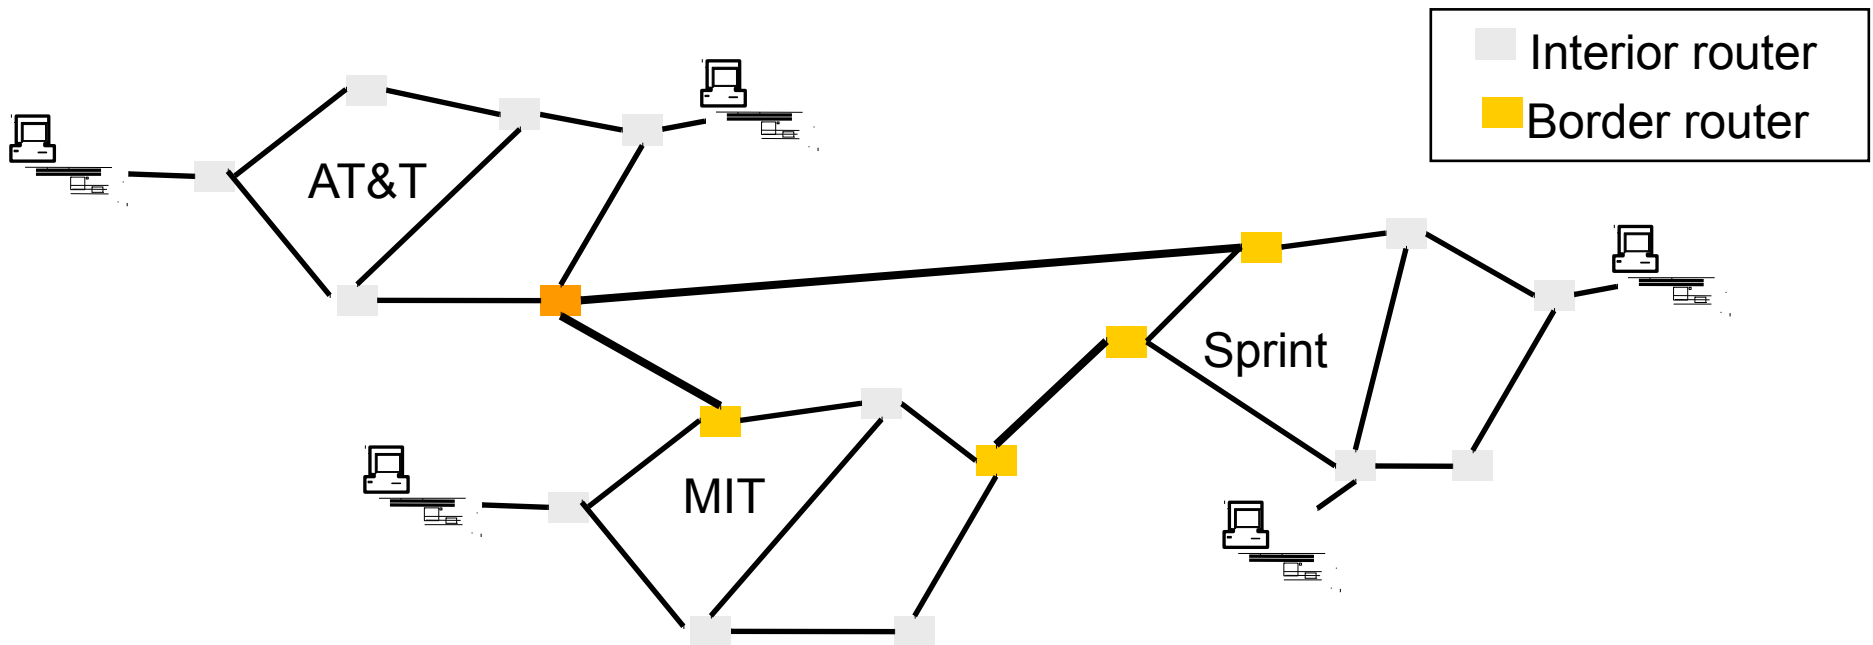



BGP and RON

Internet Routing

- Internet split into *Autonomous Systems* (ASes). BGP routing between ASes.



AS Relationships



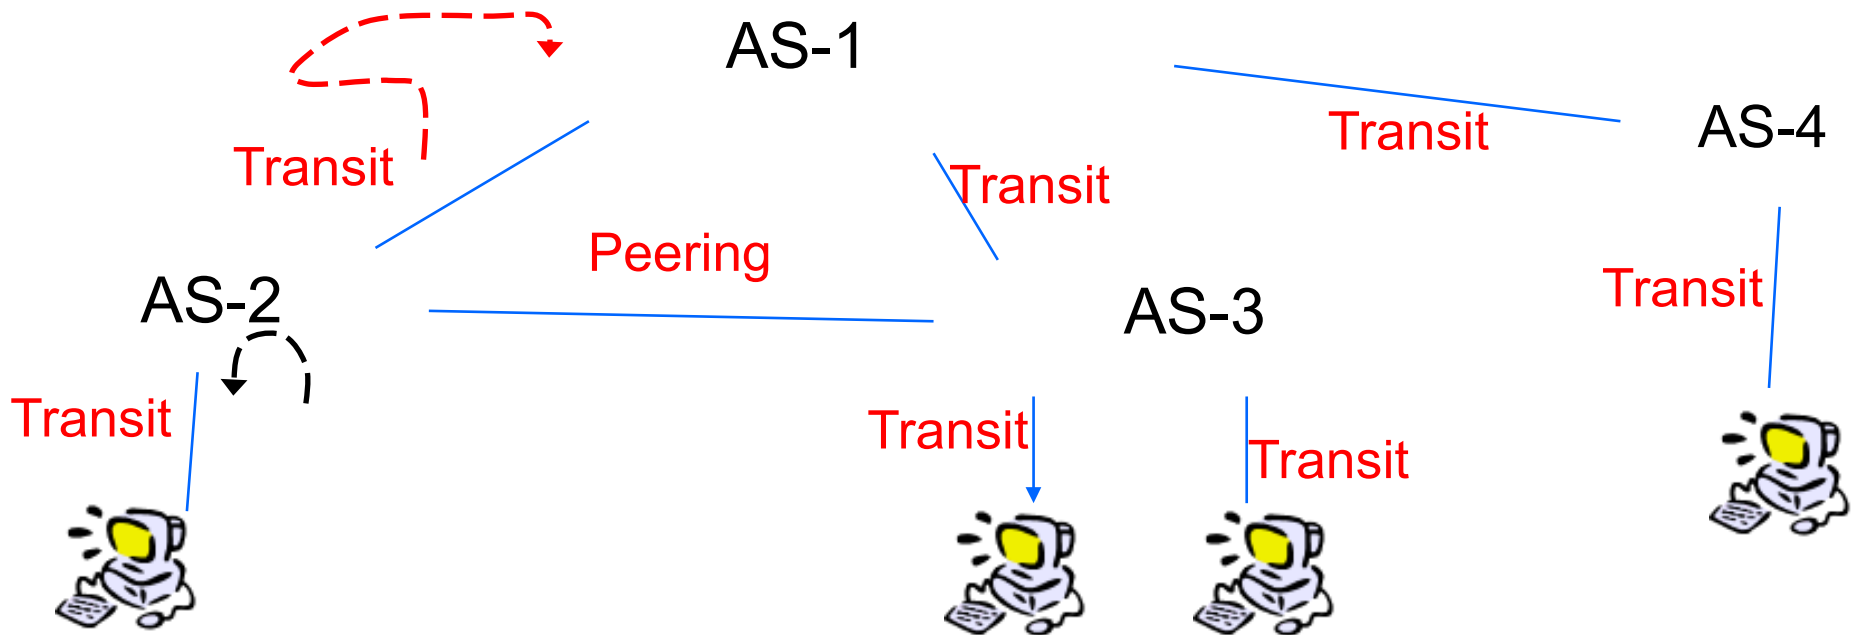
- Transit: Customer to provider, \$\$
- Peering: Peer to peer, no \$

BGP

- 2 things: Exporting (advertising) and Importing (picking) Routes
- 1 Simple Principle:
If I (an AS) take a packet, I better get...



BGP Route Export



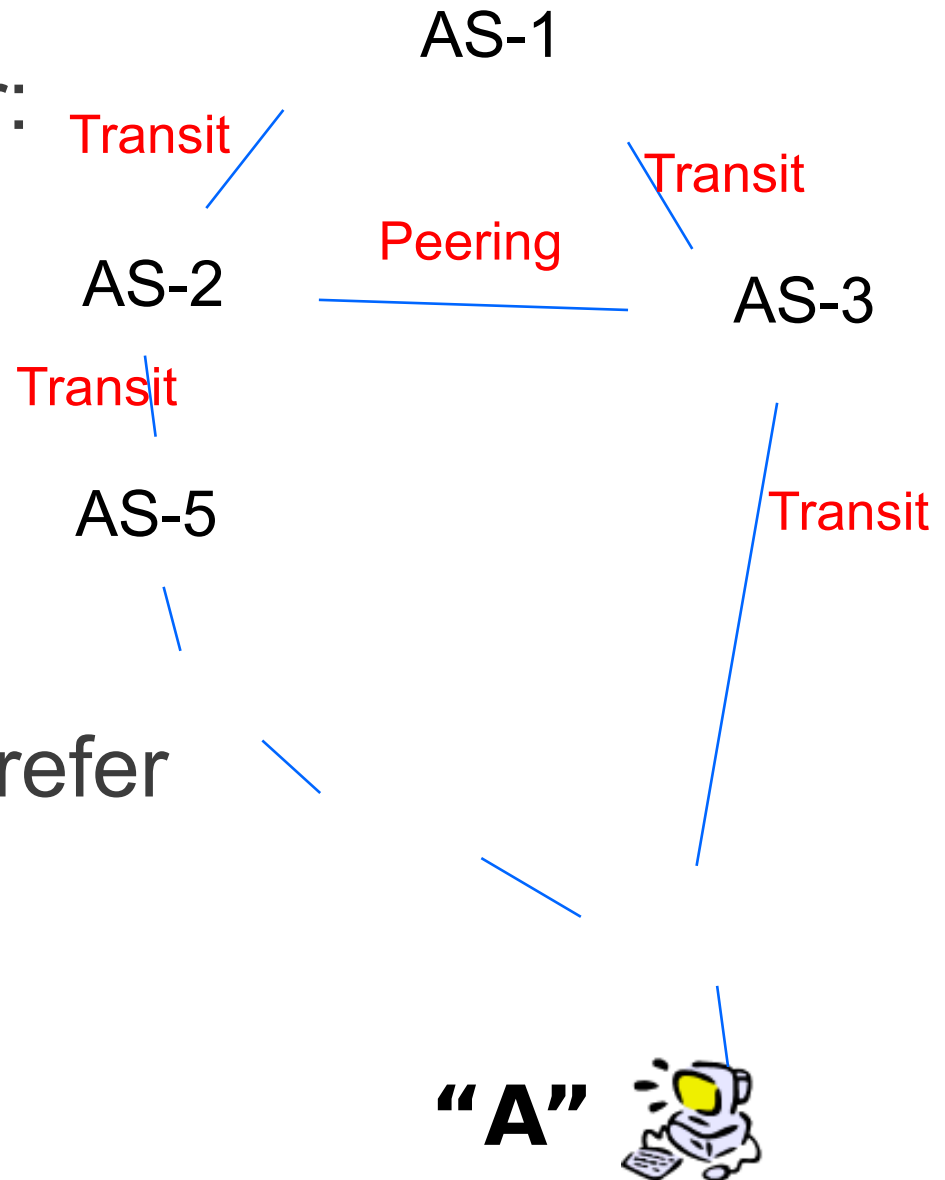
- AS tells everyone about its customers.
- AS tells customer about everyone (it knows)
- AS only tells peers about its customers.
- AS doesn't tell others about peers.

BGP Route Import

Prefer routes in order:

- Customer
- Peer
- Provider

All else equal, then prefer shorter AS-path

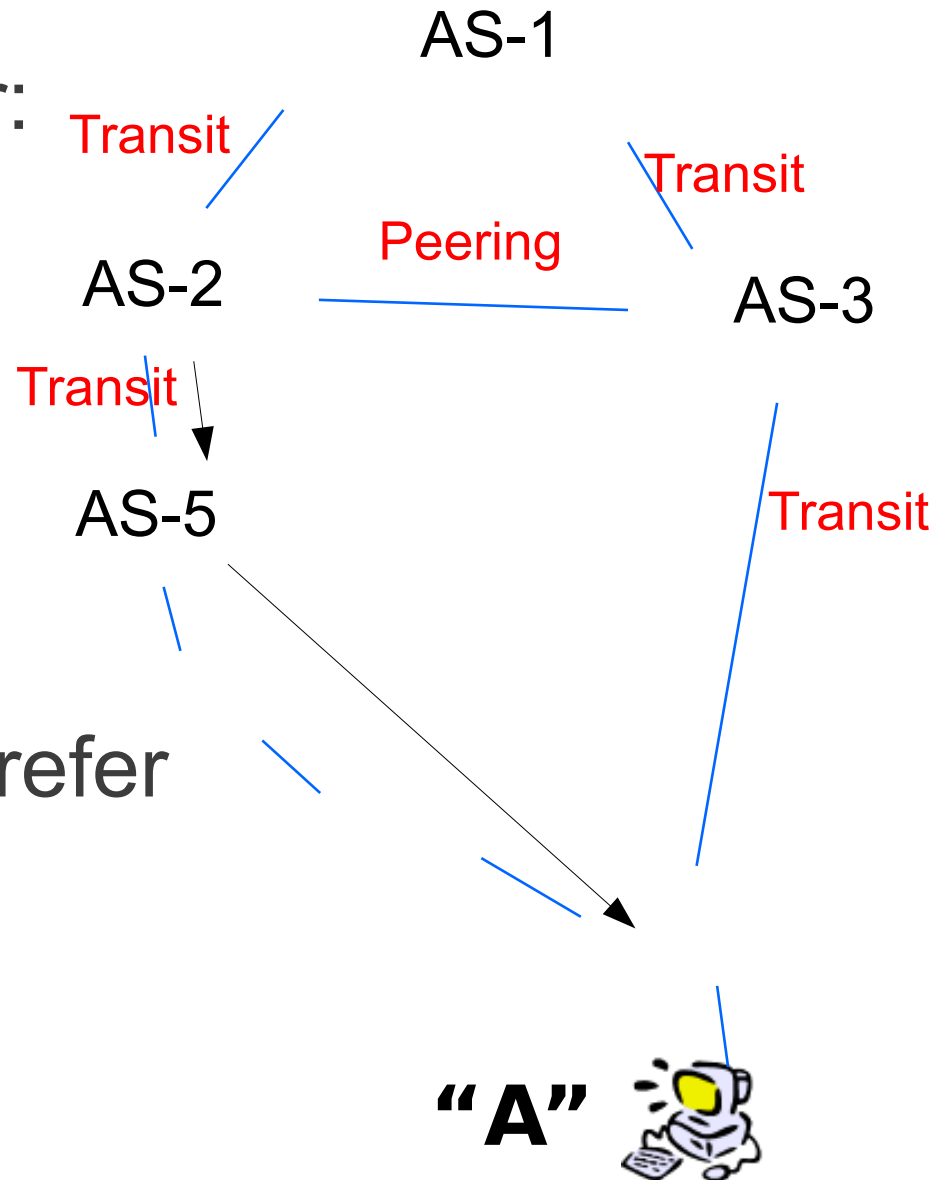


BGP Route Import


Prefer routes in order:

- **Customer**
- Peer
- Provider

All else equal, then prefer shorter AS-path




Sample Problem (2012 Q2)




- Note arrows from provider to customer, double-arrow is peering

Sample Problem (2012 Q2)



- Path AS-1 to P?

Sample Problem (2012 Q2)




- AS Path AS-1 to P?

AS-1 \rightarrow AS-3 \rightarrow AS-5


shorter path

Sample Problem (2012 Q2)




- # Paths AS-2 learns for P? Which path used?

Sample Problem (2012 Q2)




- # Paths AS-2 learns for P? Which path used?
3 (from AS-1, AS -3, and AS-4)
AS-2 → AS-4 → AS-5 (prefer customer routers)

Sample Problem (2012 Q2)



- AS-2 / AS-4 link and AS-5 / AS-6 link fail. Everyone still can reach P?

Sample Problem (2012 Q2)

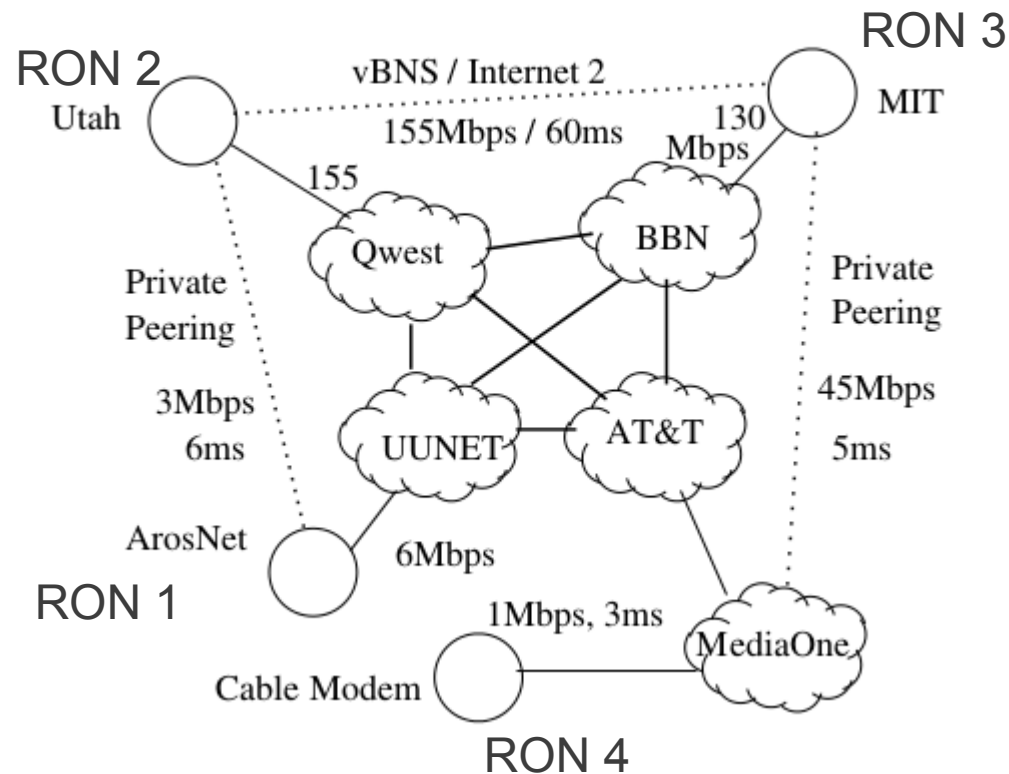


- AS-2 / AS-4 link and AS-5 /AS-6 link fail. Everyone still can reach P?

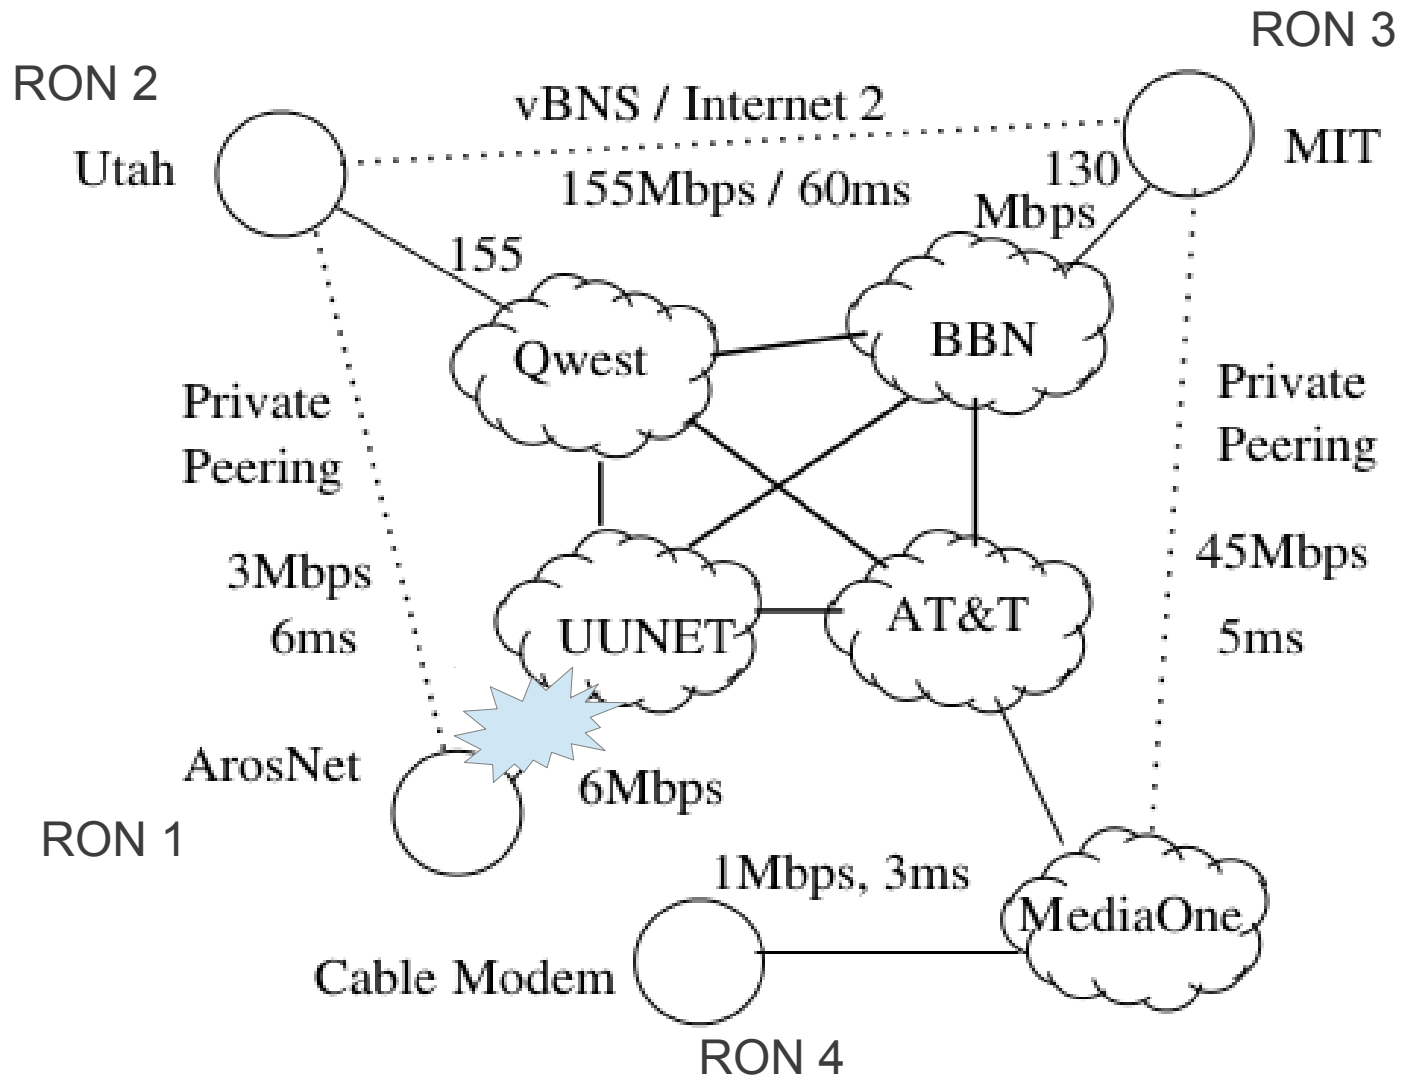
AS-7 can't! Won't know about AS3 b/c it's a peer

RON

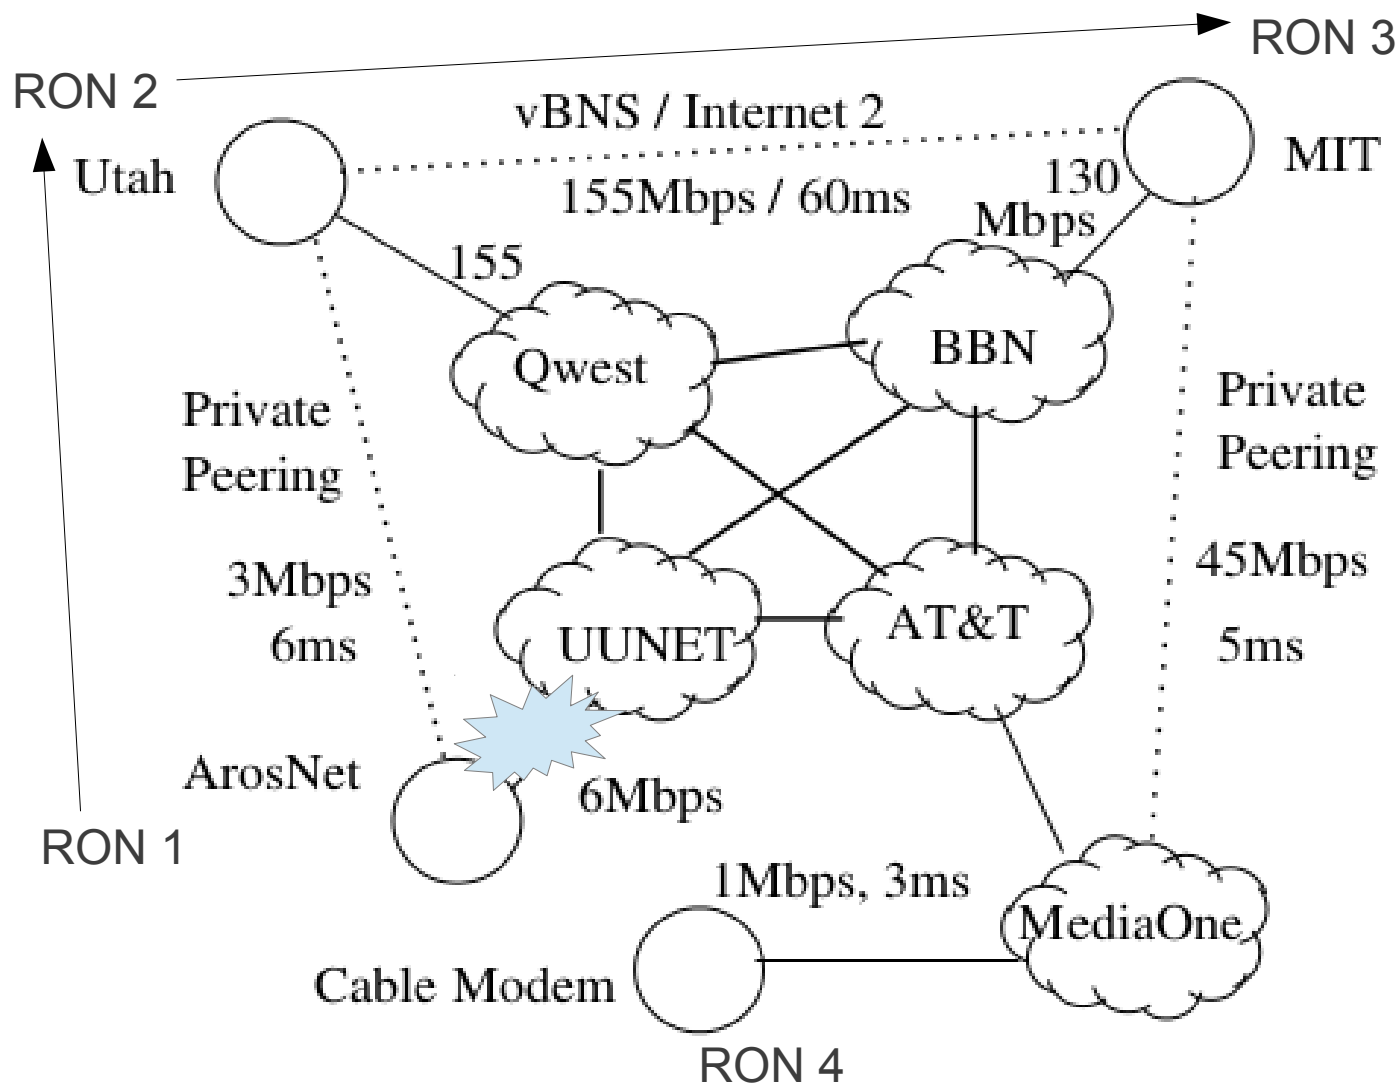
Resilient Overlay Network - built over lower-level network (in this case, the Internet and BGP).



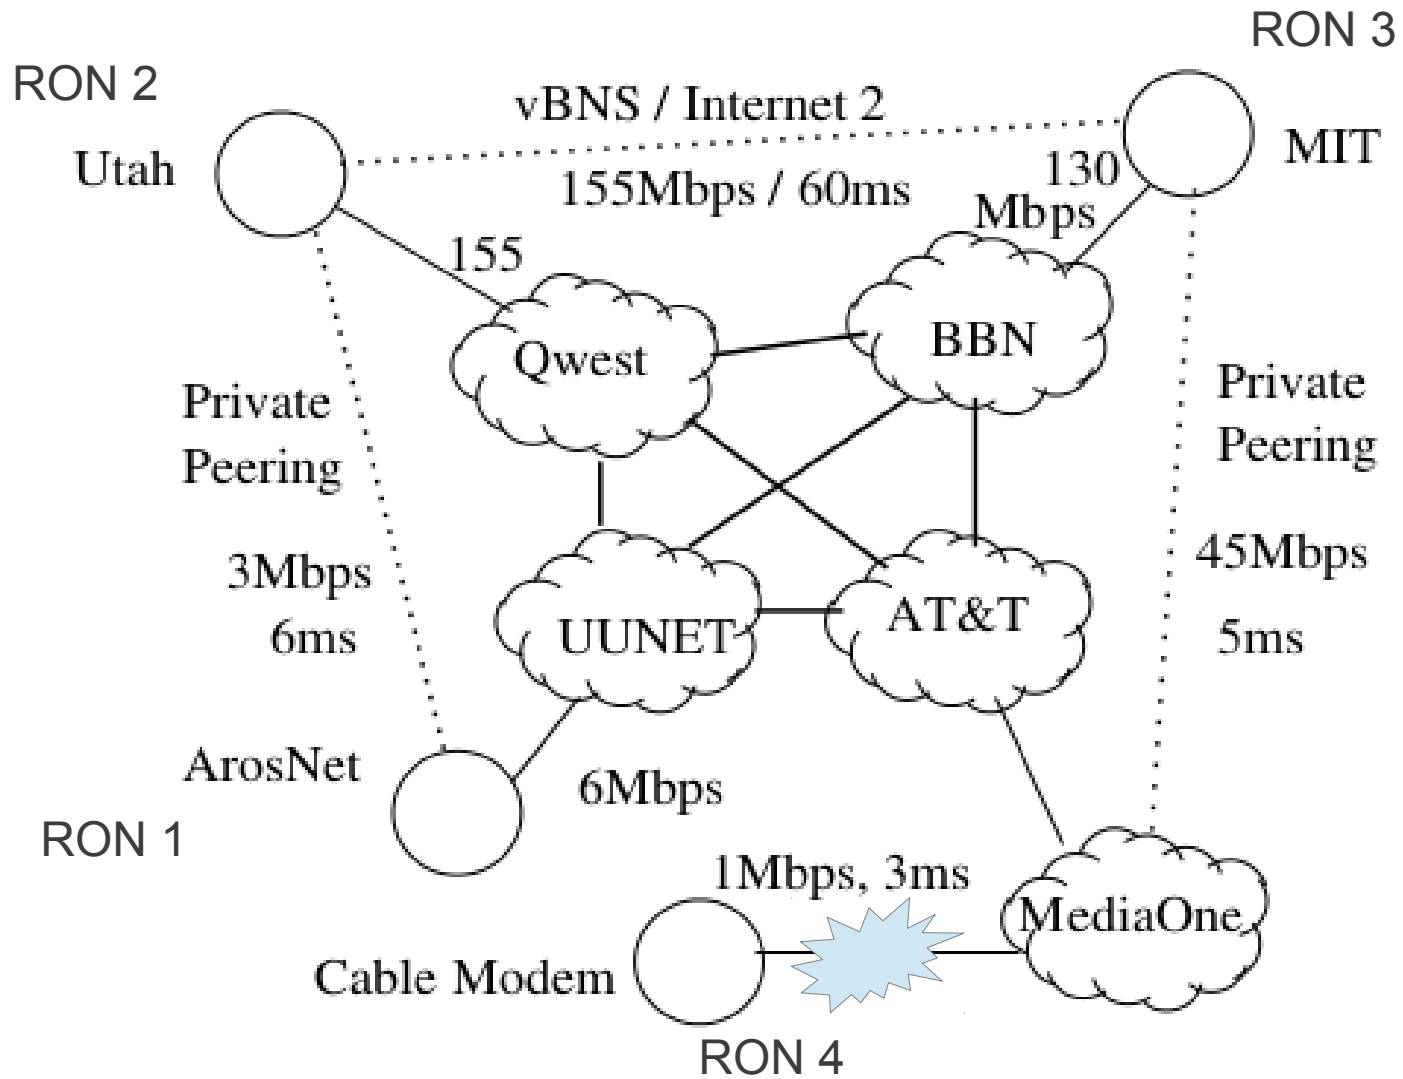
RON



RON



RON



RON Probing

- Need to know all other RON nodes, hence RON networks are small (up to 50 nodes)
- Active vs passive probing to obtain metrics
- Outage Detection (active):
 - Periodically send probe.
 - If times out, send quick sequence of probes.
 - If a threshold of probes in a row don't hear responses, mark link dead.

RON Link Metrics

- Stored in a performance database
- Latency, loss / packet drop rate, throughput (and application defined)
- Latency computed w/ a Exponential Weighted Moving Average: $rtt = a * rtt + (1-a) * (sample)$
- Loss: loss rate is set to loss rate of last $k=100$ samples.

- Throughput: estimate $score = \frac{\sqrt{1.5}}{rtt \cdot \sqrt{loss/2}}$

RON Path

- Latencies add up along path
- Loss rate: Success = $1 - \text{loss rate}$. Multiple success rates along path
- Throughput: Score depends on latency/loss rate of path.

RON Routing

- Packets have type (application dependent), used for classification
- Each class can have separate policy (ex: don't use Internet2 for commercial traffic)
- So each class has only routing table of allowed routes and their metrics.
- Routing algorithm computes best route, puts it in forwarding table.
- Experimentally, most routes are at most 1 intermediate RON node

RON Architecture

- More application integrated

