


# Lab Course „RouterLab“

## **Border Gateway Protocol (BGP)**

# Internet: „Network of Networks“



# Internet: Structure and Routing

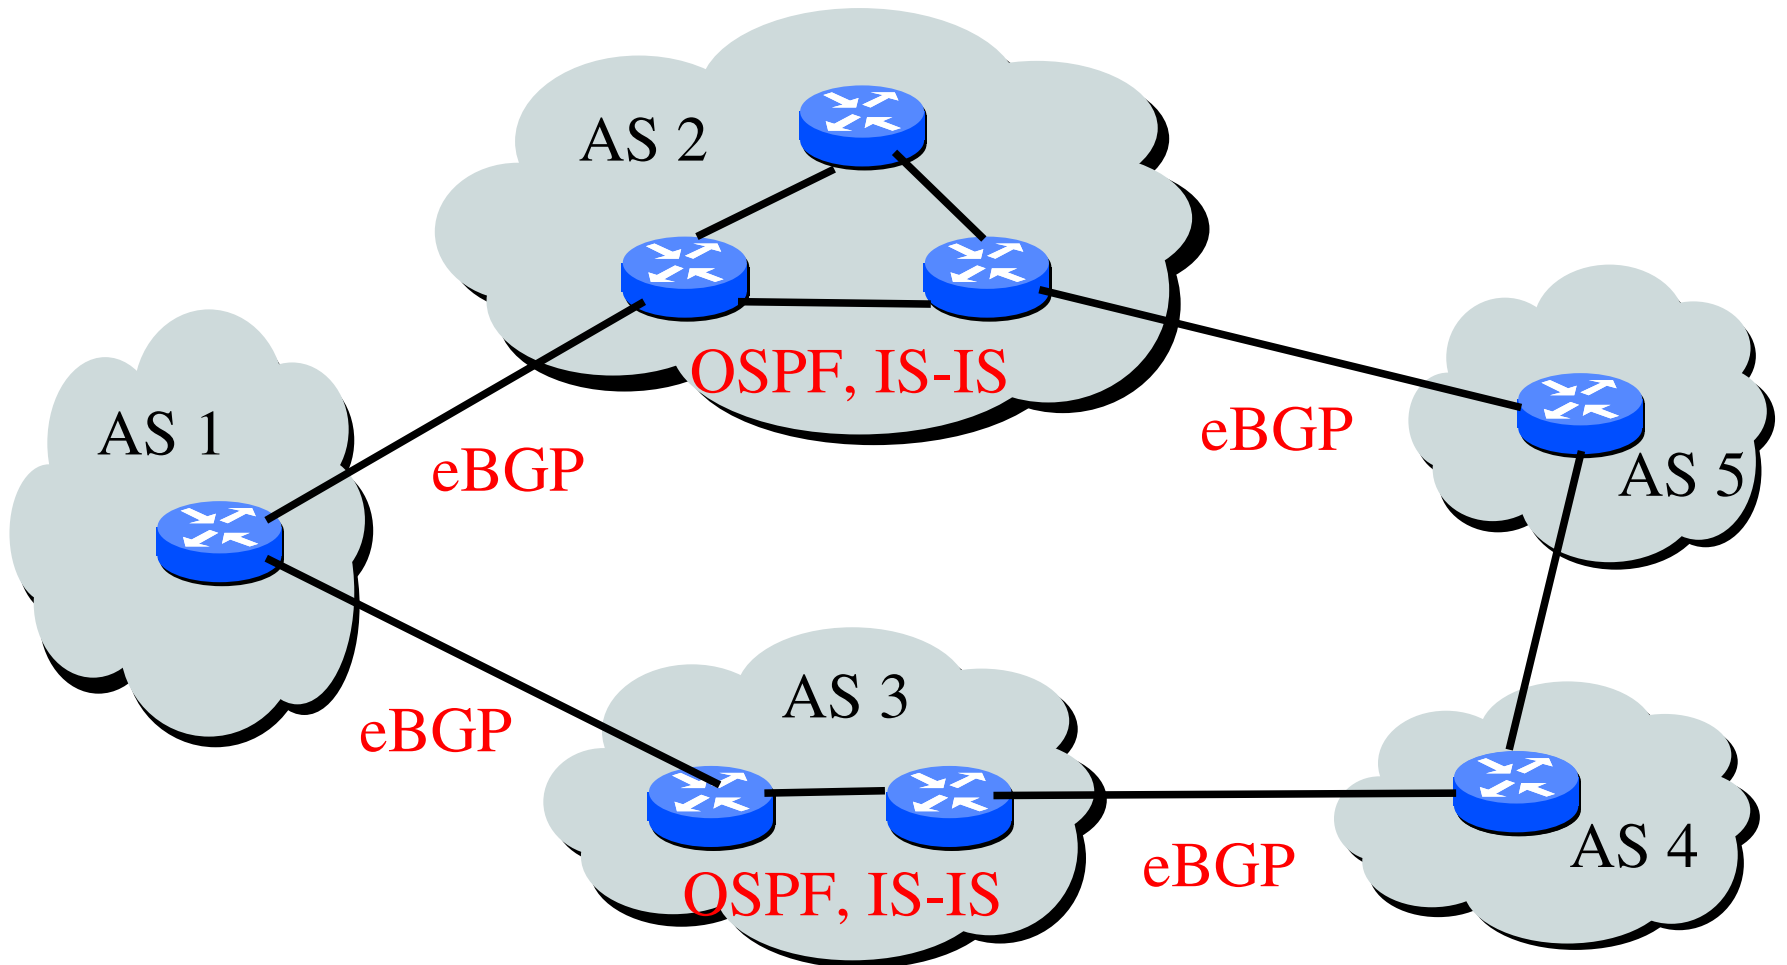
## □ **Structure:**

- > 20,000 autonomous systems (ASs)
- Examples for ASs?

## □ **Routing protocols:**

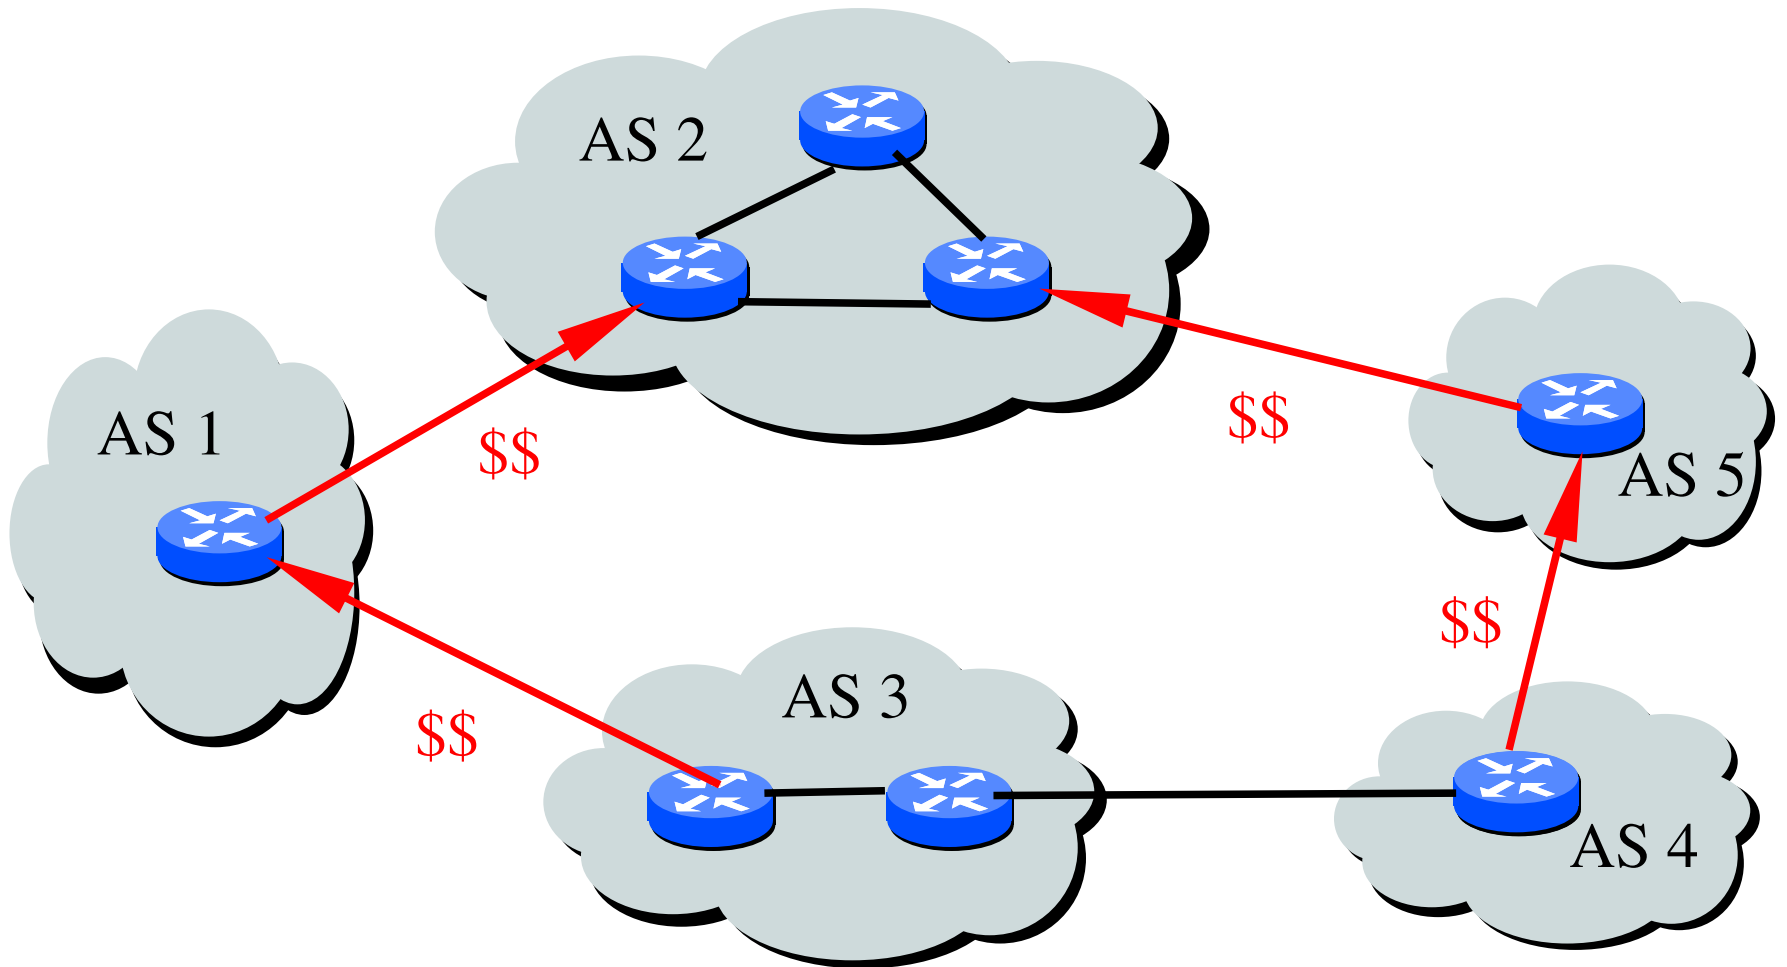
- Intra-domain: Inside Ass
  - Optimize for network performance
  - Examples: OSPF, ...
- Inter-domain: Between Ass
  - Policy-based (e.g., model customers, providers)
  - De-facto standard: *Border Gateway Protocol* (BGP)

# Inter- vs. Intra-domain routing



Note: iBGP (internal BGP) may be used between routers of same AS

# Policies: Customer-Provider



Note: Policies need to reflect business agreements ASs

# Internet inter-AS routing: BGP


- ❑ **BGP (Border Gateway Protocol):**  
*the de facto standard*
- ❑ BGP provides each AS a means to:
  1. Obtain subnet reachability information from neighboring ASs.
  2. Propagate the reachability information to all routers internal to the AS.
  3. Determine “good” routes to subnets based on reachability information and **policy**.
- ❑ Allows a subnet to advertise its existence to rest of the Internet: *“I am here”*

# BGP-4

- ❑ BGP = Border Gateway Protocol
- ❑ Is an exterior routing protocol (EGP)
- ❑ Is a policy-based routing protocol
- ❑ Is the de facto inter-domain routing protocol of today's global Internet
- ❑ Has a reputation for being complex

# BGP Basics



- ❑ Pairs of routers (BGP peers) exchange routing info over semi-permanent TCP connections: **BGP sessions**
- ❑ Note that BGP sessions do not correspond to physical links.
- ❑ When AS2 advertises a prefix to AS1, AS2 is *promising* it will forward any datagrams destined to that prefix towards the prefix.
  - AS2 can aggregate prefixes in its advertisement





# BGP Operations Simplified

**BGP Route =  
network prefix + attributes**



While connection is **ALIVE** exchange route **UPDATE** messages

# BGP messages

## Peers exchange BGP messages using TCP

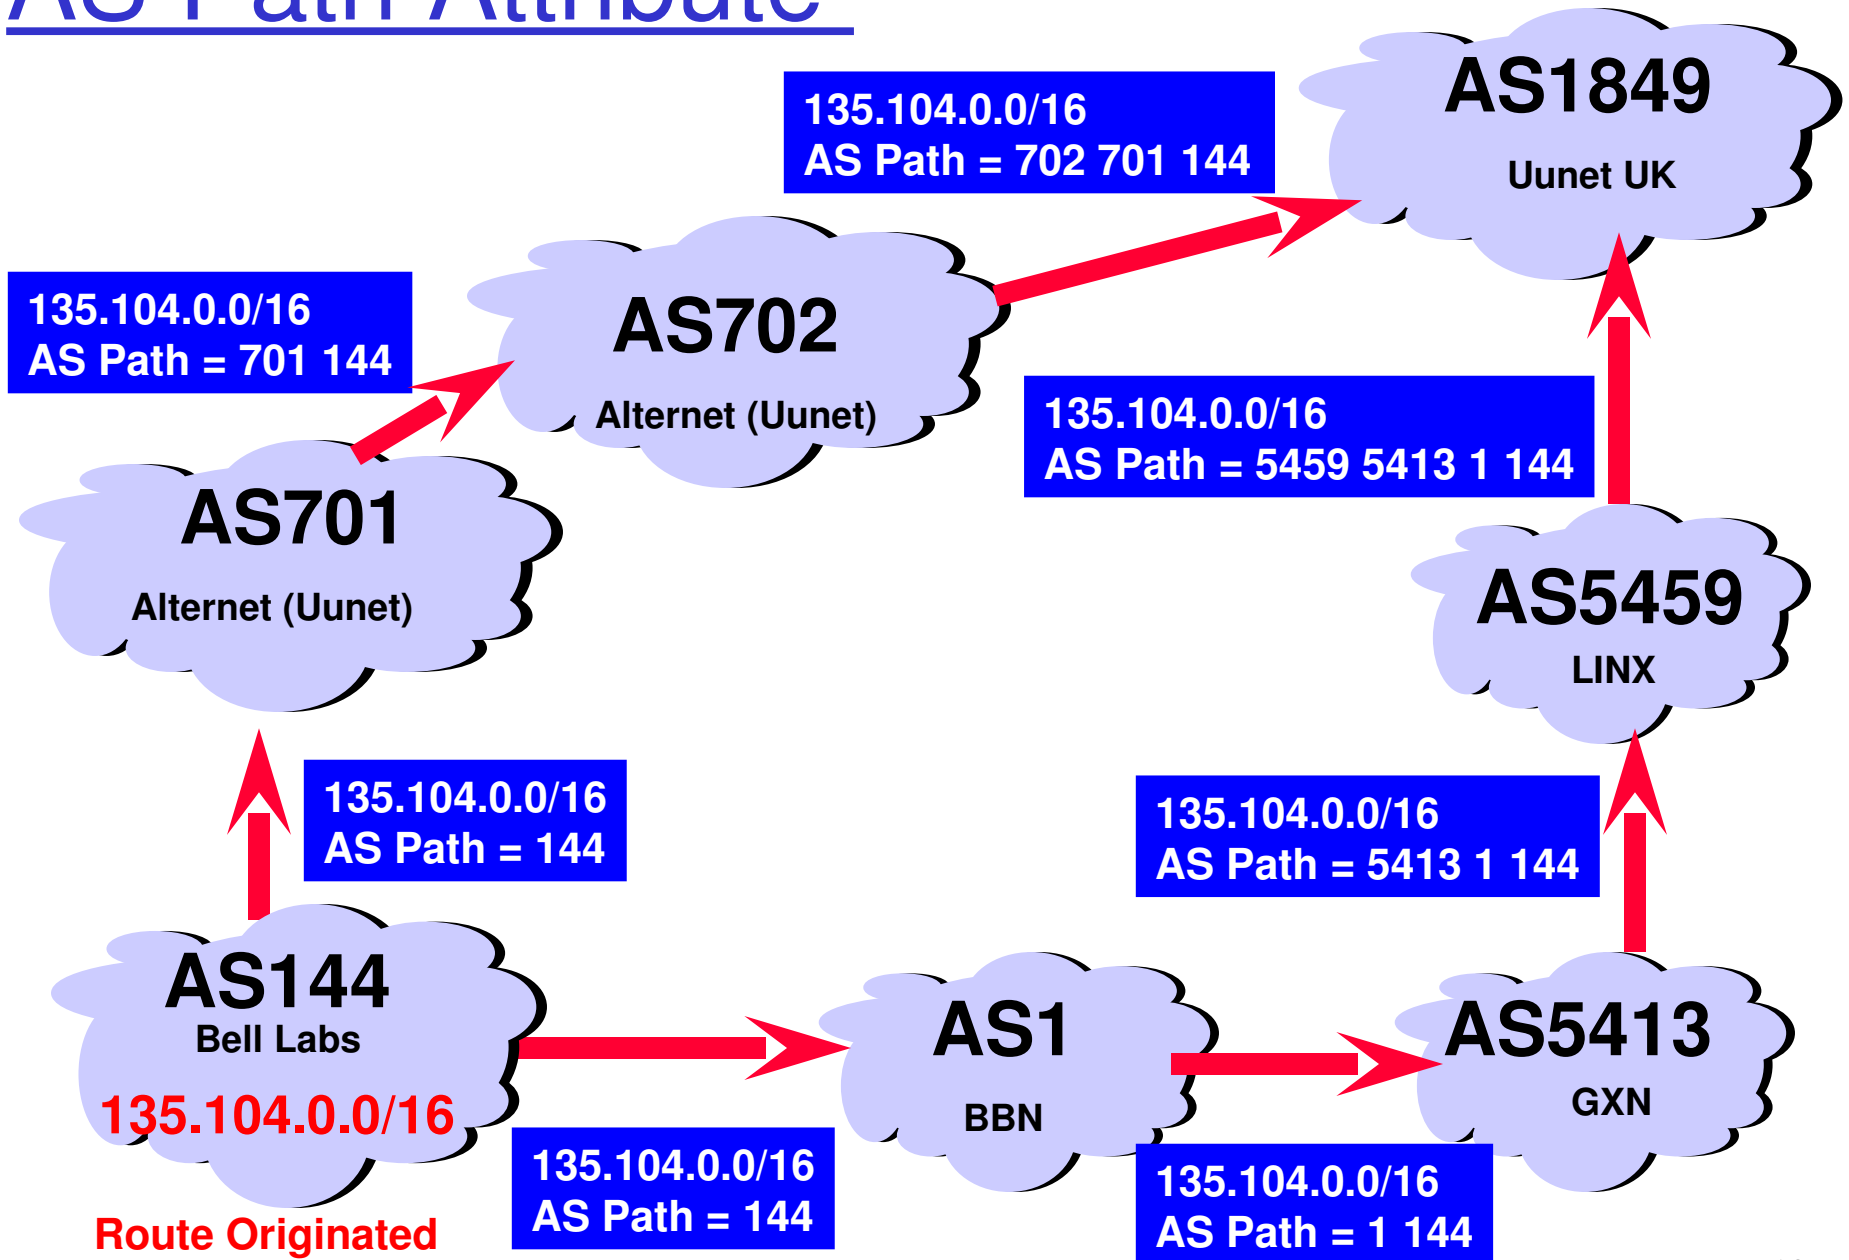
BGP messages:

- OPEN:
  - opens TCP conn. to peer
  - authenticates sender
- UPDATE:
  - advertises new path (or withdraws old)
- KEEPALIVE:
  - keeps conn alive in absence of UPDATES
  - serves as ACK to an OPEN request
- NOTIFICATION:
  - reports errors in previous msg;
  - closes a connection

# Path attributes & BGP routes

- ❑ When advertising a prefix, advertisement/update includes BGP attributes.
  - prefix + attributes = “route”
- ❑ Two important attributes:
  - **AS-PATH**: contains the ASs through which the advertisement for the prefix passed: AS 67 AS 17
    - used for loop detection / policies
  - **NEXT-HOP**: Indicates the specific internal-AS router to next-hop AS. (There may be multiple links from current AS to next-hop-AS.)


# AS Path Attribute



# BGP route selection

- ❑ Router may learn about more than one route to some prefix.
- ❑ Router must select route.
- ❑ Elimination rules:
  1. *Local preference* value attribute: policy decision
  2. *Shortest AS-PATH*
  3. ...
  - Pick route from router with lowest IP address (break tie)

# Local Preference Attribute



- AS 500 prefers path „500 200 300 400“ (higher local-preference wins)

# Routing policy

- Reflects goals of network provider
  - which routes to select for forwarding
    - prefer routes from customers
  - which routes to send to another AS
    - avoid being used as transit by your providers
  - how to manipulate the accepted routes
  - how to propagate routes through network
  - ...

# Internal BGP (iBGP)

- ❑ Same routing protocol as BGP, different application
- ❑ iBGP should be used when AS\_PATH information must remain intact between multiple eBGP peers
- ❑ Used inside Ass to keep AS path information



**Upstream  
Provider A**

**Upstream  
Provider B**

**AS100**

**AS200**

**eBGP**

**eBGP**


**iBGP**

**AS 1**

**AS 2**

**iBGP**

**eBGP**



# **BGP Configuration on Cisco and Juniper routers**

# Cisco Example - Local pref

- ❑ Configure a BGP session with neighbor-ip
- ❑ Set local-pref of 200 for all routes learned from this neighbor

```
router bgp <as number>
  neighbor <neighbor-ip> remote-as <remote AS number>
  neighbor <neighbor-ip> route-map <my_policy> in
!
route-map <my_policy> permit 10
  set local-preference 200
!
```

# Now the same on Juniper

```
routing-options {
    autonomous system <as number>
}
protocols {
    bgp {
        group <group-name> {
            peer-as <remote AS number>;
            type external; (internal for IBGP)
            neighbor <neighbor ip>;
            import <policy-name>;
        }
    }
}
policy-options {
    policy-statement <policy-name> {
        from {
            protocol bgp;
        }
        then {
            local-preference 200;
        }
    }
}
```