BGP

Attributes

Attribute Types

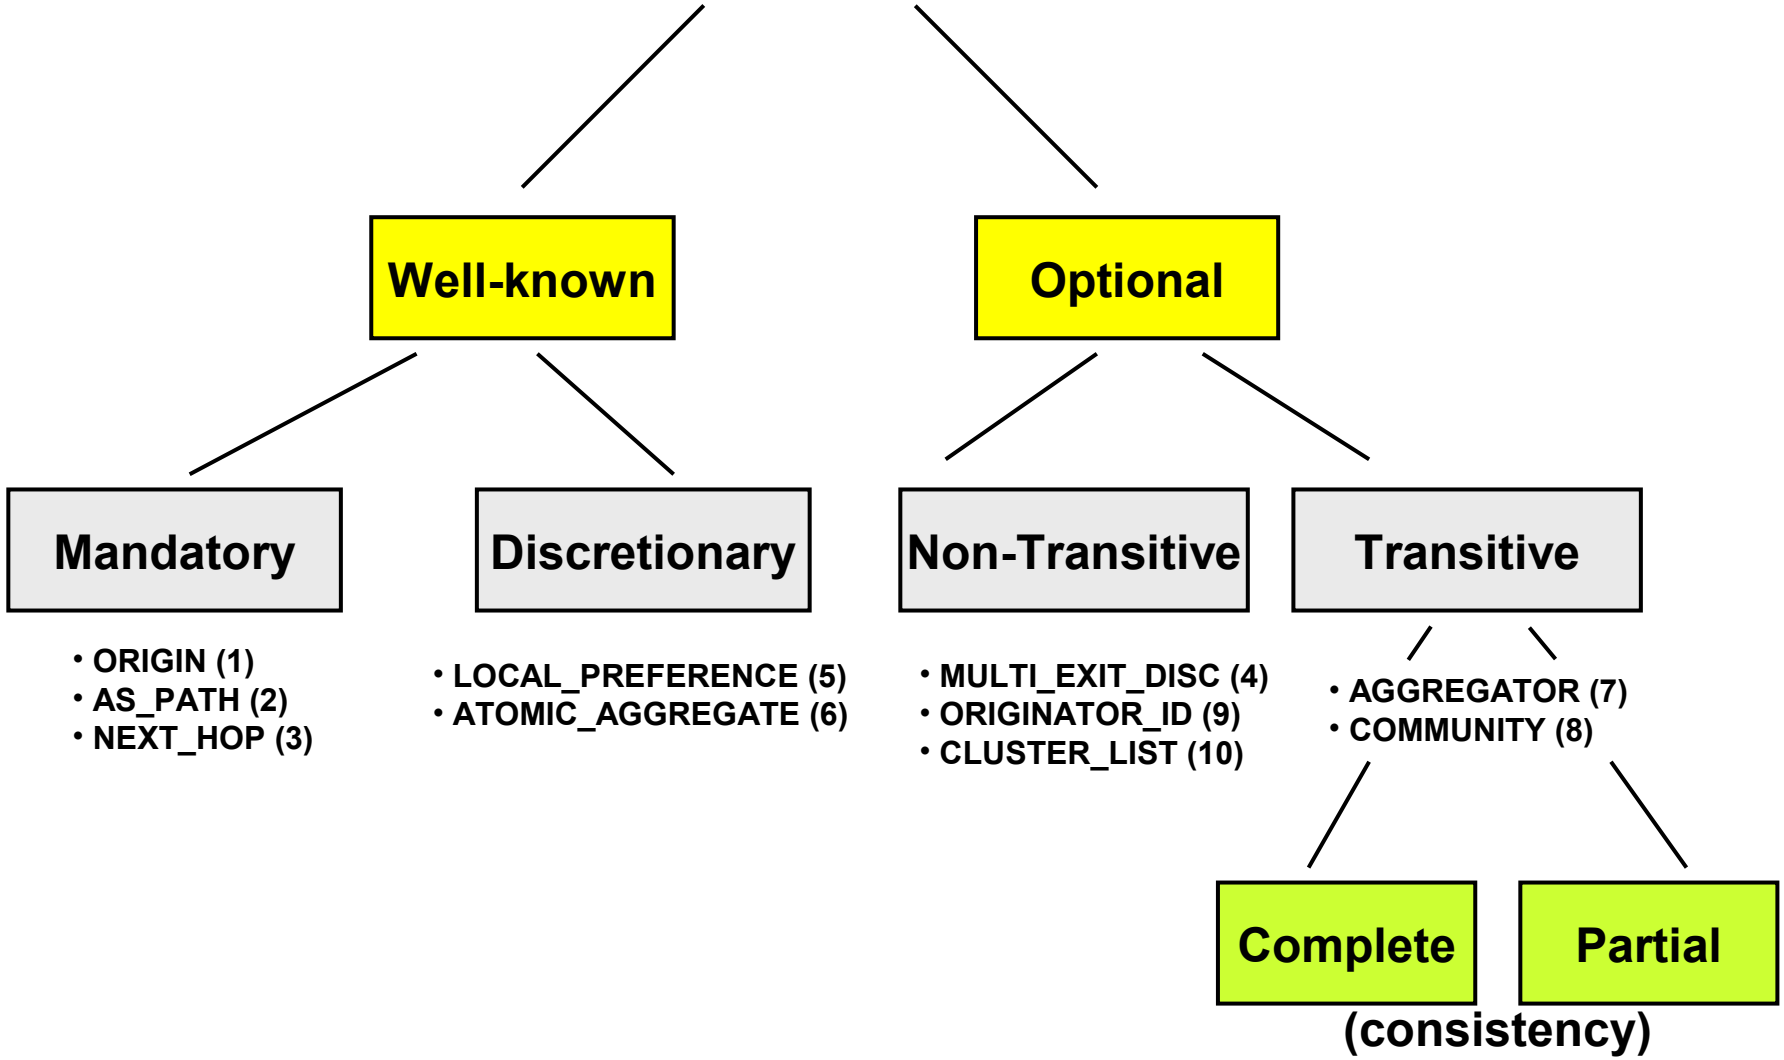

Path Attributes

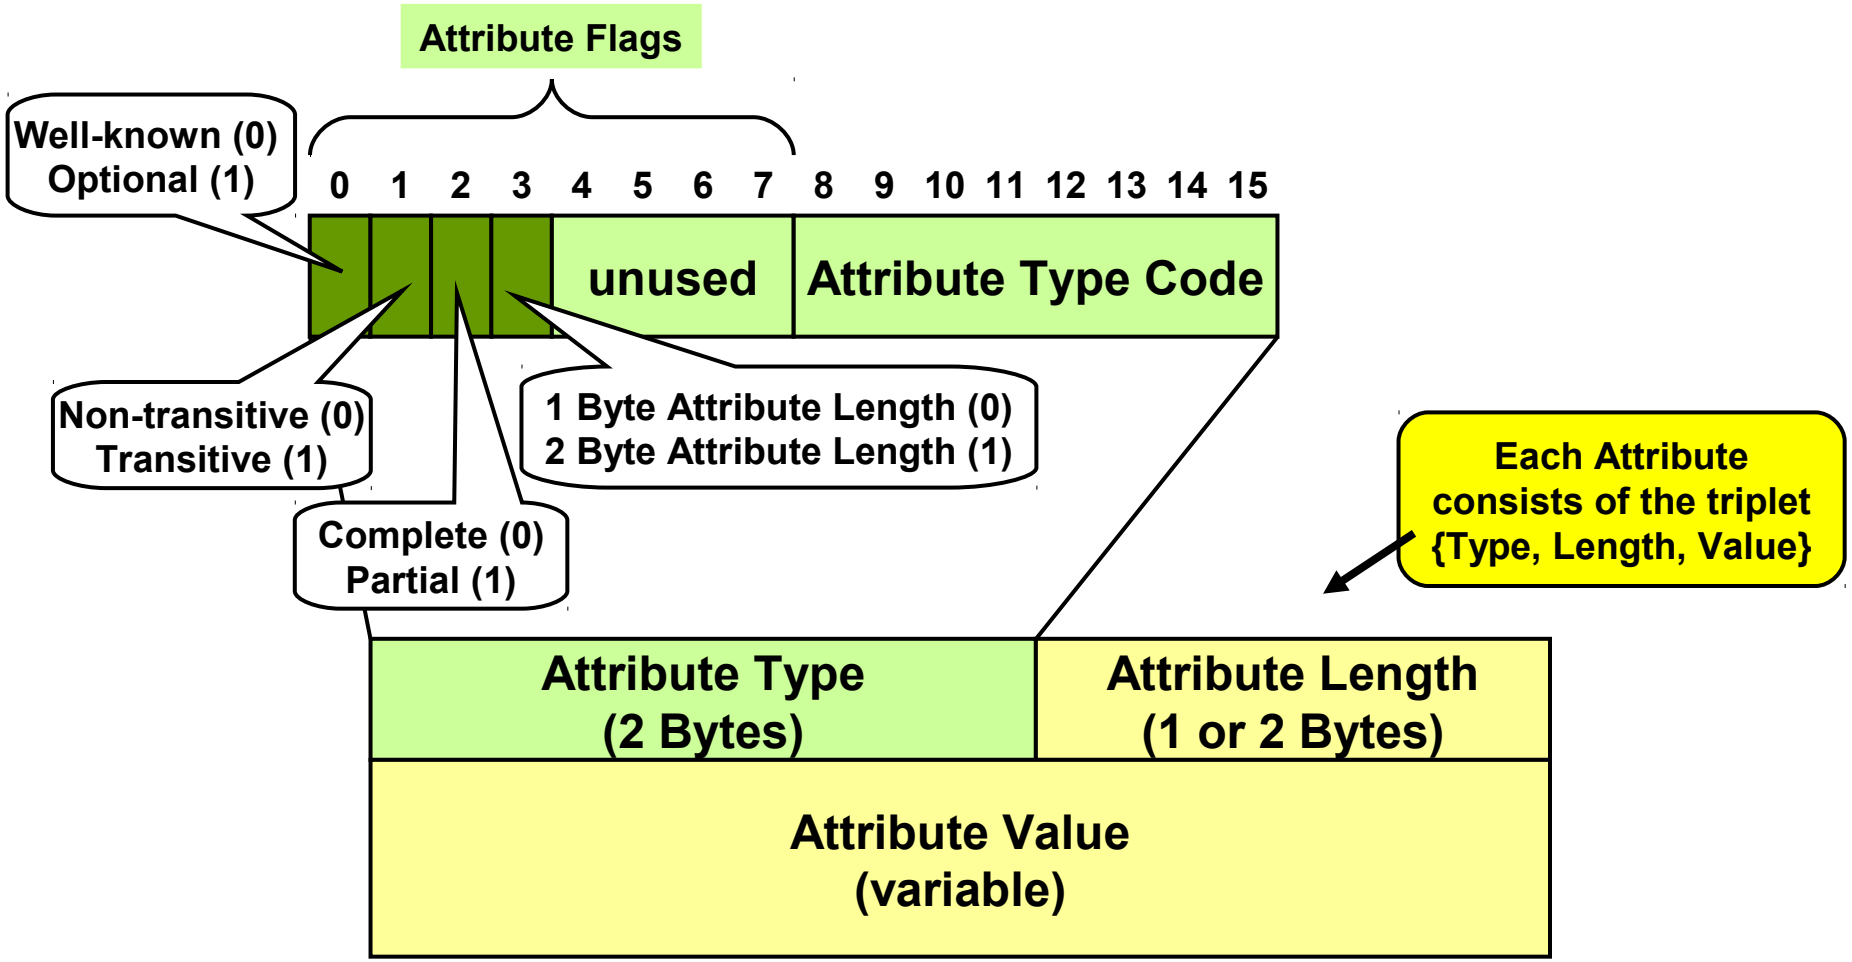

Well-known Mandatory

- AS_Path contains all ASs traversed for this route
- Next_Hop indicates the last EBGP router leading to this route
  - Not necessarily the physical next hop
- Origin indicates how this route was learned

Path Vector Protocol (1)

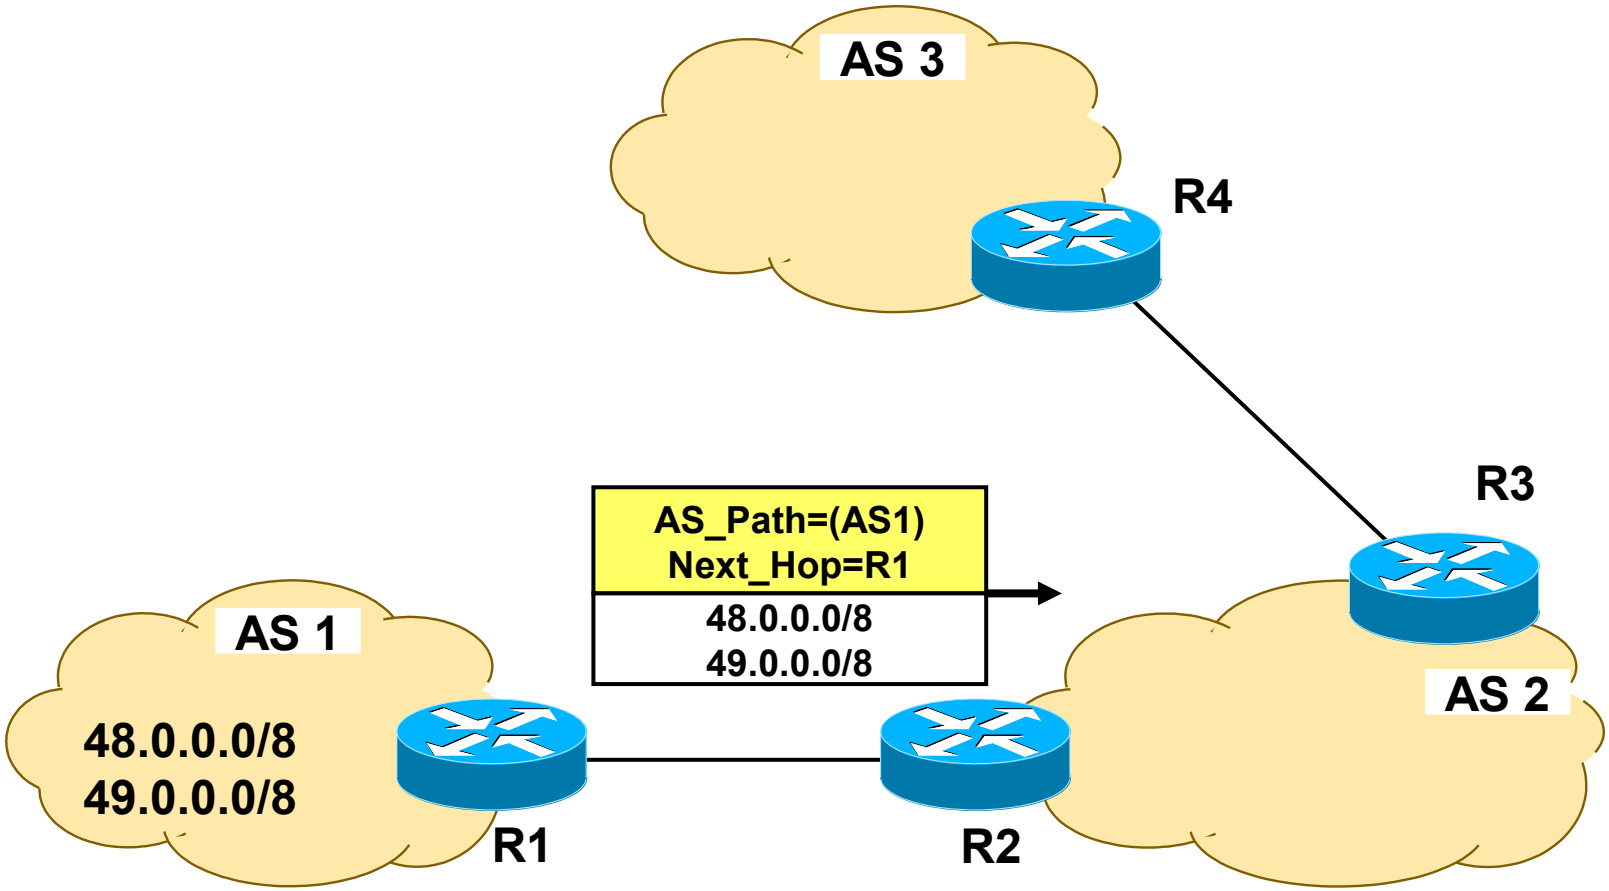

Path Vector Protocol (2)

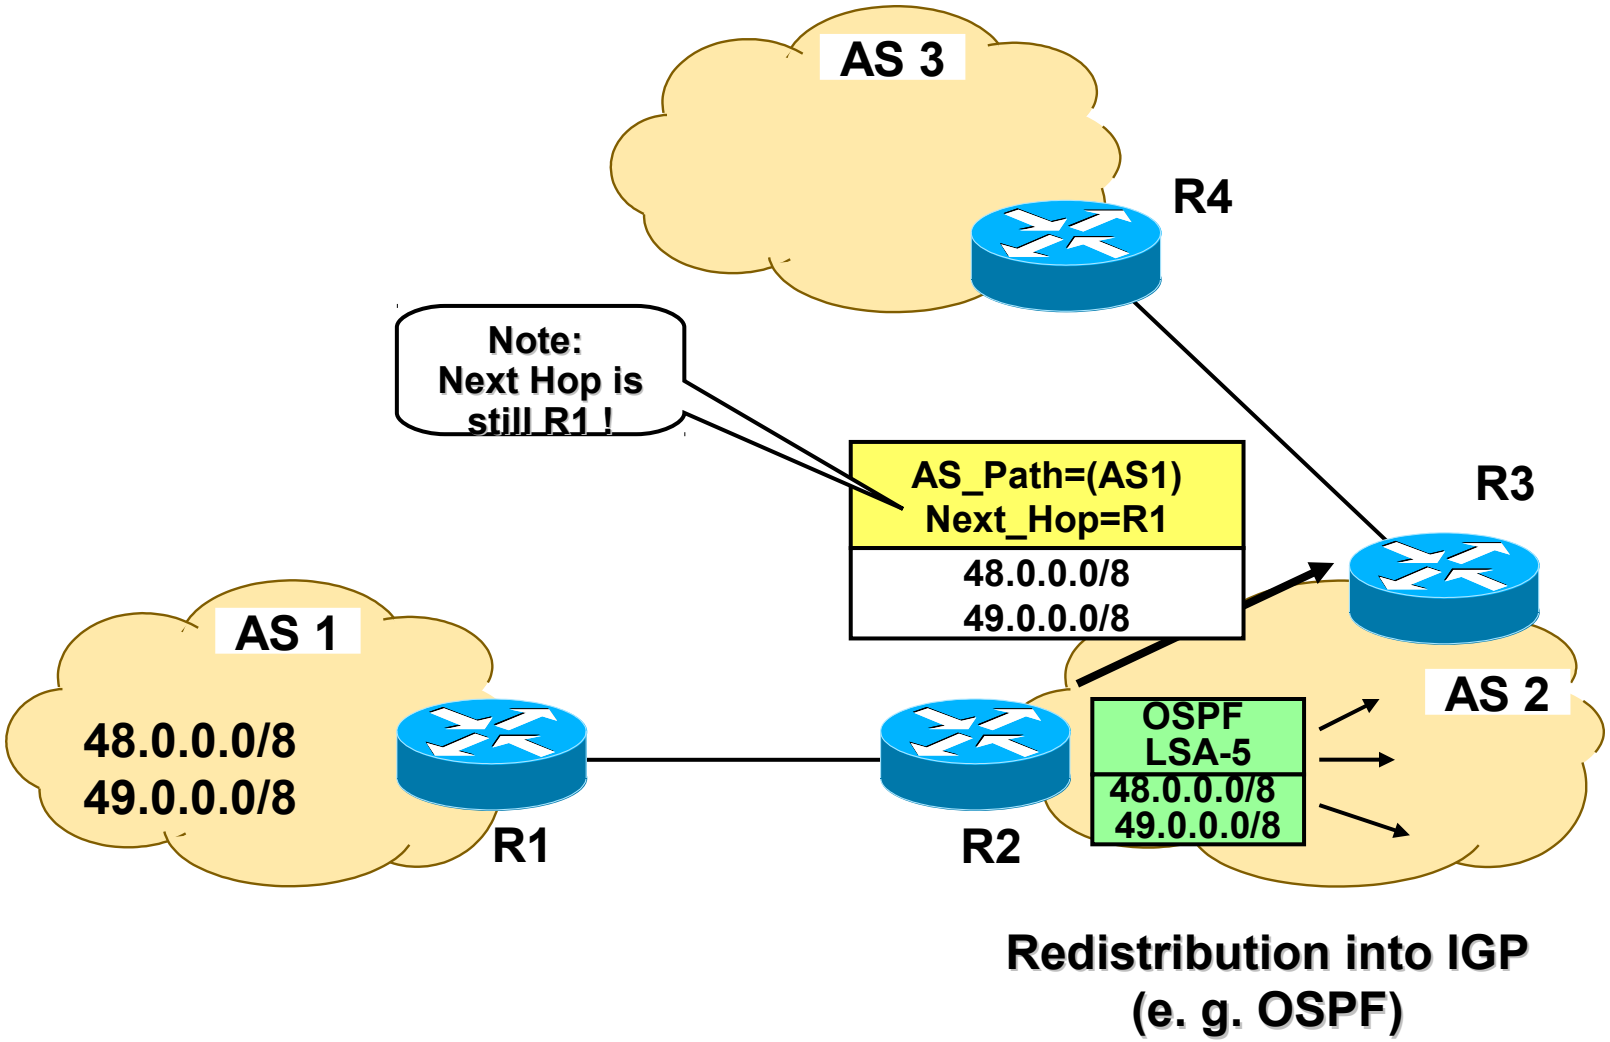

Path Vector Protocol (3)

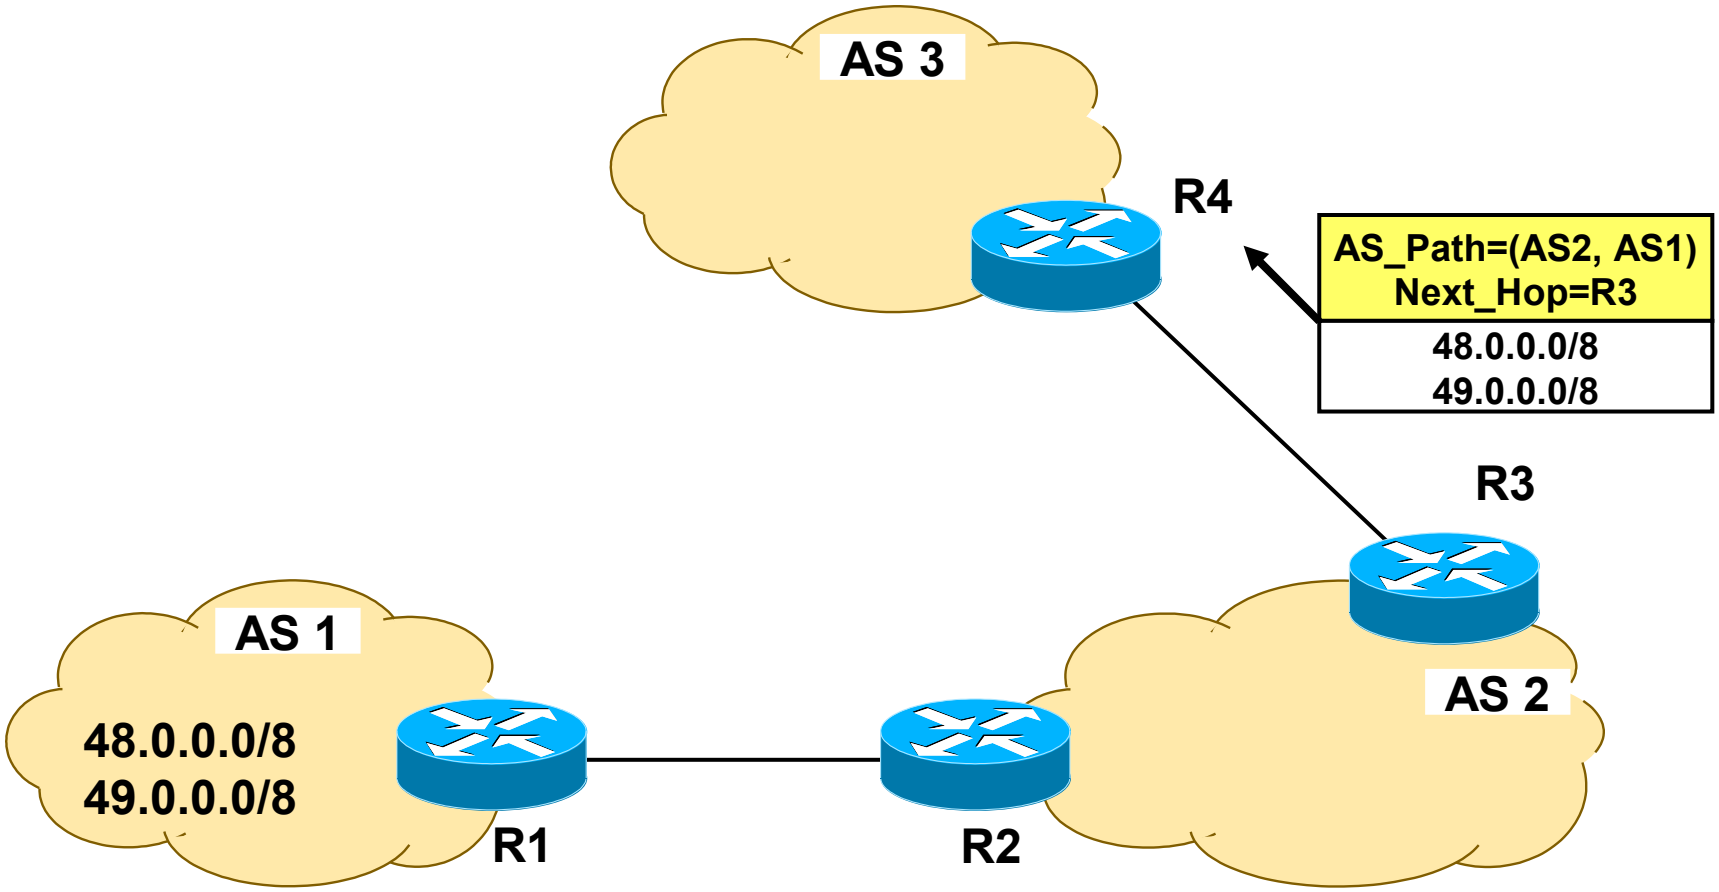

ORIGIN

1	Well-known
	Mandatory

- Value 0: IGP
  - Routes learned via network statement (NLRI is member of originating AS)
- Value 1: EGP
  - Learned via redistribution from EGP to BGP
- Value 2: INCOMPLETE
  - Learned via redistribution from IGP to BGP
  - Example: redistribute static (Cisco)

AS_PATH

2	Well-known
	Mandatory

- Composed of a sequence of AS path segments
- An AS path segment is represented by a triple
  - Path segment type (1 byte)
    - 1 = AS_Set (unordered set of ASs)
    - 2 = AS_Sequence (ordered set of ASs)
  - Path segment length (1 byte)
  - Path segment value (variable, 2 bytes per AS)

Who is NEXT_HOP?

3	Well-known
	Mandatory

- The **boundary router** that advertized the route in this AS is the next hop
  - Recursive routing table lookup might be necessary to determine the true physical next hop
- **Exception:**
  - On multi-access media (Ethernet, FDDI) always the physical next hop must be indicated

R1 and R2 have BGP session established, R3 speaks IGP only.

R2 advertises R3 as next hop to Net 30 because R3 is on the same physical media.

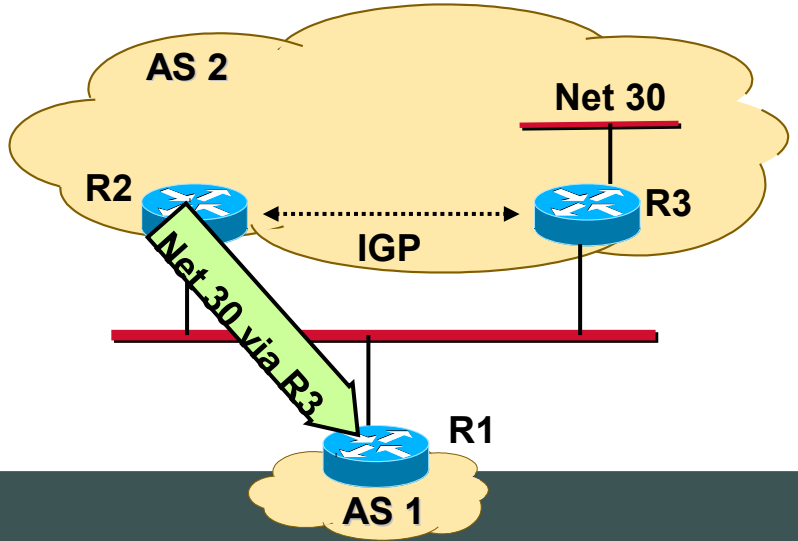

MULTI_EXIT_DISC

4	Optional
	Non-transitive

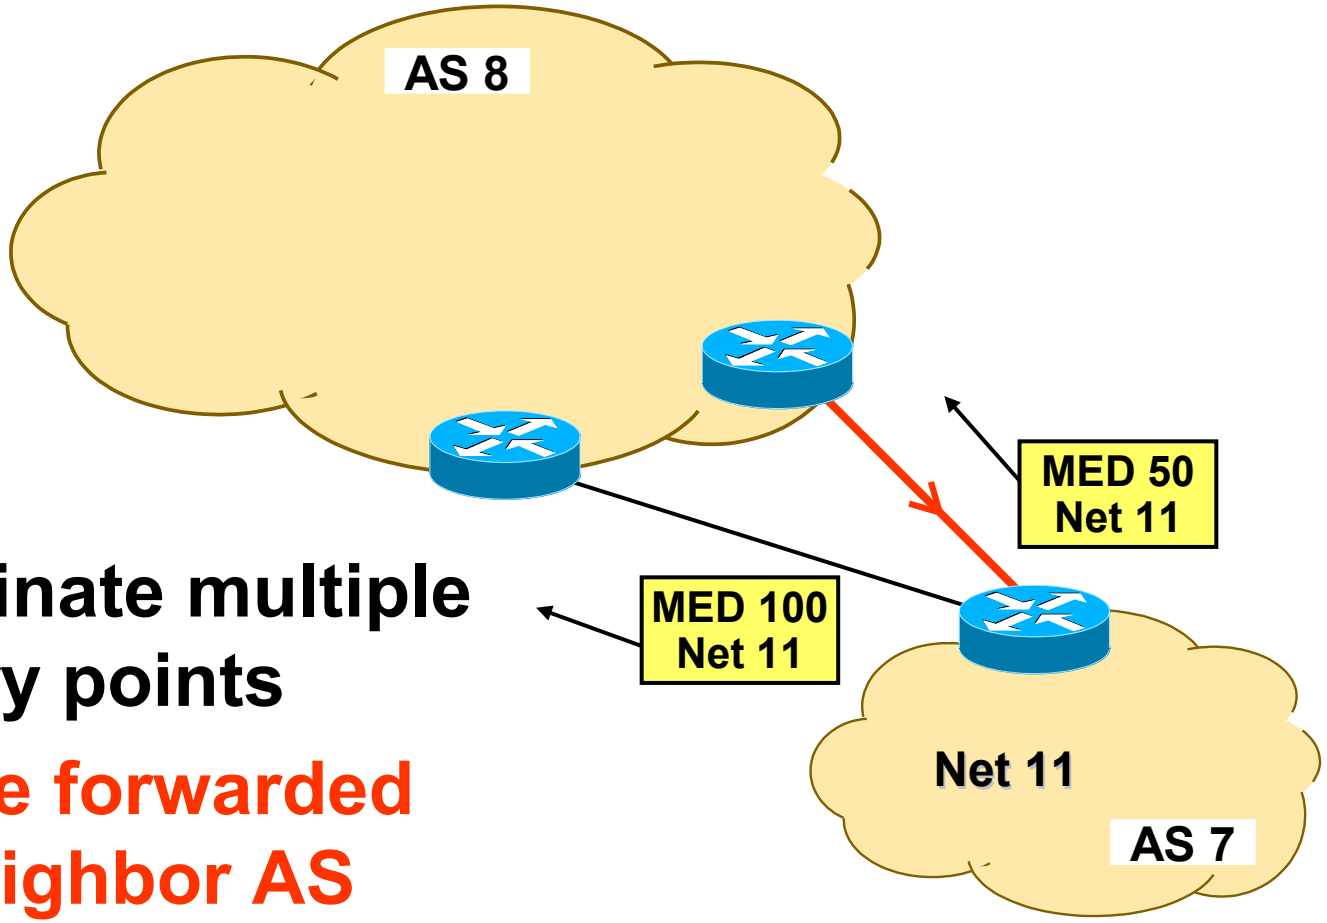

- To discriminate multiple exit or entry points
- Must not be forwarded to other neighbor AS

LOCAL_PREF

5	Well-known
	Discretionary

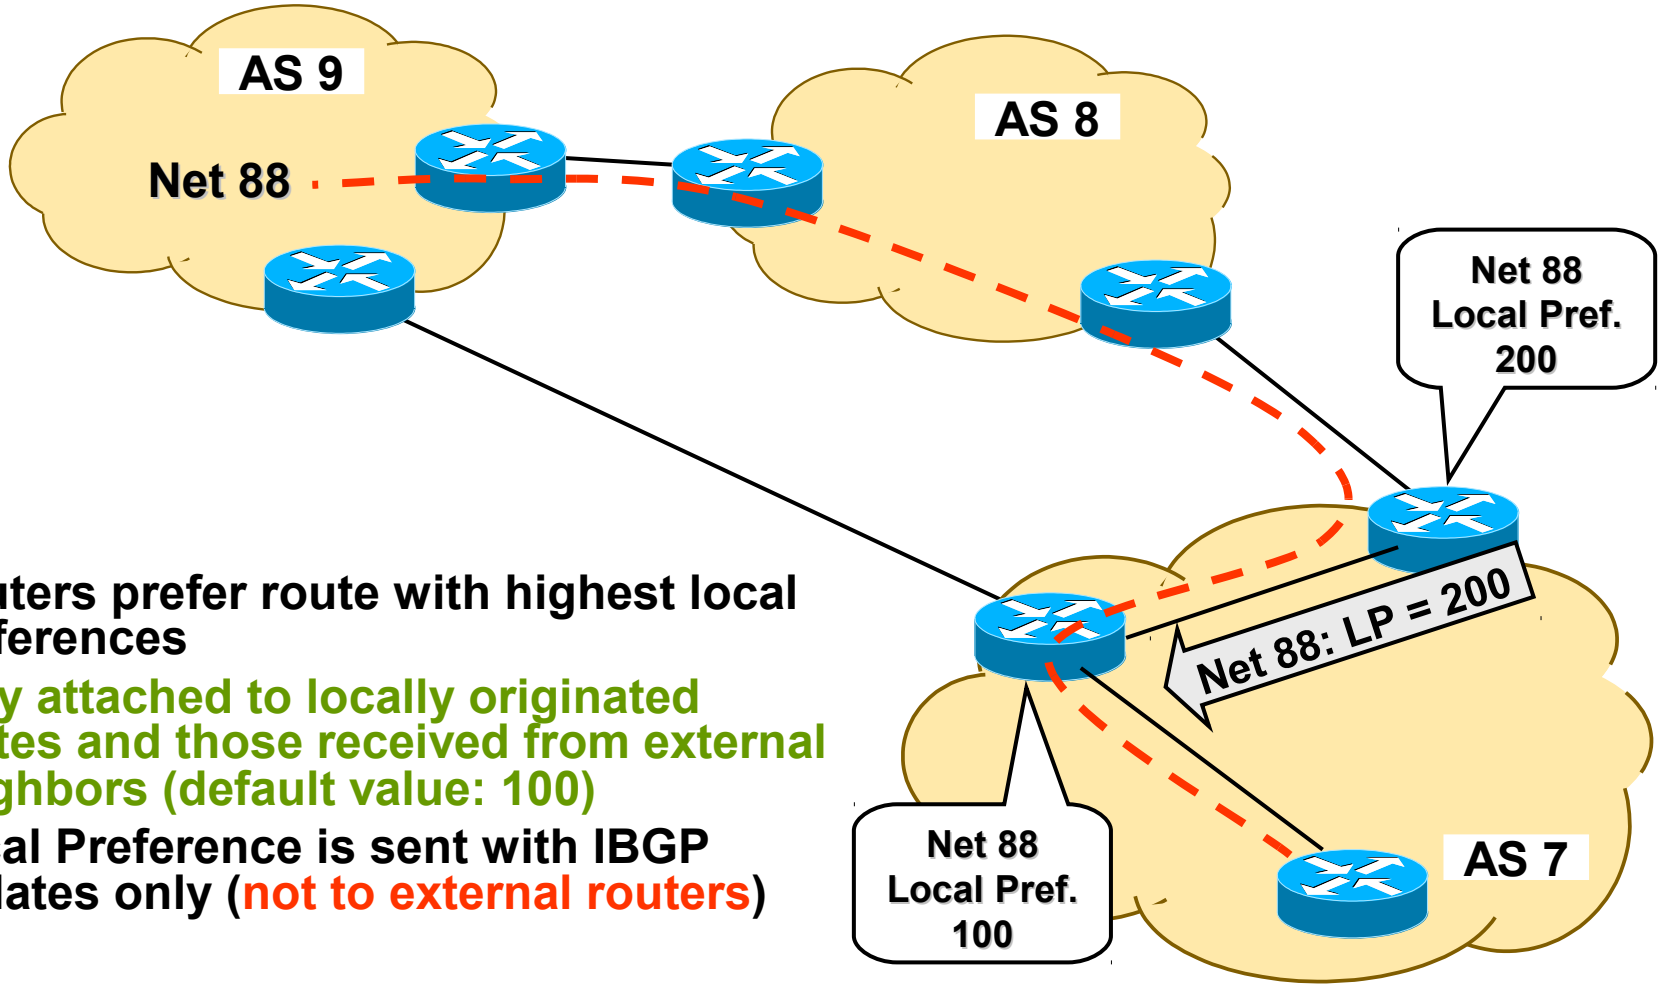

- Routers prefer route with highest local preferences
- Only attached to locally originated routes and those received from external neighbors (default value: 100)
- Local Preference is sent with IBGP updates only (not to external routers)

ATOMIC_AGGREGATE

6	Well-known
	Discretionary

- Optionally the Atomic_Aggregate attribute indicates that some BGP router made an AS aggregation
  - When selecting the less specific route on overlapping routes (rejecting the more specific route)
- Length 0

AGGREGATOR

7	Optional
	Transitive

- Contains the AS number and IP address of the BGP speaker that formed the aggregate route
- Useful for troubleshooting

COMMUNITY

8	Optional
	Transitive

- Group of destinations that share a common policy
  - Each destination could be member of multiple communities
  - Carried across ASs
- Community strings are simple policy labels
  - Any BGP router can tag routes in incoming and outgoing routing updates or when doing redistribution
  - Any BGP router can filter routes in incoming or outgoing updates or select preferred routes based on communities

Community Example (1)

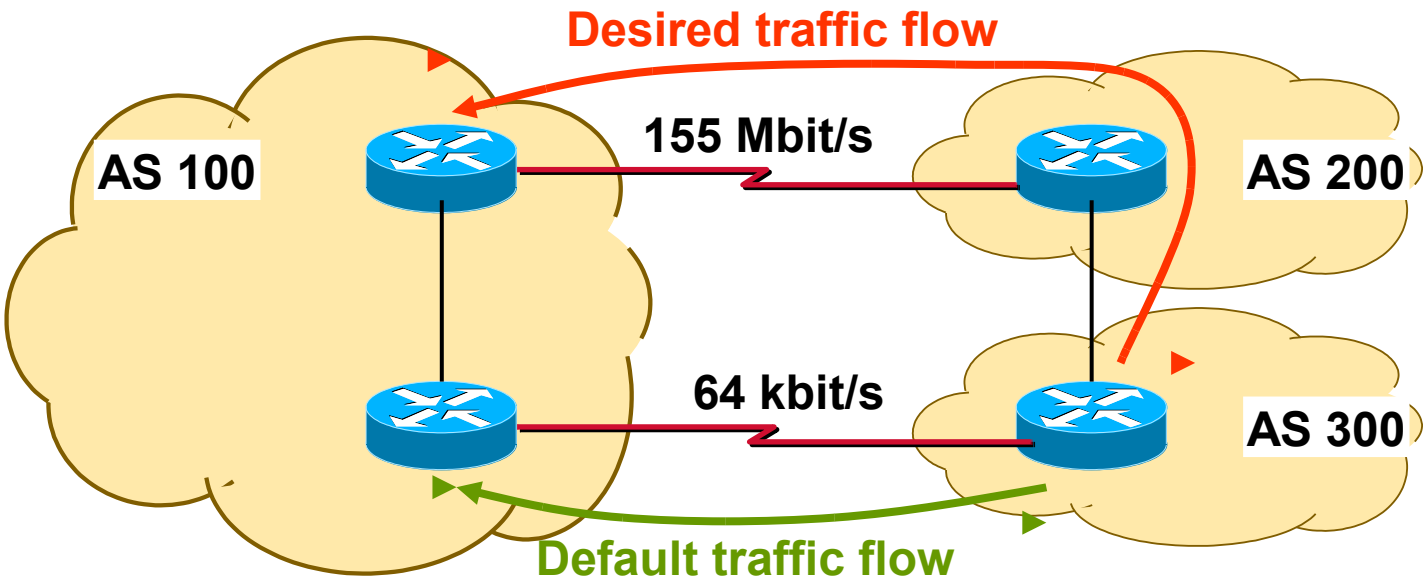

- Assume AS 100 wants AS 300 to use the 155 Mbit/s link to reach own networks
  - MED: not possible (non-transitive)
  - Local Preference: will admin of AS 300 set it?
- Best and easiest: Use community !

Community Example (2)

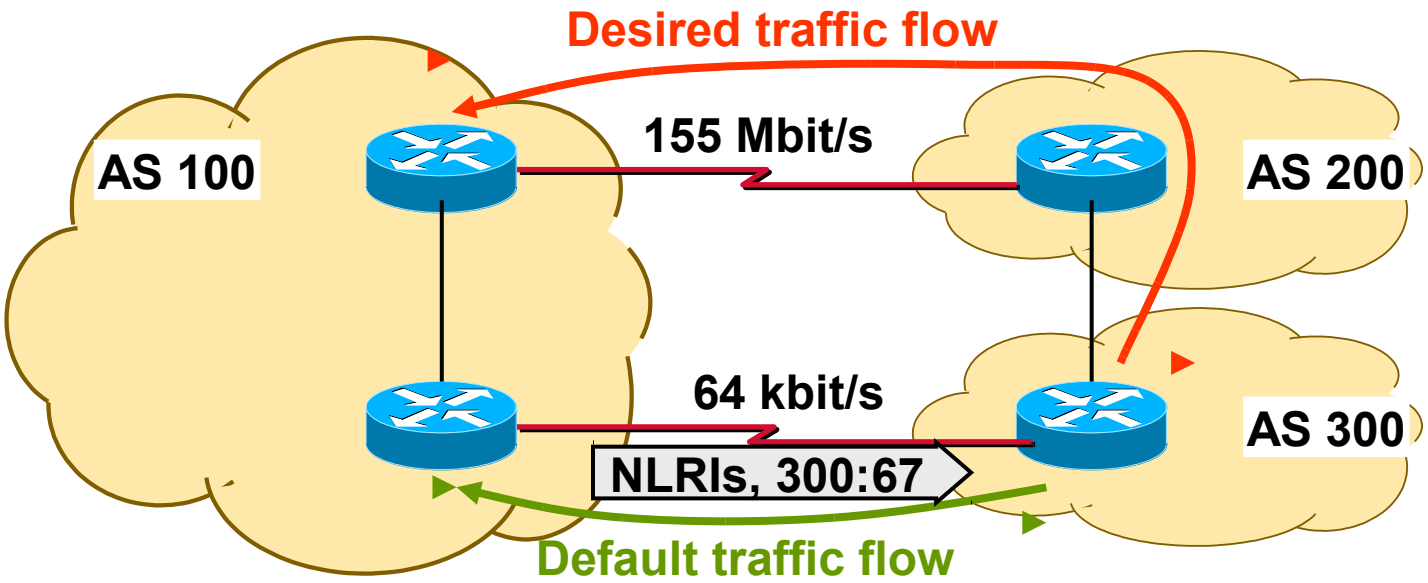

- Receiving a community string means "apply the predefined policy"
- In our example 300:67 means: "set local preference to 50"

Defining Communities

- More than one BGP community per route allowed
  - By default, communities are stripped in outgoing BGP updates
- Private range:
  0x00010000 - 0xFFFEFFFF
- Common practice
  - High order 16 bit: AS number
  - Low order 16 bit: Local significance

Well-known Communities

- Reserved ranges: 0x00000000 - 0x0000FFFF and 0xFFFF0000 - 0xFFFFFFFF
- 0xFFFFFF01 means: NO_EXPORT
  - Routes received carrying this value should not be advertised to EBGP peers, except ASs of a confederation
- 0xFFFFFF02 means: NO_ADVERTISE
  - Routes received carrying this value should not be advertised at all (both IBGP and EBGP peers)
- 0xFFFFFF03 means: NO_EXPORT_SUBCONFED
  - Routes received carrying this value should not be adverised to EBGP peers, including members of a confederation (Cisco: LOCAL_AS)

Administrative Weight (Cisco)

- No attribute – just a local parameter
- Applies only to routes within an individual router
- Number between 0 and 65535
  - The higher the weight the more preferable the route
- Initially invented to translate public routing policies (EGP)

Decision Hierarchy

1. Prefer highest weight (Cisco)
2. Prefer highest local preference
3. Prefer locally originated routes
4. Prefer shortest AS-Path
5. Prefer lowest origin code
6. Prefer lowest MED
7. Prefer EBGP path over IBGP path
8. Lowest IGP metric to next hop
9. Prefer oldest route for EBGP paths
10. Prefer path with lowest neighbor BGP router ID