Practical Considerations for Network Addressing using a CIDR Block Allocation

Håvard Eidnes [1]

Abstract

The strategies for IP address allocation are changing. The reason for this is the increasing growth the Internet has seen lately. Some short-term solutions to the strains this growth is causing have been proposed, such as CIDR and C-block address allocations. This paper will try to explain these changes, and how local network managers needing new IP addresses should plan address assignment for the benefit of all.

I. Background

During the last few years the TCP/IP protocol suite has seen a rapidly increasing use. It has always been reasonably easy for all organizations using this protocol suite to be assigned globally unique addresses; in fact, organizations have up through the years mostly been encouraged to get hold of globally unique addresses, presumably to avoid reassigning addresses when an organization needs to connect to another organization or to the world-wide Internet.

Most of this paper assumes some basic knowledge about IP addresses, for some introductory material see the sidebar on page 3 and the references cited there.

Due to the continually increasing popularity of the TCP/IP protocol suite, the Internet community is currently facing three related major problems:

- The class B address space is in danger of exhaustion. Currently more than half of the total class B address space has been handed out. The main reason for this is the lack of appropriate-sized network numbers for mid-size organizations. For most organizations, a class C network number with room for up to 254 hosts is too small, while a class B network number with room for up to 65534 hosts is way too large, resulting in poor utilization of the address space.
- The Internet routing explosion. This is a problem for those routers on the Internet which do not use a default route, and consequently have to carry a routing entry for each and every network they need connectivity to (which for the Internet basically means *all*). The current growth trend is approximately exponential, with a doubling every ten to twelve months. Clearly this cannot continue for long.
- Eventual exhaustion of the entire IP address space. With more than half the class B networks assigned and a little over 6 percent of the class C networks assigned, and the current growth in the address allocation rate, this is a real, though not imminent danger.

Clearly, if nothing is done to cope with these problems (especially the routing explosion), the globally interconnected Internet would cease to be just that — globally interconnected — as some connected institutions would not be able to attain logical connectivity to all parts of the Internet they wish to reach. If we were to run out of class B addresses, it would probably lead to more class C network numbers to be allocated, and this would clearly aggravate the routing table explosion problem, so it would have the same disastrous consequences for the Internet.

Efforts are under way to tackle these problems both on a short-term and long-term basis. For the long-term basis, the IAB/IETF have launched work to specify, standardize, develop, and deploy a new internet layer protocol to replace IP version 4 (the current version of IP) on the Internet. This work is currently under way, but it will still be a few years until we can expect to see this in commercial off-the-shelf technology and thus be able to deploy this protocol. The issues surrounding this long-term effort are largely independent of the issues surrounding the short-term solutions, which are the main focus of this paper, so I will not go into any more detail about the proposed long-term solutions here.

Of the problems mentioned above, the exhaustion of the class B address space and the Internet routing explosion are of most immediate concern. If the rate of allocation of class B addresses had continued to grow as it did around January 1993, we would have run out of class B network addresses after some 15 months[9]. Also, the current NSFnet routers (currently serving as one of the main "centers of gravity" for the Internet) are able to handle on the order of 25,000 individual network entries in their routing tables[2], see page 9 in [5]. The current NSFnet (end January 1994) has on the order of 16,000 con-

[1] SINTEF RUNIT, Havard.Eidnes@runit.sintef.no

[2] For the technically inclined, the limitation in the case of the NSFnet is really on the forwarding tables on the interface cards.

nected networks, and the number of administratively connected networks are in the order of 21,500 (i.e. approximately 5,500 networks were never announced to the NSFnet), see page 7 in [6]. Other major transit internet service providers use routers from other vendors, but they appear to have similar restrictions in their routing table sizes. With a monthly growth in the number of connected and announced networks of around 5–10%, something urgently needs to be done about this scaling problem.

II. Classless routing

The proposed solution to cope with the growth of the Internet on a short-term basis is through the deployment of two mechanisms:

- Development of new routing protocols for interdomain routing, implementing CIDR (short for "Classless Inter-Domain Routing"). Interdomain in this respect means essentially "between Internet network service providers".

- Start to assign IP network numbers based on the internet topology. This is done to be able to make effective use of CIDR routing technology.

As noted in the sidebar on page 3, routers could previously make routing decisions by extracting the "natural" network portion of an IP address and looking that up. This forces some IP routers to keep track of each and every network number connected on the Internet, and this is now straining the resources of many routers on the Internet.

A scheme has been agreed where the the concept of "class" for an IP network address is abolished, and instead routing decisions are to be taken based on a variable-length (contiguous) prefix of the IP address, thus the term "classless routing". It is hoped that this can significantly reduce the current top-level routing tables and stem their current exponential growth. The basic idea is that a network service provider which earlier announced 256 contiguous class C networks to the Internet (each with an implicit 24 bit prefix and with 1 to 254 in the third octet) could get by with only announcing a single prefix for all these networks (with only 16 significant bits). Such a prefix covering more than one traditional IP network number is commonly referred to as an *aggregate*. Also, if a network service provider "owned" 256 networks, he could announce the 16-bit prefix covering all these network numbers to the Internet, and then proceed to add customers to his network without modifying his routing announcements to the Internet at large.[3]

[3] This has the nice side-effect of also reducing the amount of "route flaps", i.e. routing announcements for individual networks which fluctuate between being reachable and unreachable.

To make this all work, network addresses have to be assigned to organizations/users based on where on the Internet they are connected. This is a major change from before, where addresses were assigned with no regard to network topology or efficiency of routing in mind. Thus, the policy for IP network number assignment has relatively recently been changed — see [10] and [11]. This policy change has as one effect that instead of assigning a single class B network address almost to anyone asking for it, such an assignment now has to be justified with appropriate documentation. In most cases, new organizations will instead receive a more appropriate-sized block of class C network numbers. In order to facilitate topology-based IP network number assignment, blocks of class C network numbers have been delegated to internet network service providers, and they can allocate sub-blocks of approriate size to their own current and prospective customers. Addresses assigned in this way, i.e. by an internet network service provider, and connected through this provider are said to be *topologically significant*.

With the new allocation strategy and no deployment of CIDR, we will see a faster increase in the total number of "classfull" routes — the reason is that the number of class C networks will rise quicker than before, since fewer class B numbers will be assigned. After CIDR deployment, however, the new allocation strategy will lead to greater benefits through the ability to construct larger aggregates and thereby reduce the number of routing table entries for the major internet transit service providers.

With that bit of background in place, I wish to turn to the main topic of this paper, which is how you as a network manager should behave under these new address allocation procedures, and how you should plan your local address assignments.

III. Apply through your provider

Under the new address allocation procedures it is of paramount importance that you choose to use topologically significant addresses if you plan to connect to the Internet (now or later). This is most easily done by getting your IP addresses from your internet service provider. If you have no immediate plans for connecting to the Internet, you have several options:

- Try to find out who would most likely be your internet service provider if you were to connect to the Internet later, and try to get IP addresses from them. Most service providers should be able and willing to hand out addresses to such prospective customers, even though you do not have immediate plans to connect to them or the Internet.

IP addresses

IP addresses are 32 bits wide, and are normally written as four decimal numbers (each representing a single octet) separated by dots, as in `129.241.1.5`. IP addresses are split in what is called the "host part" and the "network part", and traditionally all hosts having addresses with the same network part are connected to the same physical network (e.g. LAN). It is thus possible for routers to keep track of where to send a packet by inspecting only the network part of the destination address in the packet.

The size of the network and host parts of the IP addresses vary between what is referred to as class A, B, and C networks. The sizes of the host and network parts are shown in figure I.1. (Each of the four small squares represent 8 bits.)

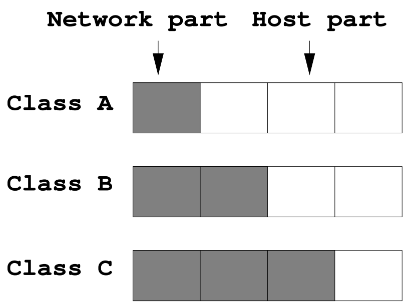

Figure I.1: Network and host part of IP addresses

In addition to class A, B, and C networks, there also exist class D network numbers which are used for IP multicast and class E and F addresses which are reserved for future use.

The division of the entire IP address space is shown in figure I.2. Note that while the class A address space takes up half of the entire address space, there are only 126 usable class A network numbers, and with the traditional routing schemes these network numbers cannot be subdivided and given to different organizations. Also note that while the class C address space is only 1/8th of the total address space, there are more than 2 million such network numbers.

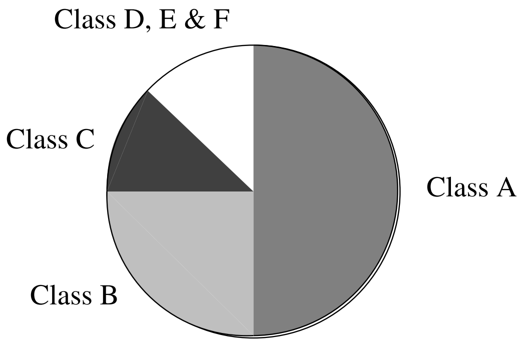

Figure I.2: Distribution of IP address space

Subnetting

As explained above, the initial model was to split IP addresses in a host and a network part. The underlying assumption is that each single physical network (e.g. LAN) would use a single network number. For some organizations this proved inadequate and also wasted a lot of addresses for those owning class B network numbers[a]. Thus, the idea of subdividing the address space at a lower level emerged, and this is done via so-called *subnetting*[13].

Subnetting is done by "extending" the network part of the address to cover part of the host portion of the address, and the bitmask covering the old network part and the "extended" network part is called the *subnet mask*. The part of the host portion of the IP address covered by the subnet mask, I will call the *subnet field*. After subnetting is put in effect for a given network, individual LANs can be given their own subnet number within a given network number. The subdivision of an IP network number into subnets is usually only visible within the subnetted network, due to limitations in older routing protocols.

[a] Few if any create single LANs containing several thousands of hosts, but instead divide their networks using routers.

- Get your IP addresses from your nearest "NIC of last resort", which is a place handing out addresses to organizations not desiring to connect to the Internet, or which do not know who they will connect through. Please note, though, that if you later should decide to connect to the Internet and have "non-provider" addresses, you will probably be asked to change your IP addresses to a set of addresses drawn from your internet network service provider's block of addresses. If you refuse to renumber your network, you may not be permitted to attain the connectivity you want. (If all organizations in the same position would refuse to renumber their networks when they connect, we are more or less back to the old way of assigning IP addresses, which is one of the important causes for the current scaling problems of the Internet.)

- If you know for certain that you will never connect to the Internet, where you get your IP addresses from is not of such great importance. Some networks will certainly never be connected to the Internet for various reasons such as security concerns.

Getting IP address space through your internet service provider is currently the only way we know to allocate topologically significant addresses. Topologically significant addresses are required for CIDR to be able to reduce the routing table sizes. Without this reduction, the danger is that the Internet will not be able to accommodate connection of new customers or user organizations. In some parts of the Internet the procedures to try to allocate topologically significant addresses have been in place for a couple of years, but it is worrying that other parts have not been equally swift in adopting these policies and have continued to assign IP network numbers with no relation to the network topology (so-called "independently assigned" network numbers). In the longer run, we will probably see that organizations already connected with independently assigned network numbers will be asked to renumber their networks to use topologically significant IP network numbers.

IV. Increase utilization

One of the problems with the previous class B allocations has been that the resulting address space utilization is much too low. If this is allowed to continue, we will likely exhaust the entire IP address space much too soon. Thus, since we now have the ability to size the address allocations more carefully, there is a desire to increase the utilization of the allocated address space. There are several techniques you as a network manager should be aware of, which is likely to play a role in achieving this goal.

IV.A. Issues

You should consider subnetting whenever possible. The Internet standard for subnetting is [13]. Some claim that it makes little sense to subnet class C networks, but I think otherwise.

First, there are the point-point links between IP routers. These physical networks will never need more than 2 IP addresses, thus allocating a whole class C network for such links is a waste of the increasingly scarce address space.

Second, there are some physical networks which will only connect a few routers — examples can be found both in LANs but this situation is perhaps more common in networks utilizing Frame Relay, SMDS or X.25, which are typically used for router-router interconnections. Usually, such interconnections are made between a fairly small number of routers, so there is address space to be saved by using subnetting.

Third, there are LANs where hosts will be directly connected. It is possible that the claim that it makes little sense to subnet class C networks is only meant to cover this case. Even if this was the case, I would still disagree. In many cases I have come across in my work, people actually plan to connect small branch offices or similar small organizations with a small number of employees to the network. If the organization only has say, 10 employees, and there are no plans to expand this significantly over the next couple of years, it makes little sense to allocate a full class C network number for this local network. In fact, as long as the estimated need for address space after two years (considering not only hosts but also routers, printers, terminal servers etc.) is less than 60, you can still subnet a class C and squeeze the utilization of your allocated address space higher. If the estimates call for more than 60 end systems after two years, it does however make more sense to allocate a whole class C network to such installations.

There are a couple of important issues you should take into consideration when doing subnetting of class C networks.

The first thing to be aware of when subnetting class C networks is the fact that when you subnet, subnets and hosts numbered "0" and "-1" (i.e. all zero or one bits in the subnet or host fields in the IP address) are reserved for special purposes in IP (broadcast and "this host/subnet" indication), and cannot be assigned to physical networks or hosts (which use that subnet mask — more on this later). Note that this precludes dividing a single class C network number into only two subnets, since both subnets in that case would have either all one or zero bits in the (single-bit) subnet field. Thus, the possible subnets of a class C network are as shown in Table 1. To give an example, the host field address ranges usable for hosts/routers in the size-62 subnets of a class C network are 65-126 and 129-190 — 0-63 cannot be used with that subnet mask since the subnet field in that case would consist of all zeroes. The usable address space in this example is illustrated in Figure 1.

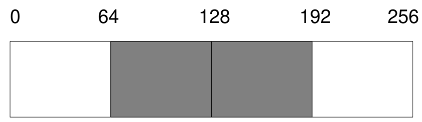

Figure 1: Usable address space with size-62 subnets.

(The `0x` prefix in table 1 is syntax borrowed from the C programming language, and is used to indicate that the following number is in base 16, i.e. hexadecimal.)

IV.B. Routing Protocols Implications

The second issue has to do with how you distribute subnets within your network. If you use a simple (if

Subnet Address Space	Usable Addresses/ subnet	Mask width	Subnet-mask
4	2	30	`0xfffffffc`
8	6	29	`0xfffffff8`
16	14	28	`0xfffffff0`
32	30	27	`0xffffffe0`
64	62	26	`0xffffffc0`

Table 1: Possible sizes for class C subnets

not simplistic. . .) routing protocol, such as e.g. RIP, you have to keep all the subnets within a subnetted network connected together with routers. What you cannot do is to connect together two subnets via (parts of) another network. Under these conditions the following address assignments are illegal:

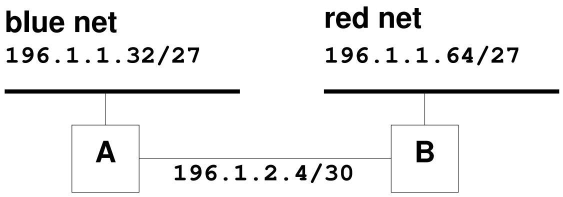

Figure 2: Illegal configuration with RIP as routing protocol

The reason is that RIP does not send subnet routing information to other networks. This is because RIP implicitly assumes that subnets are only visible within the network being subnetted, which used to be the case. In the example, router A would only receive a RIP update for 196.1.1.0 over the link from B, but since this network is also directly connected to A, A would ignore the routing update (and vice versa when A sends a routing update to B). The end result is that hosts connected to the blue net would be unable to communicate with hosts connected to the red net.

In order to correct this configuration, one would have to assign a whole size-30 subnet of 196.1.1.0 for the link between A and B. This is again due to a shortcoming in routing protocols such as RIP, which implicitly assumes that you use the same subnet mask all over a single subnetted IP network. In this case it would waste valuable address space, since 28 of the 30 valid IP addresses on the point-to-point link could never be used.

Thus, in some situations when using routing protocols such as RIP, you will not be able to assign some subnets: At the place where you want to use it, the "parent" network may not be present to preserve the continuity of the network covered by the address space of a single IP network number.

Newer routing protocols for IP (such as OSPF [14], integrated IS-IS [2] or RIP-2 [12]) remove both of these limitations. The reason is that every destination announced with these routing protocols carry with it a corresponding network mask. Thus, if you are able to use any of these newer routing protocols, you are encouraged to do so since it gives you considerably more flexibility when it comes to address assignment for subnets. There are also other reasons for using newer routing protocols, see e.g.[3].

IV.C. Variable Length Subnet Masks

If you subnet a class C network into two size-62 subnets (a 26 bit subnet mask — `0xffffffc0`), you would normally waste half of that class C network's address space since you can't use subnets 0 and -1 (under that network mask). Figure 3 illustrates the immediately usable address space with such a subnet mask (the dark shaded area).

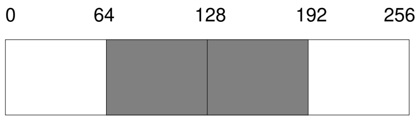

Figure 3: Usable address space with size-62 subnets.

This is also an area where a mask-based routing protocol can help increase the address space utilization, since it opens up for use of so-called variable-length subnet masks (VLSMs). With VLSMs you can use the address space otherwise taken up as subnet 0 under a 26 bit wide subnet mask (`0xffffffc0`) and subnet that address space with a wider subnet mask, say with a 29 bit wide subnet mask (`0xfffffff8`), into 8 subnets with 8 addresses each. Note that under this subnet mask, you still cannot use the subnets (or host fields) of 0 and -1, so you end up with 7 usable subnets each with room for 6 hosts each under this subnet mask for this part of the address space. You can do similarly for the address space taken up as subnet "-1" under the 26 bit wide subnet mask.[4] The address space utilized under this scheme is illustrated in Figure 4 – the lightly shaded areas are the size-6 subnets.

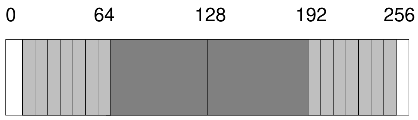

Figure 4: Using variable-length subnet masks to increase address space utilization.

[4] The reason why this becomes 7 usable subnets and not 6 is that the subnet x.x.x.56 - x.x.x.63 does not have -1 in the subnet field under a 29 bit wide subnet mask (`0xfffffff8`) — rather the subnet field becomes `0x38` (decimal 56), which has two leading zero bits.

At the current time, there is a certain lack of specific documentation available on the issues surrounding the use of VLSMs. Specifically, there is no single RFC which contains a substantially complete specification of the rules for the use of VLSMs. When it comes to the address assignment where VLSMs can be used, [15] specifies a technique which can reduce the need to renumber hosts and/or subnets, but the implications for routing are not discussed much there. The OSPF specification [14] also contains mention of VLSMs, but defers the rules for address assignment under VLSMs to another as yet non-existing document. However, this issue has received some treatment within the IETF, and it appears that if you obey the following rules, you will be on the safe side, yet still be able to utilize the address space efficiently:

- Address space where the subnet field is either 0 or -1 under one subnet mask can be used with another (wider) subnet mask.

- Under a given subnet mask, addresses where the subnet or host fields are either 0 or -1 can not be assigned to hosts or physical networks [8], section 3.2.1.3.

- The address space assigned under one subnet mask cannot be allocated under another (wider) subnet mask. Subnets can thus only have power-of-2 size.

- With some routing protocols (e.g. OSPF) it is possible to aggregate several address/mask routes into a single address/mask. When doing such aggregation it is probably wise to avoid cases which result in an address/mask combination where the subnet field (considering both the mask and the traditional class A, B, and C distinction) ends up as either 0 or -1.

Some router vendors also provide an alternative solution to the address space wastage which is particularly troublesome for point-to-point links. The feature is called *unnumbered interfaces*, and what this means is that the interfaces connecting to the point-to-point link do not have an IP address of their own, so the link does not need a separate subnet number. Instead, a router configured for unnumbered operation on a serial link uses a so-called *router ID* which is used as the IP source address e.g. when they send out routing updates, and usually this router ID is the IP address of one of the other interfaces of the router. OSPF does have provisions for handling such links, and this feature is likely to be sanctioned as recommended by the upcoming revision of the Requirements for Internet Routers document [1], [7].

It should also be mentioned that in a lot of cases, you do not need to apply all these techniques to get a reasonably high utilization of your address space. In many cases it is sufficient to subnet some C networks, where you use the same network mask within each network (the network masks may be chosen separately for each individual C network, of course). In more complex situations, utilization of the other techniques (such as OSPF and usage of VLSMs) may be required to make efficient use of the address space.

Making efficient use of the address space also means that there will be less slack than there otherwise would be, so that if you have underestimated the need for addresses on some local networks, you may end up exceeding the assigned address space. Thus, you may end up renumbering some subnets (and consequently all hosts on the subnet), and this costs valuable time and money. Since this clearly is an undesirable situation, you should take the following precautions:

- You should be careful when planning so that you minimize the number of such unfortunate surprises. This entails among other things making sure that the estimates for the number of end systems on each local network is as realistic as possible.

- You should plan for change. One component of such planning is to try not to include information about IP addresses in static configuration files etc. Use instead a naming service where you can centralize the mapping of names to IP addresses. Further, you should try to avoid using IP addresses where names can be used instead.

However, it is still in some cases unavoidable that you will have to renumber one or more subnets some time in the future. Work in the area of lessening the burden of configuring hosts is currently underway in the IETF, I am specifically referring to the work with a dynamic host configuration protocol[4]. It is my hope that this will significantly reduce the cost associated with renumbering.

I certainly realize that assigning addresses efficiently according to all these rules can at best be challenging, and sometimes problematic. A set of tools to aid in the administration of the allocated address space is sorely needed.

V. An example

Let us now try to show how these techniques can be applied by presenting a detailed walkthrough of an example. The network used in this example is shown in Figure 5, together with the number of hosts expected to be connected within two years.

This address space usage (though not the topology) is actually extracted from the DNS of a current user of a class B network and multiplied by 2 (a guess

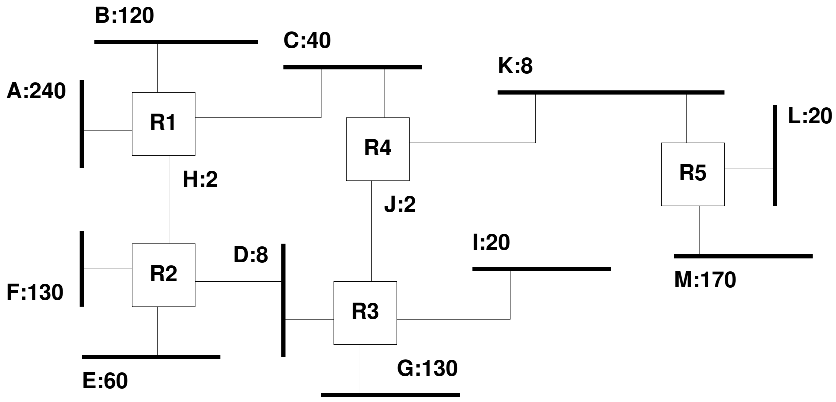

Figure 5: The example network. Number of hosts indicated for each subnet.

of 2 years growth). With the total number of hosts well below 1000 (950 to be exact), this results in address space utilization of approximately 1.4%, and clearly this is way too low to justify the allocation of a class B network number. If an organization had a class B network number with this low utilization, the right thing to do would actually be to allocate an appropriate-sized block of class C network numbers, renumber all the hosts and networks, and return the class B network number to the pool of unallocated network numbers.

Let us now look at how we can use a block of class C addresses to satisfy the address space need projected above.

The simplest method would be to allocate a single class C network to each individual physical network in the example, resulting in a need for 13 class C network numbers, and an address space utilization of 29%. This is actually quite high, especially compared to the 1.4% utilization with a class B network number. The resulting allocation would be 16 class C network numbers, though, since address space will be allocated in power-of-2-sized blocks. This results in an effective address space utilization of 23%. Even though this is quite high, it is not particularly difficult to increase this figure. I will show some alternatives to this brute-force allocation here.

Clearly, the networks A, B, F, G, and M all need an entire class C network of their own, since the largest subnet we can legally construct out of a class C network number only has room for 62 hosts. Next, there are the networks C and E which are candidates for size-62 subnets. However, these two networks are separated by a serial link (H), so the question of which routing protocol is to be used immediately appears.

Let us first take a look at how this can be solved with traditional non-mask-based routing protocols. The simple (and inefficient) method would be to allocate a single class C network for each of the networks C and E. Another alternative could be to setup the interfaces to the point-to-point link H on R1 and R2 as unnumbered interfaces, borrowing the IP addresses on C and E as router IDs. This way, we can use a single class C network, and subnet it into two size-62 subnets and hand these out to C and E. However, this solution is not without problems — consider what would happen if link H went down. The topology of the network is not physically partitioned, so all systems should still in principle be able to communicate. However, under this address assignment and routing protocol, they cannot, because we have suddenly split the network which C and E are subnets of in two. The end result is that systems on networks C and E would not be able to communicate with each other, and the hosts on each network would also be disconnected from several of the other networks in the example.

If we continue using a non-mask-based routing protocol, however, we could cure this problem by not subnetting, and instead give C and E each their own class C network number. Now we are left with the networks D, H, I, J, K, and L. All of these networks have less than 30 hosts connected to them. Thus, we can subnet a single class C into 6 usable subnets each with room for up to 30 hosts, and hand out the subnets to these physical networks. However, this solution still has the unfortunate side-effect that if the point-to-point link J goes down, the subnets D, H, and I become separated by another network (network C in the figure) from the subnets K and L, so while the physical topology still isn't partitioned, some systems would be unable to communicate.

Thus, without resorting to any special techniques (and no mask-based routing), we need 8 class C networks to assign addresses to the networks in the figure. This results in an address space utilization of 47%, which I consider to be good enough in this example. By deploying a mask-based routing protocol such as OSPF, you may reduce the needed address space down to 7 class C networks (by subnetting for E and C again, as described above), and drive the address space utilization up to 53%. At the same time

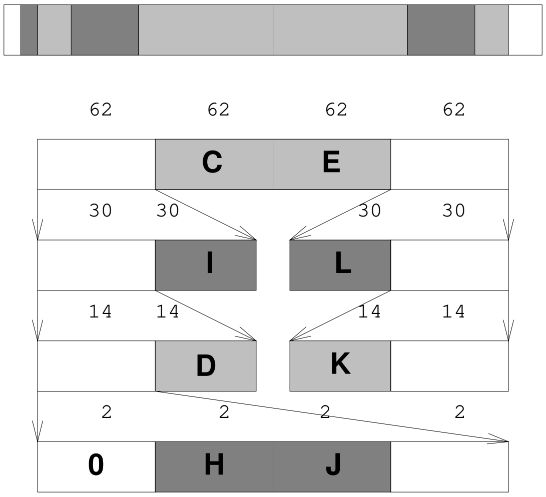

Figure 6: Using a single class C network number for 8 different physical networks, showing both placement in address space and an "exploded" view. Figures represent address space for each subnet.

you also correct the problem that even though the network isn't physically partitioned, it will be logically partitioned.

If the network in the figure was your whole network, you would actually be allocated 8 class C networks even though you only needed 7, so the actual address space utilization would still be 47% in this case. Both 47% and 53% utilization is more than enough to justify an address space allocation request, but let us for the sake of argument inspect how high we could drive the address space utilization in this example.

I now assume that OSPF is used as a routing protocol in the example network, so that we can disregard the topological placement of the individual subnets of a given network number. I will further try to utilize variable-length subnet masks to drive the utilization of the address space higher.

As before, networks A, B, F, G, and M all need their own class C network numbers. However, the remaining address space requirements can be covered with only a single class C network. The resulting assignment is depicted in Figure 6, and explained below.

Networks C and E can each be allocated a size-62 subnet of a class C network, this is under a 26 bit wide subnet mask (`0xffffffc0`). The corresponding host addresses have the following valid ranges for the last octet of the address: 65-126 and 129-190. Next, we need address space for two size-30 subnets (I and L). By utilizing VLSMs, we can use the remaining address space after allocating addresses to C and E and use a wider network mask — in this case 27 bits wide (mask `0xffffffe0`), giving the valid address ranges of 33-62 and 193-222. Next, we need two subnets of size 14 for D and K, and we can yet again use some of the remaining address space and a wider network mask — in this case 28 bits wide (mask `0xfffffff0`), and the address ranges 17-30 and 225-238 can be used. Lastly, we need two subnets for H and J of size 2, and we can again use a wider subnet mask — we use a 30 bits mask (`0xfffffffc`), giving the valid address ranges of 5-6, 9-10, 13-14, 241-242, 245-246 and 249-250. Note that we still avoid having a host or subnet field of "0" or "-1" under any given subnet mask with these allocations. Thus, we are able to use only 6 class C networks, and the resulting utilization of the address space in this case would be 62%.

Even though the gain by implementing variable-length subnet masks in this case is quite low (approximately 10%-units), the gain can be substantially larger in other cases if the number of size-62 subnets (or smaller) is larger than in the example shown here.

VI. Incentives—the stick or the carrot?

It is quite clear that there are real costs incurred by doing any of the following:

- returning class B network numbers, renumbering hosts and networks (the cost comes mostly in the form of disruption of service and the labour to do the reconfiguration)

- implementing a mask-based routing protocol will in most cases mean that you must buy or install new software, and in some cases upgrade or replace existing equipment.

- renumbering hosts and subnets internally when address assignments have to be changed, due to unexpected growth as explained above. The costs are incurred as above in disruption of service and labour of reconfiguration.

- some organizations have implemented their networks using bridges instead of (possibly multiprotocol) routers. Such networks are not always easy to handle with a multiple of class C networks unless one considers running multiple networks on the same "logical cable", which can at best be considered inefficient and sometimes problematic. The costs associated with replacing bridges with routers in such networks can be substantial.

- lastly, administrating address space under these guidelines without any automated tools to assist in this, will perhaps lead to increased administration cost (it certainly complicates matters).

Unfortunately, there are not many incentives (economic or otherwise) to make people behave for the benefit of all.

To counter the impulses to reduce these costs we could perhaps use traditional market economy mechanisms, and e.g. start charging for address space. This is however at best problematical, since address assignment authorities have a fairly widespread "customer" contact, and the increased bureaucratic burden of invoicing and accounting would be substantial. If we were to rent out address space, this would only make these matters worse. Besides, this would cause some "interesting" legal implications with respect to organizations that have already been assigned address space. One would also need to have global consensus on charging models to prevent "grey imports" of address space (which would work against the routing aggregation we are trying to achieve in the first place).

One could also consider charging for the ability to announce a separate routing entry to the Internet, but since only a fairly small portion of organizations with their own IP addresses connect to the internet (10-20%?), this will not solve the problem of total address space exhaustion.

Lastly, the Internet is not a market economy, and the way the Internet has solved similar problems before is through cooperation for the benefit of the whole community. It is, however, an open question whether the era where this technique could be applied successfully is drawing towards a close, but I think that in the short run it is our only option.

It is quite certain that the only widely available carrot we can offer as an alternative to the stick is the ability to continue to let the Internet and the usage of the Internet protocol suite grow in the coming years while a more long-term solution to the current scaling problems is being worked out.

VII. Conclusions

The Internet protocol suite is currently facing a number of scaling problems. A two-part short-term solution to some of these problems has been proposed: CIDR and modified address assignment procedures. I have outlined a number of techniques you as network manager can follow to drive the utilization of your address space upwards in order to size your address allocation request appropriately. I have also shown through an example how these techniques can be applied in practice.

References

[1] Robert Braden and Jon Postel. Requirements for internet gateways. RFC1009 – Internet standard 4, June 1987. Currently undergoing major revision within IETF.

[2] Ross Callon. Use of OSI IS-IS for routing in TCP/IP and dual environments. RFC1195 – an elective proposed Internet standard, December 1990. Currently undergoing minor revision.

[3] IAB Lyman Chapin (Chair). Applicability statement for OSPF. RFC1370 – standards track Applicability Statement, October 1992.

[4] R. Droms. Dynamic host configuration protocol. RFC1541 – Proposed Standard, October 1993.

[5] Ann Westine Cooper (ed.). Internet monthly report, September 1993. venera.isi.edu:in-notes/imr/imr9309.txt.

[6] Ann Westine Cooper (ed.). Internet monthly report, January 1994. venera.isi.edu:in-notes/imr/imr9401.txt.

[7] Frank Kastenholz (ed.). Requirements for IP routers. draft-ietf-rreq-iprouters-require-00.txt – Internet draft, December 1993. Work in progress, draft expires 26 Jun 1994.

[8] Robert Braden (ed.). Requirements for internet hosts – communication layers. RFC1122 – part of Internet standard 3 (Host Requirements), October 1989.

[9] Vincent Fuller, Tony Li, Jessica Yu, and Kenneth Varadhan. Classless inter-domain routing (CIDR): an address assignment and aggregation strategy. RFC1519 – Proposed Standard, September 1993.

[10] Elise Gerich. Guidelines for management of ip address space. RFC1466 – informational document, May 1993.

[11] Christian Huitema. Iab recommendation for an intermediate strategy to address the issue of scaling. RFC1481 – informational document, July 1993.

[12] Gary Malkin. RIP version 2 carrying additional information. RFC1388 – Proposed Standard, January 1993.

[13] Jeff Mogul and Jon Postel. Internet standard subnetting procedure. RFC950 – part of Internet Standard 5 (IP), August 1985.

[14] John Moy. OSPF version 2. RFC1247 – an elective draft Internet standard, July 1991. Currently undergoing minor revision.

[15] Paul Tsuchiya. On the assignment of subnet numbers. RFC1219 – informational document, April 1991.

Author Information

Håvard Eidnes got his M.Sc. in Computer Science from the Norwegian Institute of Technology at the University of Trondheim in 1986. He is currently employed at SINTEF, a contractual research organization with close ties to the Norwegian Institute of Technology at the University of Trondheim. Since 1988 he has worked part of his time for Uninett, the Norwegian academic and research network, as responsible for the IP network service.