Financial Incentives for Route Aggregation and Efficient Address Utilization in the Internet

Yakov Rekhter (Cisco) yakov@cisco.com

Paul Resnick (AT&T Labs-- Research) presnick@research.att.com

Steve Bellovin (AT&T Labs-- Research) smb@research.att.com

Abstract

Growth of the Internet is limited both by the ability of the Internet routing system to scale, and by the availability of IP addresses that are unique within the Internet. Unfortunately, scaleable routing and efficient address space utilization sometimes require painful renumbering of existing hosts to new addresses or otherwise undesirable address allocations. Financial incentives, in addition to the existing methods of persuasion and coercion, can motivate IP address assignment that is efficient both with respect to its suitability for aggregation (via hierarchical routing), and with respect to address space utilization. Where trade-off must be made between conflicting goals, financial incentives will permit local decisions that take into account local differences, thus leading to better choices than could be made by any centralized administrative body. This paper presents a framework for property rights and contracts so that prices can arise through natural market forces, without the need for a global authority or tax collector.

1.0 Motivations

The goal is to enable Internet growth through the smallest resource expenditures possible. The Internet consists of many independent decision-makers whose individual actions can have global effects, such as limiting future growth or imposing costs on other participants. The most effective way to induce socially responsible behavior is to introduce financial incentives that make global effects visible to individual decision makers.

1.1 Growth While Minimizing Resource Utilization

When a new host computer joins the Internet, it needs full connectivity and an unambiguous address. That is, from every existing host there must be at least one path by which packets can travel to the new host. At each step along each of those paths, a router must examine the new host's address and forward packets to the correct next hop. There are two limits, then, on growth. First, there must be an unambiguous address available for each new host. In particular, the IP address used by the new host should never be simultaneously used by any other host. Second, many routers must keep track of an extra ìrouteî, an extra entry in the table of where to forward packets.

New protocols and technology can expand these limits. For example, dynamic address allocation and Network Address Translators [RFC1631] permit reuse of addresses in some circumstances and IPv6 [RFC1883] has longer, and hence more, addresses. At least in principle, routers with more memory could keep track of extra routes. Protocol changes and additional technology, however, are costly and not always available when needed. The technology growth curve may also be unable to keep up with the growth of the Internet. Sometimes it will be either cost effective or just necessary to make more efficient use of existing resources instead.

Hierarchical route aggregation [Kleinrock:77] is one way to make more efficient use of existing resources. In particular, if all packets destined for addresses with the same prefix get forwarded to the same next hop, a router can maintain just one entry in its router table for that prefix, rather than a separate entry for each individual address. This method can be applied recursively to create larger and larger address blocks that share shorter and shorter prefixes.

Hierarchical route aggregation is widely used on the Internet today and was the main motivation behind CIDR [RFC1518, RFC1519]. Hierarchical aggregation only works, however, if traffic destined to all the addresses covered by a particular address prefix should be routed to the same next hop. Thus, the assignment of addresses to hosts must follow the connection topology of the network in order for hierarchical aggregation to be successful.

Even when address assignments match connection topology sufficiently at one point in time, itís hard to keep them aligned. Difficulties arise both from networks that grow and shrink, and from changes in network topology that occur when customers switch providers or peering arrangements change.

Consider first a growth scenario. An organization initially has some hosts to connect, but at a later time may need to connect an as-yet unknown number of additional hosts. There are three ways to handle this, as shown in Figure 1 and described below:

1. Assign the block of addresses but, in anticipation of future needs, reserve spare addresses with the same prefix. When the reserved addresses are eventually used, they all can be routed as a single aggregate. For example, the current RIPE NCC practice is to allocate a 19-bit address prefix to a new ISP, but to reserve the companion 19-bit address prefix, so that it can also be given to that ISP in the future, thus creating a single 18-bit address prefix that can all be routed to the same destination. This wastes addresses, however, since the ISP may never need the reserved addresses, and doesnít completely solve the problem, since the ISP may need more addresses than those that have been reserved.
2. When additional addresses are needed, renumber the existing hosts into a larger block that accommodates both the existing and additional hosts. The cost of renumbering the existing hosts, however, may be quite high. In many organizations, this requires manual entry of a new address in each host, and administrators may no longer remember how to enter addresses for hosts such as printers that have been running successfully for years. Horror stories abound of networks shut down for hours or even days during a changeover.
3. When additional addresses are needed, allocate them separately and advertise an extra route. This is a common practice, but it imposes a cost on the organizations that must carry the additional route.

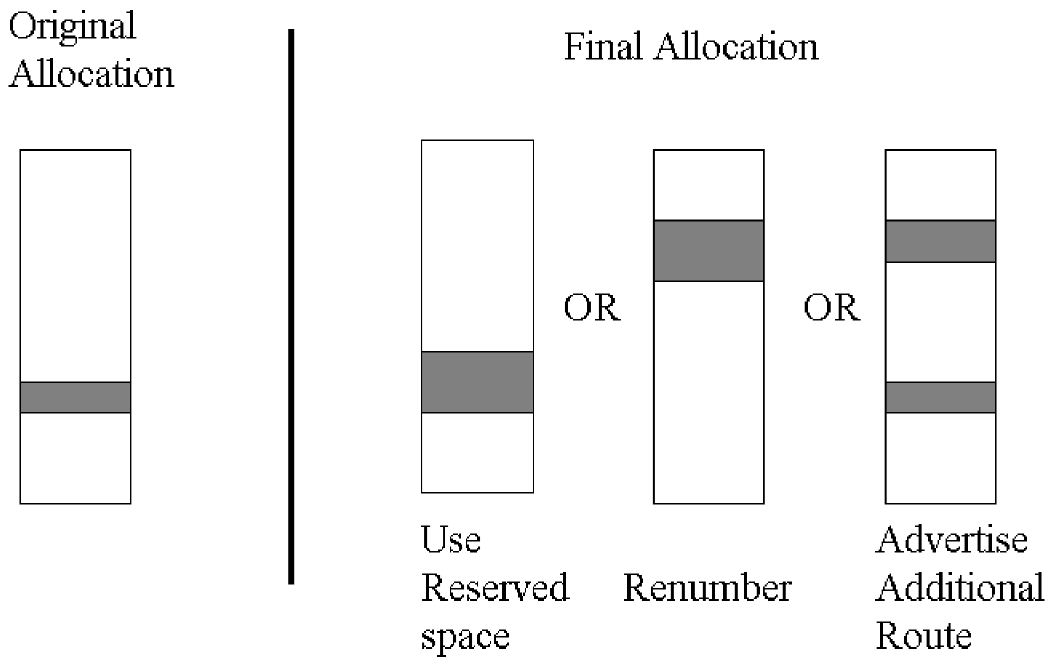

Figure 1: Three ways to accommodate growth. Each vertical block represents the entire address space; the small rectangles represent address space used by a particular organization. The first two methods enable advertisement of an aggregate route, because the addresses used are contiguous.

The best choice among the three growth options will depend on time-varying and local circumstances. If addresses are plentiful, the inefficient address space utilization of the first method may be acceptable. If renumbering costs are low, the second may be appropriate. If routers are not overloaded, the last method may be fine.

Consider also how to accommodate changes in network topology. In particular, suppose that an organization switches Internet providers. If the organization keeps the same addresses (such addresses are called "portable"), then the new provider will have to announce a separate route for those addresses, as shown in Figure 2. On the other hand, switching to addresses that the new provider can easily aggregate requires renumbering of existing hosts. Again, the best choice will involve tradeoffs between the costs of handling an additional route throughout the Internet and the costs of renumbering.

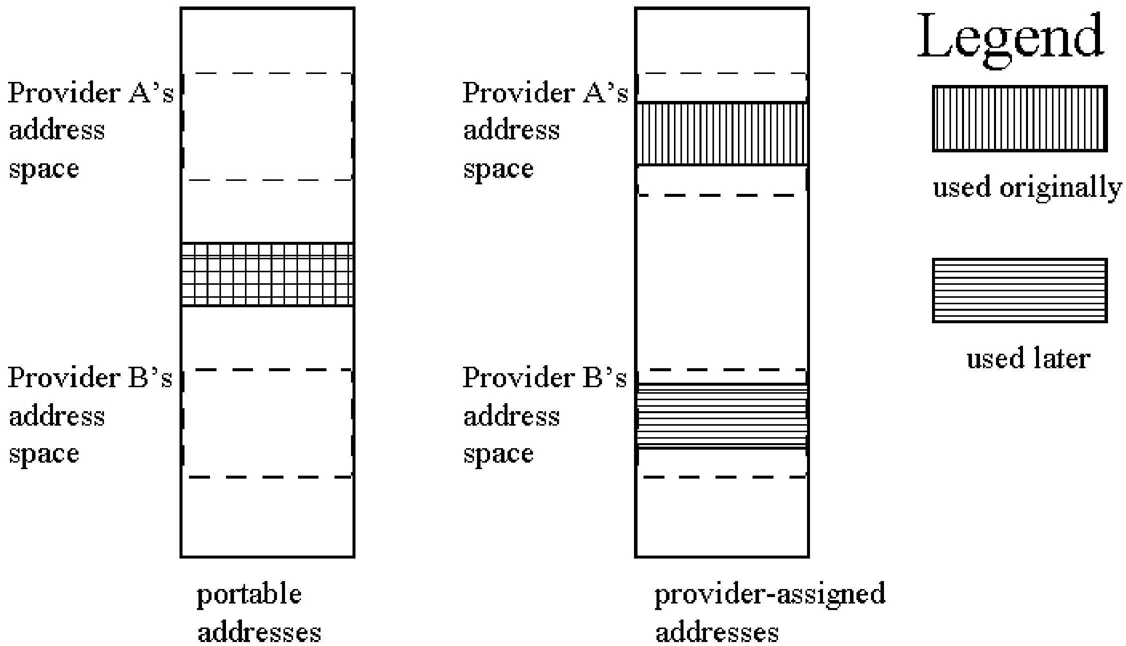

Figure 2: Portable addresses avoid renumbering, but complicate routing.

1.2 Distributed Decision Making

To complicate things further, the Internet consists of many independent participants (e.g., providers, subscribers). The self-interest of these participants may not always align with the interests of the Internet as a whole. In the growth scenario, the costs of reserving address space accrue to others who are not able to use those addresses, and the costs of additional routes are borne by routers throughout the network. The costs of renumbering, however, are borne locally. This creates an incentive for a decision maker, acting in its own local self interest, either to reserve lots of addresses or to advertise additional routes, rather than renumber. Similarly, "portable" addresses impose a cost on the rest of the Internet, but are attractive to a subscriber because they reduce the local cost of renumbering.

1.2.1 The Spirit of Cooperation

The ìspirit of cooperationî is often invoked when self-interest does not align with the larger interest. It calls on everyone to cooperate with each other and do whatever is necessary for the ìgoodness of the Internetî. This call to altruism has been remarkably effective so far: the Internet is still operational is largely due to the spirit of cooperation.

It would be unwise, however, to rely solely on this spirit, especially as the Internet grows larger and anonymous commercial ties replace the personal ties that bound network operators together in the old days. In a competitive environment, moreover, organizations that place the common interest ahead of their own self-interest will be at a disadvantage.

1.2.2 Administrative Policies

Sometimes rules, or policies, can align individual behavior with the common good despite a misalignment of interests. While reliance on policies solves some of the problems associated with the spirit of cooperation, it introduces its own set of problems.

To begin with, it is becoming more difficult to agree on policies as the Internet grows. There is no clear chain of authority for dictating policies. Consensus-based processes move slowly and may result in weak policies that are not effective.

Second, policies are difficult to enforce. Even where there is a well-recognized authority, as with the address registries today, the authorities may not have sufficient resources to investigate whether individual entities are following the rules. For example, it is difficult for registries to verify whether addresses are used for the purposes stated by the organization that requested them.

A final problem with policies is that it is difficult to make them reflect the heterogeneity of the Internet. A policy about advertising aggregate routes that is appropriate for organizations with low renumbering costs may not be appropriate for those with higher costs. Similarly, a blanket policy determining the number of addresses to allocate to a new provider does not take into account the likely growth of the provider, while an allocation policy that takes into account business plans requires sophisticated, and sometimes subjective, decision-making rules.

1.2.3 Financial Incentives

Alternatively, financial incentives can align individual self-interest with the global interest. In the scenarios above, an organization would reserve addresses or advertise additional routes only when its own benefits outweigh the global costs. As an added feature, the money collected can compensate those entities that actually bear the costs.

Decisions based on financial incentives will naturally take into account local and time-varying factors. If there are charges both for addresses and for advertisement of routes, organizations will trade off these costs against their own costs of renumbering. Since the cost of renumbering will vary between organizations, different organizations will make different tradeoffs. If the charges for route advertisements and addresses change over time, reflecting changes in scarcity, it may cause organizations to make different tradeoffs.

2.0 Charging for route advertisements

The primary goal of introducing charging for route advertisement is to limit the number of routes within the Internet routing system, while still maintaining full IP connectivity. To put it differently, the goal is to propagate each route to as few domains as possible, while maintaining connectivity among as many domains as possible.

2.1 Route ìpushî vs. route ìpullî

In order to understand the framework for charging, it is helpful to understand how routing information disseminates in the Internet. In particular, it is helpful to distinguish between the distribution of routing information that is necessary to provide connectivity, and routing information that results in improved connectivity (better paths). Figure 3 illustrates both types of route distribution.

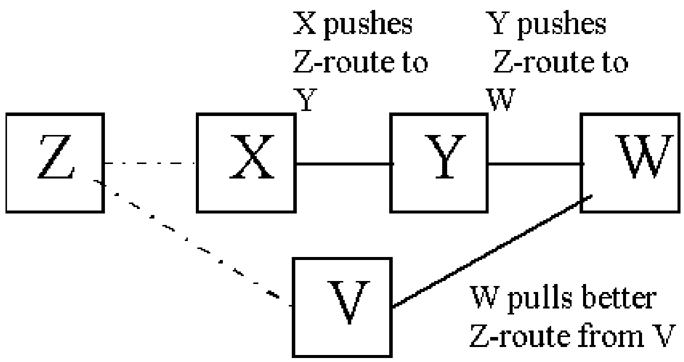

Figure 3: X has a path to Z and pushes it to its neighbor Y, which pushes it on to W. W may also choose to pull a better route to Z from V.

Route ìpushî occurs when one provider knows a path to some destination and offers to neighboring providers to forward transit traffic to that destination. Each of the neighbors can carry on the process by pushing the route to its neighbors. This is the natural way to expand the region of hosts with connectivity to a destination. In the context of hierarchical routing a route has to be pushed only until it is aggregated with other routes whose destination addresses share a common prefix. In the worst case the route may never be aggregated and has to be pushed to every provider in the default-free zone of the Internet.

Route ìpullî occurs when a provider chooses to use a better route to a destination than the route that was pushed to it, asking other providers along the new route to provide transit service. Thus, route pull is not necessary to provide connectivity. A pull route might reflect private peer-to-peer connections between two ISPs, or even two end-users. If one provider pulls a route from the other, it can use that private connection for such traffic, rather than sending it via the perhaps slower route the data would otherwise take. The decision on when to ìpullî a route is influenced by the tradeoffs between path optimality and the possibly increased volume of routing information if the pull involves disaggregating a large block of addresses that were all pushed as a single route.

2.2 Use of bilateral agreements

Both route push and pull cause resources to be consumed at providers all along a route: at every hop, a router needs to record the route advertisement in the forwarding table, and to forward packets. Yet it is only the endpoints of the route that benefit from the connectivity. Moreover, choices that the endpoints make, such as using addresses that permit route aggregation, or requesting the pull of optimal paths, influence the amount of resources consumed at the transit points. One possibility would be to devise charges whereby the endpoints directly compensate all the transit providers. Such a scheme, however, would not scale well, as it could involve settlements between every pair of providers on the Internet, even those who do not exchange packets directly.

Instead, route dissemination (and hence financial charges) can be realized as a composition of bilateral agreements between directly connected organizations. Such agreements can be made between providers and subscribers or between peers. Composition (including subcontracting) of bilateral agreements enables extension of the scope of ìpushî or ìpullî beyond those who are directly connected.. Each organization, however, need only negotiate and settle charges with those organizations directly connected to it. This follows the principle of "edge pricing" articulated in [Shenker et al.]

The Push Contract

Suppose X has a route to Z. For X to push the route to another provider Y (directly connected to X), X and Y would sign a ìcontractî of the form, ìI, Y, promise that the following providers ____, will be able to reach destination Z by sending traffic to me (Y), which I will forward to X.î Typically, Y will fulfill its promise by subcontracting with its neighbors to further push the route.

The Pull Contract

Suppose V has a route to Z. For a provider W (directly connected to V) to pull a route from V, W and V would sign a ìcontractî of the form, ìI, V, promise to forward traffic from W to Z via the following sequence of providers ___________.î This requires a bilateral agreement between V and the first organization in the sequence, which then has a subcontracting agreement with the next organization in the sequence, etc.

When an organization Y contracts to accept traffic from another organization X (through either a push or pull contract), the routing information controls what destinations X can reach through Y. This would impact the amount of traffic Y would receive from X. For example, X could push to Y just one route stating that it accepts traffic for all destinations (a default route), but that could result in X carrying a lot of transit traffic.

Both the sign and magnitude of payments accompanying a contract are open to negotiation. Since a contract involves X adding an entry to its router forwarding table and Y agreeing to carry transit traffic, the settlement payment could flow in either direction, depending on which party most wants the agreement. One typical scenario is a subscriber paying a provider to push its route to the entire Internet. The provider then disburses some of that money to its neighbors when it subcontracts with them to push the route. The recursive subcontracting provides a mechanism to disburse the subscriber's payment among all the providers who keep track of the subscriber's route and carry transit traffic to the subscriber.

2.3 Interaction with hierarchical routing

Any time routes to two or more sets of destinations are aggregated, a neighboring provider that accepts a push of the aggregated route will need to keep only one entry in its forwarding table. If memory in the forwarding table is scarce, as it is today in many Internet routers, the charge for pushing an aggregate route will be smaller than the combined charges for pushing the separate components. This will create a natural incentive for aggregating routes.

2.4 Interaction with address portability

Recall the scenario of a subscriber switching providers and considering whether to renumber to addresses in the new provider's pool or whether to keep the old addresses. Renumbering to the new provider's pool will permit the new provider to push the subscriber's route as part of an aggregate. Keeping the old addresses will force the provider to push a separate route. If router table space is scarce, the provider would have to pay extra for pushing a separate route, and will presumably pass this cost on to the subscriber in some form. One likely outcome is that subscribers who accept provider-based addressing will pay less than they do today, while those who use portable addresses and advertise additional routes will pay more than they do today. The optimal choice still will depend on local factors (such as the cost of renumbering for that subscriber), but the subscriber will take into account the costs imposed on the rest of the Internet.

3.0 Charging for Addresses

While charging for route advertisement may motivate address assignment that is efficient with respect to its suitability for aggregation, it does little (or nothing) to produce efficient address space utilization. Separate charges for addresses would motivate organizations not to reserve too many addresses for future growth and to transfer addresses that are no longer needed.

3.1 ìTransferable address ownershipî policy

Currently there are two address allocation policies in the Internet, portable and non-portable (ìaddress ownershipî and ìaddress lendingî in the terms of [RFCXXX].) Portable addresses are acquired directly from Internet Registries, including InterNIC, RIPE NCC, and APNIC. When an organization acquires portable addresses, the organization can keep the addresses when switching providers. They are not, however, transferable to another organization. The only officially sanctioned transfer is to return an address back to the Internet Registry that provided the initial allocation.

A second policy, address lending, covers non-portable addresses. Internet Service Providers lend addresses to their subscribers, but the transfers are temporary and are coupled with connectivity between the provider and its subscribers. That is, the provider lends addresses to a subscriber as part of a contract to provide routing to the rest of the Internet.

Neither of these policies is ideal for introduction of charges for addresses. The registries could impose a tax on portable addresses but there would be no market forces operating to discover the appropriate level of charges. Charges are a natural part of the service contract when a provider lends addresses to a subscriber, but then when a subscriber changes its provider it has no way to trade off the costs of renumbering against additional routing costs that come with address portability.

A third policy, ìtransferable ownershipî, is needed, which combines the best features of both policies. It adds transferability to the portable address policy: an organization can give away, barter, or sell addresses that it owns. It adds portability and indefinite duration to the lending of non-portable addresses: a subscriber can keep its addresses when switching providers.

If addresses are scarce, it is likely that money will change hands in compensation for address transfers, but no central authority need decree a price. The monetary value of addresses will create a natural incentive for minimizing their use. An organization that can, at some inconvenience, get by with fewer addresses, may choose to sell some, depending on the level of inconvenience and the going price. An organization that wants to reserve addresses for planned future growth will be able to do so, but only if it is willing to pay for the addresses, thus compensating the organization that gives them up.

If, on the other hand, addresses are not scarce, perhaps because of a transition to IPv6, the market price may drop nearly to 0. In that case, there will be no incentive for conservation but without scarcity there would also be no reason for conservation.

3.2 Scope of ownership

Transferable address ownership is defined with respect to a particular set of organizations (what we'll call a "club") that honors the ownership right. Ownership of an address implies exclusive use of that number as a destination address for routing within the club. Ownership does not imply that data packets bearing that address will be routed to a host owned by the address owner, merely that data will not be routed anywhere else. In essence, the property right results in unique address assignment within the club.

An operational definition of the address property right can be expressed in terms of a database consisting of ownership deeds for blocks of IP addresses. The club is defined as the set of organizations who agree to respect the database entries, refusing to route data (within the club) destined for a particular address except as authorized to do so by the owner of the address.

The most interesting club, of course, is the "public Internet," operationally defined as those organizations that respect the current NIC databases of address registrations. There appears to be great value in having a single Internet club, so that networks everywhere can be interconnected. Private intranets, however, would be free to form their own clubs and reuse the same IP numbers used in the public Internet. Clubs can even be interconnected using Application Layer Gateways or Network Address Translators (NATs) [RFC1631].

Note that this description of the property right immediately makes clear that no entity can exert control of the IPv4 address space without the cooperation of the Internet Service Providers who form the club. If, for example, the U.S. government claims that it owns the IPv4 address space and tries to raise revenue by auctioning it off, Internet Service Providers could simply agree en masse to form a new club with its own address allocation mechanisms.

3.3 Address ownership without connectivity

The transferable address ownership policy decouples address ownership rights from actual connectivity. Owning a (transferable) address without having global routability may have what economists call ìoption value.î That is, there is value in having the option of global routing some time in the future. For example, this would permit an organization to change its style of firewall, making formerly hidden hosts visible to the Internet, without incurring the cost of renumbering.

3.4 Interaction with hierarchical routing

The market price of address blocks is likely to be influenced by route charges. In particular, a large block of addresses that can be aggregated into a single prefix is likely to sell for more than an equivalent number of addresses that do not share a common prefix because it will be cheaper to contract for routing of the large block than the collection of smaller blocks. In fact, the mere prospect of future route charges or the prospect of some providers refusing to route to small address blocks may be sufficient to create an immediate price premium for larger address blocks.

3.5 Address Ownership Registry

If addresses are to be owned by individual parties, it is necessary to have some way to ascertain who owns - and has the right to have routed - each address. At present, registries already maintain databases of allocations; transferability merely requires a method of updating the database when an allocation is transferred. The registry database could list a cryptographic public key with each address block; all change requests would be signed by the corresponding private key. Transaction types would include rekeying, to allow for ownership changes; additionally, address blocks could be subdivided, in which case a new public key would be specified for each new block. If desired, change requests could even be accompanied by anonymous digital cash payments, to reimburse the registry for its administrative overhead.

3.6 Dangers of Hoarding

One danger in any market is the accumulation of power by one or a few players who corner the supply of a scarce resource and artificially inflate its price. While theoretically possible, this is unlikely to happen in the market for IPv4 addresses. There is a natural ceiling on the price that anyone can charge, which is determined by the cost of alternative technologies. In particular, higher prices for addresses will spur the use of NAT boxes and IPv6.

4.0 Conclusion

The current routing and addressing situation in the Internet is inadequate to support continuous growth. Space in router tables and unallocated IPv4 addresses are scarce enough to have prompted conservation measures already, but the measures are unlikely to be sufficient because the problem is a tragedy of the commons: individuals benefit from consuming these resources without directly bearing the costs. Transferable property rights to addresses and bilateral routing contracts can make the costs of common resource depletion visible to individual decision makers. The appropriate level of charges will emerge naturally in the marketplace, based on demand and the degree of resource scarcity. Pricing will not suddenly make free resources costly; it will merely expose the hidden costs, so as to encourage conservation.

5.0 Acknowledgments

We would like to acknowledge Noel Chiappa, David Conrad, John Curran, Mike OíDell, Sean Doran, Peter Ford, Geoff Huston, Tony Li, Peter Lothberg, Daniel Karrenberg, and Andrew Partan.

6.0 References

[Kleinrock:77] L. Kleinrock, K. Farouk, ìHierarchical Routing for Large Networks,î Computer Networks 1 (1977), North-Holland Publishing Company.

[RFC 1518] Y. Rekhter, T. Li, ìAn Architecture for IP Address Allocation with CIDRî.

[RFC 1519] V. Fuller, T. Li, J. Yu, K. Varadhan, ìClassless Inter- Domain Routing (CIDR): an Address Assignment and Aggregation Strategyî.

[RFC1631] P. Francis, K. Egevang, ìThe IP Network Address Translator (NAT)î

[RFC1883] S. Deering, R. Hinden, ìInternet Protocol, Version 6 (IPv6)î

[RFCXXX] Y. Rekhter, T. Li, ìImplications of Various Address Allocation Policies for Internet Routingî

[Shenker, et al] S. Shenker, D. Clark, D. Estrin, and S. Herzog, "Pricing in Computer Networks: Reshaping the Research Agenda," Telecommunications Policy, 1996, vol. 20(3).