                         Router Renumbering Guide

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

   IP addresses currently used by organizations are likely to undergo
   changes in the near to moderate term. Change can become necessary
   for a variety of reasons, including enterprise reorganization,
   physical moves of equipment, new strategic relationships, changes in
   Internet Service Providers (ISP), new applications, and the needs of
   global Internet connectivity. Good IP address management may in
   general simplify continuing system administration; a good renumbering
   plan is also a good numbering plan.   Most actions taken to ease
   future renumbering will ease routine network administration.

   Routers are the components that interconnect parts of the IP address
   space identified by unique prefixes. Obviously, they will be
   impacted by renumbering. Other interconnection devices, such as
   bridges, layer 2 switches (i.e., specialized bridges), and ATM
   switches may be affected by renumbering. The interactions of these
   lower-layer interconnection devices with routers must be considered
   as part of a renumbering effort.

   Routers interact with numerous network infrastructure servers,
   including DNS and SNMP. These interactions, not just the pure
   addressing and routing structure, must be considered as part of
   router renumbering.

Table of Contents

1. Introduction

   Organizations can decide to renumber part or all of their IP address
   space for a variety of reasons. Overall motivations for renumbering
   are discussed in [RFC2071]. This document deals with the router-
   related aspects of a renumbering effort, once the decision to
   renumber has been made.

   A renumbering effort must be well-planned if it is to be successful.
   This document deals with planning and implementation guidelines for
   the interconnection devices of an enterprise. Of these devices,
   routers have the clearest association with the IP numbering plan.

   Planning begins with understanding the problem to be solved. Such
   understanding includes both the motivation for renumbering and the
   technical issues involved in renumbering.

     1. Begin with a short and clear statement of the reason to
         renumber. Section 3 of this document discusses common
         reasons.

     2. Understand the principles of numbering in the present and
         planned environments. Section 4 reviews numbering and
         suggests a method for describing the scope of renumbering.

3.   Before the actual renumbering, it can be useful to evolve
     the current environment and current numbering to a more
     "renumbering-friendly" system. Section 5 discusses ways to
     introduce renumbering friendliness into current systems.

4.   Be aware of potential pitfalls. These are discussed in
     Section 6.

5.   Identify potential requirements for tools, discussed in
     Section 7.

6.   Evaluate the specific router mechanisms that will be affected
     by renumbering. See Sections 8 through 13.

7.   Set up a specific transition plan framework. Guidelines
     for such planning are in Section 15.

When trying to understand the interactions of renumbering on routers,
remember there different aspects to the problem, depending on the
scope of the renumbering involved. Remember that even an
enterprise-wide renumbering probably will not affect all IP addresses
visible within the enterprise, since some addresses (e.g., Internet
service providers, external business partners) are outside the
address space under the control of the enterprise.

2. Disclaimer

The main part of this document is intended to be vendor-independent.
Not all features discussed, of course, have been implemented on all
routers.   This document should not be used as a general comparison
of the richness of features of different implementations.
References here are only to those features affected by renumbering.
Some illustrative examples may be used that cite vendor-specific
characteristics. These examples do not necessarily reflect the
current status of products.

3. Motivations for Renumbering

Reasons to renumber can be technological, organizational, or both.
Technological reasons fall into several broad categories discussed
below. Just as there can be both technological and organizational
motivations for renumbering [RFC2071], there can be multiple
technological reasons.

There may not be a clear line between organizational and technical
reasons for renumbering. While networks have a charm and beauty all
their own, the organizational reasons should be defined first in
order to justify the budget for the technical renumbering. There

also may be pure technnical reasons to renumber, such as changes in
technology (e.g., from bridging to routing).

While this document is titled "Router Renumbering Guide," it
recognizes that renumbering may be required due to the initial
installation of routers in a bridged legacy network. Organizations
may have had an adequate bridging solution that did not scale with
growth. Some organizations could not able to move to routers until
router forwarding performance improved [Carpenter] to be comparable
to bridges.

Other considerations include compliance with routing outside the
organization. Routing issues here are primarily those of the global
Internet, but may also involve bilateral private links to other
enterprises.

Certain new transmission technologies have tended to redefine the
basic notion of an IP subnet. The 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 with the introduction
of new WAN technologies, especially nonbroadcast multiaccess (NBMA)
services such as frame relay. Other WAN technologies, dialup media
using modems or ISDN, also may have new routing and numbering
requirements. Switched virtual circuit services such as ATM, X.25,
or switched frame relay also interact with routing and addressing.

3.1 Internet Global Routing

Many discussions of renumbering emphasize interactions among
organizations' numbering plans and those of the global Internet
[RFC1900]. There can be equally strong motivations for renumbering
in organizations that never connect to the global Internet.

According to RFC1900, "Unless and until viable alternatives are
developed, extended deployment of Classless Inter-Domain Routing
(CIDR) is vital to keep the Internet routing system alive and to
maintain continuous uninterrupted growth of the Internet....To
contain the growth of routing information, whenever such an
organization changes to a new service provider, the organization's
addresses will have to change.

Occasionally, service providers themselves may have to change to a
new and larger block of address space. In either of these cases, to
contain the growth of routing information, the organizations
concerned would need to renumber.... If the organization does not
renumber, then some of the potential consequences may include (a)
limited (less than Internet-wide) IP connectivity, or (b) extra cost

   to offset the overhead associated with the organization's routing
   information that Internet Service Providers have to maintain, or
   both."

3.2 Bridge Limitations; 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.
   Traditional bridged networks share many of the problems of workgroup
   switches, but have additional performance problems when bridged
   connectivity extends across slow WAN links.

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

   Virtual LANs (VLAN) 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 "color" VLAN will not
   cause major changes in addressing. Many discussions of this
   technology do not make it clear that moving the same end station
   between different colors 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 need IP addresses.
   The workgroup, therefore, will need to appear as one or more IP
   subnets.

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

3.3 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. When the individual
   point-to-point VCs become separate subnets, efficient address
   utilization requires the use of /30 prefixes for these subnets. This
   typically means the addressing and routing plan must support multiple
   prefix lengths, 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 on the router, and
   may be less reliable because a single point of failure is created.
   Mechanics of these alternatives are discussed later in this section,
   but the 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 "star" applications, static routing may actually be
          preferable from performance and flexibility standpoints,
          since it does not produce routing traffic and is unaffected
          by split horizon.

   To understand how use of NBMA services affects the addressing
   structure and routers, it is worth reviewing what would appear to be
   very basic concepts of IP subnets. The traditional view is that a
   single subnet is associated with a single physical medium. All hosts
   physically connected to this medium are assumed to be able to reach
   all other hosts on the same medium, using data link level services.
   These services are medium specific: hosts connected to a LAN medium
   can broadcast to one another, while hosts connected to a point-to-
   point line simply need to transmit to the other end.

When one host desires to transmit to another, it first determines if
the destination is local or remote. A local destination is on the
same subnet and assumed to be reachable through data link services.
A remote destination is on a different subnet, and it is assumed that
router intervention is needed to reach it.

The first NBMA problem comes up when a single subnet is implemented
over an NBMA service. Frame Relay provides single virtual circuits
between hosts that have connectivity. It is quite common to design
Frame Relay services as partial meshes, where not all hosts have VCs
to all others. When the set of hosts in a partial mesh is in a
single IP subnet, partial mesh violates the local model of full
connectivity. Even when there is full meshing, a pessimistic but
reasonable operational model must consider that individual VCs do
fail, and full connectivity may be lost transiently.

There are several ways to deal with this violation, each with their
own limitations. If a specific "central" host has connectivity to N
all other hosts, that central host can replicate all frames it
receives from one host onto outgoing VCs connecting it with the (N-1)
other hosts in the subnet. Such replication usually causes an
appreciable CPU load in the replicating router.   The replicating
router also is a single point of failure for the subnet. This method
does not scale well when extended to fuller meshes within the subnet.

In a routing protocol, such as OSPF, that has a concept of designated
routers, explicit configuration usually is needed. Other problems in
using a meshed subnet is that all VCs may not have the same
performance, but the router cannot prefer individual paths within the
subnet.

One of the simplest methods is not to attempt to emulate a broadcast
medium, but simply to treat each VC as a separate subnet. This will
cause a need for renumbering. Efficient use of the address space
dictates a /30 prefix be used for the per-VC subnets. Such a prefix
often needs VLSM support in the routers.

3.4 Expansion of Dialup Services

Dialup services, especially public Internet access providers, are
undergoing 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 practice, individual users announce their address to the
access server using PPP's IP configuration option [RFC1332]. The
server may validate the proposed address against some user
identifier, or simply make the address active in a subnet to which
the access server (or set of bridged access servers) belongs.

These access server functions may be part of the software of a
"router" and thus are within the scope of this Guide.

The preferred technique [Hubbard] is to allocate dynamic addresses to
the user from a pool of addresses available to the access server.
Various mechanisms are used actually to do this assignment, and are
discussed in Section 5.5.

3.5 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) [Cansever] [Katz]. 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 work is in progress at the time of this writing,
commercial implementations based on drafts of the protocol standard
are in use.

4. Numbering and Renumbering

   What is the role of any numbering plan? To understand the general
   problem, it can be worthwhile to review the basic principles of
   routers. While most readers will have a good intuitive sense of
   this, the principles have refined in the current usage of IP.

   A router receives an inbound IP datagram on one of its interfaces,
   and examines some number of bits of the destination address. The
   sequence of bits examined by the router always begin at the left of
   the address (i.e., the most significant bit). We call this sequence
   a "prefix."

   Routing decisions are made on totalPrefix bits, which start at the
   leftmost (i.e., most significant) bit position of the IP address.
   Those totalPrefix bits may be completely under the control of the
   enterprise (e.g., if they are in the private address space), or the
   enterprise may control the lowOrderPrefix bits while the
   highOrderPrefix bits are assigned by an outside organization.

   The router looks up the prefix in its routing table (formally called
   a Forwarding Information Base). If the prefix is in the routing
   table, the router then selects an outgoing interface that will take
   the routed packet to the next hop IP address in the end-to-end route.
   If the prefix cannot be found in the routing table, the router
   returns an ICMP Destination Unreachable message to the source address
   in the received datagram.

   Assuming the prefix is found in the routing table, the router then
   transmits the datagram through the indicated outgoing interface. If
   multicast routing is in effect, the datagram may be copied and sent
   out multiple outgoing interfaces.

4.1 Categorizing the topology

   From the router renumbering perspective, renumbering impact is apt to
   be greatest in highly connected parts of "backbones," and least in
   "stub" parts of the routing domain that have a single route to the
   backbone.

```
                   Global Internet
                         ^
                         |
                         |
                 Back1------------------Back2
                    |                      |
                 +----------+          +----------+
                 |          |          |          |
                 Reg1.1------Reg1.2    Reg2.1-----Reg2.2
                 |          |          |          |
                 |          |          |          |
                 |          |          |          |
              Branch     Branch     Branch     Branch
              1.1.1 to   1.2.1 to   2.1.1 to   2.2.1 to
              1.1.N      1.2.N      2.1.N      2.2.N
```
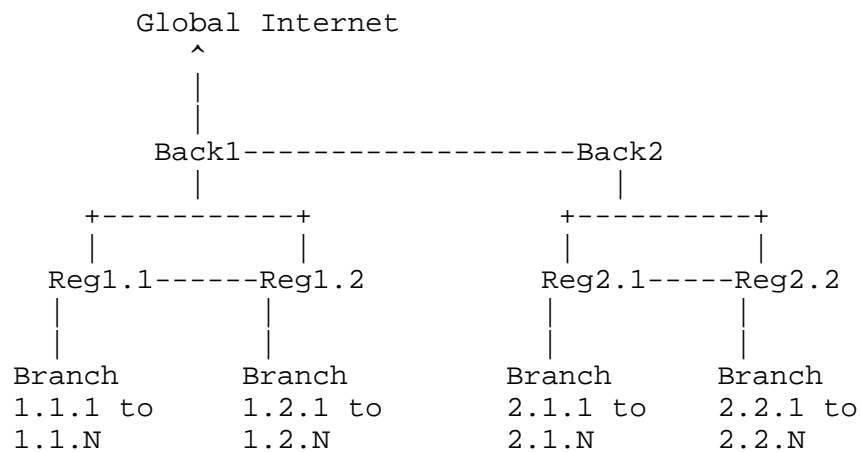
   In this drawing, assume Back1 and Back2 exchange full routes; Back1
   is also the exterior router. Regional routers (Reg) exchange full
   routes with one another and aggregate addresses to the backbone
   routers. Branch routers default to regional routers.

   From a pure topological standpoint, the higher in the hierarchy, the
   greater are apt to be the effects of renumbering. This is a first
   approximation to scoping the task, assuming addresses have been
   assigned systematically. Systematic address space is rarely the case
   in legacy networks.

4.2 Categorizing the address space

   An inventory of present and planned address space is a prerequisite
   to successful renumbering. Begin by identifying the prefixes in or
   planned into your network, and whether they have been assigned in a
   systematic and hierarchical manner.

```
    +--Unaffected by renumbering [A]
    |
    |
    +--Existing prefixes to be renumbered
    |  |
    |  |
    |  +----To be directly renumbered on "flag day"
    |  |
    |  |
    |  +----Initially to be renumbered to temporary address
    |
    |
    +--Existing prefixes to be retired
    |
    |
    +--Planned new prefixes
       |
       |
       +---totalPrefix change, no length change
       |
       |
       +---highOrderPart change only, no length change
       |
       |
       +---lowOrderPart change only, no length change
       |
       |
       +---highOrderPart change only, high length change
       |
       |
       +---lowOrderPart change only, low length change
       |
       |
       +---totalPrefix change only, changes in high and low
       |
       |
       +---highOrderPart change only, no length change
```

   Ideally, a given prefix should either be "unchanged," "old," or
   "new." Renumbering will be easiest when each "old" prefix can be
   mapped to a single "new" prefix.

Unfortunately, the ideal often will not be attainable. It may be
necessary to run parts of the new and old address spaces in parallel.

Renumbering applies first to prefixes and then to host numbers to the
right of the prefix. To understand the scope of renumbering, it is
essential to:

   1. Identify the prefixes (and possibly host fields) potentially
       affected by the renumbering operation.

   2. Identify the authority that controls the values of the prefix,
       or part of the prefix, affected by renumbering.

In a given enterprise, prefixes may be present that will be under the
complete or partial control of the enterprise, as well as totally
outside the control of the enterprise. Let us review the principles
of control over address space.

More commonly, the most significant bits of the prefix are assigned
to the enterprise by an address registry (e.g., InterNIC, RIPE, or
APNIC) or by an Internet Service Provider (ISP). This assignment of
a value in the most significant bit positions historically has been
called a "network number," when the assigned high-order part is 8,
16, or 24 bits long. More recent usage does not limit the assigned
part to a byte boundary. The preferred term for the assigned part is
a "CIDR block" of a certain number of bits [RFC1518].

The enterprise then extends the prefix to the right, creating
"subnets." It is critical to realize that routers make routing
decisions based on the total prefix of interest, regardless of who
controls which bits. In other words, the router really doesn't know
or care about subnet boundaries.

The way to think about subnetting is that it creates a longer prefix.
Even before CIDR, we collapsed multiple subnets into a single network
number advertisement sent to external routers. In a more general
way, we now think of extending the prefix to the right as subnetting
and collapsing it to the left as supernetting, aggregating, or
summarizing. Depending on the usage of subnetting or aggregation,
different prefix lengths are significant at different router
interfaces.

4.3 Renumbering Scope

   Prefixes may be taken from the private address space [RFC1918] that
   is not routable on the global Internet. Since these addresses are
   not routable on the global Internet, changing parts of private
   address space prefixes is an enterprise-local decision.

If a prefix is totally outside the control of the enterprise, it is
external, and will be minimally affected by routing. Potential
interactions of external prefixes with enterprise renumbering
include:

   1) Inadvertent alteration or deletion of external addresses
       as part of router reconfiguration.
   2) Loss of connectivity to application servers inside the
       enterprise, because the external client no longer knows
       the internal address of the server.
   3) DNS/BGP
   4) Security

Prefixes partially under the control of the enterprise may change.
The scope of this will vary depending on whether only the externally
controlled part of the prefix changes, or if part of the internally
controlled part is to be renumbered. If the length of either the
high-order or low-order parts change, the process becomes more
complex.

High-order-part-only renumbering is most common when an organization
changes ISPs, and needs to renumber into the new provider's space.
The old prefix may have been assigned to the enterprise but will no
longer be used for global routing, or the old prefix may have been
assigned to the previous provider. Note that administrative
procedures may be necessary to return the previous prefix, although
this usually will be done by the previous provider. There often will
need to be a period of coexistence between the old and new prefixes.

Low-order-part-only renumbering can occur when an enterprise modifies
its internal routing structure, and the changes only affect the
internal subnet structure of the enterprise network. This is typical
of efforts involved in increasing the number of available subnets
(e.g., for more point-to-point media) or increasing the number of
hosts on a medium (e.g., in greater use of workgroup switches).

Both the high-order and low-order parts may change. This might
happen when the enterprise changes to a new ISP, who assigns address
space from a CIDR block rather than a classful network previously
used. With a different high-order prefix length, the enterprise
might be forced to change its subnet structure.

5. Moving toward a Renumbering-Friendly Enterprise

   Renumbering affects both the configuration of specific router
   "boxes," and the overall system of routers in a routing domain. The
   emphasis of this section is on making the current enterprise more
   renumbering-friendly, before any prefixes are actually changed.

Renumbering will have the least impact when the minimum number of
reconfiguration options are needed. When planning renumbering on
routers, consider that many existing configurations may contain
hard-coded IP addresses that may not be necessary, even if
renumbering were not to occur. Part of a router renumbering effort
should include, wherever possible, replacing router mechanisms based
on hard-coded addresses with more flexible mechanisms.

Renumbering will also generally be easier if the configuration
changes can be made offline on appropriate servers, and then
downloaded to the router if the router implementation permits.

5.1 Default Routes

A well-known method for reducing the amount of reference by one
router to other routers is to use a default route to a higher-level,
better-connected router. This assumes a hierarchical network design,
which is generally desirable in the interest of scaling.

Default routes are most appropriate for stub routers inside a routing
domain, and for boundary routers that connect the domain to a single
ISP.

5.2 Route Summarization and CIDR

When routes need to be advertised, summarize as much as is practical.
Summarization is most effective when address prefixes have been
assigned in a consistent and contiguous manner, which is often not
the case in legacy networks. Nevertheless, there is less to change
when we can refer to blocks of prefixes.

Not all routing mechanisms support general summarization. Interior
routing mechanisms that do include RIPv2, OSPF, EIGRP, IS-IS, and
systems of static routes. RIPv1 and IGRP do support classful
summarization (i.e., at Class A/B/C network boundaries only).

If existing addresses have been assigned hierarchically, it may be
possible to renumber below the level of summarization, while hiding
the summarization to the rest of the network. In other words, if all
the address bits being renumbered are to the right of the summarized
prefix length, the change can be transparent to the overall routing
system.

Even when effective summarization is possible to hide the details of
routing, DNS, filters, and other services may be affected by any
renumbering.

5.3 Server References in Routers

   Routers commonly communicate with an assortment of network management
   and other infrastructural servers. Examples of these servers are
   given in the "Network Management" section below. DNS itself,
   however, may be an important exception.

   Wherever possible, servers should be referenced by DNS name rather
   than by IP address. If a specific router implementation only
   supports explicit address references, this should be documented as
   part of the renumbering plan.

   Routers may also need to forward end host broadcasts to other
   infrastructure services (e.g., DNS, DHCP/BOOTP). Configurations that
   do this are likely to contain hard-coded IP addresses of the
   destination hosts or their subnets, which will need to be changed as
   part of renumbering.

5.4 DNS and Router Renumbering

   The Domain Name Service is a powerful tool in any renumbering effort,
   and can help routers as well as end hosts. If traceroute displays
   DNS names rather than IP addresses, certain debugging options can be
   transparent through the address transition.

   Be aware that dynamically learned names and addresses may be cached
   in router tables. For a router to learn changes in address to name
   correspondence, it may be necessary to restart the router or
   explicitly clear the cache.

   Alternatively, router configuration files may contain hard-coded
   address/name correspondences that will not be affected by a change in
   the DNS server.

   Different DNS databases are affected by renumbering. For example,
   the enterprise usually controls its own "forward" data base, but the
   reverse mapping data base may be maintained by its ISP. This can
   require coordination when changing providers.

   Commonly, router renumbering goes through a transition period.
   During this transition, old and new addresses may coexist in the
   routing system. Coexistence over a significant period of time is
   especially likely for DNS references to addresses that are known in
   the global Internet [deGroot]. Various DNS servers throughout the
   world may cache addresses for periods of days.

If, for example, a given router interface may have a coexisting new
and old address, it can be appropriate to introduce the new address
as an additional A record for the new address.

DNS RR statements can end with a semicolon, indicating the rest of
the line is a comment. This can be used as the basis of tools to
renumber DNS names for router addresses, by putting a comment (e.g.,
";newaddr") at the end of the A statements for the new addresses. At
an appropriate time, a script could generate a new zone file in which
the new addresses become the primary definitions on A records, and
the old addresses could become appropriately commented A records. At
a later time, these commented entries could be removed.

Care should be taken to assure that PTR reverse mapping entries are
defined for new addresses, because some router vendor tools depend on
reverse mapping.

5.5 Dynamic Addressing

Renumbering is easiest when addresses need to be changed in the least
possible number of places. Dynamic address assignment is especially
attractive for end hosts, and routers may play a key role in this
process. Routers may act as servers and actually assign addresses,
or may be responsible for forwarding end host address assignment
requests to address assignment servers.

The most common use of dynamic address assignment is to provide IP
addresses to end systems. Dynamic address assignment, however, is
also used to assign IP addresses to router interfaces. An address
assignment server may assign an IP address to a router either in the
usual DHCP way, based on a MAC address in the router, or simply based
on the physical connectivity of the new router. In other words, any
router connected on a specific interface of the configuring router
would be assigned the same IP address.

5.5.1 Router Roles in LAN-based DHCP Address Assignment

End hosts attached to LANs often obtain address assignments from
BOOTP or DHCP servers. If the server is not on the same medium as
the end hosts, routers may need to play a role in establishing
connectivity between the end host and the address server.

If the client is not on the same medium as the address assignment
server, routers either must act as address assignment services, or
forward limited broadcasts to the location of appropriate servers.

If the router acts as an address assignment server, its database of
addresses that it can assign may change during renumbering. If the
router forwards to a DHCP or BOOTP server, it must know the address
of that server. That server address can itself change as a result of
renumbering.

While the usual perception of DHCP is that it assigns addresses from
a pool, such that assignments to a given host at a given time is
random within the pool, DHCP can also return a constant IP address
for a specific MAC address. This may be much easier to manage and
troubleshoot, especially during renumbering.

Clearly, if the DHCP server identifies end hosts based on their MAC
address, consideration must be given to making that address unique,
and changing the DHCP database if either the MAC address or the IP
address changes. One way to reduce such reconfiguration is to use
Locally-Administered Addresses (LAA) on end hosts, rather than
globally unique addresses stored in read-only memory (ROM). Using
LAAs solves the problem of MAC addresses changing when a network
interface card changes, but LAAs have their own management problems
of configuration into end systems and maintaining uniqueness within
the enterprise.

5.5.2 Router Roles in Dialup Address Assignment

There are several possible ways in which an dialup end host interacts
with address assignment. Routers/access servers can play critical
roles in this process, either to provide connectivity between client
and server, or directly to assign addresses.

Different vendors handle address assignment in different ways.
Methods include:

    1. The access server forwards the request to a DHCP server, using
        a unique 48-bit identifier associated with the client. Note
        that this usually should not be the MAC address of the access
        server, since the same MAC address would then be associated
        with different hosts.

    2. The server forwards the request to an authentication server,
        which in turn gets a dynamic address either from:

        a. internal pools
        b. a DHCP server to which it forwards

    3. The server selects an address from locally configured pools
        and provides it to the dialing host without the intervention
        of other services.

When the router contains assignable addresses, these may need to
change as part of renumbering. Alternatively, hard-coded references
to address assignment or authentication servers may need to change.

5.5.3 Router Autoonfiguration

This initial address assignment may provide an address only for a
single connection (i.e., between the new and configuring routers).
The newly assigned address may then be used to "bootstrap" a full
configuration into the new router.

Dynamic address assignment to routers is probably most common at
outlying "stub" or "edge" routers that connect via WAN links to a
central location with a configuration server. Such edge devices may
be shipped to a remote site, plugged in to a WAN line, and use
proprietary methods to acquire part or all of their address
configuration.

When such autoconfiguration is used on edge routers, it may be
necessary to force a restart of the edge router after renumbering.
Restarting may be the only way to force the autoconfigured router to
learn its new address. Other out-of-band methods may be available to
change the edge router addresses.

5.6 Network Address Translation

Network address translation (NAT) is a valuable technique for
renumbering, or even for avoiding the need to renumber significant
parts of an enterprise [RFC1631]. It is not always transparent to
network layer protocols, upper layer protocols, and network
management tools, and must not be regarded as a panacea.

In the classic definition of NAT, certain parts of the routing system
are designated as stub domains, and connect to the global domain only
through NAT functions. The NAT contains a translation mechanism that
maps a stub address to a global address. This mechanism may map
statically or dynamically.

A more general NAT mechanism often is implemented in firewall bastion
hosts, which isolate "inside" and "outside" addresses through
transport- or application-level authenticated gateways. Mappings
from a "local" or "inside" address to a global address often is not
one-to-one, because an inside address is dynamically mapped to a TCP
or UDP port on an outside interface address.

In general, if there are multiple NATs, their translation mechanisms
should be synchronized. There are specialized cases when a given
stub address appears in more than one stub domain, and ambiguity

develops if one wishes to map, say, from 10.1.0.1/16 in stub domain A
to 10.1.0.1/16 in stub domain B. In this case, both 10.1.0.1
addresses identify different hosts.   Special mechanisms would have
to exist to map the domain A local address into one global address,
and to map the domain B local address into a different global
address.

NAT can have a valuable role in renumbering. Its intelligent use can
greatly minimize the amount of renumbering that needs to be done.
NAT, however, is not completely transparent.

Specifically, it can interfere with the proper operation of any
protocol that carries an IP address as data, since NAT does not
understand passenger fields and is unaware numbers need to change.

Examples of protocols affected are:

      --TCP and UDP checksums that are in part based on the
        IP header.   This includes end-to-end encryption schemes
        that include the TCP/UDP checksum
      --ICMP messages containing IP addresses
      --DNS queries that return addresses or send addresses
      --FTP interactions that use an ASCII-encoded IP address
        as part of the PORT command

It may be possible to avoid conflict if only certain hosts use
affected protocols. Such hosts could be assigned only global
addresses, if the network topology and routing plan permit.

Early NAT experiments suggested that it needs a sparse end-to-end
traffic mapping database for reasonable performance. This may or may
not be an issue in more hardware-based NAT implementations.

Another concern with NAT is that unique host addresses are hidden
outside the local stub domains. This may in fact be desirable for
security, but may present network management problems. One
possibility would be to develop a NAT MIB that could be queried by
SNMP to find the specific local-to-global mappings in effect.

There are also issues for DNS, DHCP, and other address management
services. There presumably would need to be local servers within
stub domains, so address requests would be resolved uniquely in each
stub (or would return appropriate global addresses).

6. Potential Pitfalls in Router Renumbering

   One way to categorize potential pitfalls is to look at those
   associated with the overall numbering plan itself and routing
   advertisement, and those associated with protocol behavior. In
   general, the former case is static and the latter is dynamic.

6.1 Static

   Problems can be implicit to the address/routing structure itself.
   These can include failures of components to understand arbitrary
   prefix addressing (i.e., classless routing), reachability due to
   inappropriate default or aggregated routes, etc.

6.1.1 Classless Routing Considerations

   Among the major reasons to renumber is to gain globally routable
   address space. In the global Internet, routable address space is
   based on arbitrary-length prefixes rather than traditional address
   classes. Classless Inter-Domain Routing (CIDR) is the administrative
   realization of prefix addressing in the global Internet. Inside
   enterprises, arbitrary prefix length addressing often is called
   Variable Length Subnet Masking (VLSM) or "subnetting a subnet."

   The general rules of prefix addressing must be followed in CIDR.
   Among these are permitting use of the all-zeroes and all-one subnets
   [RFC1812], and not assuming that a route to a "Class A/B/C network
   number" implies routes to all subnets of that network. Assumptions
   also should not be made that a prefix length is implied by the
   structure of the high-order bits of the IP address (i.e., the "First
   Octet Rule").

   This ideal, unfortunately, is not understood by a significant number
   of routers (and terminal access servers that participate in routing),
   and an even more significant number of host IP implementations.

   When planning renumbering, network designers must know if the new
   address has been allocated using CIDR rules rather than traditional
   classful addressing. If CIDR rules have been followed in address
   assignment, then steps need to be taken to assure the router
   understands them, or appropriate steps need to be taken to interface
   the existing environment to the CIDR environment.

   Current experience suggests that it is best, when renumbering, to
   maintain future compatibility by moving to a CIDR-supportive routing
   environment. While this is usually thought to mean introducing a
   classless dynamic routing protocol, this need not mean that routing
   become much more complex. In a RIPv1 environment, moving to RIPv2

may be a relatively simple change. Alternative simple methods
include establishing a default route from a non-CIDR-compliant
routing domain to a CIDR-compliant service provider, or making use of
static routes that are CIDR-compliant.

Some routers support classless routing without further
configuration, other routers support classless routing but require
specific configuration steps to enable it, while other routers only
understand classful routing. In general, most renumbering will
eventually require classless routing support. It is essential to
know if a given router can support classless routing. If it does
not, workarounds may be possible. Workarounds are likely to be
necessary.

6.1.1.1 Aggregation Problems

In experimenting with the CIDR use of a former Class A network
number, it was shown in RFC1879 that CIDR compliance rules must be
enabled explicitly in some routers, while other routers do not
understand CIDR rules.

RFC 1897 demonstrated problems with some existing equipment,
especially "equipment that depends on use of a classful routing
protocol, such as RIPv1 are prone to misconfiguration. Tested
examples are current Ascend and Livingston gear, which continue to
use RIPv1 as the default/only routing protocol. RIPv1 use will
create an aggregate announcement.... The Ascend was told to announce
39.1.28/24, but since RIPv1 can't do this, the Ascend instead sent
39/8." RIPv1, like all classful interior protocols, is obsolescent.

6.1.1.2 Discontiguous Networks

Another problem that can occur with routers or routing mechanisms
that do not understand arbitrary length prefix addressing is that of
discontiguous networks.   This problem is easy to create
inadvertently when renumbering. In the example below, assume the
enterprise has been using 10.0.0.0/8 as its primary prefix, but has
introduced an ISP whose registered addresses were in 172.31.0.0/16.

Assume a RIPv1 system of three routers:

```
              10.1.0.1/16          10.2.0.1/16
                   |                    |
                   |                    |
          +-----------------------------------+
          |               Router 1            |
          +-----------------------------------+
                        | 172.31.1.1/24
                        |
                        |
                        | 172.31.1.2/24
          +-----------------------------------+
          |               Router 2            |<------OUTSIDE
          +-----------------------------------+
                        | 172.31.2.1/24
                        |
                        |
                        | 172.31.2.2/24
          +-----------------------------------+
          |               Router 3            |
          +-----------------------------------+
                   |                    |
                   |                    |
              10.3.0.1/16          10.4.0.1/16
```
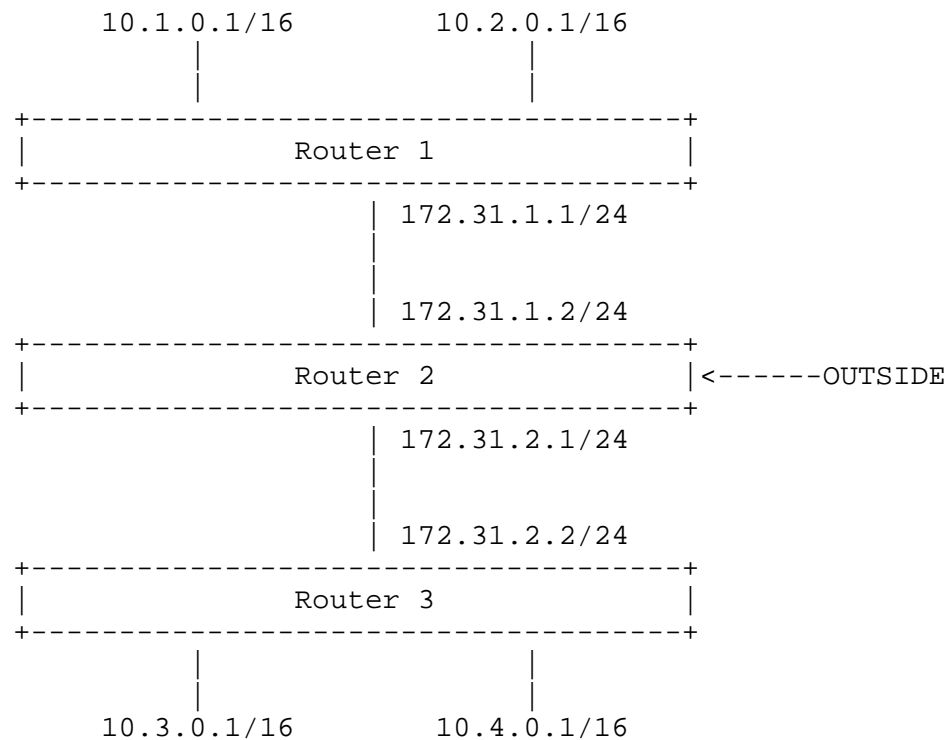
Router 1 can reach its two locally connected subnets, 10.1.0.0/16 and
10.2.0.0/16. It will aggregate them into a single announcement of
10.0.0.0/8 when it advertises out the 172.31.1.1 interface.

In like manner, Router 3 can reach its two locally connected subnets,
0.3.0.0/16 and 10.4.0.0/16. It will aggregate them into a single
announcement of 10.0.0.0/8 when it advertises out the 172.31.2.2
interface.

When Router 2 receives a packet from its "outside" interface
destined, say, to 10.1.1.56/16, where does it send it? Router 2 has
received two advertisements of 10.0.0.0 on different interfaces,
without any detail of subnets inside 10.0.0.0. Router 2 has an
ambiguous routing table in terms of the next hop to a subnet of
10.0.0.0. We call this problem, when parts of the same classful
network are separated by different networks, discontigous subnets.

Two problems occur in this configuration. Router 2 does not know
where to send outside packets destined for a subnet of 10.0.0.0.
Connectivity, however, also will break between Routers 1 and 3,
because Router 2 does not know the next hop for any subnet of
10.0.0.0.

There are several workarounds to this problem. Obviously, one would
be to change to a routing mechanism that does advertise subnets. An
alternative would be to establish an IP-over-IP tunnel through Router
2, and give this a subnet in 10.0.0.0. This additional subnet would
be visible only in Routers 1 and 3. It would solve the connectivity
problem between Routers 1 and 3, but Router 2 would still not be able
to forward outside packets. This might be a perfectly acceptable
solution if Router 2 is simply being used to connect two parts of
10.0.0.0.

Another way to deal with the discontiguous network problem is to
assign secondary addresses in 10.0.0.0 to the R1-R2 and R2-R3
interfaces, which would allow the 10.0.0.0 subnets to be advertised
to R2. This would work as long as there is no problem in advertising
the 10.0.0.0 subnets into the R2 routing system. There would be a
problem, for example, if the 10.0.0.0 address were in the private
address space but the R2 primary addresses were registered, and R2
advertised the 10.0.0.0 addresses to the outside.

This problem can be handled if R2 has filtering mechanisms that can
selectively block 10.0.0.0 advertisements to the outside world. The
configuration, however, will become more and more complicated.

6.1.1.3 Router-Host Interactions

The situation may not be as bleak if hosts do not understand prefix
addressing but routers do. Methods exist for hiding a VLSM structure
from end hosts that do not understand it. These do involve protocol
mechanisms as workarounds, but the fundamental problem is hosts'
understanding of arbitrary prefix lengths.

A key mechanism is proxy ARP [Carpenter]. The basic mechanism of
using proxy ARP in prefix-based renumbering is to have the hosts
issue an ARP whenever they want to communicate with a destination.
If the destination is actually on the same subnet, it will respond
directly to the ARP. If the destination is not, the router will
respond as if it were the destination, and the originating host will
send the datagram to the router. Once the datagram is in the router,
the VLSM-aware router can forward it.

Many end hosts, however, will only issue an ARP if they conclude the
destination is on their own subnet. All other traffic is sent to a
hard-coded default router address. In such cases, workarounds may be
needed to force the host to ARP for all destinations.

One workaround involves a deliberate misconfiguration of hosts.
Hosts that only understand default routers also are apt only to
understand classful addressing. If the host is told its major (i.e.,

classful) network is not subnetted, even though the address plan
actually is subnetted, this will often persuade it to ARP to all
destinations.

This also works for hosts that do not understand subnetting at all.
An example will serve for both levels of host understanding. Assume
the enterprise uses 172.31.0.0/16 as its major address, which is in
the Class B format. This is actually subnetted with eight bits of
subnetting -- 172.31.0.0/24. Individual hosts unaware of VLSM,
however, either infer Class B from the address value, or are told
that the subnet mask in effect is 255.255.0.0.

Yet another approach that helps hosts find routers is to run passive
RIP on the hosts, so that they hear routing updates. They assume any
host that issues routing updates must be a router, so traffic for
non- local destinations can be forwarded there. While RIPv1 does not
support arbitrary prefixes, the router(s) issuing the routing updates
may have additional capabilities that let them correctly forward such
traffic. The priority, therefore, is to get the non-local routers to
a router that understands the overall routing structure, and passive
RIP may be a viable way to do this.

It may be appropriate to cut over on a site-by-site basis [Lear]. In
such an approach, some sites have VLSM-aware hosts but others do not.
As long as the routing structure supports VLSM, workarounds can be
applied where needed.

6.1.2 MAC Address Interactions

While it is uncommon now for a router to acquire any of its interface
addresses as a DHCP client, this may become more common. When a
router so acquires addresses, care must be taken that the MAC address
presented to the DHCP server is in fact unique.

Modern routers usually support protocol architectures besides IP.
Some of these architectures, notably DECnet, Banyan VINES, Xerox
Network Services, and IPX, may modify MAC addresses of routers such
that a given MAC address appears on more than one interface. While
this is normally not a problem, it could cause difficulties when the
MAC address is assumed to be unique.

Other situations occur when the same MAC address can appear on more
than one interface. There are high-availability IBM options which
establish primary and backup instances of the same MAC address on
different physical interfaces of 37x5 communications controllers.

   Some end hosts running protocol stacks other than IP, notably DECnet,
   may change their MAC address from the globally assigned built-in
   address.

6.2 Dynamic

   Dynamic protocol mechanisms that to some extent depend on IP
   addresses may be affected by router renumbering. These include
   mechanisms that assign or resolve addresses (e.g., DHCP, DNS),
   mechanisms that use IP addresses for identification (e.g., SNMP),
   security and authentication mechanisms, etc. The listed mechanisms
   are apt to have configuration files that will be affected by
   renumbering.

   Another area of dynamic behavior that can be affected is caching.
   Application servers, directory functions such as DNS, etc., may cache
   IP addresses. When the router is renumbered, these servers may point
   to old addresses. Certain proxy server functions may reside on
   routers, and the router may need to be restarted to reset the cache.

   The endpoints of TCP tunnels terminating on routers may be internally
   identified by address/port pairs at each end. If the address
   changes, even if the port remains the same, the tunnel is likely to
   need to be reestablished.

7. Tools and Methods for Renumbering

   The function of a renumbering tool can be realized either as manual
   procedures as well as software. This section deals with functionality
   of hypothetical yet general renumbering tools rather than their
   implementation.

   General caveat: tools should know whether the environment supports
   classless addressing. If it does not, newly generated addresses
   should be checked to see they do not fall into the all-zeroes or
   all-ones subnet values.

7.1 Search Mechanisms

   Tools will be needed to search configuration files and other
   databases to identify addresses and names that will be affected by
   reorganization. This search should be based on the address inventory
   described above.

   Especially when searching for names, common search tools using
   regular expressions (e.g., grep) may be very useful. Some additional
   search tools may be needed. One would convert dotted decimal search
   arguments to binary and only then makes the comparison.

The comparison may need to be done under the constraint of a mask.
Such a comparator would also be relevant as the second phase that
looks for ipAddress and other relevant strings in MIBs.

7.2 Address Modification

The general mechanism for address modification is substitution of an
arbitrary number of bits in an address. In the simplest cases, there
is a one-to-one correspondence between old and new bit positions.

Especially when address modification is manual, it should be
remembered that the affected bits can be obscured by dotted decimal
notation. Work in, or at least consider, binary notation to assure
accuracy.

Several basic functions can be defined:

zerobits(position,length):
   clear <length> bits starting at <position>
orbits(position,length,newbits)
   OR the bit string <newbits> of <length> starting at <position>

In these examples, the index [j] is used to identify entries in the
address inventory database for existing addresses, while [k]
identifies new addresses.

Remember that these tools operate at a bit level, so the new address
will have to be converted back into dotted decimal, MIB ASN.1, or
other notation before it can be replaced into configuration files.

7.2.1 Prefix Change, No Change in Length

If the entire new prefix has the same number of bits as the old
external part, requiring no change to subnetting, the router part of
renumbering may be fairly simple. If the router configurations are
available in machine-readable form, as text files or parsable SNMP
data, it is a relatively simple task to define a tool to examine IP
addresses in the configuration, identify those beginning with the old
prefix, and substitute the new prefix bit-by-bit.

for each address[j],
   zerobits(0,PrefixLength[j])
   orbits(0,PrefixLength[j],NewPrefix[j])

Note that the host part is unaffected. Both subnet specifications
(e.g., for route advertisements) and specific host references will be
changed correctly in this simple case.

7.2.2 highOrderPart change

   Matters are slightly more complex when the change applies only to the
   externally-controlled part of the prefix, as might be the case when
   an organization changes ISPs and renumbers into the ISP's address
   space, without changing the internal subnet structure.

   The substitution process is much as the previous case, except only
   the high-order bits change:

   for each address,
      zerobits(0,highOrderPartLength[j])
      orbits(0,highOrderPartLength,newHighOrderPart[k])

7.2.3 lowOrderPart change

   Organizations might renumber only the lowOrderPart (i.e., the subnet
   field) of their address space. This might be done to clean up an
   address space into contiguous blocks prior to introducing a routing
   system that aggregates addresses, such as OSPF.

   for each address[j],
      zerobits(highOrderPartLength[j],lowOrderPartLength[j])
      orbits(highOrderPartLength[j],
          lowOrderPartLength[j],newLowOrderPart[k])

7.2.4 lowOrderPart change, Change in lowOrderPart length

   When the length of the subnet field changes in all or part of the
   address space, things become significantly more complex. A fairly
   simple case arises when the host field is consistently too long, at
   least in the affected subnets. This is common, for example, when
   address space is being recovered in an existing Class B network with
   8 bits of subnetting. Certain /24 bit prefixes are being extended to
   /30 and reallocated to point-to-point real or virtual (e.g., DLCI)
   media.

   for each address [j]
    if address is affected by renumbering
     if newLowOrderPartLength[k] > oldLowOrderPartLength[j]
      then
       zerobits(highOrderPartLength[k],newLowOrderPartLength[k])
       orbits(highOrderPartLength[k],newLowOrderPart[k])
      end

7.2.5 highOrderPart change, Change in highOrderPart length

   When the length of the high-order part changes, but it is desired to
   keep the existing subnet structure, constraints apply. The situation
   is fairly simple if the new high-order part is shorter than the
   existing high order part.

   If the new high-order part is longer than the old high order part,
   constraints are more complex. The key is to see if any of the <k>
   most significant bits of the oldHighOrderPart, which overlap the <k>
   least significant bits of the newHighOrderPart, are nonzero. If no
   bits are nonzero, it may be simply to overlay the new prefix bits.

7.3 Naming

   It is worthwhile to distinguish that a router's use of a DNS name
   does not necessarily mean that name is defined in a name server.
   Routers often contain static address to name mappings local to the
   router, so both the DNS zone files and the router configurations will
   need to be checked.

   What we first want to do is generate a list of name/address mappings,
   the mapping mechanism for each instance (e.g., static entry in
   configuration file, RR in our zone's DNS, RR in a zone file outside
   ours), the definition statement (or equivalent if the routers are
   configured with SNMP), and current IP address. We then want to
   compare the addresses in this list to the list defined earlier of
   prefixes affected by renumbering.   The intersection of these lists
   define where we need to make changes.

   Note that the explicit definition statement, or at leasts its
   variables, should be kept. In the real world, static IP address
   mappings in hosts may not have been maintained as systematically as
   are RR records in a DNS server.   It is entirely possible that
   different host mapping entries for the same name point to different
   addresses.

7.3.1 DNS Tools

   The DNS itself can both delay and and speed router renumbering.
   Caches in DNS servers both inside and outside the organization may
   have sufficient persistence that a "flag day" cutover is not
   practical if worldwide connectivity is to be kept. DNS can help,
   however, make a period of old and new address coexistence work.

   If, for example, a given router interface may have a coexisting new
   and old address, it can be appropriate to introduce the new address
   as a CNAME alias for the new address.

DNS RR statements can end with a semicolon, indicating the rest of
the line is a comment. This can be used as the basis of tools to
renumber DNS names for router addresses, by putting a comment (e.g.,
";newaddr") at the end of the CNAME statements for the new addresses.
At an appropriate time, a script could generate a new zone file in
which the new addresses become the primary definitions on A records,
and the old addresses could become appropriately commented CNAME
records. At a later time, these commented CNAME entries could be
removed.

Care should be taken to assure that PTR reverse mapping entries are
defined for new addresses, because some router vendor tools depend on
reverse mapping.

7.3.2   Related name tools

Especially on UNIX and othe that do routing, there may be static name
definitions. Such definitions are probably harder to keep maintained
than entries in the DNS, simply because they are more widely
distributed.

Several tools are available to generate more maintainable entries. A
perl script called h2n converts host tables into zone data files that
can be added to the DNS server. It is available as
ftp://ftp.uu.net/published/oreilly/nutshell/dnsbind/dns.tar.Z.
Another tool, makezones, is part of the current BIND distribution,
and can also be obtained from
ftp://ftp.cus.cam.ac.uk/pub/software/programs/DNS/makezones

See the DNS Resources Directory at http://www.dns.net/dnsrd.

8. Router Identifiers

Configuration commands in this category assign IP addresses to
physical or virtual interfaces on a single router. They also include
commands that assign IP-address-related information to the router
"box" itself, and commands which involve the router's interaction
with neighbors below the full routing level (e.g., default gateways,
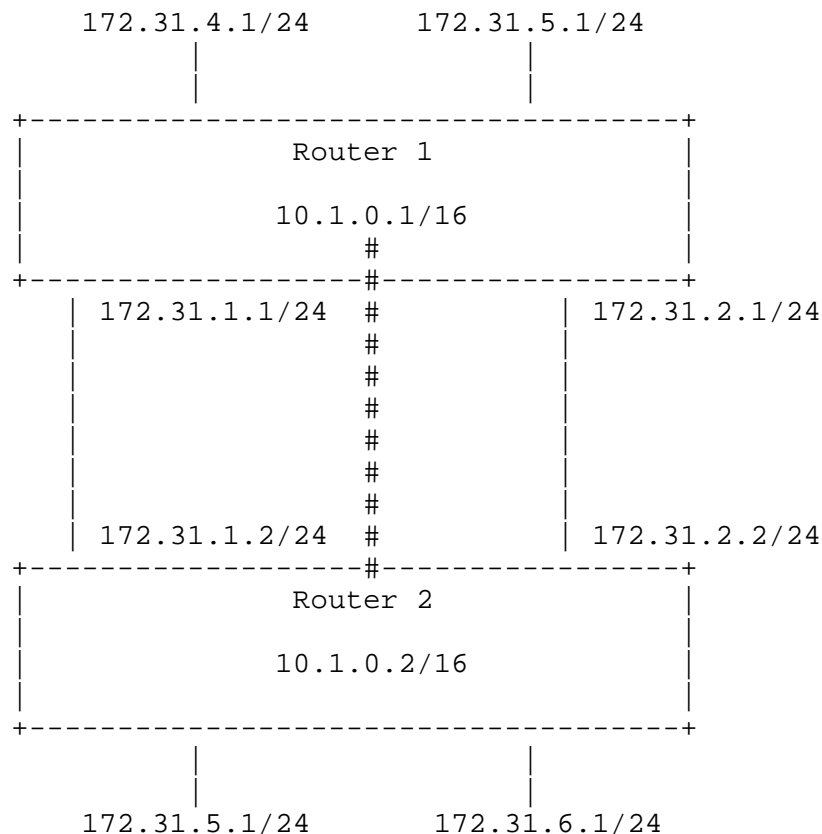ARP).

Routers may have other unique identifiers, such as DNS names used for
the set of addresses on the "box," or SNMP systemID strings.

8.1. Global Router Identification

Traditional IP routers do not have unique identifiers, but rather are
treated as collections of IP addresses associated with their
interfaces. Some protocol mechanisms, notably OSPF and BGP, need an

address for the router itself, typically to establish tunnel
endpoints between peer routers. Other applications include
"unnumbered interfaces" used to conserve address space for serial
media, a practice discussed further below.

In the illustration below, the 10.1.0.0/16 address space is used for
global identifiers. A TCP tunnel runs from 10.1.0.1 to 10.1.0.2, but
its traffic is load-shared among the two real links, 172.31.1.0 and
172.31.2.0.

```
               172.31.4.1/24      172.31.5.1/24
                     |                   |
                     |                   |
          +-------------------------------------+
          |                Router 1             |
          |                                     |
          |               10.1.0.1/16           |
          |                    #                |
          +-----------------#---------------+
          |  172.31.1.1/24  #    | 172.31.2.1/24
          |                 #    |
          |                 #    |
          |                 #    |
          |                 #    |
          |                 #    |
          |                 #    |
          |  172.31.1.2/24  #    | 172.31.2.2/24
          +-----------------#---------------+
          |                Router 2             |
          |                                     |
          |               10.1.0.2/16           |
          |                                     |
          +-------------------------------------+
                  |                   |
                  |                   |
               172.31.5.1/24      172.31.6.1/24
```

A common practice to provide router identifiers is using the "highest
IP address" on the router as an identifier for the "box." Many
implementations have a default mechanism to establish the router ID,
which may be the highest configured address, or the highest active
address.

Typical applications of a global router ID may not require it be a
"real" IP address that is advertised through the routing domain, but
is simply a 32-bit identifier local to each router. When this is the
case, this identifier can come from the RFC 1918 private address
space rather than the enterprise's registered address space.

Allowing default selection of the router ID can be unstable and is
not recommended. Most implementations have a means of declaring a
pseudo-IP address for the router itself as opposed to any of its
ports.

Changes to this pseudo-address may have implications for DNS. Even
if this is not a real address, A and PTR resource records may have
been set up for it, so diagnostics can display names rather than
addresses.

Another potential DNS implication is that a CNAME may have been
established for the entire set of interface addresses on a router.
This allows testing, telnet, etc., to the router via any reachable
path.

8.2 Interface Address

Interface addresses are perhaps the most basic place to begin router
renumbering. Interface configuration will require an IP address, and
usually a subnet mask or prefix length. Some implementations may not
have a subnet mask in the existing configuration, because they use a
"default mask" based on a classful assumption about the address. Be
aware of possible needs for explicit specification of a subnet mask
or other prefix length specification when none previously was
specified. This will be especially common on older host-based
routers.

Multiple IP addresses, in different subnets, can be assigned to the
same interface. This is often a valuable technique in renumbering,
because the router interface can be configured to respond to both the
new and old addresses.

Caution is necessary, however, in using multiple subnet addresses on
the same interface. OSPF and IS-IS implementations may not advertise
the additional addresses, or may constrain their advertisement so all
must be in the same area.

When this method is used to make the interface respond to new and old
addresses, and the renumbering process is complete, care must be
taken in removing the old addresses. Some router implementations
have special meaning to the order of address declarations on an
interface. It is highly likely that routers, or at least the
interface, must be restarted after an address is removed.

8.3 Unnumbered Interfaces

   As mentioned previously, several conventions have been used to avoid
   wasting subnet space on serial links. One mechanism is to implement
   proprietary "half-router" schemes, in which the unnumbered link
   between router pairs is treated as an "internal bus" creating a
   "virtual router," such that the scope of the unnumbered interface is
   limited to the pair of routers.

```
   |           +------------+                +------------+    |
   |           |            |                |            |    |
   |      e0   |            |s0         s0   |            |    |
   |-----------|     R1     |................|     R2     |-------|
   | 192.168.1.1 | 10.1.0.1/16|              | 10.1.0.2/16|    |
   |      /24  |            |                |            |    |
   |           +------------+                +------------+    |
```
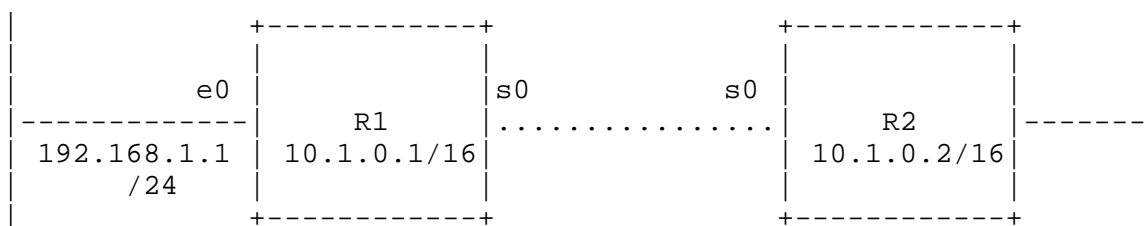
   In the above example, software in routers R1 and R2 automatically
   forward every packet received on serial interface S0 to Ethernet
   interface E0. They forward every packet on e0 to their local S0.
   Neither S0 has an IP address. R1 has the router ID 10.1.0.1/16 and
   R2 has 10.1.0.2/16.

   It is thus impossible to send a specific ping to the S0 interfaces,
   making it difficult to test whether a connectivity problem is due to
   S0 or E0. Some management is possible as long as at least one IP
   address on the router (e.g., E0) is reachable, since this will permit
   SNMP connectivity to the router. Once the router is reachable with
   SNMP, the unnumbered interface can be queried through the MIB
   ifTable.

   Another approach is to use the global router identifier as a pseudo-
   address for every unnumbered interface on a router. In the above
   example, R1 would use 10.1.0.1 as its identifier. This provides an
   address to be used for such functions as the IP Route Recording
   option, and for providing a next-hop-address for routes.

The second approach is cleaner, but still can create operational
difficulties. If there are multiple unnumbered interfaces on a
router, which one (if any) should/will respond to a ping? Other
network management mechanisms do not work cleanly with unnumbered
interface.

As part of a renumbering effort, the need for unnumbered interfaces
should be examined. If the renumbering process moves the domain to
classless addressing, then serial links can be given addresses with a
/30 prefix, which will waste a minimum of address space.

For dedicated or virtual dedicated point-to-point links within an
organization, another alternative to unnumbered operation is using
RFC1918 private address space. Inter-router links rarely need to be
accessed from the Internet unless explicitly used for exterior
routing. External traceroutes will also fail reverse DNS lookup.

If unnumbered interfaces are kept, and the router-ID convention is
used, it will probably be more stable to rely on an explicitly
configured router ID rather than a default from a numbered interface
address.

The situation becomes even more awkward if it is desired to use
unnumbered interfaces over NBMA services such as Frame Relay. OSPF,
for example, uses the IP address of numbered interfaces as a unique
identifier for that interface. Since unnumbered interfaces do not
have their own unique address, OSPF has not obvious way to identify
these interfaces. A physical index (e.g., ifTable) could be used,
but would have to be extended to have an entry for each logical entry
(i.e., VC) multiplexed onto the physical interface.

8.4 Address Resolution

While mapping of IP addresses to LAN MAC addresses is usually done
automatically by the router software, there will be cases where
special mappings may be needed. For example, the MAC address used by
router interfaces may be locally administered (i.e., set manually),
rather than relying on the burnt-in hardware address. It may be part
of a proprietary method that dynamically assigns MAC addresses to
interfaces. In such cases, an IP address may be part of the MAC
address configuration statements and will need to be changed.

   Manual mapping to medium addresses will usually be needed for NBMA
   and switched media. When renumbering IP addresses, statements that
   map the IP address to frame relay DLCIs, X.121 addresses, SMDS and
   ATM addresses, telephone numbers, etc., will need to be changed to
   the new address. Local requirements may require a period of parallel
   operation, where the old and new IP addresses map to the same medium
   address.

8.5 Broadcast Handling

   RFC1812 specifies that router interfaces MUST NOT forward limited
   broadcasts (i.e., to the all-ones destination address,
   255.255.255.255). It is common, however, to have circumstances where
   a LAN segment is populated only by clients that communicate with key
   servers (e.g., DNS or DHCP) by sending limited broadcasts. Router
   interfaces can cope with this situation by translating the limited
   broadcast address to a directed broadcast address or a specific host
   address, which is legitimate to forward.

   When limited address translation is done for serverless segments, and
   the new target address is renumbered, the translation rule must be
   reconfigured on every interface to a serverless segment. Be sure to
   recognize that a given segment might have a server from the
   perspective of one service (e.g., DHCP), but could be serverless for
   other services (e.g., NFS or DNS).

8.6 Dynamic Addressing Support

   Routers can participate in dynamic addressing with RARP, DHCP, BOOTP,
   or PPP. In a renumbering effort, several kinds of changes may made
   to be made on routers participating in dynamic addressing.

   If the router acts as a server for dynamic address assignment, the
   addresses it assigns will need to be renumbered. These might be
   specific addresses associated with MAC addresses or dialup ports, or
   could be a pool of addresses. Pools of addresses may be seen in pure
   IP environments, or in multiprotocol situations such as Apple MacIP.

   If the router does not assign addresses, it may be responsible for
   forwarding address assignment requests to the appropriate server(s).
   If this is the case, there may be hard-coded references to the IP
   addresses of these servers, which may need to be changed as part of
   renumbering.

9. Filtering and Access Control

   Routers may implement mechanisms to filter packets based on criteria
   other than next hop destination. Such mechanisms often are
   implemented differently for unicast packets (the most common case) or
   for multicast packets (including routing updates). Filtering rules
   may contain source and/or destination IP addresses that will need to
   change as part of a renumbering effort.

   Filtering can be done to implement security policies or to control
   traffic. In either case, extreme care must be taken in changing the
   rules, to avoid leakage of sensitive information. denial of access
   to legitimate users, or network congestion.

   Routers may implement logging of filtering events, typically denial
   of access. If logging is implemented, logging servers to which log
   events are sent preferably should be identified by DNS name. If the
   logging server is referenced by IP address, its address may need to
   change during renumbering.   Care should be taken that critical
   auditing data is not lost during the address change.

9.1 Static Access Control Mechanisms

   Router filters typically contain some number of include/exclude rules
   that define which packets to include in forwarding and which to
   exclude. These rules typically contain an address argument and some
   indication of the prefix length. This length indication could be a
   count, a subnet mask, or some other mask.

   When renumbering, the address argument clearly has to change. It can
   be more subtle if the prefix length changes, because the length
   specification in the rule must change as well. Needs for such changes
   may be hard to recognize, because they apply to ranges of addresses
   that might be at a level of aggregation above the explicit
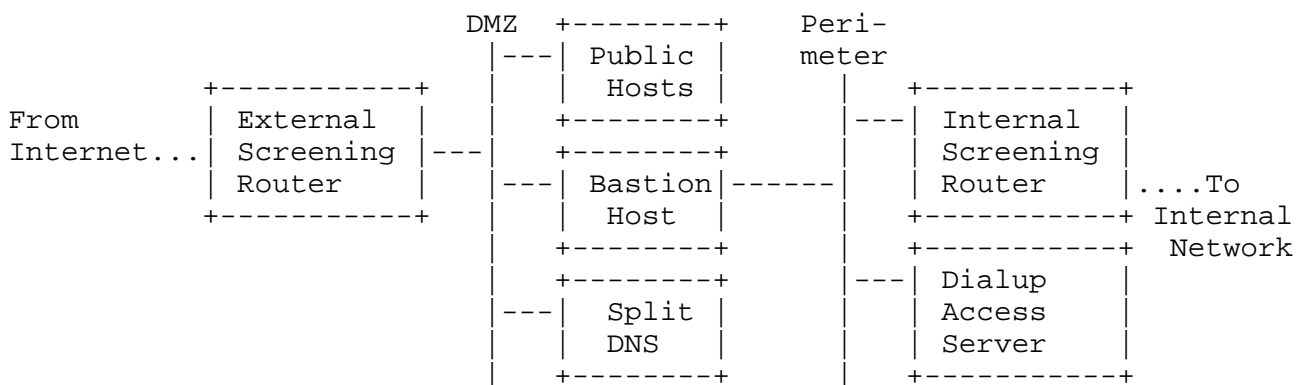   renumbering operation.

   RFC 1812 requires that address-based filtering allow arbitrary prefix
   lengths, but some hosts and routers might only allow classful
   prefixes.

9.2 Special Firewall Considerations

   Routers are critical components of firewall systems.
   Architecturally, two router functions are described in firewall
   models, the external screening router between the outside and the
   "demilitarized zone (DMZ)," and the internal screening router between
   the inside and the "perimeter network." Between these two networks
   is the bastion host, in which reside various non-routing isolation
   and authentication functions, beyond the scope of this document.

   One relevant aspect of the bastion host, however, is that it may do
   address translation or higher-layer mappings between differnt address
   spaces. If the "outside" address space (i.e., visible to the
   Internet) changes, this will mean that the outside screening router
   will need configuration changes. Since the outside screening router
   may be under the control of the ISP rather than the entrerprise,
   administrative coordination will be needed.

```
                              DMZ  +--------+     Peri-
                                   |---| Public |     meter
              +-----------+        |   | Hosts  |       |   +-----------+
From          | External  |        |   +--------+       |---| Internal  |
Internet...| Screening |---|   +--------+       |   | Screening |
              | Router    |        |---| Bastion|------|   | Router    |....To
              +-----------+        |   | Host   |       |   +-----------+ Internal
                                   |   +--------+       |   +-----------+ Network
                                   |   +--------+       |---| Dialup    |
                                   |---| Split  |       |   | Access    |
                                   |   | DNS    |       |   | Server    |
                                   |   +--------+       |   +-----------+
```

   External screening routers typically have inbound access lists that
   block unauthorized traffic from the Internet, and outbound access
   lists that permit access only to DMZ servers and the bastion host.
   The inbound filters commonly block the Private Address Space, as well
   as address space from the enterprise's internal network. If the
   internal network address changes, the inbound filters clearly will
   need to change.

   If DMZ host addresses change, the corresponding outbound filters from
   the external screening host also will need to change. Internal
   screening routers permit access from the internal network to selected
   servers on the perimeter network, as well as to the bastion host
   itself. If the enterprise uses private address space internally,
   renumbering may not affect this router.

Another component of a firewall system is the "split DNS" server,
which provides address mapping in relation to the globally visible
parts of the

9.3 Dynamic Access Control Mechanisms

Certain access control services, such as RADIUS and TACACS+, may
insert dynamically assigned access rules into router configurations.
For example, a RADIUS database "contains a list of requirements which
must be met to allow access for the user. This always includes
verification of the password, but can also specify the client(s) or
port(s) to which the user is allowed access. [Rigney]."

Configuration information dynamically communicated to the router may
be in the form of filtering rules. Effectively, this authentication
database becomes an extension of the router configuration database.
Both these databases may need to change as part of a renumbering
effort.

Another dynamic configuration issue arises when "stateful packet
screening" on bastion hosts or routers is used to provide security
for UDP-based services, or simply for IP. In such services, when an
authorized packet leaves the local environment to go into an
untrusted address space, a temporary filtering rule is established on
the interface on which the response to this packet is expected. The
rule typically has a lifetime of a single packet response. If these
rules are defined in a database outside of the router, the rule
database again is an extension of router configuration that must be
part of the renumbering effort.

10. Interior Routing

This section deals with routing inside an enterprise, which generally
follows, ignoring default routes, the rules:

    1. Does a single potential route exist to a destination?
        If so, use it.
    2. Is there more than one potential path to a destination?
        If so, use the path with the lowest end-to-end metric.
    3. Are there multiple paths with equal lowest cost to the
        destination? If so, consider load balancing.

Most enterprises do not directly participate in global Internet
routing mechanisms, the details of which are of concern to their
service providers. The next section deals with those more complex
exterior mechanisms.

10.1 Static Routes

   During renumbering, the destination and/or next hop address of static
   routes may need to change. It may be necessary to restart routers or
   explicitly clear a routing table entry to force the changed static
   route to take effect.

10.2 RIP (Version 1 unless otherwise specified)

   The Routing Information Protocol (RIP) has long been with us, as one
   of the first interior routing protocols. It still does that job in
   small networks, and also has been used for assorted functions that
   are not strictly part of interior routing. In this discussion, we
   will first deal with pure interior routing applications.

   In a renumbering effort that involves classless addressing, RIPv1 may
   not be able to cope with the new addressing scheme. Officially, this
   protocol is Historic and should be avoided in new routing plans.
   Where legacy support requirements dictate it be retained, it is
   worthwhile to try to limit RIPv1 in "stub" parts of the network.
   Vendor-specific mechanisms may be available to interface RIPv1 to a
   classless environment.

   As part of planning renumbering, strong consideration should be given
   to moving to RIPv2, OSPF, or other classless routing protocols as the
   primary means of interior routing. Doing so, however, may not remove
   the need to run RIP in certain parts of the enterprise.

   RIP is widely implemented on hosts, where it may be used as a method
   of router discovery, or for load-balancing and fault tolerance when
   multiple routers are on a subnet. In these applications, RIP need
   not be the only routing protocol in the domain; RIP may be present
   only on stub subnets. Destination information from more capable
   routing protocols may be translated into RIP updates. While it is
   generally reasonable to minimize or remove RIP as part of a
   renumbering effort, be careful not to disable the ability of hosts to
   locate routers.

   RIP is also used as a quasi-exterior routing mechanism between some
   customers and their ISPs, as a means simpler than BGP for the
   customer to announce routes to the provider.

10.3 OSPF

   OSPF has several sensitivities to renumbering beyond those of simpler
   routing protocols. If router IDs are assigned to be part of the
   registered address space, they may need to be changed as part of the
   renumbering effort. It may be appropriate to use RFC 1918 private

address space for router IDs, as long as these can be looked up in a
DNS server within the domain.

Summarization rules are likely to be affected by renumbering,
especially if area boundaries change.

Special addressing techniques, such as unnumbered interfaces and
physical interfaces with IP addresses in multiple subnets, may not be
transparent to OSPF. Care should be exercised in their use, and
their use definitely should be limited to intra-area scope.

If part of the renumbering motivation is the introduction of NBMA
services, there can be numerous impacts on OSPF. Generally, the best
way to minimize impact is to use separate subnets for each VC. By
doing so, different OSPF costs can be assigned to different VCs,
designated router configuration is not needed, etc.

10.4 IS-IS

   IP prefixes are usually associated with IS-IS area definitions. If
   IP prefixes change, there may be a corresponding change in area
   definitions.

10.5 IGRP and Enhanced IGRP

   When a change from IGRP to enhanced IGRP is part of a renumbering
   effort, the need to disable IGRP automatic route summarization needs
   to be considered. This is likely if classless addressing is being
   implemented.

   Also be aware of the nuances of automatic redistribution between IGRP
   and EIGRP. The "autonomous system number," which need not be a true
   AS number but simply identifies a set of cooperating routers, must be
   the same on the IGRP and EIGRP processes for automatic redistribution
   to occur.

11. Exterior Routing

   Exterior routes may be defined statically. If dynamic routing is
   involved, such routes are learned primarily from BGP. RIP is not
   infrequently used to allow ISPs to learn dynamically of new customer
   routes, although there are security concerns in such an approach.
   IGRP and EIGRP can be used to advertise external routes.

Renumbering that affects BGP-speaking routers can be complex, because
it can require changes not only in the BGP routers of the local
Autonomous System, but also require changes in routers of other AS
and in routing registries. This will require careful administrative
coordination.

If for no other reason than documentation, consider use of a routing
policy notation [RIPE-181++] [RPSL] to describe exterior routing
policies

11.1 Routing Registries/Routing Databases

Organizations who participate in exterior routing usually will have
routing information not only in their routers, but in databases
operated by registries or higher-level service providers (e.g., the
Routing Arbiter).

If an ISP whose previous address space came from a different provider
either renumbers into a different provider's address space, or gains
a recognized block of its own, there may be administrative
requirements to return the previously allocated addresses. These
include changes in IN-ADDR.ARPA delegation, SWIP databases, etc., and
need to be coordinated with the specific registries and providers
involved.   Not all registries and providers have the same policies.

If the enterprise is a registered Autonomous System and renumbers
into a different address space, route objects with old prefixes in
routing registries need to be deleted and route objects with new
prefixes need to be added.

11.2 BGP--Own Organization

IP addressing information can be hard-coded in several aspects of a
BGP speaker. These include:

    1. Router ID
    2. Peer router IP addresses
    3. Advertised prefix lists
    4. Route filtering rules

Some tools exist [RtConfig] for generating policy configuration part
of BGP router configuration statements from the policies specified in
RIPE-181 or RPSL.

11.3 BGP--Other AS

   Other autonomous systems, including nonadjacent ones, can contain
   direct or indirect (e.g., aggregated) references to the above routing
   information. Tools exist that can do preliminary checking of
   connectivity to given external destinations [RADB].

12. Network Management

   This section is intended to deal with those parts of network
   management that are intimately associated with routers, rather than a
   general discussion of renumbering and network management.

   Methods used for managing routers include telnets to virtual console
   ports, SNMP, and TFTP. Network management scripts may contain hard-
   coded references to IP addresses supporting these services. In
   general, try to convert script references to IP addresses to DNS
   names.

   A critical and complex problem will be converting SNMP databases,
   which are usually organized by IP address.

12.1 Configuration Management

   Names and addresses of servers that participate in configuration
   management may need to change, as well as the contents of the
   configurations they provide. TFTP servers are commonly used here, as
   may be SNMP managers.

12.2 Name Resolution/Directory Services

   During renumbering, it will probably be useful to assign DNS names to
   interfaces, virtual interfaces, and router IDs of routers. Remember
   that it is perfectly acceptable to identify internal interfaces with
   RFC1597/RFC1918 private addresses, as long as firewalling or other
   filtering prevent these addresses to be propagated outside the
   enterprise.

   If dynamic addressing is used, dynamic DNS should be considered.
   Since this is under development, it may be appropriate to consider
   proprietary means to learn what addresses have been assigned
   dynamically, so they can be pinged or otherwise managed.

   Also remember that some name resolution may be done by static tables
   that are part of router configurations. Changing the DNS entries,
   and even restarting the routers, will not change these.

12.3 Fault Management

   Abnormal condition indications can be sent to several places that may
   have hard-coded IP addresses, such as SNMP trap servers, syslogd
   servers, etc.

   It should be remembered that large bursts of transient errors may be
   caused as part of address cutover in renumbering. Be aware that
   these bursts might overrun the capacity of logging files, and
   conceivably cause loss of auditing information. Consider enlarging
   files or otherwise protecting them during cutover.

12.4 Performance Management

   Performance information can be recorded in routers themselves, and
   retrieved by network management scripts. Other performance
   information may be sent to syslogd, or be kept in SNMP data bases.

   Load-generating scripts used for performance testing may contain
   hard-coded IP addresses. Look carefully for scripts that contain
   executable code for generating ranges of test addresses. Such
   scripts may, at first examination, not appear to contain explicit IP
   addresses. They may, for example, contain a "seed" address used with
   an incrementing loop.

12.5 Accounting Management

   Accounting records may be sent periodically to syslogd or as SNMP
   traps. Alternatively, the SNMP manager or other management
   applications may periodically poll accounting information in routers,
   and thus contain hard-coded IP addresses.

12.6 Security Management

   Security management includes logging, authentication, filtering, and
   access control. Routers can have hard-coded references to servers
   for any of these functions.

   In addition, routers commonly will contain filters containing
   security-related rules. These rules are apt to need explicit
   recoding, since they tend to operate on a bit level.

   Some authentication servers and filtering mechanisms may dynamically
   update router filters.

12.7 Time Service

   Hard-coded references to NTP servers should be changed to DNS when
   possible, and renumbered otherwise.

13. IP and Protocol Encapsulation

   IP packets can be routed to provide connectivity for non-IP
   protocols, or for IP traffic with addresses not consistent with the
   active routing environment. Such encapsulating functions usually
   have a tunneling model, where an end-to-end connection between two
   "passenger" protocol addresses is mapped to a pair of endpoint IP
   addresses.   Generic Route Encapsulation is a representative means of
   such tunneling [RFC1701, RFC1702].

13.1 Present

   Renumbering of the primary IP environment often does not mean that
   passenger protocol addresses need to change. In fact, such protocol
   encapsulation for IP traffic may be a very viable method for handling
   legacy systems that cannot easily be renumbered. For this legacy
   case, the legacy IP addresses can be tunneled over the renumbered
   routing environment.

   Also note that IP may be a passenger protocol over non-IP systems
   using IPX, AppleTalk, etc.

13.2 Future

   Tunneling mechanisms are fundamental for the planned transition of
   IPv4 to IPv6. As part of an IPv4 renumbering effort, it may be
   worthwhile to reserve some address space for future IPv6 tunnels.

   While there are clear and immediate needs for IPv4 renumbering, there
   may be cases where IPv4 renumbering can be deferred for some months
   or years. If the effort is deferred, it may be prudent at that time
   to consider if available IPv6 implementations or tunneling mechanisms
   form viable alternatives to IPv4 renumbering. It might be
   appropriate to renumber certain parts of the existing IPv4 space
   directly into the IPv6 space. Tools for this purpose are
   experimental at the time this document was written.

14. Security Considerations

    Routers are critical parts of firewalls, and are otherwise used for
    security enforcement. Configuration errors made during renumbering
    can expose systems to malicious intruders, or deny service to
    authorized users. The most critical area of concern is that filters
    are configured properly for old and new address, but other numbers
    also can impact security, such as pointers to authentication,
    logging, and DNS servers.

    During a renumbering operation, it may be appropriate to introduce
    authentication mechanisms for routing updates.

15. Planning and Implementing the Renumbering

    Much of the effort in renumbering will be on platforms other than
    routers. Nevertheless, routers are a key part of any renumbering
    effort.

    Step 1--Inventory of affected addresses and names.

    Step 2--Design any needed topological changes. If temporary address
            space, network address translators, etc., are needed, obtain
            them.

    Step 3--Install and test changes to make the network more
            renumbering-friendly. These include making maximum use of
            default routes and summarization, while minimizing address-
            based references to servers.

    Step 4--Plan the actual renumbering. Should it be phased or total?
            Can it be done in a series of stub network renumberings,
            possibly with secondary addresses on core routers? Is NAT
            appropriate? If so, how is it to be used?

            What is your plan of retreat if major problems develop?
            Make a distinction between problems in the routing system
            and unforeseen problems in hosts affected by renumbering.

    Step 5--Take final backups.

   Step 6--Cut over addresses and names, or begin coexistence.

        Make needed DNS and firewall changes.
        Restart routers and servers as appropriate.
        Clear caches as appropriate.
        Remember static name definitions in routers may not be affected
          by DNS changes.
        Coordinate changes with affected external organizations (e.g.,
          ISPs, business partners, routing registries)

   Step 6--Document what isn't already documented. Make notes to help
        the person who next needs to renumber. Share experience with
        the PIER working group or other appropriate organizations.

15.1 Applying Changes

   Renumbering changes should be introduced with care into operational
   networks. For changes to take effect, it is likely that at least
   interfaces and probably routers will have to be restarted. The
   sequence in which changes are applied must be carefully thought out,
   to avoid loss of connectivity, routing loops, etc., while the
   renumbering is in process.

   See case studies presented to the PIER Working Group for examples of
   operational renumbering experience. Organizations that have
   undergone renumbering have had to pay careful attention to informing
   users of possible outages, coordinating changes among multiple sites,
   etc. It will be an organization-specific decision whether router
   renumbering can be implemented incrementally or must be done in a
   major "flag day" conversion.

   Before making significant changes, TAKE BACKUPS FIRST of all router
   configuration files, DNS zone files, and other information that
   documents your present environment.

15.2 Configuration Control

   Operationally, an important part of renumbering and continued
   numbering maintenance is not to rely on local router interfaces,
   either command language interpreter, menu-based, or graphic, for the
   more sophisticated aspects of configuration, but to do primary
   configuration (and changes) on an appropriate workstation. On a
   workstation or other general-purpose computer, configuration files
   can be edited, listed, processed with macro processors and other
   tools, etc. Source code control tools can be used on the router
   configuration files.

Once the configuration file is defined for a router, mechanisms for
loading it vary with the specific router implementation. In general,
these will include a file transfer using FTP or TFTP into a
configuration file on the router, SNMP SET commands, or logging in to
the router as a remote console and using a terminal emulator to
upload the new configuration under the router's interactive
configuration mode. Original acquisition of legacy configuration
files is the inverse of this process.

15.3 Avoiding Instability

Routing processes tend towards instability when they suddenly need to
handle very large numbers of updates, as might occur if a "flag day"
cutover is not carefully planned. A general guideline is to make
changes in only one part of a routing hierarchy at a time.

Routing system design should be hierarchical in all but the smallest
domains. While OSPF and IS-IS have explicit area-based hierarchical
models, hierarchical principles can be used with most implementations
of modern routing protocols. Hierarchy can be imposed on a protocol
such as RIPv2 or EIGRP by judicious use of route aggregation, routing
advertisement filtering, etc.

Respecting a hierarchical model during renumbering means such things
as renumbering a "stub" part of the routing domain and letting that
part stabilize before changing other parts. Alternatively, it may be
reasonable to add new numbers to the backbone, allowing it to
converge, renumbering stubs, and then removing old numbers from the
backbone. Obviously, these guidelines are most practical when there
is a distinct old and new address space without overlaps. If a block
of addresses must simply be reassigned, some loss of service must be
expected.

16. Acknowledgments

Thanks to Jim Bound, Paul Ferguson, Geert Jan de Groot, Roger Fajman,
Matt Holdrege, Dorian Kim, Walt Lazear, Eliot Lear, Will Leland, and
Bill Manning for advice and comments.

17. References

   [RFC2071] Ferguson, P., and H. Berkowitz, "Network Renumbering
   Overview: Why would I want it and what is it anyway?", RFC 2071,
   January 1997.

   [Cansever] Cansever, D., "NHRP Protocol Applicability Statement",
   Work in Progress.

   [Katz] Luciani, J., Katz, D., Piscitello, D., and Cole, B., "NBMA Next
   Hop Resolution Protocol (NHRP)", Work in Progress.

   [Hubbard] Hubbard, K., Kosters, M., Conrad, D., Karrenberg, D., and J.
   Postel, "INTERNET REGISTRY IP ALLOCATION GUIDELINES", BCP 12, RFC
   2050, November 1996.

   [RFC1631] Egevang,, K., and P. Francis, "The IP Network Address
   Translator (NAT)", RFC 1631, May 1994.

   [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., de Groot, G-J.,
   and E. Lear, "Address Allocation for Private Internets", RFC 1918,
   February 1996.

   [RFC1900] Carpenter, B., and Y. Rekhter, "Renumbering Needs Work", RFC
   1900, February 1996.

   [RPS] Alaettinoglu, C., Bates, T., Gerich, E., Terpstra, M., and C.
   Villamizer, "Routing Policy Specification Language", Work in Progress.

   [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC
   1812, June 1995.

   [Rigney] Rigney, C., Rubens, A., Simpson, W., and S. Willens, "Remote
   Authentication Dial In User Service (RADIUS)", RFC 2058, January 1997.

   [Carpenter] Message to PIER Mailing List, see PIER Archives

   [Lear] Message to PIER Mailing List, see PIER Archives

   [deGroot]   Message to PIER Mailing List, see PIER Archives

   [Wobus] "DHCP FAQ Memo",
   http://web.syr.edu/~jmwobus/comfaqs/dhcp.faq.html

18. Author's Address

   Howard C. Berkowitz
   PSC International
   1600 Spring Hill Road, Suite 310
   Vienna VA 22182

   Phone: +1 703 998 5819
   EMail: hcb@clark.net