	76.9%

of IPv4 addresses in ARIN. We verified the potential of our proposals using historical data. We demonstrated that all three of the proposed policies could effectively extend the lifetime of IPv4 for dozens of years without impeding the development of most organizations.

This paper analyzed the address allocation history of ARIN, and considers three straightforward allocation policies. For the first policy, we assert that some set of organizations (defined in this paper as two classes of address holders) have ample addresses for future use combined with significant aggregate network expertise, perverse incentives for adoption of next-generation addressing, and the potential for the greatest influence by adopting next-generation addressing. Thus the first proposal is to refuse future IPv4 allocations to that set of organization. For the second policy, we propose that IPv4 addresses be distributed to any organization but only once per organization per year, and in size-constrained allocations. This would not distinguish between organizations. The third policy proposal is to simply determine when it will be feasible to fully allocate the current pool, and allocate each year up to that limit and then cease allocating for that year.

Our result is preliminary, and future work is needed for better understanding and application of the proposals. Firstly, we gave a only rough estimate of the effectiveness of our policies, but more accurate evaluations are necessary for quantifying the effect of these policies in practice. Secondly, each of our policies has different features. Simply applying one policy may be not adequate for effective management. A combination of multiple policies will be studied in the future with more detailed projection models. Last but not least, our policies are general, and can be applied to other RIRs to resolve their address depletion issues. This is also an important future direction.

The greatest risk in these models is the assumption of equal allocation from IANA. IANA may choose to allocate based on illustrated need, thus ensuring that limited IPv4 allocation is a management choice that proves harmful to the RIR and the region it is meant to serve.

Table 6: Top 20 owners of IPv4 addresses.

Organization Name	# of Addresses
DoD Network Information Center	159501056
DDN-ASNBLK1 - DoD Network Information Center	83060992
Level 3 Communications	42936320
AT&T Internet Services	37652736
Comcast Cable Communications, Inc.	35594240
Cogent/PSI	27318784
MCI Communications Services, Inc.	25478656
AT&T WorldNet Services	22535168
Headquarters, USAISC	20205824
Cellco Partnership	19953664
Merit Network Inc.	18115328
AT&T Global Network Services	17247744
E.I. du Pont de Nemours and Co.	17175808
University of California at San Diego	17105664
HP-INTERNET-AS Hewlett-Packard Company	17049600
SITA	16983296
Massachusetts Institute of Technology	16973824
Ford Motor Company	16940544
Interop Show Network	16845824
HP-EUROPE-AS Hewlett-Packard Company	16842752

References

- [1] RFC791: Internet protocol, darpa internet program protocol specification. <http://tools.ietf.org/html/rfc791>, 1981.
- [2] IPv4 address lifetime expectancy revisited. <http://www.iepg.org/november2003/>, 2003.
- [3] ARIN whois database search. <http://ws.arin.net/whois>, 2008.
- [4] The end is near, but is ipv6? <http://lists.arin.net/pipermail/arin-ppml/2008-November/012720.html>, 2008.
- [5] IP to asn mapping - team cymru. <http://www.team-cymru.org/Services/ip-to-asn.html>, 2008.
- [6] IPv4 address report. <http://www.potaroo.net/tools/ipv4>, 2008.
- [7] Policy proposal 2008-2: Ipv4 transfer policy proposal. http://www.arin.net/policy/proposals/2008_2.html, 2008.
- [8] Policy proposal 2008-5: Dedicated ipv4 block to facilitate ipv6 deployment. <http://lists.arin.net/pipermail/arin-ppml/2008-November/012720.html>, 2008.
- [9] RIR statistics data. <ftp://ftp.arin.net/pub/stats/arin/delegated-arin-latest>, 2008.

- [10] V. Cerf and R. Kahn. A protocol for packet network intercommunication. *Communications, IEEE Transactions on [legacy, pre - 1988]*, 22(5):637–648, 1974.
- [11] C. Chatfield. Calculating interval forecasts. *Journal of Business & Economic Statistics*, 11(2):121–35, April 1993.
- [12] B. Edelman. Running out of numbers? the impending scarcity of ipv4 addresses and what to do about it. <http://www.benedelman.org/publications/runningout-draft.pdf>, June 2008.
- [13] H. Elmore, L. J. Camp, and B. Stephens. Diffusion and adoption of ipv6 in the arin region. In *WEIS '08: Proceedings of the 7th Workshop on the Economics of Information Security*, 2008.
- [14] T. Hain. A pragmatic report on ipv4 address space consumption. *The Internet Protocol Journal*, 8:2–19, September 2005.
- [15] W. Lehr, T. Vest, and E. Lear. Running on empty: The challenge of managing internet addresses. In *TPRC '08: Proceedings of the 36th Research Conference on Communication, Information and Internet Policy*, 2008.