

httpd, the first web daemon, written in Geneva

[Click here](#) for its Wikipedia entry

CERN httpd (also known as W3C httpd) was the **world's first web server software**, written in 1990 by Sir Tim Berners-Lee and his colleagues at CERN in Geneva, Switzerland.

The "d" stands for **daemon**, a Unix term for a program that runs continuously in the background, listening for and responding to incoming network requests.

Quick Facts about CERN httpd

- **The Authors:** Originally written by [Tim Berners-Lee](#), it was later co-developed with Ari Luotonen and Henrik Frystyk Nielsen.
- **The Environment:** It was written in the C programming language and initially developed on a **NeXTcube workstation** running the NeXTSTEP operating system.
- **The Port:** Berners-Lee originally mapped the server to listen on port 2784 (a number based on his parents' phone number) before it was eventually standardized to **port 80**.
- **The First Website:** On December 20, 1990, it went live to serve the world's very first website: info.cern.ch, which hosted documentation explaining how the World Wide Web worked.
- **Historical Note:** The physical NeXTcube machine used to run the server in Geneva famously had a hand-written red label on it that warned: *"This machine is a server. DO NOT POWER IT DOWN!!"*

The software was later handed over to the World Wide Web Consortium (W3C) and officially retired in July 1996 with version 3.0A, paving the way for future servers like Apache, Nginx, and modern web infrastructure.

If you are researching early internet history, let me know if you would like to explore:

- How **CERN httpd** laid the structural groundwork for the modern **Apache HTTP Server**
- The details behind the **NeXTSTEP operating system** that made this rapid development possible
- How the **first web browser** (which was also an editor) interacted with this daemon

Was it adapted by SLAC in their December website at Stanford

Yes, the **CERN httpd software was directly adapted by SLAC** (Stanford Linear Accelerator Centre) to launch the very first web server in North America on **December 12, 1991**.

The transition from Geneva to Stanford happened through a specific chain of events:

The NeXT Connection

In September 1991, SLAC physicist **Paul Kunz** visited CERN and met with Tim Berners-Lee. Because Kunz was also a fellow enthusiast of the **NeXTcube computer**, Berners-Lee was able to give him a live demonstration of the World Wide Web software directly on the NeXTSTEP operating system.

Kunz immediately recognized its potential, copied the CERN software, and brought it back to Stanford on a floppy disk.

The World's First "Killer App"

Upon returning to California, Kunz teamed up with SLAC librarian **Louise Addis** and software developer **Terry Hung**. Instead of just serving basic text files, they adapted the CERN httpd daemon to act as a front-end interface for **SPIRES** (Stanford Public Information Retrieval System)—a massive database used by global physicists to search for high-energy physics preprints and literature.

When the website officially went live in December 1991, it consisted of just three lines of text and two links:

1. A link to contact SLAC evolution teams.
2. A link that allowed users to actively **search the SPIRES database**.

Tim Berners-Lee famously called this specific adaptation the World Wide Web's **first "killer app"** because it proved the web could interactively fetch live data from massive, pre-existing databases rather than just displaying static documents.

Legacy

On December 13, 1991 (the day after it went live), Berners-Lee sent out an email to early internet newsgroups proudly announcing: *"There is an experimental W3 server for the SPIRES High energy Physics preprint database, thanks to Terry Hung, Paul Kunz and Louise Addis of SLAC."*

The original files and their 1991 timestamps were meticulously saved by SLAC staff and can still be browsed today via the official Stanford Wayback system.

**** End of Page**