deployment was in September 2003 with 14 sites, followed by rapid growth in the number of sites and the installed resources. The differing goals of the R&D grid projects and LCG with its emphasis on service stability would be the cause of considerable conflict in the ensuing years, particularly concerning the functionality and timescales of the middleware. With the deadline of LHC start-up approaching, a compromise was worked out and the Worldwide LHC Computing Grid was in round-the-clock operation at 135 sites when data began to flow in 2008 [21].

A major concern in 2000 was the network bandwidth that would be available (and affordable) between the centres and much effort was invested in working with national and international bodies that provided network capacity for science to ensure that their solutions took account of the relatively high requirements of LHC. This work paid off and by the time that LHC began operation CERN was well connected to a large, mainly fibre, international infrastructure [Highlight 9.4]. Indeed, today, fifteen years after the first proposal for the distributed model, the available bandwidth enables a much more flexible model for data distribution: instead of scheduling work to where the data is located it is now possible in many cases to move data on demand to where there are free computational resources.

9.3 Local Area Networks: Organizing Interconnection

Ben Segal

The ever closer and increasingly important interconnection of particle experimentation and accelerators with Information Technology (IT) is the main theme of this chapter. In this highlight we illustrate one essential facet: the evolution of network concepts.

Innovating in the Early Years — 1960s and 1970s

CERN rarely develops computer networking technology *per se*, but is often an early adopter of the latest technology, operating it at its technical limits. This was different, however, during the 1960s and 1970s, when CERN had no choice but to innovate due to the limited offer of suitable commercial products. During these years several networking systems were developed in-house, presented here in chronological order, and reflecting the rapid evolution of IT and its growing importance for high energy physics. A comprehensive survey is given in [22].

In the early 1960s CERN's first computer, the Ferranti Mercury, was connected by a one-kilometre *data link* to electronic experiments in the PS experimental hall, making it the first computer at CERN to analyse on-line experimental data in "real-time". The complete link interfacing was designed and built in-house.

CERN's first distributed computer system (1968–1978) was FOCUS, which linked several experiments to the Control Data Corporation mainframes in the Computer Centre [23]. The link used 2 Mbit/s twisted pair serial lines via a smaller CDC machine serving as a concentrator and front-end. This allowed physicists at their remote experiment sites to manage their files, to execute batch jobs on the mainframes, and to retrieve their results.

The next step was the OMNET network, CERN's first packet switched network (1971 to the early 80s) [24]. It eventually interconnected around 30 user minicomputers as well as the central IBM and CDC mainframes, via dedicated switching nodes. A small purpose-built CERN operating system, SMO, in a high-level language, PL-11, had to be developed for these nodes. The network protocol software, and data link hardware and interfaces were developed and built at CERN.

From 1971, twelve Remote Input Output Stations (RIOS) allowed remote batch job submission to the central CDC mainframes from remote CERN sites. Users no longer needed to transport and physically input cards and output printouts to and from the Computer Centre — a great gain in efficiency. Each minicomputer-based RIOS emulated a standard CDC remote batch terminal, but was much cheaper and more performant. Through one particularly important innovation, the human operator, who had been needed to initiate and supervise a computing job, was emulated (and replaced) by RIOS software. This strategy was successful and further expanded [25].

A technological prescient development was CERNET, a packet switched network, designed from 1975 and deployed from 1977 to 1988 [26]. CERNET offered a "Line Protocol" between switched interfaces, capable of up to 8.5 Mbit/s over short parallel connections, or serial line speeds of around 2.5 Mbit/s over twisted-pair cabling of up to 2 km (with repeaters). An "End-to-End Protocol" connected host to host applications. Once fully deployed, CERNET interconnected the central IBM and Control Data mainframes and over 100 minicomputers. The main CERNET application was the "File Access Protocol", but a "Virtual Terminal Protocol" (remote logon) accessed the main IBM system.

Not to be outdone by the experiments the accelerator builders used IT technology in innovative ways too. For example, *Accelerator Control Networks* were designed and built to control the SPS. This was based on a star topology, with a message handling computer at the centre of the star [Highlight 5.2].

Profiting from the IT Revolution: 1980s and 1990s

Starting in the 1980s, the computing and networking world was shaken by several revolutions. CERN seized this opportunity, following these developments and collaborating closely with industry in a bid to anticipate the optimal choices for hardware and software, as outlined below.

Choice of TCP/IP and Internet

Throughout the 1970s and 1980s there was a bitter contest between "US" (TCP/IP Internet) protocols and "European" (ISO-OSI) network protocols, as well as between diverse commercial network standards, Despite being in Europe and despite much pressure, CERN decided to test TCP/IP. Starting in 1984, the installation of TCP/IP on most of CERN's computing platforms was as innovative as it was audacious. By the end of the 1980s TCP/IP and the Internet had clearly become dominant and CERN could benefit fully from its courageous choice. CERN joined the global Internet in 1989 [27]. The later Web and Grid developments at CERN all depended vitally on this work [Highlights 9.2 and 9.7].

Choice of Ethernet

The first widely used Local Area Network (LAN) was the 12 Mbit/s token ring network supporting the Apollo Domain and its powerful workstations, starting in 1982. It used thin coaxial TV cable, requiring dedicated installations. Ethernet (10 Mbit/s) arrived in 1983, using different types of coaxial cables, followed by a token passing ring at 4 Mbit/s for LEP controls in 1986 [28] and much later by FDDI at 100 Mbit/s on optical fibre cables. Fortunately, Ethernet was soon recognized as the future dominant standard.

CERN's central Ethernet services

A centrally-managed Ethernet service was first set up in the mid-1980s on coaxial cables at 10 Mbit/s. Early difficulties due to the multiple network protocols present at that time had to be overcome, but the rapid evolution of Ethernet technology allowed CERN to homogenize its network infrastructure. The media (and prices!) improved dramatically and network speeds passed from 10 Mbit/s to 100 Mbit/s, and finally to Gigabit Ethernet. Today, Ethernet networking is simply a matter of choosing and installing commercial products. LHC controls are based on Gigabit Ethernet, running mostly via optical fibre [29], and distributing machine synchronization signals with sub-nanosecond accuracy [30].

Choice of Unix/Linux

Unix first appeared at CERN in 1981 but was hardly noticed in spite of its important role in connecting CERN, and later all of Switzerland, to the worldwide email and news network USENET. It was widely mistrusted compared to proprietary operating systems. Unix really came out of the closet when chosen for the Cray XMP installation in 1987. CERN and Cray worked together to develop the new UNICOS system, the first to support a full-fledged mainframe, and later to be the model for the SHIFT "distributed mainframe". Unix was chosen as the operating system for LEP controls in 1988.

While Unix had proved itself at CERN by the late 1980s, it was still for the most part a commercially supported product. The open source version, Linux, supported by the online programmer community, was mistrusted, raising fears that "we will have to fix its bugs ourselves". When finally accepted, during the final migration of SHIFT to PC/Linux platforms in the late 1990s, it rapidly became dominant.

Choice of the "SHIFT" architecture

It was a game- and life-changing choice to use commodity CPU, storage and network hardware to build a "distributed mainframe" in a scalable, flexible and economic way, starting in 1990. CERN's "Scalable Heterogeneous Integrated FaciliTy", SHIFT, was the precursor of today's cluster and cloud computing systems which completely replaced mainframes.

The networking issues for SHIFT

SHIFT's CPU, Disk and Tape server nodes had to be linked by a "Network Backplane" with, for the time, extreme throughput requirements (10 Mbit/s for a small system). Early on, it also had to deal with the high CPU power consumed by network I/O at the nodes. This was solved by the UltraNet [31] system. Later the backplane became a hybrid system containing elements of Ethernet (10 Mbit/s), FDDI (100 Mbit/s), UltraNet (600 Mbit/s) and HiPPI (800 Mbit/s) [32], until switched Gigabit Ethernet became sufficient for the task. CERN contributed to the further development, which led to the faster 6400 Mbit/s Gigabyte System Network (GSN) [33], a precursor of InfiniBand [34] now used for the data acquisition of the LHC experiments.

9.4 High-Speed Worldwide Networking: Accelerating Protocols

François Fluckiger

Public research networking created the Internet. Rarely in history did the academic community have such a dramatic and profound impact on the economy, and on society. As part of this community, the role of CERN was essential. CERN is well-known for its invention of the Web, but less so for its contribution to Internet infrastructure, central to the development of the worldwide Internet, and for its leading role in the development of high bit-rate networking.

Paving the way to the Internet and Web

In the early 1980s, CERN's external connections amounted to two minuscule leased lines operated at 9.6 kilobits per second (kbit/s), one to the Rutherford Appleton Laboratory (RAL), the other to CEA, Saclay. Ten years later, CERN had become the centre of a large star-shaped network, by far the largest internet hub in Europe. Today the LHC Optical Private Network (LHCOPN), a network connecting CERN to the LHC T1 sites operates multiple 100 Gigabit per second (Gbit/s) links. How did CERN get there?

Facing technological challenges has frequently been at the root of developing leading-edge technologies for CERN engineers. In turn, mastering a technology provides the confidence necessary to undertake very ambitious projects. Such was the case for networking, where CERN developed as early as the 1970s outstanding expertise on network protocols.

Facing the need for an open (i.e. based on non-proprietary technology), high bit-rate communication system to interconnect on-site a variety of computer brands — mainly in order to transfer files — CERN launched CERNET [26], the most ambitious networking project of its time. Communication between computers requires the use of common rules. These rules are called communication protocols. They are building blocks stacked one on top of the other in a layered structure, and are in practice implemented as modules of hardware, software or a mix of the two. They broadly divide into two types. The first type ensures the transport of raw information data between two network points, independent of its significance or use. The second type of protocol offers the end user a tangible service, e.g. email. CERNET implemented a genuine and complete suite of layered protocols, ranging from the physical and electrical plug to the CERN File Transfer Protocol running on half a dozen different brands of computers. CERNET was put into operation in 1977. This is a unique example of a complete multi-brand networking stack developed by a single organization. It was also the fastest packet-switching network of its time [Highlight 9.3].

Probably more important for the future was the type of technology adopted by CERN for one of the layers, called the packet layer. All modern networks use the packet principle where the data stream is chopped into individual pieces called packets. Packet protocols are divided into two fundamentally opposite classes: the connection-oriented class where the two end-computers must establish a connection before an information packet can be sent — the equivalent of a telephone call, and the connectionless class, where packets are independent of each other, as letters posted in a letter-box, and can be sent at any time. Connection-oriented networks check first via the call set-up mechanism that a fixed end-to-end route is available. The main advantage lies in the fact that commercial network