

A Management Briefing on

Adapting Voice

For ATM Networks:

An AAL2 Tutorial

A Management Briefing on
Adapting Voice For ATM Networks
An AAL2 Tutorial

Written by

Mike McLoughlin

John O'Neil



General DataComm

TABLE OF CONTENTS

Introduction	1
The ATM Adaptation Layer	2
ATM Adaptation Layer 1 (AAL1)	2
ATM Adaptation Layer 2 (AAL2)	3
<i>AAL2 Common Part Sub-layer (CPS)</i>	4
<i>AAL2 Service Specific Convergence Sub-Layer (SSCS)</i>	5
<i>AAL2 Protocol Efficiency</i>	6
<i>AAL2 Bandwidth Efficiency</i>	6
<i>AAL2 Voice-Over-ATM Trunking Efficiency</i>	8
Summary	10
Glossary	11
References	12

INTRODUCTION

This Paper details the protocol of the newly developed ATM Adaptation Layer 2 (AAL2), and provides examples of the use of AAL2 for Voice-Over-ATM.

A brief overview of the ATM Adaptation Layer 1 (AAL1) protocol is also included to allow differentiation between AAL1 and AAL2. Many readers will be familiar with AAL1, which has been standardized in both the ITU-T and ANSI since 1993, is incorporated in the ATM Forum specifications for Circuit Emulation Services, and is offered by several ATM equipment manufacturers. However, few know the details of the AAL2 protocol. Starting with an overview of the AAL1 protocol should simplify the process of understanding AAL2.

AAL2 had its beginnings in a contribution to Committee T1S1.5 entitled Short Multiplexed AAL (SMAAL) in September, 1995, which was authored by John Baldwin of Lucent. SMAAL was first introduced to the ITU-T at the May 1996 meeting of Study Group 13 in Geneva. At this meeting, AAL2 was initiated under the temporary name of AAL-CU — for composite user. The work on AAL-CU was given high priority within the sub-group of Study Group 13 associated with AAL development. This resulted in arguably the most rapid and stable development of any Recommendation within the ITU-T.

From inception in May 1996 to technical agreement on February 28, 1997, AAL2 was completed in the record time of 9 months. This was primarily due to a concerted effort by the ITU-T membership that had a singular goal in mind: To develop an AAL geared to the support of packetized voice and data over ATM, with full backing of the ATM Forum.

While the ITU-T was developing the protocol for AAL2, input from the ATM Forum VTOA (Voice Telephony Over ATM) working-group substantiated the critical need in the market for an AAL that fully satisfied the requirements for Voice-Over-ATM.

The cooperation between the ATM Forum, which identified market needs, and the protocol experts at the ITU-T resulted in a new AAL that is ideally suited for Voice-Over-ATM applications — AAL2.

AAL2 is defined in the ITU-T Recommendation I.363.2 that was determined at the Study Group 13 meeting in Seoul, Korea in February 1997 and will be approved at the September 1997 Study Group 13 meeting in Toronto.

For information on GDC's solution for VBR Voice-Over-ATM using AAL2, see the GDC paper "APEX Voice Service Module — Product Overview."

THE ATM ADAPTATION LAYER

The ATM Adaptation Layer (AAL) performs functions required by the user, control and management planes and supports the mapping between the ATM layer and the next higher layer. The functions performed in the AAL depend upon the higher layer requirements. In short, the AAL supports all of the functions required to map information between the ATM network, and the non-ATM application that may be using it.

Different adaptation layers exist to carry traffic as diverse as packet-based or isochronous (T1 or E1) over the ATM backbone. AALs are standardized in the ITU-T I.363.x series of Recommendations. The two most commonly implemented are AAL1 (per I.363.1), which supports isochronous transmission — circuit emulation, for example — and AAL5 (per I.363.5), which supports carrying packet data, such as Frame Relay, over ATM.

ATM ADAPTATION LAYER 1 (AAL1)

As defined in ITU-T Recommendation I.363.1, AAL1 provides the following services to the AAL user:

- Transfer of service data units with a constant source bit rate and the delivery of them with the same bit rate
- Transfer of timing information between source and destination
- Transfer of structure information between source and destination
- Indication of lost or errored information not recovered by AAL 1, if needed.

The primary application for AAL1 is circuit emula-

tion, that is, to provide a constant-bit-rate (CBR) service, enabling simplistic isochronous transports of leased-lines across the ATM backbone. To achieve this, AAL1 typically uses the ATM CBR service category definition, which specifies the Peak Cell Rate (PCR), Cell Loss Ratio (CLR), and Cell Delay Variation (CDV) necessary in the ATM network to support the application. This cell flow is independent of information contained within the service rate — i.e., cells continue to flow on the ATM virtual circuit even when there is no traffic.

The ATM Forums' Circuit Emulation Interoperability Specifications Versions 1 and 2 define the overall architecture and specification for this kind of application.

In addition to the constant cell flow using AAL1, the information payload contained within each cell is set by the basic structure of AAL1 (Figure 1). The information payload for AAL1 is always 47 octets, the basic structure used for circuit emulation. Optional structures for AAL1 add additional overhead, reduce the information payload and are used for structured circuit emulation. When assessing the use of AAL1 for Voice-Over-ATM, it is significant to note that the AAL1 protocol has the following limitations:

- Only a single user of the AAL can be supported
- Reducing delay requires significant additional bandwidth
- Bandwidth is used even when there is no traffic
- Voice is always 64K or bundles of 64K ($N \times 64$)
- No standard mechanism in the AAL1 structure for compression, silence detection/suppression, idle channel removal, or CCS (Common Channel Signaling).

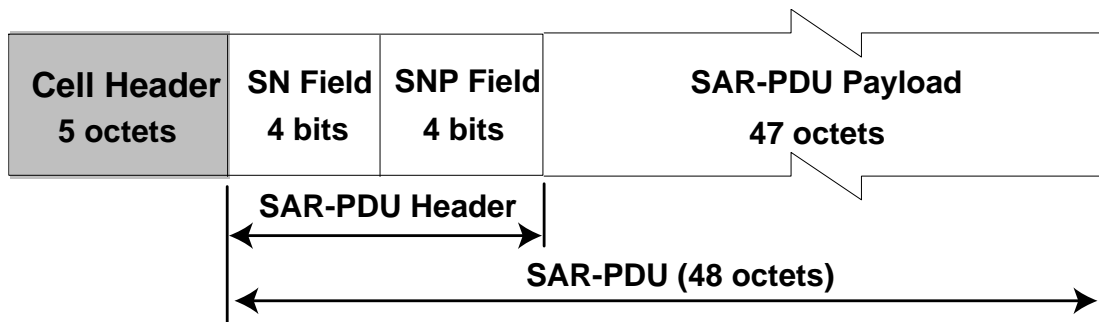


FIGURE 1 — AAL1 SAR-PDU

ATM ADAPTATION LAYER 2 (AAL2)

The basic functions of the AAL2 protocol are consistent with AAL1, in that both enhance the service provided by the ATM layer to support functions required by the next higher layer. However, AAL2 goes beyond AAL1 by defining a structure that includes functions supporting higher layer requirements neither considered or possible within the structure of AAL1.

AAL 2 provides for the bandwidth-efficient transmission of low-rate, short, and variable packets in delay sensitive applications. It enables support for both Variable-Bit-Rate (VBR) and Constant-Bit-Rate (CBR) applications within an ATM network. VBR services enable statistical multiplexing for the higher layer requirements demanded by voice applications, such as compression, silence detection/suppression, and idle channel removal. AAL2's VBR and CBR capabilities mean that network administrators can take traffic variations into account when designing an ATM network and to optimize the network to match traffic conditions.

In addition, AAL 2 enables multiple user channels on a single ATM virtual circuit and varying traffic conditions for each individual user, or channel.

The structure of AAL2 also provides for the packing of short length packets into one (or more) ATM cells, and the mechanisms to recover from transmission errors. In contrast to AAL1, which has a fixed payload, AAL2 offers a variable payload within cells and across cells. This functionality provides a dramatic improvement in bandwidth efficiency over either structured or unstructured circuit emulation using AAL1. See the GDC Paper "Adapting Voice for ATM Networks — A Comparison of AAL1 versus AAL2."

In summary, AAL2 provides the following advantages when compared with AAL1:

- Bandwidth efficiency
- Support for compression and silence suppression
- Support for idle voice channels
- Multiple user channels with varying bandwidth on a single ATM connection
- VBR ATM traffic class

The structure of AAL2, as defined in ITU-T Recommendation I.363.2, is shown in Figure 2.

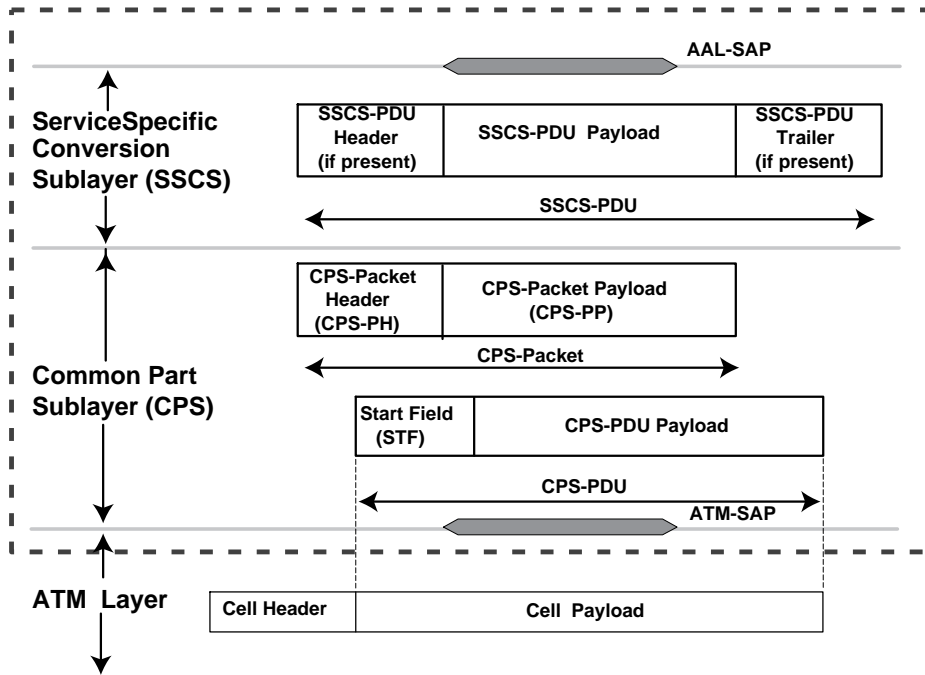


FIGURE 2 — AAL2 STRUCTURE

AAL2 is divided into two sub-layers: the Common Part Sub-layer (CPS) and the Service Specific Convergence Sub-layer (SSCS).

AAL2 Common Part Sub-Layer

Fully defined in I.363.2, the CPS provides the basic structure for identifying the users of the AAL, assembling/disassembling the variable payload associated with each individual user, error correction, and the relationship with the SSCS. Each AAL2 user can select a given AAL-SAP associated with the Quality of Service (QoS) required to transport that individual higher layer application. AAL2 makes use of the service provided by the underlying ATM layer. Multiple AAL connections can be associated with a single ATM layer connection, allowing multiplexing at the AAL layer. The AAL2 user selects the QoS provided by AAL2 through the choice of the AAL-SAP used for data transfer.

AAL2's CPS possesses the following characteristics:

- It is defined on an end-to-end basis as a concatenation of AAL2 channels.
- Each AAL2 channel is a bi-directional virtual channel, with the same channel identifier value used for both directions.
- AAL2 channels are established over an ATM layer Permanent Virtual Circuit (PVC), Soft Permanent Virtual Circuit (SPVC) or Switched Virtual Circuit (SVC).

The multiplexing function in the CPS merges several streams of CPS packets onto a single ATM connection.

The format of the CPS packet is shown in Figure 3.

Key fields of the CPS packet are the Channel Identifier (CID), the Length Indicator (LI), and the User-to-User Indication (UUI) fields. These are defined below.

CID Field: Uniquely identifies the individual user channels within the AAL2, and allows up to 248 individual users within each AAL2 structure. Coding of the CID field is shown below.

Value	Use
0	Not Used
1	Reserved for Layer Management Peer-to-Peer Procedures
2-7	Reserved
8-255	Identification of AAL2 User (248 total channels)

LI Field: Identifies the length of the packet payload associated with each individual user, and assures conveyance of the variable payload. The value of the LI is one less than the packet payload and has a default value of 45 octets, or may be set to 64 octets.

UUI Field: Provides a link between the CPS and an appropriate SSCS that satisfies the higher layer application. Different SSCS protocols may be defined to support specific AAL2 user services, or groups of services. The SSCS may also be null.




FIGURE 3 — FORMAT OF THE AAL2 CPS PACKET

AAL2 PROTOCOL EFFICIENCY

An important aspect of AAL2 is Packet Fill Delay. Packet Fill Delay allows the network operator to set a time period during which AAL2 PDUs are assembled and then segmented into cells. The setting of Packet Fill Delay allows the operator to alter the delay characteristic of voice into the ATM adaptation phase of AAL2. Different voice circuits may have different minimum delay requirements, and it is important to be able to trade off delay and efficiency within the Voice-Over-ATM environment.

Table 1 shows the relationship between the Packet Fill Delay and the AAL2 PDU payload required to support a single voice channel. For example, for a single 32K ADPCM voice channel with a Packet Fill Delay setting of 2 ms, each AAL2 PDU goes out with an 8 byte payload supporting the voice channel. If the value of Packet Fill Delay is doubled to 4 ms, then each 32K ADPCM voice channel will fill 16 bytes in every AAL2 frame before being sent into the ATM network.

Table 2 lists both PCM and 32K ADPCM channels with various Packet Fill Delay parameters to provide a basic understanding of the protocol efficiency for AAL2. However, while evaluating the protocol pay-

load and overhead may have meaning for a statistician, it is totally removed from the real benefit of AAL2, which is significant reduction in bandwidth requirements for a given application. To allow multiplexing within an AAL, additional overhead is required, but the net result is vastly improved bandwidth efficiency.

AAL2 BANDWIDTH EFFICIENCY

Due to the complexity of dealing with both fixed and statistical compression in a voice channel (for example, ADPCM and silence suppression) and the further complication of packing these voice channels into ATM cells, it is difficult to provide a simple formula to calculate the theoretical ATM bandwidth needed to support a voice service inside the ATM network. However, the following examples help to illustrate what bandwidth efficiency may be expected from an AAL2 VBR Voice Service.

Figure 5 shows how an AAL2 PDU supporting six 32K ADPCM channels with a Packet Fill Delay value of 4 ms would be structured.

Further examples of the structure of an AAL2 PDU with varying values for Packet Fill Delay are shown in Figures 6, 7 and 8.

TABLE 1 — PACKET FILL DELAY VERSUS AAL2 PDU PAYLOAD

Packet Fill Delay	CPS Header	32K ADPCM	64K PCM
2 ms	3 bytes	8 byte payload	16 byte payload
4 ms	3 bytes	16 byte payload	32 byte payload
6 ms	3 bytes	24 byte payload	48 byte payload
8 ms	3 bytes	32 byte payload	64 byte payload

TABLE 2 — AAL2 PROTOCOL EFFICIENCY

	Channel Header	Payload	Efficiency
32K ADPCM (4ms)	3 bytes	16 byte payload	84%
32K ADPCM (8 ms)	3 bytes	32 byte payload	91%
PCM (4 ms)	3 bytes	32 byte payload	91%
Default LI	3 bytes	45 byte payload	94%
Maximum LI	3 bytes	64 byte payload	96%




FIGURE 5 — SIX CHANNELS OF 32K ADPCM WITH 4 MS PACKET FILL DELAY

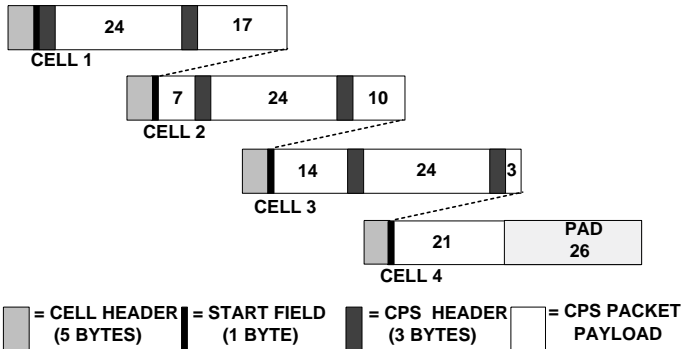


FIGURE 6 — SIX CHANNELS OF 32K ADPCM WITH 6 MS PACKET FILL DELAY

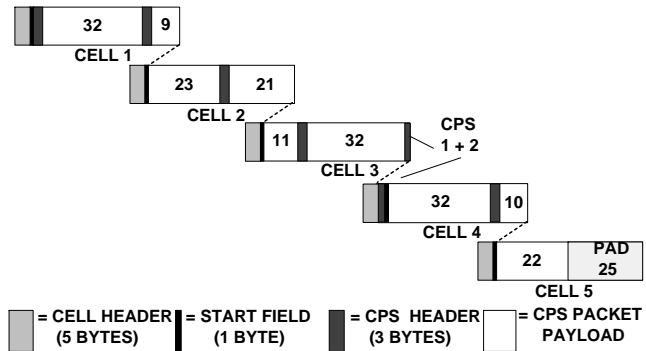


FIGURE 7 — SIX CHANNELS OF 32K ADPCM WITH 8 MS PACKET FILL DELAY

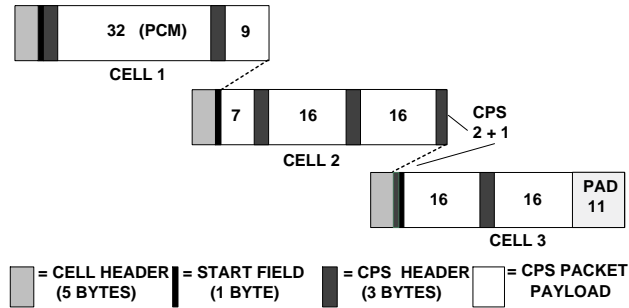


FIGURE 8 — ONE CHANNEL OF 64 K PCM AND FIVE OF 32K ADPCM WITH 4 MS PACKET FILL DELAY

TABLE 3 — BANDWIDTH REQUIRED USING AAL2 FOR SIX VOICE CHANNELS

# of Channels	Channel Rate	AAL	Packet Fill Delay (ms)	Silence Suppression	Bandwidth (Kbps)
6	64K	2	6	No	495
6	32K	2	4	No	318
6	32K	2	6	No	283
6	32K	2	8	No	265
6	64K	2	6	Yes	198
6	32K	2	4	Yes	128
6	32K	2	6	Yes	113
6	32K	2	8	Yes	106

Table 3 shows the use of AAL2 for six channels given the parameters of basic compression factor (none or 32K ADPCM), the encoding delay value, Packet Fill Delay, and support for silence suppression on or off (assuming 50% silence). The absolute bandwidth required to support the service within an ATM trunk is shown in the last column.

AAL2 Voice-Over-ATM Trunking Efficiency

Another possible way to view the efficiency of an AAL2 connection is to identify how many voice channels may be carried over a fixed bandwidth ATM trunk between ATM network elements.


Figure 9 shows the number of voice channels that can be carried over a T1 ATM trunk using AAL2. The X-axis represents the value of Packet Fill Delay, and the Y-axis shows the number of voice channels carried. Plots are shown for both 64K PCM and 32K ADPCM cases in the AAL2 frame.

Note that using 64K PCM inside an AAL2 frame, a maximum of 18 channels can be supported with a

delay of up to 8 ms. But for the 32K ADPCM encoded channels, a maximum of 35 channels may be supported.


If we include silence suppression, significant gains can be seen. Assuming that a voice circuit contains 50% silence and that 20% of all channels are idle at any one time, we see that the number of 64 K PCM channels supported by AAL2 more than doubles from 18 to 45 with 8 ms of Packet Fill Delay by adding silence detection/suppression and idle channel removal. (Figure 10) When we add 32K ADPCM compression, the T1 trunk can accommodate up to 87 high-quality, multiplexed voice channels.

Finally, if we add the fact that when these voice channels are not busy (e.g. overnight) the bandwidth being used for voice is available for other applications (remote server archiving or software download, for example), then ATM networks begin to emerge as the only viable underlying technology for efficient Wide Area Multiservice Networking.



Packet Fill (ms)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
32K —	22	28	31	33	34	34	35	35	36	36	36	37	37	37	37	37
64K —	14	16	17	17	18	18	18	18	18	18	19	19	19	19	19	19

FIGURE 9 — NUMBER OF AAL2 VOICE CHANNELS IN A DS1 TRUNK WITH VARIABLE PACKET FILL DELAY (ms) WITHOUT SILENCE SUPPRESSION



Packet Fill (ms)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
32K —	55	70	77	82	85	85	87	87	90	90	90	92	92	92	92	92
64K —	35	40	42	42	45	45	45	45	45	45	47	47	47	47	47	47

FIGURE 10 — NUMBER OF AAL2 SILENCE SUPPRESSED VOICE CHANNELS IN A DS1 TRUNK WITH VARIABLE PACKET FILL DELAY (ms)

SUMMARY

This paper introduced the reader to the new ATM Adaptation Layer 2 (AAL2) as defined in the ITU-T Recommendation I.363.2, including protocol details, examples of the use of AAL2, and unique features afforded with the completion of an SSCS for trunking.

It should provide a basic understanding of AAL2 and identify the benefits associated with this new adaptation layer. By far the most important benefit of AAL2 is the ability to substantially reduce bandwidth requirements for supporting voice on ATM networks, and the inherent flexibility to add features via the SSCS sub-layer structure.

For a detailed comparison of the benefits and improvements in efficiency afforded by using AAL2 for Voice-Over-ATM in lieu of using AAL1, see the GDC Paper “Adapting Voice For ATM Networks — A Comparison of AAL1 Versus AAL2”. After reading

this companion paper, we are sure that you will add the support of AAL2 to your list of mandatory ATM product requirements, and in all likelihood shift AAL1 from a mandatory to an optional requirement for the support of Voice-Over-ATM.

Recognizing the benefits of AAL2 for Voice-Over-ATM, GDC has not only strived for standards compliance, but also developed the Voice Service Module (VSM) as an addition to our APEX ATM Product Family to support both AAL2 and AAL1. For detailed information on this GDC product, see the paper “APEX Voice Service Module — Product Overview.”

To understand the economic benefits of GDC’s VSM, see the **TeleChoice** case studies “Voicing The Case For ATM Enterprise Networks,” “Voicing The Case For ATM Value-Added Services,” and “Voicing The Case For ATM Carrier Backbones.”

GLOSSARY

AAL1	ATM Adaptation Layer 1 defined in ITU-T I.363.1. The type of ATM adaptation principally used for circuit emulation services over an ATM network.	CLR	Cell Loss Ratio
AAL2	ATM Adaptation Layer 2 defined in ITU-T I.363.2. A new type of ATM adaptation used for variable-bit-rate Voice-Over-ATM services.	CPS	Common Part Sub-Layer
AAL5	ATM Adaptation Layer 5 defined in ITU-T I.363.5. The type of ATM adaptation principally used for frame and packet transport over an ATM network.	GDC	General DataComm
AAL-CU	ATM Adaptation Layer-Composite User	ITU-T	International Telecommunication Union - Telecommunications Standardization Sector
ADPCM	Adaptive Differential Pulse Code Modification. A compression algorithm for voice as defined in ITU-T G.726.	Kbps	Kilobits per second
ANSI	American National Standards Institute	LI	Length Indicator
ATM	Asynchronous Transfer Mode. The cell relay service that transfers mixed traffic types over a common communications medium.	ms	Milliseconds
CBR	Constant-Bit-Rate	PCM	Pulse Code Modulation. The basic modulation scheme for transporting voice channels in 64 Kbps timeslots.
CCS	Common Channel Signalling	PCR	Peak Cell Ratio
CDR	Cell Delay Variation	PDU	Protocol Data Unit
CES	Circuit Emulation Service as specified by the ATM Forum.	SMAAL	Short Multiplexed ATM Adaptation Layer
CID	Channel Identifier	SSCS	Service Specific Convergence Sub-Layer
		UUI	User-to-User Indication
		VBR	Variable-Bit-Rate
		VSM	Voice Service Module. GDC's new, leading edge VBR Voice-Over-ATM module for the APEX family.
		VTOA	Voice Telephony Over ATM. A working group of the ATM Forum

REFERENCES

ITU-T Recommendation I.363.1: B-ISDN ATM Adaptation Layer (AAL) Specification Type 1

ITU-T Recommendation I.363.2: B-ISDN ATM Adaptation Layer Type 2 Specification

ATM Forum Circuit Emulation Service Version 2 Interoperability Specification (af-vtoa-0078.000)

ABOUT THE COMPANY

General DataComm Inc. (GDC), is an international communications company headquartered in Middlebury, Connecticut, USA. General DataComm designs, markets and supports networks and networking products which integrate voice, data, image, video and LAN applications for national and multi-national companies, governments and communications service providers worldwide.

In the United States, General DataComm has sales and service offices throughout the country. Elsewhere, the company has subsidiary operations in Canada, UK, France, Mexico, Russia and Australia. To augment its vast network of international distributors, OEMs and licensees, General DataComm has regional sales and support offices in Japan, Hong Kong, Singapore, Belgium, France and Germany. General DataComm is represented on all continents, in over sixty countries around the world.

As a leading provider of networks and networking products worldwide, General DataComm offers the broadest range of products of any single communications company. These products include: ATM products, access devices (analog and digital data sets), derivation equipment (traditional multiplexers and transport management systems), and network management systems. In all product areas, General DataComm enjoys a reputation for innovation, quality and superior performance.

WORLD HEADQUARTERS
Middlebury, Connecticut USA 06762-1299
Tel: 1-203-574-1118
Fax: 1-203-738-8507
1-203-738-9518 (GDC International)
<http://www.gdc.com>

U.S. Business Systems Sales Offices

Atlanta, GA (770) 955-0682
Boston, MA (617) 622-5900
Chicago, IL (630) 261-0670
Cleveland, OH (216) 328-2044
Dallas, TX (972) 406-4800
Denver, CO (303) 782-3000
Detroit, MI (810) 549-4110
Hartford, CT (203) 574-1118
Houston, TX (713) 779-7879
Los Angeles, CA (310) 348-5200
New York, NY (212) 248-7220
Pennsauken, NJ (609) 663-0755
San Francisco, CA (510) 769-4500
Washington, DC (301) 595-0300

U.S. Telecomm Sales Offices

Atlanta, GA (770) 955-0682
Chicago, IL (630) 261-0670
Dallas, TX (972) 406-4800
Denver, CO (303) 782-3600
Detroit, MI (810) 549-4110
Honolulu, HI (808) 235-2319
Los Angeles, CA (818) 506-8897
Kirkland, WA (425) 820-1471
Minneapolis, MN (612) 935-7765
Pennsauken, NJ (609) 663-0755
San Francisco, CA (510) 769-4500
St. Louis, MO (314) 537-1333
Washington, DC (301) 595-0300

For the name of your U.S. Distributor contact:
(800) 523-1737

For 24-hour delivery, call GDC
QUICKSHIPPERS® at 1-800-432-2228

U.S. Government Sales

Washington, DC (703) 205-2960

SUBSIDIARIES

Australia

Tel: 61-2-9956-5099
Fax: 61-2-9956-5083

Canada

Tel: 416-498-5100
Fax: 416-498-0248

France

Tel: 33-1-4762-6200
Fax: 33-1-4762-9696

Germany

Tel: 49-69-950840
Fax: 49-69-5073259

Mexico

Tel: 52-5-645-2238
Fax: 52-5-645-3976

Russia

Tel: 7-812-325-1085
Fax: 7-812-325-1086

United Kingdom

Tel: 44-1189-774-8688
Fax: 44-1189-774-871

INTERNATIONAL REGIONAL OFFICES

Asia

Singapore
Tel: 65-735-2123
Fax: 65-735-6889

Hong Kong

Tel: 852-25265511
Fax: 852-25259944

China

Tel: 86-10-6500-6589
Fax: 86-10-6500-6590

Japan

Tel: 81-3-3862-1730
Fax: 81-3-3862-6356

Europe/Africa/Middle East

United Kingdom
Tel: 44-1189-774-8688
Fax: 44-1189-774-871

Latin America

Argentina
Tel: 54-1-315-6086
Fax: 54-1-312-2819

Brazil

Tel: 55-11-535-0232
Fax: 55-11-542-0547

Miami, Florida

Tel: 1-954-724-3511
Fax: 1-954-724-5397

For the name and address of your local distributor, contact your nearest area office or General DataComm International Headquarters

All specifications subject to change without notice.

© General DataComm (1997)

All Rights Reserved

® Registered trademark of General DataComm Industries, Inc.

General DataComm, GDC, and the GDC logo are trademarks of General DataComm, Inc.

All other trademarks and registered trademarks are the property of their respective owners.



Printed in USA 00570-7/97SP



General DataComm

WORLD HEADQUARTERS
Middlebury, Connecticut USA 06762-1299
Tel: 1-203-574-1118