Don't Secure Routing Protocols, Secure Data Delivery

Dan Wendlandt, Ioannis Avramopoulos, David G. Andersen, and Jennifer Rexford

1 Introduction

Internet routing and forwarding are vulnerable to attacks and misconfigurations that compromise secure communications between end-systems. With networks facing external attempts to compromise their routers [3] and insiders able to commandeer infrastructure, subversion of secure Internet communication is an ever more serious threat.

Much prior work has attempted to provide communication security by securing the routing protocols (e.g., S-BGP [10] and so-BGP [12]). We argue that solving the problem of secure routing is both harder and less effective than directly solving the core problems needed to communicate securely: end-to-end confidentiality, integrity, and availability. Secure routing protocols focus on providing *origin authentication* and *path validity*, identified as necessary by the IETF to secure BGP [7]. Unfortunately, these properties are both too little and too much:

Secure routing is too little: As we discuss further in §2, secure routing does not completely address the core problems in secure communication. For example, it cannot prevent adversaries on the communication path from eavesdropping on the data traffic; end-to-end encryption is a more secure solution. Similarly, secure routing cannot detect or prevent packet loss due to data-plane bugs, misconfigurations, or attacks.

Secure routing is too much: The mechanisms behind secure routing, both cryptographic and administrative, are painfully heavy-weight. They require hardware upgrades in the routers for cryptographic processing, time-consuming maintenance of address registries, and a new public key infrastructure (PKI).

Recognizing that a secure version of BGP will be difficult to deploy, yet provide only limited protection, we ask: what is the best division of labor between end-hosts and the routing infrastructure to provide secure, robust communication? The answer, we argue, is that the routing infrastructure must only provide *availability*, defined as the ability for a sender to find a working path to the valid destination as long as such a path exists. Endhosts must provided confidentiality and integrity as needed.

Following this model, we present Availability Centric Routing (ACR), which is based on three principles:

1. Clients learn multiple paths to a destination.
2. Clients use end-to-end integrity checks and monitor path performance to determine if a path is working.
3. Clients can change paths to find one that works.

By propagating multiple paths per destination instead of one "best path," ACR thwarts an adversary's attempt to prevent a source from hearing a valid path to a destination.

Taken together, ACR has several interesting advantages over traditional secure routing schemes:

- Availability threats involving the data plane, such as malicious drops, stray ACLs, link DoS, and transient routing issues, can be detected and avoided.
- Significant gains in resilience are achieved even if only a few interested domains cooperate.
- Adoption is simplified because no address registries, AS-level PKI, or router cryptography is required.
- Performance, usually at odds with security, also benefits from path diversity.

ACR achieves robustness by treating learned routes as possibilities, not certainties. With this approach, control-plane security that eliminates invalid routes (e.g., S-BGP) is one *optimization* for quickly finding a working path, rather than a requirement for communication security.

2 Threat Model

Reliable Internet communication can be impaired by attackers who compromise routers or hosts, or accidentally by failures, bugs, and misconfigurations. In a traditional threat model, attackers can tamper with data or impersonate identities (violate integrity), snoop on traffic (violate confidentiality), or deny service (reduce availability). In this section, we first examine why only the last of these threats—availability—must be dealt with by the routing infrastructure. We then examine in more detail the ways an attacker might attempt to deny availability to provide context for understanding the design of ACR.

Integrity can be provided end-to-end using well-known cryptographic techniques (Message Authentication Codes) along with shared secret or public key authentication schemes. Data **confidentiality** is similarly easy to protect using encryption. This leaves **availability** as the remaining threat. Unfortunately, cryptography cannot get packets across a path that drops or misdirects all traffic. [1]

[1] A more subtle threat to confidentiality is traffic analysis, which gleans information simply by observing the pattern of communication between hosts. We argue that senders who need security against traffic analysis are better served by secure techniques, such as mixnets[6], rather than by trusting the ISP infrastructure.

2.1 Malicious Routers

Control, legitimate or illegitimate, of a router grants significant power to compromise communication security.

Control Plane: An attacker can influence the *global* flow of traffic by falsifying BGP routing information. By announcing a victim's IP prefix or manipulating the AS-path, an adversary can draw traffic to its own routers, where it can observe the data, modify it, drop it, or impersonate the destination. An attacker can also prevent a portion of the Internet from hearing the valid route announcement, "blackholing" traffic to the victim. We term the falsification of routing data a "control-plane" attack. Secure BGP proposals restrict the ability of attackers to mount these attacks by providing *origin authentication* and *path validity*.

Data Plane: Despite reducing an attacker's ability to attract traffic, a secure control plane cannot prevent malicious routers or insiders that manage to be on a legitimate communication path from observing, modifying, or misdirecting traffic. Nor does control-plane security protect against packet drops, congestion, or misconfigured forwarding-level constructs such as packet filters. We term these threats "data-plane" attacks.

Because control-plane security must still be augmented with end-to-end techniques to guarantee integrity and confidentiality, we argue that the only property that the control plane must provide is availability; that is, it must guarantee that a sender will hear about a valid path to the destination if one exists.

A final threat comes from attackers who advertise unallocated or unused address space, as is sometimes done by spammers to avoid IP address blacklists [13]. While this technique contributes to spam, it is not the root cause, and we believe that the more fundamental problems of identity and incentives must be solved to effectively reduce spam. For these reasons, we do not consider authenticating "unowned" IP addresses a central requirement of routing.

2.2 Malicious End Hosts

Without access to a router, an adversary may still use end hosts to render a network link unavailable with a DoS attack. Routing protocols (secure or not) do not know what level of congestion will render a link useless for a particular application, and will not switch away even from an unusable "best path." Since link DoS provides another means for an attacker to deny availability, we argue that a routing system should allow a sender to avoid congested paths if a usable alternative exists.

We do not, however, consider DoS attacks targeted at individual hosts or applications to be in scope. Attacks not targeting links can be mitigated by end-hosts or "near-edge" network devices without support from routing.

3 Availability Centric Routing

The goal of *availability-centric routing* is to enable edge-networks or end hosts, which we collectively refer to as the network *edge*, to communicate securely even if portions of the network infrastructure are controlled by an adversary. ACR does so using four components. First, one or more transit ASes act as *availability providers* (APs) that provide the edge with multiple routes for each destination. Second, sources using ACR cryptographically verify the identity of the destination host or network, to confirm that the route they chose reaches the correct destination. Third, ACR senders securely monitor the communication performance; if performance is too poor, for whatever reason (a situation-specific definition), they signal ACR to use a different path. Fourth, the ACR senders distribute traffic over one or more paths supplied by the AP, by applying selection algorithms that quickly identify a working path with low network and computation overhead.

3.1 Multipath via Availability Providers

To provide path-choice in a legacy, single-path BGP environment, ACR includes mechanisms to advertise multiple paths for a single destination and then direct traffic onto these alternate paths. This approach is akin to proposed multipath schemes like MIRO [18]. Availability Providers give the network edge access to multiple paths via a (presumably paid) AS-level *deflection service*. Edges can avoid failures by redirecting traffic to different paths.

An availability provider maintains a *route repository* containing all routes learned from BGP peering sessions with neighboring ASes. The repository may be populated by passive BGP sniffers at peering links, or by a BGP monitoring protocol. Customers can pull routes on-demand from their AP (e.g., if their current path is not working), or subscribe to a feed of paths to particular destinations using either a custom protocol (future work) or the proposed *add-paths* extension to BGP [16].

Sources use alternate paths by tunneling packets using IP encapsulation (e.g., L2TPv3 [11]) to *deflection points* in the AP's network. Paths from the route repository include the deflection-point IP address, the encapsulation method to use, and a *deflection forwarding identifier*. This tunneling can be performed at line rate by high-end routers [8], which enable encapsulated packets to circumvent normal BGP routing using *directed forwarding*, which uses an alternate forwarding table to direct packets based on the deflection forwarding identifier in the encapsulation-layer header. After decapsulation and directed forwarding, subsequent routers forward the packet normally.

Access to the deflection service can be efficiently controlled by light-weight authentication "cookies" found in L2TPv3 and other protocols.

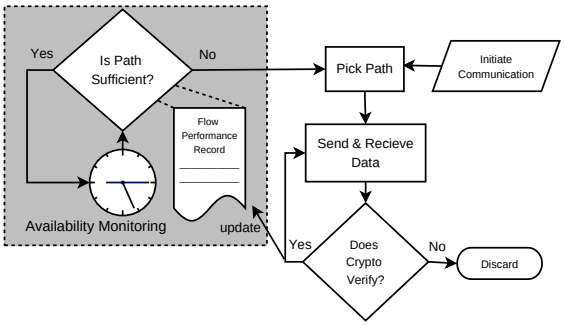

Figure 1: High-level control flow diagram of ACR.

3.2 End-to-End Integrity Check

To work, a path must connect the source to the *correct* destination. ACR allows hosts and applications to authenticate destinations in whatever way they choose, from generic approaches such as IPsec or SSL to application-specific approaches like DNSSEC. Many important protocols, including HTTP, SMTP, SSH, and SIP, already support both client and server authentication. Importantly, ACR does not require either hosts or routers to participate in a PKI. In fact, clients who contact a server may not require cryptographic authentication at all: similar to common use of HTTPS, they can dynamically establish a shared secret used to verify the integrity of all further packets.

3.3 Availability Monitoring

Detecting availability attacks requires the ability to monitor a network flow and determine whether its performance indicates that the current network path is a usable route.

In the context of Figure 1, consider a general purpose monitor within the TCP stack of an end-host using IPSec for end-to-end security. A call to *connect()* causes the path-selection component to select an initial route. TCP sends a SYN packet and sets its retransmission timer. If the timer expires before the SYN/ACK comes back, the monitor records the event and *may* change to an alternate path before retransmitting. Similar monitoring occurs for all data transfered. With TCP, the "flow performance record" consists primarily of state the protocol already keeps to manage reliable delivery, but could be augmented with retransmission or timeout counters to track recent path performance. This record must be reset each time a new path is selected, but TCP-specific behavior and state is otherwise unmodified. Received packets are verified for integrity using IPsec and are discarded if the check fails. As a result, paths with adversaries that manipulate packets are correctly recognized as unusable.

While this example monitor is simple and general, ACR can work with any type of availability monitoring the edge chooses to employ. Monitoring could be even be performed by the user (e.g., by clicking "reload" in their browser). Many applications such as VOIP clients already incorporate application-specific monitoring, and could use this information to change paths when conditions are unacceptable. Hosts could perform monitoring within the network stack, or edge routers could use a scheme similar to Listen[15] to provide simple connectivity monitoring.

The primary requirement for *secure* availability monitoring is that all monitoring decisions must only accept as input, data or network headers that are protected against tampering by an integrity check. Otherwise, an on-path adversary can falsify replies (e.g., TCP acknowledgments) to make it appear that data was correctly delivered.

3.4 Path Selection Algorithms

Path selection algorithms should quickly locate working routes, to minimize the time to recover from failures or attacks. These algorithms are triggered by the availability monitors when failures are detected. Path selection algorithms can combine topological information (e.g., BGP AS-paths) with external knowledge (e.g., AS connectivity or history of good routes) to select new candidate paths. ACR treats this information as *hints*, not truth, because the information may be stale (in the case of history) or inaccurate (in the case of data from unsecured BGP). Path selection could explore several paths in parallel to further reduce recovery time at the expense of additional bandwidth. Selection can be assisted by heuristics such as:

Static destination connectivity hints: Destinations that care about availability are likely to know their upstream connectivity. ACR can use this knowledge to give the edge with "hints" to quickly identify promising paths. BGP paths that are inconsistent with the connectivity hint from the destination receive lower priority in the path exploration process. Because their consistency is not critical (they affect only priority) static hints can be distributed ahead of time, out of band, or via replicated repositories.

Route stability heuristics: Many Internet routes, particularly those to popular destinations, are quite stable [14]. ACR could take advantage of prior work that uses historical route information to identify good paths more quickly. Unlike schemes that discard routes that fail historical tests, and so require exceptionally low "false-positive" rates, ACR will use "anomalous" routes if (and only if) they work correctly end-to-end.

Useful communication, as well as availability monitoring, realistically requires bi-directional reachability. While we describe ACR primarily from the perspective of a single source, we envision that common use of ACR would involve both communicating parties having the ability to find alternate paths.

4 ACR with Limited Deployment

In the long term, we envision ACR being used with a globally deployed multipath protocol. Yet even when only deployed by a handful of tier-1 ISPs, we demonstrate in §5 that ACR significantly improves availability in the face of routing attacks.

The key problem with limited deployment is "legacy" providers still running single-path BGP. As a result, routing attacks could render some destinations unreachable: if a destination D has only a single (legacy) provider P, then if P believes and propagates a false route for D, no availability provider would be able to reach D.

Therefore, ACR, when deployed at limited locations, requires additional light-weight control-plane countermeasures to prevent such control plane availability attacks. Before evaluating the resilience of limited ACR deployment to invalid announcements of a victim's address space (BGP hijacks) in §5, we cover two other issues related to providing availability in a legacy environment.

Resisting sub-prefix hijacks: An attacker can announce a prefix more specific than a legitimate advertisement. This attack is particularly effective because the invalid prefix propagates to all ASes and the more specific route is always used to forward traffic. If a destination is not directly connected to its availability provider, any legacy providers between the AP and the destination that hear the sub-prefix announcement will misdirect received packets to the attacker despite also having a correct but less specific route to the destination.

We propose eliminating sub-prefix attacks by emulating "flat addressing" within the limited scope of a destination and its upstream providers. If an upstream provider P agrees to accept address space from customer D as /24's and filters all incoming prefixes of greater length (as is common practice by ISPs today) no adversary can sub-prefix hijack D's address space. Peers and other customers of P are not on the path between the AP and D, so P can safely aggregate D's addresses before advertising the prefixes to these neighbors. Effectively, upstream providers accept a moderate increase in routing table size to provide increased availability for their customers, while the global routing table size remains constant.

CIDR-based hierarchical addressing, the root-cause of the sub-prefix hijacks, is also troublesome for other efforts to secure routing and forwarding. For example, sub-prefixes in forwarding tables are a primary reason that the control plane can differ from the actual path traversed by a packet, mitigating the benefit of having a secure BGP AS-Path. Similarly, prefix aggregation significantly complicates origin authentication. While we propose an incremental measure for dealing with CIDR above, ultimately we feel that the only sound architectural choice is to move toward a flat addressing model for the Internet.

Resisting deflection point hijacks: A BGP hijack could block a subscriber from reaching its AP's deflection points if the subscriber's direct upstream provider did not support ACR.[2] Fortunately, the number of deflection points is relatively small, and they are found in known locations within stably connected core networks. These properties facilitate "defensive filters" that explicitly deny route announcements for special destinations on all but a select few peering sessions[17]. Willing legacy providers can also use a simple mechanism like BGP's localpref attribute to assign static preference for AP prefix announcements heard via the links that a provider expects to use to reach the associated tier-1 provider. This defense is more flexible than simple static routing, but still mitigates hijack attacks against the availability provider.

5 Evaluation

We explore the effectiveness of ACR and its countermeasures in the context of today's Internet. In our evaluation, each path may contain at most one deflection point and only a few ASes offer deflections. Our experiments examine ACR's performance against an attacker who announces an IP prefix that belongs to a victim network.

Method: We run simulations on an AS-level graph based on July 2006 RouteViews data with AS relationships inferred using Gao's algorithm [9]. The route selection policy prefers customer-learned routes over peer-learned routes, and prefers provider-learned routes the least, with ties broken using AS-Path length. Each trial has one legitimate AS and a set of attacking ASes that all announce the same prefix. We vary the number of malicious ASes, performing 100 trials for each configuration.

Result 1: A single tier-1 availability provider significantly increases routing robustness compared to stubs using either single-path BGP or intelligent multihoming. Figure 2 charts the average reachability of the legitimate destinations versus the number of attacking ASes. The bottom line (Single-Path BGP) shows the average success rate of all stub ASes in reaching the destination using normal BGP. We simulate intelligent multihoming by testing all stub ASes with exactly five providers to see if any of their five BGP-learned routes are valid.[3] The availability providers for the Tier-1 AP data include all ten ISPs commonly thought to not purchase transit from another ISP. The results indicate the average success rate for these ISPs using deflections on all BGP-learned paths.

While intelligent multihoming sources can select from multiple paths, *only a tier-1 availability provider exposing multiple BGP-learned paths to the same destination provides strong resilience to hijacks*. ACR works so well be-

[2] This customer would have an incentive to switch to an ACR-speaking ISP, but we also believe that customers can benefit from using a "remote" (i.e., non-first-hop) availability provider (§6).

[3] A selection intended to capture stubs that have invested significantly in network availability.

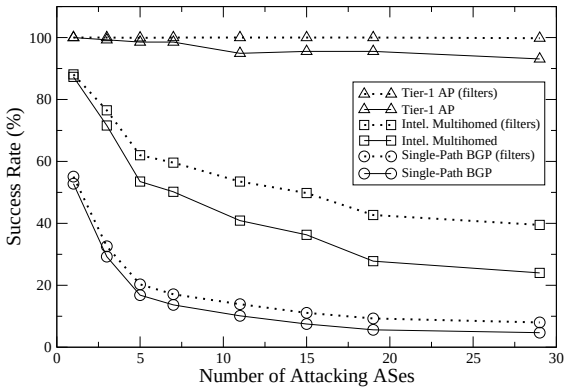

Figure 2: Success rate of sources reaching a hijacked destination when using different degrees of path diversity.

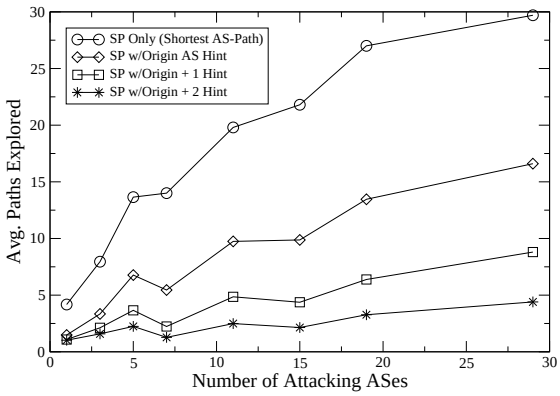

Figure 3: Number of routes explored before finding a valid forwarding path.

cause of the legacy ISP's preference for customer-learned routes, which forces an attacker to be "local" (a customer of all of a destination's providers) to prevent the AP from hearing the legitimate announcement.

Result 2: ACR's availability benefits can be further improved using easily-deployed BGP filtering local to the victim. As shown in Figure 2, adversaries are sometimes assigned to local ASes, reducing the Tier-1 AP success rate to 95% with many attackers (e.g., second from top line, far right). To defeat these adversaries, legacy ISPs can employ a tactic already common among large providers today: filtering routes from customers to accept only prefixes that the customers own and have registered. As a result, these filters block malicious advertisements by other customers. Unlike filtering to protect the legacy BGP system (which must be performed globally), these filters need only be applied locally by some of the valid destination's transit providers. The results of applying such filtering at the ISPs between the tier-1 AP and the destination are shown by the "filters" lines. The results show that *filters provide complete protection with a tier-1 AP, but provide only incremental benefit for intelligent multi-homing or single-path BGP*.

Result 3: The time to find a valid route is reasonable in the face of many adversaries, and simple connectivity hints from the destination further speed the process. Figure 3 shows the average number of paths a source must explore, averaged over all Tier-1 APs, without the benefits of destination filtering. The *Origin AS Hint* case assumes that the source knows the correct AS originating the prefix being probed, while *Origin + x Hint* indicates knowledge of all upstream providers up to x hops from the origin.

Without external topology information, ACR explores paths based only on their AS-path length. ACR must test a few paths per attacker before finding a working path, which we feel is not unreasonable. However, guiding path selection with some prior knowledge of topology provides improved efficient, requiring probing only a few paths even for large numbers of attackers. The topology hints force an adversary to pad its AS path to include the correct topology, which makes the path longer and less attractive to the short AS-path heuristic. Using these heuristics, ACR helps reduce outages to short "hiccups" in connectivity experienced while it explores new paths.

6 Deployability

ACR emphasizes low barriers to adoption: ACR simplifies deployment because it does not require cryptographic hardware in routers and because the functionality needed to perform path deflections is already widely available. Also, because parties with significant security requirements already use end-to-end security, ACR obviates the need to manage BGP authentication services, such as an AS-level PKI and address ownership registries.

ACR benefits from backward compatibility: Changing a critical part of the Internet infrastructure raises stability and reliability concerns. Because ACR runs along-side BGP, not as a replacement, operators can evaluate it on operational networks without the need for a parallel test infrastructure. Additionally, failures within ACR are isolated from BGP. As a result, unlike many secure replacements for BGP, legitimate use or misconfiguration of ACR is unlikely to result in worse reachability than provided by legacy BGP, because the single-path legacy route is still available for use.

ACR provides well-incentivized deployment: We envision deflection services being offered in two ways. First, core networks can offer deflections to customers in order to add value to their existing transit service. This could give an ISP a competitive advantage over providers that do not offer deflections: customers will receive improved resilience against attacks and gain the ability to select paths that perform better.

The second deployment mechanism is to offer a remote deflection service to customers of other providers. This

service would enable customers of legacy ISPs to gain many of ACR's benefits. This remote deflection service (in some ways, a "virtual ISP") is technically more challenging to offer, but as §5 showed, even deployment by a single large ISP can provide greatly improved attack resilience. An AP can offer remote deflection service more cheaply than normal transit service because (1) availability customers need no physical router port and (2) a tier-1 AP also receives more overall transit revenue because of increased traffic entering their network for deflections. As a result, stubs with both types of providers need not be "double-charged" for their connectivity.

7 Related Work

Secure routing has been pursued extensively in academia and industry; due to space constraints, we refer the interested reader to a recent survey of BGP security research [5]. ACR's path selection can benefit from secure routing protocols, but remains effective without them.

Popular current approaches for robust routing use overlay networks [2] or multi-home the edge [1]. While these techniques improve availability against many failures, we know of no studies that examine their resilience to deliberate attacks on the routing infrastructure. Our evaluation suggests that they cannot withstand powerful adversaries that use BGP to globally disrupt routes to a destination.

Many clean-slate source-routing architectures either do not address security (e.g., NIRA [19]), or conflict with operational practices (e.g., feedback based routing [21]) by requiring the disclosure of routing policies often guarded today by non-disclosure agreements.

Multipath interdomain routing protocols like MIRO [18] provide a foundation for communicating the multiple paths required by ACR. Recent work on router-level deflections [20] offers a complementary technique that provides finer-grained path diversity, but with less source control over how packets are deflected; ACR could leverage such techniques to help avoid adversaries within an AS. Work on Stealth Probing [4] describes a secure method for probing network paths that could serve as an availability monitor between edge-networks.

8 Conclusion

The goals of traditional secure routing (availability and communications security) can be achieved *without* securing the routing protocols. Because properties such as confidentiality and integrity can and should be provided end-to-end, this paper argues that availability is the only property that the routing system must provide. Availability, we believe, is better served by lightweight, incentive-compatible mechanisms to provide multiple paths to end hosts than by heavyweight secure routing techniques.

ACR is easier to deploy than traditional secure routing protocols and provides stronger incentives for incremental deployment. ACR can effectively defeat control-plane adversaries (§5), and by design it can circumvent more problems ("data plane" adversaries and failures) than is possible by merely securing the control plane. Because of these benefits, we believe that ACR is a worthwhile addition to the routing lexicon, regardless of whether a secure version of BGP is eventually deployed.

References

[1] A. Akella, J. Pang, B. Maggs, S. Seshan, and A. Shaikh. A comparison of overlay routing and multihoming route control. In *Proc. ACM SIGCOMM*, Aug. 2004.

[2] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient Overlay Networks. In *Proc. 18th ACM Symposium on Operating Systems Principles (SOSP)*, pages 131–145, Oct. 2001.

[3] Arbor Networks. *Infrastructure Security Survey*, 2006.

[4] I. Avramopoulos and J. Rexford. Stealth probing: Efficient data-plane security for IP routing. In *Proc. USENIX Annual Technical Conference*, May/June 2006.

[5] K. Butler, T. Farley, P. McDaniel, and J. Rexford. A survey of BGP security. Technical Report TD-5UGJ33, AT&T Labs, June 2004.

[6] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 4(2), February 1981.

[7] B. Christian and T. Tauber. *BGP Security Requirements*. IETF, Apr. 2006. Internet Draft: draft-ietf-rpsec-bgpsecrec-06.txt.

[8] P. Francios and O. Bonaventure. An evaluation of IP-based fast reroute techniques. In *Proc. CoNEXT'05*, 2005.

[9] L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Trans. Netw.*, 9(6):733–745, 2001.

[10] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (S-BGP). *IEEE JSAC*, 18(4):582–592, Apr. 2000.

[11] J. Lau, M. Townsley, and I. Goyret. Layer two tunneling protocol - version 3 (L2TPv3). RFC 3931, IETF, Mar. 2005.

[12] J. Ng. *Extensions to BGP to Support Secure Origin BGP (soBGP)*. IETF, Apr. 2004. Internet Draft: draft-ng-sobgp-extensions-02.txt.

[13] A. Ramachandran and N. Feamster. Understanding the Network-Level Behavior of Spammers. In *Proc. ACM SIGCOMM*, Aug. 2006. (to appear).

[14] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang. BGP routing stability of popular destinations. In *Proc. ACM SIGCOMM Internet Measurement Workshop*, Nov. 2002.

[15] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. Katz. Listen and Whisper: Security mechanisms for BGP. In *Proc. Symposium on Networked System Design and Implementation*, Mar. 2004.

[16] D. Walton, A. Retana, and E. Chen. *Advertisement of Multiple Paths in BGP*. IETF. Internet Draft: draft-walton-bgp-add-paths-05.txt, Expired August 2006.

[17] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. F. Wu, and L. Zhang. Protecting BGP Routes to Top Level DNS Servers. In *Proc. 23rd Intl. Conf on Distributed Computing Systems*, pages 322–331, May 2003.

[18] W. Xu and J. Rexford. MIRO: Multi-path interdomain routing. In *Proc. ACM SIGCOMM*, Sep. 2006.

[19] X. Yang. NIRA: A New Internet Routing Architecture. In *ACM SIGCOMM Workshop on Future Directions in Network Architecture*, Aug. 2003.

[20] X. Yang, D. Wetherall, and T. Anderson. Source selectable path diversity via routing deflections. In *Proc. ACM SIGCOMM*, Aug. 2006. (to appear).

[21] D. Zhu, M. Gritter, and D. Cheriton. Feedback based routing. In *Proc. HotNets-I*, Oct. 2002.