BGP

Introduction and Basic Procedures

Border Gateway Protocol (BGP)

- BGP-3
  - Was classful
  - Central AS needed (didn't scale well)
  - Not further discussed here!
  - RFC 1267
- BGP-4
  - Classless
  - Meshed AS topologies possible
  - Used today – discussed in the following sections!!!
  - RFC 1771

BGP-4 at a Glance

- Carried within TCP
  - Manually configured neighbor-routers
  - Therefore reliable transport (port 179)
- Neighbor routers establish link-state
  - Hello protocol (60 sec interval)
- Incremental Updates upon topology changes
  - New routes are updated
  - Lost routes are withdrawn
- Each route is assigned a policy and an AS-Path leading to that network
  - Using attributes

Path Vector Protocol

- Metric: Number of AS-Hops
- All traversed ASs are carried in the AS-Path attribute
  - BGP is a "Path Vector protocol"
  - Better than Distance Vector because of inherent topology information
  - No loops or count to infinity possible

BGP Database

- BGP routers also maintain a BGP Database
  - Roadmap information through path vectors
  - Attributes
- Routing Table calculated from BGP Database
- CPU/Memory resources needed

Some Interesting Numbers

- Today's Internet BGP Backbone Routers are burdened
  - About 100,000 routes (!)
  - About 10,000 Autonomous Systems
- Although excessive CIDR, NAT, and Default Routes
- Collapse expected
  - Looking for new solutions

Basic Idea of BGP is Easy !

1) BGP notifies other Autonomous Systems about reachabilities of networks

2) Each single route has attributes associated to it

3) Routers can apply policies for each route based on these attributes (e.g. filtering routes)

BGP Limitations

- Destination based routing
  - No policies for source address
- Hop-by-hop routing
  - Leads to hop-by-hop policies
  - Connectionless nature of IP
  - Mitigated through
    - Community attribute
    - Peer groups

Neighborship Establishment

- Open Message
  - BGP Version (4)
  - AS number
  - BGP Router-ID (IP address)
  - Hold Time
- Problems are indicated with Notification message

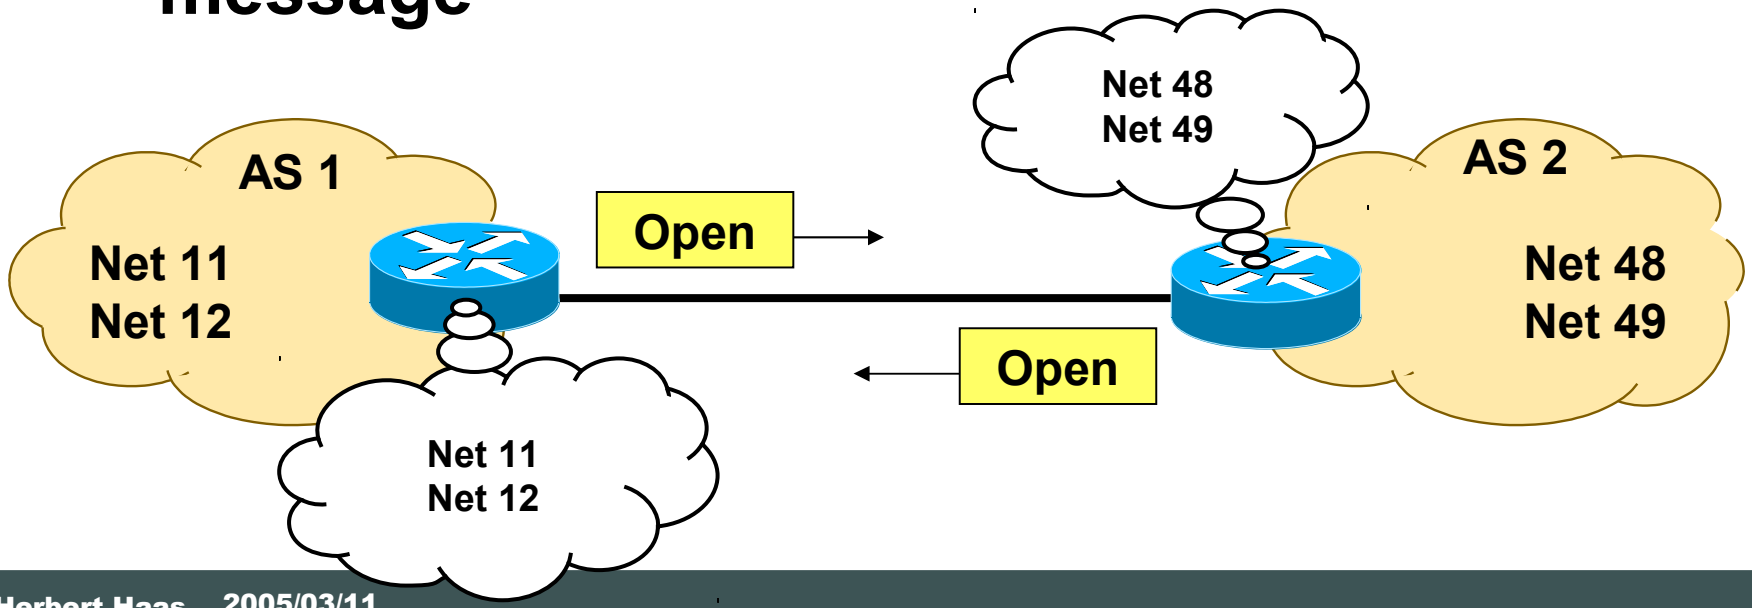

NLRI Update

- After open message, all known routes are exchanged using update messages
- Contains network layer reachability information (NLRI)
  - List of prefix and length

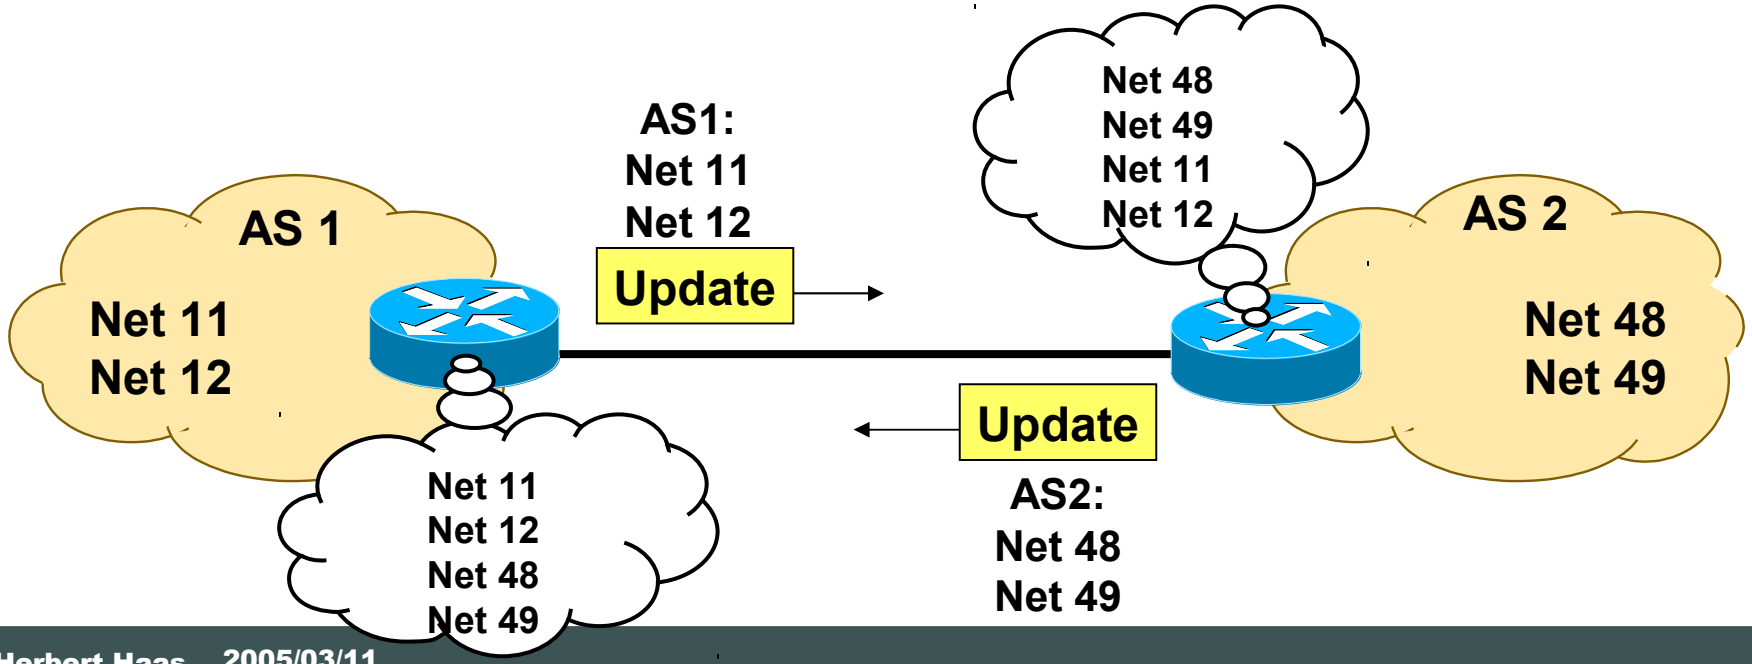

Steady State

- After Open/Update procedure, BGP is nearly quiet – *No periodic updates* !
- Only keepalive messages are sent
  - 19 Bytes
  - Per default every 60s

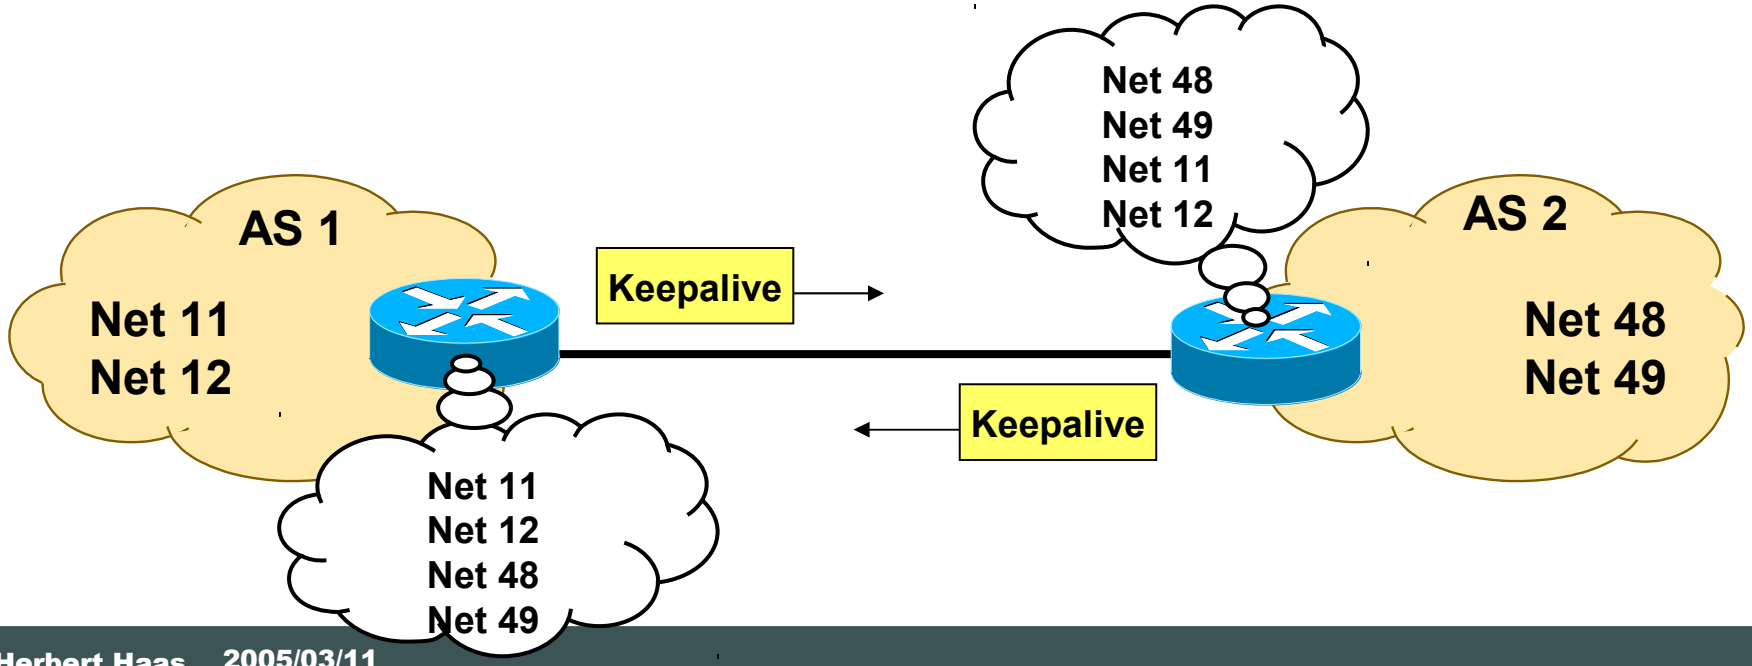

Topology Change:

- **Incremental** Updates upon topology or attribute changes
- **Withdraw** message upon loss of network

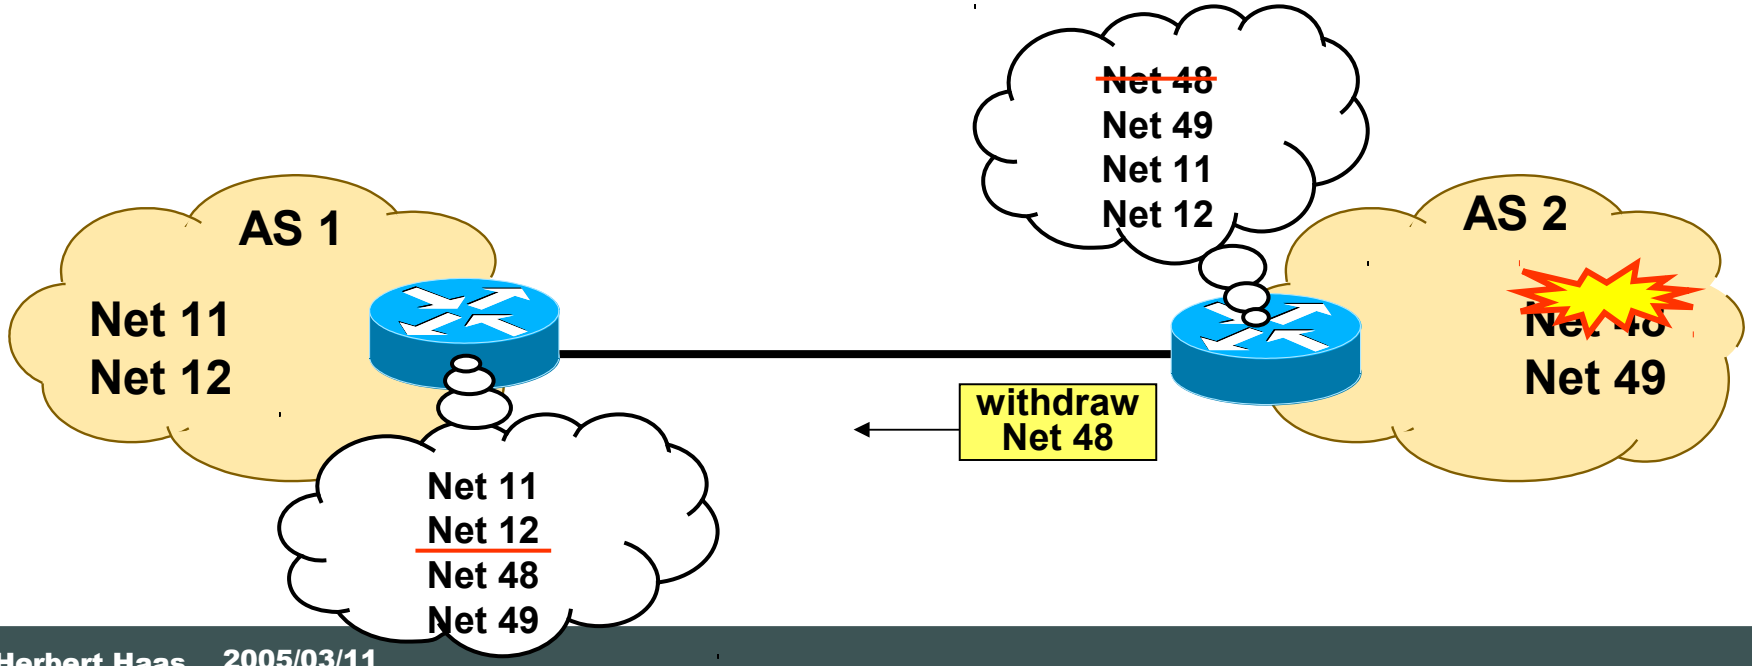

RIB

- BGP routing information is stored in RIBs
- RIBs might be combined (vendor specific)
- Only best paths are forwarded to the neighboring ASs
- Alternative paths remain in the BGP table
  - "Feasible routes" in Adj-RIB-In
  - Are used if the original path is withdrawn

BGP Routing Information Bases

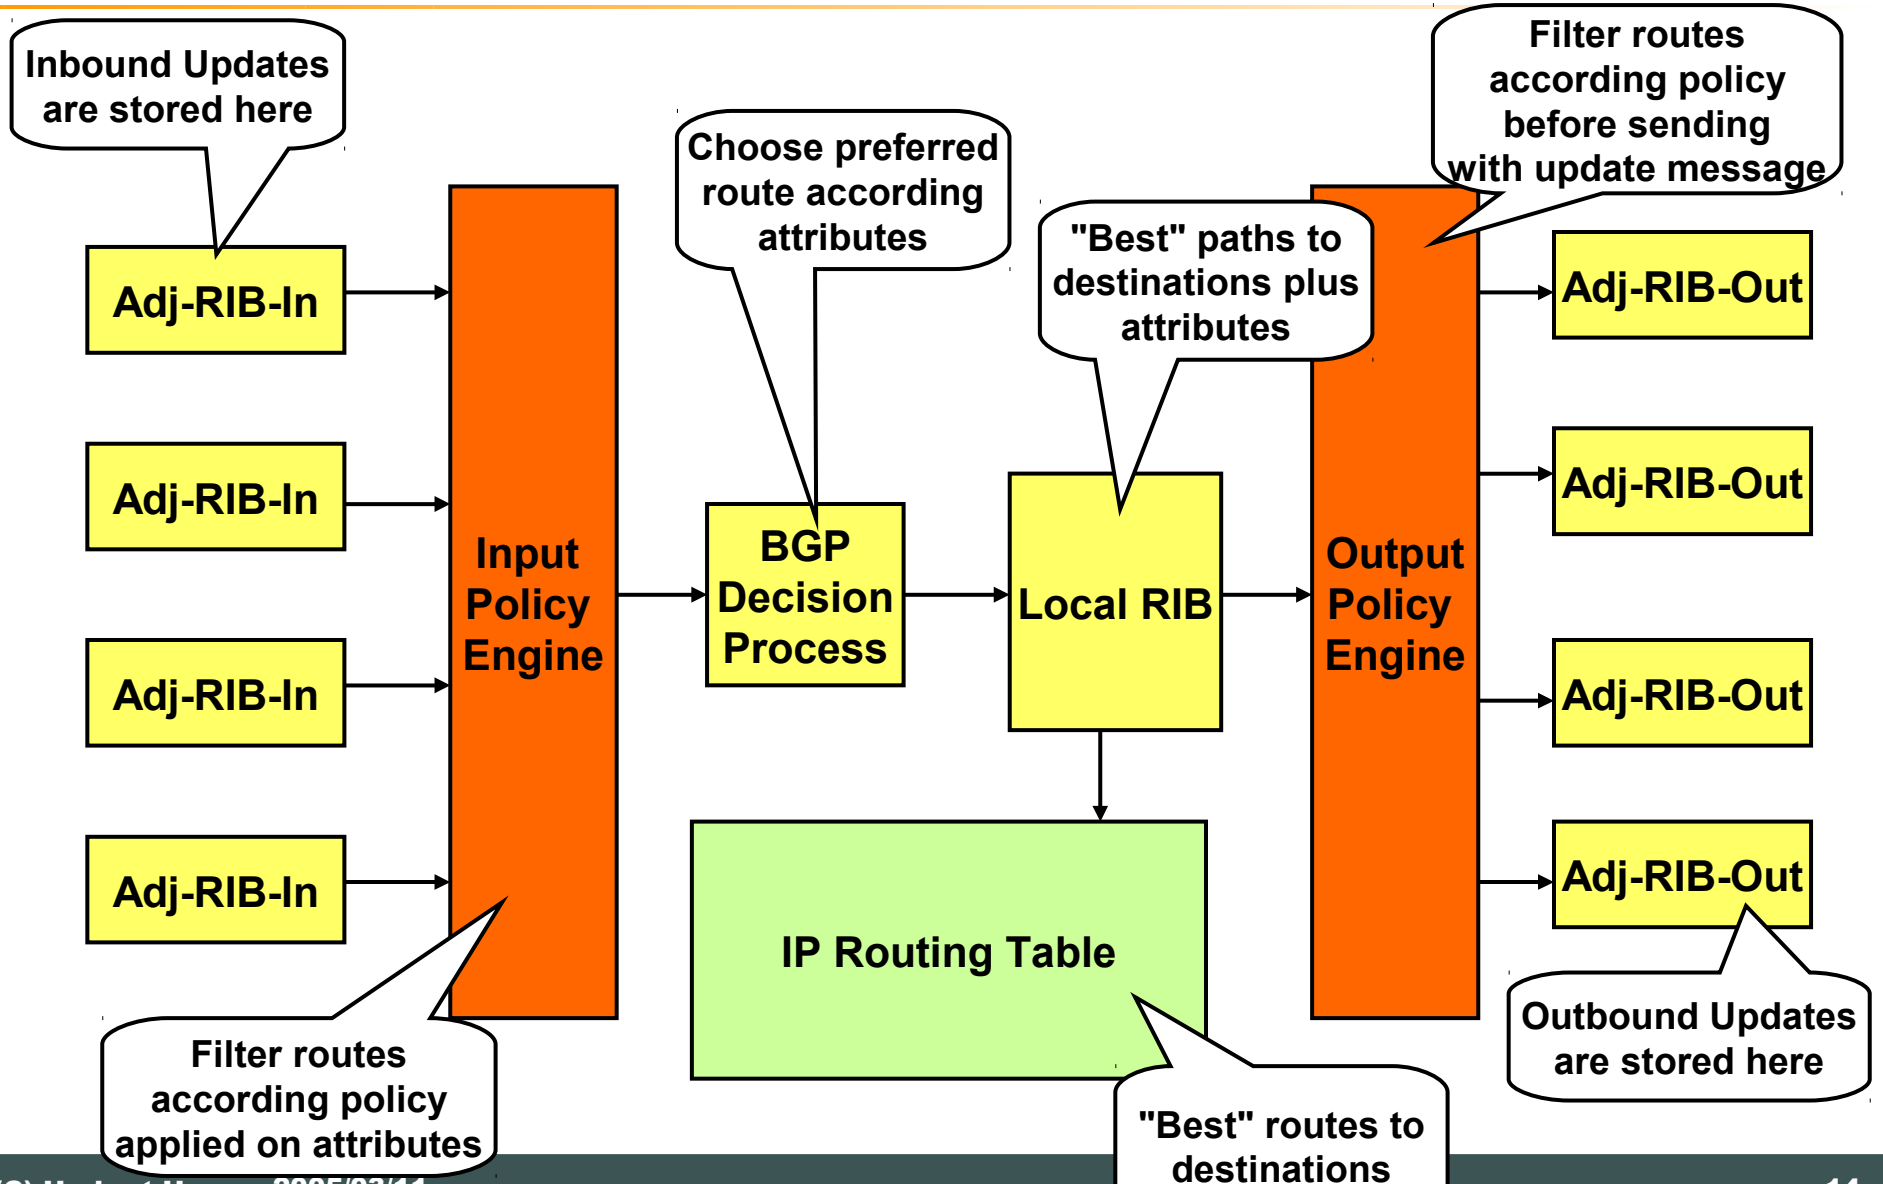

Quiz

- How many routes are maintained by BGP today?
- How many AS-numbers have been defined already?
- How long is the typical BGP convergence time?