

Flexible Forwarding

Tim Sally

<tsally2@illinois.edu>

Current State of Routing

- General idea: what to do with a packet when it arrives at an inbound interface?
- Border Gateway Protocol (BGP) and Interior Gateway Protocol (IGP) used to exchange routing information.

BGP Shortcomings

- Small networks can use complete iBGP graph, but this does not scale.
- Incomplete information results in suboptimal routing and protocol oscillations.

Some Solutions

- Replace BGP. Unlikely “flag day” is possible today.
- Use route reflectors.
- Use a Routing Control Platform (today’s paper).

Overview of Route Reflectors

- Divide routers into two classes, servers and peers.
- Forward updates from peers to servers only.
- Forward updates from servers to everyone.
- Servers should perform some logic to only forward the best routes.

Problems with Route Reflectors

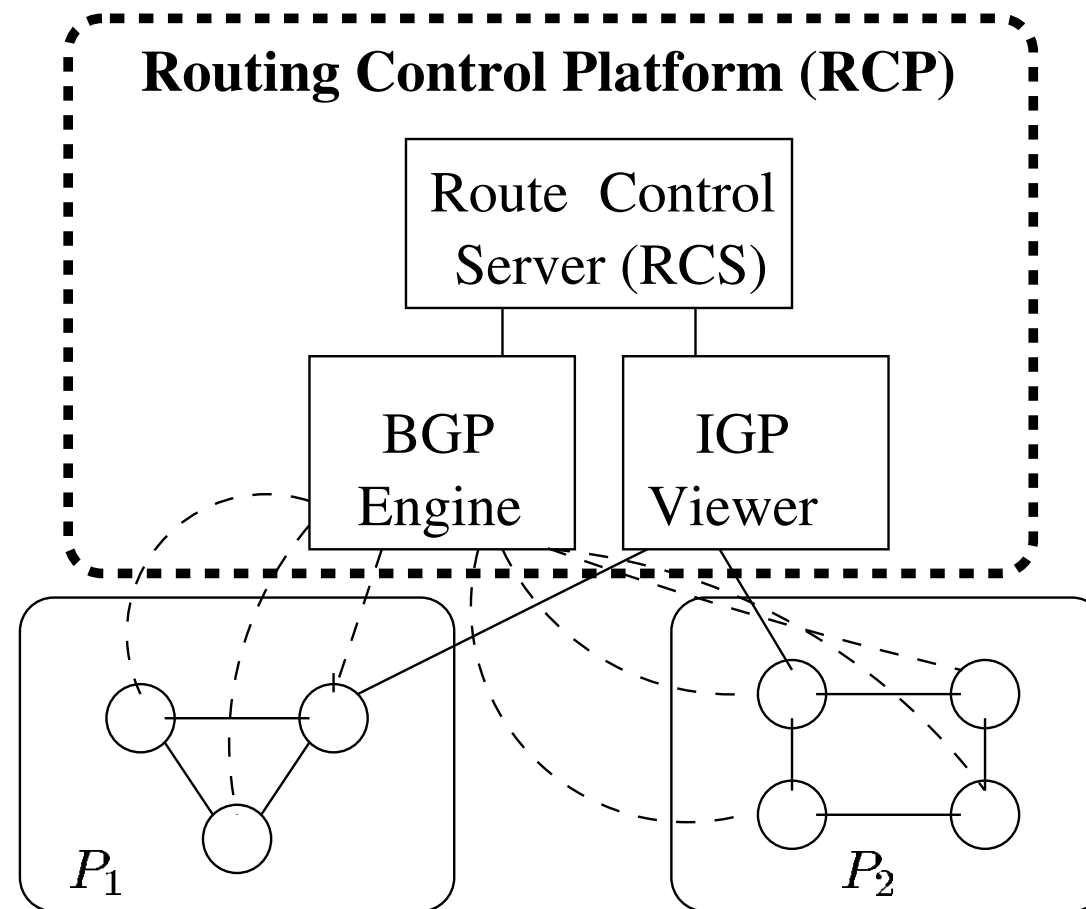
- Only an approximation of a complete graph.
- What if two peers along a path are assigned different routes by two different servers?
- We can stop inconsistent routes at the cost of efficiency with a large number of servers.

Routing Control Platform (RCP)

- A route reflector server that is connected to every router.
- Can send different routes to different routers.
- Makes the same decisions as a fully connected network.

RCP Components


- IGP Viewer to get IGP topology information.
- BGP Engine to learn BGP routes from routers and send new routes.
- Route Control Server (RCS) uses the information from above to compute optimal routes.



RCP Architecture

IGP Viewer

- Why do we even care about the IGP topology?
- Say we have two optimal BGP routes; we can pick either one.
- Pick the BGP route with the closet egress router according to IGP.



BGP Engine


- iBGP session with every router.
- RCP can learn about candidate routes and inform routers of optimal routes.
- We can forward different routes to different routers.

Route Control Server

- Must have both IGP and BGP information.
- Simply execute the BGP decision making process on behalf of all routers.

Network Operating System (NOX)

- Abstraction of network resources.
- Currently writing applications for networks is like using hardware specific assembly.



Active Networks

- Capsules (special packets) can program the behavior of the network.
- Interface is decentralized but code has to be certified by a central authority (or maybe executed in a “sandbox”).
- Trade off between capsule processing and packet forwarding at every router.

Thoughts on Flexible Forwarding

- Interoperability with the current network is a major concern.
- At the same time, we'd like to increase the number of things that are possible in the network.
- Seek to invent generalizations that cover current implementations and allow for new ones.

Thoughts on Flexible Forwarding

- Typically this involves some sort of control server (active networks take a different approach).
- Keeping any necessary logic and state at the control server minimizes changes that need to be made to the network.
- A global view of the network is often quite useful.

Thanks!