

# OSPF Goodies for ISPs

**North American Network Operators Group  
(NANOG)**

**October 1999 meeting**

**Howard C. Berkowitz**

**Gett Communications**

**[hcb@clark.net](mailto:hcb@clark.net)**

**(703)998-5819**

## The Past

- **Version 1**
  - RFC 1131 IETF Proposed Standard
- **Version 2**
  - RFC 1247 IETF Proposed Standard
  - RFC 1583 IETF Draft Standard
- **RFC 1812 requires OSPF for dynamic routing**
- **And things progressed...**

## The Near Term


- **RFC 2178**
  - **Still at Draft Standard level; still Version 2**
  - **Enhancements**
    - **NSSA**
    - **Demand Circuits**
    - **OSPF over Frame**
    - **Cryptographic authentication**
  - **TOS routing removed**
- **RFC 2328**
  - **Full Standard**

## The Standards Future

- **OSPF Version 6**
  - **Originally intended for V6 only**
  - **Made dual capable for V4 and V6, and possibly other protocols such as IPX**
  - **Generally upward compatible**
    - **better scalability for Multicast OSPF**
    - **quality of service routing**
  - **Internet Draft**

# OSPF and Single Areas

## An Area has...

- **An area ID**  **0.0.0.0 format recommended!**
- **A set of address prefixes**
  - Do not have to be contiguous
  - So a prefix can be in only one area
- **A set of router IDs**
  - Router functions may be interior, inter-area, or external

## Areas and Performance

- **General workload for routing:**

$O(\text{Prefixes})$

- **In general area<sub>k</sub> structure, workload per router:**

*$O(x)$ : "on the order of  $x$ "*

$O(\text{Prefixes}_k +$

$\text{InjectedExternals}_k)$

- **In OSPF area<sub>k</sub>, workload per router:**

$O((\text{Prefixes}_k +$

$\text{InjectedExternals}_k)) +$

$\log(\text{Routers}_k)$

## Area Sizing Guidelines

- **Rules of thumb for non-backbone area**
  - No more than 100 routers
  - No more than 50 neighbors per router
- **Decrease when media unstable**
  - Consider static/default and demand techniques
- **Decrease when large numbers of externals injected**
  - Consider if the incoming externals can be summarized or filtered



## When Might Single-Area OSPF make sense?

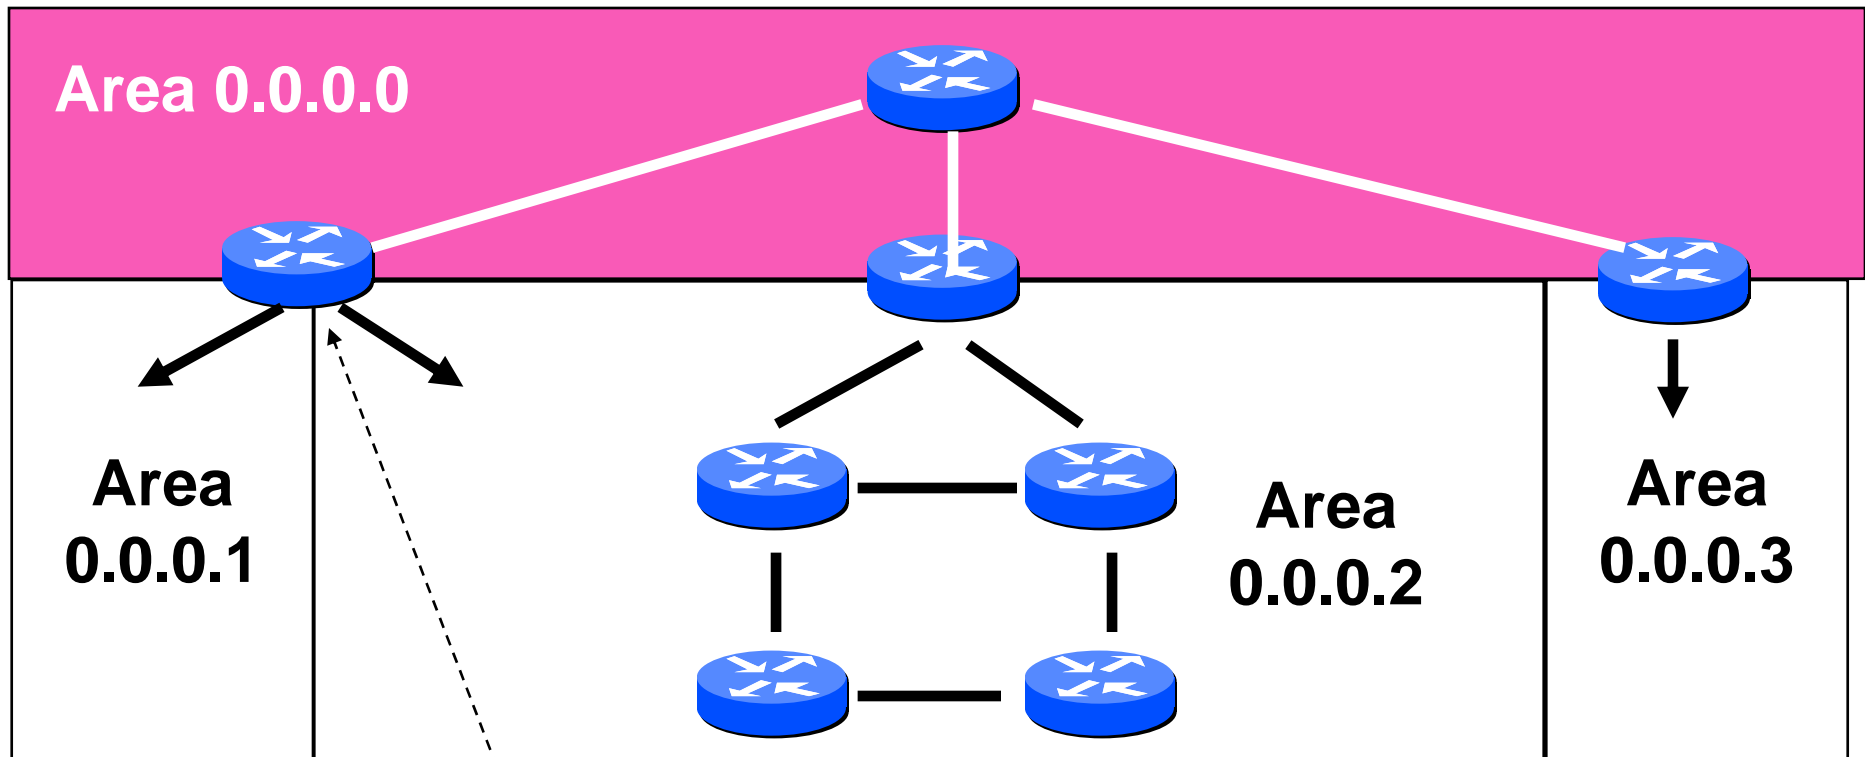
- **Fewer than 50 routers with alternate paths**
- **Needs:**
  - multivendor compatibility
  - fast convergence
  - VLSM
  - complex defaults and externals
- **No clear candidates for core**
  - OSPF power greatest with hierarchy
  - Multiple domains may be better than 1 area
  - EIGRP a possible alternative

## How Many Areas?

- **Map communities of interest (COI) to areas**
- **Add COI until number of routers exceeded**
- **Avoid >2-3 areas per ABR unless very stable**
- **Watch CPU loading in ABRs**
- **Numbering easier if #areas is power of 2**

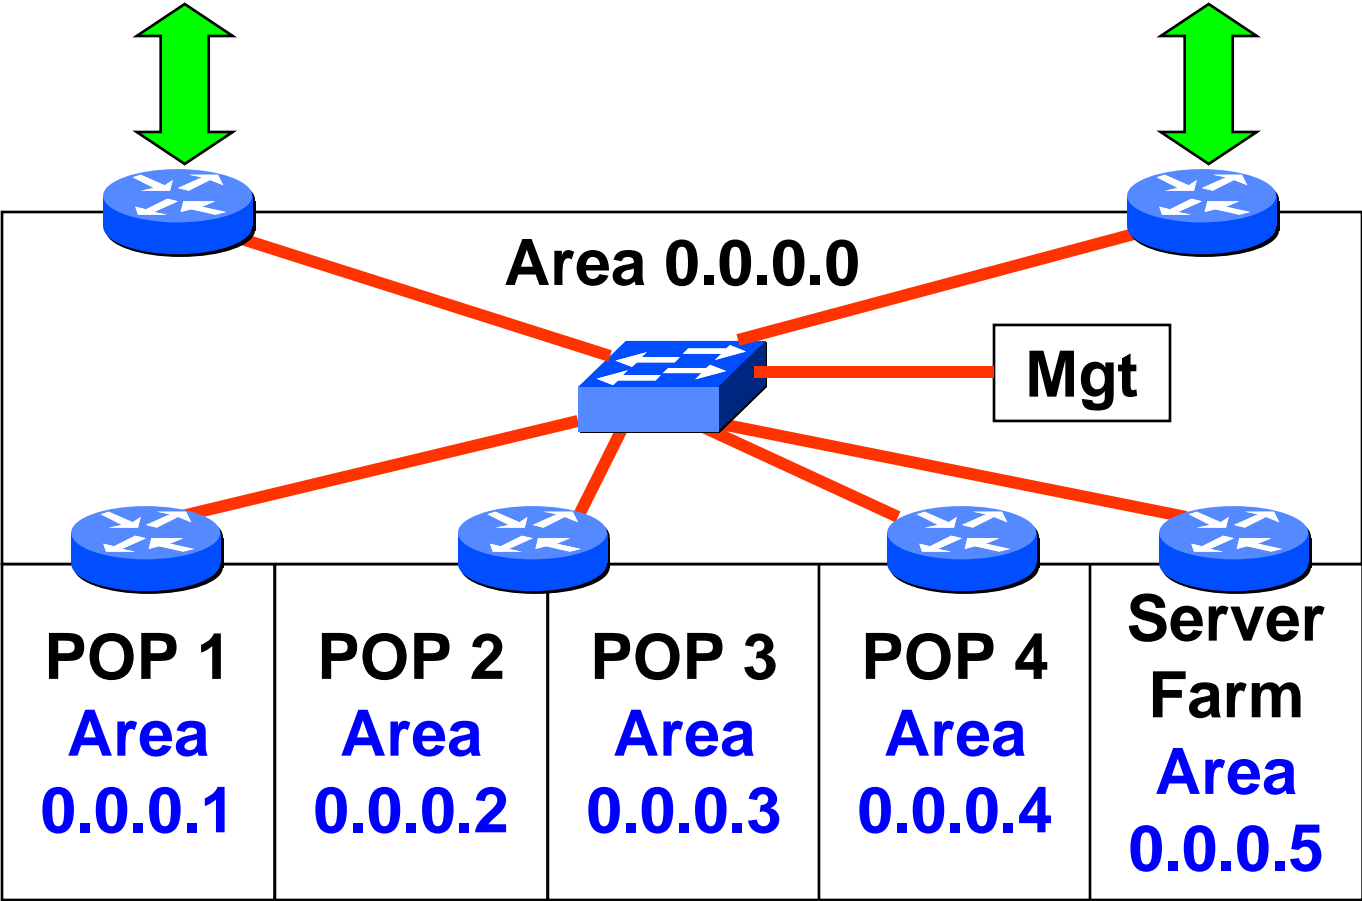
## **Single-Domain Inter-Area**

## Basic Inter-Area




**Note ABR servicing multiple nonzero areas.  
Rule of thumb -- not more than 3 nonzero areas**

# ISP OSPF 1: Implementation



# The Right Reason to Break Hierarchy Traffic management, not “redundancy”




- OSPF flow for general traffic and backup
- Static route with  $AD < 110$  for heavy traffic

## Summarization/Aggregation

- **Summarization**
  - Reduces inter-area route workload
  - Configure manually on ABR
- **Aggregation**
  - Reduces external route workload
  - Configure manually on ASBR
  - Complements stubby areas


## Multiple ABR



Both ABRs advertise the same summary when summarization is configured



# Partitioned Nonzero Areas




Both ABRs still advertise the same summary when summarization is configured. Ooops!

## **External Information**

## Externals are a Good Thing

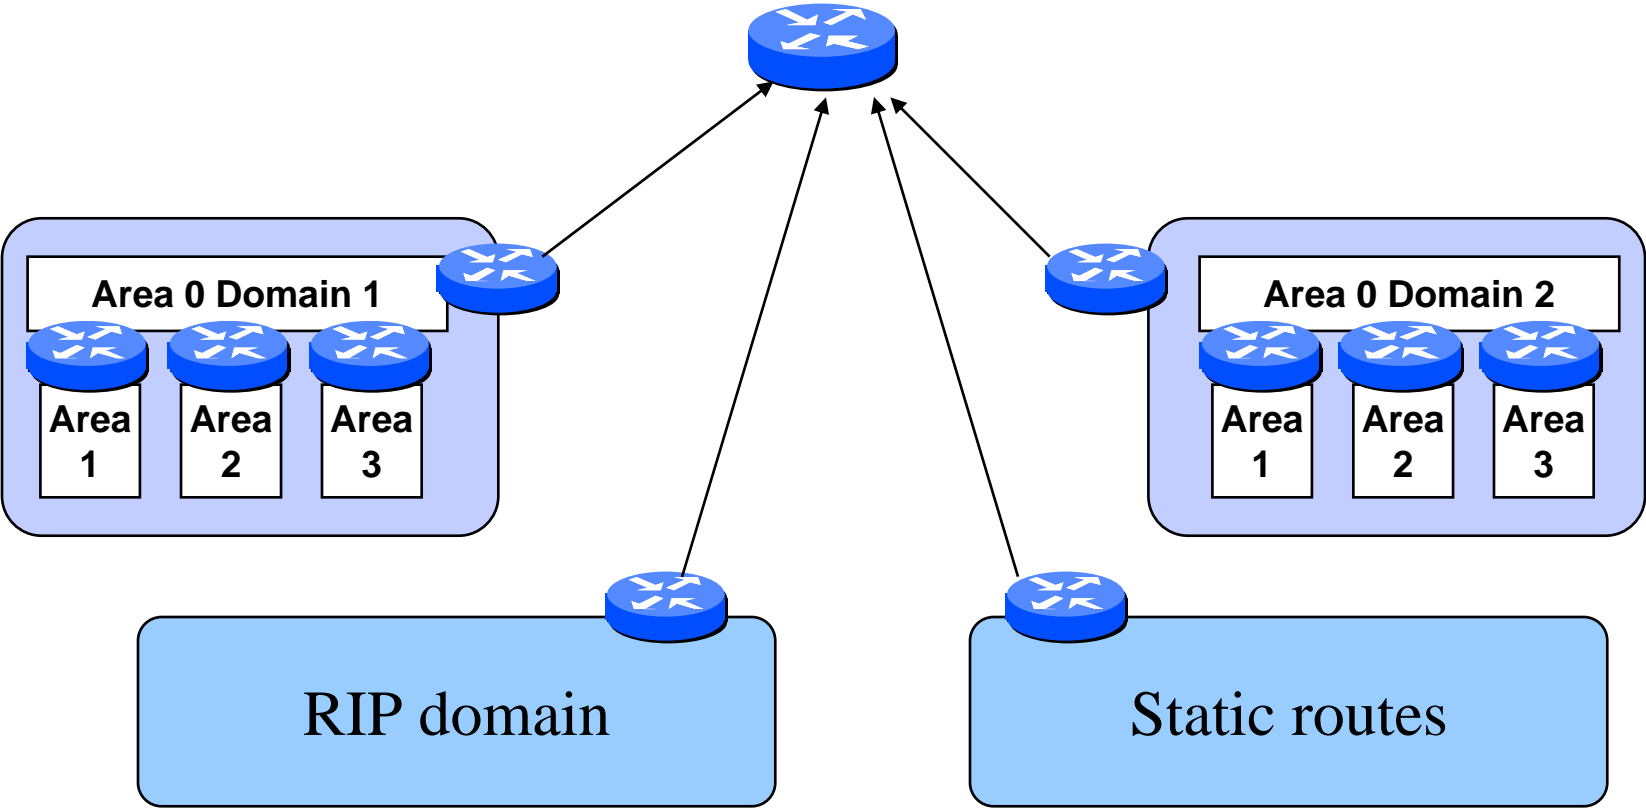
- **A way to learn about things outside your local system**
- **Allow workarounds to some awkward configurations**
- **Sources of externals include:**
  - **Other OSPF domains**
  - **Other IGPs: EIGRP, IGRP, RIP, IS-IS**
  - **BGP-4**
  - **Static routes**

# Multiple Homogeneous Domains



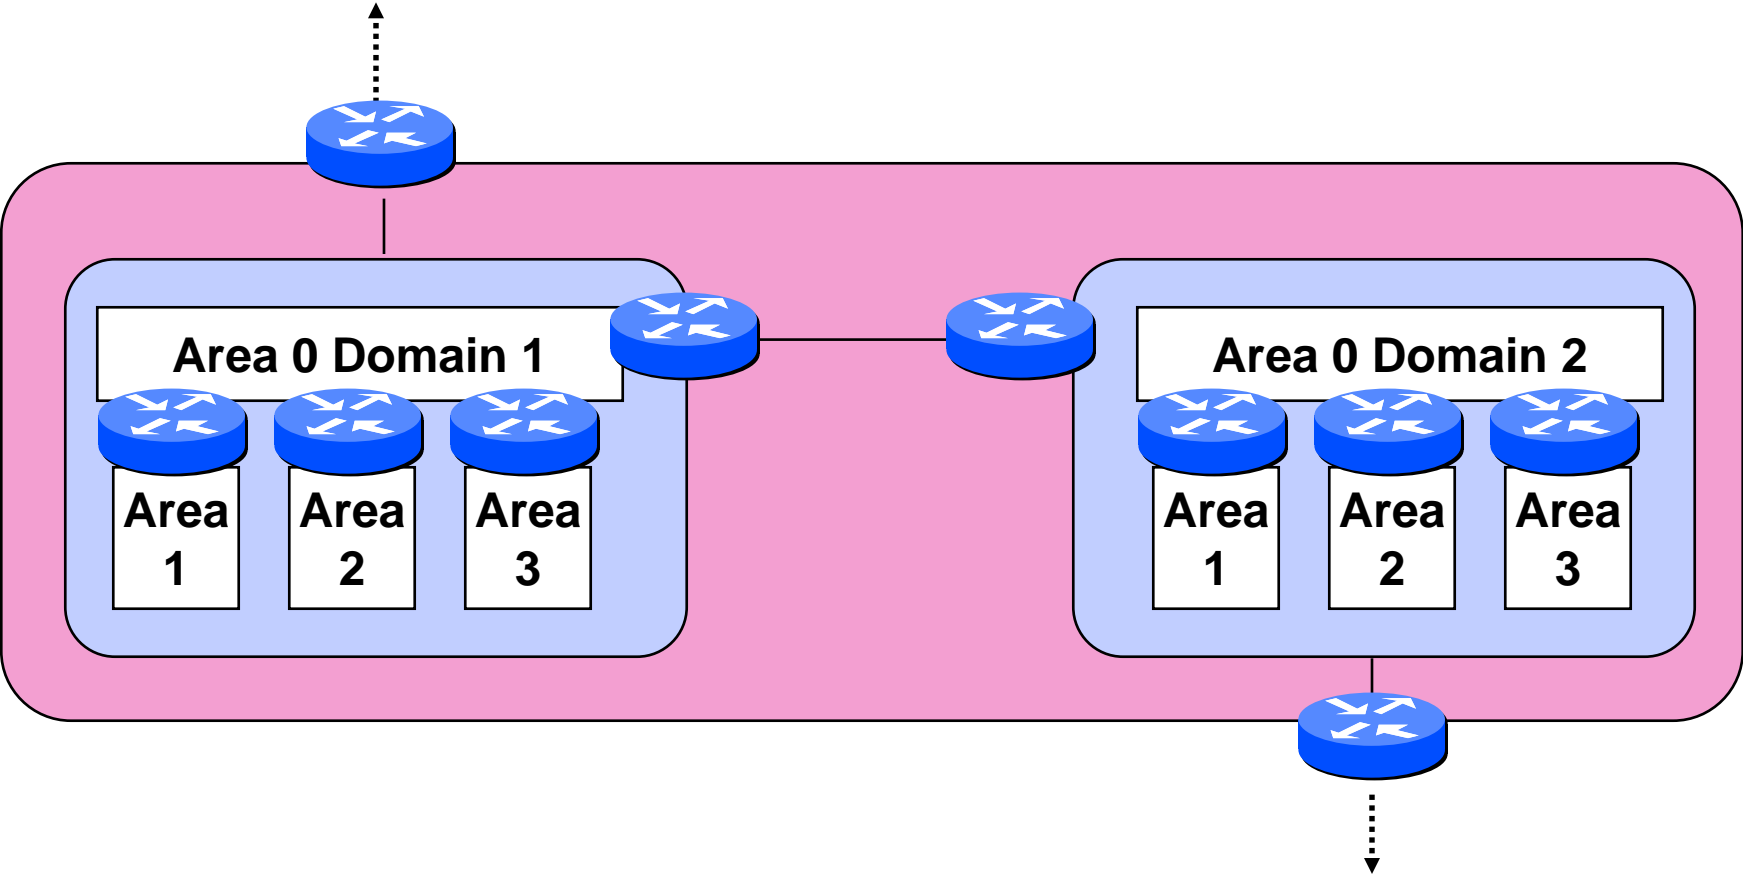
Multiple OSPF copies with different process ID

# Multiple Heterogeneous Domains



Note hierarchy!

# Multiple Autonomous Systems



BGP becomes involved

## Externals and Aggregation 1

- **A full ISP routing table has approximately 62,000 routes**
  - **But will you do anything differently if you know all of them and have a single ISP?**
  - **Multiple ISP situations call for complex OSPF and BGP design**
- **Never redistribute IGPs into BGP**
- **Restribute BGP into IGPs with extreme care**

**Question:**


**How many of your customers expect  
“automagic” configuration?**



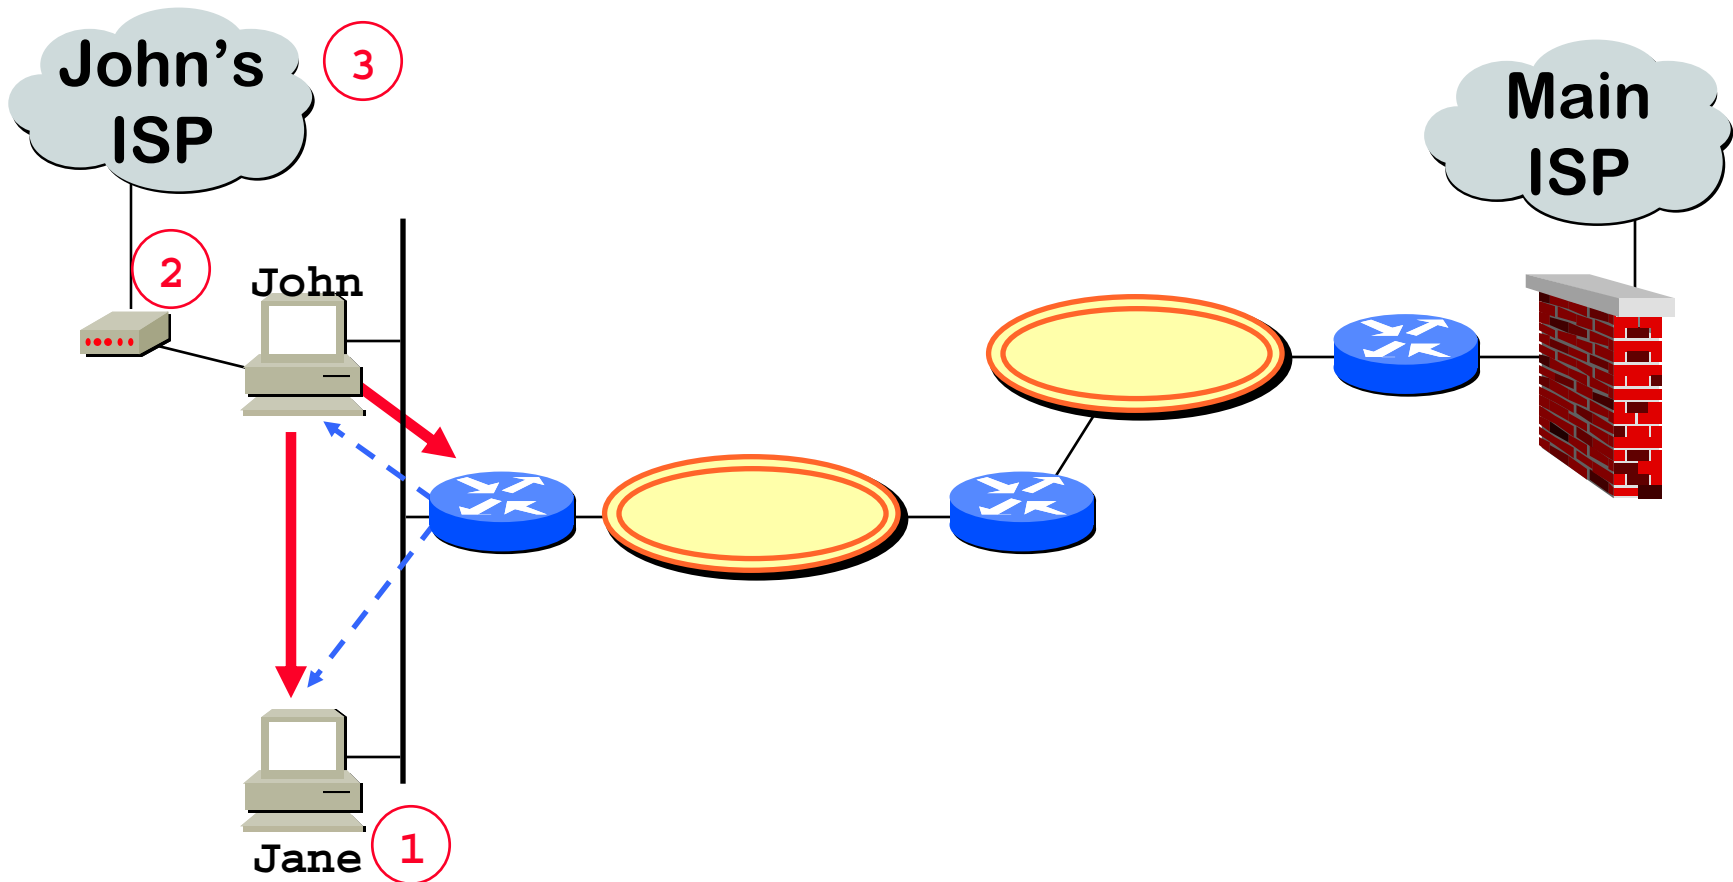
- **Customer Advertising to You with Interior Routing Protocol**

- **Generally not a good idea**
- **If you must, OSPF, EIGRP, IS-IS probably safer**
  - **RIP has too many leak risks**

- End User RIP  
Intended Usage




# • End User RIP Problems

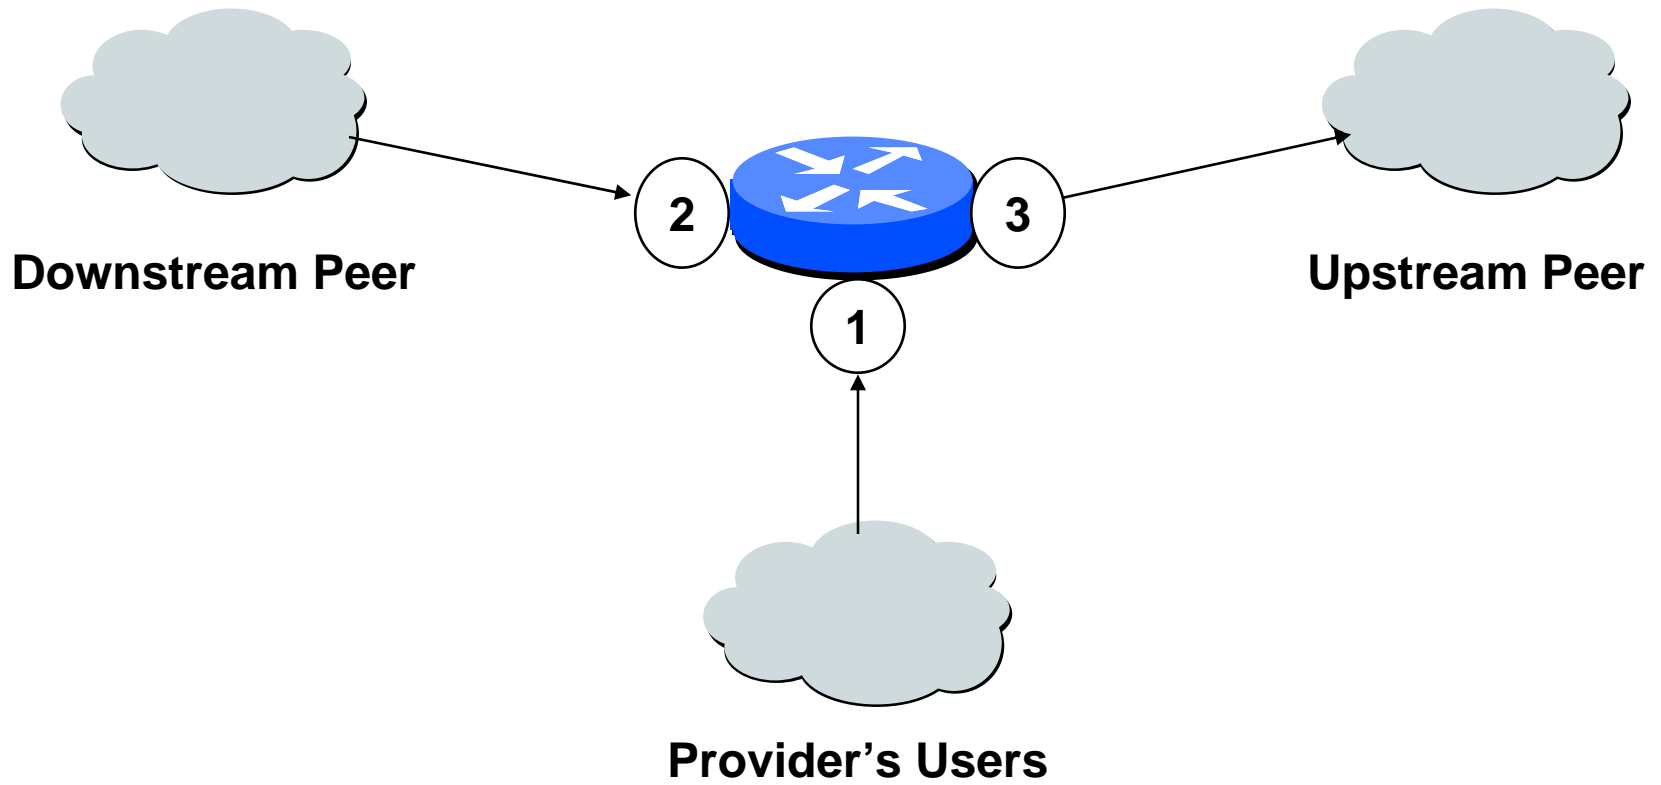


# • End User RIP

## Enterprise and ISP Protective Measures



# • Provider Filtering\



- **Enterprises Protect Themselves**

- **Sources to block**

- Internal registered space
- RFC1918
- Known rogues

- **Destinations to block**

- RFC 1918
- Firewalled internal space
- Unauthorized DMZ destinations

## Externals & Aggregation 2

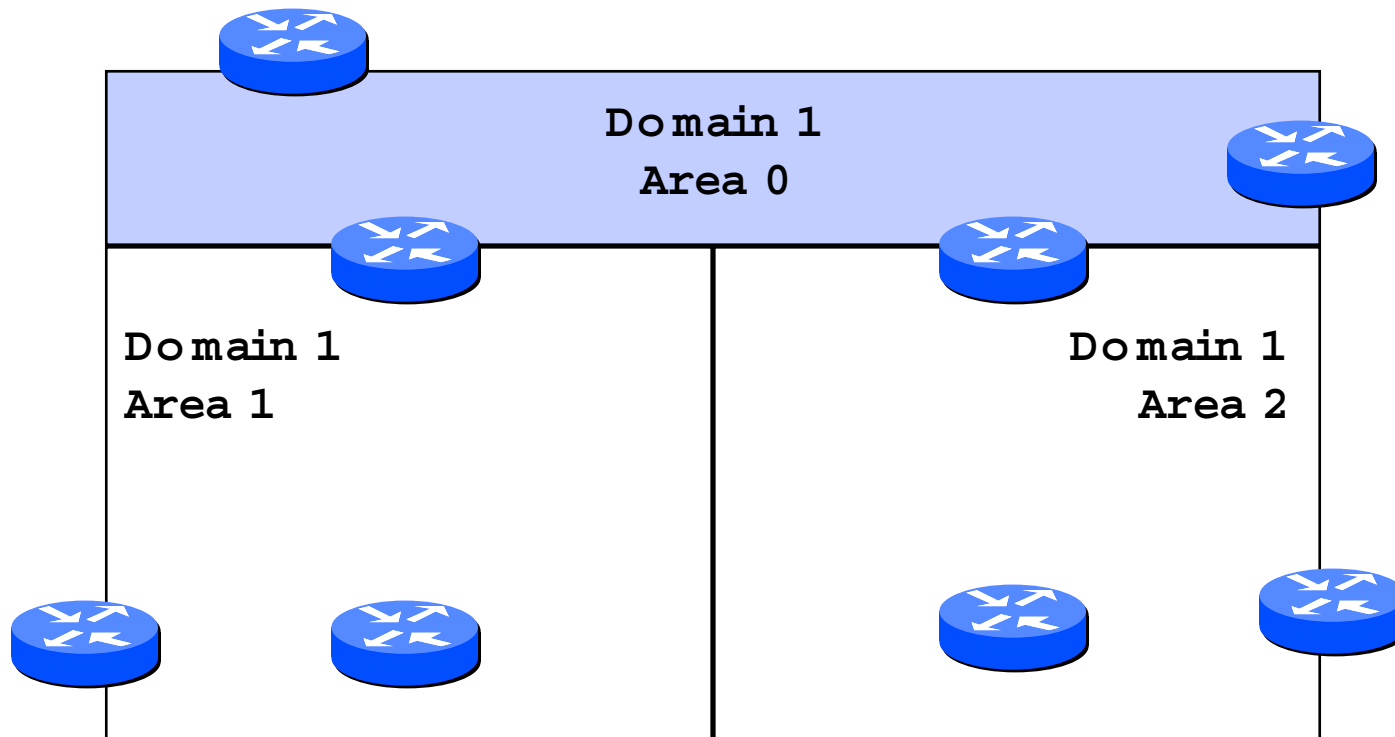
- **In an enterprise**
  - **Limit externals from subordinate domains (e.g., RIP)**
    - **Flood only in area 0 and in area with ASBR**
  - **Allow externals from Internet, peer domains to go outside Area 0**
    - **Only when there will be significant path differences**
    - **Do things with defaults where possible**

## Type 1 and Type 2 externals

- **Type 2**
  - Default type for routes distributed into OSPF
  - Exit based on external cost only
- **Type 1**
  - Needs to be set explicitly
  - Selects exit based on internal + external costs



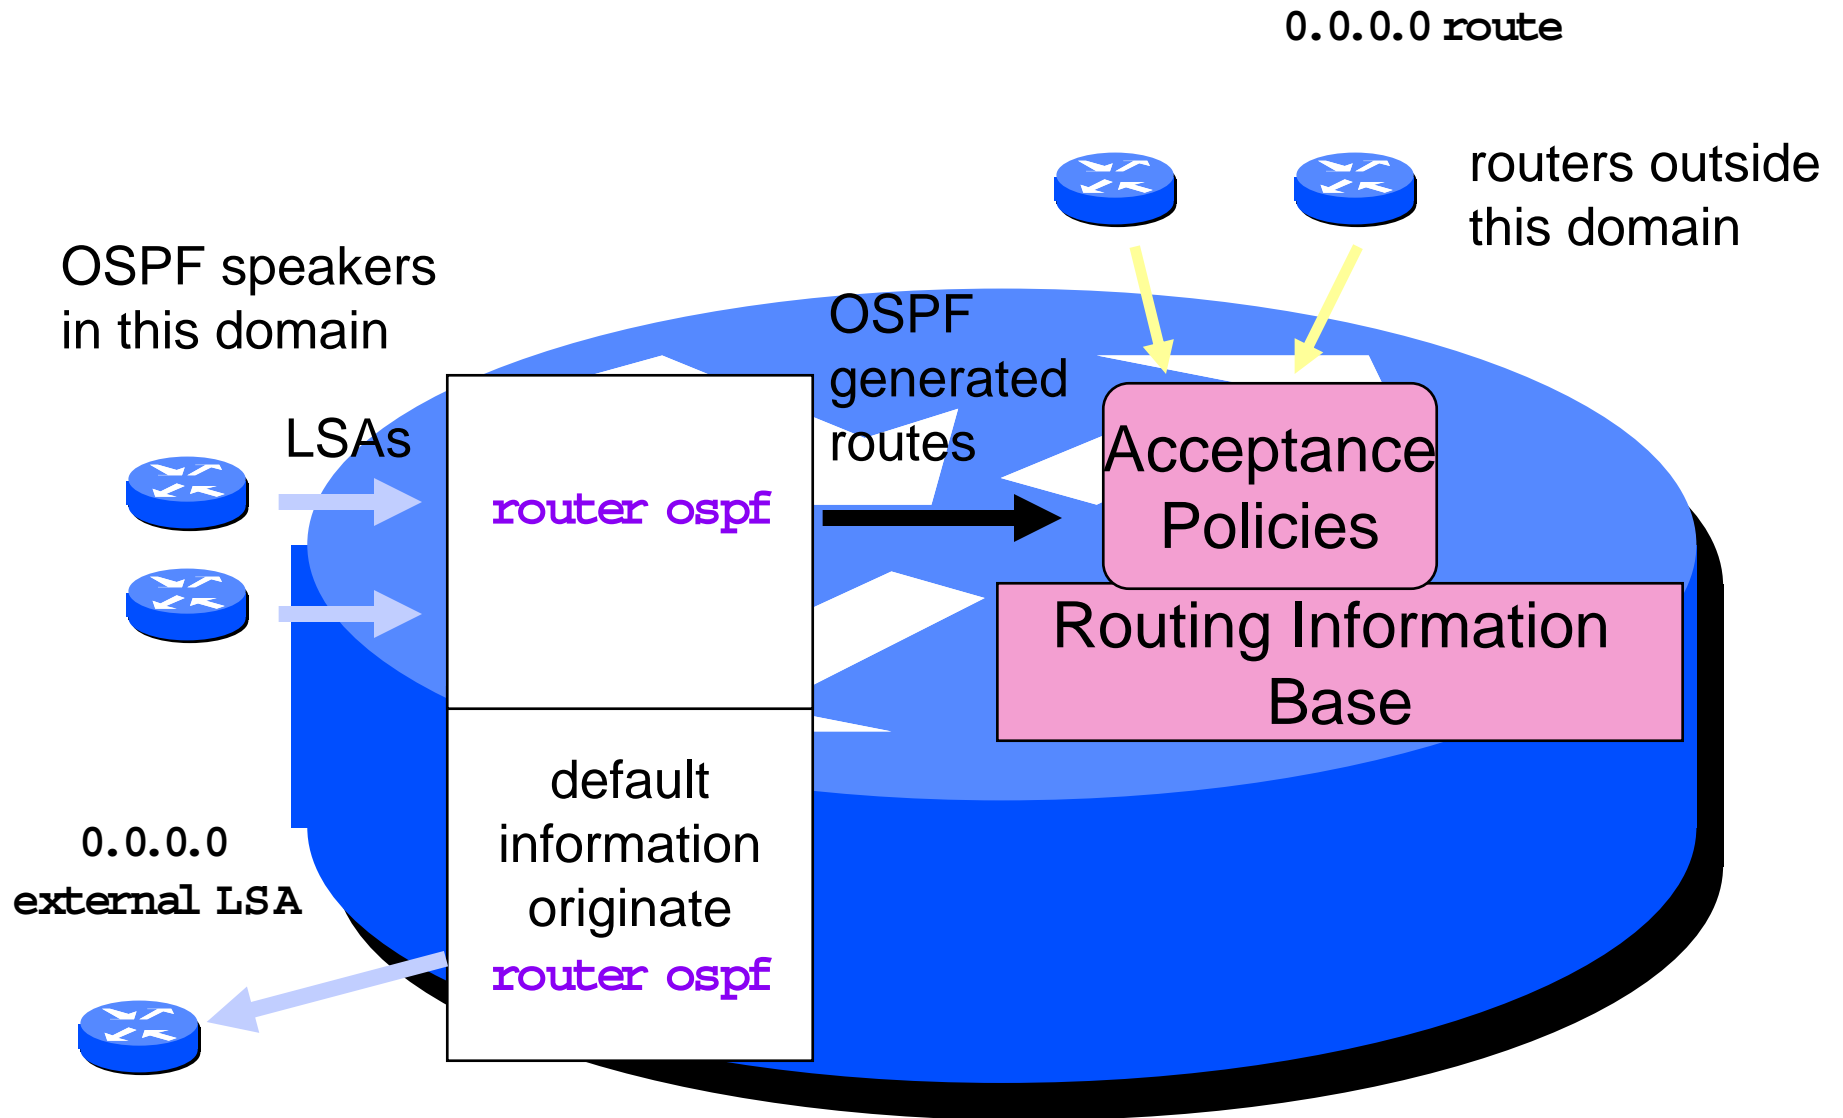
# ASBR Placement




Contrary to common opinion...you can have ASBR outside area 0.

Nonzero area with ASBR can't be stubby or totally stubby

# ASBR Default Information



# OSPF and Default Routes




OSPF Speakers

Edge routers  
(note dual-homed default)

**Stubbiness: A Means of Controlling**  
**Externals**

## Closest Exit Routing



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

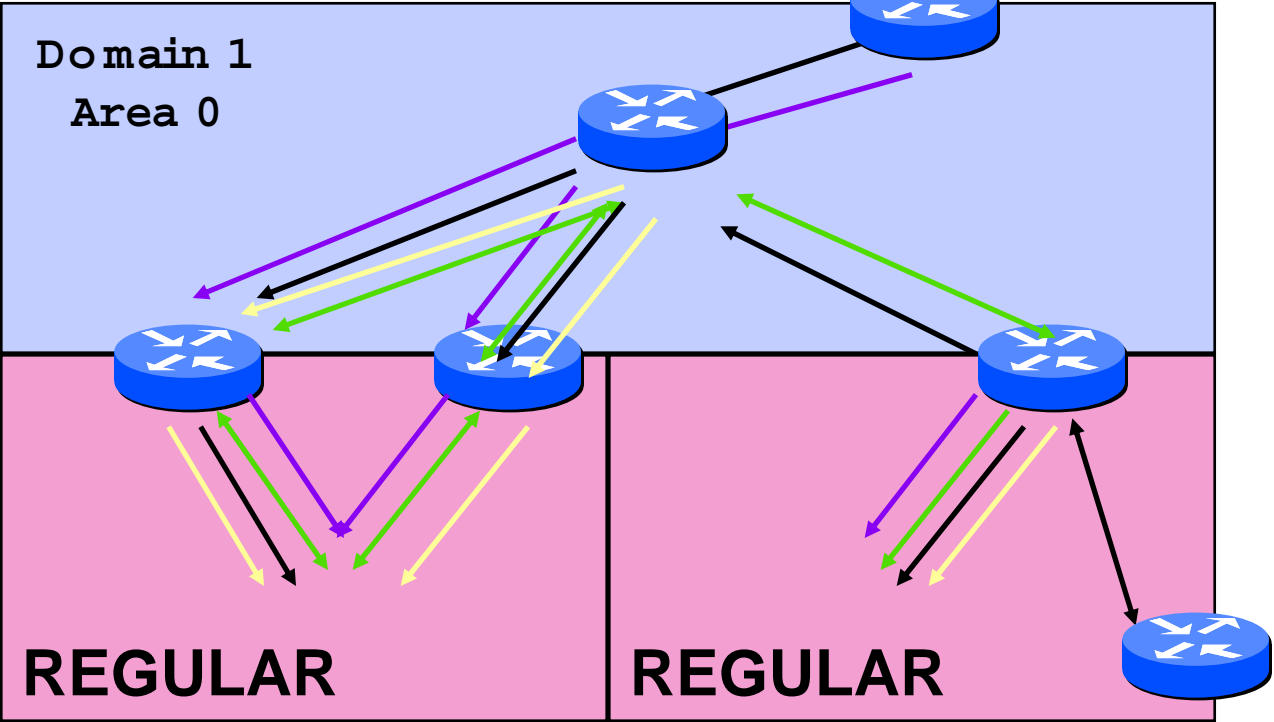
## Asymmetrical Routing

- **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**

# Regular Area

Specific internal,  
Summary external,  
Default

Specific external



## The idea of stubbiness

- **Reduce overhead in nonzero areas by reducing the number of externals**
- **Stubbiness helps only if there are significant numbers of externals**
  - **But remember any redistributed route is external**
  - **Inside an OSPF domain, stubbiness does not help**
- **External aggregation can help stubbiness**
- **Inter-area summarization complements stubbiness & aggregation for internal routes**

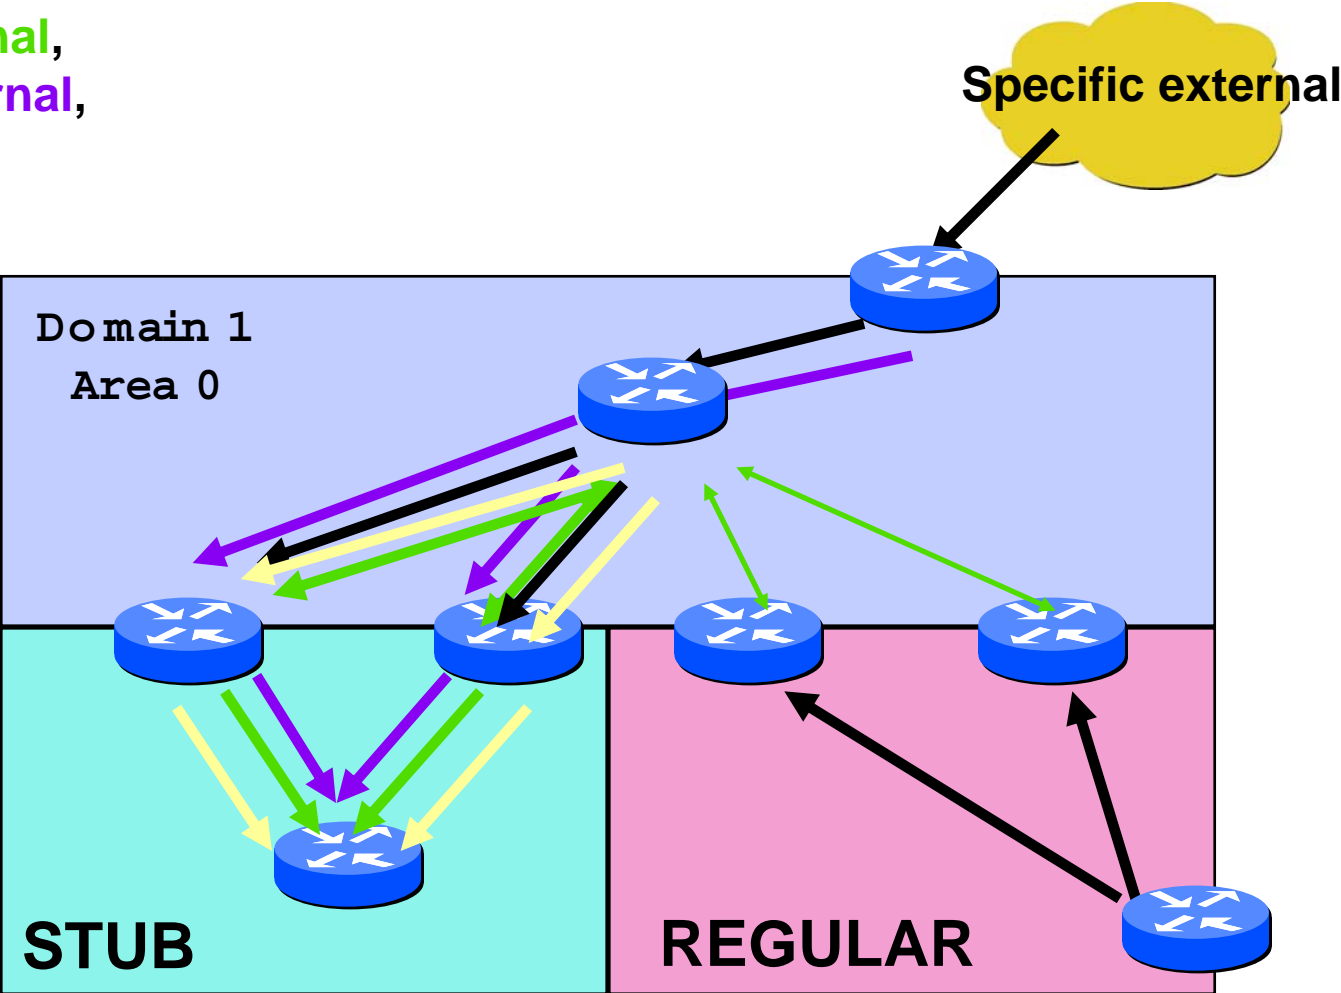


## Restrictions with Stub Areas

- **Stubby**
  - Cannot have ASBR outside Area 0
  - Potential problem during migration
  - Potential problem for multiple Internet attachment points
- **Totally stubby**
  - Single ABR; single point of failure
  - Cisco proprietary
  - Maximum traffic reduction
- **Not So Stubby**
  - Can only know externals injected into it by local ASBR

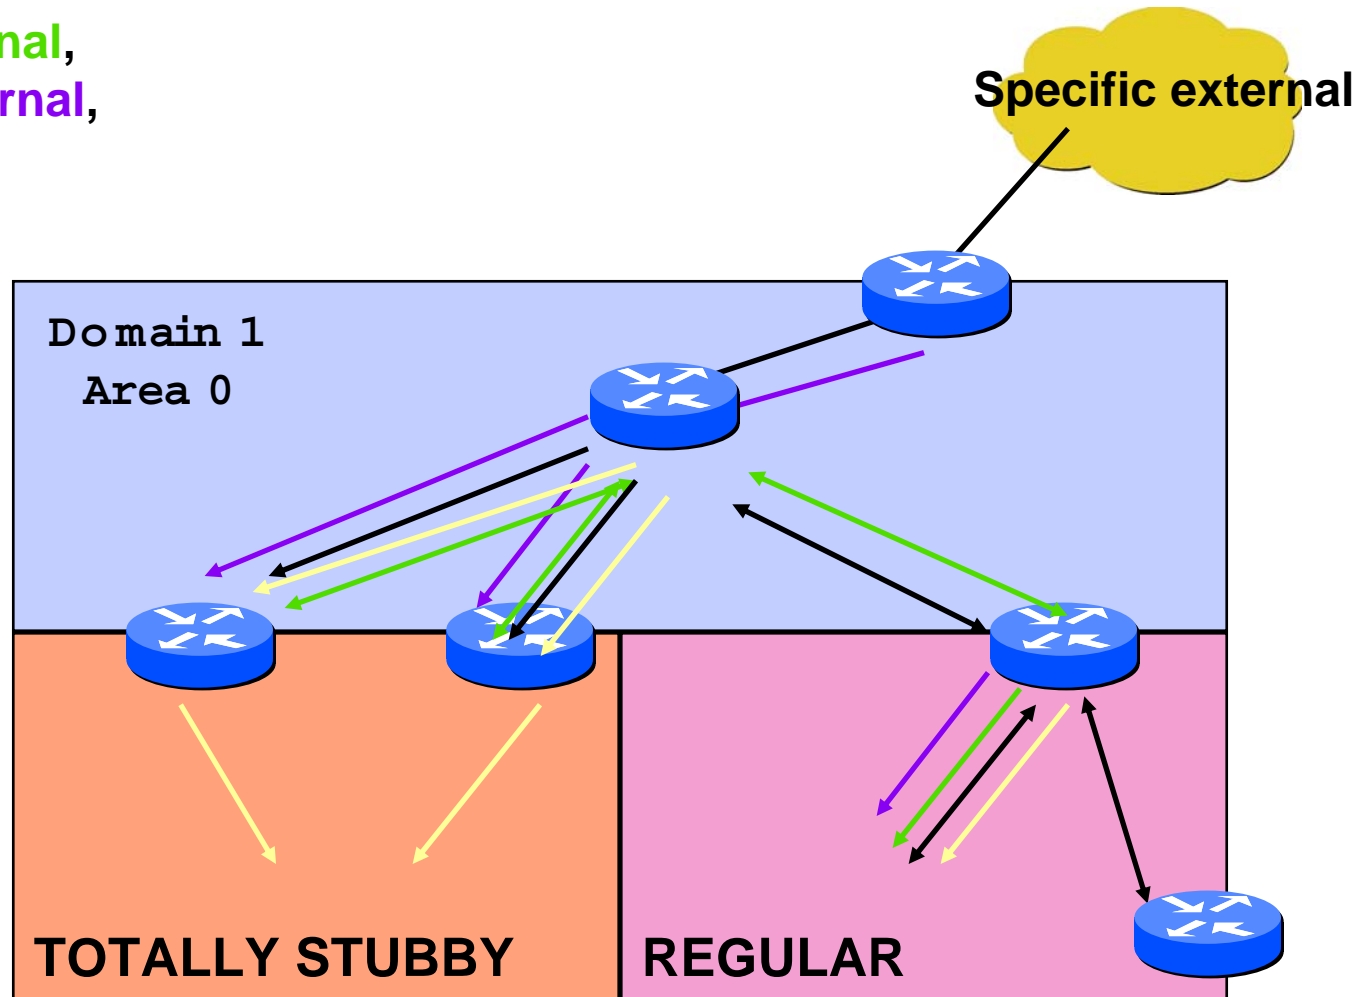
# Stubby Area

Specific internal,  
Summary external,  
Default



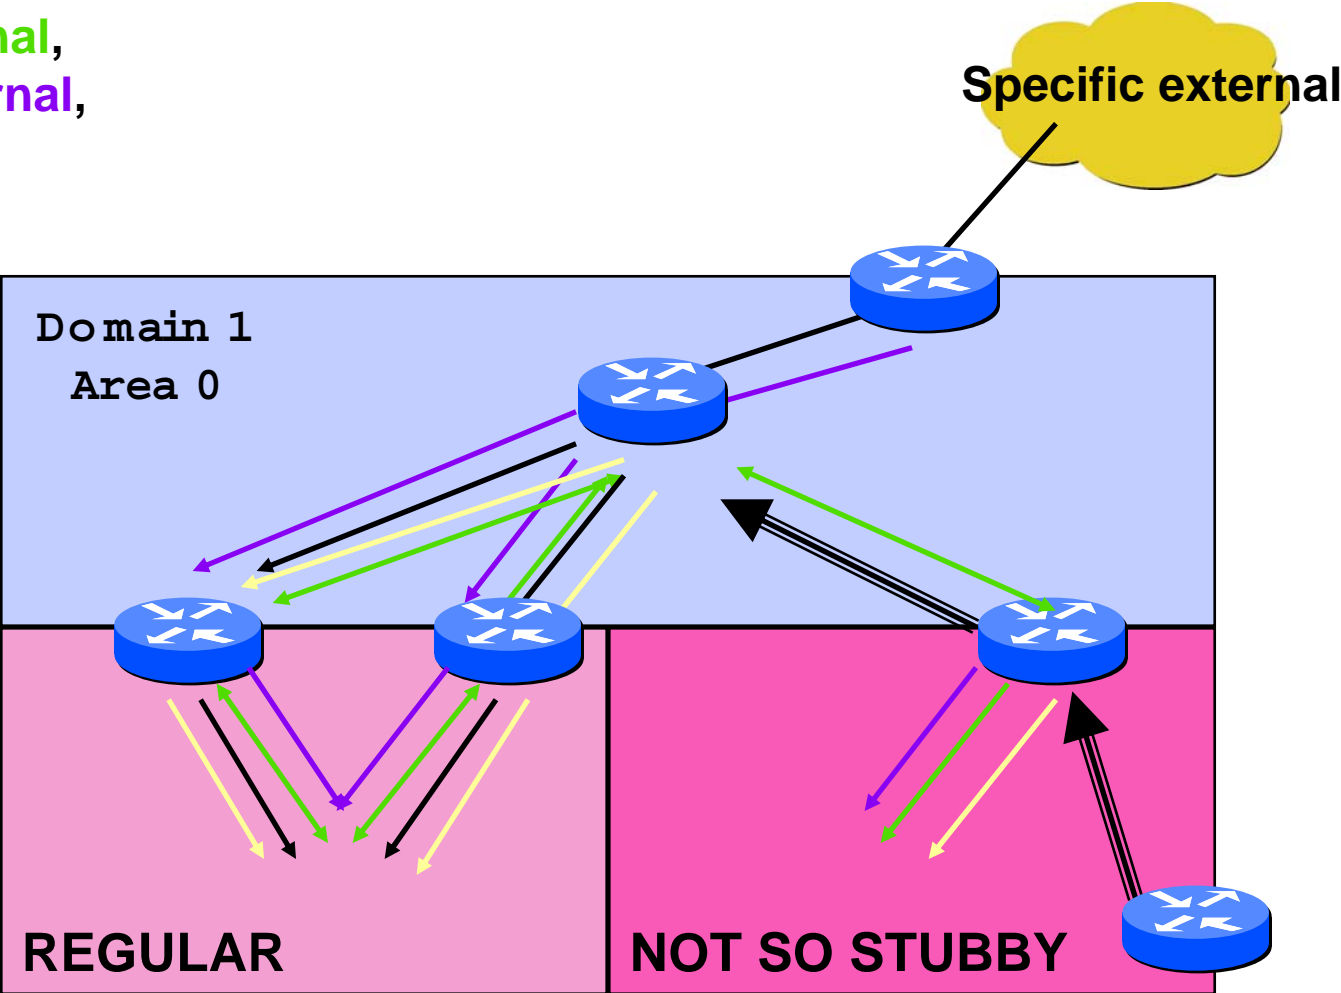
# Totally Stubby Area

Specific internal,  
Summary external,  
Default



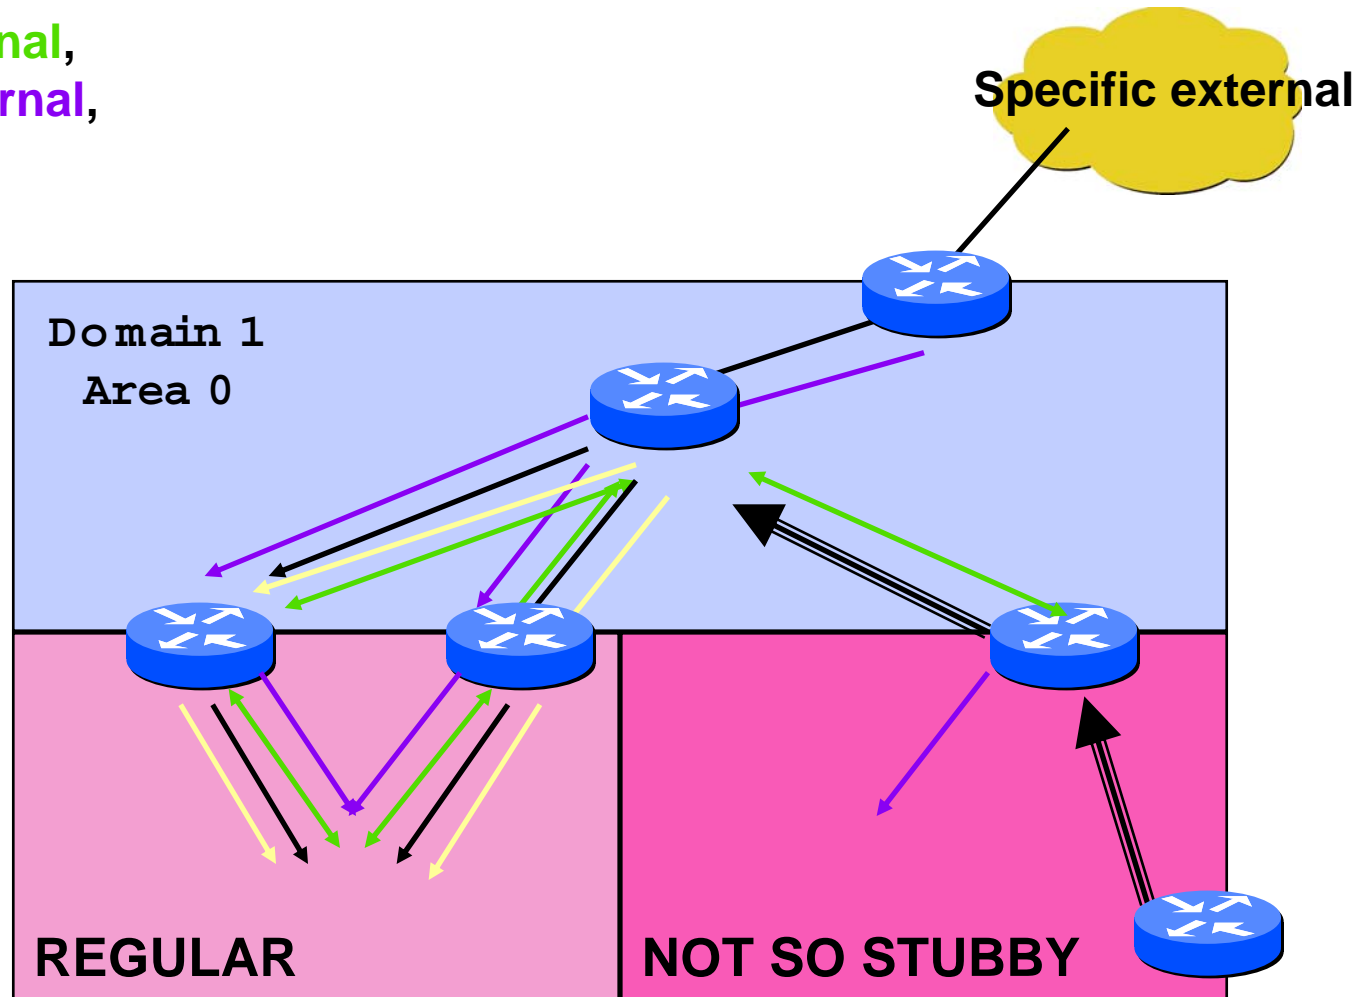
# Not So Stubby Area

Specific internal,  
Summary external,  
Default



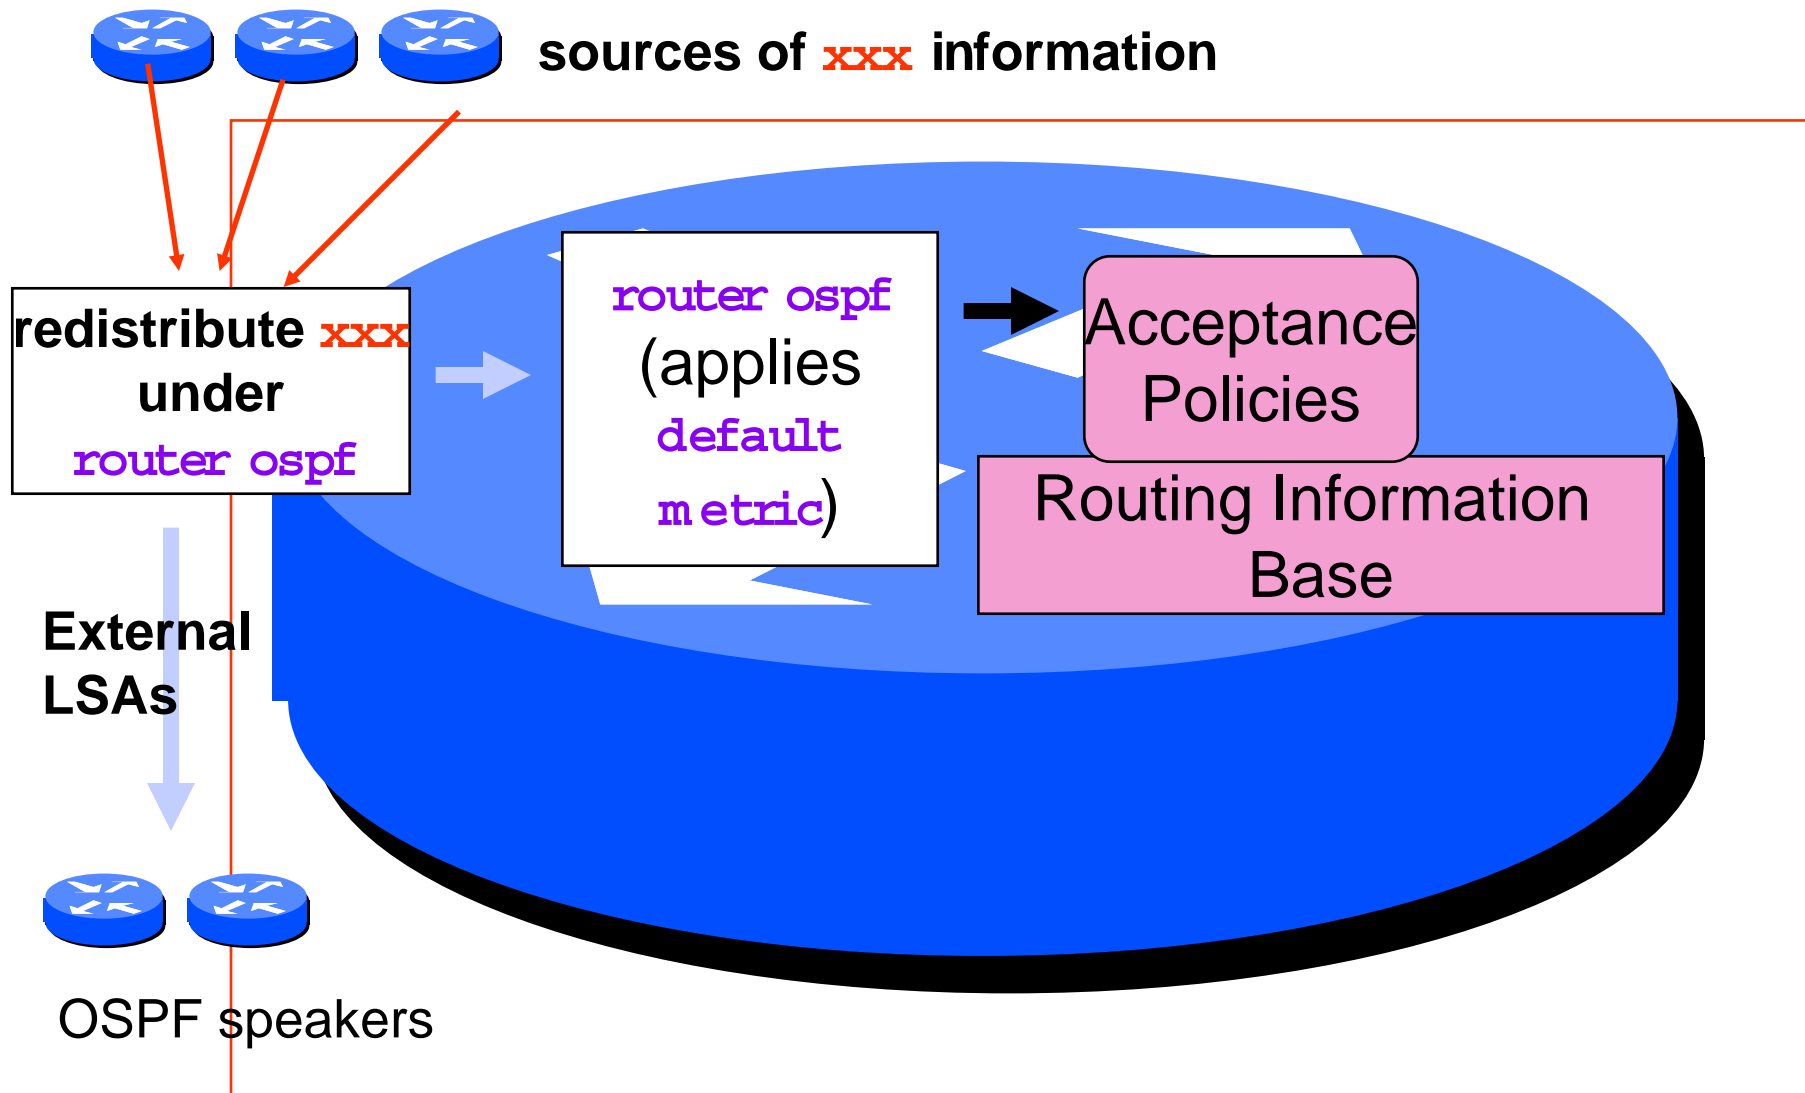
# Totally Stubby and also Not So Stubby

Specific internal,  
Summary external,  
Default

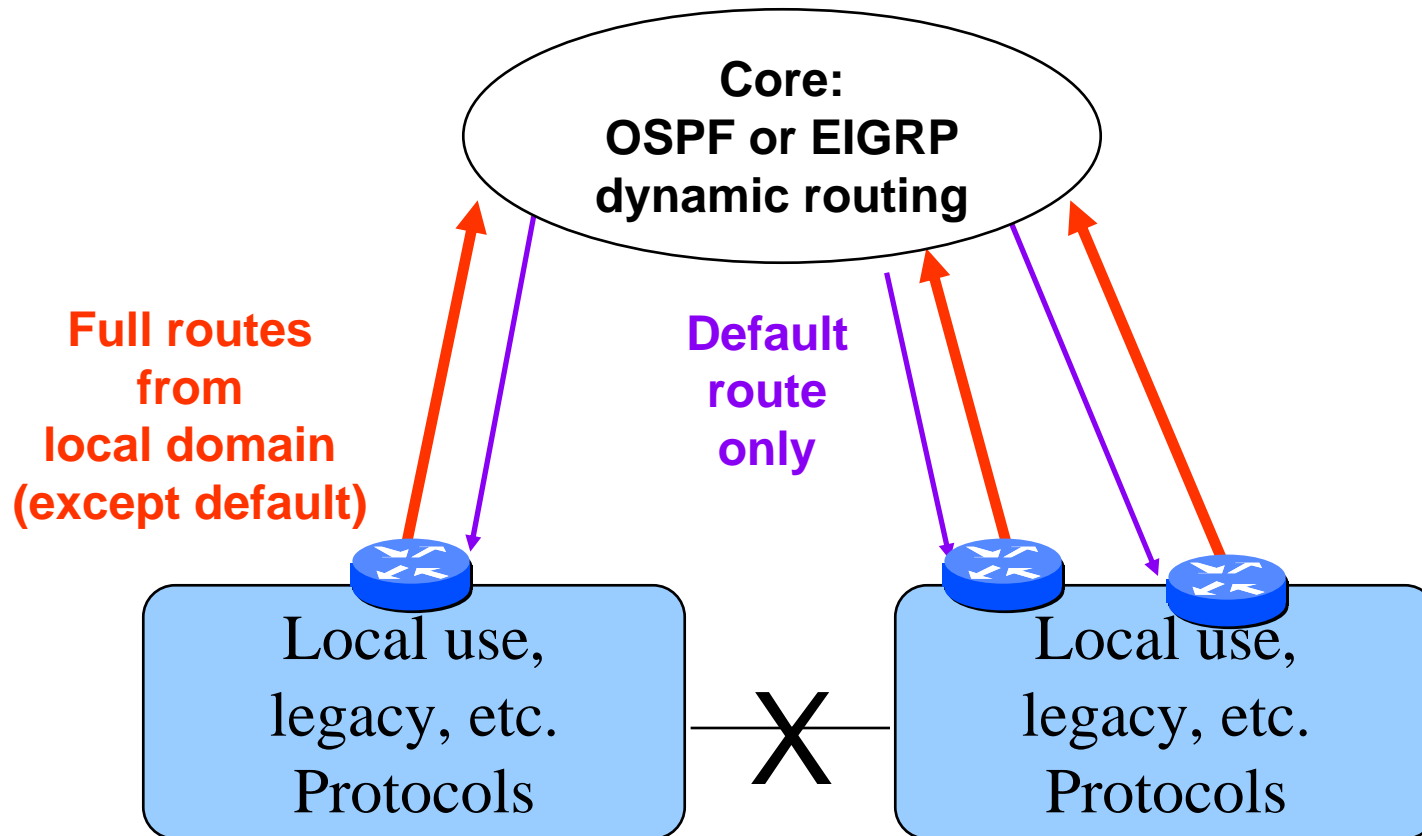


# **Redistribution**

# Redistribution = Acceptance




# Hierarchical Redistribution



**Loop-free and easy to configure**  
**Can use multiple routers**



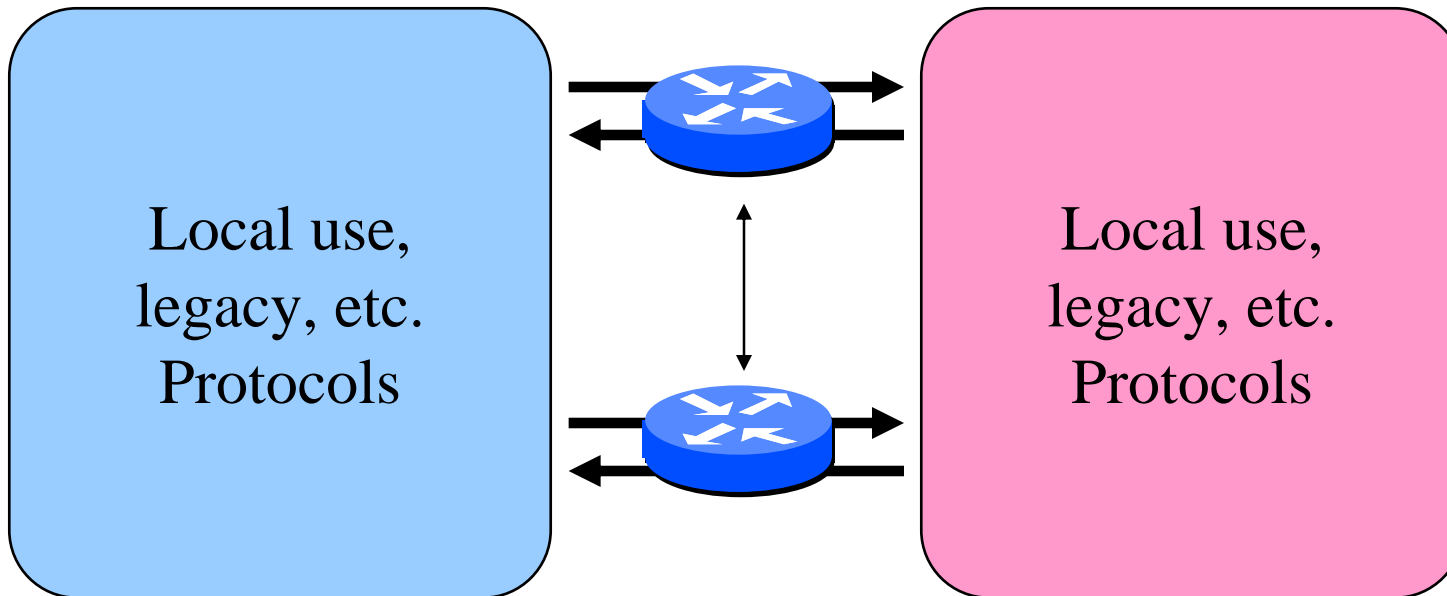
## Basic Mutual Redistribution



**Manual configuration usually needed to prevent loops**  
**Single redistributing router makes it much simpler**  
**But is a single point of failure**

# Mutual Redistribution

## Multiple Redistribution Points




**Extensive configuration usually needed to prevent loops**  
**No single point of failure**  
**Very hard to troubleshoot**

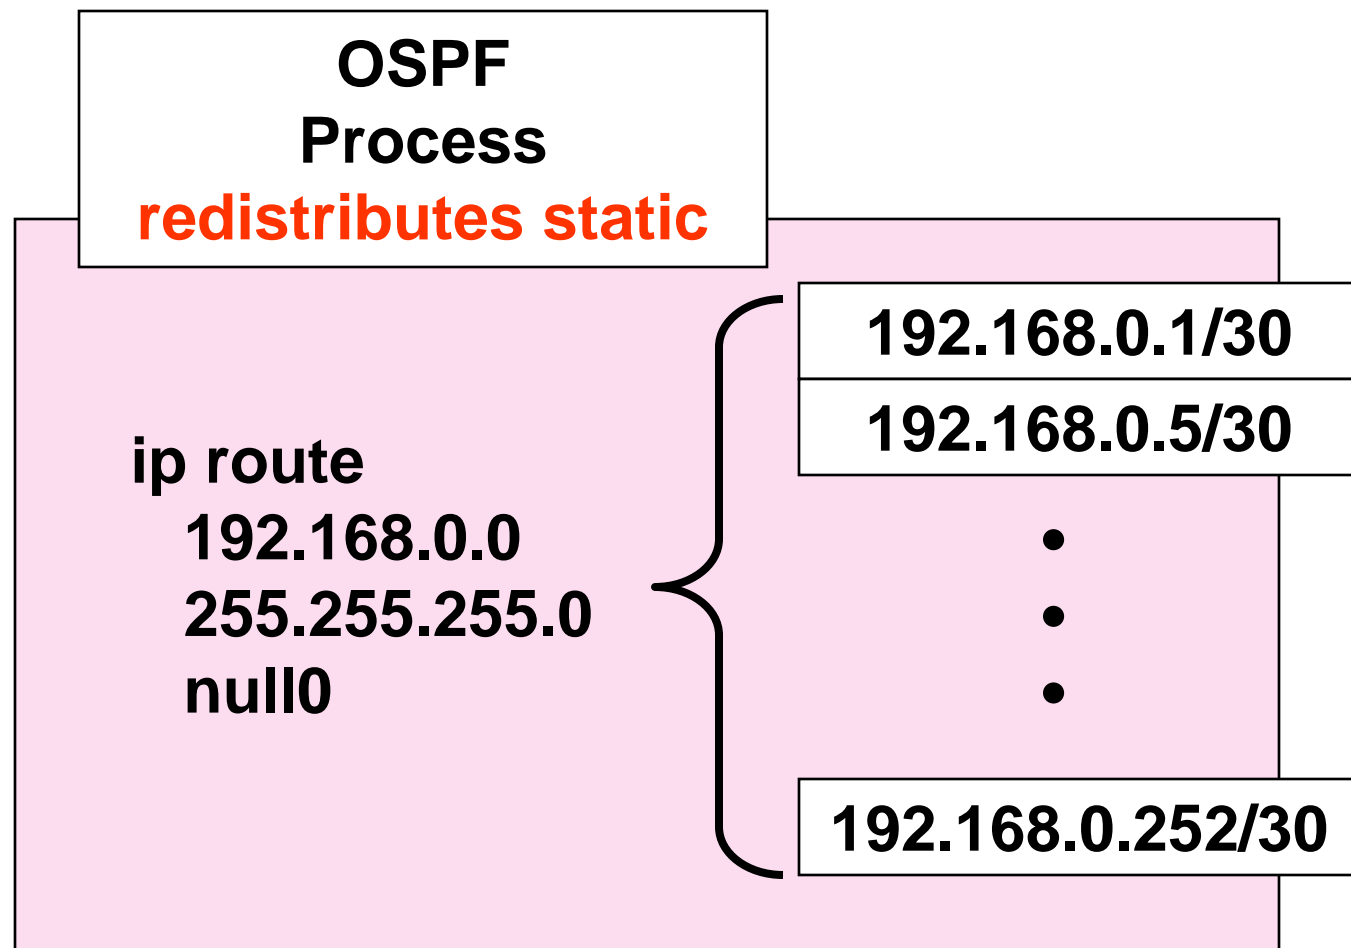
## OSPF and Default Routes

- **OSPF convention**
  - **0.0.0.0/0 (subnet mask 0.0.0.0)**
  - **Treated as external**
  - **May have different defaults in different areas**
- **Blackhole routes give pseudo-default**

# An Application of Blackhole Routes



# Inside the Dial Server



## **Complex and Useful Domain Relationships**

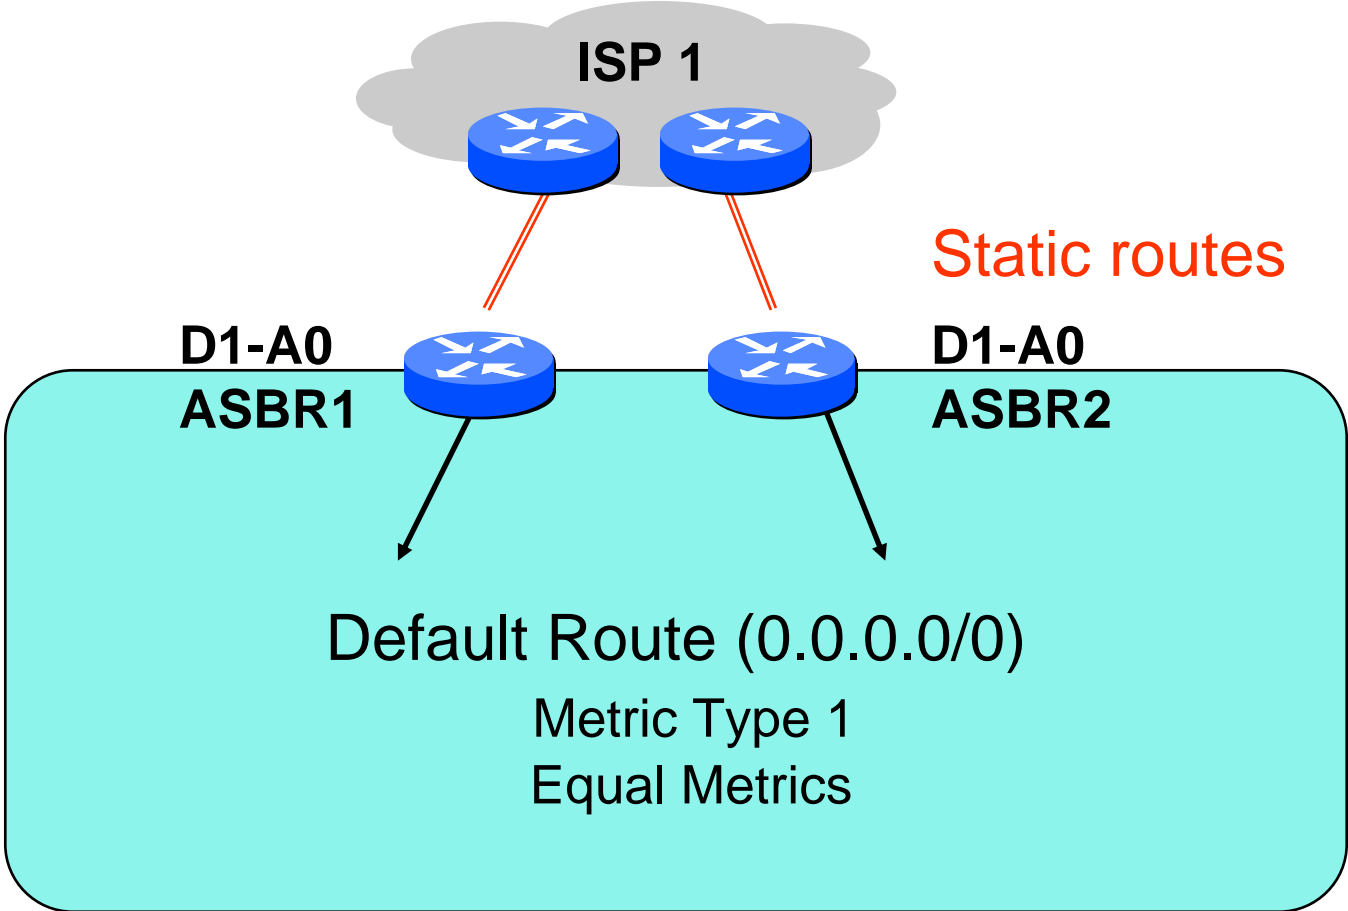
## A Trap to be Avoided:

### Overemphasizing Area 0

- **A given domain has only one Area 0**
- **But if there are scaling problems in a domain...**
- **Use multiple domains**
  - Each with their own Area 0
  - Connecting to other OSPF and non-OSPF domains via ASBRs
- **Use appropriate extensions**
  - Virtual links
  - ASBRs outside Area 0
    - NSSAs a powerful new extension

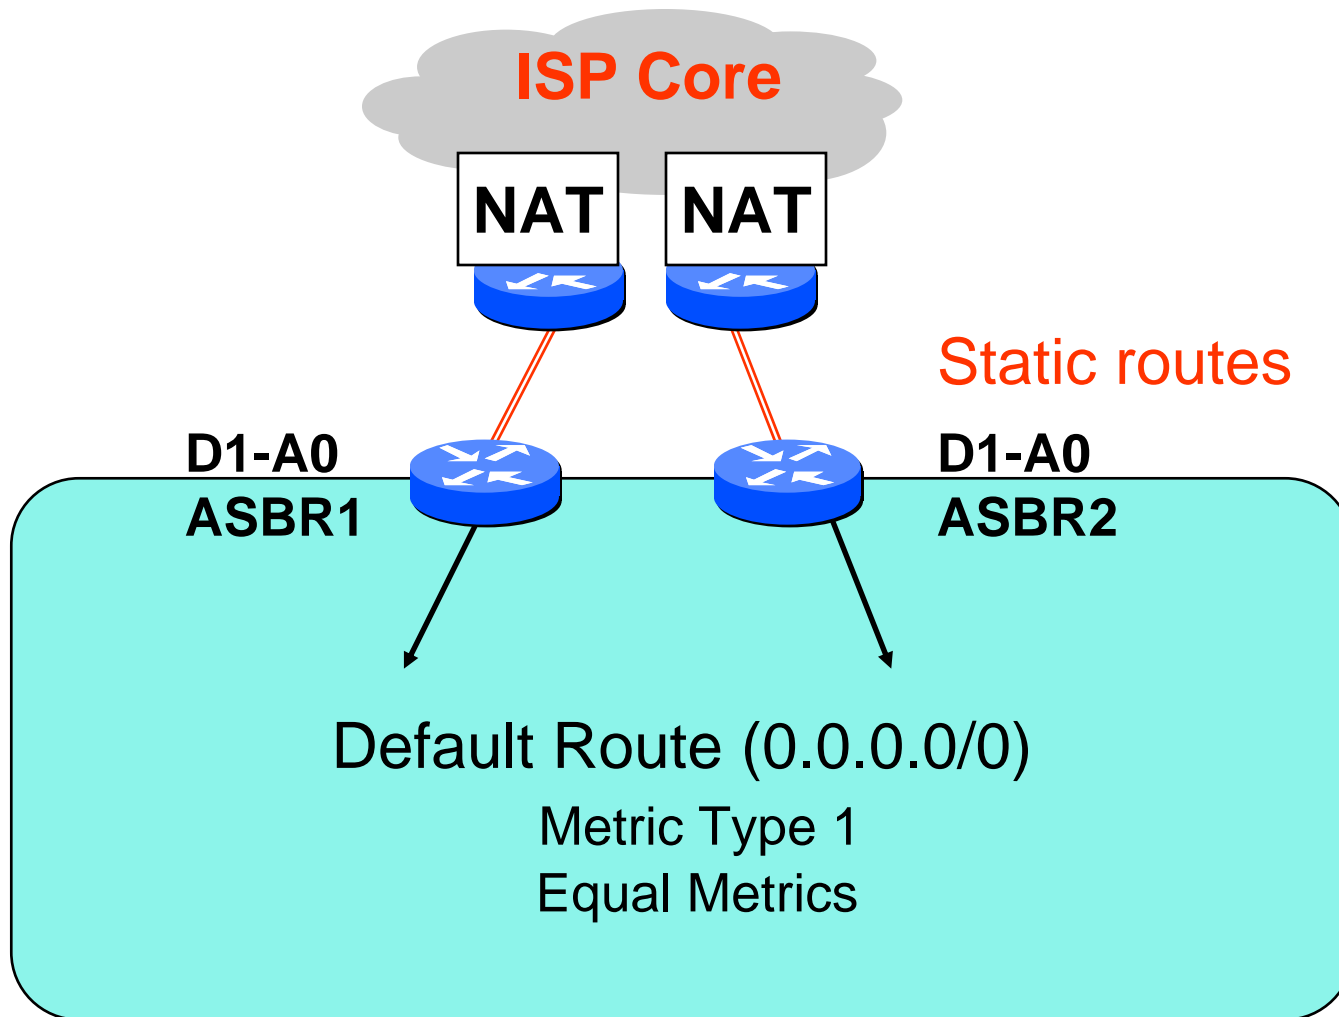
# Access to Single ISP

Closest-Exit Load Sharing, No BGP Needed



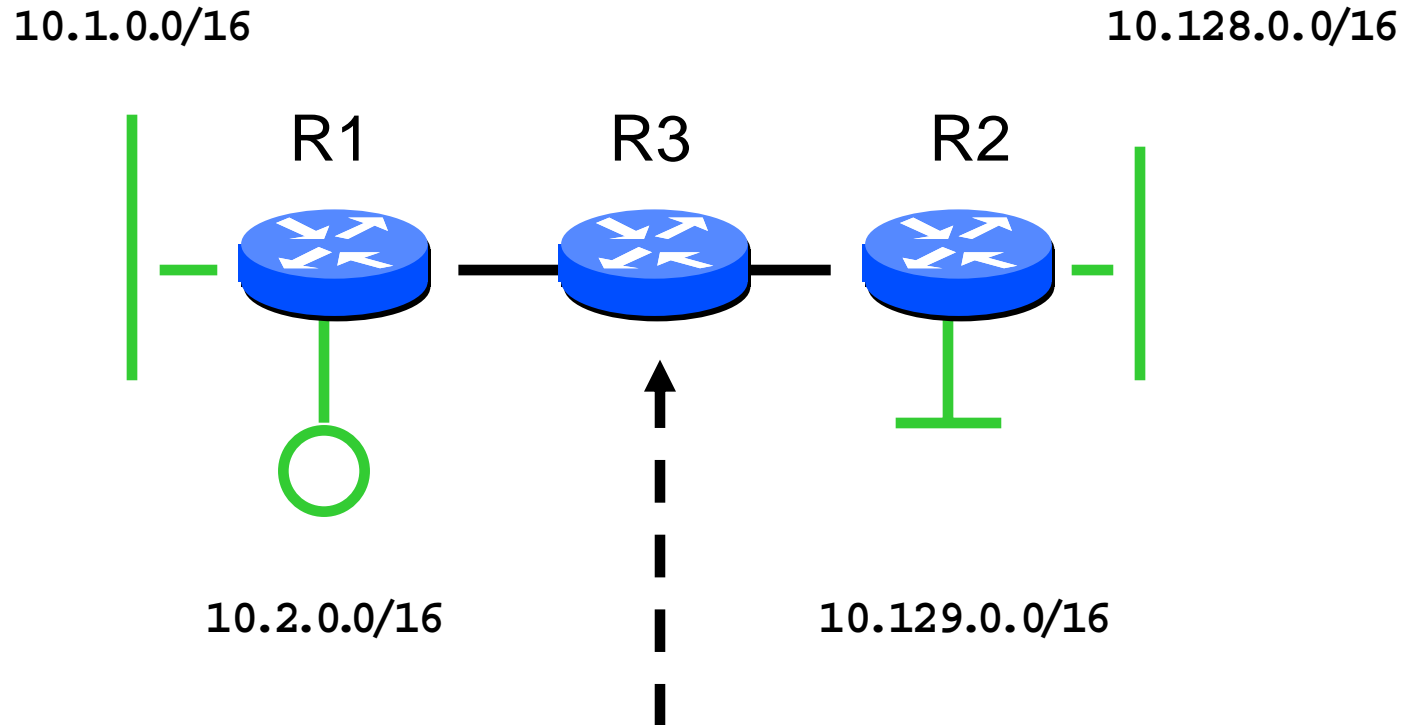


# Multihoming with Customer RFC1918 Space



# Discontiguous Subnets


## The Usual Enterprise-Oriented Picture



If R3 only knows about "network 10,"  
where does it send a packet for 10.64.0.0?

# Discontiguous Subnets

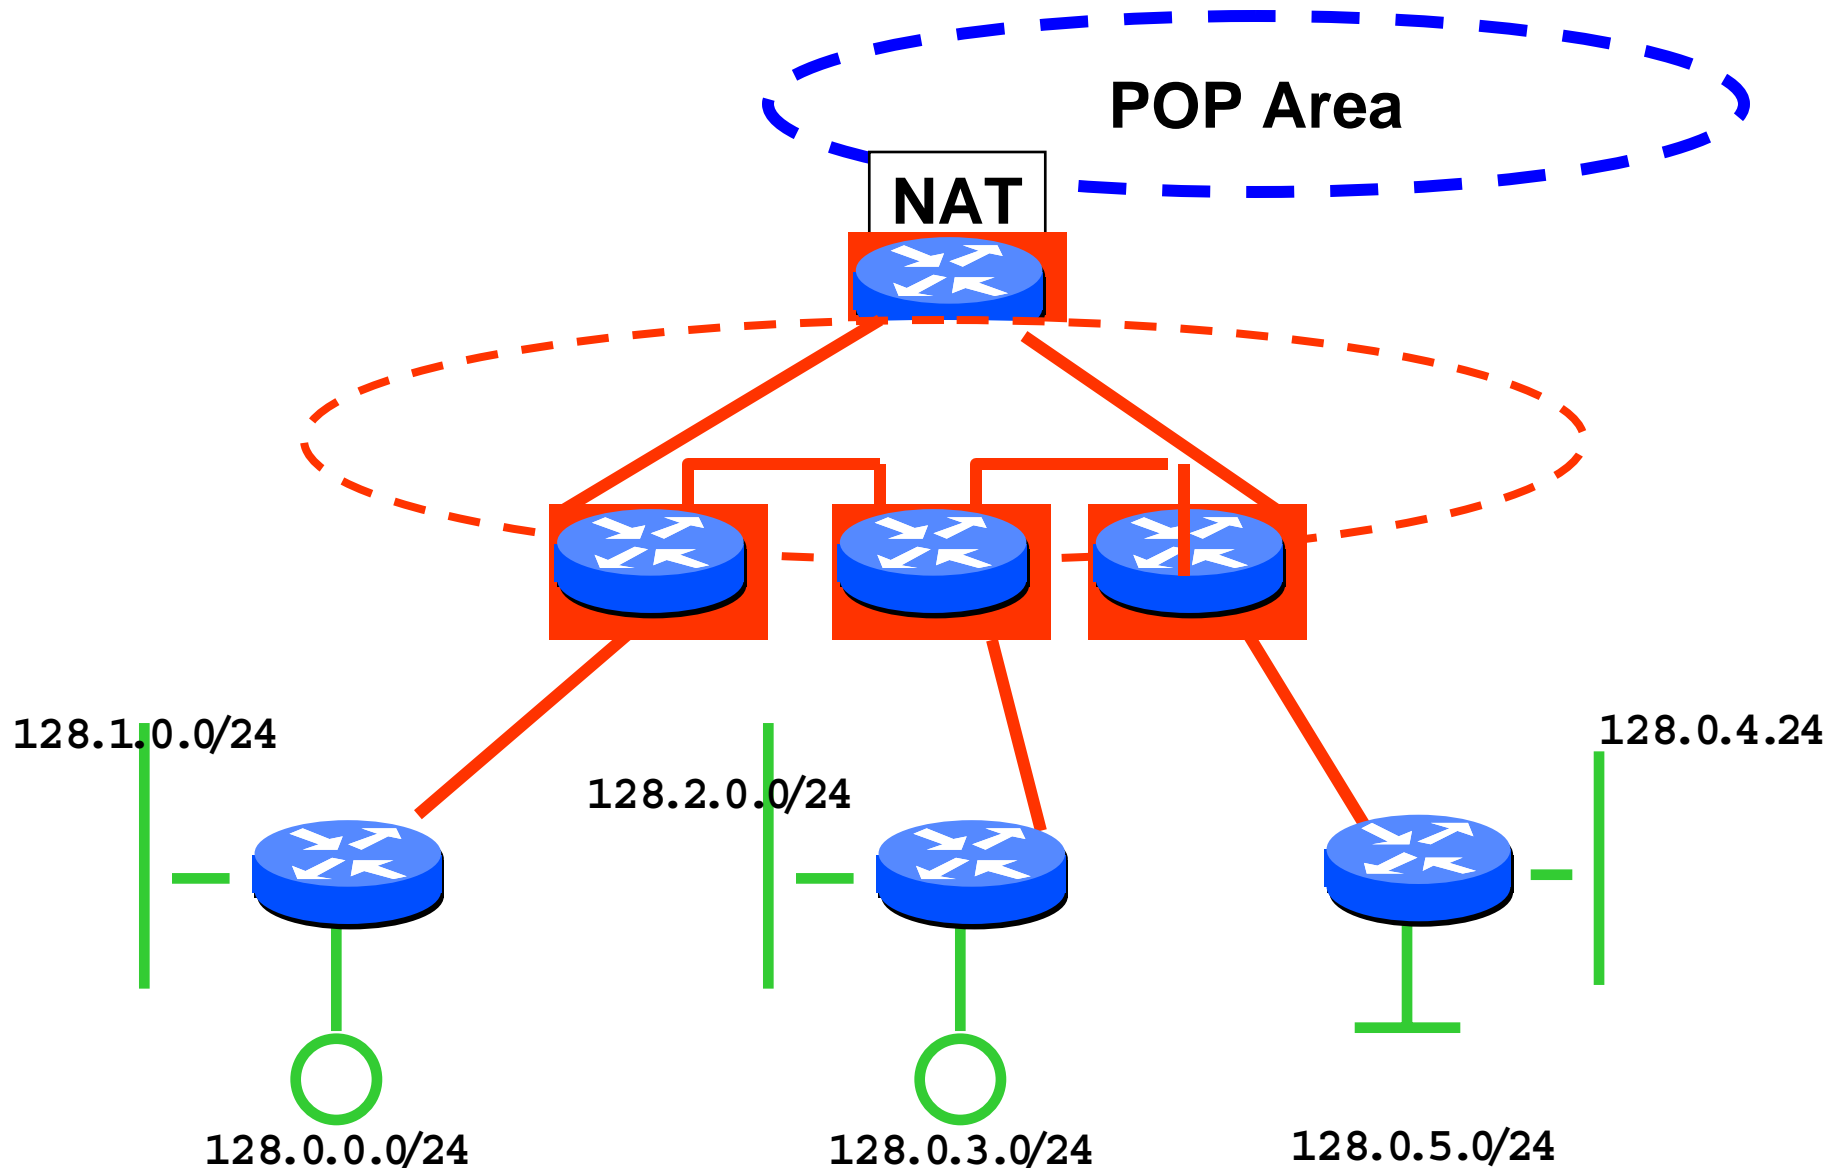
Now it's the ISP's Problem



# Healing Discontiguous Networks Techniques

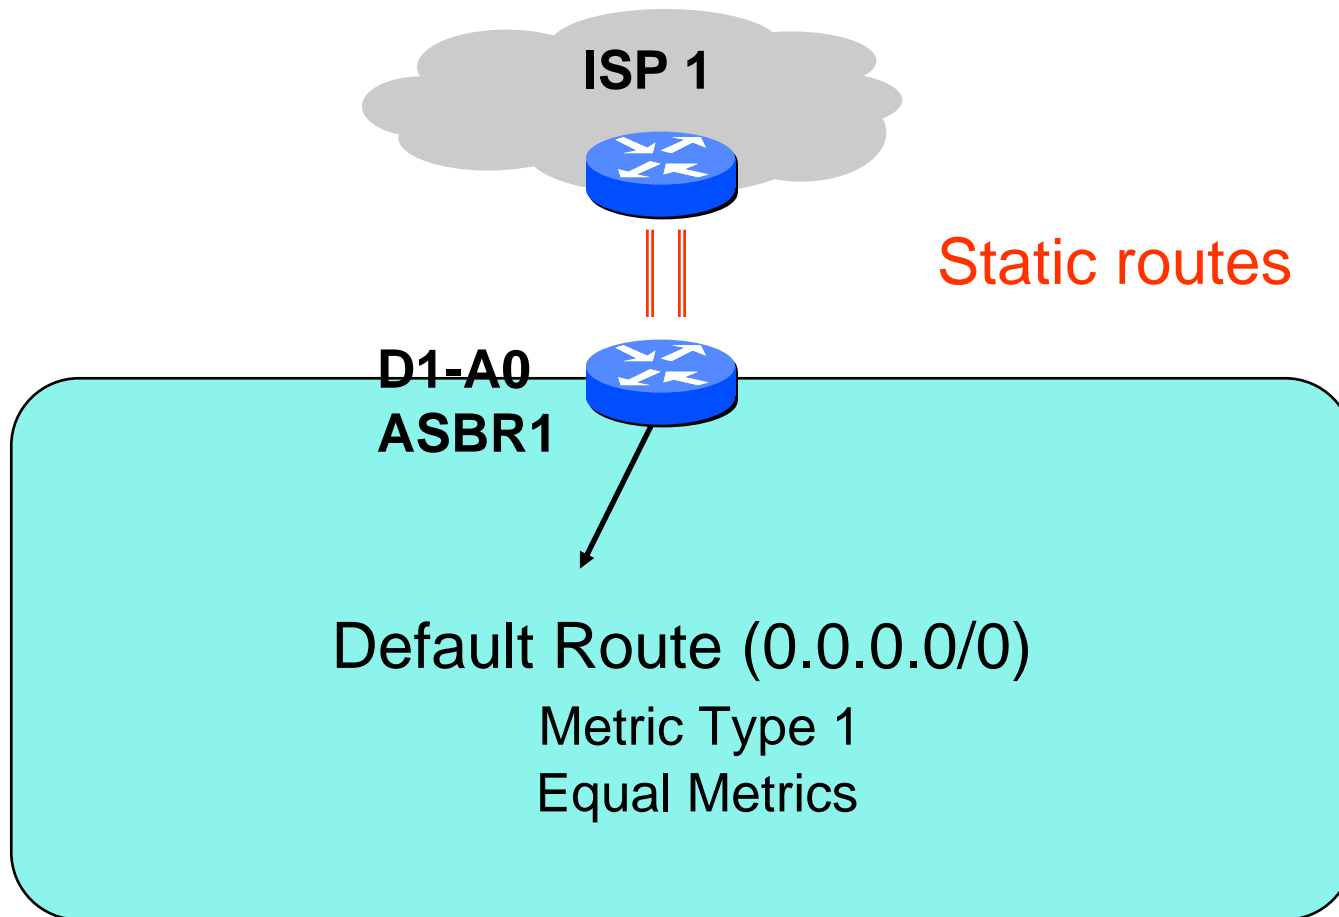
- **Secondary addressing**
- **IP over IP tunneling**
- **IP unnumbered**
- **Static routes**
- **Network address translation**
- **Best way**
  - **Move to classless environment**
  - **No more discontinuity**
- **Care needed to deal with routing table growth**
  - **Summarize where practical**
  - **Address assignment must be more careful**

# Dealing with Customers with Class

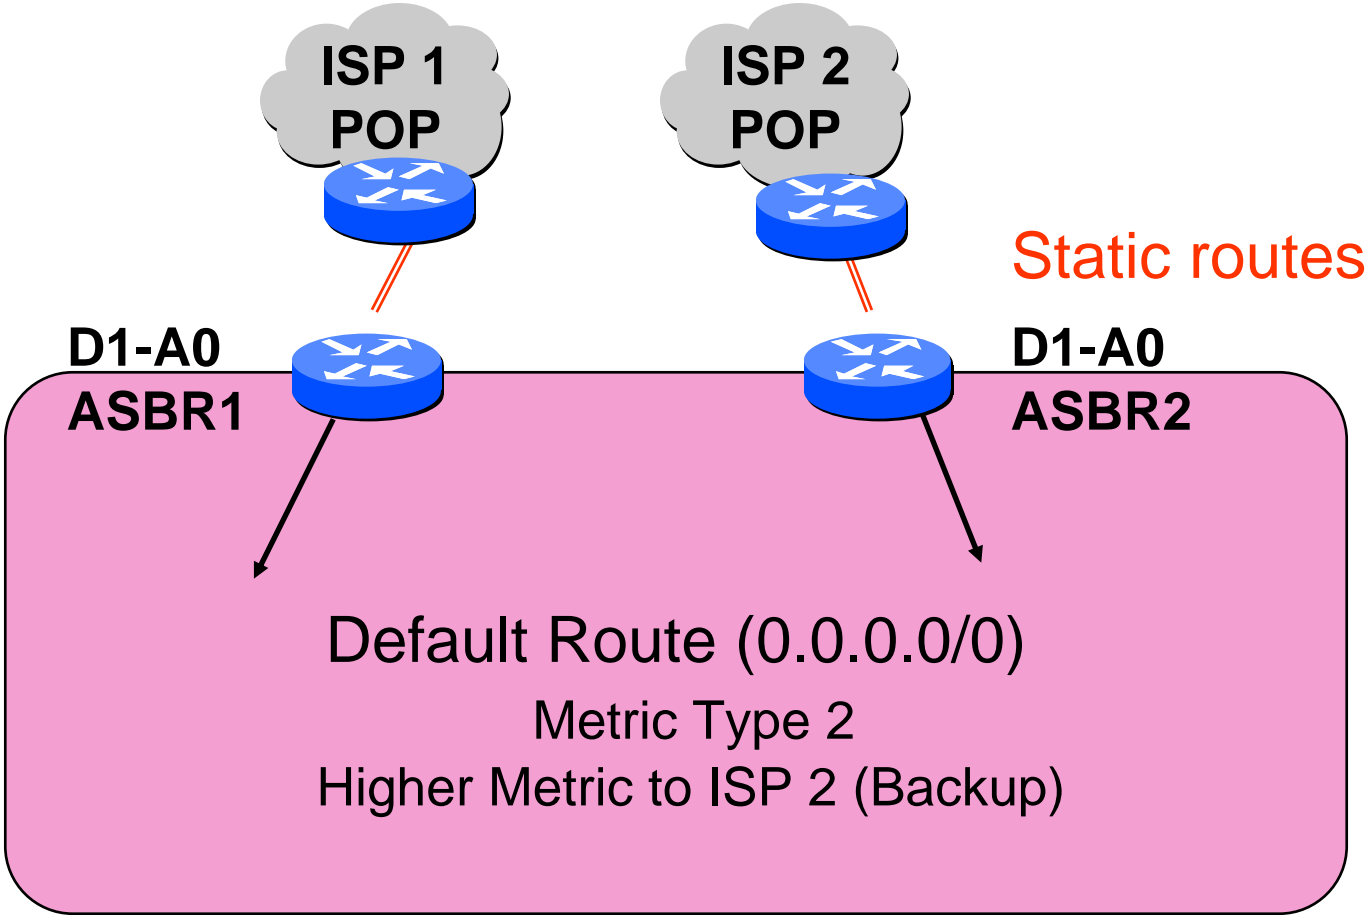


# Single ISP

## Link-Level Load Balancing



# Access to Primary & Backup ISP



# Internet Access

## with Exception Case

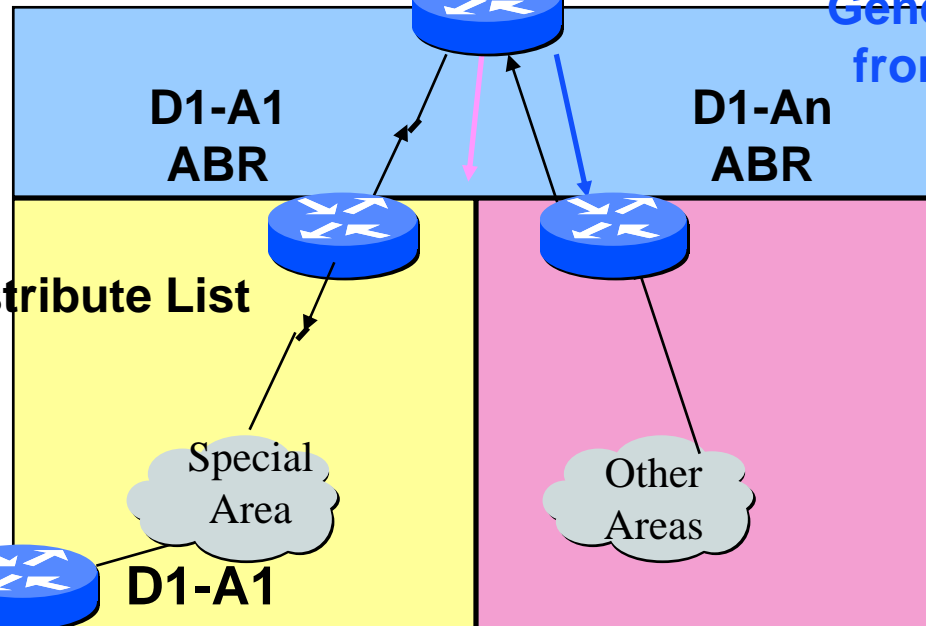
Upstream Distribute List  
blocking 0.0.0.0  
towards Area 0  
Permits all other LSAs



Static route

D1-A0  
ASBR

General default route  
from Area 0



Downstream Distribute List  
blocking 0.0.0.0  
towards Area 1

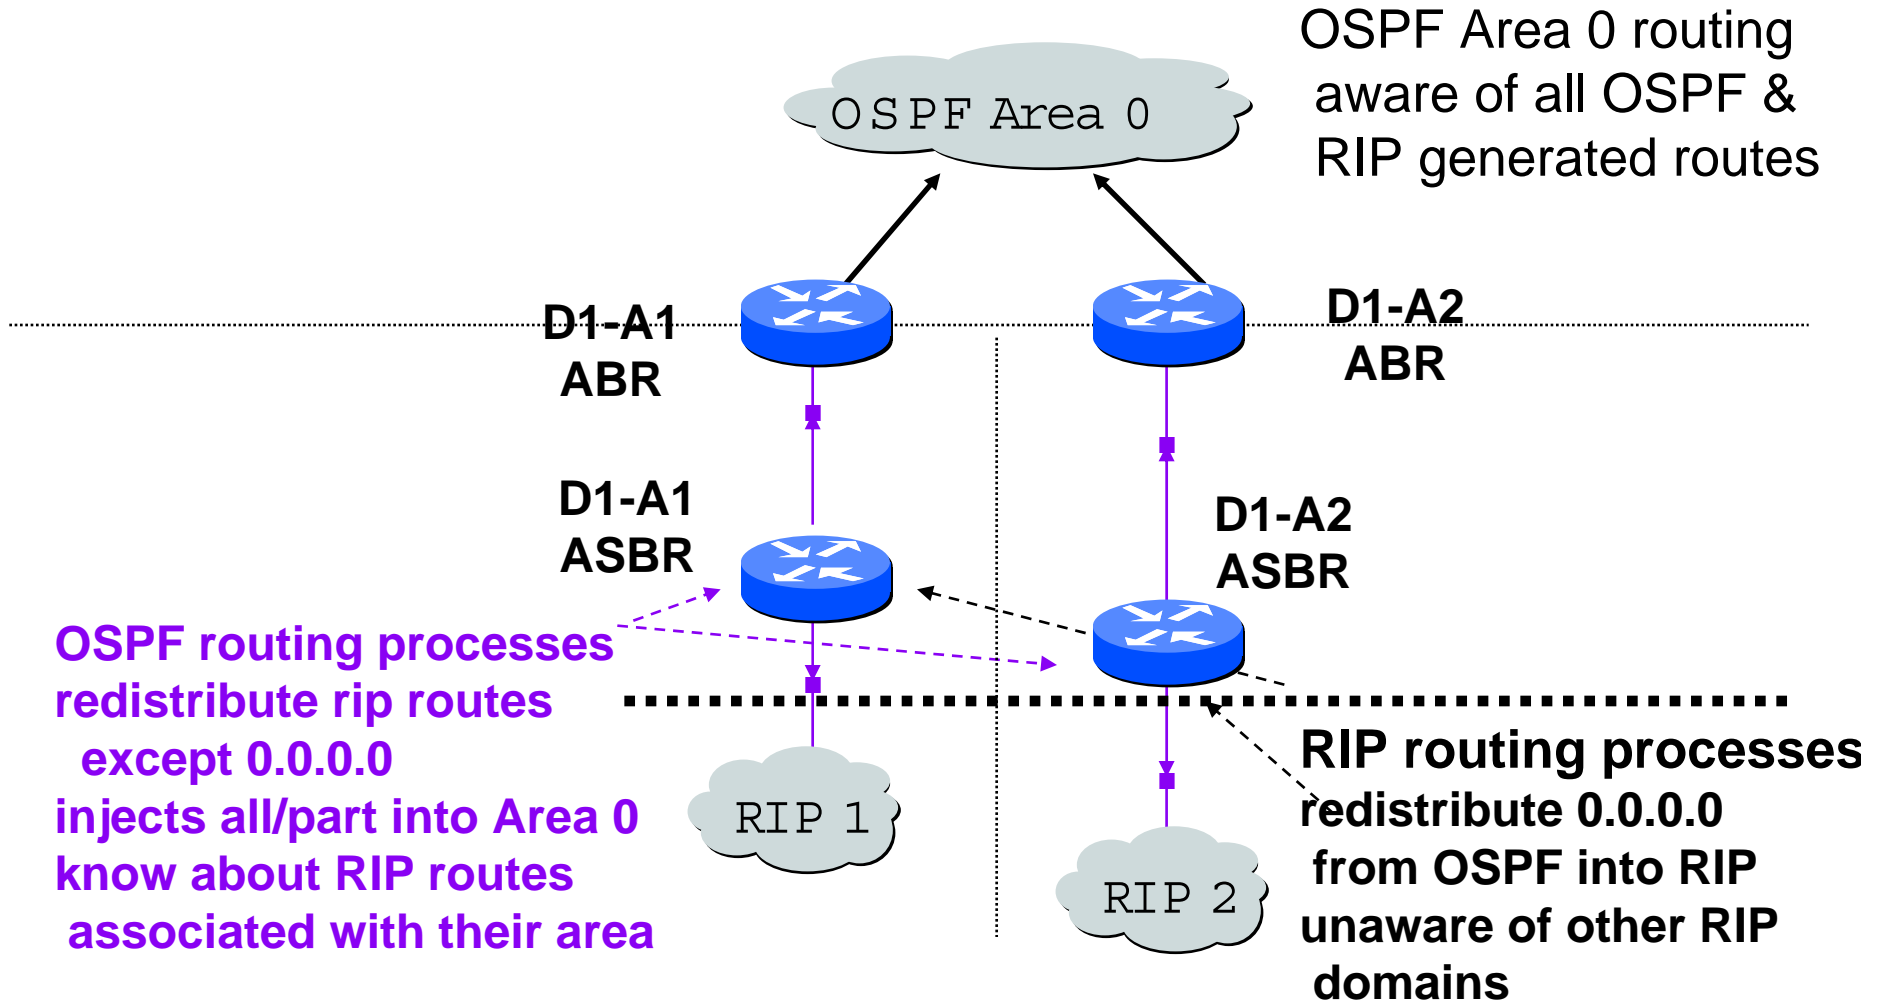


Special  
Area  
D1-A1  
ASBR

Other  
Areas



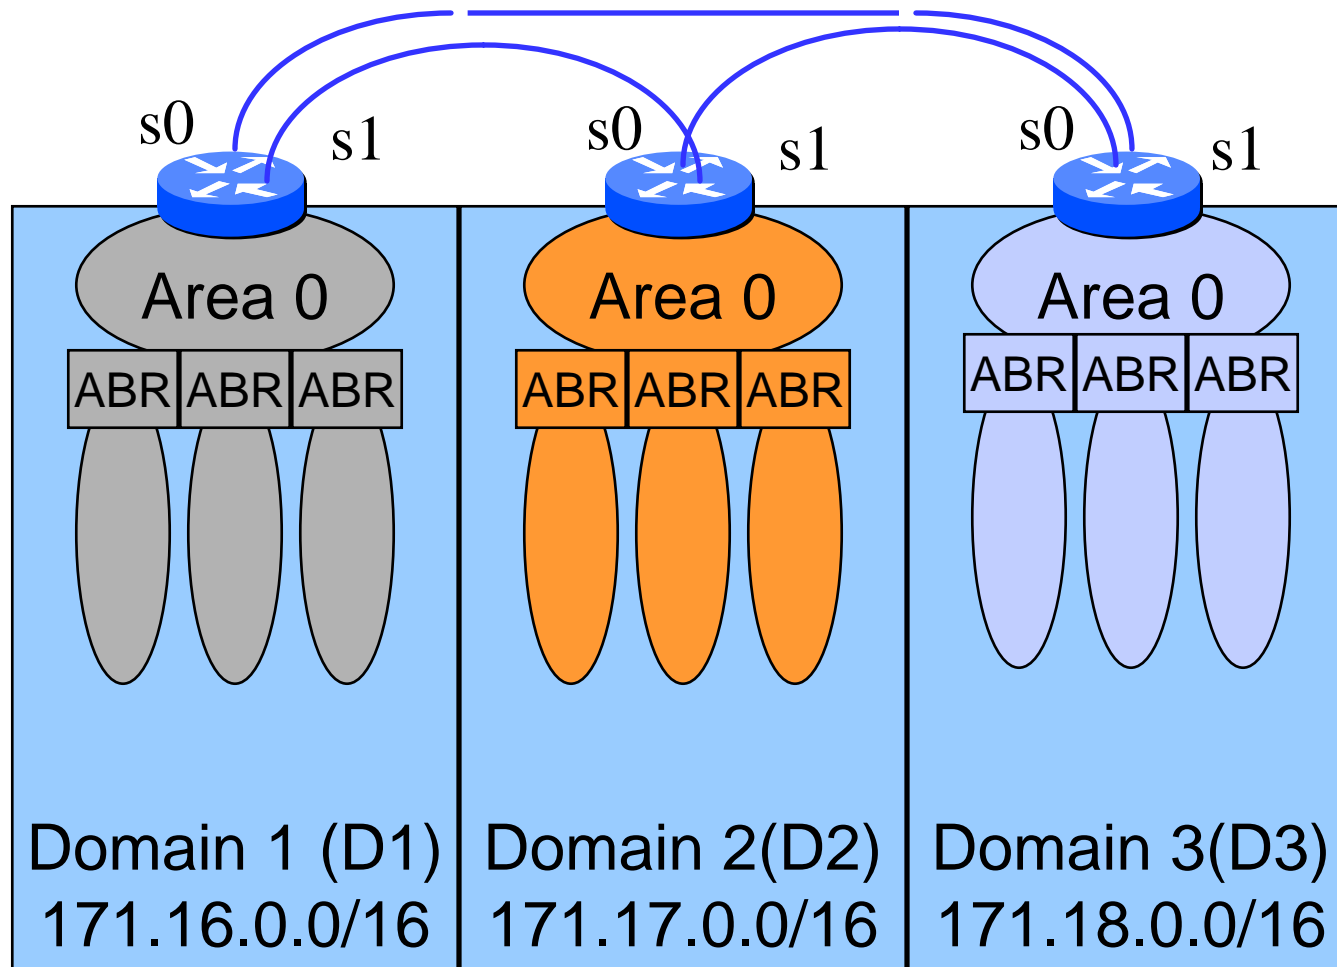
# RIP Migration




## **Backbones of Backbones**

# Default origination


Domains have clean addressing;  
static routes between



# iBGP linkage of enterprise routing domains



# eBGP linkage of enterprise routing domains



## [For Additional Information](#)

- **Current version of**  
**[http://www.ietf.org/internet-drafts/  
draft-ietf-ospf-deploy-00.txt](http://www.ietf.org/internet-drafts/draft-ietf-ospf-deploy-00.txt)**  
**[draft-berkowitz-multirqmt-01.txt](http://www.ietf.org/internet-drafts/draft-berkowitz-multirqmt-01.txt)**
- **Contact author for seminar information**

## Books

- **H. Berkowitz: *Designing Addressing Architectures for Routing and Switching (Macmillan 1998)***
- **H. Berkowitz: *Designing Routing & Switching Architectures for Enterprise Networks (Macmillan 1999)***
- **T. Thomas: *OSPF Network Design Solutions (Cisco Press, 1998)***
- **J. Moy: *OSPF: Anatomy of an Internet Routing Protocol. (Addison-Wesley, 1998)***