

The Rise and Rise of CDNs

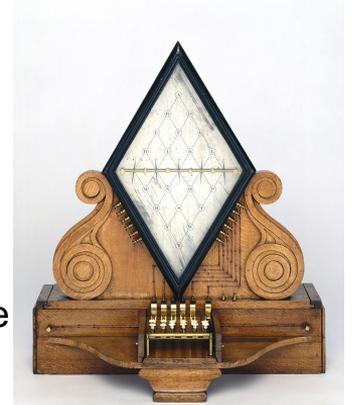
Geoff Huston AM

Chief Scientist, APNIC

Where have we come from?

The Victorian Internet

- In 1800 **Alessandro Volta** invented the battery that allowed electricity to be used in a controlled manner
- In 1820 **Hans Christian Oersted** demonstrated the connection between electric current and magnetism
- In the 1830's in the UK **William Cooke** and **Charles Wheatstone** used a system of five pointers to send text - the first use was railway signalling in the UK
- In the 1840's **Samuel Morse** developed a simpler system using a keypad to complete a circuit.
- By 1861 telegraph lines spanned the US
- By 1870 an undersea cables spanned the Atlantic



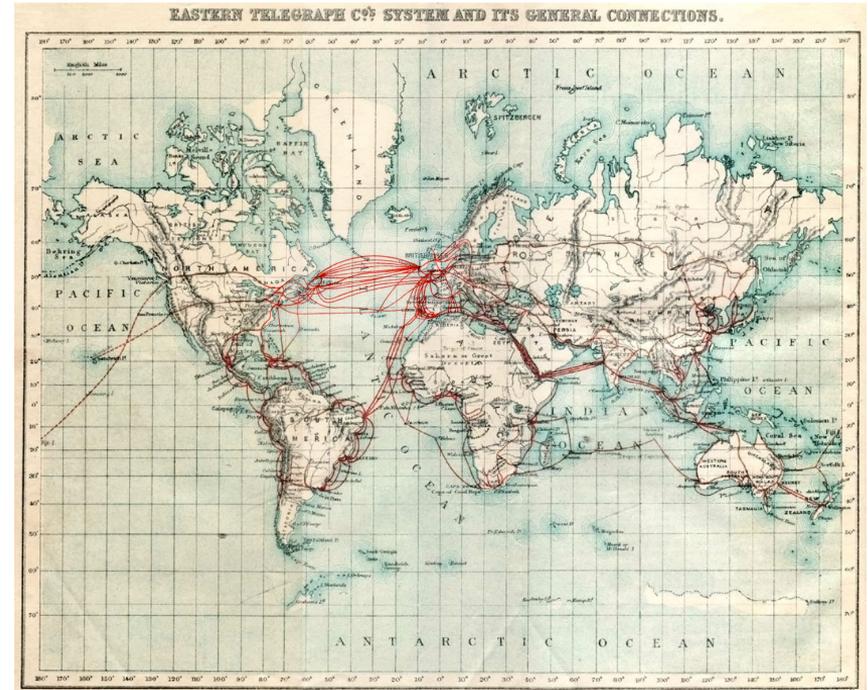
The Great Telegraph Boom

The period from the 1850's to the 1900 saw major investments in national and international telegraph cable systems

Most of the initial international investment activity was from the UK – by 1882 British companies owned and operated two thirds of the world's telegraph cables.

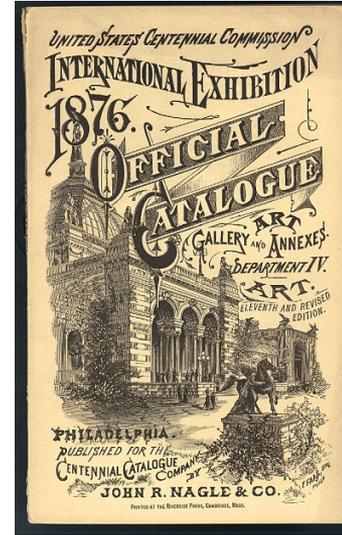
In the US newspapers expanded 5-fold in the period 1840 – 1860 as 50,000 miles of telegraph cable were installed

When combined with the railway this became an effective means for the projection of power and control – enterprises saw opportunities in extensive reach, creating private monopolies to complement the older state-sponsored monopolies



The Next Wave: the Telephone

- First shown to the world at the 1876 World Exposition at Philadelphia, its invention triggered a struggle to the death between Western Union's telegraph and Bell's telephone
- Western Union never fully appreciated that the telephone was an existential threat to the telegraph until it was simply too late.
- Thousands of regional telephone companies appeared all over the world in the following years



The Formation of the Telephone Cartel

- Theodore Vail - President of American Telephone and Telegraph (twice!) - oversaw the construction of a national monopoly masquerading as a public utility through the Kingsbury Commitment with US Congress in 1913
- AT&T divested itself of Western Union Telegraph and in return created a substantial private monopoly under the catch cry of “one policy, one system and one universal service”
- Other countries emulated this transformation from competition to national monopoly in just a few years, using existing telegraph monopoly to subsume telephone operators into public utility structures



And then the telephone story stopped

- The telephone network was constructed to match the human voice
 - And that's all it did
- Telephone service innovations were deployed early on in the story of telephone
 - Time Division Multiplexing was developed in 1870's for telegraphy (later adapted for digital telephony)
 - The Stroger Automatic Exchange was invented in 1891
 - The Fax dates back to 1843
 - And there was the Speaking Clock of course
 - And not much else!
- While telephone companies monopolized technology, they did so in a defensive stance to maintain their hold on telephony services
- Meanwhile, the phone companies' research labs invented the transistor, switching, DACs, packet networking, ciphers, radio astronomy, television, solar cells, C, Unix, microprocessors

Meanwhile...

Computer Networks

ON OFF
HALT INH
W.D.T.
AUTO RSTRT
MEMORY PRCT

1 2 3 4 5 6 7 8 9 10 11 12
13 14 15 16 17 18 19 20 21 22 23 24

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 OFF
PF1 PFH
POWER

1 2 3 4
SENSE

X B A P/Y M
OP
REGISTER

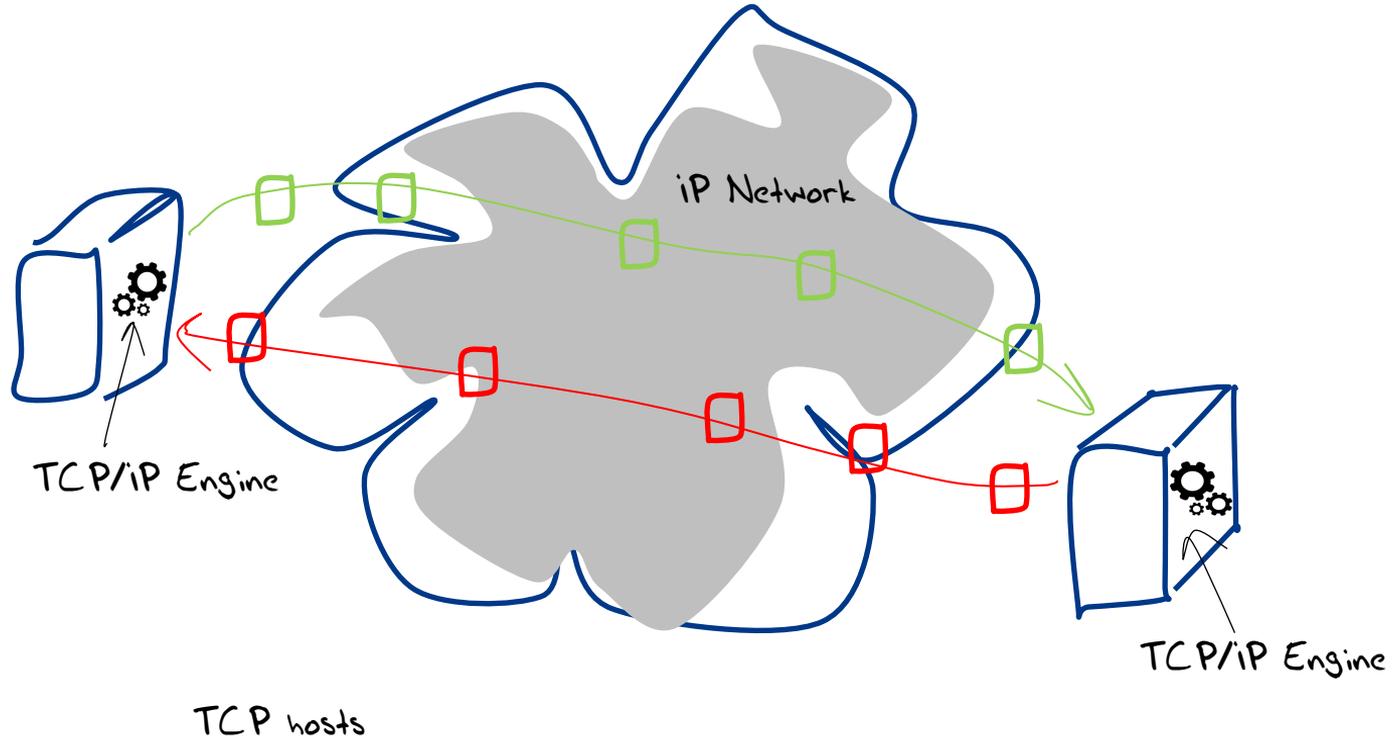
MASTER CLEAR STORE FETCH P P+1 MA RUN S1 START
OPERATION

Computer Networks

The original conceptual model for computer networks was the telephone network

- The network was there to enable connected computers to exchange data
 - All connected computers were able to initiate or receive “calls” to/from other connected computers
 - 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

1980's Internet Architecture



1980's Internet Architecture

“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

The Result was Revolutionary!

By stripping out network-centric virtual circuit states and removing time synchronicity the resultant packet 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 attached devices on the edge

Packet carriage networks were far cheaper to operate than circuit-switched networks!



(in) Equality of Networks

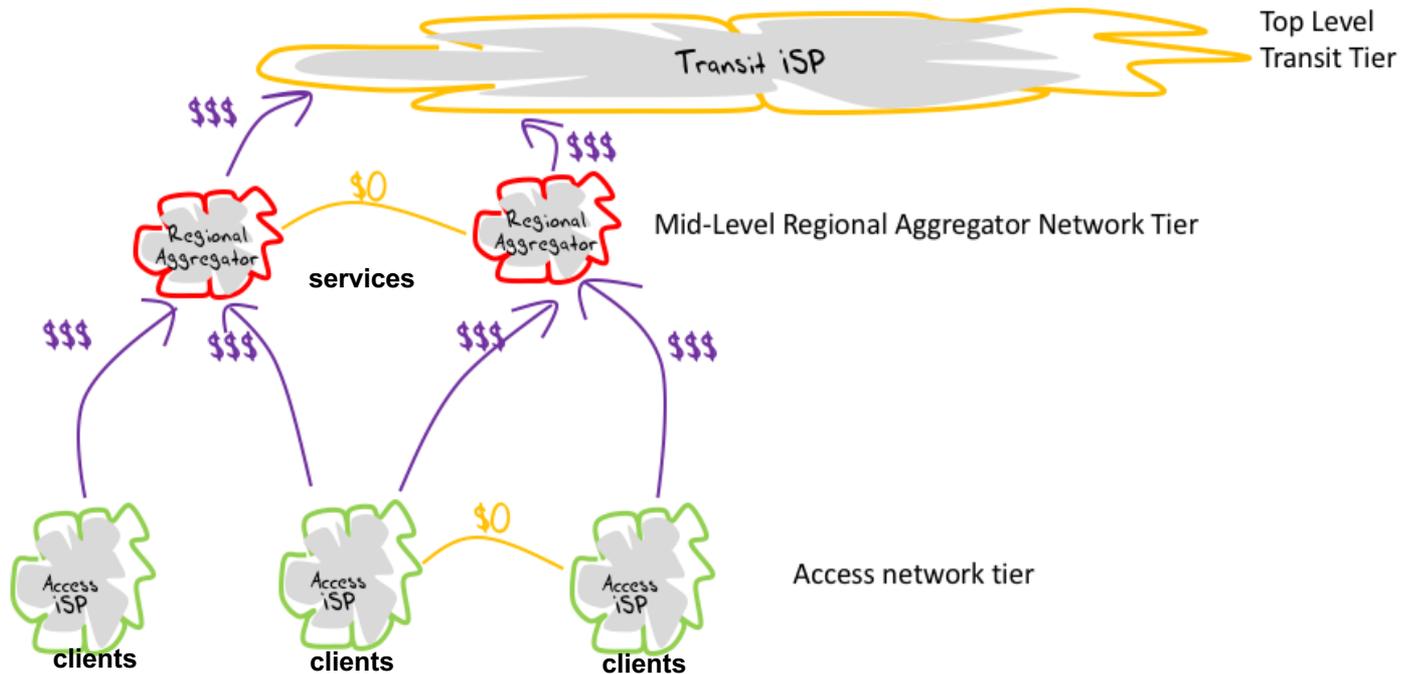
In the regulated world of national telephone operators every telephone network was “equal”

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

We also started differentiating between attached devices – some were “servers” that provided services to clients, and some (many more) were “clients” who provided no services, but access services provided by server

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

Network Role Segmentation



Enter Content

Segmenting attached computers into **clients** and **servers**

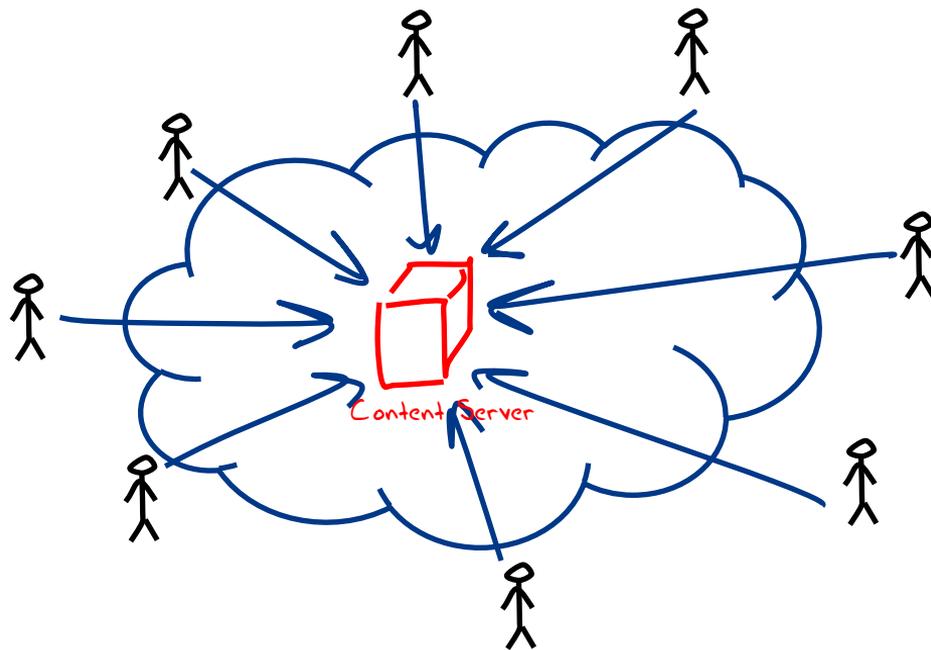
- Access networks service the needs of clients
- Clients are not directly reachable by other clients
- Clients connect to services

Content (and service) provision is the role of servers

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

- The assumption here is that there are many more clients than service points

Content Server



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 providers 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 providers should pay for carriage services, just like any other client

Content vs Carriage

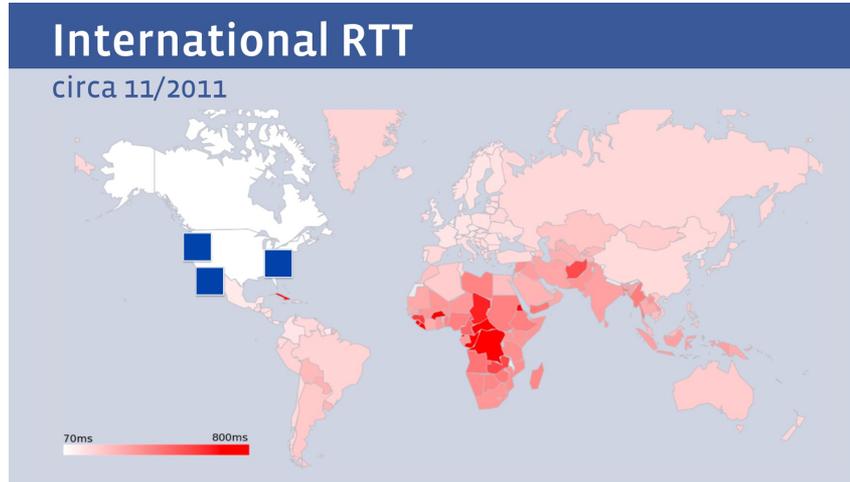
Who pays whom?

- The only reason why access networks have clients is because there are content services that clients want to use
 - Therefore, carriage providers should pay for content services
- The "peering" settlement model in the Internet – both content providers and access and network providers settle between themselves. To a carriage network, content is just another client
 - Content providers should pay for carriage services, just like any other client

The content folk resolved this fight by going "over the top" and created relationships directly with end users

The Tyranny of Distance

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

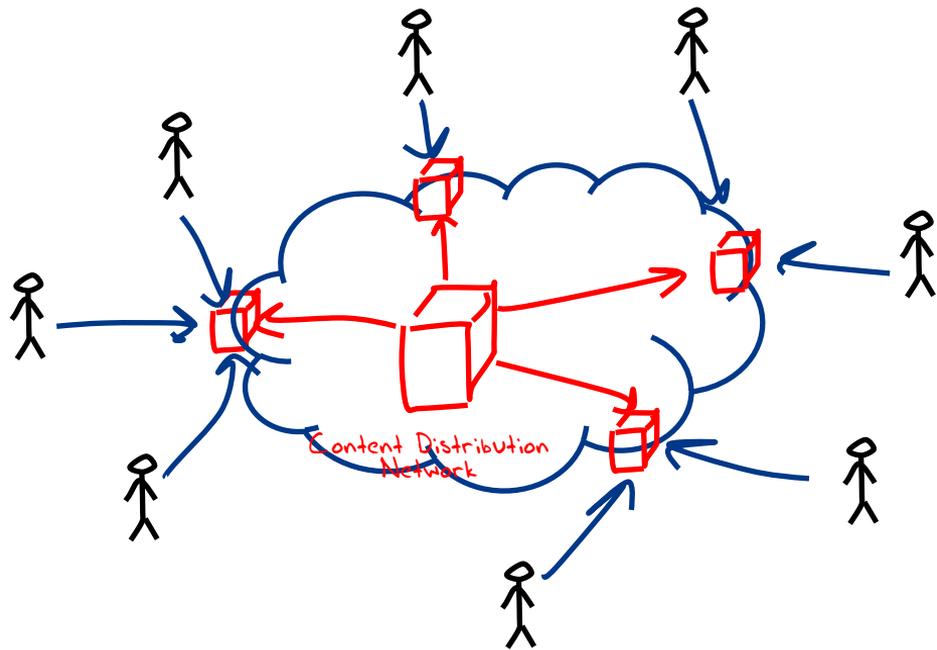


*Facebook presentation at
NANOG 68*

What's the problem?

- Having everyone pull the same content over long haul transit networks is inefficient, expensive and tediously slow
- Many carriage providers were playing with opportunistic content caching, but it relied on unencrypted content and content providers objected to the loss of “direct touch” with content consumers
- Network-level encrypted cached content can work, but it needs the participation of the content holder
- The alternative was for content caches to separately deploy at the edge of the network

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, happy customers!

Optimising Service Delivery

- Reducing the network delay between content and consumer increases the efficiency of delivery and decreases cost
- Bypassing network transit eliminates a whole bunch of odious Network Neutrality issues that have plagued many national regulatory regimes for years!
- The falling costs of storage and local delivery more than offset the costs of operating a CDN

Eliminating Distance

- The major overhead for service and content delivery is distance
 - Distance adds to the net carriage cost
 - Distance makes transport protocols less efficient
 - Distances encourages the formation of choke points
 - Distance calls in multiple providers into the e2e path, with more points of vulnerability in service delivery
- CDNs replace “on-demand” delivery with pre-provisioned local cache delivery (“just in case”)

Locality

Localized last mile delivery

- Avoids trunk carriage costs
- Removes intermediary providers
- Improves service robustness
- Allows direct "touch" between service provider and customer
- Low RTT allows efficient transport operation

Local content caches

- Exploit abundance in storage, processing and local carriage infrastructure
- Optimise cache refresh timing to scavenge otherwise uncommitted transit capacity

For the Content Provider...

CDN platforms represents a clear advantage in:

- Scalability
- Reduced service provision cost
- Reduced number of external dependencies on intermediaries
- Improved service resilience

What's not to like about
the economics of CDNs!

Part 2 - Where is this heading?

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 bunkers with global replication and DDOS hardening

Users don't reach out to content as the CDNs now bring content to users

CDNs Today

- Remember “Food Miles”?
 - The distance your food has to travel from farm gate to plate
- What about “Packet Miles”?
 - The distance your packets have to travel from platform to your device
- What is the average “Packet Mile” metric for all your content?
 - We measured this for a major ISP in Australia in the early 2000’s (pre CDNs) and the metric was ~8,000 miles for 80% of all delivered traffic
 - I suspect this has dropped to under 50 miles for the majority of that network’s customers these days because of CDNs

CDN Reach - some examples



August 2019 - <https://blog.cloudflare.com/scaling-the-cloudflare-global>

CDN Reach - some examples

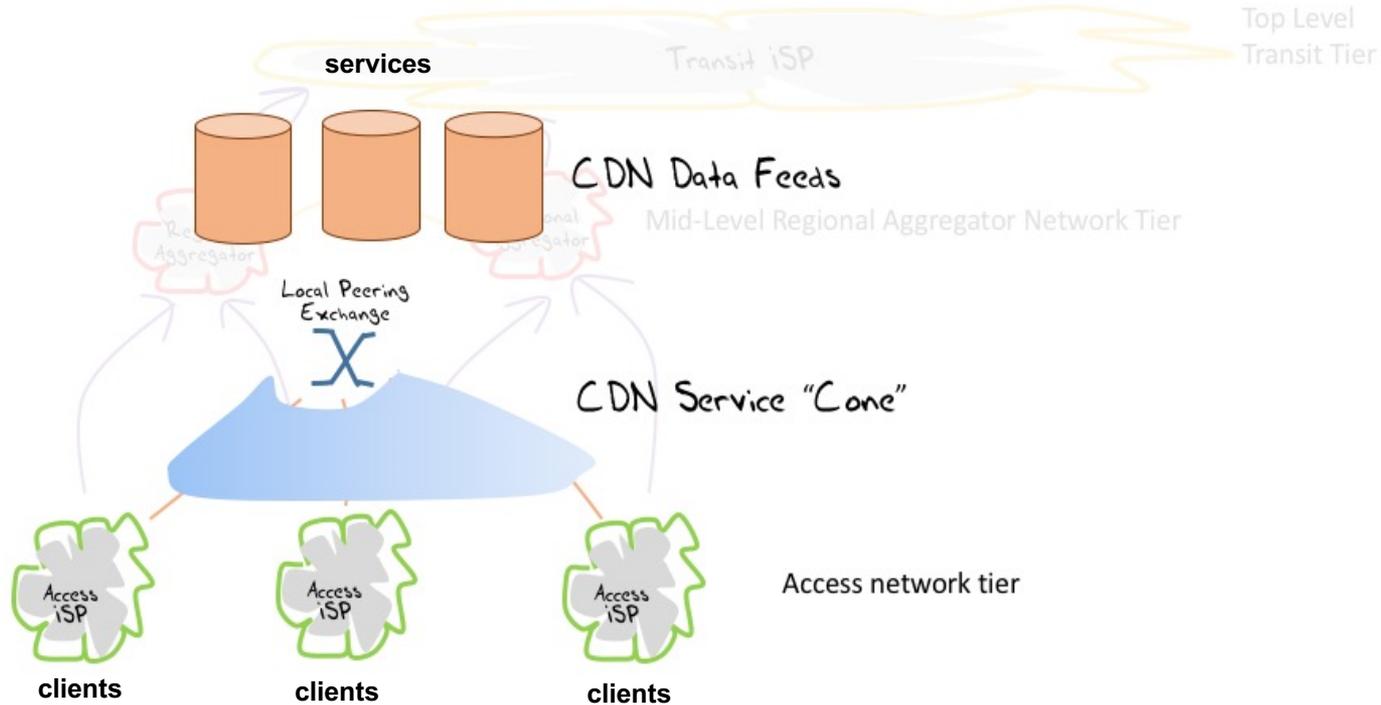


<https://www.fastly.com/network-map/>

CDN Reach - some examples



Today's Internet Architecture



Role Reversal

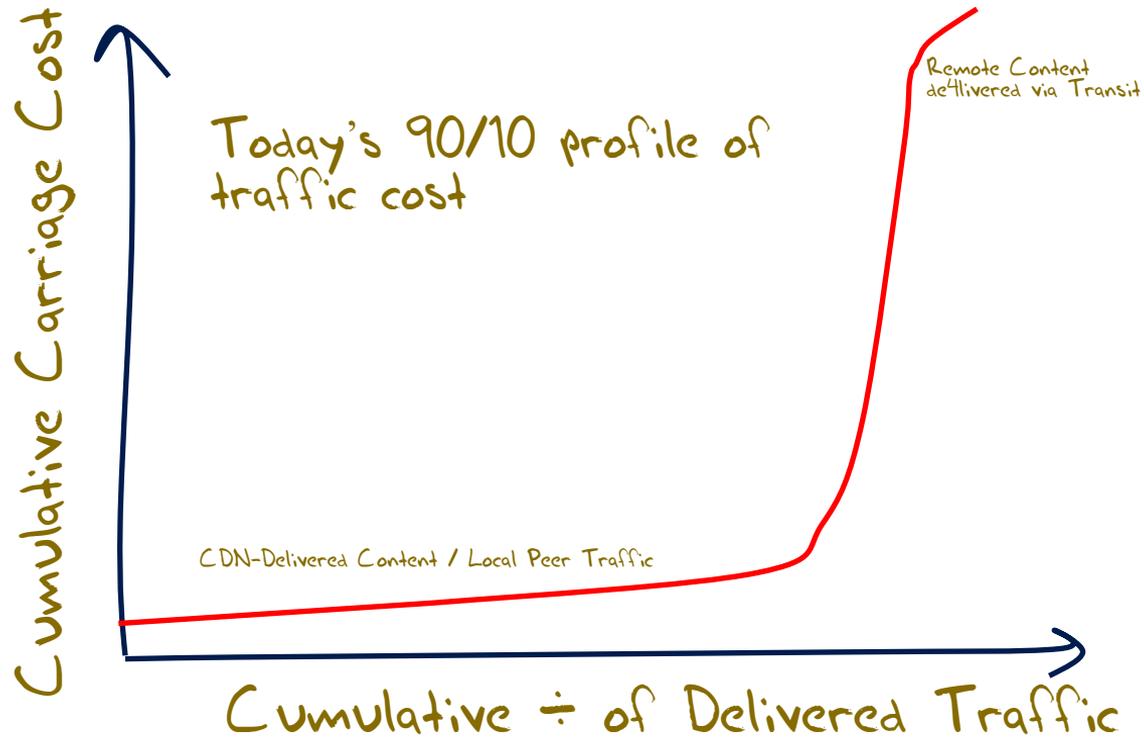
Service portals are increasingly located adjacent to users

And that means changes to the network:

- Public Networks no longer carry users' traffic to/from service portals via ISP carriage services
- Instead, Private Networks carry content to these service portals via CDN services

This shift has some profound implications for the Internet

What's important for an access Service Provider?



What's important for an access Service Provider?



When less than 10% of your traffic is 90% of your total carriage cost base then you are strongly motivated to find ways to offload that cost elsewhere!



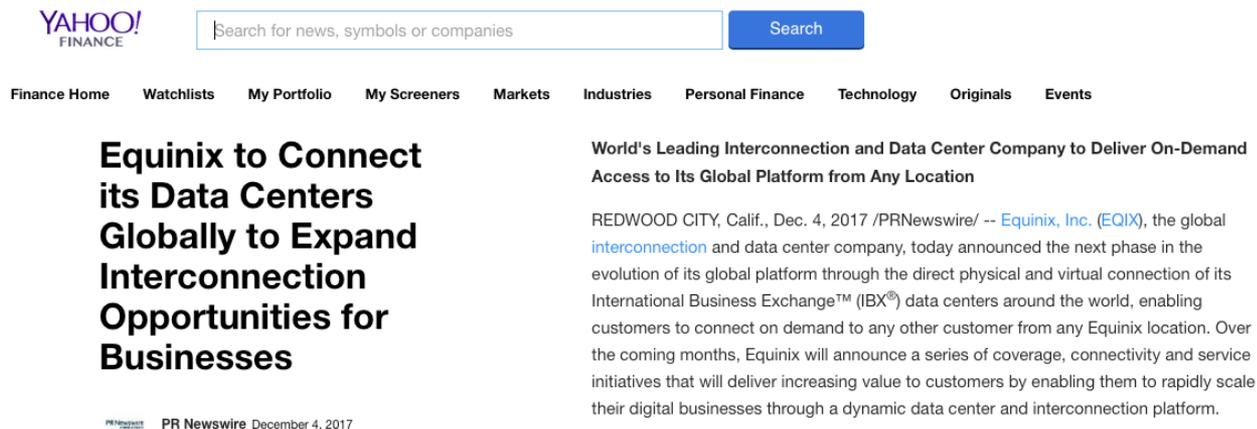
Transit?

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

Does the network fragment into a collection of localized service “cones” where all inter-zone traffic is mediated through applications and services hosted within CDNs?

Closed Transit?

We see the CDN systems reserve a carriage resource through dedicated bandwidth / wavelength / cable purchase and effectively bypass the open IP carriage infrastructure



The screenshot shows the Yahoo Finance website interface. At the top left is the 'YAHOO! FINANCE' logo. To its right is a search bar with the placeholder text 'Search for news, symbols or companies' and a blue 'Search' button. Below the search bar is a horizontal navigation menu with links: 'Finance Home', 'Watchlists', 'My Portfolio', 'My Screeners', 'Markets', 'Industries', 'Personal Finance', 'Technology', 'Originals', and 'Events'. The main content area features a large, bold headline: 'Equinix to Connect its Data Centers Globally to Expand Interconnection Opportunities for Businesses'. Below the headline is a sub-headline: 'World's Leading Interconnection and Data Center Company to Deliver On-Demand Access to Its Global Platform from Any Location'. The article text begins with 'REDWOOD CITY, Calif., Dec. 4, 2017 /PRNewswire/ -- Equinix, Inc. (EQIX), the global interconnection and data center company, today announced the next phase in the evolution of its global platform through the direct physical and virtual connection of its International Business Exchange™ (IBX®) data centers around the world, enabling customers to connect on demand to any other customer from any Equinix location. Over the coming months, Equinix will announce a series of coverage, connectivity and service initiatives that will deliver increasing value to customers by enabling them to rapidly scale their digital businesses through a dynamic data center and interconnection platform.'

Equinix to Connect its Data Centers Globally to Expand Interconnection Opportunities for Businesses

World's Leading Interconnection and Data Center Company to Deliver On-Demand Access to Its Global Platform from Any Location

REDWOOD CITY, Calif., Dec. 4, 2017 /PRNewswire/ -- [Equinix, Inc. \(EQIX\)](#), the global [interconnection](#) and data center company, today announced the next phase in the evolution of its global platform through the direct physical and virtual connection of its International Business Exchange™ (IBX®) data centers around the world, enabling customers to connect on demand to any other customer from any Equinix location. Over the coming months, Equinix will announce a series of coverage, connectivity and service initiatives that will deliver increasing value to customers by enabling them to rapidly scale their digital businesses through a dynamic data center and interconnection platform.

PRNewswire PR Newswire December 4, 2017

Open 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?

Internet Names and Addresses?

If the Internet is (or maybe soon will be) a collection of discrete CDN service ‘cones’ and a collection of applications and services, then what’s the role of IP?

Taking this one step further, why do we expect end users to pay for the maintenance of:

- A global IP address plan?
- A single global routing domain?

if all such inter-cone traffic is managed through CDNs and their service platforms?

2020's Internet Architecture

A “network” of CDNs:

- Edge devices are connected within a connected service zone defined by CDN presentation interface
- The network switching function is increasingly stateful
 - SDN and Source-Directed path specification
- Addressing is unbound from identity
- Locally-scoped routing model
- Residual transit services are provided within a service application context, managed by service applications operation within CDNs

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
- 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
- But now one-way content distribution lies at the core of today's Internet
- This content distribution role is an enterprise service framework rather than a public carriage service
- The internal parts of the carriage network are now being privatized and removed from public regulatory scrutiny (assuming that the Internet every had any such scrutiny in the first place!)

Policy?

If CDN feeder 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?

The Large and the Largest

Company	Market Cap (\$B)
Apple	2,994
Microsoft	2,795
Alphabet	1,764
Amazon	1,570
Nvidia	1,223
Meta	909
Tesla	789
Berkshire Hathaway	783
Eli Lilly	553
TSMC	539

The world's 10 largest publicly traded companies, as ranked by their market capitalization, Q3, 2024

Content Really is King

- None of these seven technology companies are a telephone company, or even a transit ISP, or even an ISP at all!
- Five 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**

Content Consolidation

- There are not thousands of content service platforms
 - There are just a few left
- And the space is dominated by a small number of dominant actors who set the rules of engagement 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.

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, with users, and even with the public sector?

Competition or Cartel?

With a small number of truly massive enterprises at the heart of the area of digital content and services, the market is one that is shaped by competition.

Or do they get to set their own terms of engagement with each other, with users, and even with the public sector?

As concerning as this might sound, it's not a novel situation!

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.

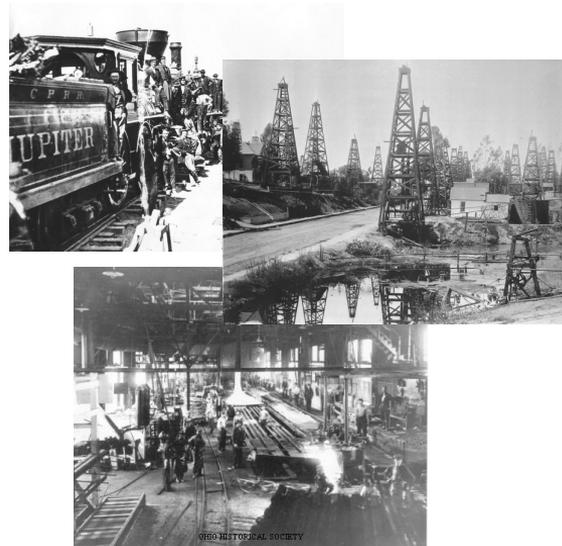
This gallery examines the leading cultural phenomenon of the 1870s and 1880s, the American Aesthetic movement, through a range of objects produced for affluent consumers. Aestheticism, rooted in the English philosophies of John Ruskin and William Morris, advanced the notion that a beautiful environment 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

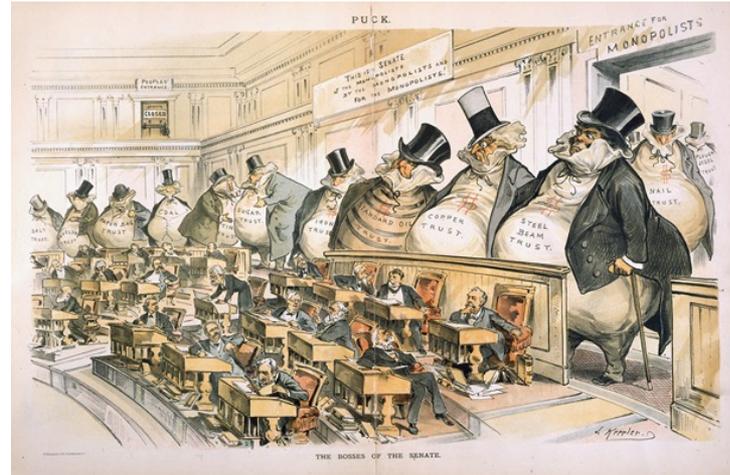
Andrew Carnegie - US Steel
John Rockefeller - Standard Oil
Theodore Vail - AT&T
George Westinghouse – Rail Brakes
Thomas Edison – General Electric
J P Morgan - Banking



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 remain dominant in their field

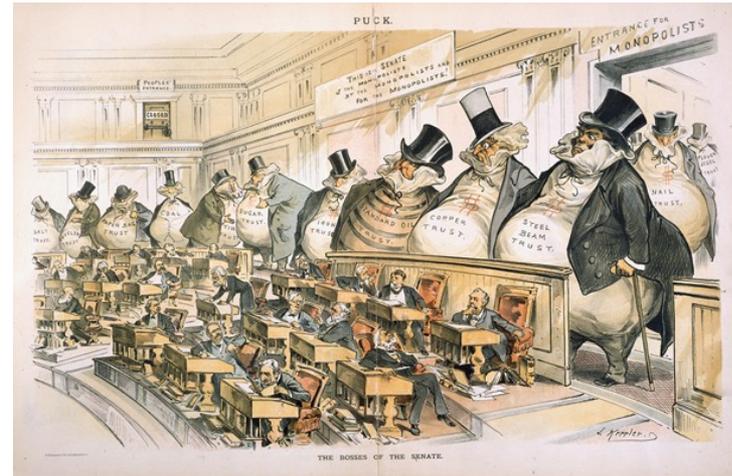


Who's Gilding Today?

Company	Market Cap (\$B)
Apple	2,994
Microsoft	2,795
Alphabet	1,764
Amazon	1,570
Nvidia	1,223
Meta	909
Tesla	789
Berkshire Hathaway	783
Eli Lilly	553
TSMC	539

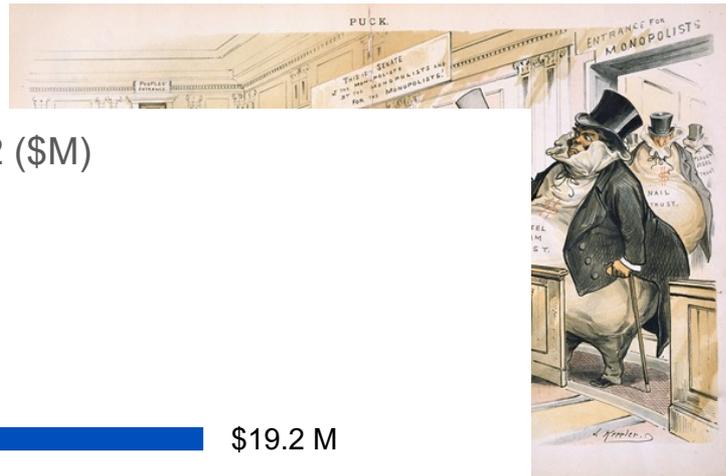
The Internet's Gilded Age

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.



The Internet's Gilded Age

At some point in the past decade or so the dominant position across the entire



US Federal Lobbying 2022 (\$M)



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 parties
-

is this the internet we were dreaming of?

By taking a dominant 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 incremental changes in carriage and communications within the Internet.

For me, the essential topic of this conversation is how we can strike a sustainable balance between an energetic private sector that has rapidly 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!