APNIC eLearning:
BGP Basics

Contact: training@apnic.net

APNIC

Overview

- What is BGP?
- BGP Features
- Path Vector Routing Protocol
- Peering and Transit
- BGP General Operation
- BGP Terminology
- BGP Attributes
- Inserting Prefixes into BGP

What is BGP?

- Border Gateway Protocol
- A Routing Protocol used to exchange routing information between different networks
  - Exterior gateway protocol
- Described in RFC4271
  - RFC4276 gives an implementation report on BGP
  - RFC4277 describes operational experiences using BGP

BGP Features

- Path Vector Protocol
- Incremental Updates
- Many options for policy enforcement
- Classless Inter Domain Routing (CIDR)
- Widely used for Internet backbone
- Autonomous systems

What is Path Vector Routing Protocol

- A path vector routing protocol is used to span different autonomous systems
- It defines a route as a collection of a number of AS that it passes through from source AS to destination AS
- This list of ASes are called AS path and used to avoid routing loop
- AS path is also used to select path to destination
- RFC 1322
  - "A path vector protocol defines a route as a pairing between a destination and the attributes of the path to that destination."

Path Vector Protocol

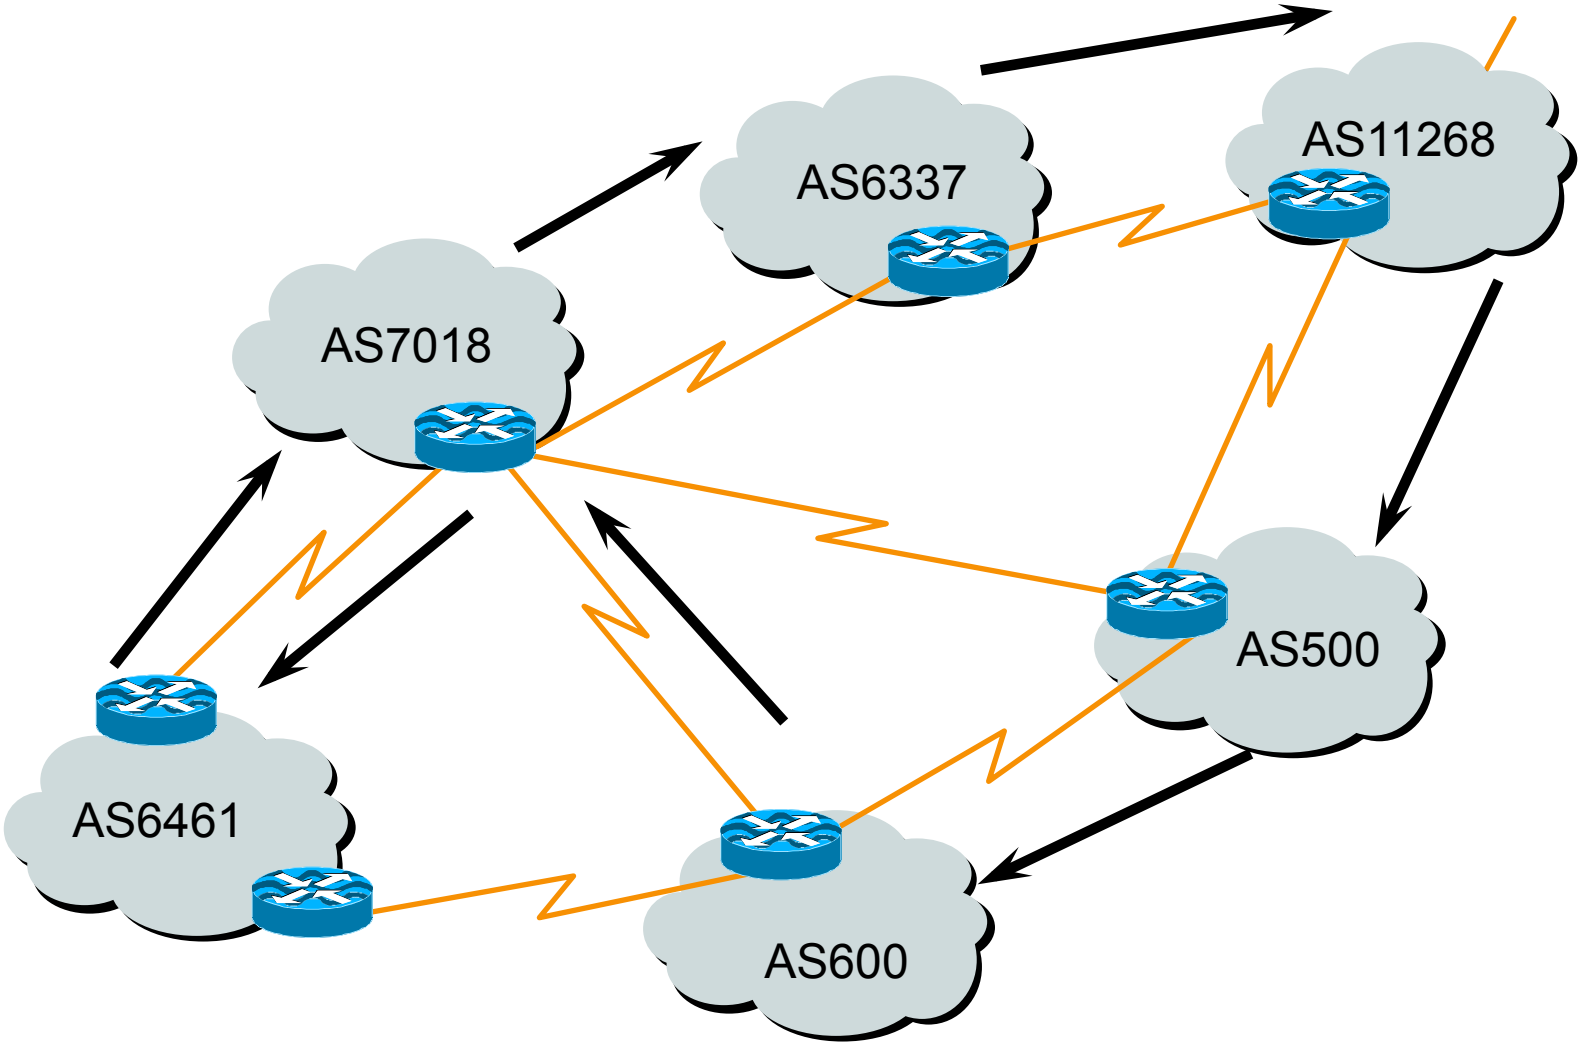

Definitions

- Transit – carrying traffic across a network, usually for a fee
- Peering – exchanging routing information and traffic
- Default – where to send traffic when there is no explicit match in the routing table

Default Free Zone

The default free zone is made up of Internet routers which have explicit routing information about the rest of the Internet, and therefore do not need to use a default route

Peering and Transit example

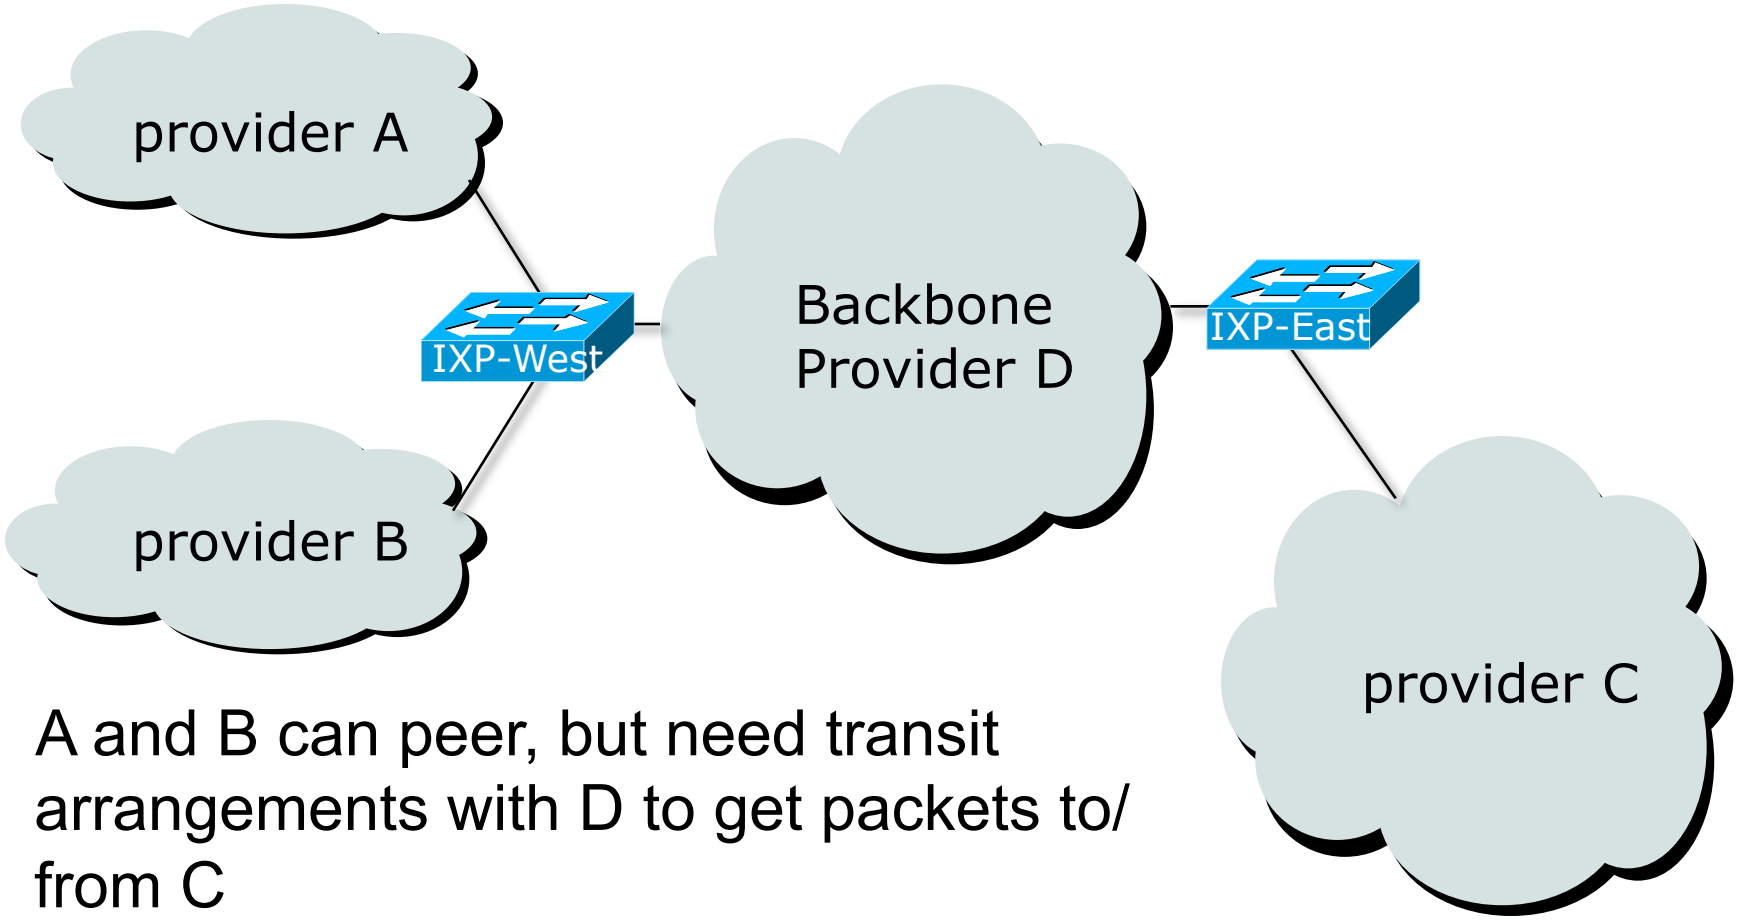

- A and B can peer, but need transit arrangements with D to get packets to/ from C

What Is An Autonomous System?

- Group of Internet Protocol-based networks with the same routing policy
  - Usually under single ownership, trust or administrative control
- The AS is used both in the exchange of exterior routing information (between neighboring ASes) and as an identifier of the AS itself
- The Autonomous System is the cornerstone of BGP
  - It is used to uniquely identify networks with a common routing policy

Autonomous System Number (ASN)

- globally unique identifiers for IP networks
- ASN uniquely identifies each network on the Internet
- allocated to each Autonomous System (AS) for use in BGP routing
- 2-byte only AS number range : 0 – 65535
- 4-byte only AS number range – represented in two ways
  - AS PLAIN: 65,536 - 4,294,967,295
  - AS DOT: 1.0 - 65535.65535

BGP General Operation

- Learns multiple paths via internal and external BGP speakers
- Picks the best path and installs it in the routing table (RIB)
- Best path is sent to external BGP neighbours
- Policies are applied by influencing the best path selection

BGP Basics

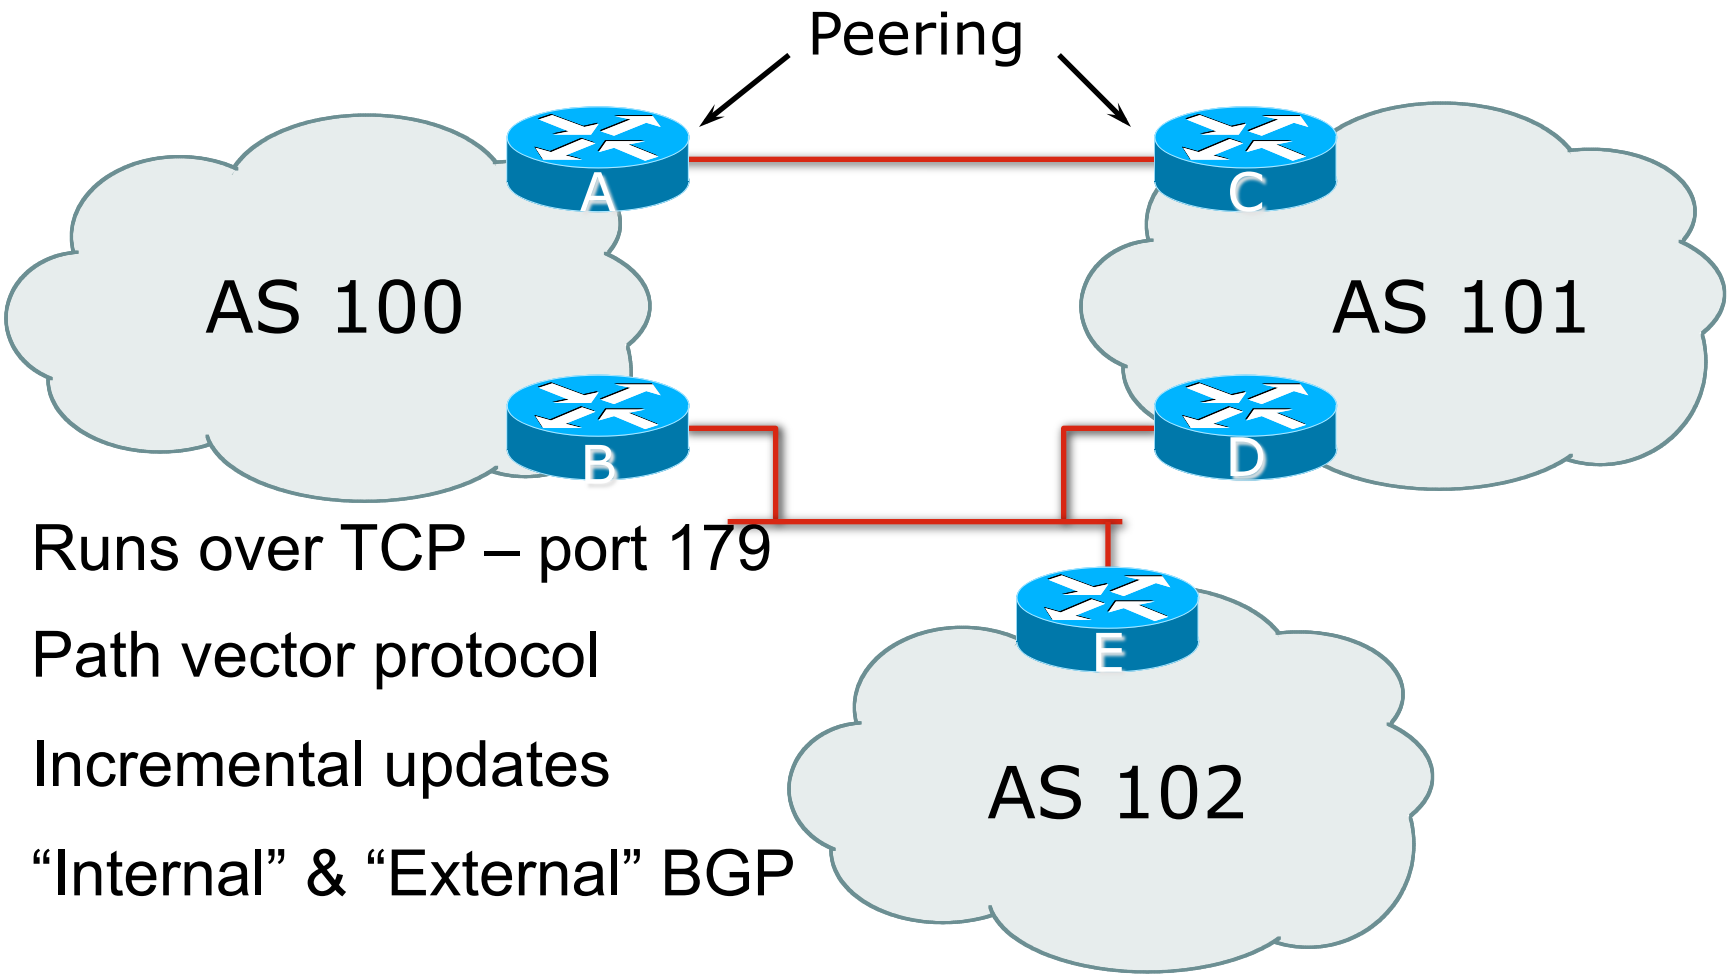

- Runs over TCP – port 179
- Path vector protocol
- Incremental updates
- “Internal” & “External” BGP

BGP Terminology

- Neighbor
  - Any two routers that have formed a TCP connection to exchange BGP routing information are called peers or neighbors
- iBGP
  - iBGP refers to the BGP neighbor relationship within the same AS.
  - The neighbors do not have to be directly connected.
- eBGP
  - When BGP neighbor relationship are formed between two peers belongs to different AS are called eBGP.
  - EBGP neighbors by default need to be directly connected.

BGP Attributes

- Well-known attributes – must be supported by every BGP implementation
- Mandatory attributes – must be included with every route entry. If one attribute is missing, it will result in an error message
  - Ex: ORIGIN, AS_PATH, NEXT_HOP, LOCAL_PREF
- Discretionary attributes – every BGP router must recognize, but they don't have to be present with every route entry
  - Ex. ATOMIC_AGGREGATE
- Optional attributes – not necessarily supported by all BGP implementations. It can be either transitive or non-transitive.
  - AGGREGATOR, COMMUNITY, MULTI_EXIT_DISC

BGP/IGP model used in ISP networks

- BGP is used internally (iBGP) and externally (eBGP)
- iBGP – used to carry some/all Internet prefixes across ISP backbone and ISP's customer prefixes
- eBGP – used to exchange prefixes with other ASes and mplement routing policy

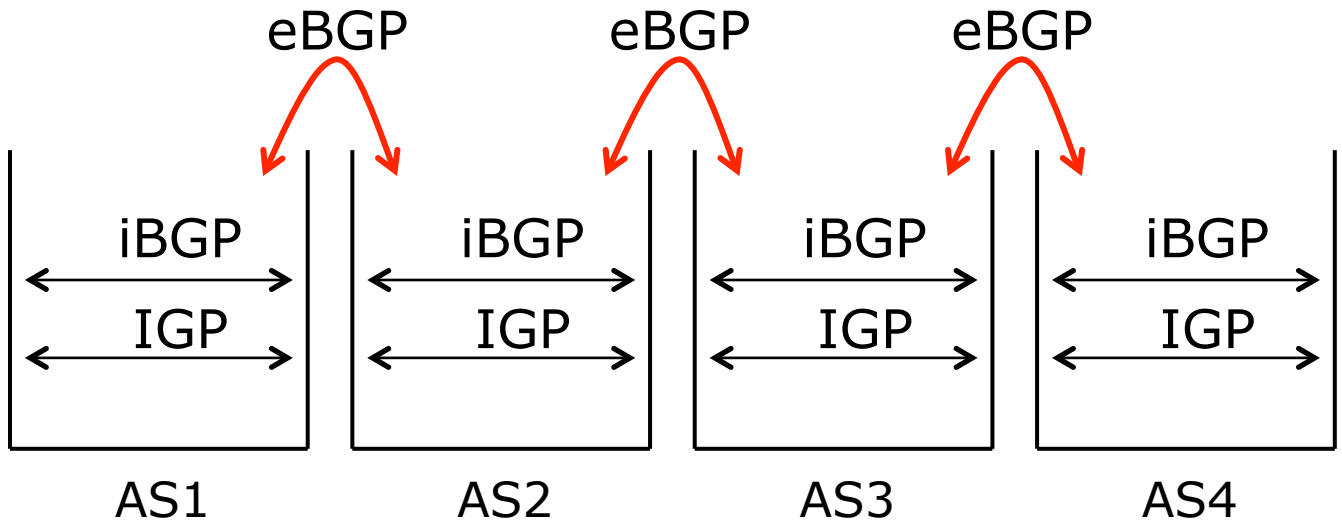

Internal BGP (iBGP)

- BGP peer within the same AS
- Not required to be directly connected
  - IGP takes care of inter-BGP speaker connectivity
- iBGP speakers must be fully meshed:
  - They originate connected networks
  - They pass on prefixes learned from outside the ASN
  - They do not pass on prefixes learned from other iBGP speakers

Internal BGP Peering (iBGP)

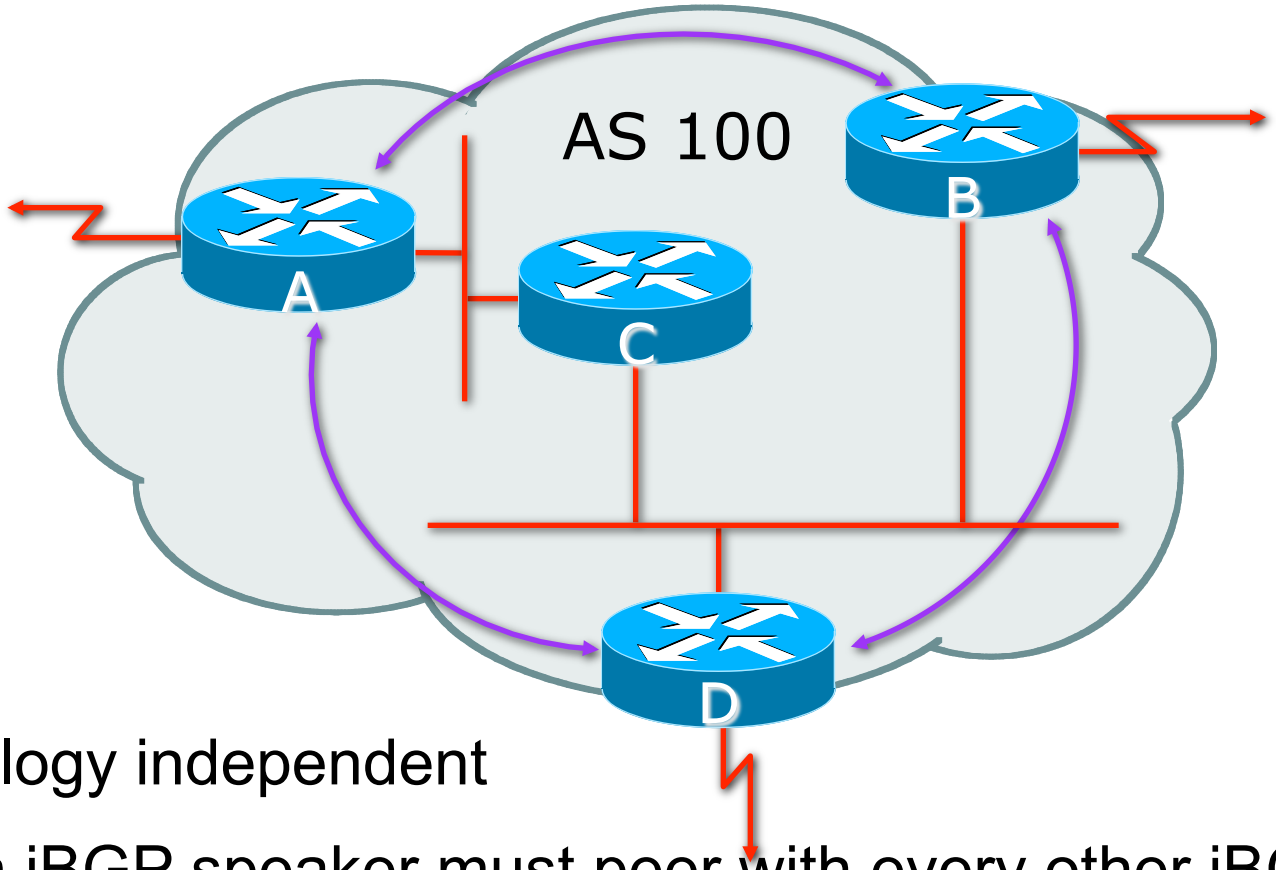

- Topology independent
- Each iBGP speaker must peer with every other iBGP speaker in the AS

Peering between Loopback Interfaces

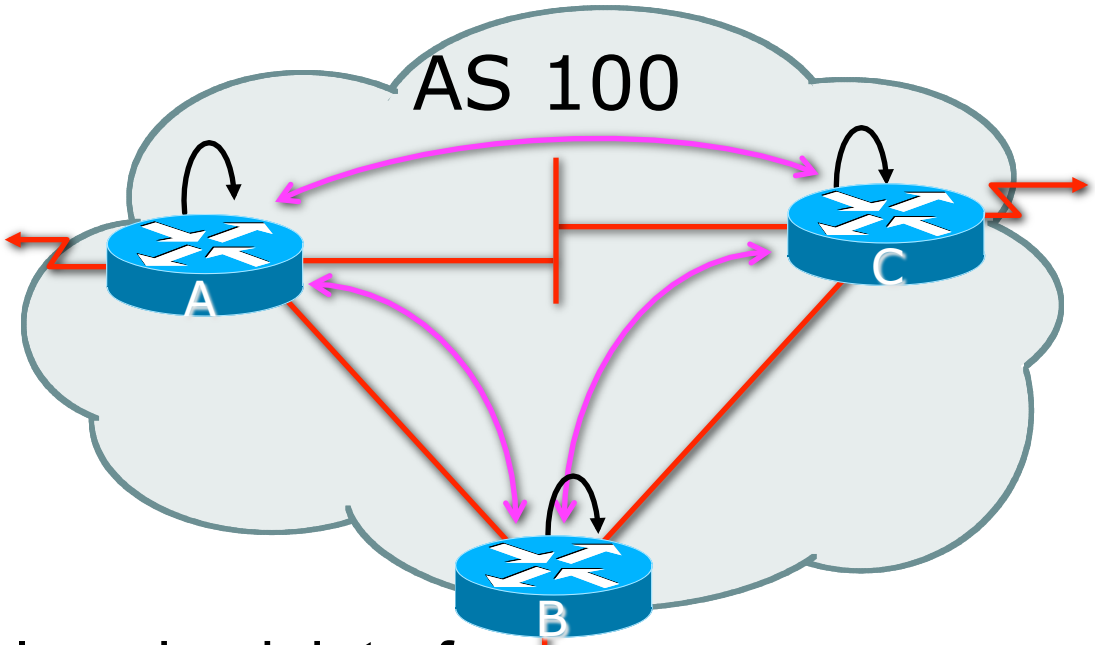

- Peer with loop-back interface
  - Loop-back interface does not go down – ever!
- Do not want iBGP session to depend on state of a single interface or the physical topology

Constructing the Forwarding Table

- BGP “in” process
  - receives path information from peers
  - results of BGP path selection placed in the BGP table
  - “best path” flagged
- BGP “out” process
  - announces “best path” information to peers
- Best path stored in Routing Table (RIB)
- Best paths in the RIB are installed in forwarding table (FIB) if:
  - prefix and prefix length are unique
  - lowest “protocol distance”

Constructing the Forwarding Table

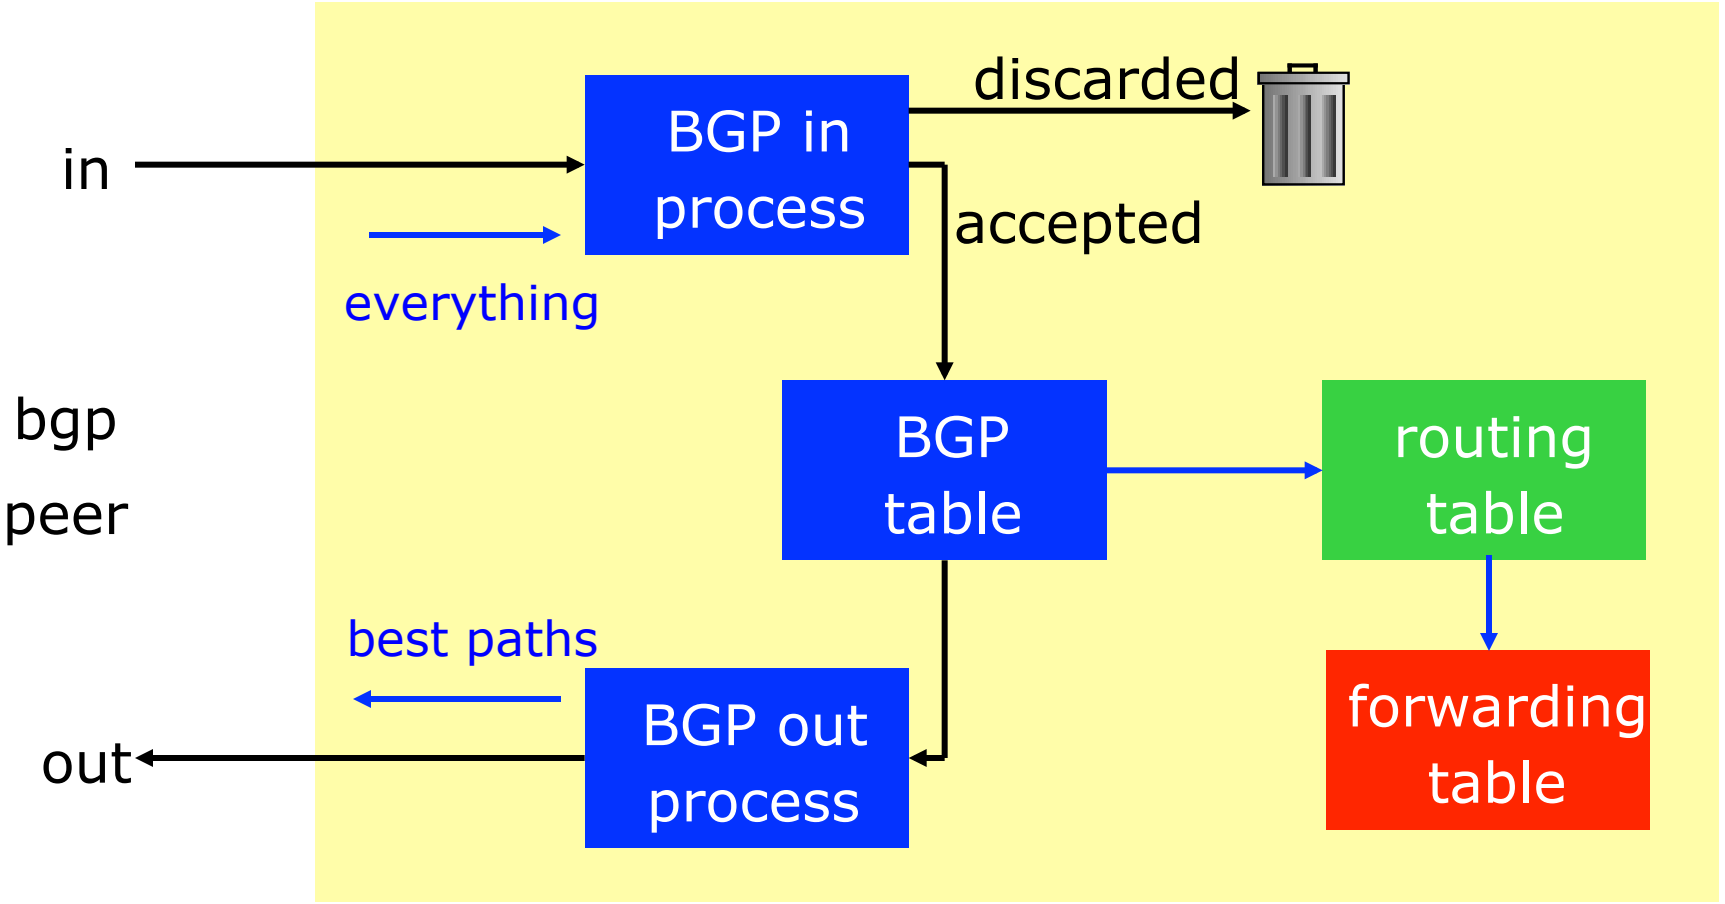

External BGP Peering (eBGP)

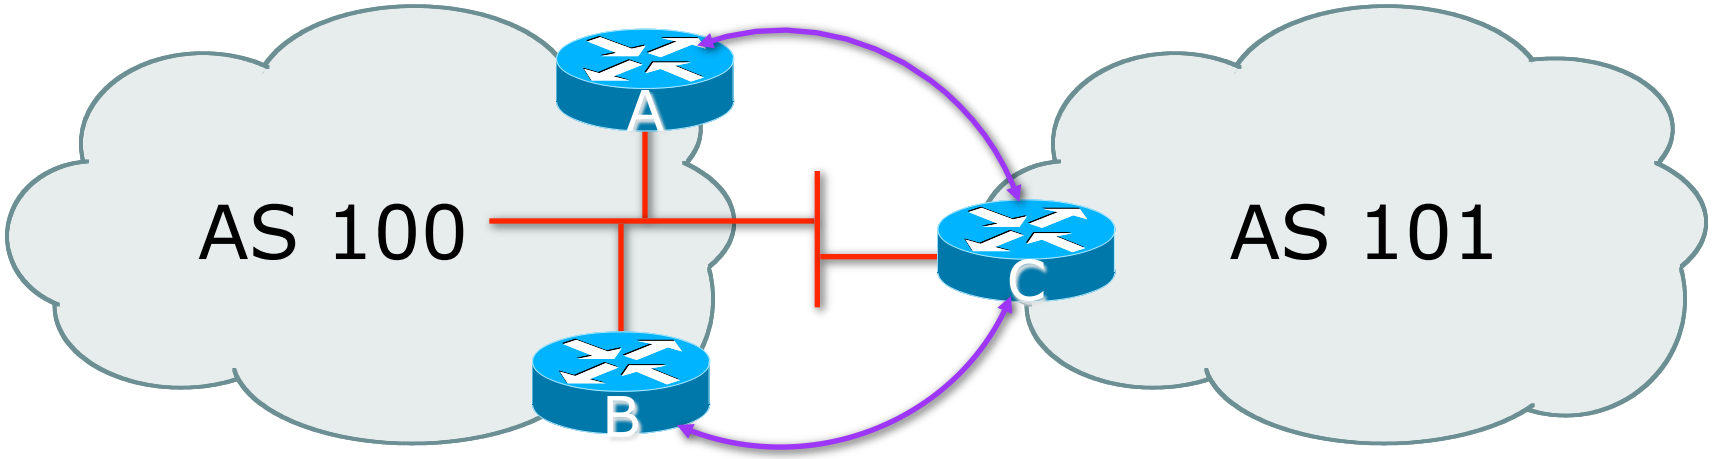

- Between BGP speakers in different AS
- Should be directly connected
- Never run an IGP between eBGP peers

Configuring BGP in Cisco IOS

- This command enables BGP in Cisco IOS:

```
router bgp 100
```

- For ASNs > 65535, the AS number can be entered in either plain notation, or in dot notation:

```
router bgp 131076
```

or

```
router bgp 2.4
```

- IOS will display ASNs in plain notation by default
  - Dot notation is optional:

```
router bgp 2.4
bgp asnotation dot
```

Configuring External BGP

Router A in AS100

```
interface ethernet 5/0
 ip address 102.102.10.2 255.255.255.240
!
router bgp 100
 network 100.100.8.0 mask 255.255.252.0
 neighbor 102.102.10.1 remote-as 101
 neighbor 102.102.10.1 prefix-list RouterC in
 neighbor 102.102.10.1 prefix-list RouterC out
!
```

ip address on ethernet interface

Local ASN

Remote ASN

ip address of Router C ethernet interface

Inbound and outbound filters

Configuring External BGP

Router C in AS101

```
interface ethernet 1/0/0
 ip address 102.102.10.1 255.255.255.240
!
router bgp 101
 network 100.100.8.0 mask 255.255.252.0
 neighbor 102.102.10.2 remote-as 100
 neighbor 102.102.10.2 prefix-list RouterA in
 neighbor 102.102.10.2 prefix-list RouterA out
!
```

ip address on ethernet interface

Local ASN

Remote ASN

ip address of Router A ethernet interface

Inbound and outbound filters

Configuring Internal BGP

Router A in AS100

```
interface loopback 0
 ip address 105.3.7.1 255.255.255.255
!
router bgp 100
 network 100.100.1.0
 neighbor 105.3.7.2 remote-as 100
 neighbor 105.3.7.2 update-source loopback0
 neighbor 105.3.7.3 remote-as 100
 neighbor 105.3.7.3 update-source loopback0
!
```

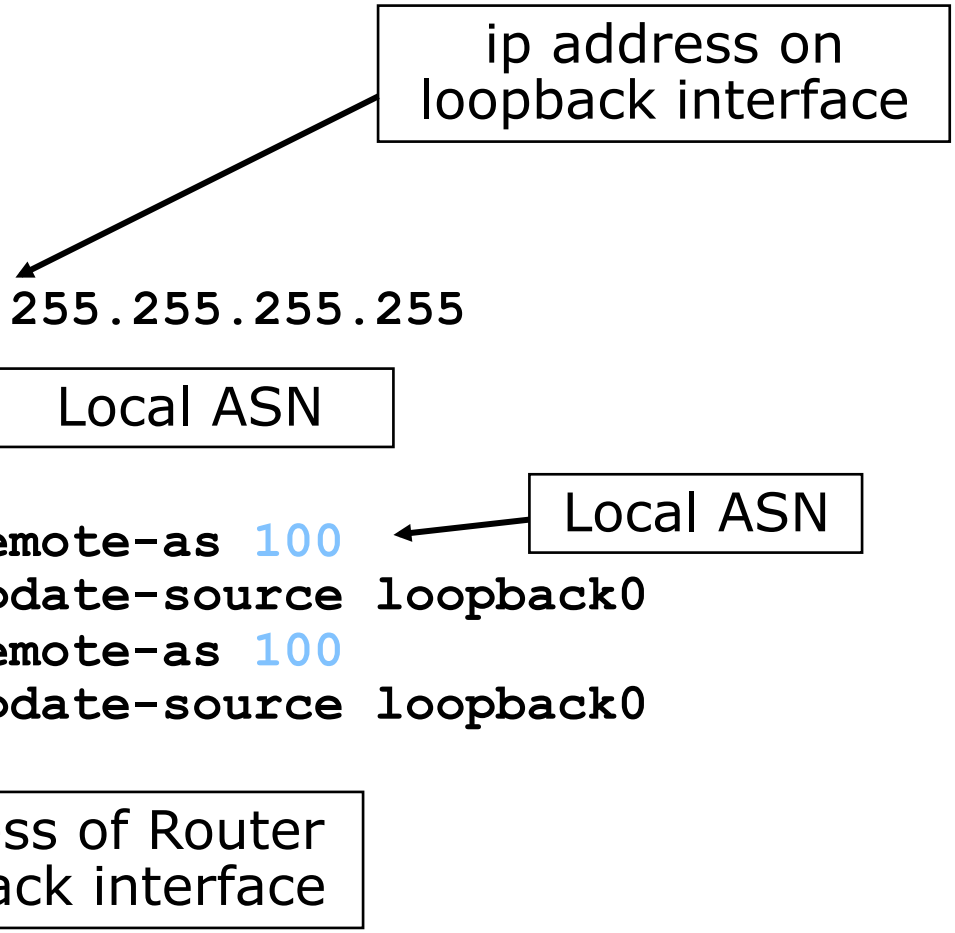

Configuring Internal BGP

Router B in AS100

```
interface loopback 0
 ip address 105.3.7.2 255.255.255.255
!
router bgp 100
 network 100.100.1.0
 neighbor 105.3.7.1 remote-as 100
 neighbor 105.3.7.1 update-source loopback0
 neighbor 105.3.7.3 remote-as 100
 neighbor 105.3.7.3 update-source loopback0
!
```

ip address on loopback interface

Local ASN

Local ASN

ip address of Router A loopback interface

Inserting prefixes into BGP – network command

- Configuration Example
  - router bgp 100
  - network 102.10.32.0 mask 255.255.254.0
  - ip route 102.10.32.0 255.255.254.0 serial0
- A matching route must exist in the routing table before the network is announced
- Forces origin to be “IGP”

Configuring Aggregation – Network Command

- Configuration Example
  - router bgp 100
  - network 102.10.0.0 mask 255.255.0.0
  - ip route 102.10.0.0 255.255.0.0 null0 250
- A matching route must exist in the routing table before the network is announced
- Easiest and best way of generating an aggregate

Summary
BGP neighbour status

```
Router>sh ip bgp sum

BGP router identifier 10.0.15.246, local AS number 10
BGP table version is 16, main routing table version 16
7 network entries using 819 bytes of memory
14 path entries using 728 bytes of memory
2/1 BGP path/bestpath attribute entries using 248 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1795 total bytes of memory
BGP activity 7/0 prefixes, 14/0 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down State/
PfxRcd
10.0.15.241     4    10       9       8       16    0    0 00:04:47        2
10.0.15.242     4    10       6       5       16    0    0 00:01:43        2
10.0.15.243     4    10       9       8       16    0    0 00:04:49        2
...
```

BGP Version

Updates sent and received

Updates waiting

Summary
BGP Table

```
Route6>sh ip bgp

BGP table version is 30, local router ID is 10.0.15.246
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i10.0.0.0/26      10.0.15.241              0    100      0 i
*>i10.0.0.64/26     10.0.15.242              0    100      0 i
*>i10.0.0.128/26    10.0.15.243              0    100      0 i
*>i10.0.0.192/26    10.0.15.244              0    100      0 i
*>i10.0.1.0/26      10.0.15.245              0    100      0 i
*> 10.0.1.64/26     0.0.0.0                  0         32768 i
*>i10.0.1.128/26    10.0.15.247              0    100      0 i
*>i10.0.1.192/26    10.0.15.248              0    100      0 i
...
```

Questions

- Please remember to fill out the feedback form
  - <survey-link>
- Slide handouts will be available after completing the survey

APNIC Helpdesk Chat

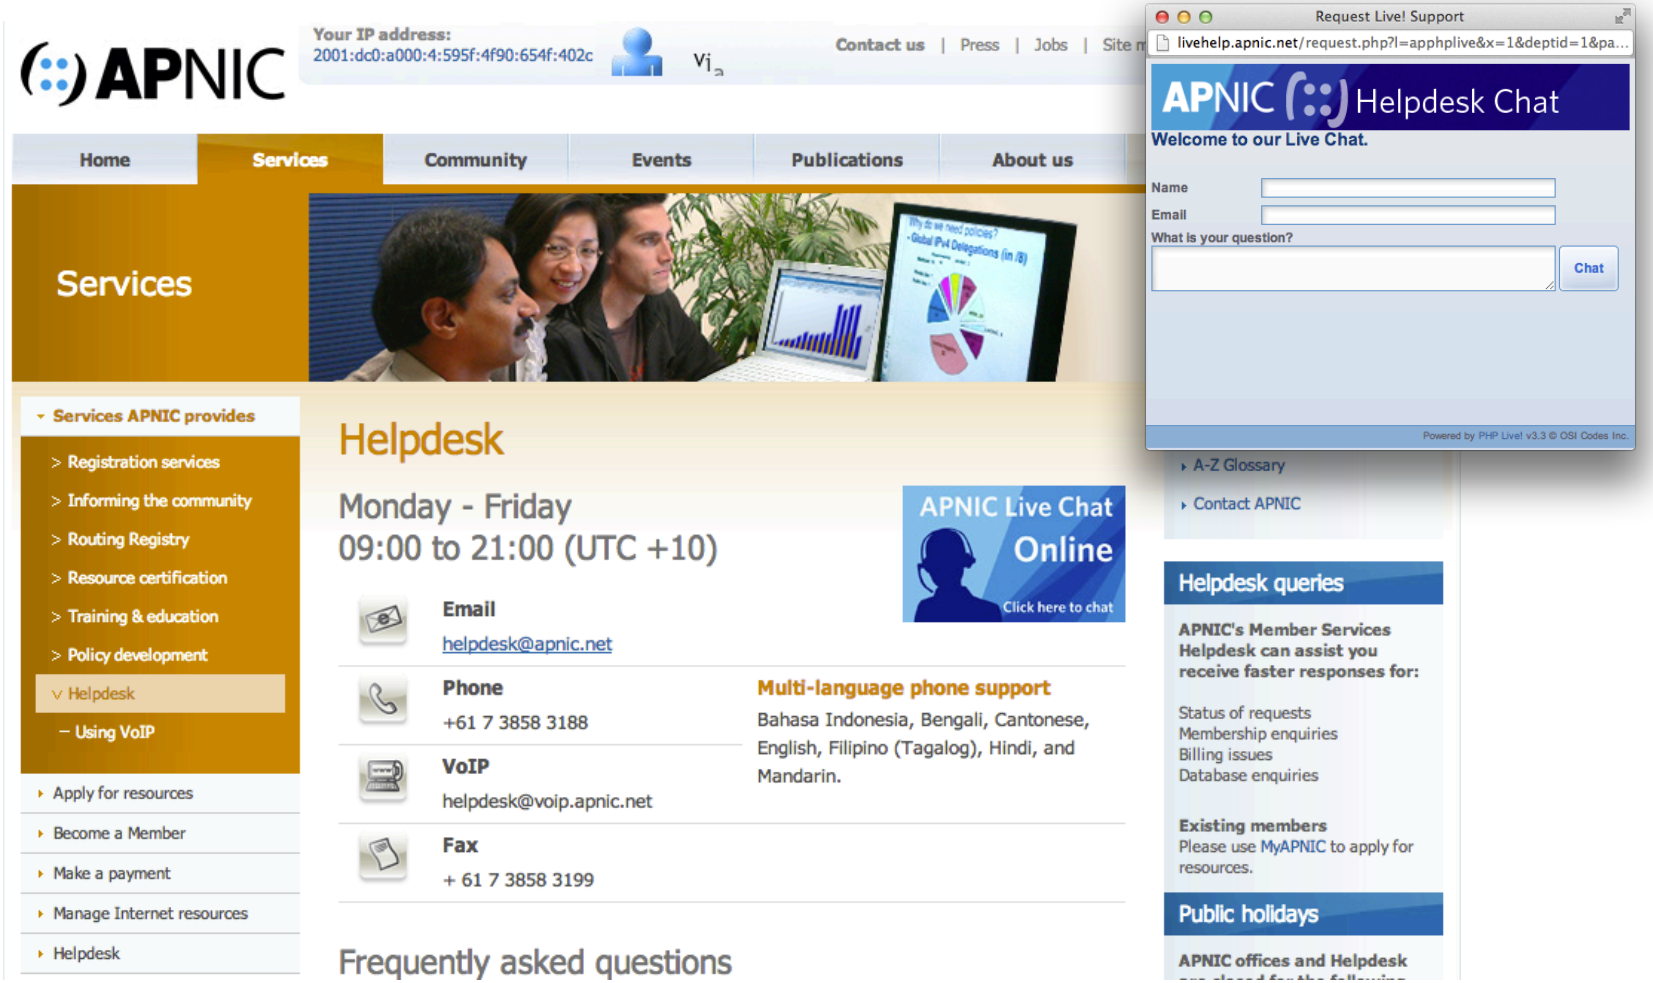

Thank you!

End of Session

APNIC