

Introduction to Peering & Internet Exchange Points

LACNIC Caribbean Regional Meeting

Monday, July 21, 2008


Bill Woodcock

Research Director

Packet Clearing House


ISP Lifecycle: Simple Aggregator

Single Transit Provider ——— IXPs



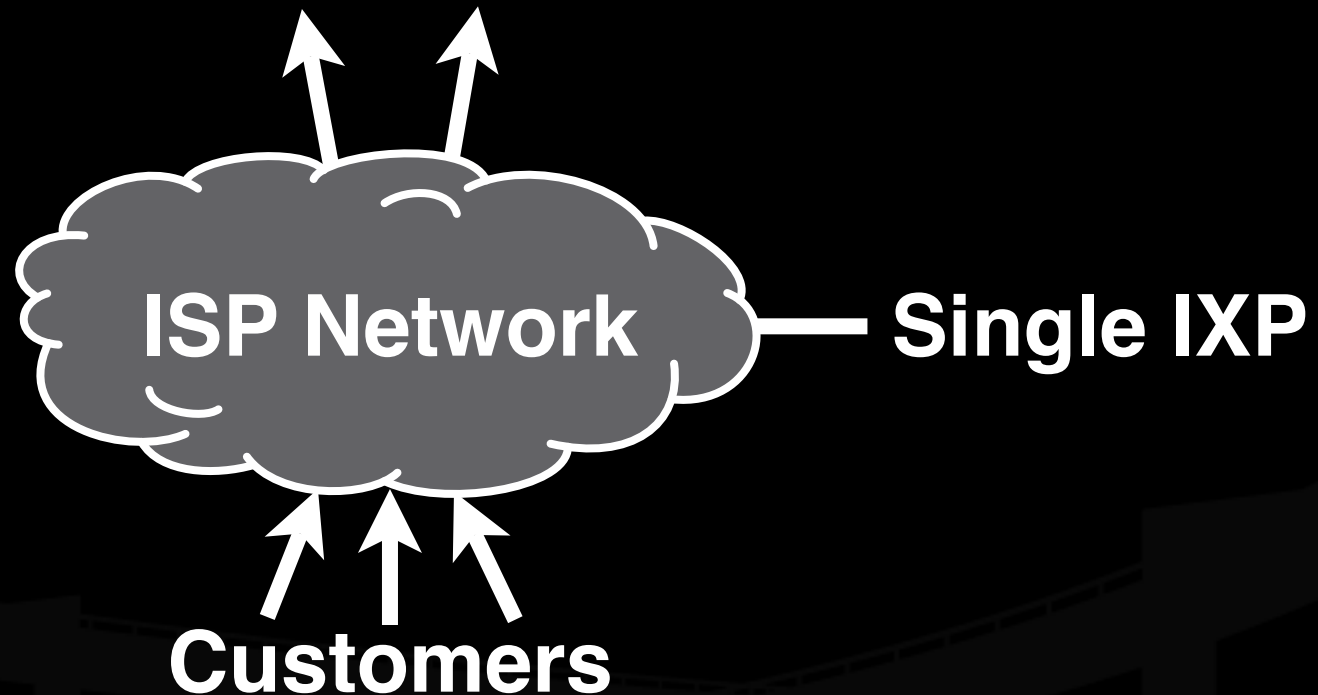
ISP Lifecycle: Redundancy and LCR

Redundant Transit Providers — IXPs




ISP Lifecycle: Local Peer

Redundant Transit Providers — IXPs




ISP Lifecycle: Network Service Provider

Redundant Transit Providers — IXP




Hot Potato Routing




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




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.

A Closer Look at The Economics of Internet Exchange Points

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.

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.

An analogy

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

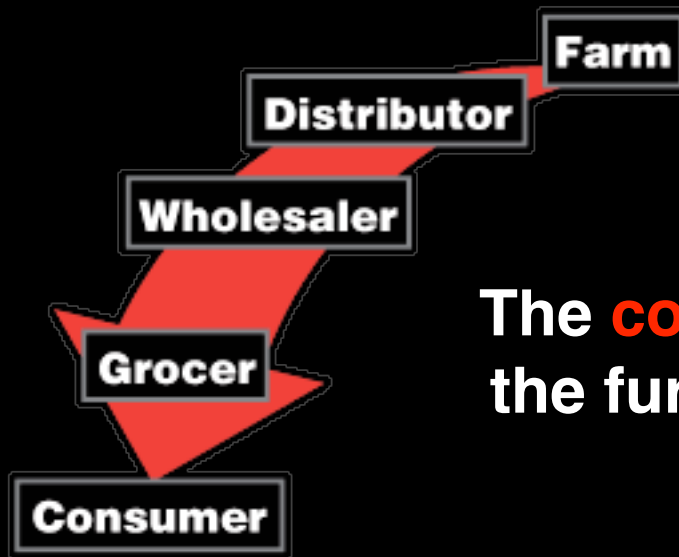
Value decreases with time & distance



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.

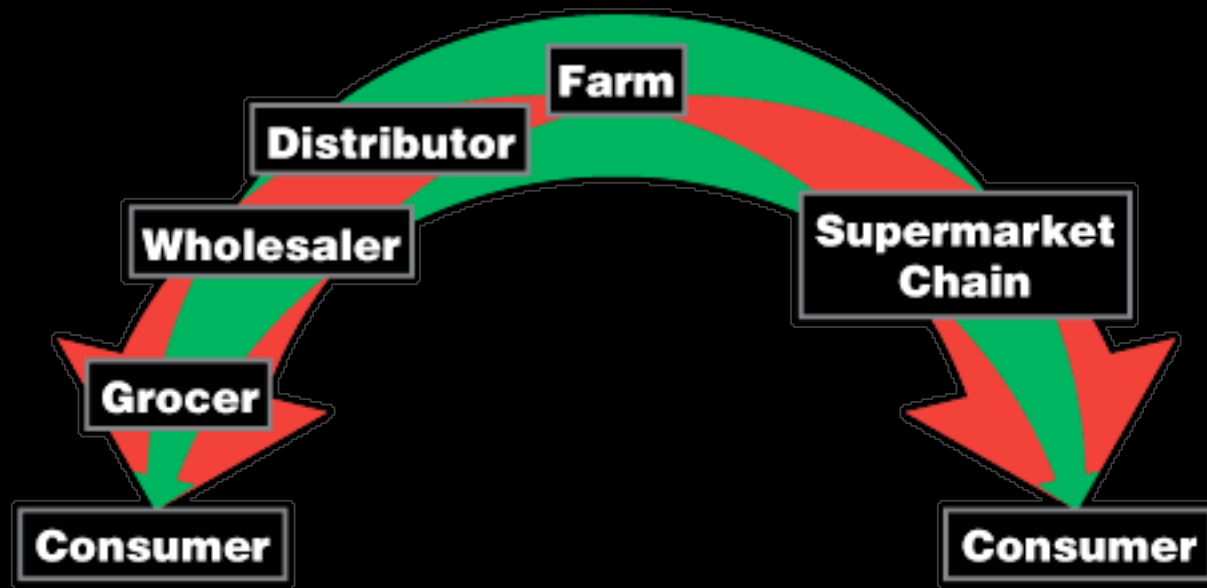
Cost increases with time & distance



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.






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.

So how do we improve things?



Bring the customer nearer an IX...



...or bring an IX nearer the customer.

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.

The irony inherent in that

An efficient IX is an ISPs lowest-cost delivery method.

In order to shift latency-sensitive traffic toward the lowest-cost delivery method, it must also be the highest-capacity pipe.

Regardless of degree of utilization.

Thus many IX connections run at low utilization: apparent inefficiency.

How to Build an Internet Exchange Point

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

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