Introduction to Routing Security

Practical Cybersecurity for Internet Operators (PCIO)

Routing Security

- Routing Security best practices have been considered vital for the Internet's operation since the mid-1990s
- Now has a catchy title (Mutually Agreed Norms for Routing Security aka MANRS), but the fundamentals are the same
- Implement the MANRS recommendations:
  - Prevent propagation of incorrect routing information
  - Prevent traffic with spoofed source addresses
  - Facilitate communication between network operators
  - Facilitate validation of routing information

MANRS

Routing Security

- Prevent propagation of incorrect routing information
  - Make sure your clients are only announcing their address space to you
    - Verify that it is indeed their address space and their ASN
    - If it is not listed in the RIR databases as theirs, do NOT route it
  - Make sure you are only announcing your and your clients' address space to upstreams and peers
  - Make sure you only accept prefixes you expect from your peers
  - Make sure you do not allow unassigned or reserved addresses and ASNs
- This means:
  - Filter ALL EBGP sessions
  - Be aware: most BGP implementations are permissive by default

Routing Security

- Prevent traffic with spoofed source addresses
  - We have had a tool for this since the late 1990s – it is called “Unicast Reverse Path Forwarding” and is implemented in the hardware of router platforms
    - Do NOT buy access routers that cannot do uRPF regardless of what the vendor tells you
    - uRPF must be implemented on all edge facing interfaces, whether access customers, datacentre customers, wireless customers,…
  - Implementing uRPF will go a long way to help global efforts to deal with the menace of denial-of-service attacks
  - How good is your network?
    - https://www.caida.org/projects/spoofer/

Routing Security

- Facilitate communication between network operators
  - Know how to get hold of the NOC of your peers and upstreams
  - Make sure you have a PeeringDB entry
    - If you have an ASN, create an entry, even if you aren't at a peering point or a data centre – it is visibility!
  - Make sure your organisation's technical contact details are up to date for your IP address and ASN holdings at your RIR (Route and AS-Objects)
  - Take part in your local network operator group/online network forum
    - Don't hide away, and then expect the world to help you when you have a problem

Routing Security

- Facilitate validation of Routing information
  - Route Origin Authorisation has been available and deployable since 2013
  - All entities with IP address holdings must sign their ROAs
    - The ROA is a cryptographically signed object linking the IP address being announced to the ASN originating it.
  - Network operators are widely encouraged to drop invalid ROAs
    - This is a route which has either the incorrect origin ASN or incorrect subnet size when compared with existing ROAs
    - Many operators drop invalids, more need to – this is called Route Origin Validation
  - This stops incorrect announcements from appearing and propagating across the Internet routing system

Practical Routing Security

- We will now look at:
  - Interior routing protocol best practices
  - Key components of BGP
  - Securing the Router
  - BGP Best Current Practices
  - uRPF
  - RTBH
  - RPKI, ROAs and ROV

ISIS vs OSPF vs EIGRP vs RIPv2

IGP BEST PRACTICES

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Making choices

- An Interior routing protocol (IGP) is used within a network so that routers in that network can find each other and distribute internal prefixes
- There are four IGPs:
  - ISIS
    - Widely used by Internet network operators, built-in multi-protocol support, the IGP of choice
  - OSPF
    - IETF standard routing protocol. OSPFv2 for IPv4, OSPFv3 for IPv6.
  - EIGRP
    - Proprietary to Cisco. Do not use unless vendor lock-in is your goal
  - RIPv2
    - Do not use, under any circumstances. Obsolete, does not scale

OSPF

- Open Shortest Path First
- Open:
  - Meaning an Open Standard, developed by IETF (OSPF Working Group)
  - OSPFv2 (RFC2328) for IPv4, OSPFv3 (RFC5340) for IPv6
- Shortest Path First:
  - Edsger Dijkstra's algorithm for producing shortest path tree through a graph
    - Dijkstra, E. W. (1959). "A note on two problems in connexion with graphs". Numerische Mathematik 1: 269–271
    - http://www-m3.ma.tum.de/foswiki/pub/MN0506/WebHome/dijkstra.pdf

IS-IS

- Intermediate System to Intermediate System
- ISO 10589 specifies OSI IS-IS routing protocol for ConnectionLess-mode Network Services (CLNS) traffic
  - A Link State protocol with a 2-level hierarchical architecture
  - Type/Length/Value (TLV) options to enhance the protocol
  - Runs on top of the Data Link Layer
  - Uses Dijkstra's SPF algorithm
- RFC 1195 added IP support (called Integrated IS-IS)
- RFC 5308 added IPv6 support

IS-IS & OSPF: Similarities

- Both are Interior Gateway Protocols (IGP)
  - They distribute routing information between routers belonging to a single Autonomous System (AS)
  - Both use Edsger Dijkstra's algorithm
- With support for:
  - Classless Inter-Domain Routing (CIDR)
  - Variable Subnet Length Masking (VLSM)
  - Authentication
  - Multi-path
  - IP unnumbered links

For Service Providers

- Which IGP should an ISP choose?
  - Both OSPF and IS-IS use Dijkstra SPF algorithm
  - Exhibit same convergence properties
  - IS-IS less widely implemented on router platforms
  - IS-IS runs on data link layer, OSPF runs on IP layer
- Why do we keep discussing the merits of each IGP?

For Service Providers

- Biggest ISPs tend to use IS-IS – why?
  - In early 1990s, Cisco implementation of IS-IS was much more stable and reliable than OSPF implementation – ISPs naturally preferred IS-IS
  - Main IS-IS implementations are more tuneable than equivalent OSPF implementations
    - Because biggest ISPs using IS-IS put more pressure on Cisco to implement “knobs” to improve performance

For Service Providers

- Moving forward a decade
  - Early Cisco OSPF implementation substantially rewritten
    - Now competitive with IS-IS in features and performance
  - Router vendors wishing a slice of the core market need an IS-IS implementation as solid and as flexible as that from Cisco
    - Those with IS-IS & OSPF support tend to ensure they exhibit performance and feature parity

How to choose an IGP?

- OSPF
  - Rigid area design – all networks must have area 0 core, with sub-areas distributed around
  - Suits ISPs with central high speed core network linking regional PoPs
- IS-IS
  - Relaxed two level design – L2 routers must be linked through the backbone
  - Suits ISPs with “stringy” networks, diverse infrastructure, etc, not fitting central core model of OSPF
  - More flexible than OSPF, but easier to make mistakes too

Considerations

- “Security”
  - IS-IS runs on link layer
  - Not possible to directly “attack” IS-IS using IP
- Not dependent on IP addressing
  - IS-IS’s NSAP addressing scheme avoids dependencies on IP as with OSPF
- “Reliability”
  - IS-IS has long been used by the majority of the world’s biggest ISPs
  - Belief that equipment vendors pay more attention to IS-IS reliability, scalability, and features

More considerations

- Migration to IPv6
  - Adding IPv6 means OSPFv2 and OSPFv3 in network
    - Two independent protocols, two sets of identical configuration
  - IS-IS simply requires the addition of the IPv6 address-family
  - Note that RFC5838 describes support of multiple address families in OSPFv3
    - Limited vendor support
    - Is not compatible with OSPFv2

Securing the IGP

- Only use IGP on internal interfaces
  - Not on customer or peer or transit interfaces
- Only use IGP for carrying internal prefixes
  - Internal point to point link addresses
  - Internal loopback (management) addresses
- Use adjacency/neighbour authentication
  - No adjacency established otherwise
- Use scalability and performance recommendations developed by the Internet industry
  - Vendor “out of the box” defaults are never sufficient

Deploying BGP with scalability and security

KEY COMPONENTS OF 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
- The Autonomous System is the cornerstone of BGP
  - It is used to uniquely identify networks with a common routing policy

Autonomous System (AS)

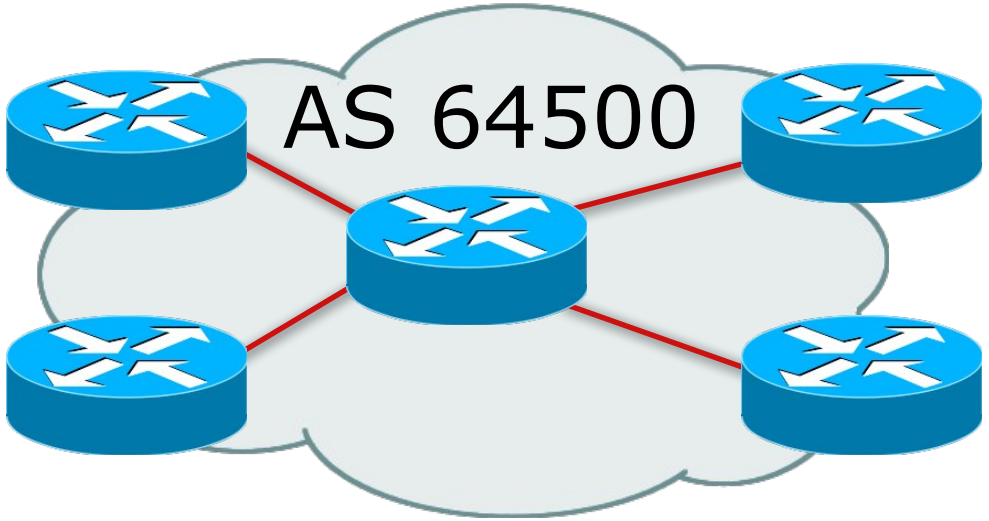

- Collection of networks with same routing policy
- Single routing protocol
- Usually under single ownership, trust and administrative control
- Identified by a unique 32-bit integer (ASN)

Autonomous System Number

32-bit range representation specified in RFC5396

Defines “asplain” (traditional format) as standard notation

Range:	
0-4294967295	(32-bit range – RFC6793)
	(0-65535 was original 16-bit range)
Usage:	
0 and 65535	(IANA Reserved)
1-64495	(public Internet)
64496-64511	(documentation – RFC5398)
64512-65534	(private use only)
23456	(represent 32-bit range in 16-bit world)
65536-65551	(documentation – RFC5398)
65552-131071	(IANA Reserved)
131072-458751	(public Internet)
458752-4199999999	(IANA Reserved/Unallocated)
4200000000-4294967294	(private use only – RFC6996)
4294967295	(IANA Reserved – RFC7300)

BGP Basics

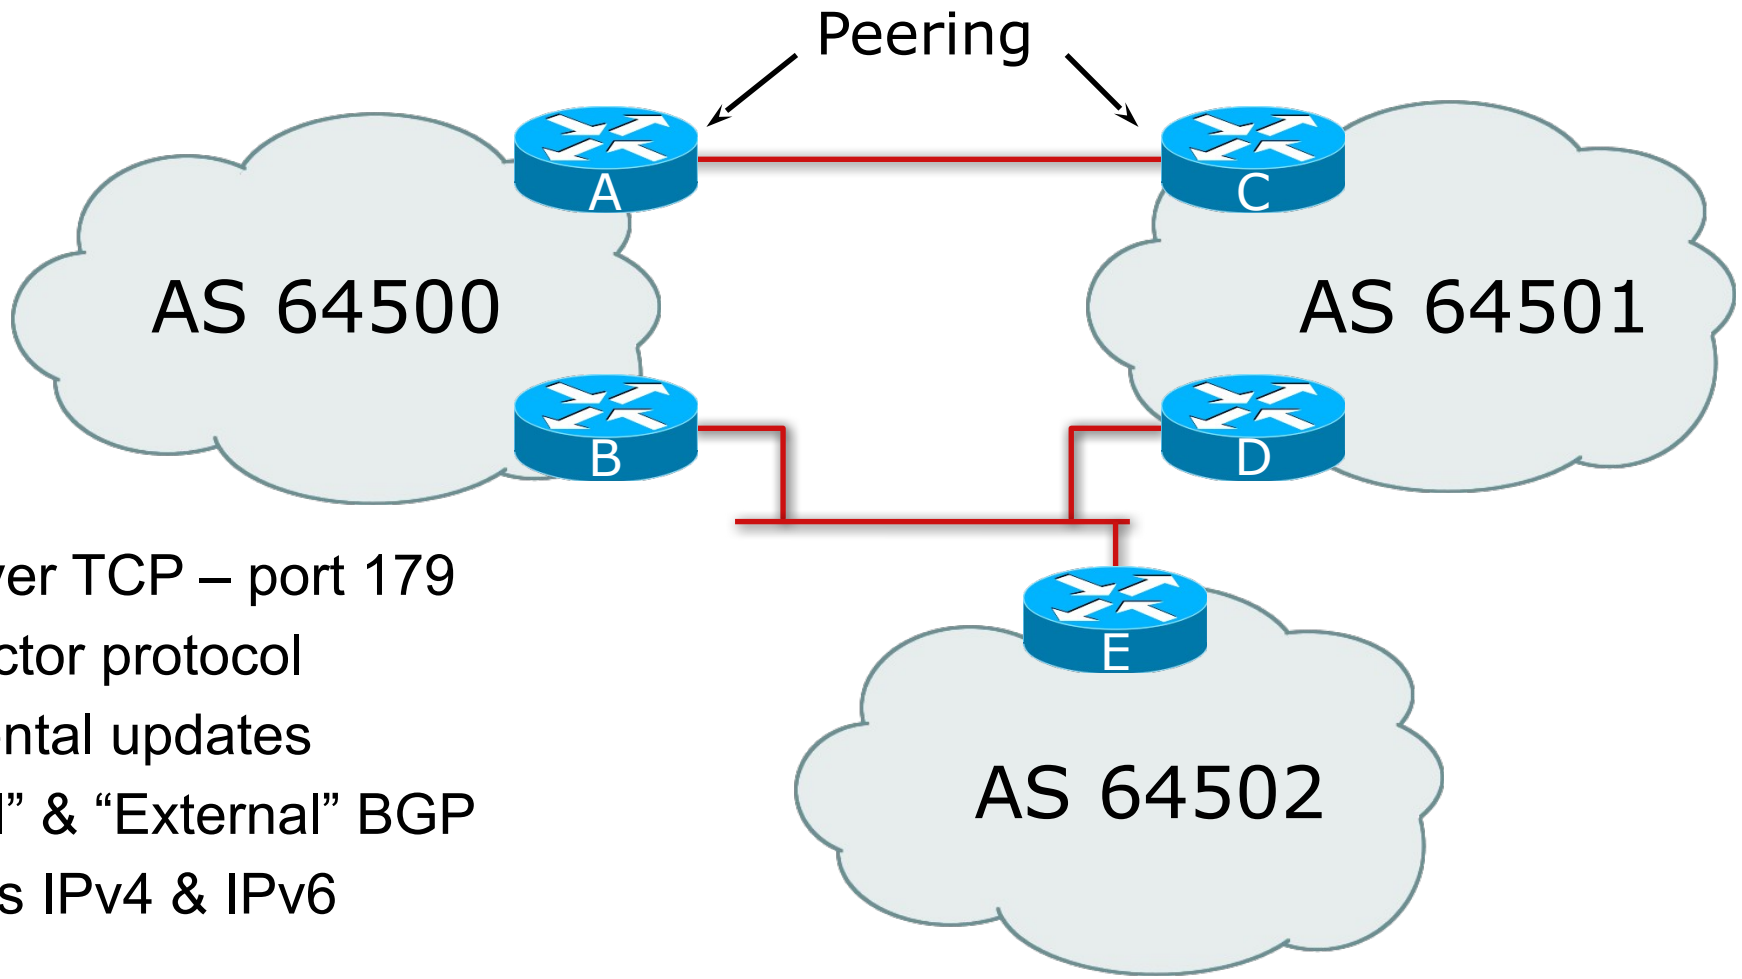

- Runs over TCP – port 179
- Path vector protocol
- Incremental updates
- “Internal” & “External” BGP
- Supports IPv4 & IPv6

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

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) if:
  - Prefix and prefix length are unique (after best path selection)
  - and
  - Lowest “protocol distance”
- Best paths in the RIB are installed in forwarding table (FIB)

Constructing the Forwarding Table

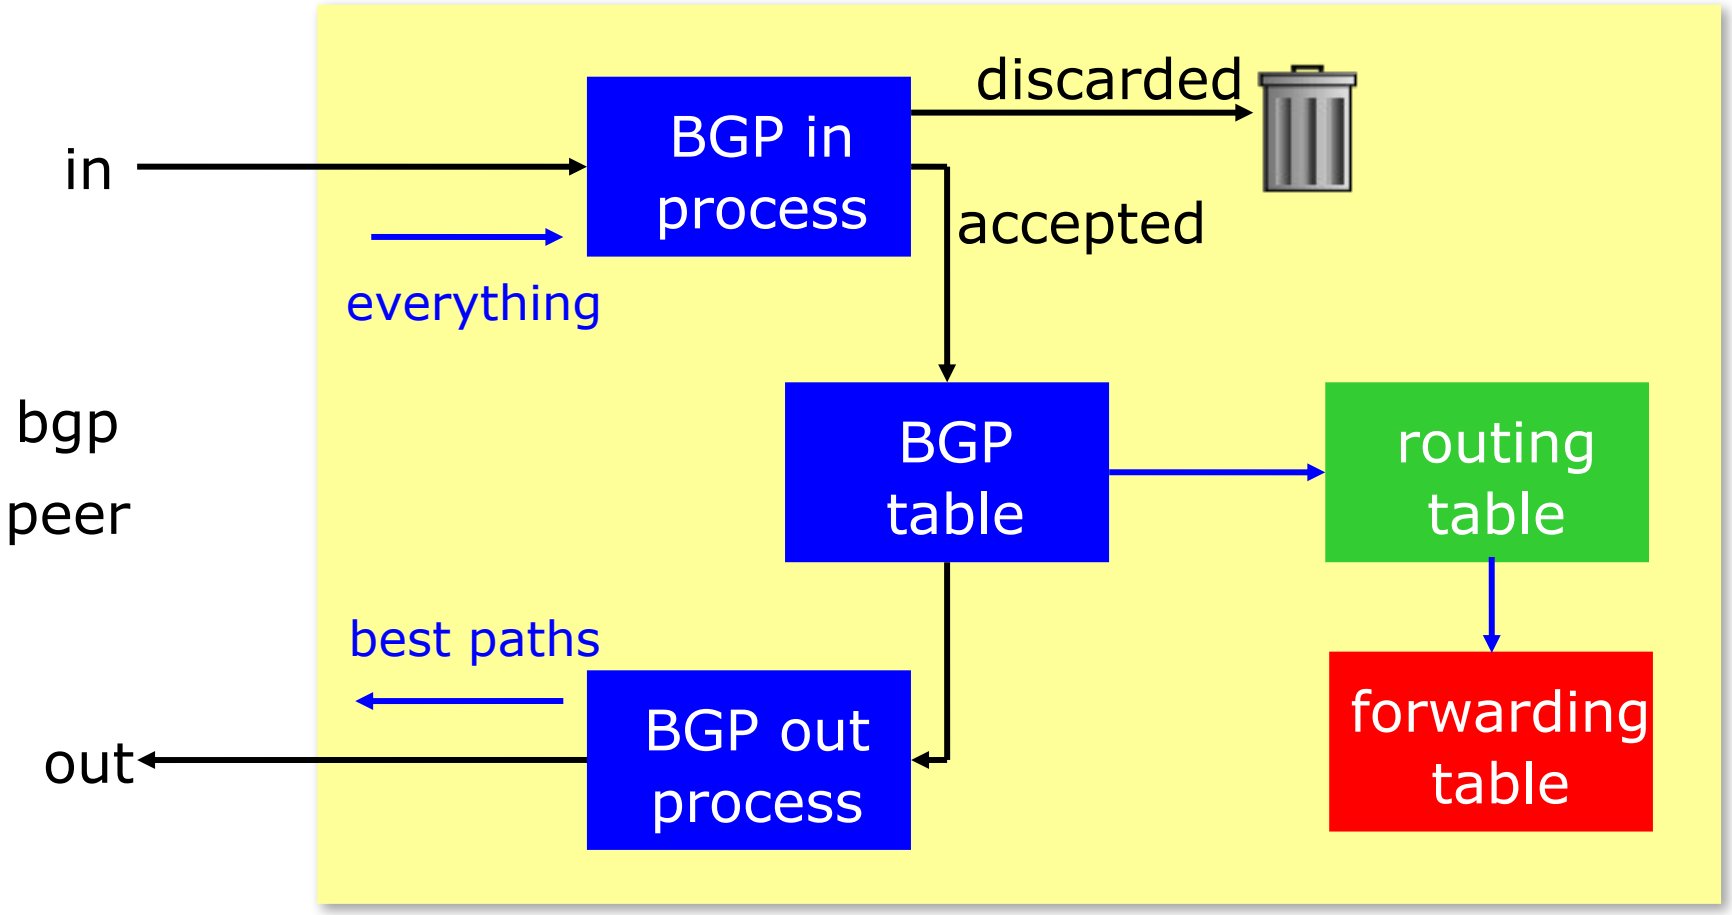

Supporting Multiple Protocols

- Independent operation
  - One RIB per protocol
    - IPv6 routes in BGP's IPv6 RIB
    - IPv4 routes in BGP's IPv4 RIB
  - Each protocol can have its own policies
- NEXTHOP
  - The IP address of the next router must belong to the same address family as that of the local router

Supporting Multiple Protocols

- Cisco IOS assumes that all BGP neighbours will exchange IPv4 unicast prefixes
  - Most other implementations do not
  - We need to remove this assumption in Cisco IOS

```
router bgp 64500
 no bgp default ipv4-unicast
```

- For operational simplicity, the desire is for:
  - IPv4 neighbours to exchange IPv4 unicast prefixes
  - IPv6 neighbours to exchange IPv6 unicast prefixes
- Failure to do this results in:
  - IPv6 neighbours appearing to be set up to exchange IPv4 unicast prefixes
  - Cluttered configuration, confusing troubleshooting and diagnosis

EBGP & IBGP

- BGP is used
  - Internally (IBGP)
  - Externally (EBGP)
- IBGP used to carry
  - Some/all Internet prefixes across network operator backbone
  - ISP's customer prefixes
- EBGP used to
  - Exchange prefixes with other ASes
  - Implement routing policy

BGP/IGP model used in service provider networks

- Model representation

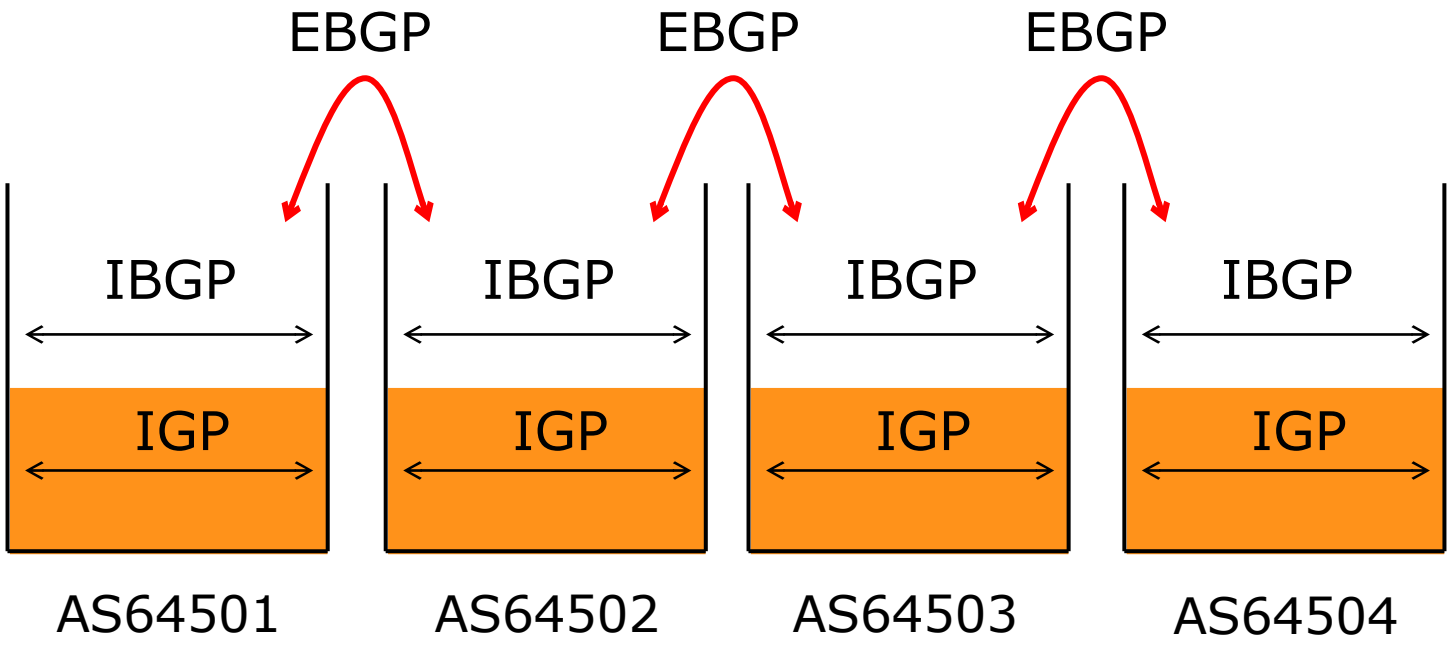

External BGP Peering (EBGP)

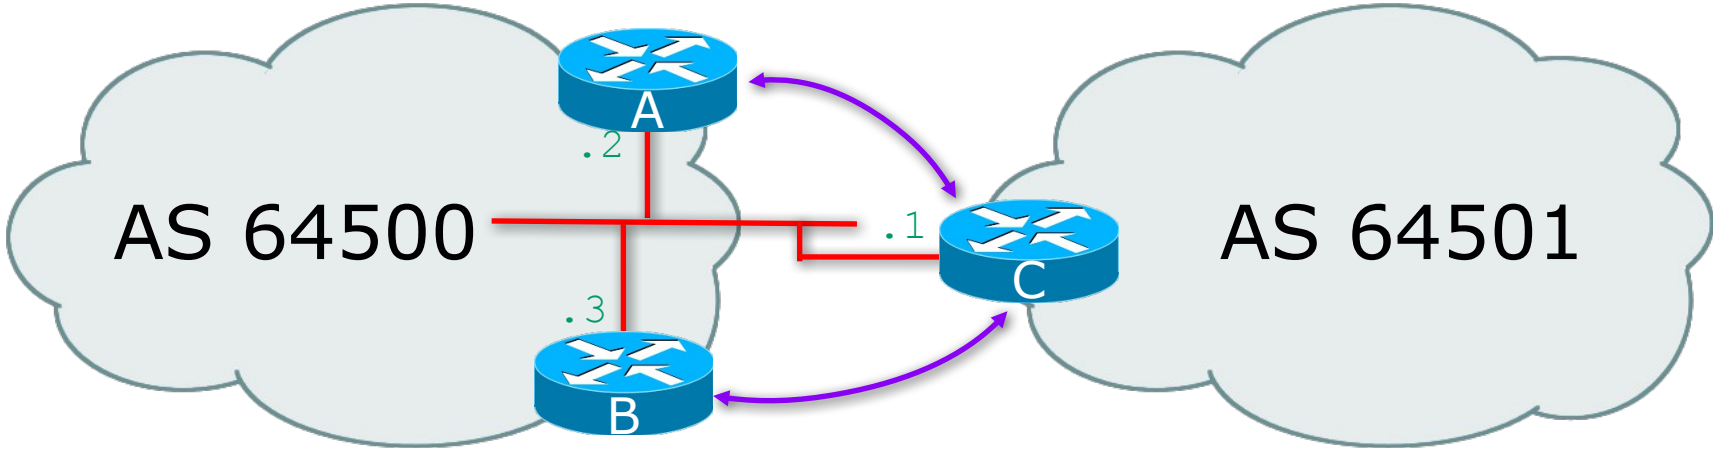

- Between BGP speakers in different AS
- Should be directly connected
- Never run an IGP between EBGP peers

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 AS
  - They do not pass on prefixes learned from other IBGP speakers

Internal BGP Peering (IBGP)

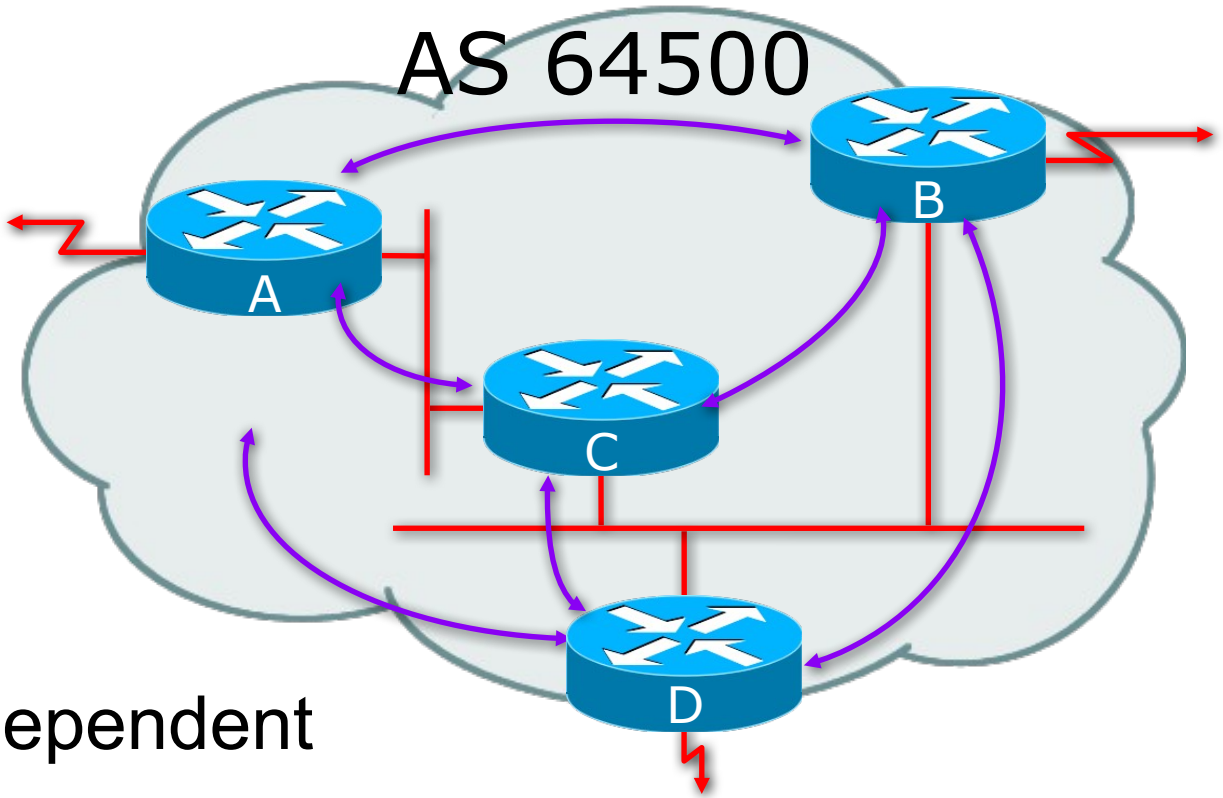

- Topology independent
- Each IBGP speaker must peer with every other IBGP speaker in the AS as per ⟷

Peering between Loopback Interfaces

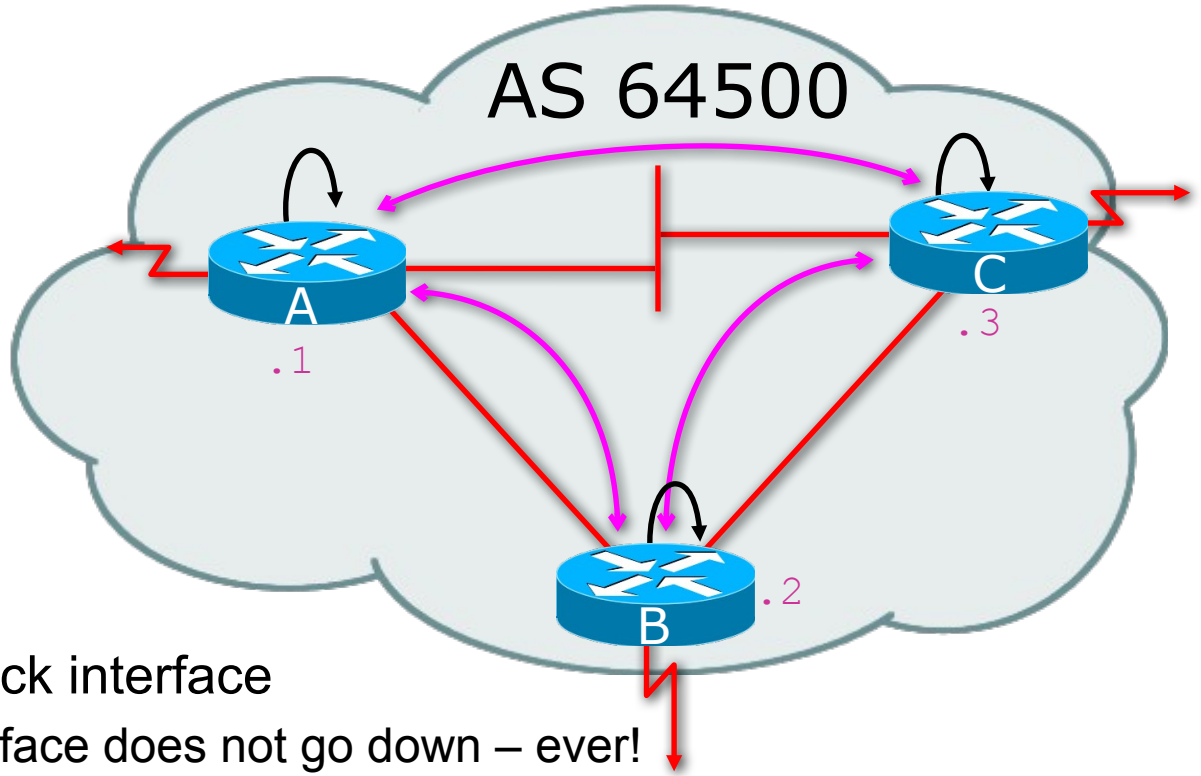

- Peer with loop-back interface
  - Loop-back interface does not go down – ever!
- Do not want IBGP session to depend on state of a single interface or the physical topology

Inserting prefixes into BGP

- Two ways to insert prefixes into BGP
  - *redistribute static*
    - Static route must exist before *redistribute* command will work
    - Care required – *redistribute* will move everything from the named protocol into BGP unless filters are applied
  - *network* command
    - A matching route must exist in the routing table before the network is announced

Practical Routing Security

- Summary:
  - MANRS
  - Interior routing protocol best practices
  - Key components of BGP
  - Securing the Router
  - BGP Best Current Practices
  - uRPF
  - RTBH
  - RPKI, ROAs and ROV

Questions/Discussion?