ECMA
EUROPEAN COMPUTER MANUFACTURERS ASSOCIATION

LOCAL AREA NETWORKS LAYERS 1 TO 4 ARCHITECTURE AND PROTOCOLS

TR/14

September 1982

Free copies of this document are available from ECMA,
European Computer Manufacturers Association
114 Rue du Rhône – 1204 Geneva (Switzerland)

ECMA

EUROPEAN COMPUTER MANUFACTURERS ASSOCIATION

LOCAL AREA NETWORKS LAYERS 1 TO 4 ARCHITECTURE AND PROTOCOLS

TR/14

September 1982

TABLE OF CONTENTS

1. GENERAL

1.1 Scope

For the purpose of compatible interconnection of data processing equipment, this Technical Report:

- describes an architecture for the Transport, Network, Link and Physical Layers for the direct interconnection of Local Area Network (LAN) (see 2);
- describes the protocols and media which are required to be used in these layers for simple basic data transfer within the boundaries of a LAN subnetwork (see 3).

Inter-layer interfaces, or any other interfaces over which the interconnection is accessed or provided, are outside the scope of this Technical Report.

This Technical Report provides a consistent technical basis for future Local Area Network standards with extended scope. Appendix C lists some candidate future extensions.

1.2 References

ISO 7498	Data Processing - Open System Interconnection - Basic Reference Model
ECMA-72	Transport Protocol
ECMA-82	Local Area Networks (CSMA/CD Baseband) Link Layer
ECMA-81	Local Area Networks (CSMA/CD Baseband) Physical Layer
ECMA-80	Local Area Networks (CSMA/CD Baseband) Coaxial Cable System

1.3 Terminology

1.3.1 Definitions

The following terms are used in this Technical Report with the meaning defined in ISO 7498.

Access point
Address
Back-pressure
Blocking
Connectionless-data-transmission
Connection-oriented
Link; Link Layer
Flow control
Layer
Link-service-access-point
Multicast addressing
Multiplexing
Network Layer
Network-service-access-point
Network-service-data-unit

Network-protocol-data-unit
Open-system-interconnection
Physical medium
Physical Layer
Protocol
Protocol-data-unit
Relay
Routing
Segmenting
Sublayer
Transport Layer
Transport-service-access-point
Transport-service-data-unit
Transport-protocol-data-unit

1.3.2 The following additional terms are defined for the purpose of this Technical Report:

1.3.2.1 <u>Datagram</u>

A connectionless-data-transmission network-protocol-data unit.

1.3.2.2 <u>Internet gateway</u>

A relay between subnetworks.

1.3.2.3 <u>Local Area Network</u>

A data communications system which allows a number of independent devices to communicate with each other, in which the communication is usually confined to a geographic area of moderate size (see also Appendix A).

1.3.2.4 <u>Subnetwork</u>

A distinct and autonomous communications subsystem at or below layer 3a.

1.3.2.5 <u>Wide area network</u>

See Appendix A.

1.3.3 The following acronyms are used in this Technical Report:

ING:	Internet gateway
LAN:	Local Area Network
LSAP:	Link-service-access-point
NSAP:	Network-service-access-point
NSDU:	Network-service-data-unit
OSI:	Open Systems Interconnection
PDU:	Protocol-data-unit
SN:	Subnetwork
TSAP:	Transport-service-access-point
TSDU:	Transport-service-data-unit
TPDU:	Transport-protocol-data-unit
WAN:	Wide area network

2. ARCHITECTURE

2.1 General

This is a basic and implementation oriented LAN transport architecture, which provides a framework within which defined protocols operate. This does not preclude the possibility of other richer or more abstract architectures.

The LAN transport architecture defined has four layers: Transport Layer, Network Layer, Link Layer and Physical Layer. These correspond to the layers of the same name in the Open System Interconnection Reference Model, ISO 7498. The corresponding layering concepts and terminology apply. The layers above the Transport Layer are outside the scope of this architecture.

The architecture includes the interconnection of subnetworks via Network Layer relays (also referred to as internet gateways). This is illustrated in Fig. 1. Only relays between LAN subnetworks are currently defined; relays between LAN and WAN subnetworks are for future study. Relays between WAN subnetworks are outside the scope.

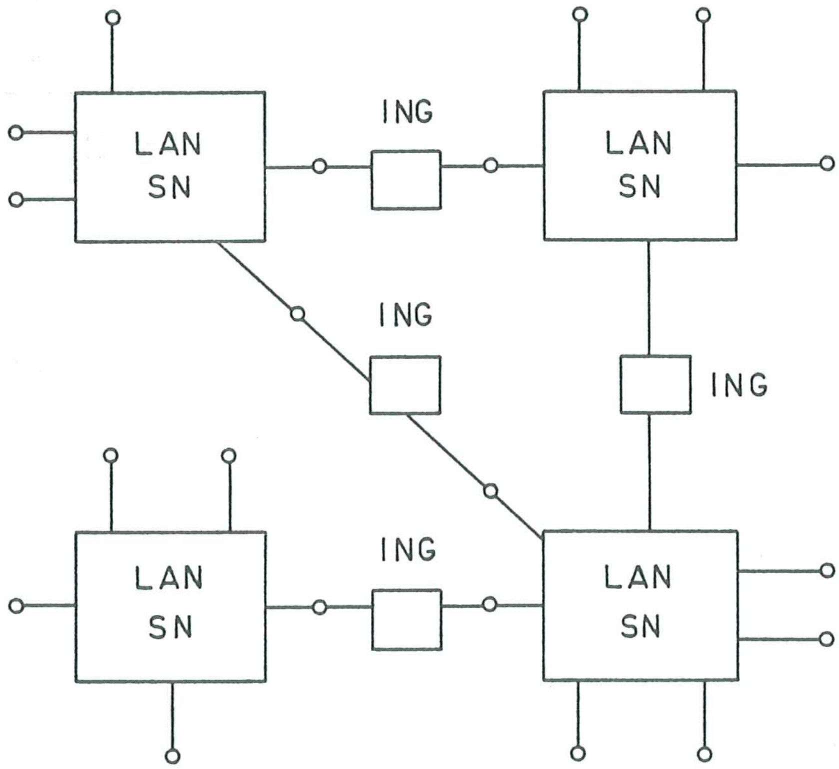

LANSN = Local Area Network Subnetwork

ING = Internet gateway

o = SN physical access point

Fig. 1 - Interconnected subnetworks

Current work within ECMA, ISO and the CCITT defines for the purpose of interconnecting subnetworks, a sublayering within the Network Layer. This LAN architecture includes only the internetwork sublayer, which is referred to as sublayer 3c. The applicability of other sublayers to the special circumstances of LANs is for further study. Fig. 2 illustrates layered communication between TSAPs over two LAN subnetworks. For communication within a single subnetwork, the layering is the same, but there is no internet gateway.

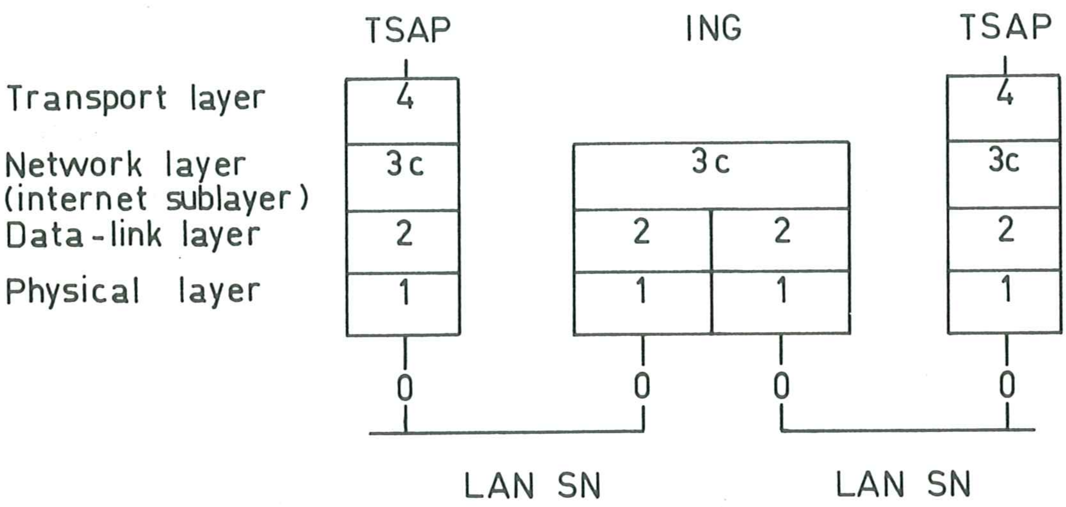

TSAP : Transport service access point

ING : Internet gateway

LAN SN: Local-Area-Network subnetwork physical medium

0 : Subnetwork physical access point

Fig. 2 - Layered model with internet gateway

The architecture is intended to be general to all types of Physical, Link, Network and Transport Layer characteristics. This initial version is specific to packet broadcast media. e.g. as in standards ECMA-80, ECMA-81 and ECMA-82. Furthermore, as indicated in the next section, this version is specific to:

- connectionless Link Layer operation,
- connectionless Network Layer operation,
- connection-oriented Transport Layer operation.

Other types of operation in these layers are for future study.

Specific provisions for voice and image transfer are for future study.

Management functions within this architecture are for future study.

2.2 Layer Functions

The Physical Layer provides mechanical, electrical, functional and procedural means for bit transmission on the physical medium of a LAN subnetwork.

The Link Layer provides the functional and procedural means for connectionless-data-transmission between access points (LSAPs) on the same LAN subnetwork.

The Network Layer provides the functional and procedural means for connectionless-data-transmission between access points (NSAPs) which are not necessarily on the same LAN subnetwork. The mechanisms used to provide this service are an inter-network datagram and associated routing functions, both positioned in sublayer 3c.

The Transport Layer provides the functional and procedural means to establish, maintain, use and terminate transport-connections between transport-service-access-points (TSAPs) in equipment attached to LANs.

2.3 Addressing

Addressing occurs in the protocols of three layers:

- Transport Layer protocol,
- Network Layer protocol,
- Link Layer protocol.

These addresses are visible in the protocols transmitted on the LAN media. Any difference of internally visible addresses within equipment attached to the LAN are outside the scope of this Technical Report.

The Link Layer addresses are used to identify and distinguish between the access-points (LSAPs) on a particular LAN. Multicase addressing is included.

The Network Layer addresses are used to identify and distinguish between access-points (NSAPs) within a site network consisting of one or more LAN subnetworks, interconnected via internet gateways.

The Transport Layer addresses are used to make further distinctions relating to, or within, the equipment accessing the LAN(s).

The addressing in each of the three layers generally include both source and destination addresses. The protocol defined for any layer may encode a null address value where the address is implicit from the operation of the lower layer(s).

The addresses used in any layer may distinguish which of several access points are used by the higher layer. This may be used to support use of different protocols at the higher layer.

The details of LAN address structures are for future study.

2.4 Routing

The equipment which is the source of a communication inserts the appropriate addresses into the protocol structure in ways defined by the applicable protocol standards.

The Link Layer destination address routes the communication to a destination equipment (DTE) on the same LAN subnetwork. By inspection of the Network Layer destination address, the equipment receiving a communication may determine whether or not it is the ultimate destination. If it is not the ultimate destination and if the equipment supports the necessary internet gateway functions, it decides where to forward the communication to next, inserts the appropriate new Link Layer header and trailer and addresses, and transmits the communication, otherwise unaltered, on the appropriate subnetwork.

When the communication reaches its ultimate destination as indicated by its Network Layer destination address, it is routed within the equipment according to its Transport Layer destination address and/or control information.

If the ultimate source and destination are on the same LAN, transmission is normally direct, and no internet gateway is involved.

2.5 Multiplexing

There may be multiple concurrent transport connections at a NSAP. This is supported by Transport Layer multiplexing mechanisms, with corresponding control information in transport protocol to distinguish between multiple connections at the same address.

For connectionless-data-transmission of the underlying layers there is no need for any explicit concept of multiplexing in these layers. Successive connectionless-data-transmissions have no formal logical relationship to one another, and multiplexing effects are always an inherent capability.

2.6 Flow Control

The general strategy for flow-control within this architecture is to queue at source (i.e. at sending transport entities). Data is generally only sent on a transport connection after the destination has signalled its readiness to receive. No more data is then sent on the connection (except for control and error recovery purposes) until the destination has again signalled its readiness to receive. Generally, each transport connection needs only a small share of the total available bandwidth (i.e. there are relatively long periods of inactivity between its transmissions), and the effect of the additional round-trip delays for flow control purposes is negligible. The opposite circumstances generally prevail on WANs, where it is typically necessary to pipeline messages into the relatively long transit delay, in order to achieve the required throughout and line utilization.

This Transport Layer strategy avoids congestion due to messages in the same connection catching up with each other. For connetionless-data-transmission in the Network Layer, there may be timers which regulate the rate at which any transport source can inject internet datagrams into the network. Other types of congestion (e.g. Link and Physical Layer) may be avoided by providing sufficient buffering, particularly in internet gateways. Additional techniques for congestion control are for future study.

Transport Layer congestion, or back-pressure at the receiving end of data transfers, results in the sending of flow-control messages which create back-pressure at the sending end (when they arrive there successfully). There may also be local back-pressure on the receiving Network Layer. See Fig. 3.

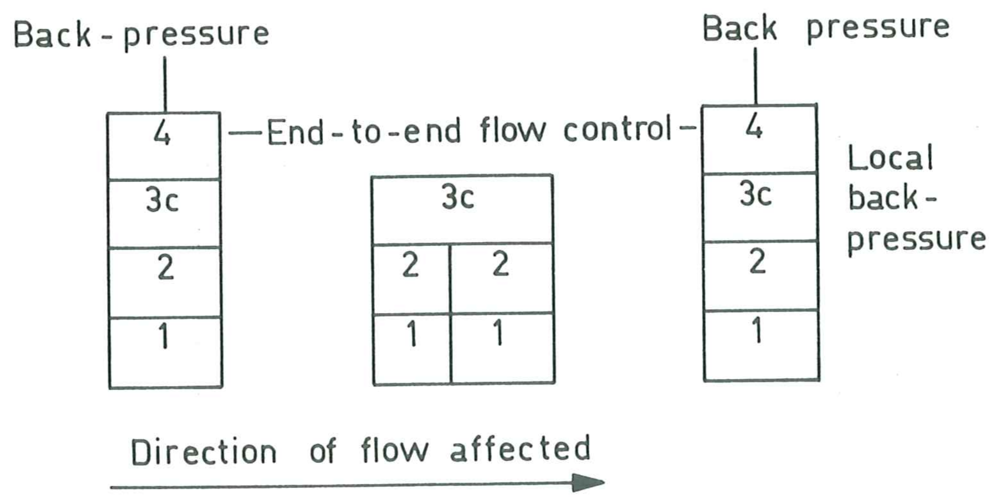

Fig. 3 - Back-pressure resulting from transport destination congestion

Network Layer congestion at gateways or final destinations does not directly result in any back-pressure in the connectionless-data-transmission case. Instead, internet datagrams are discarded at the point of congestion. See Fig. 4. This is a simple and reliable way of ensuring deadlock-free network operation.

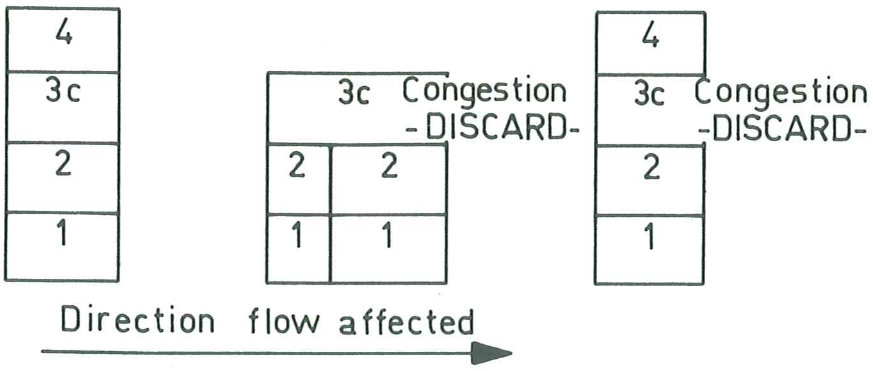

Fig. 4 - Discard resulting from network buffer congestion
(connectionless-data-transmission case)

Link Layer congestion usually results in back-pressure at the LSAP affected, with corresponding back-pressure at the NSAP(s) and TSAP(s) operating via it. Physical Layer congestion has essentially the same effects; both are media congestion. See Fig. 5.

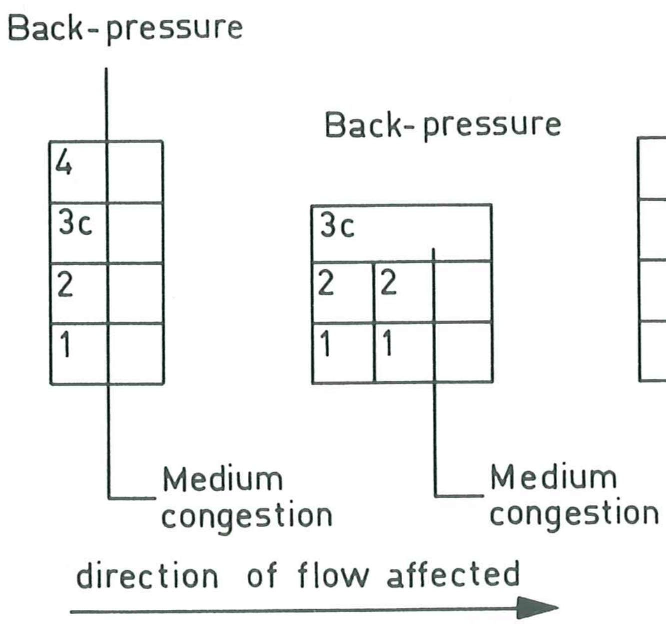

Fig. 5 - Back-pressure resulting from media congestion

2.7 Reliability

The architecture maximizes LAN transport systems reliability. This has three main elements:

- resilience,
- simplicity,
- integrity.

The resilience is maximized by the Transport Layer having its own self-contained and comprehensive error controls, with their locus in the two end-points only. This has least vulnerability to failures. A transport connection can also survive the disruption which is generally inherent in dynamic re-configuration of the underlying networks. This potential for re-configuration is another source of resilience and aids maintainability.

The simplicity of the connectionless operation of the lower layers is enhanced by them being allowed occasionally to lose, duplicate or mis-order data transmission (simplified error management, congestion control, buffer management and reconfiguration). The compensating sophistication of the Transport Layer to detect and recover from these errors is in any case generally desirable for resilience reasons. The queue-at-source flow control strategy allows simplification of transport protocol implementation. The net outcome of these simplifications is that all four layers involve less software and less state information (which are the principal sources of error with intrinsically reliable LAN media).

End-to-end data integrity is provided by the error controls in the transport layer protocol.

2.8 Protocol Structure

The architectural layering is manifested in nested layers of protocol, as illustrated in Fig. 6. For any protocol-data-unit (PDU) at a higher layer, there is always the full sequence of enclosing lower layer headers.

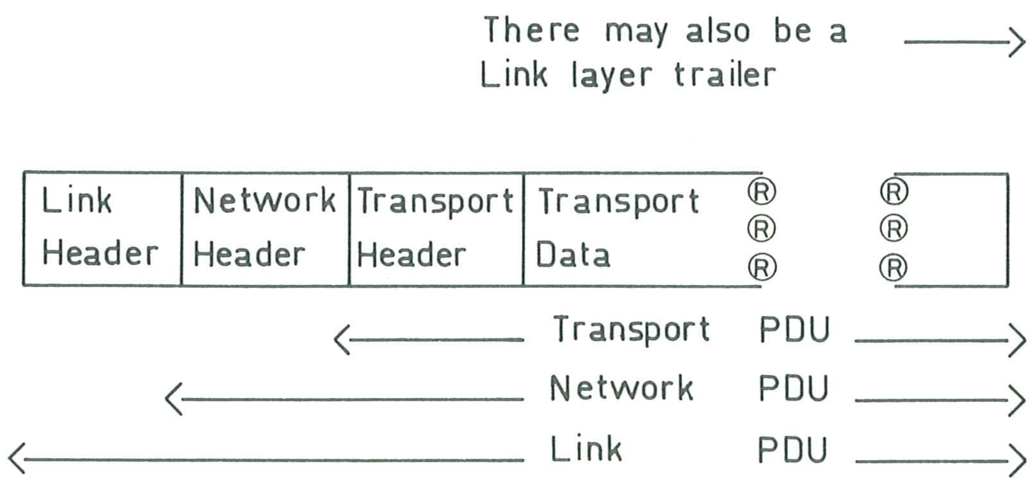

Fig. 6 - Nested protocol structure

There is data segmenting in the Transport Layer: a transport-service-data-unit (TSDU) may be subdivided into a series of smaller transport-protocol-data-units (TPDUs).

In this Technical Report, there is no data segmenting in the Network Layer or in the Link Layer. The need for segmenting in these lower layers, and the special problems of segmented connectionless-data-transmission, are for future study.

The concept of data segmenting is not applicable to the Physical Layer.

There is no data blocking in any of the four layers.

Protocols of any of the four layers may concatenate data and control information for transmission.

3. PROTOCOLS

3.1 General

Within this architecture, it is logically possible to use a wide variety of protocols. (Different LAN media generally have different Physical and Link Layer protocol characteristics. The protocols in the Network and Transport Layer tend to have a wide variety of options, as do some of the lower layer protocols). This variability is a potential source of incompatibility where interworking is required between equipment from different suppliers. To aid compatibility, this document defines "protocol sets".

NOTE 1

In this version of the Technical Report there is only one; which is named: "protocol set 1".

For each protocol set, the following shall be specified:

1. Purpose,
2. Transport protocol,
3. Network protocol,
4. Link protocol,
5. Physical protocol,
6. Physical medium,
7. Addressing.

Each protocol is usually defined by reference to a standard. Where that standard allows selection of classes, options subsets or other variability, a particular selection shall be specified. A particular way of use may be recommended.

A consistent method of addressing is generally defined for all four layers.

The forward compatibility statements give an indication of whether the protocol set will be carried forward unchanged into future versions of this Technical Report. It also provides guidance where future changes are anticipated. It represents ECMA intentions at the time of issuing this Technical Report, but is not a binding committment for future standardization.

3.2 Protocol Set 1

3.2.1 Purpose

Protocol set 1 shall be the basic set of protocols which is practicable to define for this version of this Technical Report. It shall provide for interworking between equipment from different suppliers attached to the same CSMA/CD baseband LAN subnetwork.

3.2.2 Transport protocol

The transport protocol shall be class 4 protocol defined in Standard ECMA-72 with the following mandatory restriction:

- The maximum TPDU size shall not exceed 527 octets (see note 2).

The following simplified use of the protocol is recommended but not mandatory, and shall only be implemented in ways that do not violate conformance with Standard ECMA-72:

- receiver limits maximum credit to 1 (see note 3),
- sender does not concatenate TPDUs,
- negotiate non use of the checksum option.

NOTE 2

The lower layers are required to have corresponding maximum PDU size capability (see 3.2.4). This avoids data segmenting in the lower layers.

NOTE 3

No pipeline flow control. Simplified buffer management and error recovery. It is for further study whether this would be suitable for communication over multiple subnetworks.

3.2.3 Network protocol

The Network Layer header shall be encoded as one octet, value zero. This shall be decoded as length indication field, defining the absence of further header information in this layer.

NOTE 4

It is anticipated that in the definition of a future standard for the interconnection of subnetworks, the first field of the Network Layer header would be a header length indicator.

3.2.4 Link-protocol

The Link Layer protocol shall be that defined in Standard

ECMA-82. The implementation shall be capable of receiving link-protocol-data-units with a data size of up to 528 octets (see note 5).

NOTE 5

527 octet NSDU and 1 octet Network Layer header = 528 octets.

3.2.5 Physical protocol

The Physical Layer protocol shall be that defined in Standard ECMA-81.

3.2.6 Physical medium

The physical medium of interconnection shall be that defined in Standard ECMA-80.

3.2.7 Addressing

Transport Layer addressing is not defined by this Technical Report; transport address may have any value permitted by ECMA-72.

There is no Network Layer addressing (see 3.2.3).

Link Layer addressing shall be unique within the LAN subnetwork. This does not preclude uniqueness within some wider scope. Link Layer addresses shall have any value permitted by ECMA-82.

APPENDIX A

LOCAL AREA NETWORK OVERVIEW

A Local Area Network is a data communications system which allows a number of independent devices to communicate with each other. A Local Area Network is distinguished from other types of data networks in that the communication is usually confined to a geographic area of moderate size such as a single office building, a warehouse or a campus. The network can depend on a communication channel of moderate to high data rate which has a consistently low error rate. The network is typically controlled by a single organization. This is in contrast to wide area networks which interconnect facilities in different parts of a country or are used as a public utility.

The applications environment for the Local Area Network is intended to be commercial and light industrial.

The Local Area Network is intended to have wide applicability in many environments. To this end, the Local Area Network may support applications such as:

- File transfer and access protocols
- Graphical applications
- Word processing
- Electronic mail
- Remote data base access

The Local Area Network is intended to support various data devices such as:

- Computers
- Terminals
- Mass storage devices
- Printers/plotters
- Monitoring and control equipment
- Gateways to other networks

The above lists are intended to show typical applications and devices and, as such, are not intended to be exhaustive, nor do they constitute a set of required items.

APPENDIX B

REQUIREMENTS FOR ECMA LOCAL AREA NETWORK LAYER 1 TO 4 ARCHITECTURE AND INTERCONNECTION PROTOCOLS

These are the requirements to which this Technical Report is developed.

B.1 OBJECTIVE

The objective is to provide, with the utmost urgency, a basis for an ECMA Standard which defines for Local Area Network (LANs):

1. some initial means of compatible interconnection in layers 1 to 4 of the Reference Model,
2. a consistent technical basis for future more comprehensive standardization of these layers of LANs.

B.2 PURPOSE

The purpose is two-fold:

1. to provide a strong focus around which ISO and other standardization bodies may develop the ultimate, comprehensive and definitive standards,
2. to provide a simple starting point for users and manufacturers to implement LANs, with some confidence that this is likely to be consistent with the technical content of the above ultimate standards.

B.3 CONSIDERATIONS

B.3.1 OSI

The layering concepts and terminology defined for Open System Interconnection are to be used.

B.3.2 State of the art

The techniques used should be as far as possible well proven. Existing generally available specification and standards should be applied.

B.3.3 Simplicity

This is absolutely essential.

B.3.4 No options

To maximize compatibility and simplicity, the initial standard should define one set of interconnection rules, with no variants or options.

B.3.5 <u>Restricted scope</u>

Considering B.3.2, B.3.3 and B.3.4 and to allow rapid development of the standard, some restrictions of scope have been necessary (aim to meet some of the needs now without delay, rather than more needs with more delay). This consideration should also be used to avoid areas in which agreement cannot be reached quickly.

B.3.6 <u>Future work</u>

In parallel with the publication of the first version of this Technical Report, items left for future study should be developed for inclusion in a second version at the earliest possible date.

APPENDIX C

CANDIDATE ITEMS FOR FUTURE EXTENSIONS

The following items are identified as possible candidates for future extensions:

- Inclusion of LAN/WAN gateways into the architecture (2.1).
- Possible role of sublayers 3a and 3b in LANs (2.1).
- Inclusion of circuit-switched methods into the architecture (2.1).
- Inclusion of connection-oriented store and forward internet gateways into the architecture (2.1).
- Inclusion of connectionless-data-transmission and possibly other non-connection-oriented modes of Transport Layer operation into the architecture (2.1).
- Inclusion of specific provision for voice and image transmission into the architecture (2.1).
- Inclusion of management functions into the architecture (2.1).
- Inclusion of detailed addressing architecture for all four layers (2.3).
- Inclusion of additional congestion control provisions into the architecture (2.6).
- Inclusion of Network Layer and Link Layer data segmenting into the architecture (2.8).
- Inclusion of further protocol sets (3).
- Inclusion of connection-oriented data transmission.
- Conformance specification and testing.

Other future extensions are not precluded.