A Polytheistic Approach to Secure Interdomain Routing

Ratul Mahajan
Microsoft Research

Workshop on Internet Routing Evolution and Design (2006)

The large body of research on interdomain routing security has had little real-world impact despite the immense importance of the subject.[1] I observe that the community seems to be in search of "the one" – a perfect solution in terms of security, deployment costs, deployment incentives, and transition complexity. This perfect solution is expected to be attractive to all ISPs and fit the needs of all ISPs. Many such "monotheistic" solutions have been proposed by various researchers, and new ones continue to exhaust the remaining design possibilities.

I argue that no such perfect solution exists. The security mechanisms required at large tier-1 ISPs are inherently different from those required at small stub networks, which are in turn different from those required at medium-sized ISPs. If we "stay the course," interdomain routing is likely to stay insecure. My position is that, instead of looking for one solution for the entire Internet, we should invent a broad range of solutions. While some solutions should be meant for independent deployment by individual ISPs, others should be meant for joint deployment by groups of ISPs. Individual ISPs will be free to select from this set one or more solutions that suit them best. Some ISPs may choose none of them and continue to operate like today.

How can this "polytheistic" approach to security, which is composed of independent solutions and uncoordinated deployment, be secure as a whole? Solution to this dilemma can be had from observing how the road network operates today. Drivers on the road have different cars, driving styles and skill levels, yet the network seems reasonably secure (at least in much of the developed world). The security of the road network stems from two factors. The first is *visibility:* for most accidents, it is clear who – which car and which driver – is responsible. The second is *financial disincentive:* being responsible for causing accidents leads to financial losses in the form of fines and higher insurance premiums. These factors lead to manufacturers producing safe cars and drivers driving safely and maintaining their cars in good condition.

The security of interdomain routing can be boosted by engineering these two factors into it. Visibility in this context implies that we be able to identify the sources and propagators of bad routing updates. This should not be hard to engineer in the Internet. Even today, the culprit of most routing incidents is well-known. Recent work has explored systematically providing this capability based on routing updates observed at multiple vantage points. While this research has focused on mechanisms that do not rely on network support, with network support this task should be even easier and provide precise answers.

[1] For the purposes of this note, interdomain routing is considered secure when it is not threatened by ISPs sourcing or propagating unauthorized routing updates. Protocols such as S-BGP aim for this level of security.

Financial disincentives should be created for those who source or propagate bad routing updates. One way to do so is by building on bilateral ISP contracts. ISPs should demand in their contracts that their neighbors not send them bad routing updates. Failure to comply carries financial penalties. For instance, customers can withhold part of the transit fee if their providers send bad routing updates and providers can charge more if a customer sends a bad update. Similarly, a peer ISP would need to compensate the other for sending bad updates. Like the road network, there is a possible role for insurance companies here to help ISPs deal with the financial uncertainty associated with security incidents they may accidentally cause or suffer.

In creating financial disincentives, an important difference between the road network and the proposed framework concerns the need for government regulation. The road network relies on regulation to create effective financial disincentives, but regulation is not necessary in the Internet routing context because they can be created using ISP contracts. Regulation in the road network protects the average driver who does not have contracts with other drivers on the road. Not requiring regulation to secure interdomain routing is a significant blessing because national autonomy issues make Internet-wide regulation a highly difficult issue.

With appropriate visibility mechanisms and financial disincentives in place, ISPs have direct and explicit incentives to secure their networks. There is an incentive to guard against the generation of bad updates because that leads to compensating the neighbors (or paying higher insurance premiums). ISPs also have an incentive to guard against propagating a bad update from a neighbor, especially if the compensation it receives from that neighbor is lower than what it would have to give to its other neighbors for propagating that route. In such a setting, the cost of interconnecting to an insecure network may outweigh the benefits, leading to the isolation of networks with persistent security problems. This is akin to the revocation of driving privileges.

There is now no need for ISPs to adhere to a common security solution. Instead, ISPs will select what is best for them. While a large tier-1 ISP will use appropriate mechanisms to guard against both generating and propagating bad updates, a small, simple stub network (that is confident of never generating a bad update) may chose to not deploy any security solution at all. In my opinion, this polytheistic approach is more likely to succeed at boosting the security of interdomain routing, a task where monotheistic approaches have failed in the past.

Before I conclude, I point out that, by being inspired by the road network, the proposed framework shares its security properties. It does not provide guaranteed protection to even ISPs that do their best to run a secure network. Such ISPs can still be hurt (and then compensated) by the mistakes of other ISPs, for instance, traffic intended for them may be hijacked. This dependence exists in the road network as well, where even the safest of drivers can be hurt by other careless drivers. As another similarity, the proposed framework provides to ISPs an incentive to only secure against incidents that are likely to be detected and reported; many smaller incidents may go unnoticed. Many minor violations go unnoticed and unpunished in the road network as well. This is in fact not undesirable from an economic perspective: resource should be spent primarily on preventing incidents that matter.