Securing BGP

Large scale trust to build an Internet again

Lutz Donnerhacke

db089309: 1c1c 6311 ef09 d819 e029 65be bfb6 c9cb

IKS

A protocol from better times

- A protocol from the early Internet
  - People were friendly and trustworthy
  - Internet was a warm and fuzzy place
- *BGP: protocol from admins for managers*
  - Main assumption: Routers do not lie
  - Idea 1: Announce what you have
  - Idea 2: Redistribute politically
- Inject locally, route globally

An example

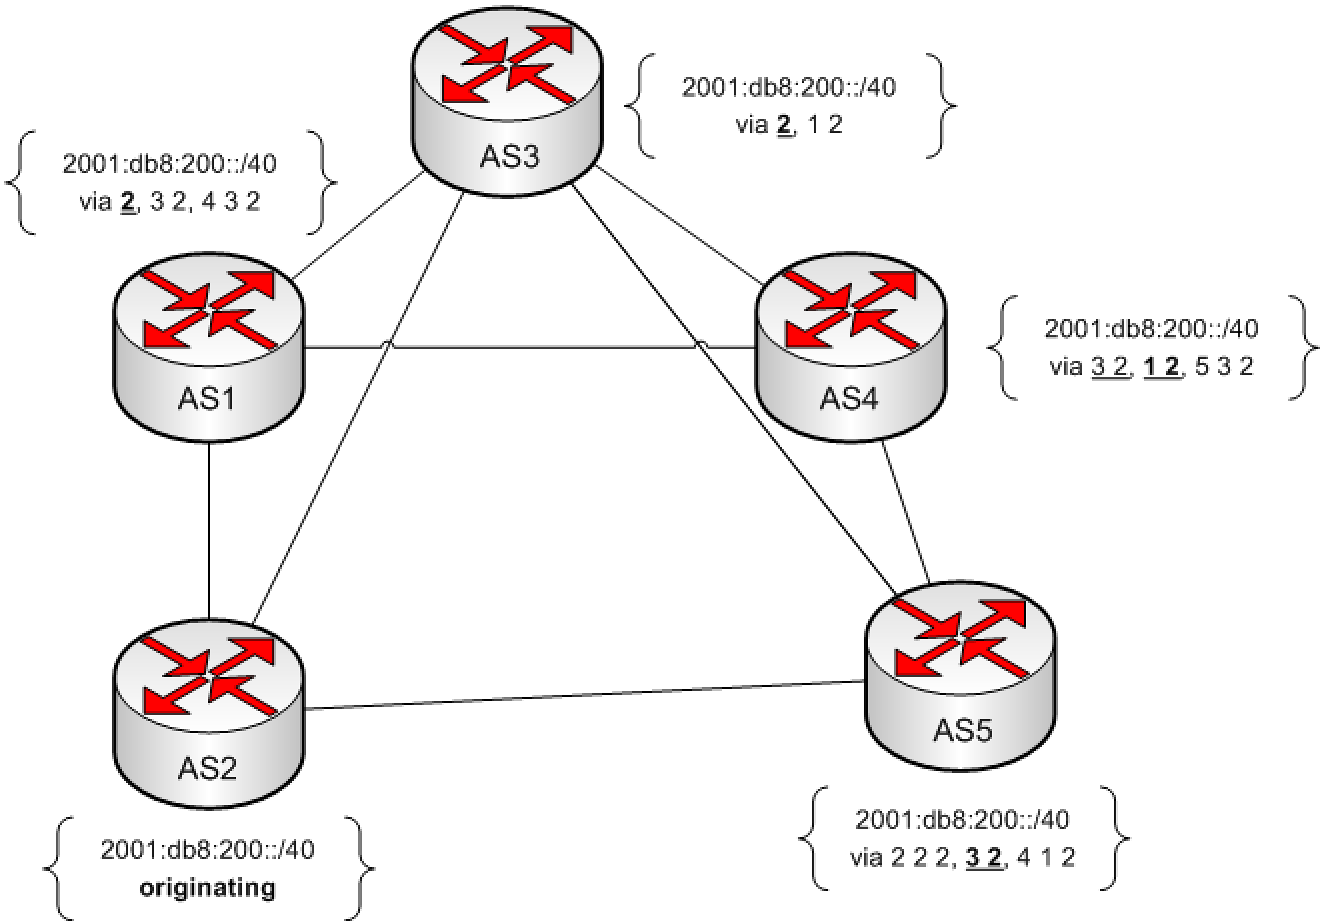

AS3
2001:db8:200::/40
via 2, 1 2
AS1
2001:db8:200::/40
via 2, 3 2, 4 3 2
AS4
2001:db8:200::/40
via 3 2, 1 2, 5 3 2
AS2
2001:db8:200::/40
originating
AS5
2001:db8:200::/40
via 2 2 2, 3 2, 4 1 2

Policy documentation

- Whois database
  - Distributed store of resource allocation
  - Database ensures correctness
- RPSL database
  - Centralized store of peering information
  - Both views of a peering: Sender / Receiver
  - Detailed peering policy incl. filter, precedence
- Software available
  - Generates router configuration

IKS

Threats to BGP

- Fat fingers
  - Announcing wrong network
  - Prepending foreign ASN
- Broken devices
  - Bitflip in memory or transit
- Commercial/criminal attacks
  - Redirect traffic (claim prefix, claim peering)
  - Inject unallocated networks (sending Spam)
- Governmental/Lawful attacks
  - Filtering traffic to protect the innocent

soBGP

- Trustworthy ISP approach
  - Transport authorisations as BGP attribute
  - Certifying assignment of a prefix by parent
- Each AS is a X.509-CA
  - Certifying injection policy per prefix (which ASNs are/is/isn't the first peerings)
  - Certifying it own peering policy with peers
- Web of trust
  - Resilience against erroneous behaviour
  - Permitting multiple hierarchies

S(ecure)-BGP

- RPKI approach
  - Transport authorisations as BGP attribute
  - Certifying allocation of prefix/ASN top-down
- Each ISP is a X.509-CA
  - Certifying injection policy: Prefixes per ASN
  - Certifying it own routers to sign redistribution
- Trust anchor management
  - Accessing various CA repositories

IKS

S-BGP operation

- Routers
  - Access external caches for object verification
  - Sign **each** update announcement
  - New hardware for storage and crypto operation
- Resource deallocation
  - Prefix updates time out => ~15 updates/s
  - Certificate and CRL times out => rsync
- Only one structure
  - Errors are disastrous
  - Ideal for LE

An other approach

- RPSL / Whois
  - Use it for non-local checks (was it allowed?)
  - No modification to BGP protocol
  - Skips gaps in deployment
  - Fails to deal with non-public policies
- Use DNSSEC ?
  - DNS as a trustworthy, distributed database
  - Routers: Offload crypto to AD-bit, caching implicit
  - Drastic RPSL simplification necessary

IKS

Comparison

Criteria	soBGP	Secure-BGP	RPSL	DNSSEC
ASN Alloc	Web of trust	RPKI	Whois	DNS
Prefix Alloc	Web of trust	RPKI	Whois	DNS
Private IP/AS	Other TA	Other TA	No	Stub zone
Router in AS	Validated	Validated	Unchecked	Unchecked
Outgoing Peer	Validated	Traced	Validated	Existence
Incoming Peer	Validated	Unchecked	Validated	Existence
Withdraw	Unchecked	Unchecked	Validated	Validated
Early scope	Many islands	Few islands	Full network	Full network
BGP protocol	Change	Change	Keep	Keep
Router HW	Change	Change	Keep	Keep
Helper Device	No	Simple Cache	Complex API	Resolver

Questions?

Why the approach is wrong

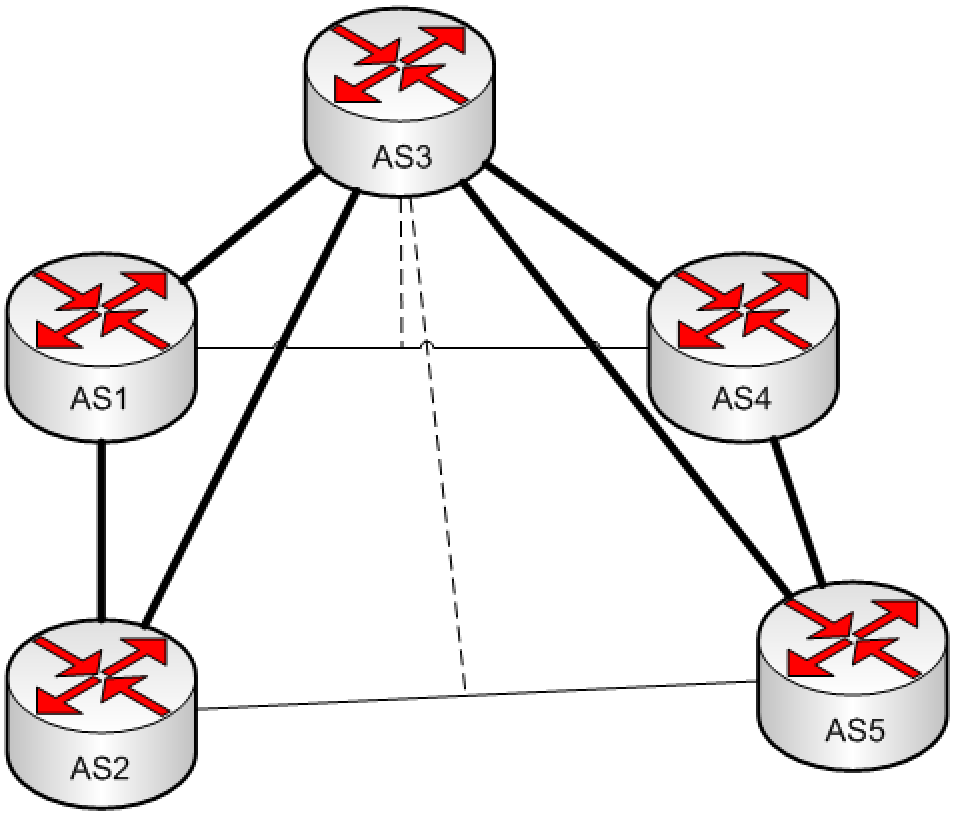

Why the approach is still wrong