CISCO SYSTEMS

IP Addressing Fundamentals

The reader-friendly explanation of how the IP address space works and how it is used

ciscopress.com

Mark A. Sportack

IP Addressing Fundamentals

Mark A. Sportack

Cisco Press

Cisco Press
201 West 103rd Street
Indianapolis, IN 46290 USA

address spaces, didn't really settle the argument. Those addresses were reserved for private networks but didn't mandate their use. With RFC 1917, the correct answer was that in the best interests of the Internet community, you shouldn't waste "real" IP addresses for an isolated network. Instead, the IETF urged you to save those addresses for use in routing across the Internet. Private IP networks, regardless of size, should use the addresses stipulated in RFC 1918 (which were previously reserved in RFC 1597, but as class-based address blocks).

The great debate had finally been settled: Don't use real IP addresses unless you need to route over the Internet. This gave the caretakers of the Internet's address space (IANA and its myriad delegate organizations) a clear directive as well as a means of enforcement.

Although that was all very well and good, and it would certainly help slow down the rate of address consumption, something still needed to be done about all the past decisions that had already been made. Quite simply, the purists of the world were sitting on a substantial number of address blocks that were being used on isolated networks! This represented a tremendous potential pool of addresses that could greatly mitigate the address crisis being experienced. The IETF and IANA would embark on an aggressive campaign to identify and reclaim unused and underused address spaces. This effort was largely successful, but, as you'll see in Chapter 13, "Planning and Managing an Address Space," it did induce some very unexpected behaviors in terms of evolving address-space management tactics!

Preserving Address Block Integrity

One of the more subtle problems plaguing the Internet during the early to mid-1990s was a side effect of address space depletion. This side effect was the rate at which the Internet's routing tables were growing. As you saw earlier in this chapter, the size of a network's routing tables are crucial to that network's end-to-end performance. Routing tables can expand for numerous reasons. In addition to the exponential rate of growth it was experiencing, the Internet's tables were expanding for three primary reasons:

- Class-based addresses still being assigned to new customers had legacy effects. RFC 1466, as you just saw, was a huge step toward minimizing such legacy effects. By converting so many class-based network addresses into classless address space, the IETF sought to make as clean a break as possible from the old, obviously inefficient, class-based IPv4 addressing.
- Customers leaving their ISPs did not return their address blocks. This forced the new ISP to support routing for those specific blocks of addresses.
- End users of the Internet were applying for their own address blocks instead of using blocks assigned to them by an ISP.

The last two items are interrelated. Rather than trying to explore them separately, it makes more sense to my twisted way of thinking to examine them from the perspective of their practical impacts on the Internet. Essentially, how well or how poorly you manage an address space is evident in how well the routes for that space aggregate or fragment. Aggregation is rolling smaller network addresses into larger network addresses without affecting routes to those networks. For example, instead of listing 10.10.1.0/24, 10.10.2.0/24, 10.10.3.0/24, and so on up to 10.10.255.0/24, you could aggregate all those network blocks into just 10.10.0.0/16. This larger block tells the rest of the world how to access all the smaller network blocks that may be defined from that /16 block.

Fragmentation is the opposite of aggregation: Network addresses are so numerically dissimilar and discontiguous as to defy aggregation. If aggregation isn't possible, the routers in the internetwork must remember a discrete route to each network. This is inefficient and can hurt the internetwork's performance.

Aggregation and fragmentation are best thought of as extremes, with infinite room for compromise in between. The politics of Internet governance have basically pitted Internet end-user organizations against ISPs in an ongoing struggle between these two idealistic extremes.

These impacts include route aggregation (and how it was damaged by rapid growth using Classical IP addressing rules) and the effects of directly registered customer address blocks. Examining these two impacts from a practical perspective will help you better appreciate how the Internet's address-space policies and rules have evolved over time and why things are the way they are.

Aggregation Versus Fragmentation

The addition of lots of new networks to the Internet would not necessarily contribute to the bloat of the Internet's routing tables. In theory, new networks (and their addresses) could be added to the Internet without increasing its routing tables if those network addresses were correlated regionally and/or by service provider. In other words, a clever approach to managing address assignment could accommodate tremendous growth with little to no impact on the Internet's routing tables simply by keeping numerically similar network addresses clumped together in a region. You could route from around the world to that region using only the highest-order bits of those related network addresses.

Although this might sound like a Utopian ideal, it is actually quite practical and realistic. To better demonstrate how such a scheme would work, consider the following example. An ISP is granted a relatively large block (let's say a Class B-sized block). It then carves this block into smaller blocks that it can assign to its customers. As far as the rest of the Internet is concerned, a single routing table entry would be required to that Class B-sized (/16) network address. This concept is known as *route aggregation*.

To see how this works, refer to Figures 5-1 and 5-2. Figure 5-1 demonstrates four ISP customers buying access to the Internet but using their own IP address space. For the sake of this example, I have used the addresses reserved in RFC 1918 instead of real addresses. The last time I used a real host address in a book, I caught an earful from the administrator of that box! Getting back to the point, even though these addresses are fictitious, they are too far apart numerically to be reliably aggregated. Thus, the ISP would have to advertise each one individually to the rest of the Internet.

Figure 5-1 *ISP Supporting Directly-Registered Customer Address Blocks*

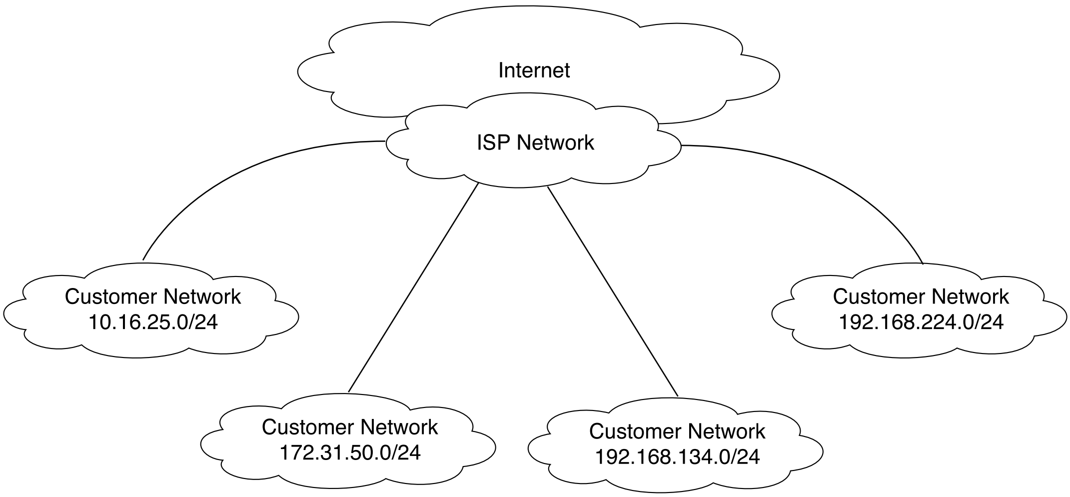

In real life, a service provider would probably have hundreds of customers, but that would make for a very cluttered illustration! Figure 5-2 shows how those customers could each use a smaller block carved from the ISP's 16-bit network address.

Figure 5-2 *Route Aggregation of an ISP Using ISP-Provided Address Blocks*

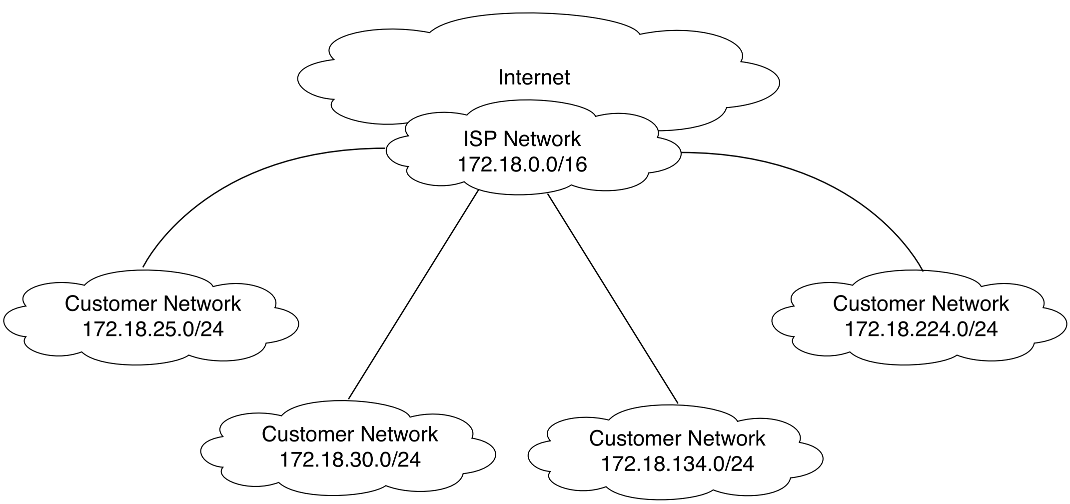

Route aggregation, in the simplistic interpretation shown in Figure 5-2, directly contributes to a reduction in the size of the Internet's routing tables. If each of the ISP's customers had its own unique network addresses, each would have to be announced to the Internet. More to the point, each network address would require its own entry in every routing table of every router in the Internet. Instead, four routing table entries can be satisfied with a single entry in the Internet, because the route through the Internet is the same for each of the four customer networks. The routes start to differ only within the network of their mutual service provider. This is the only part of the Internet that needs to track routes to the four distinct customer network addresses. Thus, the entire world uses just the first two octets of the service provider's network address to reach all its customer networks. Within that service provider's network, the third octet becomes significant for routing packets to their destination.

When viewed from this perspective, it seems perfectly logical that the right way to provide Internet addressing is via a service provider's larger address blocks. But some problems are inherent with this approach. Implementing an IP address scheme represents a potentially huge investment in planning and effort. From an Internet user's perspective, changing IP addresses is undesirable, because it forces you to make the same investment in time and effort without having anything to show for it.

Thus, although using service provider addresses does minimize routing table bloat, and makes sense for the Internet, you have to realize that benefits of this approach are asymmetrically distributed. That's a euphemistic way of saying that it is better for the Internet and ISPs than it is for the owners/operators of networks that connect to the Internet. For those entities, this approach can actually be harmful!

The danger of obtaining and using IP addresses from a service provider is that the service provider "owns" them. In effect, you are leasing the addresses for the duration of your service contract with that provider. If you wanted to change service providers (maybe you found a better deal, or maybe your provider doesn't meet your performance requirements), you would have to relinquish the "leased" addresses and obtain a new range from your new provider. Changing service providers therefore would necessitate renumbering all your IP endpoints! Renumbering becomes a very effective barrier to changing service providers.

Directly-Registered Address Spaces

It's not difficult to see why ISP customers are motivated to avoid changing their IP addresses. Renumbering endpoints is that onerous and risky a proposition! The surest way to avoid having to renumber is to "own" your IP addresses. Of course, you know that nobody really "owns" IP addresses except IANA, but ownership in this sense means that an organization would have IP addresses *registered directly in its name*. As soon as an address block is registered directly to you, it is yours forever—provided, of course, that you don't outgrow it or forget to pay the annual fee.

In the early days of the Internet, it was relatively easy to obtain directly registered address spaces. Such spaces were said to be *portable*, because they were independent of service provider address blocks. Thus, the holder of a directly registered address block enjoyed the unparalleled freedom to change service providers at will without having to worry about changing address blocks at the same time.

The drawback of this approach is, of course, the impact on the Internet's routing tables. Portability, more than any other factor, fragments large, contiguous (and therefore aggregatable) address blocks. Unfortunately for those end users, the Internet's routing tables were outpacing technology in their growth. They were becoming bigger at a faster pace than CPUs were increasing in speed. Consequently, Internet performance was deteriorating quickly, and the trend didn't look good. This was one of the facets of the impending Date of Doom that the IETF sought to obviate. One of the easy scapegoats was the portability of network addresses.

Preventing Further Fragmentation

End-user organizations highly prize portable address spaces. But they are the bane of the Internet. The IETF sought to protect the Internet by more specifically constraining address assignment practices to prevent any further address space fragmentation. This effort was documented in RFC 2050. RFC 2050 is still in effect and is also Internet Best Current Practice #12. Specifically, this document stipulated the rules and regulations regarding subassignments of IP address spaces by ISPs.

The way it works is simple. An ISP's customer—if it didn't already have directly registered address blocks of its own—could obtain addresses from its service provider. However, to preserve the integrity of large service provider address blocks, those addresses had to be surrendered when the contract for service ended. In other words, these addresses were *nonportable* and remained with the service provider with which they were registered. That way, the ISP could advertise just a single, large network address to the Internet that encompassed all its customer networks created from that large block.

But if a customer wanted to move to a different service provider to obtain a lower monthly recurring cost, it found itself in the uncomfortable position of having to quickly renumber its entire network and all its addressed endpoints. Adding insult to injury, the range it had to migrate to was another nonportable address range supplied by the new service provider. Each time the customer wanted to change providers, it had to go through the same expensive, risky, painful process of renumbering endpoints.

RFC 2050/BCP 12 doesn't do anything to mitigate this pain for end users. Rather, it compels service providers to treat all their assignable address space as nonportable, regardless of whether any given subset of it may be globally routable. If an ISP chooses to disregard RFC 2050 and let ex-customers keep their assigned space, it will eventually exhaust its address space. That ISP will find it impossible to convince its regional Registry to entrust it with more address space. An ISP without an inventory of available network addresses cannot service any new customers. RFC 2050 seeks to prevent further fragmentation of the Internet's address space (which directly increases the size of the Internet's routing tables) by giving ISPs the incentive to preserve the integrity of their existing blocks.

Rationing Directly Registered Addresses

RFC 2050, in the absence of any other policy changes, was a tweak to the nose of the Internet's end-user community. It wasn't intended to be painful. Rather, it was designed as an emergency effort to ensure the Internet's continued operability. Such austerity measures are often palatable provided that the pain is shared somewhat equitably. Because the end-user community bore the brunt of the inconvenience caused by that RFC, they had every reason to be upset. The only mitigating factor was that those organizations could shield themselves from any pain just by having their own directly registered address spaces—that is, they could until IANA started rationing new directly registered address blocks.

This is where things started getting ugly. IANA exacerbated the routability versus portability conflict when it tightened its policy on directly registered address spaces. Organizations that wanted their "own" address spaces would have to meet very stringent requirements before that privilege would be granted. The cumbersome and bureaucratic application process alone was enough to deter most would-be applicants. Those persistent enough to successfully complete an application with their Registry quickly discovered that did not guarantee their request would be granted.

Although this policy shift was a necessity caused by the impending address-space crisis, it meant that it was now almost impossible for an end-user organization to obtain its own directly registered address space. When you view this fact in conjunction with the bias against end-user organizations in RFC 2050, it becomes clear that end-user organizations bore the full brunt of the policy changes necessary to prevent the Internet's address space from collapsing.

Ostensibly, this policy shift was designed to immediately curtail the outflow of the finite supply of the remaining IPv4 addresses. This, all by itself, would buy a lot of time and forestall the crisis. However, it also forced Internet users to seriously consider alternative options for their IP addressing, such as the following:

- Using ISP-provided addresses
- Using nonunique addresses
- Using a translation function between your network and the Internet

We've already looked at why using ISP-provided addressing isn't very attractive to end-user organizations. We'll look at the other two options much more closely in the next chapter. For now, just realize that none of these were very palatable, nor as convenient as directly registered address space. The fact that these unattractive options were forced on end-user organizations created the potential for a tremendous backlash that is still being experienced today.

End-User Backlash

In some cases, the backlash has been extremely aggressive and creative. The goal is simple, and it can be summed up in just one word: portability. The extent to which end-user