



How the Internet Was Invented

By 1990 the internet moved data between computers just fine. It still had no shared way to find a document, or to link one to another. This video follows the three inventions that closed that gap, in the order they actually solved the problem rather than the order anyone announced them.

First, moving a file at all: Abhay Bhushan's RFC 114 (16 April 1971, MIT Project MAC) gave ARPANET its first File Transfer Protocol, reworked that same year in RFC 172. FTP moved a whole file between two named machines - but only if you already knew exactly where to look.

Second, browsing without knowing the name in advance: Minnesota's Mark McCahill, Farhad Anklesaria, Paul Lindner, Daniel Torrey and Bob Alberti released Gopher in 1991, built to replace a stalled 1990 proposal for a campus-wide mainframe information system. Gopher organised files into menus a user could descend - and at its peak ran on 6,958 servers worldwide. It was still a hierarchy, with no way to link one document to another.

Third, the conceptual leap: Tim Berners-Lee's 1980 Enquire notebook program at CERN, his March 1989 proposal 'Information Management: A Proposal', and Robert Cailliau's November 1990 formal proposal that got it resourced. By Christmas 1990, Berners-Lee had HTTP 0.9, HTML, the WorldWideWeb browser/editor and the first server running on his NeXT machine. info.cern.ch went public on 6 August 1991.

The video closes on why the Web displaced Gopher by the mid-1990s: NCSA Mosaic's graphical browser (January-April 1993), CERN's decision to put the Web in the public domain royalty-free (30 April 1993), and the University of Minnesota's decision to start charging commercial Gopher users a licence fee that same year - which Gopher team member Robert Alberti said 'socially killed gopher'.


The video focuses on the specific inventions and the people and institutions behind them - CERN, Berners-Lee, Cailliau, the University of Minnesota's Gopher team - rather than the web's later commercialisation.

Educational documentary. Not financial or investment advice.

Chapters

1. Before the Web
2. Moving Files: ARPANET and FTP
3. FTP's Limits
4. Browsing by Menu: Gopher is Born
5. Why Gopher Won, Then Lost
6. The Conceptual Leap: Hypertext at CERN
7. Enquire to Proposal
8. Building the Stack: HTTP, HTML and the First Server
9. info.cern.ch Goes Public
10. Three Layers, One Problem Each
11. Why the Web Displaced Gopher

1. Before the Web



By 1990, the internet already worked. Computers on opposite sides of the world could send each other data over the same network, reliably, every day.


This is a sketch of the ARPANET, the US research network that was the internet's direct ancestor, drawn in December 1969 when it connected its first four computers. By 1990, that same kind of network linked thousands of machines worldwide.


What it still could not do was let you find a document you didn't already know the location of, or link one document to another.

Three separate problems had to be solved before the internet became the web. First, how do you move a file from one machine to another at all. Second, once you can move files, how do you browse what's out there without already knowing the exact file name. Third, how do you connect one document to another, directly, so a reader can simply follow a link.

This video follows those three problems in the order they were actually solved, not the order anyone announced them.

2. Moving Files: ARPANET and FTP





The first problem was the simplest to state and the hardest to take for granted today: two separate computers, owned by two separate institutions, had no agreed way to hand a file from one to the other.

On the 16th of April 1971, Abhay Bhushan of MIT's Project MAC published RFC 114, called A File Transfer Protocol. A protocol is an agreed set of rules two computers follow so they can talk to each other reliably. An RFC is simply the document format the ARPANET, and later the internet, used to write one down and propose it. RFC 114 was the first protocol built specifically to let ARPANET machines send and fetch files from one another.

This is the ARPANET's logical map from March 1977, the network that FTP ran over. Already dozens of machines, at universities and research labs across the United States, were connected to it.

What FTP actually did was narrow and specific. You named a machine, you named a file on that machine, and the protocol moved a copy of it to you.

That structure is point to point. One machine, talking to one other machine it already knows the address and the path of, with nothing linking that file to any other file anywhere else on the network.


This is one of the ARPANET's Interface Message Processors from 1969, the specialised hardware that actually carried traffic like an FTP transfer between the machines plugged into it. Every file FTP moved in the network's first years passed through a box like this one.

Later that same year, FTP itself was reworked and formalised.

RFC 172, dated the 23rd of June 1971, set FTP out properly for the first time. A companion document, RFC 171, handled the lower-level work of actually transferring bytes between machines. Between them, the two documents gave ARPANET its first properly standardised way to move a file.

FTP solved real-sharing files between machines. What it never solved is the subject of the next chapter.

3. FTP's Limits



FTP's entire design assumes you already know exactly what you want.

To fetch a file with FTP, you need the machine's address and the file's exact name and location on that machine, in advance. There is no menu to browse. There is no way to search. And there is no way for one file to point to another – nothing inside an FTP transfer connects what you just downloaded to anything else on the network.

That gap – no way to browse, no way to link – is exactly what the next two inventions, twenty years apart, each tried to close.

4. Browsing by Menu: Gopher is Born



The first attempt to close that gap didn't try to link documents together at all. It organised them.

Gopher works as a branching menu tree. You pick an item from a list, and that choice opens another list. Choose again, and another list opens below that – descending, screen by screen, until you finally reach an actual document.

That design came out of the University of Minnesota in 1991, and it came from solving a university's own bureaucratic problem.

This is the University of Minnesota campus, home to the team that built Gopher. Gopher grew directly out of a problem this campus already had.

In 1990, university officials had proposed a Campus Wide Information System, a single mainframe database meant to hold the whole university's information. The committee behind it made little progress.

The Minnesota Computing History Project, run by the university itself, describes that proposal's progress as "a classic design-by-committee monstrosity" – years of meetings, and no working system.

This all happened in Minneapolis, Minnesota, inside the University of Minnesota's own Microcomputer Center, nowhere near CERN or MIT.

A small team built something that actually worked instead, and gave it an unglamorous but accurate name.

Gopher's protocol was first issued by the University of Minnesota's Microcomputer Center in 1991. Its credited team: Mark McCahill as project lead, Farhad Anklesaria as its primary designer, with Paul Lindner, Daniel Torrey and Bob Alberti.

A menu tree that worked, built by five people, instead of a mainframe system that didn't exist. That was Gopher's whole pitch – and for a few years, it worked extremely well.

5. Why Gopher Won, Then Lost



Gopher's menus were simple enough to run on almost anything, and that simplicity is exactly why it spread so fast.

A Gopher server needed little processing power and very little bandwidth, because it only ever sent plain text menus. Any university department, anywhere, could stand one up in an afternoon, and through 1991 and 1992, thousands of them did, across campuses worldwide.

At its peak, 6,958 Gopher servers were running worldwide. That was about two and a half times its closest rival protocol at the time, according to the University of Minnesota's own computing-history project.

That growth hid the one thing Gopher's design had never actually fixed.

Descending a menu tree only gets you to documents someone already filed under that branch. A menu has no way to connect two documents that live in completely different branches, however closely related their ideas actually are. Gopher organised the library. It still couldn't let one book cite another with a click.

6. The Conceptual Leap: Hypertext at CERN



While Gopher's menus were spreading across American campuses, a researcher at CERN, the particle physics laboratory outside Geneva, had already written down a different answer to the same problem. He had done it two years before Gopher even existed.

On the 12th of March 1989, Tim Berners-Lee circulated a document inside CERN titled Information Management: A Proposal. Its conclusion states the idea plainly: "We should work toward a universal

linked information system, in which generality and portability are more important than fancy graphics techniques."

Not a menu a user descends. A link, inside the document itself, that can point to any other document on the network.

That single sentence is the conceptual leap this video's title refers to – hypertext, a word for text that contains its own links to other text, read directly rather than reached by menu.

This is CERN's site at Meyrin, near Geneva, Switzerland, where Berners-Lee wrote the 1989 proposal and, with it, the idea the rest of this video follows to its conclusion.

CERN sits on the border of Geneva, Switzerland – a physics laboratory, not a computing company, which is exactly why this idea arrived from an unexpected direction.

A good idea on paper is not a working system, and CERN's management did not fund it immediately.

7. Enquire to Proposal



Berners-Lee's 1989 proposal did not appear from nowhere. It grew out of something he had already built, almost a decade earlier.

In 1980, Berners-Lee wrote a notebook program at CERN called Enquire. It let its user link any one entry to any other entry freely – a direct, working ancestor of the idea the 1989 proposal would later generalise to the whole network. That same linking idea reappeared later, aimed now at every document CERN held rather than one researcher's private notes. In November 1990, engineer Robert Cailliau co-wrote the formal management proposal that finally got the project resourced.

Three dates, one continuous idea. A private notebook in 1980. A proposal to CERN in March 1989. A funded project in November 1990.

Robert Cailliau's contribution was not the idea itself – that was Berners-Lee's – but turning it into a document CERN's management would actually approve and staff. Two names are on the project from this point on, not one.

With a funded project and a deadline, Berners-Lee spent the next month building the actual system the proposal had only described.

8. Building the Stack: HTTP, HTML and the First Server



Having an approved idea is not the same as having working code, and Berners-Lee built an entire working stack in a matter of weeks.

By Christmas 1990, Berners-Lee had four things built and running on his NeXT computer at CERN. HTTP 0.9, the protocol for requesting and sending a document over the network. HTML, the format for writing a document with links inside it. A program called WorldWideWeb that worked as both a browser and an editor. And the first server, later called CERN httpd.

Put together, those pieces did what no menu tree ever could.

Any document, on any server, could now contain a link straight to any other document on any other server. That is a web of documents pointing directly at each other, not a tree a user has to climb down one level at a time.

Two new words carried that idea, and both deserve a plain definition before this video uses them again.

HTTP is the set of rules a browser and a server use to ask for a page and send it back. HTML is the format of the page itself – ordinary text, with instructions marking which words are a link, and to where. Together, they are the two halves of the same job: HTTP delivers the document, HTML tells the browser what's inside it.

By the end of 1990, the pieces existed. They were still only running inside one laboratory.

9. info.cern.ch Goes Public



For the first eight months of 1991, hypertext at CERN was an internal project, visible only to people already on CERN's own network.

On the 6th of August 1991, Berners-Lee posted a short summary of the WorldWideWeb project to alt.hypertext, a Usenet newsgroup for people already interested in linked documents, and invited anyone reading it to take a look. That post made info.cern.ch, the first website, reachable from outside CERN for the first time.

Other institutions started standing up their own servers soon after, inside CERN's own research community first.

Through the rest of 1991, other European research sites began running their own web servers. The first server outside Europe went live on the 12th of December 1991, at the Stanford Linear Accelerator Center in California, built by physicist Paul Kunz after a visit to CERN that September, so SLAC's physicists could search the particle-physics archive it hosted.

What had been one laboratory's internal tool was now, however slowly, becoming a network of servers that any of them could link to.

10. Three Layers, One Problem Each



Stop here, because it is worth seeing the whole shape of the last nine chapters at once.

FTP, 1971, solved moving a file from one named machine to another. Gopher, 1991, solved browsing those files without knowing their exact names, by organising them into menus. Hypertext, HTTP and HTML, 1989 to 1991, solved linking any document directly to any other, which neither of the first two layers could do at all.

Each layer did one job the layer before it couldn't.

FTP moves a file. Gopher organises files into menus. Hypertext links documents to each other directly. Three different problems, three different answers, arriving in that order for a reason. You cannot browse what you cannot move. And you cannot usefully link what you cannot already find.

Gopher and the Web, for a few years, coexisted. Only one of them was still growing by the middle of the decade.

11. Why the Web Displaced Gopher



The Web's advantage over Gopher was always the direct link. In 1993, two separate decisions made that advantage decisive.

NCSA Mosaic, the first widely used graphical web browser, had its first pre-release on the 23rd of January 1993 and its first full release on the 21st of April 1993, built by Marc Andreessen and Eric Bina at the University of Illinois. For the first time, a web page's links could sit next to pictures, in one window, rather than a screen of plain text.

Nine days after Mosaic's full release, CERN made its own decision, and the two together set the direction for the rest of the decade.

On the 30th of April 1993, CERN published a statement putting the World Wide Web's software into the public domain, royalty-free. Anyone could now run a web server or build a browser without paying CERN anything at all.

The University of Minnesota chose the opposite path with Gopher, in the same year.

Mosaic came out of Illinois. Gopher's licensing decision came out of Minnesota. Two American university towns, far from CERN's Geneva, ended up deciding the web's next chapter between them.

In 1993, at the conference GopherCon, the University of Minnesota announced it would start charging commercial Gopher users a licensing fee. Gopher team member Robert Alberti later said the announcement, in his own words, "socially killed gopher" – even though the technology itself had not changed at all.

A free, linked, increasingly graphical web, against a menu system that had just started charging the community that built its popularity. Within two years, nearly everyone had made the same choice.

By the mid-1990s, new servers overwhelmingly ran the Web's protocols rather than Gopher's, and graphical browsers modelled on Mosaic became the ordinary way most people first went online. Gopher servers did not vanish immediately, but new growth had moved almost entirely to the Web.

Three layers, each solving the one problem the last layer left behind – and the layer that could link anything to anything else was always going to be the one the rest of the network built on.

Sources and credits

Primary sources

- IETF RFC Editor / Datatracker, RFC 114 'A File Transfer Protocol', Abhay Bhushan, 16 April 1971 - the first ARPANET file transfer protocol.
- IETF RFC Editor, RFC 171 and RFC 172, 23 June 1971 - the companion data-transfer and file-transfer protocol documents that reworked FTP.
- W3.org, Tim Berners-Lee, 'Information Management: A Proposal', 12 March 1989, fetched and quoted directly 4 October 2026.
- CERN timeline (timeline-upgrade.web.cern.ch) and physicsworld.com - the first web server outside Europe, SLAC, 12 December 1991.
- IETF, RFC 1436, 'The Internet Gopher Protocol' - confirms Gopher 'first issued by the Microcomputer Center at the University of Minnesota in 1991' and its six-person author list.
- Minnesota Computing History Project (mncomputinghistory.com/gopher-protocol/), run by the University of Minnesota - Gopher's team credits, the 1990 Campus Wide Information System proposal it replaced, the 6,958-server peak count, the 1993 licensing-fee announcement and Robert Alberti's 'socially killed gopher' quote.
- Wikipedia, 'NCSA Mosaic', cross-checked against webdesignmuseum.org and CERN's own timeline (node 793) - Mosaic's 23 January 1993 pre-release and 21 April 1993 full release.
- CERN, 'Twenty years of a free, open web' (home.cern) - the 30 April 1993 royalty-free public-domain release, corroborated by thegovlab.org and i-programmer.info.
- MinnPost, 'The rise and fall of the Gopher protocol' (2016), corroborating the 1993 licensing-fee backlash independently of the University's own history page.

Photo credits (Wikimedia Commons)

- Sketch of the ARPANET, December 1969 - the pre-web internet this video opens on: DARPA, Public domain - https://commons.wikimedia.org/wiki/File%3AA_sketch_of_the_ARPANET_in_December_1969.png
- ARPANET logical map, March 1977 - the network FTP ran over: ARPANET, Public domain - https://commons.wikimedia.org/wiki/File%3AArpanet_logical_map%2C_march_1977.png
- BBN Interface Message Processor (IMP), the ARPANET hardware that moved files between machines, 1969: ARPANET IMP, Public domain - https://commons.wikimedia.org/wiki/File%3A1969_ARPANET_BBN_IMP.jpg
- University of Minnesota campus, Minneapolis - home of the Gopher team: Unknown authorUnknown author, Public domain - https://commons.wikimedia.org/wiki/File%3AUniversity_of_Minnesota_campus%2C_1890.jpg
- CERN site, Meyrin/Geneva, Switzerland - where Berners-Lee and Cailliau built the Web: Brücke-Osteuropa, CC0 - https://commons.wikimedia.org/wiki/File%3ACERN-aerial_1.jpg

Not regulated financial advice.