

6.033 Spring 2017

Lecture #9

- **Scalable Routing**
- **Policy Routing**
- **BGP**

Internet of Problems

How do we **route** (and address) scalably, while dealing with issues of policy and economy?



BGP

How do we **transport** data scalably, while dealing with varying application demands?

How do we **adapt** new applications and technologies to an inflexible architecture?




goal of a routing protocol: allow each switch to know, for every node *dst* in the network, a route to *dst*

goal of a routing protocol: allow each switch to know, for every node *dst* in the network, a **minimum-cost** route to *dst*

goal of a routing protocol: build a routing table at each switch, such that `routing_table[dst]` contains a **minimum-cost route** to `dst`

Distributed Routing



1. Nodes learn about their neighbors via the **HELLO** protocol
2. Nodes learn about other reachable nodes via advertisements
3. Nodes determine the minimum-cost routes (of the routes they know about)



problem: neither distance-vector nor link-state routing will scale to the size of the Internet

Scalable Routing

1. **path-vector routing:** advertisements include the path, to better detect routing loops
2. **hierarchy of routing:** route between ASes, and then within an AS
3. **topological addressing:** assign addresses in contiguous blocks to make advertisements smaller




18.0.0.0, ... ,18.0.0.255
↓
18.0.0.0/24

problem: ASes also need a means to
implement **policy**

Common AS Relationships


customer/provider (“transit”)

customer pays provider for transit



peers

peers allow (free*) mutual access to each other's customers




*as long as the amount of traffic in each direction is roughly equal

Export Policies

goal: make money


customer/provider (“transit”)

providers tell everyone about themselves their customers, and tell their customers about everyone



peers

peers tell each other about their customers



Import Policies




goal: make money

customer > peer > provider

(and then a variety of other attributes when this rule isn't sufficient)

Distributed Routing

1. Nodes learn about their neighbors via the **HELLO** protocol
2. Nodes learn about other reachable nodes via advertisements
3. Nodes determine the minimum-cost routes (of the routes they know about)



does BGP scale?

- To route on the Internet means to route at an enormous scale. We deal with scale via three techniques: **path-vector routing**, a **hierarchy of routing**, and **topological addressing**.
- **BGP** provides a means for autonomous systems to do **policy routing**. While the protocol is simple, how it works in practice is enormously complex due to competing economic interests, among other things.
- Though BGP works on the Internet today, its ability (or inability) to **scale** is becoming a concern as the Internet continues to grow.