

Deployable BGP Security

Josh Karlin

Department of Computer Science
University of New Mexico
Albuquerque, NM 87131, U.S.A.

The routing protocol that connects the Autonomous Systems, BGP, is vulnerable to a number of potentially crippling attacks because it trusts unverified control plane information received from external networks. Within the last year we have seen Con-Edison hijack Panix's /16 [1], TNET hijack several /8's including 1247 more specifics of AT&T's 12/8 [2], NW Network Cable announce several /8's including nearly 2800 more specifics [3], multiple ASs (8437, 16215) announce the entirety of published dark space [4], and AS 22773 announced 128/1. We know that these events are bogus because they are so severe. Smaller attacks and misconfigurations are not so easily identified and there are likely hundreds of hijacks that we have not been able to verify.

However, it has been difficult to convince the operator community to adopt any of several proposed security solutions. The draft RFC for soBGP [5] has expired, and the community has not even started distributing AS number certificates, a prerequisite to complete verification. In this paper I discuss some of the deployment difficulties that BGP security enhancements have faced, suggest directions for future research, and propose an initial framework for a simple yet effective security improvement for BGP that can be rapidly deployed.

I. DEPLOYMENT BOTTLENECKS

A. *Incremental Effectiveness*

It appears unlikely that we will have a ubiquitously deployed security solution in the near future. Some estimates suggest that it may take 10 to 15 years. If a secure protocol requires full or substantial deployment to be effective, companies will not have incentive to adopt early because there is no guarantee that the others will also adopt. Therefore, in order to get a secure protocol solution off of the ground it needs to offer a significant benefit to early adopters.

The incremental security benefits offered by today's secure protocols are ill-defined. For instance, should SBGP [6] reject routes when ASs along the path do not run SBGP and strip the optional security data? What guarantees can be made about a path that is only partially verifiable? When is it safe to reject such routes? Can we explore the bounds upon cryptographic and anomaly based security policies under partial deployment in order to make arguments about what benefits can be realistically expected?

B. *Deployment Strategy*

If incremental deployment is a likely scenario, then we need to study the quality of available distribution mechanisms so that resources are not wasted on ineffective deployments. This information should be available by the time a protocol is ready for deployment. Distribution could begin with periphery nodes that are interested in trying new technology and where experimentation can do little harm. Alternatively, it could be deployed at random across the network. Or it could begin with the core of the AS network, which influences a high percentage of routes.

Recent work [7] suggests that core deployments are more effective than random and can provide significant security improvements with only a handful of early adopters. But many questions remain. If only a fraction of the Tier-1 ASs deploy a protocol will it be substantially weaker than if all of them deploy it? Is there a purpose to the periphery deploying the protocol at all? Perhaps the periphery should be used for testing purposes. How many and what types of ASs can we realistically expect to deploy a secure protocol? Can that same deployment maximally utilize the protocol's effectiveness or at least come close to it?

C. *Certificates and Coordination*

Both certificates and global coordination hinder the deployment of SBGP and soBGP. Certificates must first be distributed before either protocol can be used. It is unknown who would do this or when. If the authoritative parties are not willing to distribute the certificates, how else can it be done? Next, registries (distributed or centralized) require organizations to keep their portion of the database up to date. Past attempts at registries have been incomplete and riddled with stale information. What can be done to help bootstrap a clean registry? What will motivate operators to maintain it even when the protocol is only partially deployed?

II. AN ILLUSTRATION

In this section I describe a framework that illustrates a simple and rapidly deployable security improvement for BGP. It is highly effective with a small deployment, does not require operators to register their policy or prefix information, has very low operational overhead, and can be adopted via a router software upgrade.

The protocol can conceptually be split into two pieces, enforcement and monitoring. Transit ASs that would like to ensure that their customers do not route to illegitimate destinations will install the router update. Destination ASs that would like to ensure that packets destined to them arrive safely will monitor the network for aberrations.

An initial deployment of an improved Pretty Good BGP protocol could be deployed on the Tier-1 ASs. It would require the router to maintain a history of known origin ASs, prefixes, and edges as well. New edges, origins, and sub-prefixes would be depreferenced for 24 hours and if they still exist in the RIB at that time, they would be added to the history.¹ The delay period would provide monitoring ASs time to attend to events before they could cause widespread damage.

Routers that run the protocol would distribute information about the new, depreferenced routes either through a registry service like the IAR [8] or a decentralized protocol similar to soBGP. In either case, interested parties could listen (via RSS feeds from the registry or directly to the out-of-band protocol) to the stream of anomalies and use their own local peering/prefix data to filter it. Routes that disobey the local database would be flagged for the NOC to investigate and fix within the delay period.

Any change to the BGP network needs to consider its impact on security and daily operations. First, the proposed framework would not supply full path verification. But none of the proposals do at partial deployment. If registries and certificates were to eventually become available, they could run in tandem with our simple solution, trumping its preference when authenticated information is available. Second, depreferencing is not the same as rejecting, it would be possible for bogus routes to propagate but it would be unlikely to propagate through a Tier-1 which has many alternative routes. Third, each AS that delays the route would delay it for a day. It would be better if a route were only delayed once. By cooperating with the IAR, routers could determine if a route had previously been delayed. This will be explored in future work. Next, legitimate routes could be delayed, but we have shown in [7] that anomalies are rare and few legitimate paths would be affected. Finally, it is possible that adversaries could generate an arbitrary number of alarms. We believe that this problem could be solved with a simple throttling mechanism but it too remains for future work.

This framework is easy to understand, develop, and deploy. It does not require third parties to distribute certificates, ASs to register/unregister their policy and prefix information, or new hardware other than perhaps memory. Alert registries such as the IAR and PHAS already exist and could be improved upon for production use. The majority of the network would be temporarily protected by the Tier-1 delays, which would drastically cut down on the effects of misconfigurations. Finally, an AS that wished to protect itself from longer attacks would only need to subscribe to an RSS feed.

III. ACKNOWLEDGMENTS

Special thanks to Stephanie Forrest, Jennifer Rexford, and Dan Wendlandt for their insightful discussions and thoughts on this paper. The author gratefully acknowledges the support of the National Science Foundation (grants CCR-0331580 and CCR-0311686), the Homeland Security Advanced Research Projects Agency (grant 1756303), and the Santa Fe Institute.

REFERENCES

- [1] T. Underwood, "The anatomy of a leak: AS9121," *NANOG*, no. 34, May 2005.
- [2] J. Karlin, "TNet creates another large hijack," September 2006, <http://cs.unm.edu/~karlinjf/IAR/phpBB2/viewtopic.php?t=14>.
- [3] —, "A fun hijack," June 2006, <http://merit.edu/mail.archives/nanog/2006-06/msg00082.html>.
- [4] —, "AS 8437 announced a quarter of the net for half of an hour," August 2006, <http://www.merit.edu/mail.archives/nanog/msg01700.html>.
- [5] J. Ng, "Extensions to BGP to support secure origin BGP (soBGP)," *Internet Draft draft-ng-sobgp-bgp-extensions-02*, April 2004.
- [6] S. Kent, C. Lynn, and K. Seo, "Secure border gateway protocol," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 4, pp. 582–592, 2000.
- [7] J. Karlin, S. Forrest, and J. Rexford, "Pretty good BGP: Improving BGP by cautiously adopting routes," *International Conference on Network Protocols (to appear)*, November 2006.
- [8] Internet Alert Registry, <http://cs.unm.edu/~karlinjf/IAR/>.

¹New origin ASs for a prefix would be depreferenced for 24 hours if a trusted origin is not on the AS path. New sub-prefixes would be ignored for 24 hours if the less specific's origin AS is not on the path. Paths that involve new edges would be depreferenced for 24 hours.