

The Death of Transit and Beyond



Dave Crosby, <https://www.flickr.com/photos/wikidave/4044498586/>

Geoff Huston

Chief Scientist, APNIC

**This presentation is not about
any specific network details**

Or specific plans

Or particular services

Or any particular technology

Or anything like that

It's about architecture

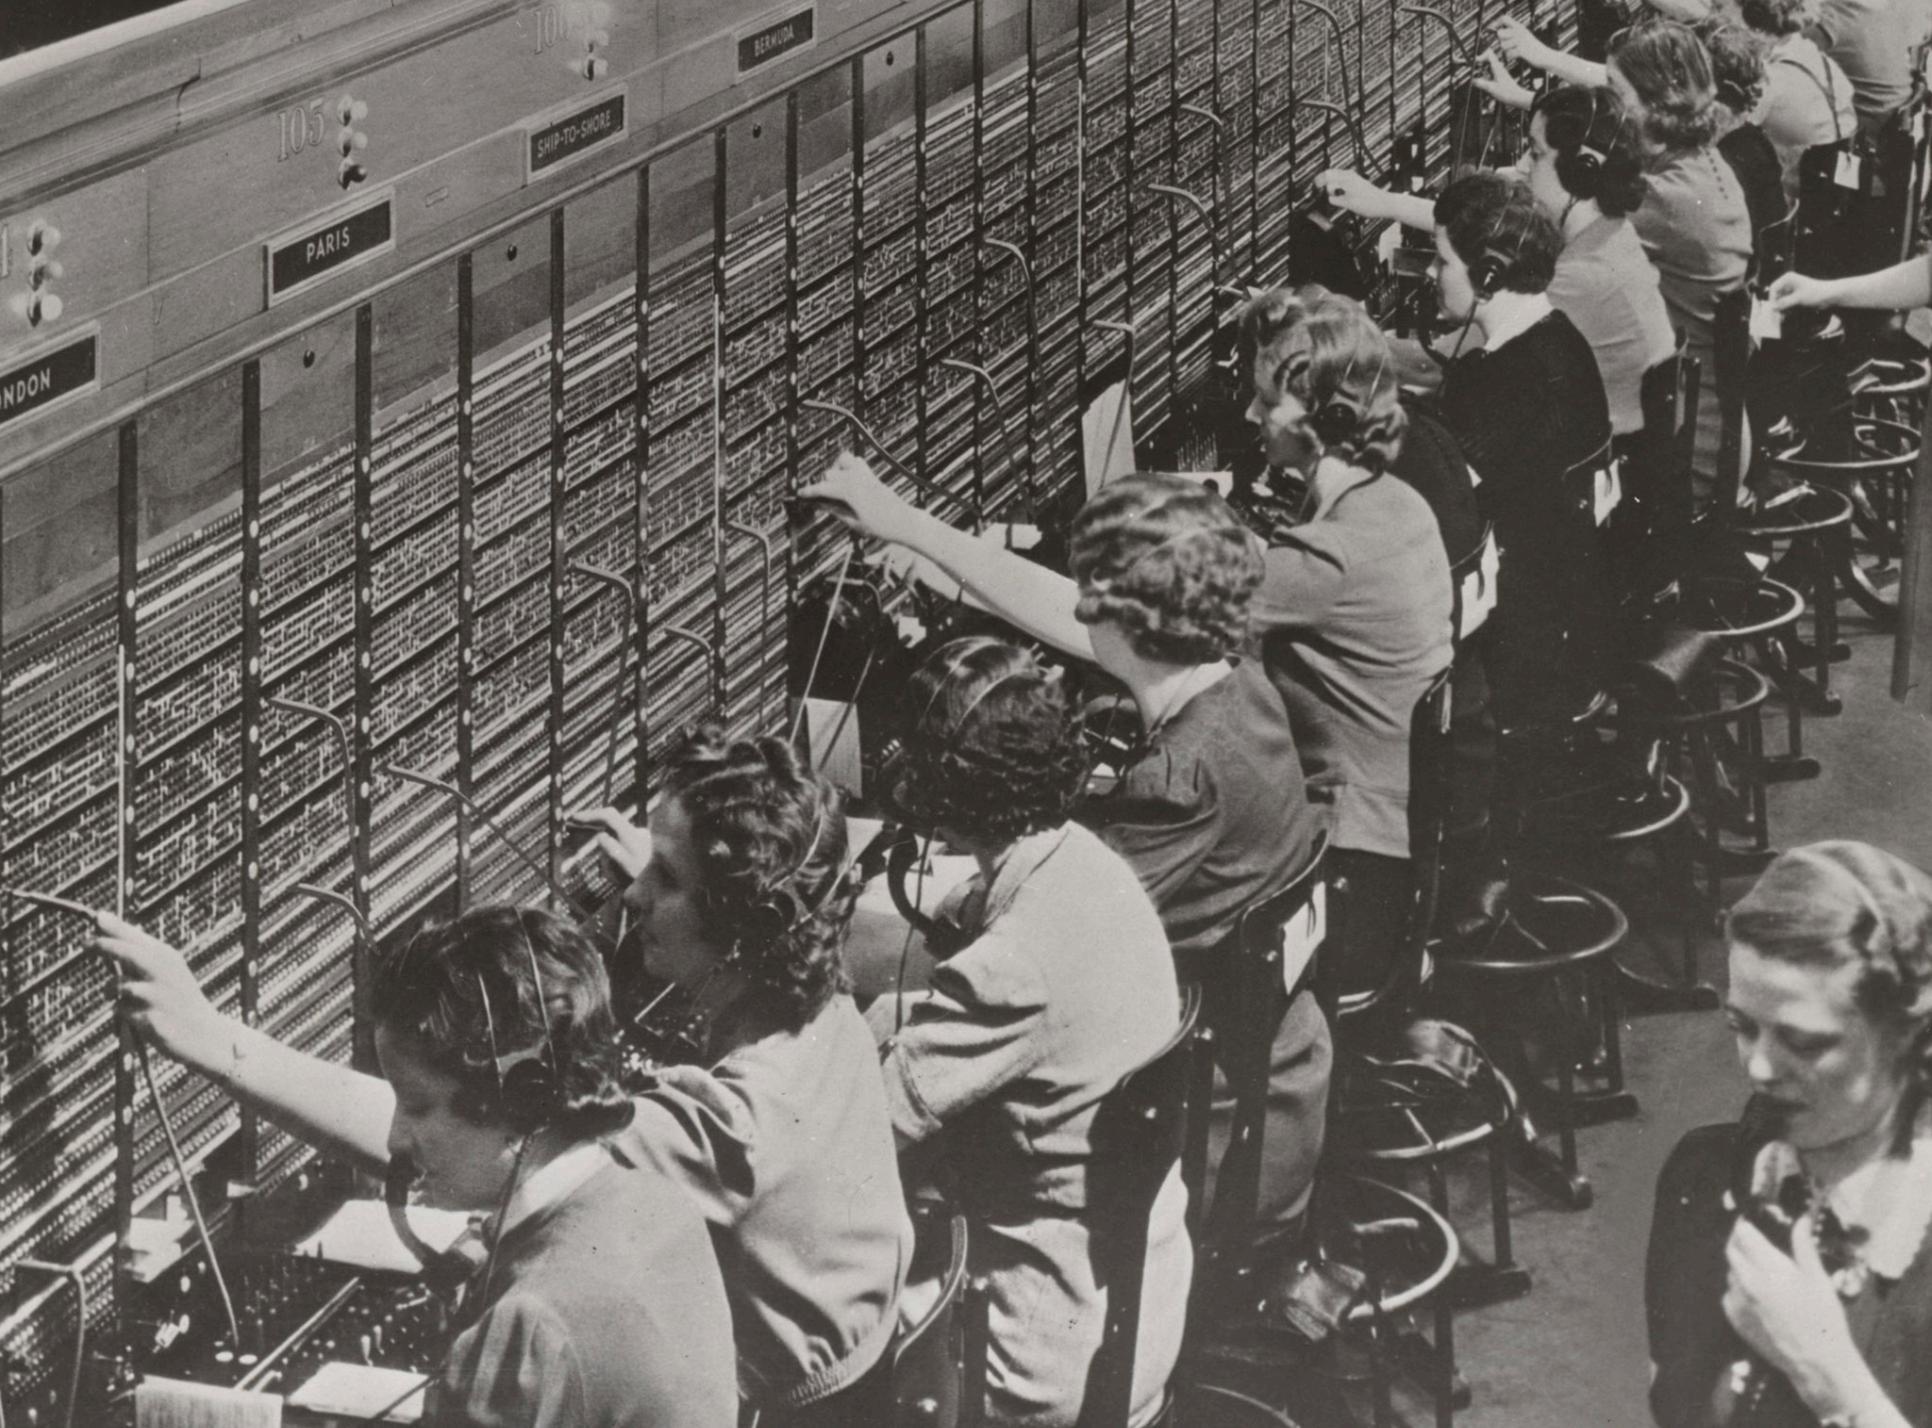
And, in particular, about the evolution of network architecture in the Internet

It's about architecture

And some thoughts about the implications of these changes in terms of public policies for the Internet

Our Heritage

The Telephone Network



Our Heritage

The Telephone Network:

The major technical achievement of the twentieth century

- Connected handsets to handsets
- The network was intentionally transparent
- Real time virtual circuit support between connected edge devices
- Network-centric architecture with minimal functionality in the edge devices

Computer Networks

ON
OFF

HALT
INH

W.D.T.

AUTO
RSTRT

MEMORY
PRCTCT

13

14

15

16

17

18

19

20

21

22

23

24

11

12

IMP STATUS PANEL

INTERFACE MESSAGE PROCESSOR

Developed for
the Advanced Research Projects Agency
by Bolt Beranek and Newman Inc.



T1

T2

T3

T4

F

I

A

C

PI

ML

EA

DP

MP

P

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

RESET

ON

PFI

1

2

3

4

X

OP

B

A

P/Y

M

MASTER

STORE

P

MA

START

OFF

PFH

CLEAR

FETCH

P+1

RUN

S1

POWER

SENSE

REGISTER

OPERATION

Computer Networks

The original concept for computer networks was like the telephone network:

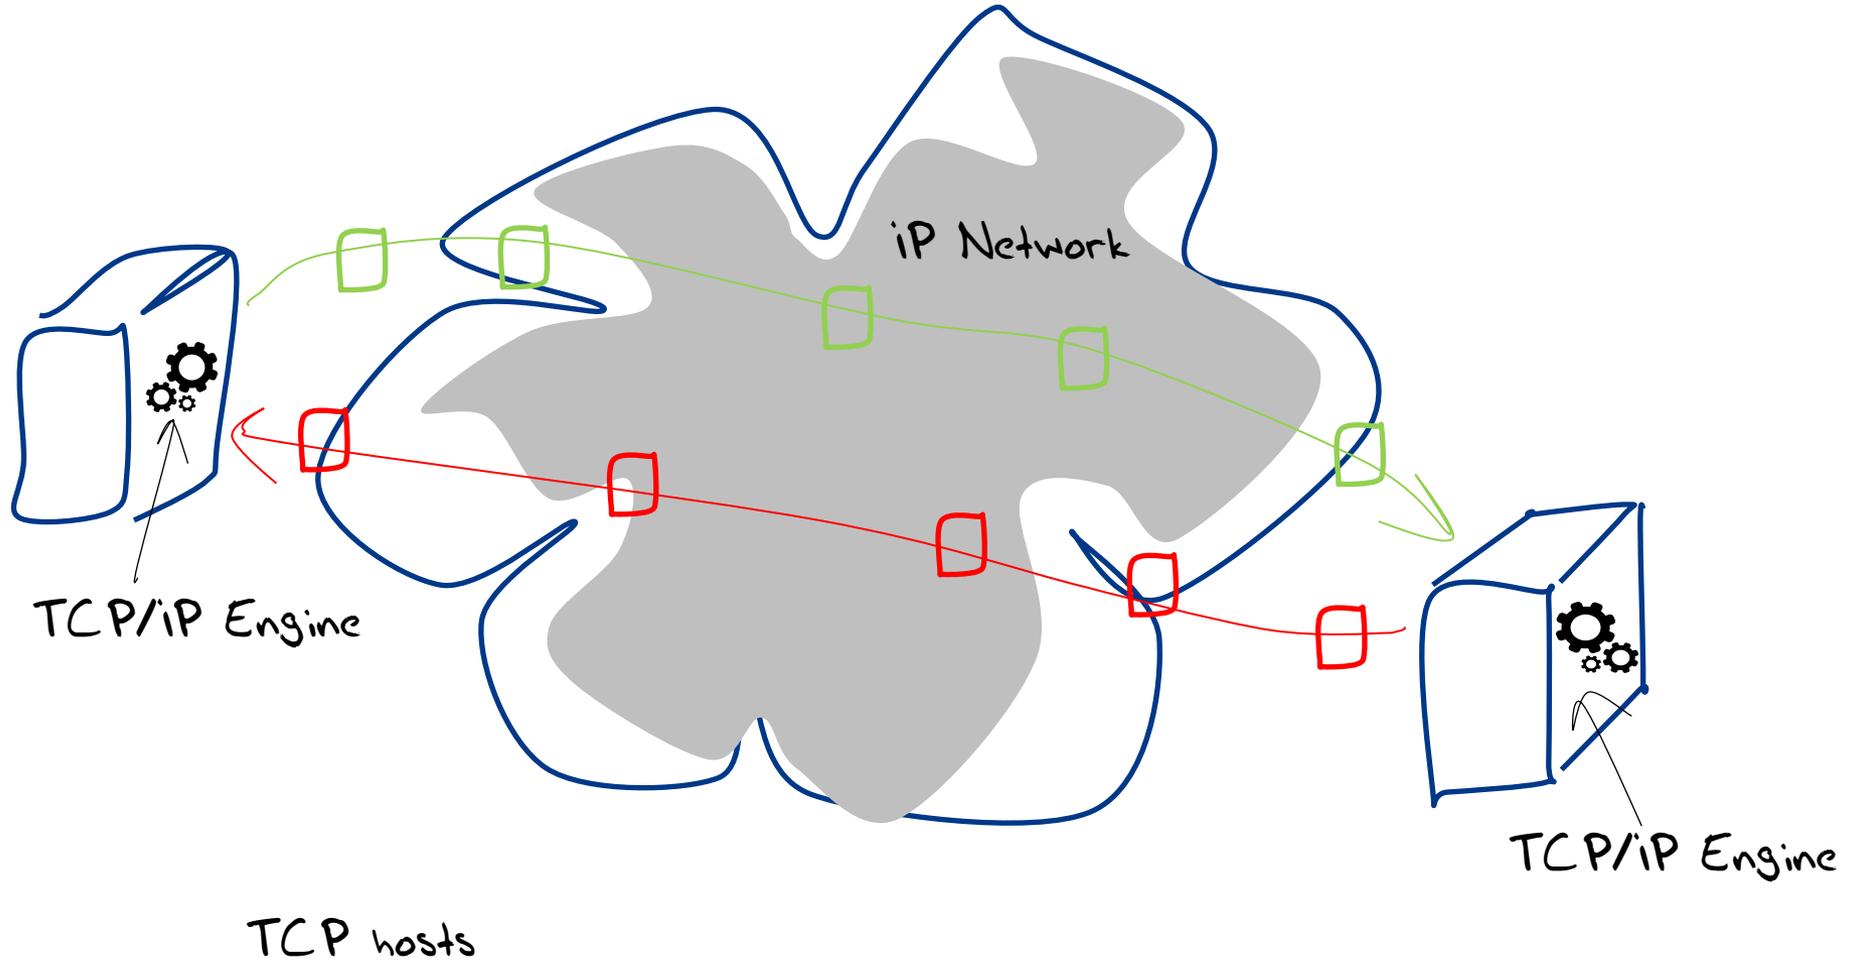
- The network was there to enable connected computers to exchange data
 - All connected computers were able to initiate or receive “calls”
 - A connected computer could not call “the network” – the network was an invisible common substrate
 - It made no difference if the network had active or passive internal elements

Internet Architecture (1980's)

“End-to-End” design:

- Connected computer to computer
- The network switching function was stateless
 - No virtual circuits, no dynamic state for packets to follow
- Single network-wide addressing model
- Single network-wide routing model
- Simple datagram unreliable datagram delivery in each packet switching element
- hop-by-hop destination-address-based packet forwarding paradigm

Internet Architecture (1980's)



The Result was Revolutionary!

- Very Simple
- Extraordinarily Cheap
- Unbelievably Efficient
- Highly Adaptable

[\[Docs\]](#) [\[txt|pdf\]](#) [\[Errata\]](#)

Updated by: [2549](#), [6214](#)

Network Working Group
Request for Comments: 1149

EXPERIMENTAL
Errata Exist
D. Waitzman
BBN STC
1 April 1990

A Standard for the Transmission of IP Datagrams on Avian Carriers

Status of this Memo

This memo describes an experimental method for the encapsulation of IP datagrams in avian carriers. This specification is primarily useful in Metropolitan Area Networks. This is an experimental, not recommended standard. Distribution of this memo is unlimited.

Overview and Rational

Avian carriers can provide high delay, low throughput, and low altitude service. The connection topology is limited to a single point-to-point path for each carrier, used with standard carriers, but many carriers can be used without significant interference with each other, outside of early spring. This is because of the 3D ether space available to the carriers, in contrast to the 1D ether used by IEEE802.3. The carriers have an intrinsic collision avoidance system, which increases availability. Unlike some network technologies, such as packet radio, communication is not limited to line-of-sight distance. Connection oriented service is available in some cities, usually based upon a central hub topology.

Frame Format

The IP datagram is printed, on a small scroll of paper, in hexadecimal, with each octet separated by whitestuff and blackstuff. The scroll of paper is wrapped around one leg of the avian carrier. A band of duct tape is used to secure the datagram's edges. The bandwidth is limited to the leg length. The MTU is variable, and paradoxically, generally increases with increased carrier age. A typical MTU is 256 milligrams. Some datagram padding may be needed.

Upon receipt, the duct tape is removed and the paper copy of the datagram is optically scanned into a electronically transmittable form.

Discussion

Multiple types of service can be provided with a prioritized pecking order. An additional property is built-in worm detection and eradication. Because IP only guarantees best effort delivery, loss of a carrier can be tolerated. With time, the carriers are self-

[\[Docs\]](#) [\[txt|pdf\]](#) [\[Errata\]](#)

Updated by: [2549](#), [6214](#)

Network Working Group
Request for Comments: 1149

A Standard for the Transmission

Status of this Memo

This memo describes an experiment with
IP datagrams in avian carriers. It is
useful in Metropolitan Area Network
recommended standard. Distributed

Overview and Rationale

Avian carriers can provide high altitude service. The connection point-to-point path for each carrier but many carriers can be used with each other, outside of early spring space available to the carriers, IEEE802.3. The carriers have an antenna system, which increases availability technologies, such as packet radio line-of-sight distance. Connect some cities, usually based upon

Frame Format

The IP datagram is printed, on a hexadecimally wrapped paper. A band of duct tape is used to seal the bandwidth is limited to the leg paradoxically, generally increases typical MTU is 256 milligrams.

Upon receipt, the duct tape is removed and the datagram is optically scanned in form.

Discussion

Multiple types of service can be



RFC-1149

April 28th, 2001

11 years ago, April 1st 1990, rfc 1149 was written. This rfc specifies a protocol for IP over avian carriers, CPIP (carrier pigeon internet protocol). In 11 years, no one has bothered to implement this important protocol stack. But happily, we don't need to wait any longer! [BLUG](#) in cooperation with [Vesta Brevdueforening](#) has given you rfc 1149 support for Linux.

Vital information

- Date: April 28 2001, 12:00
- Place: Bergen
- Addresses:
Lyngbøveien 61, contact: Audun Larsen.
Bråtet Terrasse 21, contact Kjell Haldorsen.

Details from the actual implementation day.

- [A preliminary writeup of the event.](#)
- [Log of the ping session](#)
- [Vegards pictures.](#)
- [Bjørns pictures.](#)
- [Karl Magnus' pictures.](#)

<http://www.blug.linux.no/rfc1149>

The Result was Revolutionary!

By stripping out network-centric virtual circuit states and removing time synchronicity the resultant carriage network was minimal in design and functionality

More complex functions, such as flow control, jitter stability, loss mitigation and reliability, were pushed out to the computers on the edge

Internet Devolution

In the regulated world of national telephone operators every telephone network was “equal”

But we rapidly started differentiating between Internet networks. Internet networks were not all the same.

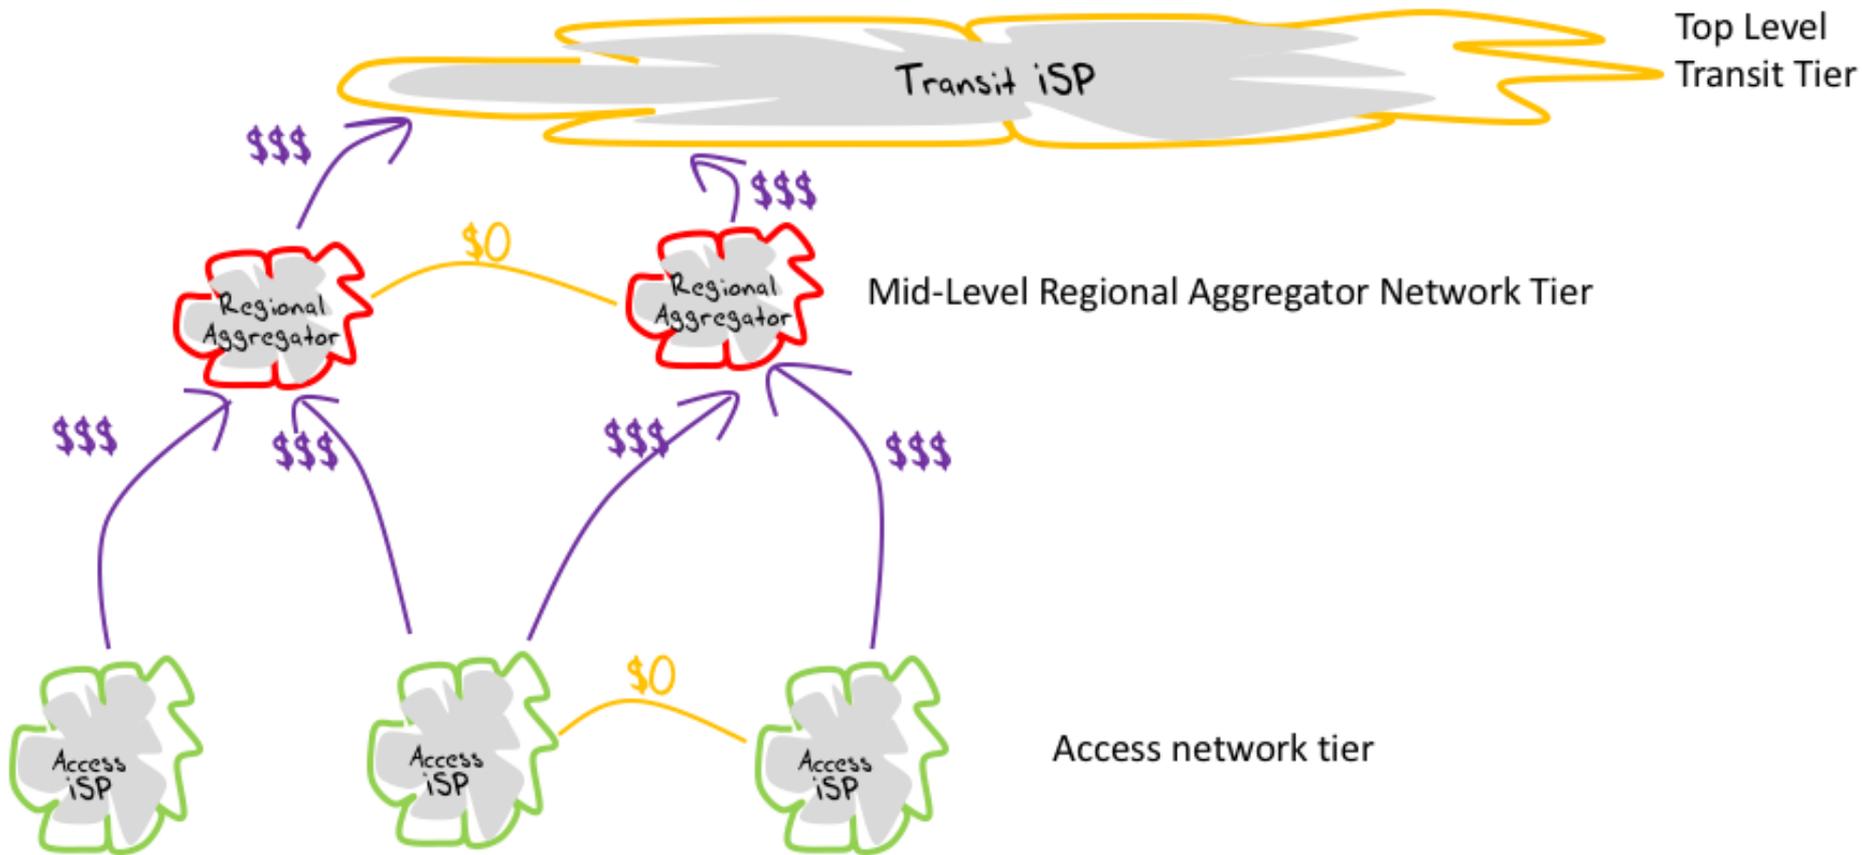
We started differentiating on roles and services and started differentiating by the flow of revenues between networks

Internet Devolution

Financial considerations of the evolving commercial Internet introduced structure of the Network Provider interaction

- Role specialization between **access networks** that serviced connection of edge devices and networks and **transit networks** that serviced interconnection of other networks
- Limited forms of financial settlement in packet networks reduced interaction to either **SKA peering** or **upstream Provider / downstream Customer**

Network Role Segmentation



Edge Role Segmentation

Breaking the edge into **clients** and **servers**

- Access networks service the needs “clients”
- Clients are not directly reachable by other clients
- Clients connect to services

The role of the network here is to carry clients to the service access point

- The assumption here is that there are many more clients than service points
- Clients pay the network for this carriage role

Content vs Carriage

Who pays whom?

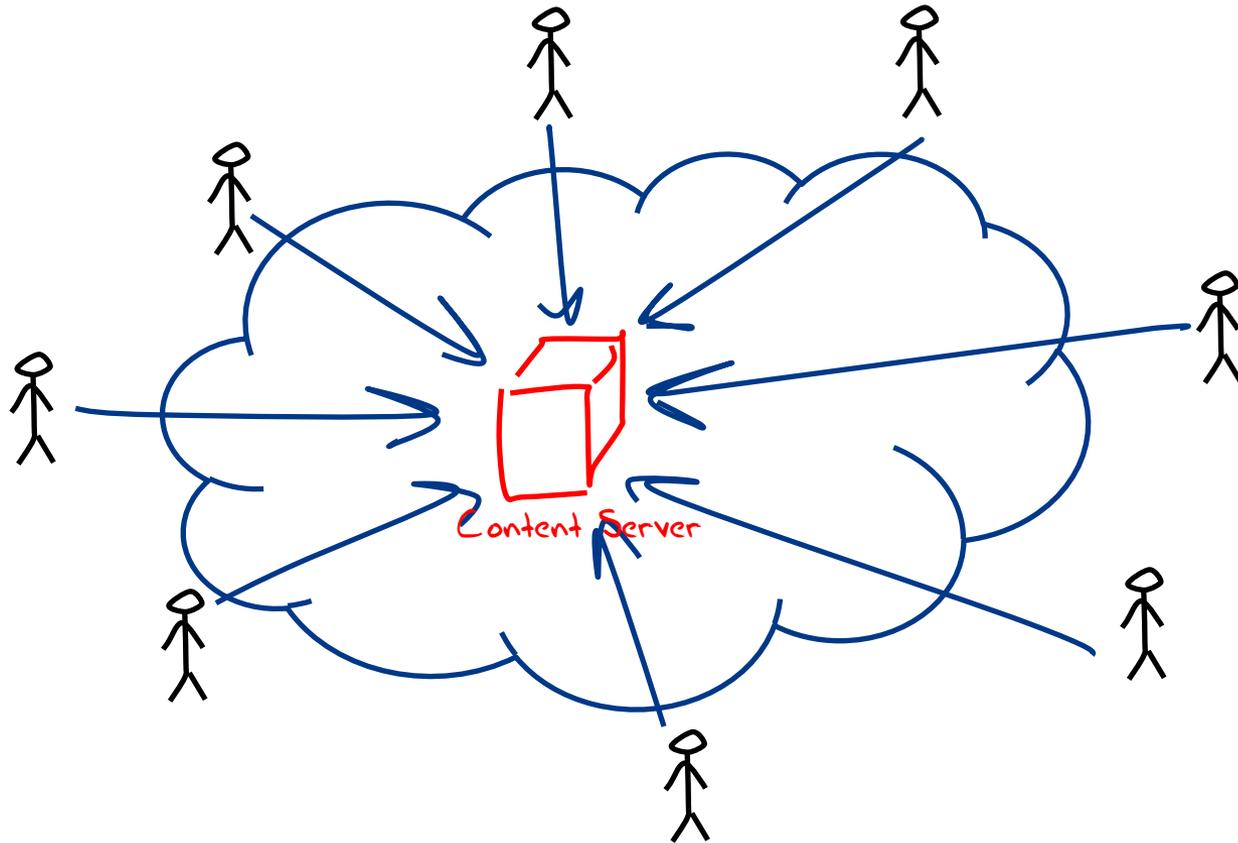
- The only reason why access networks have clients is because there are content services that clients want to access
 - Therefore carriage should pay for content
- There is no “end-to-end” financial settlement model in the Internet – both “ends” pay for access and network providers settle between themselves. To a carriage network, content is just another client
 - Content should pay for carriage, just like any other client

Content vs Carriage

Who pays whom?

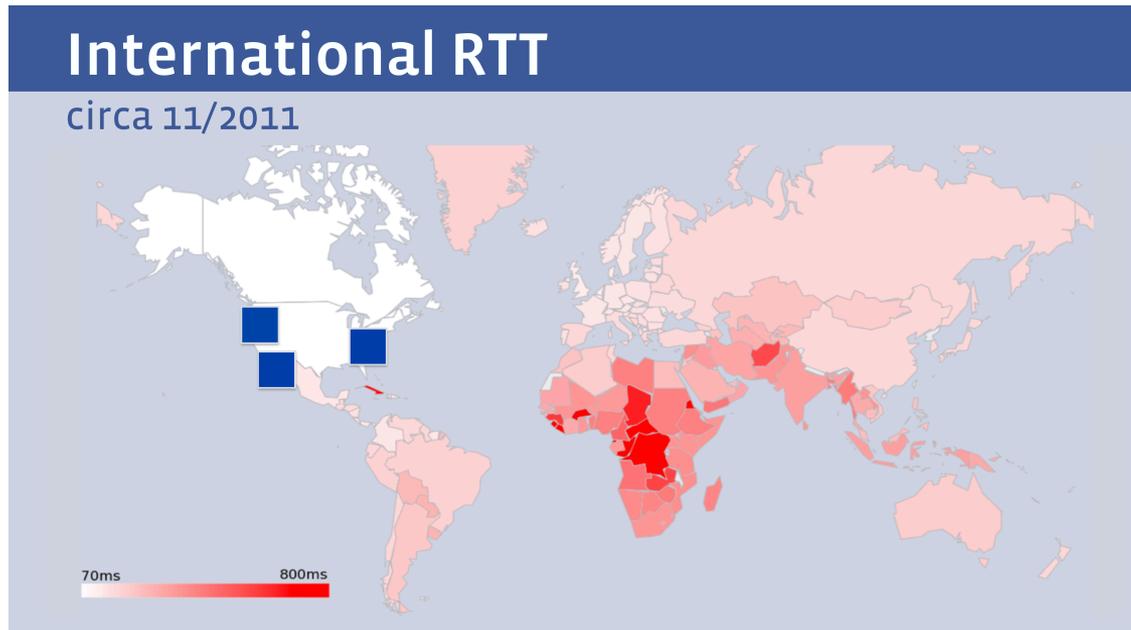
- The only reason why access networks exist is because there are no other ways to access the Internet.
The content folk resolved this by going "over the top" and created relationships directly with end users
- There is no "end-to-end" financial settlement model in the Internet – both "ends" pay for access and network providers settle between themselves. To a carriage network, content is just another client
 - Content should pay for carriage, just like any other client

Content Server



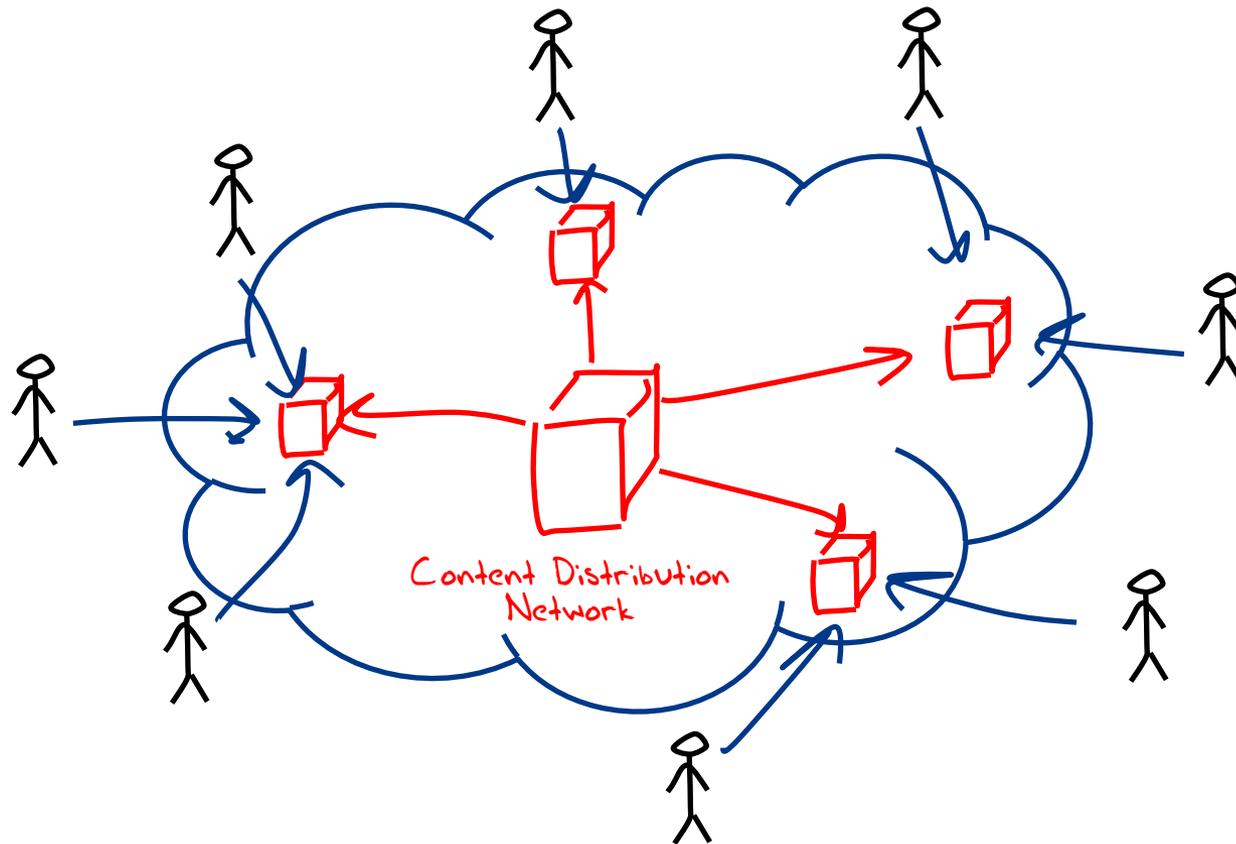
The Tyranny of Distance

But not all clients enjoy the same experience from a single service



*Facebook presentation at
NANOG 68*

Content Distribution



Let them eat data!

The rise of the Content Distribution Network

- Replicate content caches close to large user populations
- The challenge of delivering many replicant service requests over high delay network paths is replaced by the task of updating a set of local caches by the content distribution system and then serving user service requests over the access network
- Reduced service latency, increased service resilience

Role Reversal

Service portals are located adjacent to users

- Networks no longer carry users' traffic to/from service portals as ISP public carriage services
- Networks carry content to service portals as CDN private carriage services

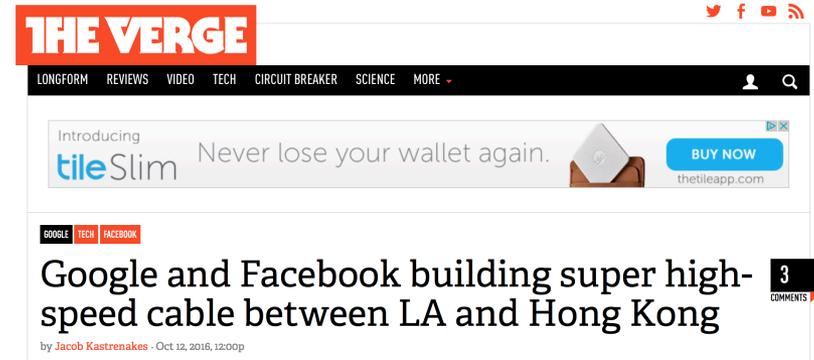
This shift has some profound implications for the Internet

Who's building now?

Almost all new submarine international cable projects are heavily underwritten by content providers, not carriers

Large content providers have huge and often unpredictable traffic requirements, especially among their own data centers. Their capacity needs are at such a scale that it makes sense for them, on their biggest routes, to build rather than to buy. Owning subsea fibre pairs also gives them the flexibility to upgrade when they see fit, rather than being beholden to a third-party submarine cable operator.”

Tim Stronge of Telegeography, January 2017

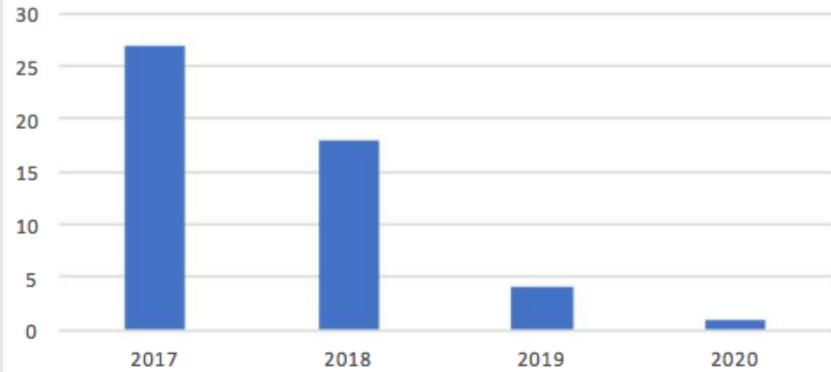


The screenshot shows the top portion of a web page from The Verge. The site's logo is in the top left, and navigation links for 'LONGFORM', 'REVIEWS', 'VIDEO', 'TECH', 'CIRCUIT BREAKER', 'SCIENCE', and 'MORE' are in a dark bar. A search icon and a user profile icon are on the right. Below the navigation is a banner for 'tileSlim' with the text 'Introducing tileSlim Never lose your wallet again.' and a 'BUY NOW' button. The main article headline reads 'Google and Facebook building super high-speed cable between LA and Hong Kong' with a '3 COMMENTS' badge. The byline is 'by Jacob Kastrenakes - Oct 12, 2016, 12:00p'.

Submarine Cables

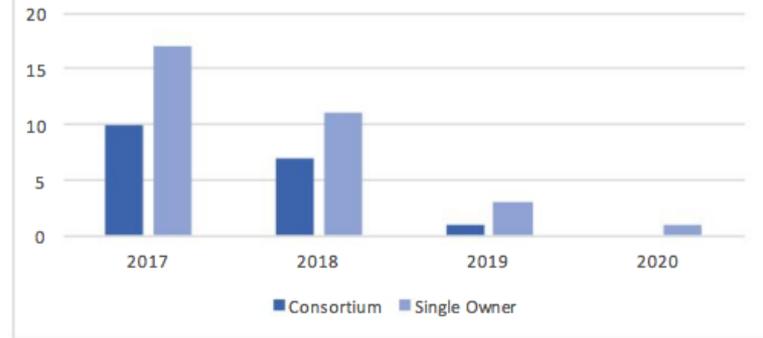
And those that are being built are now single owner cables

Systems Announced RFS
2017-2020

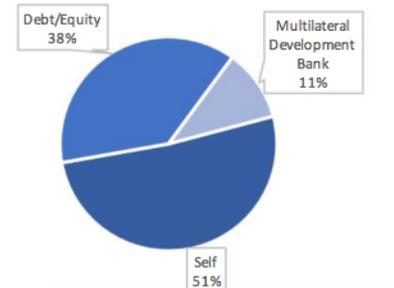


Fewer cables being built

System Ownership Type
2017-2020



System Financing Type
2017-2020



And the majority are now self-funded

Today's Internet Architecture

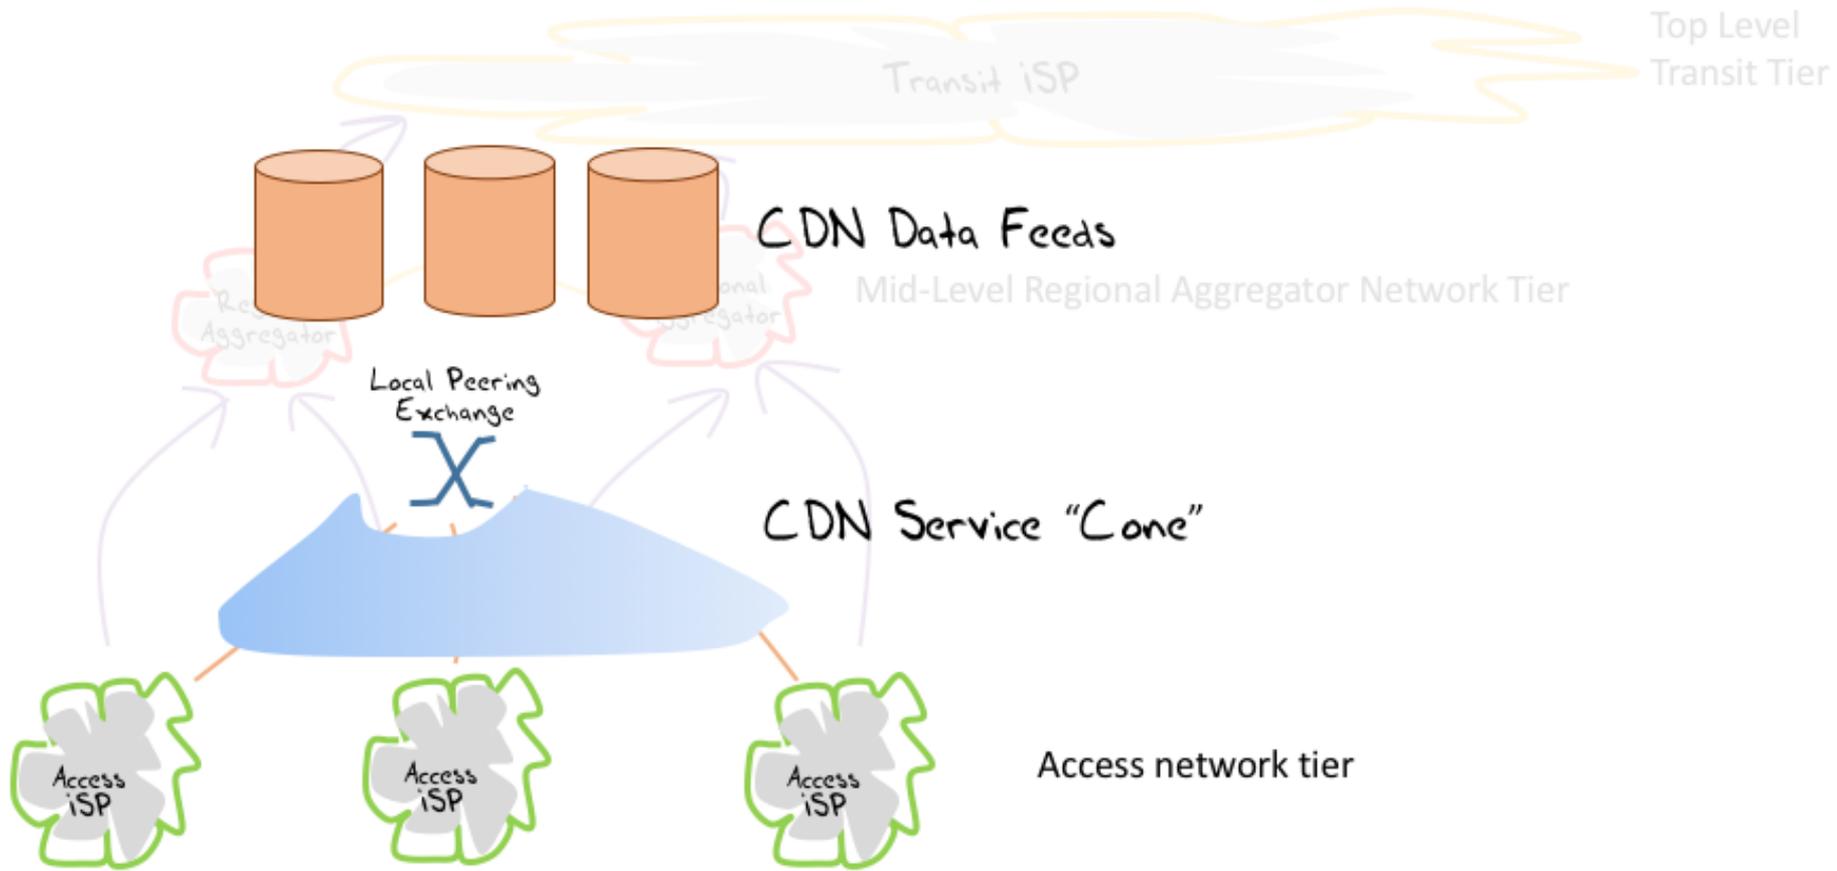
We've split the network into clients and servers

- Web servers
- Streaming servers
- Mail servers
- DNS servers

Servers and services now sit in CDN systems with global replication and DDOS resilience

Users don't reach out to content any more - the CDNs bring content to users

Today's Internet Architecture



Transit?

- If users don't send packets to users any more...
- If content is now delivered via CDNs to users via discrete service cones...
- If there is no universal service obligation...

Then why do we still need Transit Service providers?

Transit?

- Once the CDN caches sit “inside” the Edge NAT of the Access ISP then the entire wide area network becomes a marginal activity compared to the value of the content feeds!

Internet Names and Addresses?

If the Internet is (or maybe soon will be) a collection of discrete CDN service 'cones' then why do we still need :

- A global address plan (in IPv4 or IPv6)?
- A global name system?
- A single global network?

It's not just the Death of Transit ...

It's the re-purposing of the entire network

- Service provisioning sits within cloud providers and distributed data centres
- Applications that use peer-to-peer networking are now under general suspicion of dark deeds of IPR theft
- Edge computers are now acting as televisions into the clouded world of data
- The distinction between personal and public data realms is disappearing into the realm of corporately owned private data empires

Exactly where are we?

- We started this journey building a telephone network for computers to communicate between each other
- One-way content distribution lies at the core of today's Internet
- We are now far closer to a model of broadcast television or some similar form of video / data distribution
- This content distribution role is an enterprise model rather than a public service
- The internal parts of the network are now being privatized and removed from public regulatory scrutiny

Policy?

If CDN networks are private networks, and there is little residual public carriage other than last mile access networks, then what do we really mean by “public communications policy”?

In the regulatory world ‘content’ is **commerce**, not **carriage**!

Policy?

In today's Internet what do we mean in a policy sense by concepts such as:

“universal service obligation”

“network neutrality”

“rights of access” or even

“market dominance”

when we are talking about diverse CDNs as the dominant actors in the Internet?

Content is Aggregating

- There are not thousands of content service platforms
 - There are just a few hundred
- The space is dominated by a small number of massive actors who set the environment for all others

Content Consolidation

"The size and scale of the attacks that can now easily be launched online make it such that if you don't have a network like Cloudflare in front of your content, and you upset anyone, you will be knocked offline.

...

In a not-so-distant future, if we're not there already, it may be that if you're going to put content on the Internet you'll need to use a company with a giant network like Cloudflare, Google, Microsoft, Facebook, Amazon, or Alibaba.

...

Without a clear framework as a guide for content regulation, a small number of companies will largely determine what can and cannot be online.

<https://blog.cloudflare.com/why-we-terminated-daily-stormer/>

The Large and the Largest

	Company	\$B USD
	Apple	749
	Alphabet	628
	Microsoft	528
	Amazon	466
	Berkshire Hathaway	418
	Johnson & Johnson	357
	Facebook	357
	Alibaba Group	356
	Tencent	344
	Exxon Mobil	341

The world's 10 largest publicly traded companies, as ranked by their market valuation, June 2017

Content is King

- None of these seven technology companies are a telephone company, or even a transit ISP, or even an ISP at all!
- All of them have pushed aside carriage networks in order to maintain direct relationships with billions of consumers
- These valuable consumer relationships are based on content services, not carriage

Competition or Cartel?

With a small number of truly massive enterprises at the heart of the area of digital content and service is this still a space that is shaped by competitive pressures?

Or do these dominant incumbents get to set their own terms of engagement with each other and with users?

We've been here before..

We've been here before..



American Art: The Gilded Age

Mark Twain coined the phrase "the Gilded Age" in 1873. This term, with its connotations of superficiality and ostentatious wealth, has come to refer to the decades following the Civil War. During that period of rapid industrialization, the contrast between the lifestyles of so-called robber barons and average workers was enormous. The metaphor of gilded surfaces resonates in the richly decorated possessions of the ruling class, from domestic furniture to picture frames.

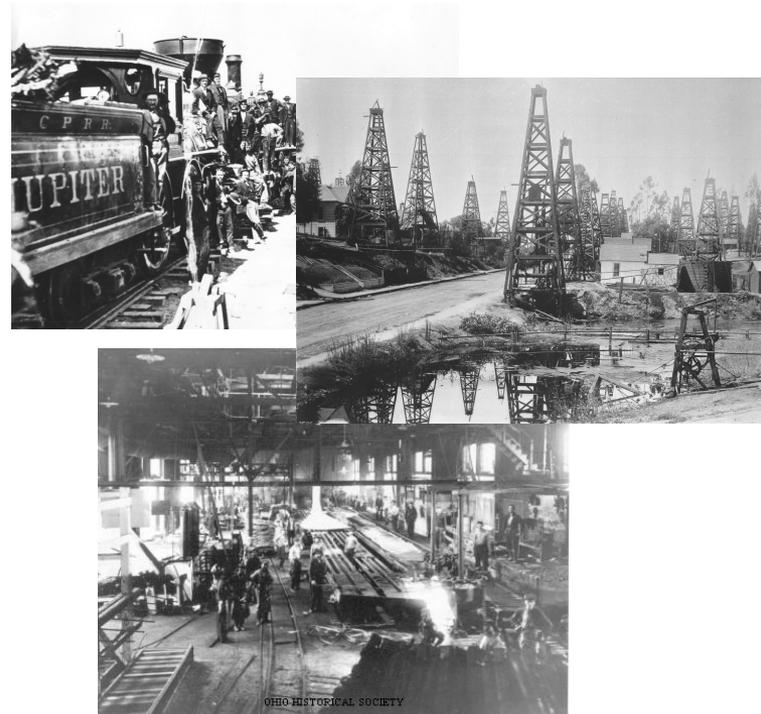
This gallery examines the leading cultural phenomenon of the 1870s and 1880s, the American Aesthetic movement, through a range of objects produced for affluent consumers. Aestheticism, rooted in the English philosophies of John Ruskin and William Morris, advanced the notion that a beautiful environment could promote moral and social reform. In the process, the Aesthetic movement helped to liberate American art and design from the confines of historicism by admitting fresh influences from foreign lands.

High Museum of Art, Atlanta

The Gilded Age

A term applied to America in the 1870 – 1890's about the building of industrial and commercial corporate giants on platforms that were a mix of industrial innovation and enterprise with elements of greed, corruption and labor exploitation

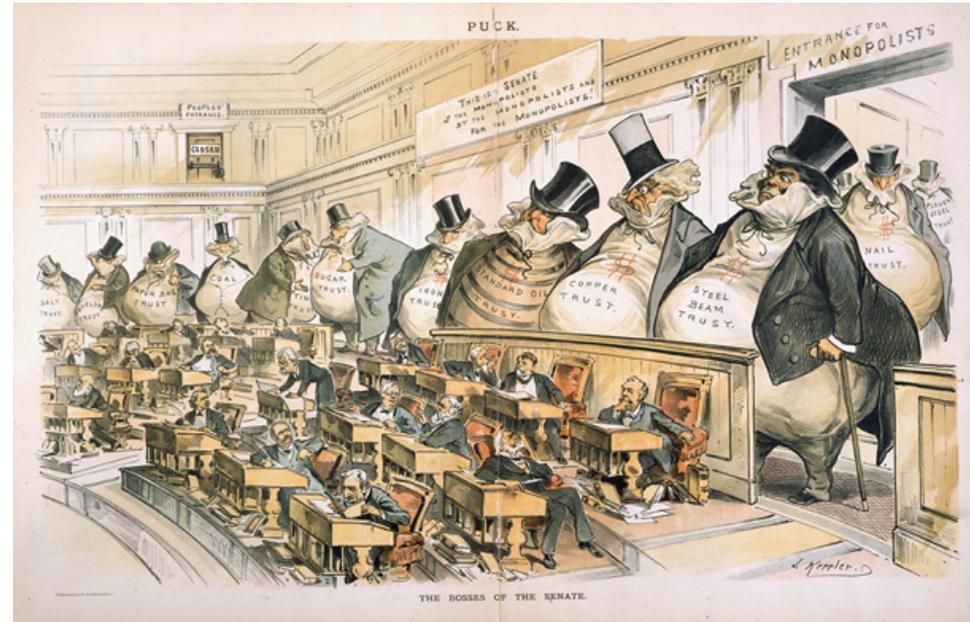
Andrew Carnegie - US Steel
John Rockefeller - Standard Oil
Theodore Vail - AT&T
George Westinghouse – Rail Brakes
Thomas Edison – General Electric
J P Morgan - Banking



The Gilded Age

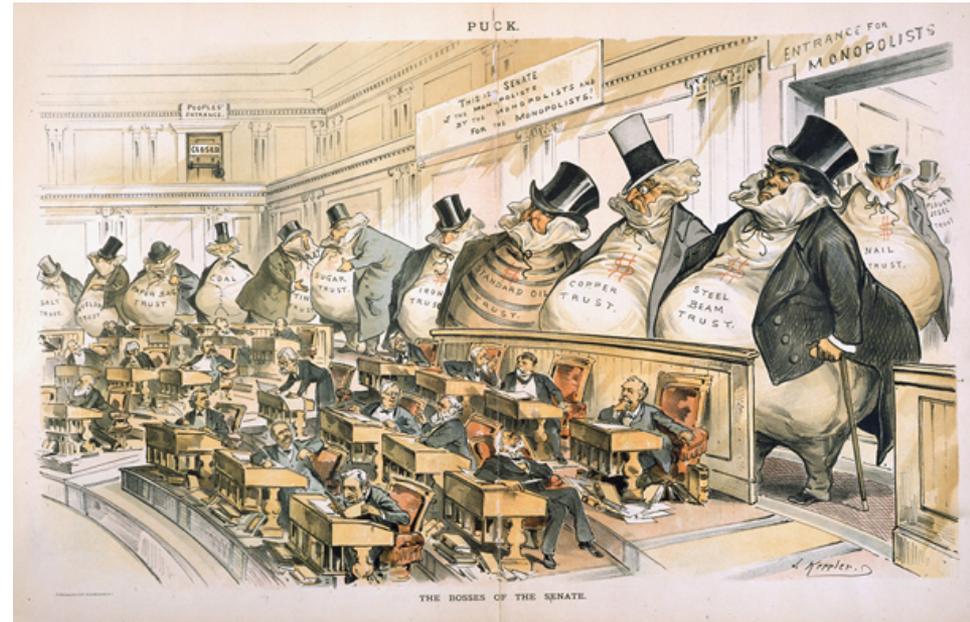
During this period in the United States the dominant position within industry and commerce was occupied by a very small number of players who were moving far faster than the regulatory measures of the day.

The resulting monopolies took the US decades to dismember, and even today many of these gilded age companies are dominant in their field



The Internet's Gilded Age

At some point in the past decade or so the dominant position across the entire Internet has been occupied by a very small number of players who are moving far faster than the regulatory measures that were intended to curb the worst excesses of market dominance by a small clique of actors.



Who's Gilding?

	Company	\$B USD
	Apple	749
	Alphabet	628
	Microsoft	528
	Amazon	466
	Berkshire Hathaway	418
	Johnson & Johnson	357
	Facebook	357
	Alibaba Group	356
	Tencent	344
	Exxon Mobil	341

The Internet's Gilded Age

These actors have enough market influence to set their own rules of engagement with:

- Users,
- Each other,
- Third party suppliers,
- Regulators and Governments

By taking a leading position with these emergent technologies, these players are able to amass vast fortunes, with little in the way of accountability to a broader common public good

The Internet's Gilded Age

These actors have enough market influence to set their own rules of engagement with:

- Users,
- Each other,
- Third

is this the internet we were dreaming of?

By taking a leading position with these emergent technologies, these players are able to amass vast fortunes, with little in the way of accountability to a broader common public good

What is this all about?

This is no longer just a conversation about changes in carriage and communications within the Internet.

It is probably not even a conversation about carriage and communications at all.

What is this all about?

The changing face of the Internet is no longer a matter of public communications, but a matter of public services.

And with this observation we are back to a more basic theme...

What is this all about?

The essential topic of this conversation is how we can strike a sustainable balance between a rapacious private sector that has amassed overarching control of the digital service and content space, and the needs of the larger society in which we all would like some equity of opportunity to thrive and benefit from the outcomes of this new digital age.

Thanks!