

A Secure and Scalable Internet Routing Architecture (SIRA)

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

1. INTRODUCTION

The Internet routing architecture faces many challenges, ranging from scaling problems, security threats, inadequate support for traffic engineering and customer multihoming, to poor fault diagnosis capabilities. A brief description of some of the problems facing the Internet follows.

Security: Internet core routers, which are essential for the routing infrastructure, are being targeted by malicious attacks such as Distributed Denial-of-Service (DDoS) attacks and software bug exploits. In addition, a malicious node can easily spoof the address of another node because Internet addresses do not convey information about *who owns that address* or *where it should originate from*. At the routing level, such prefix hijacking, commonly referred as *false origin* attack, is becoming an increasingly serious threat.

Scalability: Analysis of global routing tables from the past few years indicate that the number of customer networks (i.e., edge networks) are growing at an alarming rate compared to transit provider networks [3]. This growth in routing tables leads to expensive, complex hardware and the need for frequent replacement of core routers. In addition, due to the flat nature of Internet routing at the Autonomous System (AS) level, unstable customer AS's can introduce a lot of unnecessary updates into the global routing system, leading to delayed route convergence and even cascading router failures.

Traffic Engineering (aka Routing Policies): Internet Service Providers (ISP's) use routing policies to control the flow of traffic in order to maximize resource utilization and minimize cost. However, due to the lack of information about prefix ownership, implementing these policies is cumbersome and error prone. Furthermore, limited information about location may lead to inefficient routing decisions in hot-potato routing, as described in [4]. Customer networks, on the other hand, may want to split traffic across each multi-homed link. Unfortunately there is no means of explicitly expressing these types of policies in today's Internet.

Fault Diagnosis: At any time, routers or links may fail.

For prompt recovery, network operators need accurate information about the location of the faults and the routers involved. Given a stream of routing update messages, it is often difficult to determine what events caused the update stream. The result is poor isolation of faults that often end up with explosive number of routing updates and long convergence delays.

We argue that there are two main root causes for the above *symptoms*: mixing customers with providers and lack of important information in the address structure for making efficient routing decisions. Rather than proposing patches to existing Internet architecture, in this work, we propose a secure and scalable Internet routing architecture that fixes these symptoms at the architectural level. We should point out that our work draws heavily from previous work in this area [1, 2, 5] (see our technical report [6] for more discussion of previous work).

2. PROVIDER/CUSTOMER SEPARATION

In SIRA, we distinguish between *network customers*, who act as sources or sinks for data packets; and *network providers*, whose primary role is to provide data transit service for these customers. We observe that customer networks and provider networks grow independently, have different security threats, traffic engineering needs and fault diagnosis requirements. To recognize these fundamental differences, the SIRA design provides a *logical separation between network transit providers and network customers*. First, SIRA places customers and providers in completely *distinct routing spaces* and uses a *mapping service to bridge the two routing spaces*. For instance, provider networks combine to form the *Global Transit Network (GTN)* and maintain reachability to other providers only, whereas customer networks maintain local reachability information. There is *no routing protocol operating across the links between these two spaces*. Second, SIRA assigns separate *address spaces* to providers and customers. Finally, SIRA does *not allow* direct communication across the two spaces.

Figure 1 shows a sample SIRA session. When a source host *Src* (in customer network *S*) needs to send a packet to a destination host *Dst* (in customer network *D*), it first needs to forward the packet to one of *S*'s providers, say *A*. However, due to the separation of routing spaces, *A* does not know how to reach *Dst*. To identify the destination's provider, *Src* contacts a *Customer-to-Provider Binding (CPB)* service (not shown in figure), located in customer space (assume *Src* is bootstrapped with sufficient information about how to reach CPB servers), and retrieves a set of providers for

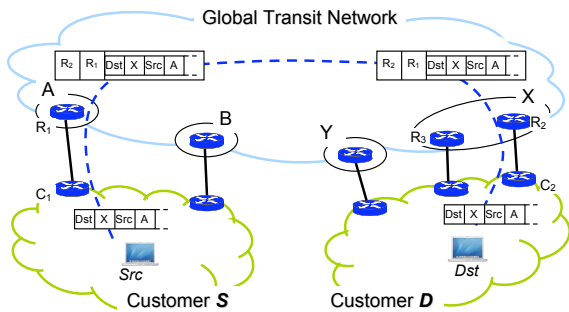


Figure 1: The SIRA Architecture

Dst, i.e., *X* and *Y*. The CPB service also carries information about *D*' preference for each of its providers. Now suppose, *Src* selects *X* (based on the preference set by *D*), it then forwards this information, along with the packet to the GTN ingress router *R*₁, via the border router *C*₁. At this point, *R*₁ needs to know the GTN egress router in *X*, to reach *Dst*. *R*₁ contacts a *Customer-to-Provider Edge Mapping* (CPEM) service (not shown in figure), located in provider network *X*, and retrieves a list of GTN egress routers to reach *D*, i.e., *R*₂ and *R*₃. The CPEM service also carries information about *X*'s preference for each of its egress routers to *D*. Now suppose *R*₁ selects *R*₂ (based on preference set by *X*), *R*₁ then encapsulates the packet in a GTN header with itself as the source and *R*₂ as the destination and forwards the packet to *R*₂. Upon receiving the packet, *R*₂ decapsulates it and sends it to *D* via *C*₂.

3. SIRA ADDRESS STRUCTURE




Figure 2: The SIRA Address Structure

SIRA also introduces a *new address structure* that embeds information of both network organizations and metropolitan location. Figure 2 shows the address structure adopted by SIRA. The four components of the address structure, i.e., the *Organization ID*, *Location ID*, *Subnet*, and *Interface ID* uniquely identify the network attachment point of a device (routers and hosts). The organization ID component identifies the organization (AS) to which the device belongs. The location ID component identifies the physical location of the device, such as continent code, country code, and metro location (e.g. longitude-latitude coordinate of the metro area). Finally, the subnet and interface ID components identify the unique subnet and identifier of the device within a network.

In addition to these four components, a set of *component relationship* bits indicate the relationship among the components (setting a bit to 1 means that the following component is *not* specific to the previous component). For example, a router connecting multiple providers at an exchange point would have *OS* = 1 and *LS* = 0. This indicates that the device is not specific to any organization, but it is located at a specific metropolitan location (i.e. it is not mobile).

4. DESIGN ISSUES

We face many issues in designing the mapping service. First, it would be infeasible for each host to maintain a flat CPB table containing all other hosts' mapping information. We believe that a hierarchical structure similar to DNS would be much more scalable, since each network only needs to maintain the CPB information of its own hosts. Second, the mapping operations cannot add much delay to data delivery. To this end, we use caching and on-demand notification to make the mapping service more efficient. Moreover, we need to consider security threats to both CPB and CPEM.

In addition, since there is no routing protocol covering the links between customers and providers, we need to handle link failures at these links. For example, in Figure 1, the link between *R*₂ and *C*₂ may go down. This event should be somehow conveyed to *Src*. In addition, since *Src* selected *X*, if there is no reachability from *A* to *X*, then this information also need to be conveyed to *Src*. Our technical report [6] provides a more complete discussion of the design issues as well as our proposed solutions.

5. BENEFITS OF SIRA

Some of the key benefits of SIRA include:

Security: SIRA separates provider networks from customer networks. End hosts cannot DDoS provider routers and routers cannot be compromised from end-hosts. Customer networks cannot announce provider prefixes. Also, the address structure permits easy detection of false origin attacks. **Scalability:** Dynamics due to failures in customer space do not trigger updates in provider space and vice-versa. In addition, the two spaces grow independently, therefore the routing tables in provider network do not grow at the same rate as current routing tables. The address structure allows easy aggregation of routes to provider networks farther away.

Traffic Engineering: The mapping service allows customers and providers to explicitly set preferences for incoming traffic. The address structure allows explicit policies based on organization and location information.

Fault Diagnosis: The organization and location information in the address structure identify the exact location of a faulty network device.

6. REFERENCES

- [1] S. Deering. Metro-Based Addressing: A Proposed Addressing Scheme for the IPv6 Internet. Presentation, Xerox PARC, July 1995.
- [2] S. Deering. The Map & Encap Scheme for Scalable IPv4 Routing with Portable Site Prefixes. Presentation, Xerox PARC, March 1996.
- [3] G. Huston. 2005 – A BGP Year in Review. APNIC 21, March 2006.
- [4] R. Mahajan, D. Wetherall, and T. Anderson. Negotiation-Based Routing Between Neighboring ISPs. In *NSDI*, May 2005.
- [5] M. O'Dell. GSE – An Alternate Addressing Architecture for IPv6. draft-ietf-ipngwg-gseaddr-00.txt, February 1997.
- [6] B. Zhang, D. Massey, D. Pei, L. Wang, L. Zhang, R. Oliveira, and V. Kambhampati. A Secure and Scalable Internet Routing Architecture (SIRA). Technical Report TR06-01, University of Arizona, April 2006.