Inter-Domain Routing

CMSC 414

November 15, 2017

Simple View of the Network

Hosts with a few routers

Nodes have

- Addresses
- *Forwarding tables*

Routing creates/maintains forwarding tables

- Start with neighbors
- Routing messages propagate info throughout network
- Adapts as network changes

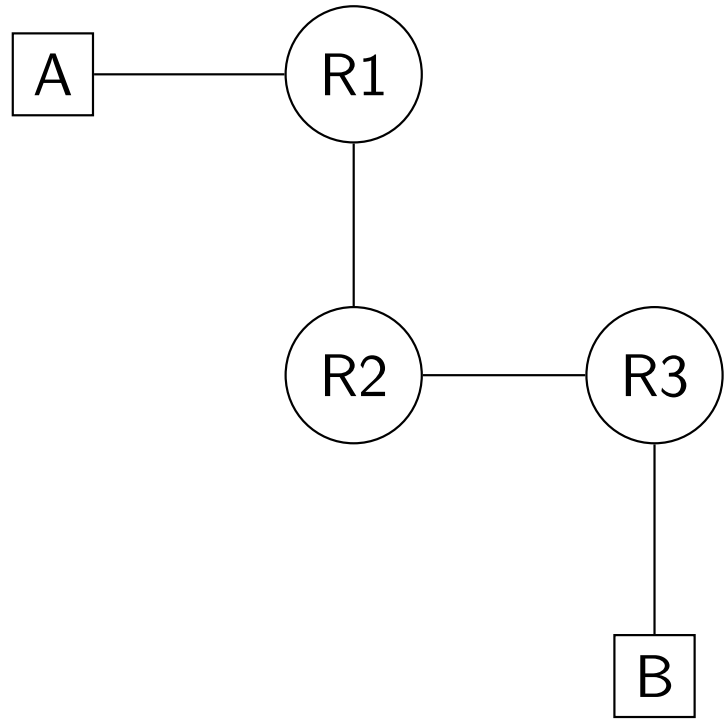

Problems with this Model

Doesn't scale

Doesn't match the way we want/need to manage networks

- Ownership of network equipment
- Policies
- Traffic engineering/prioritization

Already saw subnets

$\Rightarrow$ These are grouped into larger administrative domains

Autonomous Systems

Autonomous Systems

Someone owns the networking equipment and physical links

- Manage these as a cohesive unit
- Control what other networks we connect to and how we forward traffic to/from them

We call this grouping an Autonomous System (AS)

AS-to-AS connections go through **Gateway Routers**

Within an AS, forwarding is simply based on shortest path to destination

Between ASes, forwarding is based on the next AS that handles that IP address range

Forwarding Tables

Forwarding always based on destination IP address

Forwarding table has a next hop for *every addr in the network*

$\Rightarrow$ Grouped into CIDR blocks

$\Rightarrow$ The *next hop* is specified as an interface

Within an AS, will have fairly small ranges (or default rules)

For external addresses, large ranges will forward towards a border gateway

Organization of Autonomous Systems

One internet provider might have multiple ASes

- Northeast
- Mid-Atlantic
- Southeast
- Midwest
- South
- Southwest
- Northwest

One geographical area might have several ASes

- Verizon
- Comcast
- Time Warner
- Cornell

Border Gateway Protocol

An AS has a *globally unique AS number*

Each AS owns at least one CIDR block

An AS *must* have *at least* one **BGP speaker**
$\Rightarrow$ Often one of the border gateways

BGP speaker sends **BGP Update** messages:

- "Here's what blocks I own"
- "Here are the *AS paths* I know to specific CIDR blocks"

"AS 1 owns 1.2.0.0/16, and *advertises* (1.2.0.0/16, [1]), (3.4.5.0/24, [2,1])"

BGP Rules

BGP includes policies, usually based on financial agreements
$\Rightarrow$ Paths are only advertised if they comply with policy

Longest-Prefix Rule: The CIDR block with the longest prefix is what we use for a destination

BGP finds *shortest AS paths* to all destinations
$\Rightarrow$ Must be policy-compliant and match longest prefixes

Prefix Hijacking

BGP is *not authenticated*

Anyone can announce any prefix they like

- Neighbors might choose you for those routes
- Neighbors might propagate those routes further

Specify a longer prefix than the legitimate block
$\Rightarrow$ Might get that block's traffic sent preferentially to you

Black Hole Drop all traffic to the target range

Impersonation/Interception Analyze traffic
Maybe forward it to correct destination

Toy Example

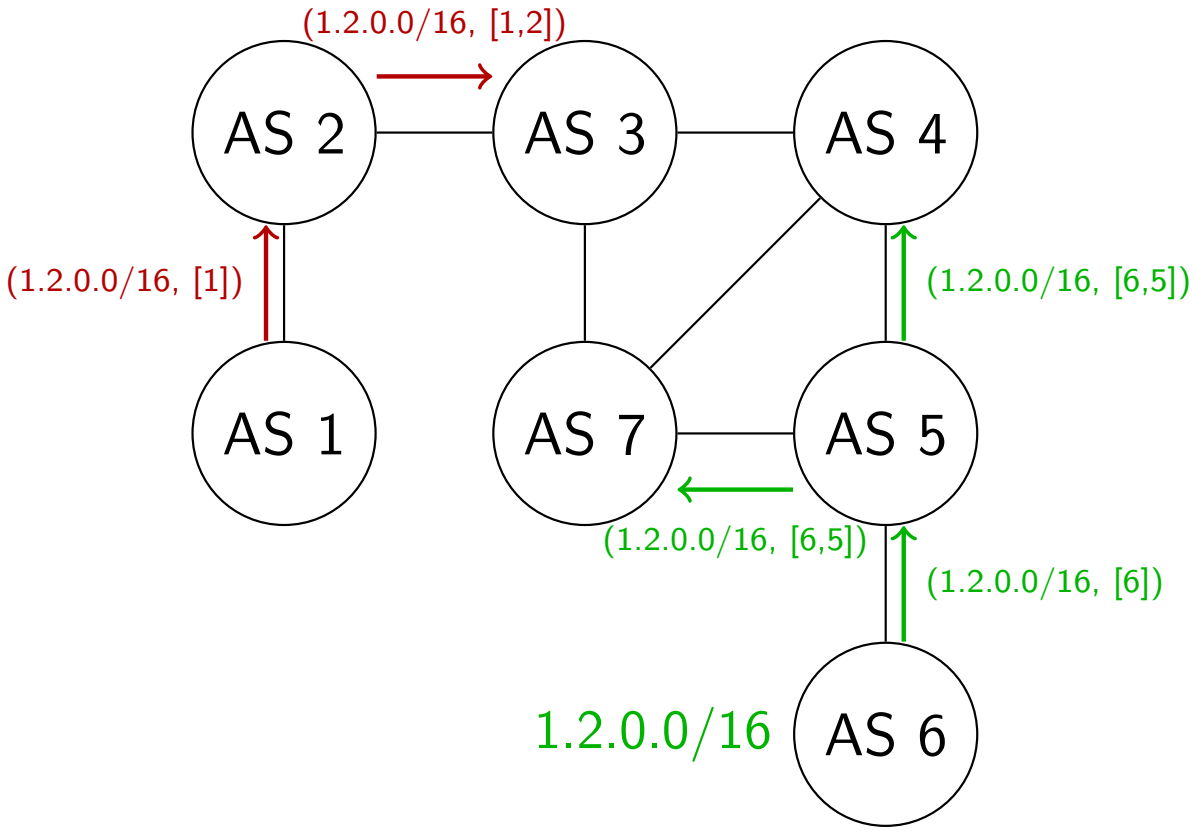

AS1 can fool AS2 and AS3 for 1.2.0.0/16

AS1 can fool everyone for 1.2.128.0/17

Real-World Examples

1997 Small ISP in FL broadcast 1-hop to everywhere
$\Rightarrow$ Took down Internet for $\sim$2hrs

2008 Pakistani gov tries to block YouTube by claiming longer-prefix CIDR block
$\Rightarrow$ $\sim$2/3 of Internet lost YouTube access for $\sim$2hrs

2010 Chinese telecom claims 1-hop to 1000's of networks (16k in US)
$\Rightarrow$ Lots of traffic goes to Beijing for 18mins
$\Rightarrow$ Supposedly an accident

Group Exercise 1

Use `get_assignment` to fork the repository `bgp`. This contains the beginnings of a toy implementation of BGP. Fill in the update propagation and processing code, and observe how routing information flows through a simple network of autonomous systems.

DoS Attacks

Malicious AS is on a non-preferred path to a target

DoS a BGP speaker for AS on best path to target
$\Rightarrow$ Neighbors *withdraw* routes

DoS ends
$\Rightarrow$ Speaker back online, re-establishes routes

Repeat the process
$\Rightarrow$ Causes **route flapping**

Flapping routes are deprioritized, to improve network stability
$\Rightarrow$ Malicious AS more likely to be on AS paths for target

Route Attribute Attacks

ASes set policies for routes, usually financial

- QoS guarantees
- Payment for transit
- etc

BGP Update messages set attribute values for breaking ties

- Path length
- Weight
- Paying customer
- etc

Bogus announcements

- Make path look shorter/longer
- Add victim AS to imply a loop

BGP Defenses

BGP is important, so people have looked at ways to secure it

Or at least prevent some bad behavior

- TTL Security Hack
- Defensive Filtering
- Authenticated Registry
- Digest for Integrity
- BGPSec

Time-To-Live Security Hack

Set TTL in BGP announcement to 255 (max allowed)

If we receive a packet with TTL $< 254 \Rightarrow$ ignore

Prevents attacks from multiple hops away

Does not defend against malicious/compromised insiders

Does not defend against tunneling-based attacks

Defensive Filtering

An AS can filter routes advertised by its customers

- If the customer doesn't own prefixes $\Rightarrow$ drop update
- Would have prevented Pakistan's YouTube attack

Customers have complex networks
$\Rightarrow$ Makes this logisitically challenging

Defensive filtering works best if everyone does it

The AS can also rewrite customers' BGP attributes to preferred values

Authenticated Registry

Can establish a public registry of accurate routing data

- Filter BGP updates to ensure consistency with this
- Can also include public keys (in a couple more slides)

Registry must be *complete*, *accurate*, and *secure*

Routing policies and topology within an organization *might be proprietary*

Digest for Integrity

MAC of TCP+BGP data per packet
$\Rightarrow$ First attempt using crypto

Can't be spoofed $\Rightarrow$ Fake routes ignored

Fits in existing TCP extension (optional behavior)

Requires *shared secret*

No confidentiality

BGPSec

Formerly call S-BGP

Uses certificates

Address attestation (claim right to a prefix)

- Hierarchical delegation up to ICANN
- Distributed out-of-band

Route attestation

- distributed within BGP update
- Signed by each AS in transit (nested signatures)

Full authentication of origins and paths

BGPSec Drawbacks

Expensive in time/storage

Every AS on a path *must* support BGPSec
$\Rightarrow$ No *incremental deployment*

On the way to becoming a standard, but

- $\sim$5% adoption globally
- less than 1% in North America (June 2015)

Group Exercise 2

In the same repository, add a malicious BGP announcement, and see how that impacts the routing decisions made by the autonomous systems in the network.