



How the Internet Goes Wrong

Jon Crowcroft,

<http://www.cl.cam.ac.uk/~jac22>



Let's look at what can go wrong

- We take the Internet for granted
- Until something doesn't work!


- Let's look at three common problems
 1. Why can't I get to a web site?
 2. Why's my download suddenly go slow?
 3. Why's my computer just got virused?



Outage!

- Aside from wires coming unplugged, or computers crashing (yours or theirs) there are several reasons you might not be able to get to a website :-
 1. Names
 2. Addresses
 3. Routes

The Wires...





Flakey hardware improved by Smart Software

- Wires get broken
 - People kick cables out
 - Turn computers off
 - Power fails
- Can we make up for this by making the whole
 - Smarter than the sum of the parts?
- Yes - control software!!

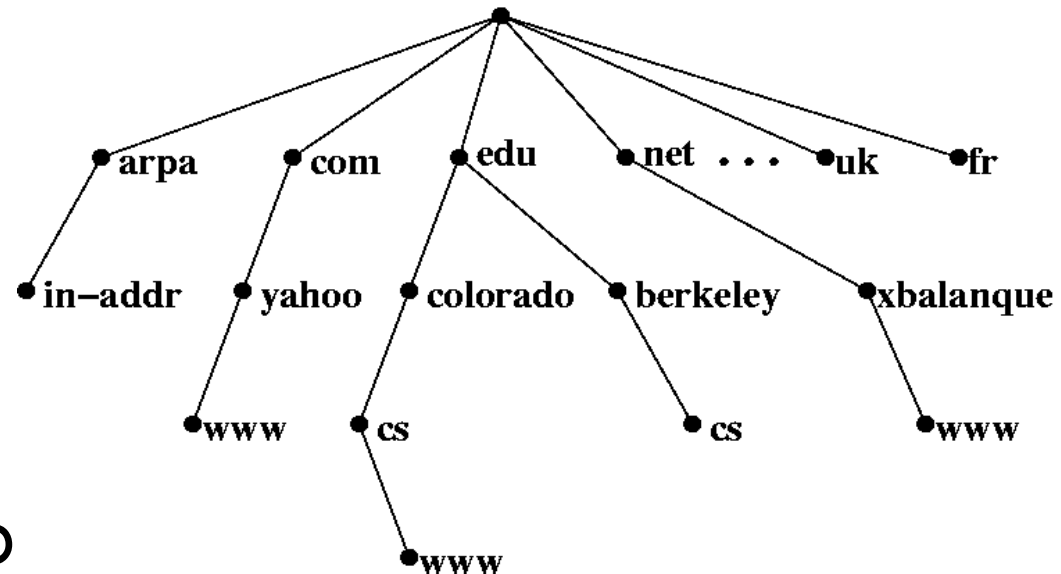
The Domain Name System

- When you type (or cut&paste) www.facebook.com, what you want...

- a "lookup" is done to

- Find the
■ IP Address

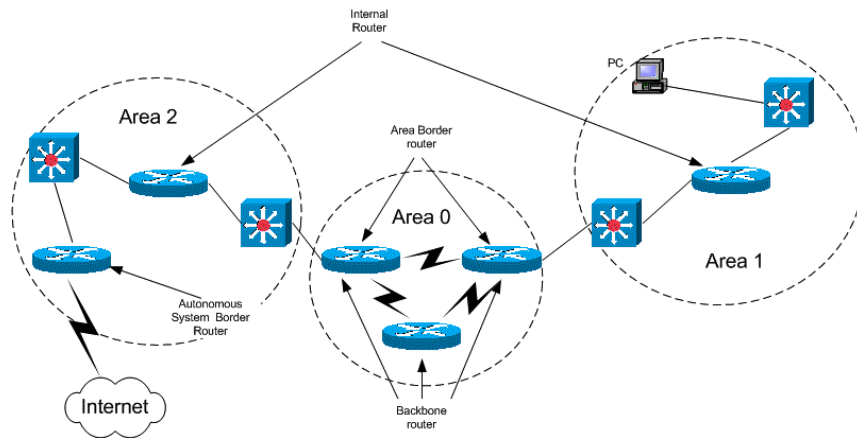
- Which is
■ **Where** it is



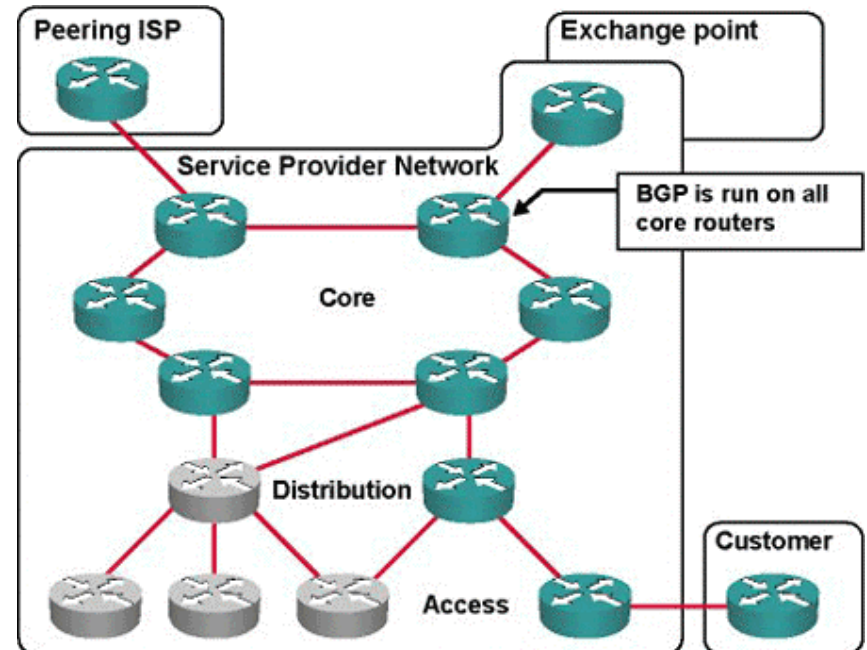
- The DNS may b
- Or you might just type something slightly wrong


Routing

- An address is where, but then you need a map and a compass to find the route




- The net does this
- For you in a
- Distributed way
- Which can go wrong!





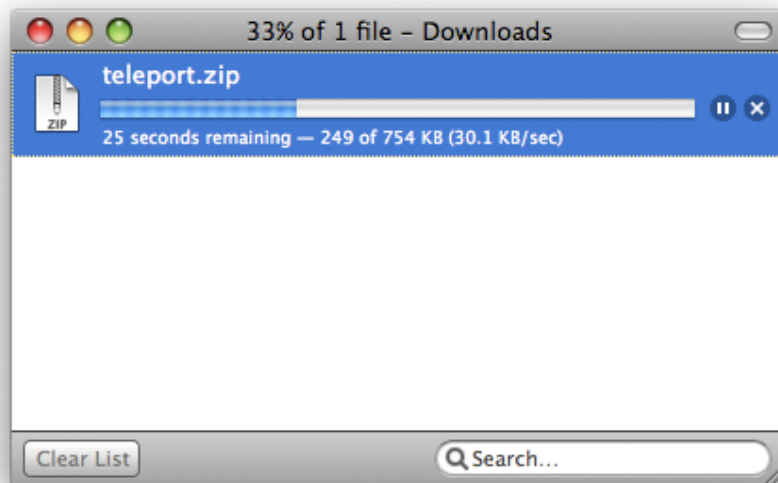
Dynamics

- Even as things change, software can keep track



Congestion

- Traffic jams can happen anywhere...on the internet too...





The Internet is shared, like roads

- Not so much like railways or flight paths
- So you have to wait your turn
- If there's a lot of users, the wait gets longer
- This is "implemented" by software in your computer which runs a *protocol*
- Called *TCP* - which cooperates with other computers implicitly to give a fair share...think about card games or anything where there are rounds...but where you can pass if you like
- It isn't exactly like that as it would take too long in a network, so instead it uses statistics



Insecurity!

- You may program your computer,
- But most the programmes you use were written by someone else (Microsoft, Apple, open source contributors)
- When you download a programme, how do you know who really wrote it, and what they really want to do with it?
- This is as true on your cell phone as it is on a notebook.
- This is true for Facebook Apps (and photo tagging) that invade your privacy.

Why do people write "malware"

- Sometimes they want to steal your ideas or your money
- But other times they want to use your computer to do things like
 - Spam
 - Botnets/ddos attacks
- Really bad guys pretend
- To be trying to help:





The Internet is quite complicated

- It isn't (usually) complex -
 - it's just made of a lot of pieces, each of which is really very simple.
 - For an "end to end" path to work
 - Properly, as expected, and to perform well
 - All the pieces have to function correctly
 - Amazingly, it does work most the time
 - Largely because we have got a lot better at designing and building computer software and hardware in the last 10-20 years
 - But there's a lot more to do still!



Highly Optimized Tolerance

- There are two possible problems that present a high risk

 Topological

 Temporal



Topology Problems

- The Internet exhibits scale-freeness
 - At many levels (link level and web level)
 - It also exhibits clustering
 - So we have small world....
 - which is good (for finding stuff)
 - But bad for attacks, due to "hub-iness"
 - Nodes of high betweenness (or spectral centrality) have to be protected/hardened
 - Its software, doh, and it's a net
 - so it isn't just thick lead walls and airgaps 😊
 - We can reboot 😊



Temporal Problems

- There are lots of synchronisation phenomena
 - Some happen all the time - the routing system is driven by clocks, for example
 - This can self-synchronise
 - The topology makes this more likely, not less
 - Bad stuff can synchronise with the routing system - an *scanning* attack can oscillate and end up blocking routing updates,
 - Leads to breaking connectivity, even if capacity, per se, wasn't in question
 - We can put in randomness to defend against this



Take Homes

- Risk if Internet Breaks is very bad indeed, if more other utilities come to depend on it for control -
- this should be prevented by legal/policy means - we need diverse networks (for energy, transport, food, knowledge), we do NOT want to couple them closely (or at all)
- The Internet itself could be made more robust/resilient, esp. to emergent bad behaviour
- People are aware of this in the tech community 😊