RPKI vs ROVER
Comparing the Risks of BGP Security Solutions

Aanchal Malhotra and Sharon Goldberg

Boston University,
aanchal4@bu.edu,goldbe@cs.bu.edu

1 Introduction

BGP, the Internet's interdomain routing protocol, lacks built-in security mechanisms to defend against unintentional misconfigurations or deliberate attacks. To defend against these attacks, the past few years have seen the deployment of the Resource Public Key Infrastructure (RPKI), a new infrastructure based on a hierarchy of authorities that provide trusted information about the interdomain routing system. RPKI is designed in a threat model where its authorities are trusted and routing is under attack. Recently, however, [2] showed that the RPKI introduces several new vulnerabilities when this threat model is flipped; namely, when RPKI authorities are misconfigured, compromised, or under attack. Moreover, the deployment of RPKI involves the development of a completely new security infrastructure (new servers, local caches, and protocols). Considering these shortcomings, and given the critical nature of the issue at hand, it is worthwhile to consider alternatives, like Route Origin Verification (ROVER), defined in [4, 6].

ROVER is a new proposal for securing routing that is based on the existing reverse DNS. While ROVER also relies on a hierarchy of trusted authorities, its designers claim that it operates in a "fail-safe" mode, where "[o]ne could completely unplug a router verification application at any time and Internet routing would continue to work just as it does today". There has been debate in Internet community mailing lists [1] about the pros and cons of both approaches. This poster therefore compares the impact of ROVER failures to those of the RPKI, in a threat model that covers misconfigurations and compromised RPKI/ROVER authorities.

2 ROVER and RPKI Primer

The RPKI and ROVER each provide a mapping from an IP prefix[1] to the Autonomous System(s) (ASes) authorized to originate (*i.e.,* claim to be the destination for) the prefixes in BGP. Both therefore protect against *prefix* and *subprefix hijacks*, where the hijacking AS originates an unauthorized IP prefix. This causes the network traffic destined for the hijacked IP prefix to flow to the hijacker's AS instead.

ROVER. ROVER leverages the existing reverse DNS to indicate which ASes are authorized to originate an IP prefix; data in the reverse zone is authenticated by standard DNSSEC signatures. ROVER adds the following to the reverse DNS: (a) A naming convention for IP prefixes of arbitrary length, and (b) two types of resource records: a Route Lock (RLOC) to opt-in or -out of ROVER, and a Secure Route Origin (SRO) to map IP prefixes in the zone to the ASes authorized to originate them in BGP. Zones in reverse DNS correspond to IP prefixes. A zone is *authoritative* for its IP

[1] IP prefix 8.0.0.0/8 has *length* 8 and *covers* 8.8.8.0/24.

prefix and subprefixes, *except* for subprefixes delegated to descendant zones (using standard DNS delegation). SRO and RLOC records are validated in the usual DNSSEC manner, up through the *validation chain* of DNSSEC records to the root of the DNSSEC hierarchy. An SRO or RLOC is *valid* if it passes DNSSEC validation.

Figure 1 shows a sample ROVER hierarchy [2]. The root for IPv4 addresses is in-addr.arpa, maintained by IANA. The root delegates 8.0.0.0/8 to a reverse zone maintained by Level3, which delegates 8.8.8.0/24 to Google and 8.34.114.0/24 to Metro Net. 8.34.114.0/24 has opted-out of ROVER; it does not have an RLOC. Meanwhile, the 8.0.0.0/8 and 8.8.8.0/24 zones each have an RLOC, and thus have opted-in to ROVER. A *BGP route* is an IP prefix π and the AS a originating it in BGP. A BGP route with π covered by 8.8.8.0/24 must have a valid SRO matching (π, a) in Google's zone. BGP routes covered by 8.0.0.0/8 *except* those covered by 8.34.114.0/24 and 8.8.8.0/24 also need matching valid SROs in Level3's zone.

RPKI. RPKI has a similar hierarchy of authorities. Each authority has a *Resource Certificate (RC)*, signed by its parent, containing its allocated IP address space and cryptographic public key. An RC can sign (a) other RCs to suballocate address space, or (b) *Route Origin Authorizations (ROAs)* (the equivalent of SROs in ROVER) to authorize an AS to originate a (sub)prefix of its address space in BGP. In the RPKI, an Regional Internet Registry (*e.g.,* the American Registry of Internet Addresses (ARIN)) would issue an RC that delegates 8.0.0.0/8 to Level3. Level3 would then issue an RC subdelegating 8.8.8.0/24 to Google and and RC delegating 8.34.114.0/24 to Metro Net. Level3's RC would issue ROAs mapping AS X to prefix Y, and AS z to prefix Y. Google's RC would similarly issue a ROA mapping 8.8.8.0/24 to AS 15169.

Validity states. ROAs from the RPKI, or SROs and RLOCs from ROVER, determine the validity state of routes a router learns in BGP. BGP Route (π, a) is *valid* if it has a matching valid SRO (in ROVER) or ROA (in the RPKI). There also are two other validity states, and the RPKI and ROVER define them differently:

- *Unknown.* In the RPKI, BGP route (π, a) is *unknown* when there is no valid *covering ROA*; a covering ROA is any ROA for a prefix that covers π. In ROVER, (π, a) is *unknown* if (a) both the RLOC for the zone that is authoritative for π, and the SRO for (π, a), are absent, or (2) any resource record fails to pass DNSSEC validation.
- *Invalid.* In RPKI, BGP route (π, a) is*invalid* when it is neither *unknown* or *valid*. In ROVER, (π, a) is *invalid* if (1) there is a valid RLOC for the authoritative zone for π, AND (2) there is no valid matching SRO for (π, a).

Routing policies. A BGP router uses its own *local policies* to decide whether to discard, or assign lower preference to, *invalid* or *unknown* BGP routes. As discussed in [2], a router that discards *invalid*/*unknown* routes has the best possible protection against attacks on BGP, but can lose connectivity to routes that become *invalid*/*unknown* as a result of misbehavior by RPKI (or ROVER) authorities.

Local caches. A router is expected to issue a ROVER query for *every* new route it learns in BGP and is stored in the local cache. Meanwhile, RPKI objects are stored in public repositories.

[2] 8.0.0.0/8 is directly allocated by IANA. For non-legacy prefixes, RIRs are also a part of reverse zone hierarchy

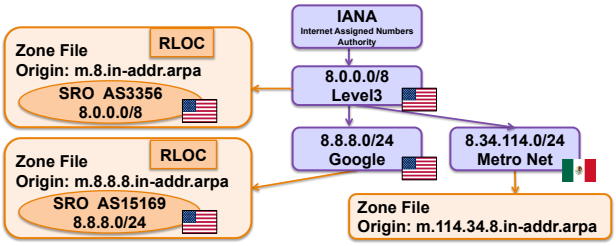

Fig. 1. Sample ROVER hierarchy

3 Attacking Paths to ROVER/RPKI

We consider an attacker that disrupts the communication path between a relying party and an RPKI/ROVER server (dropping and corrupting valid records, or injecting ill-formed records) but *cannot* forge signatures of RPKI/ROVER authorities. How does this impact the validity status of BGP routes? To answer this question, we assume the most plausible set of routing policies [2]: that a router discards *invalid* routes, and prefers *valid* routes over *unknown* routes.

Blackholing BGP routes. The assumed routing policies imply that a BGP route can be blackholed if the attacker can change its validity state from *valid* to *invalid*. [2] shows that an on-path attacker can do just that, by disrupting the delivery of ROAs from an RPKI repository during an ISP's daily update of its local cache. With ROVER, however, this is not possible; because rules for route validity status mean that an attacker that disrupts the delivery of the response (so that it *e.g.,* fails DNSSEC validation) can only cause the BGP route to become *unknown*.

Circumventing ROVER/RPKI protection of BGP. Another concern is an attacker that launches a *sub*prefix hijack on BGP, and then covers its tracks by tampering with communications with the RPKI/ROVER so that *invalid* hijacked route (that should be dropped) becomes *unknown* (and selected, since there is no *valid* route for the subprefix). This is an especially important threat, since this attacker has circumvented ROVER/RPKI's protection of BGP.

In ROVER, however, an on-path attacker can do just that. If a router issues a ROVER query for a subprefix that was hijacked in BGP, an on-path attacker can disrupt delivery of the ROVER response (so that it fails DNSSEC validation) and the *invalid* hijacked route becomes *unknown*. This highly-targeted attack affects only the hijacked route; moreover, advanced attacks on DNSSEC [3,5] mean that there is a risk that off-path attackers could launch such attacks.

Meanwhile, [2] shows that attacks of this form are less targeted with the RPKI. To transition a route (π, a) from *invalid* to *unknown*, an on-path RPKI attacker needs to disrupt the delivery of *all* the ROAs that cover prefix π. However, this would cause all routes authorized by these ROAs to go from *valid* to *unknown*, making the attack easier to detect. Moreover, the RPKI uses *manifests* to indicate which objects are stored in each repository; disrupting delivery of objects or manifests should also raise alarms.

4 Misconfigurations

We consider how misconfigurations of ROVER or RPKI authorities impact BGP route validity. [2] showed that an authority that misconfigures its ROAs can cause all routes covered by the misconfigured ROA to become *invalid*; this also includes routes for prefixes that were subdelegated to descendant authorities. ROVER avoids this. In ROVER, any ill-formed (badly signed), missing,

or mismatching SRO/RLOC record in a zone can only impact routes within the zone; the scope of the RLOC ensures that a misconfigured RLOC/SRO has no impact on other zones. However, a misconfiguration at an ancestor zone (Level3) that breaks the DNSSEC validation chain to a descendant's (Google's) SRO/RLOC records, can impact routes in the descendant zone; however, the worst that can happen is that they all become *unknown*, because of the rules for route validity status.

5 Takedowns by an Ancestor

ROVER allows an ancestor to execute very targeted blackhole attacks on BGP routes. A ROVER resolver has no way of tracking the zones delegated by an authority. Thus, an ancestor that wants a *valid* route (π, a) authorized by its descendant zone to become *invalid* (and thus unreachable in BGP), can simply respond to relevant ROVER queries with a mismatched but valid SRO (*i.e.,* mapping prefix π to some other AS a'); this attack will not impact other BGP routes. Meanwhile, [2] shows that it is much harder for a malicious RPKI authority to cause a BGP route to become *invalid* when it matches a ROA issued by a descendant RC. Attacking a specific BGP route (without harming other BGP routes as collateral damage) sometimes requires the RPKI ancestor authority to issue suspicious new RPKI objects, that could make its actions easier to detect.

6 Summary & Open Questions

ROVER provides several nice "fail-safe" characteristics. But our analysis suggest that these characteristics can be exploited, in a targeted manner, by (1) attackers that disrupt ROVER queries in order to hijack BGP routes, and (2) ROVER authorities that want to blackhole BGP routes authorized by their descendants. But there are various controversial issues that we have not addressed. For instance, [2] shows that circular dependencies are present between RPKI and BGP; are they present between ROVER and BGP? Does RPKI's approach of downloading repositories wholesale, provide better performance than the DNS point queries used in ROVER? Many questions remain open.

References

1. R. Austein. "Re: rpki vs. secure dns?", msg18. seclists NANOG Archive, June 2012. `http://seclists.org/nanog/2012/Jun/18`.
2. D. Cooper, E. Heilman, K. Brogle, L. Reyzin, and S. Goldberg. On the risk of misbehaving rpki authorities. In *HotNets*, page 16, 2013.
3. H. Duan, N. Weaver, Z. Zhao, M. Hu, J. Liang, J. Jiang, K. Li, and V. Paxson. Hold-on: Protecting against on-path dns poisoning. *SATIN*, 2012.
4. J. Gersch and D. Massey. Rover: Route origin verification using dns. In *ICCCN*, pages 1–9, 2013.
5. A. Herzberg and H. Shulman. Fragmentation considered poisonous. In *IEEE CNS*, pages 224–232, 2013.
6. C. Olschanowsky J. Gersch, D. Massey and L. Zhang, editors. *DNS Resource Records for Authorized Routing Information.* IETF Internet-Draft, February 2013. `http://tools.ietf.org/html/draft-gersch-grow-revdns-bgp-02`.