BGP101

Howard C. Berkowitz

hcb@clark.net

(703)998-5819

What is the Problem to be Solved?

- **Just configuring the protocol?**
- **Participating in the Internet**
  - **and/or running Virtual Private Networks**

A Life Cycle

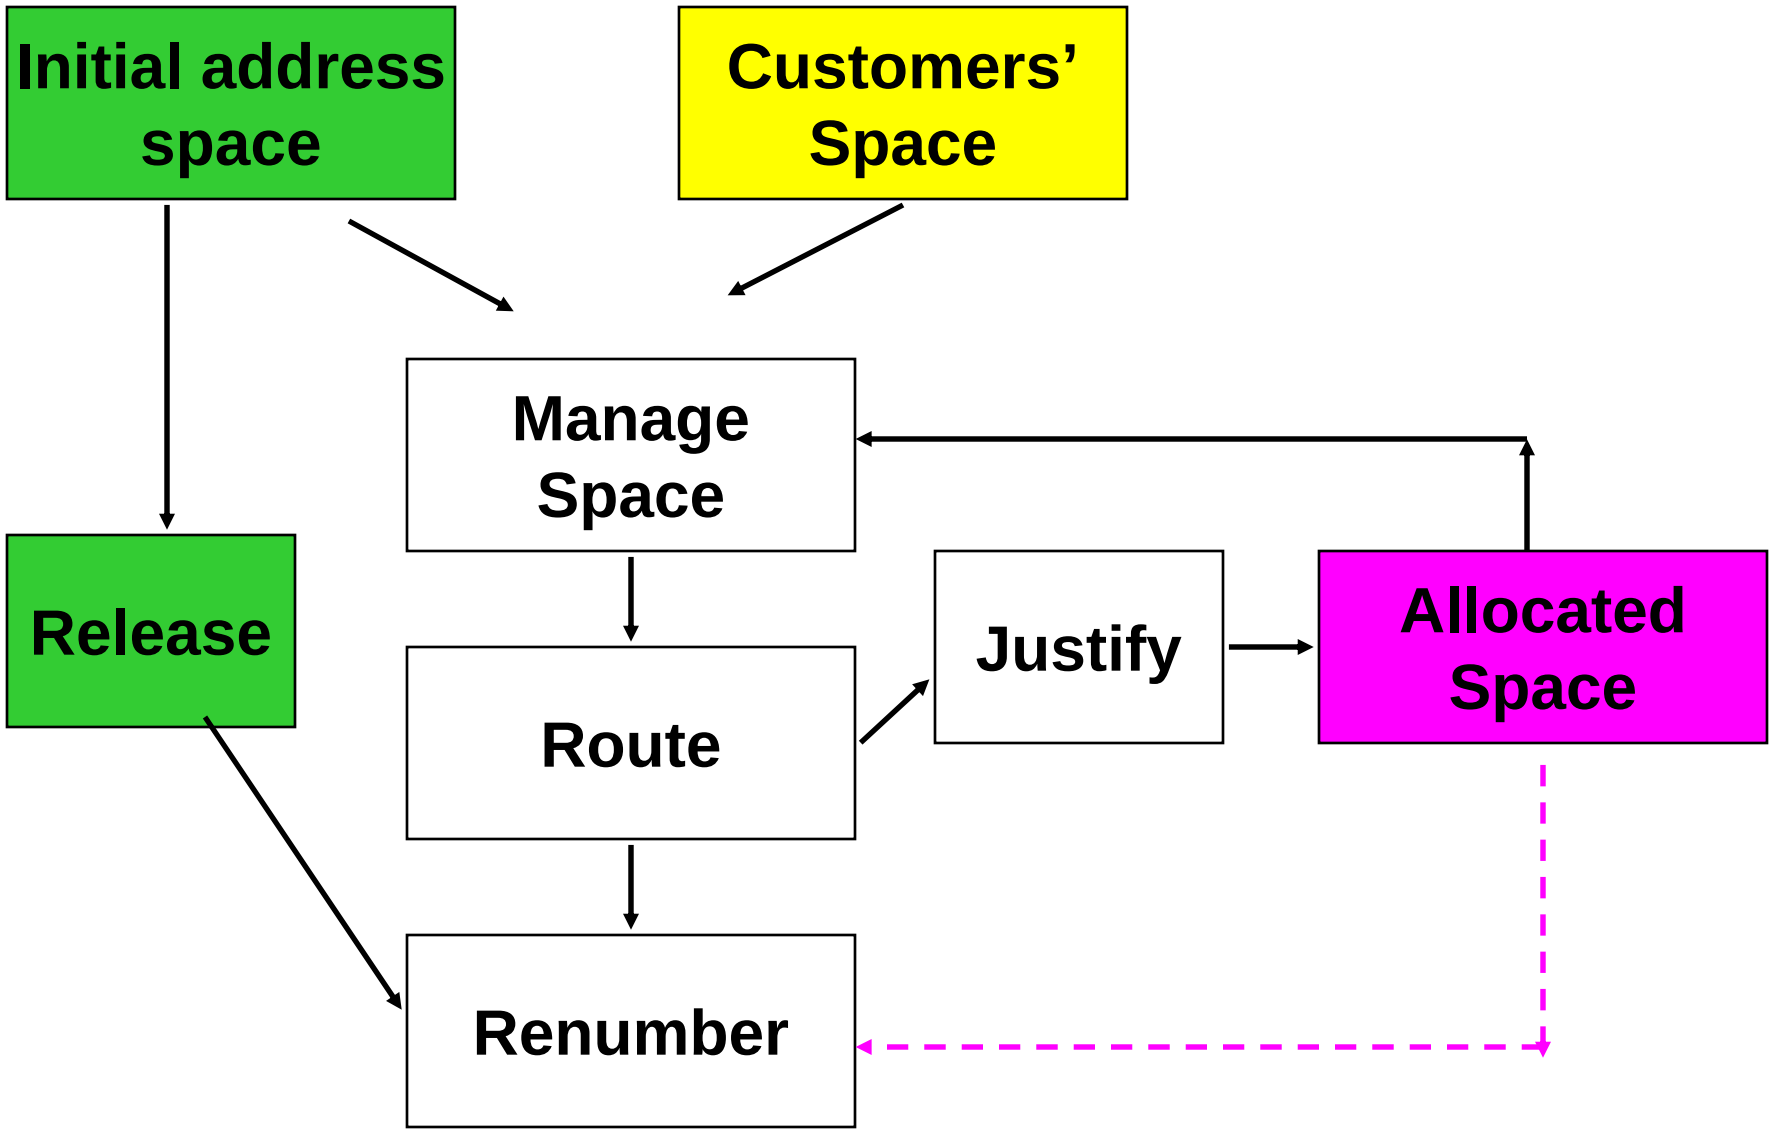
Initial address space
Customers' Space
Manage Space
Release
Route
Justify
Allocated Space
Renumber

"BGP Transmits Policies"

Policies are in Routers

- **Advertising Policies**
  - **Outbound to other AS**
  - **BGP advertisement sources**
  - **Outbound route filters**
  - **Route must be in internal routing table**
- **Acceptance Policies**
  - **Inbound AS filters**
  - **Inbound route filters**

BGP Stack

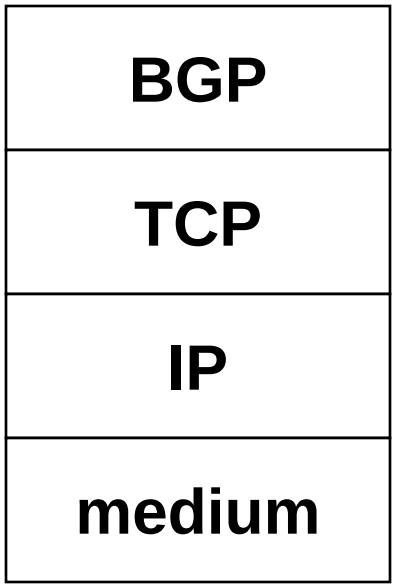

Policy vs. Protocol Flow

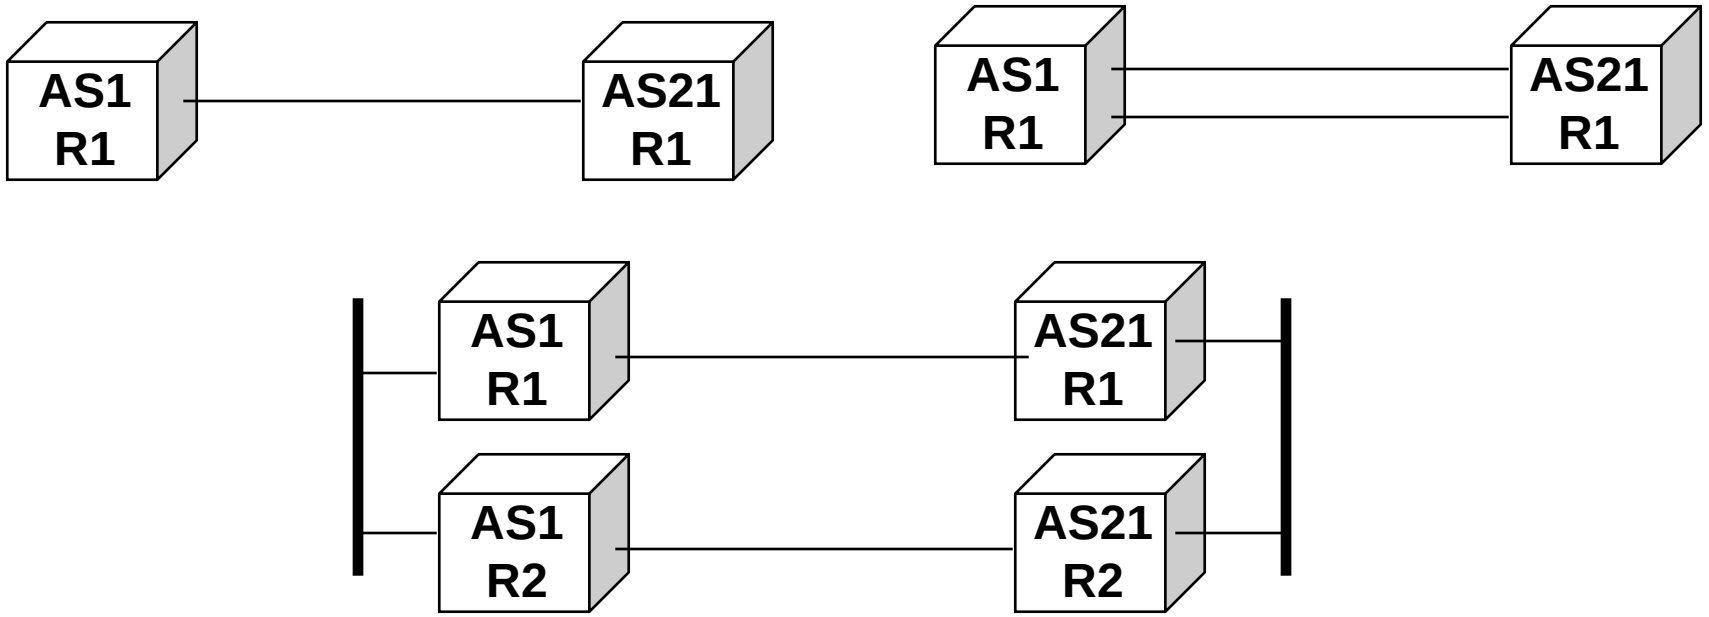

All equivalent from a policy standpoint!

AS Path

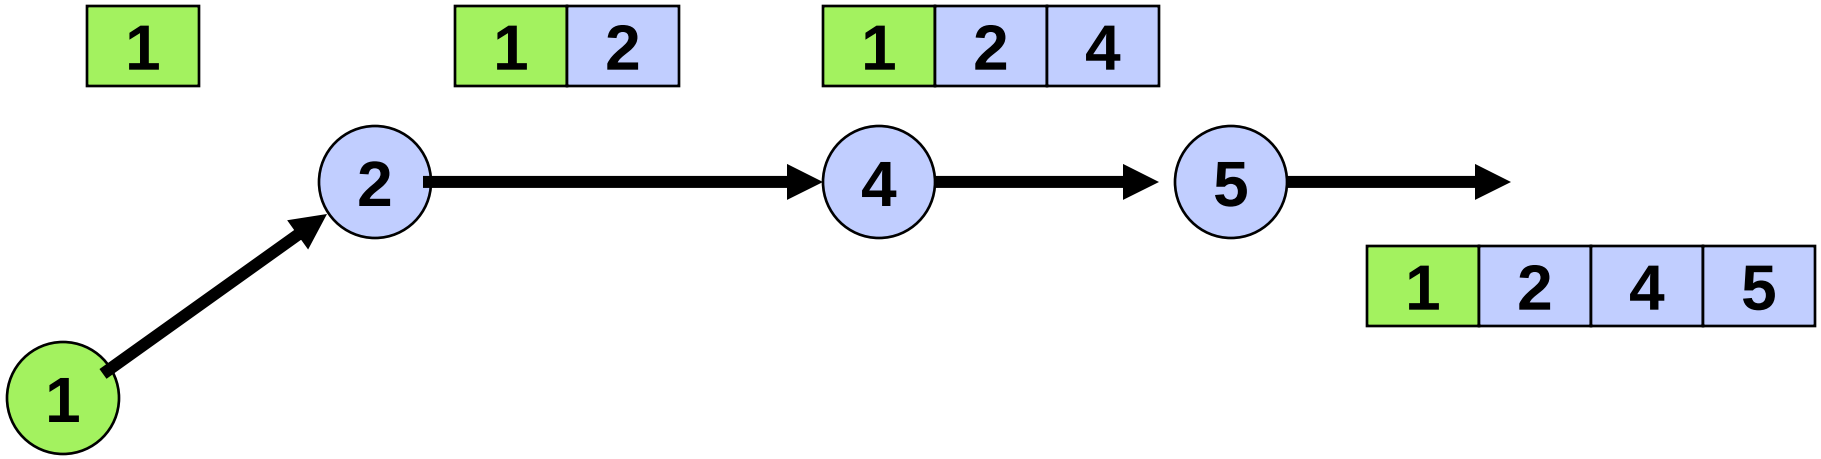

Network Layer Reachability Information (NLRI)

- Prefix
- Attributes of prefix
  - AS_PATH to reach it
  - Next hop
  - Community
- Address Aggregates
    - » Aggregate only
    - » Aggregate and all more specific
    - » Aggregate and some more specific
    - » Holes
  - • • •

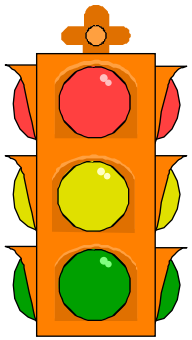

Stop!

What are you going to Advertise?

Advertising Affects

- The way the world sees you/sends to you
- Binary
  - Routes to which you provide routing
- Quantitative Preferences
  - Multi-Exit Discriminators to your Neighbors
  - AS Path Manipulation to all

Advertise

- **Routes Assigned/Allocated to You**
- **Routes Assigned/Allocated to Customers**
- **Routes for which you provide Transit**

Routes Eligible to Advertise

- Are reachable by your IGPor static routes
- Unless they are black holes
  - Which conceptiually are reachable

Do Not Advertise

- Spoofed source addresses
- Your internal addresses
- RFC1918 space
- Known rogues?
  - RBL?

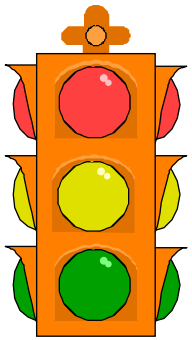

Stop!
What are you going to Accept?

Do Not Accept

- **RFC1918 source or destination**
- **Unexpected sources not assigned/allocated to peers**
- **Your internal addresses from peers**

Accept

- It depends
- Only those routes you will do something about
- Otherwise default

Definitions and Issues

Transit vs. Peer

- **Two meanings**
  - **General BGP Peering**
  - **Economic**

Peering vs. Transit Provider

- **Peer relationship**
  - **Mulual benefit customers reach one another**
  - **No monetary exchange**
  - **Each advertises customer routes**
- **Transit Provider relationship**
  - **Customer pays for service**
  - **Full routes available to customer**

Closest Exit Routing

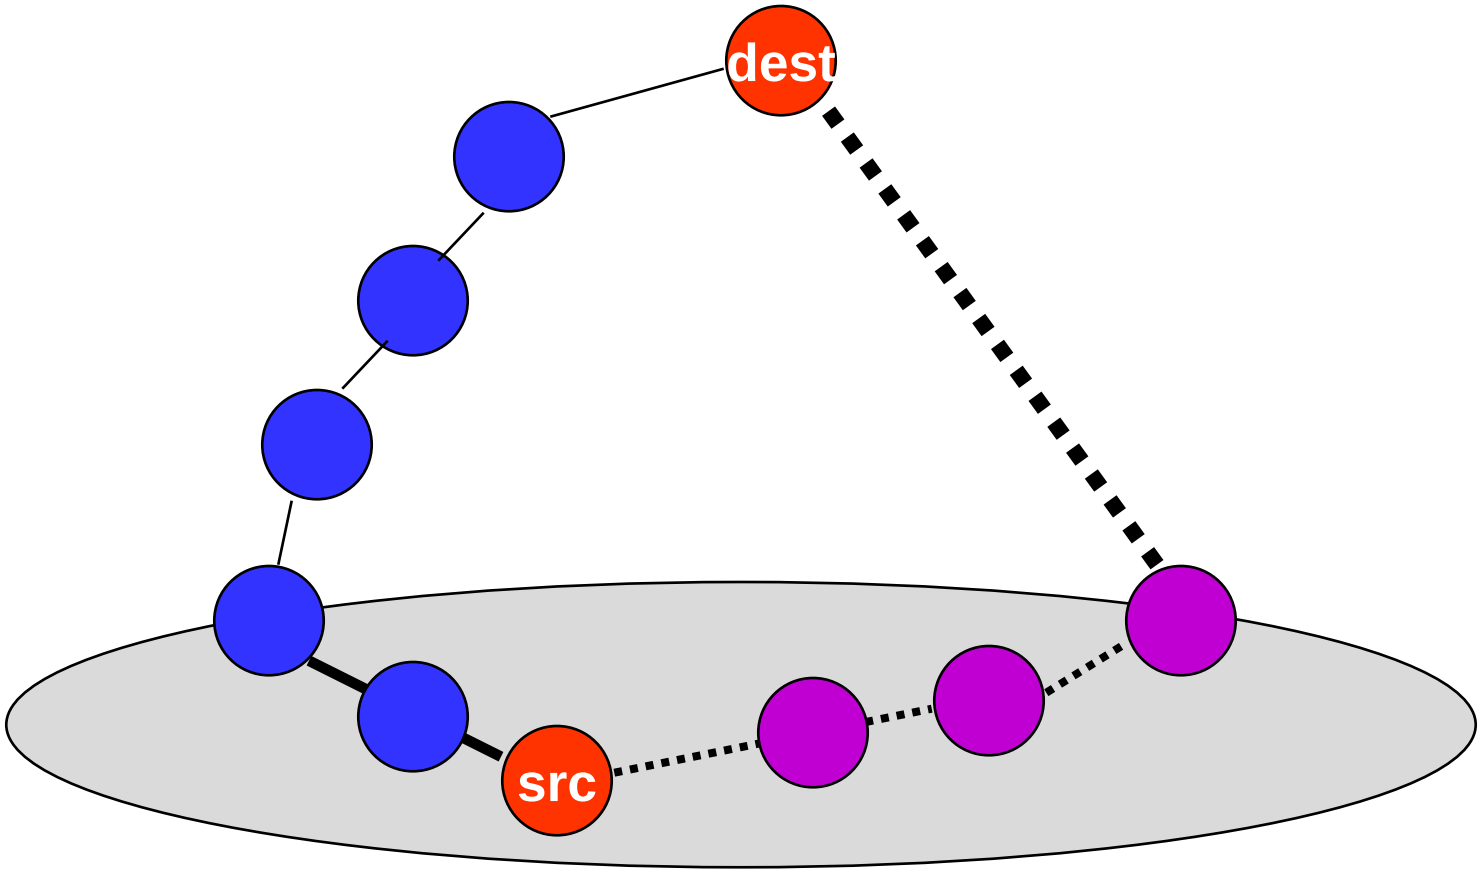

- **Paths are not optimized end-to-end**
- **Paths are optimized for each AS**

Asymmetrical Routing

- No guarantee that traffic leaving your AS at one point
- Will return at the same point
- Remember
  - Each AS in both directions makes decisions on its information

Operational Relationships 1

Addresses and Delegation

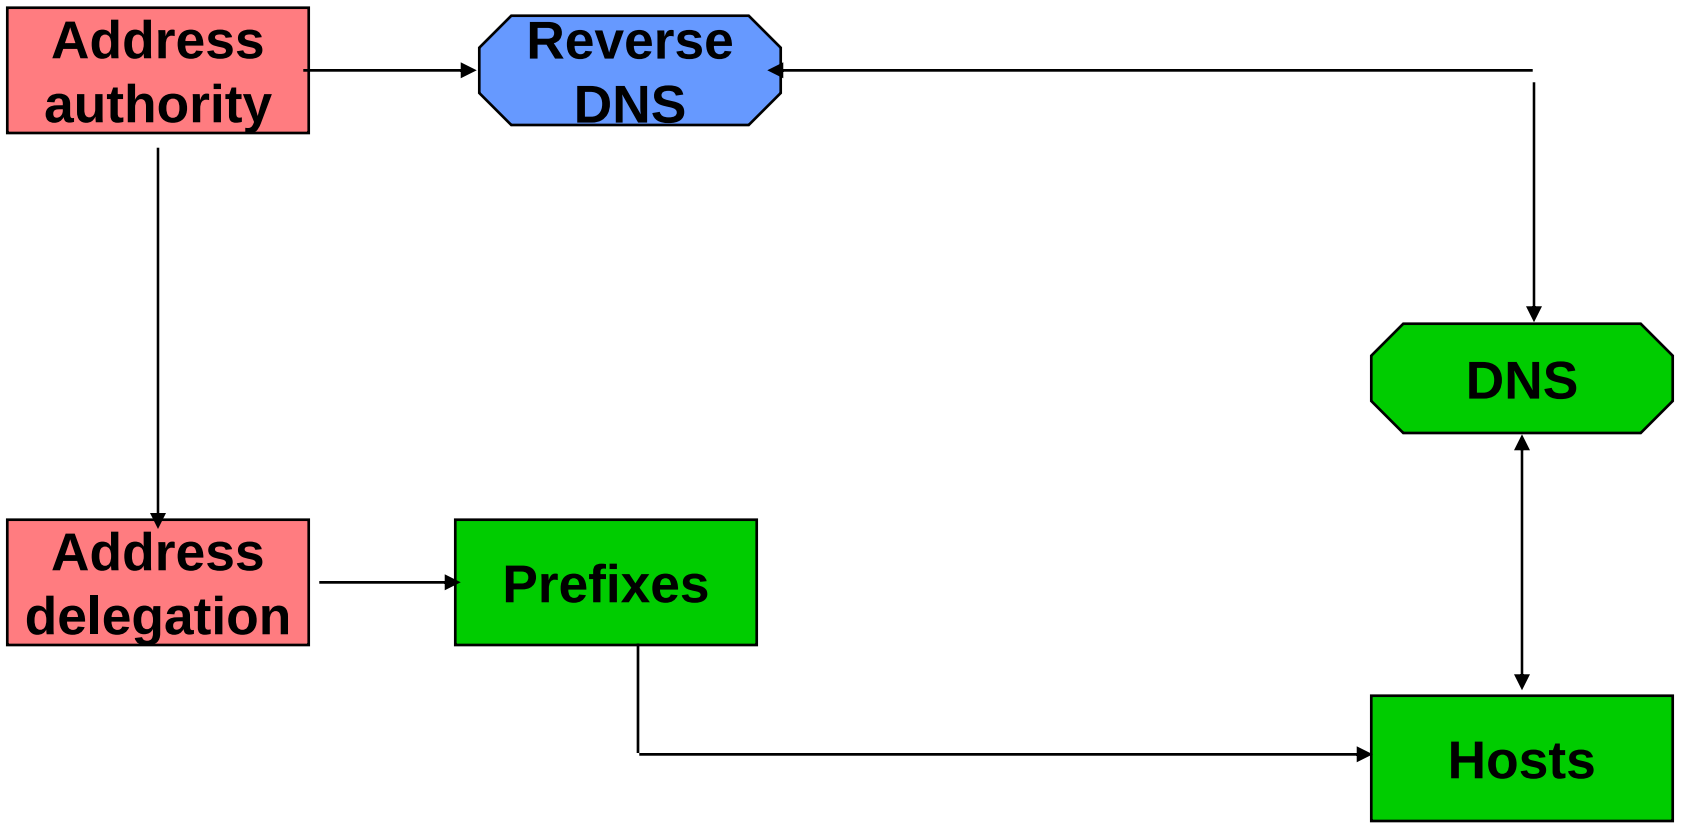

Autonomous System

- Basis of exterior routing
- AS originate routes for some prefixes they want to be visible
- AS advertise routes to one another
  - Advertisement may not contain all addresses
  - Not all advertisements need be accepted

Current AS Definition

RFC 1930

- Connected group of IP CIDR blocks
- Run by one or more network operators
- Single routing policy
  - announced to the general Internet
  - announced with BGP-4

AS Number

- 16 bit number
- Numbers assigned by registries
- Private ASNs
  - 64512 through 65535

Operational Relationships 2:

Addesses and Autonomous Systems

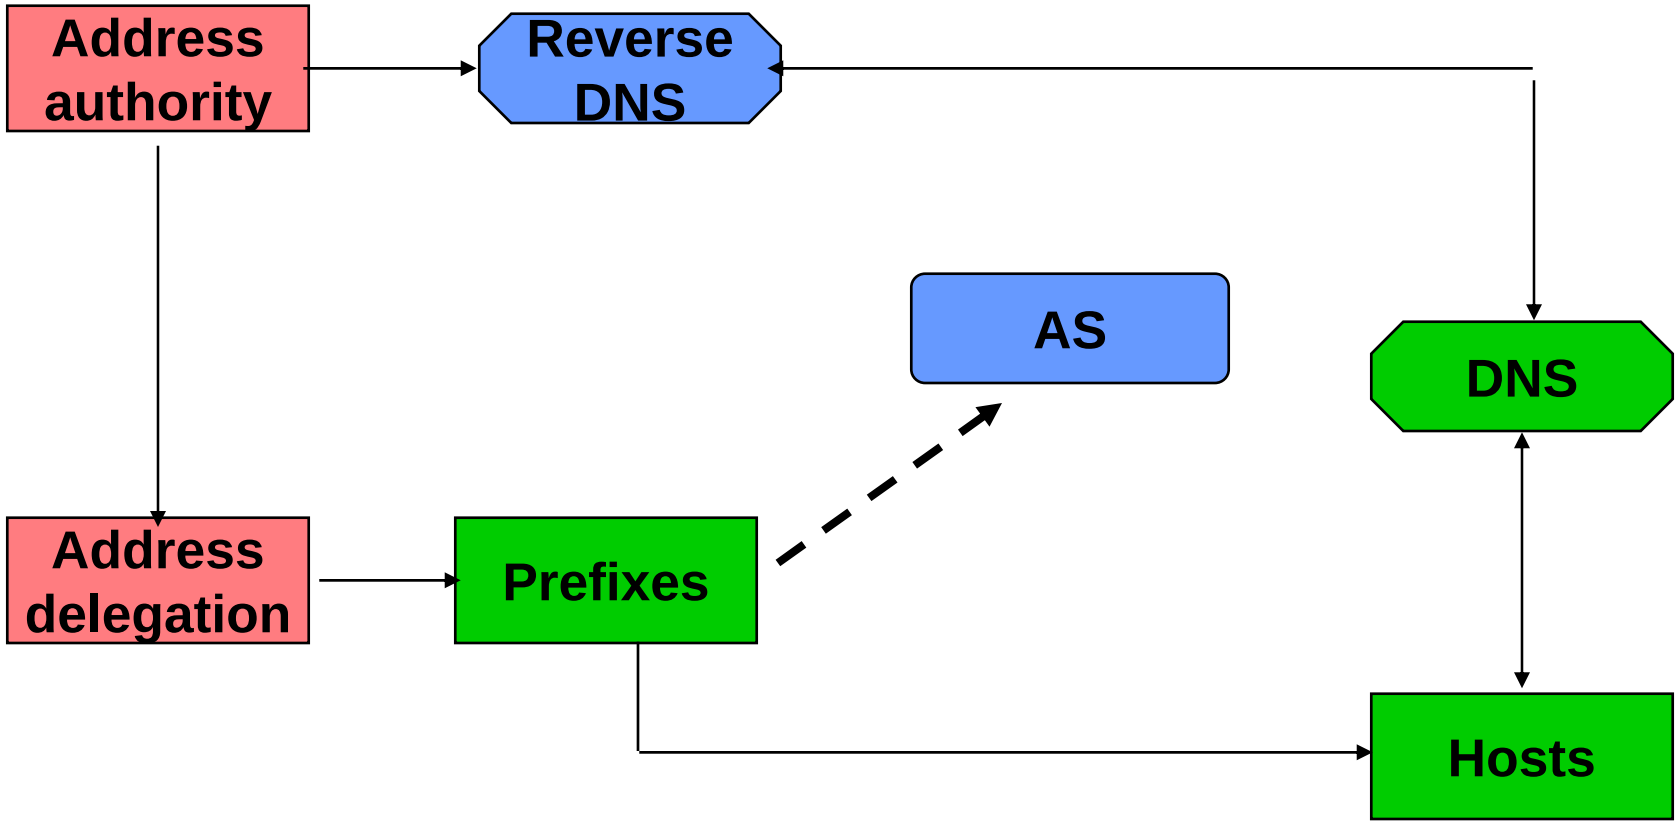

When is an AS Not Needed?

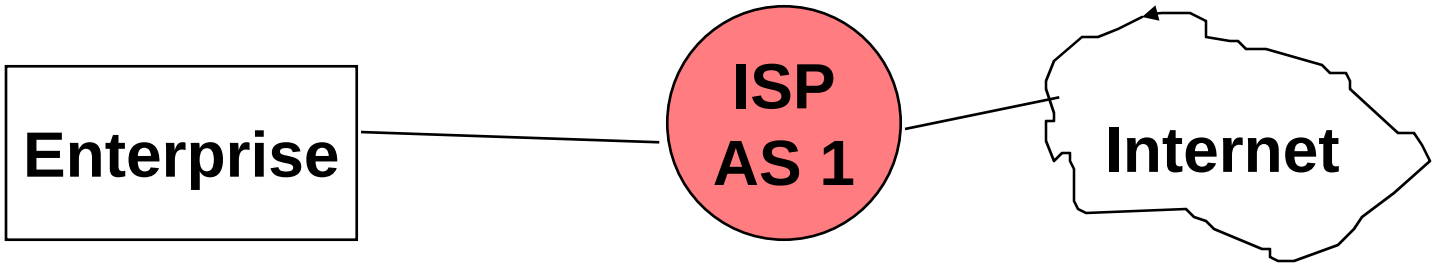

- Enterprise connects to Internet only via ISP
- Enterprise uses provider AS if it needs one at all

When is an AS Possibly Needed?

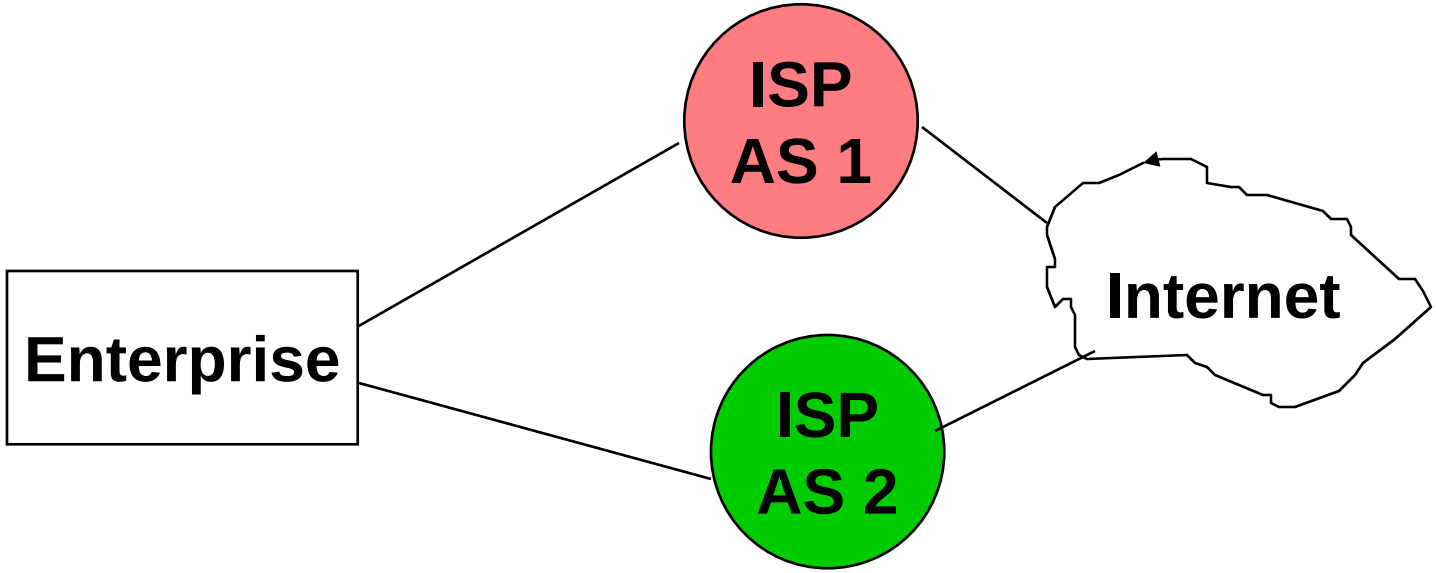

- Enterprise connects to Internet via multiple ISP
- AS may be needed for load-sharing or keepalive

When is an AS Needed?

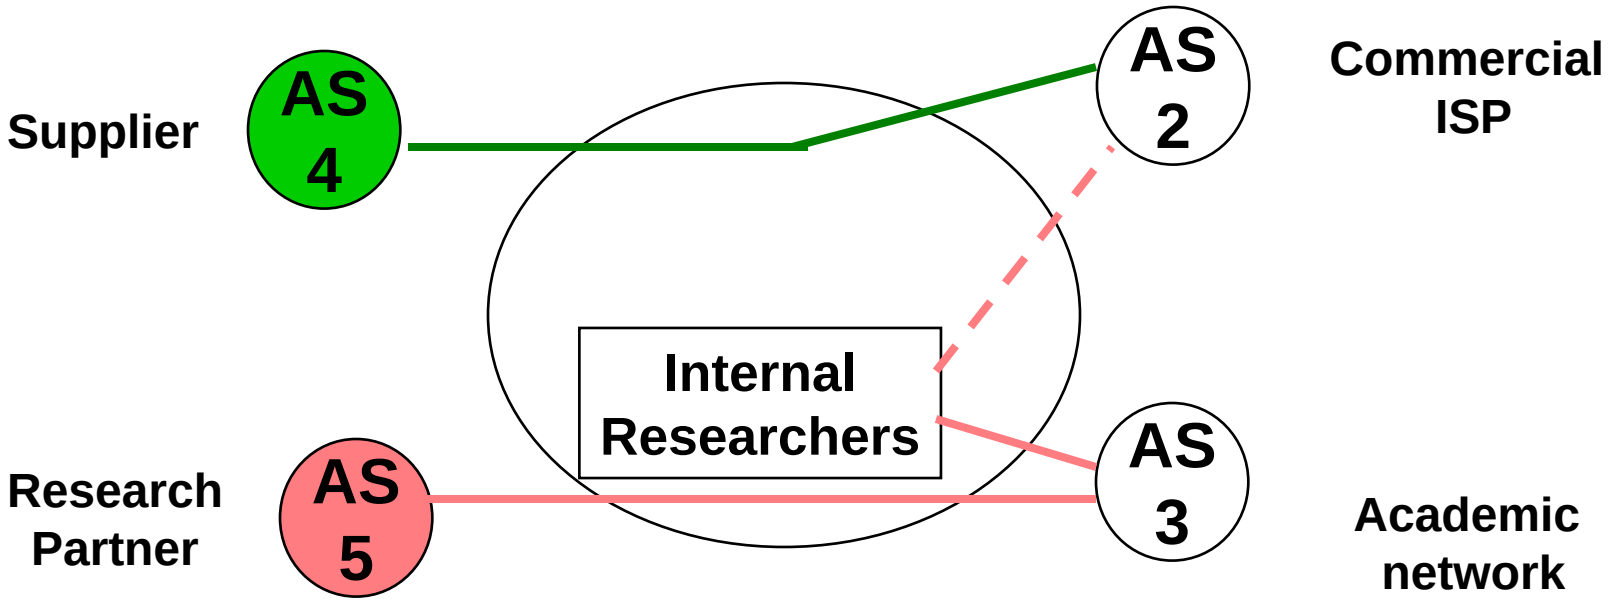

Routes/Aggregates

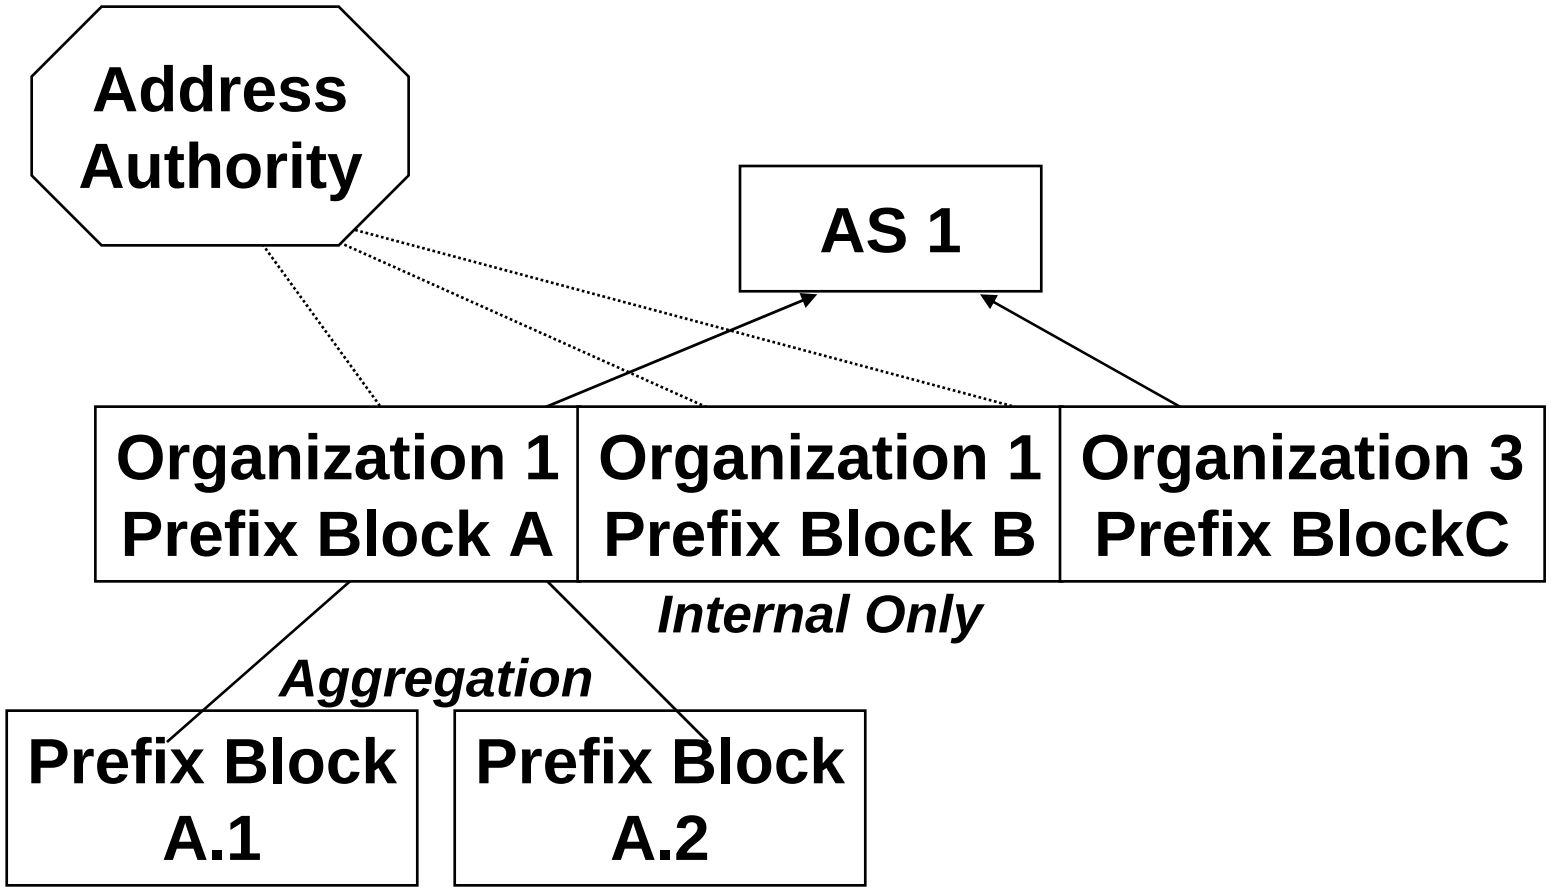

Relationships to Adjacent AS

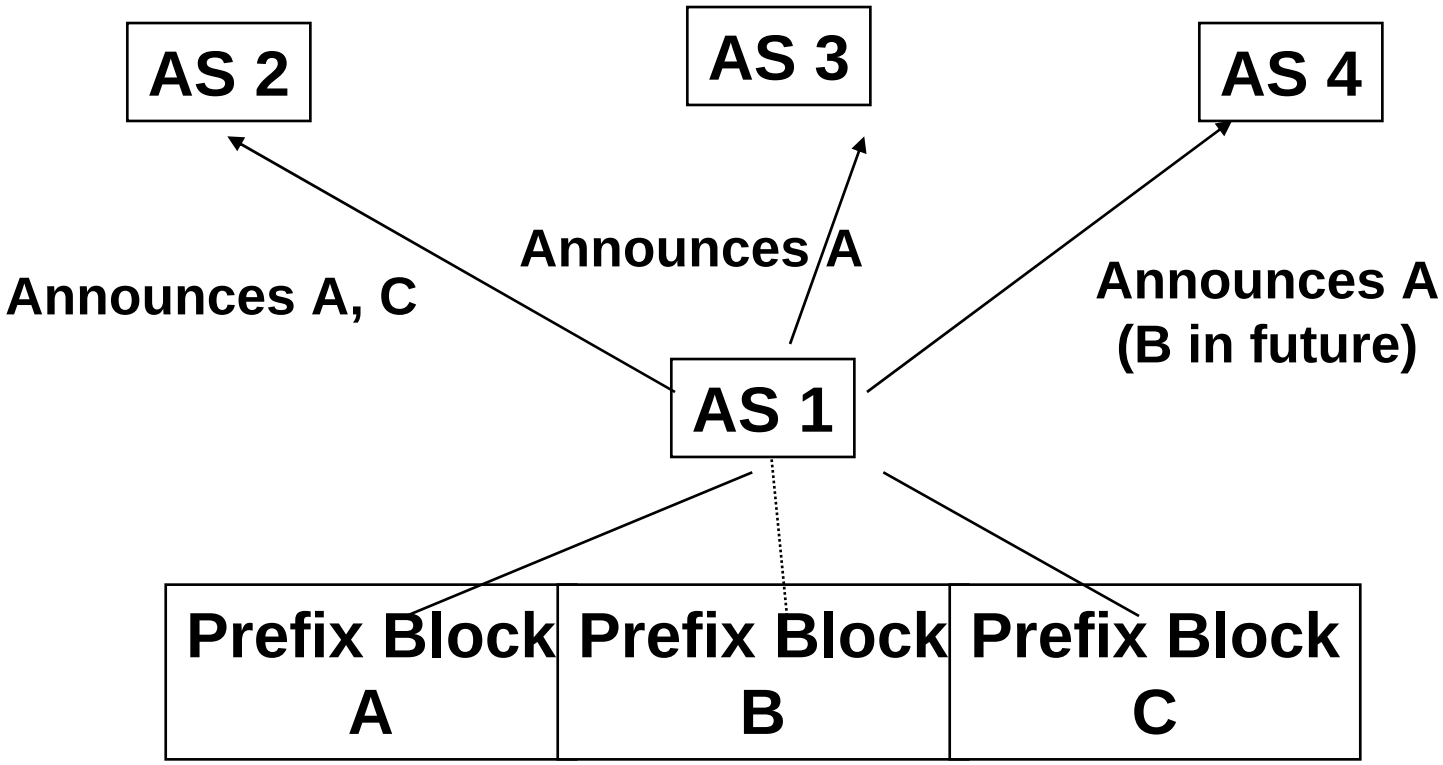

Joining the Club

- **Get/join an AS**
- **Register AS in appropriate registries**
- **Peer AS to other cooperating AS**

Establishing an AS (1)

Obtain routable address space

- **Apply to registry**
  - **RIPE, APNIC, ARIN**
  - **If immediate need for /19 or /20***
- **Obtain addresses from upstream ISP**
  - **If /19 or /20 cannot be justified**
- **Registry needs**
  - **Network design**
  - **Justification for address space**

Establishing an AS (2)

AS Number Request

- **In request to AS number registry**
  - **Administrative and technical contacts**
  - **Autonomous system name**
  - **Router description**
  - **Deployment schedule**
  - **Networks (by name) connected by the router(s)**
  - **Internet addresses of the routers**

Establishing an AS (3)

Registering in Routing Registry

- **Minimum requirements**
  - **Maintainer object**
  - **AS object**
  - **Route object (s)**

Establishing an AS (4)

Operational deployment

- **Build configuration**
  - **Policy implementation**
  - **Ingress/egress filtering**
- **Establish security procedures**
- **Start BGP connections**

Registry Issues

- Not all providers use registries
- If multihomed to any that does use registry...it's needed
- Procedure is useful as a checklist that everything is defined

Registry-Based Tools

- **Web-available Servers**
- **Routing Arbiter Tool Set**
- **Other tools**

Routing Registry Objects

- **Basic**
  - **AS**
  - **Route**
  - **Maintainer**
- **Additional**
  - **Inter-AS Network**
  - **Community**
  - **Router**

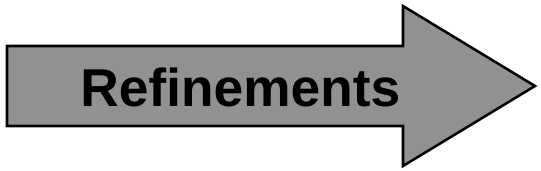

Refinements

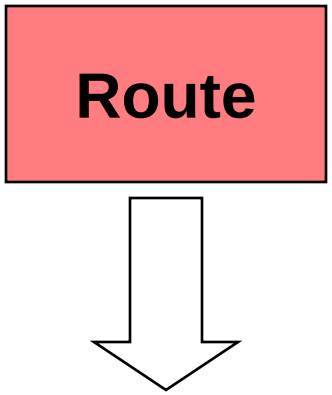

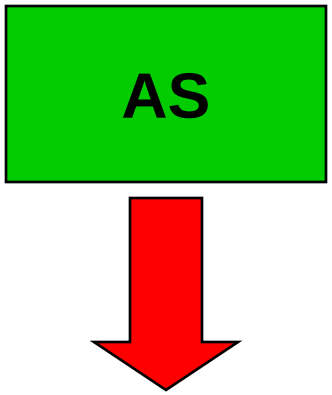

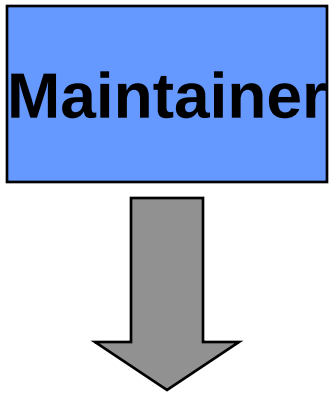

Route set
Route macro

AS set
AS macro

Person
Role

Operational Relationships 3:

Registries, Domains, etc.

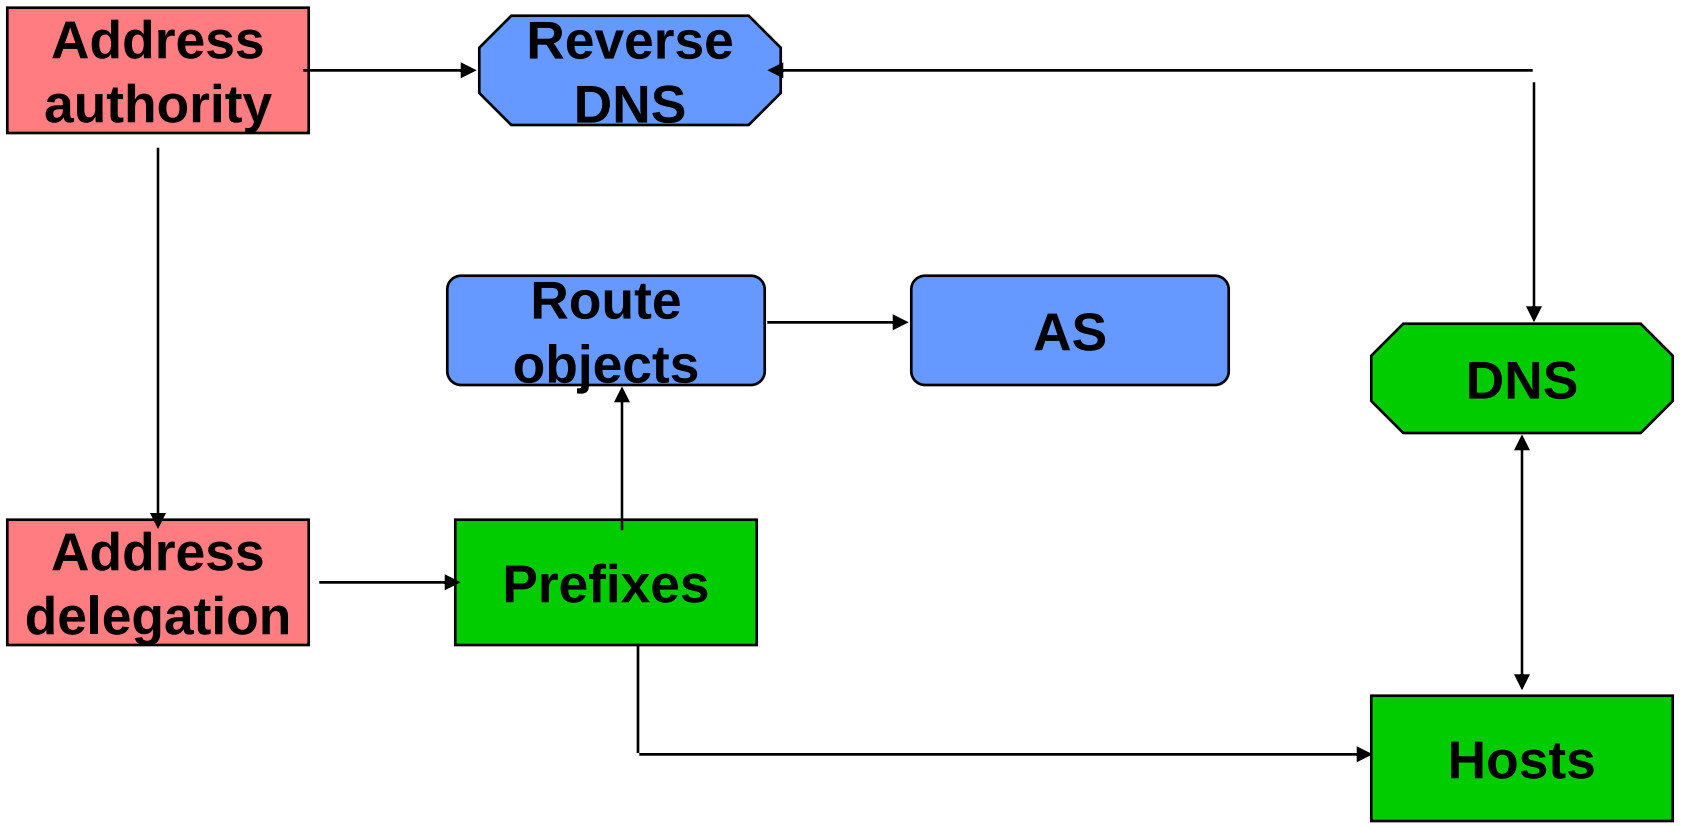

Turning it On

BGP Configuration Overview

- **Plans and policies first!**
- **Define system of BGP speakers**
- **Specific BGP speaker configuration**
  - **Identifier**
  - **BGP process**
  - **Neighbors**
  - **NLRI to advertise**
  - **Filters and other policy mechanisms**

Cisco commands used as examples

Consider

- **Register in registry**
- **Use rtconfig, etc. to generate configuration**

Router ID and loopback interface

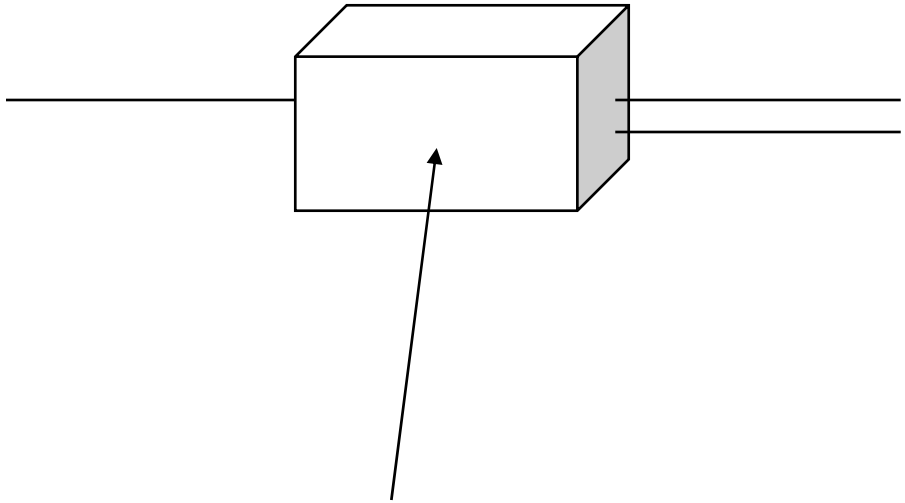

interface loopback 0
ip address 192.168.0.1 255.255.255.0

router bgp command

```
router bgp AS-number
```

- The parameter is a true AS number
- Subcommands will be needed -- minimally,
  - network
  - neighbor

`network` subcommand for BGP

`network` *`ip-address [mask]`*

- **Establishes eligibility of address ranges to be advertised**
- **But does not actually cause anything to be advertised**
  - **Different than Cisco IGP network statements**

Basic `neighbor` subcommand for eBGP

`neighbor` *`ip-address`*

`remote-as` *`as-number`*

- Identifies peer in external AS
- For eBGP, *`as-number`* is different from our own
- Will not advertise anything unless `network` statements are defined

Bilateral Exchange

Refining the Configuration

Single and Multiple Links to a Single Provider

The BGP Tunnel

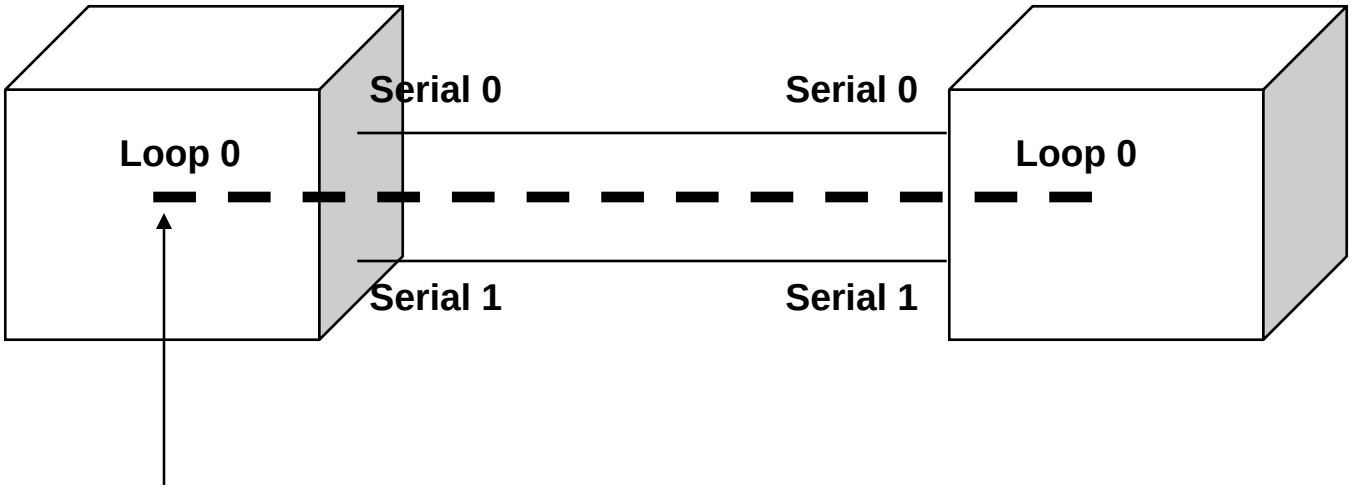

`ebgp-multihop` needed when neighbor is not on same subnet

Load Balancing 1:
IP Level to Single Provider

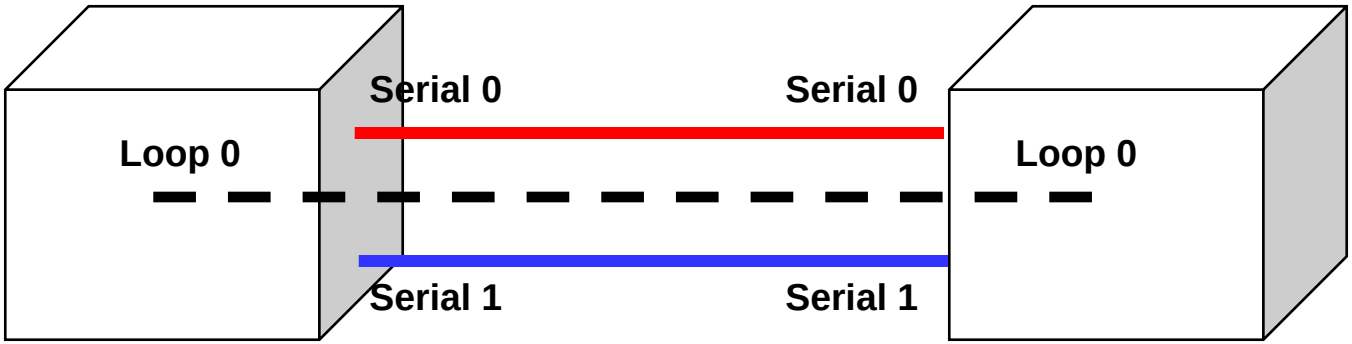

Basic IBGP

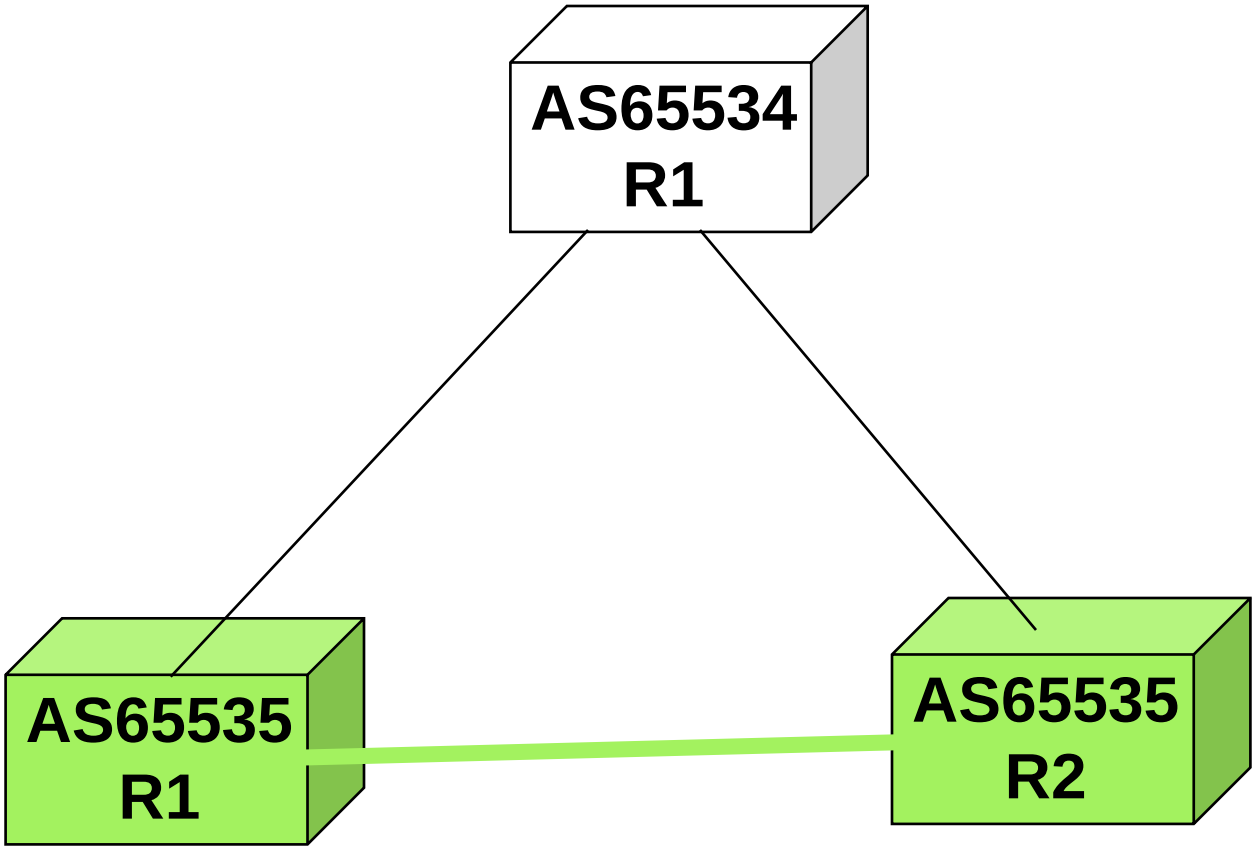

Basic `neighbor` subcommand (iBGP)

`neighbor` *`ip-address`* `remote-as` *`as-number`*

- Identifies peer in our AS
- For iBGP, *`as-number`* is our own
- Will not advertise anything unless `network` statements are defined
- Not needed in single-router configurations

Another Non-BGP Alternative

OSPF Routing Domain

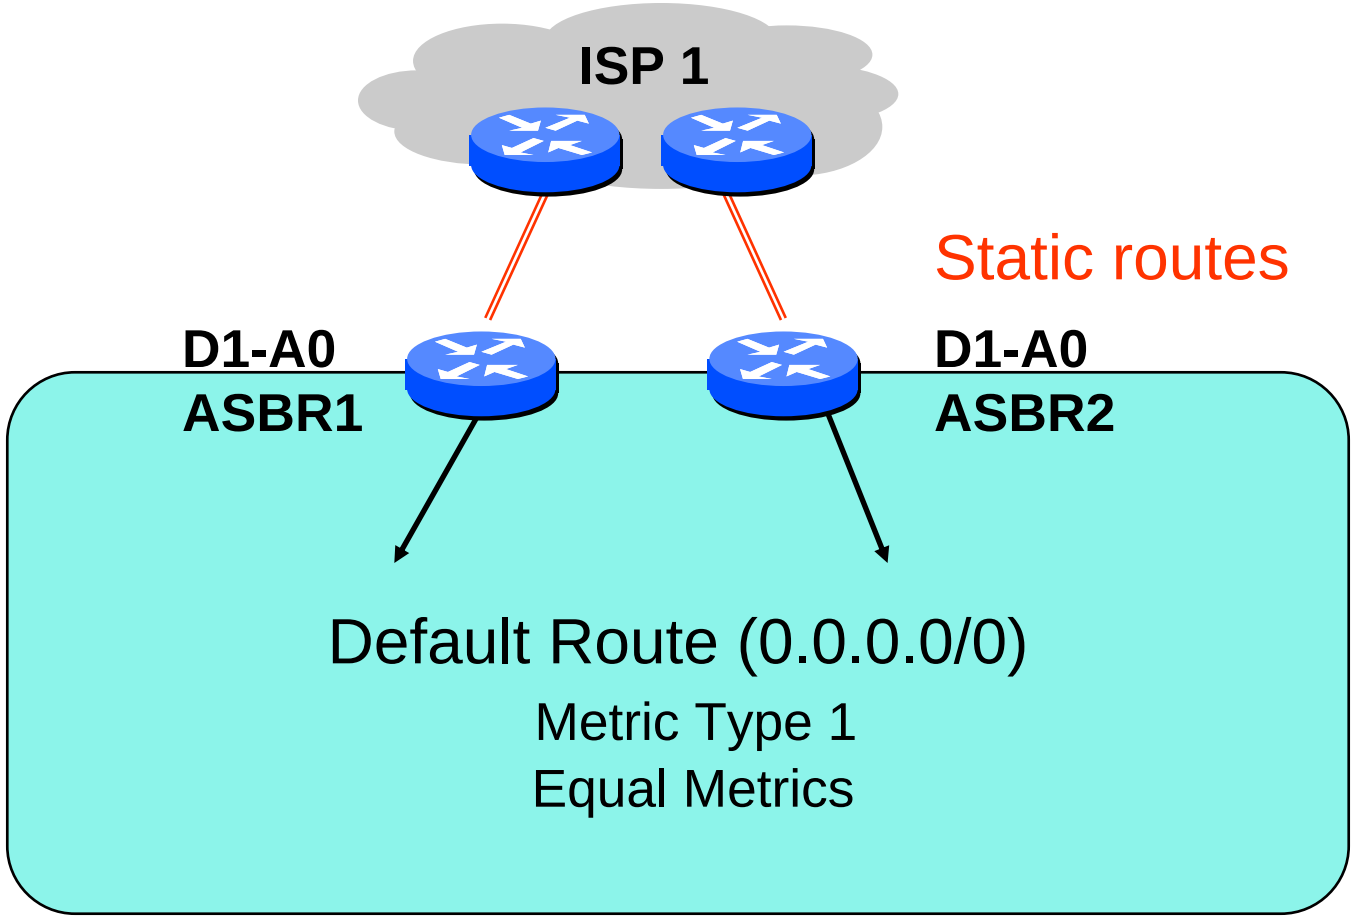

Multiple OSPF Defaults

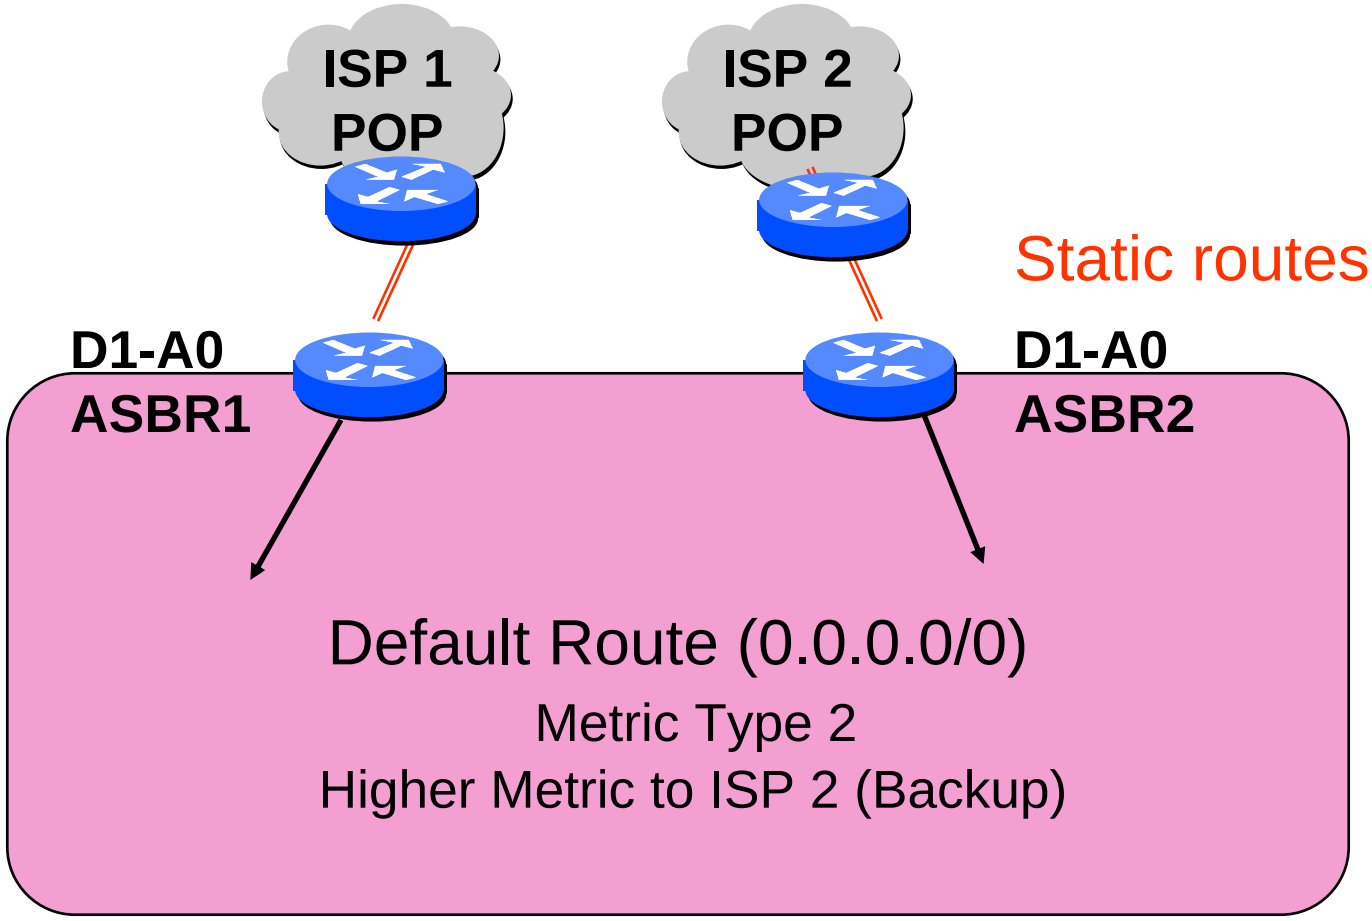

Blackhole Route

- **Establish static route to your block(s)**

```
ip route 1.2.3.4 255.255.240.0 null0
```

- **Redistribute/import into BGP**
- **Suppress more-specific prefix advertising**

Effects of Blackholing

- No route flapping outside your AS
  - If your internal routes go up or down
- Incoming traffic for specific routes that are down
  - Doesn’t match any internal route
  - Automatically discarded without concerning anyone else

BGP Path Selection

Next Hop Access

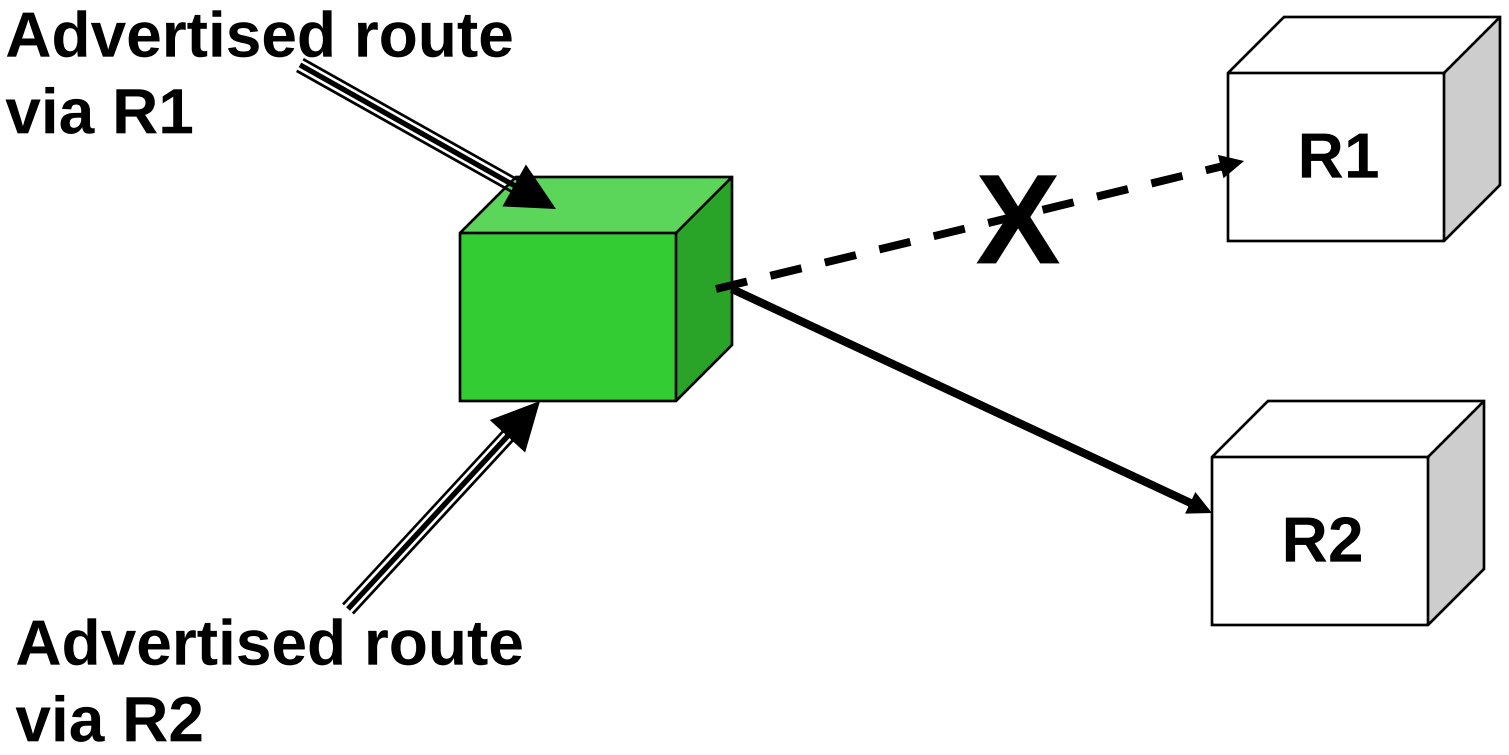

General commands

- **Pattern matching**
- **Action/set**

Scope:

MED vs. Local Preference vs. Weight

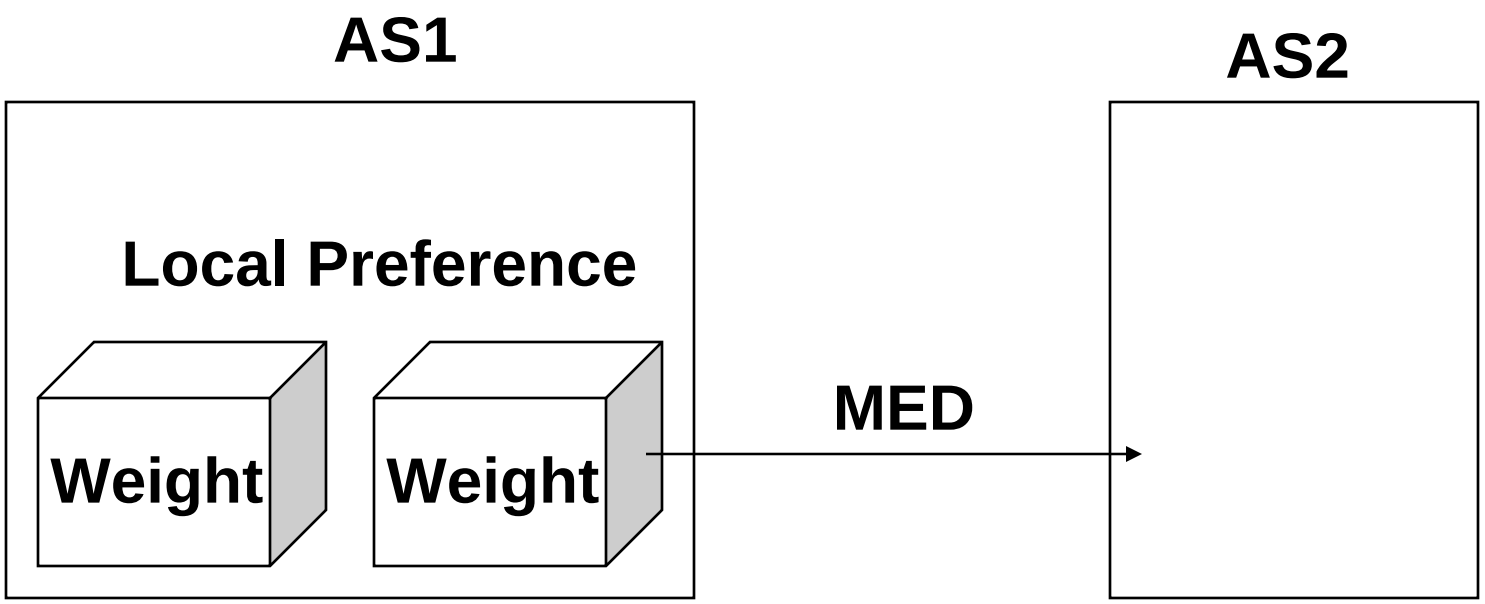

Administrative Weight

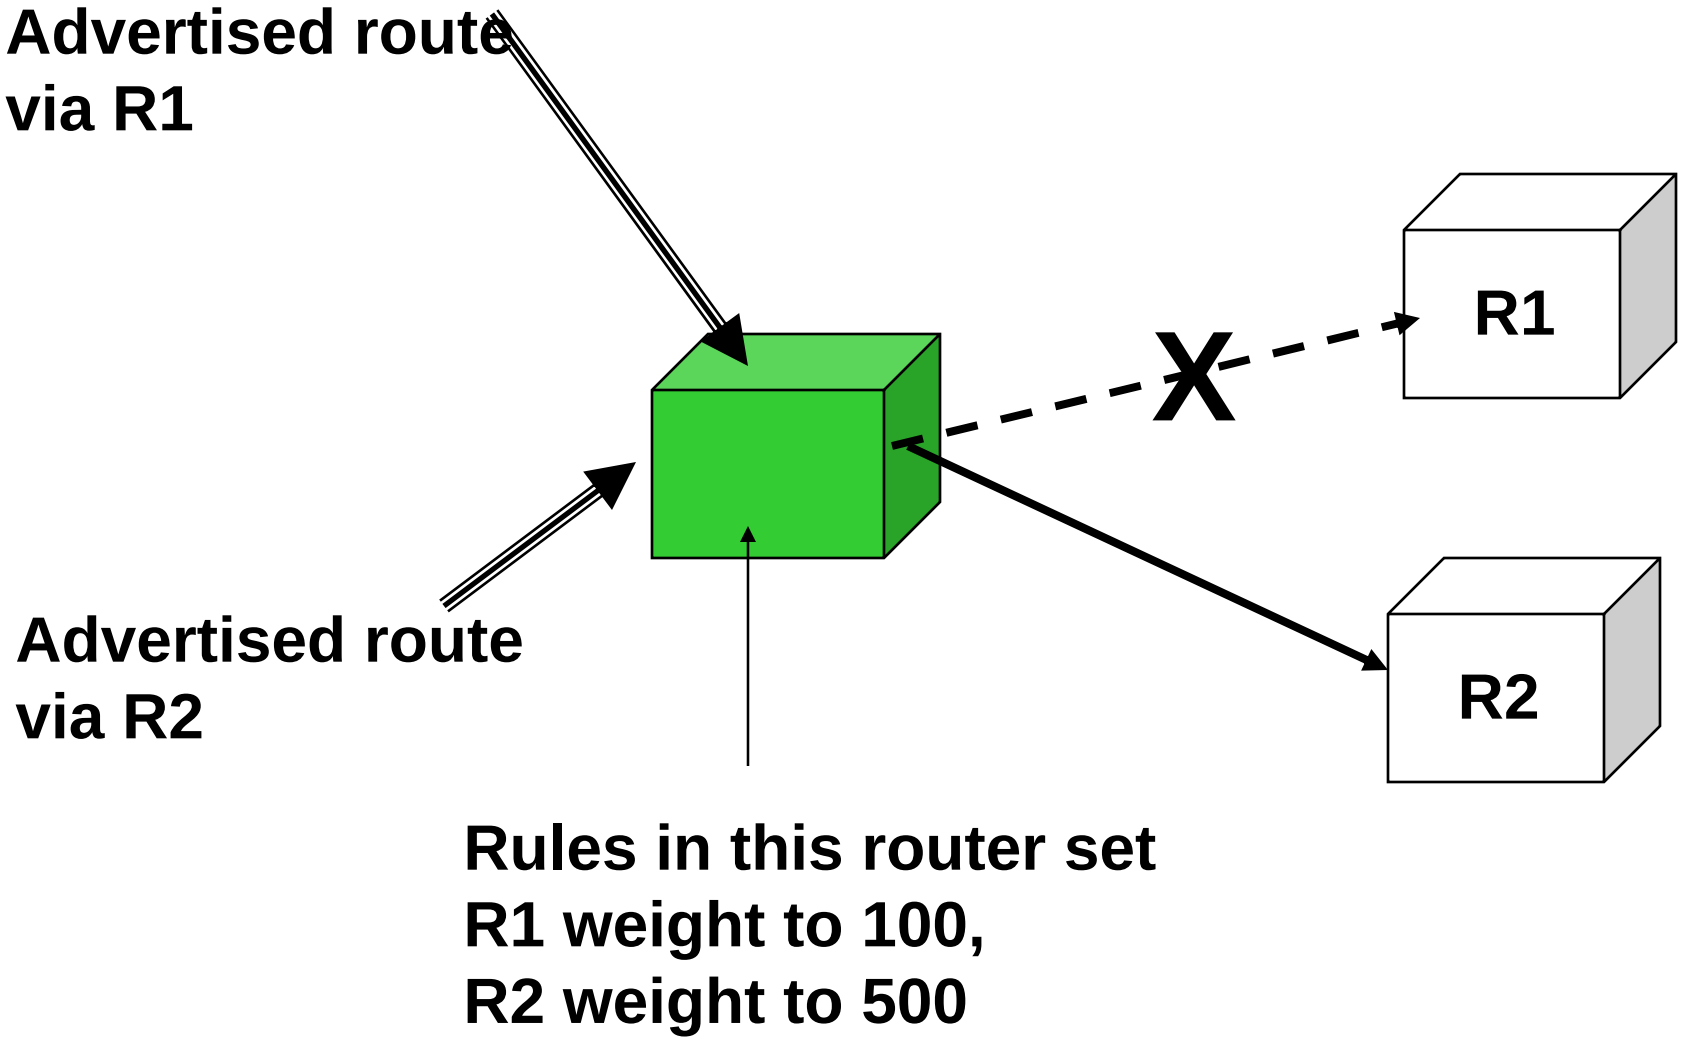

Weight example for load sharing

Default local preference 200

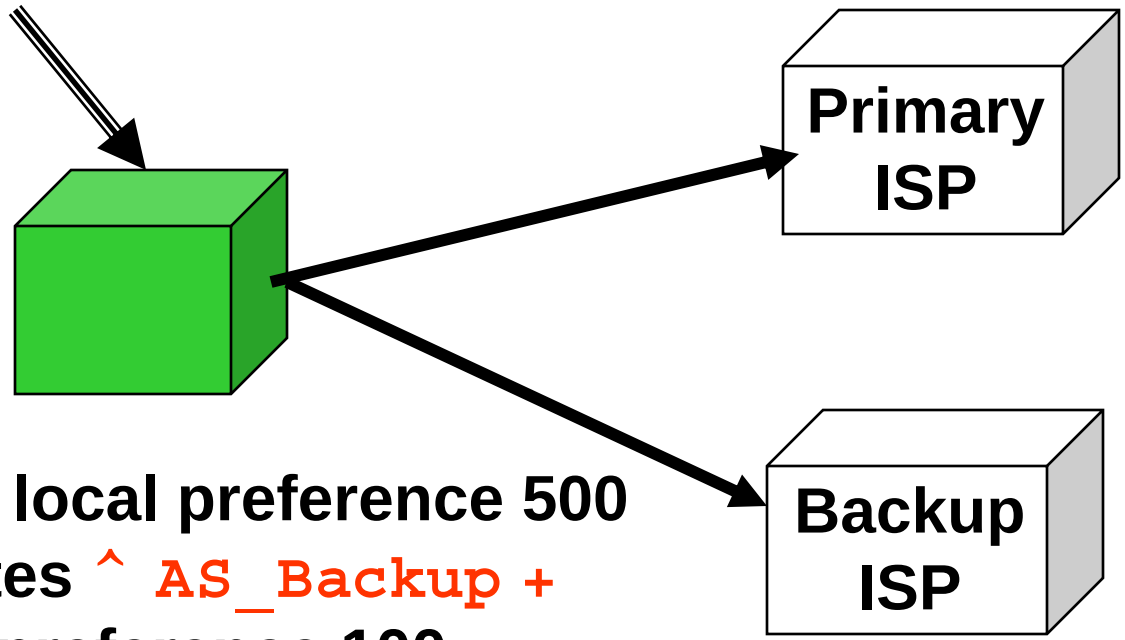

Default local preference 500
All routes ^ `AS_Backup +`
local preference 100

Tiebreaker for Equal Weight:

Local Preference

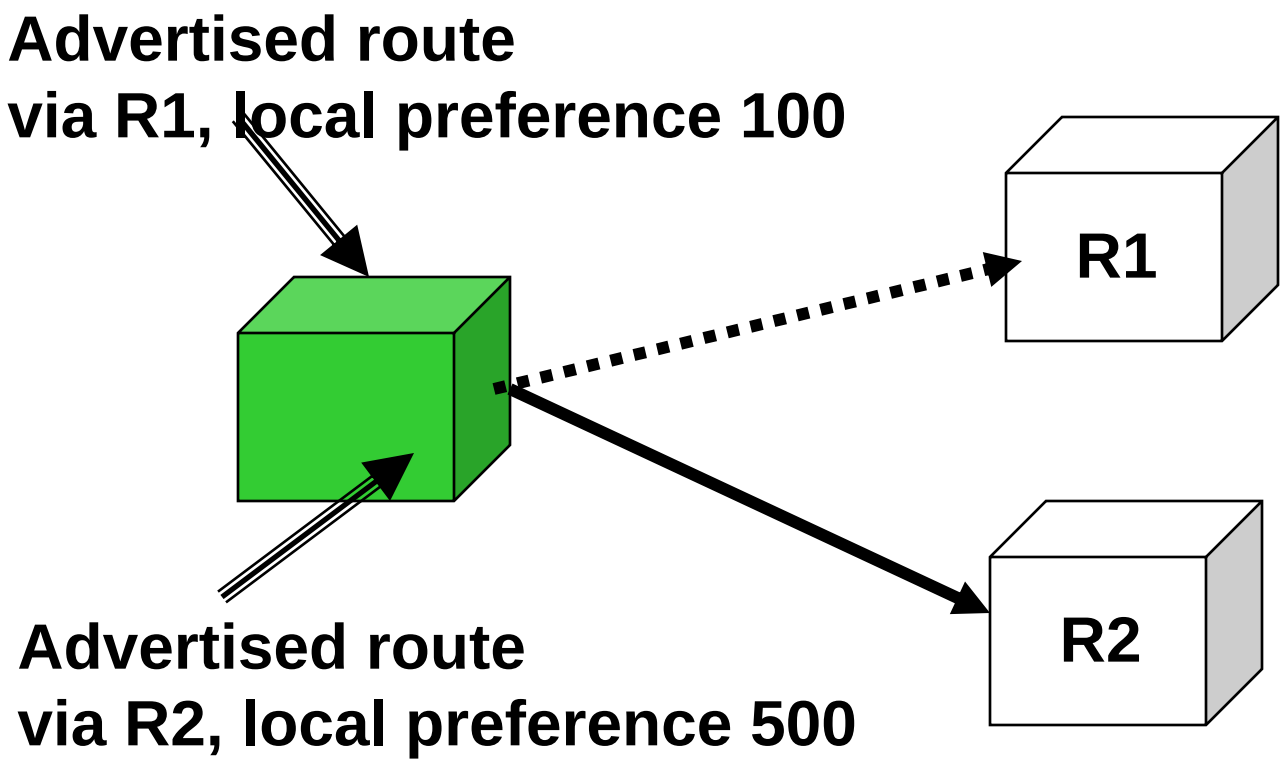

Local Preference example for load sharing

Default local preference 200

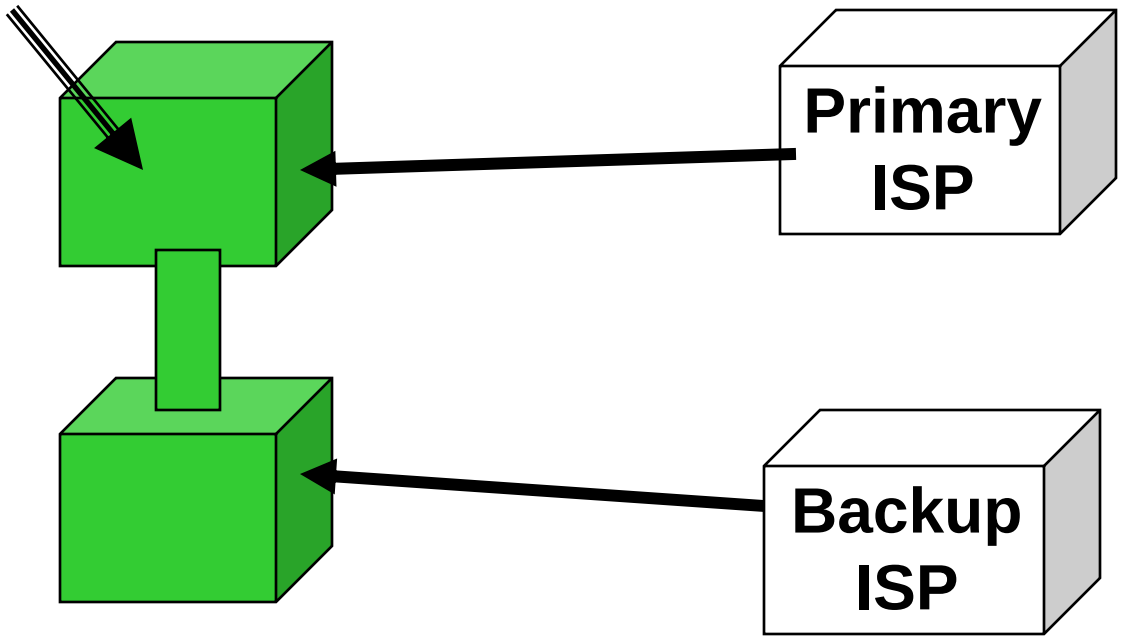

Default local preference 500

All routes ^ `AS_Backup +`

local preference 100

Mutual Backup & Private Peering

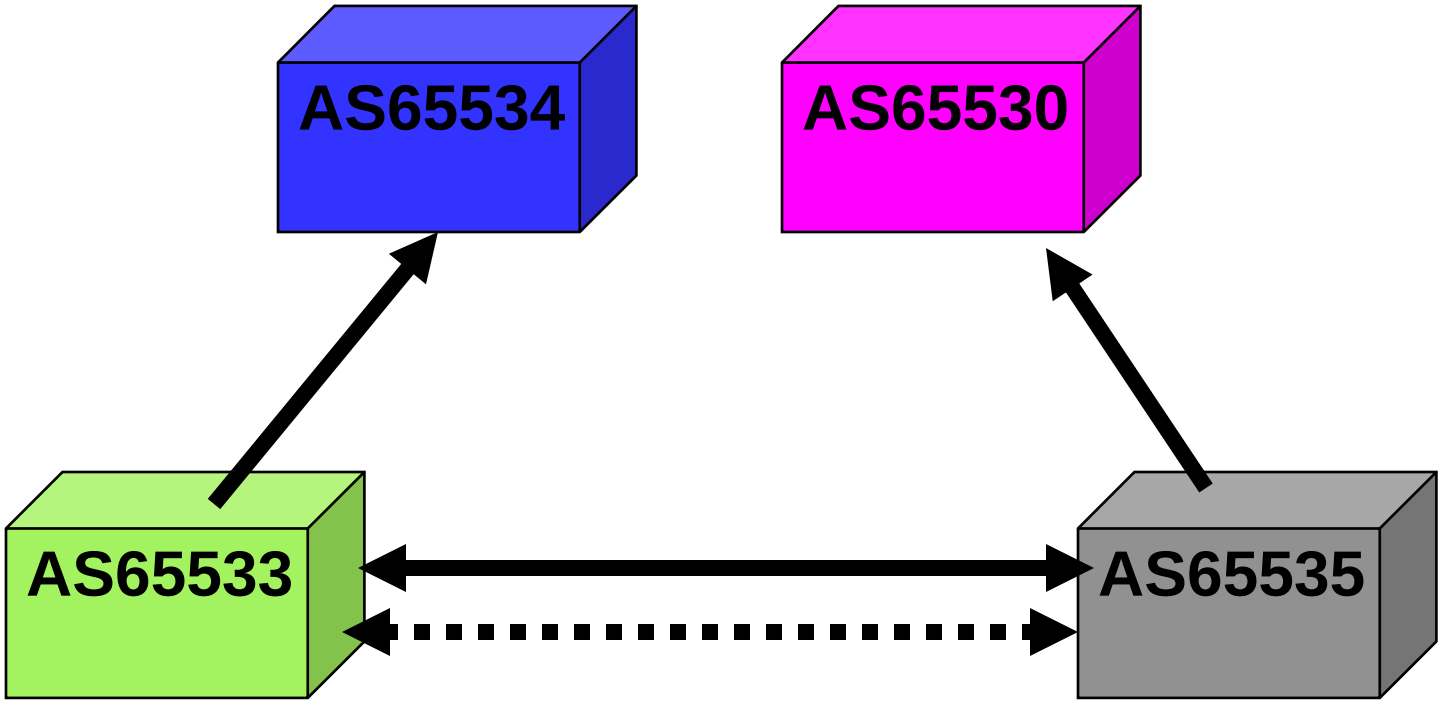

Prefer locally originated routes

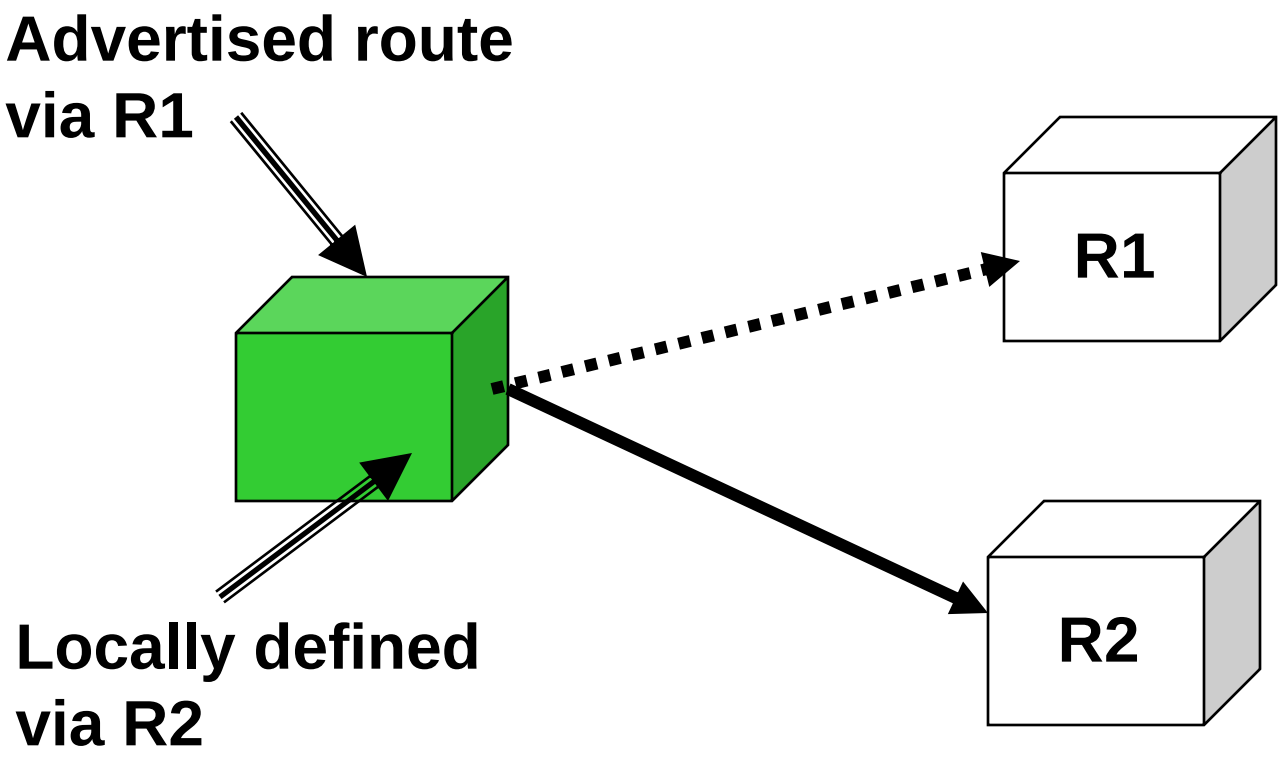

Shortest AS Path

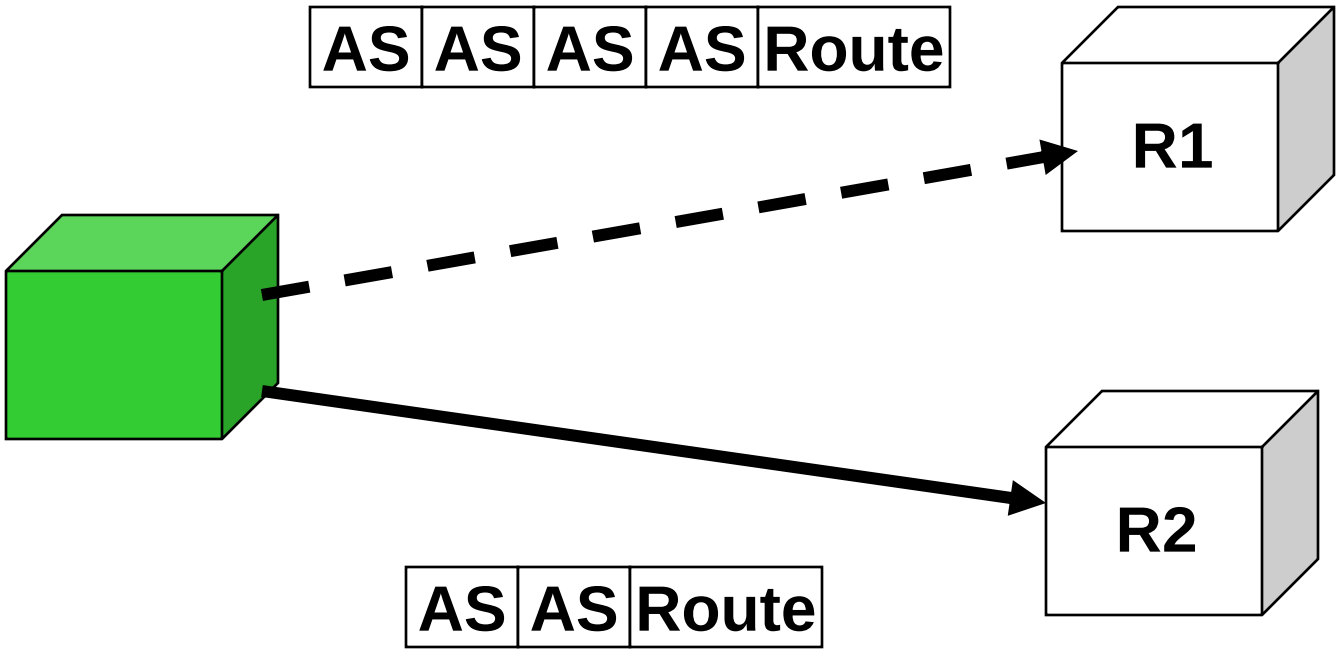

AS Path Prepending

- Applies to routes you advertise
- Makes them less attractive to others
- Increases AS_PATH length
  - your AS put in the path twice

Limitations of Prepending

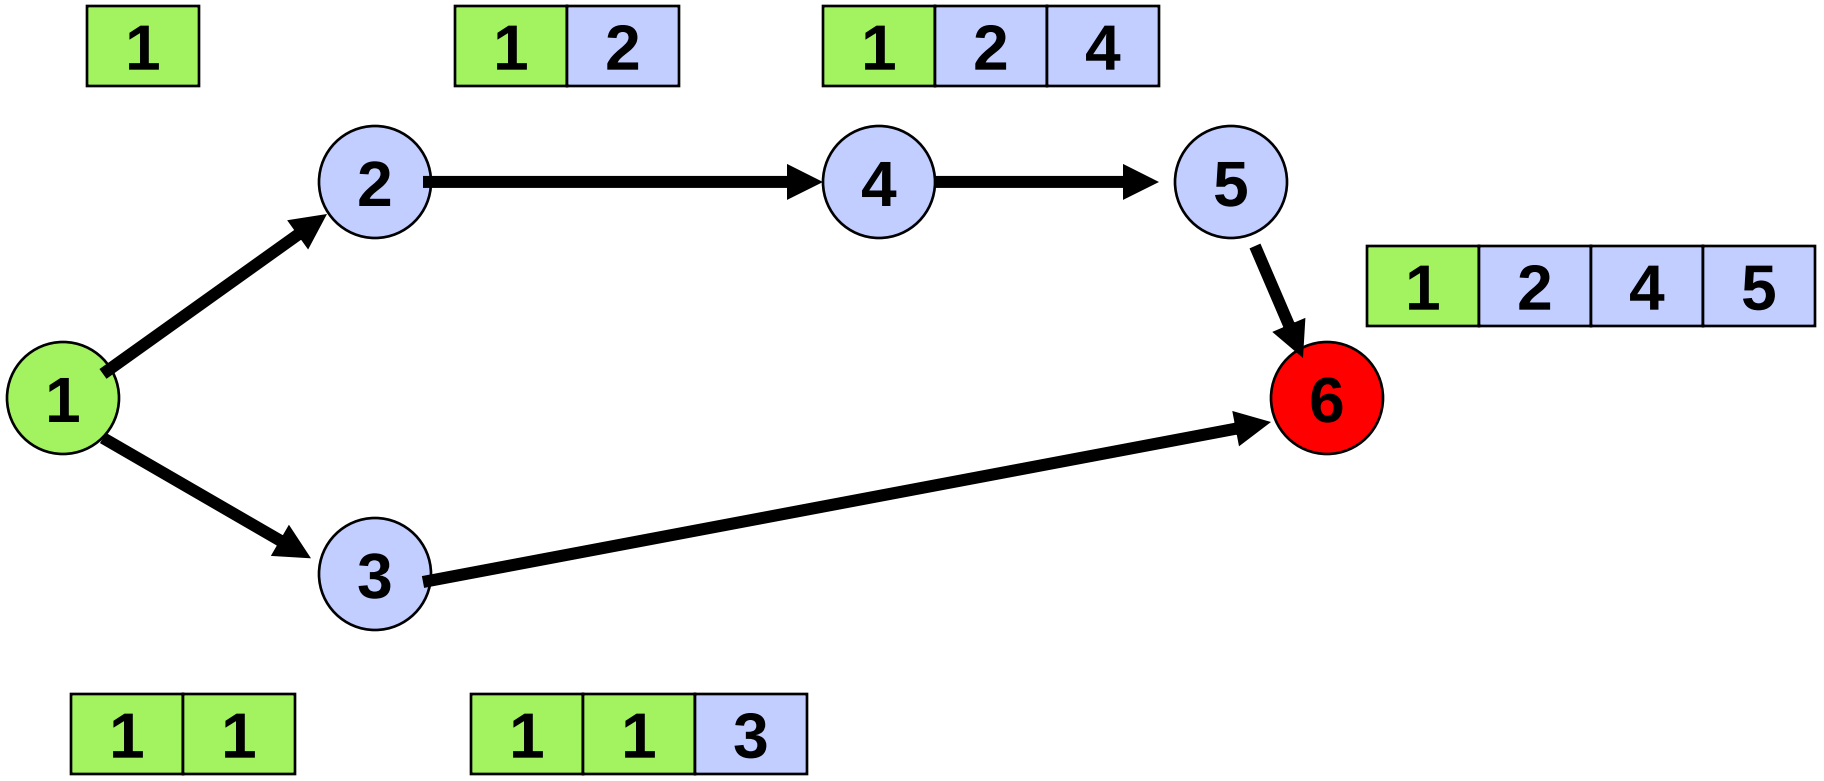

External Paths Preferred

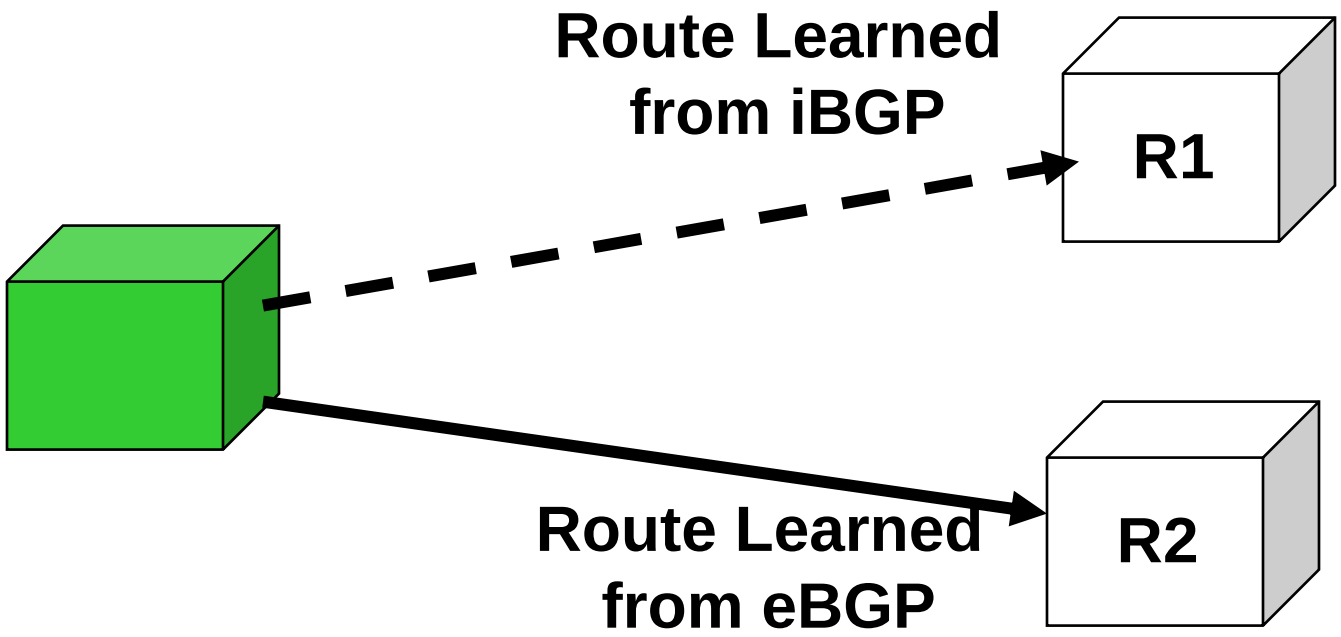

Lowest MED

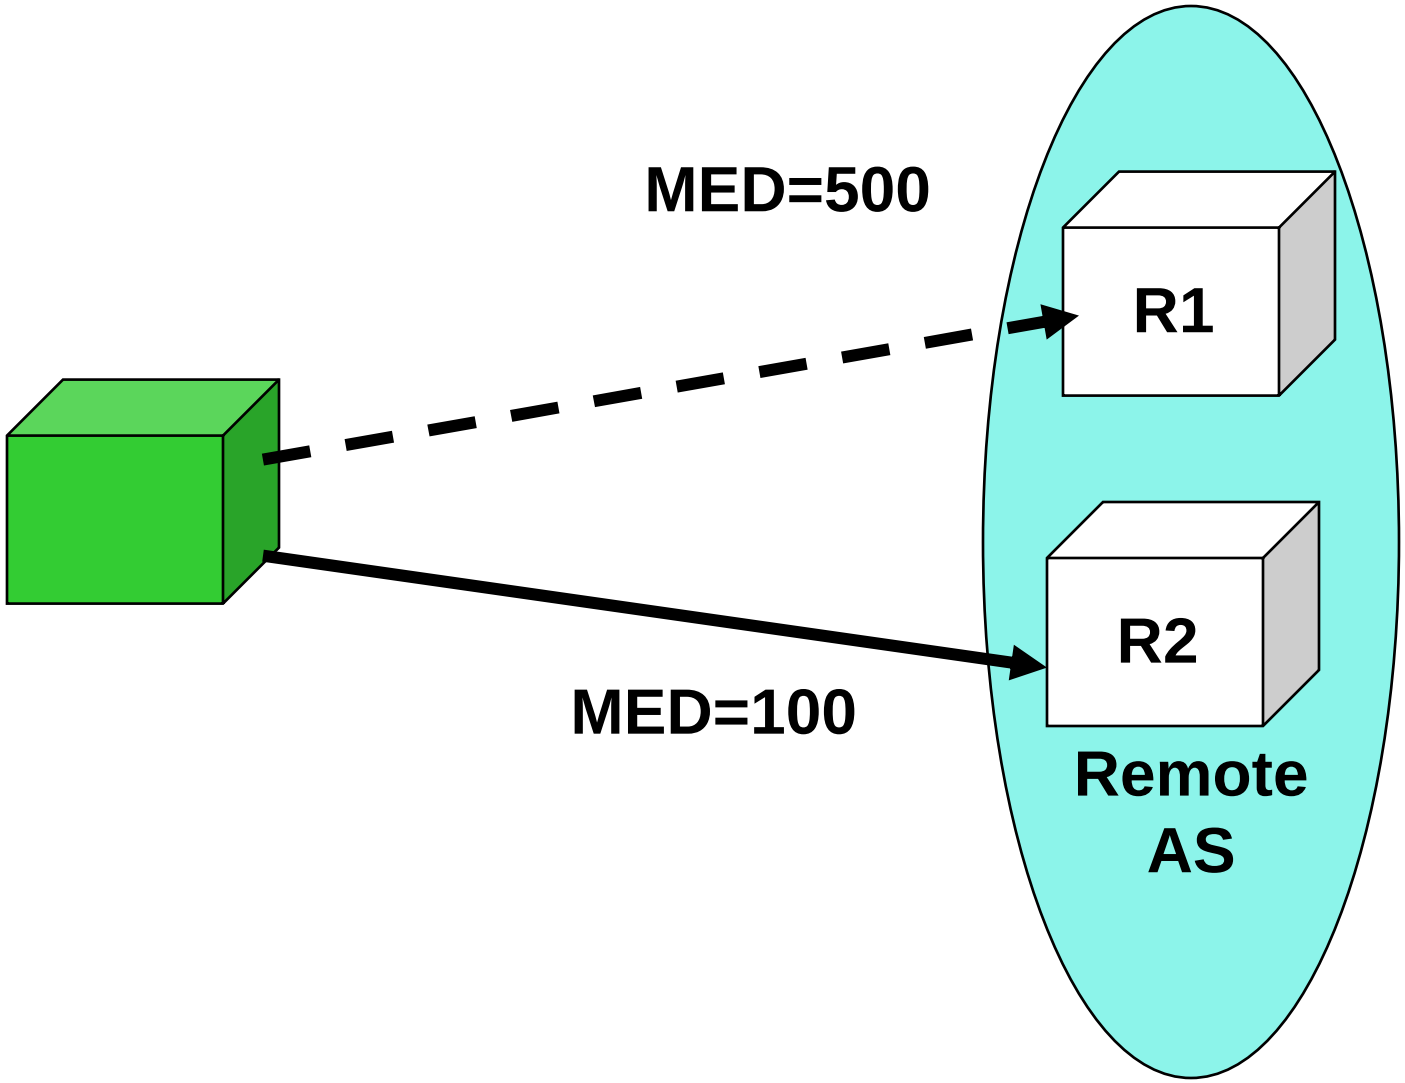

Closest Neighbor

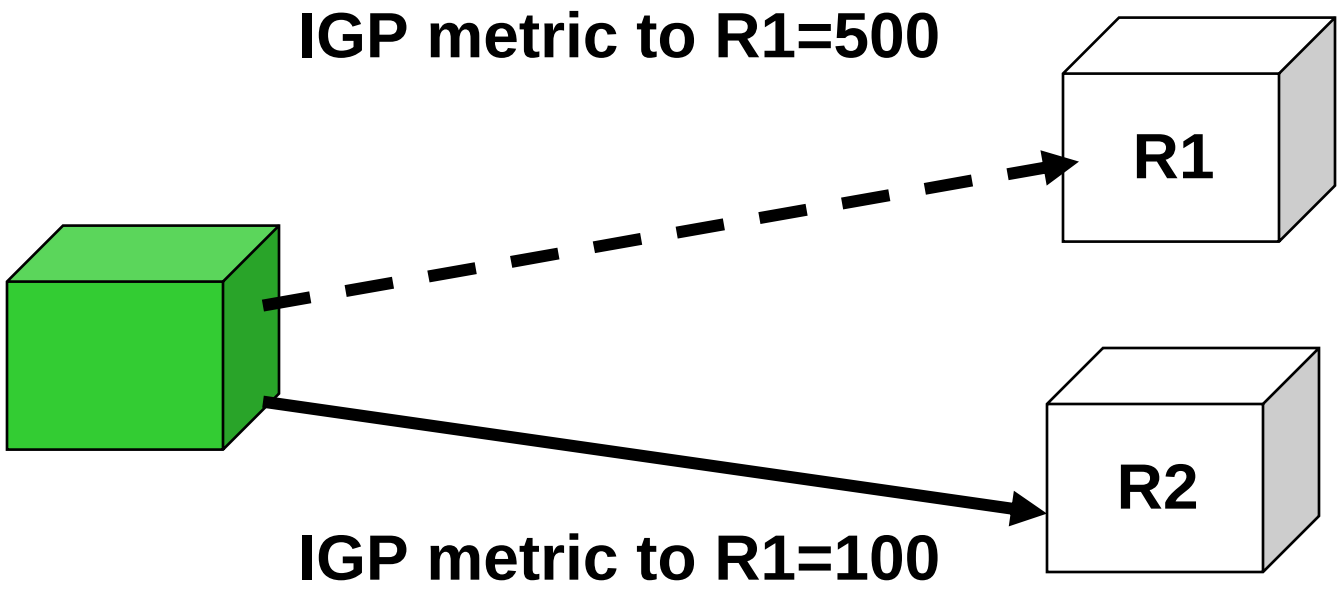

Lowest BGP router ID

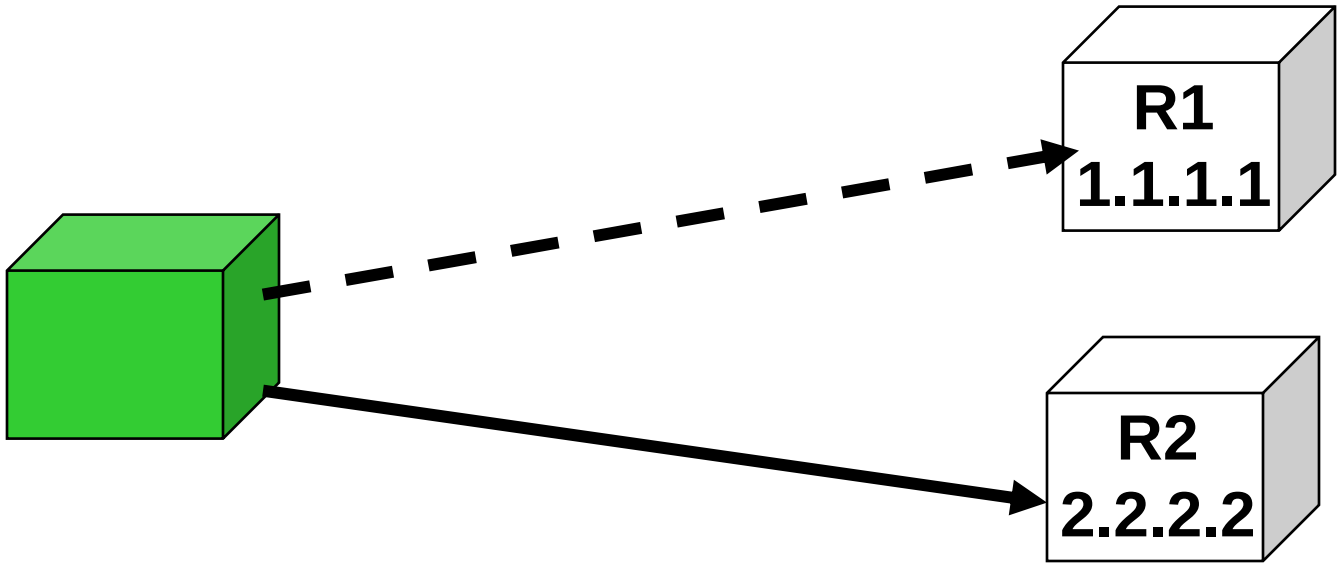

Administrative Distance

- **Preference factors for installing routes in main routing table**
  - **Cisco calls it administrative distance**
- **Will affect which sources of routing are installed**
- **Since we advertise only internally reachable routes**
  - **This affects those advertised**

Permit/Deny Filtering

- Routes can be blocked or permitted to enter BGP
- Criteria include:
  - Advertising address
  - Route address being advertised
  - AS_PATH to reach route
    - » Can use UNIX regular expressions

An ISP Topology

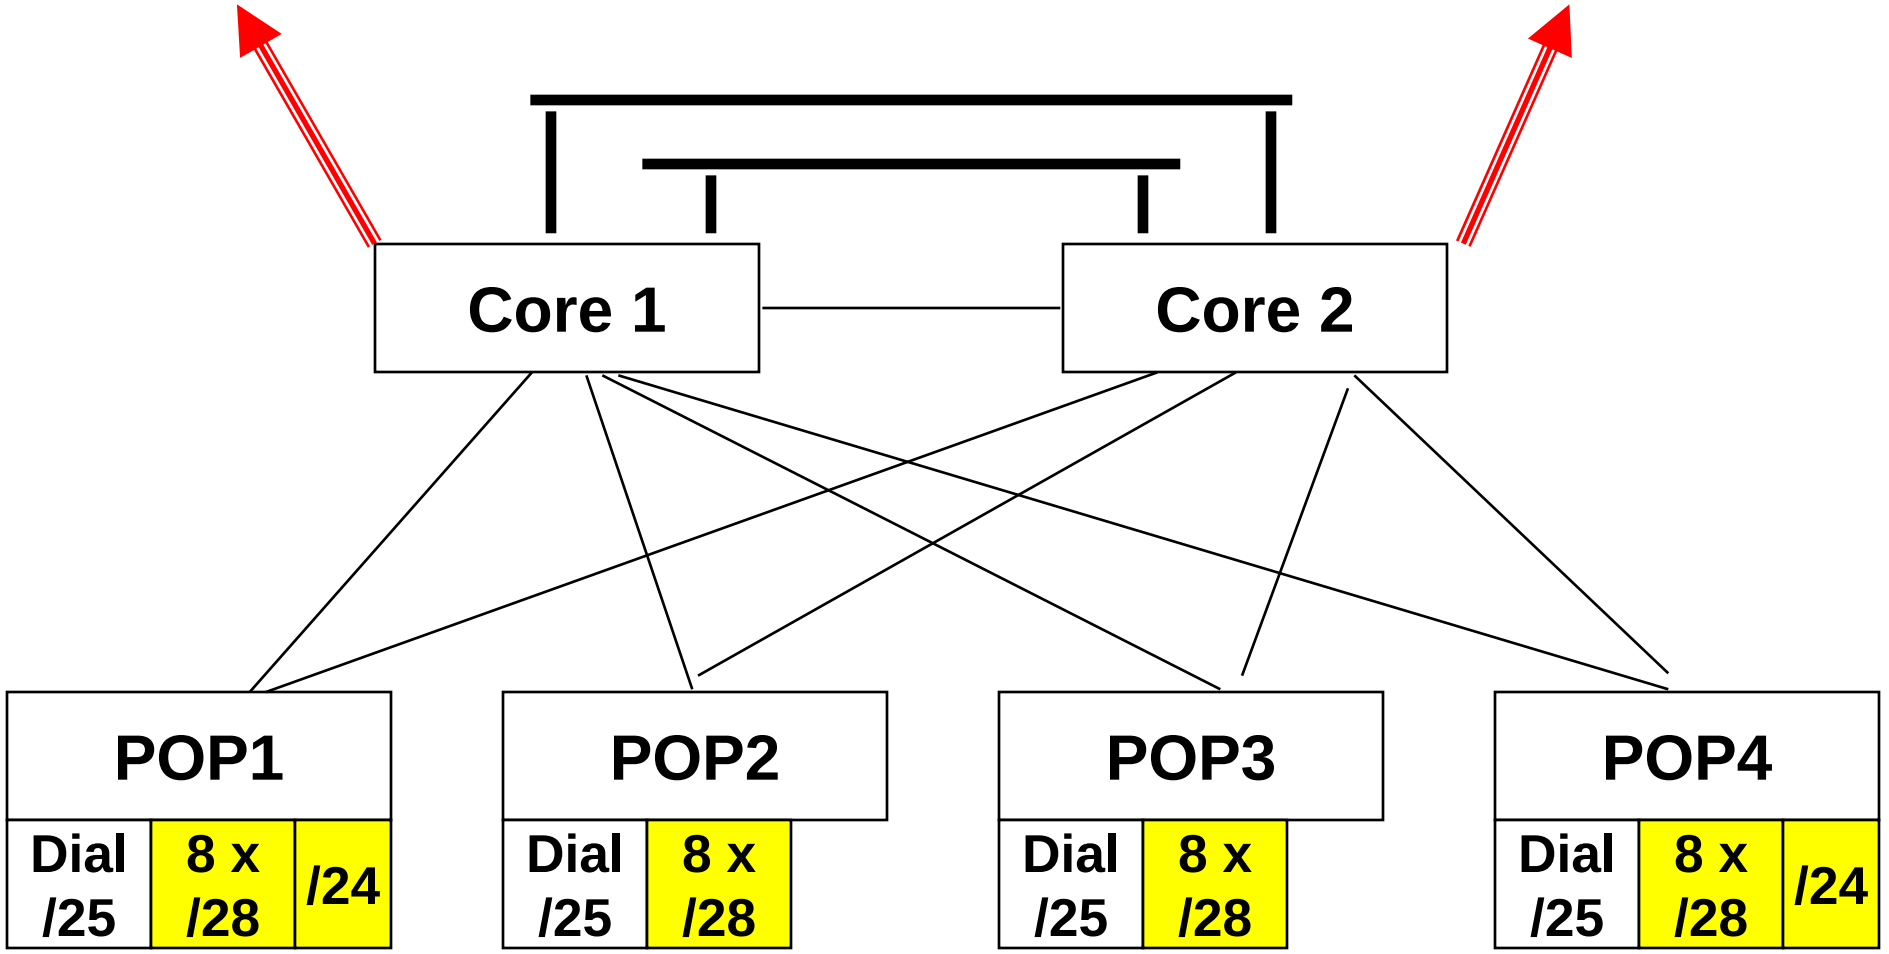

An ISP Address Plan

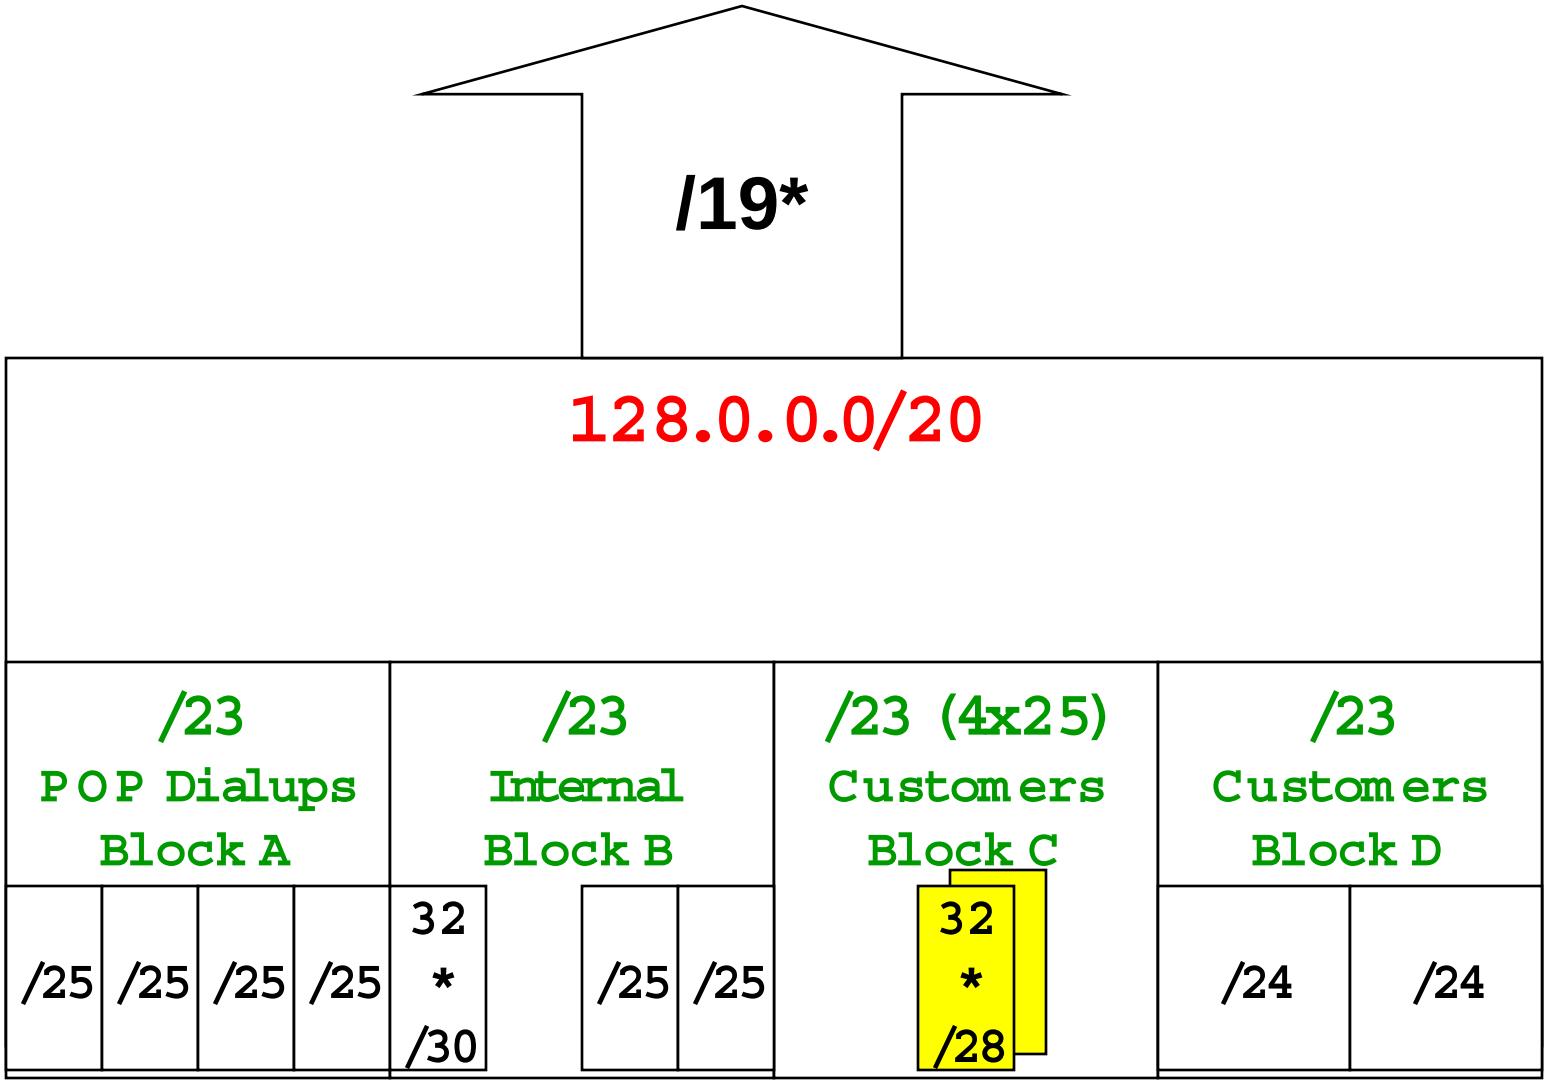

Aggregation is better than Aggravation

- **Blackhole routes for your blocks**
  - **Avoid more-specifics**
  - **Use NO-EXPORT when controlling load to upstream**
- **Encourage customers to aggregate**
  - **Proxy aggregation hard to administer**
- **Understand which blocks you can advertise**
  - **And do ingress/egress filtering**

Origination vs. Advertising

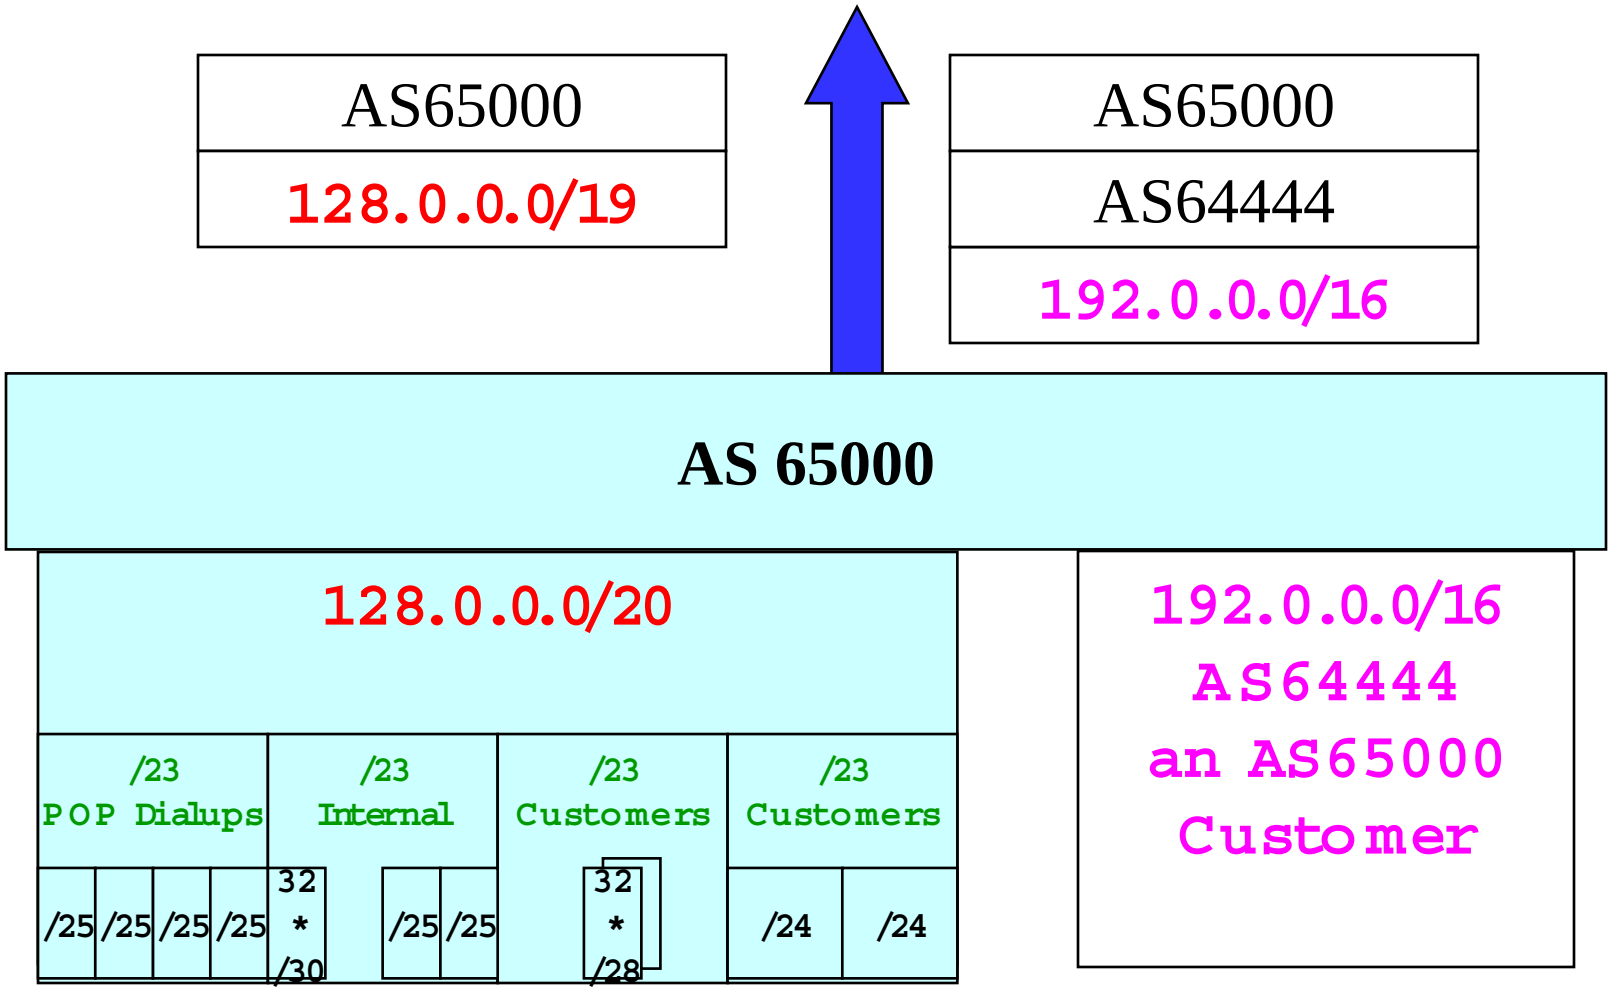

Aggregating your Own Traffic

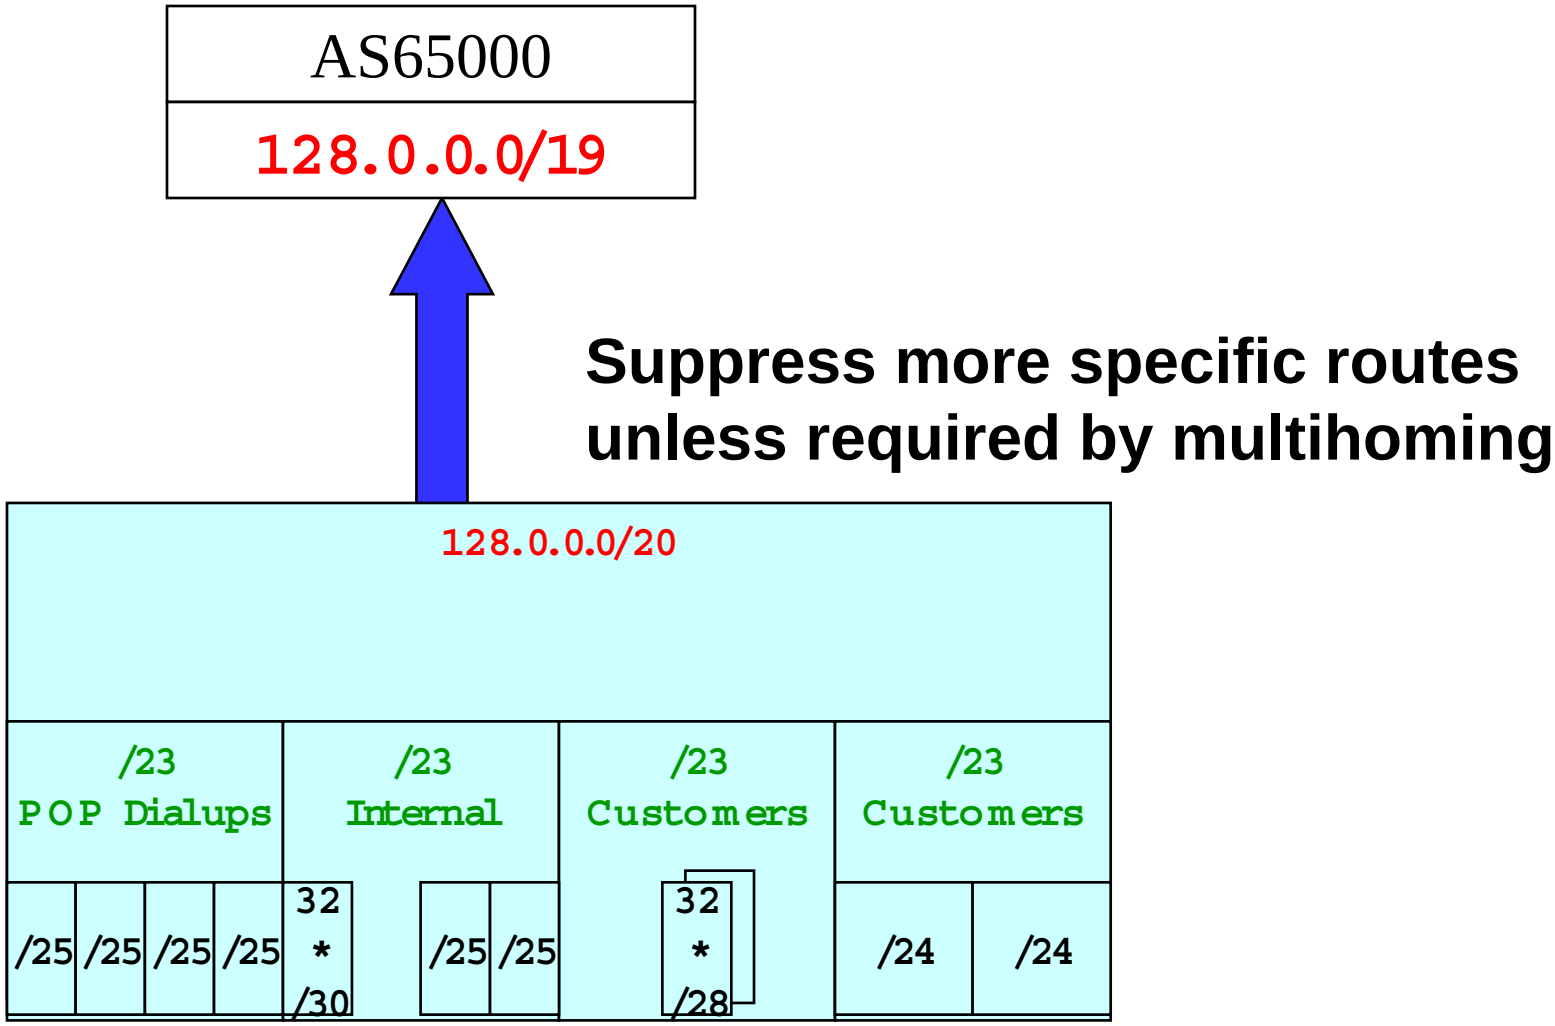

Advertising with NO-EXPORT