The Internet Ecosystem and Evolution

Contents

- The architecture of the Internet
  - Autonomous Systems (AS) and service providers
  - content/eyeball/transit AS
  - access/edge/core AS
- AS-AS business relationships
  - transit service: definition, notation
  - single-homed/multi-homed AS, PA/PI addressing

Architecture

The Internet: the “network of networks”

- ~50 thousand networks of service providers, acedemia, enterprise, governments, etc., connected into a common infrastructure
- 10 billion connected devices
- 3.5 billion users
- Beyond 100 billion USD business revenue

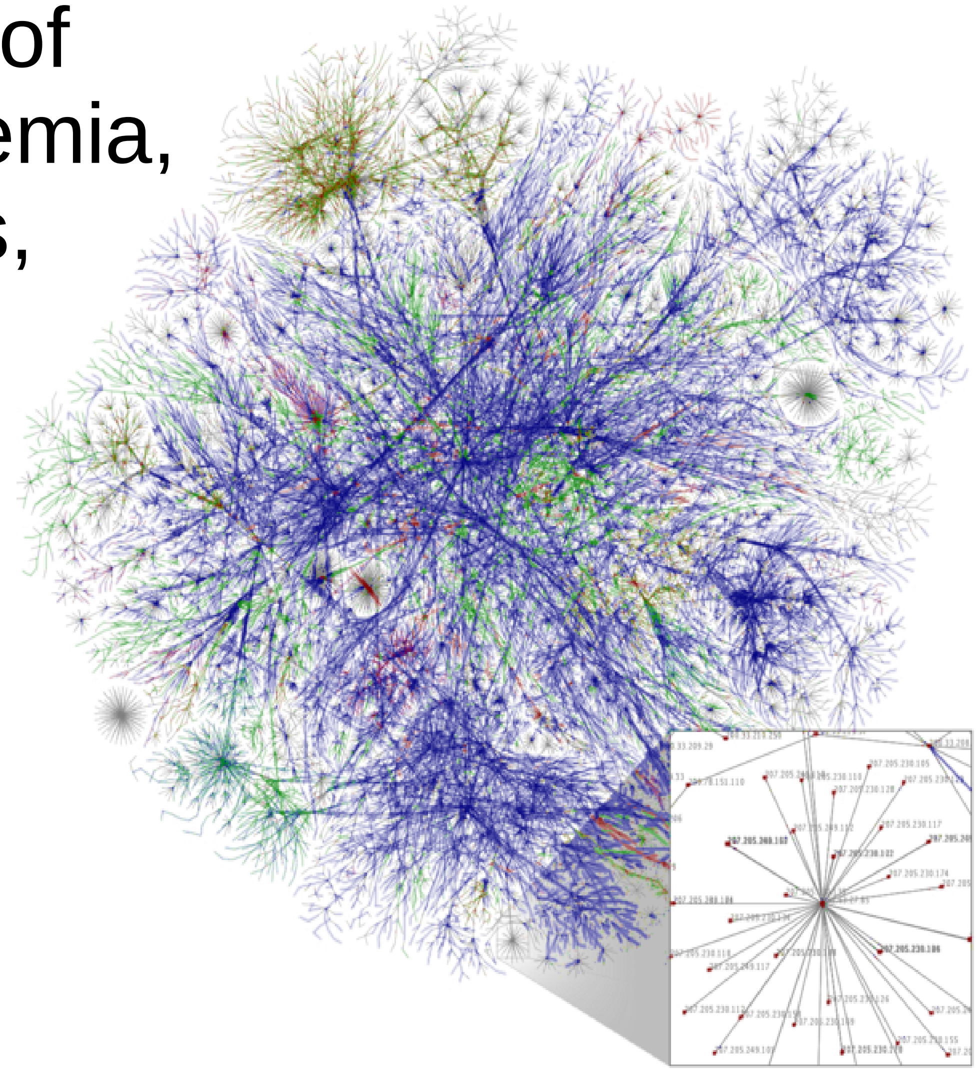

Opte Project, Wikipedia

Autonomous Systems

- The Internet connects the networks of providers, university campuses, telcos, etc., into a unified network of networks
- Every stakeholder manages and administers his/her own network autononously, independently from the rest of the actors
- **Autonomous System (AS):** a set of hosts, routers, and networks that an organization
  - owns and administers as a unit
  - exposes a unified routing policy to the rest of the Internet
- Endued with a unique 32 bit AS identifier

Autonomous Systems

- `AS1` contains routers `R1` and `R2`, switch `SW1`, and hosts `C1` and `C2`

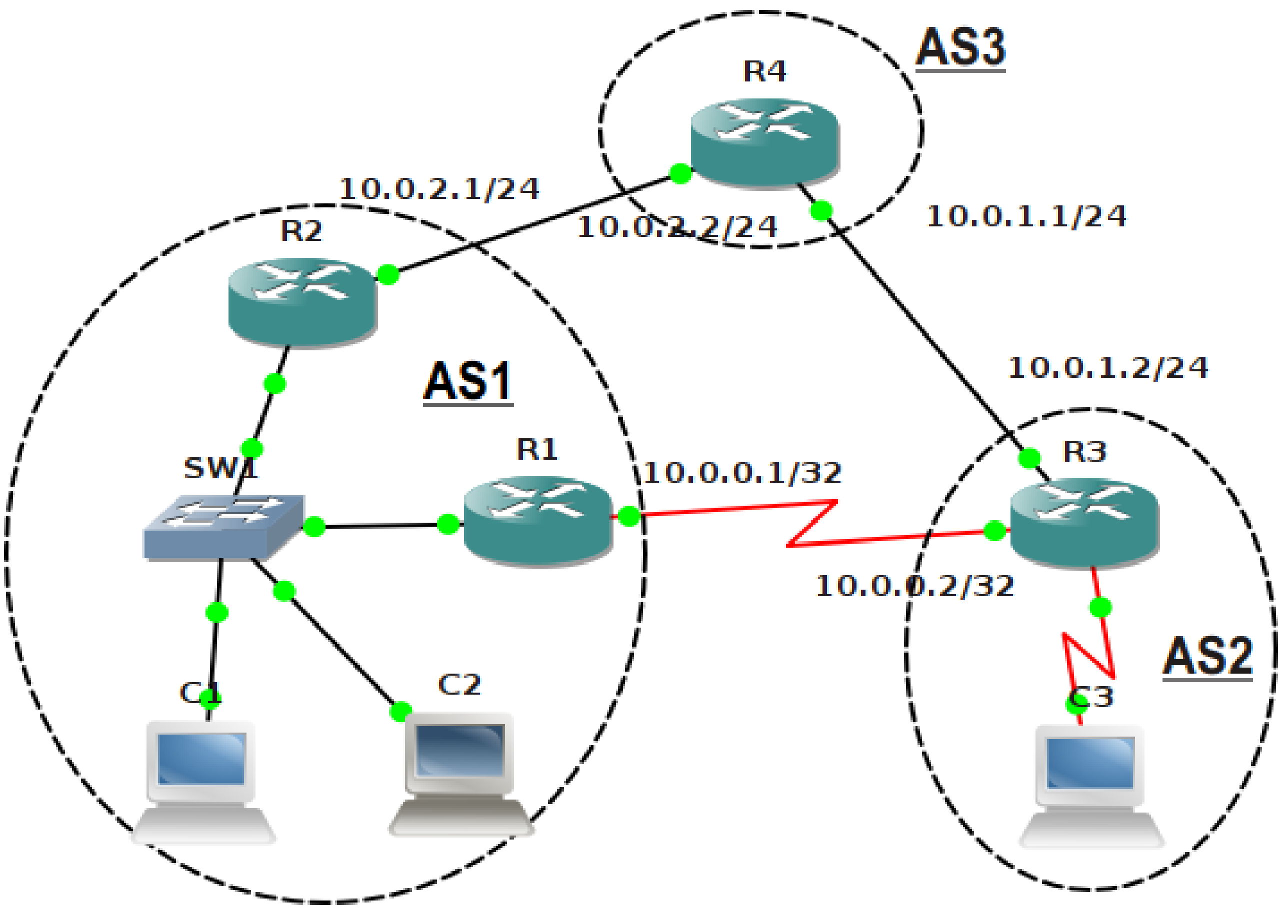

Autonomous Systems

- **ISP (Internet Service Provider):** Comcast (`AS6161`), Telenor Hungary (`AS8448`)
- **Campus:** BME (`AS2547`), Harvard (`AS11`), MIT (`AS3`), UC San Diego (`AS7377`)
- **Enterprise:** IBM (`AS547`, `AS763`), Apple (`AS714`)
- **Global service provider:** Sprint (`AS1239`, `AS1240`, `AS6211`, `AS6242`, ...), Cogent (`AS174`, `AS2149`, `AS6494`), TeliaSonera (`AS1759`)
- One ISP can provide service under multiple separate ASes...

Intra-AS and Inter-AS routing

- The Interent routing system handles an AS as a unity, so usually all IP prefixes owned by an AS get the same treatment
- Traffic can either flow though an AS between any two prefixes of two other ASes undisruptedly, or it can not flow at all
- Realized by **routing protocols**, reponsible for setting up the IP forwarding tables at routers
- Routing inside an AS (intra-domain routing):
  - **Interior-Gateway Protocol (IGP):** OSPF,RIP
- Routing between ASes (inter-domain routing):
  - **Exterior Gateway Protocol (EGP):** BGP

AS: Routing strategies

- Example: `AS2` forwards the `AS1`–`AS3` traffic
- AS is a policy unit: `AS2` forwards packets from any host/prefix of `AS1` to any host/prefix of `AS2`
- But `AS2` may block `AS1–AS4` traffic in its routers

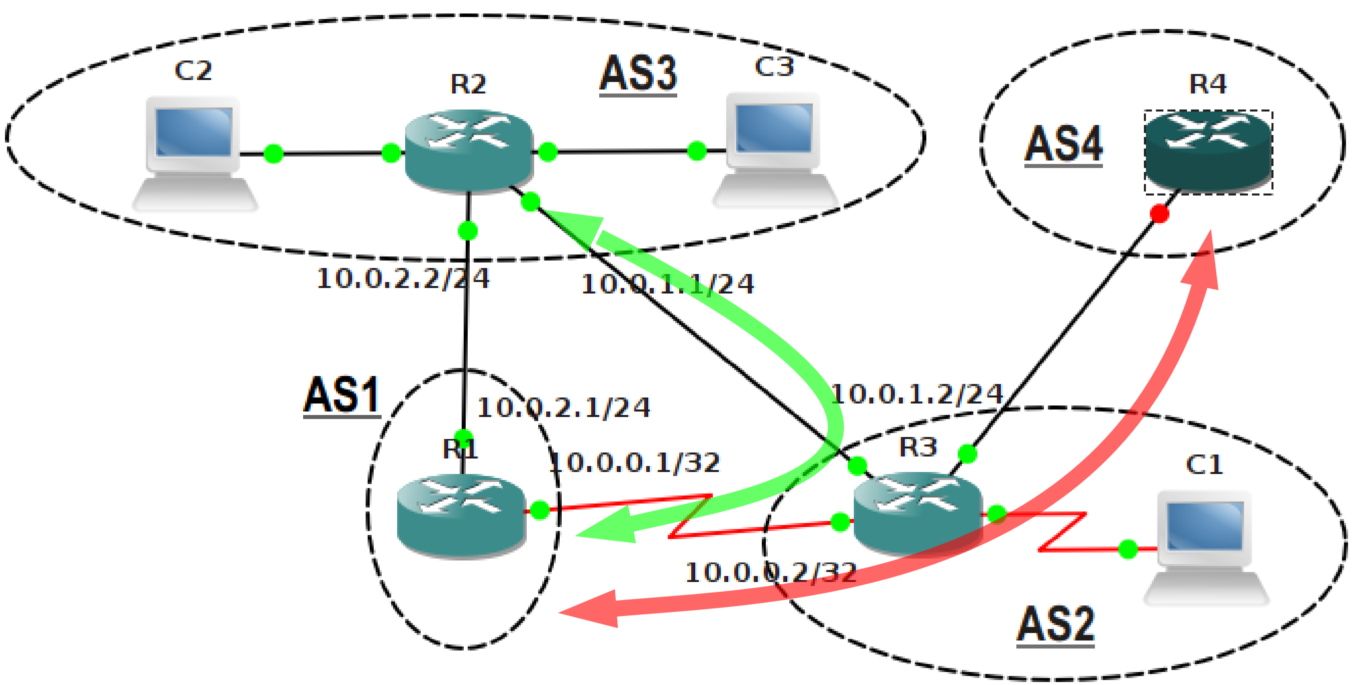

AS taxonomy 1

- **Content AS:** content provider AS
  - egress (out-from-AS) traffic is dominant
  - Youtube, Netflix, HULU
- **Eyeball AS:** content consumption
  - ISPs serving a huge base of customers
  - ingress (into-the-AS) traffic is dominent, needs small latency to the content
- **Transit AS:** global transit between the above via a world-wide network infrastucture (Cogent,Level3)
- **CDN** (Content Delivery Network): Akamai
  - content + global transit

AS taxonomy 2: Edge/Access/Core

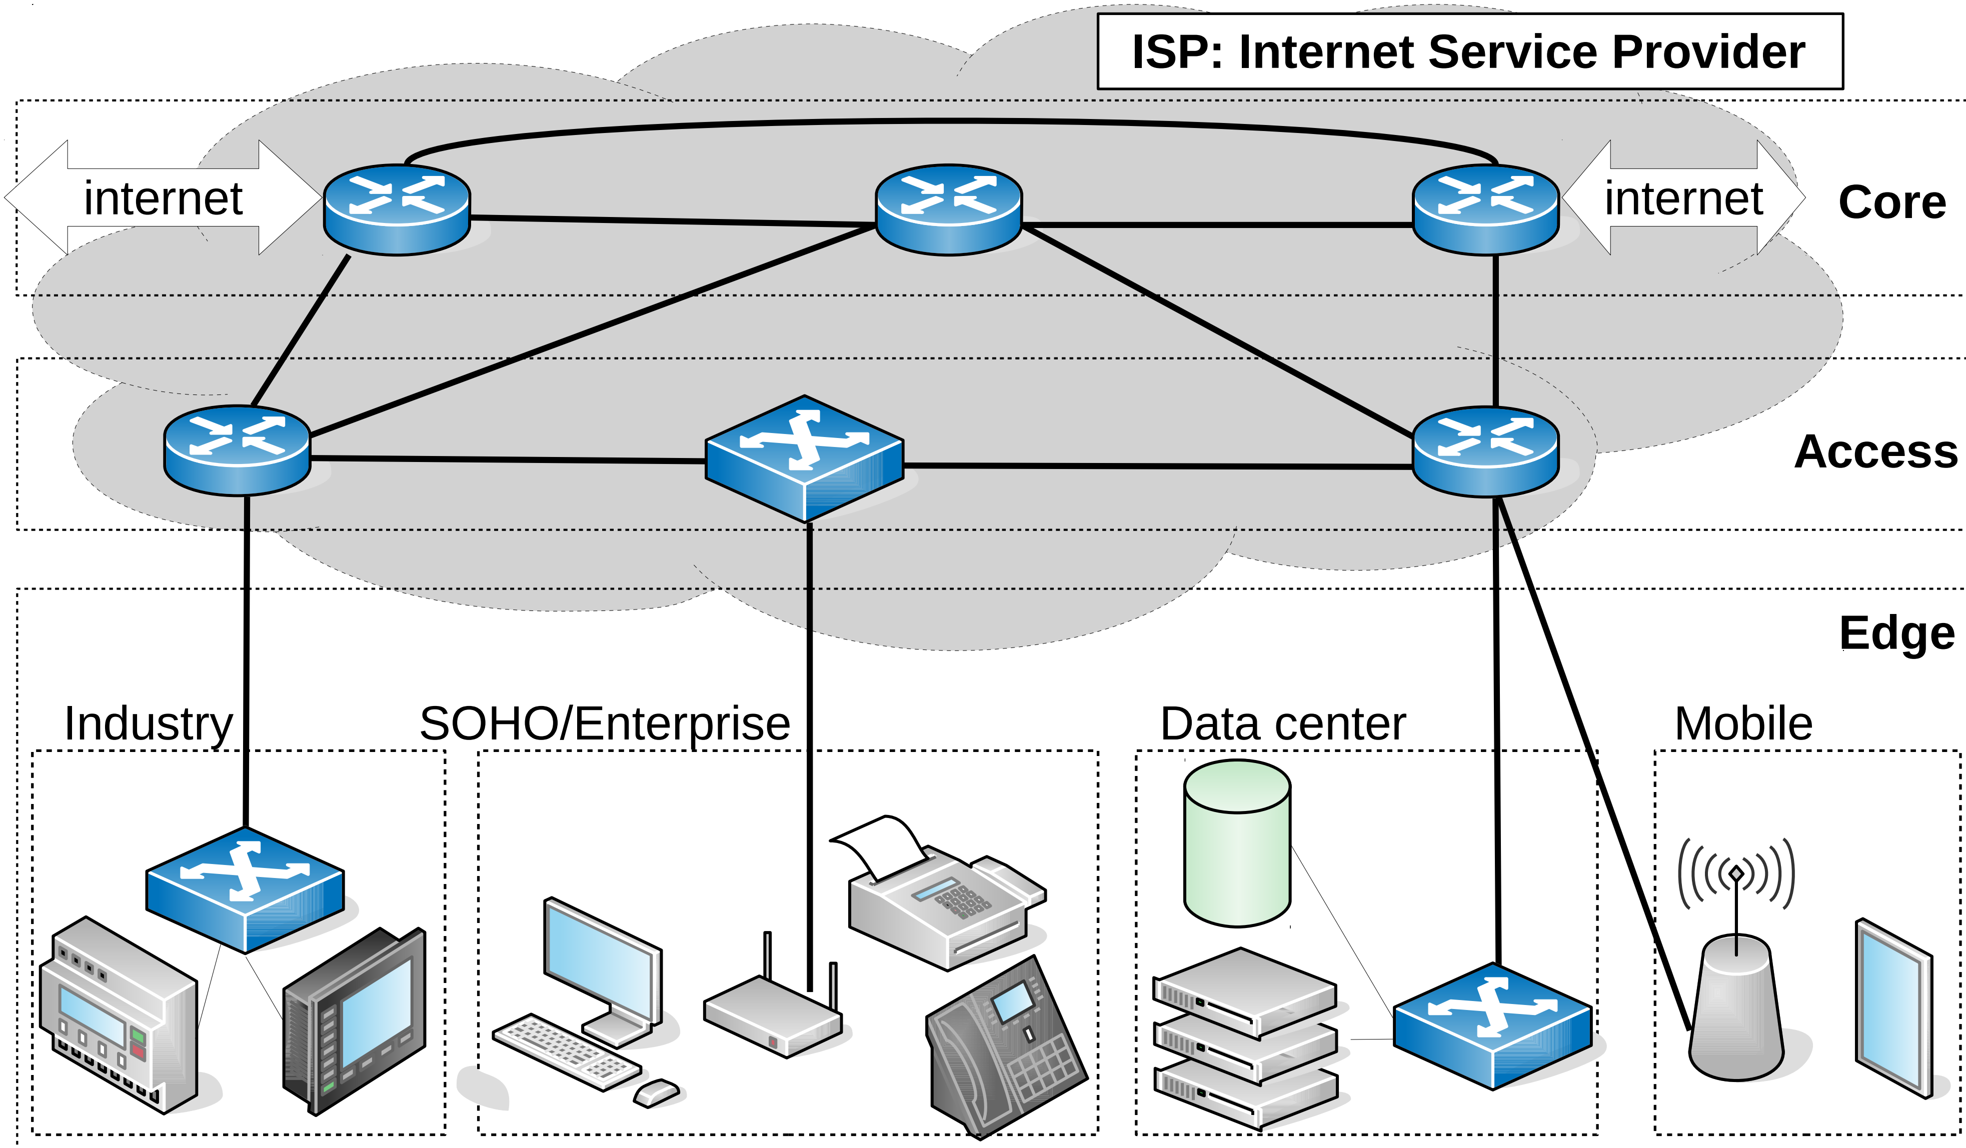

Edge

- End-devices, terminals, hosts: desktops, laptops, mobiles, PDAs, tablets, etc.
- Residential subscribers, **SOHO** (Small Office/Home Office), mobile, data center, ...
- Local networks, connected to a service provider

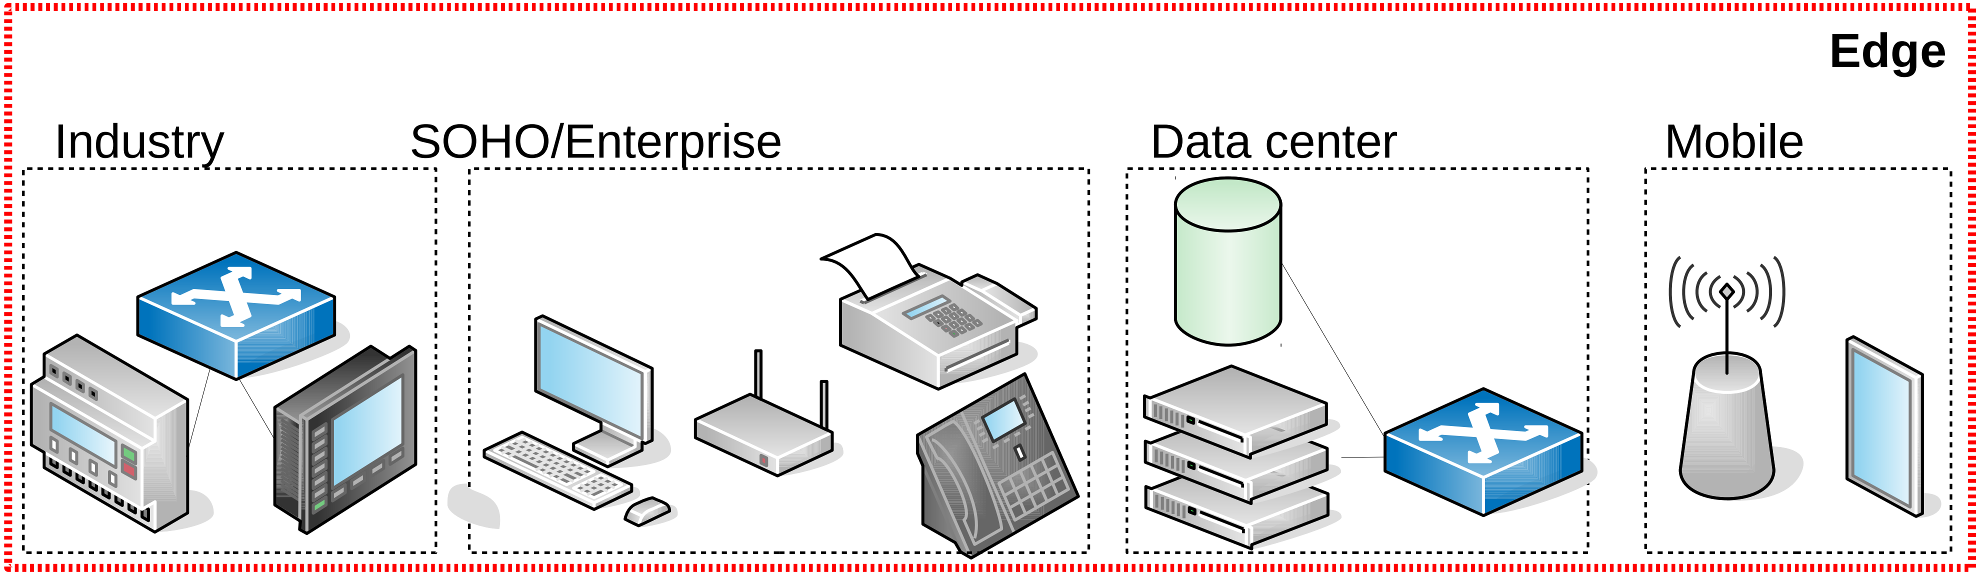

Access

- Concentrates traffic from the Edge to the Core
- **Access router:** first-hop ISP router
- **Access link:** transit medium to the access router

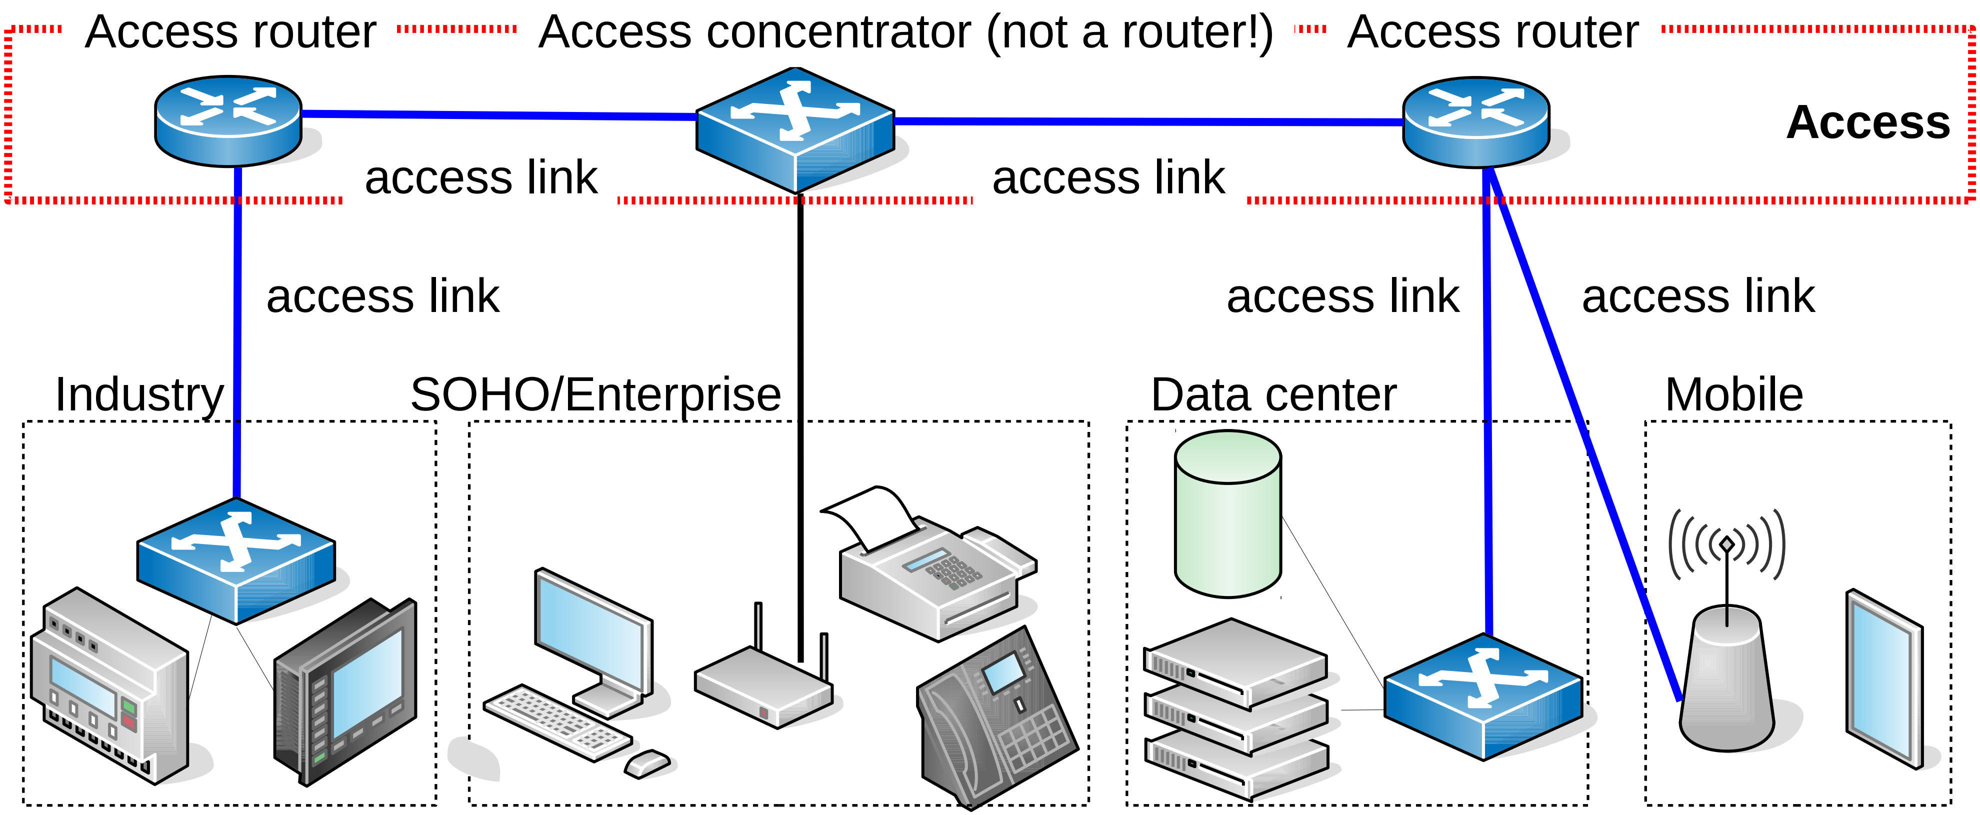

Internet access

- **CPE** (Customer-premises Equipment): a (often ISP-owned) terminal located at the customer site connecting it to the ISP network
- **Demarcation line:** the boundary between the networks of the provider and the customer

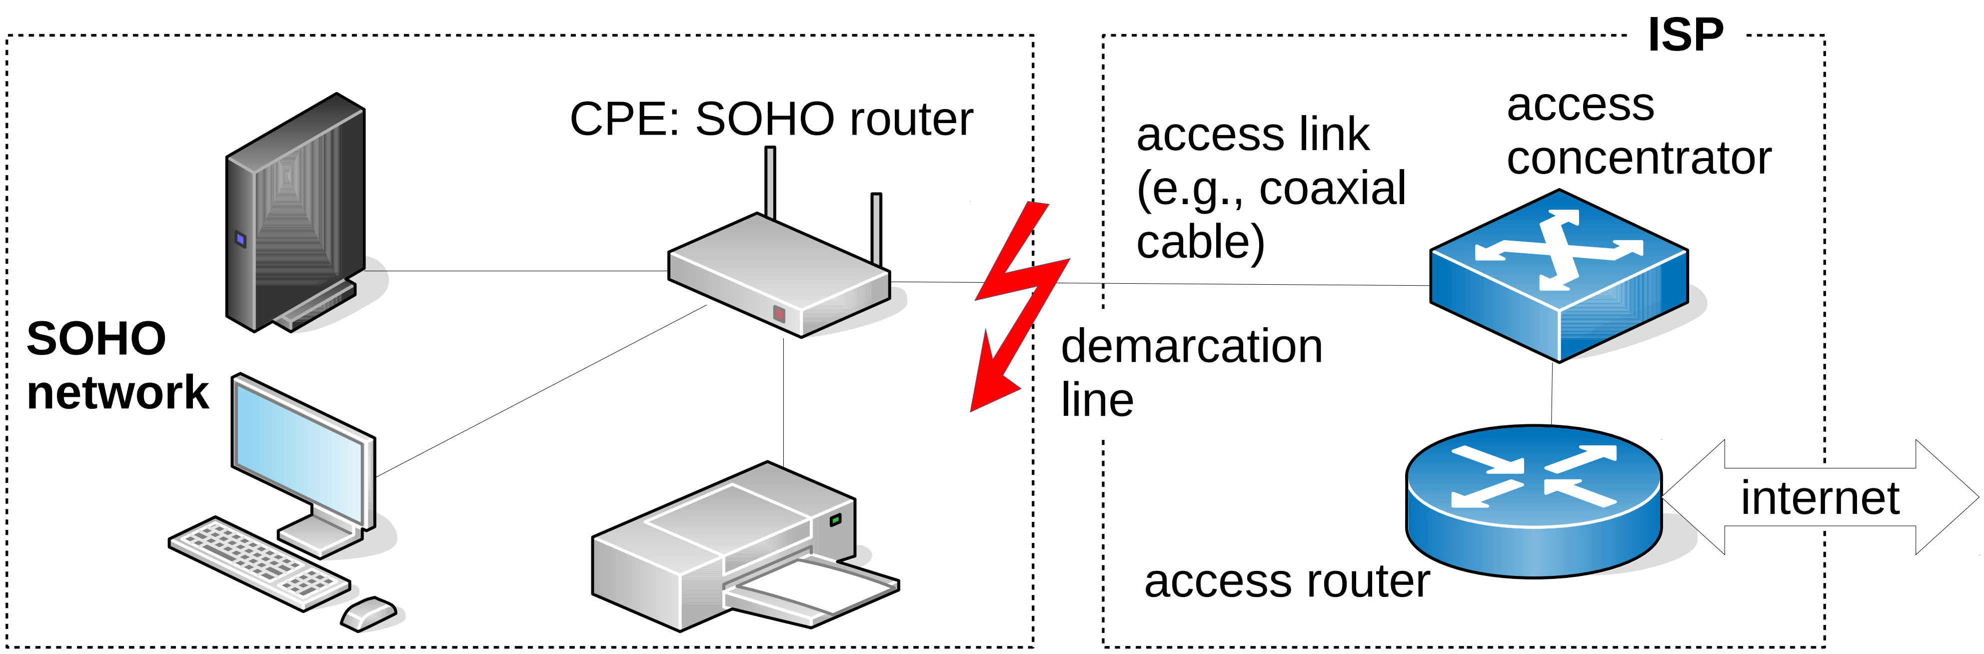

Residential/SOHO access

- Dial-up: POTS (Plain Old Telephone Service)
  - CPE: analog model
  - access link: telehone line
  - ISP/telco access: POTS subscriber card
- DSL (Digital Subscriber Line): ADSL/VDSL
  - CPE: DSL modem
  - access link: local loop, twisted pair
  - ISP/telco access: DSLAM (Digital Subscriber Line Multiplexer)

Residential/SOHO access

- Cable(TV):
  - CPE: CATV cable modem
  - link: coaxial or HFC (Hybrid Fiber Coax)
  - ISP/telco access: cable headend
- FTTx (Fiber-to-the-X): access via optics
  - FTTH (Fiber-to-the-Home): optics to CPE
  - FTTP (Fiber-to-the-Premises): optics to building, copper access from there
  - technology: PON (Passive Optical Network)

Enterprise/campus access

- **Enterprise:** business, corporate, factory,...
- **Campus:** site consisting of multiple buildings
  - education, military, government, etc.
- Example: university campus Ethernet access

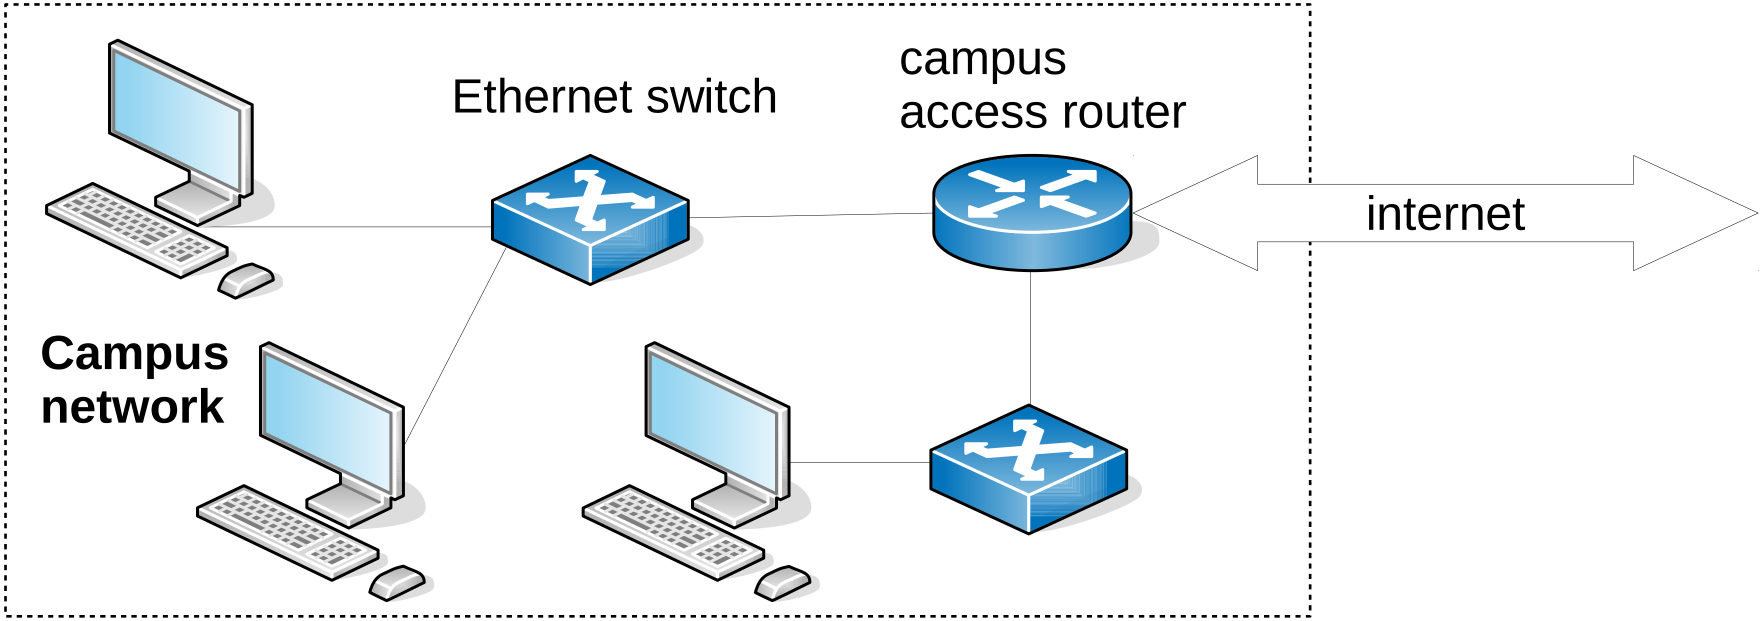

Mobile access (2G, 3G, 4G, WiMAX)

- Internet access over cellular mobile
- 2.5G: General Packet Radio Service (GPRS) packet communicaiton in on the radio interface
- Later: EDGE, UMTS, 4G, 5G, etc.

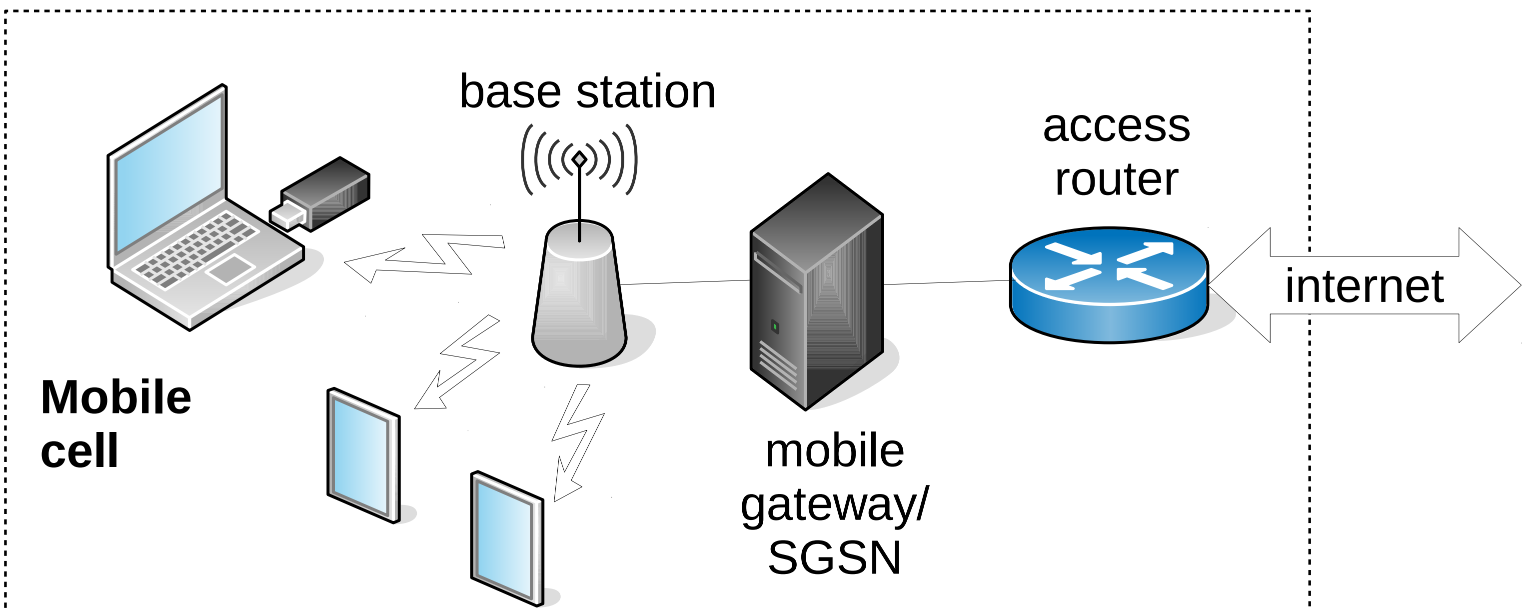

The Internet backbone: Core

- The Access concentrates traffic from the Edge to the Intrernet Core
- The core is made up of high-performance, ultra-efficient (>100Tbps) routers connected into a dense mesh topology
- The most important segment of the Internet by a large margin: our main tpoic

Cisco CRS-1, wikipedia

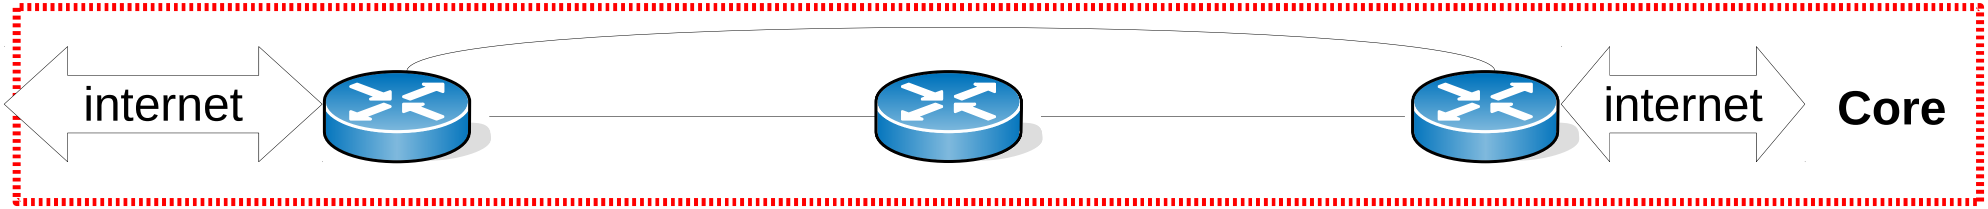

A typical AS: The ISP

- **Internet Service Provider:**
  - main profile: providing two-way Internet access to a large customer base
  - auxiliary services: domain name registration, email, cloud storage, server colocation,...
- ISP networks are made up by multiple POPs
- **POP (Point of Presence):** a set of routers, switches, etc., operated by the ISP at some physical location (e.g., a data center) to provide a service access point to its networks/services

A global ISP

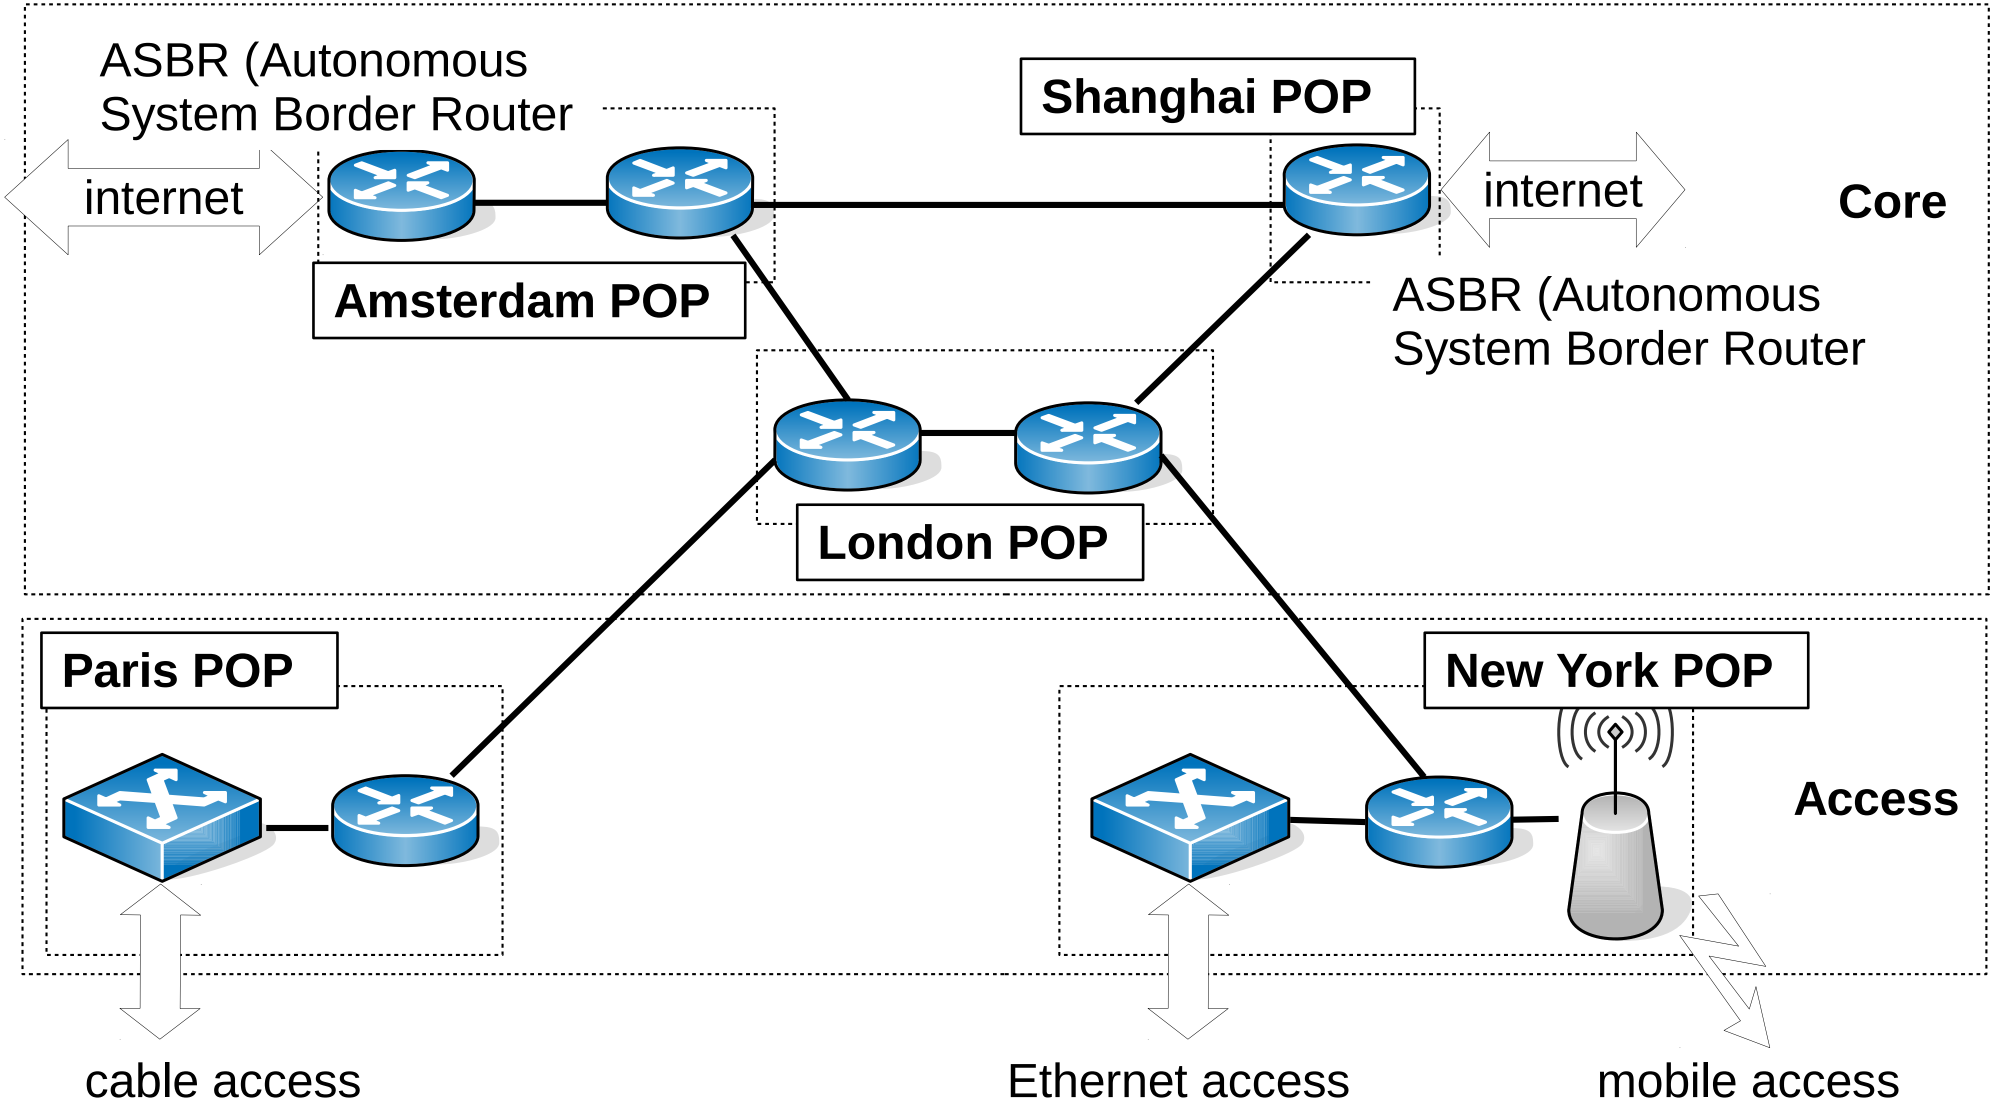
ASBR (Autonomous
System Border Router
internet
Shanghai POP
internet
Core
Amsterdam POP
ASBR (Autonomous
System Border Router
London POP
Paris POP
New York POP
Access
cable access
Ethernet access
mobile access

Internet service model

Internet: Network of ASes

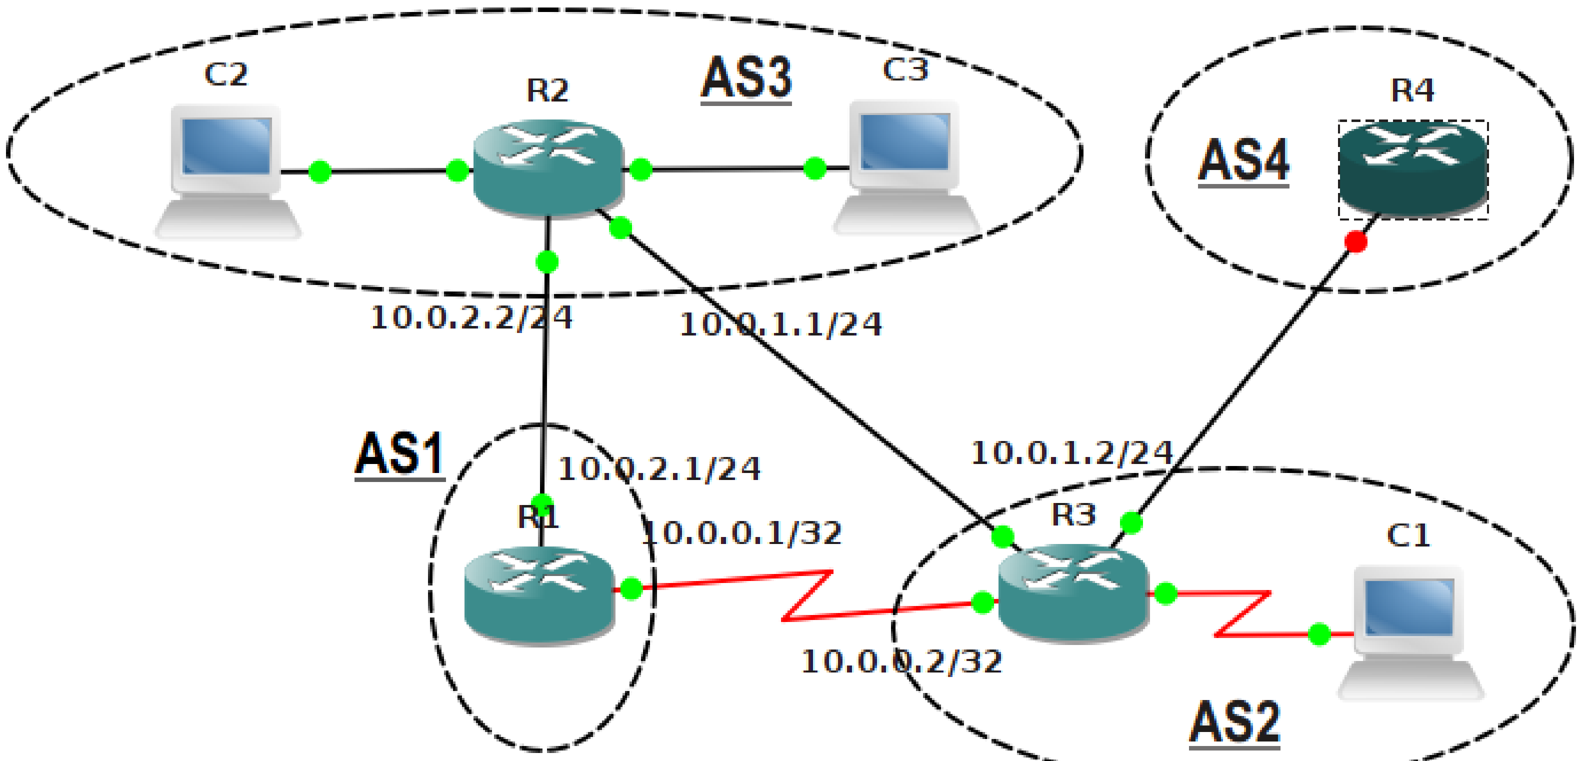

Connecting ASes

- The Internet is shaped by the way
  - **ASes physically connect** their networks and
  - ASes **route traffic along** these links
- It is the **business interests, political considerations, and reliability and security issues,** and not technical reasons like path length, that govern how ASes connect and route traffic
- **Policy routing:** packet forwarding based on policy considerations (business, political, security, etc.)
- Internet traffic does not follow shortest paths (**path inflation**)!

Routing strategies

- **Service model:** how an AS decides whether to forward or block other ASes' traffic
- For instance, `AS2` may forward `AS1–AS3` traffic but it would block `AS1–AS4` traffic entirely

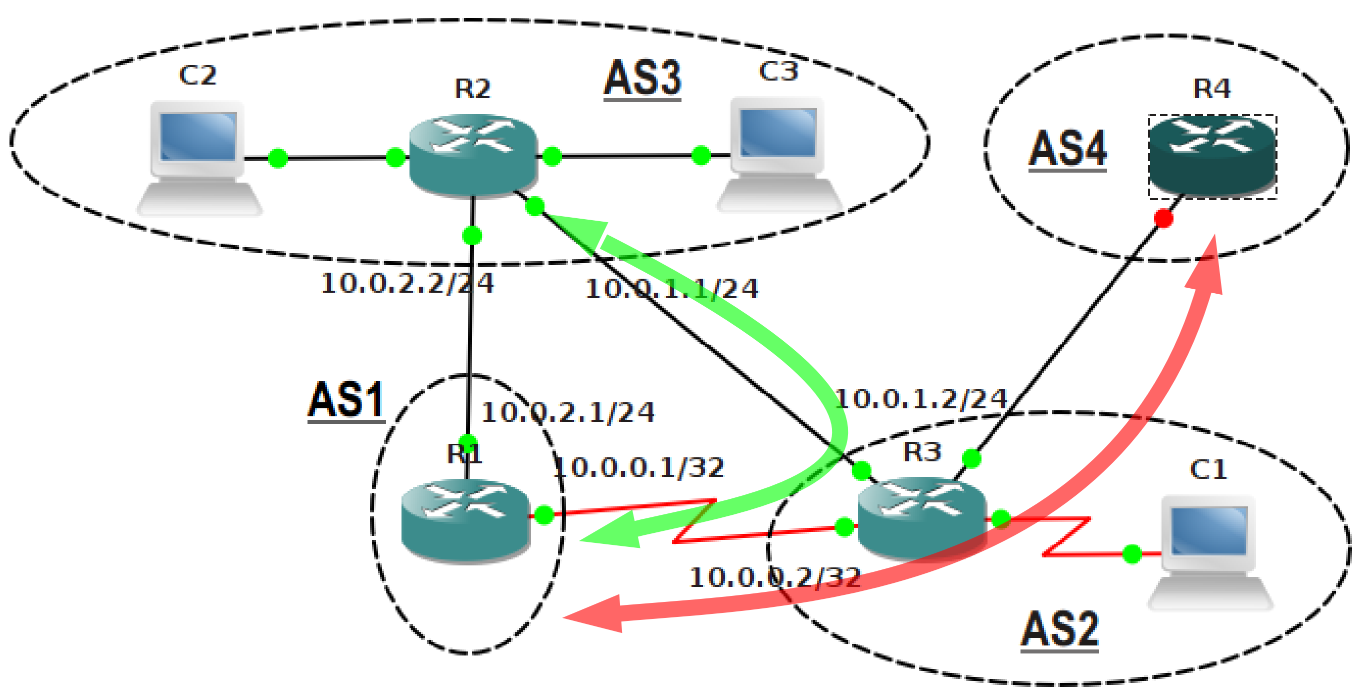

Inter-domain routing

- ASes lease IP address ranges from registries
- These ranges are announced by the routing protocol, which ensures that the announcement is communicated to other ASes
- Routers select paths

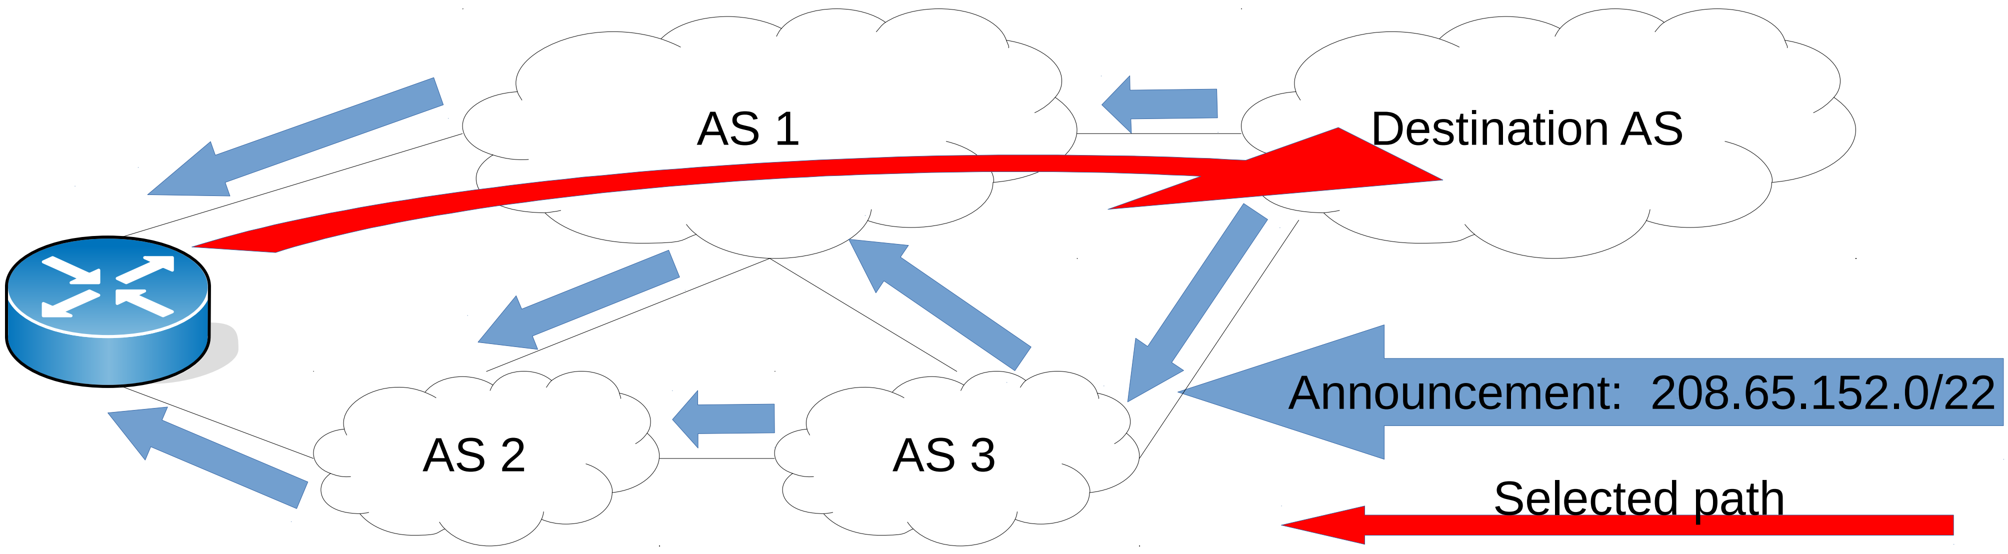

The provider-customer (transit) AS-AS relationship

The transit service

- Simplest way to connect to Autonomous Systems (ASes)
- One AS acts as a service **provider**, the other as a **customer**
- The provider is contracted
  - to deliver all traffic from the customer AS to any host connected to the Internet (**egress**), and
  - to deliver all traffic from the Internet destined to the customer (**ingress**) back to it
- **The customer is charged** by the rate of provider ↔ customer traffic flowing on the inter-AS link

The transit service

Internet Transit is the business relationship whereby an Internet Service Provider provides (usually sells) access to the global Internet.

From a high-level perspective, Internet Transit can be thought of as a pipe in the wall that says "Internet this way". Customers connect their networks to their Transit Provider, and the Transit Provider does the rest.

(DrPeering)

Transit: Addressing

- The provider can lease (delegate) an address prefix for the customer to assign IP addresses to customer hosts
- The task of the transit provider is then to
  - to distribute the BGP routing announcement originated by the customer for this prefix on the Internet,
  - to handle routing announcements from the public Internet to the customer's router, and
  - to transfer all „customer↔internet” traffic

Transit: Single-homed ASes

- Simplest way to connect an AS to the Internet: buy transit service from a single provider AS
- **Single-homed (SH) AS:** an AS that connects to the Internet through a single provider AS by buying transit service from that AS

SH AS: Schematic diagram

- **Transit relationships** are marked **by an arrow** directed from the customer to the provider
- The direction matches that of the cash-flow($$)

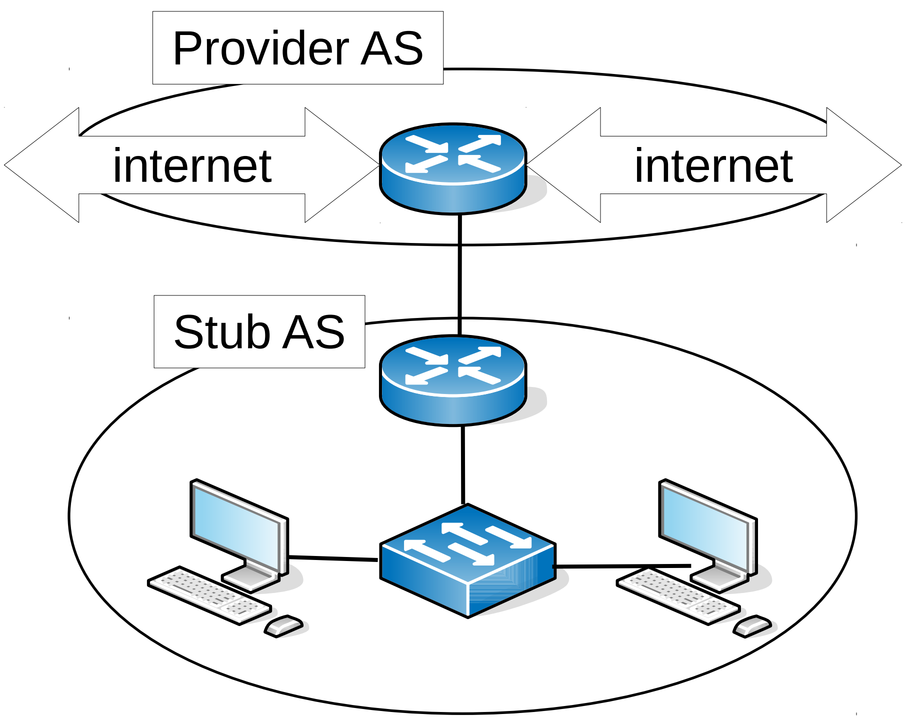

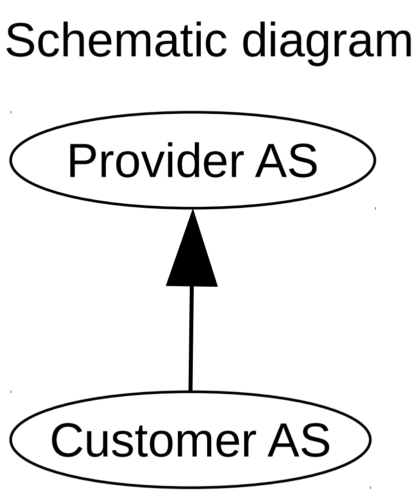

SH transit: Addressing/routing

- **Addressing:** the SH AS typically receives a subnet prefix from the provider's IP address prefix
- **Routing:** the SH AS sets the provider as default gateway
- The provider does not announce the customer's prefix separately, as that is part of its own address range, and hence of its own routing announcement

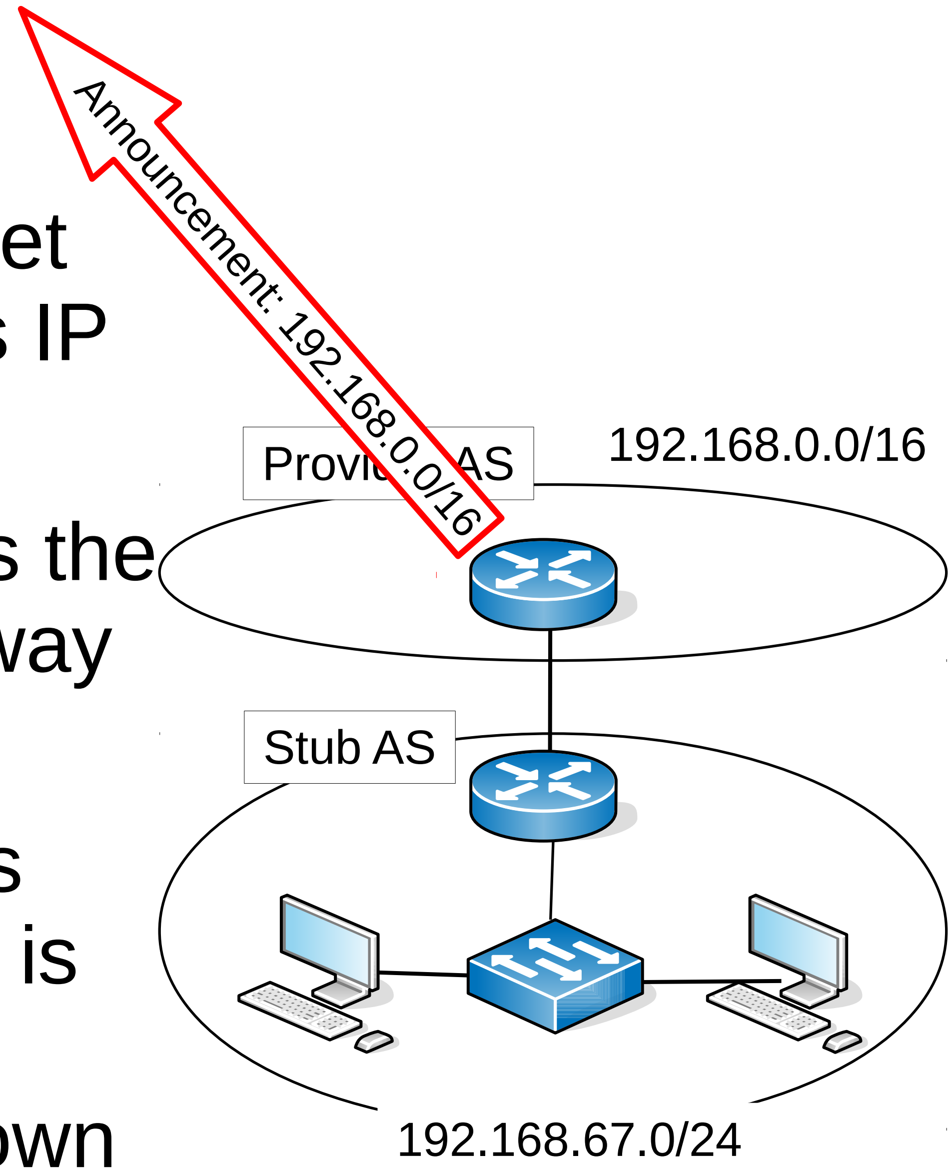

SH transit: Example

- Suppose the `Provider AS` uses the prefix `192.168.0.0/16`
- It can hand out 256 `/24` address ranges

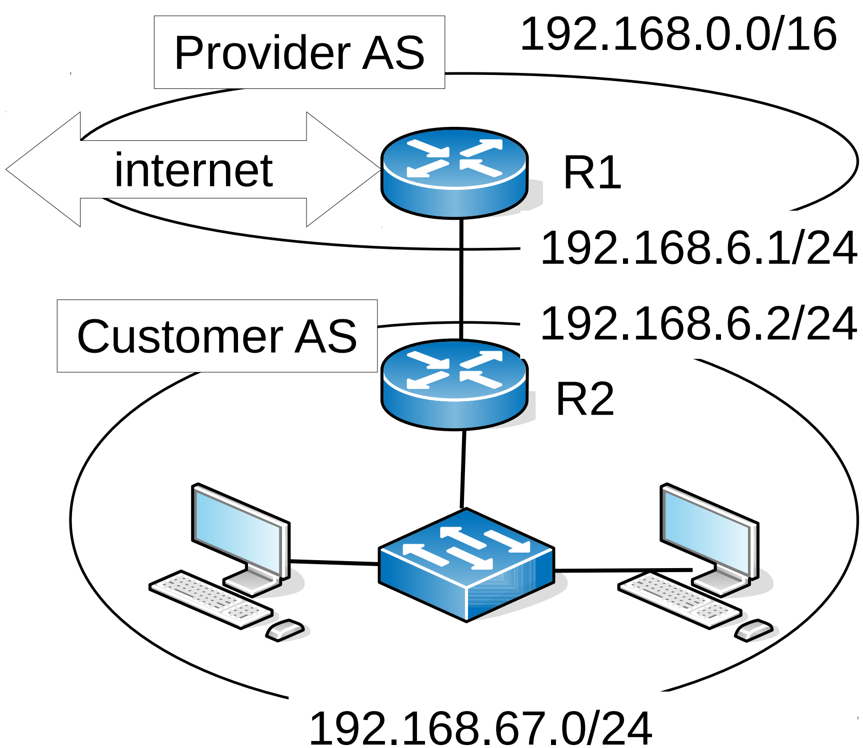

- Suppose `Customer AS` gets the prefix `192.168.67.0/24`
- Address of the border routers:
  R1: `192.168.6.1`
  R2: `192.168.6.2`

SH transit: Example

- The provider does not announce the customer's prefix separately into the routing system, as it is part of its own BGP announcement

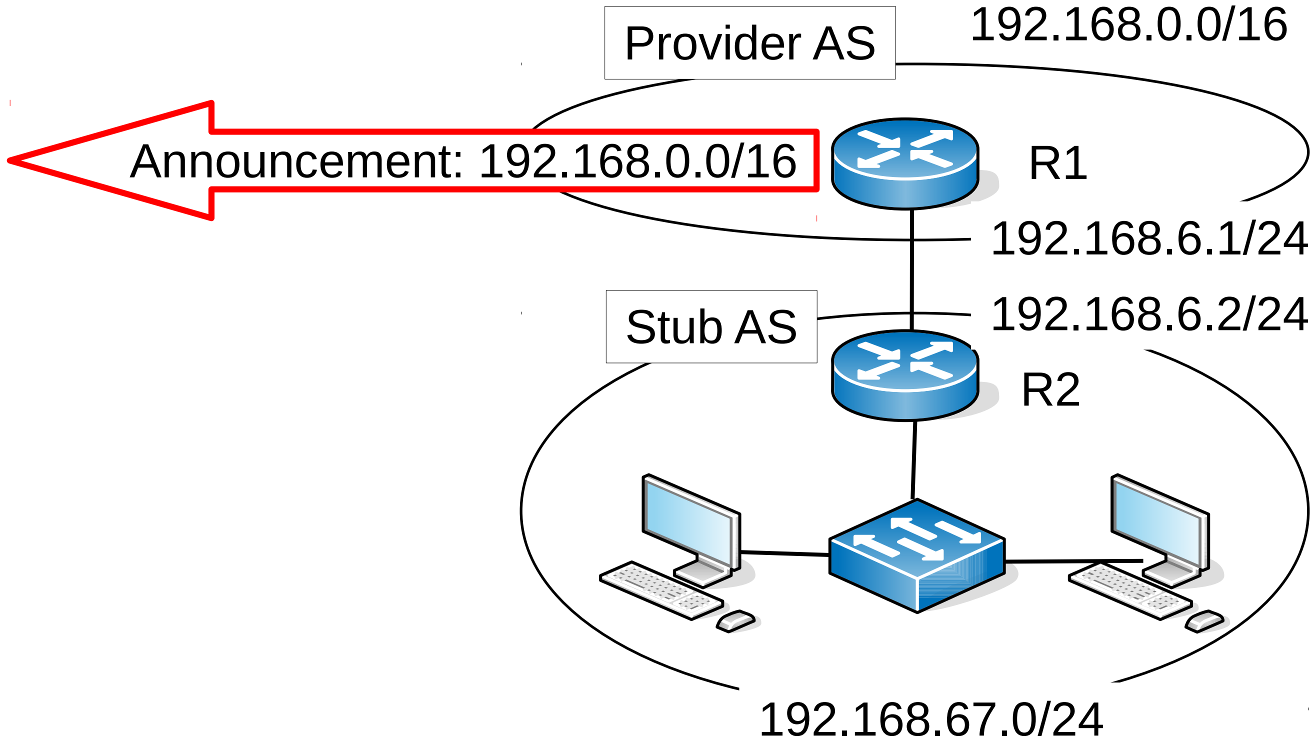

SH transit: Ingress traffic

- The provider forwards the customer's ingress traffic to it by using a **more specific** entry in the FIB of its border routers

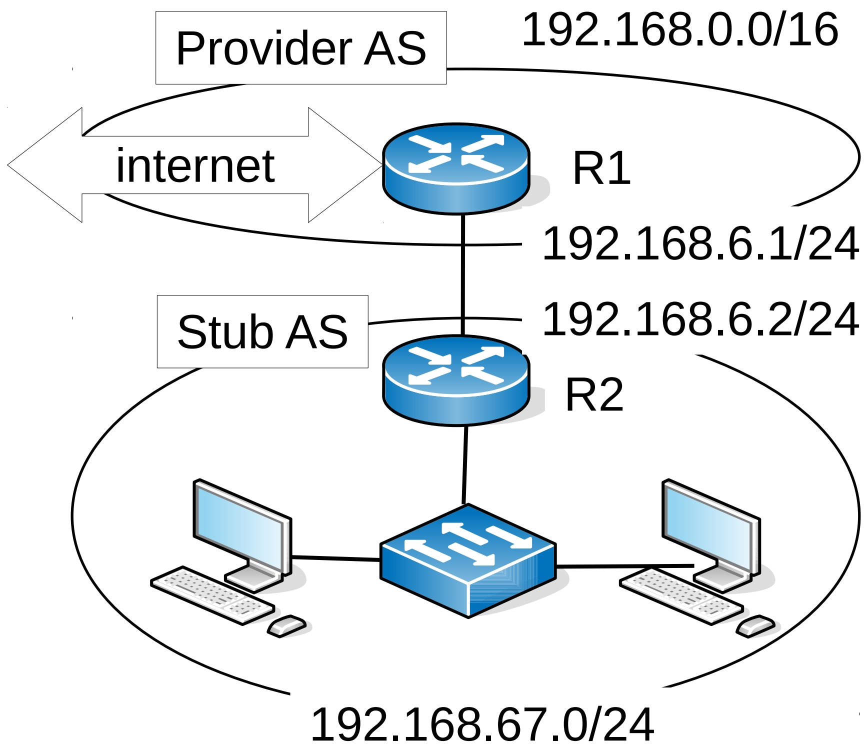

- This will apply to all traffic of the customer, but not to the provider's own traffic
- E.g., the FIB at R1

Prefix	**Next-hop**
...	...
`192.168.67.0/24`	`192.168.6.2`
...	...

SH transit: Egress

- The customer sets a **default gateway** towards the provider (less specific than any other entry)
- Forwards all traffic to the provider whose destination is not **inside** the customer AS

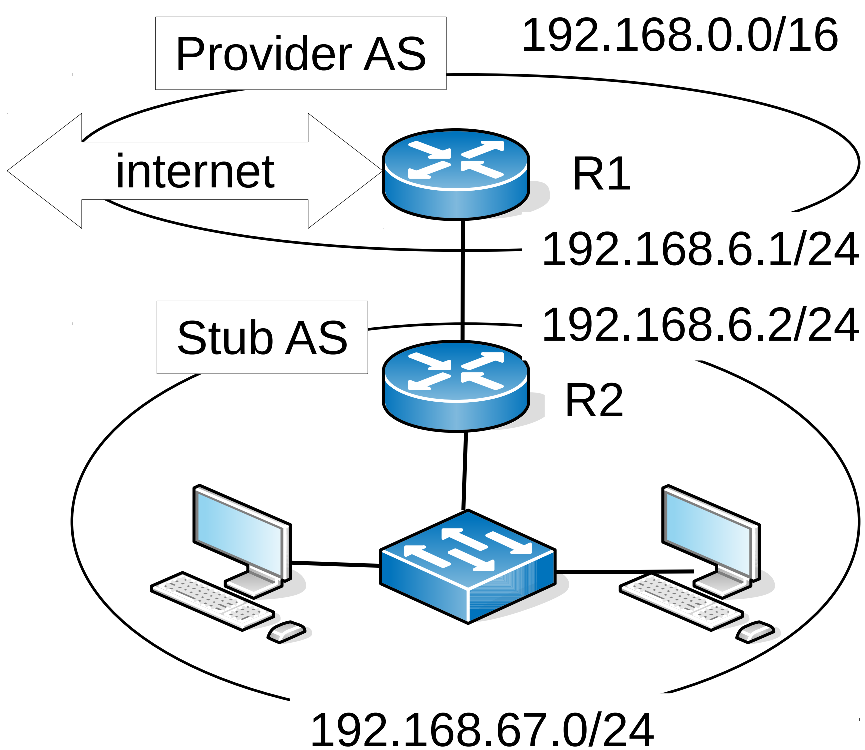

- R2 FIB

Prefix	Next-hop
0.0.0.0/0	192.168.6.1

- The SH AS does not even need to run a routing protocol

Single-homing: Pros and cons

- Advantages
  - minimal configuration and maintenance in the SH AS, does not need expert IT staff
  - does not create a separate, customer-specific routing table entry in the Internet
- Disadvantages
  - single point-of-failure, poor resilience (if the provider goes away, Internet connectivity is lost)
  - can be pricey

Multi-homed ASes

- Internet connectivity of an SH AS depends on a single provider's availability
- **Multi-homed (MH) AS:** connects to the Internet by purchasing transit from more than one providers

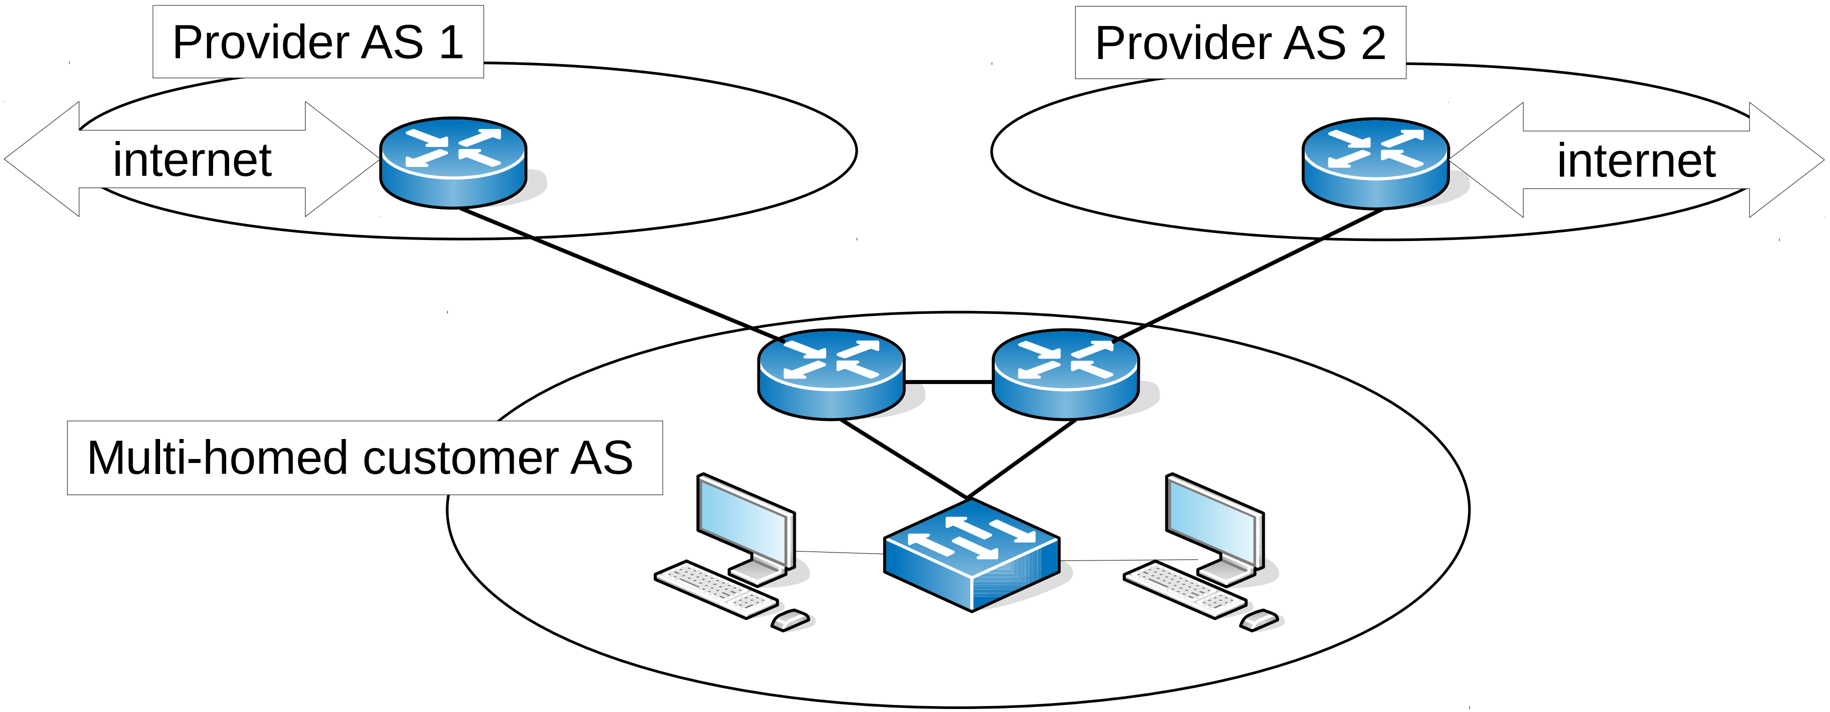

MH AS: Schematic diagram

- **Arrow** from the customer AS **to each provider AS**
- Arrow points to the direction of the cash-flow ($$)

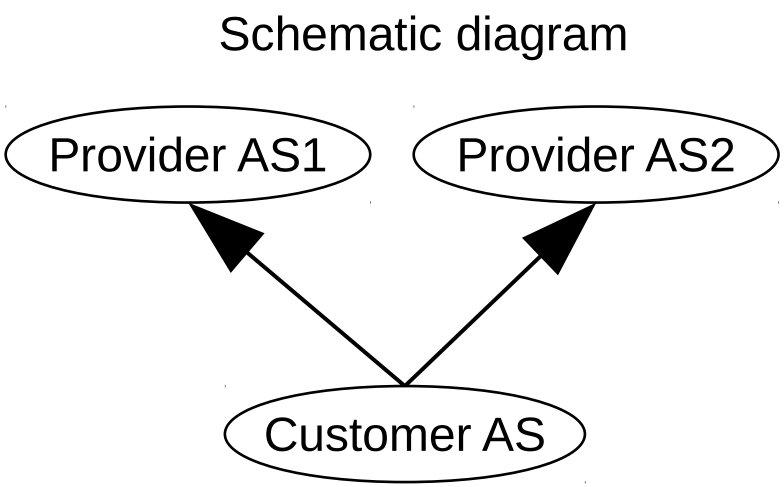

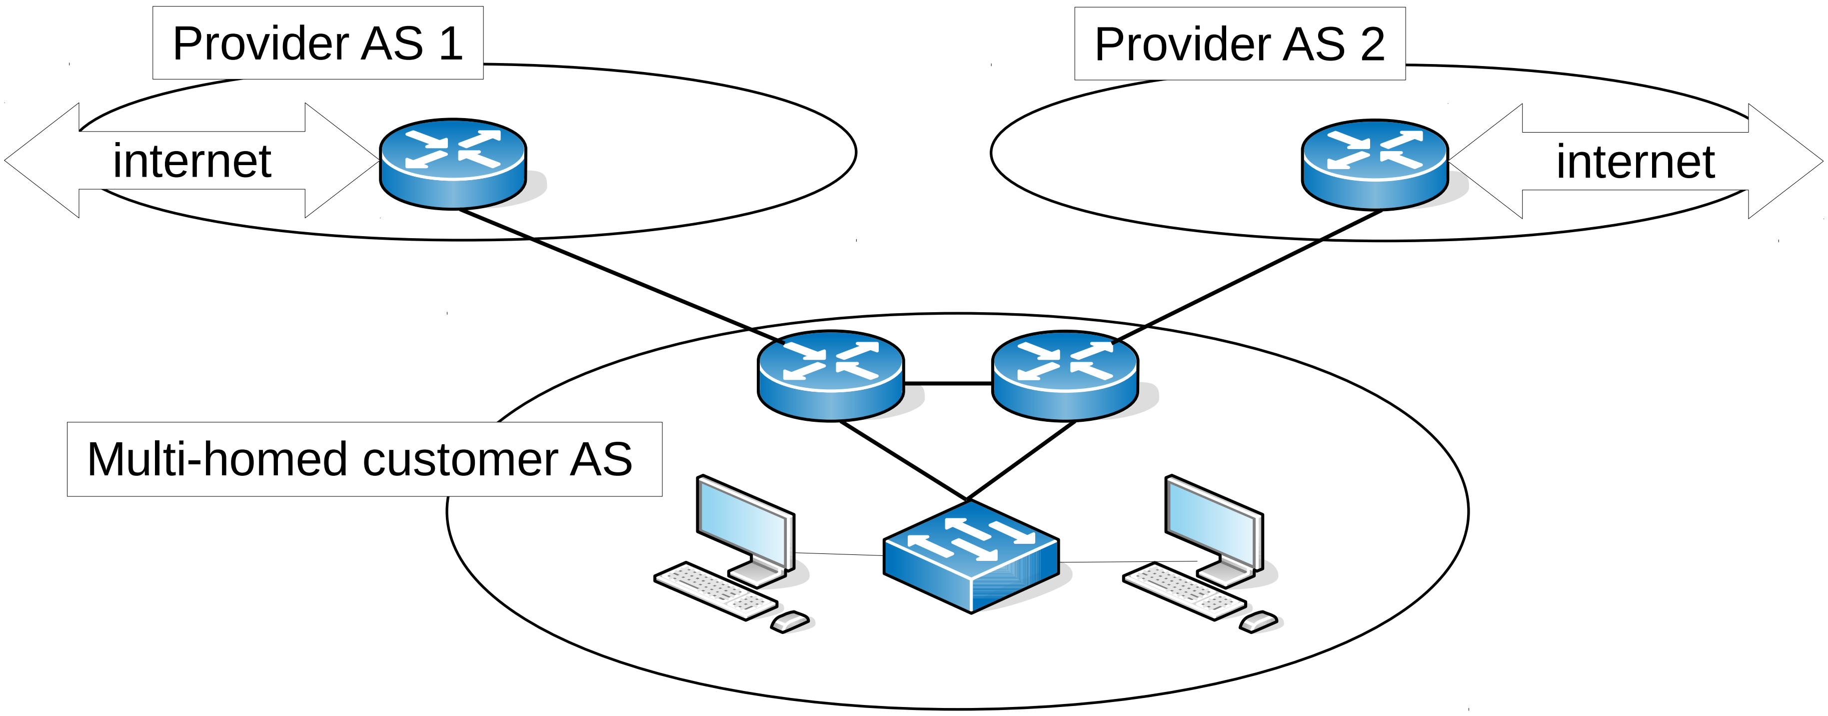

MH AS addressing: PA

- **Provider Aggregatable Addressing:** the MH AS receives a prefix from one of the providers
  - prefix aggregatable to one of the providers
  - but not aggregatable to the other!
- Customer's prefix must be announced separately

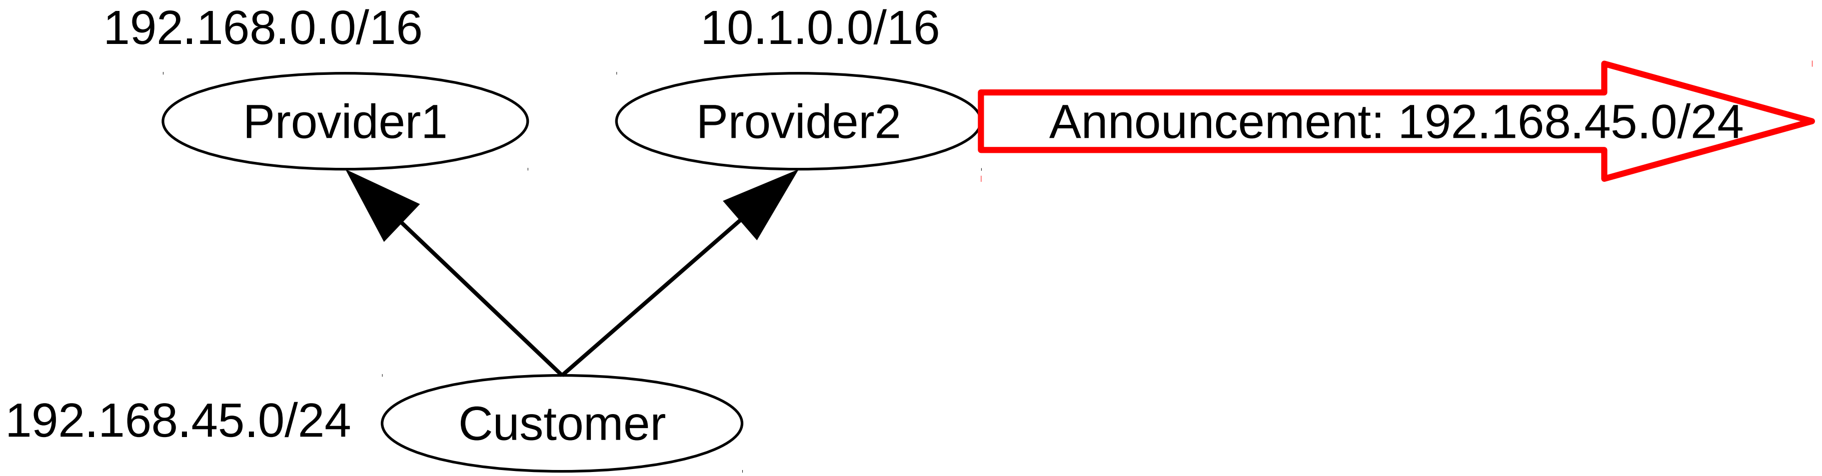

MH AS PA addressing: Example

- The prefix `192.168.45.0/24` leased by `Provider1` to `Customer1` cannot be aggregated into the prefix of `Provider2`
- `Provider2` must announce separately that it has a route to prefix `192.168.45.0/24`
- Plus one FIB entry throughout the Internet

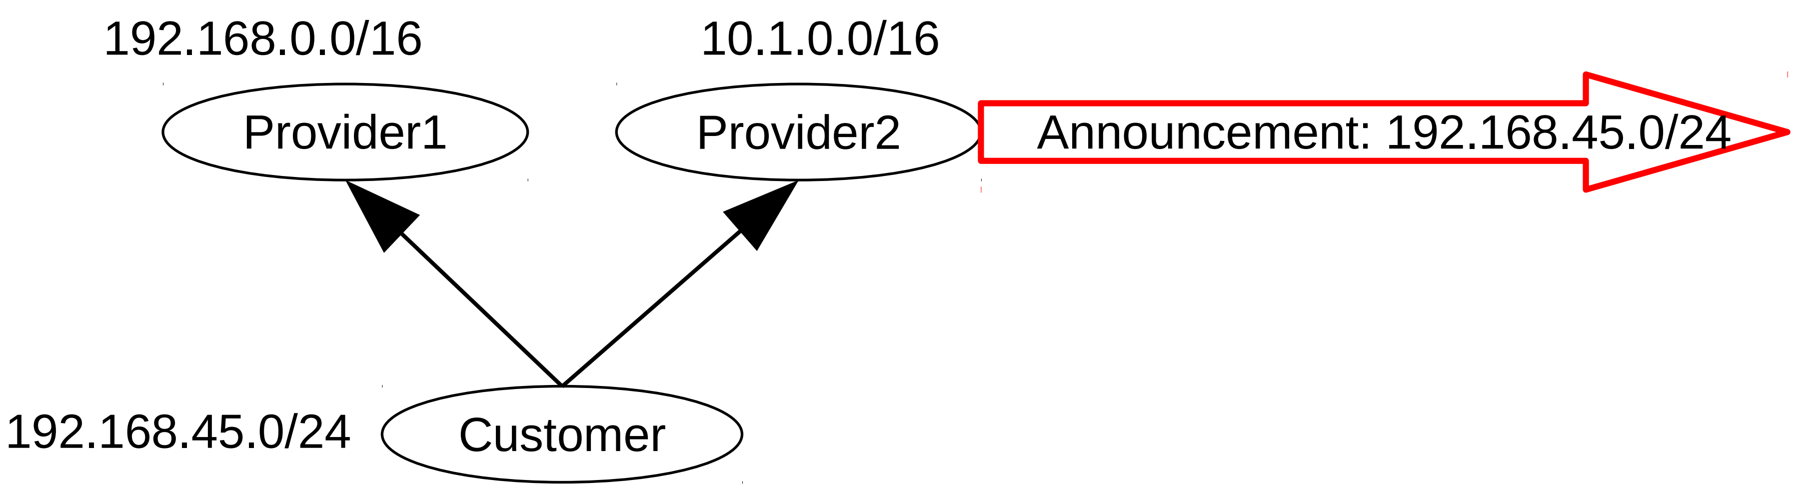

MH AS addressing: PI

- **Provider Independent Addressing:** the prefix of the customer is not aggregatable into any of the address ranges of the providers
  - providers announce the PI prefix separately
  - plus one global FIB entry

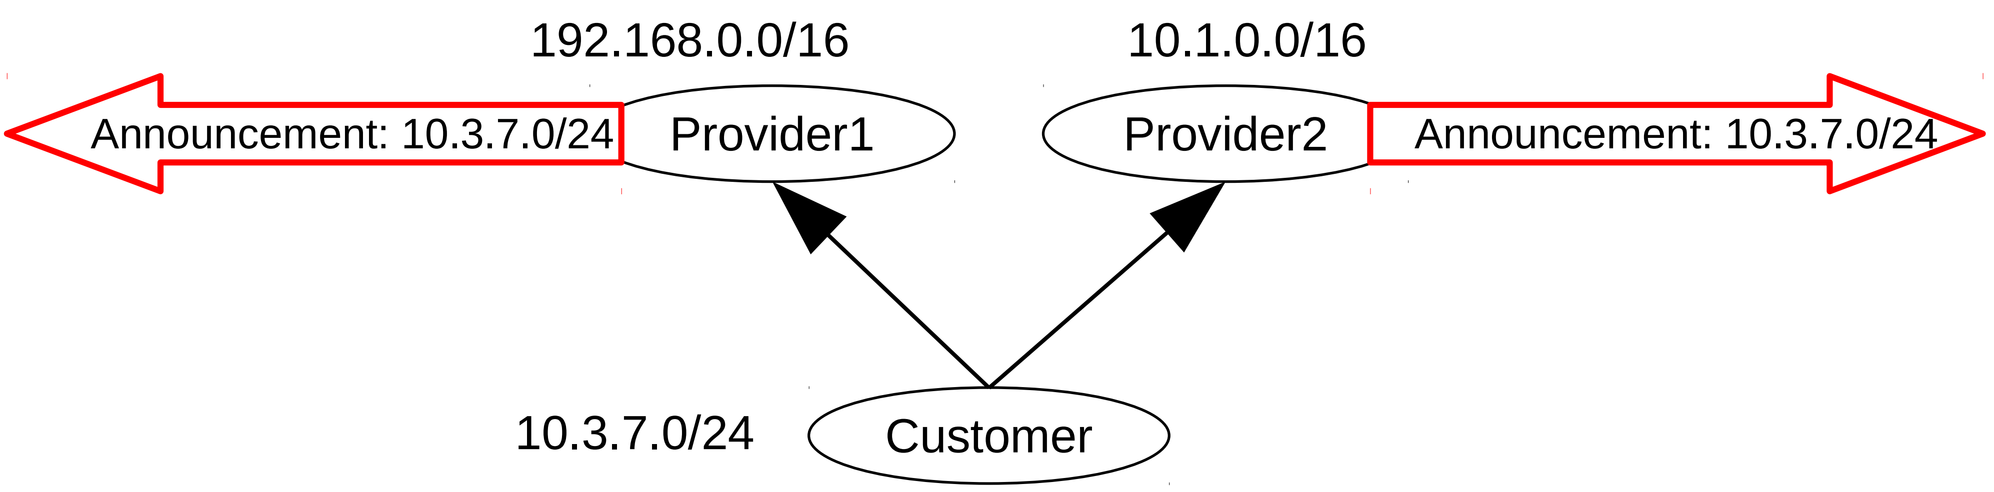

PA versus PI addressing

- If a customer uses a PA address prefix, it is tied to the provider: every interface/host in the customer's AS gets an IP address from the provider's prefix
  - difficult to change provider (**renumbering**)
  - all hosts/routers must be reconfigured
- changing providers is simpler with a PI prefix
  - no need for renumbering/reconfiguration
  - but PI addressing puts huge strain on the routing system (separate announcement)

MH transit: Example

- 3 ASes: `AS1` (prefix: `192.168.0.0/16`), `AS2` (`10.0.0.0/16`) and `AS3` (`10.1.0.0/16`)
- `AS1`: two border routers (`R1` and `R2`) and one LAN (prefix: `192.168.2.0/24`)

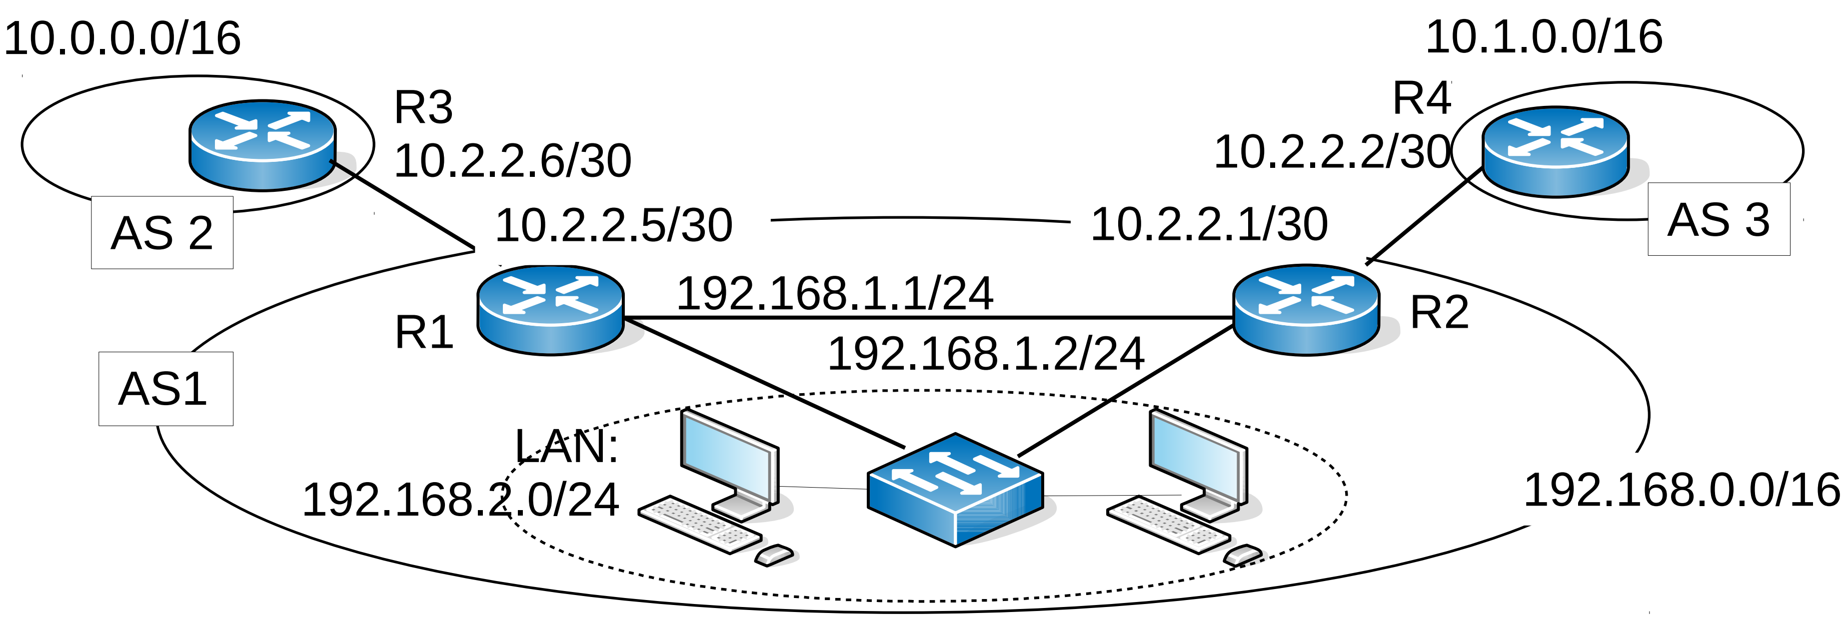

MH transit: Example

R1 FIB		
Prefix	**Next-hop**	**Note**
0.0.0.0/0	10.2.2.6	Default gateway to R3 (AS2 is primary provider)
10.0.0.0/16	10.2.2.6	Traffic to AS2 through R3 (optional, due to the default gw)
10.1.0.0/16	192.168.1.2	Traffic to AS3 through R2
192.168.2.0/24	local	LAN is local, no separate FIB entry for the two hosts

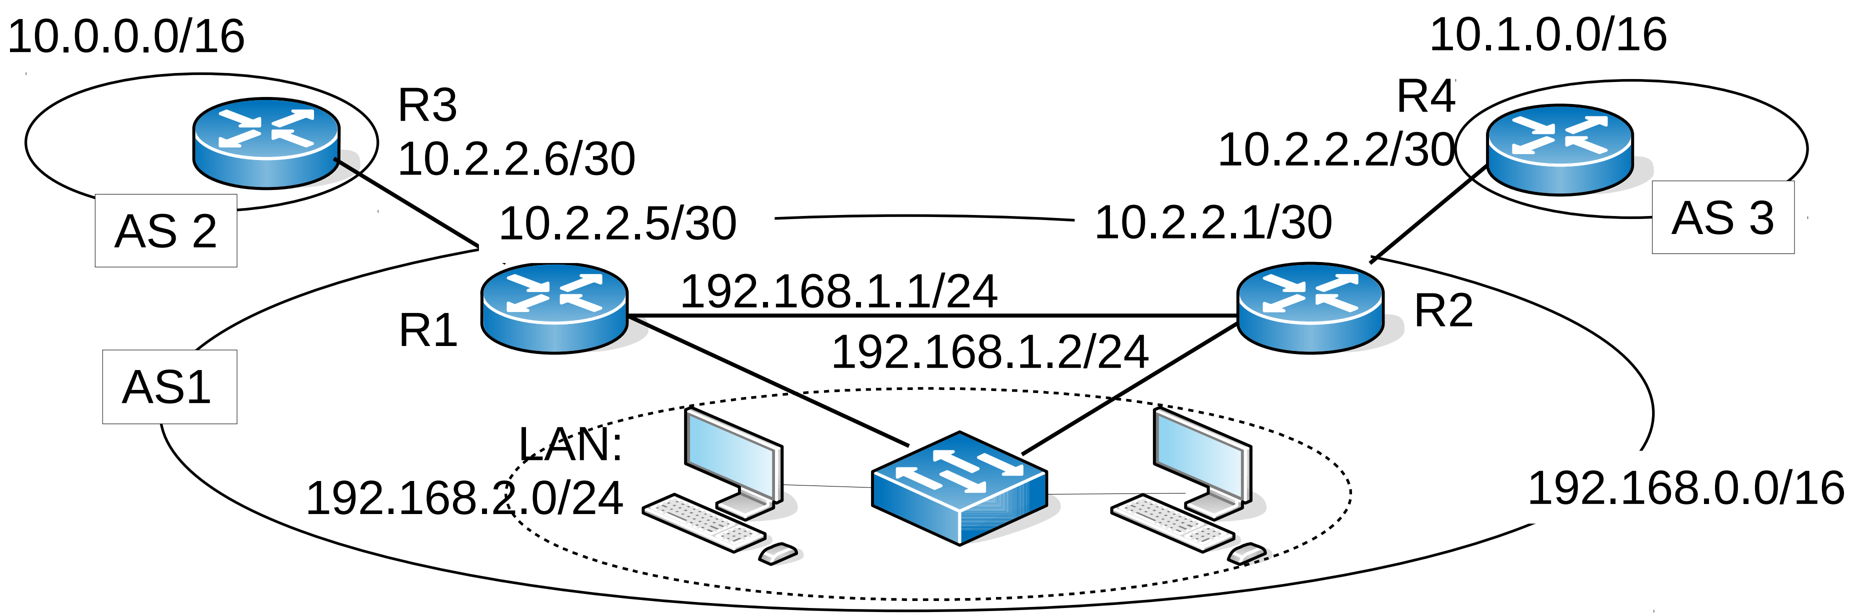

Multi-homing: Pros and cons

- Advantages of multi-homing
  - greater availability: switch to the secondary provider if connection to the primary is lost
  - cost optimization: choose less expensive provider as primary
  - load balancing between providers
- Disadvantages:
  - complex configuration: border routers need dynamic routing protocols to do the swtich-over when the primary provider dies
  - the prefix of the multi-homed AS is announced separately: increases FIB size throughout the Internet!