

# Chapter 2

## Why HEP Invented the Web?

**Ben Segal**

We are going to tell part of the story, little-known by most people, of how one of the most profound and revolutionary developments in information technology, the invention of the World Wide Web, occurred at CERN, the High Energy Physics laboratory in Geneva.

In fact one man, Tim Berners-Lee, invented the Web, not “HEP”. So our question should really be re-phrased as: “What was the influence of HEP in leading to the Web’s invention?” In the discussion that follows, we will make use of some personal recollections, partly my own but also those of Sir Tim Berners-Lee himself (“TB-L” in what follows) in his book “Weaving the Web” [1] (abbreviated below as “WtW”).

The main development and prototyping of the Web was made in 1989–91 at CERN, which by that time had become the world’s leading HEP laboratory. But the ideas for what became the Web had germinated and matured in TB-L’s mind for many years before those amazing three months in the winter of 1990–91 when the code was written and tested for all the elements composing the Web, namely a browser, a server, a naming scheme (URI, now URL), a hypertext markup scheme (HTML) and a transport protocol (HTTP).

### 2.1 Introduction

The Web’s invention, like many other such leaps, was in fact “coincidental”. It was certainly not ordered, planned or anticipated in any way by “HEP”, by CERN, or by TB-L’s programme of work there. But for that particular leap to occur, its author’s inspiration alone would not have been enough without a certain number of

---

B. Segal (✉)  
CERN, Geneva, Switzerland  
e-mail: [B.Segal@cern.ch](mailto:B.Segal@cern.ch)

supporting pre-conditions, and all these existed at that time at CERN. What were these essential items?

1. Managerial tolerance and vision.
2. A pioneering approach to certain technologies important for HEP.
3. A tradition of pragmatism in working style.
4. A movement from “proprietary” to “open” standards in computing and networking.
5. A link to the new “open-source” movement.
6. Presence of the Internet at CERN.

We will look at each of these important areas in more detail, and try to decide the importance of their influence as we go along. Let us also remember that we are discussing events which took place 20–30 years ago, the equivalent of many “generations” of technology development.

## 2.2 Item 1: Managerial Tolerance and Vision

CERN is basically a research organisation, but for physics and engineering, not for computer science. The *working style* that TB-L found in the CERN community was decidedly positive. Quoting TB-L directly: “I was very lucky, in working at CERN, to be in an environment . . . of mutual respect, . . . building something very great through collective effort that was well beyond the means of any one person – without a huge bureaucratic regime”. (WtW, p.227).

*Managerial tolerance* and support was strongly shown from the start by one CERN group leader, Mike Sendall, but much less by any higher CERN levels. Sendall was actually TB-L’s first and last manager during his CERN career. In fact, the “spare time” that TB-L used for his Web project was below the threshold of bigger bosses and Sendall’s role was crucial: without his quiet encouragement it is fair to say that the Web would probably not have been born at CERN. Later on, when more senior CERN management became involved in discussions with the Massachusetts Institute of Technology, the European Union, and other parties concerning the future of the Web and the W3C, Sendall was again a key player. But even he did not go so far as to insist that the Web’s long-term future lay at CERN.

*Vision* was absent managerially, as practically nobody recognised the actual potential of the Web project. The essential source of vision was TB-L’s own, and he held to it with consistency and determination. Thus practically no “official” CERN resources could be found to support the work, either at the start or later. Two students (Nicola Pellow and Jean-François Groff) were “poached” early on from unrelated projects, and both worked for some months with TB-L, but without serious managerial commitment. The only “planned” application of the early Web work at CERN was as a gateway allowing access from VMS, Apollo and other machines to the very popular “XFIND” and “XWHO” phone and staff look-up facilities which had been developed and were running on the IBM VM/CMS service. The success

of this gateway (1991) helped to publicise the Web's existence; even demonstrations using the primitive line mode browser convinced some people who used it that the Web perhaps had a future.

Nevertheless, the “managerial space” created by Mike Sendall allowed Tim to continue, and chance also aided Sendall's support for the project, by finding TB-L his principal collaborator, Robert Cailliau. Sendall and Cailliau had moved together from the Computing and Networks division (CN) to a new division Electronics and Computing for Physics (ECP) in 1990 and when Sendall found that Cailliau had done some independent work on hypertext, he introduced him to TB-L and encouraged their collaboration, even though collaboration across divisions was most unusual and therefore again informal.

Managerial tolerance and vision create the space for innovation.

Cailliau proved an excellent complement to TB-L, not only in the technical realm. He organised International Web Conferences from 1994 and helped to handle developing relations with the European Union. Later on, when Paolo Palazzi became Robert Cailliau's manager in ECP Division, he also recognised the potential of the Web, allocated some students to work on it, and housed the whole team including TB-L and Robert. But TB-L's own computing division CN never went so far, except for allowing TB-L to spend some time in the U.S. from 1991 to spread the word about the Web and meet collaborators there.

Outside CERN, another HEP site (SLAC) very rapidly picked up the WWW scent and started some work on it, again in an informal style based on the interest of a few individuals, Tony Johnson, Paul Kunz, Bebo White and Louise Addis. The SLAC work offered access via the Web to physics reprints hosted on a large database called SPIRES, thus attracting many users. Other HEP sites followed slowly, but with no organised momentum.

We conclude that there was just enough managerial tolerance and vision to allow the work to get started and make some progress, but that the future of the WWW in HEP was never assured.

### **2.3 Item 2: Pioneering Approach to Certain Technologies**

The HEP community, run by physicists and engineers, is not afraid of taking a leading role to develop technologies that it considers essential. Examples include vacuum and low-temperature techniques, magnet and accelerating cavity design, tunneling and other civil engineering methods, as well as some areas of computer technology including high speed data links, large storage systems and computer clusters. Nevertheless, any decisions to develop CERN's own solutions were always

dependent on the prevailing staff and financial situation. Many periods were characterised by a “Buy, don’t develop” attitude, or – in more straitened times – even by attempts to outsource computer operations or software development.

Pioneering approach allows direct access to innovative technologies before these are available off-the-shelf.

Fortunately, at the start of the 1990s, the atmosphere was reasonably open in the computer field. In fact two major revolutions were ongoing: the transition to Internet protocols both within and outside CERN, and the move from mainframes to distributed computing using networked clusters of Unix workstations – the “SHIFT” project [2]. This loosening of earlier more centralised attitudes to physics computing had some beneficial effects for the Web project, opening both minds and working practices toward a truly worldwide approach.

Nevertheless, in the end, the laboratory’s concern to prevent any distraction from its mission to obtain approval for the LHC project made it impossible to keep WWW or even part of the W3C at CERN, even though financial support for this had been offered by the European Union in 1994.

## 2.4 Item 3: A Tradition of Pragmatism in Working Style

CERN and the HEP community often favour “engineering” solutions in software development, and are sometimes criticised for this. But this tradition helped TB-L, who made several very important pragmatic choices while developing the Web.

First was his extremely powerful but simple scheme for what he called “universal naming”, leading to the notation for today’s URL. The entire subject of name spaces at that time was arcane and acrimonious, blocking progress towards any practical consensus. Even the Internet Engineering Task Force (IETF), the Internet’s normally rational and productive standards body, took a negative and parochial view of TB-L’s proposal for a standard “Universal Document Identifier”, objecting to its name “universal” and finally accepting only the watered-down “Uniform Resource Locator” or URL. But this was sufficient for the Web’s purposes.

Another good example was TB-L’s acceptance of broken hypertext links, which radically aided the practical implementation and growth of the Web. This put him into considerable conflict with the main community of computer scientists then working with hypertext systems. He put it concisely in his book (WtW, p. 30): “Letting go of that need for consistency was a crucial design step that would allow the Web to scale. But it simply wasn’t the way things were done”. This simplification was analogous to the liberating role played by IP in the success of the Internet (using a connection-less, best-effort lower network layer, in contrast to rival connection-oriented systems such as X.25 which accumulated too much state to scale up

easily). In fact a similar sort of “religious war” was fought between the hypertext traditionalists and the emerging WWW community as that already ongoing between the X.25 (and OSI) community and the Internet developers. Internet and WWW proponents both won their respective battles against these powerful opponents.

Pragmatic “engineering” approach as opposed to “theoretical Computer Science” approach to deliver results quickly and allow evolution of the concept.

Another example of TB-L’s pragmatism was his design of HTML, based on the pre-existing standard markup language SGML but drastically simplified. “When I designed HTML for the Web, I chose to avoid giving it more power than it absolutely needed – a ‘principle of least power’, which I have stuck to ever since”. (WtW, p. 197).

His choice of programming language for the Web prototype was also significant, namely Objective C, considered by many at that time as too crude and not sufficiently structured. The choice of the NexT platform and its NexTStep environment was also extremely unorthodox but very astute, reducing the development time to a minimum (three months!), but ruffling feathers at CERN. In fact, purchasing the two NexT machines for TB-L’s development work was another administrative feat accomplished by Mike Sendall.

## 2.5 Item 4: Proprietary Versus Open Standards

The 1980s saw enormous changes in attitudes and practice as the transition occurred from closed proprietary software and hardware systems to open and standardised systems. Typical were: the replacement of proprietary computer hardware architectures (IBM, DEC, etc.) by a small number of microprocessor-based systems; the replacement of diverse and incompatible operating systems by Unix; the increasing availability of programming languages like C which allowed portable application and system development; and the replacement of proprietary and incompatible networking media and protocols by Ethernet LAN’s and Internet protocols.

Open standards have been one of the “enabling” elements of the Information Technology revolution we have witnessed in the past twenty years.

TB-L learned and benefited from many of these advances when he took the major role in implementing a system for Remote Procedure Call (RPC) linking

many computer systems at CERN, in collaboration with myself and others in the mid 1980s. RPC was a new paradigm which allowed computing procedures required by one system to be executed on other(s) to increase power and flexibility. The RPC project's implementation, largely carried out by TB-L himself, was an extraordinary accomplishment. Developed between 1985 and 1989, it allowed graceful cooperation between many extremely diverse computer systems of all sizes involved in the physics experiments and control systems associated with CERN's new LEP accelerator. Messages passed over a wide variety of physical networks and bus systems, and the system supported a very wide range of operating systems, from microprocessors and minicomputers to the largest Computer Centre mainframes. This was Berners-Lee's practical introduction to the world of distributed computing with its plethora of competing standards, but with a set of tools allowing its harmonisation and mastery. His ability both as a visionary and as a powerful implementor was fully utilised and became apparent to his close colleagues as a result.

## 2.6 Item 5: Link to the New Open-Source Movement

The Web idea was based on a vision of *sharing* among groups of people all over the world – sharing ideas, blueprints, discussions, and of course code.

Tim had a particular knowledge of and interest in what is today known as the Open Source movement, but which in the 1980s was just getting started. His awareness of the open source community had led him to the decision to find outside programming resources for WWW, once it had become clear to him that CERN would not provide sufficient support for a serious effort to develop a browser that would run on Windows PC's and mainline Unix desktop systems. He placed some prototype code on the Usenet group alt.hypertext in August 1991 – probably the first time that code written at CERN was published in this way. Very soon this stimulated many browser development efforts including Viola, Erwise, Arena, Lynx ... and of course Mosaic (later to evolve into Netscape, Mozilla and Firefox). The level of interest in the Web increased worldwide and exponentially from that moment. In TB-L's own words: "From then on, interested people on the Internet provided the feedback, stimulation, ideas, source-code contributions and moral support that would have been hard to find locally. The people of the Internet built the Web, in true grass-roots fashion". (WtW, p. 51).

The Open Source community built the Web as we know it, and in return they got a very effective tool to foster the development of their community.

TB-L was one of very few people at CERN who knew about such things as Gnu Public Licences (GPL). He realised that trying to licence the WWW commercially would be a kiss of death (as it proved to be for the rival Gopher system which rapidly disappeared when licence fees were imposed for it). From 1992 onward he tried to get CERN to grant GPL status for the WWW code, but later realised that even this would be too restrictive and finally in April 1993 obtained CERN's permission for release of the code with no conditions attached.

## **2.7 Item 6: Presence of the Internet at CERN**

This section is somewhat more detailed than those above, as the establishment of Internet protocols at CERN is an area in which the present author played a major role. I have thus permitted myself a more personal approach (including the use of the first person in some parts). The text is based on an article I first wrote in 1995 [3].

### ***2.7.1 In the Beginning: The 1970s***

In the beginning was – chaos. In the same way that the theory of High Energy Physics interactions was itself in a chaotic state up until the early 1970s, so was the so-called area of “Data Communications” at CERN. The variety of different techniques, media and protocols used was staggering; open warfare existed between many manufacturers' proprietary systems, various home-made systems (including CERN's own “FOCUS” and “CERNET”), and the then rudimentary efforts at defining open or international standards. There were no general purpose Local Area Networks (LANs): each application used its own protocols and hardware. The only really widespread CERN network at that time was “INDEX”: a serial twisted pair system with a central Gandalf circuit switch, connecting some hundreds of “dumb” terminals via RS232 to a selection of accessible computer ports for interactive log-in.

### ***2.7.2 The Stage is Set: Early 1980s***

To my knowledge, the first time any “Internet Protocol” was used at CERN was during the second phase of the STELLA Satellite Communication Project, from 1981–83, when a satellite channel was used to link remote segments of two early local area networks (namely “CERNET”, running between CERN and Pisa, and a Cambridge Ring network running between CERN and Rutherford Laboratory). This was certainly inspired by the ARPA IP model, known to the Italian members of the STELLA collaboration (CNUCE, Pisa) who had ARPA connections; nevertheless

the STELLA Internet protocol was independently implemented and a STELLA-specific higher-level protocol was deployed on top of it, not TCP. As the senior technical member of the CERN STELLA team, this development opened my eyes to the meaning and potential of an Internet network protocol.

CERN's leading role in European Networking slows down its uptake of Internet in the first half of the '80s.

In 1983, for the first time at CERN, a Data Communications (DC) Group was set up in the computing division (then “Data-handling Division” or “DD”). Before that time, work on computer networking in DD had been carried out in several groups: I myself belonged to the Software (SW) Group, which had assigned me and several others to participate in DD’s networking projects since 1970. All my work on STELLA had been sponsored in this way, for example. The new DC Group had received a mandate to unify networking practices across the whole of CERN, but after a short time it became clear that this was not going to be done comprehensively. DC Group decided to leave major parts of the field to others while it concentrated on building a CERN-wide backbone network infrastructure. Furthermore, following the political currents of the time, they laid a very formal stress on ISO standard networking, the only major exception being their support for DECnet. PC networking was ignored almost entirely; IBM mainframe networking (except for BITNET/EARN), as well as the developing fields of Unix and workstation-based networking, all remained in SW Group. So did the pioneering work on electronic mail and news, which made CERN a European leader in this field. In fact, from the early 1980s until about 1990, CERN acted as the Swiss backbone for Usenet news and gatewayed all Swiss e-mail between the EUnet uucp network, BITNET, DECnet and the Internet. As these were precisely the areas in which the Internet protocols were to emerge, this led to a situation in which CERN’s support for them would be marginal or ambiguous for several years to come, as the powerful DC Group neglected or opposed their progress.

It was from around 1984 that the wind began to change.

### ***2.7.3 TCP/IP Introduced at CERN***

In August 1984, I wrote a proposal to the SW Group Leader, Les Robertson, for the establishment of a pilot project to install and evaluate TCP/IP protocols on some key non-Unix machines at CERN including the central IBM-VM/CMS mainframe and a VAX VMS system. The TCP/IP protocols had actually entered CERN a few years earlier, inside a Berkeley Unix system, but not too many people were aware of that event. We were now to decide if TCP/IP could indeed solve the problems

of heterogeneous connectivity between the newer open systems and the established proprietary ones. We also proposed to evaluate Xerox's XNS protocols as a possible alternative. The proposal was approved and the work led to acceptance of TCP/IP as the most promising solution, together with the use of "sockets" (pioneered by the BSD 4.x Unix system) as the recommended API.

In early 1985 I was appointed as the "TCP/IP Coordinator" for CERN, as part of a formal agreement between SW Group (under Les Robertson) and DC Group (under its new leader, Brian Carpenter). Incorporating the latter's policy line, this document specifically restricted the scope of Internet protocols for use only within the CERN site. Under no circumstances were any external connections to be made using TCP/IP: here the ISO/DECnet monopoly still ruled supreme, and would do so until 1989.

Internet is finally introduced at CERN in the second half of the 80's.

Between 1985 and 1988, the coordinated introduction of TCP/IP within CERN made excellent progress, in spite of the small number of individuals involved. This was because the technologies concerned were basically simple and became steadily easier to buy and install. A major step was taken in November 1985 when the credibility of the Internet protocols as implemented within CERN was sufficient to convince the management of the LEP/SPS controls group that the LEP control system, crucial for the operation of CERN's 27 km accelerator LEP then under construction, should use TCP/IP. This decision, combined with a later decision to use Unix-based systems, turned out to be essential for the success of LEP. The TCP/IP activity in LEP/SPS included a close collaboration with IBM's Yorktown Laboratory to support IP protocols on the IBM token ring network that had been chosen for the LEP control system.

Other main areas of progress were: a steady improvement of the TCP/IP installations on IBM-VM/CMS, from the first University of Wisconsin version (WISCNET) to a later fully-supported IBM version; the rapid spread of TCP/IP on DEC VAX VMS systems, using third-party software in the absence of any DEC product; and the first support of IBM PC networking, starting with MIT's free TCP/IP software and migrating to its commercial descendant from FTP Software. All this was accompanied by a rapid change from RS232 based terminal connections to the use of terminal servers and virtual Ethernet ports using TCP/IP or DEC-based protocols. This permitted either dumb terminals or workstation windows to be used for remote log-in sessions, and hence to the use of X-Windows. In particular, starting from 3270 emulator software received from the University of Wisconsin and developed by myself and others for Apollo and Unix systems, a full-screen remote log-in facility was provided to the VM/CMS service; this software was then further developed and became a standard way for CERN users to access VM/CMS systems world-wide.

In 1988 CERN installs its first “supercomputer” running Unix and of course the TCP/IP stack.

Nevertheless, as late as September 1987, DD’s Division Leader would still write officially to a perplexed user, with a copy to the then Director of Research: “The TCP-IP networking is not a supported service.” This again illustrates the ambiguity of the managerial situation, as these words were written at essentially the same time as another major step forward was made in the use of Unix and TCP/IP at CERN: the choice to use them for the new Cray XMP machine instead of Cray’s well-established proprietary operating system COS and its associated Cray networking protocols. Suddenly, instead of asking “What use is Unix on a mainframe?” some people began to ask: “Why not use Unix on everything?”. It is hard to realise today how provocative such a question appeared at that time.

The Cray represented CERN’s first “supercomputer” according to U.S. military and commercial standards and a serious security system was erected around it. As part of this system, in 1987 I purchased the first two Cisco IP routers in Switzerland (and probably in Europe), to act as IP filters between CERN’s public Ethernet and a new secure IP segment for the Cray. I had met the founder of “cisco systems”, Len Bosack, at a Usenix exhibition in the U.S. in June 1987 and been very impressed with his router and this filtering feature. Cisco was a tiny company with about 20 employees at that time, and doing business with them was very informal. It was hard to foresee the extent to which they would come to dominate the router market, and the growth that the market would undergo.

### ***2.7.4 Birth of the European Internet***

In November 1987 I received a visit from Daniel Karrenberg, the system manager of “mcvax”, a celebrated machine at the Amsterdam Mathematics Centre that acted as the gateway for all transatlantic traffic between the U.S. and European sides of the world-wide “Usenet”, the Unix users’ network that carried most of the email and news of that time using a primitive protocol called “uucp”. Daniel had hit on the idea of converting the European side (“EUnet”) into an IP network, just as major parts of the U.S. side of Usenet were doing at that time. The news and mail would be redirected to run over TCP/IP (using the SMTP protocol), unnoticed by the users, but all the other Internet utilities “telnet”, “ftp”, etc. would become available as well, once Internet connectivity was established. Even better, Daniel had personal contacts with the right people at the Internet Network Information Center (NIC) who would grant him Internet connect status when he needed it. All he was missing was a device to allow him to run IP over some of the EUnet lines that were using X.25 – did this exist? I reached for my Cisco catalogue and showed him the model number

he needed. Within a few months the key EUnet sites in Europe were equipped with Cisco routers, with the PTT's, regulators and other potential inhibitors none the wiser. The European IP network was born without ceremony.

### **2.7.5 CERN Joins the Internet**

In 1988, the DC Group in DD Division (later renamed CS Group in CN Division) finally agreed to take on the support of TCP/IP, and what had been a shoestring operation, run out of SW Group with a few friendly contacts here and there, became a properly staffed and organised activity. John Gamble became the new TCP/IP Coordinator; he had just returned from extended leave at the University of Geneva where he had helped to set up one of the very first campus-wide TCP/IP networks in Europe. A year later, CERN opened its first external connections to the Internet after a “big bang” in January 1989 to change all IP addresses to official ones. (Until then, CERN had used an illegal Class A address, Network 100, chosen by myself).

In 1990 CERN becomes the largest Internet site in Europe and the endpoint of the U.S.-Europe Internet link.

CERN's external Internet bandwidth flourished, with a growing system of links and routers. Concurrently with the growth of the new European IP network (later to be incorporated as “RIPE” within the previously ISO-dominated organisation “RARE”), many other players in Europe and elsewhere were changing their attitudes. Prominent among these was IBM, who not only began to offer a good quality mainframe TCP/IP LAN connection product of their own but also began to encourage migration of their proprietary BITNET/EARN network towards IP instead of the much more restricted RSCS-based service. They even began a subsidy programme called EASINET to pay line charges for Internet connection of their European Supercomputer sites of which CERN was one. In this way, the principal link (1.5 Megabit/sec) between Europe and the U.S. was located at CERN and funded by IBM for several years during the important formative period of the Internet.

By 1990 CERN had become the largest Internet site in Europe and this fact, as mentioned above, positively influenced the acceptance and spread of Internet techniques both in Europe and elsewhere. The timing was perfect for providing TB-L with a platform from which to launch the World Wide Web.

### 2.7.6 *The Internet and Commerce*

Finally, in 1991, the U.S. Congress passed legislation permitting commercial use of the Internet. This had been forbidden previously, using the argument that the Internet infrastructure had been paid for by the U.S. taxpayer and hence should not allow “for-profit” use on behalf of industry or other groups worldwide. Clearly, without such a change of attitude, the Web’s level of success would have been totally different, with its use restricted to research and non-profit organisations in the same way that the Internet had been before this date. Again, a coincidence in time and awareness had made it possible for the Worldwide Web to become truly worldwide and truly revolutionary.

## 2.8 Conclusions and Lessons Learned

If any simple conclusion can be drawn from the above, it is that major research advances such as the emergence of the Web depend more on chance meetings, coincidence and unexpected insights than on planned programmes or top-down analysis, even though these latter items are of course also needed.

Consequently, the best possible thing that a management structure can do if it aims to foster major inventiveness is to adopt a *hands-off attitude* to its creative personnel, together with methods for detecting and nurturing the existence of truly creative people within its ranks, and providing protection for their activities. I sometimes refer to this process as “making space” for research talent. Apart from our own example of CERN in the case of the Web (where *just enough* space was created!), I could cite as past examples certain IBM and AT&T research laboratories which in their heydays fostered research of Nobel Prize quality in several instances. At the present time, the best example is Google Inc. which encourages its employees to spend a certain percentage of their paid time in the pursuit of their own innovative ideas.

## 2.9 Outlook for the Future

To predict the future of such a dynamic phenomenon as the World Wide Web is a perilous task. Its creation, as we have stressed above, was dependent on chance and coincidence to a major extent. The Web is in fact a perfect example of what N. N. Taleb would identify as a “Black Swan” [4]. As Taleb puts it: “History does not crawl, it jumps”. The later development of the Web’s own technology has perhaps been less unpredictable, guided as it has been by the World Wide Web Consortium led by TB-L himself, but the larger lines of media and social development which have been triggered by the Web’s presence have been very difficult to foresee.

These include Facebook, YouTube, Cloud Computing, and of course the whole panoply of search engines and their evolving algorithms. In conjunction with the last 20 years of prodigious expansion in computing power, storage capacity and network bandwidth, some radically new problem-solving techniques have emerged, which could not have been envisaged previously and which therefore look today like Black Swans. A striking example is “Google Translate”. Many decades of intense research and investment in the field of automatic natural language translation using analytic linguistic techniques have yielded far less success than expected. But recently Google has applied a totally unexpected approach to this problem: by using only massive text-matching and database power, they have achieved remarkable success over a wide range of target languages. So we must conclude that in this age of very rapid technology development, predictions of the future are even less likely to succeed than they have in the past.

## References

1. Berners-Lee, T., Fischetti, M.: Weaving the Web : The Original Design and Ultimate Destiny of the World Wide Web by its Inventor. Harper San Francisco, San Francisco (1999)
2. Baud, J.P., et al: SHIFT, the Scalable Heterogeneous Integrated Facility for HEP Computing. In: Proceedings of International Conference on Computing in High Energy Physics, Tsukuba, Japan. Universal Academy Press, Tokyo, March 1991
3. Segal, B: A Short History of Internet Protocols at CERN CERN computer newsletter No.2001-001, section “Internet Services”, April 2001
4. Taleb, N.N.: The Black Swan : The Impact of the Highly Improbable. Random House Publishing Group, NY (2007)



<http://www.springer.com/978-3-642-23156-8>

From the Web to the Grid and Beyond  
Computing Paradigms Driven by High-Energy Physics  
Brun, R.; Carminati, F.; Galli-Carminati, G. (Eds.)  
2012, XX, 360 p., Hardcover  
ISBN: 978-3-642-23156-8