# Peering 101

William B. Norton

NANOG 45
Tutorial

Derived from UC Berkeley talk 10/21/08

v0.2

# My Promise to you

- You will understand (roughly) how the Internet is interconnected
- Specifically, you will understand and apply the following terms:
- 1) Internet Transit
- 2) Internet Peering
- 3) Internet Peering Ecosystem, Tier 1 ISPs, Tier 2 ISPs; their position and motivations, and the role of Internet Exchange Points

# Building Blocks

- Part I - Definitions of Transit and Peering
- Part II - Application of definitions: The Internet Peering Ecosystem
- Part III - The Theoretical Framework behind Internet Exchange Points

This is a Discussion Talk

Quizzes scattered throughout

# Who am i?

- William B. Norton (Bill Norton)
- 1998-2008 (former) Co-Founder and Chief Technical Liaison for Equinix, Inc.
- 1994-1998 North American Network Operators Group (NANOG) Chair
- - Internet Operations Researcher: Authored Industry White Papers...

**Interconnection Strategies for**

DRAFT 2.8

**Internet S**

DRAFT 1.4

DRAFT 1.1

**The Foll**



DRAFT 1.5

W. B. Norton

# Video Internet: The Next Wave of Massive Disruption to the U.S. Peering Ecosystem (v1.5)

William B. Norton <wbn@equinix.com>


## Abstract

In previous research we documented three significant disruptions to the U.S. Peering Ecosystem as the Cable Companies, Large Scale Network Savvy Content Companies, and Tier 2 ISPs started peering openly. By peering content directly with eyeballs, they effectively bypassed the Tier 1 ISPs resulting in improved performance, greater control over the end-user experience, and overall lower operating costs.

This paper predicts a new wave of disruption that potentially dwarfs this previous redirection of Internet traffic. Short video clip web sites, full length motion pictures, and television shows are now available via streaming to on-line devices and via downloading to iPods. More sites are coming on-line High quality movies from independent producers are being distributed via peer-to-peer methods. We observe these flash crowd effects and the larger movie file sizes as the crest of the first wave of significant incremental load on the Internet.

The majority of this paper details four models for Internet Video Distribution (Transit, Content Delivery Networks, Transit/Peering/DIY CDN, Peer2Peer) across three load models. The cost models include network and server equipment along with pricing models for various distribution methods. Over one hundred walkthroughs of this paper have led to stepwise refinements of the models and insights into why one would prefer or not prefer one model over the other.

The summary of the paper is a comparison of these video distribution techniques in terms of $-per-video units from the Video Service Provider perspective. We highlight cascading obstacles preventing large scale delivery of video traffic using commodity transit in a single location. The CDN solution and the multi-site Transit with Peering solution bypass some of these obstacles, while the peer-2-peer solution, while controversial, yields (by far) the lowest cost solution from the video service provider perspective.


## Previous Wave of Evolution of the U.S. Peering Ecosystem

The U.S. Internet Peering Ecosystem went through three significant disruptions in or about 2001[1]:

1. **Cable Companies Peer.** The North American cable companies' Internet transit provider (@Home) went bankrupt in 2001, forcing the cable companies to build out and manage their own multi-gigabit-per-second Internet infrastructure with only 30 days notice. With peer-2-peer traffic representing 40% to 60% of their transit bill, they quickly recognized the benefits of peering[2] that traffic directly with each other.
2. **The Large Scale Network Savvy Content Providers** entered into the Peering Ecosystem as their traffic volume grew into the ten's of gigabits-per-second. By engaging in peering directly with the Tier 2 ISPs, both groups were able to improve performance and lower their transit expenses while enhancing and increasing control over the end-user experience.

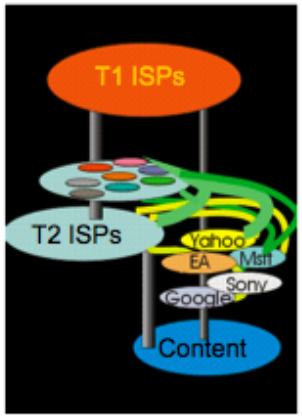


**Figure 1 – 2001 U.S. Internet Peering Ecosystem Evolution**

3. **MSOs peer with Content.** Since then, the cable companies peer directly with each other and with the large scale content companies. Several of these content companies have

[1] See "The Evolution of the U.S. Internet Peering Ecosystem" for a more detailed discussion of this.

[2] The term "Peering" is the reciprocal (and usually free) exchange of access to each others customers.





Tier 1 ISPs don't pay for transit.
minimize their interconnection
ding sufficient interconnection
t their customer base and their



**Peering Ecosystem**

mich.edu>

e Internet[1] consists of millions of network
s (routers, servers, workstations, etc.) operated
wide variety of network operators, content
ers, and end users. We will call these operators
etwork devices "Internet players". These
rs" are independent (but interconnected)
ers of a system we will call an "Ecosystem";
player has a definable role, an associated set of
ations and corresponding behaviors. We will
in this paper only on the core of the Internet
stem : the network operators that make up the
et *Peering* Ecosystem.

**finition:** The "Internet Peering Ecosystem"
ts of a community of loosely affiliated network
tors that interact and interconnect their networks
ious business relationships.

ere are many Internet Peering Ecosystems
d the world, each with their own set of network
tors collectively providing Internet access to the
et Region.

ere are vast differences between each region of
nternet Peering Ecosystem[2], so we model the
l Internet Peering Ecosystem as a loosely
ed set of "Internet Regions" each with its own
g Ecosystem.

Internet Regions

order for an ISP to provide access to the global
et, it must get attached to the Global Internet. It
o with either Transit[3] or Peering[4] relationships,
ombination of both within an Internet Region.

convention, when we refer to the big 'I' Internet we
an the global Internet.

second follow-on paper, we document some of the
tional Peering Dynamics, detailing some general
teristics and some unique characteristics of the Japan,
lia, Hong Kong and Singapore Peering Ecosystems.
o touches on several dynamics unique to the
tional Internet Peering environment.

nition: **Transit** is a business relationship whereby
e ISP provides (usually sells) access to the Global
ernet.

nition: **Peering** is a business relationship whereby
o ISPs provide reciprocal access to each others'
tomers. This is typically a free exchange of traffic.

11/19/2003
Comments Welcome to <wbn@equinix.com>

# Part I: Definitions of Transit and Peering

# Def: The Internet is a network of networks.

# Def: ISP sells access to the Internet, so...

# An ISP must itself get attached to an ISP already attached to the Internet.

Def: ‘Transit’ is service whereby one ISP sells access to the Internet.

“A port in the wall that says ‘Internet this way’”


metered
95th percentile...

Q?Transit Billing...

# Billing Internet Transit: 95th Percentile

1 month of 5 min Samples=$v_n$ - $v_{n-1}$

Highest Monthly Sample


36 hrs/month to burst for free

95th Percentile measure

Lowest Monthly Sample


**Challenge**: How can you GAME this system?

# Free Transit!!
# 35hrs burst to 26 ISPs

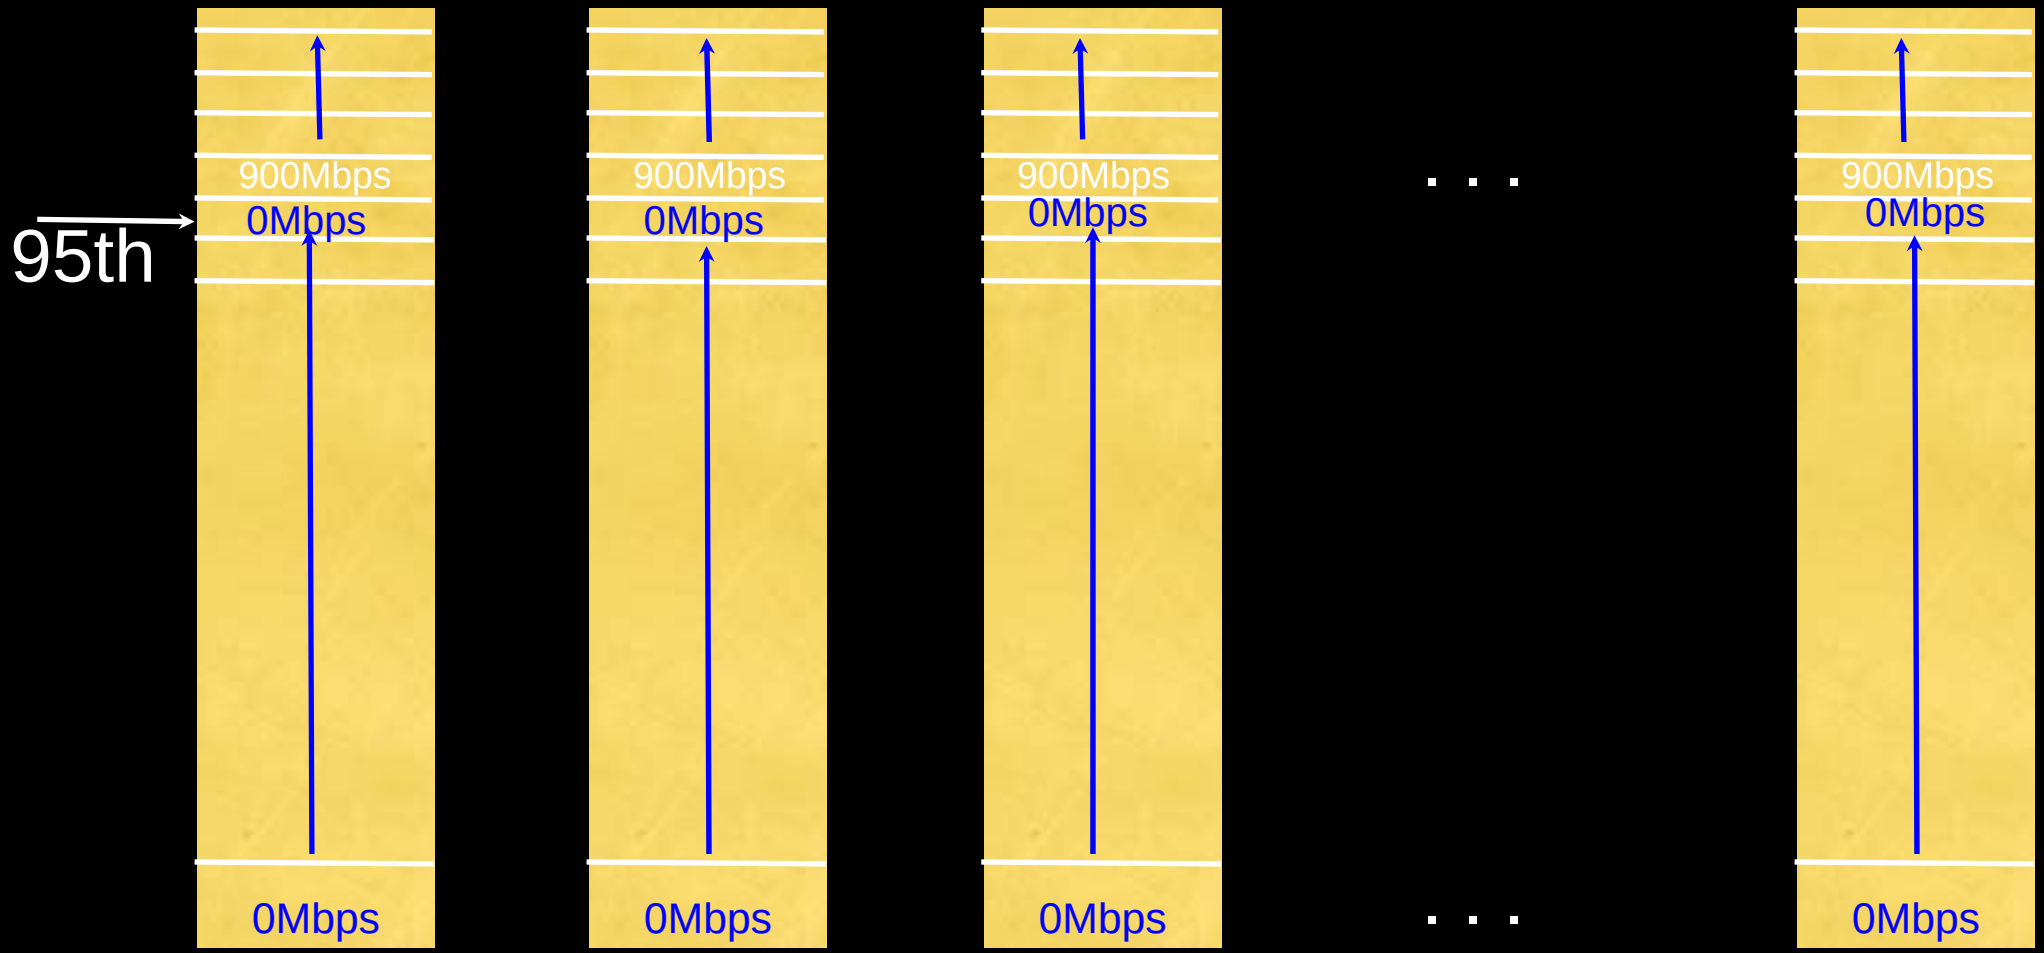


Enter minimum commits..

# Minimum Commits: Tiered Transit Pricing

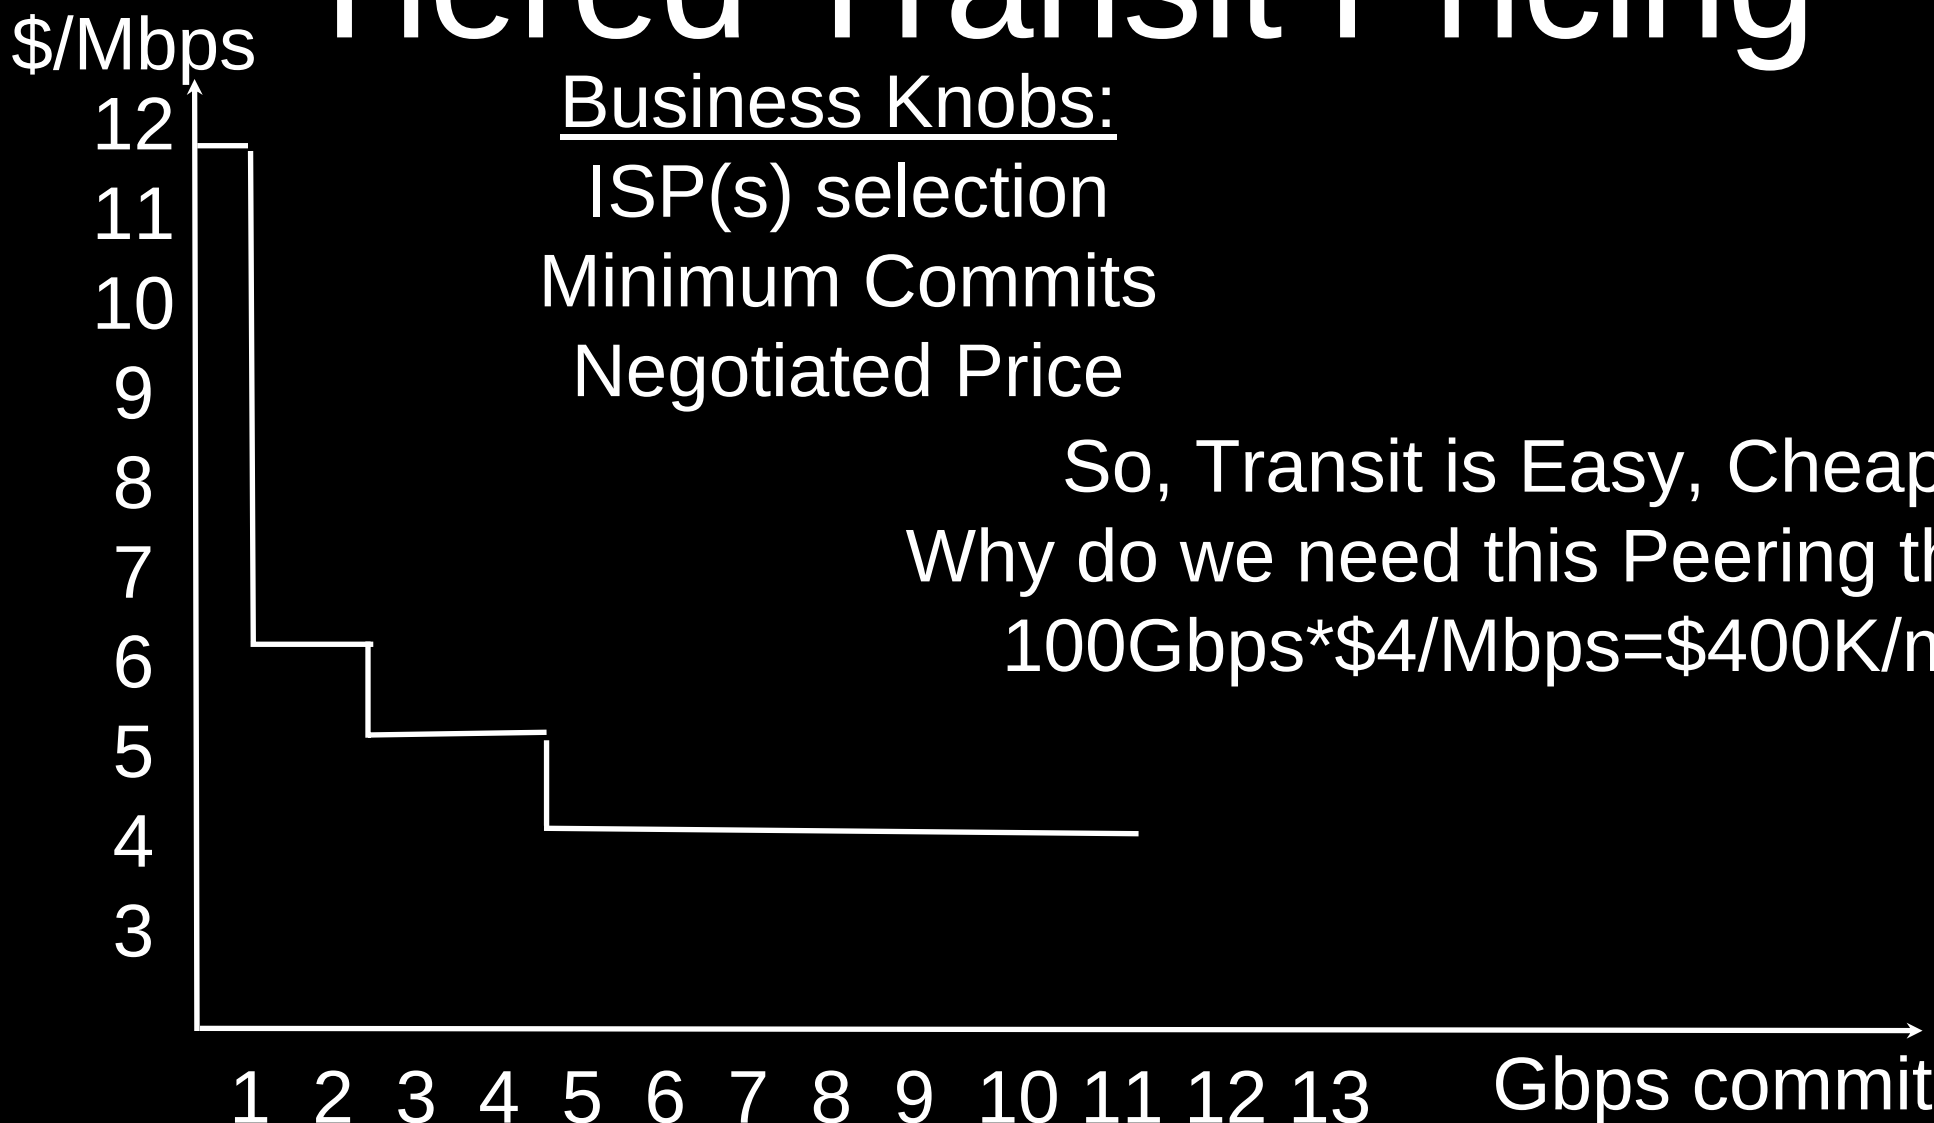


Business Knobs:

ISP(s) selection

Minimum Commits

Negotiated Price

So, Transit is Easy, Cheap

Why do we need this Peering thing?

100Gbps*$4/Mbps=$400K/mo

Source: 2008 NANOG Discussions: $4/Mbps!

# **Def: 'Peering'** is a reciprocal exchange of access to each others customers.

<u>Motivations to peer</u>

1) Cost Savings
2) Performance Benefits
3) Additional revenue
   3a) ABOV
   3b) AMZN

All traffic except red traffic goes this way

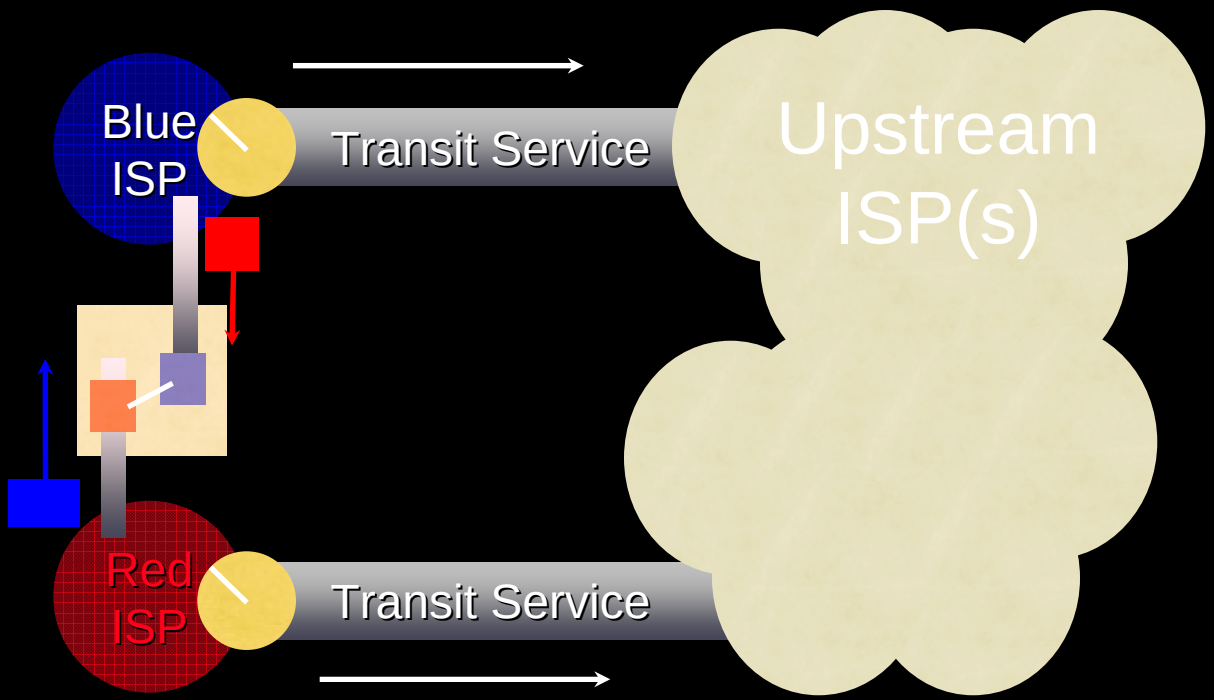


All traffic except blue traffic goes this way

# Two key points about peering

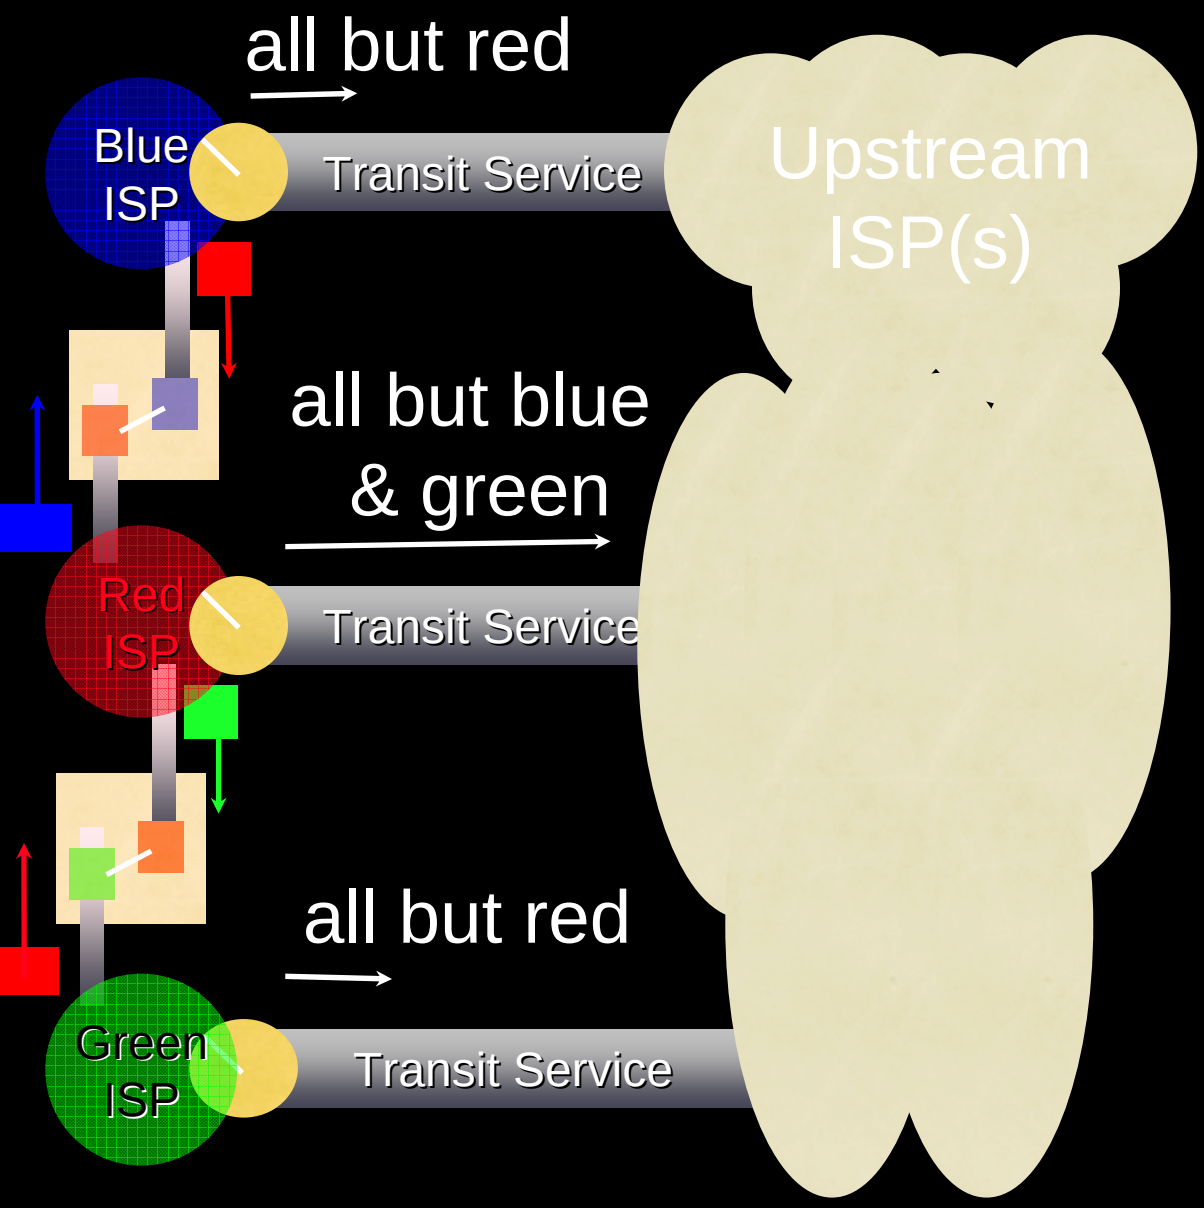


1) Peering is *not* *Transitive*

2) Peering is *not a perfect substitute* for Transit

# Part II - The Internet Peering Ecosystem

## From 30,000 feet - applying the definitions

Armed with these definitions..

Def: Global Internet Peering Ecosystem consists of a set of interconnected internet regions (countries).

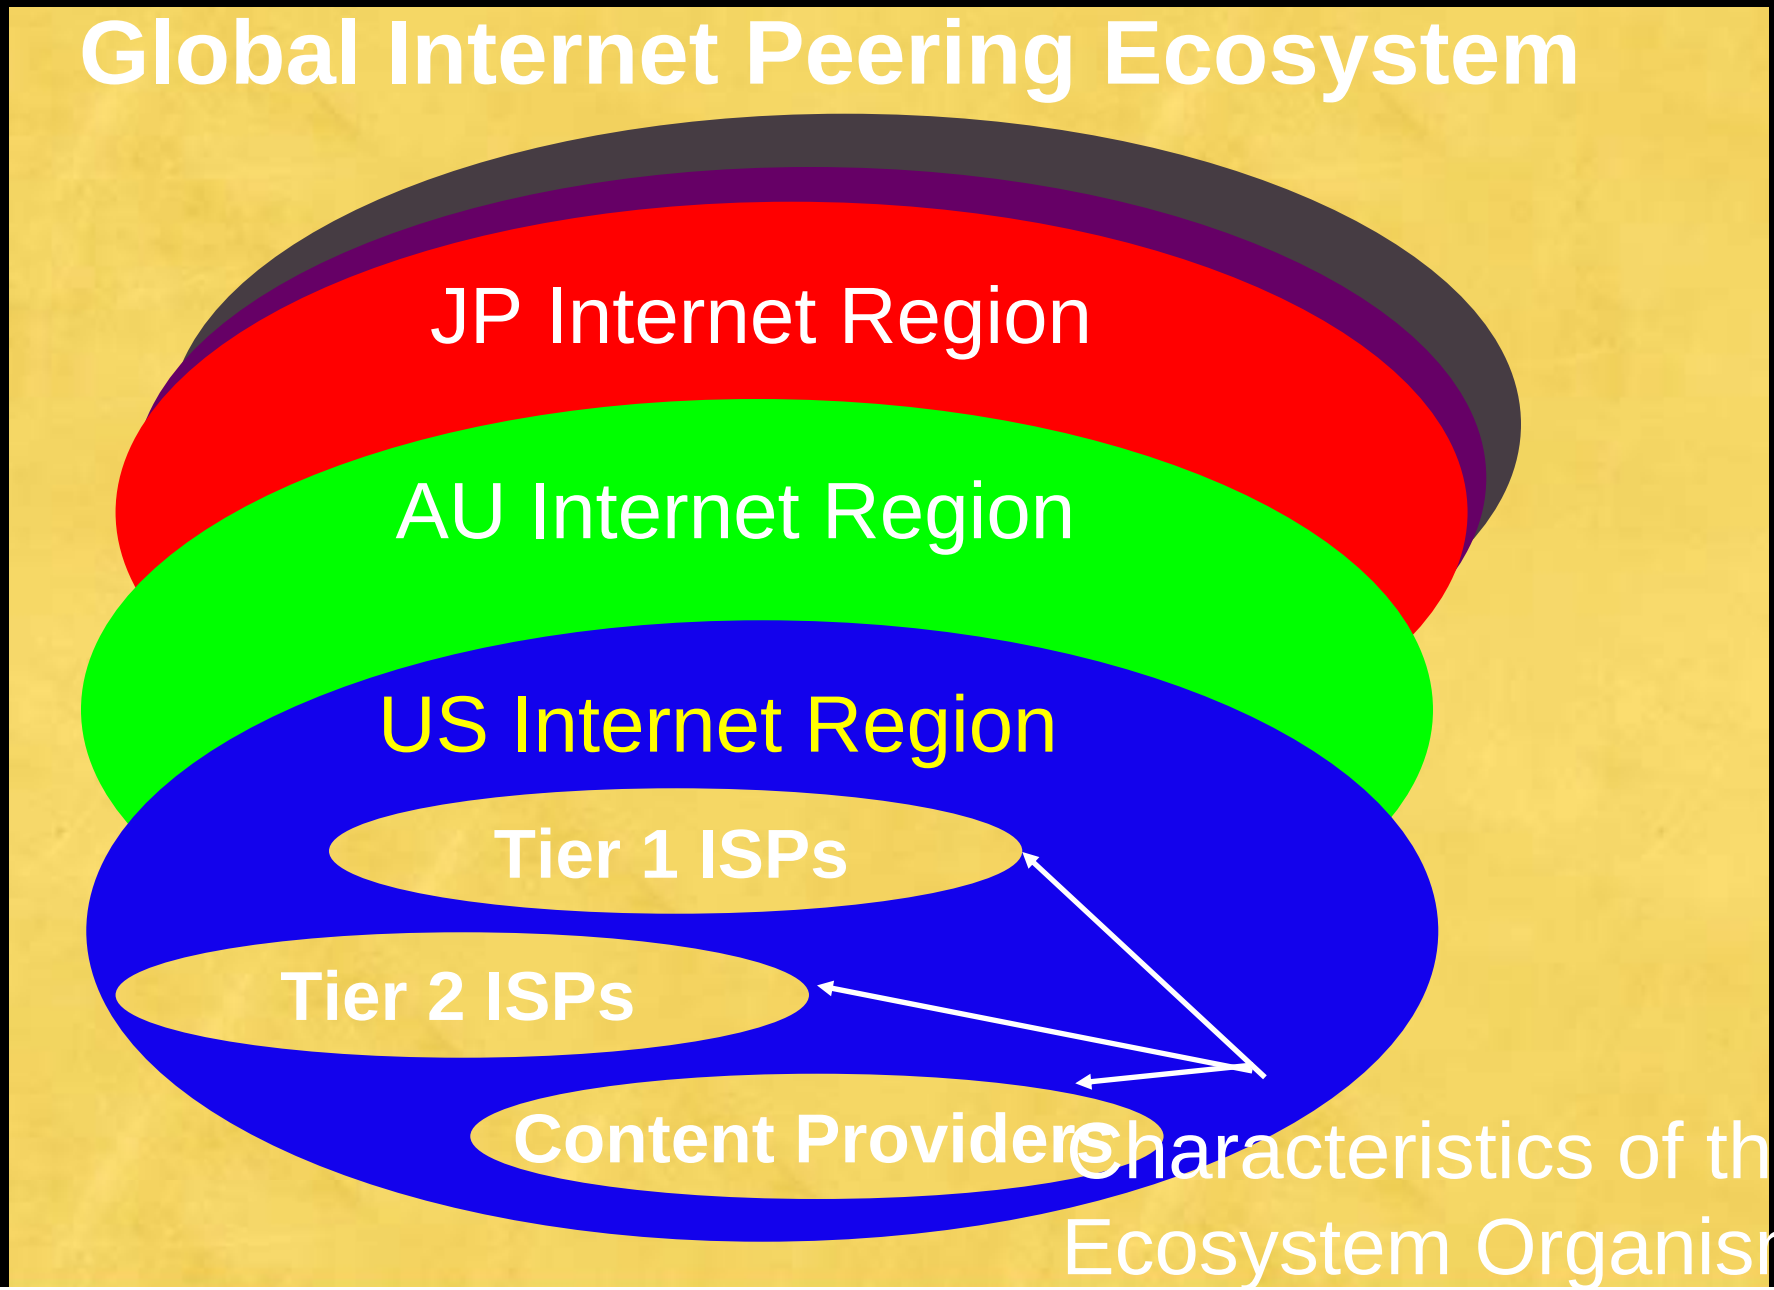

# Ecosystem Member: Tier 1 ISP

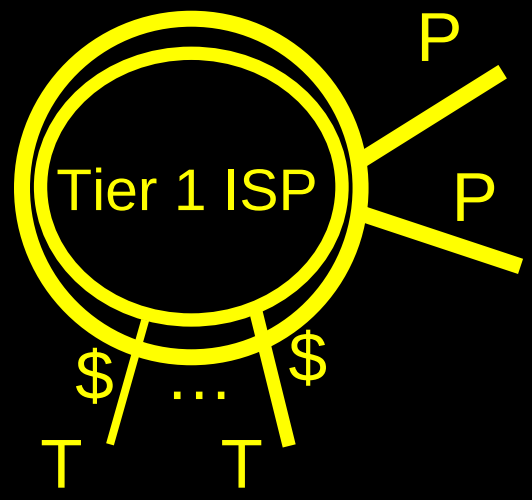


**Def: A Tier 1 ISP is an ISP that has access to the ENTIRE Internet Region Routing Table Solely via Peering Relationships**

(Doesn't buy transit from anyone to reach any destination in the Internet Region.)

Motivation: Is NOT motivated to Peer in region to reduce transit fees, Is NOT motivated to peer with anybody else.

Behavior: "Restrictive" Peering *def: Policy

# Ecosystem Member: Tier 2 ISP

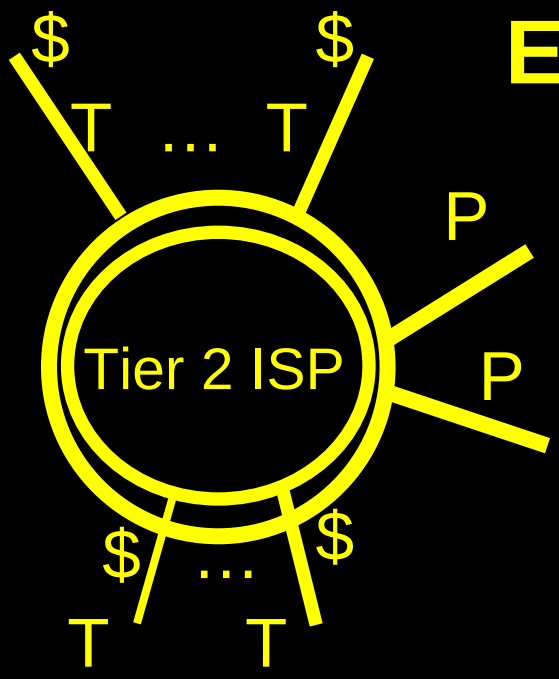


**Def: A Tier 2 ISP is an ISP that has to purchase Transit to access some part of the Internet Region.**

Motivation: Is motivated to Peer in region to reduce transit fees.

Behavior: "Open" Peering or "Selective" Peering Policy Active in Peering Forums

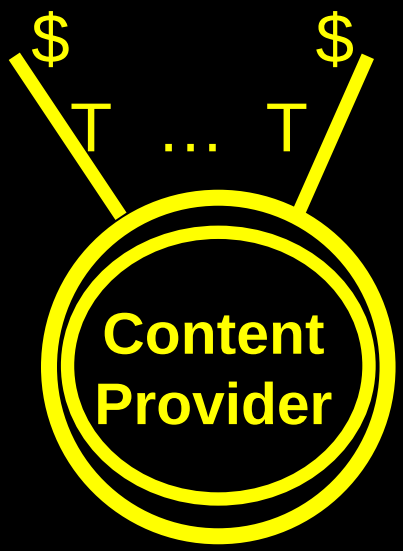


# Content Providers

**Def: A Content Provider focuses on content development and does not Sell access to the Internet.**

Motivation: SLAs w/well known ISP

Behavior: “No Peering” Policy

# Generic Peering Ecosystem

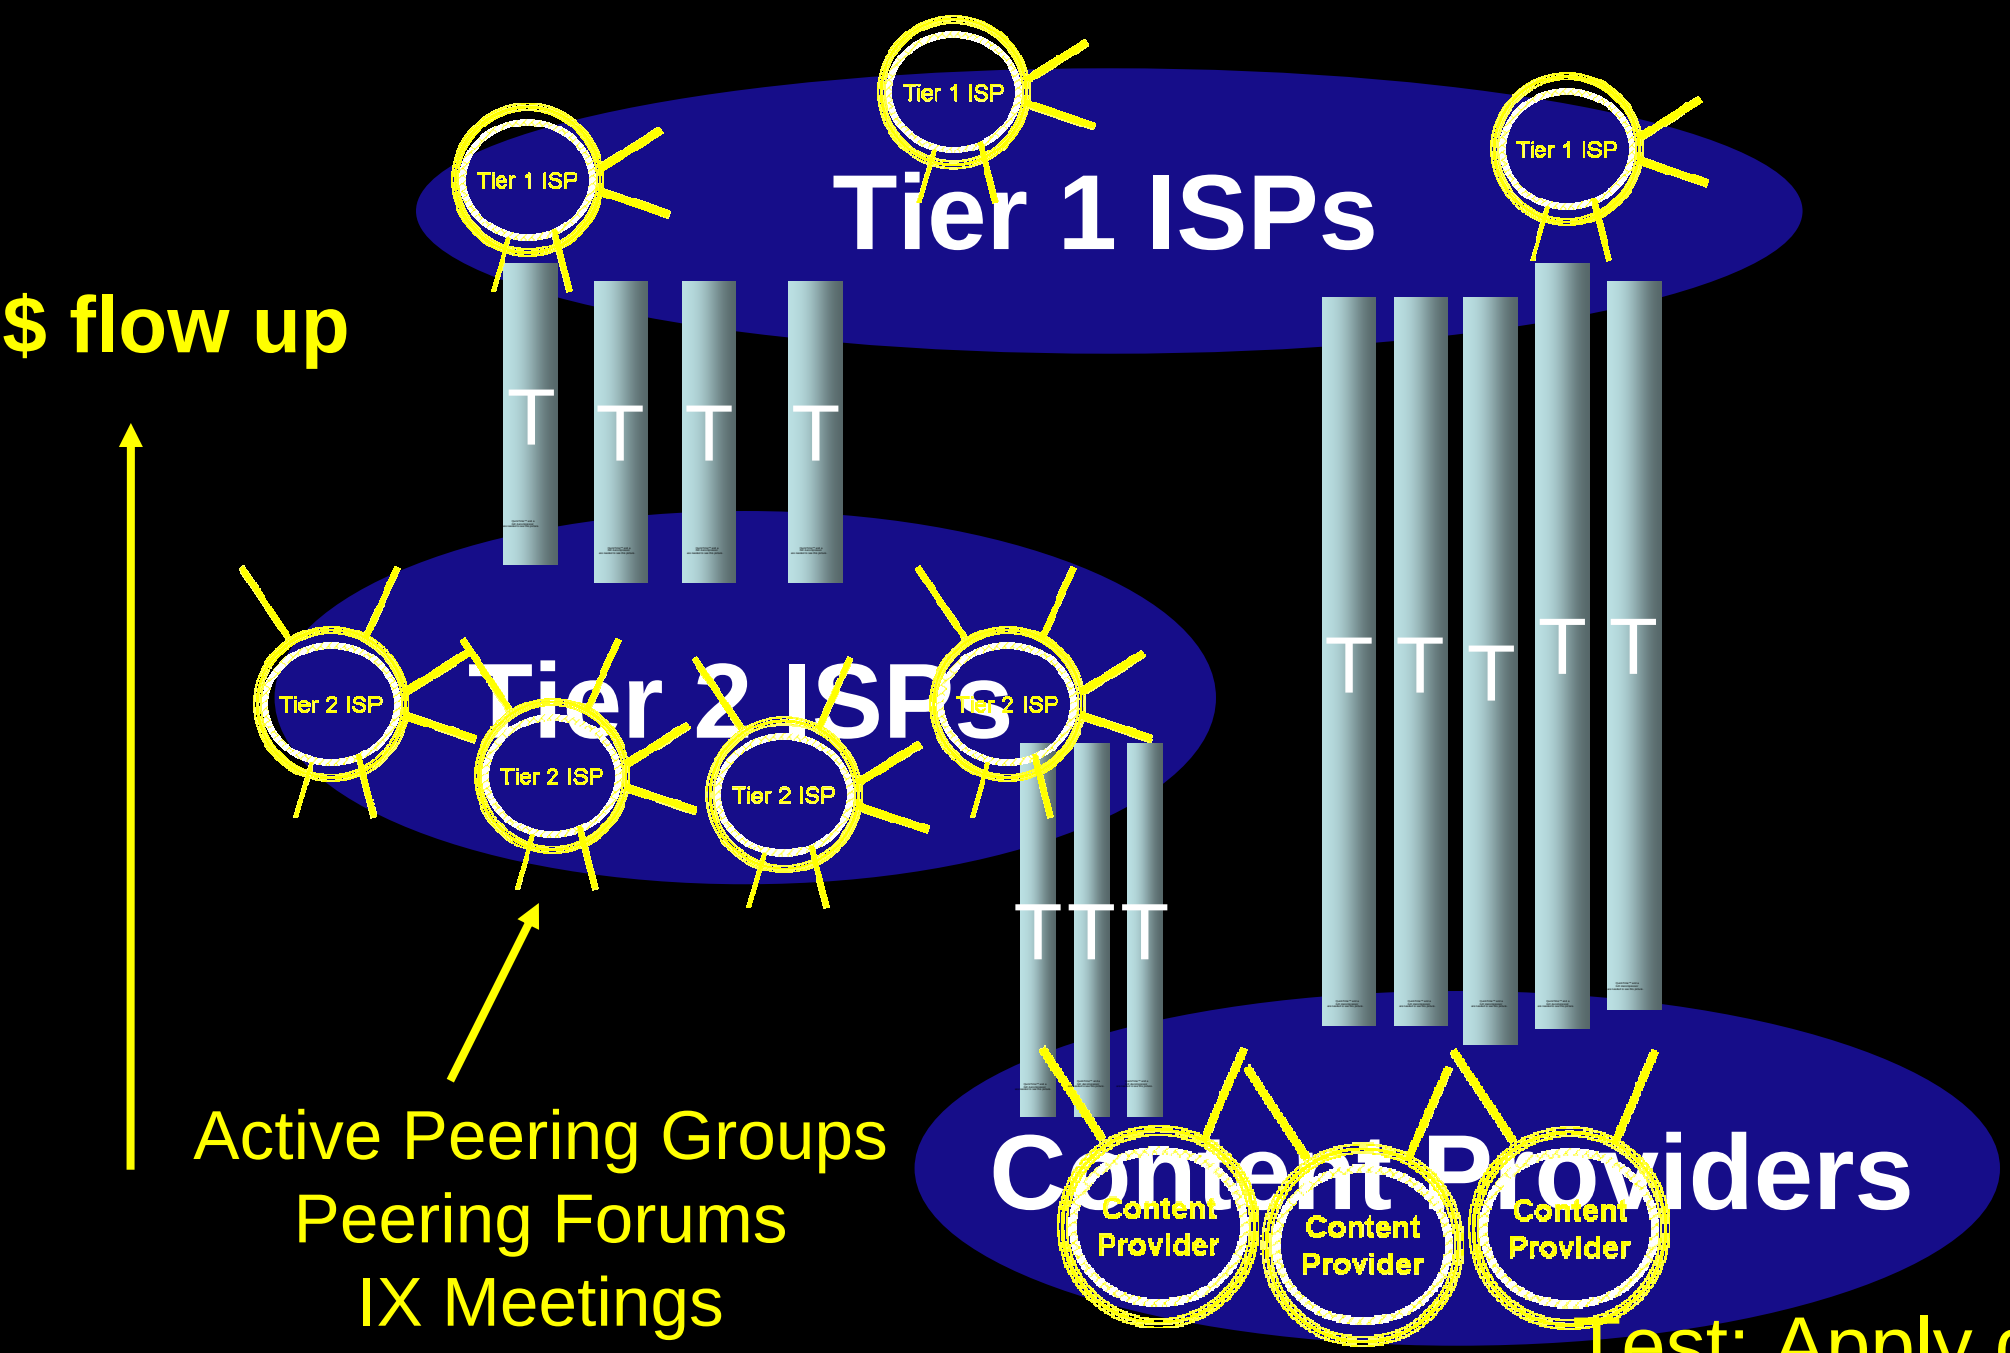

# Quiz

2) Definition of Peering:

_____

1) Definition of Transit:

_____

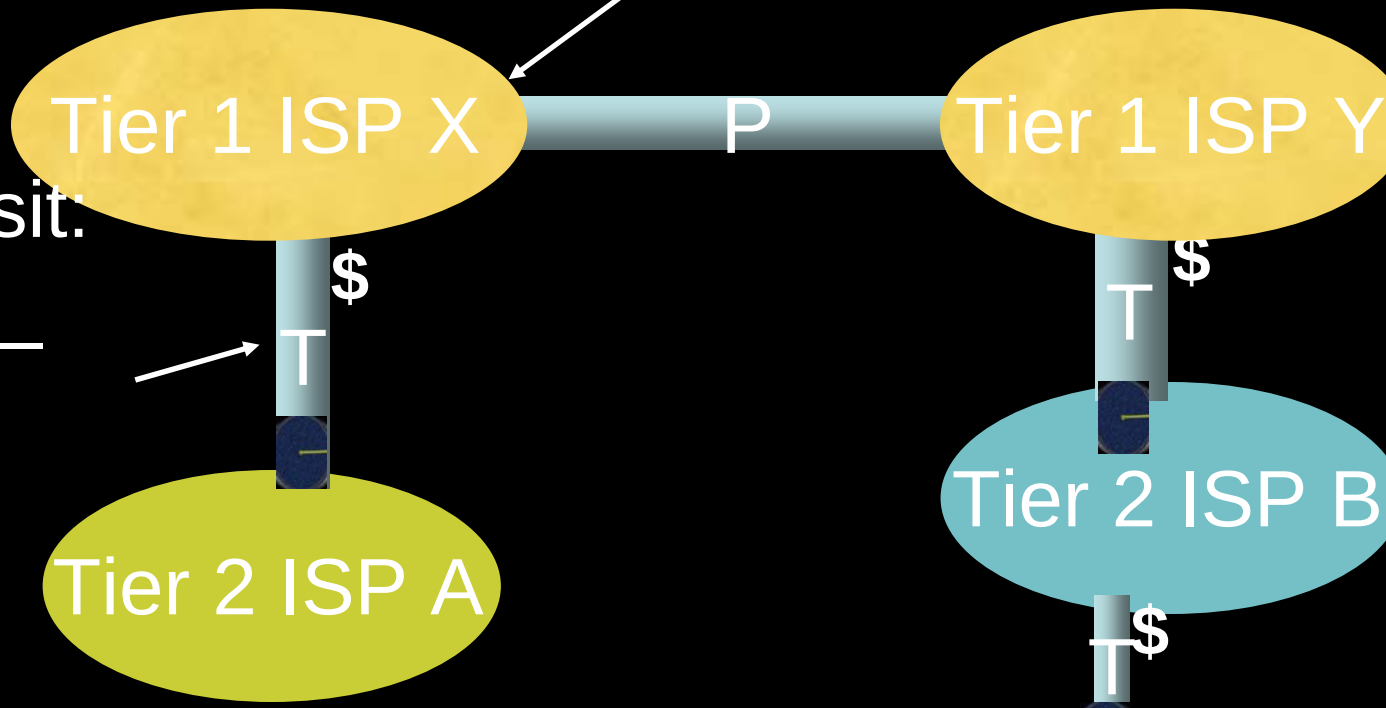


3) Definition of an “Open” Peering Policy:

_____

4) Definition of a “Selective” Peering Policy:

_____

5) Definition of a “Restrictive” Peering Policy:

# Apply Defs: Peering Dynamics & Motivations

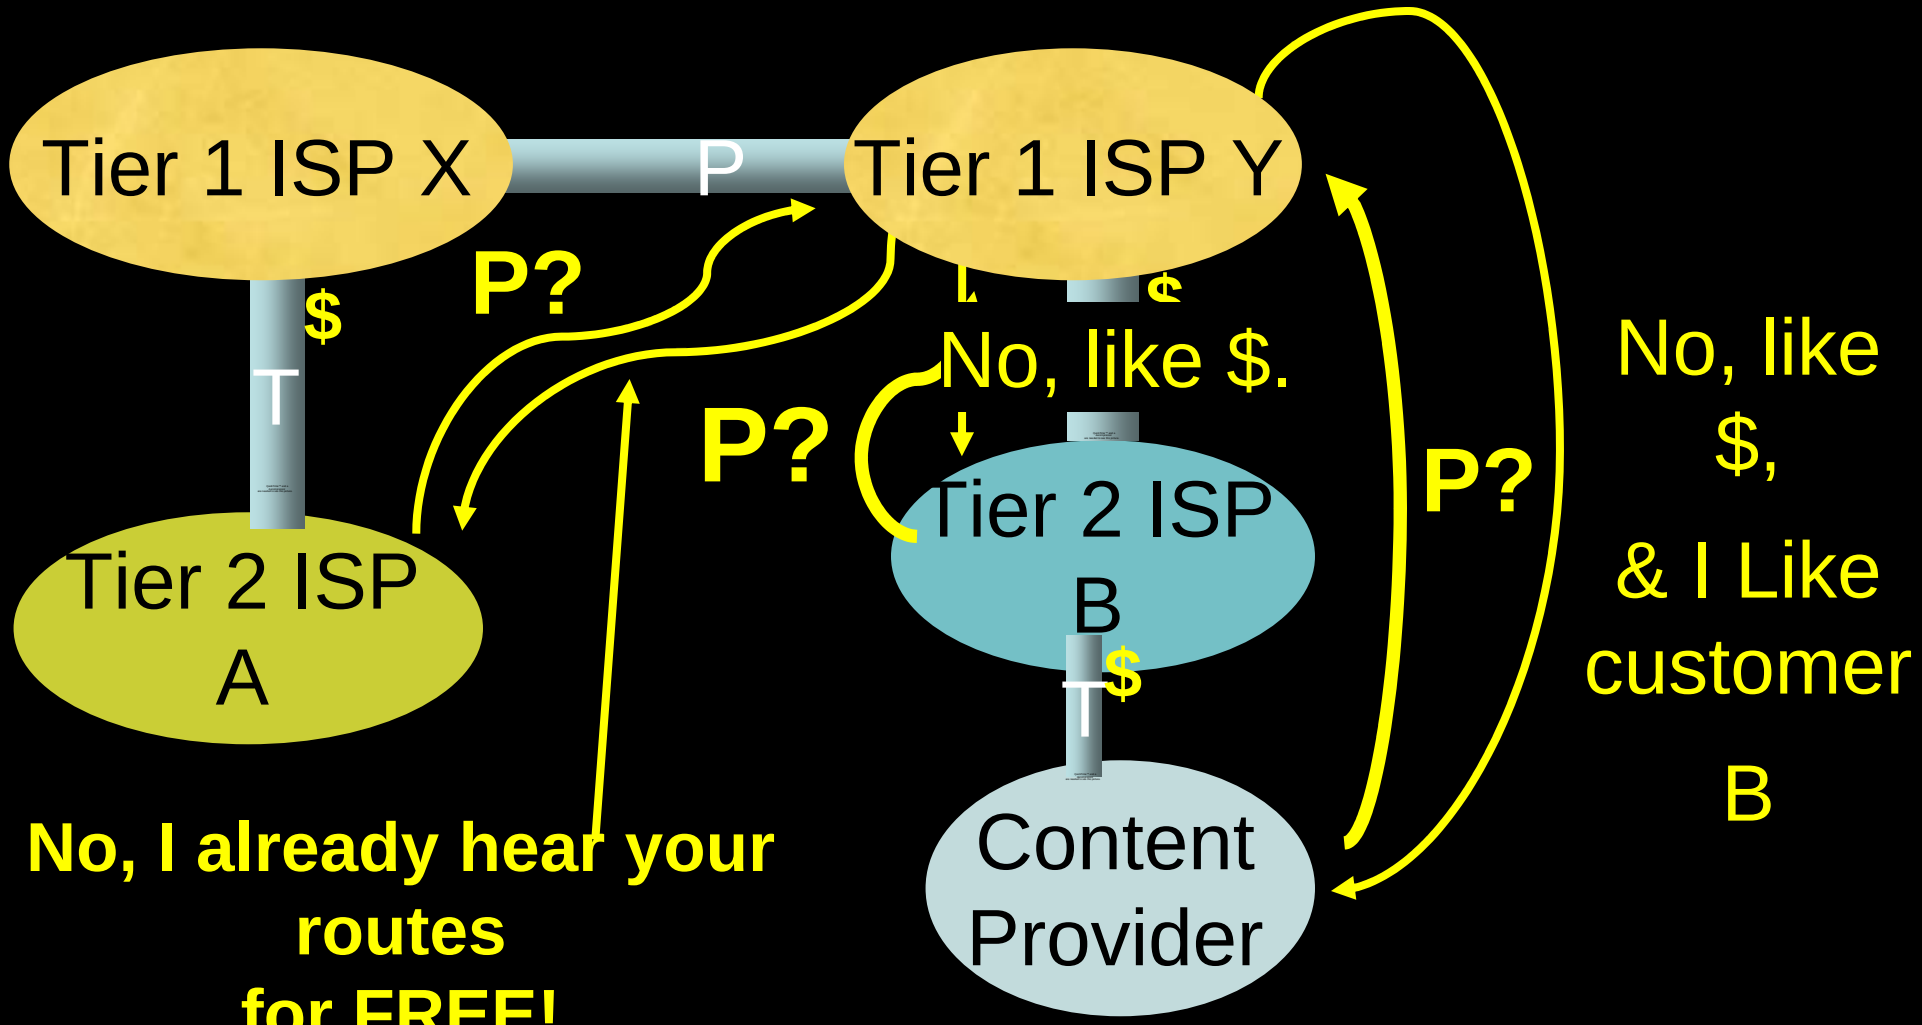


**Synch Point:**

**You have all the defs needed to predict behavior in the Peering Ecosystem.**
**You should be able to answer the question at hand.**

# Evolution of the U.S. Peering Ecosystem

Illustrative of dynamics

Applies definitions

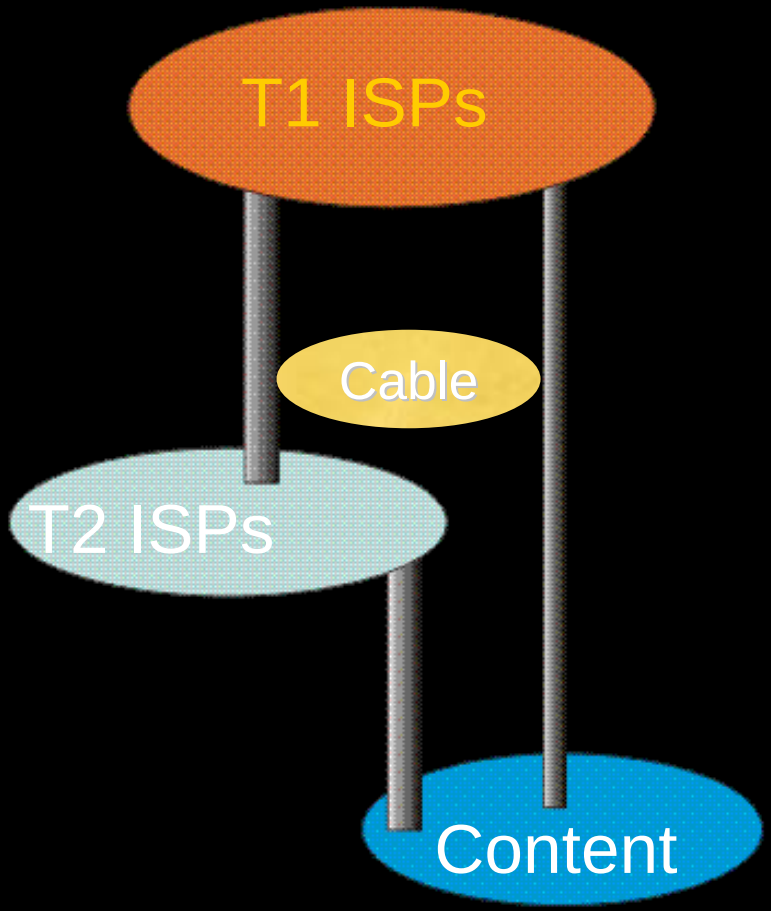


# U.S. Evolution #1
# Cable Companies Peer

## Significant Evolution…

1) Volume of traffic is huge

2) Cable Cos  Open Peering

3) “Kazaa Effect” amplifies peering benefits

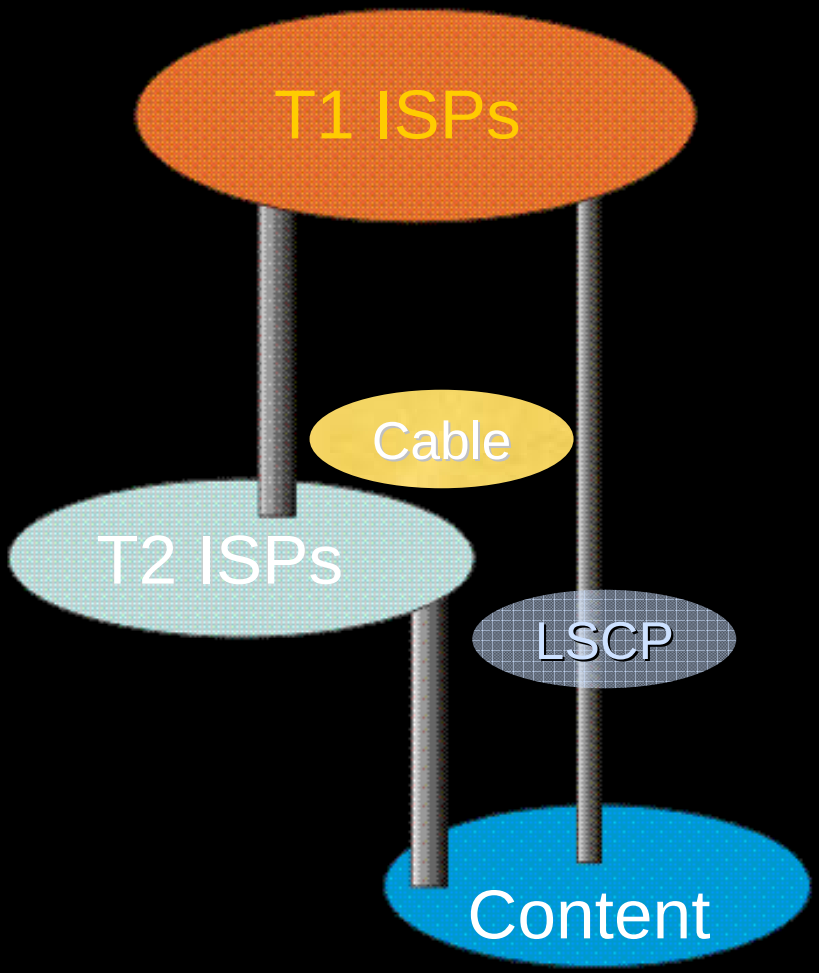


# U.S. Evolution #2
## Large Scale Content Players Peer

### Significant Evolution…

1) Volume of traffic is huge

2) Content is Open Peering

3) Improves End-User Experience

4) Leading Players are paving the way

…need to move out of Bankrupt colo anyway…

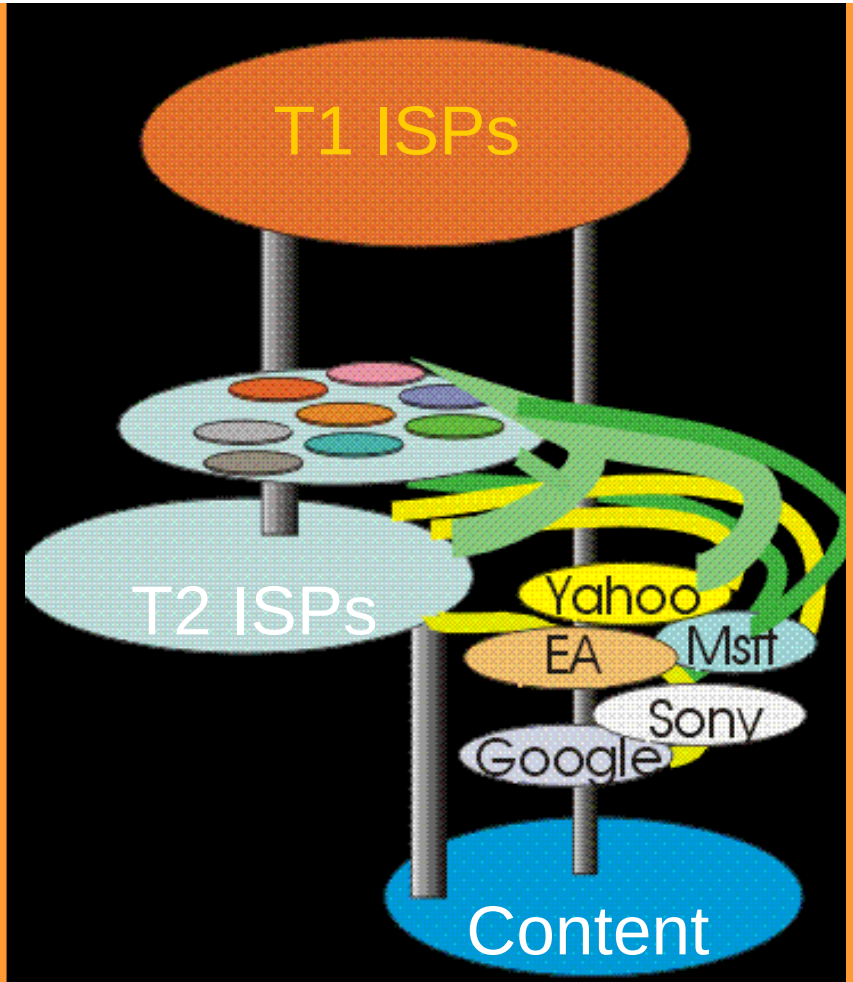


# U.S. Evolution #3
# Cable Cos Peer w/Large Scale Content Players

## Significant Evolution…

1) Volume of traffic pulled away from T1s is huge

2) Reduces perceived need for T1s (for local delivery anyway)

3) T1s still needed for distance

→ Content Literally right on the

# Internet Exchange Points

## A Theoretical Framework

# IX Network Externality

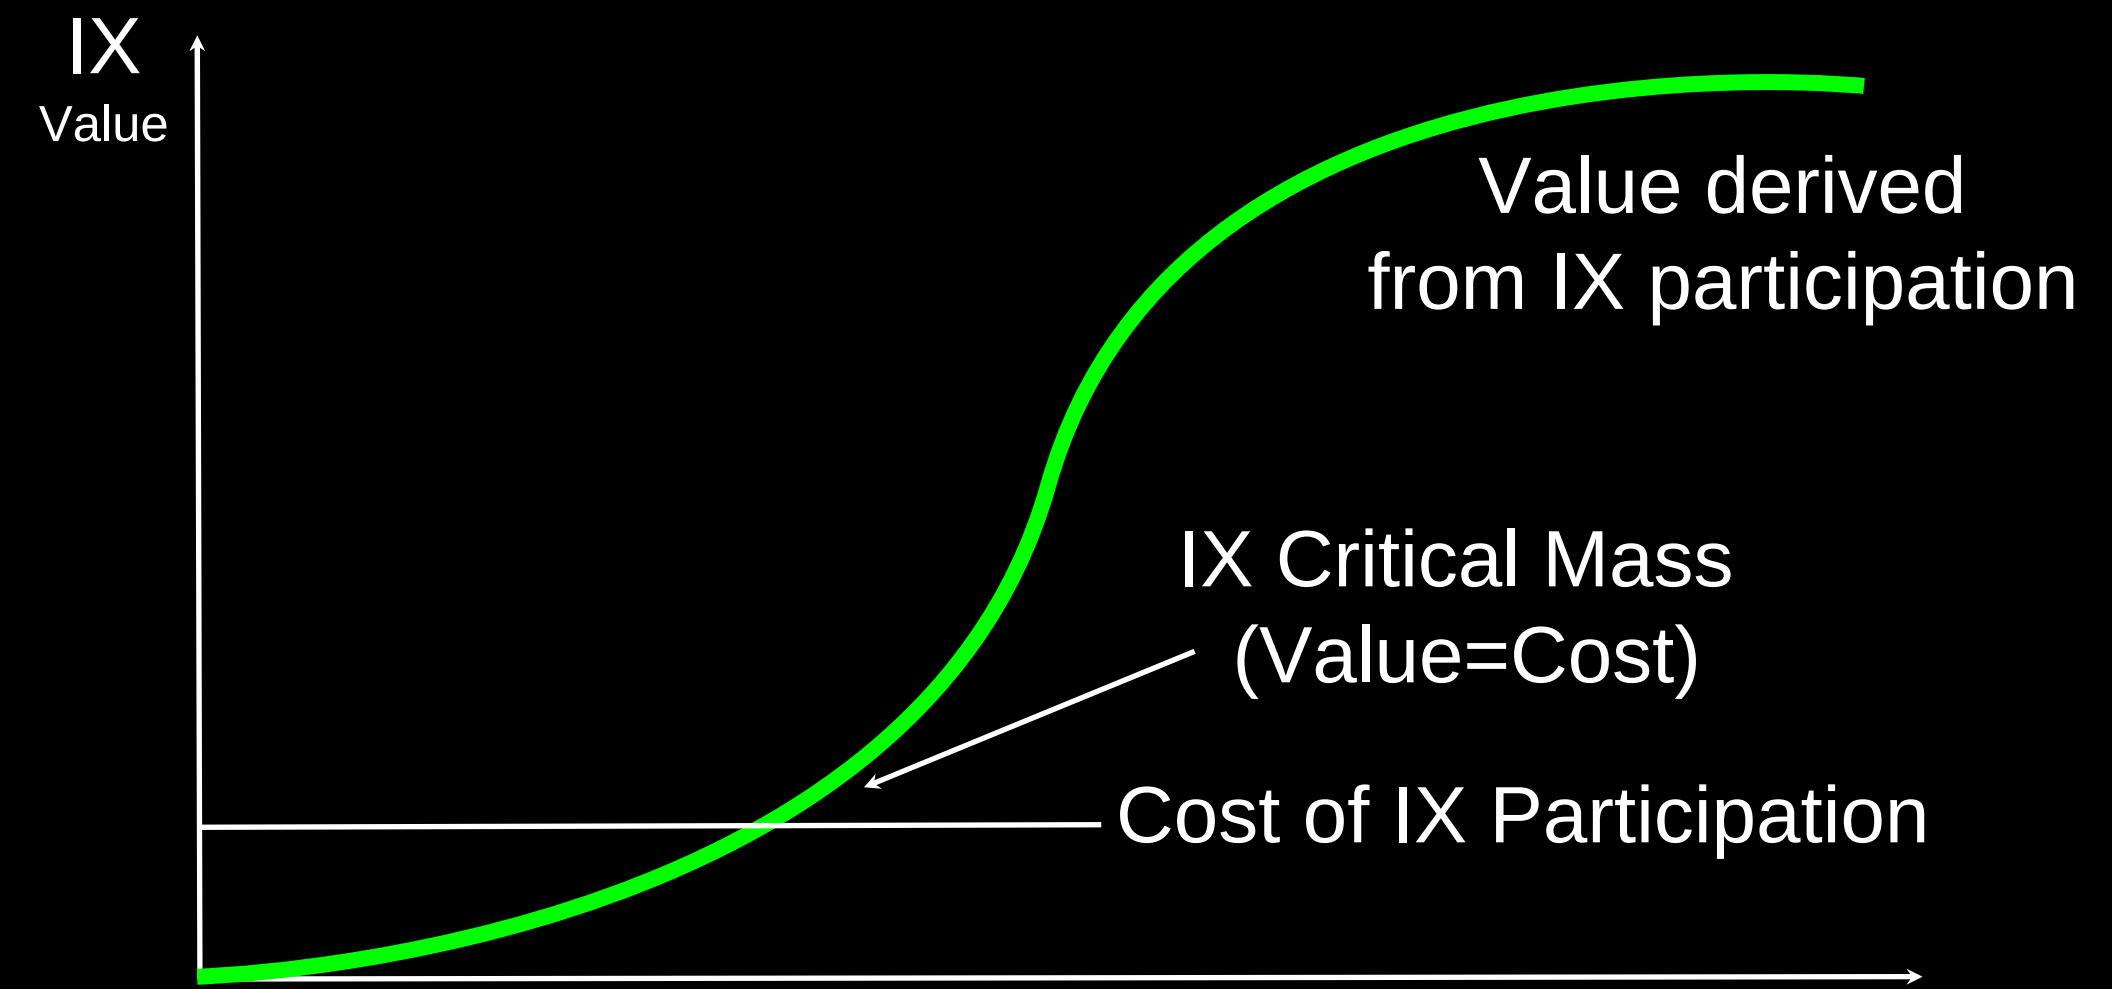


The Startup Hump

Discussion Here

# Asked IX Operators

- How did you get to critical mass?
- Europe: ISP consortium starts it
- Commercial Company targets key ISPs
- Drop price of Participation
- Equity
- Evangelize, host content (BW sales as lure in)
- Find new large volume target peering customers (Video)

Source: The Art of Peering: the IX Playbook

# Models of IXes

- IX separate from colo (Europe
- IX owns colo (US)

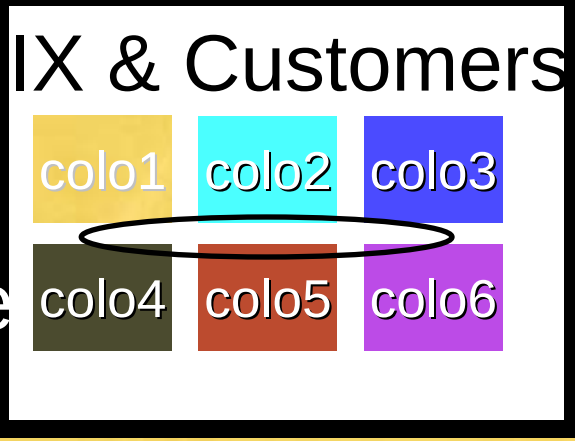


- LINX/AMS-IX/DE-CIX model
- US Equinix/PAIX/NOTA mod

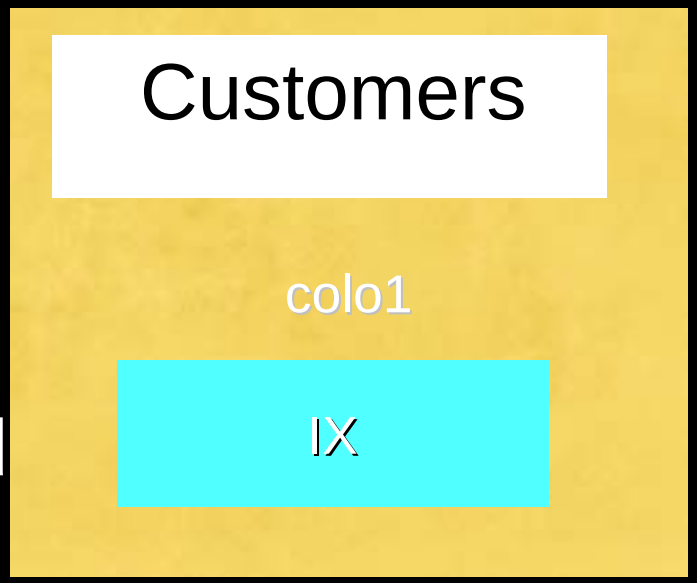

# Neutrality

- ISP-neutrality
- Carrier-Neutrality
- Carrier and Colo neutrality
- Why is this important - Turkish Internet example

# Peering Math

| | | |
|---|---|---|
| Colo | $1,500 | per mo |
| IX Port | $3,500 | per mo |
| Total Peering Costs | $5,000 | |

| Gbps | Transit $/Mbps | Peering |
|---|---|---|
| 1 | $5.00 | $5.00 |
| 2 | $4.50 | $2.50 |
| 3 | $4.50 | $1.67 |
| 4 | $4.50 | $1.25 |
| 5 | $4.30 | $1.00 |
| 6 | $4.30 | $0.83 |
| 7 | $4.30 | $0.71 |
| 8 | $4.00 | $0.63 |

# Peering v Transit

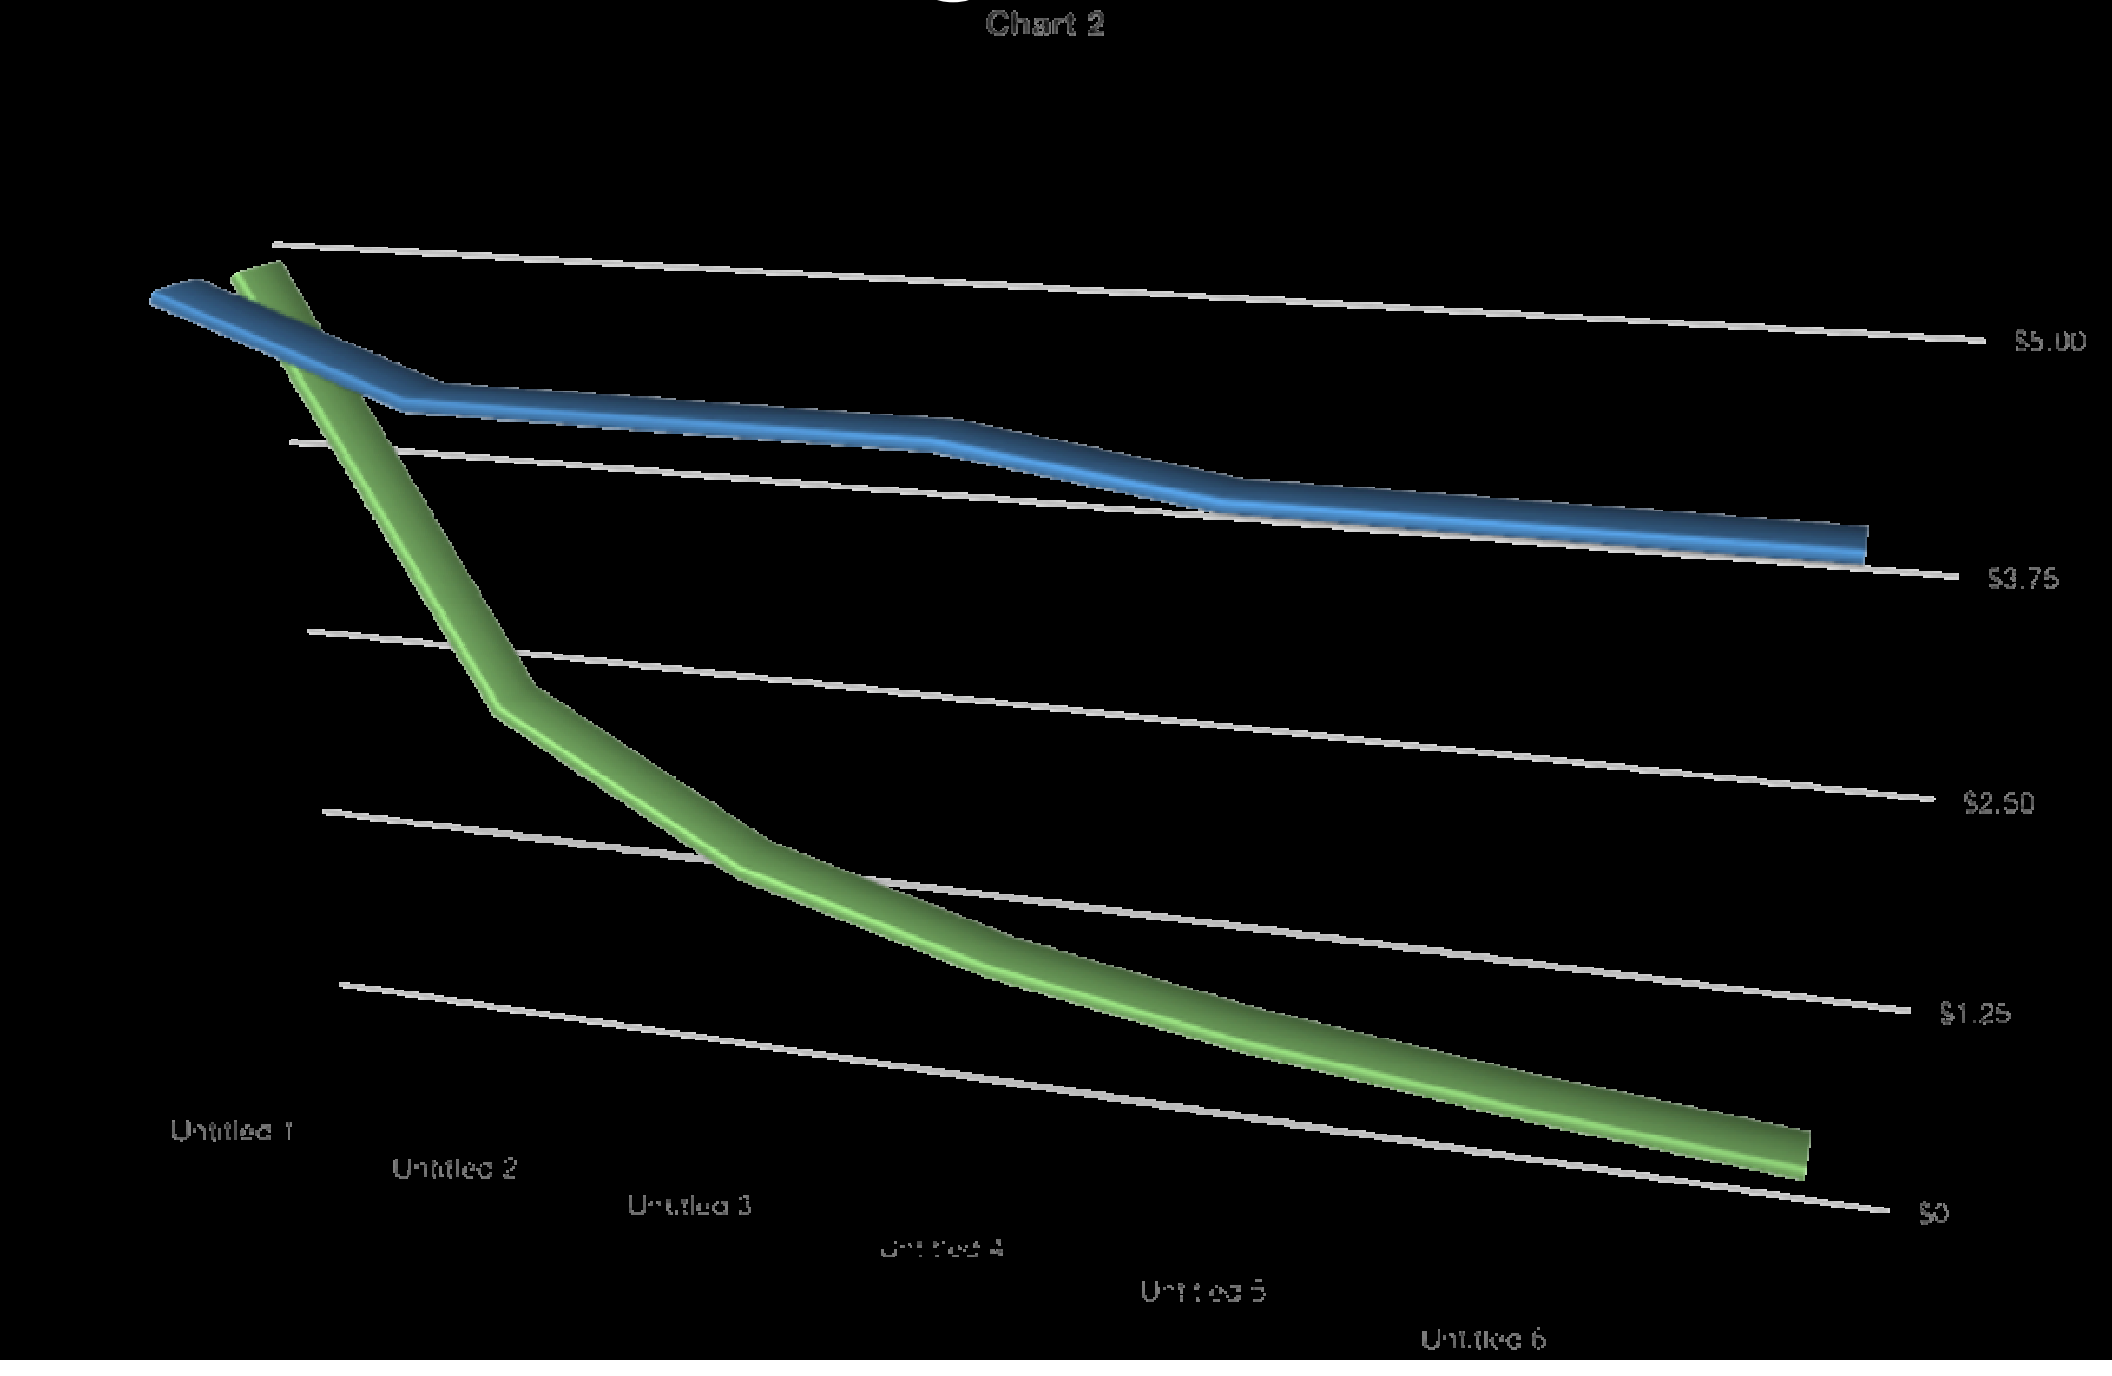

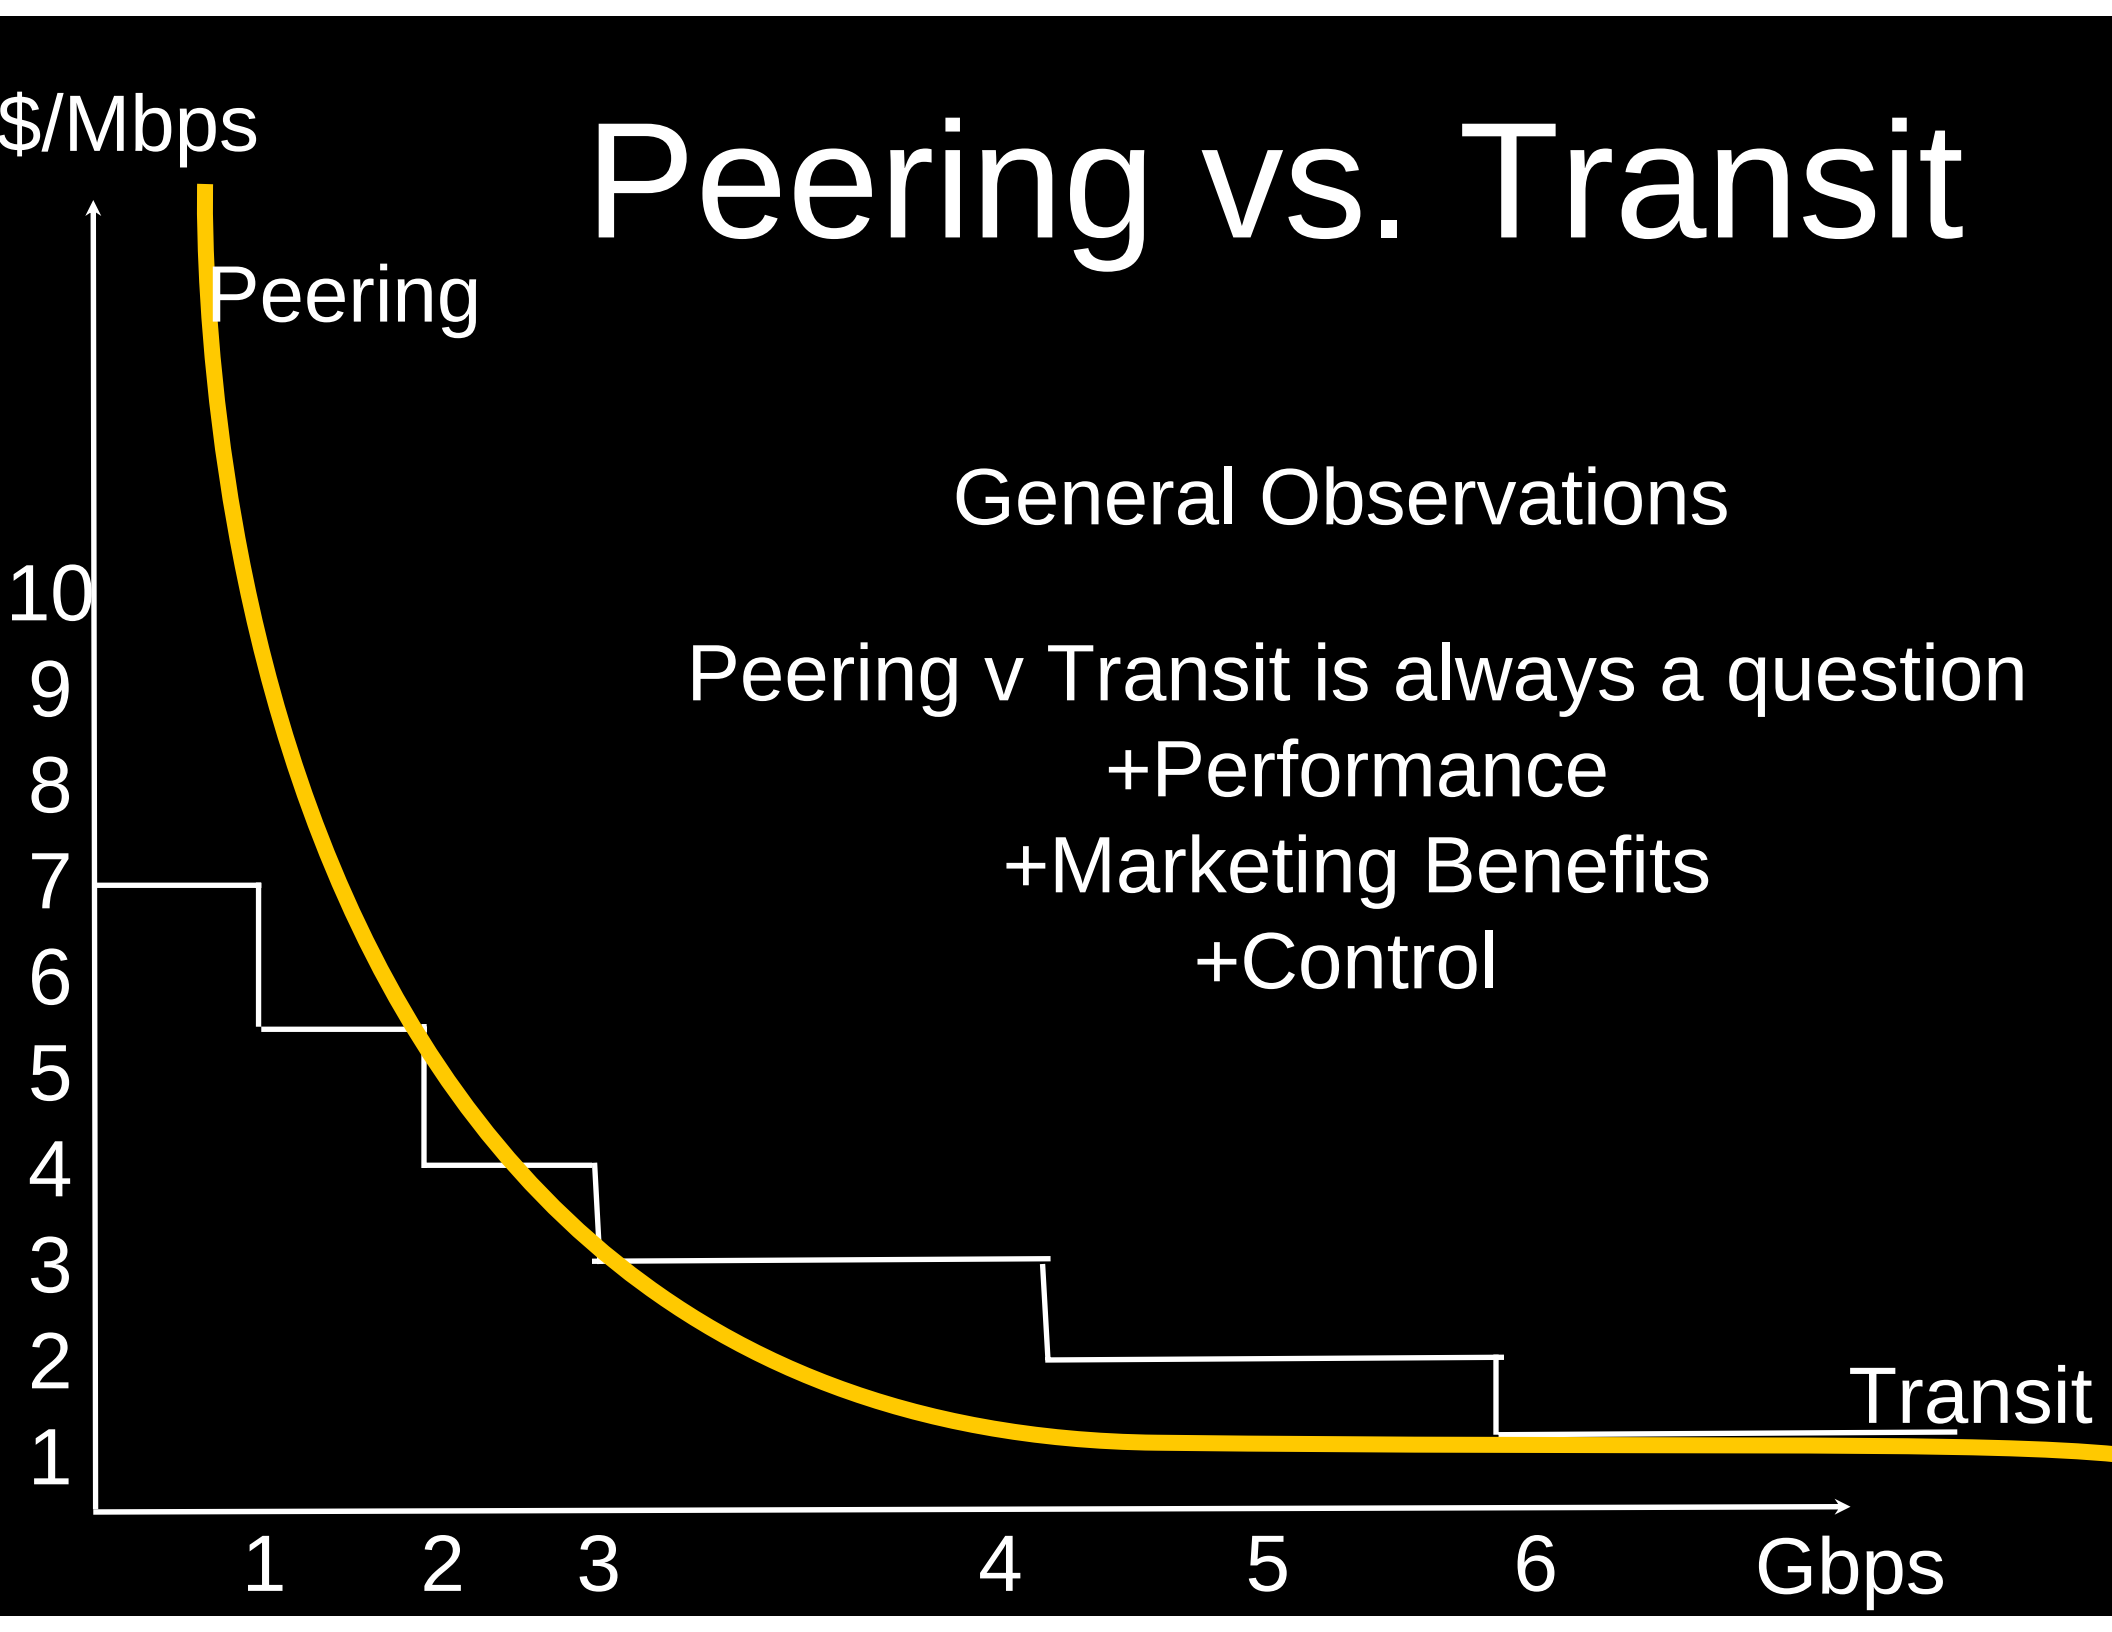
$/Mbps
Peering vs. Transit
Peering
General Observations
Peering v Transit is always a question
+Performance
+Marketing Benefits
+Control
10
9
8
7
6
5
4
3
2
1
Transit
1
2
3
4
5
6
Gbps

# PeeringDB

- How do I contact an ISP for Peering?
- PeeringDB - http://www.peeringdb.com
- Face-to-face
- E-mail
- Phone Calls
- Internet Relay Chat
- Introductions
- IX Operator staff

# Sign into Peering DB

- Browse the pages

# Common Peering Prerequisites

- 24/7 NOC
- Multiple geographically diverse locations
- consistent announcements
- single AS
- traffic volume minimums
- Not be a customer

# Conclusion

- This was an overview (Peering 101)
- concepts and common lexicon
- NANOG is an opportunity to have face-to-face discussions with potential peers
- White Papers on the net:
  - Google ‘william b. norton’