Introduction to Peering & Internet Exchange Points

LACNIC Annual Meeting

Monday, May 25, 2009

Bill Woodcock
Research Director
Packet Clearing House

ISP Lifecycle: Simple Aggregator

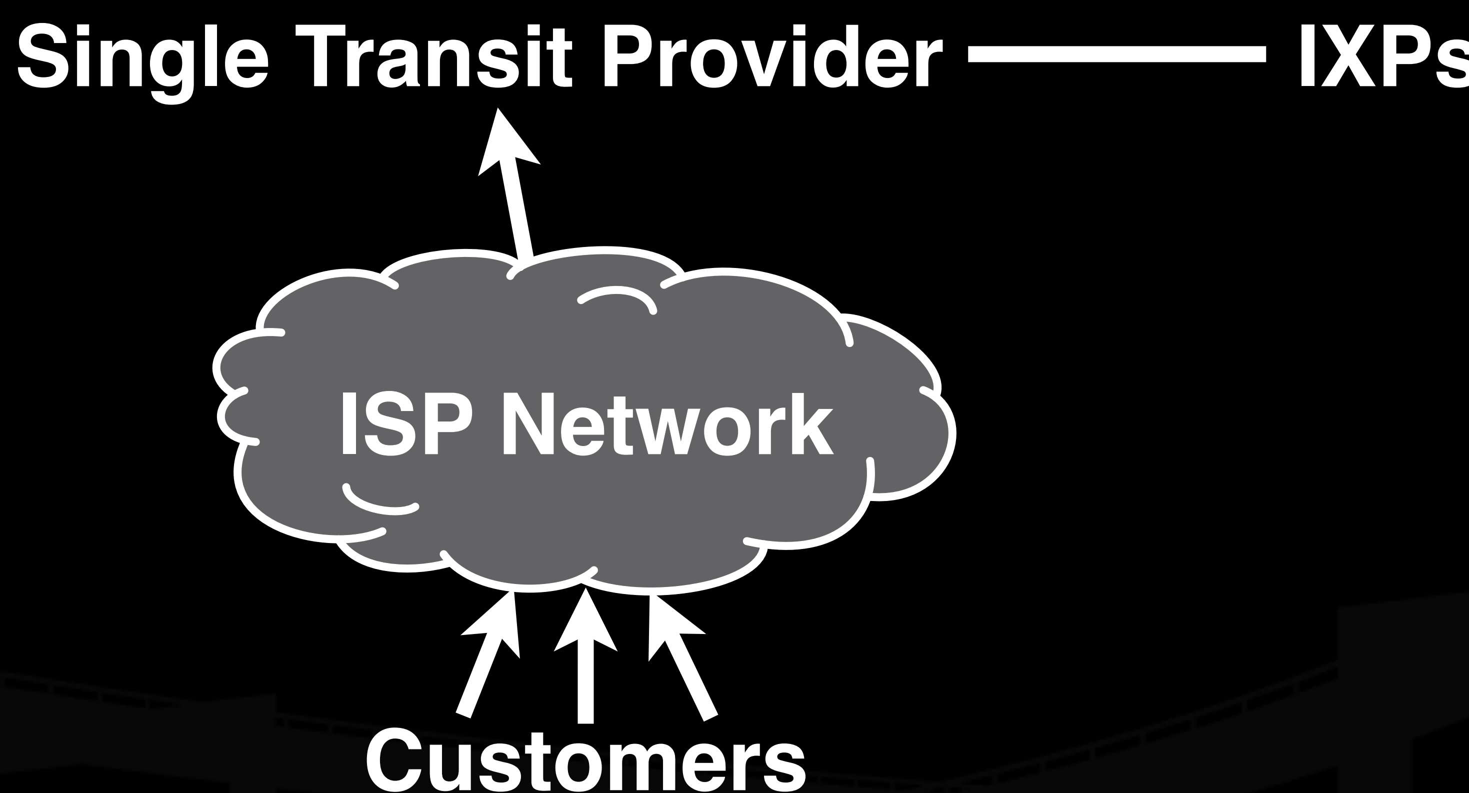

Packet Clearing House

ISP Lifecycle: Redundancy and LCR

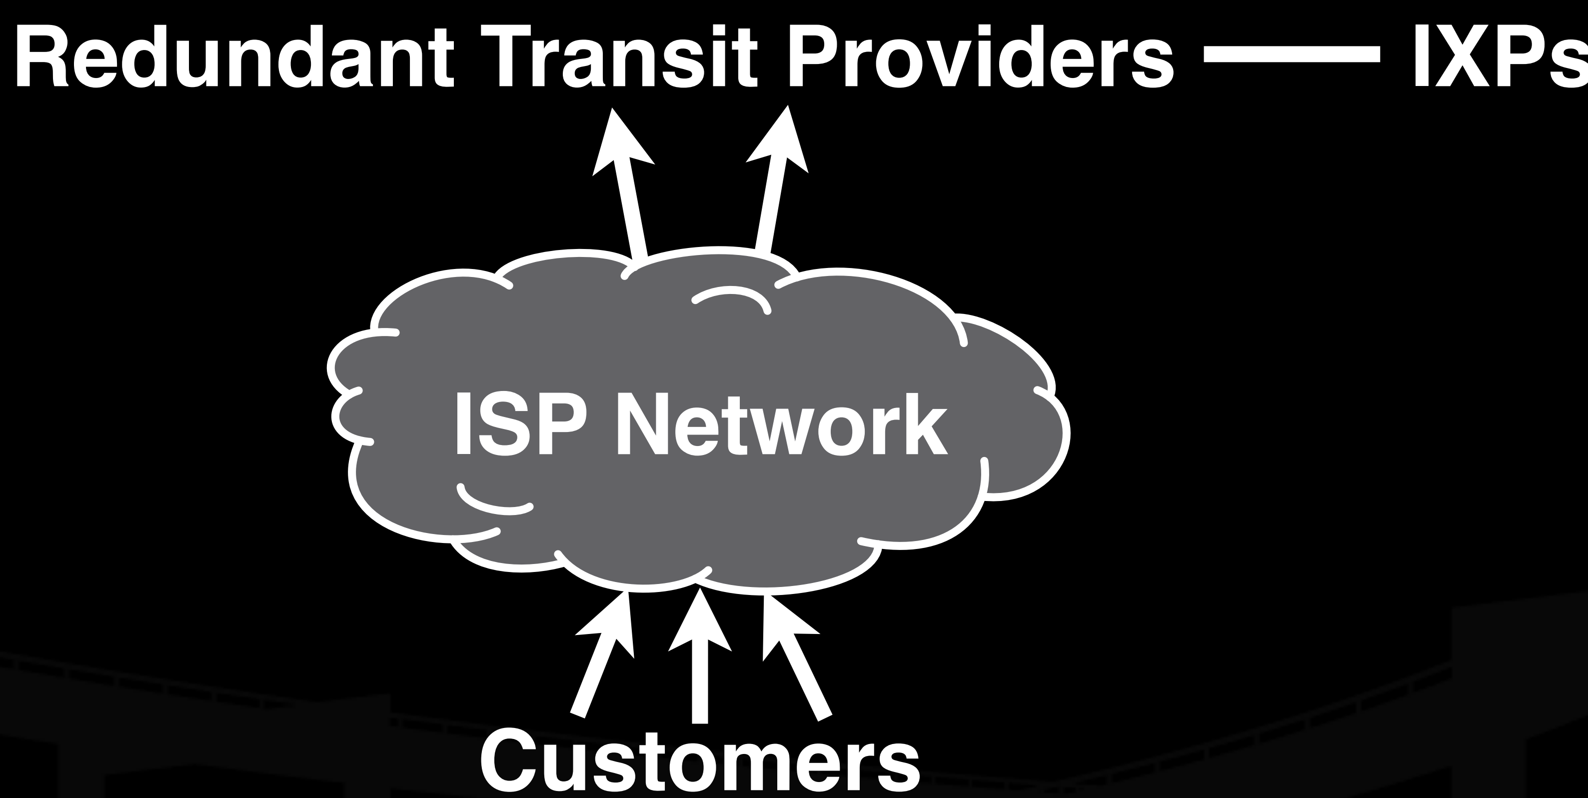

ISP Lifecycle: Local Peer

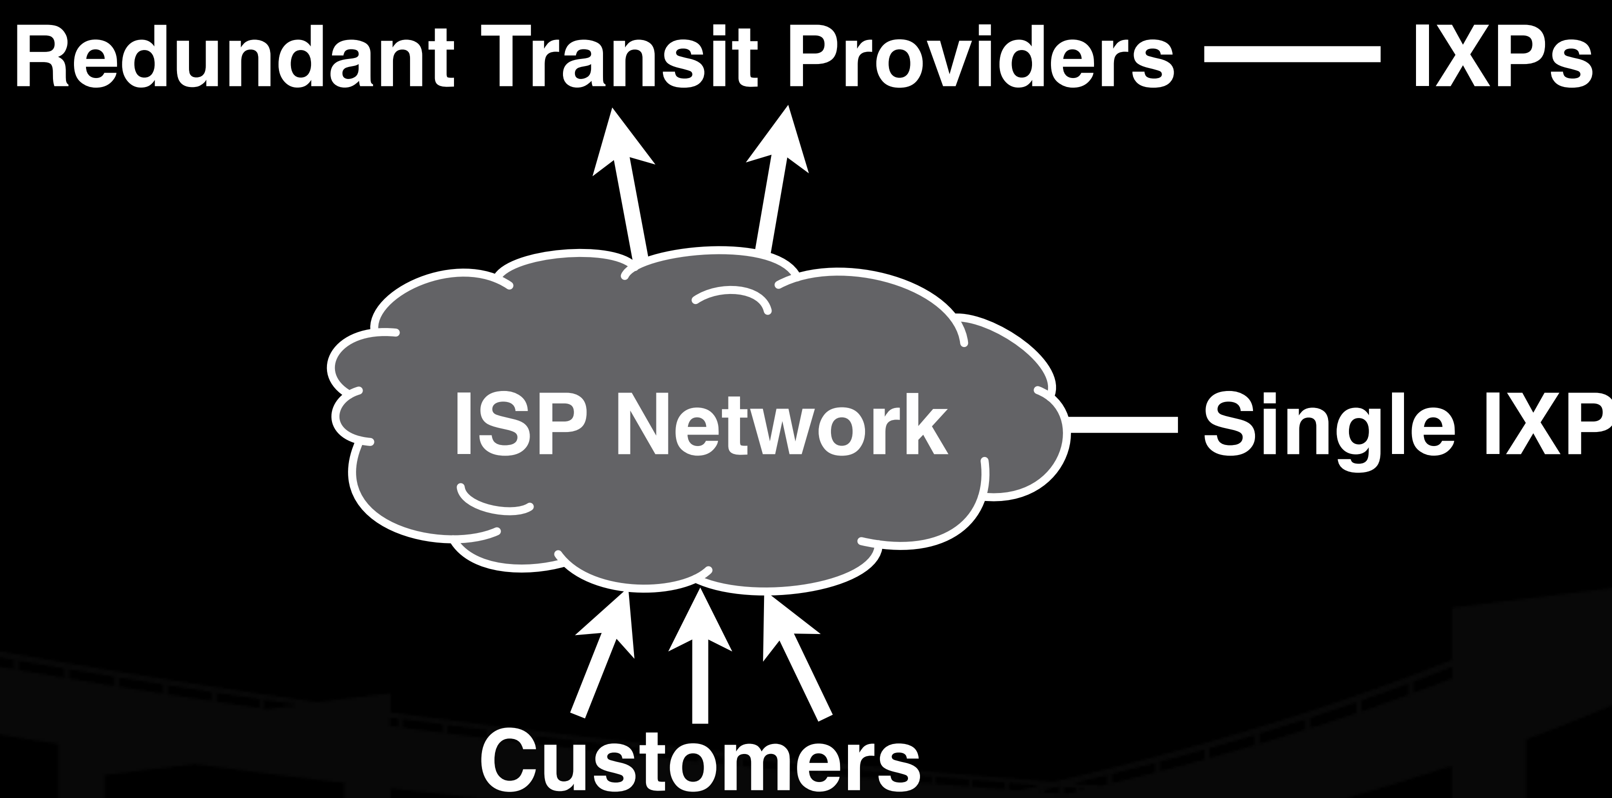

ISP Lifecycle: Network Service Provider

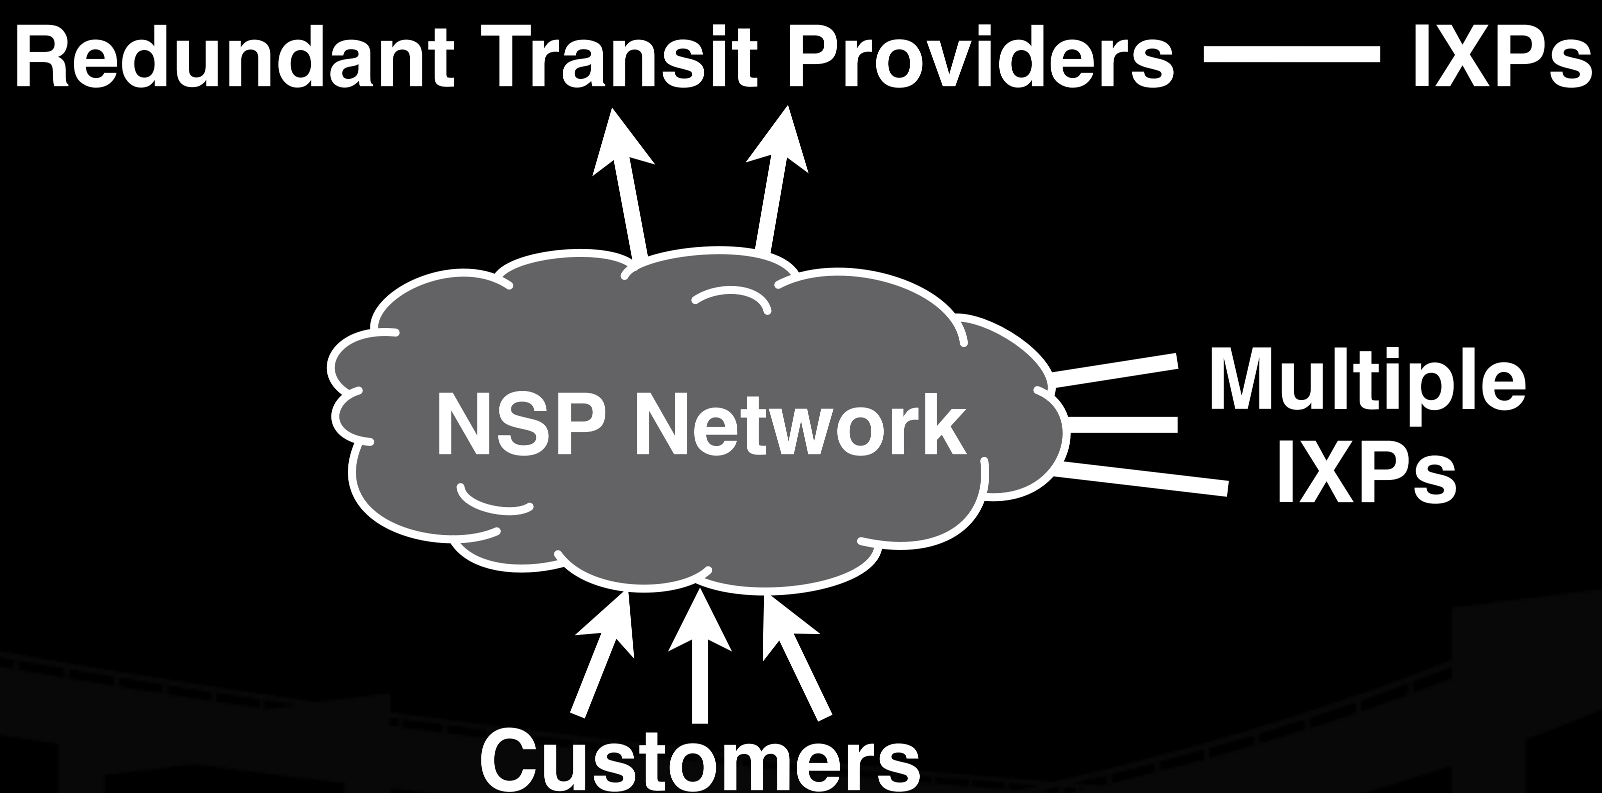

Hot Potato Routing

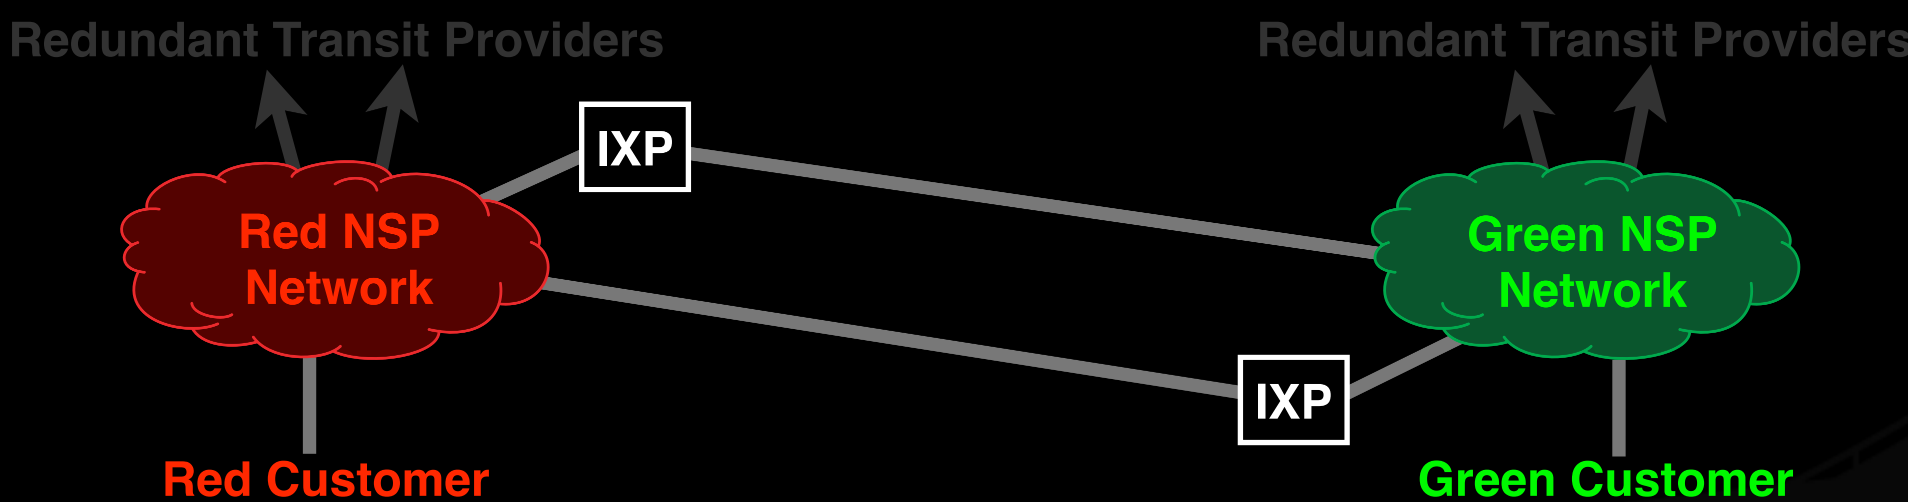

Packet Clearing House

Hot Potato Routing

Red Customer sends to Green Customer via Red NSP

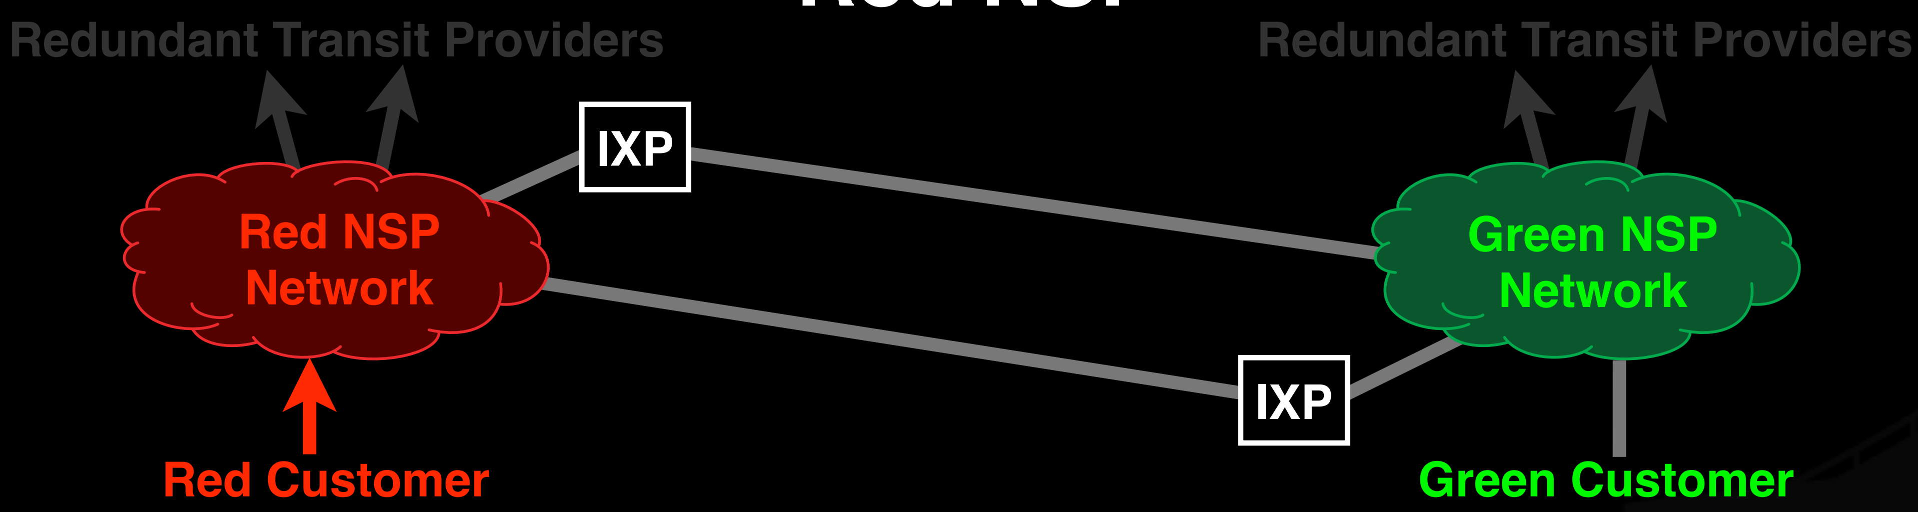

Hot Potato Routing

Red NSP delivers at *nearest* IXP

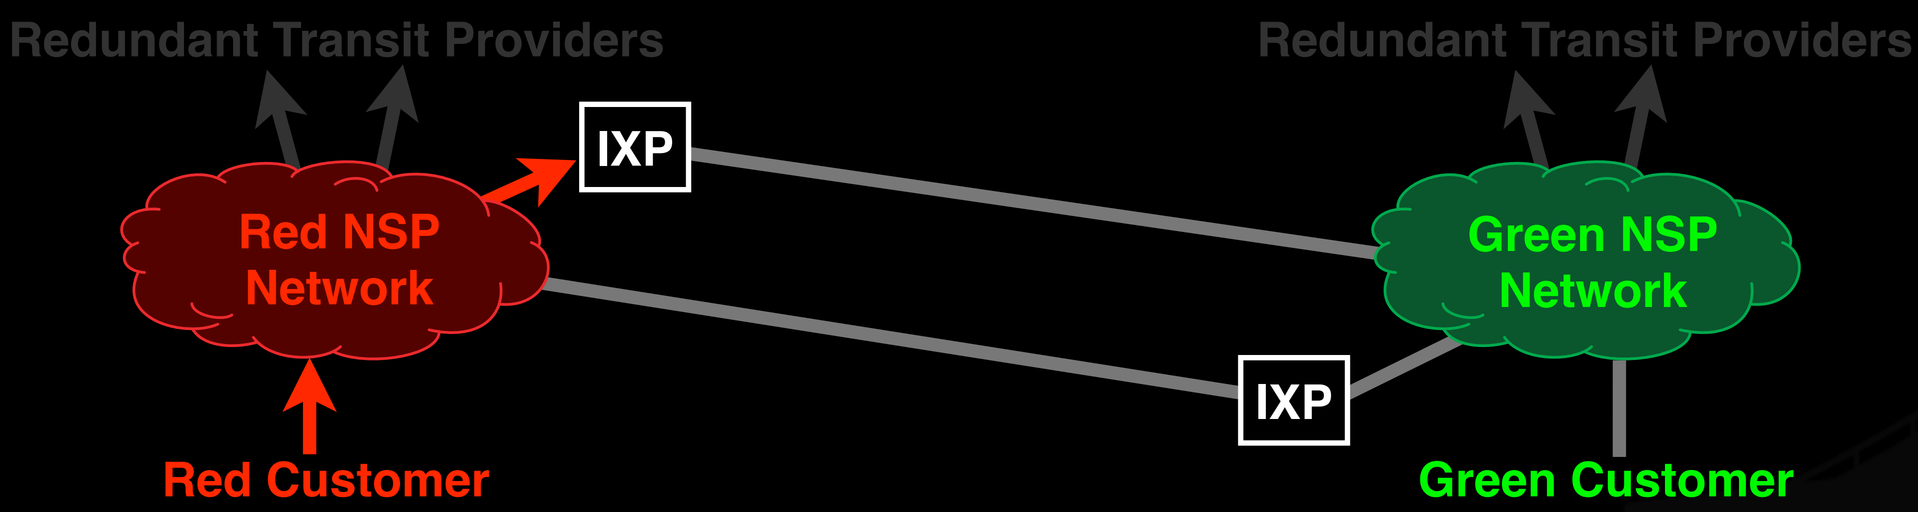

Hot Potato Routing

Green NSP backhauls from distant IXP

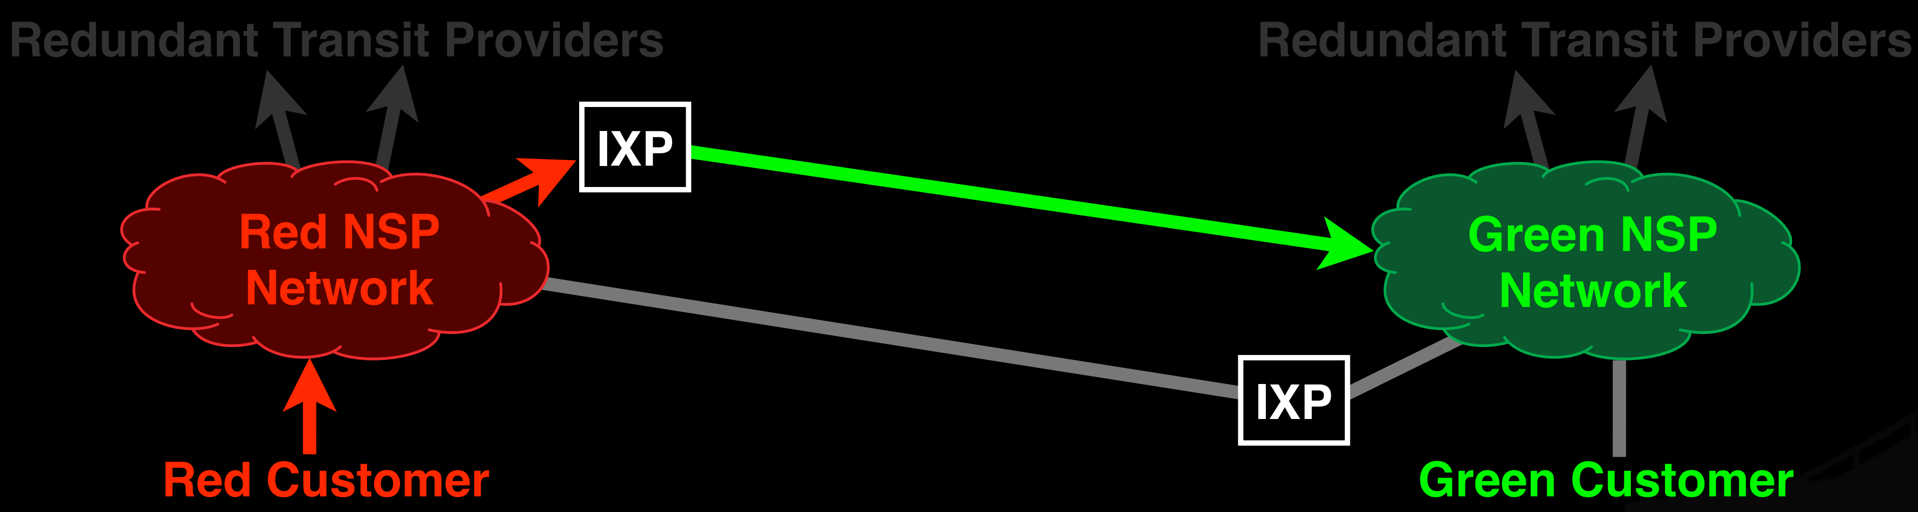

Hot Potato Routing

Green ISP delivers to Green Customer

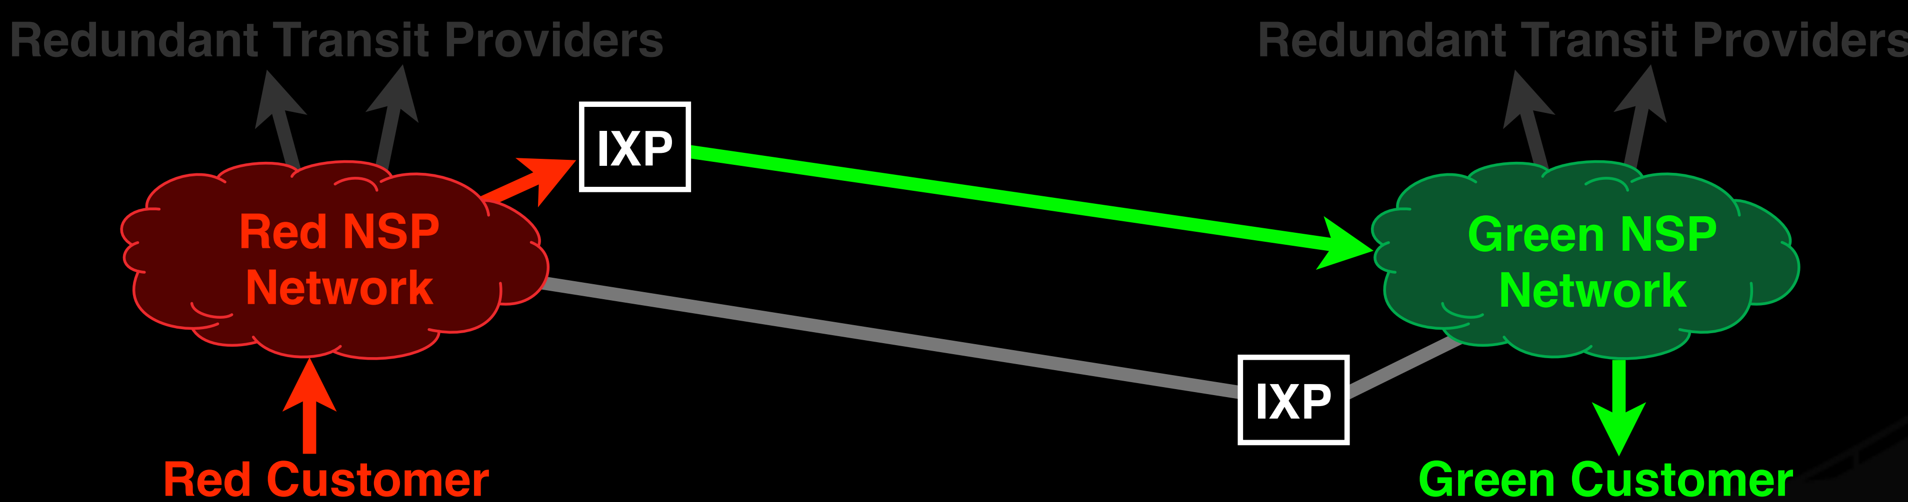

Hot Potato Routing

Green Customer replies via Green NSP

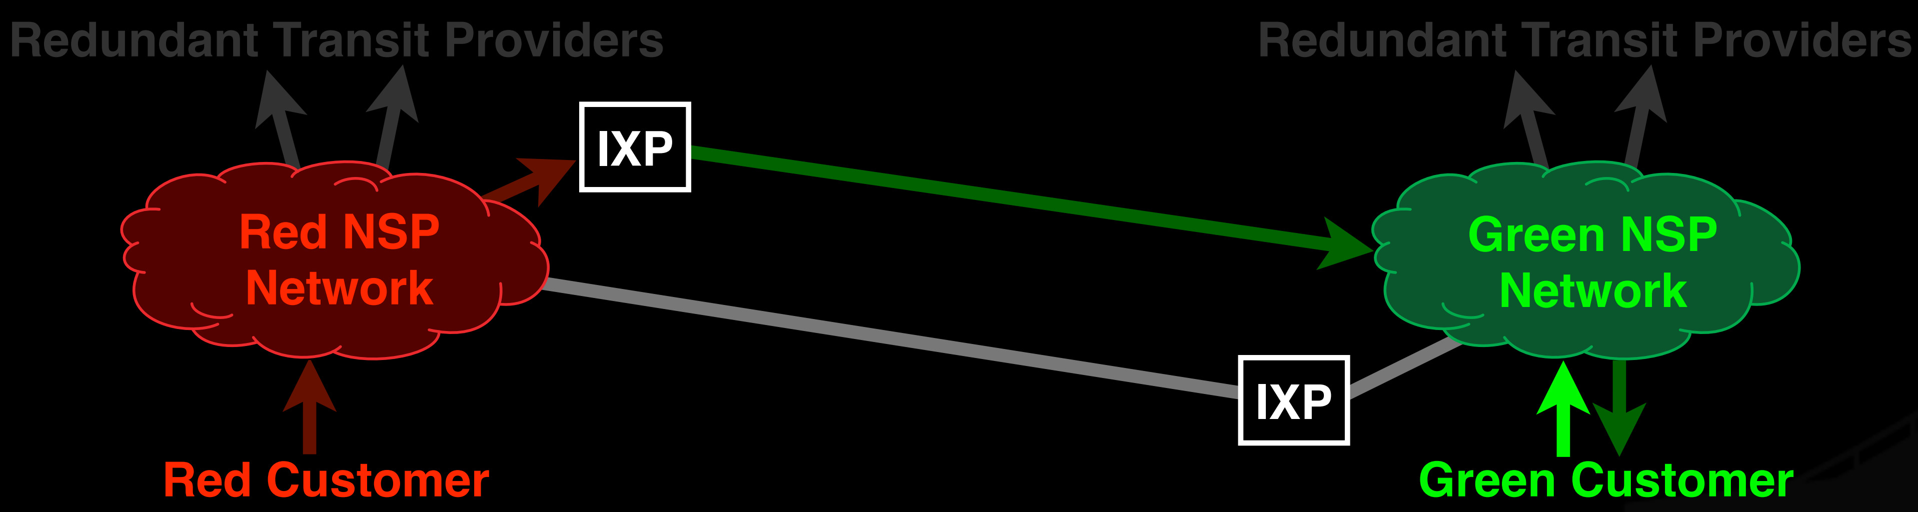

Hot Potato Routing

Green NSP delivers at nearest IXP

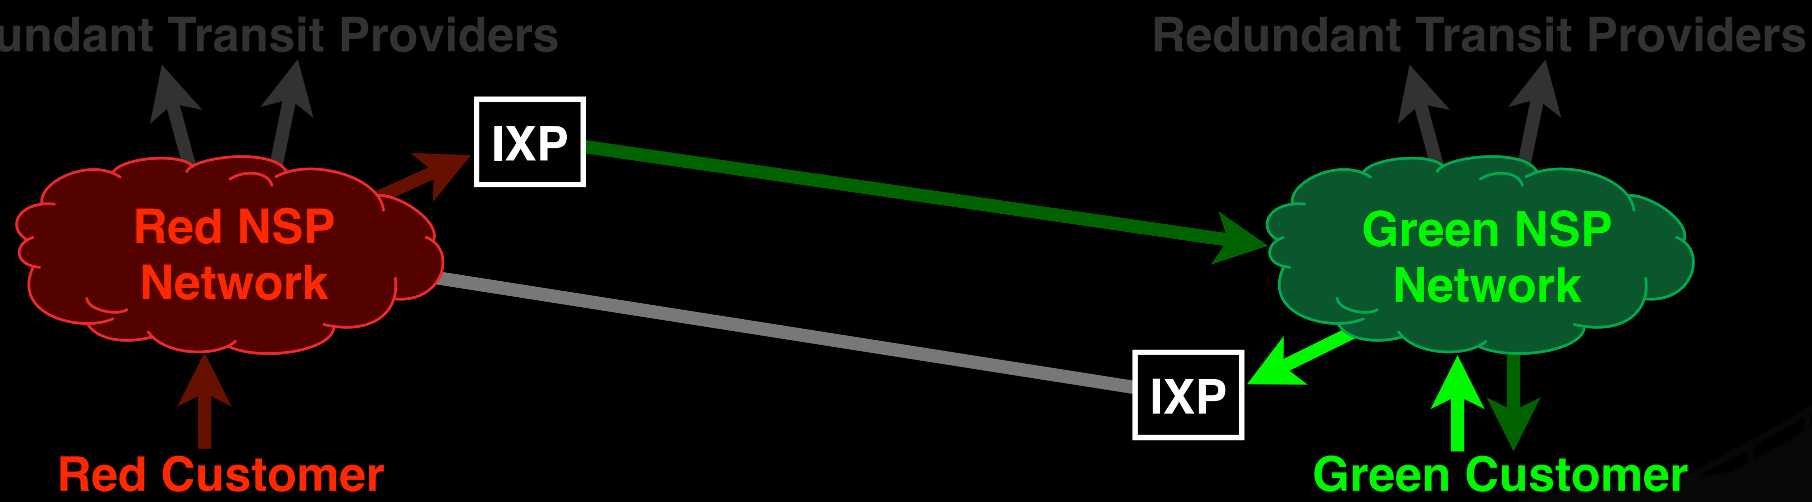

Hot Potato Routing

Red NSP backhauls from distant IXP

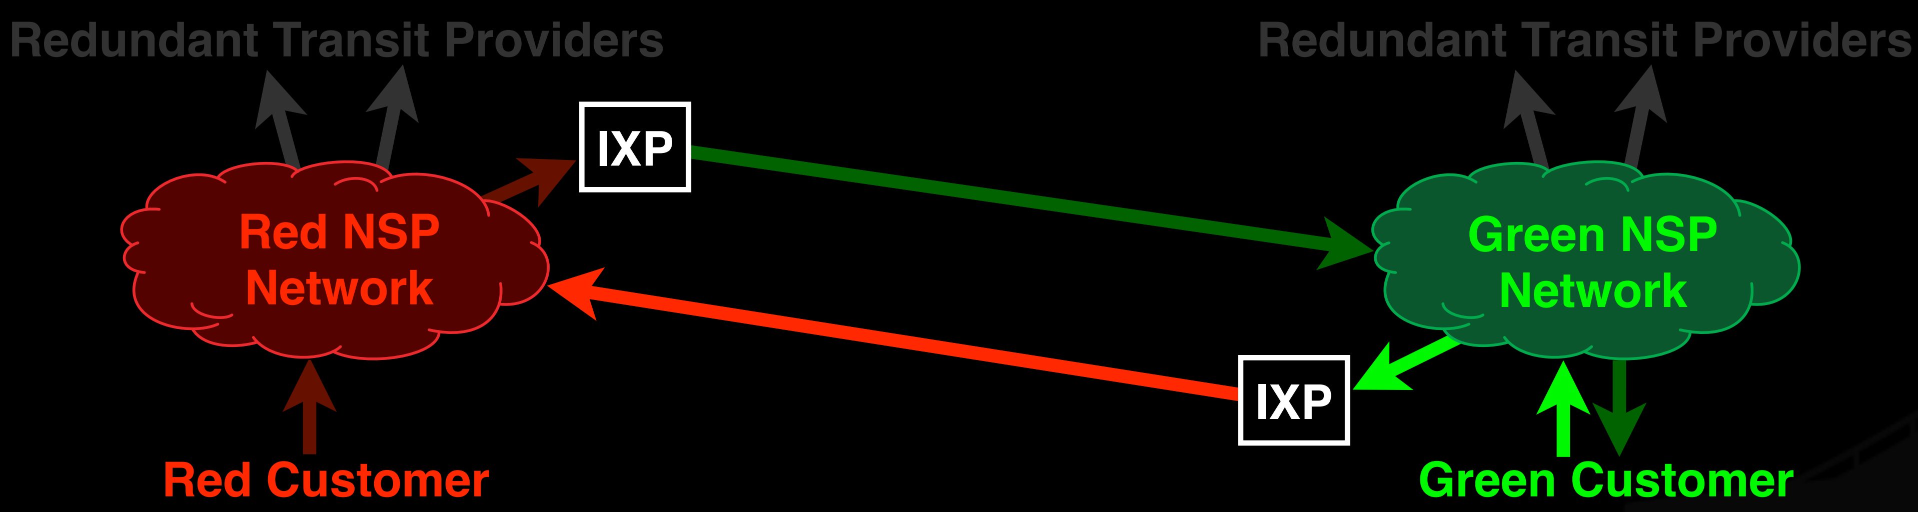

PCH Packet Clearing House

Hot Potato Routing

Red NSP delivers to Red Customer

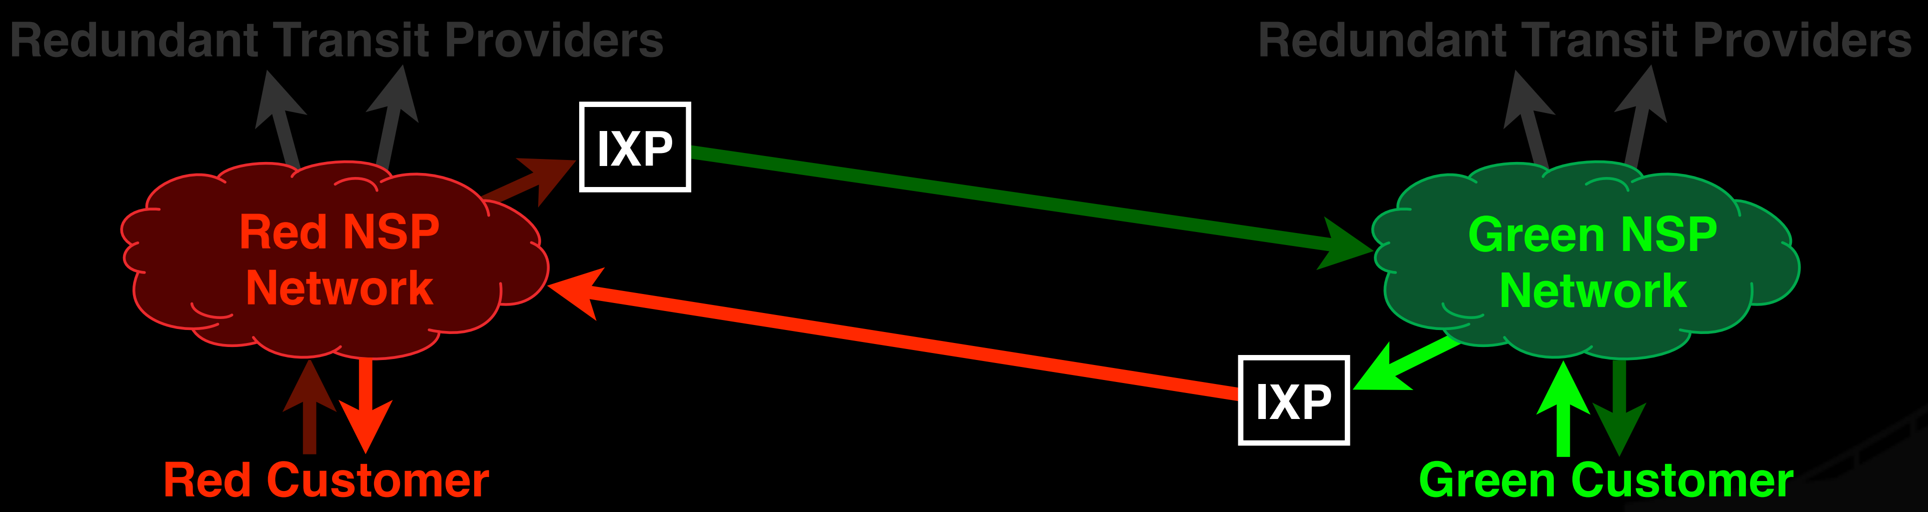

Hot Potato Routing

Red Network is responsible for its own costs

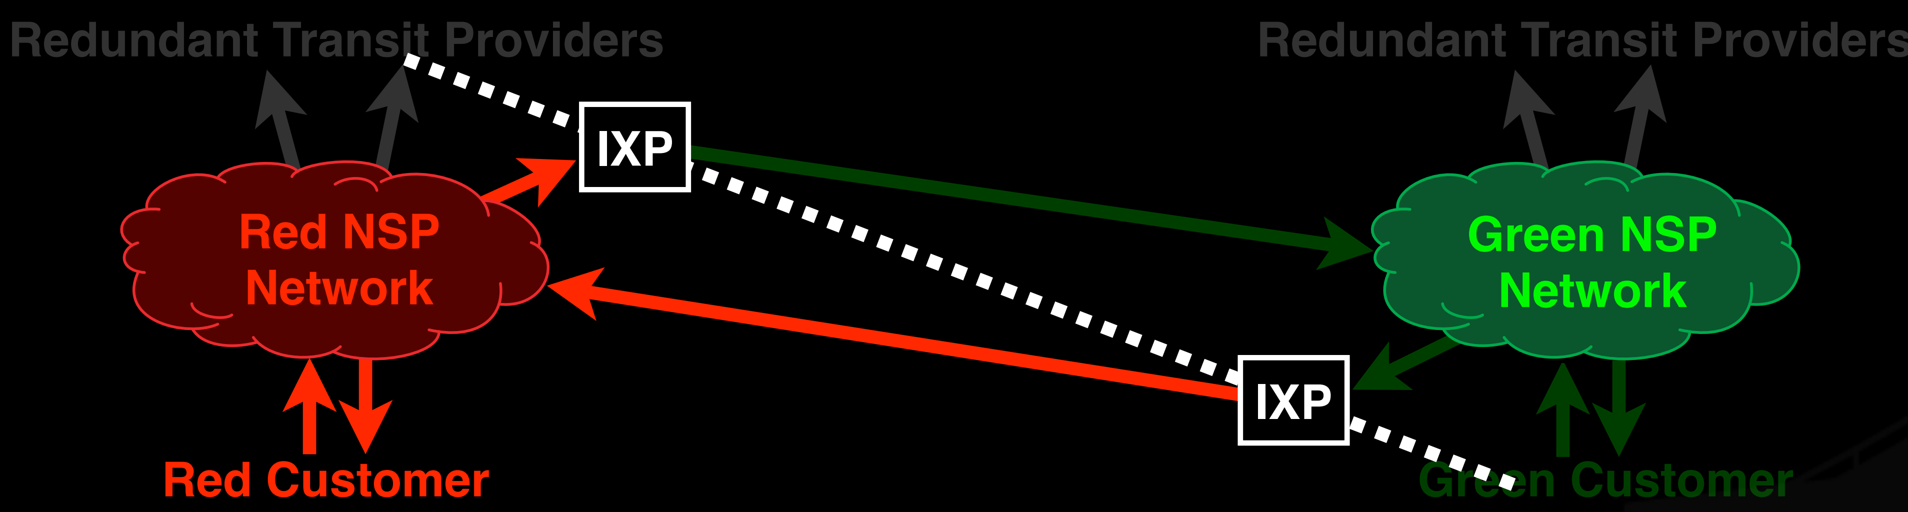

Hot Potato Routing

Green Network is responsible for its own costs

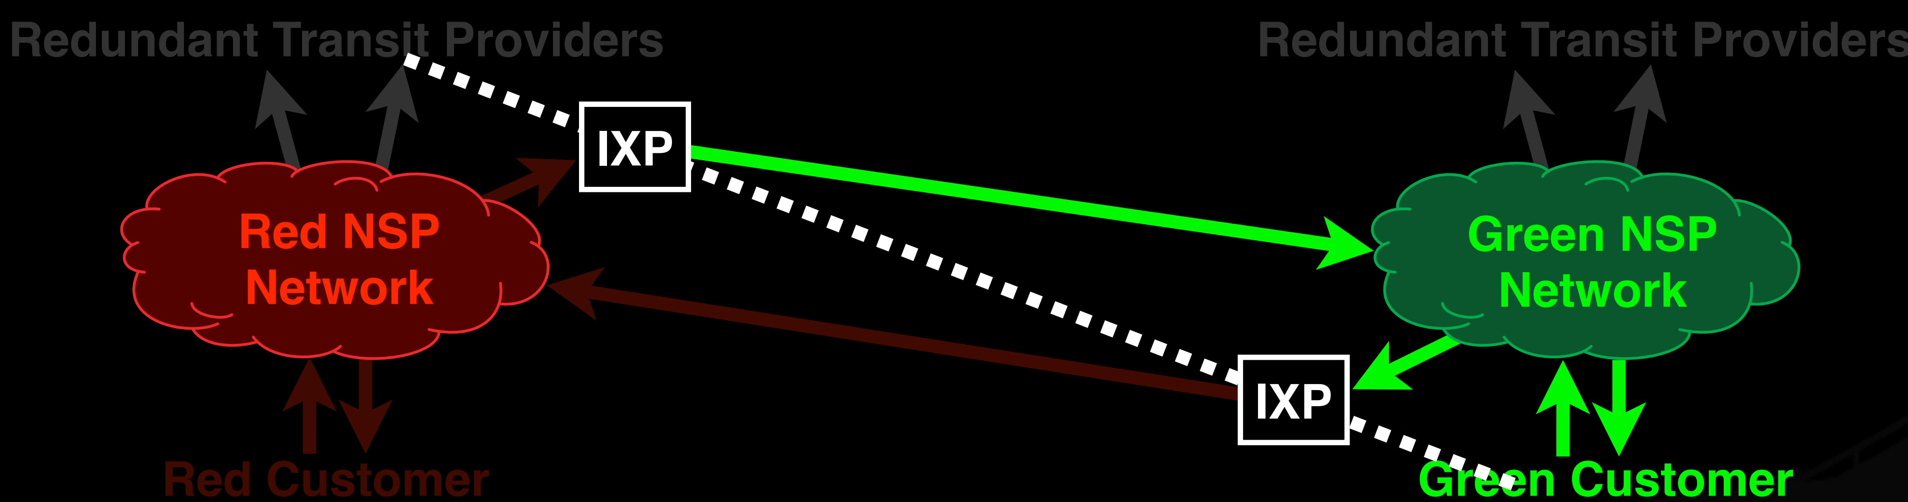

Hot Potato Routing

Symmetry: Fair sharing of costs

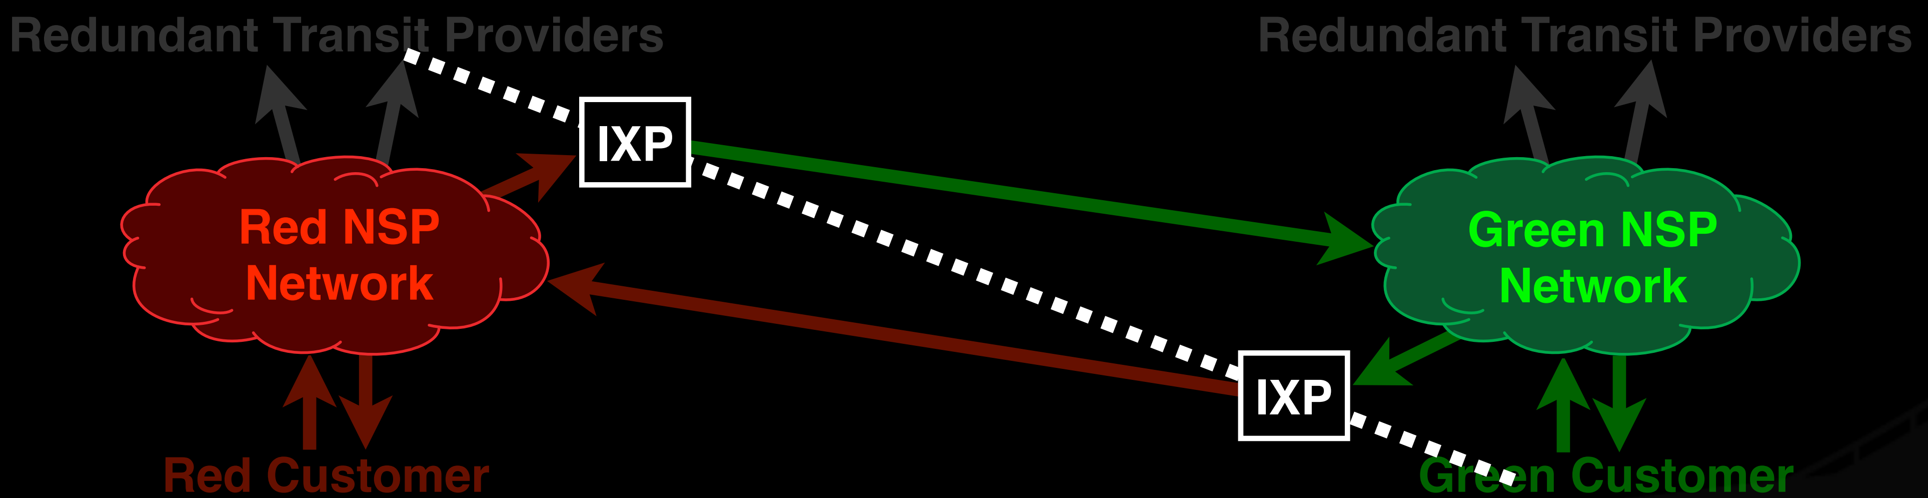

The efficiency of the Internet depends upon this principle:

For any two parties who wish to exchange traffic, there must be a pair of exchanges, one near each party.

The manifestation of this inefficiency:

Countries which haven’t yet built Internet Exchange Points disadvantage themselves, and export capital to countries that already have.

One More Catch

Of multiple possible traffic paths, customers will choose the one with the largest capacity, not the one with the lowest cost or lowest utilization.

One More Catch

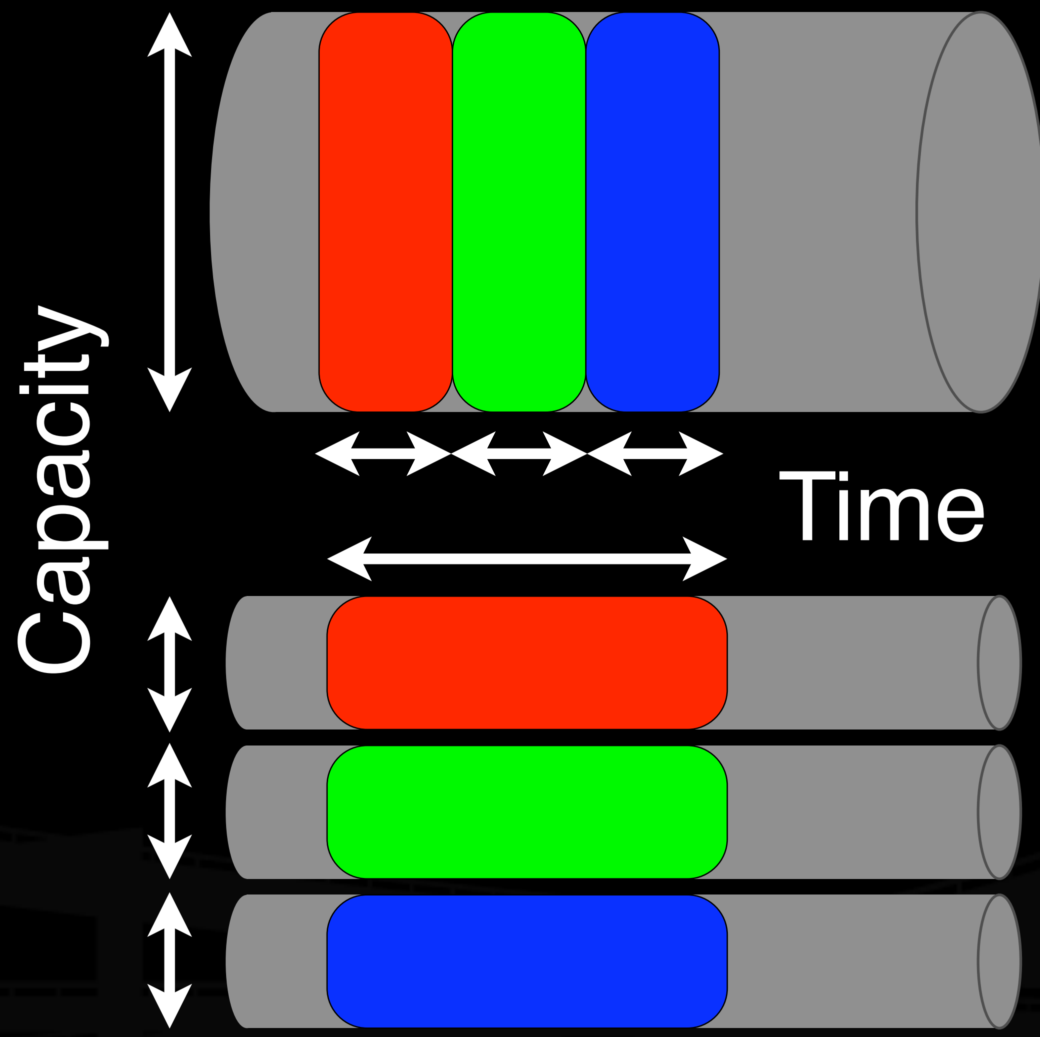

One More Catch

Therefore, every ISP must make sure that their lowest-cost connection is also their highest-capacity connection, even if that means very low initial utilization.

Otherwise, an IXP will fail to grow, and will not produce enough bandwidth to overturn existing market dynamics.

http://pch.net/ixpdir

PCH
Packet Clearing House

Logged in as Bill Woodcock

search pch.net

Home | About | Purpose | Technology | Resources | Calendar | Sponsors | Contact

INTERNET EXCHANGE DIRECTORY

- Hide inactive IXPs
- Show summary views
- Submit a change
- Feature/Bug report

Region	Country	City	Internet Exchange Name	Participants	Traffic	Prefixes	Established	URL
Europe	Germany	Frankfurt	Deutscher Commercial Internet Exchange	235	601G	2455762	May 1995	
	Netherlands	Amsterdam	Amsterdam Internet Exchange	312	592G	5117671	29 Dec 1997	
	United Kingdom	London	London Internet Exchange	316	441G	2833874	8 Nov 1994	
Asia-Pacific	Korea	Seoul	Korea Internet Data Center	200	300G		Oct 1999	
	Japan	Tokyo	JPNAP Tokyo I, Otemachi	58	180G		2001	
		Tokyo	Japan Internet Exchange	111	124G		Jul 1997	
Europe	Spain	Madrid	Espana Internet Exchange	47	112G		1997	
	Hungary	Budapest	Budapest Internet Exchange	47	98.7G		1996	

http://pch.net/ixpdir

PCH
Packet Clearing House

Logged in as Bill Woodcock

search pch.net

Home | About | Purpose | Technology | Resources | Calendar | Sponsors | Contact

INTERNET EXCHANGE DIRECTORY

- Hide Inactive IXPs
- Show summary views
- Submit a change
- Feature/Bug report

Region	Country	City	Internet Exchange Name	Participants	Traffic	Prefixes	Established	URL
Europe	Germany	Frankfurt	Deutscher Commercial Internet Exchange	235	601G	2455762	May 1995	
	Netherlands	Amsterdam	Amsterdam Internet Exchan	312	592G	5117671	29 Dec 1997	
	United Kingdom	London	London Internet Exchange	316	441G	2833874	8 Nov 1994	
Asia-Pacific	Korea	Seoul	Korea Internet Data Center	200	300G		Oct 1999	
	Japan	Tokyo	JPNAP Tokyo I, Otemachi	58	180G		2001	
		Tokyo	Japan Internet Exchange	111	124G		Jul 1997	
Europe	Spain	Madrid	Espana Internet Exchange	47	112G		1997	
	Hungary	Budapest	Budapest Internet Exchang	47	98.7G		1996	

(participants) 235 ISPs participate at Deutscher Commercial Internet Exchange. By participants, it ranks #1 in Germany, #4 in Europe, and #4 in the world.

(traffic) Deutscher Commercial Internet Exchange peaks at 601G of traffic each day. By traffic, it ranks #1 in Germany, #1 in Europe, and #1 in the world.

http://pch.net/ixpdir

Google

PCH
Packet Clearing House

Logged in as Bill Woodcock

search pch.net

Home | About | Purpose | Technology | Resources | Calendar | Sponsors | Contact

INTERNET EXCHANGE DIRECTORY

- Hide inactive IXPs
- Show summary views
- Submit a change
- Feature/Bug report

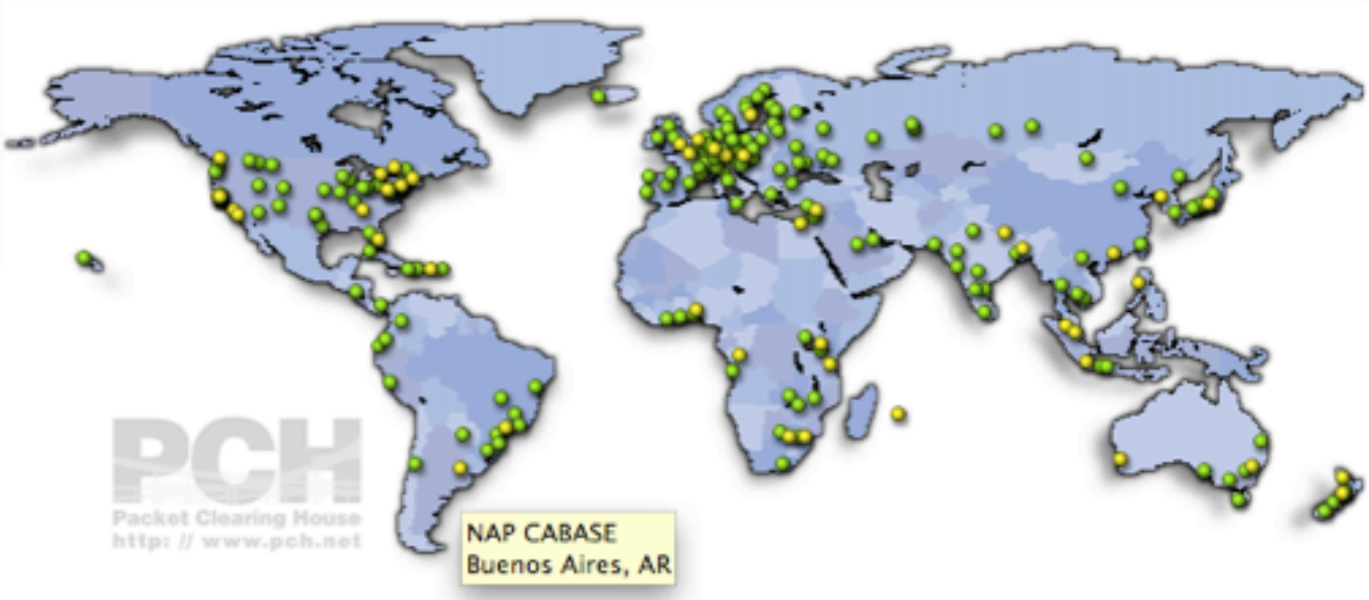

Region	Country	City	Internet Exchange Name	Participants	Traffic	Prefixes	Established	URL
Europe	Germany	Frankfurt	Deutscher Commercial Internet Exchange	235	601G	2455762	May 1995	
	Netherlands	Amsterdam	Amsterdam Internet Exchange	312	592G	5117671	29 Dec 1997	
	United Kingdom	London	London Internet Exchange	316	441G	2833874	8 Nov 1994	
Asia-Pacific	Korea	Seoul	Korea Internet Data Center	200	300G		Oct 1999	
	Japan	Tokyo	JPNAP Tokyo I, Otemachi	58	180G		2001	
		Tokyo	Japan Internet Exchange	111	124G		Jul 1997	
Europe	Spain	Madrid	Espana Internet Exchange	47	112G		1997	
	Hungary	Budapest	Budapest Internet Exchange	47	98.7G		1996	

http://pch.net/annual-growth

Google

Internet Exchange Point Growth

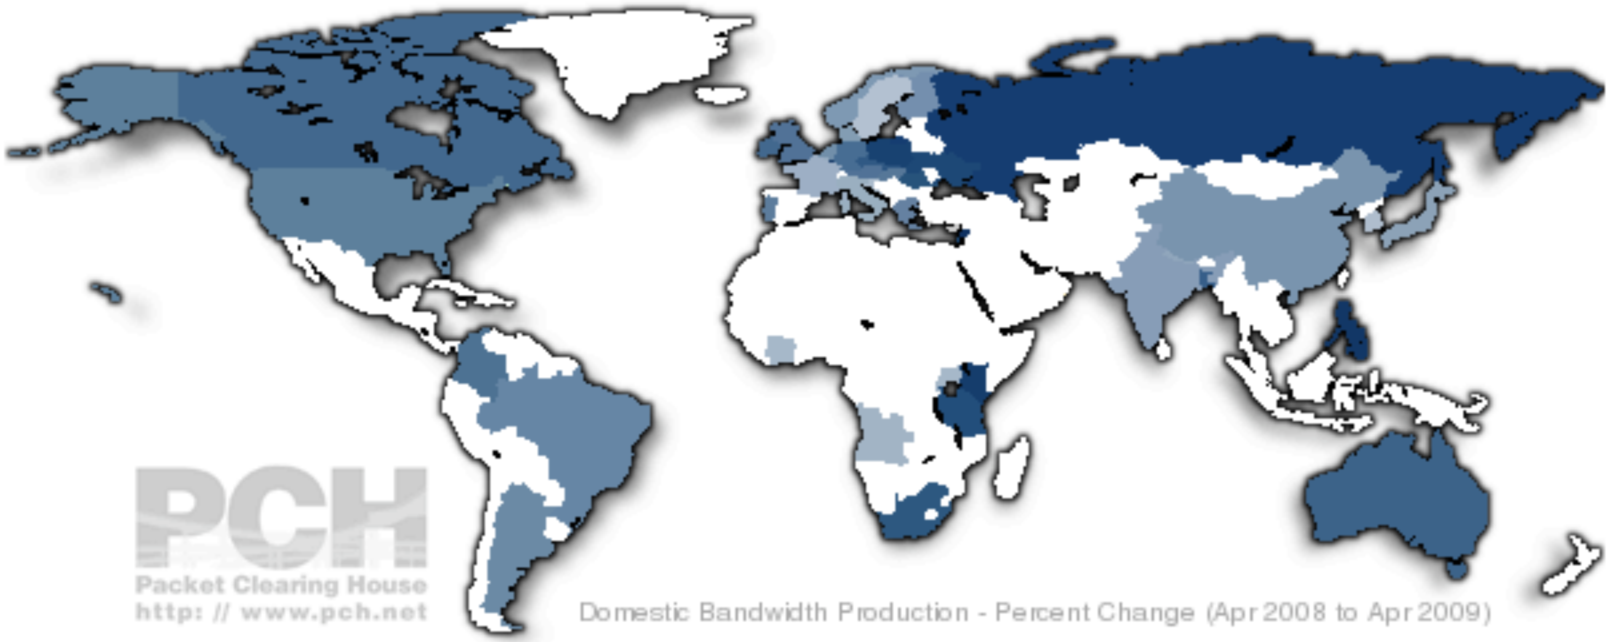

Domestic Bandwidth Production - Percent Change (Apr 2008 to Apr 2009)

	Internet Exchange Points				Domestic Bandwidth Production			
Country	Apr 2008	Apr 2009	Net Change	Percent Change	Apr 2008	Apr 2009	Net Change	Percent Change
Lebanon	1	1			4.73M	17.6M	+12.9M	+272%
Luxembourg	1	1			43.7M	155M	+111M	+254%
Philippines	2	2			4.05M	13.2M	+9.15M	+226%
Kenya	1	1			14M	43.6M	+29.6M	+211%
Russia	9	10	+1	+11%	27.3G	82.6G	+55.3G	+203%
Poland	5	5			24.4G	65.3G	+40.8G	+167%
Israel	1	1			795M	2.08G	+1.29G	+162%
Ukraine	5	5			16.7G	40.9G	+24.2G	+145%
Tanzania	2	2			2.07M	5.01M	+2.94M	+142%
Romania	2	2			9.41G	21.9G	+12.5G	+133%

http://pch.net/annual-growth

Internet Exchange Point Growth

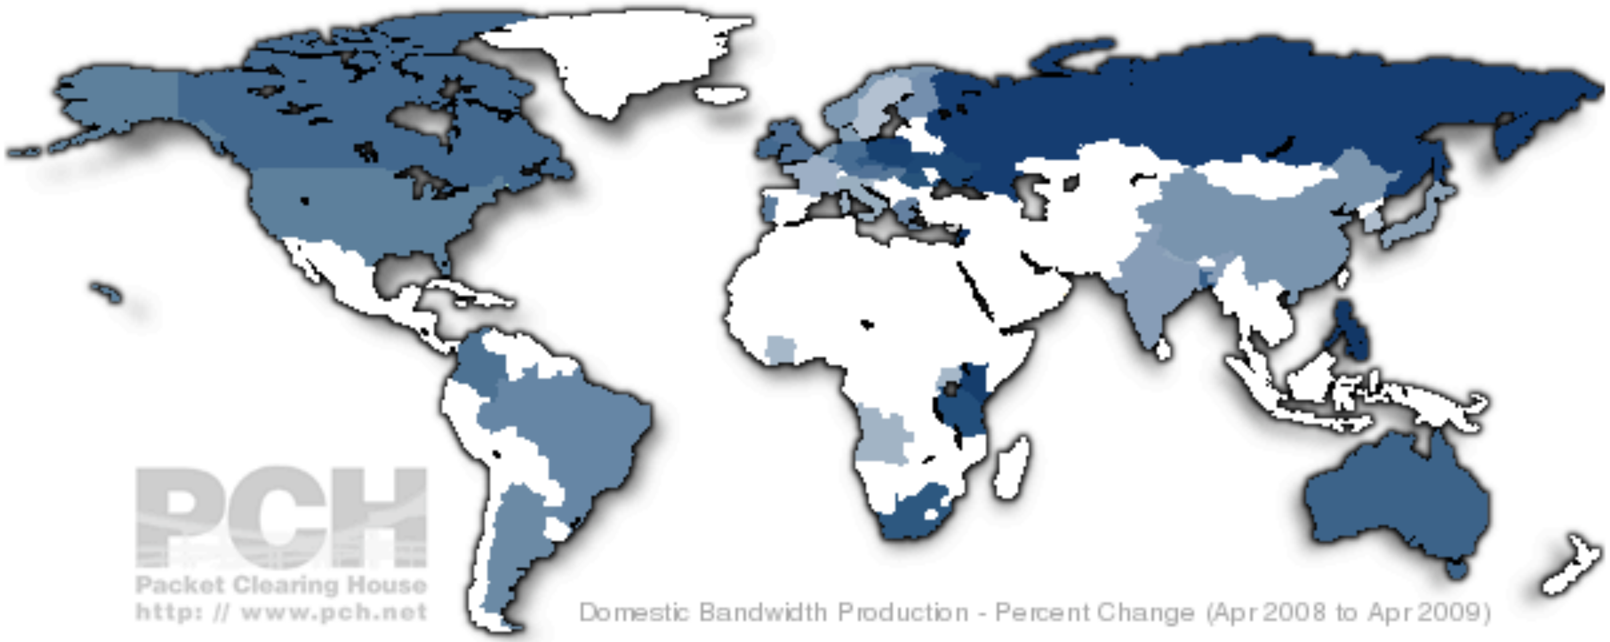

Region	Internet Exchange Points				Domestic Bandwidth Production			
	Apr 2008	Apr 2009	Net Change	Percent Change	Apr 2008	Apr 2009	Net Change	Percent Change
Africa	20	21	+1	+5%	211M	450M	+239M	+113%
Asia-Pacific	69	73	+4	+6%	922G	1.05T	+128G	+14%
Europe	119	123	+4	+3%	1.88T	2.75T	+873G	+46%
Latin America	23	24	+1	+4%	18.3G	25.7G	+7.48G	+41%
North America	87	88	+1	+1%	229G	332G	+103G	+45%
Total	**318**	**329**	**+11**	**+3%**	**3.05T**	**4.16T**	**+1.11T**	**+27%**

IXP growth by country | IXP growth by region | Auto reporting | IXP with IPv6 subnets | Root server locations

A Closer Look at The Economics of Internet Exchange Points

Packet Clearing House

Tools for thinking about Internet Exchanges in economic terms

What are we, as ISPs, selling?

- The right to modulate bits.
- That right is a perishable commodity.
- Where do we get the potentially-modulatable bits?

The right to modulate bits

Any Internet connection is a serial stream of time-slices.

Each time-slice can be modulated with a binary one or zero, one bit.

Each customer purchases potentially-modulatable bits at some *rate*, for example, 2mbps, which is 5.27 trillion bits per monthly billing cycle.

Packet Clearing House

That’s a perishable commodity

The quality (as opposed to quantity-per-time) characteristics of an Internet connection are *loss*, *latency*, *jitter*, and *out-of-order delivery*.

Loss increases as a function of the number and reliability of components in the path, and the amount of contention for capacity.

Latency increases as a function of distance, and degree of utilization of transmission buffers by competing traffic sources.

Jitter is the degree of variability in loss and latency, which negatively affects the efficacy and efficiency of the encoding schemes which mitigate their effects. Jitter increases relative to the ratio of traffic burstiness to number of sources.

Out-of-order delivery is the portion of packets which arrive later than other, subsequently-transmitted packets. It increases as a function of the difference in queueing delay on parallel paths.

All of **these properties become worse with time and distance**, which is a reasonable definition of a perishable commodity.

So where do we get the bits?

The value of the Internet is communication.

The value is produced at the point at which communication occurs between two ISPs, and it is transported to the customers who utilize it.

Thus, all the bits we sell come from an Internet exchange, whether nearby, or far away.

Packet Clearing House

An analogy

Let’s look at another perishable commodity with more readily observed economic properties... **Bananas**.

Packet Clearing House

Value decreases with time & distance

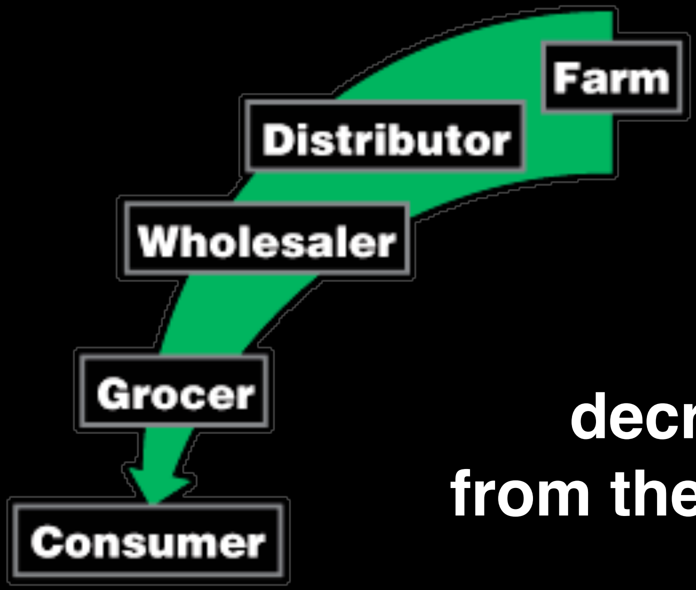

The value of a banana decreases, the further it gets from the farm which produced it.

The shelf-life which the consumer can expect decreases, and eventually it becomes overripe, then rotten.

Packet Clearing House

Cost increases with time & distance

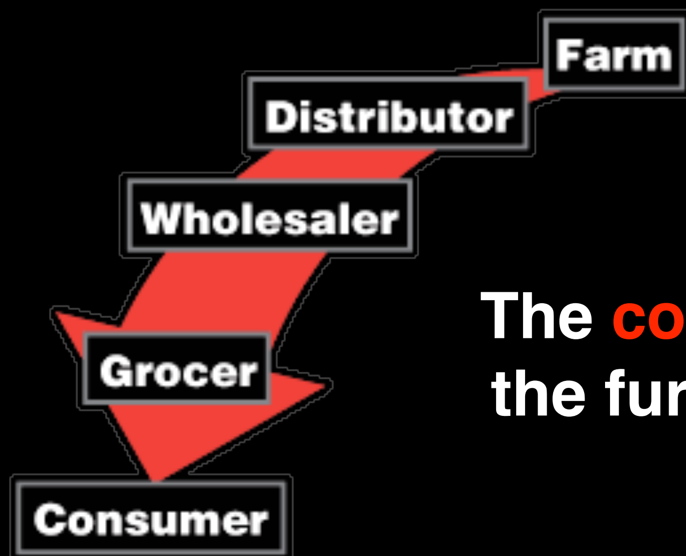

The cost of a banana increases, the further it gets from the farm which produced it.

Salaries and hourly labor, warehouse leasing, diesel fuel, truck amortization, loss and spoilage, insurance, and other factors contribute additively.

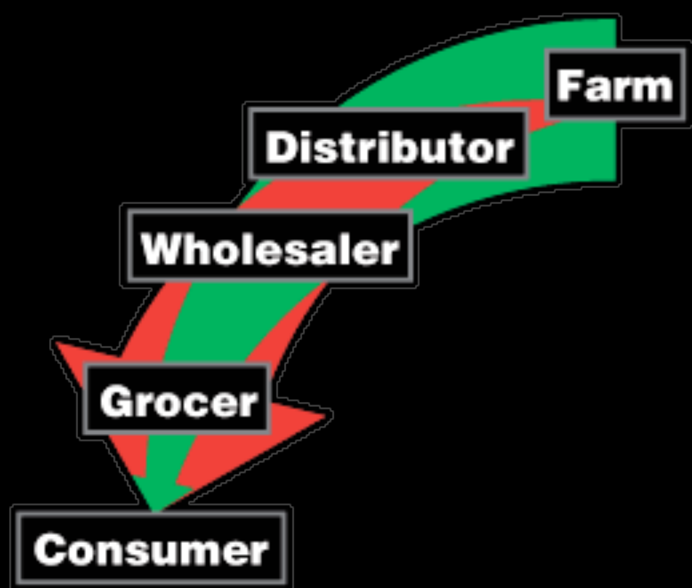
Farm
Distributor
Wholesaler
Grocer
Consumer

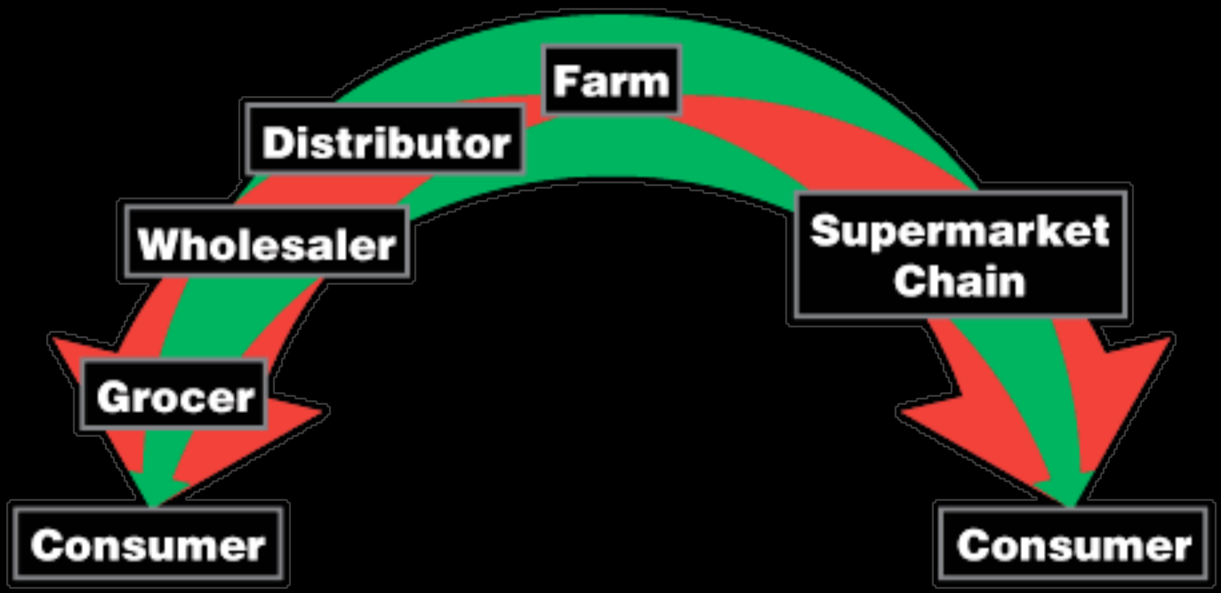

In a competitive environment, retail price is limited by competition, so time and distance influence the price more than the number of middlemen.

The problem is the same:

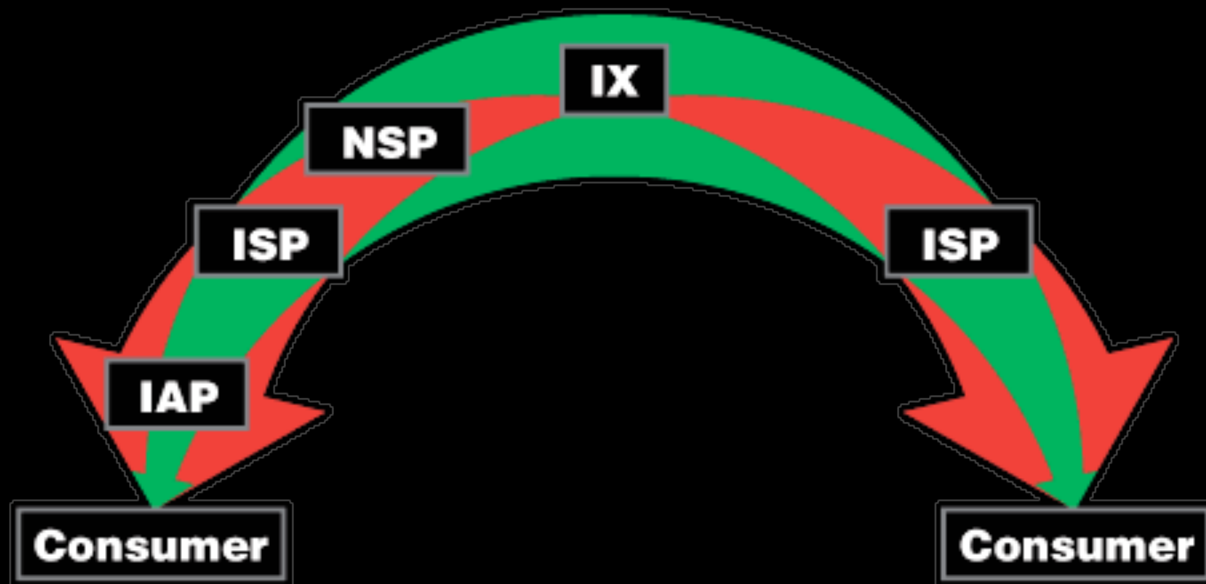

ISPs form a delivery chain, bringing perishable bits to the consumers who purchase them.

Packet Clearing House

So how do we improve things?

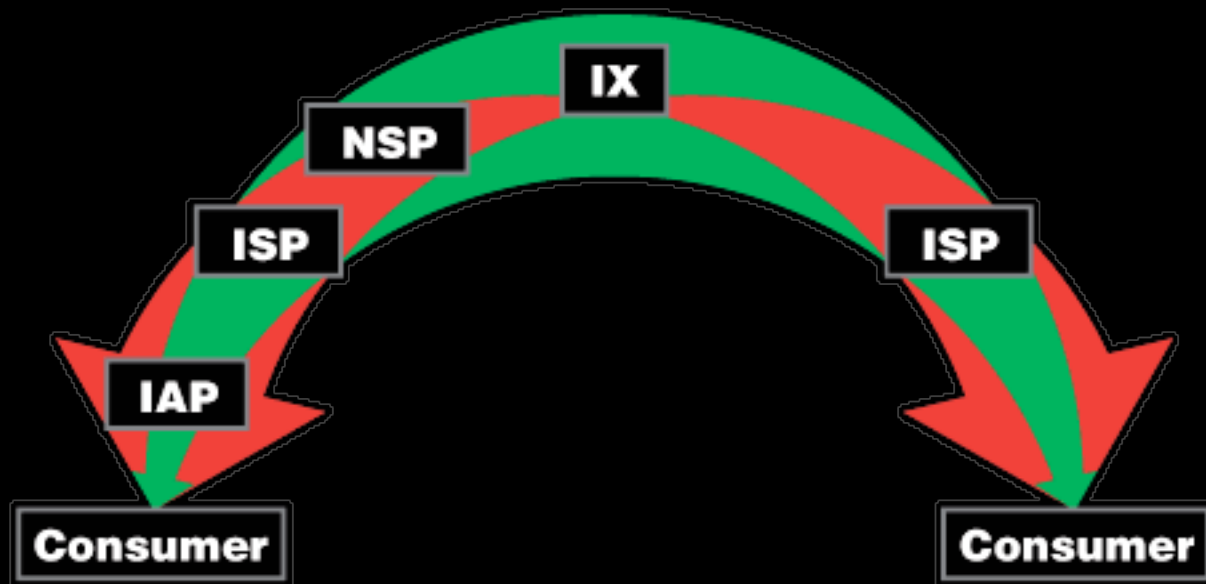

Bring the customer nearer an IX...

High cost
Low value

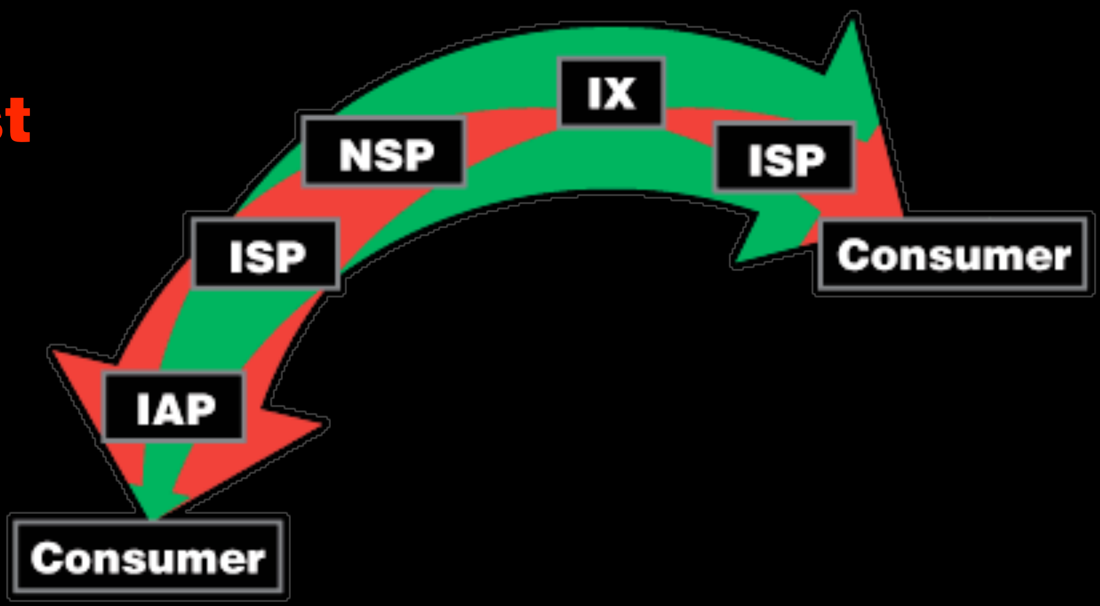

Low cost
High value

...or bring an IX nearer the customer.

Packet Clearing House

So how do we recognize a successful exchange?

The purpose of an IX is to lower participating ISPs' average per bit delivery costs (APBDC).

A cheap IX is probably a successful one.
An expensive IX is always a failure.
Reliability is just hand-waving by salespeople.

Packet Clearing House

How to Build an Internet Exchange Point

Packet Clearing House

Determining Need

Sufficient end-user base?

No existing facility to build upon?

Sufficient degree of locally-destined traffic?

Geographic Location

User population

Fiber facilities or rights-of-way

Founding participants

Density

Centralized in one room

Campus of adjacent buildings

MAN

Frame / ATM / SMDS / MPLS cloud

Building Management

Telco hotel

University computing or telecommunications facility

City emergency services facility

In-Building Facilities

Pathways
Power
Cooling
Access and security

Services

Switch fabric
Crossconnects
Route-server
Remote hands
NTP
Web caching

Business Structure

Incorporated or unincorporated?

Staffed or volunteer?

Non-profit or for-profit?

Cooperative or external ownership?

Cost-recovery (predictive or actuals), ad-hoc, or market pricing?

Policies

BLP, MLPA or MMPLA?

Mandatory looking-glass?

Routing and switch-port information public or members-only?

Disclosure in the event of security problems, failures, or mistakes?

Extensible switch fabric?

Thanks, and Questions?

Copies of this presentation can be found
in Keynote, PDF, QuickTime and PowerPoint formats at:

http:// www.pch.net / resources / tutorials / ix-construction

Bill Woodcock
Research Director
Packet Clearing House
woody@pch.net