Internet Multi-Homing Problems: Explanations from Economics

Richard Clayton

Computer Laboratory, University of Cambridge
JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom
`richard.clayton@cl.cam.ac.uk`

Abstract

Companies seeking to ensure that their Internet connection is resilient often purchase services from multiple providers. This leads them inexorably towards having their IP address range visible in the global routing table, increasing the resource usage of every Internet router. Since this is essentially 'free', yet impacts the cost and stability of every router in the world, this is a classic 'tragedy of the commons'. There is little prospect of change in the IPv4 world, but there is a chance to fix the problem as IPv6 is rolled out. Unfortunately, SHIM6, the engineering solution chosen to solve this issue in IPv6, will only be effective if universally adopted, and there are no short-term incentives to prefer SHIM6 over a duplication of the IPv4 arrangements. Incentives could be artificially introduced by requiring payment for adding multi-homed address space to the global routing table – a naïve estimate of the actual cost being \$77 000 per routing prefix. However, it would be almost impossible to ensure the substantial revenues involved are correctly redistributed to those bearing the costs.

1 Introduction

The increasing reliance of all sizes of business on Internet connectivity is leading them to seek resilient methods of ensuring that they are never disconnected. Ironically, this resilience is creating instability within the Internet, and, for reasons that economists will instantly recognise, current attempts at solutions are failing to be effective.

The growth of email use in companies has been extraordinarily rapid. For example in the UK, a 1998 survey [15] found only a quarter of small companies using email (and in two thirds of them, only 10% of employees used email regularly). By 2002 a survey [11] of marketing and procurement managers in the auto/electrical component manufacturers, financial services and telecommunications industries didn't even mention if any company wasn't using email – it was just assumed that within this industry

May 24, 2009

sector they would. The 2002 survey was more concerned to show that email was now second in importance to the telephone for both buyers and suppliers. Usage has continued to grow, and access speeds have become faster, so that by 2006, an OFCOM survey of SME businesses found that 84% had an Internet connection, and only 20% of those were still using dialup.

Companies are now increasing their dependence on the Internet by migrating their telephone usage to VoIP (Voice over IP) services, so that their voice traffic shares the same link as their Internet traffic. Recent surveys, such as the 2008 annual OFCOM Communications Market report, show VoIP usage remaining very low with just 20% of users making one or more calls a month. However, this is mainly measuring Skype usage by individuals, whereas the companies being considered in this paper would purchase integrated telecoms products, for which there are few reliable statistics.

As companies discover that they cannot operate without a working Internet connection, they will insist upon resilience. The obvious solution, to purchase connectivity from more than one Internet Service Provider (ISP), turns out to be complicated, as will now be explained.

2   How Internet routing works

As is well understood, machines connected to the Internet have a unique 'IP address'. When machines communicate, routers inspect each of the packets they forward to pick out the destination IP address and send the packet over an appropriate link to a router that is, in some sense, 'closer' to where the packet is to be finally delivered.

Internet address space is allocated to ISPs in a hierarchical manner by the five Regional Internet Registries (RIRs), ARIN, RIPE, LACNIC, APNIC and AFRINIC. The ISPs are also allocated AS (Autonomous System) numbers by the RIRs, which are used to group together their allocations of address space for which they will have a consistent routing policy. The ISPs operate routers which communicate with their neighbours using BGP (the Border Gateway Protocol). These routers learn which 'routes' their neighbours are aware of, where a route consists of a 'route prefix' (the first n bits of a block of IP address space, along with the value of n) and an 'AS path' which indicates the AS's which must be traversed to reach the AS that owns the address block.

In the absence of any overriding local configuration, a router chooses which neighbouring router to send a packet to on the basis of two rules: first it picks the 'most specific' route prefix (the one with largest value of n, representing the smallest enclosing address block). The router then picks the shortest AS path from amongst competing advertisements of that prefix. The reason for selecting the shortest path is the obvious one of getting packets to their destination as efficiently as possible. The reason

for the 'most specific' rule is to simplify route announcements; an ISP can announce a large address block such as a /16 (where the prefix length n is 16), without having to split this up into separate chunks if a subset of the address space, such as a single /19 ($n = 19$, one eighth the size), is to be routed differently.

For a multi-homed company to fully benefit from the resilience of having multiple connections to the global Internet, it must use a fixed set of IP addresses, and the traffic will then arrive over whichever path is shortest and still working. From the description above, it can be seen that for a customer to use the same set of IP addresses with two ISPs, it is necessary for this block of address space to be announced by both providers. There isn't strictly a necessity for the customer to have their own AS, but this is generally seen as the 'clean' way to operate. It has the advantage to the customer that they can more easily change providers, it simplifies configuration for all concerned, and it permits remote systems to check some security properties of the announcement.

Therefore, in practice, for a customer to be multi-homed they will need to obtain an AS of their own; operate a BGP-speaking router (or ask a provider to run it for them); and announce their route prefix to their connectivity providers, so that it will become known to the rest of the world. Hence, an entirely local decision to arrange for resilience has, of necessity, a global impact because the route prefix will be recorded in the 'global routing table' that each and every router must construct to know where to send packets.

3 The 'Global Routing Table'

The size of the global routing table has been a matter of concern for many years. Routers need to keep the table in memory for instantaneous access; which has proved to be a problem for older router architectures where adding memory is expensive or even impossible past a certain limit. Furthermore, inter-router traffic grows along with the size of the table.

There is a specific concern about apparently unnecessary entries, where for example a provider splits some address space in two, and advertises two adjacent /19 blocks rather than a single /18. The CIDR report [4] tracks these occurrences, and at present the global routing table would reduce by 37% if all possible aggregations occurred.

Aggregation is of course impossible if address space is fragmented, e.g.: when a new allocation of address space to an ISP is not adjacent to their existing space. Fragmentation may also occur by choice, because the ISP wants to avoid congestion by splitting the traffic to different parts of their network over multiple ingress paths. Nonetheless a great deal of fragmentation is unnecessary and aggregation is often possible. Social pressure, exemplified by the weekly publication of the CIDR report, has helped to

reduce the number of unnecessary announcements. The importance of this social pressure can be found in a 2001 survey paper [8], where the observation is made that there are visible dips in the upward trend immediately after IETF meetings where the issue of routing table size was discussed.

Growth of the routing table has usually been exponential [8], and the current trend is a growth of about 25% per annum, with the May 2009 size being just under 300 000 prefixes. The growth is caused by new allocations of IP address space (as new people connect to the Internet), traffic engineering schemes to balance the load and avoid congestion, and route prefixes that are only present to permit multi-homing.

A 2005 study by Meng et al. found that around 45% of prefixes were 'covered', viz: they were more specific prefixes for other routes; and they ascribed 44% of these to multi-homing; i.e. around 20% of the entire global routing table is present solely because of multi-homing [12]. Furthermore, Bu et. al found that the number of multi-homing prefixes (along with prefixes that were present for load balancing reasons) was growing faster than the routing table as a whole [5].

There has been a similar growth in AS number allocations, with about 31 000 currently in active use, and another 15 000 allocated but not yet in use on the public Internet [7]. Growth is presently a steady 5 000 or so per annum. Since AS numbers were originally 16-bit values, this would have meant exhaustion in 2011 or so, and so the BGP protocol has been re-engineered to permit the use of 32-bit AS numbers [16] and support for this will be universal by the beginning of 2010.

Further evidence of the role of multi-homing can be seen by examining the amount of address space advertised per AS. Since AS numbers are generally allocated in order (albeit they are passed to the RIRs in lumps which are then used up at different rates), the higher the AS number the more recently it has been issued. Additionally, most of today's ISPs have existed for many years (albeit seldom under the same name, or management).

Therefore we would expect ISPs to have low AS numbers and large amounts of address space, but higher AS numbers will have been allocated to multi-homed companies who use a small amount of address space. Examining a scatter plot of the address space announced by each AS (see Figure 1) we see that our prediction is borne out, and most of the high AS numbers (past 20 000) have very small amounts of address space, whereas many of the low AS numbers (particularly below 5 000) have considerably more.

Besides the impact on the size of the global routing table, multi-homing companies share a further unfortunate characteristic in that they are more volatile. Meng et al. observed [12] that covered prefixes (i.e. the category into which multi-homed companies fall) were more likely to be announced and then later withdrawn. Each time a route prefix appears or disappears, then all of the world's routers have to recalculate their version of the global routing table, a resource intensive task. When many prefixes are announced
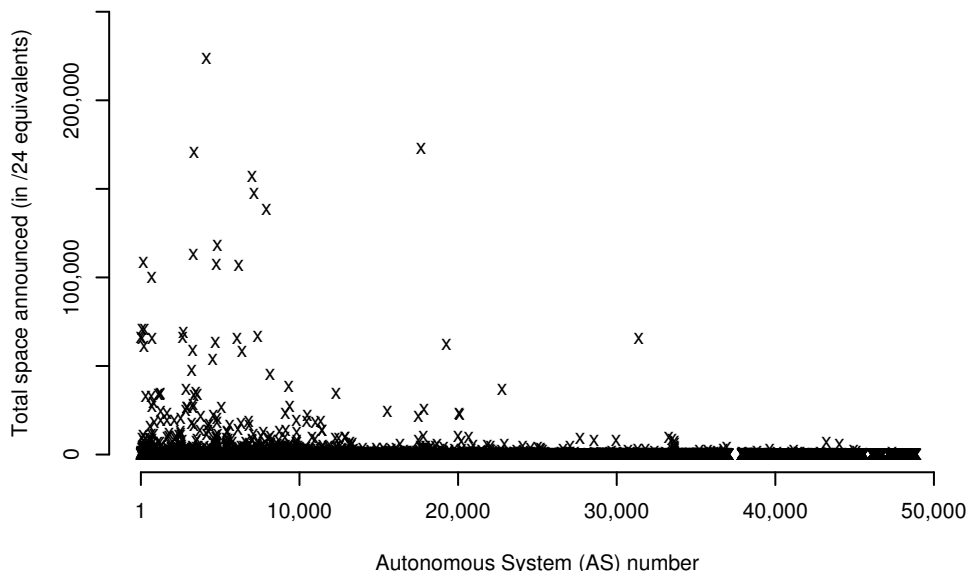
Figure 1: Size of IPv4 address space announced (in /24 (256 address) equivalents) plotted against the AS number making the announcement.

or withdrawn over a short period of time it can be several minutes before the routers catch up with the changes and are routing packets normally again. Thus the existence of the extra routes is contributing to overall instability and adversely affecting 'availability' world-wide.

Economists will not find it hard to see parallels with other scenarios. Individual ISP customers choose whether or not to become multi-homed by considering a local cost/benefit analysis, rather than assessing the cumulative impact on the size of the global routing table, or the need to re-engineer the entire BGP infrastructure to cope with 32-bit AS numbers. This is essentially Hardin's 'Tragedy of the Commons' played out in a high-tech setting [6].

4   IPv6

Although it is too late within the IPv4 protocol to prevent local multi-homing decisions having global impact, one might hope that this mistake will not be repeated in IPv6. But the outlook is gloomy.

The problem facing IPv6 is well understood in the community. In August 2003, RFC3582 [1] "Goals for IPv6 Site-Multihoming Architectures" explained the need for multi-homing and set out some clear goals, including scalability ("A new IPv6 multihoming architecture should scale to accommodate orders of magnitude more multihomed sites without imposing un-

reasonable requirements on the routing system.") and limited cooperation ("A multihoming strategy may require cooperation between a site and its transit providers, but should not require cooperation (relating specifically to the multihomed site) directly between the transit providers"). The IPv4 multi-homing system was assessed against the RFC3582 considerations by Abley et al. in RFC4116 and found wanting [2].

Furthermore, Savola and Chown [14] provided some indications of the sort of scalability required, calculating (extremely simplistically) that if there were 1000 multi-homed firms per million people this would result in a million extra routing prefixes if the current multi-homing scheme was continued. They also drew attention to the risk that major network failures could result in bursts of 100 000 simultaneous BGP updates – a considerable workload.

Their paper went on to survey the new architectures being proposed in 2005. They distinguish:

- Host-centred proposals, where the hosts have multiple IP addresses, one for each link to the Internet. The hosts must arrange to communicate these addresses to the other end of connections, who then select which address to use.

- Modifications to the transport layer to allow dynamic changes to IP addresses within the TCP protocol (or the replacement of TCP with some other protocol such as SCTP). They did not believe there was much enthusiasm for this.

- Use of the 'Mobile IPv6' mechanisms to permit the link to the Internet to change. This posed some difficulty, not least because a key security mechanism of Mobile IPv6 is that when bindings change a check is made to ensure that this is agreed to by communicating with the old address – but if the link to the old address has just failed then this is impossible.

- Schemes that break the binding between identification and location. The Host Identity Protocol (HIP) is one such, using cryptographic hashes to link identifiers at the transport layer and address values, however this requires too many changes to be viable. Another, LIN6, has been patented and this has prevented serious consideration.

- Schemes that propose geographic allocation of IPv6 addresses. These fix the aggregation problem because most customers would multi-home with geographically close providers. However, the Internet isn't wired up in this way, and it is unclear how country level links, carrying significant volumes of traffic, would be funded.

As can be seen, the assessment made of these proposals was basically just testing their engineering elegance; with the addition of a small amount of commonsense thinking about how Internet peering actually works.

Around the same time, Lear documented the issues that ought to be considered in RFC4219 [10]. He set out 45 questions, all of which related to technical aspects of possible solutions. He failed to ask what the prospects were of getting a solution deployed in the real-world, perhaps because it was widely believed that there would be no actual choice about that.

The proposal that eventually emerged from amongst the various competing ideas to be taken forward was SHIM6, a host-oriented scheme.

4.1 SHIM6

In the SHIM6 design, connections are made using the TCP protocol in the normal way, but if more than a few packets are exchanged (and so the overhead appears to be worthwhile), the multi-homed host will tell the other end of the connection about any other IPv6 addresses on which it can be reached. If the connection subsequently fails, then the other end will use these fallback values, and tag the packets to indicate this has happened.

The higher stack levels will be unaware of the changed IPv6 address values because the receiver detects the tag and fixes up the packets to contain the original address, hiding the link failure. The 'fixing-up' layer is implemented as an add-on within the network stack's IP layer, hence the name, which is not an acronym, but is chosen because 'shim' is a common jargon word for modules that add functionality to a network stack layer.

The SHIM6 protocol is complex, not least because it must be secure against fraudulent announcements of which IPv6 addresses are valid ways of making contact. The main description covers 136 pages, along with another 125 pages of related documents. Admittedly, some of the pages are filled with justifications for architectural choices and reasons why parts of the design are the way that they are, but it is still a significant undertaking to implement the protocol. By comparison, the base document for TCP (RFC793) is 85 pages.

This complexity is compounded by the protocol still being (in May 2009) a work-in-progress, with the descriptions remaining at the 'Internet-Draft' stage, rather than having become stable 'Standards Track' RFC documents. This strongly suggests that SHIM6 will not be widely implemented and universally deployed in the near future, if at all.

Unfortunately, the way that SHIM6 works means that if it is to provide any resilience, then both ends of a connection must be using it. Thus, for its benefits to be fully enjoyed by a multi-homed site, it must be universally deployed. Naturally, there is a clear incentive for the multi-homed site to upgrade their machines to use the new protocol. However, the incentive for others is entirely absent, which means that even if SHIM6 turns out to be

straightforward to deploy, there is no obvious reason for people to bother.

Hence, especially in the short term, we must expect multi-homed IPv6 sites to use the same scheme as in IPv4, viz: obtaining their own AS number, and adding their route prefix to the global routing table. Since these sites will now no longer benefit from SHIM6 there will no longer be any incentive even for them to deploy it. With no advantages to early adopters, and those who might use it settling for another approach entirely, it is difficult, at the time of writing this paper, to see the protocol catching on.

4.2   Co-operating ISPs

Although, as discussed above, RFC3582 [1] ruled out solutions that require cooperation between transit providers, this could in fact offer a way forward.

In practice, multi-homed companies will be purchasing service from a small number of ISPs in their geographic region. These ISPs could cooperate by arranging that all of the multi-homed customers they shared with a particular competitor were place within a single block of address space whose prefix was announced by both ISPs. When connectivity via one of the ISPs failed, the other ISP (where there was no problem) would then announce a more specific route for the customer, so that all of the traffic flowed through the working connection.

Whilst there were no connectivity problems this would markedly reduce the number of prefixes in the global routing table, and the extra routes added in the event of local failures would not be a huge burden. However, IP address space management would be far from simple – in regions where there were dozens of competing ISPs there would have to be hundreds of different blocks of shared address space.

So although this approach could conceivably be made to work, there would be considerable costs involved in arranging the necessary cooperation between the ISPs. In addition, the scheme would almost certainly require customers who changed suppliers to renumber to another block of IP address space. Since renumbering is of itself disruptive, this might suit the ISPs (because there would be a disincentive for customers to leave) but it must be presumed that the customers would not choose such an arrangement if others were on offer.

Hence although co-operation might be desirable, without creating some disincentives to the existing method of multi-homing, it is most unlikely to be adopted.

5   Discouraging growth in the global routing table

One way to prevent unjustified growth in the global routing table would be to be charge people for entries. Provided that the charge was correctly set, this could fairly recompense those whose resources are being consumed by

companies choosing to become multi-homed in the current manner. In fact, there are existing mechanisms which could be used for this purpose, because adding a route is not quite as free as has been suggested so far.

The Regional Internet Registries (RIRs) currently fund their activities by charging members for their services. For example, RIPE NCC (the RIR for Europe, the Middle East and parts of Central Asia) splits their membership up by size, from "large" though "medium" to "small", charging €5 500/annum to the large members, and €1 300/annum to small ones. The size is determined by a complex formula that assesses how many AS numbers and blocks of IP address space have been allocated, and how long ago this allocation was initially made.

Therefore, should a company wish to become multi-homed, they could join RIPE in their own right – which would cost them €2 300 in the first year and €1 300 thereafter. However, if they were to obtain their space via an existing member then that member might well pay nothing more by becoming a little 'bigger', but even it the new customer pushed them over a charging boundary, the amortised cost over all of their customers would only be a handful of Euro each.

So there is a small financial disincentive to creating new multi-homed sites. However, the actual worldwide cost of coping with the extra prefix is substantially more than a few thousands of Euros. We can estimate what this cost might be by calculating the total current cost of providing routing, and dividing this down by the 300 000 route prefixes currently in the global routing table. Unfortunately, this estimate can only be made very roughly, because of a lack of detailed numbers.

One rough and ready approach is to consider the topmost tier of network providers, those who do not have 'transit providers', but only mutual peering relationships. There are currently 13 such, each of which will have around 10 000 routers costing say $100K each (i.e. $13 billion of kit between them). The next tier down, which have complete meshes within regions, are about 10 times as many, albeit around 10 times smaller, but with cheaper hardware their routers cost them in total around $8 billion. Finally there are the stub systems, around 30,000 of these, but with just a handful of $30K routers each: for roughly another $2 billion.

Hence the total infrastructure cost can be estimated to be very roughly $23 billion. This is in line with estimates of yearly sales of $12.8 billion [9], given that routers need regular replacement as traffic (and the global routing table) grows. Dividing this down gives a cost per prefix of $77 000.

Of course, this is only one way of calculating the cost of adding a route prefix. The actual costs of any particular prefix is either zero (the general case where it makes no difference) or occasionally the cost of an entire new router (when an old one can no longer cope). Furthermore, new routers may be purchased anyway to handle greater amounts of traffic – and being newer designs they may cope with bigger routing tables as a matter of course.

Hence other calculations are certainly possible. But the real difficulty in trying to take this approach is not how much should be charged, but the lack of any obvious way to distribute this money to subsidise the people purchasing and running the routers. If the money is equally shared 'per router' then if the $77 000 figure is correct, by purchasing an AS and a cheap router you would actually get given money! If routers are not counted equally then money should flow to tier 1 providers from the multi-homed edge systems; but it would be extremely hard to prevent 'gaming' of the system by misrepresenting how many routers are actually needed and hence should be subsidised.

The conclusion must be that there doesn't seem to be any practical way of charging for routes at the present time; but the disparity between the straw man figure of $77 000 and the few thousand Euro, that would be the most that would currently be paid, underlines the point that multi-homed customers are consuming expensive resources but are not having to pay anything like the full cost.

6   Related work on the economics of protocols

Ozment and Schechter specifically looked at the issue of bootstrapping the adoption of Internet protocols, their focus being specifically on security protocols [13]. They developed a formal model, and considered strategies that might lead to protocol adoption.

Only a few of their strategies would work for SHIM6. "Global Mandate" would correspond to having some way of fining people who did not deploy the protocol, which would be unrealistic. "Partial Mandate" is inapplicable because there is no 'tipping point' after which deploying SHIM6 would be an obvious choice. "Bundling" is also inapplicable at present because SHIM6 does not give any other benefits – although if there was more commonality with the 'Mobile IPv6' protocols that might change. Their "Facilitating Sub-network Adoption" strategy might be viable if multi-homed companies were able to work with the subset of the whole Internet with whom they wanted to have reliable long duration connections; that is, they don't need the whole Internet to use SHIM6, just certain parts of it. "Coordination" also seems inapplicable, but "Subsidization" might well be the best way forward – those who stood to lose most from a ever growing IPv6 global routing table could invest in ensuring that SHIM6 was incorporated into standard network stacks, and hence became widely adopted.

The real problem is that SHIM6 may make engineering sense (albeit, given its complexity, that could be debated), but the economics of its deployment has hardly been considered within the IETF. In contrast, within the totally unrelated area of email spam control, economic arguments have come to be seem as absolutely key when evaluating proposals.

It is extremely common for new anti-spam solutions to be proposed which would only work if universally deployed, which have no benefits for early adopters, which assume that spam senders would not change their behaviour, or that senders of legitimate email would be delighted to pay extra for the privilege. Proposals with such failings are routinely dismissed by the anti-spam community and no substantial work put into experimenting with them.

This type of security economics analysis is widely used within forums such as the IETF Anti-Spam Research Group (ASRG). It is not presently described in any formal academic papers, but, as is the way of these things, has been quite beautifully encapsulated in the widely circulated "Why Your Anti-spam Solution Won't Work" [3] which, although written to amuse, is of immense practical use in summarising what is wrong with a new proposal. Almost none of its points are technical. The emphasis is on economic, legal and philosophical objections – as well as the occasional medical issue, since imaginative new methods for killing spammers are seldom painful enough.

7   Conclusions

As uninterrupted access to the Internet becomes central to the day-to-day operation of companies, they are seeking ways to make that access more resilient. Purchasing connectivity from multiple ISPs gives resilience, but to fully realise the benefits when one of the connections fails, it is necessary for every router in the world to learn of the existence of their particular block of IP address space. The cost of this is out of all proportion to what is actually being paid by the company – a modern day "Tragedy of the Commons".

SHIM6, the engineering fix for this within the upcoming IPv6 protocol is complex, yet to be finalised, and offers no special benefits to early adopters. There is little reason to believe that it will be rapidly and universally deployed. This means that the current exponential growth of the global routing table in IPv4 is likely to be replicated in IPv6.

Security Economics has already begun to permeate the way in which we evaluate other protocols, such as anti-spam schemes. It is clearly well past time that proposals for new network layer protocols were considered in a similar manner. One way of achieving this would be for the IETF to require an "Economics Considerations" section within all standards track RFC documents. Sections on "IANA Considerations" and "Security Considerations" are already mandatory.

Social pressure has had a significant effect on the growth of the global routing table so far. This may continue to be the most effective (and by far the cheapest and simplest) mechanism to rely upon. The way forward may be for multi-homing of small customers using global routing announcements to cease to be seen as a legitimate engineering solution.

It can only be a matter of time until a major ISP does a deal with a

competitor to offer multi-homing to ten thousand of their biggest business customers, with a managed BGP-speaking router and a block of address space bundled into their offering. When that happens, they may agree to cooperate in announcing the address space as set out in 4.2 above. If not, and their initiative is popular enough in the marketplace to grow the global routing table by 30% almost overnight; we may see a rapid change away from current laissez faire attitudes.

Acknowledgements

References

[1] J. Abley, B. Black and V. Gill: Goals for IPv6 Site-Multihoming Architectures. RFC3582, IETF, August 2003.

[2] J. Abley, K. Lindqvist, E. Davies, B. Black and V. Gill: IPv4 Multihoming Practices and Limitations. RFC4116, IETF, July 2005.

[3] Anonymous: Why Your Anti-spam Solution Won't Work". 2004. `http://craphound.com/spamsolutions.txt`

[4] T. Bates et. al: CIDR Report. `http://www.cidr-report.org`

[5] T. Bu, L. Gao and D. Towsley: On characterizing BGP routing table growth. Computer Networks, 45(1), 2004, pp. 45–54.

[6] G. Hardin: The Tragedy of the Commons, Science, 162(3859), 1968, pp. 1243–1248.

[7] G. Huston: The 16-bit AS Number Report. `http://www.potaroo.net/tools/asn16/`

[8] G. Huston: Analyzing the Internet BGP Routing Table. Internet Protocol Journal, 4(1), 2001, pp. 2–15.

[9] Infonetics Research: Great year for service provider router market, but 4Q showed signs of downturn. Press Release, 27 February 2009. `http://www.infonetics.com/pr/2009/router-switch-market-highlights.asp`

[10] E. Lear: Things Multihoming in IPv6 (MULTI6) Developers Should Think About. RFC4219, IETF, October 2005.

[11] S. Leeka, P.W. Turnbulla and P. Naudé: How is information technology affecting business relationships? Results from a UK survey. Industrial Marketing Management, 32(2), 2003, pp. 119–126.

[12] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, L. Zhang: IPv4 Address Allocation and the BGP Routing Table Evolution. ACM SIGCOMM Computer Communication Review, 35(1), 2005, pp. 71–80.

[13] A. Ozment and S. Schechter: Bootstrapping the Adoption of Internet Security Protocols. Fifth Annual Workshop on the Economics of Information Security (WEIS2006), 26–28 June 2006.

[14] P. Savola and T. Chown: A Survey of IPv6 Site Multihoming Proposals. Vol 1, Proceedings of the 8th International Conference on Telecommunications, ConTEL 2005, IEEE, 2005, pp. 41–48.

[15] A.A. Sillince, S. MacDonald, B. Lefang and B. Frost: Email Adoption, Diffusion, Use and Impact Within Small Firms: A Survey of UK Companies. International Journal of Information Management, 18(4), 1998, pp. 231–242.

[16] Q. Vohra and E. Chen: BGP Support for Four-octet AS Number Space. RFC4893, IETF, May 2007.