

SDX: A Software-Defined Internet Exchange

Nick Feamster
Georgia Tech

Jennifer Rexford, Scott Shenker, Russ Clark,
Ron Hutchins, Dave Levin, Josh Bailey

Problems with Interdomain Routing

Difficult to manage, troubleshoot, and debug.

- Security
 - BGP does not prevent a network from making arbitrary announcements
 - The forwarding path might not match the AS path
- Policy
 - Policies are too coarse-grained
 - Contracts result in market inefficiencies
- Stability
 - Even with stable inputs, BGP might not converge
 - BGP routes can oscillate within a single AS (e.g., route flaps)

A (Partial) Wish List for Interdomain Routing

- Better Peering
 - Peering for specific applications
 - More efficient pricing tiers, as opposed to “blended rate” pricing
- Better Control Over End-to-End Performance
 - “Remote control” peering: Content provider can affect route selection along the path, closer to access network/customer
- Better Security
 - Automatically prefer routes that have a higher reputation score (e.g., from hijack alert systems)
 - Incorporate checks for consistent route advertisement at peering points

The Promise of SDN

- SDN has reshaped many types of networks
 - Data Centers
 - Individual backbone networks
 - Others: Campus, Enterprise, Home, Cellular
- **What about interdomain routing, the protocol which has received so much attention for being so “broken”?**

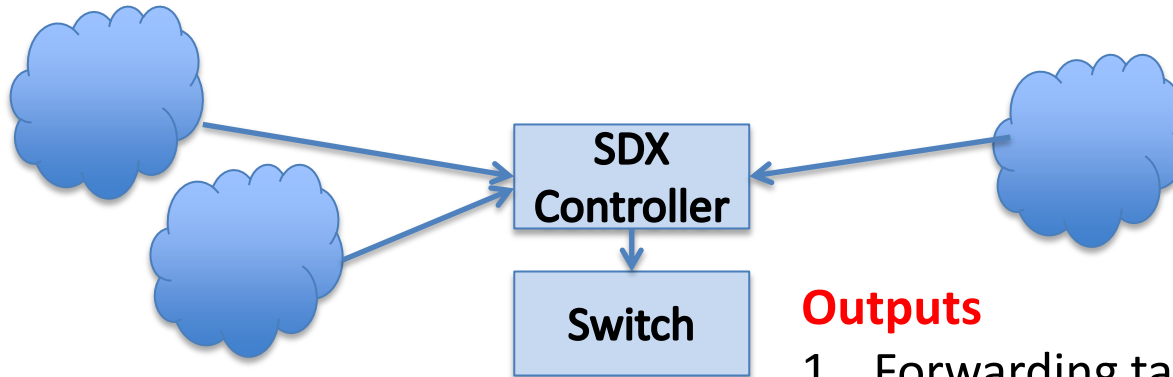
Proposed solution: A Software-Defined Internet Exchange (SDX)

SDX Controller Architecture: Inputs and Outputs

Like a route server, but with the additional capability of custom, per-peer route selection, and packet rewriting.

Inputs

1. Routes (via BGP) per IP prefix (including attributes like price, etc.)
2. Selection function



Outputs

1. Forwarding table entries in switch: One or more entries per AS that satisfy the selection for that AS
2. Packet rewriting (e.g., of destination IP address)


SDX Controller Architecture: Operation

- **Step 1:** Controller at exchange receives
 - BGP routes from all peers at the exchange
 - Auxiliary information (e.g., pricing, reputation, etc.)
- **Step 2:** Participant at exchange runs a function that executes at the controller to select route (and optionally rewrite packets).

Two possible architectures:

- One controller clearinghouse
- One controller per AS

Current Status



- Deployment at 55 Marietta Street in Atlanta, GA (SNAP)
- Two servers:
 - Floodlight controller
 - Virtual machine/network host
- Two OpenFlow switches:
 - Brocade
 - NEC
- Connectivity
 - 56 Marietta (TelX)
 - Southern Crossroads
 - Georgia Tech (via SOX)
 - Experimental rack at 55 Marietta

Ongoing Work

- **May 2013: Building the SDX**
 - Finish setting up basic connectivity between controller and SDN switches in exchange
 - Set up Mininet on servers in exchange
 - Basic BGP route exchange
- **July/August 2013: Using the SDX**
 - Start exploring use cases

Challenges: Building the SDX

- **Scaling:** Switch must perform per-AS forwarding, which causes state explosion in the forwarding table.
- **Controller architecture**
 - **Isolation:** How to ensure that each AS can apply route selection independently?
 - **Incremental deployment:** What happens when some exchanges are BGP, others SDX?
 - **Distributed computation:** How to perform route computation across multiple exchange points?
 - **Programming models:** Who is the programmer? (the ISP at the IXP, the content provider, etc.) What is the evaluation environment at the controller?

Challenges: Using the SDX

- Application-specific peering
- Avoidance routing (LIFEGUARD, Pathlets, etc.)
- Time-of-day peering/routing
- Balancing load across servers and data centers
- Secure routing
 - Route preference based on external inputs
 - Enforcement of export and preference policies

Summary

- Interdomain routing continues to be plagued by problems with security and manageability.
- An **SDN-based exchange** is promising for both fixing these problems and presenting new opportunities.
- Many research challenges remain, both for building the exchange and for using it.