Improving Internet Routing Robustness

ftp://engr.ans.net/pub/slides/ndss/feb-1999/

Curtis Villamizar <curtis@ans.net>

- How does Internet routing work?
- How noticeable are routing related outages?
- Why have routing attacks not occurred?
- What measures are being taken today?
- Approaches to improve external routing robustness.

How does Internet (Big-I) routing work?

- The Internet is broken up into Autonomous Systems (AS)
  - An AS is an administrative boundary. Each ISP has 1 or more AS. Major ISPs have 10s of AS.
- Within an AS an Interior Gateway Protocol (IGP) is used.
  - OSPF and IS-IS are used exclusively. No major ISP uses RIP-2 or EIGRP in their backbones.
  - Routers within an AS are under common administration.
- External routes (inter-AS routes) are carried by BGP-4
  - IBGP is used within an AS to carry external routes.
  - EBGP is used to exchange routes among adjacent peers.
  - Inter-AS routing (EBGP) reflects routing policy that is determined by business relationships.

Some characteristics of Internet Routing

1. Physical topology is relatively static.
2. Logical topology (AS boundaries and adjacencies) is relatively static.
3. Routing is highly dynamic.
4. There are just under 4,000 AS numbers in use.
5. The global routing table has about 60,000 routes.
6. Typical days see peaks on the order of thousands of announcements in a fifteen minute measurement interval.
7. Last week (for example) saw a peak interval over 6,000 and numerous adjacent intervals in the 2,000-4,000 range.
8. These characteristics must be considered when examining scaling properties of proposals.

Characteristics of Internet route selection

1. Routing protocols implicitly prefer more specific routes unless explicit configuration per route dictates otherwise.
2. Regardless of all other path attributes, if a more specific route is accepted, traffic will be directed toward that route.
3. If an erroneous route covers critical hosts such as DNS servers or WWW servers, a denial of service can occur.
4. Many providers accept all routes from peers, with minimal filtering but lower BGP LOCAL_PREF (which does not override a more specific route from overriding a less specific).
5. A least one major provider does not filter route announcements from its customers.
6. Misconfiguration by a customer can cause widespread denial of service for a specific prefix if there are no sanity filters.
7. Router software error or radical misconfiguration can cause an outage for a wide range of prefixes.

How noticeable are routing related outages?

- Outages can be widespread and can get wide press coverage.
- For example consider the incident on April 25, 1997.

  http://www.news.com/News/Item/0,4,10083,00.html?latest "Router glitch cuts Net access" By Nick Wingfield, Staff, CNET News.com

  http://www.wired.com/news/technology/story/3442.html "Net Outage: The Oops Heard 'Round the World" by Michael Stutz, Wired

  http://www.merit.edu/ipma/press/death.html Other articles on this are listed on the IPMA "Death of the Internet" page.
- More than a year earlier a similar incident occurred. The Internet seems to average one every few years.
- Smaller incidents are occurring much more frequently.

Why have routing attacks not occurred?

1. In a routing based outage false routing information is injected into the global routing data. Since many sites log routing activity malicious action would be too easily traced.

2. The impact of a routing based attack would be limited to a denial of service.

3. A routing attack in progress can be contained with the installation of a route filter and completely neutralized if the filter is at or near the source.

4. The combination of little effect (short term denial of service only) and high risk (too easily traced) is probably what is preventing any malicious activity.

What measures are being taken today?

1. IGP protocols use peer to peer authentication, usually based on MD5, but sometimes based on simple password. Snooping IGP exchanges is difficult.
2. IBGP typically uses at least MD5 authentication within IBGP itself. TCP/MD5 is also used, specifically to address a potential BGP denial of service (RFC2385).
3. Often no authentication is used over EBGP though some use TCP/MD5 between peers. These are usually switched interfaces and EBGP uses a TTL of 1. Using TCP/MD5 is a better practice.
4. The amount of sanity checking on external route announcements ranges from close to none to only accepting specific prefixes from specific peers.

Approaches to improve external routing robustness.

- Information Storage
  1. DNS - zone transfer on request, expire and refresh timers.
  2. IRR - current: centralized database with full mirrors; moving toward: distributed with exchange of deltas
- Authorization Model
  1. DNS - simple delegation hierarchy, authorization per DNS zone based on signature of zone file
  2. IRR - hierarchies on AS, IP address, and routes relying on both AS and IP address, authorization per object based on hierarchy of maintainer objects
- Verifying Route Announcements
  1. Sanity filters applied to EBGP peers
  2. Signatures on route origination only
  3. Signatures at each BGP exchange

Possible Approaches

- Current Proposals
  1. DNS distribution with origin signatures
  2. DNS distribution with full AS path signatures
  3. IRR distribution with BGP filters on peers
- Worth Considering
  1. IRR distribution with per ISP selection of filters, origin signatures, or full AS path signatures

Distributed Routing Registry

- Repositories do not need to trust each other.
- Repositories agree to common authorization rules and use common authentication methods.
- For any given object there is exactly one repository in which the object can be created and modified unless delegated.
- An object and any beneath it in the hierarchy may be delegated to another repository.
- Transactions are flooded and each repository can recheck the authentication and authorization of incoming transactions.
- Repositories and mirrors will have a complete copy of the set of repositories through processing the deltas to the database.
- A small number of flooding adjacencies are needed.
- Scales according to the rate of change of the database.
- For details, consult RPS WG internet-drafts.

Initial Object Submission and Redistribution

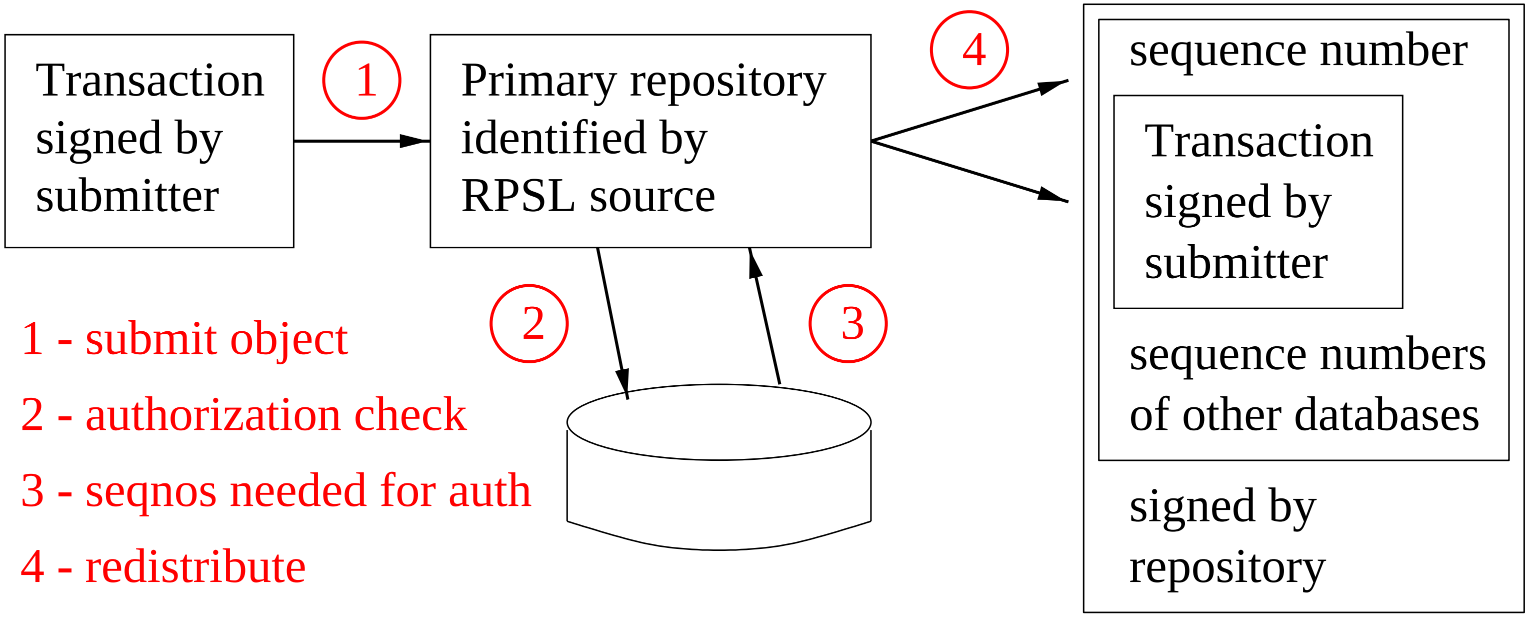

Transaction signed by submitter

Primary repository identified by RPSL source

sequence number

Transaction signed by submitter

sequence numbers of other databases

signed by repository

1 - submit object

2 - authorization check

3 - seqnos needed for auth

4 - redistribute

If the authorization check requires objects from other repositories, then the sequence numbers of the local copies of those databases is required for mirrors to recheck the authorization.

Alternate Initial Object Submission

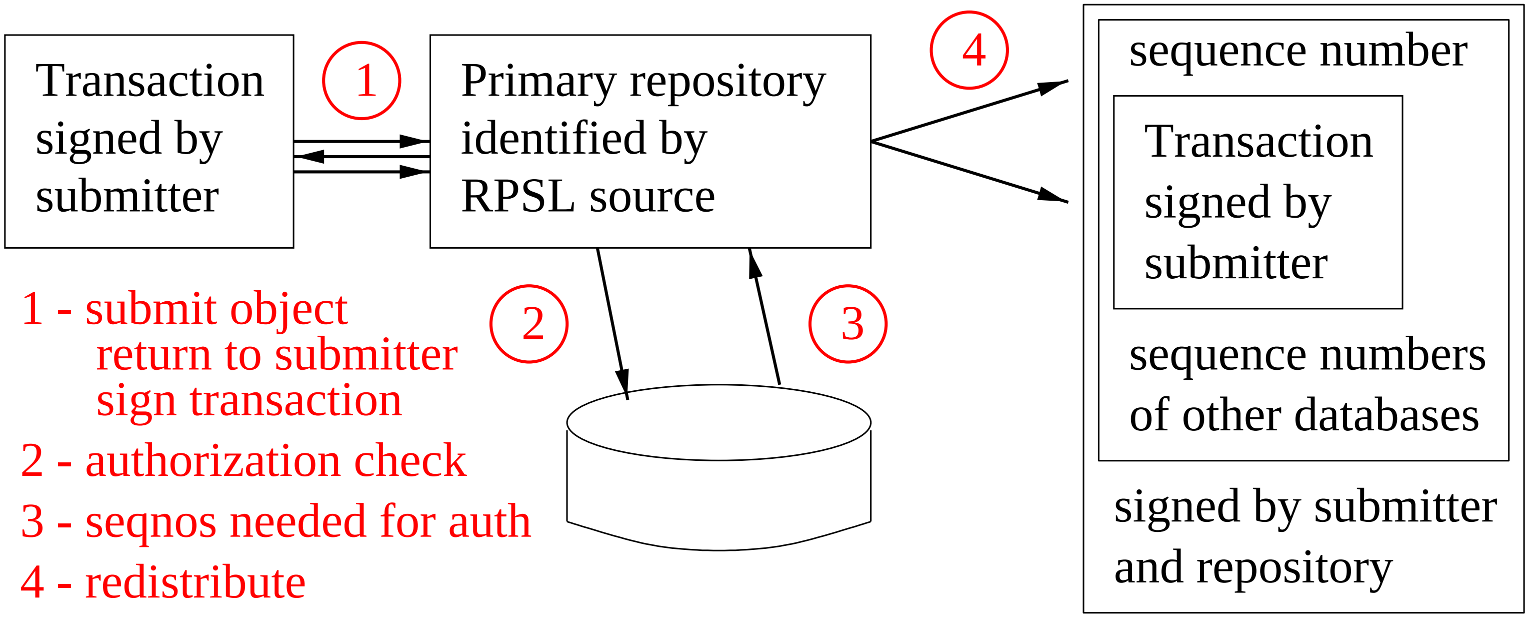

Transaction signed by submitter

Primary repository identified by RPSL source

sequence number

Transaction signed by submitter

sequence numbers of other databases

signed by submitter and repository

1 - submit object
return to submitter
sign transaction

2 - authorization check

3 - seqnos needed for auth

4 - redistribute

Note: the submitter is protected against the possibility of the repository replaying a submission later. This method is not in the cuurent draft.

Further Transaction Redistribution

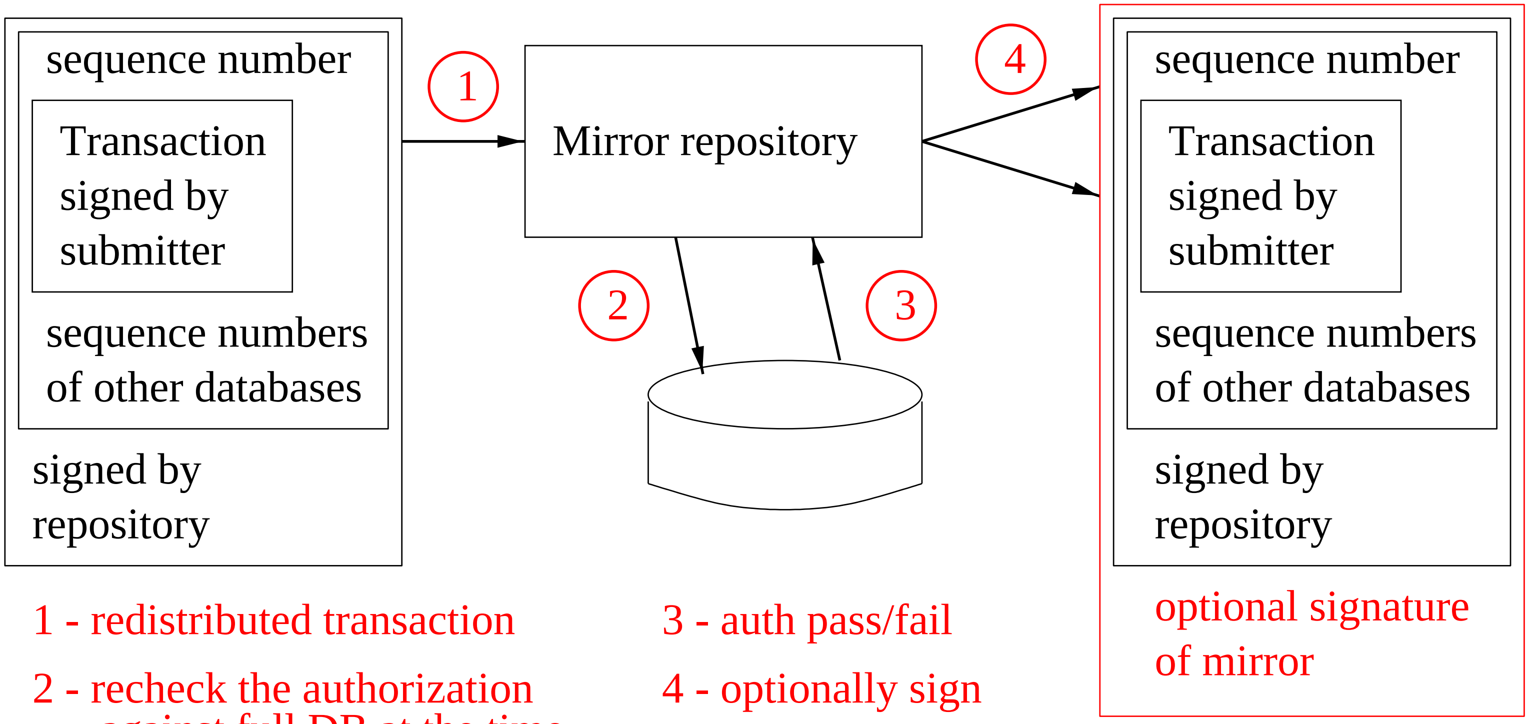

sequence number

Transaction signed by submitter

sequence numbers of other databases

signed by repository

1

Mirror repository

2

3

4

sequence number

Transaction signed by submitter

sequence numbers of other databases

signed by repository

optional signature of mirror

1 - redistributed transaction

2 - recheck the authorization against full DB at the time of transaction using seqnos

3 - auth pass/fail

4 - optionally sign then redistribute

If the authorization check was repeated, the mirror may optionally add a signature before passing the transaction any further.

Redistribution to Lightweight Mirrors

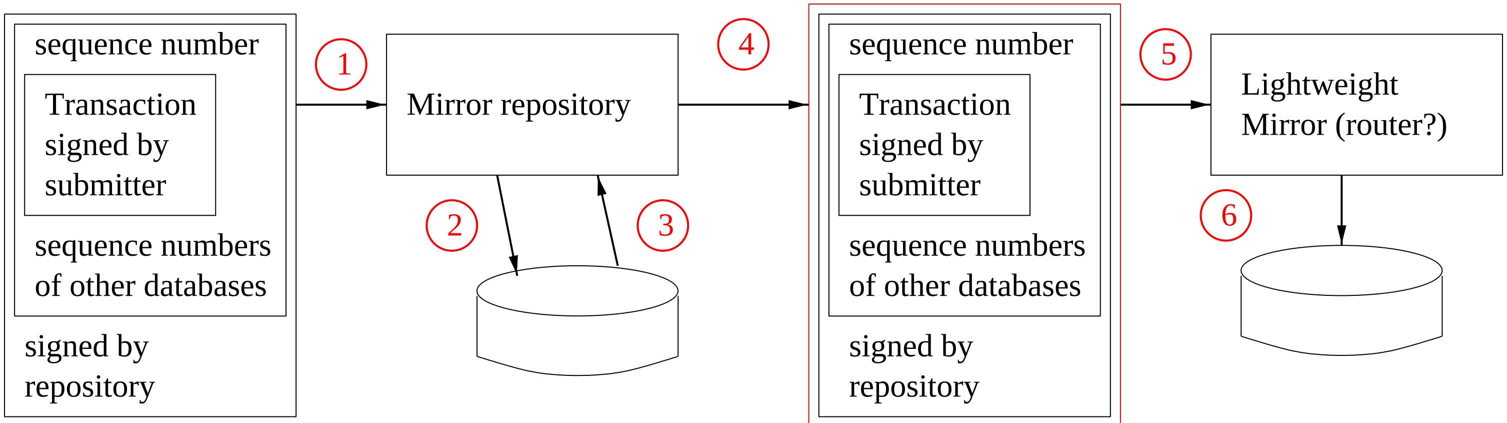

sequence number

Transaction signed by submitter

sequence numbers of other databases

signed by repository

1

Mirror repository

2

3

4

sequence number

Transaction signed by submitter

sequence numbers of other databases

signed by repository

optional signature of mirror

5

Lightweight Mirror (router?)

6

1 - redistributed transaction

2 - recheck the authorization against full DB at the time of transaction using seqnos

3 - auth pass/fail

4 - sign and redistribute

5 - just check mirror signature

6 - apply change with no authorization check

The lightweight mirror must trust the mirror from which it gets a feed. This is a safe assumption if the two are under the same administration (the mirror providing the feed is a host owned by the same ISP who owns the routers). The lightweight mirror simply checks the signature to insure data integrity.

Optional Commit and Confirm

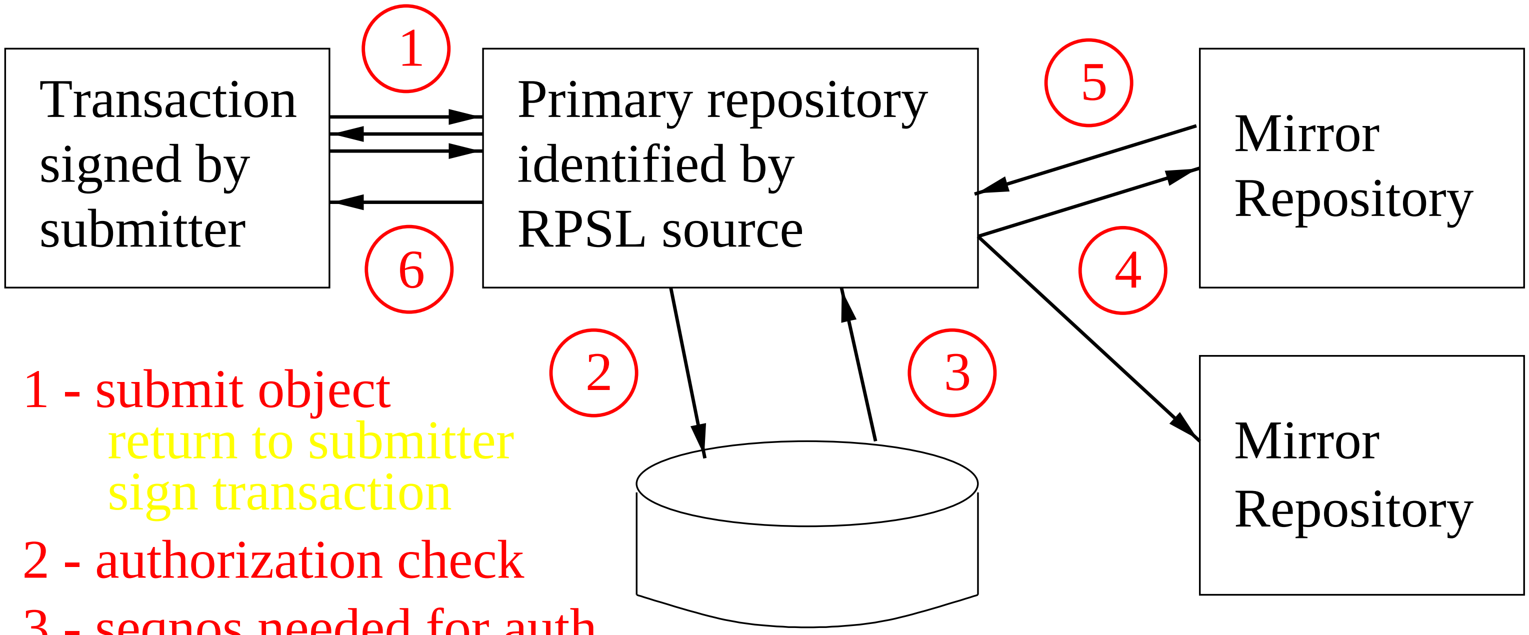
Transaction signed by submitter
Primary repository identified by RPSL source
Mirror Repository
Mirror Repository
1
2
3
4
5
6

1 - submit object
return to submitter
sign transaction

2 - authorization check

3 - seqnos needed for auth

4 - redistribute

5 - mirror confirms

6 - repository confirms

In the event of a disk crash, the repository has already successfully flooded the transaction before sending the confirm back to the submitter. If a mirror is under the same administration, the repository can recover from disk and roll forward the transactions from the mirror before resuming operation.

Signatures on Route Origination Only

1. The BGP route originator signs the BGP route.
2. Determine public key for an IP prefix.
3. When receiving a route the public key for the prefix is determined and the signature is verified.

Full AS Path Signature

1. The BGP route originator signs the BGP route.
2. At each exchange of a BGP route a signature is added by the advertising router indicating what AS the route was advertised to.
3. When receiving a route the public key for the prefix and the public key for each AS must be determined and the signature verified for the originator and each AS in the AS Path.

Filtering Routes from EBGP Peers

- AS numbers and IP addresses are assigned hierarchically.
- Intention to announce a route can be registered with authorization from both the AS and the IP address holder.
- AS registrations may include AS adjacencies and policy.
- All registry changes must follow a set of authorization rules which include authentication requirements (with public keys where authorization specifies cryptographic authentication)
- Complete registry mirrors can (and should) repeat authorization and authentication checks.
- Filters may be constructed from registry information. The type of filter used by ISPs may vary.

Types of Filters with Registry Based Approach

1. As a minimum measure, an ISP can filter their direct customers announcements using a (relatively small) list of prefixes based on registered routes. This prevents becoming the source of a denial of service.
2. As a protection to direct customers, filters may be added which deny more specific routes for any customer routes. These filter list can become quite lengthy.
3. If registration of routes reaches a critical mass, the announcements of routing peers can be limited to specific prefixes with preferences assigned according to stated policies along the path. These filters can also become quite lengthy.
4. Filters may assign preferences to specific AS paths if adjacencies and policy of the complete path is documented in the database. This form of filtering is not yet in use.

Pros and Cons of Origin Signature

- Advantages:
  1. Prevents most or all accidental attacks seen today.
- Disadvantages:
  1. Replay attacks and accidental replay
  2. Scaling wrt number of routes and announcements.
     - One signature verifications per route received.
- Deployment:
  - None.

Pros and Cons of Full Path Signature

- Advantages:
  1. Origination and path is authenticated.
- Disadvantages:
  1. Scaling wrt number of routes and announcements.
     - Two or more signature verifications per route received.
  2. Route aggregation removes signatures of originator of more specifics or increases BGP overhead dramatically.
- Deployment:
  - None in Internet. Elsewhere?

Pros and Cons of Registry and Sanity Filters

- Advantages:
  1. No signature checks per route when a route is received.
  2. Frequency of cryptographic authentication check is logical topology change (database update). Changes to registry information are infrequent (on the order of a few hundred per day) and rarely needs to be reflected in real time.
  3. BGP filtering is implemented in currently shipping routers.
- Disadvantages:
  1. Origin and path is sanity checked but not authenticated.
  2. Filters are expensive in terms of router resources, though less expensive than signature verification.

Deployment of Registry and Sanity Filters

- There are 5 closely cooperating major registries (ANS, CANET, CW, RADB, RIPE) referred to as the IRR.
- There are many private ISP registries and a number of registries wishing to join the IRR.
- Common policy description defined by RFC2280.
- Distributed registry and common authorization and authentication model is specified (IETF drafts) and being implemented. Source will be freely distributed.
- Numerous providers use the IRR or a private registry to configure router filters based on their own routing policy.
- Critical mass has not been reached. Route objects are about 90% populated. Adjacencies and tools necessary to base local policy on policies of entire routed path are not available.

Common Limitations

1. Partial origin deployment yields routes with no originating signature or origination of unregistered routes.
2. Partial transit deployment or cooperation yields incomplete signature chain or AS with no stated policy.
3. Security compromise along the transit path results in denial of service in either case. A complete signature chain in this case provides no assurance of traffic delivery and therefore limited security advantage.

Among Proposals, Which is Better?

- This may be a question of applicability.
  - Full AS path signatures may be preferable for smaller higher security networks who prefer signatures though limited security advantage is offered.
  - Sanity filters may be more applicable for the global Internet where scaling is critical.
- Both types of information, AS adjacency and policy, and public keys, can be held in the routing registry.

Summary

1. Signatures on BGP AS Path offers security advantage over filtering. In the filtering model there is often no assurance that the downstream AS is filtering and filters are not commonly applied against the full AS path. A complete signature chain provides a clear positive indication.
2. Either originator only signature or filtering offer a substantial improvement in routing robustness over doing nothing.
3. Filters offer scalability over the signature techniques that is critical in very large deployments such as the global Internet.
4. It may be that the two/three types of solution are applicable in different situations.
5. All of the information needed for either approach can be distributed using the routing registries with the addition of an optional public key per AS and route.

References

ftp://engr.ans.net/pub/slides/ndss/feb-1999

http://www.ietf.org/

RFC-2280 (RPSL)

draft-ietf-rps-rpsl-v2-01.txt

draft-ietf-rps-auth-02.txt

draft-ietf-rps-dist-01.txt

draft-ietf-rps-dbsec-pgp-authent-00.txt