Who Owns the Internet?

Jason Prado

Who Owns the Internet?

by Jason Prado. I don't own the Internet.

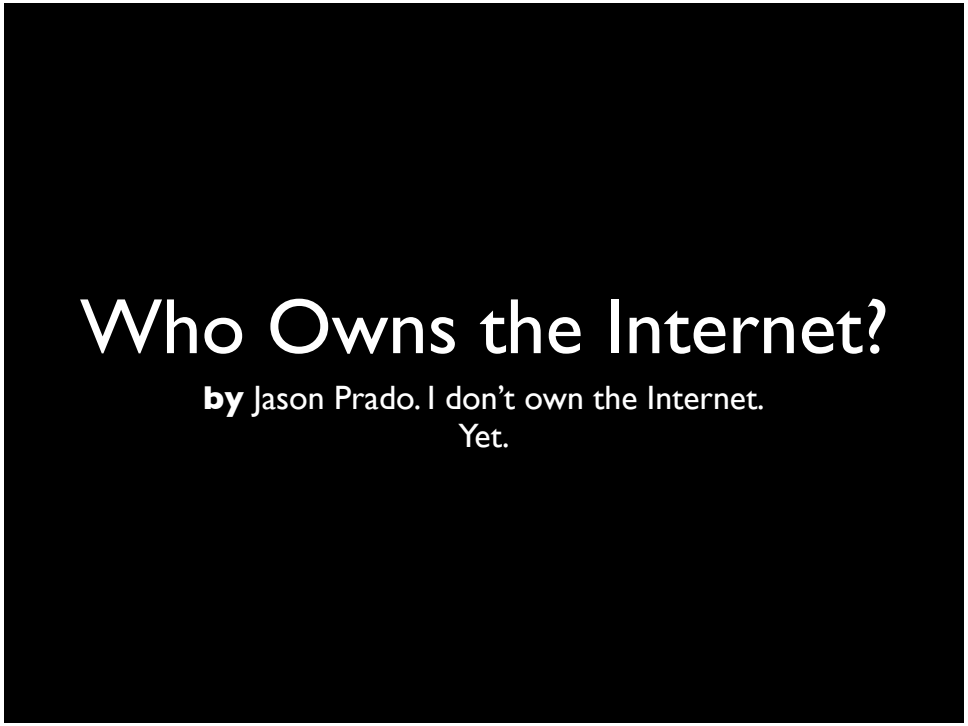

speaking of jason prado, who am i and why do i think i know a lot about the internet?

in increasing order of importance,

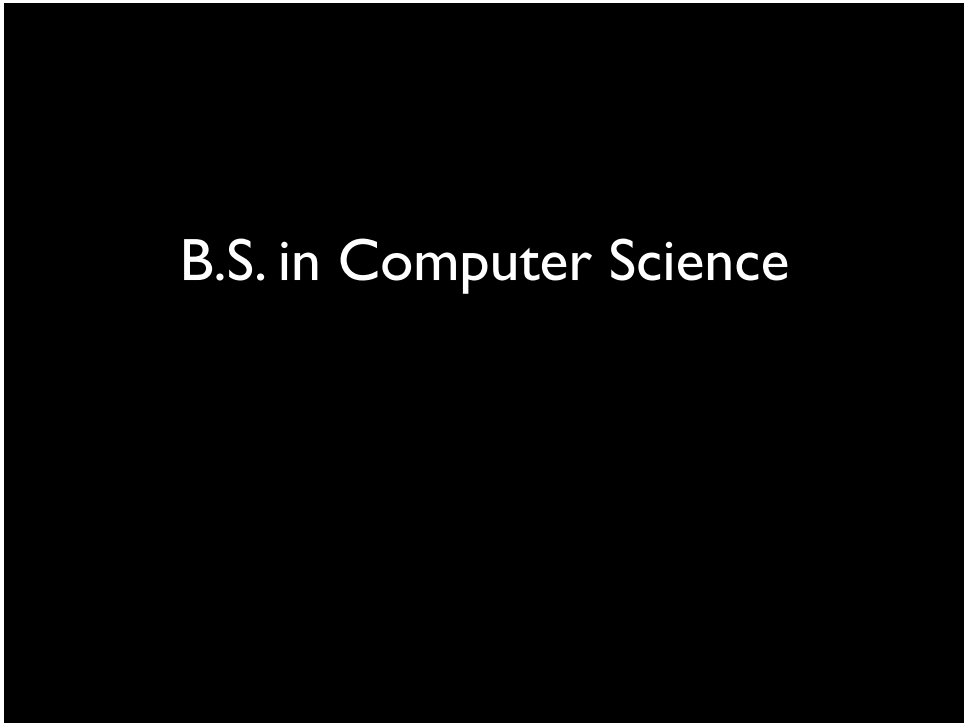

from Stanford University

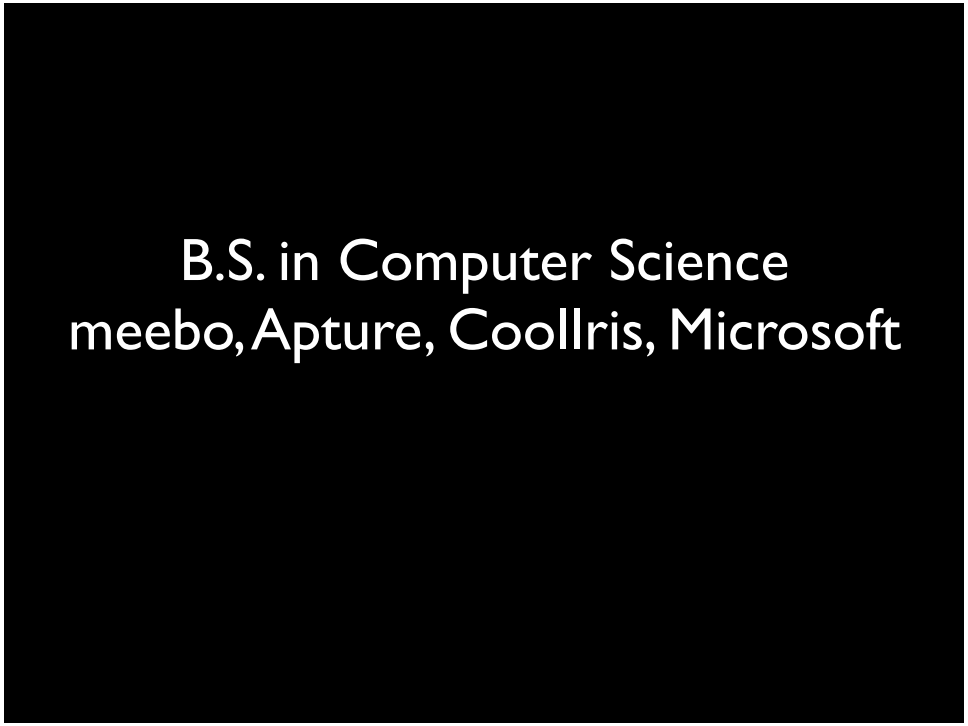

first hire to 490,000th hire

B.S. in Computer Science
meebo, Apture, CoolIris, Microsoft
everythingIsTheBest, LLC

B.S. in Computer Science
meebo, Apture, CoolIris, Microsoft
everythingIsTheBest, LLC
Serious Internet Enthusiast

So, what happens in your browser? A lot.
A lot of electrons moved over a lot of wires.
A lot of people own these wires.
The electrons form a signal.
A lot of people agreed on how those signals are organized into data that can be read by a computer.
Then that data is read and understood by the operating system and web browser.
A lot of people agreed on the protocols that are used to read that data.
There are layers upon layers of meaning in the electrical signal across that wire.
And everyone has to agree on how things work on every level, or else traffic just doesn't flow.

Technical problems: how do you move data from one side of the world to the other?

Also not so technical problems.

Authority and Ownership are not necessarily technical problems, but we'll look at some technical solutions to them.

We'll see these issues raised in two aspects of the internet.

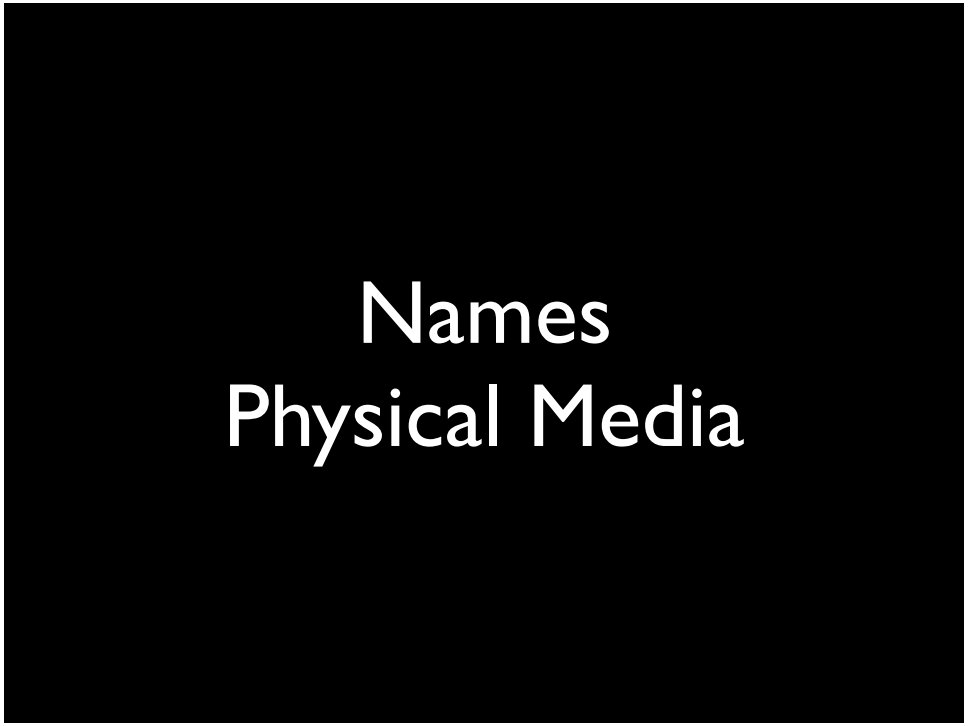

Names are important because every machine on the internet must be uniquely addressable.

Physical media is interesting because wires are tangible goods that are owned by companies, but all their value comes from the data moved over them, which is not material.

Basics

IP Addresses

We're going to assume that every machine on the internet has its own unique IP address.

All data transmitted on the internet is sent to an IP address, that's how it knows where to go.

IP Addresses

Me: 98.237.252.207

We'll talk more about IP addresses later

If I need a message to get from me to google, I send it from my IP to google's IP.

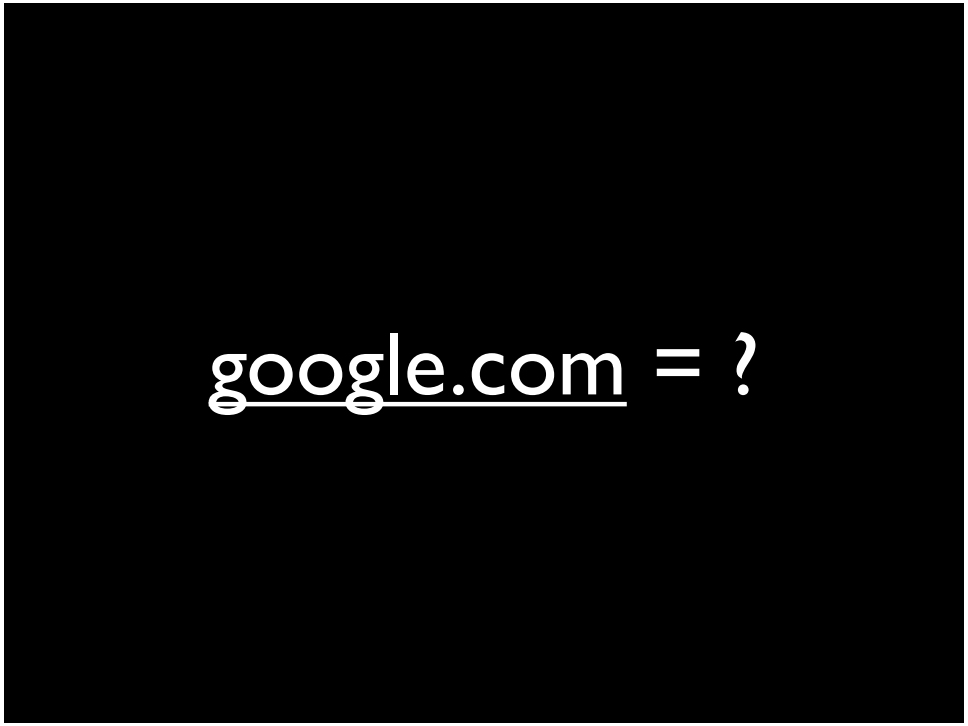

Back to the browser. I typed in google.com. I don't really know how to send a message to google.com.

I'm going to need to resolve google.com to an IP address.

We will do this via...

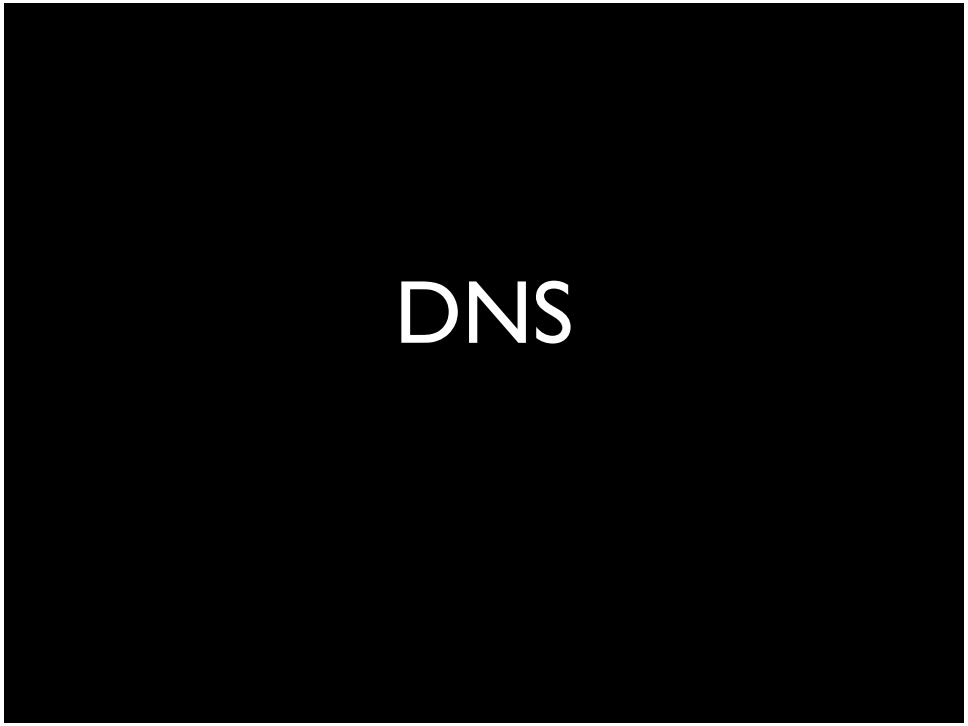
DNS

DNS
Domain Name System

I am told that about a hundred years ago, when you needed someone's phone number you consulted a large book that mapped names to phone numbers.

This would be nice to have on the internet, but it's a little more complicated.

DNS is considered critical to the functioning of the internet. It is part of the Internet Suite, a collection of such critical technologies. To give some perspective, the world wide web and email are not considered important enough to be in the IP suite.

DNS is a big deal.

Your computer needs to ask someone for google.com's IP, but who?

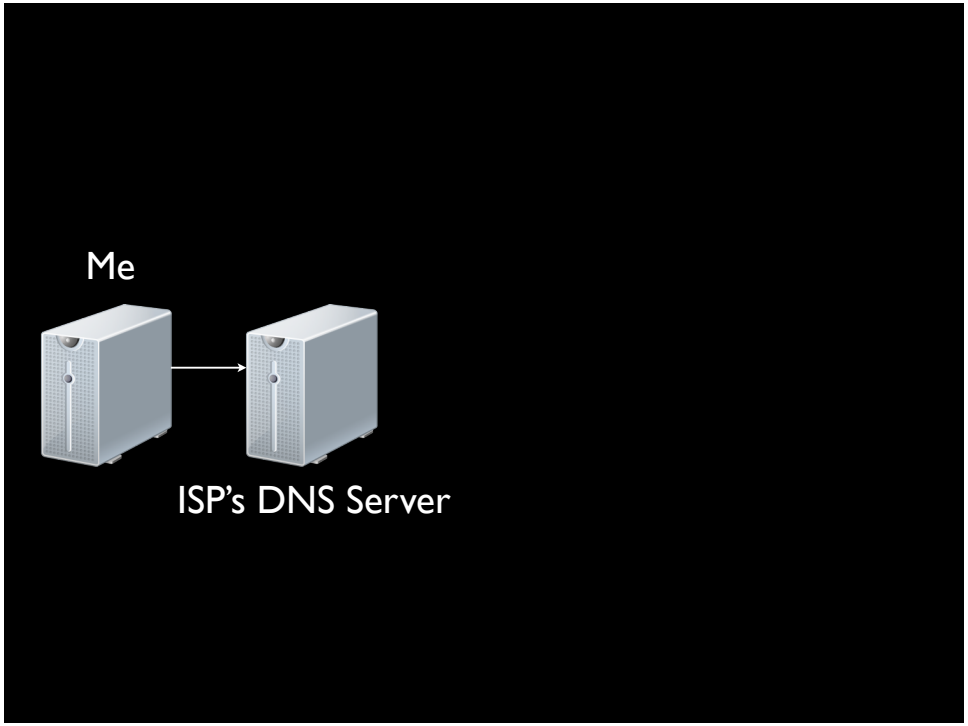

Your computer knows at least one valid IP address out there in the world. It is probably set up automatically by your ISP.

It is a DNS server. A machine on the internet that will answer your question. What is the IP for google.com?

But it doesn't know.

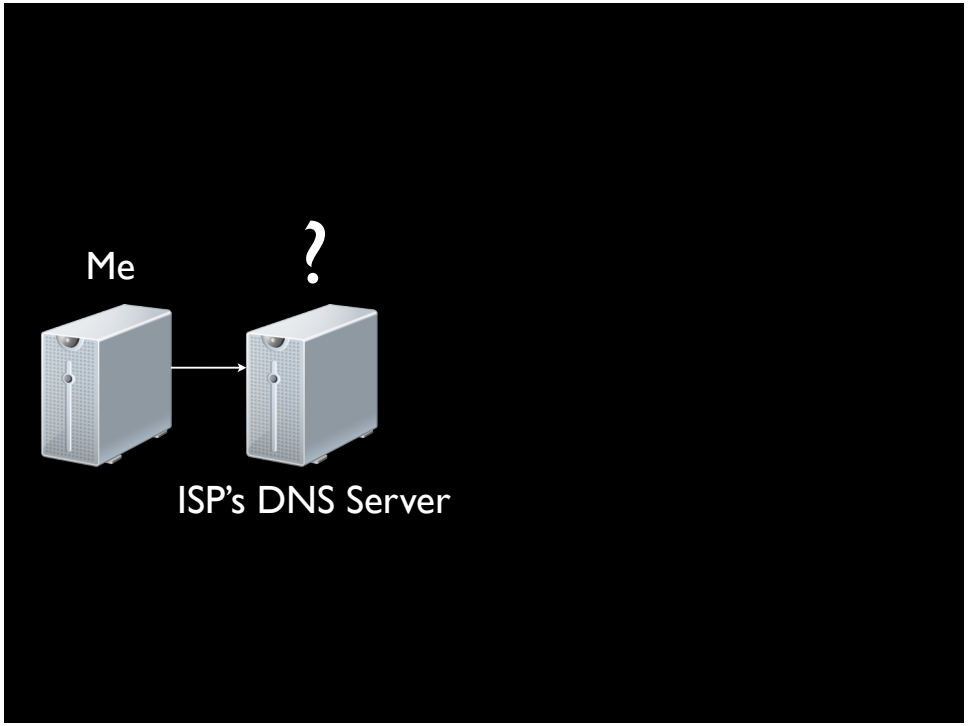

So it has to ask someone else. Who will it ask?

It parses google.com. google and com. The last part of it is .com. The DNS server knows that someone out there knows more about .coms than it does, and they probably know who to ask about google.com.

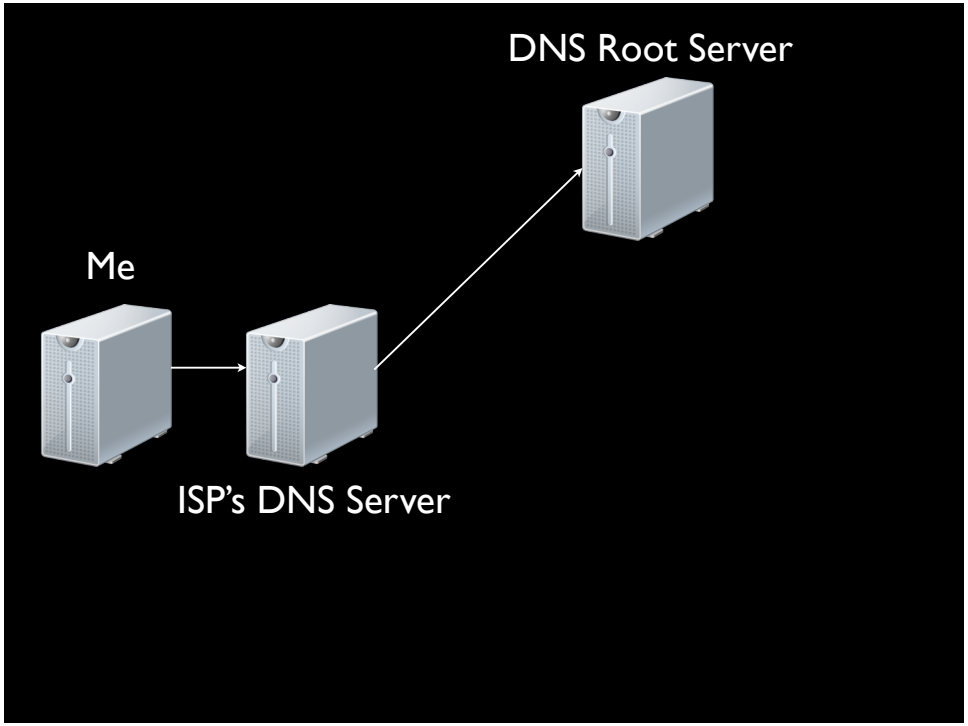

That someone is a root server.

A root server will tell another DNS server a little bit more information. In this case, it will refer it to another server to ask about .com's. It would refer it to a different server for .uk or .cn.

Root servers are very interesting. There are 13 of them in the world. They are the ultimate authority on name resolution on the internet.

So here we have our first questions about authority. Who owns the root servers? What are their motivations?

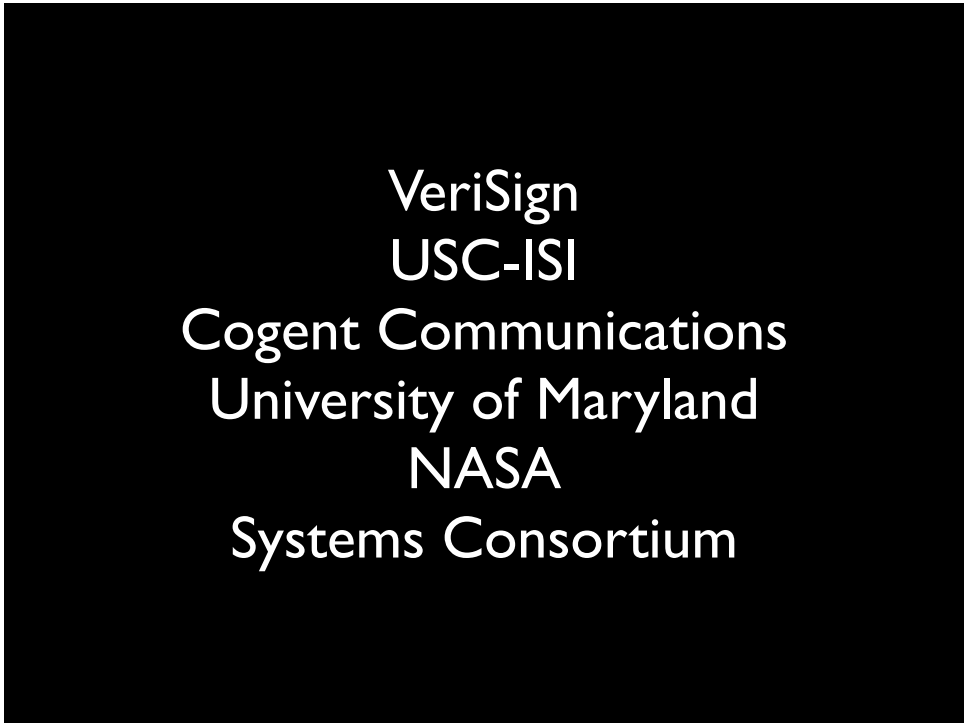

These are the maintainers of the root servers.

Some are corporations.

Some are public institutions.

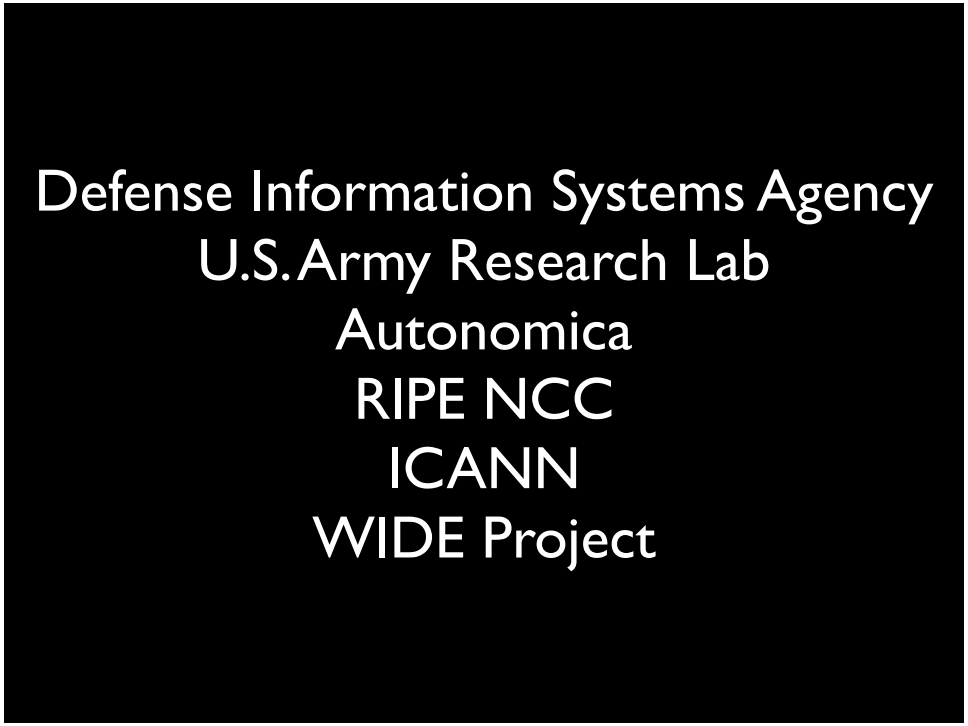

Some are for-profit.

Some are not-for-profit.

Some are military

To look at a few, one organization is RIPE NCC.

To look at a few, one organization is RIPE NCC.

RIPE NCC is a not-for-profit whose mission is to support the infrastructure of the internet.

They are funded by their members, for whom they provide various services.

They approach the task with a spirit of stewardship over the internet.

VeriSign is a for-profit corporation. They make a lot of money off the internet and find it in their interest to give back.

Their brand is boosted by the prestige of having a hand in the root infrastructure of the internet.

They get a seat at meetings also.

The University of Maryland hosts a root server as a public service. Like some other hosts, they receive grants to keep their root server running.

Government agencies and defense organizations host root servers as a public service and because the infrastructure of the internet is now critical to military and government operation.

As an aside, what happens if all 13 root servers were attacked at once, physically or via a cyberattack?

Well, the 13 servers are actually 200 servers masquerading as 13. And they're all around the world. And they're backed by fat pipes. So that's unlikely.

In 2002, allegedly half of all the servers were taken down for several hours. No users noticed.

If, somehow, they were all taken offline for 48 hours, then with no interference the internet would grind to a halt. In reality, even if that happened, diligent engineers would be able to work around it in time.

So, don't worry.

The servers are coordinated by ICANN.

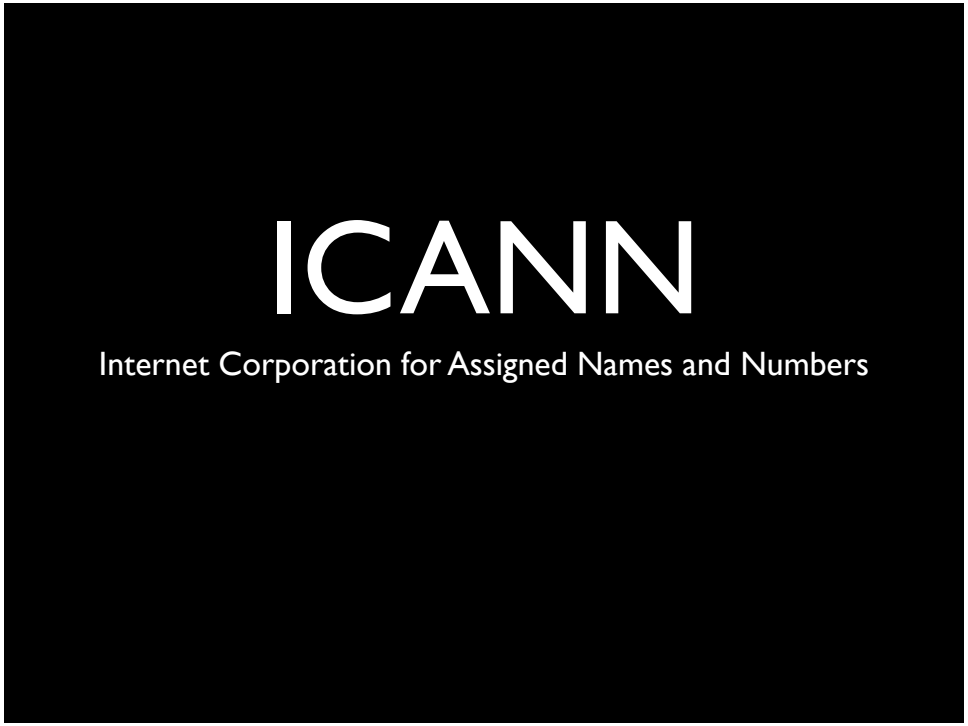

ICANN is a not-for-profit corporation.

They have control over all Top-Level Domains. .com, .uk, .cn, all of them.

Vint Cerf, the “creator of the internet”, was the Chairman of the ICANN board until 2007.

Members of the board are put through a nomination process, and the board takes input from advisors and the public.

You might remember hearing about their public debate process a few years ago when there were efforts to approve a .xxx TLD.

Here you have it, an American corporation that has an incredible amount of power over the internet. They seem to have our best interests in mind, though.

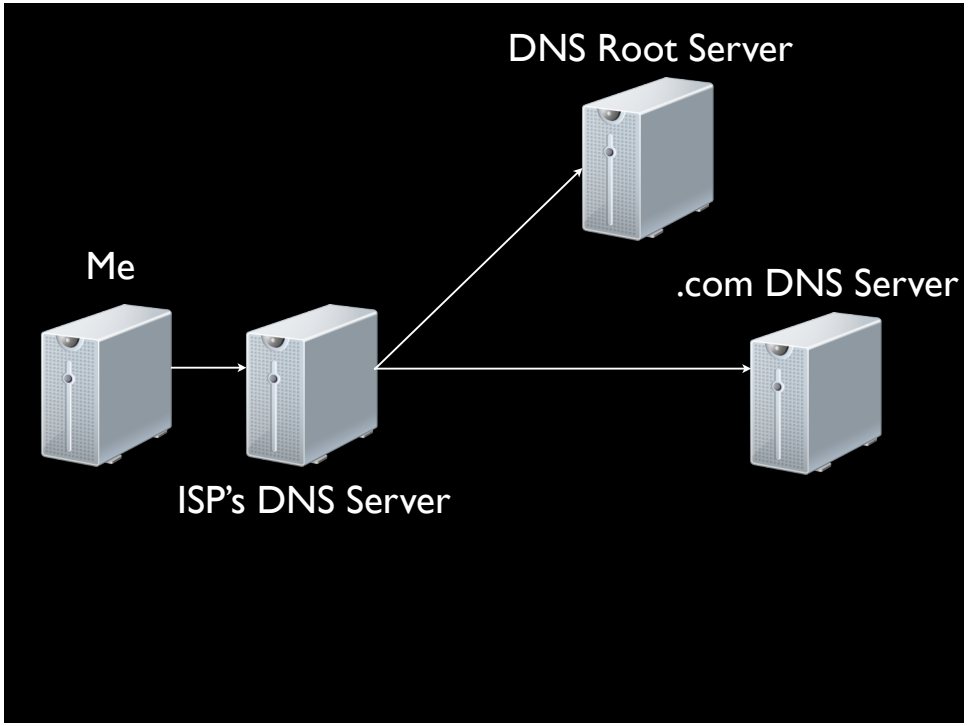

The root server has told us what DNS server is authoritative for all .coms. Guess who that is?

VeriSign has the power to tell you where to go to look up google.com They control .coms. They are very powerful.

They are a for-profit corporation. So how do you think they use this power?

They make a lot of money.

To be precise, the US Department of Commerce is granted power over .com by ICANN. Then the DoC contracts out administration of the TLD to VeriSign.

One source I found put the price of a domain at $50, with $35 going to VeriSign and $15 going to a government fund.

When you buy a domain you are largely buying an entry in VeriSign's DNS records.

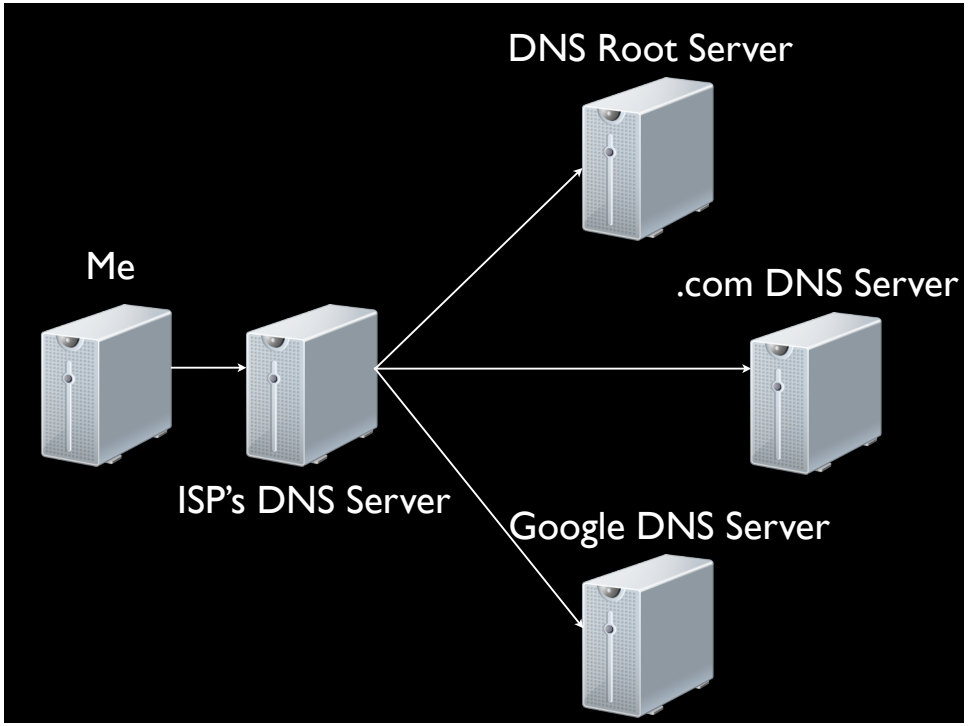

The root server pointed you to VeriSign's server.

VeriSign's server points to Google's DNS server. This is the DNS server Google specified when they purchased google.com from VeriSign.

Your ISP's DNS server asks Google's DNS server about google.com and finally gets an IP back, which it gives to you.

It will then hold on to that information in case you ask again right away.

Why is the DNS system so complicated?

Why a distributed database?

Why so many organizations involved?

Why talk to 4 servers to get a DNS lookup?

Centralization is Bad

The alternative is centralization.

We could just politely ask Google to run all of DNS. And they would say yes. Just ask Google for a lookup.

But do we trust Google? Do we trust Microsoft?

And doing so represents a single point of failure. And the internet abhors a single point of failure.

There's a lot of requests and they have to be fulfilled. Reliability it seriously important.

No one entity can be trusted to run the DNS system. Rather, this balance of entities all hold a stake in the DNS system's continued function.

DNS follows this distributed principle. Email, built atop DNS, does as well. If we agree on DNS, we can do a lot.

Twitter does not care about being distributed. Facebook.com is the only way to message people on Facebook. If they disappear tomorrow, so do all the data you've put into it. This should bother you.

IP Addresses

Now that we know who is authoritative with respect to DNS, you might be wondering who controls the IP addresses themselves.

This question is somewhat simpler. Again, ultimately ICANN and its subsidiaries does.

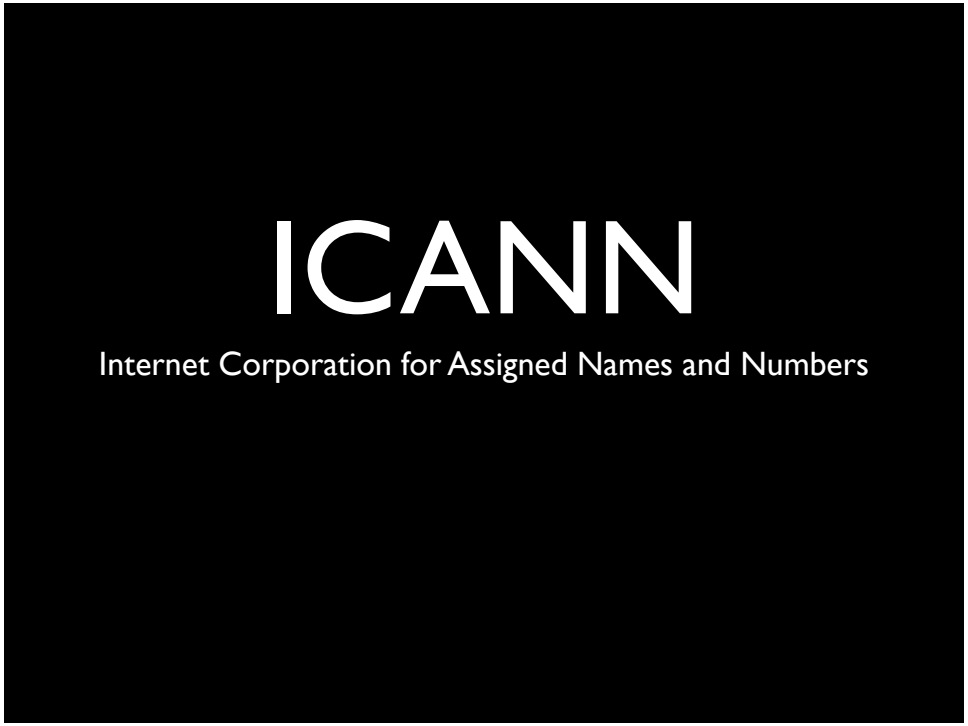

Particularly the IANA subsidiary of ICANN.

They delegate control over IP addresses to five regional not-for-profits.

In the US, we get IP addresses from ARIN.

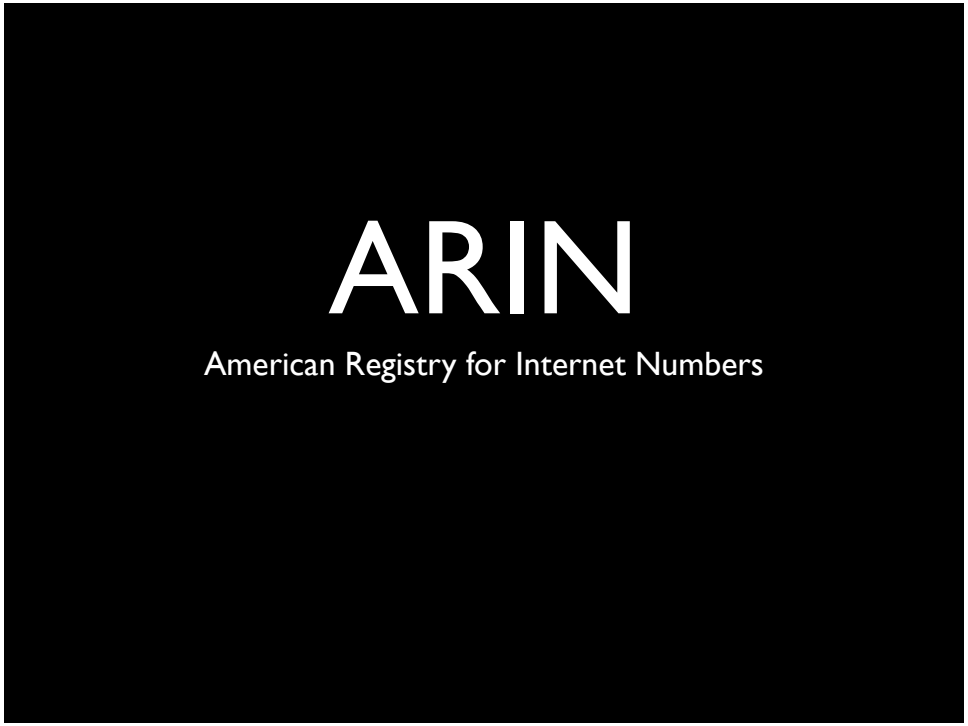

An ISP buys a block of IP addresses from ARIN at about $1000/yr for a few thousand IPs.

Then you pay your ISP your monthly fee and they lease an IP to you.

Names

That's how naming works on the internet.

Names are an intangible aspect of authority and ownership on the internet

The last thing I'd like to cover is the physical media of the internet.

Physical Media

The cables in the ground.

The first important thing to consider here is that no one entity owns enough cables to get data from your house to google.com

Routing

Your ISP, say Comcast, probably only owns or leases cables only within Seattle. To leave Seattle, they have to put data on someone else's cables

Data gets across the country via routing.

A single packet of data probably takes a dozen hops to get from here to California. Each hop is a machine, a router, in a place where cables join.

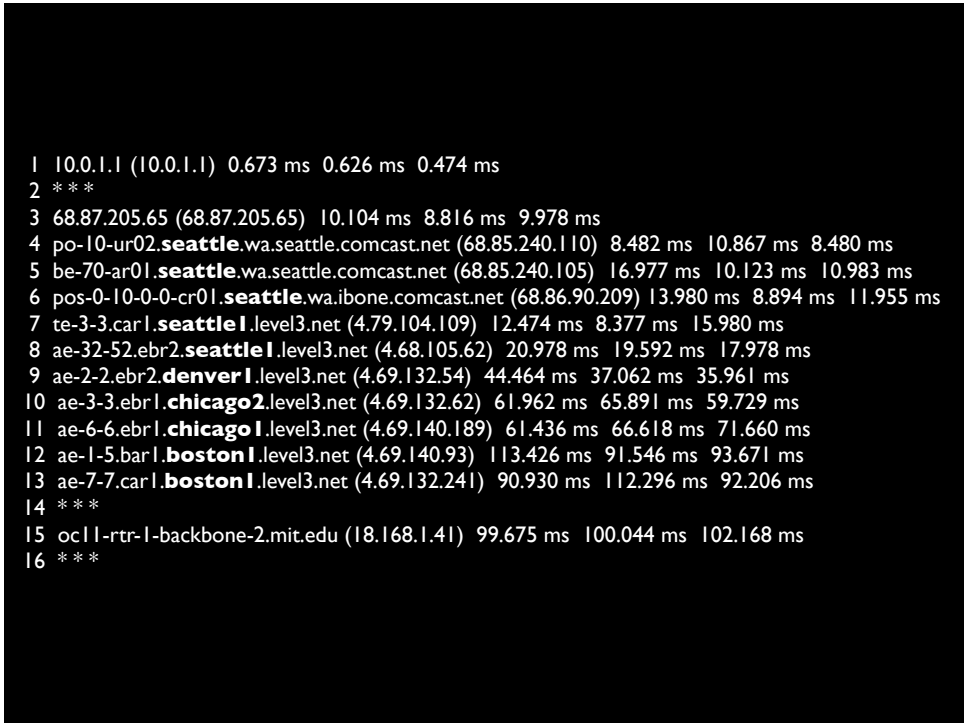

Here's a sample traceroute of a packet moving from my computer in Seattle to a computer at MIT.

See the physical locations of the data? It bounces around Seattle for a while, then hops across the country and lands in Boston.

A local ISP owns or leases the cable within a city. So who owns these hops in between?

Level 3
Verizon
AT&T
Sprint

These are some companies that have long runs of fiber in the ground.

What are their motivations? They're corporations, so they are getting paid somewhere.

Let's walk through a transaction to see who pays who.

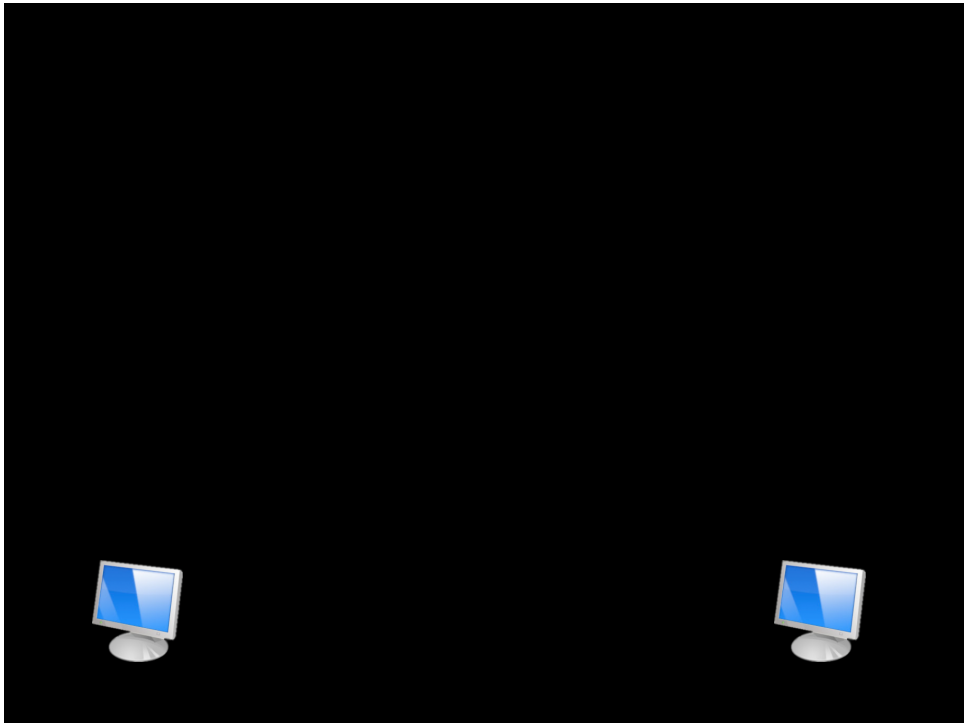

I want to download an MP3 via Napster from someone across the country.

I have an ISP and they have an ISP. Suppose they're fairly small ISPs.

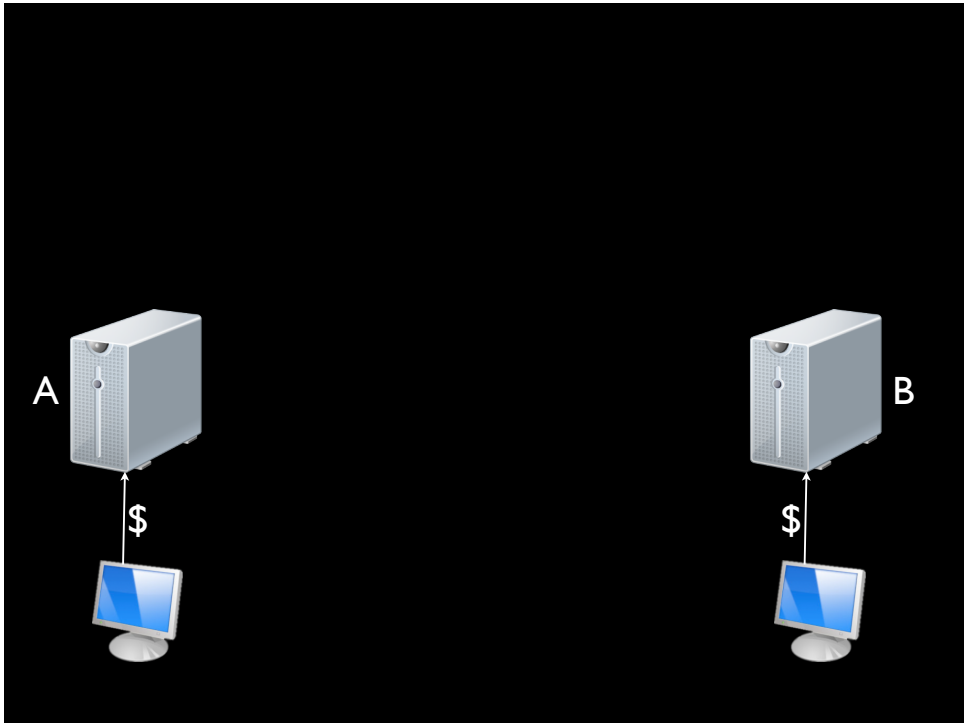

I have ISP A and they have an ISP B

We each pay them a monthly fee

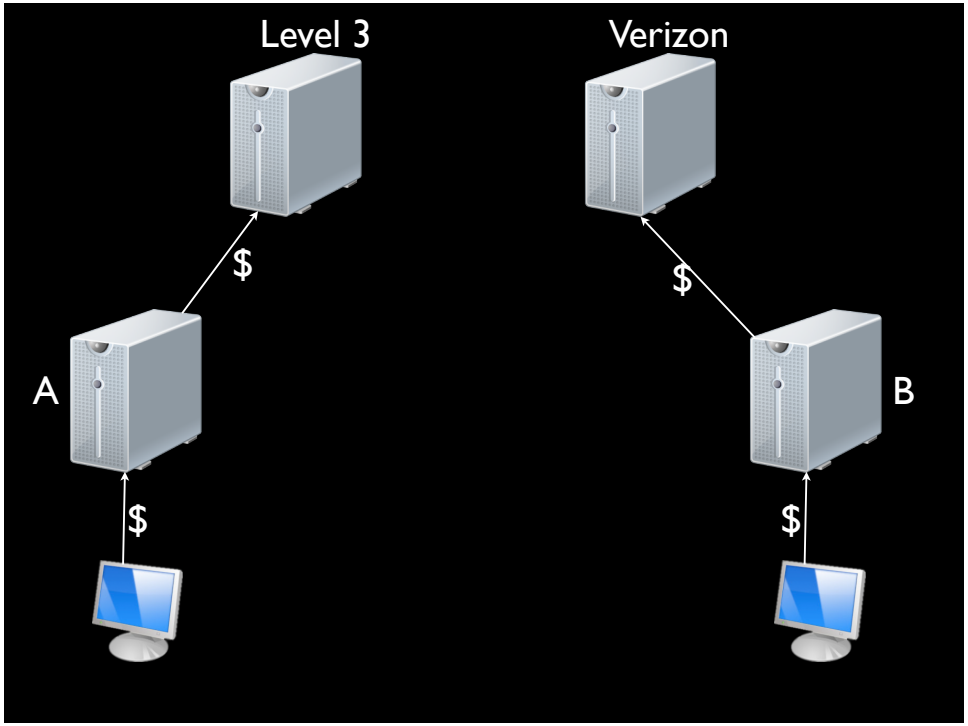

They only have cables in their respective cities. A and B have no idea how to get data to each other.

But they peer with backbone providers. Level 3 and Verizon.

If ISPs A and B are sufficiently small enough, they pay these bigger providers for access.

So what are they paying for? They pay for the right to plug their cables in at the same building as these bigger providers. Then their respective routers can talk to each other and move data.

This is called paying transit. It's a peering agreement wherein one smaller network pays a larger network for a connection.

What about bigger providers though? Who pays who in that case?

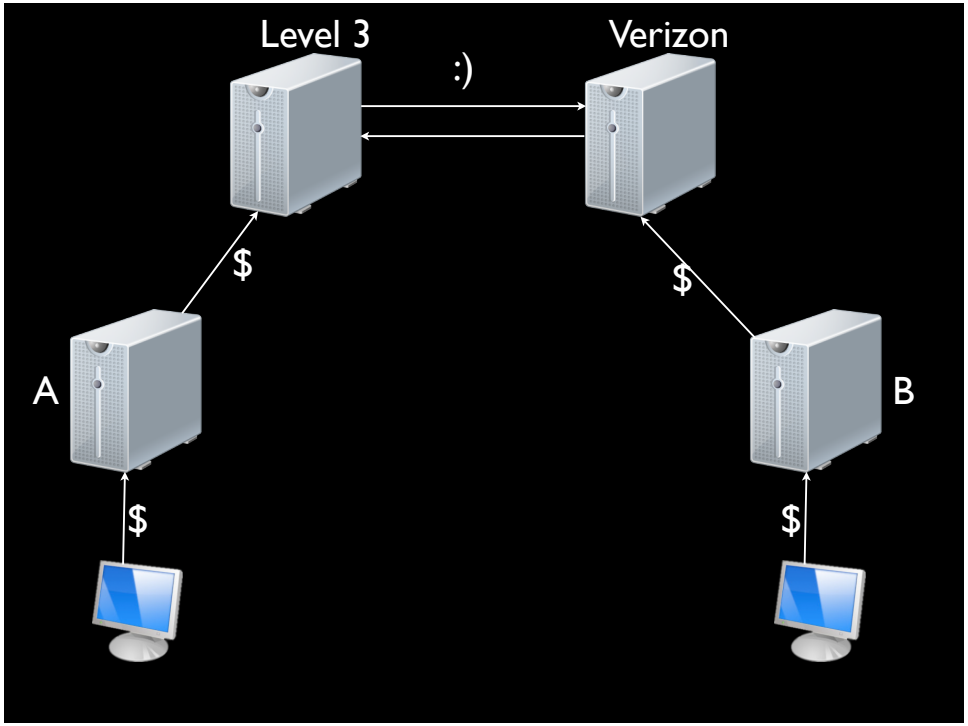

They have a true peering agreement. They connect at an Internet Exchange, an IX, but no money changes hands.

They do this because they are in the Tier 1 provider club.

Tier 1

Level 3, Verizon, and about a dozen other providers are Tier 1 providers.

The definition of a Tier 1 provider is a provider that does not have to pay transit to reach any other provider. They carry the bulk of internet traffic and essentially all traffic over long distances.

Any network connecting to a Tier 1 provider has to pay transit to get their traffic on the network.

This means small ISPs that put customers on the internet pay bigger providers. But also ISPs that host websites pay providers an amount roughly proportional to the amount of bandwidth they consume.

That's the way it traditionally works, but this is changing.

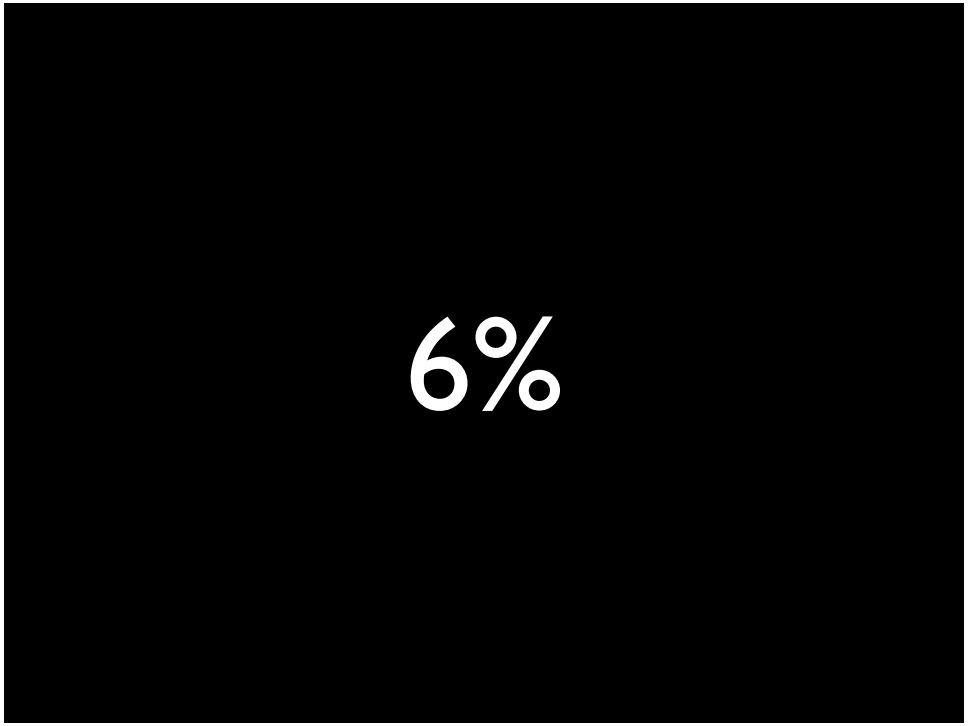

Things are different now because YouTube is responsible for about 6% all internet traffic.

ISPs are consolidating and so are large websites. Large websites are larger and more important than ever.

They use a lot of bandwidth, but they've made themselves an indispensable part of the internet. ISP customers demand access to YouTube, so ISPs demand access to it from larger providers, and Tier 1 providers strike friendlier and friendlier deals with YouTube.

So friendly, in fact, that YouTube no longer pays a bandwidth bill. That's down from a rumored million dollars a day in bandwidth.

And it's possible to imagine a future wherein Level 3 pays YouTube some fee for better and faster access to their content. Then they could advertise that they provide the best access to the largest content provider on the internet.

This is very, very new. I'm not sure what it means. I'm not sure that anyone knows. But it is a shift in power. 5 or 10 years ago, companies with fat pipes held a lot more power than they do today. Now, content providers are gaining more power.

other:
This could be a good thing-- If a morally bankrupt backbone or ISP tried to do something shady, say limit bandwidth for Google Voice, then Google could strike back and deprioritize YouTube on their network. Content providers have leverage against providers.

But maybe it's a bad thing. Maybe Google could demand tribute from every ISP in the world and threaten to block it if they don't.

There are not currently laws regulating these behaviors.

Who Owns the Internet?

The internet is just its interconnections and the policies that keep it running.

The low levels we've looked are controlled by a mixture of public and private interests. There are stewardship organizations, governments, and corporations.

Remember that everyone has to agree on a lot of things for the internet to work at all. Organizations with very disparate interests have managed to come together to keep the internet running.

Luckily the balance struck between these interests seems to generally work.

questions?