Internet Exchange Point Workshop

MENOG-16

March 22, 2016

Nishal Goburdhan

Packet Clearing House

ISP Lifecycle: Simple Aggregator

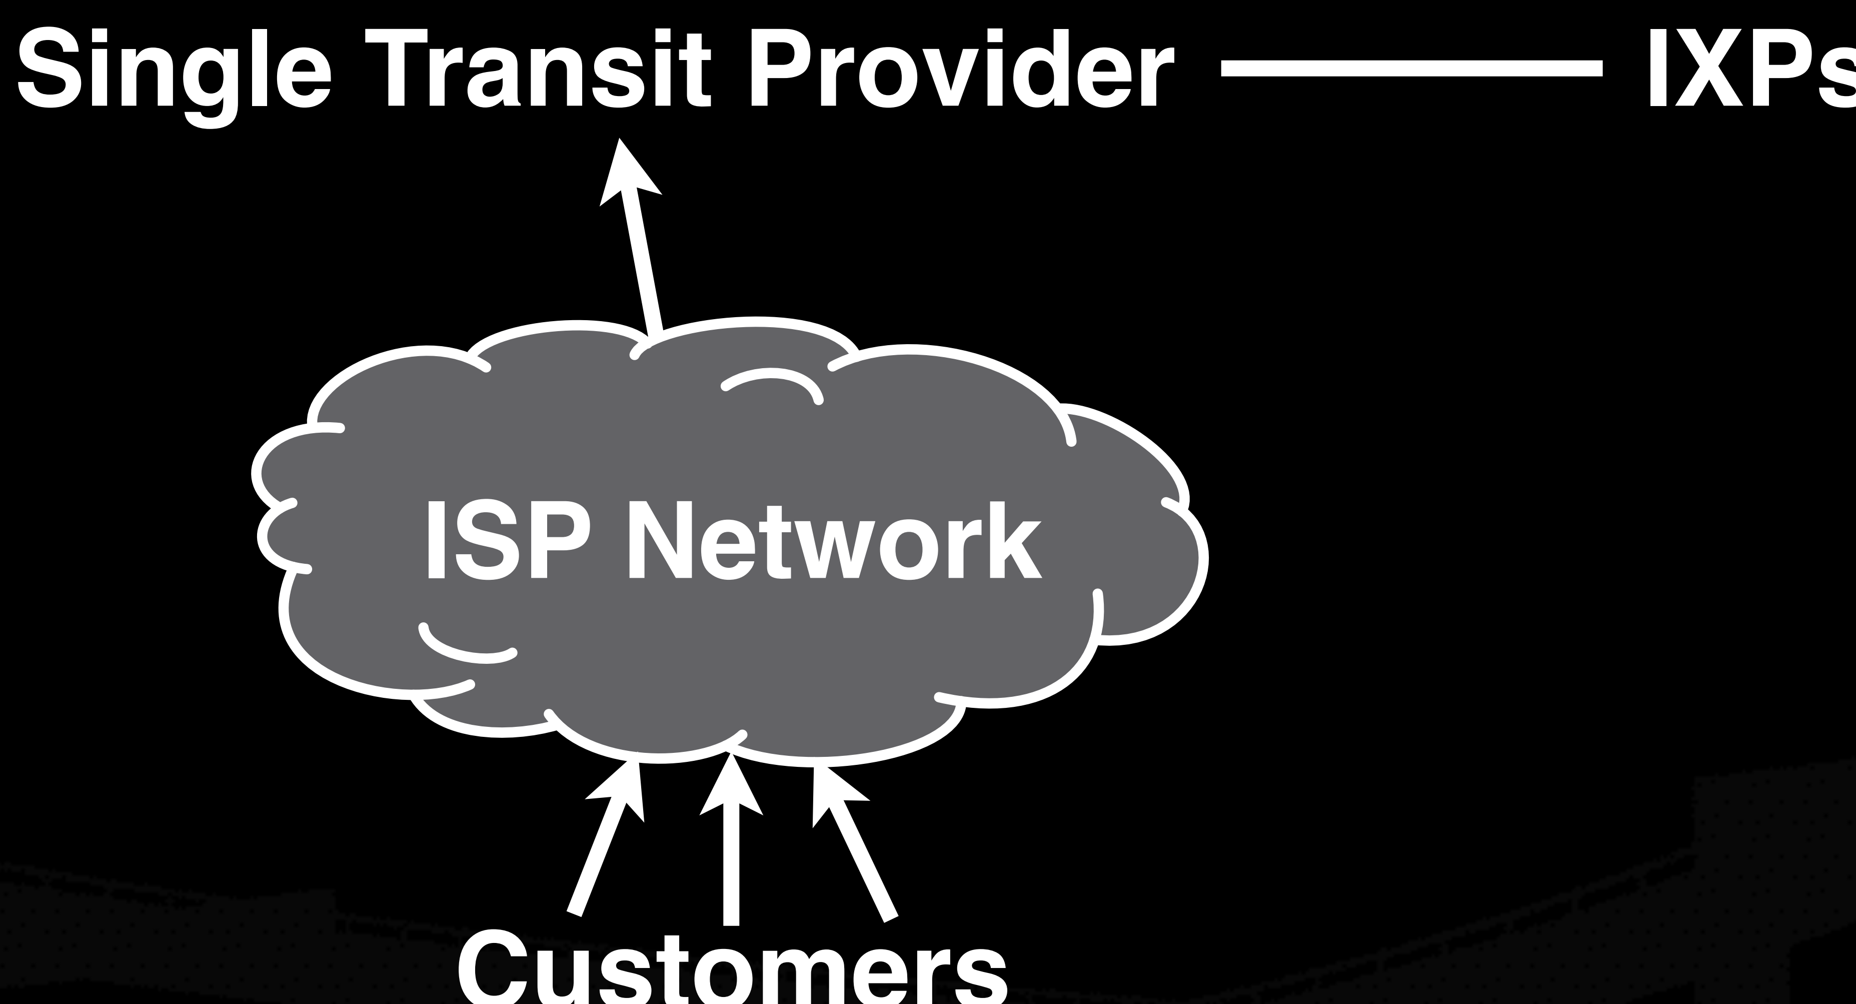

ISP Lifecycle: Redundancy and LCR

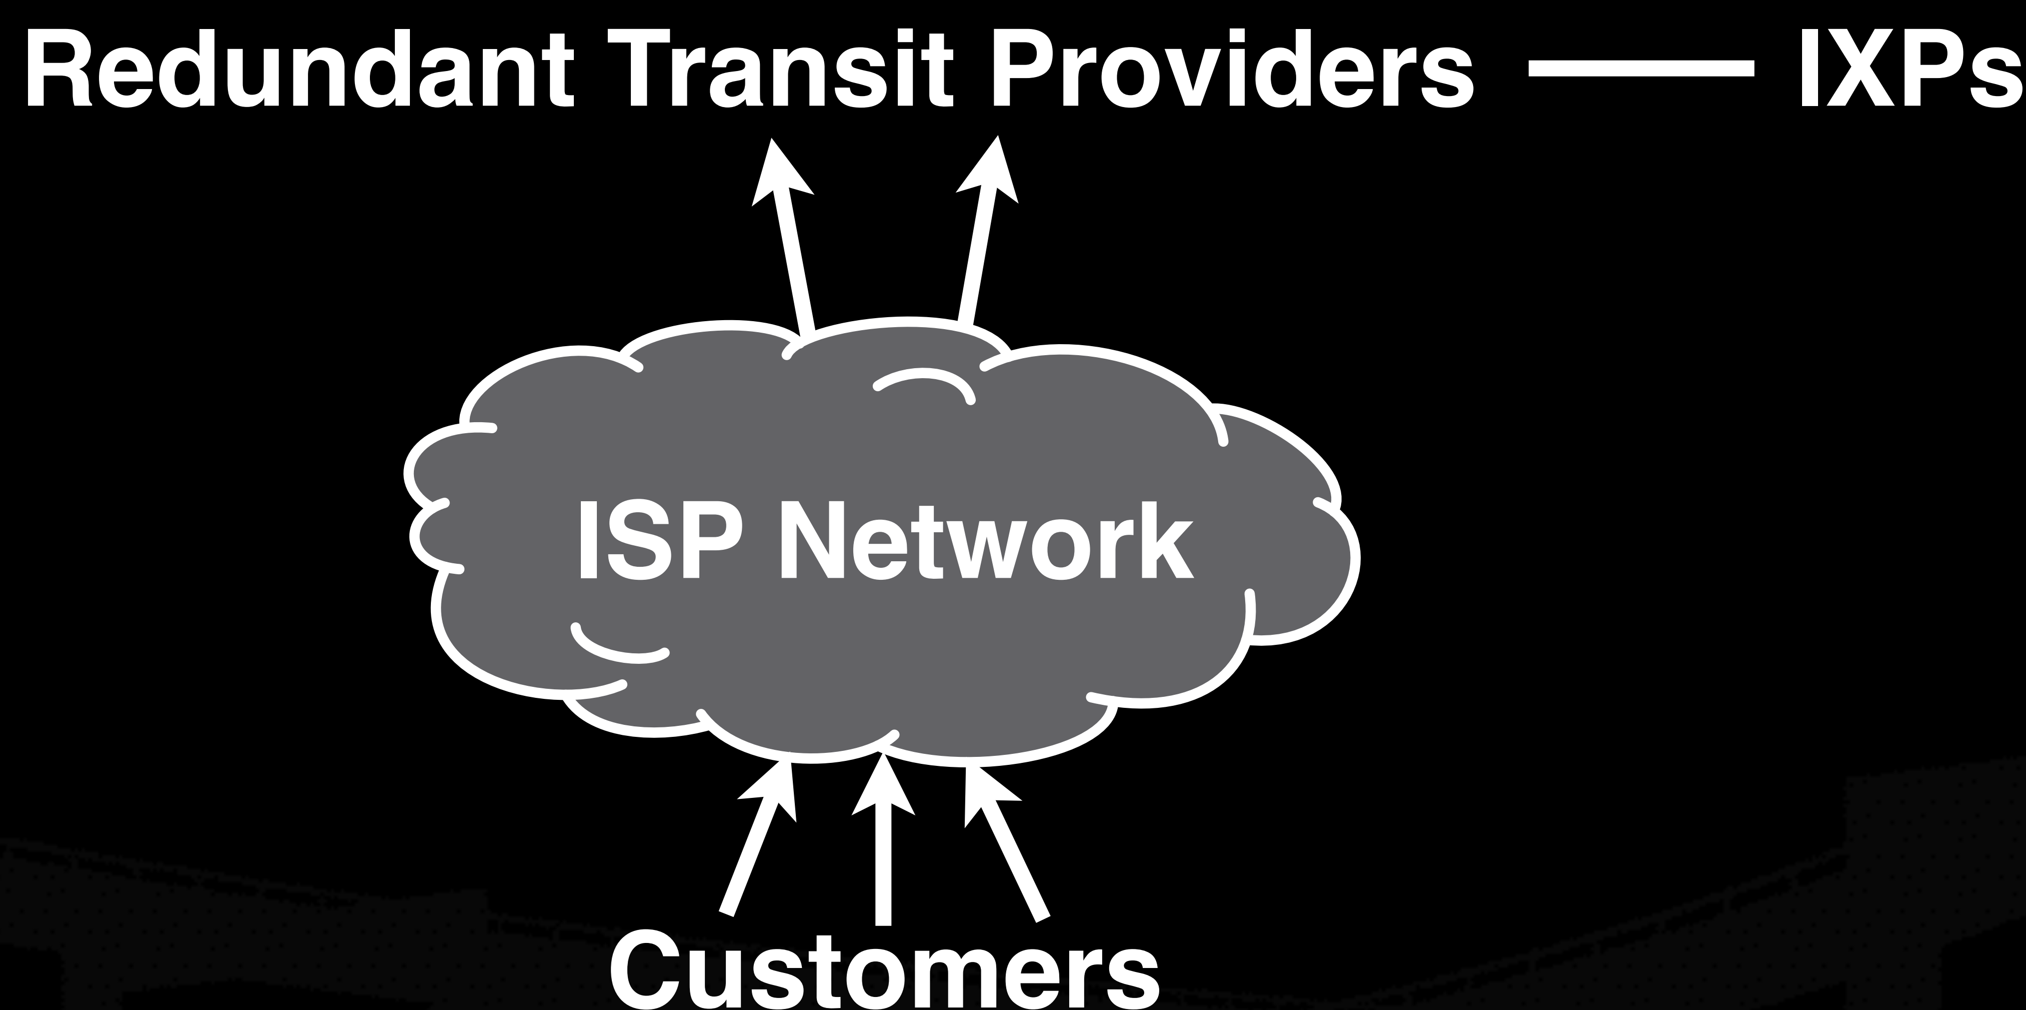

ISP Lifecycle: Local Peer

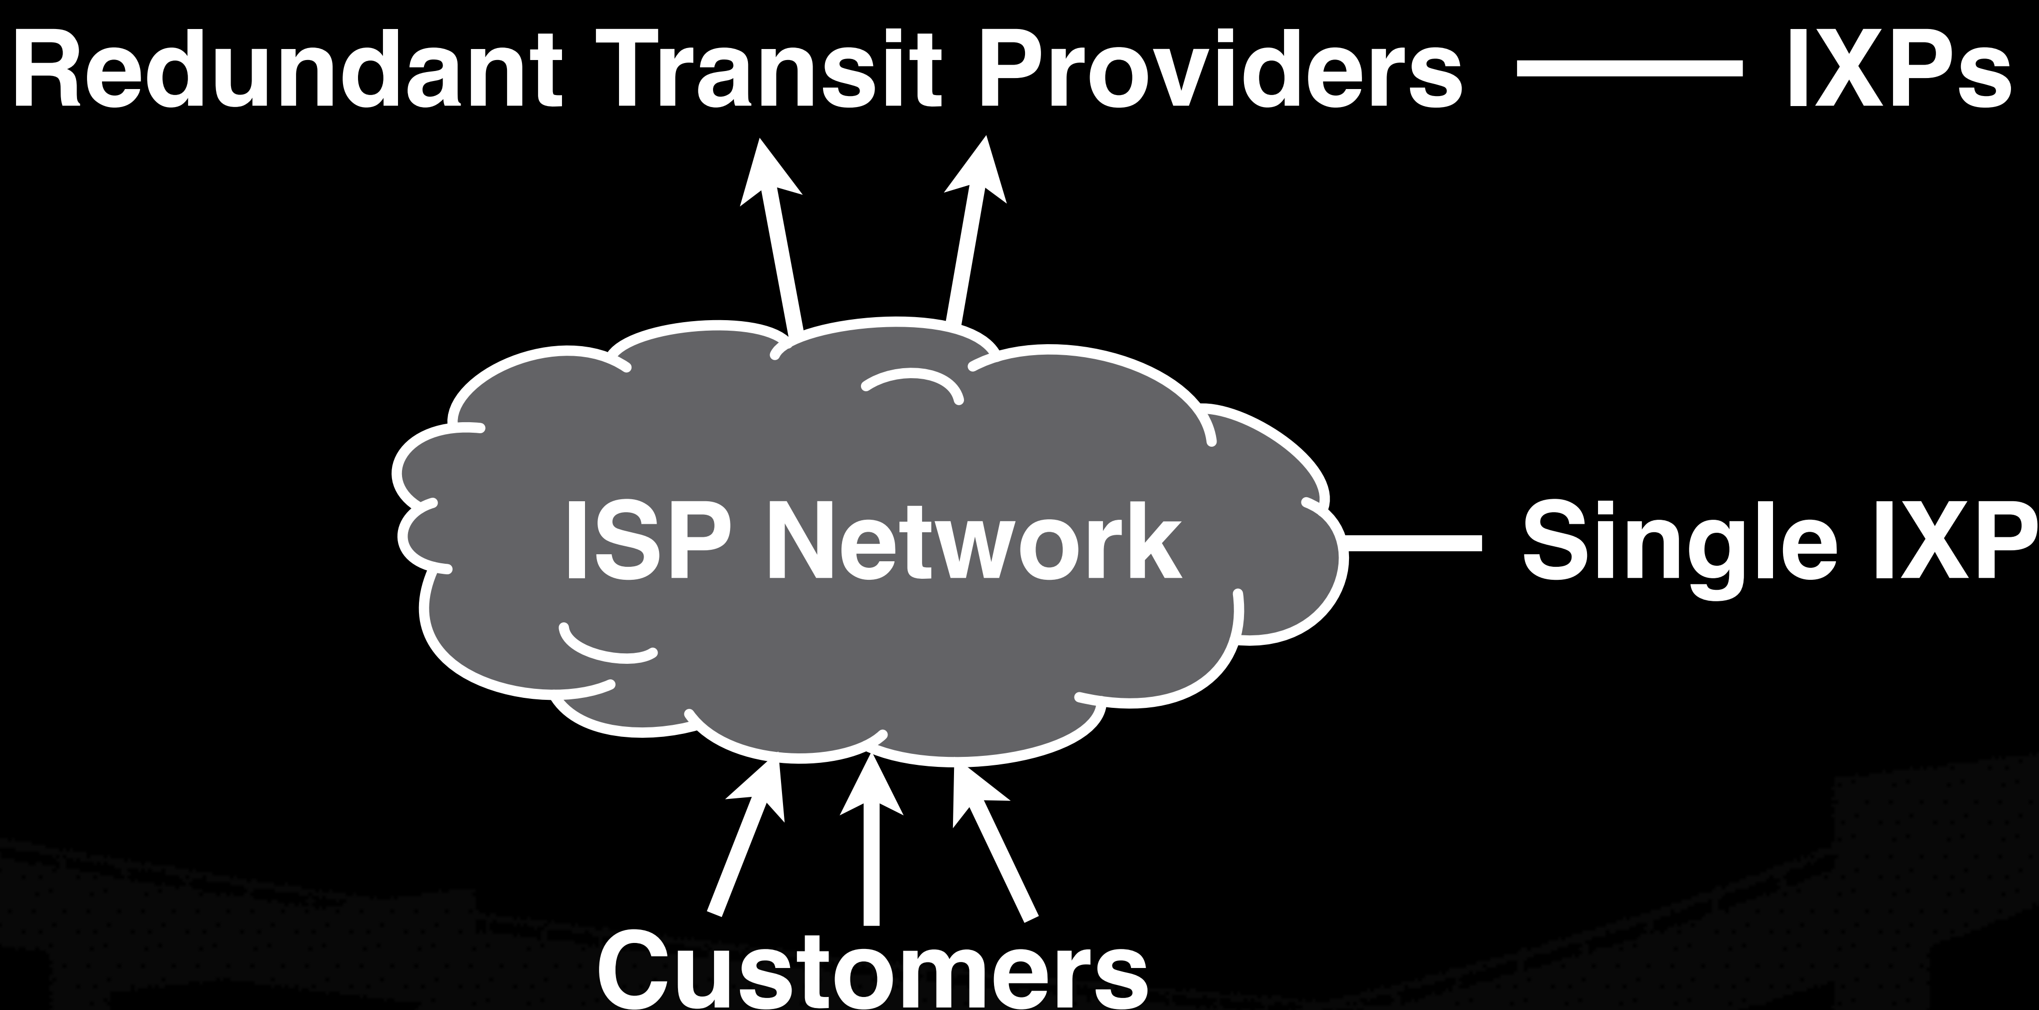

ISP Lifecycle: Network Service Provider

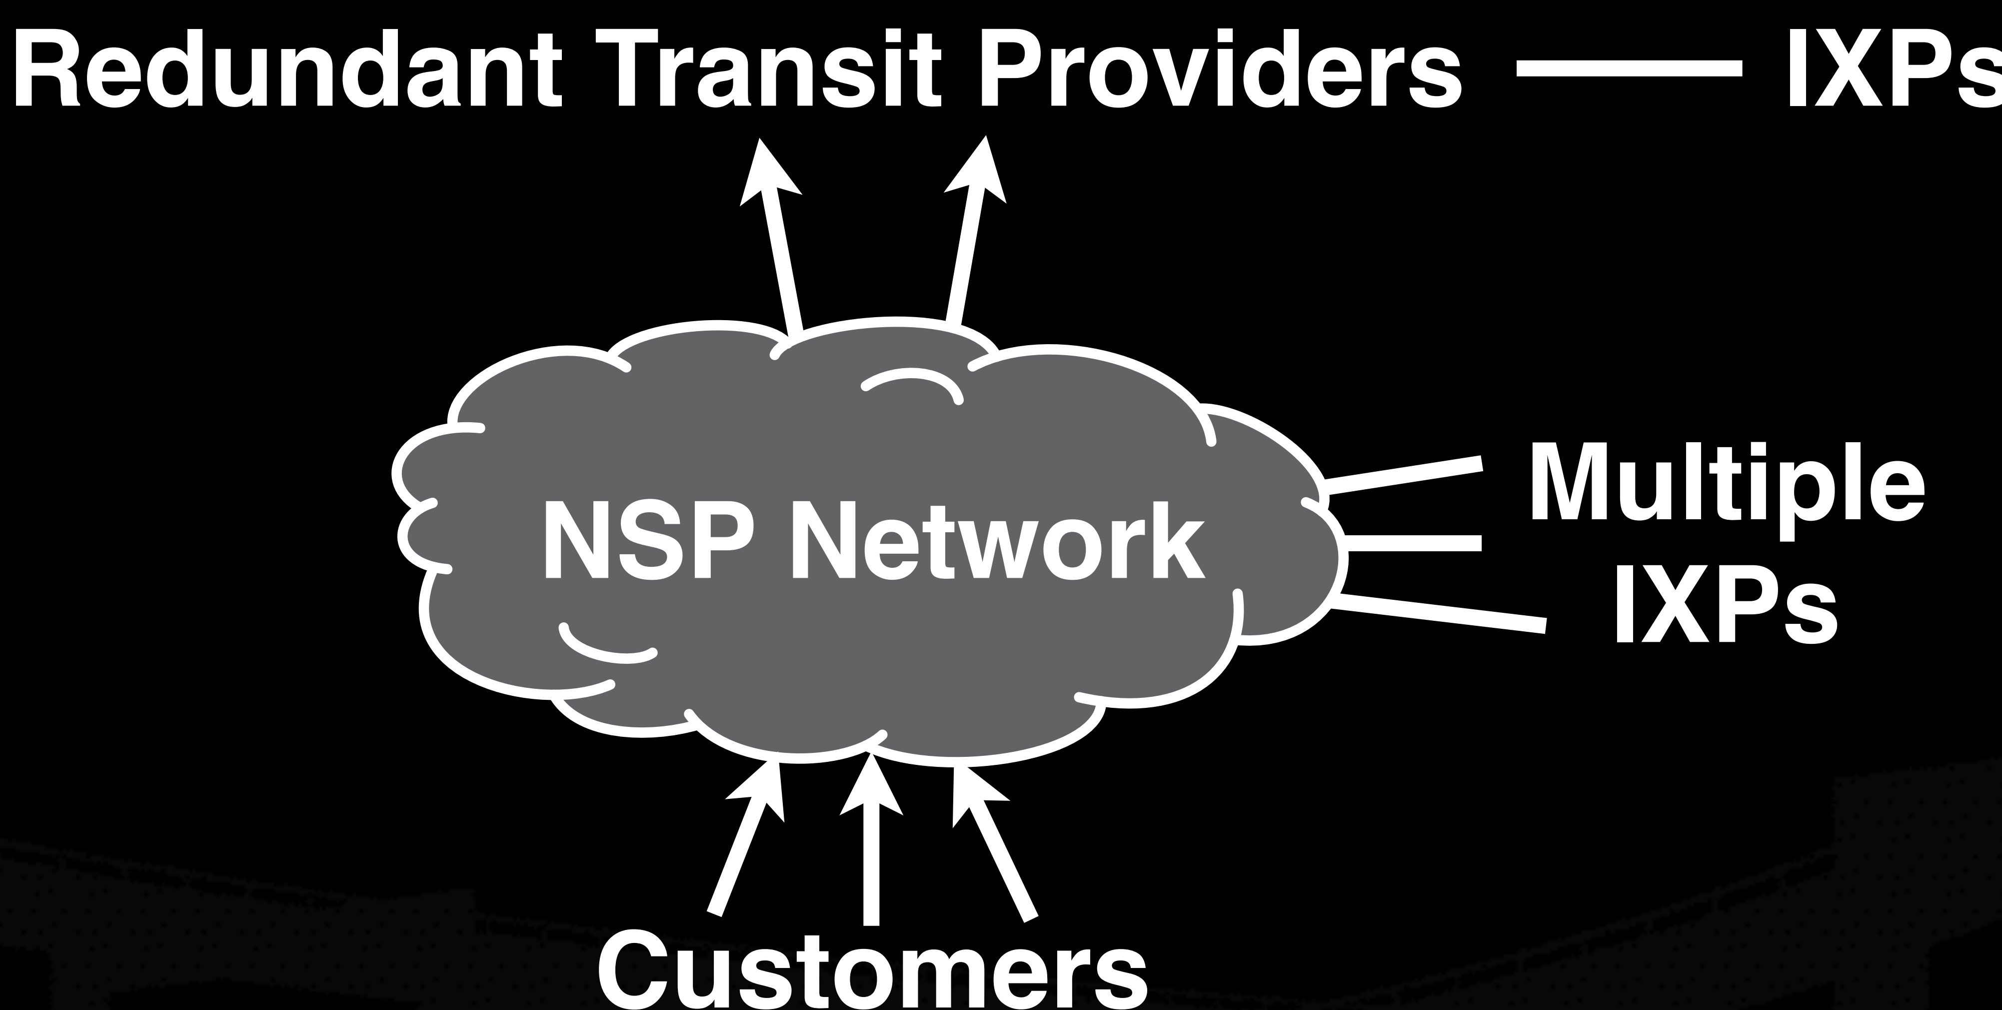

Hot Potato Routing

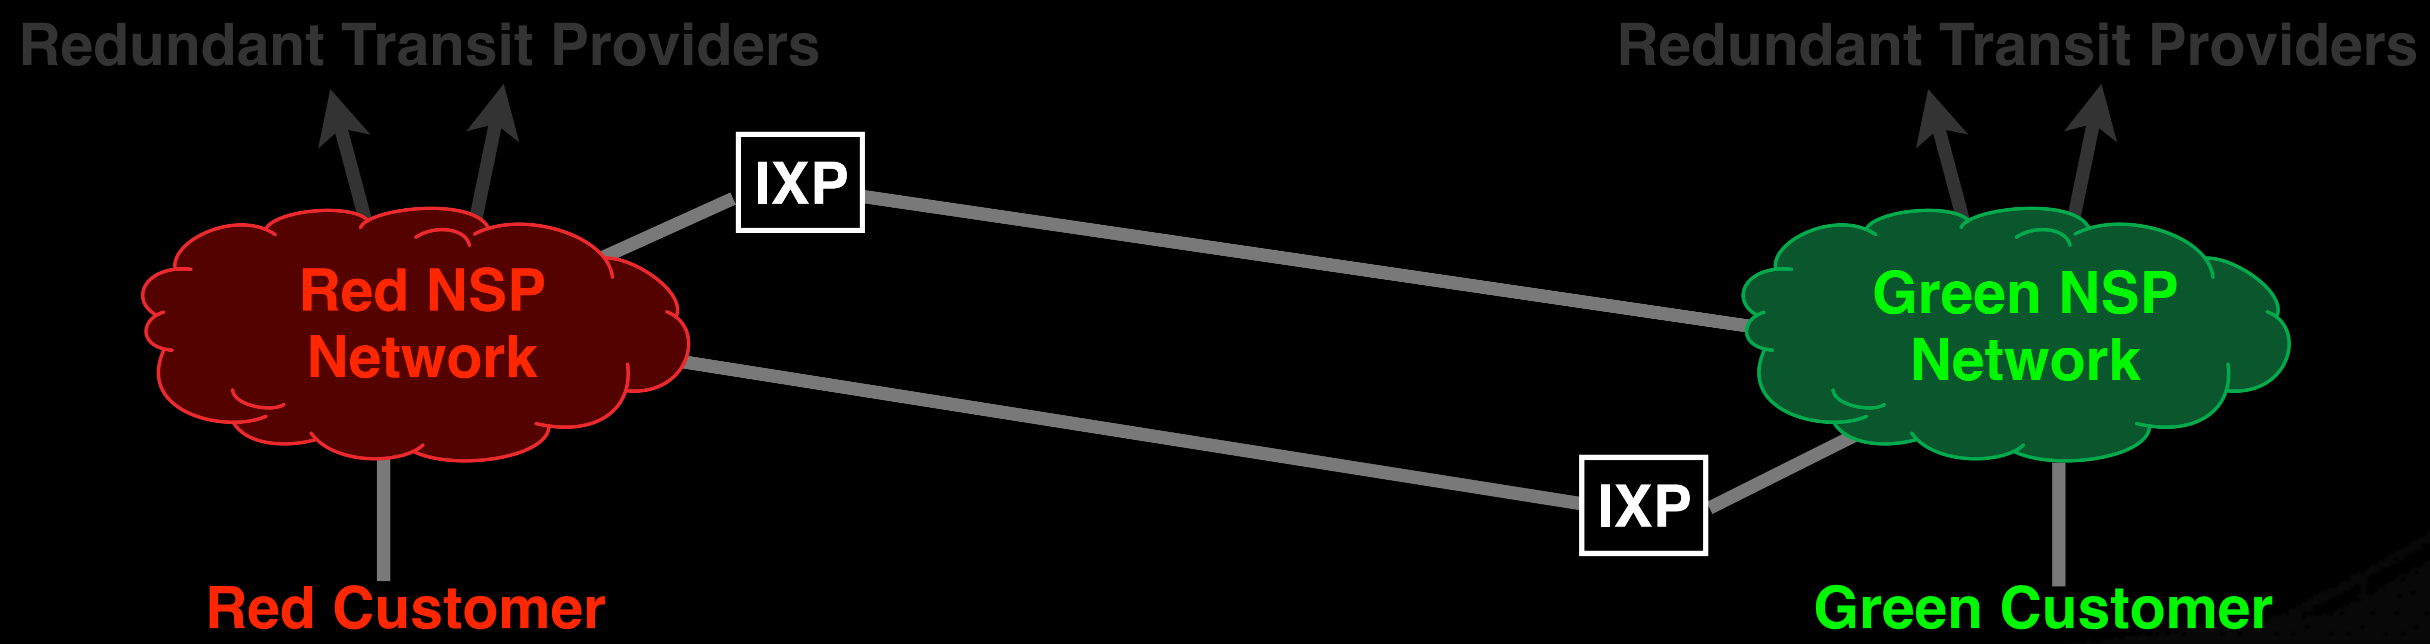

Hot Potato Routing

Red Customer sends to Green Customer via Red NSP

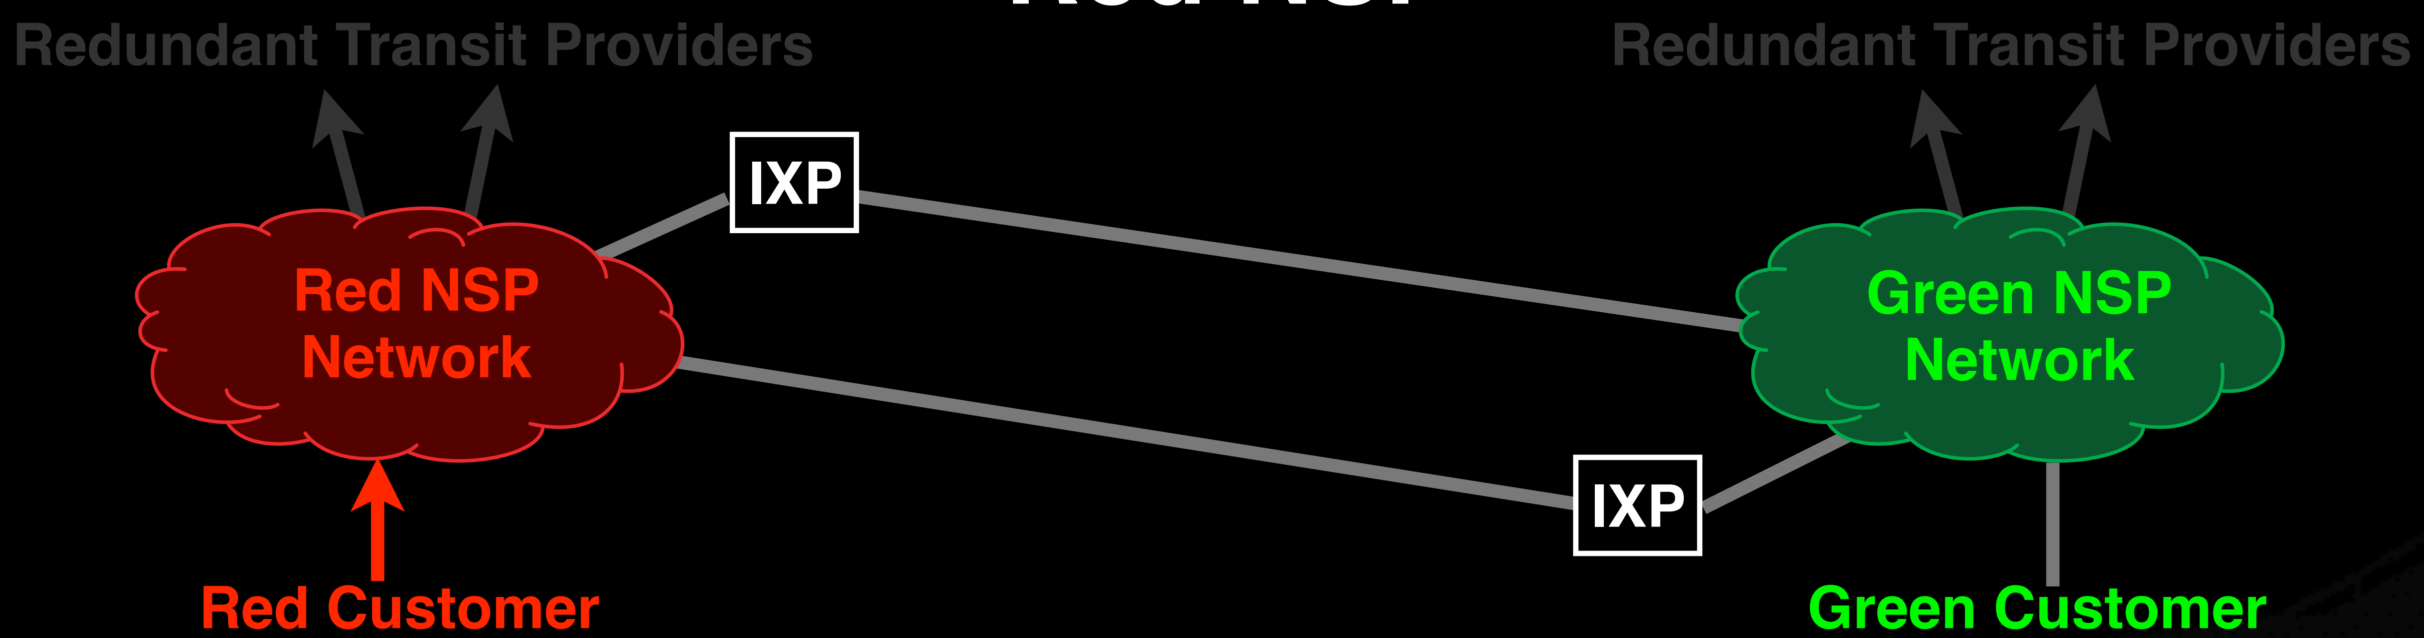

Hot Potato Routing

Red NSP delivers at *nearest* IXP

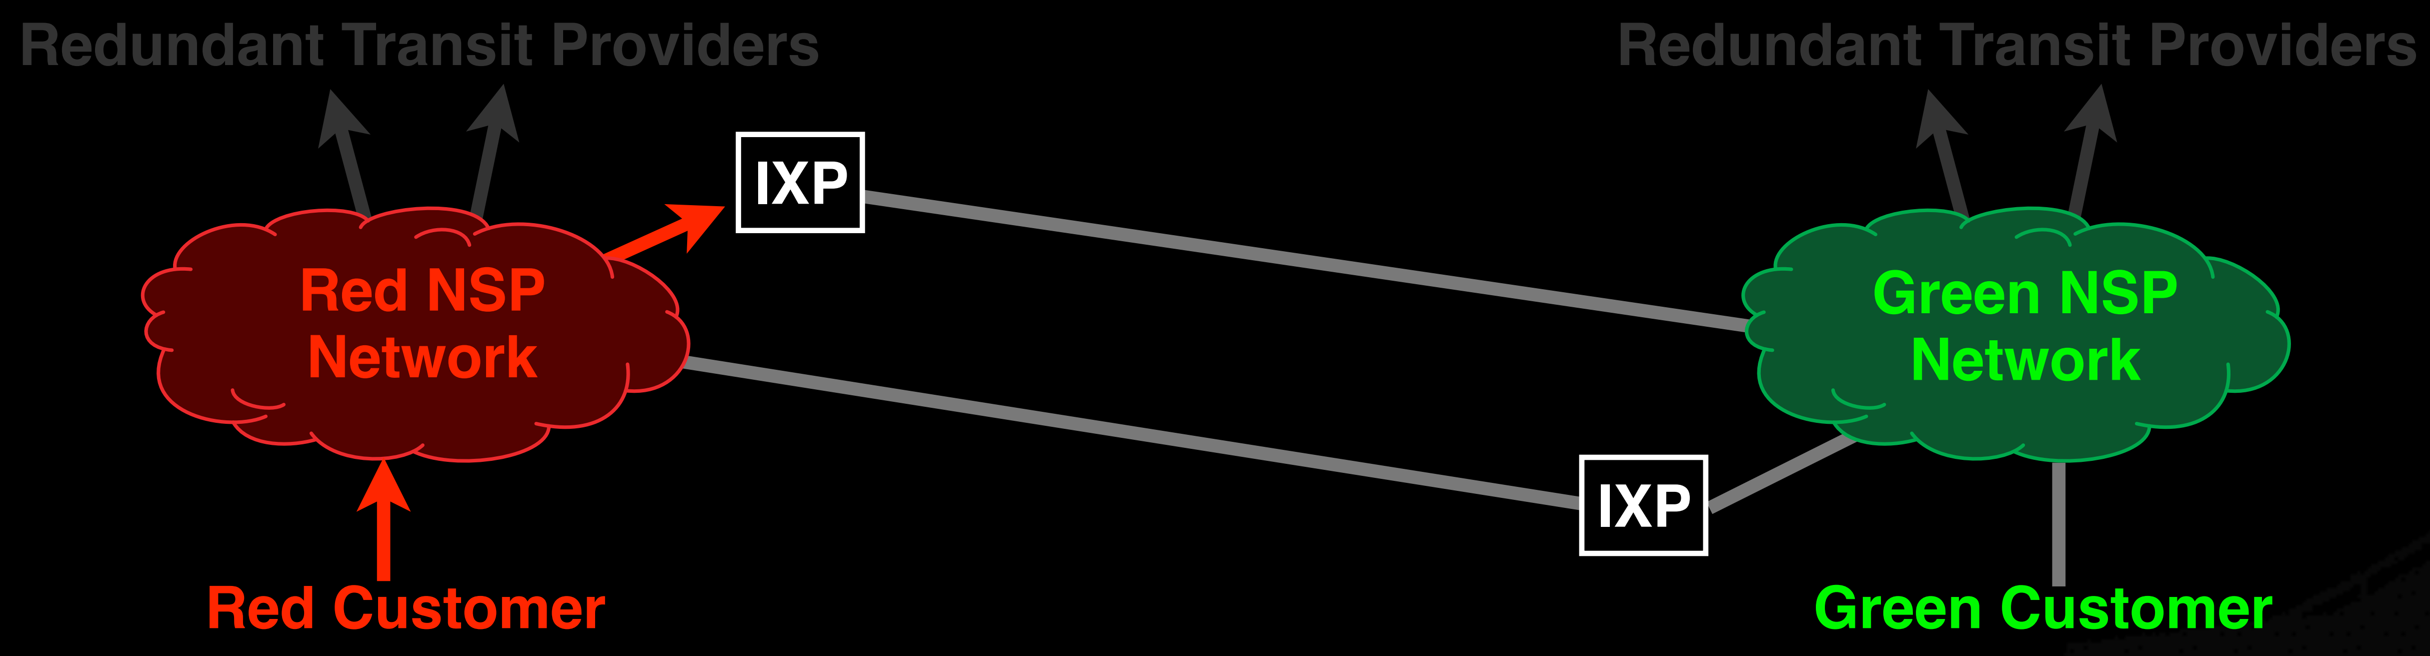

Hot Potato Routing

Green NSP backhauls from distant IXP

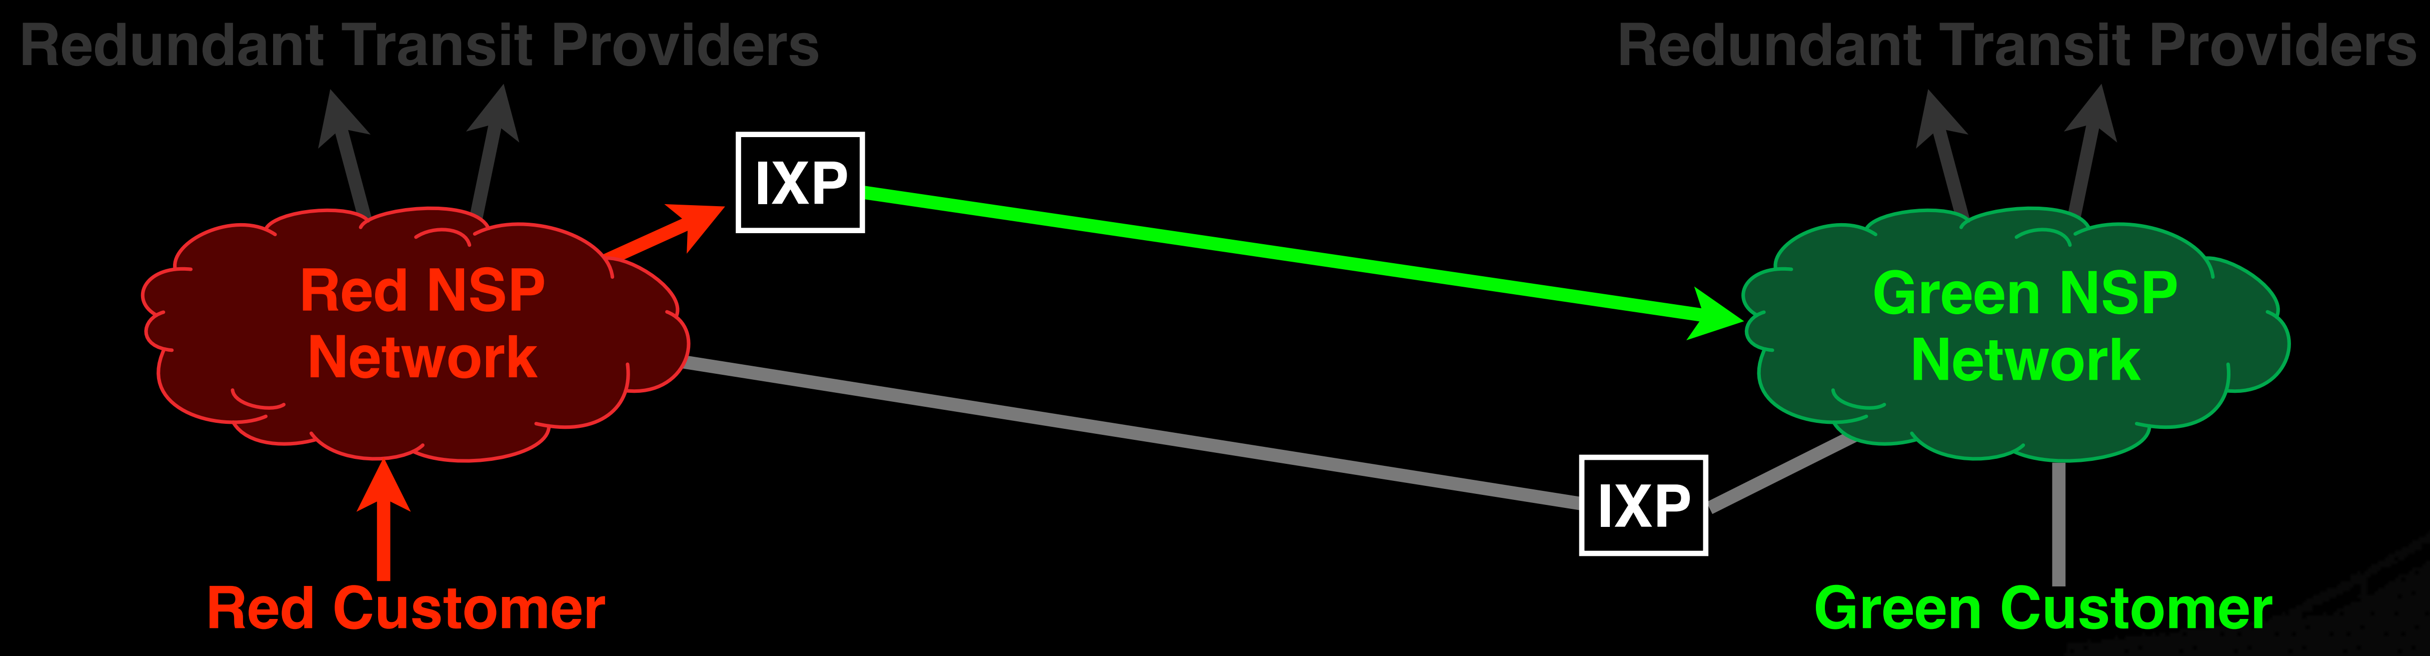

Hot Potato Routing

Green ISP delivers to Green Customer

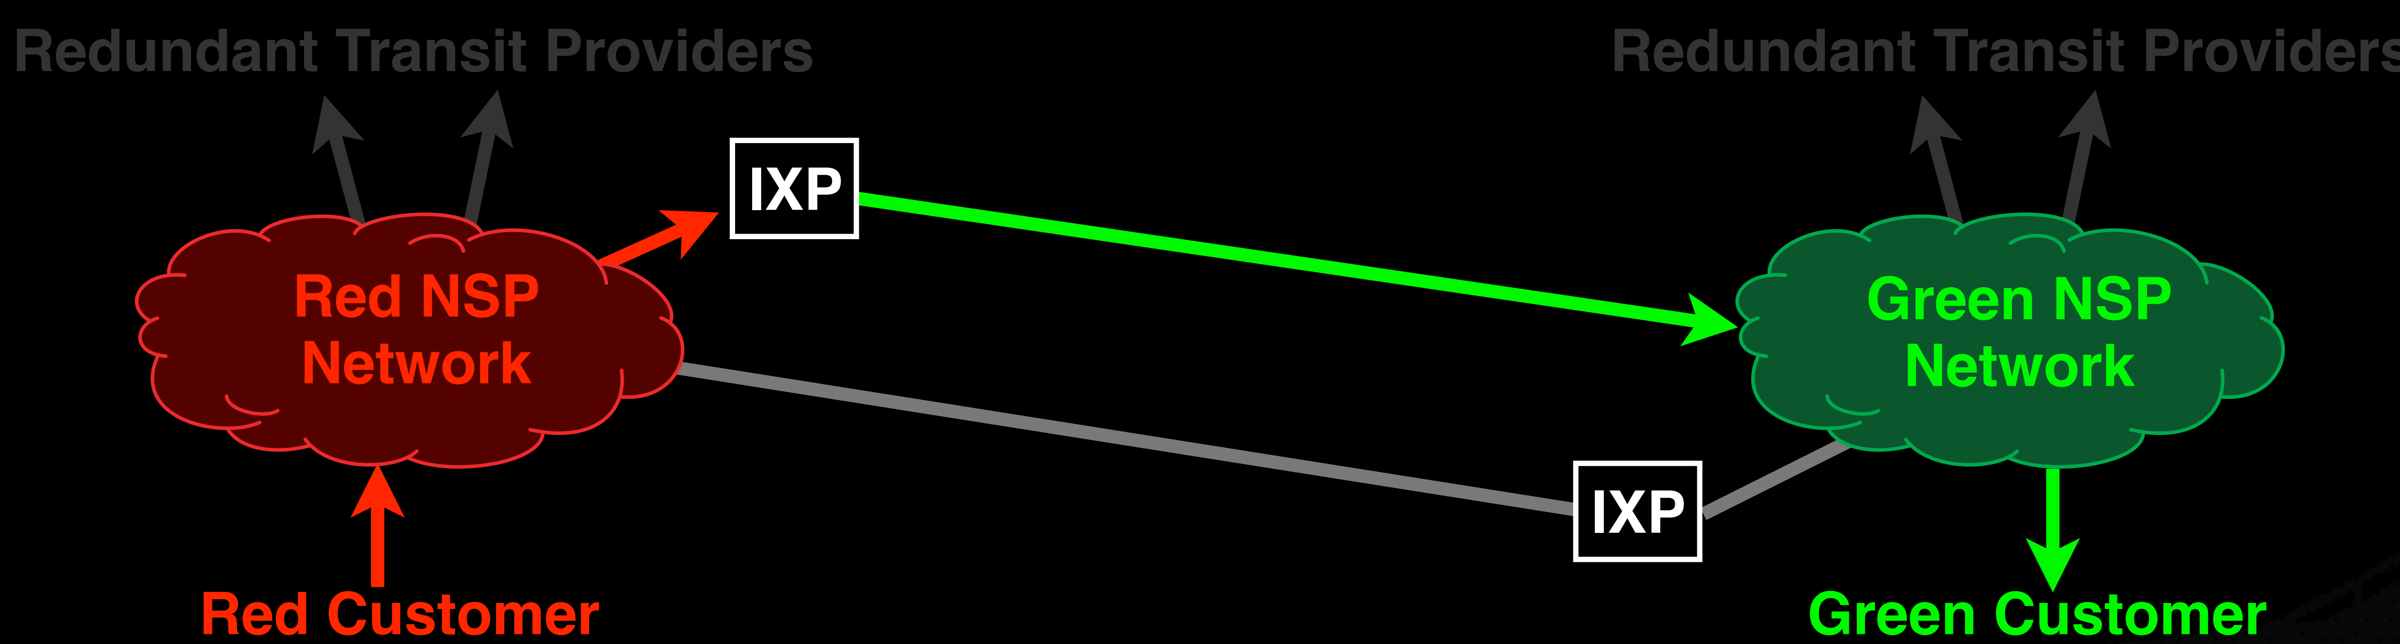

Hot Potato Routing

Green Customer replies via Green NSP

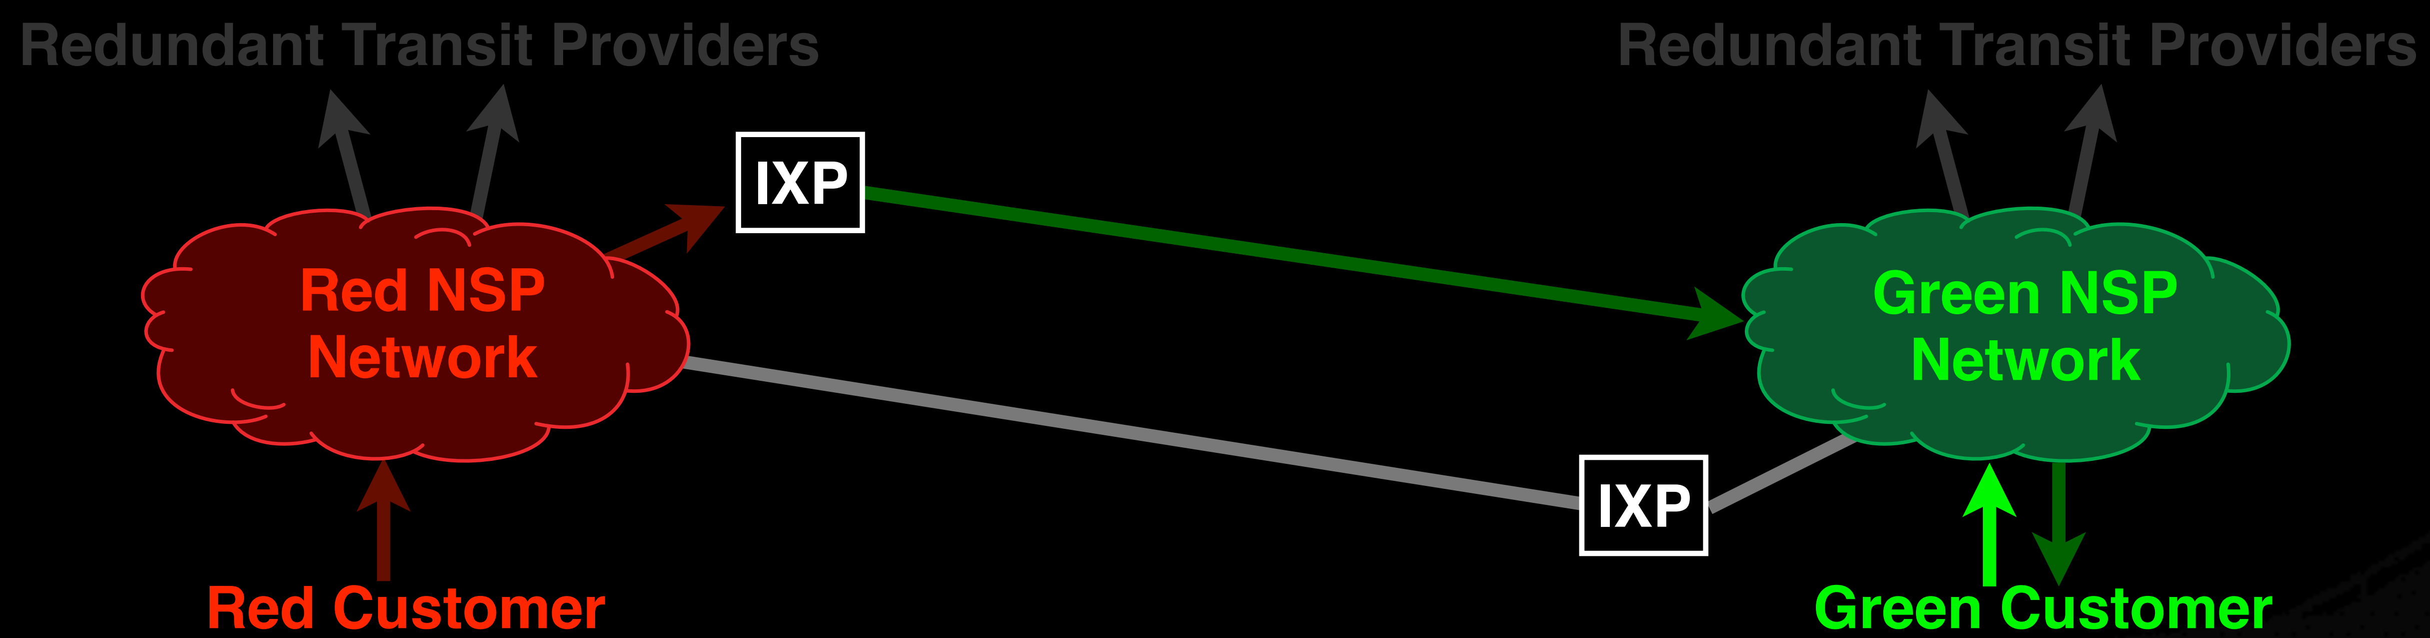

Hot Potato Routing

Green NSP delivers at nearest IXP

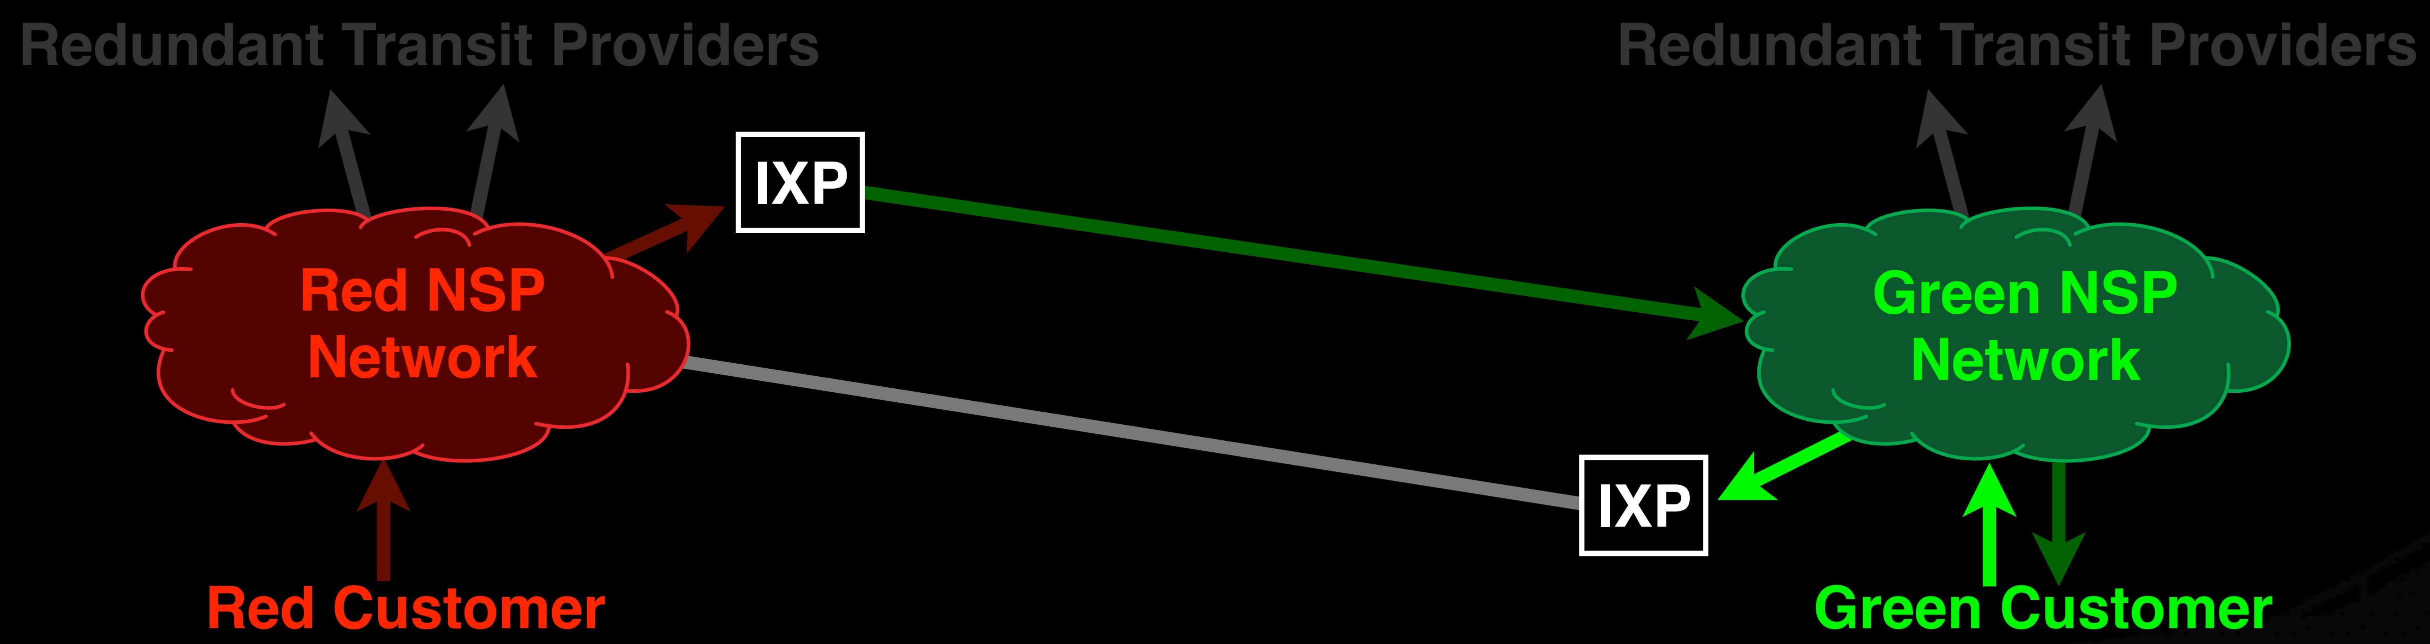

Hot Potato Routing

Red NSP backhauls from distant IXP

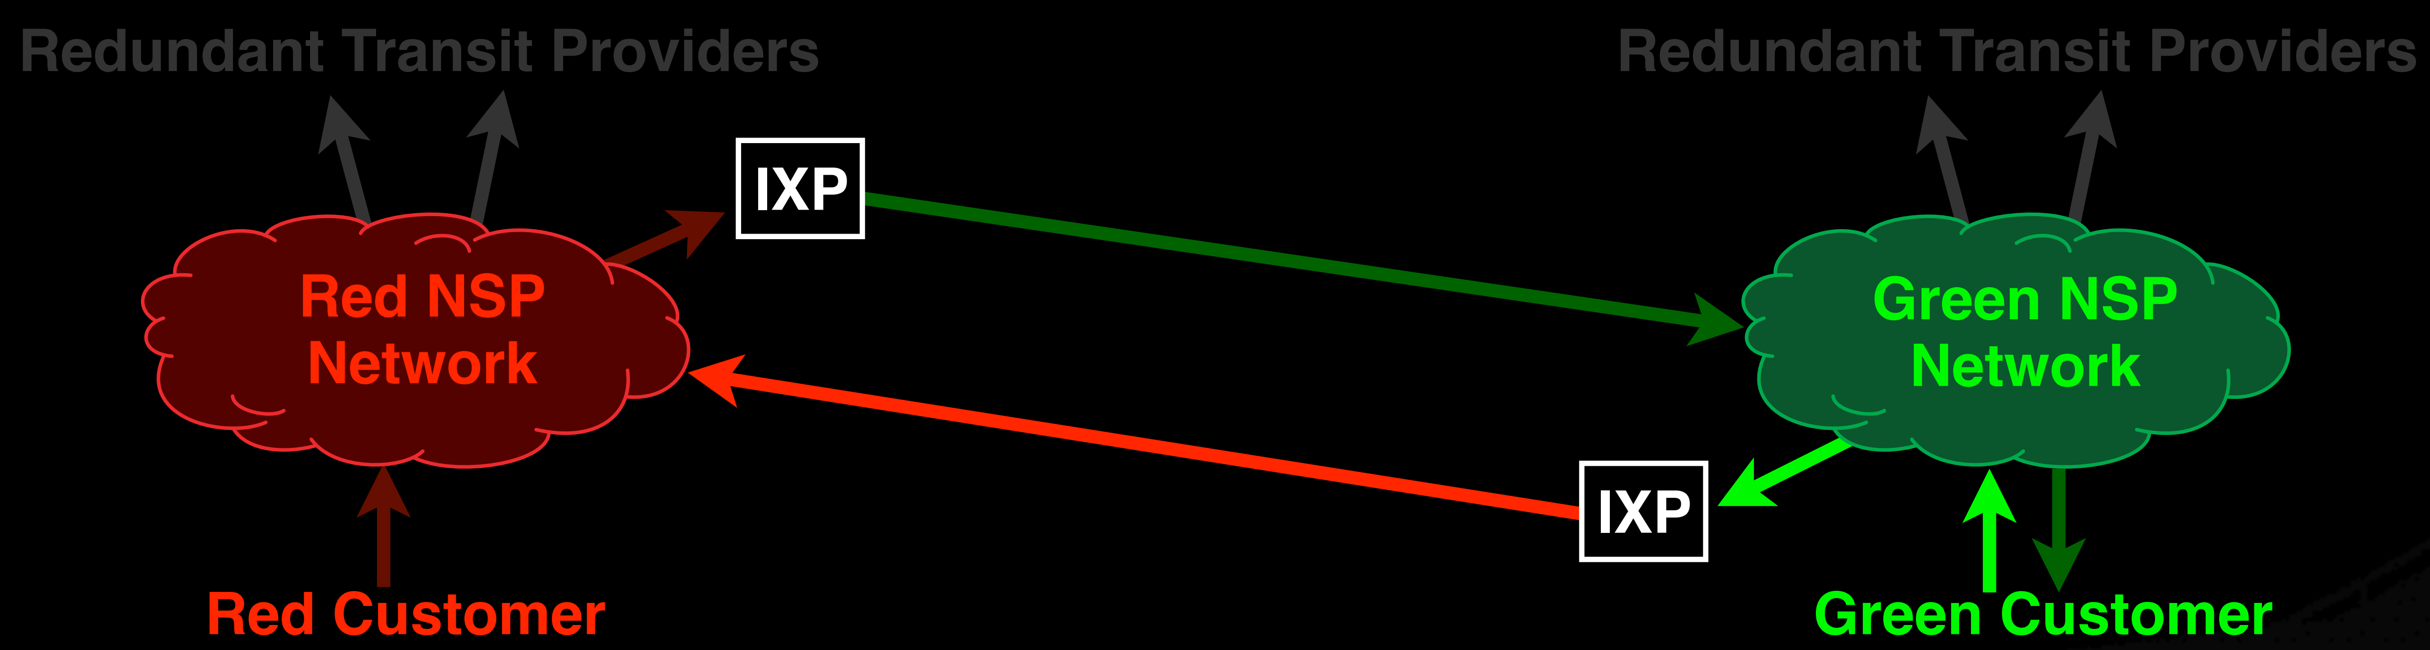

Hot Potato Routing

Red NSP delivers to Red Customer

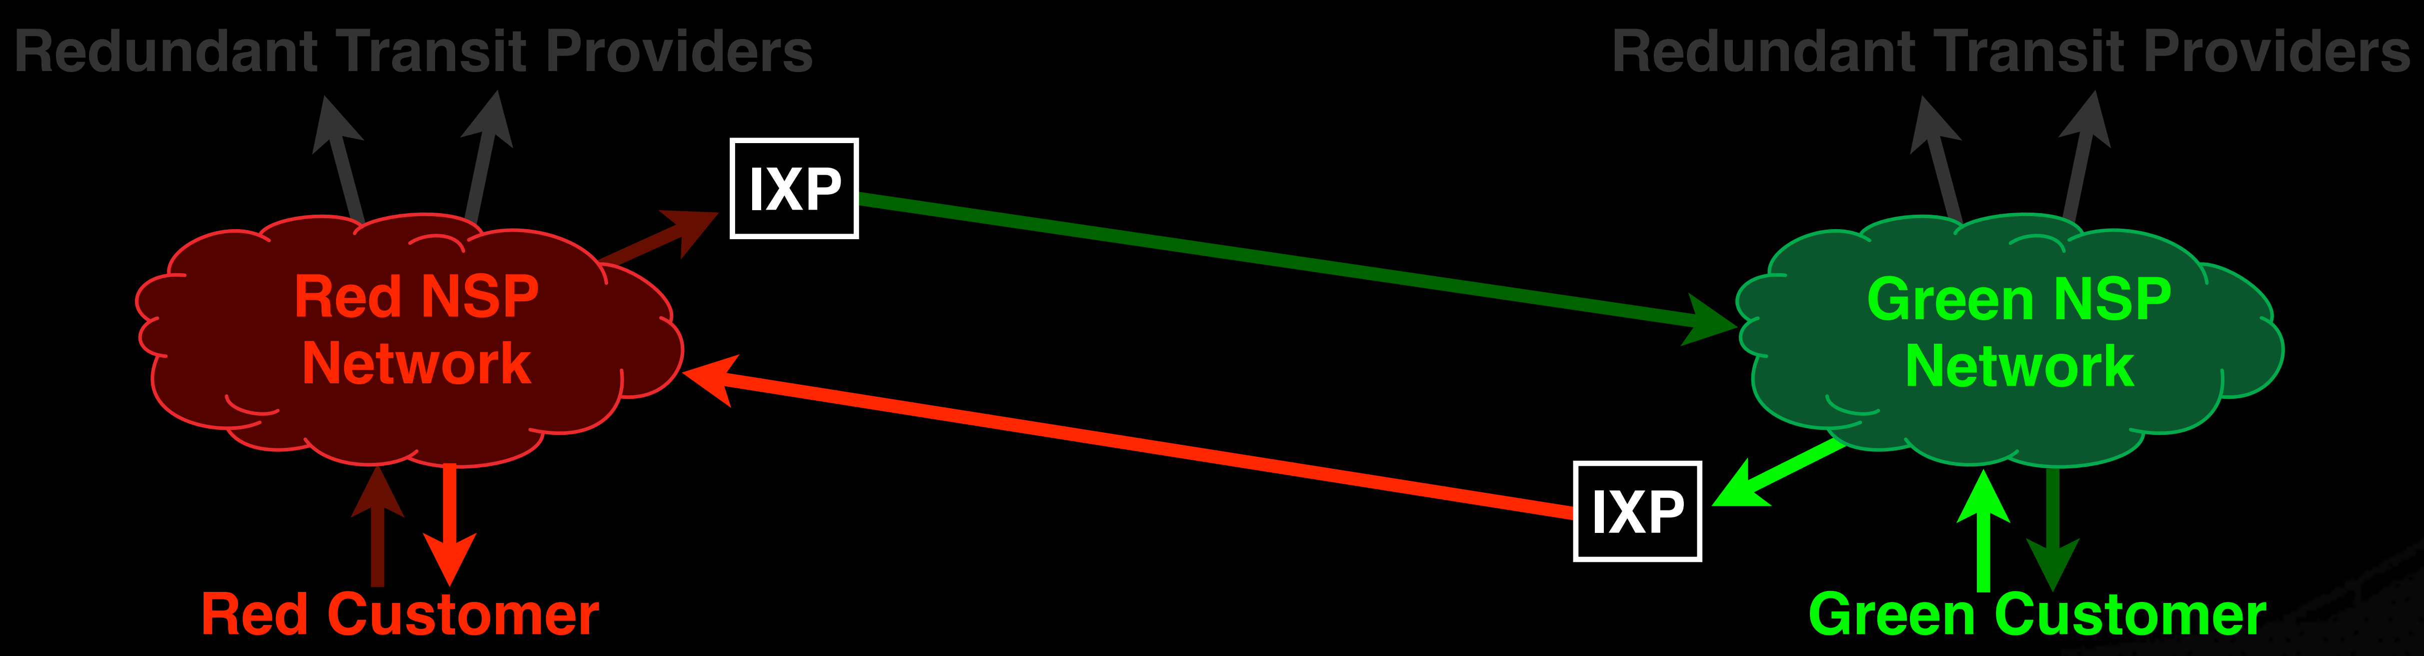

Hot Potato Routing

Red Network is responsible for its own costs

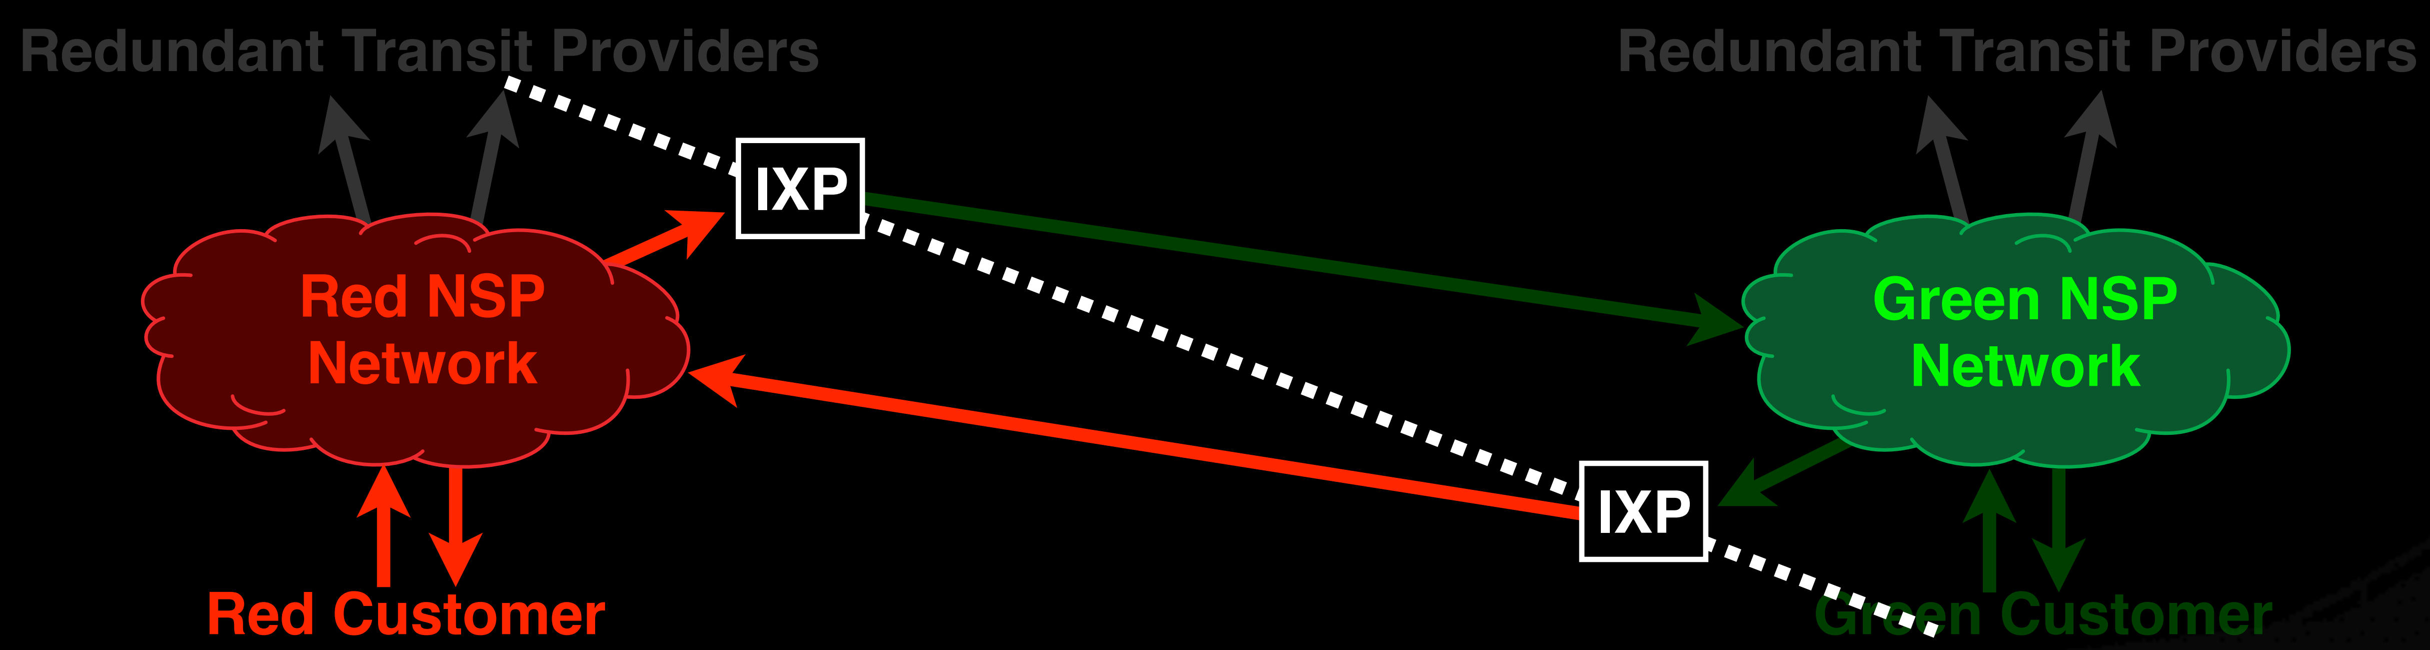

Hot Potato Routing

Green Network is responsible for its own costs

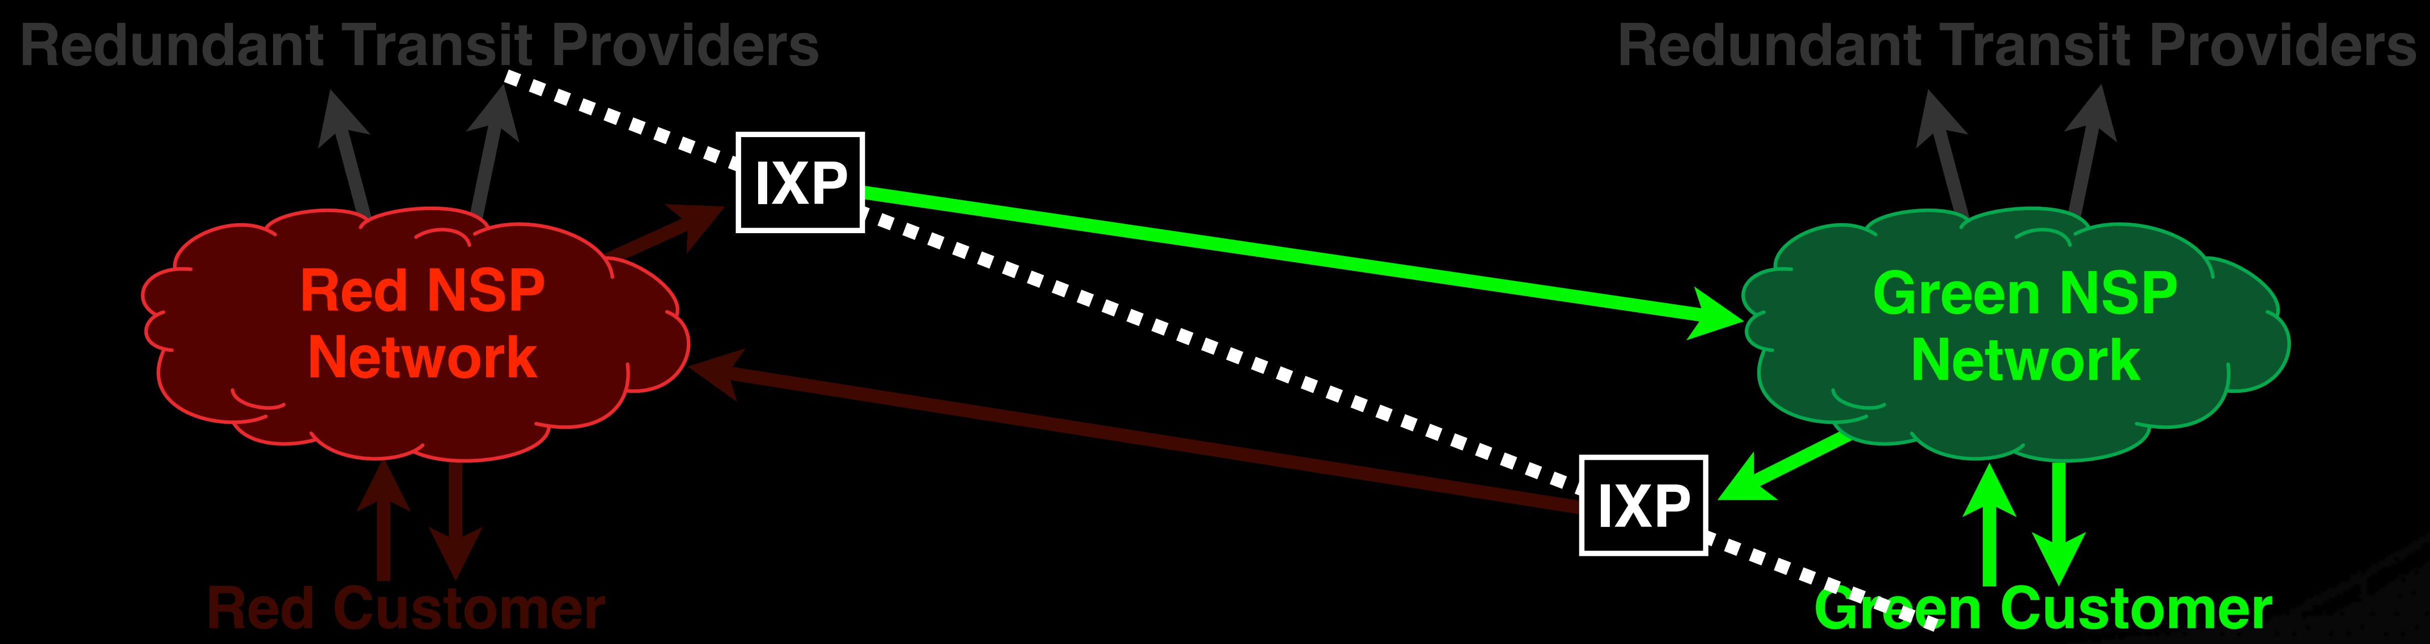

Hot Potato Routing

Symmetry: Fair sharing of costs

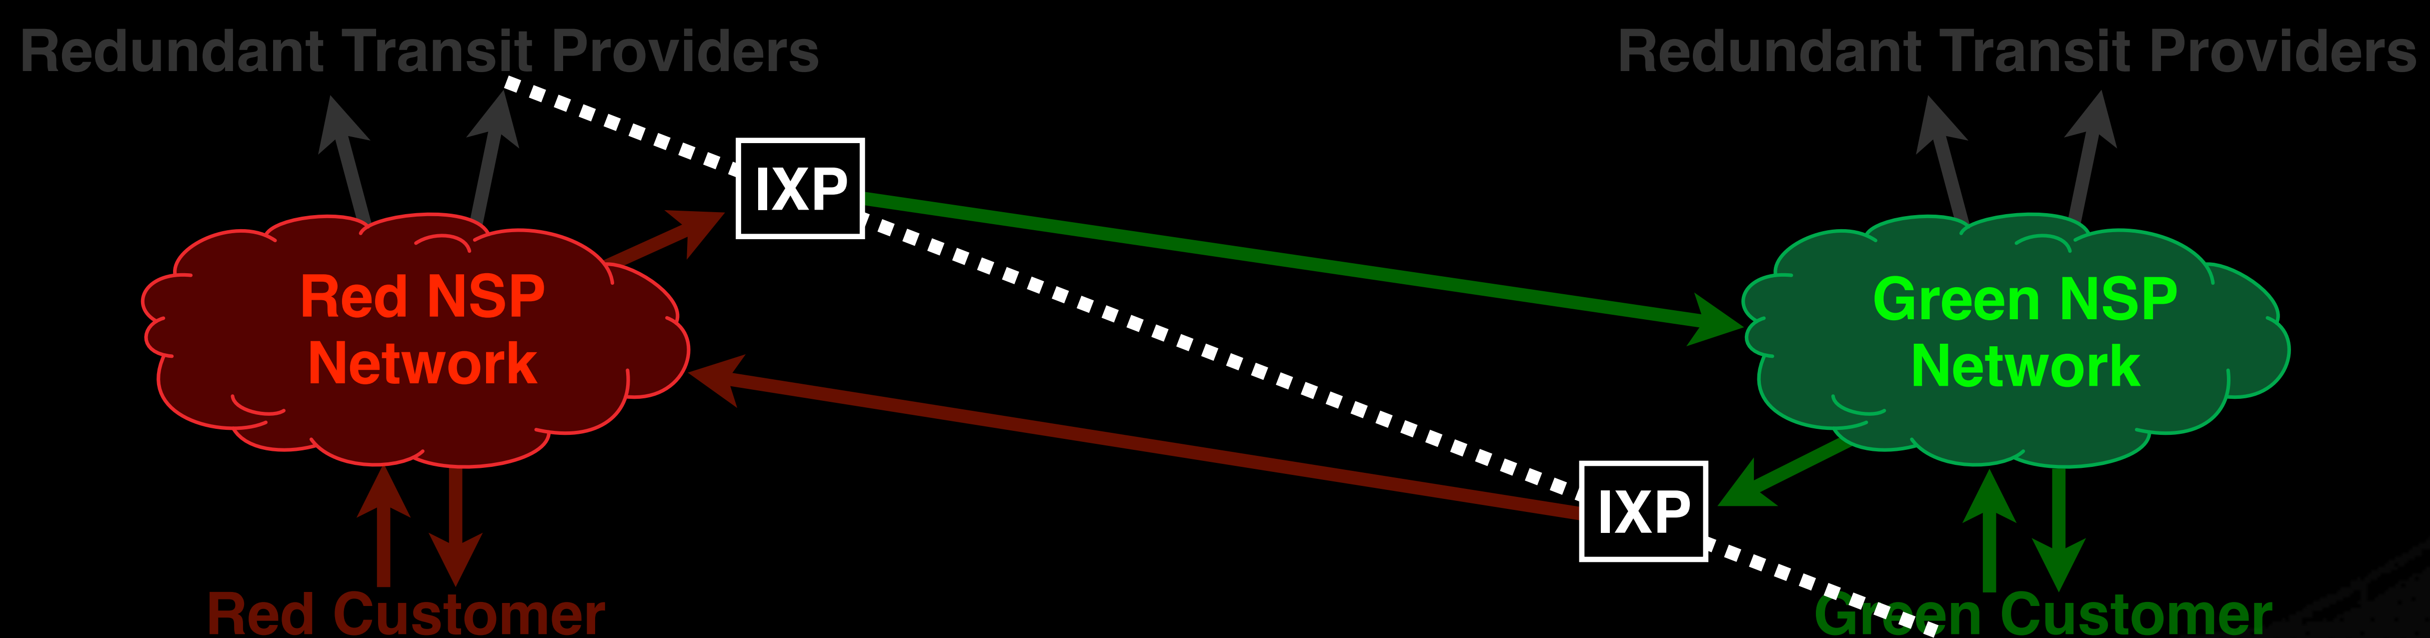

The old circuit-switched networks have dubbed the Internet financial model "bill and keep"

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.

Packet Clearing House

Distribution of IXPs

Half of all countries still have no IXP, while others have dozens.

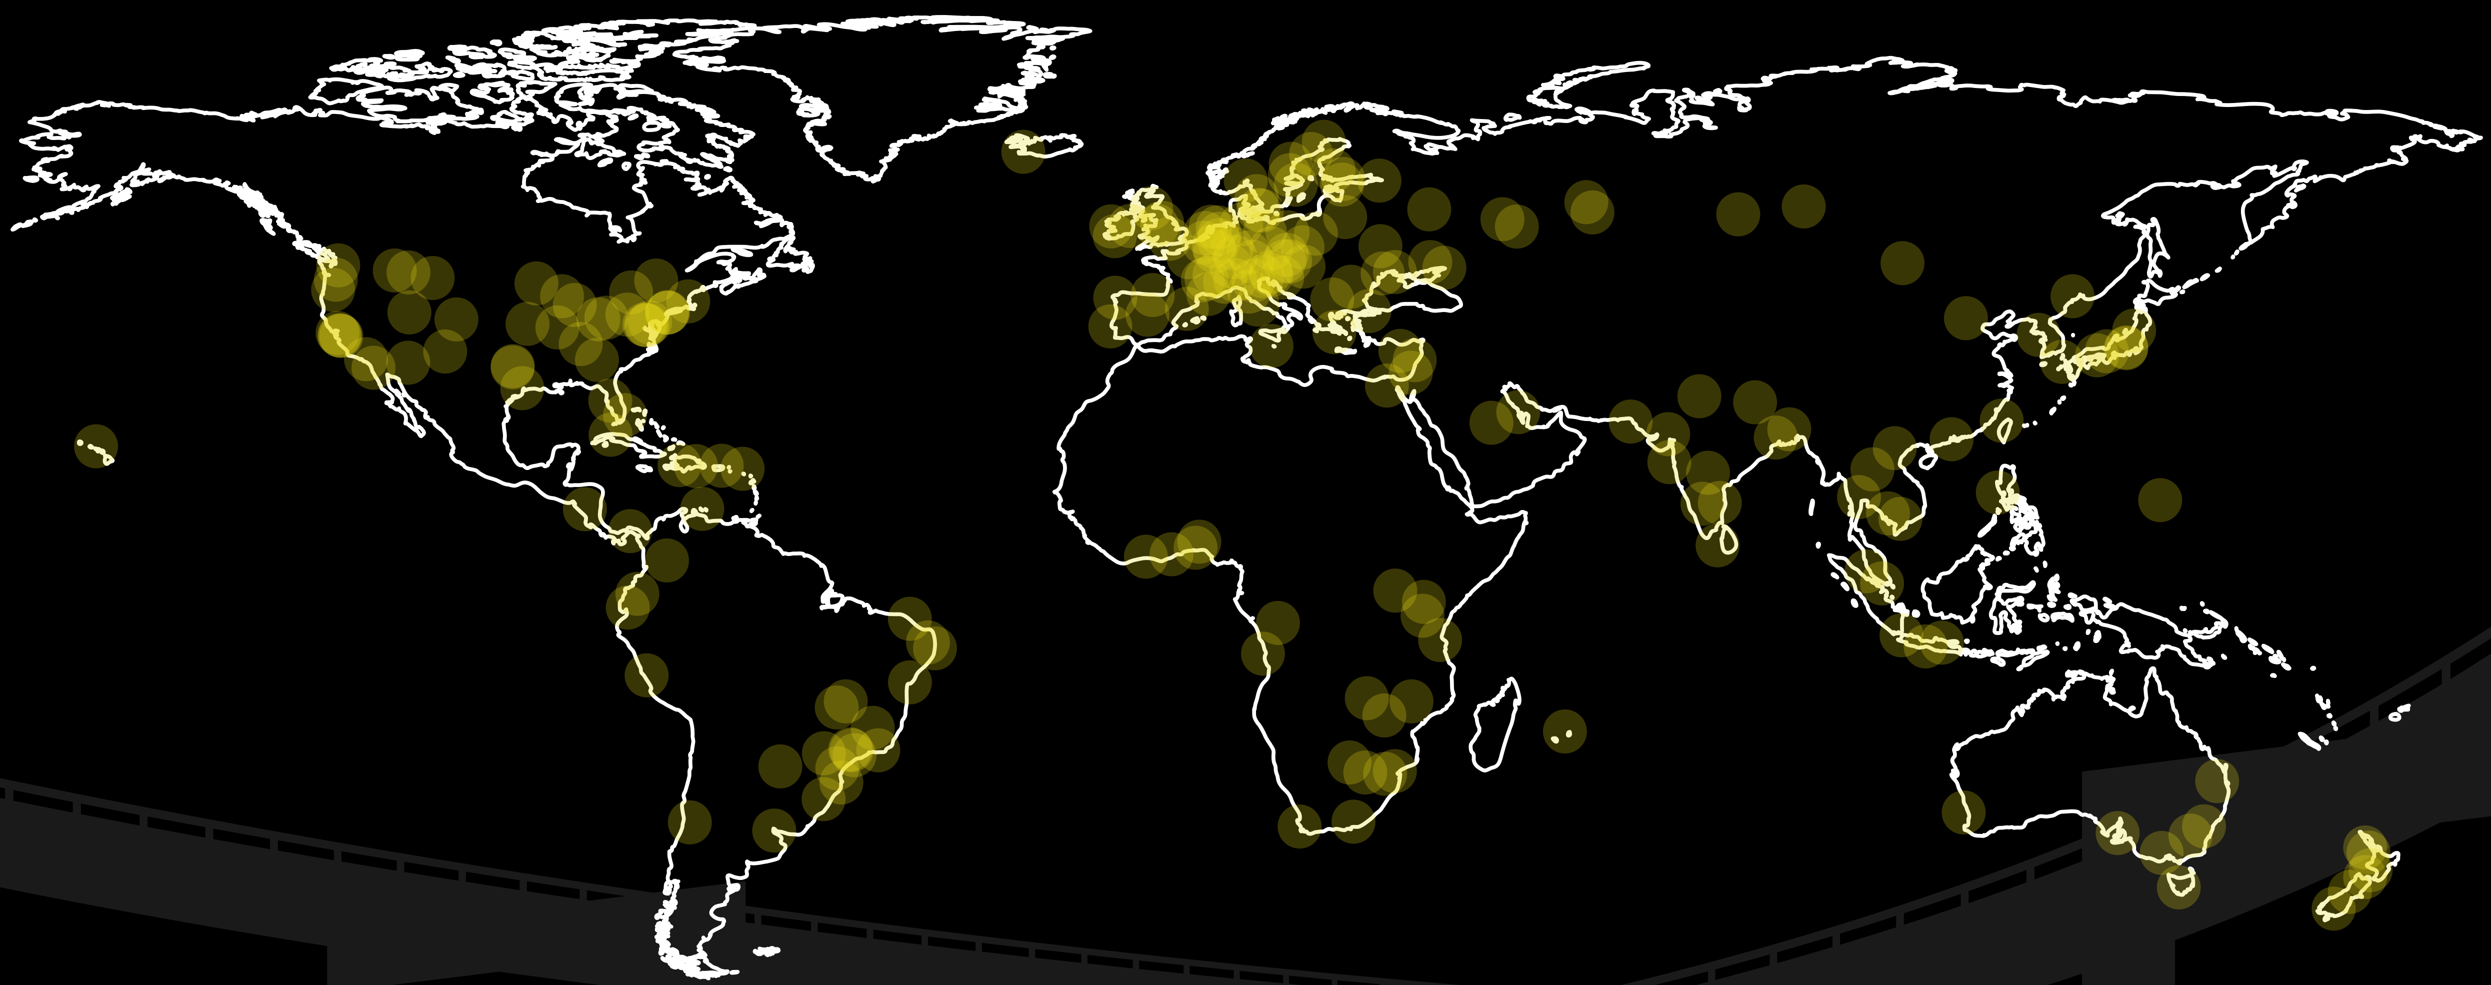

2011 Peering Survey

We conducted a survey of ISPs between October 2010 and March 2011, and analyzed 142,210 peering agreements.

86% of all ISPs represented, in 96 countries

99.51% of peering agreements required no written contract

99.73% of peering agreements had symmetric terms

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**.

Packet Clearing House

Value decreases with time & distance

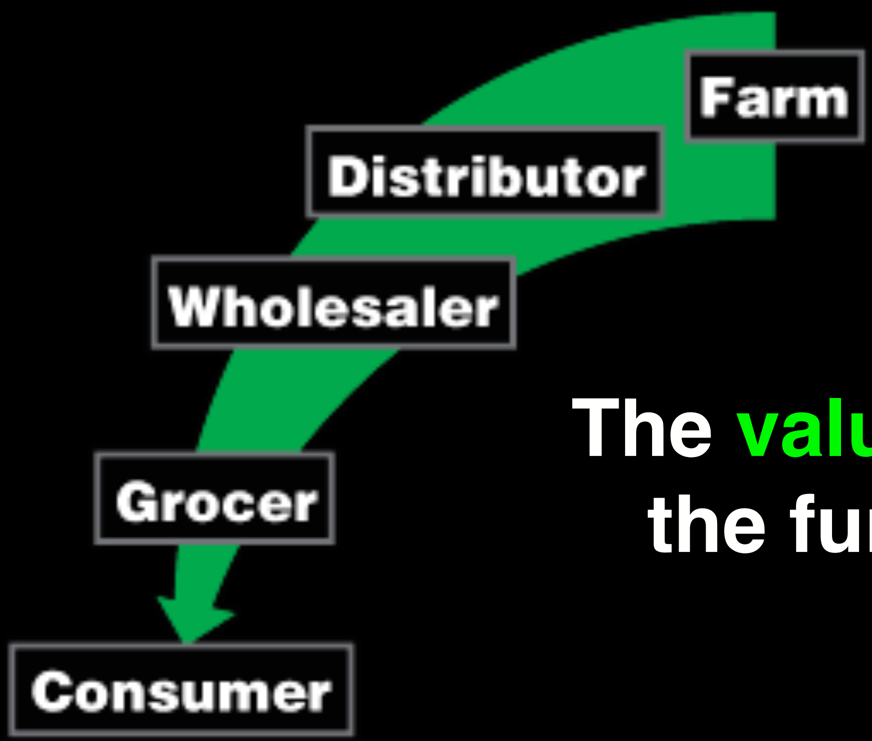

The value of a banana decreases, the further it gets from the farm which produced it.

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

Packet Clearing House

Cost increases with time & distance

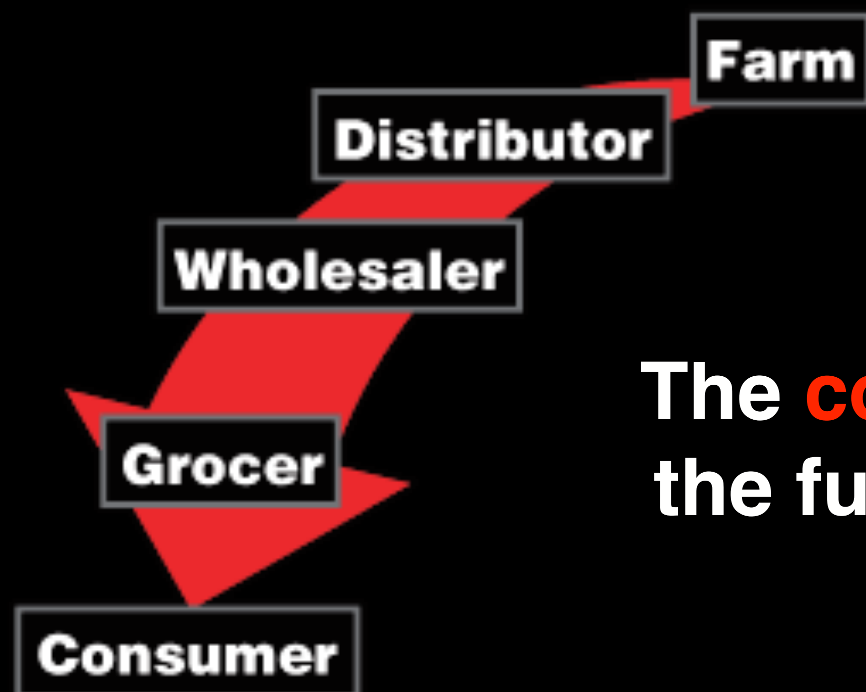

The cost of a banana increases, the further it gets from the farm which produced it.

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

PCH
Packet Clearing House

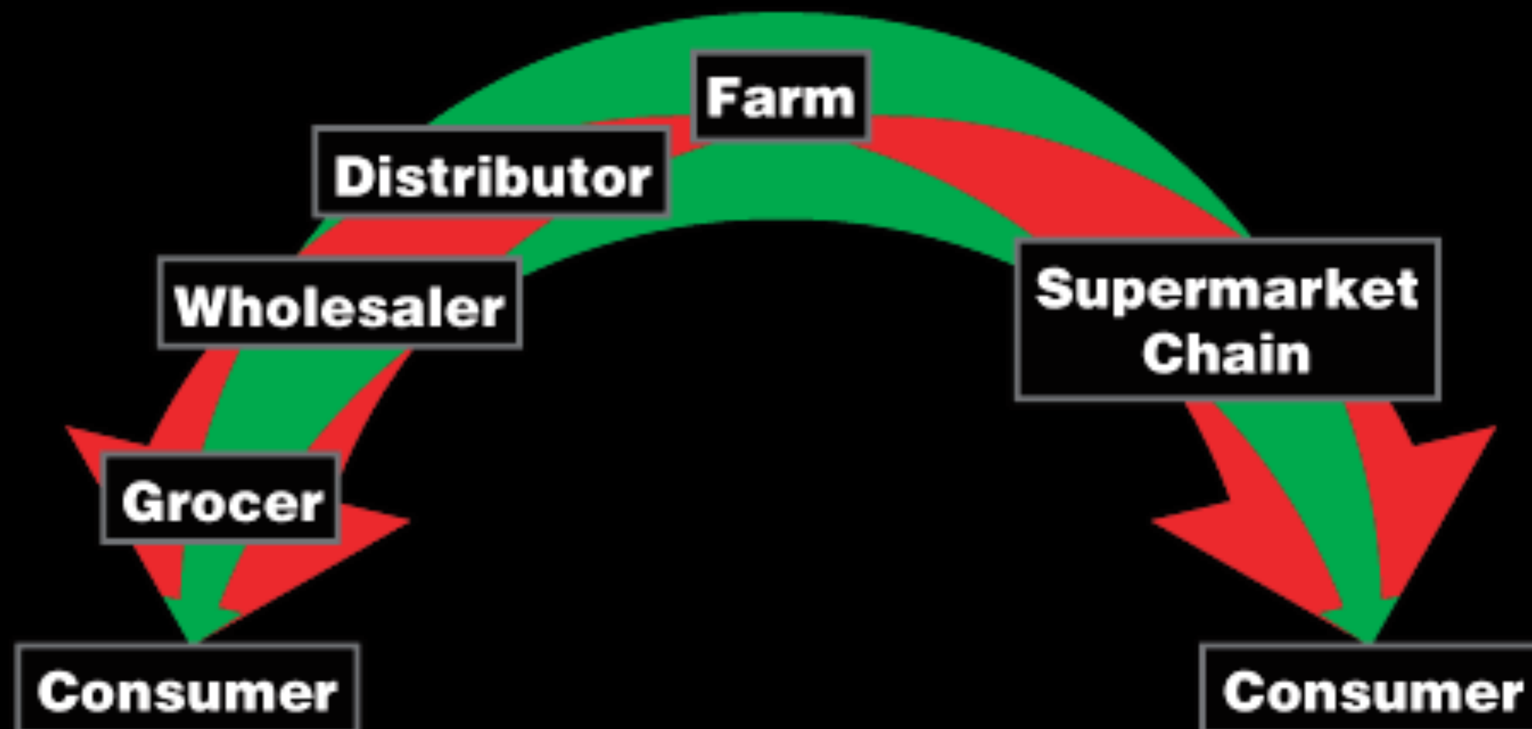

In a competitive environment, retail price is limited by competition, so time and distance influence the price more than the number of middlemen.

The problem is the same:

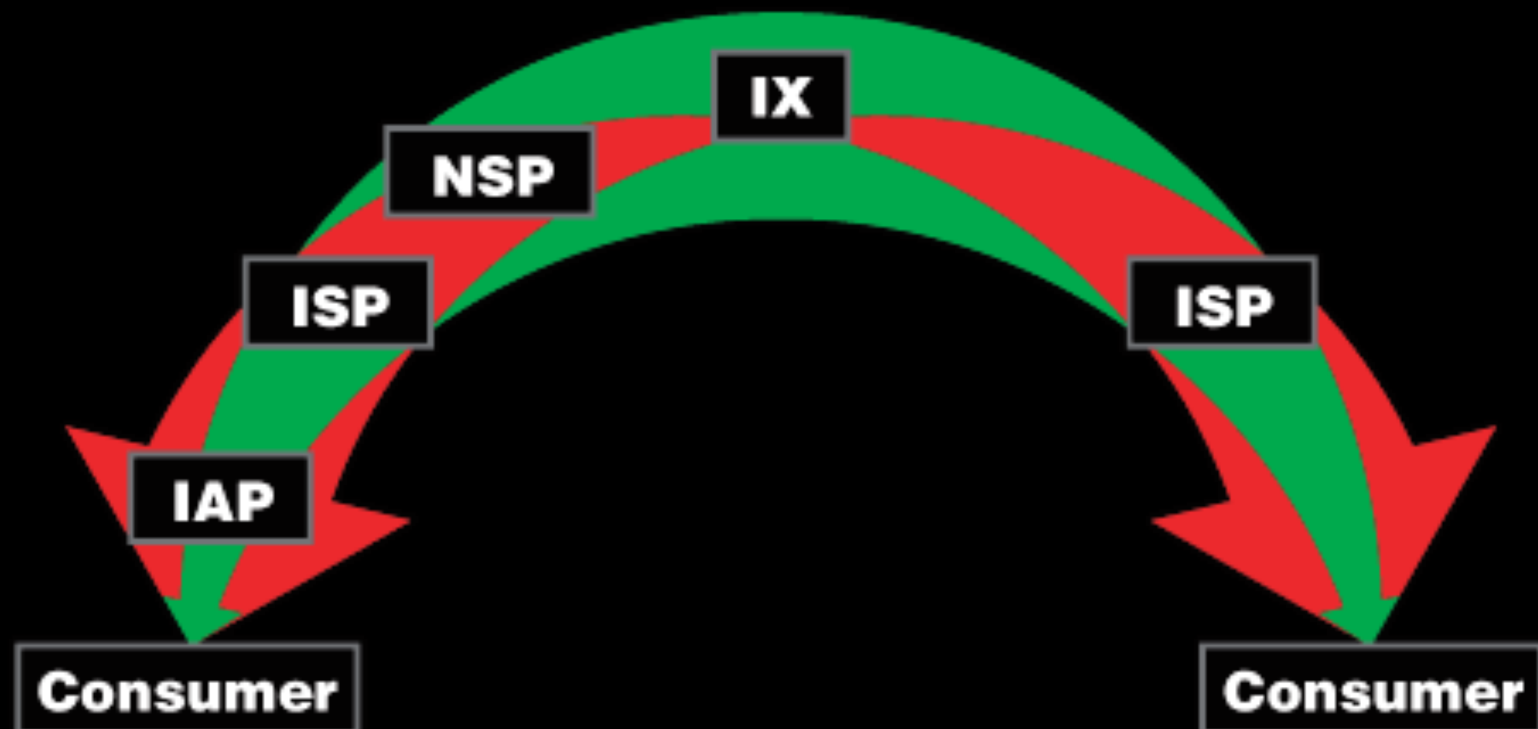

ISPs form a delivery chain, bringing perishable bits to the consumers who purchase them.

So how do we improve things?

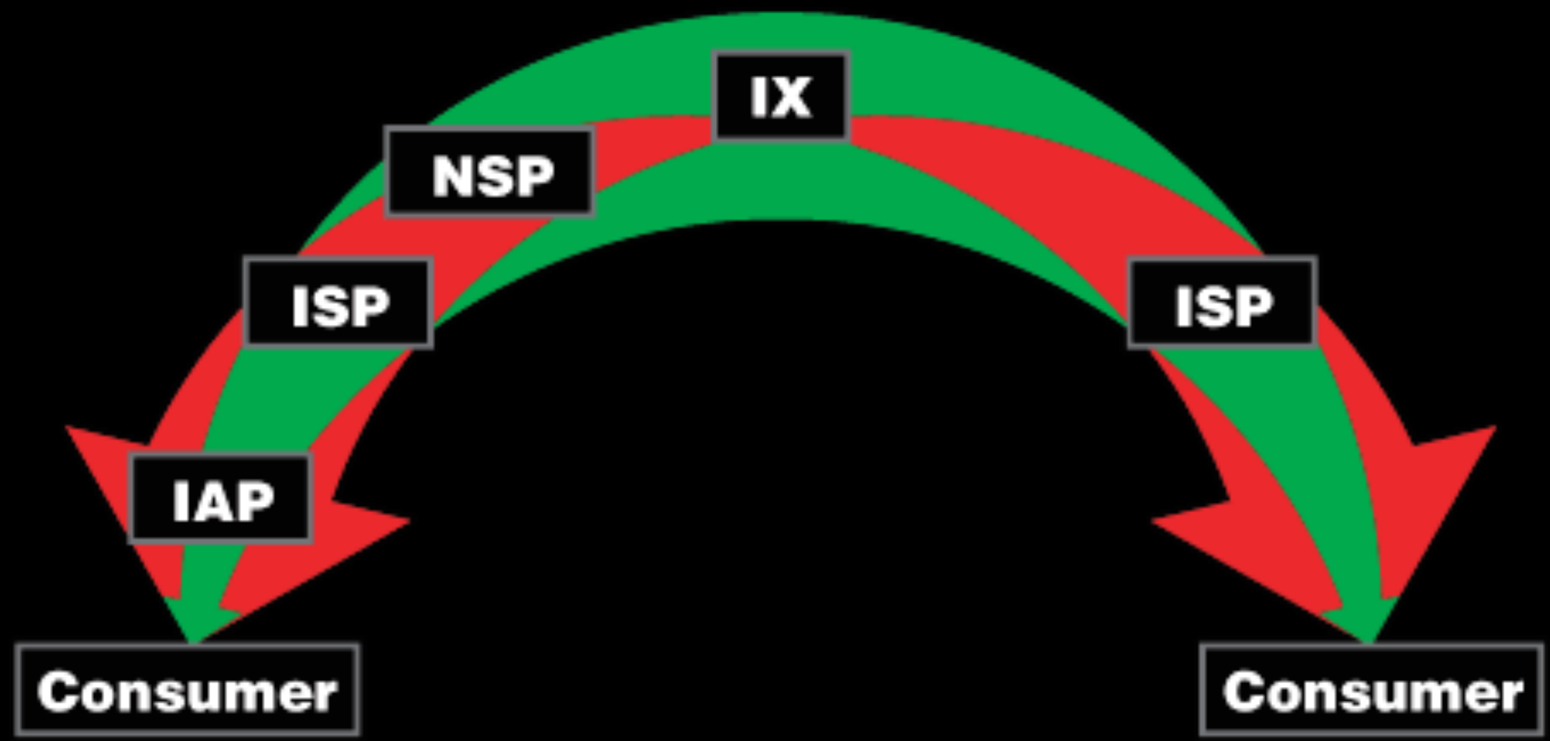

Bring the customer nearer an IX...

High cost
Low value

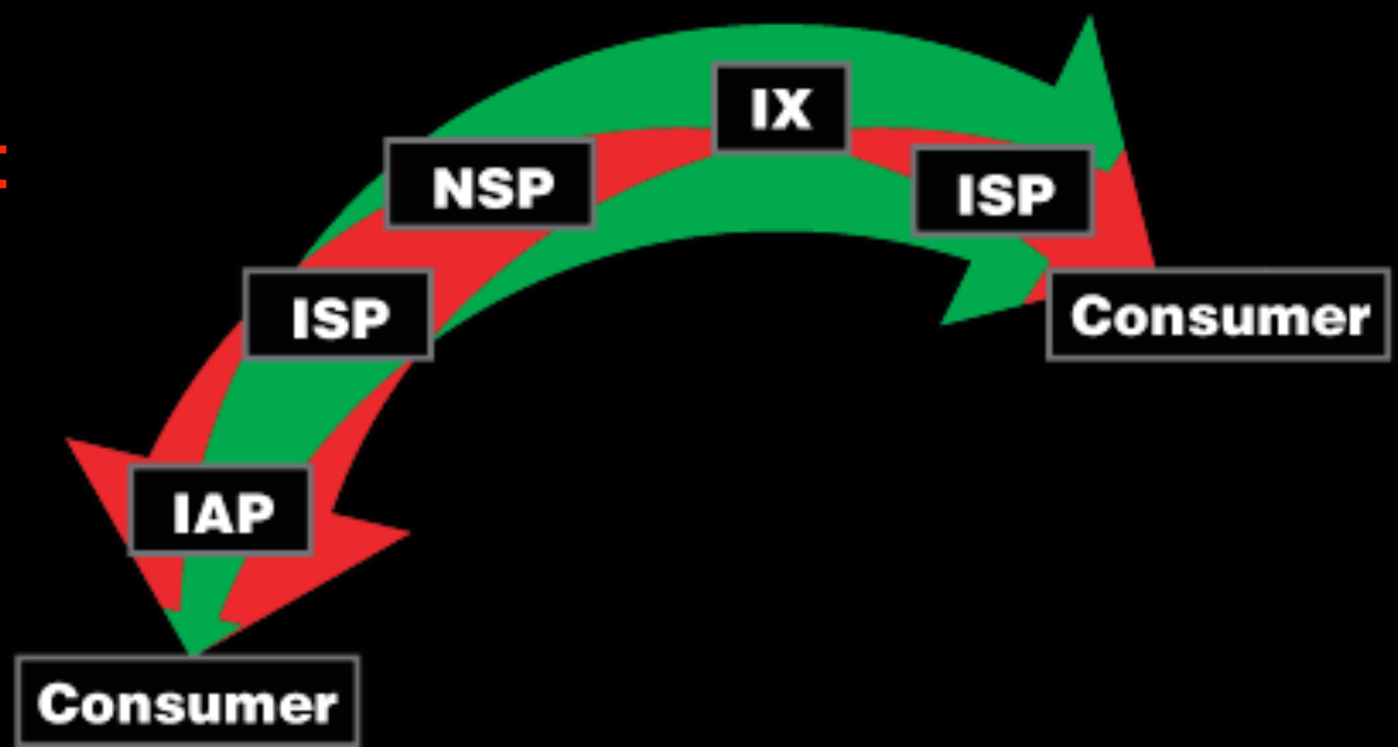

Low cost
High value

...or bring an IX nearer the customer.

Packet Clearing House

speed * distance = cost

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.

An Apparent Contradiction Resolved

In order to optimise the performance and profitability of Internet transit provision, users must be incentivised to select services reachable through peering, rather than through transit.

Therefore, peering circuits must be larger than transit circuits, even if that means that they operate at much lower utilization.

1bps utilization of a 1bps circuit, 100% full, 1 second

Average 1 second to completion

Packet Clearing House

1bps utilization of a 10bps circuit, 10% full, 1 second

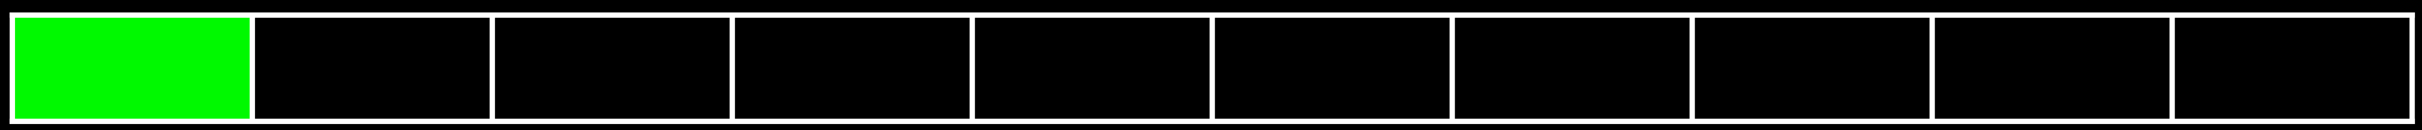

Average 0.1 seconds to completion

Packet Clearing House

1bps utilization of a 10bps circuit, 100% full, 1 second

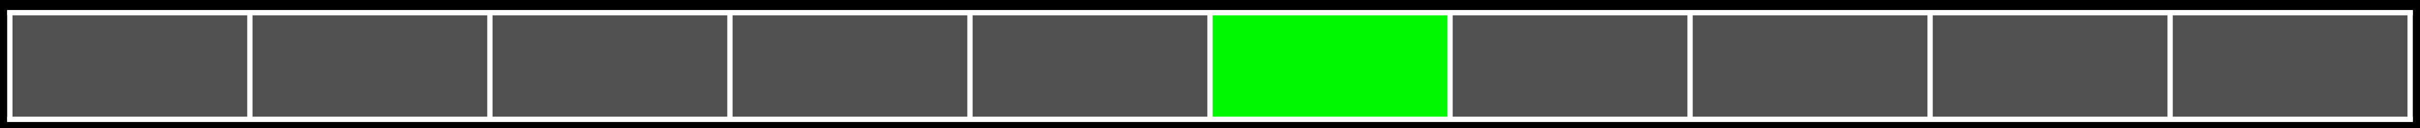

Average 0.55 seconds to completion

1bps utilization of a 100bps circuit, 1% full, 1 second

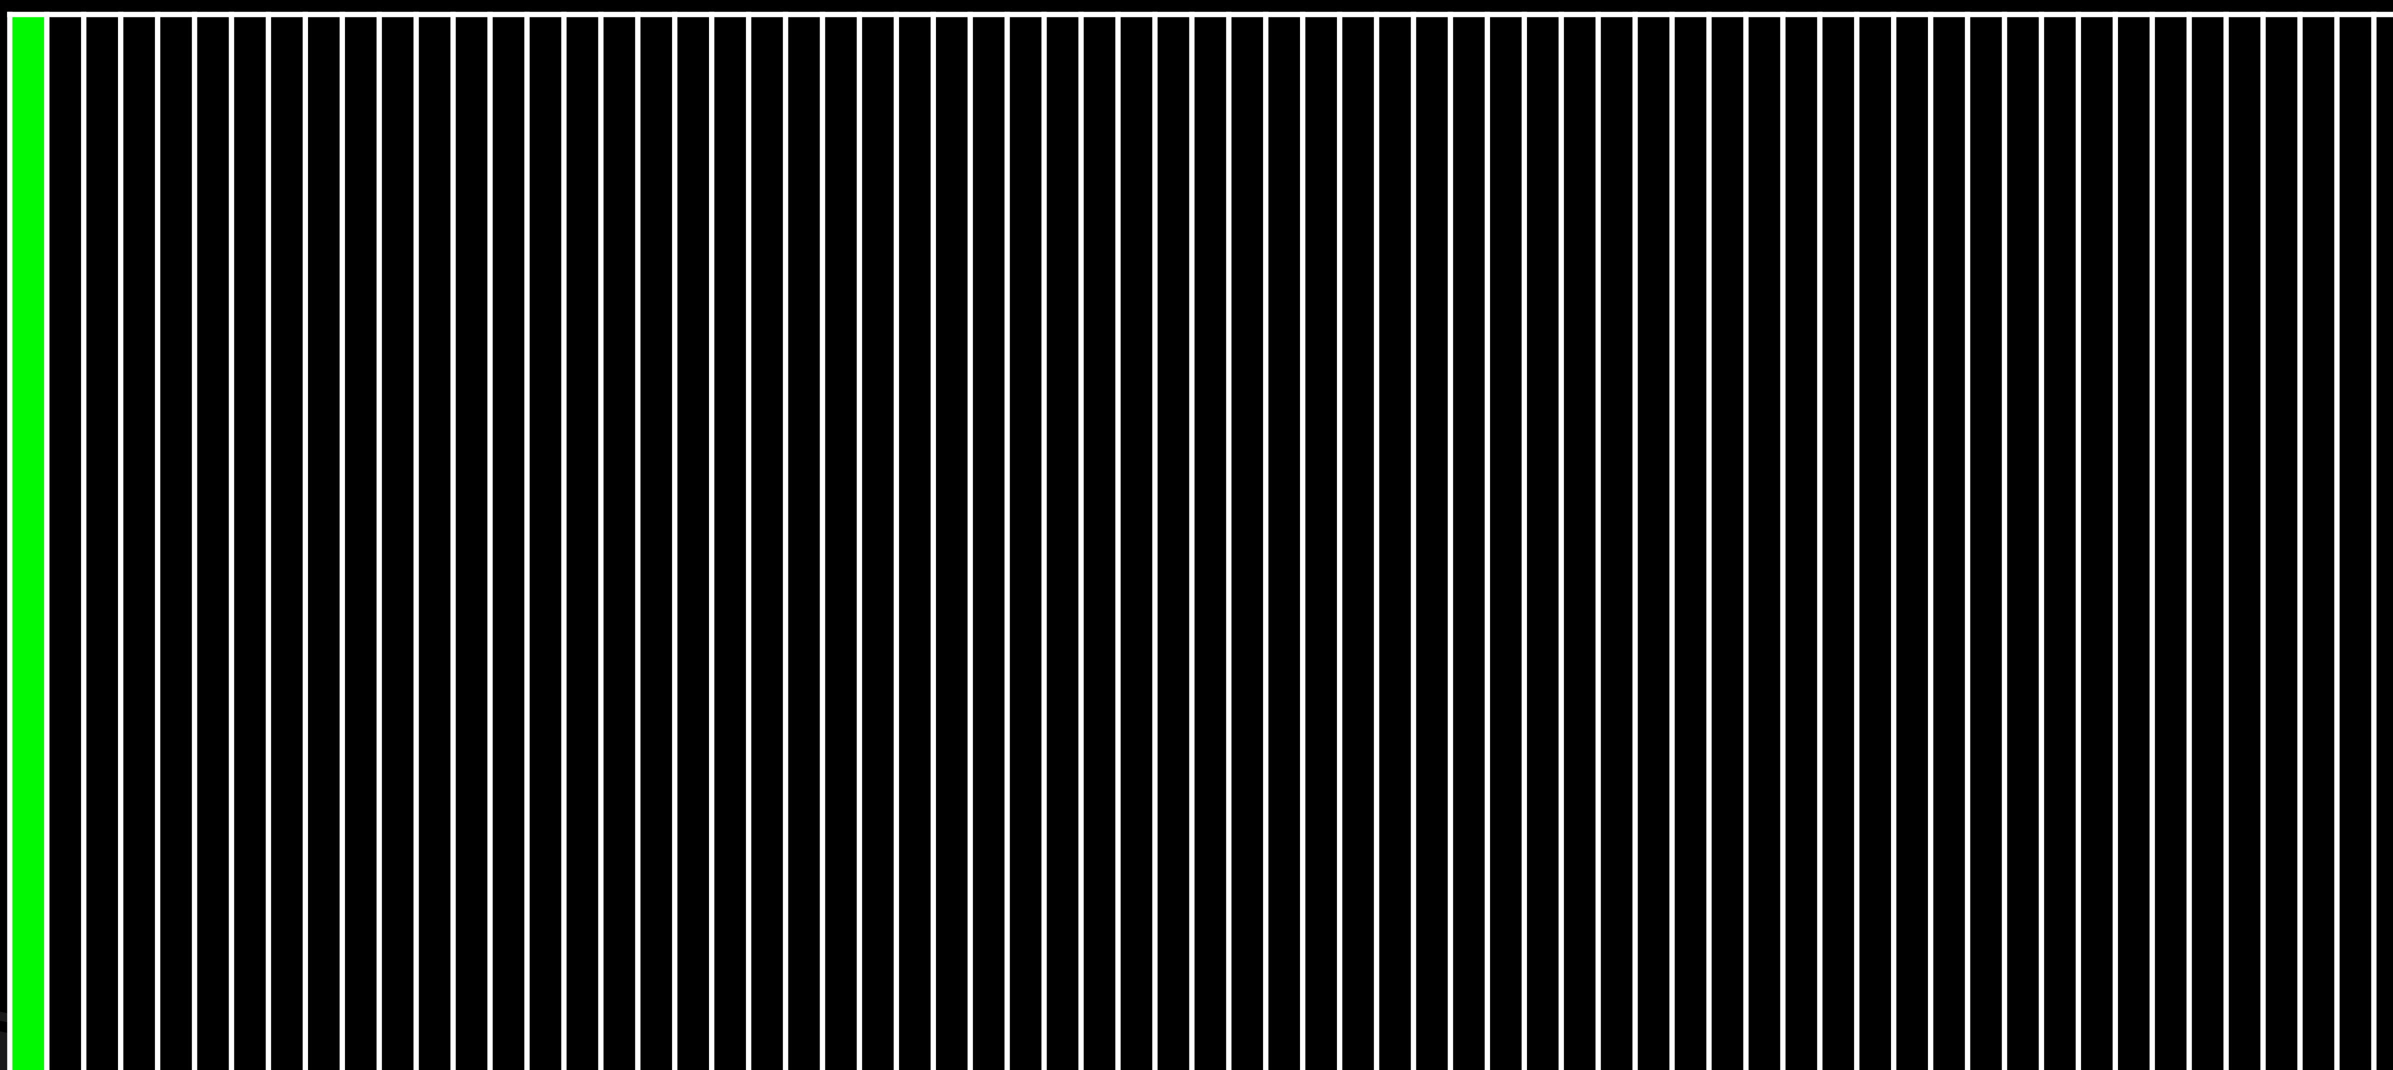

Average 0.01 seconds to completion

Flawed Logic:

10% of our traffic can be offloaded at a local Internet exchange. Therefore we need a circuit to the exchange that's one tenth as large as our transit circuit.

8bps utilization of 10bps and 100bps circuits, 80% full, 1 second

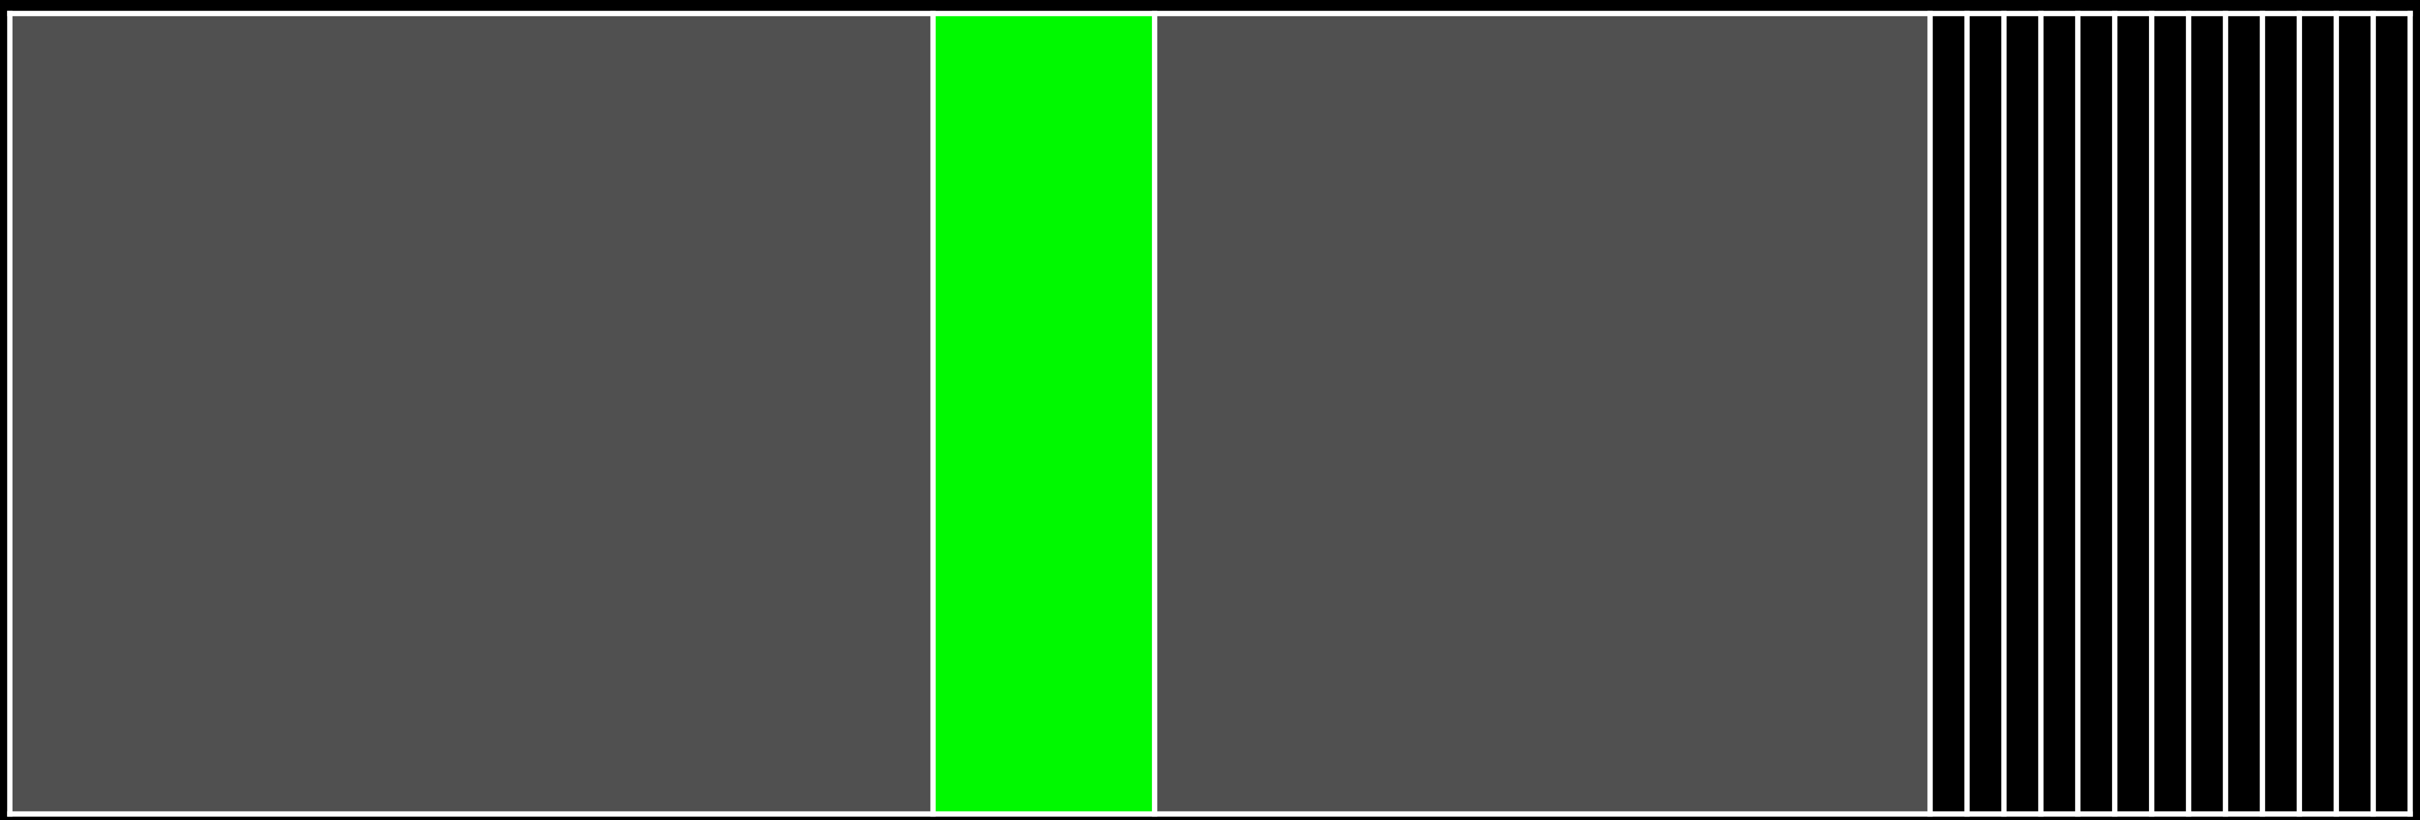

Average 0.48 seconds to completion

Average 0.8 seconds to completion

This Discourages Users

Users will always select services available over the largest-capacity circuit, not the least-utilized circuit, because that choice minimizes their wait-to-completion.

Packet Clearing House

The Lesson to ISPs:

Be sure that your largest circuit corresponds with your lowest cost path.

If 10% of the traffic in a well-engineered network is going to an IXP, the circuit it flows over will be more than 90% empty.

That's not a problem.

IXPs Improve Quality

...but only if networks are engineered to align users' incentives with operators' costs.

Packet Clearing House

Effect of an IX on Transit Purchase

Sometimes people assume that the introduction of low-cost bandwidth from a local IX is a literal substitute for high-cost bandwidth from existing transit providers, and that this will result in a reduction of total costs. This is not true, and stems from trying to view ISP economics in terms other than APBDC and exponential growth.

Effect of an IX on Transit Purchase

In fact, transit contracts tend to be constrained to fixed terms, are not subject to cost-effective early cancellation, and are time- and labor-intensive to initiate.

Effect of an IX on Transit Purchase

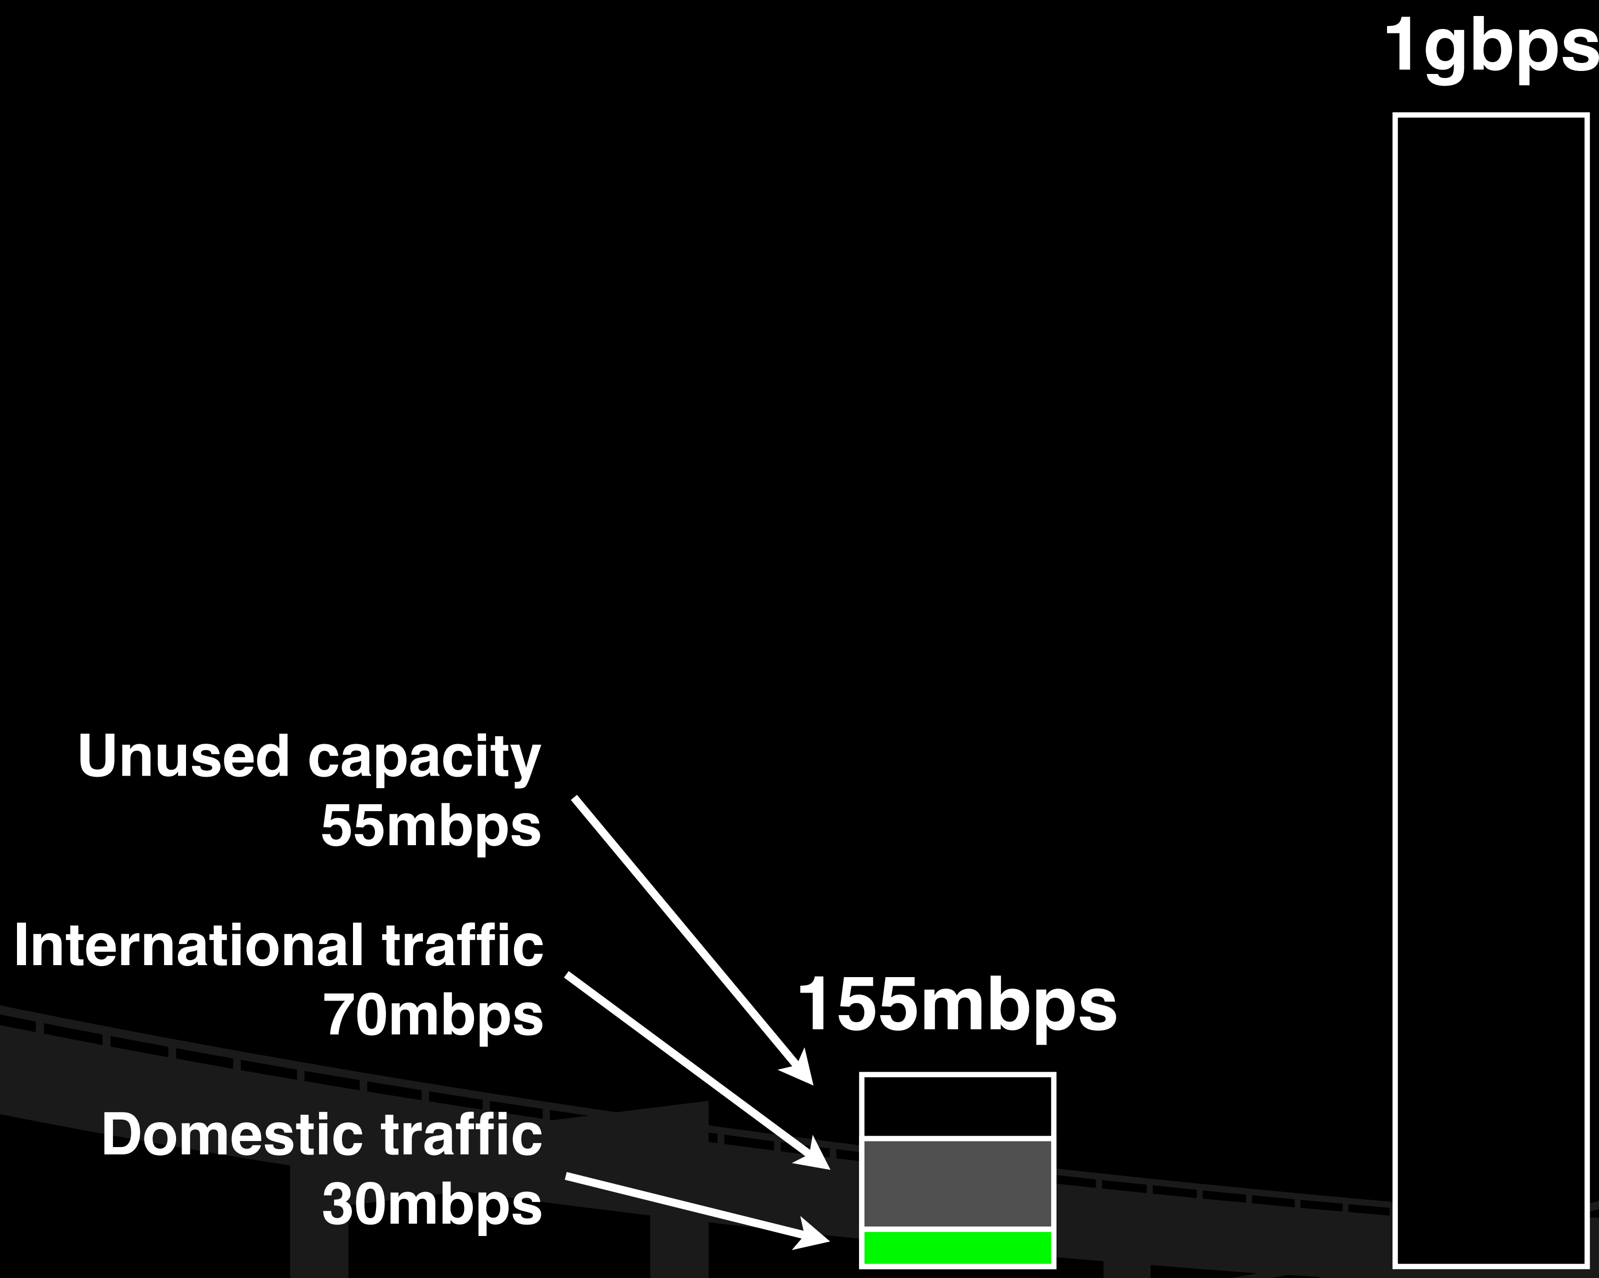

Transit Bandwidth Prior to Domestic Peering

Effect of an IX on Transit Purchase

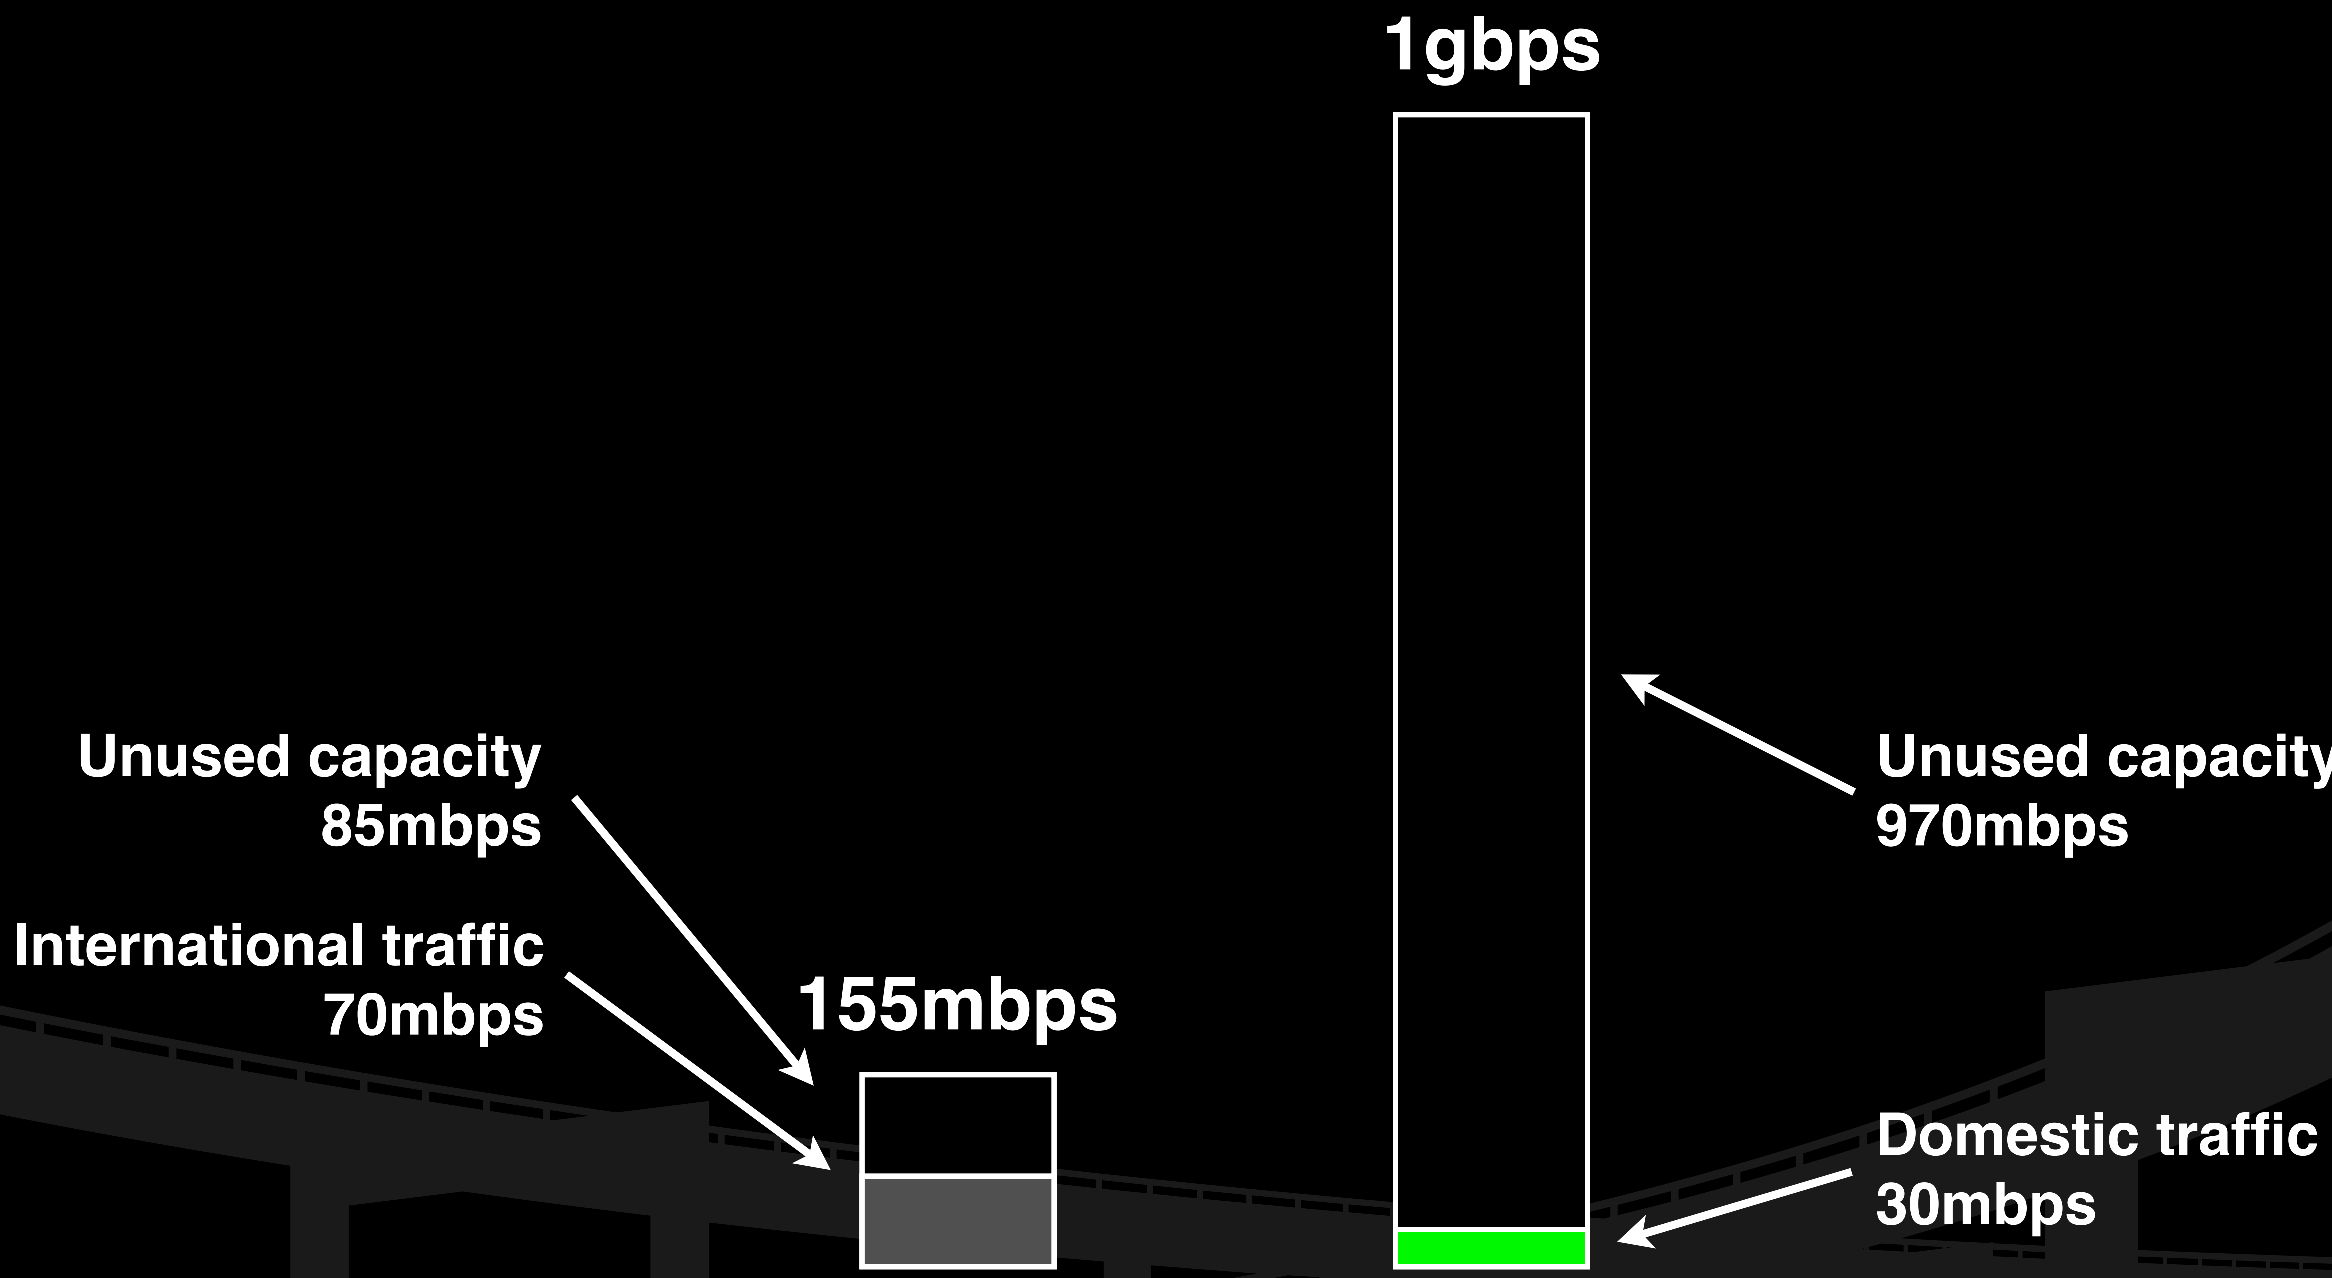

Transit Bandwidth After Domestic Peering

PCH
Packet Clearing House

The Virtuous Cycle of Upgrades

The Internet has doubled in size every ten and a half months for the past thirty years. Keeping up with this exponential growth is a process of addressing each revealed bottleneck and moving on to the next in a continuous virtuous cycle of upgrades, eventually returning to each bottleneck many times.

Packet Clearing House

The Virtuous Cycle of Upgrades

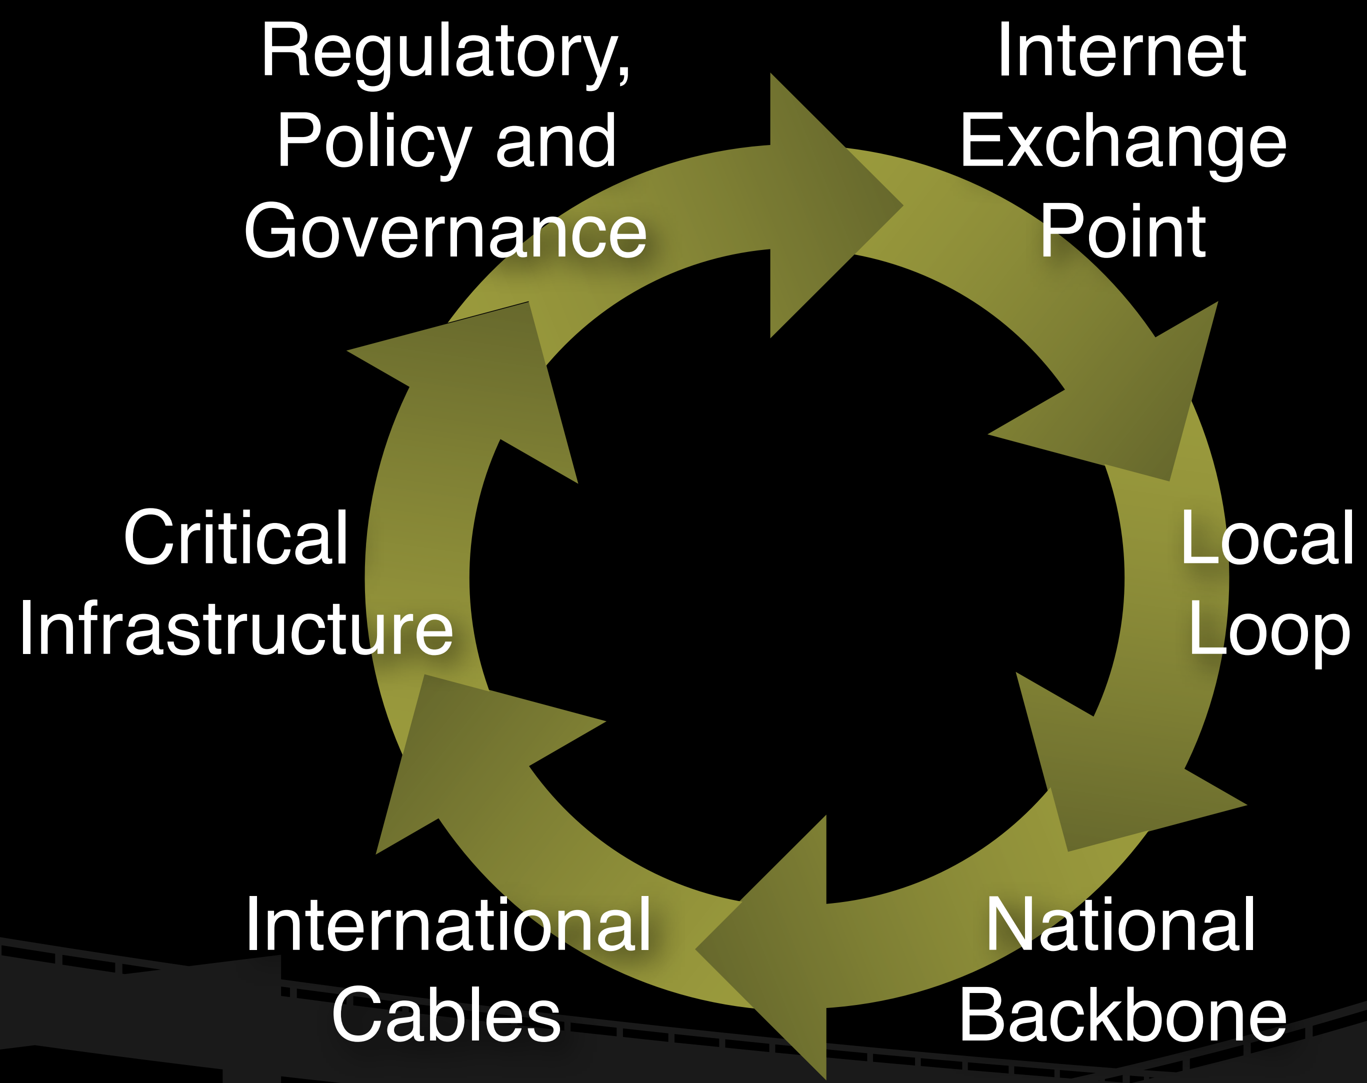

The Virtuous Cycle of Upgrades

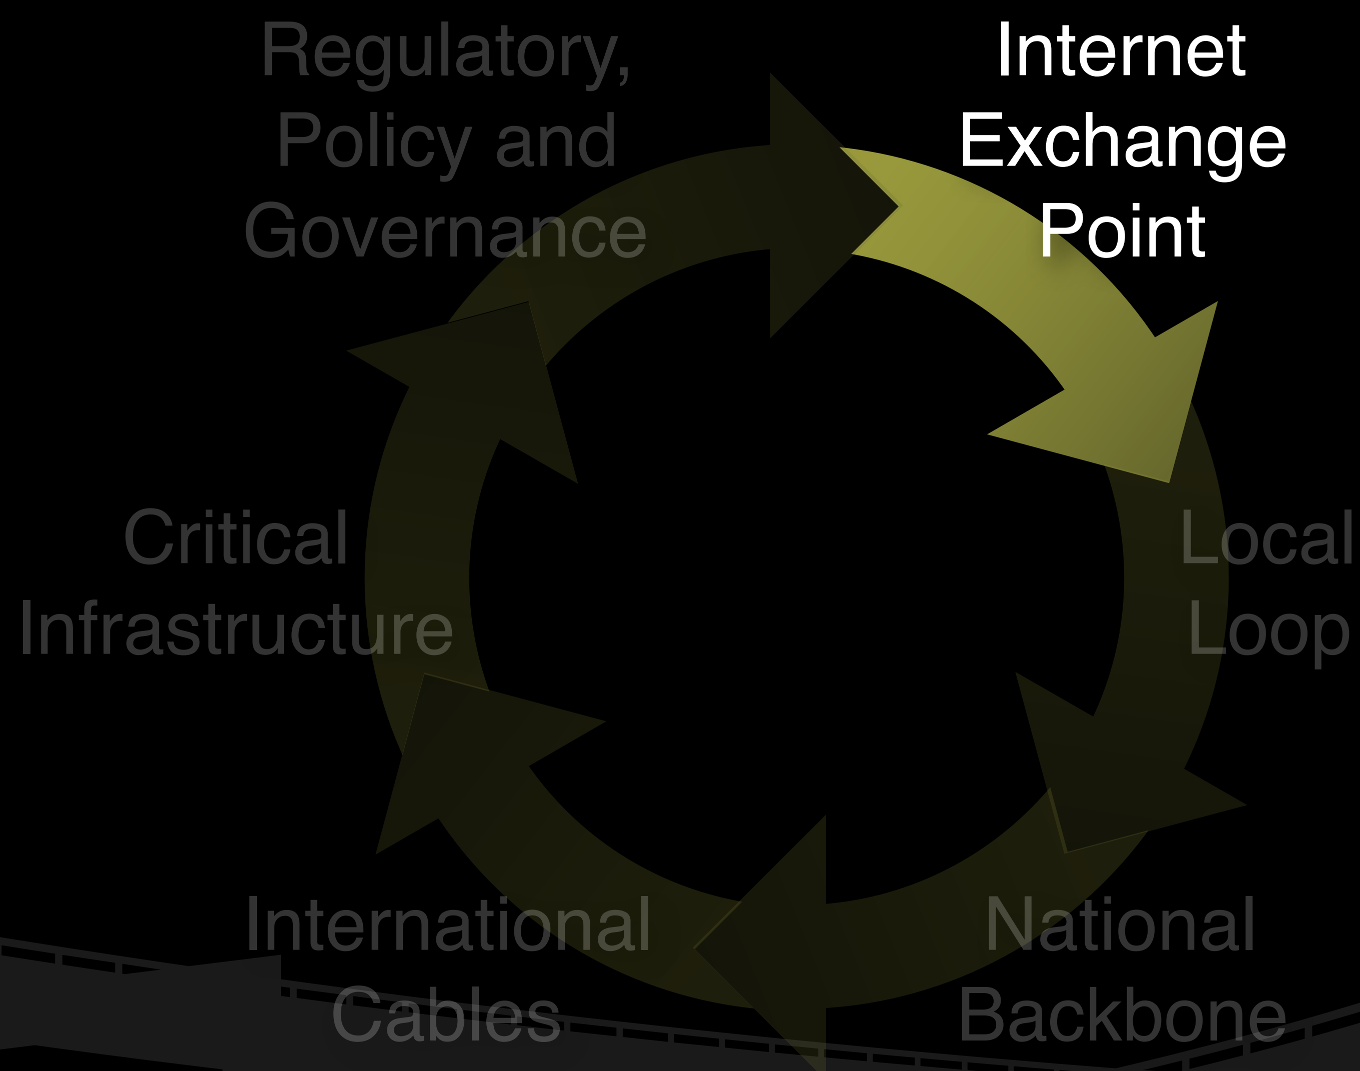

IXPs provide a domestic source of Internet bandwidth that is not dependent on International transit or transport.

The Virtuous Cycle of Upgrades

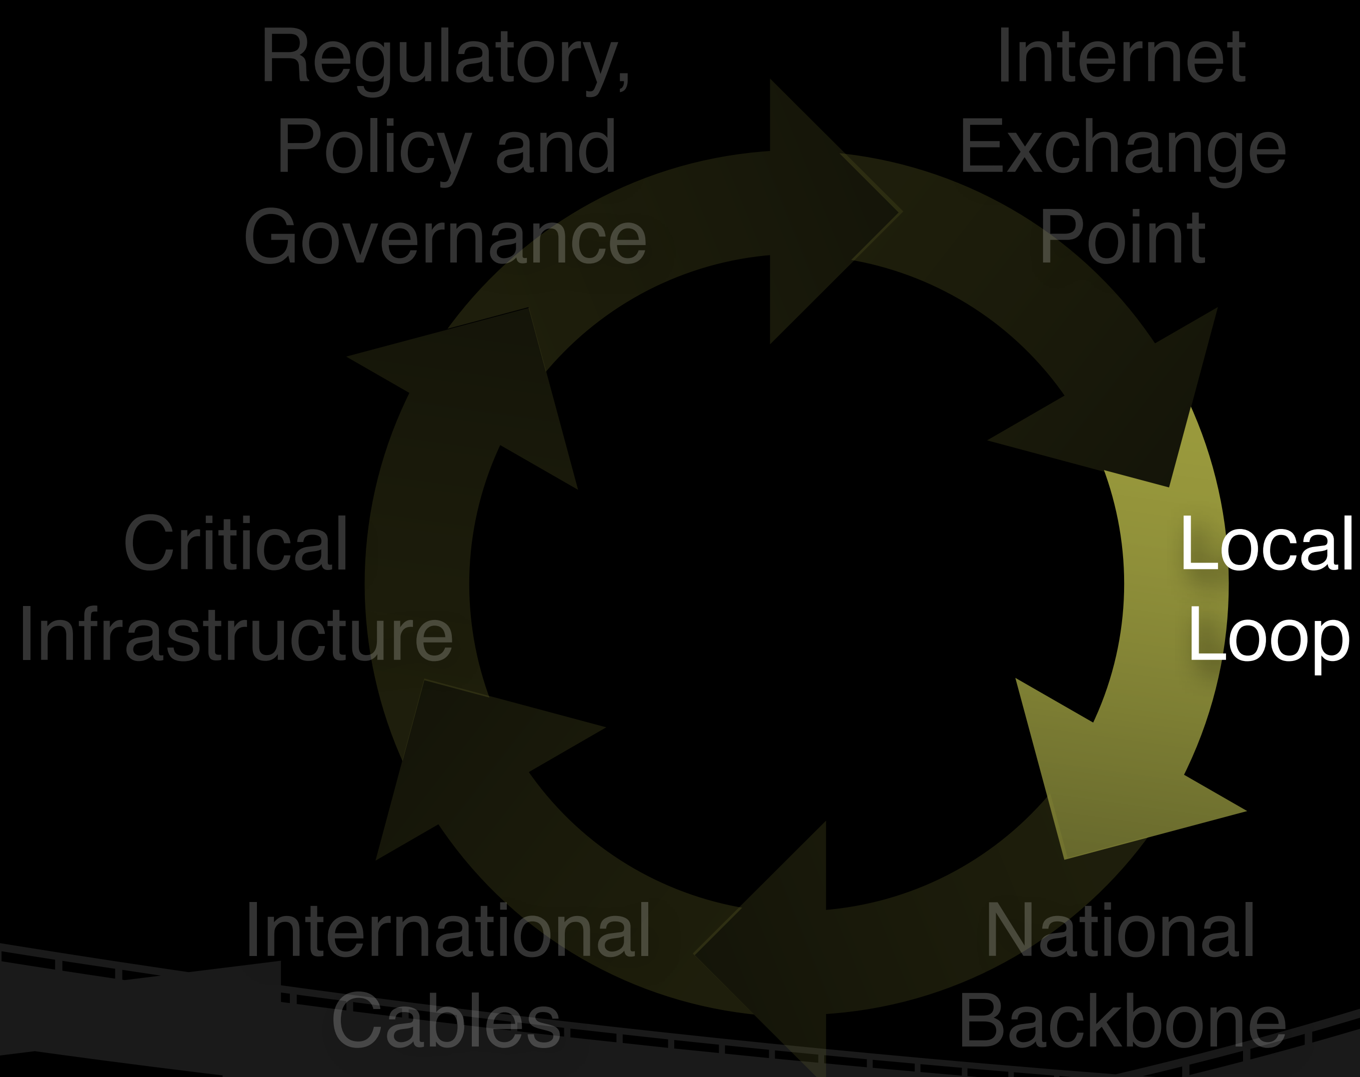

Local loops (or "last miles" of copper, fiber, or wireless must be available to connect networks to customers and resources like IXPs, critical infrastructure, and international cables.

Packet Clearing House

The Virtuous Cycle of Upgrades

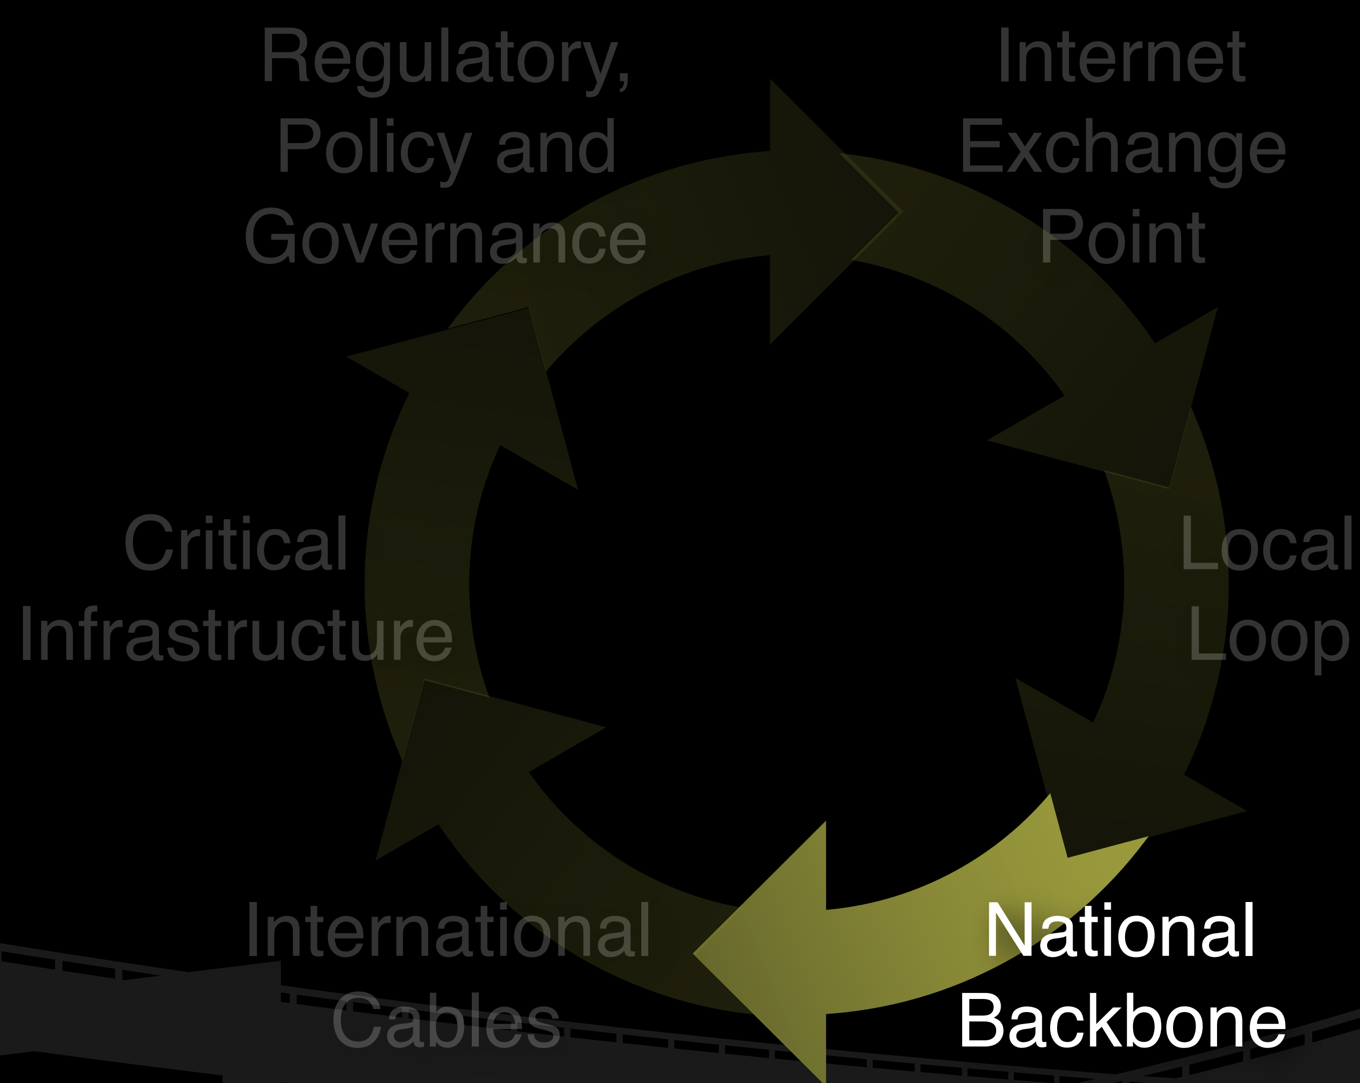

Long-haul backbone fiber must be available to interconnect IXPs in one city with customers in another. They form the web of circuits between a country's cities.

The Virtuous Cycle of Upgrades

International cable systems must be available for domestic ISPs to bring traffic in from foreign IXPs, and for foreign ISPs to receive traffic from domestic IXPs.

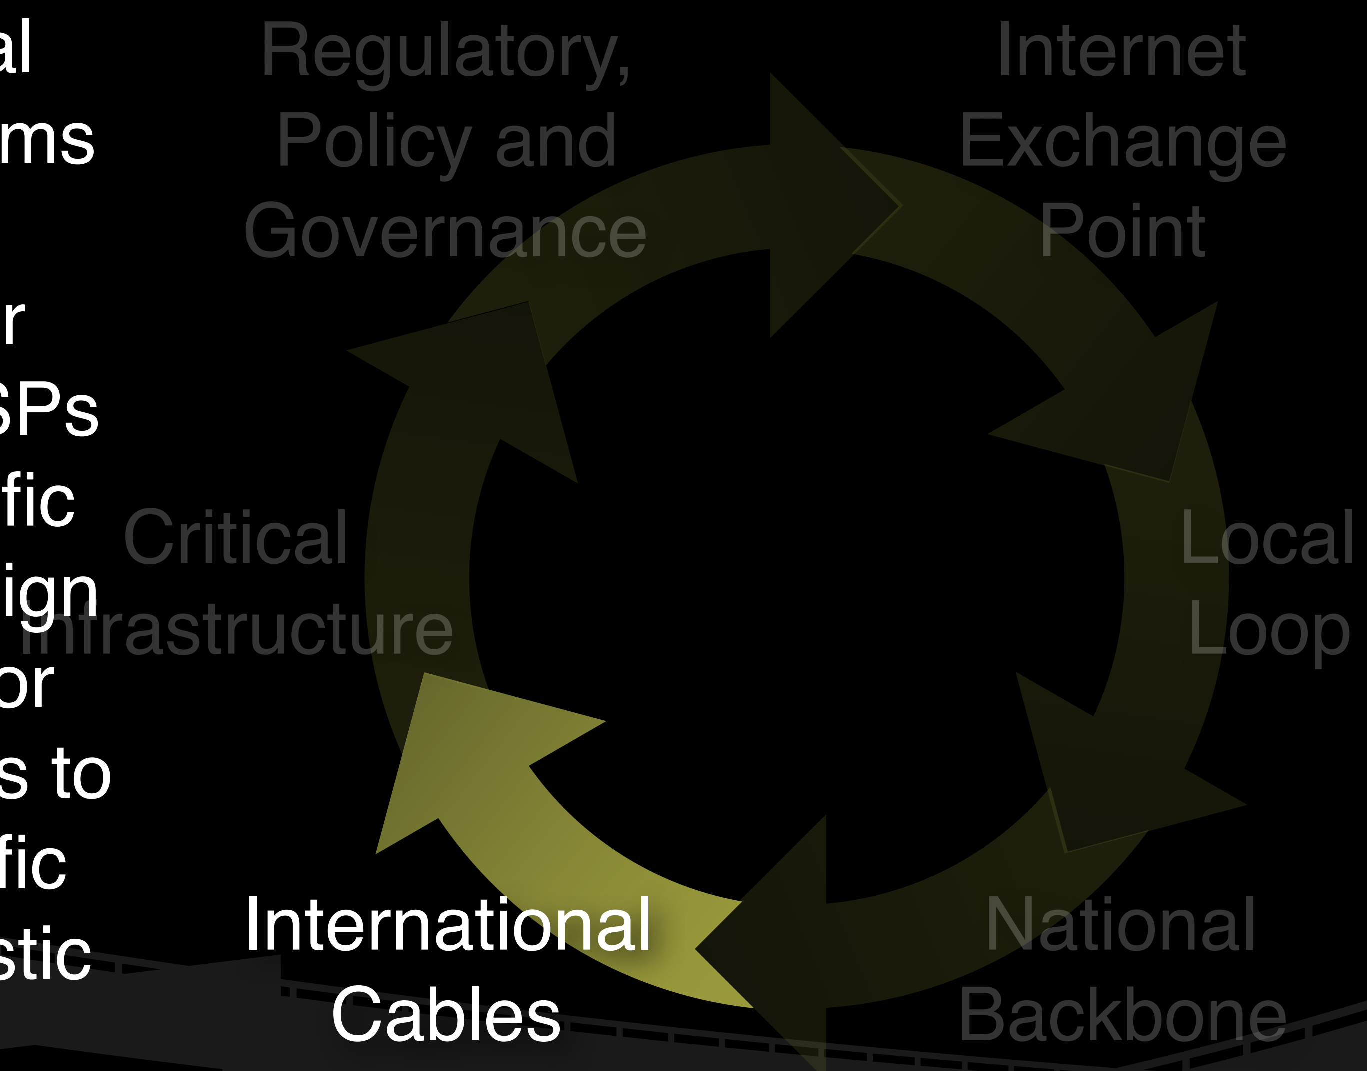

The Virtuous Cycle of Upgrades

Critical infrastructure like root and ccTLD nameservers, CERTs, and law enforcement intercept must be speedy enough to keep up with the rest of the network.

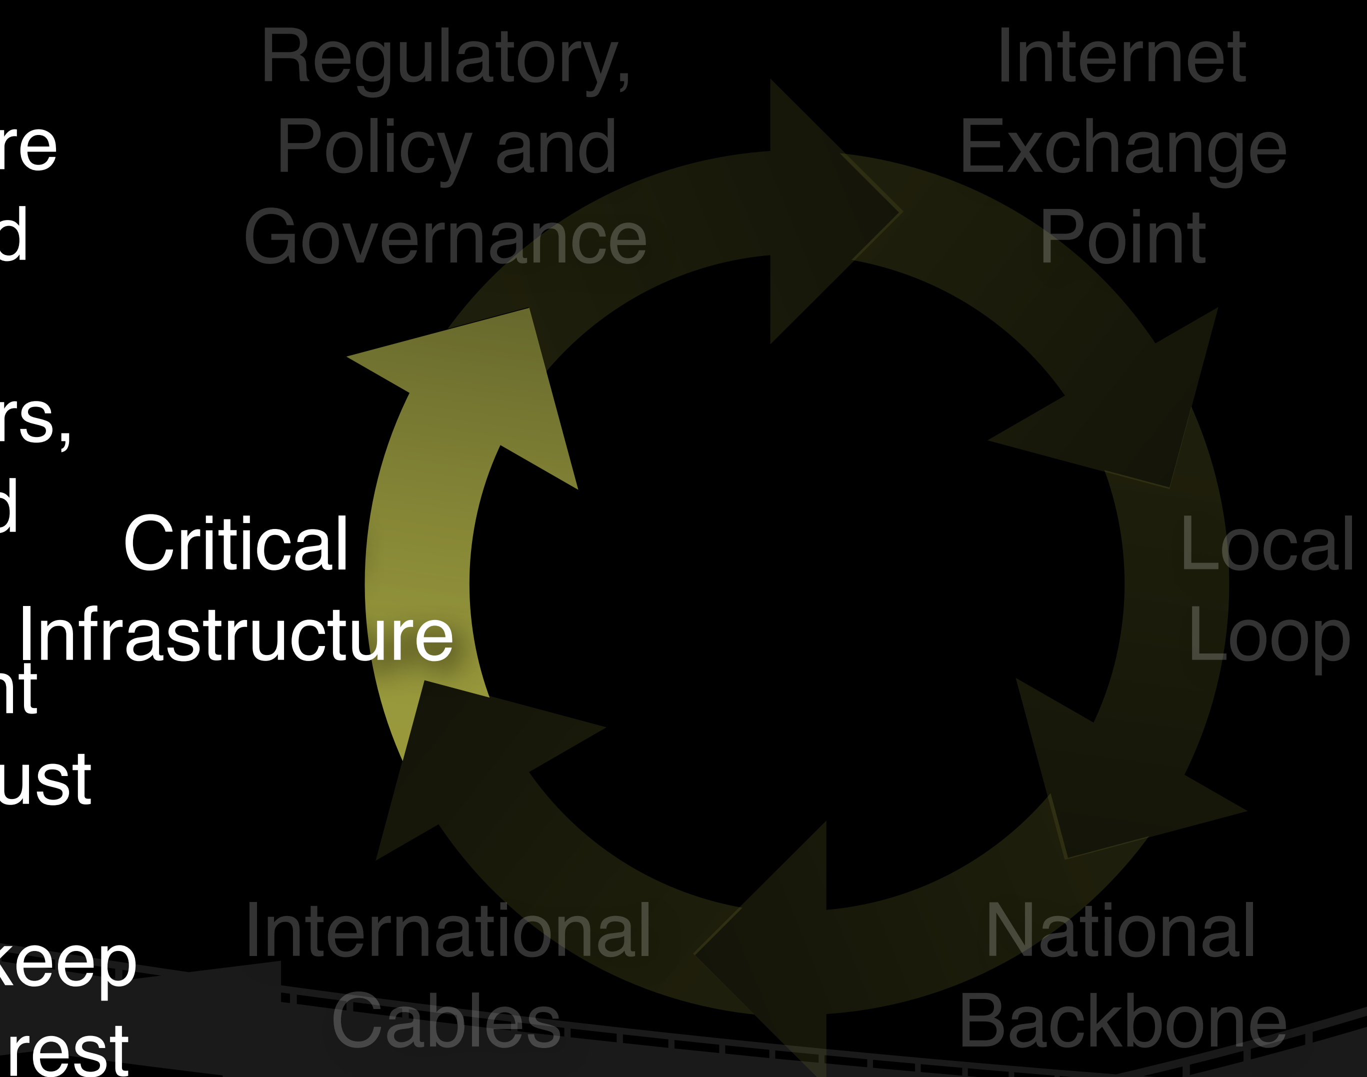

Packet Clearing House

The Virtuous Cycle of Upgrades

The public policy environment must support and encourage new market entrants and competition among existing players to ensure continuous price/performance improvement.

Packet Clearing House

The Virtuous Cycle of Upgrades

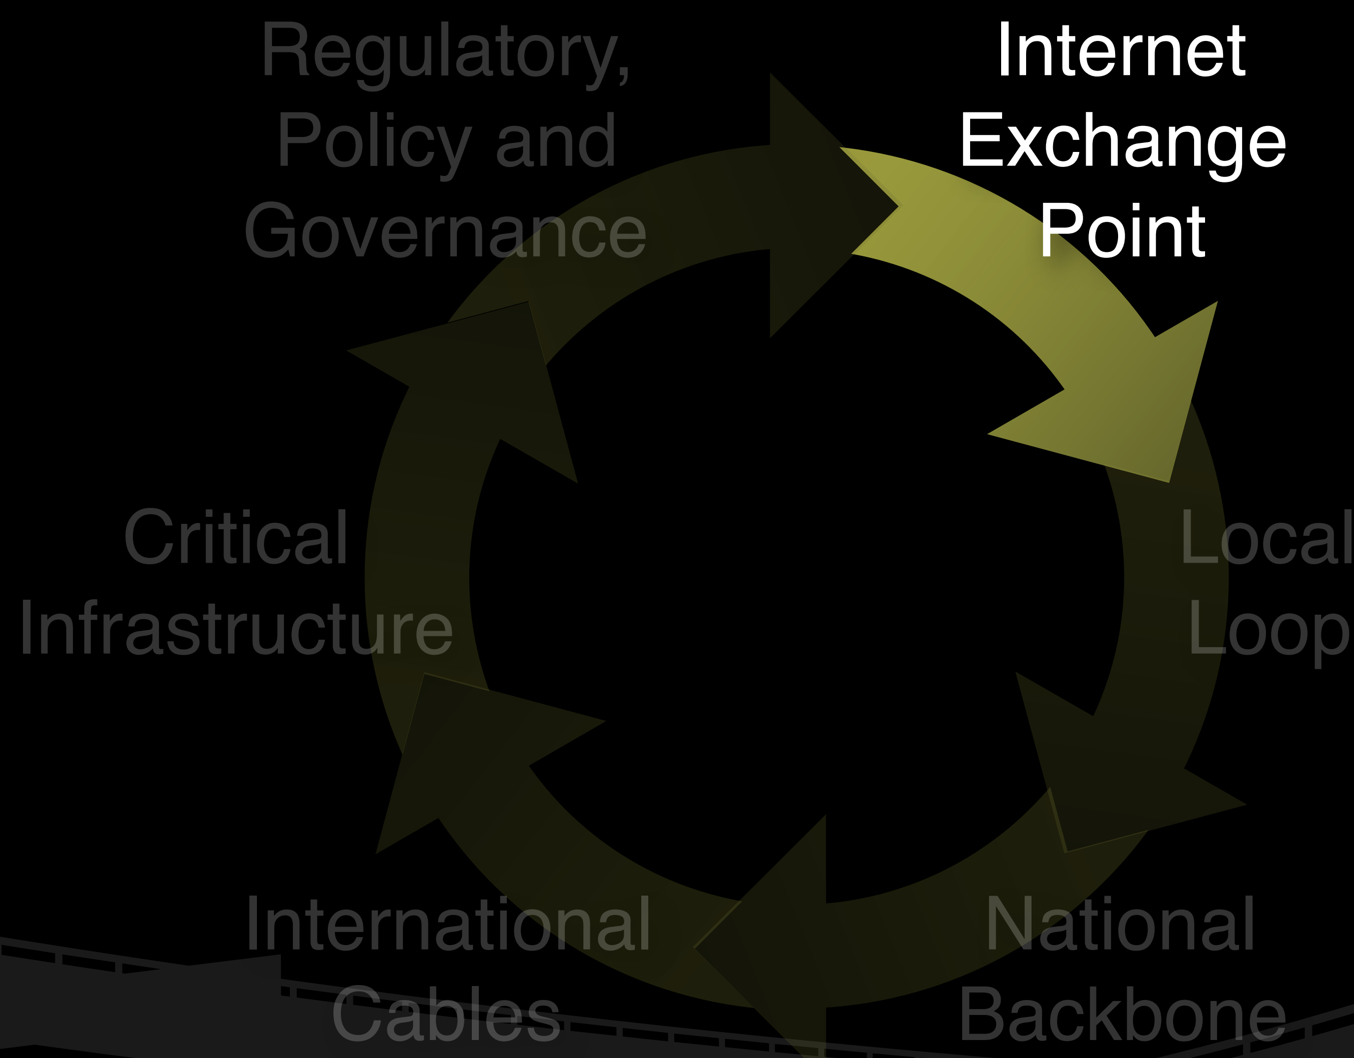

...by which time we're ready to upgrade our IXP again.

And so on, around the circle.

Growth of IXP Physical Plant

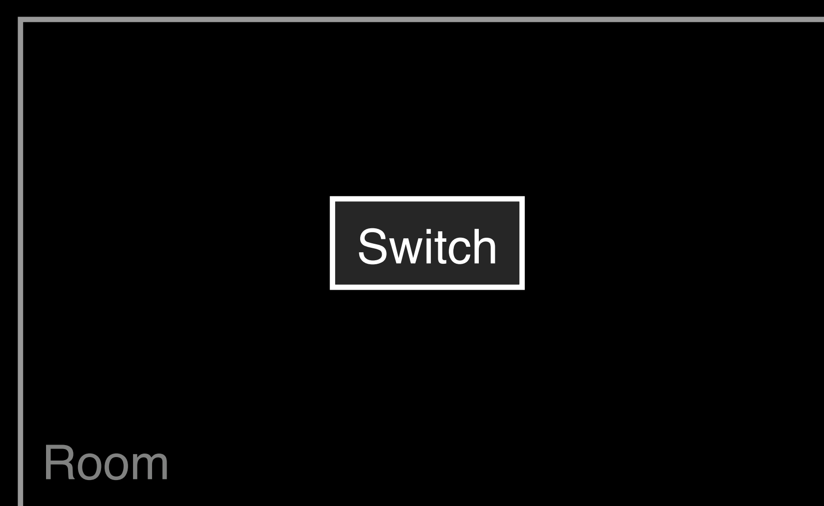

Growth of IXP Physical Plant

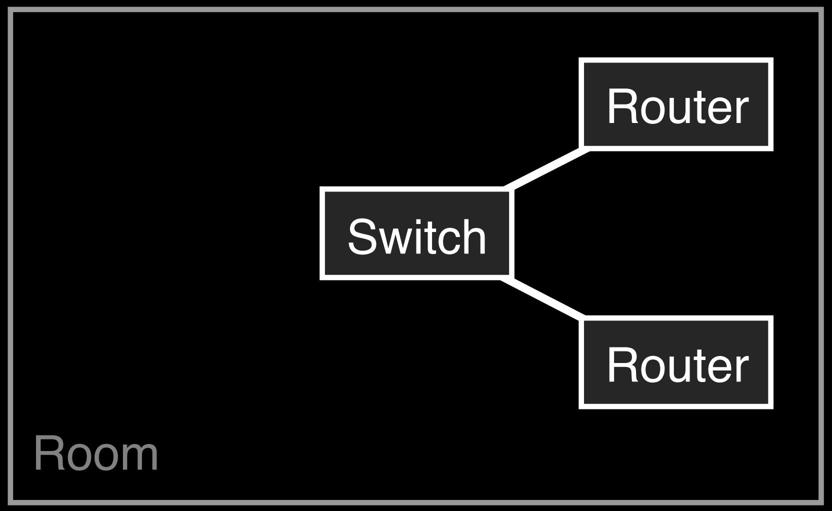

Growth of IXP Physical Plant

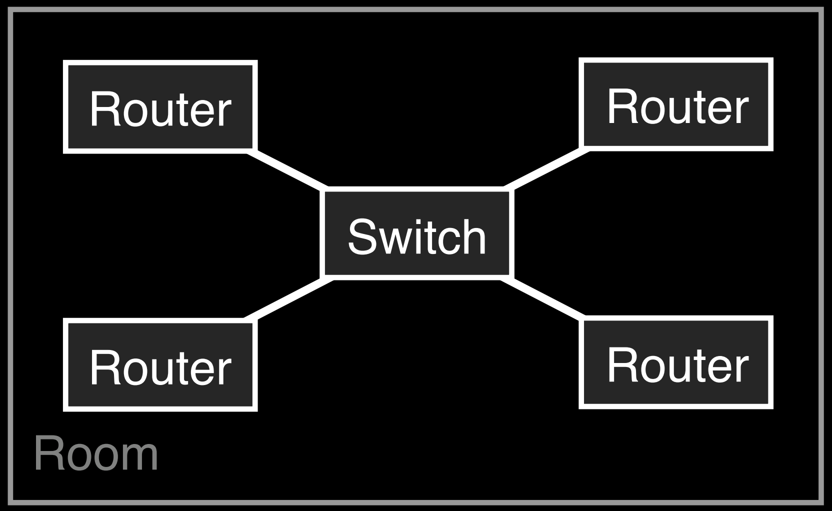

Packet Clearing House

Growth of IXP Physical Plant

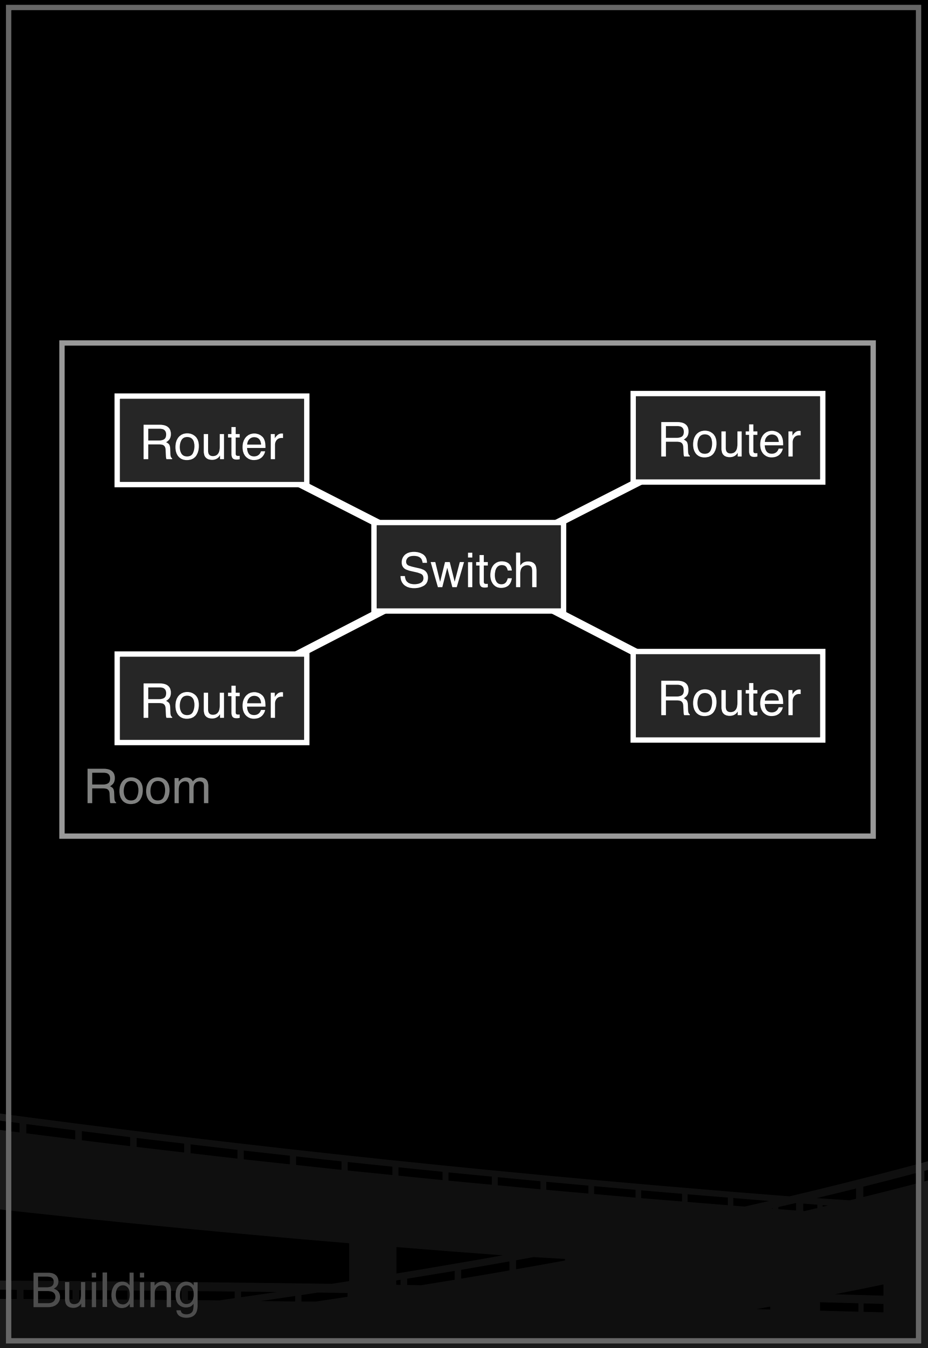

Packet Clearing House

Growth of IXP Physical Plant

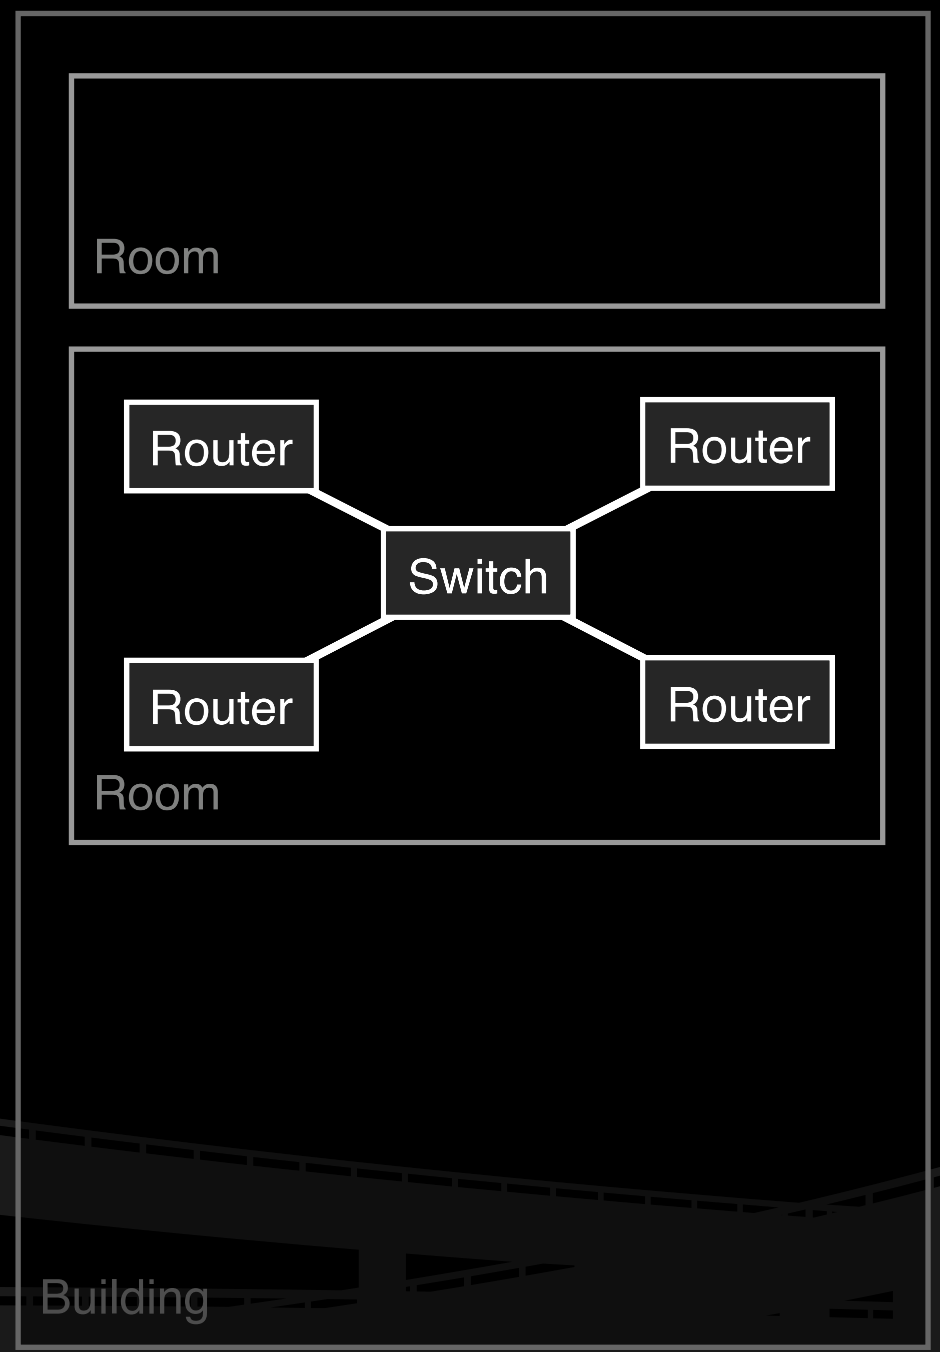

Growth of IXP Physical Plant

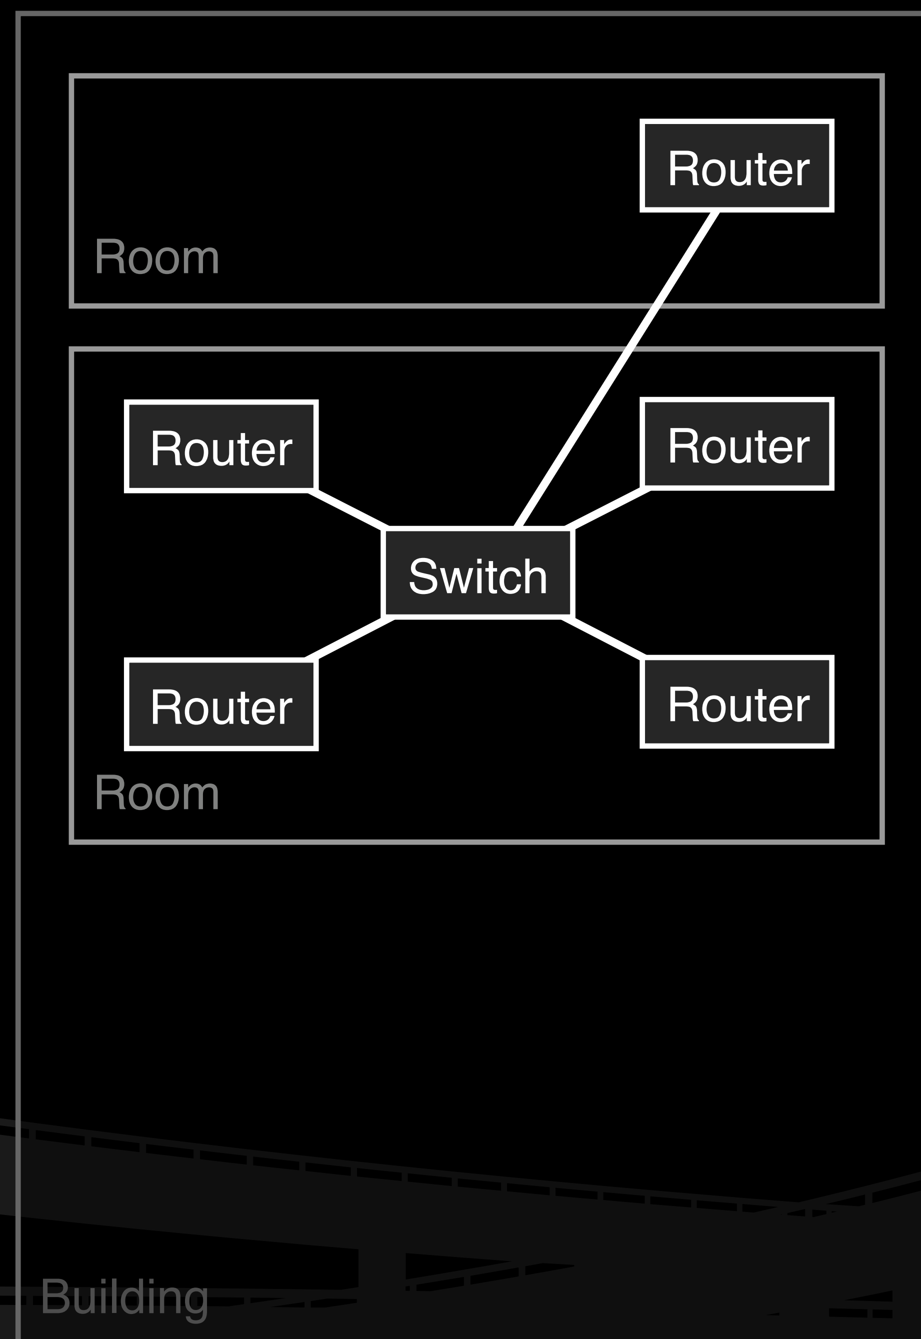

Growth of IXP Physical Plant

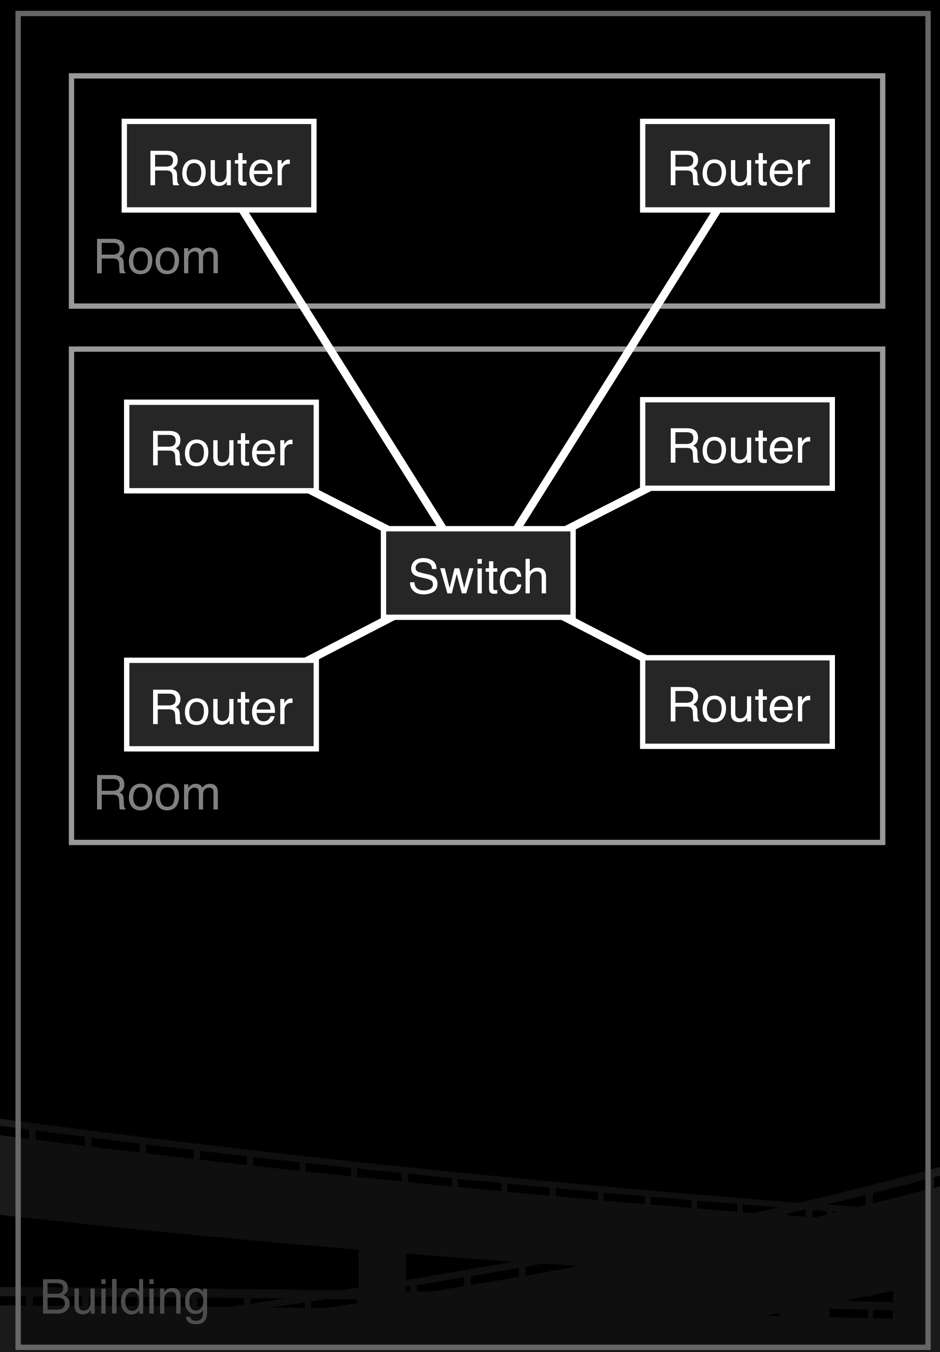

Packet Clearing House

Growth of IXP Physical Plant

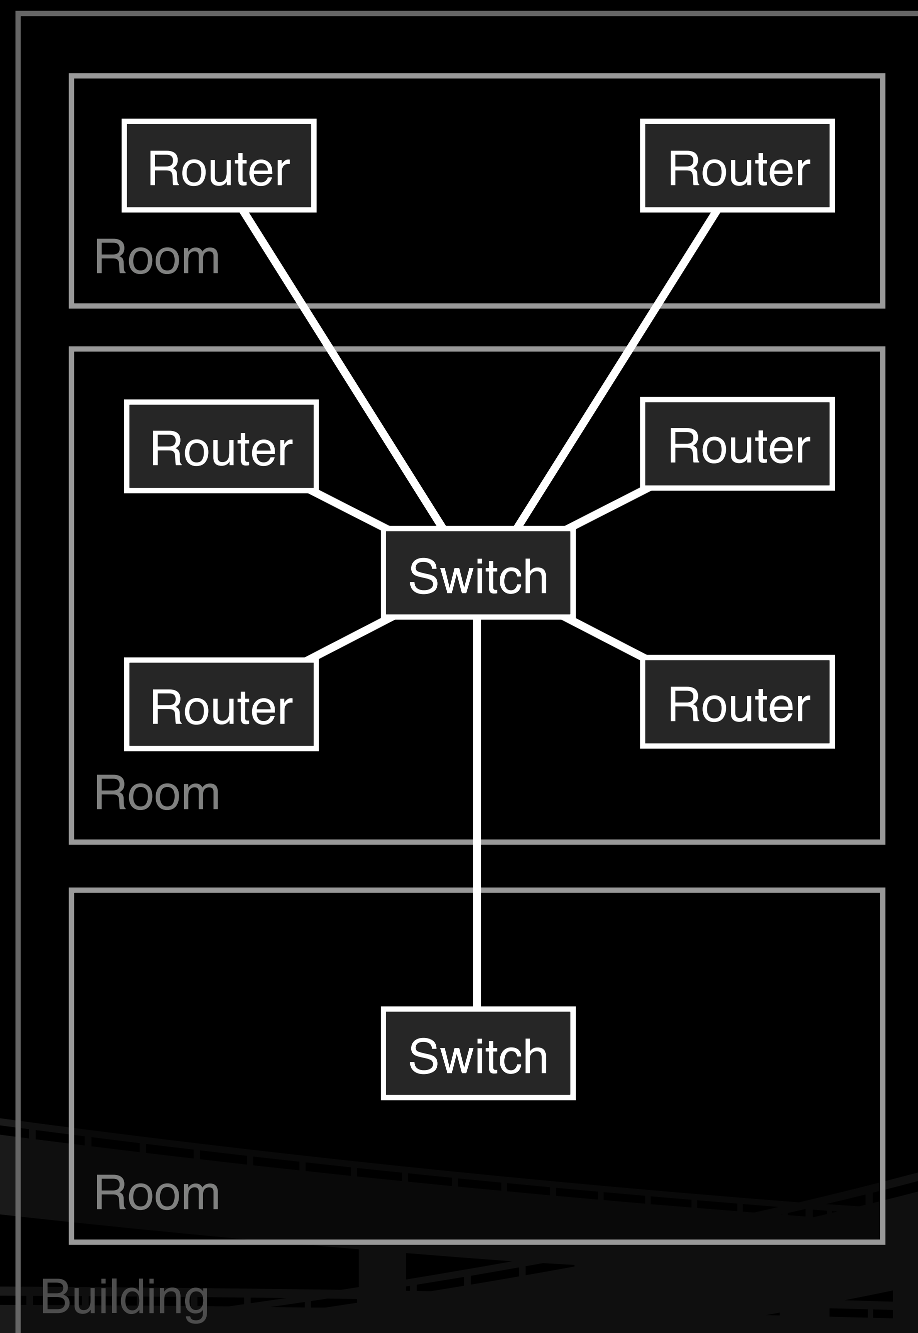

Growth of IXP Physical Plant

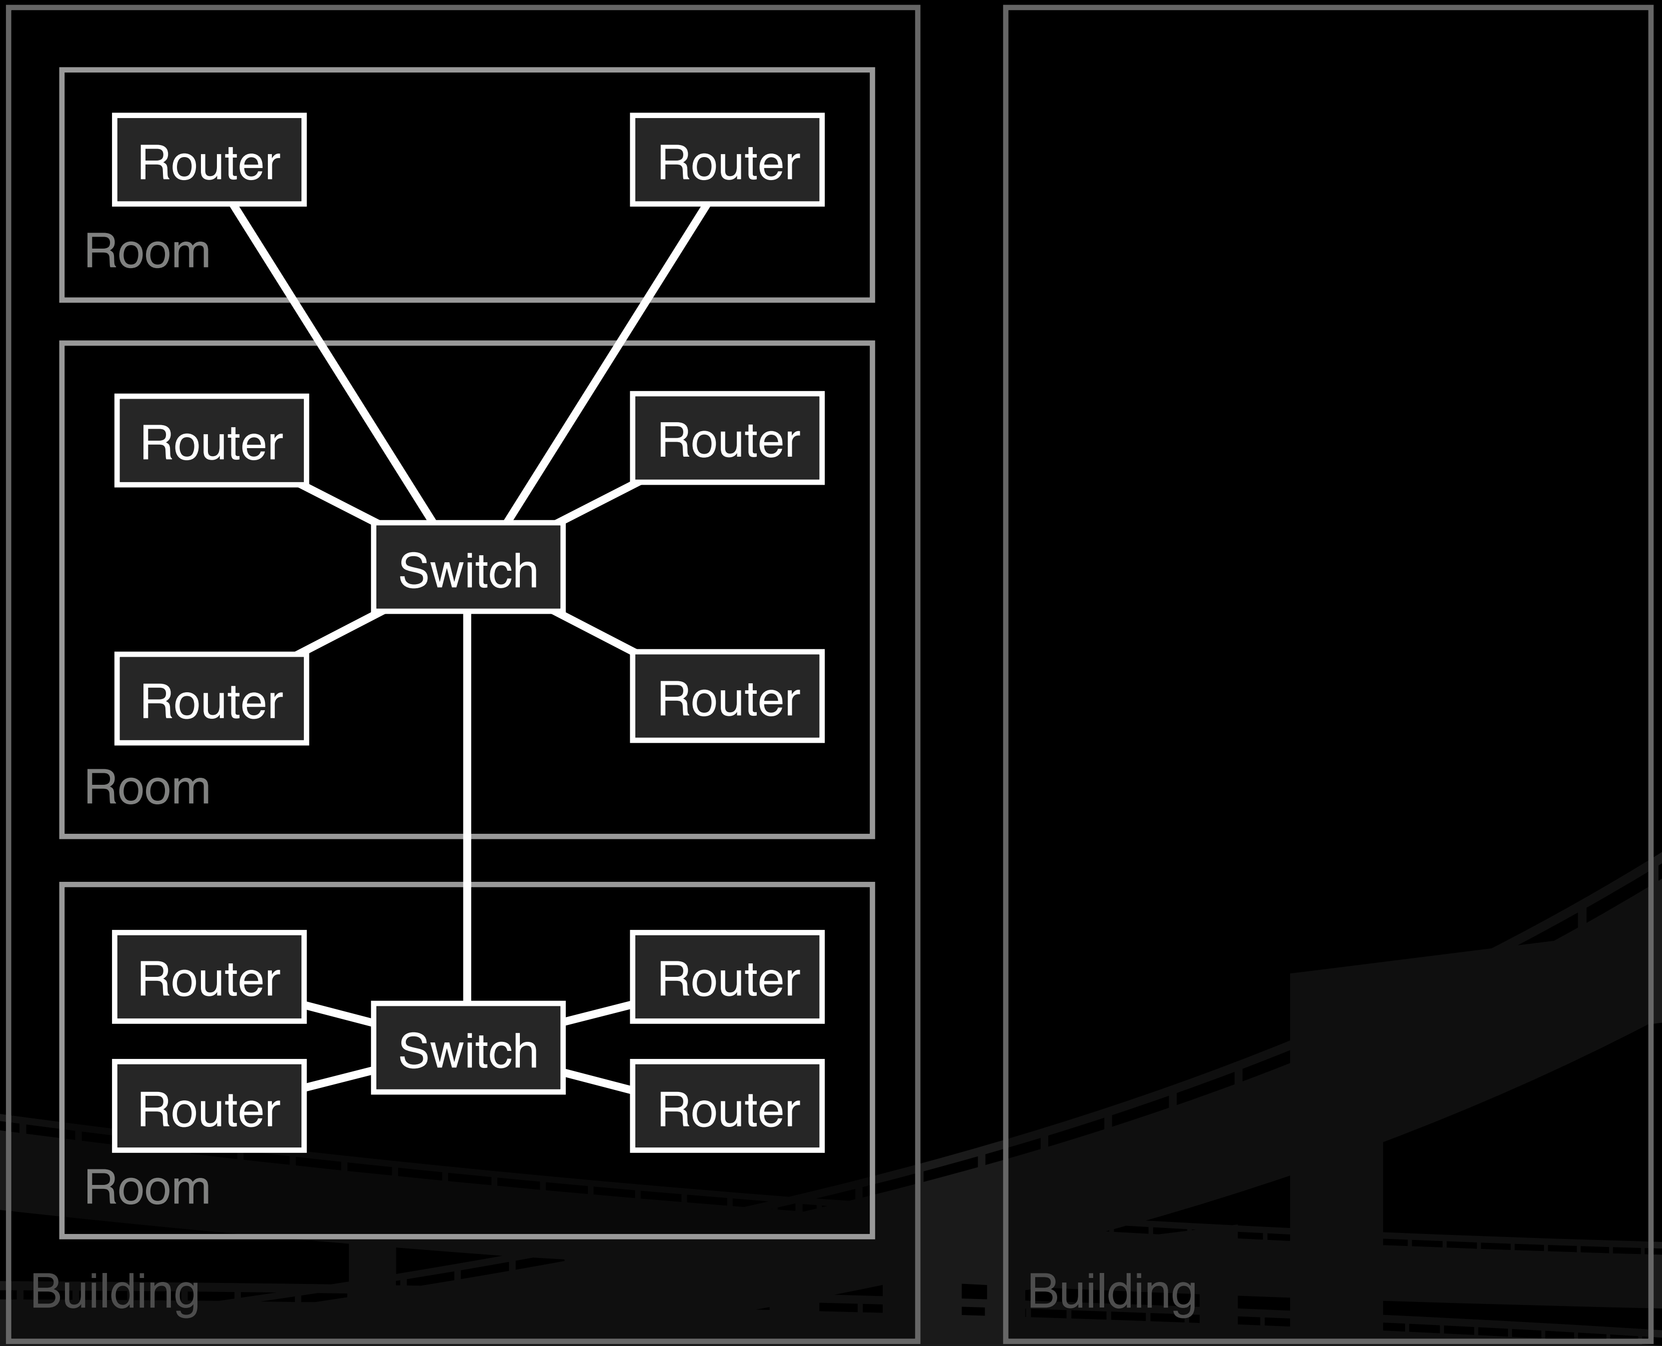

PCH Packet Clearing House

Growth of IXP Physical Plant

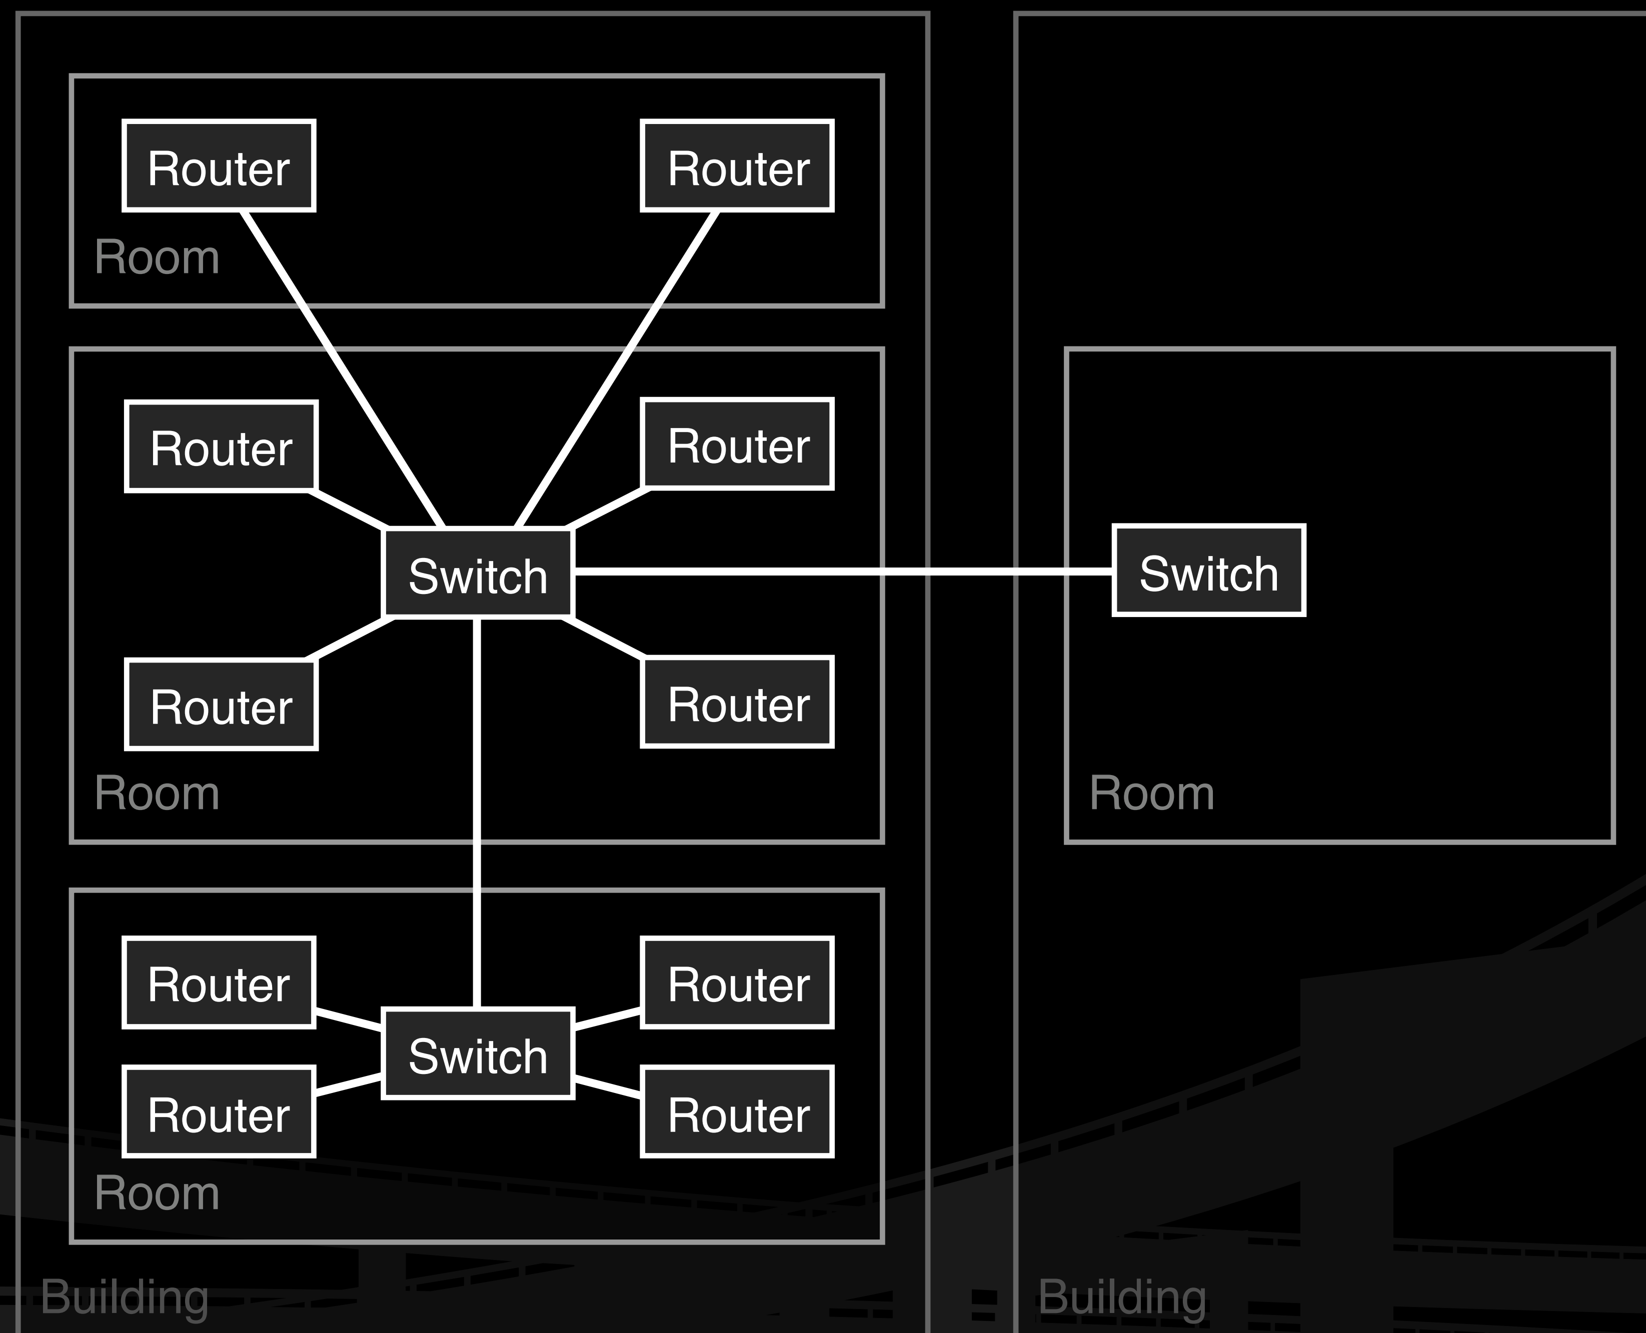

Growth of IXP Physical Plant

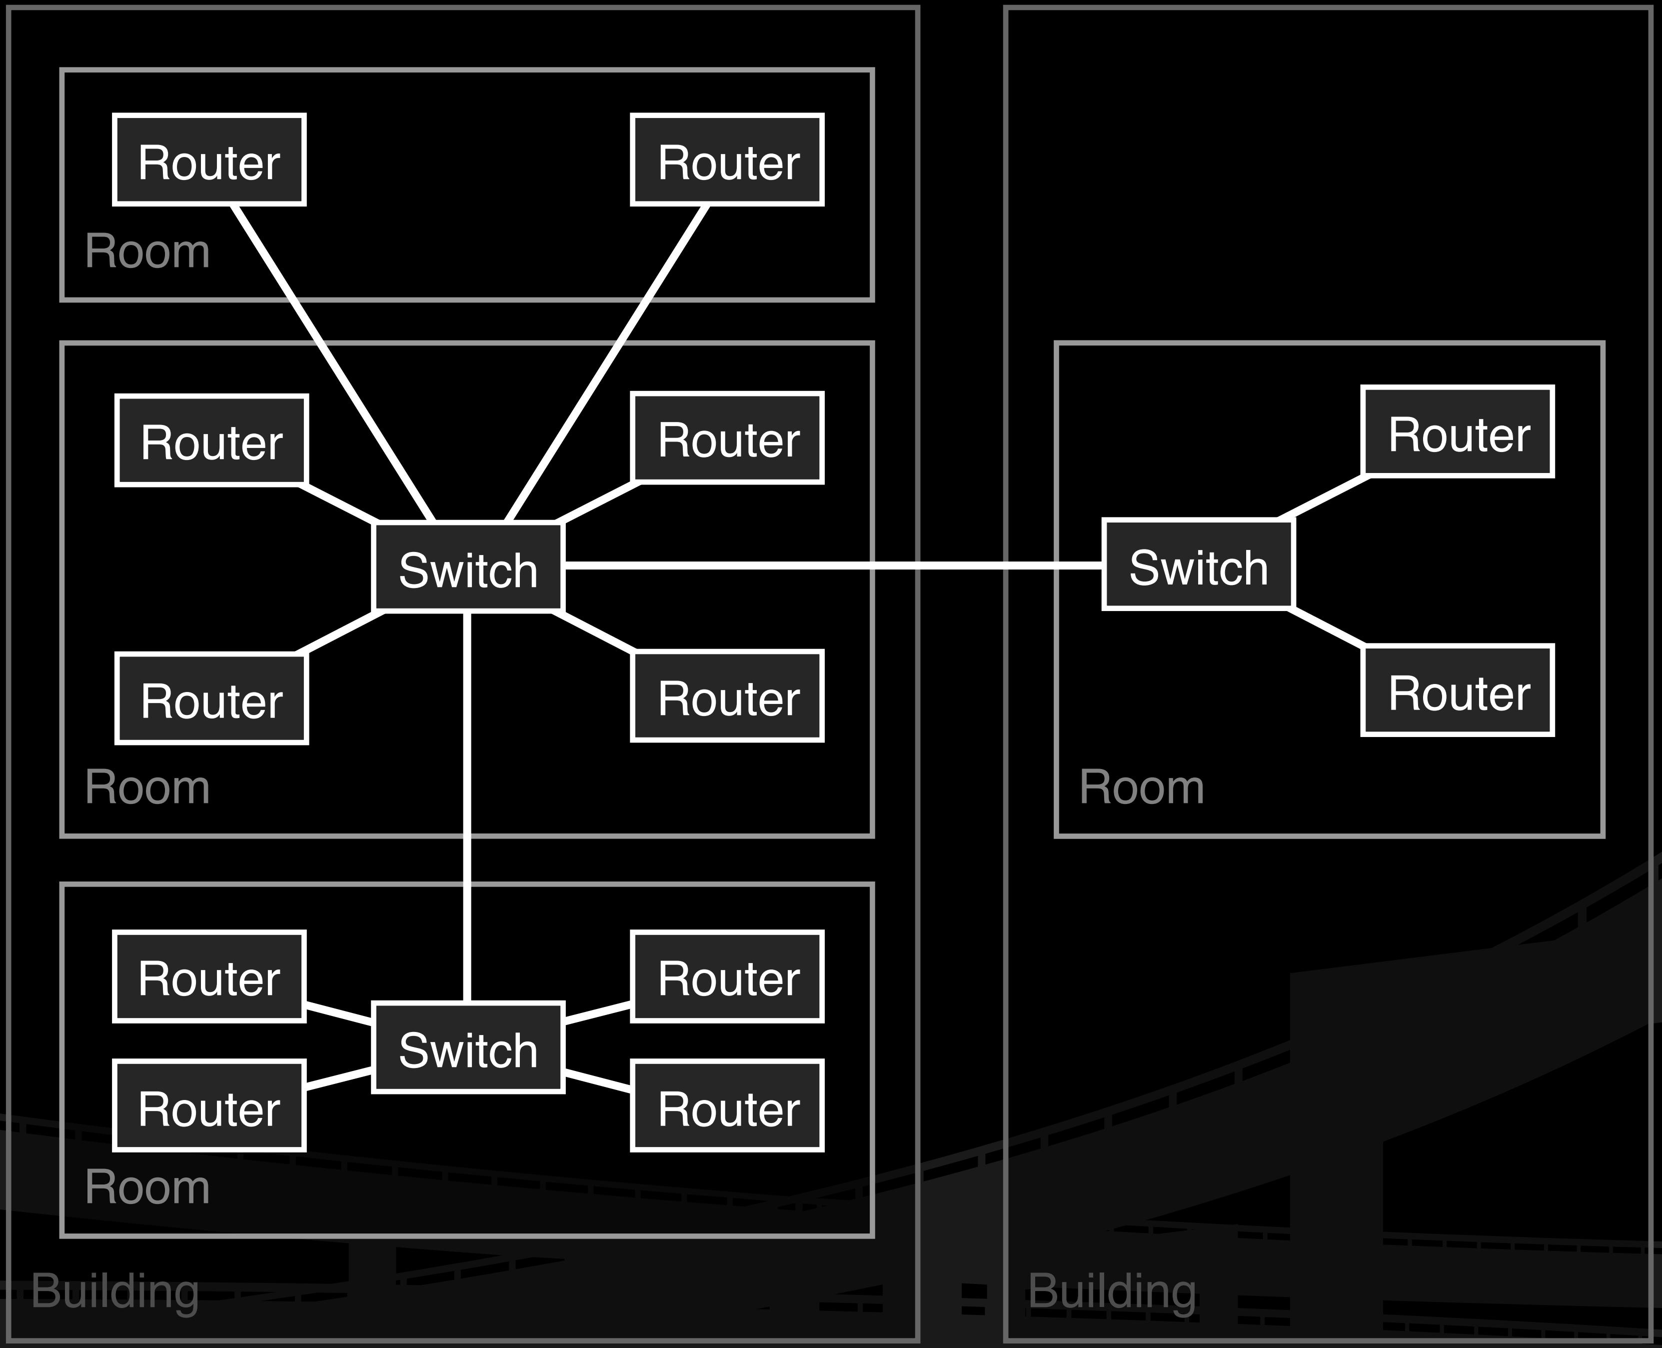

Growth of IXP Physical Plant

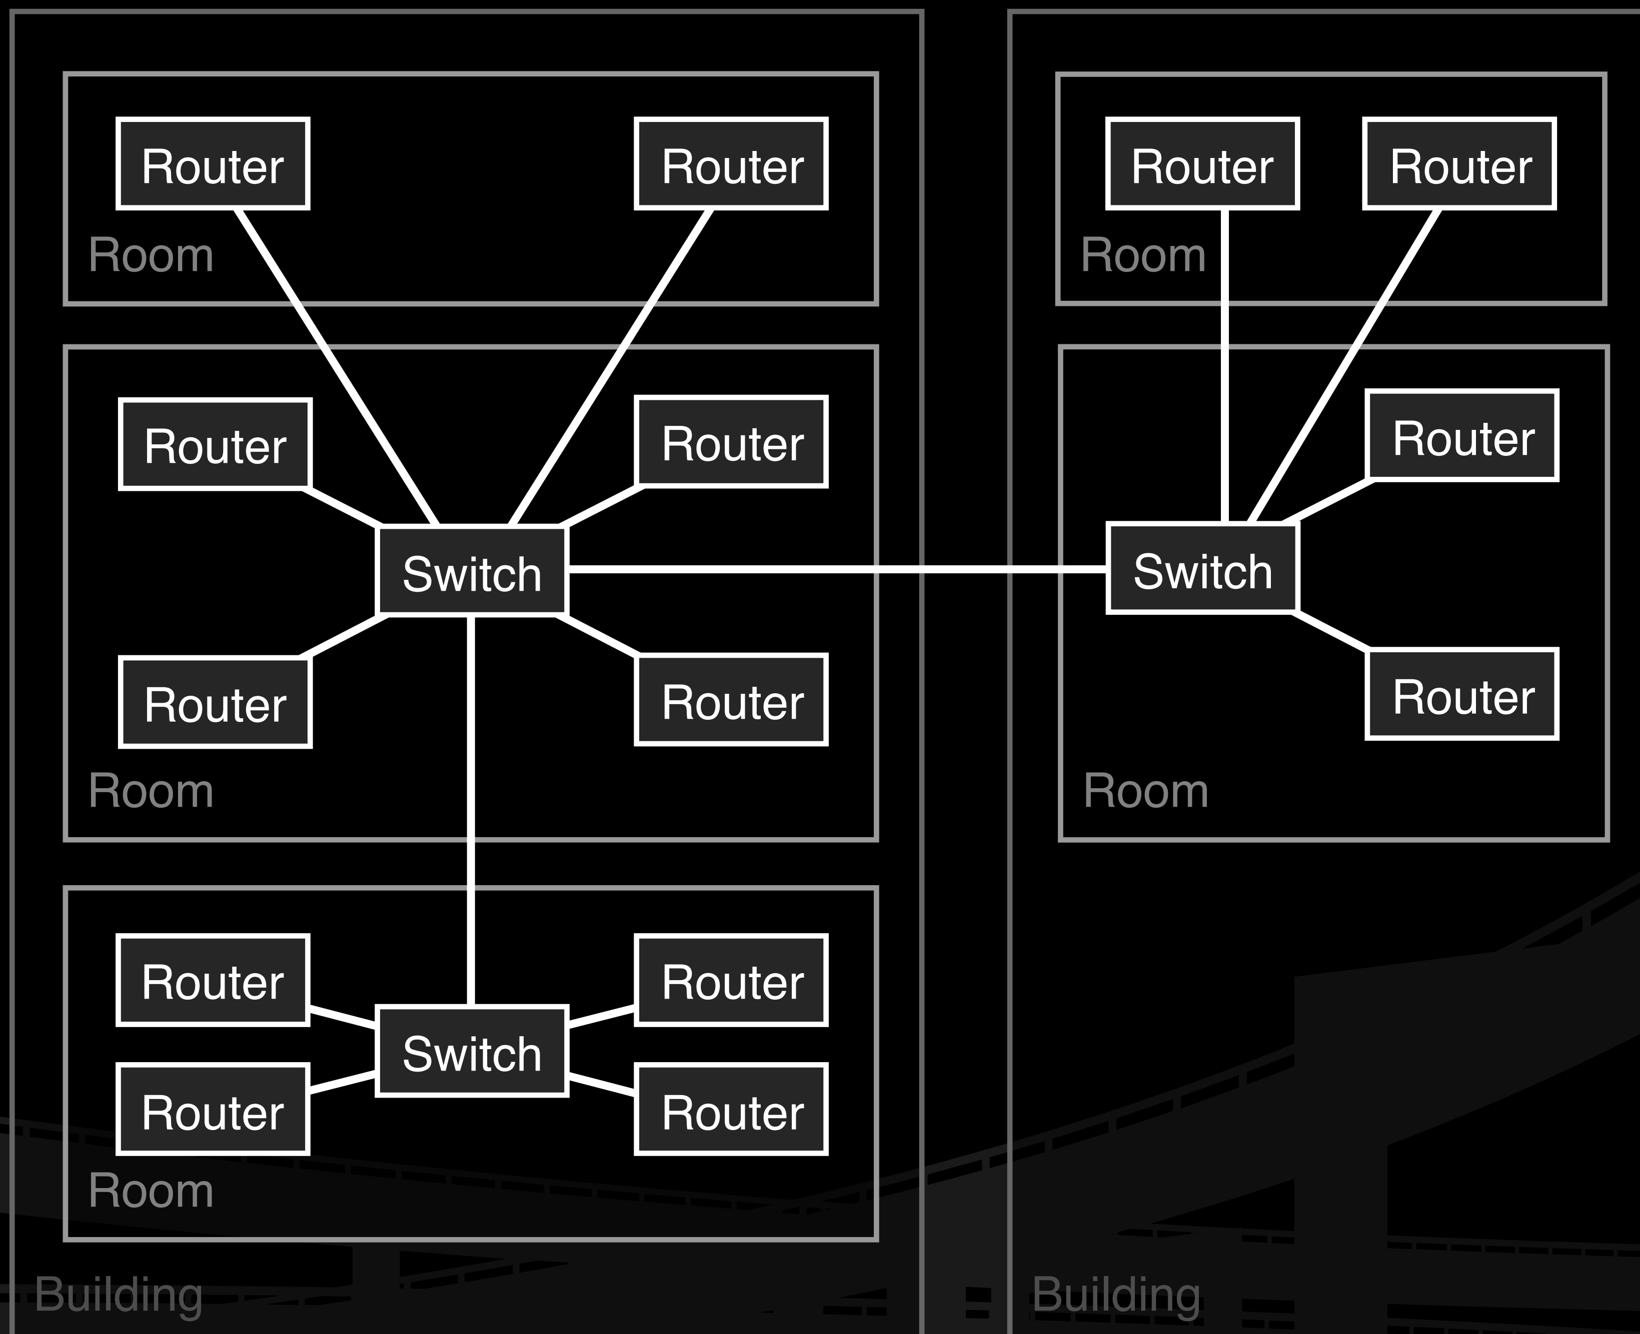

Growth of IXP Physical Plant

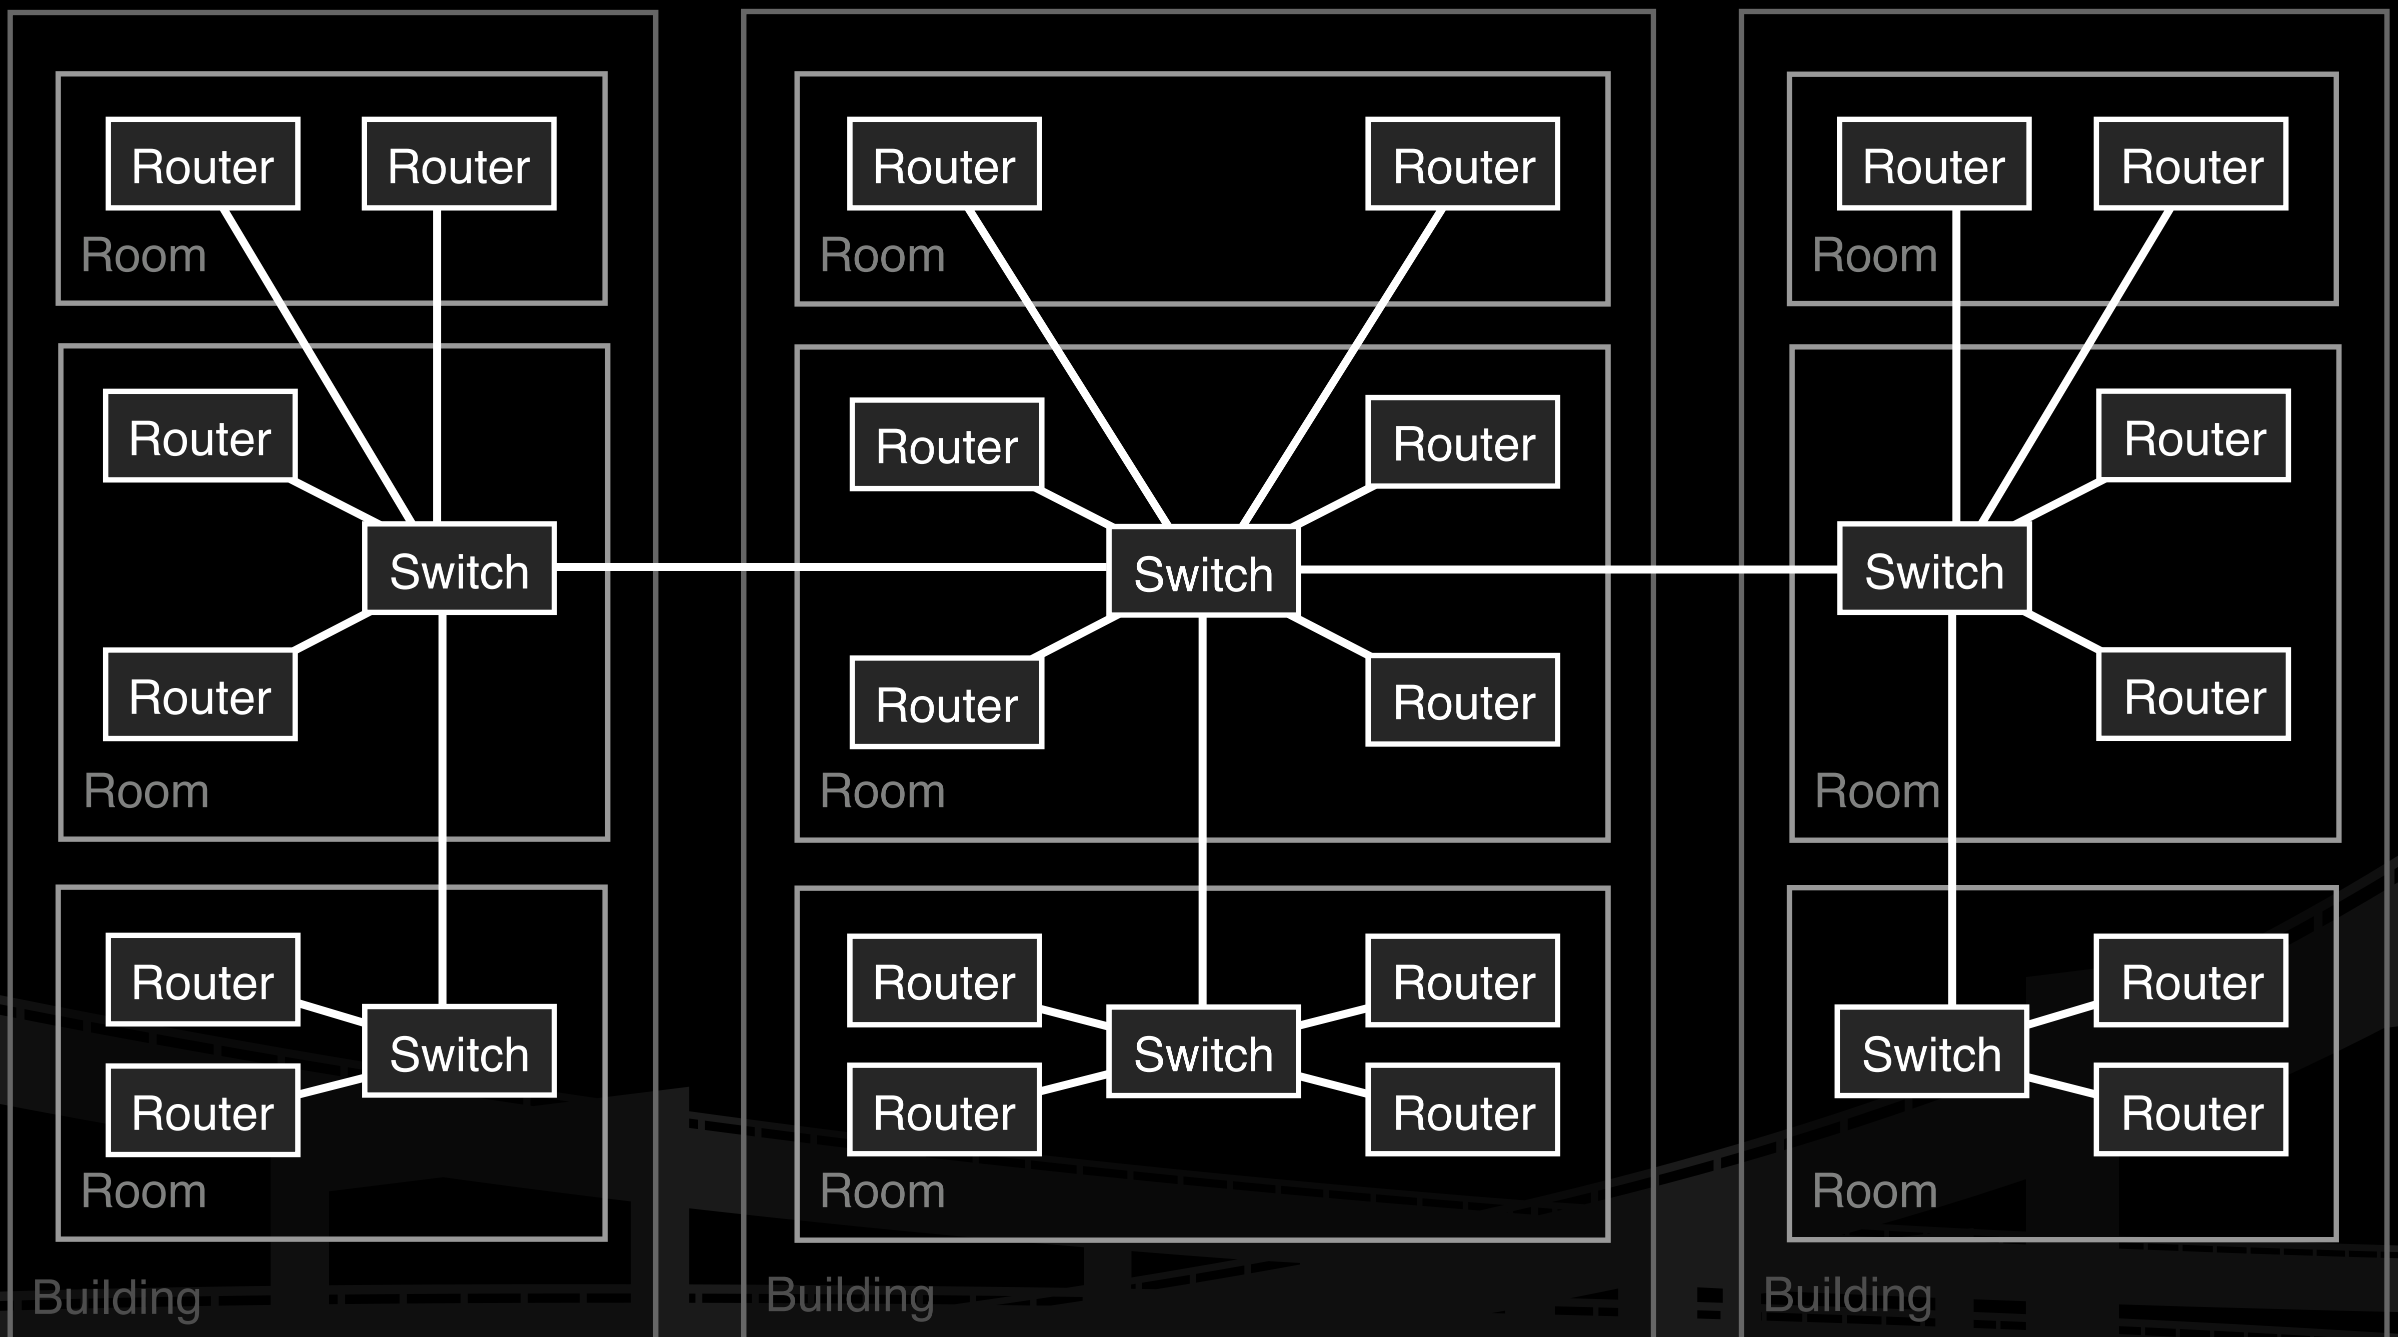

Layer 2 vs. Layer 3

L2 vs L3

L2 (peering over an ethernet switch) is ALWAYS cheaper than peering over a router. Since peering is meant to reduce the APBDC increasing the cost of peering infrastructure is counter-intuitive to supporting the development of the IXP

L2 vs L3

Disincentive to use new technologies like IPv6.

Location, Location, Location !!!