Internet Routing: Separating Customers from Providers

Beichuan Zhang
bzhang@cs.arizona.edu

Vamsi Kambhampati
vamsi@cs.colostate.edu

Daniel Massey
massey@cs.colostate.edu

Ricardo Oliveira
rveloso@cs.ucla.edu

Dan Pei
peidan@research.att.com

Lan Wang
lanwang@memphis.edu

Lixia Zhang
lixia@cs.ucla.edu

September 2006

ABSTRACT

To address the serious challenges in scalability, stability and security facing the global routing infrastructure today, we propose a new routing architecture SIRA[1], which separates Internet providers and customers to different address and routing spaces. This separation insulates the core routing infrastructure from the growth, dynamics and security threats generated by edge customers, promotes site multi-homing by eliminating provider-dependent addresses, and provides multi-homed customers an explicit channel to express their preferences among their providers. We also identify the major challenges in engineering the proposed seaparation, namely how to design a robust and secure mapping service, and how to bridge the two separate routing spaces.

1 A FUNDAMENTAL PROBLEM: GLOBAL ROUTING SCALABILITY

The Internet has been a great success, however the developing business climate in the Internet is testing the limits of the current network architecture. In particular, the pervasive multihoming practice has made a profound impact on the scalability of the current routing and address architecture. A multihomed customer can be reachable through whichever provider that remains functioning after network failures; in the absence of failures, the customer can use its multiple provider connectivity to maximize its locally defined goals such as performance, throughput, or cost. Therefore, business users buy Internet service from multiple providers for improved Internet availability[2]. However, being reachable through any of its providers implies that the customer must be visible in the global routing table. The customer may even split its address prefix into multiple longer ones to do load balancing on incoming traffic.

[1]SIRA: Scalable Internet Routing Architecture.

[2]The report "Impact of the 2003 Blackouts on Internet Communications" by Renesys shows that this is an effective approach [5].

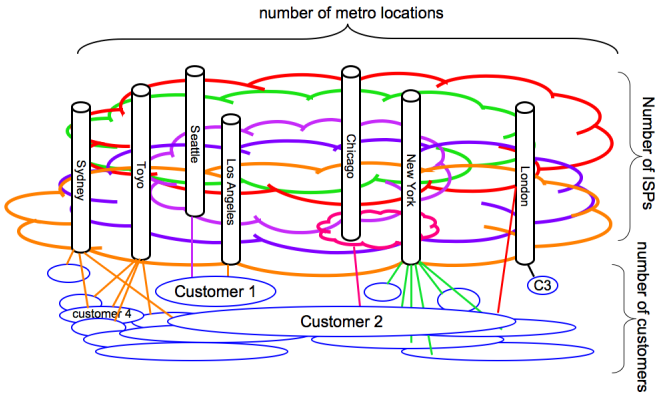

Figure 1: This figure depicts the major factors driving the global routing table size: the number of ISPs, the number of metro locations each ISP has a presence, and most importantly the number of multihomed customer networks which is not only orders of magnitude larger than the other two factors but also increasing rapidly.

We attempt to capture the scaling challenge in Figure 1: as an example, it shows that all global ISPs, and some regional ISPs, have a presence in New York; a large number of customers connect to the Green ISP at New York; similarly (not shown in the figure) a large number of customers also connect to other ISPs in the same location. An entry in the global routing table is required to be able to reach each New York customer through any of the ISPs it connects. Ideally the global routing table size should be proportional to the number of ISPs, in addition it may also be a function of the number of major metro locations where ISPs interconnect. With multihoming and traffic engineering, however, the current routing table size is driven by the product of the number of ISPs, the number of metro locations, and the number of customers connected at each {ISP, metro}.

Furthermore, Internet users in general desire provider-independent (PI) prefixes in order to avoid renumbering when changing providers. The resistance to renumbering is not only due to the burden of renumbering all the

routers and hosts on a customer site, but more importantly due to the fact that the IP addresses have been embedded deeply in many other places: network security policy configurations, network applications, network accounting systems, to name only a few. The use of PI prefixes also naturally fits the multihoming practice, as a customer network can simply ask each of its providers to announce its PI prefix, or fragments of it. These customer desires *directly conflict* with that of the providers who strive to keep the global routing table size moderate and stable through the use of provider-allocated (PA) prefixes that can be aggregated for routing scalability.

Another major factor regarding routing scalability is the amount of update messages routers must process in real time. Because of the *flat* nature of the Internet routing, a routing flap to any destination can trigger routing updates to be propagated through the entire Internet, even when no one communicates with the unstable destination before its connectivity recovers. Both our own measurements and that of others have shown that the overwhelming majority of BGP updates are generated by a very small number of sources, most of them being small edge networks [11, 14].

The above discussion can be summarized into one fundamental problem in today's routing architecture: there exists fundamental conflicts of interests between the Internet service providers (ISPs) and customers regarding the Internet address allocations. ISPs would like to see all customers using topologically aggregatible PA prefixes to make the Internet routing system scalable with ever increasing user population, while Internet users desire PI prefixes to avoid renumbering, and perform multi-home with various degrees. Multihoming essentially destroys topology-based prefix aggregation, no matter from where the users received their prefix allocations.

How did such fundamental conflicts arise? Why didn't the design of the existing Internet architecture foresee such conflicts and avoid them? What can be an effective solution to today's routing scalability problem? In the rest of this paper we attempt to answer the first questions in Section 2 through offering a quick review of the Internet evolution, and then sketch out our proposed solution to the problem. Our solution eliminates the fundamental conflicts in IP address allocations between ISPs and customers by putting them on two separate address spaces, and we present the advantages in doing so in Section 4 and the open challenges in Section 5.

2 How We Got Where We Are

The fundamental goal of the original IP design was to interconnect all packet switched networks that might be based on different communication technologies, so that packets can be delivered from any IP box to any other IP boxes [4]. *Gateways* were invented to interconnect networks of different communication technologies. At the time the basic Internet architecture was sketched out [3], resilience against physical failures was the primary concern, and all the gateways ran in a single routing domain and they were expected to forward packets for all their neighbors. The design noted the different ownerships of networks and gateways, however there was no contractual relationships but cooperation between neighboring networks.

As the Internet expanded rapidly to a large number of institutions in early 80's, it was quickly realized that a flat routing architecture could no longer keep up with the increasing network scale and management complexity. The concept of *Autonomous System (AS)* was developed and routers are grouped into ASes [15]. Each AS is under a single administrative control and runs its own routing protocol internally; a standard inter-domain routing protocol, BGP, runs between ASes. We note three properties of the AS concept. First, the BGP routing is flat at the AS level and all ASes play a part, no matter whether an AS is a global ISP or a small user site. Second, it enables route aggregation of large networks: One can route traffic according to detailed topological granularity inside an AS to optimize performance, and export to neighbor ASes aggregated prefixes to make the global routing more scalable[3]. Third, the AS is a unit in route management and control but is hidden from the routing table. One cannot tell, by looking at individual prefixes, which prefix belongs to which AS. This has been the basic routing architecture till today.

The commercialization of Internet in the late 80's changed the landscape completely. Most importantly, it created a market for global data delivery service. Naturally ASes began differentiation driven by economic forces. The most prominent distinction is the one between *customer networks* and *provider networks*. Customer networks serve end users directly and are consumers of the global data delivery service. Provider networks have the sole purpose of delivering packets *for a charge*. In return, they hold a *contractual obligation* to the customer networks for providing packet delivery service.

As the Internet penetrates into every corner of our society, the economic forces drive every AS to optimize its connectivity for its own purposes. Today's customer networks (e.g., university campuses, global enterprises) have different business models, different growth trends, and fundamentally different requirements than service provider networks (e.g., AT&T). These differences resulted in the failure of route aggregation, hence the explosion of the global routing table size.

We believe that the failure to accommodate the distinc-

[3]However this also implies that one's internal table can be substantially bigger than the external BGP routing table.

tion between customer networks and provider networks in today's routing and addressing architecture is the root cause of the scaling problems facing us today. It is time to evolve the routing architecture once again to keep up with the evolving Internet.

3 The Separation of Customers and Providers

Let us examine the three major routing scale factors in Figure 1 again: to cross the Internet backbone, a packet needs to be forwarded towards the destination ISP *P* and the metro location *L*. The information of destination customer is not needed until the packet has reached {P, L}. However the first two factors, the number of ISPs and the number of metro locations, have been relatively stable, it is the third factor that has been driving the growth of the global routing table. Thus we believe one effective solution is to eliminate the third factor from the global routing system, that is to separate out customer networks from the global routing system.

SIRA distinguishes between *Customer Networks (CNs)*, which act as sources or sinks for data packets, and *Provider Networks (PNs)*, whose primary role is to provide data transit service at the inter-domain level. SIRA puts all the PNs in one address space, and all the CNs in a separate address space. A PN address identifies a connection point to the Internet backbone. A CN address identifies an interface of a customer host. In other words, one may call PN address space the *locator* space, and CN address space the *identifier* space.

In the locator space, provider networks interconnect to form the *Global Transit Network (GTN)*. A provider routing protocol (e.g., such as today's BGP) runs between GTN routers to maintain reachability among all the PNs. Individual customer networks, on the other hand, operate in the identifier space and use a customer routing protocol (e.g., such as today's OSPF) that maintains routes to reach internal subnets and its immediate neighbors (its providers or other directly connected customer networks). The two different address and routing spaces are interconnected by those links that connect customer networks to their providers. *No* routing protocol runs across the links between PN and CN routers. How to handle the failures of these links will be discussed in Section 5.

End-to-End Data Delivery In SIRA, end-to-end data delivery across GTN is achieved by encapsulating customer packets in a GTN packet header, with the source address as the GTN entry router and the destination address as the GTN exit router. A mapping service is needed to map a destination customer address to the corresponding exit router(s) address in the provider space. Figure 2 illustrates a typical packet forwarding example. When a source host *Src* (in customer network *S*) sends a

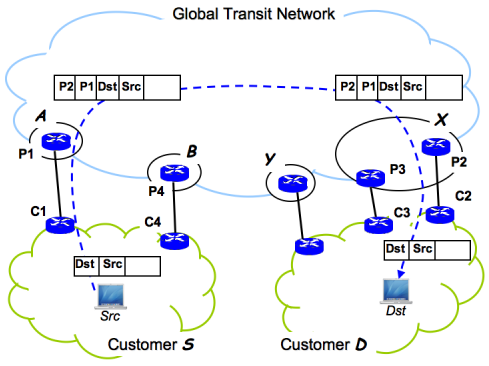

Figure 2: A Sample SIRA Session from Src to Dst

packet to a destination host *Dst* (in customer network *D*), the packet will be forwarded to one of *S*'s providers, say *A*. The mapping service maps the address *Dst* to GTN exit router (P_2) that connects to *D*. The ingress GTN border router P_1 then encapsulates the packet with its own address as the source and P_2's address as the destination, and forwards it to P_2. Upon receiving the packet, P_2 will decapsulate it and send it to *D*. Viewed from customer networks, the GTN is a single logical hop connecting all customer networks.

On the surface the encapsulation step in crossing GTN may bear a resemblance to NAT (Network Address Translation); however, in reality SIRA differs from NAT in fundamental ways. SIRA assigns unique and provider-independent addresses to all hosts in customer space, thus they can be reached in the face of individual provider failures. In SIRA, any customer host can *directly* talk to any other customer host by simply putting the destination address in the packet. Therefore SIRA helps reinstall the end-to-end transparency model in the Internet.

4 Benefits from the Separation

The major benefits of separate CN and PN address spaces, or identifier and locator spaces, can be summarized below.

Routing Scalability and Stability In SIRA, GTN routing is only concerned with reachability among PNs. Not only the number of PNs should be much smaller than that of CNs, but more importantly its growth trend is expected to be much slower. Our measurement shows that the number of transit ASes is only about 20% of the total ASes in today's Internet, and the number of transit ASes grows at 1/5 of the rate of all the ASes. Each PN may announce multiple prefixes depending on its size and traffic engineering practice; however, this number is not related to the number of CNs it supports, or how its CNs may be multihomed or do load balancing. These prefixes are *topologically aggregatable*.

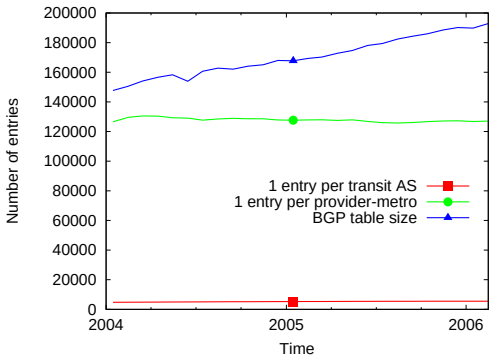

Figure 3: Routing Table Size Growth

Figure 3 shows estimated SIRA routing table size as compared to the current BGP table size, which is represented by the top blue curve in the graph. From Jan. 2004 to Feb. 2006, BGP table sizes increased from below 150,000 to above 190,000. The middle green curve is an estimated upper bound on the SIRA routing table, assuming that each tier-1 AS represents a PN which announces one prefix per metropolitan area in the world, each tier-2 or tier-3 AS represents a PN that announces one prefix per metro in its own country, and each tier-4 PN announces 2 prefixes[4]. The bottom red curve is the number of transit ASes of all the tiers, representing an estimated number of PNs. We believe that GTN's routing table size should be somewhere between the green and red curves and is likely to remain more or less constant over time.

Because of the separation of CNs from GTN, routing dynamics occurring inside CNs or at the border (between CNs and GTN) will have no impact on the routing stability inside GTN. Also, since the size of GTN is expected to be much smaller than the number of ASes today, routing convergence would be substantially faster than that of today's BGP.

Site Multihoming and Traffic Engineering Because CNs are on a address space separate from that of PNs, naturally the entire address space is provider-independent. Customers can change providers freely without renumbering their networks, and can subscribe to as many providers as they want with no negative impact on the global routing table. As a result, SIRA removes roadblocks for customers to adopt multi-homing, which improves the reliability of their Internet connectivity.

In addition to enhanced network reliability, customers may also want to fully utilize the parallel connectivities provided by multihoming. Since the address space separation between CNs and PNs introduces the need for a mapping service, we can utilize this mapping service for effective traffic engineering support. In addition to the basic goal of mapping a customer address to that of its providers, customers can inject into the mapping service additional *policy* information to facilitate the selection of provider address among multiple alternatives. For example, customers can specify preferences about which provider to use for incoming traffic. In the example shown in Figure 2, *Dst* may want to split its incoming traffic between its two providers X and Y. For instance, it can specify in its mapping entry a preference of receiving 60% traffic via provider X and 40% traffic via provider Y. The sender *Src* learns *Dst*'s preference through the mapping service. It can now make an informed decision based on both the receiver's and its own connectivity and preferences, taking full advantage of multihoming. In the current Internet, on the other hand, there is no effective way for a receiver to influence the incoming traffic paths except announcing longer prefixes and prefix splitting, which is one of the main causes of today's routing scalability problems.

Security Enhancement Because SIRA puts all customer hosts on an address space separate from that of backbone routers, all user data packets are encapsulated when they cross the backbone. As a result, compromised hosts in the customer space no longer have direct access to the provider infrastructure. Attackers can still use compromised hosts within a customer network to DDoS the local GTN *border* routers, however such attacks only make a local impact and are relatively easy to deal with. Attackers may also use compromised hosts from multiple customer sites (e.g., a botnet) to DDoS the routing infrastructure by flooding packets to some remote customer destinations. However, given the GTN topology is opaque to end users, attempting to DDoS any specific component in the provider topology becomes more difficult. Although SIRA does not eliminate any specific security threat, it raises the barrier against malicious attacks targeted at the global routing infrastructure.

The encapsulation of customer packets also makes it easy to trace attack packets back to the GTN ingress router even if they have spoofed source addresses, since the encapsulation header records the addresses of the GTN entry and exit routers. In today's Internet, some providers follow the recommended practice and configure border routers to check the source address of packets coming from their customers. However, not all providers implement such ingress filtering, as they do not perceive a direct benefit for the deployment cost.

Although SIRA makes it difficult to gain unauthorized access to GTN routers, we do not expect GTN to be

[4]We classified the ASes into different tiers using the method described in [16]. We define a metropolitan area as a city with a population of 250,000 or greater and obtained the population data from [2].

free of malice. Routers in GTN may still get compromised, and providers in GTN may belong to unknown parties of different interests. Thus, it is necessary for SIRA to develop effective mechanisms to detect compromised routers and misbehaving PNs *within GTN*. However, we believe that compromised user hosts are a major source of malicious attacks, and SIRA raises the barrier against such attacks. Furthermore, the locator space is expected to be substantially smaller compared to today's Internet, making it much easier to detect attacks and diagnose problems.

Incremental Deployment In SIRA, customer packets traverse GTN via encapsulation. Since packet encapsulation can be performed between any pair of IP boxes, it allows ISPs to move to the new locator space one by one. Given customer renumbering is difficult if not impossible, let us assume that the current IP address space will eventually become the identifier space in SIRA. When the first ISP moves to the new locator space, it can encapsulate all packets traversing through itself at corresponding entry and exit routers. When a second ISP joins the locator space, if it is connected to the first one, the two create a partial GTN.

5 CHALLENGES

With all the benefits from separating CN and PN address spaces, the separation also raises a few challenges in the overall system design and deployment. The essential ones include how to design the mapping service between customer and provider address spaces, how to handle link failures between GTN and CNs, how to measure GTN properties by customers, and how to draw the line between CNs and PNs. This section discusses these open issues and the pros and cons of different design approaches.

5.1 The Mapping Service

The basic functionality of the mapping service is, given a destination customer address, it should return a destination provider address so that the packet can be encapsulated and forwarded across the GTN. The mapping service can also be augmented to include traffic engineering preference information of the receiving network. The mapping can be done within the source CN or at the GTN entry router. We classify various ways of implementing this mapping service into the following two types.

One approach is to disseminate the mapping information to every CN or every GTN entry router. The dissemination can be done in multiple ways. One possibility is to attach the information to the GTN routing updates, however doing so would bring back the scalability problem associated with the number of CN attachment points. Another possibility is to run a separate dissemination protocol among GTN routers to propagate mapping information, however it suffers from the same scalability problem. Yet another possibility is to build an overlay network to broadcast or multicast the information. The common advantage of these schemes is that lookup can be done locally at source CNs or GTN entry routers, therefore the mapping does not incur significant delay in packet forwarding. The disadvantage is that any change of customer-provider mapping must be *proactively* propagated *globally*, even if the change may not affect any data traffic. Given the number of customer networks grows at a rapid rate, the dissemination system itself faces a scalability challenge.

The second approach is to provide the mapping service by distributed servers in a way similar to DNS system, and let hosts in CNs, or GTN entry routers, query the servers for the information needed to forward each packets. The advantages of this approach are that changes are made to local servers, rather than being proactively propagated globally, and that individual responsible parties can selectively enhance *their own* mapping service through faster servers or more replications. The disadvantage is that the query will add extra delay to packet forwarding. Caching and prefetching popular mapping entries may provide effective performance improvement.

It should be emphasized that the mapping service is a necessary cost for the gains from the separated provider (locator) and customer (identifier) spaces. Up to now, at least in theory, packet delivery relies only on the intra- and inter-domain routing to work correctly. SIRA introduces a new dependency, the mapping service, which represents a system cost in providing the mapping servers, a performance cost in query delays, and a target for potential attacks. At the same time, we would also like to point out that the introduction of DNS 20 years ago could be considered remotely analogous: DNS introduced a new dependency in data delivery (DNS name to IP address translation), the cost in providing DNS servers and lookup delays, and a target for potential attacks, which have occurred frequently in recent years. Yet the gain from DNS is so essential that all the cost is considered necessary tradeoffs. We believe the same arguments can be made for the new mapping service, which can also potentially leverage on the existing DNS infrastructure to minimize additional costs for the new service.

Securing the mapping service is essential for SIRA to succeed. Two types of attacks are of particular concern: denial of service attacks and response modification attacks. By disabling access to the mapping service for a given customer network, one can deny packet delivery to that network. To make the mapping service resilient to DoS attacks, data can be widely replicated and

cached so that knocking down one or a few servers does not disrupt packet delivery for a customer, as DNS has demonstrated. The modifications to the mapping replies is also a shared problem with DNS, which may require cryptographic authentication protection. Non-crypto solutions include (1) querying the information from multiple servers for mutual checking, assuming man-in-middle attackers cannot hijack all the queries or replies, and (2) periodic queries to monitor the mapping service in an effort to detect any false data.

5.2 Handling Border Link Failures

SIRA creates separate customer and provider routing spaces where any topological change in a customer network is handled by the local customer routing protocol, and any change in GTN is handled by the GTN routing protocols. However, a link between a customer and its provider, for example the link between provider router P_2 and customer router C_2 in Figure 2, is not part of either the provider or the customer routing domain. Thus when this link, or the router at the CN side, fails, no routing update would be generated in GTN. This can be viewed as an advantage as it provides the isolation of edge dynamics from GTN; this also introduces a challenging problem in assuring packet delivery.

Consider the problem of sending from *Src* to *Dst* in Figure 2. Assume the packet is forwarded from *Src* to GTN ingress router P_1 in provider *A*. When the mapping service shows that P_2 is the egress router for the destination network, P_1 forwards the packet to P_2. Assume that the link between P_2 and C_2 fails. By the SIRA design, provider *A* is not informed of the failure. However when P_2 receives the packet, it cannot forward the packet onto the destination. At this point P_2 may look for alternate route to *Dst* (in this case P_2 could re-encapsulate the packet and forward to P_3) or drop the packet and send an ICMP "Destination Network Unreachable" message to P_1. In the first option, the only remaining routes to the destination may be via some alternate provider (such as *Y* in the figure) and providers may not be willing to perform the added work of finding and using the alternate paths. In the second option, the sender will not learn of the failure until the ICMP message is received. It is also possible to have a combination of both approaches: by default P_2 sends a notification message if the link to the destination CN has failed; however for a premium price, P_2 may also forward packets along alternate routes as a value-added service.

5.3 Measuring GTN Properties by Customers

The separation makes the GTN network appear as a single hop to customers. A customer may choose which provider to use when entering the GTN and perhaps which provider to use when exiting the GTN, but the internal workings of the GTN are not exposed to the customer. This is distinct from current Internet where any host may attempt to use tools such as `traceroute` to gather information on the path between the source and destination hosts. One challenge in SIRA is how to expose information that is legitimately needed by the customers without violating the separation property.

There are two general approaches to providing customers with information about the GTN. In the first approach, although customers cannot directly address, or send packets to, routers in GTN, GTN routers can still send packets to customer hosts. For example, to support a `traceroute`-like tool, a customer host H1 sends special "Trace Request" packet to a destination customer host H2, and each router along the path can generate a path report message to send to H1. As with the current Internet, individual PN routers may or may not choose to process such requests. A second, and perhaps more practical, approach is for GTN to provide a monitoring service where any user can obtain data about the general status of the GTN. In the current Internet, one can consult monitoring projects such as Oregon RouteViews [1] to view the BGP routes from multiple locations towards a specific prefix, or Internet Health Report provided by Keynote which displays the packet delivery delay among ISPs [12]. In our model the GTN can provide similar monitoring services and customers can use these services to infer specific GTN properties.

5.4 How to Draw the Line Between Customers and Providers

Last but not the least question SIRA must answer is how to draw the line between CNs and PNs. As a rule of thumb, if a network appears in the middle of some AS-PATH in today's BGP routing, at least part of the network will become a PN in SIRA; otherwise it should be a CN. For example, a university network will be a CN; it may be connected to multiple providers but it only provide data delivery service for users of the university. It is also possible that two large CNs may set up a connection in between to swap traffic directly. There may exist gray areas, where a network *N* may directly serve end users as well as provide transit services for some other users. If none of *N*'s users have connections to other networks, *N* can be lumped together with all its users and viewed as a CN. Otherwise it must split its transit routers to a PN.

6 Related Work

A number of research efforts have been devoted to Internet scalability and security problems. HLP [17] compartmentalizes the lower tiers in Internet's topological hierarchy into *separate* regions, thereby improving the scalability and stability in today's inter-domain routing sys-

tem. HLP shares a common goal with SIRA, in that it isolates edge instability from the backbone core. However, it does not address the scaling issue caused by today's pervasive multihoming practice.

Handley *et al.* [9] propose to put client and server hosts in different address spaces to prevent DoS attacks. In their design, clients are only allowed to initiate connections to servers, and servers can only connect directly with clients. In a follow-up work, Greenhalgh *et al.* [8] tackle the same problem by creating points of encapsulation and filtering in special subnets called *server-nets*. They use IP-in-IP encapsulation to enable tracing of attack flows, and move the packet filtering upstream, close to the attack sources. SIRA differs from [8] in its systematic address space separation between customers and providers, which enhances both routing scalability and Internet backbone security.

Greenberg *et al.* [7] propose a clean slate design for Internet routing based on four fundamental layers: decision, dissemination, discovery and data. Their architecture decouples the control plane from the data plane for improved network security and manageability. However, all the components are in the same address space and the design does not directly address the routing scalability problem.

In [13], O'Dell proposed a new routing design for IPv6, commonly known as GSE, which had the greatest influence on SIRA. The basic idea in [13] is to divide IPv6's 16-byte address into two parts, with the lower $16 - N$ bytes being the End System Designator (ESD), and the higher N bytes (called Route Goop, or RG) being used for inter-AS routing. The GSE design hides a customer site's RG from internal hosts, which is filled in when packets exit the customer site. This late binding can offer several benefits similar to those provided by SIRA's separation between customers and providers. Unfortunately, the proposal was not adopted and the design was never finished. Prior to GSE, Hinden and Deering also proposed the use of separate address spaces for providers and customers [10, 6], and the use of tunnels over the provider space to carry packets from source customer networks to destination customer networks. These proposals stopped as sketched ideas only; however, they inspired our investigation which led to the SIRA design.

7 Conclusion

This paper presented a key concept in the SIRA design, putting customers in a separate address space from that of providers. We also identified major issues raised by this design and articulated different approaches in the solution space. We believe that SIRA can provide significant advantages in routing scalability and raise the bar for security threats. Our ongoing work includes analyzing engineering trade-offs in a complete realization, and designing a new address structure for the locator space by incorporating location information into provider-based addresses.

References

[1] Advanced Network Technology Center and University of Oregon. The RouteViews project. `http://www.routeviews.org/`.

[2] T. Brinkhoff. City population statistics [cited 2006-01-29]. http://www.citypopulation.de.

[3] V. Cerf and R. Kahn. A Protocol for Packet Network Intercommunication. *IEEE Trans. Comm.*, 22(5):637–48, May 1974.

[4] D. Clark. The Deisgn Philosophy of the DARPA Internet Protocols. In *ACM SIGCOMM*, pages 106–114, 1988.

[5] J. H. Cowie, A. T. Ogielski, B. Premore, E. A. Smith, and T. Underwood. Impact of the 2003 Blackouts on Internet Communications. Technical report, Renesys Corporation, 2003.

[6] S. Deering. The Map & Encap Scheme for Scalable IPv4 Routing with Portable Site Prefixes. Presentation, Xerox PARC, March 1996.

[7] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang. A Clean Slate 4D Approach to Network Control and Management. *SIGCOMM CCR*, 35(5):41–54, 2005.

[8] A. Greenhalgh, M. Handley, and F. Huici. Using Routing and Tunneling to Combat DoS Attacks. In *USENIX SRUTI'05*, 2005.

[9] M. Handley and A. Greenhalgh. Steps towards a DoS-resistant Internet architecture. In *ACM FDNA '04*, pages 49–56, 2004.

[10] R. Hinden. New Scheme for Internet Routing and Addressing (ENCAPS) for IPNG. *RFC 1955*, 1996.

[11] G. Huston. 2005 – A BGP Year in Review. APNIC 21, March 2006.

[12] KEYNOTE. Internet health report. `http://www.internetpulse.net/`.

[13] M. O'Dell. GSE - An Alternate Addressing Architecture for IPv6. February 1997.

[14] R. Oliveira, R. Izhak-Ratzin, B. Zhang, and L. Zhang. Measurement of Highly Active Prefixes in BGP. In *IEEE GLOBECOM*, November 2005.

[15] E. C. Rosen. Exterior Gateway Protocol (EGP). RFC 827, 1982.

[16] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Kat. Characterizing the Internet Hierarchy from Multiple Vantage Points. In *IEEE INFOCOM*, 2002.

[17] L. Subramanian, M. Caesar, C. T. Ee, M. Handley, Z. M. Mao, S. Shenker, and I. Stoica. HLP: A Next Generation Inter–domain Routing Protocol. In *ACM SIGCOMM*, 2005.