BGP

Internet Routing

- Internet splits into *Autonomous Systems*(ASes)
- BGP is for Inter-AS Routing

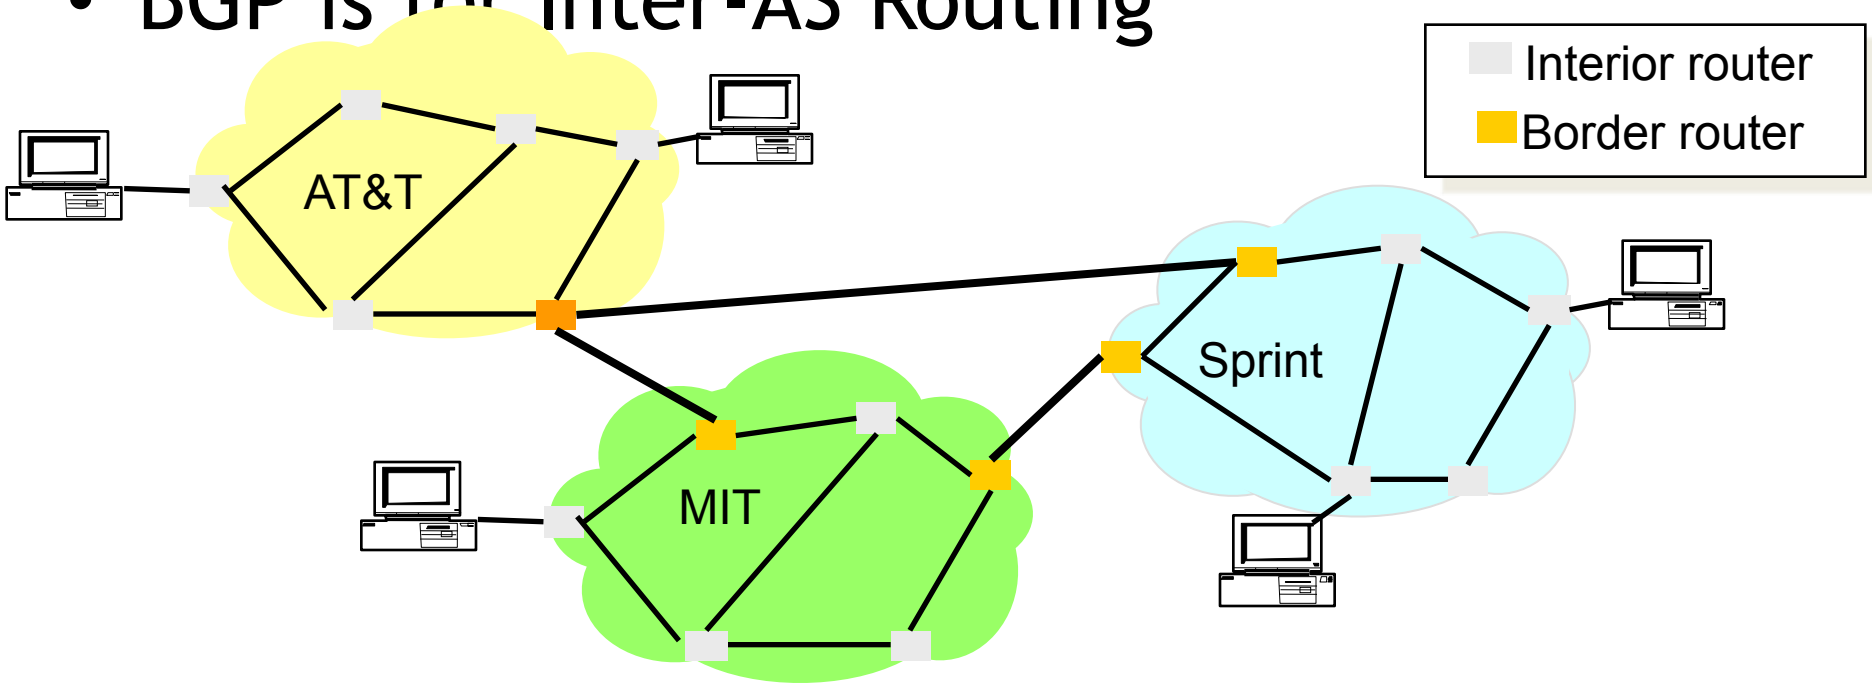

Inter-AS Relationships

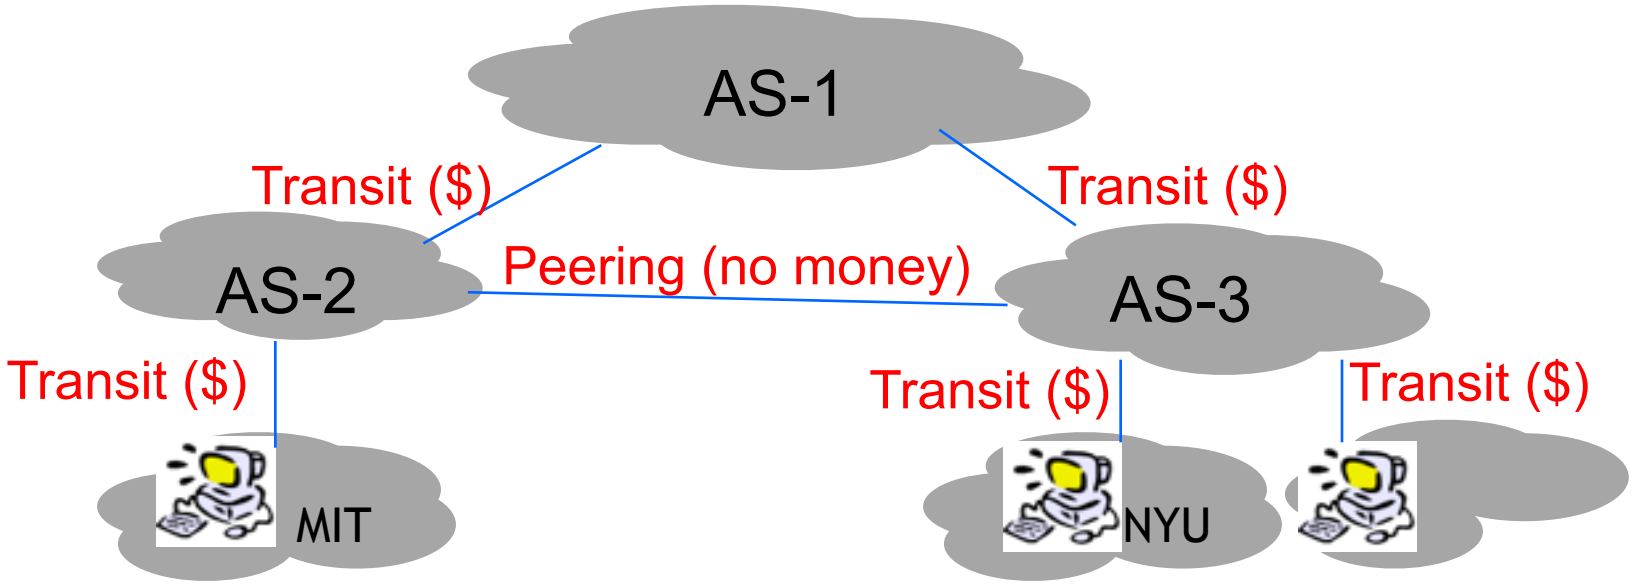

- Transit: $$ (service provider, customer)
- Peering: No $

BGP

- 2 parts:
  - Exporting (Incoming policy: Telling others about my routes)
  - Importing (Outgoing policy: Learning routes from others)
- 1 Simple Principle:

An AS does not transit traffic unless it makes money of it

Exporting Routes (Incoming Policy)

- Two rules:
  - AS tells everyone about routes to its customers
  - AS tells customers about every routes it knows
- Policies: To Whom an AS should tell which routes

To Whom:	Tell What:
Providers	*Only routes to customers*
Peers	*Only routes to customers*
Customers	*Every routes AS knows*

Importing Routes (Outgoing Policy)

- For each dst. prefix, AS picks its preferred route from those in its routing table as follows:
  - Prefer route from **Customer > Peer > Provider**
  - Then, prefer route with shorter AS-Path
- Rule 1 supersedes rule 2
  - Making money is always most important!

Sample Problem (2012 Q2)

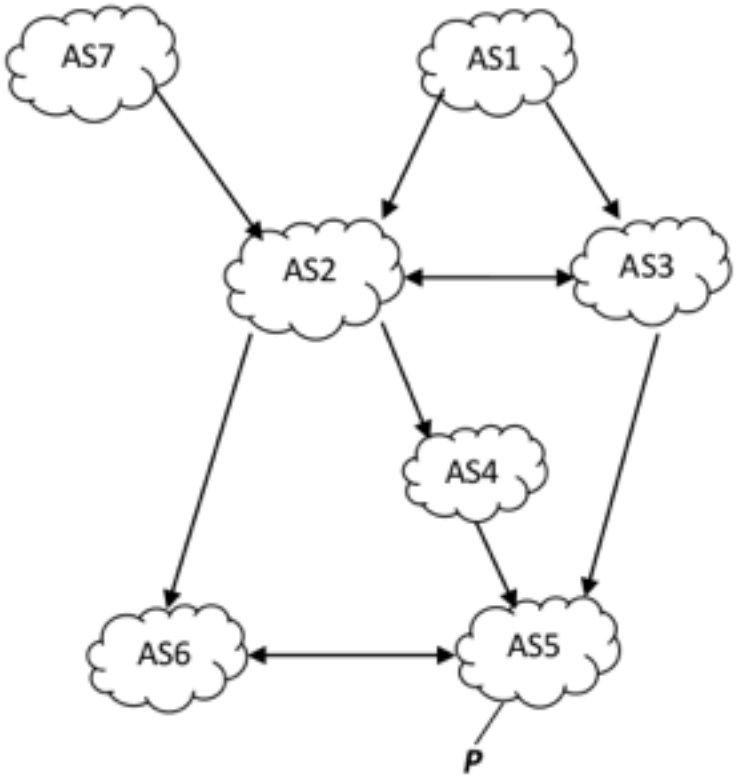

- Single Arrow: from provider to customer
- Double Arrow: peering

Sample Problem (2012 Q2)

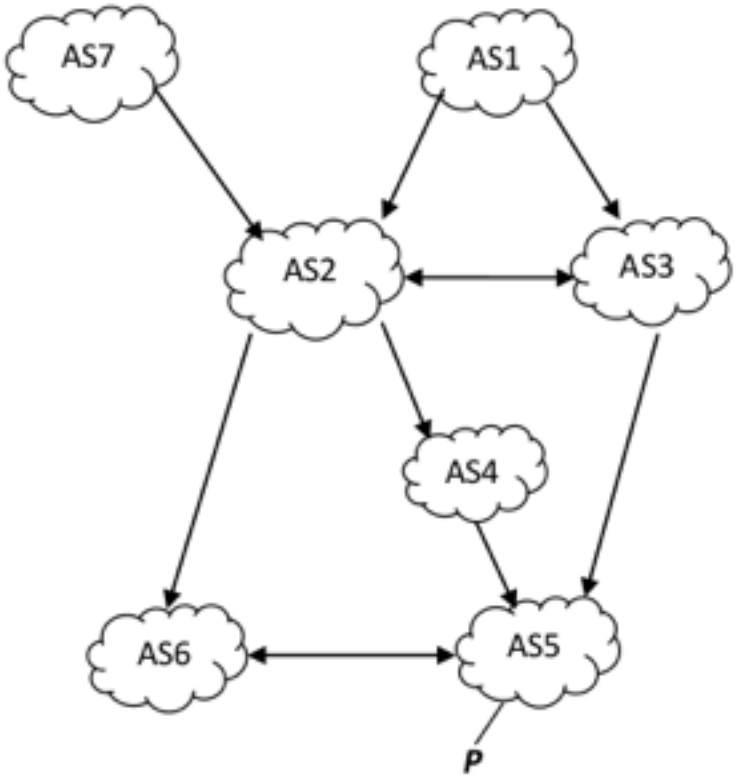

- **Q: Path AS1 to P?**
- AS1's routing table:
  - AS1-AS3-AS5-P
  - AS1-AS2-AS4-AS5-P
- Both of them are customer routes: pick the shortest!
- A: AS1-AS3-AS5-P

Sample Problem (2012 Q2)

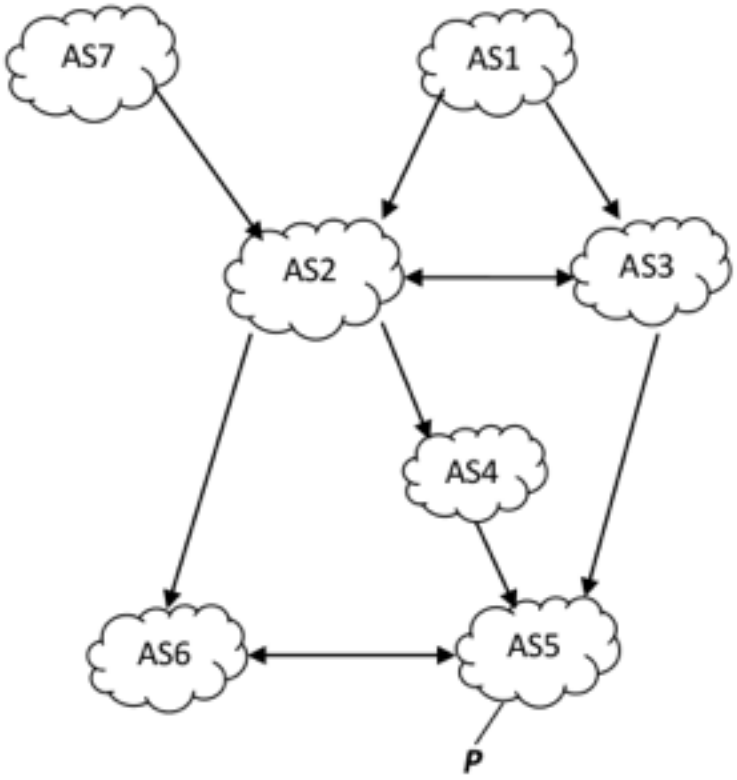

- **Q: What routes does AS2 learn for P?**
- A: 3 routes
  - AS2-AS4-AS5-P
  - AS2-AS3-AS5-P
  - AS2-AS1-AS3-AS5-P
- Preferred routes:
  - AS2-AS4-AS5-P

Customer > Peer > Provider

Sample Problem (2012 Q2)

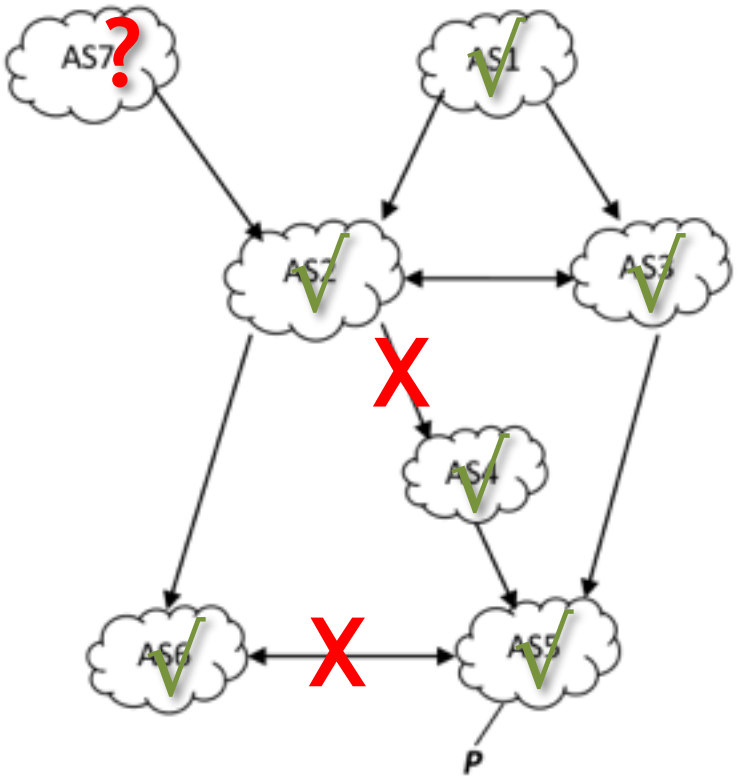

- Q: **If link AS2-AS4, AS5-AS6 goes down, who can reach P now?**
- A: Everyone except AS7

RON

RON

- Overlay network
  - Built on top of Internet (thus BGP)
- Application layer: application-specific
- Runs only at limited number of RON points
  - Doesn't scale well (N^2 probing / routing table)

Major Benefits over BGP

- Performance
  - BGP does not have a sense of “bad links”
- Better Recovery

This path is not allowed in BGP, but can be discovered in RON

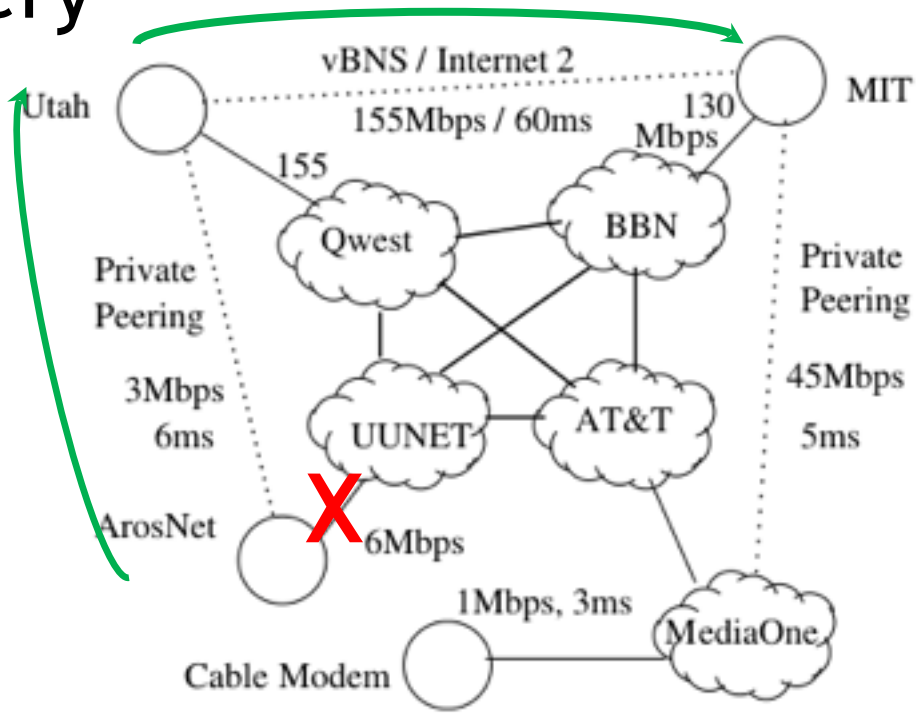

How Does RON Detect Outage?

- Active probing & passive observation
- Outage detection:
  - 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.

How Does RON Monitor Performance?

- Latency, packet loss, throughput, etc.
- Estimate link performances by defined metrics
- Combine link metrics to path metrics:
  - Latency: adding up
  - Loss rate: multiplying
  - Throughput: combination of latency / loss rate.

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

Sample Problem (2012 Q6)

True or False:

The authors' experiments determined that outages occur only in the inner core of the Internet, and never on last-hop links such as cable-modem access links.

X

The experiments weren't designed to evaluate last-hop failures, and the second dataset did encounter failures RON couldn't route around.

Sample Problem (2012 Q6)

True or False:

When an application requests a path optimized for high throughput, RON uses a path scoring metric based on combinations of measured loss rate and latency.

✓

Sample Problem (2012 Q6)

True or False:

In order to route around transient failures related to BGP reaction time in the core Internet, a RON network should include nodes at locations served by many different ISPs.

√

Often, different ISPs have independent failures. So it is very unlikely that two ISPs fail at the same time.

Sample Problem (2012 Q6)

True or False:

RON demonstrated that choosing between paths based on application-centric metrics such as latency and loss rate often improves application performance, and concluded that BGP should enact these same link measurements in order to scale up RON's performance gains to the entire Internet.

X

Sample Problem (2012 Q6)

True or False:

RON demonstrated that choosing between paths based on application-centric metrics such as latency and loss rate often improves application performance, and concluded that BGP should enact these same link measurements in order to scale up RON's performance gains to the entire Internet.

√

X

Some people might complain about the fact that RON violate (AS) policies; could force ASes to be transit providers when they don't want to be.