# An Operational Perspective on Routing Security

Geoff Huston
Chief Scientist, APNIC

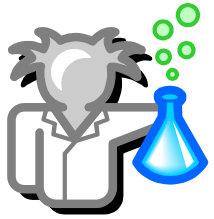

# On the Internet…

# there are many ways to be bad!

- Enlist a bot army and mount multi-gigabit DOS attacks
  - Extortion leverage and general mayhem
- Port Scan for known exploits
  - General annoyance
- Spew spam
  - Yes, there are still gullible folk out there!
- Mount a fake web site attack
  - And lure victims
- Mount a routing attack
  - And bring down an entire region / country / global network!

# If I were bad (and greedy)…

**I'd attack routing.**

- Through routing I'd attack the DNS
- Through the DNS I'd lure traffic through an interceptor web server
- And be able to quietly collect user details

Welcome to today's online fraud industry

# If I were really bad (and evil)…

## I'd attack routing.

- Through routing I'd attack:
  - the route registry server system
  - the DNS root system
  - trust anchors for TLS and browser certificates
  - isolate critical public servers and resources
  - overwhelm the routing system with spurious information

And bring parts of the network to a complete chaotic halt

# What's the base problem here?

- Routing is insecure
- Routing is built on sloppy mutual trust models
- Routing auditing is a low value activity that noone performs with any level of thoroughness
- We have grown used to lousy solutions and institutionalized lying in the routing system
- And because instances of abuse are relatively infrequent we are prepared to tolerate the risk of having an insecure routing system

# Routing is a shared problem

- It's a tragedy of the commons situation:
  - Nobody can single-handedly apply rigorous tests on the routing system
  - And the lowest common denominator approach is to apply no integrity tests at all
  - It's all trust and absolutely no defence

# So we need routing security

like we need motherhood, clean air and clean water

- But what does this “need” mean beyond various mantras, noble intentions and vague generalities about public safety and benefit?
  - Who wants to pay for decent security?
  - What’s the business drivers for effective security?
  - How do you avoid diversions into security pantomimes and functionless veneers?
- Can you make decent security and also support “better, faster and cheaper” networked services?

# Risk Management

- Adding operational security measures is not about being able to create and maintain absolute security. Its about a pragmatic approach to risk mitigation, using a trade-off between cost, complexity, flexibility and outcomes

- Its about making an informed and reasoned judgment to spend a certain amount of resources in order to achieve an acceptable risk outcome

# Threat Model

Understanding routing threats:

- What might happen?
- What are the likely consequences?
- What's my liability here?
- How can the consequences be mitigated?
- What's the set of cost tradeoffs?
- Does the threat and its consequences justify the cost of implementing a specific security response?

# Threat Response

- Collective vs unilateral responses to security threats
  - Should I trust noone else and solve this myself?
  - How much duplication of effort is entailed?
  - Is the threat a shared assessment?
  - Can we pool our resources and work together on a common threat model?
  - What tools do we need?
  - Are there beneficial externalities that are also generated?
  - Who wants to work with me?
  - What's the framework for collective action?

*When will you stop asking all these bloody annoying questions and just tell me what to do!*

# Routing Security

Protecting **routing protocols** and their operation

- Threat model:
  - Compromise the topology discovery / reachability operation of the routing protocol
  - Disrupt the operation of the routing protocol

Protecting the **protocol payload**

- Threat model:
  - Insert corrupted address information into your network's routing tables
  - Insert corrupt reachability information into your network's forwarding tables

# Threats

- Corrupting the routers' forwarding tables can result in:
  - Misdirecting traffic (subversion, denial of service, third party inspection, passing off)
  - Dropping traffic (denial of service, compound attacks)
  - Adding false addresses into the routing system (support compound attacks)
  - Isolating or removing the router from the network

# Operational Security Measures

- Security considerations in:
  - Network Design
  - Device Management
  - Configuration Management
  - Routing Protocol deployment

- Objectives:
  - Mitigate potential for service disruption
  - Deny external attempts to corrupt routing behaviour and corrupt routing payload

# Basic Network design

Isolate your network at the edge:

- Route all traffic at the edge
- NO sharing LANs
- NO shared IGPs
- NO infrastructure tunnels

Isolate your customers from each other:

- NO shared access LANs

Isolate routing roles within the network:

- Exterior-facing interface routers
- Internal core routers

# Configuration Tasks - Access

- Protecting routing configuration access
  - ssh access to the routers
  - filter lists
  - user account management
  - access log maintenance
  - snmp read / write access control lists
  - protect configurations
  - monitor configuration changes

- Protecting configuration control of routers is an essential part of network security

# Configuration Tasks - BGP

- Protecting BGP
  - Protect the TCP session from intrusion
  - Minimize the impact of session disruption on BGP.
  - Reduce third party dependencies to a minimum
  - Monitor and check all the time

# Configuration Tasks - BGP

Basic BGP configuration tasks:

- No redistribution from iBGP into the IGP
- Use session passwords and MD5 checksums to protect all BGP sessions
- For iBGP use the local loopback address as the nexthop (next-hop-self)
- Use filter lists to protect TCP port 179
- Use maximum prefix limiting (hold mode rather than session kill mode preferred)
- Use maximum as path limiting
- Use a silent recovery from mal-formed Updates
- Use eBGP multi-hop with care (and consider using TTL hack)
- Align route reflectors with topology to avoid iBGP traffic floods

Operating BGP:

- Use soft clear to prevent complete route withdrawals
- Use BGP session state and BGP update monitors and generate alarms on session instability and update floods

# Configuration Tasks - BGP

- Check your router config with a current best practice configuration template
  - Rob Thomas' template at http://www.cymru.com/Documents/secure-bgp-template.html is a good starting point
  - Remember to regularly check the source for updates if you really want to using a static bogon list

# Protecting the Payload

- How to increase your confidence in determining that what routes you learn from your eBGP peers is authentic and accurate
- How to ensure that what you advertise to your eBGP peers is authentic and accurate

- Manage your routes!
  - validate your customer's routes using registry information
  - filter your peers using route registries

# Customer Routes

- Authenticate customer routing requests:
  - Check validity of the address
    - Own space – validate request against local route object registry
    - Other space – validate request against RIR route object database registered POC
      - This is often harder than it originally looks!
  - Adjust explicit neighbor eBGP route filters to accept route advertisements for the prefix
  - Apply damping filters

# Exchange Peer Routes

- Higher level of mutual trust
- Accept peer routes  - apply local policy preferences
- Filter outbound route advertisements according to local policy settings
- Use max prefix with “discard-over-limit” action (if available)

# Upstream Routes

- One-way trust relationship
- Apply basic route filters to incoming route advertisements
  - RFC 1918 routes
  - own routes (?)

# Even so…

After all this effort, its not all that good is it?

# The Current State of Routing Security

Is pretty bad.

- This is a commodity industry that is not really coping with today's level of abuse and attack
  - Incomplete understanding
  - Inadequate resources and tools
  - Inadequate information
  - Inadequate expertise and experience

Can we do better?

# Routing Security

- The basic routing payload security questions that need to be answered are:
  - **Who** injected this address prefix into the network?
  - Did they have the necessary **credentials** to inject this address prefix? Is this a valid address prefix?
  - Is the forwarding path to reach this address prefix **trustable**?

- What we have today is a relatively fuzzy insecure system that is vulnerable to various forms of disruption and subversion
  - While the protocols can be reasonably well protected, the management of the routing payload cannot reliably answer these questions

# One approach…

- The use of authenticatable attestations to allow automated validation of:
  - the authenticity of the route object being advertised
  - authenticity of the origin AS
  - the binding of the origin AS to the route object

- Such attestations used to provide a cost effective method of validating routing requests
  - as compared to the today's state of the art based on techniques of vague trust and random whois data mining

# Taking a further step…

- Attestation validation to be a part of the BGP route acceptance / readvertisement process as a strong local selection preference
  - The use of a Route Origin Attestation that can validate the authenticity of the prefix and the validity of the originating AS

# What would also be good…

- A mechanism to check the validity of a received AS path:
  - Does the path represent a viable forwarding path through the network to reach the destination?
  - Has the Update Message itself traversed every element in the path?

# And what should be retained…

- BGP as a “block box” policy routing protocol
  - Many operators don’t want to be forced to publish their route acceptance and redistribution policies.
- BGP as a “near real time” protocol
  - Any additional overheads of certificate validation should not impose significant delays in route acceptance and re-advertisement
- BGP as a “simple” protocol
  - simple to configure, easy to operate

# Status of Routing Security

- We are nowhere near where we need to be
- We need more than “good routing housekeeping”
- We are in need of the adoption of basic security functions into the Internet’s routing domain
  - Injection of reliable trustable data
    - Address and AS certificate injection into BGP
    - Use a PKI for address “right-of-use”
  - Explicit verifiable trust mechanisms for data distribution
    - Adoption of some form of certification mechanism to support validated routing protocol information distribution

# Status of Routing Security

- It would be good to adopt some basic security functions into the Internet's routing domain
  - Certification of Number Resources
    - Who is the current controller of the resource?
  - Explicit verifiable trust mechanisms for data distribution
    - Signed routing requests
    - Adoption of some form of certificate repository structure to support validation of signed routing requests
    - Have they authorized the advertisement of this resource?
    - Is the origination of this resource advertisement verifiable?
  - Injection of reliable trustable data into the protocol
    - AS path validation in BGP

# Current Activities

- Some interest in this activity from a variety of public and private sector players (and still a lot of the typical security scepticism)
- Take previous work on various forms of secure BGP protocols (sBGP, soBGP, pgBGP, DNSRRs) and attempt to develop a common architecture for securing the Internet's routing system
- IETF Working Group on Securing Inter-Domain Routing active in standardizing elements of a secure routing framework
- RIR activity in producing resource "title" certificates to as an adjunct to their registry data

# Current Steps in Securing Routing

- PKI infrastructure support for IP addresses and AS numbers
- Certificate Repository infrastructure
- Operational tools for near-line validation of signed routing requests / signed routing filter requests / signed entries in route registries
- Defining the validation elements of a routing system
- Validation of information presented in BGP Updates

# Concerns

- Any security mechanism has to cope with partial deployment
  - Which means that the basic conventional approach of “what is not provably good must be bad” will not work
  - Which means that AS path validation is going to be very challenging indeed
- Which implies that a partially “secure” environment is more expensive but no more secure than what we have today

# Concerns

- Concentration of vulnerability
  - If validation of routing information is dependant on the availability and validity of a single root trust anchor then what happens when this single digital artifact is attacked?
- But can you successfully incorporate diversity into a supposed secure framework?
  - This is challenging!

# Security only works in practice if:

we can make secure mechanisms cheaper, easier, more robust, and more effective than existing practices

- Security as an added cost product feature has been a commercial failure in the Internet
- We need to understand how to deploy secure mechanisms that can reduce operational costs and bolt security features into the basic fabric of the Internet

# Thank You

Questions?