Securing IXP Connectivity

Mike Jager

synack

APNIC 34
Phnom Penh, Cambodia
August 2012

IXP 101

- Shared layer 2 network
  - (ethernet)
- IXP operator assigns IP addresses
- IXP members stand up BGP sessions
  - between each other - bilateral
  - to route servers (run by IXP operator) - multilateral
- Routes exchanged via BGP
- Packets flow across IXP

IXP 101

- Peering reduces requirement for transit (save $$)
- Keeps local traffic local
  - increased bandwidth between peers
  - reduced latency/jitter
- Results in one port with many peers
- Cheaper/easier than many ports with one peer

IXP attack methods

- Transit: dedicated, private, point-to-point circuit
- Intermediary routers must forward IP packets across the Internet to you - ie, attacker must manipulate all intermediate RIBs

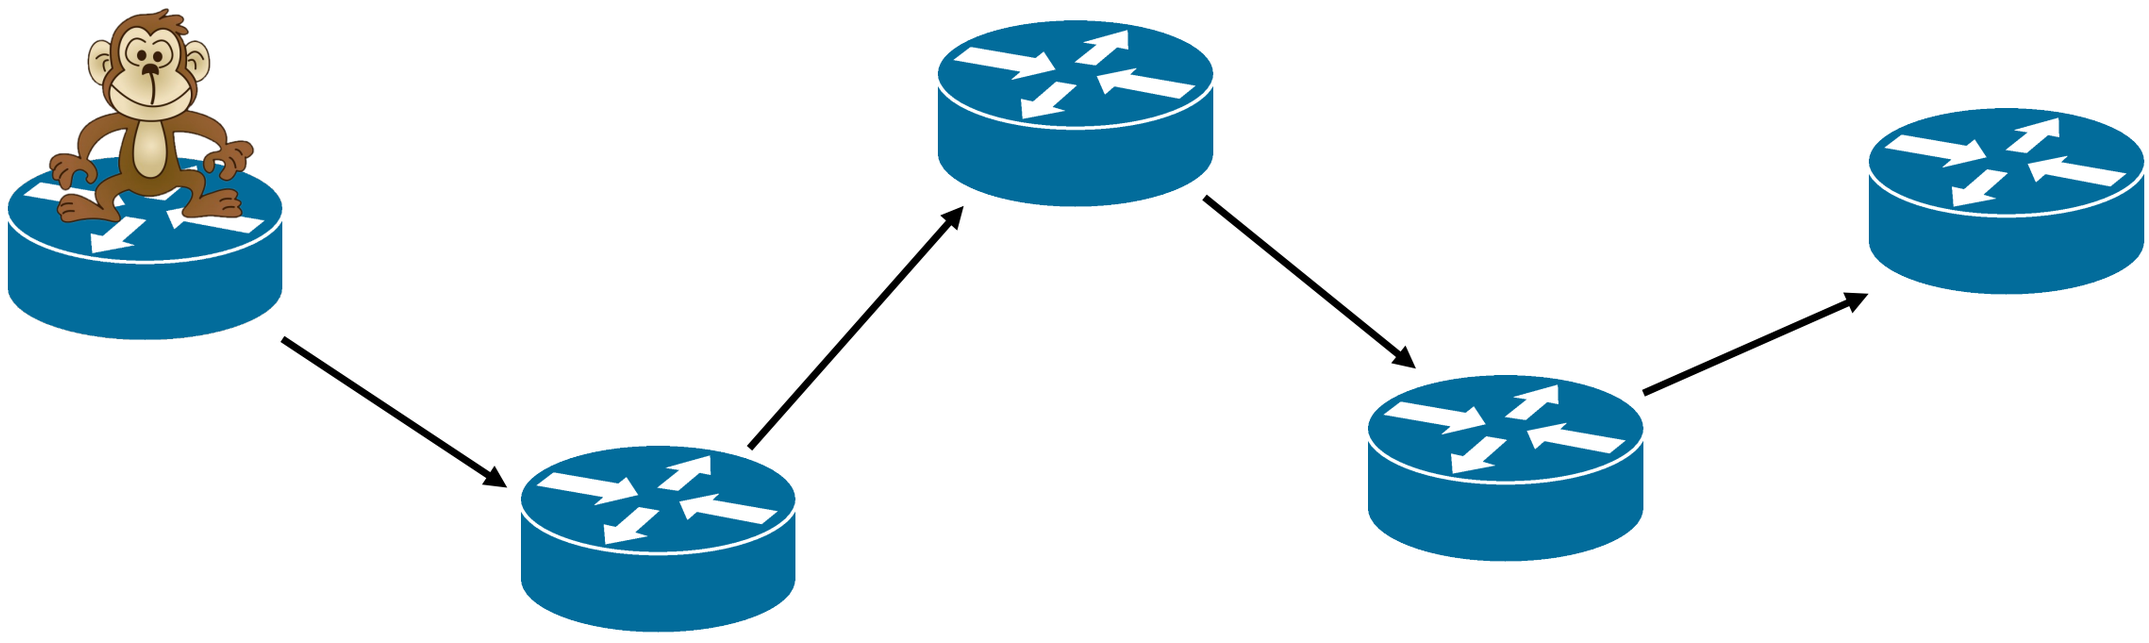

IXP attack methods

- IXP enables any IXP member to send any packet directly into the border of your network

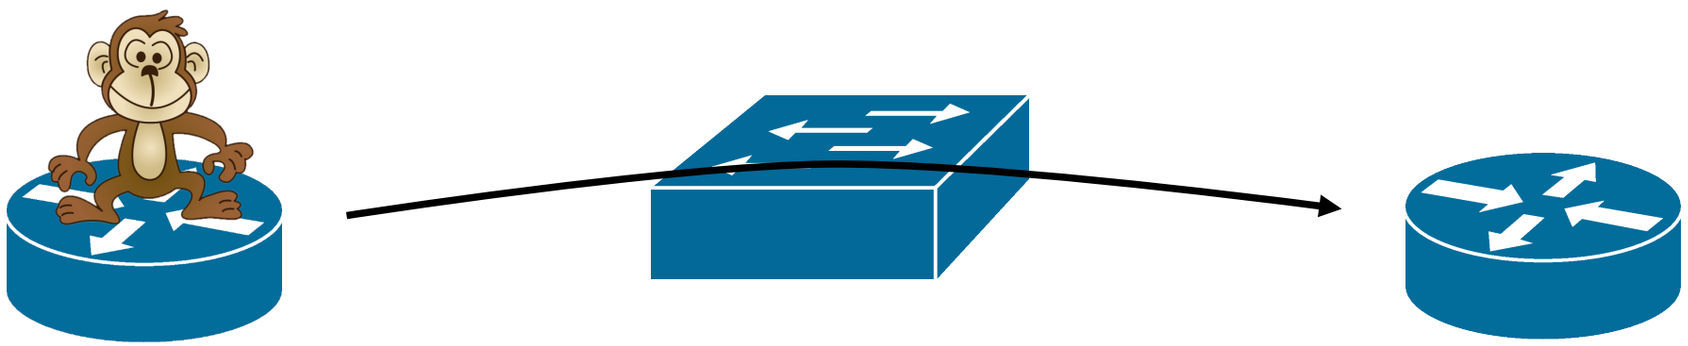

- L2 frame destined to your router's IXP interface MAC address will enter that port
- Routers attached to shared layer 2 networks are more vulnerable to receiving malicious packets

IXP attack methods

- IXP enables any IXP member to see any packet leaked by your network into the IXP

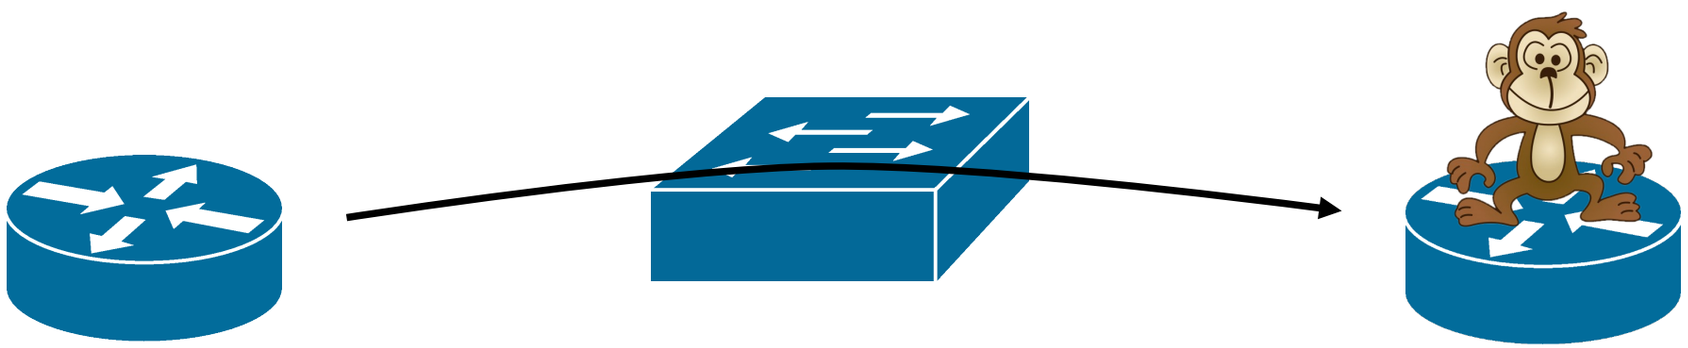

- Routers attached to shared layer 2 networks are more at risk of providing valuable information to attackers

IXP attack methods

- Unhygienic routers
- Network capacity theft
  - outbound
  - inbound
  - bi-directional
- BGP manipulation

IXP hygiene

- IXP is a switched ethernet
- Assuming unicast exchange traffic only:
  - unicast packets between member networks
  - non-unicast packets required for IXP operation
    - broadcast ARP for IXP IPv4 addresses
    - multicast IPv6 NS/NA for IXP IPv6 addresses

IXP hygiene

- Do not want:
  - auto-config: DHCP, IPv6 SLAAC, broadcast TFTP
  - interior routing protocols: OSPF, IS-IS, EIGRP, RIP
  - layer 2 bits: STP, leaking >1 MAC, VTP, keepalive
  - layer 3 bits: proxy ARP, ICMP redirects
  - multicast: PIM, IGMP, MLD
  - network discovery: CDP, LLDP, EDP
  - ?!?!: DECNET MOP

IXP hygiene

- DHCP/SLAAC/IGP: manipulation of RIB
- TFTP: configuration replacement on reload
- layer 2: VLAN database manipulation; trigger IXP port security -> disabled port
- proxy ARP/ICMP redirect: (probably) unintended manipulation of other members' next hop selection
- CDP/LLDP/EDP: reconnaissance of your network
- legacy on-by-default protocols: indicative of poor management practice

IXP hygiene - proxy ARP

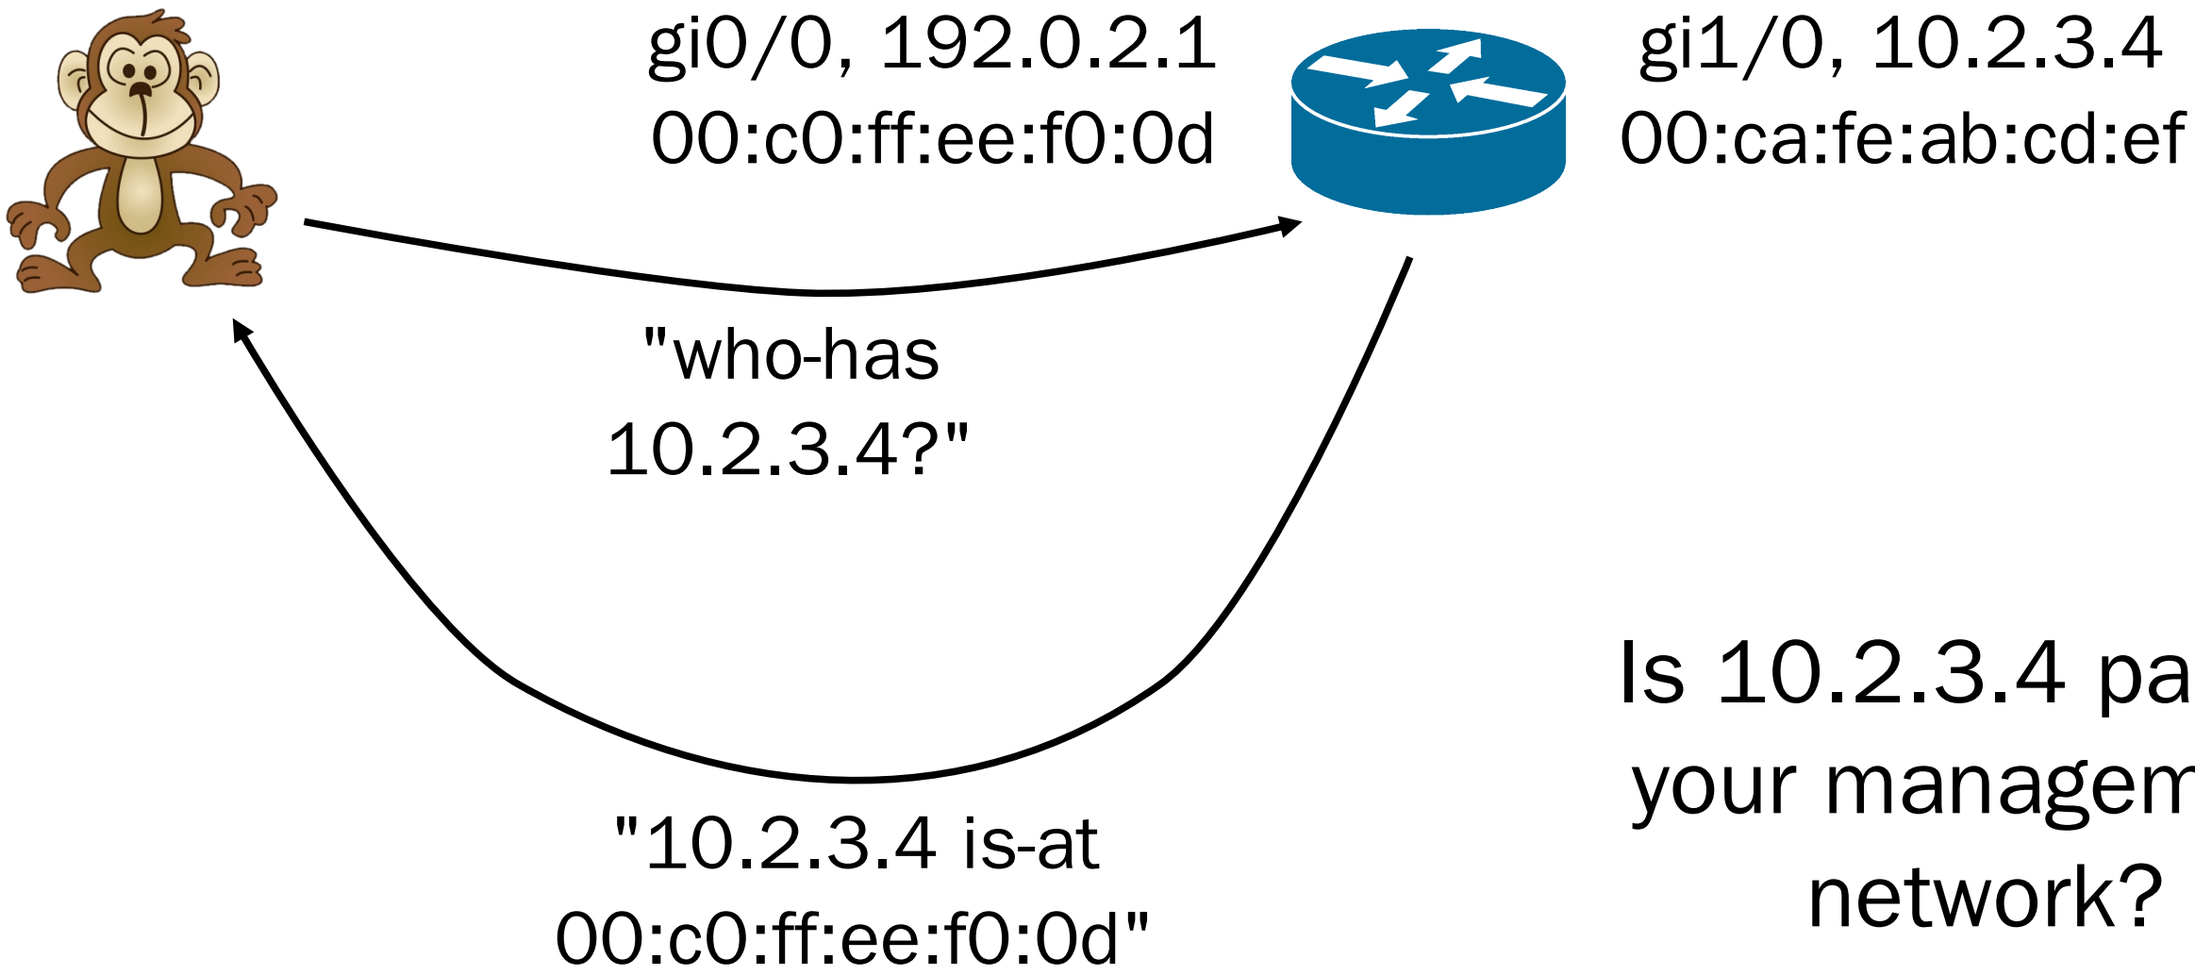

Is 10.2.3.4 part of your management network?

Outbound theft

- Normally, "best" path from RIB installed into FIB
- Next-hop IP address resolved, destination MAC discovered, frame put on wire
- Attackers can ignore RIB and choose arbitrary next-hop

Outbound theft

- Directly into AS Y that wont peer with AS M

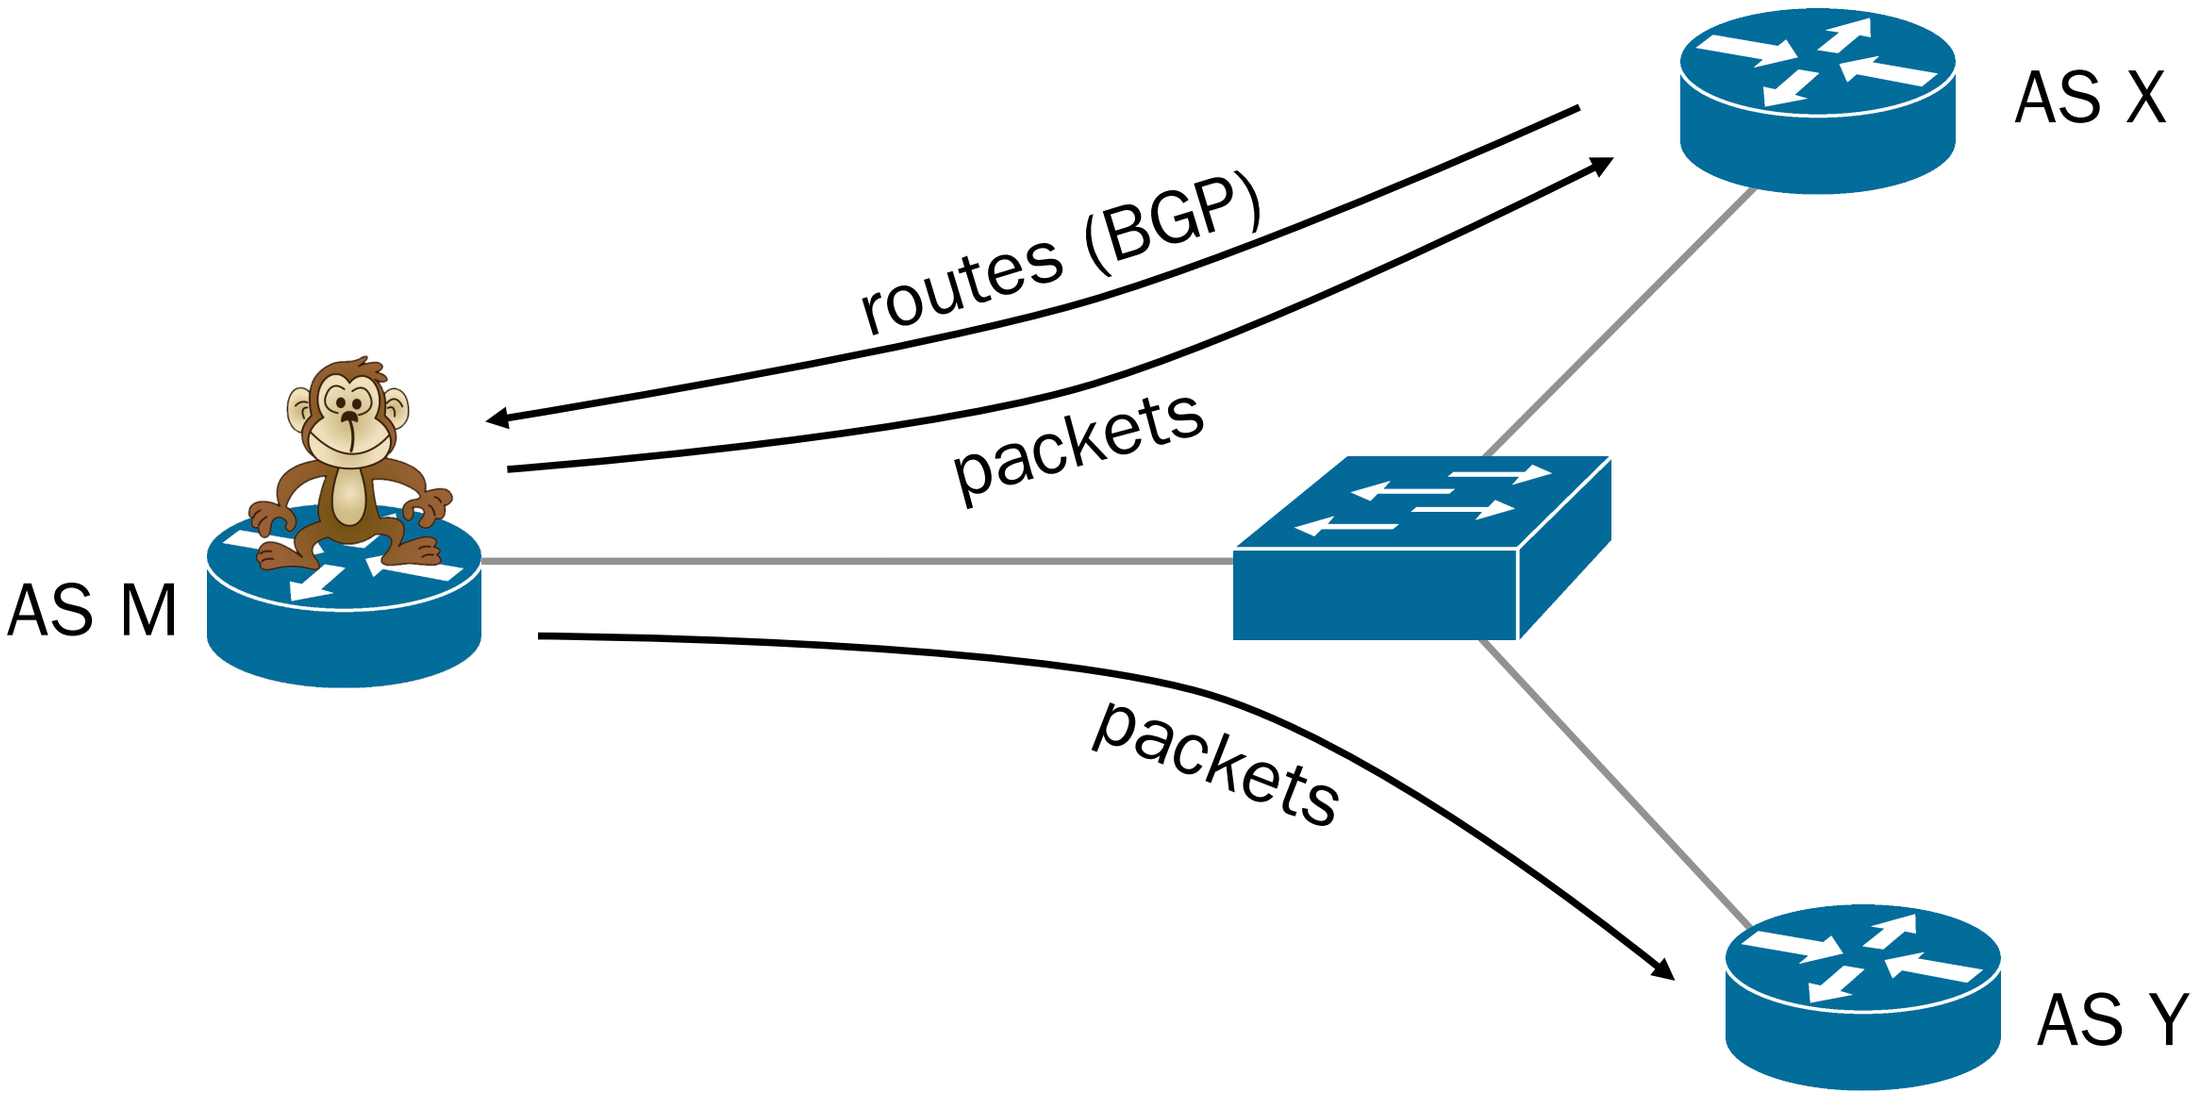

Outbound theft

- Into AS X to reach AS Y that wont peer with AS M

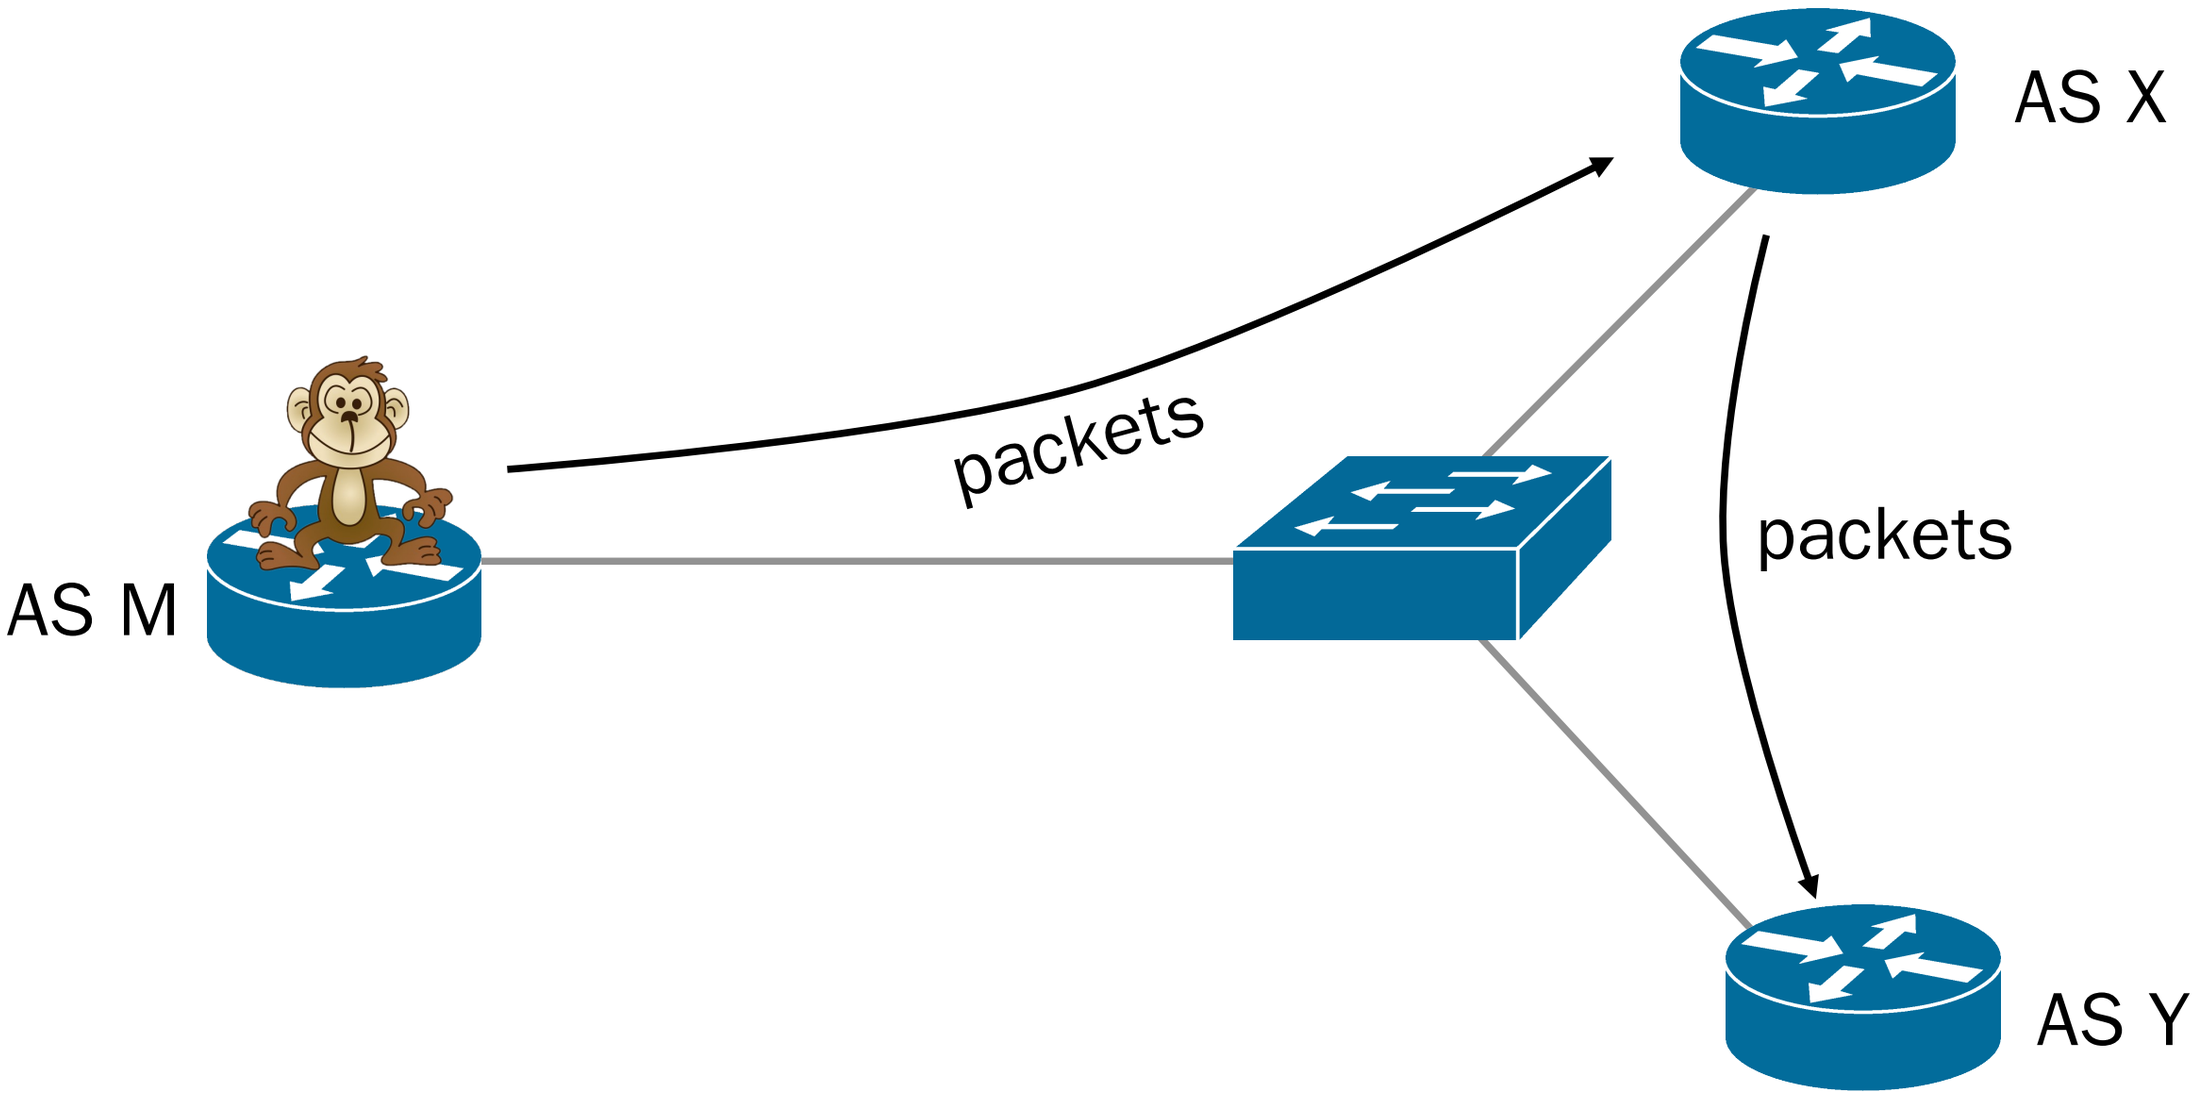

Outbound theft

- Into AS X to reach AS Y that is not present at the same IXP as AS M
  - utilise AS X's private peering link with AS Y

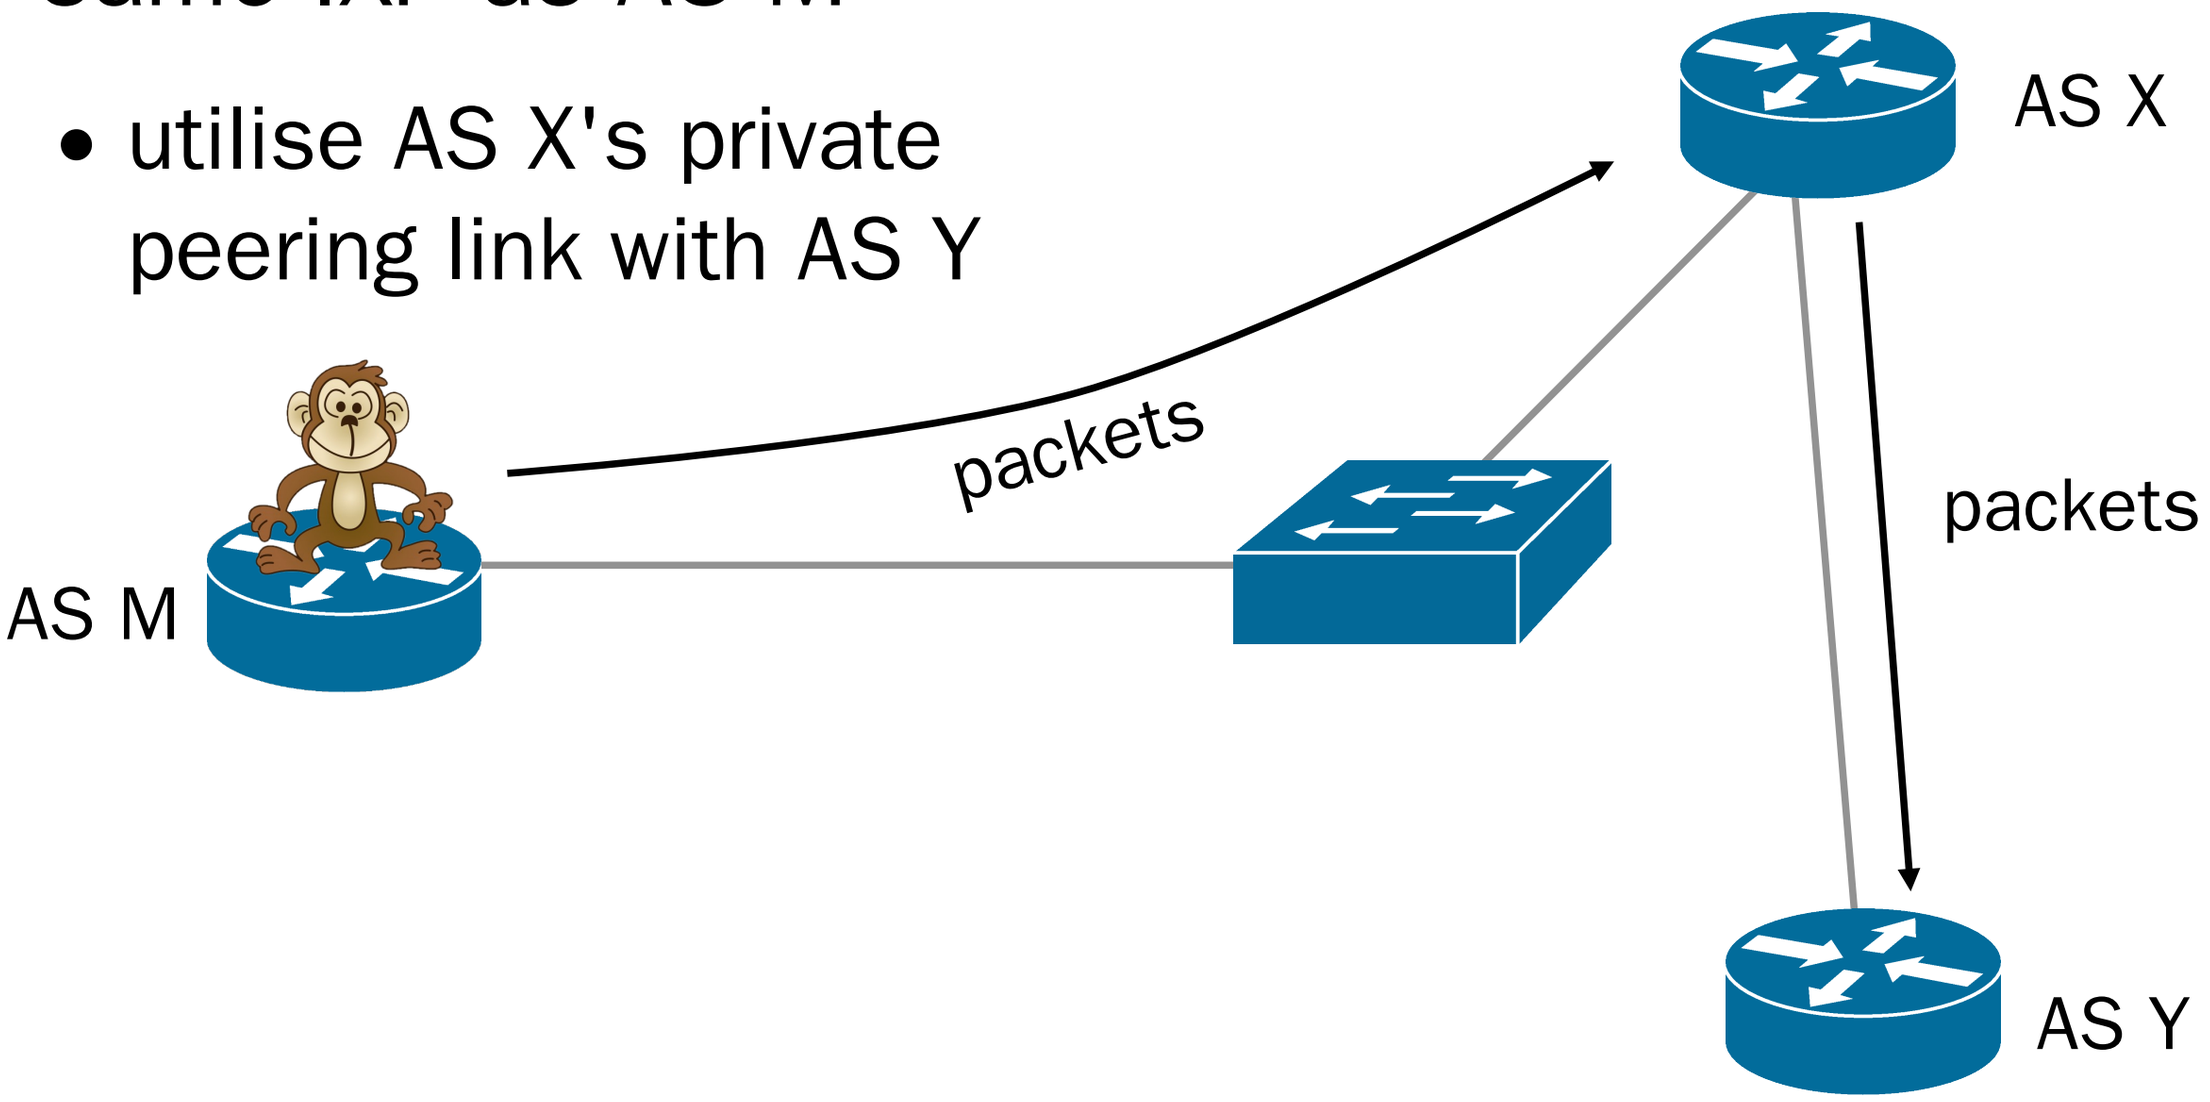

Outbound theft

- Into AS X to reach ASes at a different IXP

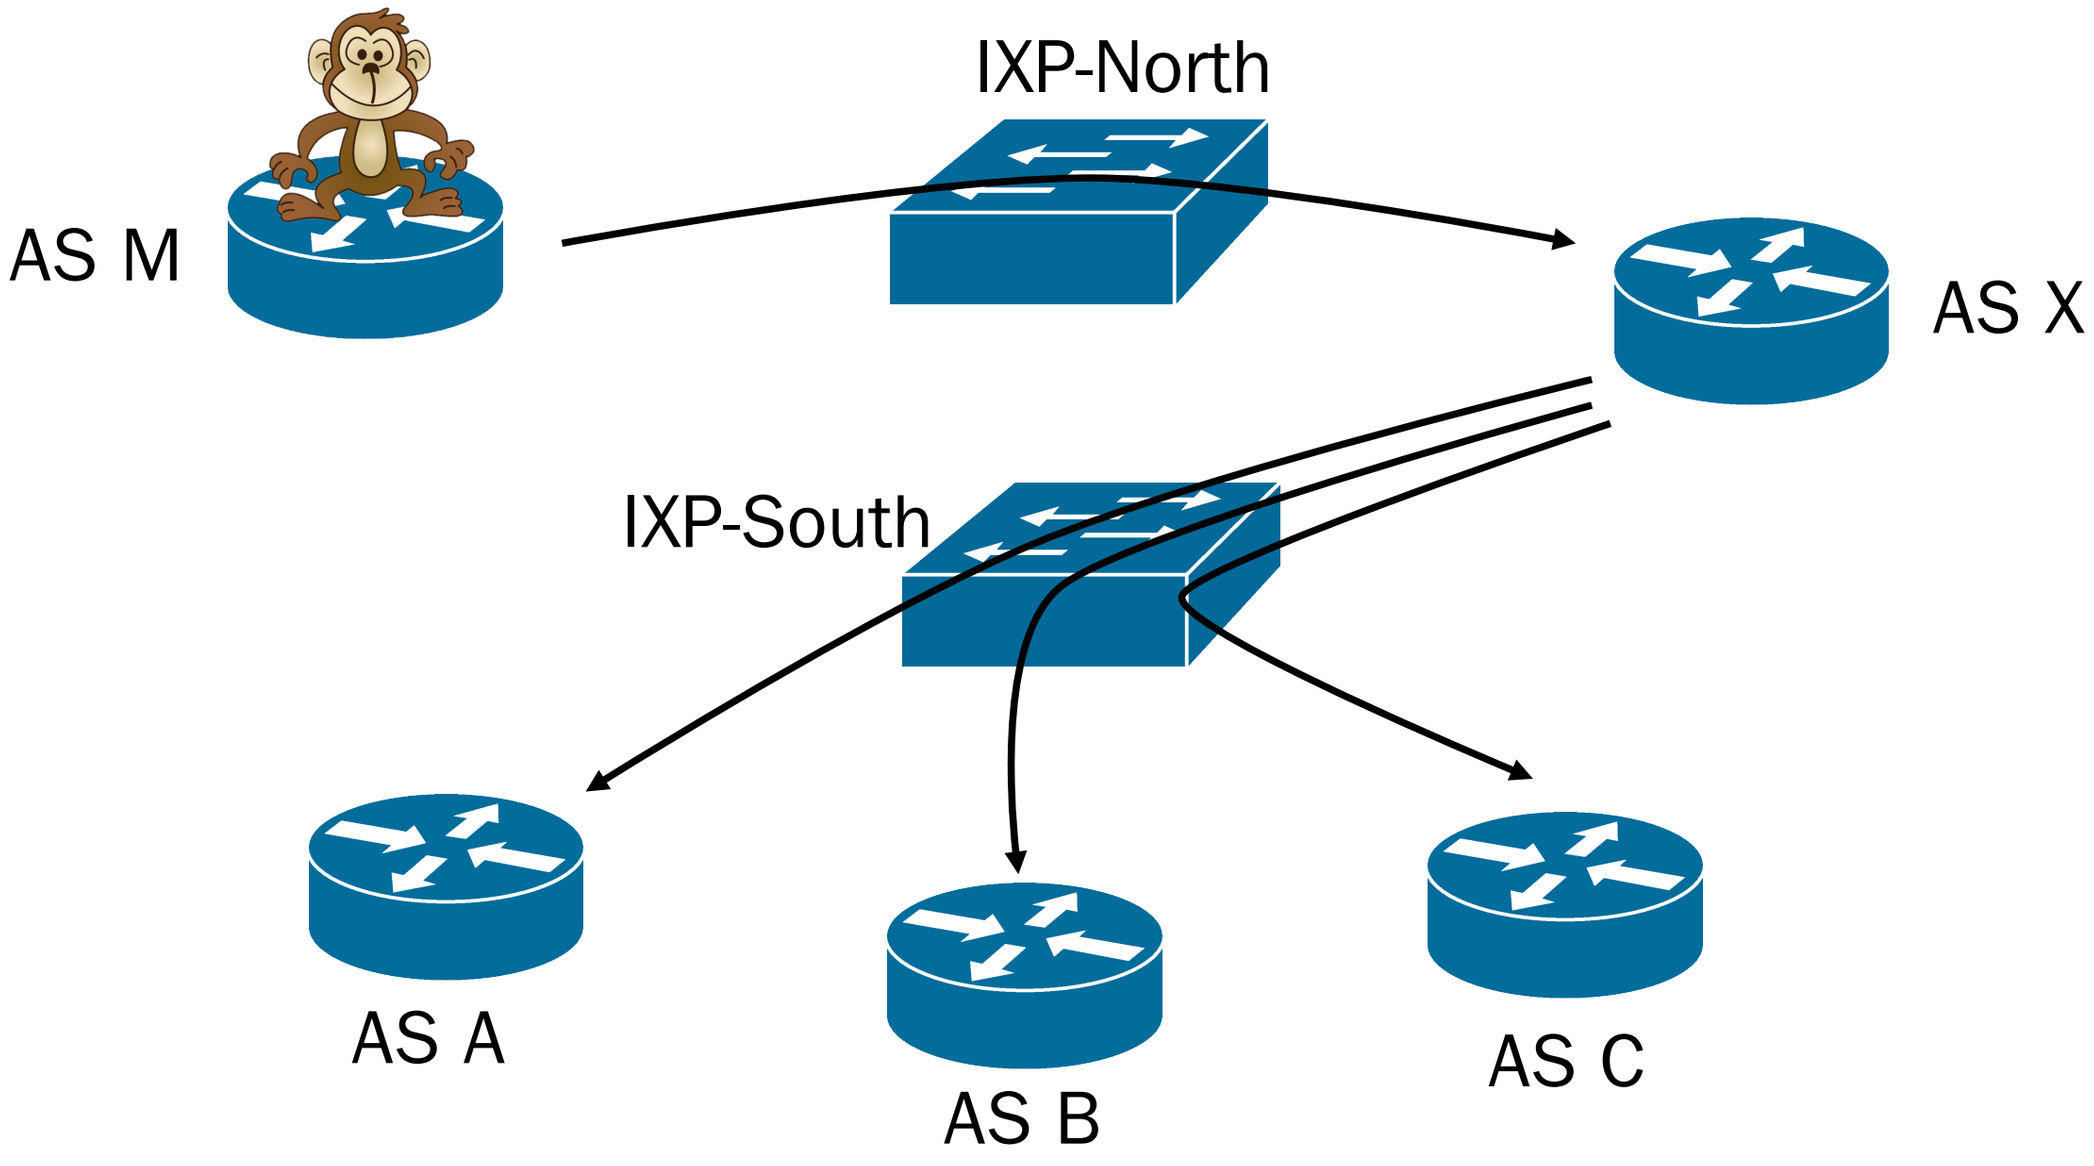

Outbound theft

- Into AS X to reach the Internet at large

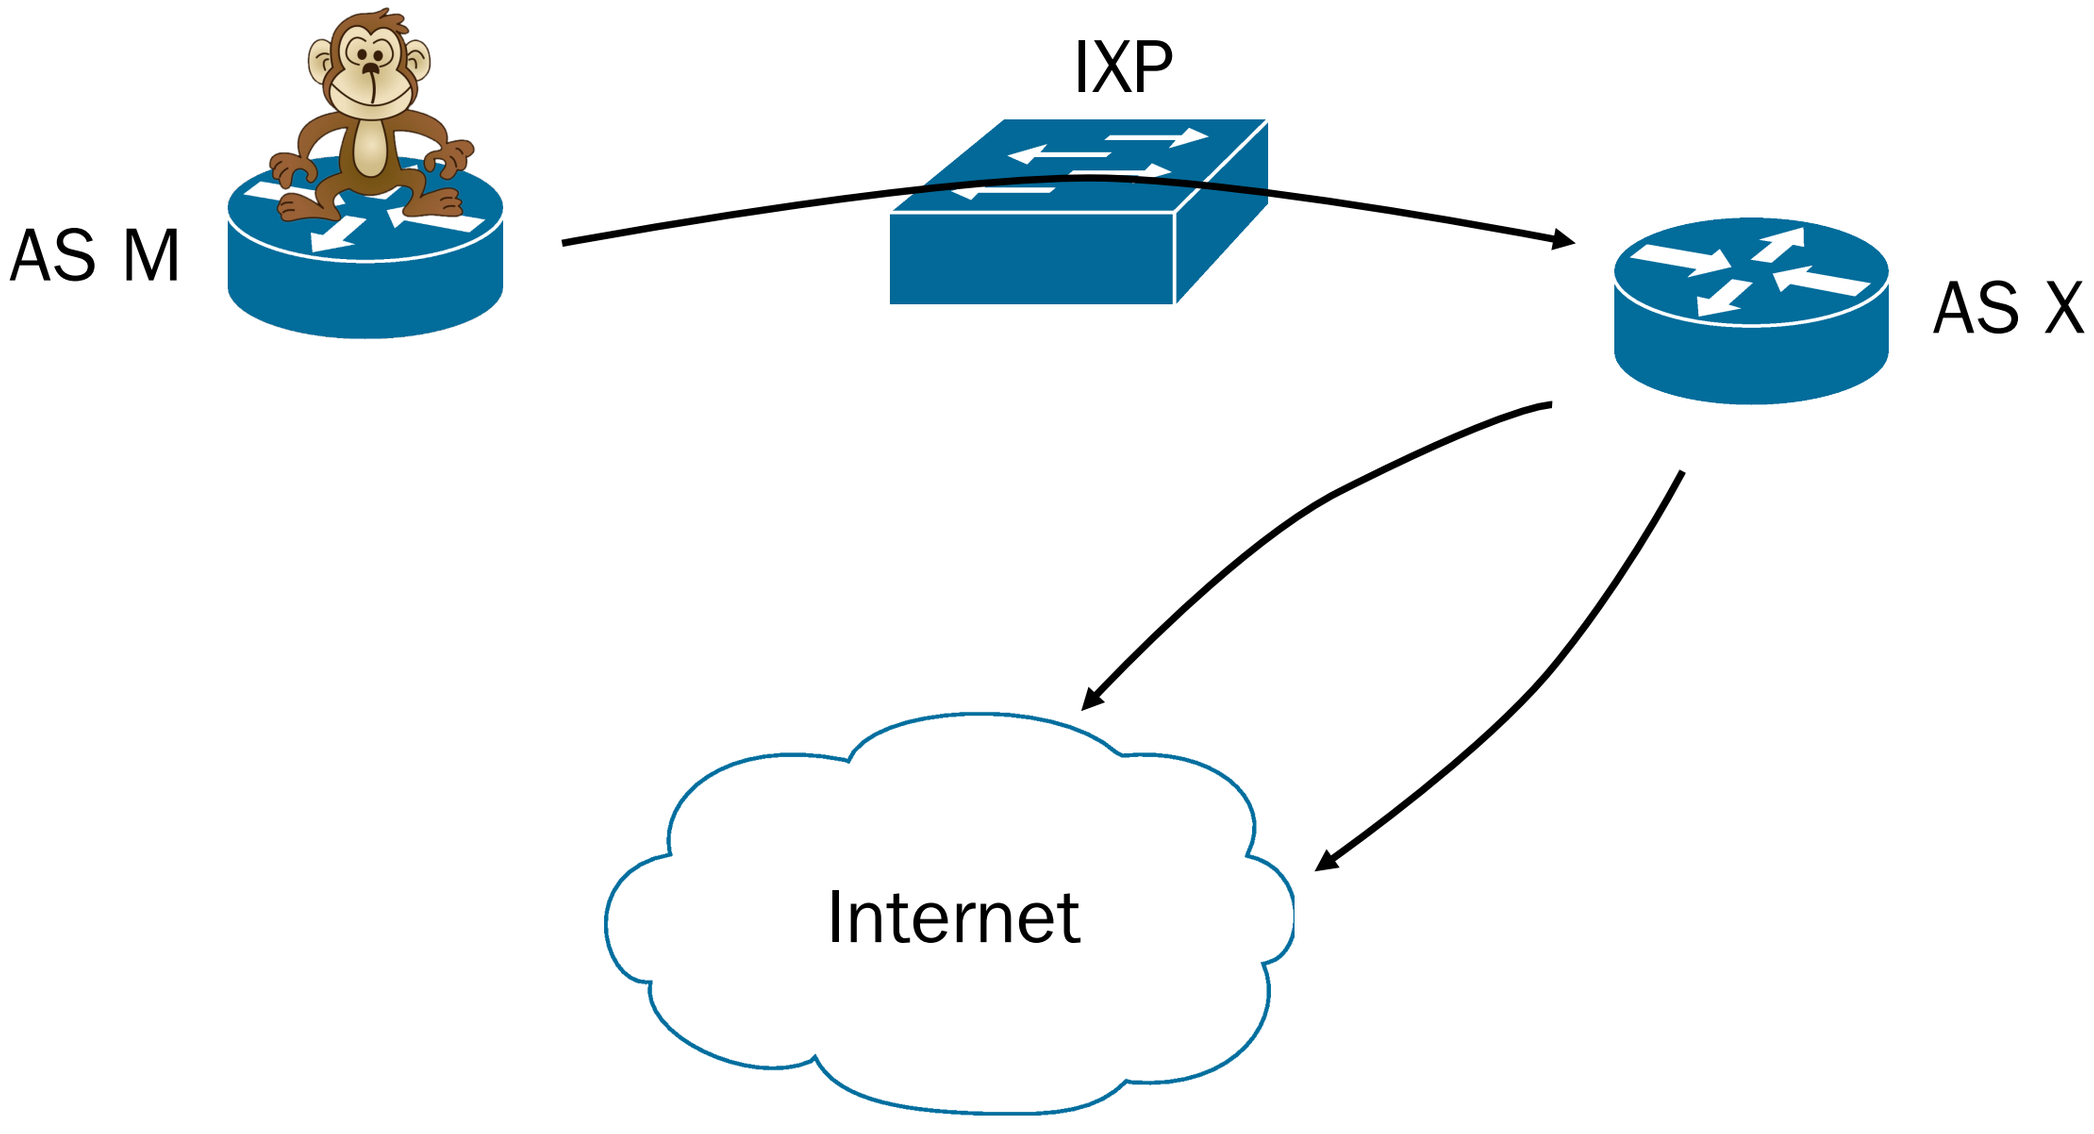

New Zealand

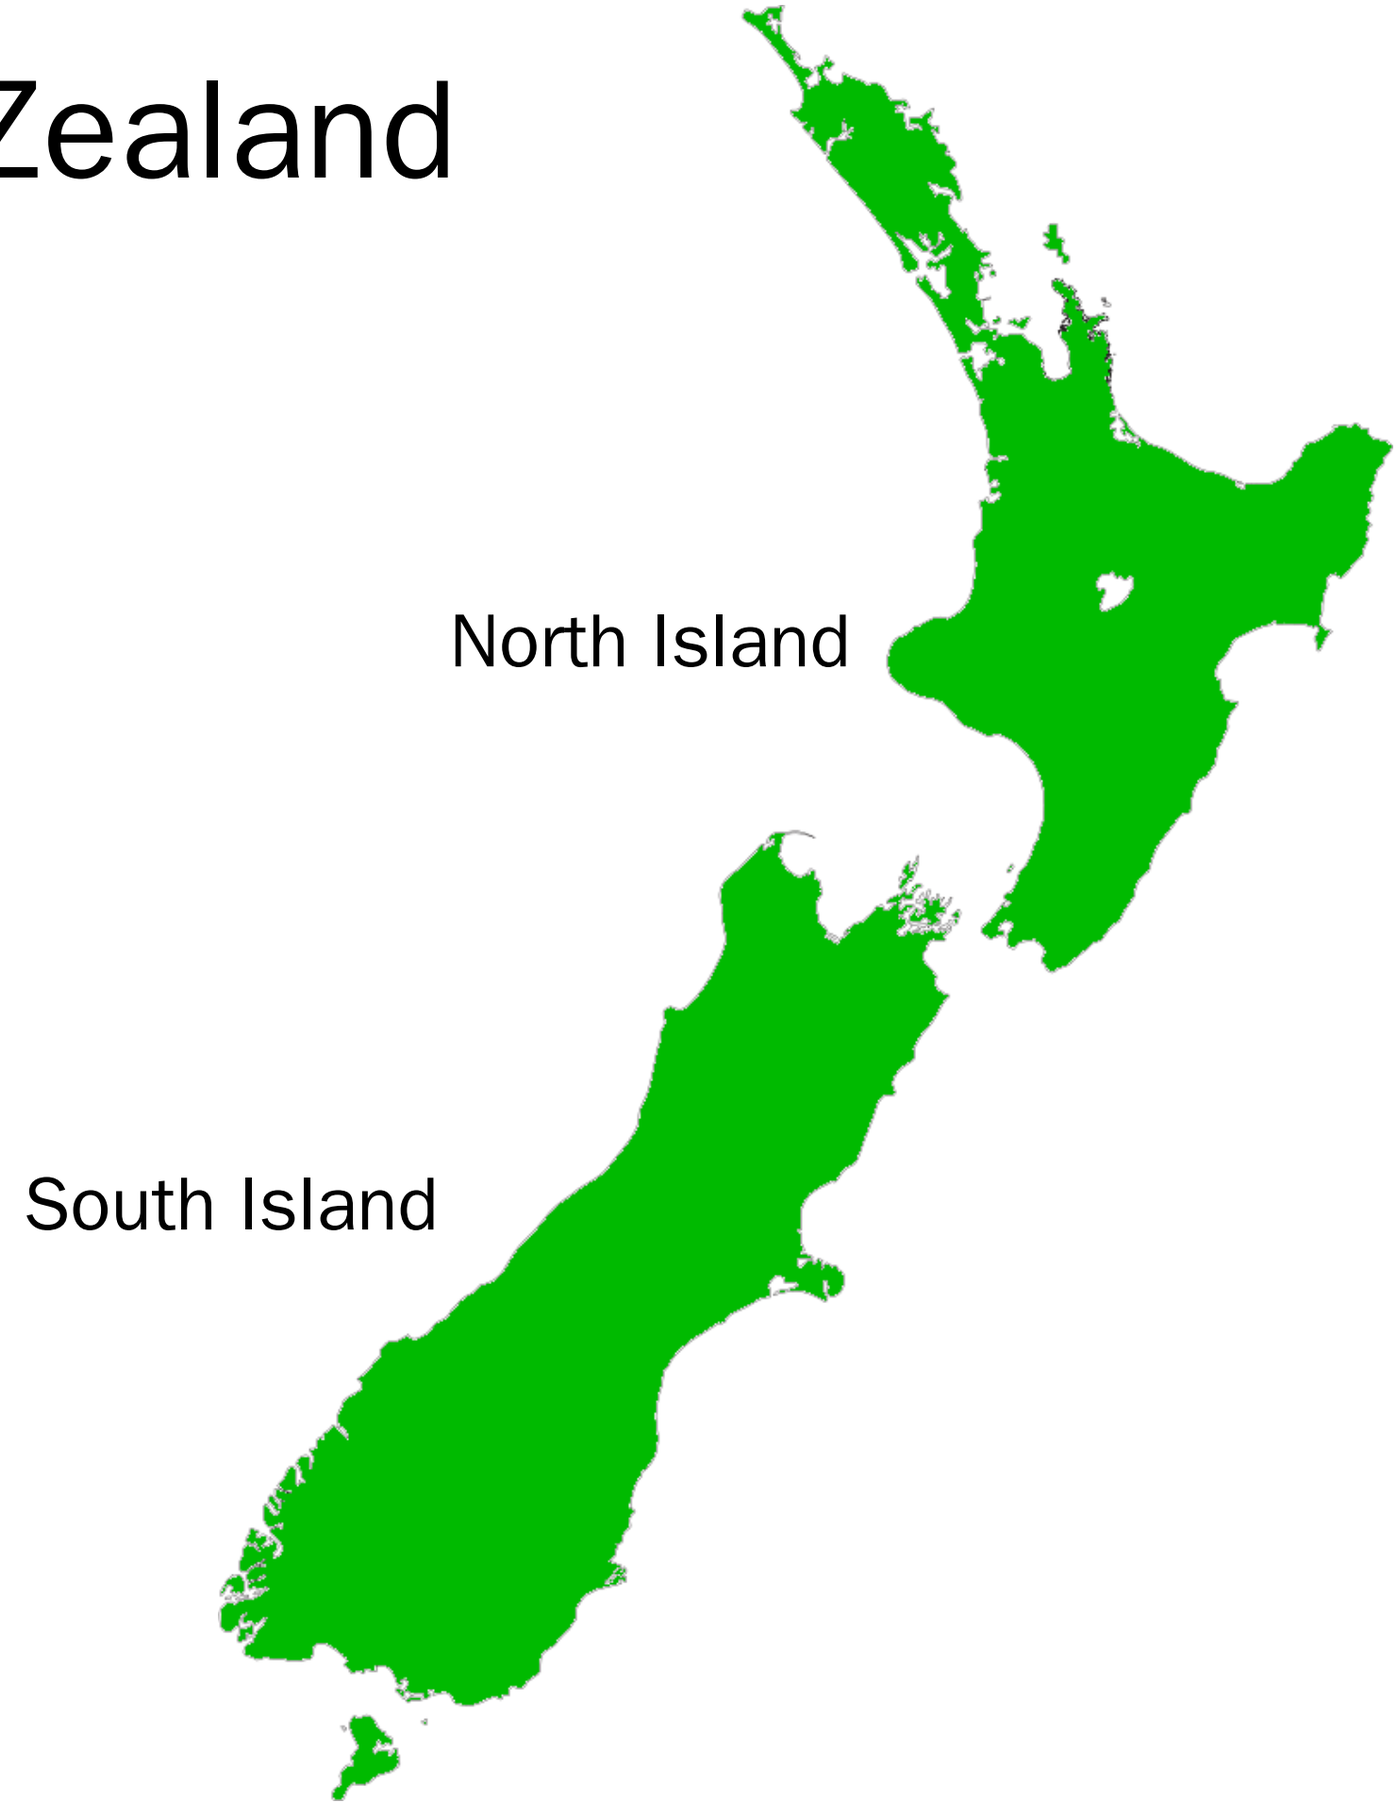

7 IXPs:

- 3CIX
- APE
- CHIX
- DPE
- HIX
- PNIX
- WIX

CITYLINK

www.citylink.co.nz

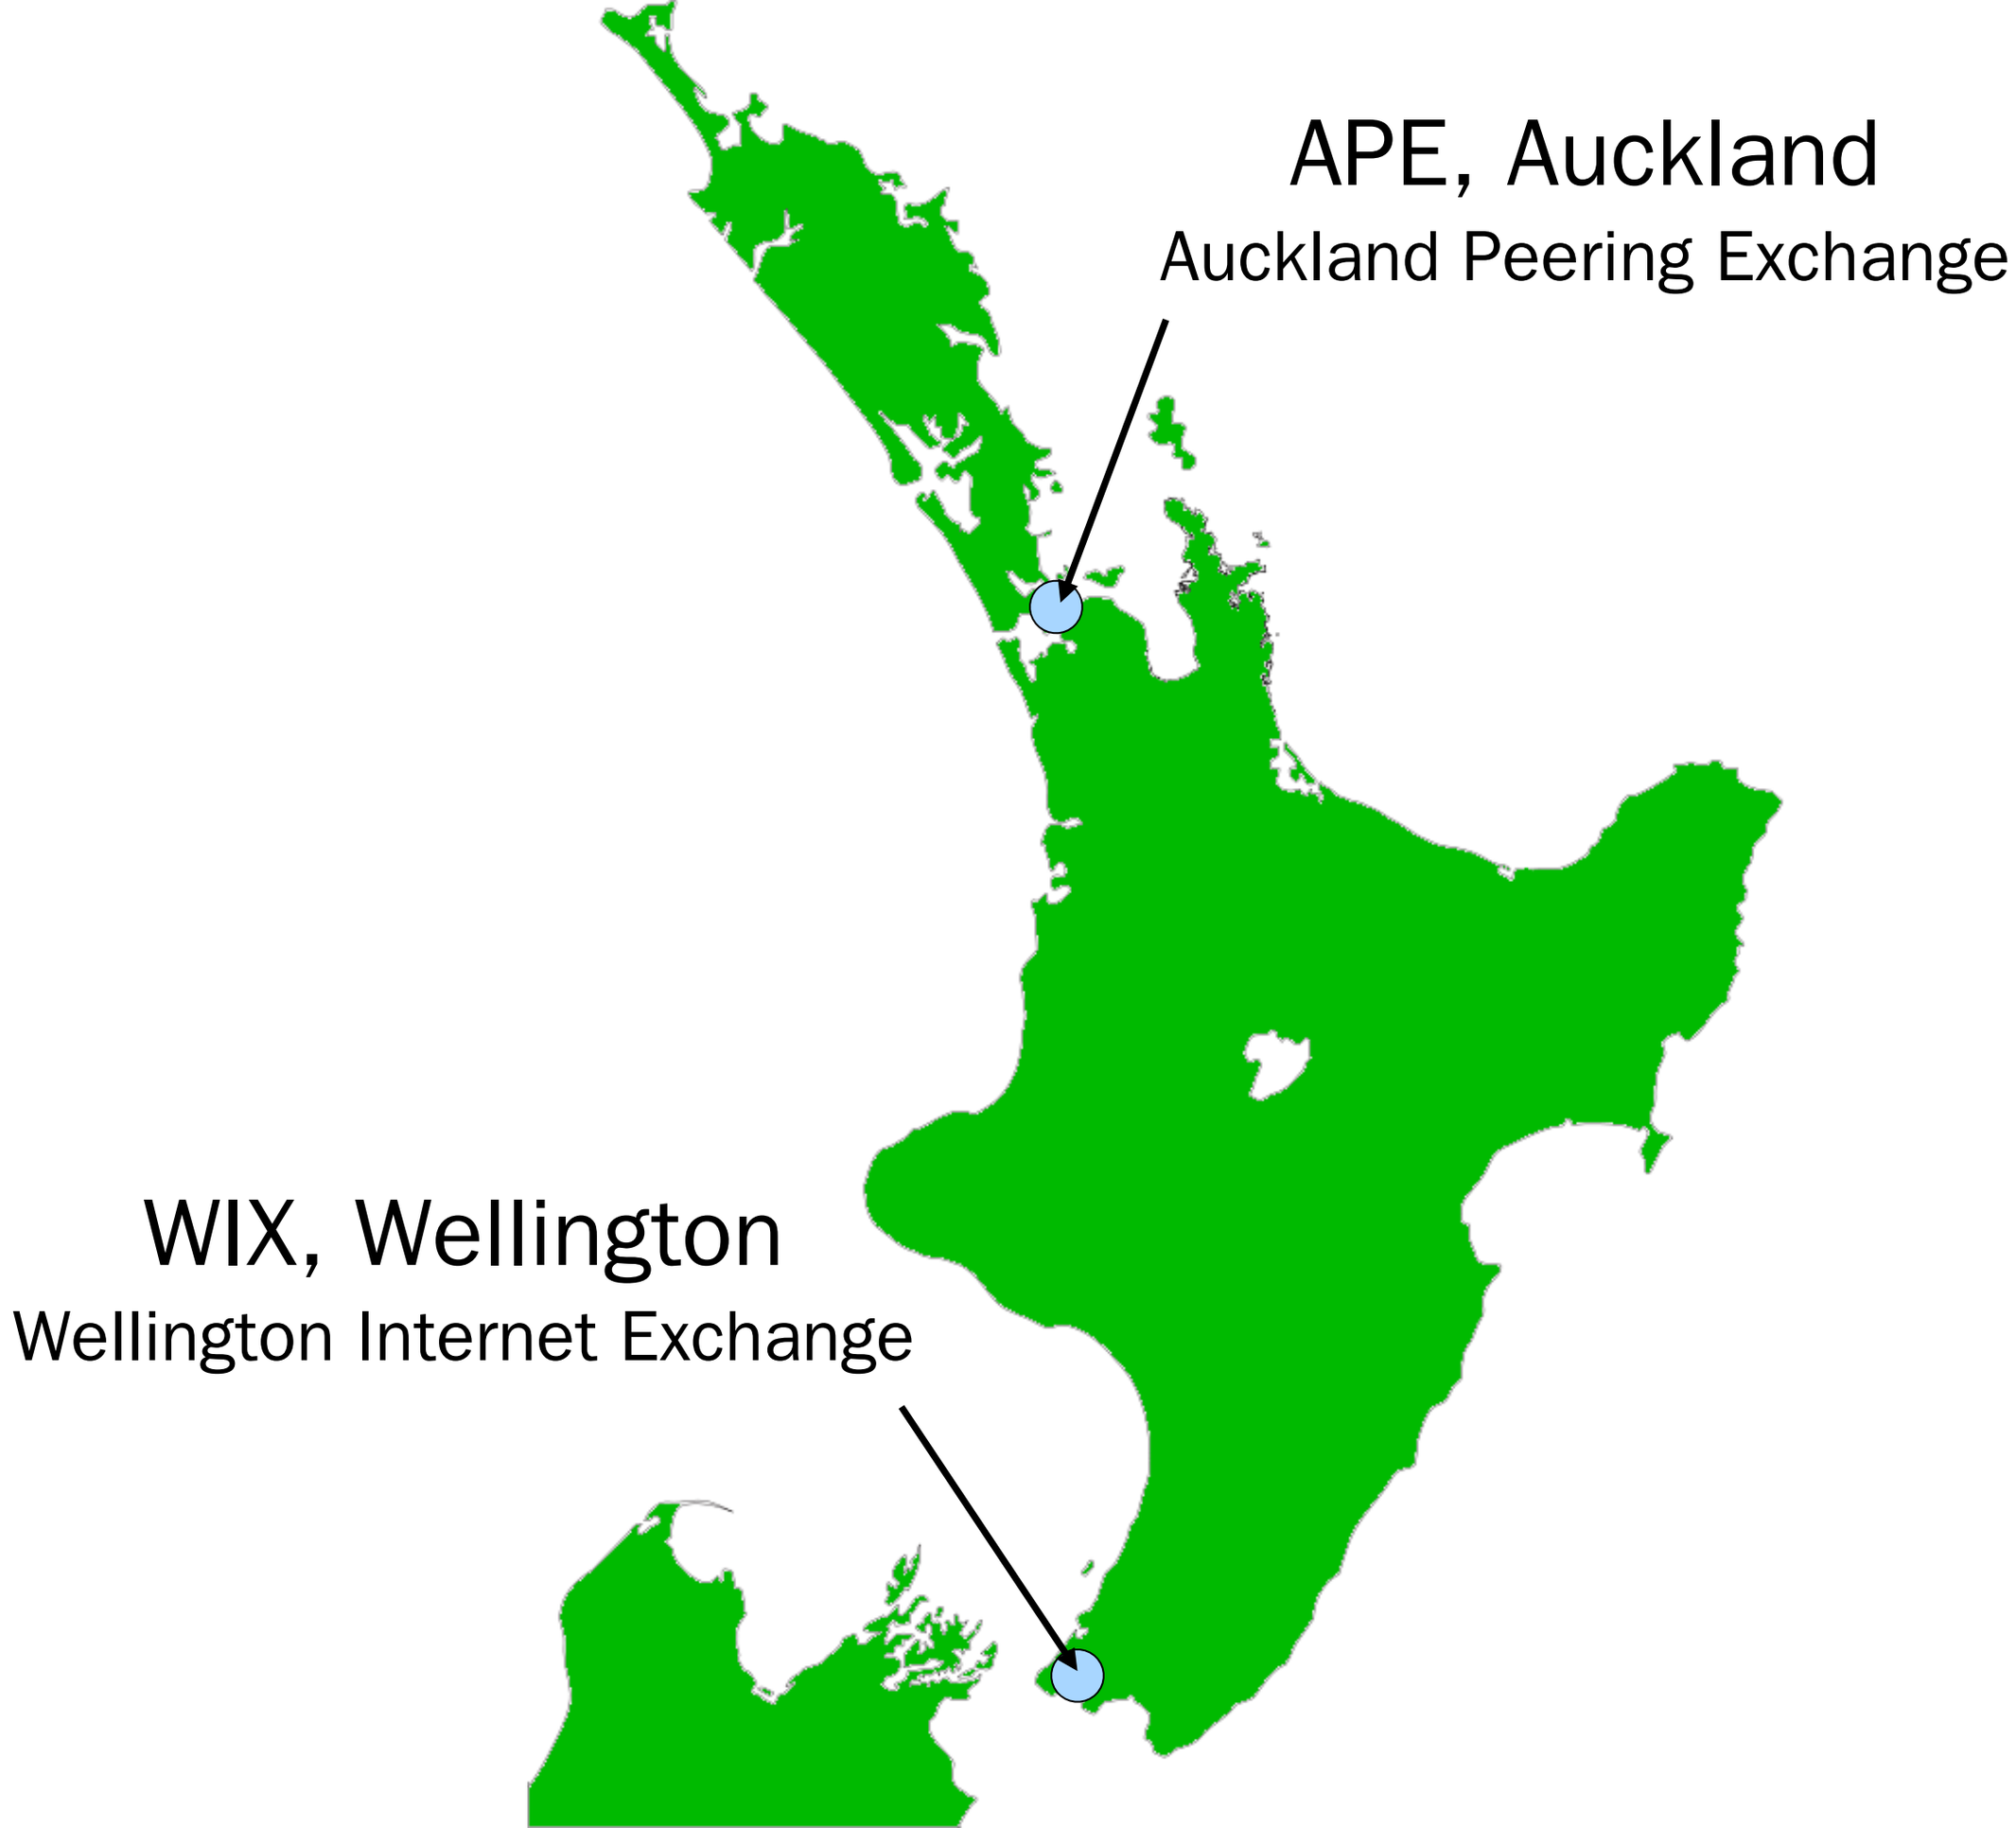
APE, Auckland
Auckland Peering Exchange
WIX, Wellington
Wellington Internet Exchange

Outbound theft

- APE:
  - 90 devices respond to ARP scan of IXP /24
  - 43 hosts carry packet to international destination
- WIX:
  - 112 devices respond to ARP scan of WIX /23
  - 60 hosts carry packet to international destination

Inbound theft

- Packets can be destined to:
  - IXP interface address
  - addresses being advertised via BGP

Inbound theft

- Packets destined to IXP interface address requires far end to have a route to the IXP prefix
- Announcing an IXP prefix across an AS boundary is generally not a good idea
  - eg, if you announce it upstream...
    - and your upstreams announce it to their peers/upstreams
    - who announce it to yet more networks
    - etc
  - ...you're providing free transit for the IXP prefix

Inbound theft

- Prefixes for 31 IXPs originated by 29 ASes

  (still work in progress. IXP data from PeeringDB, RIB data from Route Views LINX)

  - Alpes Adria Internet eXchange
  - Balkan Internet Exchange
  - Bulgarian Internet eXchange
  - Caribbean Internet Exchange
  - Catalunya Neutral Internet Exchange
  - CoreSite - Any2 DC (and Northeast)
  - CoreSite - Any2 Denver
  - CoreSite - Any2 Silicon Valley
  - DRFortress Exchange
  - ECIX Berlin
  - ECIX Duesseldorf
  - Equinix Ashburn Exchange
  - Equinix Internet Exchange New York
  - Equinix Internet Exchange Palo Alto
  - Equinix Los Angeles Exchange
  - Equinix Zurich
  - Espana Internet Exchange
  - Grazer Internet eXchange
  - Groningen Internet Exchange
  - Hong Kong Internet Exchange
  - iAdvantage Internet Exchange
  - London Network Access Point
  - Matrix Cable Internet eXchange
  - Mongolian Internet Exchange
  - NAP Of The Americas
  - Netnod
  - Northwest Access Exchange
  - San Diego NAP
  - Slovenian Internet Exchange
  - Stuttgarter internet eXchange
  - Vienna Internet Exchange

Inbound theft

- Packets destined to IXP interface address requires far end to have a route to the IXP prefix
  - don't announce the IXP prefix outside your AS!
- Packets sourced from non-IXP address requires far end to have a route to that address reachable via the IXP
  - critical to not speak IGP on IXP interfaces
  - manipulate upstream provider's RIB

Inbound theft

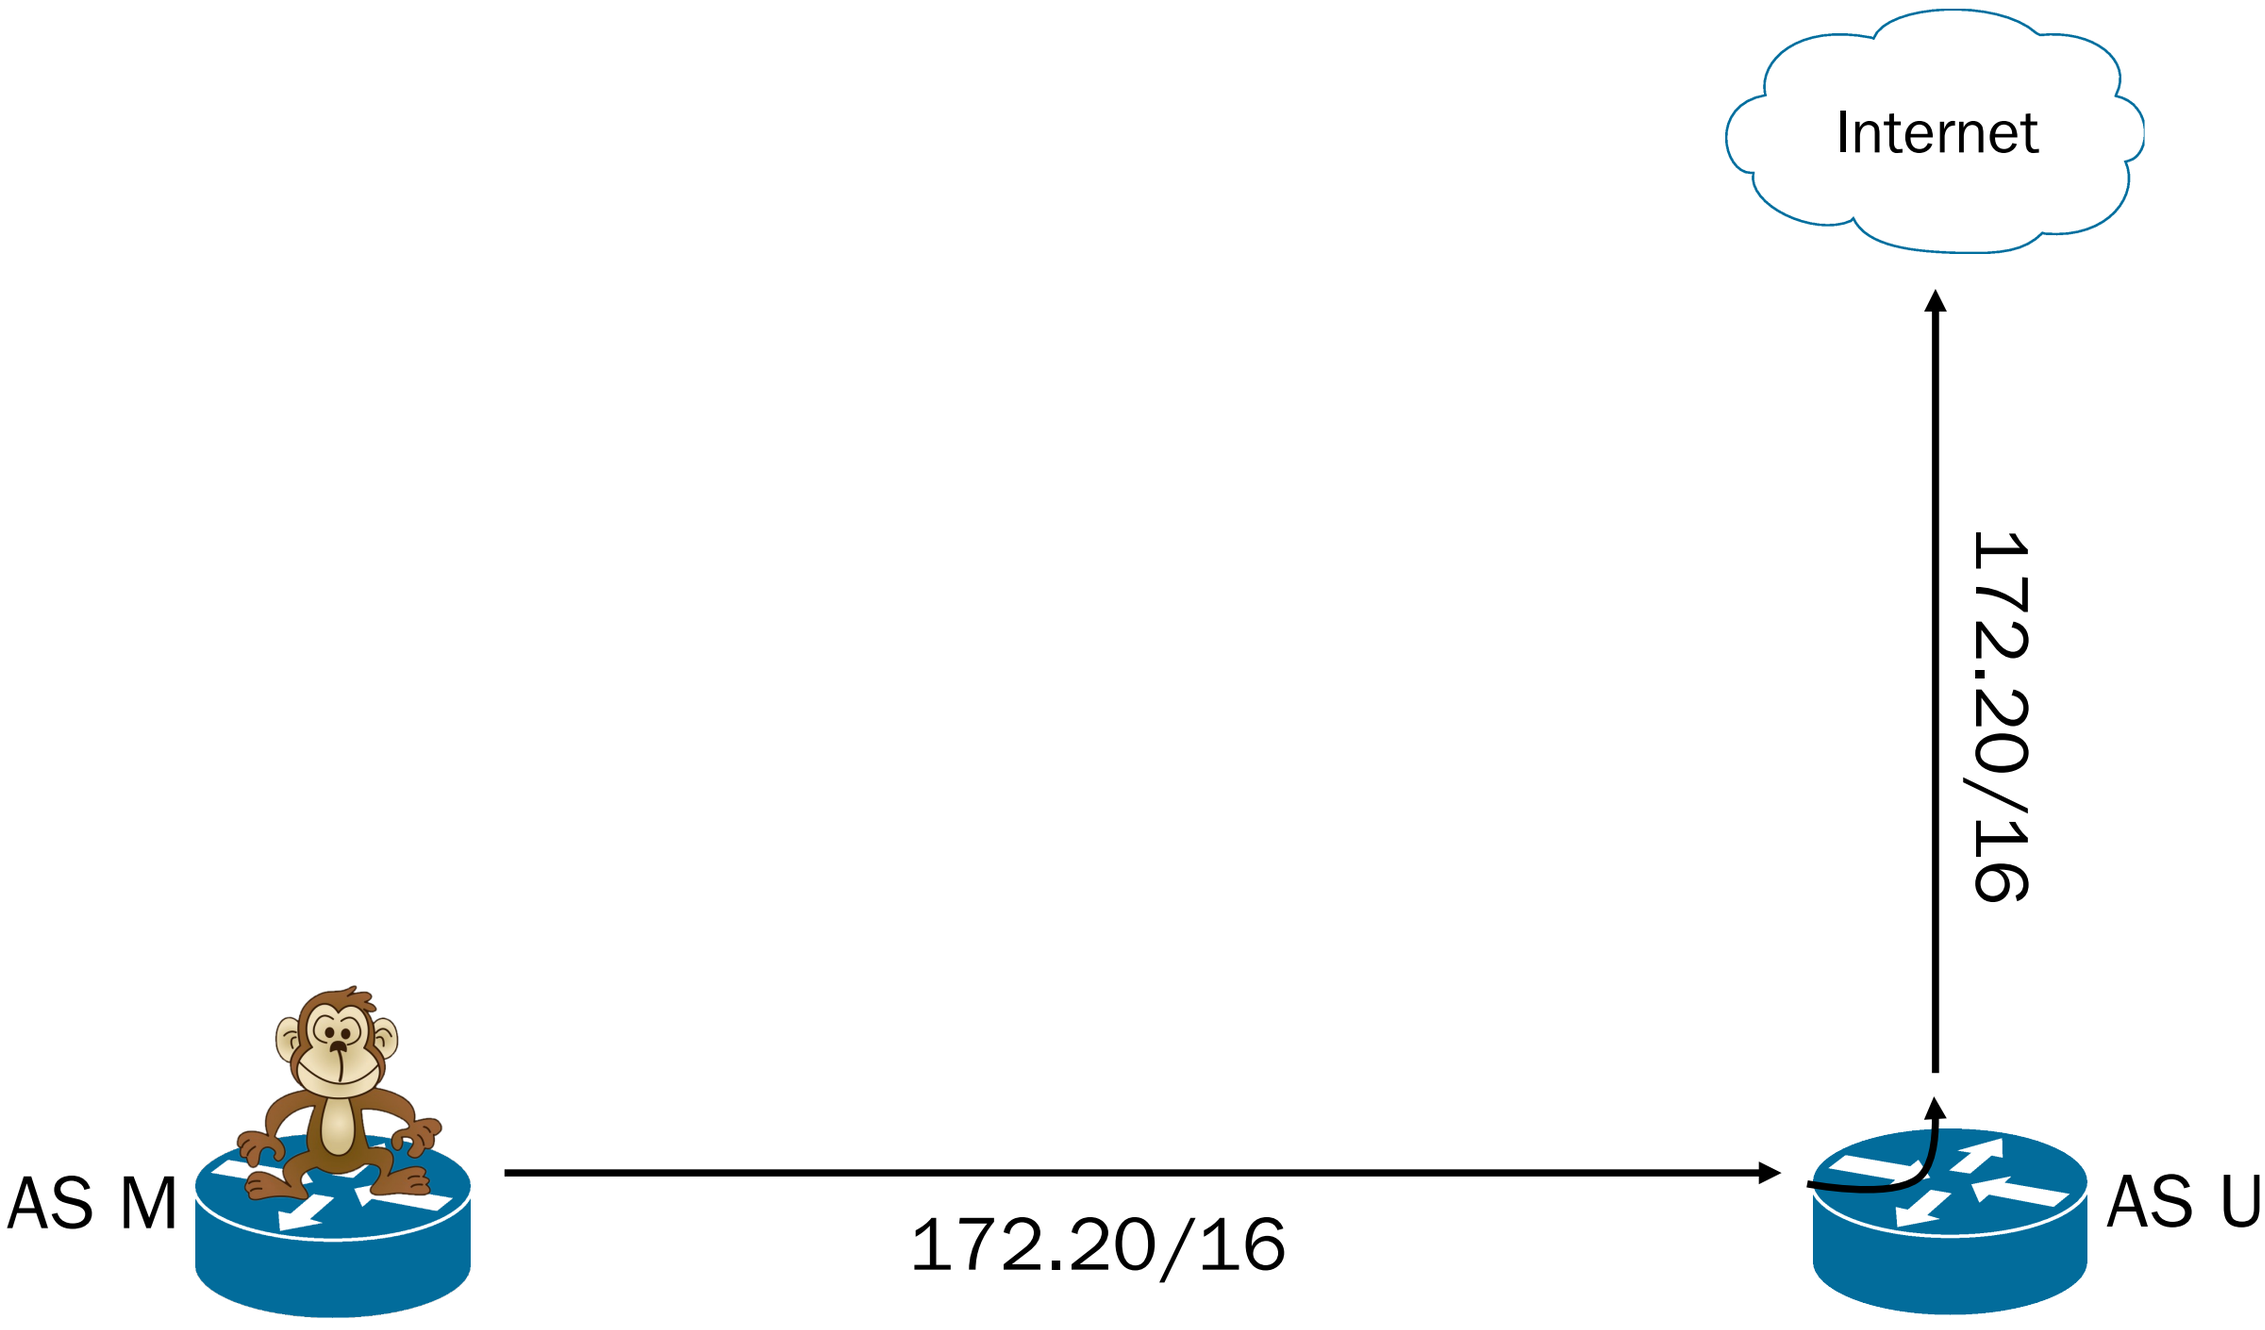

Inbound theft

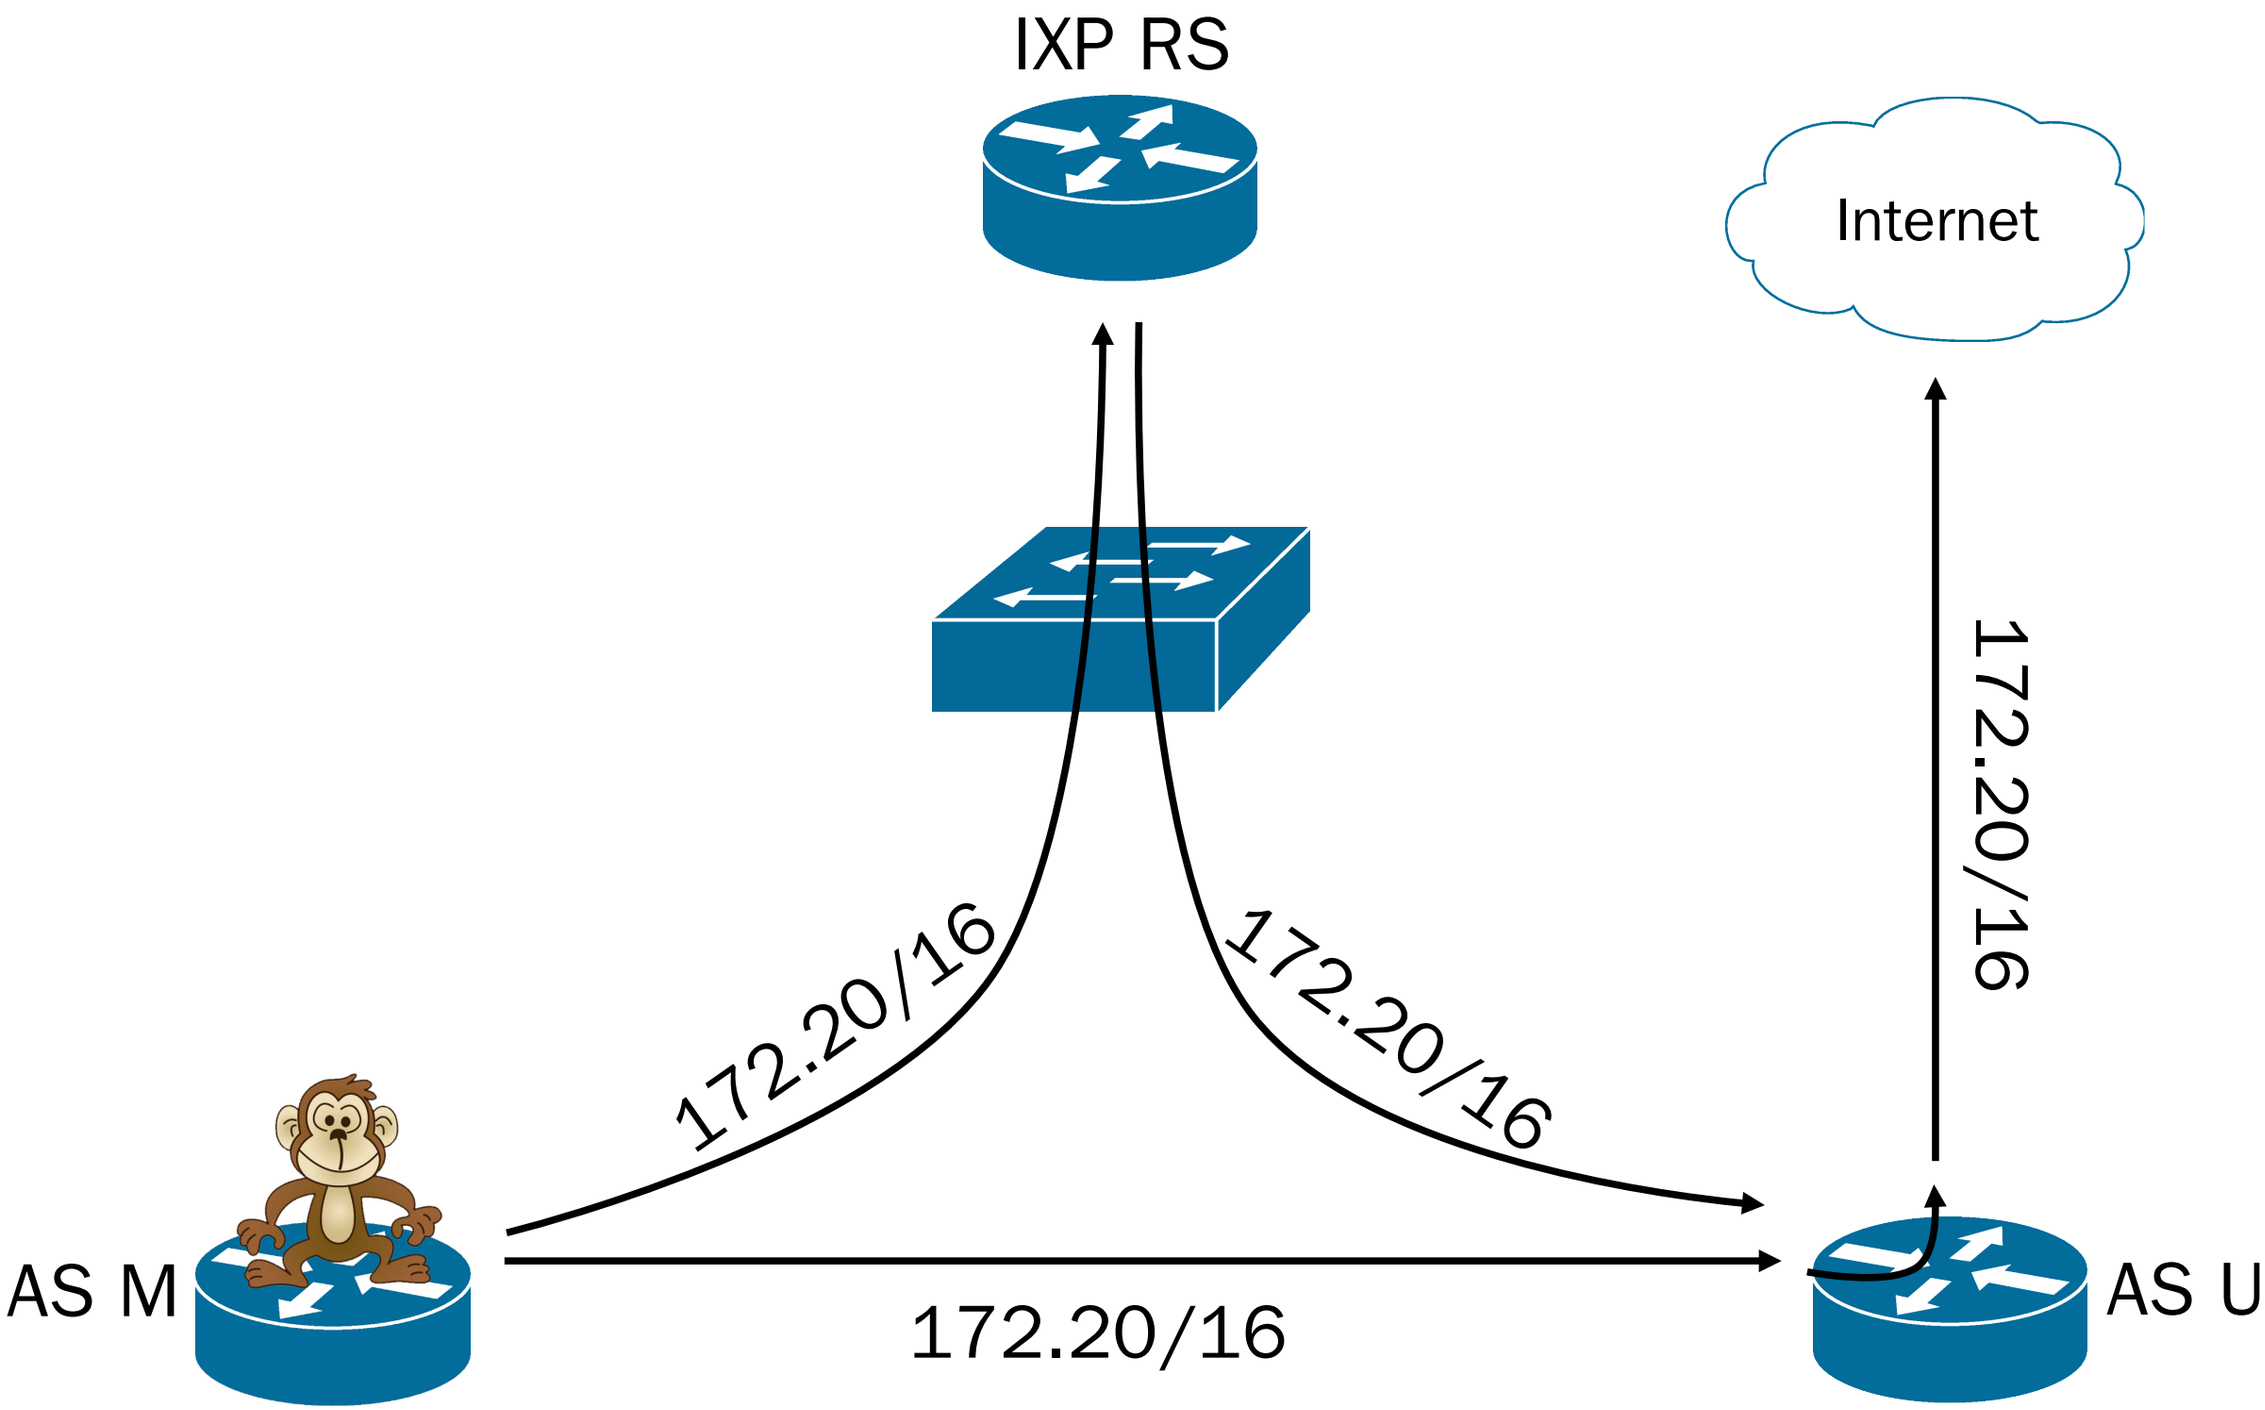

Inbound theft

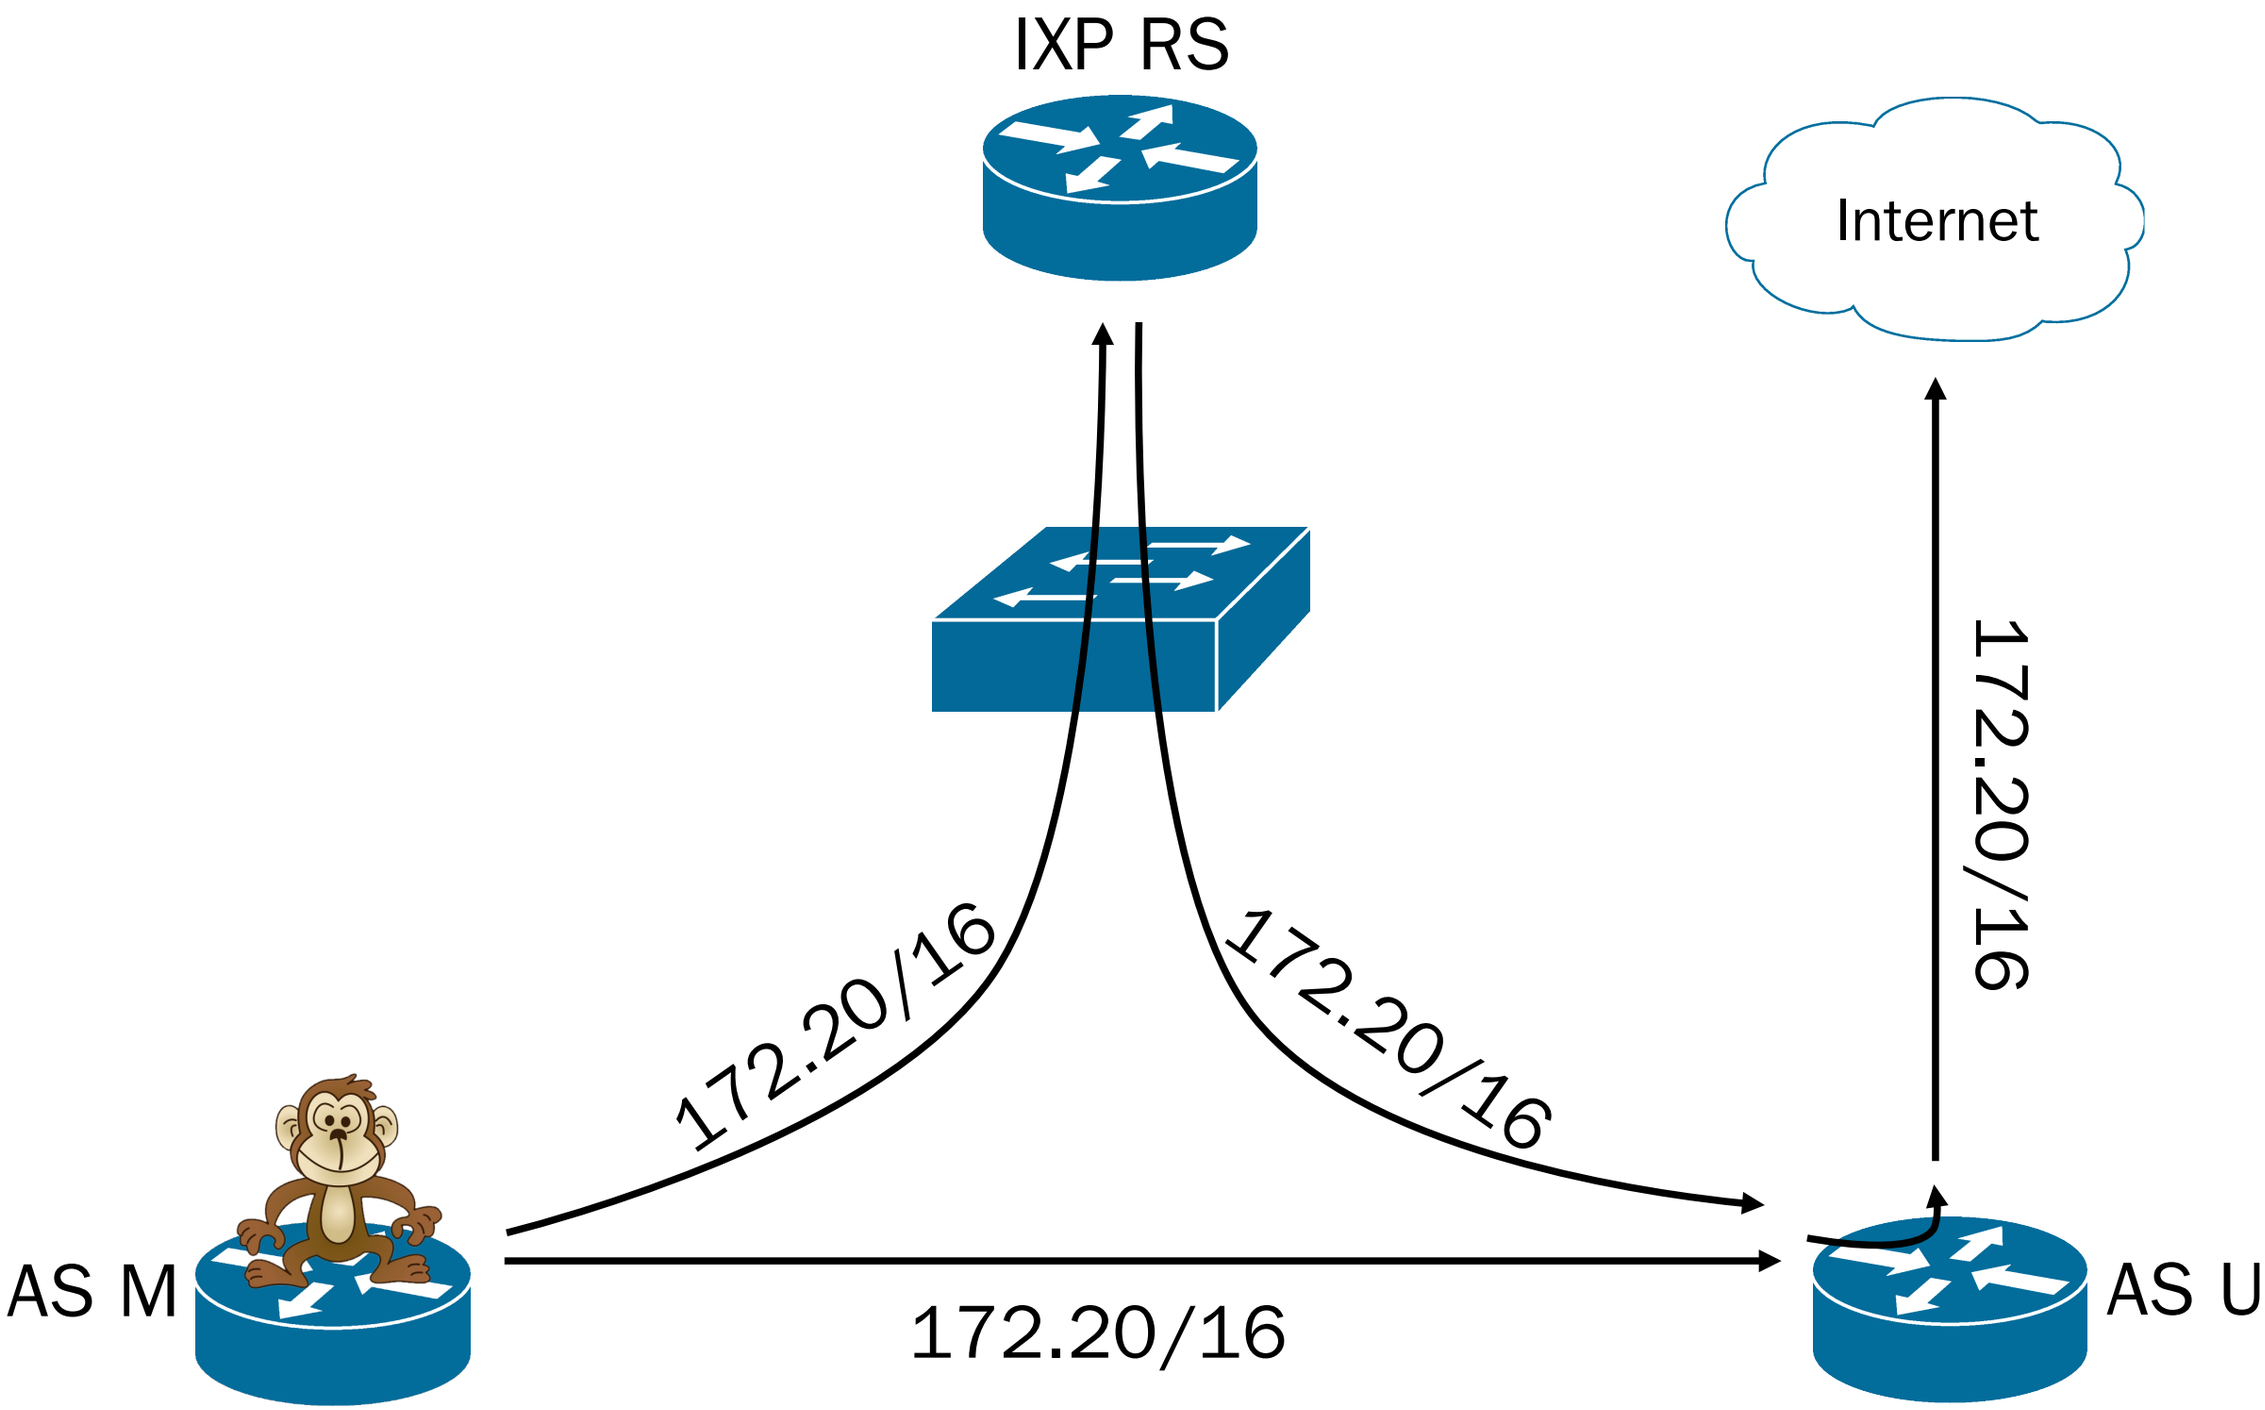

Inbound theft

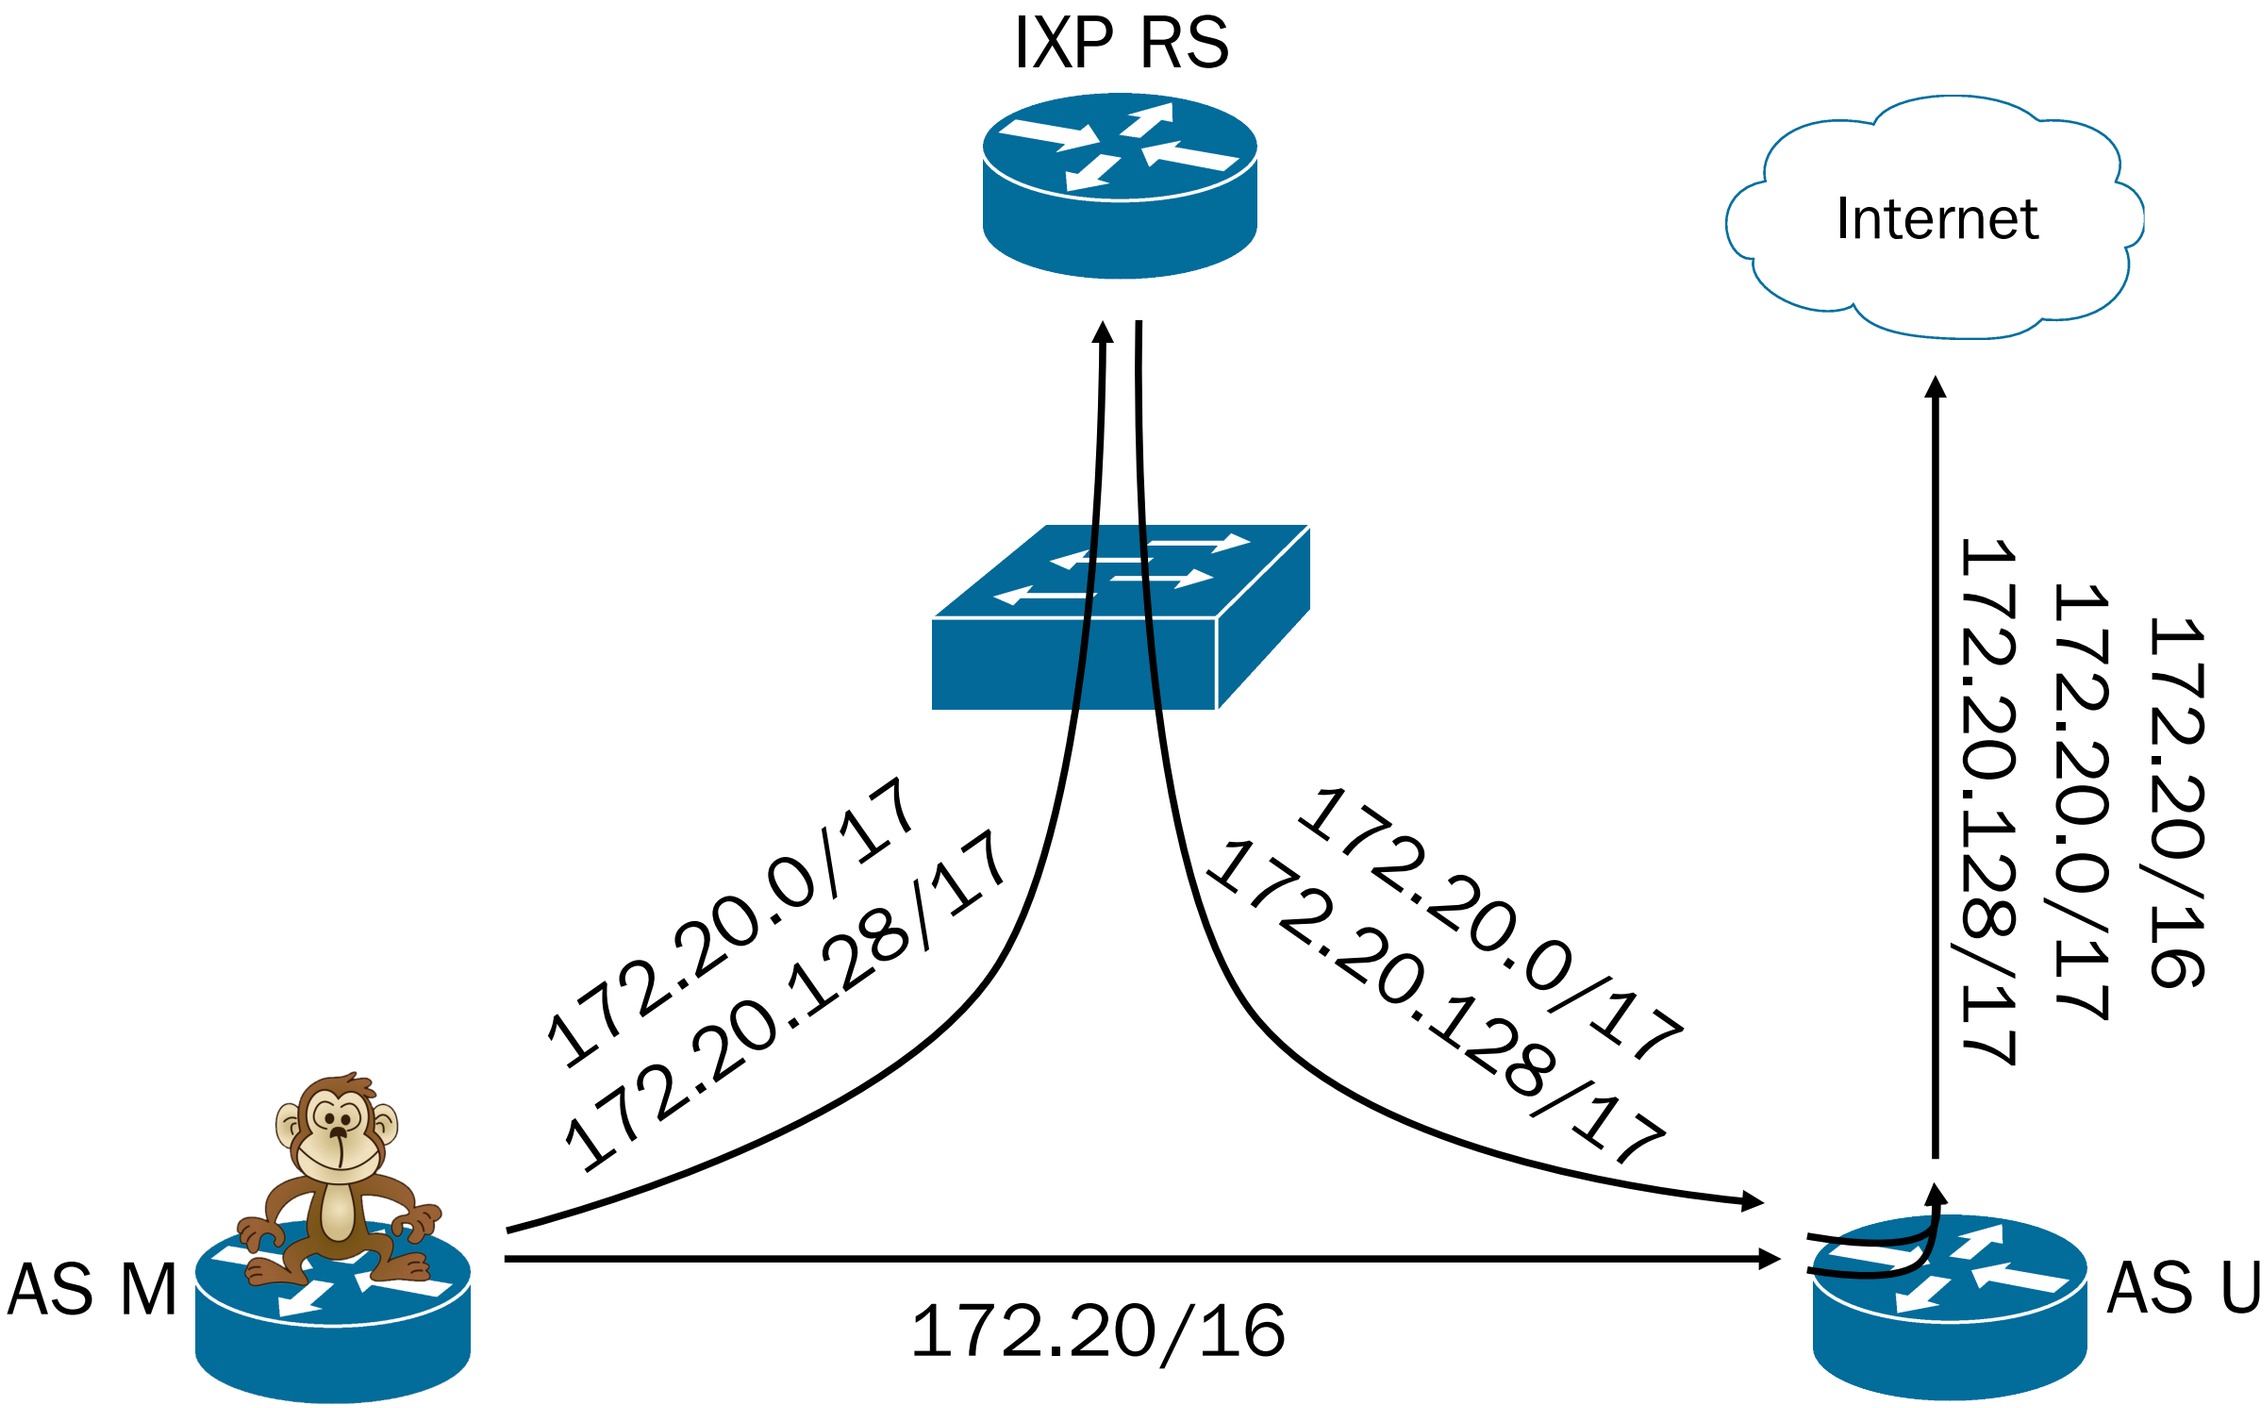

Bi-directional theft

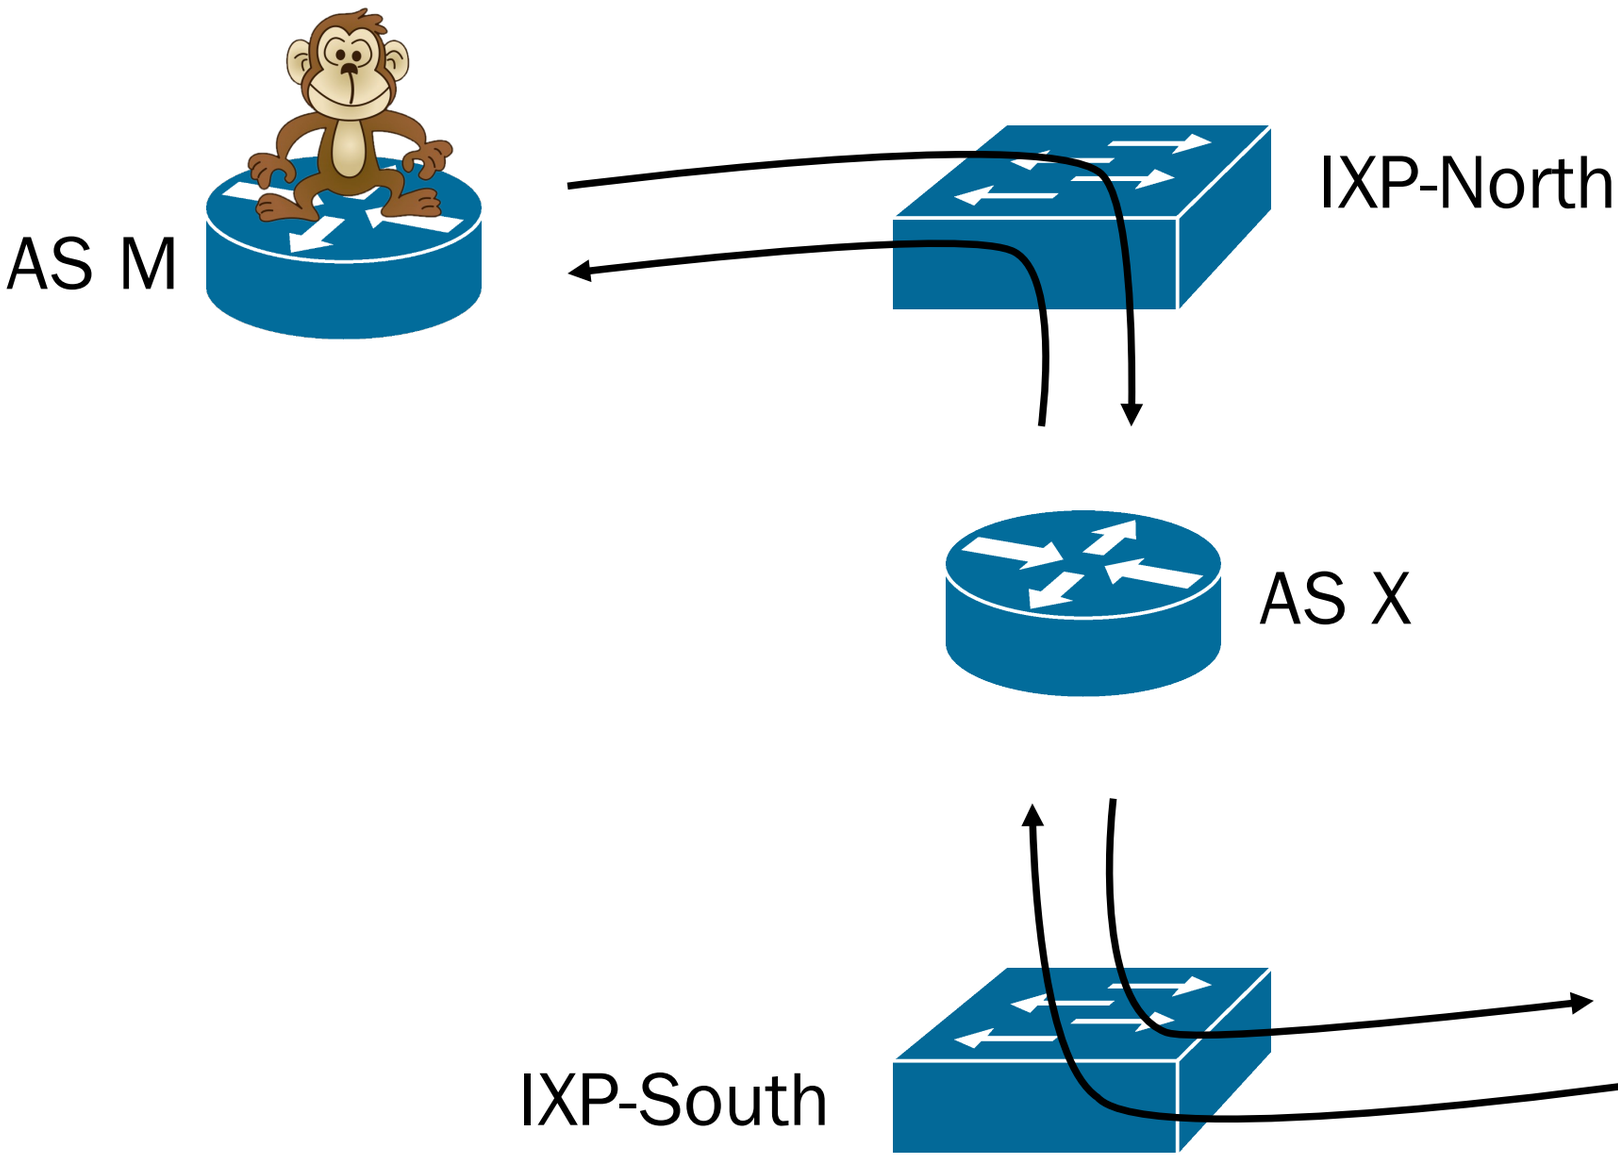

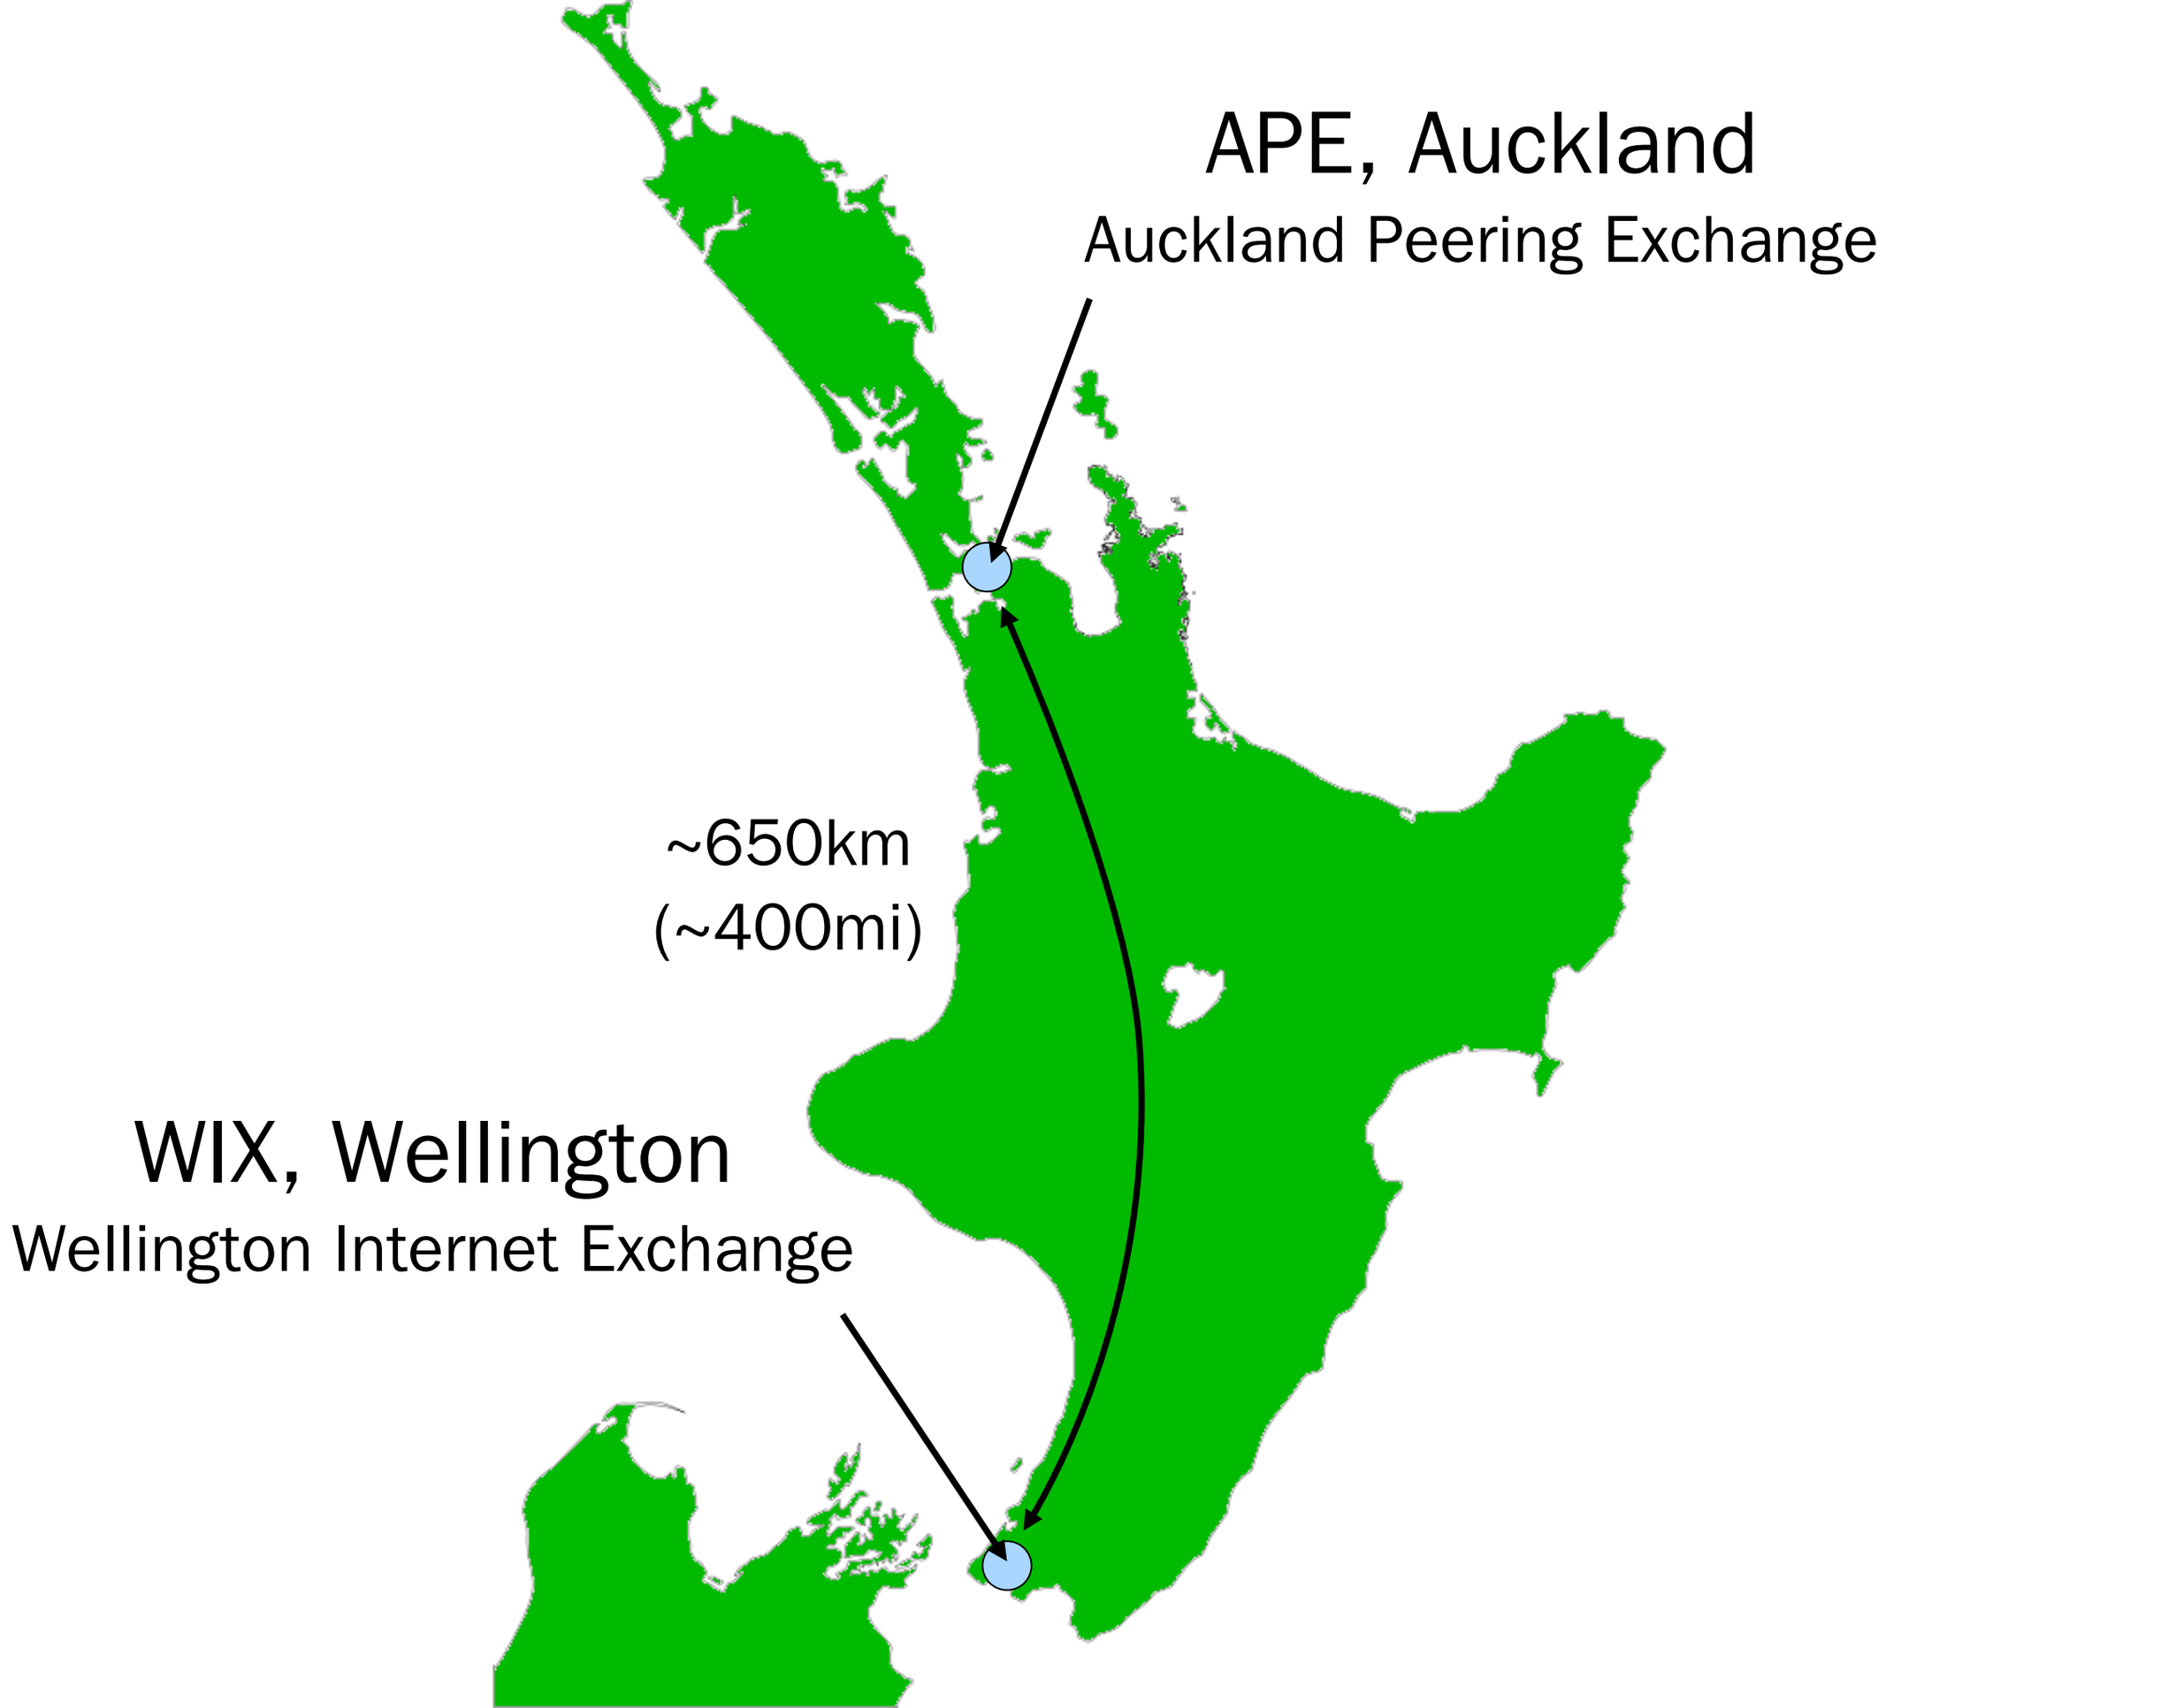
APE, Auckland
Auckland Peering Exchange
~650km
(~400mi)
WIX, Wellington
Wellington Internet Exchange

Bi-directional theft

- APE:
  - 90 devices respond to ARP scan of IXP /24
  - 43 hosts carry packet to WIX /23
- WIX:
  - 112 devices respond to ARP scan of WIX /23
  - 64 hosts carry packet to APE /24

BGP manipulation

- All the usual BGP best-practices still apply!
  - particularly important when connected to IXPs
- Transit providers have incentive not to misbehave
- Peers, especially via RS where contractual agreement between member networks may not be in place, may not have any such incentive
  - or may simply make a mistake!
- BGP lies more likely to result in problems when they are heard across peering than transit

BGP manipulation

- Next-hop attribute
- Prefix hijacking
- Default routes (in either direction)
- Full/partial GRT leakage (in either direction)
- Customers announcing more-specific networks from their own prefixes may result in transit theft (see previously)

Attack mitigations

- For unhygienic routers:
  - http://ams-ix.net/config-guide
  - don't rely on vendor defaults
  - know what your routers are doing
  - if a feature isn't specifically required, turn it off
  - no unrestricted automatic discovery/configuration!
  - no no no IGP on peering interfaces!

Attack mitigations

- Make sure you're only forwarding packets that you want to forward
  - probably only want packets destined for networks that you're advertising via BGP
  - in most cases, only want packets destined for your network, and networks you sell transit to
    - (the packets you're being paid to forward)

Attack mitigations

- Method depends on what your network does, how complicated it is, and your budget
- Either:
  - prevent unwanted packets entering your network
    - apply packet filters on your peering router(s) to only allow packets destined for "valid" destinations through
  - ensure unwanted packets can't get anywhere
    - (or at least anywhere expensive)
    - ensure that your peering router(s) only contain "correct" routes in RIB

Attack mitigations

- "Stub" network
  - smaller networks, only one or two routers in total
  - not many IP customers (and not much churn)
  - unlikely to have dedicated peering router
- Router probably carries default route
- Apply packet filter on IXP interface so only packets destined for your network (and customers) are allowed in

Attack mitigations

- Network with dedicated peering router
  - modifying IXP interface packet filters may be too much work as customers come and go
    - but other things need modification when customers come and go anyway, so automate it as part of your provisioning and deprovisioning processes
- Ensure router carries only:
  - your prefixes/customer prefixes
  - prefixes learned from peers at IXP
- No default route!

Attack mitigations

- Larger network, complex routing policies
- Multiple routing tables
- Ensure that your IXP interface is in a VRF that only contains:
  - your prefixes/customer prefixes
  - prefixes learned from peers at IXP

Attack mitigations

- BGP manipulation:
  - Prefix filter
  - AS_PATH filter
  - maximum prefix limits
  - driven out of IRR where possible
- Probably better to accidentally reject a legitimate route via IXP than to accept a broken one

Questions?

Mike Jager
mike@mikej.net.nz