CS-435
spring semester 2020

Network Technology & Programming Laboratory

University of Crete
Computer Science Department

Stefanos Papadakis

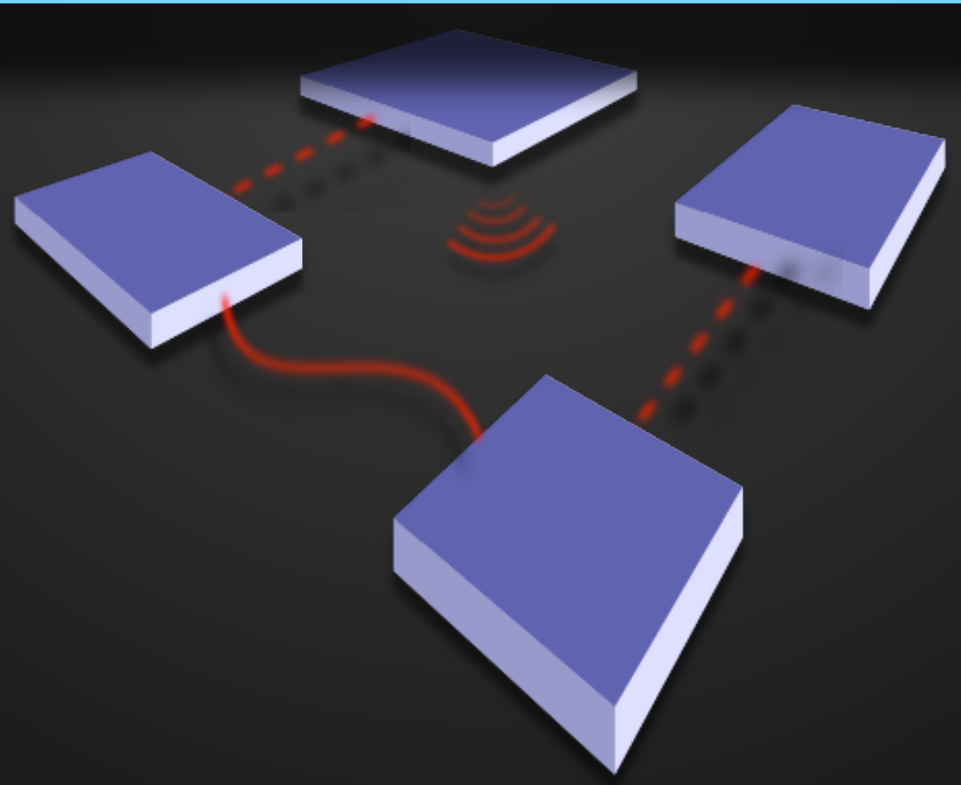

CS-435

Lecture #09 preview

- Dynamic Routing Protocols
- IGPs
- EGPs

Autonomous Systems

- Collection of IP routing prefixes
- Common routing policy to the Internet
- One or more network operators
- Acronym: AS

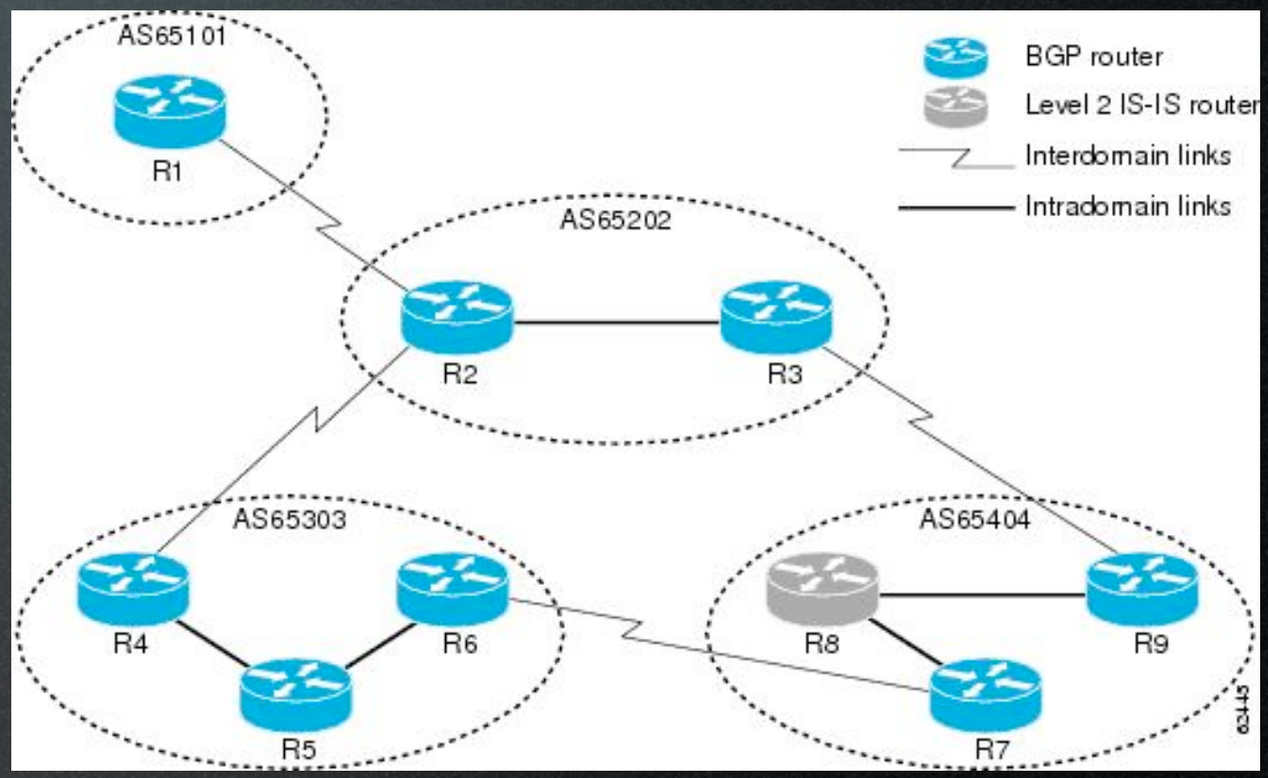

Autonomous Systems

- Internet Assigned Numbers Authority (IANA)
  - Regional Internet Registries (RIRs)
- Autonomous System Number (ASN)
  - Unique ASN allocated to each AS (network ID)
  - ASN
    - old: 16bit
    - new: 32bit
      - 0.y y=16bit same as old
      - 1.y reserved

Autonomous System Number

16bit ASN

- 0 reserved - may be used to identify non-routed networks
- 58368-64495 reserved by the IANA
- 64496-64511 reserved for documentation and sample code
- 64512-65534 designated for private use (allocated to the IANA)
- 65535 reserved

32bit ASN

- 0-65535 same as in 16bit
- 65536-65551 reserved for documentation and sample code
- 65552-131071 reserved
- 4294967295 reserved

Autonomous System Number

- RIPE NCC (total numbers on 1/1/2011):
  - 21025 ASNs
    - 810 are 32bit
- assigned in 2010: 2437 ASNs
  - 1854 16bit
  - 583 32bit

Autonomous System Number

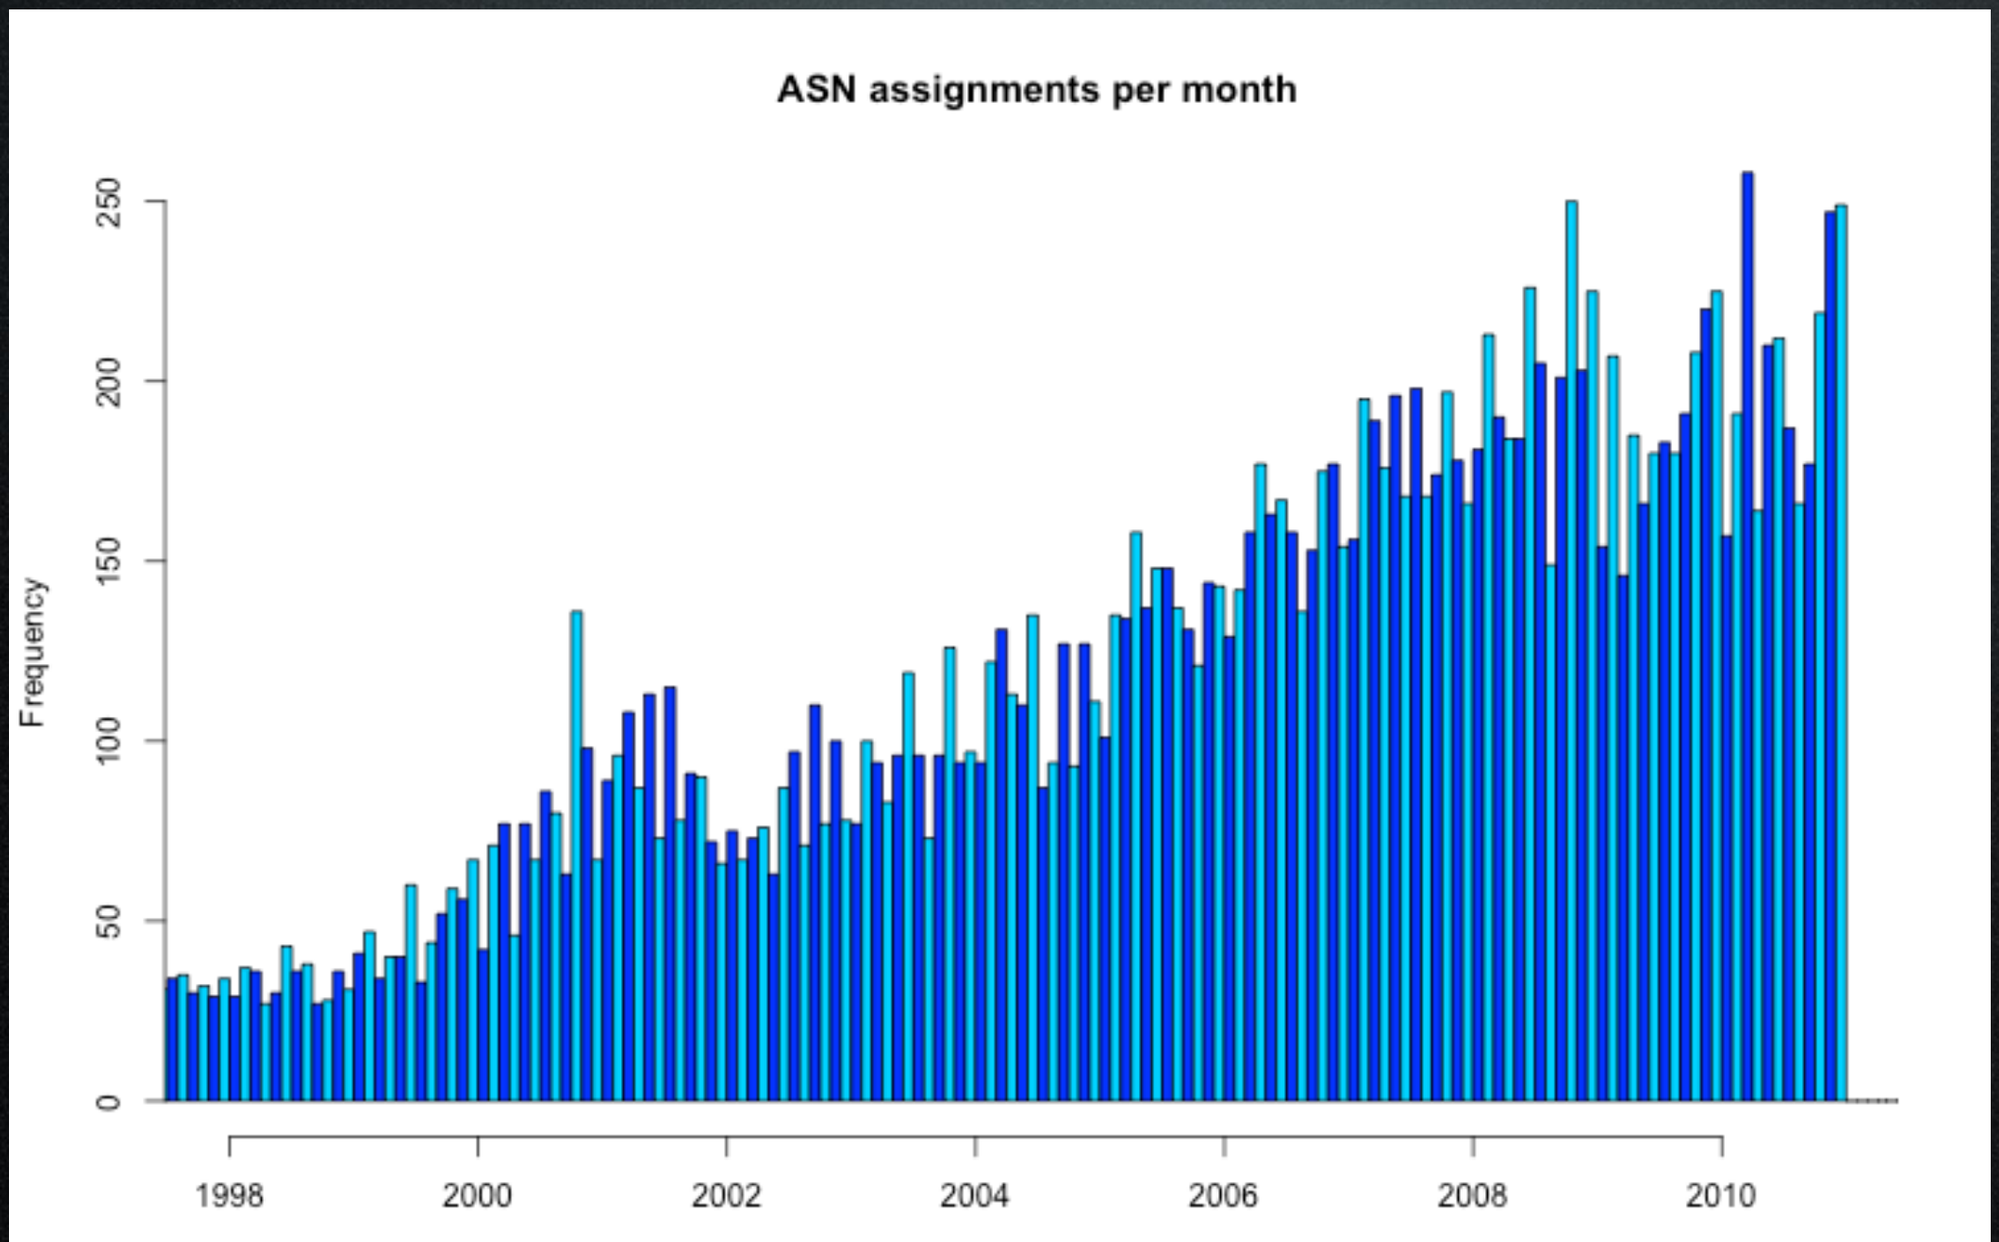

RIPE NCC

Autonomous System Number

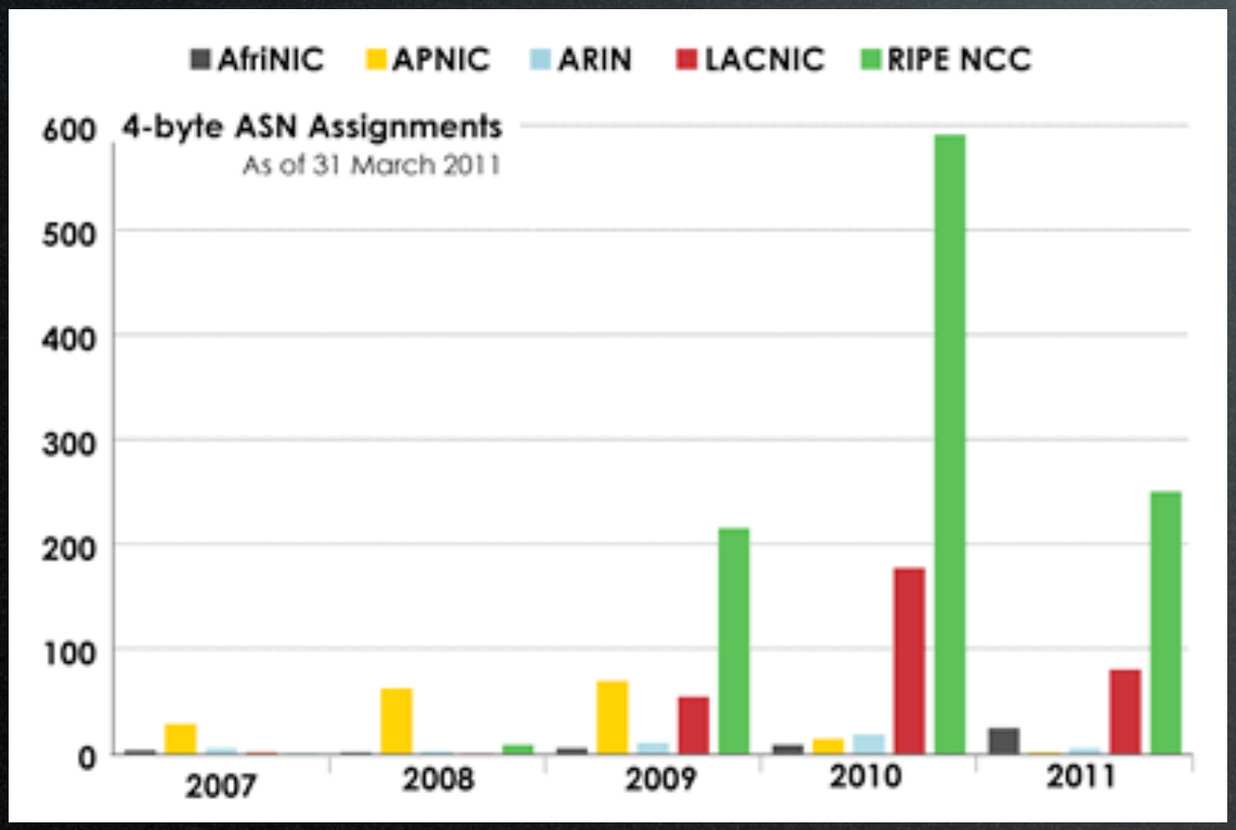

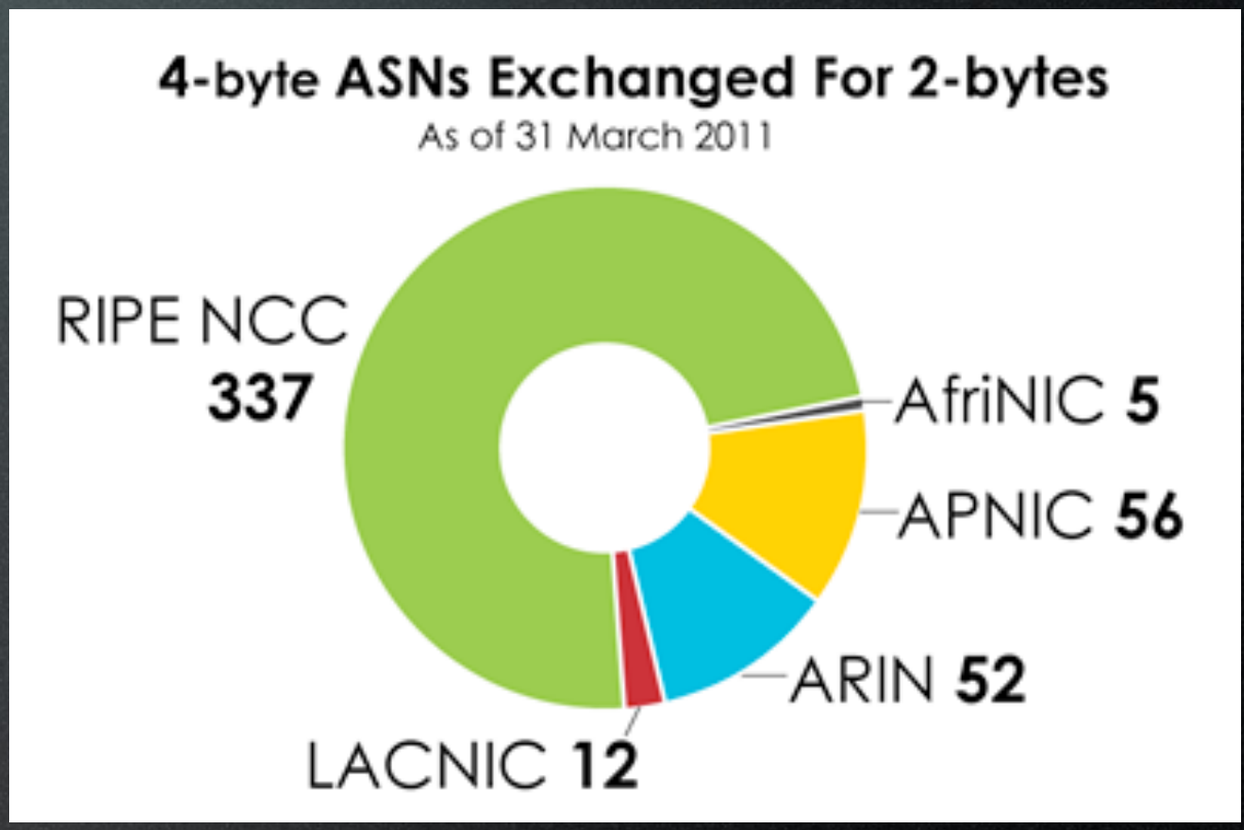

32bit ASN

Regional Internet Registries

- RIR organization manages allocation & registration of:
  - IP addresses
  - ASNs
- 5 RIRs:
  - African Network Information Centre (AfriNIC)
  - American Registry for Internet Numbers (ARIN)
  - Asia-Pacific Network Information Centre (APNIC)
  - Latin America and Caribbean Network Information Centre (LACNIC)
  - Réseaux IP Européens Network Coordination Centre (RIPE)

Regional Internet Registries

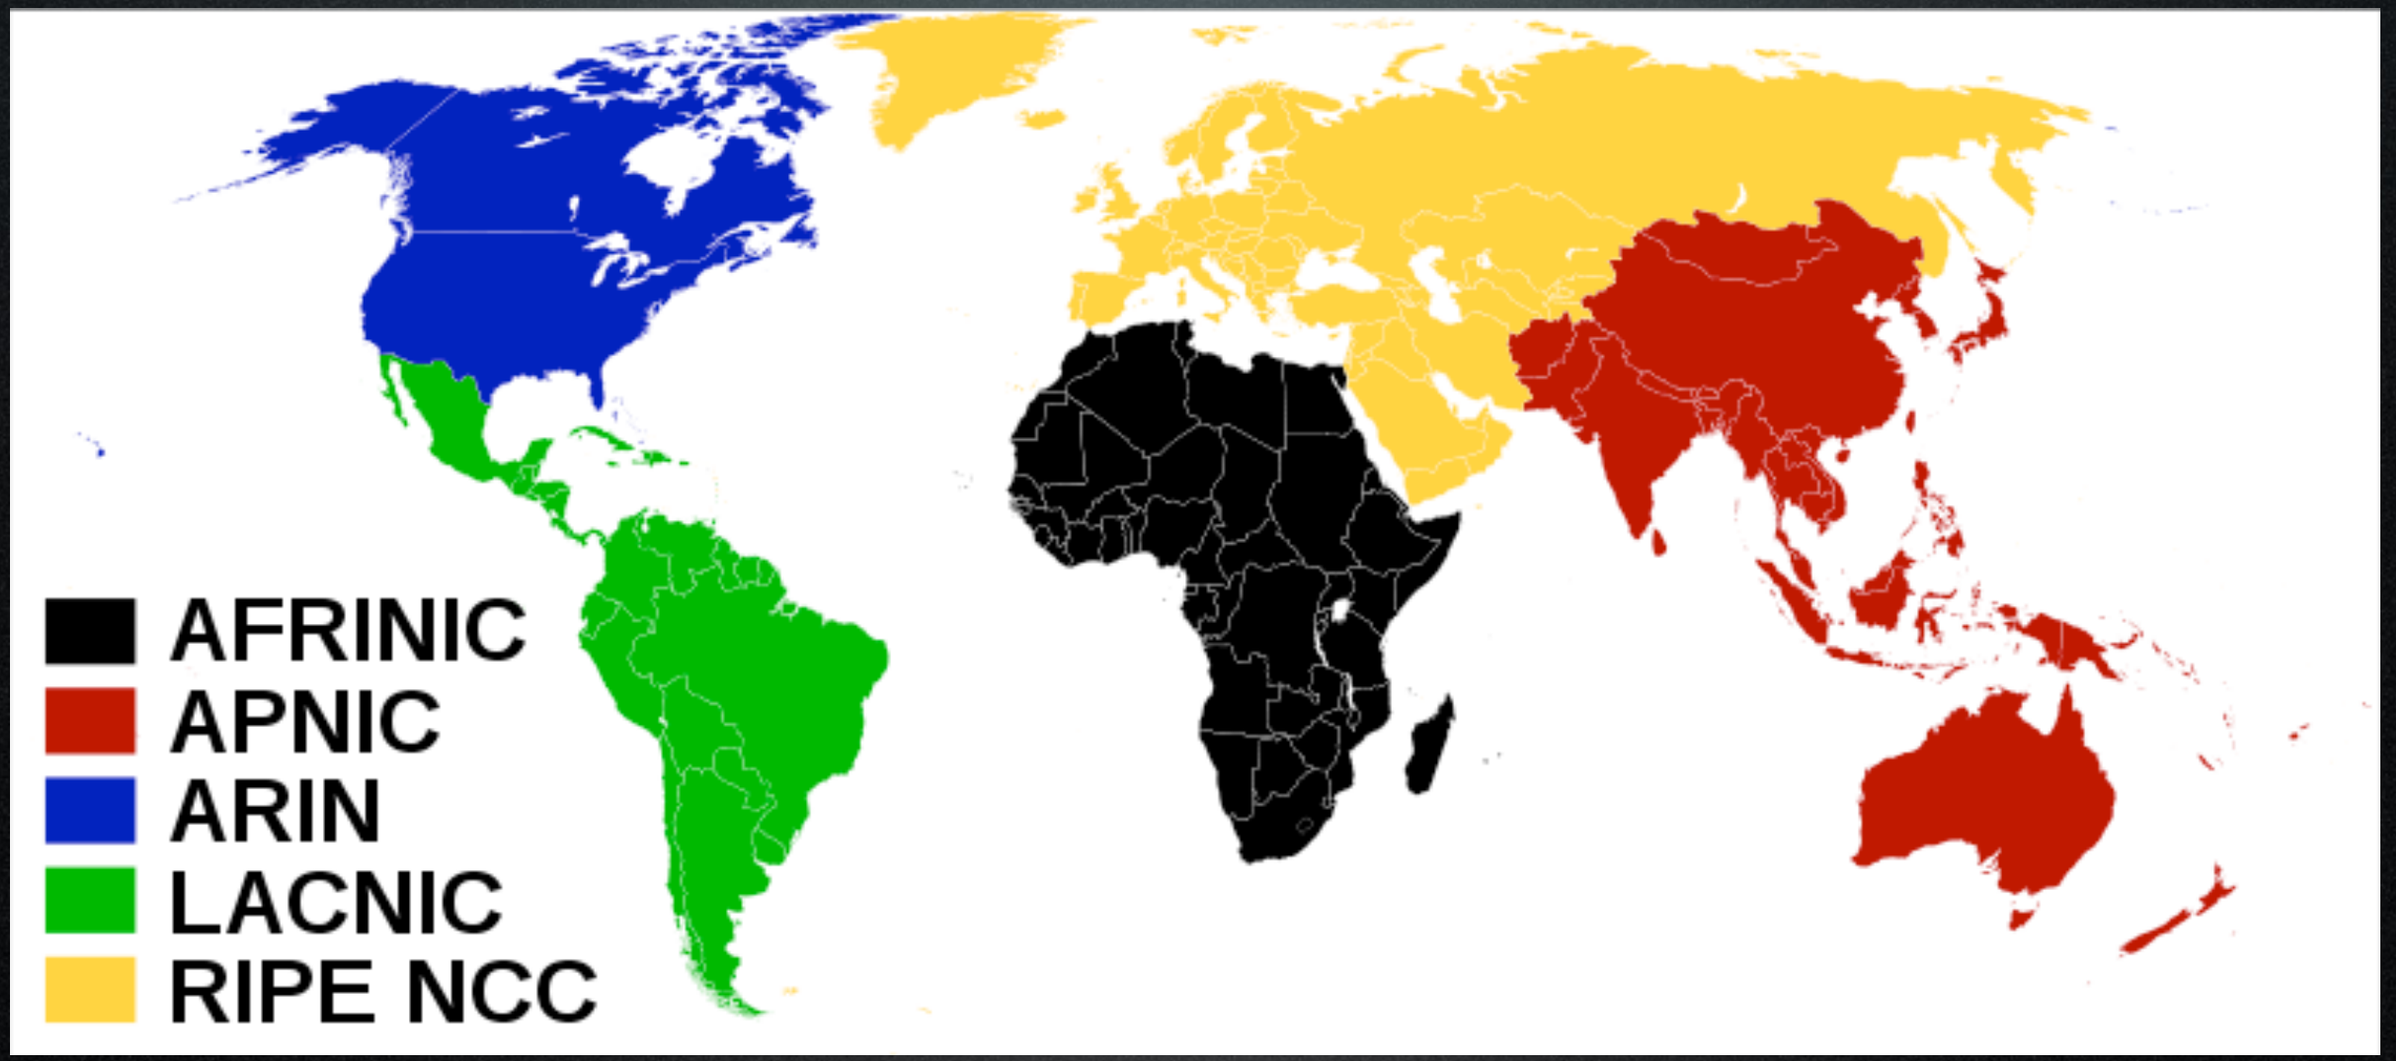

Number Resource Organization

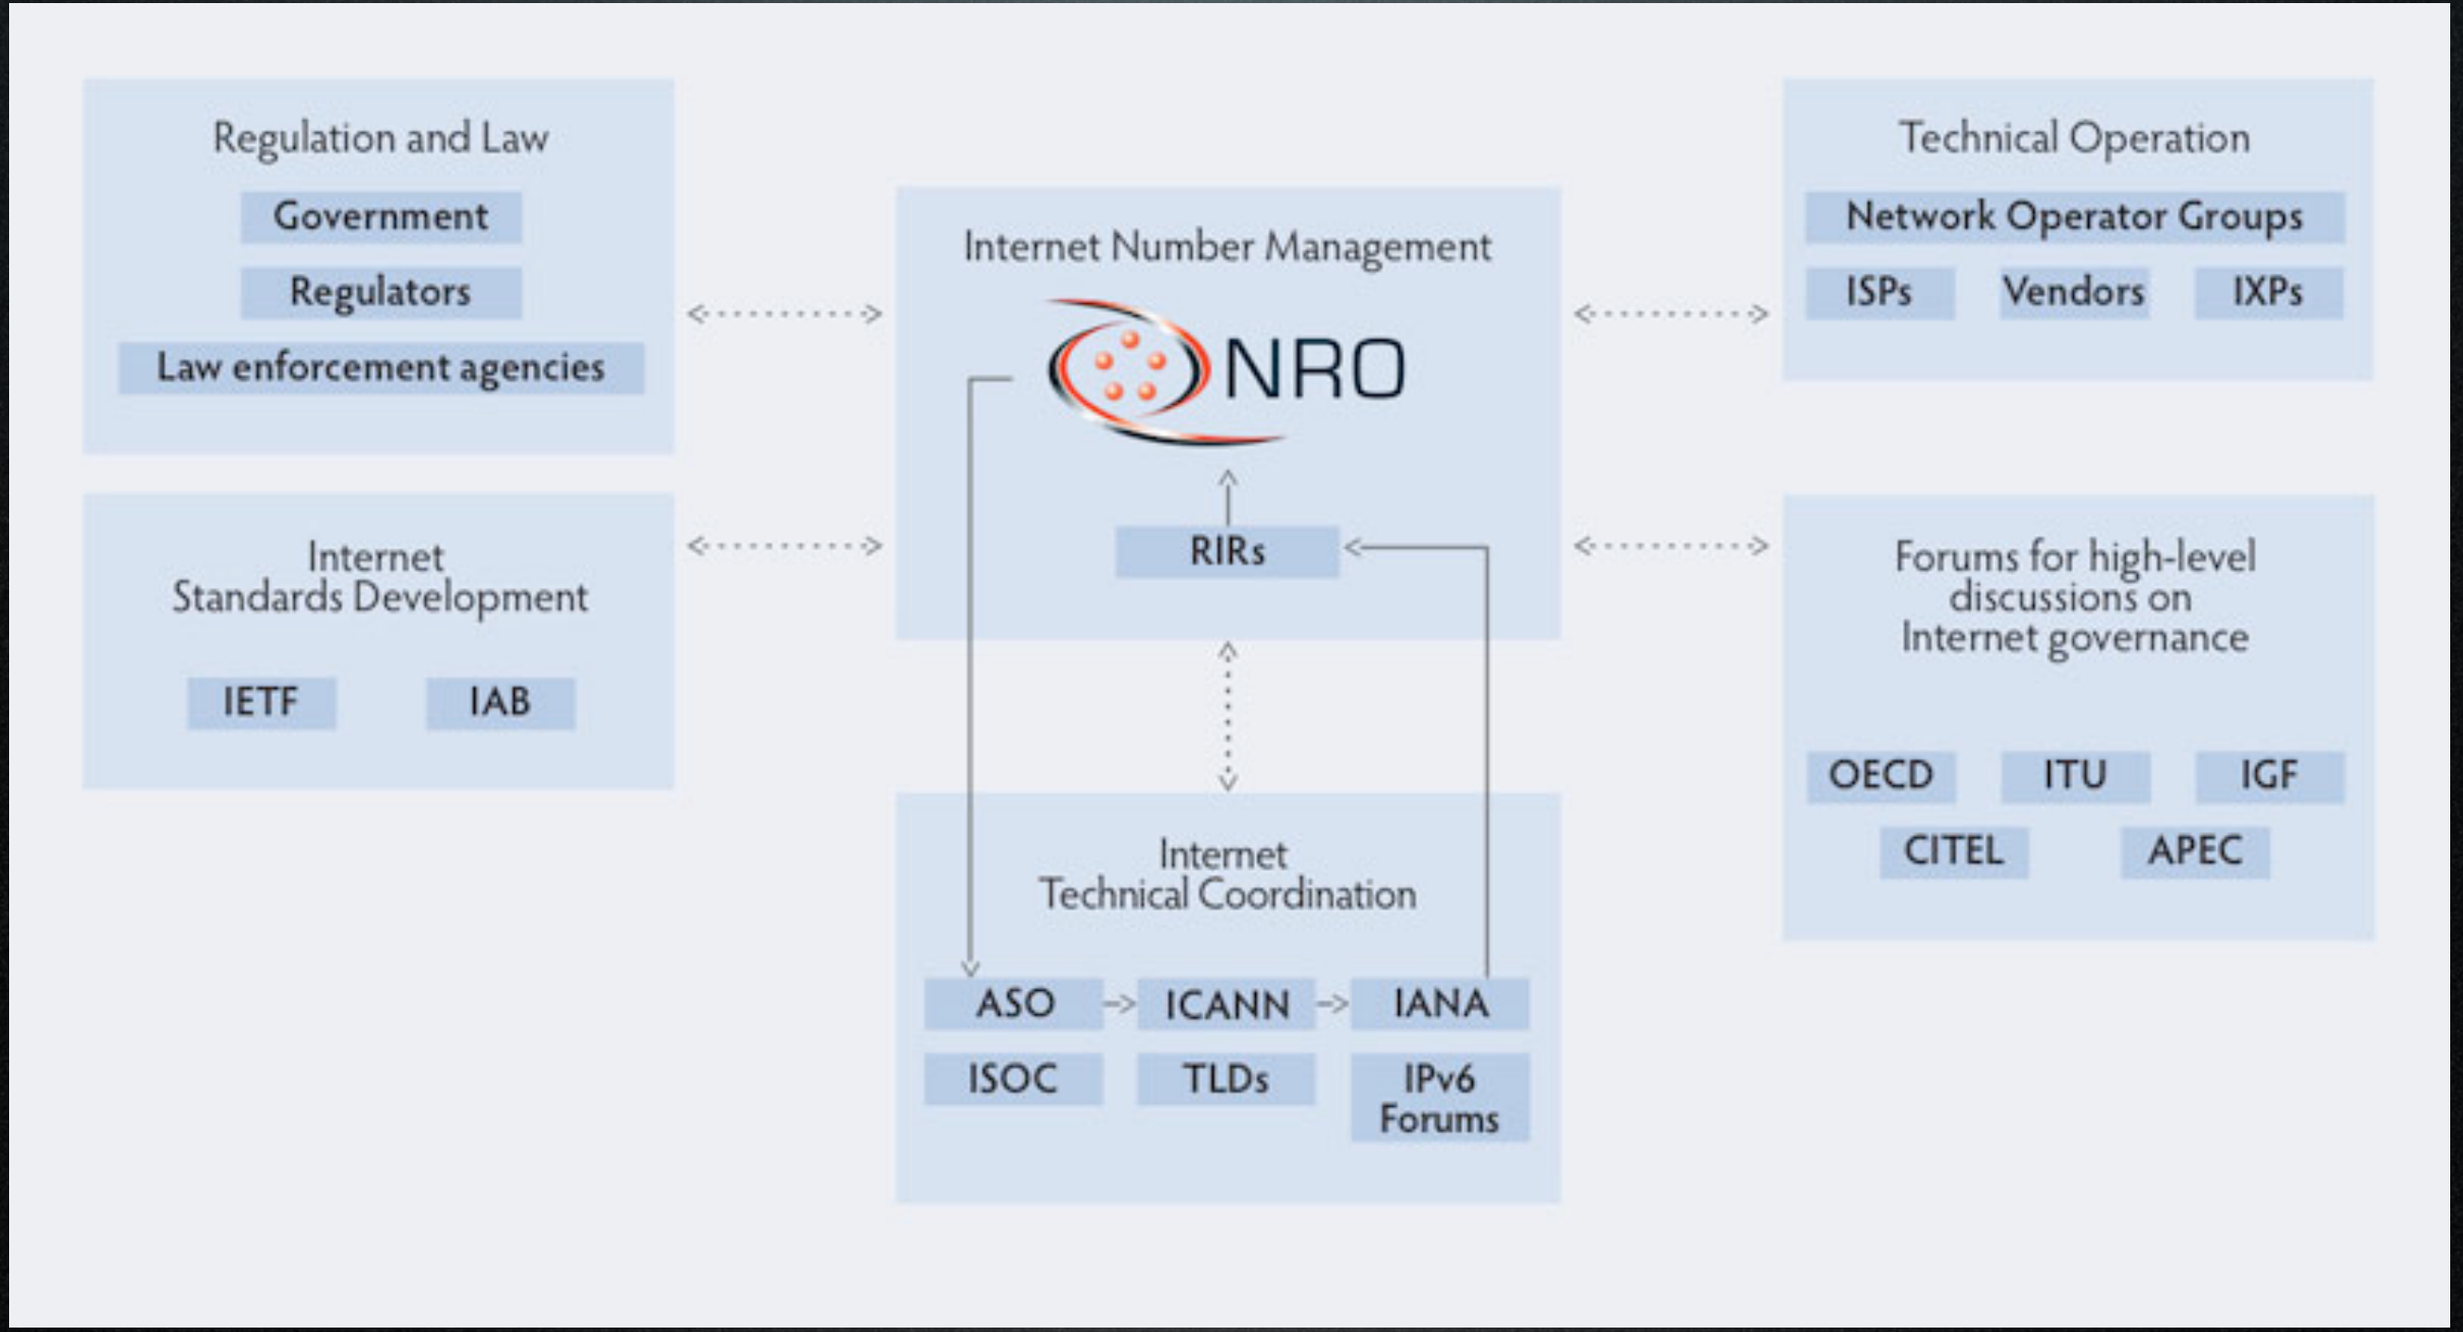

numbering allocations delegation

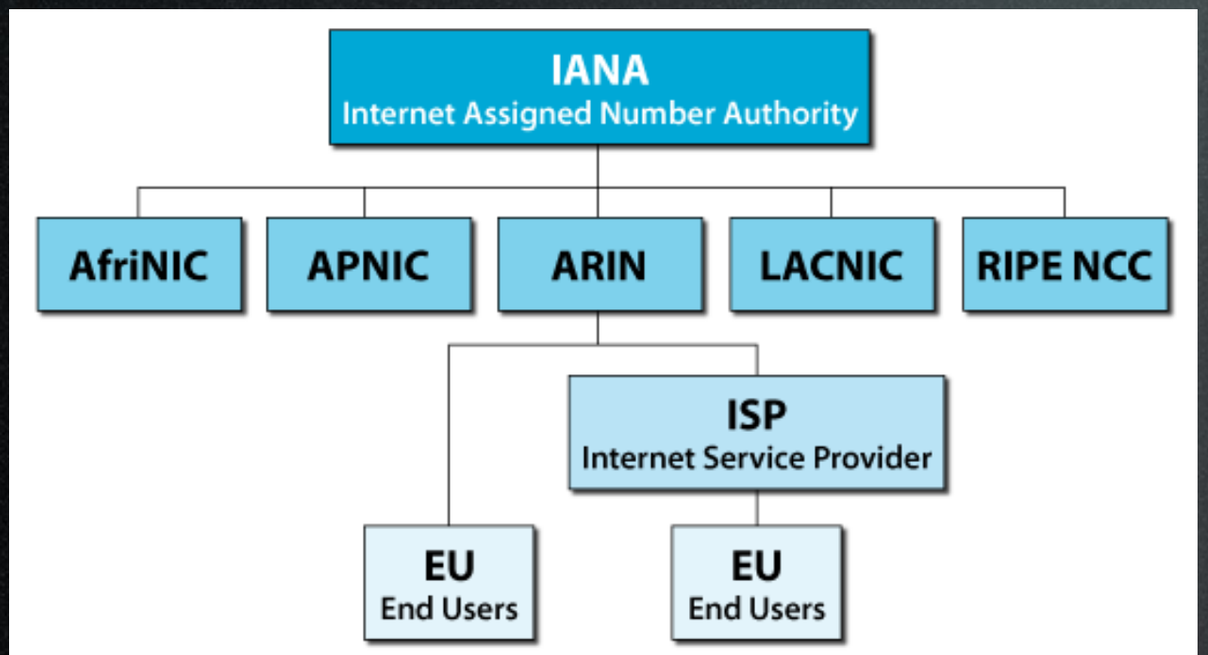

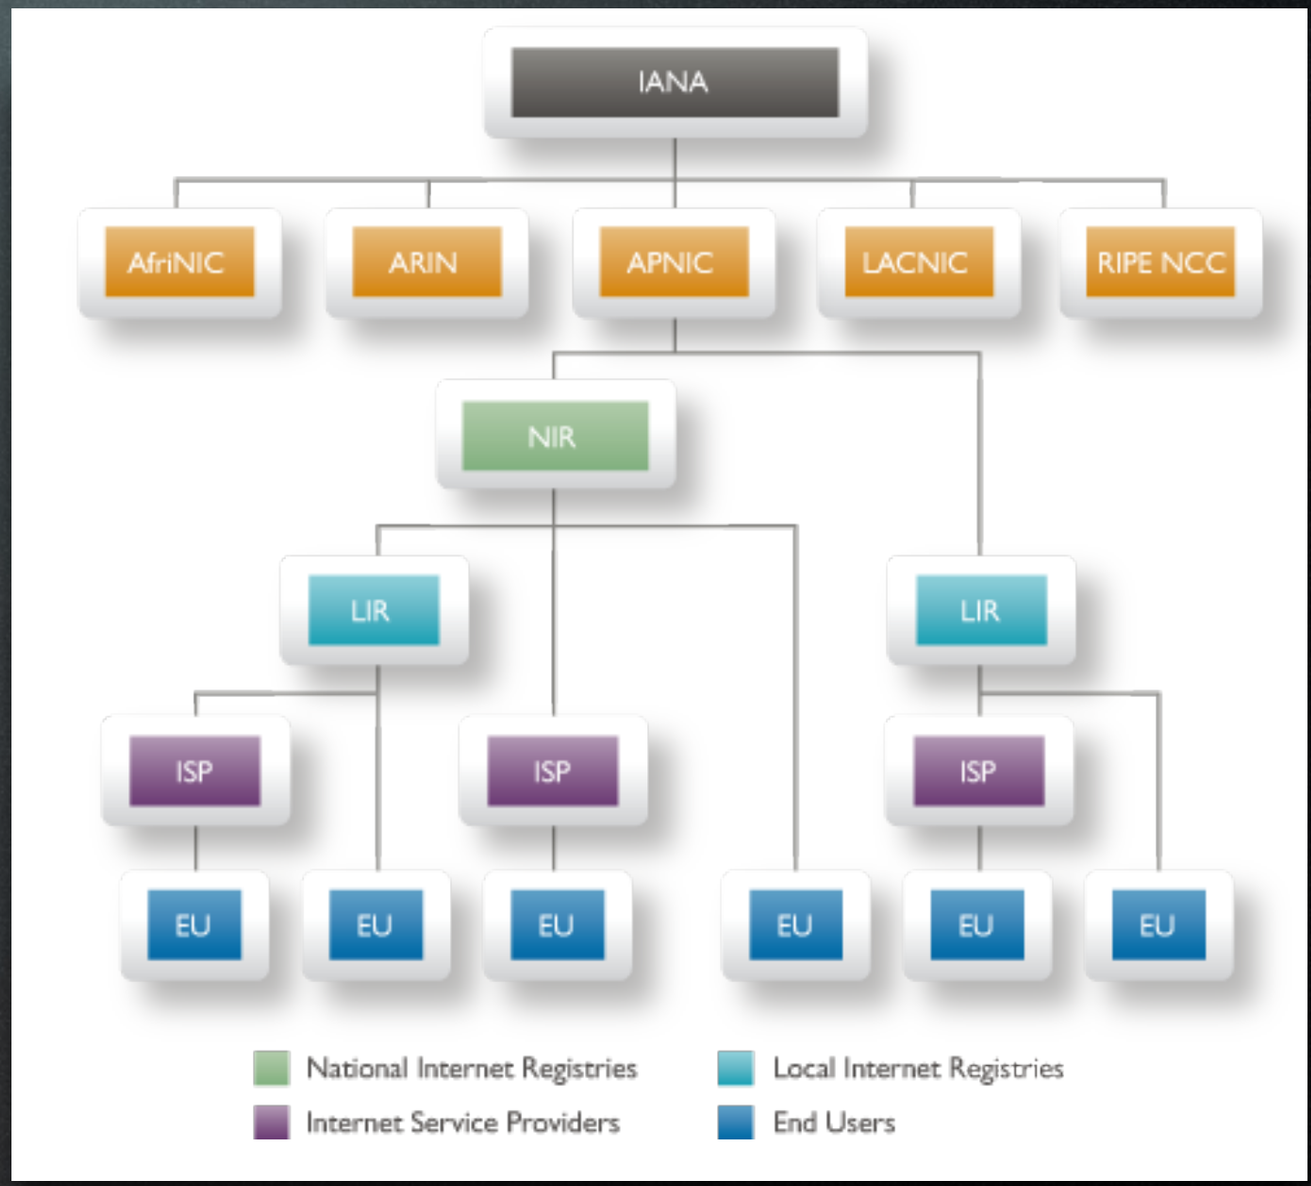

IP routing

- Forwarding / Routing Table
- Static routing
  - manual setup of routes
- Dynamic routing
  - automatic setup of routes

Routing

- correct delivery at reasonable level of performance
- paths with no cycles
- functions:
  - topology update
  - route computation
  - switching

Routing Protocols

- Distance Vector
  - Link cost
  - minimum total cost
  - table exchange
- Link-State
  - Nodes graph
  - shortest path
  - flooding
- Path Vector
  - similar to distance vector
  - table exchange between **speaker** nodes
  - no metrics

Routing Protocols

- Intra-domain / Interior Gateway Protocols (IGPs)
  - routing within an AS
  - ignores the Internet outside
- Inter-domain / Exterior Gateway Protocols (EGPs)
  - routing between ASs
  - assumes that the Internet consists of interconnected ASs

Routing Protocols

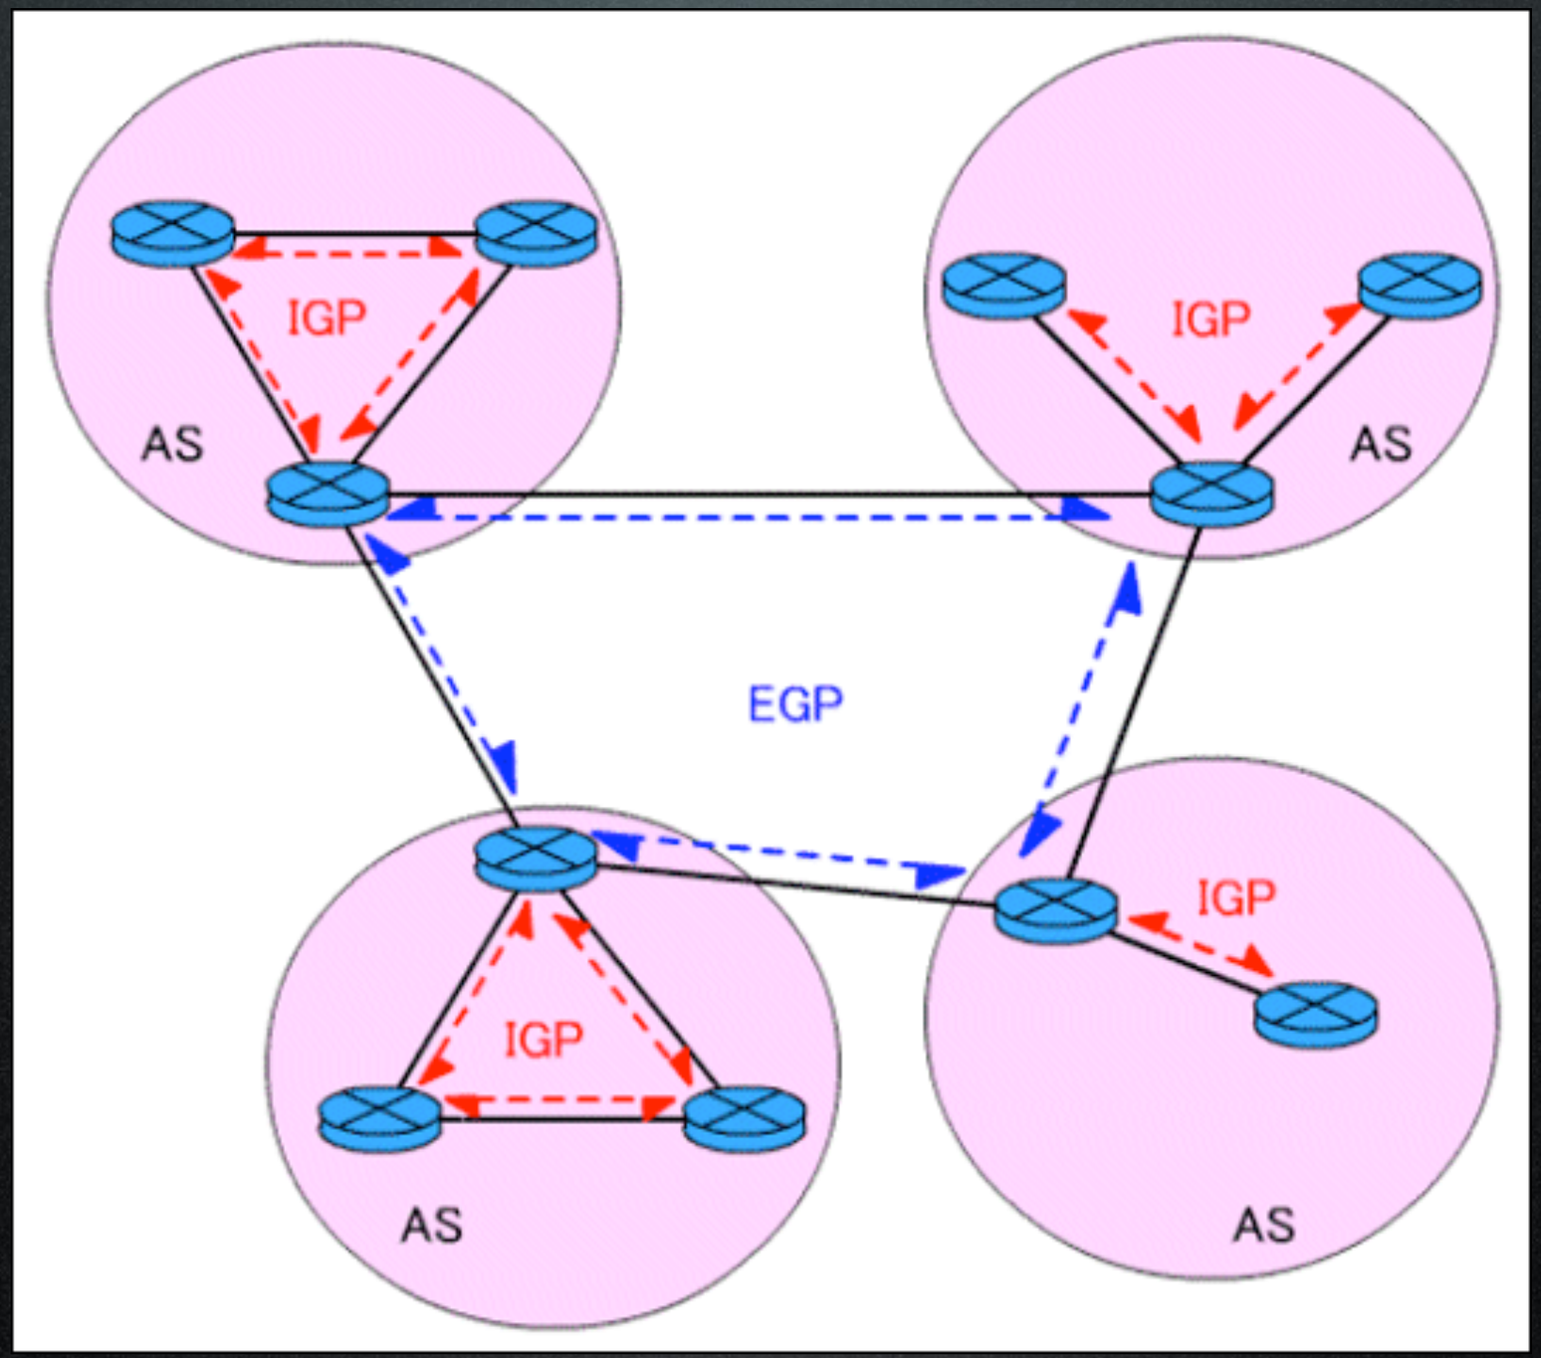

Routing Protocols

- Intra-domain / Interior Gateway Protocols (IGPs)
  - RIP
    - IGRP (Cisco)
    - EIGRP (Cisco)
  - OSPF
    - IS-IS
- Inter-domain / Exterior Gateway Protocols (EGPs)
  - BGP
    - EGP3

Routing Information Protocol
RIP

- distance-vector
- Bellman-Ford algorithm
- metric: hop count
- maximum number of hops: 15
- hop count 16: infinite
- uses UDP

RIP

- RIP version 1
  - classful routing - no variable length subnet masks
    - impossible to have subnets with different sizes
  - broadcast messages
  - maximum 16 hop counts
- RIP version 2
  - CIDR support
  - multicast messages to adjacent routers (224.0.0.9)
  - maximum 16 hop counts
- RIPng
  - IPv6 support
  - maximum 16 hop counts

RIP

- normal routing updates: every 30s
- triggered updates:whenever the metric for a route changes
- after 3 minutes routes are set to 16 (inf) if not updated
- takes a long time to stabilize
- count to infinity problem
- maximum path is 15 hops

Dynamic Routing Protocols

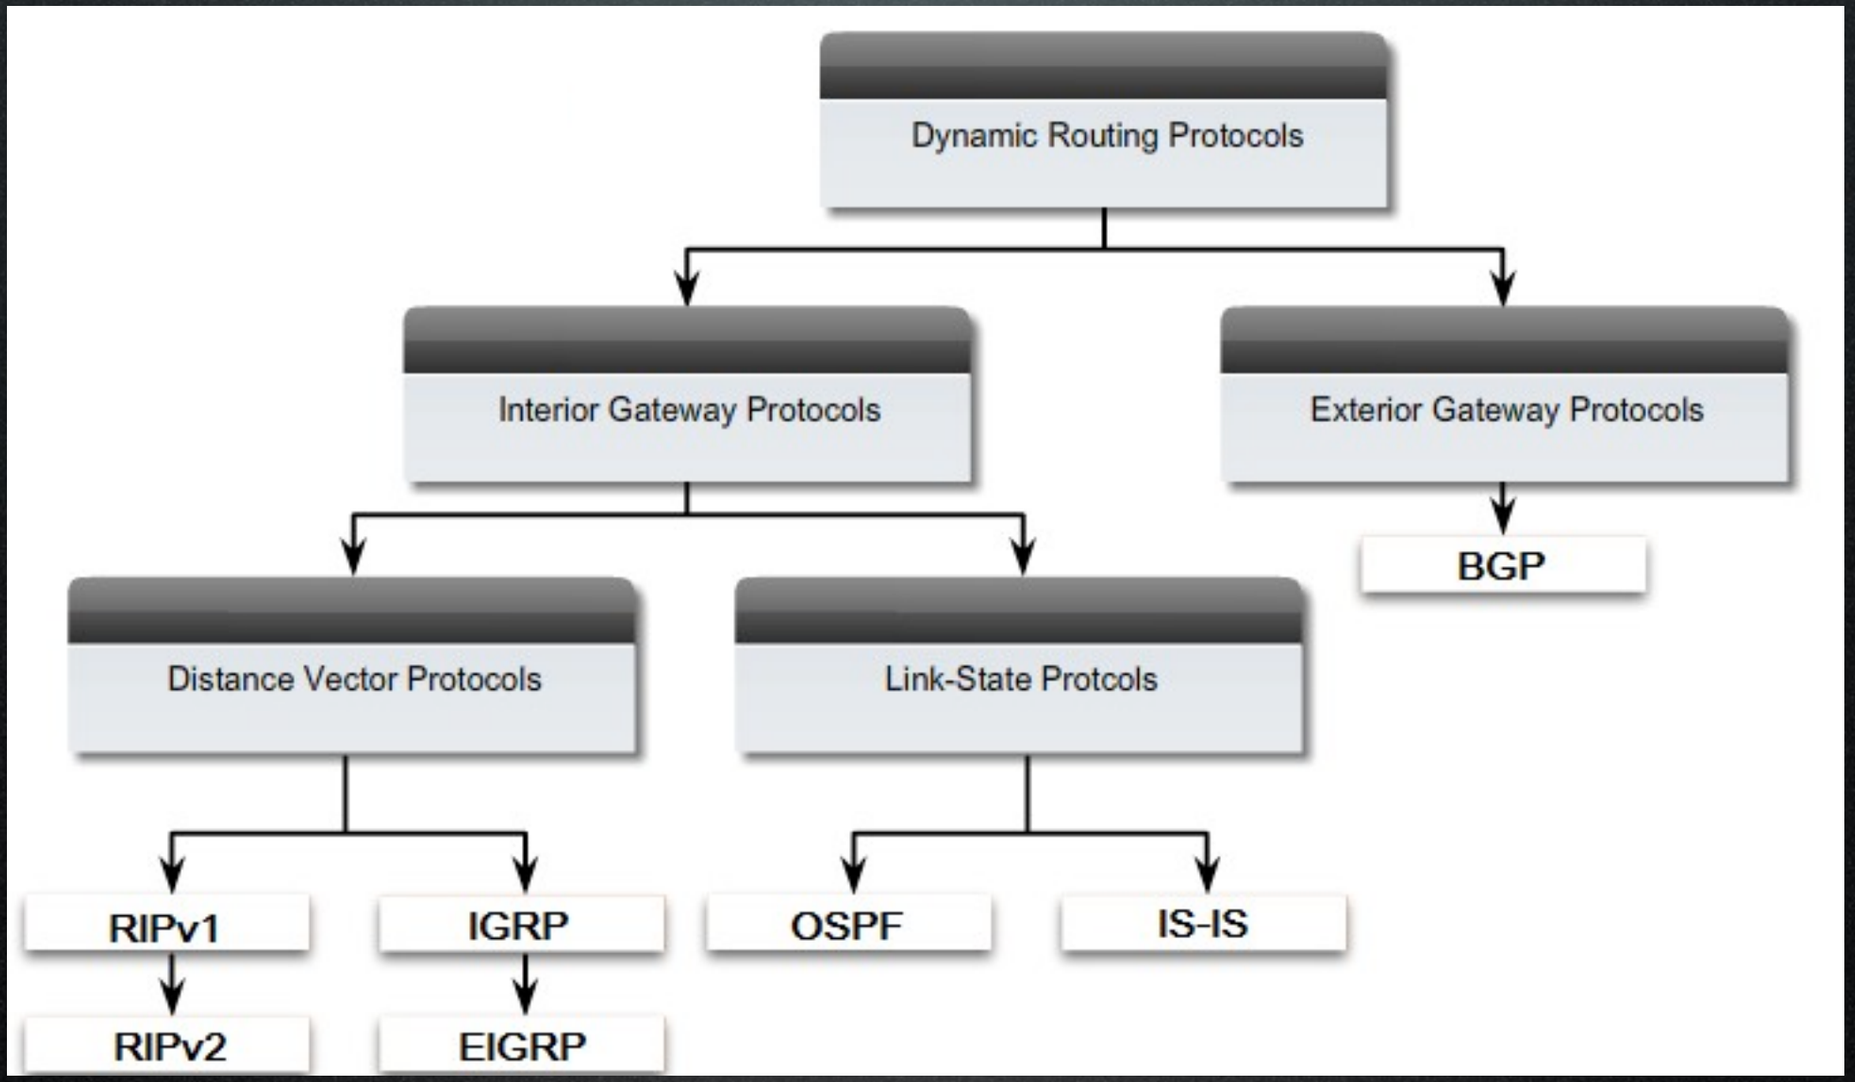

RIP vs OSPF

PROPERTIES	RIPv2	OSPFv2
IPv6 version / Update	RIPng	OSPFv3
Developed	1993 (standarized in 1998)	defined in 1998 (v3 in 2008)
Type	Distance Vector	Link-state
Administrative Distance	120	110
Max. Hop Count	15	No limit
Network Size	Small	Small to Large
Algorithm	Bellman-Ford	Dijkstra
Convergence	Slow	Fast
Variable Length Subnet Masks (VLSM)	Yes in v2, not v1	Yes
Routing Updates	Full Table every 30 seconds	LSA generated upon any routing change
Bandwidth Usage	uses more bandwidth	uses less bandwidth
Areas / Boundaries	No	Yes
Metric	Hop Count	Cost

Open Shortest Path First
OSPF

- link state routing protocol
- layer 3 protocol, runs on top of IP (IP datagrams protocol #89)
- version 2: IPv4
- version 3: IPv6
- routes packets within a single AS

Open Shortest Path First
OSPF

- Dijkstra's algorithm
- shortest path tree
- metric: link RTT, throughput, reliability
- supports CIDR
- Link State Database (LSDB) of entire network
  - tree-image of entire topology
  - refreshed periodically through flooding (multicast to 224.0.0.5 /224.0.0.6)

OSPF

- Areas - subdivisions of network
  - maintains a separate link state database
  - connecting router: summarized info towards the rest of the network
    - the topology of an area is unknown outside of the area
- 32bit - dot decimal notation (i.e. 0.0.0.1)
- Area 0: backbone/core area
  - all other areas must connect directly to Area 0
- Area Border Router: connects to backbone

OSPF

- Use of own IP protocol 89 & multicast for flooding
- OSPF router types:
  - Area border router (ABR)
  - Autonomous system boundary router (ASBR)
  - Internal router (IR)
  - Backbone router (BR)

OSPF

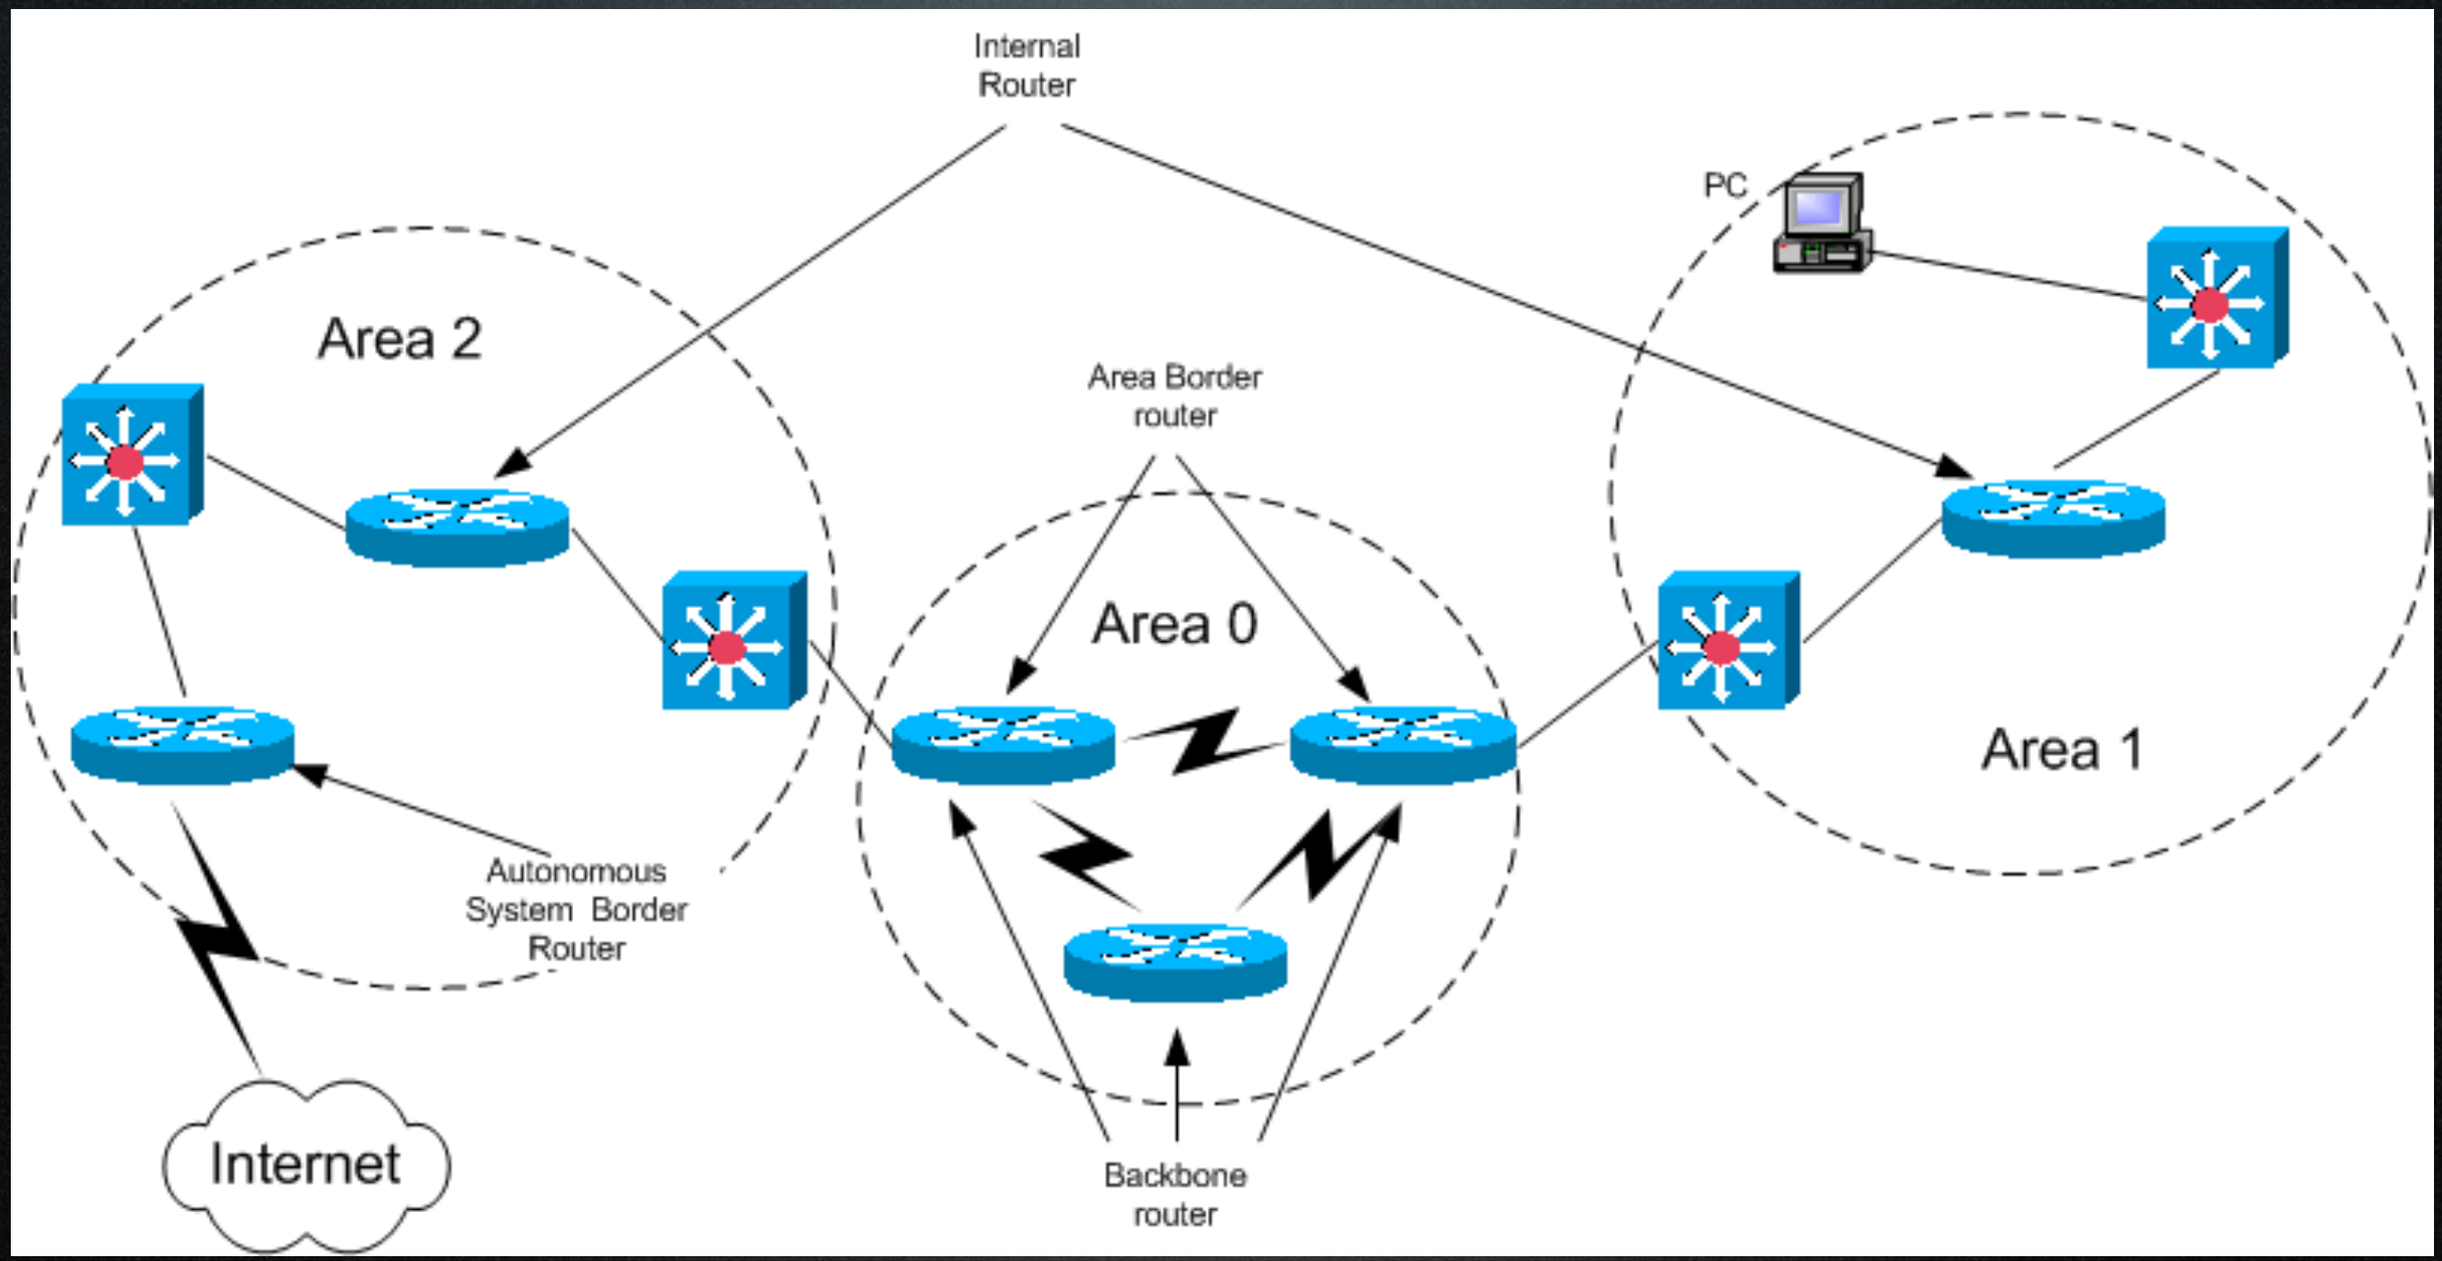
Internal
Router
PC
Area 2
Area Border
router
Area 0
Area 1
Autonomous
System Border
Router
Internet
Backbone
router

Intermediate System to Intermediate System
IS-IS

- the de facto standard for large service provider network backbones
- OSI network layer protocol
- does not use IP
- different usage of “areas” than in OSPF
  - levels: 1(intra area); 2(inter area); 1-2(both)
- lighter than OSPF -> scales better
- routes packets within a single AS

OSPF vs IS-IS

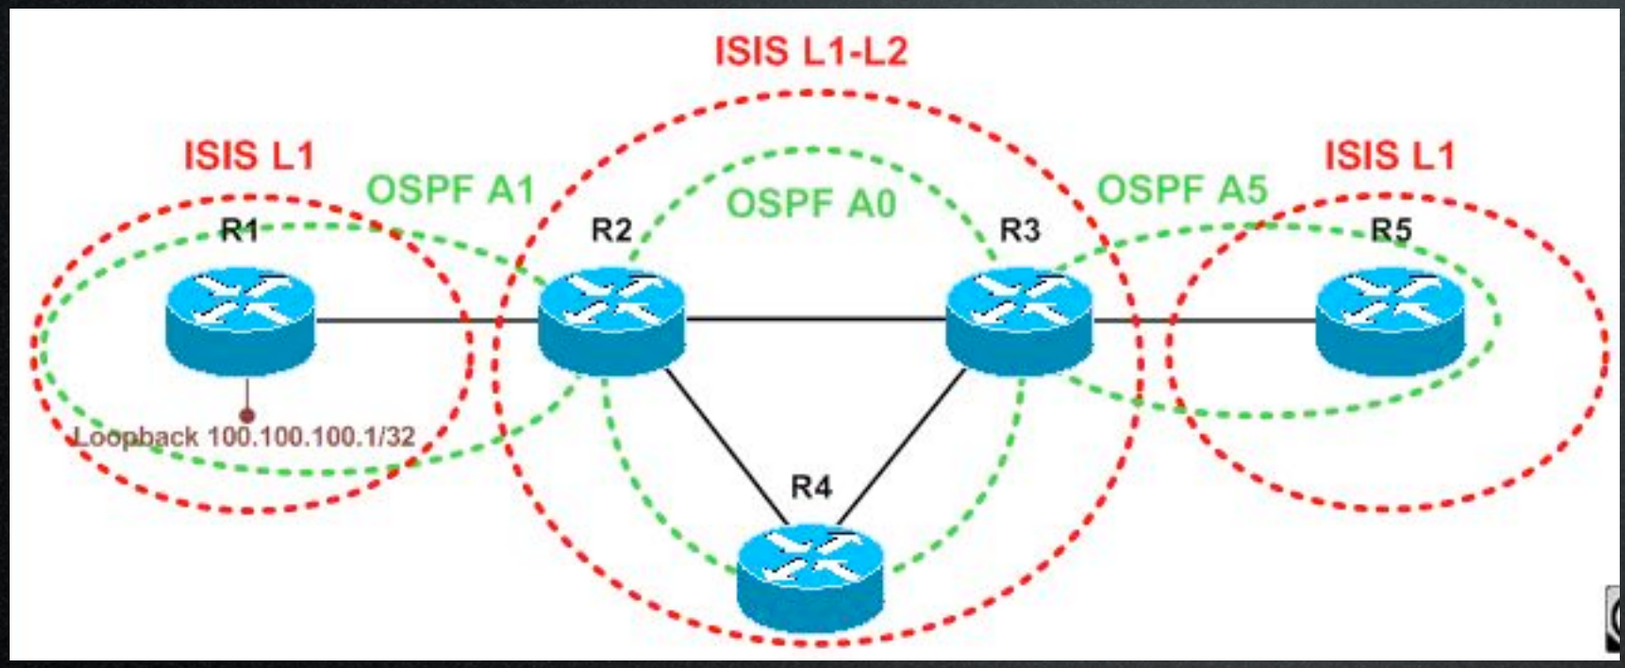
ISIS L1-L2
ISIS L1
OSPF A1
OSPF A0
OSPF A5
ISIS L1
R1
R2
R3
R5
R4
Loopback 100.100.100.1/32

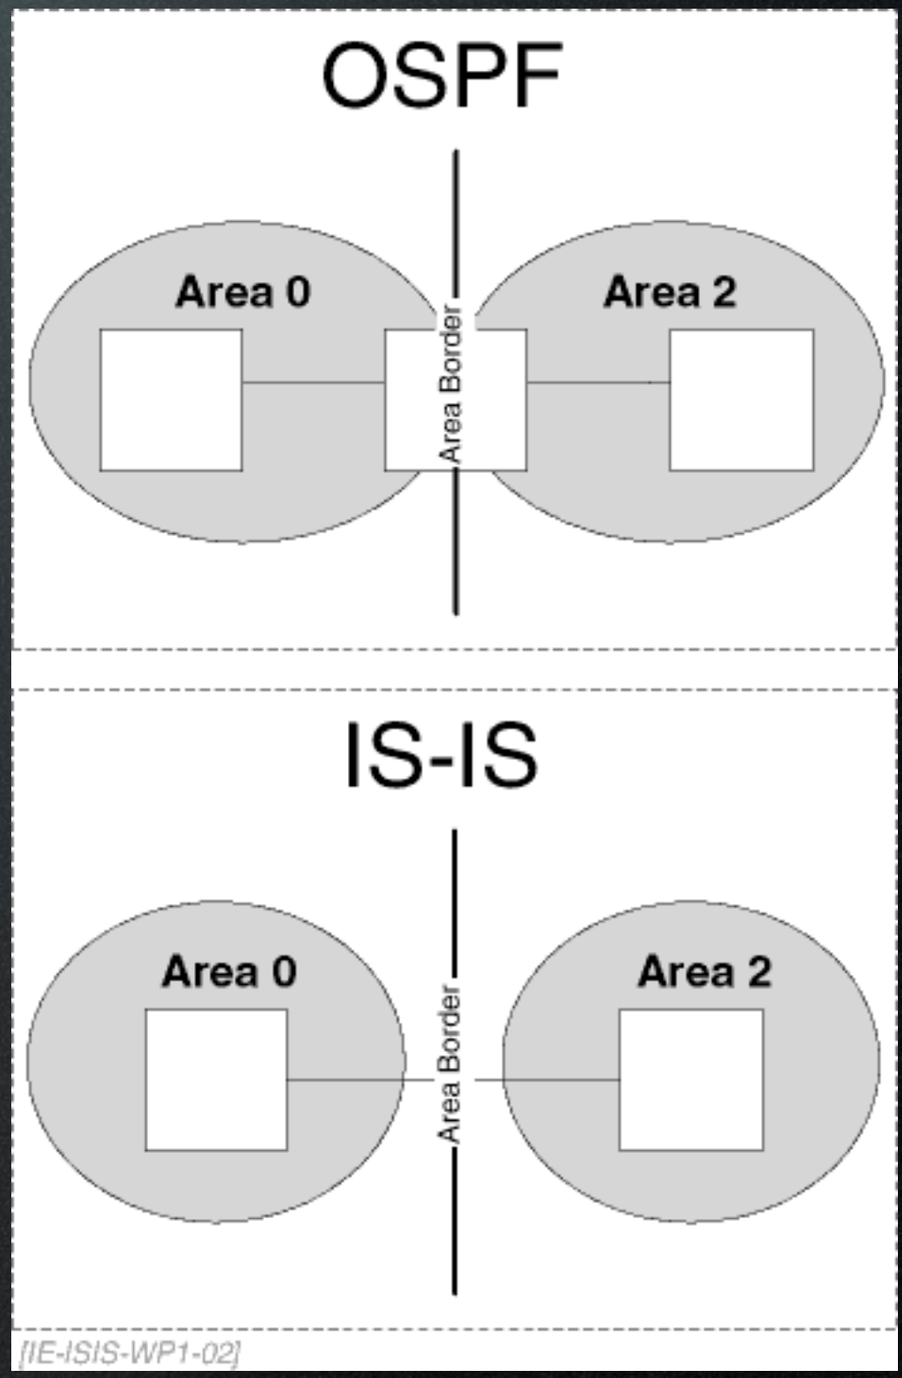
OSPF
Area 0
Area Border
Area 2
IS-IS
Area 0
Area Border
Area 2
[IE-ISIS-WP1-02]

IS-IS

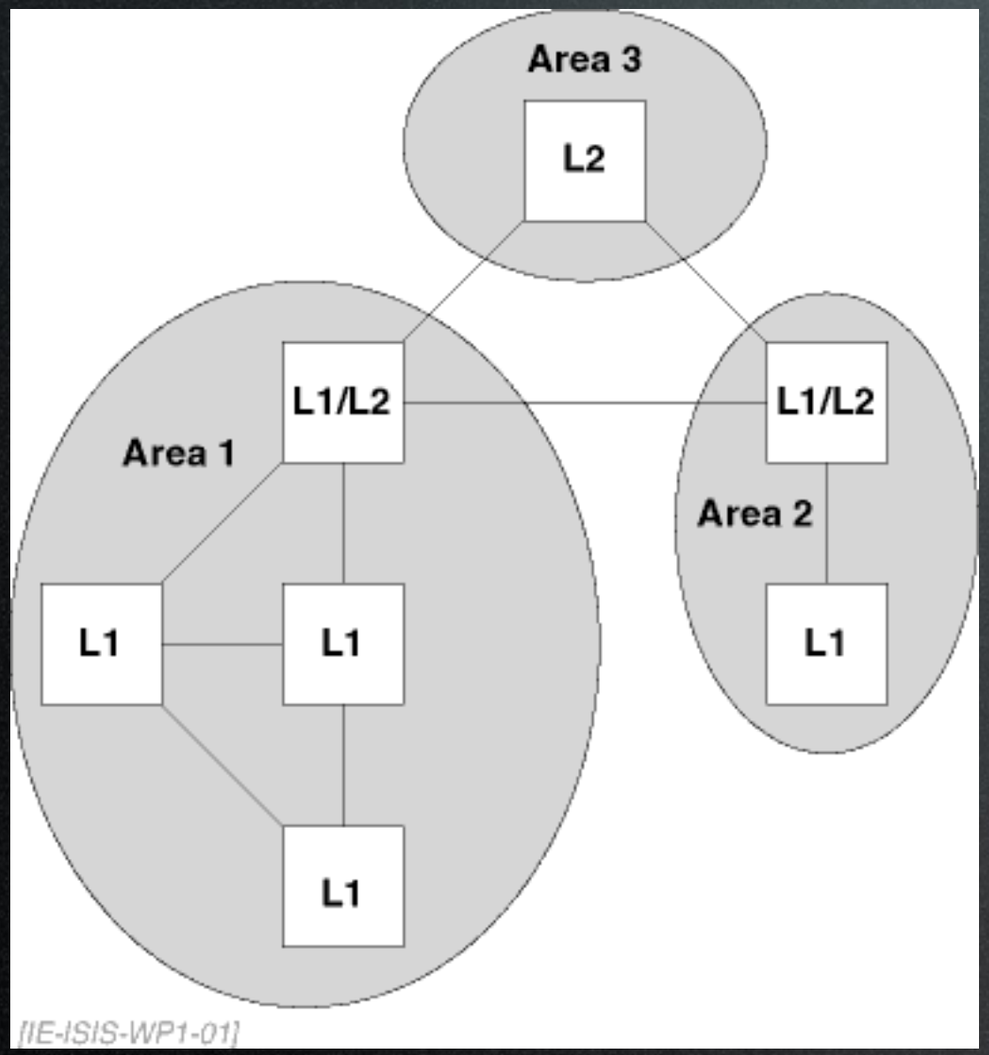
Area 3
L2
L1/L2
L1/L2
Area 1
Area 2
L1
L1
L1
L1
[IE-ISIS-WP1-01]

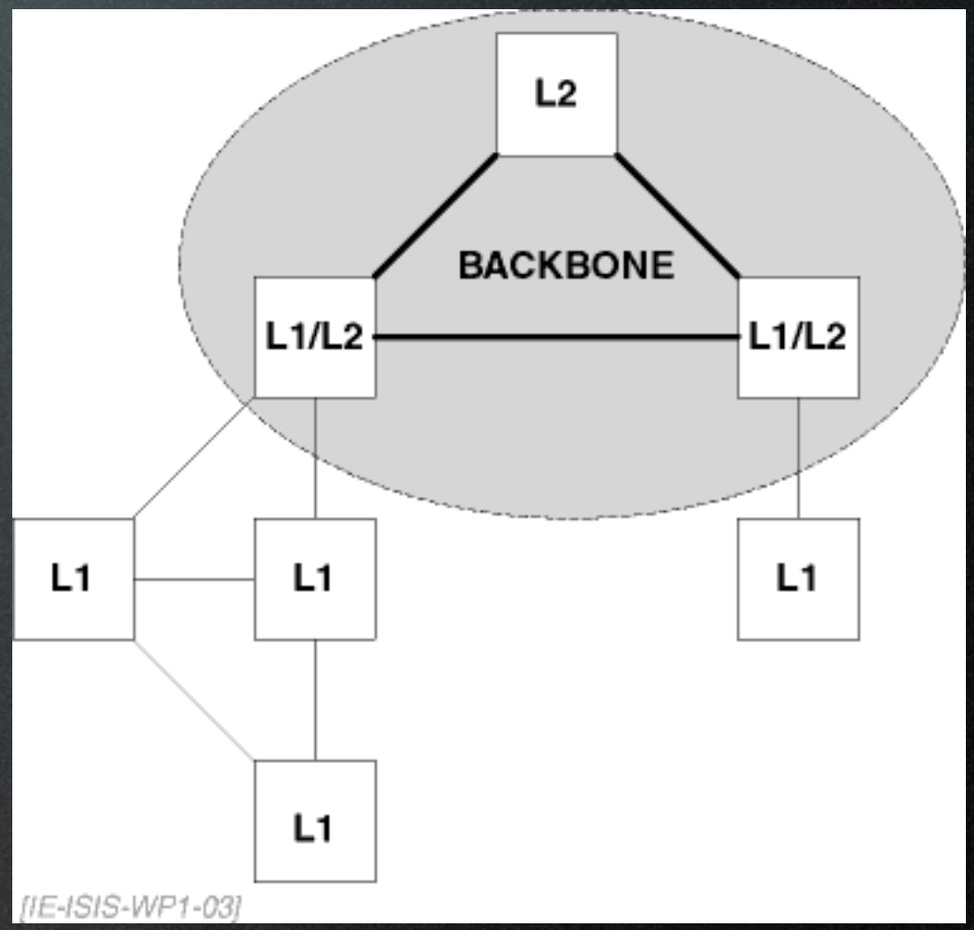
L2
BACKBONE
L1/L2
L1/L2
L1
L1
L1
L1
[IE-ISIS-WP1-03]

Border Gateway Protocol
BGP

- path vector protocol
- table of IP networks/prefixes
- CIDR
- route aggregation
- most ISPs **must** use BGP
- multihoming:
  - single ISP (multiple access points)
  - multiple ISPs

BGP

- manually configured neighbors/peers
- TCP
- 60sec keep alive messages
- hold time before removing peer
- incremental updates
- full table exchange only when (re)establish peering relationship

BGP

- IBGP (internal/interior BGP inside an AS)
- EBGP (between ASs)

- Border/Edge routers (exchange information between ASs)
- Global routing table - over 300000 entries

BGP

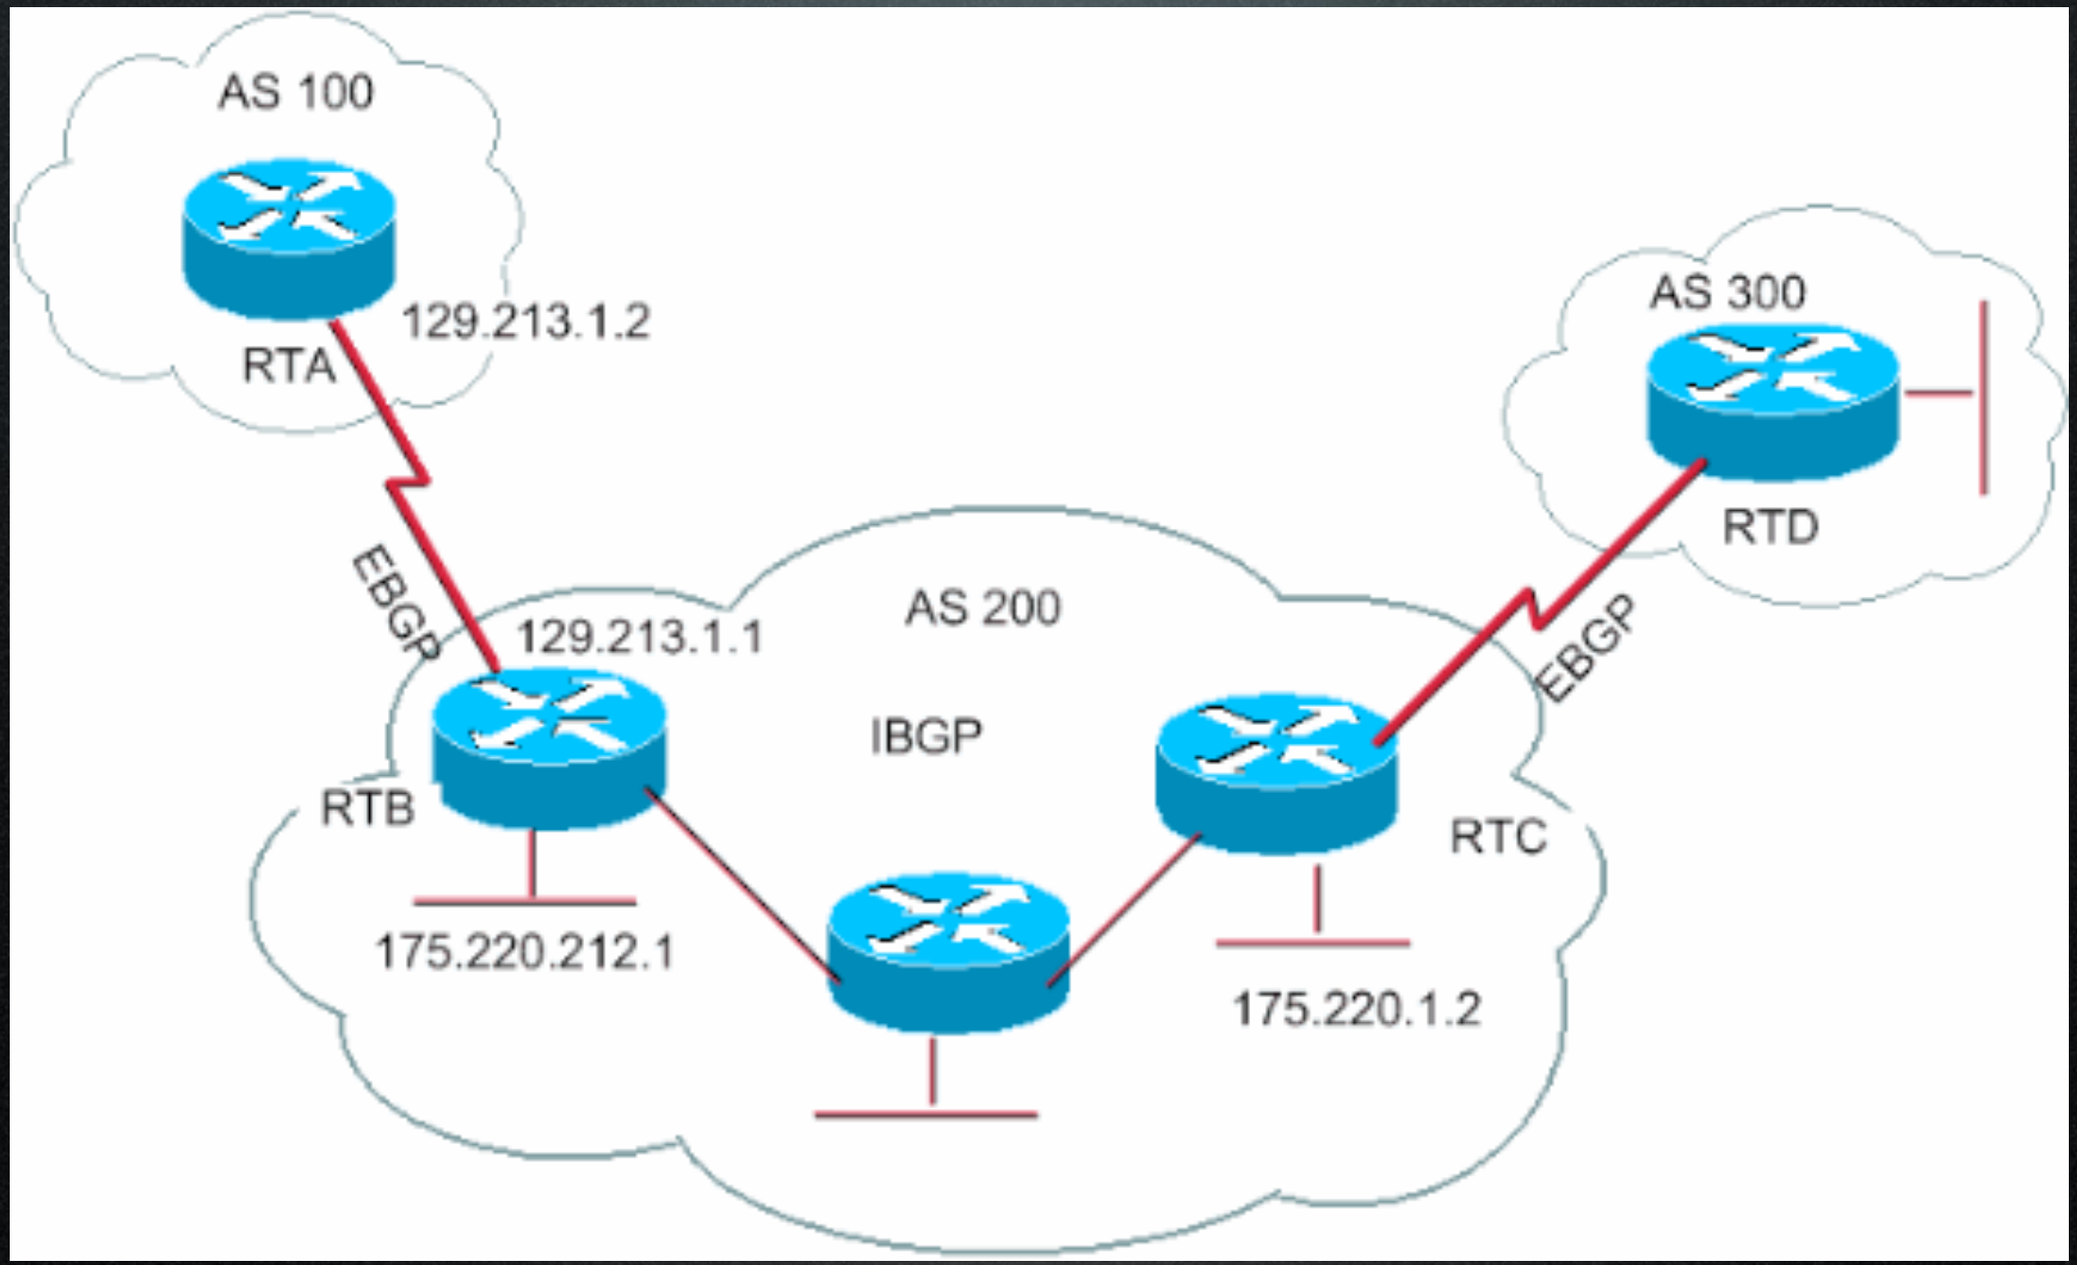
AS 100
RTA
129.213.1.2
EBGP
AS 300
RTD
EBGP
AS 200
129.213.1.1
IBGP
RTB
RTC
175.220.212.1
175.220.1.2

BGP

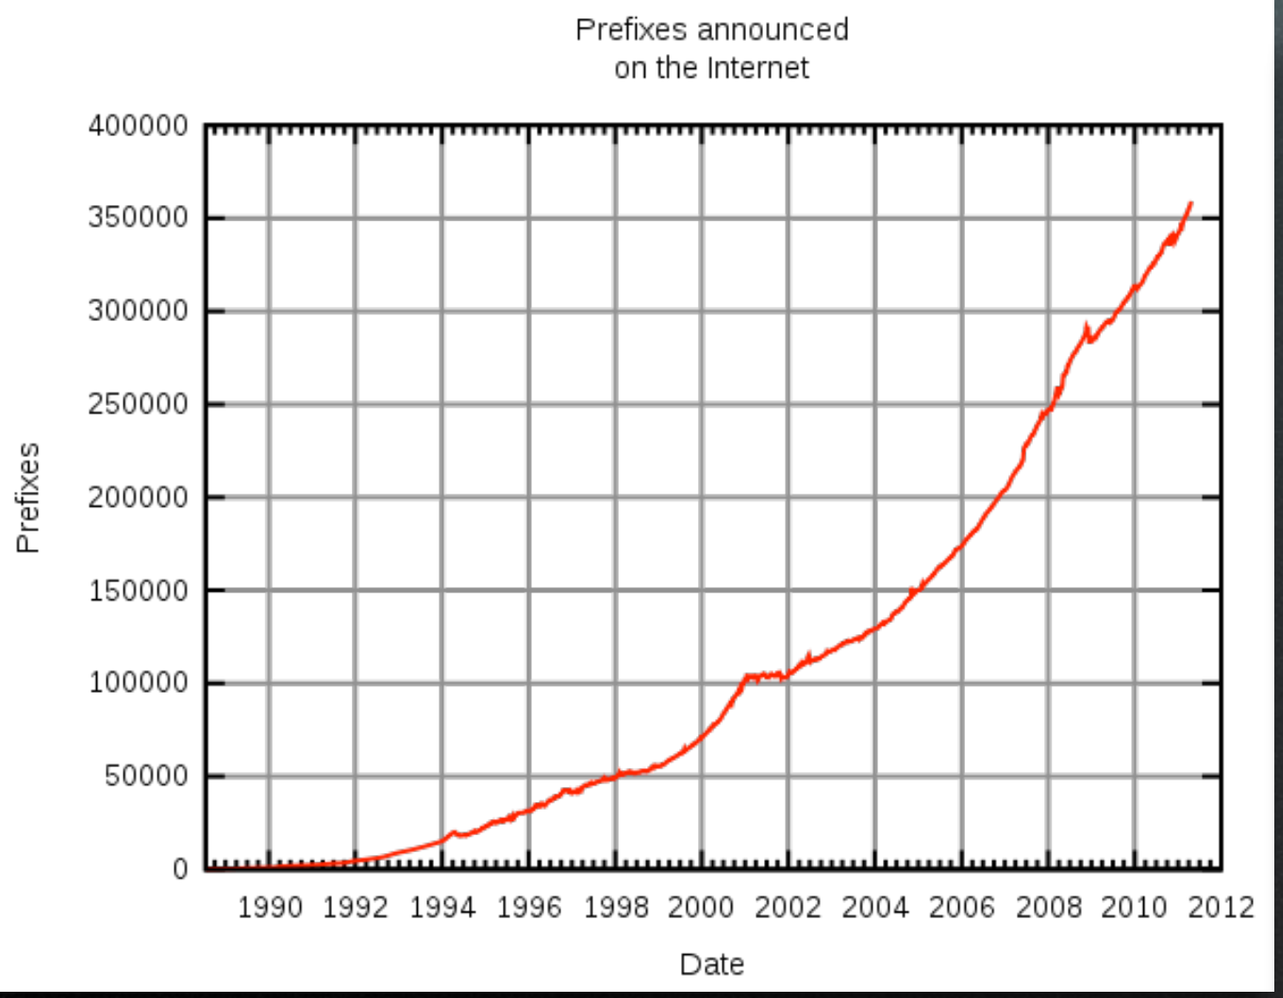
Prefixes announced
on the Internet
Prefixes
400000
350000
300000
250000
200000
150000
100000
50000
0
1990 1992 1994 1996 1998 2000 2002 2004 2006 2008 2010 2012
Date

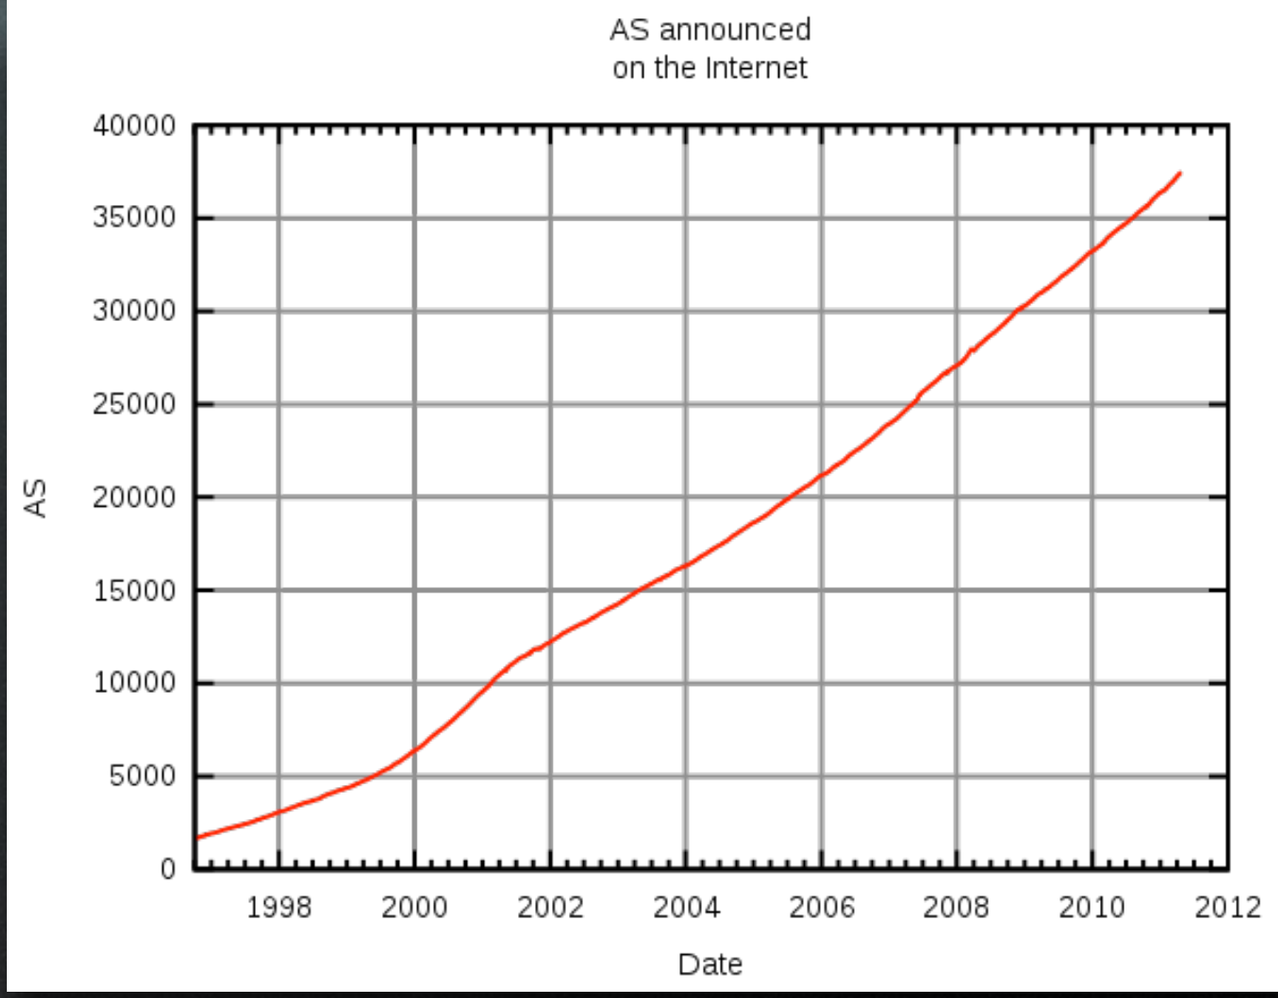
AS announced
on the Internet
AS
40000
35000
30000
25000
20000
15000
10000
5000
0
1998 2000 2002 2004 2006 2008 2010 2012
Date

BGP multihomming

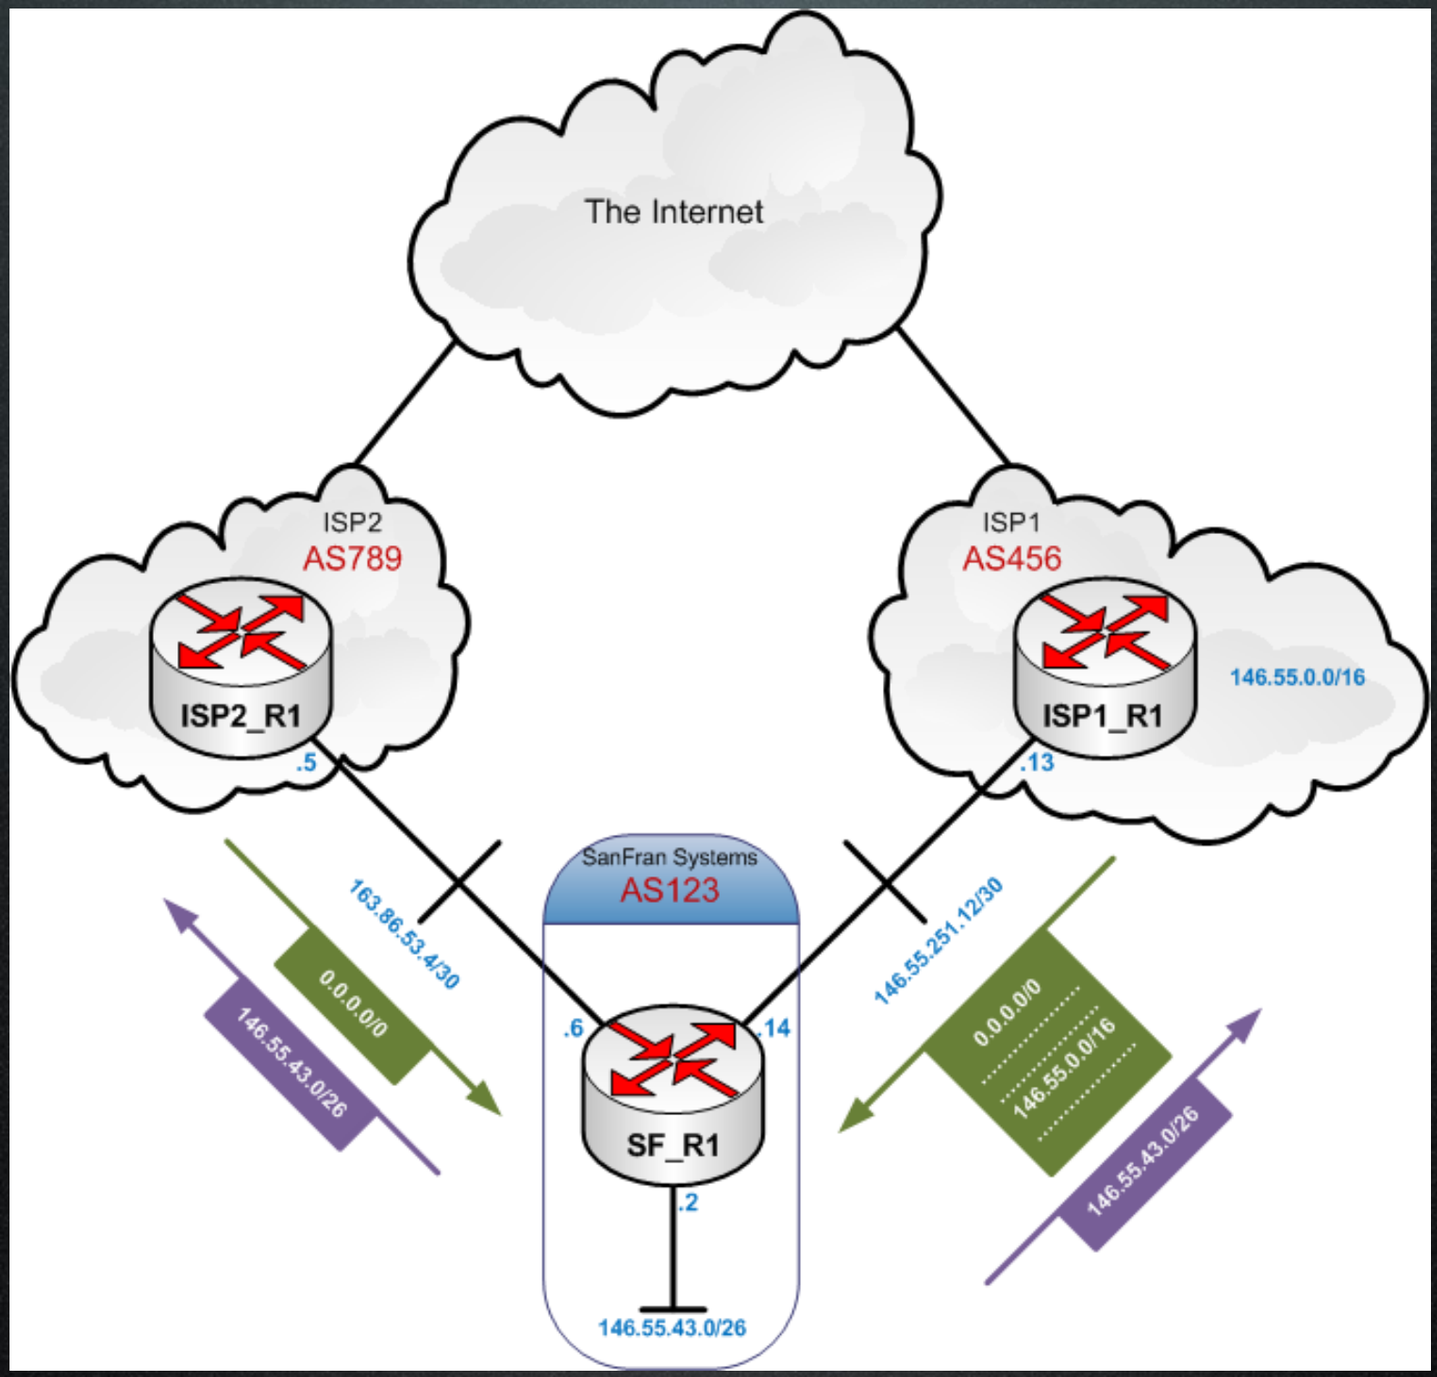

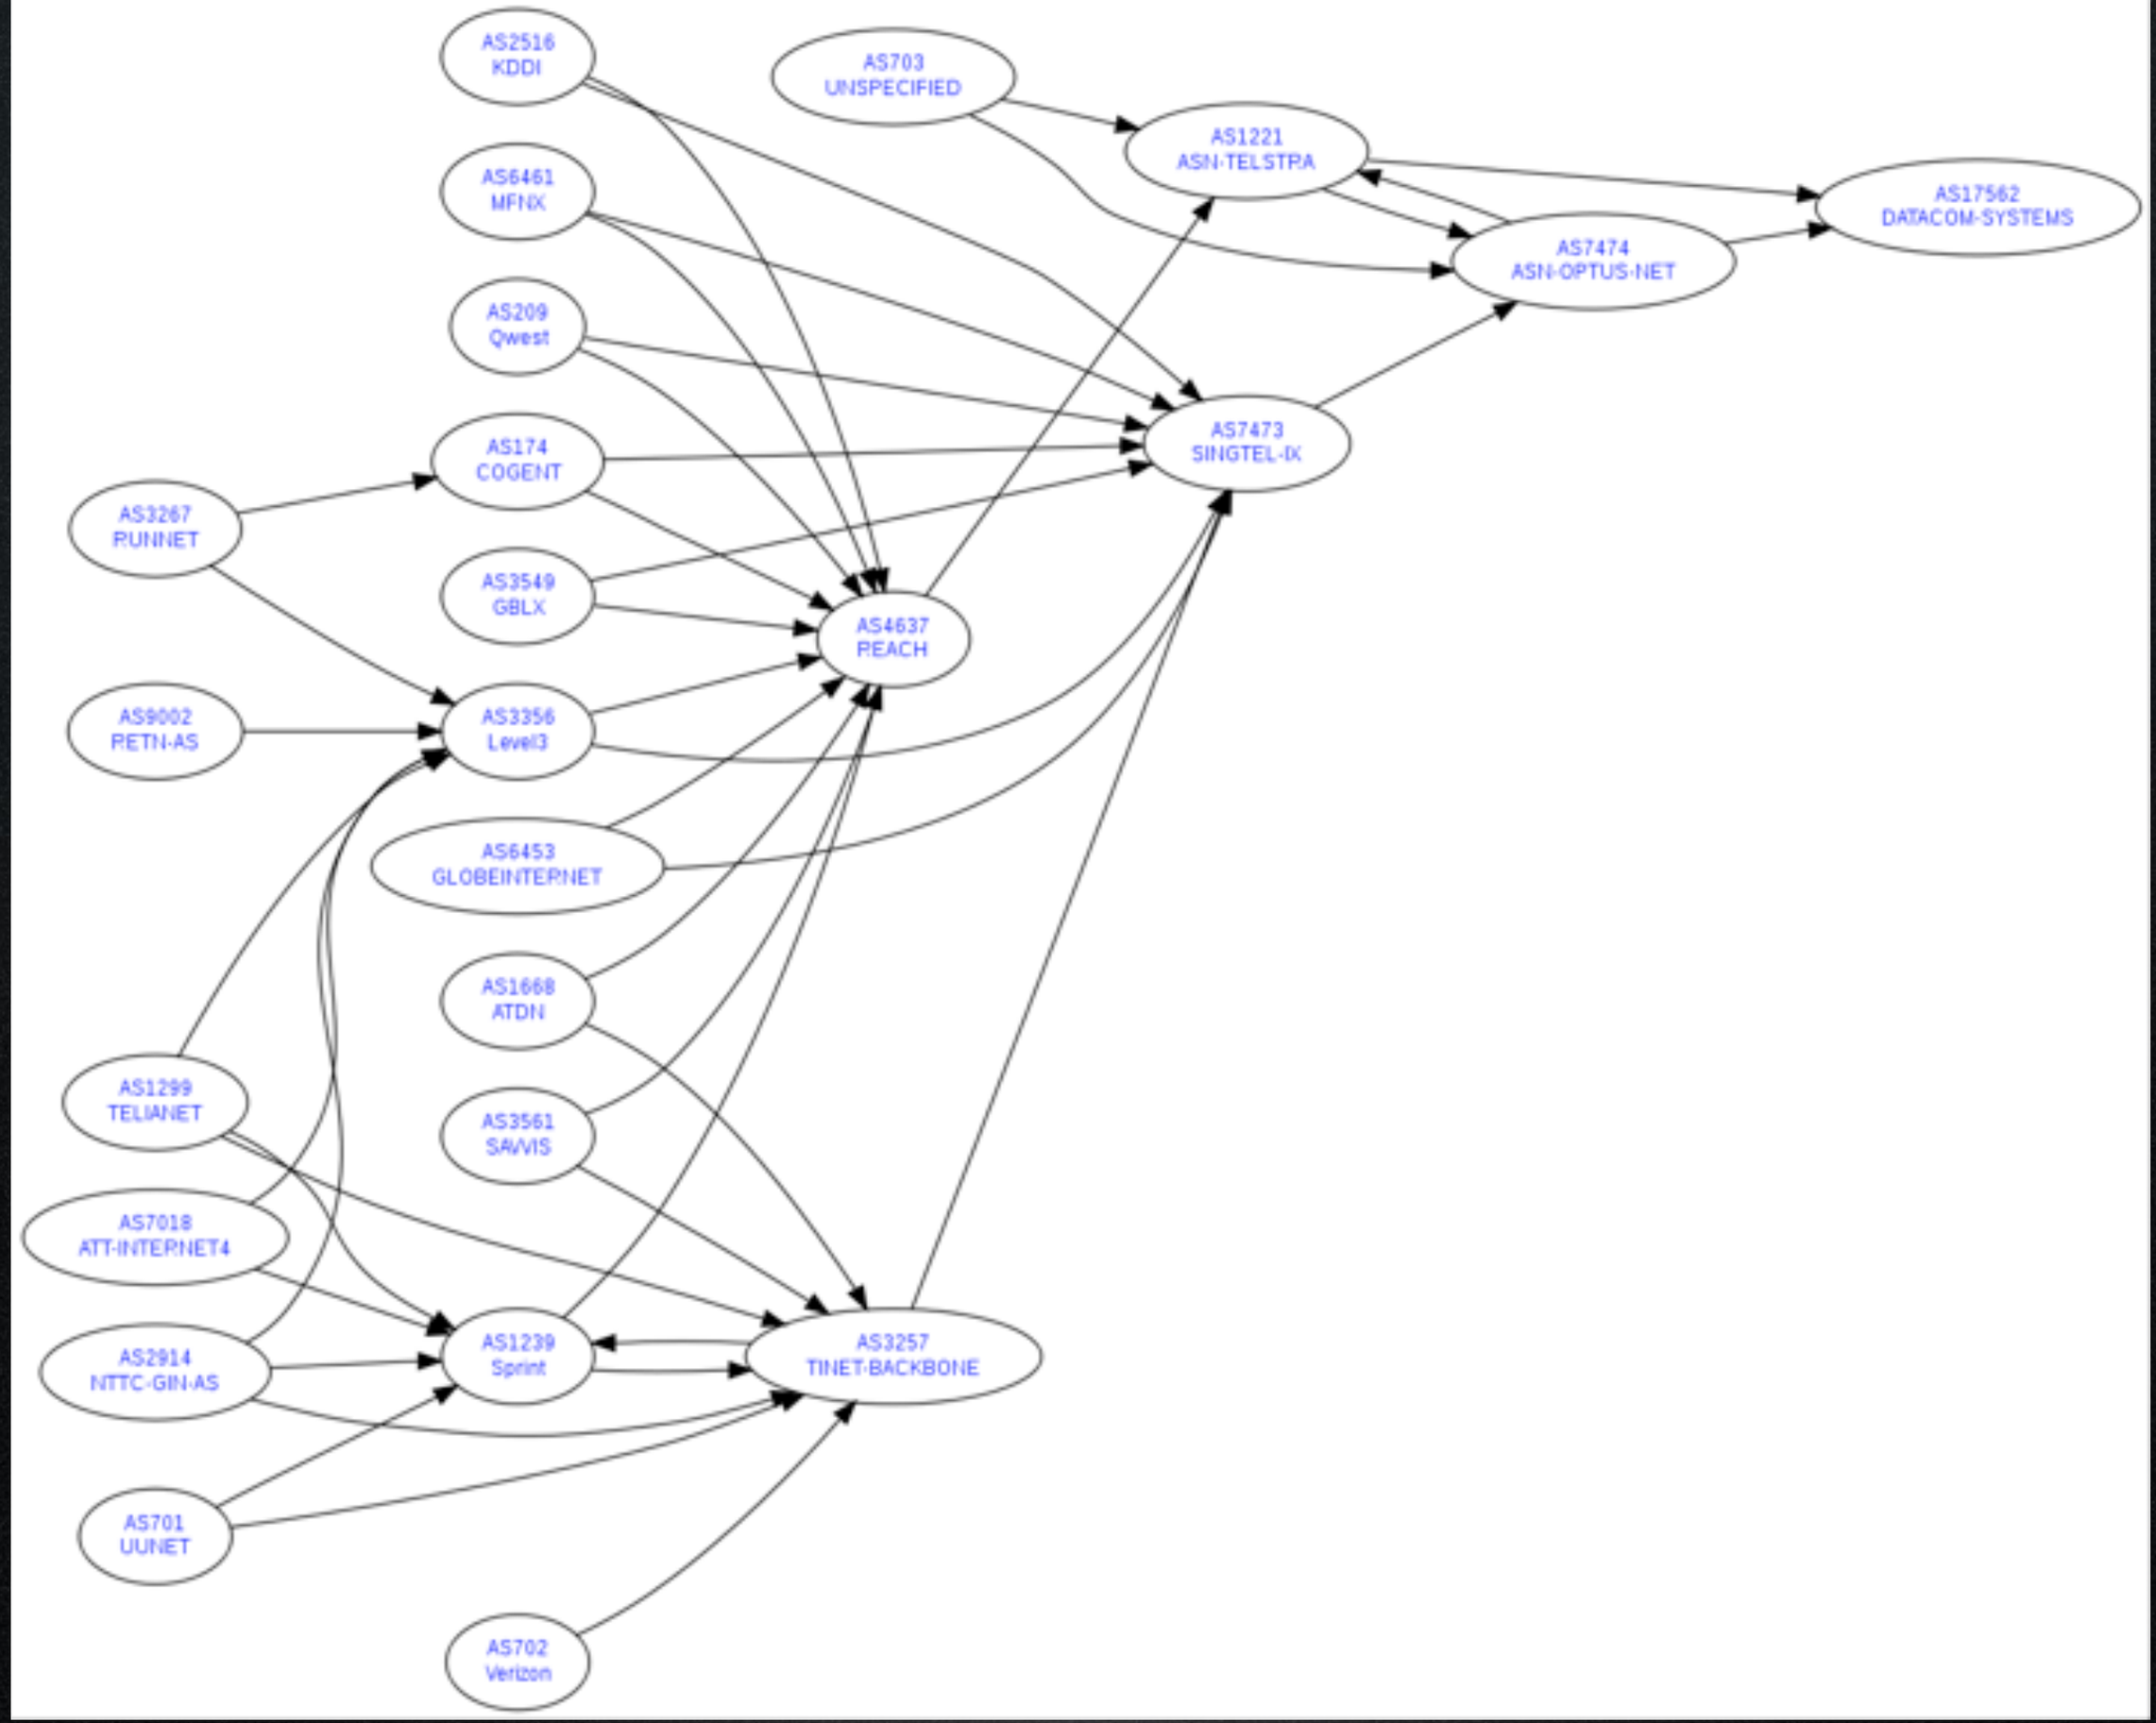
AS2516
KDDI
AS703
UNSPECIFIED
AS1221
ASN-TELSTRA
AS17562
DATACOM-SYSTEMS
AS6461
MFNX
AS7474
ASN-OPTUS-NET
AS209
Qwest
AS7473
SINGTEL-IX
AS174
COGENT
AS3267
RUNNET
AS3549
GBLX
AS4637
REACH
AS9002
RETN-AS
AS3356
Level3
AS6453
GLOBEINTERNET
AS1668
ATDN
AS1299
TELIANET
AS3561
SAVVIS
AS7018
ATT-INTERNET4
AS2914
NTTC-GIN-AS
AS1239
Sprint
AS3257
TINET-BACKBONE
AS701
UUNET
AS702
Verizon