

Gcom®
Knowledgebase
Reference

Table of Contents

Table of Contents	2
Bisync and 3270	10
Does Gcom Bisync implement 3270?.....	10
Can Gcom Bisync be either the terminal or host end of 3270?	10
Bisync and 3780	11
Does Gcom Bisync implement 3780?.....	11
Can Gcom Bisync be either the terminal or host end of 3780?	11
Bisync and adapters	12
Can I run Gcom Bisync on Gcom's T1/E1 adapter?	12
Bisync and APIs.....	13
What is the API to Gcom Bisync?	13
Bisync and standards	14
To what standards does Gcom Bisync conform?	14
Bisync and Zengin	15
Is Gcom Bisync compatible with Zengin (Japanese banking Bisync)?	15
Bisync, multiple lines, and other protocols	16
Can I run Gcom Bisync on one line and other protocols on other lines?	16
Cables and T1	17
How do I create a T1 crossover cable?	17
Frame Relay and APIs	18
What is the API to Gcom Frame Relay?.....	18
Frame Relay and DLCIs	19
Can I use a DLCI as a WAN link for TCP/IP?	19
Does Gcom Frame Relay support four-byte DLCIs?	19
How many DLCIs can I run on each Frame Relay line?.....	19
Can I run other protocols, such as LAPD, on a DLCI?	19
Frame Relay and LMI	20
What LMI standards does Gcom Frame Relay implement?	20
Can Gcom Frame Relay run without LMI?	20
Frame Relay and standards	21
To what standards does Gcom Frame Relay conform?	21
Frame Relay, multiple lines, and other protocols	22
Can I run Gcom Frame Relay on one line and other protocols on other lines?	22
Gcom Management Console	23
How do Gcom Management Console (GMC) stack, SyncSockets Daemon, and application configurations relate to each other?	23

GPA	24
Why are there silos on the GPA 2G boot-up splash screen?	24
Can I change the GPA host name to reflect our subsystem Id?	24
GPAAs and GPAsi	25
Why pay for a dedicated GPA 2Gs? Why not develop and test on a GPA 2Gsi?	25
HDLC and APIs	26
What is the API to Gcom HDLC?	26
HDLC and balanced mode	27
Does Gcom HDLC operate in balanced mode?	27
HDLC and extended address mode	28
Can Gcom HDLC run in extended address mode?	28
HDLC and extended frame types	29
Does Gcom HDLC support all the HDLC commands and responses?	29
HDLC and extended sequence mode	30
Can Gcom HDLC run in extended sequence mode?	30
HDLC and frame length restrictions	31
Does Gcom HDLC have any frame length restrictions?	31
HDLC and links	32
Can I control the link setup from my application program?	32
HDLC and multiple stations	33
Can I run multiple stations on a single physical line?	33
HDLC and raw interface	34
What is a raw interface?	34
What is raw HDLC?	34
HDLC and SREJ	35
Does Gcom HDLC support SREJ?	35
HDLC and standards	36
To what standards does Gcom HDLC conform?	36
HDLC, multiple lines, and other protocols	37
Can I run multiple lines of HDLC?	37
Can I run Gcom HDLC on one line and other protocols on other lines?	37
HDLC, primary, and secondary	38
Does Gcom HDLC operate as a primary?	38
Does Gcom HDLC operate as a secondary?	38
LAPB and APIs	39
What is the API to Gcom LAPB?	39
LAPB and extended sequence mode	40

Does Gcom LAPB support extended sequence mode?	40
LAPB and links	41
Can I configure the link addresses?	41
Can I control the link setup from my application program?	41
LAPB and SREJ	42
Does Gcom LAPB support SREJ?	42
LAPB and standards.....	43
To what standards does Gcom LAPB conform?.....	43
LAPB and timers.....	44
Does Gcom LAPB implement the T2 timer?	44
LAPB, DTE, and DCE.....	45
Can Gcom LAPB run as either DTE or DCE?	45
Can Gcom LAPB adaptively choose to be DTE or DCE?	45
LAPB, multiple lines, and other protocols.....	46
Can I run multiple lines of LAPB?	46
Can I run Gcom LAPB on one line and other protocols on other lines?.....	46
LAPB, TEST, and XID	47
Does Gcom LAPB support TEST and XID?	47
LAPD and APIs.....	48
What is the API to Gcom LAPD?	48
LAPD and links	49
Can I configure the link addresses?	49
Can I control the link setup from my application program?	49
LAPD and multiple stations	50
Can I run multiple LAPD stations on one line?	50
LAPD and SREJ	51
Does Gcom LAPD support SREJ?	51
LAPD and standards	52
To what standards does Gcom LAPD conform?	52
LAPD and timers.....	53
Does Gcom LAPD implement the T2 timer?.....	53
LAPD, DTE, and DCE	54
Can Gcom LAPD run as either DTE or DCE?	54
Can Gcom LAPD adaptively choose to be DTE or DCE?	54
LAPD, multiple lines, and other protocols	55
Can I run multiple lines of LAPD?	55
Can I run Gcom LAPD on one line and other protocols on other lines?.....	55

LAPD, TEST, and XID 56
 Does Gcom LAPD support TEST and XID? 56
 QLLC and APIs..... 57
 What is the API to Gcom QLLC? 57
 QLLC and call attempts 58
 QLLC - Can I control the amount of time between call attempts? 58
 QLLC and incoming calls..... 59
 Can Gcom QLLC listen for incoming calls? 59
 QLLC and outgoing calls 60
 Can Gcom QLLC initiate outgoing calls? 60
 QLLC and SNA 61
 Can I run SNA on top of Gcom QLLC?..... 61
 QLLC and standards 62
 To what standards does Gcom QLLC conform? 62
 QLLC, primary, and secondary..... 63
 Can Gcom QLLC act as a primary?..... 63
 Can Gcom QLLC act as a secondary? 63
 SDLC and acknowledgement delays 64
 Does Gcom SDLC have an acknowledgment delay?..... 64
 SDLC and APIs..... 65
 What is the API to Gcom SDLC? 65
 SDLC and extended sequence mode 66
 Can Gcom SDLC run in extended sequence mode? 66
 SDLC and frame length restrictions..... 67
 Does Gcom SDLC have any frame length restrictions? 67
 SDLC and links 68
 Can I control the link setup from my application program? 68
 SDLC and multiple stations 69
 Can I run multiple stations on a single physical line? 69
 SDLC and polling..... 70
 Does Gcom SDLC implement full-duplex polling? 70
 SDLC and standards 71
 To what standards does Gcom SDLC conform? 71
 SDLC, multiple lines, and other protocols 72
 Can I run multiple lines of SLDC? 72
 Can I run Gcom SDLC on one line and other protocols on other lines?..... 72
 SDLC, primary, and secondary 73

Does Gcom SDLC operate as a primary?	73
Does Gcom SDLC operate as a secondary?.....	73
SDLC, TEST, and XID.....	74
Does Gcom SDLC support TEST and XID?	74
SNA and 3270	75
Does Gcom SNA have 3270 client software?.....	75
Does Gcom SNA have a TN3270 server?	75
SNA and 5250	76
Does Gcom SNA have a TN5250 server?	76
SNA and APIs.....	77
What is the API to Gcom SNA?	77
SNA and encryption.....	78
Does Gcom SNA support cryptography, data stream compression, etc.?	78
SNA and file transfer	79
What file transfer protocols does Gcom SNA support?	79
SNA and HLLAPI.....	80
Does Gcom SNA implement HLLAPI?	80
SNA and hosts.....	81
Can Gcom SNA act as the host end of SNA?	81
SNA and LU 6.2.....	82
What node types does Gcom SNA support for LU 6.2?	82
Does Gcom SNA LU 6.2 support CP (Control Point) - CP sessions?	82
SNA and printer support.....	83
How does Gcom SNA support printing?	83
SNA and standards	84
To what standards does Gcom SNA conform?	84
SNA and XID	85
Does Gcom SNA support XID exchanges?	85
SNA, LUs, and PUs	86
What SNA LU types does Gcom SNA support?	86
What SNA PU types does Gcom SNA support?.....	86
Can I use any LU type on any PU type?.....	86
SNA, multiple lines, and other protocols	87
Can I run multiple lines of SNA?	87
Can I run Gcom SNA on one line and other protocols on other lines?.....	87
Solaris 8 and kernels	88
How do I boot 32- and 64-bit kernels?	88

Booting the 32-bit Kernel.....	88
Booting the 64-bit Kernel.....	88
How do I use 32- and 64-bit kernels?	89
SyncSockets and TCP.....	90
Can I use pure TCP to interface to the SyncSockets Daemon without using SyncSockets headers?	90
TN Server and debug	91
How can I see debug information for the TN Server's operation?	91
TN Server and file locations	92
Where are the default configuration files located?.....	92
Where are the default log files located?	92
TN Server and multiple servers.....	93
How do I configure multiple servers?.....	93
TN Server and RFPs	94
What RFPs are supported by the Gcom TN server?	94
TN Server, 5250, and 3270	95
Can I run both 5250 and 3270 protocols simultaneously?	95
I am having problems with echo while running X3270. What should I do?	95
TN Server, LUs, and PUs.....	96
How do I configure a specific host LU for access by one user?	96
How many LU/PU combinations can the server support?	96
Token Ring and adapters	97
Which adapters does Gcom Token Ring support?	97
Token Ring and operating systems.....	98
On which operating systems does Gcom support Token Ring?.....	98
Token Ring and other protocols	99
What is the Link Layer protocol used with Gcom Token Ring?	99
Can I run SNA on top of Gcom Token Ring?.....	99
Can I run TCP/IP out the same Token Ring adapter?	99
X.25 and APIs.....	100
What is the API to Gcom X.25?	100
X.25 and channels.....	101
Does Gcom X.25 support incoming only channels?	101
Does Gcom X.25 support outgoing only channels?.....	101
X.25 and extended sequence mode.....	102
Does Gcom X.25 operate in extended sequence mode?	102
X.25 and incoming calls.....	103

Can Gcom X.25 bar incoming calls?	103
Can Gcom X.25 route incoming calls to different applications depending upon the DTE address?	103
Can my application choose the line from which to receive a call?	103
How do I get the application to listen for incoming calls?	103
How does an application listen for incoming calls on a specific port on a synchronous serial adapter?	103
How does an application listen for a call to a specific address?.....	103
X.25 and negotiation.....	104
Does Gcom X.25 support flow control negotiation?.....	104
How do I configure X.25 to negotiate?.....	104
Should I use negotiation?	104
What if the other end refuses to negotiate and clears my calls if there are negotiation facilities in the call packet?	104
What settings do you use for a no-negotiation situation?.....	104
X.25 and outgoing calls	105
Can Gcom X.25 route outgoing calls to different lines depending on the DTE address?	105
Can my application choose the line on which to send a call?	105
How do I specify the called address when the application makes a call?	105
How does an application force outgoing calls to a specific port on a synchronous serial adapter?	105
X.25 and packets.....	106
Does Gcom X.25 assemble incoming M-bit packets?	106
Does Gcom X.25 break outgoing messages into packets?	106
Does Gcom X.25 support registration packets?	106
How do I know the packet size the other end can handle?	106
How do I prevent being overrun by packets I can't handle?	106
How does X.25 know the application has consumed a packet and the other end can send another packet?	106
How many data packets can be sent without acknowledgment?	106
What is the difference between a call packet and call accept packet?	106
What is the size of a data packet?	106
How do I handle non-standard X.25 packet sizes, such as used by AX.25?	106
How do I keep the calling address out of the call accept packet?	107
X.25 and PAD.....	108
Does Gcom X.25 implement the host end of a PAD?	108
Does Gcom X.25 implement the terminal end of a PAD?	108
Does Gcom X.25 implement X.3/X.28/X.29 PAD?	108
X.25 and reverse charge calls	109

Does Gcom X.25 manage reverse charge calls? 109

X.25 and standards 110

 To what standards does Gcom X.25 conform? 110

X.25 and user groups 111

 Does Gcom X.25 support closed user groups? 111

X.25 and bits 112

 Does Gcom X.25 support the D-bit? 112

 Does Gcom X.25 support the M-bit? 112

 Does Gcom X.25 support the Q-bit? 112

X.25, DTE, and DCE 113

 Can Gcom X.25 run as either DTE or DCE? 113

X.25, multiple lines, and other protocols 114

 Can I run Gcom X.25 on one line and other protocols on other lines? 114

 Can I run multiple lines of X.25? 114

 Can I run X.25 over any Link Layer protocol, not just LAPB? 114

 Can Gcom X.25 operate as a WAN link for TCP/IP? 114

X.25, PVCs, and SVCs 115

 Does Gcom X.25 support PVCs? 115

 Does Gcom X.25 support SVCs? 115

Bisync and 3270

Does Gcom Bisync implement 3270?

Yes

Can Gcom Bisync be either the terminal or host end of 3270?

Yes

Bisync and 3780

Does Gcom Bisync implement 3780?

Yes

Can Gcom Bisync be either the terminal or host end of 3780?

Yes

Bisync and adapters

Can I run Gcom Bisync on Gcom's T1/E1 adapter?

No - The T1/E1 adapter is HDLC only.

Bisync and APIs

What is the API to Gcom Bisync?

User applications use Gcom's SyncSockets® API to interface to the Gcom SyncSockets Daemon.

The SyncSockets Daemon may be running on the same system as the user application or on an external Gcom Protocol Appliance accessed via Ethernet.

Connection configurations for the SyncSockets Daemon specify per-connection parameters for Bisync and are set up using the Gcom Management Console.

The SyncSockets Daemon manages the protocol-specific part of the connection, letting you program in a protocol-independent manner.

Gcom provides sample applications in source code form. For more information, see the [SyncSockets User Guide](#).

Bisync and standards

To what standards does Gcom Bisync conform?

- IBM General Information – Binary Synchronous Communications
- IBM Doc #GA27-3004-2
- IBM 3274 Control Unit Description and Programmer's Guide
- IBM Doc #GA23-0061-2

Bisync and Zengin

Is Gcom Bisync compatible with Zengin (Japanese banking Bisync)?

Yes

Bisync, multiple lines, and other protocols

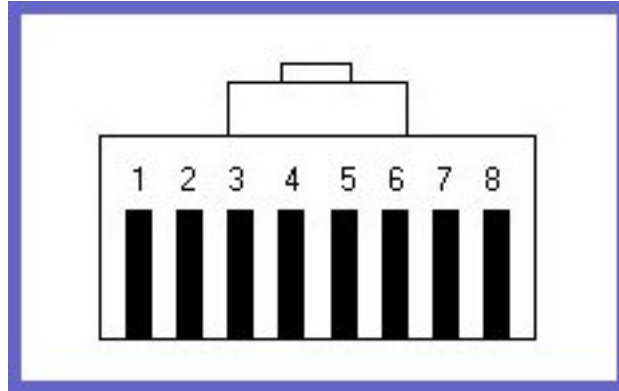
Can I run Gcom Bisync on one line and other protocols on other lines?

Yes - All Gcom protocol stacks are configured independently of each other.

Cables and T1

How do I create a T1 crossover cable?

The following diagram shows how to construct a T1 crossover cable, which can be used for loopback testing.



Pin 1	↔	Pin 4
Pin 2	↔	Pin 5
Pin 3	↔	Pin 3
Pin 4	↔	Pin 1
Pin 5	↔	Pin 2
Pin 6	↔	Pin 6
Pin 7	↔	Pin 7
Pin 8	↔	Pin 8

Frame Relay and APIs

What is the API to Gcom Frame Relay?

User applications use Gcom's SyncSockets® API to interface to the Gcom SyncSockets Daemon.

The SyncSockets Daemon may be running on the same system as the user application or on an external Gcom Protocol Appliance accessed via Ethernet.

Connection configurations for the SyncSockets Daemon specify per-connection parameters for Bisync and are set up using the Gcom Management Console.

The SyncSockets Daemon manages the protocol-specific part of the connection, letting you program in a protocol-independent manner.

Gcom provides sample applications in source code form. For more information, see the [SyncSockets User Guide](#).

Frame Relay and DLCIs

Can I use a DLCI as a WAN link for TCP/IP?

Only on Linux

Does Gcom Frame Relay support four-byte DLCIs?

Yes - Two-byte and four-byte DLCIs are configurable.

How many DLCIs can I run on each Frame Relay line?

The maximum allowed by the protocol

Can I run other protocols, such as LAPD, on a DLCI?

Yes - You can run any Gcom Link Layer protocol on top of a frame relay DLCI.

Frame Relay and LMI

What LMI standards does Gcom Frame Relay implement?

- LMI
- Annex D

Can Gcom Frame Relay run without LMI?

Yes - All DLCIs are considered active when running without LMI.

Frame Relay and standards

To what standards does Gcom Frame Relay conform?

- ANSI T1.606
- LMI
- Annex D
- T1.602 (LAPD) Core Services

Frame Relay, multiple lines, and other protocols

Can I run Gcom Frame Relay on one line and other protocols on other lines?

Yes - All Gcom protocol stacks are configured independently of each other.

Gcom Management Console

How do Gcom Management Console (GMC) stack, SyncSockets Daemon, and application configurations relate to each other?

These three kinds of configurations are distinct but can work in concert with one another.

- Stack configurations
Stacks refer to legacy protocol stacks. These include SNA, X.25, Bisync, Frame Relay, and various other Link Layer protocols. You must build and configure a legacy protocol stack to connect to a legacy network. The specific configuration options vary depending on the properties of the network and the device you use to access it.
- SyncSockets Daemon configurations
The SyncSockets® Daemon (SSD) is an application that mediates connectivity between user applications and a legacy protocol stack. If your user application uses the TCP-based SyncSockets API, you must configure at least one SSD connection.
- Application configurations
You can define application-specific command lines, including whatever options/switches can be used with the commands. Once these commands and options are defined, you can configure and save specific instances of commands. A few example application commands are included with the standard GMC installation.

After you configure a stack and an SSD, you assign them to one or more virtual ports. You may then start the ports as a group or individually.

GPA

Why are there silos on the GPA 2G boot-up splash screen?

Gcom's home office is in Savoy, Illinois, USA, near the campus of the University of Illinois.

Both the engineering college and the computer engineering college at the University are ranked in the top five in the USA. In 2005, the Institute for Higher Education at Shanghai Jia Tong University in the Peoples Republic of China ranked the University as the 25th best university in the world and the 19th best university in the nation. The University is also home to the National Center for Supercomputing Applications (NCSA) – where the web browser was invented. You can find more University of Illinois ranking information at <http://www.publications.uiuc.edu/info/rankings.html>

A number of Gcom employees are University of Illinois graduates, including its founder, Dave Grothe.

The silos are on a research farm owned by the University. The splash screen photograph shows them catching the golden rays of the sun during sunset. It is not easy to see, but the silos are actually a dark royal blue – they look black in the photo. The University of Illinois school colors are orange and blue.

Incidentally, the sunset photograph on shutdown splash screen was taken at about the same time as the silo photograph, just looking towards the sunset instead away from it.

Bottom line: We're located on the prairie, but we're NOT just a bunch of hicks!

Can I change the GPA host name to reflect our subsystem Id?

Yes – You can assign a name via DNS administration.

GPAAs and GPAsi

Why pay for a dedicated GPA 2Gds? Why not develop and test on a GPA 2Gsi?

- Testing on a development system frequently fails to find problems that occur when software is loaded on another system. Therefore, standard industry practice is to use separate development and test systems.
- Testing the USB pen drive distribution process requires a development system and target system; they cannot be the same system.
- A GPA 2Gds provides an environment compatible to that of the GPA 2Gsi and the GPA 2G, eliminating the need for a developer to manage system variables such as OS and Compiler versions.
- The GPA 2Gds is an open development platform on which a developer can load preferred tools. The GPA 2Gsi and the GPA 2G are closed production platforms; loading development tools on them voids the Gcom warranty.

HDLC and APIs

What is the API to Gcom HDLC?

User applications use Gcom's SyncSockets® API to interface to the Gcom SyncSockets Daemon.

The SyncSockets Daemon may be running on the same system as the user application or on an external Gcom Protocol Appliance accessed via Ethernet.

Connection configurations for the SyncSockets Daemon specify per-connection parameters for Bisync and are set up using the Gcom Management Console.

The SyncSockets Daemon manages the protocol-specific part of the connection, letting you program in a protocol-independent manner.

Gcom provides sample applications in source code form. For more information, see the [SyncSockets User Guide](#).

HDLC and balanced mode

Does Gcom HDLC operate in balanced mode?

Yes - It can operate multiple balanced stations on a single line.

HDLC and extended address mode

Can Gcom HDLC run in extended address mode?

Yes - It can operate with any length of address.

HDLC and extended frame types

Does Gcom HDLC support all the HDLC commands and responses?

Yes - Gcom has implemented extensions to the DLPI interface to represent the extended frame types, such as RIM, SIM, and RD.

HDLC and extended sequence mode

Can Gcom HDLC run in extended sequence mode?

Yes - This is a per-station option.

HDLC and frame length restrictions

Does Gcom HDLC have any frame length restrictions?

Only those supported by the surrounding operating system - Some STREAMS environments limit buffers to 4096 bytes; others have no limitations.

HDLC and links

Can I control the link setup from my application program?

Yes - The application program receives connection indications via the API, or it issues connection requests via the API.

HDLC and multiple stations

Can I run multiple stations on a single physical line?

Yes. An HDLC primary can obviously operate multiple stations on a single line. Gcom's HDLC secondary can also operate multiple stations on one line, each with its own address. Gcom's HDLC balanced mode can operate multiple balanced stations on a single line, each with its own pair of addresses.

HDLC and raw interface

What is a raw interface?

A raw interface lets an application send data to a communication line without using any protocol. When data arrives on the communication line, it is forwarded to the application without any processing.

You can configure a communication line for asynchronous data, Bisync data, and HDLC framed data.

What is raw HDLC?

Raw HDLC is a raw mode SyncSockets connection that passes data on a communication line to/from an application without performing any protocol processing. There are few extra bytes that actually appear on the line. The layout of an HDLC frame is:

Field	Size	Description
flag	1	Byte 0x07E indicating the beginning of the frame
Data	2-n	Data (always at least two bytes)
CRC	2	Error detection
flag	1	Byte 0x07E indicating the end of the frame

A raw HDLC connection reads/writes the Data bytes. The flags and CRC are thrown away on input and created on output. The application need not worry about generating the CRC. Any frame received with a bad CRC is thrown away by Gcom communication hardware.

HDLC and SREJ

Does Gcom HDLC support SREJ?

Yes - This is a per-station option.

HDLC and standards

To what standards does Gcom HDLC conform?

- ISO 7809
- ISO 4335
- ISO 3309
- ANSI X3.66-1979

HDLC, multiple lines, and other protocols

Can I run multiple lines of HDLC?

Yes

Can I run Gcom HDLC on one line and other protocols on other lines?

Yes - All Gcom protocol stacks are configured independently of each other.

HDLC, primary, and secondary

Does Gcom HDLC operate as a primary?

Yes - It can poll multiple secondaries on one line.

Does Gcom HDLC operate as a secondary?

Yes - This is the mode most often used to communicate with IBM hosts.

LAPB and APIs

What is the API to Gcom LAPB?

User applications use Gcom's SyncSockets® API to interface to the Gcom SyncSockets Daemon.

The SyncSockets Daemon may be running on the same system as the user application or on an external Gcom Protocol Appliance accessed via Ethernet.

Connection configurations for the SyncSockets Daemon specify per-connection parameters for Bisync and are set up using the Gcom Management Console.

The SyncSockets Daemon manages the protocol-specific part of the connection, letting you program in a protocol-independent manner.

Gcom provides sample applications in source code form. For more information, see the [SyncSockets User Guide](#).

LAPB and extended sequence mode

Does Gcom LAPB support extended sequence mode?

Yes - This is a per-line option.

LAPB and links

Can I configure the link addresses?

Yes - You can use non-standard link addresses.

Can I control the link setup from my application program?

Yes. The application program receives connection indications via the API, or it issues connection requests via the API. When LAPB is used in conjunction with Gcom's X.25 packet level, link setup is automatic and is managed by the packet level protocol module.

LAPB and SREJ

Does Gcom LAPB support SREJ?

Yes - This is a per-line option.

LAPB and standards

To what standards does Gcom LAPB conform?

- CCITT X.25 LAPB 1980, 1984, 1988
- ISO 7776

LAPB and timers

Does Gcom LAPB implement the T2 timer?

Yes - It can be configured in milliseconds or completely disabled.

LAPB, DTE, and DCE

Can Gcom LAPB run as either DTE or DCE?

Yes - This is a per-line parameter.

Can Gcom LAPB adaptively choose to be DTE or DCE?

Yes. Under a per-line option, LAPB probes the other end and changes roles as needed. It also informs the X.25 packet level of its role so the packet level can also change roles.

LAPB, multiple lines, and other protocols

Can I run multiple lines of LAPB?

Yes

Can I run Gcom LAPB on one line and other protocols on other lines?

Yes - All Gcom protocol stacks are configured independently of each other.

LAPB, TEST, and XID

Does Gcom LAPB support TEST and XID?

Yes - This is a per-line option.

LAPD and APIs

What is the API to Gcom LAPD?

User applications use Gcom's SyncSockets® API to interface to the Gcom SyncSockets Daemon.

The SyncSockets Daemon may be running on the same system as the user application or on an external Gcom Protocol Appliance accessed via Ethernet.

Connection configurations for the SyncSockets Daemon specify per-connection parameters for Bisync and are set up using the Gcom Management Console.

The SyncSockets Daemon manages the protocol-specific part of the connection, letting you program in a protocol-independent manner.

Gcom provides sample applications in source code form. For more information, see the [SyncSockets User Guide](#).

LAPD and links

Can I configure the link addresses?

Yes - This is a per-instance option.

Can I control the link setup from my application program?

Yes. The application program receives connection indications via the API, or it issues connection requests via the API. When LAPB is used in conjunction with Gcom's X.25 packet level, link setup is automatic and is managed by the packet level protocol module.

LAPD and multiple stations

Can I run multiple LAPD stations on one line?

Yes - LAPD includes a multiplexing module that splits out frame traffic based on the content of the address field.

LAPD and SREJ

Does Gcom LAPD support SREJ?

Yes - This is a per-instance option.

LAPD and standards

To what standards does Gcom LAPD conform?

CCITT Q.921

LAPD and timers

Does Gcom LAPD implement the T2 timer?

Yes - It can be configured in milliseconds or completely disabled.

LAPD, DTE, and DCE

Can Gcom LAPD run as either DTE or DCE?

Yes - This is a per-instance option.

Can Gcom LAPD adaptively choose to be DTE or DCE?

Yes. Under a per-line option, LAPD probes the other end and changes roles as needed. It also informs the X.25 packet level of its role so the packet level can also change roles.

LAPD, multiple lines, and other protocols

Can I run multiple lines of LAPD?

Yes

Can I run Gcom LAPD on one line and other protocols on other lines?

Yes - All Gcom protocol stacks are configured independently of each other.

LAPD, TEST, and XID

Does Gcom LAPD support TEST and XID?

Yes - This is a per-instance option.

QLLC and APIs

What is the API to Gcom QLLC?

User applications use Gcom's SyncSockets® API to interface to the Gcom SyncSockets Daemon.

The SyncSockets Daemon may be running on the same system as the user application or on an external Gcom Protocol Appliance accessed via Ethernet.

Connection configurations for the SyncSockets Daemon specify per-connection parameters for Bisync and are set up using the Gcom Management Console.

The SyncSockets Daemon manages the protocol-specific part of the connection, letting you program in a protocol-independent manner.

Gcom provides sample applications in source code form. For more information, see the [SyncSockets User Guide](#).

QLLC and call attempts

QLLC - Can I control the amount of time between call attempts?

Yes - This is specified in the configuration file.

QLLC and incoming calls

Can Gcom QLLC listen for incoming calls?

Yes. This is specified in the configuration file, along with the address pattern to list on. You can also listen for calls from a particular remote DTE address.

QLLC and outgoing calls

Can Gcom QLLC initiate outgoing calls?

Yes - This is specified in the configuration file, along with the address of the remote DTE.

QLLC and SNA

Can I run SNA on top of Gcom QLLC?

Yes

QLLC and standards

To what standards does Gcom QLLC conform?

IBM Doc #GA27-3345-2 - The X.25 Interface for Attaching SNA Nodes to Packet-Switched Data Networks General Information Manual

QLLC, primary, and secondary

Can Gcom QLLC act as a primary?

Yes - This is an option in the configuration file.

Can Gcom QLLC act as a secondary?

Yes - This is an option in the configuration file.

SDLC and acknowledgement delays

Does Gcom SDLC have an acknowledgment delay?

Yes - This is useful for a full-duplex secondary to give the application time to respond to an I-frame in a single polling cycle.

SDLC and APIs

What is the API to Gcom SDLC?

User applications use Gcom's SyncSockets® API to interface to the Gcom SyncSockets Daemon.

The SyncSockets Daemon may be running on the same system as the user application or on an external Gcom Protocol Appliance accessed via Ethernet.

Connection configurations for the SyncSockets Daemon specify per-connection parameters for Bisync and are set up using the Gcom Management Console.

The SyncSockets Daemon manages the protocol-specific part of the connection, letting you program in a protocol-independent manner.

Gcom provides sample applications in source code form. For more information, see the [SyncSockets User Guide](#).

SDLC and extended sequence mode

Can Gcom SDLC run in extended sequence mode?

Yes - This is a per-station option.

SDLC and frame length restrictions

Does Gcom SDLC have any frame length restrictions?

Only those imposed by the surrounding operation system - Some STREAMS environments limit buffers to 4096 bytes; others have no limitations.

SDLC and links

Can I control the link setup from my application program?

Yes - The application program receives connection indications via the API, or it issues connection requests via the API

SDLC and multiple stations

Can I run multiple stations on a single physical line?

Yes. An SDLC primary can obviously operate multiple stations on a single line. Gcom's SDLC secondary can also operate multiple stations on one line, each with its own address.

SDLC and polling

Does Gcom SDLC implement full-duplex polling?

Yes - The primary can send data to multiple stations while polling an individual station.

SDLC and standards

To what standards does Gcom SDLC conform?

IBM Doc #GA27-3093-2 - IBM Synchronous Data Link Control, General Information

SDLC, multiple lines, and other protocols

Can I run multiple lines of SLDC?

Yes

Can I run Gcom SDLC on one line and other protocols on other lines?

Yes - All Gcom protocol stacks are configured independently of each other.

SDLC, primary, and secondary

Does Gcom SDLC operate as a primary?

Yes

Does Gcom SDLC operate as a secondary?

Yes

SDLC, TEST, and XID

Does Gcom SDLC support TEST and XID?

Yes. TEST frames are passed up to the application. XID frames are passed as well, but can be analyzed as SNA XID frames when SDLC is configured to do so.

SNA and 3270

Does Gcom SNA have 3270 client software?

No

Does Gcom SNA have a TN3270 server?

Yes

SNA and 5250

Does Gcom SNA have a TN5250 server?

Yes

SNA and APIs

What is the API to Gcom SNA?

User applications use Gcom's SyncSockets® API to interface to the Gcom SyncSockets Daemon.

The SyncSockets Daemon may be running on the same system as the user application or on an external Gcom Protocol Appliance accessed via Ethernet.

Connection configurations for the SyncSockets Daemon specify per-connection parameters for Bisync and are set up using the Gcom Management Console.

The SyncSockets Daemon manages the protocol-specific part of the connection, letting you program in a protocol-independent manner.

SNA and encryption

Does Gcom SNA support cryptography, data stream compression, etc.?

The user application must encrypt or compress the data stream. The API has flags to set so the appropriate bits are set in the RH (request header).

SNA and file transfer

What file transfer protocols does Gcom SNA support?

Any file transfer protocol that a user application supports - Gcom's SNA Server code provides the SNA transport, but does not implement any particular file transfer protocol.

SNA and HLLAPI

Does Gcom SNA implement HLLAPI?

No - Most TN client packages implement HLLAPI.

SNA and hosts

Can Gcom SNA act as the host end of SNA?

Yes

SNA and LU 6.2

What node types does Gcom SNA support for LU 6.2?

LEN (Low End Node) type for LU 6.2 - Gcom is currently working on an END Node support.

Does Gcom SNA LU 6.2 support CP (Control Point) - CP sessions?

No

SNA and printer support

How does Gcom SNA support printing?

SNA supports the following LU types:

- LU 1 – Intelligent Printer Data Stream (IDPS)
- LU 3 – 3270 data stream – printers

Printers attached (either directly or via LAN connection) to the device running Gcom SPS software can be used as SNA printers as long as Gcom printer software and CUPS are also installed on the Gcom SPS device.

Printers attached to a TN client are supported through the TN protocol.

SNA and standards

To what standards does Gcom SNA conform?

IBM Systems Network Architecture Formats and Protocol Specifications

SNA and XID

Does Gcom SNA support XID exchanges?

Yes:

- XID format or format 1 for non-LU 6.2 connections
- Format 3 for LU 6.2 connections

SNA, LUs, and PUs

What SNA LU types does Gcom SNA support?

- LU 0: Application-to-application connections
- LU 1: 3270 intelligent printer support
- LU 2: 3270 terminal sessions
- LU 3: 3270 data stream printer support
- LU 4: 5250 printer support
- LU 7: 5250 terminal support
- LU 6.2 peer-to-peer application support (APPC/APPN)

What SNA PU types does Gcom SNA support?

- For LU 0, 1, 2, 3: PU type 2
- For LU 4, 7: PU type 1
- For LU 6.2: PU type 2.1

Can I use any LU type on any PU type?

No

SNA, multiple lines, and other protocols

Can I run multiple lines of SNA?

Yes

Can I run Gcom SNA on one line and other protocols on other lines?

Yes - All Gcom protocol stacks are configured independently of each other.

Solaris 8 and kernels

How do I boot 32- and 64-bit kernels?

Solaris 8 runs with either a 64-bit kernel or a 32-bit kernel.

Gcom software prior to version 1.51 does not work on the 64-bit kernel. The earlier Gcom software uses a 32-bit driver and does not work with the 64-bit Solaris 8 kernel.

Unfortunately, the Solaris 64-bit kernel does not generate an error message when it attempts to run a 32-bit driver; the driver simply does not function.

Users of Gcom software version 1.51 or later do not have this problem. The later versions of the software install and run correctly on either version of the kernel.

It is possible to install Gcom software on a machine running a 64-bit kernel and not notice anything until you try to run the Gcom software and it does not work.

Booting the 32-bit Kernel

Logged in as root, open a terminal window and type in the following:

```
# su -
# cd /
# halt -y
```

The system responds with something similar to the following:

```
syncing file systems... done
Program terminated
{2} ok
```

Type the following:

```
{2} ok boot kadb -D kernel/unix
```

The machine responds by rebooting.

Log in as root. Open up a terminal window and type the following:

```
# su -
# cd /
# eeprom boot-file=kernel/unix
# reboot
```

At this point, the Solaris 8 machine boots using the 32-bit kernel by default.

Booting the 64-bit Kernel

To change back to the 64-bit kernel, type the following.

```
# su -
# cd /
# halt -y
```

The system responds with something similar to this:

```
syncing file systems... done
Program terminated
{2} ok
```

Type the following:

```
{2} ok boot kadb -D kernel/sparcv9/unix
```

The machine responds by rebooting.

Login as root. Open up a terminal window and type the following:

```
# su -
# cd /
# eeprom boot-file=kernel/sparcv9/unix
```

```
# reboot
```

At this point, Solaris 8 machine boots using the 64-bit kernel by default.

How do I use 32- and 64-bit kernels?

The installation procedure for Gcom Protocol Suite version 1.51 or later installs software for both 32-bit or 64-bit mode, and uses symbolic links to point to the appropriate version depending upon the kernel running at the time of installation. If the 32-bit kernel is running, the links point to the 32-bit software. If the 64-bit kernel is running, the links point to the 64-bit software.

To switch between 32-bit and 64-bit versions of the software, simply boot the desired kernel.

When Solaris boots, it activates either the 32-bit or the 64-bit version of the Gcom drivers depending upon the mode of operation of the kernel (32- or 64-bit).

The Gcom utility programs (Gcom_dump, etc) and the Gcom API libraries are installed in both 32-bit and 64-bit form. When the system boots, a Gcom-provided script forms symbolic links to the correct set of commands and API libraries for the mode in which the system was booted.

If you switch from 32-bit to 64-bit, or vice versa, rebuild your application program and relink with the appropriate Gcom API libraries.

When building your application for 32-bit mode, include the following compilation options for the Sun Workshop C compiler: `-Xa`

When building your application for 64-bit mode, include the following compilation options for the Sun Workshop C compiler: `-Xa -xarch=v9 -xcode=abs64 -xregs=no%appl`

When linking your application program, always use the same Gcom API library name, e.g., `/usr/lib/gcom/npiaapi.a`. This file will be either 32-bit or 64-bit code depending on which mode the kernel was most recently booted.

If you try to link 32-bit code and 64-bit code together, the Sun Workshop linker prints error messages informing you of that fact. If you experience such errors, your application code and the Gcom API libraries are not compatible. This can happen if you:

- Change the compilation options for your application and rebuild but do not boot the Solaris kernel in the new mode.
- Reboot the kernel in the new mode but do not rebuild your application with compatible options then you will see these errors.

Gcom's experience thus far is that applications that use the Gcom API libraries to access the protocol stacks, when built as 32-bit code, function correctly with 64-bit drivers. The opposite, however, does not hold. A 64-bit application does not function correctly with 32-bit drivers.

When running on a dual mode processor, such as an Ultra SPARC 60 or 80, Gcom has found that code runs up to 30% faster in 32-bit mode than in 64-bit mode. We do not have any test results for the new 64-bit only SPARC CPUs. They may run faster than a 32-bit machine at the same clock rate, but we have no data to support this.

If your application program functioned in 32-bit mode but seems to not work in 64-bit mode, consult Sun's documentation - at <http://docs.sun.com/ab2/coll.45.13/SOL64TRANS/@Ab2TocView?Ab2Lang=C&Ab2Enc=iso-8859-1> - on 64-bit computing to ensure your application is written in such a manner that it is portable between 32-bit and 64-bit mode.

SyncSockets and TCP

Can I use pure TCP to interface to the SyncSockets Daemon without using SyncSockets headers?

No - Certain information in the SyncSockets® headers is necessary for applications to properly receive messages delimited by legacy protocol configurations.

TN Server and debug

How can I see debug information for the TN Server's operation?

In the configuration file for the server, set parameters in the following manner:

- `tndmon_trace = 1`
- `tnsvr_trace = 1`
- `tnpars_trace = 1`
- `tnopts_trace = 1`
- `netif_trace = 1`
- `snaif_trace = 1`
- `exit_trace = 1`
- `luman_trace = 1`
- `tnmain_trace = 1`
- `tne_trace = 1`
- `snaapi_log_options = 0x7d`

TN Server and file locations

Where are the default configuration files located?

The distribution disk places the configuration file in `/usr/lib/gcom/sna`. The default files are:

- `svrconf.lnx.3270`
- `svrconf.lnx.5250`
- `luconfig.3270`
- `luconfig.5250`

The `svrconf.lnx.*` files configure the telnet server. The `luconfig.*` files configure the mapping between client LUs and host LUs.

Where are the default log files located?

The log files are placed in `/usr/spool/gcom`. Two log files can be enabled:

- `Api.log` contains operational debug information for the telnet server.
- `Lumanager.sts` contains information about LU activation/deactivation.

TN Server and multiple servers

How do I configure multiple servers?

You need two `svrconf.lnx` files and two `luconfig` files. In the `svrconf.lnx`, set a unique `port/PU/luconfig_file` for each server. Then create a second `start.tnsvr` script. For example:

```
svrconf.lnx.server1 -> set port 10000, set luconfig_file ->
luconfig.server1, set PU (PU number from host)
start.sna.server1 -> set -f/usr/lib/gcom/sna/svrconf.lnx.server1
svrconf.lnx.server2 ->
set port 10001, set luconfig_file -> luconfig.server2, set PU (PU
number from host)
start.sna.server2 -> set -f/ usr/lib/gcom/sna/svrconf.lnx.server2
```

TN Server and RFPs

What RFPs are supported by the Gcom TN server?

- RFP 1576 - TN3270 Current Practices
- RFP 1647 - TN3270 Enhancements
- RFP 1646 - TN3270 Extensions for LUsername and Printer Selection

TN Server, 5250, and 3270

Can I run both 5250 and 3270 protocols simultaneously?

Yes. You need only configure one server for 3270 and one for 5250.

I am having problems with echo while running X3270. What should I do?

In your configuration file, add or set the following parameter: `using_x3270_client = 1`

TN Server, LUs, and PUs

How do I configure a specific host LU for access by one user?

You need the LU number from the host that the user will access. In the `luconfig` file, under the `[SPECIFIC_NAMES]` header, add an entry that uses that LU number. The label is what will be entered by a user when connecting to the telnet server. Make sure the selected LU is not part of a `[GENERIC_NAMES]` group.

How many LU/PU combinations can the server support?

This is limited only by the amount of memory in your server computer. 3270 limits 254 LUs per PU. 5250 limits 16 LUs per PU.

Token Ring and adapters

Which adapters does Gcom Token Ring support?

- The IBM 16/4 ISA and PCI adapters on Linux
- Any adapter for which QNX has a driver on QNX

Token Ring and operating systems

On which operating systems does Gcom support Token Ring?

Linux and QNX

Token Ring and other protocols

What is the Link Layer protocol used with Gcom Token Ring?

LLC-II

Can I run SNA on top of Gcom Token Ring?

Yes - Token Ring is intended for use with SNA.

Can I run TCP/IP out the same Token Ring adapter?

Yes

X.25 and APIs

What is the API to Gcom X.25?

User applications use Gcom's SyncSockets® API to interface to the Gcom SyncSockets Daemon.

The SyncSockets Daemon may be running on the same system as the user application or on an external Gcom Protocol Appliance accessed via Ethernet.

Connection configurations for the SyncSockets Daemon specify per-connection parameters for Bisync and are set up using the Gcom Management Console.

The SyncSockets Daemon manages the protocol-specific part of the connection, letting you program in a protocol-independent manner.

Gcom provides sample applications in source code form. For more information, see the [SyncSockets User Guide](#).

X.25 and channels

Does Gcom X.25 support incoming only channels?

Yes - Over the full channel range

Does Gcom X.25 support outgoing only channels?

Yes - Over the full channel range

X.25 and extended sequence mode

Does Gcom X.25 operate in extended sequence mode?

Yes - This is a per-line option.

X.25 and incoming calls

Can Gcom X.25 bar incoming calls?

Yes

Can Gcom X.25 route incoming calls to different applications depending upon the DTE address?

Yes. The application can specify a wildcarded listen address when performing the NPI bind operation. Listen patterns for different applications must be different from one another.

Can my application choose the line from which to receive a call?

Yes - The application can restrict its listen pattern to calls that arrive on a particular line.

How do I get the application to listen for incoming calls?

In the SyncSockets® configuration file, set `listen` to 1.

How does an application listen for incoming calls on a specific port on a synchronous serial adapter?

In the SyncSockets® configuration file, set `line_number` to the port.

How does an application listen for a call to a specific address?

In the SyncSockets® configuration file, set `listen_address` to the desired address.

Note: You can use wildcards to listen for calls to a range of addresses. You can also use data masks to listen for specific values in the X.25 data field.

See X.25 SyncSockets SS_OP_LISTEN Message in the [SyncSockets User Guide](#) for more information.

X.25 and negotiation

Does Gcom X.25 support flow control negotiation?

Yes, via the packet size and window size facilities in the call setup packets. It contains a variety of options to limit this negotiation because (far too) many X.25 implementations do not support these negotiations.

How do I configure X.25 to negotiate?

The default mode of operation for Gcom X.25 does negotiation (as prescribed in the standard).

Should I use negotiation?

Yes - Your system will be more robust. If someone changes the X.25 configuration at the other end, the connection still works.

What if the other end refuses to negotiate and clears my calls if there are negotiation facilities in the call packet?

Set the following bit-encoded x2 options:

Option Byte	Option Bit	Option Set No.	Description
x2_12	0x40	G.23	Suppress flow control facilities in call packets
x2_16	0x80	G.56	Suppress flow control facilities in call request packet only.

What settings do you use for a no-negotiation situation?

Set the following bit-encoded x2 options:

Option Byte	Option Bit	Option Set No.	Description
x2_12	0x40	G.23	Suppress flow control facilities in call packets.
x2_16	0x40	G.55	Call accept flow control facilities depend on call request packet.
x2_16	0x80	G.56	Suppress flow control facilities in call request packet only.
x2_17	0x01	G.57	Suppress calling address in call accept packet.

X.25 and outgoing calls

Can Gcom X.25 route outgoing calls to different lines depending on the DTE address?

Yes. The configuration file contains routing entries with wildcarded address patterns used for this purpose. There are also cost factors that can be configured - the call router orders potential routes from least to most costly.

Can my application choose the line on which to send a call?

Yes - The call routing can be overridden with an explicit directive to send the call on a particular line.

How do I specify the called address when the application makes a call?

In the SyncSockets® configuration file, set the `remote_address` value - This becomes the called address. (The `local_address` in the SyncSockets configuration file becomes the calling address.)

In the sample below, 3300201 is the called address and 4400444 is the calling address:

```
line_number = 2
receipt_confirmation = 1
local_address = "4400444"
use_fast_select = 0
user_data = ""
fast_select_restricted = 0
reset_handling = 0
remote_address = "3300201"
listen = 0
user_data_mask = ""
accept_with_address = 0
max_message_size = 256
facilities = ""
logsize = 64000
logopts = 0
logfile = "gpi.log"
```

How does an application force outgoing calls to a specific port on a synchronous serial adapter?

In the SyncSockets® configuration file, set `line_number` to the port.

X.25 and packets

Does Gcom X.25 assemble incoming M-bit packets?

No - Packets are passed to the application with the M-bit set as it was received from the X.25 line.

Does Gcom X.25 break outgoing messages into packets?

Yes. The application can give the API a pointer to a buffer of any size. The API breaks the buffer into a packet sequence using the M-bit. The packet size used is the size negotiated at call setup.

Does Gcom X.25 support registration packets?

Yes - You can configure which facilities are negotiable and the initial settings for the negotiations.

How do I know the packet size the other end can handle?

You can't!

The X.25 standard specifies a way for two ends to agree on packet and window size:

- One end sends a call packet specifying a preferred packet size and window size.
- The other end either:
 - Rejects the call (with a call clear).
 - Or accepts the call and specifies an acceptable packet size and window size.

So a call could specify a window of 7 2048-byte packets; and the other end could accept a window of 2 256-byte packets.

All X.25 implementations must accept a packet size of 128 and a window size of 2. The standard negotiates toward this default. So a window of 7 2048-byte packets could be accepted as 2 128-byte packets; but will never be accepted as 7 32-byte packets.

How do I prevent being overrun by packets I can't handle?

The transmission stops as soon as the window closes.

How does X.25 know the application has consumed a packet and the other end can send another packet?

When the application reads a SyncSockets® data message, the SyncSockets Daemon informs X.25. X.25 then notifies the other end (using an RR packet) that it can send another packet.

How many data packets can be sent without acknowledgment?

From 1 to 7 - This is what window size specifies; however, you can configure extended mode X.25 which can go from 1 to 127.

What is the difference between a call packet and call accept packet?

A *call packet* is used to request a connection from the other end. The other end indicates acceptance of the connection by sending a *call accept packet* (or a call clear packet if it doesn't accept).

What is the size of a data packet?

From 32 to 2,048 bytes

How do I handle non-standard X.25 packet sizes, such as used by AX.25?

X.25, including Gcom X.25, does not operate with packet sizes other than a power of two. So, for example, 128 and 256 are valid packet sizes, but 240 is not.

Gcom X.25 looks at the packet size configured by the user and internally uses the largest power of two not greater than the configured packet size. Thus, if you configure packet size of 240, Gcom X.25 uses 128 as the true X.25 packet size.

In CCITT and ISO X.25, packet size negotiations are performed using the exponent of the power of two that you want for your packet size. Thus, if you want 256-byte packet sizes, the negotiated number is 8, not 256. This is essentially why packet sizes must be a power of two.

Standard X.25 also specifies that data packets that have the M-bit set must be full packets, unless the packet also has the D-bit set. That would imply that you cannot have a data packet with M-bit 1, D-bit 0, and 240 bytes of data.

If you need to interact with a network or another DTE that is operating with non-power of two packet sizes, say 240 bytes, do the following.

- Configure the Gcom X.25 for 256-byte packets. This allows 240-byte packets from the other end to be received with no packet size error.
- Set the G.58 packet level *performed* option (Gcom 2.2 and later version). This causes Gcom X.25 to relax the restriction that data packets with the M-bit set must be full packets. This allows the receipt (and sending) of data packets with 240 bytes and the M-bit set. (The value 240 is just an example. Other non-full sizes work just as well.)
- The application code must refrain from sending any packet larger than 240 bytes of data. This prevents the other end from rejecting the packets as exceeding their imagined 240-byte maximum length.
- If the application needs to send larger complete packet sequences, it must carve them up into 240-byte packets, set the M-bit in all but the final packet, and send them using SyncSockets® or NPI API. The SyncSocket connection must have the "Inhibit Packetizing" parameter set to "Yes".

How do I keep the calling address out of the call accept packet?

Set the following bit-encoded x2 option:

Option Byte	Option Bit	Option Set No.	Description
x2_17	0x01	G.57	Suppress calling address in call accept packet.

X.25 and PAD

Does Gcom X.25 implement the host end of a PAD?

Yes, using the `Gcom_hpad` program that runs as a daemon

Does Gcom X.25 implement the terminal end of a PAD?

Yes, using the interactive program `Gcom_tpad`

Does Gcom X.25 implement X.3/X.28/X.29 PAD?

Yes, using an application program

X.25 and reverse charge calls

Does Gcom X.25 manage reverse charge calls?

Yes. There are options to disable acceptance of reverse charge calls. There are also options to force reverse charging for all outgoing calls. The application program can examine the facilities of the incoming call (including reverse charging) and either accept or clear the incoming call.

X.25 and standards

To what standards does Gcom X.25 conform?

- CCITT X.25 1980, 1984, and 1988
- ISO 8208 with the exception of the data mode timers, which are not implemented

X.25 and user groups

Does Gcom X.25 support closed user groups?

Yes

X.25 and bits

Does Gcom X.25 support the D-bit?

Yes. The D-bit is made visible at the API. Also, the D-bit modification facility for DCE mode is implemented.

Does Gcom X.25 support the M-bit?

Yes - The M-bit is made visible at the API.

Does Gcom X.25 support the Q-bit?

Yes - The Q-bit is made visible at the API.

X.25, DTE, and DCE

Can Gcom X.25 run as either DTE or DCE?

Yes. This is a per-line option. The packet level can also be made to change roles depending on the results of role negotiation by the link layer.

X.25, multiple lines, and other protocols

Can I run Gcom X.25 on one line and other protocols on other lines?

Yes - All Gcom protocol stacks are configured independently of each other.

Can I run multiple lines of X.25?

Yes

Can I run X.25 over any Link Layer protocol, not just LAPB?

Yes - It is immaterial to X.25 whether these link layers are LAPB, LAPD or even SDLC stations.

Can Gcom X.25 operate as a WAN link for TCP/IP?

No

X.25, PVCs, and SVCs

Does Gcom X.25 support PVCs?

Yes - Over the full channel range

Does Gcom X.25 support SVCs?

Yes - Over the full channel range