



**FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA**

**CENTRO DE INFORMACIÓN Y DOCUMENTACIÓN
" ING. BRUNO MASCANZONI "**

El Centro de Información y Documentación Ing. Bruno Mascanzoni tiene por objetivo satisfacer las necesidades de actualización y proporcionar una adecuada información que permita a los ingenieros, profesores y alumnos estar al tanto del estado actual del conocimiento sobre temas específicos, enfatizando las investigaciones de vanguardia de los campos de la ingeniería, tanto nacionales como extranjeras.

Es por ello que se pone a disposición de los asistentes a los cursos de la DECFI, así como del público en general los siguientes servicios:

- **Préstamo interno.**
- **Préstamo externo.**
- **Préstamo interbibliotecario.**
- **Servicio de fotocopiado.**
- **Consulta a los bancos de datos: librunam, seriunam en cd-rom.**

Los materiales a disposición son:

- **Libros.**
- **Tesis de posgrado.**
- **Publicaciones periódicas.**
- **Publicaciones de la Academia Mexicana de Ingeniería.**
- **Notas de los cursos que se han impartido de 1988 a la fecha.**

En las áreas de ingeniería industrial, civil, electrónica, ciencias de la tierra, computación y, mecánica y eléctrica.

El CID se encuentra ubicado en el mezzanine del Palacio de Minería, lado oriente.

El horario de servicio es de 10:00 a 14:30 y 16:00 a 17:30 de lunes a viernes.



**FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA**

A LOS ASISTENTES A LOS CURSOS

Las autoridades de la Facultad de Ingeniería, por conducto del jefe de la División de Educación Continua, otorgan una constancia de asistencia a quienes cumplan con los requisitos establecidos para cada curso.

El control de asistencia se llevará a cabo a través de la persona que le entregó las notas. Las inasistencias serán computadas por las autoridades de la División, con el fin de entregarle constancia solamente a los alumnos que tengan un mínimo de 80% de asistencias.

Pedimos a los asistentes recoger su constancia el día de la clausura. Estas se retendrán por el período de un año, pasado este tiempo la DECFI no se hará responsable de este documento.

Se recomienda a los asistentes participar activamente con sus ideas y experiencias, pues los cursos que ofrece la División están planeados para que los profesores expongan una tesis, pero sobre todo, para que coordinen las opiniones de todos los interesados, constituyendo verdaderos seminarios.

Es muy importante que todos los asistentes llenen y entreguen su hoja de inscripción al inicio del curso, información que servirá para integrar un directorio de asistentes, que se entregará oportunamente.

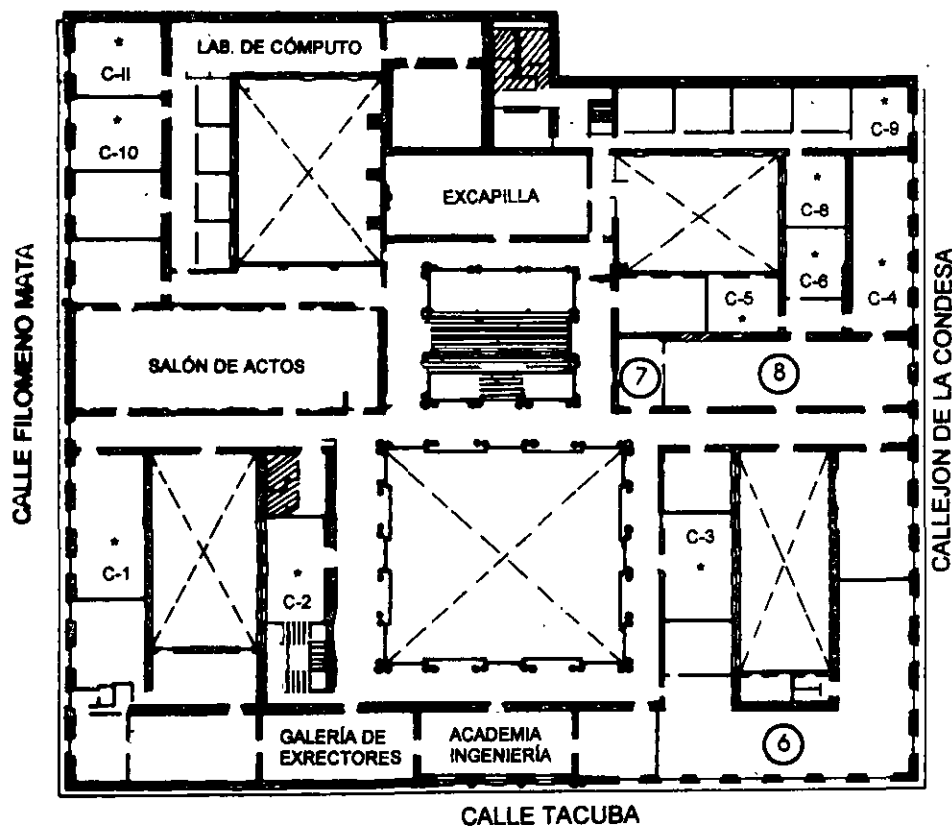
Con el objeto de mejorar los servicios que la División de Educación Continua ofrece, al final del curso deberán entregar la evaluación a través de un cuestionario diseñado para emitir juicios anónimos.

Se recomienda llenar dicha evaluación conforme los profesores impartan sus clases, a efecto de no llenar en la última sesión las evaluaciones y con esto sean más fehacientes sus apreciaciones.

Atentamente

División de Educación Continua.

PALACIO DE MINERÍA



1er. PISO

GUÍA DE LOCALIZACIÓN

1. ACCESO
2. BIBLIOTECA HISTÓRICA
3. LIBRERÍA UNAM
4. CENTRO DE INFORMACIÓN Y DOCUMENTACIÓN "ING. BRUNO MASCANZONI"
5. PROGRAMA DE APOYO A LA TITULACIÓN
6. OFICINAS GENERALES
7. ENTREGA DE MATERIAL Y CONTROL DE ASISTENCIA
8. SALA DE DESCANSO

SANITARIOS

* AULAS




DIVISIÓN DE EDUCACIÓN CONTINUA
FACULTAD DE INGENIERÍA U.N.A.M.
CURSOS ABIERTOS


DIVISIÓN DE EDUCACIÓN CONTINUA



PALACIO DE MINERIA



PLANTA BAJA



MEZZANINNE



DIVISION DE EDUCACION CONTINUA
FACULTAD DE INGENIERIA, UNAM
CURSOS ABIERTOS



CURSO: *CC070 Redes de Alto Desempeño: Fast y Giga Ethernet, FDDI-II, "Switching"...*
FECHA: *23 al 27 de octubre del 2000*

EVALUACIÓN DEL PERSONAL DOCENTE
 (ESCALA DE EVALUACION: 1 A 10)

CONFERENCISTA	DOMINIO DEL TEMA	USO DE AYUDAS AUDIOVISUALES	COMUNICACION CON EL ASISTENTE	PUNTUALIDAD
Ing. Saúl S. Magaña Cisneros				

Promedio _____

EVALUACIÓN DE LA ENSEÑANZA

CONCEPTO	ICALIF
ORGANIZACION Y DESARROLLO DEL CURSO	
GRADO DE PROFUNDIDAD DEL CURSO	
ACTUALIZACION DEL CURSO	
APLICACION PRACTICA DEL CURSO	

Promedio _____

EVALUACIÓN DEL CURSO

CONCEPTO	ICALIF
CUMPLIMIENTO DE LOS OBJETIVOS DEL CURSO	
CONTINUIDAD EN LOS TEMAS	
CALIDAD DEL MATERIAL DIDACTICO UTILIZADO	

Promedio _____

Evaluación total del curso _____

Continúa...2

1. ¿Le agradó su estancia en la División de Educación Continua?

SI

NO

Si indica que "NO" diga porqué:

2. Medio a través del cual se enteró del curso:

Periodico <i>La Jornada</i>	
Folleto anual	
Folleto del curso	
Gaceta UNAM	
Revistas técnicas	
Otro medio (Indique cuál)	

3. ¿Qué cambios sugeriría al curso para mejorarlo?

4. ¿Recomendaría el curso a otra(s) persona(s) ?

SI

NO

5. ¿Qué cursos sugiere que imparta la División de Educación Continua?

6. Otras sugerencias:




**FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA**

**DIPLOMADO EN REDES DE COMPUTADORAS
Y TELECOMUNICACIONES**

MODULO IV

**REDES DE ALTO DESEMPEÑO:
FAST Y GIGA ETHERNET, FDDI-II, "SWITCHING",
ATM Y FRAME RELAY**

OCTUBRE DEL 2000




DIVISION DE EDUCACION CONTINUA DE
LA FACULTAD DE INGENIERIA

DIPLOMADO EN REDES DE
COMPUTADORAS (LAN, WAN y GAN)

Redes de Alto Desempeño :
FAST y GIGA ETHERNET, FDDI-II,
ATM y FRAMERELAY

Coord. Académico: Ing. Saúl S. Magaña Cisneros



DIPLOMADO EN REDES DE
COMPUTADORAS (LAN, WAN Y GAN)

REDES DE ALTO DESEMPEÑO: FAST y GIGA
ETHERNET FDDI-II, SWITCHING, ATM y
FRAMERELAY

Presentación

Coord. Académico: Ing. Saúl S. Magaña Cisneros

Redes de Alto Desempeño: FAST y GIGA ETHERNET, FDDI-II, "SWITCHING", ATM Y FRAME RELAY

Temas

1

FAST ETHERNET y
GIGA ETHERNET

2

FDDI, FDDI II

3

SWITCHES

4

ATM

5

FRAME RELAY

6

APLICACIONES

TEMARIO

1.- FAST ETHERNET y GIGA ETHERNET

- ⊕ Introducción
- ⊕ Características de 100BaseT
- ⊕ Estándares y Normalización
- ⊕ Tipos de cableado
- ⊕ Características de los dispositivos Fast-Ethernet
- ⊕ Características de los dispositivos GIGA-Ethernet
- ⊕ Redes Conmutadas
- ⊕ Alternativas de implementación

2.- FDDI, FDDI II

- ⊕ Introducción
- ⊕ Fibras ópticas
- ⊕ Backbones
- ⊕ Antecedentes de FDDI y FDDI-II
- ⊕ Características
- ⊕ Funcionamiento
- ⊕ Dispositivos
- ⊕ Normalización

3.- SWITCHES

- ⊕ Introducción
- ⊕ Características
- ⊕ Tecnologías Store and Forward y Cut-Through
- ⊕ Switches ATM y Switches Ethernet

4.- ATM

- ⊕ Introducción
- ⊕ Componentes
- ⊕ Servicios
- ⊕ Estructura de la celda
- ⊕ Modelo B-ISDN
- ⊕ Niveles de adaptación, convergencia y físico
- ⊕ Aplicaciones y casos de estudio

5.- FRAME RELAY

- ⊕ Tecnologías antecesoras
- ⊕ Terminología y funcionamiento
- ⊕ Estructura de frame
- ⊕ Administración de la congestión
- ⊕ Técnicas de reducción de tráfico
- ⊕ Interfaces de administración local
- ⊕ Estándares
- ⊕ Aplicaciones y casos de estudio

6.- APLICACIONES

- ⊕ Redes Virtuales
- ⊕ Redes Multimedia
- ⊕ Video Conferencia
- ⊕ Integración total de Redes; LAN=MAN=WAN=GAN

Redes de Alto Desempeño: FAST y GIGA ETHERNET, FDDI-II, 'SWITCHING', ATM Y FRAME RELAY

Temas

1 FAST ETHERNET y
GIGA ETHERNET

2 FDDI, FDDI.II

3 SWITCHES

4 ATM

5 FRAME RELAY

6 APLICACIONES

PRESENTACION

La constante evolución en las tecnologías de las redes de cómputo, las comunicaciones y las telecomunicaciones ha tenido como objetivo central incrementar su rendimiento, esto es crear **REDES DE ALTO DESEMPEÑO** para poder satisfacer las nuevas necesidades de los usuarios: Transmisión de grandes volúmenes de información, de datos, de voz, de vídeo a altas velocidades y cubriendo grandes distancias, permitiendo que el que hacer del hombre en este campo, cada día acorte el tiempo, mejore la seguridad en sus aplicaciones e incremente su productividad.

Las nuevas redes de alto rendimiento serán conformadas por enlaces locales basados en Fast y Giga Ethernet o FDDI-II, las comunicaciones entre redes estarán sustentadas en los servicios de **Cell Relay** que derivan en la tecnología **ATM** y los enlaces remotos soportados por Freme-Relay, todo integrado en poderosos switches de niveles 2 y 3. Estas tecnologías de vanguardia, el día de hoy nos permiten alcanzar velocidades de transmisión de 622 mbps y "backplanes" de 4Gbps y su desempeño se seguirá incrementando

La marcada evolución en la tecnología de las Redes a ido acompañada de un alto desarrollo en los medios comunicación como hoy lo son los enlaces basados en fibras ópticas y cables telefónicos de altas velocidades como FTP y UTP niveles 6 y 7 además de los servicios ofrecidos por las compañías telefónicas como **ISDN** y **B-ISDN** los cuales nos ofrecen integración de múltiples servicios (**voz, datos, imagen y sonido**) gracias a sus amplios anchos de banda. Combinando ambas innovaciones, surgen fuertemente a partir de 1995, las **REDES DE ALTO DESEMPEÑO** que definen a las Redes de cuarta generación.

Las **REDES DE ALTO DESEMPEÑO** con sus elementos de comunicación, implican una serie de tecnologías y arquitectura modemas y avanzadas, que generan la necesidad del conocimiento y dominio de las mismas, y esto es imperante. Se requiere por lo tanto, de especialistas y ejecutivos

bien capacitados y bien informados respectivamente, para un soporte técnico y toma de decisiones adecuados en este profundo y apasionante campo de las Redes. Conscientes esta necesidad, ofrecemos este curso como un módulo más del Diplomado, y/o como una oportunidad de actualización, tratando de lograr los siguientes

OBJETIVOS

Introducir a los participantes en las tecnologías de los Servicios Integrados de Redes Digitales de Banda Ancha (B-ISDN) y dar a conocer los nuevos estándares de las tecnologías de redes de alto desempeño

Que el participante conozca los antecedentes y conceptos de las tecnologías Fast y Giga Ethernet, FDDI-II, "SWITCHING", CELL RELAY, ATM y FRAME RELAY para poder aplicarlo para la toma de decisiones o la implantación de estas tecnologías, según las condiciones del mercado mexicano.

A QUIEN VA DIRIGIDO

A todos aquellos profesionales y profesionistas que por sus necesidades laborales, estén involucrados con las Redes de Cómputo y requieran actualizarse en las Redes de Alto Desempeño, y a los Ejecutivos que necesiten bases técnicas en su responsabilidad de toma de decisiones.


REQUISITOS

Los participantes que estén sustentando el diplomado haber cursado al menos los módulos I y II

Para los participantes que tomen este modulo como un curso abierto; es necesario tener un buen nivel en microcomputacion y conocimientos avanzados en redes de computadoras y comunicaciones

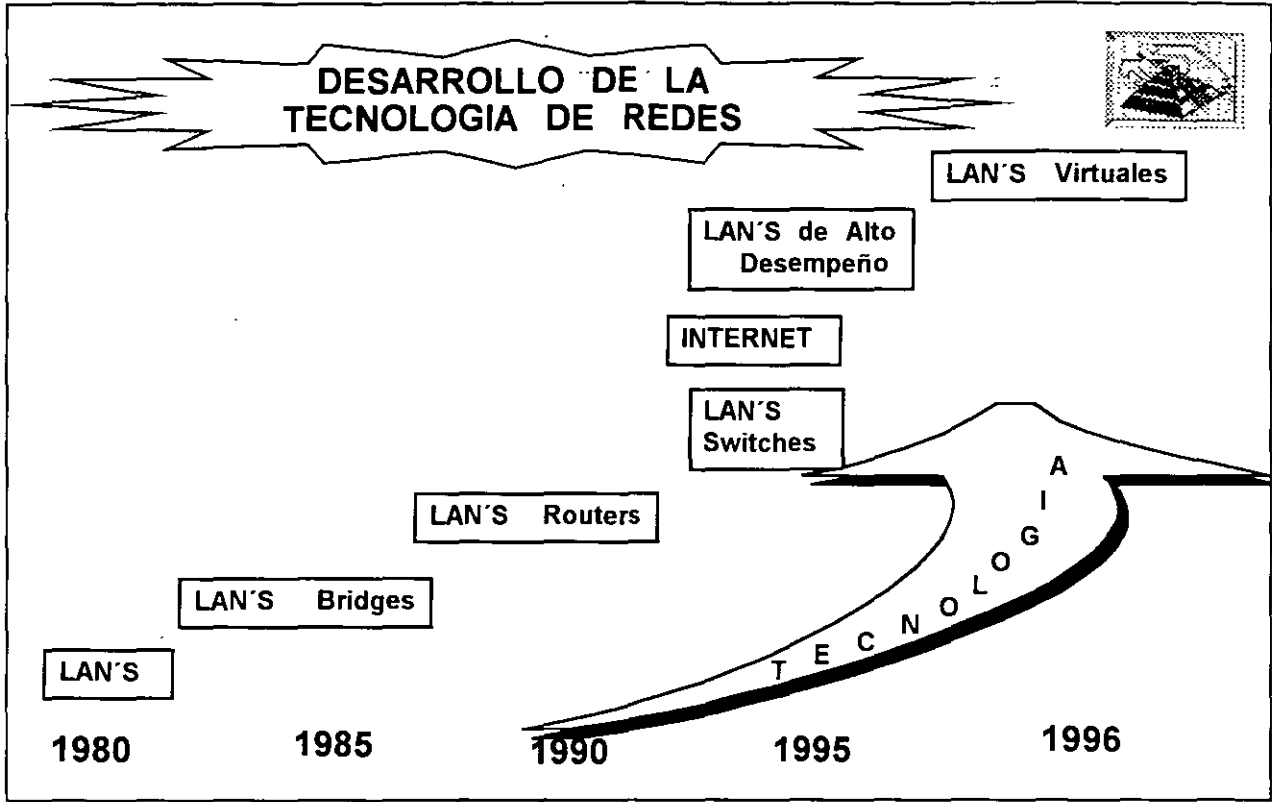


NUEVAS TECNOLOGÍAS DE REDES DE COMPUTADORAS



1995

Notas:




Notas:

Redes de alto desempeño



- FDDI, FDDI - II
- FAST ETHERNET
- TECNOLOGIA SWITCHING
- ATM
- FRAME RELAY
- B - ISND



- REDES VIRTUALES
- REDES MULTIMEDIA VIDEOCONFERENCIAS

REDES

LAN = MAN = WAN = GAN

Notas:



COMUNICACION DIGITAL

BANDA BASE

BANDA ANCHA


Notas:



BANDA BASE

Características:

- Un solo canal
- Bajo costo
- Se modula y demodula la señal
- Utilizada por los estándares actuales de REDES locales




Notas:



BANDA ANCHA

Características:

- Varios Canales Paralelos
- Multiplexaje por Frecuencia
- Un canal de Transmisión
- ← Un Canal de Recepción



Notas:



SERVICIOS CONMUTADOS DE ALTA VELOCIDAD

Alta Velocidad:

- ☐ ISDN Integrated Service Digital Network
- ☐ B-ISDN Broadband-Integrated Service Digital Network


Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



ISDN

Acceso a los servicios de telecomunicaciones sin ISDN




Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



ISDN

Acceso a los servicios de telecomunicaciones con ISDN




Notas:



ISDN

Acceso Básico




Notas:



ISDN

Acceso Primario



Notas:



ISDN Velocidades

Canal	Velocidad de Transmisión	Asociado A
B	64 Kbps	ISDN
D	16 Kbps y 64 Kbps	ISDN
E	64 Kbps	ISDN
H0	384 Kbps = 6B	BISDN
H11	1536 kbps = 24B	BISDN
H12	1920 Kbps = 30B	BISDN
.		
.		
.		
H4	120 a 140 Kbps	BISDN

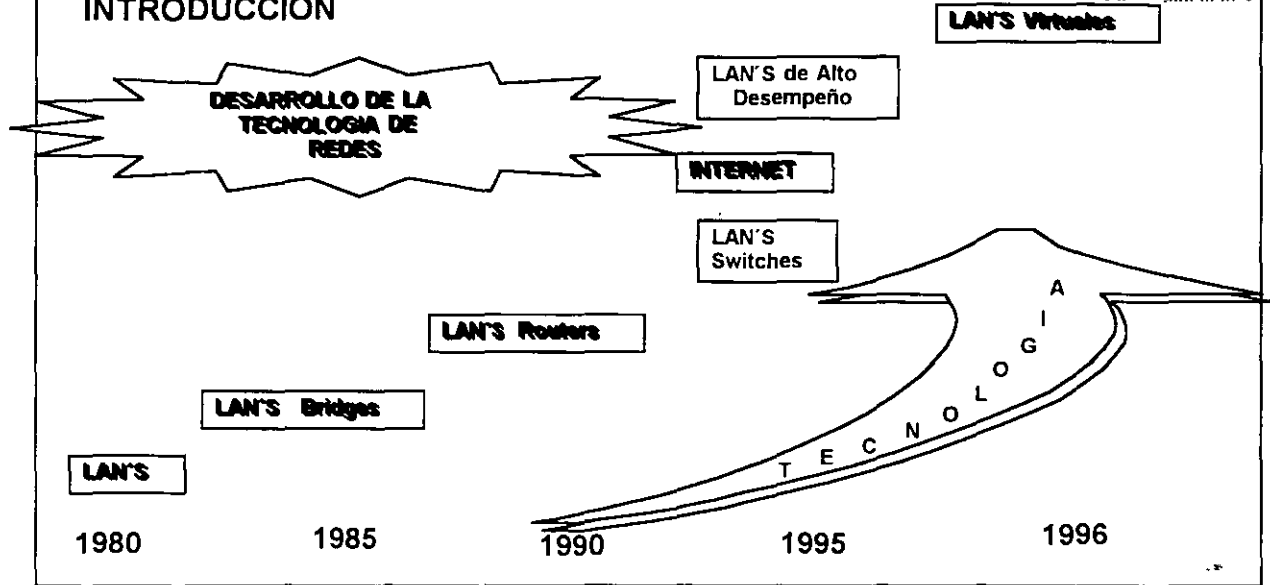
Ejemplo: Canal 23B+D = 23X64 Kbps + 64 Kbps

Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



INTRODUCCION




Notas:



Redes de alto desempeño

- FDDI, FDDI - II
- FAST ETHERNET
- TECNOLOGIA SWITCHING
- ATM
- FRAME RELAY
- B - ISDN



- REDES VIRTUALES
- REDES MULTIMEDIA
VIDEOCONFERENCIAS

REDES

LAN = MAN = WAN = GAN

Notas:



B-ISDN. Estándares

- ☐ En 1988 se establece la recomendación I.121 del CCITT.
- ☐ En 1990 el grupo de estudio XVIII aprueba 13 recomendaciones básicas, entre ellas:
 - ☞ Aspectos generales de B-ISDN
 - ☞ Servicios específicos de Red
 - ☞ Características fundamentales de ATM
 - ☞ Aplicaciones ATM
 - ☞ Operación y mantenimiento de los accesos a B-ISDN
- ☐ A partir de 1992, se han generado nuevas recomendaciones y grupos de estudio, entre ellas la I.113 de vocabulario y términos.

Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



ORGANIZACIONES INVOLUCRADAS EN LA ESTANDARIZACION DE -ISDN

A nivel mundial

CCITT	Comité Consultivo Internacional de Telegrafía y Telefonía
ISO	International Standards Organization

En Europa

CEPT	European Conference of Posts and Telecommunications Administrations
ETSI	European Telecommunications Standards Institute

En Estados Unidos

ANSI	American National Standard Institute
EIA	Electronic Industries Association
BELLCORE	Bell Communications Research

Notas:



B-ISDN.- INTRODUCCION

Diseñada para soportar conmutación de acuerdo a la demanda y conexiones en banda ancha tanto permanentes como semipermanentes para las aplicaciones punto-a-punto y punto-a-multipunto.

Soporta servicios de conmutación de circuitos y de conmutación de paquetes, aplicaciones "single media", "mixed-media" y "multimedia".

Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



BISDN .- CARACTERISTICAS

Conexiones conmutadas por demanda en Banda Ancha

- Permanentes
- Semipermanentes

Aplicaciones

- Punto a punto
- Punto a multipunto

Notas:



BISDN .- CARACTERISTICAS

Modos de Conmutación

- Paquetes
- Circuitos

Naturaleza de Servicios

- "Connection - oriented"
- "Connectionless"

Configuraciones

- Unidireccionales
- Bidireccionales

Notas:



BISDN. Características

Tráfico

- ☐ Velocidad constante CBR
(Constant Bit Rate)
 - ☞ Sin negociación de velocidad

- ☐ Velocidad variable VBR
(Variable Bit Rate)
 - ☞ Con negociación de velocidad

Notas:



BISDN CARACTERISITCAS

- ☐ **Conmutación por demanda**

- ☐ **Conexiones permanentes y semimermanentes**
 - ☞ **Punto a Punto**
 - ☞ **Punto a multipunto**

- ☐ **Conmutación de paquetes y conmutación de circuitos**
 - ☞ **Single media**
 - ☞ **Mexed media**
 - ☞ **Multimedia**
 - ☞ **“Conection less” y “Conection-oriented”**
 - ☞ **VBR y CBR**

Notas:



ISDN.- TERMINOLOGIA:

- ☐ Grupos Funcionales.
- ☐ Puntos de referencia.

Notas:



ISDN.- TERMINOLOGIA:

Grupos Funcionales.

- ☞ Terminadores de Red 1 (NT1).
Funciones equivalentes a las del nivel 1 del modelo de referencia OSI.
- ☞ Terminadores de Red 2 (NT2)
Funciones equivalentes a las de los niveles 1, 2 y 3 del modelo OSI.
- ☞ Equipo Terminal (TE)
Teléfonos digitales, Equipos terminales de datos y estaciones de trabajo que integran voz y datos.

Notas:




ISDN.- TERMINOLOGIA:

Grupos Funcionales.

- ☞ **Equipo terminal tipo 2 (TE2)**
Equipo terminal con interfaces no-ISDN
- ☞ **Adaptador terminal (TA)**
Grupo funcional que incluye las funciones para conectar equipo TE2 dentro de ISDN.

Notas:



ISDN.- TERMINOLOGIA:

Puntos de Referencia:

R: Interface funcional entre un grupo TE2 y un TA.

T: Interface entre el equipo NT2 y el NT1.

S: Interface entre equipos de usuario como pueden ser los TA o los TE1 y el equipo NT2.

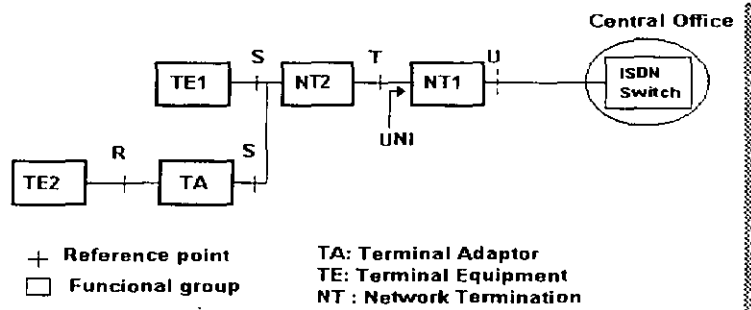
U: Interface del lado de la red del equipo NT1.

Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



ISDN.- TERMINOLOGIA:



Notas:



ISDN.- EQUIPO

Canales de Acceso:

- ☞ Canal B: 64Kbps para voz, datos en conmutación de circuitos o datos en conmutación de paquetes (B= bearer "portadora")
- ☞ Canal D: 16 ó 64Kbps para señalización, control o información del cliente en paquetes (D=delta).
- ☞ Canal H: 384Kbps (H0), 1,536Mbps (H11) ó 1,920 Mbps (H12) para teleconferencias, datos en alta velocidad o audio de alta calidad.

Notas:




ISDN.- EQUIPO

UNI: User Network Interface

- ☞ **Basic Rate Access (o BRI basic rate interface).**
Interface de usuario que provee 2 canales B y un canal D (2B+D).
- ☞ **Primary Rate Access (o PRI primary rate interface)**
Interface de usuario que provee 23 canales B y un canal D (23B+D).
- ☞ **Para canales H se prevee que en el futuro se utilice una interface de red tipo H+D.**

Notas:




DIPLOMADO EN REDES DE
COMPUTADORAS (LAN, WAN Y GAN)

REDES DE ALTO DESEMPEÑO: FAST y GIGA
ETHERNET FDDI-II, SWITCHING, ATM y
FRAMERELAY

Fast Ethernet y Giga Ethernet

Fast Ethernet

REQUERIMIENTOS DE ALTA VELOCIDAD Y SOLUCIONES PROPUESTAS.



Día con día, cada vez más usuarios de PC's se agregan a las redes. Al final de 1994 solo el 40% de las PC's en el mundo estaban conectadas en redes. Al mismo tiempo, la tecnología estaba logrando avances significativos como el lanzamiento comercial de el INTEL PENTIUM y tecnologías como POWER PC, tecnologías de sistemas de almacenamiento en disco duro avanzadas que decrementaban los costos, con el objeto de dar potencia a aplicaciones de redes basadas en PC's de propósito crítico, aplicaciones que hasta recientemente han sido posibles solo en un mainframe.

La capacidad de las PC's ha crecido en forma exponencial, al igual que las aplicaciones que corren en éstas, por lo que las tecnologías para conectar las PC's entre si, empiezan a ser un factor determinante en la funcionalidad de las redes locales.

Aunque no todos los usuarios requieren una red con capacidad de 100 mbps. muchas aplicaciones "lan-intensive" ya empujan los 10 mbps existentes y pueden beneficiarse con la tecnología actual de 100 mbps

Surgieron aplicaciones de datos intensivos como multimedia, trabajo en grupo y bases de datos cliente-servidor, que pronto harán de los 100mbps parte crítica de la mayoría de las Lan's.

Así mismo, como los servidores de red son ahora mas poderosos, han sido reubicados de conexiones locales a centrales de datos, donde necesitan conexiones de alta velocidad a 100 mbps al "backbone" para proporcionar capacidad centralizada al costo óptimo.

¿Que tecnología está mejor situada dentro del crecimiento de los requerimientos de alta velocidad de las redes de hoy?



La respuesta depende del usuario y de las necesidades de la red. **FAST ETHERNET** es una excelente alternativa por las siguientes razones:

ventajas de Fast Ethernet

- ◇ Alto rendimiento.
- ◇ Tecnología basada en estándares.
- ◇ Migración a costo aceptable con máximo aprovechamiento del equipo ya existente (infraestructura de cableado, sistemas de administración de red etc...)
- ◇ Soporte de los principales vendedores en todas las áreas de productos de red.
- ◇ Costo óptimo.

↳ Alto rendimiento.

Una de las mejores razones para cambiar a fast ethernet para grupos de trabajo, es la disponibilidad de manejo de ambas demandas agregadas, de una red multiusuario y el excesivo tráfico ocasionado por el alto desempeño de las PC's y las sofisticadas aplicaciones empleadas. Fast Ethernet es la solución óptima para grupos de trabajo.

↳ Tecnología basada en estándares.

Fast Ethernet está diseñada para ser la evolución más directa y simple de ethernet 10 base-T, la clave de su simplicidad es que fast ethernet usa csma/cd definido en el media access control.

El 100 base-T es una versión escalada del (M.A.C.), usado en ethernet convencional, sólo que más rápido, es la misma tecnología robusta, confiable y económica usada por 40 millones de usuarios hasta hoy, lo que es más, la misma compatibilidad entre 10 base-T y 100 base-T permite la fácil migración a conexiones de alta velocidad sin cambiar el cableado, depurando técnicas de administración de red y más.

Adicionalmente, ambas tecnologías ofrecen ambientes compartidos con conexiones ethernet compartidas o conmutadas permitiendo 10 o 100 mbps a todas las estaciones conectadas al hub, esto es ideal para grupos de trabajo de tamaño mediano con incrementos de demanda de ancho de banda ocasionales, ethernet compartido delibera el ancho de banda a un costo muy bajo.

Ambientes conmutados proveen el máximo ancho de banda para cada puerto conmutado del hub. Para grupos de trabajo grandes con demanda agregada que excede los 100 mbps, ethernet conmutado es la mejor solución.

↳ Costo efectivo de migración.

Como el protocolo natural de 10 base-T, virtualmente no cambia en fast ethernet, éste puede ser introducido fácilmente en ambientes de ethernet estandar. la migración es simple y económica en muchos aspectos importantes.



- ◊ Las especificaciones de el cableado para red 100 base-T permiten a fast ethernet correr en la mayoría de cableados comunes en ethernet, incluso categorías 3,4 y 5 de utp, stp y fibra óptica.
- ◊ Experiencia administrativa. los administradores pueden relevar en ambientes 100 base-T con herramientas de análisis de red familiares.
- ◊ La administración informática se traduce fácilmente de ethernet a 10MBPS a redes fast ethernet lo que significa recapacitación mínima del personal de administración y mantenimiento de la red.

Software de administración. Las redes fast ethernet pueden ser administradas con un protocolo simple como smnp.

Soporte de software. El software de aplicación y manejo de redes no cambia en redes 100 base-T.

Migración flexible. Adaptadores autosensibles de velocidad dual pueden correr a 10 ó 100 mbps en el medio existente, al igual que los concentradores con 10 100 mbps permiten el cambio dependiendo de la transmisión que se esté realizando

↳ Soporte de los principales fabricantes.

Fast ethernet es soportado por más de 60 fabricantes importantes, incluyendo empresas líder en adaptadores, conmutadores, estaciones de trabajo y empresas de semiconductores como 3Com, SMC, Intel, Sun Microsystems y Synoptics que empezaron a comercializar productos interoperables a fines de 1994.

Estas empresas son miembros de la Fast Ethernet Alliance (FEA), un consorcio cuyo objetivo es acelerar la tecnología fast ethernet a través de la Norma 802.3 del IEEE. Además la FEA estableció procedimientos de prueba y estándares para asegurar la interoperabilidad para los fabricantes de productos 100 Base-T.

↳ Valor óptimo.

Como la estandarización progresa rápidamente y los productos estarán disponibles por una gran variedad de fabricantes, el precio/desempeño de fast ethernet estará regido por la competitividad de las tecnologías de alta velocidad.

Al principio, los precios de fast ethernet superaban 10 veces el desempeño por menos de la mitad del costo por conexión. Ahora los precios están casi a la par de la tecnología de 10 Base-T y aún tienen las ventajas sobre otras tecnologías no ethernet.

↳ La tecnología tras fast ethernet.



Fast ethernet es una extensión del estandar existente 802.3 del IEEE, la nueva tecnología usa el mismo control (Media Access Control), de 802.3 conectado a través de otro control (Media Independent Interface), a otros tres controles de nivel físico, la especificación de M.I.I., es similar a la AUI de 10 mbps y proporciona una sola interface que puede soportar transceivers externos con alguna de las especificaciones 100 Base-T.

100 base-T soporta tres especificaciones: 100 baseTx, 100 base T4 y 100 base Fx, el estándar 100 base-T, también define una interface para concentrador universal y una interface de manejo.

En el diseño del MAC para 100 base-T, el IEEE reduce el tiempo de transmisión de cada bit, del MAC de 10 mbps de csma/cd multiplicado por un factor de 10 proporcionando turbo velocidad al paquete. Desde que el MAC está especificado de manera independiente de la velocidad, la funcionalidad en el formato del paquete no cambia, la longitud, el control de errores y la información de manejo son idénticos a 10 Base-T.

↳ Alternativas de cableado.

- ◊ 100 base-T soporta 3 especificaciones físicas
- ◊ 100 Base Tx: Cable UTP o STP de un par trenzado eia 568 o categoría 5 para datos.
- ◊ 100 Base T4: Cable UTP de 4 pares trenzados para voz y datos categoría 3, 4 ó 5.
- ◊ 100 Base Fx: sistema estándar de 2 fibras ópticas.

La flexibilidad de estas especificaciones permite a 100 base-T, implementar un ambiente de cable 10 Base-T virtual, permitiendo a los usuarios conservar la infraestructura de cableado mientras emigran a fast ethernet.

Las especificaciones 100 base Tx y 100 Base T4, juntas cubren todas las especificaciones de cableado que existen para redes 10 Base-T, las especificaciones fast ethernet pueden ser mezcladas e interconectadas a un hub como lo hacen las especificaciones 10 Base-T.

100 Base Tx está basado en la especificación PMD (Physical Media Dependent), desarrollada por el ansi x3t9.5, éste combina el MAC escalado con los mismos chips del transceiver y el PHY desarrollados para FDDI y CDDI. Como estos chips están disponibles y el estándar de señalización está completo, 100 Base-T ofrece una solución de tecnología aprobada y basada en estándares y soporta ambientes de cableado 10 Base-T.

100 Base-T permite transmisión a través de cable UTP 5 instalado virtualmente en las redes nuevas.

100 Base T4 es una tecnología de señal desarrollada por 3Com y otros miembros de Fast Ethernet Alliance para manejar las necesidades de cableado UTP 3 instalado en la mayoría de las antiguas redes basadas en 10 Base-T, esta tecnología permite a 100 Base-T correr sobre cableados UTP 3, 4 ó 5 permitiendo a las redes con cableado UTP 5 moverse a la tecnología de 100 Base-T sin tener que recablear.



100 Base FX es una especificación para fibra, ideal para grandes distancias o BackBones o ambientes sujetos a interferencia eléctrica.

↳ Auto-Negociación 10 / 100 MBPS

Para facilitar la migración de 10 a 100 MBPS el estándar 100 Base-T, incluye un sensor automático de velocidad, esta función opcional permite transmitir a 10 o 100 MBPS con comunicación automática disponible en ambos casos.

Auto-Negociación es usado en adaptadores 10 / 100 MBPS este proceso se da fuera de banda sin interposición de señal, para comenzar, una estación 100 Base-T advierte sus capacidades enviando un barrido de pulsos de prueba para verificar la integridad del enlace llamados FAST LINK PULSE, generados automáticamente al encender el equipo.

Si la estación receptora es un hub con capacidad 10 Base-T únicamente, el segmento operará a 10 MBPS, pero si el hub soporta 100 Base-T, este será censado por el FLP y usará el algoritmo de auto-negociación para determinar la mayor velocidad posible en el segmento, y enviar FLP's al adaptador para poner ambos dispositivos en modo 100 Base-T.

El cambio ocurre automáticamente sin intervención manual o de software, (una RED o un segmento de RED puede ser forzado a operar a 10 MBPS a través de un manejo de mayor jerarquía, aunque éste sea capaz de trabajar a 100 MBPS, si así se desea.)

↳ REGLAS DE TOPOLOGIA.

Fast Ethernet preserva la longitud crítica de 100 metros para cable UTP, como resultado del MAC escalado de la interface Ethernet.

Otras reglas topológicas de 100 MBPS son diferentes de las reglas Ethernet.

La figura 3 ilustra la clave de las reglas topológicas 10 Base-T y muestra ejemplos de como éstas permiten la interconexión en gran escala.

La máxima distancia en cable UTP es 100 metros igual que en 10 Base-T.

- ◇ En UTP se permiten máximo 2 concentradores y una distancia total de 205 mts.
- ◇ En topologías con un solo repetidor un segmento de fibra óptica de hasta 225 metros, puede conectarse a un backbone colapsado.
- ◇ Conexiones MAC to MAC, Switch to Switch, o End Station to Switch, se usan segmentos de hasta 450 mts., de fibra óptica bajo 100 Base FX.



- ◇ Para distancias muy largas una versión completamente duplex de 100 Base FX puede ser usada para conectar dos dispositivos a más de 2 KM de distancia.

Al principio, estas reglas topológicas pudieron parecer restrictivas , pero ahora en las redes con backbone, que usan fibra óptica , concentradores y/o ruteadores o puentes, Fast Ethernet puede ser fácilmente implementado en redes de gran escala o corporativas.


↳ ETAPAS DE MIGRACION.

La migración hacia fast ethernet está determinada en etapas, permitiendo al Administrador de la RED emigrar fast ethernet cuando y donde lo necesite.

Aquí tenemos una secuencia típica.

- ◇ Determine el tipo de cableado instalado, si este es categoría 5, se usan adaptadores 100 Base TX, las categorías 3 ó 4 requieren adaptadores 100 Base-T4
- ◇ Instale adaptadores de velocidad dual 10 /100 MBPS en PC's nuevas; para prepararse a la migración de la nueva tecnología, las PC's deben estar configuradas con adaptadores de velocidad dual, entonces podrán soportar ethernet compartido, ethernet conmutado, fast ethernet y aún fast ethernet conmutado.
- ◇ Instale concentradores 100 Base-T conforme el número de PC's se incremente, o conforme el tráfico de la RED empiece a crecer, comience la migración con hubs de velocidad dual, use un puente 10 / 100 MBPS para nodos que trabajen aún con 10 Base-T.
- ◇ Instale hubs conmutados 10 / 100 MBPS para las PC's que ya existen en la RED, para usarse con las PC's que no requieren tanta velocidad de comunicación, que además, necesitan conectarse a backbones o servidores a alta velocidad, el único cambio requerido en las conexiones ethernet 10 Base-T compartido a los puertos conmutados 10 /100 MBPS.
- ◇ Extienda 100 Base-T a los backbones. Conecte los grupos de trabajo y servidores a un backbone de alta velocidad, un puente o un ruteador con capacidad fast ethernet.





DIPLOMADO EN REDES DE
COMPUTADORAS (LAN, WAN Y GAN)

REDES DE ALTO DESEMPEÑO: FAST y GIGA
ETHERNET FDDI-II, SWITCHING, ATM y
FRAMERELAY

FDDI, FDDI - II



FDDI

Fiber Distributed Data Interface


Red anillo Token-Passing 100 Mb/s con redundancia.
(ANSI-X3T9)

Anillo principal = Conexión Punto a Punto entre
nodos para transmisión de datos
Anillo Secundario = Transmisión de datos/respaldo
del anillo principal en caso de
falla

FDDI provee comunicaciones por conmutación de
paquetes y transmisión de datos en tiempo real.



Notas:





FDDI

- * FDDI emplea una codificación 4B/1. tasa de transmisión a 100 Mb/125 Mhz 80% de eficiencia en el ancho de banda
 - * ETHERNET Y TOKEN-RING emplea una codificación Manchester
 - * Tasa de transmisión - ETHERNET: 10Mb/s-20 Mhz
- TOKEN-RING 16Mb/s-32 Mhz
- 50% DE EFICIENCIA EN EL ANCHO DE BANDA



Notas:




FDDI

FDDI: VS TOKEN - RING 16 MB/S	
* Reloj distribuido recuperación de errores	{ Monitor Activo
* Doble anillo	{ Anillo Sencillo
* Rotación del "TOKEN"	{ Sistema de reservación por prioridad
* Uso de Fibra Optica	{ Uso de Par Trenzado/Fibra Optica



Notas:




FDDI

TOKEN-PASSING ofrece una transmisión de datos más eficiente, ya que conforme aumenta el tráfico se requiere un mayor ancho de banda. TRT 85 %.

CSMA/CD Resulta más eficiente cuando se utiliza un menor ancho de banda.



Notas:






FDDI

- * FDDI Ofrece hasta 1000 conexiones físicas (500 Estaciones) y una distancia total de 200 Km. de extremo a extremo.
- * La distancia máxima entre nodos activos es la de 2 Km
- * Fibras Opticas empleadas:
 - A) Fibra tipo unimodo. con gran ancho de banda (GHz) y largas distancias (20-30 Km)
 - B) Fibra tipo multimodo. Fibras con nucleo 50-62.5 Micras y Medianas distancias (10-20 Km.) a 1300 nanómetros.



Notas:





FDDI

CONSIDERACIONES

Manejo

SMT (Interface SNMP)


Estadística de las estaciones reset. Soporte para deshabilitar.

300KM-180 Miles

El control es crítico para las Redes de gran tamaño y capacidad.




Notas:





FDDI: BACKBONES TOPOLOGIAS



Notas:



FDDI

ESTACIONES

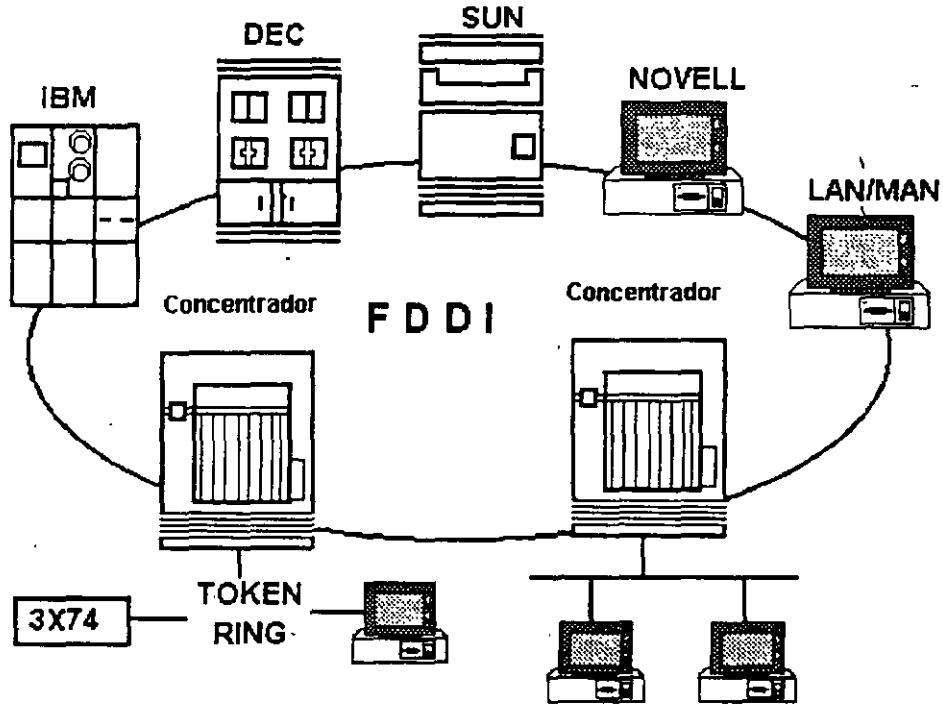
Tipo Clase A: Se conecta directamente al anillo doble

Tipo Clase B: Se conectan al concentrador puertos múltiples en Red estrella o Estaciones con posibilidad de conexión sencilla. Los concentradores pueden ser conectados en cascada.


Notas:



USANDO HOSTS CON FDDI



Notas:



FIBRAS OPTICAS

Hasta hace cerca de una década, las comunicaciones fueron realizadas a través de medios como cable coaxial o cable telefónico, Desde hace algunos años y ahora más fuerte que nunca se introduce un nuevo medio de comunicación: **las fibras ópticas**.

El uso de la luz como un medio de comunicación no es nuevo. El fuego fué usado como señal de comunicación en los amaneceres de la historia humana. La clave Morse fue utilizada particularmente en comunicaciones de una embarcación a otra usando espejos para reflejar la luz y transmitir señales.

En 1860 Alejandro Graham Bell demostró la transmisión de voz usando espejos.

Estos vibraban debido a las ondas sonoras generadas por la voz, de manera que la luz reflejada por los espejos era modulada por el sonido. La luz modulada en el receptor era enfocada en una lámina de Selenio, la resistencia de la lámina y su respectiva corriente variaba con los cambios de intensidad de la luz incidente. Esta corriente se aplicaba a un dispositivo parecido a un altavoz moderno.

Todos estos métodos dependían del medio ambiente y solo cubrían distancias pequeñas y para aplicaciones visuales en línea directa, en 1960 con la invención del láser, el interés por la comunicación luminosa tomo fuerza, aunque, contando con el láser, los métodos de comunicación por luz al aire libre seguían dependiendo del ambiente y limitados en alcance.

El primer intento para transmitir a larga distancia a través de fibra de vidrio fue realizado en 1966, pero las excesivas impurezas de la fibra de vidrio generaban grandes pérdidas de energía de la luz que viajaba a través de ésta. La transmisión seguía limitada en distancia, además de que el tamaño de los lasers con que se contaba en aquel tiempo hacían muy difícil el acoplamiento de la energía luminosa en las fibras de manera eficiente.

Con el desarrollo del diodo láser, del diodo LED, y más tarde la introducción de alta pureza, llegó la era de la comunicación por fibra: transmisión a largas distancias sin la necesidad de reamplificar la señal.

La historia del desarrollo de la tecnología de fibra óptica se centra en aplicaciones de comunicación y desarrollo e investigación gubernamental, los avances mas significativos se lograron recientemente en la década de los 70's y los 80's, aunque la teoría general de la propagación de la luz se desarrolló a lo largo de muchos años de investigaciones intentos y fracasos.

Una fibra óptica es una delgada varilla transparente hecha de vidrio o plástico puro, a través del cual la luz puede propagarse con una pérdida de señal muy baja, la estructura de una fibra óptica moderna consiste en el tubo de vidrio delgado recubierto por otro material con distintas



características ópticas, éste evita que la señal que viaja a través de la fibra óptica se refracte fuera de la misma ocasionando pérdidas en la señal.

El uso de fibra óptica para transmitir señales de comunicación tiene muchas ventajas importantes sobre los medios de comunicación convencionales:

- ◊ La baja pérdida en la energía de la señal.
- ◊ La baja tasa de distorsión en los pulsos de la señal transmitida
- ◊ El ancho de banda es mucho mayor que en UTP o coaxial.
- ◊ No es susceptible de ruido o interferencia eléctrica o electromagnética.
- ◊ Es muy segura, no es posible "robarse" la señal de la fibra óptica.
- ◊ Soporta ambientes hostiles, contaminación, salinidad, humedad o radiación. Es inmune.
- ◊ No existe una conexión eléctrica entre receptor y transmisor.
- ◊ El costo de la fibra óptica es casi el mismo que el del cable coaxial.
- ◊ Las velocidades de transmisión son muy altas.

Recientes desarrollos han permitido fibras ópticas con 0.2 dB de atenuación por kilómetro, además de los desarrollos de equipos para trabajar con fibra óptica con capacidad de operación de hasta 1 Ghz y mas de 3000 canales de comunicación individuales.

Las fibras ópticas se clasifican en dos tipos: unimodo y multimodo. Llamadas así por el número de modos de propagación de la longitud de onda de operación

Fibra multimodo

Es un tipo de fibra en la cual hay más de un modo de propagación de señal. Van desde las que tienen dos modos hasta cientos de modos de propagación. Las aplicaciones típicas de estas fibras son la telecomunicación con anchos de banda de 1 a 2 Ghz, cableado de inmuebles, con anchos de banda de 500 a 1000 Mhz y enlaces donde la potencia y el ancho de banda son necesarios, generalmente 50 a 100 Mhz son suficientes.

Fibra unimodo

La fibra unimodo es fabricada con los mismos materiales y bajo los mismos procesos que las fibras multimodo, la diferencia es el tamaño del centro de la fibra que es mas pequeño y la cantidad de impurezas que es diferente a la fibra multimodo, hace la diferencia de características de operación.



Las siguientes tablas ofrecen un panorama general de características

↳ Dimensiones

Fibra óptica Tipo	diámetro del núcleo (micras)	diámetro del revestimiento (micras)	longitud de onda (Nanometros)
unimodo	8.10	125	1300,1500
multimodo	50	125	850,1300

↳ cuadro comparativo de atenuación.


Medio de comunicación Tipo	Longitud de onda o Frecuencia	Atenuación (dB / Km.)
COAXIAL	100 Mhz	61
Fibra Óptica Multimodo	850 Nm	2.4 - 3.2
Fibra Óptica Multimodo	300 Nm	1.0 - 1.5
Fibra Óptica Unimodo	1300 Nm	menor a 0.5
Fibra Óptica Unimodo	1300 Nm	menor a 0.25

↳ Distancias máximas cubiertas por un segmento de línea de comunicación


Medio de comunicación Tipo	Distancia máxima sin repetidor (Mts) (Rango dinámico típico 35 dB)
COAXIAL	570
Fibra óptica Multimodo a 850 Nm	10, 000
Fibra óptica Multimodo a 1300 Nm	20, 000
Fibra óptica Unimodo a 1300 Nm	60, 000
Fibra óptica Unimodo a 1550 Nm	120, 000

ASPECTO DE LA FIBRA OPTICA






existe una gran variedad de presentaciones para fibras ópticas dependiendo de las aplicaciones.



CABLE DE FIBRA OPTICA PARA ESTRUCTURA



TUBO DE FIBRA OPTICA DE USO INDUSTRIAL

↳ CONECTORES DE FIBRA OPTICA.

Son dispositivos de unión, que realizan la función de acoplamiento entre dos fibras ópticas o en los extremos de éstas, permitiendo un fácil manejo, instalación y mantenimiento de la fibra óptica.



FDDI-5

Los parámetros que definen la calidad de un conector para un sistema de transmisión dado son los siguientes:

- ◊ Pérdida por inserción.
- ◊ Facilidad para su ensamble y montaje.
- ◊ Estabilidad al ambiente.
- ◊ Confiabilidad.
- ◊ Inserción de perturbaciones al sistema.
- ◊ Costo.

Aunque normalmente es imposible optimizar todos los parámetros, la elección de un conector es el resultado de un balance de necesidades específicas, debe tenerse el cuidado no solo de seleccionar el conector adecuado, sino que también debe ponerse especial atención en el momento del manejo y ensamble de los conectores.

FDDI

Las nuevas tecnologías de interconexión de redes tienden al uso de la fibra óptica, como medio de comunicación, tiene una capacidad de transmisión de datos y de seguridad muy altas. Las fibras ópticas pueden soportar transmisiones de varios cientos de Mbps. Los cableados por medio de fibra óptica pueden soportar grandes distancias sin necesidad de repetidores, además de ser un medio inmune a la interferencia electromagnética.

Los costos de conexión con fibra óptica son típicamente altos, pero podemos esperar que estos precios bajen significativamente en los próximos años.

Ya existen en el mercado, proveedores que cuentan con las tarjetas necesarias para poder realizar conexiones con fibra óptica para las topologías Ethernet y Token Ring.

Muchas compañías están optando por la fibra óptica por diversas razones, entre ellas está la velocidad de transmisión de la que es capaz. Por ejemplo, **FDDI**¹ soporta velocidades de transmisión de hasta 100 Mbits por segundo. En comparación con Ethernet que transmite a 10 Mbits por segundo o Token Ring que transmite a 4 ó 16 Mbits por segundo.

El comité 802.6 de la IEEE ha adoptado estándares para redes de área metropolitana, y el American National Standards Institute ha desarrollado los estándares FDDI y FDDI-II .

Además, la fibra óptica tiende a ser más segura que el cableado de cobre. Una red interconectada por medio de fibra óptica puede trabajar cerca de equipo eléctrico altamente sensible sin interferir uno con el otro. Un cable de fibra óptica entre dos edificios no atraerá rayos como el cable de cobre.

¹ Fiber Distributed Data Interface



Al hablar de redes interconectadas por medio de fibra óptica, generalmente se está hablando de FDDI, diversos productos capaces de soportar FDDI han estado saliendo lentamente al mercado y se han dejado ver en diversas exposiciones de computadoras.

Como Token Ring, FDDI usa una topología con forma de anillo y un Token eléctrico para pasar el control de la red de una estación a otra, más no es compatible con Token Ring.

La mayor parte de las redes actuales con FDDI usan un doble anillo en donde cada nodo se une a los dos anillos independientes, transmitiendo los datos en sentidos opuestos. Esta configuración mejora la velocidad de transmisión así como la confiabilidad de la red, pero es muy caro.

Hasta ahora, FDDI se ha usado para interconectar PC's de alta velocidad o estaciones de trabajo con redes, o bien como *backbone* para interconectar estaciones más lentas, de igual manera que una carretera une los diferentes pueblos. Conectarse a FDDI es caro, dado el alto costo de los componentes ópticos, así como el costo del transreceptor y los integrados necesarios para FDDI. Debido a sus características de ancho de banda, la fibra óptica se usa principalmente para backbones (que es un segmento que une varias redes locales) .

Existe también FDDI-II que es una segunda versión de FDDI que nos permite transmitir voz y video además de datos. De manera distinta a FDDI que tiene un reloj corriendo de manera independiente, FDDI-II tendrá un marco de 125 microsegundos, permitiendo ser sincronizado con la red de comunicaciones

3
4
:



FAQ'S sobre FDDI y FDDI-II

Q. What does FDDI stand for?

Fiber Distributed Data Interface

Q. What is the difference between FDDI and FDDI-II?

Both FDDI and FDDI-II runs at 100 M bits/sec on the fiber.
FDDI can transport both async and sync types of frames.
FDDI-II has a new mode of operation called Hybrid Mode.
Hybrid mode uses a 125usec cycle structure to transport isochronous traffic, in addition to sync/async frames.
FDDI and FDDI-II stations can be operated in the same ring only in Basic mode.

Q. What is the name of the standards and where can I get them?

ANSI X3T9.5 standards

American National Standards Institute
1430 Broadway, New York, NY 10018, USA
Attention: Sales Dept.

- IEEE Standards

IEEE Service Center
445 Hoes Lane, Piscataway, NJ 08855, USA

- X3T9.5 Documents

Global Engineering Documents
(USA) 1-800-854-7179

Q. What are other good sources of printed information?

- FDDI Technology and Applications: Edited Mirchandani and Khanna
- Handbook of Computer Communications Standards Vol 2: By Stallings
- Call up DEC to ask for the free FDDI tutorial book
- Dig up 1986-1992 issue of IEE Local Computer Network Conference

Q. I've heard that FDDI uses a token passing scheme for access arbitration, how does this work?

A token is a normal FDDI frame with a fixed format.
The station waits until a token comes by, grabs the token, transmits the the frames and release the token. The amount of frames that can be transmitted is determined by timers in the MAC protocol chips.

[You really need a diagram for the station and/or topology.]

Q. I've heard that FDDI is a counter-rotating ring, what does this mean?

FDDI is a dual ring technology. And each ring is running in the opposite direction to improve fault recovery.

Q. What is a dual ring of trees?

See the diagram.




Q. What is dual homing?

When a DAS is connected to two concentrator ports, it is called dual-homing. One port is the active link, where data is transmitted and the other port is a hot standby. The hot standby will constantly testing the link and will kick in if the active link failed or disconnected. The B-port in a DAS is the active port and the A-port is the hot-standby.

Q. What is a DAS?

DAS (Dual Attach Station) is a station with two peer ports (A-Port and B-Port). The A-port is going to the B-Port of another DAS, and the B-port is going to connect to the A-Port the yet another DAS. ie:




Q. What is a SAS?

SAS (Single Attach Station) is a station with one peer port (S-Port). It is usually connected to the M-Port of a concentrator.

Q. What is a wrapped ring?

When a link in the dual-ring is broken or not connected, the two adj ports connecting to the broken link will be disconnected and the both stations enter the wrap state.



Q. Do I need a concentrator port for each workstation, or can workstations be chained together?

Usually you will need a concentrator port (M-Port) to connect each SAS. DAS can be hooked up to the main rings or concentrator port(s).

Q. If I use a concentrator, what are the advantages/disadvantages?

Advantages: Fault tolerance. When a link breaks, the ring can be segmented. A concentrator can just bypass the problem port and avoid most segmentations. It also gives you better physical planning. Usually people prefer tree physical topology. Generally star configuration of a concentrator system is easier to troubleshoot.



Disadvantages: A concentrator represents a single point of failure.
There may also be more costly.

you can build a tree as deep as you want. We have

Q. Can I cascade concentrators? Are there limitations as to how many?

Yes. And dual-rings of concentrator here connecting machine rooms and wiring rooms. And from there we connect to other concentrators to different offices. Then we have a concentrator in the lab to different machines. There is a maximum of 500 stations on an FDDI LAN.

Q. What is a bypass and what are the issues in having or not having one?

Bypass is a (\$600-\$1200) device that is used to skip a station on the ring if it is turned off. Therefore, you don't need to use concentrator to avoid the segmentation problems. One problem with them is that they increase the db loss of the fiber, so you can't have too many of them (3 activated in a row maximum, I believe).

Q. What are the minimum/maximum distances on fiber runs?

no min, 2 km max for multimode fiber. 20 km max (may be as high as 60km, we're not sure) for single mode fiber.
500 m for the new Low Cost Fiber.

Q. What are the types of fiber that are supported?

Multimode (62.5/125 micron graded index multimode fiber)
and other fiber like 50/125. 85/125. 100/140 allowed
Single mode (8-10 micron)
The new Low Cost Fiber (plastics?) standard.

Q. I've hear of FDDI over Copper, what type of cable does this scheme use?

Type 1 STP - distance between connections must be less than 100 m
Category 5 UTP - distance between connections must be less than 100 m
(The ANSI standard for STP and UTP is incomplete, but a number of companies are already shipping proprietary twisted-pair solutions until the standard is completed, which is expected later this year.)

? Q. Is there any advantage to seperating the fiber pairs (will the ring work better if only one strand is broken on a DAS connection?)

Q. I have ethernet, can I bridge/route between the 2 topologies?

Yes. But from what we are hearing some protocols are having problems. Only TCP/IP is handling frame fragmentation correctly. (See below).
It should also be noted that frame fragmentation will not work for DECNET, IPX, LAT, Appletalk, NETBEUI etc.
IP is the only protocol that has a standard method of fragmenting.
Other protocols destined for Ethernet Lans must stay below the 1500 MTU.



Q. I've heard that there is a frame length difference, what are the issues and problems here?

FDDI frames has a max size of 4500 bytes and Enet only 1500 bytes. Therefore your bridge or router needs to be smart enough to fragment the packets (eg into smaller IP fragments). Or you need to reduce your frame size to 1500 bytes (of data).

Q. What does an FDDI frame look like?

PA	Preamble (II) (8 or more Idle symbol pairs)
SD	Starting Delimiter (JK) (J followed by K control symbol)
FC	Frame Control (nn) (Tell you if it is a token, MAC frame, LLC frame, SMT frame, frame priority, sync or async)
DA	Destination Address (nn) (6 bytes of MAC Address in MSb first format)
SA	Source Address (nn) (6 bytes of MAC Address of this station)
INFO	Information field (nn) (Variable Length. Usually starts with LLC header, then SNAP field, then the payload eg IP packet)
ED	Ending Delimiter (T) (one T control symbol)
FS	Frame Status (EAC) (Three symbols of status of Error, Address_match, and Copied. Each symbol is either SET or RESET. eg If EAC == RSS, then then frame has no error, some station on the ring matched the DA, and some station on the ring copied the frame into its buffer.

Q. So FDDI is 100 Megbits per second, what is the practical maximum bps?

Depends. You can get aggregate usage up to 95Mbit/s with no problem. But 75Mbps is pretty good. Actually, this question depends so much on how you construct your test, what equipment you use, etc, that the best idea is to let the user decide.

Q. What happens when I bridge between a 100 Mbps FDDI and a 10Mbps ethernet if the FDDI traffic destined for the ethernet gets above 8 Mbps? 10 Mbps?

After the buffer fills Frames start dropping. This is not a problem unique to FDDI however. Consider ethernet to T1, or multiple ethernets to a single ethernet.

Q. What is the latency across a bridge/router? (Yes I know that different vendors are different, but what is a the window?)

No idea.



Q. Are there FDDI repeaters?

Yes. But it is not a standard yet. A group in the ANSI committee is looking into making FDDI repeater a standard. Other companies like ODS has something like single mode to multimode converter.

Q. What type of test and trouble shooting equipment is available for FDDI?

Digital Technology Inc (DTI), W&D, HP, and Tekelec all sell FDDI analyzers. The Sniffer from Network General also has a module that works with the NPI FDDI Cards. SGI has a nice looking ringmap program. IBM has a product called DatagLANce. Most Ethernet tools will also work with FDDI in the protocol level. Also a optical time domain reflectometer (TDR) is recommended for db loss checking and distance measurements, though it has been reported that an FDDI link tester is less expensive and will do the job.

Q. What about network station management? Does FDDI support SNMP?

Yes. There is a FDDI-SNMP MIB translation from the SNMP working group.

Q. What is a beaconing ring? Does FDDI beacon?

Beacon is a special frame that FDDI MAC sends when something is very wrong. When Beaconing for a while, SMT will kick in trying to detect and solve the problem.

Q. How about interoperability, does one manufacture's equipment work with others?

Just like any networking products, Ethernet, Token, FDDI, ATM, there is a possibility that one vendor does not work with another. But most of the equipment shipping today is tested at InterOp, UNH or ANTC, are this is the equipment that will meet the minimum interoperability requirements. Ask the vendor what type of testing they did and ask them to ship you a system for field trial before you pay big bucks for it.

Q. Can I interface FDDI to a PC (ISA Bus), PC (EISA Bus), PC (Micro channel Bus), Macintosh, Sun workstation, DECstation 5000, NEXT computer, Silicon Graphics, Cisco router, WellFleet router, SNA gateway (McData), other?

Yes. I am not sure if NeXT has any FDDI adaptor software, but there are ~5 different NuBus FDDI cards in the market. But FDDI adaptors are available for all other buses or vendors.

Q. What is the maximum time a station has to wait for media access. What type of applications care?

$MaxTime = \sim(\#of\ stations * T_neg)$

(T_neg ist the negotiated target token rotation time)

Usually this won't happened. It is only a very very heavily loaded ring but the station be waiting for that long. If this is the case, then change the T_request of the station to some lower value (eg 8 msec).



Q. Can I bridge/route TCPIP, SNA, Novell, Sun protocols, DecNet, Banyan Vines, Appletalk, X windows, LAT?

Yes for IP, Novell, DecNet, X windows.
Don't know about the others.

Q. What are the applications that would use FDDI's bandwidth?

Basically anything will be at least a bit faster. From NFS to images transmission. Even if a single station cannot take advantage of the 100M bit/sec, the aggregate bandwidth will help a lot if your Ethernet is saturated. However, note that though FDDI has higher bandwidth than ethernet, the signals travel at the same speed. The propagation of a signal on the transmission line is the same for ethernet, token ring, and FDDI.

Q. What are the effects of powering off a workstation on a DAS or SAS connection?

Depends. Let's do SAS first, it is easier. If a SAS is connected to a concentrator, then the concentrator will bypass the SAS connection using an internal data path. If the DAS is connected to a concentrator, then the concentrator will also bypass the DAS. If the DAS is connected to the trunk rings without using an optical bypass switch, then the trunk ring will wrap. If multiple stations power off on the trunk rings, then the ring will be badly segmented. Now if the DAS is using an optical bypass switch, the switch will kick in and prevent the ring from wrapping.

Q. What are the effects of disconnecting the fiber on a DAS or SAS connection?

SAS connecting to concentrator:

Same as above.

DAS dual-home to a concentrator:

If A-port fiber breaks, no effect on B port since A port is a backup port. (And SMT will NOT send out alert msg.)

If B-port fiber breaks, A-port will kick in, complete PCM and be used as the primary connection.

DAS on trunk rings, with no optical bypass:

If one fiber breaks, then the ring will wrap.

If both fibers break, ring will wrap, station won't be communicate.

DAS on trunk rings using optical bypass:

If one fiber between bypass and the next station breaks, then the ring will wrap.

If both fibers between bypass and the next station break, ring will wrap, station won't be able to communicate.

If one fiber between bypass and the host station breaks, then the ring will wrap.

If two fiber between bypass and the host station breaks, then the ring will wrap.



Q. What is one recommended topology?

Connect backbone concentrators and ring monitors to the trunk rings, and connect all the workgroup concentrators and users stations to the backbone concentrators. Connect bridges and routers to backbone concentrators using dual-homing.

Q. What is Graceful Insertion? Should I demand it from my vendors?

Graceful Insertion is a method to insert a station (or a tree) in a concentrator without losing any data frames (and not going into Ring_Non_Op mode). The theory goes as Graceful Insertion can minimize ring non_op and losing frame, therefore it saves you transmission timeout of lost frame in upper layer protocol (eg TCP) and retransmission effort. The following is the counter argument: Graceful Insertion can hold up the ring for more time that the FDDI ring non-op recovery time. And Upper layer protocol is designed to perform frame recovery and retransmission anyway. And no vendor can guarantee 100% Graceful Insertion anyway. Should I get Graceful Insertion in my concentrators? If it is free, take it. You are going to get ring_op no matter what (eg insertion in the trunk ring and station power down).

Q. Is there a Graceful De-insertion?

No.

Q. Can you name a few FDDI Concentrator vendors?

IBM, Optical Data System, SynOptics, Cabletron, DEC, Chipcom, NPI, Synernetics, 3Com, Interphase, Ungermann-Bass, Timeplex, Crescendo/Cisco, Sumitomo etc ... (vendors feel free to email me to be included here)

Q. Can I run FDDI on electrical cable?

DEC is already sell a FDDI link that runs on coax.
ANSI is currently finishing up the TP-FDDI Standard for running FDDI on twisted-pair media (Category 5 Cable).
ANSI is also working on a standard (long term TP working group) to run FDDI on telephone cable. [Please comment.]
IBM and a group of vendors (SynOptics, National Semiconductor ...) promote SDDI that runs FDDI on Shielded Twisted-Pair cable. (this is incomplete), there is much work being done on FDDI over various types of electrical cable, most notably twisted pair.

Q. What does SMT stand for? What does it do? Do I need it?

Station Management (SMT). It is part of the ANSI FDDI Standards that provides link-level management for FDDI. SMT is a low-level protocol that addresses the management of FDDI functions provided by the MAC, PHY, and PMD. It performs functions like ring recovery, frame level management, link control, etc. Every stations on FDDI need to have SMT. The latest version of the SMT standard is version 7.3, but most vendors ship products with SMT version 6.2.



Q. Who supports FDDI-II?

National Semiconductor Corp, IBM, Apple Computer, XDI, Alpha Inc, etc

Q. Who is working on Synchronous frame type utilization?

Alpha, IBM, and many more companies. Try to contact scoop4@aol.com and warren@lgev2.vnet.ibm.com. They are working with a group of companies to define the usage of SYNC frame in FDDI-I rings.

Q. Can I connect two Single attach stations together and form a two stations ring without a concentrator?

yes. You can do that if both stations support the S-S port connection. Most vendors support the S-S connections.

Q. What are ports? What are the different type of ports?

A port is the basically the fiber optic connector on the card. FDDI SMT defines 4 types of ports (A, B, M, S). A dual-attach station has two ports, one A-port and one B-port. A single attach station has only one port (S-port). A concentrator will have many M-port for connecting to other stations' A, B or S-ports.

Q. What are the port connection rules?

When connecting DASs, one should connect the A-port of one station to the B-port of another. S-port on the SAS is to connect to the M=port on the concentrators. A and B-port on DASs can also connect to the M-port of concentrator. But M-ports of the concentrator will not connect to each other.

In more detail, SMT suggested tthe following rules:


	A	B	M	S
A	-	+	+	-
B	+	-	+	-
M	+	+	X	+
S	-	-	+	-

==> '+' is the preferred connection

==> '-' connection has possible problems, and a vendor can choose to disable that connection in the default configuration

==> 'X' indicates a legal connection and will be rejected





DIPLOMADO EN REDES DE
COMPUTADORAS (LAN, WAN Y GAN)

REDES DE ALTO DESEMPEÑO: FAST y GIGA
ETHERNET FDDI-II, SWITCHING, ATM y
FRAMERELAY

ATM

ATM and Cell Relay Service

1.1 Introduction

1.1.1 Background

Asynchronous transfer mode (ATM), as the term is used in current parlance, refers to a high-bandwidth, low-delay switching and multiplexing technology that is now becoming available for both public and private networks. ATM principles and ATM-based platforms form the foundation for the delivery of a variety of high-speed digital communication services aimed at corporate users of high-speed data, LANs interconnection, imaging, and multimedia applications. Residential applications, such as video distribution, videotelephony, and other information-based services, are also planned. ATM is the technology of choice for evolving broadband integrated services digital network (B-ISDN) public networks, for next-generation LANs, and for high-speed seamless interconnection of LANs and WANs. ATM supports transmission speeds of 155 Mbits/s and 622 Mbits/s, and will be able to support speeds as high as 10 Gbits/s in the future. Networks operating at these speeds have been called gigabit networks. As an option, ATM will operate at the DS3 (45 Mbits/s) rate; some proponents are also looking at operating at the DS1 (1.544 Mbits/s) rate. While ATM in the strict sense is simply a Data Link Layer protocol, ATM and its many supporting standards, specifications, and agreements constitute a platform supporting the integrated delivery of a variety of switched high-speed digital services.

Cell relay service (CRS) is one of the key new services enabled by ATM. CRS can be utilized for enterprise networks that use completely private communication facilities, use completely public communication facilities, or use a hybrid arrangement. It can support a variety of evolving corporate applications, such as desk-to-desk videoconferencing of remote parties, access to remote multimedia video servers (for example, for network-based client/server video systems), multimedia conferencing, multimedia massaging, distance learning, business imaging (including CAD/CAM), animation, and cooperative work (for example, joint document editing). CRS is one of three "fastpacket" technologies, that have entered the scene in the 1990s [the other two are frame relay service and Switched Multimegabit Data Service (SMDS)]. A generic ATM platform supports all of these fastpacket services (namely, it can support cell relay service, frame relay service, and SMDS), as well as circuit emulation service.

1993 saw the culmination of nine years of ATM standards-making efforts. Work started in 1984 and experienced an acceleration in the late 1980s and early 1990s. With the ITU-TS (International Telecommunications Union Telecommunication Standardization) standards and the ATM Forum implementers' agreements, both of which were finalized in 1993, the technology is ready for introduction in the corporate environment. In particular, a user-network interface (UNI) specification that supports switched cell relay service as well as the critical point-to-multipoint connectivity, important for new applications, has been finalized (multiservice UNIs are also contemplated). In 1993, the ATM Forum also published a broadband intercarrier interface (B-ICI) specification; this specification is equally critical for wide-area network (WAN) inter-LATA service. At press time, a variety of vendors were readying end-user products for 1994 market introduction; some prototype products have been on the market since the early 1990s. A number of carriers either already provide services or are poised to do so in the immediate future.

A key aspect of B-ISDN in general and ATM in particular is the support of a wide range of data, video, and voice applications in the same public network. An important element of service integration is the provision of a range of services using a limited number of connection types and multipurpose user-network interfaces. ATM supports both nonswitched permanent virtual connections (PVCs) and switched virtual connections (SVCs). In a PVC service, virtual connections between endpoints in a customer's network are established at service subscription time through a provisioning process; these connections or paths can be changed via a subsequent provisioning process or via a customer network management (CNM) application. In SVC, the virtual connections are established as needed (that is, in real time) through a signaling capability. ATM supports services requiring both circuit-mode and packet-mode information transfer capabilities. ATM can be used to support both connection-oriented (e.g., frame relay service) and connectionless services (e.g., SMDS).

1.1.2 Course of Investigation: applying ATM to enterprise networks

This book is aimed at corporate practitioners who may be interested in determining how they can deploy ATM and cell relay technology in their networks at an early time and reap the benefits. The purpose of this first chapter is to provide an overview of key ATM/cell relay service concepts. These concepts will be revisited in more depth in the chapters that follow.

The book has four major segments: (1) platform technology applicable to all B-ISDN services, (2) cell relay service, (3) interworking and support of basic multimedia, and (4) use of ATM in corporate enterprise networks. Table 1.1 provides a roadmap of this investigation.

The text is not a research monograph on open technical issues related to ATM, such as traffic descriptors, ingress/egress traffic policing, object-oriented signaling, etc. A literature search undertaken in the spring of 1993 showed that about 5000 papers and trade articles have been written on ATM in the previous nine years, including Refs. 7 through 15. The purpose of this book, therefore, is to stick to the facts and avoid unnecessary hype. There are a few books already available, but these tend to focus on protocol issues. This text aims at a balance between standards, platforms, interworking, and, most important, deployment issues.

In summary, a network supporting cell relay service accepts user data units (called cells) formatted according to a certain layout and sends these data units in a connection-oriented manner (i.e., via a fixed established path), with sequentiality of delivery, to a remote recipient (or recipients). Every so often a cell may be dropped by the network to deal with network congestion; however, this is a very rare event. The user needs a signaling mechanism in order to tell the network what he or she needs. The signaling mechanism consists of a Data Link Layer capability (where the Data Link Layer has been partitioned into four sublayers) and an application-level call-control layer. ATM switches and other network elements supporting cell relay service can also support other fastpacket services. If the user wishes to use ATM to achieve a circuit-emulated service, certain adaptation protocols in the user equipment will be required. Other adaptation protocols in the user equipment are also needed to obtain fastpacket services over an ATM platform. ATM supports certain operations and maintenance procedures that enable both the user and the provider to monitor the "health" of the network. Figure 1.1 is a physical view of an ATM network.

A glossary of some of the key ATM and related concepts, based on a variety of ATM standards and documents, is given in Table 1.2

1.1.3 Early corporate applications of ATM

Table 1.3 depicts some of the proposed applications for ATM/cell relay service.

TABLE 1.1 Areas of Investigation In This Text

1. ATM and cell relay service: an overview
 2. ATM platform aspects and ATM proper
 3. ATM Adaptation Layer
 4. Signaling
 5. Cell relay service-a formal definition
 6. Cell relay service-traffic and performance issues
 7. Support of fastpacket services and CPE
 8. ATM interworking: support of basic multimedia
 9. Third-generation LANs
 10. Network management
 11. Typical user equipment and public carrier service availability
 12. How to migrate a pre-ATM enterprise network to CRS
-

1.2 Basic ATM Concepts

1.2.1 ATM protocol model: an overview

ATM's functionality corresponds to the Physical Layer and *part* of the Data Link Layer of the Open Systems Interconnection Reference Model (OSIRM). This protocol functionality must be implemented in appropriate user equipment (for example, routers, hubs, and multiplexers) and in appropriate network elements (for example, switches and service multiplexers). A *cell* is a block of information of short fixed length (53 octets) that is composed of an "overhead" section and a payload section (5 of the 53 octets are for overhead and 48 are for user information), as shown in Fig. 1.2. Effectively, the cell corresponds to the Data Link Layer frame that is taken as the atomic building block of the cell relay service. The term *cell relay* is used because ATM transports user cells reliably and expeditiously across the network to their destination. ATM is a transfer mode in which the information is organized into cells; it is asynchronous in the sense that the recurrence of

cells containing information from an individual user is not necessarily periodic.

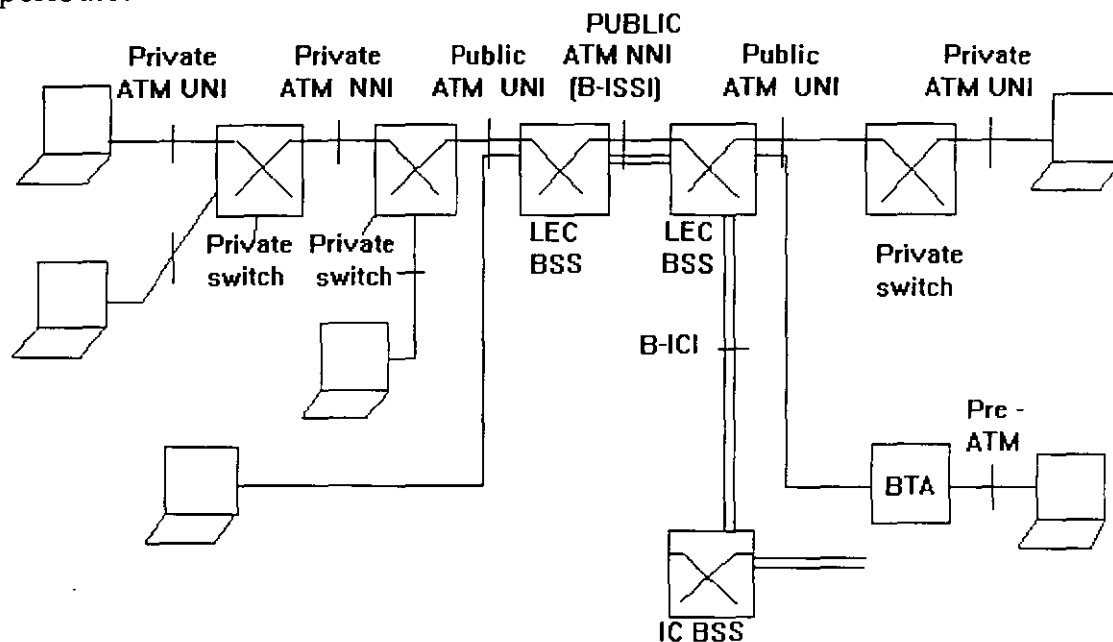


Figure 1.1 A physical view of an ATM/CRS private/public network. BSS = broadband switching system (B-ISDN switch), BTA = broadband terminal adapter, B-ISSI = broadband interswitching system interface; BICI = broadband intercarrier interface; LEC = local exchange carrier; IC = interexchange carrier.

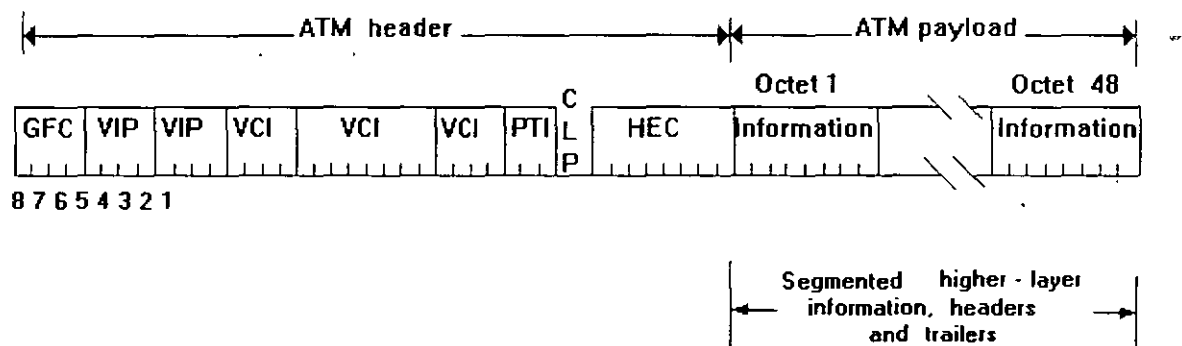


Figure 1.2 ATM cell layout

The ATM architecture utilizes a logical protocol model to describe the functionality it supports. The ATM logical model is composed of a User Plane, a Control Plane, and a Management Plane. The *User Plane* with its layered structure, supports user information transfer. Above the Physical Layer, the ATM Layer provides information transfer for all applications the user may contemplate; the ATM Adaptation Layer (AAL), along with associated services and protocols, provides service-dependent functions to the layer above the AAL.

TABLE 1.2 Glossary of Key ATM Terms

AAL	A layer that adapts higher-layer user protocols (e.g., TC/IP, APPN) to the ATM protocol (layer).
AAL connection	An association established by the AAL between two or more next higher layer entities.
Asynchronous time-division multiplexing	A multiplexing technique in which a transmission capability is organized in a priori unassigned time slots. The time slots are assigned to cells upon request of each application's instantaneous real need.
Asynchronous transfer mode	A transfer mode in which the information is organized into cells. It is asynchronous in the sense that the recurrence of cells containing information from an individual user is not necessarily periodic.
ATM Layer connection	An association established by the ATM Layer to support communication between two or more ATM service users (i.e., between two or more next higher layer entities or between two or more ATM management entities) The communication over an ATM Layer connection may be either it is bidirectional or unidirectional. When it is bidirectional, two VCCs are used. When it is unidirectional, only one VCC is used.
ATM Layer link	A section of an ATM Layer connection between two adjacent active ATM Layer entities (ATM entities).
ATM link	A virtual path link (VPL) or a virtual channel link (VCL).
ATM peer-to-peer connection	A virtual channel connection (VCC) or a virtual path connection (VPC).
ATM traffic descriptor	A generic list of traffic parameters that can be used to capture the intrinsic traffic characteristics of a requested ATM connection
ATM user-user connection	An association established by the ATM Layer to support communication between two or more ATM service users [i.e., between two or more next-higher-layer entities or between two or more ATM management (ATMM) entities]. The communication over an ATM Layer connection may be either bidirectional or unidirectional. When it is bidirectional, two VCCs are used. When it is unidirectional, only one VCC is used.
Broadband	A service or system requiring transmission channels capable of supporting rates greater than the Integrated Service Digital Network (ISDN) primary rate.
Call	An association between two or more users or between a user and a network entity that is established by the use of network capabilities. This association may have zero or more connections.
Cell	ATM Layer protocol data unit.
Cell delay variation	A quantification of variability in cell delay for an ATM Layer connection.

TABLE 1.2 Glossary of Key ATM Terms (continued)

Cell header	ATM Layer protocol control information.
Cell loss ratio	The ratio of the number of cells "lost" by the network (i.e., cells transmitted into the network but not received at the egress of the network) to the number of cells transmitted to the network
Cell transfer delay	The transit delay of an ATM cell successfully passed between two designated boundaries
Connection	The concatenation of ATM Layer links in order to provide an end-to-end information transfer capability to access points.
Connection admission control (CAC)	The procedure used to decide if a request for an ATM connection can be accepted based on the attributes of both the requested connection and the existing connections
Connection endpoint (CE)	A terminator at one end of a layer connection within a SAP
Connection endpoint identifier (CEI)	Identifier of a CE that can be used to identify the connection at a SAP.
Corresponding entities	Peer entities with a lower-layer connection among them.
Header	Protocol control information located at the beginning of a protocol data unit.
Layer connection	A capability that enables two remote peers at the same layer to exchange information.
Layer entity	An active element within a layer.
Layer function	A part of the activity of the layer entities
Layer service	A capability of a layer and the layers beneath it that is provided to the upper-layer entities at the boundary between the layer and the next higher layer.
Layer user data	Data transferred between corresponding entities on behalf of the upper-layer or layer management entities for which they are providing services
Multipoint access	User access in which more than one terminal equipment (TE) is supported by a single network termination.
Multipoint-to-multipoint connection	A collection of associated ATM VC or VP links and their associated endpoint nodes, with the following properties (1) All N nodes in the connection, called endpoints, serve as root nodes in a point-to-multipoint connection to all of the $(N - 1)$ remaining endpoints. (2) Each of the endpoints on the connection can send information directly to any other endpoint [the receiving endpoint cannot distinguish which of the endpoints is sending information without additional (e.g., higher-layer) information].
Multipoint-to-point connection	A multipoint-to-point connection where the bandwidth from the root node to the leaf nodes is zero, and the return bandwidth from the leaf node to the root node is nonzero.
Network node interface (NNI)	The interface between two network nodes.

TABLE 1.2 Glossary of Key ATM Terms (continued)

Operation and maintenance (OAM) cell	A cell that contains ATM Layer Management (LM) information. It does not form part of the upper-layer information transfer.
Peer entities	Entities within the same layer.
Physical Layer (PHY) connection	An association established by the PHY between two or more ATM entities. A PHY connection consists of the concatenation of PHY links in order to provide an end-to-end transfer capability to PHY SAPs.
Point-to-multipoint connection	A collection of associated ATM VC or VP links, with associated endpoint nodes, with the following properties: (1) One ATM link, called the root link, serves as the root in a simple tree topology. When the root node sends information, all of the remaining nodes on the connection, called Leaf Nodes, receive copies of the information. (2) Each of the leaf nodes on the connection can send information directly to the root node. The root node cannot distinguish which leaf is sending information without additional (higher-layer) information. (3) The leaf nodes cannot communicate with one another directly with this connection type.
Point-to-point connection	A connection with only two endpoints
Primitive	An abstract, implementation-independent interaction between a layer service user and a layer service provider or between a layer and the Management Plane
Protocol	A set of rules and formats (semantic and syntactic) that determines the communication behavior of layer entities in the performance of the layer functions
Protocol control information (PCI)	Information exchanged between corresponding entities, using a lower-layer connection, to coordinate their joint operation.
Protocol data unit (PDU)	A unit of data specified in a layer protocol and consisting of protocol control information and layer user data.
Relaying	A function of a layer by means of which a layer entity receives data from a corresponding entity and transmits them to another corresponding entity.
Service access point (SAP)	The point at which an entity of a layer provides services to its layer management entity or to an entity of the next higher layer.
Service data unit (SDU)	A unit of interface information whose identity is preserved from one end of a layer connection to the other
Source traffic descriptor	A set of traffic parameters belonging to the ATM traffic descriptor used during the connection setup to capture the intrinsic traffic characteristics of the connection requested by the source.

TABLE 1.2 Glossary of Key ATM Terms (*continued*)

Structured data transfer	The transfer of AAL user information supported by the CBR AAL when the AAL user data transferred by the AAL are organized into data blocks with a fixed length corresponding to an integral number of octets.
Sublayer	A logical subdivision of a layer.
Switched connection	A connection established via signaling.
Symmetric connection	A connection with the same bandwidth value specified for both directions.
Traffic parameter	A parameter for specifying a particular traffic aspect of a connection.
Trailer	Protocol control information located at the end of a PDU.
Transit delay	The time difference between the instant at which the first bit of a PDU crosses one designated boundary and the instant at which the last bit of the same PDU crosses a second designated boundary.
Unstructured data transfer	The transfer of AAL user information supported by the CBR AAL when the AAL user data transferred by the AAL are not organized into data blocks.
Virtual channel (VC)	A communication channel that provides for the sequential unidirectional transport of ATM cells.
Virtual channel connection (VCC)	A concatenation of VCLs that extends between the points where the ATM service users access the ATM Layer. The points at which the ATM cell payload is passed to or received from the user of the ATM Layer (i.e., a higher layer or ATM management entity) for processing signify the endpoints of a VCC. VCCs are unidirectional.
Virtual channel link (VCL)	A means of unidirectional transport of ATM cells between the point where a VCI value is assigned and the point where that value is translated or removed.
Virtual path (VP)	A unidirectional logical association or bundle of VCs
Virtual path connection (VPC)	A concatenation of VPLs between virtual path terminators (VPTs). VPCs are unidirectional.
Virtual path link (VPL)	A means of unidirectional transport of ATM cells between the point where a VPI value is assigned and the point where that value is translated or removed.

In approximate terms, the AAL supplies the balance of the Data Link Layer not included in the ATM Layer. The AAL supports error checking, multiplexing, segmentation, and reassembly. It is generally implemented in user equipment but may occasionally be implemented in the network at an interworking (i.e., protocol conversion) point. The *Control Plane* also has a layered architecture and supports the call control and connection functions. The *Control Plane* uses AAL capabilities as seen in Fig. 1.3; the layer above the AAL in the Control Plane provides call control and connection control.

TABLE 1.3 Possible early applications of ATM in real environments (partial list)

Application	Advantages of ATM use	Associated true-to-life business issues
WAN interconnection of existing enterprise network	High bandwidth; switched service	Unknown cost; geographic availability; equipment availability
WAN interconnection of existing LAN, especially FDDI (fiber distributed data interface) LANs	High bandwidth, switched service	Unknown cost; geographic availability
WAN interconnection of mainframe and supercomputer channel	High bandwidth; only service that supports required throughput (200 Mbits/s); switched service	Unknown cost; geographic availability; equipment availability
WAN interconnection of ATM-based LANs	High bandwidth; switched service; multipoint connectivity	New application, not widely deployed; unproven business need; unknown cost; geographic availability
Support of distributed multimedia	High bandwidth; switched service; multipoint connectivity	New application, not widely deployed; unproven business need; unknown cost; geographic availability
Support of statewide distance learning with two way video	High bandwidth; switched service; multipoint connectivity	New application, not widely deployed, unproven market; other solutions exist; unknown cost; geographic availability
Support of videoconferencing (including desktop video)	High bandwidth, switched service; multipoint connectivity	Not widely deployed; unproven market; other solutions exist, particularly at lower end (e.g , 384 Kbits/s H.200 video); unknown cost; geographic availability
Residential distribution of video (video dial tone)	High bandwidth; switched service; multipoint connectivity	Unproven market., other solutions exist, particularly CATV, expensive for this market; needs MPEG II (Motion Picture Expert Group) hardware; geographic availability

It deals with the signaling necessary to set up, supervise, and release connections. The *Management Plane* provides network supervision functions. It provides two types of functions: Layer Management and Plane Management. Plane Management performs management functions related to the system as a whole and provides coordination among all planes, Layer Management performs management functions relating to resources and parameters residing in its protocol entities. See Fig. 1.3. (The various protocols identified in this figure will be discussed at length later.)

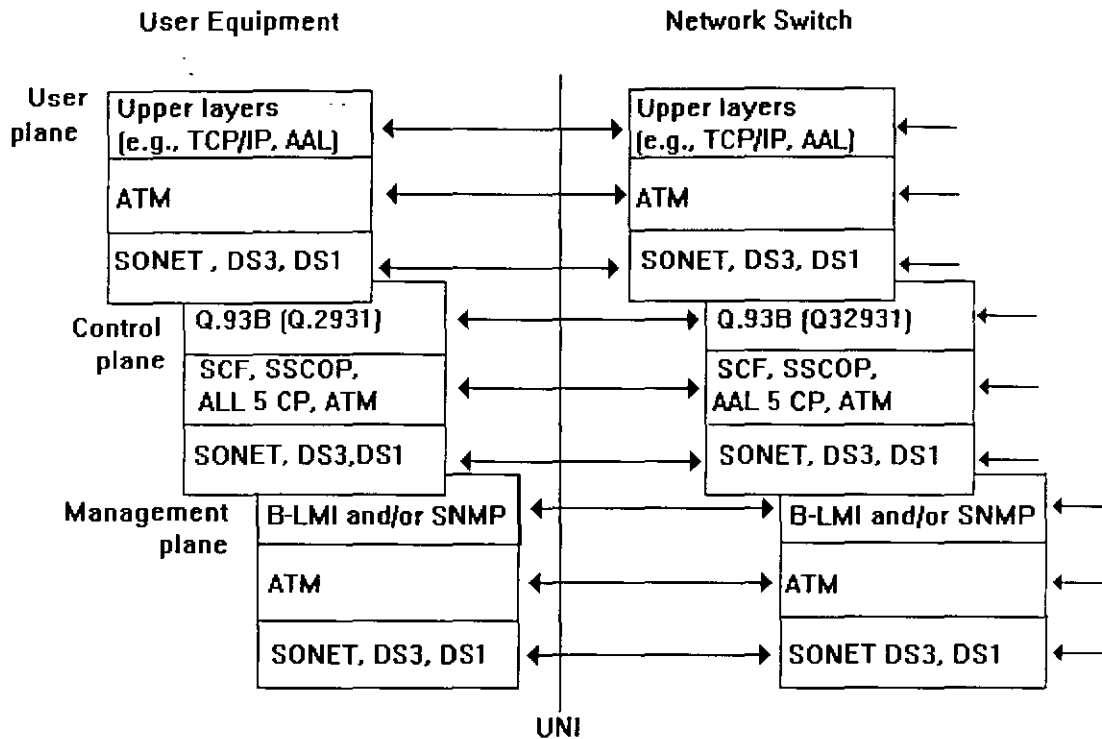


Figure 1.3 Planes constituting the ATM protocol model.

As noted in this description, four User Plane protocol layers are needed to undertake communication in an ATM-based environment:

1. A layer below the ATM Layer, corresponding to the Physical Layer. The function of the Physical Layer is to manage the actual medium-dependent transmission. Synchronous Optical Network (SONET) is the technology of choice for speeds greater than 45 Mbits/s.
2. The ATM Layer (equating approximately, for comparison, to the upper part of a LAN's medium access control layer), which has been found to meet specified objectives of throughput, scalability, interworking, and consistency with international standards. The function of the ATM layer is to provide efficient multiplexing and switching, using cell relay mechanisms.
3. The layer above the ATM Layer, that is, the AAL. The function of the AAL is to insulate the upper layers of the user's application protocols [e.g., TCP/IP (Transmission Control Protocol/Internet Protocol)] from the details of the ATM mechanism.
4. Upper layers, as needed. These include TCP/IP, IBM APPN, OSI TP, etc.

Several layers are needed in the Control Plane. Early PVC service users do not need the signaling stack in the Control Plane (this situation is analogous to the early PVC frame relay environment). SVC service needs both an information transfer protocol stack and a companion signaling protocol stack.

ATM is intended to support a variety of user needs, including highspeed data, video, and multimedia applications. These applications have varying quality of service (QOS) requirements. For example, video-based services have stringent delay, delay variation, and cell loss goals, while other applications have different QOS requirements. Carriers are proposing to support a number of service classes in order to tailor cell relay to a variety of business applications. In particular, there have been proposals to support a “guaranteed” and a “best efforts” class.

1.2.2 Classes of ATM applications.

Two main service categories of ATM have been identified (from the network point of view): (1) interactive broadband service and (2) distributive broadband service. See table 1.4.

1.2.3 Virtual connections

Just as in traditional packet switching or frame relay, information in ATM is sent between two points not over a dedicated, physically owned facility, but over a shared facility composed of virtual channels. Each user is assured that, although other users or other channels belonging to the same user may be present, the user's data can be reliably, rapidly, and securely transmitted over the network in a manner consistent with the subscribed quality of service. The user's data is associated with a specified virtual channel. ATM's “sharing” is not the same as a random access technique used in LANs, where there are no guarantees as to how long it can take for a data block to be transmitted: in ATM, cells coming from the user at a stipulated (subscription) rate are, with a very high probability and with low delay, “guaranteed” delivery at the other end, almost as if the user had a dedicated line between the two points. Of course, the user does not, in fact, have such a dedicated (and expensive) end-to-end facility, but it will seem that way to users and applications on the network. Cell relay service allows for a dynamic transfer rate, specified on a per-call basis. Transfer capacity is assigned by negotiation and is based on the source requirements and the available network capacity. Cell sequence integrity on a virtual channel connection is preserved by ATM.

Cells are identified and switched by means of the label in the header, as seen in Fig. 1.2. In ATM, a *virtual channel* (VC) is used to describe unidirectional transport of ATM cells associated by a common unique identifier value, called the *virtual channel identifier* (VCI). Even though a channel is unidirectional, the channel identifiers are assigned bidirectionally. The bandwidth in the return direction may be assigned symmetrically, or asymmetrically, or it could be zero. A *virtual path* (VP) is used to describe unidirectional transport of ATM cells belonging to virtual channels that are associated by a common identifier value, called the *virtual path identifier* (VPI). See Fig. 1.4.

VPIs are viewed by some as a mechanism for hierarchical addressing. In theory, the VPI/VCI address space allows up to 16 million virtual connections over a single interface; however, most vendors are building equipment supporting (a minimum of) 4096 channels on the user's interface. Note that these labels are only locally significant (at a given interface). They may undergo remapping in the network; however, there is an end-to-end identification of the user's stream so that data can flow reliably. Also note that on the network trunk side more than 4096 channels per interface are supported.

Figure 1.5 illustrates how the VPI/VCI field is used in an ATM WAN. Figure 1.6 depicts the relationship of VPs and VCs as they might be utilized in an enterprise network.




Figure 1.4 Relationship of VCs VPs

TABLE 1.4 Broadband Service Supported by ATM/Cell Relay

Interactive services *Conversational services* provide the means for bidirectional communication with real-time, end-to-end information transfer between users or between users and servers. Information flow may be bidirectional symmetric or bidirectional asymmetric. Examples: High speed data transmission, image transmission, videotelephony, and videoconferencing.

Messaging services provide user-to-user communication between individual users via storage units with store-and-forward, mailbox, and/or message handling (e.g., information editing, processing, and conversion) functions. Examples: Message handling services and mail services for moving pictures (films), store-and-forward image and audio information.

Retrieval services allow users to retrieve information stored in information repositories (information is sent to the user on demand only). The time at which an information sequence is to start is under the control of the user. Examples: Film, high-resolution images, information on CD-ROMs, and audio information.

Distributive services Distribution services without user individual presentation control provide a continuous flow of information that is distributed from a central source to an unlimited number of authorized receivers connected to the network. The user can access this flow of information without having to determine at which instant the distribution of a string of information will be started. The user cannot control the start and order of the presentation of the broadcast information, so that depending on the point in time of the user's access, the information will not be presented from its beginning. Examples: broadcast of television and audio programs.

Distribution services with user individual presentation control provide information distribution from a central source to a large number of users. Information is rendered as a sequence of information entities with cyclical repetition. The user has individual access to the cyclically distributed information, and can control the start and order of presentation. Example: broadcast videography.

1.3 ATM Protocols: An Introductory Overview

Figure 1.7 depicts the cell relay protocol environment, which is a particularization of the more general B-ISDN protocol model described

earlier. The user's equipment must implement these protocols, as must the network elements to which the user connects. Some of the key functions of each layer are described next.

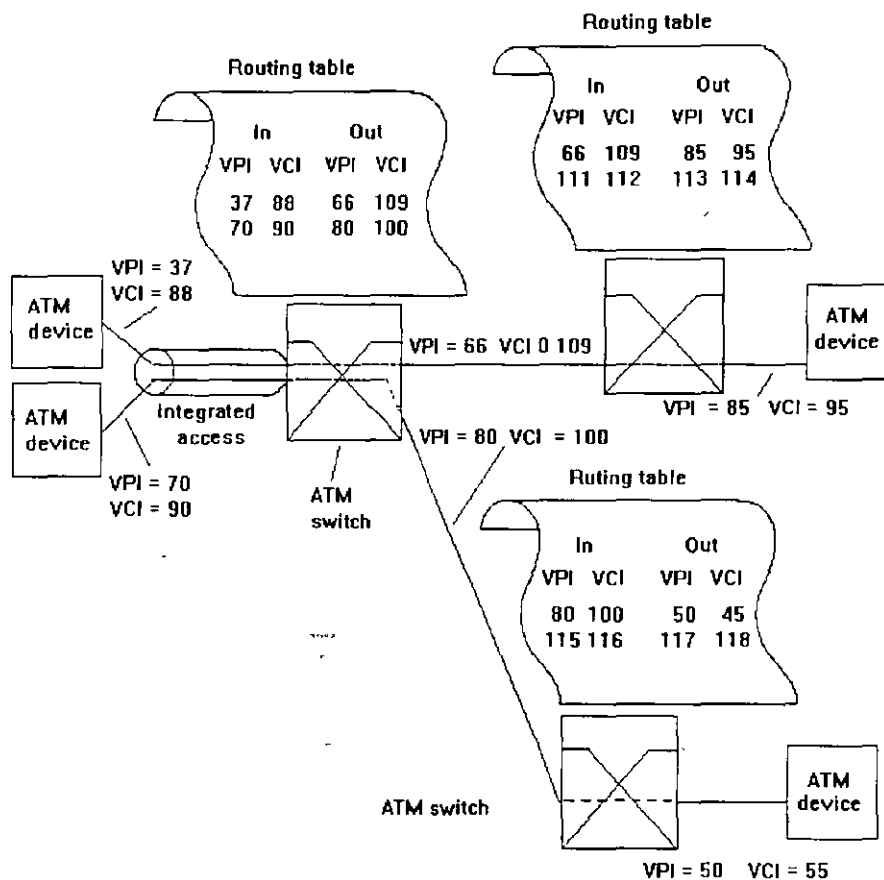


Figure 1.5 Illustrative use of VPIs and VCIs.

1.3.1 Physical Layer functions

The Physical Layer consists of two *logical* sublayers: the Physical Medium-Dependent (PMD) Sublayer and the Transmission Convergence (TC) Sublayer. The PMD includes only physical medium-dependent functions. It provides bit transmission capability, including bit transfer, bit alignment, line coding, and electrical-optical conversion. The Transmission Convergence Sublayer performs the functions required to transform a flow of cells into a flow of information (i.e., bits) that can be transmitted and received over a physical medium. Transmission Convergence functions include (1) transmission frame generation and recovery, (2) transmission frame adaptation, (3) cell delineation, (4) header error control (HEC) sequence generation and cell header verification, and (5) cell rate decoupling.




Figure 1.6 Example of use of VPs and VCs in an enterprise network (broadband switches not shown for simplicity). *Note:* VPs and VCs can be preprovisioned (PVCs) or on-demand (SVC with signaling).

The transmission frame adaptation function performs the actions that are necessary to structure the cell flow according to the payload structure of the transmission frame (transmit direction) and to extract this cell flow out of the transmission frame (receive direction). In the United States, the transmission frame requires SONET envelopes above 45 Mbits/s. Cell delineation prepares the cell flow in order to enable the receiving side to recover cell boundaries. In the transmit direction, the payload of the ATM cell is scrambled. In the receive direction, cell boundaries are identified and confirmed, and the cell flow is descrambled. The HEC mechanism covers the entire cell header, which is available to this layer by the time the cell is passed down to it. The code used for this function is capable of either single-bit correction or multiple-bit error detection. The transmitting side computes the HEC field value. Cell rate decoupling includes insertion and suppression of idle cells, in order to adapt the rate of valid ATM cells to the payload capacity of the transmission system.

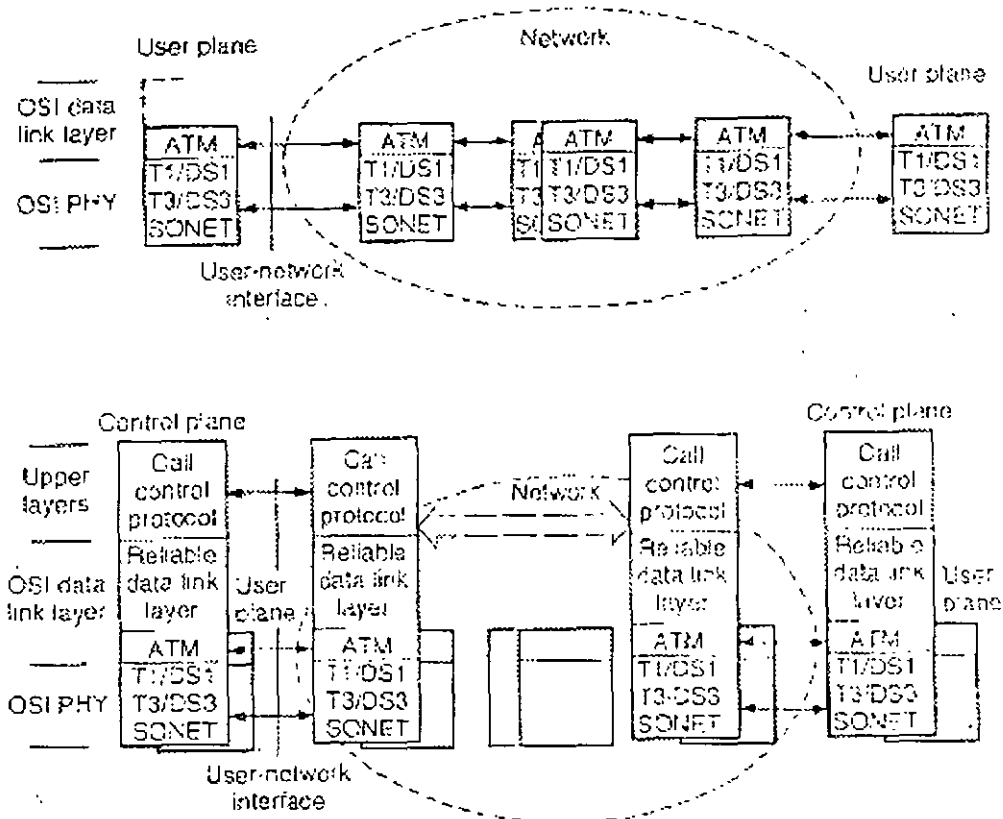


Figure 1.7 CRS environment, protocol view. Top: User Plane (information flow)
Bottom: Control Plane (signaling)

The service data units crossing the boundary between the ATM Layer and the Physical Layer constitute a flow of valid cells. The ATM Layer is unique, that is, independent of the underlying Physical Layer. The data flow inserted in the transmission system payload is physical medium-independent; the Physical Layer merges the ATM cell flow with the appropriate information for cell delineation, according to the cell delineation mechanism.

The transfer capacity at the UNI is 155.52 Mbits/s, with a cell-fill capacity of 149.76 Mbits/s because of Physical Layer framing overhead. Since the ATM cell has 5 octets of overhead, the 48-octet information field equates to a maximum of 135.631 Mbits/s of actual user information. A second UNI interface is defined at 622.08 Mbits/s, with a service bit rate of approximately 600 Mbits/s. Access at these rates requires a fiber-based loop. Other UNIs at the DS3 rate and perhaps at the DS1 rate are also being contemplated in the United States. The DS1 UNI is discussed in the context of an electrical interface (T1); so is the DS3 UNI.

1.3.2 ATM Layer functions

ATM supports a flexible transfer capability common to all services, including connectionless services (if these are provided). The transport functions of the ATM Layer are independent of the Physical Layer implementation. As noted, connection identifiers are assigned to each link of a connection when required and are released when no longer needed. The label in each ATM cell is used to explicitly identify the VC to which the cells belong. The label consists of two parts: the VCI and the VPI. A VCI identifies a particular VC link for a given virtual path connection (refer to Fig. 1.6). A specific value of VCI is assigned each time a VC is switched in the network. With this in mind, a VC can be defined as a unidirectional capability for the transport of ATM cells between two consecutive ATM entities where the VCI value is translated. A VC link is originated or terminated by the assignment or removal of the VCI value.

The functions of ATM include the following

Cell multiplexing and demultiplexing. In the transmit direction, the cell multiplexing function combines cells from individual VPs and VCs into a noncontinuous composite cell flow. In the receive direction, the cell demultiplexing function directs individual cells from a noncontinuous composite cell flow to the appropriate VP or VC.

Virtual path identifier and virtual channel identifier translation. This function occurs at ATM switching points and/or cross-connect nodes. The value of the VPI and/or VCI field of each incoming ATM cell is mapped into a new VPI and/or VCI value (this mapping function could be null).

Cell header generation / extraction. These functions apply at points where the ATM Layer is terminated (e.g., user's equipment). The header error control field is used for error management of the header. In the transmit direction, the cell header generation function receives cell payload information from a higher layer and generates an appropriate ATM cell header except for the HEC sequence (which is considered a Physical Layer function). In the receive direction, the cell header extraction function removes the ATM cell header and passes the cell information field to a higher layer.

For the UNI, as can be seen in Fig. 1.2, 24 bits are available for cell routing: 8 bits for the VPI and 16 bits for the VCI. Three bits are available for

payload type identification; this is used to provide an indication of whether the cell payload contains user information or network information. In user information cells, the payload consists of user information and, optionally, service adaptation function information. In network information cells, the payload does not form part of the user's information transfer. The header error control field consists of 8 bits.

The initial thinking was that if the cell loss priority (CLP) is set by the user (CLP value is 1), the cell is subject to discard, depending on the network (congestion) conditions. If the CLP is not set (CLP value is 0), the cell has higher priority. More recent thinking proposes not making use of this bit on the part of the user (i.e., it must always be set to 0 by the user).

ATM is discussed further in Chap. 2.

1.3.3 ATM Adaptation Layer

Additional functionality on top of the ATM Layer (i.e., in the ATM Adaptation Layer) may have to be provided by the user (or interworking) equipment to accommodate various services. The ATM Adaptation Layer enhances the services provided by the ATM Layer to support the functions required by the next higher layer. The AAL function is typically implemented in the user's equipment, and the protocol fields it requires are nested within the cells' payload.

The AAL performs functions required by the User, Control, and Management Planes and supports the mapping between the ATM Layer and the next higher layer. Note that a different instance of the AAL functionality is required in each plane. The AAL supports multiple protocols to fit the needs of the different users; hence, it is service-dependent (namely, the functions performed in the AAL depend upon the higher-layer requirements). The AAL isolates the higher layers from the specific characteristics of the ATM Layer by mapping the higher-layer protocol data units into the information field of the ATM cell and viceversa. The AAL entities exchange information with the peer AAL entities to support the AAL functions.

The AAL functions are organized in two logical sublayers, the Convergence Sublayer (CS) and the Segmentation and Reassembly Sublayer (SAR). The function of the CS is to provide the AAL service to the layer above it; this sublayer is service-dependent. The functions of the SAR are (1) segmentation of higher-layer information into a size suitable for the information field of an ATM cell and (2) reassembly of the contents of ATM cell information fields into higher layer information.

Connections in an ATM network support both circuit-mode and packet-mode (connection-oriented and connectionless) services of a single medium and/or mixed media and multimedia. ATM supports two types of traffic: constant bit rate (CBR) and variable bit rate (VBR). CBR transfer rate parameters for on-demand services are negotiated at call setup time. (Changes to traffic rates during the call may eventually be negotiated through the signaling mechanism; however, initial deployments will not support renegotiation of bit rates.) CBR transfer rate parameters for permanent services are agreed upon with the carrier from which the user obtains service. This service would be used, for example, to transmit real-time video. VBR services are described by a number of traffic-related parameters (minimum capacity, maximum capacity, burst length, etc.). VBR supports packet like traffic (e.g., variable-rate video, LAN interconnection, etc.). The AAL protocols are used to support these different connection types.

In order to minimize the number of AAL protocols, however, a service classification is defined based on the following three parameters: (1) the timing relation between source and destination (required or not required), (2) the bit rate (constant or variable, already discussed), and (3) the connection mode (connection-oriented or connectionless). Other parameters, such as assurance of the communication, are treated as quality of service parameters, and therefore do not lead to different service classes for the AAL. The five classes of application are:

Class A service is an on-demand, connection oriented, constant-bit rate ATM transport service. It has end-to-end timing requirements. This service requires stringent cell loss, cell delay, and cell delay, variation performance. The user chooses the desired bandwidth and the appropriate QOS during the signaling phase of an SVC call to establish a Class A connection (in the PVC case, this is prenegotiated). This service can provide the equivalent of a traditional dedicated line and may be used for videoconferencing, multimedia, etc.

Class B service is not currently defined by formal agreements. Eventually it may be used for (unbuffered) compressed video.

Class C service is an on-demand, connection-oriented, variable-bitrate ATM transport service. It has no end-to-end timing requirements. The user chooses the desired bandwidth and QOS during the signaling phase of an SVC call to establish the connection.

Class D service is a connectionless service. It has no end-to-end timing requirements. The user supplies independent data units that are delivered by the network to the destination specified in the data unit. SMDS is an example of a Class D service.

Class X service is an on-demand, connection-oriented ATM transport service where the AAL, traffic type (VBR or CBR), and timing requirements are user-defined (i.e., transparent to the network). The user chooses only the desired bandwidth and QOS during the signaling phase of an SVC call to establish a Class X connection (in the PVC case, this is prenegotiated).

Three AAL protocols have been defined in support of these User Plane applications: AAL Type 1, AAL Type 3/4, and AAL Type 5. Type 1 supports Class A, Type 3/4 supports Class D, and Type 5 supports Class X. It appears that the computer communication community (e.g., LAN and multiplexing equipment) will use AAL Type 5. Additionally, the ATM service likely to be available first (and the one supported by evolving computer equipment vendors) is Class X (that is, cell relay service).

Note that two stacks must be implemented in the user's equipment in order to obtain VCs on demand (i.e., SVC service) from the network. With this capability, the user can set up and take down multiple connections at will. The Control Plane needs its own AAL; there has been agreement to use AAL 5 in the Control Plane. Initially only PVC service will be available in the United States. In this mode, the Control Plane stack is not required, and the desired connections are established at service initiation time and remain active for the duration of the service contract. Also note that AAL functions (SAR and CS) must be provided by the user equipment (except in the case where the network provides interworking functions). Additionally, the user equipment must be able to assemble and disassemble cells (i.e., run the ATM protocol).

AAL is discussed further in Chap. 3. Signaling is discussed in Chap. 4.

1.4 Multiservice ATM Platforms

SMDS and frame relay PVC are currently available fastpacket services. SMDS is a high-performance, packet-switched public data service being deployed by the Regional Bell Operating Companies (RBOCs), GTE, and SNET in the United States. SMDS is also being deployed in Europe. Frame relay PVC is a public data service that is widely available today and is expected to be deployed by all RBOCs and most interexchange carriers by the end of 1994.

Frame relay SVC should be available in the 1994 - 1995 time frame. ATM is a switching and multiplexing technology that is being embraced worldwide by a wide spectrum of carriers and suppliers. This new technology can switch and transport voice, data, and

video at very high speeds in a local or wide area. What is the relationship of SMDS and frame relay to ATM?

SMDS and frame relay are carrier services, whereas ATM is a technology, as indicated at the beginning of this chapter. ATM will be used by carriers to provide SMDS, frame relay, and other services, including cell relay service (a fastpacket service based on the native ATM bearer service capabilities). Customers who deploy SMDS or frame relay now will be able to take advantage of the benefits of ATM technology without changing the services they use as carriers upgrade their networks to ATM. The customer's investment in SMDS or frame relay equipment and applications is thus preserved.¹⁷

SMDS is based on well-defined specifications and provides switched, LAN-like transport across a wide area.¹ SMDS service features include a large maximum packet size, an addressing structure that enables data transfer among all SMDS customers, the ability to send the same SMDS packet to several destinations by specifying one address (group addressing), address screening, and strict quality of service values. As ATM technology is deployed within public carrier networks, SMDS service features will not change. The current SMDS interface between the customer and the network uses an access protocol based on the IEEE 802.6 standard. As ATM technology is deployed, this existing SMDS interface will be maintained. The published requirements for ATM switching and transmission technology specify that the existing well-defined SMDS communications interface with the customer must be supported by ATM. When a carrier introduces ATM-based switching systems, customers need not see any effect on their SMDS service. Any technology conversion will be made within the carrier networks. Thus, customers reap the benefits of the latest technology development, while maintaining a consistency and continuity in the service they already employ. Because ATM and IEEE 802.6 technology are both cell based and have the same size cells, such conversion will be facilitated.¹⁷

In addition, with the introduction of ATM, SMDS can be combined with other services over a new ATM multiservice communications interface. In this case, the communications interface between the customer and the network is based on ATM protocols for all the services on the multiservice interface, including SMDS and frame relay service. This combination was foreseen in the development of ATM standards. In fact, AAL 3/4 (the ATM Adaptation Layer for SMDS) was specifically designed by ITU-T to carry connectionless services like SMDS. Figure 1.8 depicts the typical platform configuration for carrier-provided ATM-based services.

With its large capacity and multiservice capability, ATM provides SMDS with a faster and more scalable technology platform whose cost can be shared among multiple services. SMDS, along with frame relay PVC, is encouraging the use of high-speed, wide-area public networking in the United States. SMDS and frame relay provide ATM

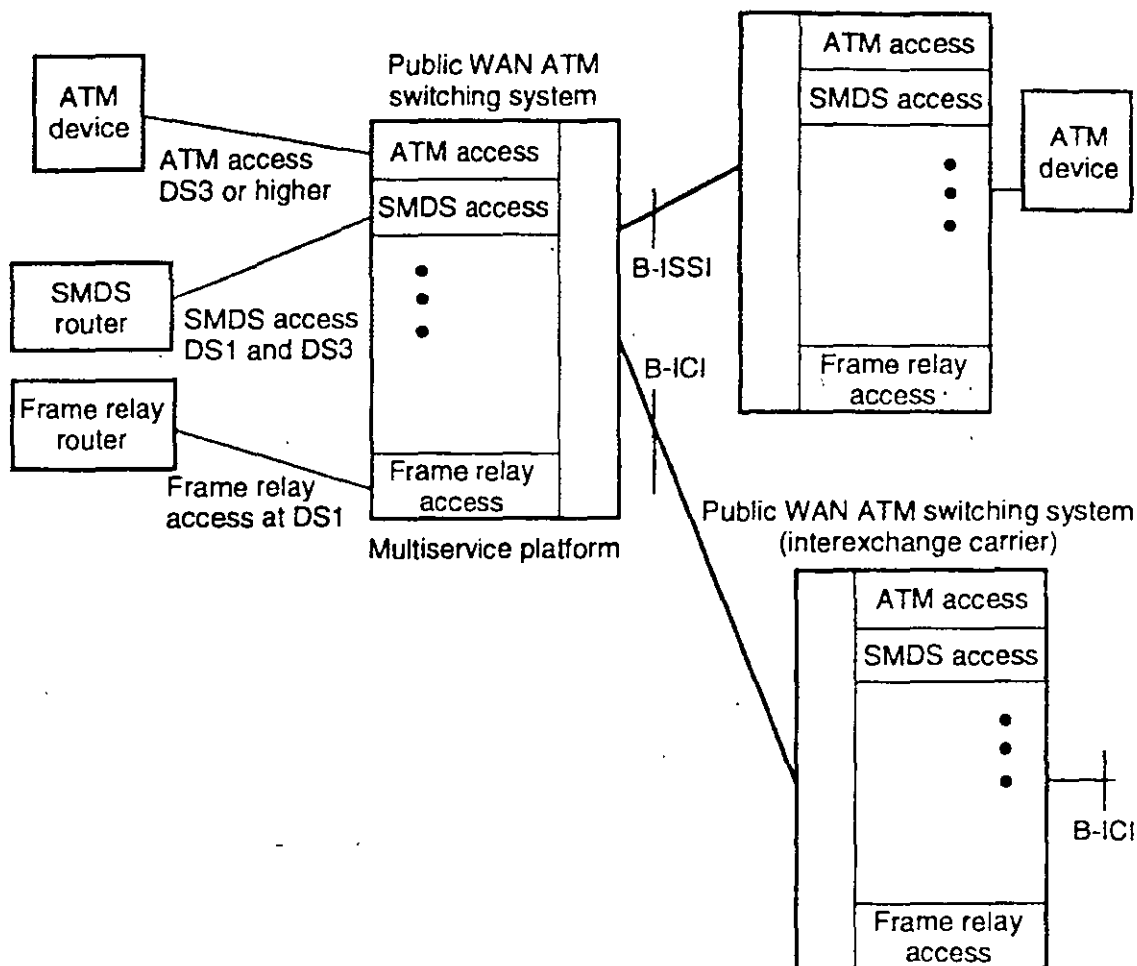


Figure 1.8 Multiservice broadband switching system. B-ISSI = broadband interswitching system interface; B-ICI = broadband interexchange carrier interface.

with significant revenue-producing services that will justify its deployment and allow users and carriers to benefit from the multiservice technology platform.

Frame relay PVC will be a key low-cost, low-overhead broadband data service available in public networks for at least the rest of this decade. The service is currently provided via both frame and ATM switching platforms; ATM simply provides a faster, more scalable platform, as discussed, for SMDS.¹⁸ It appears that frame relay PVC access rates will probably not be extended beyond DS3 (currently, the standards and the deployed services only cover speeds up to 2.048 Mbits/s). This presents the PVC-oriented customer with the possibility of needing to interwork emerging cell relay PVC service with frame relay PVC service. If a user requires PVC service at access speeds of DS3 and above (for example, to aggregate traffic), it is likely that the user will use the ATM cell relay PVC service. This is because customer premises equipment with high-speed wide-area interfaces (e.g., routers) will use ATM technology, thus making cell relay PVC a good choice. As new applications are developed that require these speeds, it is likely that cell relay

PVC service will need to interwork with the users' large installed base of lower-speed wide-area networks for years to come. To meet this need to interwork, the Frame Relay Forum, the ATM Forum, and standards bodies are working on specifications to assure the smooth interworking of these services (ITU-T I.555, in particular).

For the same reasons that carriers are choosing ATM technology (i.e., speed and flexibility), workstation, computer, hub, and LAN manufacturers are turning to ATM for their next-generation networking needs. This is happening because current networks based on Ethernet, FDDI, etc., have limitations when handling the multimedia communications (video, voice, and data) that will flow among future workstations in a network. These manufacturers see global multimedia communications among devices as essential. To meet these networking needs, future workstations and computers will transport user information in ATM cells. Public carriers will offer cell relay service that will transport ATM cells across metropolitan area networks (MANs), across WANs, and internationally as networks evolve. Cell relay service is targeted initially toward high-end users with multimedia needs to transport video, voice, and data across their WANs. When ATM technology extends from the desktop and throughout the network, cell relay service will join SMDS and frame relay as another service that data communications managers can use to support evolving high-bandwidth corporate applications.

Cell relay service is described in Chaps. 5 and 6. Additional aspects of fastpacket are covered in Chap. 7.

1.5 Commercial Availability of ATM Equipment and Network Services

As with any other service, at least three parties are needed to make this technology a commercial reality (if any of these three parties fails to support the service, the service will not see any measurable commercial deployment): (1) carriers must deploy the service, (2) equipment manufacturers must bring user products to the market,* and (3) users must be willing to incorporate the service in their networks. (Some observers add two more forces: agencies supporting R&D and standardization, and the trade press to "educate" the end users.) The early phases of ATM research, including all of the work already accomplished in standards organizations (that is, the topics treated in Chaps. 2 through 10 of this book), cover the first item. The industry activity discussed briefly below and in Chap. 11 covers the second item. The user analysis that will follow (not covered

*In order for item 1 to occur, some vendors must bring out network products; this point refers to user products (see ISDN switches versus availability of cost-effective terminal adapters).

in this book), where users assess applicability, cost, support of embedded base, and manageability, all of it in situ, in their own environment (rather than in a multicolor brochure), covers the third item.

The paragraphs to follow describe industry activities that show encouraging signs of the acceptance of cell relay as a commercially viable networking technology. However, as with all new technologies, there are a number of potential hurdles and roadblocks that can delay or deter its success. History has shown that in spite of industry standards, interoperability problems can exist if different manufacturers implement subsets (or supersets) of the required networking features. Networking hardware may precede the availability of software applications designed to exploit the networking power of ATM, and this may slow user acceptance of cell relay. In addition, advances in existing technologies (e.g., the emergence of "fast" Ethernet) may extend the life cycle of existing products and slow the acceptance of new technologies. These challenges must be met to make ATM cell relay a long-term commercial success.

Vendors are in the process of bringing products to the market. By 1994 there already were several vendors of ATM hubs and a dozen vendors of ATM workstation plug-ins. Some equipment vendors are building stand-alone premises switches; others are adding switching capabilities to their hubs and at the same time are developing ATM adapter cards for workstations to allow them to connect to the hub. Some are also working on bridge-router cards for ATM hubs that enable Ethernet LANs to connect to ATM. About three dozen vendors had announced firm equipment plans by publication time. Over 320 companies have joined the ATM Forum, which is an organization whose goal is to expedite and facilitate the introduction of ATM-based services. PC/workstation cards are expected to become available for about \$1000 per port, although the initial cost was in the \$2800-5000 range.

Carriers are deploying broadband switching systems (BSSs) based on ATM technology to support a variety of services. As noted earlier, ATM is designed to be a multi-service platform. For example, frame relay and SMDS will be early services supported on these platforms; another early service is cell relay service, which allows users to connect their ATM equipment using the native ATM bearer service.

Early entrants, including Adaptive, AT&T Network Systems, Cabletron, Digital Equipment Corporation, Fore Systems, Fujitsu, GDC, Hughes, Newbridge, Stratacom, Sun, SynOptics, and Wellfleet, were demonstrating ready or near-ready products for a variety of user networking needs in 1994. The first products were targeted to the local connectivity environment, but WAN products are also expected soon. Additionally, about a dozen vendors have working carrier-grade switching products.

Hubs and switches to support the bandwidth-intensive applications listed earlier, such as video, are becoming available. Typical premises switches now support 8 to 16 155-Mbits/s ports over shielded twisted pair or multimode fibers [lower speeds (45 or 100 Mbits/s) are also supported]. Some systems can grow to 100 ports. Typical backplane throughput ranges from 1 or 2 Gbits/s, up to 10 Gbits/s. A number of these products support not only PVC but also SVC; some also support multipoint SVC service. Products already on the market (e.g., from Hughes LAN, Synoptics, Newbridge, Adaptive, Fore Systems, etc.) are priced as low as \$1500 per port. Some of the hubs also act as multiprotocol routers, either (1) accepting ATM devices internally for WAN interconnection over SMDS and frame relay networks, (2) accepting ATM devices internally for WAN interconnection over a cell relay network, or (3) accepting traditional devices internally for WAN interconnection over a cell relay network (these are stand-alone ATM multiprotocol routers).

One major push now is in the network management arena. Users need the capability to integrate the support of ATM products into the overall enterprise network, specifically the corporate management system. Some typical features recently introduced include automatic reconfiguration of virtual connections in case of failure, loopback support, performance and configuration management, and Simple Network Management Protocol (SNMP) functionality [with private management information base (MIB) extensions].

Interface cards for high-end workstations (e.g., SPARCstation) are also appearing (e.g., Synoptics, Adaptive, etc.). These typically support 45 Mbits/s (DS3) on twisted-pair cable and 100 or 155 Mbits/s on multimode fiber, consistent with the ATM Forum specification. Some even support prototype 155-Mbits/s connectivity on shielded twisted pair. These boards are already available for as little as \$1250.

Specifically for WAN cell relay service, Sprint has already demonstrated a prototype service operating at the DS3 rate. A three-phase approach has been announced publicly by the company. Phase 1 (1993) entails frame relay interconnectivity with local exchange carriers, Phase 2 (1993–1994) supports PVC cell relay service at the DS3 rate, and Phase 3 (1994–1995) enhances the Cell Relay Service to 155 Mbits/s. AT&T, Wiltel, BellSouth, NYNEX, and Pacific Bell have also announced deployment plans for ATM platforms and for cell relay service. There is strong support for the introduction of cell relay service at the local level. Now users can expect public cell relay service in a number of key metropolitan areas.

In addition to the international and domestic standards, additional details and clarifications are needed to enable the deployment of the technology. To this end, in 1992, Bellcore completed generic require-

ments that suppliers need in order to start building ATM equipment that will enable the BOCs to offer PVC cell relay services. Work on generic requirements for ATM equipment that provides SVC cell relay was completed at Bellcore in 1994. In particular, Bellcore has already published (preliminary) requirements to define nationally consistent cell relay PVC exchange and cell relay PVC exchange access services, including

“Cell Relay PVC Exchange Service,” 1993 [CR PVC exchange service is a public cell relay intra-LATA service offering from local exchange carriers (LECs)]

“Cell Relay PVC Exchange Access CRS (XA-CRS),” 1993 [a PVC XA-CRS is provided by an LEC to an interexchange carrier (IC) in support of the IC’s inter-LATA cell relay PVC offering]

“Cell Relay SVC Exchange Service,” 1993

The Framework Advisories, Technical Advisories, and Technical Requirements can be used by (1) LECs interested in providing nationally consistent cell relay PVC exchange service to their customers, (2) suppliers of ATM equipment in the local customer environment (e.g., ATM LANs, ATM routers, ATM DSUs, ATM switches), and (3) suppliers of ATM equipment in LEC networks.

The development of nationally consistent LEC cell relay (as well as an exchange access cell relay) service is critical to provide a consistent set of service features and service operations for customers who will want to use the service on a national basis. The following phases of nationally consistent service have been advanced. It is possible that LECs may be offering “pre-nationally consistent” cell relay PVC to meet customers’ near-term demand for the service in the late 1993–early 1994 period. These carriers are expected to support a nationally consistent cell relay PVC exchange service at some point thereafter.

- Phase 1.0: Nationally consistent cell relay PVC exchange service based on a core set of service features by the fourth quarter of 1994. The core set is proposed to be a subset of the preliminary generic requirements published by Bellcore in 1993.
- Phase 2.0: Nationally consistent cell relay PVC exchange service based on generic requirements published by Bellcore in 1994 by the second quarter of 1995. Phase 2.0 builds on the capabilities of Phase 1.0 and supports expanded capabilities in some areas, such as traffic management, congestion management, and customer network management.

- Phase 3.0: This will see the initial support of a cell relay SVC exchange service in mid to late 1995 based on generic requirements expected to be published in 1994.

Figure 1.9 depicts the set of Bellcore generic requirements in support of ATM, SMDS, cell relay, and frame relay.¹⁷ These are just some of the key documents that form the foundation for ATM. Standards bodies such as the ITU-T and ANSI (American National Standards Institute) T1S1, and industry bodies such as the ATM Forum and the Frame Relay Forum also publish related documents.

1.6 Typical Examples of Cell Relay Usage in an Enterprise Context

1.6.1 Front-end and back-end usages

Cell relay/ATM is being contemplated at the local-area network level as well as the wide-area network level. Several approaches have been followed by vendors:

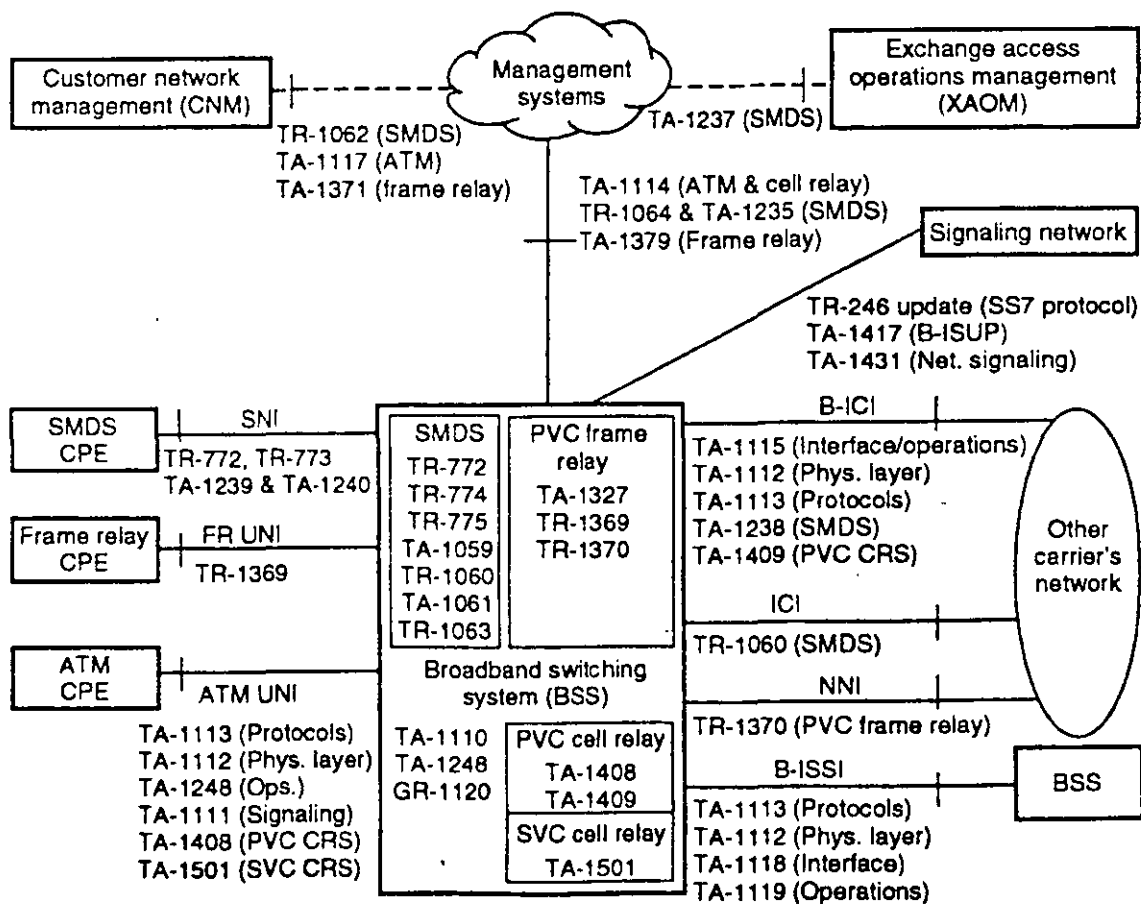


Figure 1.9 ATM, SMDS, cell relay, and PVC frame relay generic requirements.

1. TR-NWT-00246, Bellcore Specification of Signaling System 7, B-ISUP, Issue 2, December 1993.
2. TR-TSV-000772, Generic System Requirements in Support of SMDS, May 1991.
3. TR-TSV-000773, SMDS Requirements, Objectives, and Interfaces, Revision 1, December 1993.
4. TR-TSV-000774, SMDS Operations Technology Network Element Generic Requirements, Issue 1, March 1992, Supp. 1, March 1993.
5. TR-TSV-000775, Usage Measurement Generic Requirements In Support of Billing for Switched Multi-Megabit Data Service, Issue 1, June 1991.
6. TA-NWT-001248, Generic Operations Requirements for Broadband Switching Systems, Issue 2, October 1993.
7. TA-TSV-001059, Generic Requirements for SMDS Networking, Bellcore, Issue 2, August 1992.
8. TR-TSV-001060, Switched Multi-Megabit Data Service Generic Requirements for Exchange Access and Intercompany Serving Arrangements, Issue 1, December 1991, and Revision 1, August 1992; Revision 2, March 1993.
9. TA-TSV-001061, Operations Technology Network Element Generic Requirements in Support of Inter-Switch and Exchange Access SMDS, Issue 1, May 1991.
10. TR-TSV-001062, Generic Requirements for SMDS Customer Network Management Services, Bellcore, Issue 1, March 1992.
11. TR-TSV-001063, Operations Technology Network Element Generic Requirements in Support of Exchange Access SMDS and Intercompany Serving Arrangements, Issue 1, March 1992; Revision 1, March 1993.
12. TR-TSV-001064, SMDS Phase 1 Operations Information Model, December 1993.
13. TA-NWT-001110, Broadband ISDN Switching System Generic Requirements, Issue 2, July 1993.
14. TA-NWT-001111, User to Network Access Signaling Requirements, July 1993.
15. TR-NWT-001112, Broadband ISDN User to Network Interface and Network Node Interface Physical Layer Generic Criteria, July 1993.
16. TA-NWT-001113, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 2, July 1993.
17. TA-NWT-001114, Generic Requirements for Operations Interfaces Using OSI Tools: Broadband ISDN Operations, Issue 2, October 1993.
18. TA-NWT-001115, Broadband InterCarrier Interface (B-ICI) Requirements, September 1993.
19. TA-NWT-001117, ATM Customer Network Management (CNM), September 1993.
20. TA-TSV-001118, Broadband InterSwitching System Interface (B-ISSI) and Network Generic Requirements, July 1993.
21. TA-NWT-001119, B-ISSI Operations, December 1993.
22. GR-1120-CORE, Guide to Generic Requirements for Usage Information to Support Billing for ATM Broadband Networking, Issue 1, December 1993.
23. TA-NWT-001235, Exchange Access SMDS Operations Interface Model, April 1993.
24. TA-TSV-001237, A Framework for High Level Generic Requirements for SMDS Exchange Access Operations Management Services, July 1993.
25. TA-TSV-001238, SMDS 155 Mbps ATM B-ICI, December 1992.
26. TA-TSV-001239, Low Speed SMDS Access via Data Exchange Interface (DXI), June 1993.
27. TA-TSV-001240, Frame-Based Access to SMDS via SRI, June 1993.
28. TA-NWT-001248, B-ISDN Network Operations Criteria, Issue 2, October 1993.
29. FA-NWT-001327, Frame Relay NE Operations Functional Requirements, Bellcore, Issue 1, December 1992.
30. TR-TSV-001369, Frame Relay (PVC) Exchange Service Definition, May 1993.
31. TR-TSV-001370, Exchange Access Frame Relay (PVC) Service Definition, May 1993.
32. TA-TSV-001371, Frame Relay (PVC) Customer Network Management Service, September 1993.
33. TA-NWT-001379, Frame Relay Network Operations Using OSI, July 1993.
34. TA-TSV-001408, Generic Requirements for Exchange PVC Cell Relay Service, Issue 1, August 1993.
35. TA-TSV-001409, Generic Requirements, Issue 1, November 1993.
36. TA-NWT-001417, B-ISUP Generic Requirements, Issue 1, February 1994.
37. TA-NWT-001431, CCS Network Signaling Specification Supporting B-ISDN Generic Requirements, Issue 1, May 1994.
38. TA-NWT-001501, Generic Requirements for Exchange SVC Cell Relay Service, December 1993.

Figure 1.9 (Continued)

1. Use of ATM technology between traditional local or remote LAN hubs; Fig. 1.10 shows a case of interconnection of remote hubs. (The LAN hubs are implicit in the figure.)
2. Introduction of ATM cards on traditional routers for access to a public cell relay service (see Fig. 1.11).
3. Introduction of ATM-based LAN hubs, extending ATM all the way to the desktop, for front-end applications (see Fig. 1.12).
4. Development of private-enterprise ATM switches to support generic corporate networking.
5. Development of carrier-grade multiservice ATM switches (also known as broadband switching systems) to support services such as cell relay service, frame relay service, and SMDS.
6. Development of related equipment (for example, Fig. 1.13 depicts usage in a channel extension environment).

Some industry proponents expect to see Fortune 1000 users passing the majority of their LAN-to-WAN traffic through premises-based ATM switches by 1997. Approximately 50 percent of the ATM traffic in these companies is expected to be in support of LAN interconnection, for LANs serving traditional business applications, and for traditional enterprise data applications, such as mainframe channel extension; the other 50 percent of the traffic is expected to be split fairly evenly among application supporting real-time video, imaging, real-time voice, and multimedia.

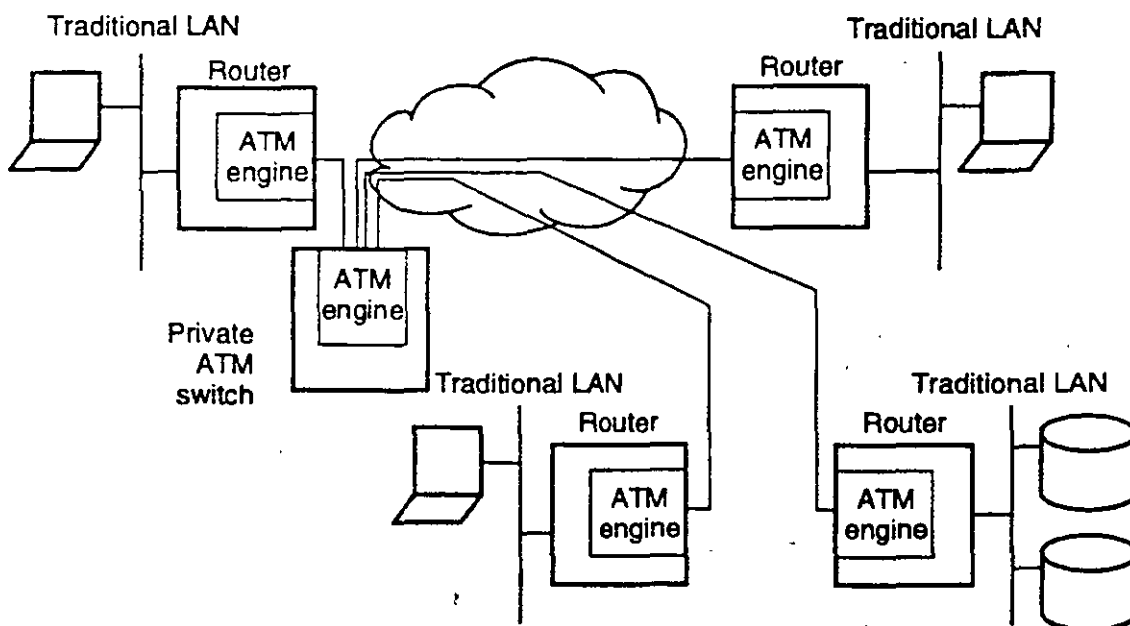


Figure 1.10 Private ATM technology to interconnect dispersed LAN hubs. ATM engine = the logic implementing ATM, control, and, optionally, user plane protocols.

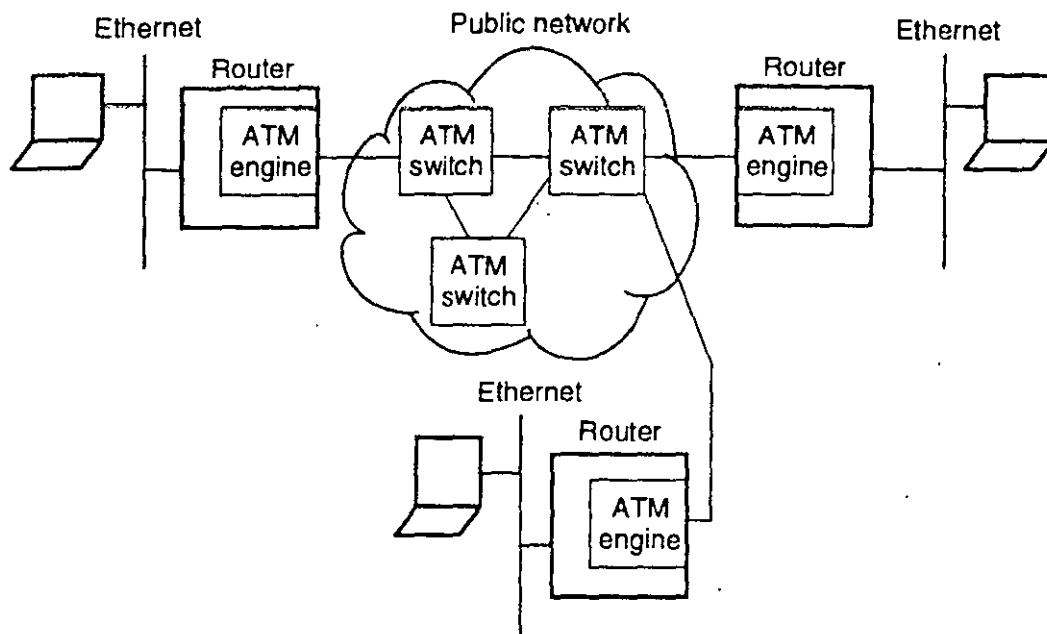


Figure 1.11 Routers used in conjunction with a public cell relay service. ATM engine = the logic implementing ATM, control, and, optionally, user plane protocols.

Figure 1.14 depicts a typical "full-blown" ATM/cell relay arrangement for both WAN and LAN applications. This supports ATM to the desktop for such applications as desk-to-desk videoconferencing and multimedia. Figure 1.15 depicts an example of the protocol machinery across a router/public switch arrangement that is expected to be a common deployment scenario in client/server environments. Figure 1.16 depicts an example in




Figure 1.12 ATM to the desktop.

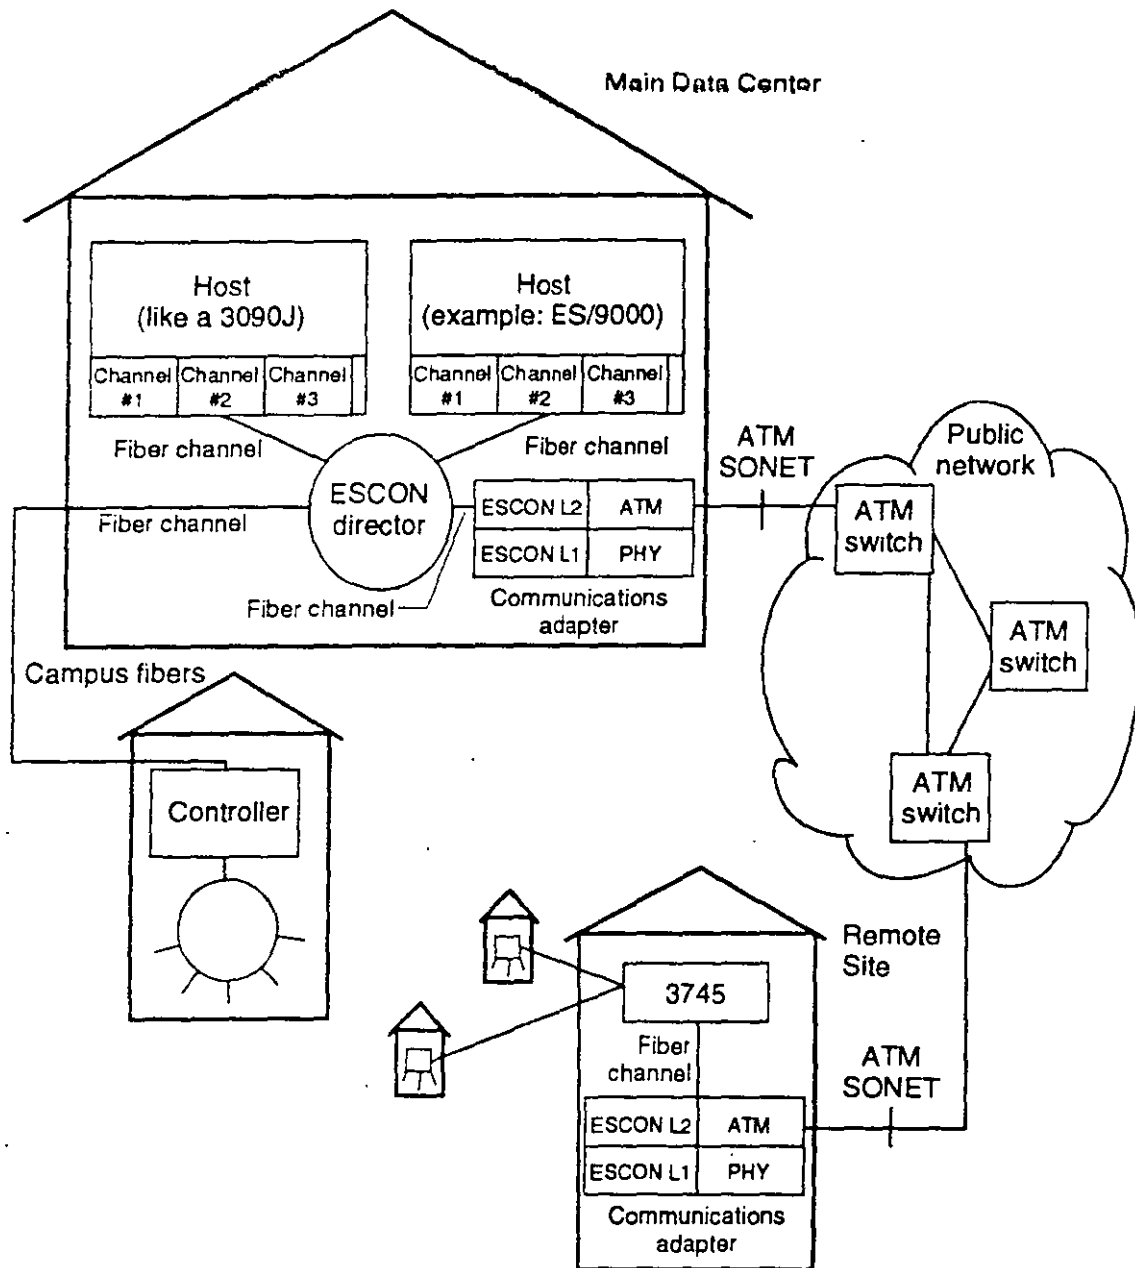


Figure 1.13 Channel extension via ATM services.

a videoconferencing application, also from a protocol point of view. Figure 1.15 shows an example in a corporate network supporting business imaging.

Figure 1.18 depicts a more complete enterprisewide use of cell relay service, while employing a public WAN CRS network. For this example, ATM-ready workstations and devices connected to an ATM-based hub with ATM WAN router capabilities (the router could also be a separate device) can get direct access to the ATM WAN. Some of the hub and router vendors are taking this path to the market. The figure also shows that traditional LAN users can employ an ATM-ready router to obtain the benefit of cell relay WAN services without having to replace their




Figure 1.14 Example of usage of ATM in an enterprise network.

desktops or in-house wiring. It also depicts another route to the market, followed by some of the more sophisticated multiplexer manufacturers: The multiplexer can connect traditional data devices, mainframe channels, and video to a cell relay WAN network by supporting ATM on the trunk side. Some of these multiplexers also support traditional LANs on the house side over a frame relay interface. (*Note: Carrier-deployed ATM "service nodes" in close proximity of the user location but on the*

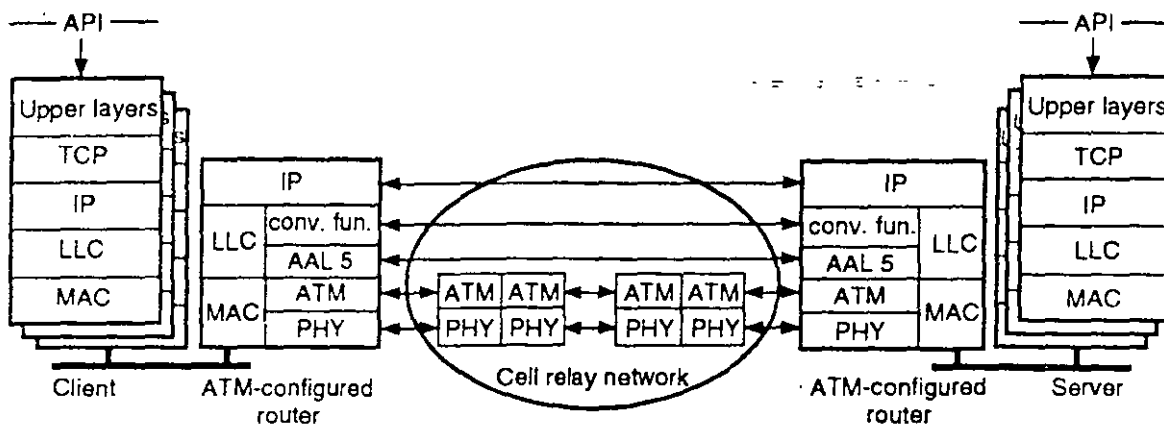


Figure 1.15 Typical corporate application from a protocol-stack point of view. conv.fun. = convergence function.

network side of the interface support these same services plus LAN emulation service.)

Figure 1.19 depicts some user applications of cell relay service in the case where the user wants to develop a private ATM/cell relay service WAN. Note the need to (1) install privately managed switches, (2) use dedicated high-speed WAN lines, and (3) backhaul remote locations to a remote switching site. Public cell relay service may prove less demanding in terms of users' responsibility. Hybrid arrangements are also possible.

1.6.2 Client/server issues

The client/server architecture being put in place in many organizations is truly distributed in the sense that the corporate user has access to data regardless of where the data are located, be they on a system in another campus, another city, another state, or another continent. Client/server applications require extensive interchange of data blocks, often entailing multiple transactions. Low end-to-end delay is critical in making client/server computing possible.¹⁹

Applications requiring large transfers (e.g., 50–100 kbits) are not unusual in these environments, particularly for imaging video, and

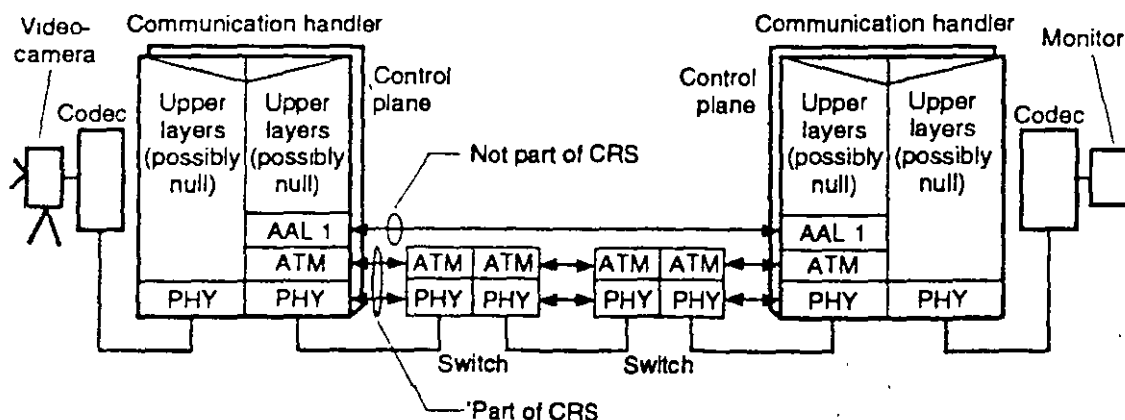


Figure 1.16 Example of video application over ATM/cell relay arrangement.




Figure 1.17 Use of ATM/cell relay to support imaging.

multimedia applications (the last two applications also have stringent delay sensitivities). A 100-MByte data unit across the application programming interface (API) running on a remotely located LAN-resident server is segmented into approximately 60 Ethernet frames. Each Ethernet frame is then segmented into approximately 30 cells by an ATM-configured router for delivery over a public cell relay network.

Some wish to clarify the implications of the interplay between the network (or private ATM switch) performance in terms of cell loss/mutilation, response time, latency, and the end-to-end error correction protocols (e.g., included in TCP). For example, if one of the 29 cells that

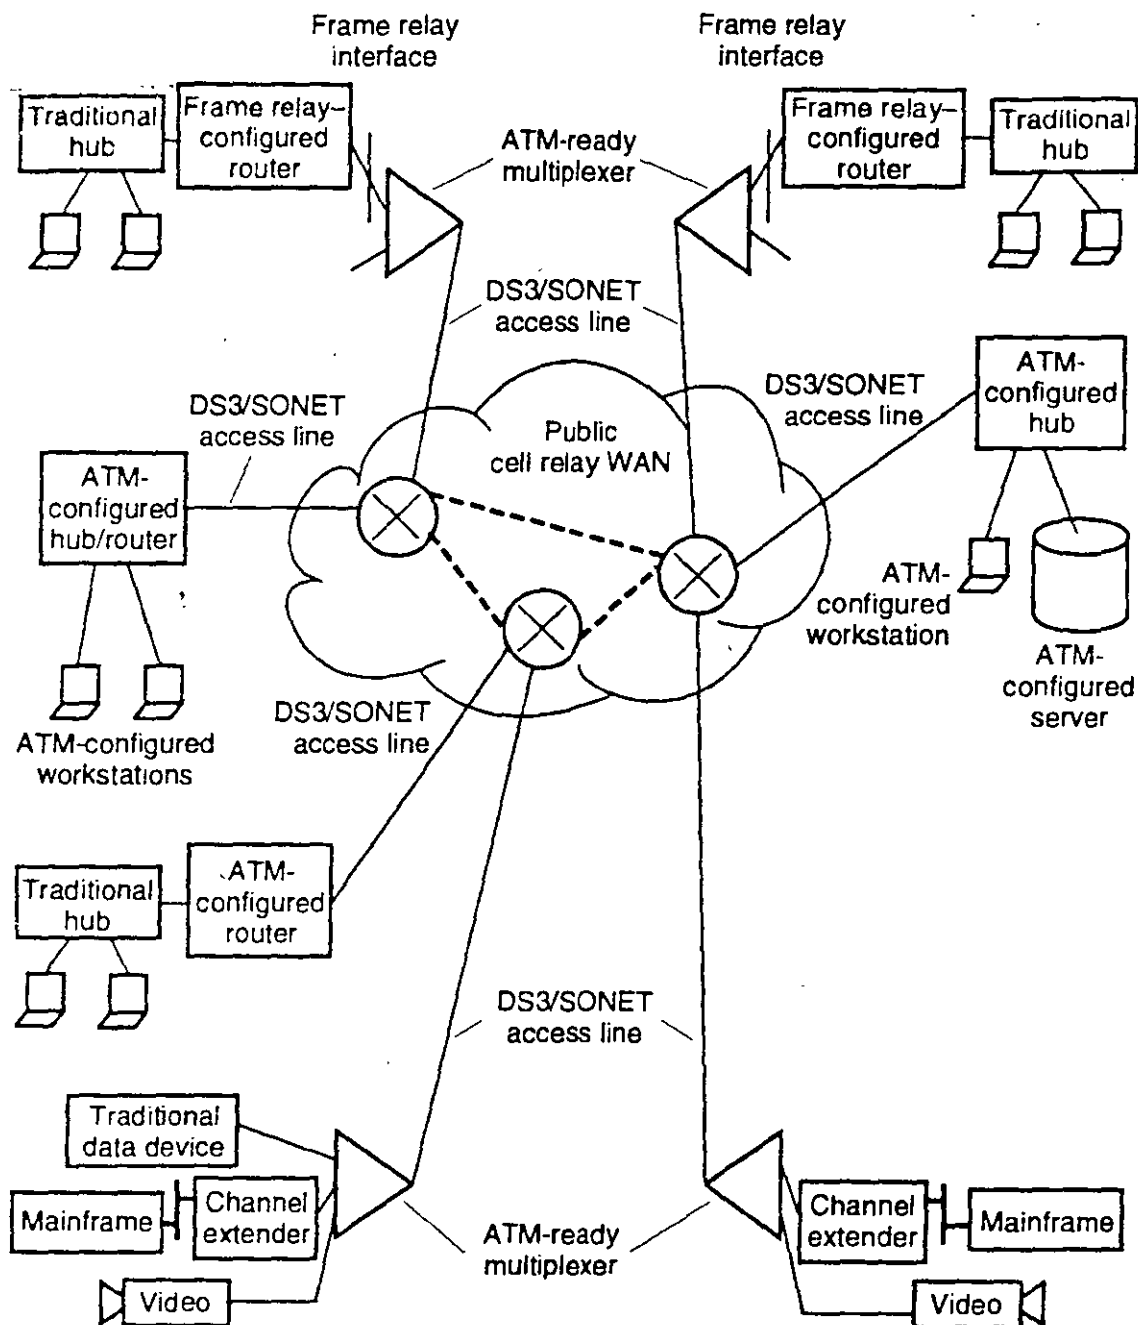


Figure 1.18 CRS to support enterprise networking in WAN applications.

made up a frame is lost, the entire frame (30 cells) needs to be retransmitted by TCP. Under heavy user load as well as coterminous ATM switch overload (whether public or private), the combination of client/server architecture and ATM communication could result in degradation, saturation, or instability. A number of simulation-based studies have shown that, when properly engineered, the network should behave as expected.

Chapter 9 covers ATM-based LANs, while Chaps. 11 and 12 cover other details pertaining to the deployment of ATM in users' environments.

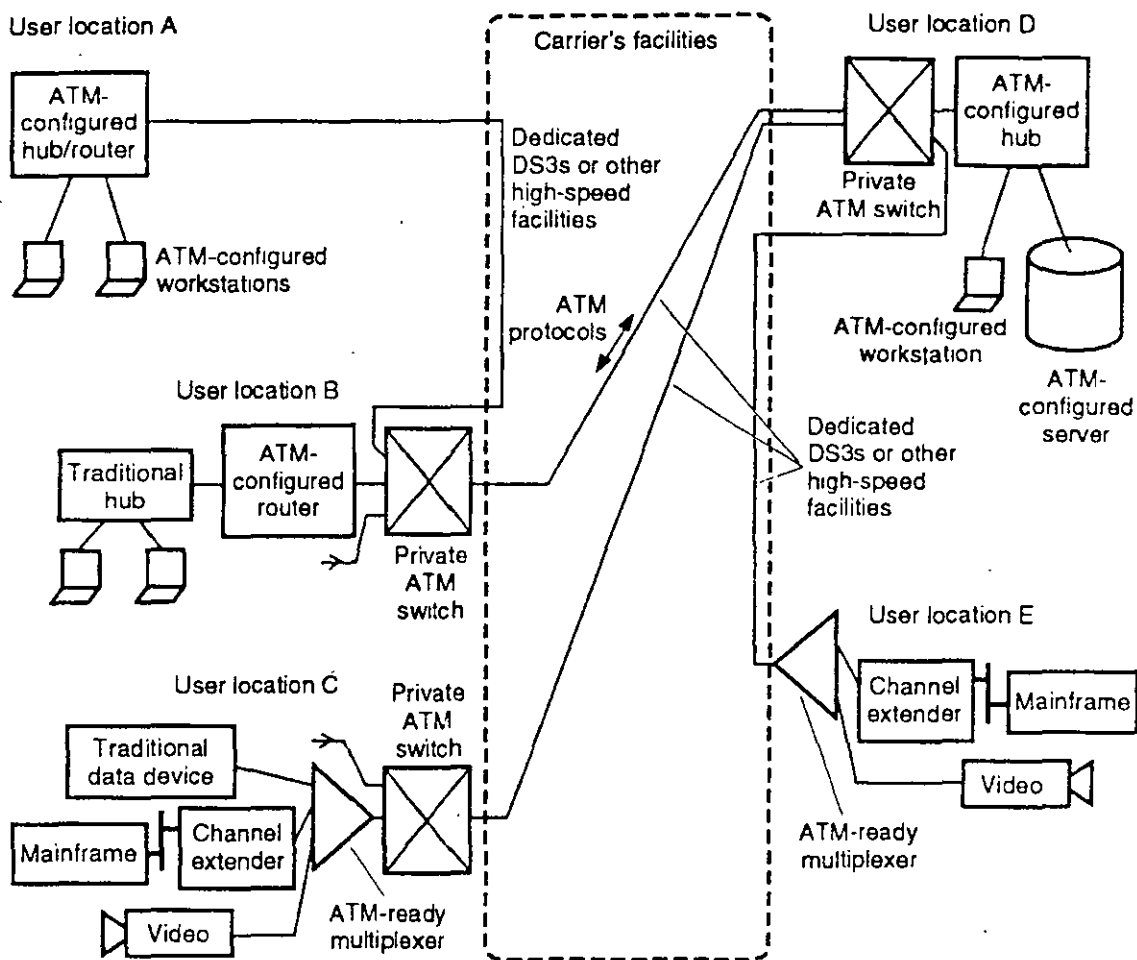


Figure 1.19 CRS to support enterprise networking in WAN applications (private network).

1.7 The Value of Standards

It is a well-known fact that standards benefit not only an industry but an entire economy. Many industries would not have arisen (e.g., the VCR industry, the CD audio industry, television, radio, etc.) if it were not for standards. Standards make a level playing field, fostering competition; this is in contrast to vendor proprietary approaches, where only those vendors have access to a market or have disproportionate control of it. However, for a standard to be effective, it must be widely available, without restrictions on promulgation, discussion, commentary, proliferation, distribution, and duplication. In our opinion, a standard is not an open standard if it is restricted, copyrighted, or patented, if it represents someone's intellectual property, or if it is "owned" by someone (sounds mighty close to a proprietary system to us!) because all of these factors frustrate the exact purpose for which the standard aims to exist (or has a reason to exist). There is much discussion at large about "free trade," "free movement of information," and "lack of censorship."

Standards are developed by industry consensus. This means that representatives from many companies, typically several dozen, have input into the standard. These proceedings can go on for years, and the representatives of these companies travel to many meetings and invest company resources back home to work on technical issues, prepare contributions, review contributions, and act as editors, chairs, etc. In the end, no one individual or institution should be able to claim ownership. There must be a free flow of specification information. Developers must be able to obtain copies. Programmers must be able to use the material. Documentaries must be able to write down the standard and comment on how they implemented various aspects. Educators must be able to discuss the standard and promulgate it to users. Otherwise, such a standard may go nowhere, as many examples of voluminous standards from the (late) 1980s illustrate.

Given this philosophical imperative, and in spite of the less than eloquent case made in these terse paragraphs, we have taken the approach of discussing here, in this text, the dozens of standards that support cell relay service and ATM, regardless of their source. In the end, all stand to benefit from such open and uninhibited discussion at the birth of this new technology. Since this book is only a brief synopsis of the estimated 15 cubic feet of standards material that forms the basis for ATM (ITU-T, ANSI T1S1, ATM Forum, Frame Relay Forum, Bellcore, and other documents), the reader is constantly referred to the original documents for the full-scale detail. In particular, developers, who stand to benefit commercially from their efforts, should definitely refer to the original documentation for the necessary level of detail. The purpose of this book is strictly pedagogical and for the end user. Each of the more than 100 documents alluded to earlier can be obtained from the original source for \$100 or less.

References

1. D. Minoli, *Enterprise Networking—Fractional T1 to SONET, Frame Relay to BISDN*, Artech House, Norwood, Mass., 1993.
2. D. Minoli, *1st, 2nd, and Next Generation LANs*, McGraw-Hill, New York, 1994.
3. D. Minoli, *Imaging in Corporate Environments: Technology and Communication*, McGraw-Hill, New York, 1994.
4. D. Minoli and B. Keinath, *Distributed Multimedia: Through Broadband Communication Services*, Artech House, Norwood, Mass., 1994.
5. D. Minoli et al., *ATM Layer Bearer Service/Cell Relay Service Extended Stage 1 Description for Public Service Offerings, T1S1.5/93-021*, February 1993.

6. T1S1.5/93-52, *Broadband Aspects of ISDN Baseline Document*, T1S1 Technical Subcommittee, August 1990, Chief Editor: Erwin Fandrich.
7. D. Minoli, "The New Wide Area Technologies: SMDS and B-ISDN," *Network Computing*, pp. 88ff., August 1991.
8. D. Minoli, "Understanding ATM—Part 1," *Network Computing*, pp. 128ff., Oct. 15, 1992.
9. D. Minoli, "Understanding ATM—Part 2," *Network Computing*, pp. 156ff., Nov. 15, 1992.
10. D. Minoli, "Third-Generation LANs," *UNIX Expo 92 Proceedings*, Bruno Blemheim Inc., Fort Lee, N.J., 1992.
11. D. Minoli, "Third Generation LANs," *Proceedings of Texpro 1993*, Pacific Bell, San Francisco, April 1993.
12. D. Minoli, "Cell Relay and ATM," WAN Insert to *Network Computing and Communications Week*, pp. 22 ff., August 1993.
13. D. Minoli, "Wide Area Networking for Multimedia?," WAN Insert to *Network Computing and Communications Week*, pp. 60ff., August 1993.
14. D. Minoli, "Broadband Integrated Services Digital Network," Datapro Communications Series: *Broadband Networking*, Report #2890, April 1992.
15. D. Minoli, "ATM and Cell Relay Concepts," Datapro Communications Series: *Broadband Networking*, Report #2880, April 1992.
16. J. T. Johnson, "Applications Catch Up to ATM," *Data Communications*, pp. 41–42, July 1993.
17. F. Gratzner and S. Walters, "ATM and Fast Packet Services—Perfect Together," *Bellcore DIGEST*, vol. 10 (6), pp. 3ff., 1993.
18. D. Minoli, "Designing Scalable Networks," *Network World Collaboration*, pp. 17ff., January 10, 1994.
19. D. Minoli, *Analyzing Outsourcing: Reengineering Information and Communication Systems*, McGraw-Hill, New York, 1995.

Asynchronous Transfer Mode

As noted in Chap. 1, ATM is a new transport and switching technology that can be used in a variety of telecommunications and computing environments. ATM is a cell-based technology, designed to support user applications requiring high-bandwidth, high-performance transport and switching. This chapter provides a summary description of the peer-to-peer ATM protocol at the user-network interface in support of cell relay service and other ATM capabilities. It describes functionality in the User Plane, thereby enabling a PVC service. The addition of Control Plane support enables the user to obtain an SVC service; the operation of the ATM Layer in the Control Plane is nearly identical to that of the User Plane (the Control Plane functionality is discussed in Chap. 4). Some aspects of the underlying transport mechanism are also briefly covered at the end of the chapter.

A description of general aspects of the access interface(s) between the user and the network is followed by a description of the protocol across such an interface. The protocols and related requirements are associated with two functional OSIRM layers: the Data Link Layer and the Physical Layer. Figure 2.1 depicts this peer-to-peer protocol view of the service. Figure 2.2 depicts communication through a set of network peers. As described in ITU-T Recommendation X.210, *Open Systems Interconnection, Layer Service Definition Conventions*,¹ the service defined at the Data Link Layer also relies on the capabilities of the Physical Layer. This view of cell relay service in general and of the ATM protocol in particular establishes requirements on what an entity in the ATM Layer (whether the entity is in the network or in the user's equipment),




Figure 2.1 Peer entities across the user-network interface. TC = Transmission Convergence Sublayer; PMD = Physical Medium-Dependent Sublayer.

defined in ITU-T Recommendation I.361, *B-ISDN ATM Layer Specification*,² and in T1S1.5/92-410, *Broadband ISDN—ATM Layer Functionality and Specification*,³ expects the remote peer entity to support. The physical aspects of the UNI supporting cell relay service are based on the B-ISDN UNI defined in ITU-T Recommendation I.432, *B-ISDN User-Network Interface—Physical Layer Specification*⁴ and on the ATM Forum's *UNI Specification*⁵ for public UNIs. This discussion only provides an overview; the reader interested in additional details should consult Refs. 6 and 7.

This chapter only covers the interface between user equipment and a public network; intra-CPE interfaces (for example, for ATM-based LANs), although similar in many respects to the interface between the CPE and the network, are not addressed. Table 2.1 depicts some of the key ITU-T standards in support of ATM in general and the peer-to-peer cell relay protocol in particular.

2.1 Access Interface

This section defines the concept of access interface. This is accomplished by defining an access reference configuration, functional entities (groups), and logical reference points.

An *access reference configuration* for B-ISDN is defined in ITU-T Recommendation I.413, *B-ISDN User-Network Interface*.⁸ This configu-

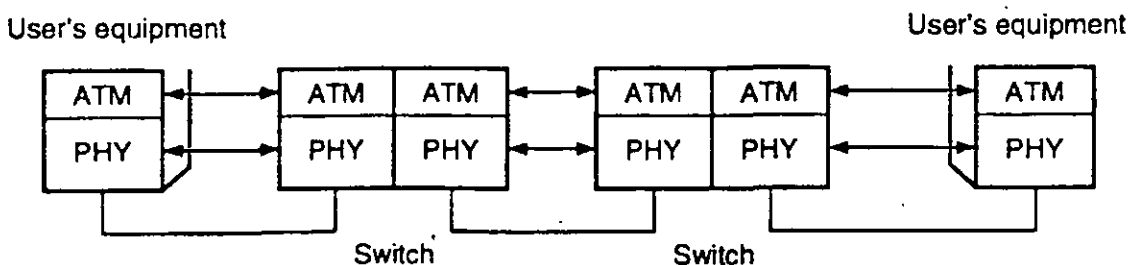


Figure 2.2 Cascaded ATM entities.

TABLE 2.1 Key ITU-T Standards in Support of ATM

F.811	B-ISDN Connection-Oriented Bearer Service
F.812	B-ISDN Connectionless Bearer Service
I.113	B-ISDN Vocabulary of Terms
I.121R	Broadband Aspects of ISDN [Basic Principles and Evolution]
I.150	B-ISDN ATM Functional Characteristics
I.211	B-ISDN Service Aspects
I.311	B-ISDN General Network Aspects
I.321	B-ISDN Protocol Reference Model and Its Applications
I.327	B-ISDN Functional Architecture Aspects
I.356	Quality of Service Configuration and Principles
I.361	B-ISDN ATM Layer Specification
I.362	B-ISDN AAL Functional Description
I.363	B-ISDN AAL Specification
I.371	Traffic Control and Resource Management
I.374	Network Capabilities to Support Multimedia
I.413	B-ISDN UNI
I.432	B-ISDN UNI Physical
I.555	Interworking with Frame Relay
I.555	Interworking with ISDN
I.610	B-ISDN OAM Principles
I.cls	Support for Connectionless Data Service on B-ISDN
Q.93B (now Q.2931)	B-ISDN Call Control
Q.SAAL 1 and 2 (now Q.2110 and Q.2130)	Signaling AALs [Q.2110, Service-Specific Connection-Oriented Protocol (SSCOP); Q.2130, Service-Specific Coordination Function (SSCF)]

ration forms the basis for the definition of access interfaces supporting cell relay service.

Functional entities are logical abstractions of functions typically found in network equipment and in users' equipment, also known as customer premises equipment (CPE). Public network switch-termination functions are modeled by the broadband line terminator/exchange terminator (B-LT/ET) functional group. The CPE is modeled by the broadband network termination 2 (B-NT2) functional group; NT2 functions include concentration, switching, and resource management. Broadband network termination 1 (B-NT1) functions support line termination, line maintenance, and performance monitoring. The broadband terminal

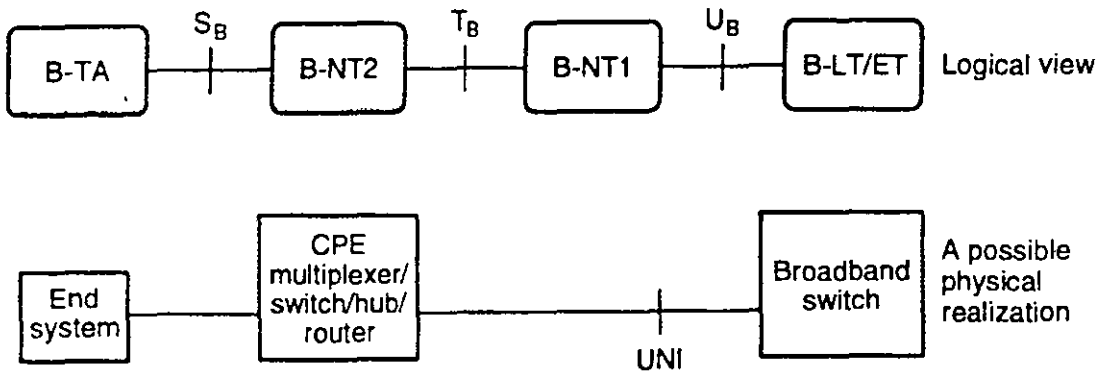


Figure 2.3 B-ISDN access reference configurations. B-TA = broadband terminal adapter.

equipment, such as a workstation, is modeled by the broadband terminal equipment (B-TE) functional group.

Logical reference points are defined between B-ISDN functional entities. T_B is the logical reference point between a B-NT2 and a B-NT1. U_B is the logical reference point between a B-NT1 and a B-LT/ET. In this description, the UNI is associated with the U_B reference point. See Fig. 2.3.

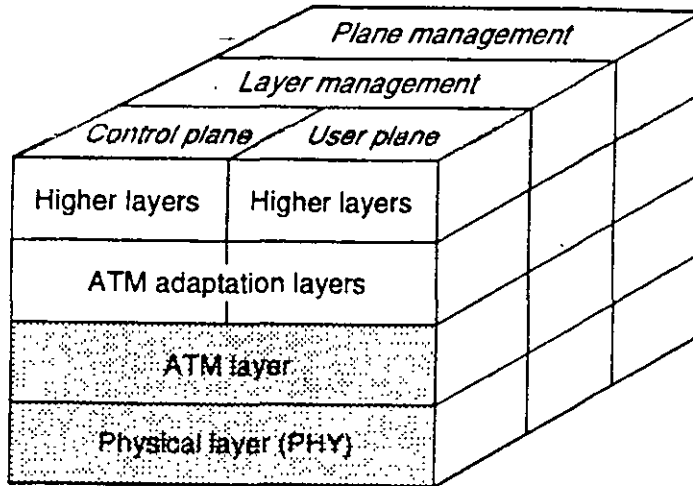
Note: This description only covers the case where there is a single B-NT2 (however, several B-TEs may be connected to the B-NT2). The case where the B-NT2 is null and there are several B-TEs connected to a single UNI is not addressed in the initial view of ATM services in the United States.

2.2 ATM-Level Protocol

2.2.1 Overview

UNI protocols define the way in which users communicate with the public network for the purpose of accessing the service provided by the network. Figure 2.4 illustrates the B-ISDN Protocol Reference Model, which is the basis for the protocols that operate across the UNI (this is another common way to represent the protocol model of Fig. 1.3). The B-ISDN Protocol Reference Model is described in ITU-T Recommendation I.121. This model is made up of three planes, already discussed in Chap. 1: the User Plane, the Control Plane, and the Management Plane. Table 2.2 provides a summary of the functions supported by each plane.

The UNI specified at this level includes the functions associated with the User Plane at the Physical Layer and the ATM Layer. The Physical Layer provides access to the physical medium for the transport of ATM cells. It includes methods for mapping cells to the physical medium (i.e., the Transport Convergence Sublayer) and methods dependent on the physical medium (i.e., the Physical Medium-Dependent Sublayer). The ATM layer provides for the transport of cells between end-user locations. An ATM cell contains a header that contains control information, iden-



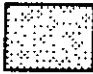
 Protocol layers relevant to cell relay service

Figure 2.4 B-ISDN protocol reference model.

tifies the type of cell, and contains routing information that identifies a logical channel (i.e., a VPC or a VCC) over which the cell is to be forwarded.

The interactions of each protocol layer with other layers and with its own layer management are described in terms of primitives. Primitives describe abstractly the logical exchange of information and control

TABLE 2.2 Functions of Various Planes of the Protocol Model

User Plane	Provides for the transfer of end-user information. It consists of the Physical Layer and the ATM Layer. The model also includes ATM Adaptation Layers and higher layers necessary for each end-user application. (Because these layers are specific to each application, they are not part of the cell relay service described here and in Chap. 5.)
Control Plane	Provides for the transfer of information to support connection establishment and control functions necessary for providing switched services. The Control Plane shares the ATM and Physical Layer with the User Plane. Also, it contains AAL procedures and higher-layer signaling protocols. The Control Plane is discussed in Chap. 5.
Management Plane	Provides for operations and management functions and the capability to exchange information between the User and the Control Planes. The Management Plane is made up of the <i>Layer Management</i> (for layer-specific management functions such as detection of failures and protocol abnormalities) and the <i>Plane Management</i> (for management and coordination functions related to the complete system). The Management Plane is discussed in Chap. 10.

through a service access point, while not imposing any constraint on the implementation. Figures 2.5, 2.6, and 2.7 depict some aspects of this protocol machinery.

2.2.2 ATM Layer

The ATM Layer provides for the transport of fixed-size cells between end-user locations. It is implemented in users' equipment (workstations, routers, private switches, etc.) and in network equipment. ATM cells from end users are forwarded across virtual connections through the public network. These connections are provided at subscription time or in real time via signaling (as described in Chap. 4). The ATM Layer also provides multiplexing functions to allow the establishment of multiple connections across a single UNI.

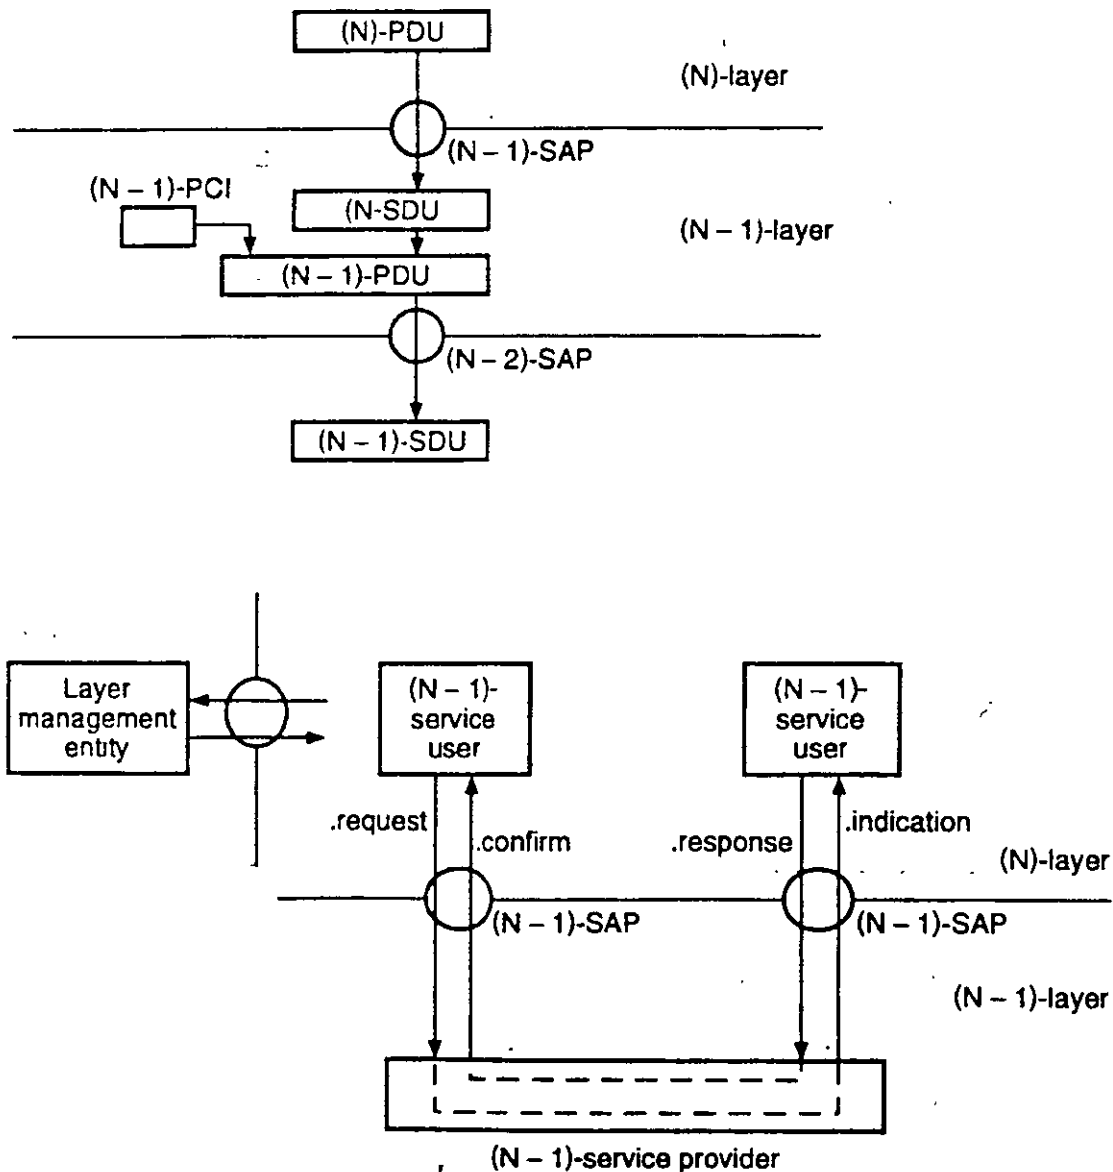


Figure 2.5 SAPs (top) and primitives (bottom). SAP = service access point; PDU = protocol data unit; SDU = service data unit; PCI = protocol control information.

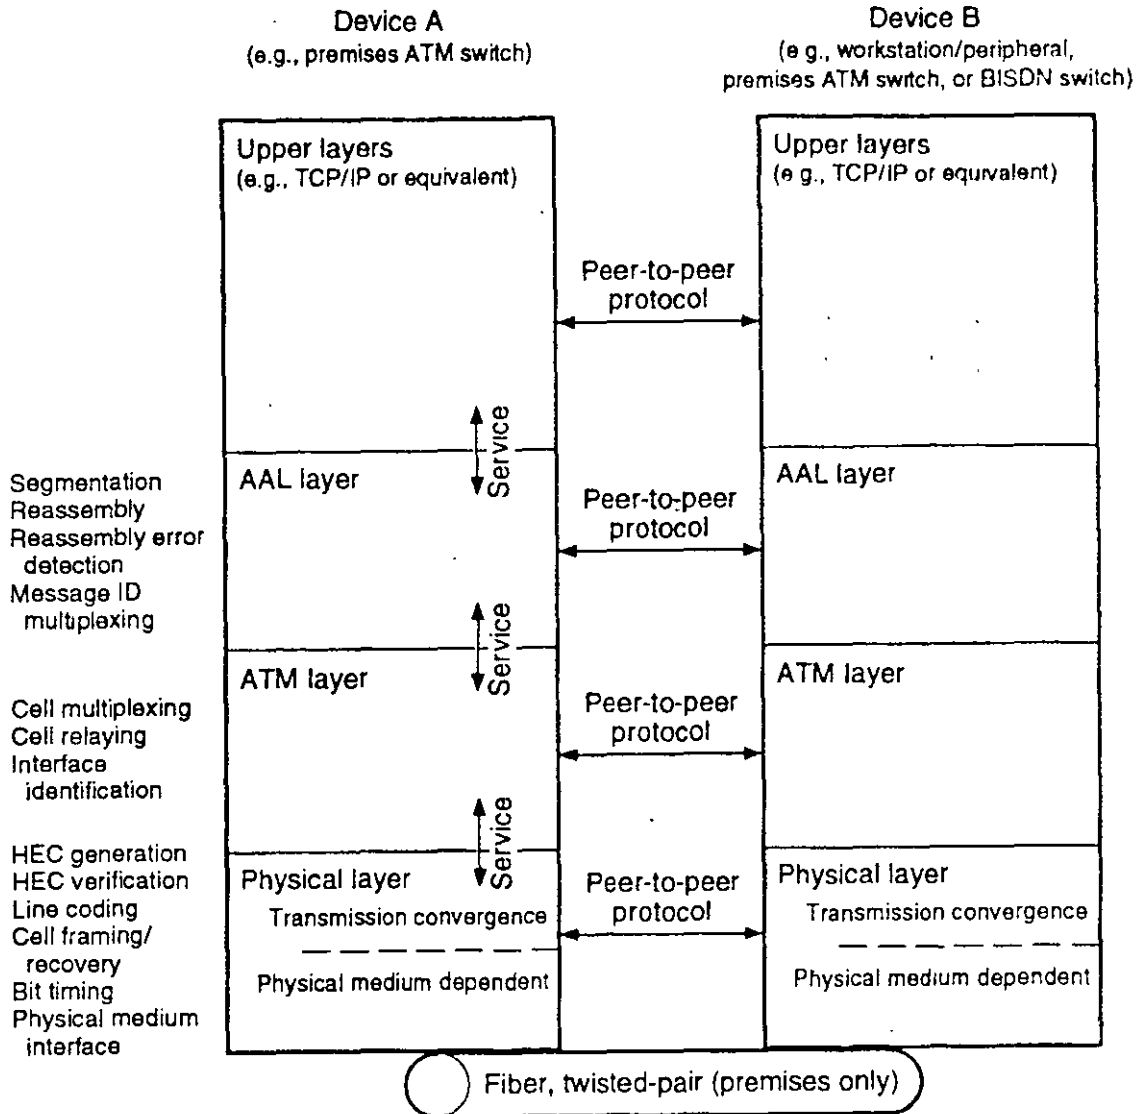


Figure 2.6 ATM protocols.

Service provided to the upper layer. The ATM-Layer service is based on fixed-size ATM service data units which consist of 48 octets. It provides for the transparent transfer of ATM SDUs between communicating peer upper-layer entities. To accomplish this, the ATM Layer generates a 53-octet ATM cell by prepending a 5-octet header to the ATM SDU. The header contains routing and protocol control information. The interaction between the ATM Layer and its service users is implemented by the primitives shown in Table 2.3.

Service expected from the lower layer. The ATM Layer expects the Physical Layer to support the transparent transport of ATM cells between peer ATM entities. The exchange of information between the ATM Layer and the Physical Layer is implemented by the primitives shown in Table 2.4. The PHY-SDU parameter in these primitives contains the 53-octet cell to be transmitted between peer ATM entities.

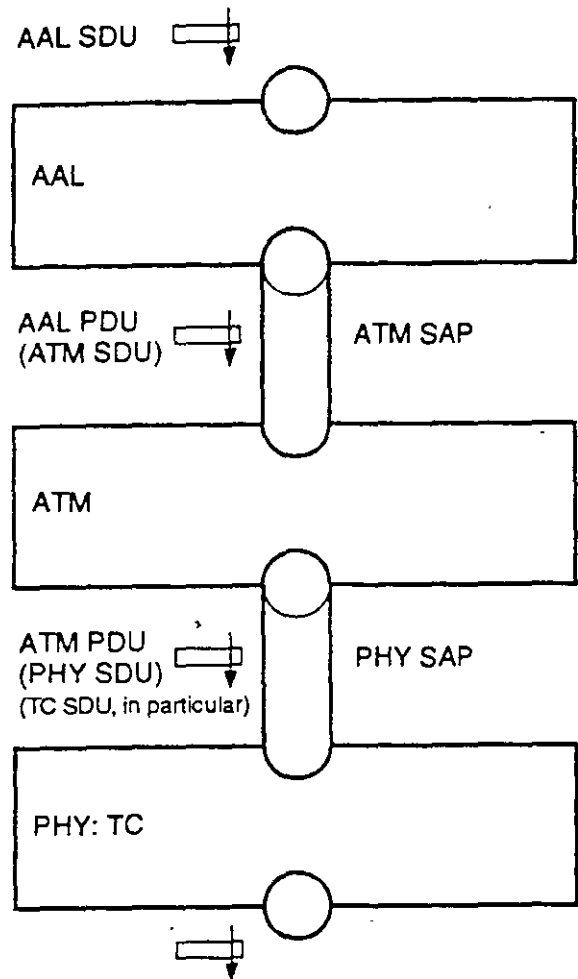


Figure 2.7 Pertinent ATM SAPs.

ATM cell format. The ATM cell format used across the UNI is shown in Fig. 2.8 (which is another way of looking at Fig. 1.2). Table 2.5 describes the meaning of the fields.

- End-to-end operations administration and maintenance capabilities need to be supported. For VPs, operation functions are supported via specially marked ATM cells, which are transmitted over VCs with specific VCI values* (these are known as F4 flows). For VCs, operation functions are supported via cells marked with an appropriate codepoint in the Payload Type Indicator field (these are known as F5 flows). The functions supported are shown in Table 2.6. Figure 2.9 illustrates the difference between these two OAM flows.

Table 2.7 provides the encoding for the PTI field. Code point 100_B (B = binary) indicates a segment OAM F4 cell flow used to monitor the status of a segment within the virtual connection. Code point 101_B indicates an end-to-end OAM F5 cell flow used to monitor the status of a connection end to end. Code point 110_B is reserved for future traffic control and resource management procedures.

*VCI is 4 for end-to-end operations and 3 for segment information.

TABLE 2.3 ATM Layer Primitives

ATM-DATA.request (ATM_SDU, Submitted_Loss_Priority, Congestion_Indication, SDU_Type)	Used to request transmission of an ATM SDU across a VPC or VCC to a peer entity
ATM-DATA.indication (ATM_SDU, Received_Loss_Priority, Congestion_Indication, SDU_Type)	Used by the ATM Layer to indicate to the service user the arrival of an ATM cell

Description of parameters:

ATM_SDU: The 48 octets of information to be transferred by the ATM Layer between peer communicating upper-layer entities.

Submitted_Loss_Priority: The relative importance of the ATM_SDU contained in this primitive. Two values are possible. A value of "high" indicates that the resulting ATM cell has higher (or equivalent) loss priority than a cell with a value of "low." A high value may be translated to a cell loss priority value of 0 in the cell header. Similarly, a low value may be translated to a CLP value of 1 in the cell header.

Congestion_Indication: This parameter indicates whether this cell has passed through one or more network nodes experiencing congestion. It has two values: True or False.

SDU_Type: This parameter indicates the type of SDU to be transferred between peer upper layer entities. It can take only two values, 0 and 1, and its use is as determined by the higher layer. For example, AAL Type 5 sets SDU_Type to 1 to indicate the last cell of a frame. In other words, this field is currently used by the AAL Type 5 Common Part protocol to distinguish between cells that contain the last segment of an AAL Type 5 Common Part PDU and those that do not. AAL Type 1 and AAL Type 3/4 always set the bit to 0.

Received_Loss_Priority: This parameter indicates the CLP field marking of the received ATM_PDU. Two values are possible. A value of "high" indicates that the received ATM cell has higher (or equivalent) loss priority than a cell with a value of "low." A high value may be translated to a cell loss priority value of 0 in the cell header. Similarly, a low value may be translated to a CLP value of 1 in the cell header.

ATM Layer procedures. This section summarizes the functions performed by ATM layer entities.

ATM sending procedures. These procedures are performed by an ATM entity to send ATM cells to a peer ATM entity. The procedures are organized according to the categories of functions performed by the ATM Layer.

ATM layer connections. As described earlier, the ATM service is provided by means of virtual connections. For the PVC cell relay service, connections are established at subscription time. For SVC service,

TABLE 2.4 Physical Layer Primitives

PHY-DATA.request (PHY_SDU)	Requests the Physical Layer to transport an ATM cell between peer ATM entities over an existing connection.
PHY-DATA.indication (PHY_SDU)	Indicates to the ATM Layer that an ATM cell has been received over an existing connection.

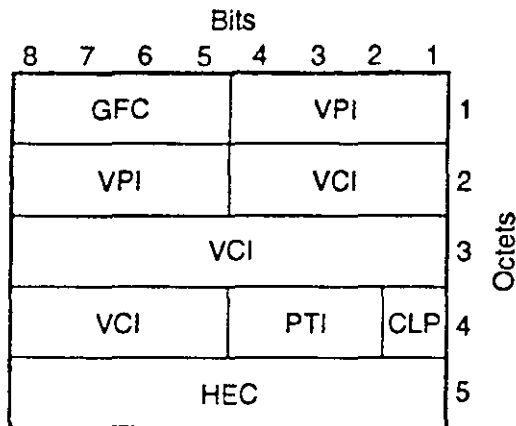


Figure 2.8 ATM cell format.

connections are established by a signaling mechanism. As will be seen in Chap. 4, about one dozen parameters need to be specified to describe a connection (for example, called party, bandwidth, quality of service, etc.).

Cell rate decoupling. A sending ATM entity must add unassigned cells to the assigned cell stream to be transmitted, so that a continuous cell stream matching the line rate of the UNI is provided to the Physical Layer. This is necessary in order for the Physical Layer to perform adequate cell delineation functions. Unassigned cells are empty cells which have the first 4 octets of the cell header encoded as depicted in Fig. 2.10. Unassigned cells do not carry information. Therefore, they must be extracted at the receiving ATM entity and not passed to the upper layer.

Loss priority indication. Traffic management functions may use tagging as a way to control traffic entering the network across the UNI. The network may choose to tag cells that violate a traffic descriptor for the connection by setting the CLP bit to 1. If cell discarding is necessary, these cells would be discarded first. Some traffic management procedures are discussed in Chap. 6.

ATM receiving procedures. This section describes the procedures an ATM entity executes when receiving an ATM cell to ensure its proper processing. These procedures include the provision for sequenced processing of ATM cells which arrive across a virtual connection.

Sequenced ATM processing. ATM cells received across a virtual connection must be processed in sequence to ensure adequate service to the higher layers.

Cell validation procedures. The cell validation procedures determine whether a received cell is an unassigned cell and detect invalid header patterns. These procedures also detect cells received with inactive VPI/VCI values (e.g., VPI/VCI values which identify inactive connections). Unassigned cells and cells found to be in error are discarded.

TABLE 2.5 ATM Cell Fields

Generic Flow Control (GFC)	The 4-bit GFC field has only local significance and may be used to provide standardized local functions at the customer site (e.g., passive bus support); the field is ignored and may be overwritten by the public network.
Virtual Path Identifier/Virtual Channel Identifier	The 24-bit VPI/VCI field indicates the virtual connection over which a cell is to be forwarded. The number of connections needed across the UNI is less than 2^{24} , therefore, only some bits of the VPI and VCI subfields are used. Those bits are called <i>allocated bits</i> , and all other bits in the VPI/VCI field are set to 0. A VPI value of 0 is not available for user-to-user virtual path identification. Similarly, a VCI value of 0 is not available for user-to-user virtual channel identification.
Payload Type Indicator (PTI)	The 3-bit PTI field indicates whether the cell contains user information or layer management information. Code points 000 to 011 indicate user information; these PTI values identify two types of end-user information and whether the cell has experienced congestion (the two types of information are used by the end-user application). For user data, the public network does not change the SDU_Type indicated by the PTI field. The public network can, however, change the PTI value from Congestion_Experienced = False to Congestion_Experienced = True. Code points 100 to 111 identify different types of operations flows. See Table 2.7.
Cell Loss Priority	This 1-bit field allows the user to indicate the relative cell loss priority of the cell. The network may attempt to provide a higher cell loss priority (or equivalent) for cells marked with high priority than for cells marked with low priority. The current view is to only let the user set CLP to the value 0.
Header Error Control	The 8-bit HEC field is used by the Physical Layer to detect transmission errors in the cell header and in some cases for cell delineation.

Cell discrimination based on PTI value. A receiving ATM Layer entity processes cells according to the type of payload they contain as indicated by the value in the PTI field. User cells (PTI values 000–100) are forwarded across the appropriate virtual channel. If neces-

TABLE 2.6 Layer Management Functions Included in Cell Relay Service

Fault management functions	Alarm surveillance: AIS (alarm indication signal) Alarm surveillance: FERF (far-end receive failure; now known as remote defect indicator) Connectivity verification: cell loopback continuity check
Performance management functions	Forward monitoring Backward reporting Monitoring/reporting
Activation/deactivation	Performance monitoring Continuity check

sary, PTI values may be modified to indicate whether the cell experienced congestion.

Layer Management cells (PTI values of 101–111) are used to provide various operations flows to support functions like performance monitoring and trouble sectionalization. CPE supporting the UNI is not required to support these operations flows. However, network equipment must support them so that it can interface with end-user equipment supporting these functions. (This topic is revisited in Chap. 10.)

2.2.3 Layer Management

There are two types of interactions between the ATM entity and the ATM Management entity. One interaction is for the exchange of local information between these two entities. The primitives are shown in Table 2.8 (the parameters are not shown for simplicity). The other interaction is for peer-to-peer communication between ATM Management entities. The primitives for this interaction are shown in Table 2.9. For more details, refer to Ref. 2, 5, or 6. (This topic is revisited in Chap. 10.)




Figure 2.9 OAM F4 and F5 flows.

TABLE 2.7 PTI Code Points


PTI code point	Meaning
000	User data—SDU_Type 0, no congestion experienced
001	User data—SDU_Type 1, no congestion experienced
010	User data—SDU_Type 0, congestion experienced
011	User data—SDU_Type 1, congestion experienced
100	Segment OAM F5 flow cell
101	End-to-end OAM F5 flow cell
110	Reserved for future traffic control and resource management functions
111	Reserved for future use

2.2.4 Physical Layer

Although the emphasis of this chapter is on the ATM Layer, a brief discussion of the underlying Physical Layer is also provided. Figure 2.11 depicts some of the key Physical Layer protocols supported.

As noted, the Physical Layer is made up of two sublayers: the Transmission Convergence Sublayer and the Physical Medium-Dependent Sublayer. The TC Sublayer “maps” the cell stream to the underlying framing mechanism of the physical transmission facility and generates the required protocol control information for the Physical Layer (e.g., SONET overhead octets). It also generates the HEC. The PMD Sublayer deals with the electrical or optical aspects of the physical interface (e.g., timing, power, jitter).

The UNI providing the service’s access interface includes the physical characteristics of facilities that provide actual realizations of the U_B reference point. In practical terms, this access interface specifies the means and characteristics of the connection mechanism between CPE supporting cell relay service and a LEC’s switch providing the same service. UNIs are specified by characteristics such as physical and electromagnetic/optical characteristics, channel structures and access



A: This bit is available for use by appropriate ATM layer function.
X: This bit is a don't care bit.

Figure 2.10 First four octets of cell header for unassigned cells.

TABLE 2.8 ATM Management Primitives for Local Communication

ATMM-MONITOR.indication	Issued by an ATM Layer Management entity to deliver the content of an ATM_PDU received by the ATM entity, to facilitate an OAM function
ATMM-ASSIGN.request	Issued by an ATM Layer Management entity to request the establishment of an ATM link
ATMM-ASSIGN.confirm	Issued by an ATM Layer Management entity to confirm the establishment of an ATM link
ATMM-REMOVE.request	Issued by an ATM Layer Management entity to request the release of an ATM link
ATMM-REMOVE.confirm	Issued by an ATM Layer Management entity to confirm the release of an ATM link
ATMM-ERROR.indication	Issued by an ATM Layer Management entity to indicate an error and invoke appropriate management actions
ATMM-PARAMETER-CHANGE.request	Issued by an ATM Layer Management entity to request a change in a parameter of the ATM link

capabilities, user-network protocols, maintenance and operations characteristics, performance characteristics, and service characteristics.

The physical access channel for ATM-based fastpacket services such as cell relay service supports one of the following access rates: 622.080 Mbits/s (future); 155.520 Mbits/s; 44.736 Mbits/s; 1.544 Mbits/s (perhaps in the future). The corresponding channel signal formats are STS-12c (Synchronous Transport Signal Level 12, concatenated), STS-3c, DS3 (Digital Signal Level 3), and DS1.

Physical-Layer mappings. The mapping of cells onto the DS1, DS3, and SONET STS-3c has also been defined.⁹ Some key aspects of how cells are inserted over the underlying framing mechanism are discussed below.

TABLE 2.9 ATM Management Peer-to-Peer Primitives

ATMM-DATA.request (ATM_SDU, Submitted_Loss_Priority, PHY_CEI(s))	Issued by an ATM Layer Management entity to request transfer of a management ATM_SDU
ATMM-DATA.indication (ATM_SDU, Received_Loss_Priority, PHY_CEI, Congestion_Indication)	Issued to an ATM Layer Management entity to indicate the arrival of a management ATM_SDU

Note: CEI is the connection endpoint identifier.

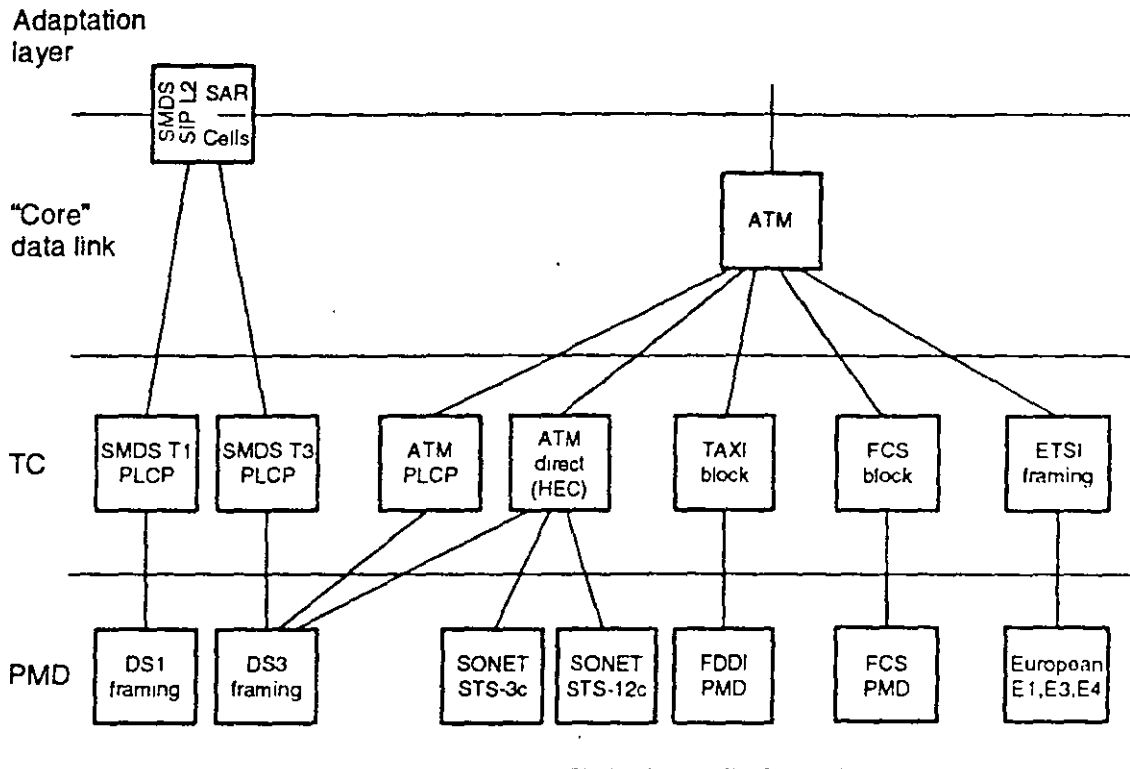


Figure 2.11 Key Physical Layer protocols supported. PLCP = Physical Layer convergence procedure; FCS = fiber channel standard; FDDI = fiber distributed data interface; ETSI = European Telecommunications Standards Institute.

The challenge at the receiving end is to extract the cell from the underlying frame, that is, to establish cell boundaries.

Mapping of ATM cells into 1544-kbit/s DS1 frame. Frame format. The multi-frame structure for the 24-frame multiframe as described in ITU-T Recommendation G.704 is used. The ATM cell is mapped into bits 2 to 193 (i.e., time slots 1 to 24 described in Recommendation G.704) of the 1544-kbit/s* frame, with the octet structure of the cell aligned with the octet structure of the frame (however, the start of the cell can be at any octet in the DS1 payload; (see Fig. 2.12).

Cell rate adaption. The cell rate adaption to the payload capacity of the frames is performed by the insertion of idle cells, as described in ITU-T Recommendation I.432, when valid cells are not available from the ATM Layer.

Header error control generation. The Header Error Control value is generated and inserted in the specific field in compliance with ITU-T Recommendation I.432.

Scrambling of the ATM cell payload (optional). As an option, the ATM cell payload (48 bytes) can be scrambled before it is mapped into the 1544-kbit/s signal. In the reverse operation, following termination

*As of press time, however, standards for the delivery of ATM over a DS1 access were still being investigated.




Figure 2.12 Direct mapping of cells onto DS1 frame (example).

of the 1544-kbit/s signal, the ATM cell payload is descrambled before being passed to the ATM Layer. The self-synchronizing scrambler with the generator polynomial $x^{43} + 1$ is used.

Cell delineation. Cell delineation is performed using the header error control mechanism as defined in ITU-T Recommendation I.432. This direct mapping approach means that the algorithm parses 5 octets on the fly until a 5-octet boundary is found through the HEC procedure. Once the header boundary is found, the rest of the cell boundary is established by counting 48 additional octets.

Cell header verification and extraction. The cell header verification is performed in compliance with ITU-T Recommendation I.432. Only valid cells are passed to the ATM Layer.

Mapping of ATM cells into 44,736-kbit/s DS3 frame

Frame format. The multiframe format at 44,736 kbits/s, as described in ITU-T Recommendation G.704, is used.

Two mappings are available:

1. Physical Layer Convergence Protocol (PLCP)-based mapping of ATM cells, derived from SMDS principles
2. A direct (HEC-based) mapping, established in 1993

This discussion focuses on PLCP, since the direct mapping is similar to the DS1 mapping.

The ATM PLCP defines a mapping of ATM cells onto existing 44,736-kbit/s facilities. The DS3 PLCP consists of a 125- μ s frame within a standard 44,736-kbit/s payload. Note that there is no fixed relationship between the PLCP frame and the 44,736-kbit/s frame; i.e., the PLCP can begin anywhere inside the 44,736-kbit/s payload. The PLCP frame, Fig. 2.13, consists of 12 rows of ATM cells, each preceded by 4 octets of overhead. Nibble stuffing is required after the twelfth cell to fill the 125- μ s PLCP frame. Although the PLCP is not aligned with the 44,736-kbit/s framing bits, the octets in the PLCP frame are nibble-aligned with the 44,736-kbit/s payload envelope. Nibbles begin after the control bits (F, X, P, C, or M) of the 44,736-kbit/s frame. The stuff bits are never used in the 44,736-kbits/s, i.e., the payload is always inserted. The reader interested in a detailed explanation of the DS3 framing format may refer to Ref.10 or other material. Octets in the PLCP frame are described in the following sections.

Cell rate adaption. The cell rate adaption to the payload capacity of the PLCP frame is performed by the insertion of idle cells, as described in ITU-T Recommendation I.432, when no valid cells are available from the ATM Layer.

PLCP (1 octet)	Framing (1 octet)	POI (1 octet)	POH (1 octet)	PLCP payload (53 octets)	
A1	A2	P11	Z6	First ATM cell	
A1	A2	P10	Z5	Second ATM cell	
A1	A2	P09	Z4	Third ATM cell	
A1	A2	P08	Z3		
A1	A2	P07	Z2		
A1	A2	P06	Z1		
A1	A2	P05	X		
A1	A2	P04	B1		
A1	A2	P03	G1		
A1	A2	P02	X		
A1	A2	P01	X	Eleventh ATM cell	(13 or 14 nibbles)
A1	A2	P00	C1	Twelfth ATM cell	Trailer

Figure 2.13 PLCP frame. POI = path overhead indicator; POH = path overhead; BIP-8 = bit interleaved parity-8; X = unassigned (receiver to ignore). [Note: Order and transmission of all PLCP bits and octets are from left to right and top to bottom. This figure shows the most significant bit (MSB) on the left and the least significant bit (LSB) on the right.]

Header error control generation. The HEC generation is based on the algorithm described in ITU-T Recommendation I.432.

Cell delineation. Since the cells are in predetermined locations within the PLCP, framing on the 44,736-kbit/s signal and then on the PLCP is sufficient to delineate cells.

Cell header verification and extraction. The cell header verification is consistent with ITU-T Recommendation I.432. Only valid cells are passed to the ATM Layer.

PLCP overhead utilization. The following PLCP overhead bytes/nibbles are activated across the UNI:

- A1: Frame alignment
- A2: Frame alignment
- B1: PLCP path error monitoring
- C1: Cycle/stuff counter
- G1: PLCP path status
- Px: Path overhead identifier
- Zx: Growth octets
- Trailer nibbles

Frame alignment (A1, A2). The PLCP framing octets use the same framing pattern: A1 = 11110110, A2 = 00101000.

PLCP path error monitoring (B1). The BIP-8 field supports path error monitoring, and is calculated over a 12×54 octet structure

consisting of the POH field and the associated ATM cells (648 octets) of the *previous* PLCP frame.

Cycle/stuff counter (C1). The cycle/stuff counter provides a nibble-stuffing opportunity cycle and length indicator for the PLCP frame. A stuffing opportunity occurs every third frame of a three-frame (375- μ s) stuffing cycle. The value of the C1 code is used as an indication of the phase of the 375 μ s stuffing opportunity cycle, as follows:

C1 code	Frame phase of cycle	Trailer length
11111111	1	13
00000000	2	14
01100110	3 (no stuff)	13
10011001	3 (stuff)	14

Notice that a trailer containing 13 nibbles is used in the first frame of the 375 ms stuffing opportunity cycle. A trailer of 14 nibbles is used in the second frame. The third frame provides a nibble-stuffing opportunity. A trailer containing 14 nibbles is used in the third frame if a stuff occurs. If it does not, the trailer will contain 13 nibbles.

PLCP path status (G1). The PLCP path status is allocated to convey the received PLCP status and performance to the transmitting far end. This octet permits the status of the full receive/transmit PLCP path to be monitored at either end of the path.

Path overhead identifier (P00–P11). The path overhead identifier (POI) indexes the adjacent path overhead (POH) octet of the PLCP.

Growth octets. These are reserved for future use. The receiver ignores the values contained in these fields.

Trailer nibbles. The content of each of the 13 or 14 trailer nibbles is 1100.

Other Mappings. Other mappings have been defined. Direct mappings for E1, DS2, and STS-3c are available.⁴

References

1. ITU-T Recommendation X.210, *Open Systems Interconnection, Layer Service Definition Convention*, Geneva, Switzerland, 1989.
2. ITU-T Recommendation I.361, *B-ISDN ATM Layer Specification*, Geneva, Switzerland, June 1992.
3. T1S1.5/92-410, *Broadband ISDN—ATM Layer Functionality and Specification*, August 1992.
4. ITU-T Recommendation I.432, *B-ISDN User-Network Interface—Physical Layer Specification*, Geneva, Switzerland, June 1992.

5. ATM Forum, *ATM User-Network Interface Specification*, Version 3.0, August, 1993.
6. Bellcore, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements*, TA-NWT-001113, Issue 2, July 1993.
7. ANSI T1.ATM-1993, *Broadband ISDN—ATM Layer Functionality and Specification*, New York.
8. ITU-T Recommendation I.413, *B-ISDN User-Network Interface*, Geneva, Switzerland, 1991.
9. ITU Draft Recommendation G.804, *ATM Cell Mapping into Plesiochronous Digital Hierarchy*, Geneva, Switzerland, February 1993.
10. D. Minoli, *Enterprise Networking—Fractional T1 to SONET, Frame Relay to BISDN*, Artech House, Norwood, Mass., 1993.

ATM Adaptation Layer

3.1 Introduction

As discussed in the previous two chapters, the Protocol Reference Model applicable to both the User Plane and the Control Plane (see Fig. 3.1) is divided into three protocol layers: the *Physical Layer*, the *ATM Layer*, and the *AAL and Service-Specific Layers*.

- The Physical Layer provides the ATM Layer with access to the physical transmission medium. Its functions include transmission of bits across the physical medium, timing recovery, line coding, cell delineation, cell scrambling and descrambling, and generation and checking of the header error control.
- The ATM Layer provides for the transport of ATM cells between the endpoints of a virtual connection. It is the basis for native cell relay service as well as other services. ATM cells are delivered across the network in the same sequence they are received from the CPE.
- The AAL maps the upper-layer data into cells for transport across the network. The Service-Specific Layers perform application-dependent processing and functions.

This chapter focuses on AAL protocols. As noted, the AAL performs the functions necessary to adapt the capabilities provided by the ATM Layer to the needs of higher-layer applications using CRS or other ATM-based services.¹⁻⁴ AALs are typically implemented in end user equipment, as shown, for example, in Fig. 1.16, but can also (occasionally) be found in the network, as seen later. The functions of the AAL include segmentation and reassembly of the higher-layer data units and mapping them into the




Figure 3.1 Protocol reference model.

fixed-length payload of the ATM cells. Effectively, AAL protocols allow a user with some preexisting application, say using TCP/IP, to get the benefits of ATM. To date, three AAL protocol types have been standardized: AAL Type 1 for circuit emulation (or CBR) services, and AAL Type 3/4 and AAL Type 5 for VBR services. A number of service-specific parts have also been standardized. For many years “AAL” meant segmentation/reassembly and error detection only. With the recent inclusion of service-specific functions into the AAL, the functionality has been significantly increased. Two examples of service-specific parts are briefly discussed at the end of this chapter. In AAL Type 1, 1 octet of the cell payload is reserved for control; the remaining 47 octets are utilized for user information. AAL Type 3/4 reserves 4 octets of each cell payload for control use. AAL Type 5 provides all 48 octets of each cell (except for the last cell of a higher-layer packet; see Sec. 3.5.2) for user information.

Note: In this discussion, the term *user* is employed consistent with protocol parlance, unless noted otherwise. Namely, it represents the (protocol) entity just above the AAL Layer; it does not refer to the ultimate user of the (corporate) network. Such a corporate user would access ATM through the top of the protocol stack, e.g., via an application such as E-mail over TCP/IP over ATM.

Recall, for positioning, as we proceed, that AAL provides the balance of capabilities to “fill out” part, but not all, of the Data Link Layer in the OSIRM. Typically the stack {AAL, ATM, PHY} runs just under the Logical Link Control of a traditional LAN, or directly under TCP/IP in an ATM-based LAN or ATM-based WAN.

The novice reader may choose to skip this chapter on first reading; alternatively, the reader may read the first few sections to understand what the AAL aims at doing, without concentrating on how it does it.

3.2 AAL Model

Architecturally, the AAL is a layer between the ATM Layer and the “service layer” (the service layer is shown in Fig. 3.5). The purpose of the ATM Adaptation Layer is to provide the necessary functions to support the service layer that are not provided by the ATM Layer. The functions

provided by the AAL depend upon the service. VBR users may require such functions as PDU delimitation, bit error detection and correction, and cell loss detection. CBR users typically require source clock frequency recovery and detection and possible replacement of lost cells.

Figure 3.2 depicts the positioning of the AAL in the context of the corporate user equipment. AAL capabilities can also be used at an interworking point in the carrier's network, as shown in Fig. 3.3 (this topic is reexamined in Chap. 7). Figure 3.4 shows a classification of services that has been used for specifying ATM Adaptation Layers for different services.

Five AAL protocol types to support the following services are covered in this chapter:

- CBR service using the AAL 1 protocol
- VBR service using the AAL 3/4 Common Part protocol
- VBR service using the AAL 5 Common Part protocol
- Frame relay service (the Frame Relay Service-Specific AAL protocol, which utilizes the AAL 5 Common Part protocol)
- UNI signaling service (the UNI Signaling AAL protocol, which utilizes the AAL 5 Common Part protocol)

The AAL for VBR services consists of two parts: a Common Part (CP) and a Service-Specific Part (SSP). The SSP is used to provide those additional capabilities, beyond those provided by the CP, that are necessary to support the user of the AAL. For some applications the SSP may be "null"; in these cases, the user of the AAL utilizes the AAL Common Part (AALCP) directly. For all AAL types, the AAL receives information from the ATM Layer in the form of 48-octet ATM service data units (ATM_SDU). The AAL passes information to the ATM Layer in the form of a 48-octet ATM_SDU. Figure 3.5 depicts some of the more common protocol arrangements.

Section 3.3 discusses the AAL description for Class 1 (e.g., circuit emulation services), and Sec. 3.4 discusses the AAL description for Class

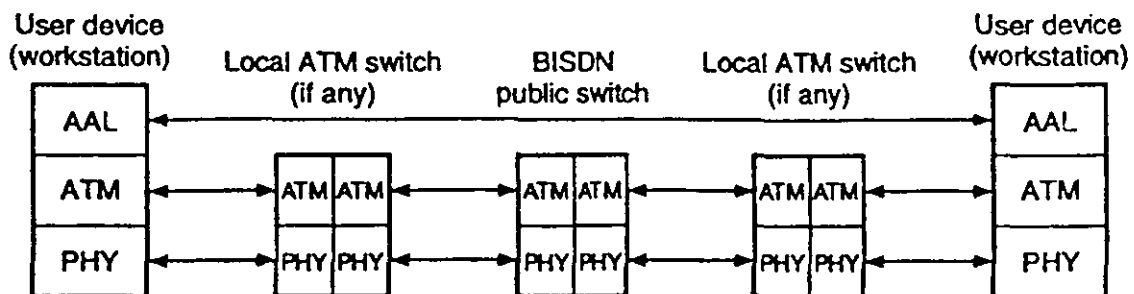


Figure 3.2 The positioning of AAL in CPE.

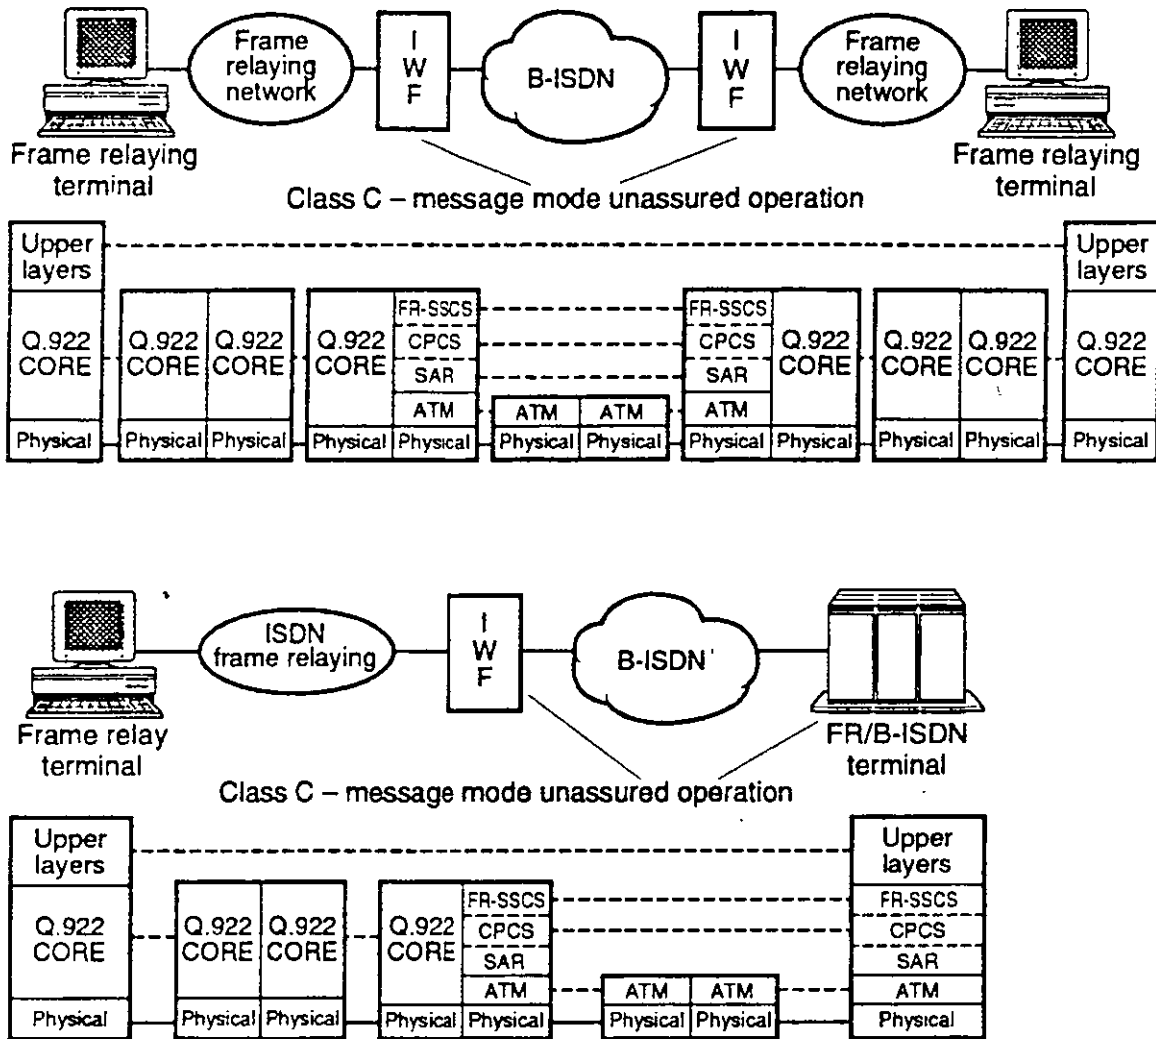


Figure 3.3 Use of AAL protocols at interworking points.

3/4 (e.g., connectionless data services, such as SMDS). Maximum commonality between Class 4 and Class 3 (e.g., connection-oriented data services) AALs has been sought, and people now refer to this AAL as AAL 3/4. The AAL specification for Class 2 services (e.g., variable-bit-rate video services) may occur at a future date. Section 3.5 describes AAL 5, Sec. 3.6 covers the Frame Relay Service-Specific AAL, and Sec. 3.7 briefly covers the signaling AAL.

3.3 AAL Type 1

3.3.1 Overview

One of the services possible with an ATM platform is emulation of a dedicated line (typically at 1.544 or 45 Mbits/s). This type of service is also known as Class A or CBR service. To support CBR services, an adaptation layer is required in the user's equipment for the necessary

Attributes	Class 1	Class 2	Class 3	Class 4
Timing between source and destination	Related		Nonrelated	
Bit rate	Constant	Variable		
Connection mode	Connection-oriented			Connection-less

Figure 3.4 Classification of services for AAL specification. Examples of services: Class 1, circuit emulation; Class 2, variable bit rate video; Class 3, connection-oriented data; Class 4, support of connectionless data transfer; Class X, unrestricted.

functions that cannot be provided by the ATM cell header. Some characteristics and functions that may be needed for an efficient and reliable transport of CBR services are identified below.

Ideally, CBR services carried over an ATM-based network should appear to the corporate user as equivalent to CBR services provided by the circuit switched or dedicated network. Some characteristics of these CBR services are

1. Maintenance of timing information
2. Reliable transmission with negligible reframes
3. Path performance monitoring capability

CBR services with the above characteristics can be provided by assigning the following functions for the CBR Adaptation Layer:

1. Lost cell detection
2. Synchronization
3. Performance monitoring

(These functions may not be required by all the CBR services.)

Therefore, the CBR AAL performs the functions necessary to match the service provided by the ATM Layer to the CBR services required by its service user. It provides for the transfer of AAL_SDUs carrying information of an AAL user supporting constant-bit-rate services. This layer is service-specific, with the main goal of supporting services that

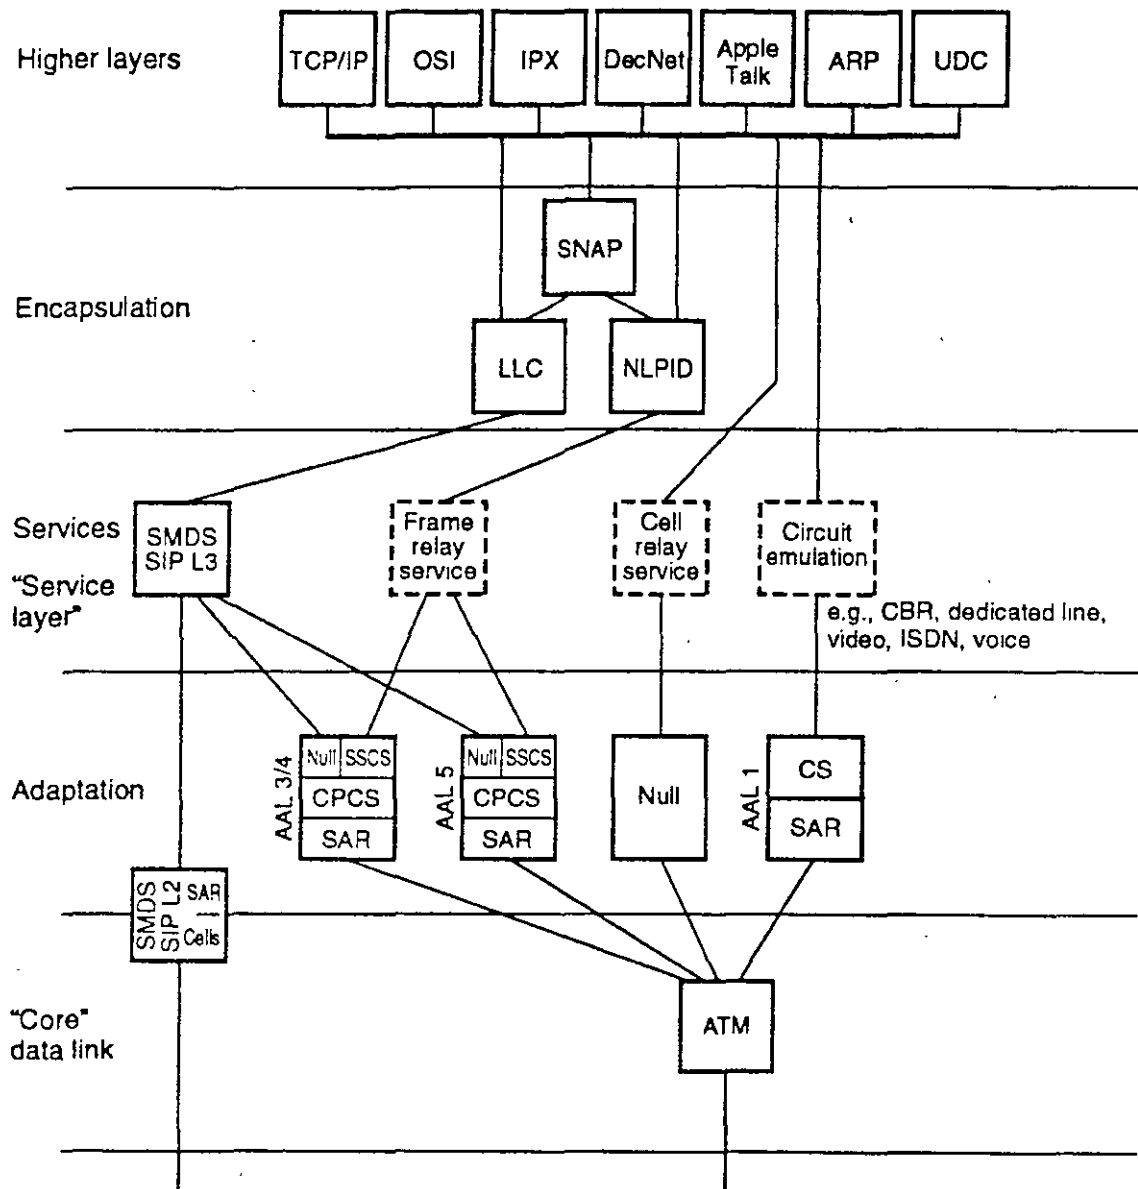


Figure 3.5 Support of user applications. CPCS = common part CS; SSCS = service-specific CS; LLC = logical link control; SNAP = Subnetwork Access Protocol; NLPID = Network Layer Protocol ID.

have specific delay, jitter, and timing requirements, such as circuit emulation. It provides timing recovery, synchronization, and indication of lost information.

The AAL 1 functions are grouped into Segmentation and Reassembly Sublayer functions and Convergence Sublayer functions. The existing agreements in ITU-T Recommendation I.363 and the ANSI CBR AAL Standard³ provide two basic modes of operation for the CBR AAL:

- Unstructured data transfer (UDT)
- Structured data transfer (SDT)

When the UDT mode is operational, the AAL protocol assumes that the incoming data from the AAL user are a *bit stream* with an associated bit clock. When the SDT mode is operational, the AAL protocol assumes that the incoming information is *octet blocks* of a fixed length (such as an $n \times 64$ kbit/s channel with 8-kHz integrity) with an associated clock. While the SDT mode of operation has not been completely specified in the standards, a substantial enough body of agreements exists to assume that by the end of 1994 a complete SDT mechanism will be defined.

3.3.2 CBR AAL services

AAL Type 1 services and functions. The CBR AAL functions are grouped into two sublayers, the SAR Sublayer and the Convergence Sublayer. The SAR is responsible for the transport and bit error detection (and possibly correction) of CS protocol control information. The CS performs a set of service-related functions. It blocks and deblocks AAL_SDUs, counting the blocks, modulo 8, as it generates or receives them. Also, it maintains bit count integrity, generates timing information (if required), recovers timing, generates and recovers data structure information (if required), and detects and generates indications to the AAL management (AALM) entity of error conditions or signal loss. The CS may receive reference clock information from the AALM entity which is responsible for managing the AAL resources and parameters used by the AAL entity. The services provided by AAL Type 1 to the AAL user are

- 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 the source and the destination
- Transfer of structure information between the source and the destination
- Indication of lost or errored information that is not recovered by AAL Type 1, if needed

Specifically, the functions are:

1. Segmentation and reassembly of user information
2. Handling of cell delay variation
3. Handling of cell payload assembly delay
4. Handling of lost and misinserted cells
5. Source clock recovery at the receiver
6. Recovery of the source data structure at the receiver

7. Monitoring of AAL-PCI for bit errors
8. Handling of AAL-PCI bit errors
9. Monitoring of the user information field for bit errors and possible corrective actions

SAR functions. The SAR functions are

- Mapping between the CS_PDU and the SAR_PDU (the SAR Sublayer at the transmitting end accepts a 47-octet block of data from the CS and then prepends a 1-octet SAR_PDU header to each block to form the SAR_PDU).
- Indicating the existence of a CS function (the SAR can indicate the existence of a CS function; the use of the indication mechanism is optional).
- Sequence numbering (for each SAR_PDU payload, the SAR sublayer receives a sequence number value from the CS).
- Error protection (the sequence number and the CSI bits are protected).

A buffer is used to handle cell delay variation. When cells are lost, it may be necessary to insert an appropriate number of dummy SAR_PDUs. Figure 3.6 depicts the AAL Type 1 frame layout.

Convergence Sublayer functions. The functions of the CS are

- Handling of cell delay variation for delivery of AAL_SDUs to the AAL user at a constant bit rate (the CS layer may need a clock derived at the S_B or T_B interface to support this function).
- Processing the sequence count to detect cell loss and misinsertion.
- Providing the mechanism for timing information transfer for AAL users requiring recovery of source clock frequency at the destination end.

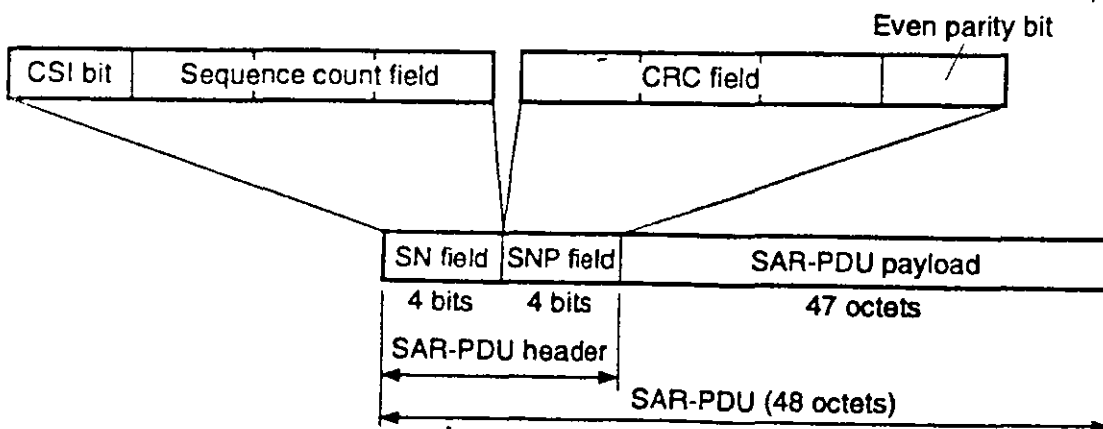


Figure 3.6 AAL Type 1 frame layout. SN = sequence number; SNP = sequence number protection; CSI = Convergence Sublayer indication.

- Providing the transfer of Structure information between source and destination for some AAL users.
- Supporting forward error correction (particularly for video)

For those AAL users that require transfer of structured data [e.g., 8-kHz structured data for circuit-mode bearer services for 64-kbit/s-based ISDN (see Chap. 8)], the Structure parameter is used. This parameter can be used when the user data stream to be transferred to the peer AAL entity is organized into groups of bits. The length of the structured block is fixed for each instance of the AAL service. The length is an integer multiple of 8 bits. An example of the use of this parameter is to support circuit-mode services of the 64-kbit/s-based ISDN. The two values of the Structure parameter are

Start. This value is used when the DATA is the first part of a structured block, which can be composed of consecutive data segments.

Continuation. This value is used when the value Start is not applicable.

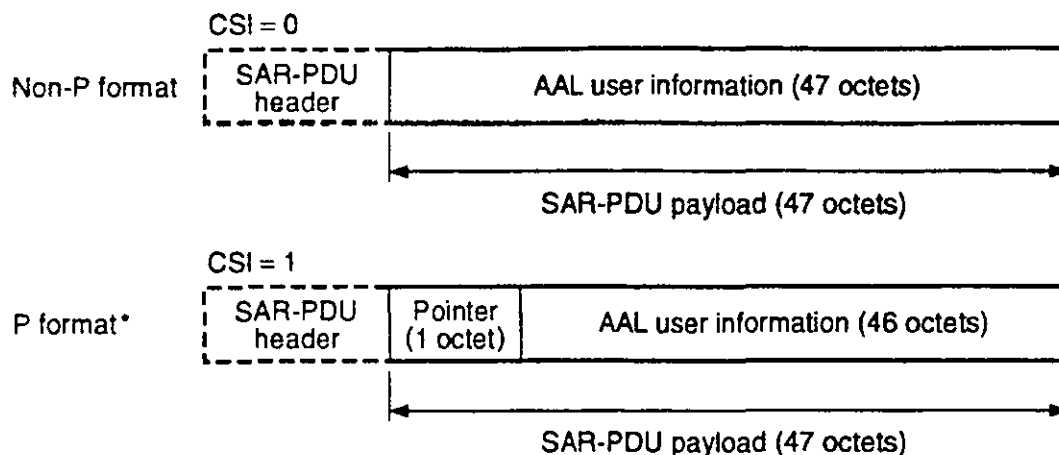
The use of the Structure parameter depends on the type of AAL service provided; its use is agreed upon prior to or at the connection establishment between the AAL user and the AAL.

I.363 notes that “for certain applications such as speech, some SAR functions may not be needed.” For example, I.363 provides the following guidance for CS for voice-band signal transport [which is a specific example of CBR service (see Chap. 8)]:

- *Handling of AAL user information.* The length of the AAL_SDU (i.e., the information provided to the AAL by the upper-layer protocols) is 1 octet (for comparison, the SAR_PDU is 47 octets).
- *Handling of cell delay variation.* A buffer of appropriate size is used to support this function.
- *Handling of lost and misinserted cells.* The detection of lost and inserted cells, if needed, may be provided by processing the sequence count values. The monitoring of the buffer fill level can also provide an indication of lost and misinserted cells. Detected misinserted cells are discarded.

P and non-P formats. The 47-octet SAR_PDU payload used by CS has two formats called non-P and P formats, as seen in Fig. 3.7. These are used to support transfer of information with Structure.

Note that in the non-P format, the entire CS_PDU is filled with user information.



* Used when the SAR-PDU SN = 0, 2, 4, or 6

Figure 3.7 Non-P and P formats.

Partially filled cells. I.363 notes that SAR_PDU payload may be filled only partially with user data in order to reduce the cell payload assembly delay. In this case, the number of leading octets utilized for user information in each SAR_PDU payload is a constant that is determined by the allowable cell payload assembly delay. The remainder of the SAR_PDU payload consists of dummy octets.

Clocking Issues. Besides the UDT/SDT issues discussed earlier, the other basic CBR service attribute that determines the AAL functionality required to support a service is the status of the CBR service clock:⁵

- Synchronous
- Asynchronous

Since the service clock is assumed to be frequency-locked to a network clock in the synchronous case, its recovery is done directly with a clock available from the network. For an asynchronous service clock, the AAL provides a method for recovering the source clock at the receiver. Two methods are available, the synchronous residual time stamp (SRTS) method and the adaptive clock method. The SRTS method is used to recover clocks with tight tolerance and jitter requirements, such as DS1 or DS3 clocks. The adaptive clock recovery method has not been described in enough detail to determine what types of service clocks are supported [presumably less accurate clocks with looser low-frequency jitter (i.e., wander) specifications] or what, if any, added agreements are needed. However, since adaptive clock recovery is common in user equipment, this method is assumed to be available.

The support of DS1 and DS3 CBR service

- Uses the entire 47-octet information payload available with the basic CBR AAL protocol.
- Uses the UDT mode of operation.
- Uses the SRTS method of timing recovery, if the service clock is asynchronous.
- Maintains bit count integrity by inserting the appropriate alarm indication signal for the service supported as a DS1 and DS3 error control measure.

3.3.3 CBR AAL mechanism

The CBR AAL provides its service over preestablished AAL connections. The establishment and initialization of an AAL connection is performed through the AALM. The transfer capacity of each connection and other connection characteristics are negotiated prior to or at connection establishment (the CBR AAL is not directly involved in the negotiation process, which may be performed by management or signaling). The AAL receives from its service user a constant-rate bit stream with a clock. It provides to its service user this constant-rate bit stream with the same clock. The CBR service clock can be either synchronous or asynchronous relative to the network clock. The CBR service is called synchronous if its service clock is frequency-locked to the network clock. Otherwise, the CBR service is called asynchronous.

The service provided by the AAL consists of its own capability plus the capability of the ATM Layer and the Physical Layer. This service is provided to the AAL user (e.g., an entity in an upper layer or in the Management Plane). The service definition is based on a set of service primitives that describe in an abstract manner the logical exchange of information and control. Functions performed by the CBR AAL entities are shown in Table 3.1.

The logical exchange of information between the AAL and the AAL user is represented by two primitives, as shown in Table 3.2.

Service expected from the ATM Layer. The AAL expects the ATM Layer to provide for the transparent and sequential transfer of AAL data units, each of length 48 octets, between communicating AAL entities over an ATM Layer connection, at a negotiated bandwidth and QOS. The ATM Layer transfers the information in the order in which it was delivered to the ATM Layer and provides no retransmission of lost or corrupted information.

TABLE 3.1 Functions Performed by CBR AAL

Detection and reporting of lost SAR_PDUs	Detects discontinuity in the sequence count values of the SAR_PDUs and senses buffer underflow and overflow conditions.
Detection and correction of SAR_PDU header error	Detects bit errors in the SAR_PDU header and possibly corrects a 1-bit error.
Bit count integrity	Generates dummy information units to replace lost AAL_SDUs to be passed to the AAL user in an AAL-DATA.indication.
Residual time stamp (RTS) generation	Encodes source service clock timing information for transport to the receiving AAL entity.*
Source clock recovery	Recovers the CBR service source clock.
Blocking	Maps AAL_SDUs into the payload of a CS_PDU.
Deblocking	Reconstructs the AAL_SDU from the received SAR_PDUs and generates the AAL-DATA.indication primitive.
Structure pointer generation and extraction	Encodes in a 1-octet structure pointer field at the sending AAL entity the information about periodic octet-based block structures present in AAL-DATA.request primitives. The receiving AAL entity extracts the structure pointer received in the CS_PDU header field to verify locally generated block structure.

*Refer to Ref. 3 for a description of the time stamp mechanism.

Interactions between the SAR and the Convergence Sublayer. The logical exchange of information between the SAR and the Convergence Sublayer is represented by the primitives of Table 3.3.

Interacting with the Management Plane. The AALM entities in the Management Plane perform the management functions specific to the AAL. Also, the AALM entities, in conjunction with the Plane Management, provide coordination of the local interactions between the User Plane and the Control Plane across the layers.

The AAL entities provide the AALM entities with the information required for error processing or abnormal condition handling, such as indication of lost or misdelivered SAR_PDUs and indication of errored SAR_PDU headers.

TABLE 3.2 Primitives for CBR AAL

AAL-DATA.request (AAL_SDU, Structure)	This primitive is issued by an AAL user entity to request the transfer of an AAL_SDU to its peer entity over an existing AAL connection. The time interval between two consecutive AAL-DATA.request primitives is constant and a function of the specific AAL service provided to the AAL user.
AAL-DATA.indication (AAL_SDU, Structure, Status)	This primitive is issued to an AAL user entity to notify the arrival of an AAL_SDU over an existing AAL connection. In the absence of error, the AAL_SDU is the same as the AAL_SDU sent by the peer AAL user entity in the corresponding AAL-DATA.request. The time interval between two consecutive AAL-DATA.indication primitives is constant and a function of the specific AAL service provided to the AAL user.

Description of parameters:

AAL_SDU: This parameter contains 1 bit of AAL user data to be transferred by the AAL between two communication AAL user peer entities.

Structure: This parameter is used to indicate the beginning or continuation of a block of AAL_SDUs when providing for the transfer of a structured bit stream between communicating AAL user peer entities (structured data transfer service). The length of the blocks is constant for each instance of the AAL service and is a multiple of 8 bits. This parameter takes one of the following two values: Start and Continuation. It is set to Start whenever the AAL_SDU being passed in the same primitive is the first bit of a block of a structured bit stream. Otherwise, it is set to Continuation. This parameter is used only when SDT service is supported.

Status: This parameter indicates whether the AAL_SDU being passed in the same indication primitive is judged to be nonerrored or errored. It takes one of the following two values: Valid or Invalid. The Invalid value may also indicate that the AAL_SDU being passed is a dummy value. The use of this parameter and the choice of the dummy value depend on the specific service provided.

TABLE 3.3 SAR Primitives

SAR-DATA.invoke (CSDATA, SCVAL, CSIVAL)	This primitive is issued by the sending CS entity to the sending SAR entity to request the transfer of a CSDATA to its peer entity.
SAR-DATA.signal (CSDATA, SNCK, SCVAL, CSIVAL)	This primitive is issued by the receiving SAR entity to the receiving CS entity to notify it of the arrival of a CSDATA from its peer CS entity.

Description of parameters:

CSDATA: This parameter represents the interface data unit exchanged between the SAR entity and the CS entity. It contains the 47-octet CS_PDU.

SCVAL: This 3-bit parameter contains the value of the sequence count associated with the CS_PDU contained in the CSDATA parameter.

CSIVAL: This 1-bit parameter contains the value of the CSI bit.

SNCK: This parameter is generated by the receiving SAR entity. It represents the results of the sequence number protection error check over the SAR_PDU header. It can assume the values of SN-Valid and SN-Invalid.

3.4 ATM Adaptation Layer Functions for VBR (or Bursty Data) Services

As seen in Fig. 3.5, AAL functions for VBR services such as SMDS and frame relay consist of a set of core functions and a set of optional functions. This AAL is now commonly referred to as AAL Type 3/4. As an example, SMDS over ATM uses AAL Type 3/4. The purpose of the ATM Adaptation Layer Type 4/3 Common Part (CPAAL3/4) protocol is to support the upper-layer data transfer needs while using the service of the ATM Layer. This protocol provides for the transport of variable-length frames (up to 65,535 octets in length) with error detection. The CPAAL3/4 provides service over preestablished connections. Termination of a CPAAL3/4 connection also coincides with termination of an ATM Layer service. The establishment and initialization of a CPAAL3/4 connection is performed by interaction with CPAAL3/4 Layer Management entities. There is a dual view of the AAL3/4 Layer.

1. View in terms of Service-Specific Parts and Common Part, as shown in the left-hand side of Fig. 3.8. Core functions are required by all bursty data applications; these functions are known as CP. Optional SSPs are selected as needed. For some applications the SSP is null, implying that the user of the AAL3/4 Layer utilizes the Common Part directly.

2. View in terms of a combination of SAR, the Common Part of the Convergence Sublayer, and SSP, as shown in the right-hand side of Fig. 3.8. SAR and the Common Part of the Convergence Sublayer taken together make up the CP; the Common Part of CS and SSP together form the CS. In other words, the Convergence Sublayer has been

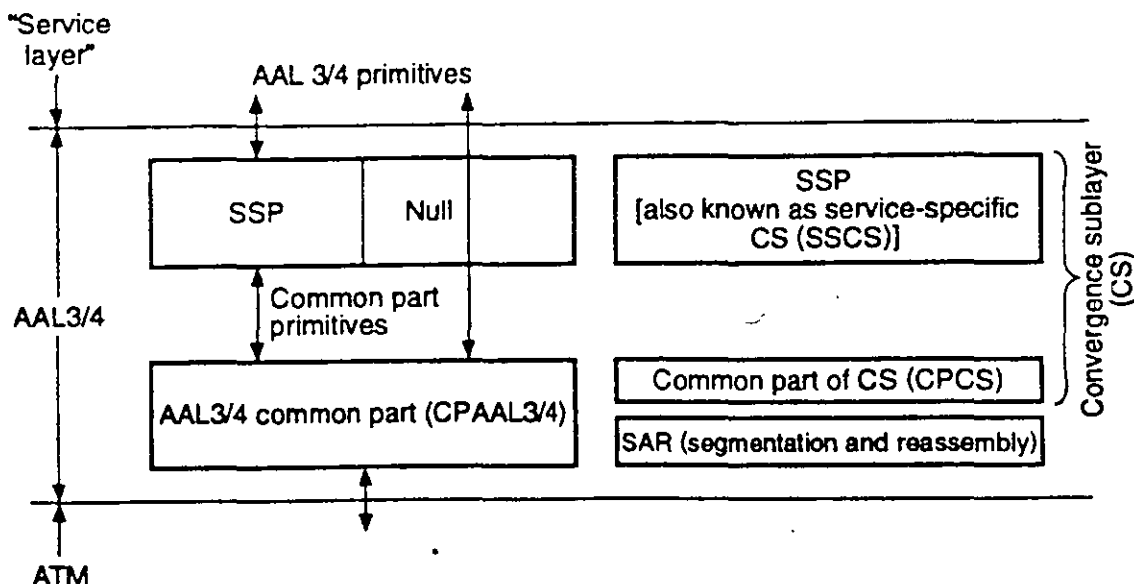


Figure 3.8 Model of AAL3/4. Left: CP/SSP view; right: CS/SAR view.

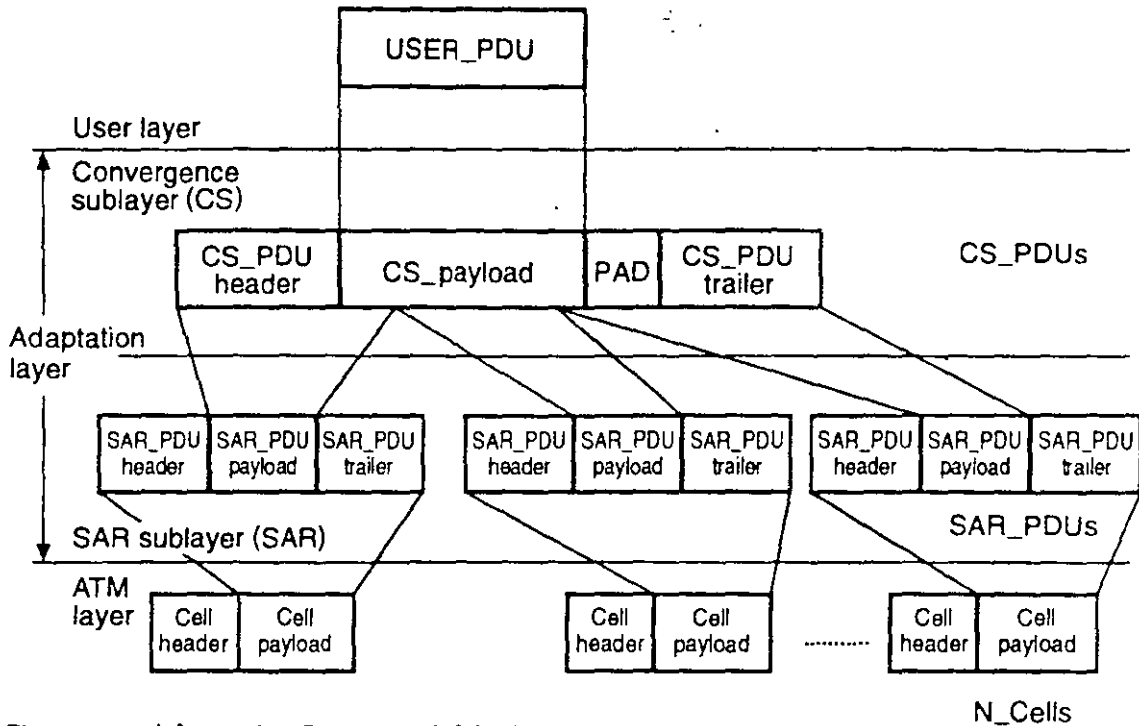


Figure 3.9 Adaptation Layer model for bursty data services.

subdivided into the Common Part CS (CPCS) and the Service-Specific CS (SSCS). In this view, functions are provided by the operation of two logical sublayers, the CS and the SAR. Figure 3.9 shows the operation of AAL3/4 in terms of the PDUs.

The SAR Sublayer is common to all VBR services using AAL3/4, whereas the Convergence Sublayer provides additional, service-specific functions (note that some VBR services may use AAL5). The functions of the Common Part are clearly common by definition. In addition to this, achieving the maximum commonality in the Convergence Sublayer protocol for bursty data services has also been an objective, as implied in Fig. 3.5. For these services, the user presents a variable-size PDU for transmission across the ATM network. The transmission is accomplished by using fixed-length cells to transport data in ATM, as discussed in Chap. 2. At the receiving end of the ATM connection, the user layer receives the PDU that has been reassembled by the SAR and CS protocols.

The discussion that follows looks at AAL3/4 first from a CP point of view (the left-hand model in Fig. 3.8), then from the SAR point of view (the right-hand side of Fig. 3.8). As noted, the functions of the CPAAL3/4 in this view have been grouped into two sublayers: CPAAL3/4 Segmentation and Reassembly (CPAAL3/4_SAR) and CPAAL3/4 Convergence Sublayer (CPAAL3/4_CS). The CPAAL3/4_SAR deals principally with the segmentation and reassembly of data units so that they can be

mapped into fixed-length payloads of the ATM cells, while the CPAAL3/4_CS deals mainly with checking missassembled CPAAL3/4_CS_PDUs.

CPAAL3/4 Layer Management is responsible for the following capabilities: assignment of the CPAAL3/4 association necessary for the establishment of CPAAL3/4 connections between peer CPAAL3/4 entities, resetting the parameters and state variables associated with a CPAAL3/4 connection between peer CPAAL3/4 entities, and monitoring performance for the quality of the ATM connection service provided through notification of errors.

3.4.1 Services provided to the upper layer

The CPAAL3/4 provides, on behalf of its user, for the sequential and transparent transfer of variable-length, octet-aligned CPAAL3/4_SDUs from one corresponding CPAAL3/4 peer to one or more CPAAL3/4 peers. The service is unassured: CPAAL3/4_SDUs may be lost or corrupted. Lost or corrupted CPAAL3/4_SDUs are not recovered by the CPAAL3/4. As an option, corrupted CPAAL3/4_SDUs may be delivered to the remote peer with an indication of the error (this option is known as corrupted data delivery option).

Specifically, the functions performed by the CPAAL3/4 are⁶

- Data transfer between CPAAL3/4 peers
- Preservation of CPAAL3/4_SDUs (delineation and transparency of CPAAL3/4_SDUs)
- CPAAL3/4_SDU segmentation
- CPAAL3/4_SDU reassembly
- Error detection and handling (detects and handles bit errors, lost or gained information, and incorrectly assembled CPAAL3/4_SDUs)
- Multiplexing and demultiplexing (optional multiplexing of multiple CPAAL3/4 connections or interleaving of CPAAL3/4_CS_PDUs)
- Abort (termination of task in case of partially transmitted/received CPAAL3/4_SDUs)
- Pipelining (forwarding PDUs before the entire PDU is received)

This layer provides its user two services:

1. Message-mode service: In this service mode, the CPAAL3/4_SDU passed across the CPAAL3/4 interface is exactly equal to one CPAAL3/4 interface data unit (CPAAL3/4_IDU), as seen in Fig. 3.10.

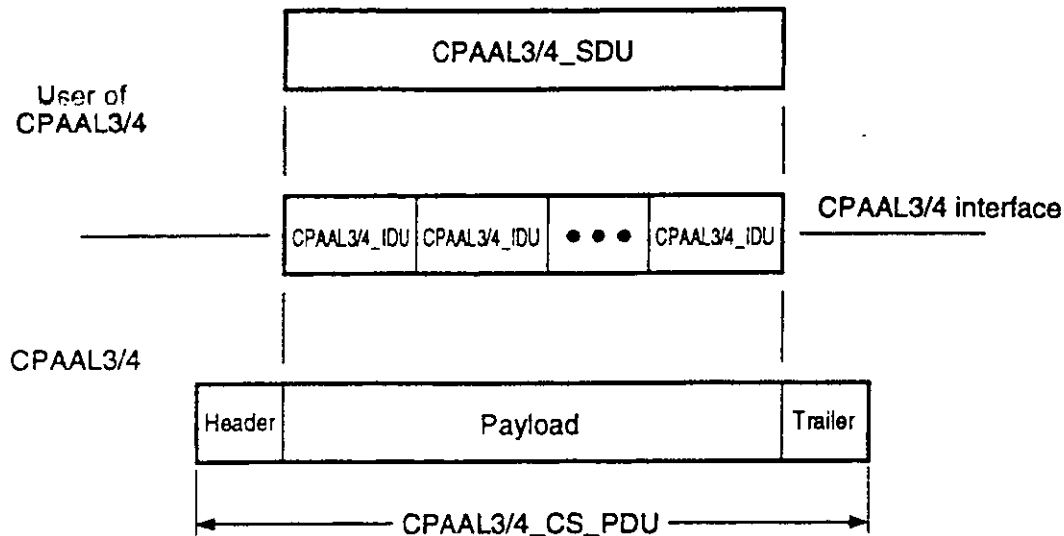


Figure 3.11 Streaming-mode service.

the .signal to an .indication. Table 3.4 provides additional information on these primitives.

Services from the ATM Layer. The CPAAL3/4 expects the ATM Layer (discussed in Chap. 2) to provide for the transparent and sequential transport of 48-octet CPAAL3/4 data units (that is, CPAAL3/4_SAR_PDUs) between communicating CPAAL3/4 peers over preestablished connections at a negotiated QOS. The information is transferred to the ATM Layer in the order in which it is to be sent, with no retransmission of lost or corrupted information.

Interaction with CPAAL3/4 Management entities. Management information is exchanged using five management primitives. See Ref. 4 for details.

3.4.2 SAR Sublayer functions

There is a single SAR function for all bursty data services. Hence, the SAR control fields that appear in each cell payload must be the same, regardless of the service and whether or not the fields are used by a particular application. A single SAR for these services leads to lower overall costs for equipment providers and network providers, and hence for end users (e.g., diagnostic generation, testing, and maintenance are simpler when only a single SAR function is used for all services).

The SAR control fields include the following:⁶

Segment_Type field to identify the cell payload as being beginning of message (BOM), continuation of message (COM), end of message (EOM), or only a single-segment message (SSM).

TABLE 3.4 CPAAL3/4 Primitives

CPAAL3/4-UNITDATA.invoke	Issued by a CPAAL3/4 entity to request the transfer of a CPAAL3/4_IDU over an existing CPAAL3/4 connection. This IDU is not subject to any flow control and is always transmitted. The transfer of the IDU is subject to the service mode being used (message versus streaming).
CPAAL3/4-UNITDATA.signal	Issued to a CPAAL3/4 entity to indicate the arrival of a CPAAL3/4_IDU over an existing CPAAL3/4 connection.
CPAAL3/4-U-ABORT.invoke	Issued by a CPAAL3/4 entity using streaming-mode service to request the termination of a CPAAL3/4_SDU that has been partially transferred. The issue of this primitive also causes the generation of an abort message by the CPAAL3/4 to its peer entity if the transmission of the message has already started. (This primitive is not used in message mode.)
CPAAL3/4-U-ABORT.signal	Issued by a CPAAL3/4 entity using streaming-mode service to indicate the termination of a partially delivered CPAAL3/4_SDU by instruction from its peer entity. (This primitive is not used in message mode.)
CPAAL3/4-P-ABORT.signal	Issued by a CPAAL3/4 entity using streaming-mode service to indicate to its user that a partially delivered CPAAL3/4_SDU is to be discarded because of the occurrence of some error; it has local significance. (This primitive is not used in message mode.)

Description of parameters:

ID (Interface data): This parameter contains the interface data unit (CPAAL3/4_IDU) exchanged between CPAAL3/4 entities [it may be the entire CPAAL3/4_SDU (message mode) or segments (streaming mode)].

M (more): Used only in streaming mode to indicate whether the CPAAL3/4_IDU communicated in the ID parameter contains the ending segment of the CPAAL3/4_PDU (=0) or does not (=1).

ML (maximum length): Used only in streaming mode to indicate the maximum length of the CPAAL3/4_SDU; it has values from 0 to 65,535.

RS (reception status): Indicates that the CPAAL3/4_IDU delivered may be corrupted.

LP (loss priority): Indicates the loss priority assigned to the CPAAL3/4_SDU. Two levels of priority are supported, but how to map this parameter to and from the ATM_Submitted_Loss_Priority (discussed in Chap. 2) has not yet been worked out.

CI (congestion indication): Indicates the detection of congestion experienced by the received CPAAL3/4_SDU.

Sequence_Number field to improve the reassembly error detection process.

Message_ID (M_ID) field, which, for connectionless services, allows for the collection of the cell payloads that make up a CS PDU.

Cell Fill field that allows the identification of the fill within a cell payload. It can be used to locate the last octet in the end of message cell. The last octet in the EOM cell could also be identified from the length field associated with the PDU; additionally, data pipelining could be provided by a series of partially filled single-segment message cells. However, in the latter case, significant additional processing is required to reconstruct the original data unit compared with the case where partial fills are indicated by a cell-associated length field.

Error Control field which provides error detection capabilities across the adaptation header and the information payload. The error check is made across all 48 octets irrespective of whether the cell is fully or partially filled.

On transmission, the process is used by the sending CPAAL3/4 entity. The SAR Sublayer accepts variable-length CPAAL3/4_CS_PDUs from the Convergence Sublayer and maps each CPAAL3/4_CS_PDU into a sequence of CPAAL3/4_SAR_PDUs, by placing at most 44 octets of the CPAAL3/4_CS_PDU into a CPAAL3/4_SAR_PDU payload, along with additional control information, described below, used to verify the integrity of the CPAAL3/4_SAR_PDU payload on reception and to control the reassembly process. The sending CPAAL3/4 entity transfers the CPAAL3/4_SAR_PDUs to the ATM Layer for delivery across the network.

On reception, CPAAL3/4_SAR_PDUs are validated, and the user data in the CPAAL3/4_SAR_PDU (note that a CPAAL3/4_SAR_PDU can be partially filled) are passed to the Convergence Sublayer.

3.4.3 Convergence Sublayer functions

On transmission, the Convergence Sublayer accepts variable-length user protocol data units (USER_PDUs) from the service layer. The Convergence Sublayer prepends a 32-bit header to the USER_PDU, then appends from 0 to 3 pad octets to the USER_PDU to build it out to an integral multiple of 32 bits. Next, it appends a 32-bit trailer to the concatenated header, USER_PDU, and pad structure. This collection (the header, USER_PDU, pad, and trailer) is referred to as a CPAAL3/4_CS_PDU. The header and trailer fields are used to detect loss of data and to perform additional functions as required by the service user. After appending the trailer, the Convergence Sublayer passes the CPAAL3/4_CS_PDU to the SAR Sublayer for segmentation and then transmission.⁶

On reception, the Convergence Sublayer validates the collection of CPAAL3/4_SAR_PDU payloads received from the SAR Sublayer by using the information contained in the Convergence Sublayer header and trailer. It removes the pad octets, if any, and presents the validated CPAAL3/4_CS_PDU payload to the user (i.e., the service layer).

3.4.4 SAR Sublayer fields and format

The SAR Sublayer functions are implemented using a 2-octet adaptation header and a 2-octet adaptation trailer. The header and trailer, together with 44 octets of user information, make up the payload of the ATM cell. The sizes and positions of the fields are given in Fig. 3.12. The use of the error control field for error detection is mandatory. The 10-bit CRC has the capability of single-bit error correction over the 48 octets. If the underlying transmission system produces single-bit errors, error correction may be applied at the receiver.

Figure 3.12 shows the CPAAL3/4_SAR_PDU components of the Adaptation Layer, which include a SAR_PDU_Header and an SAR_PDU_Trailer. These two fields encapsulate the SAR_PDU_Payload, which contains a portion of the CPAAL3/4_CS_PDU.

The SAR_PDU_Header is subdivided into three fields: a Segment_Type field, a Sequence_Number field, and a Message Identification (MID) field. The SAR_PDU_Trailer is subdivided into two fields: a Payload_Length field and a Payload CRC field. Details of the purpose and encoding of each subfield follow.⁶

Segment_Type subfield. The 2-bit Segment_Type subfield is used to indicate whether a CPAAL3/4_SAR_PDU is a BOM, COM, EOM, or SSM. Table 3.5 shows the encodings for the Segment_Type subfield.

Sequence_Number subfield. Four-bits are allocated to the SAR_PDU Sequence_Number (SAR_SN) subfield, allowing the streams of

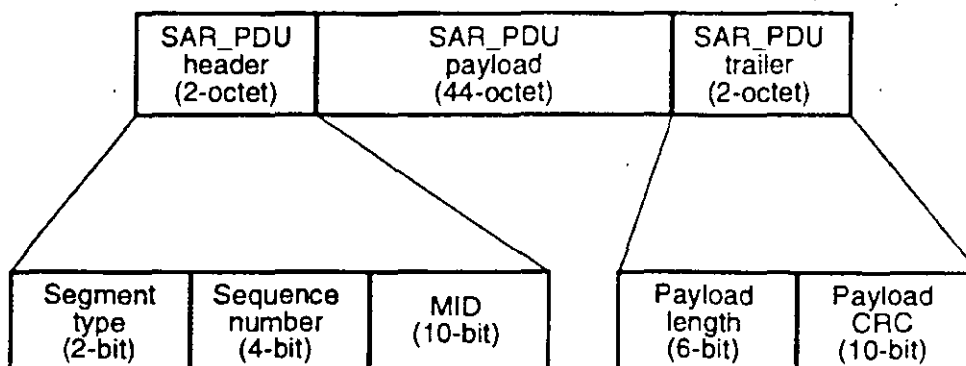


Figure 3.12 CPAAL3/4_SAR_PDU Sublayer format of AAL. MID = message identifier, or multiplexing identifier.

TABLE 3.5 Encoding of the Segment_Type Subfield

Segment_Type	Encoding
BOM	10
COM	00
EOM	01
SSM	11

CPAAL3/4_SAR_PDUs and CPAAL3/4_CS_PDUs to be numbered modulo 16. The SAR_SN is set to all 0s for the first CPAAL3/4_SAR_PDU associated with a given CPAAL3/4_CS_PDU (i.e., the BOM). For each succeeding CPAAL3/4_SAR_PDU of that CPAAL3/4_CS_PDU, the SAR_SN is incremented by 1 relative to the SAR_SN of the previous CPAAL3/4_SAR_PDU of the CPAAL3/4_CS_PDU. When reassembling a CPAAL3/4_CS_PDU, a state variable is maintained that indicates the value of the next expected SAR_SN for the CPAAL3/4_CS_PDU. If the value of the received SAR_SN differs from the expected value, the CPAAL3/4_SAR_PDU is dropped, the partially reassembled errored CPAAL3/4_CS_PDU is discarded, and any following CPAAL3/4_SAR_PDUs associated with this corrupted CPAAL3/4_CS_PDU are dropped.

The use of this function allows the detection of most consecutive losses of COM cells as soon as the following COM or EOM cell of the CPAAL3/4_CS_PDU is received. If the number of COMs of a given CPAAL3/4_CS_PDU that is lost is an integer multiple of 16, the SAR_SN cannot detect them. Therefore, the use of the length field at the CS Sublayer is still required to detect any modulo 16 consecutive losses of CPAAL3/4_SAR_PDUs that may occur during situations like network congestion or protection switching events.

In addition, the use of this function will allow for immediate detection of most cases of cell insertion.

The use of Sequence_Number to detect situations in which two CPAAL3/4_CS_PDUs are inadvertently merged into one and the resulting length matches the length field in the CPAAL3/4_CS_PDU trailer is weak. This is due to the fact that this error event requires that the lengths of the original CPAAL3/4_CS_PDUs be the same. This implies that the same number of CPAAL3/4_SAR_PDUs will probably be required to transport two CPAAL3/4_CS_PDUs. Therefore, the SAR_SNs of the received CPAAL3/4_SAR_PDUs will probably be consecutive, and so the SAR Sublayer will not detect this error event. As a result, the use of the Etag at the CS Sublayer is still required.

Message Identification (MID) subfield. The 10-bit MID subfield is used to reassemble CPAAL3/4_SAR_PDUs into CPAAL3/4_CS_PDUs. All CPAAL3/4_SAR_PDUs of a given CPAAL3/4_CS_PDU will have the same MID. Note that this provides the basis for reassembly of discrete connectionless packets. Use of this subfield as the basis for a multiplexing or reassembly capability for connection-oriented services is for further study.

Payload_Length subfield. The 6-bit Payload_Length subfield is coded with the number of octets from the CPAAL3/4_CS_PDU that are included in the current CPAAL3/4_SAR_PDU. This number has a value between 0 and 44 inclusive. This subfield is binary coded with the most significant bit left-justified. BOM and COM cells take the value 44; EOM cells take the values 4, 8, ..., 44; SSM cells take the values 8, 12, ..., 44.

SAR_PDU_Payload. The CPAAL3/4_CS_PDU is left-justified in the SAR_PDU_Payload of the CPAAL3/4_SAR_PDU. Any part of the SAR_PDU_Payload that is not filled with CS information shall be coded as zeros.

Payload_CRC subfield. The 10-bit Payload_CRC subfield is filled with the value of a CRC calculation that is performed over the entire contents of the CPAAL3/4_SAR_PDU payload, including the SAR_PDU_Header, the SAR_PDU_Payload, and the SAR_PDU_Trailer. The CRC-10 generating polynomial has the capability of single-bit error correction over the CPAAL3/4_SAR_PDU. The following generator polynomial is used to calculate the Payload_CRC:

$$G(x) = x^{10} + x^3 + 1$$

The CRC remainder is placed in the CRC subfield with the most significant bit left-justified in the CRC subfield.

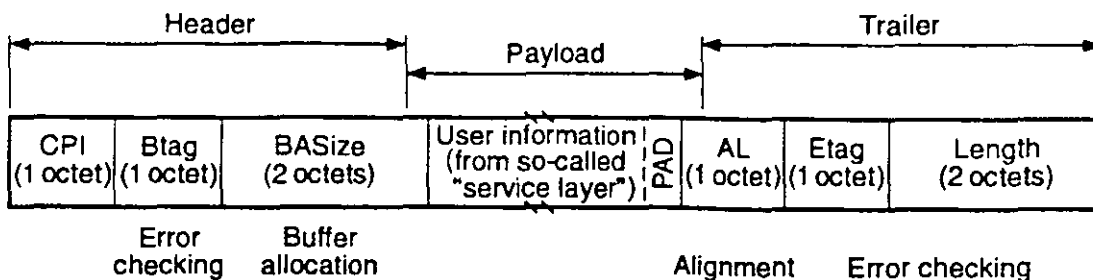


Figure 3.13 CPAAL3/4_CS_PDU Sublayer format of AAL.

3.4.5 Convergence Sublayer fields and format

Figure 3.13 depicts the Convergence Sublayer format of the AAL3/4.

There are two Adaptation Layer control fields: the CS_PDU_Header and the CS_PDU_Trailer, both of which are 4 octets long. The CS_PDU_Header and CS_PDU_Trailer encapsulate the user's protocol data units (USER_PDU). In addition, there may be from 0 to 3 pad octets added to align the CPAAL3/4_CS_PDU with a 32-bit boundary.

The CS_PDU_Header is subdivided into three fields: an 8-bit Common Part Indicator field, an 8-bit Beginning Tag (Btag) field, and a 16-bit Buffer Allocation size (BAsize) field. Likewise, the CS_PDU_Trailer is also subdivided into three fields: an 8-bit filler field, an 8-bit End Tag (Etag) field, and a 16-bit Length field.⁶

Common Part Indicator subfield. The 8-bit Common Part Indicator (CPI) subfield is used to identify the message type, i.e., to interpret subsequent fields for the CPAAL3/4-CS functions in the CPAAL3/4_CS_PDU header and trailer. It also indicates the counting unit for the values specified in the BAsize and Length fields.

CS_PDU Header—Btag subfield. For a given CPAAL3/4_CS_PDU, the same value appears in the 8-bit Btag field of the CS_PDU_Header and in the Etag field in the CS_PDU_Trailer. This allows the identification of a BOM segment and an EOM segment, and hence all intervening COM segments, as belonging to the same CPAAL3/4_CS_PDU. This correlation is required to implement segment loss detection over a CPAAL3/4_CS_PDU. As each CPAAL3/4_CS_PDU is transmitted, the Etag value is changed so that the entire range of Etag field values (0 to 255) is cycled through before reuse to aid in this segment loss protection.

BAsize subfield. The 16-bit Buffer Allocation size (BAsize) subfield is used to predict the buffer requirements for the CPAAL3/4_CS_PDU. Therefore, it must be greater than or equal to the true CPAAL3/4_CS_PDU length. This field is binary coded with the most significant bit left-justified in the subfield. If message-mode service is being provided, the BAsize value is encoded to be equal to the length of the USER_PDU field contained in the CPAAL3/4_CS_PDU Payload field. If streaming-mode service is being provided, the BAsize value is encoded to be equal to the maximum length of the CPAAL3/4_SDU.

USER_PDU field. The variable-length USER_PDU field contains user information. It contains the CPAAL3/4_SDU. It is octet aligned, as it is

limited in length to the value of the BAsize field multiplied by the value of the counting unit (as identified in the CPI field).

Pad Field. The Pad field consists of 0, 1, 2, or 3 octets set to zero, so that the CPAAL3/4_CS_PDU is padded out to a 32-bit boundary.

AL. This 8-bit subfield is used to achieve 32-bit alignment in the CPAAL3/4_CS_PDU trailer. This is strictly a filler octet and does not contain any additional information.

Etag subfield. The 8-bit Etag subfield in the CPAAL3/4_CS_PDU trailer has the same value as the Btag subfield in the corresponding CPAAL3/4_CS_PDU header. As was mentioned earlier, the Btag and Etag subfields in the CS_PDU_Header and CS_PDU_Trailer are correlated in order to detect segment loss and misassembly. This field is binary coded with the most significant bit left-justified.

Length subfield. The 16-bit Length subfield specifies the length, in octets, of the USER_PDU (that is, the length of the user information contained in the CPAAL3/4_CS_PDU Payload field). This field is binary coded with the most significant bit left-justified in the subfield. It is used in conjunction with the Btag and Etag fields for the purpose of detecting misassembled CPAAL3/4_CS_PDUs.

3.5 AAL Type 5

The goal of the AAL Type 5 is to support, in the most streamlined fashion, those capabilities that are required to meet upper-layer data transfer over an ATM platform. The AAL Type 5 Common Part (CPAAL5) protocol provides for the transport of variable-length frames (1 to 65,535 octets) with error detection (the frame is padded to align the resulting PDU with an integral number of ATM cells). A length field is used to extract the frame and detect additional errors not detected with the CRC-32 mechanism. ANSI had a Letter Ballot for AAL Type 5 Common Part at press time, and ITU-TS had a draft version of I.363 (Section 6); approval was expected.

The Convergence Sublayer has been subdivided into the Common Part CS (CPCS) and the Service-Specific CS (SSCS), as shown in Fig. 3.14. Different SSCS protocols, to support specific AAL user services or groups of services, may be defined. The SSCS may also be null, in the sense that it provides only for the mapping of the equivalent primitives of the AAL to CPCS and vice versa. SSCS protocols are specified in separate Recommendations, not in, say, ITU-T I.363. This discussion

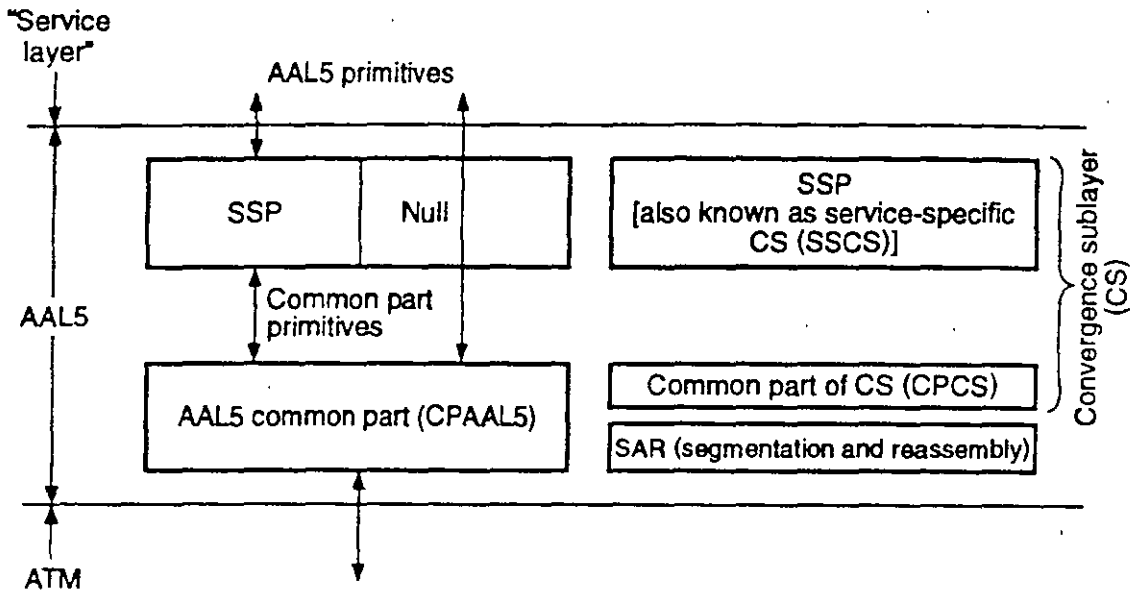


Figure 3.14 Structure of AAL Type 5.

therefore focuses on CPCS and SAR. Notice that $CPAAL5 = SAR + CPCS$. Also see Fig. 3.15.

3.5.1 Service provided by CPAAL5

The Common Part of AAL Type 5 provides the capability to transfer the CPAAL5_SDU from one CPAAL5 user to another CPAAL5 user through the ATM network. During this process, CPAAL5_SDUs may be corrupted or lost (in this case, an indication of the error is provided). Corrupted or

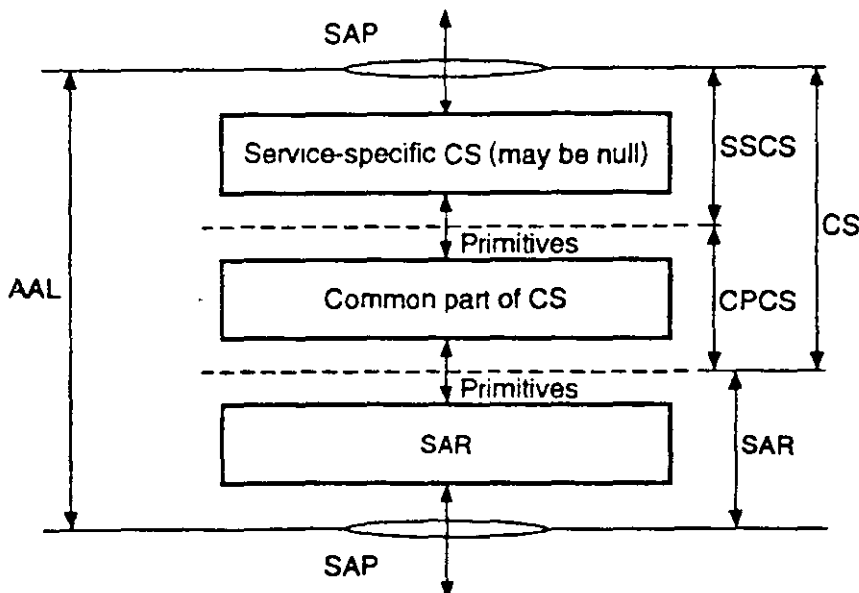


Figure 3.15 Another view of the structure of AAL Type 5.

lost CPAAL5_SDUs are not recovered by CPAAL5. CPAAL5 supports a message mode and a streaming mode. The message-mode service, streaming-mode service, and assured and nonassured operations as defined below for CPAAL5 are identical to those defined for AAL Type 3/4.

1. *Message-mode service.* The CPAAL5_SDU is passed across the CPAAL5 interface in exactly one Common Part AAL interface data unit (CPAAL5_IDU). This service provides the transport of fixed-size or variable-length CPAAL5_SDUs.
 - a. In the case of small fixed-size CPAAL5_SDUs, an internal blocking/deblocking function in the SSCS may be applied; it provides the transport of one or more fixed-size CPAAL5_SDUs in one SSCS_PDU.
 - b. In the case of variable-length CPAAL5_SDUs, an internal CPAAL5_SDU message segmentation/reassembling function in the SSCS may be applied. In this case, a single CPAAL5_SDU is transferred in one or more SSCS_PDUs.
 - c. Where the above options are not used, a single CPAAL5_SDU is transferred in one SSCS_PDU. When the SSCS is null, the CPAAL5_SDU is mapped to one CPCS_SDU.
2. *Streaming-mode service.* The CPAAL5_SDU is passed across the CPAAL5 interface in one or more CPAAL5_IDUs. The transfer of these CPAAL5_IDUs across the CPAAL5 interface may occur separated in time. This service provides the transport of variable-length CPAAL5_SDUs. Streaming-mode service includes an abort service by which the discarding of an CPAAL5_SDU that has been partially transferred across the AAL interface can be requested.
 - a. An internal CPAAL5_SDU message segmentation/reassembling function in the SSCS may be applied. In this case, all the CPAAL5_IDUs belonging to a single CPAAL5_SDU are transferred in one or more SSCS_PDUs.
 - b. An internal pipelining function may be applied. It provides the means by which the sending CPAAL5 entity initiates the transfer to the receiving CPAAL5 entity before it has the complete CPAAL5_SDU available.
 - c. Where option *a* is not used, all the CPAAL5_IDUs belonging to a single CPAAL5_SDU are transferred in one SSCS_PDU. When the SSCS is null, the CPAAL5_IDUs belonging to a single CPAAL5_SDU are mapped to one CPCS_SDU.

Both modes of service may offer the following peer-to-peer operational procedures:

- *Assured operations.* Every assured CPAAL5_SDU is delivered with exactly the data content that the user sent. The assured service is provided by retransmission of missing or corrupted SSCS_PDUs. Flow control is provided as a mandatory feature. The assured operation may be restricted to point-to-point AAL connections.
- *Nonassured operations.* Integral CPAAL5_SDUs may be lost or corrupted. Lost and corrupted CPAAL5_SDUs will not be corrected by retransmission. An optional feature may be provided to allow corrupted CPAAL5_SDUs to be delivered to the user (i.e., optional error discard). Flow control may be provided as an option.

Description of AAL connections. The CPAAL5 provides the capability to transfer the CPAAL5_SDU from one AAL5-SAP to another AAL5-SAP through the ATM network. CPAAL5 users have the ability to select a given AAL5-SAP associated with the QOS required to transport that CPAAL5_SDU (for example, delay- and loss-sensitive QOS).

The CPAAL5 in nonassured operation also provides the capability to transfer the CPAAL5_SDUs from one AAL5-SAP to more than one AAL5-SAP through the ATM network.

CPAAL5 makes use of the service provided by the underlying ATM Layer. Multiple AAL connections may be associated with a single ATM-Layer connection, allowing multiplexing at the AAL; however, if multiplexing is used in the AAL, it occurs in the SSCS. The AAL user selects the QOS provided by the AAL through the choice of the AAL5-SAP used for data transfer.

Primitives for the AAL. These primitives are service-specific and are contained in separate Recommendations on SSCS protocols.

The SSCS may be null, in the sense that it provides only for the mapping of the equivalent primitives of the AAL to CPCS and vice versa. In this case, the primitives for the AAL are equivalent to those for the CPCS but are identified as CPAAL5-UNITDATA.request, CPAAL5-UNITDATA.indication, CPAAL5-U-Abort.request, CPAAL5-U-Abort.indication, and CPAAL5-P-Abort.indication, consistent with the primitive naming convention at an SAP.

Primitives for the CPCS of the AAL. As there is no SAP between the sublayers of the AAL5, the primitives are called .invoke and .signal instead of the conventional .request and .indication to highlight the absence of the SAP.

CPCS-UNITDATA.Invoke and CPCS-UNITDATA.signal. These primitives are used for data transfer. The following parameters are defined:

- **Interface data (ID).** This parameter specifies the interface data unit exchanged between the CPCS and the SSCS entity. The ID is an integral multiple of 1 octet. If the CPCS entity is operating in message-mode service, the ID represents a complete CPCS_SDU; when operating in streaming-mode service, the ID does not necessarily represent a complete CPCS_SDU.
- **More (M).** In message-mode service, this parameter is not used. In streaming-mode service, this parameter specifies whether the interface data communicated contains a beginning/continuation of a CPCS_SDU or the end of a complete CPCS_SDU.
- **CPCS loss priority (CPCS-LP).** This parameter indicates the loss priority for the associated CPCS_SDU. It can take only two values, one for high priority and the other for low priority. The use of this parameter in streaming mode is for further study. This parameter is mapped to and from the SAR-LP parameter.
- **CPCS congestion indication (CPCS-CI).** This parameter indicates that the associated CPCS_SDU has experienced congestion. The use of this parameter in streaming mode is for further study. This parameter is mapped to and from the SAR-CI parameter.
- **CPCS user-to-user indication (CPCS-UU).** This parameter is transparently transported by the CPCS between peer CPCS users.
- **Reception status (RS).** This parameter indicates that the associated CPCS_SDU delivered may be corrupted. This parameter is utilized only if the corrupted data delivery option is used.

Depending on the service mode (message- or streaming-mode service, discarding or delivery of errored information), not all parameters are required.

CPCS-U-Abort.invoke and **CPCS-U-Abort.signal.** These primitives are used by the CPCS user to invoke the abort service. They are also used to signal to the CPCS user that a partially delivered CPCS_SDU is to be discarded by instruction from its peer entity. No parameters are defined. These primitives are not used in message mode.

CPCS-P-Abort.signal. This primitive is used by the CPCS entity to signal to its user that a partially delivered CPCS_SDU is to be discarded because of the occurrence of some error in the CPCS or below. No parameters are defined. This primitive is not used in message mode.

Primitives for the SAR sublayer of the AAL. These primitives model the exchange of information between the SAR sublayer and the CPCS.

As there is no SAP between the sublayers of the AAL5, the primitives are called `.invoke` and `.signal` instead of the conventional `.request` and `.indication` to highlight the absence of the SAP.

`SAR-UNITDATA.invoke` and `SAR-UNITDATA.signal`. These primitives are used for data transfer. The following parameters are defined:

- **Interface data (ID).** This parameter specifies the interface data unit exchanged between the SAR and the CPCS entity. The ID is an integral multiple of 48 octets. It does not necessarily represent a complete SAR_SDU.
- **More (M).** This parameter specifies whether the interface data communicated contains the end of the SAR_SDU.
- **SAR loss priority (SAR-LP).** This parameter indicates the loss priority for the associated SAR interface data. It can take on two values, one for high priority and the other for low priority. This parameter is mapped to the ATM Layer's submitted loss priority parameter and from the ATM Layer's received loss priority parameter.
- **SAR congestion indication (SAR-CI).** This parameter indicates whether the associated SAR interface data has experienced congestion. This parameter is mapped to and from the ATM Layer's congestion indication parameter.

3.5.2 Functions, structure, and coding of AAL5

Functions of the SAR Sublayer. The SAR Sublayer functions are performed on an SAR_PDU basis. The SAR Sublayer accepts variable-length SAR_SDUs which are integral multiples of 48 octets from the CPCS and generates SAR_PDUs containing 48 octets of SAR_SDU data. It supports the preservation of SAR_SDUs by providing for an "end of SAR_SDU" indication.

SAR_PDU structure and coding. The SAR Sublayer function utilizes the ATM-Layer-user-to-ATM-Layer-user (AUU) parameter of the ATM Layer primitives to indicate that a SAR_PDU contains the end of a




Figure 3.16 SAR_PDU format for AAL5. [Note: The payload type (PT) field belongs to the ATM header. It conveys the value of the AUU parameter end-to-end.]

SAR_SDU. A SAR_PDU where the value of the AUU parameter is 1 indicates the end of a SAR_SDU; a value of 0 indicates the beginning or continuation of a SAR_SDU. The structure of the SAR_PDU is shown in Fig. 3.16.

Convergence Sublayer. The CPCS has the following service characteristics.

- Nonassured data transfer of user data frames with any length measured in octets from 1 to 65,535 octets.
- The CPCS connection will be established by management or by the Control Plane.
- Error detection and indication (bit error and cell loss or gain).
- CPCS_SDU sequence integrity on each CPCS connection.

Functions of the CPCS. The CPCS functions are performed per CPCS_PDU. The CPCS provides several functions in support of the CPCS service user. The functions provided depend on whether the CPCS service user is operating in message or streaming mode.

1. *Message mode service.* The CPCS_SDU is passed across the CPCS interface in exactly one CPCS-IDU. This service provides the transport of a single CPCS_SDU in one CPCS_PDU.
2. *Streaming mode service.* The CPCS_SDU is passed across the CPCS interface in one or more CPCS-IDUs. The transfer of these CPCS-IDUs across the CPCS interface may occur separated in time. This service provides the transport of all the CPCS-IDUs belonging to a single CPCS_SDU into one CPCS_PDU. An internal pipelining function in the CPCS may be applied which provides the means by which the sending CPCS entity initiates the transfer to the receiving CPCS entity before it has the complete CPCS_SDU available. Streaming-mode service includes an abort service by which the discarding of a CPCS_SDU partially transferred across the interface can be requested.

Note: At the sending side, parts of the CPCS_PDU may have to be buffered if the restriction "interface data are a multiple of 48 octets" cannot be satisfied.

The functions implemented by the CPCS include:

1. *Preservation of CPCS_SDU.* This function provides for the delineation and transparency of CPCS_SDUs.
2. *Preservation of CPCS user-to-user information.* This function provides for the transparent transfer of CPCS user-to-user information.

3. *Error detection and handling.* This function provides for the detection and handling of CPCS_PDU corruption. Corrupted CPCS_SDUs are either discarded or optionally delivered to the SSCS. The procedures for delivery of corrupted CPCS_SDUs are for further study. When delivering errored information to the CPCS user, an error indication is associated with the delivery. Examples of detected errors would include received length and CPCS_PDU Length field mismatch including buffer overflow, an improperly formatted CPCS_PDU, and CPCS CRC errors.
4. *Abort.* This function provides for the means to abort a partially transmitted CPCS_SDU. This function is indicated in the Length field.
5. *Padding.* A padding function provides for 48-octet alignment of the CPCS_PDU trailer.

CPCS structure and coding. The CPCS functions require an 8-octet CPCS_PDU trailer. The CPCS_PDU trailer is always located in the last 8 octets of the last SAR_PDU of the CPCS_PDU. Therefore, a padding field provides for a 48-octet alignment of the CPCS_PDU. The




Figure 3.17 CPAAL5_PDU.

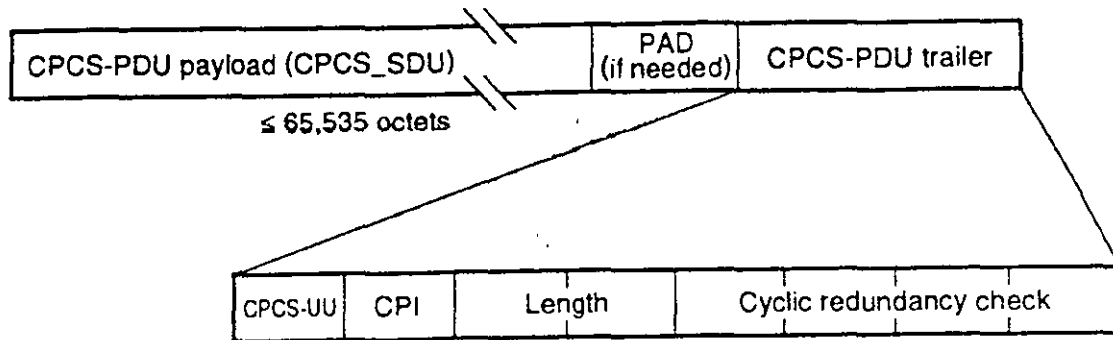


Figure 3.18 CPAAL5_PDU, another view.

CPCS_PDU trailer, the padding field, and the CPCS_PDU payload make up the CPCS_PDU.

The coding of the CPCS_PDU conforms to the coding conventions specified in 2.1 of Recommendation I.361. See Figs. 3.17 and 3.18.

1. *CPCS_PDU payload.* The CPCS_PDU payload is the CPCS_SDU.
2. *Padding (Pad) field.* Between the end of the CPCS_PDU payload and the CPCS_PDU trailer, there will be from 0 to 47 unused octets. These unused octets are called the padding (Pad) field; they are strictly used as filler octets and do not convey any information. Any coding is acceptable. This padding field complements the CPCS_PDU (including CPCS_PDU payload, padding field, and CPCS_PDU trailer) to an integral multiple of 48 octets.
3. *CPCS User-to-User Indication (CPCS-UU) field.* The CPCS-UU field is used to transparently transfer CPCS user-to-user information.
4. *Common Part Indicator (CPI) field.* One of the functions of the CPI field is to align the CPCS_PDU trailer to 64 bits. Other functions are for further study. Possible additional functions may include identification of Layer Management messages. When only the 64-bit alignment function is used, this field is coded as zero.
5. *Length field.* The Length field is used to encode the length of the CPCS_PDU payload field. The Length field value is also used by the receiver to detect the loss or gain of information. The length is binary coded as number of octets. A Length field coded as zero is used for the abort function.
6. *CRC field.* The CRC-32 is used to detect bit errors in the CPCS_PDU. The CRC field is filled with the value of a CRC calculation which is performed over the entire contents of the CPCS_PDU, including the CPCS_PDU payload, the Pad field, and the first 4 octets of the CPCS_PDU trailer. The CRC field shall contain the 1s complement of the sum (modulo 2) of

- a. The remainder of $xk(x^{31} + x^{30} + \dots + x + 1)$ divided (modulo 2) by the generator polynomial, where k is the number of bits of the information over which the CRC is calculated.
- b. The remainder of the division (modulo 2) by the generator polynomial of the product of x^{32} and the information over which the CRC is calculated.

The CRC-32 generator polynomial is:

$$G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$

The result of the CRC calculation is placed with the least significant bit right-justified in the CRC field.

As a typical implementation at the transmitter, the initial content of the register of the device computing the remainder of the division is preset to all 1s and is then modified by division by the generator polynomial (as described above) of the information over which the CRC is to be calculated; the 1s complement of the resulting remainder is put into the CRC field.

As a typical implementation at the receiver, the initial content of the register of the device computing the remainder of the division is preset to all 1s. The final remainder, after multiplication by x^{32} and then division (modulo 2) by the generator polynomial of the serial incoming CPCS_PDU, will be (in the absence of errors)

$$C(x) = x^{31} + x^{30} + x^{26} + x^{25} + x^{24} + x^{18} + x^{15} + x^{14} + x^{12} + x^{11} + x^{10} + x^8 + x^6 + x^5 + x^4 + x^3 + x + 1$$

3.6 Frame Relay Service-Specific AAL

The Frame Relay Service-Specific ATM Adaptation Layer Convergence Sublayer (FR-SSCS) is positioned in the upper part of the ATM Adaptation Layer; it is located above the CPAAL5, as shown in Figs. 3.19 and 3.20. It is an example of an SSP. The purpose of the FR-SSCS protocol at an ATM CPE (that is, user's equipment) is to emulate the Frame Relaying Bearer Service (FRBS) in an ATM-based network (Fig. 3.19). On network nodes, the FR-SSCS is used for interworking between an ATM-based network and a Q.922-based Frame Relaying Network (Fig. 3.20).

The FR-SSCS protocol provides for the transport of variable-length frames with error detection. The FR-SSCS provides its service over

*This discussion is based on Ref. 4.

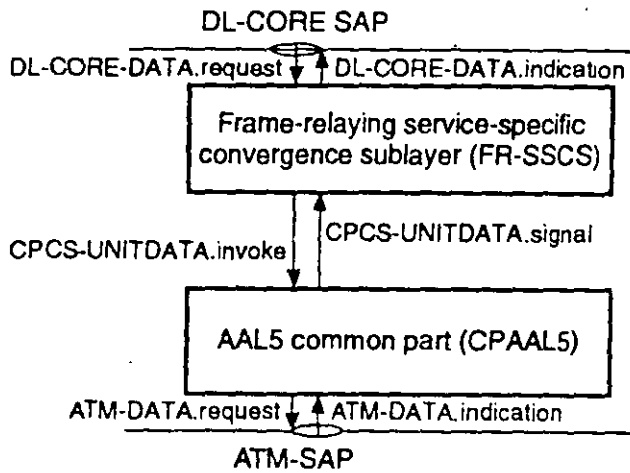
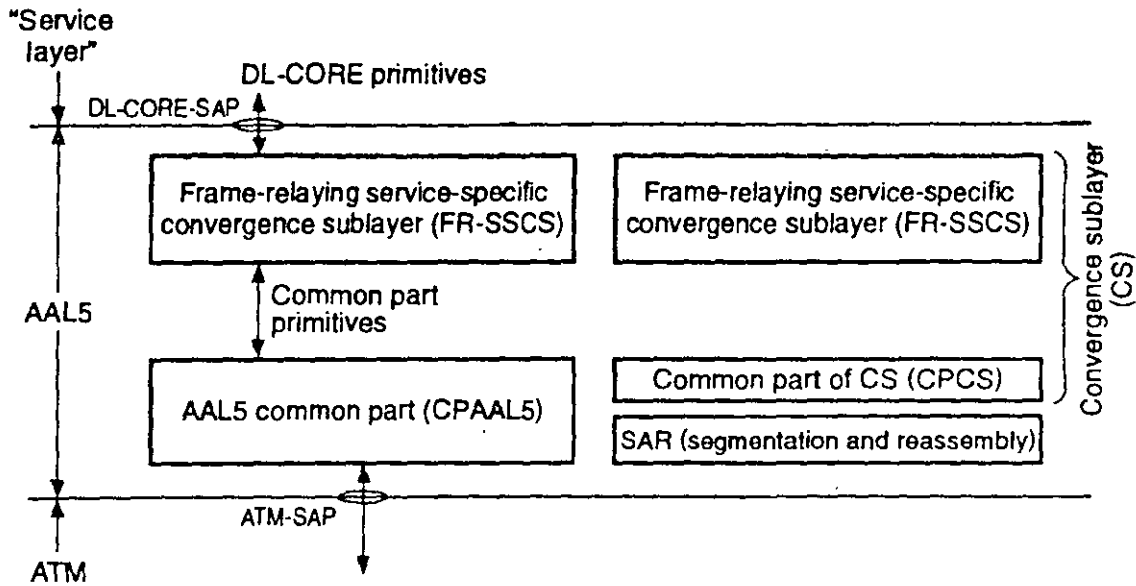


Figure 3.19 AAL5 for interworking of frame relay and ATM (in CPE).

preestablished connections with negotiated traffic parameters. An FR-SSCS connection represents the segment of an end-to-end frame relay (FR) connection over B-ISDN. At an ATM-based B-TE, the FR-SSCS connection is terminated at the point of termination of the FR-SSCS service and represents one end of the FR connection. Optionally, multiplexing may be performed at the FR-SSCS, allowing various FR-SSCS connections to be associated with a single CPAAL5 connection (and with the corresponding ATM connection). FR-SSCS connections within a CPAAL5 connection are uniquely identified by data link connection identifiers (DLCIs). The establishment (or provisioning) and initialization of an FR-SSCS connection is performed by interaction with FR-SSCS Layer Management (MFR-SSCS) entities. The traffic parameters of each FR-SSCS connection are determined at the time of its estab-

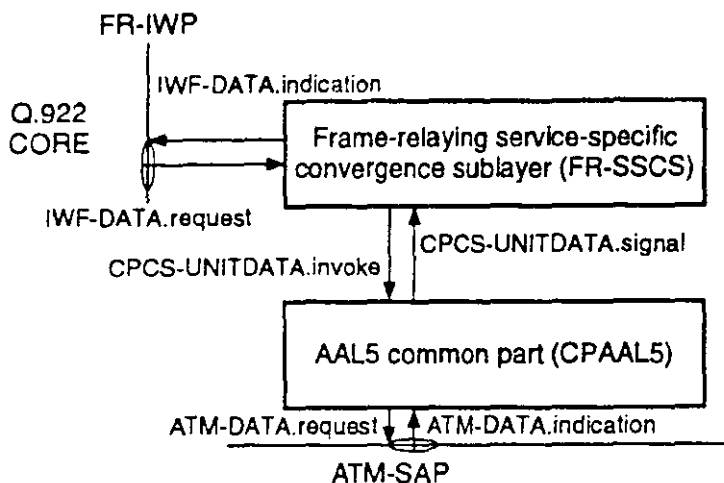


Figure 3.20 AAL5 for interworking of frame relay and ATM (in a network element supporting an interworking function). FR-IWP = frame relay interworking point.

lishment. The negotiated traffic parameters are bounded by the ATM Layer/CPAAL5 connection characteristics.

The FR-SSCS can indicate to its user that the receiver FR_SSCS_PDU has experienced congestion (forward congestion) or that an FR_SSCS_PDU traveling in the opposite (sending) direction has experienced congestion (backward congestion). The FR-SSCS allows for two discard eligibility priorities. The FR-SSCS user can request the discard eligibility (loss priority) associated with each FR_SSCS_SDU. The FR-SSCS uses the CPAAL5 message-mode service without the corrupted data delivery option and preserves the FR_SSCS_SDU sequence integrity.

The MFR-SSCS is responsible for the following actions: assignment of the FR_SSCS association necessary for the establishment or provisioning of FR-SSCS connections between peer FR-SSCS entities, resetting the parameters and state variables associated with a FR-SSCS connection when required, releasing the association created for a FR-SSCS connection between peer FR-SSCS entities, and performance monitoring of the quality of the FR-SSCS connection service provided through notification of errors (i.e., FR_SSCS_PDU discards resulting from errors in the FR_SSCS_PDU).

Service provided by the FR-SSCS. The FR-SSCS provides services to (1) the core service user (upper layer) at ATM-based B-TEs or (2) the Q.922-CORE Data Link Layer (Q.922-DLL) on network nodes at interworking functions (IWFs) points. Only item (1) is covered here.

The FR-SSCS provides the capability to transfer variable-length octet-aligned FR_SSCS_SDUs from one or more FR_SSCS users. The FR-SSCS Sublayer preserves the FR_SSCS_SDU sequence integrity

within an FR-SSCS connection. During this process, FR_SSCS_SDUs may be lost or corrupted. Lost or corrupted FR_SSCS_SDUs are not recovered by the FR-SSCS. The FR-SSCS uses CPAAL5 message-mode service without the corrupted data delivery option.

FR-SSCS functions. The functions provided by the FR_SSCS include

Multiplexing/demultiplexing. This function provides for the optional multiplexing and demultiplexing of FR-SSCS connections into a single CPAAL5 connection. The number of FR-SSCS connections supported over a CPAAL5 connection is defined at connection establishment or provisioning. The default number of FR-SSCS connections when multiplexing is not supported is 1. Within a given FR-SSCS connection, sequence integrity is preserved.

Inspection of the FR_SSCS_PDU length. This function inspects the FR_SSCS_PDU to ensure that it consists of an integral number of octets and to ensure that it is neither too long nor too short.

Congestion control. These functions provide the means to notify the end user that congestion avoidance procedures should be initiated,

TABLE 3.6 DL-CORE Primitives

DL-CORE-DATA.request (DL_CORE_User_Data, Discard_Eligibility, DL_CORE_Service_User_Protocol_Control_Information)	This primitive is received from the FR-SSCS user to request the transfer of an FR_SSCS_SDU over the associated FR-SSCS connection.
DL-CORE-DATA.indication (DL_CORE_User_Data, Congestion_Encountered_Backward, Congestion_Encountered_Forward, DL_CORE_Service_User_Protocol_Information)	This primitive is used to the FR-SSCS user to indicate the arrival of an FR_SSCS_SDU from the associated connection.

Description of parameters:

DL_CORE_User_Data: This parameter specifies the FR_SSCS_SDU transported between the FR-SSCS user and the FR-SSCS. This parameter is octet-aligned and can range from 1 to a maximum of at least 4096 octets in length.

Discard_Eligibility: This parameter indicates the loss priority assigned to the FR_SSCS_SDU. Two levels of priority are identified: High and Low. A value of High indicates that the FR_SSCS_SDU may experience a better quality of service with respect to loss (i.e., minimal loss) than if the Discard_Eligibility parameter were set to Low.

DL_CORE_Service_Protocol_Information: This parameter specifies a 1-bit FR-SSCS/Q.922-DLL user control information to be transparently transferred between FR-SSCS/Q.922-DLL users.

Congestion_Encountered_Backward: This parameter indicates that an FR_SSCS_SDU has experienced congestion in the opposite (sending) direction, and therefore that an FR_SSCS_SDU sent on the corresponding connection may encounter congested resources. This parameter may take on two values: True or False. A value of True indicates that an FR_SSCS_SDU has experienced congestion in the opposite (sending) direction of the connection.

Congestion_Encountered_Forward: This parameter indicates that the received FR_SSCS_SDU has experienced congestion. This parameter may take two values: True or False. A value of True indicates that the FR_SSCS_SDU has experienced congestion.

where applicable (congestion control forward and congestion control backward). In addition, the functions provide the means for the end user and/or the network to indicate what frames should be discarded in a congestion situation.

Primitives. The information exchanged between the FR-SSCS and its user (for ATM-based B-TEs) is modeled by the primitives of Table 3.6 (which are the same DL-CORE primitives in Annex C of ITU-T Recommendation I.233.1).

Services expected from the CPAAL5. The FR-SSCS expects the CPAAL5 to provide the capability to transfer variable-length (from 3 to a maximum of at least 4100 octets) octet-aligned FR_SSCS_SDUs, with error detection and in sequence, between communicating FR-SSCS entities. Lost or corrupted FR_SSCS_PDUs are not expected to be recovered by the CPAAL5. Multicast services, derived from the ATM Layer, are expected.

The FR-SSCS entity expects the CPAAL5 to provide each FR_SSCS_PDU (CPAAL5_SDU) with the CP_Congestion_Indication (True or False) set to the value of the Congestion_Indication received by the ATM Layer with the last ATM_SDU conforming to the CPAAL5_SDU; and with the CP_Loss_Priority set to either Low, if any of the ATM_SDUs conforming to the CPAAL5_SDU was received with the Received_Loss_Priority parameter set to Low, or High otherwise.

The FR-SSCS entity passes each FR_SSCS_PDU (CPAAL5_SDU) with the CP_Loss_Priority set to the value of the Discard_Eligibility parameter received from the upper layer or the Q.922-DLL (High or Low), the CP_Congestion_Indication (True or False) always set to False, and the User_User_Indication parameter always set to zero.

3.7 Signaling ATM Adaptation Layer (SAAL)

This section describes the Signaling ATM Adaptation Layer (SAAL) for use at the UNI. SAAL is used in the Control Plane. (This topic could also have been treated in the next chapter, but it was decided to include it here with other AALs.)

The SAAL resides between the ATM Layer and Q.2931 in the user's equipment, specifically in the software implementing the Control Plane (i.e., the signaling capability). The purpose of the SAAL is to provide reliable transport of Q.2931 messages between peer Q.2931 entities (e.g., ATM switch and host) over the ATM Layer. The SAAL is composed of two sublayers, a Common Part and a Service-Specific

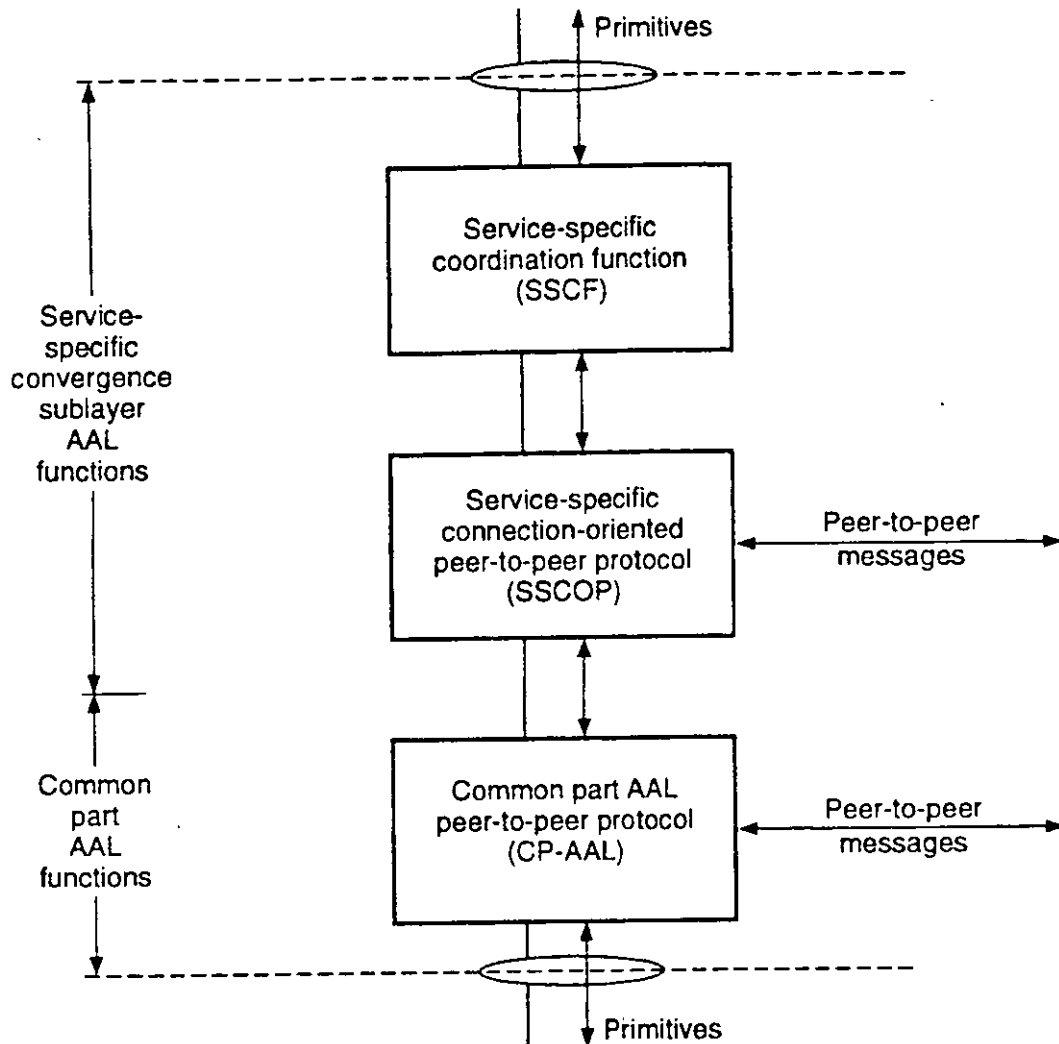


Figure 3.21 SAAL structure. (Note: This figure represents the allocation of functions and is not intended to illustrate sublayers as defined by OSI modeling principles.)

Part. The Service-Specific Part is further subdivided into a Service-Specific Coordination Function (SSCF) and a Service-Specific Connection-Oriented Protocol (SSCOP). Figure 3.21 illustrates the structure of the SAAL.⁵

The SAAL for supporting signaling uses the protocol structure illustrated in Fig. 3.21. The Common Part AAL protocol provides unassured information transfer and a mechanism for detecting corruption of SDUs. The AAL Type 5 Common Part protocol is used to support signaling. The AAL Type 5 Common Part protocol is specified in Draft Recommendation I.363.

The SAAL for supporting signaling at the UNI uses the AAL Type 5 Common Part protocol, discussed above, as specified in Ref. 7 with minor amendments.⁸

The Service-Specific Connection-Oriented Protocol (SSCOP) resides in the Service-Specific Convergence Sublayer (SSCS) of the SAAL. SSCOP is used to transfer variable-length service data units (SDUs) between users of SSCOP. SSCOP provides for the recovery of lost or corrupted SDUs. SSCOP is specified in ITU-T Recommendation Q.2110.⁹

The SAAL for supporting signaling utilizes SSCOP as specified in Q.2110.⁹

An SSCF maps the service of SSCOP to the needs of the SSCF user. Different SSCFs may be defined to support the needs of different AAL users. The SSCF used to support Q.93B at the UNI is specified in ITU-T Recommendation Q.2130.¹⁰

The external behavior of the SAAL at the UNI appears as if the UNI SSCF specified in Q.2130¹⁰ were implemented.

References

1. CCITT I.362, *B-ISDN AAL Functional Description*, Geneva, Switzerland, 1992.
2. CCITT I.363, *B-ISDN AAL Specification*, Geneva, Switzerland, 1992.
3. ANSI T1.BCR-199x, *Broadband ISDN—ATM Adaptation Layer for Constant Bit Rate Services Functionality Specification*, New York, Nov. 13, 1992.
4. Bellcore, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements*, TA-NWT-001113, Issue 2, July 1993.
5. B. Kittams, Bellcore, personal communication, May 1993.
6. TIS1.5/93-52, *Broadband Aspects of ISDN Baseline Document*, TIS1 Technical Subcommittee, August 1990, Chief Editor: Rajeev Sinha. Reissued February 1993, Chief Editor: Erwin Fandrich.
7. CCITT Document TD-XVIII/10 (AAL5), "AAL Type 5, Draft Recommendation Text for Section 6 of I.363," Geneva, Switzerland, Jan. 29, 1993.
8. ATM Forum, *ATM User-Network Interface Specification*, Version 3.0, August 1993.
9. ITU Document DT/11/3-28 [Q.SAAL1 (now Q.2110)], *Service-Specific Connection-Oriented Protocol (SSCOP) Specification*, Geneva, Switzerland, May 17, 1993.
10. ITU Document DT/11/3-XX [Q.SAAL2 (now Q.2130)], *Service-Specific Connection-Oriented Protocol (SSCOP) Specification*, Geneva, Switzerland, May 17, 1993.

Cisco LAN Switching Products: Catalyst 5000

The Catalyst Family of Multilayer Switches: Optimized for Flexible Wiring Closets

The first member of the Catalyst family, introduced in March 1994, addressed the increasing needs of client/server applications by boosting the performance of multisegment hubs and dedicating bandwidth to servers. In less than one year, the Catalyst 1200 acquired a 20 percent share of the switching market, according to market projections. The Catalyst 5000, introduced in March 1995, represents the next-generation switching system for the wiring closet with dedicated bandwidth to the desktop. The Catalyst family gives users the scalability to build large switched internetworks with multilayer intelligence.

The Cisco Catalyst 5000: The Next Generation of Multilayer LAN Switching


The Catalyst 5000 is a modular switching platform that will meet the ever-changing needs of today's high-performance, bandwidth-intensive, multiple-media network switching applications. Dedicated bandwidth is delivered to users through multiple-media switching options that encompass 10-Mbps Ethernet, 100-Mbps Ethernet, and ATM, with future capability for switched FDDI and switched Token Ring.

As networks migrate to higher-speed technologies such as Fast Ethernet and ATM, the long-term viability of the Catalyst 5000 is enhanced by five slots that provide configuration flexibility. The first slot contains a Supervisor Engine that enables aggregate switching and switch management. The remaining four slots support a growing combination of switching modules that include switched 10BaseT/10BaseFL, switched 100-Mbps Fast Ethernet, and an ATM LAN emulation module. This mix of interface and backbone modules allows the Catalyst 5000 to integrate the broadest range of environments in the industry, providing a cohesive network solution and a migration path to switched ATM-based networks while leveraging infrastructure investments.

ATM backbone access is achieved through an ATM LAN emulation module that allows applications based on standard protocols such as TCP/IP, Novell NetWare, DECnet, and AppleTalk to run unchanged over ATM networks. An innovative design executes ATM protocols as onboard processor-executable code, ensuring compatibility with future software versions.

Multiple Switching Options

Cisco Systems
Catalyst 5000
serves the needs of
multiple-media
network switching
applications with
options that include
10-Mbps Ethernet,
100-Mbps Ethernet,
switched Token
Ring, switched
FDDI, and ATM.



Because many of today's organizations experience significant personnel relocations every year, the Catalyst 5000 gives users the flexibility to support the formation of VLANs within and between Catalyst 5000 switches and across the internetwork, spanning routers and ATM. The architecture will scale to support up to 1024 VLANs and can be defined and maintained across platforms through ATM or 100-Mbps links.

Unlike traditional shared hubs, the Catalyst 5000 architecture is designed and optimized for switching. The Catalyst 5000's switching backplane operates at 1.2 gigabits per second (Gbps) and provides nonblocking performance for all switched 10-Mbps Ethernet interfaces. The internal switching architecture of the Catalyst 5000 supports multiple media options including Ethernet, Fast Ethernet, 100BaseVG-AnyLAN, Token Ring, FDDI, and ATM. A dedicated management bus provides distributed access to all switching modules for monitoring performance, controlling configuration and VLANs, and updating operational software for each module.

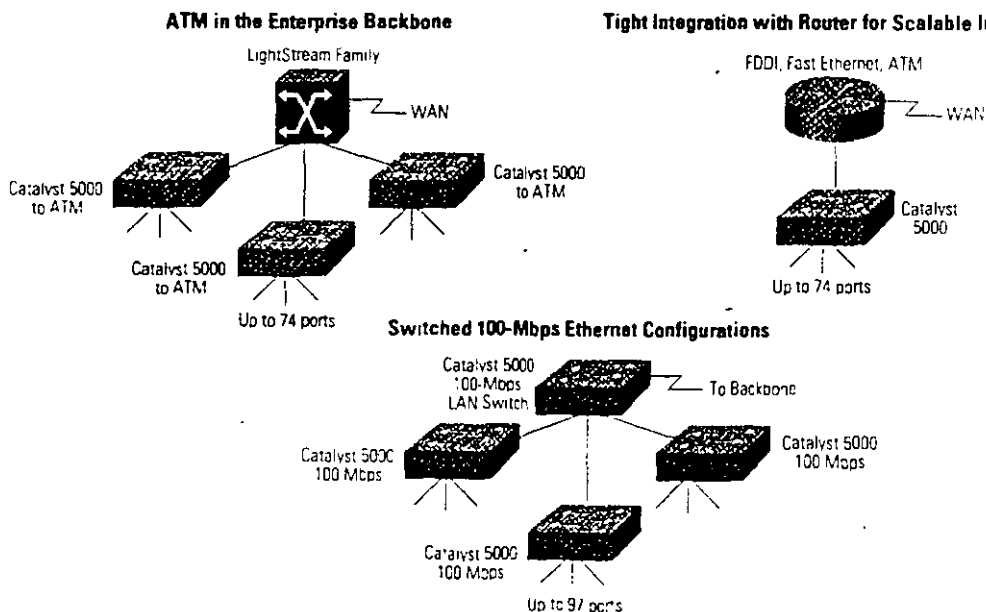
Unique Traffic Management

Support for traffic management by switches is one of the most important attributes of a scalable switched internetwork, and the Catalyst 5000 includes several key traffic management features. A large, 192-kilobyte (KB) buffer ensures adequate port buffering for workgroup applications without dropping information during peak traffic periods. Tri-level priority on the backplane ensures that delay-sensitive applications receive the necessary priority on a port-by-port basis.

With increasing support for high-performance, high port density, multimedia solutions, the Catalyst 5000 architecture has the foundation to support the growing needs of business networks today and into the future.

Catalyst 5000 Offers Maximum Port Density

The Catalyst 5000 platform features a highly scalable architecture that provides maximum performance to large workgroups.



Cisco LAN Switching Products: Catalyst 5000

	Feature	Function	Benefit
Basic Platform	Modular, five-slot chassis	Supports required supervisor engine and four additional modules	Allows users flexibility to add, or mix and match modules as needed
	Fits standard 19-inch rack, rack-mounting hardware and cable guides included	Compatible with existing wiring closet rack; mounts at front or rear of chassis	Ensures ease of installation in wiring closet and data center
	Fault-tolerant power supplies	Offers one or two 376-watt power supplies with auto-sensing	Increases reliability
	Switch interface modules	Offers wide variety of interface choices including 10-Mbps and 100-Mbps Ethernet and ATM; Token Ring and CDDV/FDDI to be added in future releases	Accommodates all LAN and ATM technologies
	Specialized Supervisor Engine module has hardware-based switching engine, bridge address table for a maximum of 16,000 MAC addresses, two 100-Mbps Fast Ethernet uplinks, and a network management processor	Delivers over 1 million pps, low latency, low-cost switching for desktop and backbone applications, full local and remote management	Provides a completely manageable switching system that delivers high-performance switching for even the most demanding workgroup applications
	Hot-swap of power supplies and modules	Enables users to add, move, or change modules any time without bringing down the Catalyst 5000 or losing connectivity with other network devices	Increases reliability by reducing downtime
	Internal architecture [®]	Single switching backplane accommodates Ethernet and ATM, simultaneously (FDDI and Token Ring in future releases)	Delivers a single-system solution for all current desktop switching needs with a migration path to ATM
	1.2-Gbps switching backplane	Forwards more than 1 million 64-byte Ethernet packets (equates to nonblocking architecture for 100 10-Mbps Ethernet interfaces at wire speed)	Designed to meet the demands of a fully populated, 10-Mbps system with all interfaces operating at wire speed (each interface communicates with only one other interface)
	Three levels of priority on the data-switching backplane	Accommodates technologies that implement prioritization schemes, such as ATM, 100VG-AnyLAN, and Token Ring; enables user to define higher priority on a per-interface basis, allows any interface to reach the highest priority when its buffer reaches capacity	Accommodates bursty traffic and prevents higher-layer protocol time-outs by giving the interface the highest priority on the bus until traffic is reduced, allows managers to prioritize data transmission to critical resources such as servers or remote routers
Custom application-specific integrated circuit (ASIC) development	Uses a custom ASIC controller on each Ethernet interface	Lowers cost per switched interface to the user; enables state-of-the-art features such as VLANs	
Standards-Based Interfaces	Switched 10-Mbps Ethernet	Offers choice of Category 3 unshielded twisted pair (UTP) cable (10BaseT) and fiber-optic cable (10BaseFL), supports both full- and half-duplex operation	Eases integration into existing installed base of 10-Mbps Ethernet hubs and adapters by preserving wiring hardware and application investment, leverages existing UTP support; allows long-distance runs with fiber across a campus, increases capacity with full-duplex operation
	24 interfaces of 10BaseT per module	Uses female RJ 21 connectors	High density, low cost per interface connection into existing telco-configured wiring
	12 interfaces of 10BaseFL on fiber per module	Uses female RJ-21 connectors	Higher density than competitors
	Switched 100 Mbps Ethernet	Offers choice of Category 5 UTP cable and, in the future, fiber-optic cable for 100BaseFX and Category 3 UTP for 100BaseTX and 100BaseVG-AnyLAN; supports both full- and half-duplex operation; supports auto-sensing between 100-Mbps Ethernet	Ensures compatibility with a large number of 100-Mbps Ethernet hubs and adapters; uses cost-effective UTP for desktop connectivity; allows long-distance runs with fiber across a campus in full-duplex mode, increases capacity with full-duplex mode


Catalyst 5000 Features and Benefits

The Catalyst 5000 is a strategic switching platform that will support the growing bandwidth needs of today's workgroups. It offers users a strategic foundation to shift from shared media hubs to switching in the wiring closet.

	Feature	Function	Benefit
Standards-Based Interfaces (con't)	12 interfaces of 100BaseTX per module	Uses female RJ-45 connectors	Provides lower-cost, 100-Mbps interface to routers, hubs, and adapters
	FDDI connection to backbone, switch, or router	Provides scalable, high-speed connection to server; offers connection to choice of multimode (MIC female), single mode (ST female), or Category 5 UTP (RJ-45 female) cabling	Ensures interoperability and compatibility with existing FDDI networks
	ATM backbone connection to ATM switch	Provides scalable, high-speed connection; offers connection to choice of multimode or single mode fiber	Connects to ATM backbone
	Onboard hardware diagnostics and LEDs	Status LED on each module shows successful completion or minor and major failure of power-up diagnostic; Link Good LED shows status of 10BaseT and 100BaseTX interface; Switch Load LEDs show backplane data bus utilization	Enables easy visual troubleshooting
Virtual LANs	1024 VLANs supported	Ensures high number of switched VLANs are available for enterprise network; enables users to select interfaces on multiple system-wide switches to create a VLAN; allows LANs to be multiplexed between switches using Fast Ethernet, FDDI, and ATM	Eases network administration by enabling users to be logically grouped together regardless of physical interface location for performance and security considerations; provides VLAN capability without forcing users to invest in new backbone technology
	Inter Switch Link (ISL)	Supports VLANs between switches using any Fast Ethernet interfaces for cost-effective point-to-point desktop switching; can operate at full-duplex (100 Mbps) over low-cost copper and long-distance fiber connections	Leverages cost-effective Fast Ethernet technology to provide VLANs across an organization's network; offers increased Fast Ethernet capacity at an incremental cost
Network Management	Full SNMP management (Ethernet MIB, ILMI MIB, FDDI MIB, Bridge MIB, AToMIC, MIB II, and system extensions)	Enables Catalyst 5000 to be managed from an SNMP-based management station	Eases management from installed network management platforms
	Local (out-of-band) management	Supervisor has an EIA/TIA-232 interface for modem or console terminal connection	Manages Catalyst 5000 from directly attached terminal, modem, or PC
	In-band (Telnet) management	Accessible through any switched or ATM interface	Manages Catalyst 5000 from anywhere in the network
	Management console	Command line interface	Provides easy-to-use ASCII text interface that requires no special applications
	NVRAM	Preserves configuration information	Eases switch configuration and maintenance
	Flash memory for TFTP download/upload of operating software	Remotely downloads new revisions of operating system without hardware change	Reduces cost of administering software upgrades by providing centralized network management capability
	Status LEDs	Allows user to visually monitor operation of power supplies, fans, switches, and backbone interfaces	Eases switch diagnosis and troubleshooting
Switch Management Applications	CiscoView application	Provides intuitive, GUI interface that supports chassis physical view, configuration, performance monitoring, and troubleshooting	Simplifies switched internetwork management; provides integrated management solution with one tool for determining system status
	VLAN management application	Intuitive GUI for adding new users, moving users between wiring closets, changing users' VLAN associations	Reduces training time and ensures that administrator will be able to manage VLANs

High-Performance Scalable Networking with Routed ATM

Right-Sizing the Bandwidth in Today's Networks





High-Performance Scalable Networking with Routed ATM

Contents

Executive Summary	2
Planning for High Bandwidth Demand	2
Customer Guidelines	2
High-Performance Scalable Networking	2
Today's Collapsed Backbone Architecture	4
HPSN Stage 1: Enhancing the Collapsed Backbone with Additional Horizontal and Vertical LANs	5
Overcoming Constraints on Downlinks	5
Port Grouping for Virtual LANs and Workgroups	5
Creating Routing Clusters with Switching	6
HPSN Stage 2: Creating High-Speed Downlinks to Increase Bandwidth and Reduce Segmentation	6
Route Caching for Scalable Routing	7
LinkSwitch for Connections Between Segments	7
High-Speed Downlinks Using a Switching Hub and Bridge/Router	7
ATM as a Downlink Technology	8
CellBuilder for ATM Conversion	8
Integrating Meshed Campus Backbones with ATM Switching	9
HPSN Stage 3: Enhancing the Collapsed Backbone with Routed ATM	9
Route Determination for ATM End Systems	9
Managing the Network	10
Today's Foundation for HPSN Migration	10
Summary	11
Glossary	12
For More Information	12

Executive Summary

This paper explains what 3Com's High-Performance Scalable Networking (HPSN) is, how it works, and ways managers can use it to migrate their networks to higher bandwidths without assuming the cost and risky "tear-it-down" approach.

Specifically, the paper focuses on how managers can use HPSN to enhance building and campus backbones in order to increase performance in existing floor plans

without the need for major relocations or major network redesigns. It also discusses the benefits of HPSN, such as the ability to scale network performance to meet specific business requirements, and the ability to extend the reach of today's LAN resources by dramatically reducing the cost and complexity of providing remote and WAN connections across the enterprise.

Planning for High Bandwidth Demand

Enterprise networks linking desktop and mobile computer clients with servers and other computing resources are critical to the information flow within many of today's companies.

However, demand for more capacity is beginning to strain the capabilities of these networks. As current technologies are pushed to their limits, MIS managers are under pressure to provide users with more bandwidth while continuing to preserve and optimize existing investments.

A network manager planning to meet bandwidth demand by expanding an existing network or designing a new one must answer a number of critical questions:

- Which high-speed technologies will best meet network requirements for various parts of the network now and in the future?
- How should this technology be implemented—with mixed-media hubs, switches, multiplexers, routers, or a combination of products?
- Can higher bandwidth be provided for backbones, servers, and client PCs at a reasonable cost, and without introducing unnecessary complexity?
- Will users of new technologies be able to communicate transparently with existing LAN users?
- How can the network be optimized for new client-server applications?
- And, most importantly, how can this migration to higher performance be accomplished without disrupting the existing network and sacrificing productivity?

Customer Guidelines

In formulating its performance migration strategy, 3Com followed certain guidelines articulated by network managers in diverse companies and organizations:

- Deploy new technology only in the parts of the network where it is needed.
- Migrate the network in a series of steps at minimal incremental cost.
- Implement new capabilities by building on earlier enhancements with no loss of functionality.
- Maintain seamless connectivity throughout the configuration.
- Simplify the task of managing the network.

3Com assimilated these guidelines into a comprehensive and balanced strategy that is both technically sound and responsive to long-term requirements.

High-Performance Scalable Networking

With 3Com's HPSN strategy, network planners can:

- Build on both current and emerging LAN technologies to meet present and future needs.
- Implement lower-latency connectivity systems in the workgroup and campus backbone for better performance.
- Scale network performance to meet specific business requirements.
- Extend the reach of today's LAN resources by dramatically reducing the cost and complexity of providing remote and WAN connections across the enterprise.

As shown in Figure 1, the HPSN approach applies to all portions of the network—the workgroup, personal office, remote office, building/campus backbone, and WAN. And HPSN builds on the full range of 3Com's network products.

In building/campus backbones, HPSN provides a step-by-step migration to a high-bandwidth environment. This environment can accommodate various high-speed technologies, including Fiber Distributed Data Interface (FDDI), 100-Mbps Fast Ethernet, and Asynchronous Transfer Mode (ATM). HPSN emphasizes innovations that can be added to products 3Com is shipping now.

Among these products are:

- NETBuilder II® bridge/router
- LANplex™ 6000 and LANplex 5000 intelligent switching hubs
- LinkBuilder® 3GH internetworking hub
- LinkBuilder MSH™ multi-services hub
- LinkBuilder FMS™ and FMS II stackable hubs
- LinkBuilder TR Token Ring hub
- LinkBuilder FDDI concentrator

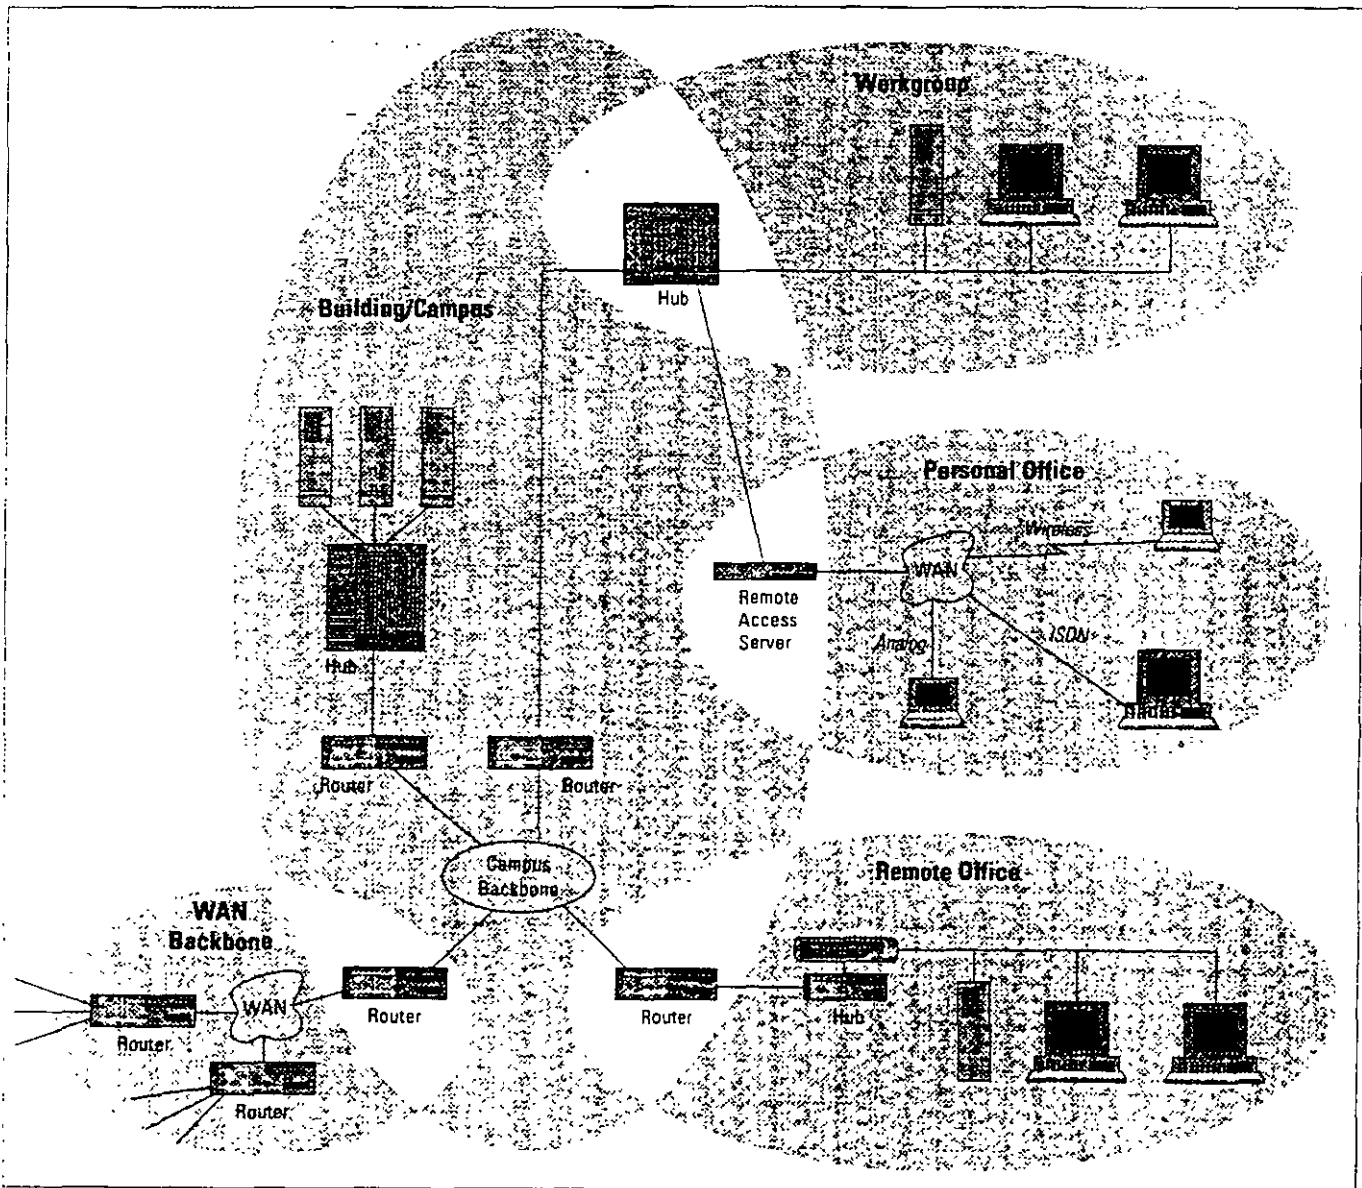


Figure 1 3Com's High Performance Scalable Networking Strategy

Today's Collapsed Backbone Architecture

As part of the evolution from a single LAN per building to separate LANs on every floor, most forward-looking network managers are reconfiguring their distributed networks to collapsed backbones. A collapsed backbone configuration avoids having to put a router or switch on each floor. Instead, each floor's horizontal LAN segments are repeated across a vertical downlink to a single router, which is usually located in the basement along with a group of high-end servers known as a "server farm."

In a typical configuration, the collapsed backbone is a star configuration with network nodes connected by unshielded twisted-pair (UTP) wiring to hubs on each floor. The floor hubs are interconnected vertically through one of the fiber pairs in a bundle often contain-

ing 12 fiber-optic interrepeater link (FOIRL) fiber pairs. The server farm LAN segments are directly attached to the collapsed backbone router in the basement.

This arrangement collapses the network backbone onto the high-speed backplane of a router. (An intelligent switching hub may also be used for this purpose, but for the sake of brevity most examples in this paper will assume that the device is a bridge/router.) In the case of 3Com's NETBuilder II bridge/router, data moves approximately 80 times faster than it would on a distributed Ethernet backbone, and eight times faster than on an FDDI backbone. The hubs in each floor's wiring closet continue to concentrate the LAN floor segments, but networking intelligence and complexity now reside in the basement with the collapsed backbone router.

3Com's HPSN strategy uses the collapsed backbone model for an economical, three-stage migration to scalable performance.




Figure 2. Collapsed Backbones with Multiple LAN Segments on the Third Floor

HPSN Stage 1: Enhancing the Collapsed Backbone with Additional Horizontal and Vertical LANs

As the demand for bandwidth grows, a performance bottleneck can quickly result if all the users on one floor share a single LAN. A collapsed backbone can split users across multiple LANs because each new horizontal LAN segment can extend vertically to a collapsed backbone router port using a separate fiber downlink. This effectively scales the bandwidth of the vertical cabling infrastructure in proportion to the number of horizontal LAN floor segments.

Figure 2 depicts two possible configurations of this type, one using a NETBuilder II bridge/router as the backbone device and one using a LANplex intelligent switching hub. In each configuration, three horizontal LAN segments are deployed on the third floor. This increases floor bandwidth by a factor of three, each with its own downlink to the router.

Overcoming Constraints on Downlinks

Additional vertical downlinks do not usually require new cabling installation, since most buildings have spare fiber pairs in each floor's bundle. But one constraint on the addition of downlinks is the router's port capacity. The manager must make sure that the collapsed backbone device has enough ports to handle both current and near-term future downlinks from the floors. Of course, multiple collapsed backbone routers can always be located together.

Another limiting factor on the number of downlinks is increased complexity. Segmentation improves performance, but it also means there are more LANs to manage. For example, if IP is used as a network layer protocol, each new segment requires its own IP network number, complicating administration and depleting the organization's allotment of IP numbers.

To reduce complexity, the manager could, for example, group the downlinks associated with the three third-floor segments and connect them to a bridge, which in turn would be connected to a port on the router or switching hub. Bridging the three grouped LAN segments in Figure 2 into one logical workgroup requires only a single IP network number, and the router insulates this group of LANs from the others.

However, this solution requires an extra bridge, adding to the delay in the vertical infrastructure. Furthermore, unless the port connecting the external bridge to the router is equivalent in speed to the three downlinks, the bridge traffic may become congested. A further consideration is that adding a high-speed link between the external bridge and router increases cost.

Port Grouping for Virtual LANs and Workgroups

These problems with bridging may be resolved by adding a *port grouping* feature to the collapsed backbone router, providing the bridging function internally between the three downlinks.

Since the "port" to the router is internal, there is no additional delay because its speed is proportional to the speed of the three downlinks. The three grouped LAN segments are referred to as a *virtual LAN*. A *virtual workgroup* is defined by the collection of nodes or end systems attached to the grouped LAN segments.

A virtual LAN that takes up more than one port on the bridge/router looks like a single LAN to the network. But because the administrator can still route traffic between virtual LANs, port grouping retains the advantages of full multiprotocol routing. Furthermore, since port grouping is provided in a router, techniques like proxy address resolution protocols (ARPs) can reduce broadcast/multicast traffic within the virtual LAN.




Figure 3. Example of Virtual Workgroups within a Building

Figure 3 shows an example of multiple virtual workgroups. In this example, the engineering, marketing, and finance groups are kept separate (insulated by a "firewall") to isolate data resources and manage traffic between these virtual LANs. Using port grouping, all of the engineering LAN segments are combined into a single virtual workgroup, even though they are physically divided into three segments spread across two floors. The virtual workgroup is assigned a single IP network number instead of three different numbers.

Creating virtual workgroups using a single IP network number (because they are on the same virtual LAN) mitigates the complexity of segmentation on the floors (because there is no need to change end-system addresses). It also saves management time—for instance, workgroups can be “soft” configured in the NETBuilder II without changes to the physical plant.

Creating Routing Clusters with Switching

When network traffic proliferates between the buildings in a campus environment, the result is an increase in demand for bandwidth on the campus backbone. Network managers often respond by connecting buildings with a single, shared 100-Mbps FDDI network.

As the network grows, however, each building may have hundreds of devices contending for bandwidth on the single, shared interbuilding backbone. Another problem is that users and workgroups often need to be relocated within the campus, while the server they primarily communicate with remains in a central server farm in another building. As a consequence, the campus backbone becomes a traffic bottleneck.

To accommodate growth and alleviate congestion, a switching hub such as 3Com’s LANplex can be used to create *router clusters*—private, high-speed switched links to each building. LANplex devices allow each building to have a private link comprised of either a 10-Mbps Ethernet or 100-Mbps FDDI. Another approach is to use ATM cell switches and the ATM DXI interface on the routers. The ATM DXI interface will accommodate speeds up to 52 Mbps., full duplex. Figure 4 shows three alternative campus backbone configurations that can yield a significant improvement in performance compared to a single FDDI backbone link.

HPSN Stage 2: Creating High-Speed Downlinks To Increase Bandwidth and Reduce Segmentation

Increasing bandwidth within the workgroups means increasing LAN segmentation on the floors. But, eventually, the administrator runs out of spare

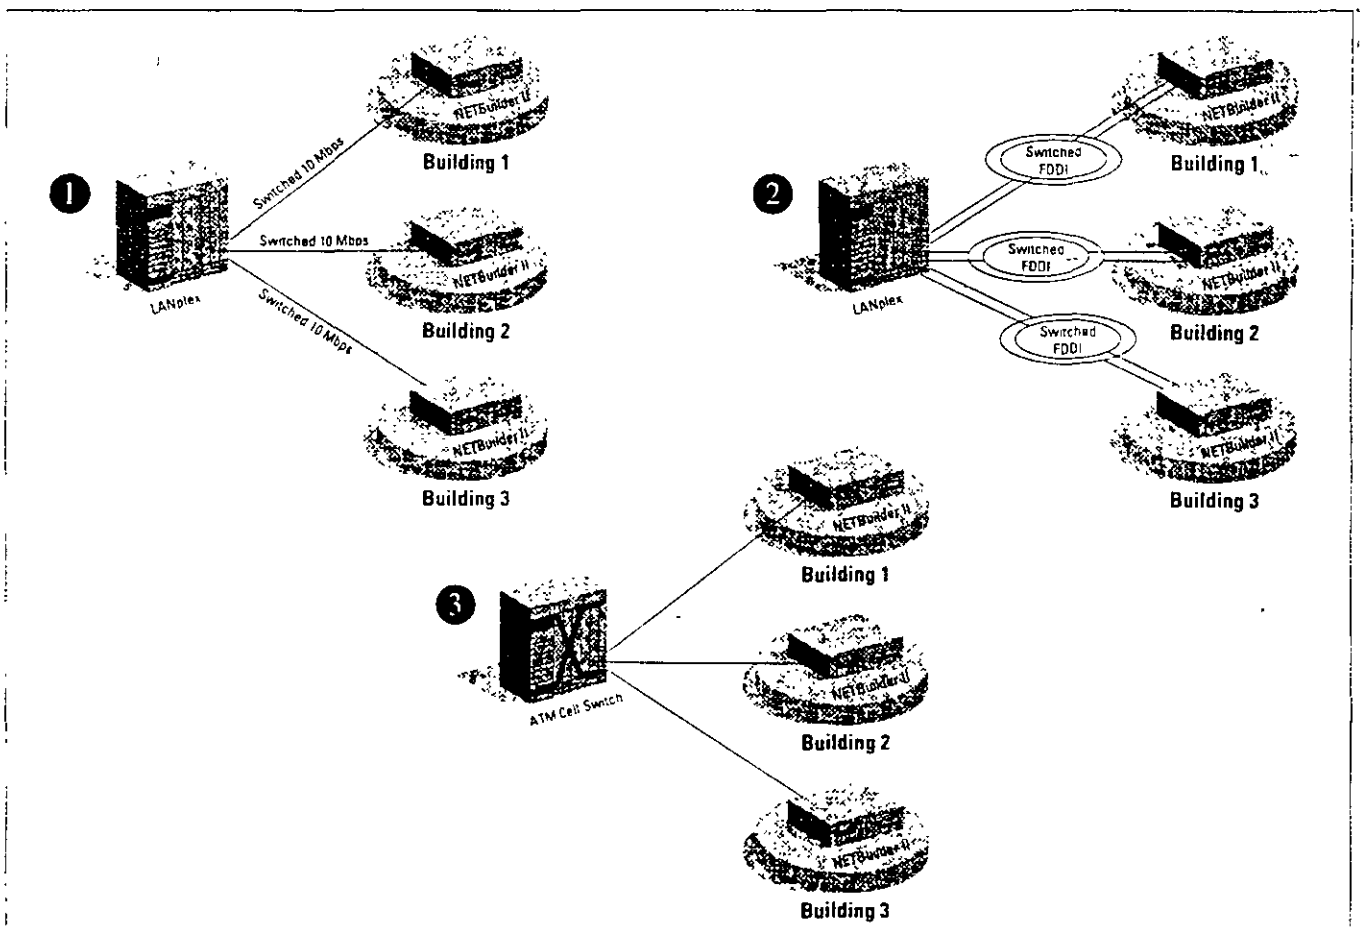


Figure 4 Three Alternatives for Implementing Campus Router Clusters with LANplex Intelligent Switching Hubs

fiber cabling in the building riser, or uses up all the physical ports on the router.

This dilemma may be resolved by using a single high-speed downlink to replace multiple slower LAN segment downlinks. The manager can continue to increase overall network performance by means of additional horizontal floor segmentation without the need to change hardware or software at each desktop.

To support high-speed downlinks, the per-port frame processing performance of the collapsed backbone needs to be increased by approximately a factor of 10. There are numerous ways to accomplish this performance scaling. All of them require distributing some level of routing functionality within the collapsed backbone router.

Route Caching for Scalable Routing

The preferred way to scale performance is to distribute only the simpler high-performance frame forwarding logic to the port interface cards, otherwise known as the *port switching engines*, and to centralize the complex route determination logic in a *central routing engine*.

This is termed an "advise-and-consent" approach because the first time a destination is "seen" by a port switching engine, the central routing engine determines the route and tells the switching engine how to forward subsequent frames with the same destination. The port switching engine thus performs the forwarding operation with the advice and consent of the central routing engine.

The port switching engine remembers the routing information in a *route cache*. Route caching adopts many of the caching principles used to speed memory access in mainframe virtual memory caching schemes, but with one significant advantage: Each switching engine is responsible for routing only the frames from end systems associated with the attached downlinks. Therefore, each port switching engine "sees" only a few routes compared to the total number of routes available to the central switching engine. Also, from the perspective of the port switching engine, the routes the end systems use do not change frequently.

As a result, the switching engine's cache hit ratio (the rate at which references to the faster-access cache memory are successful) over a 24-hour period is likely to be very close to 100 percent. With route caching, the frame-forwarding capacity scales proportionally to the number of high-speed downlinks.

LinkSwitch for Connections Between Segments

High-speed downlink support also requires some level of internetworking on the floor for attaching Ethernet and Token Ring LAN segments. The challenge is to move simple, low-cost internetworking

functionality—a form of LAN switching that 3Com calls LinkSwitch™ technology—to the floor without giving up the full-function benefits of a collapsed backbone router.

LinkSwitch is an extension of the advise-and-consent technique. It behaves like a NETBuilder II port switching engine for its attached LAN segments. Like the port switching engine, if it knows the route, it handles the forwarding of frames independently with the advice and consent of the central routing engine in the collapsed backbone router.

LinkSwitch technology is planned for release as a set of modules in the LinkBuilder MSH in late 1994.

High-Speed Downlinks Using a Switching Hub and Bridge/Router

Figure 5 shows an example of a unique high-speed downlink solution using a LANplex switching hub and a NETBuilder II bridge/router. The LANplex provides FDDI downlinks configured as three FDDI segments, with each segment defining a separate workgroup. Each high-end server in the server farm is attached to the FDDI segment associated with its primary workgroup by means of bridge-per-port or FDDI concentration within the switching hub.




Figure 5. High-Speed Downlinks from Ethernet LANs through a LANplex to a NETBuilder II

The NETBuilder II provides full-function routing between the three FDDI segments and an FDDI campus backbone. The result is three extremely high-speed workgroups within a building with complete firewall protection between them and the campus backbone. Performance is scaled by distributing intraworkgroup frame-forwarding to the LANplex switch, while assigning the much more complex route determination logic to the NETBuilder II.

Today, FDDI is the only standard high-speed LAN technology that can be deployed as a downlink for interconnecting both Ethernet and Token Ring LAN segments. However, in the near future 100-Mbps Fast Ethernet will also be available. Fast Ethernet will support low-cost, 100-Mbps workgroup and downlink connections over existing Ethernet cabling.

ATM as a Downlink Technology

One issue that will eventually have to be resolved with regard to LAN downlinks is the fact that all the segments switched into the downlink must use the same network number, or they must have a full-function router at both ends. Certainly multiple downlinks may be used, but even this approach will present a density problem as the number of LAN segments increases.

The 155-Mbps multimode fiber interface specified by the ATM Forum is an ideal technology for enhancing the speed of downlinks. ATM offers a number of advantages to managers looking for a high-bandwidth alternative that can handle advanced network applications and growth in the future.

Because the identity of a LAN segment can be retained by mapping it to an individual virtual channel, a single ATM downlink can forward frames from LAN segments associated with multiple network numbers. ATM allows for considerably more LAN segmentation without using up fiber cabling pairs. A single ATM link will easily support 15 to 30 Ethernet or 10 to 20 16-Mbps Token Ring LANs on each floor.

If each LAN segment is mapped to a different virtual channel within the downlink, every segment can be identified by the router. NETBuilder II can now perform virtual channel grouping, allowing the administrator to create virtual LANs just as though each segment had its own downlink. The ATM downlink, which uses existing multimode fiber-optic cabling in the building riser, may be implemented by simply adding new modules to the collapsed backbone router, and adding new modules to chassis-based hubs or new units to hub stacks.

LAN downlinks transmit variable-size frames, and consequently have variable delay—known as *latency*. ATM downlinks overcome latency by segmenting frames into short, fixed-length blocks called cells. Data, real-time voice, and video transmissions can all be transmitted together through an ATM pipeline.

Latency can also be minimized at the router once the frames have been segmented into cells. Since all the routing information is normally contained in the first cell, the frame-forwarding decision does not have to wait until all the cells have been received, especially if the destination port is also an ATM interface. 3Com calls this cut-through routing technique *stream routing*. When stream routing is combined with route caching, it is possible to reduce latency in the network significantly.

CellBuilder for ATM Conversion

3Com will support ATM downlinks using CellBuilder™ technology within LinkBuilder MSH and LinkBuilder FMS hubs, converting Ethernet and Token Ring frames into ATM cells for transmission across the ATM downlink. When cells are received from the ATM downlink, 3Com's CellBuilder technology performs the reverse process, reassembling the ATM cells into LAN frames for transmission to the stations.* Figure 6 illustrates this conversion process for an ATM downlink on a building backbone:




Figure 6. CellBuilder Support for LAN Segments on an ATM Downlink in a Building

* CellBuilder uses the ATM Forum standard multimode fiber User Network Interface (UNI) to exchange cells across the ATM downlink. And it uses the International Telecommunication Union ITU-TS standard ATM adaptation AAL5 protocol for ATM segmentation and reassembly.

Integrating Meshed Campus Backbones with ATM Switching

Forward-looking network planners may choose to lay the foundation for ATM downlink and desktop connections by deploying ATM switching in the campus interbuilding backbones as a way to create router clusters. The NETBuilder II bridge/router's multimode fiber ATM module interface supports meshed campus backbones with speeds of 155 Mbps. ATM provides low-latency, high-speed switching between buildings, while NETBuilder II and LANplex switches segment existing LANs.

This campus backbone will have a very high aggregate bandwidth, with data rates measurable in tens of gigabits per second. A bandwidth of this magnitude could serve up to 20,000 Ethernet nodes.

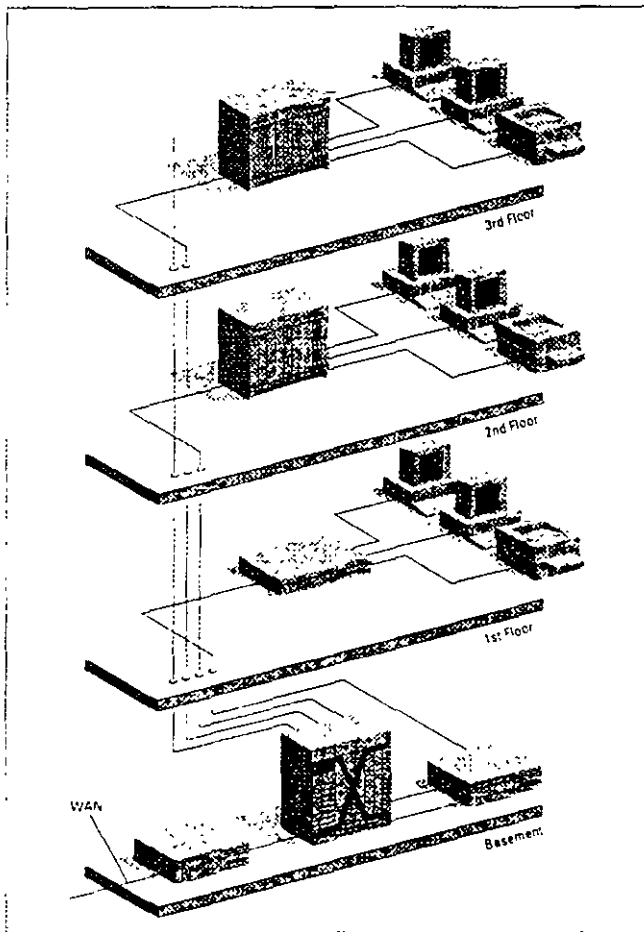


Figure 7. Using ATM Switching to Cross-Connect LAN Segments, and Router Clusters to Build an ATM Campus Backbone

HPSN Stage 3: Enhancing the Collapsed Backbone with Routed ATM

The standards-compliant ATM downlink described in the previous section can be connected to an ATM cell switch, giving the administrator the ability to cross-connect each virtual channel, and therefore each LAN segment, to a specific router or router port. Such an arrangement is shown in Figure 7.

Cross-connecting LAN segments improves performance because the traffic load is shared across routers. Splitting traffic across routers and hubs also has the advantage of protecting the network against component failures. Furthermore, "homing" of segments is possible—for example, the LAN segments making up a virtual workgroup can all be directed, or homed, to a single router. That way, the administrator can change the homing segments without having to alter the physical network.

Route Determination for ATM End Systems

There are two major functional components of an ATM switch. The first is *cell forwarding*, or digital cross-connect logic, which to date has received most of the attention. In general, once a virtual channel has been established, ATM switches have exceptional high-speed, low-latency cell forwarding capabilities.

The other functional component of an ATM switch is *route determination*, which sets up the virtual channels. In current LAN environments, route determination is handled automatically by the routing engine in the collapsed backbone router. This level of automatic functionality is also needed in the ATM environment.

Since end systems use many different protocol stacks—common examples being TCP/IP, IPX, AppleTalk®, and DECnet®—the route determination function needs to be protocol-dependent. The easiest way to accomplish this is by adding ATM to the extensively developed multiprotocol route-determination logic residing in the switching engine of the collapsed backbone router. Standardized methods for route determination are expected to emerge from work being conducted by the ATM Forum on LAN emulation, and by the Internet Engineering Task Force (IETF) on routing over ATM.

Figure 8 illustrates how ATM route determination works. An ATM end system establishes a virtual channel to the collapsed backbone router and transmits its route determination frames (for example, ARP frames for IP) across this virtual channel.




Figure 8. ATM Switching with the ATM End System Functioning as a Switching Engine and the Router Acting as a Routing Engine

If the destination system identified in the route determination frame is attached to the same switch (that is, it has a virtual channel to the router), the routing engine in the router helps set up a virtual channel between the two end systems for direct communication. Otherwise, the routing engine helps set up a virtual channel between the requesting ATM station and the collapsed backbone router, which in turn forwards frames to and from the end system. This end system could be either at the local LAN or at a remote site.

From an architectural perspective, then, ATM switching is an extension of the advise-and-consent technique: A central routing engine works in tandem with a switching engine to optimize network throughput.

Managing the Network

HPSN implements 3Com's Transcend™ family of network applications. The Transcend management architecture goes beyond individual devices to control logical systems made up of all the devices in a particular workgroup, building, campus, remote office, or WAN. The administrator can manage logically connected groups of nodes as a single entity, rather than having to correlate information from hundreds of dispersed devices.

The architecture takes advantage of 3Com's breadth of current products and SmartAgent™ intelligent device agents, while also supporting Simple Network Management Protocol (SNMP) compliant products from other vendors. More intelligence can be added in the future to automate management tasks.

Because an SNMP Management Information Base (MIB) for ATM is being defined by the Internet Engineering Task Force, incorporating ATM in Transcend applications will be straightforward. Transcend applications run on most popular management platforms, including Sun's SunNet™ Manager and Hewlett-Packard's Open View.®

Today's Foundation for HPSN Migration

3Com's HPSN migration strategy involves an evolution of the network, not just the upgrading of individual network components. As each new capability is added to one area of the network, complementary capabilities are already in place in other areas. By orchestrating network change in this way, you can maximize your return on investment for each improvement. 3Com is uniquely positioned to support this migration process because the 3Com product line spans the entire network—from local workgroups, to campus backbones, to remote users.

Equally critical to cost-effective performance migration is the ability to expand the capability of the network by adding new functionality to the devices that have already been installed. As this paper points out, 3Com products shipping today have this expansion capability. These future-proofed products allow the network manager to establish the foundation for performance migration while still protecting network investments, both now and in the years to come.

3Com

Oklahoma Telemedicine Network:

Rural Healthcare Enters a Bold New Era



*A 3Com
Super
Profile*

The Power of a New Medical Technology

Experts in the field note that the image-oriented nature of medicine makes the application of telemedicine the ideal prescription for improving rural health-care and lowering costs. In essence, the network becomes directly involved in the diagnostic process, which has a significant bearing on the cost and timeliness of subsequent treatment.

Telemedicine is also in step with the growing emphasis on preventive care, a component found in most broad health-care reform proposals. By allowing the rural patient's primary care physician to supervise and direct care-giving locally, rather than referring the patient immediately to a regional hospital, telemedicine aids in the early detection of illnesses and helps speed diagnosis. And the convenience of local care encourages patients to visit their local doctor more often for routine check-ups—bringing expert care-givers at the regional hospital into the loop at the earliest diagnostic stages.

Recent results from a telemedicine system in Georgia demonstrate the technology's promise: patient transfers from rural facilities to regional or urban centers were reduced by 85 percent, saving patients both time and money, and more importantly, reducing delays in treatment.

Telemedicine is simply an application of a wide-area networking infrastructure. A basic telemedicine system incorporates image-gathering and image-transmission hardware including wide-area links between remote sites and larger regional centers.

The Oklahoma Telemedicine Network, designed and installed by systems integrator CPI/MicroAge, relies on 3Com's high-performance routers at the larger regional hospitals. Each rural hospital is equipped with a 3Com remote office router for handling transmissions over the network's T-1 lines.

The radiology application configured by CPI/MicroAge at each rural site relies on a Sun® workstation and a Kodak® image scanner that digitizes x-rays and transfers them to a color screen with a 2,000 x 2,500 pixel resolution in about five minutes.

At each rural hospital a 3Com hub provides wiring concentration as part of 3Com's SuperStack™ system architecture for building completely stackable, cost-effective networks. Apple Macintosh® computers at each site carry basic administrative productivity applications, as well as Lotus Notes® software for inter-hospital communications. And 3Com's Transcend™ software provides the necessary centralized management for the entire network.

Evolution of the Oklahoma Telemedicine Network

The Oklahoma Telemedicine Network (OTN) began as a state-funded pilot project at the University of Oklahoma's Health Sciences Center.

In that initial test, six rural hospitals were linked to the University to validate the telemedicine concept. Interest grew—and so did the number of rural hospitals seeking to join the network, bringing the total to 38. In addition, more than 15 regional hospitals are also participating.

The full rollout of the OTN is receiving federal support. Seeing telemedicine as a viable economic development tool for rural communities, the U.S. Department of Commerce allocated \$4.5 million in block grants to fund the first two years of the project. Money that typically was dedicated to building local roads and bridges is committed instead to supporting a new kind of infrastructure investment—a rural information highway.

These grants are supporting the rural hospitals' involvement, while the regional medical centers are participating at their own cost. Within two years, savings from the network are expected to cover the costs of involvement for the rural hospitals—as they gain the ability to treat more patients locally.

The three primary goals of the network's initial application were access, decision support and usefulness, said Gene Hopper, director of the Oklahoma Telemedicine Network.

"Rural providers are isolated, and medicine changes rapidly," she said. "They wanted more support for the decisions they were making about treatment. And rural professionals are trying to provide excellent care with minimal resources. A teleconferencing tool that physicians would use once a week did not meet the rural hospital's primary need."

In short, rural healthcare providers needed access to information wherever it resided—and access on their own terms. “With these fundamental needs in mind,” Hopper said, “radiology was the logical first choice among applications.”

Without the network, rural patients can expect to wait three to five days before their x-ray film can be read by a board-certified radiologist at the regional hospital. The OTN promises to reduce that turnaround time to just 15 minutes in emergency situations and less than an hour in non-life-threatening situations.

But Hopper added that tele-radiology is “just the tip of the iceberg. We’ll look into tele-cardiology, tele-dermatology and many other medical applications,” she said. “We can also use it for education, and access to enormous information resources like pharmaceutical databases and literature services.”

Basic electronic communication via e-mail is viewed as another important advantage, raising the level of communication between colleagues and rural agencies. And in managed-care environments, where cost savings are key to the organization’s success, the network has

potential for increasing administrative efficiencies and helping hospital officials evaluate costs and procedures.

“What we can do with this infrastructure is unlimited,” Hopper said. “Our funding, however, is not—so we started with radiology.”

Building the Oklahoma Telemedicine Network: Initial Project Implementation

In the first phase of the OTN project, 38 smaller rural hospitals will connect to larger, full-service, regional healthcare facilities and to each other via a statewide WAN over T-1 lines. First to go on-line is the radiology-imaging application, which will reduce the time needed for proper evaluation of x-rays from several days to about 15 minutes.

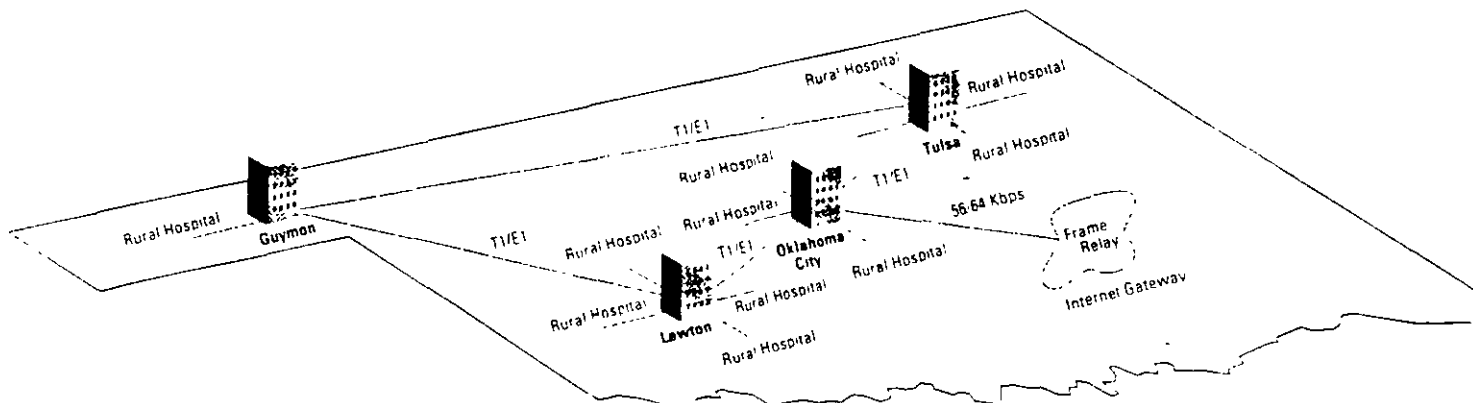
Each rural site has a 3Com router, 3Com hub, a Sun workstation, Kodak image scanner and Macintosh computers. The 3Com solutions are part of the innovative SuperStack system of completely stackable remote-site and departmental networks. SuperStack system solutions include wiring concentration, bridging, routing, LAN switching, redundant power

supplies and SDLC conversion—all designed to be economical and easy to install and manage. And 3Com’s Boundary Routing[®] system software centralizes complexity in 3Com NETBuilder II routers at hub sites where support is provided. So access can be extended to up to 10 times as many sites as with traditional routers—without any added administrative demands.

Building the Oklahoma Telemedicine Network: Future Plans

Additional applications are expected to be made available as users become familiar with the system. As Hopper noted, nearly any image-based diagnostic procedure can be carried over the network. At the same time, with traditional barriers removed, the network should introduce a new level of competition for certain medical services.

“Now geographics are moot,” noted David Blankenship, project manager for CPI/MicroAge. “Anyone in Oklahoma can provide the x-ray reading service. So it’s going to foster commerce while lowering costs—and it’s going to change the way hospitals operate.”



In this environment, the rural hospital could drive expansion of the network just as much as the regional hospitals. Blankenship predicts.

"Regional hospitals realize there are a lot of services they could provide via this infrastructure," he said. "For instance, after-hours cardiac monitoring to a 40-bed rural hospital that can't afford to staff the facility."

"We anticipate the regional hospitals being creative with new applications," he added. "Then the OTN will mushroom at an exponential rate—in sites and applications."

Exploring the Possibilities

In the future, the OTN infrastructure will provide easier, faster access to vital patient records, no matter which network facility the patient chooses to visit. This application can reduce paperwork, thus allowing healthcare professionals to devote more time to treatment, rather than chasing records. Individual physicians' offices will also be able to link to the network, using solutions such as 3Com's AccessBuilder™ for dialing in to gain easy network access.

Other partners involved in the OTN project include: Access Radiology, Apple Computer, AT&T, Kodak Health Imaging Systems and Lotus Development.

The Evolution of Telemedicine

The first hurdle in the path of telemedicine had nothing to do with technology, notes Dave Swartz, a computer engineer and seasoned telemedicine consultant. Healthcare providers themselves were almost universally skeptical.

"Historically, telemedicine was based on synchronous protocols—video or voice," he said. "Two-way video was the norm."

But live consultations between doctors at two video conference-capable locations usually proved expensive and inefficient. The cost for the bandwidth was high, and the systems made poor use of doctors' time.

"Did you ever try getting a doctor on the phone?" Swartz said. "Try getting two doctors on the TV monitors."

So, while deploying fiber throughout the University of Oklahoma medical campus, Swartz sought out medical services and technologies that could benefit from a high-speed, wide-area network implementation—imaging.

The images could be stored and forwarded for reading at the radiologist's convenience—and at low cost. In this case, the network truly created new possibilities instead of imposing limitations. Other image-based applications also make sense—cardiology, ophthalmology, pathology, dermatology, ultrasound—any case where a computer could capture video images or even audio.

Swartz also noted that the infrastructure established for these medical services can do double-duty providing access to the Internet, on-line medical libraries and countless other educational resources.

For telemedicine networking hardware in Oklahoma, where he has consulted on the OTN's development, Swartz has recommended 3Com solutions.

"3Com is playing a pivotal role in making sure we can install this network today—and then upgrade in the future without throwing everything out," he said.

The future of multimedia is tied to Asynchronous Transfer Mode (ATM), he explained, so it's important to ensure a smooth upgrade to ATM as standards become firm.

"With the 3Com router, all you'll have to do is load up an ATM card," he said.

ATM makes interactive desktop video a real possibility and can expand the potential of telemedicine networks exponentially. 3Com's ATM leadership and expanded ATM product line will provide its customers with a ready source of solutions for making those improvements a reality.

"Our goal is to build systems of the future," said Swartz. "In Oklahoma, they're well on their way to doing that."

Telemedicine network promises better and faster care for less money

In recent years, landmark developments in medical technology have changed the face of healthcare across the United States. The effects of many of these breakthroughs however, have had relatively little impact in many rural communities, which are too far removed from urban centers where the most sophisticated technologies are supported by larger populations and greater economies of scale.

A cornerstone of most reform initiatives today is expanded access—putting more people in touch with better care. Today, in rural America, the State of Oklahoma, CPI/MicroAge, 3Com and other solutions suppliers are joining forces to bring that goal within reach—through an application of wide-area networking technology known as telemedicine.

Telemedicine allows seasoned medical specialists to pay “electronic house calls” to even the smallest rural hospitals and clinics. Network links can send x-rays and other clinical images great distances in a matter of minutes, allowing rural patients and their doctors to tap medical resources and experts once reachable only through lengthy travel.

The Oklahoma Telemedicine Network (OTN)—believed to be the largest in the nation, has continued to be a trailblazer in this rapidly-evolving network application. Developed under the auspices of the Health Sciences Center at the University of Oklahoma, the OTN is moving healthcare reform forward today in the state’s rural areas.

And at the heart of this network are proven, practical remote-site networking solutions from 3Com.

Representing a revolution in the delivery of rural healthcare, the OTN promises to accomplish many of the primary objectives for overall healthcare reform:

- **Improved care, as rural residents gain access to experienced specialists hundreds of miles away, without leaving their home communities.**
- **Lower costs for treatment.**
- **Less time between diagnosis and delivery of treatment.**
- **Expanded opportunities for training and continuing education of rural medical professionals.**



El Grupo Financiero Probusa (GFP), es un Caso de Exito Internacional para Cisco Systems de México, dado que es la primera Institución Financiera en integrar Tecnología ATM en su operación de Redes.

Con el fin de dar a conocer en el campo práctico las ventajas de las aplicaciones de la Tecnología ATM, entrevistamos al Ingeniero Jorge Macías, Subdirector de Telecomunicaciones y al Ingeniero Fernando Krasovsky, Gerente de Implantación de Proyectos del GFP

Los antecedentes de equipo Cisco en el GFP datan desde principios de 1993, cuando este grupo financiero desarrolló la estrategia de consolidar en un solo centro de cómputo la operación de la totalidad de sus sucursales y oficinas operativas, lo que representó un alto grado de complejidad en términos de la integración y conectividad de distintas plataformas y redes. A través de ruteadores Cisco del tipo AGS se hizo posible resolver estos problemas, además de mejorar los niveles de disponibilidad mediante el uso de medios redundantes en puntos críticos, aprovechando las características de comunicación atómica propias de esta tecnología.

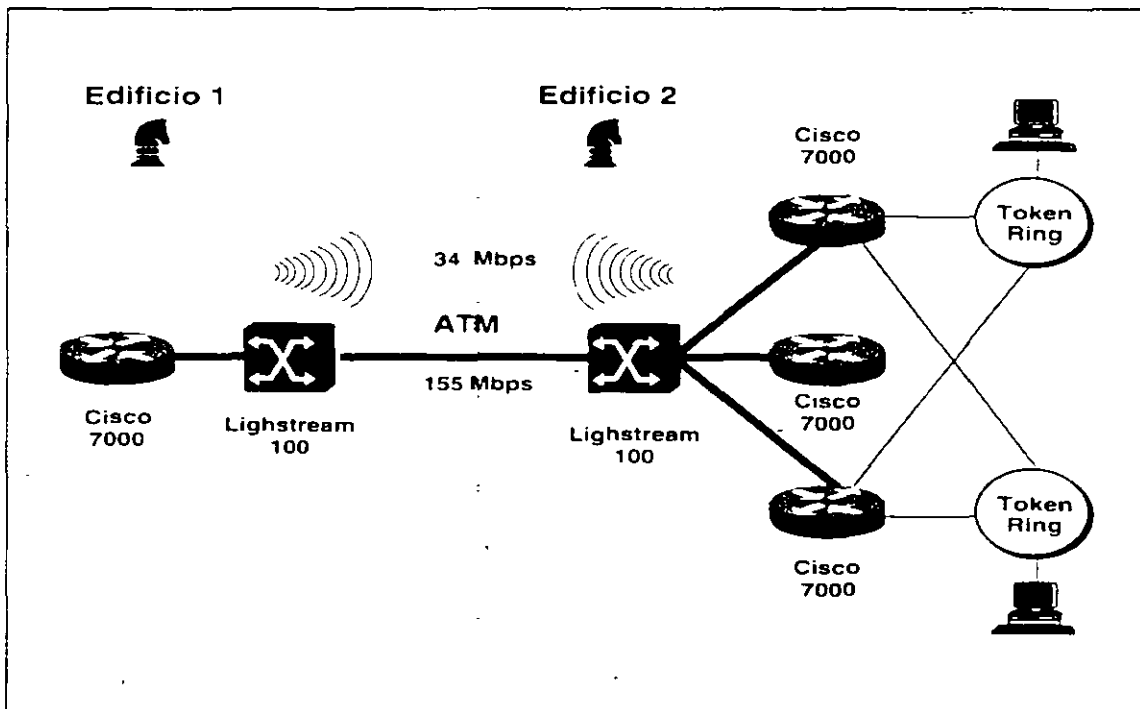
La red de Probusa, es decir, la red principal o backbone, tiene comunicación

entre sus sucursales y grandes plazas por medio de RDI y satélite. A su vez, cuentan con rutas de respaldo por cada uno de sus enlaces, y están basadas por completo en Ruteadores Cisco. La finalidad de la Tecnología ATM en la red, es la de comunicar a altas velocidades (155 Mbps) los dos edificios corporativos del Grupo Financiero: el Montes Urales I y el nuevo edificio Montes Urales II, a través de un backbone. Ambos corporativos tienen comunicación ATM por medio de dos Switches A100 Cisco y Ruteadores Cisco 7000. Como todas las aplicaciones se comunican bajo ATM, y la pérdida de los enlaces puede ser bastante crítica, se ha colocado una trayectoria principal de fibra óptica con un respaldo por micro-ondas, sobre el mismo concepto de ATM.

Por su parte el esquema que tienen funcionando los Ruteadores 7000 hacia los Switches ATM, está establecido por la tarjeta ATM Interface Processor. El computador central es un 9121, con sistemas AS/400 y comunicación con el computador Tandem de la Bolsa Mexicana de Valores (BMV)

Primer esquema de funcionamiento

El esquema que se planteó fue de dos anillos por piso, cada uno con dos Token Rings de diferentes Ruteadores 7000. Cada anillo cuenta, además, con una llegada por fibra óptica y la otra por cobre. Esto muestra



que se cuenta con un esquema de redundancia completo de ruteadores, fibra, cobre y anillos, en caso de que se perdiera cualquier ruta.

A fin de disponer de una redundancia completa en todos los anillos Token Ring, los ruteadores dividen sus cargas para todo el tráfico de datos. Esta facilidad la da el Stand by Router, para que en el caso de que haya algún problema con uno de los ruteadores, él o los otros, puedan soportar la comunicación de todo el edificio. Podemos resumir que el Ruteador que soporta las dos rutas, automáticamente switchea de uno al otro, sin perder la información que manejan los usuarios, en el caso de que falle una de ellas

Esto es de gran importancia para el GFP, ya que cuenta con un número aproximado de tres mil LU's conectadas en red, hacia el procesador central 9121. Las LU's no son exactamente usuarios. La diferencia radica en que cada usuario utiliza de dos a tres sesiones para tener diferentes aplicaciones en forma simultánea. Podemos agregar que las LU's son unidades direccionables que permiten a los usuarios finales comunicarse entre sí y tener acceso a los recursos de la Red SNA.

Hoy en día, todo el esquema de comunicación hacia el edificio Montes Urales II - cuyo objetivo es el de integrar a las diferentes empresas del GFP en un solo edificio -, es a través de RDI con respaldo en microondas y/o satélite en las diferentes sucursales

Por qué ATM con Cisco

La razón por la cual Proburza está innovando una arquitectura ATM, se debe a su búsqueda por la mejor tecnología para el Grupo Financiero. ATM representa para GFP la mejor inversión en cuanto a tecnología y rentabilidad. Aunque anteriormente se analizó una propuesta para instalar Frame Relay, se convino en que ATM está más a la vanguardia que la tecnología anterior

Sin tener que estar haciendo cambios en aplicaciones y considerando la infraestructura existente, en Proburza simplemente se incorporaron los equipos ATM a producción.

Aunque todavía no se cuenta con la aplicación liberada en video, en Proburza se está consciente de la necesidad de contar en un futuro con aplicaciones que puedan manejar video, voz y datos de manera conjunta. Por lo tanto el GFP ha preparado la infraestructura necesaria para darle soporte a sus clientes tanto externos como internos. Con la aplicación de imagen, se



Los Ings Jorge Macías H. y Fernando Krasovskiy S. del Grupo PROBURSA

está planeando manejar videokonferencias con los clientes del GFP para juntas de trabajo, capacitación, asesorías financieras, etc

Actualmente las aplicaciones del GFP son cliente-servidor

Todo el ambiente de su red de área amplia está migrando al esquema TCP/IP y se están eliminando todos protocolos no ruteables en la red. El sistema operativo que se maneja, está en dos plataformas Windows NT y OS/DOS

Las aplicaciones que el GFP maneja hacia el computador central 9121, el AS 400 y la Tandem de la Bolsa Mexicana de Valores, las accesa desde el nuevo edificio a través del backbone de ATM. Las ventajas más significativas que proporciona ATM en el backbone son: la velocidad y la confiabilidad

Al principio el GFP probó varias marcas de ruteadores. Sin embargo Cisco Systems de México apareció como la mejor opción, por todas las facilidades con que cuenta su equipo y el soporte técnico que les brinda en conjunto con Red Uno. Otra de las ventajas de la tecnología Cisco es su escalabilidad, el acceso a nuevas versiones y el sistema de monitoreo centralizado con que cuenta a través del software Cisco Works. Éste ha resultado la herramienta ideal para administrar, controlar y configurar toda la red de datos en su seguridad, porque proporciona diferentes parámetros cuando surge algún problema, permitiendo solucionarlo dónde y cuando se presente.

El plan final del GFP es el tener en cada sucursal nueva un Ruteador y una Red. En un futuro cercano se planea que su integración total sea a través de ATM

Por último, puntualizaron que quien tenga el mayor número de servicios, será el líder del mercado financiero. ▲▲



ATM Protocolos

☐ Nivel Físico

☞ " Physical Medium-Dependent " (PMD)

☞ " Transmission Convergence " (TC)

Notas:



ATM Protocolos

▣ Sub nivel PMD

- ☞ Transferencia de bits
- ☞ Alineación de bits
- ☞ Codificación en línea
- ☞ Conversión Electro-óptica

Notas:



ATM Protocolos

▣ Subnivel TC

- ✓ Generación, transmisión y recuperación del Frame
- ✓ Adaptación de Frame
- ✓ Delineación de la celda
- ✓ Cabecera de control de error (HEC)
- ✓ Desacoplamiento de la taza de celdas

Notas:



ATM Protocolos

☐ Nivel ATM

Las funciones de transporte del nivel ATM son independientes del nivel físico, razón por la cual se requieren identificadores de conexión, los cuales indican los canales virtuales y la ruta virtual que debe seguir cada celda.

Notas:



ATM Protocolo

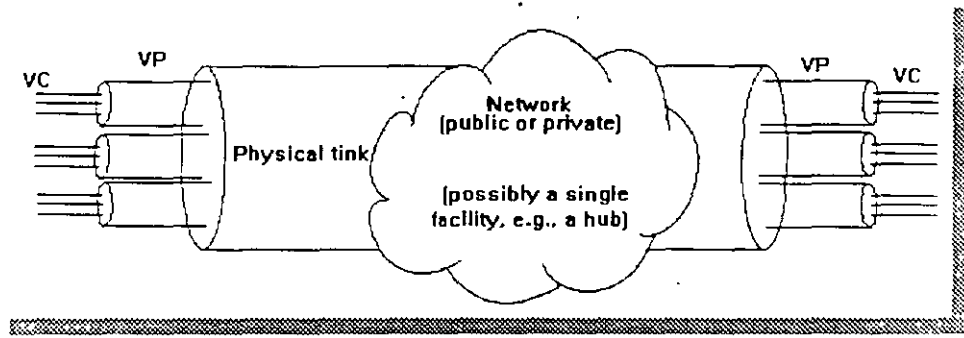
▣ VC

Un canal virtual es el enlace que esté entre 2 puntos, está conformado por un VCI (Virtual Channel Identifier) y un VPI (Virtual Path Identifier), donde el VCI determina enlace particular VC para una determinada ruta virtual identificada con un VPI.

Notas:



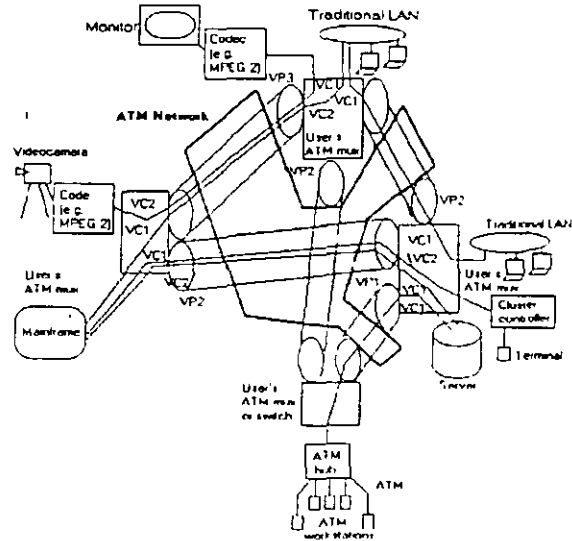
ATM Canales y Rutas Virtuales



Notas:



ATM Canales
y
Rutas Virtuales



Notas:



ATM Protocolos

☐ Funciones del nivel ATM

- ☞ Multiplexaje y demultiplexación de celdas
- ☞ Traducción de VCIs y VPIs
- ☞ Generación y extracción de la cabecera de la celda

Notas:



ATM Protocolos

☐ Ruteo de las Celdas

Dentro de la UNI existen 24 bits destinados al ruteo de la celda:

☞ 8 bits para asignar un VPI

☞ 16 bits para asignar un VCI

3 más se destinan para la descripción del contenido de la zona de carga, esto es se determina si el contenido es información de la red o del usuario, y 8 más se destinan para el HEC.

Notas:



ATM Protocolos

☐ Nivel de adaptación ATM

También llamado " AAL ", es el encargado de brindar el soporte necesario a los niveles superiores, es decir de convertir los datos de niveles superiores en bloques de información para las celdas.

Notas:



ATM Protocolos

☐ Nivel AAL

Es requerido por los 3 planos, el de control, el de administración y el del usuario aunque las funciones que realiza dentro de cada una de ellos son diferentes.

Notas:



ATM Protocolos

☐ Nivel AAL

Sus funciones están organizadas en dos subniveles

☞ Convergencia (CS)

☞ Segmentación y reensamblaje (SAR)

Notas:



ATM Protocolos

☐ Nivel AAL

Tiene 5 clases de aplicación o servicio:

- ☞ Clase A: orientado a conexión, CBR
- ☞ Clase B: no definido
- ☞ Clase C: orientado a conexión, VBR
- ☞ Clase D: orientado a No-conexión
- ☞ Clase X: orientado a conexión CBR o VBR definidos por el usuario

Notas:



ATM Protocolos

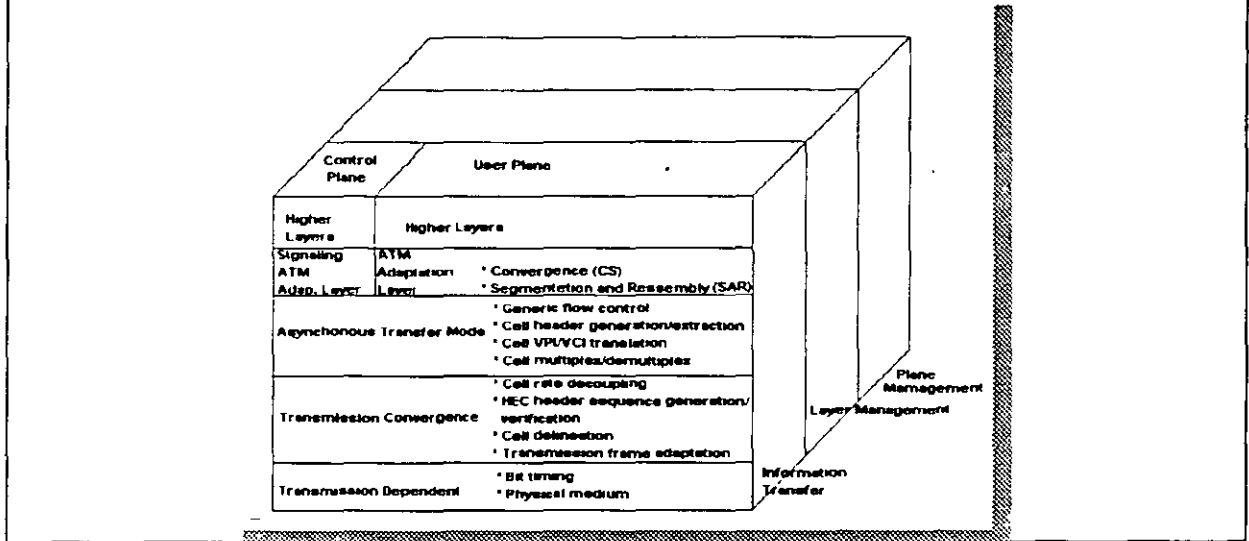
☐ Nivel AAL

Según su aplicación se definen los siguientes tipos en el plano del usuario:

- ☑ AAL tipo 1 para clase A
- ☑ AAL tipo 3/4 para clase D
- ☑ AAL tipo 5 para clase X

Notas:

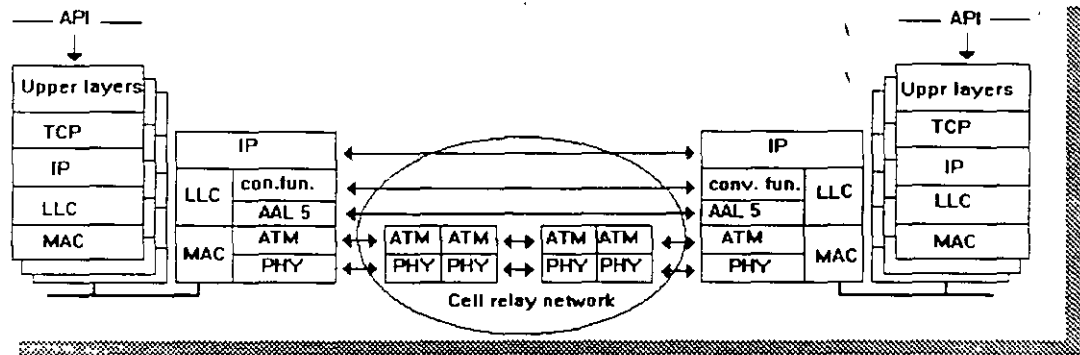
TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



Notas:



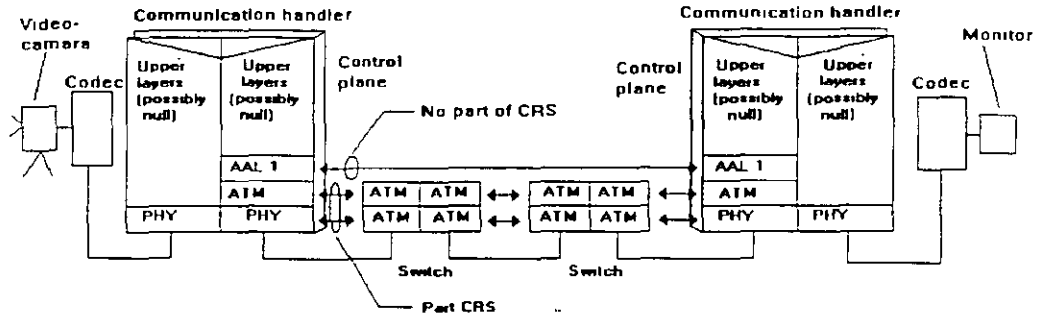
ATM Ejemplo



Notas:



ATM Ejemplo



Notas:

REDES DE ALTO DESEMPEÑO

5.- FRAME RELAY



Frame Relay Technology and Services

Frame relay¹ concepts, technologies, and services are receiving considerable attention in the vendor, carrier, and user communities. The combined frame relay and SMDS market is expected to exceed \$1.2 billion by 1995, making these important new fields. This chapter provides a detailed view of many of the underlying frame relay principles, issues, and concepts.² Some basic frame relay concepts were described in Chapter 1, which the reader may want to review at this juncture. Frame relay can be deployed in a private network, or, as already alluded to at the end of the preceding chapter, the service can be obtained from a public network. Each approach has advantages and disadvantages within the framework of a corporate enterprisewide network. In the following, the term "frame relay" refers generically to either the service or the supporting technology, depending on the context: usually "frame relay service" refers to a public carrier service, while "frame relay technology" implies platforms for private network solutions.

This chapter aims at answering questions such as When does frame relay make sense? Should the user deploy a private frame relay network or employ carriers' services? When is PVC service adequate and when is SVC desirable? Is LAN traffic leaving a high-throughput router really bursty? When is it better to use other technologies? Basic frame relay concepts are introduced and the inter-relationship with cell relay is discussed (Section 11.2). Benefits of frame relay in both private and public networks are identified (Section 11.3). Frame relay standards are surveyed (Section 11.4). Steps for deployment of the technology in corporate networks are described (Section 11.5). Carrier services and equipment availability are surveyed (Sections 11.6, 11.7, and 11.8).

¹A frame in this context is a data link layer construct. The "frames" discussed in Chapters 2 through 9 were physical layer constructs. There is no relationship between the two concepts. The term "relaying" is used by CCITT.

²Some portions of this chapter are based on AT&T product literature [11.1]

11.1 BACKGROUND

Frame relay is a recently introduced multiplexed data networking capability supporting connectivity between user equipment (routers and nodal processors/fast packet switches) and between user equipment and carriers' frame relay network equipment (i.e., switches). The frame relay protocol only supports data transmission over a connection-oriented path; it enables the transmission of variable-length data units over an assigned virtual connection.³ Compared to traditional packet-switched services, frame relay can reduce network delays, provide more efficient bandwidth utilization, and decrease communication equipment cost. Traditional packet services typically introduce a 200-ms network delay or more (40 to 60 ms per hop to handle error correction and control on a hop-by-hop basis), whereas with frame relay that network delay can be reduced to about 20 to 40 ms [11.2].

As is the case in X.25, frame relay standards specify the user interface to a device or network supporting the service. Namely, it specifies the UNI. This interface is called frame relay interface (FRI). A FRI supports access speeds of 56 kbps, N x 64 kbps, and 1.544 Mbps (2.048 Mbps in Europe) [11.3–11.5]. Some vendors are attempting to extend the speed to 45 Mbps. The service can be deployed (1) in a point-to-point link fashion between two routers, (2) using customer-owned frame relay nodal processors (frame relay switches which employ cell relay on the trunk/NNI side⁴), and (3) using a carrier-provided service. Table 11.1 provides a summary of key features of a frame relay network [11.6]. Figure 11.1 depicts the technology at the logical level.

11.1.1 Usage of the Technology—An Overview

Frame relay functions supporting the FRI need to be added to user equipment such as routers, T1 multiplexers, FEPs, PADs and so on, in order to be connected to a private or public frame relay network. At this time, frame relay technology is being applied mostly to LAN interconnection environments. Existing routers are easily upgraded to support the FRI. If commercially successful, public frame relay service may supplant X.25 service in the U.S. as a public switched service by the mid-1990s. The frame relay market was expected to reach \$210 million by 1993 and \$850 million by 1995 [11.7]. For comparison, the SMDS market is expected to reach \$500 million by 1995, and BISDN (for data applications) would reach \$500 million by 1997 [11.8].

In some cases, asynchronous terminals may also be supported by a frame relay network; synchronous SNA terminals or other devices (e.g., a front-end

³Some early trials of N x 64-kbps video were undertaken in 1992.

⁴Instead of using cell relay on the trunk side, some switches use frame relay instead. In the long run, most switches will probably migrate to the cell relay NNI.

Table 11.1
Summary of Key Features of a Frame Relay Network

- Standardized by CCITT and ANSI standards
- Only "core" functions are provided by the network
- Network does not guarantee delivery of data
- Protocols in user's equipment is responsible for retransmitting data that is lost, misrouted, or discarded by the network because of congestion
- Frames are transported transparently (only label, congestion bits, and frame check sequence are modified by network)
- Network detects (but does not correct) transmission, format, and operational errors
- Network does not acknowledge or retransmit frames
- Delivers frames in sequence

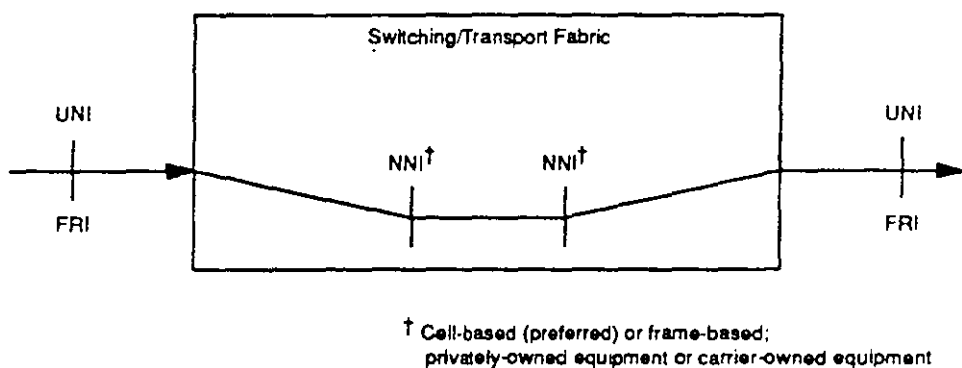


Figure 11.1 A logical view of frame relay communication.

processor) can also use frame relay. In these cases, an appropriate frame relay handler (similar to a PAD but only supporting a layer 2 FRI to the network) is required. In terms of quality of service, the delay incurred by the needed protocol conversion to support these devices (which is similar to the delay through a PAD in an X.25 environment) is not eliminated by frame relay technology, except possibly to reduce queueing time toward the switch, given lower network congestion. The delay through such a frame relay handler is determined by the access speed of the user's line, plus the handler's processing time. For example, if the user's frame from a synchronous terminal contained 262 octets and the access line was 9,600 bps, the initial frame relay "framing" delay would be 219 ms; the "deframing" delay would also be 219 ms. If the access speed is 56 kbps, the figure would be 37 ms. This framing/deframing delay is in addition to the frame relay network delay.⁵

⁵In a cell-based switching/transport platform, as soon as a frame starts to arrive at a switch, it is immediately reduced to cells "on-the-fly."

From the user perspective, any delay is important regardless of whether the delay is generated by the frame relay network or by the access apparatus to the frame relay network (whether in a PC, a terminal adapter, or a PAD-like device). This last issue was never properly appreciated by the packet equipment vendors, and is the reason why packet technology did not see much penetration in the mission-critical synchronous networks of the 1980s.

In order to get maximum benefit from frame relay without having to incur large communication or equipment charges (i.e., for dedicated T1 links between sites, or for the deployment of user-owned frame relay nodal processors), the service needs to be provided by a carrier. The seven BOCs and several value-added carriers have announced frame relay services in the U.S. Carrier networks based on frame relay provide communications at up to 1.544 Mbps (in the U.S.), shared bandwidth on demand, and multiple user sessions over a single access line. The use of a router equipped with a frame relay interface over a dedicated end-to-end T1 link is not economically advantageous compared to a non-frame relay solution, and, in fact, may affect response time. A carrier, on the other hand, can multiplex the traffic of one user with that of other users, and can therefore pass back to the users the economic advantages of bandwidth sharing, much the same way X.25 or FT1 provides such economic efficiencies. Without carriers or private switches, dedicated T1 links between two sites to be interconnected are needed, regardless of the protocol used over the link. However, it should be noted that when using an IXC or VAN service, the user needs a dedicated T1 or 56-kbps link to the IXC's or VAN's POP. If the LEC serving the user's location provides the service, the dedicated T1 is required only to the serving CO.

A second way to benefit from frame relay is to use it in conjunction with a frame relay nodal processor (variously known as "corporate backbone network switch," "fast-packet backbone multiplexer," or "frame relay handler"), first described in Chapter 6. Some corporate networks already deployed this technology in the late 1980s in the form of fast packet multiplexers. In this case, the service can often be cost-effective, since the user can obtain from the backbone bandwidth on demand, rather than on a preallocated (and inefficient) basis. The "saved" bandwidth is then available to other users of the same backbone, in theory minimizing the amount of new raw bandwidth the corporation needs to acquire from a carrier in the form of additional T1 or FT1 links. Without a nodal processor using cell relay principles, dynamic bandwidth allocation is not easily achievable. Some processors use frame relay on the NNI; fine-grain multiplexing is more difficult, particularly in mixed-media and multimedia applications (next-generation LANs use cell principles to support multimedia, as discussed in the previous chapter). In the private network application, the user leases from a carrier private lines between the remote devices and the nodal processors, and between the nodal processors; the user employs frame relay to statistically multiplex traffic in a stan-

standardized way, in order to achieve better utilization of the (now common) transmission resources. The nodal processors must be housed in selected user locations.

Frame relay supports bursty traffic at medium speeds. Consequently, many of the applications now advanced by vendors for frame relay services are feasible more in view of the increased throughput and reduced network latency compared to X.25 networks than any other new intrinsic feature of frame relay.

11.1.2 Some Motivations for Frame Relay Services

As indicated, the current major application of frame relay is for LAN interconnection. A combination of recent trends has forced LAN managers to investigate new approaches in order to provide connectivity at reasonable price:

- In 1989, only two out of ten corporate terminals, PCs, and workstations were connected to a LAN. By 1993 eight out of ten terminals will require to be connected to a LAN [11.9].
- Business shifts to accommodate the "Information Age" have forced companies to rely in greater measure on data collection, processing, and distribution. For many such companies, the ratio of terminals to employees is approaching 1, and in some cases it is even exceeds 1. Chapter 1 already discussed some of these trends.
- In an effort to be more efficient, find a cheaper work force and reach global markets, companies are moving toward distributed data processing. Employees are distributed in smaller work groups located around town or around the country, closer to the resources and to the markets. Yet, more than ever, they need to be tightly connected through a reliable and easy-to-use enterprisewide network.

These trends have led to the following sequence of events:

1. Introduction of more terminals.
2. Introduction of more LANs to support the increased number of terminals.
3. Interconnection of these LANs, including those that have emerged in geographically dispersed buildings around the country. In 1991, an estimated 65% of the top 5,000 U.S. companies had LAN interconnection needs extending beyond a building; that number is expected to grow to 85% by 1995 [11.10, 11.11].

The traditional WAN approach of connecting a few LANs with routers over dedicated point-to-point lines is no longer adequate in an environment of many remote LANs. Some actual networks can have as many as 600 or 800 routers [11.12]. LAN managers have sought solutions that reduce the number of dedicated lines in order to keep transmission costs down, and at the same time increase

flexibility and make network management easy. It almost sounds like the perfect WAN environment for packet switching technology, but not exactly. Packet switching allows users to be easily added and interconnected while following open international standards. But packet switching has traditionally been slow and the throughput has been limited.⁶ These limitations are becoming more accentuated with the new LAN applications, which may involve graphics, multimedia, desktop publishing, bulk file transfer, and other data-intensive requirements. This has led to the development of two solutions specifically aimed at LANs: frame relay and SMDS (SMDS is discussed in Chapter 12).

Frame relay is an early 1990s technology allowing users of multiple routers to connect them in an effective manner. It follows open standards and improves throughput, while at the same time reducing the end-to-end delay (throughput can be increased up to three orders of magnitude). When used in a private network employing a cell relay platform, it allows the LAN manager to rapidly meet the evolving high-speed LAN interconnection needs of the corporation, and do so in a cost-effective manner. The same can be said when using a public network frame relay service.

As discussed in Chapter 1, the business trend is toward interconnection of all company resources into a seamless enterprisewide network. However, such interconnection can become prohibitively expensive, unless it is done correctly. Corporations also see the emergence of new LAN applications in the 1990s, which must be supported by the enterprisewide network. New high-bandwidth applications dictate the introduction of new high-capacity digital services and technologies in the corporate network. See Table 11.2, compiled from a variety of sources, including [11.10, 11.11]. High end-to-end throughput, low latency, cost-effective bandwidth on demand, and any-to-any connectivity are the order of the day. A major evolution in the way corporations connect their computers and the ever-ubiquitous PC is already evident in progressive companies. Now, new equipment and new communications services allow corporations to redesign their networks and save money, while at the same time increase their capabilities and work force productivity. The key to achieving these communication goals in a private network environment is frame relay over a cell relay platform, or a high-capacity public switched service like frame relay (and SMDS).

The new technologies needed to support the evolving corporate environment, however, cannot be introduced in a vacuum. It would be easy to deploy an optimal state-of-the-art network when the LAN manager could throw away everything and start completely from scratch. In times of cost-containment, though, the LAN

⁶Packet switching and other statistical multiplexing schemes do fulfill the role of supporting error-free transmission of asynchronous traffic from "dumb" terminals (or devices and PCs emulating "dumb" terminals), which have no error protection of their own—a frame relay service would be a mismatch in this environment.

Table 11.2
Interconnection Needs of LANs

	1991	1995
LANs per typical company	11	22
Growth in LANs	—	100%
Percentage requiring interconnection outside a building	65%	85%
Growth in interconnected LANs	—	160%
Link speeds		
• ≤ 9.6 kbps	25%	7%
• $9.6 \leq \text{link} \leq \text{DS0}$	30%	23%
• $\text{DS0} \leq \text{link} \leq \text{DS1}$	21%	25%
• $\text{DS1} < \text{link} \leq \text{DS3}$	20%	30%
• $> \text{DS3}$	4%	15%

manager is being asked to improve the efficiency of the interconnection network in a judicious and effective manner. The transition steps to the new communication environment depend on the network currently in place in the company. Four generations of corporate networks have been deployed in the past decade. Some users moved from generation to generation; other users leapfrogged one or two generations. Others were forced to retain a network until the payback could be achieved and then some. Usually a network stays in place for a period corresponding to the useful life of the equipment, which can be 5 to 8 years. A short review of the four generations of corporate networking assists discussion of how frame relay evolved and what problems it solves; it can also establish the point of departure for the transition that the LAN manager needs to undertake in order to implement frame relay.

First-Generation Corporate Networks. This phase saw the introduction of unintegrated nationwide networks, which typically employed low-speed analog lines to support discrete mission-critical corporate functions. LANs were just being introduced in companies. Connectivity among LANs, for the few progressive companies attempting it at that time, used its own point-to-point transmission facilities. Different departmental data applications (e.g., a mainframe payroll application and a minicomputer supporting marketing) used separate networks. Not only was this solution expensive because of the duplicate transmission costs, but it was also difficult to manage and to grow. A number of mission-critical networks in place today still conform to this architecture. These companies have found that until now a backbone network was not cost-effective.

Second-Generation Corporate Networks. This phase saw the introduction of T1 multiplexers and supporting digital transmission facilities. The data applications were aggregated over a common backbone network, improving network manage-

ment, simplifying the topology and reducing the communications cost. A few traditional analog lines are enough to justify the cost of a high-speed digital link, as discussed elsewhere in this book, making this transition a popular upgrade in the recent past. Voice traffic was also carried by the backbone. One of the shortcomings of this approach, however, is that the LAN interconnection traffic, now growing, usually remained separate, perpetuating the problem of overlay networks. This was typically due to restrictions of the byte-interleaved multiplexer (namely, the inability to support dynamic traffic for bursty users) and interface problems. A number of mission-critical networks in place today still conform to this architecture.

Third Generation Corporate Networks. In the immediate past, T1 multiplexers have started to support LAN interconnection traffic. The traffic is assigned a fixed amount of TDM bandwidth over the corporate backbone network. Although this approach to LAN interconnection had several advantages compared to the previous arrangement, it also had a number of disadvantages. Consider N nodes with high peak-to-average (bursty) traffic needing to be supported by an enterprise network. A mesh network providing full interconnection between key nodes may have been installed in many companies, as shown in Figure 11.2. This arrangement can be expensive due to the number of communications links. For example, five locations require 10 T1 links and six locations require 15 links. The addition of a new backbone node also requires the introduction of many new links. Less than fully

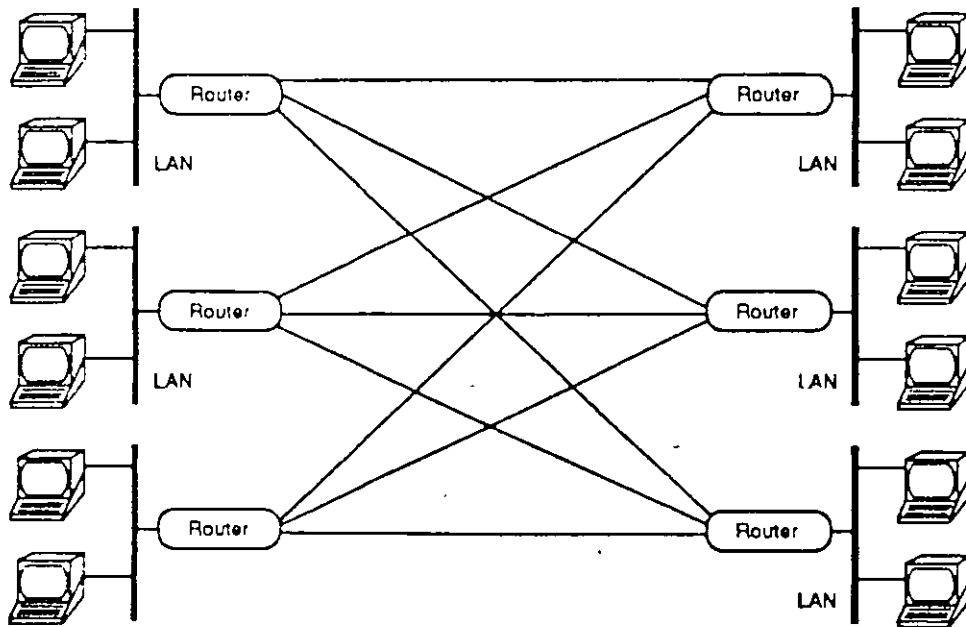


Figure 11.2 Traditional LAN connectivity.


interconnected router networks are usually not the best answer to the interconnection requirement, since this tandem arrangement affects the end-to-end delay and complicates network management. Additionally, and perhaps equally important, the bandwidth is not efficiently allocated by using the TDM techniques common to the equipment supporting this type of network. As a short-term solution, companies sought to keep the number of designated first-tier locations needing full interconnection down to a small number, typically between three and six, thereby limiting the number of required links. Many companies have these types of networks today, particularly for mission-critical applications.

Some LAN managers, understanding the intuitive advantage of packet switching, chose to rely on a private (or public) packet-switched network to interconnect the multitude of user routers. These packet networks typically introduce a router-to-router delay on the order of 200 ms or more. This delay is due to (1) protocol processing at intermediary nodes and (2) the hop-by-hop error correction and control used by packet networks. In addition to the delay, the throughput of these networks is not sufficient to support today's applications. That is why a new technology is needed.

Fourth Generation Corporate Networks. Figure 11.2 depicted a common configuration of user environments of the recent past. The figure makes the challenge of a LAN manager obvious—what is needed is a technology for high end-to-end throughput, low latency, cost-effective bandwidth on demand, and any-to-any connectivity. The restrictions discussed above of many of the existing unintegrated networks, or of the integrated networks using TDM technology, has led to the development of the frame relay concept, frame relay standards, and supporting frame relay hardware. Such frame relay solutions are now emerging and are composed of the following three components (see Figure 11.3):

- A high-throughput nodal processor built from the ground up, unaffected by TDM restrictions, which supports high-speed switching to facilitate high end-to-end throughput. This nodal processor was called a "corporate backbone network switch" in Chapter 6; as discussed there, it was initially used in the T1 multiplexer context, but is now making a strong appearance in the LAN context. This nodal processor can also be provided by a carrier, supporting a "switching" function at the CO.
- Standard high-speed interfaces to the bridge/routers to facilitate the interconnection of equipment from a variety of vendors. This open (standard) interface based on the concept of packet switching is a frame relay interface.
- Standard high-speed interfaces between nodal processors, supporting cell relaying and switching.

Since corporate resources are increasingly being deployed on LANs, including the users, the mainframes, and the data bases, and since routers have taken on the function of the T1 multiplexers in a number of environments, this intercon-



FRNP Frame Relay Nodal Processor

Figure 11.3 Use of frame relay nodal processors for a private frame relay network.

nection solution is the one being contemplated for the enterprisewide networks of the 1990s by many progressive companies. Fewer and shorter T1 or FT1 links can be used between the users and the nodal processor, thereby reducing communication costs.

11.2 BASIC FRAME RELAY AND CELL RELAY CONCEPTS

Before addressing how frame relay technology can be deployed in an enterprisewide multirouter environment, we must address some of the technical details [11.13]. Table 11.3 provides a miniglossary of key terms. As this discussion proceeds, it should be remembered that all high-speed lines used in frame relay, either for access or between switches, are unchannelized FT1, T1, or T3 lines, discussed in the rest of the book (T3 may be used on the trunk side).

11.2.1 Frames

A frame is a block of user data, as created by the data link layer (layer 2). It consists of a flag, a header, an information field, and a trailer. Different data link

Table 11.3
 Miniglossary of Frame Relay Terms

Access Rate. The rate of the access channel employed by the user's equipment (measured in bits per second). The speed of the access channel determines how rapidly the end-user can send data to the nodal processor or network.

American National Standards Institute. An organization that accredits groups developing U.S. standards required for commerce. One such group is the Exchange Carriers Standards Association, which developed the T1.606, T1.617, and T1.618 for frame relay (also see Chapter 1).

Asynchronous Transfer Mode. A packet switching technique developed by CCITT which uses packets of fixed length, resulting in lower processing and higher speeds. Also known as cell relay (see Chapter 10).

Bandwidth. The communications capacity (measured in bits per second) of a transmission line or a path through a network.

Backward Explicit Congestion Notification Indicator. A bit in the frame set by the network to notify the user's equipment that congestion avoidance procedures should be initiated in order to limit the amount of traffic injected into the network or sent to the nodal processor. The field is set in a frame going in the opposite direction of the congestion (i.e., it is sent to the origination). It is similar to a "slowdown" signal.

Bursty Traffic. Traffic where the ratio of the maximum intensity to the average intensity is very high (≥ 10). Typical of some LAN environments.

Consultative Committee on International Telephony and Telegraphy. A United Nations organization which develops international standards and interfaces for telecommunications. The frame relay standards are based on underlying CCITT standards (also see Chapter 1).

Cell. A fixed-length packet of user data (payload) plus an overhead. A cell is usually small, being 53 octets or less.

Cell Relay. A high-bandwidth, low-delay switching and multiplexing packet technology required to implement a frame relay network in an efficient manner. Trunk transmission technique used by nodal processors. Also known as ATM.

Committed Information Rate (CIR). Specifies the amount of bandwidth guaranteed to a user between any two points. CIR can be as high as the access rate. If the CIR is exceeded, the frame relay device can send the data, but it should set the DE bit to indicate that the data can be discarded if necessary.

Core Functions. Data link layer functions supported by frame relay. Core functions include frame delimiting, alignment, and transparency; frame multiplexing/demultiplexing using the address field; and detection of transmission errors.

Data Link Connection Identifier (DLCI). A field in the frame indicating a particular logical link over which the frame should be transmitted. The field has local significance, since it can be changed by the nodal processors as the frame traverses a single-node network (the input DLCI is mapped to an output DLCI). Multinode networks may "pipeline" cells to the network edges. In this case, virtual channel identifiers are used and remapped. Access DLCIs are assigned by the network manager, while trunk VCIs are allocated dynamically. Toward the network, the nodal processor associates each VCI with the physical address of the trunk over which the frame needs to be transmitted to reach its ultimate destination. Toward the user, VCIs are associated with the physical line supporting the DLCI identifying the user.

Data Link Layer. OSI/RM layer 2 functionality, responsible for reliable transmission over a single communication link. It combines data bits into a block called a "frame" and adds a frame check sequence to allow detection of bit errors at the remote point.

Table 11.3
Continued

Discard Eligibility Indicator. A field in the frame set by the user's equipment to indicate that the frame can be discarded if needed in case of congestion, in order to maintain the committed throughput.

Fast Packet. A term for various streamlined packet technologies, now synonymous with cell relay. Supports reduced functionality compared to X.25 packet switching, so that it can operate at much higher speeds.

Forward Explicit Congestion Notification Indicator. A field in the frame set by the network or nodal processor to notify downstream equipment and/or the destination equipment that congestion avoidance procedures should be initiated. The field is set in a frame in the direction of the destination. It is similar to a "hold-on" signal for received frames, as well as a destination "slowdown" signal for traffic from the destination.

Frame. A block of user data, as created by the data link layer. It consists of a flag, a header, an information field, and a trailer.

Frame Relay. A 1990s packet-based high-speed technology that provides for dynamic bandwidth allocation with high throughput and low delay to support the increasing amount of bursty traffic in the corporate environment. Frame relay defines a standardized format for data link layer frames transmitted over a network of interconnected LANs.

Frame Relay Assembler/Disassembler. A device or capability allowing non-frame relay terminals, typically in a non-LAN environment (e.g., SNA) to be carried in a frame relay network.

Frame Relay interface. A standardized interface between customer equipment and a nodal processor or a frame relay network. A two-layer protocol stack interface capability implemented at both endpoints of a link.

Error Correction. In frame relay, error correction and retransmission are done in the user equipment. The network can detect errors, but the correction is relegated to the end systems.

Link Access Procedure F. The data link layer protocol used in frame relay. It is specified by ANSI T1 618-1991 and is similar to CCITT Q 922 Annex A. It is a slimmed-down protocol supporting core functions only. It is based on ISDN's LAP-D protocol.

Link Access Procedure F-Core. That subset of LAP-F used in frame relay.

Local Management Interface (LMI). A specification for use by frame relay products which defines a method of exchanging status information between the user device and the network. It is used to manage PVCs and is specified in ANSI's T1 617

Multicast. An LMI option that allows a frame relay device to broadcast frames to multiple destinations.

Nodal Processor. A frame relay processor is a switch that "connects" users, facilitating any-to-any connectivity. Connections are accomplished in real time over the PVC (eventually over an SVC). Tables are maintained by the node to facilitate the connections.

Open Systems Interconnection Reference Model. A model for data communications interconnection which maps functions necessary for undertaking orderly communication to one of seven hierarchical layers

Permanent Virtual Circuits (PVC). A logical link or path between the originating and terminating routers. No resources are allocated to the link unless data is actually being sent. The link is set up by the administrator and remains in place for however long it is needed (days, months, or years).

Router. A device operating at the network layer of the OSIRM, used by a LAN to access other LANs across a variety of WANs.

Table 11.3
Continued

Systems Network Architecture. A network architecture used in IBM networks in support of mission-critical functions. Originally, the architecture was strictly hierarchical and employed front-end processors and cluster controllers. It is now moving toward a peer-to-peer architecture, supporting LAN access to the mainframe.

Switched Virtual Circuit. A virtual circuit that is set up on a call-by-call basis. A future frame relay service, of particular importance to public frame relay networks.

Time-Division Multiplexing. A traditional method of combining multiple simultaneous channels over a single transmission path by assigning discrete time slots to each channel. It results in inefficient bandwidth allocation in bursty environments.

Virtual Circuit. A logical connection established through a frame relay or packet network. Frames or packets are routed through the network in an order-preserving transfer. The connection is similar to a dedicated line between the endpoints.

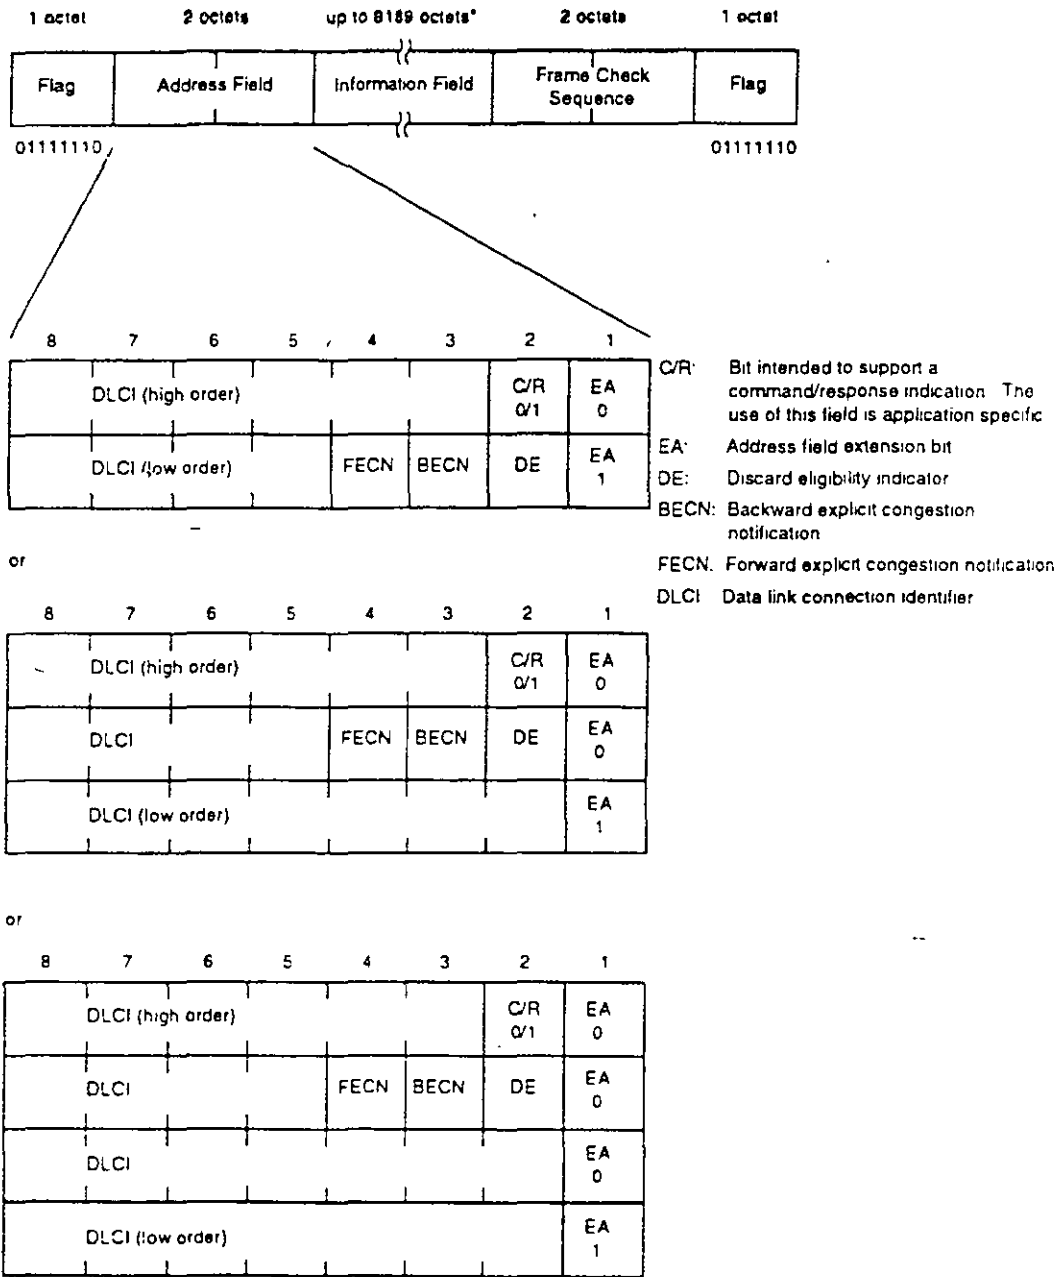
Virtual Circuit Identifier (VCI). A label used by a cell switch to identify cells belonging to a given user. VCIs have local significance (see Chapter 10).

layers create different frames; differences manifest themselves in terms of the fields, their positions, and their lengths. The logical link control sublayer of a LAN creates a frame of particular interest, since it is the frame that is involved in the transmission of data over a network of interconnected LANs. Frame relay service has a specific frame format, described below and expanded upon in Section 11.4.2.

11.2.2 Frame Relay

Frame relay is a new packet-based high-speed technology. It provides for dynamic bandwidth allocation with high throughput and low delay to support the increasing amount of bursty traffic in the corporate environment. Frame relay defines a standardized format for data link layer frames, which are transmitted over a network of interconnected LANs or over a public network. A frame relay frame is assembled by user equipment and is interpreted by frame relay nodal processors or, in cases where there are no processors, by the remote router. The frame relay frame is shown in Figure 11.4. Frame relay is based on the 1988 and 1992 CCITT standards and recent ANSI extensions, clarifications, and refinements for the U.S. market, particularly for private nonswitched access. There is now wide vendor support of frame relay standards.

Frame relay may be thought of as a streamlined version of X.25 that can be implemented on or integrated onto a router. Figure 11.5 depicts this simplification [11.4, 11.14]. Streamlining is accomplished by stripping away all of the X.25 network layer (layer 3), adding a statistical multiplexing capability via individually addressed frames to the data link layer, and reducing the functionality of layer 2 by removing



*ANSI specifies 4096 but some vendors may use the higher value; minimum length, 262.

Figure 11.4 Frame and address field format (LAP-F/ANSI T1.618-1191; CCITT Q.922 Annex A).

error correction and retransmission capabilities. Error detection is retained and errored frames are discarded by the frame relay network. Frame relay can provide both a PVC and a SVC service.

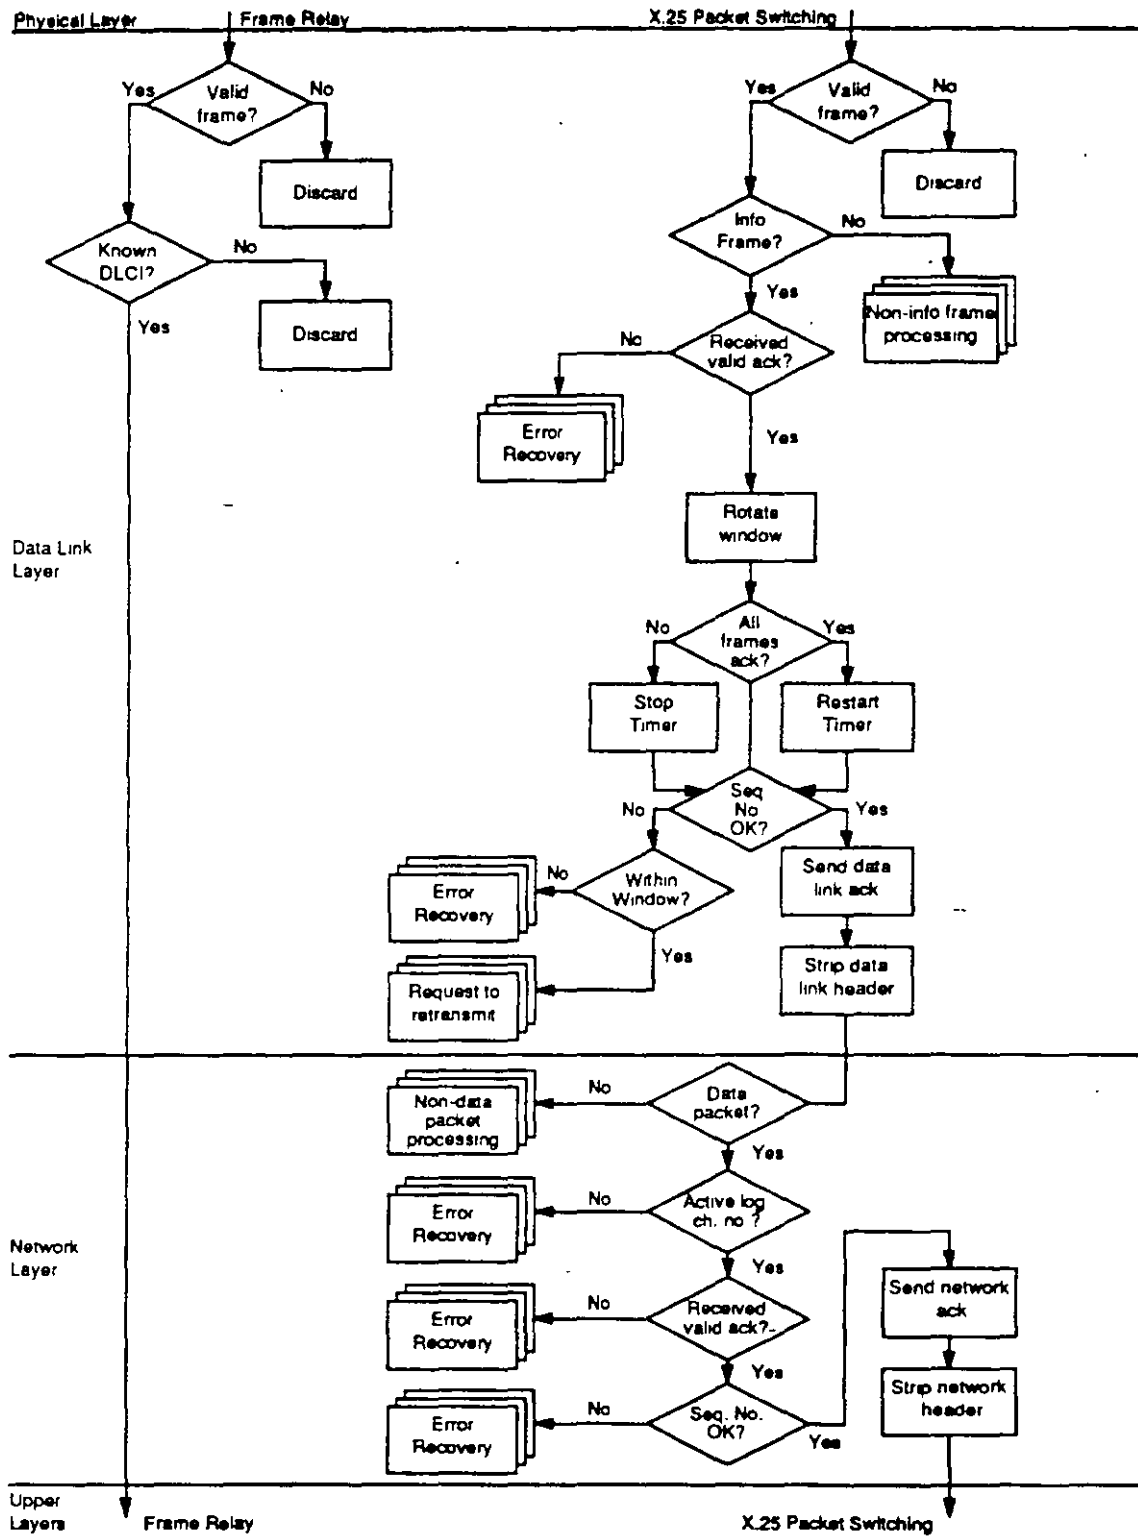


Figure 11.5 Comparison of frame relay protocol state diagram with the one for X.25.

Initially, frame relay was developed as an ISDN packet service, with a logically separate control plane for SVC management and user plane. In the control plane, all signaling capabilities for call control, parameter negotiation, etc., would be based on a set of protocols common to all ISDN services. As currently evolving at the U. S. commercial level, frame relay service does not support a logically separate control plane; in fact, no call setup mechanisms are supported, even "in-band" (as would be the case for X.25 SVC packet switching). Public SVC frame relay may become available in 1994 or 1995. Currently, there is no network equipment to support user-to-network or network-to-network SVC signaling, and there is no SVC user equipment. This implies that the service now only supports PVCs predefined by the network administrator, and that all frames follow the same route to the destination.

Permanent virtual circuits establish a fixed path through the network so that a message or file can be reassembled quickly by the receiving end. Frames are passed across the network with minimal processing by the network nodes. All of the bandwidth on the physical path of the frame is available for the duration of the frame. The result is a high-speed, low-delay, bandwidth-on-demand network well suited for LAN-to-LAN traffic. Frame relay has the advantage of providing line consolidation and, hence, reduced equipment expenditures. It has the inherent multiplexing capability where one physical access can support up to 1,024 logical connections (in actual implementations, however, the number of logical channels is usually much smaller, say, 32 for physical and/or performance reasons [11.2, 11.15]). The PVC approach does not support an addressing apparatus adequate to support a switched service.

11.2.3 Frame Relay Interfaces

Like X.25, frame relay specifies the interface between customer equipment and the network (i.e., the UNI), whether the network is public or private. This interface specification is described in generalities in CCITT Recommendation I.122 of 1988. I.122 describes the unacknowledged order-preserving transfer of data units from the network side of a user-network interface to the network side of the other user-network interface. A frame relay interface is a two-layer protocol stack capability implemented at both endpoints of a link (i.e., by the user equipment and by the network's nodal processor). Frame relay interfaces rely on (1) the existing intelligence of end-user equipment, such as routers, to run the protocol; (2) today's higher quality digital transmission facilities; and (3) error detection, correction, and recovery at the higher end-system layers (transport layer, or even at the ultimate application level) [11.13].

The frame format for data transfer is based on a subset of Q.921 (LAP-D), but extended with the flow control fields. The protocol is now known as Link

Access Procedure F-Core (LAP-F Core) and is defined in ANSI T1.618-1991; it is also defined in CCITT's Q.922 Annex A, adopted in 1992. Q.922/LAP-F (ISDN Data Link Layer Specification for Frame Mode Bearer Services) is a full data link layer protocol in its own right; it was adopted in 1991. Frame relay uses the subset called LAP-F Core. LAP-F functions like windowing and error correction are not included in the core subset [11.16]. A 1992 CCITT protocol extended from Q.931, called Q.933, is to be used to support SVC service.

As implied in Figure 11.1, a network platform is needed to route and forward frames received over the UNI conforming to the FRI specification. This platform is composed of one or more nodal processors (whether owned by the user or by a carrier).

11.2.4 Error Correction

In frame relay, error correction and retransmission are done in the user equipment. The network can detect errors, but the correction is relegated to the end-systems. Error conditions include lost, duplicated, misdelivered, discarded, and out-of-sequence frames; recovery from these error conditions must be performed by the user's equipment, which must be appropriately configured to support these tasks. This does not require any additional functionality, which most intelligent equipment, like LAN routers, has today. Furthermore, with today's higher quality digital transmission facilities and the migration to fiber, it is unlikely that many frames will be received in error, requiring end-to-end retransmission. Error-prone circuits of the past necessitated complex error checking and recovery procedures at each node of a network. The X.25 packet standards assume that the transmission media is intrinsically error-prone, and in order to guarantee an acceptable level of end-to-end quality, error management is performed at every link by a fairly sophisticated but resource-intensive data link protocol, as illustrated in Figure 11.5. With a high-quality fiber-based communication infrastructure becoming commonplace, many of the error correction and retransmission capabilities of X.25 can be safely eliminated [11.17].

Since error correction and flow control are handled at the endpoints, frame relay expedites the process of routing packets through a series of switches to a remote location by eliminating the need for each switch to check each packet and correct those in error. This error treatment increases performance and reduces bandwidth requirements, which in turn can reduce communications costs [11.18, 11.19].

In the past, when transmission errors were common, it was not efficient to require the transport layer (whose job it is to guarantee ultimate end-to-end reliability) to keep track of unacknowledged PDUs. Instead, the data link layer, closer to where the problem had its roots, was responsible for the correction task. It turns

out in the final analysis that when the probability of error over a link is relatively high, it is better to do error correction on a link-by-link basis (i.e., at the data link layer), as measured by the amount of network bandwidth required to successfully send a PDU (although it may, in fact, have been faster to do it end to end, as we show below). When the probability of error is low, it is better to do error correction end to end (i.e., at the transport layer). In other words, for the same amount of network bandwidth, the PDU gets delivered faster by doing the error management end to end; in addition, the nodal processors can be cheaper, since they need to undertake fewer tasks.

Tables 11.4, 11.5, and 11.6 provide a numerical example of a simplified model to illustrate the point (the reader may choose to skip ahead to Section 11.2.5, if desired). In each case, a three-link path is studied. In the first case, the probability of link success is 0.9, in the second case it is 0.99, and in the third case it is 0.999 (the same probability for all three links is used for simplicity). The tables (almost precisely) compute the expected bandwidth to deliver a correct PDU with link-by-link correction and with end-to-end correction. The expected delay to deliver a correct PDU with link-by-link correction and with end-to-end correction is also computed. Before discussing briefly how the tables are derived, let us focus on the results. The following summary emerges:

	<i>Expected Bandwidth (l-b-l)</i>	<i>Expected Bandwidth (e-t-e)</i>	<i>Expected Delay (l-b-l)</i>	<i>Expected Delay (e-t-e)</i>
Link success probability: 0.9	3.30674	4.02831	0.66135	0.53711
0.99	3.03027	3.09182	0.60606	0.41224
0.999	3.00300	3.00902	0.60060	0.40120

As the probability of successful transmission over a link increases, the expected bandwidth approaches three units in the link-by-link case: one unit for the first link, one unit for the second link, and one unit for the third link. When the probability is lower, the expected bandwidth goes up because of the required retransmissions. (Some PDUs will require no retransmissions, while others will require a few retransmissions. On the average, 3.3 units of bandwidth would be required for the 0.9 case; since the example is fairly realistic, the "overhead" bandwidth is rather small. To illustrate the point more pedagogically it would be necessary to use a failure value of 0.5.) Note that, in the end-to-end correction case, considerably more expected bandwidth to successfully transmit a PDU would be required when the probability of link failure is relatively high (4.02 units versus 3.30 units for the link-by-link case). However, as the probability of success increases to 0.999, effectively the same expected bandwidth is required by both methods (3.009 units versus 3.003 units). More bandwidth is required in the end-to-end correction case when the BER is high, because the distant node would have to go

Table 11.4
 Delay and Throughput Comparison When Probability of Successful Path Transmission Is 0.9

Probability of transmission success over a link: 0.9

Transmission delay (units): 0.1

Node delay for error management (units): 0.1

<i>Link 1</i>	<i>Link 2</i>	<i>Link 3</i>	<i>Total Units of Bandwidth Required for Link-by-Link Correction</i>	<i>Probability</i>	<i>Expected Bandwidth (Link-by- Link)</i>	<i>End-to-End Delay</i>	<i>Expected Delay</i>
s	s	s	3	0.72900	2.18700	0.600	0.43740
f.s	s	s	4	0.07290	0.29160	0.800	0.05832
s	f.s	s	4	0.07290	0.29160	0.800	0.05832
s	s	f.s	4	0.07290	0.29160	0.800	0.05832
f.s	f.s	s	5	0.00729	0.03645	1.000	0.00729
f.s	s	f.s	5	0.00729	0.03645	1.000	0.00729
s	f.s	f.s	5	0.00729	0.03645	1.000	0.00729
f.f.s	s	s	5	0.00729	0.03645	1.000	0.00729
s	f.f.s	s	5	0.00729	0.03645	1.000	0.00729
s	s	f.f.s	5	0.00729	0.03645	1.000	0.00729
f.f.f.s	s	s	6	0.00073	0.00437	1.200	0.00087
s	f.f.f.s	s	6	0.00073	0.00437	1.200	0.00087
s	s	f.f.f.s	6	0.00073	0.00437	1.200	0.00087
f.f.s	f.s	s	6	0.00073	0.00437	1.200	0.00087
f.f.s	s	f.s	6	0.00073	0.00437	1.200	0.00087
s	f.f.s	f.s	6	0.00073	0.00437	1.200	0.00087
...				0.995814	3.306744		0.66135
<i>1st Pass</i>	<i>2nd Pass</i>	<i>3rd Pass</i>	<i>Total Units of Bandwidth Required for Link-by-Link Correction</i>	<i>Probability</i>	<i>Expected Bandwidth (Link-by- Link)</i>	<i>End-to-End Delay</i>	<i>Expected Delay</i>
s.s.s			3	0.72900	2.18700	0.4	0.29160
s.s.f							
s.f.s							
f.s.s							
s.f.f							
f.s.f							
f.f.s							
f.f.f							
>any f	s.s.s		6	0.19756	1.18535	0.8	0.15805
>any f	any f	s.s.s	9	0.05354	0.48185	1.2	0.06425
>any f	any f	any f	12	0.01451	0.17411	1.6	0.02321
...				0.99461	4.02831		0.53711

s = success f = failure

Table 11.5
 Delay and Throughput Comparison When Probability of Successful Path Transmission Is 0.99

Probability of transmission success over a link: 0.99

Transmission delay (units): 0.1

Node delay for error management (units): 0.1

<i>Link 1</i>	<i>Link 2</i>	<i>Link 3</i>	<i>Total Units of Bandwidth Required for Link-by-Link Correction</i>	<i>Probability</i>	<i>Expected Bandwidth (Link-by- Link)</i>	<i>End-to-End Delay</i>	<i>Expected Delay</i>
s	s	s	3	0.97030	2.91090	0.600	0.58218
f.s	s	s	4	0.00970	0.03881	0.800	0.00776
s	f.s	s	4	0.00970	0.03881	0.800	0.00776
s	s	f.s	4	0.00970	0.03881	0.800	0.00776
f.s	f.s	s	5	0.00010	0.00049	1.000	0.00010
f.s	s	f.s	5	0.00010	0.00049	1.000	0.00010
s	f.s	f.s	5	0.00010	0.00049	1.000	0.00010
f.f.s	s	s	5	0.00010	0.00049	1.000	0.00010
s	f.f.s	s	5	0.00010	0.00049	1.000	0.00010
s	s	f.f.s	5	0.00010	0.00049	1.000	0.00010
f.f.f.s	s	s	6	0.00000	0.00001	1.200	0.00000
s	f.f.f.s	s	6	0.00000	0.00001	1.200	0.00000
s	s	f.f.f.s	6	0.00000	0.00001	1.200	0.00000
f.f.s	f.s	s	6	0.00000	0.00001	1.200	0.00000
f.f.s	s	f.s	6	0.00000	0.00001	1.200	0.00000
s	f.f.s	f.s	6	0.00000	0.00001	1.200	0.00000
				0.999996	3.030279		0.60606
<i>1st Pass</i>	<i>2nd Pass</i>	<i>3rd Pass</i>	<i>Total Units of Bandwidth Required for End-to-End Correction</i>	<i>Probability</i>	<i>Expected Bandwidth (End-to- End)</i>	<i>End-to-End Delay</i>	<i>Expected Delay</i>
s.s.s			3	0.97030	2.91090	0.4	0.38812
s.s.f							
s.f.s							
f.s.s							
s.f.f							
f.s.f							
f.f.s							
f.f.f							
>any f	s.s.s		6	0.02882	0.17291	0.8	0.02306
>any f	any f	s.s.s	9	0.00086	0.00770	1.2	0.00103
>any f	any f	any f	12	0.00003	0.00031	1.6	0.00004
				1.00000	3.09182		0.41224

s = success f = failure

Table 11.6
 Delay and Throughput Comparison When Probability of Successful Path Transmission Is 0.999

Probability of transmission success over a link: 0.999							
Transmission delay (units): 0.1							
Node delay for error management (units): 0.1							
<i>Link 1</i>	<i>Link 2</i>	<i>Link 3</i>	<i>Total Units of Bandwidth Required for Link-by-Link Correction</i>	<i>Probability</i>	<i>Expected Bandwidth (Link-by- Link)</i>	<i>End-to-End Delay</i>	<i>Expected Delay</i>
s	s	s	3	0.99700	2.99101	0.600	0.59820
f,s	s	s	4	0.00100	0.00399	0.800	0.00080
s	f,s	s	4	0.00100	0.00399	0.800	0.00080
s	s	f,s	4	0.00100	0.00399	0.800	0.00080
f,s	f,s	s	5	0.00000	0.00000	1.000	0.00000
f,s	s	f,s	5	0.00000	0.00000	1.000	0.00000
s	f,s	f,s	5	0.00000	0.00000	1.000	0.00000
f,f,s	s	s	5	0.00000	0.00000	1.000	0.00000
s	f,f,s	s	5	0.00000	0.00000	1.000	0.00000
s	s	f,f,s	5	0.00000	0.00000	1.000	0.00000
f,f,f,s	s	s	6	0.00000	0.00000	1.200	0.00000
s	f,f,f,s	s	6	0.00000	0.00000	1.200	0.00000
s	s	f,f,f,s	6	0.00000	0.00000	1.200	0.00000
f,f,s	f,s	s	6	0.00000	0.00000	1.200	0.00000
f,f,s	s	f,s	6	0.00000	0.00000	1.200	0.00000
s	f,f,s	f,s	6	0.00000	0.00000	1.200	0.00000
...				1	3.003003		0.60060
<i>1st Pass</i>	<i>2nd Pass</i>	<i>3rd Pass</i>	<i>Total Units of Bandwidth Required for End-to-End Correction</i>	<i>Probability</i>	<i>Expected Bandwidth (End-to- End)</i>	<i>End-to-End Delay</i>	<i>Expected Delay</i>
s,s,s			3	0.99700	2.99101	0.4	0.39880
s,s,f							
s,f,s							
f,s,s							
s,f,f							
f,s,f							
f,f,s							
f,f,f							
>any f	s,s,s		6	0.00299	0.01793	0.8	0.00239
>any f	any f	s,s,s	9	0.00001	0.00008	1.2	0.00001
>any f	any f	any f	12	0.00000	0.00000	1.6	0.00000
...				1.00000	3.00902		0.40120

s = success f = failure

back all the way to the origination (several hops away) and tie up resources across the network. For example, if a PDU has successfully traversed two hops and then fails on the last hop, the resources required for a correct reception under the hop-by-hop method would be $B + B + (B + B)$; under the end-to-end method it would be $(B + B + B) + B + B + B$.

In each case, the time required to successfully send a correct PDU decreases as the probability of correct transmission over a link increases. However, the end-to-end correction case was (in this case) always superior in terms of speed (0.66 units for link-by-link versus 0.54 units at 0.9, and 0.60 units for link-by-link versus 0.40 at 0.999; individuals engineering packet-switched networks have trade-off bandwidth efficiency in favor of end-to-end response time). The data make clear that if the probability of link error is low, it is possible to relegate error correction to the endpoints of the network without negatively impacting the throughput, while substantially improving response time. The additional delay in a traditional packet-switched network is incurred in having to manage error situations at each hop instead of just doing it once at the endpoints, as illustrated in Figure 11.5. (This more efficient use of bandwidth is not by itself responsible for the increased throughput of a frame relay network; the increased throughput facilitating T1 access is due to the fact that faster switches are used.)

What this analysis should also make clear is that, in single-switch frame relay applications, as some vendors are suggesting for private frame relay networks, the advantages gained by relegating the error management to the endpoints are minimal, if they exist at all. The path and the link are nearly identical in an environment where the frame relay routers are connected through a single nodal processor, and are exactly identical when two frame relay routers are connected directly with a T1 line. In fact, the overhead incurred in segmenting and reassembling a 1,500-octet Ethernet frame into 70 or so cells may practically wash away any gain from streamlining the error management procedure (at least in a single-node frame relay network).

Some explanation of the model of Tables 11.4, 11.5, and 11.6 follows. Let p be the probability of successful transmission over a single link. In the link-by-link error procedure, a PDU is successfully transmitted if it is successful over the first, second, and third link. The probability of this event is $p \times p \times p = p^3$. In this case, one unit of bandwidth is expanded over the first link, one unit over the second link, and one unit over the third link, for a total of three units. The delay is $(T + P) + (T + P) + (T + P)$, where T is the transmission time and P is the protocol processing time. The model proceeds by looking at all (in actuality, the most significant) other events. For example, there could be an unsuccessful transmission over the first link, followed by a successful retransmission and two other transmissions. The probability of this is $(1 - p) \times p \times p \times p = (1 - p)p^3$. In this case, two units of bandwidth are expanded over the first link, one unit over the second link, and one unit over the third link, for a total of four units. The delay is $\{(T +$

$P) + (T + P)] + (T + P) + (T + P)$. Other cases are shown in the table (99.4% or more of all cases are accounted for in the tables). Finally, the expected bandwidth value is calculated as

$$PR_{\text{case1}} \times \text{Bandwidth}_{\text{case1}} + PR_{\text{case2}} \times \text{Bandwidth}_{\text{case2}} + \dots$$

and the expected delay is calculated as

$$PR_{\text{case1}} \times \text{Delay}_{\text{case1}} + PR_{\text{case2}} \times \text{Delay}_{\text{case2}} + \dots$$

where PR_x = Probability of combined event x . The second part of the table calculates the same factors when the error processing is done end to end. Here a transmission is successful only if successful over the first, second, and third links. The probability of this is $p \times p \times p = p^3$. In this case, one unit of bandwidth is expanded over the first link, one unit over the second link, and one unit over the third link, for a total of three units. The delay is $(T) + (T) + (T + P)$, where T is the transmission time and P is the protocol processing time (in this case, P could be incurred by the end-user equipment, but it still impacts the response time; in addition, P here is smaller than in the previous case). The model proceeds by looking at all (in actuality, the most significant) other events. Any link failure leads to an end-to-end retransmission. Events such as s,s,f or s,f,s or f,s,s or s,f,f or f,s,f or f,f,s or f,f,f (s = success, f = failure) fall into this category (in practical terms, an errored frame is not actually transmitted—the prescription shown describes an *a posteriori* probabilistic characterization). Each of these cases would be followed by a second phase, which, it is hoped, would be of the form s,s,s. The probability of this is $(1 - p^3) \times p \times p \times p$. In this case, as many as three units of bandwidth are expanded in the first case, plus three units in the second phase, for a total of six units. The delay is $2 \times [(T) + (T) + (T + P)]$. Other cases are shown in the table (99.4% or more of all cases are accounted for in the tables). Finally, the expected bandwidth usage and delay are calculated as described above.

Public frame relay networks must be designed with quality of service in mind. Some of the parameters being discussed are

$$\text{Ratio of nondelivered PDUs to total PDUs} \leq 10^{-4}$$

$$\text{Ratio of errored PDUs to total PDUs} \leq 10^{-13}$$

$$\text{Ratio of misdelivered PDUs to total PDUs} \leq 10^{-8}$$

$$\text{Ratio of duplicate PDUs to total PDUs} \leq 10^{-9}$$

(Today's frame relay networks do not yet meet these goals.) If the network is not properly engineered from a traffic perspective (i.e., insufficient trunk bandwidth is provided) and unreliable flow control procedures are used, the number of network-discarded frames could become significant; this issue will be reexamined later. In addition to the quality of service measures with reference to error conditions,

carriers aim for an end-to-end delay of about 250 ms per average frame (1,000 octets) over a DS0 access line and 20 ms over a DS1 access line. Carriers have the opportunity of tariffing a level of service or network delay [11.2]. Since frame relay will not work too well over noisy lines (e.g., in some countries outside the U.S.), it will not have the international reach of traditional packet switching.

11.2.5 Frame Relay Processors

A frame relay processor is a switch that "connects" users, facilitating any-to-any connectivity. Connections are accomplished in real time over the PVC (the PVC itself, however, had to be previously established). The frame relay interface is only a definition of what the data stream into the frame relay network looks like. Equipment in the form of nodal processors is needed in the network (private or public) to make the frame relay concept a reality. Like a packet switch, a frame relay nodal processor supports a virtual connection. Tables are maintained by the node that tell the node the physical port on which an incoming frame must be transmitted. For users terminating on the same node, the frames are directly sent to the destination by checking the address and determining which physical port needs to receive the data. For users terminating on two different nodes, the data must be sent over the appropriate trunk to the destination node for ultimate delivery. Centralized administration of the backbone network routing tables and the natural port sharing and multiplexing attributes of frame relay make network growth manageable and simple. According to some observers, the annual demand for fast packet frame relay technology will surpass TDM-based T1 multiplexer systems by 1995 [11.20].

11.2.6 Frame Relay Networks

Frame relay networks can be private, public, or hybrid. A network consists of (1) user equipment supporting the frame relay interface, (2) one or more frame relay processors owned by the user or a carrier, and (3) communication links between the users and the nodal processors and between the processors (links between the nodal processors are owned by the carrier in a public network). The user equipment typically consists of appropriately configured LAN routers. The nodal processors interpret the frame and transmit them (using cells or, in some cases, frames), making the concept of frame relay a reality. Figure 11.6 shows an example of a (public) frame relay network: frames traverse a fixed PVC path through the network, although transmission resources (including bandwidth) are not dedicated to each virtual connection.

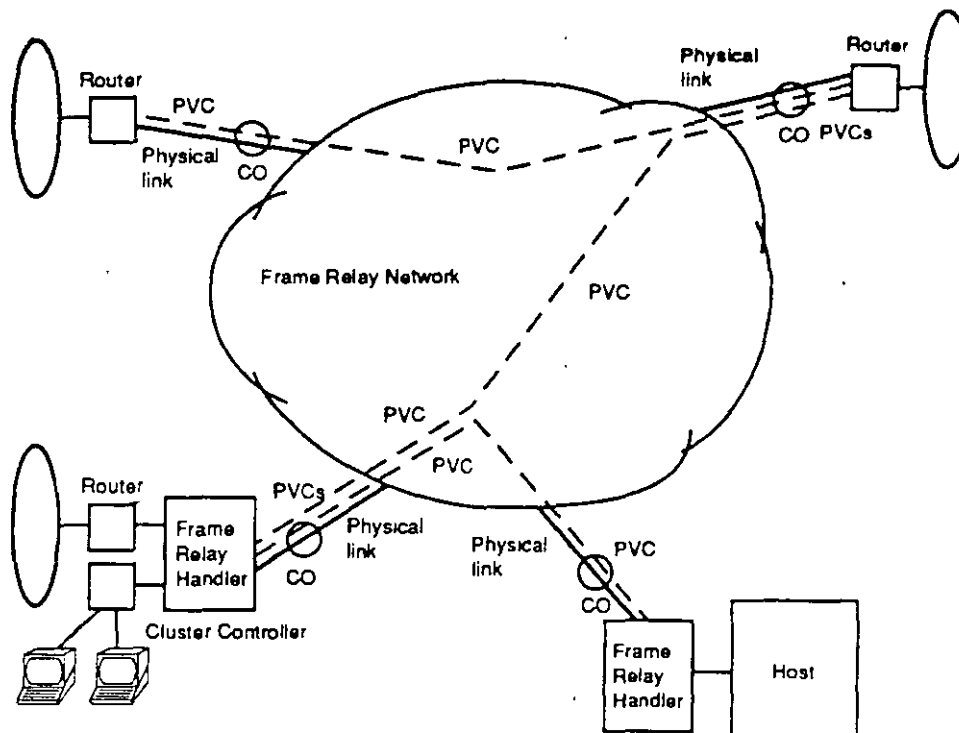


Figure 11.6 A (public) frame relay network.

11.2.7 Transmission Mode

Frame relay is a connection-oriented technology. Traditionally, CCITT, the original frame relay standardization body, has pursued a connection-oriented philosophy. Connection-oriented service involves a connection establishment phase, a data transfer phase, and a connection termination phase. A logical connection is set up between end-systems prior to exchanging data. These phases define the sequence of events, ensuring successful data transmission. Sequencing of data, flow control, and transparent error handling are some of the capabilities typically inherent in the service. The call setup phase (as would be the case in the SVC environment) adds some delay to each call, but it facilitates dynamic connectivity. For today's permanent virtual circuit-based frame relay, setup is done once by the system administration on behalf of the user. The PVC approach implies the allocation of some resources—like table entries—regardless of the real-time user traffic condition). Since the PVC is established at subscription time, there is no need for real-time signaling in this type of service (there may be status signaling, but this is unrelated to the establishment of the channel).

In a connectionless service, such as SMDS, each data unit is independently routed to the destination; no connection-establishment activities are required, since each data unit is independent of the previous or subsequent one. Connectionless-mode service provides for unit data transfer without regarding the establishment or maintenance of connections. This is advantageous in interenterprise applications. Each unit of data contains the addressing information and the data itself. The responsibility of ensuring that the message gets at the other hand is shifted up from the data link layer to higher layers, where the integrity check is done only once, instead of being done at (every) lower layer. Connectionless communication is now a very common technique, and is found, for example, in LANs. Since SMDS is also connectionless, the two technologies can interwork in an optimized fashion.

11.2.8 PVC Establishment in Private Networks

The backbone frame relay processors typically have a centralized network management terminal to provision connections. The manager specifies the endpoints (i.e., the two routers for which a PVC is desired). The network management system will then automatically build a path between the nodes (and, hence, the endpoints) and inform all nodes in the network of the route. Some processors require manual entry of the entire routing path in the various tables. This path will be used for all subsequent transmission between the specified endpoints. The manager can also specify alternate logical/physical routers to deal with node or trunk failure (user access line failure cannot be dealt with by this method).

11.2.9 Frame Relay Protocol Stack and Protocols

Figure 11.7 depicts a typical frame relay network protocol architecture. These stacks must be implemented in the user equipment and in the nodal processors in order to implement frame relay. In the example, there are two PC users on two geographically separate LANs. These LANs would access the frame relay node via routers configured to terminate the frame relay interface. There are two PC users on the two remote LANs. Three network nodes have been provisioned to logically interconnect the end-user equipment via permanent virtual circuits. Nodes 1 and 3 terminate the end-user equipment directly over a link with a frame relay interface. They must support segmentation functions like CS and SAR (discussed in Chapter 10) in order to accommodate cell-based transmission within the network.

Standards work for frame relay started in 1986; work accelerated in 1989, after the publication of the first CCITT frame relay standards. CCITT's Q.922 and ANSI's T1.618, T1.617 Annex B, and T1.617 Annex D describe the UNI. Transfer of PDUs is based on Core Aspects of LAP-F protocol (ANSI T1.618). LAP-F equates to Q.922 and to the older "I.441* Core" defined in the 1988 version of

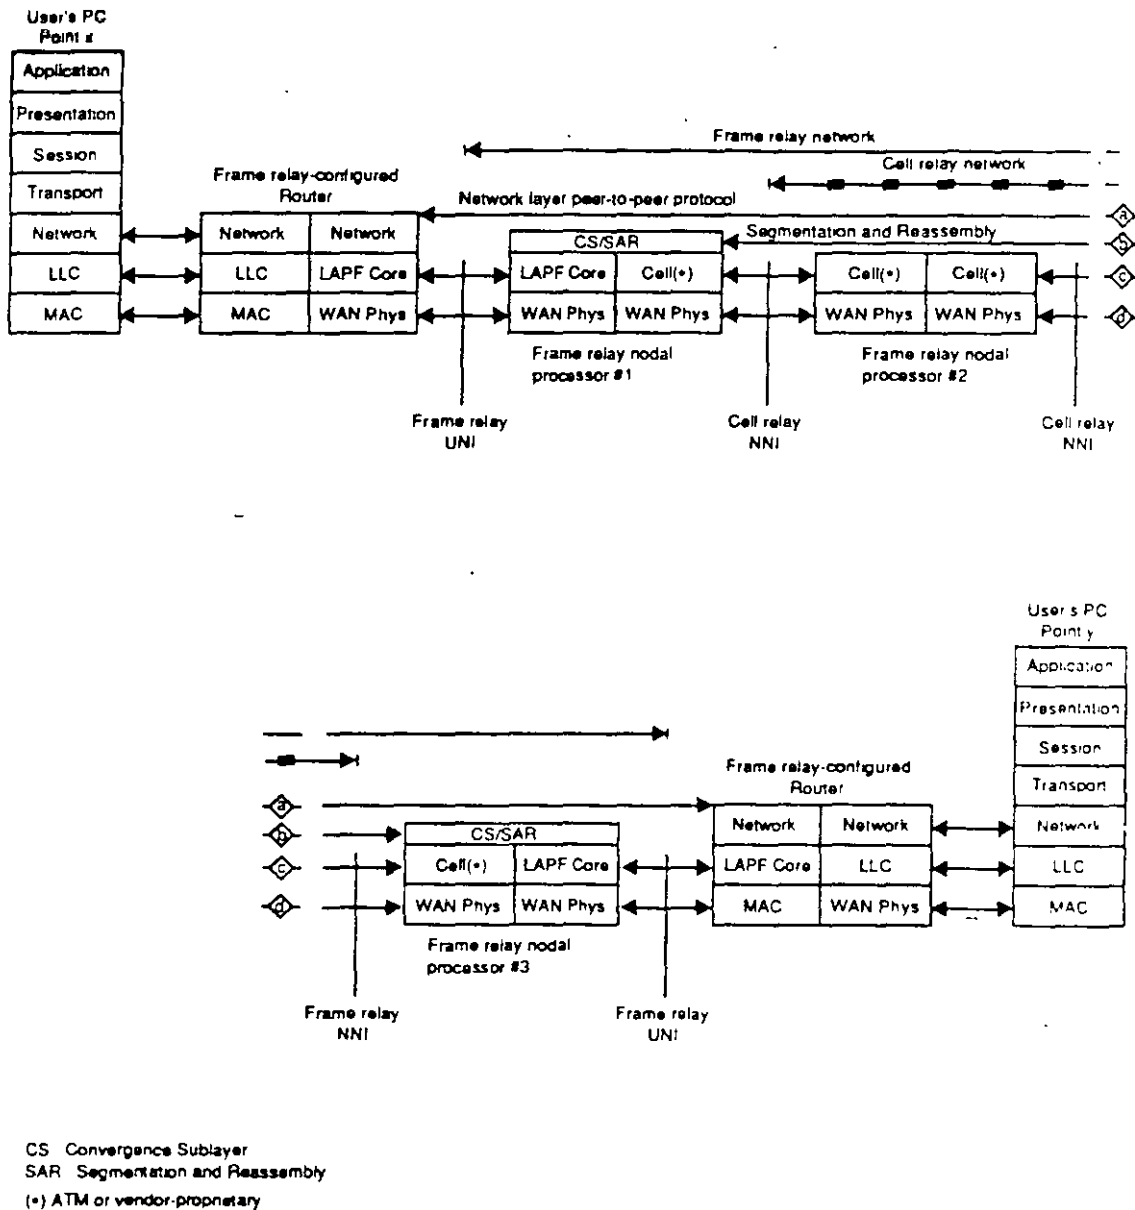


Figure 11.7 Protocol stacks in a frame relay/cell relay network

I.122. PVC management functions are included in T1.617 Annex D; many features of the local management interface (LMI) specification were initially proposed by vendors and by the frame relay forum.

Table 11.7 summarizes the status of the standards. T1.606 provides a description of the frame relay service. It was approved in 1990 by ANSI. The equivalent CCITT recommendation, I.233, was in the final stages of approval. T1.606 Adden-

Table 11.7
Frame Relay Standards, January 1992

<i>Standard</i>	<i>ANSI Number</i>	<i>Status</i>	<i>CCITT Number</i>	<i>Status</i>
Framework			I.122	Published
Service Description	T1.606	Published	I.233	Final
Congestion Management	T1.606 Addendum	Approved	I.370	Final
Data Transfer Protocol	T1.618	Approved	Q.922 Annex A	Final
Access Signaling	T1.617	Approved	Q.933	Frozen

dum describes congestion management. It was approved in November 1991. The equivalent CCITT recommendation, I.370, was in the final stages of approval at press time for the March 1992 CCITT meeting. T1.618, based on core aspects of the LAP-F protocol, describes the data transfer protocol at the UNI. The standard was approved in 1991. The equivalent CCITT recommendation, Q.922 Annex A, was approved in March 1992. T1.617 describes access signaling. It was approved in 1991. The equivalent CCITT recommendation, Q.933 was approved in March 1992. T1.617 Annex B describes management of PVCs on a channel that supports a mix of PVC and SVC services. T1.617 Annex D provides key PVC management functions. The LMI of Annex D makes possible for the network to notify the end-user of the addition, deletion, or presence of a PVC at a specified UNI (any such information received on a UNI applies to that particular UNI). Areas requiring standardization include NNI protocols and interoffice signaling.

The minimum information field allowed by the protocols is 1; this implies that there are no restrictions on how small the frame is. A total of 1,021 PVCs per UNI are supported. Logical channel 0, 1, and 1023 are reserved; channel 1023 is used to send link layer management messages from the network to the user's device; other logical channels (up to 45) may be reserved by some carriers.

11.2.10 Transmission Mechanism Across a Frame Relay Network

When using the frame relay interface, the router on a LAN selects the required remote router by specifying the permanent virtual circuit via a data link connection identifier contained in the frame relay frame it builds prior to transmitting the data (the identifier is originally assigned by the network administrator). If the system is well designed, there should be no segmentation of the LAN frames into multiple frame relay frames, although this could happen in theory, adding delay and overhead. The nodal processor accepts the frame it receives on one of its incoming ports, segments it into cells while appending a sequence number for remote-switch cell-to-frame reassembly, and delivers it over the trunk connecting to that remote switch. Initially, trunk interfaces used a "packet-like" protocol; more recently,

products are moving in the direction of CCITT's ATM cells standards. The switch must segment incoming frame relay frames for delivery through the cell format, because these frames can be long while cells are much shorter; for example, CCITT's standards specify cells with 48 octets for the payload and five extra cells for overhead, as discussed in Chapter 10. As indicated, the network does not worry about error correction. Nodal processors do, however, check the frame check sequence (FCS⁷) code of a received frame. If the frame is found in error, it is dropped without further processing (refer again to Figure 11.5).

Every network needs to ensure that traffic is routed reliably from the source to the destination. In a frame relay network, routing of the frames from the various routers is determined by the DLCI of the frame on a given user-network interface. Nodes use the DLCI to determine the frame's destination. The DLCI is not an address of the destination, since it may change as the frame travels through the network (i.e., the DLCI has local significance only). Instead, it identifies the logical connection between an element in the network and the next element in the network (i.e., endpoint and nodal processor, and nodal processor and endpoint; routing between nodal processors is accomplished through the VCI). See Figure 11.8 for an example. The routing table entries for permanent virtual circuit service are populated via the network management system, and routing is not determined on a "per-call" basis as in X.25 SVC service. In the example of Figure 11.7

- The network layer in the PC at location x (typically part of the TCP/IP stack) looks in the routing table for the address associated with the destination application, known at the sending end by some logical name, say, y. The table indicates that the local router must be specifically addressed for the selected destination.
- Upon reception of the frame, the router checks its routing tables to determine the local DLCI needed to be appended to the frame in order to reach remote destination y.
- The router's data link layer places the information in a frame relay frame and sends it to node 1, with the DLCI label properly appended to the information.
- Node 1 recognizes the DLCI associated with an existing logical path through the network. If the frame is not in error and it has a valid DLCI, it is segmented into cells which are subsequently identified by a node-assigned VCI and other SAR mechanisms (some nodal processors forward entire frames without segmentation; the advantages of cell relay NNIs over frame relay NNIs are discussed in Section 11.2.14). The cells are sent on to node 2 and from there to node 3. Otherwise, it discards the frame.

⁷This acronym and the Fiber Channel Standard acronym introduced in Chapter 1 clearly refer to different concepts.

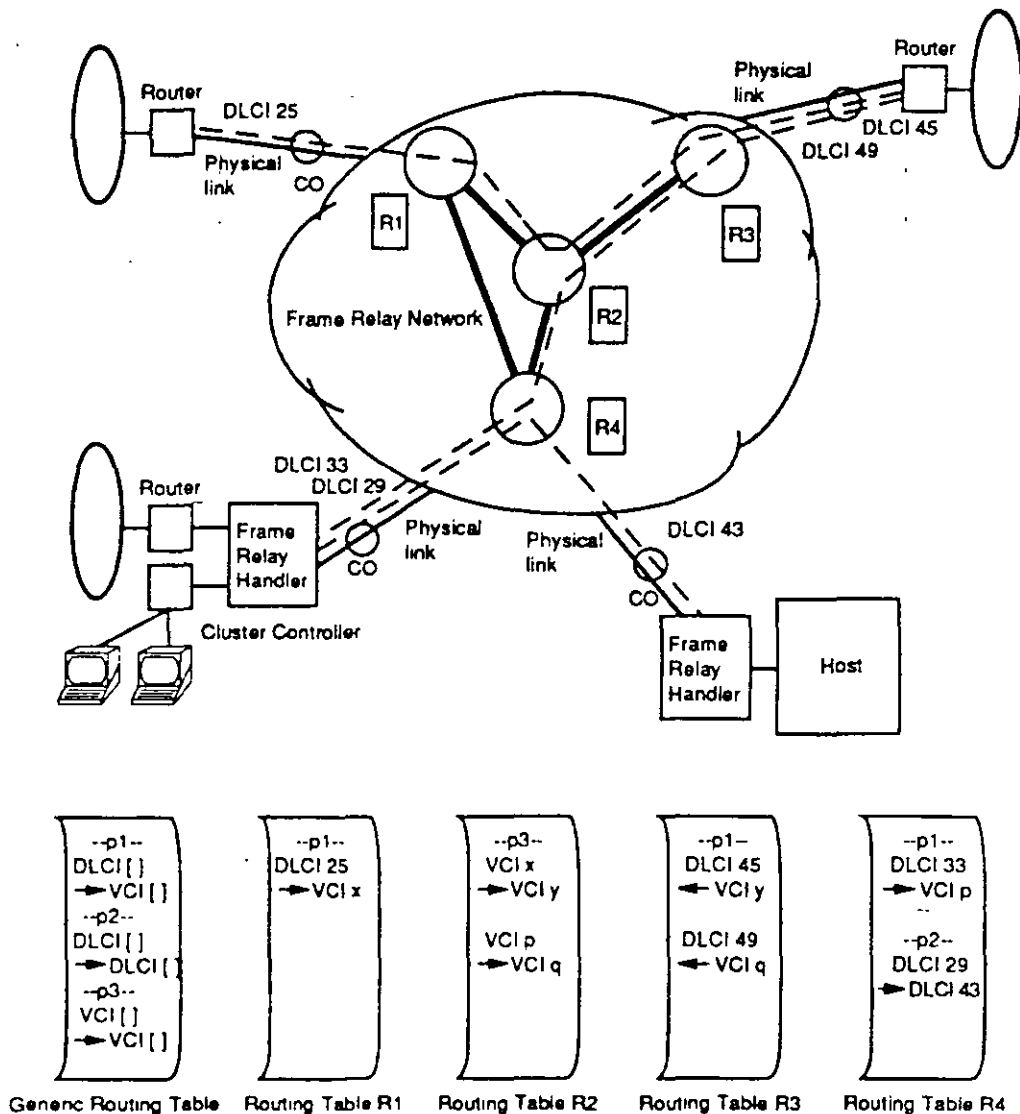


Figure 11.8 Mapping of DLCIs in a frame relay network.

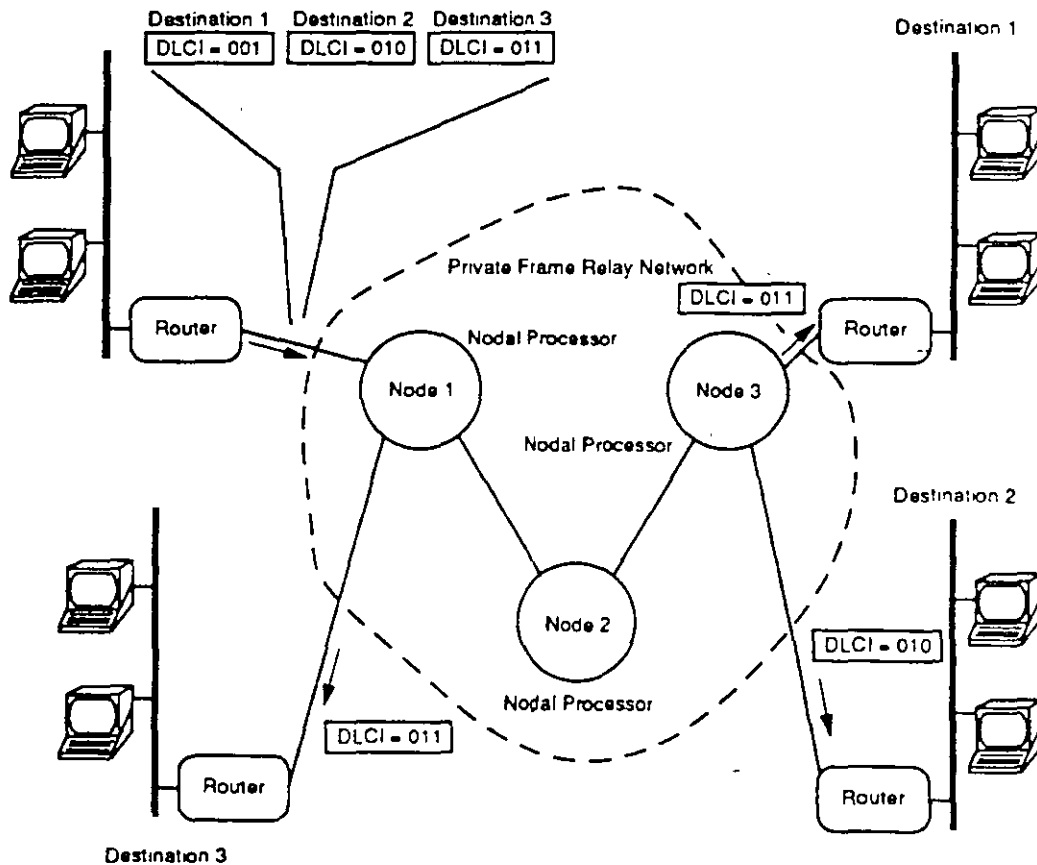
- Node 3 reassembles the cells into the frame using the VCI and other SAR mechanisms; the node then passes the frame over the access line that terminates in the equipment supporting application y.
- Upon receipt, the router forwards the information to the PC. In turn, the data is sent from the PC's data link layer to application y via the transport, session, and presentation layers.

The nodal processors do not have to read the variable-length frame to achieve switching; instead, the DLCI is sufficient to allow the edge processors to make the

necessary routing decisions. Figure 11.9 shows one physical realization of the protocol architecture of Figure 11.7 in order to illustrate routing of frames through a private frame relay network. The DLCI may be reused by virtual circuits that do not share one or both endpoints.

In fact, Figure 11.8 implies that the routing is more complex than the short discussion given above, since there is an interplay between DLCI, the cell's VCI (or equivalent vendor-proprietary indicator), and ultimate trunks. Three aspects of routing exist:

- An association between the locally significant DLCI and the cell's VCI (and the other way around). This occurs at origination and destination nodal processors.



Note: In actuality, DLCIs only have local significance and can be altered during transmission

Figure 11.9 Routing in a private frame relay network.

- Remapping of a VCI to another VCI. As discussed in Chapter 10, this occurs whenever there are ATM switching points (a specific VCI has no end-to-end significance if the virtual channel connection is switched; the VCI could remain the same end to end if the virtual connection is provided on a semi-permanent basis). This occurs at intermediate nodes.
- Association between a local DLCI and a DLCI of a user connected to the same nodal processor.

This in turn raises several questions pertaining to vendors' implementation of the frame relay/cell relay processors:

1. Are tandem nodal processors allowed, or must nodal processors be connected with direct trunks?
2. Does a tandem nodal processor have to reassemble cells back into frames, or are the cells relayed ("pipelined") directly as needed?
3. How does a nodal processor treat an incoming frame destined for a user directly connected to the same processors? Namely, is segmentation required?

These questions have a critical impact on the end-to-end delay of the frame relay network. Just the initial segmentation and the remote reassembly can already be significant; any intermediary reassembly impacts the grade of service further. Figure 11.7 depicted a scenario where the frame is segmented by the first processor handling it (node 1), and then sent downstream to a tandem processor (node 2), which accepts cells as such and transmits them along individually and discretely, without intermediary reassembly (Figure 12.3 shows an example of a segmentation process similar to the one discussed here). The frame is reassembled only by the destination node (node 3). Note that Figure 11.7 did not show SAR/CS function at node 2. This would happen if the nodal processor followed cell relay/switching ATM principles; such a processor would typically serve a variety of end-user streams, some of which could be digitized video, some could be digitized voice, and some could be frame relay information. Notice that, at the very least, the use of tandems implies having to incur the trunk transmission time twice. It is conceivable that if a nodal processor does not follow cell relay/switching principles, each frame must be assembled and disassembled by each nodal processor in the path.

11.2.11 Congestion Management

Users, LAN managers in particular, may worry about migrating traffic away from dedicated interrouter links they have used until now and onto a network based on high-speed packet technology. However, this is not an insurmountable problem, since frame relay has (in principle) a way to manage and control congestion [11.21]. The frame relay network composed of the nodal processors, private or public,

attends to this by first using congestion notification strategies and then by selectively discarding frames when needed to relieve congestion. Congestion control mechanisms are utilized to treat users fairly and to protect the network and users by localizing the congestion within the network.

The congestion notification takes place when a network node determines that it is becoming congested. It sets the forward explicit congestion notification (FECN) bit in the frames as it sends them to the destination router (see Figure 11.4). It also sets the backward explicit congestion notification (BECN) bit in the frames destined for the source router. Upon receipt of these frames, the source and destination routers are expected to initiate procedures to throttle back the traffic offered to the network. If congestion continues to increase despite using congestion notification, the network will begin to discard eligible frames and will put the congestion localization procedures into effect. The network of nodal processors selects frames for discard by looking at the discard eligibility (DE) bit in each frame to see if it has been set by the router. If it is set, then the network discards the associated frame. These procedures continue until the congestion subsides.

One issue, however, is if and how the router can enforce throttling back to the PCs originating the traffic. Hence, the important question to ask about a frame relay router, a nodal processor, and a carrier service, is whether or not the full congestion control apparatus specified by the standard is implemented in each of these devices. Congestion in public frame relay networks will be discussed in a later section.

11.2.12 Quality of Service

Multirouter networks using frame relay interfaces provide for proper frame sequencing and minimize the likelihood of misdelivered frames. This is accomplished by the nodal processors by using the connection-oriented PVC service. The same predetermined logical path is used by the nodes for all frames using the same DLCI on a given access interface. Recovery from errored frames is accomplished by the end-user equipment, since the network will detect and discard all errored frames.

11.2.13 Cell Relay

A cell is a fixed-length packet of user data (payload) plus an overhead, usually small, of 53 bytes or less. Cell relay is a high-bandwidth, low-delay switching and multiplexing packet technology (discussed in Chapter 10) which is required to implement a frame relay network in an efficient manner, particularly for mixed-media and multimedia applications. The international cell relay standard, ATM, was also discussed at length in the previous chapter. With cell relay, information

to be transferred is packetized into fixed-size cells. The term "cell relay" and the term "cell switching" are both used by the BISDN community [11.22]. (Some distinctions are possible, although we use the term interchangeably: cell relay can refer to an environment of PVCs where cells are simply relayed along the same path according to some static incoming-to-outgoing trunk association; cell switching can refer to a SVC environment where cells are dynamically switched according to a nearly-real-time incoming-to-outgoing trunk association created by the user via a signaling process.) Vendors tend to use the term "cell relay switch" (or node) when their equipment does not implement the CCITT ATM standard, but a proprietary standard. If the CCITT standard is implemented, they typically refer to the equipment as an "ATM switch."

Cells are identified and switched by means of a VCI/VPI label in the header. A number of functions of the layer 2 protocol are removed to the edge of the backbone, while "core capabilities" are supported directly by the cell switches, in addition to layer 1 functions (clocking, bit encoding, physical medium connection). Cells allocated to the same connection may exhibit an irregular recurrence pattern, since cells are filled according to the actual demand. Cell relay allows for capacity allocation on demand, so the bit rate per connection can be chosen flexibly. In addition, the actual "channel mix" at the interface can change dynamically. The cell header (such as the ATM's header) typically contains a label and an error detection field; error detection is confined to the header. The label is used for channel identification, in place of the positional methodology for assignment of octets, inherent in the traditional TDM T1/T3 systems. Cell relay is similar to packet switching, but with the following differences: (1) protocols are simplified and (2) cells (packets) have a fixed and small length, allowing high speed switching nodes; switching decisions are straightforward and many functions are implemented in hardware. Cell relay is critical to the deployment of frame relay, and only those nodal processors implementing it give the users the full advantages of the new technology.

One complication of using cell relay at the NNI instead of using frame relay at the NNI has to do with network discard options. A packet-based frame switching nodal processor (e.g., Netrix, BT Tymnet, and others) can discard a frame found to be in error or, in case of overload, a frame designated as eligible for discard by the user. In fast packet/cell relay platforms, the frame loses its identity in transit (since it is pipelined and only reassembled at the remote nodal processor, not an intermediary processor). The issue then arises of what to "throw away" in case of congestion: although a frame might have been segmented into, say, 30 cells, throwing away 30 random cells might imply corrupting the integrity of 30 frames, not (just) one frame, as might have been the intention of the network. As a practical solution, manufacturers of cell-based nodal processors put greater emphasis on designing their processors to avoid a congestion state rather than on how to deal

with congestion after it arises (systems with frame relay-based NNIs tend to do the opposite). Hence, these processors have enough buffering to absorb user's input data during an interval of network congestion instead of having to start forwarding of that data into the network, just to find later that some cells were lost, necessitating some remedial action (this is analogous to the airline industry principle of holding planes on the ground if congestion exists rather than launching the plane and then having to hold it while in transit).

11.2.14 Cell Relay Platforms for Frame Relay

"Frame relay" switches, already identified as nodal processors, need to be deployed in order to properly allocate bandwidth on a dynamic basis (alternatively, this can be achieved by using the frame relay facilities of a carrier). It is critical that a frame relay nodal processor support a dynamic view of the data being transferred through it; otherwise, the user will not obtain the full benefit possible with the technology. Without a cell-based switch, dynamic bandwidth allocation is not easily achievable. Figure 10.19 clearly indicated three modes of deploying frame relay in a corporate environment. The simplest way is to upgrade the routers with a frame relay board and retain the existing point-to-point infrastructure. This approach does not provide any consequential advantage over the existing environment [11.3, 11.4].

Frame relay describes an interface specification; nodal processor equipment vendors can still use proprietary internal protocols. This is similar to the X.25 case, where packet switches support a standardized interface, but use internal transport, routing, and flow control protocols. This forces a user wanting to establish a private network to use the equipment from the same vendor throughout the network. By contrast, the cell switching technology specified in ATM is open by design.

Many customers deploy high-capacity circuits to meet peak traffic (and performance goals); however, DS1 lines used exclusively for data are reported by some to be only loaded at 15% or less [11.2, 11.23, 11.24]. Dynamic bandwidth allocation requires the incorporation of cell relay in the nodal processor to handle communication over the trunks (another way would be to employ a frame switching nodal fabric, but the granularity or efficiency of the multiplexing can be significantly lower). Dynamic bandwidth allocation is done by designing the nodal processor from the ground up and eliminating any fixed-bandwidth constraints imposed by a TDM nodal architecture. No internal blocking should be allowed in the switch, and queuing must be eliminated or at least minimized. Vendors which have experimented with these architectures over the past few years are in a position to incorporate these ground-breaking architectures in the products they manufacture. TDM and cell relay can be viewed at two ends of a spectrum: it is not possible to take full advantage of cell relay if the node has internal and/or external TDM structural

restrictions. Because of the efficient multiplexing possible with cell relay, dynamic bandwidth allocation, so important to bursty users, particularly in the LAN router environment, is achieved.

Users with a mix of bursty traffic may find it advantageous to upgrade T1 equipment that uses time-division multiplexing to cell relay platform. The drawback of traditional TDM techniques is that users must allocate portions of the T1 circuit to individual channels, each supporting transmission of a specific data source. Since that bandwidth is allocated to only one user, it remains unused when it is not needed by that one user. Simply retrofitting a circuit-switched TDM-based T1 multiplexer with frame relay UNIs does not deliver the intrinsic benefits of frame relay, any more than simply replacing a standard router with one supporting frame relay while still using a point-to-point T1 line would. With circuit switching systems, the user has to preallocate some (or all) bandwidth to the frame relay service, whether that bandwidth will be used or not. An efficient utilization of the technology over a private backbone network requires a nodal processor with internal fast packet technology, namely, a processor which employs cell relay technology. In this case, letting all applications compete for the backbone bandwidth allows them to access the entire bandwidth when anyone has data to transmit, not only on the trunk side but also on the access side, since frame relay supports multiple PVCs on one physical link. On the other hand, a frame relay application on a circuit-switched multiplexer can only access some fraction of the total bandwidth.

When a network is properly designed, the full bandwidth of the frame relay interface can be available to any application that requires it for relatively long-duration bursts of data, as may be the case for interconnected LANs. These applications may require that the network nodes support bursts occupying the full access bandwidth for intervals of up to 10 seconds or more in order to support transfer of large files or interactive traffic.

There are economic advantages of using the combination of frame relay access and a cell-based backbone network. Using frame relay technology in conjunction with a cell-based backbone multiplexer as an upgrade of an existing private corporate backbone can be cost-effective, since the user can obtain from the backbone needed bandwidth on demand, rather than on a fixed (and inefficient) basis. The "saved" bandwidth is then available to other users of the same backbone, in theory minimizing the amount of new raw bandwidth the firm needs to acquire from a carrier in the form of additional T1 or FT1 links. As an alternative strategy, the service from a carrier can be used. Although nodal processors can also support non-frame relay traffic (e.g., voice or video), the two technologies together, cell relay and frame relay, promise to increase throughput between locations that have large amounts of bursty traffic.

One may wonder why it is beneficial to utilize segmentation of a frame into many (up to 133) cells and, consequently, why a cell-based platform is superior to a frame switching technology in the nodal processor. The explanation follows.

Frame relay is a data-only protocol; it is intended to support HDLC-type traffic (e.g. LAN packets). Its main focus is on data services. Cell relay (fast packet) switches, on the other hand, can also handle voice and video. For example, voice "frames" may be as small as one octet. Therefore, should a user's need be strictly LAN interconnection, then a frame switching technology with FRIs on both the access and on the trunk side, might, in fact, be superior in terms of performance. However, if the user also contemplates supporting voice and video, the best solution is to use a cell relay platform that supports FRI for LAN access, some other access protocol for voice and for video, and a cell method on the trunk side. As covered in the previous chapter, the ATM cell procedure is being introduced under B-ISDN to support all media, including voice, data, and video. Multiplexer vendors view frame relay as an access protocol; the cell relay/fast packet backbone is viewed as giving the user better control over the quality of service of the path and facilitating a mix of traffic [11.15].

11.3 BENEFITS OF FRAME RELAY

11.3.1 Deployment Approaches

About 75% of large (Fortune 1500) companies had a dozen or more bridges interconnected via DS1 lines in 1990 [11.8]. That number is likely to reach 100% by 1993. However, as discussed earlier, private networks based on dedicated lines tend to become impractical when there is a large number of remote data sources/sinks generating bursty traffic. The number of links grows quadratically with the number of sites to be interconnected. In addition, the interconnection capacity needs to be higher; this increase in the speed is often dictated by applications requiring more data to be transacted, as well by the number of users of the service of interconnection [11.17, 11.25–11.28]. This implies that fairly expensive links are required.

Five classes of solutions are available:

1. Instead of connecting all routers with a fully interconnected network, some routers are connected in tandem. While this reduces the number of links, it introduces extra end-to-end delay and increases nodal processing (requiring more machine cycles).
2. Deploy a private frame relay network using frame relay nodal processor(s). Instead of physical point-to-point links, this approach only requires connecting the routers to the nodal processor(s) with a single physical link. Connection between various routers is accomplished with PVCs (illustrated in Figure 11.3).
3. Use a PVC-based carrier-provided frame relay network. Instead of many physical point-to-point links, this approach only requires connecting the rou-

ters to the carrier's switch with a single physical link. Connection between various routers is accomplished with PVCs that are established at service subscription (see, for example, Figure 11.6).

4. Use a hybrid configuration. A cluster of sites employ private frame relay; other sites take advantage of public frame relay services.
5. Use a SVC-based carrier-provided frame relay network. Instead of many physical point-to-point links, this approach only requires connecting the routers to the carrier's switch with a single physical link. Connection between various routers is accomplished as needed by establishing a real-time SVC, which is in existence only for the duration of the session. Figure 11.10 illustrates this approach.

The evolution in the private environment involves using nodal processors which provide FRIs to the routers and use cell relay/ATM technology between nodes. (As indicated, it would be technically possible to also use frame relay techniques between the switches, as, in fact, two frame relay routers connected by a dedicated T1 link use, and as some vendor architectures based on traditional packet engines do, but this approach has not seen major commercial realization.) Although frame relay remains a connection-oriented service, there are still advan-

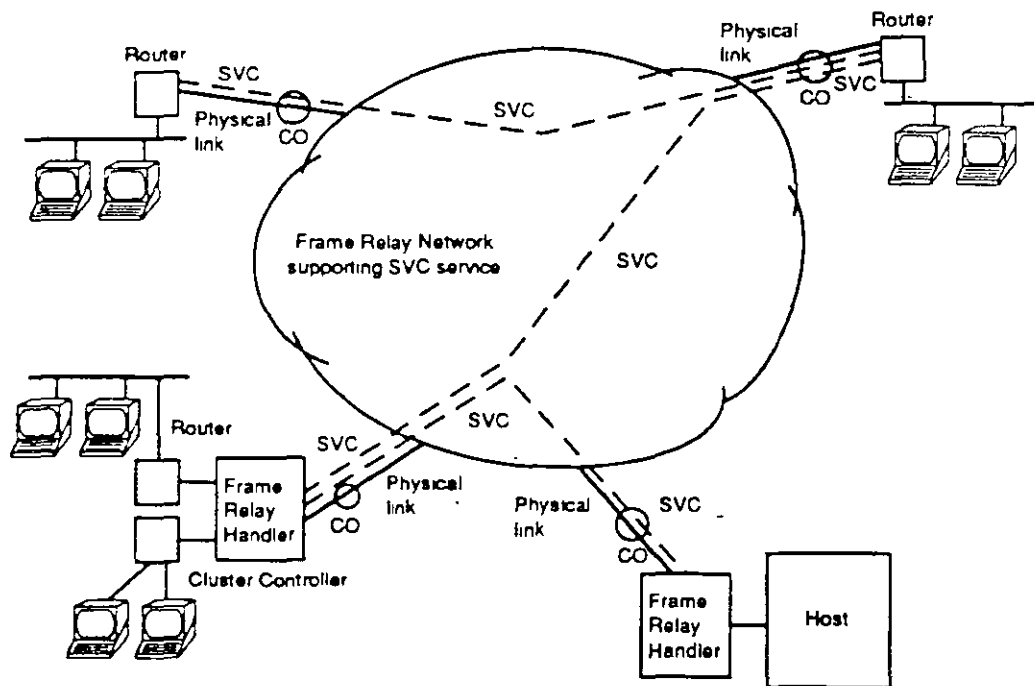


Figure 11.10 A SVC-based frame relay network.

tages in connecting LANs through frame relay rather than X.25 packet switching. In particular, when using X.25 service, routers had to encapsulate LAN traffic in the X.25 packet, imposing substantial performance degradation. Frame relay, on the other hand, incurs little overhead and allows for a number of protocols to be transported transparently [11.21].

In view of the growth in the population of LANs, carriers are readying themselves to provide public PVC-based frame relay data services that support high-capacity access/throughput, coupled with the universal access, survivability, economies of scale, and efficiency available through resource sharing. Chapter 1 provided information on the time frame of frame relay deployment in corporate networks. SVC-based frame relay can have some advantages, but it also has some limitations. First, the service may become available only later in the decade. Second, a user needing to send data to some remote user on another LAN may not be willing to incur the call setup time each time a session is required. The way some people have gotten around the setup time issue in packet-switched networks is to use long-duration SVCs; these are set up once and kept active for an appropriate amount of time, such as a day.

Some users may deploy hybrid frame relay networks. These users could use their own frame relay backbone connecting major sites and use a public frame relay network to connect secondary sites. Interworking issues have to be resolved before this approach can be realized in practice.

11.3.2 Benefits of Frame Relay in Private Networks

In the business and economic landscape of the 1990s it is prudent for the communication manager to look at networking solutions that will not have to be discarded after a couple of years to keep up with network growth or higher speed networking needs or technologies. Some nodal processors now on the market only support data. Other nodal processors support data, voice, and video. Because nodal processors based on cell switching utilize backbone facilities better than existing static channel banks or circuit switching T1 multiplexers (and also existing X.25 switches), the deployment of these mixed-media nodal processors in a private network benefits users that need to connect LANs over integrated backbones supporting a variety of other traffic. Users with LAN traffic only may choose data-only nodal processors. The financial advantage of a frame relay network becomes more marked when the number of routers is high (half a dozen to a dozen, or more) and when the distances between routers is considerable (hundreds or thousands of miles—if the routers are all located within a small geographic area, such as a city, a county, or a LATA, the economic advantage of elimination lines is less conspicuous). Table 11.8 summarizes some of the benefits.

Table 11.8
Some Benefits of Frame Relay

Port and link sharing
Bandwidth on demand
High throughput and low delay
Ease of network expansion
Ease of transition from existing router networks
Ease of transition from any existing network
Cohesiveness and symbiosis with LANs
Simplified network administration
Standards-based
Economic advantages (in a variety of situations)

Port and Link Sharing

Among the advantages of frame relay is its ability to statistically multiplex frames from multiple LANs at one location onto a single user network interface and associated communication link. Frames going to multiple destinations can share the same router port. The frame relay interface to the nodal processor provides for the end-user equipment the capability to place frames destined for different network endpoints onto the same network access line by using the DLCI mechanism. This accomplishes port sharing and allows each frame to have use of the entire bandwidth of the access line when there is a frame to be sent. Further efficiencies are gained on the backbone network interconnecting the nodal processors by combining the traffic from multiple routers onto the network trunks using efficient cell/ATM protocols. Instead of having to purchase more expensive multiport routers that otherwise would be needed, simpler point-to-point routers can be used.

Bandwidth on Demand

All of the bandwidth on the frame relay access interface can be available to the end-user system when it needs to transmit data across the network. The nodal processor can be optioned to accept, under conditions of slack, all the incoming traffic from one user up to the full access speed. Alternatively, the nodal processor can be optioned to accept up to some prenegotiated rate less than the full access speed, but more than the average user requirement. For example, the access line could be a T1 facility; the user's average input could be 128 kbps. The nodal processor could be optioned to accept an instantaneous input (over a short horizon, say, 10 seconds) of 512 kbps.

Improved Use of Bandwidth

Dynamic bandwidth allocation reduces the aggregate backbone transmission bandwidth the manager needs to secure from a carrier, which would otherwise be needed without it. Additional transmission resources contribute to a direct increase in transmission cost. For example, if six users need a maximum of 512 kbps, two T1s would be required under TDM, although their average rate may only be 128 kbps. With frame relay, one T1 ($5 \times 128k + 512k$) should suffice if the traffic is truly random. The upgrade of an existing backbone network with a private frame relay network can save, according to some early users, 20% of the total network bandwidth [11.29].

High Throughput and Low Delay

Since all of the bandwidth is available, high throughput is possible. Minimal delay is encountered within the backbone network, since there is little protocol processing required with frame relay. Cell-based nodal processors, particularly those employing the latest high-power microprocessors, are fast. The switching decisions based on the cell header are simple and direct. For example, some studies have shown that with a private X.25 with 56-kbps access, it took 4 minutes to transmit a benchmark file; with a frame relay network based on a T1 backbone and accessed with a 56-kbps line, the file could be transmitted in 45 seconds [11.29].

Ease of Network Expansion

Network expansion is straightforward with frame relay. Adding a new router to the network requires only the assignment of an access port on the network node, and the interconnection of the router with the network nodal processor via the appropriate transmission facility. The interconnection of the new router with the existing routers is accomplished by logically provisioning the network using a centralized network management system. The cell-based protocols used by vendors today could lend themselves to migration to the standard ATM cell format. This migration will permit the nodal processors to support some of the new high-speed services being developed by carriers. Thus, both frame relay and access to these higher speed services can be supported on the same backbone.

Ease of Transition From Existing Router Networks

Existing routers typically need only a software upgrade to implement the frame relay interface. Once this is done, the routers can be re-homed on the new frame relay backbone.

Ease of Transition From Any Existing Network

It was indicated above that many users still have networks that were put in place a few years ago. A frame relay network can easily be deployed, no matter what the user baseline is, as will be discussed in more detail later.

Cohesiveness and Symbiosis With LANs

Frame relay technology is similar in some aspects to the data link layer discipline of a LAN. Since the data need to leave the LAN and travel over a WAN, it is desirable to use a WAN technology that has an affinity with the LAN technology. This minimizes the amount of protocol conversion/remapping which would otherwise be needed.

Simplified Network Administration

Several recent studies have indicated that, when considering the true corporate cost of communication, 30% to 50% of the network expense corresponds to operation and administration efforts, commonly known as network management. Any tool or system that improves the way network management is done is a welcome and cost-saving feature. Administration in frame relay can be performed from a central network management and administrative system. Moves, changes, and additions to the network are typically handled through an automatic permanent virtual circuit provisioning capability within the system.

Standards-Based

The frame relay PVC UNI is an accepted and stable ANSI and CCITT standard, with wide support from both user equipment and network system vendors.

Vendor Support

Over three dozen vendors support frame relay. These vendors include router manufacturers, T1 multiplexer vendors, PAD developers, nodal processor and switch providers, and carriers. Frame relay routers cost from \$400 to \$15,000, depending on vendor and features, compared to a standard router. Nodal processors cost from \$20,000 to \$50,000, depending on vendor and features.

Saving Communications Costs With Nodal Processors

Perhaps the most significant benefit is that the use of frame relay over a cell relay platform can, in the right circumstances, save money for the multirouter LAN manager. Private networks based on dedicated lines become expensive when there is a large number of remote data sources/sinks generating bursty traffic. Initially, users may have employed dedicated lines operating at 19.2 kbps, 56 kbps, FT1, and T1 speeds to interconnect LANs; this increase in the speed is often dictated by applications requiring more data to be transacted, and by the number of users using the service of interconnection. As the number of LANs grows, this proliferation of T1 lines becomes impractical from both a cost and management perspective.

The topic of economics is always a complex issue. A detailed example of an economic analysis in the presence of nodal processors in an enterprise network was provided in Chapter 6, which the reader may wish to review at this point. It was shown that a private frame relay network can be cost-effective compared to a full mesh network at the same link speed. To undertake a cost analysis of a private frame relay network, the LAN manager should first calculate the network cost with traditional connectivity and then the cost of using frame relay technology. The process starts by determining the location of the sites to be interconnected. Sites can be identified by vertical and horizontal (V&H) coordinates. The V&Hs allow the manager to obtain the distance of all sites and, hence, the length of the required communication links. A T1 (or FT1) local loop must be costed out using the local exchange carrier's tariffs; these tariffs may be different at each site. Then the cost of the access facility between the serving CO and the interexchange carrier's POP must be calculated. Both the distance and the tariff may be site-dependent. Finally, the cost of the set of required long-distance T1 links can be obtained using the interexchange carrier's tariff.

The cost of the frame relay alternative is calculated as follows (refer to Figure 11.3 as a guide). One or more centrally located sites are chosen where the nodal processors will be located; the V&Hs of the nodes are noted (this choice may be subject to an optimization procedure). The cost of the nodal backbone network is determined by deciding what the required nodal connectivity will be, and then by costing out the transmission facilities (this will involve a T1 or FT1 loop, an access facility to the POP, the long-distance trunks, and the remote access and loops). The cost of the router access subnetwork is calculated next. This involves first determining which nodal processor each router will be homed to. Then the cost of the transmission link between the router and the nodal processor is calculated (this also will involve a T1 or FT1 loop, an access facility to the POP, the long link, and the remote access and loops). The (amortized) cost of the nodal processors and the routers' upgrade must also be included. The total cost is obtained by adding all of these factors.

Table 11.9 assesses the effectiveness of frame relay as a function of the geographic scope. This example is based on the baseline network discussed in Section 6.5, which should be reviewed at this point if there are questions about topologies, backbone mileages, etc. As can be seen from this table, the savings due to frame relay become less significant as the network gets geographically smaller: a nationwide (private) frame relay network costs 0.40 times as much as a mesh point-to-point solution, while a statewide network costs 0.61 times as much as a mesh point-to-point solution, implying less dramatic savings.

Table 11.10 undertakes a similar study, where there is only one nodal processor with no backbone (as some vendors of private frame relay networks are suggesting, particularly in the LAN interconnection context). Figure 11.11 depicts the topology, showing a centrally located nodal processor; all assumptions are similar to the previous case. The following conclusion emerges: for national networks, the one-node frame relay solution is slightly more expensive compared to the three-node solution; for regional and statewide networks the one-node solution is somewhat less expensive compared to the three-node solution ($\leq 20\%$ less).

Figure 11.12 depicts a one-node solution where the node is collocated with one of the routers (in a central location). Here, one less access line is required. However, as the figure shows, most other locations need longer access lines compared to the case of Figure 11.11. In this particular case, the total mileage is slightly reduced (from 4,800 for the baseline case of Figure 11.11, to 4,600), implying that there would be a small decrease in the total cost. The decrease is composed of three factors: (1) less IXC mileage, reducing the cost by an amount proportional to the mileage charge times the difference in mileage; (2) since there is an IXC "ramp up" on the T1 tariff of approximately \$2,100 (for the first mile), this expense disappears when one line is eliminated; and (3) one LATA line (premises to POP) is eliminated. The national, regional, large-state, and medium-state numbers are \$58,050, \$44,250, \$30,450, and \$23,550, respectively; this is an 8% to 10% reduction compared to the previous case. It should be noted, however, that this saving will become less important, diminish, and, in fact, even disappear as the number of routers increases, if these routers are widely dispersed.

If there were several routers clustered in one location, collocation of the nodal processor at that location would superficially appear beneficial, because multiple lines could be eliminated from that location to the centrally located nodal processor. However, since frame relay allows multiple PVCs on a single physical line, this saving is more apparent than real. Figure 11.13 shows one example (which we worked out on a paper plate with ruler—but we could as well have used trigonometry). The results depend on many factors: are the routers located on a circular path, an elliptical path? How many routers are collocated? The example demonstrates that, in fact, it would be better to locate the nodal processor at a central location. In Case A of the figure, the total mileage would be 10M (M = miles) if the nodal processor were centrally located, and 12M if it were placed in one router.

Table 11.9
Cost-Effectiveness of a Three-Node Five/Six Router Network as a
Function of the Network Geographical Scope

National Network (all IXC distances of baseline multiplied by 1.5)							
<i>Case*</i>	<i>Loops</i>	<i>Cost of POP</i> <i>Access</i>	<i>IOC</i> <i>Cost</i>	<i>Routers</i>	<i>FPS</i>	<i>Backbone</i> <i>Cost</i>	<i>Total</i>
A with T1s	\$4,500	\$6,200	\$139,200	\$ 833	\$ 0	\$ 0	\$150,733
B with T1s	\$6,750	\$9,300	\$205,200	\$1,000	\$ 0	\$ 0	\$222,250
C with T1s	\$2,250	\$3,100	\$ 31,800	\$ 917	\$2,350	\$20,700	\$ 61,117
D with T1s	\$2,700	\$3,720	\$ 41,400	\$1,100	\$2,350	\$20,700	\$ 71,970

Regional Network (e.g., northeast network) (Baseline—See Chapter 6 for assumptions)							
<i>Case</i>	<i>Loops</i>	<i>Cost of POP</i> <i>Access</i>	<i>IOC</i>	<i>Routers</i>	<i>FPS</i>	<i>Backbone</i> <i>Cost</i>	<i>Total</i>
A with T1s	\$4,500	\$6,200	\$100,800	\$ 833	\$ 0	\$ 0	\$112,333
B with T1s	\$6,750	\$9,300	\$148,800	\$1,000	\$ 0	\$ 0	\$165,850
C with T1s	\$2,250	\$3,100	\$ 25,200	\$ 917	\$2,350	\$16,200	\$ 50,017
D with T1s	\$2,700	\$3,720	\$ 32,400	\$1,100	\$2,350	\$16,200	\$ 58,470

Statewide Network, Large State (All IXC distances of baseline divided by 2)							
<i>Case</i>	<i>Loops</i>	<i>Cost of POP</i> <i>Access</i>	<i>IOC</i>	<i>Routers</i>	<i>FPS</i>	<i>Backbone</i> <i>Cost</i>	<i>Total</i>
A with T1s	\$4,500	\$6,200	\$ 62,400	\$ 833	\$ 0	\$ 0	\$ 73,933
B with T1s	\$6,750	\$9,300	\$ 92,400	\$1,000	\$ 0	\$ 0	\$109,450
C with T1s	\$2,250	\$3,100	\$ 18,600	\$ 917	\$2,350	\$11,700	\$ 38,917
D with T1s	\$2,700	\$3,720	\$ 23,400	\$1,100	\$2,350	\$11,700	\$ 44,970

Statewide Network, Medium State (All IXC distances of baseline divided by 4)							
<i>Case</i>	<i>Loops</i>	<i>Cost of POP</i> <i>Access</i>	<i>IOC</i>	<i>Routers</i>	<i>FPS</i>	<i>Backbone</i> <i>Cost</i>	<i>Total</i>
A with T1s	\$4,500	\$6,200	\$ 43,200	\$ 833	\$ 0	\$ 0	\$ 54,733
B with T1s	\$6,750	\$9,300	\$ 64,200	\$1,000	\$ 0	\$ 0	\$ 81,250
C with T1s	\$2,250	\$3,100	\$ 15,300	\$ 917	\$2,350	\$ 9,450	\$ 33,367
D with T1s	\$2,700	\$3,720	\$ 18,900	\$1,100	\$2,350	\$ 9,450	\$ 38,220

A = Five routers without frame relay

B = Six routers without frame relay

C = Five routers with frame relay

D = Six routers with frame relay

(*)Refer to corresponding example in Chapter 6 for all assumptions and topologies.

Table 11.10
 Cost-Effectiveness of a One-Node Five-Router Network as a
 Function of the Network Geographical Scope

National Network (All IXC distances of baseline multiplied by 1.5)							
<i>Case*</i>	<i>Loops</i>	<i>Cost of POP</i>		<i>IOC Cost</i>	<i>Routers</i>	<i>FPS</i>	<i>Total Cost</i>
		<i>Access</i>					
A with T1s	\$4,500	\$6,200		\$139,200	\$833	\$ 0	\$150,733
C with T1s	\$2,250	\$3,100		\$ 55,200	\$917	\$783	\$ 62,250
Regional Network (e.g., northeast network) (Baseline)							
<i>Case</i>	<i>Loops</i>	<i>Cost of POP</i>		<i>IOC</i>	<i>Routers</i>	<i>FPS</i>	<i>Total Cost</i>
		<i>Access</i>					
A with T1s	\$4,500	\$6,200		\$100,800	\$833	\$ 0	\$112,333
C with T1s	\$2,250	\$3,100		\$ 40,800	\$917	\$783	\$ 47,850
Statewide Network, Large State (All IXC distances of baseline divided by 2)							
<i>Case</i>	<i>Loops</i>	<i>Cost of POP</i>		<i>IOC</i>	<i>Routers</i>	<i>FPS</i>	<i>Total Cost</i>
		<i>Access</i>					
A with T1s	\$4,500	\$6,200		\$ 62,400	\$833	\$ 0	\$ 73,933
C with T1s	\$2,250	\$3,100		\$ 26,400	\$917	\$783	\$ 33,450
Statewide Network, Medium State (All IXC distances of baseline divided by 4)							
<i>Case</i>	<i>Loops</i>	<i>Cost of POP</i>		<i>IOC</i>	<i>Routers</i>	<i>FPS</i>	<i>Total Cost</i>
		<i>Access</i>					
A with T1s	\$4,500	\$6,200		\$ 43,200	\$833	\$ 0	\$ 54,733
C with T1s	\$2,250	\$3,100		\$ 19,200	\$917	\$783	\$ 26,250

A = Five routers without frame relay


C = Five routers with frame relay

*Refer to corresponding example in Chapter 6 for all assumptions on tariff and to Figure 11.11 for baseline topology.

location (Case A'). What happens if some routers are clustered? If separate lines to a central location (Case B) were used, it would still take 10M of circuit; however, since multiple PVCs can be put on a single link (assuming that the performance issue was appropriately handled), 6M of circuits is sufficient (Case C). Locating the nodal processor at the location with several routers (Case B') only cuts the circuit length to 7M, which is more than with the centrally located nodal processor.

It is difficult to draw general conclusions about the cost-effectiveness of private frame relay networks (except that they are cheaper than full mesh networks), because the problem is highly multidimensional (50 to 100 dimensions or more).

Part A
 Ten access lines
 Total circuit mileage: 12,800 miles



Part B: Adding a router
 Five T1 access lines
 and one nodal processor
 Total circuit mileage: 4,800 miles




Figure 11.11 Economics of a single-node frame relay network.

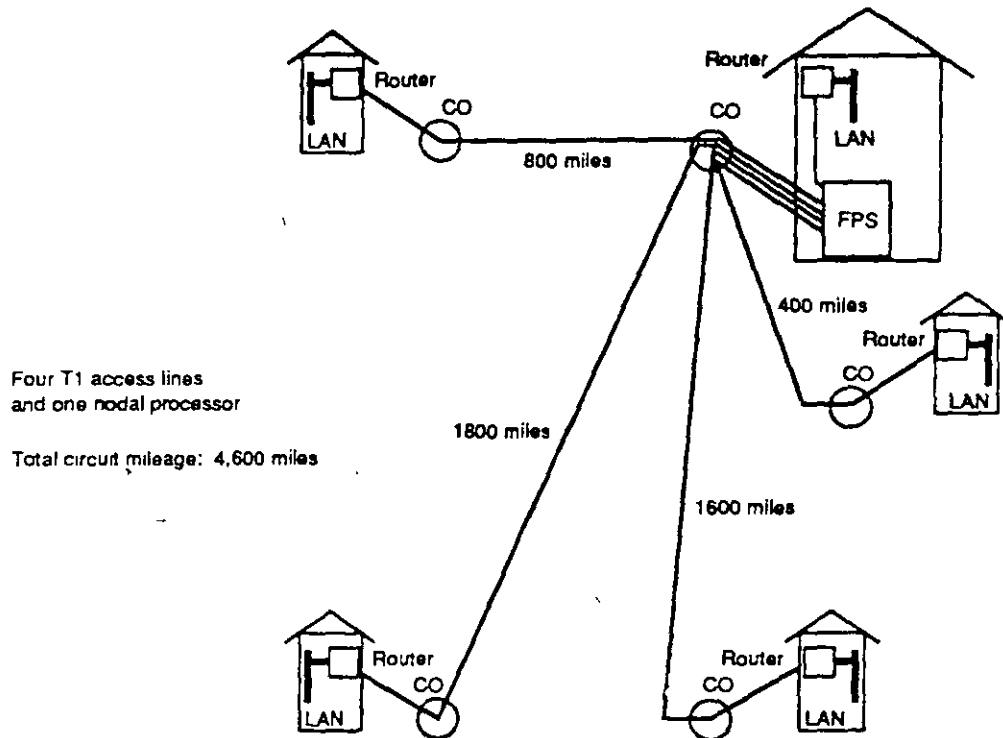


Figure 11.12 Collocating a single node with a site with a router.

Also, in addition to transmission costs, some of the other factors that LAN managers and network designers take into account in selecting a network architecture include network reliability, network availability, ease of network management, compatibility with open international standards, ease of network upgradeability, initial costs, migration costs, growth capabilities for both traffic and sites, integration with embedded base, and vendors' technology support. However, recurring transmission charges continue to be a visible component of any calculation assessing the desirability of a network redesign. Where is a frame relay network particularly cost-effective? In trying to draw some general conclusions, the answer is when one or more of the following apply:

- There is a large number of remote sites (half a dozen or more). The larger the number of sites, the more cost-effective frame relay will be.
- The remote sites are highly dispersed (at least regionally or nationally). The higher the combined network mileage, the more cost-effective frame relay will be. This implies that national-scope networks are reasonably suited to a private frame relay technology.
- The traffic is highly bursty. This occurs when traffic leaving the router is small and occurs in just a few instances during the day (not all traffic leaving

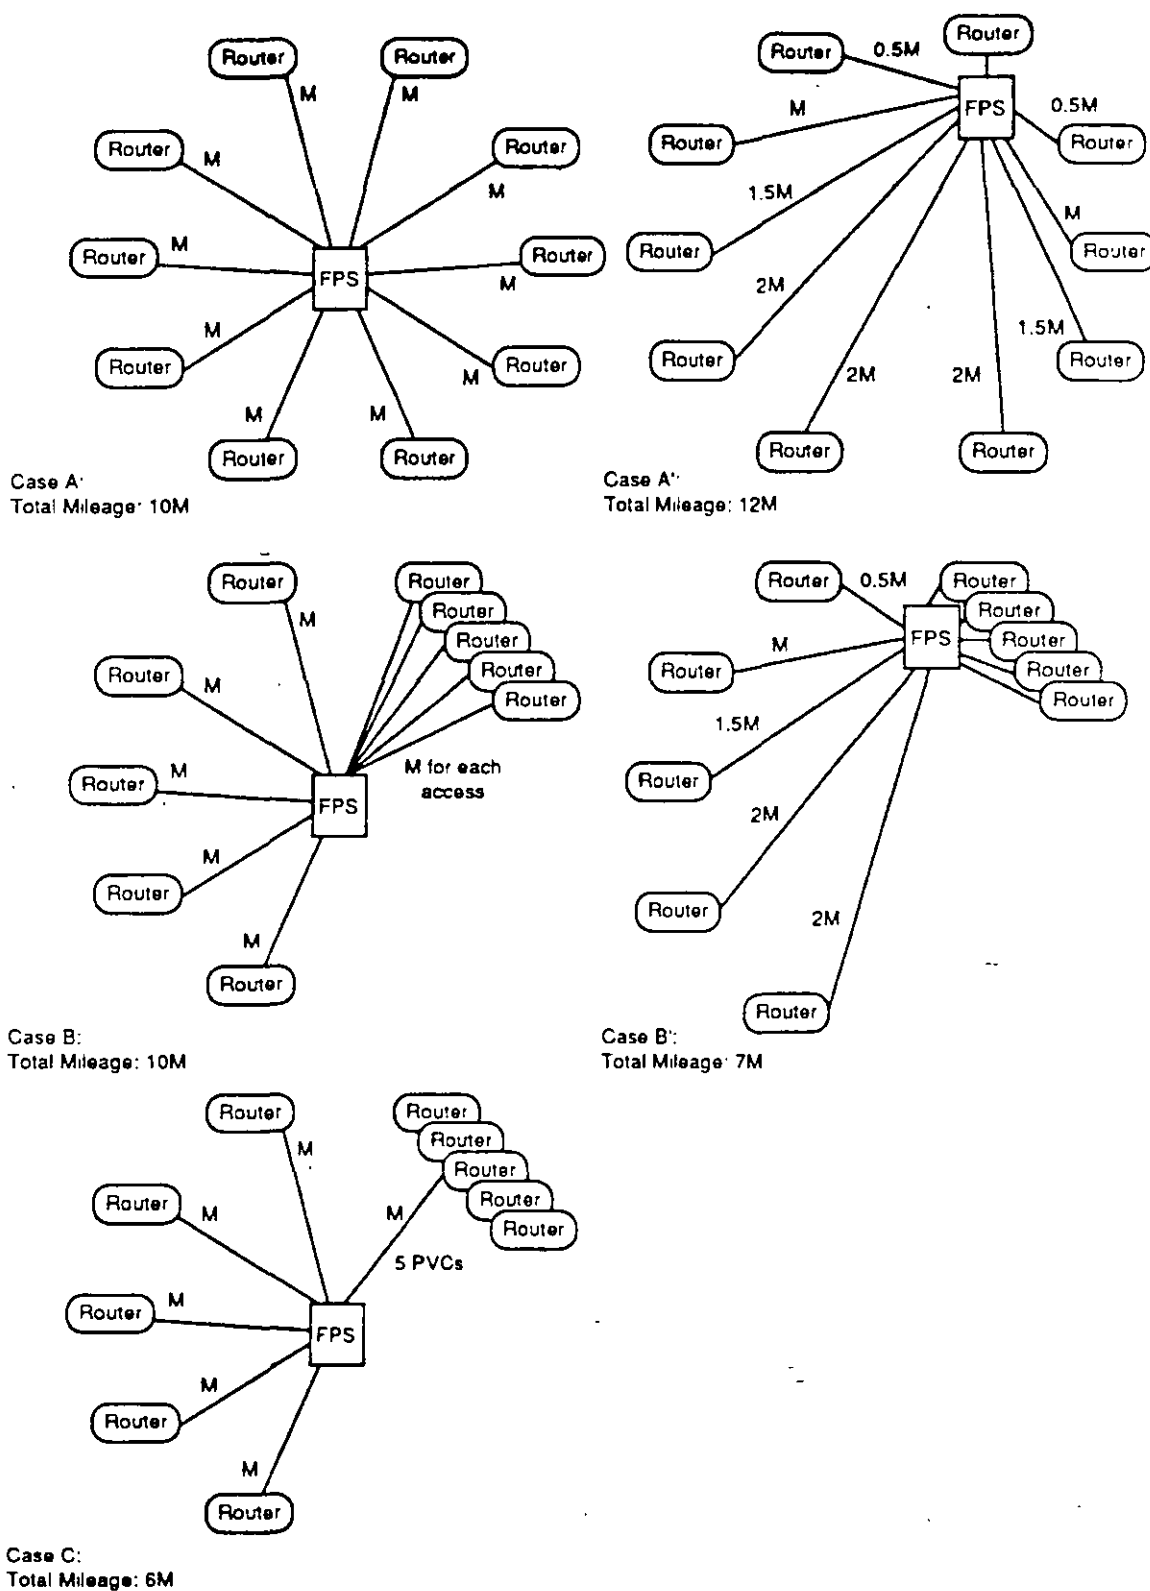


Figure 11.13 The geometry of locating the nodal processor.

- a router is bursty—as the utilization of the router approaches 100% the traffic will become more predictable). Alternatively, this occurs when the applications transact large amounts of data at discrete instances (e.g., file transfer).
- There are multiple LANs at a site which, for whatever reasons, are not interconnected with each other, and yet all need to reach the network. This takes advantage of frame relay's ability to place multiple logical channels over a single physical channel. If there are D remote destinations and N unconnected LANs, $D \times N$ virtual channels are required (if the various local LANs were already interconnected with bridges, then the number of required virtual channels is only D).
 - New sites/routers are added to the network with relatively high frequency (say, once every six months or more frequently).
 - The links between the routers have relatively low speed (FT1), and more bandwidth appears to be required. Upgrading the FT1 mesh topology links to full T1 facilities may be very expensive. Nodal processors can increase throughput for less money than would otherwise be needed.

From a carrier's perspective, frame relay service will impact private line services the most; less impact is expected on public packet networks (since these either address themselves to lower speeds, or to international destinations) and on SMDS services (since these provide higher speed, are connectionless, and support true switching capabilities).

11.3.3 Benefits in Public Networks

Some carriers and vendors have made commitments to frame relay, others carriers have made plans to deploy cell relay, and several carriers are pursuing both technologies (including the seven BOCs). Some view the two approaches as complementary, others as competitive. Frame relay service and cell relay service are designed to meet different objectives, and hence have evolved in different directions.* A categorization in the public network environment is as follows [11.30]:

- Frame relay is a medium- to high-speed (DS0-DS1) data interface for private networks which is being implemented at this time. Some observers believe that frame relay may in fact have market importance at the DS0 level.
- Cell relay/switching is a high- or very-high-speed switching service capable of supporting public BISDN and SMDS networks. Cell switching supports 155-Mbps, 622-Mbps, and eventually higher SONET/SDH rates.

*"Cell relay service" refers to providing a cell/BISDN UNI, not a cell in the NNI, as we have discussed so far. (In the private network environment, corporate backbone network switches supporting LAN applications typically use a FRI UNI and a cell NNI.)

Frame relay carriers provide intra-LATA, inter-LATA, and international service. U.S. frame relay networks can be classified as private networks (discussed above), IXC networks, VAN networks, and LEC networks. Some of the benefits of using public frame relay networks are covered next.

Major Reduction in Transmission Costs

Based on current tariffs, the transmission cost can be reduced as much as 70% compared to a mesh network. This topic will be revisited in a later section.

Low Startup Cost

The only expenditures in beginning to use a public network are the upgrade of the routers to support the FRI; this can be accomplished for about \$1,000 per router. Some carriers even supply a frame relay-ready router to get the user going. In other cases, the networks provide frame relay PAD functions, so that traditional devices (such as SNA terminals) can be supported directly. In contrast with private frame relay networks, there are no expenses for nodal processors and the communication backbone infrastructure.

Ability to Support a Variety of User Equipment

LANs, terminals, front-end processors, and even X.25 equipment can be supported by the public networks.

Ability to Transmit Instantaneous Bursts Exceeding the Throughput Class

At the establishment of a PVC, the user can select a throughput class. A public frame relay network allows the user to exceed, on an instantaneous basis, the selected class (up to the maximum access speed) without further negotiation with the network. If the network has spare capacity at that point, it will transport these additional bursts. For example, if the throughput class (also called "committed information rate") is 512 kbps, and the user has a T1 access line, short-duration bursts up to 1.544 Mbps can be presented to the network. A few vendors have announced plans to offer frame relay products supporting access speeds of 45 Mbps (these include Coral Network Corporation, Newbridge, and StrataCom).

Multiple service providers may have to be involved when frame relay services cross LATA or national boundaries. Although standardization of frame relay protocols makes the interworking between local exchange carriers, interexchange carriers, and international carriers feasible in principle, administrative, billing, and

operational issues make the delivery of a nationwide multicarrier service a nontrivial effort, likely to take time [11.17].

11.3.4 Other LAN Interconnection Solutions

It was pointed out in Chapter 1 that many services could be applied to a communication problem, and that the number of such solutions is increasing. In addition to private frame relay, a number of other evolving technologies could be applied to the multirouter LAN interconnection problem, implying that the user should not feel compelled to instantly redesign the network to deploy frame relay. However, each approach has advantages and disadvantages.

The fact that direct connections with many dedicated T1 lines are expensive could be mitigated with the use of FT1 lines in lieu of the T1 lines. While this solution lowers the transmission cost, it also greatly impacts performance, since a router link no longer has access to the 1.544-Mbps bandwidth when it needs to send an instantaneous burst; instead, it only has access to a fraction, which could be as low as 1/24th. Another option would be to use switched T1.

The use of public frame relay would bring some of the benefits of private frame relay. Initially, however, the public service may be limited only to the major cities (40 by the end of 1993). Then, unless the local exchange carriers also support the service in the access segment, a dedicated T1 line to the interexchange carrier's POP will be required; this could be expensive, although, in some cases, the frame relay carriers absorb the cost. In addition, there will be usage charges, which are not present in the private network solution. Network management will also be more difficult, although capabilities are being put in place.

SMDS is also available for LAN interconnection. SMDS supports a UNI at 45 Mbps (T3); this may be appropriate for CAD/CAM and other imaging applications. T3 service, however, requires the installation of a fiber to each LAN location, unless CO-based multiplexing of T1 lines into T3 lines is used.

TDM-based T1 multiplexers supporting a traditional backbone could also be used, but in order to guarantee the grade of service to a very bursty user, a large portion of bandwidth must be statically allocated to each router; this would accommodate short, intensive bursts. The problem with this approach is that the large amount of allocated bandwidth is not utilized, except on a short basis, and yet cannot be made available to any other user when not being put to useful work. This results in the need for more transmission bandwidth, contributing to a direct increase in transmission cost. In some cases, however, this bandwidth may in fact be available for "free" and could therefore be used. This could be the case, for example, where a user replaced five T1 lines for a T3 line costing just as much, making 23 T1 lines available for additional usage.

As a specific example, assume that a user had three major sites with three multiplexers, all of which are connected with four T1 lines, each costing, say, \$3,000

a month. Assume that more applications are coming on board, requiring some additional bandwidth. One option would be to replace the existing multiplexers with nodal processors and retain the four T1 lines; better bandwidth utilization may suffice to carry the new applications: This would cost \$150,000 (\$50,000 each), but would also need staff retraining and new management tools. A second approach would be to retain the three multiplexers, upgrade them for T3 usage (say, \$15,000), and replace each of the four T1 lines with a T3 line costing, say, \$15,000 a month. Here the incremental communication cost would be \$9,000 a month, implying that it is cheaper to retain the TDM multiplexers for up to 15 months ($\$9,000 \times 15 + 15,000$) compared to a nodal processor replacement. In addition to the fact that vast amounts of additional "free" bandwidth is available, no staff retraining and no new management tools are necessary.

11.4 FRAME RELAY PROTOCOLS AND STANDARDS

This section provides more details on frame relay standards.

11.4.1 CCITT View

One of the goals of the recent CCITT work has been to align some of the available data communications protocols and offer recommendations for a set of efficient network services that can then be built upon by user equipment. One aspect of these new services is the separation of the control information from the user information into logically separate (but not necessarily physically separate) paths, as is the case in ISDN. Another aspect of the goal was to simplify the network protocols. Simplification, as provided by frame relay, allows the realization of services that are superior in terms of delay and throughput than existing services, since there is much less per-frame processing on the part of the network.

In most existing networks (e.g., X.25 networks, SNA networks, and analog voice networks), there is no clear end-to-end distinction between the logical control path and the data path. A close coupling between information and control limits the flexibility needed to support new services and new signaling and transport needs. Separation, the goal of frame relay as originally conceived, has the following benefits [11.31]:

- There is the potential for the integration of signaling for voice, data, and other media. This is important for future multimedia services.
- Since the information path does not have to support control, its logic can be substantially simplified. This implies that the hardware will be cheaper and faster.
- Independent optimization of the two paths can be accommodated.

The major characteristics of ISDN's frame relay are out-of-band call control

and link layer multiplexing. Under ISDN, all the new packet services, particularly the SVC services, have the following characteristics:

1. All control procedures, if needed, are performed in a logically separate manner (channel) using protocol procedures that are integrated across all telecommunications services. Consequently, Recommendations I.430 and I.431 provides the layer 1 protocol for the control channel; Recommendations I.441 and I.451 are extended as the layer 2 and 3 protocols, respectively (LAP-F/Q.922 and Q.933). In the case of PVCs, no real-time call establishment is necessary and any parameters are agreed on at subscription time.
2. The data transfer procedures share the same layer 1 functions based on Recommendations I.430 and I.431. The data transfer may use any channel on which the user implements at least the lower part (the core functions) of I.441* (LAP-F). I.441* is the generic protocol terminology of I.122-1988 (namely, I.441 appropriately extended to frame relay).

The separation can occur in a number of ways, including (1) on a physically separate interface, and (2) on another logical channel within the same interface (e.g., a time slot or the D-channel).

CCITT I.122 recognizes two frame relay implementations: a switched implementation under the auspices of ISDN, using the CCITT Q.933 protocol for call setup, and a PVC implementation. The PVC does not require call setup and call termination, but is obviously not as efficient in resource utilization as SVC. I.122 is an access standard; on the trunk side no restriction is imposed (same as in ISDN). As discussed, the trunk side is typically cell-based.

The term *relay* implies that the layer 2 data frame is not terminated and/or processed at the endpoints of each link in the network, but is relayed to the destination, as is the case in a LAN. In contrast with X.25-based packet switching, in frame relay the physical line between nodes consists of multiple data links, each identifiable by information in the data link frame. Unlike the (X.25-based) X.31 packet-mode services, frame relay services (SVC in particular) integrate more completely with ISDN services because of the out-of-band procedures for connection control.

In X.25, multiplexing is achieved through the use of logical packet layer channels; hence, the network layer provides switching. In frame relay, switching is accomplished at the data link layer, and link layer multiplexing is used in the user's plane to facilitate sharing of bandwidth among multiple users. Switching in the data link layer is achieved by binding the DLCIs to routing information at intermediary nodes to form a set of network-edge to network-edge logical paths [11.31]. Multiplexing is done through the statistical multiplexing of different data link connections on the same physical channel, as specified in LAP-F Core/Q.922. Frame relay service is based on the frame structure originally employed by the ISDN D-channel

LAP-D, which provides statistical multiplexing of different user data streams within the data link layer (layer 2).

Put slightly differently, a feature of frame relay is to have the virtual circuit identifier, currently implemented in the network layer of X.25, positioned at the data link layer so that switching can be accomplished more easily. In the X.25 environment, when a data call is established the virtual circuit indicator is negotiated and used for the duration of the call to route packets through the network. In a layered protocol environment, layer $n + 1$ protocol information is enveloped inside layer n information. The network layer routing indicator is enveloped within the layer 2 headers/trailers, which must be processed before it can be exposed. This processing involves more than just stripping the header/trailer; for example, it involves error detection and correction. In LANs, the routing of the data units is accomplished directly at layer 2; the data frames are supplied with a 48-bit destination address, which is readily available and which is used to physically route the data to the intended destination. Also, there is no error recovery in a LAN as a packet flows by a station on its way along the bus or ring. In frame relay, only the lower sublayer of layer 2, consisting of such core functions as frame delimiting, multiplexing, and error detection, are terminated by a network at the user-network interface. The upper procedural sublayer of layer 2, with functions such as error recovery and flow control, operates between users on an end-to-end basis. In this sense, a user's data transfer protocol is transparent to a network.

Limiting layer 2 functionality to the core functions implies that the user's FRI functions can be implemented in hardware rather than in software, improving throughput/delay characteristics at the interface. Frames with error are identified and discarded, and the network boundary entities or, more commonly, user equipment are expected to recover via upper layer protocols (with cleaner fiber-based circuits, BER is much improved). The data link layer core functions are

- Frame delimiting, alignment, and transparency.
- Frame multiplexing/demultiplexing using the address field.
- Inspection of the frame to ensure that it consists of an integer number of octets prior to zero bit insertion or following zero bit extraction.
- Inspection of the frame to ensure that it is neither too long nor too short.
- Detection of transmission errors.


Frame relay implements only the core functions on a link-by-link basis; the other functions, particularly error recovery, are done on an end-to-end basis. Indeed, the capabilities provided by the transport layer protocol accommodate this transfer of responsibilities to the boundaries of the network. On the user side, beyond the frame relay interface with the network, the user can employ any end-system-to-end-system protocol.

Protocol standardization work followed the publication of I.122 in 1988. Additional supporting standardization was needed before the service could be offered in a carrier/vendor-independent fashion. As initially defined by CCITT, core functions do not include flow control. The addendum to ANSI's T1.606 now defines congestion management strategies; it covers both network and end-user mechanisms and responsibilities to avoid or recover from periods of congestion. Additional standards remained to be developed in 1992 and beyond, particularly in support of interconnection of frame relay networks from different carriers (i.e., national and/or international interworking) and SVC service.

Family of Services

I.122-1988 describes a family of frame relay services. The purpose of defining a family of services, instead of a single service, was to provide a degree of flexibility in order to choose the best service based on the requirement of the application. Elements of this family are distinguished by the difference in degree of protocol support. Another way of looking at this is the different levels of protocol termination at the network edges after call establishment. Figure 11.14 depicts different protocol breakpoints, or points at which a network can terminate the protocols in support of the requested bearer service [11.31].

CCITT, in Recommendation I.122 ("Framework for providing additional packet mode bearer services"), describes three frame relay services.⁹ Refer to



UP: Upper Part

Note: I.441* now has materialized into Q.922/LAP-F

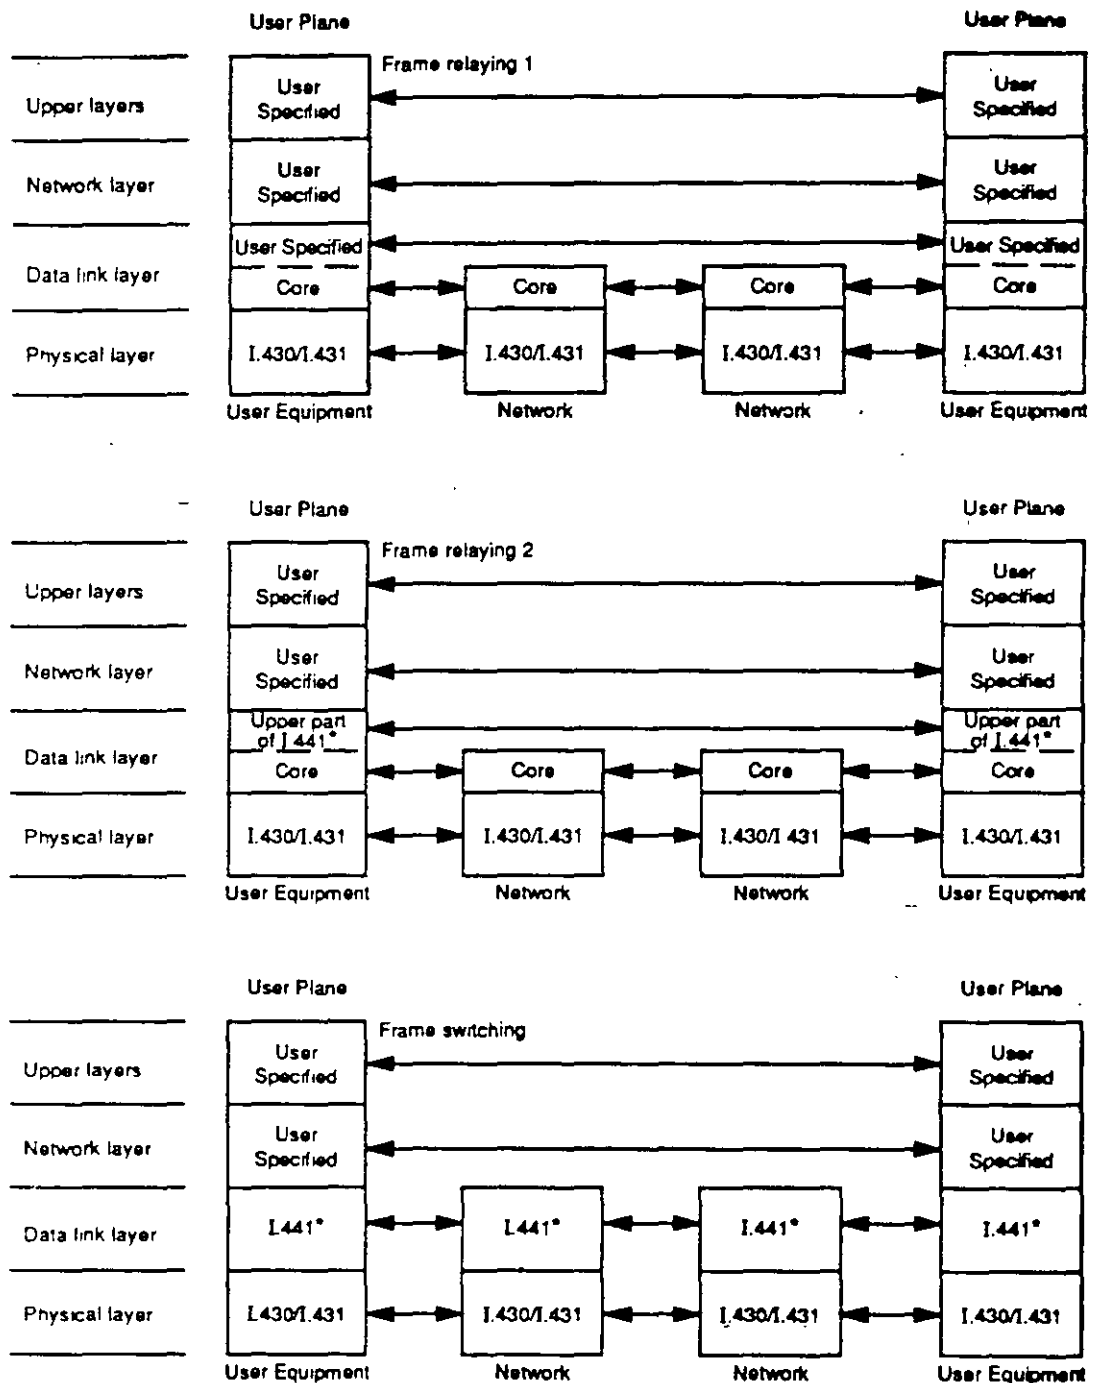
Figure 11.14 I.122 frame relay protocol specification.

⁹The 1992 version of I.122 no longer distinguishes between FR-1 and FR-2 services.

Figure 11.15, which retains the original I.122 protocol terminology (I.441* is I.441/Q.921, extended to cover I.122 requirements; this is Q.922).

1. Frame relaying 1 (FR-1) (no functions above core data link functions are terminated by the network; if needed, such functions are terminated only end to end). The basic service provided is the unacknowledged transfer of frames from S/T network boundary to S/T network boundary. Any user-selected end-to-end data link layer above the core functions can be used. More specifically:
 - It preserves frame order as given at one S/T reference point if and when the frames are delivered at the other end. (Since the network does not terminate the upper part of I.441*/LAP-F, sequence numbers are not kept by the network; networks should be implemented in a way that, in principle, frame order is preserved.)
 - It detects transmission, format, and operational errors.
 - Frames are transported transparently (in the network); only the address and FCS field may be modified (some bits being defined in the address field for congestion control may also be modified).
 - It does not acknowledge frames (within the network).
2. Frame relaying 2 (FR-2) (no functions above the core data link functions are terminated by the network; I.441* (i.e., LAP-F) upper functions are terminated only at the end points). The basic service provided is an unacknowledged transfer of frames from S/T to S/T reference point. The upper part of I.441* is used end to end; however, the network only supports the core functions. More specifically:
 - It preserves frame order as given at one S/T reference point if and when the frames are delivered at the other end. (Since the network does not terminate the upper part of I.441* (i.e., LAP-F), sequence numbers are not kept by the network; networks should be implemented in a way that, in principle, frame order is preserved.)
 - It detects transmission, format, and operational errors.
 - Frames are transported transparently in the network; only the address and FCS field may be modified.
 - It does not acknowledge frames (within the network).
 - Normally, the only frames received by a user are those sent by the distant user.
3. Frame switching: the full Recommendation I.441* (i.e., LAP-F) protocol is terminated by the network. The user's data link layer protocol must be I.441* (i.e., LAP-F), and is fully terminated by the network (only the network layer and the upper layers are end to end).


In summary, Figure 11.16, from I.122, shows the partition of the data link layer in the frame relay environment. For both FR-1 and FR-2, the network sup-



Notes

- (1) Core functions are now defined in LAP-F Core (Q.922 Annex A) - They are based on LAP-D principles
- (2) I.441* - Q.922

Figure 11.15 Frame relay 1, frame relay 2, and frame switching, original I.122 view.



	User Equipment		Network Equipment	
	Px	Py	Pw	Pz
Frame relay 1	I.441* Core	User specified	I.441* Core	No partner
Frame relay 2	I.441* Core	I.441* Upper	I.441* Core	No partner
Frame switching	I.441* Core	I.441* Upper	I.441* Core	I.441* Upper

I.441* Core = LAP-F Core
 I.441* Upper = LAP-F *minus* LAP-F Core

Figure 11.16 Protocols stacks for various frame relay services

ports only the "core" aspects of the data link protocol I.441* (i.e., "Core Part of I.441*" or "LAP-F Core"). The user's equipment in FR-1 has a protocol partner with the network supporting the "Core Part of I.441." What the equipment supports end to end above core aspects is a user's option. Hence, the "remainder" of the data link layer functions above the core functions and the upper layers need to be defined by a set of user-provided peer-to-peer protocols. In FR-1, the network has no knowledge of the end-to-end protocol. The user's equipment in FR-2 terminates the full data link protocol (i.e., I.441*, which is composed of the "Core Part of I.441*" plus "Upper Part of I.441*"). The user equipment must have a protocol partner with the network supporting the "Core Part of I.441," and it must have a protocol partner end to end supporting the balance of the data link layer, namely, the "Upper Part of I.441*" (upper layers are user-defined). In frame switching,

the user equipment must have a full protocol partner with the network supporting the entire data link layer, I.441*.

The core functions are sufficient to transfer data during the data transfer phase (i.e., after the call has been established either in real time—SVC—or by an administrative process—PVC); only frames with valid format and valid address are delivered. Data link layer functions not specified by the frame relay service (FR-1, FR-2, or frame switching), as well as the network and upper layer functions, are transparent to the network, being implemented end to end in the end-systems. For example, in addition to the data link layer multiplexing, which is provided by the network over the UNI, a user may also choose to perform network layer multiplexing. This implies that a given frame relay connection supports data for multiple end-users; this multiplexing, however, is transparent to the network [11.31].

At the UNI (seen from the network), there are no significant differences between FR-1 and FR-2. Differences are visible, however, to the end-systems' network layer: depending on the data link layer used, different OSI services are provided to the network layer. In FR-2 and frame switching, the network layer services are specified by Q.922; for FR-1, the data link layer service is specified according to the user's choice of protocol. Because of this choice, there can be differences in performance between FR-1 and FR-2.

To use a frame relay network, the user's protocol-specific frames are encapsulated in the Q.922 Annex A frames, as shown in Figure 11.17. Any data link layer protocol with error recovery (HDLC, SDLC, LAP-B, LAP-D, LLC) can be encapsulated and transmitted over the network. Such encapsulation must be done by the user's equipment.

11.4.2 ANSI Frame Relay Standardization Efforts

Several documents have recently been issued by ANSI in reference to frame relay service in the U.S. [11.32–11.37]. These were identified earlier in Table 11.7. Two key standards are T1.606 and T1.618.

The data transfer phase of the frame relay bearer service is defined in T1.606-1990. This document specifies a framework for frame relaying service in terms of user-network interface requirements and internetworking requirements [11.38]. Both interworking with X.25 and interworking between frame relaying service is included in this standard.

The protocol needed to support frame relay is defined in T1.618-1991 (LAP-F Core). The protocol operates at the lowest sublayer of the data link layer and is based on the core subset of T1.602 (LAP-D). The frame relay data transfer protocol defined in T1.618/LAP-F Core is intended to support multiple simultaneous end-user PVCs, possibly using different protocols within a single physical

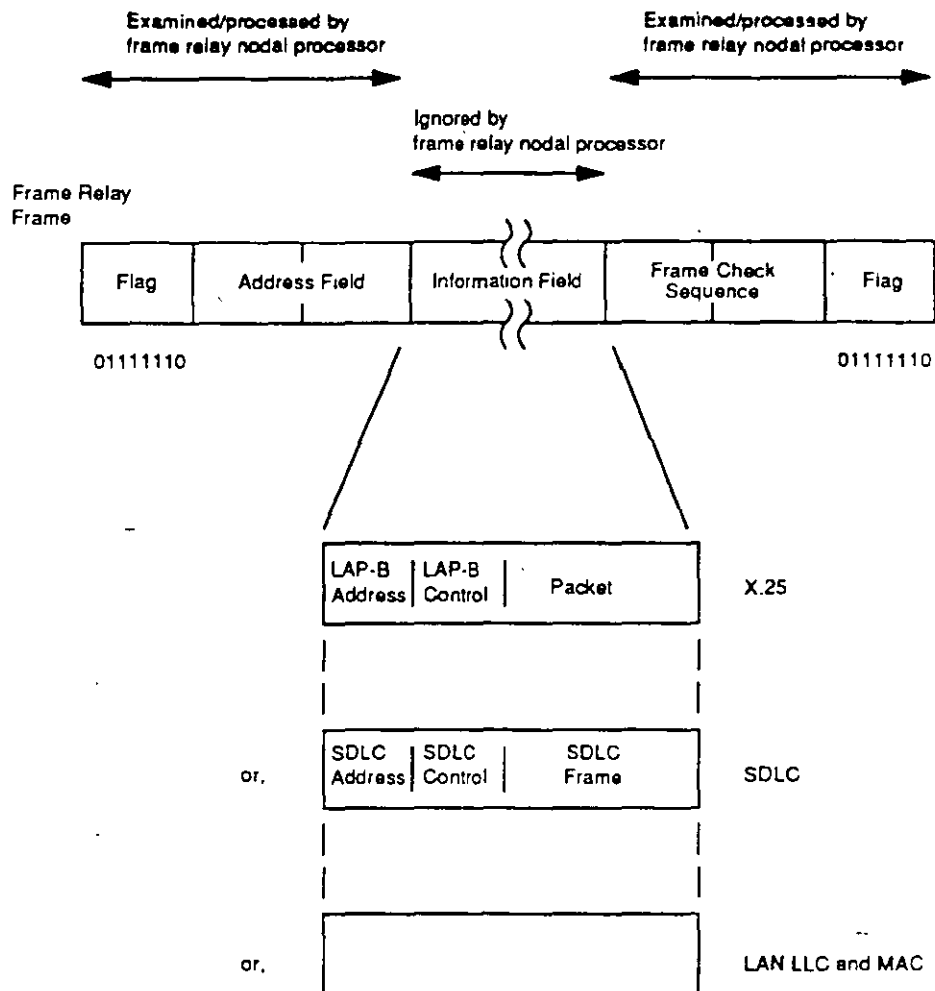


Figure 11.17 Encapsulation of user's protocols within frame relay.

channel. This protocol provides transparent transfer of user data and does not restrict the contents, format, or coding of the information, or interpret the structure.

Frame Relay Frame Structure

The frame relay frame format was shown in Figure 11.4. The field shown in the figure are described below.

Flag Sequence. All frames start and end with the flag sequence consisting of one 0 bit followed by six contiguous 1 bits and one 0 bit. The flag preceding the address field is defined as the opening flag. The flag following the FCS field is defined as

the closing flag. The closing flag may also serve as the opening and must be able to accommodate reception of one or more consecutive flags.

Address Field. The address field (more precisely, routing label) consists of at least 10 bits over two octets, as illustrated in Figure 11.4, but may optionally be extended up to four octets. To support a larger DLCI address range, the three-octet or four-octet address fields may be supported at the user-network interface or the network-network interface based on bilateral agreement.

Control Field (C/R). There is no control function for frame relay core services. The field is not used by the network and is passed transparently between user equipment for application-specific uses. This bit is used in protocols such as LAP-D to indicate that the frame is a command or a response.

FECN. This bit is set to 1 by the network to notify the user receiving the frame that the frame has been delivered through a congested path in the network. This implies that insufficient network resources are available to continue handling the traffic at the current rate. Two actions could ensue (depending on the user's equipment capabilities):

1. The inbound traffic, if any, from the destination (i.e., the traffic going in the opposite direction of the received frame) should be temporarily reduced.
2. The destination should be willing to enter a "hold-on" or "wait" state, since traffic may arrive at longer intervals than otherwise expected.

BECN. This bit is set to 1 by the network to notify the user that traffic sent in the opposite direction to the frame with the bit set may pass through a congested path. Consequently, the sending equipment should reduce its inbound traffic to the destination, if there is any. Figure 11.18 depicts the operation of the FECN and BECN.

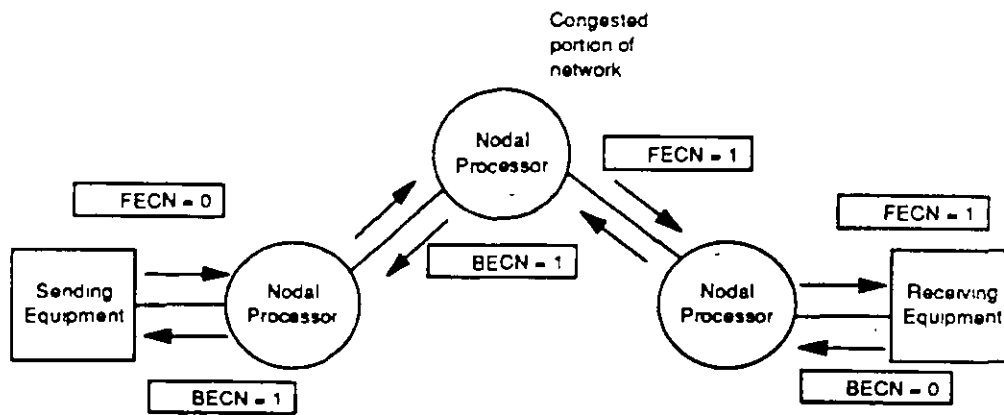


Figure 11.18 FECN and BECN action.

EA. EA is used as an expansion bit indicating that the DLCI is longer than 10 bits.

DE. DE is the discard eligibility bit set by the user to inform the network that in case of congestion this frame can be dropped before other frames not so indicated are touched.

Frame Relay Information Field. The frame relay information field follows the address field and precedes the frame check sequence. The contents of the user data field consists of an integral number of octets (no partial octets). The default information field size to be supported by networks is 262 octets; other values are negotiated between users and networks and between networks. The support of a maximum value of 8,189 octets is suggested for applications such as LAN interconnection to prevent the need for segmentation and reassembly by the user equipment (however, the usage of a cell-based nodal processor runs counter to this philosophy). The frame length can be variable. Table 11.18, shown later, depicts some of the maximum frame lengths supported by various vendors. Since the 16-bit FCS specified for frame relay can detect errors in frames of lengths up to 4,096 octets, some are recommending that only this maximum be actually allowed; otherwise the network cannot even detect errored frames [11.6].

Frame Checking Sequence Field. The FCS field is a 16-bit CRC sequence used to determine the integrity of the information.

Transparency. A transmitting data link layer entity must examine the frame content between the opening and closing flag sequences (address, frame relay information, and FCS fields), and must insert a 0 bit after all sequences of five contiguous 1 bits (including the last five bits of the FCS) to ensure that a flag or an abort sequence is not simulated within the frame. A receiving data link layer entity must examine the frame contents between the opening and closing flag (five contiguous 1 bits).

Order of Bit Transmission. The octets are transmitted in ascending numerical order. Inside an octet, bit 1 is the first bit to be transmitted.

Invalid Frames. An invalid frame is a frame that

1. Is not properly bounded by two flags (e.g., a frame abort), or
2. Has fewer than five octets between flags (note: if there is no information field, the frame has four octets and the frame will be considered invalid), or
3. Contains more than 8,193 octets between flags, or
4. Does not consist of an integral number of octets prior to 0 bit insertion or following 0 bit extraction, or
5. Contains a frame check sequence error, or
6. Contains a single octet address field, or
7. Contains a data link connection identifier that is not recognized by the network.

If the frame received by the network is too long, the network may either

1. Discard the frame.
2. Send part of the frame toward the destination user and then abort the frame,
or
3. Send the frame toward the destination user with invalid FCS.

Selection of one or more of these behaviors is an option for designers of frame relay network equipment, and is not subject to further standardization. Users cannot not make any assumption as to which of these actions the network will take. In addition, the network may optionally clear the frame relay call if the number or frequency of excessively long frames exceeds a network-specified threshold. Invalid frames are discarded without notification to the sender. No action is taken as a result of that frame.

Frame Abort. Receipt of seven or more contiguous 1 bits is interpreted as an abort, and the data link layer ignores the frame currently being received.

11.4.3 Industry Efforts

1990 saw a number of vendors backing an interim joint frame relay specification in an effort to ensure some degree of interoperability of new products then being developed. Digital Equipment Corp., Cisco Systems, Inc., Northern Telecom, Inc., and StrataCom, Inc., jointly developed the frame relay specification on which product development could be based until national and international standards become available [11.39, 11.40]. Eventually, over 65 vendors agreed to support this de facto standard [11.41]. More complete ANSI/CCITT standards are now available. In fact, most aspects of this interim specification found their way into the ANSI standards. The need to offer interoperable frame relay products is critical, and vendors realize that users may not be willing to deploy technologies that lock them in with systems that could become obsolete in a year or two. Agreement on frame relay implementation specifications facilitates the emergence of equipment from a variety of vendors, allowing flexibility in user choices [11.40]. Vendors are trying to avoid the implementation problems that were experienced in the early 1980s when X.25 packet switching products started to enter the market. Incompatible implementations of X.25 still abound to this day.

The early joint specification was based on the ANSI standard, but it had some additional management features and broadcasting [11.40]. For example, it included capabilities for congestion control; it also supported automatic reconfiguration of devices with a frame relay interface and the ability to detect faults. Features included [11.42]

1. Support for a global addressing convention to identify a specific end-device.

2. Multicast capability to send frames to all devices that belong to a "multicast group."
3. Flow control for preventing congestion in a frame relay network.
4. Extensions to the LMI.
5. Asynchronous status updates (asynchronous notification by the network to the user's equipment of a change in DLCI status).

This specification defined these enhancements in the form of a new protocol and a new set of messages to make the configuration and maintenance of PVCs easier. The protocol describes a LMI which is applicable between the network and the user's equipment (i.e., at the UNI). The LMI transfers messages that provide notification by the network to the user of the presence of an active DLCI, notification of the removal or failure of a DLCI, and real-time monitoring of the status of the physical and logical link between the network and each user device. In other words, the LMI solves the issue of a "keep-alive signal" between the network and the user's equipment. It also provides capabilities for downloading logical link addresses from the network to the user's equipment. Also, as indicated, a multicast facility for ease of address resolution by bridges and routers is included [11.43]. (Additional aspects of LMI are discussed in the next section.) These features are now included in the ANSI standards.

This vendor cooperation led to another development. On 15 July 1991, the Frame Relay Forum held its initial annual meeting. At that time, 52 companies joined the Forum; membership has increased since then. The Frame Relay Forum was formed to promote the acceptance and implementation of frame relay based on national and international standards. Membership in the nonprofit organization is open, and organizations may participate either as voting members or as observers [11.44]. The Forum has three working groups:

1. Market Development and Education.
2. Technical.
3. Interoperability and Testing.

The Market Development and Education Committee has as a goal the development of the market for frame relay products, services, and applications. The Technical Committee provides a liaison to the standards groups and related technical organizations, such as ANSI/ECSA, CCITT, ETSI, and the Internet Engineering Task Force. The Interoperability and Testing Committee aims at promoting efficient and effective methods of testing and certification of frame relay conformance and interoperability. It works with manufacturers of test equipment, with public frame relay carriers, and with third-party test laboratories. The Forum has adopted an implementer's agreement which identifies the guidelines vendors should follow in developing frame relay equipment. It also has contracted with the NIUF to develop a software test set based on the implementer's agreement, so that

prospective vendors, carriers, and users can undertake conformance testing. Approximately 500 items need to be tested to verify compatibility.

The importance of conformance and conformance testing cannot be over-emphasized. Already, in early 1992, carriers testing frame relay equipment were reporting that “many vendors have improperly implemented frame relay protocols” [11.45]. Frame Relay Forum efforts underway at press time included network-to-network interface implementation agreements, SVC specification, multiprotocol interconnection of data terminals, and, possibly, a standard for packetized voice over a frame relay network.

11.4.4 Carrier-Specific Extensions and LMI

Many portions of the vendors’ extensions for network management, particularly the LMI’s local in-channel signaling, have subsequently been incorporated in the ANSI standards (ANSI T1.617 Annex D, Additional Procedures for PVC’s Using Unnumbered Information Frames). The LMI specification describes a protocol and associated procedures operating at the UNI to handle network management functions. The features of a network that supports LMI include notification to the user of the addition, deletion, and presence of a PVC in the network, and notification to the user of end-to-end availability of a PVC [11.6]. Vendors are working on implementing support of Annex D. In addition, a standard to support X.25 over a public frame relay network has evolved. The LMI protocol consists of an exchange of messages between the user and the local access node of the network.

The LMI protocol is based on a polling scheme—the user’s equipment (router) polls the network to obtain status information for the PVCs defined over a given UNI interface. The user device issues a Status Enquiry message and the network responds with a Status message. Figure 11.19 provides an illustration of the process. The LMI uses a connectionless data link protocol based on Q.921/LAP-D, making the procedure easy to implement. At layer 3, Q.931 messages are used, as in ISDN.

Annex D of T1.617 specifies procedures for the following tasks:

- Addition or deletion of a PVC.
- Status determination (availability/unavailability) of a configured PVC.
- Local in-channel signaling for link reliability errors.
- Local in-channel signaling for link protocol errors.

Data Link Layer

The LMI data link layer conforms to a subset of LAP-D. Only unnumbered information frames are used. The poll bit is set to 0, and the control field is coded as 00000011. The DLCI is set to 0 (see Figure 11.20).

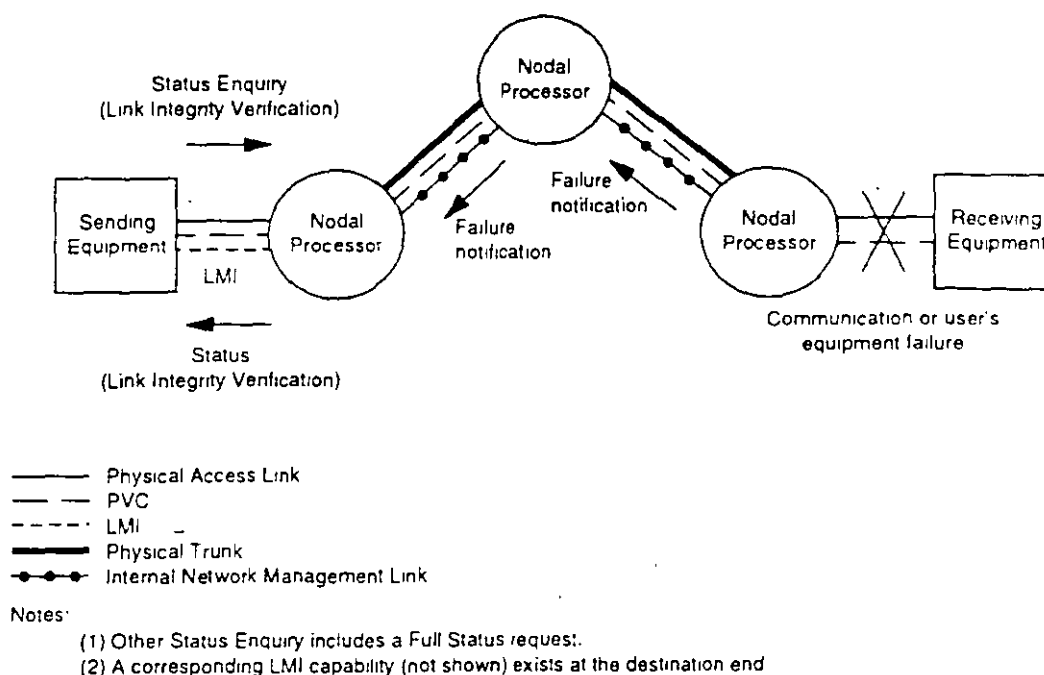


Figure 11.19 LMI environment.

The DLCI is specified in bits 3 through 8 of the second octet, and bits 5 through 8 of the third octet (the leftmost bit is bit 8; the rightmost bit is bit 1). The message field must contain the LMI Protocol Discriminator set to 00001001 in the LAP-D frame; it is used by the user-network call control to distinguish this message from other messages. The Call Reference is set to the dummy 00000000. A Locking Shift field is also required; it is used to identify codesets (currently only codeset 5 is supported).

Management Layer

This layer consists of two facets: (1) the format of the message field, including Information Elements; and (2) the message functional description.

An entire LMI message always fits an entire LAP-D frame. The Information Elements have specific formats. The formats are specified by the bit mappings for various functions (these are not further described here; see, for an example, [11.6, 11.36]).

The Link Integrity Verification Status Enquiry from the user and the Status message from the network allow both the user and the network to determine link reliability errors (physical faults) and protocol errors. The Full Status Report has

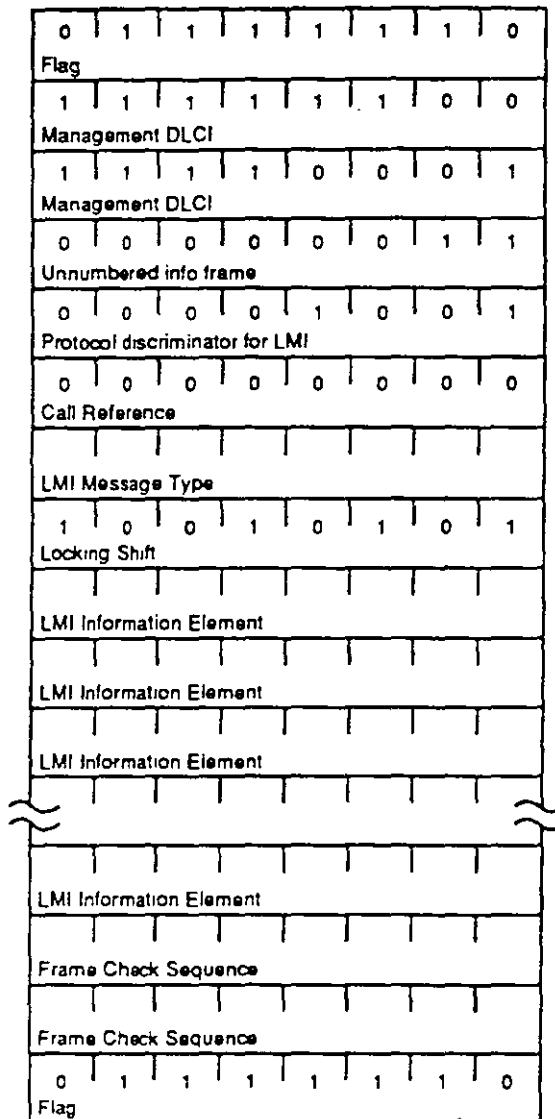


Figure 11.20 Data link layer format of LMI messages.

a PVC Status Information Element that allows the user to detect the addition of a PVC, the deletion of a PVC, the availability of a configured PVC, and the unavailability of a configured PVC. A user's frame relay device (e.g., a frame relay capable router) periodically issues a Status Enquiry message for the network's Full Status Report to determine when a PVC has become active or inactive. The reports are exchanged using DLCI 0. Full Status Reporting (PVC Status and Link Integrity Verification Information Element) is employed to report communication or remote user equipment failure to the local user. This procedure can also be used to signal a trunk or nodal processor failure.

The LMI messages and some related Information Elements are shown in Table 11.11. These procedures are driven by a set of parameters that are established at subscription. Table 11.12 depicts some key parameters. Additional procedural details, not covered here, are required to undertake the network management functions.

On the topic of network management, it is worth noting that protocol analyzers supporting frame relay were beginning to appear in 1992 (from a few vendors; however, they were initially rather expensive (\$15,000 range).

11.5 IMPLEMENTING FRAME RELAY IN A PRIVATE CORPORATE NETWORK

Users of dedicated LAN internetworking links may want to examine traffic loads to determine if frame relay and cell relay/fast packet will be economically beneficial. Users with little LAN interconnection traffic but with considerable traditional data traffic may be better off using a TDM-based T1 multiplexer, while those with higher LAN volumes may want to replace TDM multiplexers with processors (or multiplexers) supporting frame relay over a cell relay platform.

To maximize the benefit of frame relay technology in a private network environment without having to incur large communication charges (i.e., for dedicated T1 links between sites), the service needs to be provided by a backbone network configured with nodal processors that support dynamic bandwidth allocation via cell relay. The use of a router equipped with a frame relay interface

Table 11.11
LMI Messages

Messages:	
STATUS	Sent from the network to user device in response to a Status Enquiry. Has Message Type field of 01111101.
STATUS ENQUIRY	Used by the router or frame relay device to request status information. Actual configuration and status information is contained in the Information Elements. Has Message Type field of 01110101.
Information Elements:	
REPORT TYPE	Used to indicate either the type of enquiry requested by the user's frame relay device or the contents of the Status message returned by the network. It can be a Full Status or a Link Integrity Verification only.
LINK INTEGRITY VERIFICATION	Used to exchange sequence numbers between network and user equipment on a periodic basis to indicate to each other that they are active and operational.
PVC STATUS	Present in a Status message and is sent by the network to notify the user's frame relay device of the configuration and status of an existing PVC; the PVC is identified at the LMI UNI by the DLCI.

Table 11.12
LMI Parameters

Full Polling Cycle: This parameter describes the number of polling cycles between Full Status Reports. It is set by the user and has range of 1 to 255, with a default value of 6.

Error Threshold: Number of reliability or protocol errors before a PVC or a user device is declared inactive. It is set by both the network and the user and has a range of 1 to 10, with a default value of 3.

Monitored Events Count: This parameter specifies the size of the window that is employed by the network or user to determine if a PVC or user device is active. After a PVC or device is declared inactive, the network waits a number of successful poll cycles specified by this parameter before it is declared active again. It has a range of 1 to 10, with a default value of 4.

Link Integrity Verification Timer: This parameter indicates how frequently the user should send a Status Enquiry. It is set by the user. It has a range of 5 to 30 seconds, with a default value of 10.

Polling Verification Timer: This parameter indicates the interval of time the network should wait between Status Enquiry messages; if no messages are received, the network posts an error. It is set by the network. It can range from 5 to 30 seconds and has a default value of 15 seconds.

over a dedicated T1 link is not advantageous compared to a traditional non-frame relay solution. Some early users of frame relay took this route, but they are now finding that the nodal processor is an integral component of a dynamic bandwidth network: a backbone network can multiplex the traffic of one user with that of other users, realizing the economic advantages of bandwidth sharing, much the same way an X.25 private packet network provided such economic efficiencies for low-bandwidth users.

Therefore, (1) the availability of a cell backbone and (2) the addition of frame relay interface capability to user's equipment (usually with a plug-in card plus appropriate software) will facilitate deployment of the new technology for LAN interconnection usage within a corporation. Each user device will require only one physical connection to the network instead of multiple connections. In addition, data transmission over these permanent virtual circuits can vary dynamically as needed (up to the maximum access speed, i.e., 1.544 Mbps).

11.5.1 Implementation Steps

It is straightforward to migrate from the current router network configuration to a frame relay-based network solution. There are two main areas that need to be addressed:

- Network nodes.
- Router upgrades.

Backbone Networking Nodes—Installation

It is necessary to install nodes that support frame relay user-network interfaces and use cell-based backbones. Migration from the existing network to the frame relay configuration can be done in an organized, step-by-step fashion. This will minimize disruptions to LAN applications and end-users by permitting the changes to be made on a scheduled basis [11.13].

Router Upgrades to Support the Frame Relay Interface

Upgrade of the router is needed to implement the frame relay interface to the network node. This is usually in the form of a low-cost software upgrade offered by most router vendors. Costly hardware replacement is not usually necessary, since the existing communication chips on the routers are typically reusable for frame relay. Even more significant is the fact that the end-user applications do not have to be modified to accommodate frame relay.

11.5.2 Migration From Existing Baseline

Different users find themselves in different situations. Some still have unintegrated networks without backbones (generation 1). Others have a classical backbone network for inquiry/response applications, but the LAN traffic is not integrated (generation 2). Some have a TDM-based backbone network which provides fixed bandwidth to most applications of the enterprise, including LANs (generation 3). Frame relay over cell relay can be beneficial to all three classes of users. Naturally, each network has different levels of migration and immediate payback by undertaking this transition.

Unintegrated Networks Without Backbones (Generation 1)

Users of these networks stand to get the major quantum advantage from frame relay. First, many discrete low-speed lines are replaced with fewer high-quality T1 lines, which in itself can be cheaper and easier to manage. Secondly, the advantages of dynamic bandwidth allocation reduce the transmission bandwidth that would otherwise be needed; additional transmission resources contribute to a direct increase in transmission cost. To migrate to a frame relay network, the user needs to deploy the necessary number of nodal processors, upgrade the terminal equipment for frame relay (this could be done using a terminal server on a LAN and

then using a frame relay router), and install the high-speed and backbone transmission infrastructure.

Classical Backbone Networks, LAN Traffic Not Integrated (Generation 2)

These networks benefit from the introduction of frame relay because bandwidth can be better utilized, postponing or even eliminating the need to upgrade the transmission lines to either multiple T1s or T3s. In fact, it may even be possible to replace some T1s lines with less expensive FT1 lines. To migrate to a frame relay network, the user needs to replace the TDM-based fixed-allocation multiplexers with nodal processors and connect the LANs to the same structure. Usually, the transmission facilities making up the backbone network remain in place, eliminating expensive installation charges for communication upgrades.

TDM-Based Backbone With Fixed Bandwidth (Generation 3)

These networks are the easiest to upgrade by simply replacing the TDMs with frame relay hardware. The network runs better and is more efficient.

Some Evolving Issues

Two important issues need to be fully resolved before the introduction of frame relay services in mission-critical applications can be fully rationalized. These issues affect private networks but are also important in public networks. They are network management and congestion control.

Users need to be able to monitor traffic, establish PVCs, obtain management reports, undertake fault management, do traffic engineering, rearrange existing PVCs, and so on. Nodal processors supporting private frame relay networks come with a variety of network management interface tools, but may or may not implement the full Annex D LMI apparatus. However, public services may not match this level of network management richness in terms of front-end functions like graphics, reports, menu-driven commands, and so on. Users are also looking to integrate the LAN and WAN management system.

Congestion control remains a critical issue. Congestion results when the combined request for bandwidth from all users exceeds what the network can provide. Total network bandwidth is ultimately determined by the number and size of the trunks between the carrier's or user's nodes. Congestion becomes more likely as the number of subscribers increases. Some argue that "when congestion starts to occur, people will have significant problems . . . users' expectations for frame relay are too high" [11.46].

Vendors' initial approaches to the congestion issue has not satisfactorily solved the problem the way the implementation of the full ANSI apparatus end to end would. For example, some provide large buffers in the nodal processors for storing frames that cannot be immediately sent. However, networks such as SNA retransmit data if it is not acknowledged within a given time interval; hence, the delay seen by the FEP because of the buffering can cause it to send more data, which is exactly the opposite of what is needed in the congestion state. Others use some of the congestion techniques employed in X.25. However, this does not go to the source of the problem, which is the throttling back of the input traffic. Others deal with the problem by over-engineering the network (reportedly, this includes BT North America, Sprint Data Group, and MCI Communications [11.46]). This approach is not cost-effective for private network solutions.

With the mechanism provided in the frame relay standard, nodal processors can send notifications to the attached routers and other devices to slow them down. The router in turn has to be able to inform the end-user generating the traffic (such as a user, a host, or a file server) to slow down. According to observers, end-to-end cooperation is 2 or 3 years away (i.e., it will be achieved in 1993 to 1994).

11.5.3 Topologies and Support of Non-LAN Traffic

Equipment is appearing on the market to connect 3270 SNA and Bisync terminals to a frame relay network. See Figure 11.21 for an example of this application. Users want to be able to combine SNA traffic with other traffic over a WAN using a common technology like frame relay [11.29]. Any savings in transmission could be neutralized by the need to maintain two or more separate networks, staffs, management tools, etc. [11.47]. Users want to support an enterprisewide network




Figure 11.21 Use of frame relay in an SNA environment.

with as few technologies as possible; hence, the issue of whether frame relay can support multiple corporate applications emerges. While many users are migrating to LAN-based SNA configurations, facilitating the direct usage of frame relay, some SNA traffic remains on the large embedded base of traditional cluster controllers.

A number of vendors are introducing standalone frame relay adapters to support non-LAN traffic (see Figure 11.22). With these PAD-like systems, SNA multidrop lines between the the IBM FEP and the remote cluster controllers can be replaced with frame relay PVCs. Other vendors are incorporating the adaptation function directly in the nodal processors. SDLC frames are passed across the network in a predetermined PVC by assigning the destination of the frame on a per-port basis. Some public networks also provide PAD-like functions.

11.5.4 Enterprisewide Use of Frame Relay

This section looks at frame relay from an enterprisewide perspective. Because equipment based on frame relay over a mixed-media cell relay platform utilizes backbone facilities better than existing circuit switching T1 multiplexers, frame relay benefit users that want to connect LANs over integrated backbones while supporting a variety of other traffic (to take advantage of resource sharing). But users that simply want or need to provide high-speed links between remote LANs may be better off using FT1, T1, FT3, or even T3 links [11.19]. According to some observers, most users need to transport a mix of data, voice, and video; hence they may find it difficult to cost-justify building a pure frame relay network solely dedicated to LAN traffic [11.19]. More expensive nodal processors also support voice and video.

Two views on frame relay penetration exist: those who see frame relay deployed mostly in private networks, and those who believe carriers will make major inroads. A 1991 study found that 37% of Fortune 1000 companies interviewed were planning to use public frame relay services, 24% were planning to use private frame relay, 24% use hybrid networks, and the balance (15%) were not sure. Given the outsourcing trends discussed elsewhere in this book and the plethora of reasonably priced carrier frame relay services appearing on the market, public and/or hybrid application of the technology may in fact be the route to frame relay deployment. Table 11.13 summarizes possible strategies.

Figure 11.23 depicts a number of traditional LAN interconnection methods [11.25]. Part A of the figure shows a T1 line totally dedicated to routers. Part B of the figure shows a typical arrangement where a fixed portion of bandwidth from a T1 multiplexer is employed for LAN usage; this is typically 56/64 kbps. Part C shows a sophisticated T1 multiplexer which, includes an integrated bridge; a fixed portion of bandwidth on the T1 multiplexer is used. This usage of a T1 multiplexer

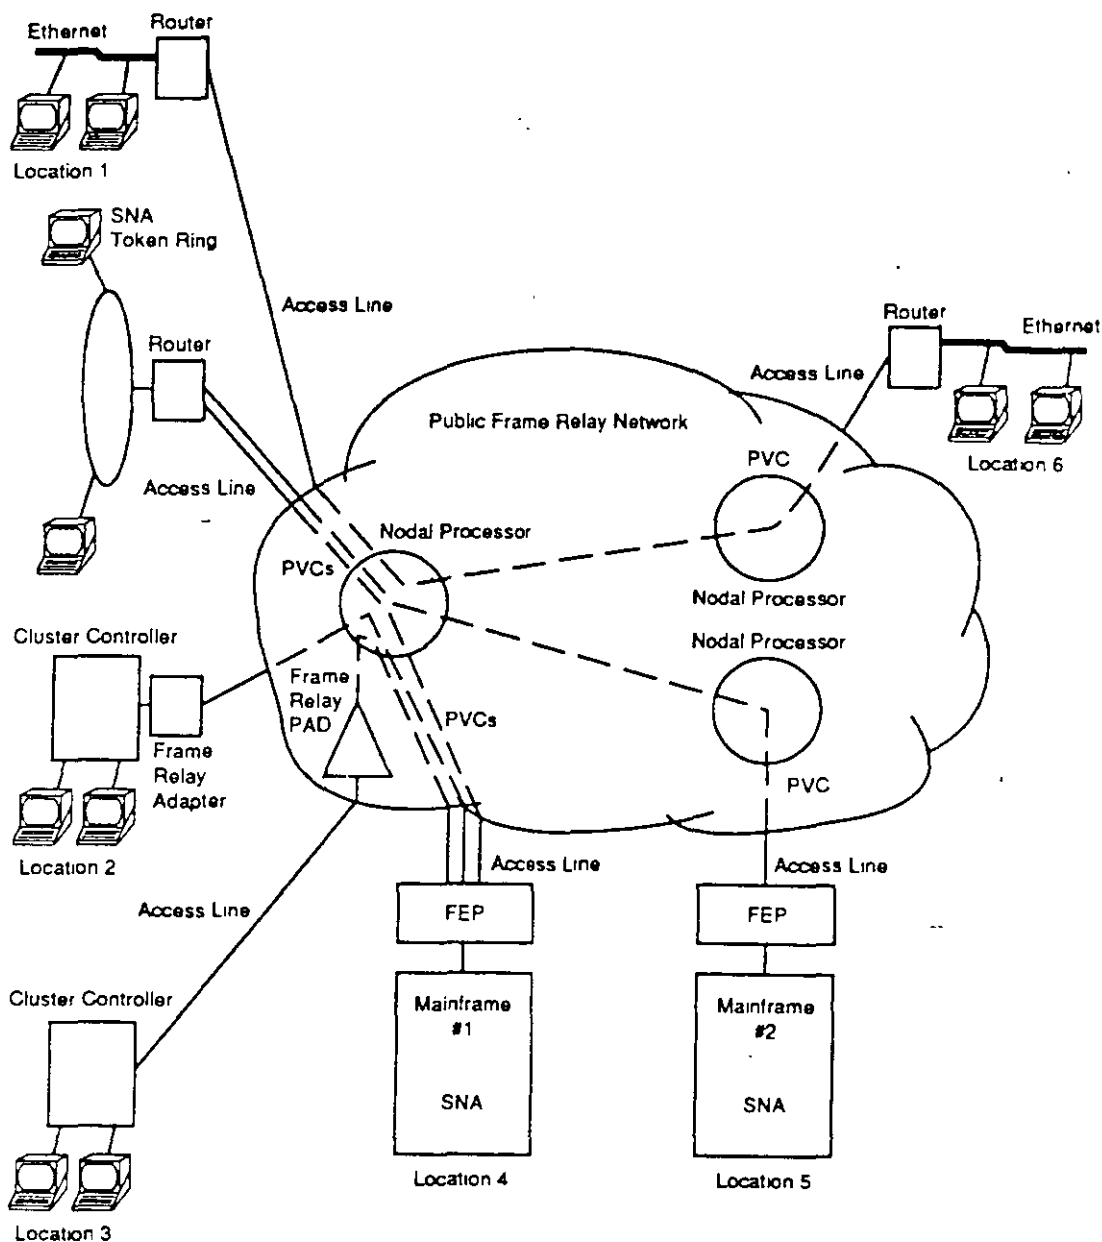


Figure 11.22 Supporting an enterprisewide network with frame relay.

was already discussed in Chapter 6. Note that three logically distinct components are required: a router, a multiplexer, and a line dedicated end to end.

Figure 11.24 shows some examples of LAN interconnection options using private network frame relay technology. Part A shows the use of a T1 line dedicated to a new router system that incorporates frame relay. Part B shows the case where a fixed portion of bandwidth from a T1 multiplexer is employed to connect a router

Table 11.13
Possible Strategies for Deployment of Frame Relay

Private Network Implementation

- Over a point-to-point line, connecting two routers directly
- Single-node data-only processor supporting LAN traffic
- Single- or multiple-node mixed-media processor(s) supporting enterprisewide networking

Public Network Implementation

- Data-only service for LANs or other devices (through PADs)

Hybrid Network Implementation

- Data-only service with private processors, while using public network to reach secondary sites
 - Mixed-media environment with private processors, while using public network to carry data
-

system which incorporates frame relay. Part C is a diagram of a T1 multiplexer which includes an integrated frame relay card but not a router; a fixed portion of bandwidth from the T1 multiplexer is employed. These three scenarios are likely to represent the early usage of the technology. Note that, as in Figure 11.23, three logically distinct components are required: a frame relay configured router, a multiplexer, and a line dedicated end to end.

Figure 11.25 shows other examples of possible interconnection options using frame relay. Part A shows a T1 multiplexer which includes an integrated router which uses frame relay; a fixed portion of the T1 bandwidth is employed. Part B depicts a situation where various streams run into a multiplexer where the trunk side uses frame relay (pursued mostly by packet switch vendors). Part C is the same as the previous case, but the trunk side uses cell relay and the trunk bandwidth is managed in fast packet mode. Here is where frame relay starts to offer advantages.

Figure 11.26 depicts a more sophisticated usage of frame relay. Part A demonstrates a private network using frame relay networkwide to achieve efficiency. PADs may be required to support non-LAN devices. A separate network for voice and video is required. Part B depicts the use of a mixed-media nodal processor, which also supports nondata applications. Part C of the figure shows a public frame relay network where multiple users share the network. PADs may be required. A separate network for voice and video is generally required. In this "optimal case," the user uses a router that implements the frame relay interface specification; but instead of obtaining a high-capacity line dedicated end to end, the user only gets the high-capacity line to the CO or POP (at both ends). By connecting to the carrier frame relay service, the carrier provides the multiplexing, releasing the users from that investment [11.2]. Note parenthetically that if the two endpoints terminate on the same CO (e.g., if they are in relative proximity within a city), then the bandwidth saving advantage disappears. When connected with a carrier frame relay

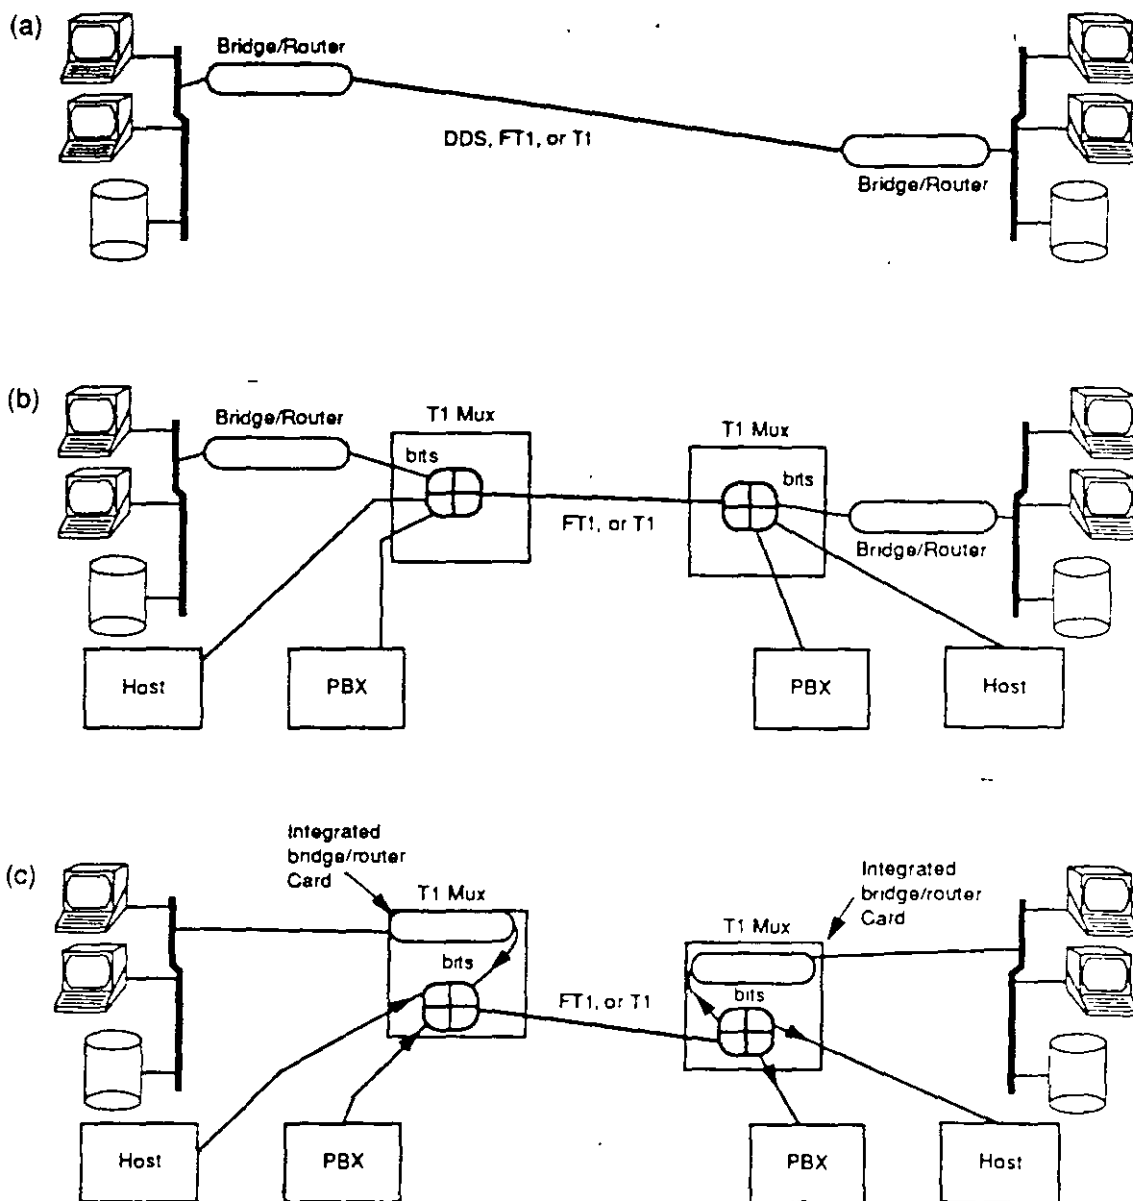


Figure 11.23 Traditional LAN interconnection methods (a) T1 line dedicated to bridge/router system, (b) a fixed portion of bandwidth from a T1 mux is employed; (c) T1 mux includes an integrated bridge, a fixed portion of bandwidth from a T1 mux is employed.

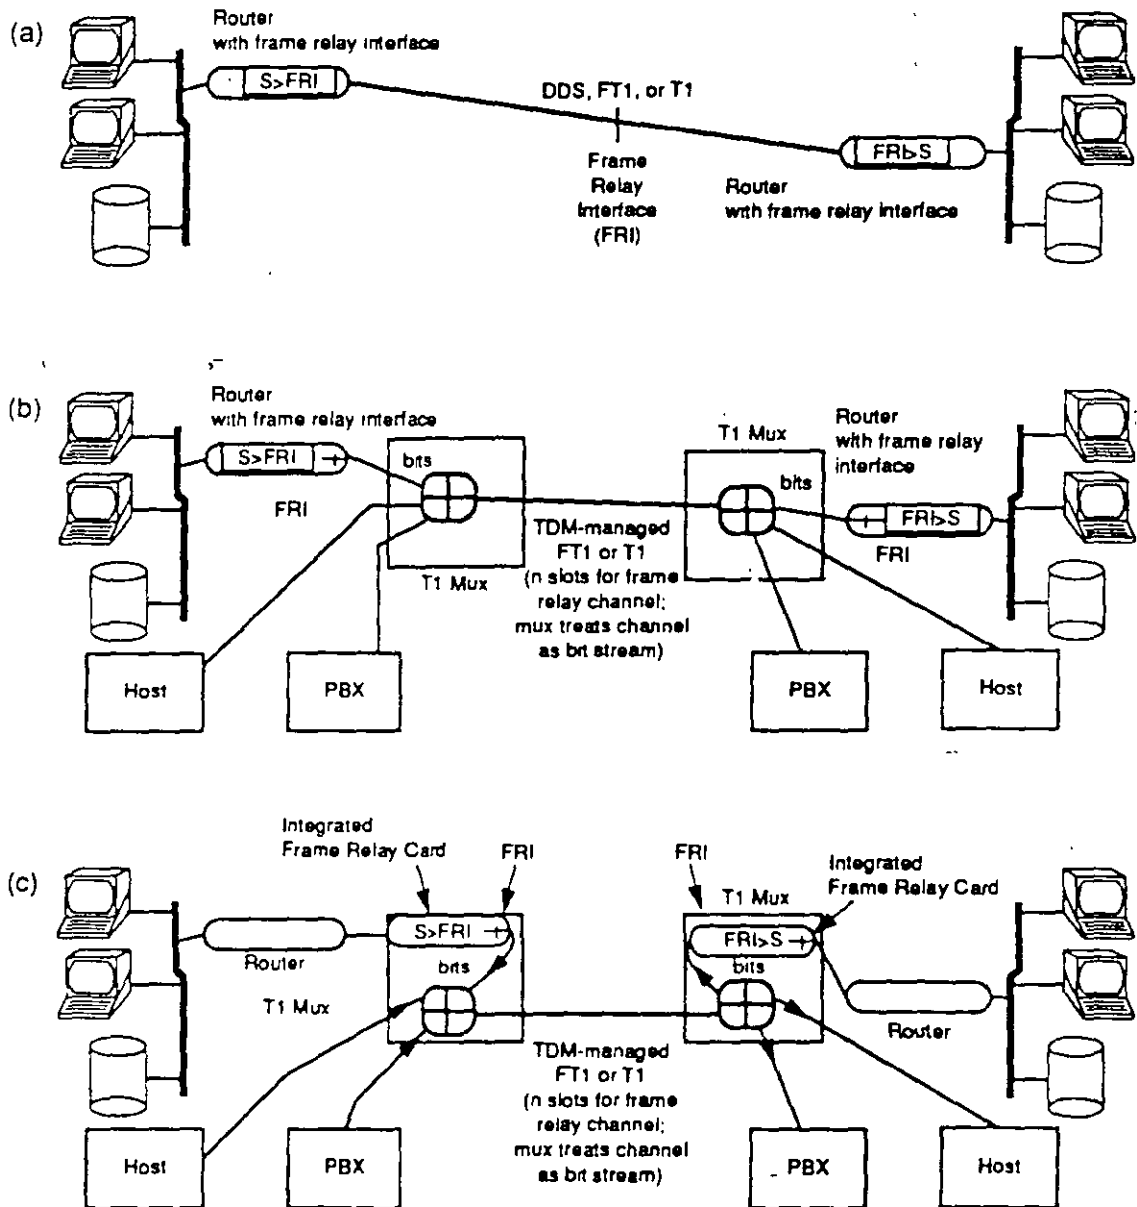


Figure 11.24 LAN interconnection options using frame relay: (a) T1 line dedicated to bridge/router system which incorporates frame relay; (b) a fixed portion of bandwidth from a T1 mux is employed to connect a bridge/router system which incorporates frame relay, (c) T1 mux includes an integrated frame relay card but no a bridge, a fixed portion of bandwidth from a T1 mux is employed. (S: Standard)

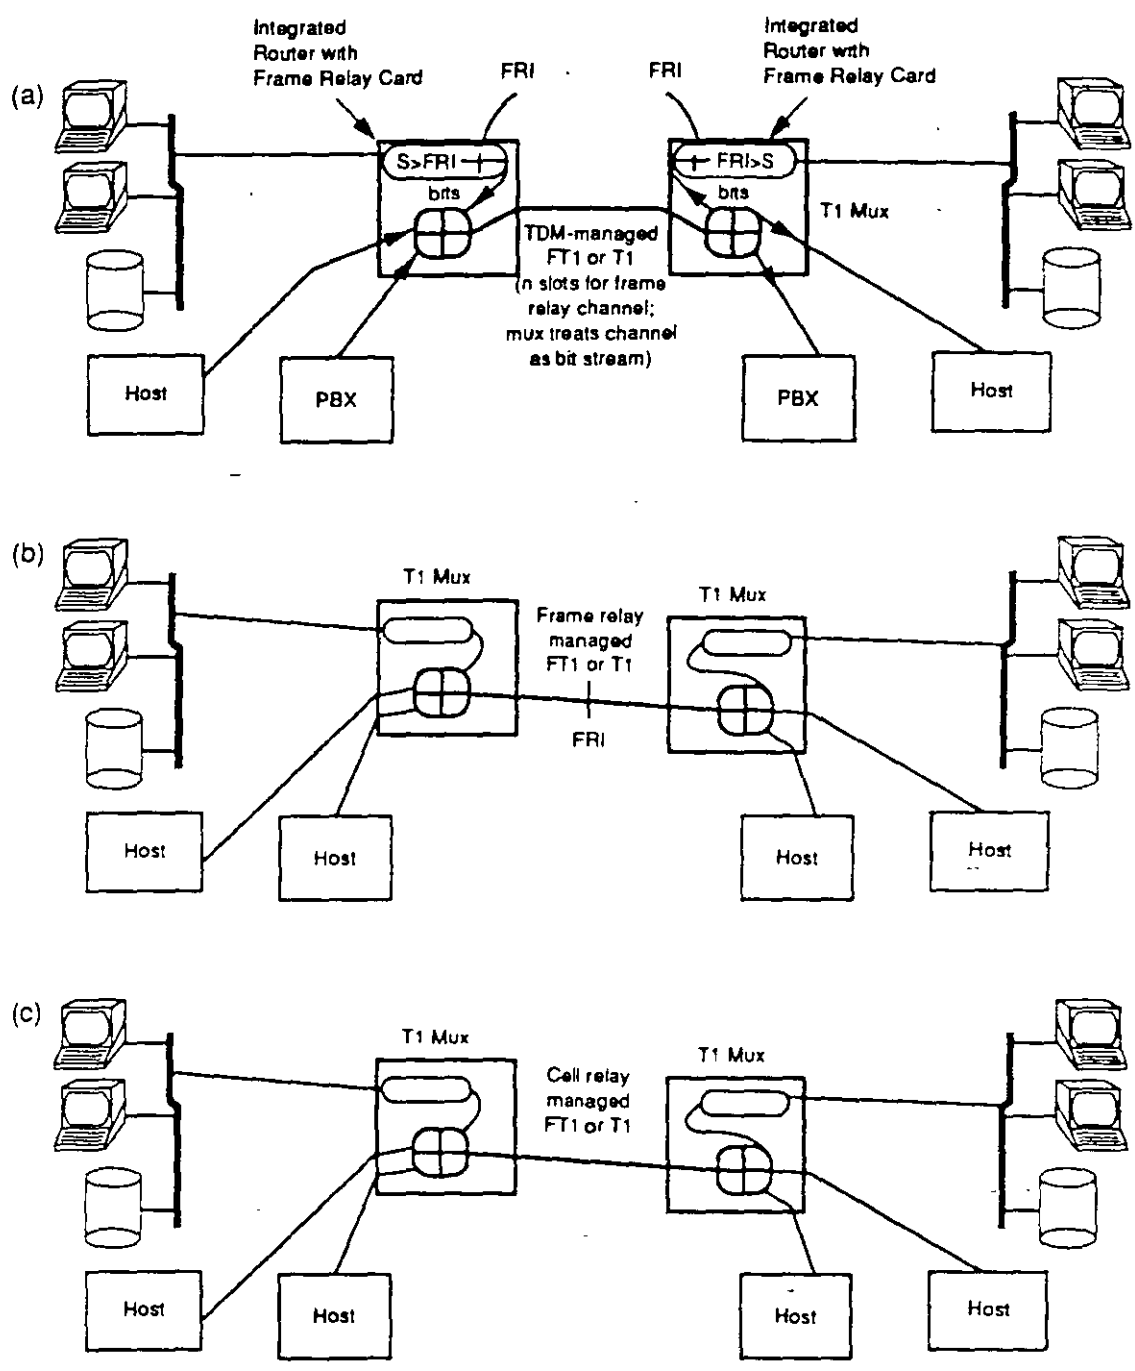


Figure 11.25 LAN interconnection options using frame relay: (a) T1 mux includes an integrated bridge which uses frame relay, a fixed portion of bandwidth from a T1 mux mux is employed; (b) various streams run into the frame relay-configured mux, the trunk side uses frame relay (this configuration used mostly by packet-based architectures—see Section 11.8.3); (c) same as previous case, but the trunk side uses cell relay.

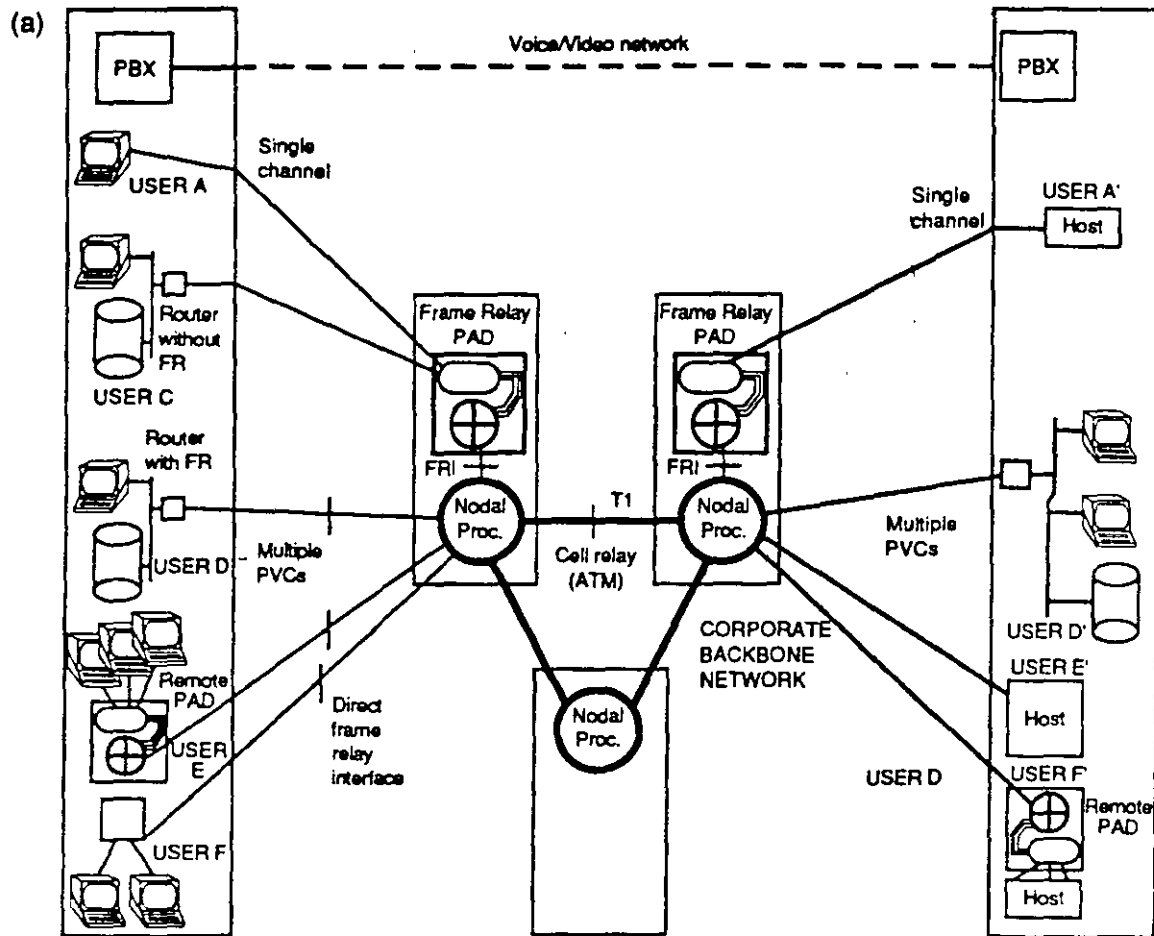


Figure 11.26 Interconnection options using frame relay: (a) a private network utilizes frame relay to achieve efficiency. PADs may be required. A separate network for voice and video is required

service, the routers see no difference compared to a private line. One of the advantages of this arrangement (but also shared by traditional packet switching and SMDS) is that if any part of the interoffice network fails, the carrier may be able to automatically recover or reroute. If this is done in real time, the user would be unaware of the failure event.

11.5.5 Practical Comparison of Interconnection Technologies

Frame relay fits in a continuum between private lines, SMDS, and BISDN services. Some users are planning to incorporate frame relay technology in their private

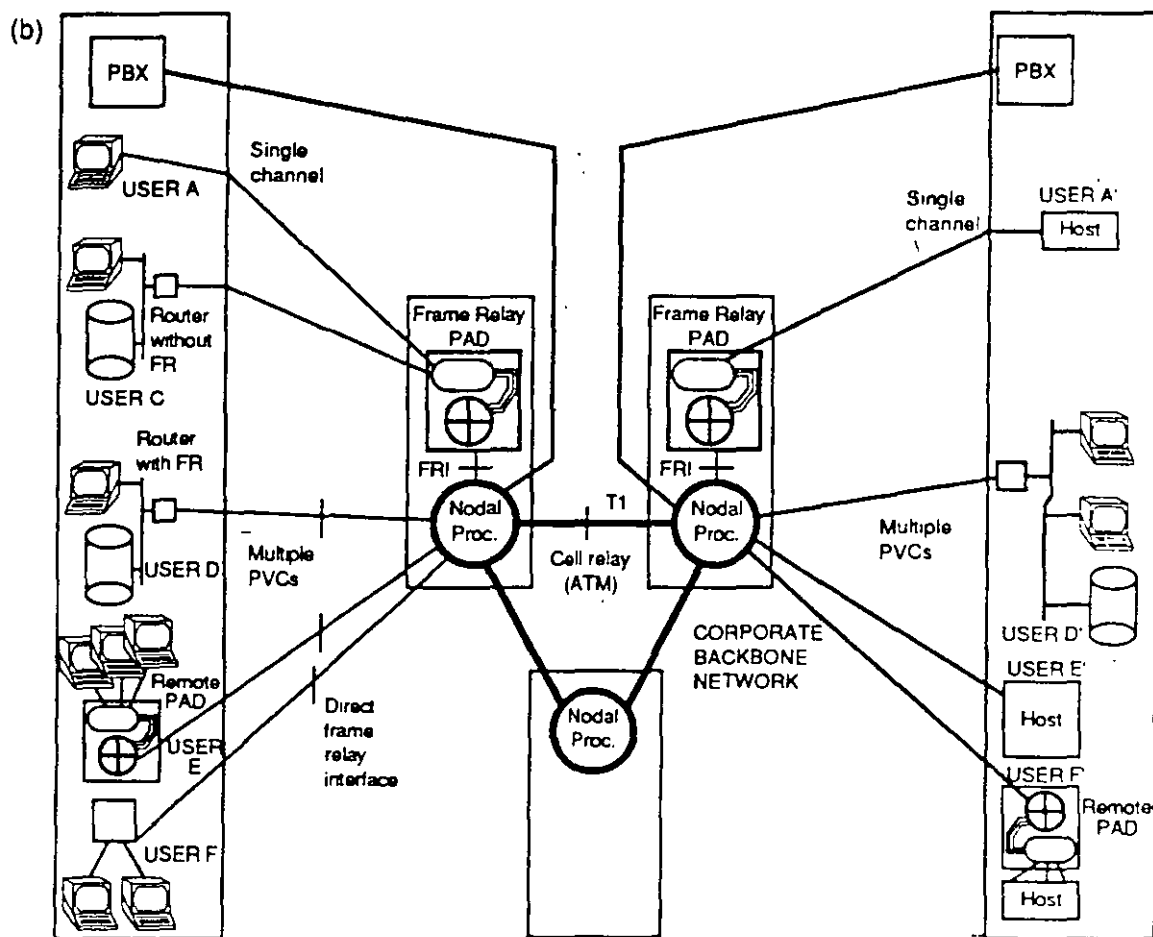


Figure 11.26 (Continued) Interconnection options using frame relay: (b) a private network utilizing mixed-media nodal processors.

networks. In the public arena, the progression of services in terms of complexity and availability will be frame relay, SMDS, and ATM/BISDN. Experts predict that it is likely that frame relay technology may be deployed in the same way that X.25 was: first on large private networks and then with carriers. Table 11.14 summarizes the frame relay/cell relay environment by highlighting the UNI/NNI characteristics.

The evolution toward SMDS seems clear. While routers have been quoted as passing in the neighborhood of 10,000 to 20,000 packets per second, the latest generation of bridges and routers now beginning to become available process 50,000 to 500,000 packets per second [11.8, 11.48, 11.49]. This means that while frame relay may be adequate for some LAN internetworking applications, other applications may need higher speeds, as provided by SMDS. Example of these applications include CAD/CAM, medical imaging, heavy-use desktop publishing, and

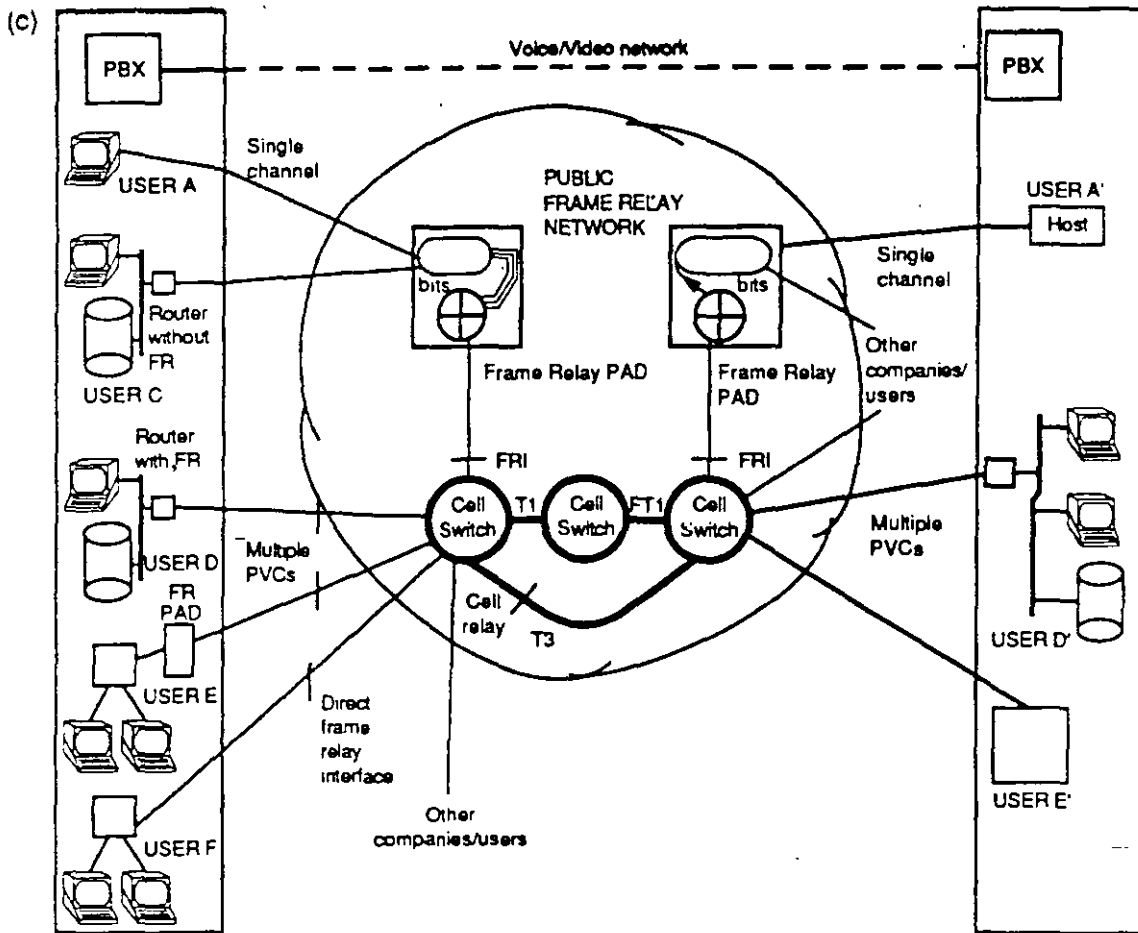


Figure 11.26 (Continued) Interconnection options using frame relay: (c) use of a public frame relay network to achieve efficiency, PADs may be required. Multiple users share the network. A separate network for voice and video is required.

animation. FDDI systems may become more prevalent now that the FDDI standards are practically complete and given that FDDI may actually be deliverable over twisted-pair. In addition, work has been underway to allow FDDI to interwork with SONET, implying that there may be an impetus to their introduction (i.e., the user does not require dedicated fiber, but can use facilities from the public network). This in turn may require high-throughput internetworking. It is not clear that a 1.544-Mbps service can bridge LANs operating at 100 Mbps. For some users, FDDI rates are too low (e.g., in supercomputer environments, discussed in Chapter 1).

At the pure technical level, since frame relay is a connection-oriented technology and LANs are connectionless, the ideal way to interconnect LANs is with

Table 11.14
Characterization of Various High-Speed Technologies

Network	UNI	NNI
Private data-only nodal processors	FRI	TDM
	FRI	Cell
	FRI	FRI*
	FRI	ATM**
Private voice/data nodal processors	FRI plus voice, video, and other data interfaces	TDM
	FRI plus voice, video, and other data interfaces	Cell
	FRI plus voice, video, and other data interfaces	ATM**
Public frame relay networks	FRI	Carrier-internal (cell or ATM)
SMDS	SNI	Carrier-internal
BISDN/cell relay service	ATM	ATM

*Not commercially implemented

**Beyond 1992 to 1993

a connectionless network-based service (such as SMDS) [11.50]. Also, it is desirable to avoid needing to develop entire technologies, and deploy networks which cater to a single application (e.g., just for LAN interconnection). Frame relay, as currently being standardized and deployed by carriers, is designed for data communications only, as a long overdue improvement of traditional X.25 packet switching. Cell relay (BISDN UNI) is specifically designed to support the sophisticated mix of services likely to be present in an organization of the 1990s: data, voice, facsimile, high-quality image and graphics, integrated messaging, and video.

Table 11.15 compares X.25, TDM multiplexers, native frame relay, frame relay over a fast packet switch platform, SMDS, and ATM from a service perspective (also see [8.22]).

Some users are reportedly concerned that the push for deployment of frame relay is coming from vendors rather than from network managers and users. Some users characterize frame relay as "more hype than necessity," since existing equipment can answer equally well the needs of stream traffic and data traffic with high autocorrelation (such as in file transfer) [11.51]. The promises of "seamless" LAN

interconnection cannot be delivered in full by frame relay because of the speed limitations, and because it is a connection-oriented technology.

11.6 FRAME RELAY EQUIPMENT

In a private frame relay network, the nodal processor is the most critical component. With a low-capacity processor, frame relay will not support the required throughput. A cell relay-based platform with effective network management tools is the type of equipment end-users are looking for.

A high-throughput nodal processor built from the ground up, unaffected by TDM restrictions, which supports high-speed switching to facilitate high end-to-end throughput, low latency, and any-to-any connectivity is required to derive the advantages that frame relay promises. A fast internal processor must be used to sustain the switching at the level required by the new routers now reaching the market and by the data-intensive user applications.

The nodal processors must support standard high-speed interfaces to the routers to facilitate the interconnection of equipment from a variety of vendors. This open frame relay interface should support a full T1 rate in order to properly interwork with existing router systems now deployed on dedicated T1 lines (some processors do not support a full T1). It is important that an adequate number of PVCs per frame relay interface be supported. A restrictive number of PVCs defeats the link and port sharing benefits of frame relay.

The nodal processors must support standard high-speed interfaces between nodal processors to provide cell relay and switching. The flexibility of being able to support fractional T1 or full T1 rates for the trunks is necessary in order to fine tune the network to the actual traffic patterns of the corporation. Generally, not all locations in a company have the same incoming and/or outgoing traffic volumes. Hence, the ability to be able to utilize a mix of T1s and FT1 trunks is an important cost-saving feature. Usually it is better to use outboard CSUs so that the LAN manager can optimize the investment needed to obtain the appropriate link management features without duplication. The choice of the CSU can be linked with the T1 channel at hand: for example, a link may or may not support B8ZS, and so the CSU can be chosen appropriately. In addition, the failure of the CSU, possibly incapacitating a path, can be mitigated by the use of a spare CSU, which is more difficult to do when the CSU is integrated with other hardware. In addition, a nodal processor should not impose topological constraints in terms of the number of nodes which can be supported.

Not every user device in an existing user network can be retrofitted with a \$1,000 frame relay board. A nodal processor should, therefore, support devices such as asynchronous terminals, synchronous terminals, and X.25 streams for those

situations where the frame relay interface is not available or will be installed at a future date.

A sophisticated network management capability is required to facilitate PVC establishment and to undertake all the necessary monitoring functions so important in mission-critical enterprise networks. A centralized system with access to the entire network through a distributed architecture is desirable. Graphical workstations with windows and user-friendly interfaces are a clear advantage. A rich feature set for fault, performance, accounting, security, and configuration management is an important business advantage.

Since the state of the art is not going to stand still, the nodal processor must be able to grow with new needs, features, and technologies. Some examples are the ability to migrate to BISDN, support SVCs, and deploy more data-intensive network management facilities in support of tighter control. The issue of congestion control is critical in order for the LAN-manager to guarantee a grade of service to the user community. A nodal processor should support the full ANSI congestion mechanism in order to achieve this goal.

11.7 CARRIER SERVICES

Several carriers now provide or plan to provide public frame relay services. Not only is it important that the service be available from a carrier, but it is also critical that the service be tariffed in a competitive way if users are to make investments for migration to the new technology. This section examines some issues pertaining to the public service.

11.7.1 Congestion Control Issues for Public Networks

As indicated, in frame relay the entire bandwidth, up to the maximum access speed, can be made available to a single user during peak periods. A problem may arise in the network if many users require this bandwidth simultaneously, as might be the case when LANs from multiple organizations (or departments within an organization) are terminated on the network. The frame relay network must be able to detect any overload condition and quickly initiate corrective actions.

Congestion control (also known as flow control) is already needed in traditional public packet networks, but in a frame relay network its need is more critical due to the performance objectives of the latter, and the greater access speed. In X.25 networks, the access speed is normally much lower than the speed and capacity of the backbone. It is unlikely that a single device would ever monopolize the backbone. In a LAN interconnection/frame relay environment, the routers seen as an ensemble may transmit a combined rate which might approach the capacity

of the backbone itself. A single router may flood the backbone; this in turn will starve other circuits of bandwidth.

Temporary conditions of overload occur in any well-utilized network. Networks which never experience temporary overloads may in fact be underutilized. Over-engineering, however, is not a desirable way to handle congestion control because such an approach is not cost-effective. Ignoring the issue of congestion is also undesirable, since, in effect, it means not capitalizing on the full potential of frame relay. In private networks, transmission costs are a major component of any design evaluation, and most of the benefits of frame relay technology are lost if implementing it demands the leasing of excessive amounts of bandwidth [11.21]. The challenge is not how to preclude any temporary congestion, but how to react to it when it occurs. Over-engineering or, better yet, relying on statistical averaging to obtain the most efficient utilization of deployed resources may be an approach that is viable in a public network environment, given the large population of potential users.

The ANSI standards specify explicit congestion control notification bits and a congestion notification control message. The important fields in the address portion of the frame relay format are the FECN, BECN, and DE, described earlier. In the ANSI standard, each of the individual virtual circuits in a frame relay connection (if the user and/or topological implementation calls for multiple PVCs over a physical link) can be independently throttled back. To be fair, the sources that contribute the most to the congestion should be slowed down the most, while sources contributing less traffic should be slowed down less. Hence, the network must be able to identify which PVCs over a physical link or, beyond the access portion, in the network are responsible for monopolizing resources.

Both the user's equipment and the switch should be able to respond to congestion control actions implied by the congestion control fields. For example, during periods of heavy load, the network could signal the user's equipment, by setting the congestion bit, to reduce the traffic arrival rate; when the overload situation dissipates, the opposite action could be achieved by setting the congestion bit back to normal. In some situations, the user's equipment could be overloaded; for example, a LAN gateway may be servicing another user and may not be able to absorb heavy loads of traffic coming from the network. Here, the user's equipment must be able to throttle the network.

The ANSI standards also provide for a DE capability to discard some frames if the initial congestion control actions do not correct the situation. The network should not be designed to discard frames indiscriminately: it is fairer to discard frames from the users who contributed the most to the congestion. If the implementation supports the DE field, this can be accomplished equitably, since the user's equipment can indicate which frames should be discarded first. The DE capability makes it possible for the user to temporarily send more frames than it is allowed on the average. The network will forward these frames if it has the

capacity to do so; but if the network is overloaded, frames with the DE bit set will be discarded first [11.21].

Some network/equipment vendors may implement a simple flow control procedure, rather than the full ANSI capability. For routers incapable of implementing the control mechanism of the ANSI Annex D specification, a simplified X-on/X-off form of flow control is allowed by LMI. The optional flow control limits transmission in the direction of the network, but not the reciprocal way. In the view of observers, while this approach is useful, backbone frame relay networks must also implement the full ANSI mechanisms; otherwise, the network will not be able to control effectively overloads from these devices.

Implicit Congestion Notification (to the transport layer of the ultimate user equipment, i.e., the PC) occurs when the user's end-to-end protocol determines that data been lost. Actions to deal with Implicit Congestion Notifications usually take higher priority than Explicit Congestion Notifications. The former is normally handled by the ultimate equipment; the latter is handled first by the router and subsequently by the ultimate equipment. The network may indicate to the user's router that the data may be about to traverse a congested path by the FECN/BECN bits previously discussed. The user response to these congestion notifications is dependent on the type of notification and the frequency in which they are received [11.6].

To reduce oscillations possibly due to transient congestion conditions, a congestion monitoring period (CMP) can be established by the user's router to track the frequency of Explicit Congestion Notifications received. This CMP is typically defined as four times the round trip delay through the network. The CMP starts upon receipt of a frame with the BECN or FECN bit set, or if the logical link is currently recovering from a congestion state. In a windowing environment, two window rotations may be used to measure the CMP instead of four times the round-trip delay. The user's router receiving the FECN bit set in half or more of the frames received during the CMP should start throttling data in the direction of the received frame. Since data at any given time is typically weighted in the direction opposite of the frame with the BECN bit set, the BECN indication is likely to occur less frequently than the FECN indication. The user's equipment should therefore start throttling data in the opposite direction of the received frame when the first indication of BECN is received [11.6].

During data transfer, one of the following four states is active. Typical carrier-suggested actions are [11.6]

1. Data throttling due to Implicit Congestion Notification. When a frame has been lost, as seen from the end-to-end protocols, the data flow should typically be reduced by approximately one-fourth of current flow. Data should not be throttled below the minimum end-to-end protocol flow (e.g., minimum window size).

2. Data throttling due to Explicit Congestion Notifications. When data has not been lost during the CMP, and the criteria for FECN or BECN frequency during the CMP has been fulfilled (i.e., half or more of the received frames have the FECN bit set, or one or more of the received frames have the BECN bit set), then the data flow should be reduced by approximately one-eighth of the current flow. Data should not be throttled below the minimum end-to-end protocol flow (e.g., minimum window size).
3. Data flow recovery. If the criteria for FECN or BECN frequency has not been fulfilled during the CMP (i.e., fewer than half of the received frames have the FECN bit set, or no more received frames have the BECN bit set), then the data flow should be gradually returned to normal flow at a rate of one-sixteenth of the normal end-to-end protocol flow.
4. Normal data flow. No congestion notification occurs and data throttling is not necessary (i.e., no congestion action is taken).

11.7.2 Class of Service Parameters

Carriers are specifying various class of service parameters for the PVC frame relay service. These include:

- Committed burst size (CBS). This is the maximum amount of user data (in bits) that the network agrees to transfer, under normal conditions, during one second.
- Excess burst size (EBS). This represents the maximum amount of uncommitted data exceeding the CBS that the network will attempt to deliver during one second.
- Committed information rate (CIR). This represents the user's throughput that the network commits to support under normal network conditions. CIR is measured in bits per second.
- Committed rate measurement interval (CRMI). This is the time interval during which the user is allowed to send information at the CBS rate or at the CBS + EBS rate.

See Figure 11.27 for a graphical interpretation. These quantities are important, since they are the basis of the services the carriers provide and for the supporting tariffs. Frame relay carriers will enforce the subscribed CBS, EBS, and CIR in the network in order to meet the grade of service. The user must allocate some minimum CIR to every possible device-to-device relationship (i.e., PVC); this implies that frame relay service, as currently available, is not the optimal solution to interenterprise applications (where SMDS may be).





Figure 11.27 Traffic arrival and treatment in a frame relay network.

11.7.3 An Example of Designing Networks With Public Frame Relay

A study of frame relay tariffs at press time revealed that each frame relay carrier had a different pricing scheme. Not only are these pricing schemes complicated, but a reliable comparison between services is difficult. It is almost impossible to generalize about the cost of frame relay services from one carrier to another,

especially when using published prices (some carriers avoid publishing generic tariffs; while nondominant carriers are not obligated to publish tariffs, such publication would certainly help the user choose a service/carrier) [11.52, 11.53].

While some carriers offer flat pricing, others offer pricing based on the number of user locations, the amount of bandwidth, and distance between the carrier's POP and the user's location. Some carriers sum the bandwidths defined on all the network PVCs (whether actually in use or not). Some add a surcharge for any data that needs to be delivered over a user channel exceeding 1,800 miles (presumably this is related to the fact that the propagation time slows down the delivery of the data to the user, implying added network responsibility). Many have access line charges, although some hide (absorb) that cost.

A published comparison among three carriers for service in four cities (Chicago, New York, Dallas, and Los Angeles) is shown in Table 11.16 [11.52]. The table shows that there is a lot of variability in the cost, and a rational comparison is difficult.

One conclusion that does emerge is that frame relay service is cheaper than fully interconnecting all locations with point-to-point high-speed digital lines. A public frame relay network generally costs about a third of a fully interconnected mesh network. Assuming that the carrier has a service POP in all LATAs where the user has traffic sources/sinks, the cost-effectiveness of the frame relay solution increases as the number of sites to be connected increases. In addition, 56- and 64-kbps frame relay services are universally cheaper than comparable X.25 services, which frame relay can replace in a number of situations (e.g., LAN interconnection).¹¹¹

Table 11.17 compares a public frame relay network with FT1 effective throughput (the physical access line may in fact have to be a T1 line), a traditional mesh FT1 network, and a private one-node frame relay network. Figure 11.28

Table 11.16
A Cost Comparison Between Frame Relay Services (January 1992)

CompuServe	\$23,140
Sprint Standard	\$12,260 to \$13,370 (depending on usage volume)
Sprint Reserved	\$36,300
Sprint Hybrid	\$19,920
Wiltel	\$19,620 (estimated)
Coverage: Chicago, New York, Dallas, Los Angeles	
Access (physical T1): \$1,300	
Access (logical): 1.024 Mbps	
PVC: 512 kbps	

¹¹¹Some carriers report that many users in fact employ the service at the 64-kbps rate.

Table 11.17
 Typical Frame Relay Costs; Public Network Covering New York,
 San Francisco, Atlanta, Dallas, and Chicago

<i>Configuration*</i>	<i>WlTel CIR</i>	<i>Sprint Data Standard Rate (no guarantee)</i>	<i>Sprint Data Reserved (guarantee)</i>	
256 kbps access 256 kbps throughput	\$ 4,527	\$ 4,950	\$ 5,600	
1.024M access 256 kbps throughput	\$ 9,011	\$10,700	\$11,350	
1.024M access 1.024 M throughput	\$15,354	\$10,700	\$17,950	
Mesh Dedicated FT1 Network*				
	<i>Mileage</i>	<i>FT1/64</i>	<i>FT1/128</i>	<i>FT1/256</i>
Chi-SF	1,860	\$ 893.62	\$ 1,674.92	\$ 3,149.99
Chi-NY	710	\$ 502.62	\$ 927.42	\$ 1,746.99
Chi-Atl	720	\$ 506.02	\$ 933.92	\$ 1,759.19
Chi-Dal	800	\$ 533.22	\$ 985.92	\$ 1,856.79
NY-SF	2,580	\$1,138.42	\$ 2,142.92	\$ 4,028.39
NY-Atl	940	\$ 580.82	\$ 1,076.92	\$ 2,027.59
NY-Dal	1,370	\$ 727.02	\$ 1,356.42	\$ 2,552.19
Atl-Dal	820	\$ 540.02	\$ 998.92	\$ 1,881.19
Atl-SF	2,230	\$1,019.42	\$ 1,915.42	\$ 3,601.39
Dal-SF	1,480	\$ 764.42	\$ 1,427.92	\$ 2,686.39
Total		\$7,205.60	\$13,440.70	\$25,290.10
Private Frame Relay Network*				
	<i>Mileage</i>	<i>FT1/256</i>	<i>FT1/512</i>	<i>T1</i>
Chi-SF	1,860	\$3,149.99	\$ 5,606.21	\$13,560.00
Chi-NY	710	\$1,746.99	\$ 3,099.21	\$ 6,660.00
Chi-Atl	720	\$1,759.19	\$ 3,121.01	\$ 6,720.00
Chi-Dal	800	\$1,856.79	\$ 3,295.41	\$ 7,200.00
Total (transmission)		\$8,512.96	\$15,121.84	\$34,140.00
Total (with amortized node)		\$9,512.96	\$16,121.84	\$35,140.00

*InterLATA costs only
 Press time tariffs, subject to change

depicts the topology of this example. A frame relay network is much cheaper than a mesh network; for the example shown (five cities), the frame relay service at 256 kbps of throughput is only 15% of the cost of a mesh network. This is what was meant earlier when it was stated that "in order to get the maximum benefit from frame relay technology without having to incur large charges, the service needs to

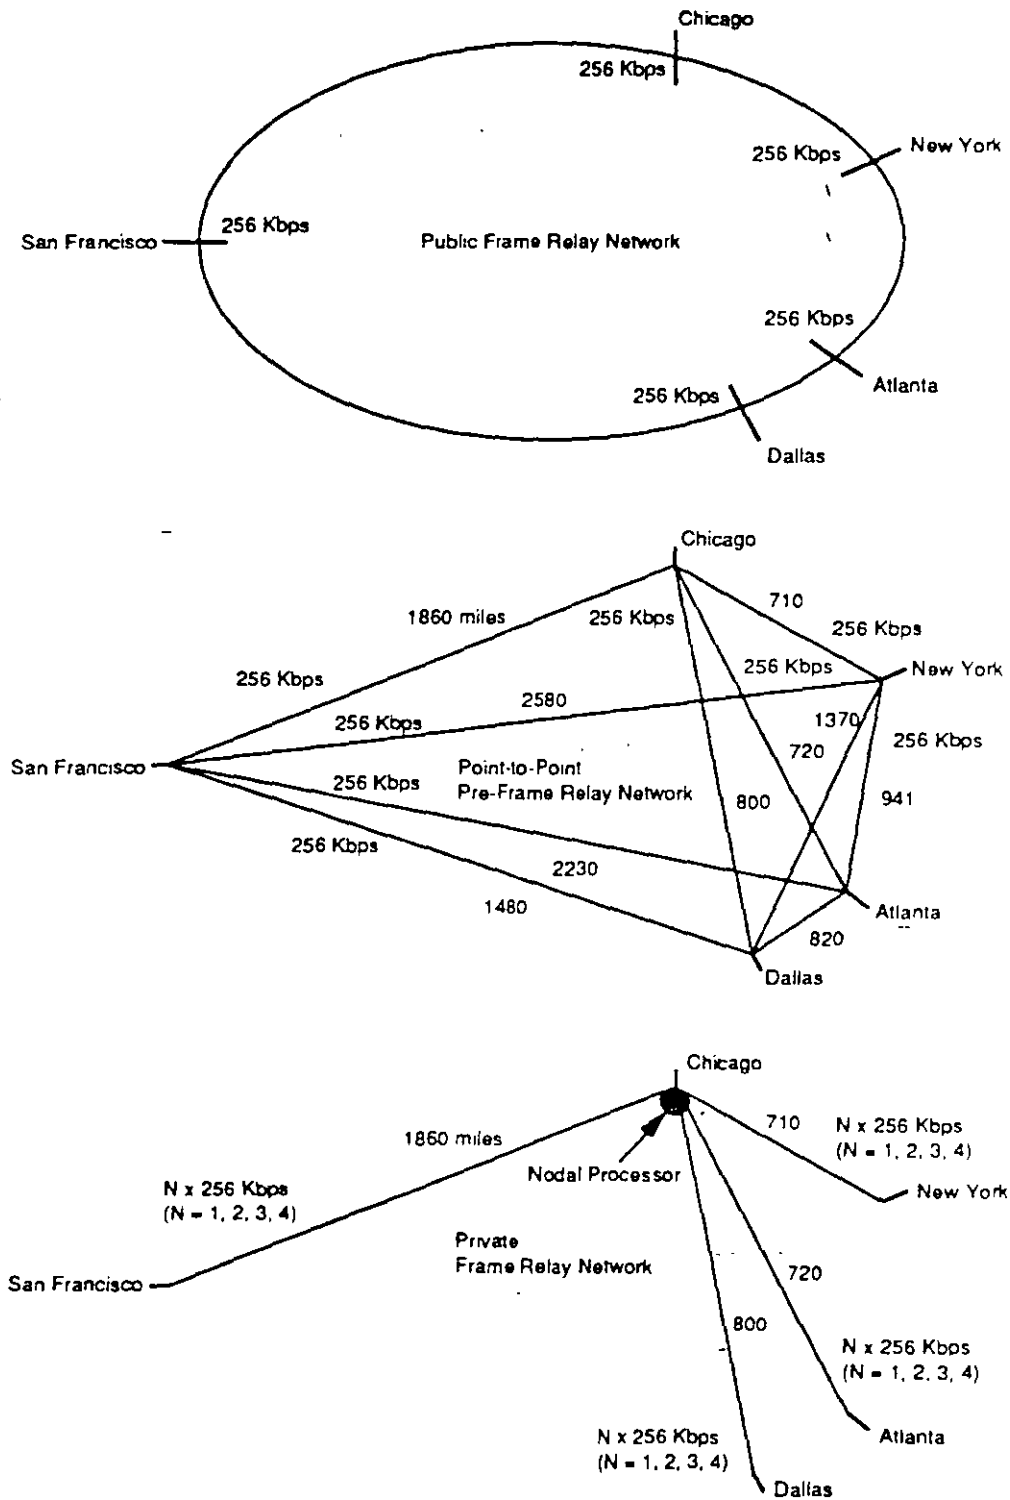


Figure 11.28 Comparing a public frame relay network with a mesh network and with a private frame relay network.

be provided by a carrier." The public frame relay network is only 25% (or less) of the cost of an appropriately configured private frame relay network. Note that in the public frame relay network, the throughput for each PVC from San Francisco, for example, could be 265 kbps. This implies that the private frame relay version must use a FT1 link to the node (which in this example was placed in Chicago), which should be 1,024 kbps as a "conservative" design, or at least 512 kbps as an "average" design.

One issue not clear from Table 11.17 is quality of service. In the full mesh network, the end-to-end delay approximately equals one transmission time. For example, if the mesh network used FT256 and the user's (ethernet) frame was 1,500 octets, then the delay would be 0.047 seconds. In the one-node private frame relay network, the delay would be 0.104 seconds, since the transmission time must be incurred twice, and there is nodal protocol processing delay (which we have assumed at 0.010 seconds). If two backbone nodes must be traversed (and it is assumed that the backbone link is also 256 kbps, the nodal protocol processing delay is 0.010, and the frame-to-cell and cell-to-frame assembly is 0.020 seconds), the total end-to-end delay would be 0.181, approaching the notorious delay incurred through a satellite link.

This example should make clear what this entire book has tried to do: there is no uniquely superior answer to a corporate networking problem. Each solution has advantages and disadvantages. A mesh network is more expensive, but the grade of service is better. A public frame relay network is cheaper, but there is more network delay, the service may not be available at all sites, and dedicated T1 access lines are still required. A private frame relay network is cheaper than a mesh network, while costing more than a public network; this solution, however, requires the user to purchase new equipment and to manage it. Another factor to take into consideration is the cost of the "access." If the carrier has a POP in the LATA(s) in question, that cost equals the cost of a T1 facility between the user's location and the POP. If the carrier only has a few nodes across the country, as is currently the case, the user may have to incur the cost of the T1 line to reach the switch; this could be hundreds of miles (some carriers pick up the cost of the access up to some distance).

The author is of the opinion that a practitioner may be hard pressed to try to rationalize why Company X (which may be profiled in a trade press magazine, or described by colleagues) used a given technology. Likely, Company X used a technology because of (1) how well a vendor made the case for the technology they sell, or (2) some senior manager in the company was "sold" by a trade press article which highlighted the advantages of a technology without ever describing its drawbacks (as is the practice), or point out the fact that the utility is highly dependent on the user's specific environment (ultimately, trade press magazines are influenced by the companies supporting them through their advertisement dollars).

11.8 FRAME RELAY PRODUCT AVAILABILITY

Vendors started to embrace frame relay technology in 1990, and equipment was appearing in 1991. As of press time, at least three dozen vendors have announced frame relay equipment and/or services [11.54]. For some vendors, such as those offering internetworking products, adding frame relay support may require a simple software upgrade of the hardware, since bridges and routers are already based on packet architectures. The same HDLC chips currently used on the communication side can be micro-programmed for frame relay [11.25]. The first wave of frame relay products must at least provide support for the access protocol, congestion management, and the status of PVCs. From a user's perspective, in order to deploy the equipment in the critical path of the corporation's ability to conduct business, robust and sophisticated network management capabilities must also be in place. This section provides a partial survey of some frame relay products in order to reinforce the fact that the technology is quickly materializing and that users can begin to study if and how frame relay can truly benefit their bottom lines in terms of decreasing their communications budget. This information will evolve over time.

11.8.1 T1 Multiplexers and Nodal Processor Manufacturers

Vendors of T1 multiplexers based on circuit switching TDM architectures need more work to transition to frame relay than vendors already supporting fast packet switching. These vendors need to add a cell engine to support frame relay in an effective manner; some have done so, while others are in the process of doing so. See Table 11.18, which provides a variety of other product information (based partially on [11.55]). Two approaches were used in the early 1990s as a short-term solution, short of a total architectural redesign. The first approach is to offer frame relay modules, or boards, for existing circuit-switched multiplexers. The second approach is to use a front-end frame relay developed by another vendor or strategic partner. With near-term solutions, the T1 multiplexer may typically only allocate a definite amount of bandwidth for frame relay support, and there may be performance and throughput problems. In the long term, traditional T1 equipment will have to be redesigned to incorporate fabrics which can exploit fully the advantages of cell switching. An important consideration is congestion control. Some vendors have experience in this arena, and others may not. In particular, vendors of packet switching equipment have dealt with this issue for years; vendors of T1 multiplexers have generally not had a need to deal with it. These products are typically used for private frame relay networks, although the more sophisticated equipment (e.g., StrataCom's IPX) can also be used to build public networks.

StrataCom's IPX Fast Packet multiplexer has supported a cell relay engine since the mid-1980s [11.20]. To support frame relay, the IPX required a software

Table 11.18
Partial List of Frame Relay Nodal Processors

Vendor	Product	Min	Max	Access Lines			NNI*
		Frame, Octets	Frame, Octets	64 kbps	512 kbps	T1	
Amnet	FRNS 7000	5	1,600	63	32	8	Arpanet datagram
Dowty	FPX 2000	1	4,096	120	80	40	FR
Hughes Network	FRS 9000	1	2,100	384	128	32	FR
NET		7	2,112	220	150	150	FR
Netrix	#1-ISS	1	4,096	300	48	16	FR
Newbridge	3600 IFS	1	8,200	30	4	1	FR
NTI	S/DMS	5	2,106	14,000	2,448	612	cell/ATM
	DPN-100	1	2,048				NTI's UTP
StrataCom	IPX 32	5	4,506	80	80	20	cell
Telematics	NET-25	5	4,096	64	16	16	Telematics' TNP
Timplex	Frame Server	5	1,600	12	12	12	FR
US Sprint	TP4900	1	8,189	528	66	22	FR

*FR = Frame relay

upgrade and frame relay cards (while routers typically already have HDLC cards, T1 multiplexers usually do not). Early support included Cisco routers [11.39]. The frame relay card accepts frame relay frames and segments them into 24-byte cells that can be transmitted over the StrataCom's proprietary T1 backbone. Users are not forced to dedicate bandwidth to the frame relay services a priori and on a preallocated basis. Each frame relay board (dubbed FRI-1M) consists of a V.35 interface with four ports and costs \$12,000. Initially the UNI was not supported at the full T1/E1 rate (it supported 1.024 Mbps), but as of 1992 these access rates are supported (using boards dubbed FRI-2M, which cost \$14,000) [11.56]. Carriers reported to use StrataCom's equipment include AT&T, WilTel, CompuServe, National Telecom Corp. (Canada), and Telecom Finland. The IPX switch also supports voice, and is therefore a mixed-media nodal processor (private implementations can support voice and data, but, to date, public implementations using the IPX only support data). Between the end of 1990 and the end of 1991, StrataCom sold 2,000 frame relay ports [11.57]. In 1990, StrataCom and DEC announced an equity agreement, resulting in DEC's worldwide distribution of StrataCom's IPX systems [11.58].

AT&T's BNS-1000 Fast Packet Switch is a multiport data-only switch supporting frame relay on the DTE user side and cell relay on the network side. ANSI and CCITT standards are supported on the access side (T1.606). Access rates can be as high as a full T1 or E1 (2.048 for European operation), or can be a standard subrate [11.59]. Preprovisioned PVCs' operation through the BNS-1000 makes network administration simple and eliminates the delay associated with call setup,

which is otherwise needed. Standard physical access-side interfaces are supported, including CCITT V.35, EIA RS-449, and EIA RS-232. At the upper protocol layers, the node transparently supports TCP/IP and other LAN industry standards. BNS-1000 nodes connect with other nodes over carrier-provided T1 or E1 links, or private fiber. Cell relay is used on the network side. When wide-area configurations require multiple nodes, a node-to-node maintenance channel up to 44 kbps is available to support the user's management requirements. The nodal processor can be configured as a frame relay switch in an existing multiple router environment; alternatively, when used in conjunction with the AT&T LCS200 Network Router and LCS100 Network Gateway products, it can be configured as a complete virtual private network platform for wide-area LAN interconnection. The switch achieves high reliability using both hardware and software redundancy for call processing. Automatic alternate routing is supported on the backbone side. In the event of failure or high incidence of fault occurrence on links between nodes, the Session Maintenance feature automatically detects trunk failures and reroutes traffic to alternate trunks, using previously unassigned bandwidth. Existing and rerouted traffic can share the same trunk. The process of detection, bandwidth negotiations, and route switching is accomplished within 10 seconds. Rerouted traffic can be moved back to its original path when the faulty link is restored. The switch supports over 30,000 endpoints simultaneously (15,000 two-way connections), over a private network, in a nonblocking mode. It can switch and forward 44,000 packets per second. The hardware is scalable in terms of the number of frame relay interfaces the individual nodes support in modules of four ports [11.1]. The BNS-1000 is aimed at private frame relay networks; a switch for public networks, the BNS-2000, is also available from the manufacturer, although the emphasis of the latter is on SMDS.

Network Equipment Technologies was reported to be looking at designing a frame relay interface for its Integrated Digital Network Exchange (IDNX) T1 multiplex product. The company was planning to offer first a proprietary IDNX board that incorporates the functions of a router and a high-performance packet switch to support direct LAN connections on the T1 multiplexer. The card would later be adapted to support frame relay [11.50]

GDC has made public commitments to a frame relay interface to the Megamux TMS T1 multiplexer. In addition to the new interface, the equipment internal bus was to be enhanced to support both a circuit-switched as well as a packet-switched architecture. It was planning on combining elements of TDM/circuit switching with frame relay and cell switching [11.20].

Newbridge has announced frame relay support through a new Distributed Communications Processor module of its 3600 MainStreet Bandwidth Manager T1 multiplexer. The product formats data from attached LANs into the frame relay format and passes the data unchanged over the circuit-switched private backbone. Initially, backbone bandwidth was allocated among TDM data, voice applications,

and frame relay data in a predetermined way rather than dynamically (this is true of all circuit-switched T1 multiplexers).

Motorola Codex has also announced frame relay support in its product line: the 6290 Series T1 Multiplexer, the 6525 Packet Switch, and the 6507 multifunction PAD. The frame relay interface for the 6290 is implemented using a two-card set: a four-port V.35 interface and a frame relay PAD. The cost of upgrading an existing node is in the \$20,000 range, while the cost of a new 6290 equipped with frame relay interfaces starts at around \$40,000. The 6290 can be managed using an OSI-based system. The 6525 packet switch can grow in 6-port increments up to a total of 48 ports; like most switches, it supports both dialup asynchronous terminals and access over a dedicated line. Adding the frame relay interface to an existing 6525 enables users to create X.25 subnetworks that feed into the 6290 fast packet backbone. Beginning in 1991, the frame relay interface became standard equipment on the packet switch; the switch upgrade costs in the neighborhood of \$5,000. The frame relay software supports up to 32 logical links over a single physical connection. The 6507 PAD supports ports individually operating distinct protocols, including Bisync/SNA, frame relay, and asynchronous dialup. The PAD can be connected to either the 6525 packet switch or the 6290 fast packet switch, and it costs in the neighborhood of \$2,000 [11.15].

Timeplex announced support of frame relay in both its internetworking and multiplexing product line [11.12]. The FrameServer System can be used either as a standalone frame relay nodal processor or in conjunction with the Link multiplexer family. The processor is quoted at \$14,000 to 25,000. A frame relay capability for the router product line has also appeared, allowing routers to connect to a public or private frame relay network. SNA traffic can also be consolidated for transport over the WAN frame relay network. The capability costs in the \$1,000 range [11.12].

Other vendors with frame equipment include [11.50] Coral Network Corporation and Hughes Network Systems; this list is likely to grow over time.

11.8.2 Router Manufacturers

Many router vendors now support frame relay interfaces, including 3Com, ACC, AT&T, Cisco Systems, CrossComm Corp., DEC, Hughes, Proteon, RAD, Sun Microsystems, SynOptics, Timeplex, Vitalink Communications, and Wellfleet [11.60]. Some routers can be upgraded using a frame relay software module; these range in price from \$750 to \$4,000. Some routers support both frame relay and SMDS. Most routers support ANSI LMI (Annex B and/or D).

11.8.3 Packet Switch Manufacturers

Traditional packet switch manufacturers are also positioned, in theory, to support frame relay, but they need to upgrade their switches to support higher speeds. These vendors have not made major breakthroughs in speeds in the past decade, and some observers question their ability to respond to the new environment. Bursty applications in LAN interconnection require DS1 or DS3 speeds to achieve optimal operation in today's environment of "network computing," file transfers, graphics, and decision support systems (such as spreadsheet applications). This type of equipment tends to use FRI at the NNI, rather than cell relay (future migration is possible). Some vendors are described below.

Northern Telecom announced a frame relay interface for the DPN-100 packet switch. The switch can be used to support hybrid public/private networks. A frame relay capability for Northern's CO switches, DataSPAN, is also available. DataSPAN's offering has been developed to retain compatibility with the installed base of CO switching equipment (both local and toll offices). DataSPAN is not an adjunct frame relay switch, which could introduce OAM&P complexities for the carrier. Instead, the frame relay fabric can be integrated on an existing switch, sharing common equipment, interfaces, and operations systems. DataSPAN is based on the Link Peripheral Processor of a DMS SuperNode. The Link Peripheral Processor serves a variety of functions, including ISDN D-channel packet handler, Signaling System 7 message processor, and frame relay handler. Any DMS-100, DMS-200, or DMS-250 switch in the network can be upgraded to DMS SuperNode. To add frame relay service, appropriate interface cards are put in front of, and new software is put into, the Link Peripheral Processor [11.2, 11.17]. Northern's implementation provides the PVC version of the service, but an ISDN/SVC version is under development [11.61]. Each frame relay interface can accept an unchanneled DS1 signal or a channelized DS1 representing 24 individual 56 kbps. On the trunk side, DataSPAN operates in a cell switching mode. The user's message arrives at the switch in a frame conforming to the frame relay specification; the switch segments the frame into cells and transmits them across the interoffice facilities. At the remote end, the cells are reassembled into frames while guaranteeing order preservation. For the applications requiring speeds in the 9.6-kbps range, DataSPAN may be connected to a 56-kbps service (DDS or ISDN). These interfaces are supported via standard DMS SuperNode Copper peripherals: the lower speed circuits are multiplexed into channelized DS1s and connected to units on the Link Peripheral Processor. Many LECs and IXC's are equipping their DMS SuperNodes with Link Peripheral Processors to implement Signaling System 7 capabilities [11.18, 11.62]. A trial with NYNEX for the frame relay interface on the DMS-100 and DMS-250 was planned for 1991.

BBN Communications was bringing out in 1992 a high-end packet switch supporting frame relay. The new T/300 Packet Switching Node supports up to 77



**FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA**

***DIPLOMADO EN REDES DE COMPUTADORAS
Y TELECOMUNICACIONES***

MODULO IV

***REDES DE ALTO DESEMPEÑO:
FAST Y GIGA ETHERNET, FDDI-II, "SWITCHING",
ATM Y FRAME RELAY***

OCTUBRE DEL 2000

-
- [11.33] *Addendum to T1.606—Integrated Services Digital Network (ISDN)—Architectural Framework and Services Description for Frame-Relaying Bearer Service*, ANSI T1.606-1990 Addendum, American National Standards Institute, New York.
- [11.34] *Digital Subscriber Signaling System No. 1—Layer 3 Signaling Specification for Circuit Switched Bearer Service*, ANSI T1.607-1991, American National Standards Institute, New York.
- [11.35] *Signaling Specification for Frame Relay Bearer Service*, ANSI T1.617-1991, American National Standards Institute, New York.
- [11.36] *Additional Procedures for PVC's Using Unnumbered Information Frames*, ANSI T1.617-1991 Annex D, American National Standards Institute, New York.
- [11.37] *Integrated Services Digital Network (ISDN)—Core Aspects of Frame Protocols for Use with Frame Relay Bearer Service*, ANSI T1.618-1991, American National Standards Institute, New York.
- [11.38] *Network World*, 10 September 1990, p. 1.
- [11.39] *Network World*, 15 October 1990, p. 68.
- [11.40] Joint News Release: Northern Telecom, Digital, StrataCom, Cisco Systems, September 4, 1990.
- [11.41] *Broadband Networking News*, Phillips Publishing, November 1991.
- [11.42] *Frame Relay Specification with Extensions, Revision 1.0*, DEC, Northern Telecom, StrataCom, Cisco Systems, 18 September 1990.
- [11.43] B. Schultz, "Frame Relay Taking Shape," *Communications Week*, 19 December 1990, p. 28.
- [11.44] T.H. Jones, "The Frame Relay Forum," *Business Communications Review*, October 1991 (Frame Relay Insert).
- [11.45] J. T. Mulqueen, "Kinks Show Up in Frame Relay Gear," *Communications Week*, 6 January 1992.
- [11.46] P. Desmond, "Net Control Issues Cloud Frame Relay," *Network World*, 24 December 1991, pp. 1 ff.
- [11.47] B. Schultz, "Routing SNA Over Frame Relay Networks," *Communications Week*, 4 November 1991, p. 21.
- [11.48] J. M. McQuillan, "Broadband Networks—The End of Distance," *Data Communications*, June 1990.
- [11.49] M. Molloy, "Wellfleet Offers Up Gigabit-Speed Router," *Network World*, 9 September 1991, pp. 1 ff.
- [11.50] P. Desmond, "Frame Relay Standards Get T1 Mux Makers' Backing," *Network World*, 11 June 1990, p. 2.
- [11.51] B. Schultz, "Jury Out on Frame Relay," *Communications Week*, 19 December 1990, p. 23.
- [11.52] E. Herman, "Pricing: The Hardest Part," *Communications Week*, 2 December 1991, pp. 1 ff.
- [11.53] P. Desmond, "New Frame Relay Pricing Lets Users Size up Services," *Network World*, 14 October 1991, pp. 1 ff.
- [11.54] LAN Newsletter, Information Gatekeepers, October 1991.
- [11.55] *Broadband Networking News*, Phillips Publishing, February 1992.
- [11.56] P. Desmond, "StrataCom Mux to Support Higher Speed Frame Relay," *Network World*, 16 December 1991, pp. 4 ff.
- [11.57] StrataCom Press Release, 16 December 1991.
- [11.58] StrataCom/DEC Press Release, 18 June 1990.
- [11.59] B. Schultz, "AT&T Network Systems Intros Products for Connecting LANs," *Communications Week*, 14 October 1991, p. 29.
- [11.60] StrataCom Press Release, 26 February 1990.
- [11.61] P. Desmond, "NTI Airs ISDN Interface for Frame Relay Module," *Network World*, 21 October 1991, p. 19.
- [11.62] *Communications Week*, 24 September 1990, p. 67.



TECNOLOGIAS "FAST-PACKET "

- CRS (Cell Relay Service)
- Frame Relay
- SMDS (Switched Multimegabit Data Service)

Notas:



FRAME RELAY

Se estima que el mercado conjunto de Frame Relay y SMDS (Switched Multimegabit Data Service) sobrepasan los \$1.2 billones de dólares en 1995 en los Estados Unidos.

Notas:



FRAME RELAY

Frame Relay es una nueva tecnología de alta velocidad basada en paquetes, con asignación dinámica del ancho de banda, alto rendimiento y bajo retardo, para soportar el incremento del tráfico de información en ambientes corporativos. Frame Relay define un formato estandarizado para los "Frames" de nivel de "data link", los cuales son transmitidos sobre una red de LANs interconectadas o sobre una red pública de datos. Un "Frame" es ensamblado por el equipo terminal del usuario y es interpretado por los procesadores nodales de Frame Relay o en su caso por ruteadores remotos.

Notas:



FRAME RELAY

Características:

- ✓ Soporta solo transmisión de datos
- ✓ La transmisión se establece por unidades de datos de longitud variable ("Frames").
- ✓ Se establecen conexiones virtuales
- ✓ Comparado con la tecnología tradicional de conmutación de paquetes, reduce significativamente los retardos en la transmisión.
- ✓ Más eficiente en el uso del ancho de Banda
- ✓ Decrementa costos en los equipos de comunicación

Notas:



FRAME RELAY

Eficiencia:

En los servicios tradicionales de conmutación de paquetes los retardos que se introducen en la red, son de 200 ms o más. En redes Frame Relay se reducen de 20 a 40 ms.

Notas:

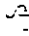
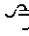
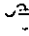
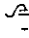
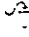


FRAME RELAY

Velocidad:

 FRI.- Frame Relay Interface

Soporta velocidades de:

-  56 Kbps
-  N x 64 Kbps
-  1,544 Mbps
-  2,048 Mbps (Europa)
-  45 Mbps (Algunos fabricantes)


Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



FRAME RELAY

Tecnología anterior:



Se requiere de 9 canales T1


Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



FRAME RELAY

Tecnología Frame Relay



Se requiere un canal T1 o FT1

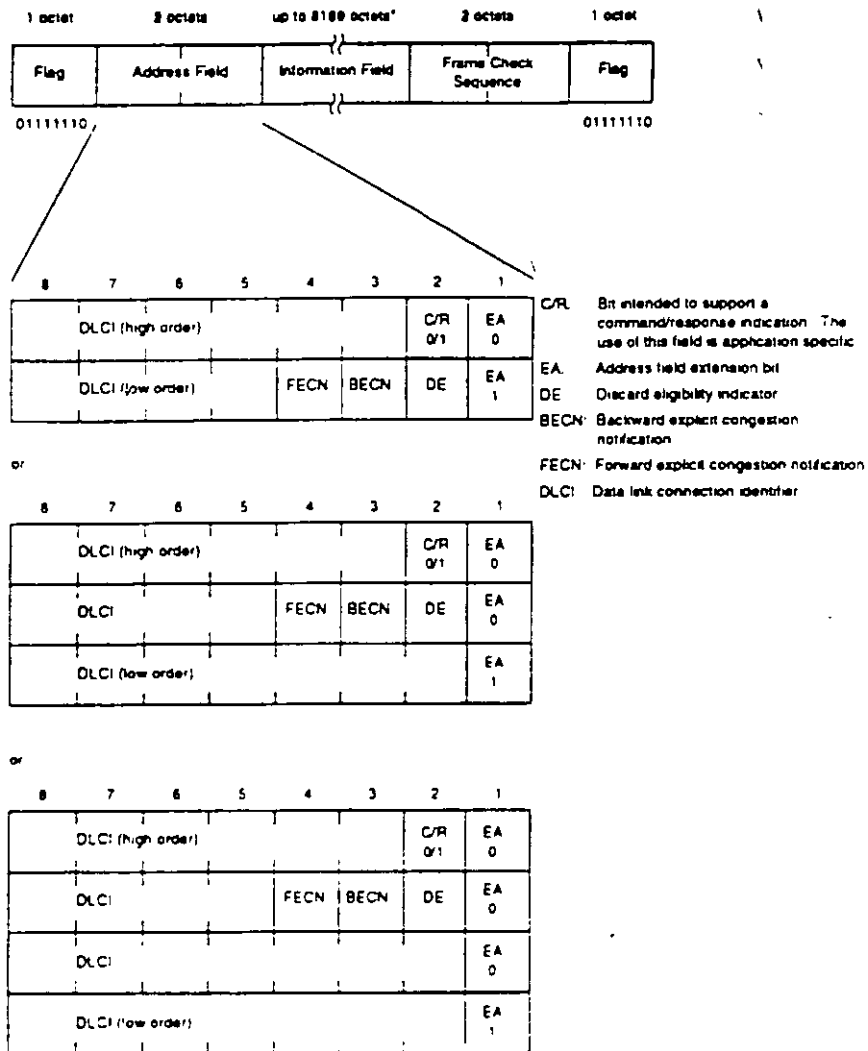
Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



FRAME RELAY

Frame y Formato del campo de direcciones



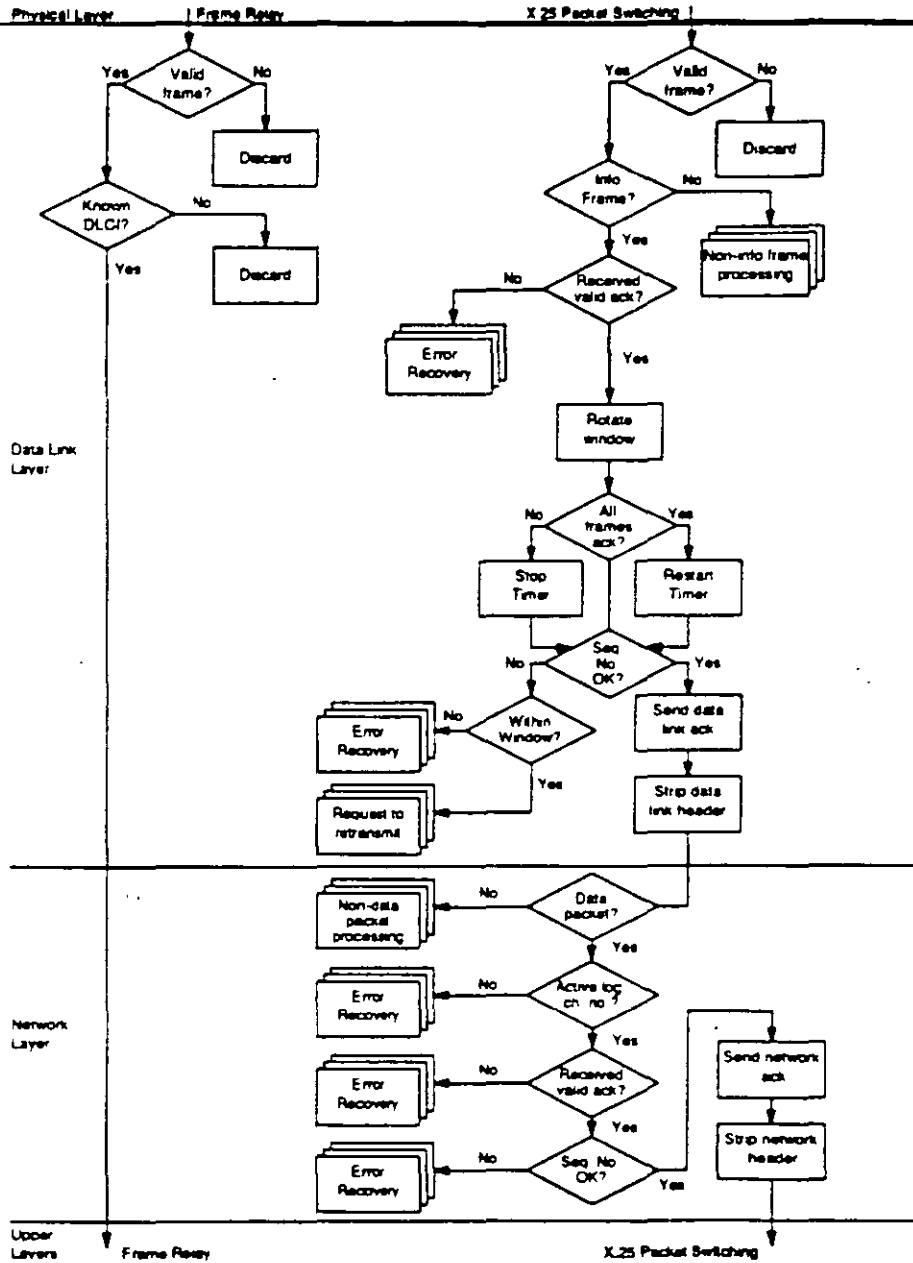
*ANSI specifies 4096 but some vendors may use the higher value minimum length 262

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



FRAME RELAY

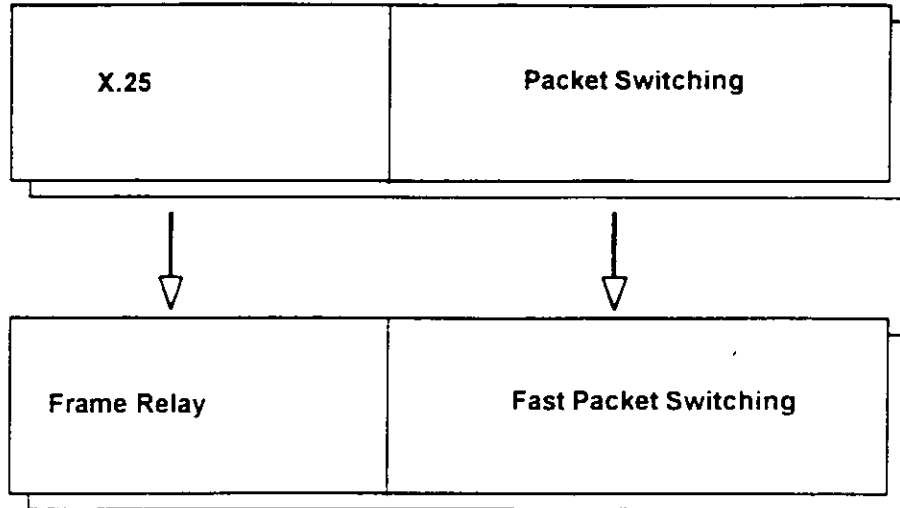
Comparación con X.25



TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



X.25 → FRAME RELAY




Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA




EVOLUCION FRAME RELAY

1990 - 1992



1993 - 1995



Notas:



FRAME RELAY

Corrección de Errores:

Quando en la red se detectan errores la corrección es relegada a los sistemas terminales. Las condiciones de error incluyen pérdida, duplicación, descartación, pérdida de secuencia de los frames. La forma de recuperación de errores puede medir el rendimiento de los equipos.

Estadísticamente se ha comprobado que la probabilidad de falla en un frame es del orden del 0.9%.

Notas:



FRAME RELAY TIPOS DE SERVICIOS

☰ SVC Semipermanet/Switched Virtual Circuit

✓ Asignación dinámica de rutas

☰ PVC Permanet Virtual Circuit

✓ Configuración estáticas de rutas

Notas:



FRAME RELAY

Conceptos:

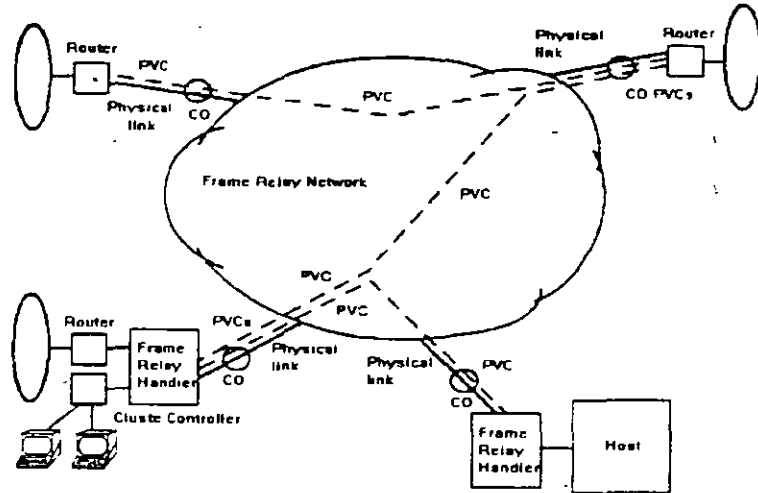
- ☐ Frame Relay Processors
- ☐ Frame Relay Nodal Processors
- ☐ Frame Relay Networks
 - √ Pública
 - √ Privada
 - √ Híbrida

Notas:



FRAME RELAY

Red Pública

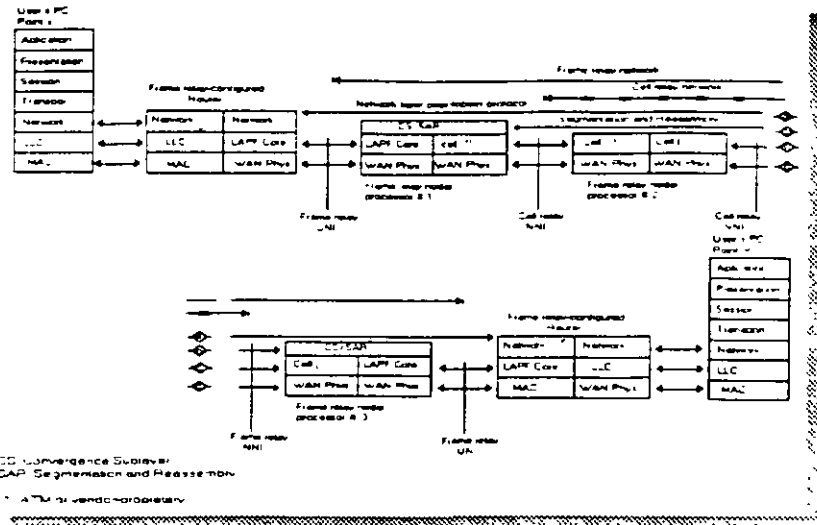


Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



FRAME RELAY



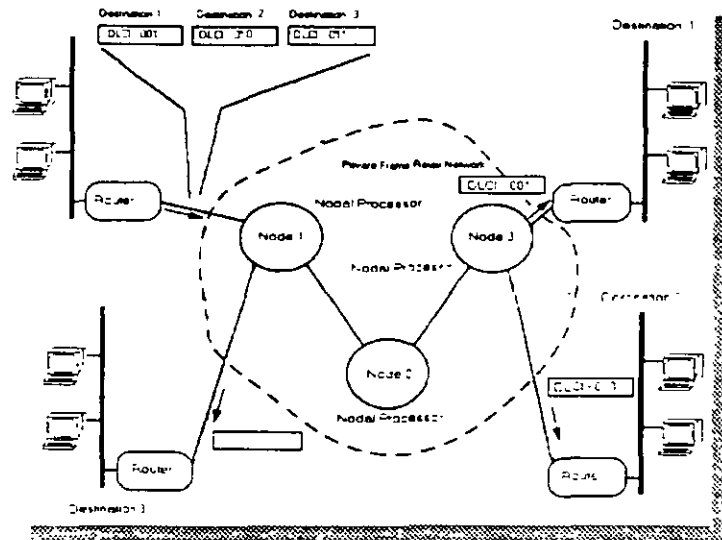
Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA



FRAME RELAY

Red Pública



Notas:

**DIVISIÓN DE EDUCACIÓN CONTINUA DE LA FACULTAD DE INGENIERÍA
UNIVERSIDAD NACIONAL AUTÓNOMA DE MÉXICO**

**Diplomado en Redes de Computadoras y Telecomunicaciones
(Tecnologías de la Información)**

Examen del Curso Redes de Alto Desempeño
Impartido del 23 al 27 de Octubre del 2000

Objetivo: Migrar la red de computadoras que se ilustra en el esquema 1-a una red que aproveche la tecnología de vanguardia, para establecer una red de ALTO RENDIMIENTO obteniendo la mejor relación COSTO/BENEFICIO.

Especificaciones: En el esquema 1 se tienen los siguientes elementos

Edificios: 1,2,3 ubicados contiguamente en la Cd. de México, existe Conectividad directa entre ellos.

Computadoras:

- A Computador de arquitectura ISA, interface de red Ethernet 10BaseT. 30 computadoras por piso.
- B Computador que se requiere adquirir e integrar a la red. 6 computadoras nuevas por piso.
- C Servidor de arquitectura EISA 1 interface de red 10BaseT
1 interface de red 10Base5
- D Hubs, concentradores Ethernet 10BaseT 16 puertos 10BaseT
1 puerto AUI
1 puerto BNC
2 concentradores por piso
- E Segmentos de cable UTP-5
- F Segmento de RG-11 como Backbone
- G Ruteadores con 2 puertos Ethernet
2 slots libres.
- H Ruteador concentrador 4 puertos 10BaseT
1 puerto serial
1 puerto AUI
2 slots libres
- I Canal E0

J Red publica X.25

K Enlace a Internet

L Servidor de nombres de E-mail

N Ruteador con 2 puertos Ethernet, 1 puerto serial, 1 slot libre

1.- Sobre el esquema 1 indicar las velocidades máximas de:

E: _____

F: _____

I: _____

M: _____

2.- Indicar las características adecuadas de:

Computadoras B: _____

Modelo : _____ Costo:\$ _____

3.- Justificar e indicar los cambios de equipo o configuración de los servidores:

Modificación de hardware: _____

o
Cambio de equipo : _____

4.- Ilustrar en el esquema 2, la red integral informática en operación con tecnología de alto desempeño

- * El usuario final debe tener acceso a toda la red.
- * El usuario final debe tener acceso a Internet
- * Las aplicaciones serán Cliente/Servidor
- * Los enlaces entre todos los edificios deberán ser a la mayor velocidad posible.
- * Indica en el esquema 2 como quedaría la migración de red a tecnología alto desempeño.
- * Piso 3- Aplicaciones de misión crítica

de

- * Mostrar sobre el esquema las velocidades máximas de los enlaces
- * Indicar con nomenclatura los elementos que conforman la red
- * UTILIZAR UNICAMENTE EQUIPO DE LA LISTA DE PRECIOS ANEXA, HACIENDO REFERENCIA AL MODELO.

5.- Requisar el pedido e indicar el costo total del proyecto.

NOTAS:

- El objetivo determina obtener la mejor relación costo/beneficio, esto es

- 1) Cubrir las necesidades
- 2) Al menor costo
- 3) No sobre pasar el presupuesto asignado

- Presupuesto asignado \$ _____
- Se pueden solicitar los servicios profesionales de un asesor

Fecha de inicio ___/___/___/

Hora Inicial ___:___

Fecha de aprobación ___/___/___/

Hora Final ___:___

Equipo: _____

R.F.C. _____

Nombre: _____

Tel: _____

Organismo: _____

Fax: _____

Equipo: _____

R.F.C: _____

Nombre: _____

Tel: _____

Organismo: _____

Fax: _____

Equipo: _____

R.F.C. _____

Nombre _____

Tel. _____

Organismo _____

Fax: _____

LISTA DE PRECIOS

Ejercicio de evaluación

MODELO	PRECIO USD	DESCRIPCION
PC1	\$ 1,500.00	Computadora Pentium III 450 Mhz,133Mb RAM, Slot PCI, Puertos USB,Monitor SVGA, Disco duro de 6 GB.
PC2	\$ 2,000.00	Computadora Pentium II 500 Mhz, 64 Mb RAM, 1 slot PCI, Puertos USB, Monitor SVGA, Disco duro de 8 GB.
SV1	\$ 6,500.00	Servidor Pentium III 550Mhz, 128Mb RAM, 3slot PCI, Monitor SVGA , Disco duro SCSI de 9 Gb, CD-ROM, Unidad de cinta y Hot Swap
SV2	\$ 8,000.00	Servidor Pentium III 600Mhz., 515 KB cache, 133Mb RAM, Monitor SVGA, Disco duro SCSI de 9Gb. CD-ROM, unidad de cinta y Hot Swap.
CT1	\$ 900.00	Concentrador de 16 puertos 10BaseT
CT2	\$ 3,000.00	Concentrador de 16 puertos 100/10BaseT, 1puerto AUI
CT3	\$ 1,800.00	Concentrador de 16 puertos 10BaseT, 1 puerto AUI, 1 puerto BNC,compatible SNMP.
CT4	\$ 3,600.00	Concentrador de 16 puertos 100/10BaseT, compatible SNMP
RT1	\$ 3,500.00	Router con 2 puertos Ethernet, 1 puerto ISDN
RT2	\$ 2,500.00	Router con 2 puertos Ethernet, 1 puerto serial
RT3	\$ 3,500.00	Router con 4 puertos Ethernet, 1 puerto ISDN
RT4	\$ 3,000.00	Router con gabinete, 4 slots y fuente de poder
RT5	\$ 5,500.00	Router con gabinete, 8 slot fuente de poder
SW1	\$ 3,500.00	Switch Ethernet con 16 puertos 10BaseT, 1 puerto AUI,1 puerto BNC, 1 puerto 100TX
SW2	\$ 3,500.00	Switch Fast Ethernet con 8 puertos 100/BaseT, 1 puerto AUI
SW3	\$ 7,500.00	Switch Fast Ethernet con 8 puertos 100/BaseT, 1 puerto AUI,1 puerto ATM

SW4	\$	3,200.00	Switch Fast Ethernet con 8 puertos 100/BaseT, 1 puerto AUI, SF
SW5	\$	2,800.00	Switch Fast Ethernet con 8 puertos 100/BaseT, 1 puerto AUI, CT
SW6	\$	6,000.00	Switch ATM, 4slot, chasis y fuente de poder, backplane 2 GB
SIMM1	\$	600.00	SIMM de memoria 8 Mb para PC.
TJA1	\$	2,200.00	Tarjeta para ruteador 2 puertos Fast Ethernet, 1 Puerto ATM.
TJA2	\$	1,500.00	Tarjeta para ruteador FRI.
TJ3	\$	2,500.00	Tarjeta para ruteador 4 puertos FDDI
TJ4	\$	5,500.00	Tarjeta para switch ATM 4 puertos ATM
TJ5	\$	900.00	Tarjeta para switch ATM 1 FRI
TJ6	\$	2,000.00	Tarjeta para switch ATM 4 puertos FDDI.
TJ7	\$	2,500.00	Tarjeta para switch ATM 4 puertos Fast Ethernet
RNT1	\$	800.00	Renta mensual de puerto FR sw una red publica Frame
RNT2	\$	2,000.00	Renta mensual proveedor de servicios Internet puerto de ruteador B-ISDN.
ICI	\$	100.00	Hora de asesoría en diseño de redes de alto desempeño
T1	\$	130.00	Tarjeta de red 10BaseT
T2	\$	150.00	Tarjeta Fast Ethernet, PCI 100BaseT

