CAIDA Ark IXP Active Measurement at an Internet Exchange Point



William Herrin Email: herrin@caida.org

Introduction – Ark IXP

- Internet exchange points give us a different view of the Internet than Ark nodes at the network edge
- But they don't work like the edge
 - -We only see part of the Internet
 - $-\ensuremath{\mathsf{We}}$ see multiple versions of part of the Internet
 - Asymetric: just because we see them doesn't mean they see us

Requirements

- Want to measure each IXP peer individually
 - When two peers offer a route we want to investigate both, not just the "best" one
- Want to understand when the return path can't reach us through the peer
 - E.g. when a downstream takes a partial table and defaults to an ISP not at the exchange
- Can't use the IXP-assigned address for measurement —May not be routeable beyond the exchange

Requirements

- Play nice on the IXP LAN
 - -Single IPv4, IPv6 and MAC addresses
 - -Even though we want to talk to each peer separately
 - -Without requiring them to configure multihop BGP

The Plan

- Deploy a server, not a Raspberry Pi
 - Make virtual Ark nodes in containers, one for each peer
 Announce an IPv4 /24 and an IPv6 /48
- One peer to one Ark container
 - -Source-route the BGP connection based on the peer's IP
 - Use NAT tricks to make the containers all talk to their respective peers as the server's primary IP address
 - –Use Linux "mangle" iptables to adjust the IP ttl to allow the extra hop through the host to the container
 - -FRR instance for each peer in each container

The Plan

- Assign each virtual ark node container 1 IP address from the /24 or /48
 - Scamper in the container uses that address to originate measurements
 - -Limits measurements to the routes in the table no default
 - Address identifies the peer to which a measurement was sent
 - -When the host receives the return, route it back to scamper in the container

The Plan

- The host (not the containers) has a default route to a full Internet transit provider.
- Also does BGP
 - -Announces less-specific prefixes covering the /24 and /48
 - Draws traffic sent to peers whose response fails to return to the peer
 - Doesn't disrupt the routing at the peer they have a more specific prefix

IXPs wanted!

 Looking for four volunteer IXPs who'd like to like to add to the general good in the world by supporting science.