The Internet Protocol Journal

March 2009 Volume 12, Number 1

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

You can download IPJ back issues and find subscription information at:
www.cisco.com/ipj

ISSN 1944-1134

From The Editor

IP Version 4 address exhaustion and migration to IP Version 6 continues to be the focus of many Internet-related organizations and events. The *Regional Internet Registries* (RIRs), still debating what will happen as the IPv4 address pool runs out, are developing policies for how to manage address-block transfers between address holders. One potential result of the address shortage is that a *market* (official or otherwise) will develop for the buying and selling of IPv4 addresses. In our last issue, we brought you the first in a two-part series of articles entitled "The End of Eternity," by Niall Murphy and David Wilson. Part Two, included in this issue, discusses what a market-based IP trading exchange might look like.

IP address allocation, transfers, and even the potential trading market for addresses is ultimately dependent on a reliable and trusted registry for this information. The RIRs have been working on a way to ensure that information about *IP Number Resources* (that is, IPv4 addresses, IPv6 addresses, and *Autonomous System* [AS] numbers) are securely stored and distributed so that users of such information can be assured that it is authentic. The underlying technology is a *Resource Certificate Public Key Infrastructure* (RPKI), and it is described in our second article by Geoff Huston.

The Internet technical community is discussing the so-called *identifier/locator split* as a major change to the Internet architecture. The IETF is developing several proposals, including the *Locator Identifier Separation Protocol* (LISP) discussed in our March 2008 issue. In this issue we look at another proposal, the *Host Identity Protocol* (HIP). The article is by Andrei Gurtov, Miika Komu, and Robert Moskowitz.

You will notice that our back cover has a new look. This layout is not the result of any creative design urges, but rather a change in U.S. Postal Service regulations regarding the placement of the subscriber address label. I guess the Internet isn't the only place where addressing is a major topic.

As always, your comments, suggestions, and contributions are welcome, including Letters to the Editor, Book Reviews, and of course full-length articles. Our Call for Papers is included on page 35. Contact us by e-mail at ipj@cisco.com

—Ole J. Jacobsen, Editor and Publisher
ole@cisco.com

The End of Eternity
Part Two: Address Space Trading and the Routing Table

by Niall Murphy, Google, and David Wilson, HEAnet

In our last article[0], we wrote about the onset of scarcity and the problems that are likely to ensue as a result. We characterized the problem we face as the *gap*, the length of time between the end of IPv4 plenty and the beginning of a universally reachable IPv6 Internet. Noting that any solution should either make the gap shorter, by bringing forward full IPv6 deployment, or make it less painful, by reducing the pressure of IPv4 scarcity, we propose that the fairest, most neutral way to encourage networks out of IPv4 while providing help for those who need it is to introduce a market-based IP address trading exchange. Let us explore now how such a system could work.

Possible Market Structures: Advantages and Drawbacks

An exchange could be set up and operated in many ways. Our preference, however, is for such a service to be run by the existing, trusted, and stable *Regional Internet Registries* (RIRs). Not only are they experienced in maintaining the values that the community as a whole wants to see maintained—fairness and neutrality, transparency, etc.—the RIRs are also in an excellent position to establish the *quality* of prefixes traded in an exchange, having excellent service contracts and history with members. Furthermore, the RIRs are unlikely to be made available for onward sale or transfer to other organizations with "different values," and would maintain their traditionally community-focused policy-making apparatus. They would also be in a position to act quickly to coordinate and assume responsibility if given sufficient authority by the membership.

It does not have to be an RIR, of course: we *could* set up another industry body, but it would take valuable time and require a new governance model. We could also outsource the whole thing to any professionally run auction-handling site, but for such a fundamental change in how we do things, it seems wise to keep it under direct control. Finally, the psychology of continuity is important; if organizations are used to dealing with the RIRs, it provides an important perception of stability to keep them as the interface to getting new addresses.

As with our previous article, we emphasize again that the RIRs have provided excellent service in focusing the consensus of the community in a form that can be passed back to governments and other stakeholders, both external and internal.

The shield provided by the RIRs, protecting the members from the outside and protecting the members from themselves, has worked well for three reasons:

- First, RIR consensus is widely seen to broadly reflect the wishes of their communities as a whole because of the extremely low barrier to representation—in essence anyone who cares can attempt to influence policy, and no formal attempt is made to weigh one set of opinions over another. As a result, RIR policy is a lowest common denominator that is in general free from many of the more partisan stances usually found in the telecommunications arena, leading to greater credibility outside the RIR system, and greater credibility within, because the oppression of a minority by the majority within the context of policy formation is very difficult.
- Secondly, possessing that credibility has led to repeated success for the RIRs in the arena of disseminating and explaining policies outward, and they have therefore reinforced the confidence their members have in them.
- Finally, the RIRs are also comparatively financially easy to run; in the *Réseaux IP Européens* (RIPE) region, fees are by no means excessive given the ratio of customers to addresses; they are observed and validated by RIPE *Network Coordination Centre* (NCC) members, and any competing industry body would have to duplicate not only all the previously mentioned activities, but also the large working surplus that allows the RIRs to ensure stability through more turbulent times. Or to put it another way, "it's open, it works, and it's cheap." We would recommend that any significant extension to the RIR authority, such as running an exchange as proposed, should endeavour to preserve as many of these properties as possible.

So if RIRs are to be the point of contact and policy making, how might such an exchange operate? We have a few guidelines from a relatively new field of economics, called *Market Design Theory*[21], that might help to inform our choices. Firstly, we must have *thickness*: we must have enough traders (both buyers and sellers) entering the market, such that the populace at large can be assured that if they need to perform a transaction, the exchange is the place to do it, rather than private trades. (Private trades, although they enable liquidity, have the disadvantages that the WHOIS database is not maintained, that policy cannot be centralized, that prefix de-aggregation can occur arbitrarily, and so on.) We should avoid *congestion:* so many participants that it becomes difficult to trade. Finally, we must have *safety:* the assurance that if a transaction is engaged in, it will complete, and buyers will receive what they want.

Although other properties exist, those are the main ones required for the exchange to operate successfully. On thickness, we think it is clear that attracting buyers in a time of scarcity will not be a problem. The problem will be attracting sellers from such constituencies as have them available (old *Internet Assigned Numbers Authority* [IANA]-allocation holders, dot-com failures, and so on). An open question is whether the exchange can do more to attract sellers than the monetary reward for selling would do on its own; more meaningful incentives for them are difficult to determine. Overall, congestion does not seem likely to be a concern, given that the RIR model most usefully supports only membership-based participation initially. (Furthermore, our guess is that the "product" will be quite homogenous, so performing trades will presumably be mostly a matter of determining price.)

Let us return to the question of prefix *quality*. The single most important measure of quality of a prefix, the attribute without which the prefix is useless, is *uniqueness*. One must be assured that the prefix one holds is acknowledged as being held by oneself, and that *Internet Service Providers* (ISPs) will accept its announcement from *no other* parties.

From a plentiful pool, where prefixes have no cost other than the service charge of the registry, ensuring uniqueness is perhaps not a simple task, but it is a relatively uncontroversial one. When scarce, prefixes become valuable and will be given a cash value, either officially or by other means. ISPs will then have a business reason to break with consensus on routing filters, as we discuss later in more detail; but regardless, prefixes allocated from the IANA free pool generally have an impeccable heritage and do not vary greatly in usability. There are, of course, the natural delays in having new /8s incorporated into routing filters across the world. Those delays do have real effects, but the recipient of these prefixes usually has good reason to believe that a) these problems will be corrected over time, and b) everyone else in the same /8 will have the same problem.

In the new paradigm, each prefix must be carefully examined by the recipient to test that it is uniquely held by the proffering organization, and the recipient will presumably have a further interest in its routability and membership in blacklists. The quality problem arises in both private and public trades; if the RIRs implemented a quality test, that would be yet another advantage of centralization to the benefit of everyone.

Closely associated with prefix quality is the question of *safety*. Again the RIRs are in an excellent position to provide the necessary support for good-faith transactions, certification of prefixes being the primary mechanism, although various other possibilities (such as membership controls) might also exist.

More pertinently, pricing of the goods traded in such an exchange is an important question. Various natural calculations might support the calculation of address costs, including but not limited to average revenue per address, operational costs averaged over all addresses held, and so on. Our primary contention here is that the RIRs should not engage in price setting directly. Doing so would at the very least invite regulation. There may be a case for placing caps on trades as an antispeculation measure, but that requires further analysis.

What exactly the "goods" are in this case also needs consideration. Our preference is that what is traded is the right to use a prefix, rather than a prefix itself. Quite apart from the inherent oddness in selling a 32-bit integer (with 5-bit netmask), we should avoid the land registry model, where all the previous history of a prefix must be checked before sale. We need the RIR to intermediate itself and provide quality evaluation services rather than leaving it up to the end buyer. We should also not be selling rights to use prefixes of fixed sizes. The exchange needs to offer a spread of lengths in order to meet the needs of all potential customers.

You Say You Want a Revolution

To be sure, a change in the perceptual or legal status of IP addresses is a revolution in how we do things. The ramifications of IP addresses becoming property, or even acquiring intermediate states with property-like title rights, are manifold and they involve sweeping changes. Suddenly things that had no value have a clear public worth. Will organizations then be compelled to list addresses on their books as an asset? Could they then be taxed on them? What would such a tax rate be? Could organizations not actually using the asset (say, the RIRs) avoid this charge? Would transfers entail a taxable operation? These questions are significant and difficult. The right thing for the community is almost undoubtedly that IP addresses do not become simple property, but rather have (at a minimum) transfer and sale rights associated with them. In this way we could enable liquidity without complications, and avoid introducing extra complications at a difficult time. But it is unclear whether regulatory authorities will see it this way without the correct guidance.

The change in legal status of IP addresses is not the only violent change that could be unleashed by exhaustion. Consider, for example, the potential for litigation led by both new entrants unable to acquire an allocation to fulfill their business plan and incumbents seeking to either cause confusion (as an anticompetitive measure against just about anyone) or to try to disrupt any fragile consensus about how the last allocations play out. Leaving aside the question of whether simple prudence would recommend or deprecate such a move, there is a very clear risk of attempted litigation affecting the outcome of the end game.

However, one of the major benefits of a market is that it allows the RIRs to maintain a hands-off approach while still making it at least theoretically possible for an organization to get an independent allocation. The community can be doing all that it realistically can to continue the flow of IPv4, in terms of creating conditions fostering its dissemination, while being seen to be doing such, rather than simply running out of ideas and giving up. It could, of course, be seen—not unfairly—that participating in the transition to a market mechanism might amount to the effective transference of title to those who happened to be in the room at the time of exhaustion, an effective "insider privatization."

Yet, if a market does not emerge, it is hard to see how any new entrants can have a business plan not directly dependent on incumbents. Although there are plenty of incumbents who would value having more address space to continue their business over the cash value of their addresses, so rendering entrance to the market impossible, there are plenty of other organizations that have only ever used a portion of their first allocation and would theoretically be well motivated to disburse these addresses accordingly.

To avoid exceptional attention from regulatory authorities, and to prevent the exchange from failing, we should design the exchange to deter in a systematic way the misbehavior of markets: speculation, hoarding, cartels, price fixing, and regional disadvantage should all be made as difficult as possible within the context of running a limited-membership market.

If we define *speculation* as short-term dealing with no expectation of use, we may be able to limit this kind of behavior naturally as a consequence of the membership-based participation inherent in the RIR model, and as a function of the periodic nature of routing filter generation. Increasing the price with short-term speculation disincentivizes the end purchaser with a use expectation from actually buying the prefix, because there will be a time delay before it can be used; therefore the purchaser with no use expectation will find it more difficult to find a buyer if the price rises to unreasonably high levels.

Hoarding, defined as long-term speculation with no use expectation, is bad for the exchange in that thickness is reduced, but also bad for the hoarder because the long-term value of the asset should decrease, in line with the increase in deployment of IPv6.

The formation of *cartels* would actually be quite a practical difficulty, especially under the closer attention likely to be paid to the exchange by competition authorities. Notwithstanding the coordination difficulties, we are inclined to say again that enough buyers should help to control this problem sufficiently to make the exchange work.

Regional disadvantage is, however we look at this situation, a problem. If scarcity is likely to lead to some monetary value being placed on address space, we face a vista where regional disadvantage can only be reduced, not eliminated. The inequality is, ultimately, one of the most compelling reasons to minimize the length of the transition period, and it would benefit us all to do so. Some measures go part way toward alleviating the problem. For instance, regional cooperation can help—in a market, if buyers cooperate and bulk buy, the threshold for organizations that would otherwise be facing a prohibitive barrier to entry would be reduced.

If we do not have a globally accessible exchange, it does not necessarily mean that the organizations will simply fail, entrenching the regional inequality, but they may respond by trying to fulfill their customer requirements by means of private, uncoordinated trading, with all the problems that entails.

We note that it is probably best to structure the actual trades as *auctions,* rather than facilitated marketplace transactions. When quality is asserted, one prefix is much like another—at least compared to prefixes of a similar size—and treating them as a commodity in this way facilitates the enforcement of policies on a centralized basis.

Drawbacks of a Market

Many cautionary tales about the operation of markets exist. Irrational exuberance, long-lasting depressions, fraudulent or exploitative behavior of all kinds—all of these effects, either enabled or supported by market mechanisms, are well known. Do we have any reason to believe either that these consequences will be not serious in our particular domain or that we have any new way of preventing them from happening?

In truth, we have no particular reason to believe that they won't happen, but there is a structural reason to believe that they might not matter to the exclusion of all else: the worse the situation becomes in the IPv4 marketplace, the more incentive there is to move to IPv6. To that extent, the market might be considered as providing a somewhat self-regulating reason for transition. Of course, we can put various mechanisms in place to help mitigate unstable behavior, as we suggested previously, but ultimately this is a fundamentally new way of doing things that we are ill equipped to understand the full consequences of.

Perhaps the largest drawback, outside of the practical difficulties in getting IPv4 addresses to organizations, is the philosophical impediments that come inherent with switching to a market-based model for allocation. Although a market cannot be said to rule out the consensus model that has turned out well for the Internet community, it also cannot be said to fully support it. This change may be a cultural one we find difficult to reverse, and it might undermine any future attempt by the community to try to differentiate itself on governance model.

Even though we have proposed the market model in good faith, as an attempt to meet the needs of new entrants and existing organizations—and as a boost to the faster deployment of IPv6—if it proves to be a failure in meeting those needs, there may be no more credible strategies left if governments insist on action. That in itself might represent even larger, more unpredictable change for the industry.

Effects on the Routing Table

Another inescapably important question is what will happen to the *Default Free Zone* (DFZ) routing table. A world in which address blocks transfer without the aggregating procedures of the RIRs is naturally a cause for concern, and when needs-based allocation comes to an end, a change in the rate of growth does seem inevitable. We can, however, make some observations that might reassure us, to some extent, that the rate of growth will not be calamitous.

First, as we go from a time of address plenty to address scarcity, one can assume that the ongoing fulfilled demand for address space will be no greater than it is now. Hence, the future growth in the number of prefixes in the routing table—regardless of prefix length—would seem to have an upper limit consistent with the number of allocations by RIRs to *Local Internet Registries* (LIRs) at the moment. This limit is still a multiple of the current curve, because we lose the benefit of the aggregation function performed by LIRs, but it suggests that we will at least not face an order-of-magnitude step change as a result of a disorderly competition.

Then there is the question of the routability of smaller prefixes. There is, at the moment, a *de facto* longest prefix size of around /24 that has close to universal reachability on the general Internet. One might assume that this prefix size will grow inexorably during and after exhaustion, as existing space is broken up into smaller and smaller blocks. Implicit in that assumption is the notion that such block sizes will be adequate for users and worthwhile for ISPs to route; we should probably not rely on networks "making do" with smaller and smaller chunks of address space.

Simultaneously, inexorably growing prefix lengths in the DFZ can only come about because of operator action. In particular, although there is a rough consensus in DFZ operators at the moment that /24 is routable and /25 is not, this policy is not a consensus-approved policy of the RIRs or the IETF. Each operator makes its own decision, based on its own customer needs, its own network, and the expectation of routability with other networks.

Reachability, therefore, depends on ISPs cooperating, and universal reachability depends on ISPs cooperating universally. An ISP may well choose to carry smaller prefixes on behalf of its customers, but unless this practice becomes widespread, no expectation can be made of universal reachability, and the practice will remain a minority one conducted by cooperating ISPs, as occasionally happens from time to time today, and this situation will little affect the size of the routing table for those involved.

Is there a competitive advantage to the largest of the ISPs in investing in very large routers that can carry many millions of prefixes, more than the smaller ISPs can support? If there were, it could perhaps lead to a concentration of power in the tier-one providers (who, as inevitable parts of any lengthy path across the Internet, have the greatest influence on the *de facto* longest routable prefix.) This situation could perhaps be true if routers are price-limited by the supportable number of prefixes, but this characteristic is typically a secondary one at worst. Routers are grouped by the bandwidth they can support, and priced accordingly; a 100-Mbps router that can support a million prefixes will certainly be more expensive than a 100-Mbps router that can support only ten thousand, but there is an order of magnitude step from either router to a router that can support 10 Gbps.

Inaction Leads to Harm

In fact the argument that the effect on the routing table will be unsustainable is opposed to the argument that there may not be adequate liquidity to sustain the market. It is true that we could find ourselves in the latter position, and so the effect of this system on reducing the problem (characterized as "the gap") will be smaller than we might like—but, as a best-effort scenario, not negligible, particularly in regard to showing good stewardship of the resource to potential outside influence. Compared to any other proposal, and particularly compared to voluntary release or a locking down of the address space, we think that this way is the best way to assure that we make available what liquidity there is.

It is difficult to see any model—even an idealized one—that could possibly service the run rate while maintaining aggregatability. The sparse allocation model used by the RIRs is dependent upon the continued availability of large, clean blocks of space, that is, /8s from IANA. With this address plenty comes freedom in our choice of policies, and with that freedom comes relatively quick consensus.

Post-exhaustion, the space will not be plentiful, and regardless of whether a monetary cost is attached, it will no longer be free. At this point, the legitimacy of the consensus of the RIR fora becomes critical. It is a fiercely defended bottom-up process. As the legitimacy of policies in the *Domain Name System* (DNS) world comes from consensus to abide by a single `root.cache`, so the legitimacy of policies in routing comes from general agreement on route filters and the authenticity of data in the RIR WHOIS databases.

We have also learned from the DNS world what happens to operational consensus when the resource becomes in some way valuable. Although the current RIR meetings are able to come to decisions that roughly reflect the consensus of the operational Internet, the necessarily tougher decisions forced upon us will challenge those who participate directly in policy making to reach conclusions that will satisfy operators who are not present. In principle it should not be necessary to account for those who do not represent themselves, but when the legitimacy of our policies is derived from their operational choices, the burden rests on us to ensure that our processes are truly representative.

If we are unsuccessful in doing so, or indeed if we choose to maintain the status quo, we cannot assume that the policies implemented on the operational Internet will themselves remain static. It is already the case that ISPs will work together, as is their entitlement, to agree to route prefixes for the benefit of their mutual customers. It is not unusual for one ISP to accept the announcement from a customer of a subnet of another ISP's address space. This decision is one for those ISPs to make about their own operational environments.

If we choose not to endorse a particular short-term solution to depletion, it falls upon ISPs themselves to find a way to continue their business operations, and resolve their customers' problems. If they cannot get address space from themselves, it will be their *duty* to their customers to get routable address space from somewhere—by negotiating, if necessary, with their peers and upstream providers to change the definition of "routable address space." Ultimately we may assume that if we do not provide a solution to the industry, the industry will invent one—or several competing ones.

Because we assert that the solution that best solves this problem is an address space trading exchange, we may well end up getting one—but one (or more) that is private, and out of sight of our existing policy-making structure. Worse still, competing exchanges would not have access to the RIRs data, and so would not be in a position to assure the quality of a prefix—a situation that could threaten all transactions.

Without exaggerating, it is likely that what we do in response to this crisis will determine the architecture of the Internet for a long while to come. Although we are reminded of Woody Allen's quote wherein he "... hope[s] mankind has the wisdom to choose correctly... between utter hopelessness and total extinction[22, 23]," there are, as we have outlined, measures we can take to survive the coming storm. They are not beautiful solutions. They are not how we have traditionally done things, or even how we would like to do things. Adopting them will almost certainly result in someone being worse off than if we had simply done nothing. But they represent, to our minds, the best, most realistic chance to avoid widespread difficulties and the loss of many of the principles we in the networking community hold dear, to ourselves and in our institutions. Let us begin this process now.

Acknowledgements

The authors would like to gratefully acknowledge help and support from Léan Ní Chuilleanáin, Emma Apted, and David Malone for diligent editing.

References

[0] Murphy, Niall and Wilson, David, "The End of Eternity Part One: IPv4 Address Exhaustion and Consequences," *The Internet Protocol Journal,* Volume 11, No. 4, December 2008.

[1] **ftp://ftp.ietf.org/ietf-online-proceedings/94dec/area.and.wg.reports/ipng/ale/ale-minutes-94dec.txt**

[2] **http://tools.ietf.org/html/rfc2008**

[3] Hain, Tony, "A Pragmatic Report on IPv4 Address Space Consumption," *The Internet Protocol Journal,* Volume 8, No. 3, September 2005.

[4] **http://playground.sun.com/ipv6/doc/history.html**

[5] **http://ipv4.potaroo.net**

[6] **http://www.ripe.net/ripe/meetings/ripe-55/presentations/murphy-simlir.pdf**

[7] **http://www.isoc.org/educpillar/resources/ipv6_faq.shtml**

[8] **http://www.ietf.org/internet-drafts/draft-narten-ipv6-statement-00.txt**

[9] **http://www.apnic.net/meetings/24/program/sigs/policy/presentations/el-nakhal-prop-051.pdf**

[10] **http://www.ripe.net/ripe/policies/proposals/2007-06.html**

[11] **http://www.switch.ch/pki/meetings/2007-01/namebased_ssl_virtualhosts.pdf**

[12] For example, **http://h.root-servers.org/128.63.2.53_2.html versus http://h.root-servers.org/h2_5.html**

[13] **http://www.ripe.net/ripe/meetings/ripe-55/presentations/vegoda-reclaiming-our.pdf**

[14] A "smooth and convenient" dialing plan for India. **http://www.mycoordinates.org/indias-phone-june-06**

[15] **http://en.wikipedia.org/wiki/UK_telephone_code_misconceptions**

[16] **http://code.google.com/p/simlir/**

[17] **http://www.ripe.net/docs/ripe-407.html#membership**

[18] **http://www.ripe.net/ripe/policies/proposals/2007-03.html**

[19] **http://www.ripe.net/ripe/policies/proposals/2007-06.html**

[20] **http://www.ripe.net/ripe/policies/proposals/2007-07.html**

[21] **http://kuznets.fas.harvard.edu/~aroth/alroth.html**

[22] Woody Allen, "Side Effects," 1980.

[23] Woody Allen through (most famously) Stephen Hawking, **http://www.cnn.com/2006/WORLD/asiapcf/07/04/talkasia.hawking.script/index.html**

[24] **http://icann.org/en/announcements/proposal-ipv4-report-29nov07.htm**

[25] **http://www.ripe.net/ttm/**

[26] **http://www.ripe.net/ripe/tf/enhanced-cooperation/index.html**

[27] **http://www.nro.net/documents/nro18.html**

[28] **http://www.ripe.net/maillists/ncc-archives/im-support/2004/index.html**

NIALL MURPHY holds a B.Sc. in Computer Science and Mathematics from University College Dublin. While in university, he founded the UCD Internet Society, which provided Internet access to approximately 5000 students. He went on to work for (and found) various organizations: the .IE domain registry, Club Internet (now Magnet Entertainment), Ireland On-Line, Enigma Consulting, Bitbuzz, and Amazon.com. He is currently in Site Reliability Engineering at Google. He is the coauthor of numerous articles, some RFCs, the O'Reilly book *IPv6 Network Administration,* and is a published poet and keen amateur landscape photographer. E-mail: **niallm@avernus.net**

DAVE WILSON holds a B.Sc. in Computer Science from University College Dublin, not coincidentally from around the same time as Niall. He has worked at HEAnet, the Irish National Research & Education Network, for more than 10 years, maintaining an involvement with RIPE and with the pan-European research network Géant. Dave is a member of the ICANN Address Supporting Organization Address Council; he helped to found the Irish IPv6 task force, which has the support of the national government there. E-mail: **dave.wilson@heanet.ie**

Resource Certification

by Geoff Huston, APNIC

Opinions vary as to what aspect of the Internet infrastructure represents the greatest common vulnerability to the security and safety of Internet users, but it is generally regarded that attacks that are directed at the network infrastructure are the most insidious, and in that case the choice is probably between the *Domain Name System* (DNS) and the interdomain routing system.

The question of how to improve the robustness of these functions has been a longstanding topic of study. For the DNS it appears that there is convergence on *Domain Name System Security Extensions* (DNSSEC) as the technical solution to securing DNS resolution operations, and the focus of attention in this space has shifted from technical behavior to topics relating to operational deployment. It has been a difficult time for DNSSEC and to say that there is an end in sight may well be premature at this stage, but there are definite signs of progress in this space. The same cannot be said of progress with securing routing, and particularly in securing interdomain routing. Here much remains to be done in order to achieve reasonable consensus on what technical measures to adopt, let alone the second step of study of how such measures could be deployed across the Internet.

The IETF's approach to addressing the topic of securing interdomain routing has followed a conventional IETF path. The first step has been to consider the nature of various vulnerabilities that exist within today's interdomain routing system and then develop a set of requirements that should be addressed in any solution space, without necessarily defining what such a solution may be. When the enumeration of requirements achieves a suitable level of consensus from the community, it is then possible to commence work on standardizing solutions. In the case of securing interdomain routing, the first steps were undertaken in *Birds of a Feather* (BOF) sessions and in the subsequently formed *Routing Protocol Security Requirements* (RPSEC) Working Group. This work is almost complete, and apart from some definitive statement relating to a requirement for securing the *Autonomous System* (AS) Path attribute in *Border Gateway Protocol* (BGP), the set of requirements for securing interdomain routing is now in an almost final state[1]. The task of the *Securing Inter-Domain Routing* (SIDR) Working Group is to standardize technologies that can meet these requirements.

So where does "Resource Certification" fit in?

Public Key Cryptography

One commonly used security technology is *Public Key Cryptography,* a technique that is easily explained. The approach uses a pair of keys, A and B. Anything enciphered with key A can be deciphered only with key B, and conversely, and knowledge of the value of one key does not lead to discovery of the value of the other key. Key A is kept as a closely guarded secret, whereas key B is openly published. If I want to send you a message that only you can decipher and read, I should encrypt it using your public key. If I want to send you a message that only I could have sent (nonrepudiation), then I will generate a digital signature of the message using my private key. That way any attempts to alter the message will also be detectable.

This latter approach, of using keys to generate digital signatures of messages, lies at the heart of DNSSEC, because DNSSEC adds public keys and digital signatures to the DNS. A DNS query can generate a response that lists both the DNS answer and the digital signature of that answer. The DNS can also be queried to retrieve the public key used to sign all the components of that zone, so that the digital signature can be verified and the query agent can be assured that the response is a genuine one. But how can the key itself be verified? In DNSSEC the hierarchical nature of the DNS itself is exploited by having each zone "parent" sign the keys of its delegated "children." So the zone key can be verified by retrieving the parent's signature across that zone key, and so on to the root of the DNS. As long as the query agent knows beforehand the value of the public key used to sign the root zone of the DNS, and as long as DNSSEC is used universally, all DNS responses can be verified in DNSSEC.

Although this approach works in the interlocked hierarchical structure of the DNS, when we turn our attention to securing the use of IP addresses and AS numbers in the context of interdomain routing, there is no comparable hierarchy to exploit. In such cases a common solution is to turn to *Digital Certificates*.

Digital Certificates are digitally signed public attestations by a certification authority that associate a subject's public key value with some attribute of the subject. A typical application is in identity certification, where the certification authority is attesting that the holder of the private key whose matching public key is provided in the certificate has met the authority's certification criteria to be identified by a particular name. Digital certificates are useful in that they can reduce the number of trust points in a security domain, so that each member of the domain does not have to validate identity and exchange public keys with every other member of the domain, but can undertake a single transaction with a certification authority that is trusted by all the members of the domain. As long as every member of the domain carries the public key of the certification authority and can access all issued digital certificates, then the members of the domain can verify each other's attestations and digital signatures.

Of course digital certificates are used for far more than attestations of identity, and can encompass the authority to perform specific tasks, undertake particular roles, or grant permissions and right-of-use authorities. It is this latter use case that is relevant to resource certification.

Resource Certificates

A Resource Certificate is a conventional X.509 certificate that conforms to the *Public Key Infrastructure Working Group* (PKIX) profile (RFC 5280) with one critical component, namely a certificate extension that lists a collection of IP number resources (IPv4 addresses, IPv6 addresses, and AS numbers)[17].

These certificates attest that the certificate issuer has granted to the entity represented by the certificate subject a unique "right-of-use" of the associated set of IP number resources listed in the certificate extension, by virtue of an associated resource allocation. The unique "right-of-use" concept mirrors the resource allocation framework, where the certificate provides a means of third-party validation of assertions related to resource allocations[2].

By coupling the issuance of a certificate by a parent *Certification Authority* (CA) to the corresponding resource allocation, a test of the validity of a certificate, including the IP number resource extension, can also be interpreted as validation of that resource allocation. Signing operations that descend from that certificate can therefore be held to be testable, under the corresponding hierarchy of allocation. In other words, if you received your address block from a particular *Regional Internet Registry* (RIR), then only that RIR can issue a Resource Certificate for you that includes your public key and the allocated number resources. Anything you sign using your private key can be verified through the RIR's issued certificate.

Unlike certificates that relate to attestations of identity, Resource Certificates are not necessarily long-lived. When an additional allocation action occurs, the associated Resource Certificate is reissued with an IP number resource extension that matches the new allocation state. In the case of a reduction in allocated resources, the previously issued certificates are explicitly revoked when the new certificate is issued. In other cases there is no explicit revocation of the older certificates.

The intention here is that any instrument signed by the subject's private key that relates to an assertion of resource control, whether it is a protocol message in a routing protocol or an administrative request to an *Internet Service Provider* (ISP) to route a prefix or as assertion of title over the "right-of-use" of a number resource, can be validated through the matching public key contained in the certificate and the IP number resource that is enumerated in this certificate. The Resource Certificate itself can be verified in the context of a Resource Certificate *Public Key Infrastructure* (PKI).

The Resource Certificate Public Key Infrastructure

The *Resource Certificate Public Key Infrastructure* (RPKI) describes the structure of the certification framework used by Resource Certificates. The intent of the RPKI is to construct a robust hierarchy of X.509 certificates that allows relying parties to validate assertions about IP addresses and AS numbers, and their use.

The structure of the RPKI as it relates to public use of IP number resources is designed to precisely mirror the structure of the distribution of addresses and ASs in the Internet, so a brief description of this distribution structure is appropriate. The *Internet Assigned Numbers Authority* (IANA) manages the central pool of number resources. The IANA publishes a registry of all current allocations. The IANA does not make direct allocations of number resources to end users or *Local Internet Registries* (LIRs), and instead allocates blocks of number resources to the RIRs. The RIRs perform the next level of distribution, allocating number resources to LIRs, *National Internet Registries* (NIRs), and end users. NIRs perform allocations to LIRs and end users, and LIRs allocate resources to end users (Figure 1).

Figure 1: Address Distribution Hierarchy for the Internet

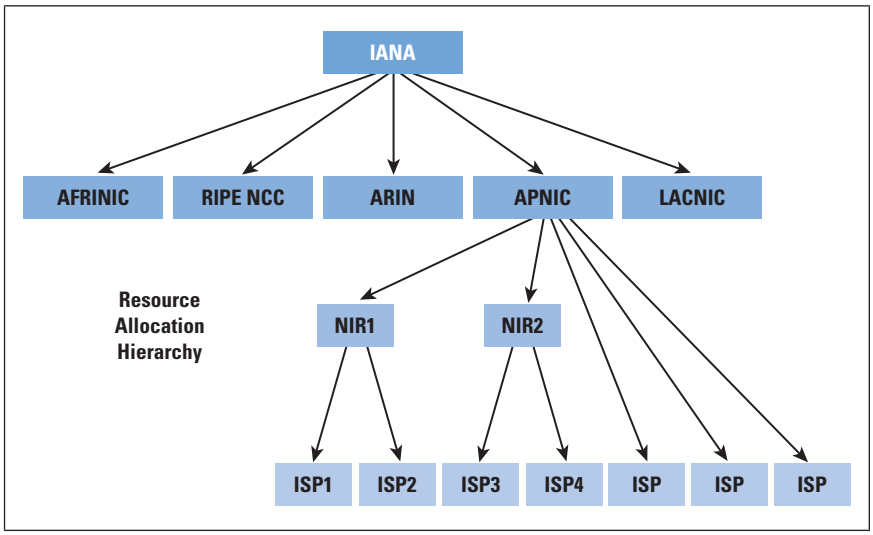

The RPKI mirrors this allocation hierarchy. One interpretation of this model would send the IANA manager a root RPKI key, and using this key the IANA would issue a self-signed "root" certificate, and also issue subordinate certificates to each of the RIRs, describing in the resource extension to the certificate the complete set of number resources that have been allocated to that RIR at the time of issuance. The certificate would also hold the public key of the RIR and would be signed by the private key of the IANA. Each RIR would issue certificates that correspond to allocations made by that RIR, where the resource extension to those certificates lists all the allocated resources, and the certificate includes the public key of the recipient of the resource allocation, signed with the private key of the RIR. If the recipient of the resource allocation is an LIR or an NIR, then it too would also similarly issue resources certificates (Figure 2).

Figure 2: RPKI Resource Certificate Hierarchy

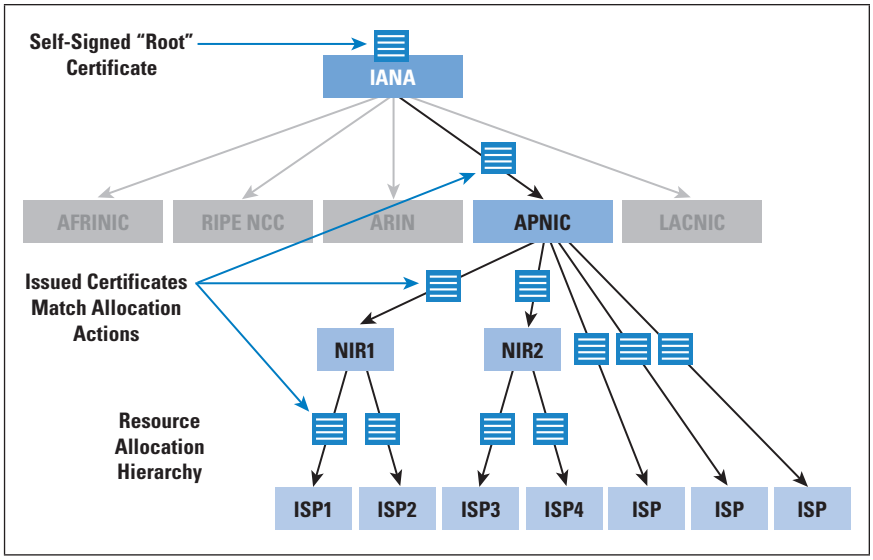

The common constraint within this certificate structure is that an issued certificate must contain a resource extension that contains a subset of the resources that are described in the resource extension of the issuing authority's certificate. This requirement corresponds to the allocation constraint than a registry cannot allocate resources that were not allocated to the registry in the first place. One implication of this constraint is that if any party holds resources allocated from two or more registries, then it will hold two or more Resource Certificates in order to describe the complete set of its resource holdings.

Validation of a certificate within this RPKI is similar to conventional certificate validation within any PKI, namely establishing a chain of valid certificates that are linked by issuer and subject from a nominated trust anchor CA to the certificate in question. The only additional constraints in the RPKI are that every certificate in this validation path must be a valid Resource Certificate, and the IP number of resources described in each certificate must be a subset of the resources described in the issuing authority's certificate.

Within this RPKI all Resource Certificates must have the IP addresses and AS resources present, and marked as critical extensions. The contents of these extensions correspond exactly to the current state of IP address and AS number allocations from the issuer to the subject.

Any holder of a resource who can make further allocations of resources to other parties must be able to issue Resource Certificates that correspond to these allocations. Similarly, any holder who wishes to use the RPKI to digitally sign an attestation needs to be able to issue an *End Entity* (EE) certificate to perform the digital signing operation.

For this reason all issued certificates that correspond to allocations are certificates with the CA capability enabled, and each CA certificate is capable of issuing subordinate CA certificates that correspond to further sub-allocations and subordinate EE certificates that correspond to a generation of digital signatures on attestations.

The RPKI makes conventional use of *Certificate Revocation Lists* (CRLs) to control the validity of issued certificates, and every CA certificate in the RPKI must issue a CRL according to the nominated CRL update cycle of the CA. A CA certificate may be revoked by an issuing authority for numerous reasons, including key rollover, the reduction in the resource set associated with the subject of the certificate, or termination of the resource allocation. To invalidate the authority or attestation that was signed by a given EE certificate, the CA issuing authority that issued the EE certificate simply revokes the EE certificate.

Resource Certificates are intended to be public documents, and all certificates and objects in the RPKI are published in openly accessible repositories. The set of all such repositories forms a complete information space, and it is fundamental to the model of securing the public Internet interdomain routing system that the entire RPKI information space is available. Other uses of the RPKI might permit use of subsets, such as the single chain from a given end-entity certificate to a trust anchor, but routing security is considered against all known publicly routable addresses and AS numbers, so all known resource certification outcomes must be available. In other words the intended use of the RPKI in routing contexts is not a case where each relying party may make specific requests for RPKI objects in order to validate a single object, but one where each relying party will perform a regular sweep across the entire set of RPKI objects in order to ensure that the relying party has a complete picture of the RPKI information space.

This aspect of the RPKI represents some interesting challenges, in that rather than having a single CA publish all the certificates produced in a security application at a single point, the RPKI permits the use of many publication points in a widely distributed fashion. Each CA can issue RPKI objects and publish them using a locally managed publication point. It is incumbent upon relying parties to synchronize a locally managed cache of the entire RPKI information space at regular and relatively frequent intervals.

For this reason the RPKI has introduced an additional mechanism in its publication framework, namely the use of a "manifest" to allow relying parties to determine whether they have been able to retrieve the entire set of RPKI published objects from each RPKI repository publication point, or if there has been some attempt to disrupt the relying party's access to the entire RPKI information set.

It also implies that the RPKI publication point access protocols should support the efficient function of a synchronization comparison, so that a locally managed cache of the RPKI need only call for the uploading of those objects that have been altered since the previous synchronization operation.

Signed Attestations and Authorities

The underlying intent of digital certificates, and Resource Certificates in particular, is in terms of supporting a transitive trust relationship that allows a relying party to verify the authenticity of a signed artefact through verification of the signer's key using the PKI. So the obvious question is: what artefacts are useful to sign?

Much of the motivation for Resource Certificates has come from a desire to underpin efforts in securing aspects of interdomain routing. This effort goes well beyond securing the individual point-to-point connection used between BGP speakers, and refers to the matter of verifying the authenticity of the payload of the BGP protocol exchange. The specific question that may be posed is: how can a BGP speaker validate the authenticity of the route object being presented to it?

The approach being studied by the SIDR Working Group is to use structured attestations, where, like the digital certificate itself, the attestation is structured in an ASN.1 digital object, and this object is signed using a signing formation that is itself a piece of structured ASN.1, namely the *Cryptographic Message Syntax* (CMS)[18].

The first of these attestations relates to the ability to verify the authenticity of the "origination" of an interdomain routing object. This verification refers to the address prefix and the originating AS, and the questions that this verification function is intended to answer include:

- Is this a valid address prefix and AS number? Have these resources been allocated through the IP number resource allocation process?
- Has the holder of the title of "right-of-use" for the address prefix authorized the AS holder to originate a routing advertisement for this prefix?

Here an address holder is authorizing a particular ISP to generate a route announcement for its particular address prefix. In this case the prefix holder would generate an EE Resource Certificate with the IP number resource extension spanning the set of addresses that match the address prefixes that are the intended subject of the routing authority, and place validity dates in the EE certificate that correspond to the intended validity dates of the routing authority.

The signed authority document would contain the AS number that is being authorized in this manner, a description of the range of prefixes that the prefix holder has authorized, and the EE certificate. The document would be signed by the EE certificate private key using a CMS signing structure. The resultant object is published in the RPKI distributed publication repository as a *Routing Origination Authorization* (ROA). A relying party can validate the ROA by checking to ensure that the digital signature in the ROA is correct, indicating that the authority document has not been tampered with in any way since it was signed, that the resources in the associated EE certificate encompass the prefixes specified in the document, and the EE certificate itself is valid in the context of the RPKI by verifying that there is an issuer-subject chain of valid certificates that link one of the relying party's nominated trust anchors to the EE certificate.

The ROA itself is valid as long as the signing EE certificate is valid. To withdraw the authority prior to the expiration of the EE certificate, the ROA publisher can simply revoke the EE certificate, leading to the concept of "one-off-use" EE certificates in the RPKI, where a key pair and a corresponding EE certificate are generated in order to sign a single attestation or authority. If the authority's lifetime is extended, the authority is reissued with a new EE certificate and a new digital signature, and, as noted, the authority can be prematurely terminated through revocation of the EE certificate, so at no stage is there a need to reuse the original signing private key. After the private key is used to sign this object, the key is destroyed, alleviating to some extent the key management load.

In any security system knowledge of what is authorized is helpful, but knowledge of what has not been authorized is perhaps even more helpful. For ROAs there is an analogous situation to DNSSEC, where DNSSEC is most effective from a client's perspective after the entire DNS space is DNSSEC signed. Where there are gaps in the DNSSEC signing chains the client is left in an uncertain state regarding the verification outcomes of the unlinked DNS sub-hierarchies. The same could apply to ROAs, in that in an environment where not every originated route object has a published ROA, the absence of a ROA does not necessarily indicate an unauthorized route origination. If one of the objectives of this study is to define a framework that can unambiguously identify the unauthorized use of IP number resources in routing (route "hijacks") even in a world where ROAs are used in a piecemeal fashion, then one possible refinement to the ROA model is the introduction of a comparable negative authority, the *Bogon Origination Attestation* (BOA).

In this case the prefix holder generates a signed attestation, or BOA, in a similar manner to the ROA, but does not provide any originating AS. Instead the BOA refers to "all originating ASs," and has the semantic interpretation that any use in the routing space of this address prefix described in the BOA, or any more specific address prefix, should be regarded as unauthorized and the route should be discarded.

Although this process makes the detection of route hijacks more direct in a world of piecemeal use of ROAs, there is now the added complication of having both "positive" and "negative" authorities. The proposed resolution of this dilemma is to use a relative priority rule that ROAs take precedence over BOAs, so that if a valid ROA and a valid BOA both exist that describe the origination component of a route, then the route can be regarded as authorized.

It should be noted, however, that at this stage these concepts are "work in progress," and are part of the SIDR Working Group's agenda of study, and the working group has not as yet reached any consensus regarding the decision to advance these proposals onward along the Internet Standards Process.

Also on the near-term horizon for SIDR is examining approaches to secure the AS path in BGP updates. The RPSEC Working Group has explored two approaches in this space. One involves an incremental multiple signature technique that allows a receiver of a BGP update to verify that the AS path described in the update is matched by a sequence of interlocking AS digital signatures using the RPKI. At the same time that an AS adds its own AS to the AS path prior to further *External Border Gateway Protocol* (eBGP) propagation of the route update, the AS would digitally sign over an analogous sequence of AS signatures. This approach allows a receiver to perform a match of the AS sequence in the AS path with the AS number sequence identified in the AS signature block. A match here would indicate that the BGP update has indeed been sequentially passed along the sequence identified by the AS path. This approach was originally proposed in the *Secure BGP* (sBGP) design[21] and has attracted some comment related to the computation overhead associated with the application and validation of these AS path signature sequences. An alternative approach has been one that is described by RPSEC as being less rigorous, and refers to a "feasibility" check, which checks to ensure that each pair of ASs represented in the AS path has an associated verifiable assertion of inter-AS adjacency that is digitally signed by both ASs.

It should also be noted that this activity of addressing aspects of improving the robustness of interdomain routing has some previous context. In many parts of the Internet, some degree of routing integrity is managed through the use of *Internet Routing Registries* (IRRs) and the publication of routing policies through the use of *Routing Policy Specification Language* (RPSL) objects.

Although opinions vary as to the robustness of the security offered by the IRR approach, at the very least it can mitigate some weakness in the routing system through the use of a "second check" that can be used to filter the information that is being provided in a BGP feed.

The weaknesses in the IRR system tend to relate to the consistency, completeness, and authenticity of the IRR data, and in many cases the trust in the integrity of the data relies on the admission practices of the IRR itself, and individual data objects cannot be verified by clients of the IRR. One possible way to address this situation has been through the use of *Routing Policy System Security* (RPSS) measures, but the adoption of these measures has not been widespread, and the question still remains for the client that even if an IRR object was authenticated upon admission, it does not mean that when the object is subsequently used by an IRR client the information reflects the current situation, and the information could well be invalid or not reflect the current policies of the author of the IRR object.

One possible approach being considered by the SIDR Working Group is to implement the RPSS authentication models using object signing in the context of the RPKI. For example, the RPSS assumption that routes should be announced only with the consent of the holder of the origin AS number of the announcement and with the consent of the holder of the address space implies in RPSS that both parties should authorize the entry of a *route object* into the IRR. Translating this stipulation into an analogous model using the RPKI would require that a route object be signed with the digital signatures of both the AS holder and the address space holder, and a IRR client can verify this route object at the time of use by verifying both digital signatures. Either the address space holder or the AS holder can revoke authorization by revoking the EE certificate used to sign the route object, and the verification is independent of the particular IRR that has published the route object. It is also a possibility that the IRR itself can be folded into the RPKI distributed publication repository framework, because there is no particular requirement in such an environment for a disparate collection of IRRs with their own partial collections of routing policy information, although at this stage this discussion is heading into the realm of more advanced speculation about the potential for application of Resource Certificates and digital signatures to RPSL and the IRR framework.

Putting Resource Certificates into Context

Resource Certificates and the associated RPKI represent a major part of any effort to construct a secure interdomain routing framework. An RPKI, even partially populated with signed information, allows BGP speakers to make preferential selections to use routing information where the IP address block and the AS numbers being used are recognized as valid to use, and the parties using these IP addresses and AS numbers are properly authorized to so do. The RPKI can also be used to identify instances of unauthorized use of IP addresses and attempts to hijack routes.

However, the RPKI represents only one part of a larger framework of securing interdomain routing, and the next step is that of applying the RPKI to the local BGP processing framework. There is also the need to move beyond validation of route origination and look at the associated topic of validation of the AS path, and potentially to consider the most challenging task, of attempting to validate whether the initial forwarding decision associated with a route object actually represents the correct first hop along a usable forwarding path for packets to reach the network destination.

The concerns here include not only a consideration of what can be secured and validated, but matters of scalability and efficiency in terms of deployment cost. The various approaches to routing security studied so far offer a wide variety of outcomes in terms of the amount of routing information that is validated, the level of trust that can be placed in a validation outcome, and the overheads of generating and validating digital signatures on routing information. The next step appears to include the task of establishing an appropriate balance between the overheads of operating the security framework and the extent to which efforts to disrupt the routing system can be successfully deflected by such measures.

The RPKI has been designed as a robust, simple framework. As far as possible existing technologies and processes have been exploited, reflecting to some extent a level of conservatism of the routing community and the difficulty in securing widespread acceptance of novel technologies.

References and Further Reading

The following documents provide further detail about the IETF work on resource certification. The Internet Drafts listed here are still a "work in progress," and although they are reflective of the areas of activity of the SIDR Working Group, they do not necessarily represent finished work.

Internet Drafts

Requirements:

[1] B. Christian, T. Tauber, eds., "BGP Security Requirements," work in progress, Internet Draft, **`draft-ietf-rpsec-10.txt`**, November 2008. *The report of the consensus outcomes of the RPSEC Working Group in enumerating the requirements for securing interdomain routing. The outstanding topic in this report remains in the area of AS path validation and the level of requirement associated with the two approaches described in the report.*

Architecture:

[2] M. Lepinski, S. Kent, "An Infrastructure to Support Secure Internet Routing," work in progress, Internet Draft, **`draft-ietf-sidr-arch-04.txt`**, November 2008. *An overview of the RPKI approach, describing the RPKI, the distributed repository structure, and common operations.*

Resource Certificates:

[3] G. Huston, G. Michaelson, R. Loomans, "A Profile for X.509 PKIX Resource Certificates," work in progress, Internet Draft, **`draft-ietf-sidr-res-certs-15.txt`**, November 2008. *The specification of the Resource Certificate.*

RPKI Repository Structure:

[4] G. Huston, G. Michaelson, R. Loomans, "A Profile for Resource Certificate Repository Structure," work in progress, Internet Draft, **`draft-ietf-sidr-repos-struct-01.txt`**, October 2008. *A description of the proposed distributed publication repository structure for the RPKI, including contents, access protocols, and object name conventions.*

[5] R. Austein et al., "Manifests for the Resource Public Key Infrastructure," work in progress, Internet Draft, **`draft-ietf-sidr-rpki-manifests-04.txt`**, October 2008. *A specification for repository manifests. Manifests are signed constructs that describe all the objects currently loaded into a repository publication point, and are used by relying parties as a means of ensuring that a local RPKI repository cache is correctly synchronized against the authoritative original publication point.*

[6] G. Huston, R. Loomans, B. Ellacot, R. Austein, "A Protocol for Provisioning Resource Certificates," work in progress, Internet Draft, **`draft-ietf-sidr-rescerts-provisioning-03.txt`**, August 2008. *A proposed protocol for use between a subject and a certificate issuer to ensure that certificate requests, the IP number resource allocation state, and the issued certificate status are correctly synchronized. This synchronization extends the conventional certificate request model into a transaction protocol that also includes the ability to perform certificate revocation requests and status queries from the subject.*

RPKI Signed Objects:

[7] M. Lepinski, S. Kent, D. Kong, "A Profile for Route Origin Authorizations (ROAs)," work in progress, Internet Draft, **`draft-ietf-sidr-roa-format-04.txt`**, November 2008. *The specification of the syntax for signed ROAs.*

[8] G. Huston, T. Manderson, G. Michaelson, "A Profile for Bogon Origin Attestations (BOAs),"work in progress, Internet Draft, **`draft-ietf-sidr-bogons-02.txt`**, October 2008. *The specification of the syntax for signed BOAs.*

[9] G. Huston, G. Michaelson, "Validation of Route Origination in BGP Using the Resource Certificate PKI," work in progress, Internet Draft, **`draft-ietf-sidr-roa-validation-01.txt`**, October 2008. *The specification of the semantics of ROAs and BOAs and the manner in which these objects may be interpreted in terms of the integration of these origination security credentials onto a BGP route-selection process.*

Certificate Policy and Practice Statements:

[10] K. Seo, R. Watro, D. Kong, S. Kent, "Certificate Policy (CP) for the Resource PKI (RPKI)," work in progress, Internet Draft, **draft-ietf-sidr-cp-04.txt**, November 2008. *A description of the certificate policy that applies to all certificates issued within the RPKI framework.*

[11] D. Kong, K. Seo, S. Kent, "Template for an Internet Registry's Certification Practice Statement (CPS) for the Resource PKI (RPKI)," work in progress, Internet Draft, **draft-ietf-sidr-cps-irs-04.txt**, November 2008. *A template for the Practice Statement used by Internet Registries (IRs) to describe their operational practices in the issuance and management of Resource Certificates.*

[12] D. Kong, K. Seo, S. Kent, "Template for an Internet Service Provider's Certification Practice Statement (CPS) for the Resource PKI (RPKI)," work in progress, Internet Draft, **draft-ietf-sidr-cps-isp-03.txt**, November 2008. *A template for the Practice Statement used by ISPs to describe their operational practices in the issuance and management of Resource Certificates.*

Individual Submissions:

[13] G, Huston, G. Michaelson, "A Profile for AS Adjacency Attestation Objects," work in progress, Internet Draft, **draft-huston-sidr-aao-profile-00.txt**, September 2008. *The specification of the syntax for a pairwise inter-AS routing adjacency attestation.*

[14] R. Kisteleki, J. Boumans, "Securing RPSL Objects with RPKI Signatures," work in progress, Internet Draft, **draft-kisteleki-sidr-rpsl-sig-00.txt**, October 2008. *The specification of the addition of RPKI digital signatures to RPSL Objects in the context of an Internet Route Registry.*

[15] T. Manderson, G. Michaelson, "RPKI Repository Retrieval Mechanism," work in progress, Internet Draft, **draft-manderson-sidr-fetch-00**, October 2008. *A proposed mechanism to use the manifest as the basis for performing a synchronization operation between a local RPKI cache and a source point.*

RFCs:

[16] D. Cooper et al., "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," RFC 5280, May 2008.

[17] C. Lynn, S. Kent and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers," RFC 3779, June 2004.

[18] R. Housley, "Cryptographic Message Syntax (CMS)," RFC 3852, July 2004.

[19] C. Alaettinoglu, et al., "Routing Policy Specification Language (RPSL)," RFC 2622, June 1999.

[20] C. Villamizar et al., "Routing Policy System Security," RFC 2725, December 1999.

Other Documents:

[21] Kent, S., "Securing BGP: S-BGP," *The Internet Protocol Journal,* Volume 6, No. 3, September 2003.

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. The author of numerous Internet-related books, he is currently the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of the Internet Society from 1992 until 2001. E-mail: **gih@apnic.net**

Host Identity Protocol: Identifier/Locator Split for Host Mobility and Multihoming

by Andrei Gurtov and Miika Komu, Helsinki Institute for Information Technology, and Robert Moskowitz, ICSAlab

A host and its location are identified using *Internet Protocol* (IP) addresses in the current Internet architecture. However, IP addresses can serve only as short-term identifiers because a considerable amount of hosts are *portable* devices and they change their IP addresses when moved from one network to another. Short-term identifiers disrupt long-term transport layer connections, such as Internet phone calls, and make locating the peer host more difficult. Therefore, mobility and multihoming are hard to implement securely in the present Internet. Upon changing an IP address, the host must prove to its peers that it is the same entity they communicated with before, requiring the use of cryptographic identities.

Another challenge the Internet faces is due to the fact that deployed protocols in the Internet are prone to *Denial-of-Service* (DoS) attacks. Substantial memory state can be created before the communicating peer is authenticated. Impersonation attacks are possible because IP addresses are relatively easy to forge. Because of difficulties in configuring *IP Security* (IPsec) for users, most Internet traffic is still transmitted in plaintext, making it easy for attackers to collect passwords or lists of visited websites, for example, in public *Wireless Local-Area Networks* (WLANs). As the IPv6 protocol is seeing gradual deployment, interoperating traditional IPv4 applications with new IPv6 applications remain a challenge.

The so-called *identifier/locator split* is recognized by the *Internet Engineering Task Force* (IETF) community as a next big change in the Internet architecture. Although the problem has been known for a long time[17], it has only recently started to get sufficient attention. Developments in public key cryptography and increased computational resources of hosts enables the use of cryptographic mechanisms to securely handle identities. Several proposals are under consideration in the IETF, including the *Locator Identifier Separation Protocol* (LISP)[16] for the network-based and the *Host Identity Protocol* (HIP) for the host-based approach. LISP focuses on improving scalability of the routing system, whereas HIP provides secure end-to-end mobility and multihoming. Therefore, the two proposals are complementary rather than competing.

HIP Architecture

The HIP architecture[1,2] uses the identity/locator split advantage to address Internet architecture challenges in an integrated approach. HIP was proposed by Bob Moskowitz in 1999 and since then has been under active development in the IETF Working Group and *Internet Research Task Force* (IRTF) Research Group.

HIP enables host mobility and multihoming across different address families (IPv4 and IPv6), offers end-to-end encryption and protection against certain DoS attacks, allows moving away from IP address-based access control to permanent host identities, and restores end-to-end host identification in the presence of several addressing domains separated by *Network Address Translation* (NAT) devices.

HIP separates the identity of a host from its location. The location of the host is bound to IP addresses and used for routing packets to the host in the same way as in the current Internet architecture. However, transport and application layers use *host identity,* consisting of the public key component of a private-public key pair. Each host is responsible for creating one or more public/private key pairs to provide identities for itself. Because the host identities are based on public key cryptography, they are computationally difficult to forge. Host identities are location-independent identifiers that allow a mobile host to preserve its transport layer connections upon movement. On the other hand, the host identity can be used for looking up the current location of a host because the host identity is a long-term identifier. A client host obtains the host identity of a server typically from the *Domain Name System* (DNS)[7] or a *Distributed Hash Table* (DHT). However, the infrastructure may not support this DHT in certain scenarios, such as in peer-to-peer and temporary environments. In such cases, *opportunistic* HIP can be used for contacting a peer without prior information of the identity of the peer. Opportunistic HIP is based on a "leap-of-faith," meaning that it is prone to man-in-the-middle attacks for the initial connection. It is similar to the *Secure Shell* (SSH) *Protocol,* where the public key of the server is added to the known host list after the first connection.

The problem of certifying the keys in *Public Key Infrastructure* (PKI) or otherwise creating trust relationships between hosts has explicitly been left out of the HIP architecture, because it is expected that each system using HIP may want to address it differently. For mere mobility and multihoming, the systems can work without any explicit trust management, in an opportunistic manner.

All other parties use the host identifier, that is, the *public key,* to identify and authenticate the host. Typically, a host identifier is a 128-bit-long bit string, the *Host Identity Tag* (HIT), as shown in Figure 1. A HIT is constructed by applying a cryptographic hash function over the public key. The introduction of new endpoint identifiers changes the role of IP addresses. When HIP is used, IP addresses become pure topological labels, naming locations in the Internet. One benefit of this identity/locator separation is that hosts in private address realms (behind NATs) can name each other in a unique way with HITs. A second benefit is that the hosts can change their IP address without breaking transport layer connections of applications and rely on HIP to manage host mobility; the relationship between location names and identifiers becomes dynamic.

To start communicating through HIP, two hosts must establish a HIP association. Known as the HIP *Base Exchange* (BEX)[3], this process consists of four messages (I1, R1, I2, and R2) transferred between the initiator and the responder. After BEX is successfully completed, both hosts are confident that private keys corresponding to host identifiers (public keys) are indeed possessed by their peers. Another purpose of the HIP base exchange is to create a pair of IPsec *Encapsulated Security Payload* (ESP) *Security Associations* (SAs), one for each direction. HIP uses IPsec ESP *Bound End-to-End Mode* (BEET)[4,9] to provide data encryption and integrity protection for network applications.

Figure 1: HIP Architecture

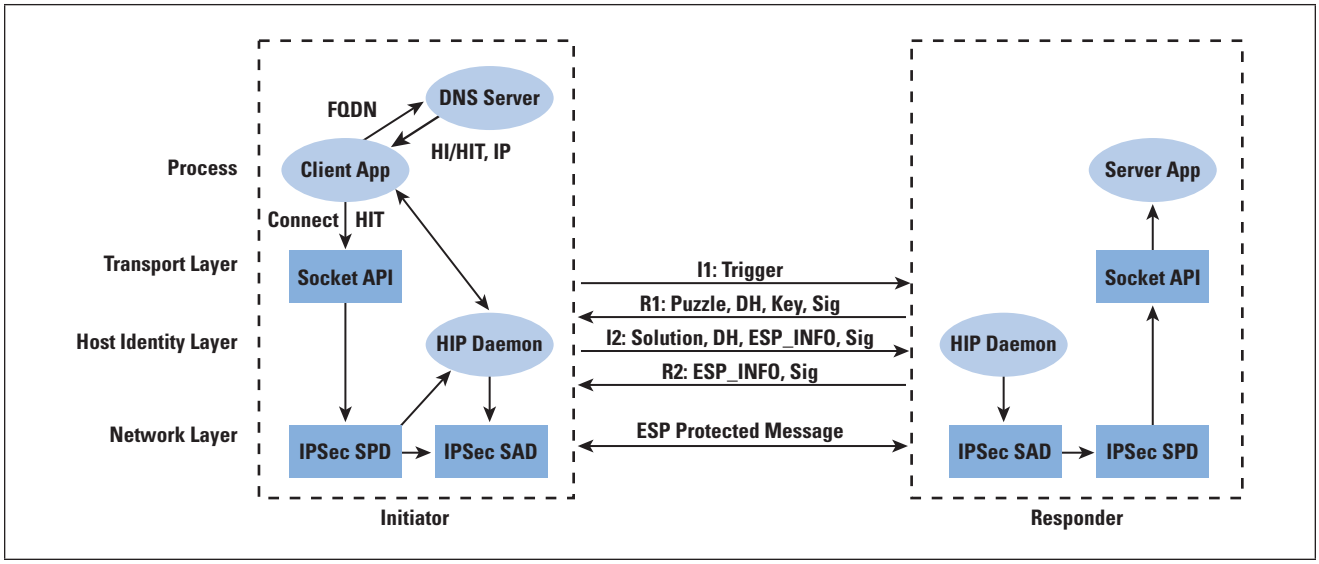

Because neither transport layer connections nor security associations created after the HIP base exchange are bound to IP addresses, a mobile client can change its IP address (that is, upon moving, because of a *Dynamic Host Configuration Protocol* [DHCP] lease or IPv6 router advertisement) and continue to transmit ESP-protected packets to its peer. HIP supports such mobility events by implementing an end-to-end three-way UPDATE signaling mechanism[8] between communicating nodes. HIP multihoming uses the same mechanisms as mobility for updating the peer with a current set of host IP addresses.

A rendezvous server[6] provides a mechanism to locate a host, for example, when two communicating hosts move simultaneously. To employ a rendezvous mechanism, a host first must perform a registration procedure[5], which is an extended version of the HIP base exchange.

The HIP control packets as well as ESP-encapsulated data packets have difficulties in going through NAT applications and firewalls. To traverse NAT, HIP uses *User Datagram Protocol* (UDP)-based encapsulation provided by the *Interactive Connectivity Establishment* (ICE) protocol.

It enables two hosts located behind NAT to communicate through a Rendezvous server. Bob Moskowitz suggests an alternative approach, where HIP always uses IPv6 for end-to-end communication and the *Teredo* protocol is employed to traverse NAT instances in IPv4 networks if native IPv6 connectivity is not available.

Most Internet applications can run unmodified over HIP[10], although only HIP-aware (new) applications using the extended socket interface can take better advantage of the new features that HIP provides. As HIP secures application data traffic with IPsec that is located logically "deep" within the networking stack, the challenge is to provide proper and understandable security indicators to the user to convince the user that the connection, for example, to a banking website, is secured. Such indicators can be developed as extensions to applications (for example, a security plug-in to the *Firefox* browser) or within a hostwide HIP management utility that controls all applications.

HIP provides a network layer alternative to using *Secure Sockets Layer/Transport Layer Security* (SSL/TLS) for application security, which has its benefits and drawbacks. HIP is a generic solution that should work for any transport protocol, whereas until recently TLS supported only TCP. HIP enables host mobility and multihoming, which is not supported by TLS. TLS runs on top of TCP, leaving it vulnerable to various TCP attacks; for example, using spoofed *reset* (RST) packets or DoS attacks with SYNs. Applications must be designed explicitly to use TLS, whereas HIP can provide security as an add-on to existing traditional applications. On the other hand, TLS does not have a problem with traversing traditional middleboxes such as NATs and firewalls that need special attention for HIP. Both protocols share the characteristic of endorsing host identity. TLS relies on certificates issued by one of the known Certification Authorities, whereas HIP can use *Domain Name System Security Extensions* (DNSSEC)[18] or a PKI infrastructure.

There are currently three open-source interoperating HIP implementations. *OpenHIP* from Boeing runs on Linux, Windows, and Mac OS, whereas *HIP on Linux* (HIPL) runs on Linux and Symbian, and *HIP for Inter.net* from Ericsson runs on FreeBSD and Linux. Several testbeds are deployed based on HIP, including the Everett Boeing factory[11], the P2PSIP pilot in Finland[14], and Wi-Fi P2P Internet Sharing Architecture in Germany[12]. Ericsson NomadicLab and TeliaSonera have demonstrated using HIP for transparent IPv4 and IPv6 handovers, mobile router, simultaneous multiaccess, and the use of proxy for traditional hosts[13,15].

Acknowledgements

We are grateful to Pekka Nikander, Tom Henderson, and others in the IETF and the *Internet Research Task Force* (IRTF) community who were encouraging and contributing to the development of HIP. We thank Andrey Khurri for the figure on HIP architecture and Henry Sinnreich for encouraging us to write this article.

We also thank members of InfraHIP II project for comments helping to improve this article.

References

[1] Moskowitz, R. and Nikander, P., "Host Identity Protocol Architecture," RFC 4423, May 2006.

[2] Gurtov, A., *Host Identity Protocol (HIP): Towards the Secure Mobile Internet,* ISBN 978-0-470-99790-1, Wiley and Sons, June 2008.

[3] Moskowitz, R., Nikander, P., Jokela, P. and Henderson, T., "Host Identity Protocol," RFC 5201, April 2008.

[4] Jokela, P., Moskowitz, R. and Nikander, P., "Using the Encapsulating Security Payload (ESP) Transport Format with the Host Identity Protocol (HIP)," RFC 5202, April 2008.

[5] Laganier, J., Koponen, T. and Eggert, L., "Host Identity Protocol (HIP) Registration Extension," RFC 5203, April 2008.

[6] Laganier, J. and Eggert, L., "Host Identity Protocol (HIP) Rendezvous Extension," RFC 5204, April 2008.

[7] Nikander, P. and Laganier, J., "Host Identity Protocol (HIP) Domain Name System (DNS) Extension," RFC 5205, April 2008.

[8] Nikander, P., Henderson, T., Vogt, C. and Arkko, J. "End-host Mobility and Multihoming with the Host Identity Protocol," RFC 5206, April 2008.

[9] Nikander, P. and Melen, J., "A Bound End-to-End Tunnel (BEET) Mode for ESP," Internet Draft, Work in Progress, **draft-nikander-esp-beet-mode-09**

[10] Henderson, T., Nikander, P. and Komu, M., "Using the Host Identity Protocol with Legacy Applications," RFC 5338, September 2008.

[11] Boeing, "Secure Mobile Architecture (SMA) for Automation Security," **http://www.isa.org/wsummit/presentations/Boeing-NGI_SMA_Automation_Security_Vancouver_ISA_presentationtemplates_7-23-07.ppt**

[12] Heer, T., Götz, S., Weingärtner, E. and Wehrle, K., "Secure Wi-Fi Sharing on Global Scales," in Proceedings of the 15th International Conference on Telecommunication (ICT), St. Petersburg, Russian Federation, IEEE, 2008. **https://www.ds-group.info/members/heer/publications-tobias-heer/pdfs/HeerEtAl2008.pdf**

[13] Jokela, P., Ylitalo, J., and Salmela, P., "HIP Mobile Router Demo," March 2007.
http://www.ietf.org/proceedings/07mar/slides/HIPRG-3.pdf

[14] Koskela, J., Heikkila, J. and Gurtov, A., "A Secure P2PSIP System with SPAM Prevention," Poster at ACM Mobicom, September 2008.

[15] Korhonen, J., Mäkelä, A., and Rinta-aho, T., "HIP Based Network Access Protocol in Operator Network Deployments," in First Ambient Networks Workshop on Mobility, Multiaccess, and Network Management (M2NM'07), Sydney, Australia, October 2007.

[16] Meyer, D., "The Locator Identifier Separation Protocol (LISP)," *The Internet Protocol Journal,* Volume 11, No. 1, March 2008.

[17] Saltzer J., "On The Naming and Binding of Network Destinations," RFC 1498, September 1992.

[18] Gieben, M., "DNSSEC: The Protocol, Deployment, and a Bit of Development," *The Internet Protocol Journal,* Volume 7, No. 2, June 2004.

[19] Sinnreich, H., "Letter to the Editor," *The Internet Protocol Journal,* Volume 11, No. 3, page 37, September 2008.

ANDREI GURTOV received M.Sc and Ph.D. degrees in Computer Science from the University of Helsinki, Finland. He presently is Principal Scientist, leading the Networking Research group at the Helsinki Institute for Information Technology, focusing on distributed system security and next-generation Internet architecture. He co-chairs the IRTF research group on HIP and teaches as an adjunct professor at Helsinki University of Technology. He is a regular visitor of the ICSI Center for Internet Research (ICIR) at Berkeley. Andrei has co-authored more than 50 publications, including a book, research papers, patents, and RFCs. He can be reached through the webpage: **http://www.hiit.fi/~gurtov**

MIIKA KOMU received his M.Sc. from Helsinki University of Technology and continues his studies as a postgraduate student. He is working as a full-time researcher and software engineer at Helsinki Institute for Information Technology. He is an active IETF participant and co-author of RFC 5338. Miika is an open source advocate and martial arts fan. E-mail: **miika.komu@hiit.fi**

ROBERT MOSKOWITZ is senior technical director for ICSA Labs and is an active member in the IAB, IETF, and IEEE. At ICSA Labs, Moskowitz leads the IPsec product and system certification program. Prior to the ICSA, he led the adoption of the world's largest IPsec network deployment servicing the automotive industry. As a former co-chair of the IPsec Working Group, Moskowitz provided a user set of multivendor, multipolicy, and multiuser requirements that galvanized many of the debates on the use of IPsec. A contributing editor for *Network Computing Magazine,* Moskowitz is currently helping define the new security component for the 802.11 standard. E-mail: **rgm@htt-consult.com**

Fragments

Allocation Policy for the Remaining IPv4 Address Space Ratified by ICANN

On 6 March 2009, the *International Corporation for Assigned Names and Numbers* (ICANN) Board ratified the *Global Policy for the Allocation of the Remaining IPv4 Address Space*. The policy requires ICANN to reserve one /8 for each *Regional Internet Registry* (RIR) from the *Internet Assigned Numbers Authority* (IANA) free pool. This has been done. The remainder of the implementation will be done once the IANA free pool has been fully allocated to RIRs. There are currently 32 unallocated unicast IPv4 /8s. 27 are in the IANA free pool and five are reserved under the Global Policy for the Allocation of the Remaining IPv4 Address Space.

On 4 February 2009, the Chair of the *Address Supporting Organization Address Council* (ASO AC) forwarded the Proposed Global Policy for the Allocation of the Remaining IPv4 Address Space for ratification by the ICANN Board. On 5 March 2009, the ASO AC submitted advice in full support of the proposal to the ICANN Board. This proposed global policy had been submitted to the ASO AC by the Executive Council of the *Number Resource Organization* (NRO) on 3 December 2008, and adopted by the ASO AC on 8 January 2009. Each RIR community individually discussed the policy and approved its adoption via its own policy development process. The policy text is published on the ICANN web site at:

http://www.icann.org/en/general/allocation-remaining-ipv4-space.htm

ISOC's Trust and Identity Initiative

The Internet Society's *Trust and Identity Initiative* recognizes that in order to be trusted, the Internet must provide channels for secure, reliable, private, communication between entities, which can be clearly authenticated in a mutually understood manner. The mechanisms that provide this level of assurance must support both the end-to-end nature of Internet architecture and reasonable means for entities to manage and protect their own identity details.

A *trusted* Internet takes into account security, transaction protection, and identity assertion and management. Given the network dependence on unique numbers and the escalating amount of geolocation data being gathered, the privacy implications of the current Internet represent a significant and growing concern. Trust must be a primary design element at every layer of the architecture, and in some cases, existing elements may need to be redesigned or improved to meet emerging requirements.

In late 2007, the ISOC Board of Trustees held an intensive retreat to consider ISOC's role in identifying and pursuing trust and identity issues. The report arising from that meeting, "Trust and the Future of the Internet,"[1] forms the basis of ISOC's current long term strategic initiative.

The Trust and Identity initiative focuses on the following major research programs:

- *Architecture and Trust:* This research program investigates the implementation of open-trust mechanisms throughout the full cycle of Internet research, standardization, development, and deployment.
- *Current Problems and Solutions and Trust:* This research program investigates the mitigation of the social, policy, and economic factors that may hinder development and deployment for trust-enabling technologies.
- *Identity and Trust:* This research program investigates the elevation of identity to a core issue in network research and standards development. ISOC is taking a lead role in reviewing the current Internet architecture and the model of Internet development and deployment. This includes active engagement with participants within the traditional ISOC sphere, as well as with the research, enterprise, and end-user communities. We offer the kind of support for research that enhances and facilitates trust and collaboration with the standards community and that advances the most interesting outcomes of that research.

ISOC is reaching out to the businesses and end users that rely on the Internet to exchange sensitive data. Their needs and concerns inform both our baseline research agendas and ongoing standards and development work. ISOC continues to support the advancement of current technical solutions and best practices through our existing programs.

[1] "Trust and the Future of the Internet," **http://www.isoc.org/isoc/mission/initiative/docs/trust-report-2008.pdf**

[2] "Trust and Identity Initiative" brochure, **http://www.isoc.org/pubs/isoc/docs/trust.pdf**

Call for Papers

The Internet Protocol Journal (IPJ) is published quarterly by Cisco Systems. The journal is not intended to promote any specific products or services, but rather is intended to serve as an informational and educational resource for engineering professionals involved in the design, development, and operation of public and private internets and intranets. The journal carries tutorial articles ("What is...?"), as well as implementation/operation articles ("How to..."). It provides readers with technology and standardization updates for all levels of the protocol stack and serves as a forum for discussion of all aspects of internetworking.

Topics include, but are not limited to:

- Access and infrastructure technologies such as: ISDN, Gigabit Ethernet, SONET, ATM, xDSL, cable, fiber optics, satellite, wireless, and dial systems
- Transport and interconnection functions such as: switching, routing, tunneling, protocol transition, multicast, and performance
- Network management, administration, and security issues, including: authentication, privacy, encryption, monitoring, firewalls, troubleshooting, and mapping
- Value-added systems and services such as: Virtual Private Networks, resource location, caching, client/server systems, distributed systems, network computing, and Quality of Service
- Application and end-user issues such as: e-mail, Web authoring, server technologies and systems, electronic commerce, and application management
- Legal, policy, and regulatory topics such as: copyright, content control, content liability, settlement charges, "modem tax," and trademark disputes in the context of internetworking

In addition to feature-length articles, IPJ contains standardization updates, overviews of leading and bleeding-edge technologies, book reviews, announcements, opinion columns, and letters to the Editor.

Cisco will pay a stipend of US$1000 for published, feature-length articles. Author guidelines are available from Ole Jacobsen, the Editor and Publisher of IPJ, reachable via e-mail at `ole@cisco.com`

The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED

PRSRT STD
U.S. Postage
PAID
PERMIT No. 5187
SAN JOSE, CA

The Internet Protocol Journal is published quarterly by the Chief Technology Office, Cisco Systems, Inc.
www.cisco.com
Tel: +1 408 526-4000
E-mail: ipj@cisco.com