IP ROUTING

Metcalf's Law

How much is a network worth?

- Approximation: 1 unit for each person a person can communicate with
  - The more people I can talk to, the more I value the network.
- N people in the network $\Rightarrow$
  - network is worth N^2 "units"
- Network value scales as N^2, (not N) is called Metcalf's law

Implications for Peering

- Simple model of network value implies peering should often happen
- What is the increase in value to each party's network if they peer?
  - Want to compute change in value, ΔV
  - Take larger network value and subtract old
- $\Delta V_1 = N_1(N_1+N_2) - (N_1)^2 = N_1 N_2$
- $\Delta V_2 = N_2(N_1+N_2) - (N_2)^2 = N_1 N_2$

Symmetric increase in value

- Simple model shows net increase in value for both parties
- Both network's values increase is equal!
  - Smaller network: a few people get a lot of value
  - Larger network: a lot get a small value.
- Helps explain "symmetric" nature of most peering relationships, even between networks of different sizes

Takeovers

- Instead of peering, what if the larger network acquires the smaller one?
  - suppose it pays the value for the network too
- $\Delta V = (N_1+N_2)^2 - (N_1)^2 - (N_2)^2 = 2(N_1N_2)$
  - Captures twice as much value by acquisition as peering
- An incentive to not peer
  - E.g. to force a sale or merger, allowing larger network to capture a greater value than by peering

Reasons not to Peer

- Asymmetric Traffic
  - More traffic goes one way than the other
  - Peer who carries more traffic feels cheated
  - Hassle
- Top tier (big) ISPs have no interest in helping lower tier ISPs compete
  - The "Big Boys" all peer with each other at no/little cost
- Harder to deal with problems without strong financial incentive

A lower tier strategy

- Buy transit from big provider
- Peer at public exchange points to reduce transit cost
- Establish private point-to-point peering with key ISPs
- When you're big enough, negotiate peering with transit provider

BGP and Traffic

- Network engineering
  - Estimate traffic matrix
  - Tune network for performance
- Stability assumptions for estimation, tuning
- Reality:
  - Inter-domain connectivity grown rapidly
  - Large # of BGP entries, changes
  - Can result in unstable Traffic Matrix
  - Can be bad for performance

Important BGP attributes

- LocalPREF
  - Local preference policy to choose "most" preferred route
- Multi-exit Discriminator
  - Which peering point to choose?
- Import Rules
  - What route advertisements do I accept?
- Export Rules
  - Which routes do I forward to whom?

BGP Operations (Simplified)

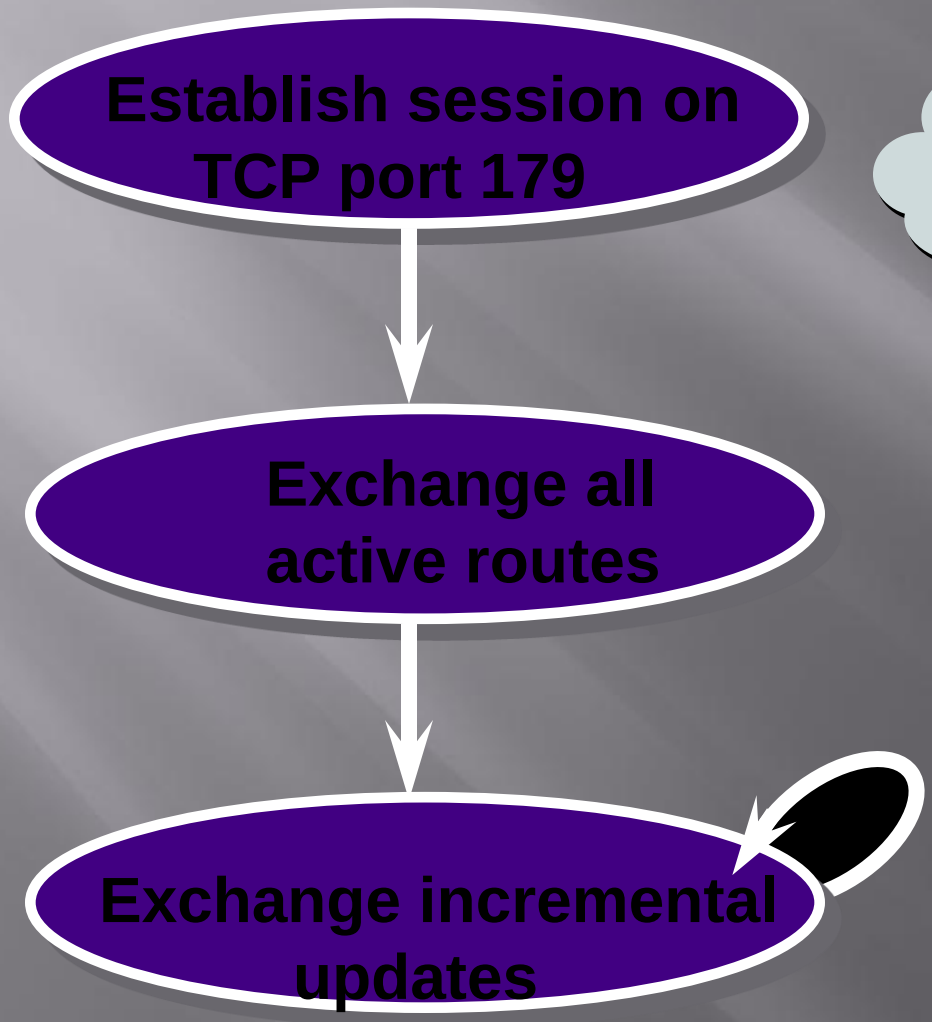

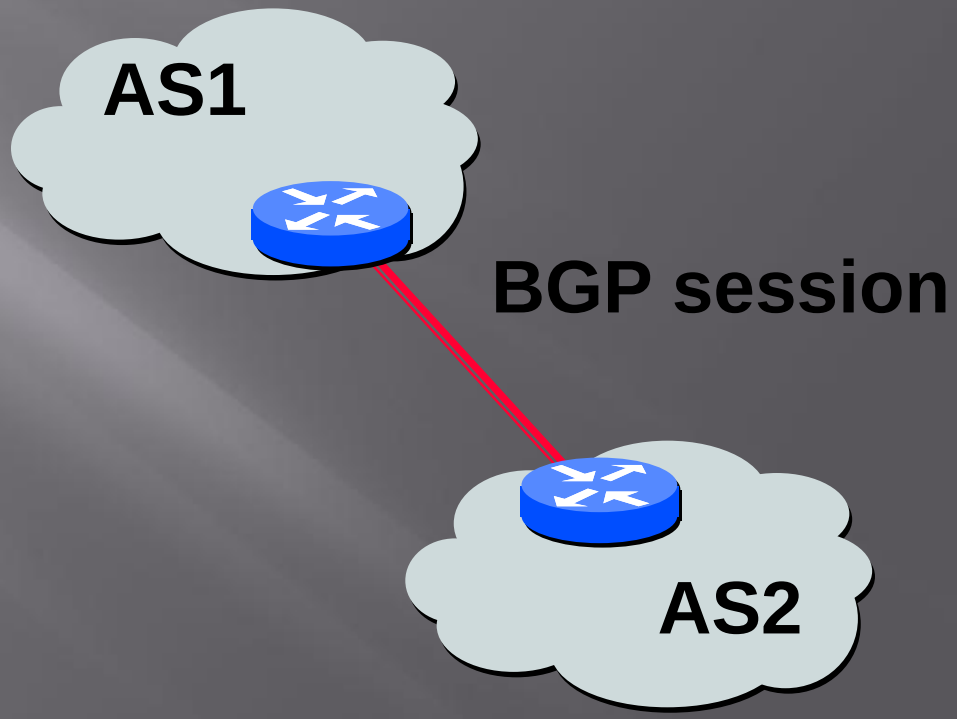

While connection is ALIVE exchange route UPDATE messages

Four Types of BGP Messages

- **Open** : Establish a peering session.
- **Keep Alive** : Handshake at regular intervals.
- **Notification** : Shuts down a peering session.
- **Update** : Announcing new routes or withdrawing previously announced routes.

announcement

=

prefix + attributes values