

Internet Exterior Routing Protocol Development: Problems, Issues, and Misconceptions

Tim Bass, SAIC, Center for Information Protection

Abstract

Global routing in the Internet continues to have scalability problems which underscore weaknesses in the design and implementation of the various TCP/IP exterior routing protocols. This article explores the historical design and development relative to the decision-making process in the specification and implementation of Internet external routing protocols, in particular discussing problems associated with provider-based address space allocation.

This interim solution in no way constrains the selection of the next generation addressing and routing technology. — RFC 1367, October 1992 [1]

In 1977, Kleinrock and Kamoun [2] published a detailed discussion of hierarchical routing in large internetworks. In their landmark paper, “Hierarchical Routing for Large Networks,” they discussed the problems associated with scalability and the length of the route forwarding table in large packet-switched networks. They concluded that hierarchical routing was a requirement for very large networks. Routing tables require less information for destination networks *far* from the source relative to hop count or another nearness metric. Detailed routing information, one entry per destination, is required for *near* nodes. The construction of logical sets in the routing hierarchy is not related to geographical distances. Instead, administrative clustering of routing domains becomes the technical basis for designing large scalable internetworks.

Problems and issues with scalability in packet-switched networks were well documented prior to 1982, the publication date of Rosen’s Internet Request for Comment, “Exterior Gateway Protocol” (EGP) [3]. In this publication, Rosen discussed hierarchical routing and the limitations of EGP clearly and concisely, cautioning the reader that EGP did not “constitute a network routing algorithm.” Unable to predict the explosive growth of the Internet, the National Science Foundation (NSF), Internet Engineering Task Force (IETF), Internet Activities Board (IAB), Internet Protocol (IP) service providers, and router vendors embarked on a development track redesigning EGP in favor of the Border Gateway Protocol (BGP). Policy-based routing dominated requirements for global scalability in large packet switched networks.

Further inspection confirms that the early IP exterior routing protocol development track by Internet developers and router vendors was dominated by concerns about policy-based routing [4–6]. It was not until 1992 that the engineering focus in the IETF apparently began to shift from policy-based routing to global network scalability [7–9].

A short synopsis of IP exterior routing protocol development is presented in the next section, highlighting EGP, BGP,

and policy-based routing in the Internet. The section after that introduces the paradigm shift from a concentration on policy-based routing to an emphasis on global IP scalability. The focus then returns to the BGP-3→BGP-4 paradigm shift associated with classless interdomain routing (CIDR). Important quotations in the Internet Requests for Comments (RFCs) are highlighted as the discussion develops.

The fifth section presents a generic global hierarchical routing model which is introduced as a basis for discussion relative to both provider and geographic IP address allocation paradigms. The discussion highlights important statements in the Internet RFCs related to the topic, and the developed model demonstrates weaknesses in provider-based address space allocation. The sidebar shows a brief summary of historical call routing in the pre-1984 Bell System and the present telephony hierarchical model. At the end of the article, the presented models, relative analysis, and business issues provide a foundation for discussing the problems with a hierarchical routing model for the Internet based on the current implementation of provider-based IP address space aggregation.

Historical Protocol Development Summary

Exterior Gateway Protocol does not in itself constitute a network routing algorithm ... is NOT intended to provide information which could be used as input to a completely general area or hierarchical routing algorithm. It is intended for a set of autonomous systems which are connected in a tree, with no cycles. It does not enable the passing of sufficient information to prevent routing loops if cycles in the topology do exist. — RFC 827, October 1982 [3]

Early in the 1980s, the ARPANET¹ increased in size dramatically, and it became necessary to adopt a hierarchical clustering paradigm [11]; therefore, Rosen [3] published a new routing paradigm, External Gateway Protocol, or EGP, in 1982 as RFC

¹ The Defense Advanced Research Project Agency network consisted of 57 nodes in December 1975 [2].

827. He explicitly stated that EGP was not a generic hierarchical routing protocol and cautioned developers by including strong statements highlighting the limitations of EGP.

Policy routing, motivated by NSF acceptable use policies, became an increasingly imperative topic [6], and to overcome the tree-topology limitations of EGP by pruning routing loops in a meshed AS² graph, a new exterior routing protocol, Border Gateway Protocol, was proposed in RFC 1105 [33]. Loughheed and Rekhter state the proposed functionality of BGP:

“an inter-autonomous system routing protocol ... routing loops may be pruned and policy decisions at an AS level may be enforced.” — RFC 1105, June 1989 [12]

BGP was specified and intended for a meshed, multihomed global Internet [13, 14]. According to the authors of RFC 1164, BGP

“... provides a high degree of control and flexibility for doing interdomain routing while enforcing policy and performance constraints and avoiding routing loops.”

— RFC 1164, June 1990 [13]

EGP was replaced by BGP [12–14] and further evolved to BGP-3 [10, 15] and BGP-4 [16]. BGP, as an exterior routing protocol, was not designed to dynamically exchange routing policies. BGP was primarily designed to provide full global network reachability routing information. This includes the full transit path of autonomous systems (ASs) and manually configurable policy-based routing. BGP exchanges reachability information with other BGP speakers. This includes information on the full AS path to advertised networks constructing a graph of AS connectivity.

The core effort of the BGP development track was concentrated around the development of a centralized Internet managed by manually configurable routing policies. Policy-routing concerns dominated the practical implementation and design of IP exterior routing protocols.

In 1989 there was increased interest in restricting routing information in the NSFNET backbone [4–6] due to NSF policy and bilateral peering agreements between both regional and peer networks. According to RFC 1092, these policies and agreements made policy-based routing “imperative” [6].

Kumar [17] analyzed existing and proposed policy-based routing protocols, including BGP, IDRP, and others, and concluded:

“... there is a serious mis-match in the level of control desired ... and that being offered. ...to match the expectations of AD³ administrators, performance goals, and scalability ...” [17]

BGP allowed for policy-based routing, but these policies were determined by network administrators and enforced according to rules defined in router configuration files, and were not a part of the protocol [18]. BGP configuration files allowed the router to choose between different AS paths when more than one alternative path was possible. Routing policies were primarily related to political, economic, or security concerns [18].

Huitema [11] discusses the NSF acceptable use policy, or AUP, as necessary due to the fact that NSFNET was supported by public subsidies; therefore, commercial users of the Internet were not allowed, by NSF policy, to use the NSFNET backbone for commercial or nonacademic IP traffic.

Based on this policy, Huitema states, some researchers

would pick an IP AS path based on the nature of the conversation, traversing NSFNET if the nature of the exchange was an academic experiment, but another AS path if the same users were discussing commercial contracts [11]. This evolved to requirements for more complex and sophisticated policy-based routing. Hence, NSF policies for acceptable use of NSFNET appear to have been a major force behind the efforts to develop policy routing capabilities in the 1980s and early 1990.

There were other organizations, including the U.S. Department of Defense and other government organizations, interested in policy-based routing. Steenstrup [19] published RFC 1478 in 1993, describing the interdomain policy routing architecture, or IDPR.⁴ IDPR was designed for heterogeneous internetworks with “tens of thousands” of administrative domains and complex routing policies, further defining transit policies based on:

- Access restrictions
- Quality of service (QoS)
- Economic considerations.

IDPR was a major divergence from the BGP development track, the *de facto* standard of the Internet, and was not adopted by industry. One primary reason for this was that IDPR was incompatible with the BGP development track, which was operational, mature, and commercially supported by major IP router vendors with running code on the Internet. In addition, the IETF considered the complications of IDPR QoS routing enhancements less attractive relative to the simpler BGP paradigm.

From the open systems interconnection (OSI) development track, interdomain routing protocol, or IDRP, has been proposed and developed [11, 20], and is built on the CIDR model with enhancements. IDRP is quite similar to BGP [11] and has merged with the next development iteration of BGP-4. IDRP added enhancements include the aggregation of ASs into “confederations.” This article does not discuss IDRP relative to policy-based routing, but it does discuss a key element of IDRP, the IP address space aggregation paradigm. The reader is referred to the references for further discussion of IDRP.

A strong polemic can be formulated around the following premise: policy-based routing issues dominated IP exterior routing protocol development and detracted from the development of a more robust scalable global routing paradigm. The number of networks in the Internet had increased so dramatically that routers would run out of forwarding table space or grind to a halt processing routing updates. The explosive growth of the Internet forced a paradigm shift, the NSFNET AUP was abandoned, and scalability issues began to dominate the Internet community.

A Paradigm Shift to Scalability Issues

“Sub-allocations of the block ... is beyond the scope of this paper.” — RFC 1366, Oct. 1992 [8]

In July 1992 the IETF considered an IP addressing allocation plan [1] documented in RFC 1366 [8] describing a scheme to super-net unallocated class C IP address space into eight geographical blocks, including Europe, North America, Central and South America, and the Pacific Rim. Class C address space previously allocated was defined as *multiregional*. The remaining three proposed blocks were defined for “flexibility” in consideration of geographical regions previously allocated.

Suballocations within geographic blocks were specified to

² Autonomous system: a network of routers under a single administrative authority managing interior routing of their network, peering with other ASs using an exterior routing protocol.

³ Administrative domain: maps to autonomous system [11].

⁴ Sometimes confused with IDRP, Interdomain Routing Protocol.

be provided within a contiguous eight-bit-wide boundary [8] to allow for easy binary bit-masking and suballocations. RFC 1366 did not, however, discuss or specify the more complex and far-reaching effect of how IP registries would manage the suballocations of these newly defined large geographic blocks. Suballocation or downstream IP address space provisioning, remained for the most part unspecified.

RFC 1338 [21] on IP address *supernetting* was published in June 1992 prior to the July IETF BOF, with the recommendation:

“The NIC should begin to hand out large blocks of class-C addresses to network service providers. ... Service providers will further allocate power-of-two blocks of class-C addresses from their address space to their subscribers.”
— RFC 1338, June 1992 [21]

RFC 1338 was a major paradigm shift to establish a provider-based addressing routing hierarchy that was multi-tiered on binary boundaries. This would become relevant in RFC 1366, not as a provider-based addresses hierarchy, but as a geographical hierarchy.

Kleinrock and Kamoun [2, 22] demonstrated that a hierarchical clustering paradigm was sufficient to meet the scalability criteria of large packet-switched networks. The supernetting scheme in RFC 1338 specified a new direction for the exterior routing protocol development track. In fact, the RFC confirmed another iteration of the policy-based BGP development track.

With the new RFC 1338-style provider-based *supernetting* it was possible to create multiple hierarchical tiers, and most tiers were envisioned by RFC authors to be IP service providers [21]. RFC 1338 did not acknowledge the paradigm shift, explicitly stating:

“This solution does not change the Internet routing and addressing architectures.” — RFC 1338, June 1992 [21]


In actuality, the shift to a provider-based hierarchical cluster was a major change in the Internet routing and addressing architecture. RFC 1366 explicitly stated that aggregation would be accomplished by geographical registration authorities. These guidelines were marginalized by proponents of provider-based IP address space allocation, including large IP providers. A major paradigm shift had occurred in the Internet; provider-based address space allocation was the new model, and BGP would evolve to BGP-4 [16], incorporating the RFC 1338 paradigm. However, for this shift to occur, the technique for supernetting-subnetting the IP address space required modification. This new feature was called *classless interdomain routing*.

BGP-4 and Classless Interdomain Routing

The IAB supports the actions taken by the IANA and the InterNIC, router vendors, and the network operators ... to address the scaling problem ... — RFC 1481, July 1993 [23]

The paradigm of provider-based supernetting [21] in the Internet was the basis for the next evolution of BGP. CIDR [16, 24, 25] was the beginning of a new paradigm for IP service providers.

RFC 1366 guidelines discussed geographic address space management. However, the actual implementation was provider-based IP address blocks. This specific CIDR implementation provided numerous opportunities for large, pre-existing IP service providers, especially organizations with existing relationships with the NSF, to position themselves in



■ Figure 1. Example: meshed top-level superstructure.

the upper tiers of a new CIDR-based global routing hierarchy.

Ford *et al.* [26] discuss BGP-4 and CIDR from the perspective of a global Internet based on a spanning-tree model with a core NSFNET backbone and regional providers' default routing to NSFNET,⁵ stating:

“... NSFNET is hierarchical ... Only a few domains require full routing information of the Internet. ... Ideally, sites will get their IP addresses from their network service

provider ... it does not mandate a strict routing hierarchy.” [26]

To the contrary, BGP-4 with provider-based CIDR allocation requires, and in fact tacitly mandates, a strict hierarchical graph based on upstream service provider IP address allocation. This, in practical business terms, requires mid-level providers and users to adhere to a restrictive policy with regard to IP address space allocation, ownership, and portability [25, 27–31].

Stevens, discussing BGP prior to CIDR, stated:

“... the Internet is then viewed as an arbitrary interconnection of transit, multi-homed and stub ASs.” [18]

The above statement is no longer true with provider-based CIDR because arbitrary interconnections are restricted and actively discouraged. Interconnections are based on a hierarchy of large blocks of classless address space and downstream IP address space allocation. In the following section, we demonstrate why CIDR, as implemented by BGP, is a quasi-rigid hierarchical routing paradigm. Furthermore, based on a simple generic internetworking model, it will be suggested that future hierarchical routing models must not be constrained by a provider-based aggregation paradigm.

Provider-Based IP Address Allocation

Choosing a hierarchical structure is difficult, and it often becomes difficult to change a hierarchy once one has been established.” — Comer, 1991 [32]

Tanenbaum [33] describes hierarchical routing as routers divided into *regions*, where each router in a region knows how to route to every node in their region, but having no knowledge of the internal structure of other regions. This is consistent with Kleinrock and Kamoun's [2] discussion of hierarchical routing as well as the AS/AD concept.

For hypothetical purposes of discussion let us assume a *fully meshed world model* (Fig. 1) with six of the seven continents all fully meshed with huge backbone pipes. The continental regions will be designated as numbers which have no significance, numbered only for the drawing program, counterclockwise:

- 1 Asia
- 2 Europe
- 3 North America
- 4 South America
- 5 Africa
- 6 Australia

Also, assume that Antarctica is not meshed into the global backbone and hence has been omitted. Further assume that the integration of all Asian nations into one routing region is feasible. These assumptions allow the number of meshed global regions to be limited to six and constrains the number of global transit links.

Let us further assume that the entire IP address space is

⁵ See Fig. 3 in [26].

classless; then we enjoy the engineering artistic license to allocate all the classless IP address space to six regional registries corresponding to the six regions of the fully meshed world. The global exterior routes now contribute a handful of routes to the global regional gateways.

The global regional registries may now allocate their classless IP address blocks to *child* registries, who allocate from their subblocks to their children, who further suballocate. In addition, assume that every child is unsymmetrically meshed with its parent and siblings; and every region has at least one fully operational interior, or perhaps a new *trans-regional* routing protocol. This is one possible routing model of many such paradigms which might be created.

Now, let us leave the model above in place and return to the pragmatic *real world*, where networking is a very dynamic process with very demanding children. In fact, all children demand to switch parents easily, quickly, and efficiently. Children become unhappy with their parents and want new parents for a myriad of reasons. Perhaps other parents provide better services or live in regions where the children wish to relocate.

"The greatest degree of negotiating clout will lie with users who generate large traffic volumes and can migrate to other suppliers ..." [34]

For purposes of discussion, let us assume five generations of relations and that each generation is a large service provider. After the original allocation of IP address space the simplified regional routing topology resembles the example hierarchy of Fig. 2.

After some time, provider h , P_h , is attracted to the services offered by P_d . P_h is well established and has two major customers who are large IP service providers. Figure 3 illustrates the new business and routing hierarchical relationships.

It can be observed a priori that IP address space allocated from $P_1 \rightarrow P_a \rightarrow P_h$ continues as address space $P_h \rightarrow P_p \rightarrow P_x$. It also follows that when P_h makes the business decision to discontinue service with P_a and become a customer of P_d ,




Figure 3. Provider P_h changes to provider P_d .




Figure 4. Provider P_i changes to provider P_d .




Figure 2. An example of regional network provider relations.

based on the IETF guidelines previously cited, P_h and all downstream customers of P_h , including other providers who may have large customer bases themselves, must renumber their networks.

As a final plausible business example, P_i is also impressed with the services and pricing structure of P_d and, following P_h , decides to become a customer of P_d . Figure 4 illustrates this change by P_i to P_d from P_e . In this example, assume that P_i is moving their network management center from a geographic location in region 1 to another location within region 1, from Japan to Hong Kong, for example.

It follows, as in the first example, that downstream customers of P_i , including other providers who have large customer bases, must renumber their IP networks.

Address Translation and Renumbering

With the implementation of the supernetting paradigm of RFC 1338 in the Internet, provider-based address allocation, IP customers who prefer to change service providers are required to either renumber their networks or implement IP address translation schemes. RFC 1631 [35] introduced the concept of network address translation (NAT) to mitigate the problems associated with IP AS mobility, as illustrated in Figs. 2-4. In addition to address translation, software has been proposed, and software tools are under development, to automate address space renumbering tasks.

Both address translation and renumbering have practical problems. This can be seen a priori in Fig. 3 when P_h changes to P_d from the original upstream service provider, P_a . In this case, all the downstream IP customers of P_h must either renumber their networks or translate based on the allocated address space. In the provider-based paradigm, from Fig. 2, P_h , P_p , P_q , P_x , and P_y all are subelements of the address allocation of P_a . Hence, it is an impractical administrative task to renumber or translate numerous regional providers, ASs, and countless end-user customers.

Discussions in the IETF [28, 36] normally consider the special case of renumbering when, for example, P_y changes to provider P_w from P_p . In this special case, the impact of renumbering may be dramatic, but limited, affecting a single lower-tier IP service provider or stub AS. However, in the general case, when any large provider supporting other IP providers in a meshed global model desires to reposition itself with a new service provider, economic and business issues will overshadow technical considerations. Renumbering proves impractical and costly from numerous perspectives.

Address translation is problematic as well. For example, numerous Transmission Control Protocol (TCP)/IP protocols use checksums to detect transmission errors; checksums must be recalculated for most translated datagrams. Translating large regions, for example, when P_h chooses to move to P_d in Fig. 3, would impose a heavy processing burden on busy routers performing regional address translation. Under this architecture, scalability mitigation is shifted from the physical size of the route forwarding table to processing

speed. This causes slower packet forwarding and increased router latency.

The simplified model presented does not illustrate the more complex and pragmatic technical problems of multi-homed IP providers meshed with other upper- and lower-tier ASs. Security schemas encrypting application-level protocols (i.e., FTP and TELNET) do not work with NAT. Kerberos authentication also has problems relative to network address translation [36]. In addition, potential nontechnical business and political issues associated with organizational identification with IP address numbers may surface in the future. These tangential issues will not be discussed in this article.

Provider-Based Routing Discussion

"This memo provides information for the Internet community. ... the basic components of the plan ... management of Internet address space and ... aggregation of routing information." — RFC 1481, July 1993 [23]

Provider-based IP address space management with CIDR is a quasi-rigid routing hierarchy. If we compare this routing paradigm with the old Bell System routing hierarchy (see sidebar), a provider-based IP address space hierarchy is as inflexible as the antiquated Bell System of five hierarchical classes.

Historical Call Routing Hierarchy

One of the best places to find an evolving and scalable geographical hierarchical routing model is in global telephony. Briley [40] describes the Bell System hierarchy of the 1980s, and both the domestic U.S. numbering plan and the world numbering zones. Similar to the hypothetical meshed model presented in Fig. 1, the International Consultative Committee for Telephone and Telegraph (CCITT) developed a world numbering plan for international call routing. The world was divided into nine zones, and each country was assigned a code, up to three digits, beginning with the digit corresponding to their geographic location in the world numbering zone.

In the United States prior to 1984, under the old Bell System, domestic call routing was based on a five-level switch hierarchy:

- Class 5 end office (EO)
- Class 4 toll center (TC)
- Class 3 primary center (PC)
- Class 2 sectional center (SC)
- Class 1 regional center (RC)

The end office served customers directly, and the toll center was the first stage of toll aggregation. Toll centers routed calls to primary centers, then on to sectional centers. Regional centers were the top of the call routing tier. There were restrictions imposed on the call hierarchy, known as the one-level limit rule [40]: "The switching function performed for the first routed traffic by the switching system at either end of a high-usage trunk group may be of the same class number of switching function, or differ by at most one."

Switching in the pre-1984 five-level Bell System utilized a trial-and-error algorithm, trying successive trunks and then switching to another trunk if a busy group was encountered. The old five-level Bell System hierarchy was converted to a simple two-level hierarchy after the introduction of dynamic nonhierarchical routing (DNHR) in the 1980s. DNHR algorithms dynamically changed the routing strategy to be a function of time and loading [40], similar in concept to IP routing protocol metrics in use today. Class 5 central offices in the pre-1984 hierarchy become equal peers, referred to as local exchange carriers (LECs) [41]. LECs are analogous to local Internet service providers. Similar parallels may be made between interexchange carriers (IXCs) at points of presence (POPs) and Internet transit providers peering at network access points (NAPs).

In practical terms, business and political agendas were the constraining factors regarding service provider mobility; the addressing-numbering schema was not a limiting factor.

Earlier, with a few simple illustrations, we demonstrated how large IP service providers must readdress their networks in the event they choose another IP service provider. In addition, the closer the provider is to the top of the CIDR routing hierarchy and the larger their downstream IP customer base, the greater the effects of renumbering. This configuration change affects downstream providers and thousands of networks. In large internetworks, networks which may span many corporate, geographic, and political domains, renumbering ASs or translating network addresses may be technically feasible. However, renumbering may be impractical and cost-prohibitive, and these constraints bring about an anti-competitive environment where large established upper-tier IP providers have a significant advantage in the marketplace.

Consumers of telecommunications services benefit tremendously when all users, including established and mid-tier IP service providers, can change their service providers quickly and effectively. In this environment, the competitive quality of telecommunications services is causal to a more economic and robust telecommunications marketplace. One of the key attributes of a competitive telecommunications market, relative to the business climate, is an environment having no restrictions on entry to, or exit from, the marketplace, in which new suppliers can readily enter the marketplace [37]. The Internet and global IP internetworking should be no exception.

Concluding Remarks

The early development track for an IP exterior routing protocol concentrated on policy-based routing in a spanning-tree topology and was not a scalable meshed routing paradigm. Many solutions were proposed, and many are cited in the references, which would have been a dramatic departure from the spanning-tree dependency and unscalable paradigm of BGP development. This development track is causing growing dependence on a quasi-rigid routing hierarchy in the Internet based on provider-based address allocation.

Ironically, many of the important Internet RFCs cited here actually state that provider-based address space allocation with CIDR is not a strict routing hierarchy and does not affect the IP routing architecture. The reader is referred to the referenced IETF documents for completeness.

One fact requires emphasis: The implementation of BGP-4 and application of CIDR to provider-based address space management was a major paradigm shift for the Internet, one that dominates routing and routing policies in the Internet today.

Global IP service providers must have the same degree of freedom to enter, exit, reorganize their networks, and change service providers without limitations created by restrictive IP routing architectures and addressing schemes. This freedom benefits the IP consumer by promoting network growth, value-added services, and lower costs. The downstream provider-based IP addressing paradigm is causal to an anti-competitive telecommunications environment.

Recently, the organization responsible for IP

registry has proposed charging significant "registration fees" for IP address space [38]. This proposal was a natural and predictable iteration of the problematic BGP development track, which has not yet been abandoned by the IETF, large IP service providers, and router vendors. If the economic model of requiring registration fees for IPv4 address space is applied to IPv6, IP address space registration will become a billion dollar global industry.

Global IP routing can and must be accomplished without creating a tangible commodity from the IP numbering system. Excellent technical proposals have been presented, many cited in the references. New proposals in the IETF include a new scalable routing architecture for the Internet, NIMROD [39].

NIMROD has the difficult design goal of representing routing information at numerous levels of abstractions via distributed databases scalable to very large dynamic internetworks. In addition, NIMROD is being designed to support QoS and policy routing requirements. As Comer elegantly cautioned in 1991, changing an established routing hierarchy is very difficult. Implementing NIMROD or a similar scalable IP routing paradigm remains an extremely challenging global issue.

Acknowledgments

This article would not have been possible without a pointer to Kleinrock and Kamoun's paper by Noel Chiappa during lively discussions in the CIDRD WG. A special thanks to Jon Crowcroft for reviewing this article and suggesting a review of telephony call routing. Numerous reviewers have made excellent suggestions which have been extremely helpful and very much appreciated.

References

- [1] C. Topolcic, "Schedule for IP Address Space Management," Internet RFC 1367, Oct. 1992.
- [2] L. Kleinrock and F. Kamoun, "Hierarchical Routing for Large Networks," *Comp. Networks*, North-Holland, vol. 1, no. 3, 1977, pp. 155-74.
- [3] E. Rosen, "Exterior Gateway Protocol," Internet RFC 827, Oct. 1982.
- [4] H. Braun, "The NSFNET Routing Architecture," Internet RFC 1093, Feb. 1989.
- [5] D. Estrin, "Policy Requirements for Inter Administrative Domain Routing," Internet RFC 1125, Nov. 1989.
- [6] Y. Rekhter, "EGP and Policy Based Routing in the New NSFNET Backbone," Internet RFC 1092, Feb. 1989.
- [7] P. Gross, and P. Almquist, "IESG Deliberations on Routing and Addressing," Internet RFC 1380, Nov. 1992.
- [8] E. Gerich, "Guidelines for Management of IP Address Space," Internet RFC 1366, Oct. 1992.
- [9] C. Huitema, "An Experiment in DNS Based IP Routing," Internet RFC 1383, Dec. 1992.
- [10] V. Lougheed and Y. Rekhter, "A Border Gateway Protocol (BGP-3)," Internet RFC 1267, Oct. 1991.
- [11] C. Huitema, *Routing in the Internet*, Englewood, NJ: Prentice Hall, 1995.
- [12] V. Lougheed and Y. Rekhter, "A Border Gateway Protocol (BGP)," Internet RFC 1105, June 1989.
- [13] J. Honig *et al.*, "Application of the Border Gateway Protocol in the Internet," Internet RFC 1164, June 1990.
- [14] V. Lougheed and Y. Rekhter, "A Border Gateway Protocol (BGP)," Internet RFC 1163, June 1990.
- [15] P. Gross, and Y. Rekhter, Eds., "Application of the Border Gateway Protocol in the Internet," Internet RFC 1268, Oct. 1991.
- [16] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," Internet RFC 1654, July 1994.
- [17] B. Kumar, "Models, Implementations and Design Options for Inter-Domain Policy Routing Protocols," *IEEE CompEuro Proc.*, 1992, pp. 481-86.
- [18] R. Stevens, *TCP/IP Illustrated*, vol. 1, Reading, MA: Addison-Wesley, 1994.
- [19] M. Steenstrup, "An Architecture for Inter-Domain Policy Routing," Internet RFC 1478, June 1993.
- [20] Y. Rekhter, "Routing in Communications Network," *Inter-Domain Routing: EGP, BGP, and IDRP*, M. Steenstrup, Ed. Englewood Cliffs, NJ: Prentice Hall, 1995.

- [21] V. Fuller *et al.*, "Supernetting: An Address Assignment and Aggregation Strategy," Internet RFC 1338, June 1992.
- [22] F. Kamoun and L. Kleinrock, *Stochastic Performance Evaluations of Hierarchical Routing for Large Networks*, *Computer Networks*, vol. 3, North-Holland, 1979, pp. 337-53.
- [23] C. Huitema, "IAB Recommendation of an Intermediate Strategy to Address the Issue of Scaling," Internet RFC 1481, July 1993.
- [24] V. Fuller *et al.*, "Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy," Internet RFC 1519, Sept. 1993.
- [25] Y. Rekhter and T. Li, "An Architecture for IP Address Allocation with CIDR," Internet RFC 1518, Sept. 1993.
- [26] P. Ford, Y. Rekhter, and H. Braun, "Improving the Routing and Addressing of IP," *IEEE Network*, May 1993, pp. 10-15.
- [27] B. Carpenter and Y. Rekhter, "Renumbering Needs Work," Internet RFC 1900, IAB, Feb. 1996.
- [28] Y. Rekhter, "CIDR and Classful Routing," Internet RFC 1817, Aug. 1995.
- [29] Y. Rekhter and T. Li, "Implications of Various Address Allocation Policies for Internet Routing," Internet RFC 2008, Oct. 1996.
- [30] Y. Rekhter, "Routing in a Multi-Provider Internet," Internet RFC 1787, Apr. 1995.
- [31] K. Hubbard *et al.*, "Internet Registry IP Allocation Guidelines," Internet RFC 2050, Nov. 1996.
- [32] D. Comer, *Internetworking with TCP/IP*, vol. I, Englewood, NJ: Prentice Hall, 1991.
- [33] A. Tanenbaum, *Computer Networks*, 3rd ed., Englewood, NJ: Prentice Hall, 1996.
- [34] R. Frieden, *International Telecommunications Handbook*, Boston, MA: Artech House, 1995.
- [35] P. Francis and K. Egevang, "The IP Network Address Translator (NAT)," Internet RFC 1631, May 1994.
- [36] P. Francis, "Minutes of the Network Address Translation BOF," Interim meeting report, Bellcore, July 1993.
- [37] C. Kennedy, *An Introduction to U.S. Telecommunications Law*, Boston, MA: Artech House, 1994.
- [38] InterNIC, "American Registry for Internet Numbers Proposal," <http://rs.internic.net/ar/n/>, Network Solutions, 1996.
- [39] I. Castineyra, N. Chiappa, and M. Steenstrup, "The Nimrod Routing Architecture," Internet RFC 1992, Aug. 1996.
- [40] B. Briley, *Introduction to Telephone Switching*, Bell Tel. Labs., Reading, MA: Addison-Wesley, 1983.
- [41] B. Bates, and D. Gregory, *Voice & Data Communications Handbook*, New York: McGraw-Hill, 1996.

Additional Reading

- [1] Z. Avramovic, "Policy Based Routing in the Defense Information System Network," *Proc. IEEE MILCOM*, vol. 3, 1992, pp. 1210-14.
- [2] E. Chen and T. Bates, "An Application of the BGP Community Attribute in Multi-Home Routing," Internet RFC 1998, Aug. 1996.
- [3] D. Estrin and K. Obraczka, "Connectivity Database Overhead for Inter-Domain Policy Routing," *Proc. IEEE INFOCOM*, vol. 1, 1991, pp. 265-78.
- [4] E. Estrin, M. Steenstrup, and G. Tsudik, "A Protocol for Route Establishment and Packet Forwarding Across Multidomain Internet," *IEEE/ACM Trans. Networking*, vol. 1, no. 1, Feb. 1993.
- [5] J. Hawkinson and T. Bates, "Guidelines for Creation, Selection, and Registration of an Autonomous System (AS)," Internet RFC 1930, Mar. 1996.
- [6] A. Noll, *Introduction to Telephones and Telephone Systems*, Norwood, MA: Artech House, 1991.
- [7] W. Nurmi, "Telephoning: Making the Connection Between Switching and Routing," *IEEE Potentials*, Oct./Nov. 1995, pp. 25-28.
- [8] R. Perlman, "Hierarchical Network and the Subnetwork Partition Problem," *Comp. Networks and ISDN Sys.*, vol. 9, June 1985, pp. 287-303.
- [9] P. Tsuchiya, *The Landmark Hierarchy: A New Hierarchy for Routing in Very Large Networks*, ACM Publication, MITRE Corp., 1988.
- [10] Z. Wang and J. Crowcroft, "A Two-Tier Address Structure for the Internet: A Solution to the Problem of Address Space Exhaustion," Internet RFC 1335, May 1992.

Biography

TIM BASS (bass@silkroad.com) is a technical director with SAIC, Center for Information Protection. He graduated from Tulane University in 1987 with a B.S.E. with Departmental Honors in electrical engineering. He played a principal role in building the SprintLink Integrated Network Management Center and the Sprint Managed Router Network organization in 1993. Recently, he completed the design and implementation of the original Base Network Control Center (NCC) prototype for the United States Air Force and designed and built over 20 network control centers for USAF bases worldwide. His current interests are TCP/IP performance analysis, high-speed internetworking, wireless networks, and network security issues.