CERN-Data Handling Division
DD/88/7
J. Joosten
April 1988

BRIDGED ETHERNETS AT CERN AND BEYOND

(To be presented at the 3rd Symposium on Electronics for Telecommunications, Porto, 4-6 May 1988)

DD-vf

Bridged Ethernets at CERN and beyond.

J. Joosten
CERN, 1211 Geneva 23, Switzerland

Abstract

This paper will address the evolution of the general purpose internal network at the European Laboratory for Particle Physics (CERN), which is currently based on Ethernets interconnected with MAC-level bridges.

This architecture has been selected because the site is too large to be covered by a single Ethernet and because many different protocols are being used. Furthermore local traffic does not cross a bridge, hence the potential combined capacity is very high.

With the start-up of the Large Electron-Positron collider, the requirements for site-wide, long distance communication will grow dramatically. The need for a site-wide 100 Mbit/s backbone is apparent. Some ways of tackling this problem will be described.

The management of such a widely distributed bridged network with about 600 systems connected to it and still expanding at a high rate is a difficult task, especially in the rapidly evolving environment that confronts us.

The available management tools will be described as well as plans for an integrated network management system, using a data base containing all the relevant data on the network and its users.

I Introduction

When the first Ethernet was introduced at CERN in 1983, general purpose networking of computers was generally done with Cernet [1]. Cernet is a site-wide, mesh-structured, packet switching network, designed and built at CERN in 1975. The communication hardware was designed for line speeds of 2.5 Mbit/s. It had however been conceived in such a way that the use of the CCITT G703 transmission standards of 2 and 8.5 Mbit/s [2] could be easily introduced at a later stage, giving Cernet a practically unlimited span.

The communication subnet provides a datagram service, and software installed on the connected machines provides the services for Transport and higher levels. Thus connectivity was obtained between mainframes (IBM, CDC Cyber) and over 100 minicomputers (VAX, PDP, HP, Norsk Data).

The trend to personal computers, workstations, terminal concentrators etc., together with the emergence of international standards for Local Area Networks (LANs) [3] to which these devices tend to be connected, forced us to solve the problem of internetworking physically distinct LANs (because of distance) and the use of different protocols on these LAN's (DECnet, ISO/OSI, TCP/IP, XNS etc.).

Since no commercial products were originally available to solve our problem, it was decided to use the existing Cernet infrastructure as a vehicle (backbone) for interconnecting the LANs in a transparent (protocol indepedent) way.

The idea was to be able to migrate gracefully to a high performance backbone when available.

II Cernet as a backbone for Local Area Networks

A project was started with the aim of connecting the different Ethernets via MAC-level bridges [4] [5] to the Cernet backbone. By 1987, 13 Ethernets distributed over the CERN site were interconnected via such bridges called Frigates. [6]

The switching capacity of a Cernet node, to which an Ethernet segment is connected through a Frigate is about 170 packets/s. Although not much for today's requirements, it was hoped that it would be sufficient till a new technology based on a fast (100 Mbit/s FDDI) ring would be available. This proved however to be wrong for the following reasons:

- The explosive growth of DECnet: a factor five in 12 months (from 40 to 200 nodes) and its generous use of multicasts, mainly due to lack of active DECnet management.
- The multicast multiplication factor: Since a multicast issued on one segment had to be forwarded to all other segments, each multicast packet had to be multiplied by a factor 12 (for 13 segments). This is done in the Cernet node to which the Frigate is connected. It is easy to understand that this would cause the Cernet nodes to be overloaded, long before real data traffic could saturate the system.
- Progress on FDDI proved to be much slower than anticipated and essentially no products can currently be obtained to satisfy our requirements.

A quick solution had to be found to handle this crisis. The only choice proved to be the introduction of one Ethernet segment as a backbone for all the inter-Ethernet traffic on the site.

III The Intermediate Back-Bone

An Ethernet segment was set up as a backbone to which all other segments were to be connected via commercial high performance bridges. This solution called the Intermediate Back-Bone (IBB) gives a throughput of about 7 Mbit/s, which is hopefully sufficient till the 100 Mbit/s backbone becomes available.

The introduction of an IBB could of course not be done overnight and a migration plan had to be made consisting of the following steps:

- All the segments within 1.5 km from the backbone segment in the computer centre (called High Rate Segment, see figure 1) would be gradually connected via a fibre-optic repeater and a local bridge. To achieve this, a fibre-optic infrastructure had to be installed on the part of the CERN site surrounding the computer centre. This infrastructure is now partly available and two of the most busy segments are now connected like this to the HRS.

- Meanwhile, segments that were more than 1.5 km away were to be connected with long distance bridges using 2 Mbit/s channels of the TDM infrastructure that is progressively being installed on the site.
  Initially these connections are based on two Frigates connected back-to-back, using the Cernet transmission scheme at 2.5 Mbit/s, or TDM at 2 or 8.5 Mbit/s. These Back-to-Back Frigates (BBF's) have each a aggregate throughput of 500 packets/s.
  Where more throughput is required high performance commercial long distance bridges, which are now becoming available will be gradually installed. As is shown in figure 1, already one such a bridge is operational.

- For performance reasons a Medium Rate Segment was introduced to which only Frigate based equipment is connected. The reason is, that these devices use CPU time for filtering, meaning that their throughput is affected if the segment is heavily loaded with traffic that is not meant to be forwarded.

Is is hoped that by the end of the year, all the Cernet connected bridges are replaced by fibre-optic connected bridges, BBF's and commercially available long distance bridges, running at least at 2 Mbit/s.

Medium term evolution

A Large Electron-Positron Collider (LEP), with a diameter of about 8.5 km, is nearing completion. There will be four underground locations (pits) along the ring where the particles will be made to collide and where very large experiments will be set up. Some of these locations are about 10 km away in the country side.

Each experiment will have a number of computers for data acquisition and control requiring Ethernet connectivity. Taking into account that most LEP experiments will also have computing equipment in the annex of the computer centre (called the 'barn'), it is planned to set up there a new segment to which the the four regions will be connected with remote bridges using 2 Mbit/s TDM links. This segment forms a sort of mini backbone for the LEP experiments and will be connected with a bridge to the HRS (backbone) segment, which provides the connectivity with the rest of the CERN site. Figure 2 shows the probable final topology of the IBB exercise.

It should be noted that this topology is not perfect since each LEP experimental group would like to have a (private) triangular high performance connection between their offices, their computers in the barn and their experiment in the pit. However there are some technical, practical and financial problems:

- Bridges cannot be used in a triangular topology, because this would create a loop.
- The TDM infrastructure is not yet covering the entire site. It is therefore not (yet) possible to make a remote bridge connection between an experiment and the offices of the physicists.
- The above mentioned offices form already part of the site-wide Ethernet that is shared with other people. This has to be taken into account when adding new "private" segments.
- It is unpractical and expensive to have several parallel Ethernets in the same physical locations.

The situation will be reviewed in due time.

IV Ethernet cable infrastructure

Providing site-wide Ethernet connectivity is not a trivial matter for a big site as is the case for CERN, with many buildings distributed over a large, complex geographical area.

Taking into consideration that management and trouble shooting is best done in a star shaped topology, an effort has been made to shape the Ethernet infrastructure in primary and secondary stars. This has been done as follows:

All the major segments are connected via bridges to the central backbone segment and this forms the primary star. (In practice it is a set of two to three interconnected segments.) In this way no major segment will be used as a hop for another major segment.

On a main segment, which is generally a yellow coaxial cable of length up to 500 meters, secondary stars are formed by installing in strategic places multi-port repeaters from which Thin Wire Ethernet segments fan out to the offices, laboratories etc. Each Thin Wire segment services a restricted area which forms a topological unit. It is important to limit the scope of one Thin Wire segment, because it can be easily disrupted and should therefore be restricted to service a small self-contained user community.

At the secondary stars one tends also to connect bridges that have to isolate segments with a lot of local traffic (for a cluster for instance).

There are currently about 70 buildings with access to Ethernet. It is estimated that most of the site will have proper access to Ethernet in 18 months time. This time takes already into account that some of the cabling is done in parallel.

It is of course impossible for a central team to know all the details about the Thin Wire infrastructure in all the buildings. Therefore the job of a First Line Interface Person has been created. Every region should have such a person, who is responsible for the proper operation of the Thin Wire Ethernet segments and who should be able to solve certain problems before deciding to escalate.

V Network usage

Because of the multi-vendor and multi-national characteristics of the laboratory and its large community of visitors, it became clear very soon that a large variety of systems using different communication protocols would be used at CERN. The decision was therefore taken to connect all the Ethernets with MAC-level bridges which are transparent to protocols.

DECnet

There are currently about 250 DECnet nodes at CERN and DECnet is currently the most widespread single application on our network. On the segment with the central services it represents 20% of the packets.

Local Area VAX Clusters (LAVC)

There are some ten LAVC's on the site. Since these systems can generate a considerable amount of traffic and even generate traffic if no work is being done, the rule was established that they would be shielded with a bridge to keep them really local. This meant also blocking the multicast traffic such clusters produce for internal usage.

Seen the distributed nature of the certain activities, this rule is already being challenged. A solution based on a concept called superclusters which uses DEC's Distributed File System (DFS) to share data between clusters is under study and should provide the advantages of a LAVC, without the drawbacks [8]. On one general purpose office segment 75% of the packets were seen to be of the LAVC type.

Local Area Transport (LAT) protocol

This protocol is being used between a VAX/VMS system and a DECserver. There are some 60 (8 port) terminal servers of this kind operational and the number is growing fast.

In fact, since the Gandalf terminal switch is reaching saturation with about 2000 terminals connected to it, the only expansion for terminal connections is via servers connected to Ethernet or via the digital telephone system. The latter is however for practical and economic reasons not yet used apart from some specific cases.

The DECservers are also used in the so-called reverse LAT mode, allowing a user to enter the server over Ethernet and to exit via a terminal port which is connected to Index, a terminal multiplexer of a mainframe etc.

40% of the packets on the central services segment is LAT traffic.

TCP/IP

TCP/IP, being a de facto standard, gets more and more used at CERN. The major applications are in PC type systems, certain terminal concentrators, access to large host systems (Ultrix, Unicos, VM, VMS), as well as for Apollo systems when not using the Domain network.

It is also used in the LEP control networks, based on an interconnected set of token rings (IS 8802.5 standard) [7]. These token rings are connected to the site-wide Ethernets via IP gateways. In contrast to MAC-level bridges, gateways give a better protection against unwanted interference, which is essential for an accelerator control system. Furthermore the 8802.3 and 8802.5 standards are sufficiently different to defeat an attempt to bridge on the MAC-level.

A major future use of TCP/IP will probably be for workstations for physics analysis, which need to have remote access to large data sets.

There are currently around 250 IP addresses allocated and this number is growing rapidly. On the segment with the central services computers 25% of the packets are TCP/IP.

XNS

XNS got introduced at CERN as part of a turnkey system. It provides connectivity between 13 Daisy CAE workstations.

Kinetics-AppleTalk

It is possible to communicate between two AppleTalks over Ethernet via an interface called Kinetics FastPath. This device can also function as an IP gateway, providing TCP/IP (TELNET and FTP) functionality for a Macintosh on AppleTalk. One can also connect a Macintosh directly to Ethernet and use AppleTalk over Ethernet (called EtherTalk) or TCP/IP.

Unfortunately, these systems make generous use of broadcasts, which become a problem if many such devices get connected to a very large bridged Ethernet system.

There are many Macintoshes at CERN, but only a few are connected to Ethernet. The situation could change rapidly. For efficiency it is recommended to connect groups of Macintoshes to an AppleTalk network with one gateway to Ethernet.

PC Clusters

Clusters of PC's are being gradually introduced. Systems from several manufacturers are installed and it is hoped that one can end up with one recommended system. PC clusters tend to use their own specific protocols, adding to the variety of protocols already being used. In contrast to LAVC's, the PC clusters currently at CERN do not load the network when not being used and are currently not restricted to a local segment.

Remote Procedure Calls

A software package based on the Remote Procedure Call (RPC) technique has been developed at CERN [9] and it is being used to communicate over Ethernet between VME based test systems called VALET [10] and a variety of peripheral servers.

MAID

Since a considerable amount of high level software was developed for Cernet applications, it was judged effective to use this software over Ethernet. This has been achieved by introducing the so-called Transport MAID (Manager Automatic Internet Datagram) service. This protocol is essentially used on the central Ultrix service to access the Cernet filemanager on the IBM MVS system via Ethernet.

ISO Transport Protocol Class 4 (TP4)

Since CERN is committed to the use of appropriate international standards, an effort has been made to introduce ISO protocols. It is however clear that in spite of commitment to ISO protocols by essentially every serious manufacturer, very few packages are currently available. A few applications are nevertheless running at Cern. One is on a PC cluster running MSnet on top of TP4. Another set of applications is on small systems used for monitoring and control in experiments and accelerator control, allowing these systems to communicate with a VAX or an Apollo.

Network File System (NFS)

NFS has been recently introduced at CERN and is being used between Unix systems. One application is between the central Ultrix service and certain VAXstations. Another is between a Cray and SUN workstations on a private secure Ethernet with strictly controlled access from the the rest of the world (by means of a special IP gateway).

Network management traffic

A certain amount of traffic is generated by monitoring and network management activities. It is mainly consisting of command/response type of traffic.

Summarised, it can be concluded that we are living in a multi-protocol environment and that this will continue to be so for some time to be. It is even likely, that in spite of standardisation, there will always be special protocols around, which are tailored to be very efficient for specific applications.

It is also fairly certain, that on our bridged Ethernet certain protocols are being used, of which we do not know the existence. If they do not cause trouble we might eventually know about their existence, but because of the MAC-level bridging approach it does not really matter.

VI Network Monitoring and Management

Monitoring and management are crucial to the good operation of a network. Although we started out with very poor tools, good progress has been made over the last year, and we are not completely blind as to what is happening on the site-wide bridged Ethernets and the evolution of traffic patterns.

The plan is to concentrate the network management in one computer system dedicated to this task. The heart of the management system should be a data base in which all the relevant data about the topology and the systems connected to the network is present.

Using tools of which some are described below and others that are to be produced yet, it should be possible to permanently watch the behaviour of the network and give an alarm on abnormal conditions.

The definition of what is an abnormal situation will probably never be written in stone and will also depend on the degree of control one wants to have over the network.

The network management system should provide an operator interface suitable for day-to-day operational problems. It should also provide engineers and managers the information they need to do their job, i.e. trouble shooting, analysis and planning.

Data base

A data base has been set up in which all the information on the Ethernet infrastructure is being captured. Information on the physical lay-out of the Yellow and Thin Ethernet cables, how they are interconnected with repeaters and bridges, the position on a cable of transceivers and fan-out units is all entered in the data base. Relevant information on the host systems connected to Ethernet, should also be entered in the data base, i.e. its type; location; Ethernet, IP, DECnet addresses etc.

A considerable amount of information is already available in files kept by certain individuals and is to be copied into the data base. In the end private files should be extracted from the data base, being the focal point of all relevant information on the network.

To keep the data base up to date in the very dynamic and relatively uncontrollable environment we are living in will never be 100% successful. Even so, the benefits will still be considerable.

Since the data base does not provide a visual picture of a cable lay-out, a set of drawings showing the basic topology is made in parallel.

Frigate management

Remote Frigate management allows for control and monitoring of the remaining 8 Frigates. It is possible to remotely collect the statistics that they produce and when done on a regular basis, the evolution of traffic on a segment will help to anticipate the need for future topology changes. Although advance knowledge about future connections to a segment is important, regular traffic measurement is the final judge.

The Frigate statistics are also an excellent means of detecting all sorts of abnormal conditions on a segment, that are normally not seen by commercial bridges, because they do not look at it or because they do not get into trouble seen their high throughput.

It is therefore planned to equip all the segments with cheap single board systems that run the Frigate monitoring and statistics software. These devices will be remotely accessible from the management machine.

Whenever a Frigate gets reloaded, it sends a packet to the network management machine, which reports the incident and sends back the correct date and time. Having the correct date and time in a distributed computer system is very important for analysing problems.

Back-to-Back Frigates are managed in an identical way.

Remote Bridge Management Software (RBMS)

RBMS is a DEC product to manage their bridges, of which there are about ten installed at CERN. This package allows to control the topology of a set of interconnected bridges as well as the setting of certain parameters. It can also collect statistics from the bridges. Another feature is that one can interfere with the forwarding address table in the sense that specific addresses can be blocked, meaning that certain packets will not be forwarded. Unfortunately it is impossible to freeze an address table, meaning that a bridge does not stop learning.

LAN Traffic Monitor (LTM)

This DEC software product operates in combination with a bridge set up as a monitor, providing detailed statistics for an Ethernet segment. It gives utilisation figures (peak and average); it tells what are the most busy systems on a segment etc. It is therefore an excellent tool for finding out about the general utilisation of a segment. It is however not intended to analyse the behaviour of protocols and the inter-packet timing for which other commercially available monitors are available.

Multicast traffic monitoring

One of the first monitoring programs written, collected all the multicast packets on the network and produced a statistics about the frequency at which they were issued by the systems on the network. This proved to be very useful to locate abnormal usage. The program is written in assembler and executes in a VME system.

TCP/IP monitoring

Another software package has been written recently that collects all the packets circulating on the segment to which the network management machine is connected. An analysis is made of each IP packet and a table is build, that correlates the Ethernet source and

destination addresses with the IP addresses. This list can then be compared with the list of IP addresses distributed to the users, giving a means to find inconsistencies, errors etc.

Type field and multicast address distribution

For internal applications, it was decided to distribute Type fields in the 5000 – 5FFF hexadecimal range. This is convenient for monitoring and trouble shooting. In fact it would have been useful, if such a group would have been internationally agreed, such that no overlaps with commercial products would ever arise.

Similarly, it is important, to give multicast addresses to applications that need it. This will prevent people from using broadcasts, that are seen by all the systems that are connected. A program has been written to generate these addresses in relation to the hash filter available in certain Ethernet chips.

VII Review of problems and incidents

Although MAC-level bridging is a very elegant technique, it is prone to disruption by faulty behaviour of any host on the system. It is however remarkable, that the number of serious problems on the bridged Ethernets has been very limited up till now, especially if one takes into consideration the impossibility at CERN to rigidly control who connects what to Ethernet. The comparison with the mains plug can be made, with the difference that the electricity distribution has fuses.

Some notable problems are described below:

- A TCP/IP worm on the first of April 1987. All the machines running TCP/IP became crazy and their TCP/IP drivers had to killed. It was never found out what happened, but it could have been what follows.
- A certain system put a broadcast address as a source in an IP packet (probably due to bad initialisation); The addressed service responded with a packet having a broadcast destination address; A TCP/IP system with routing capabilities received the packet and decided it was not for him; It retransmitted it with a broadcast destination address.
  One needs only two such routers to create a big problem. Obviously, a packet should not be routed to the network from where it came.
- Once in a while a certain machine floods the network with a burst of multicasts. Since multicasts are forwarded by bridges, such a flood affects the whole network. Sometimes the cause is found and corrected.
- A badly set-up PC pretended to be every TCP/IP service on the site. Dependent on order of arrival of the response to an ARP (Address Resolution Protocol) packet issued by a user, he got to the real service or not.
- By continuously sending packets with random source and destination addresses, a system was able to overflow the address tables in all the bridges.
- There exist systems that communicate very well between each other till one system gets switched off: the other system starts to output packets at a high rate till it gets a response, which might take a very long time.
- There was an application where packets were sent to a system that was not supposed to respond. This confused the bridges however, because they never saw the particular destination address as a source and could therefore not learn on which segment that system resided. The result was that these packets got broadcast to all the segments on the site.
- At several occasions loops were created in the network. This caused a lot of trouble in many systems and some crashed even. The reason for the creation of these loops was twofold: human errors (due to faulty test set-ups in the network laboratory), but also by bugs in a commercial bridge.

The major danger to a bridged Ethernet originates from badly set-up software, errors in poorly tested software packages and software that has not been developed with a large network in mind. We see for instance industrial products that make indiscriminate use of broadcasts, something that affects every system on the network.

One can conclude that apart from occasional mishaps our system of bridged Ethernets is working well. However, there is a reason to worry as the system is growing and the concept of gateways as "fire doors" has been proposed to limit the damage in case of disruption.

VIII Future requirements

With LEP in mind, for which four very large experiments are being prepared, an estimate has been made for the peak data rates that a future backbone has to provide. The outcome was that a peak 100 Mbit/s throughput was required.

In contrast to traffic patterns in conventional networks, most of the data has to travel over the longest distances, which are in the order of 10 km.

FDDI

The obvious choice would be to use the FDDI standard proposed by ANSI X3T9.5 [11]. It was hoped that industrial products would be soon available that allow to connect Ethernets on the MAC-level over FDDI as well as provide high performance interfaces to mainframe computers. Unfortunately, progress has not been as fast as hoped. By introducing our Intermediate Back-Bone we bought hopefully enough time.

Even when FDDI becomes available, we shall have to address several problems:

- Since optical fibres in the LEP area are a scarce resource, dedicating fibres to FDDI is discouraged. The policy is to use the fibres for TDM according to the G703 specifications of the CCITT [2].
- Since FDDI is specified for a maximum distance of two km between active stations, there is also a problem of covering LEP distances that are longer. However, it seems that there is a proposal to remove this limitation.
- An FDDI ring, spanning all the important sites, would be about 50 km long. Common sense tells that it is not a desirable situation that traffic between a set of mainframes in a computer centre would have to travel such a distance. It would make more sense to restrict the FDDI ring to a small physical area and connect remote Ethernets with high speed long distance bridges using the TDM infrastructure.
- In the not to distant future there will probably be several FDDI rings on the site and they need to be interconnected in a way that guarantees redundancy.

It is to early to make a final decision on these topological problems. Meanwhile work is being done to solve some of the above mentioned problems, i.e. a means of passing FDDI traffic over a TDM channel. This is done by encapsulating the 125 Mbit/s FDDI bitstream in a frame at 139.264 kbit/s, which is one of the CCITT standard bitrates. At the output of the TDM channel the original FDDI bitstream is then recovered again. [12].

Long distance bridging of Ethernet to FDDI or between two FDDI rings, is currently under study at CERN and it is hoped that European industry can be encouraged to take an interest in such a project.

Fast point-to-point links

Another problem that has to be addressed is the requirement of certain LEP experiments for fast long distance point-to-point links between non-identical computers. The ultimate requirement is for a sustained throughput of 1 Mbyte/s. There is work going on in industry to provide such products, but currently available products are very limited and not based on our preferred standards. A FDDI ring connecting two mainframes together would be interesting. Such a connection should use the above mentioned FDDI/G703 converter to cross the CERN site.

IX Conclusions

The site-wide bridged Ethernet system, which is now regularly being used by about 600 hosts, is functioning well. It is expected that the introduction of the Intermediate-Back-Bone, based on one central Ethernet segment will give us sufficient bandwidth, till new products in the 100 Mbit/s range become available.

The network usage is however evolving rapidly and new situations may arise that will force us to review our plans. Fortunately, we have now some tools available to monitor the evolution of the traffic and that should allow us to adapt the network infrastructure, hopefully in time.

X Acknowledgements

In this paper the valiant work of a number of people has been reported. In writing this, I also took advantage of the knowledge and expertise of my colleagues in the Communications group and some other persons in the DD Division, for which I thank them very much. For the errors I take the entire responsibility.

References

[1] J. M. Gerard, "CERNET – A High-speed Packet-Switching Network" CERN 81–12, Dec. 1981

[2] "Digital Networks, Transmission Systems and Multiplexing Equipment, Recommendations G701–G956" CCITT red book, volume III–Fascicle III.3, 1985

[3] International Standards Organisation, International Standards IS 8802/3, IS 8802/4, IS 8802/5.

[4] IEEE Computer Society, "MAC and MAC-Bridge Service Definition", document H12, Feb. 1985

[5] "Local Area Networks–MAC sublayer Interconnection (MAC bridging)" ISO/TC 97/SC 6, 1987-01-30.

[6] J. Joosten and C. Piney, "CERN's approach to MAC-level bridging of LANs", EFOC/LAN 87 Proceedings, published by IGI Europe, Basel.

[7] P.G. Innocenti, "The LEP control system", CERN SPS/87-41 (ACC), Presented at the Europhysics Conference on Control Systems for Experimental Physics, Villars sur Ollon, 28.9-2.10.1987

[8] J. Shiers, Cern, DD Division, private communication.

[9] T Berners-Lee, "Experience with Remote Procedure Call in data acquisition and control", IEEE Transactions on Nuclear Science, August 1987 Volume NS-34 Number 4.

[10] Y. Perrin et al, "The VALET-plus, a VMEbus based microcomputer for physics applications", IEEE Transactions on Nuclear Science, August 1987 Volume NS-34 Number 4.

[11] Fiber Distributed Data Interface, ANSI X3T9.5 committee proposals PHY, PMD, MAC, SMT.

[12] A. Leite, INESCN (Portugal) and CERN, "FDDI over a CCITT 140 Mbit/s TDM channel", Report to be publised.

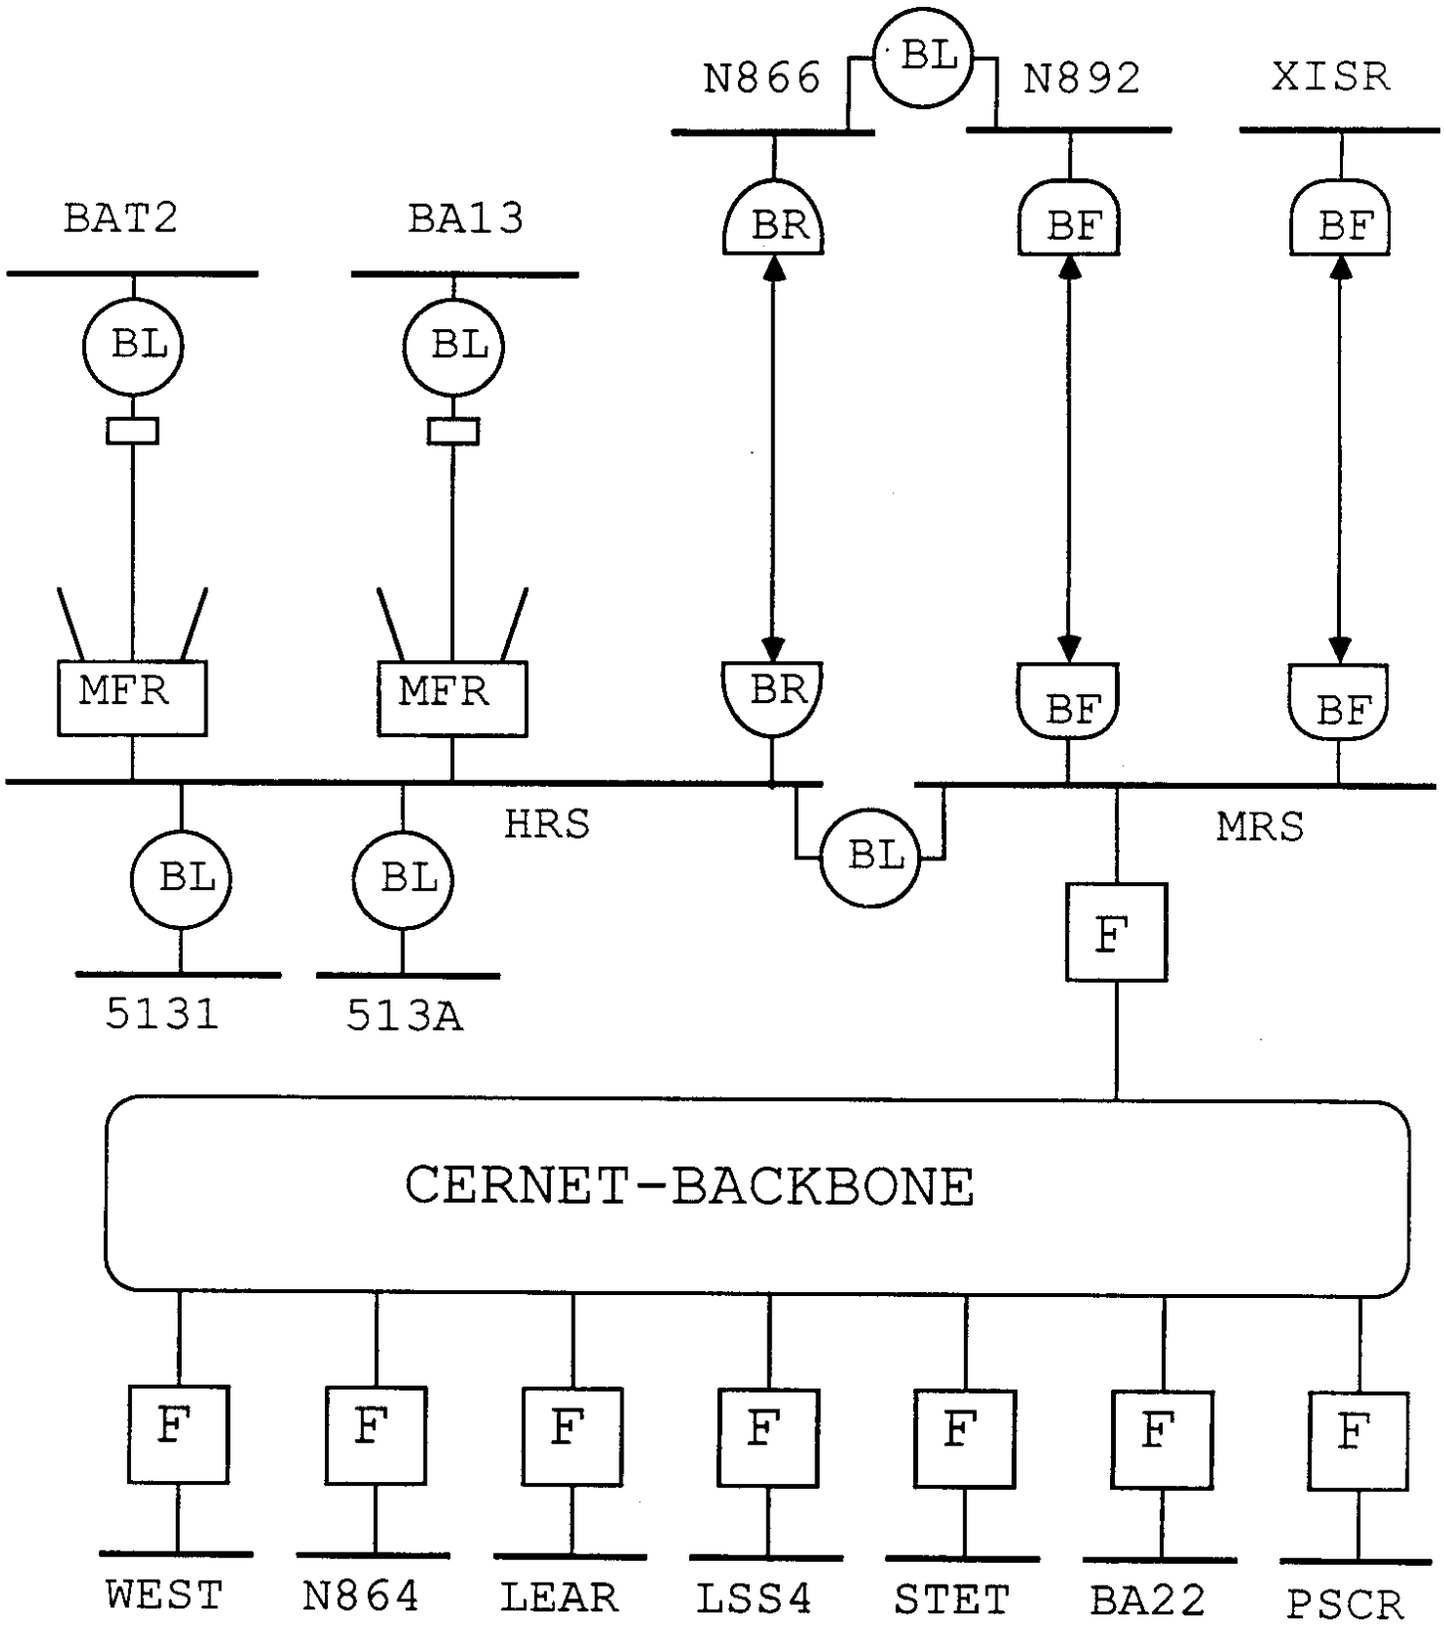

HRS: HIGH RATE SEGMENT
MRS: MEDIUM RATE SEGMENT
FIBRE-OPTIC TRANSCEIVER
MFR: MULTI-PORT F.O. REPEATER
BF: BACK-TO-BACK FRIGATE
BL: LOCAL BRIDGE (10 Mbit/s)
BR REMOTE BRIDGE (2 MBIT/S)
F: FRIGATE

FIGURE 1: BRIDGED ETHERNET TOPOLOGY 1988-03-15

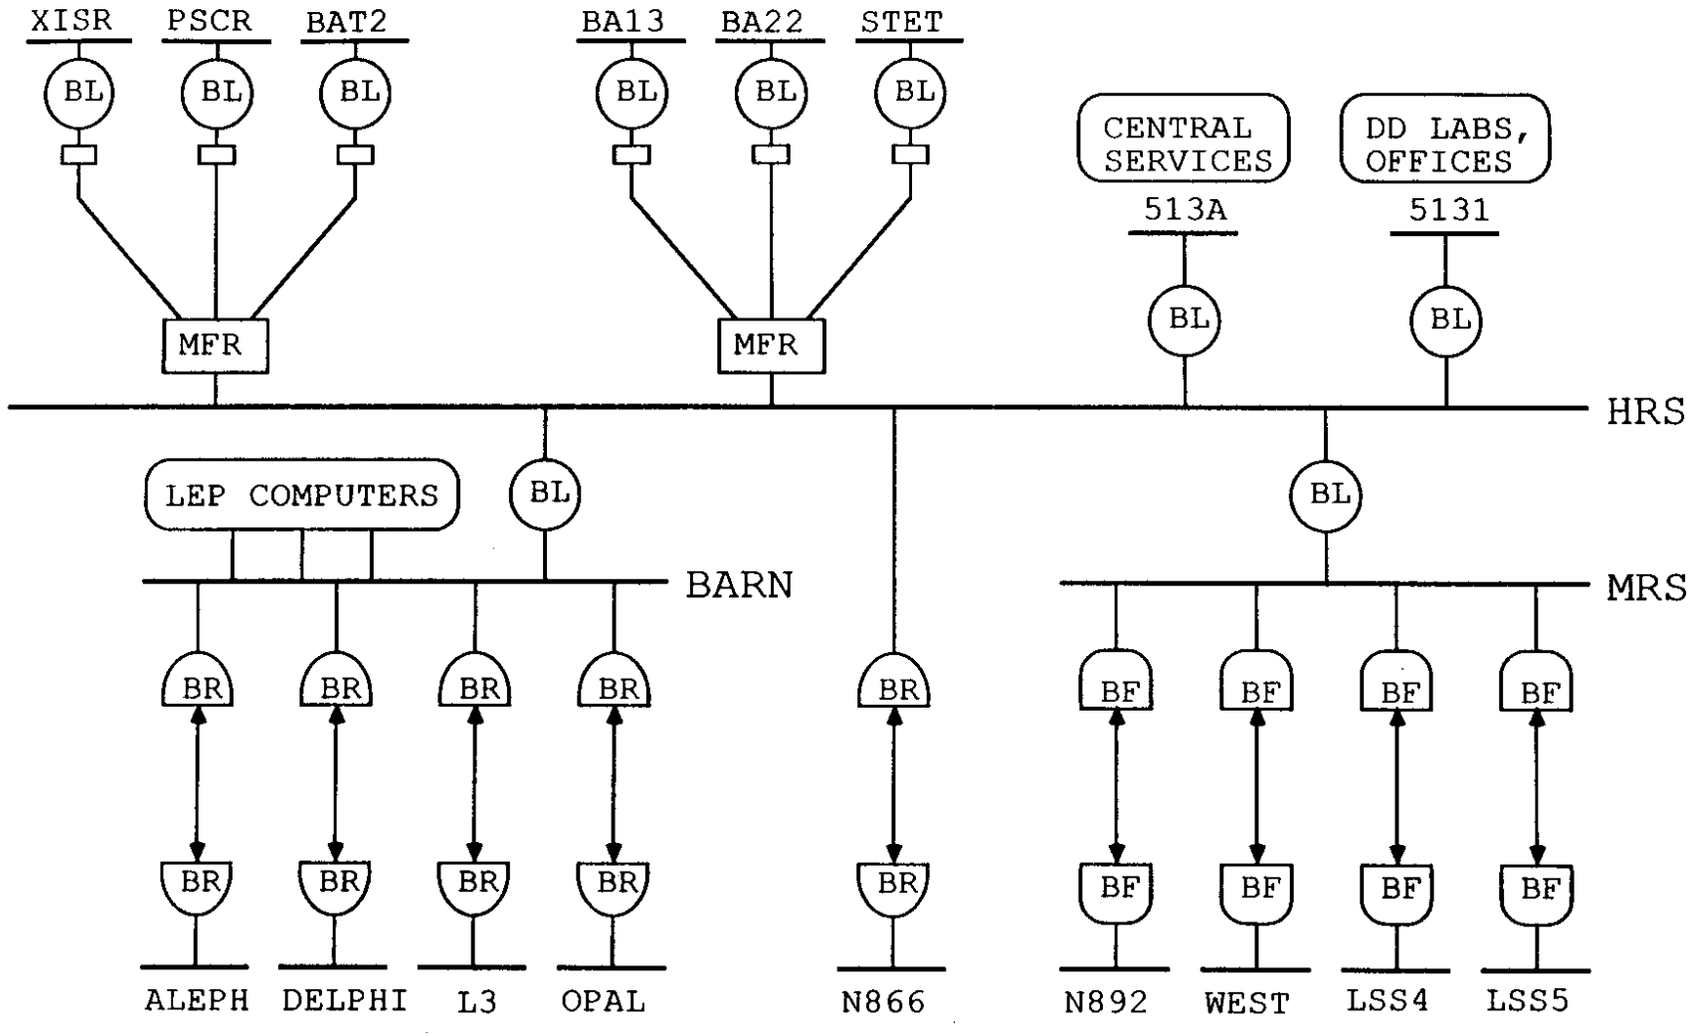

MFR: MULTI-PORT F.O. REPEATER

: FIBRE-OPTIC TRANSCEIVER

HRS: HIGH RATE SEGMENT

MRS: MEDIUM RATE SEGMENT

BR: REMOTE BRIDGE (2 Mbit/s TDM)

BL: LOCAL BRIDGE (10 Mbit/s)

BF: BACK-TO-BACK FRIGATE

FIGURE 2: PROJECTED TOPOLOGY END 1988