

Workable Routing Security

Steven M. Bellovin
smb@cs.columbia.edu
Columbia University

That routing security is important has long been known [2, 8]. The problem was described as urgent in a National Academies study [10]. Since then, there have been many proposals, including [7, 6, 1, 9, 4, 5]. None have been accepted by the community, let alone deployed.

There are many reasons for this, including simple politics and turf-building. Rather than explore the reasons for failure in detail, we instead set forth some principles that a future secure routing architecture should follow. In a nutshell, it should *work*.

The word “work” encompasses a wide range of notions. Among these are realistic assumptions, deployability, economics, operability, and (of course) security. We explore each in turn.

Note carefully that some of the points we make are, if taken to their logical conclusions, mutually contradictory. Delicate compromises will be needed. This in turn suggests the need for not just fundamental research but also good taste in engineering.

Realistic Assumptions

The real world is not a simple place. Assumptions about the control of autonomous systems¹ must match reality: any given individual or company may control more than one AS, overtly or covertly, and may co-operate or collaborate with other AS owners, even those that are nominal competitors.

Deployability

Whatever scheme is used must be deployable. 25 years of TCP/IP and the Internet have taught certain lessons about that. Systems that require heavy, up-front infrastructure and/or centralized authorities are harder to deploy. Even if a central registry is conceptually necessary, it helps if the cost — in hours, interactions, and complexity, as well as in

¹We do not assume that a future routing architecture will include AS’s of today’s type. That said, it seems very likely that there will be some such notion, if only because enterprises and ISPs will wish to control their own routing.

dollars, euros, or zorkmids — grows with the net. IP addresses are a good example. Once upon a time, they were handed out by one person, on request; today, elaborate justifications are needed, to an ISP or RIR, demonstrating an adequate utilization level, numbering plan, etc. Schemes that can be built from the edges in are the most successful.

Economics

The role of economics is two-fold. First, the cost of the defenses must be commensurate with the potential loss avoided. Spending \$1 to save \$1,000,000 is obviously worthwhile; conversely, spending \$1,000,000 to avert a \$1 loss is dubious. The difficulty here is that neither the magnitude of the potential loss nor its probability are clear.

More subtly, the cost burden must fall on those who benefit from it. Schemes that require significant expenditures by, say, all end systems, while protecting the ISP, are not likely to catch on.

A corollary of this is that organizations will need to impose policy constraints, in order to minimize their own costs or maximize their own revenues. Any secure routing scheme has to permit local policies to be implemented.

Operability

The operational cost must be acceptably low. This refers not just to direct monetary cost or complexity — which itself implies higher cost — but also to dealing with failures. More or less by definition, a secure routing scheme requires that some advertisements be rejected. False rejections, due to protocol, implementation, or administrative flaws, must be handled expeditiously. Schemes where rectification of errors is harder will tend to drive customers away from the secure network. This may imply that reliance on a central authority is unwise.

As with direct economic costs, the operational cost must be borne by those who benefit; more generally, communications should flow along the paths of direct business relationships. Providers are not beholden to random strangers, and have no incentive to correct errors on their behalf.

Security

It may seem redundant to list security as a requirement for a secure routing system. Nevertheless, it is important to focus on it. “Security” can have many meanings and many different shades. What does it mean for a routing protocol to be “secure”? How powerful is the enemy? What percentage of packets must flow as intended? Must the solution be secure against short-cut attacks? Link-cutting attacks [3]? How much traffic overhead can we afford, to send multiple copies of packets by different paths?

Once we have a definition, we can assess the strength of a proposed solution. Ideally, this will be done formally. Even there, though, we must be careful; cryptographic protocols that have been proven secure have later been cracked by people who made slightly different assumptions [11].

References

- [1] T. W. an dEvangelos Kranakis and P. van Oorschot. Pretty secure BGP (psBGP). In *Proceedings of the Symposium on Network and Distributed System Security*, February 2005.
- [2] S. M. Bellovin. Security problems in the TCP/IP protocol suite. *Computer Communications Review*, 19(2):32–48, April 1989.
- [3] S. M. Bellovin and E. R. Gansner. Using link cuts to attack Internet routing, 2003. Draft.
- [4] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin. Working around BGP: An incremental approach to improving security and accuracy of interdomain routing. In *Proceedings of the IEEE Network and Distributed System Security Symposium*, February 2003.
- [5] Y. Hu, A. Perrig, and M. Sirbu. SPV: Secure path vector routing for securing BGP. Portland, OR, August 2004. Proceedings of ACM SIGCOMM.
- [6] S. Kent, C. Lynn, J. Mikkelsen, and K. Seo. Secure border gateway protocol (S-BGP) – real world performance and deployment issues. In *Proceedings of the IEEE Network and Distributed System Security Symposium*, February 2000.
- [7] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (Secure-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, April 2000.
- [8] R. Perlman. *Network Layer Protocols with Byzantine Robustness*. PhD thesis, M.I.T., 1988.
- [9] Russ White, ed. Architecture and deployment considerations for Secure Origin BGP (soBGP), April 2004. draft-white-sobgparchitecture-00.txt.
- [10] F. B. Schneider, editor. *Trust in Cyberspace*. National Academy Press, 1999.
- [11] D. A. Wagner. Cryptanalysis of a provably secure crt-rsa algorithm. In *Proceedings of the ACM Computer and Communications Security (CCS) Conference*, October 2004.