ATNP WG3/SG3 (Upper Layer Architecture)

Draft Upper Layer Guidance Material for CNS/ATM-1 Package

(Presented by WG3/SG3)

REVISION STATUS

Version 1.0 First draft Guidance Material for Upper Layers (Gold Coast)

TABLE OF CONTENTS

1. INTRODUCTION 1

1.1 SCOPE OF DOCUMENT 1
1.2 OBJECTIVES 1
1.2.1 The ATN Upper Layer Environment 2

2. ARCHITECTURE 2

2.1 SCOPE 3
2.2 THE BASIC ARCHITECTURAL CONCEPT 3
2.2.1 Scope 3
2.2.2 Background 3
2.3 BENEFITS OF THIS APPROACH 4
2.4 UPPER LAYER CONCEPTS AND STRUCTURE 5
2.4.1 Scope 5
2.4.2 Functions of the Upper Layers 6
2.4.3 Upper Layer Structure 7
2.5 UPPER LAYER SERVICES AND PROTOCOLS 9
2.6 ENCODING RULES AND DATA COMPRESSION 10
2.6.1 Scope 10
2.6.2 Standard Encoding Rules 10
2.7 NAMING, ADDRESSING AND REGISTRATION 12
2.7.1 Scope 12
2.7.2 Naming and Addressing Guidance 12
2.7.3 Naming 13
2.7.4 Addressing 15
2.7.5 Name - Address Mapping 16
2.7.6 Selectors in the ATN 16
2.7.7 Registration Issues 17

3. IMPLEMENTATION ISSUES 18

3.1 INTRODUCTION AND SCOPE 18
3.2 EVALUATION OF OSI UPPER LAYER PROTOCOL OVERHEADS 18
3.2.1 Basis of Overhead Analysis 18
3.2.2 Basic Assumptions 19
3.2.3 Encoding Options 19
3.2.4 Name and Address Options 19
3.2.5 Presentation Context 20
3.2.6 User Data 20

4. CNS/ATM-1 PACKAGE GUIDANCE MATERIAL 20

4.1 INTRODUCTION 20
4.1.1 Motivation of the Work 20
4.1.2 Architectural Guidance Material 25
4.2 DIALOGUE SERVICE 27
4.2.1 Introduction 27
4.2.2 Connection Mode 27
4.2.3 D-ABORT 27
4.2.4 Security 28
4.2.5 Mapping to transport 28
4.3 CONTROL FUNCTION (AIR/GROUND) 28
4.4 SESSION 31
4.4.1 Session Defect Report 31
4.5 PRESENTATION 31
4.5.1 Presentation Defect Report 31
4.5.2 Presentation Fast Byte Guidance Material (OSIEFF) 31

4.6 ACSE 34
4.6.1 Discussion of differences in ACSE editions *34*
4.6.2 ACSE Primitive Flow Diagrams *34*
4.7 NAMING AND ADDRESSING 39
4.7.1 Implementation of ULA Construction of Titles and Addresses *39*
4.7.2 Guidance on CNS/ATM-1 Naming and Addressing *39*

5. SARPS DEFECT REGISTER 40

1. Introduction

1.1 Scope of document

This document is intended to describe the architectural framework for the standardisation of ATN Upper Layers. It covers the following areas:

- Supporting upper layer stacks
- Upper layer overhead comparisons
- Encoding rules and data compression, for upper layer PCI
- Naming, addressing and registration
- Application layer structure and service elements
- The use of transport services

The following sections draw together existing material on each of these areas, as well as providing additional summaries and analysis not found elsewhere.

1.2 Objectives

The aim of this document is to define the general communications architecture for ATN upper layer(s) (i.e. everything above the ATN Transport Service) and to provide reference material to aid the development and implementation of the upper layers.

The basic aim of this document is to define a set of architectural principles which will allow ATN Applications to be constructed in a standard way. This "building block" approach has many well-known advantages, including:

- the duplication of effort associated with designing and debugging similar functionality for many different application types is minimised
- the same type of design problem would otherwise have to be repeatedly solved for each new application
- the productivity of designers, programmers, system engineers and testers is increased, as they only have to deal with a single architecture
- the certification effort is eased, as experience is gained with previously accredited modules.

This work should stabilise and document the architectural basis for standardisation in ICAO, in particular considering:

- The approach to upper layer stack selection

- Preferred data encoding schemes
- Naming and Addressing
- Application Layer Structure

1.2.1 The ATN Upper Layer Environment

The OSI/ATN environment as it relates to upper layers and ATN Applications is depicted in Figure 1.1. This architecture effectively embraces two areas of standardisation:

a) Communications profiles of the upper layers to deliver the communications requirements for proposed applications; such profiles should be based on existing or new (defined by the ATN community) standards.

b) Application specifications, defining the messages and message sequence rules for applications to meet specific operational requirements, as set down in a completed application SARPs.

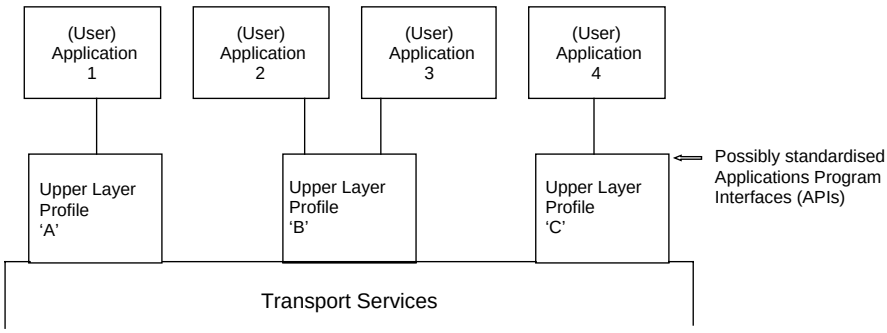

Figure 1.1 : ATN Host Software Environment

In OSI terms, the DLA as defined above will comprise the LOCAL part of an OSI Application Process (AP). The OSI AP will also include an Application Entity, comprising the Application Service Elements and Application Service Objects (ASEs and ASOs - see 2.3.4) appropriate to the specified profile.

Each DLA Application Process shall be allocated a name by which the particular application, version and message structure can be recognised, and which can be used to indicate the application support capability of a particular aircraft or ground station. The name may be used by the Context Management Application (CMA).

2. Architecture

2.1 Scope

This section examines the basic ATN Application architecture to provide a framework in which the upper layer profiles are defined. The services provided by upper layers are described for reference.

2.2 The Basic Architectural Concept

2.2.1 Scope

This section proposes generic communication services for use on the ATN. It describes the proposed services and suggests existing application layer standards which may be used to implement them. It also examines the application of the generic communication services to the ATN and considers some of the special constraints imposed on the ATN by the communications subnetworks.

2.2.2 Background

The framework for the standardisation of the upper layers puts forward the proposal that a set of standardised common communication services should be provided on which user applications, or so called Data Link Applications, would be constructed. These communication services will be used by the user applications to exchange information and, where thought appropriate, the interactions (i.e. message exchanges) which take place over these services would also be standardised in terms of the functions offered by the specific service.

The adoption of this framework means that the standardisation process may be subdivided into two parts; firstly, the standardisation of a common set of communication services, and secondly the standardisation of the DLA which uses those services. The aims of this framework are:

a) to keep to a minimum the number of standard application services

b) to use existing OSI application layer standards wherever possible, thus removing the need to define, standardise and conformance test new application layer standards

c) to tailor some of the service profiles to the underlying restrictions of some of the low bandwidth air-ground subnetworks in the ATN, but to use recognised application profiles wherever possible.

It is planned to define a limited set of communication service profiles for use in the aeronautical domain, to provide applications with access to the ATN. Each profile constitutes an upper layer "protocol stack" definition which when implemented provides the appropriate functionality in the selected upper layers.

Profiles are defined by selecting valid combinations of protocol standards and forming valid subsets in such a way as to deliver a specific level of service to the applications. A number of such profiles will be defined to provide wide applicability for the differing

upper layer support requirements of different applications. Collectively, these profiles will provide a well-defined set of services which can be utilised when designing and implementing particular ATN applications. This does not imply that it is necessary or even desirable to implement the complete set of selected upper layer profiles on all end-systems. Subsets of the full set can be selected, to provide appropriate levels of functionality to meet the requirements of different classes of applications.

Inherent in this discussion is the desire to use standardised protocols whenever possible. Many of the necessary protocol standards already exist, and profiling is already being undertaken to produce International Standardized Profiles (ISPs). However, functionality or performance requirements may exist which are not satisfied by existing upper layer protocols. In such circumstances, it may be necessary to develop specialised upper layer protocol definitions, within the framework of the OSI reference model, for use in the aeronautical domain.

The adopted framework separates the communications profiling from the application standardisation and tries to standardise at the communications level only a small number of generic communications classes, which would be appropriate for use by a wide range of applications.

Each DLA will be defined by a SARPs document. The specification document will include complete message definitions, including encoding rules, sequencing rules, exception conditions and temporal relationships to be met by the application. It will also specify what Quality of Service (service characteristics) are required from the underlying communications, so as to allow selection of the appropriate ATN UL profile and Data Link.

2.3 Benefits of this approach

The following benefits result from this approach:

a) It allows the separation of application specific and communication specific functions.

b) The certification of communications and applications software may be carried out separately.

c) It allows the use of a small number of standard communication services, based on distinct modes of interaction, by a large number of DLAs.

d) It does not require the definition of new application contexts, presentation transfer syntaxes,... whenever a new DLA is introduced or an existing DLA modified.

e) Some DLA communication requirements may be satisfied by existing ASEs and profiles.

f) The communication services may in some cases be based on COTS (Commercial Off The Shelf) products.

g) It does not require the implementation of application specific interactions (eg timeouts, message sequencing rules) within the communication service.

In the ATN Protocol Architecture, OSI application entities provide communications services to ATN applications. The service boundary between the application entities and the ATN applications is an abstract interface which may or may not be realised as an exposed interface in a real implementation. ATN profiles are defined to lie on the service-provider side of this boundary (see figure 2.1).

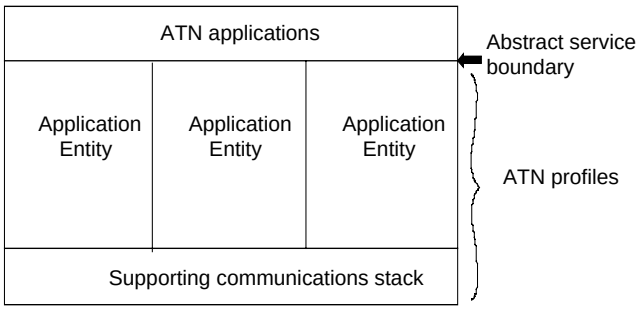

Figure 2.1 : Position of ATN Profiles within an End System

ATN applications will be defined to use the services of the selected ATN profiles. The communication aspects of these applications will be defined in terms of:

a) the semantics and structure of the information to be exchanged (‘messages’)

b) the rules governing the dialogue between parties (message sequencing)

c) requirements imposed on the underlying communications services (quality of service, etc.)

ATN Applications will be standardised as ICAO SARPs which will specify which subset of the ATN profiles is appropriate for the particular application, and if more than one subset is possible, under which circumstances one or the other would be selected. These choices map onto upper layer stack selections for the ATN application.

2.4 Upper Layer Concepts and Structure

2.4.1 Scope

This section deals with the upper layer model. It covers the existing ASE structure and also considers the applicability of the revised OSI Application Layer Structure (ALS) model in the 2nd edition of ISO/IEC 9545 (also known as Extended Application Layer Structure or XALS) in terms of ASOs and associated Control Functions.

2.4.2 Functions of the Upper Layers

The service currently offered by the ATN Internet is a low-level communications service which corresponds to the OSI Transport Service defined in ISO 8072. Although it is possible to construct ATN Applications which make direct use the raw transport service, such applications will not benefit from the "common building block" approach.

It is therefore envisaged that a set of functions which add value to the raw ATN transport service will be standardised, and these will provide high-level services to the specific ATN Applications. Since it is a fundamental principle of the ATN that it adopts the protocols defined in the international standards for Open Systems Interconnection (OSI), it is logical to look to the standards defining the upper layers of the OSI model to provide the required value-added functions.

Before examining further alternatives, this section considers what the standard OSI upper layer (Session, Presentation and Application) protocols offer as added value on top of the raw transport service. It is these features that need to be incorporated in any ATN Application, or rejected as being unnecessary for a particular requirement.

One important, static function is to define formats and encodings for data interchange, in an unambiguous and open way, i.e. independent of any particular hardware bit-ordering or word-size conventions.

The features listed in the following table are taken from a paper by Peter Furniss. They deliberately use inexact language and (sometimes) avoid technical terms as it is the effective role, not the specified function that is important:

	Upper Layer Feature	Comment
a)	The first octet of each session protocol data unit (PDU) identifies what part of the conversation is being sent	Something will <u>always</u> be needed to distinguish broad categories of message that all travel on the same supporting service (in this case, on T-DATA)
b)	Session and presentation connects include selectors that can be used for upper-layer routing.	The upper-layer selectors are redundant if the upper-layer implementation is integrated
c)	Presentation layer is primarily concerned with presentation context - the identification and negotiation of the abstract and transfer syntax of the application data. There are several aspects to this:	
	i) the presentation context identification of each piece of application data allows different message types (APDUs) to be distinguished, without fear of ambiguity, even if two messages from different sets (abstract syntaxes) happen to have the same bit-pattern when encoded.	This is only necessary when there are multiple abstract syntaxes (but, since ACSE is considered to be in the application layer, there always are).

	ii) the presentation context identification can (usually) be used to distinguish separate streams of dialogue - because each stream will usually be using a separate set of message types	This works provided all the ASEs involved are different, but becomes rather complicated otherwise. Something that is not concerned with message type identification is needed for the general case of identifying the appropriate recipient of a message. Assuming that the message type tells you who it is for does not always work (try sending an anonymous birthday card to twins!)
	iii) abstract syntax negotiation states/agrees which sets of message types will be used on the association	See also f): the application and presentation contexts have some overlap of function. Since an application context definition identifies which ASEs, and thus which sets of message types, will be involved, abstract syntax negotiation can be no more than a parameterisation of the application context, and will usually be unnecessary.
	iv) transfer syntax negotiation (for each particular presentation context) allows alternative representations of the same message to be offered	This is probably the most useful feature of the upper-layers, and the least used at present. It offers considerable scope for extensible, conformant optimisation in the processing OR the transmission of application protocols.
d)	Presentation provides bracketing around the encodings of application data, so they need not contain their own length fields	Anything whose encoding is not self-delimiting will require this feature. Conversely, if the supporting layers provide the feature, it need not be duplicated in the application protocol (i.e. the application protocol need not be self-delimiting)
e)	ACSE provides naming fields (AP-title, AE-qualifier and invocation identifiers) that can be location-independent	The ability to carry the names of the communicating entities explicitly, rather than inferring them from the addressing (via reverse lookup) is likely to be needed generally, but cannot always achieve what is hoped for. It cannot be relied on for effective authentication of the sender, though it can give a weak identification. Depending on how fast directory updates take place, it may allow a subsequent return call to the same entity on a different lower layer address. But it certainly does allow an un-authenticated assertion of who the other party is.
f)	ACSE carries an application context identifier that identifies what the communication will be about, and what should be assumed about the intents and capabilities of the partners.	In a particular instance, an application context will either be completely redundant or absolutely essential! Often the entity that is listening on a particular lower layer address will only support one application context. If a call is made to that address, the only advantage in communicating the application context explicitly is to identify "wrong numbers". In other cases, the entity may be capable of several different things, and the application context is needed to tell it what the call is about.
g)	ACSE identifies which of its APDUs is being used (they all have different tags)	Since each ACSE APDU has (at present) a unique mapping to Session, this is redundant. However, if different ACSE APDUs could travel on the same immediately supporting service, the distinction is needed - in fact it has taken on part of the role of the session SI field in a).
h)	With its recent extension, ACSE can carry security information	Where needed, and perhaps especially combined with the naming fields, carriage of security information will be vital
j)	All three protocols have facilities for version negotiation and, in the connect-request, extensibility rules that allow old implementations to ignore fields they do not recognise	Inventing a non-extensible protocol can be expected to be the last action of a doomed design team.

2.4.3 Upper Layer Structure

The fact that the upper layers (Session, Presentation and Application) are divided into three layers in the OSI reference model, and that the Application Layer is divided into discrete service elements, does not imply that this abstract division should be visible in real world implementations.

To assist in application design, a standardized application layer architecture has been developed as ISO/IEC 9545 Application Layer Structure (ALS). The architecture specifies the existence of the following major application layer components:

- *Application Process (AP)*: an element within a real Open System which performs the processing for a particular application. An example of an AP is an X.400 software package, of which some elements will be responsible for OSI communication and other elements will have responsibilities beyond the scope of the OSI environment, for example user interfaces and document filing systems. The AP can thus be considered to be partially in the OSI environment and partially in the local system environment. It may be modelled as a set of application entities (AEs).

- *Application Entity (AE)*: an aspect of an application process pertinent to OSI, and is the means by which application processes exchange information using defined application protocols and the underlying presentation service. An example of an application entity is an X.400 Message Transfer Agent (MTA) or Message Store (MS).

- *Application Service Object (ASO)*: the generic term for an object which performs some communications related task. For example, a file transfer system such as an FTAM initiator is an ASO. An AE is a particular kind of ASO, the outermost ASO. An ASO can contain further ASOs, which are thus recursively defined.

- *Application Service Element (ASE)*: An AE can be broken down further into a number of ASEs, each of which provide a set of OSI communication functions for a particular purpose. An ASE is a particular kind of ASO which is elemental and therefore cannot be subdivided. An ASE may be thought of as a leaf node in a tree of ASOs. ASEs may provide general purpose functions which can be used by a number of applications. ASEs in general are defined in separate standards. For example, the *Association Control Service Element (ACSE)* is used to create and release associations between applications, the *Remote Operation Service Element (ROSE)* is used to support general-purpose enquiry-response type interactions, and the *Reliable Transfer Service Element (RTSE)* is used to ensure that a unit of information is completely transferred once and only once between communicating applications.

- *Control Function (CF)*: this exists within an ASO to coordinate the use of the different services provided in the ASOs and ASEs, and also the use of external services such as the presentation service. It provides a mapping of the ASO to the subordinate ASOs and ASEs which it contains.

AEs, ASOs and ASEs are abstract representations of some part of an application. They cannot perform any action without first being invoked; this is equivalent to having a computer program that cannot do anything until it is run. An *AE invocation (AEI)* can be thought of as an AE that is running; similarly an *ASO-invocation (ASOI)* and *ASE-invocation (ASEI)* are ASOs and ASEs that are running.

2.5 Upper Layer Services and Protocols

Application layer standards define the services and protocols supported by ASEs. Some of these standards define services which are common to a range of application layer standards, and which can be used as "building blocks" when defining new applications. The most significant of these are summarised below; others define the specific services and protocols supported by particular application layer standards such as electronic messaging and file transfer.

Association Control Service Element (ACSE)

The ACSE service allows one application-specific AE (for example, an X.400 MTA) to request that an association be established with a remote AE (in this example, an adjacent X.400 MTA) for the purpose of information transfer. ACSE will attempt to establish the association using the supporting OSI layers, and will report the success or failure of this attempt to the requesting AE. ACSE also provides for the orderly and abrupt release of the association once established.

The connection-oriented ACSE service and protocol are defined in ISO 8649 (CCITT X.217) and ISO 8650 (CCITT X.227) respectively. The connectionless ACSE protocol is defined in ISO 10035 (CCITT X.237), in which the request to establish an association, the information to be exchanged and the request to release the association are all transferred together.

The following ACSE functional units are defined:

- Kernel
- Authentication
- Application Context Negotiation

ASO-Association Control Service Element (AACSE, or "A2CSE")

There is currently a project under way within ISO/IEC JTC 1/SC 21 to modify the ACSE standards to support the revised ALS model more directly by allowing associations between named ASOs to be set up explicitly and to perform the context demarcation functions required for ASO-associations

Remote Operations Service Element (ROSE)

ROSE provides a general-purpose 'request-response' service for use by application-specific ASEs. It is therefore used by interactive-type applications in which one ASE requests that another ASE performs an operation, with the performing ASE reporting the outcome of that operation to the requesting ASE. For example, ROSE is used by the X.400 remote User Agent (UA) to interrogate and retrieve messages from the Message Store.

The ROSE service and protocol are both defined in ISO 13712 (CCITT X.218 for the service and X.228 for the protocol). Connectionless ROSE is defined in ISO/IEC 13712-2/PDAM 1.

2.6 Encoding Rules and Data Compression

2.6.1 Scope

This section covers the encoding of information for interchange between heterogeneous systems, and possible compression mechanisms. There are three areas where data compression could be carried out:

- in application specific user data
- in the presentation layer for APDU overhead reduction
- at the link layer where standard compression techniques could be used (eg V.42bis)

Abstract Syntax Notation One (ASN.1) is a standard notation for representing data structures in an implementation-independent way[1]. The standards for ASN.1 have recently been updated. There exist a number of standardised encoding rules which enable data structures expressed in ASN.1 to be encoded as a machine-independent octet stream which is well suited to data interchange. Most current OSI upper layer protocols are specified using ASN.1 definitions, and they explicitly specify the Basic Encoding Rules (BER) for encoding their protocol data units (PDUs). The use of the Packed Encoding Rules (PER) might be more efficient but would not conform to current OSI standards. This issue needs to be investigated, with the possible outcome of requesting alternative encoding support in ISO standards.

2.6.2 Standard Encoding Rules

Before an application transmits data to a peer on another machine, the data must be encoded in some form. The data will already be stored on the local system and encoded in some local format, but this is not necessarily acceptable to the peer application. It may, therefore, have to be re-encoded before being transmitted. On receiving the data, the peer application may have to reformat it into some local form suitable for local use.

1 The fact that it is called ASN number 1 emphasises that other syntax notations are foreseen in the future.

The data transferred between the two systems uses some common encoding for transfer, which is the subject of this section.

Before the definition of the presentation layer (e.g. in X.400 1984 and in Internet applications) the encoding rules had to be built into the application itself. Thus there was no possibility of choosing a more efficient encoding without changing the definition of the application. One of the functions of the presentation layer is to negotiate the set of transfer encodings to be used.

The application need not be concerned with the encoding rules, it simply deals with information in its local syntax, which must be a realisation of the abstract syntax in which the application is defined. ASN.1 (Abstract Syntax Notation number one) is commonly used to define the syntax of information in an abstract form, divorced from the concrete syntax of the encoding rules. The ASN.1 standards describe the notation which is used in the specification of existing ASEs, and which can be used to develop new ASEs or to express application user data.

Several sets of standard encoding rules have been defined or proposed, including the following:

Basic Encoding Rules (BER)

BER is the original set of ASN.1 encoding rules, assumed by many existing OSI standards for their transfer syntax. The encoding is based on an octet-aligned TLV (type-length-value) approach. The Basic Encoding Rules provide a number of ways to encode ASN.1, giving the encoding entity a choice of methods of encoding lengths, and in some cases a choice of order. BER makes no attempt to change the encoding of long octet strings in order to reduce the size.

Packed Encoding Rules (PER)

PER has been designed to minimise the bits transmitted. The encoding rules work on the basis that every bit that is not needed to transmit useful information is not sent. Thus for example, a Boolean is transmitted as a single bit rather than as an octet. Tags are generally not included and other ways of encoding this information are used instead, for example in a SEQUENCE with OPTIONAL and DEFAULT elements a single bit index is usually all that is needed to indicate the presence or absence of each element. There are four variants, namely:

- Packed encoding of a single ASN.1 type (basic aligned)
- Packed encoding of a single ASN.1 type (basic unaligned)
- Packed encoding of a single ASN.1 type (relay-safe-canonical aligned)
- Packed encoding of a single ASN.1 type (relay-safe-canonical unaligned)

The ALIGNED variety of PER inserts padding bits to align some elements to octet boundaries, whereas the UNALIGNED version does not. In the basic aligned and basic unaligned variants, the number of bits is minimised, but it is assumed that both encoder and decoder have precise foreknowledge of the abstract syntax being encoded, including all constraints. Thus, the flexibility of BER to handle unknown data types by skipping over them is lost. In some cases, the use of PER may result in a larger encoding and decoding processing overhead, compared with BER.

Distinguished and Canonical encoding rules (DER and CER)

For each ASN.1 type, these encoding rules select only one encoding from those allowed by the BER, eliminating all of the sender's options. The bit efficiency is therefore comparable with BER. Two main reasons for using these in preference to BER are: (a) some data integrity algorithms require the information to be encoded in exactly the same way whatever the implementation, and (b) skinny stack implementations which include the application layer will require a standard encoding of information so that pattern matching techniques can be used to scan incoming PDUs for known bit sequences. The main difference between DER and CER is that DER uses definite length encoding and CER uses indefinite length. CER is best chosen when the information is large and not all in main memory, whereas DER is best chosen when the information is all in memory (and thus it is easy to determine the size).

Lightweight encoding rules (LWER)

This is a proposed set of rules which would minimise the overhead of encoding and parsing the resulting transfer data streams in real time. Since most applications work on the basis of transferring data that is contained in a data structure in memory, LWER is designed to transfer data structures. The ASN.1 representation of the PDU will be transformed into a data structure in a systematic way. Provided that both applications use this structure and provided they have a similar architecture (in particular the same word size), the data can be loaded into memory and extracted from memory with minimal effort. The number of bits transferred will probably be higher than for BER. The intended environment is a high bandwidth communications path.

2.7 Naming, addressing and registration

2.7.1 Scope

This section contains material on upper layer naming and addressing which is not covered in the ATN Manual (2nd edition).

2.7.2 Naming and Addressing Guidance

Naming refers to the identifier which must be assigned to any information object which may need to be referred to during information processing. Addressing refers to the physical location of a resource. To quote ISO/IEC TR 10730:

> *"Naming and addressing mechanisms are an essential aspect of OSI. Real Open Systems, even while being fully conformant to OSI protocols in all seven layers, may well be unable to set up a dialogue because of inconsistencies among their naming and addressing policies."*

2.7.3 Naming

Information objects which may need to be named include:

- application processes
- managed objects
- ATN application message types

AP, AE and Context Naming

In a given end system, application processes (APs) are the elements which perform the information processing for particular user applications. They are identified by AP-titles, which must be unambiguous throughout the OSI environment (OSIE), and thus require a Registration Authority to allocate names.

When APs in different end systems need to cooperate in order to perform information processing for a particular user application, they include and make use of communication capabilities which are conceptualised as application entities (AEs). An AP may make use of communication capabilities through one or more AEs, but an AE can belong to only one AP.

Application entities are each named in terms of a unique Application Entity Title (AET), which consists of an AP-title and an AE Qualifier.

AP-titles and AE-qualifiers may be assigned either an attribute-based name form or an Object Identifier name form. When an AP-title is allocated an attribute-based name form, all of the associated AE-qualifiers must also be assigned an attribute-based name form; when an AP-title is allocated an Object Identifier name form, all of the associated AE-qualifiers must also be assigned an Object Identifier name form.

It may at times be necessary to make a distinction between the various invocations of a given AP running concurrently on an Open System. Thus it is possible, say, to have a pool of generic application servers, and when a server is allocated to perform a particular task, it is identified via its Invocation-identifier. This is done through the use of AP-invocation-identifiers which must be unambiguous only within the scope of the AP, and thus do not have to be registered.

Similarly, it may be necessary to distinguish between the various invocations of a given AE running concurrently as part of a given AP. This is done through the use of AE-invocation-identifiers which must be unambiguous within the scope of the {AP-invocation, AE} pair and thus do not have to be registered.

For communication purposes, AE-invocations have to handle one or more Application Associations. These can be identified by Application-Association-Identifiers, which need only be unambiguous within the scope of the cooperating AE-invocations, and thus do not have to be registered.

In connection-mode communications, the AET can be used in called, calling and responding application address parameters in A-ASSOCIATE (and RT-OPEN) service primitives. The ACSE service provides for the optional specification of an AET value by its component values (AP-title and AE-qualifier) in A-ASSOCIATE primitives. Similarly, the RTSE service (which itself makes use of ACSE) provides for the optional specification of an AET value by its component values in RT-OPEN primitives.

The calling/called AP title identifies the AP that contains the requester/acceptor of the A-ASSOCIATE service. The AE qualifier identifies the particular AE of the AP that contains the requester or the acceptor of the A-ASSOCIATE service. The AP and AE Invocation-identifiers identify the AP invocation and AE invocation that contain the requester or the acceptor of the A-ASSOCIATE service. The presence of each of these addressing parameters is defined in ISO standards as a user option.

Other information objects which must be named are the Application Context name and Presentation Context identifier which are exchanged at connect-time.

Managed Object naming

The selection of managed object (MO) instances for management operations is key to remote system management. CMIP supports the identification of MO instances by reference to the underlying Object Class and the Object Instance:

a) The Object Class is identified in CMIS primitives by either a global form or a local form of identifier (Object Identifier or Integer, respectively). Where the local form is used, the permissable Object Class identifiers are defined as part of the application context in which they are used. The local form is outside the scope of profiles AOM 11 and AOM 12.

b) A specific Object Instance is selected in CMIS primitives either by a global Distinguished Name (DN), a local DN, or a non-specific form. The DN is defined in the Directory standards; it consists of a series of attribute value assertions, each of which is an attribute type (Object Identifier) and an attribute value (any type). A local DN is that part of the DN that is necessary to identify the MO unambiguously within the context of the communication between the open systems. The non-specific form is defined to be an Octet String, and is outside the scope of profiles AOM 11 and AOM 12.

Thus, the parameters to identify a specific managed object can be many octets in length. It is likely that a more efficient mechanism will be required for the ATN. The containment hierarchy for ATN has not yet been defined, so it is not possible to define the exact length of the relevant DNs.

Naming of ATN Application Message Types

If a limited set of ASEs is used the problem arises as to how a user application which receives an incoming datastream can interpret the datastream, ie. how does it know how the data is encoded, and how to process it. A number of solutions are possible, including the following:

a) user applications could be identified by AP title (and hence by Presentation address). This has the disadvantage that the overhead of a new association would be incurred each time a different ATN application was invoked.

b) a standard header could be defined for all user data, which would identify the relevant ATN application (cf. the "UNB" header in an EDIFACT interchange). This implies that the header, even if defined in ASN.1, must always be encoded the same way. In practice, to minimise the message overhead, the header would be defined as a bit-pattern. The actual message contents would follow the standard header field, using any encoding mechanism defined by the ATN application. For maximum flexibility, the standard header could indicate which encoding algorithm has been applied (eg BER, PER, DLAC, etc).

Note 1: ACSE specifies a mandatory Application Context Name. This is an Object Identifier that in effect specifies the set of ASEs to be used. ACSE also specifies a mandatory presentation context definition list which should contain an entry for the presentation context identifier for each ASE in the application entity, including ACSE.

Note 2: the user data that passes across the presentation service interface consists of ASN.1 objects - associated with each of these objects is a presentation context identifier. In the applications being outlined in this paper we need to keep the number of presentation context identifiers to a minimum - in fact to just one; that of ACSE presentation context.

2.7.4 Addressing

According to ISO 7498-3 (Basic Reference Model - Naming and Addressing Addendum), a Presentation Address comprises a globally unique NSAP address, plus local Transport selector, Session selector and Presentation selector. Each (N)-selector identifies one or more (N)-SAPs at a given layer (N). The principles of are illustrated in figure 2.7.

Figure 2.7 : Naming and Addressing Concepts

The only mandatory addressing parameters in the A-ASSOCIATE (and RT-OPEN) service primitives are the calling, called and responding Presentation Addresses. These are passed transparently by ACSE (and RTSE) to the Presentation Service.

According to ISO/IEC DISP 11183 (Management Communications Protocols), in those cases where upper layers address selectors are used, the following maximum length constraints shall be observed and conformant implementations shall parse and handle values of the selectors of up to these maxima:

Session selector:	16 octets
Presentation selector:	4 octets

2.7.5 Name - Address Mapping

Any OSI layer entity or application process can be named via a title. This must be translated into an address by means of a directory function either at Application or Network Layer. Each AE is attached to one or more PSAP and hence the AET is associated with the corresponding Presentation Address(es). The AET is mapped onto a Presentation Address by means of an Application Layer directory function. The Application Layer directory function provides a mapping from an AET into the PSAP address required to access the referenced application entity.
The use of selectors is a local function and there may in practice be a direct correspondence between application entity titles and TSAP address or NSAP address.

2.7.6 Selectors in the ATN

Session service users are addressed at a Session Service Access Point (SSAP), where a session address consists of two parts - a session selector (Ssel) and a transport address. Likewise presentation service users are addressed at a Presentation Service Access Point (PSAP), where a presentation address consists of a presentation selector (Psel) and a session address. The various components are simple strings of zero or more octets and are meaningful only at a particular network address. An Application Process (AP) address can be defined as:

AP Address = NSAP Address + Tsel + Ssel + Psel

It follows that the values chosen for the Presentation and Session selectors is a local matter, although the actual values chosen must be made known to any potential ‘remote’ users - even if they are null. In the case of null presentation and session selectors the value of the TSAP can be used to convey information, such as the ATN application to be used.

For ATN upper layer addressing, the use of Session and Presentation selectors is not recommended. The use of Transport selectors is discussed in the section on the Transport Layer.
Upper layer addressing in ATN shall use null Session and Presentation selectors.

2.7.7 Registration Issues

It has been agreed that the ISO 9834 registration scheme should be put forward for adoption by ICAO.

A number of situations have been identified where object identifiers (OIDs) are being interchanged; some of these are registered elsewhere, some will need registration by ICAO. A given object should only be assigned one OID, ie. it should only be registered once (either by ICAO or by some other organization).

ICAO is in the process of setting up a registration authority under ISO. ICAO Working Groups will need to register information objects including ASOs, ASEs, Application Titles, Presentation Contexts through such a registration authority. It may also be necessary to set up a registration authority for Distinguished Names, as used by the Directory service and by systems management.

IATA, in developing the AOP, has adopted the ICD (or international) form of network address (NSAP) using ICD 0027, which is jointly administered by ICAO and IATA. Traditionally, IATA has assigned globally unique identifiers to international organisations, which they use for both store-and-forward and transactional communications based on IATA standard protocols. These identifiers include, among others, the two-character airline codes assigned to IATA members and other organisations.

3. Implementation Issues

3.1 Introduction and Scope

This section examines the protocol overhead of various upper layer profiles. The term "upper-layer OSI protocols" is itself ambiguous. Sometimes it used to mean ACSE, Presentation and Session (sometimes referred to as A/P/S) - the protocols that are used directly or indirectly by application-layer protocols (the fact that ACSE is itself nominally in the application-layer adds to this confusion). Sometimes it means everything above the transport layer, including the application layer protocols. Most of the discussion in this section concerns the supporting upper-layers; at least some of it can be applied to application protocols proper. It draws on the analysis of [2.6].

There are so many options available in the different layers that a thorough and complete analysis is not possible in a document of this size. Assumptions have had to be made in order to simplify the analysis. These assumptions are based on three premises:

- The functions in the stack will be selected in order to optimise the use of the bandwidth available.
- The majority of the overhead will be due to association set up, release and data transfer. Cases where the association is not accepted and cases where other upper layer services are used are therefore not considered.
- As the transport and network overheads are also significant, the aim is to minimise the number of PDUs transferred; therefore, when the protocol allows it, user data will be piggy backed on top of other PDUs.

3.2 Evaluation of OSI Upper Layer protocol overheads

3.2.1 Basis of Overhead Analysis

The use of supporting upper layers by a real application may be specified in a profile, rather than explicitly by the application base standard. For example FTAM [ISO 8571] has an associated profile specifying the use of ACSE, Presentation and Session [ISP 10607-1]. Such profiles may specify that non-mandatory items in the ACSE, Presentation and Session base standards are made mandatory for conformance to the profile. However, in this section the assumption is that the requirements of the application are based on the conformance clauses of these base standards.

The estimates below optionally include the most basic of upper layer addressing in the form of a NULL byte for Session and Presentation selector addresses.

For the purposes of this document a very simple OSI Application is considered which will use ACSE, Presentation and Session at their most basic. Note that moving common functionality from these layers into the applications themselves may cause the applications to be far from simple - which may cause a net gain in the number of bytes of data crossing the Transport-Session boundary.

A detailed evaluation of the overheads incurred by use of the OSI upper layer protocols is given in Annex C. In making this analysis, the following assumptions are made:

3.2.2 Basic Assumptions

a) "ACSE" refers to the first edition of ISO 8650 (1988), plus Amendment 1 - Authentication during association establishment (1990). "AACSE" refers to the draft text to extend ACSE to support ASO-associations.

b) Session Version 2 is used.

c) Unless otherwise stated, all OPTIONAL items are omitted from PDU size calculations.

3.2.3 Encoding Options

d) When the ASN.1 Basic Encoding Rules (BER) are used for ACSE and Presentation headers, the length field is mostly the definite variant (short or long) option.

e) ASN.1 Packed Encoding Rules (PER) in this paper refers to the Basic Unaligned variant of ISO/IEC CD 8825-2. By imposing limits (bounds) on value ranges in the abstract syntax, use can be made of constrained and semi-constrained whole number coding (for lengths and integers).

f) ASN.1 type EXTERNAL is taken to be as defined in ISO/IEC DIS 8824-1 (1992).

g) When a SEQUENCE OF SEQUENCE is encoded by ASN.1/PER it is assumed that the lower bound is zero and the upper bound is 255. In this way just one octet is required to specify the component count.

3.2.4 Name and Address Options

h) Where object identifiers (OIDs) are used, it is assumed that they are constrained such that no more than five arcs will be present and the component values will be small. This enables the OID to be encoded in 4 octets, with a length field of 1-octet (for BER) or 2-bits (for PER). The OID values assumed are listed in section 5 of Annex C.

i) An AP-title is assumed to be an OID and an AE-qualifier is assumed to be a small value Integer (less than 128). The effect of including addressing fields in Called and Calling Identification fields in ACSE connect PDUs is illustrated in section 6 of Annex C.

j) Both the Presentation Selector and Session Selector have a value of NULL. Note that CULR-1 specifies a maximum of 4 octets; however 1 octet will provide 255 different values for each selector.

3.2.5 Presentation Context

k) When AACSE is used, even the presentation-context-definition-list in Presentation CP and CPA PDUs are omitted 2 (AACSE provides this information in its own PDUs).

l) It is assumed that in the case of BER encoding the default presentation context definition will use an abstract syntax of ACSE (Version 1) and a transfer syntax of ASN.1/BER. Basic communication applications do not require default Presentation Context in CP and CPA [CULR-3].

m) It is assumed that Presentation Context Identifiers are small valued integers (less than 128). One presentation context only is specified, this is for ACSE. If the application needs to specify a distinct abstract syntax, then the addition of this adds 15 octets (when selecting ASN.1/BER) or 17 octets (when selecting ASN.1/PER), as illustrated in Annex C.

n) Only one transfer syntax (ASN.1/BER or ASN.1/PER as appropriate) is negotiated.

3.2.6 User Data

o) Presentation user data is simply-encoded-data, that is an OCTET STRING.

p) The effect of application user data is shown as an option. User data is sent as an OCTET STRING and so its structure will be the subject of bilateral agreement.

4. CNS/ATM-1 Package Guidance Material

4.1 Introduction

4.1.1 Motivation of the Work

4.1.1.1 Introduction

The original discussion for the CNS/ATM-1 ULA is presented.

The section discusses the proposed ULA in terms of its relation to two other proposals. These proposals are the application-over-transport proposal, and the current OSI solution.

As requirements for an ATN ULA are extracted, they are stated as Rn statements. After the two current approaches are discussed, the ATN ULA is presented. The Rn statements are then collated and ULA satisfaction of those requirements are stated in complementary Un statements.

[2] Strictly the presence of the presentation-context-definition-list field is mandatory, as ISO 8823 Section 6.2.2.7 states "This shall be a list containing one or more items."

4.1.1.2 Applications over transport solution

Current work, for example, the ground-breaking RTCA DO-212, DO-219, DO-223 series of applications, is specified directly over the transport layer. This leads naturally to the question. Why should one have a ULA at all? If one can send messages over a transport provider, what else is required?

As one reads, e.g., DO-219 the RTCA Two Way Data Link (TWDL) specification, certain characteristics of that application design process naturally emerge. For example, message processing in terms of message identification, message urgency and message response characteristics are all specified. Clearly, the application must handle duplicate messages. Each application needs to perform this function separately. Second, message urgency and response requires the application to implement an urgency scheme to handle multiple messages correctly.

As one reads further in DO-219, the full encoding of each message in terms of Abstract Syntax Notation 1 (ASN.1) and Packed Encoding Rules (PER) is specified. Clearly, without a ULA, an application over transport must itself fully specify the complete encoding of each message. We note that DO-219 found it necessary to import a great deal of material on encoding into its specification. We also note that DO-212 encoding is specified in terms of bit maps. If DO-212 were upgraded to ASN.1/PER for consistency and transmission savings, that section of the standard would have to be rewritten.

The first point, then, is that examination of present RTCA specifications has ULA requirements. However, as each application is constructed, it must itself build and specify each application element that it requires, even if that application element exists in a standardized form, or exists elsewhere. Upper layer requirements exist. The choice is a use of a standardized ULA and associated tools, or specification and implementation of upper layer elements individually for each application. Reusability of existing work is an important goal. This leads to shorter development, smaller code, and certification credit.

R1. ULA Requirements (message accountability, encoding) exist for all cited applications, regardless of whether they are formalized in a ULA.

R2. Common application service elements are an important ULA contribution

The application requires an orderly termination of the application dialogue. A transport service alone cannot perform an orderly termination. If a data (DT) transport protocol data unit (TPDU) is sent, and then a disconnect request (DR) TPDU is sent, clearly the delivery of the DT is not guaranteed. Indeed, a DT can have been acknowledged by the Transport Service Provider (TS-Provider), and still not delivered to the Transport Service User (TS-User) if the DR arrives. Clearly, this requirement must thus be handled above transport.

The definition of when message exchange begins and ends can only be defined by the cooperating applications. If this is not done in a current application, the transport layer

cannot help. The orderly termination of the message exchange is a requirement for every current application. This message exchange is generally referred to as an association. It is crucial to the discussion to realize that an application over transport makes the association synonymous with the transport connection.

R3. Every application has a requirement for association control (orderly initiation and termination of message exchange)

An important consequence of this is that every application must open a transport connection. Thus, for example, the Context Management (CM) dialogue consists of exactly one uplink and one downlink. Still, an entire transport connection setup and teardown is required to support the one exchange. Clearly, aircraft running multiple applications to one air traffic control (ATC) center would benefit immensely from multiplexed associations over a single transport connection. The conservation of bandwidth by multiplexing aircraft-center applications is a worthwhile benefit.

R4. Multiple applications running over a single transport connection is an important goal.

The security requirements of the initial ATN applications merit consideration. The hole in initial authentication requirements for RTCA applications is well-known. For example, the TWDL application implements extensive procedures for transition to the Next Data Authority (NDA), but the aircraft simply responds to the first Connection Request (CR) it receives as its initial data authority. Security requirements are not specified in current application-over-transport standards.

R5. There exist CNS/ATM-2 requirements authentication security requirements for each application.

The previous hole might be addressed as follows. If it were required that CM executed before CPDLC, then there is a start to a guarantee that the entity calling the aircraft actually got the address from CM.

R6. Currently specified applications have implicit requirements for application control in terms of application-to-application interaction and application ordering.

This section has considered certain design aspects of applications over transport. Certain ULA aspects present in these applications were discussed and certain ATN ULA requirements were discussed.

4.1.1.3 Classic OSI solution

The ATN standards have consistently stated that the ATN is fully based on International Organization for Standardization (ISO) standards for Open System Interconnection (OSI). For the ATN to succeed, international voluntary standards must be followed, rather than proprietary specifications.

R7. The ATN is an OSI standard architecture.

ISO has specified a seven-layer reference model. At the upper layers, application, presentation, and session protocols are specified. Certain aspects of the protocol are not workable for the ATN. For example, a full upper layer protocol establishment requires 98 octets. This is not workable for air-ground subnetworks.

R8. Current OSI Upper Layer protocol overhead is unacceptable for air-ground links.

It is also clear no current ATN air-ground applications require any of the services offered by the session protocol.

R9. There are no current air-ground requirements for the OSI session protocol (ISO 8327, edition 2) functional units.

The ISO presentation protocol similarly has greater functionality than required. For example, there is no need for negotiation of complete defined context sets (DCS) for ATN applications.

R10. The OSI presentation protocol (ISO 8823, edition 2) must be tailored for ATN use.

This short discussion indicates that the classic OSI protocols are untenable in terms of protocol overhead and also in terms of extra protocol functionality.

4.1.1.4 The ATN ULA

The ATN ULA has three aspects. These are as follows:

ISO 9545, edition 2 specifies the extended application layer structure (ALS). This allows modular construction of software by specification of application service elements (ASEs). An ASE is implemented as a software module. These are combined into Application Service Objects (ASOs). Interaction between ASEs and ASOs are mediated by a control function (CF).

ISO 8649, edition 2 and ISO 8650, edition 2 specify the Application Control Service Element (ACSE) needed to support the ALS. ACSE allows the establishment of associations over transport connections.

Amendments to ISO 8823 and ISO 8327 specify efficient Presentation Protocols and Session Protocols. The amendments specify minimal functionality protocols to indicate protocol functionality.

The ATN ULA uses the ATN transport service.

4.1.1.5 Review of ULA requirements conformance

The requirements for the ULA are listed below.

R1. ULA Requirements (message accountability, encoding) exist for all cited applications, regardless of whether they are formalized in a ULA.

U1. The ULA offers a formal method of ASO specification. This also promotes software reusability. Certification authorities have reacted positively to the use of ALS for certification credit, such that ASOs do not have to be recertified.

R2. Common application service elements are an important ULA contribution.

U2. In the ULA, ASEs can be formally specified. They can be formally checked for correctness. They can then be combined in a modular fashion.

R3. Every application has a requirement for association control (orderly initiation and termination of message exchange).

U3. The association control service element (ACSE) offers a tool for association control for every application.

R4. Multiple applications running over a single transport connection is an important goal.

U4. ACSE directly supports the higher level associations. All extensions to ACSE are bit-wise backwards compatible with previous editions of ACSE.

R5. There are CNS/ATM-2 authentication security requirements for each application.

U5. ACSE has an authentication element which is an important element in a security architecture.

R6. Currently specified applications have requirements for application control in terms of application-to-application interaction and application ordering.

U6. The use of the CF allows precise specification of application interaction and application ordering.

R7. The ATN is an OSI standard architecture.

U7. When the ATN SARPs are published, every ATN ULA element will be an OSI standard. Every ULA enhancement is consistent with the OSI architecture.

R8. Current OSI Upper Layer protocol overhead is unacceptable for air-ground links.

U8. The upper layer efficiency standards reduce the upper layer establishment overhead from 98 octets to one.

R9. There are no current air-ground requirements for the OSI session protocol (ISO 8327, edition 2) functional units.

U9. The session functional units are not in the ATN ULA. If and when required, the session functional units will be recast.

R10. The OSI presentation protocol (ISO 8823, edition 2) must be tailored for ATN use.

U10. The OSI efficiency measures do not require the use of the current presentation protocol.

4.1.2 Architectural Guidance Material

4.1.2.1 Description of the Upper Layer Architecture

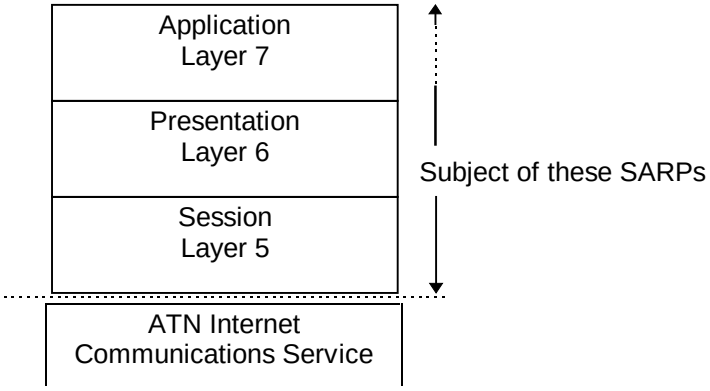

Figure 4.1: Conceptual view of the scope of the UL SARPs

Note -- In the picture above, it is useful to note the effect of the supporting layer efficiency enhancements.
First, the connect request at the application layer is mapped to the connect request of the supporting upper layers.
Second, the data request at the application layer is mapped directly to the transport data request.
Third, there is no congruent mapping for the release request. It must be handled by the Control Function using data requests (to carry the ACSE release) and abort requests (to actually clear the connection after the graceful release.)

4.1.2.2 Description of the Application Layer

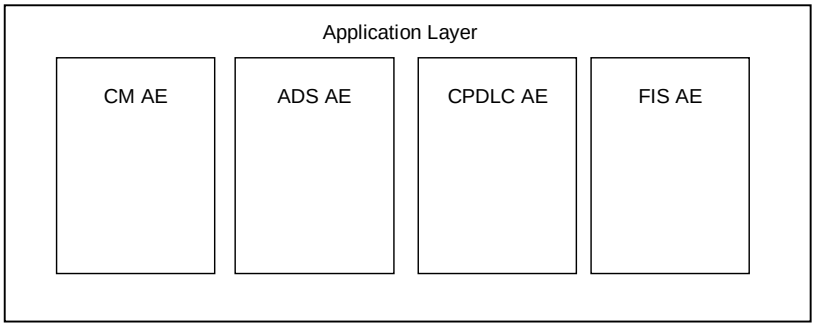

Figure 4.2: Conceptual view of Application Layer

Note -- It is important to note the architectural decisions embodied in the above picture.
Each application is embodied in an AE.
Each AE contains an ATN ASE, which is the communications element responsible for an ATN application. The internal structure of the ATN ASE may be of arbitrary complexity. The dialogue service is the ATN ASE's view of the ATN ULA.
Thus, the type of the AE is the same as the type of the ASE. That is, an ADS AE will contain only an ADS ASE. Thus, on connection establishment, the peer title can be completely constructed. The AE-title is provided by the peer ID value in the dialogue service, and the peer AP-qualifier (e.g., CPC) is the same as yours is.
There is no architectural capability for multiple instances of the same ATN ASE within the same AE. A requirement for another instance of an ASE (e.g., the CM contact procedure), requires spawning another AE. This implies that the ATN ASE will generate and manage only one dialogue over the lifetime of the ASE.
There is no interaction inside an AE between ASEs of different types, since these are not provided by the architecture. Any requirement for interaction between applications occurs in user space through, e.g., global data structures.
As implied in the figure, there is no upper-layer addressing. All addressing of the AE-type is complete with the TSAP address. All upper layer selectors are necessarily nil.
The application context name is also a simple construct, since the CNS/ATM-1 ASE list can be inferred from the AP-title. The CNS/ATM-1 application context name encodes only the application version.

4.1.2.3 Description of the AE Structure

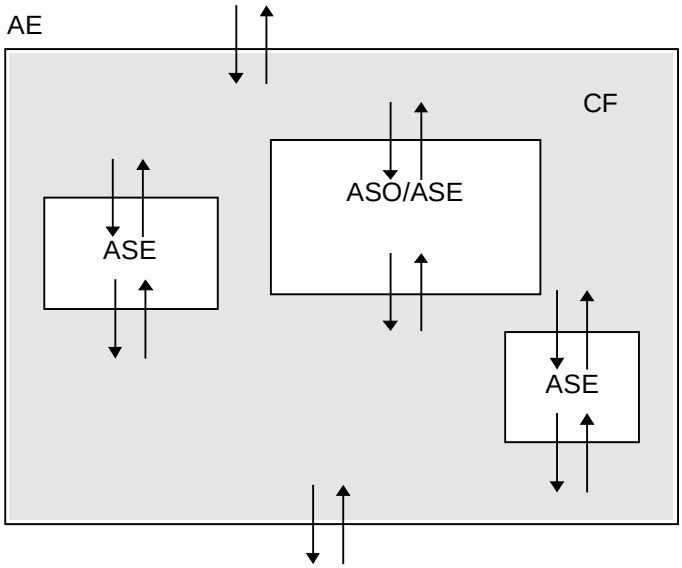

Figure 4.3: Generic Application Entity structure

The internal picture of the AE indicates that the AE comprises several ASEs. The AE picture is completed by the presence of an upper and lower service boundary.
Each ASE picture is similar in form, with an upper and lower service boundary. The job of the Control Function is solely to map inputs and outputs to and from ASE and AE service boundaries. By definition, the CF cannot produce protocol; in that case, an ASE is required. 18:31 observe a state machine.
The confirmed data service element (CDSE) is still undergoing requirements analysis. The CDSE is shown as an example of the inclusion in the ATN ULA of a common ASE constructed for use by all ATN applications.
In CNS/ATM-1, the upper AE service boundary is identical with the upper ASE service boundary. In CNS/ATM-1 all AE inputs are mapped to the ATN ASE.

4.2 Dialogue Service

4.2.1 Introduction

The dialogue service is a service that allows a user to bind to an association, send data, and unbind from the association.

4.2.2 Connection Mode

The dialogue service provides a connection mode upper layer service to the application. There is no connectionless mode upper layer architecture in CNS/ATM-1. Thus, there is no use of the connectionless transport protocol in CNS/ATM-1.

4.2.3 D-ABORT

The D-ABORT whether the abort is from above or below the Application Entity (AE). The D-ABORT contains an additional parameter (*Originator*) which identifies who requested the abort. This parameter can have one of the following abstract values:

- User - Indicates that the user of the AE requested the abort; or
- Provider - Indicates that the either the local or the remote AE requested the abort.

It is important to note that the D-P-ABORT indication now uniquely indicates that the underlying communications service provider has aborted.

4.2.4 Security

No application in the CNS/ATM-1 package utilizes the D-START security parameter, hence its implementation is not required in the CNS/ATM-1 package.

4.2.5 Mapping to transport

The ATN ULA makes extensive use of the user-data capability of the transport service. The ULA attempts to map the AARQ (containing user data) to the T-CONNECT. If this is not possible, based on user-data size, the T-CONNECT is sent, and the T-DATA is used to convey the AARQ (including user data). The implementor is advised to consult the PDU calculations in the present Guidance Material, but generally if the DS-User places more than five octets of user-data in the D-START, the D-START is mapped to the T-CONNECT + T-DATA.

The details of the mapping of D-START values (i.e., Residual Error Rate (RER) and Traffic Type) to transport values is for urgent resolution.

4.3 Control Function (Air/Ground)

The Control Function (CF) mapping function is described in the following figure. The five threads to implement the CF are described.

The D-START is mapped to the A-ASSOCIATE, which is mapped to the P-CONNECT by ACSE. The P-CONNECT is then mapped to the P-CONNECT (which is the T-CONNECT plus fast byte).

The D-P-ABORT occurs as an indication (‘up’) only. The P-P-ABORT is mapped to the A-P-ABORT. The A-P-ABORT is mapped to the D-P-ABORT.

The D-ABORT is mapped to the A-ABORT. The A-ABORT is mapped to the P-U-ABORT. The P-U-ABORT with user-data is mapped to the P-DATA which is followed by a P-U-ABORT. The P-U-ABORT without user-data is mapped to the P-U-ABORT.

The D-END is mapped to the A-RELEASE, which becomes a P-RELEASE. The P-RELEASE is mapped to the P-DATA.

The D-DATA is mapped to P-DATA, which in CNS/ATM-1 is syntactically equivalent to T-DATA.

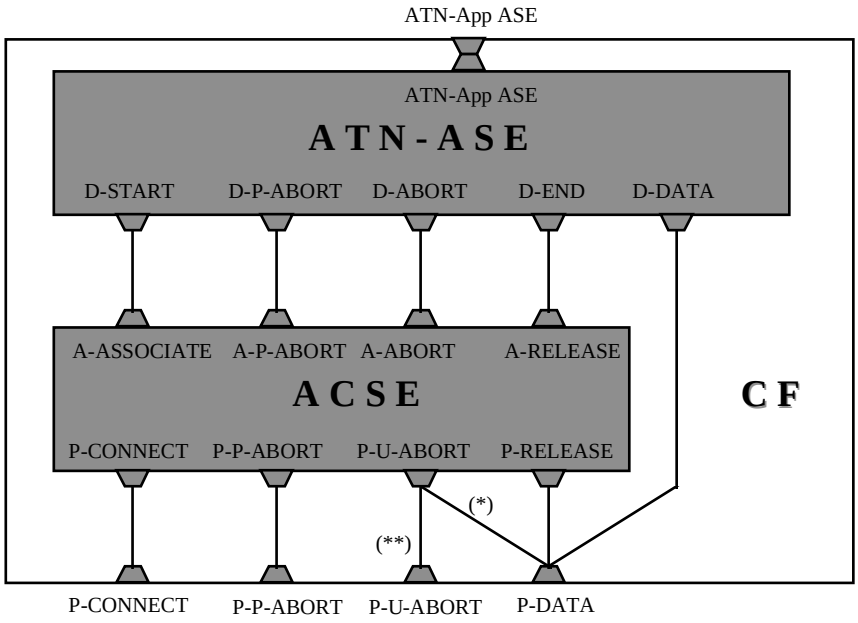

(*) user data in the abort
(**) no user data in the abort

4.4 Session

4.4.1 Session Defect Report

If the AARQ fits in the T-CONreq, but the AARE does not fit in the T-CONcnf, the session state machine will interpret the lack of user-data in the T-CONcnf as a refusal of fast-byte capability. A continue bit should be defined in the session short accept, meaning that user-data follows the accept.

4.5 Presentation

4.5.1 Presentation Defect Report

When a D-STARTrsp- produces a P-CONrsp-, the presentation state machine has an empty transition. The transition (P-CONrsp-, Await S-P-CONrsp) should be [SCR, STAI0]

4.5.2 Presentation Fast Byte Guidance Material (OSIEFF)

Null functionality at a layer refers to the case where no functionality is required of a layer during the data transfer phase but where OSI compatibility and compliance are required. This is the case which is most clearly applicable to the ITU-T applications which use “short stacks”, to permit much greater OSI compatibility while at the same time allowing efficient communications. While it is possible to use the normal OSI layer protocol to signal that null functionality shall be required in the data phase, in certain instances, it is also possible to use a different protocol which is considerably more efficient (in terms of byte efficiency and, possibly, connection-establishment efficiency) to perform the negotiation. The term “fast byte” has been employed, as a convenient mnemonic, to refer to the insertion of a single byte PCI at connection establishment to signal that no further PCI will flow for that instance of communication. The use of the fast byte at a layer therefore serves to provide a service mapping between the layer above it and the layer below it.

NOTE — Early discussions on the fast byte concept considered the possibility of using the byte — theoretically only a bit — to serve as a *placeholder* so that it was possible, at some point during the instance of communication, to use the normal layer protocol to re-negotiate the use of some layer functionality where none was required earlier. At least for the upper layers, such a dynamic re-negotiation of layer functionality is for further study.

Thus, if a transport layer fast byte is exchanged, the layer service remains the same, i.e., the transport fast byte is a different version of the transport protocol with a one-to-one mapping of the network services to the transport services. In other words, by using the transport fast byte, one gets a QoS which is only as good as that provided by the underlying network service. The lower layer fast bytes are particularly useful for cases where the applications are communicating over a single data link, as is the case with ISDN access signalling.

For the upper layers, the typical, normal OSI implementation requires a 13 to 20 octet overhead on a single presentation data value (pdv) using the presentation and session data transfer services. This overhead is necessary to identify the state of the communication (i.e., that it is the data transfer phase as opposed to, say, the release phase), and to identify the pdv as belonging to a particular presentation context. Clearly, a null PCI optimization for the data phase implies a reduction in the layer

service available to the application. For instance, in this case where all the application data is carried directly as user-data of the transport service, there is no guarantee that an encoded application pdu will not resemble a session spdu; therefore, null PCI for the session data transfer phase implies that it is not possible to distinguish session spdus from application PCI. Therefore, it is not possible to use the orderly release facility of the session layer, though, of course, the application protocol can be defined to perform this function. Similarly, null PCI for presentation data transfer implies that there can only be one presentation context for the application pdus, whose abstract transfer syntaxes are known a priori. Thus, reducing the upper layer functionality inherent in the null functionality data phase restricts the range of applications that can use this optimization.

This loss of functionality must be reflected to the user at the service interface. For the session and presentation layers, the layer services are bundled together in groups known as functional units. At this time, orderly release of the session connection is provided as a part of the mandatory kernel functional unit. The use of null encoding for the data phase requires that the users have negotiated the use of a new functional unit, the so-called no orderly release functional unit, which removes the orderly release from the kernel functional unit.

NOTE — The orderly release capability would more logically be a functional unit separate from the kernel; the new "negative" functional unit provides compatibility with the current specificatons that require the (non-negotiable) kernel to be indivisible.

To this end, ITU-T Rec. X.pfbs defines the pass-through access to the session service, in particular the (new) no orderly release functional unit. As the presentation layer uses the session layer services for release of the presentation connection, there is no reduction to the presentation services. Thus efficiency optimizations available at the presentation layer are new protocol options, i.e., alternative, efficient PCI and procedures.

ITU-T Rec. X.pfbp define two protocol options at the presentation layer that greatly reduce the quantity of presentation PCI in cases where the presentation user's requirements for presentation functionality are limited. The null-encoding protocol option provides an alternative presentation protocol option for data transfer with zero PCI which can be negotiated at connection establishment only if one of the following conditions described below is true:

a) the presentation context definition list contains precisely one item in which the abstract syntax name is known to the responding presentation protocol machine by bilateral agreement; or

b) the presentation context definition list is empty and the default context is known by bilateral agreement; or

c) the presentation context definition list is empty and the abstract syntax of the default context is known to the responding presentation protocol machine by bilateral agreement and is specified in ASN.1.

NOTE — It may be possible in the future to negotiate the null-encoding protocol option for efficient data transfer using the presentation protocol defined in ITU-T Rec. X.226 (1994) | ISO/IEC 8823-1:1994. It is left for further study to define an alternative version of the presentation protocol encoded using PER which will permit byte-efficient presentation negotiation of the full set of presentation functionality.

In addition, it is possible to use another protocol option, the short-encoding option, which defines encodings for some presentation PPDUs which are considerably shorter than the current ones if *both* conditions d) and e) described below are true:

d) the calling and called presentation selectors are null; and

e) the presentation-requirements parameter in the P-CONNECT service includes the kernel functional unit only.

The short-encoding protocol option allows the negotiation of the encoding rule which shall be used as the transfer syntax of the application PCI belonging to the single presentation context (which may be the default context) from one of BER, the aligned and unaligned variants of PER or a “transparent” encoding which is understood by bilateral agreement.

ITU-T Rec. X.sfbs specifies the no orderly release functional unit, whose selection by the session user indicates that the user has no requirements for orderly release of the session connection. Thus, either the application protocol has been chosen to perform this function, or the application association (which is one-to-one with the underlying session connection) is released by disconnecting the transport connection or by an abortive release of the session connection. The selection of this functional unit by the initiating session user permits the initiating session protocol machine to offer the use of the null-encoding protocol option on the established session connection. The responding session protocol machine can accept this option if the responding session user has selected only (and nothing other than) the kernel, full-duplex and no orderly release functional units for use on the connection.

ITU-T Rec. X.sfbp describes how the negotiation of the null encoding protocol option can be done using the protocol options field of the session establishment SPDUs defined in ITU-T Rec. X.225 (1995) | ISO/IEC 8327-1:_____. However, ITU-T Rec. X.sfbp also defines the possibility of using the short-encoding protocol option for the establishment SPDUs, which define a one byte PCI for these SPDUs which are distinct from the leading octet of the current SPDUs, which provides a byte-efficient negotiation of the null-encoding protocol option provided that there is no session layer addressing information required to be exchanged, i.e., the session selectors are null.

It is expected that the short-encoding protocol option will be used in conjunction with the transport connection set-up to achieve interworking with current implementations and, for the case where the responder also implements this protocol option, achieve an improvement in round-trip efficiency by setting up the upper layer connections concurrently with the transport connection. This is achieved as follows: the SHORT CONNECT SPDU — which is the short-encoding version of the current session CONNECT SPDU — is sent as user data of the T-CONNECT request service primitive. This requires that the SHORT CONNECT SPDU plus any accompanying user data meet the 32 octet limitation on the size of the transport user data.

NOTE — The transport protocol class 0 does not permit the carriage of user data. Therefore, for this scenario to work, the transport protocol clsss 4 should be available at both ends, or the transport fast byte protocol should be employed.

Current session implementations ignore any user data on the T-CONNECT indication primitive, or, at worst, disconnect the transport connection. Thus, absence of any user data on the T-CONNECT confirm primitive is a signal to the initiating session protocol machine that the responder is an implementation of the current standards. If the responding session entity implements the short-encoding protocol option, the SHORT ACCEPT SPDU is sent as user data of the T-CONNECT response service primitive, and its receipt by the initiating session protocol machine completes the session connection establishment in tandem with the transport connection establishment.

Of course, the short-encoding option may be used with the T-DATA service for the case where an already established transport connection is assigned to the session

connection. Interworking is not fully achieveable as there is no guarantee that the responding session entity, if based on the current standards, will send a REFUSE SPDU to signal a protocol error, which is what a short-encoding for an SPDU would be.

4.6 ACSE

4.6.1 Discussion of differences in ACSE editions

As the ACSE is the major part of available software in the ULA implementation, a description is provided of the evolution of the ACSE standard. The editor's preface to ISO 8649, Service Description, was amended three times from the first edition to the second edition. The three amendments are 1) Peer-entity Authentication during Association Establishment, 2) Connectionless ACSE Service, and 3) Application Context Name Negotiation. There were also three technical corrigenda (Tcs) cited. The most important of the TCs resolved a defect wherein EXTERNAL events could affect ACSE sequencing and state machine. The EXTERNAL events were added to ACSE in a TC to answer a defect that pointed out that a session resynchronization could purge (destroy) the session finish or disconnect that the A-RELEASE is traveling on. Now, clearly over the ATN ULA of strict fast-byte supporting upper layers none of this matters, since there are no external events. A classic implementation, though, must put in a few EXTERNAL hooks at the bottom of its state machine when things go wrong in classic session/presentation.

ISO 8650-1, edition 2, has two Amendments and four TCs noted in the Editor's Preface. AM1 is Authentication, AM2 is Application Context Name negotiation, and TC2 is the 'EXTERNAL' TC.

The changes in the ACSE protocol specification are roughly similar, e.g., for TC2, "A-RELEASE procedure is disrupted if P-RESYNCHRONIZE, P-U-EXCEPTION-REPORT, or P-EXCEPTION-REPORT primitives occur on the association".

The state machine also adds two stimuli for EXTERN-1 and EXTERN-2. These stimuli cause the ACPM to return to the Associated state from one of the Attempting Release state.

The discussion indicates that the CNS/ATM-1 ULA requires none of the changes that distinguish ACSE, edition 1 from ACSE, edition 2.

The two changes to ACSE that are required are the requirements to encode the ACSE PDUs in PER, and to map the P-RELEASE to P-DATA.

4.6.2 ACSE Primitive Flow Diagrams

The role of the CNS/ATM-1 Control Function (CF) is to modulate the interaction of the ATN ASE and the ACSE. To that end, description of the characteristic flows involving ACSE is provided in the figures below.

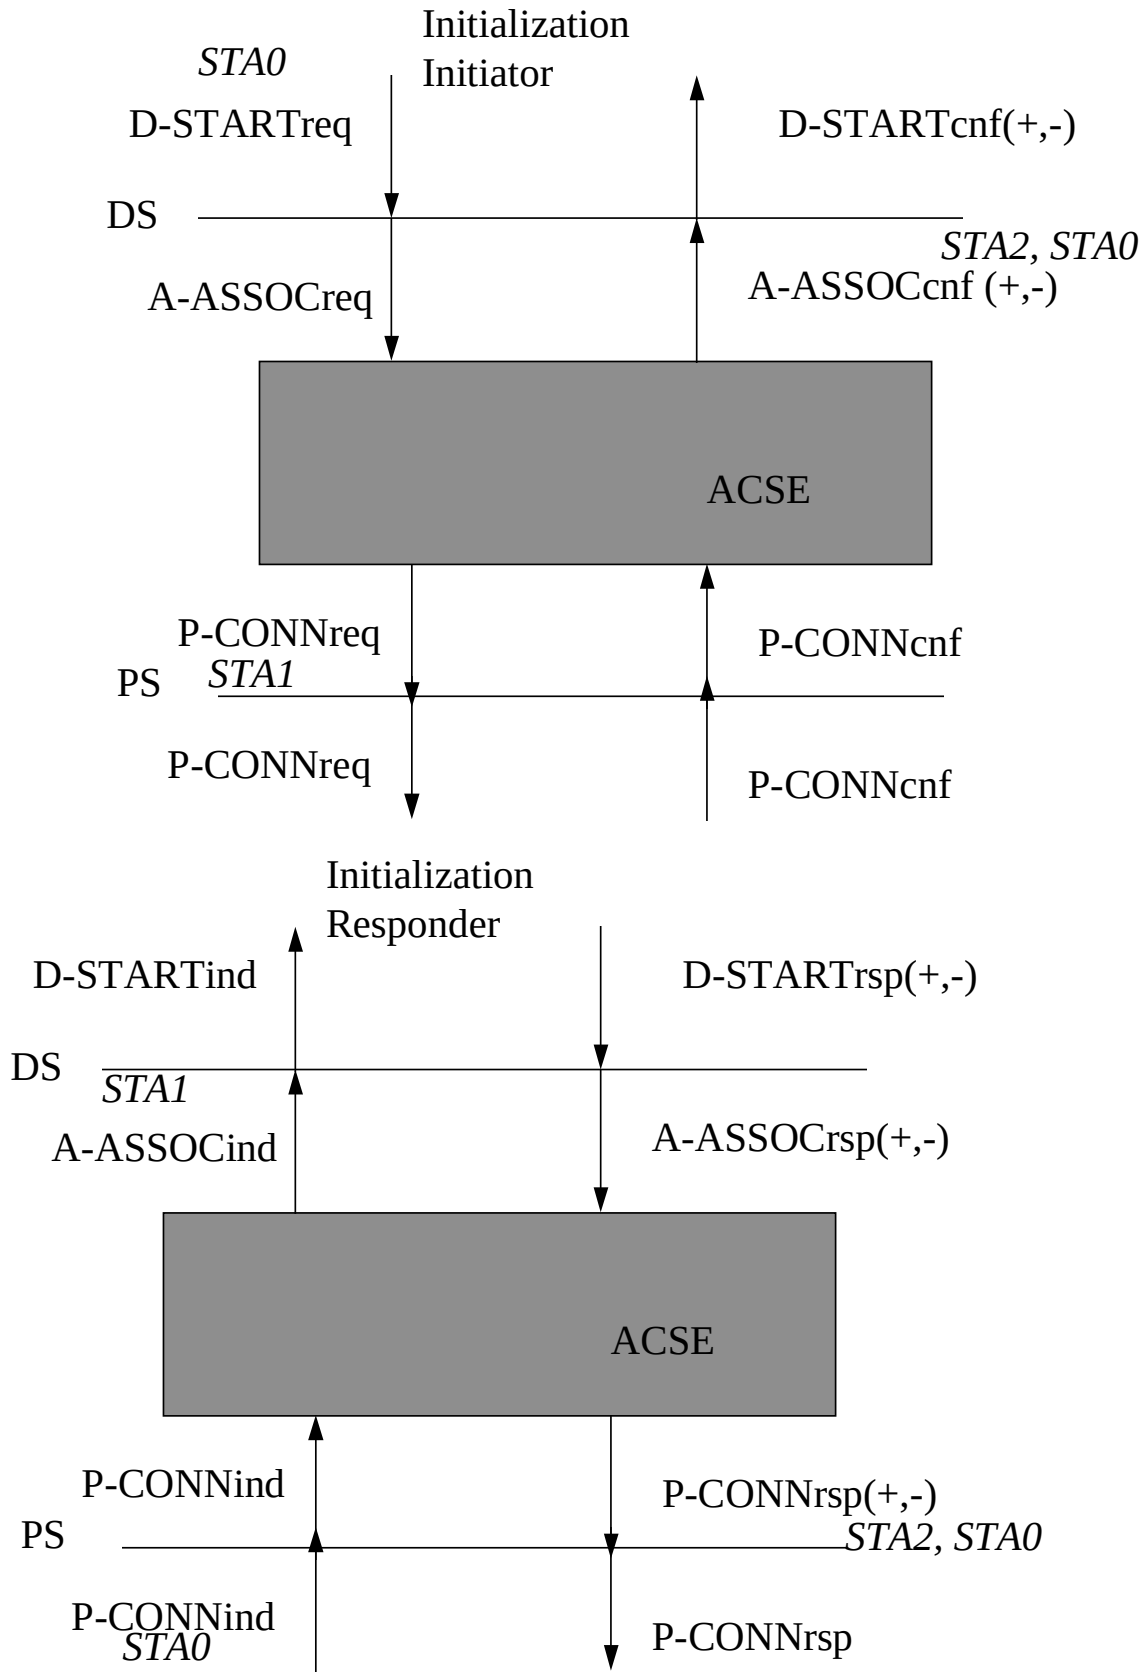
Initialization
Initiator
STA0
D-STARTreq
D-STARTcnf(+,-)
DS
STA2, STA0
A-ASSOCreq
A-ASSOCcnf (+,-)
ACSE
P-CONNreq
P-CONNcnf
PS
STA1
P-CONNreq
P-CONNcnf
Initialization
Responder
D-STARTind
D-STARTrsp(+,-)
DS
STA1
A-ASSOCind
A-ASSOCrsp(+,-)
ACSE
P-CONNind
P-CONNrsp(+,-)
PS
STA2, STA0
P-CONNind
STA0
P-CONNrsp

Data Transmission
Data Receipt

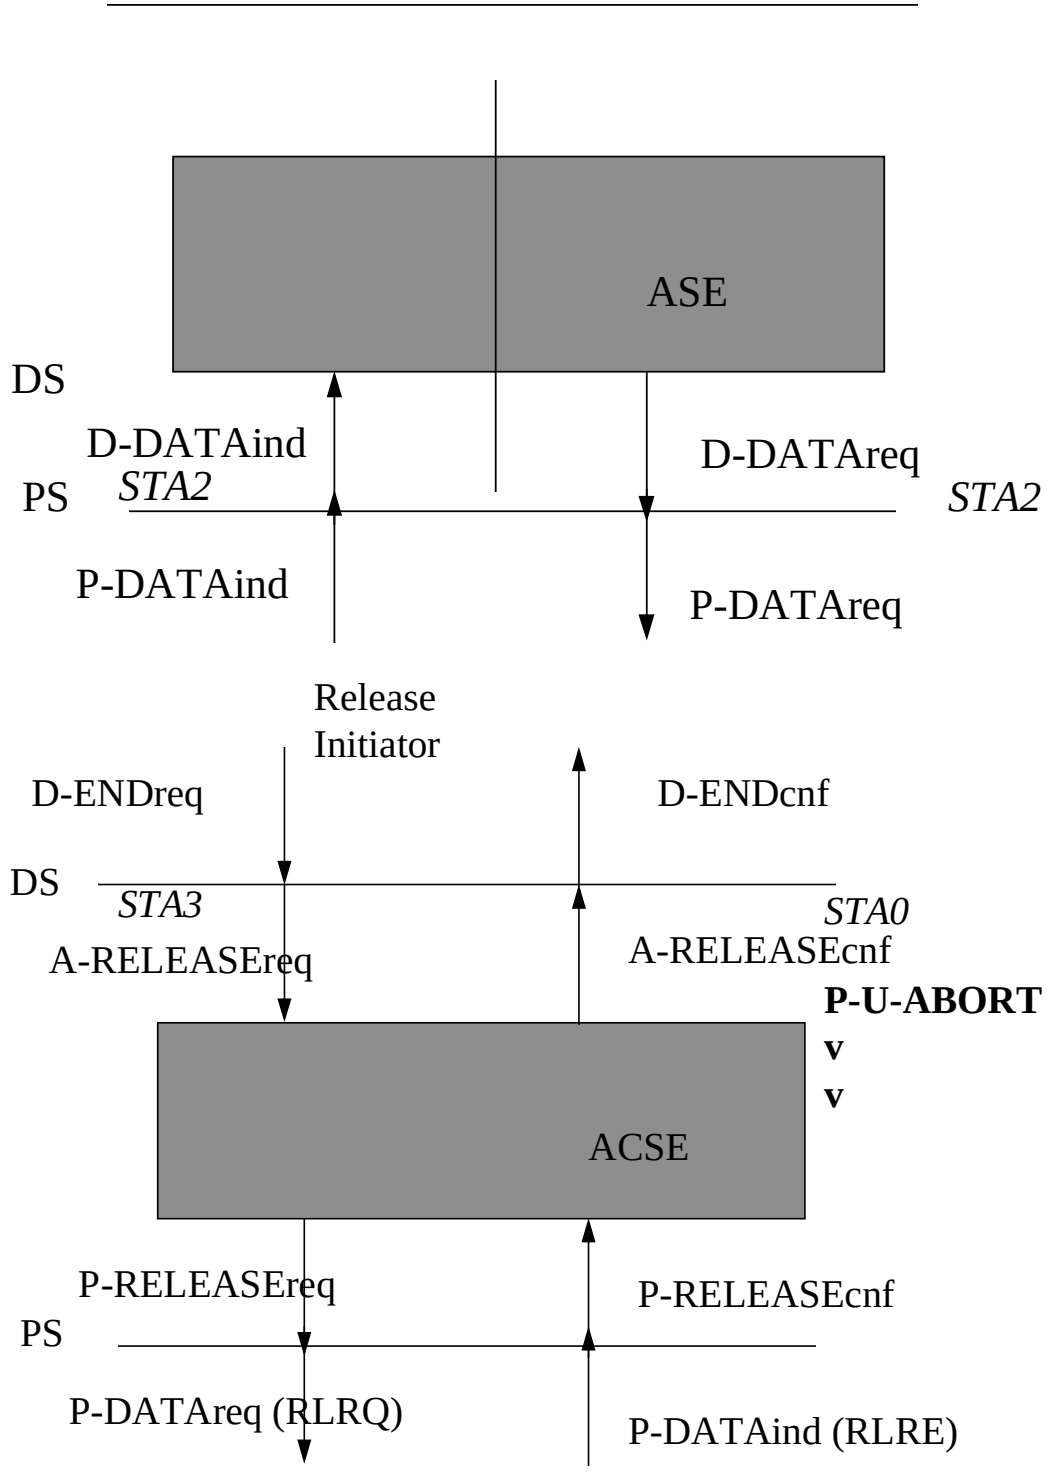

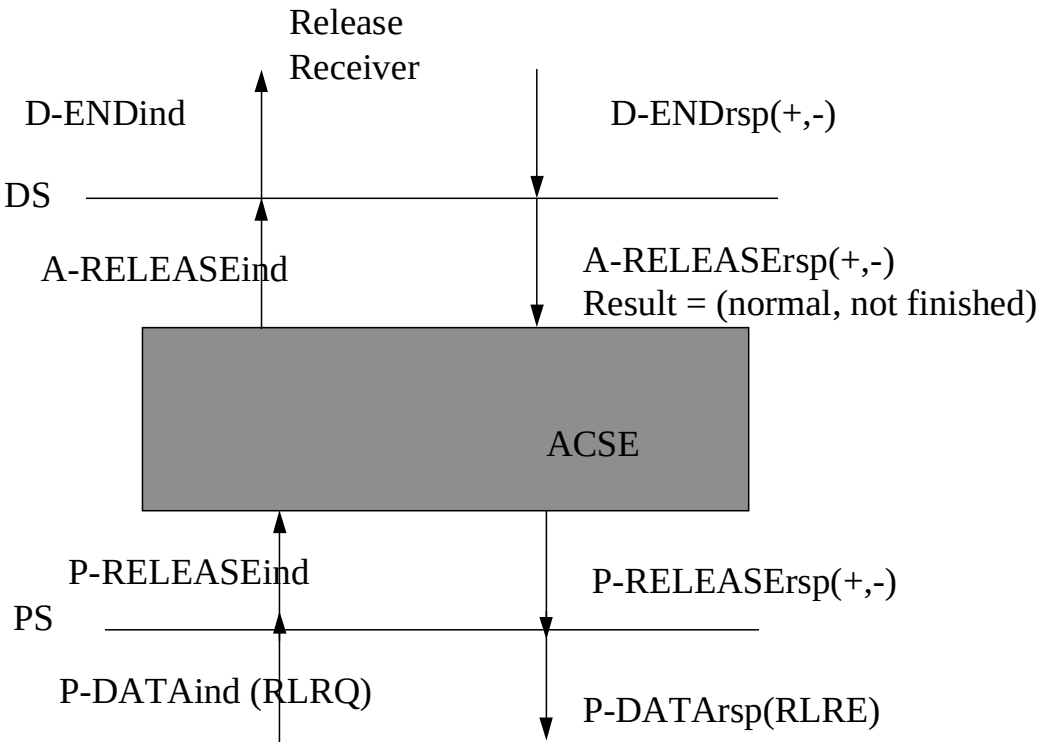

Release Collision - Initiator

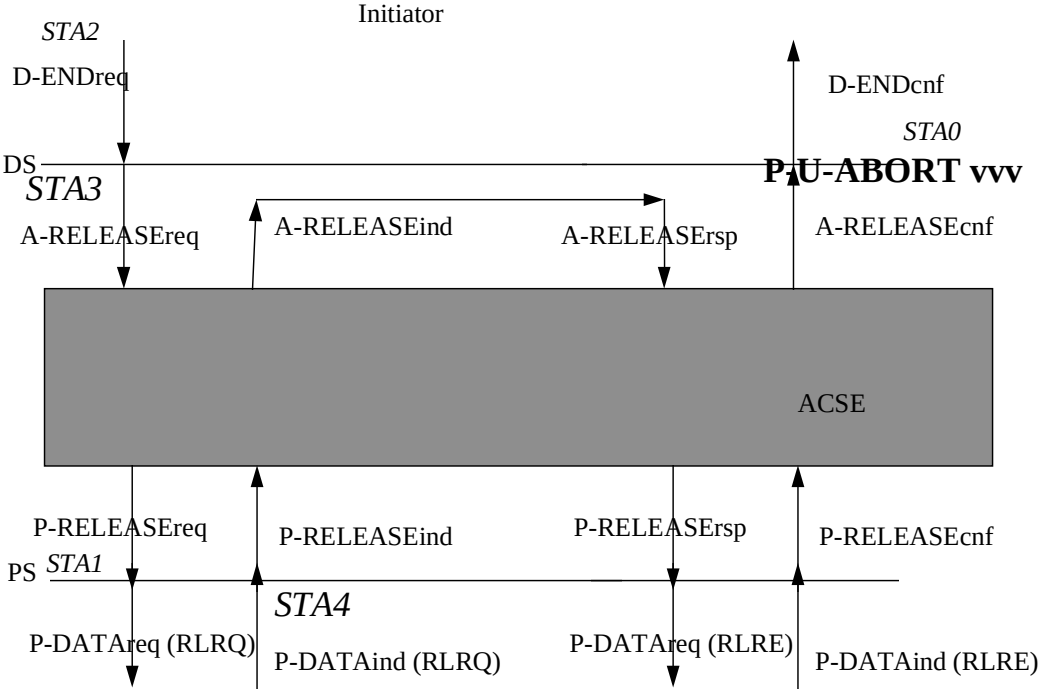

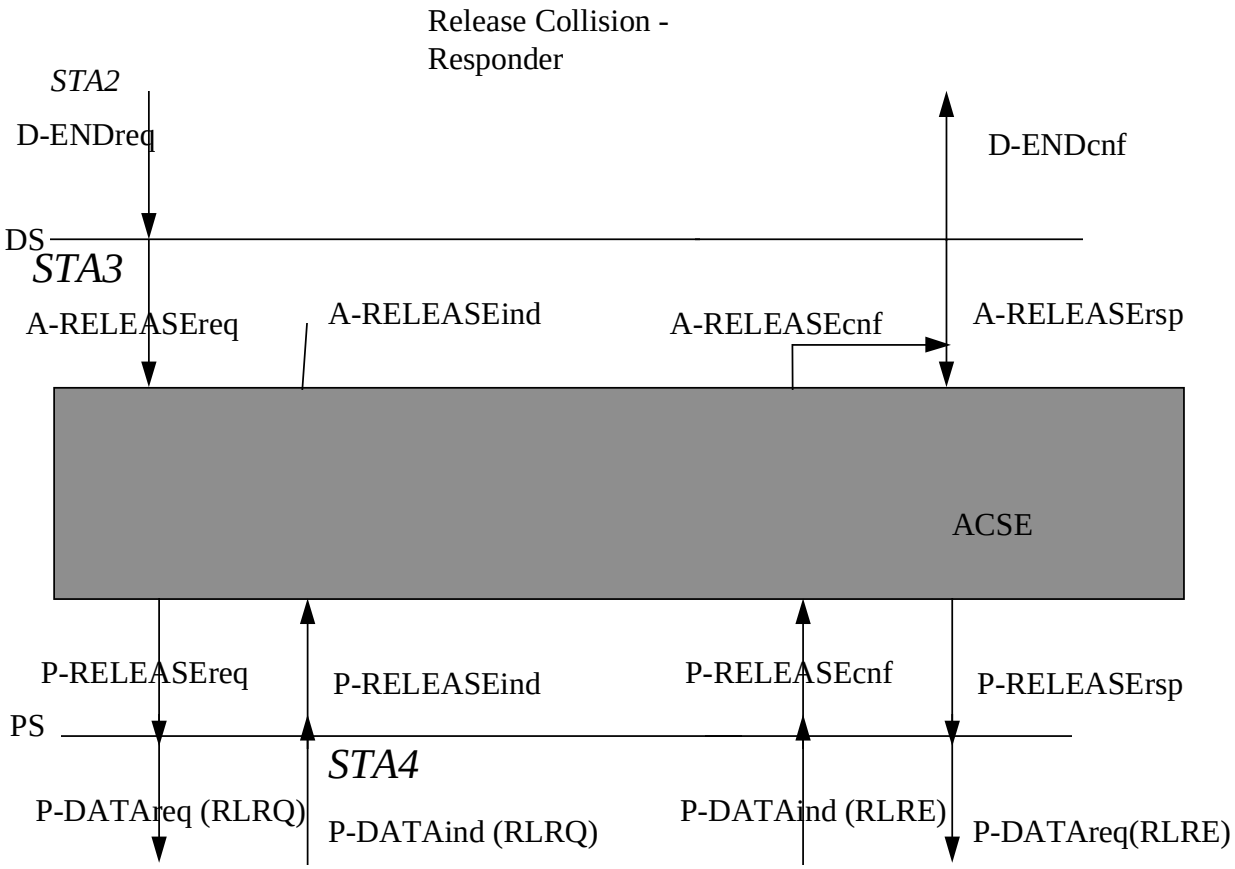
Release Collision -
Responder
STA2
D-ENDreq
D-ENDcnf
DS
STA3
A-RELEASEreq
A-RELEASEind
A-RELEASEcnf
A-RELEASErsp
ACSE
P-RELEASEreq
P-RELEASEind
P-RELEASEcnf
P-RELEASErsp
PS
STA4
P-DATAreq (RLRQ)
P-DATAind (RLRQ)
P-DATAind (RLRE)
P-DATAreq(RLRE)

4.7 Naming and Addressing

4.7.1 Implementation of ULA Construction of Titles and Addresses

The TSAP comprises the IDP (3 octets), the RDP (5 octets), the LDP (ARS (3) + 9 = 12 octets) and the SEL (0-2 octets) fields.

The fields are supplied as follows:
IDP = a priori
RDP = CMA (long TSAP - short TSAP) (Response)
ARS = D-START Called Peer ID (For air initiated requests only)
-- Note work is being done to determine what is required for the ground.
LDP - ARS + SEL = CMA short TSAP (Response)

The application name comprises the AE-title = AP-title + AE-qualifier

The fields are supplied as follows

AP-title is derived from the D-START Called Peer ID (Request)
AE-qualifier is derived from the CMA ASN.1 type APName

The CMA returns TSAP fragments and AE-titles (AP-titles + AE-qualifiers). The TSAP fragments are turned into complete PSAP addresses as follows:

a) If the TSAP fragment comprises only the local domain part (LDP), the Routing Domain Part (RDP) is restored from the previous TSAP fragment - name pair.

b) The Initial Domain Part (IDP) is restored to the prefix of the TSAP

c) No presentation or session selectors are added to the TSAP address.

This results in a complete application name - address mapping.

The received application context is mapped by to the appropriate name and address by matching ATN ASE-type and presentation address.

4.7.2 Guidance on CNS/ATM-1 Naming and Addressing

The following figure shows the registration tree for the ATN CNS/ATM-1 package and shows how it would be expanded to cope with future developments.

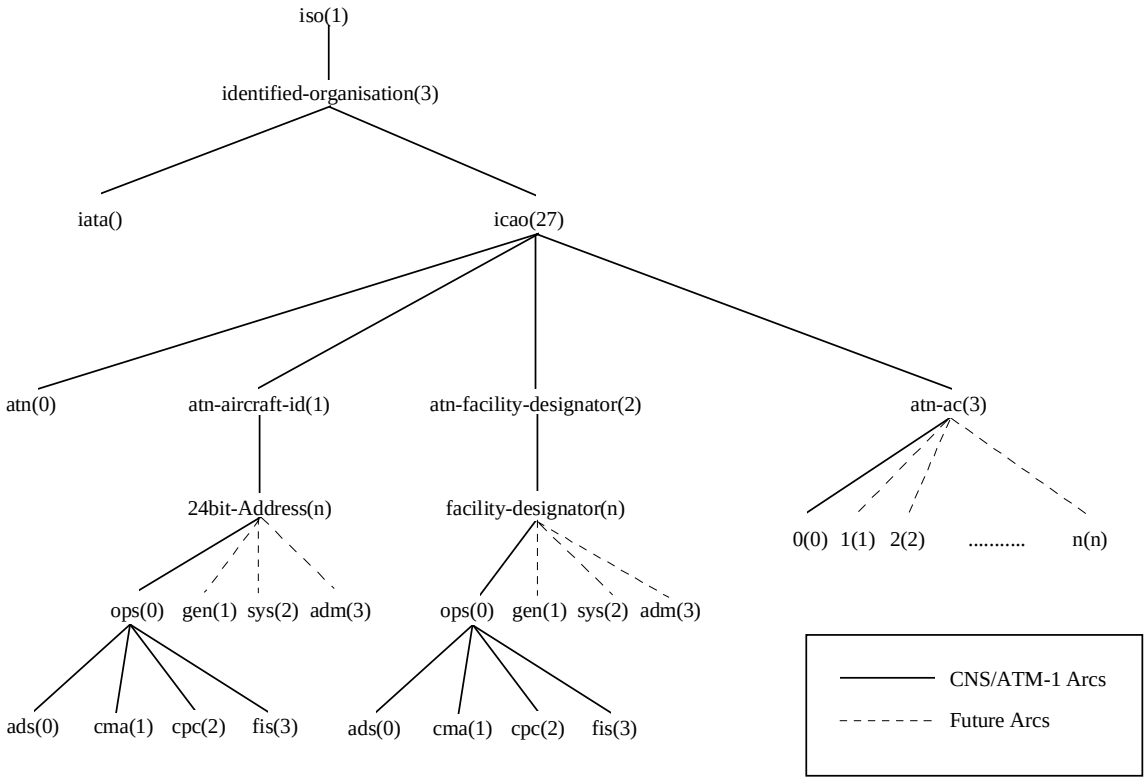

5. SARPs Defect Register

1. In Chapter 1 and Chapter 3, there are guidance figures that use CDSE, which is a CNS/ATM-2 ASE.

2. In Chapter 2, reference is made to the four-character ICAO location identifier, whereas ADSP now proposes use of the eight-character ICAO location identifier.

3. In Chapter 3, the same state is variously called ASSOCIATE PENDING and ASSOCIATION PENDING.

4. In Chapters 4 and 5, the ITU-T defect reports shall also be implemented.

5. In Chapter 6, AE-Title Name Form is described as (Receiver (ISO M, ATN X)).

6. In Annex A, the ICAO and IATA arcs are shown as separate, although they are lexically identical (i.e., they both use the 027 arc). Thus, there is no provision for AINSC naming in the ATN tree.

7. In Annex A, please add SMA (system management application) as AE-qualifier value 8.