Network Working Group
INTERNET-DRAFT
Expire in six months

T. Li
Juniper Networks
Y. Rekhter
cisco Systems

June 1996

Towards a Cost Model for Routing and Addressing

<draft-li-piara-cost-model-00.txt>

Status of this Memo

Introduction

The IP address space is a fixed size, fully recyclable resource which must be shared amongst the entire Internet community in order to achieve global connectivity. Similarly, the routing table entries in the routers that comprise the backbone are a scarce resource. Squandering either of these resources can lead to an Internet in which some systems have significantly reduced reachability.

There are a variety of mechanisms which could be employed to ensure that neither address space nor routing table space is needlessly wasted. Some have proposed [Rekhter, Resnick] that a market based approach is a practical and reasonable allocation of these scarce resources.

This memo is a preliminary investigation into a cost model for routing and addressing services, in an attempt to understand the interaction between the address market and the routing table entry market. As this is a preliminary model, it makes some significant simplifications. It's hoped that this memo will also interest those more skilled in the art of cost modeling to correct and refine this model.

The Intuitive Model

We begin by constructing the model based on the costs that the user sees when provisioning Internet service. We make the simplifying assumption that the user needs to provision a singly homed, non-transit domain (a.k.a., a stub network) that can be addresse entirely by a single prefix (i.e., a contiguous power of two range of addresses). Note that the following analysis can be extended for multiple prefixes by extending the cost functions to take vectors of prefixes as arguments.

In purchasing such an Internet service, we posit that there are three significant costs: the cost of bandwidth, the cost of addressing, and the cost of routing. We'll proceed by describing these components intuitively and then turn to a more formal description. The cost of bandwidth includes the cost of the link, both in installation and periodic charges, and the hardware necessary to support both ends of the link (e.g., routers, CSU/DSU's, cables, etc.). There will also be an administrative component to support the link and other non-addressing and non-routing services, which we will consider to be part of the cost of bandwidth in order to simplify this analysis. Note that this implies that the cost of bandwidth is a function of the bandwidth requested and the service provider that it's requested from. Also note that we make the simplifying assumption that bandwidth from all providers is identical, implying a consistent level of service from all competitors in a local market.

The cost of addressing is the amount that the user must pay to borrow a particular prefix from an address broker. The bandwidth provider may or may not be an address broker. This cost would include any charges for the administrative costs of lending out the prefix, such as DNS PTR record maintenance. Thus, the cost of addressing is a function of the prefix and the address broker. It's likely that the addressing cost is charged periodically.

The cost of routing is the amount that the user must pay to the provider for routing information to the remainder of the Internet. This may involve making configuration changes in a series of routers, in a routing registry, or in the administrative databases of several domains. In the common case where the bandwidth provider is also the address broker, it is likely that the cost of routing is amortized across all customers sharing the particular address block that contains the user's prefix and that the cost of routing is a periodic charge.

Note that not all of these costs may be apparent today. This is not to say that these costs do not exist, just that they're not charged for separately.

The Formal Model

In this section, we describe all of the above in a more formal manner. First, we let the cost of bandwidth be a function Cbw : BW X Prov -> $, where BW is the space of possible bandwidths, Prov is the space of possible providers, and $ is the space of monetary cost. We will also write this as Cbw(bw, P), where bw is the bandwidth and P is the provider.

Second, we let the cost of addressing be a function Ca : Pref X Brok -> $, where Pref is the space of possible address prefixes, and Brok is the space of address brokers. We will also write this as Ca(p, B). Note that the space of address brokers may intersect with the space of service providers.

Third, we let the cost of routing be a function Cr : Pref X Prov -> $. We will write this as Cr(p, P). We assume that Cr is not sensitive the size of prefix p, but it is sensitive to the allocation of prefix p. Thus, it costs the same to route a /16 or a /8 prefix. However, a /16 from another provider may cost more to route that a /16 where P = B.

Finally, we let the total cost be the function C : BW X Prov X Brok X Pref -> $. We will write this as C(bw, P, B, p). Further, as this is simply the sum of the previous three functions, we have:

$$C(bw,P,B,p) = Cbw(bw,P) + Ca(p,B) + Cr(p,P) \qquad (1)$$

Some Observations

We make the simplifying assumptions that the market always has perfect information and that the consumer will always minimize the overall cost C.

Observation 1: If Ca(p,B) is insensitive to the size of prefix p, then address space will be squandered.

In this situation, there is no negative feedback to the user for wasting address space. As there is potential benefit in having future address space available, it is in the user's best interest to overstate their addressing needs. Subsequently, demand for addresses will increase, and Ca(p,B) will rise. If it rises sufficiently, then those who have borrowed an address will subdivide their unused addresses, becoming brokers themselves, thus resulting in a Ca(p,B)

which is sensitive to the size of the prefix. This is to say that an insensitive Ca(p,B) is self-correcting.

Corollary 1: If Ca(p,B) is 0, then address space will be squandered.

Observation 2: If Ca(p,B) << C(bw,P,B,p), then address space will be squandered.

By reasoning similar to that of Observation 1, the marginal cost of unnecessary address space must exceed the user's marginal benefit of such space, given the user's cost sensitivity or address space will be wasted. For example, consider the case where Ca is only 0.01% of C, and doubling the size of the prefix results in doubling Ca. The resulting Ca is only 0.02% of C, which is not a significant deterrent to wasting addresses.

Corollary 2: If Ca(p,B) << Cr(p,P), then address space will be squandered.

Observation 3: We call a prefix which is borrowed from the service provider a 'local prefix'. We call a prefix which borrowed from a broker who is not the service provider a 'foreign prefix'. If Cr(p,P) is insensitive to whether a prefix is local or foreign, then routing table entries will be squandered.

Note that we also make the simplifying assumption that proxy aggregation is not effective. In the above scenario, if the local and foreign prefixes are identical in cost, then it the user will optimize wholly based on the independent costs of bandwidth and address space, obtaining address space from any address broker, regardless of the possibilities of aggregation. Normal entropy at this point will eventually flood the backbone routing tables. Note that a Cr(p,P) which is insensitive to foreign prefixes is also self correcting as it will increase demand on routing table entries, thereby encouraging aggregation.

Corollary 3: If Cr(p,P) = 0 for a foreign prefix, routing table entries will be squandered.

Observation 4: If Cr(p,P) << C(bw,P,B,p) then routing table entries will be squandered.

By reasoning similar to observation 2.

Corollary 4: If Ca(p,B) >> Cr(p,P), then routing table entries will be squandered.

Corollary 5: To encourage conservation, Ca(p,B) and Cr(p,P) must be

proportional to Cbw(bw,P), and Ca(p,B) ~= Cr(p,P).

From observations 2 and 4 and equation (1), it follows that increasing Cbw(bw,P) must also increase Ca(p,B) and Cr(p,P). From corollaries 2 and 4, it follows that Ca(p,B) and Cr(p,P) must be roughly equal.

Observation 5: Within the U.S., Cbw(bw,P) is predominantly a function of the length of the circuit and bandwidth selected.

Hardware costs are generally a small fraction of the line costs involved in providing bandwidth, usually due to periodic or usage charges. These charges are a function of the length of the circuit (or distance called) and the bandwidth of the circuit. The granularity of the length of the circuit varies depending on the exact media (e.g., Frame Relay may be sold for a flat rate anywhere within a LATA. Local loops may be charged by the mile.) Outside the U.S., the local tariff structure may have significant political distortions.

Corollary 6: To encourage conservation within the U.S., Ca(p,B) and Cr(p,P) must be a significant function of the length and bandwidth of the circuit.

Note that this is a non-intuitive result which will be difficult to justify to most customers. We suspect that many providers will not even attempt to do so, instead charging a flat rate for routing services, or basing the charge on prefix length. If this occurs, corollary 5 implies that those who are paying the most for Internet services will see little economic incentive for conservation.

It's also interesting to observe that an address broker who is not also the service provider has a particularly difficult situation. For the broker to be an agent of conservation, she must charge varying amounts based on the cost of bandwidth as charged by the service provider. This coupling of costs between different entities has possibly severe legal and logistic implications, not the least of which is the liability of anti-trust action for price fixing. As a result, it seems as if an address broker can never actually price address space in a manner that is consistent with address space conservation.

Conclusions

This memo has attempted to posit a simple cost model for Internet addressing and routing in order to help understand the possible markets for address space and routing services. The model focused on the costs as seen by an end user. The implications of the model are

rather disturbing. If the end user does see costs that would encourage conservation, then the costs are non-trivial and are proportional to the cost of the bandwidth that he purchases. Further, in such a situation, there appears to be no feasible role for an independent address broker, resulting in users who have to get address space from their service provider.

Alternatively, if a 'natural' market for address space and routing services develops, then the costs for these services are independent of the other costs of provisioning Internet services. The implication is that those who are not cost sensitive will not be well motivated to conserve scarce resources.

We further observe that perfect allocation of scarce resources cannot occur unless there is a perfectly accurate measurement of each users need for address space and routing services. The willingness to pay for a resource is only slightly correlated with true need. Thus, any scheme which depends on cost for resource allocation is inherently flawed and can at best provide a first order approximation. The only alternative to such a cost based scheme is political allocation, which has a variety of its own problems.

Acknowledgements

Tony Li's contribution to this work was supported in part by cisco Systems, Inc.

References

[Rekhter] Rekhter, Yakov, "Charging for Routes", Presentation at the 35th IETF, March 1996

[Resnick] Resnick, Paul, "Suggestions for Market-Based Allocation of IP Address Blocks", Internet Draft, draft-ietf-cidrd-mktbased-alloc-00.txt, Feb. 1996

Authors' Addresses

Tony Li
Juniper Networks, Inc.
101 University Ave Ste 240
Palo Alto, CA 94301
Phone: +1 (415) 614 4145
Email: tli@jnx.com

Yakov Rekhter
cisco Systems, Inc.
170 Tasman Dr.

San Jose, CA 95134
Phone: (914) 528-0090
Email: yakov@cisco.com