PCH
Packet Clearing House

Looking Glasses

Steve Gibbard
Packet Clearing House
http://www.pch.net/

PCH
Packet Clearing House

In this talk

- General Looking-glass information:
  - What is a Looking-glass
  - How to use a Looking-glass
  - Types of looking-glass data
  - Why use looking-glasses
- The PCH Looking-glass – my interest in this topic
- How to do your own.

PCH
Packet Clearing House

What is a looking-glass?

- Either a web page or just a router that allows users to look at a network's routing information.
  - See routing information from various network vantage points.
  - Generally permits all or a subset of "show ip bgp" commands.

PCH
Packet Clearing House

How to use a Looking Glass

- Telnet to it:
  - Best known example is route-views.oregon-ix.net.
  - Takes standard Cisco “show ip bgp” commands.
- Web
  - Most other looking-glasses.
  - Web interfaces specify what commands you can use.
  - More user-friendly
  - Less scary to operators

Routeviews screen shot

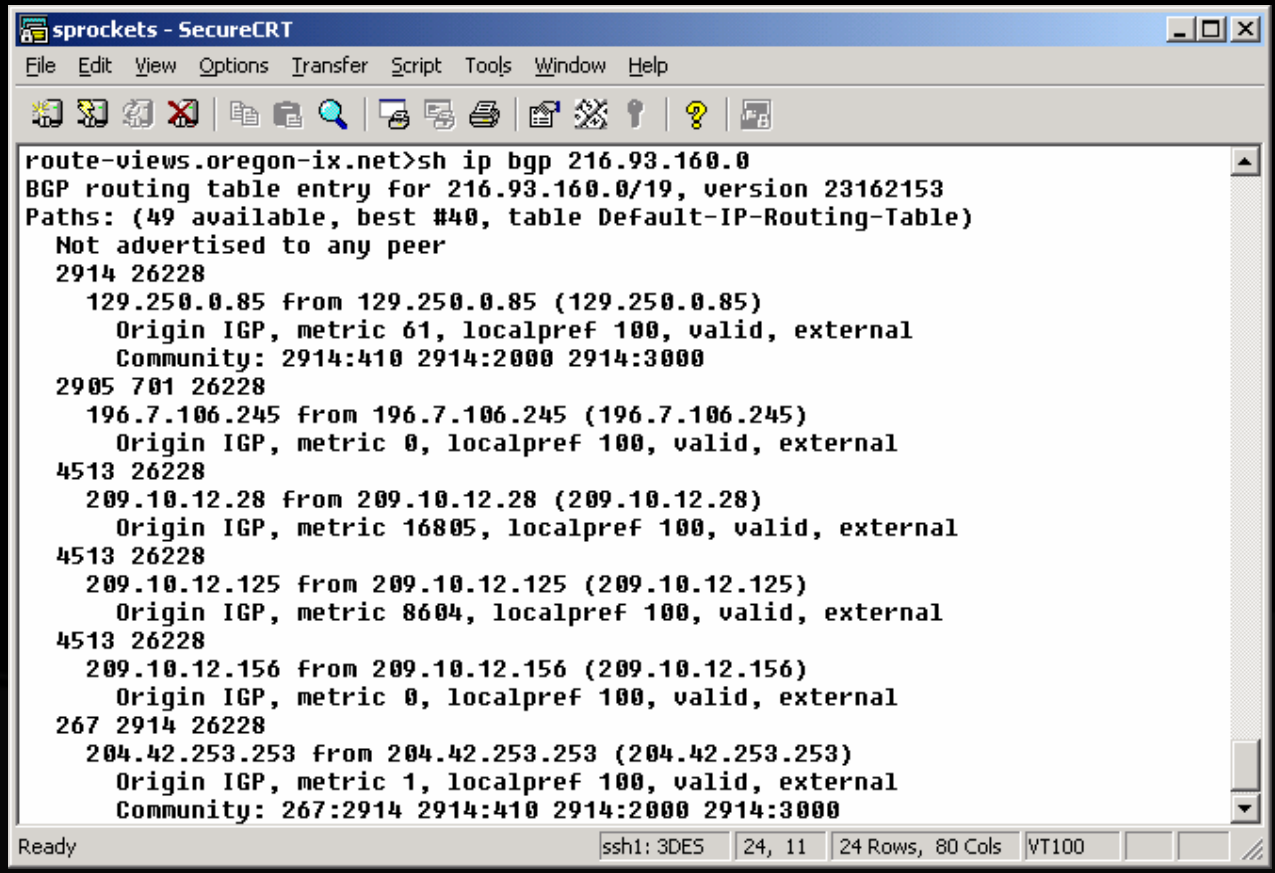

Web-based looking-glass

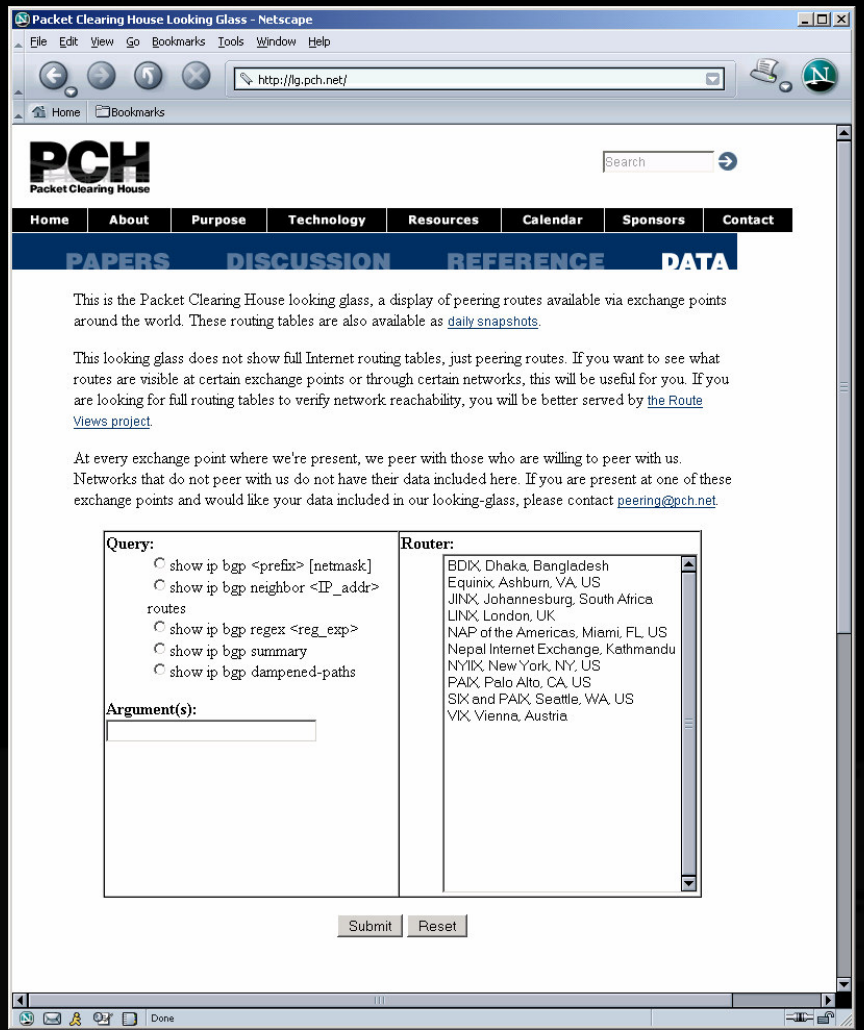

PCH
Packet Clearing House

Types of looking-glass data

- Full routes from one network.
  - What does my network look like to one other network?
  - Why can't this network's users get to me?
  - Same questions, for other networks.
  - Nitrous.digex.net was historical example of this.

PCH
Packet Clearing House

Types of looking-glass data (cont.)

- Full routes from many networks.
  - Variations on the above theme, but more information in one place.
  - Useful for seeing if a route is being generally propagated.
  - Shows only each network's view of the best route to a destination, as announced to the route collector.
  - Route-Views project is the best known example of this.

Special purpose looking-glasses

- Designed to display some specific type of information.
- What we do at PCH:
  - Only collect peering routes; no full routing tables.
  - More on this later.

Special purpose looking-glasses…

- Really cool commercial stuff:
  - Renesys
    - http://www.renesys.com
    - Visual maps of the AS tree
    - Historical and real time alerting of BGP announcements.
    - Todd Underwood does lots of presentations on this.

Special purpose looking-glasses…

- RIPE:
  - archives routing announcements and withdrawals too, but without as nice an interface.

PCH
Packet Clearing House

How to find looking-glasses

- If you're looking for information on a specific network, you could ask them.
- Huge list at http://www.traceroute.org/
- Looking Glass WIKI at http://www.bgp4.net/

PCH
Packet Clearing House

Why use a looking-glass?

- When you change routing announcements. Make sure your change worked, and that the world is seeing your network the way you want it to be seen.
  - Check looking-glasses of your upstream providers. Make sure they're seeing your announcements.
  - Check other looking-glasses and make sure your routes are visible.
  - Look at how much you've been flapping, whether your routes have been dampened, etc.

PCH
Packet Clearing House

Why use a looking-glass? (cont.)

- Troubleshooting routing issues. Why are people complaining that they can't get to your network?
  - Check whether your routes are still being seen. Is there a problem with your announcements? Did somebody make a filter change that's blocking you?
  - Do your routes look the same from everywhere, or are they being seen inconsistently?

Why use a looking-glass? (cont.)

- See how well connected a network is, and what you'd get by buying transit from it, or peering with it.
  - Do they have lots of diversity in their routing, or does everybody see them through the same transit AS?
  - Maximum diversity is not always optimal, and being single homed is not always horrible, but it's something to be aware of.

Why use a looking-glass? (cont.)

Seeing how well connected a network is…

- Do the ASes you're going through to get to some other network make geographic sense?
  - This is less of an issue in North America than in some other places.
  - Going through another continent to get across town is rarely a good thing.
  - Some ASes are pretty spread out so AS path isn't always an indication of geography
  - This is really a topic for another paper.

The PCH looking-glass

- Questions we keep being asked:
  - Aren't there enough route collectors already?
  - Why should I peer with yet another route collector?
- Answer: We're doing something a bit different.

The PCH looking-glass (cont.)

- A a route collection network, providing current and historical data.
- Peering data:
  - We only collect peering data. This gives us a view of what routes are available by peering at exchange points, or with specific peers.
    - Useful for network planning and research.
    - Probably not so useful for troubleshooting.

The PCH looking-glass (cont.)

- We collect that data from a lot of places.
  - Roughly 30 exchange points with our equipment.
  - Not all are fully operational yet.
  - As of mid-December, we have 293 peering sessions with 205 different ASes.

The PCH looking-glass (cont.)

- Mapping
  - What's connected where?
- Traffic/peering analysis:
  - Using a non-peered routing table for traffic analysis leads to missed peering paths.
  - We've got a view of what gets announced at which exchange points.

Our data could be more complete

- We peer with those who will peer with us.
- Some networks object to peering with us.
  - They can already see us through transit, and don't want more peering.
  - We don't have much traffic.
- Maybe this is ok. We can perhaps assume that those who don't peer with us don't peer openly, and don't really count for peering analysis.

How to look at our data

- Real time: http://lg.pch.net
- Archives of full tables: http://archive.pch.net.
  - Archive server isn't working very well right now, but we're working on that.

PCH
Packet Clearing House

Please peer with us

- It's easy. We don't ask for anything except your peering routes, so you can treat us like any other peer.
- It makes our data more complete.
- It lets your potential peers see what they could get by peering with you.
- If you peer with our looking glass, we'll also give you peering with our anycast DNS network.

PCH
Packet Clearing House

Added bonus: Anycast

- If you peer with our looking glass, we'll also give you peering with our anycast DNS network.
  - We host several TLDs, including some big ones.
  - We host some servers that will soon become anycast copies of I-Root.
  - This isn't much traffic, but it can make a big difference to your network's reliability.

PCH
Packet Clearing House

Methodology: What we do

- Cisco 1760 router as collector.
  - Can only pass around 20 Mb/s of traffic, but that's overkill for a route collector.
  - Holds 192 MB of memory, enough for lots of routes, and even full tables.
- Connect to lots of exchanges around the world – 4 U standard install includes collector, anycast router, anycast server, and switch.

PCH
Packet Clearing House

Methodology (cont.)

- Peer with every network that's willing to peer with us.
  - Send out lots of e-mail, see who replies.
  - Spend lots of time talking to peering coordinators at conferences and elsewhere.
  - Even easier if we're helping set up the exchange.
- Take only peering routes.

Methodology: Software

- Slightly modified version of RANCID looking-glass as user interface.
  - Looking glass CGI currently running on one location. Looking into anycasting it.

PCH
Packet Clearing House

If you want to do your own

- If you want to use a stand-alone route collector:
  - The 1760 works very nicely.
  - A Zebra or Quagga box would give you more flexibility, but might be more work to maintain.
- You could also point the looking-glass software at a production router, if you wanted to show the production router's view of the world.

PCH
Packet Clearing House

Looking-glass software

- RANCID software is available at http://www.shrubbery.net/rancid
- There's a good list of other looking glass packages on http://www.traceroute.org.
- CGI on web server will need access to log into routers.

Thanks!

- Contact information:
  - Steve Gibbard
  - scg@pch.net
  - http://www.pch.net/
- More Looking-glass information:
  - http://lg.pch.net/
  - http://www.traceroute.org/
  - http://www.bgp4.net/