LEP to an *fcc-ee*

A bridge too far?

DPHEP workshop, June 2021

See also:
https://indico.cern.ch/event/949986/

Jamie.Shiers@cern.ch

Introduction

- This talk attempts to summarize (a summary of) key official documents, presentations and papers regarding LEP computing (preparation, execution)
- **Why? To capture this information before too late**
  - **Of interest to possible future machines???**

- **2020 ESPP:**
  - ***The vision is to prepare a Higgs factory, followed by a future hadron collider***
  - ***Given the unique nature of the Higgs boson, there are compelling scientific arguments for a new electron-positron collider operating as a Higgs factory***
  - ***An electron-positron Higgs factory is the highest priority next collider***
  - ***Further development of internal policies on open data and data preservation should be encouraged, and an adequate level of resources invested in their implementation***

- 2020: ~1/2 way between end of LEP and start of fcc-ee
- (Most) LEP data still available and used – can we / should we keep it this way?

Sources (Scanned in 2015)

1. Computing at CERN in the LEP era (May 1983 – available at https://cds.cern.ch/record/98880/files/)
2. The MUSCLE report: The Computing Needs of the LEP experiments (CERN Data Handling Division DD/88/1 https://cds.cern.ch/record/184457/files/CM-P00059881.pdf)
3. Computing at CERN in the 1990s (July 1989 – https://cds.cern.ch/record/206085/files/Computing%20at%20CERN%20in%20the%201990s.pdf)

- **The somewhat (sadly) incomplete CHEP record – try to identify some key transitions plus point out some significant gaps!**

- **It is not really possible to do justice to all of these in a short presentation – please look at the complete report (and / or the original sources) for full details!**

“The Green Book” – May 1983

- There was no formal “ESPP” that led to LEP but an e^+e^- machine was widely considered to be the “next priority” at CERN (machine was approved in 1981)
  - Main WGs of LEP Computing Planning Group (LCPG):
    1. Data acquisition & monitoring
    2. Networks and telecommunications
    3. General Computing Services & Facilities
  - The LCPG made 22(!) recommendations (see notes / paper)
- Just over 6 years before LEP **HUGE** uncertainties in all areas: Computing, Storage, Networking, Costs, …

CERN in 1983 – 6 years before data!

- Telex to communicate with sites
  - Not e-mail, not even FAX
- No central interactive services
  - VM/CMS; VAX/VMS and workstations to come
  - *"… it is not just the top of the market (PERQ and APOLLO) which could be of interest, but also low end (DEC, IBM, and **why not APPLE?**)".*
- Huge uncertainties in network technologies, cost and bandwidth
  - *Cambridge ring, token ring, token bus, … even Ethernet*
  - *Heterogeneity and gateways planned and built (MINT, GIFT, …)*
- Costs: 320KCHF for 8MB mainframe memory
- (Pyjama paper – not cut sheet A4!)

Offline Data Processing

- *In the LEP context, support for offline data processing will be most meaningful in the area of well defined support packages, of which the* ***GEM memory management*** (system?) *is an example, and in the area of well defined support facilities, of which the* ***Patchy source code management system*** *is an example.*
  Other areas which are growing in importance for physicists, and where a concentrated support effort might be needed in order to provide the required facilities, ***include advanced interactive graphics*** **(what other kind is there?)**, ***database management, and user transparent network interfaces***.

- **Beware of making predictions, particularly about the future…**

Green Book – Summary

- A huge amount of work went into its preparation
- A lot of background work long since lost
- Many assumptions – such as structure and layout of the CERN Program Library – not described
- (And we've "lost" a lot of the source code anyway…)
- But it set the scene for offline computing at CERN – at least until LEP startup, after which "**all hell broke loose**"

The MUSCLE Report – Jan ’88

- Meeting to Understand the Specific Computing needs of the Lep Experiments: CERN DD/88/1 (A DOI?)
  - Chapter 1 – Introduction
  - Chapter 2 – General elements of LEP computing
  - Chapter 3 – Data storage
  - Chapter 4 – On-site links
  - Chapter 5 – Estimating the requirements for LEP computing
  - Chapter 6 – Computing outside CERN
  - Chapter 7 – Distribution of the computing activities
  - Chapter 8 – The impact of “private” mainframes
  - Chapter 9 – Resources
  - Chapter 10 – Recommendations
  - Appendix – Comparisons with other experiments (e.g. UAn, HERA)

- ***"Many of us… are only now starting to understand the enormity of LEP computing".***

- ***"Physicists should not find themselves in a situation where they have to come to CERN to obtain adequate computing resources".***

  - **Analysis and simulation centres? T1s / T2s?**

➢ **A *MUSC fcc-ee E*?**

Computing at CERN in the 90s

- AKA “the new green book”
- Very comprehensive report, covering far more than computing for LEP
- Impact can still be seen today – 30 years later!
  - Computing for Experiments;
  - CERN Data Networking Requirements in the 90s;
  - Accelerator Computing Requirements in the 90s;
  - Computing for Engineering at CERN ditto;
  - MIS Computing ditto;
  - Theory Computing ditto.
- 10 recommendations and a “computing plan”:

1. Increase computing resources and review in the context of the CERN Scientific Programme
2. HEP-CCC should evolve ... with a strong mandate to coordinate European HEP Computing
3. Coordination of mainframe production resources across European sites
4. “Private” computing facilities should be carefully controlled
5. **Major Components of the Computing Plan (5 years)**
6. 5-year staffing plan (more, of course);
7. Alternative Staffing methods (ditto);
8. Training (see 2020 EPPSU!);
9. Advanced development projects (ditto);
10. The next steps that “we”, the CERN management should define for the next 5 years.

- Long-distance networking (2Mbits / s?)
- An evolution to more Distributed Computing (FDDI?)
- Data Storage and Handling (250K cartridges, 1TB disk)
- Processor Power (600 CERN Units = 370/168 or 8600)
  - ***A pilot Unix service on central IBM asap (Unix & VBS!!!)***

- Computing at the Experiments
- Software diversity – *common software environments*
  - ***HEPCCC should establish a software support committee***

- MIS
- Computing for Engineers
- Accelerator Control
- Computing for Theorists

Summary of “coloured books”

- A significant fraction of a decade went into the preparation of these 3 documents
  - Say 2 years for each of the green books, 1 year for the MUSCLE report

➢ **But LEP data taking was only just over a decade!**

- Would a lighter weight approach have worked?
- Such as the on-going MEDDLE / FOCUS meetings during LEP running?
- How long can one expect (IT) technology predictions to be valid for anyway?

➢ **(BTW, LHC-related technology predictions already lost)**

CHEP for LEP

- Some papers can be found through CDS; others through InSPIREHEP; some "open access" journals ask for license / payment
- No unique "collection"
  - (International) Conference on High Energy (and Nuclear) (Physics) (try all permutations)
  - Also inspirehep.net "CHEP" (These lead to different results ☹)
- No agreement on which was the 1st conference!)
- **MANY missing papers, MANY poorly (illegibly in some cases) scanned, some REDACTED!!!**

➢ **JUST WHAT IS THIS? (Its not preservation)**

- **Can probably be recovered by scanning (properly) and OCR-ing existing "hard copies", ResearchGate and other sources**
- **Plus POLICY for future CHEPs to preserve "born digital" presentations, papers, "templates", agendas etc etc**

CHEPs (mainly) with LEP content

- CHEP 1985 – Paolo Zanella, DD Division Leader
- CHEP 1987 – see Science Direct archive!
- CHEP 1989 – David Williams, CN Division Leader
- CHEP 1991 – Les Robertson, SHIFT and WLCG project leader
- CHEP 1992 – SSC, LHC, OO, …
- CHEP 1994 – Data Handling for LEP & future.. R. P. Mount
- **CHEP 1995 (last one with multiple LEP papers – but not online!)**
- CHEP 1997 (no LEP papers found)
- CHEP 1998 (ditto)
- CHEP 2000 (2 DELPHI papers; one on **"LTDP"**)
- **DPHEP 2009 – joint LEP paper on LTDP**
- CHEP 2012 – first (and last) concurrent DPHEP "track"
- **DPHEP "Full Costs of Curation" workshop**

CHEP 1985 (Amsterdam)

- ***"Even the most reluctant physicists are now convinced that computers are critical for the success of their experiments".***
- ***"To satisfy the computing needs of … the LEP era, a straightforward increase in overall capacity is not enough. New facilities for data acquisition, storage, processing and communications are also required".***
- ***"FORTRAN is probably the only perennial standard that will never be questioned"***.

Paolo Zanella

CHEP 1987 (Asilomar)

- **DEVELOPMENT OF SOFTWARE FOR ALEPH USING STRUCTURED TECHNIQUES**
- USE OF THE ADAMO DATA MANAGEMENT SYSTEM WITHIN ALEPH
- PAW - TOWARDS A PHYSICS ANALYSIS WORKSTATION
- SOFTWARE ISSUES FOR LARGE DETECTORS (not directly related to LEP)
- TWO YEARS OF REAL PROGRESS IN EUROPEAN HEP NETWORKING: A CERN PERSPECTIVE
- DATABASE SYSTEMS FOR HEP EXPERIMENTS
- **PERSONAL COMPUTERS IN HIGH ENERGY PHYSICS**

Gottfried Kellner on s/w in HEP

- *"Software systems are very often badly designed, badly or not documented, delivered too late, very costly to modify and unstable against upgrades. Our own environment, i.e. High Energy Physics, is no exception.*

- *Why is this so? It is certainly true that concepts of software engineering were only introduced some 15 years ago. The tendency is to write code immediately — because one knows what one wants —and think afterwards how to modify it to include other options. Software writing is still an art — some people are very good at it but they don't (or can't) explain to other people how they do it."*

David Quarrie on PCs

- The first suggestion that I know of regarding PCs – as opposed to workstations (green books) for physics-related work
  - Pointed out the numerous deficiencies at the time

- A decade or so before other “first proposals”?

CHEP 1989

- *"It is rather amazing that the computing for LEP appears to be in such good shape. But much work remains ahead of us. The data volume that we are faced with is one* (to) *two orders of magnitude more than we have faced before."*

➢ **But already there were several talks where the mainframe's role was greatly reduced…**

- ***THE OPAL EVENT SERVER***
  - This paper presented a hybrid approach, using a CERN mainframe and workstations and other computers both in the local and wide area. Notwithstanding the relatively low bandwidth of the network connections available at the time, they were still sufficient for messaging and the retrieval / distribution of small event samples.
- ***THE ALEPH EVENT RECONSTRUCTION FACILITY: PARALLEL PROCESSING USING WORKSTATIONS***
  - This paper described one of an increasing number of the use of multiple workstations where previously a mainframe would have been used. In this case the features of the VMS operating system were used to advantage and similar solutions, based either on VMS or increasingly Unix, were a feature of the early days of LEP.
- ***OVERVIEW OF THE ESSENTIAL TOOLS***
  - In this model, the mainframe assumes the role of data server, feeding and receiving data from multiple workstations. Whilst this might work well as a centralized solution, it leaves remote physicists somewhat ostracized. By coupling similar systems at a small number of remote sites (file server, DB server, workstations) and ensuring synchronization between the sites, for which only modest networking is required, this problem is solved.

CHEP 1991

- Is the role of the mainframe terminated? (DOW)
  - A paradigm “shift” since CHEP 1989
- Computing at LEP (Manuel Delfino)
  - IMHO practice vs theory (MUSCLE)

- **MISSING PAPERS:**
  - **SHIFT: The Scalable Heterogeneous Integrated Facility for HEP computing**
  - **The network and communications environment for a large collaboration by Mount**
  - **Reality of software engineering in high energy physics**
  - **The L3 database system Database management and distributed data in high energy physics: Present and future**
  - **The DELPHI off-line packages: Difficulties and trends**

CHEP 1992

- This was the CHEP were OO and the use of non-HEP, “standard” tools came to the fore
- Also where WWW won the conference prize
  - And not PAW++, nor the CERNettes…
- LEP talks:
  - Offline analysis system of OPAL (SON)
  - OPAL MC system
  - Data Storage and Network in Offline Analysis of ALEPH (Ronald Hagelberg)
  - Distributed analysis using MAW (Adrian Buijs)
  - DELFARM (Michelotto), SHIFT (Panzer), …

➢ Distributed and / or widely distributed systems

CHEP 1994

- This conference featured a paper *"Data handling at CERN for LEP and future experiments"* (p. 245) by Mount, R P (but no online version).
  - Recently, when archiving my AFS / DFS / EOS space I found a copy of Richard's paper:
    a) What was it doing there?
    b) How does one submit it to the official record?
- (74 contributions can be found through InspireHEP, ***some*** of which have associated PDF: https://inspirehep.net/conferences/971233)

Why even bother with LEP?

- The total data volume is now trivial but **preserving** the software, documentation, meta-information and knowledge certainly is not. (Not to mention planning documents, papers, presentations and so forth)

- There are those that argue:

  - **If / when a new *fcc-ee* (or other Higgs factory) is built, its precision and statistics will soon make the LEP data totally redundant;**
  - **Even if not, can there be any information left that we failed to extract at the time, when we had fully trained active collaborations?**

➢ **But can this data, this knowledge, and this experience really be of no value, even in the interim?**

➢ **(At a minimum, this should be some sort of reality check for those planning offline computing for future machines)**

“My Humble Opinion” on LTDP

- **Preservation is an activity that cannot be left to individuals but needs to be part of an overall policy of the (virtual) organisations involved.**
- **A “policy” is not just one phrase in the ESPP…**
- **And it must be born out by (measured) practice….**
  - Much information, data and knowledge has already been lost regarding the once flag-ship LEP experiments (& LHC!!!)
  - Consider also the many pre-LEP ones, where shadows remain in the Particle Data Book and in the HepData database.

➢ **Whilst in theory it is never too early to consider LTDP, in practice it often turns out to be too late.**

COMPUTER COMMUNICATIONS AT CERN

B.E. Carpenter
DD Division, CERN, CH-1211 Geneva 23

CERN has been a leading user of data communications technology for many years, and yet still needs to extend and improve its facilities. A major review of needs was carried out in 1982, consolidation has taken place since then, and a general policy for the future has now been proposed. This is based on the OSI model, on the short-term use of non-standard protocols, and to a large extent on a backbone network and gateways.

1. INTRODUCTION

CERN has been near to the state of the art in data communications for more than a decade. Major achievements include the site-wide general-purpose CERNET [1] and the special-purpose controls network [2], followed more recently by the STELLA satellite experiment [3]. In the LEP era we must solve the problems of multiple sites spread over many kilometres, and the interconnection of systems from many manufacturers, and provide facilities for physicists based in hundreds of institutes and universities world-wide. Thus data communications, in the most exciting phase of its history, is an area where CERN is investing scarce human and financial resources with high priority.

This paper first briefly reviews current services, before outlining our general policy and strategy for meeting the challenge facing us. Two technical appendices outline the various standard and ad hoc protocols of interest to CERN.

2. CURRENT SERVICES

The largest system today is the INDEX circuit-switching network for ASCII terminals. This now has about 1400 terminal connections and 1000 computer ports, interconnected by three central nodes (to be replaced by two larger units during 1985) and six subsidiary nodes. In addition, we have some 30 internal modem connections, and a similar number of external modem lines.

The variety of these external connections is too great to detail here. Two are of particular interest: the IBM-sponsored EARN network and the connection to the Swiss PTT X.25 network Telepac and to private X.25 networks via international leased lines. These complementary services give access to many academic institutes in Europe and North America.

The internal packet switching network, CERNET, has about 110 computers on a mesh of 18 switching nodes and 2Mb/s links. The greatest user is the "emulator farm" processing collider data in full flow, its average data rate is about 80 kilobytes/s on all lay-out lines, consisting of transfers between the IBM 3081 and three dedicated PDP/11s. Over 20 Norsk-Data computers are on the SPS and PS controls networks.

During 1984, Ethernet grew from a small pilot project to a serious installation. There are now at least seven Ethernets on site; that in the Computer Centre has about 20 machines connected to it.

UAB-LFAE 91-07

Computing at LEP

M. Delfino
Laboratori de Física d'Altes Energies
Universitat Autònoma de Barcelona
E-08193 Bellaterra (Barcelona) Spain

(Invited talk at the Computing in High Energy Physics 91 conference, Tsukuba City, Japan, March 11-15, 1991)

A general review of computing at LEP is given, with emphasis on event reconstruction, data flow and storage, and physics analysis. The description is from the point of view of the strategies used to solve the problems; the reader is referred to the parallel session talks in CHEP91 for technical details and information which is specific to each of the four LEP detectors.

CERN-COMPUTING & NETWORKS DIVISION
CN/91/2
March 1991

HEPnet in Europe: Status and Trends

François FLUCKIGER

OCR-ing can improve quality!

1. Introduction

Although the use of large disk subsystems in Particle Physics data analysis has become common-place in the last 10 years, the random access features of disks have rarely been used. Most of the time, disks were used as magnetic tape emulators, with the advantage of simultaneous access by many users and elimination of the delay of mounting a tape. This approach was used, for example, by the MAC experiment at PEP [1], and then became more fomally implemented as the "tape staging" extensions to the IBM VM and VAX VMS operating systems. Some operating systems, from CDC for example, have supported tape staging for decades. In this paper, we describe the exploitation of true random access on large disks for event analysis of the ALEPH detector at LEP. The random access features were designed and implemented by V. Blobel in the input/output (I/O) modules of the BOS memory management package [2]. It should be noted that although the description always refers to disk files, all of the features are also available for magnetic tape cartridges staged to disk.

2. Why Random Access ?

High Energy Physics data are usually arranged in files, each corresponding to a period