

The Death of Transit and Beyond

Geoff Huston

Chief Scientist, APNIC



This presentation is not about details

- Or specific plans
- Or particular services
- Or any particular technology
- Or anything like that

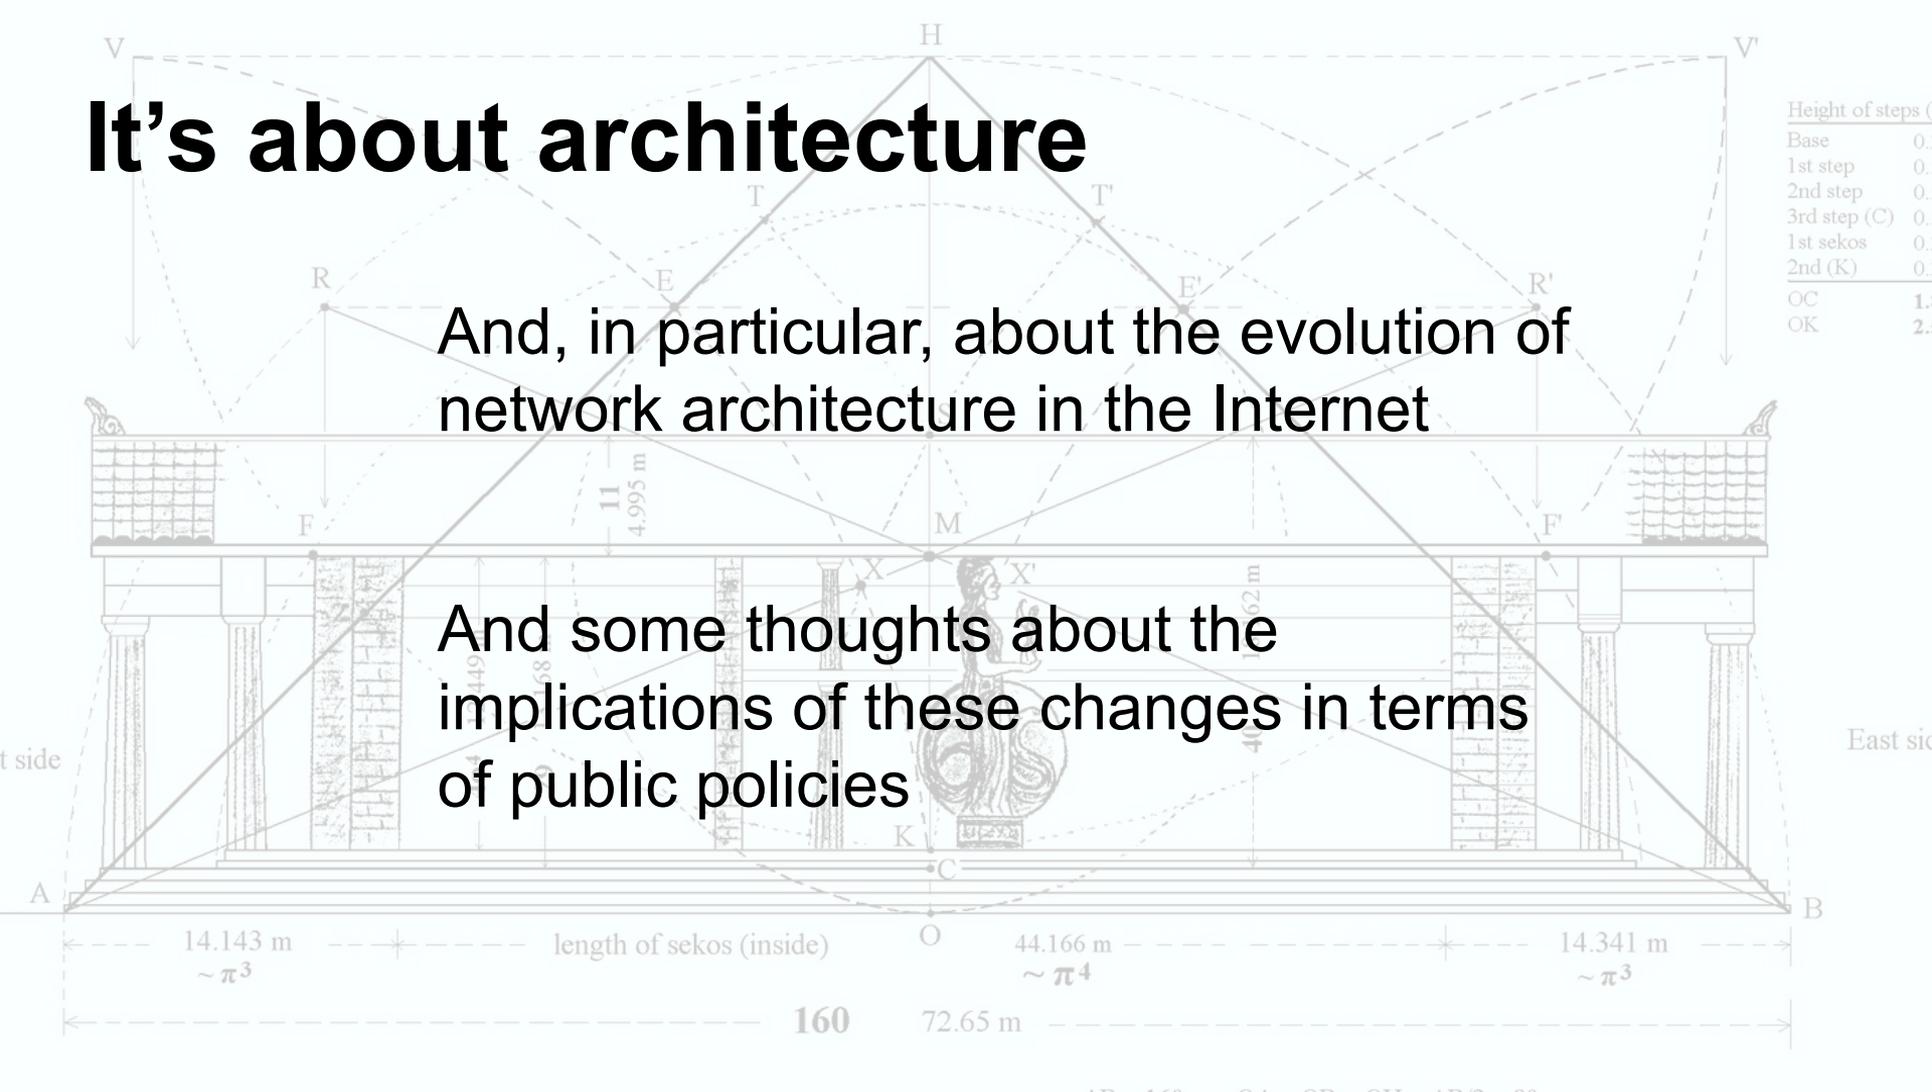


It's about architecture

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

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

Height of steps (m)	
Base	0.0
1st step	0.1
2nd step	0.2
3rd step (C)	0.3
1st sekos	0.4
2nd (K)	0.5
OC	1.0
OK	2.0



14.143 m
 $\sim \pi^3$

length of sekos (inside)

44.166 m
 $\sim \pi^4$

14.341 m
 $\sim \pi^3$

160

72.65 m

West side

East side

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

The rise of digital networks for computers

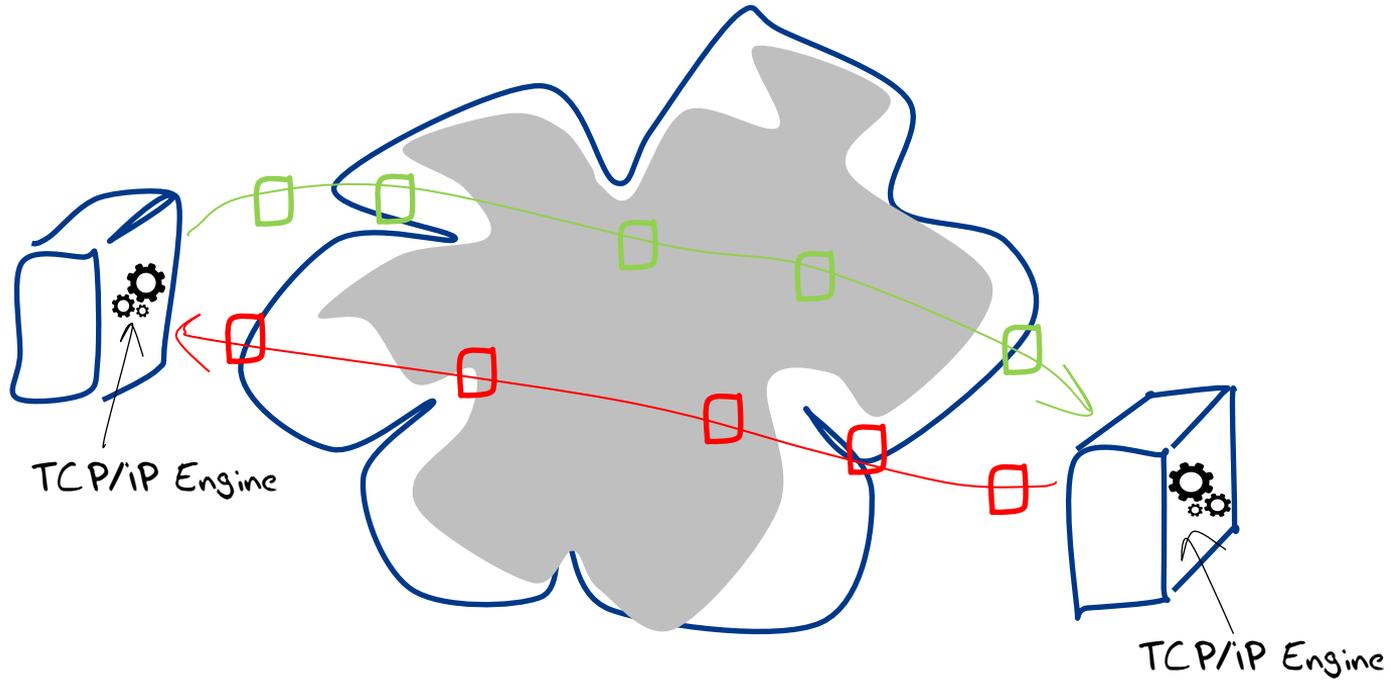
- No requirement for real time synchronous transparency
- Provided the opportunity to drop synchronicity and virtual circuits and use packetised comms
 - Packet Data Networking allowed us to realise opportunities for higher efficiencies and lower costs
- Connected computer to computer
- The network was intentionally transparent
- Edge computers were variously both clients and servers

The Internet Architecture (c1980'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





The Result was Revolutionary!

- Very Simple
- Extraordinarily Cheap
- Unbelievably Efficient

By stripping out network-centric virtual circuit states and removing time synchronicity the resultant carriage network was minimal, while more complex functions such as flow control and reliability were pushed out to the computers on the edge



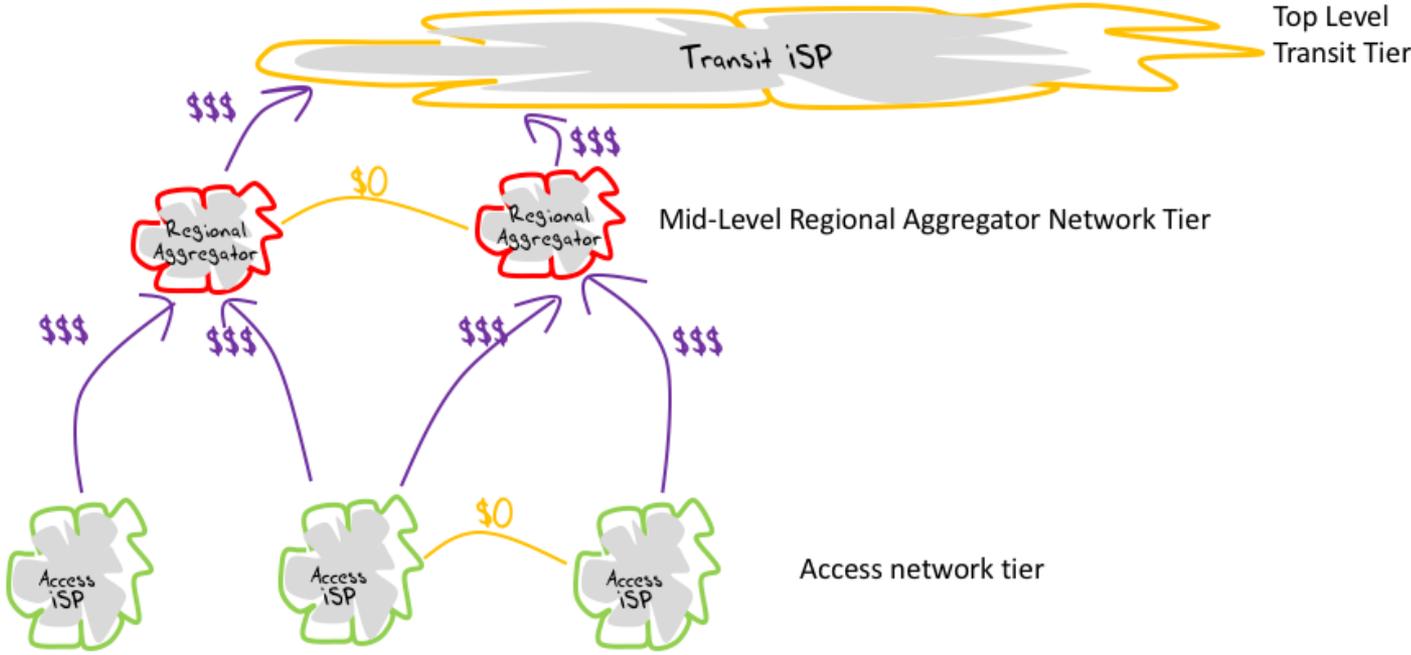
Internet Evolution

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

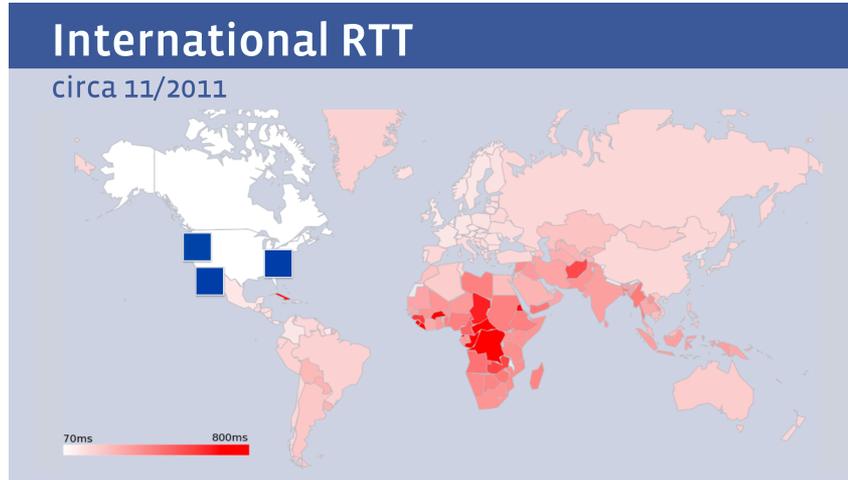
Who pays whom?

- The only reason why access networks have a “top” relationship with content providers is because there are content services that clients pay for.
 - Therefore, content providers should pay for carriage.
- *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



The Tyranny of Distance

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



*Facebook presentation at
NANOG 68*



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 small 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

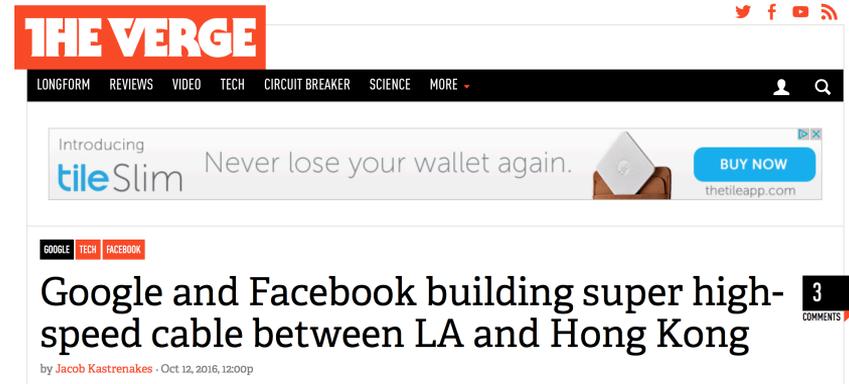


Who's building now?

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



The screenshot shows the top portion of a web browser displaying a 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 black bar. A social media bar with icons for Twitter, Facebook, YouTube, and RSS is in the top right. Below the navigation is a white banner for 'tileSlim' with the text 'Introducing tileSlim Never lose your wallet again.' and a 'BUY NOW' button. The main article title is 'Google and Facebook building super high-speed cable between LA and Hong Kong' with a '3 COMMENTS' badge. The byline reads 'by Jacob Kastrenakes · Oct 12, 2016, 12:00p'.



Today's Internet Architecture

We've split the network into clients and servers

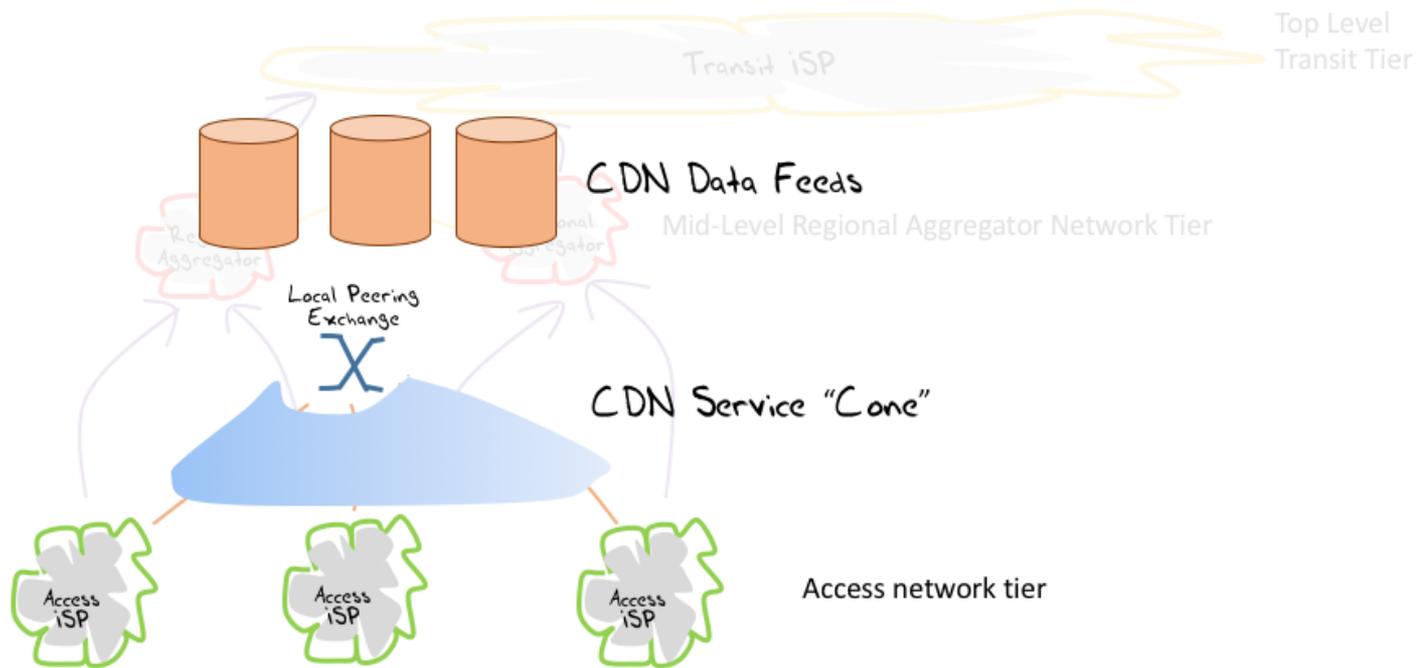
- Web servers
- Streaming servers
- Mail servers
- DNS servers

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!
- If the Internet is (or maybe soon will be) a collection of discrete private CDN service ‘cones’ then why do we still need :
 - A global address plan?
 - 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



Where are we?

- We started this journey building a telephone network for computers to communicate between each other
- But today 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?



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.

** look it up on Google!*



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 party suppliers
- Regulators

Was this the internet we were dreaming of?

In 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



Thanks!

