

The IP Addressing Issue

0 Status of this memo

This draft document will be submitted to the RFC editor as an informational Document. Distribution of this memo is unlimited. Please email comments to jnc@lcs.mit.edu or fax comments to (804) 898-7663.

1 Introduction

The packet layer of the IP architecture is about to enter a period of stress caused by deficiencies in the IP address. This stress is caused by a number of inter-related problems. This note describes these problems, lists some suggested solutions, and discusses pros and cons of each of those solutions.

2 The Problems

The IP address structure is facing three problems; some widely appreciated, and some not.

2.1 Exhaustion of Class B Address Space

The first, and one that has received a fair amount of notice in the technical community, is that the supply of class B network numbers is being used up quite rapidly. Around 16 thousand of these network numbers exist in the current address structure; xxxx thousand have already been assigned. Moreover, a simplistic statistical analysis indicates that the rate of growth closely fits an exponential curve, and if that curve continues to be followed, the existing stock of class B addresses will be used up in yy months.

Of course, the curve is really an S-curve, and will top out at some point, but it is not clear when. While the spread of the Internet in the U.S. academic community is slowing as coverage becomes ubiquitous (i.e. it is itself an S-curve), there are two largely untapped population factors which will keep the growth up for a while to come. The first is the spread of the Internet within the U.S. government and within U.S. companies; while there is currently some coverage in these groups, it has not yet reached the levels of the academic community, but can be expected to do so over time. The second is the spread of the Internet to the international community, where this three phase growth pattern will probably repeat.

2.2 Exhaustion of Entire IP Address Space

This is a short term version of a longer range problem, which is that the entire IP address space is being used up. At first glance, this seems unlikely, since the 32 bit space allows for approximately 4 billion distinct addresses.

However, a great deal of structure is imposed on this 32 bit field, in an attempt to allow agents inspecting an address to discover something about where in the network that address is. Imposing this structure does not allow the address space to be used efficiently, since there are many unused bit combinations. For example, if it is known that all addresses that start with the bit pattern 00010010 are at MIT (and all packets whose destination address starts with that pattern are sent to MIT), then unless there are 2^{24} hosts at MIT (which there are not), all unused combinations starting with that pattern are effectively wasted. There is no way to avoid this without a fundamental change to IP addressing and routing.

It is unclear when the IP address space as a whole will be used up, assuming we continue in something like the current fashion. The IP address space is divided into different classes, and the usage rate in the various classes differs. As mentioned above, one class is already in some danger of running out. If the limit above is not to be the finale, presumably some reallocation of classes would be necessary, but how long this will give us depends on how complete the reallocation is (more efficient schemes are more complex), as well as future growth patterns, which are not well understood.

Even if the address space were used efficiently (i.e. all possible bit combinations were valid addresses), if the Internet continues to grow and is not supplanted by an alternate networking technology, there will eventually be more hosts that can be addressed in this many bits (unlikely as this may seem at the moment).

2.2 Inutility of the Current IP Address Structure

Finally, and most important, the routing mechanisms themselves are breaking down due to the sheer number of different destinations they are required to track. This is in fact a far more serious problem (both in terms of immediate failures in the network as well as long range intractability) than the two above, although it may not be as well known or as apparent to the broader community.

The difficulty stems from the fact that the Internet is effectively a single level hierarchy. (For the purposes of this analysis, subnets can be ignored, although they did help a great deal.) The objects in terms of which routing is done are IP networks; each network is visible throughout the entire Internet (modulo local policy controls), and traffic cannot be routed to a destination in a network unless that network and something about its location are known to the source of the traffic and points in between. Simply examining a network number (as opposed to an entire address) will tell you nothing

about where in the network that network is.

The problem is that any routing technology is limited in the number of discrete destinations it can track. The limits come in storage space to keep databases and processing time to do computations based on them, as well as bandwidth usage to send information around. Even systems which cost $O(\text{zzzzz})$ in the number of destinations (and most routing technologies aren't even that good) will have limits. The exact number of destinations at which a particular routing technology tops out differs, but as a general proposition the best systems of which we know now seem to have limits in the low tens of thousands with the current levels of hardware performance.

The solution to the problem is to adopt an address with more structure in it; this will allow groups of networks (and eventually groups of groups) to be represented as single objects for the purposes of routing, which will reduce the costs of routing to a manageable burden. Once again, there is no way to handle this without a fundamental change to IP addressing and routing; better hardware will not solve the problem as the size of the Internet is increasing faster than the capabilities of the hardware.

As a side note, the existing routing technology in the Internet has a very short useful lifetime. The problems of EGP are well known; in addition to not being able to handle cycles in the topology, since it sends the entire routing table as a single packet, it has a practical limit in the low thousands of networks. BGP will remove these two limits, but has little support for real policy routing, and in any case is subject to the ultimate size limits above inherent in the existing IP addresses. As far as IGP's go, the new link state protocols OSPF and dual IS-IS are very good protocols, but have the same problems as BGP. For example, OSPF can probably barely handle (if at all) the 2^{14} class B network numbers currently allowed (the lack of definity is because the actual limits depend on the speed and memory size of the router), so increasing the number of class B networks by a factor of 10 (assuming this were somehow possible within the IP address space) would not really help.

3 Possible Solutions

A number of solutions have been proposed to solve the 'IP addressing problem'. However, most of them do not fully address all three of the problems outlined above. In some cases, this is deliberate; the proposals were put forward as temporary fixes to allow more time to be spent on a more complete solution; in other cases, the author was perhaps not aware of the full dimensions of the problems.

Two self-evident conclusions can be drawn from the list of problems above. First, any proposal which does not solve all the problems listed above is not a real solution, but simply a temporary

patch, and is only worth considering if the extra time provided is needed and worth the cost.

Second, given the interrelation between addressing and routing, a more satisfactory solution almost certainly lies in considering and solving problems in the two areas together. Given that routing is by far the harder problem, the address structure chosen should be designed in light of the requirements of the final routing architecture; put another way, the address structure should be designed to make the job of the routing as easy as possible. Different address architectures can make a great deal of difference in the difficulty of routing them; to design an address structure without reference to the routing system that will provide the paths for the traffic is most unwise given the extreme technical challenges posed by the current requirements on routing in the Internet.

In any case, it is worth going down the list, explaining each, and listing the pros and cons of each. It is unfortunately not possible to go into detail on each solution, since that would require a long paper on each one. Note that most of these only address the address space limitations, not the routing problems, which, as have been pointed out, are actually the most severe.

In general, most solutions propose the adoption of some sort of new address; schemes vary as to how (or if) the new address is to be carried in the packet. Some propose to keep the existing IP packet format, while others propose to modify it.

Note also that most of the proposals are not mutually exclusive; one can take parts of one and mix it with parts of another to provide more complete solutions. For example, some of the ideas in the last proposed solution (such as the completion of Open Routing to handle new and more complex addresses and use of the existing 32 bit address as a UID) could obviously be used in other proposals, and several of the other solutions appear as pieces of or options in the last proposal.

3.1 Increase the Number of Class B Networks

The first (and only non-radical) proposed solution is to increase the number of class B network numbers available, either by allocating half the class A numbers to class B, or part of the class C address space.

The advantages of this are that it will require no changes in most hosts, and will extend the life of the current addresses substantially, since class B is the only one nearing exhaustion.

The disadvantages are that it is only a short term solution, it is not a solution to the most pressing problem (the third) anyway, it is not a solution to the address space exhaustion, it will require

changes to all routers, and, if the class A subvariant is used, to some hosts (which persist in ignoring the requirements in the Host Requirements and understand the structure of IP addresses).

Nonetheless, this solution might be of use as an interim measure if more time is needed to implement a final plan.

3.2 Reformat the Existing IP Address

The second proposed solution is to redefine the interpretation of the existing 32 bit field to make it more useful. The primary proposal here is the 'bottoms up' proposal. Briefly, this proposal contemplates redefining the interpretation of the IP address to allow a multi-level hierarchy in which (through use of masks and assignment of extra bits to each level as needed) each level can grow efficiently, as needed. A way to perform policy routing using this address interpretation is also provided.

The advantages are that this does make efficient use of the address space, and does not require changes to hosts and many routers. It also improves the routing situation.

The disadvantage is that this particular proposal does not remove the ultimate limit on the size of the IP address; the proposal as written uses the existing 32 bit addresses. (Clearly, one could vary this to use new addresses, in which case it would fall into one of the classes below.)

3.3 Make IP Addresses Non-Globally Unique

The third proposed solution is to change IP so that addresses are not globally unique, but only unique in a single AS. This effectively corresponds to the scenario seen in a number of protocol families where distinct catenets with overlapping address spaces are glued together. This is usually undertaken to avoid renumbering an entire catenet when two catenets which developed separately are joined, rather than to expand a single catenet, but the details are the same whatever the motivation. In any case, experience in other protocol families with this solution is a useful guide.

The advantages are that it requires no change to the hosts, and it also avoids changes to any non-border routers.

The disadvantages, depending on the subvariant, are that either an overall size limit still remains, or a mechanism equivalently complicated to that of some of the more complete proposals must be developed to do the routing among AS's, etc. (This is an example of the adage that a problem swept under the carpet will always pop up somewhere else.)

In one subvariant (seen where distinct catenets with overlapping address space assignments are joined, but probably not

useful to the Internet), distant AS's are permanently mapped into the address space of the local AS, but in a place different from their 'native' address. Thus, both the third (routing difficulties) and second (total size limit) problems still exist.

In a different subvariant, either not all AS's can be mapped in, or external destinations are dynamically mapped into the address space of the local AS. In this case, a larger (and probably more complex) address must be adopted for use between AS's, routing must be designed to route it, and mechanisms developed to do the dynamic mapping. Depending on the details, this might look very similar to the next scheme.

In yet another subvariant, the AS's are joined at their edges not by packet level routers, but by application level gateways. This is in a sense a variant of the one above, except that the more complex address is the host name. The difficulty here is that the applications must in most cases be modified to pass the identification of the ultimate destination on start-up (in-band, since there is no out-of-band channel in TCP); very few (such as electronic mail) currently do this.

3.4 Define a New IP Address

The fourth proposed solution is to define a new IP address, which would be carried in an IP option in existing IP packets, perhaps with a pointer to the location of the option carried in a class E address. In one variant of this idea (similar to the one above) the AS number would be the extension (so that the existing IP address is again not unique across the entire Internet), and would be carried in an option. In another, the 'bottoms-up' addressing scheme would be used.

The advantages are that this removes the limits on the number of addresses, and this allows a solution to the routing problem, although one is not proposed, except in the 'bottoms-up' scheme.

The disadvantage is that all the hosts must be modified to generate the addresses in this fashion, and the overall size of the new address (assuming it is a multi-level address to make routing easier) is limited by the maximum amount of free space in the IP header. (If the hosts are not modified, but the new address is added by some agent, then this solution turns into the previous one.) An additional disadvantage in the case where the AS number is the extension is that this still only provides one extra level of hierarchy, which in the long run will not be enough.

3.5 Define a New IP Packet

This is a more radical attack on the problem, and since a 'clean slate' is available, the proposals differ substantially. The

chief advantage of this approach is that other problems with IP can be solved at the same time, but these are outside the scope of this discussion. The chief disadvantage is that all the hosts must be modified to generate the new format packets, but some schemes include interoperable transition plans to ease this.

A number of proposals include this step as an option, to make the new addresses less of a kludge, or to provide extra capabilities.

3.6 Integrated New Host Identifier, Address and Routing.

The sixth proposal is a multi-stage solution which attacks a number of problems at once. It is in some respects related to the two above, since it contemplates a new IP address and (eventually) a new packet format, but it differs from them in introducing a new concept into the IP architecture (the Host Identifier) as well as tight coupling to a routing architecture.

It envisions creation of a new IP address, of varying length with a varying number of levels, upgrading the Open Routing protocol to handle these new addresses, conversion of interpretation of the existing 32 bit addresses to UID's for Host Identifiers, and, in the long run, a final step to allow the system to contain more than 2^{32} distinct nodes. One possibility for the latter is to make the UID's locally significant only (using some of the mapping techniques laid out in the second sub-variant of 3.3); the other is a new IP packet format with 64 bit UID's.

The advantages of this proposal are varied. First, the structure of the new addresses can be oriented toward the main goal of making the routing easier. Second, a UID means that a number of problems with machines with more than one address or changing address can be attacked. Third, the initial retention of the existing address as the UID means that hosts and routers do not have to be changed right away, and the change to hosts is small; the address can be constructed by the first 'new-style' router (although eventually for efficiency the hosts should do this directly). Fourth, the existing **32 bit space can be used with maximal efficiency, delaying the date at** which the exhaustion of this space must be tackled. Fifth, if the 64 bit UID path is chosen, and the phaseover started before the 32 bit space is used (so that the UID of any host in the new 64 bit space is simply the zero extension of its UID in the 32 bit space), there will not be any cases of new and old style hosts which cannot communicate due to the inability of the old address space to name hosts in the new space, which is the usual cause of problems in conversions.

The disadvantages of this proposal are that the first 'new-style' router will have to determine (and add, or otherwise retain) the address which corresponds to that destination UID, which is a repeated small cost, and in the long run all the hosts will have to be changed.

3.7 Use ISO

There are a number of sub-variants in this option.

One possibility is to use just the ISO address; this is effectively a new address scheme, as described in 3.4, but the address would be in common with ISO. This would allow 'packet wrapping' in a simple algorithmic way (since no complex tables would be needed to map addresses), but it is not clear if this is useful. Packet translation would also be possible, with the same caveats.

Another is to use the ISO packet layer, but retaining the existing stream and above protocols; this is effectively a new packet format, as described in 3.5. This would unify the two catenets at the packet level, but this is probably not a big advantage given the multi-protocol router technology. Hosts from different suites would still not be able to interoperate without application or other gateways.

A third is to keep the TCP suite of applications, but to run them above the ISO stream protocols. This is little different from the previous scheme in its effects, but would have more far-reaching effects in terms of host software.

The last is a complete conversion to the ISO suite. This would have problems during the (likely lengthy) transition which are identical to the ones we see now with interoperation; service gateways are an imperfect (and in some ways crippling) solution, and general translation has proven impractical.

4 Conclusions

It is possible to draw some initial conclusions as to which of the possible classes of solutions is to be preferred.

To begin with the last, a conversion to ISO, while alluring, is not currently a useful option, for reasons both political and technical.

Reflection in both camps on the complex political situation between supporters of the two protocols has led to a strategy that actually appears to have some advantages. Basically, both sides agree that the overall goal is to create the best possible packet data architecture. Given that, a plausible case can be made that that end is actually better served by the existence of two competing efforts, provided that no energy is wasted in fruitless political combat between adherents of the two camps. As long as all the effort is directed towards improvement of the two protocols suites, the end result will be better than the result of a single effort, especially if each side feels free to inspect, learn and borrow from the work of the other.

In the context of this strategy, and also with reference to the technical status of the ISO family, adoption of an ISO conversion strategy to answer the problems of the Internet is clearly inappropriate. Both protocol families face problems in handling growth, especially given the choice of the ISO designers in making administrative concerns the main spring of their address design, rather than topology and abstraction. The continuing existence of the two different standards (down to the address/packet layer, where the real problems lie), each with radically different means of creating designs, is necessary to create the best answer to these challenges. In addition, the ISO architecture is still lacking some necessary pieces (such as an inter-domain routing protocol) for use in replacement of the Internet.

As noted above, the first two options are not really suitable. Since neither removes the straightjacket of the existing IP address (in length, if not in structure), any solution based on them would be short term at best.

Ruling out the third option (and all its subvariants) is a little harder to rigorously justify (since strictly speaking they can meet all the requirements), but it seems clear in the light of engineering experience that the existence of a single global system in which each host has a unique tag is extremely valuable and robust. Previous experiences with mapping solutions indicate that while they can eventually be made to work, they are inferior in many ways. Many other goals (such as security) are much easier to tackle in such a system. The last subvariant (with application level gateways) will crimp development and deployment of new applications, since lack of direct packet level connectivity will require the creation of application gateways as well as the applications themselves before a new application can spread. One of the chief advantages of the IP system (as compared with the various local solutions arrayed around the NCP-based ARPANet which preceeded it) is the direct packet access to all corners of the network.

The choice thus comes down to the middle three options of a new address, packet format, or the integrated rework of the packet layer. This is not really a choice between differing approaches, but simply a choice of how expansive a rework is desired or feasible.

Given that major changes are going to be necessary in any case, and given that a conversion/interoperation plan with minimal up-front costs is available, it seems likely that the best course for the long-term is the third; a complete and integrated rethinking and reworking of the basic addressing and routing facet of the packet level of the architecture.