Position Paper: Operational Requirements for Secured BGP

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

John Ioannidis
ji@cs.columbia.edu
Columbia University

Randy Bush
randy@psg.com
IIJ

We are all aware that routing security is a problem. There were warnings in the literature more than 15 years ago [Bel89, Per88]. In 1999, a National Research Council study called it one of the crucial vulnerabilities of the Internet [Sch99]. There have been fully worked-out solutions for several years [KLS00, KLMS00] and even running code. There have been notorious Internet outages, such as the "AS 7007" incident, due to routing configuration errors. Routers owned by major ISPs have been compromised by hackers. There are even reports of spammers engaging in route hijacking [Ao03]. Despite that, there is not only no production-quality code, there is not even a standard.

From a security perspective, S-BGP [KLS00, KLMS00] is a very complete solution. Certificates from authoritative sources are used to identify the owners of autonomous systems (ASs) and IP address blocks. Each AS (or even router) hop is digitally signed; anyone who wishes can verify that the path is genuine. However, it still does not address issues such as detecting policy violations, or incorrect propagation of route announcements or withdrawals. Moreover, there is no consensus for it. Nor is there any consensus for other solutions, such as [Rus04, adKv05]. We suggest that a major reason is that insufficient attention has been paid to operator concerns.

There have been a number of objections raised to the various proposals. Some are inherent to any possible solution; there will indeed be new ways that connectivity can fail, for example. This is inevitable; a purpose of *any* secure routing mechanism is to cause certain claimed routes to be rejected. If the originator does not use the proper credentials in the proper form, the route will indeed be ignored. This is as it should be. In a less justifiable vein, misconfiguration of a router could cause it to reject a valid secure routing announcement. Something that today is relatively innocuous—a router with the wrong idea of the current date—could cause certificates to be rejected as expired. Other failure modes include buggy software or problems with the PKI structure. In short, secured BGP adds complexity and fragility. We understand and accept this; we also assert that after a short learning process, the total number of failures due to routing problems will decrease, since large classes of routing configuration errors will be eliminated.

Another major concern is the financial cost of secure routing. This cost can be broken down into several categories: capital expenses, ongoing operational expenses, and database creation costs.

The last represents an area where money can profitably be spent now. The database of address assignments is dirty, especially for very old assignments in the so-called "swamp space" [LAB01]. Any conceivable secure routing initiative is going to require authoritative address block ownership data; we should start cleaning it up now. (We also note that a very good case can be made for U.S. and non-U.S. government funding for this effort. The benefit of the cleanup cannot easily be attributed to any of the current RIRs, though they have already done a lot to clean up the data; furthermore, many of these assignments took place at a time when the Internet was still largely government-funded.)

There is more variability among different secure routing proposals when it comes to capital expenses. In a practical sense, some would *require* hardware add-ons in order to perform the necessary digital signature calculations. The chips themselves might be relatively inexpensive; the necessary engineering, board space, router upgrades, etc., will not be. Persuading ISPs of the necessity of such spending is problematic, since the potential for profit seems limited. Still, the big expense will be in the initial deployment; we can assume that the necessary hardware functionality will follow the usual Moore's Law curve and rapidly become insignificant. The important thing is to leave the necessary hooks now.

Operational expenses are a thornier problem. While the details will differ somewhat among the different proposals, the expense of issuing credentials for address blocks is inherently an ISP and RIR responsibility. These are the entities that assign addresses; no one else has authoritative information. Development of a "CA in a box" will help; even with that, however, operation of such a CA will have to be integrated with existing operational support systems. Furthermore, there is likely to be an ongoing customer care expense.

Some proposals, notably [Rus04], require publication of routing policies. While many policies can be intuited by examining public routing measurements [Gao00], some major ISPs have strongly resisted publishing such data in the past. There are only two choices here: either a solution must be adopted that does not require such publication, or

operators must be persuaded that they in fact lose very little by publishing; there is a long history of poor success with the latter.

One often overlooked problem in routing security analyses problem is that of *anomaly containment*. Because of the egalitarian nature of the BGP protocol, incorrect information that originates with any AS, no matter how small or large, can quickly propagate throughout the Internet and cause disruptions [ZPW+01]. Most anomalies that result from incorrect data being introduced into the routing process by badly-configured or compromised routers could be stopped one hop away if everybody in the Internet were running a secured routing protocol. However, this is unlikely to happen, so the part of the Internet that is actually running more advanced security mechanisms has to protect itself from the rest.

Finally, there is the problem of phased deployment. It is obviously impossible to deploy any solution instanteously; none of us have powerful enough magic wands. The deployability characteristics of the different solutions may be a key distinguishing factor. An ideal solution would permit phased deployment, not just within the Internet but within an ISP. Islands of secured BGP should be able to do partial verification of each others' routing advertisements, even when connected to or separated by ISPs that do not support new security mechanisms. Larger contiguous regions should be correspondingly more secure. A mechanism must exist that will permit authoritative determination of whether or not a particular routing announcement should have been secured.

We have deliberately not discussed the security properties of the different proposals. That said, they do differ in the kinds of attacks they can deflect. It makes little sense to go to great expense and trouble deploying a solution that will soon be useless against more capable adversaries, including ones who are capable of compromising critical routers. We thus suggest that a good initial solution is one that can easily be upgraded to handle increased threats. For example, a scheme that had tunable knobs for the number signatures on an AS path could be set to "low" initially, thus saving the verification expense (and the hardware that that would require); as the threat level increased, more signatures could be added and verified. In the same vein, perhaps an AS could include a security level in its routing announcements; this would be used to control the signing and verifying behavior of downstream routers.

References

[adKv05] Tao Wan an dEvangelos Kranakis and P.C. van Oorschot. Pretty secure BGP (psBGP). In *Proceedings of the Symposium on Network and Distributed System Security*, February 2005.

[Ao03] Xuhui Ao. Report on DIMACS Workshop on Large-Scale Internet Attacks. Technical report, Rutgers University, November 2003.

[Bel89] Steven M. Bellovin. Security problems in the TCP/IP protocol suite. *Computer Communications Review*, 19(2):32–48, April 1989.

[Gao00] Lixin Gao. On Inferring Autonomous System Relationships in the Internet. IEEE Global Internet Symposium, November 2000.

[KLMS00] Stephen Kent, Charles Lynn, Joanne Mikkelson, and Karen 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.

[KLS00] Stephen Kent, Charles Lynn, and Karen Seo. Secure border gateway protocol (Secure-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, April 2000.

[LAB01] Craig Labovitz, Abha Ahuja, and Michael Bailey. Shining Light on Dark Address Space. Technical Report Arbor Networks Tech Report, November 2001.

[Per88] Radia Perlman. *Network Layer Protocols with Byzantine Robustness*. PhD thesis, M.I.T., 1988.

[Rus04] Russ White, ed. Architecture and deployment considerations for Secure Origin BGP (soBGP), April 2004. draft-white-sobgparchitecture-00.txt.

[Sch99] Fred B. Schneider, editor. *Trust in Cyberspace*. National Academy Press, 1999.

[ZPW+01] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang. An analysis of BGP multiple origin as (MOAS) conflicts. Proceedings of ACM SIGCOMM Internet Measurement Workshop, 2001.