

# **Internet mini-cores**

Local communications in the Internet's "spur" regions

**Steve Gibbard**

**Packet Clearing House**

# Current Internet Structure

- Well connected core
- Less well connected spurs
- In the core, connectivity is good:
  - Lots of fiber
  - Lots of redundancy
  - Lots of cheap bandwidth
  - The ability to send large amounts of data quickly between urban areas of the “developed world” can now be taken for granted.

## Current Internet structure (cont.)

- In the rest of the world:
  - Many ISPs in a region are likely to have connectivity to the core, but generally few connect to each other.
  - Connections to the core can go a very long way, sometimes via satellite.
  - Even “local” connectivity uses these connections.
  - Long distance connectivity is expensive and unreliable.
  - In these “spur” regions, the Internet often doesn’t work very well.

# Packet world tour...

```
1 <10 ms 10 ms <10 ms gw.conference.sanog.org [169.223.0.1] -- Kathmandu
2 270 ms 160 ms 100 ms gw-pck-sp.wlink.com.np [202.79.55.9]
3 40 ms 50 ms 40 ms gw.wlink.com.np [202.79.40.1]
4 231 ms 80 ms 110 ms mahesh.wlink.com.np [202.79.32.60]
5 892 ms 1131 ms * 216.236.105.33
6 1643 ms 1221 ms 1533 ms 69.88.1.189 -- Honolulu
7 3105 ms 1442 ms 1962 ms 216.236.111.25
8 1041 ms 912 ms 1432 ms hnl-edge-01.inet.qwest.net [67.129.94.1]
9 1222 ms 1322 ms 1131 ms bur-core-01.inet.qwest.net [205.171.13.169] -- Los Angeles
10 1062 ms 1031 ms 1022 ms lap-brdr-01.inet.qwest.net [205.171.213.46]
11 1432 ms * 2563 ms 203.208.168.185
12 1743 ms 1552 ms 921 ms 203.208.168.221
13 2784 ms 851 ms 1102 ms 203.208.182.133
14 1542 ms 1672 ms 1643 ms 203.208.172.138 -- Singapore?
15 1572 ms 1222 ms 1342 ms 203.208.145.38
16 1251 ms 1122 ms 1432 ms 203.208.140.133
17 1432 ms 1542 ms * 203.208.140.133
18 * * * Request timed out.
19 1713 ms 1602 ms * 202.52.242.65 -- Nepal
20 1683 ms 1742 ms 1533 ms 202.52.242.65
```

Trace complete.

# Example traceroute with peering

```
1 gw.sfo.gibbard.org (216.93.185.185) 0.214 ms 0.161 ms 0.173 ms
2 border-core2-ge6-0.sfo2.servepath.net (69.59.136.17) 0.659 ms 0.218 ms
0.639 ms
3 paix.pch.net (198.32.176.249) 2.941 ms 3.808 ms 2.408 ms
4 host.paix.pch.net (206.220.231.245) 2.521 ms * 2.468 ms
```

# Different from traditional phone networks

- Traditional phone networks:
  - Big cost advantage to making local calls.
  - True even as definitions of local have shifted.
  - Local phone calls tend to be pretty reliable; international calls are often a different story.
  - Few people notice when international phone networks break.
- The Internet:
  - “Distance is dead:” Local and long distance communications cost the same.
  - This is widely touted as a feature, and sometimes is.
  - Local communication becomes less reliable, more expensive, and slower than long distance communication.

# Examples

- Costs:
  - Urban US: Lots of traffic is local. Cost is around \$100 per Mb/s.
  - Northwest Montana (rural US): Not much local traffic. Transit cost is \$1,000 per Mb/s.
  - Kathmandu, Nepal:
    - International transit: \$5,000 per Mb/s.
    - For ISPs that peer, local traffic is \$50 per Mb/s.
  - Perth, Western Australia:
    - Transit: \$500 per Mb/s
    - Local traffic via peering: \$1 per Mb/s.
- Reliability:
  - Sri Lanka:
    - Fiber cut in harbor.
    - Outage of “Internet and international phone service.”

## Proposed new model

- Nothing wrong with the current core, for the parts of the world it covers.
- The rest of the world shouldn't have to send everything through it.
- A better model would be to have lots of "regional cores."
- Long distance circuits should be reserved for long distance traffic.
- Data sent between neighbors should not go to other continents.
- If it's going to replace the traditional phone network, local Internet connectivity needs to be as reliable.

## How to get there:

- Keep local traffic local:
  - Local exchange point.
    - All ISPs should have access to local peering.
    - This connectivity does not need to be direct. Buying transit from somebody who peers locally is sometimes sufficient.
    - Scales well. No hard limit on participants in the market.
  - Monopoly transit provider
    - Keeps traffic local – until somebody decides to compete with it.
    - Doesn't have much incentive to improve service, or lower costs.

## Exchange points aren't enough

- Keeping local traffic local doesn't help, if what you need to talk to isn't local.
  - Connectivity at layer 3 doesn't help if you're cut off from DNS.
  - Even with local DNS, Hotmail (or whatever) may not be local.
  - To be self-sufficient, a region needs its own "critical services."

# What is a critical service?

- DNS
  - Root.
  - Local ccTLD.
  - Any other zones in local use.
  - Use of domains without local DNS should be avoided.
- E-mail:
  - Local ISP's mail server is presumably safe.
  - Local Equivalent of Yahoo or Hotmail?
- VOIP
  - SIP server.
  - VOIP to PSTN gateways.
- What else?
  - Is Google a critical service?
  - Something for content providers to think about.

# Progress

- Exchange points being built in lots of places.
- Local TLD operators are hosting in their own regions.
  - With a local exchange point, this helps. Without one, it doesn't do much.
- Root servers are becoming more distributed.
- Local content providers are starting to host content locally in some places.

## More needs to be done

- Many regions still don't have local exchange points.
  - Without an exchange, other locally hosted services are of little value.
- Johannesburg and Jakarta are the only developing areas with root DNS servers (according to [www.root-servers.org](http://www.root-servers.org)).
- .com/.net footprint is still very small, as are many other gTLDs.

## Documentation required

- Internet users aren't conditioned to think of locations of services.
- "The local service is faster than the far away service" is easy to understand.
- Services for which location doesn't noticeably affect performance are the real reliability "gotchas." ISPs can help with this.

## Caveats

- This shouldn't be seen as an attack on long-distance communication.
- The ability to communicate easily over long distances is a very good thing.
- It just shouldn't have to be depended on for local communications.

# Thanks!

- Contact info:
  - Steve Gibbard
  - Packet Clearing House
  - [scg@pch.net](mailto:scg@pch.net)
  - <http://www.pch.net>