

Incentives for Deploying Secure Routing

Dan Wendlandt - Carnegie Mellon University - dwendlan@cs.cmu.edu

The Internet’s routing and forwarding infrastructure is vulnerable to attacks and misconfigurations that threaten the ability of hosts to communicate. Attacks that inject false information into BGP or compromise a router on the data path can impair data confidentiality & integrity, and disrupt a destination’s network availability. Proposals to secure BGP have existed for nearly a decade (s-BGP was begun in 1996), yet researchers, vendors, and operators have failed to reach any consensus on a solution.

1 The Need for Incentivized Adoption

Why have past solutions, while technically correct, failed to gain significant traction? I argue that by emphasizing security benefits provided under universal deployment and cryptographic performance, past secure routing proposals failed to pay enough attention to other key factors inhibiting deployment. These deployment road-blocks include: (1) limited and ill-defined security benefits throughout partial deployment; and (2) security benefits *entirely* reliant on a cumbersome to create and maintain registry and PKI to authenticate prefix ownership.

Both factors contribute to a “chicken or the egg” problem for secure BGP protocol adoption. ISPs do not want to assume the risk of investing time and money to upgrade router hardware/software, create and maintain registries, and participate in a PKI until they believe that such steps will provide commensurate benefits by helping them selecting better outgoing routes and improving global reachability to their own prefixes. As a result, few customers are pressuring their upstream providers or router vendors to make any of the proposed solutions a reality. Recent work on modeling the adaptability of BGP [1] takes an important first step by considering past proposals with respect to the benefits of early adopters, but ignores the second issue of relying on accurate registries, likely because all past proposals examined in the study required an address registry and PKI.

Taking an alternate approach, instead of asking “what must be authenticated to secure the BGP protocol?”, we consider the more basic question of “how can hosts communicate information reliably and securely over the Internet?”. We find that considering secure communication despite an insecure or partially secure control plane introduces a new design space with promising deployment incentives.

2 Multi-Path Routing & End-to-End Security

Secure communication requires protecting data confidentiality, data integrity, and the availability of the communication channel. Most security-sensitive traffic today is already protected against the first two threats end-to-end, either using SSL built into popular protocols (e.g., HTTPS) or with IPsec or SSL tunnels protecting legacy traffic. *This leaves availability as the only security property that requires support from the routing and forwarding infrastructure.* As a result, we consider a simplified design space where the infrastructure has a single goal: to protect against availability attacks by allowing a sender to discover multiple paths and use any available and non-malicious path to reach the destination.

We refer to this approach as Availability-Centric Routing (ACR)[3], because it focuses on achieving only availability in the routing infrastructure, while leaving confidentiality and integrity to be performed completely end-to-end. With ACR, routers do not need address registries and a PKI to perform cryptographic verification of route announcements. Instead, routers can supply multiple paths to the source, which can then use already common end-to-end security mechanisms (e.g., SSL certificates) to verify they are communicating with the correct destination. ACR can provide strong robustness because an attacker must now prevent a source from hearing any valid route to the destination in order to deny availability.

Our analysis [3] demonstrates that even a single tier-1 ISP exposing alternate BGP-learned routes would provide sufficient path diversity to avoid nearly all routing attacks. This diversity could be sold as a value-added service, at little cost to the tier-1 ISP. Multi-path forwarding over IP (similar to that in MIRO[4]) is straightforward for ISPs to offer, even to non-directly connected customers, because newer router hardware supports the required IP-in-IP encapsulation at line-rates. As a result, stub ASes that receive a “feed” of alternate routes to select from can achieve vastly improved routing robustness, even if the majority of the Internet (including their own providers) has not

adopted a secure routing scheme. Additional benefits provided by ACR include the ability to use multiple paths to route around data plane attacks or find paths with better latency/capacity/loss (a rare example when both security and performance benefit from the same architecture).

Our evaluation of ACR also demonstrates that the number of paths a source must explore remains reasonable, even in the face of many attackers. However, using inexpensive, easily implemented heuristics (e.g., first trying paths that have worked in the past) or a partially populated secure topology database can speedup this process by prioritizing the exploration of promising routes. In this way, *control plane security can be seen as an optimization, rather than a requirement for secure communication.*

3 A Progressively Secure Control Plane

ACR is an example of how secure communication is possible despite an insecure control plane, yet even this approach (along with most secure routing proposals) clearly benefits from better data on control plane topology. Therefore, we ask: what is needed to “jump-start” better control plane security?

Insecure topology data provides an interesting means of avoiding the “chicken or the egg” problem in BGP security adoption. If sources begin making routing decisions based on useful yet imperfect knowledge, destinations have an incentive to supplement any inaccurate information with *secure* and correct information in order to improve its own reachability.

As an example, consider a scheme like PGBGP [2] in which sources flag prefix/AS origination pairs never seen before as “suspicious” and temporarily lowers that route’s local-preference (this allows the valid destination to notice and mitigate an attack before it impairs reachability). If a destination wanted its new route accepted immediately, it could choose to participate in a secure topology database where updates are vetted for authenticity. Thus, the destination has a clear understanding of the benefits of participating in the secure routing scheme and the control plane becomes progressively more secure, until eventually historical data may no longer be required.

Using the past existence of a route as a weak “proof-of-validity” is interesting in that it highlights the trade-off between the willingness of a source to accept new routes and the source’s vulnerability to attacks. For security-sensitive destinations, a source may prefer to keep an older but known-good route instead of using a potentially better performing but unproven new path. This trade-off, however, cannot be made by an upstream provider on behalf of all customers (as is required with single-path BGP). Therefore, similar to the case of using end-to-end authentication with ACR, *the usefulness of unauthenticated control plane information seems most powerful when used along with multipath routing.* This helps stub ASes avoid the case where each upstream offers a single path, none-of-which are desirable according to the stub’s security metrics.

In addition to insuring that a participant’s security benefits are always match their participation costs, any partial control plane security proposal must also :

1. Minimize demands on ISP resources by making requirements that closely map to operational reality.
2. Avoid introducing new, even temporary, vulnerabilities into the routing system.
3. Provide a clear path to a final solution that offers highly reliable security.

That is to say, the partial cure should not be worse than the disease and must not simply be a hack to provide limited improvement. Our recent work suggests that multipath routing systems provide a compelling, and relatively unexplored, design space within which to consider these requirements.

References

- [1] H. Chan, D. Dash, A. Perrig, and H. Zhang. Modeling adoptability of secure BGP protocol. In *Proc. ACM SIGCOMM*, Sep. 2006.
- [2] J. Karlin, S. Forrest, and J. Rexford. Pretty good BGP: Improving BGP by cautiously adopting routes. In *Proc. International Conference on Network Protocols*, Nov. 2006.
- [3] D. Wendlandt, I. Avramopoulos, D. Andersen, and J. Rexford. Don’t secure routing protocols, secure data delivery. <http://www.cs.princeton.edu/~jrex/papers/acr.pdf>.
- [4] W. Xu and J. Rexford. MIRO: Multi-path interdomain routing. In *Proc. ACM SIGCOMM*, Sep. 2006.