Internet Routing Validation

John “JI” Ioannidis

ji@research.att.com

AT&T Labs – Research

Joint work with:

Bill Aiello, Steve Bellovin, Matt Blaze, Tim Griffin,

Howard Karloff, Patrick McDaniel (AT&T),

Geoffrey Goodell (Harvard),

Avi Rubin (Johns Hopkins University)

BGP Listens to Many Peers

- Each router receives multiple announcements for each destination.
- But: announcements are not authenticated.
  - We don’t even always know who is allowed to advertise a prefix!
- Anyone can announce (almost) any prefix.
  - Maliciously.
  - Accidentally.
- Frequent source of problems.
- Best case: more routing data than necessary.
- Usual case: blackholed traffic.
- Extreme case: redirect traffic for intercepting.

BGP Chooses Among Many Paths

- Each router receives multiple announcements for each destination.
- Uses *path attributes* to select the best path.
- But: path attributes are not authenticated either.
- AS changing path attributes can disrupt routing.
  - Cause suboptimal paths to be taken.
    - Or paths where an adversary is listening!
  - Interfere with policy decisions.
  - Cause parts of the network to become unreachable.

BGP Is All About Policy

- The main goal of BGP is to arrive at routes that satisfy policy.
- ISPs need a way of checking that others are abiding by their advertised policies.
- The Internet Routing Registry (IRR) project aims to provide this global policy knowledge.
- But:
  - Private peering agreements are usually confidential.
    - And of no interest to non-participants.
  - Updating the RRs is not done in real time.
    - So the registry does not always reflect current policy.
  - There is still no way to check whether BGP updates **received** abide by the policies of all the **intermediate** ASes.

BGP Hides Information

- When something goes wrong, you are trying to infer what went wrong from what you are seeing in bgp data.
- Having a richer channel to convey that information will allow us to figure out whether what we are seeing is indeed an anomaly or it is according to what should be happening.
- “Root-cause analysis”.

Enter IRV

- **Internet Routing Verification**.
- This is an effort to unhide the information.
- Each AS maintains a [distributed,replicated] Routing Verifier.
- An IRV is a repository for:
  - Current policy.
  - Current routing state.
- Other ASes may query the IRV, subject to access controls.

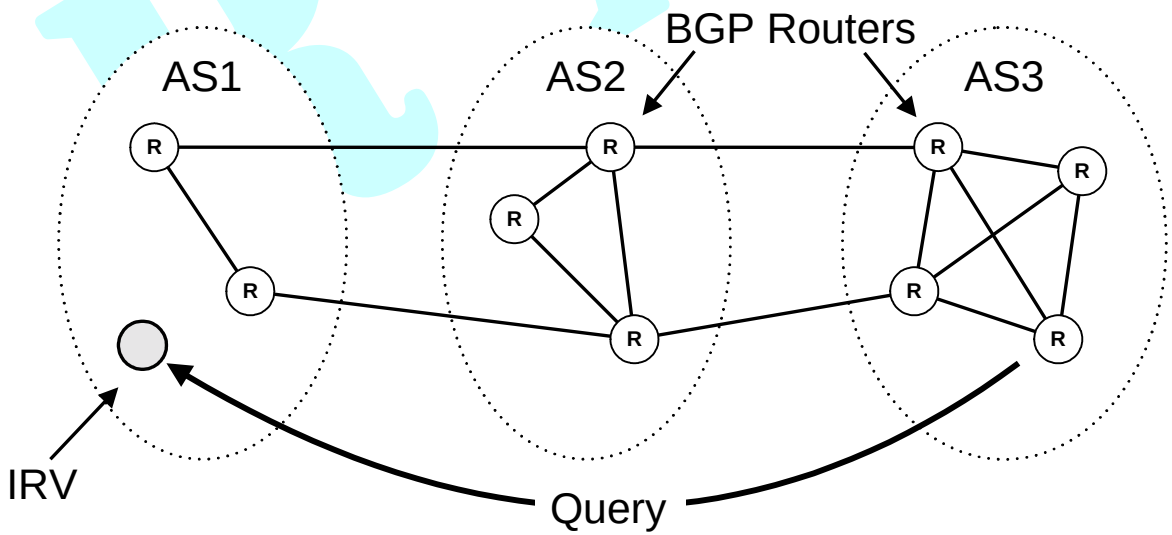

IRV Has Policy Data

- The main function of an IRV is to keep up-to-date policy data.

- Policies may be re-imported from the RRs in RPSL.
- New policies may be written in either RPSL or XML.
- Policies are stored in XML.
  - Schemata are still being worked on.
- Other ASes query the IRV to consult/verify policy information.
  - IRV has a query protocol.
  - Based in Xquery.
- The IRV becomes the canonical repository for an AS's policy information.

IRV Has Current BGP Data

- The IRV keeps peering sessions with all its BGP routers:

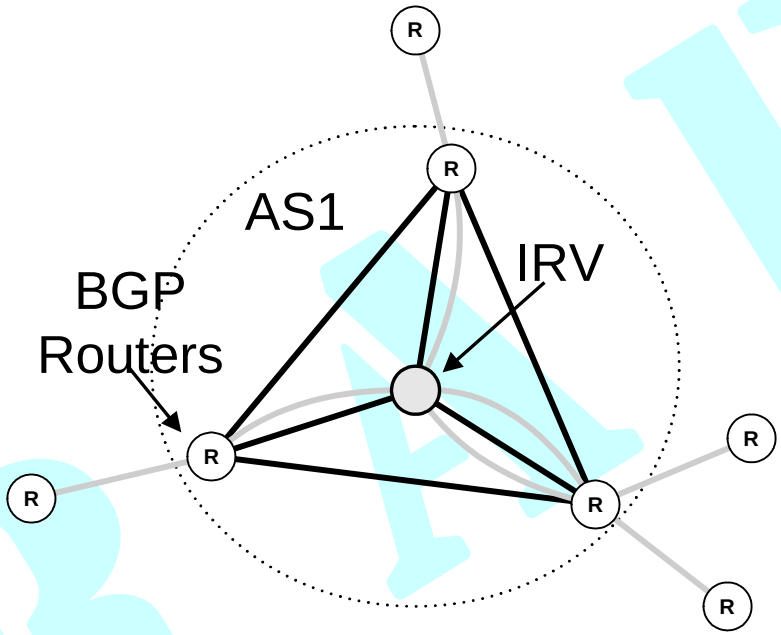

- It maintains all the routes that the AS receives and announces (in XML, queriable subject to access control).
- It can also digest SNMP data from the routers.

IRV Has Configuration Data

- The IRV also maintains router configurations.
- Parts of the configuration can be automatically translated **into** policy.
- Parts of the configurations can be automatically generated **from** policy.

Origin Verification

- Any router can query the IRV responsible for an AS to verify that it is indeed originating a prefix.
  - Subsumes S-BGP Address Attestation.

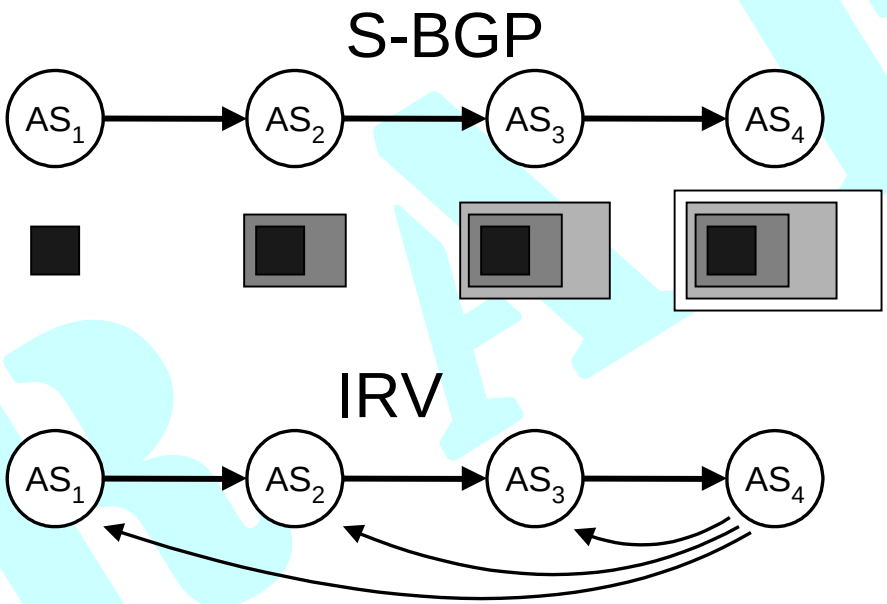

- Even in the absence of rigid public key infrastructure, this can yield benefits.
- ISPs can verify that their announcements are reaching other parts of the net.

Path Verification

- Query IRVs in the ASes listed in the AS_PATH.
- Verify that each AS announced the prefix to the following AS in the AS_PATH.
  - Subsumes S-BGP Route Attestation.
- Verify that the attributes are consistent with policy.
- ISPs can also verify that their announcements are not being corrupted.

Access Control

- IRV data are subject to access control.
  - Private peering information is not available to just everybody.
- Different granularities/levels of access depending on requester.
- Some parts of this could be a service.
  - Public service.
  - For-profit service.
  - I'll-show-you-mine-if-you-show-me-yours service.
- Not a big issue – we pretty much know how to solve it!

Paper at NDSS'03

```
@inproceedings{irv-ndss03,
        author = {Geoffrey Goodell and William Aiello and Timothy Griffin and
                  John Ioannidis and Patrick McDaniel and Aviel Rubin},
        title = {{Working Around BGP: An Incremental Approach to Improving
                  Security and Accuracy of Interdomain Routing}},
        booktitle = {{Symposium on Network and Distributed Systems Security}},
        city = {San Diego, CA},
        month = {February},
        year = 2003,
        url = "http://www.tla.org/papers/irv-ndss03.pdf"
}
```

Joint work with Steve Bellovin, Matt Blaze, Howard Karloff, Fabian Monrose.

Summary

- IRV provides an asynchronous way to verify BGP data against **policy**, **configuration**, and **current routing state**.
- XML-based.
- No router modifications needed.
- Incremental deployment.
- Value increases as more ISPs adopt it.

Future (current!) Work

- Lots of open questions.
- That's why we're here!

- What would it take for ISPs to consider deploying it?
- Interaction with soBGP?
  - We are working on this.
- How well will it scale in practice?
- Will it provide something people need (we hope so)?