

Current BGP Security Issues

LACNIC XII

Danny McPherson

danny@arbor.net

Agenda

- Overview
- Transport Connection Protection
- The Real Root of the Problem
- Route Advertisement Validation
- What to do...

Overview

- BGP is the *de facto* protocol for inter-domain routing on the Internet – used to convey destination reachability to peers
 - **prefix** – set of destinations (e.g., 10.0.0.0/8)
 - **attributes** (e.g., AS_PATH, MED, Origin , etc.)
- Large number of loosely interconnected routing domains, represented as **autonomous systems (AS)**, make up global routing system
- The Internet operates in “routing by rumor” mode
 - each AS operates autonomously
 - implicitly trusts what it hears from peers
 - employs that information to reach the destinations reported
 - passes the reachability information on to it’s adjacent peers

Transport Connection Protection

- BGP runs over TCP port 179
 - Therefore susceptible to TCP-based attacks
 - TCP Reset Attacks
 - TCP Hijacking
 - TCP Window Congestion
 - Etc..
- But...
 - BGP is a stateful protocol – usually, you either want to disrupt the session or hijack it
 - Flapping sessions are as good (or better) as down sessions if Denial of Service is goal

Protecting BGP Transport Connections

- BGP was motivation for several TCP protection mechanisms
 - Most popular is TCP MD5 Signature Option
 - MD5 proven weak but this is most widely deployed
 - TCP-AO (authentication option) work happening in IETF to replace MD5 – in TCPM WG
 - Can also run over IPSEC
- BCP for BGP Transport protection today
 1. TCP MD5 Signature Option on all sessions
 2. Infrastructure ACLs
 - Deny everything to *internal* BGP speaker addresses at perimeter
 - Only allow eBGP peers to send packets to adjacent node
 3. Generalized TTL Security Mechanism
 - Accept packets only *if* $TTL > n$, where n is number of hops to peer (usually 1)

The Root of the Problem

- Still today, no formally-verifiable authoritative database exists for who is authorized to originate (or transit) what prefixes in the Internet routing system – or who owns what address space, for that matter
- Absent this, you cannot secure the routing system - or the data plane (e.g., inter-domain anti-spoofing mechanisms)
- Always horse/cart issue – no one had incentive to invest in database infrastructure; this incentive model has morphed (with the help of impending IPv4 Exhaustion and IPv6 ‘on the wire’ incompatibility)
- Development of RPKI underway, driven by RIRs and various government entities (e.g., US DHS) concerned with availability and security of Internet

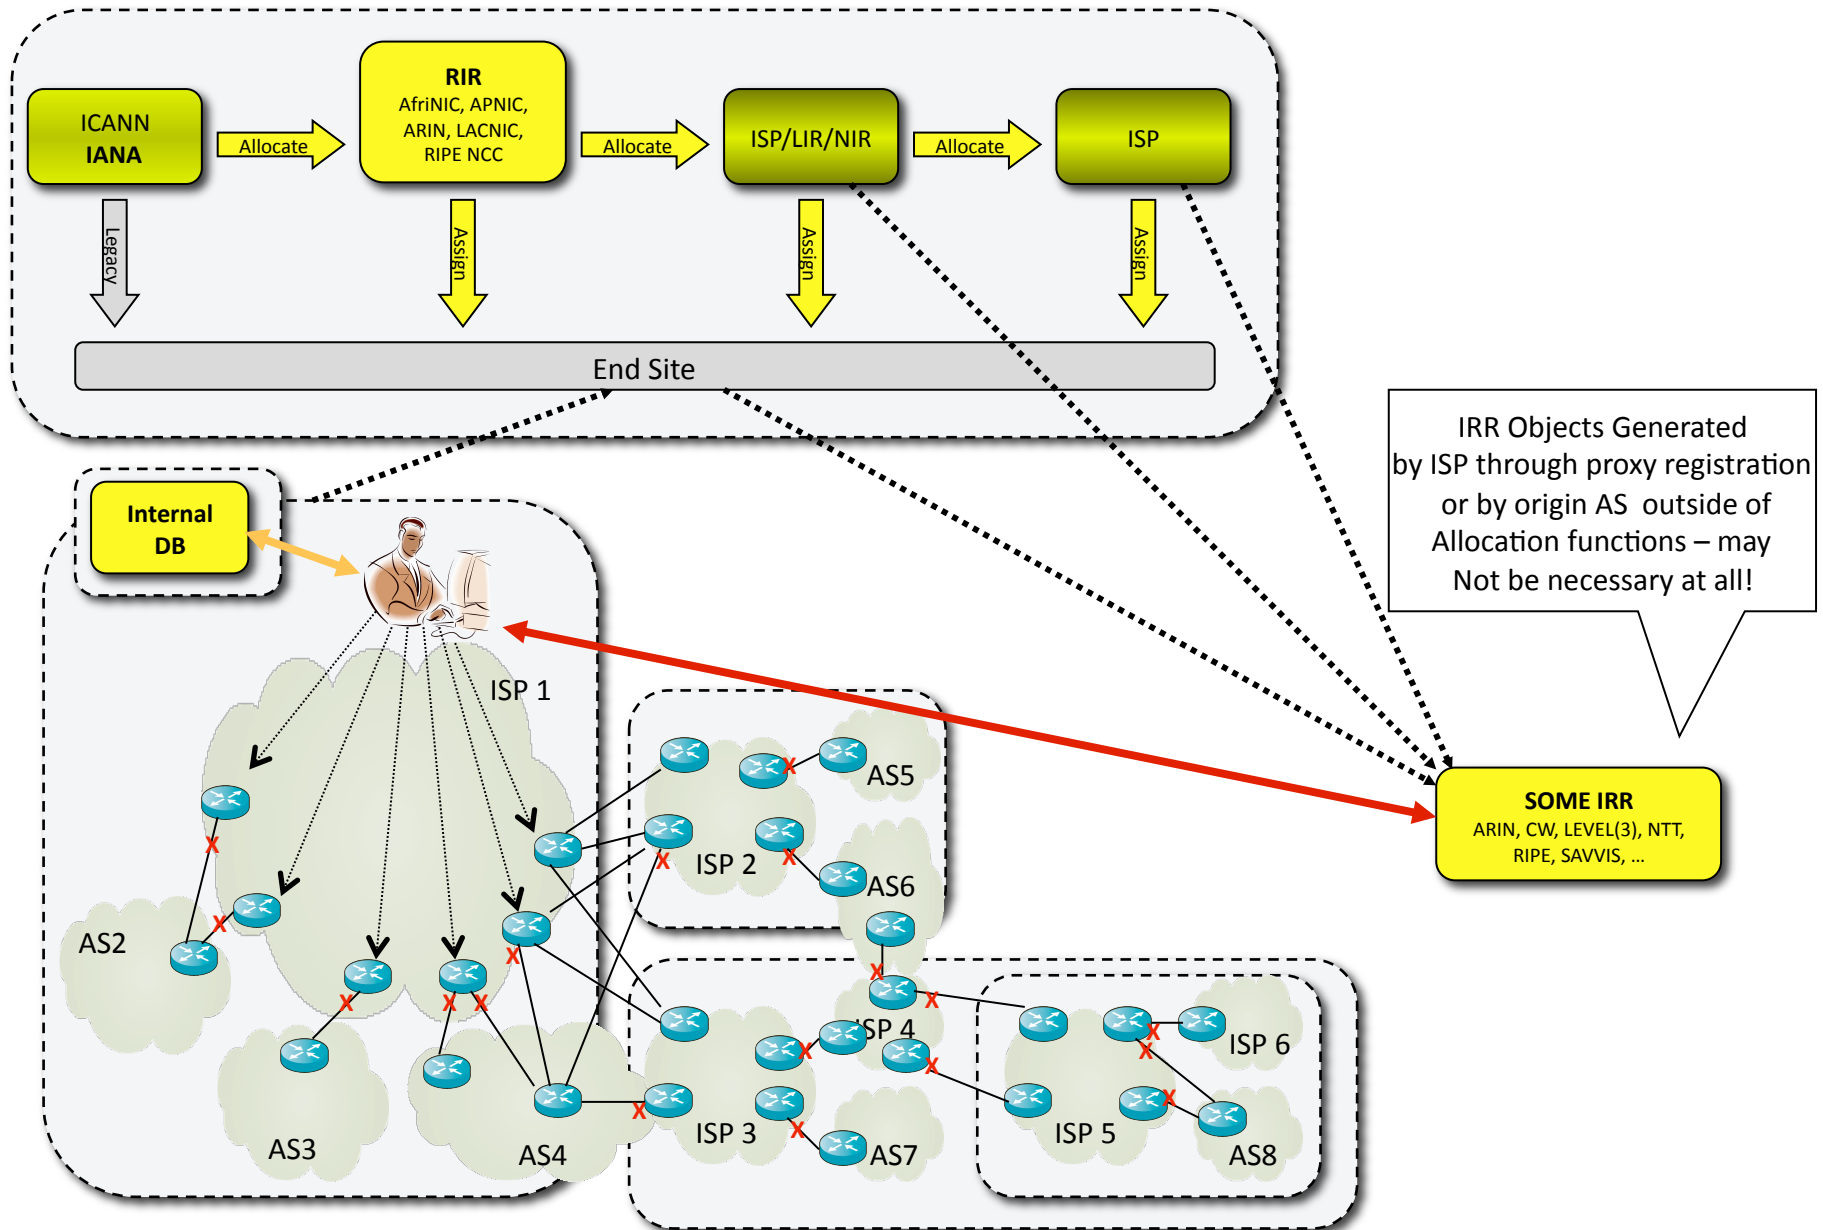
The Threat?

- Route hijacks
 - Routing **always** prefers the longest prefix
 - Then the shortest AS path
 - Scoping attacks topologically is trivial
 - E.g., Youtube & Pakistan Telecom
 - Malice or in error – same effect, mostly
- Route table leaks
 - ISPs often prefer customer routes over peer routes; if customer leaks routes they may become transit between ISPs
- Deaggregation
 - E.g., AS7007, hard to recover, router FIBs blown out, etc..
- Projected 1M++ meat computers with keyboard access to BGP speaking routers

Why, again?

- In general, there is NO tie-in from what's allocated by RIRs, LIRs, etc.. and what's routed
- RIRs have no operational role in today's routing system
- Internet Routing Registries (IRRs) challenged
 - Insecure update mechanisms
 - No one ever deletes objects (feasible path enumeration problems)
 - Some proxy register objects based on routing system state
- RIPE-181, then RPSL in IRRs used by some ISPs to generate custom route filter policies (route objects and as-sets for customer prefix filter generation), enumerate prefixes to be accepted from peer
- Very little (no?) explicit inter-ISP prefix filtering today
- Some generic policies by some ISPs, e.g.:
 - prefix length-based – deny le /25
 - deny bi-lateral peer AS in AS_PATH from transit customer
- The application of inter-domain filtering and the security of the routing system has only deteriorated in the last 15 years

IP Address Allocation, Assignment, and IRR Conceptual Model



Enter RPKI

- RIRs would maintain Resource PKI infrastructure that links prefixes with authorized origin AS via formally verifiable mechanisms
- Could be employed out of band (OOB) to generate prefix lists, to validate OOB and then install, or directly inband via secure routing protocol such as soBGP or SBGP down the road

More on RPKI

- Secure Inter-Domain Routing (SIDR) working group and RPKI today only concerned with Origin validation, route leaks and other similar threats to plausible
 - BGPSEC Requirements WG couldn't agree on path validation model
- Trust Anchors key discussion point currently
 - Who holds the keys (RIRs, IANA, both?)
- *Note: fundamental change to Internet routing architecture, necessary evil, perhaps – **trading autonomy for security***

BCPs for this..

1. If you speak BGP - follow IETF SIDR working group (and participate)
2. Engage with your RIR about what they're doing in this area
3. Register and keep your IRR data and your transit customer data up to date
4. Ask about security of your IRR
5. Automate explicit prefix filtering on all customers based on IRR data

In closing...

- Lots more to talk about in this area, but time was short..
- Please email with with any questions or for pointers, etc.. danny@arbor.net
- Thanks much for your time!
- EOF