



Oral History of Jean-François Groff

Interviewed by:
Marc Weber

Recorded: November 11, 2009
Mountain View, California

CHM Reference number: X5612.2010

© 2009 Computer History Museum

Marc Weber: I'm Marc Weber of the Computer History Museum, and I'm here today with Jean-François Groff, who is a Web pioneer at CERN. And I thank you for joining us.

Jean-François Groff: Thank you.

Weber: So if you could start just by giving a little bit of your background and what brought you to CERN.

Groff: Well, I happen to be French, and in those days, France still had a mandatory military service. And this is how I came to CERN, actually to use my brains instead of using my feet. So I was a telecom engineer. I had just finished, just graduated from the telecom school in Paris. And I came to CERN to be a "co-operant." I was supposed to work on software that was analyzing data from the particle detectors. And that was very boring, because by the time I arrived, everybody had already done the job I was supposed to do. So I met Tim Berners-Lee, who was working next door, and started to work with him on the World Wide Web.

Weber: Would you back up a little bit and just say where you were born and when you got interested in computers.

Groff: I started in computers very early. My father happened to be a computer engineer as well. He was working in a corporation called Saint-Gobain, which is a big French company manufacturing construction materials. And so we he was installing mainframe computers there. He actually even wrote an emulator of a former version of the operating system of IBM so he could run older programs on the new machines. So he introduced me to computers very early on, I'd say when I was seven or eight years old. And then I started programming on personal computers at home. I had an Amstrad, and then I moved up to Atari, and then I moved up to-- I don't know, Sun Microsystems, probably.

Weber: And so then you went to get your degree in telecommunication.

Groff: Yes, that's right. Yeah.

Weber: So you ended up at CERN, and how did you meet Tim Berners-Lee?

Groff: Tim happened to be a neighbor, actually. I was introduced to Tim by Ben Segal. Ben Segal is the man who introduced the internet at CERN, with a big battle against the OSI and X.25 protocols. So introducing the IP protocol, that was very nice. So we had the internet by the time I arrived at CERN, that was late 1990. And they had the internet running there for about a year. So it was one of the biggest nodes in Europe, one of the best-connected nodes. Which means they probably had one megabit per second of worldwide connectivity. And Tim had been working on RPC systems, that is a way for computers from different brands to talk to each other, to send commands to each other. So it was a standardized API for calling a function on a computer from another computer. And after that work, he had started to build a hypertext system on top of these networking concepts. And that was the World Wide Web. I arrived at CERN on December 6, 1990. And you may remember from Tim's recollections that he completed his first NeXT editor for the Web at Christmas, 1990. So I saw the very, very early working product from his vision. And then I helped him to make this software run on other kinds of machines, because obviously, a hypertext system that would be running only on the NeXT would only have limited interest.

Weber: Do you remember your first impressions of Tim?

Groff: Very enthusiastic, brilliant. Most people didn't understand what he was up to. He spoke very fast, that didn't help. But the enthusiasm kind of carried people with him anyway, that's how he convinced his boss to let him do what he wanted to do. Although that was always a nonofficial project at CERN, including my participation. I mean, it took eight months until we convinced my boss to officially assign me to the World Wide Web.

Weber: And who was your boss?

Groff: My boss was called Peggy Rimer. She was the wife of Tim's boss, who was Mike Sendall. Yeah, that helps.

Weber: And so for the first eight months, you were working mostly on the project for her?

Groff: I was supposed to be working for, as I said, for some software regarding the -- I'm not going to launch into an explanation of what CERN does, that would take forever. I was working on physics software, basically. I mean, data acquisition software, filtering, porting some pre-existing software from [the] Pascal language to the C language, making it work on new machines, porting it from VMS to UNIX. I had a pretty strong background in UNIX. I had translated the reference book of the C language, "The C Programming Language," by Kernighan and Ritchie. So I was qualified to do that kind of job. But the Web excited me immediately, because I had been studying telecommunications and I always had the passion for classification, information retrieval, databases. And what Tim had written first on paper, his famous proposal, and then the prototype he had actually implemented very quickly -- that struck me as the right way to do things. So I was immediately taken to helping him. And we worked very, very well together.

Weber: You shared an office with him, right?

Groff: Not immediately, no. I used to be in my office. He was in another office a few doors on the same corridor. And then yeah, after we were officially put together, I mean, after I was officially assigned to work with him, that was August, '91, then we moved to another office above the computer center in Building 31. So originally, I was in Building 513 and just two rooms away from Tim.

Weber: And then you moved into room 007, think.

Groff: That's right, the famous room.

<crew talk>

Weber: Could you say, "And then we moved into room 007," please.

Groff: Yeah, so at that time, when I went to work officially with Tim, we moved to Building 513, [floor 1,] room 007, on the first floor, yeah, just above the computer center.

Weber: But Ben Segal introduced you in the cafeteria or something?

Groff: I don't remember.

Weber: Okay. So what was the first task that you did for Tim on the Web?

Groff: Oh, on the first day, he showed me his notes on paper, and he asked me if I had any suggestions. So I took them home, I spent the night looking at that and making notes. And I gave them back to him the next day and he said, "Oh, you actually understand what I'm doing." I said, "Yeah, and it's cool." So that was the first thing I did, yeah. I gave some suggestions, notably about the links being in one direction only, being only forward links. I thought it would be nice to have backward links as well. Of course, all the hypertext people thought it would be nice. But it turned out that forward links were one of the things that made the Web scale very well. Another suggestion was that we could have links pointing to various versions of the information. That was actually implemented in the code, but nobody used it in browsers. It would have been too complex anyway.

Weber: But explain how that would work in practice.

Groff: In practice, when you look at the Webpage and click on the link, instead of going directly to the destination, for some links, you could have a choice. You could say for example, print this document or send it to a friend.

Weber: You're talking about just type--

Groff: You could do an action on top of the link.

Weber: But that would be separate from typed links.

Groff: It could be used for typed links as well, yeah. Typed link saying 'give me the text version of this or give me the picture version.'

Weber: So explain sort of the menu of different links you wanted people to have. You would go over a link, and then you would see what?

Groff: It's whatever the author wanted to specify. There are cases in technical documentation, for example, which was a big application of the Web in the early days. As an aside, we wrote the Web's documentation with the Web itself, the whole files that explain how the Web worked. So for example, in technical documentation, when you point to a function name, you might want to go to the specification of that function, or you might want to go to the source code. It's up to the author to say what he wants to link to, we're just providing a mechanism to allow this.

Weber: So the first thing, you read the proposal, you made suggestions, and then just tell the story.

Groff: I can't remember now what kind of suggestions I made. I just know that Tim liked what I said, and we worked together. He also was in need of people who could help him, because as you know, CERN was not giving him officially any resources. It's good enough that CERN was letting him do what he wanted.

Weber: But then the first concrete tasks that you did for him?

Groff: You mean programming? I don't remember what was really the first one. I don't remember. I could tell you a few things that I did in the early months. One thing I did that was very funny was we had the line mode browser, which was very crude, but working. So people could look at the Web without having a Web browser. Because there were no Web browsers except the NeXT thing, okay? So the line mode browser was just a program running on the terminal where you could go to the home page and it would be displayed as text with a crude style sheet mechanism. And near each hypertext link, you had a number. And if you typed that number, enter, you would jump to the destination of the link. And of course, we had some commands to manage the history and to go back and forward, and also to go to a URL directly by typing it. I don't remember if we had bookmarks, but we certainly had the history, that's for sure. So that line mode browser was the way for people to discover the Web without having a NeXT. It was meant as a temporary system while we were working on graphical browsers for the PC, for the Mac, and for the X Windows System. [?? History then, ??] it so happens that these graphical browsers were developed by teams outside of CERN because we could never get the resources to hire a couple of people to write that. It's that simple. So this line mode browser, it was a bit of a chicken and egg problem, because to use it, you had to download the software first and install it and possibly compile it. So what I did in order to let people access the Web without installing any software, on the NeXT cube that was running the first Web server, I modified the login process, it's a UNIX process that responds to login requests and does the authentication of a user. So that when somebody would login to that computer from the internet, if he provided a username we checked his identity. And if he didn't provide a username, we would send him directly into the Web, into the home page. So that was a little hack that I'm proud of. So with that in place, we could go to conferences and tell people if you want to learn more about the Web, you just to a terminal on any machine connected to the internet and you type **telnet info.cern.ch**, and bang, you would be welcomed with the project page and then you could navigate through the documentation and download the software if you wanted.

Weber: So at that point you could download the line mode browser as well as the NeXT.

Groff: Yes. Anything. And the documentation and the library.

Weber: But the telnet interface gave you the same functionality as the line mode browser, effectively.

Groff: The telnet was a way to run the line mode browser remotely, basically. Instead of installing it on your machine, you would be looking at this from a server somewhere.

Weber: Right. So there was no need to install it on your machine at that point.

Groff: But it had a perverse effect. Because a lot of people who discovered the Web after we talked about it, discovered the Web through the line mode browser and therefore they only saw the network hypertext reading of information. That was interesting enough that people weren't really interested in the

key feature of the NeXT browser, which was also a hypertext editor. And I think looking back, this very fact was one of the factors that none of the early browsers, after what we did at CERN, none of these browsers was read/write. We had to wait until commercial products, such as Netscape Gold, in '95 or '96, to have an editing capability. And still, it was not as smooth, because you would have to switch from reading mode to editing mode. Whereas on the NeXT machine, you could type as you were reading. It was just seamless. But that was lost.

Weber: Well, describe a bit the use of the NeXT browser editor and what the vision was for combining those features.

Groff: Well, the NeXT editor was, we call it now a prototype, but in fact, it was more advanced than what came out later. And I think it's because the NeXT machine itself was more advanced. Basically, in order to display text with various fonts and page layout, there was on the NeXT machine a text object that you could grab and customize. So what Tim did is take this text object and add the hypertext capability to it. But all the text editing and formatting was already in that object. So we have to thank the team of engineers up [at] NeXT who provided that. And we didn't have such a text widget readily available on the X Windows System or on the PC. So people who wrote browsers for these other platforms had more hurdles to go through. It was much easier for them to write a read-only browser. So yeah, the read-write [capability] on the NeXT, you would just power it on and it would start on your home page. And the home page was not located on a server somewhere else. It was your own document residing on your local machine, where you would write down whatever you wanted. So it's like a notebook, basically. It was a notebook. And so you could create as many pages as you wanted on your local machine. These pages could be served to other people if you published them, if you gave people rights to read them as well. So it was [a] distributed editing documentation system, basically.

Weber: You and the early users, what did you typically put on those pages?

Groff: Oh, we were basically writing the documentation for the Web. And then also once people started to fire up Web servers in various places, we kept a primitive library of all the servers in the world. That didn't last long because a couple years later they were sprouting up too fast. But then the folks at Mosaic kept a "what's new" list, which was hugely popular. And that went on for a year or two until it was saturated as well. And then what then? Yahoo came along and had a structured directory. And then you had search engines, etcetera, etcetera.

Weber: But going back to the early days, so you could also have multiple people working on the same document.

Groff: Yeah, definitely. It was, however, restricted to people sharing documents on the local network. We didn't have the protocol to authenticate and save documents on any node on the net. This would come only many years later, and it's called WebDAV. We didn't have that. The HTTP protocol in the 1.0 version didn't have support for that. I mean, it had support for creating documents remotely, okay? The "put" method, that was here, and the "delete" method. But we didn't have the support for authenticating users and giving permissions, except by checking the IP address or very crude mechanisms.

Weber: Was that planned?

Groff: It was planned, but it was a big can of worms. Because nobody wrote a read/write editor, nobody had to really go head to head against that problem. So.

Weber: And so just go through some of the main things you remember doing in the early months.

Groff: Oh well, my main task during my days at CERN working on the Web software, my main task was porting all the software libraries, I mean the software components that were on the NeXT system into a universal code library that was written in C, it's the "libwww." It didn't even have a name at the beginning, which is why in some history books, you say, "Oh, libwww was released in November of '92." No, it wasn't, you know? It was running since February '91, it just didn't have that name. So we had to take the networking components that were building on top of sockets and off the DNS system. We had the page rendering system, the parsing of HTML, and also all the URL mechanisms, history list, all that was abstracted into one software library as a package, as a toolset basically. And then in August '91, I think when we announced the World Wide Web, we also said "You can use that toolset and build whatever you want with it." So including interfaces to existing information systems. So some people took that and built interfaces to databases or to help systems. Some people took that and wrote browsers. And so yeah, the rest is history.

Weber: But in that roughly eight months, your main task was creating the code library.

Groff: Creating the library, yeah, definitely yeah. And maintaining it and helping people also. There was a lot of evangelizing. Helping people who subscribed to our mailing lists, helped them understand how the Web technology worked and how to implement things based on it, either to use it and provide content. We were really pushing people to provide content, because that's one of the areas where competing systems had an advantage. For example, Gopher was very easy to set up. If you had a [INAUDIBLE] file somewhere, you could plug in a Gopher server and your files would be visible from anywhere on the net. Whereas installing a Web server was a bit more difficult. And if you wanted the Web to be really interesting, you not only had to install the software to serve the files, but you also had to create hypertext files to describe your content. Of course, it was then much richer, because you had the hypertext navigation. And some people wrote automatic converters from let's say documentation systems, and creating links on the fly as people were browsing.

One of the very first dynamic Web services was the CERN Phonebook. The Phonebook was an application that was running on an IBM mainframe. We had to specially adapt the code library with function names that were no more than eight characters long, because that's all the IBM linker could understand. But then, we plugged into the Phonebook program. And so the way it worked is you would come to the Phonebook home page, which said type a name, type a query, just like a search engine today. And you would say, "Weber" and enter. And then your browser would send a query to that server saying "search?Weber" for example. And that query would be passed on to the IBM system. And this bridge was made with a scripting language called Rexx, which is similar to what Perl became a few years later on UNIX, a very nice language. And so this Rexx program would run on the mainframe, and the results would come back as preformatted text. And then we would extract relevant information from that search result and turn them into hyperlinks. And so we would create an HTML page on the fly and give you the results. So if you had let's say five people called Weber, we would build a list with all those names, and you could click individually on each name, which would trigger another query to grab details about this particular Weber. And that was in certainly the first half of '91, that was working. I believe it was the first dynamic information server.

Weber: And that was [INAUDIBLE]?

Groff: This was programmed jointly between Tim Berners-Lee and Bernd Pollermann who was an IBM [INAUDIBLE] writer at CERN, yes.

Weber: And in the process of creating the code library in C, was this mostly you or you with Tim?

Groff: It was Tim and I, yeah, mostly.

Weber: So 50/50 roughly?

Groff: How can I say?

Weber: Not obviously.

Groff: We were working day and night.

Weber: What was a typical day like in that period?

Groff: I was driving an old Audi 80 to CERN in the snow. And yeah. I had an IBM workstation with a big screen. And I was running my whole life in Emacs, the text editor. That could do everything except coffee. So Emacs, I had a window with the code, another window with the documentation, another window with the line mode browser itself. And we would just code things. And Tim was on his NeXT machine. I didn't have a NeXT machine, there was only one. I mean, there was a second one, but it was used in another department. But my job was to bring things out of the NeXT into other systems, so I didn't need one. And so yeah, we were trying to make things compatible and understandable by future programmers. I think it was very important to design the protocols and to design the functions so that they could be easy to use for someone who wanted to write a specific server or extensions to a system. And then after I left, some people took over the maintenance of that library, and it evolved for a number of years. For the first five years I think that all browsers in the world inherited the code from that library. Then, you know, as people sometimes do a clean break and rewrite everything, it was wiped out. But the legacy-- yeah, excuse me.

Weber: Cello was the only one that was written from scratch in that period? But otherwise, yes.

Groff: You mean Cello didn't use the CERN library at all? Wow, that's impressive.

Weber: But all the others did, including Mosaic, as I understand.

Groff: Yes.

Weber: So the first, up until August, which was when you publically released the library, correct? You were in the other building, but you were spending time in Tim's office?

Groff: Yeah, I was working in the evenings. Well, not in his office, but I was working on the code. Yeah.
CHM Ref: X5612.2010 © 2009 Computer History Museum Page 8
of 24

Weber: And where were you living?

Groff: Of course, it became more intense when I went officially there.

Weber: Where were you living at the time?

Groff: I was living in Saint-Genis, a small town across the border, just five kilometers from CERN.

Weber: And working regular hours?

Groff: Of course not, I never did that in my life.

Weber: So a typical day, you might be there until four in the morning or something, right?

Groff: Yeah. Because in the daytime, we had to go skiing.

Weber: Oh yeah. Did you have a girlfriend at the time?

Groff: Did I ever what?

Weber: Did you have a girlfriend at the time?

Groff: Yes. Yes.

Weber: You had a social life, in addition.

Groff: Yeah. What are all these questions?

Weber: Any other significant tasks that you did in that first six or eight months?

Groff: What was it? No. [???] Well, the other big thing was explaining the Web to people. Not just in the beginning, but afterwards, in '92, '93, I helped Tim and Robert, going to conferences. And sometimes, I went by myself, just to explain. We had to share the load, because everybody wanted to hear about the Web, and we were just three people. What do we do?

Weber: You remember the first conference you went to?

Groff: I'm not sure it's the first. But there's one I remember very well was in Pisa in Italy, I think it was in '92. And I went alone. And I drove middle of the night to reach that conference. And gave a speech about the Web, explained how it worked, and everybody was enthusiastic. It was a good memory, yeah. But on the other hand, you had many people, when you explained the Web, who said it'll never scale. It is interesting to see, 20 years down the road, see how far it scaled. Because if you look at the HTTP protocol as it stands now, it has almost not changed since the very early implementation. Of course, the

internet scaled up dramatically. And I believe that if we didn't have the Web, we wouldn't know what to do with all this internet bandwidth.

Weber: And when did you get involved with the virtual library? You started that, right?

Groff: As I said, we kept a catalog of the interesting sites that were coming up.

Weber: Okay, so that's what evolved into the virtual library?

Groff: Yeah, yeah. The virtual library was named that way only after I left.

Weber: Oh, okay.

Groff: But it was great to have a list of good sites to demo. We always were looking for interesting content. And suddenly, you had let's say one guy in a biochemistry department somewhere understood the Web. And suddenly, one month later, you had 50 Web servers about biochemistry, because he told all his colleagues. That was very interesting. You had nothing in a particular subject area, and then from one day to the next, you had lots of stuff, just because it's the individual process of discovering the Web at one place, and then vroom, it was sprouting.

Weber: But the early virtual library had the categories, like it did later?

Groff: Yeah.

Weber: So those are basically card catalog categories, right? Same as you would find in a card catalog?

Groff: It wasn't formal at all. It was whatever exists will-- if there are more than ten servers in a particular subject then we'll add the subject headline. There was no attempt at seriously cataloging stuff. I think Yahoo were the first ones who really did a strict ontology on the Web, and that was the base of their success. It was extremely successful, because it was well cataloged.

Weber: And the idea for the CERN library was to have subject matter experts maintain each part? Is that correct?

Groff: The idea of the CERN library?

Weber: You wanted to have subject matter experts.

Groff: Who, me? No. Not me.

Weber: No, no.

Groff: I didn't manage that.

Weber: But the virtual library, who was meant to expand and increase it?

Groff: I don't know. I didn't manage the virtual library. Yeah, well okay I remember using the virtual library, but I was not managing it. It's true that yeah, there were some people who volunteered to maintain parts of the library that related to their field of expertise. But this was never organized as a cooperation.

Weber: And Robert Cailliau, tell me about how much did you work with him in that early period?

Groff: Oh, a lot, a lot. Yeah. Excuse me, I'm just thinking of something else, because you asked what else I did. One of the things was, it didn't go anywhere but it was an interesting concept at the time. It was to make the Web run on other things than the internet, on other things than IP. There was, for example at CERN, well, you mentioned Robert. Robert was evangelizing the Web with the physics community. And the physics people, they were addicted to the Digital Equipment machines, to the VAX operating system, or sorry, VMS operating system. And Digital had its own networking protocol called DECnet, and all the physics labs in the world were interconnected through DECnet. CERN was a bit of an outsider, I mean of a pioneer, having the internet [in addition to] DECnet. Of course, they were big DECnet nodes as well. So Robert had asked me to implement the HTTP protocol over DECnet instead of over IP. That was quite a challenge. And it was in the early days of '92. I did that, and then nobody used it because TCP/IP just was adopted also by DEC. I remember we had to install, in order to compile the Web on these DEC VAX VMS machines, we had to install a third-party IP stack made by Wollongong Software, an Australian firm. So we installed the IP drivers on the IP software stack. And then on top of this stack, we could use the standard socket interface. Yeah, it was fun. It's similar to what happened with PCs later, because basically, when we introduced the first PCs under Windows, when we wanted to install the Web there, we had to install a TCP/IP stack. There were several vendors. I remember one called Chameleon. And we had to install that. When we went to install the internet for someone in my company afterwards, we had to not only install a modem, but also install TCP/IP on their PC and then install the Web on top of that. That was quite an adventure.

Weber: But theoretically, the Web was agnostic about the network it ran over.

Groff: Yes. I even studied a way of letting the Web run on asynchronous networks so that we could fetch a document with a URL that would actually be a phone number. At the time in France, they were transitioning from the RTC telephone network [(RTC is a French abbreviation for « Réseau Téléphonique Commuté, the traditional analogue telephone system called PSTN in English)] to the ISDN telephone network [(a digital telephone line)]. And one of the key advantages of ISDN is that you could establish a phone connection instantly instead of waiting several seconds for a plain old telephone connection. So we thought, hey we could do hypertext with that. Imagine there's a server somewhere which doesn't have the internet, but it has an ISDN plug, which was commonplace. We could just dial that ISDN number instead of dialing an IP address, and the server would respond, send the document, and close the connection. It would work. That remained at the study stage.

Weber: But it could have run over OSI, or anything that had--

Groff: Yeah, certainly.

Weber: And so talk about the different roles of Tim, Robert, and yourself. And then Nicola Pellow you didn't mention.

Groff: Well, I think this is well known history. Well, Tim had the original idea and impetus. But more than the idea, he had the capacity to implement things. I think that's what sets him apart. Robert was working on hypertext systems as well. And he had also some talent in implementing. But Tim was light years ahead. So when they teamed up, it was very creative, because they exchanged ideas. And they had already specified everything. When I arrived, I helped with the implementation part a lot. I didn't help much in the design itself.

Weber: But Nicola Pellow had already, she did the line mode browser and had left right before you came, is that right?

Groff: Yeah. Nicola Pellow was a student from England who I think Tim had hired to help him write a line mode browser. I think the line mode browser was-- well, all the underlying components, the networking components, were written by Tim, as I said, because it was his background in RPC procedure calls. And I'm not sure, but I think Nicola wrote the user interface layers, so the thing that would display the text and let people type numbers to go back and forth. And when I took over that piece of software, basically we modularized it. We kept the interface separately and we plugged it into the standard C library. So it was a stepping-stone towards the architecture of Web browsers and servers, where you have to generate components and then each implementer is free to do what he wants to adapt to his own machine. But Nicola left, yeah, I mean, I arrived and she left. We just worked one month together, so I can't tell much about what she did precisely.

Weber: So then after that, it was really the three of you were the main people working on the Web at that point.

Groff: We were the only ones working full time on the Web, yeah. There were some people at CERN who were helping. For example, we mentioned Bernd Pollermann, or Dietrich Wiegandt who was a great guy. He was a UNIX administrator at CERN. He helped us set up archives and he gave us all the bandwidth we wanted, all the storage space we wanted on the machines. It was very good, very helpful.

Weber: And I guess Ben Segal supported it, but he wasn't--

Groff: Oh, Ben Segal, by that time, he was already working on the next generation of stuff. He was instrumental in inventing storage arrays and cluster computing. It was fantastic what happened at CERN. I think the computer museums should ask CERN about artifacts of that period, because there were many things that were invented at CERN and that are not now known outside the physics community.

Weber: Yeah. I've started to talk to Ben about that.

<crew talk>

Weber: And then actually, we would like you to talk a little bit about CERN.

<crew talk>

Weber: --did at CERN.

Groff: Oh, that's my quote.

Weber: Yeah. And that had to do with the feudal nature of computing at CERN, I believe. So anyway, if you could address both those.

Groff: Yeah. CERN, in 1990, was-- I don't know how to say it in English. The <French word> [?? chaudron ??], the place where you make the best food, you know? Because you pour all the ingredients and you have the fire underneath, and it will just mix properly. Because you have -- it's dedicated to researching physics, so you have of course brilliant people from physics theory and from physics experiment. And the people from the physics experiment, they need things that don't exist in order to conduct those experiments. So they have to go to industry and ask, "Can you do a device that will be 10 times better?" Or "Can you do a magnet that will be 10 times stronger?" Or "Can you do an electronic system that will process things 1,000 times faster?" So it really pushes technology to its limits. Although CERN has no technological goal in itself. But it was the locus where many things were invented or perfected. For example, Georges Charpak got the Nobel Prize in physics in '92, I believe, for an invention that he did in '73, which was an electronic detector of particle movement. And that is now used every day in brain scanners. So that's a type of really how CERN is the example of how fundamental research can lead to unexpected results in practical applications. And if you ask a theoretical physicist about the Web, he doesn't care. And that's why they didn't give us any funding. That's normal, you know, it's not their job. However, the rest of the world cares.

Weber: But CERN also had many different types of computer systems.

Groff: Yeah, that's right. The world of computing was balkanized. Each major computer company had its own standards for networking, for programming, for APIs. And at the end of the '80s, we were starting to see a little bit of interoperability between computers. It was still a very new idea. It was just like when object-oriented programming was new as well. So the concept of cross-compiling was new. I remember whenever we installed, we received a new machine from Sun or from DEC or from IBM, the first thing we did was rip out all the programming tools that were bundled with it, and we replaced it with the new series of tools, the GCC compiler, because it was cross-platform, and we could compile from one machine with a target that would be any machine. And also, if something was wrong in the software itself, we could fix it, it was open source software. It was very important. So this culture really intersected, this computer culture from the USA was influential on the way we approached software development for the Web. It had to be open source from the start. It took us a few years to get CERN management to approve and release it formally, but that was one of the big success factors of the Web. Some people from Digital Equipment and IBM had visited CERN and asked whether CERN would ask for licenses on the Web. And apparently, if CERN had asked for any licenses, then these companies would have gone ahead and developed their own software. So you would have had a balkanized Web. You would have had the IBM Web, and the DEC Web and the Microsoft Web. Microsoft tried that, by the way. They wanted -- AOL was a pre-existing system, at CompuServe as well. When Microsoft said, "We're going to build a worldwide network called MSN, it was not supposed to be universal, it was supposed to be Microsoft only, and then they, you know, thankfully changed their plans and made it fully web compatible -- web based. I shouldn't say, web-compatible, web-based is different, you know. So gradually, in the early days, one of the key things of the

web was that we had devised a system that could encompass the other systems. The URL had this protocol string at the beginning, so, HTTP means, we're going to use hypertext transfer protocol to transmit data. But if we said, FTP, we were going to use the pre-existing FTP protocol. If we said, gopher:, then we were going to use the Gopher protocol. WAIS, Prospero, Archie, Telnet, there were plenty of protocols, and we were just implementing them, adding them to the software library, and any browser could use those protocols. So we were integrating any data that was published under any network information system that could be viewed through the web, because hypertext is a richer model than all the other models, so it could represent anything. And so that was the case for I'd say, five years, and then there was a turning point where, suddenly, people were only creating information for the web, and they were not integrating information from the other systems. They were phasing out the other systems. That's when the web became totally dominant, and that's between, let's say, '94 and '97. In '97, it was over. There was nothing else, there was just the web, and e-mail.

Weber: So what do you think the most important factors were? You've been discussing some of them, but for making the web succeed where the others didn't.

Groff: Well... From a features standpoint, it's the URL, that's the most fundamental thing, being able to have a representation of a document that can be easily written on the back of an envelope, that can be spoken on the phone. That's extremely important. I remember when we designed the protocols, the HTTP, the URL, the HTML, there was a naming standard for documents, it was called, ASN.1, Abstract Syntax Notation, System Notation, I don't know. It was very cumbersome, you know? You could identify uniquely any document, but the identifier itself was two or three lines long, because it was a universal thing. And well, that wouldn't work, you know? Thankfully we had the DNS, it's already a big first step, where you could say, you could say, **weber.com**, and it would go to that site. And behind the DNS, we tacked on whatever file hierarchy you wanted. So you could organize your website by department, you could organize it like a file system, you could organize it by function. The URL made all this very, very flexible. So that's, to me, the URL is the key feature of the web, you know, not HTML. HTML is the visible thing, but it's not very important. And HTML has changed a lot, by the way.

Weber: And how about ease of installation, in the sense that someone could fairly casually, it may be tougher than Gopher, but they could fairly casually set up a web server. They didn't have to buy licenses or get training, or buy equipment.

Groff: Yeah, well that's being open source, you know? We gave the tools, you could download them and play with them as you wanted. The limiting factor in installing a web server is that you need it to have a permanent connection to the internet and only universities had that. Even corporations didn't have that, not -- few of them, and for individuals, forget it. So that's how the hosting business came to exist, because some people said, "Wait a minute, we're going to install servers on places where we have a fixed lease line to the internet, which costs a fortune, so you know, to recoup a part of our costs, we're going to let people host their sites on there. And you see how this plays out with the, not having an editor. Because if you're going to host your files on a remote server, then you don't need a local editor, you know, you need to prepare your pages in advance, and then you upload them to the host, okay, and on the host, they're being served. So it was not meant to be that way, okay? Imagine an alternate universe where technology would have made, you know, permanent internet connections, let's say, ADSL was available on the phone lines in '95, instead of being available in 2000. Then everybody would have had a server at home, everybody. And there would have been a great demand for integrated editing tools and sharing tools, because the very computer on which you would type the information would be the same computer that

would serve the information. That would, keeping with the spirit of the internet, which is a peer-to-peer network.

Weber: So describe that alternative universe a bit. What would it look like?

Groff: I'll just-- so each computer at home or in a company has its own server bundled, you know, when you install the web, it's a reader and it's a server. And there's an authentication linked to either that machine, or to user accounts on that machine, and whenever you create a document, you decide if it should be private, if it should be published to your group of peers, let's say people in your company, or your friends network, or if it should be public, visible for everybody. And this is all administered on your local machine, so you don't have to set the permissions on the web host somewhere, okay? On the web host, it's more or less all-or-nothing. So I suppose some people would have written intuitive tools to manage permissions and groups, just like today, you go to Facebook, and you manage your preferences on who can read what about your life, except these are centralized on Facebook's servers, whereas in this alternate universe, it would be fully distributed on each computer or each PDA you have, would manage that for you. So some people would love that world to exist. It's very hard to go back and do it, you know?

Weber: And what would people have done on that system? How would it have changed the way they did daily tasks?

Groff: I think, with good design of user interfaces, it would have been much, much easier to publish things and to share things. I mean, for example, something like Wikipedia would have emerged five years before it really did, but the usage would be the same, I think, you know, the people are doing on the internet all they want to do. The tools are there. It would have been structured differently. I don't think it would have fundamentally changed the pace of adoption of the net. It's also a question of generation, you know. People who were born with computers have a natural tendency to use them.

Weber: But it would have extended to things that are now personal information management functions, right? You would have done your own to-do list and things like that.

Groff: Yeah, well today there's-- it's very hard to manage what you publish or what you don't publish, and it's very confusing. So it's often an all or nothing proposition, you know? You keep your data on your private computer, or you publish it on the public platform. It's a complex issue, you know. I'm sure that, yeah, if you had had your own servers, people would have written tools that enable you to better manage your personal information, so perhaps socially the impact would have been a bit different, but it's really hard to say, you know.

Weber: But when you and Tim and Robert were talking about these sorts of things, the universe you just described is how you wanted it to unfold, then, or thought it would unfold.

Groff: Yeah, definitely, yeah.

<crew talk>

Jon Plutte, producer: I'm sorry. Yeah, I wanted to pick up on something that you said, which was just sort of the fact, the spirit of the internet, which is an intriguing thought to me, and it also--

Groff: Oh, I can explain that.

Jon: But what's interesting is, I'm picking up that there's a philosophical under -- a backbone to this whole thing. And I'd really like to talk about, what is the spirit of the internet, and what was that philosophical backbone that helped you develop this?

Groff: Okay, yes, certainly. The internet-- so before the web, okay, before all that, the internet was revolutionary in two ways. One of them is that it enabled to interconnect various kinds of computers that were not speaking the same language previously, so it's a universal way for computers to communicate. And then, on top of that, we build applications for humans to communicate. So that's one revolution. And the second revolution is that, on the internet, every machine, every host has the same rights. Every participant has equal rights. So if you get an IP address on the net, the principle was that you could talk to any other IP address on the network, unhampered by permissions or where you're sitting, or, okay, you would be limited by the speed of your connectivity, but you should not be limited by artificial barriers in the software itself. And if you ask Vint Cerf, that's what he will tell you. The internet is a peer-to-peer network. Every machine has the same rights. And, over time, as the internet was commercialized, it kept that openness technically, but it lost it socially. That is, there's what I already said, that if you publish something, you cannot host it yourself, you have to host it somewhere else, and there is also the fact that, depending on what kind of connection you have, your content may not be prioritized properly. It is the whole political issue now, not net neutrality. I believe the net will keep its neutrality, but the big factor that has already killed this equality on the technical level, is the network address translation, is it N-A-T, NAT. This little box you have at home that is the bridge between your home network and your internet connection. Okay, a modem, a router, so this router performs a procedure called NAT, which pretends to the outside world, that you only have one IP address, whereas in your inside network, the various computers in your home, your PDA, your iPhone, your printer, they all have different IP addresses, so in a truly neutral net, all your home devices could be individually speaking to any home device of someone else in the world. But with the Network Address Translation, you can't. You have to open a special gateway in there, a special port. That's what, you know, the peer-to-peer software are avoiding. They have clever ways to, yeah, to share files through the net. But the NAT was also necessary because there are not enough IP addresses, so it's a chicken and egg problem. So yeah, I'm not sure we will ever go back to a fully open internet on the technical level. I'm not even sure it's desirable, you know, because if your computer is wide open, I mean, it means also anybody can go into your computer, so in a sense, it's good that you have this little firewall at the entry of your home.

Weber: And does this have consequences socially and -- I mean, because I think John is asking also about, does this affect society at large, whether you have a more open net, or not?

Groff: Well I think Jon's question was, okay, how does the web's openness relate to the internet's openness, right?

Jon: Well it was also a bigger question, as you were developing the web, that you were aware of all this, that you've been talking about, and I think you talked a little bit about the philosophical underpinnings of the internet, but what about when you created the web, were you thinking about that?

Groff: We were definitely thinking that publishing should not be reserved to publishers, but anybody could publish, you know, what he had to say, so we provided a tool. So, in that sense, it relates to Gutenberg if you go very far in the past. So... and the web definitely made it cheaper and easier for people to broadcast their views. But before the web, people had found ways to broadcast their views on the internet, you know, without hypertext. I'm thinking of Usenet, that was a fantastic forum of discussion and freedom. So people just use the tools that they have. With the web, we tried to give a tool that would be as unconstrained as possible, and it's amazing, you know, what people have done with it. Yeah, good.

Weber: Talk a little bit about multimedia on the web, and what the plans were for it.

Groff: Well, multimedia, okay, hypertext is text which is not constrained to be linear. You have text, you can read it from beginning to end, and you can jump from one place to another place. That's hypertext. And then hypermedia was the same idea, but with anything that is not necessarily text. So, for example, in hypermedia, hypermedia video system, is what, I think, in the video industry they call non-linear editing, so you shoot a video and the base track is stored in a fixed place, but just by making pointers to various pieces of the video, that's how you do the editing. So you don't actually copy bits and pieces of the video like you had to do with film or with tape for audio, but you just edit the raw material through pointers. The raw material remains in place. Then after that, okay, you can produce a copy that will have your edited parts, and that's another story. So that's hypermedia. The web was agnostic with regards to what type of media is being transmitted over the net, all right? So, obviously we could transmit text in various encodings, so that includes hypertext, that includes HTML, but we could transmit images, we could transmit sound, we could transmit video, and we could make pointers to any document in any format. So then it became a question of which user interface tools were provided to deal with hypermedia, and... what was the question?

Weber: In the early--

Groff: Your question was, was the web designed for multimedia, I guess, and my answer is clearly yes. But again, it was a platform that enabled things, and then, if people wrote tools that used that platform, for example, the web has been used for transmitting audio files in a streaming way, but it has also been used as the way to jump into the middle of an audio file or a video file. The basic protocol allowed that a long time ago. It's just a matter, then of user interface, which is a big deal in itself.

Weber: Wasn't there also a lot of talk about format negotiation being able to adapt the level of multimedia to the type of browser you were using?

Groff: Yeah, well the idea was, when you request a document, the person serving the document, or rather the machine serving the document on behalf of a person, should know a little bit about what the client can display. Because obviously if you're looking at the document from a text-based terminal, you cannot send a diagram, and if you're looking on a color screen, you should send the color picture. If you're looking on the black and white screen, you should send a lower resolution picture. So there was a... we-- Yeah, there was a... how do you say, a standard? No, a protocol as part of the HTTP protocol, there was a way for the client to tell the server, "Oh, by the way, I'm able to display images of that resolution, and I'm able to play sound, or not." And then the server could react by sending the appropriate document. In practice it was not used very much. It was used quite a bit in the 90s, because many computers had limiting factors in their way to display graphics or process sound. Nowadays, any computer can do anything basically, so we're not looking at those restrictions any more. However, it has

been used to good effect in mobile computing, where the first mobile terminals, they would say, "Oh, by the way, I'm a mobile terminal, so then the server would send a restricted version of the page. Not restricted, a simplified version of the page.

Weber: So you mentioned that Tim Berners-Lee spoke quickly.

Groff: Very quickly.

Weber: So how did you used to get him to slow down?

Groff: No, I just had to make my brain run faster, I guess. I over-clocked my brain, and that worked.

Weber: Didn't you sometimes ask him to do presentations in another language, perhaps?

Groff: No, no, no.

Weber: You told me years ago you'd ask him to speak in French, because he spoke more slowly.

Groff: Oh yeah? Possibly.

Weber: And you were there for the very first web t-shirt?

Groff: Well, actually, Robert had found a t-shirt printing shop in Geneva, and he had designed, as you know, the world wide web logo with the green Ws, so he had a test run of t-shirts made, and I believe he printed just three, you know, for himself, and Tim and I. And this one just had w.w.w, world wide web, written and then he did another run of t-shirts that said, "Let's share what we know." And he made several dozens of those. And yeah, I had one of the original three, but unfortunately it was lost to garbage.

Weber: And in that-- well, actually, go a little bit into what you did after the summer of '91, just finishing up your time at CERN.

Groff: Oh, I was fully involved in the web, you know, it's even after my contract was finished at CERN, I kept working for the web a little bit. That was, let's say, spring '92, I-- that's the time I did the port from UNIX to -- from IP to DECnet. And then I worked also -- I started working on projects to implement the web in other institutions, for example, the CERN Library, not the software library, the library of books, wanted to have a web interface to their catalogue system, which was a huge system called Alice, I believe. And, well, for some political or budget reason it couldn't be done at CERN, but the same Alice, or Aleph system was installed in the National Library of Denmark, and so they sent me out to Denmark to build a web interface to that system. And that's how I started my career in web consulting. I believe this contract in July '92, was the first paid consulting for web technologies. Yeah, I found the contract recently, with my colleagues in Copenhagen. I have it, signed.

Weber: So the milestone, summer of '91, the code library was up.

Groff: We released-- that's where we released the code, let's say, a bit more publicly. You know, prior to that, we were showing it as demonstrations, and people were free to install their stuff and we were helping them. In August we announced it on various newsgroups, alt.hypertext, and that's where you see the archives now. It's the key date. There are several birthdays of the web, of course. There's '89, when Tim wrote his proposal, there's 1990 when we completed the next software, there is August '91 when we announced it on alt.hypertext, and then there's probably '93, when we put the software in the public domain, when CERN agreed to put the software into the public domain. So that's already four birthdays, so we get to celebrate almost every year, we celebrate a web milestone of some sort. Journalists come back and say, "Oh, it's the 15 years of the web. Hey, but last year it was the 20 years of the web. Which web are you talking about?" We did a pretty big celebration in February 2009 at CERN, for the 20 years of the web, web being the initial proposal, so it's a piece of paper that says: "this is how the web will work". And actually, yeah, one year later it worked. It's good.

Weber: So after the summer of '91, you stayed at CERN for how long, officially?

Groff: Well I was officially there, I think until April '92, and then I got, I think, I got a two month temporary contract to do something else, and they just couldn't -- They wanted to keep me, but they couldn't find a way to officially keep me. There were plenty of regulations, you know, negotiations between governments, because if you take so many people from France, because it's a military program, then you shouldn't, you know, these people shouldn't have an unfair advantage because they're already there, so basically the people who did cooperant at CERN had to be kicked out. And then I worked, as I said, I did this contract for the National Library of Denmark during the summer, and then I went to work for Steve Jobs. I went to work for NeXT Computer.

Weber: But in Geneva.

Groff: Yeah, in Geneva, essentially. I came to California quite often as well.

Weber: But you were still spending time with Tim and Robert, though.

Groff: Yeah, of course, I was helping to evangelize, but well, by that time they had hired all the people to maintain the code. I was not active in maintaining that code and expanding it any more. And then that's when Marc Andreessen arrived. He arrived in, I think in '92, he started to write Mosaic. I was already out of CERN for more than six months at that time, so we didn't really cross.

Weber: So, do you want to talk a little bit about what you did then, afterward with, just briefly, your web services firm.

Groff: Yeah, well I had-- I worked for NeXT for a while, and then NeXT stopped producing hardware, so they offered me a job here in the R&D department in California, but I didn't want to move to California, so I started my own company in Geneva, and built web systems for many, many companies there. Yeah, that's it.

Weber: And during the period you were working at CERN, can you tell a little bit of how it worked, like, day to day, between-- you were sharing an office with Tim, you were meeting with Robert daily, or--

Groff: No, Tim was at the other end of campus. I mean, Robert. Robert was at the other end of campus. And he was, he was mostly promoting the web. With this he's trying to set up a case -- interesting cases where the web could be used, getting people interested in the web, preparing conferences, etcetera. He was not very much involved with technical development, although he had written a - - he had written a prototype browser for the Macintosh, which he called Samba, but that was never very successful. It was hard, you know, programming for the Macintosh was very hard.

Weber: And within CERN, was this a well-known project?

Groff: Oh, it certainly got known, yeah, because... there were very useful resources to physicists that were available through the web, so, [I] don't remember exactly who was doing that, but I know that, at some point, the web browser, the line mode browser was installed as part of the standard on every machine at CERN. So you could go to the web, it was part of standard procedure to go to the web and look up something. So, in that sense, yes, it was very much known. But they didn't have a sense that it was a world wide web. It was more interesting to them as an internal documentation system, which happened to be connected to some other physics institute and physics databases, but no, I think the worldwide impact, that only came after Mosaic was popularized, and after the momentum came to the [United] States, clearly.

Weber: Now, and you talked a little bit about putting support for WAIS and FTP and other-- encapsulating other protocols. That was quite early, right?

Groff: Oh yeah, that was-- I remember, we were adding a new protocol to the web library every month. "Oh, this guy has done a new protocol, great, let's see how we can implement it. Blah-blah-blah. Yeah, it's done. Okay, let's try it."

Weber: That was in '91, '92?

Groff: Yeah, yeah, along those times, yeah. I remember implementing Hyper-G, you know, it was, I think, created in Austria.

Weber: And which did you perceive-- Gopher was a competitor, in a sense, right?

Groff: Oh, there was no competition. They were friends. Absolutely not. Actually, we were surprised because the simultaneousness of it, you know. It's often like that when there's something is in the air, something visible. You've got independent inventions in various places. And Gopher was basically an independent invention of the HTTP protocol. The Gopher protocol was extremely similar to HTTP. You know, it was a one-for-one mapping, it just had different names, because, you know, they did their thing. So it was very easy for us to map the Gopher protocol, to have a web browser speak the Gopher protocol so it could display documents from a Gopher server. So it was extremely easy, it took, maybe, two days of programming. But what was missing in Gopher was the hypertext dimension, which was both a curse and a blessing. It was a blessing because people could install Gopher server without worrying about hypertext, so, at one point, Gopher had three percent of the NSF backbone traffic. It was a huge success, whereas the web was only one percent. And two years later it was, you know, the web that had, you know, five percent or ten percent, I don't know. We should find one of those graphs of the NSF backbone traffic. You know, NSFNET was shut down in 1995, I think. So I think until '95, we must be able to find

statistics that were published monthly. We knew what percentage of the bandwidth is used by which protocol. The dominant things were FTP and e-mail, obviously, and then the Gopher went like this, and the web went like uh.. whoo!

Weber: But you saw these other protocols as being something you would just incorporate and support, rather than competitors.

Groff: Exactly, yeah, we would incorporate them. And we thought there would be a diversity of protocols. Frankly, I was surprised that HTTP became the dominant protocol and I was surprised that HTML became the dominant language.

Weber: Yeah, talk a little bit about whether--

Groff: There's only the URL, where I think it had to be this one. There had to be one standard, just like the DNS, it had to be the dominant name system. You could only have one, you know, you could not have a Balkanized DNS. Some people tried that. They were kicked out of the internet really quickly. So it's good we have a uniform URL system, but in terms of protocols... and today, if you look, what are we, 20 years down the road? The protocols on the net, HTTP is absolutely not dominant, you know? The telephony protocols are here, the video transmission protocols, Skype, what have you. You know, there's plenty of things happening on the net, so there was a period in time where most people were... were not seeing the difference between the internet and the web. You know, get on the net, that means have a browser, and go through hypertext. And now this has changed, because people are chatting, they're using the phone, they're doing video, so HTTP is like a universal glue that allows you to go to a site and do things, but more and more, people, they install specific applications, the explosion of peer-to-peer technology, it doesn't use HTTP at all. And that's good. The internet is resilient, and I'm sure people will invent more protocols.

Weber: So it's coming back, in a sense, to that vision.

Groff: It's true that the protocol to grab a document or make a query somewhere, and get a result, this is definitely HTTP, it's totally dominant in that space.

Weber: And coming from France, you're obviously familiar with Minitel, and...

Groff: Hello, Minitel, yes. <laughs>

Weber: So talk a bit about how the web felt with being someone who had grown up around a web-like system.

Groff: You have to understand that Minitel was the biggest, the most widely used information system by the general public. Nothing came close until, I would say, '98. And the Minitel started in '81, it was at its peak usage in terms of population penetration in, let's say in '90, so that all predates the web. It predates the commercialization of the internet and there were millions of people using the Minitel daily.

Weber: And you had grown up using it casually, right?

Groff: Yeah, naturally. I was studying... of course I was using it at home, because we had... the Minitel was given the terminal, the Minitel terminal, I think you have one in the museum. It was given to every household as a replacement for the paper phone book to save on paper, you know? Of course it cost a fortune initially, but it was a case of seeding the population for other usage, which then made a fortune for the telecom industry. So that Minitel device was on every household and everybody was using it as a replacement for the phone book. It was much faster than using a computer, because you turned it on, it was instantly connected to the phone book system, so you didn't have to wait for the computer to boot, or for the modem to dial or something. It was really, really fast.

Weber: So to what extent did you... was Minitel a kind of mental model for what you're doing with the web?

Groff: It wasn't the mental model for the technology, because the technology was completely different. The similarities in the markup language, maybe for display, yes.

Weber: No, but the use.

Groff: But the usage, yeah, the popularity... with the Minitel, I knew that something like the web could be a mass market thing, whereas the internet, in the 90s, in the early '90s, was still seen as something that would never go out of academia, or maybe some corporate use, but that's all, you know? It was not planned that it would reach the masses. I think the web contributed immensely to bringing the internet to the masses. It gave a purpose to having a modem and paying for [an] internet connection. But I knew it was possible, because I had seen the Minitel work, yeah.

Weber: And did you think the web might go in somewhat the same directions, commercial and that's...

Groff: I don't understand the question.

Weber: Did you think the web might go in some of those same directions, into a commercial system, with, you know, payment and authentication and all of that?

Groff: Ah, you mean... okay, yeah, you have to explain then, that the Minitel, besides the phone book application, which was free, when you dialed into a service, such as the train schedule, or the newspaper, you would pay by the minute. Okay, it was a telecom initiative, so you would dial a number and the number was 3615, 36-15, trente-six, quinze. <laughs> And then after that number, you would give the name of a service, so, you know, if you gave, SNCF, trente-six, quinze SNCF, that's the National railroad system, so you would go to the train schedule. And you would pay by the minute. So the time it took you to actually, you know, look at the page, and type the name of your origin and destination cities, and look at the schedules, and then you will hang up, well that's what you would pay for. It was a very simple model, and it worked well. At the beginning of the ISP phenomenon, some ISPs were charging depending on usage. You know, you could have ten cents every time you send an email, and I don't know, one dollar per megabyte of data. There were some that charged that way, and there were some others that charged a flat fee. It turned out that the ones who charged a flat fee got the most customers, and after a few years, everybody charged a flat fee. Okay, it could have been different, but I think that's what the market wanted then. That's the best, so... Does that answer your question? Is that what you meant?

Weber: I guess what I'm getting at is, did you, Tim [Berners-Lee] and Robert [Cailliau] foresee some of the commercial uses and the popular uses, and you, from a French perspective, did you also, you know, have a unique insight because of Minitel as to what could happen?

Groff: I didn't have a unique insight. You know, when you use the web, then you obviously... it's obvious that it's going to be popular to share all kinds of data and documents, and to provide all kinds of service on top of that document infrastructure. So I wouldn't say it's a special insight, it's just obvious. Then, of course, people who... there are many people who saw the web, and who had done business on the Minitel, and they thought it was a nightmare, because, on the web, you have to provide content for free, whereas on the Minitel, you know, the newspapers were making money just by providing the news feeds on the Minitel system. It was [a] significant portion of the newspaper revenue, and when they suddenly had to hire people to build websites and pay for servers to serve something, and they could have no revenue besides banner ads, they were... they freaked out, you know? That's why they held onto the Minitel for a very long time. Minitel was still drawing a lot more traffic than the internet in '99, then it really...-

Weber: In France, yeah.

Groff: It really... in France, the Minitel was knocked down, not by the internet, because I said, the Minitel terminal in many ways was more convenient than having a computer and having the web with a modem. It definitely was more convenient. I remember when Jef Raskin came visit me in France in '98, Jef Raskin being one of the inventors of the Macintosh. I showed him the Minitel face-to-face with a PC connected to the 'net. And he said, "Wow, this thing is impressive, you know, it's much faster," even though the technology was much, much lower than even an average PC in '95. But what knocked down the Minitel is the... is the broadband access. As soon as ADSL came into the picture, let's say, 2000, 2001, I got my first DSL line in France in 2001, that really prompted people to buy computers more massively, and because it had this instant connection capability there was no advantage to the Minitel anymore. And over the next two or three years, the usage dropped. And, Minitel was finally shut down, I think early this year, in 2009. The last services were shut down.

Weber: So what do you think of the Computer History Museum?

Groff: It's brilliant. I've just visited the Visible Storage exhibit. I learned a lot of things. So these machines from the 1950s or '60s where, at every generation of hardware, you had to invent new ways to store data, or to input data, or... there was a timesharing computer from 1961 where you could actually type commands and get a reply, get a result immediately on a piece of paper. That was revolutionary. I can imagine the people being used to bring their programs to the computer room and wait until the result arrived hours later, and now they could do it interactively, brilliant, yeah. It's going to be a lot more difficult, I think, for you to show... to communicate the excitement around software and around networking. I think it's quite a challenge. Good luck.

Weber: Did you ever think that some of the work you did might be in a museum?

Groff: Yeah, I'm thinking when did I think that, but we knew it was historical, certainly two or three years after the early web servers, we knew it was a historical significant development.

Weber: Was there a particular moment that that came home to you, an experience?

Groff: I don't know. No, it was just a fact of life. It was here all along, you know. Because we had to, even after I left CERN, and after I wasn't involved in evangelizing the web anyway, I ran my little business creating web applications but I still had to explain to people what the web could do. There were places where we installed the first web server and connected it to databases and it was revolutionary. People realized it was a new way of computing and of networking, yeah.

Weber: That's good, great.

Groff: Enough for today?

Weber: Well thank you very much.

Groff: See you in 20 years? <laughs>

END OF INTERVIEW