Report on 2nd Operator Roundtable on Routing Security

13-14 December 2011, NH Grand Hotel Krasnapolsky, Amsterdam

Introduction

The Internet Society Technology Roundtable Series are recurring events on various topics that the Internet Society organizes to facilitate the development and adoption of technologies that are beneficial for the Internet as a whole. These roundtables offer a venue for the open discussion of technology issues among concerned parties to facilitate collaboration in ways that might not be feasible in other venues.

The focus of this one-and-a-half day roundtable was routing security. The purpose of the roundtable was to provide a venue for the open exchange of information on operational and technical issues, share experience, do a gap analysis and identify possible incremental steps in improving routing security, while maintaining the big picture. This is a continuation of the dialog the Internet Society began in 2009 in order to better understand from the operators what their needs and stumbling blocks are and where ISOC can be helpful.

Welcome and goal setting

The meeting began with a round of introductions followed by a short session of setting the context for the discussions. A few objectives of the meeting were identified, including reaching common understanding of risks and benefits, a common vision and possible future focus areas. The attendees were asked to focus on more immediate problems and steps, while keeping long-term goals in mind, to look at the issues from a risk perspective and think of things we need to have in place to make incremental further steps.

Routing security - Where are we? Status update

Benno Overeinder presented results of stocktaking of inter-domain routing security produced by NlnetLabs. The stocktaking consisted of an online survey, running from February to March 2010, and a series of one-on-one interviews with network engineers/experts. One of the key findings is the almost omnipresent deployment of session security techniques like MD5, BGP TTL hack, or BCP 38 (anti-spoofing). However, only 45% of the respondents saw an improvement, while 45% did not saw any improvement, and another 10% considered it even counter productive. Prefix, AS path, and max prefix filtering are also almost everywhere applied (from the survey - 87%). Filtering is seen as a very effective way of securing the routing infrastructure: 80% do see improvement, 17% no improvement and 3% - counter productive. From the interviews, it became clear that the first concern with network operators is network stability. Stability and reachability is where money can be earned. Security is only an issue when stability and reachability are put at risk. Also, with

the deployment of routing security methods, there is an asymmetric cost–benefit relation. With investments in inter-domain routing security, the investing organization does not benefit but other networks do (the tragedy of the commons, but reverse).

The level of routing security awareness relates to the size of the network: large networks have more NOC staff with security expertise. Most incidents seen in inter-domain routing are mistakes, “fat fingering”, although large attacks are not spoken about in public. With RPKI deployment eminent, there are still some concerns with network operators, among which the single authoritative trust anchor, the costs of certificates, and instability and/or vulnerability of the RPKI infrastructure. In the RIPE region there was quite some discussion on this topic, resulting in a voting on the continuation of RPKI and its use in routing. Although the majority was in favor of the continuation of development of the RPKI infrastructure and service, there was significant percentage of those who were against.

Benno’s presentation was followed by the attendees sharing their experience and practices in dealing with routing security issues in their networks, as well as thoughts about routing security on general. The state-of-the-art approach is strict filtering of customers, use of RPF, bogon filtering, more relaxed policy with peers and transit (maxprefix).

Many providers are actively using IRRs for building filters as part of their provisioning system, some of the providers also use it for building filters for peers (along with other available information), but it is hard, since information is incomplete. Provides are using BGPmon (http://www.bgpmon.net), Cyclops (http://cyclops.cs.ucla.edu) and IS Alarms, formerly known as MyASN (http://www.ripe.net/is/alarms/). Some operators reported about their plans to begin testing RPKI-based solutions.

Many operators mentioned that they were deploying best current practices within the constraints of the operations budget: it was difficult to justify a separate budget line for routing security. In some cases these solutions are reviewed from a threat perspective, but a formal risk analysis rarely takes place.

It was noted that software bugs (*e.g.* malformed BGP attributes) and environmental factors sometimes present bigger risks to the stability of the network than, say prefix hijacking. Several attendees noted that more testing and better handling of erroneous conditions were required (*e.g.* a full session reset in cases of malformed BGP attributes).

Positive and Negative Factors affecting progress (Challenges)

This session began with reflections on the last roundtable on routing security organized by the Internet Society back in 2009. Many issues and challenges identified back then seem to be still relevant: lack of accurate and complete data, difficult business case, lack of tools, operational practices are customer checks, rarely peers. Short-term needs and priorities (2-4 years) identified at that roundtable included origin protection (non-conflicting validation and resolution of revocation), open software tools, path protection without protocol changes, as well as clean IPv6 data. The long-term ones (3-6 years) were: path protection, hardware changes, clean IPv4 data.

While good progress has been made in the IETF SIDR working group, a lack of progress in the network operations in this area was obvious.

The group engaged in the discussion of challenges and stumbling blocks that prevent the progress.

It was noted that doing things specifically for security benefits is very difficult for lack of a business case. Most of the attendees indicated network stability and performance as the priority, so if security is coming as a by-product while solving these paramount issues (*e.g.* by deploying new platforms or applications) that gives it best chances to get deployed as well. In some cases a certain business case can be found, for instance, RPKI may simplify the ISP provisioning system and improving routing security may eliminate the need for deaggregation, when it is used to protect against prefix hijacking.

Also, because securing routing is an end-to-end problem, an operator can only have control over a very small part of it. Because it is end-to-end one needs to rely on others, which makes it hard to establish a decent security policy.

Attendees agreed that availability of accurate and complete data is crucial for improvements in routing security. Another important building block is tools that allow a uniform way of publishing, collecting, processing this information and applying it to ISP's routing policy. Both are lacking at the moment. Lack of standards and standard practices in the area of configuration was also mentioned.

It was noted that factual data about routing security incidents is missing: level of instability caused by BGP vulnerabilities is speculative and there is no solid analysis. Such analysis is challenging, because the BGP data itself doesn't inform about the intent, but at the same time some approximation can serve two main purposes: provide a trend analysis (are we getting better or worse?), and be a good foundation for risk management (also at the management level). Some suggestions included signature-based or learning-based anomaly detection within the BGP table by establishing a baseline and then flagging for something that is out of sync with it.

The group spent some time in discussing the RPKI technology. Some attendees observed that RPKI brings some of the missing qualities - trustworthy data, a standard format, storage and access methods, although by itself the RPKI doesn't make the data complete. Some of the noted challenges in the deployment of RPKI by the operators included lack of understanding how RPKI interferes with policies (or, more generally, how RPKI data could be applied), changes with regards to ownership and responsibility - making routing infrastructure beholden to systems infrastructure teams. There were more general concerns whether the RPKI is the right trust model for this data. It was also observed that RPKI data alone is not enough - it doesn't reflect ISP relationships: policy/intent, how ones traffic is going to be handled by a peer (peer's policy).

Some of the attendees voiced an opinion that BGP4 doesn't meet a broad array of demands by todays and future applications: voice and video, named data networking, etc., and security is just one of them. Trying to meet these requirements by adding features to BGP4 will have an opposite effect with regards to its scalability, resilience, performance and costs. To address these requirements in the longer term a more comprehensive solution might be needed – a next version of BGP. The group did not explore this area further.

A separate session was dedicated to looking at the existing tools and infrastructure, specifically into the IRRs.

Eric Osterweil delivered a talk on IRR Efficacy that was intended to outline the role that Internet Routing Registries (IRRs) historically used to play in network operations, and what role they may be able to play today. In addition, the talk focused a lot on the details of what limited their utility in the past (and led to the decline in their use), and what has changed that re-enables this today, and if the roundtable attendees felt that they could reemerge in operations in today's Internet. There was a subtle notion that a lot of the security concerns that people are addressing in new systems work was already addressed through proper use of IRRs (modulo the limitations they previously faced).

The talk broadly classified each of the set of problems that IRRs faced as being in one of the following general categories: accuracy and integrity of IRR data, operation of IRR infrastructure, historical BGP and hardware constraints, security considerations, and policy and privacy Issues. The talk then went on to discuss how all of the ``major'' limitations (with one notable exception) in each of these categories have been addressed and overcome in the last 10 years. The major missing piece that still remains is an operationally deployed resource certification framework. The talk gave a rough outline of requirements that a resource certification infrastructure should meet.

During the presentation, the discussion spent a lot of time refining the specifics in the IRR presentation. Then, at the end, the audience discussed a number of issues about how IRR data could be salvaged (considering issues like rot), and if a reboot were to happen, could it be useful, etc. Ultimately, most of the audience seemed to agree that many of the reasons that led to the demise of the IRR system may not exist today, and a fresh look and some consideration might be warranted.

Main takeaways and focus areas

Correct and complete data is very important

Independent of technology used routing policy and decisions should be based on data that includes:

- assertions about address allocations (and ASN assignments) holdership
- assertions about routing policy related to a particular prefix

Usefulness of these data depends on how correct, complete and authoritative they are. Different trust models may be considered (*e.g.* a web of trust, a PKI), but there is a need for some sort of resource certification. RPKI is a step forward, the infrastructure and the ability to requests certificates and create ROAs is offered by some RIRs as a pilot service, while others are further along in delivering services that can be operationalized.

It was noted that the IRR/RPSL is still the only standard way to express and share ISP's routing policy and that the utility of the IRR can be increased by integrating it with some form of resource certification framework.

Focus on stability and performance

Network stability, resilience and performance are the main priorities for the operators. Requirements for routing security are somewhat embedded in these goals, but if considered alone it is a hard sell – the costs are significant and not all benefits directly contribute to the objectives. Risks are considered low, but may go up in the future. The most effective way to deploy routing security measures is as a by-product of improvements in stability, resilience and performance.

Lack of tools and common practices

Attendees observed that there is lack of agreed uniform way to define, publish and process providers routing policy. Everyone is designing their own approach, and although some building technology blocks exist (RPKI/ROA, IRR/RPSL) this also results in absence of open flexible toolsets. This creates a relatively high threshold to the deployment of the routing security measures.

Factual data about incidents is missing

Although BGP threats and vulnerabilities are well known, level of instability caused by BGP vulnerabilities is speculative and limited to high-profile cases. There is lack of solid analysis and statistical data. Improvements in this area can provide a better foundation for risk analysis and inform the corporate decision making process. These data can also provide monitoring of the situation – are we getting better or worse, - and facilitate trend analysis.

BGP fragility and missing features

Besides BGP vulnerabilities that are being addressed in the IETF SIDR WG, there is a separate class of vulnerabilities related to implementations, where malformed BGP attributes may cause significant problems. The group's opinion was that there needed to be additional focus on testing and probing individual BGP implementations to identify bugs and vulnerabilities. This could be a combination of security research within academia, the white hat community, router vendor internal testing and carrier certification testing.

Some attendees mentioned that inter-peer error reporting is insufficient (for more details see draft-ietf-grow-ops-reqs-for-bgp-error-handling).

Some of the attendees voiced an opinion that in order to support future applications on the Internet a more comprehensive solution might be needed – a next, "clean slate" version of an inter-domain routing protocol.

Attendees

Nina Bargisen, TDC
David Freedman, ClaraNet
Wes George, TWC
Jac Kloots, Surfnet
Olaf Kolkman, NLnetLabs
Peter Lötberg, DT
Jared Mauch, Neustar
Arnold Nipper, DE-CIX
Eric Osterweil, Verisign
Benno Overeinder, NLnetLabs
David Roy, Orange/FT
Ruediger Volk, DT

Leslies Daigle, Internet Society
Matthew Ford, Internet Society
Phil Roberts, Internet Society
Andrei Robachevsky, Internet Society

Internet Society
Galerie Jean-Malbuisson 15
CH-1204 Geneva, Switzerland
Tel: +41 22 807 1444
Fax: +41 22 807 1445
http://www.isoc.org

1775 Wiehle Ave. Suite 201
Reston, VA 20190, USA
Tel: +1 703 439 2120
Fax: +1 703 326 9881
Email: info@isoc.org