

Frequently Asked Questions (about Corelatus Hardware and Software)

Revision: 1.15

The latest version, in both PDF and HTML formats, is available from
<http://www.corelatus.com/>

Contents

1	Questions about the operating system	2
1.1	How can I shut down a GTH?	2
1.2	Why are log timestamps during boot zero?	2
2	Questions about Ethernet and IP	2
2.1	Why is the GTH sending UDP packets to port 9 on my network?	2
2.2	Why is putting both ethernet interfaces on the same subnet not recommended?	2
2.3	Why can't I reach a GTH on the other side of a gateway?	3
2.4	Does the GTH support Ethernet speed auto-negotiation?	3
2.5	What happens if I connect a 100Mbit Ethernet port to a switch which does not support auto-negotiation?	3
3	Questions about G.703 E1/T1 Layer 1	4
3.1	What do the E1/T1 Layer 1 (L1) states mean?	4
3.2	What do the E1/T1 L1 Error counters represent?	5
3.3	Can you explain, briefly, how sync works in PDH networks?	6
3.4	What happens when there is a slip?	8
3.5	How does sync work on the GTH?	8
3.6	How is "ppm" related to packet loss rate?	8
3.7	Why is Ethernet "immune" from slips?	9
4	Questions about Layer 2 Signalling	10
4.1	Can you summarise the basic layer 2 (L2) functions in the transmission systems used by GTH hardware?	10
4.2	Why does MTP-2 produce data even when L1 indicates LOS?	11
4.3	Why does my application occasionally receive MTP-2 packets with length=0?	11

4.4	What do the internal error messages relating to signalling sockets mean?	11
5	Hardware Questions	12
5.1	How can I identify Corelatus hardware?	12
5.2	How do I make a straight Ethernet cable?	13
5.3	How do I make a crossed Ethernet cable?	14
5.4	How do I make a crossed E1/T1 cable?	15
6	Historical Hardware questions	16
6.1	What are the differences between hardware revisions and models?	16
6.2	What do the rotary switches on IEB modules do?	16
7	Logistics Questions	16
7.1	How do I recycle end-of-life systems and packaging?	17
7.2	What are the packaging dimensions?	17
7.3	Which customs classification number do you use?	17

1 Questions about the operating system

1.1 How can I shut down a GTH?

Turn the power off. The GTH does not require any special preparation before power is turned off.

The GTH can be rebooted using the API `<reset>` command.

1.2 Why are log timestamps during boot zero?

The realtime clock is set during the boot process, so all items logged before the clock is set show the time since boot instead of "wall time":

```
Jan  1 00:00:06 (none) syslog.info klogd: Frame Relay RX driver
Jan  1 00:00:07 (none) syslog.info klogd: stream_dev
Aug 12 14:03:10 (none) user.info fpga: Programming fpga
Aug 12 14:03:11 (none) daemon.info klogd: Power 1 on
Aug 12 14:03:11 (none) daemon.info klogd: Power 2 off
```

In the above example, the clock was set immediately before the "Programming fpga" entry.

2 Questions about Ethernet and IP

2.1 Why is the GTH sending UDP packets to port 9 on my network?

GTHs send one packet per minute on all activated ethernet interfaces if nothing has connected to the WWW server (at port 8888) or the API (port 2089).

This is useful if you have a GTH and don't know what its IP address is. A network sniffer (such as tcpdump or wireshark) will show these 'chirp' packets.

Wireshark is available for most operating systems from <http://www.wireshark.org/>.

2.2 Why is putting both ethernet interfaces on the same subnet not recommended?

Whenever the OS sends a packet, it looks at its routing tables to decide which interface to send the packet on. If you have two interfaces on the same subnet, the first interface which matches the subnet will be used to send *all* the packets, which is probably not what you wanted.

References: <http://www.netsys.com/sunmgr/1996-09/msg00083.html>

RFC 2328 (OSPF)

<http://support.microsoft.com/support/kb/articles/q244/2/68.asp>

There are various (trouble-prone) ways to work around this restriction.

2.3 Why can't I reach a GTH on the other side of a gateway?

By default, GTH systems are configured without a default gateway (route). You can only reach a GTH which is on the same subnet as the host.

A default gateway can be configured using API commands.

2.4 Does the GTH support Ethernet speed auto-negotiation?

Current model (GTH 2.0, shipping since 2006) hardware supports Ethernet auto-negotiation on both Ethernet interfaces. Depending on the capabilities of the Ethernet port the GTH is connected to, it run either 10 or 100Mbit/s, either half- or full-duplex.

2.5 What happens if I connect a 100Mbit Ethernet port to a switch which does not support auto-negotiation?

The GTH will sense the link speed and run at that speed (either 10 or 100Mbit/s). The GTH will assume half-duplex. This behaviour is required by IEEE 802.3:2000.

Some 10/100Mbit/s Ethernet devices, such as high-end switches, allow auto-negotiation to be disabled and the link speed/duplex setting to be manually configured. Several possible configurations will result in problems:

Remote port configuration	GTH will use	Comment
auto-negotiate	100Mbit/s full-duplex	Good (recommended)
100Mbit/s full-duplex	100Mbit/s half-duplex	Not recommended (see below)
100Mbit/s half-duplex	100Mbit/s half-duplex	Acceptable
10Mbit/s full-duplex	10Mbit/s half-duplex	Not recommended (see below)
10Mbit/s half-duplex	10Mbit/s half-duplex	Acceptable

The **not recommended** configurations result in network performance problems, including late collisions.

References:

IEEE 802.3:2000 (The "Fast Ethernet" specification)

[CISCO tech note 17053: "Troubleshooting Cisco Catalyst Switches to NIC Compatibility Issues"](#)

[comp.dcom.lans Ethernet FAQ](#)

3 Questions about G.703 E1/T1 Layer 1

3.1 What do the E1/T1 Layer 1 (L1) states mean?

LOS Loss of signal. There is no signal strong enough to detect on the incoming line. In North America, this is often called a “red alarm”.

Typical causes:

- no cable plugged in, or cable plugged in to the wrong port
- incorrectly wired cable
- incorrect monitor point construction (i.e. the monitor point's attenuation lies outside the nominal 17-23dB range)
- The GTH system not configured for monitoring even though the input signal is -20dB from nominal.

LFA Loss of frame alignment. There is a signal on the line, but E1 (or T1) framing cannot be recovered.

Typical causes:

- Incorrect L1 configuration, e.g. the incoming line may be T1 but the GTH is configured for E1, or the incoming line may be doubleframe but the L1 is configured for multiframe.
- Attempting to monitor a signal which is neither E1 nor T1. For instance, attempting to monitor an HDSL modem on its analog line port instead of its E1 port will result in LFA.

LMFA Loss of multiframe alignment. There is a signal on the line and framing can be recovered, but multiframe framing cannot be recovered.

Typical cause: incorrect L1 configuration

AIS Alarm indication signal. The E1 receiver in the GTH has detected a fault on the signal it is receiving (too many 1-bits in a frame). This is often called a “blue alarm” in North America.

Typical cause: incorrect L1 configuration

RAI Remote alarm indication. The E1 receiver in the equipment being monitored (i.e. not the GTH) has detected an LOS, LFA or LMFA condition in the signal it is receiving. This is often called a “yellow alarm” in North America.

Typical causes **when using a monitor point**:

- Incorrectly wired or incorrectly installed monitor point which interrupts the signal on the line being monitored.
- Fault in the line being monitored
- Fault in the transmitter on the *other* pair of the line being monitored.
- Fault in the receiver on the *other* pair of the line being monitored.

Typical causes **when using full-duplex (normal operation)**:

- Incorrectly wired cable
- GTH L1 parameters incorrectly configured, e.g. monitoring enabled, or configured for T1 when attached to an E1 system.

3.2 What do the E1/T1 L1 Error counters represent?

The GTH maintains a number of error counters for each receiver. The counters can be used to help diagnose a number of L1 problems.

Slip The number of positive and negative slips. Slips are covered in more detail in later questions.

Typical causes:

- The operator's PDH network is not maintaining correct internal sync. Understanding the operator's sync topology is a pre-requisite to troubleshooting this problem.
- GTH is (mis-)configured to use its internal frequency reference when it should be configured as a sync slave to one of the E1/T1 inputs.
- GTH is unable to find a sync source, for instance because the framing configuration is incorrect. One example is configuring GTH for multiframe when the link is actually doubleframe.

Frame Error The incoming signal did not conform to the expected framing.

Typical causes:

- Incorrect framing configuration, e.g. GTH configured for multiframe, whereas the link is actually doubleframe.
- Incoming signal is not an E1 signal, e.g. attempting to monitor a DSL line can give this symptom.
- Incoming signal contains too much noise to decode correctly.

Code violation The incoming signal did not conform to the expected signal coding rules.

Typical causes:

- Incorrect coding configuration, e.g. GTH configured for HDB3 coding (used on all E1 links) but the incoming signal is actually T1 (which uses B8ZS coding). In most cases, this will cause hundreds or thousands of code violations per second.
- Incoming signal is not an E1 signal, e.g. attempting to monitor a DSL line can cause this symptom.
- Excessively weak incoming signal. Use a handheld signal tester or an oscilloscope to verify that the signal strength is as expected.
- Incorrect monitoring configuration, e.g. GTH is configured for -20dB monitoring, but the incoming signal is not attenuated.

- Line noise/interruption. A small number of code violations is normal. In some installations, this can be as low as five or ten code violations per link per month. Error rates greater than a few tens of violations per day are usually worth investigating.

CRC error The Layer 1 data integrity check (a CRC check) failed. The counter shows how many times the check failed.

This counter is only used when the GTH is configured for multiframe framing.

Typical causes:

- Line noise. A small number of CRC errors is normal. In some installations, this can be as low as five or ten per link per month.
- Incorrect framing configuration, e.g. GTH configured for multiframe, link is actually doubleframe.

3.3 Can you explain, briefly, how sync works in PDH networks?

E1 (and T1) lines are part of the 'PDH' network. PDH stands for plesiosynchronous digital hierarchy, meaning that the whole network runs with "almost" synchronised clocks.

Every switch in a PDH network must have a source for its 8kHz sync. Possible sync sources are:

- An internal oscillator
- An E1 cable
- A special sync input

PDH requires all switches in a network to run "almost" synchronised. In practice, this is done by specifying master/slave sync relationships between the switches. A simple, correctly structured network is shown in figure 1:




Figure 1: A simple, correct, sync network

The lines indicate E1 cables between switches. The red lines with arrows indicate a master/slave sync relationship: switch 4 uses switch 2 as its sync master. Switch 1 has no master, it is the sync reference for the whole network.

There are many ways to mis-configure networks such that sync will not work.




Figure 2: Two sync masters

In figure 2, neither switch 1 nor switch 2 has a master. Both will generate their frequency reference internally. In practice, these two frequency references will be different, resulting in slips on E1s 1-2, 2-3 and 3-4. This is **bad**. There will be no slips on 1-3 and 2-4 since the switches on either side of those links are synchronised.




Figure 3: All switches unsynchronised

Figure 3 shows another **bad** configuration. None of the switches are synchronised. There will be slips on all of the E1s.

Figure 4 is the **worst** of all the examples. It shows a circular sync relationship, where each of the switches chases a sync source which ultimately originates from itself. Such setups are unstable and result in unpredictable behaviour. A common result of such a configuration is that the frequency migrates to the extreme of the sync range the equipment can handle, thus taking it outside the frequency limits allowed by G.703.




Figure 4: A circular sync relationship

The ITU-T standard *G.781* contains further information about synchronisation hierarchies in PDH networks.

3.4 What happens when there is a slip?

A slip occurs when a transmitter emits frames at a different rate to the rate the receiver accepts the frames. When the transmitter's rate is greater than the receiver's, the receiver deals with the situation by discarding one frame every so often.

Conversely, when the transmitter's rate is lower than the receiver's the receiver duplicates a frame every so often.

On E1, a frame is 32 octets of data, one octet per timeslot.

On voice circuits, a slip results in an audible 'pop'. On signalling circuits it results in corrupted packets. In MTP-2, a slip will (almost) always result in one corrupted packet.

3.5 How does sync work on the GTH?

This is described in detail in section 7.2 of the API manual ("Sync sources"). By default, the GTH will automatically select an E1/T1 input to synchronise to. The GTH will adjust its internal frequency to match the chosen interface, as long as the E1/T1 line's frequency is within the limits specified by G.703. (+/- 50ppm)

The GTH's sync behaviour can also be manually configured, either to use a particular line as its sync source, or to use its internal oscillator. Example: to set the sync source to PCM2A:

```
<set name="sync">
  <attribute name="source" value="pcm2A"/>
</set>
```

The stability of the GTH's internal oscillator is +/- 1 ppm, within the normal temperature range. This stability determines the performance in situations where the sync source is temporarily lost, in practice it means the GTH can tolerate sync loss for approximately two minutes without having to slip.

3.6 How is "ppm" related to packet loss rate?

ppm means parts-per-million. Thus if one system runs at 8000Hz and another at 8000.2Hz, they differ by 250ppm. The time between slips is given by $1 / (d \times R)$, where 'd' is the frequency error and R is the frame rate.

Examples:

d	R	Time between slips
1ppm	8000 frames/s	125 seconds
50ppm	8000	2.5 seconds

In MTP-2, packets are transmitted continuously, so a slip will (almost) always damage a packet. On an "idle" MTP-2 link, the wire has nonstop 5 octet FISUs, a 50ppm frequency error should give about 24 damaged FISUs per minute. These are counted as ESUs.

On a "busy" MTP-2 link, the average packet length will be longer than 5 octets, resulting in fewer ESUs per minute.

3.7 Why is Ethernet "immune" from slips?

10Mbit/s and 100Mbit/s ethernet links re-synchronise at the start of each transmitted packet.

In most applications, this means there are never slips. In applications which transmit fixed-rate data, e.g. a VOIP system, slips still occur. They are usually dealt with in an application-specific way. VOIP systems often use "silence compression/expansion" to hide the slips.

4 Questions about Layer 2 Signalling

4.1 Can you summarise the basic layer 2 (L2) functions in the transmission systems used by GTH hardware?

	Ethernet/IP	E1/HDLC	E1/ATM
What happens when there is no data to send?	Electrical silence (no voltage on line)	MTP-2: send FISU, LAPD/FR: send FLAG	Send "idle" cell or "unassigned" cell
How does the hardware find the start of a packet?	Watch for the transition from electrical silence to a "preamble" sequence which MUST come before every packet. The preamble sequence is 80 bits of 101010... It has the secondary function of synchronising clocks.	Find a flag (bit pattern 01111110) followed by a non-flag. Bit stuffing guarantees that FLAG never occurs inside a packet.	Next packet (cell) always comes exactly 53 octets after start of current. At startup, we find the start of a packet by just starting somewhere and computing the header CRC. If it's wrong, we move forward one octet and try again. When we find a correct CRC, we move forward 53 octets and see if that header is also correct.
What are the alignment requirements?	(Not relevant)	No alignment requirements at all, not even octet alignment.	Cells are octet aligned, but not frame aligned. A cell must always start on an octet boundary. A cell will usually NOT start on a frame boundary.
How are bit errors detected?	32-bit CRC after every frame (packet)	16-bit CRC after every signal unit (packet)	Cell header (5 octets) protected by an 8-bit CRC. AAL0: No payload protection. AAL2: No payload protection (but 5 bit CRC on secondary header). SSCS sublayer <i>may</i> add a 10-bit CRC. AAL5: 32-bit CRC on payload
How is packet length limited?	Ethernet frame is about 1.5k max. TCP can reassemble frames with no length limit. UDP limits to 64k	MTP-2: 279 octets. LAPD: 265 octets by default Frame relay: generally 1.6k	AAL0: Cells are always 53 octets. AAL2: 45 octets by default (64 octets optional) AAL5: 64k

4.2 Why does MTP-2 produce data even when L1 indicates LOS?

The LOS indication is a warning; it means "the signal is now so weak that correct data extraction cannot be guaranteed". LOS does not itself disable the incoming data stream; the decision whether or not to terminate L2 if L1 indicates LOS is left to the application.

4.3 Why does my application occasionally receive MTP-2 packets with length=0?

With ESU filtering disabled (ESU="yes"), the GTH delivers *all* signal units (packets), including signal units which are less than the minimum length. Some examples of input bitstreams which cause signal units with a zero-byte length:

Four-bit signal unit (zero bytes!): 0111 1110 **1011** 0111 1110

One-bit signal unit (zero bytes!): 0111 1110 **1011** 1111 0

Instantly aborted signal unit (also zero bytes): 0111 1110 1111 1111

The same applies to LAPD and Frame Relay.

4.4 What do the internal error messages relating to signalling sockets mean?

When sending signalling data, for instance in an MTP-2 monitoring application, some information about sockets going down is logged. This is intended for internal use, but they might be of some use when solving problems. The possible messages are:

```
killed sock N fd M (poll() event: error/input), errno P
```

There are two possible reasons for this message.

1. The socket was closed remotely, either by the application or the OS.
2. Input was received for transmission to a timeslot configured for monitoring.

```
killed sock N fd M (write failed), errno P
```

This happens if an attempt to write a socket failed with an error. Some values of *P* indicate specific problems:

Code	Cause
32	The reading end of the socket has closed
104	The remote OS reset the socket (typical behaviour when a remote program is terminated or a server is rebooted)
110	The remote connection timed out. This is typical of an ethernet network failure, e.g. a cable has been unplugged.

Most other values do not indicate an error, in particular 0, 4 and 11 are normal.

In all cases, preceding or following error messages may give provide extra information.

5 Hardware Questions

5.1 How can I identify Corelatus hardware?

There are several identification codes in Corelatus products.

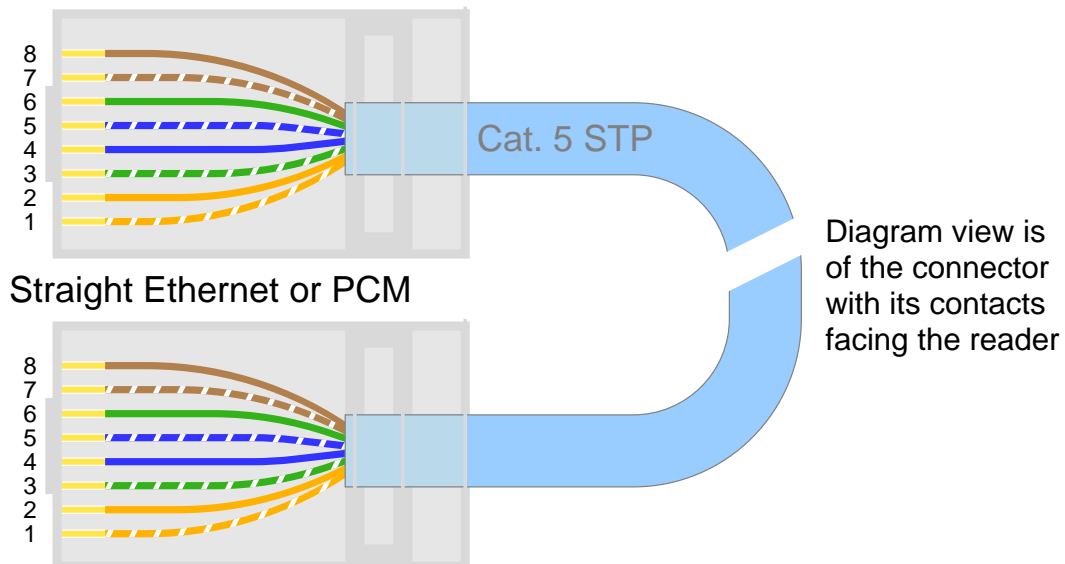
Name	Location	Purpose
Chassis serial number (CSN)	On a metallic silver self-adhesive on the left side of the grey steel chassis, immediately following the text "Serial:".	Uniquely identifies each chassis. At delivery time, Corelatus supplies a list of cards in each chassis.
Ethernet MAC address	Factory-programmed on GTH modules. Can be viewed via the in-built webserver (on the "ethernet" page)	Uniquely identifies each ethernet interface.
ROM ID	Factory-programmed on GTH and IEB modules. Can be viewed via the in-built webserver	Uniquely identifies each GTH or IEB module.

Corelatus maintains a database which cross-references all of the above information for each module.

5.2 How do I make a straight Ethernet cable?

A straight cable is most often used to connect an Ethernet interface to an Ethernet hub or switch. Shielded CAT-5 cable (STP) must be used, the connectors are RJ45.

End A Pin	Cable Wire Colour	End B Pin
1	orange-white	1
2	orange	2
3	green-white	3
4	blue	4
5	blue-white	5
6	green	6
7	brown-white	7
8	brown	8




The wire colours have been chosen to follow one of the schemes allowed in EIA/TIA 568 A/B.

5.3 How do I make a crossed Ethernet cable?

A crossed Ethernet cable can be used to connect a GTH directly to a server's Ethernet port, without using a hub or switch. This is sometimes also called a crossover cable. Shielded CAT-5 cable must be used, the connectors are RJ45.

End A		Cable Wire Colour	End B	
Pin	Name		Pin	Name
1	TX+	orange-white	3	RX+
2	TX-	orange	6	RX-
3	RX+	green-white	1	TX+
4	Termination	blue	4	Termination
5	Termination	blue-white	5	Termination
6	RX-	green	2	TX-
7	Termination	brown-white	7	Termination
8	Termination	brown	8	Termination

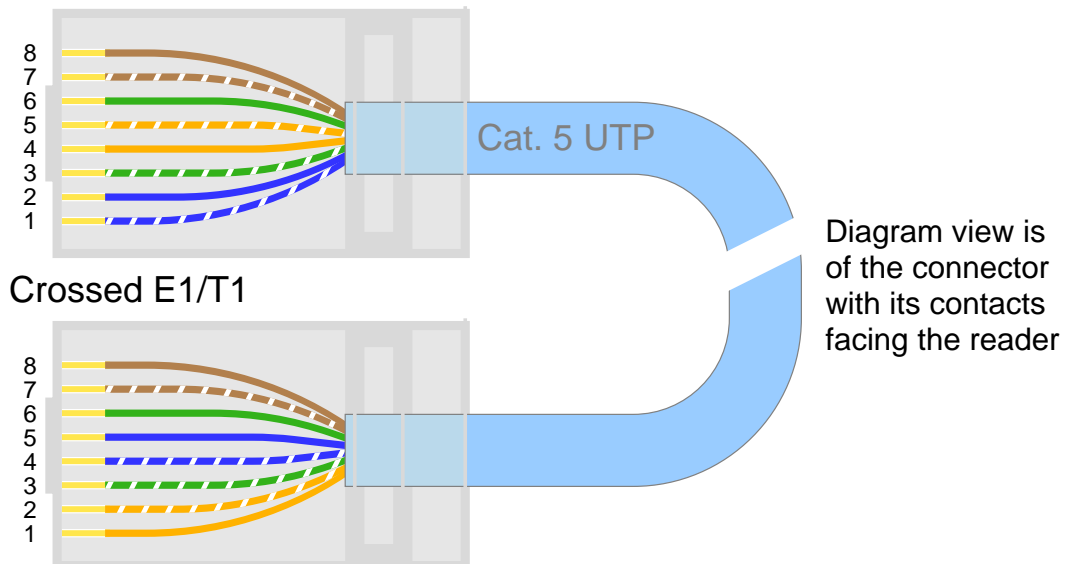


5.4 How do I make a crossed E1/T1 cable?

A crossed E1/T1 cable can be used to loop back an E1/T1 port, i.e. make a GTH transmit to itself. The connector is RJ45.

E1 lines are specified for 120 ohm cables, so a CAT-5 or CAT-5e cable (100 ohm) is out-of-spec and thus not recommended for anything beyond lab use.

End A		Cable	End B	
Pin	Name	Wire Colour	Pin	Name
1	RX	orange-white	4	TX
2	RX	orange	5	TX
3	Ground	green-white	3	Ground
4	TX	blue	1	RX
5	TX	blue-white	2	RX
6	Ground	green	6	Ground
7	Ground	brown-white	7	Ground
8	Ground	brown	8	Ground



6 Historical Hardware questions

6.1 What are the differences between hardware revisions and models?

Model	GTH 2.0	GTH 1.x
Shipping years	2006 onwards	2001–2006
Ethernet ports	2x10/100	1x10/100, 1x10
E1/T1 ports (duplex)	8	4
E1/T1 ports (monitoring)	16	4
MTP-2 monitoring	64 timeslots	32
MTP-2 Annex A monitoring	4 channels (124 timeslots)	2 channels
LAPD monitoring	64 timeslots	32
Frame relay monitoring	64 channels, up to 256 timeslots (16Mbit/s)	32 channels, 4Mbit/s
ATM monitoring	6 channels, 12Mbit/s	2 channels, 4Mbit/s

GTH 1.0, GTH 1.1 and GTH 2.0 modules use the same software API.

In addition to the above modules, Corelatus also shipped an expansion module, the IEB 1.0, from 2003–2006. The IEB module provided 12 E1/T1 ports, but relied on an adjacent host GTH 1.1 module for processing power and ethernet.

6.2 What do the rotary switches on IEB modules do?

IEB modules each have one rotary switch. The rotary switches on the modules are correctly set at the factory. If IEB systems are expanded in the field, it may be necessary to alter the rotary switch.

The rotary switch on *IEB* modules identifies the module to the *GTH* it is connected to.

IEB position in chassis	Rotary Switch Setting
IEB module adjacent to GTH module	arrow points at '1'
IEB module furthest from the GTH	arrow points at '2'

The rotary switch present on some manufacturing runs of *GTH* modules does not affect the *GTH* in any way. It is factory preset to zero for compatibility with future software releases.

7 Logistics Questions

7.1 How do I recycle end-of-life systems and packaging?

Packaging material consists of cardboard and polyethylene plastic, both of which are suitable for recycling. Packaging may be returned to Corelatus, in which case items in good condition will be re-used.

Corelatus hardware at end-of-life may also be returned to Corelatus Stockholm for disassembly and recycling of materials.

7.2 What are the packaging dimensions?

A fully-equipped 19" chassis, ready for mounting in a rack:

482mm x 342mm x 42 mm (WxDxH)
6kg

A ready-to-ship Corelatus plywood crate packed with two fully-equipped chassis:

610mm x 450mm x 204mm (WxDxH)
18kg

7.3 Which customs classification number do you use?

We use the customs classification number 85173000 (Telephonic or telegraphic switching apparatus) and country of origin code 30 (Sweden).