

Network Working Group
Request for Comments: 2071
Category: Informational

P. Ferguson
cisco Systems, Inc.
H. Berkowitz
PSC International
January 1997

Network Renumbering Overview:
Why would I want it and what is it anyway?

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

The PIER [Procedures for Internet/Enterprise Renumbering] working group is compiling a series of documents to assist and instruct organizations in their efforts to renumber. However, it is becoming apparent that, with the increasing number of new Internet Service Providers (ISP's) and organizations getting connected to the Internet for the first time, the concept of network renumbering needs to be further defined. This document attempts to clearly define the concept of network renumbering and discuss some of the more pertinent reasons why an organization would have a need to do so.

Table of Contents

| | | |
|----|--------------------------------------|----|
| 1. | Introduction | 2 |
| 2. | Background | 2 |
| 3. | Network Renumbering Defined. | 3 |
| 4. | Reasons for Renumbering. | 3 |
| 5. | Summary. | 12 |
| 6. | Security Considerations | 12 |
| 7. | Acknowledgments. | 12 |
| 8. | References | 13 |
| 9. | Authors' Addresses | 14 |

1. Introduction

The popularity of connecting to the global Internet over the course of the past several years has spawned new problems; what most people casually refer to as "growing pains" can be attributed to more basic problems in understanding the requirements for Internet connectivity. However, the reasons why organizations may need to renumber their networks can greatly vary. We'll discuss these issues in some amount of detail below. It is not within the intended scope of this document to discuss renumbering methodologies, techniques, or tools.

2. Background

The ability for any network or interconnected devices, such as desktop PCs or workstations, to obtain connectivity to any potential destination in the global Internet is reliant upon the possession of unique IP host addresses [1]. A duplicate host address that is being used elsewhere in the Internet could best be described as problematic, since the presence of duplicate addresses would cause one of the destinations to be unreachable from some origins in the Internet. It should be noted, however, that globally unique IP addresses are not always necessary, and is dependent on the connectivity requirements [2].

However, the recent popularity in obtaining Internet connectivity has made these types of connectivity dependencies unpredictable, and conventional wisdom in the Internet community dictates that the various address allocation registries, such as the InterNIC, as well as the ISP's, become more prudent in their address allocation strategies. In that vein, the InterNIC has defined address allocation policies [3] wherein the majority of address allocations for end-user networks are accommodated by their upstream ISP, except in cases where dual- or multihoming and very large blocks of addresses are required. With this allocation policy becoming standard current practice, it presents unique problems regarding the portability of addresses from one provider to another.

As a practical matter, end users cannot assume they "own" address allocations, if their intention is to be to have full connectivity to the global Internet. Rather, end users will "borrow" part of the address space of an upstream provider's allocation. The larger provider block from which their space is suballocated will have been assigned in a manner consistent with global Internet routing.

Not having "permanent" addresses does not mean users will not have unique identifiers. Such identifiers are typically Domain Name System (DNS) [4] names for endpoints such as servers and workstations. Mechanisms such as the Dynamic Host Configuration Protocol (DHCP) [5]

can help automate the assignment and maintenance of host names, as well as the 'borrowed' addresses required for routing-level connectivity.

The PIER Working Group is developing procedures and guidelines for detailed renumbering of specific technologies, such as routers [6]. PIER WG documents are intended to suggest methods both for making existing networks prepared for convenient renumbering, as well as for operational transition to new addressing schemes.

Also, in many instances, organizations who have never connected to the Internet, yet have been using arbitrary blocks of addresses since their construction, have different and unique challenges.

3. Network Renumbering Defined

In the simplest of definitions, the exercise of renumbering a network consists of changing the IP host addresses, and perhaps the network mask, of each device within the network that has an address associated with it. This activity may or may not consist of all networks within a particular domain, such as FOO.EDU, or networks which comprise an entire autonomous system.

Devices which may need to be renumbered, for example, are networked PC's, workstations, printers, file servers, terminal servers, and routers. Renumbering a network may involve changing host parameters and configuration files which contain IP addresses, such as configuration files which contain addresses of DNS and other servers, addresses contained in SNMP [7] management stations, and addresses configured in access control lists. While this is not an all-inclusive list, the PIER working group is making efforts to compile documentation to identify these devices in a more detailed fashion.

Network renumbering need not be sudden activity, either; in most instances, an organization's upstream service provider(s) will allow a grace period where both the "old" addresses and the "new" addresses may be used in parallel.

4. Reasons for Renumbering

The following sections discuss particular reasons which may precipitate network renumbering, and are not presented in any particular order of precedence. They are grouped into reasons that primarily reflect decisions made in the past, operational requirements of the present, or plans for the future.

Some of these requirements reflect evolution in the organization's mission, such as a need to communicate with business partners, or to work efficiently in a global Internet. Other requirements reflect changes in network technologies.

4.1 Past

Many organizations implemented IP-based networks not for connectivity to the Internet, but simply to make use of effective data communications mechanisms. These organizations subsequently found valid reasons to connect to other organizations or the Internet in general, but found the address structures they chose incompatible with overall Internet practice.

Other organizations connected early to the Internet, but did so at a time when address space was not scarce. Yet other organizations still have no requirement to connect to the Internet, but have legacy addressing structures that do not scale to adequate size.

4.1.1 Initial addressing using non-unique addresses

As recently as two years ago, many organizations had no intention of connecting to the Internet, and constructed their corporate or organizational network(s) using unregistered, non-unique network addresses. Obviously, as most problems evolve, these same organizations determined that Internet connectivity had become a valuable asset, and subsequently discovered that they could no longer use the same unregistered, non-unique network addresses that were previously deployed throughout their organization. Thus, the labor of renumbering to valid network addresses is now upon them, as they move to connect to the global Internet.

While obtaining valid, unique addresses is certainly required to obtain full Internet connectivity in most circumstances, the number of unique addresses required can be significantly reduced by the implementation of Network Address Translation (NAT) devices [8] and the use of private address space, as specified in [9]. NAT reduces not only the number of required unique addresses, but also localizes the changes required by renumbering.

It should also be noted that NAT technology may not always be a viable option, depending upon scale of addressing, performance or topological constraints.

4.1.2 Legacy address allocation

There are also several instances where organizations were originally allocated very large amounts of address space, such as traditional "Class A" or "Class B" allocations, while the actual address requirements are much less than the total amount of address space originally allocated. In many cases, these organizations could suffice with a smaller CIDR allocation, and utilize the allocated address space in a more efficient manner. As allocation requirements become more stringent, mechanisms to review how these organizations are utilizing their address space could, quite possibly, result in a request to return the original allocation to a particular registry and renumber with a more appropriately sized address block.

4.1.3 Limitations of Bridged Internetworks

Bridging has a long and distinguished history in legacy networks. As networks grow, however, traditional bridged networks reach performance- and stability-related limits, including (but not limited to) broadcast storms.

Early routers did not have the speed to handle the needs of some large networks. Some organizations were literally not able to move to routers until router forwarding performance improved to be comparable to bridges. Now that routers are of comparable or superior speed, and offer more robust features, replacing bridged networks becomes reasonable.

IP addresses assigned to pure bridged networks tend not to be subnetted, yet subnetting is a basic approach for router networks. Introducing subnetting is a practical necessity in moving from bridging to routing.

Special cases of bridging are realized in workgroup switching systems, discussed below.

4.1.4 Limitations of Legacy Routing Systems

Other performance problems might come from routing mechanisms that advertise excessive numbers of routing updates (e.g., RIP, IGRP). Likewise, appropriate replacement protocols (e.g., OSPF, EIGRP, S-IS) will work best with a structured addressing system that encourages aggregation.

4.1.5 Limitations of System Administration Methodologies

There can be operational limits to growth based on the difficulty of adds, moves and changes. As enterprise networks grow, it may be necessary to delegate portions of address assignment and maintenance. If address space has been assigned randomly or inefficiently, it may be difficult to delegate portions of the address space.

It is not unusual for organizational networks to grow sporadically, obtaining an address prefix here and there, in a non-contiguous fashion. Depending on the number of prefixes that an organization acquires over time, it may become increasingly unmanageable or demand higher levels of maintenance and administration when individual prefixes are acquired in this way.

Reasonable IP address management may in general simplify continuing system administration; a good numbering plan is also a good renumbering plan. Renumbering may force a discipline into system administration that will reduce long-term support costs.

It has been observed "...there is no way to renumber a network without an inventory of the hosts (absent DHCP). On a large network that needs a database, plus tools and staff to maintain the database." [10] It can be argued that a detailed inventory of router configurations is even more essential.

4.2 Present

Organizations now face needs to connect to the global Internet, or at a minimum to other organizations through bilateral private links.

Certain new transmission technologies have tended to redefine the basic notion of an IP subnet. An IP numbering plan needs to work with these new ideas. Legacy bridged networks and leading-edge workgroup switched networks may very well need changes in the subnetting structure. Renumbering needs may also develop due to the characteristics of new WAN technologies, especially nonbroadcast multi-access (NBMA) services such as Frame-Relay and Asynchronous Transfer Mode (ATM).

Increased use of telecommuting by mobile workers, and in small and home offices, need on-demand WAN connectivity, using modems or ISDN. Effective use of demand media often requires changes in numbering and routing.

4.2.1 Change in organizational structure or network topology

As companies grow, through mergers, acquisitions and reorganizations, the need may arise for realignment and modification of the various organizational network architectures. The connectivity of disparate corporate networks present unique challenges in the realm of renumbering, since one or more individual networks may have to be blended into a much larger architecture consisting a different IP address prefix altogether.

4.2.2 Inter-Enterprise Connectivity

Even if they do not connect to the general Internet, enterprises may interconnect to other organizations which have independent numbering systems. Such connectivity can be as simple as bilateral dedicated circuits. If both enterprises use unregistered or private address space, they run the risk of using duplicate addresses.

In such cases, one or both organizations may need to renumber into different parts of the private address space, or obtain unique registered addresses.

4.2.3 Change of Internet Service Provider

As mentioned previously in Section 2, it is increasingly becoming current practice for organizations to have their IP addresses allocated by their upstream ISP. Also, with the advent of Classless Inter Domain Routing (CIDR) [11], and the considerable growth in the size of the global Internet table, Internet Service Providers are becoming more and more reluctant to allow customers to continue using addresses which were allocated by the ISP, when the customer terminates service and moves to another ISP. The prevailing reason is that the ISP was previously issued a CIDR block of contiguous address space, which can be announced to the remainder of the Internet community as a single prefix. (A prefix is what is referred to in classless terms as a contiguous block of IP addresses.) If a non-customer advertises a specific component of the CIDR block, then this adds an additional routing entry to the global Internet routing table. This is what is commonly referred to as "punching holes" in a CIDR block. Consequently, there are usually no routing anomalies in doing this since a specific prefix is always preferred over an aggregate route. However, if this practice were to happen on a large scale, the growth of the global routing table would become much larger, and perhaps too large for current backbone routers to accommodate in an acceptable fashion with regards to performance of recalculating routing information and sheer size of the routing table itself. For obvious reasons, this practice is highly discouraged by ISP's with CIDR blocks, and some ISP's are making this a contractual

issue, so that customers understand that addresses allocated by the ISP are non-portable.

It is noteworthy to mention that the likelihood of being forced to renumber in this situation is inversely proportional to the size of the customer's address space. For example, an organization with a /16 allocation may be allowed to consider the address space "portable", while an organization with multiple non-contiguous /24 allocations may not. While the scenarios may be vastly different in scope, it becomes an issue to be decided at the discretion of the initial allocating entity, and the ISP's involved; the major deciding factor being whether or not the change will fragment an existing CIDR block and whether it will significantly contribute to the overall growth of the global Internet routing tables.

It should also be noted that (contrary to opinions sometimes voiced) this form of renumbering is a technically necessary consequence of changing ISP's, rather than a commercial or political mandate.

4.2.3 Internet Global Routing

Even large organizations, now connected to the Internet with "portable" address space, may find their address allocation too small. Current registry guidelines require that address space usage be justified by an engineering plan. Older networks may not have efficiently utilized existing address space, and may need to make their existing structures more efficient before new address allocations can be made.

4.2.4 Internal Use of LAN Switching

Introducing workgroup switches may introduce subtle renumbering needs. Fundamentally, workgroup switches are specialized, high-performance bridges, which make their main forwarding decisions based on Layer 2 (MAC) address information. Even so, they rarely are independent of Layer 3 (IP) address structure. Pure Layer 2 switching has a "flat" address space that will need to be renumbered into a hierarchical, subnetted space consistent with routing.

Introducing single switches or stacks of switches may not have significant impact on addressing, as long as it is understood that each system of switches is a single broadcast domain. Each broadcast domain should map to a single IP subnetwork.

Virtual LANs (VLANs) further extend the complexity of the role of workgroup switches. It is generally true that moving an end station from one switch port to another within the same VLAN will not cause major changes in addressing. Many overview presentations of this

technology do not make it clear that moving the same end station between different VLANs will move the end station into another IP subnet, requiring a significant address change.

Switches are commonly managed by SNMP applications. These network management applications communicate with managed devices using IP. Even if the switch does not do IP forwarding, it will itself need IP addresses if it is to be managed. Also, if the clients and servers in the workgroup are managed by SNMP, they will also require IP addresses. The workgroup, therefore, will need to appear as one or more IP subnetworks.

Increasingly, internetworking products are not purely Layer 2 or Layer 3 devices. A workgroup switch product often includes a routing function, so the numbering plan must support both flat Layer 2 and hierarchical Layer 3 addressing.

4.2.4 Internal Use of NBMA Cloud Services

"Cloud" services such as frame relay often are more economical than traditional services. At first glance, when converting existing enterprise networks to NBMA, it might appear that the existing subnet structure should be preserved, but this is often not the case.

Many organizations often began by treating the "cloud" as a single subnet, but experience has shown it is often better to treat the individual virtual circuits as separate subnets, which appear as traditional point-to-point circuits. When the individual point-to-point VCs become separate subnets, efficient address utilization requires the use of long prefixes (i.e., 30 bit) for these subnets. In practice, obtaining 30 bit prefixes means the logical network should support variable length subnet masks (VLSM). VLSMs are the primary method in which an assigned prefix can be subnetted efficiently for different media types. This is accomplished by establishing one or more prefix lengths for LAN media with more than two hosts, and subdividing one or more of these shorter prefixes into longer /30 prefixes that minimize address loss.

There are alternative ways to configure routing over NBMA, using special mechanisms to exploit or simulate point-to-multipoint VCs. These often have a significant performance impact, and may be less reliable because a single routing point of failure is created. Motivations for such alternatives tend to include:

1. A desire not to use VLSM. This is often founded in fear rather than technology.
2. Router implementation issues that limit the number of subnets or interfaces a given router can support.
3. An inherently point-to-multipoint application (e.g., remote hosts to a data center). In such cases, some of the limitations are due to the dynamic routing protocol in use. In such "hub-and-spoke" implementations, static routing can be preferable from a performance and flexibility standpoint, since it does not produce routing protocol chatter and is unaffected by split horizon constraints (namely, the inability to build an adjacency with a peer within the same IP subnetwork).

4.2.5 Expansion of Dialup Services

Dialup services, especially public Internet access providers, are experiencing explosive growth. This success represents a particular drain on the available address space, especially with a commonly used practice of assigning unique addresses to each customer.

In this case, individual users announce their address to the access server using PPP's IP control protocol (IPCP) [12]. The server may validate the proposed address against some type of user identification, or simply make the address active in a subnet to which the access server (or set of bridged access servers) belongs.

The preferred technique is to allocate dynamic addresses to the user from a pool of addresses available to the access server.

4.2.6 Returning non-contiguous prefixes for an aggregate

In many instances, an organization can return their current, non-contiguous prefix allocations for a contiguous block of address space of equal or greater size, which can be accommodated with CIDR. Also, many organizations have begun to deploy classless interior routing protocols within their domains that make use of route summarization and other optimized routing features, effectively reducing the total number of routes being propagated within their internal network(s), and making it much easier to administer and maintain.

Hierarchical routing protocols such as OSPF scale best when the address assignment of a given network reflects the topology, and the topology of the network can often be fluid. Given that the network is fluid, even the best planned address assignment scheme, given time, will diverge from the actual topology. While not required, some

organization may choose to gain the benefit of both technical and administrative scalability of their IGP by periodically renumbering to have address assignments reflect the network topology. Patrick Henry once said "the tree of liberty must from time to time be watered with the blood of patriots." In the Internet, routing trees of the best-planned networks need from time to time be watered with at least the sweat of network administrators. Improving aggregation is also highly encouraged to reduce the size of not only the global Internet routing table, but also the size and scalability of interior routing within the enterprise.

4.3 Future

Emerging new protocols will most definitely affect addressing plans and numbering schemes.

4.3.1 Internal Use of Switched Virtual Circuit Services

Services such as ATM virtual circuits, switched frame relay, etc., present challenges not considered in the original IP design. The basic IP decision in forwarding a packet is whether the destination is local or remote, in relation to the source host's subnet. Address resolution mechanisms are used to find the medium address of the destination in the case of local destinations, or to find the medium address of the router in the case of remote routers.

In these new services, there are cases where it is far more effective to "cut-through" a new virtual circuit to the destination. If the destination is on a different subnet than the source, the cut-through typically is to the egress router that serves the destination subnet. The advantage of cut-through in such a case is that it avoids the latency of multiple router hops, and reduces load on "backbone" routers. The cut-through decision is usually made by an entry router that is aware of both the routed and switched environments.

This entry router communicates with a address resolution server using the Next Hop Resolution Protocol (NHRP) [13]. This server maps the destination network address to either a next-hop router (where cut-through is not appropriate) or to an egress router reached over the switched service. Obviously, the data base in such a server may be affected by renumbering. Clients may have a hard-coded address of the server, which again may need to change. While the NHRP protocol specifications are still evolving at the time of this writing, commercial implementations based on drafts of the protocol standard are in use.

4.3.2 Transitioning to IP version 6

Of course, when IPv6 [14] deployment is set in motion, and as methodologies are developed to transition to IPv6, renumbering will also be necessary, but perhaps not immediately mandatory. To aid in the transition to IPv6, mechanisms to deploy dual- IPv4/IPv6 stacks on network hosts should also become available. It is also envisioned that Network Address Translation (NAT) devices will be developed to assist in the IPv4 to IPv6 transition, or perhaps supplant the need to renumber the majority of interior networks altogether, but that is beyond the scope of this document. At the very least, DNS hosts will need to be reconfigured to resolve new host names and addresses, and routers will need to be reconfigured to advertise new prefixes.

IPv6 address allocation will be managed by the Internet Assigned Numbers Authority (IANA) as set forth in [15].

5. Summary

As indicated by the Internet Architecture Board (IAB) in [16], the task of renumbering networks is becoming more widespread and commonplace. Although there are numerous reasons why an organization would desire, or be required to renumber, there are equally as many reasons why address allocation should be done with great care and forethought at the onset, in order to minimize the impact that renumbering would have on the organization. Even with the most forethought and vision, however, an organization cannot foresee the possibility for renumbering. The best advice, in this case, is to be prepared, and get ready for renumbering.

6. Security Considerations

Although no obvious security issues are discussed in this document, it stands to reason that renumbering certain devices can defeat security systems designed and based on static IP host addresses. Care should be exercised by the renumbering entity to ensure that all security systems deployed with the network(s) which may need to be renumbered be given special consideration and significant forethought to provide continued functionality and adequate security.

7. Acknowledgments

Special acknowledgments to Yakov Rekhter [cisco Systems, Inc.], Tony Bates [cisco Systems, Inc.] and Brian Carpenter [CERN] for their contributions and editorial critique.

8. References

- [1] Gerich, E., "Unique Addresses are Good", RFC 1814, IAB, July 1995.
- [2] Crocker, D., "To Be 'On' the Internet", RFC 1775, March 1995.
- [3] Hubbard, K., Kouters, M., Conrad, D., Karrenberg, D., and J. Postel, "INTERNET REGISTRY IP ALLOCATION GUIDELINES", BCP 12, RFC 2050, November 1996.
- [4] Mockapetris, P., "Domain Names - Concepts and Facilities", and "Domain Names - Implementation and Specification", STD 13, RFCs 1034, 1035, November 1987.
- [5] Droms, R., "Dynamic Host Configuration Protocol", RFC 1541, October 1993.
- [6] Berkowitz, H., "Router Renumbering Guide", RFC 2072, January 1997.
- [7] Case, J., Fedor, M., Schoffstall, M., and J. Davin, "A Simple Network Management Protocol (SNMP)", STD 15, RFC 1157, May 1990.
- [8] Egevang, K., and P. Francis, "The IP Network Address Translator (NAT)", RFC 1631, May 1994.
- [9] Rekhter, Y., Moskowitz, R., Karrenberg, D., de Groot, G-J., and E. Lear, "Address Allocation for Private Internets", RFC 1918, February 1996.
- [10] Messages to PIER list on CERN renumbering; Brian Carpenter, CERN. Available in PIER WG mailing list archives.
- [11] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, October 1993.
- [12] McGregor, G., "The PPP Internet Protocol Control Protocol (IPCP)", RFC 1332, May 1992.
- [13] Luciani, J., Katz, D., Piscitello, D., and Cole, B., "NBMA Next Hop Resolution Protocol (NHRP)", Work in Progress.
- [14] Deering, S., and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, December 1995.

[15] IAB and IESG, "IPv6 Address Allocation Management", RFC 1881, December 1995.

[16] Carpenter, B., and Y. Rekhter, "Renumbering Needs Work", RFC 1900, February 1996.

9. Authors' Addresses

Paul Ferguson
Cisco Systems, Inc.
1875 Campus Commons Road
Suite 210
Reston, VA 22091

Phone: (703) 716-9538
Fax: (703) 716-9599
EMail: pferguso@cisco.com

Howard C. Berkowitz
PSC International
1600 Spring Hill Road
Vienna, VA 22182

Phone (703) 998-5819
Fax: (703) 998-5058
EMail: hcb@clark.net

