



Scaling IXPs

Philip Smith

E2 Workshop, AfNOG 2007



Objectives


- To explain scaling options within the IXP
- To introduce the Internet Routing Registry at the IXP



IXP Scaling Techniques

- Route Collector
- Route Server
- Internet Routing Registry

Introduction to Route Collectors





Route Collector Background

- What is a Route Collector?
- Features of a Route Collector
- Purpose of a Route Collector
- IXP Design with a Route Collector



What is a Route Collector?

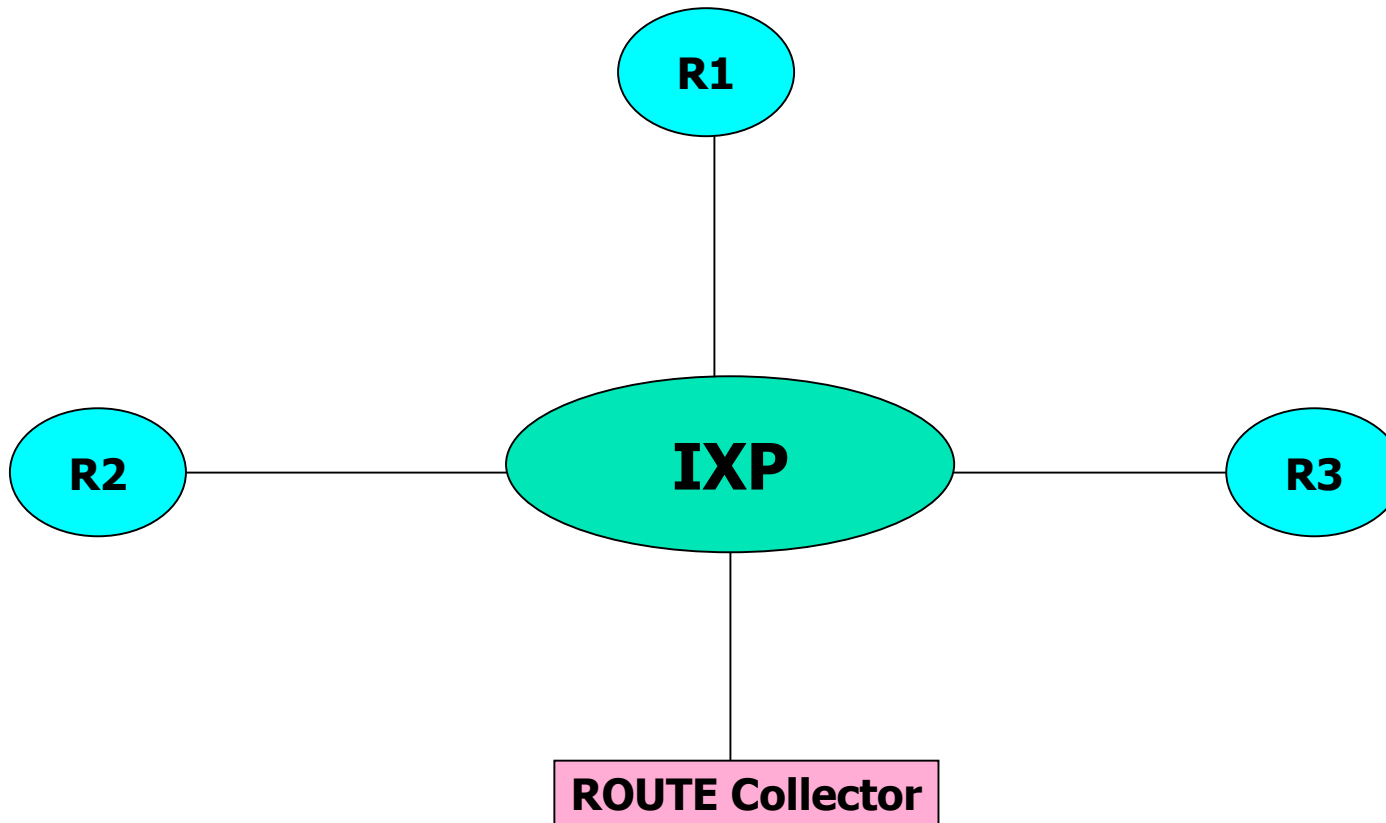
- Usually a router or Unix box running BGP
- Gathers routing information from service provider routers at an IXP
- Does not forward packets



Purpose of a Route Collector

- To provide a public view of the Routing Information available at the IXP
 - Useful existing members to check functionality of BGP filters
 - Useful for prospective members to check value of joining the IXP
 - Useful for the Internet Operations community for troubleshooting purposes
 - E.g. www.traceroute.org

Route Collector at an IXP





Route Collector Requirements

- Router or Unix system running BGP
- Peers eBGP with every IXP member
 - Accepts everything; Gives nothing
 - Uses a private ASN
 - Connects to IXP Transit LAN
- “Back end” connection
 - Second Ethernet globally routed
 - Connection to IXP Website for public access



Route Collector Implementation

- Most IXPs now implement some form of Route Collector
- Benefits already mentioned
- Great public relations tool
- Unsophisticated requirements
 - Just runs BGP



Introduction to Route Servers

Route Collector plus more



Route Server Background

- What is a Route Server?
- Features of a Route Server
- Advantages of using a Route Server
- Exchange Point Design with a Route Server



What is a Route Server?

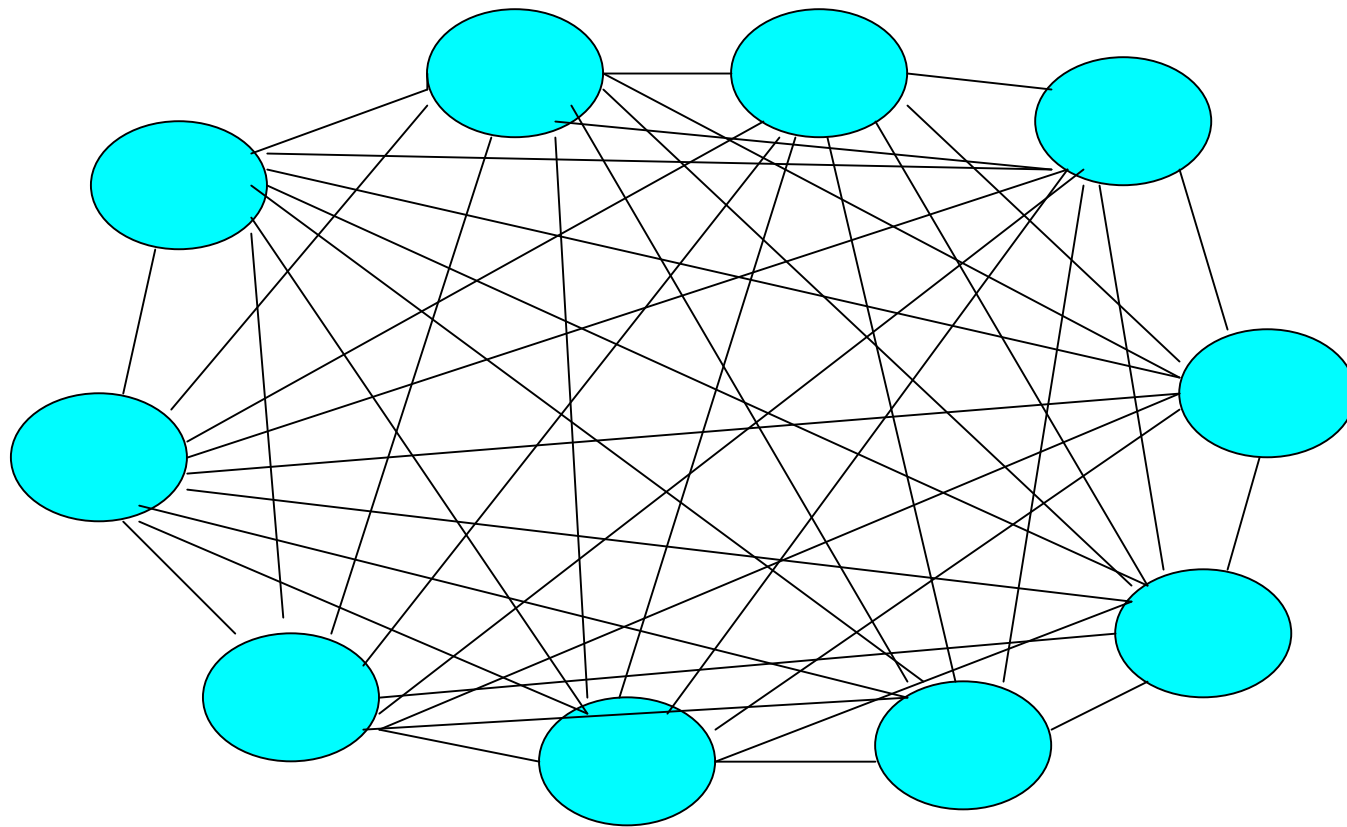
- All the features of a Route Collector
- But also:
 - Announces routes to participating IXP members according to their routing policy definitions
- Implemented using the same specification as for a Route Collector



Features of a Route Server

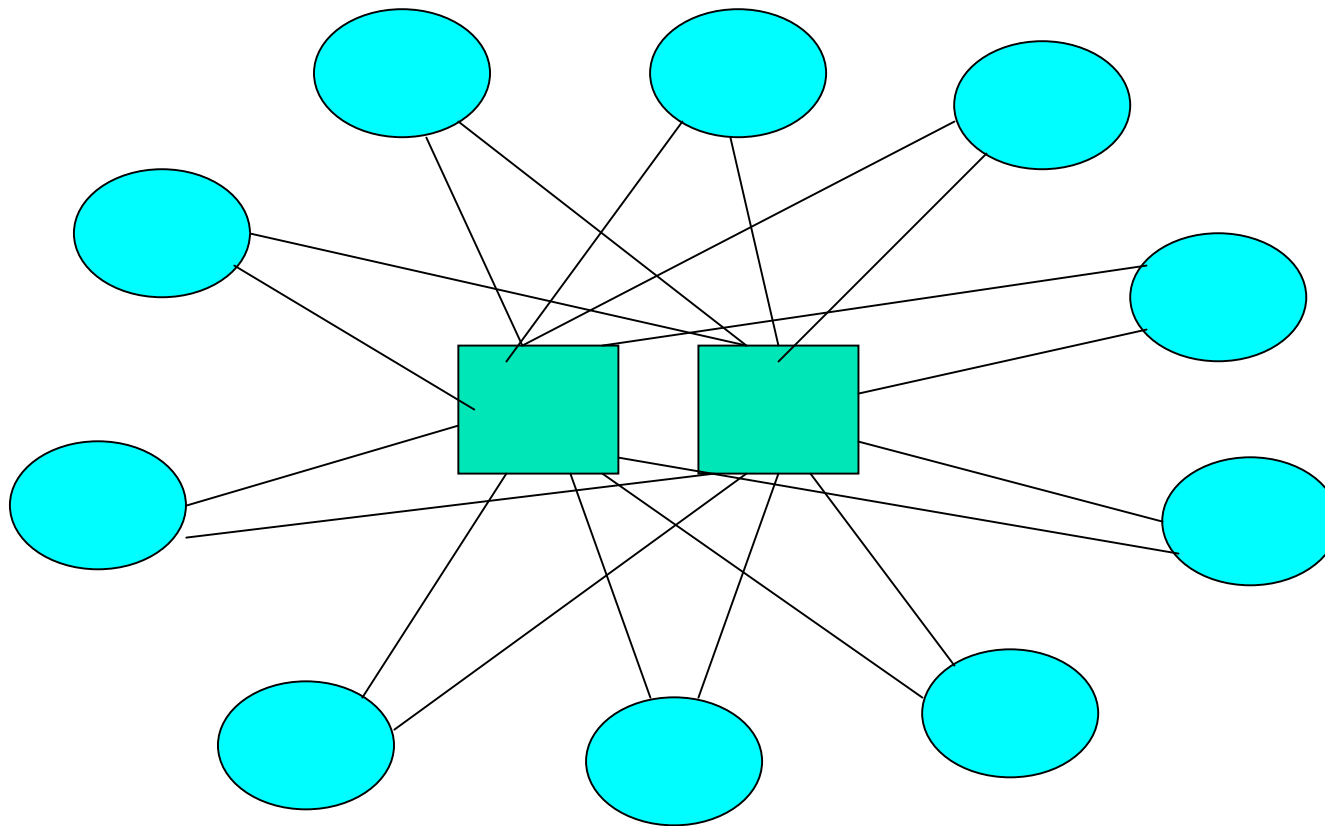
- Helps scale routing
- Simplifies Routing Processes on ISP Routers
- Insertion of RS Autonomous System Number in the Routing Path
- Uses Policy registered in IRR (optional)

Diagram of N-squared Mesh

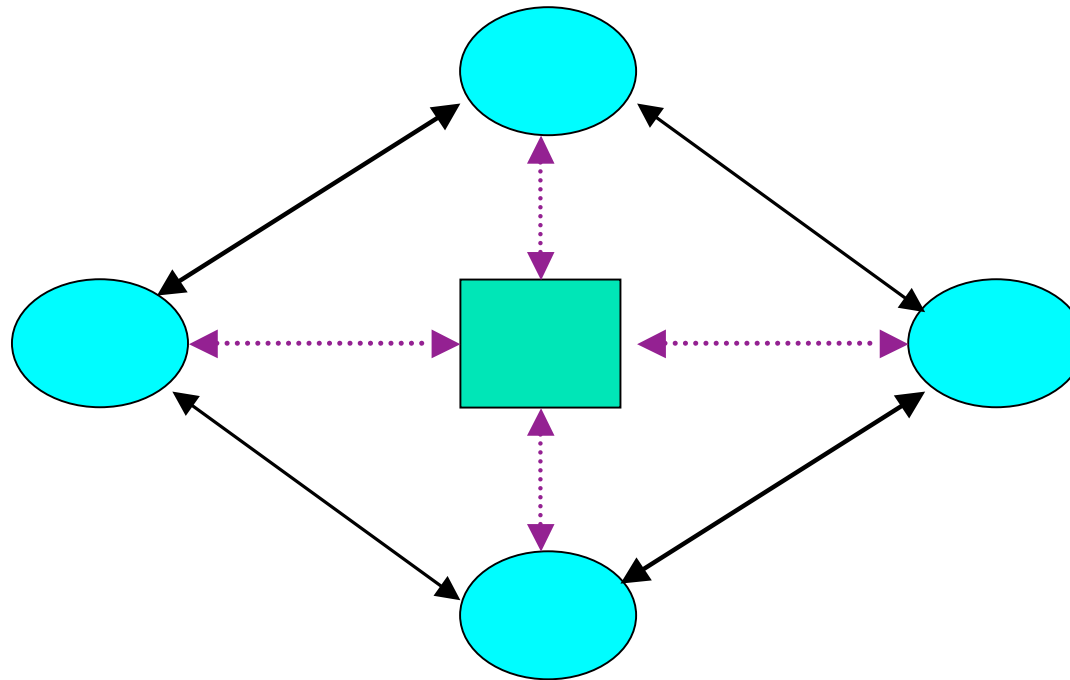




With the Route Servers



RS based Exchange Point Routing Flow



TRAFFIC FLOW

ROUTING INFORMATION FLOW



Advantages of Using a Route Server

- Helps scale Routing
- Separation of Routing and Forwarding
- Simplify Routing Configuration Management on ISPs routers
- Enforce Good Routing Engineering
- Helps prevent the spread of bogus routing information!



Disadvantages of Using a Route Server

- ISPs can lose direct policy control
 - Peer with all ISPs, want to or not
- Completely dependent on 3rd party for configuration and troubleshooting
- Insertion of RS Autonomous System Number in the Routing Path

- These are major disadvantages
 - They often outweigh the advantages



Peering with the Route Servers

- Any ISP attached to an IXP can peer with the Route Servers
- ISP must register their policy in the Internet Routing Registry
 - Most IXPs who provide the RS facility also provide a local IRR for policy registration
- Must use BGP



Things to think about...

- Would using a route server benefit you?
 - Helpful when BGP knowledge is limited
 - Avoids having to maintain a large number of eBGP peers
 - But can you afford to lose policy control?



Introduction to the IRR

The Internet Routing Registry



What is the Routing Registry

- Contact names, email addresses and telephone numbers for an AS
- Routing policy for an AS (what other ASes does it connect to, which routes do they exchange)
- Information about routes (most important is which AS originates the route)
- Several other types of information



What is the Routing Registry?

- Distributed database collectively known as Internet Routing Registry (IRR)
 - APNIC, RIPE, ARIN, RADB, etc
 - <http://www.irr.net/docs/list.html>
- Providers register routing policy
- Used for planning, debugging and generating backbone router configs



What is the Routing Registry?

- Can be used by anyone worldwide
 - debugging
 - configuring
 - engineering routing
 - addressing



What happens if I don't use the IRR

- Routing Horror Stories
 - AS7007
 - announcing bogus routes
- Inconsistent policy at network borders
 - Peers and upstreams need physical notification of policy changes
 - Mistakes easily made



So, I need to use the database because.....

- Filters generated off the IRR protect against inaccurate routing information
- Makes troubleshooting and debugging easier
- Keep track of policy
- Security
- **Filter! Filter! Filter!!**



Why Bother using the IRR?

- View of global routing policy in a single cooperatively maintained database
- To improve integrity of Internet's routing
- generate router configs
 - protect against inaccurate routing info distribution
 - verification of Internet routing
- Many providers require that you register your policy (or they won't peer with you)



Describing Policy

- Use the policy languages to describe your relationship with other Peers
 - routes importing
 - routes exporting
 - specific policies
 - interfaces, MEDs, communities
- register routes
 - with origin AS



Querying the Database

- `whois -h whois.ripe.net AS702`
- `whois -h whois.ripe.net AS1849-MAINT`
- `whois -h whois.ripe.net 158.43.0.0`



How to Register your IRR policy

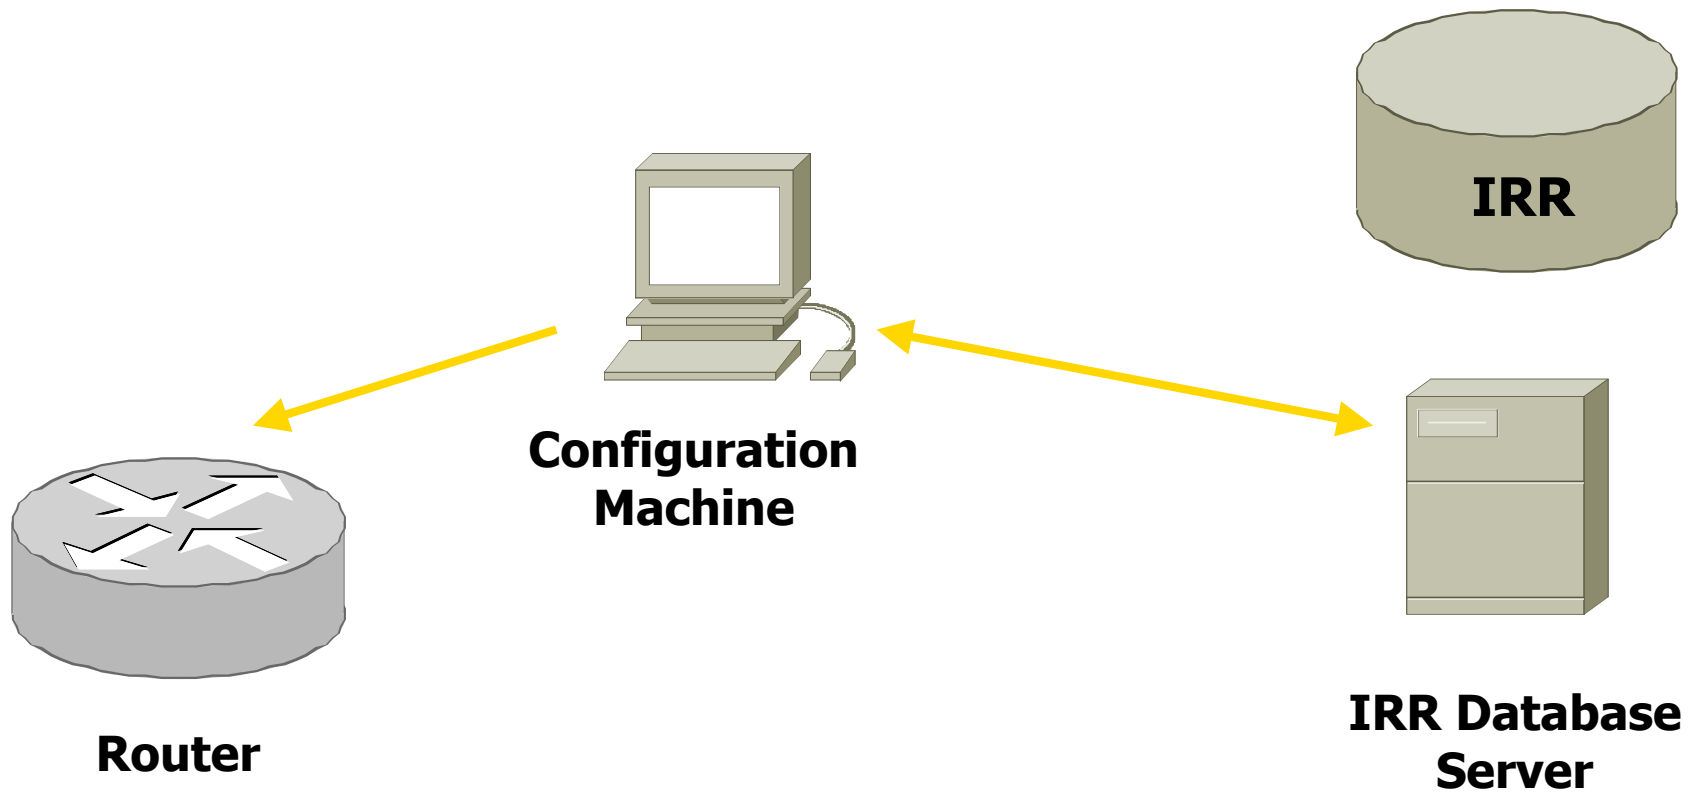
- Register one or more maintainers
- Register AS and policy information
- Register Routes
- Describes your import and export policy
 - At the very least, provides contact information



Router Configuration

- Currently configs by hand - slow and inaccurate
- Configuring routers using the IRR
 - lots of tools available!!!
 - IRRToolSet maintained by ISC
 - route and Aspath filters.
 - Import and export
- Filtering is a good thing...

Router Configuration





How do I use the IRR to generate configurations

- Tools available to generate config files for most BGP implementations
- IRRToolSet
 - <http://www.isc.org/sw/IRRToolSet/>
 - Started off as RAToolSet as a project of ISI
 - Moved to RIPE NCC custodianship and became IRRToolSet
 - Enhanced to support RPSL (RFC2622)
 - Now maintained by ISC



How do I participate?

- Set up your own registry
 - Private for your ISP?
 - Community for the region?
 - Download the software (from ISC)
- Use one of the many public IRR systems
 - Ask AfriNIC to set one up?



Things to think about...

- How would you register your policy?
 - Try to describe it in an aut-num object
- How would registering your policy benefit you? The community?