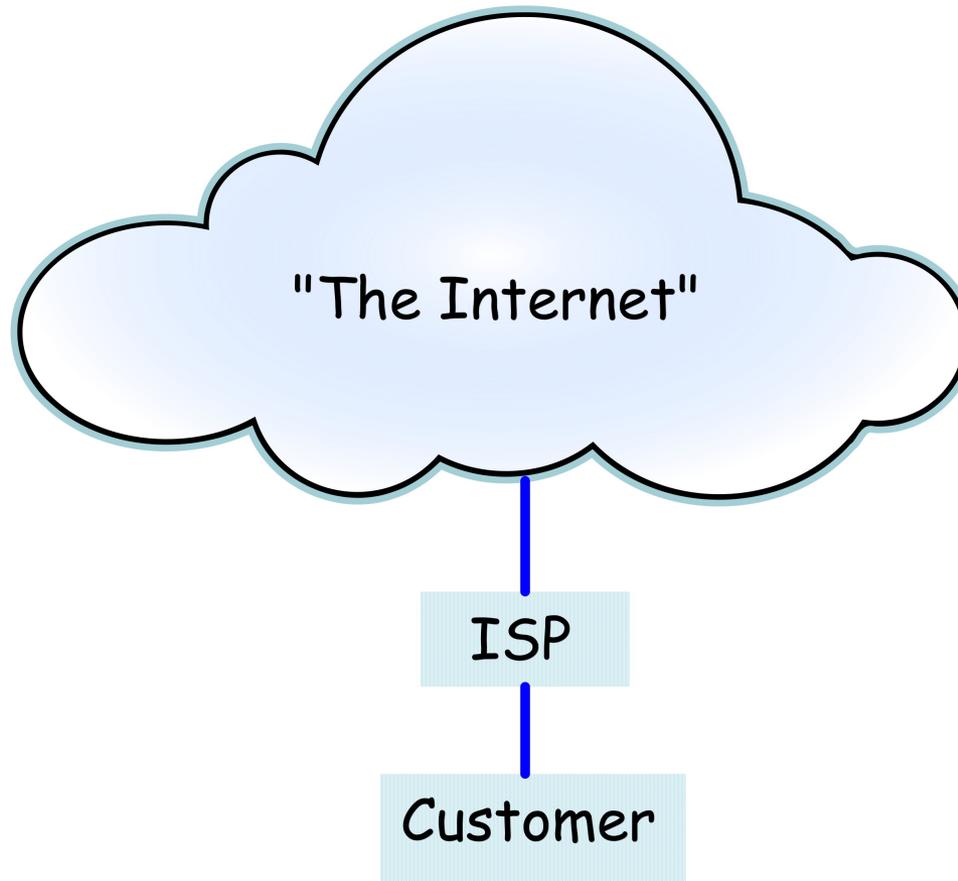


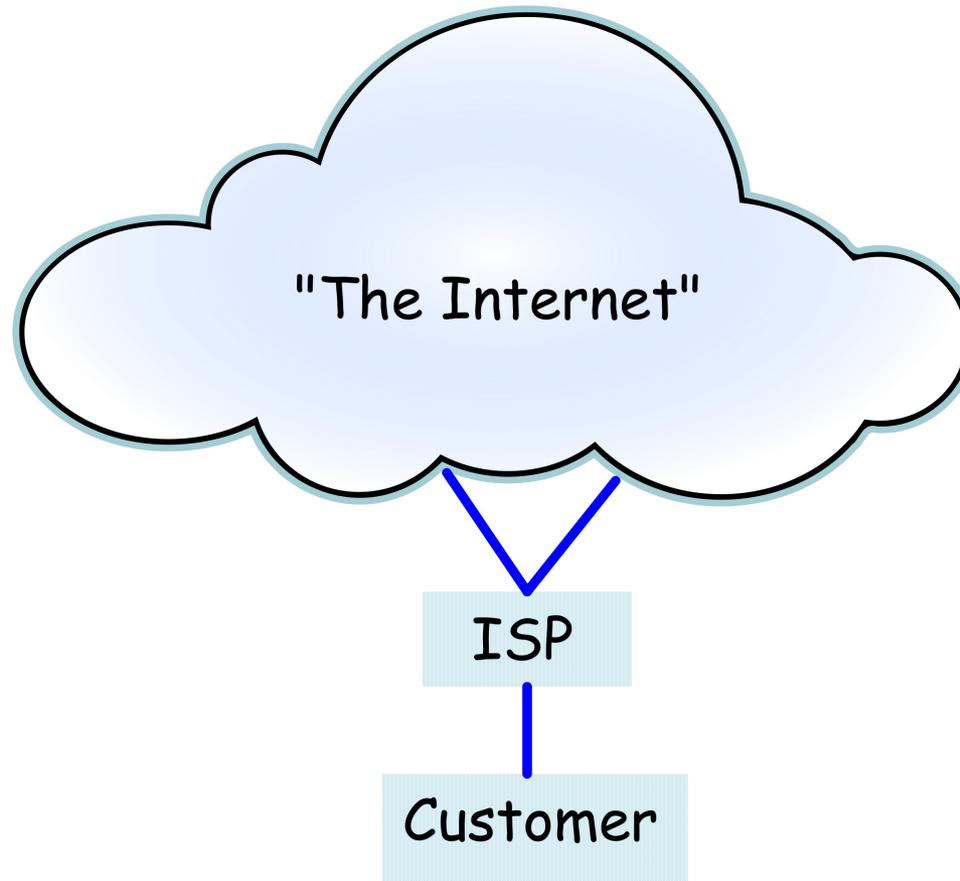
Internet Peering

Why, How, Where, ...

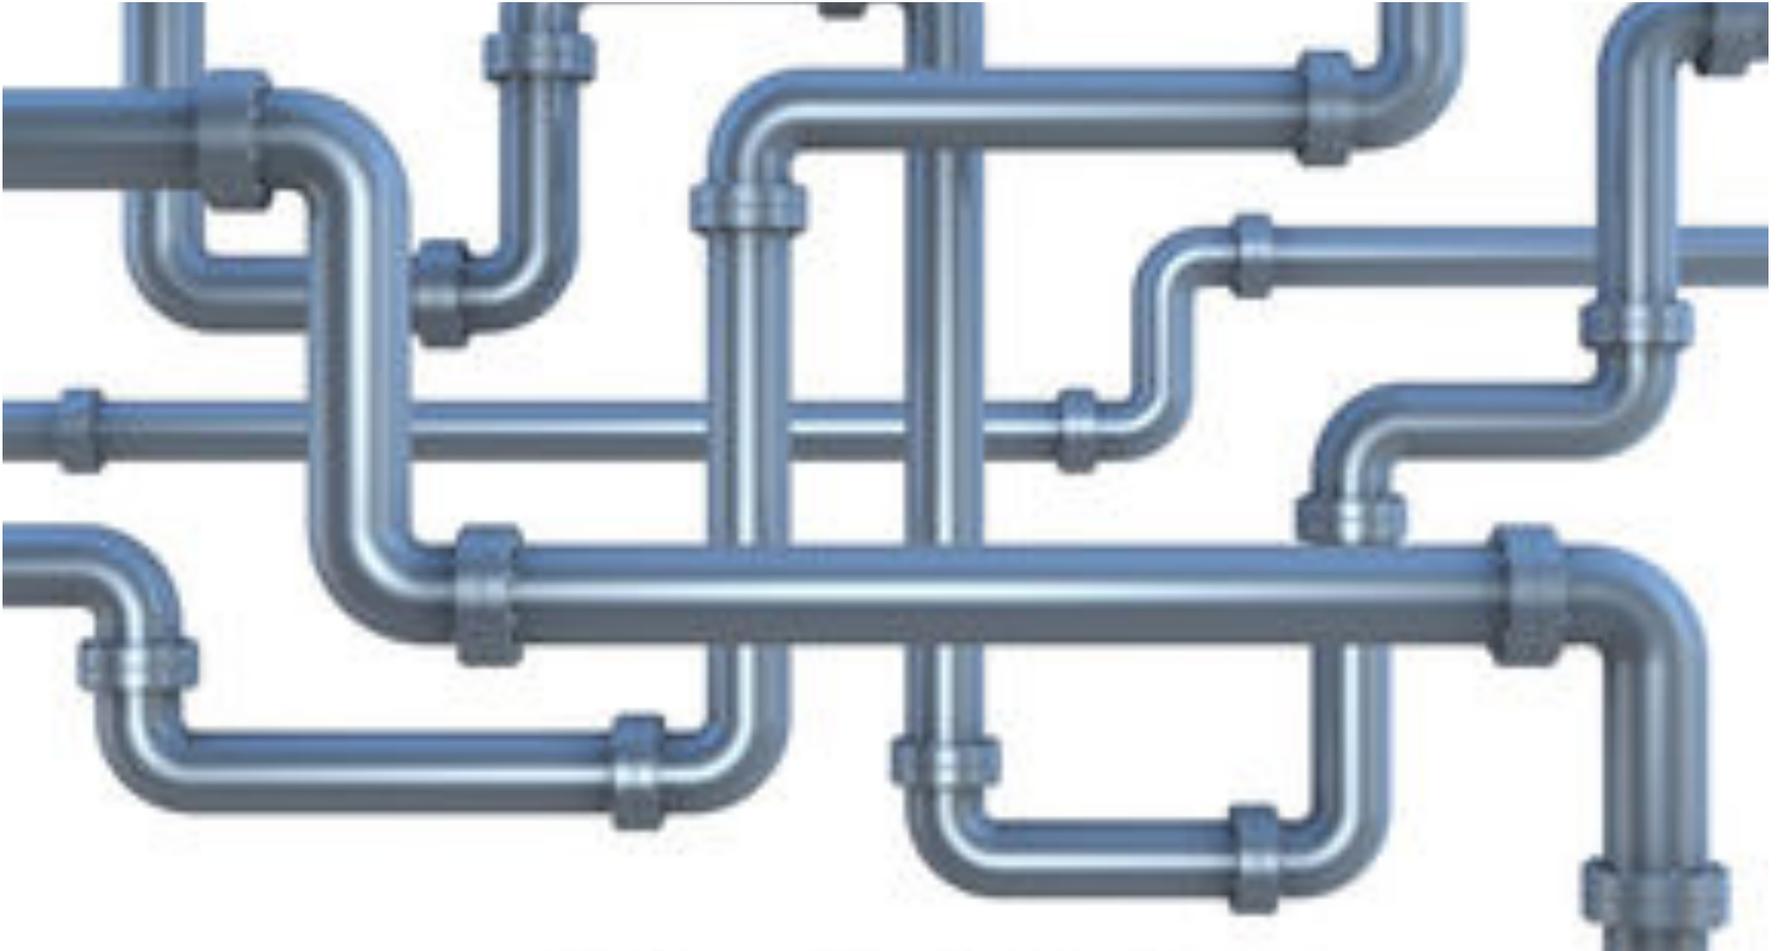
Customer's Expectation



Or maybe this



But it's really just...



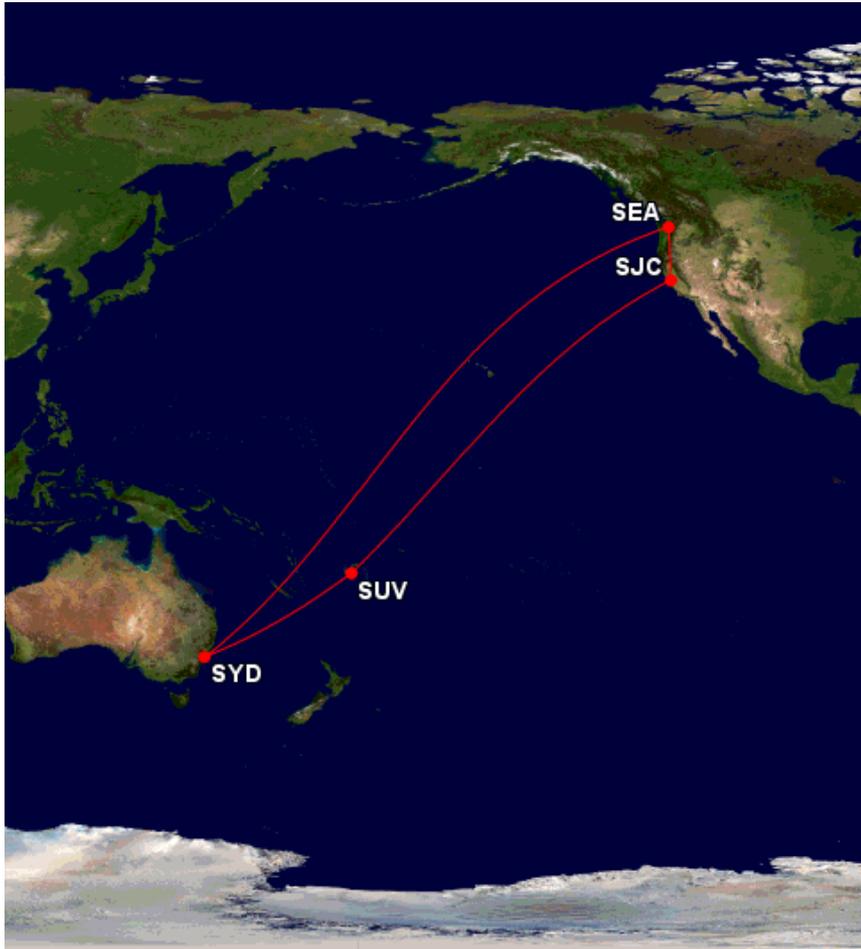
Until this happens



Or this

```
2 core1-vl400.vcc.kidanet.com.fj (113.20.64.49) 327.221 ms 204.824 ms 12.070 ms
3 202.170.41.85 (202.170.41.85) 1.564 ms 2.537 ms 1.383 ms
4 202.170.33.2 (202.170.33.2) 2.826 ms 2.738 ms 1.563 ms
5 gi0-2-1-4.rcr21.b001848-1.sjc01.atlas.cogentco.com (38.122.92.249) 123.644 ms 123.736 ms ...
6 be2063.ccr21.sjc01.atlas.cogentco.com (154.54.1.161) 124.323 ms
  be2095.ccr22.sjc01.atlas.cogentco.com (154.54.3.137) 124.578 ms
  be2063.ccr21.sjc01.atlas.cogentco.com (154.54.1.161) 123.759 ms
7 be3144.ccr41.sjc03.atlas.cogentco.com (154.54.5.102) 124.418 ms 124.695 ms
  be3142.ccr41.sjc03.atlas.cogentco.com (154.54.1.194) 123.785 ms
8 zayo.sjc03.atlas.cogentco.com (154.54.10.194) 126.692 ms 125.425 ms 124.132 ms
9 ae16.cr2.sjc2.us.zip.zayo.com (64.125.31.14) 126.694 ms 123.851 ms 124.828 ms
10 ae27.cs2.sjc2.us.eth.zayo.com (64.125.30.232) 142.824 ms 142.947 ms 142.736 ms
11 ae3.cs2.sea1.us.eth.zayo.com (64.125.29.41) 142.369 ms 142.763 ms 142.015 ms
12 ae28.mpr1.sea1.us.zip.zayo.com (64.125.29.105) 142.880 ms 144.592 ms 142.519 ms
13 64.125.193.130.i223.above.net (64.125.193.130) 162.471 ms 163.139 ms 162.358 ms
14 xe-1-0-1.pe2.brwy.nsw.aarnet.net.au (202.158.194.120) 163.443 ms 162.016 ms 163.059 ms
15 ae9.bb1.a.syd.aarnet.net.au (113.197.15.57) 162.210 ms 163.574 ms 162.243 ms
16 ge-1-1-0.bb1.a.suv.aarnet.net.au (202.158.194.226) 198.100 ms 197.932 ms
```

What's wrong with this picture?



- Fintel customer in Suva
- Accessing content at the University of the South Pacific in Suva
- Packet travels $> 25,000\text{km}$
- Physical distance $< 10\text{km}$
- Adding latency
- Possibly jitter too
- Using expensive submarine capacity

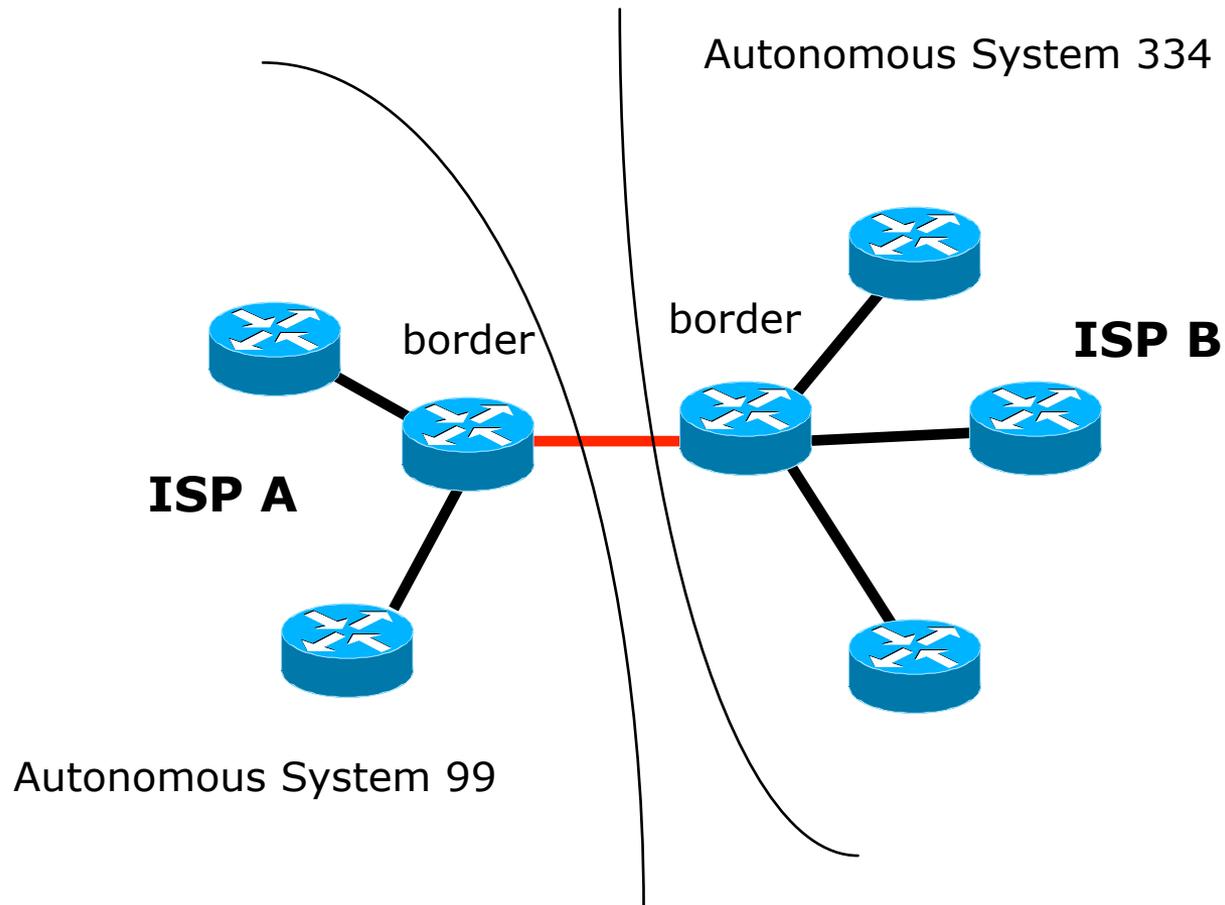
Interconnection, aka Peering

- Connection to a “peer” network
 - Exchange of traffic to customers of each peer
- Requires a circuit to the peer
 - Cost based on capacity of the link
 - May also require a cross connect in a data centre
 - Could be fixed cost or more likely monthly recurring fee
- Traffic is settlement free
- Cost is the same if zero bytes exchanged or link saturated
 - Don't saturate the link, customers will be grumpy 😊

We compete, why interconnect?

- International Connections...
 - If satellite, RTT is around 550ms per hop
 - Compared to local traffic < 10ms round trip
- International bandwidth
 - Costs significantly more than domestic bandwidth
 - Don't congest it with local traffic
 - Wastes money
 - Harms overall performance (end-user experience)

Private Interconnect



What is a Private Interconnect?

- Local (loop) connections
 - Using leased line, dark fibre, microwave, ...
- Can't be in a customer/transit relationship
- Sharing customer & infrastructure routes only
 - Routes that generate revenue for you
- Share costs
 - Two circuits, pay for one each

Results of Peering

- Both parties save money
- Local traffic stays local
- Costs typically fixed rather than volume based
- Better performance, better QoS, ...
- Expensive international bandwidth available for actual international traffic
- Everyone is happy (except submarine cable and satellite owners)

Scaling peering

- Just repeat process?
- Private peering means that each ISP has to buy circuits to every other peer (perhaps 2 for redundancy)
 - For (n) peers each peer needs (n-1) half circuits
 - Eg 5 peers => 4 half circuits

Why use an Internet Exchange Point?

- Private peering relies on just the two parties making best use of the circuit
- Multiple peers => multiple circuits
- Using an Internet Exchange Point every participant only has to buy just one whole circuit
 - From their premises to the IXP fabric
- Connecting via an IXP maximises the opportunity to fill that circuit
 - Peak traffic isn't the same across all peers
 - Also can improve latency performance between peers where traffic volume wouldn't justify a dedicated circuit

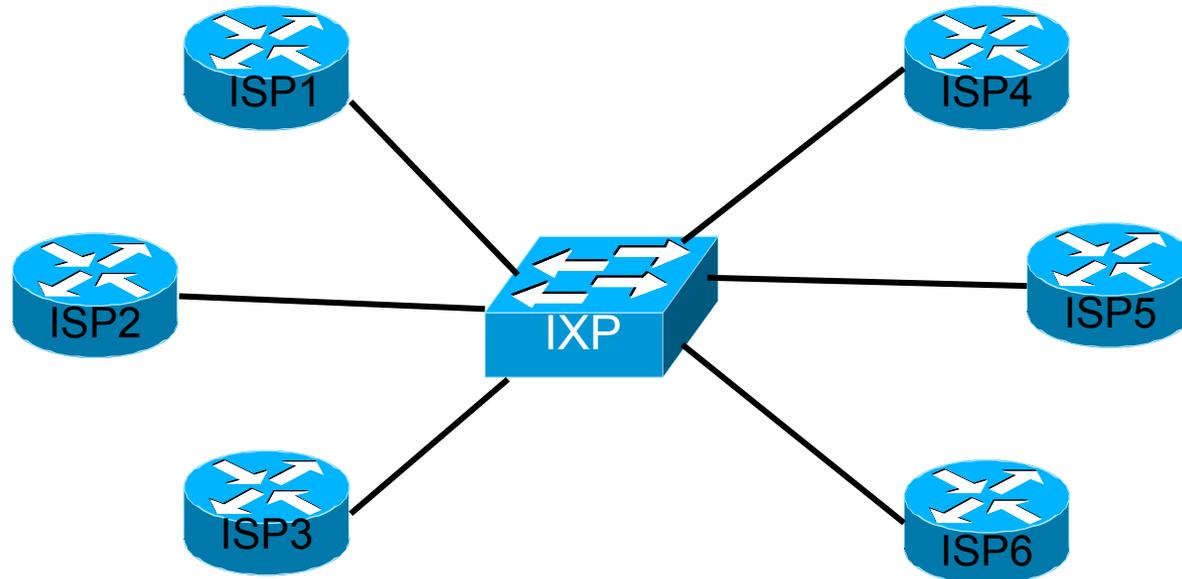
Internet eXchange Point (IXP)

- Need a location or facility that ISPs can access and can connect to each other over a common shared media
 - Eg: Ethernet switch
- Should be a NEUTRAL venue
- Needs to have multiple telco circuit providers and/or allow any licenced provider to install services
- Needs controlled environment & access

Internet eXchange Points

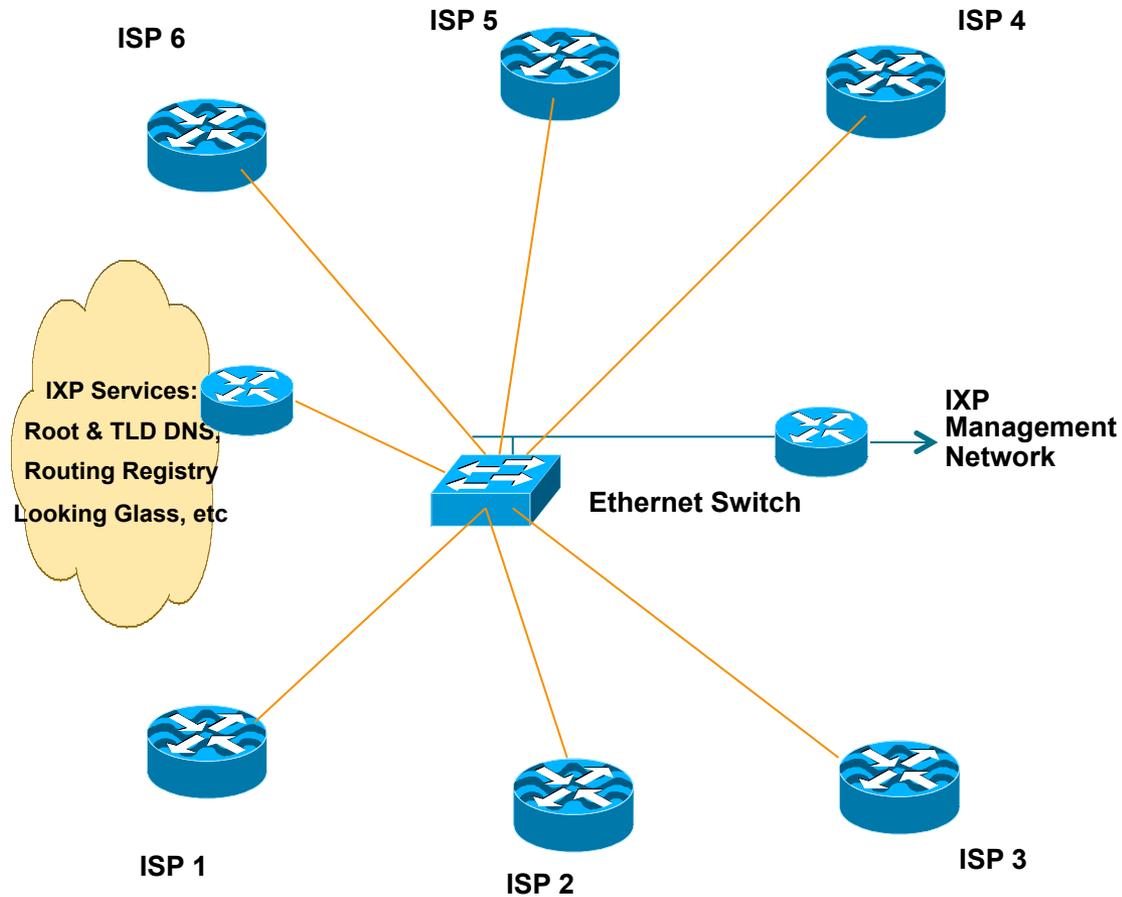
- Variety of shapes and sizes
 - Commercial
 - Community
 - Tbps to Mbps
 - Single location or Metropolitan Area scoped
 - Purely a traffic exchange
 - Value added services
- Layer 2 exchange point
 - Ethernet Switches (100Gbps/10Gbps/1Gbps/100Mbps)

Internet eXchange Point

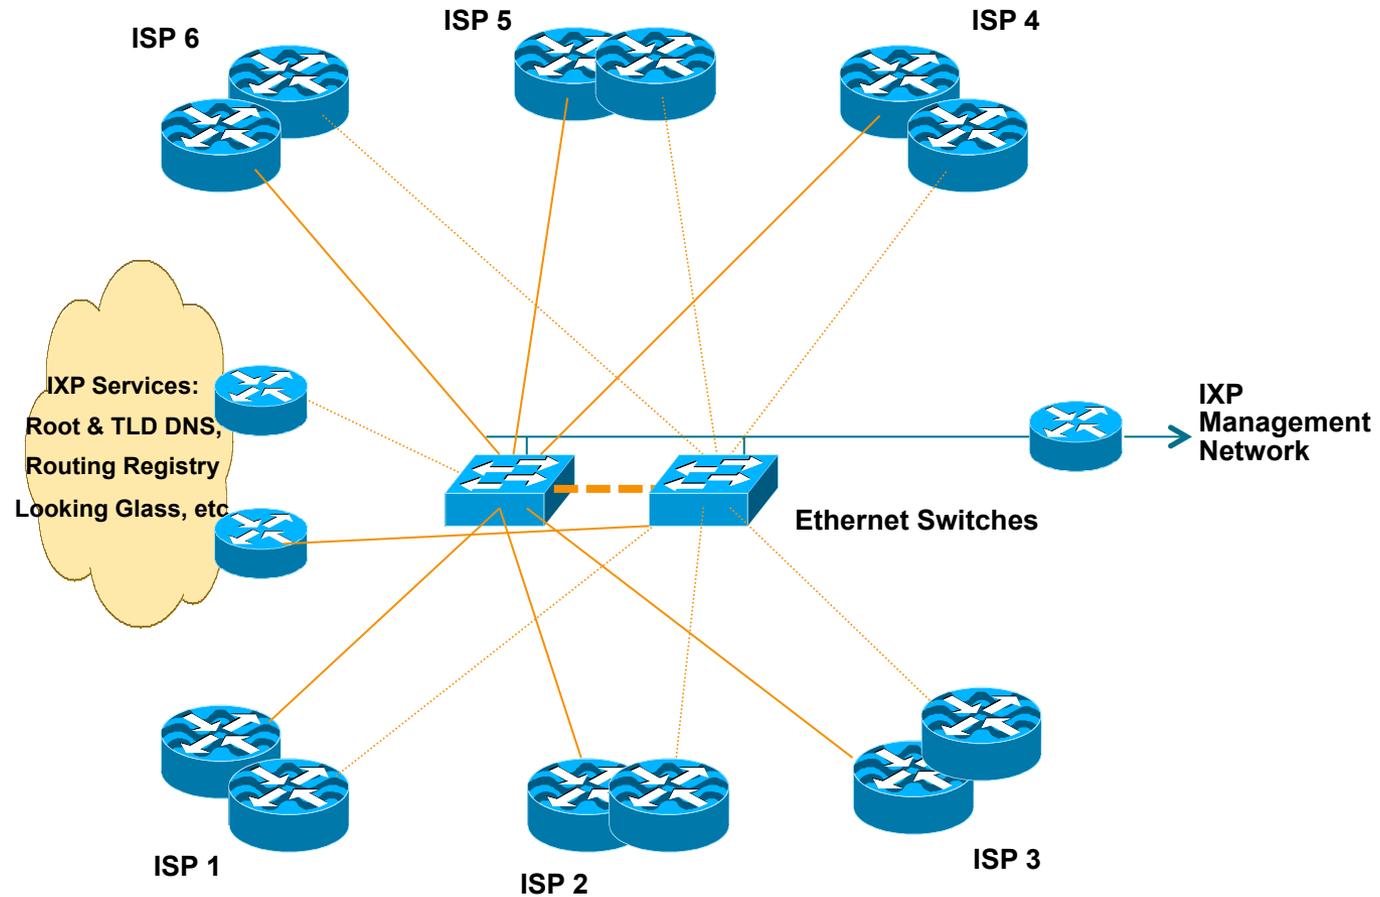


- Border routers in different Autonomous Systems

Typical Exchange



Or with more redundancy



Internet Exchange Point

- Two switches for redundancy
- ISPs use dual routers for redundancy or load sharing
- Offer services for the “common good”
 - Internet portals and search engines
 - DNS Root & TLDs, NTP servers
 - Routing Registry and Looking Glass

Internet Exchange Point

- Requires neutral IXP management
 - Usually funded equally by IXP participants
 - 24x7 cover, support, value add services
- Secure and neutral location
- Configuration
 - Private address space if non-transit and no value add services
 - Otherwise public IPv4 (/24) and IPv6 (/48, /56, /64)
 - ISPs require AS, basic IXP does not

Internet Exchange Point

- Network Security Considerations
 - LAN switch needs to be securely configured
 - Management routers require AAA authentication, vty security
 - IXP services must be behind router(s) with strong filters

