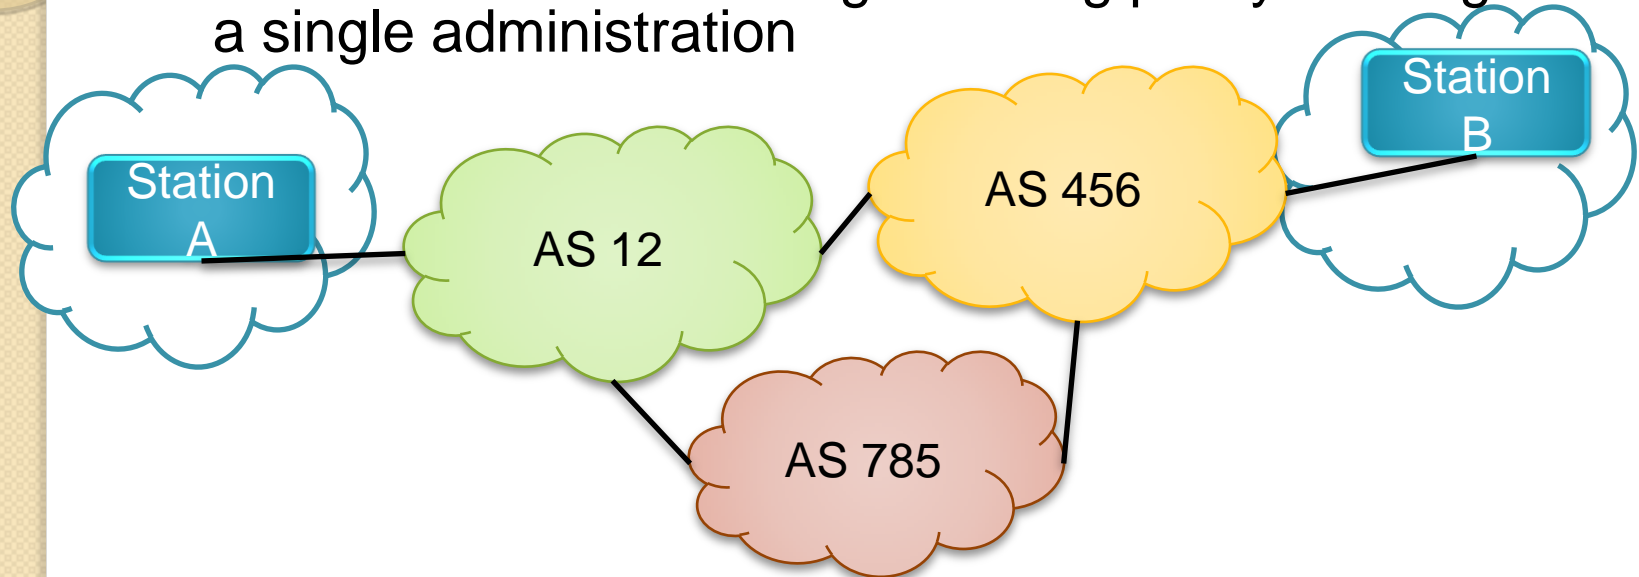




Inter-Domain Networks

Inter-domain Structure

- Autonomous system (AS):
 - set of routers have a single routing policy running under a single administration



- From outside, AS looks like a single entity – identified by a AS#
- Routing information between ASs exchanged by the **Border Gateway Protocol (BGP)**


Inter-domain Structure...

- Segregating the network into AS
 - controls the expansion of routing tables
 - provide a structured view of the Internet
 - AS could have independent routing policies
- Problems:
 - End-to-end quality of service

Inter-domain Structure...

Stub AS:

- An AS is “stub” when it reaches outside networks via a single point
- Stub AS is single-homed with respect to its provider
- Stub AS numbered from a “private pool” 65412-65535
- Stub AS need not learn any AS-level routes – just use a default route to provider AS




W and Y are stub AS domains

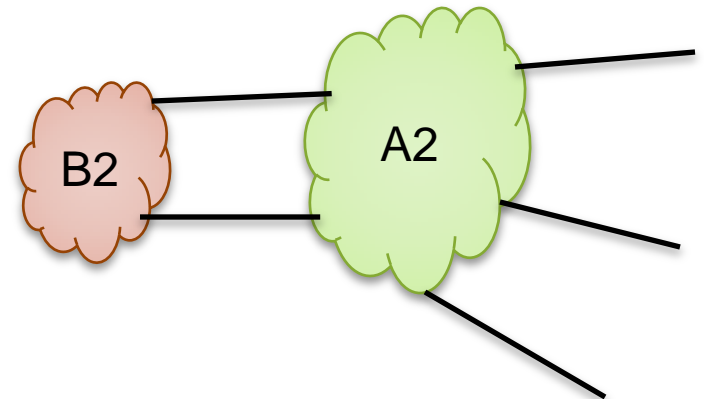
Inter-domain Structure...

Multihomed AS

- AS is multihomed if it has more than one exit point to outside
- AS can be multihomed to a single or multiple providers

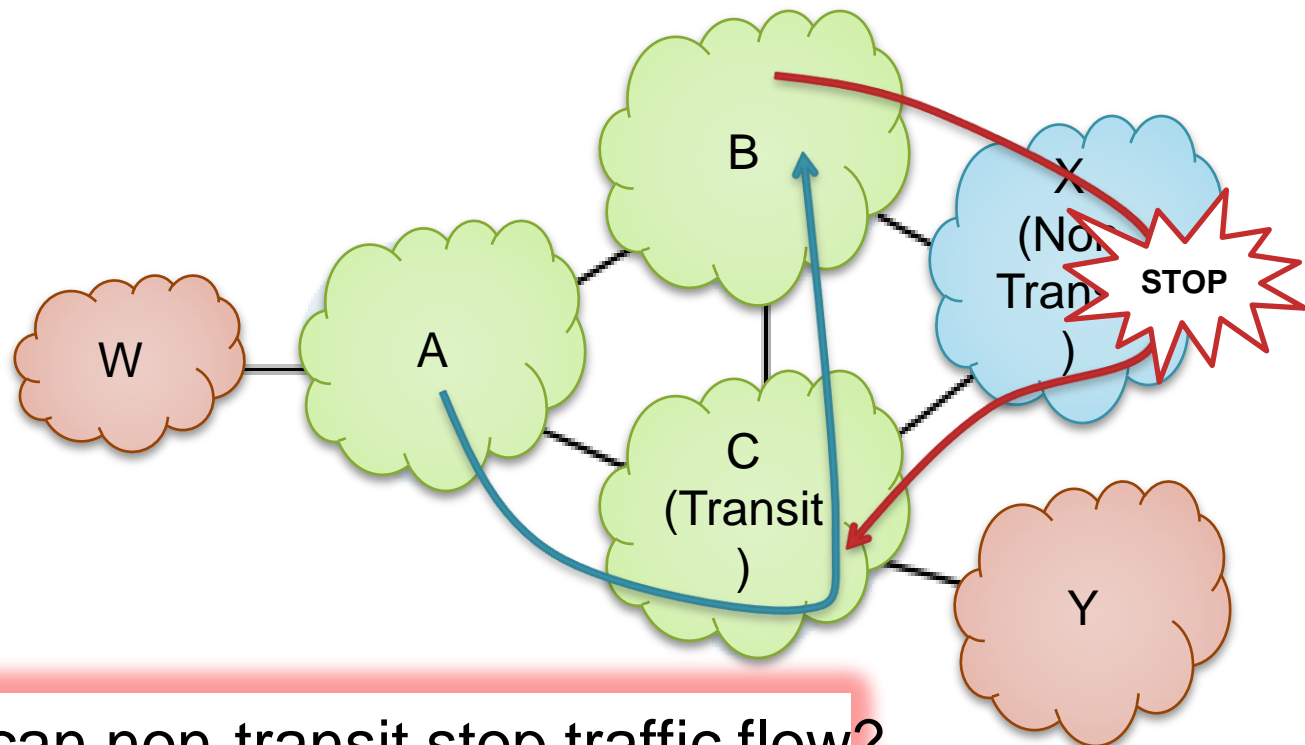


A, B, C, and X are multihomed AS domains



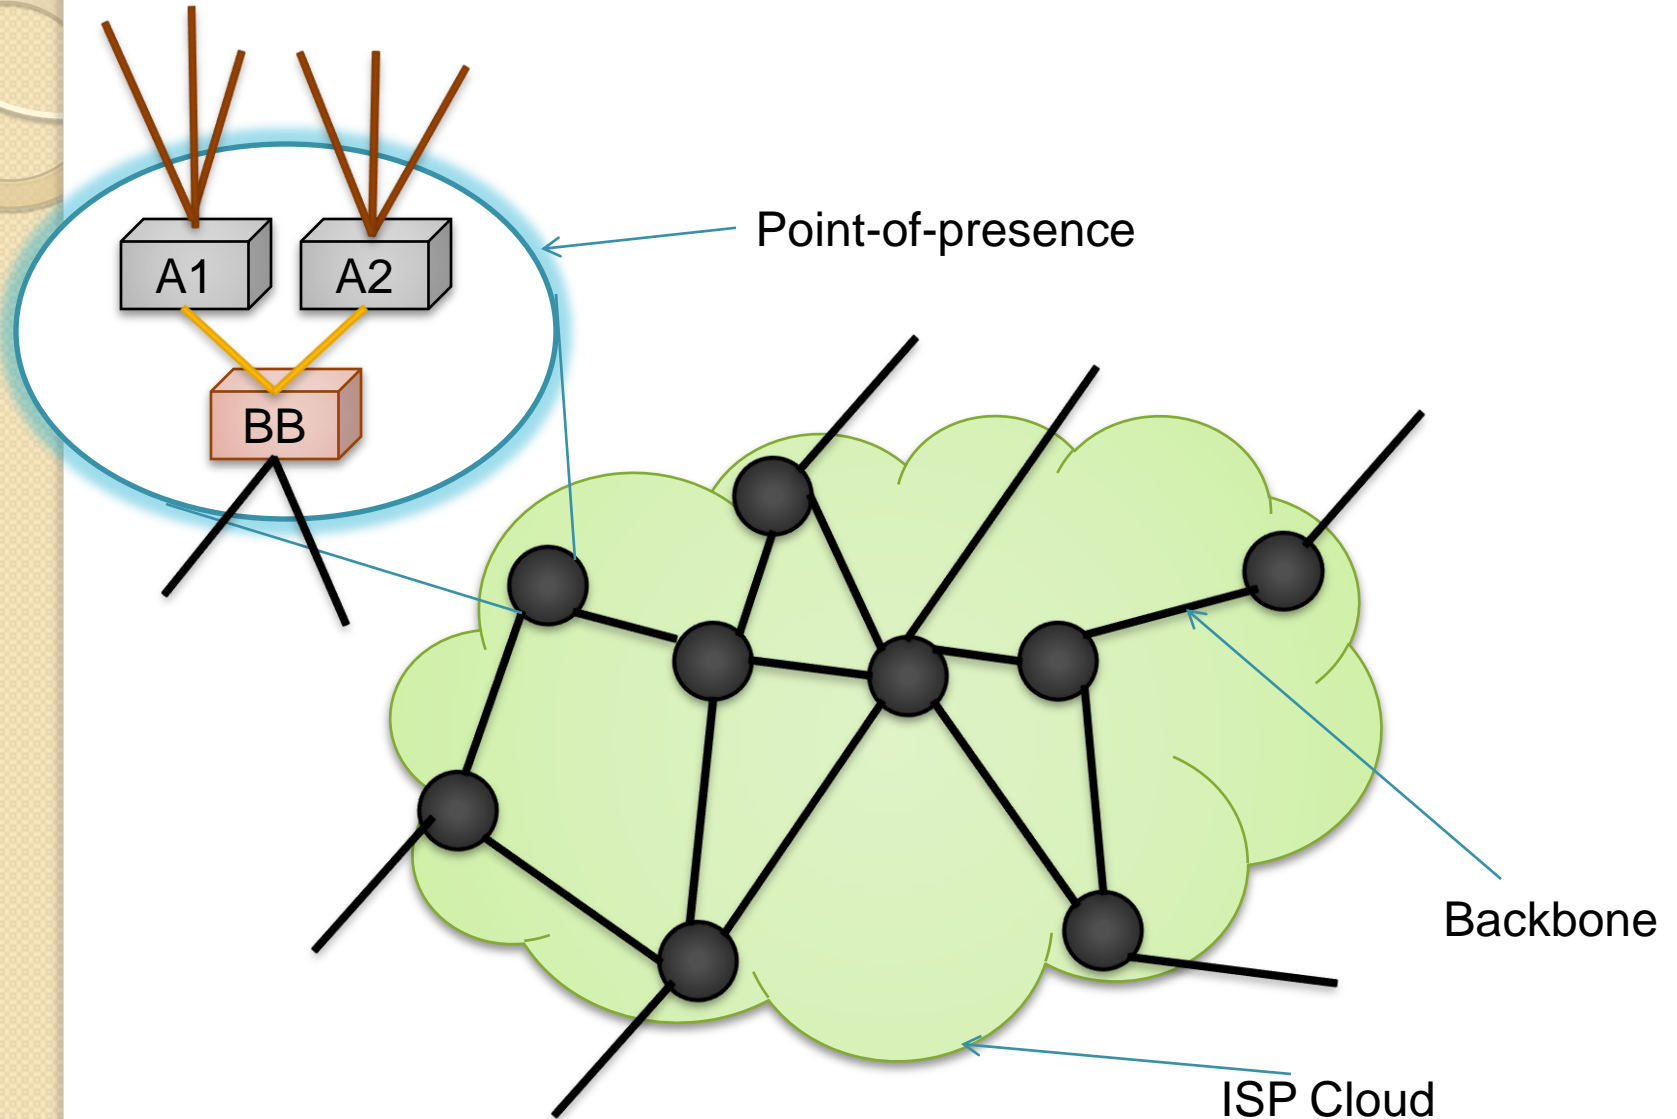
Inter-domain Structure...

- **Multi-homed** domains can be split
 - **Transit** – allow other traffic to pass through
 - **Non-transit** – does not carry traffic for others!

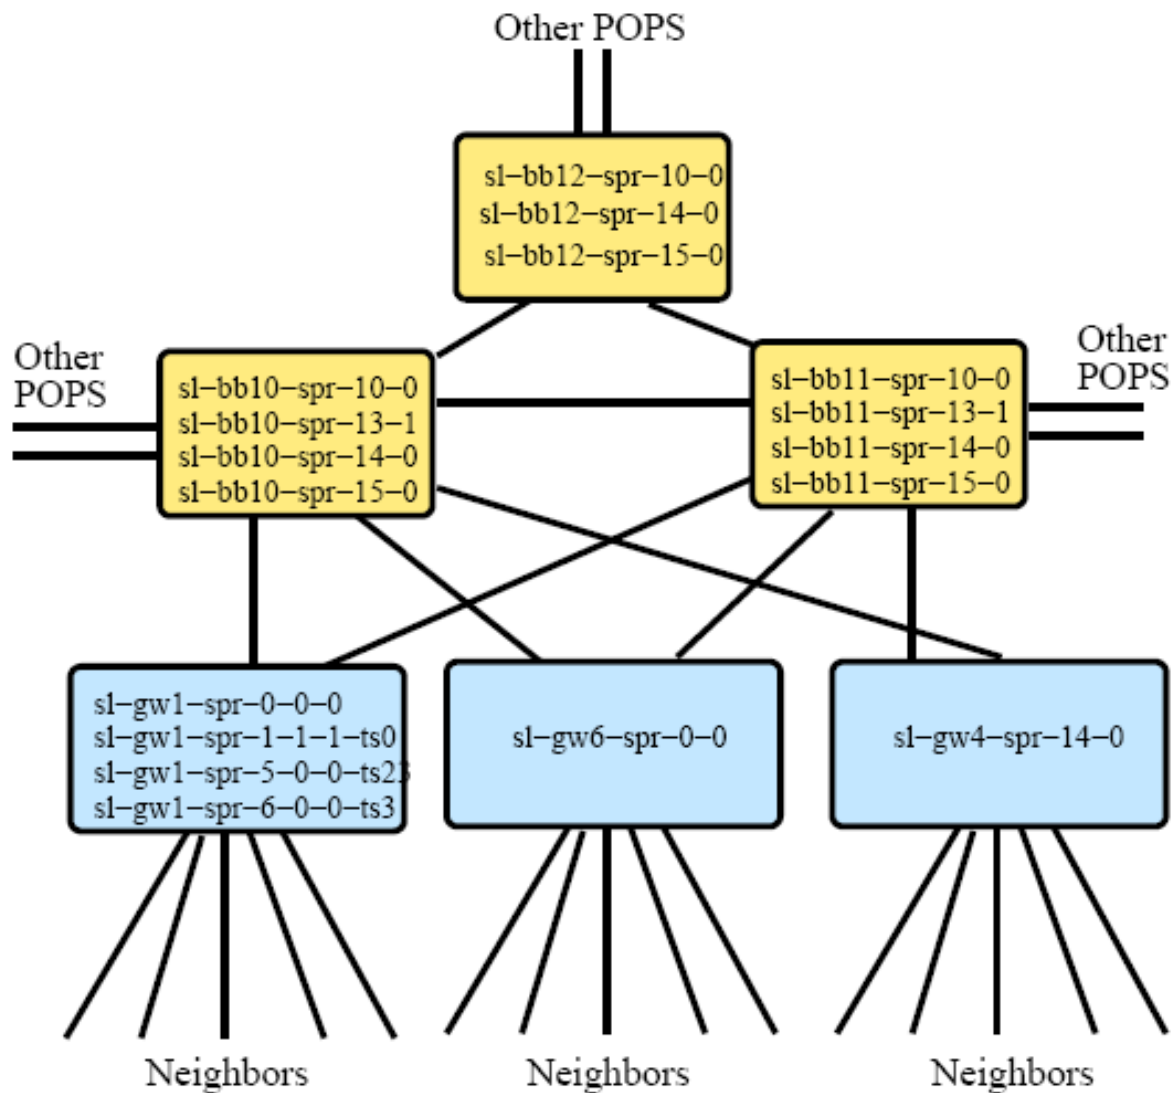


How can non-transit stop traffic flow?

More on Inter-Domain Structure



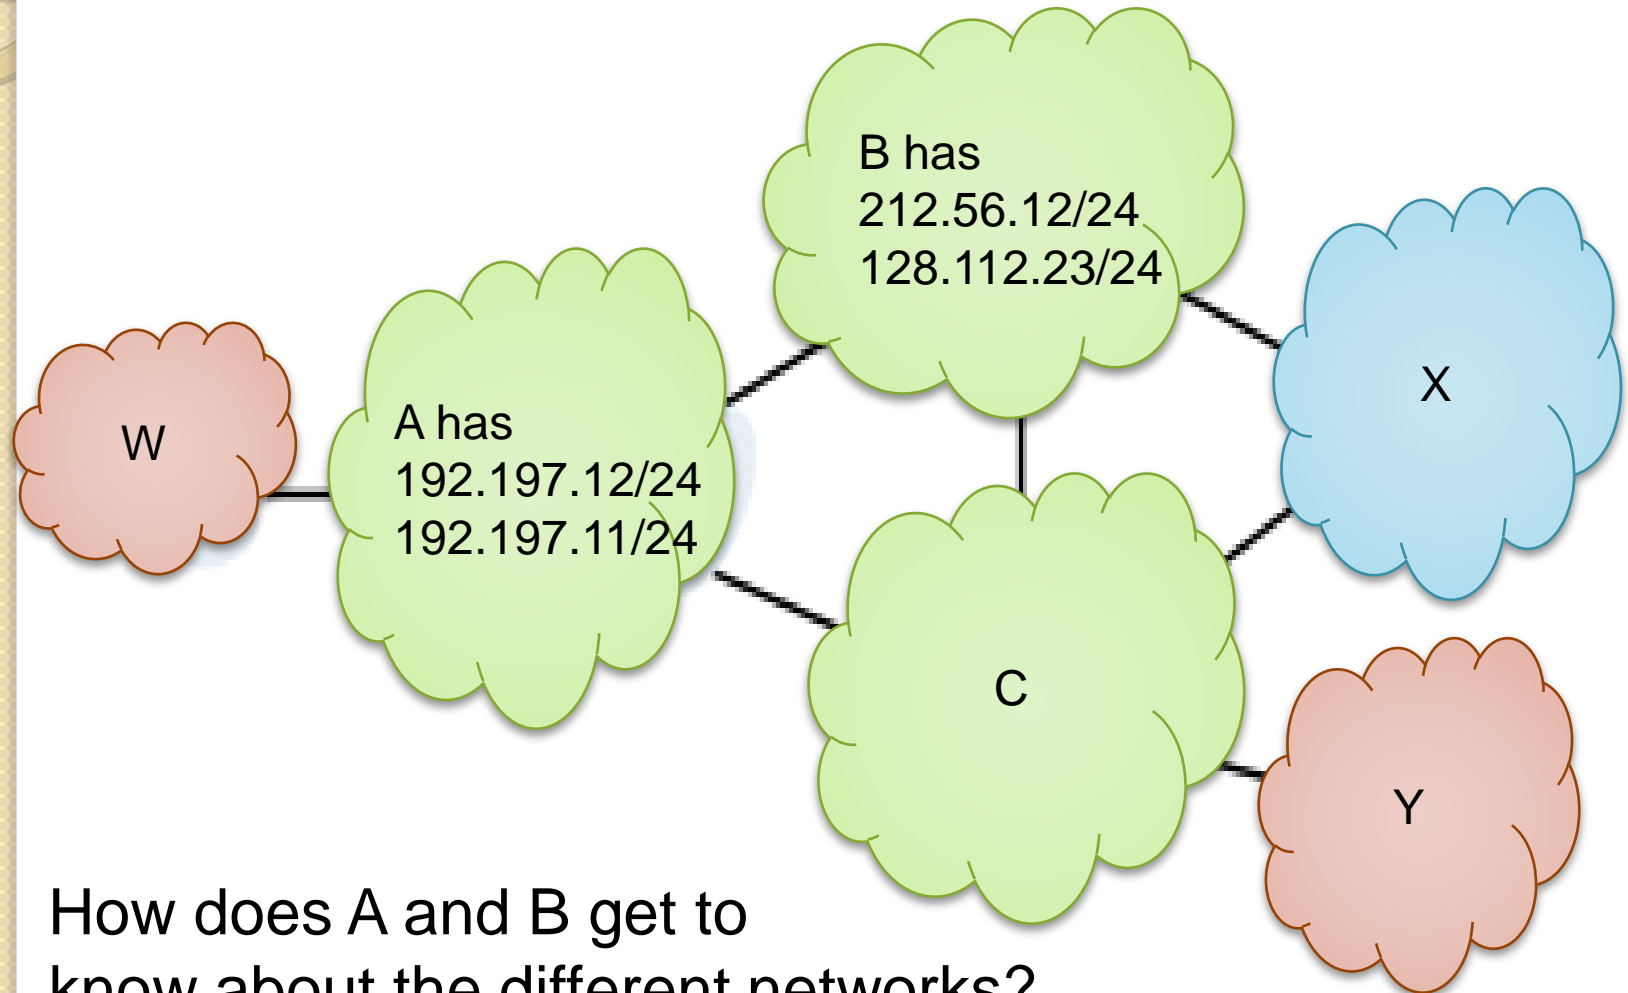
More on Internet Structure




More on Internet Structure

AS	Name	ISP		with customer & peer		POPs
		Routers	Links	Routers	Links	
1221	Telstra (Australia)	355	700	2,796	3,000	61
1239	Sprintlink (US)	547	1,600	8,355	9,500	43
1755	Ebone (Europe)	163	300	596	500	25
2914	Verio (US)	1,018	2,300	7,336	6,800	121
3257	Tiscali (Europe)	276	400	865	700	50
3356	Level3 (US)	624	5,300	3,446	6,700	52
3967	Exodus (US)	338	800	900	1,100	23
4755	VSNL (India)	11	12	121	69	10
6461	Abovenet (US)	367	1,000	2,259	1,400	21
7018	AT&T (US)	733	2,300	10,214	12,500	108

Inter-Domain Routing Basics



Inter-Domain Routing Basics



Inter-domain Routing Basics


- Border Gateway Protocol (BGP)
 - a **path vector protocol** used to carry routing information between AS domains
 - path refers to a sequence of AS#s – indicating the path traversed
- Two BGP routers (peers) exchange NLRI peers



Inter-domain Routing Basics

AS2 is a transit domain for this advertisement

NLRI given in IP prefixes (CIDR style)



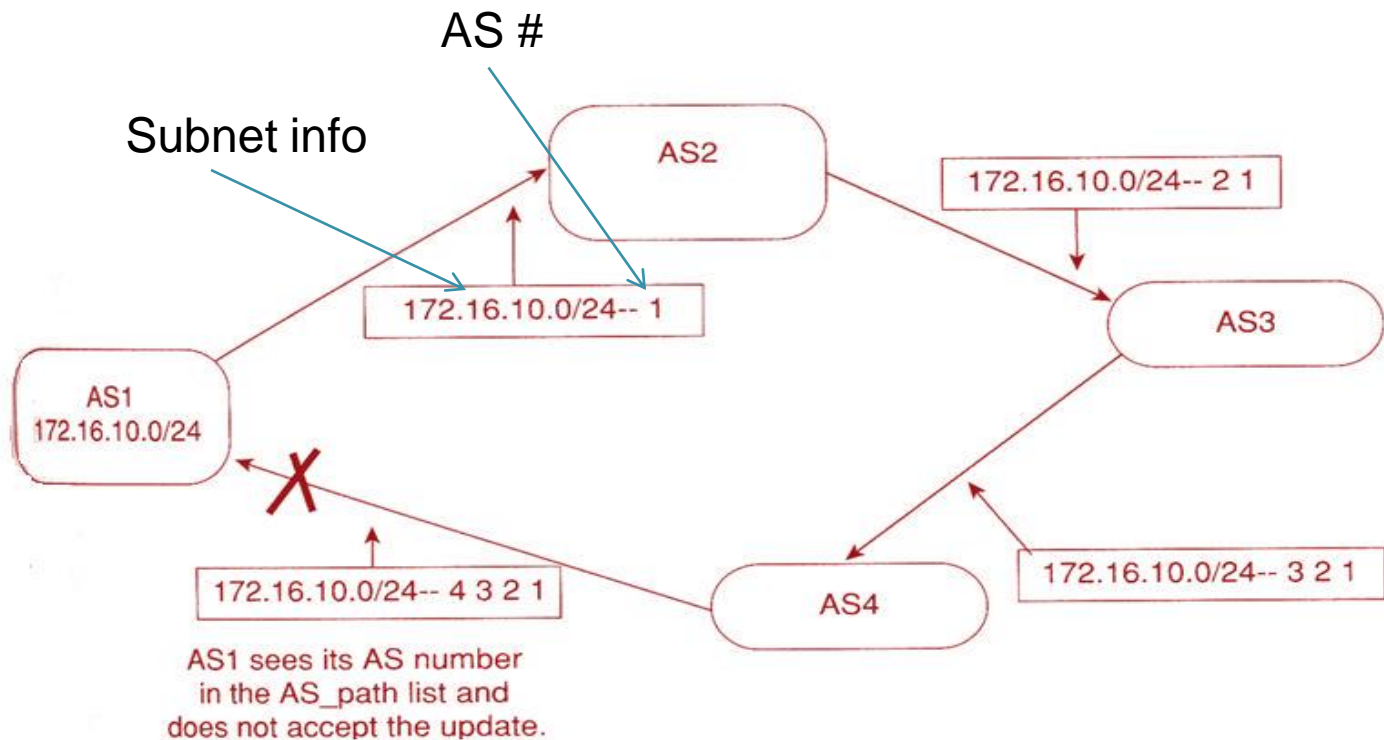
NLRI – network layer reachability information

Inter-domain Routing Basics...

- NLRI includes
 - **Subnet** that is reachable
 - AS_path (sequence of AS#s need to be traversed to reach the **Subnet**)
- AS originating route adds its own AS#
- Each AS transmitting the advertisement prepends its own AS#
 - Originating AS at the end of the AS_path
- AS_path used for loop detection and prevention


Simple Example of BGP Routing

- If route is advertised to the AS that originated it (loop), the AS_path attribute will contain the AS#, the AS will reject the route



Another BGP Example

- AS T has 2 class C nets
 - 197.8.0/24 and 197.8.1/24 – aggregated by a 23-bit prefix 197.8.0/23
- ASs X and Y that use AS T as transit and are allocated
 - 197.8.2/24 and 197.8.3/24, respectively



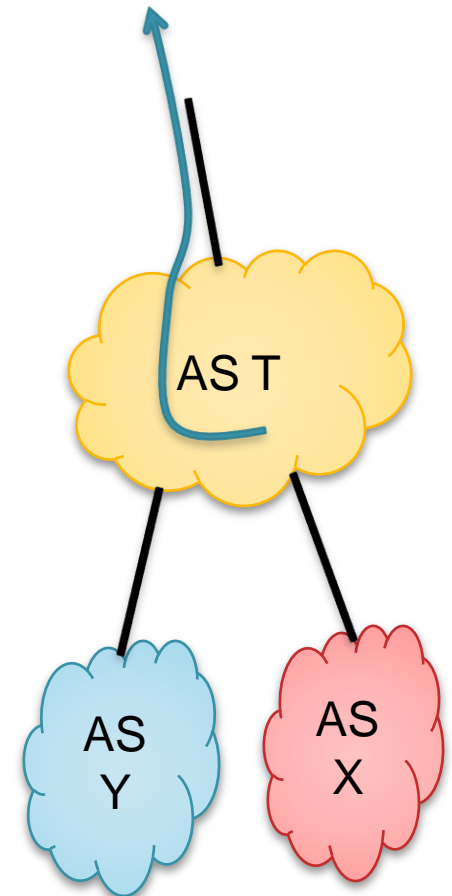
Another BGP Example...

- Can we **aggregate** to reduce the number of advertisements?
- With route aggregation
 - Path A: reaches 197.8.0/22
- What is the AS path?
 - we cannot just list T, loop detection need the complete path
 - listing a complete path like T, X, Y is misleading -- implies a three hop path

Another BGP Example

- AS path attribute into two components:
 - **ordered** list -- AS sequence
 - **unordered** set -- AS set
 - Path: (Sequence (T), Set (X, Y))

Path A:
197.8.0/22 (T {X,Y})



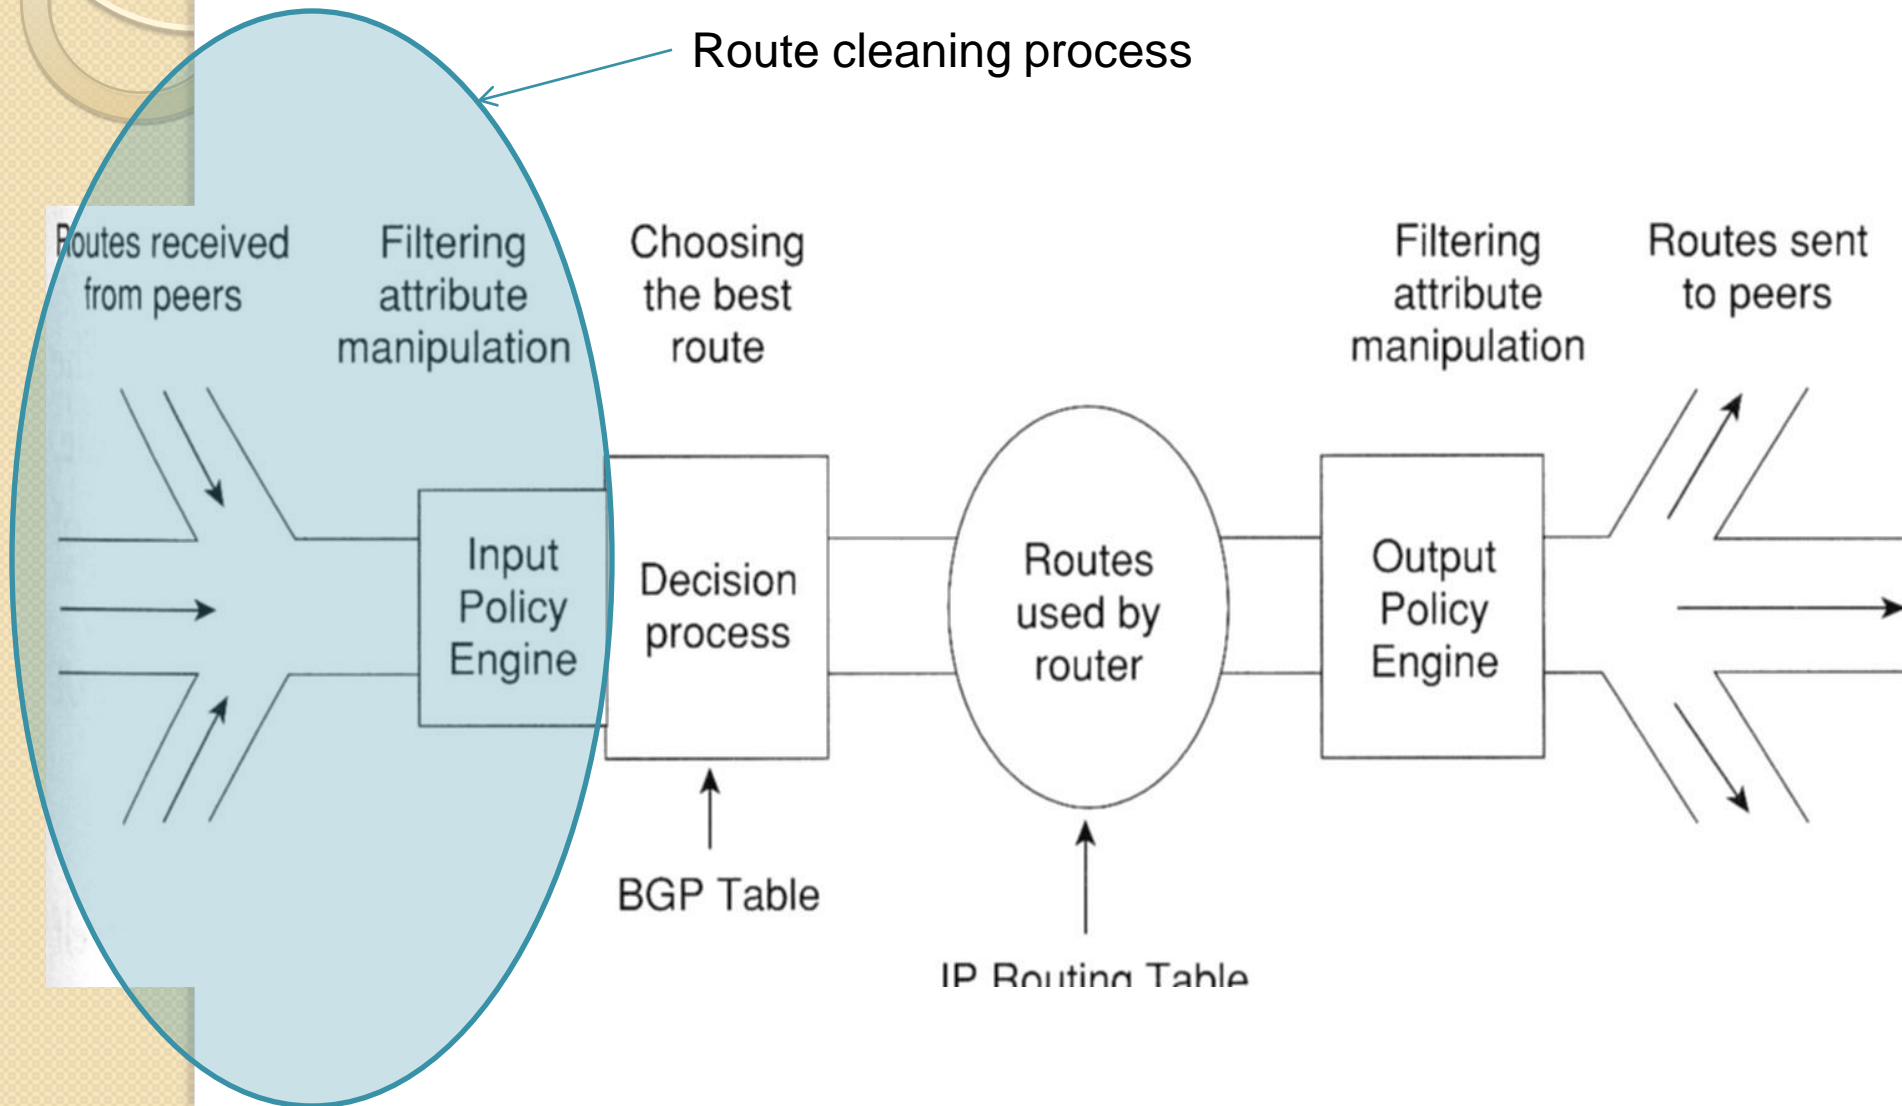
Path Vectors

- If Z wants to forward this path to one of its neighbors, it will place its own AS # in the front
 - Path: (Sequence (Z, T), Set (X, Y))
- Sequence and set components are used for loop detection
- Rule for path aggregation:
 - sequence components should be the intersection of all sequences
 - set of components contain all the ASs mentioned in any of the paths to aggregate yet are not present in the aggregated sequence

BGP Routing Process

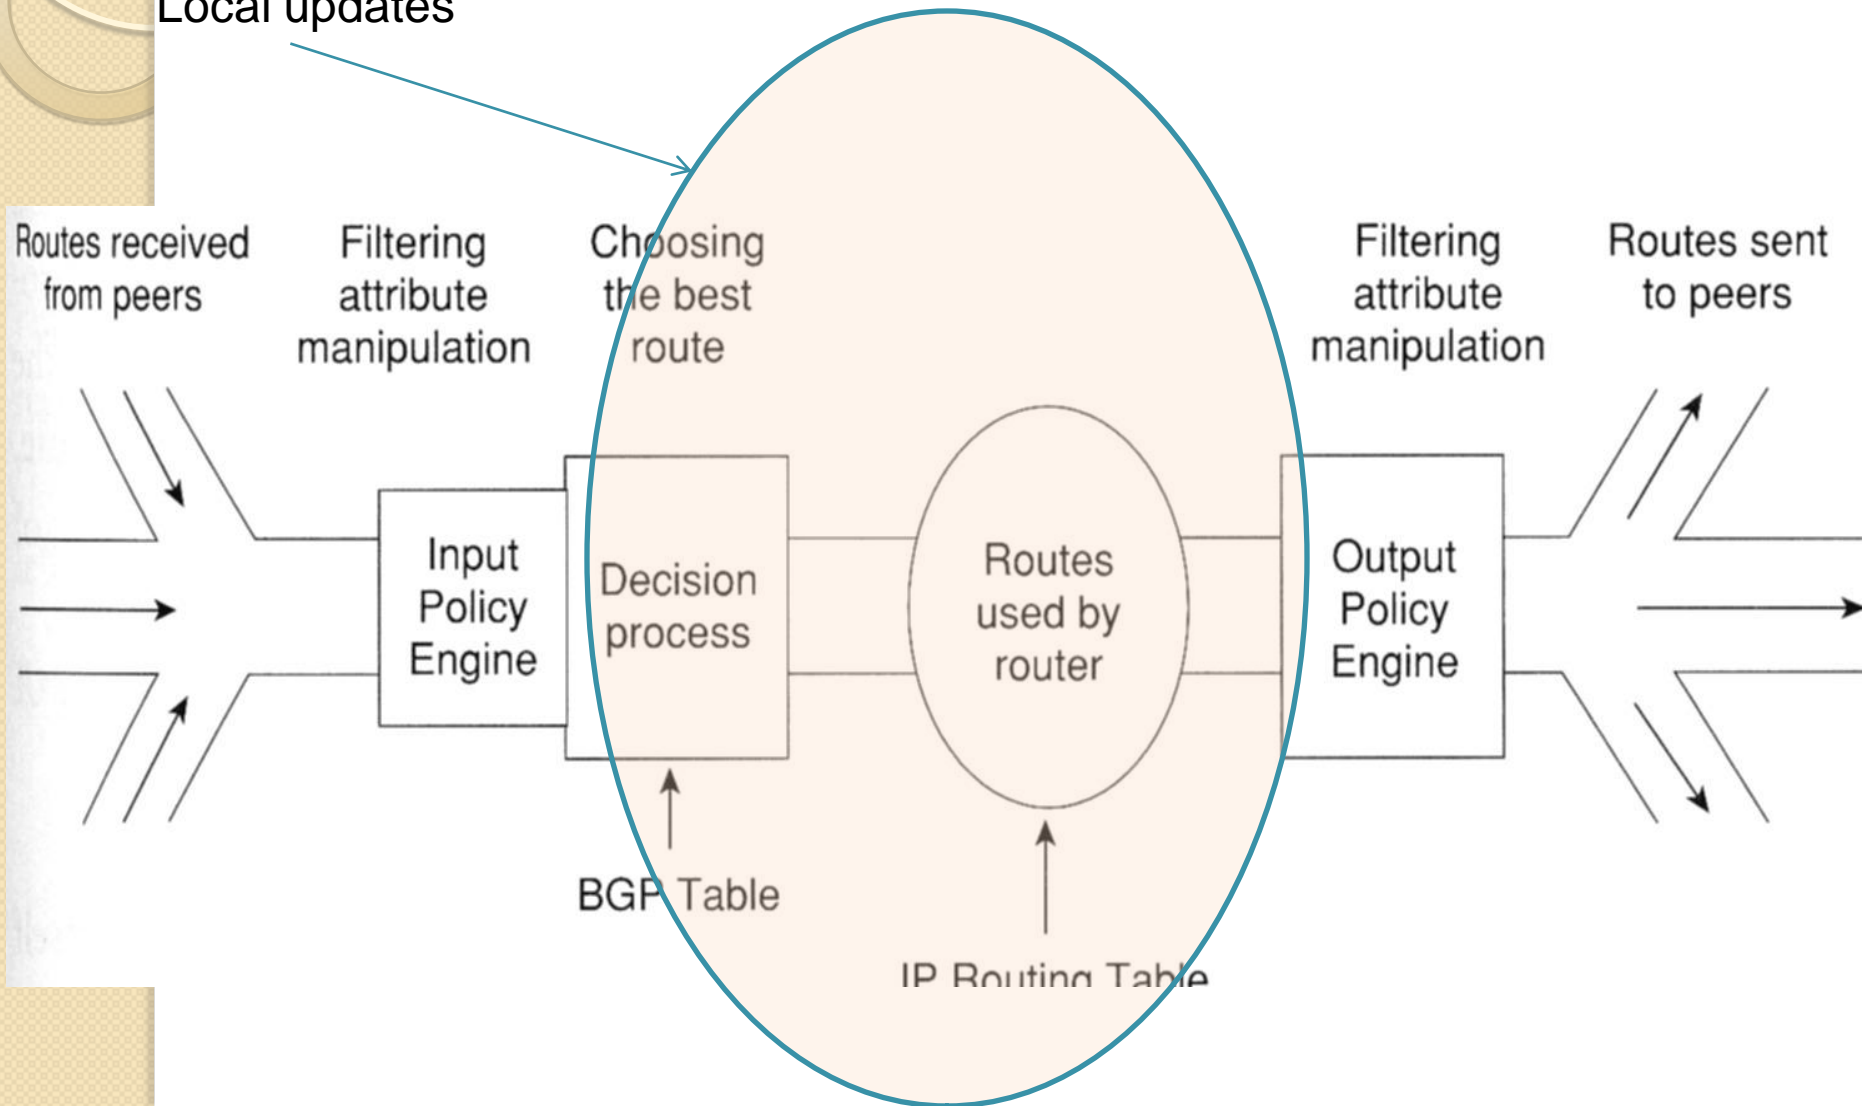
- Rough overview
 - Exchange BGP advertisements
 - Clean incoming advertisements using local policies
 - Select best reachability information and update local (intra-domain) routing information
 - Advertise new reachability to other domains

BGP Routing Process

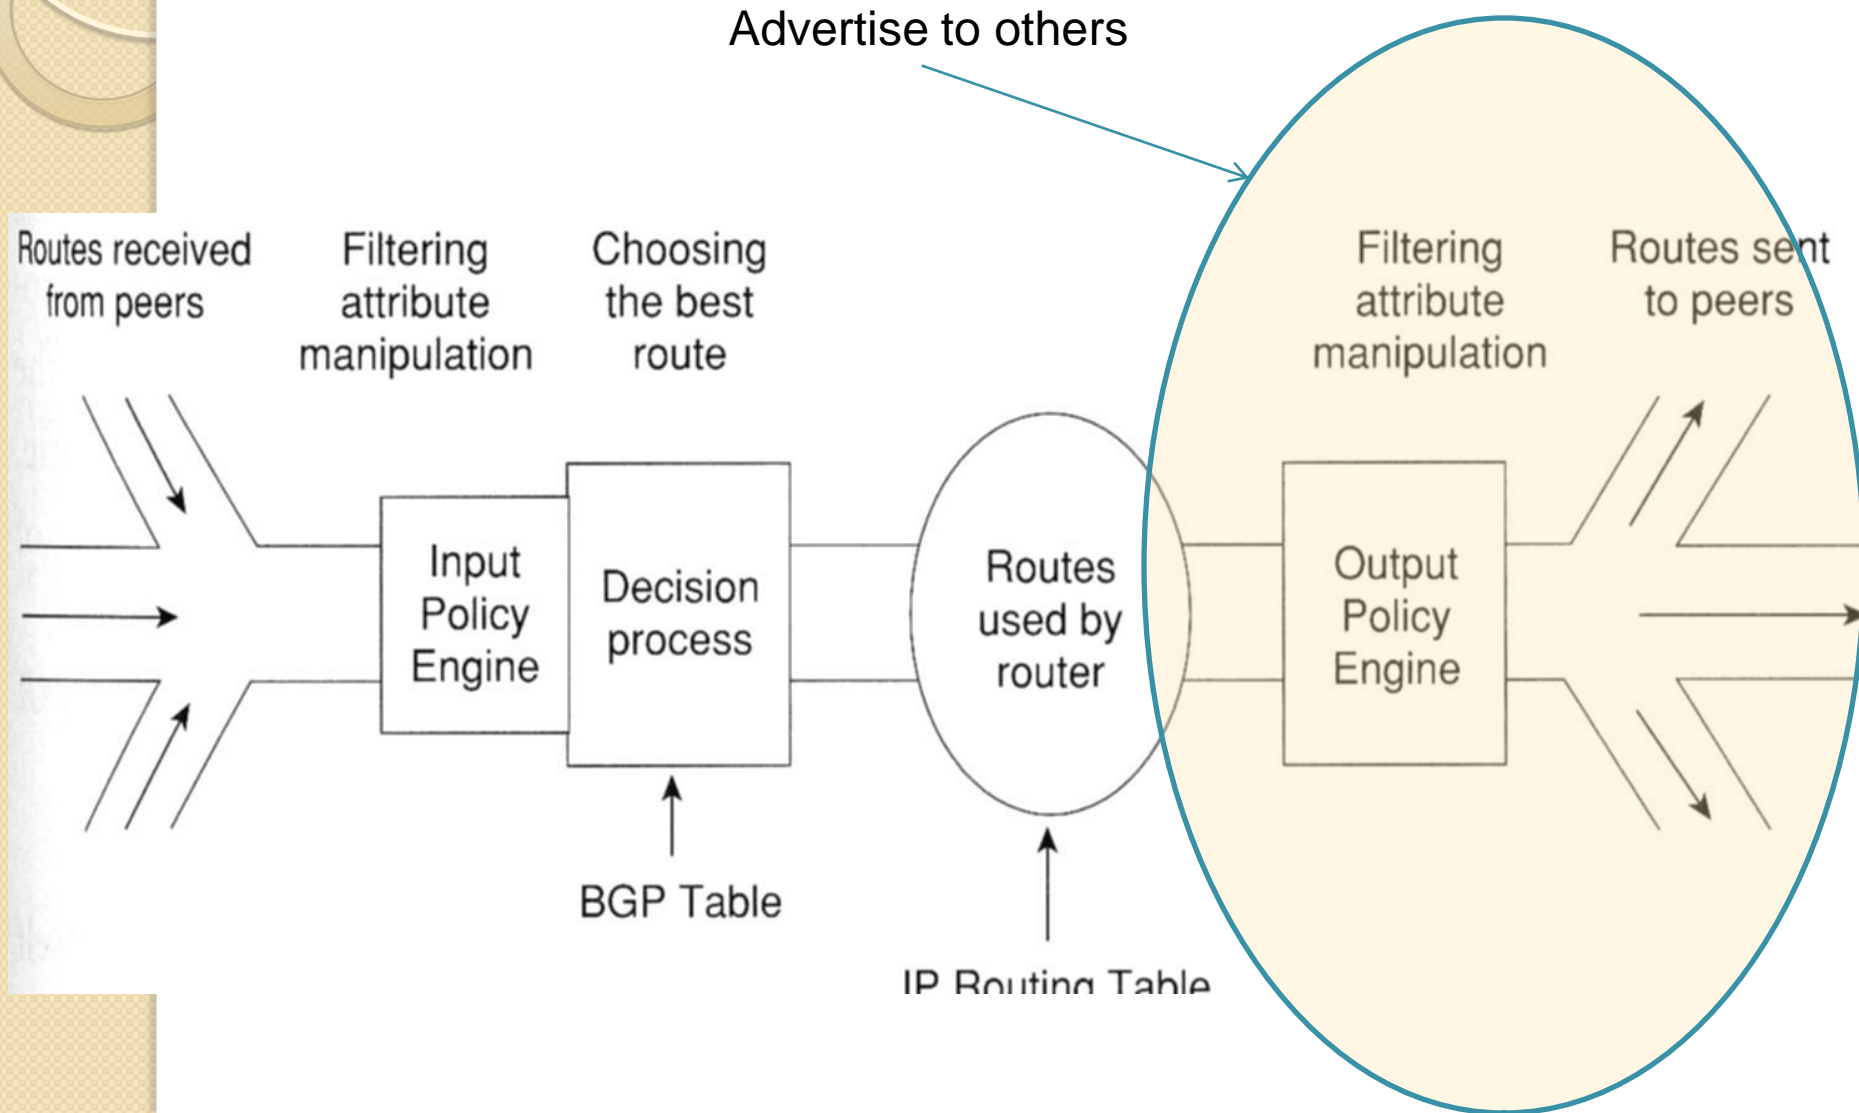


BGP Routing Process

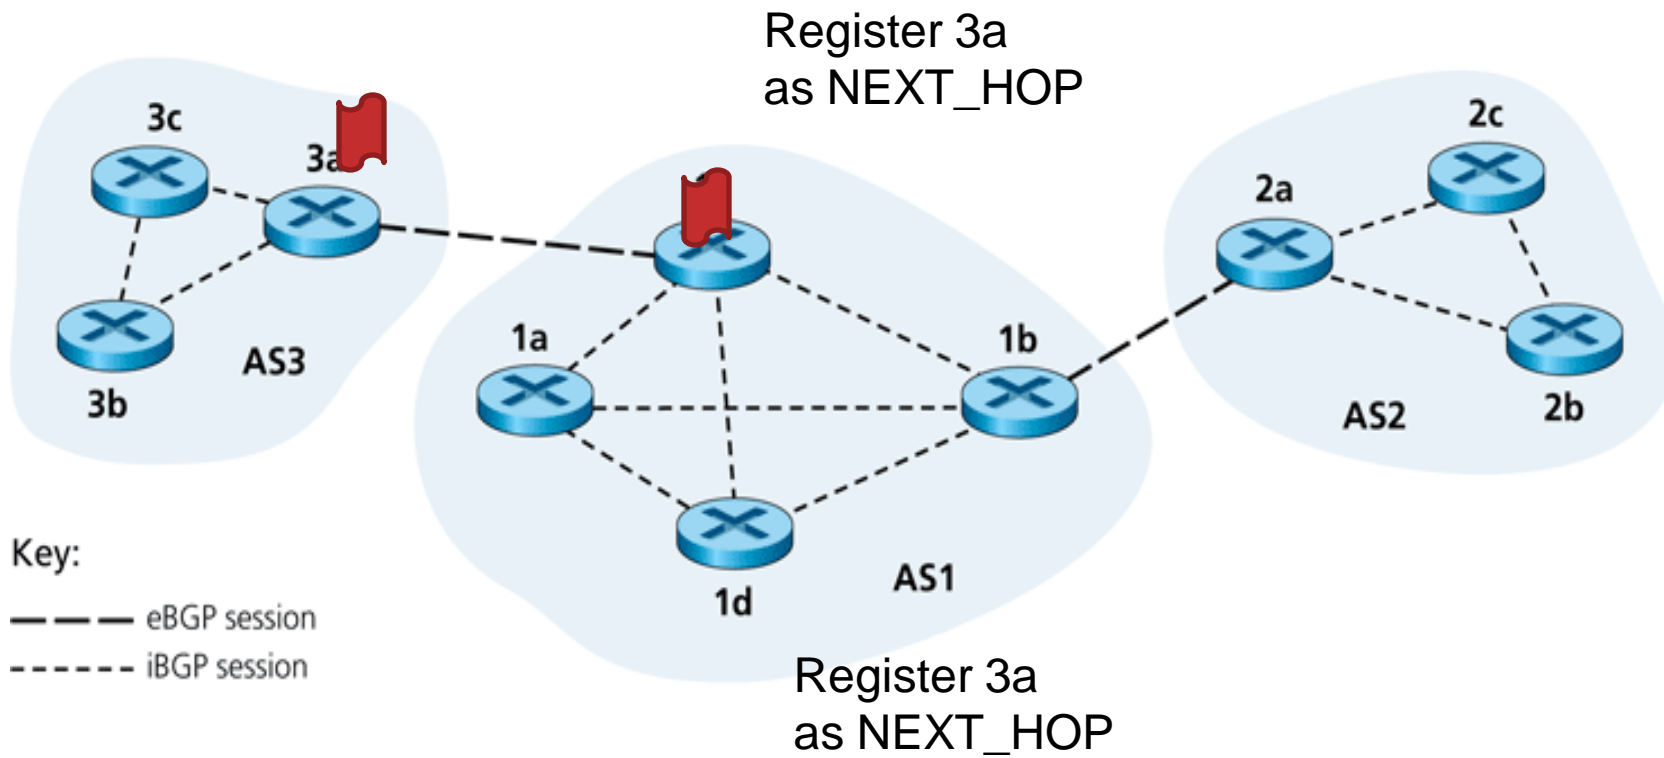
Local updates



BGP Routing Process



Next_Hop



Hot Potato Routing

- Off-load inter-domain traffic as fast as possible
 - Find the shortest path to the destination domain

Hot-potato routing,
Use least resource usage within domain

