Acquisition Modules and Platinum Firmware

Technical Manual

Document No. MAN-EAM-0003

Designed and manufactured by
Güralp Systems Limited
3 Midas House, Calleva Park
Aldermaston RG7 8EA
England

Issue E February 2014

Table of Contents

1 Preliminary Notes

1.1 Proprietary Notice

The information in this document is proprietary to Güralp Systems Limited and may be copied or distributed for educational and academic purposes but may not be used commercially without permission.

Whilst every effort is made to ensure the accuracy, completeness and usefulness of the information in the document, neither Güralp Systems Limited nor any employee assumes responsibility or is liable for for any incidental or consequential damages resulting from the use of this document.

1.2 Cautions and Notes

Cautions and notes are displayed and defined as follows:

Caution: A yellow triangle indicates a chance of damage to or failure of the equipment if the caution is not heeded.

Note: A blue circle indicates indicates a procedural or advisory note.

1.3 Manuals and Software

All manuals and software referred to in this document are available from the Güralp Systems website: www.guralp.com unless otherwise stated.

1.4 Conventions

Throughout this manual, examples are given of command-line interactions. In these examples, a fixed-width typeface will be used:

```
Example of the fixed-width typeface used.
```

Commands that you are required to type will be shown in bold:

```
Example of the fixed-width, bold typeface.
```

Where data that you type may vary depending on your individual configuration, such as parameters to commands, these data are additionally shown in italics:

```
Example of the fixed-width, bold, italic typeface.
```

Putting these together into a single example:

```
System prompt: user input with variable parameters
```

1.5 A note on terminology

Scientists and engineers from different disciplines often use different terminology to describe similar concepts. The following terminology is used consistently throughout this document.

1.5.1 Sensor

A “sensor” is an accelerometer, velocimeter or other transducer (e.g. geophone or hydrophone) with an analogue output - i.e. where a continuously varying voltage is used to represent the magnitude of the quantity being measured.

A sensor cannot generally be used as a standalone component: it forms a sub-assembly of an instrument (see below).

1.5.2 Instrument

An “instrument” is the assembly of sensors, control electronics, casing and connectors. An example of an instrument is the CMG-3T seismometer shown on the right.

A digital instrument is one that combines one or more sensors with a digitiser; their part numbers include a 'D'. A 3TD, for example, is a 3T instrument combined with a DM24 digitiser.

Within this document, the term “digital sensor” is used in two contexts: when discussing digital inputs (which may be connected to either digitisers or digital sensors) and when describing configuration items which apply to both the digitiser modules embedded within digital sensors and to stand-alone digitisers.

1.5.3 Digitiser

A “digitiser” is an electronic device designed to accept analogue inputs from one or more sensors and, using sampling techniques, convert these analogue signals into streams of numerical data, which are then stored or transmitted digitally.

An example of a digitiser is the CMG-DM24 shown in the image above.

2 Equipment Overview

2.1 Introduction

The range of Güralp acquisition modules include the:

- Embedded Acquisition Module (CMG-EAM);
- Data Communications Modules (CMG-DCM), now superseded;
- Network Appliance Module (CMG-NAM);
- Data Acquisition Systems; and
- Integrated Instruments.

All of these units are Linux-based devices but, in general, no Linux knowledge is required in order to make full use of them. The use of Linux provides a high degree of flexibility: additional functionality can often be added on request – contact Güralp Systems for further information.

2.2 Platinum Firmware

All acquisition modules use Güralp Systems' Platinum firmware for configuration and control of the following:

- Data acquisition
- Data processing
- Data recording
- Data forwarding via serial port or over IP networks using a variety of protocols such as: SEEDlink, CD1.1, WIN, QSCD (Quick Seismic Characteristic Data) and GSMS (Güralp Seismic Monitoring System)
- Network communication over Ethernet, Modem, Wireless and Bluetooth, as required.

The firmware is accessed through a web interface or command-line interface, as detailed in section 3 on page 20.

2.2.1 Important information about build 10,000 and above

Different versions of Platinum firmware are identified by their build number. This manual covers builds numbers greater than 10,000. Build 10,000 and all subsequent builds differ significantly from previous versions and the important notes in the firmware upgrade chapter (section 5.1 on page 51) should be read before upgrading from earlier versions. Users choosing to remain at an earlier build should continue to refer to MAN-EAM-0001, which describes the operation of Platinum build 3801.

2.3 Platinum systems

2.3.1 Embedded Acquisition Module

The Embedded Acquisition Module (EAM) range consists of data recording, communications and control modules available in various cases and form factors suitable for deployment in the field. It is compatible with all Güralp digitisers and instruments.

Multiple digitisers and instruments can be attached and controlled by a single module, with data being recorded to a removable hard disk (on peli-cased modules) or internal or external flash, either as a standalone recorder or as a backup for data communications.

The EAM has 100BASE-TX and 10BASE-T Ethernet, up to 8 serial ports for connecting to external devices and a USB port for use with external storage. Cylindrical versions have additional ports, including a USB port that can be connected directly to a PC for access to the internal storage.

The EAM can be supplied as standalone, borehole- and sensor-integrated variants. Other options include tamper-line monitoring, external power control and an authentication module.

2.3.2 Data Communications Module

The CMG-DCM is the precursor of the CMG-EAM. It is a versatile Linux-based module for storing and transmitting digitised data. CMG-DCMs were originally supplied with different firmware, which is no longer supported. Platinum firmware, as used on CMG-EAMS, has been ported to the CMG-DCM platform and all users are recommended to upgrade their CMG-DCMs to run Platinum firmware. An article on our web site, http://www.guralp.com/upgrading-cmg-dcm-units-with-legacy-firmware-to-platinum-firmware/, has full details of the upgrade procedure.

2.3.3 Network Appliance Module

The CMG-NAM is a rack-mountable device intended for use as a data concentrator in seismic networks. It provides more interfaces, processing power and storage than the CMG-EAM. The CMG-NAM is intended for use in a data centre and consumes more power than the CMG-EAM, which was designed specifically to be a low power device.

2.3.4 Data Acquisition Systems

The Data Acquisition Systems are range of products combining analogue-to-digital converters (digitisers) with a CMG-EAM in a single package.

For example, the DM24S24EAM combines four DM24 digitisers and an EAM for connection to up to eight triaxial or twenty-four uniaxial instruments.

They are available in various different package options, including Peli-case and steel or aluminium tubes.

2.3.5 Integrated instruments

Integrated instruments combine seismic sensors and a data acquisition system in a single package and are ideally suited for environments where rapid deployment is required.

Most integrated instruments have their own manuals, which are used in combination with that for the digitiser and this publication.

2.4 Typical Acquisition Modules

Stand-alone CMG-EAM

Stand-alone CMG-EAM

Cylindrical DAS

24 Channel DAS

CMG-NAM

Integrated instruments (seismometer or accelerometer, digitiser and EAM)

2.5 Ports

This section lists the ports (external connectors) found on Platinum systems.

Not all ports listed in this section are available on all devices. A typical NAM, for example, only has power and Ethernet ports while the CMG-5TCDE integrated instruments have Data, GPS, Ethernet, GPIO and USB.

Note: Refer to section 16.5 on page 273 for information on connector pin-outs.

2.5.1 Ports A, B, C...

The 10-pin data input ports accept serial data from digitisers for processing by an acquisition module. They can also be used for other functions, as listed in the description of the Power/Data port, below.

2.5.2 SENSOR ports

CMG-DAS units have one or more 26-pin connectors for attaching analogue instruments. They provide power and control signals to the instruments and accept analogue data from the sensors.

The number of input connectors depends on the model of the digitiser.

2.5.3 Ethernet

The 6-pin Ethernet port is a 10BASE-T /100BASE-TX Ethernet connection, referred to as `eth0` within the user interface. The supplied cable supports connection to a hub, switch or router. If direct connection to a PC or laptop is required, an optional cross-over cable can be ordered.

2.5.4 GPIO

The 12-pin GPIO (***G***eneral ***P***urpose ***I***nput/***O***utput) port fulfils three functions:

- It provides a serial console to the EAM, which can be used for monitoring, configuration and control. This is permanently configured to run at 38,400 baud;
- It provides USB access from a PC or laptop to the internal FLASH storage for data collection (use of this feature is described in section 11.3.3 on page 161); and
- It provides a number of tri-state lines which can be used to control or monitor external equipment. One application is as tamper detection lines, which can be connected to external switches and monitors as part of a secure installation.

2.5.5 USB

The 6-pin USB port allows connection of an external USB storage device for data collection. It is also possible to perform firmware upgrades using this port in situations where internet access is not available – see section 5.3.3 on page 59 for details.

2.5.6 GPS

The 10-pin GPS port allows connection of a GSL GPS receiver for use as a synchronisation source for time-stamping seismic data.

2.5.7 Power/Data

The 10-pin Power/Data Out port is a power input and also a general-purpose serial port which can be used for GCF output (suitable for serial connection to Scream), PPP network connections, inbound GCF (from a digital instrument, for example), NMEA functions, TCP serial conversion, a modem answering service or as a recorder to store and forward data from any instrument with a serial output.

2.5.8 POWER

The 4-pin Power port, where fitted, is an alternative power input. It can be used as a permanent power input in situations where the Power/Data port is only used occassionally.

2.6 Typical Applications

2.6.1 Autonomous remote data-logger

In this application, depicted below, a CMG-EAM is used to collect data from a digital instrument (or analogue instrument and digitiser) and store it on its

hard drive. The low power consumption and high storage capacity of the CMG-EAM makes it ideal for this purpose.

Where appropriate, the battery supply can be augmented with a solar panel. The CMG-EAM is capable of interfacing with and monitoring many types of solar charge controller.

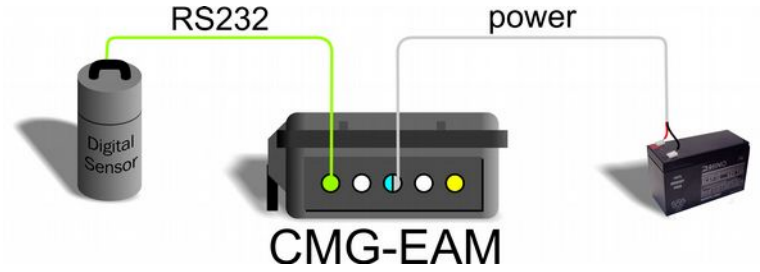

If it is desired to contact an acquisition module for monitoring or urgent data download purposes, the unit can be fitted with a GPRS or satellite modem, allowing remote connectivity.

2.6.2 Protocol Converter

An acquisition module can be deployed as a protocol converter: the wide variety of output formats and connectivity options make it ideal for this application. In the illustration below, a digital instrument's GCF output is retransmitted as SEEDLink data over Ethernet.

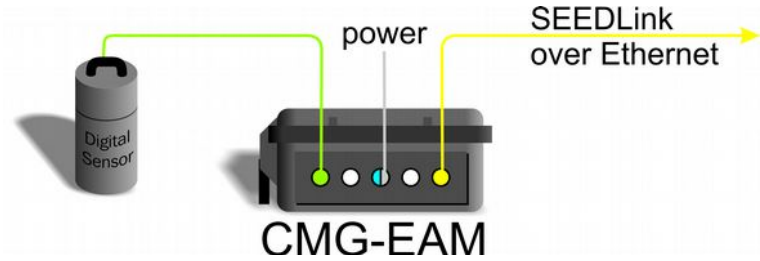

2.6.3 Array Concentrator

The acquisition module can combine the data from many instruments in an array and retransmit them over a single link (serial or network).

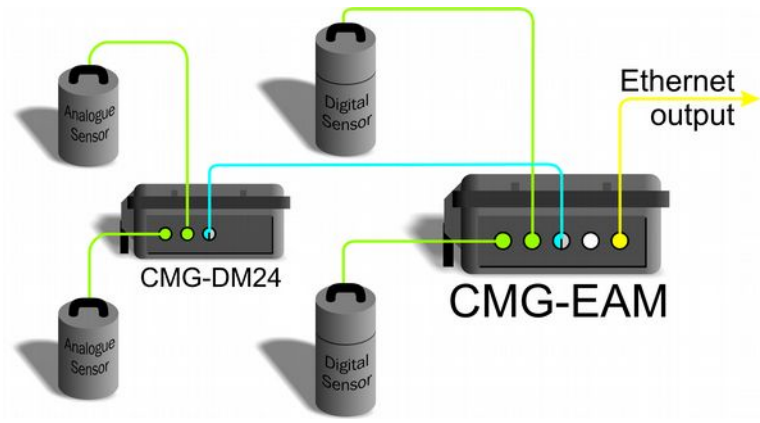

If the output link is over a network, all external serial ports of a standard CMG-EAM (including the "DATA OUT" port) are available for connection to digitisers or digital instruments, allowing up to eighteen channels to be aggregated. An arbitrary number of CMG-EAMs may be chained together, allowing for even more extensive arrangements.

2.6.4 PPP Networking

A serial output from the CMG-EAM can be used for point-to-point protocol (PPP) networking. This protocol allows full internet access to the device over a serial link. Operators can access the web page of the acquisition module for configuration and monitoring. If other Ethernet devices are present at the deployment site, the CMG-EAM can function as a router, passing their traffic over the PPP link. The output from other serial devices can also be passed over the PPP link by use of the built-in serial-to-TCP converter (see section 7.8 on page 97 for details).

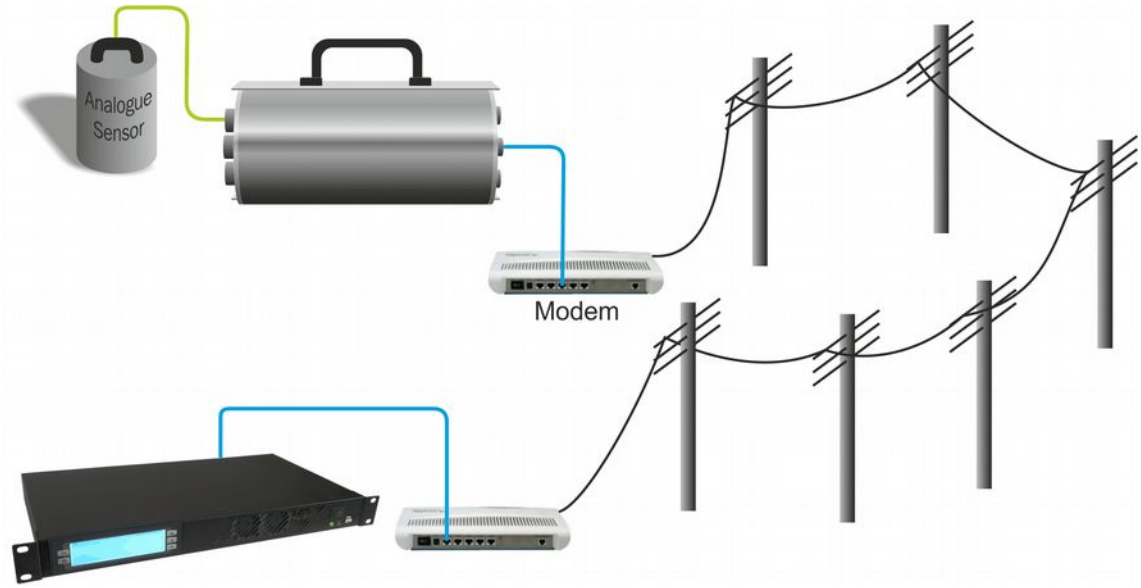

2.6.5 Resilient Networking

Platinum firmware includes a number of ways to implement network resilience. For example, the GSTM protocol (for communication between Platinum units) allows data to be routed over a low-cost but unreliable DSL network with automatic switch-over to a higher-cost satellite link only when the DSL network is unavailable. The failed link is regularly re-tried and, when communication is re-established, the data are re-routed back to the lower cost link.

The CMG-NAM acquires data from Scream servers (e.g. CMG-6TD or CMG-3ESPCDE) and the data are stored locally on an optional RAID disk array with up to 2 TB capacity.

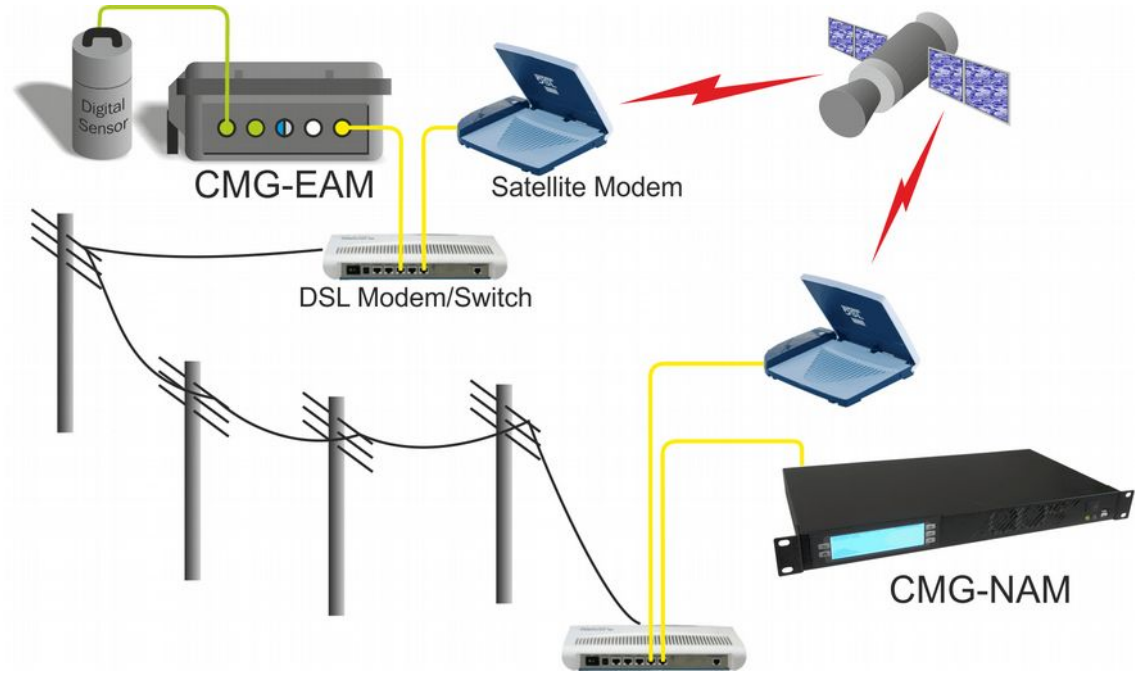

The CMG-NAM can also act as a data server to remote clients supporting GCF (SCREAM server), EARTHWORM (via scream2ew), ANTELOPE (via Guralp2orb), CD1.0/CD1.1 and SEEDLINK.

It is also possible to use data filtering (by channel name and/or sample rate) in such a way that, should it become necessary to use the higher-cost link, only high priority data (e.g. samples resulting from an activated trigger) are sent across this link while lower priority data are enqueued until the low-cost link becomes available again.

2.6.6 CD1.1 Networking

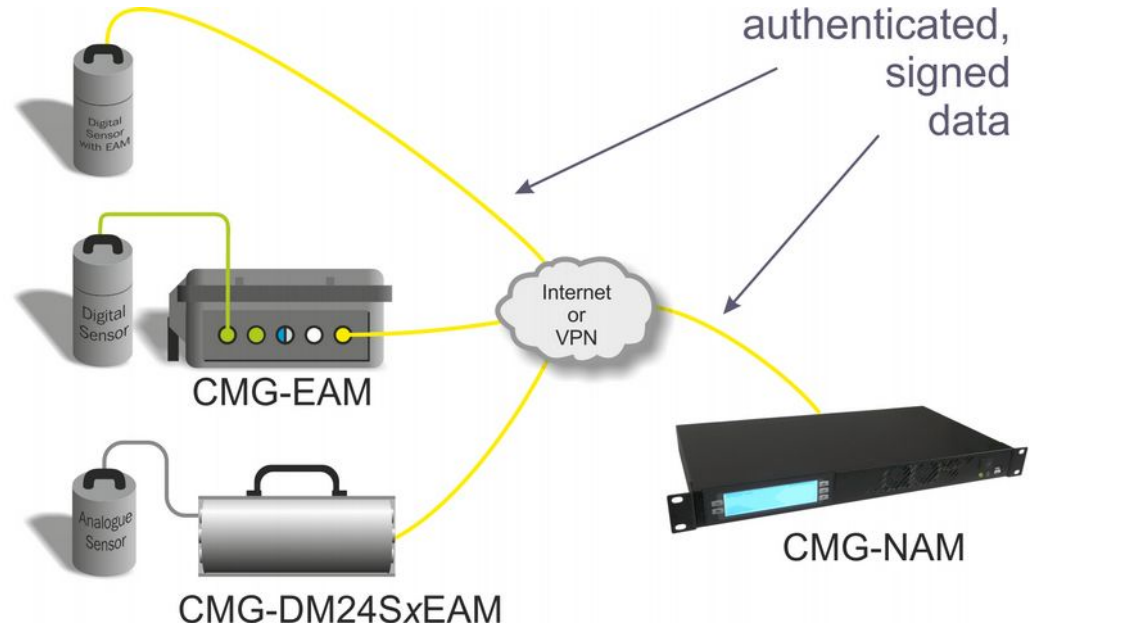

Platinum firmware has support for CD1.1 frame generation and forwarding with strong authentication provided by an optional embedded Spyrus hardware encryption device, allowing CMG-EAMs and CMG-NAMs to form the basis of a secure CD1.1 network.

Data transmitted in CD1.1 format use strong authentication and digital signatures to ensure that the received data are exactly those transmitted: i.e. that they have not been tampered with during transmission.

Platinum firmware contains many facilities to support CD1.1. These are documented in a separate manual, MAN-EAM-1100, which is available on request from support@guralp.com.

3 Initial set-up

3.1 Introduction

All acquisition modules except the CMG-NAM can be configured and monitored either over an Ethernet network or via a serial (RS232) line and are provided with at least one network and one serial port. Because the CMG-NAM is designed for use in a data centre, it only has a network port.

The actual number of network or serial ports is dependant on the exact model of the acquisition module. CMG-EAMs in peli-cases are supplied with a dedicated console connector located under the lid.

If using a network, the acquisition module can be accessed using a web browser or, in character mode, using ssh. Instructions for connecting to the network port are given in section 3.2 on page 20.

If you prefer to use serial communications, the module can be accessed using a terminal emulator. Instructions for connecting via the serial port are given in section 3.3 on page 30.

Note: We recommend use of the web interface over a network for general configuration and operation.

3.2 Connecting to the network port

To use the network port, you must first set up a network address. Some networks need manual configuration (normally referred to as “static” addressing); others use the Dynamic Host Configuration Protocol (DHCP) to allow a DHCP server to automatically assign network addresses. If no DHCP server is present, many systems will fall back to a randomly-generated “link-local” address (known by Microsoft as Automatic Private IP Addressing, or APIPA). Before you can access an acquisition modules over a network, you must set (for static addresses) or discover (if you use DHCP) its IP address.

Note: If you are setting up a unit in the laboratory for subsequent deployment in the field, you can set up the final network address using the web interface and over-ride it with a temporary, static network address using the command line. The web-configured address will take effect when the unit is next rebooted.

3.2.1 DHCP-assigned addresses

Acquisition modules are supplied configured for DHCP. If your network uses DHCP to assign addresses, simply connect the acquisition module to the network and wait a few minutes for the process to complete. Your network administrator should then be able to tell you the address that has been

assigned to the acquisition module. If your module has an LCD status display (CMG-NAMs and some integrated instruments), the assigned address will be displayed in it. The LCD display shows lots of information so you may have to wait a short time until the IP address scrolls into view.

When using DHCP, it is recommended that the DHCP administrator allocates a fixed IP address to the acquisition module's MAC address in order to avoid unexpected address changes. The MAC address is displayed by the `ip` command – in the example in section 3.2.1.1 on page 21, it is `00:50:c2:40:54:75`.

If you cannot learn the IP address in this way, there are three methods available to discover which address has been allocated.

3.2.1.1 Address discovery – serial connection

You can connect via a serial port (as described in section 3.3 on page 30) and issue the `ip` command:

```
eam999 ~ # ip addr show eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc
pfifo_fast qlen 1000
   link/ether 00:50:c2:40:54:75 brd ff:ff:ff:ff:ff:ff
   inet 192.168.0.101/24 brd 192.168.1.255 scope global eth0

   inet6 fe80:250:c2ff:fe40:5475/64 scope link
   valid_lft forever preferred_lft forever
eam999 ~ #
```

The key things to look for here are the adapter status and the IP address. The first line of the output should contain the word `UP`, confirming that the adaptor has been enabled. The IP address that has been assigned is shown on the line beginning `inet` - in this case, it is `192.168.0.101` (with a netmask of 24 bits indicated by `/24`).

Note: With an IP version 6 network, the IP address will be on a line beginning `inet6`. In practice, most networks today are still IPv4, as in the above example.

3.2.1.2 Address discovery – Scream's “Detect servers” tool

Start with the EAM turned off. Then, from Scream's “Network control” window, select the “My Client” tab. Right-click in the server list-box and select “Detect servers...” - Scream will then start monitoring DHCP traffic on the local network. Power up the EAM and allow it a minute to boot. When Scream notices a DHCP negotiation with an appropriate MAC address, it will display the allocated IP address in the window. You can add the EAM to Scream's list of servers by clicking the appropriate button.

3.2.1.3 Address discovery – GCFPing

Scream's "GCFPing" feature sends a specially formatted broadcast packet to all hosts on the local network. Any GCF servers that see this packet should respond with a GCF acknowledgement packet (GCFACKN). Scream displays the IP addresses associated with all acknowledgement packets that it receives.

To use this feature, add a new UDP server with an IP address of 255.255.255.255. Right-click the server and select GCFPING from the context menu. A window will appear as shown below if the ping packet is acknowledged. In the control pane (at the bottom right of the main window), a GCFACKN line will be printed for every address that responded to the GCFPING. In the example below, you can see responses from 192.168.42.135, 192.168.42.98 and 192.168.42.139. These are the addresses of all GCF servers listening on port 1567 on the local network.

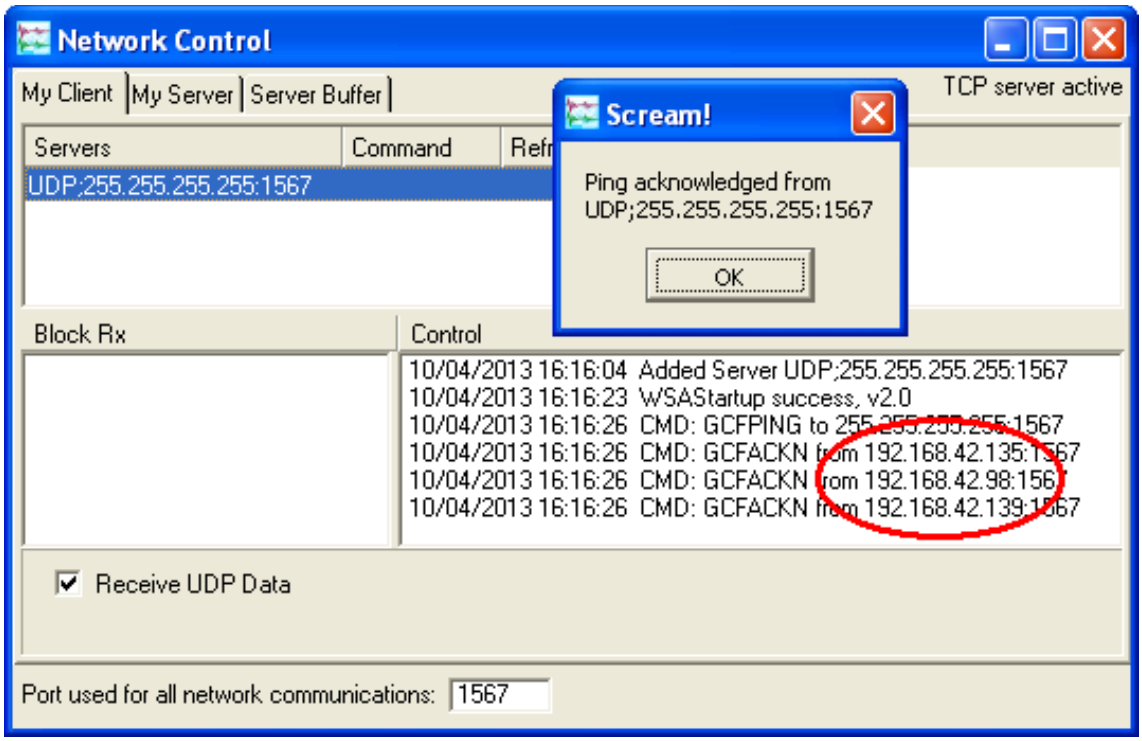

3.2.2 Link-local addresses

Many systems, when configured for DHCP, will generate a random address if no DHCP server is present. This is known as a "link-local" or APIPA address. For IPv4 networks, it will be in the range 196.254.0.0 to 196.254.255.254 (i.e. on the 196.254.0.0/16 network). For IPv6 networks, it will be in the fe80::/10 network. The random, host-specific part of the address is derived from the (unique) MAC address, so there are unlikely to be conflicts between addresses of systems in networks with small or medium numbers of hosts.

This is useful when, for example, visiting a remote EAM: A laptop can be plugged directly into the network port of the EAM (using a cross-over cable, if

necessary) and, provided both systems are set to use DHCP, both will assign themselves addresses on the same network. If the laptop is running Scream, you can add a server using the link-local network's broadcast address, 196.254.255.255, and start communicating immediately. If the address of the acquisition device is required (for, say, web access), this can be read from Scream's control window, or in the acknowledgement window resulting from a GCFPING.

The acquisition device will search for a DHCP server every minute and, should one become available, it will ask it for a new address.

3.2.3 Assigning a static IP address

If you wish to configure a static IP address, you must first connect to the command line via a serial port. This process is described in section 3.3 on page 30.

3.2.3.1 Assigning a static IP address using net-setup

Once logged in, issue the following command:

```
eam999 ~ # net-setup
```

Note: This command relies on the EAM understanding what type of terminal emulator you are using. If the display is corrupted and not usable, set the TERM variable (see section 4.3.1 on page 35), or simply power-cycle the EAM and use the ip command (see section 3.2.3.2 on page 25) instead.

The following screen is displayed:

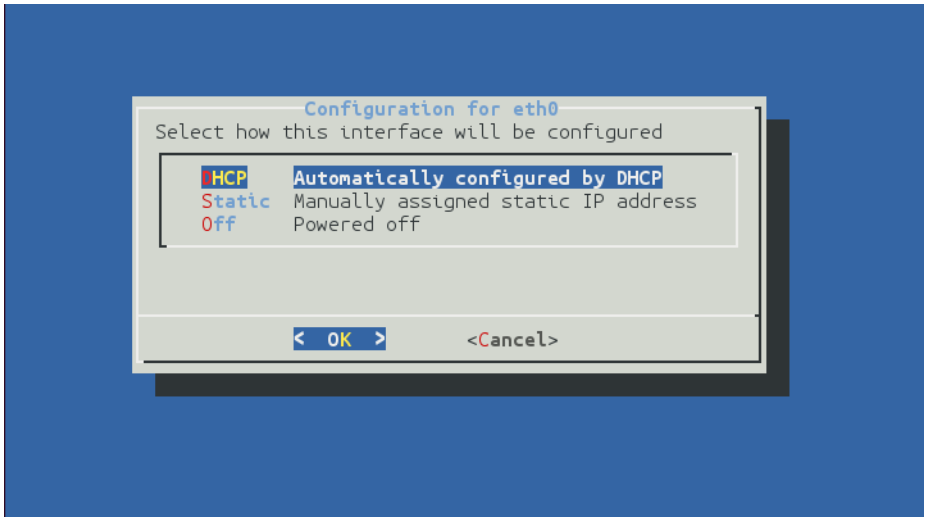

Using the ↑ and ↓ keys (or, on some systems, the mouse), select whether you wish to use DHCP or static addressing. You can, alternatively, key D to select DHCP or S to select static addressing. Use the ENTER key to confirm

your choice. If you do not wish to make a change, use the key to select "Cancel" and then key ENTER to close the net-setup tool and return you to the command prompt.

If you select static addressing, the following screen is displayed:

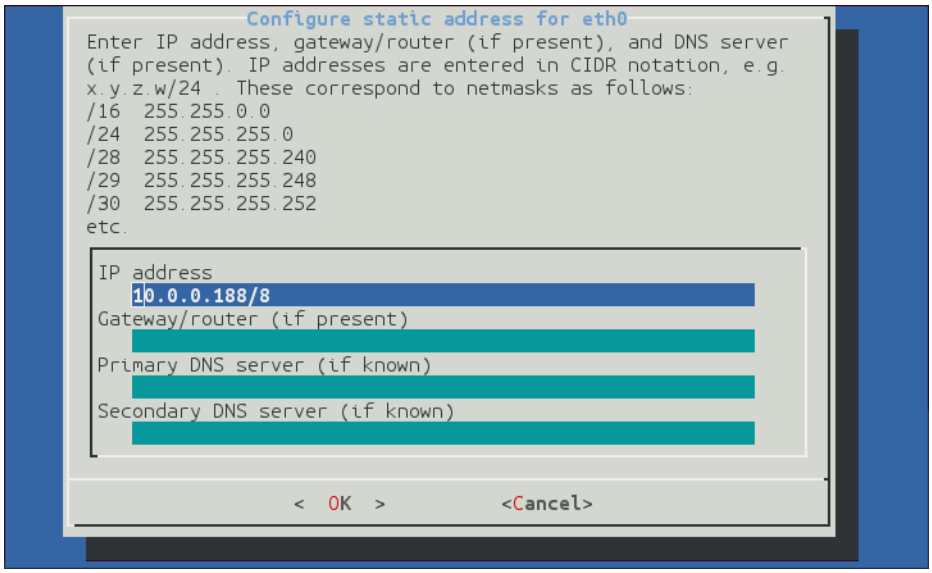

The IP address field must be populated with a valid IP address in CIDR notation. If you know the netmask but not the corresponding CIDR notation, use the information on screen as a guide or search the web for an on-line converter.

If your network has a router which acts as a gateway to the Internet or to other networks, use the key (or, on some systems, the mouse) to move to the "Gateway/router" field and enter the address of the gateway in standard, dotted-quad notation (i.e. `10.0.0.1`).

If your network has a DNS (domain-name service) server, sometimes called a name-server, use the key (or, on some systems, the mouse) to move to the "Primary DNS server" field and enter the address of the name-server in standard, dotted-quad notation (i.e. `10.0.0.5`).

Note: If your network is connected to the Internet but you do not know the name of your DNS server, ask your Internet Service Provider for the correct address to use or enter `8.8.8.8`, which is a free, public DNS server operated by Google Inc.

Use the ENTER key to confirm your choice and reconfigure the network. If you do not wish to make a change, use the key to select "Cancel" and then key ENTER to close the net-setup tool and return you to the command prompt.

3.2.3.2 Assigning a static IP address using the `ip` command

The `ip` command is an alternative to `net-setup`. You may wish to use it if

- you want to configure a temporary address without updating the configuration files
- you cannot use the `net-setup` utility for any reason
- you are very familiar with the linux command line

Log in as normal and then issue the following command:

```
eam999 ~ # ip addr add 192.168.0.1/24 dev eth0
```

replacing the example IP address (`192.168.0.1/24` in the example above) with the required value. It must be specified in CIDR format, where the actual address is followed by the number of bits of the network mask. The above example uses 192.168.0.1 with a netmask of 255.255.255.0 (24 bits of network address). A PC connected to this network could communicate with the acquisition module if it was configured to use an IP address of (for example) 192.168.0.2 with a matching netmask of 255.255.255.0.

Note: IP addresses assigned using this method will be lost if the unit is rebooted. To permanently assign an IP address, use `net-setup` (see section 3.2.3.1 on page 23), the web interface (section 7.1 on page 70) or gconfig (section 7.1 on page 70).

If you wish to connect to the acquisition module from a PC, they must either both be on the same physical network and have the same network address (usually the first three numbers of the IP address) or be able to connect to each other via routers.

In the latter case, you will need to tell the acquisition module the address of its default router (also known as the gateway). Issue the command:

```
eam999 ~ # ip route add default via 192.168.0.254
```

substituting the address of your network's default router in place of the example address (`192.168.0.254`) shown.

If you wish to be able to access your acquisition module across the Internet, perform firmware upgrades or access GSL remote support, you will also need to configure a default router as described in the preceding paragraph.

Note: Both the static IP address and any route configured in this way are temporary and will persist only until the acquisition module is rebooted or powered off. Refer to section 7.1 on page 70 for information about configuring permanent static IP addresses and routes.

If you wish your acquisition module to initiate connections with remote systems across the Internet, or to be able to access firmware upgrades, you need to configure a DNS server (also known as a name-server). If you do not know the address of your DNS server, your Internet Service Provider (ISP) will be able to tell you. You can also use 8.8.8.8, which is a free, public DNS server operated by Google Inc. Enter the command

```
eam999 ~ # echo "nameserver 8.8.8.8" >> /etc/resolv.conf
```

substituting the address of the required DNS server in place of the example address (`8.8.8.8`) shown in the example above.

3.2.4 Connecting to the web interface

The Platinum firmware on all acquisition modules provides a web interface for configuration and control of the module and connected equipment. While there are other methods of connecting to the modules, the web interface is recommended.

Once the IP address of the acquisition module has been set or determined, enter the it into the address bar of a web browser to connect to the module's web interface . The examples below are for an EAM address entered into Firefox and Internet Explorer:

On versions of Platinum from release 10,000, the web interface will initially show a status display and a brief menu. There is an option to log in on that menu. Click on the link and enter the default user-name of `root` and password of `rootme`.

If you are connecting to the acquisition module over a network that you consider insecure (such as the internet), it is recommended that you use the HTTPS (secure HTTP) protocol, which uses TLS to encrypt the link. Simply change the **`http://`** prefix to **`https://`** in the browser's address bar. Most browsers will complain that the certificate cannot be verified: This is not a problem: simply press the “accept” button to proceed. The link will then be encrypted and nobody will be able to “sniff the wire” in an attempt to discover passwords and other data.

Once connected and logged in, you will be presented with the main summary screen and a much larger menu. The summary screen contains general information about the status and health of the connected modules and equipment.

The exact contents and layout of this screen will vary depending on the configuration of both the acquisition module and of any attached devices.

See section 4on page 33 for information on using the web interface.

3.2.4.1 Connection trouble-shooting

If the browser fails to connect, the most likely explanation is that the machine running the browser does not have working network communications to and from the acquisition module. This can be verified by "**ping**ing" the IP address of the browser from the command line of the acquisition module:

```
eam999 ~ # ping -c3 192.168.0.2
PING 192.168.0.2 (192.168.0.2): 56 data bytes
64 bytes from 192.168.0.2: seq=0 ttl=63 time=2.284 ms
64 bytes from 192.168.0.2: seq=1 ttl=63 time=1.129 ms
64 bytes from 192.168.0.2: seq=2 ttl=63 time=1.944 ms
--- 192.168.42.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 1.129/1.785/2.284 ms
eam999 ~ #
```

To resolve this class of problem, ensure that the cables are functioning (which can be verified by checking the diagnostic lights on most switches/hubs) and double-check that the PC and acquisition module are on the same subnet (which means the CIDR suffices must match and the first sections of the IP addresses – as defined by the CIDR suffices - must also match). The website http://en.wikipedia.org/wiki/IP_address has some useful information for those for whom sub-networking is unfamiliar.

3.2.5 Connecting to the command line using SSH

SSH (**S**ecure **SH**ell) is the most flexible way to control an acquisition module, but it is less friendly than using the web interface. It is possible to configure more advanced operations using SSH but the majority of control and configuration tasks can be achieved most easily through the web interface.

SSH is shipped as standard with most Linux distributions and is available for Windows as part of the free terminal emulation package PuTTY, which is available from http://www.chiark.greenend.org.uk/~sgtatham/putty/

3.2.5.1 SSH connections using the ssh program

To use SSH, you must know or discover the IP address of the unit, as described in the previous section. Once you have the IP address, issue the **SSH** command on the PC you are using:

```
mypc$ ssh root@192.168.0.1
```

Replace `192.168.0.1` with the IP address of the acquisition module.

The first time you use SSH to connect to a host, you will be asked to verify the "host key". This is normal but, if you are ever asked this again, it means that either the host key of the acquisition module has changed – perhaps because of a firmware upgrade – or there is a network address conflict or, worse, a security problem on your network.

```
user@mypc:~$ ssh root@192.168.0.1
The authenticity of host '192.168.0.1 (192.168.0.1)' can't
be established.
RSA key fingerprint is
62:a6:70:29:d4:1a:db:5a:75:6e:96:13:54:f5:a9:d9.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '192.168.0.1' (RSA) to the list
of known hosts.
root@192.168.0.1's password:
eam999 ~ #
```

You will be prompted for a password; the default password is **rootme**. Note that no characters will be echoed to the screen as you type the password.

Once connected, you will be presented with a shell prompt which is ready to accept commands.

When you are finished with your SSH session and want to disconnect, enter **exit** at the command line, or type Ctrl + D. There are a number of escape sequences for controlling the session, all of which begin with a tilde ('**~**') so, if you need to send a tilde character to the acquisition module, type two tildes consecutively. For more information, see the section on "Escape Characters" in the manual at http://man-wiki.net/index.php/1:ssh

Note: If you plan to use ssh regularly to communicate with a acquisition module, you can configure the system to bypass the password prompt when logging in from a list of pre-authorised computer/user combinations. This involves generating a unique key-pair (for the user and PC which will access the acquisition module) and then copying the public half of the key-pair to the acquisition module. This can be more secure than passwords and is fully documented at http://suso.org/docs/shell/ssh.sdf. For details about uploading your keys to the CMG-EAM, see section 7.6 on page 90.

3.2.5.2 SSH connections using PuTTY

To use PuTTY, you must first know or discover the IP address of the unit, as described in the previous sections. Once you have the IP address, start PuTTY by choosing it from the "Start" menu or double-clicking on its icon. You will be presented with the following screen:

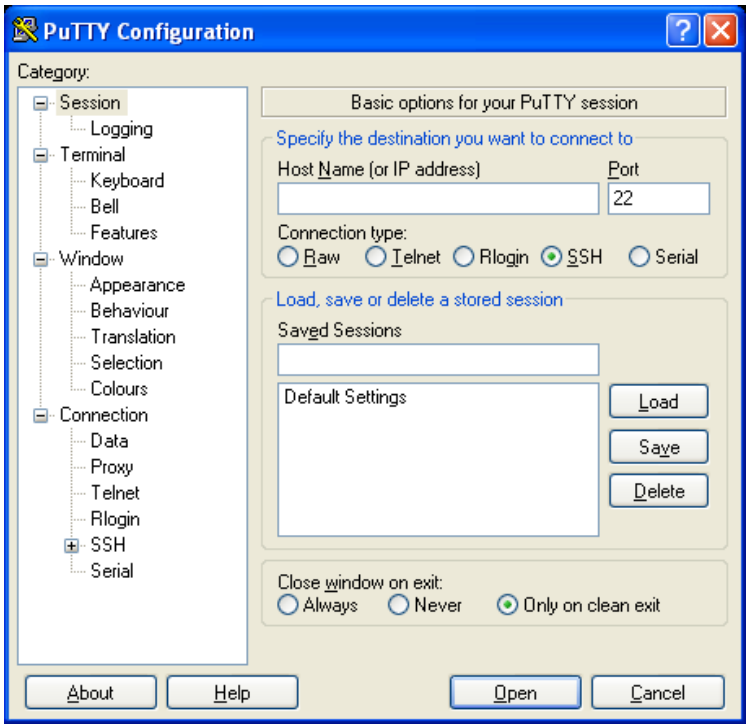

Enter the IP address of the acquisition module into the **Host Name (or IP address)** field, check that SSH is selected as the **Connection type** and then click the Open button.

The first time you use SSH to connect to a host, you will be asked to verify the "host key":

This is normal (simply click Yes to dismiss the dialogue) but, if you are ever asked this again, it means that either the host key of the acquisition module has changed – perhaps because of a firmware upgrade – or there is a network address conflict or, worse, a security problem on your network.

You will now be prompted for a login name: the default is `root`. Type this at the prompt and press the ENTER key. You will next be prompted for a

password; the default password is `rootme`. Note that no characters will be echoed to the screen as you type the password.

Once connected, you will be presented with a shell prompt which is ready to accept commands. The shell prompt contains the serial number of the acquisition module.

When you are finished with your SSH session and want to disconnect, type "exit" at the command line, or Ctrl + D.

PuTTY allows you to save multiple sessions, along with a default login identity and screen colours for each. See the PuTTY manual for more details.

Note: If you plan to use ssh regularly to communicate with a acquisition module, you can configure the system to bypass the password prompt when logging in from a list of pre-authorised computer/user combinations. This involves generating a unique key-pair (for the user and PC which will access the acquisition module) and then copying the public half of the key-pair to the acquisition module. This can be more secure than passwords and is fully documented at http://suso.org/docs/shell/ssh.sdf. For details about uploading your keys to the CMG-EAM, see section 7.6 on page 90.

3.3 Connecting to the Serial Port

A number of acquisition modules have dedicated a 9-pin 'D' console port connector located under its lid. It can be connected via a serial (RS232) cable to a PC. See section 16.5.5 on page 277 for the pin-out.

Some acquisition modules without a dedicated console port connector are supplied with 'Y' cables that connect to the DATA OUT port, your PC and a power supply.

Acquisition modules with a GPIO connector are provided with a blue serial cable, terminating in a female DE9 connector. This gives access to the module's console.

Once you have connected the serial cable, you can run either Scream or a terminal emulator to interact with the acquisition module.

3.3.1 Using Scream

Select `Terminal...` from the `File` menu. A window will open, from where you can select the correct serial port:

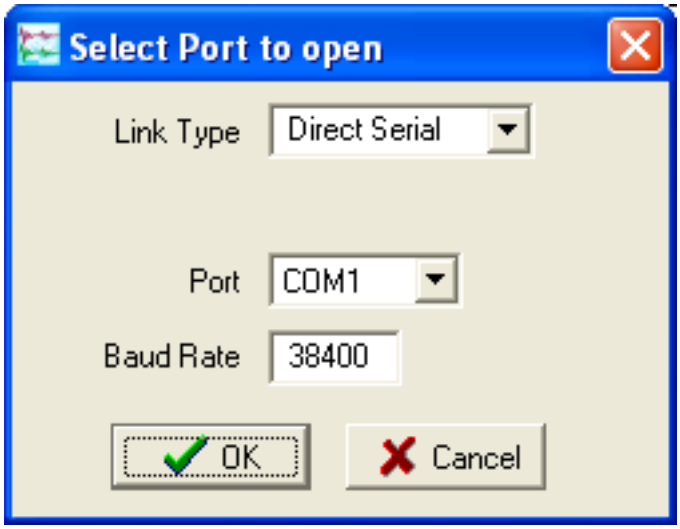

Set the Baud Rate to 38400, as shown, and ,once the acquisition module and computer are communicating properly, an emulation window will open and you will see the Login: prompt. If it does not appear immediately, press the ENTER key a few times.

Note: If a terminal session has just been closed, it can take up to ten seconds for a new session to start.

3.3.2 Using a terminal Emulator

You are free to use whatever terminal emulation software you wish. Common choices are applications such as minicom on Linux or PuTTY on Microsoft Windows. See section 16.2 on page 262 for more information.

Configure your emulation software to use the correct serial port. Set the line speed (Baud rate) to 38,400 and the communication parameters to eight data bits, no parity bits and one stop bit. This combination is commonly referred to as “8-N-1”. Turn off all hardware flow control (RTS/CTS and/or DSR/DTR should not be used). Turn off all software flow control (XON/XOFF should not be used).

Once the emulator is connected, you will see a `Login:` prompt. If it does not appear immediately, press the ENTER key a few times.

Note: If a terminal session has just been closed, it can take up to ten seconds for a new session to start.

3.3.3 Logging in

Acquisition modules are shipped with a default user-name of **root** and password of **rootme.** (Additional users with controlled capabilities can be added if required).

Enter the user name and password. Note that nothing will display on the screen when typing the password. You will then be presented with a shell prompt, which will accept commands as shown in the image below:

```
eam999 login: root
Password:
eam999 ~ #
```

The output may vary slightly due to the configuration of the unit. In particular, the acquisition module name, as displayed in the prompt (`eam999` in this example), will be different.

4 Platinum Overview

4.1 Introduction

All key acquisition module configuration and control tasks can be carried out either from the web interface or from the command line.

The web interface presents some additional options not available from the command line. Some of these are merely short-cuts into the main configuration system while others offer additional monitoring and diagnostic facilities. Configuration of attached digitisers can be carried out using the web interface but not from the command line (although access to the digitiser's command line is available).

The command-line interface also supports a number of advanced facilities which are not available via the main configuration system: these are mostly diagnostic tools which are not required for normal operation.

4.2 Using the web interface

The web interface is split into two frames:

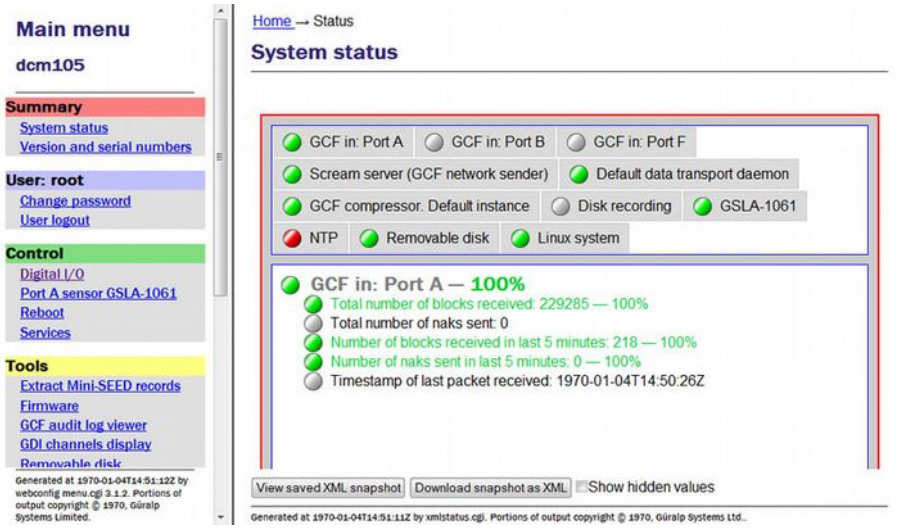

The left-hand frame contains the system ID above a menu while the main frame displays sub-menus, input forms and display screens.

4.2.1 Navigation aides

To help with navigation, the main frame displays a series of links indicating the current position in the configuration-system hierarchy:

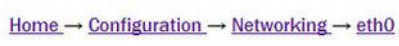

These links are commonly known as a “bread-crumb trail”. The “Home” link will return the user to the system home page shown above.

4.2.2 Display options and form submission

Many web forms display a series of buttons at the bottom of the form:

The Config home button returns you to the configuration homepage. This page can also be accessed using the **All options** link in the main menu.

The Help button displays on-line help in blue, interleaved with the form fields.

The Expert button changes the display mode to show extra form fields that do not normally need to be changed. These 'expert' settings are all described in the relevant section of the manual for that screen. When in expert mode, the button changes to Simple so that the display can be returned to the default mode.

The Submit button sends the updated settings to the acquisition modules. In many cases, the changes are accepted immediately; some, however, require a reboot of the modules or attached device.

User input to the forms is validated after submission. Where invalid parameters have been detected, an error message is displayed in red below the appropriate field:

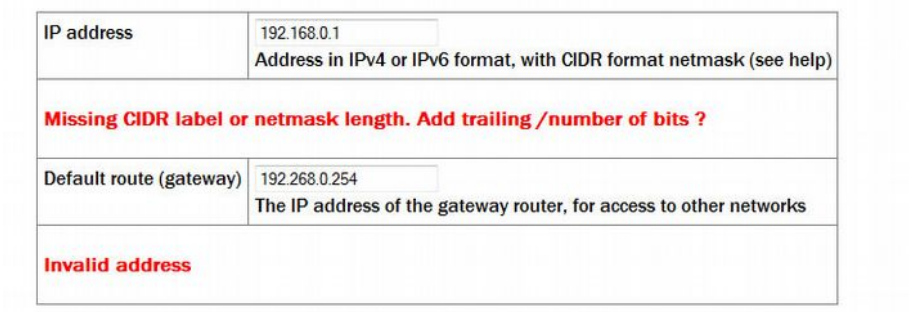

On pages where there are tabbed pages, the number of errors is displayed on the first tabbed screen:

Throughout the remainder of this manual, screen-shots of the configuration system's web interface will normally omit the left-hand pane, as in the two illustrations above.

4.2.3 Navigation instructions in the manual

In this manual, instructions on how to reach a particular screen are displayed like this:

Configuration → Interfaces → eth0 - Primary wired network interface

In this example, to reach the page indicated, scroll to the “Configuration” section of the main menu then click on the “Interfaces” link". From the list that appears in the main frame, click on the “eth0 - Primary wired network interface” link.

In some cases there are two possible paths to a particular screen. Both are shown in the manual.

4.3 Using the command-line configuration system

All of the configuration facilities available under the “All options” item (in the main menu of the web interface) are also available from a text-based GUI tool called `gconfig` (**G**üralp **Config**urator). This can be accessed either by using a serial link or, over the network, by using ssh.

Connection via a serial link is discussed in section 3.3 on page 30 and connecting over Ethernet is described in section 3.2 on page 20. The use of ssh is covered in section 3.2.5 on page 27.

4.3.1 Using graphical interfaces from the command line

Some Platinum applications use a system called “ncurses”, which allows graphical interfaces to be implemented on text-only terminals. This requires the applications to know the type of terminal from which they are being accessed. The terminal type is stored in an environment variable called **`TERM`**, which is queried with the command

```
eam999 ~ # echo $TERM
vt100
eam999 ~ #
```

(note the use of the $ sign when accessing the value of this variable) and set with the command

```
eam999 ~ #: export TERM=vt100
```

No spaces should be used around the '=' sign.

Platinum is aware of around thirty different terminal types and uses the “terminfo” system to support them (so you can add your own types, if you need). Files describing each terminal type are stored under the directory (folder) `/usr/share/terminfo` in sub-directories named after the initial letter of the terminal name.

Some settings for specific applications are given in the following table.

Emulator in use	TERM setting
ssh under Unix	No action required - the ssh protocol sets the TERM environment variable automatically.
PuTTY (in ssh mode) under Windows	No action required - the ssh protocol sets the TERM environment variable automatically. PuTTY will default to **xterm** emulation, allowing the mouse to be used within gconfig.
Minicom under Unix	Minicom emulates a vt100-style terminal and automatically maps the keystrokes and display sequences for the actual terminal you are using, so the default **TERM** setting of ***vt100*** is correct.
HyperTerminal under Windows	Using the File menu option "Settings", ensure that the terminal type is set to ***VT100***. HyperTerminal will then emulate a ***vt100*** style terminal, which will match the default **TERM=*vt100*** on the acquisition module.
Scream	Scream versions before 4.5 did not support the required screen-drawing control codes so an upgrade to version 4.5 or later is recommended.

These settings will provide the best results for the listed applications. Note that, when connecting with SSH from an **xterm** window, for example, use of the mouse for menu navigation is supported.

4.3.2 Using gconfig

When you enter the **gconfig** command, the initial screen looks like this:

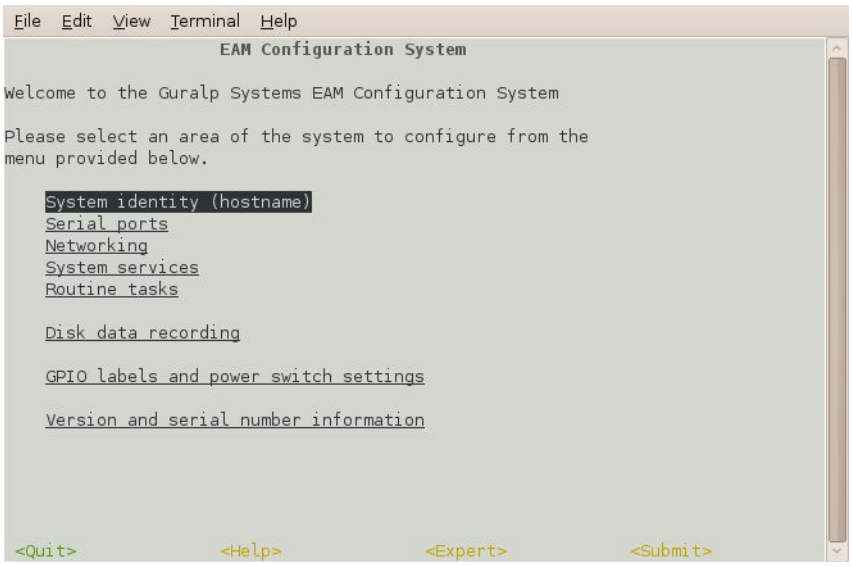

The gconfig interface can be navigated entirely using the keyboard but, if you use **xterm** (or your terminal emulator supports an "xterm" mode) you can use

your mouse to select menu items, input fields and items from drop-down menus. The scroll-wheel is not currently supported, so you need to use the keyboard to access second and subsequent pages of multi-page forms.

The gconfig interface displays menus and forms. The screen-shot above is an example of a menu. The following table shows the navigation keys that are available for use with menus:

Keystroke	Used for...
→, ↓ or Tab	Cursor to next item
← or ↑	Cursor to previous item
Enter	Select item under cursor
Page Up or F6	Move to the next page of a multi-page menu
Page Down or F5	Move to the previous page of a multi-page menu
F1	Go to the home menu or, if there already, exit `gconfig`.

The four words enclosed in chevrons at the bottom of screen act as push-buttons:

They are always present at the bottom of each gconfig screen, although they change slightly according to the context. To invoke the action associated with any of them, move the cursor to it and press Enter.

In the example, all options other than `<Quit>` are disabled and this is indicated by the colour coding. If you select an item from a menu which leads to a sub-menu, the `<Quit>` option changes to `<Home>` and invoking it will then take you back to the top-level menu.

Selecting a menu item will lead you either to another menu or to a form. Forms are composed of text entry fields, check boxes or drop down menus (all described below).

To move between fields, use the ↑, ↓ or Tab keys.

4.3.3 Text entry fields

Text-entry fields are those into which the operator can type textual parameters. Text-entry fields are identified by surrounding braces: **[** and **]**

To edit the contents of a text-entry field, move the cursor on the field and start typing. Characters typed are always inserted at the current cursor position i.e. existing characters are never over-typed. The whole field is shown with a black background, as seen below, and the cursor is identifiable as a pale block.

The ← and → keys move the cursor within the field and the ← key deletes characters to the left of the cursor.

4.3.4 Check-boxes

Check-boxes offer the operator a “yes/no” choice. Check-boxes are identified by surrounding guillemots: '{' and '}'.

To change the setting of a check-box, move the cursor on the field and use the space key to toggle between selected and not selected.

When the value is “yes”, “enabled” or otherwise selected, the field is shown with an 'X' in it:

When the value is “no”, “disabled” or otherwise de-selected, the field is shown as blank:

4.3.5 Drop-down menus

Drop-down menus are where the operator must choose one option from a list. Drop-down menus are identified by surrounding chevrons: '**<**' and '**>**'.

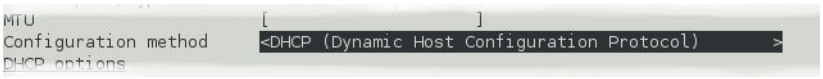

To change the setting of a drop-down menu field move the cursor on the field and activate the menu by using the ↵ key.

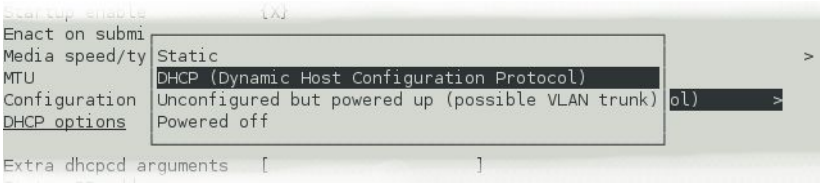

While the menu is active, you can move between options using the the ↑, ↓, ←, → and ⇄ keys: the currently selected option is shown highlighted. When the desired option is selected, press the ↵ key again to confirm the choice and de-activate the menu.

4.3.6 Using forms

Most of the configuration forms have on-line help available. This can be turned on for the current page by using the <Help> button displayed at the bottom of the form. The help text will appear in blue, interleaved with the form itself. Help can also be activated using the F2 key.

Many of the configuration forms have two modes, “simple” and “expert”. They display, by default, only the parameters most likely to be used. For example, the network interface form does not normally display IP aliasing parameters. When access is required to these additional features, they can be displayed by using the <Expert> button displayed at the bottom of the form. They can be hidden again by using the <Simple> button. It is also possible to toggle between simple and expert mode using the F3 key.

Some forms are too large to fit in a single page. In this case, an indicator appears at the top right of the screen. For example, the network interface configuration form, in expert mode, takes three pages to display:

To move to subsequent pages, you can use either the Page Down or the F6 key.

To return to previous pages, you can use Page Up or F5.

When all required fields have been set to the desired values, the changes can be enacted by using the <Submit> button or the F4 key.

It is possible to leave any form by using the <Home> button (or the F1 key) but, if you have made changes to the contents of any fields, the system will warn you that you will lose your changes if you continue.

It offers an opportunity to return to the form:

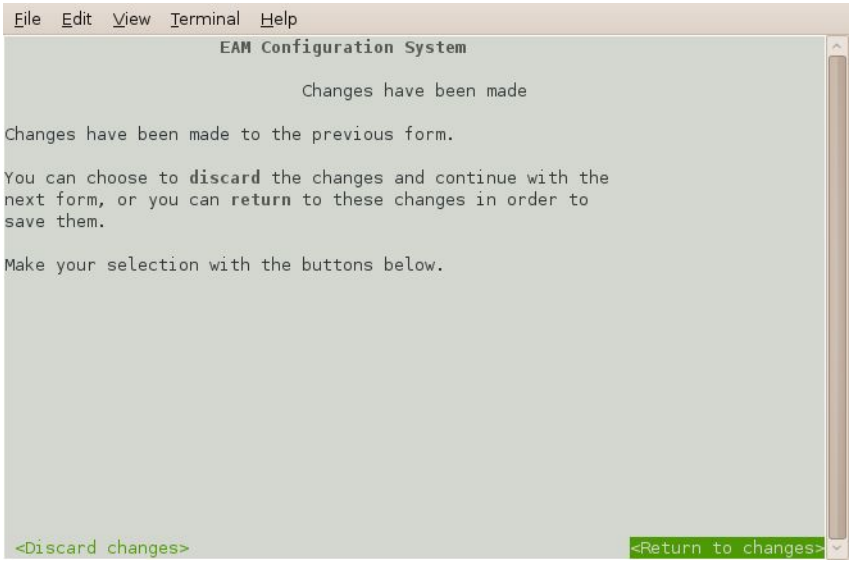

Select `<Discard changes>` to ignore any parameters that you have altered and continue to the home menu or select `<Return to changes>` to review the form and, if desired, submit it before navigating away from it.

The following table provides a summary of all keystrokes that can be used when filling or navigating forms:

Keystrokes used in forms	
Keystroke	**Used for...**
↓ or ⇄	Cursor to next field or, if in an activated drop-down menu field, next item in drop-down menu
→	Cursor to next field or if in a text field move edit cursor one character to the right
↑	Cursor to previous field or, if in an activated drop-down menu field, previous item in drop-down
←	Cursor to previous field or if in a text field move edit cursor one character to the left
↵	Activate or deactivate list field or bottom-line button
space	Toggle check-box
F1	Go to the home menu or, if there already, exit **`gconfig`**.
F2	Display help text
F3	Show “expert mode” fields
F4	Submit the current form
Page Down or F6	Move to the next page of a multi-page form
Page Up or F5	Move to the previous page of a multi-page form

4.4 Configuration Management

Platinum has a comprehensive configuration management system that allows both complete configurations and individual classes of configuration information, such as data processing and networking, to be saved individually and merged during restoration.

Note: This system does not save the configuration of any attached ADC modules (DM24 or CD24). Only the Platinum configuration is stored.

This feature can be very useful when multiple acquisition modules are to be configured for a project. In a typical array with a central communications hub, only two data processing configurations need be created: one for the hub and one for an array element. The latter can then be copied from acquisition module to acquisition module, avoiding having to configure each unit individually. Network configurations need be created for each element of the array and for the hub but these can all be created and stored on a single acquisition module. If the complete set of stored configurations is then copied to each machine and to any "hot spares", then every acquisition module becomes rapidly interchangeable: all that is required to deploy a unit is to restore the correct data processing configuration (hub or element) and then restore the appropriate network configuration.

Configuration files can also be backed up and stored on different sites to provide a disaster management resource.

Arbitrary pairs of stored configurations may be compared to reveal what changed between them.

Note: It is recommended that configuration backups are only restored onto acquisition units with hardware and firmware revision identical to those on the unit on which it was saved. Failure to adhere to this policy could result in unexpected behaviour or, in extreme cases, an unusable system.

Note: Libraries of stored configurations should be updated after each firmware upgrade.

4.4.1 Automatic saving of configurations

The system checks every hour for configuration changes and, if any have been made in the preceding hour, an automatic configuration backup is performed. This provides a useful audit trail and a convenient means to return to previous configurations.

Automatically saved configurations can be distinguished from manually saved configuration by inspecting the file-names displayed on the Configuration Save/Restore web page. Automatically saved configurations have file-names of the the form

`auto-`*`nnnnnnnn`*`-`*`hostname`*

where *nnnnnnnn* is an automatically incremented sequence number and *hostname* is the host-name of the acquisition unit. Configurations saved manually can have arbitrary names; the name is chosen by the user when the configuration is saved. The default file-name is generated from the date, time and hostname, in the form

`YYYY-MM-DD_hhmm_hostname`

4.4.2 Saving a configuration

To save a configuration backup from the web interface, select:

Configuration → Save/Restore

The following screen is displayed:

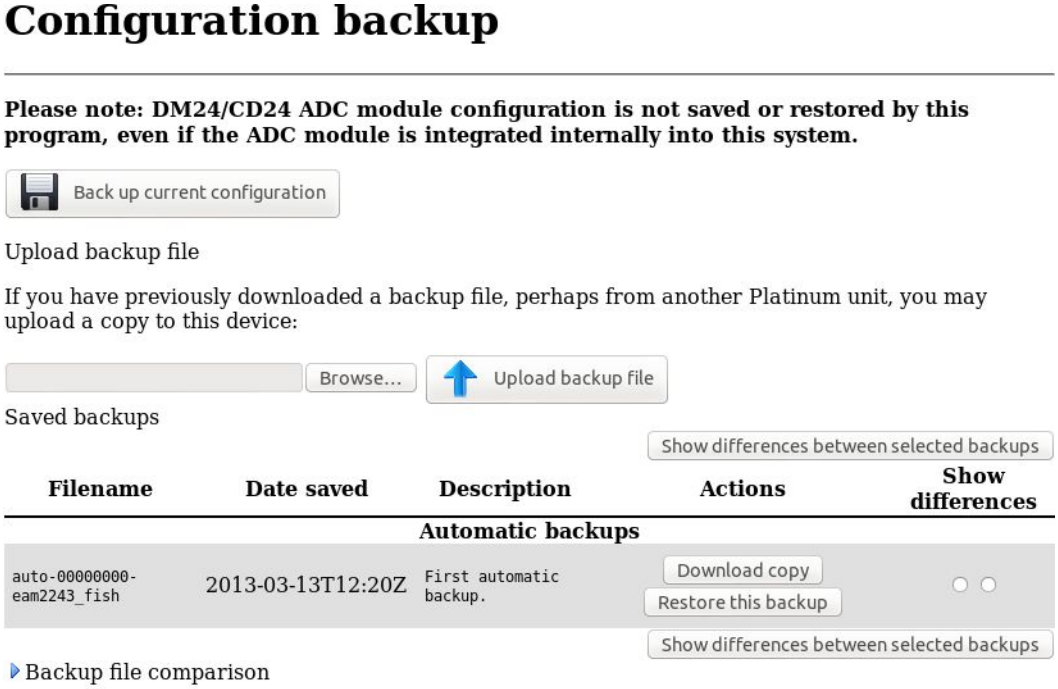

This screen offers a back-up button, an upload facility and a list of stored configurations, each with a download button and a restore button.

Clicking the back-up button displays the following pop-up window:

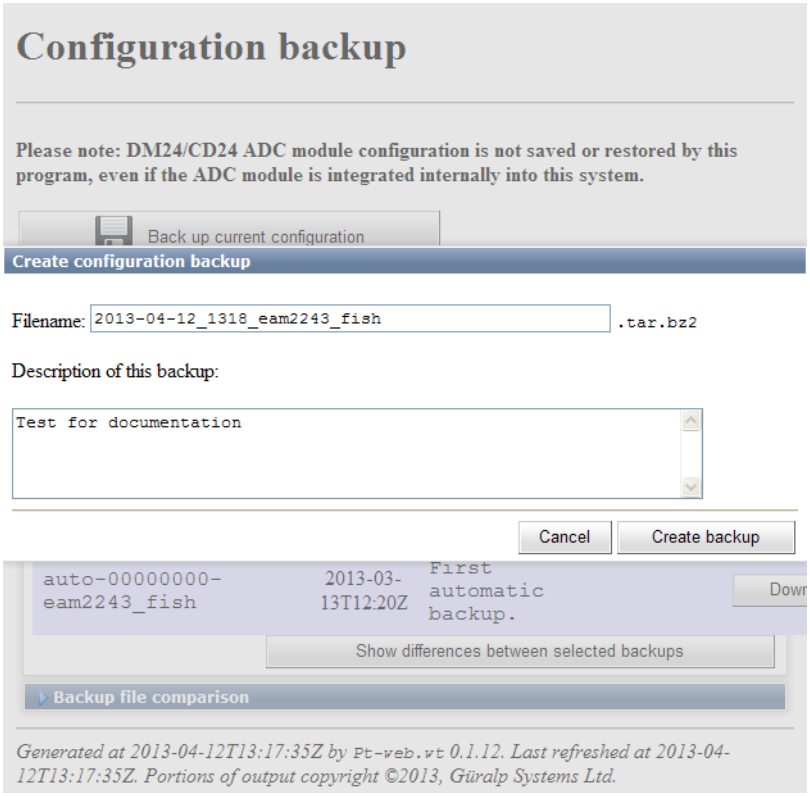

A default file-name is suggested but this can be over-written if desired. You can also enter a description of the back-up, which will be displayed in the list of stored configurations. Click Create backup to save the current configuration to the internal storage. The new stored configuration will appear in the list of stored configurations.

To save a configuration backup from the command line, enter

```
config-backup --backup --comment 'user description' file
```

or, in abbreviated form

```
config-backup -b -C 'user description' file
```

replacing `user description` with your own description of the backup, within single quotation marks (`'`). The configuration will be saved to the file specified by *file*.

4.4.3 Downloading a saved configuration

Using the web interface, choose the required configuration from the list of saved backups and click the associated Download copy button. The backup will be downloaded using your browser's normal download mechanism.

Saved configurations can also be downloaded using file transfer tools like scp and sftp. Manually saved configuration backups can be found in

`/var/lib/config-backup` and automatically saved configuration backups can be found in `/var/lib/config-backup.auto`.

4.4.4 Uploading a saved configuration

Backups that have been saved to a PC can be restored to any acquisition module – not just the one from which it was originally downloaded. This provides a useful mechanism for maintaining a centralised configuration library.

To upload a previously stored configuration backup, click the Browse... button. Your operating system's standard file selection dialogue will open. Navigate to the required file and select it, then click the Upload backup file button. The uploaded configuration will then appear in the list of saved backups and can be restored, if required, as described in the following section.

Saved configurations can also be uploaded using file transfer tools like scp and sftp. They should be copied to `/var/lib/config-backup`.

4.4.5 Restoring a configuration

The same screen in the web interface is also used for restoring a configuration. Simply select the required saved configuration from the list of saved backups and click on the Restore this backup button.

The following screen will be displayed:

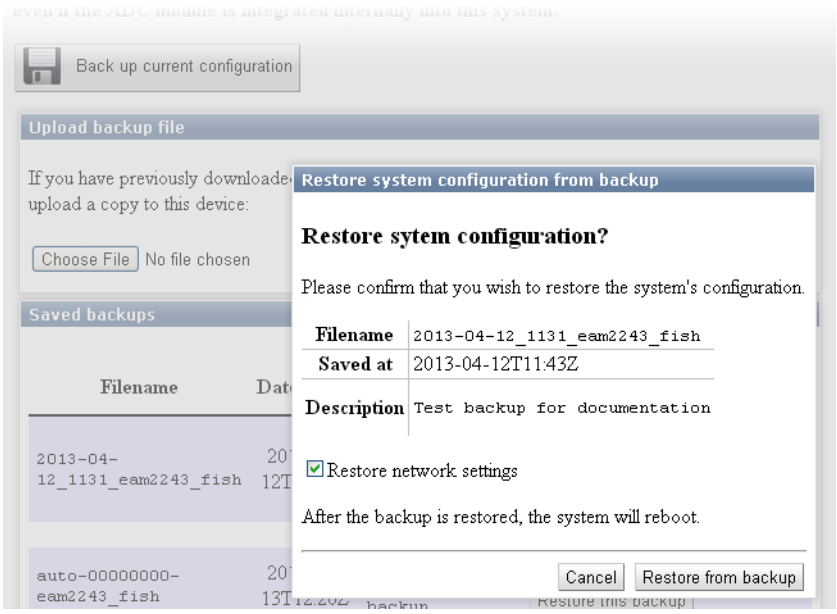

The date and time at which the backup was saved is shown; in the example above, this is the 12th of April, 2013 at 11:43. The file name and description of the backup are also given.

The “Restore network settings” check-box allows you to choose whether or not to restore the network configuration or not. You should select this if, for

example, you are configuring a spare unit to replace an existing system that has been damaged. You should leave it unselected if, for example, you are cloning a configuration across an array of systems (where each would need their own, unique, network settings).

If you are happy that you wish to proceed, click Restore from backup. The backup will be restored and the system will then automatically reboot so that the restored configuration can take effect.

Note: If you you are working on a remote system and have selected "Restore network settings", this action may sever your communication link. Ensure that you understand the consequences before proceeding.

To restore a configuration backup from the command line, enter

```
config-backup --restore path-to-file
```

or, in abbreviated form

```
config-backup -r path-to-file
```

replacing path-to-file with the full path (directories plus file-name) of the previously saved configuration backup file.

If you wish to avoid overwriting the network settings, use

```
config-backup --restore --no-network path-to-file
```

or, in abbreviated form

```
config-backup -r -N path-to-file
```

The system will automatically reboot after restoring the configuration so that the new settings can take effect. If you don't require the system to automatically reboot, add `--no-reboot` or `-R` to the options, as in

```
config-backup --restore --no-reboot path-to-file
```

or, in abbreviated form

```
config-backup -r -R path-to-file
```

4.4.6 Comparing configurations

If more than one saved configuration backup is present on a system, any two can be compared to show the differences. This is useful both to understand what changes have been made to a system and also to compare two different systems. To compare two systems, save a configuration backup on one, download it to a PC, upload it to the second system (without restoring it) and then use the comparison feature to reveal the differences between the two.

Each stored configuration is displayed with two radio-buttons next to it, as shown below:

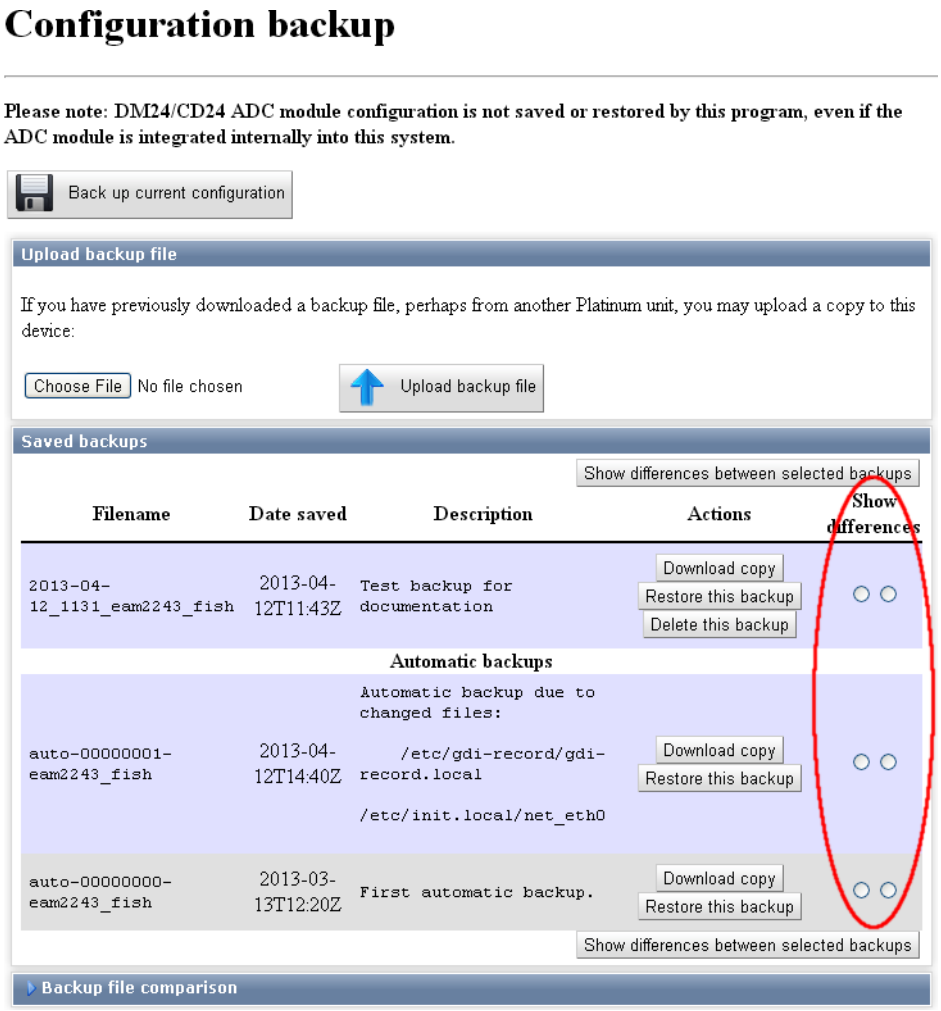

To compare two saved configuration backups, select the left-hand-side radio-button next to one of them and the right-hand-side radio-button next to the other, and then click Show differences between selected backups.

The differences are displayed in Linux `diff` format, with colour-coding used to identify any lines that belong to one file but not the other. Detailed explanation of the possible entries is beyond the scope of this document. A simple introduction to the output from the `diff` command is available at http://lowfatlinux.com/linux-compare-files-diff.html.

To compare two saved configuration backups from the command line, enter the command

```
config-backup --diff path-to-file1 path-to-file2
```

or, in abbreviated form

```
config-backup -d path-to-file1 path-to-file2
```

replacing *path-to-file1* and *path-to-file2* with the full paths (directories plus file-name) of the previously saved configuration backup files to be compared.

4.4.7 Deleting saved configurations

It is possible to delete configuration backups that have been manually saved as well as those that have been uploaded. Automatic configuration backups cannot be deleted: the system will retain the last fifty stored backups and delete any older ones, using the directory cleaner (see section 14.5.1 on page 246).

To delete a stored configuration backup, simply click the Delete this backup button next to the unwanted backup, as highlighted in the screen-shot below.

Stored configuration backups can be deleted from the command line by simply deleting the relevant file in `/var/lib/config-backup` (for user-generated backups) and `/var/lib/config-backup.auto` for automatically saved configuration backups. For example, the command

```
rm /var/lib/config-backup/2013-04-12_1131_MyEAM.tar.bz2
```

will delete a backup performed at 11:31 UTC on the 12th of April, 2013 on a system called "MyEAM".

It is possible to examine the comments embedded in a manually saved configuration backup file using the command

```
config-backup --examine path-to-file
```

or, in abbreviated form,

```
config-backup -x path-to-file
```

Either of the commands will display the comments which were entered by the user (in the "Description" field of the web interface or using `--comment` from the command line) when the backup was created.

4.4.8 Transferring backups between systems

If the command-line program, `config-backup`, is invoked in save or restore mode without a filename argument, standard input or standard output is used, as appropriate. This allows the command to be used in pipe-lines, as in the following examples.

A configuration file can be created on a Platinum system and immediately copied to a linux PC (without being saved on the Platinum system) with a single command:

```
ssh root@platinum config-backup --backup > local.backup
```

where *platinum* is the hostname of the CMG-EAM or other Platinum system.

Likewise, the configuration of a running system can be copied to another with commands like:

(from a Linux PC):

```
ssh root@platinum1 config-backup --backup | \
ssh root@platinum2 config-backup --restore --no-network
```

(from the target system):

```
ssh root@source config-backup -b | config-backup -rN
```

where source is the hostname of the system to be cloned. The argument `-N` prevents the network configuration on the target machine from being over-written.

4.4.9 Technical details

Stored configuration backups are gzipped tar volumes containing all configuration files from the host system. They are each around 33 kilobytes in size.

Automatically-saved configurations are stored in the directory `/var/lib/config-backup.auto` and a system-supplied directory-cleaner instance deletes all but the most recent fifty files. Manually saved configuration backups are stored in `/var/lib/config-backup`.

The command-line interface, `config-backup`, displays full usage details if invoked with `--help`:

```
Program usage:

   config-backup --backup [options] [filename]
   config-backup --restore [options] [filename]
   config-backup --examine [options] [filename]
   config-backup --diff filename1 filename2


Options:

 -h, --help     Display this screen.
 -V, --version  Display program version number.
 -b, --backup   Perform a configuration backup.
 -r, --restore  Restore configuration from a backup.
 -x, --examine  Examine a backup file (prints comment).
 -d, --diff     Compare differences between 2 backups.
 -q, --quiet    Inhibit progress messages.
 -s, --syslog   Log to syslog rather than screen.


Backup options:

 -C, --comment <text>  Comment for this backup.


Restore options:

 -R, --no-reboot       Do not reboot after a restore.
 -N, --no-network      Do not restore network settings
                            (hostname, IP addresses).
```

5 Platinum Firmware Upgrades

Platinum firmware is regularly updated to provide extra features, improve performance and, occasionally, to correct errors. The upgrade process is normally fast, simple and can be carried out remotely using either the web or command-line interface (but see the important notes, below).

Note: This procedure does not upgrade the firmware of connected or embedded digitisers, which should be upgraded using the web interface or the `dm24-upgrade` command as documented in section 8.2.3 on page 115.

The upgrade process makes use of the rsync protocol which uses an elegant and efficient algorithm to, effectively, transfer only the differences between revisions; even within individual files. This significantly reduces the time required compared to traditional upgrade methods.

5.1 Important notes regarding build 10,000

5.1.1 Significant changes at build 10,000

All systems that run Platinum, excluding the CMG-NAM, are based around ARM microprocessors. Platinum relies on the Linux operating system and a number of utility programs which are made freely available by the ARM development community. In 2010, this community agreed a new specification for the way that programs interact with the operating system and each other. This specification is know as the Application Binary Interface, or ABI. The new ABI was called the "EABI" and the original one was retrospectively renamed the "OABI", for *old* ABI. The two paradigms cannot coexist and all new community development will be carried out using the newer EABI.

For this reason, starting with build 10,000, Platinum development also moved to the EABI.

Using the EABI requires the use of a different compiler and, to simplify development, a similar compiler change has been made for CMG-NAM and CMG-NAM64 development.

The decision to use the EABI and the new compiler has one significant consequence – an upgrade from a build before 10,000 to a current build is vulnerable to interruptions in a way that no other upgrade is. This has implications for systems installed in remote locations.

5.1.2 Systems installed in remote locations

Programs using the EABI cannot interoperate (on the same processor) with programs using the OABI. There are also problems when mixing programs

built using the two different compilers. For this reason, it is important that an upgrade from build 3801 to build 10,000 or later is not interrupted.

The upgrade mechanism has, until now, been extremely tolerant of interruptions and it has been normally sufficient to simply restart an interrupted upgrade process. This will also be true for future upgrades, once build 10,000 or higher has first been installed.

For upgrades from build 3801 to build 10,000 or later, however, it is highly recommended that an upgrade to an acquisition module only be carried out in a situation with a stable power supply, a reliable internet connection and where you have physical access to the unit.

5.1.3 Procedures for upgrades spanning build 10,000

5.1.3.1 CMG-EAMs

An extra step must be taken when upgrading from build 3801 to the latest build and special attention must be paid to the amount of free disk space available. Please see this article on our web site for complete information and detailed instructions:

www.guralp.com/articles/20110309-EAM-EABI-upgrade

5.1.3.2 CMG-NAMs and CMG-NAM64s

After upgrading from build 3801 to the latest build, you may see several `FATAL: kernel too old` messages. In order to resolve this issue, please log in to the command line as root and run the command:

```
/sbin/manual-post-upgrade
```

If, despite all precautions, the upgrade is interrupted at a crucial time, attempt first to restart the upgrade. If this fails, please follow the recovery instructions available from our web site at:

www.guralp.com/articles/20110318-NAM-recovery

5.1.3.3 CMG-DCMs

When upgrading a CMG-DCM, the upgrade process should be carried out twice. The second run will resolve problems with broken symlinks. During the first upgrade run, you may see several errors about "*`Directory not empty`*" and references to broken files. These are normal and the second upgrade invocation is designed to resolve these problems and leave the system in a consistent state.

5.2 Determining the current firmware level

To determine the current firmware version from the web interface, select:

Summary → System status

and click on the "Linux system" tab. The screen displays the current version information:

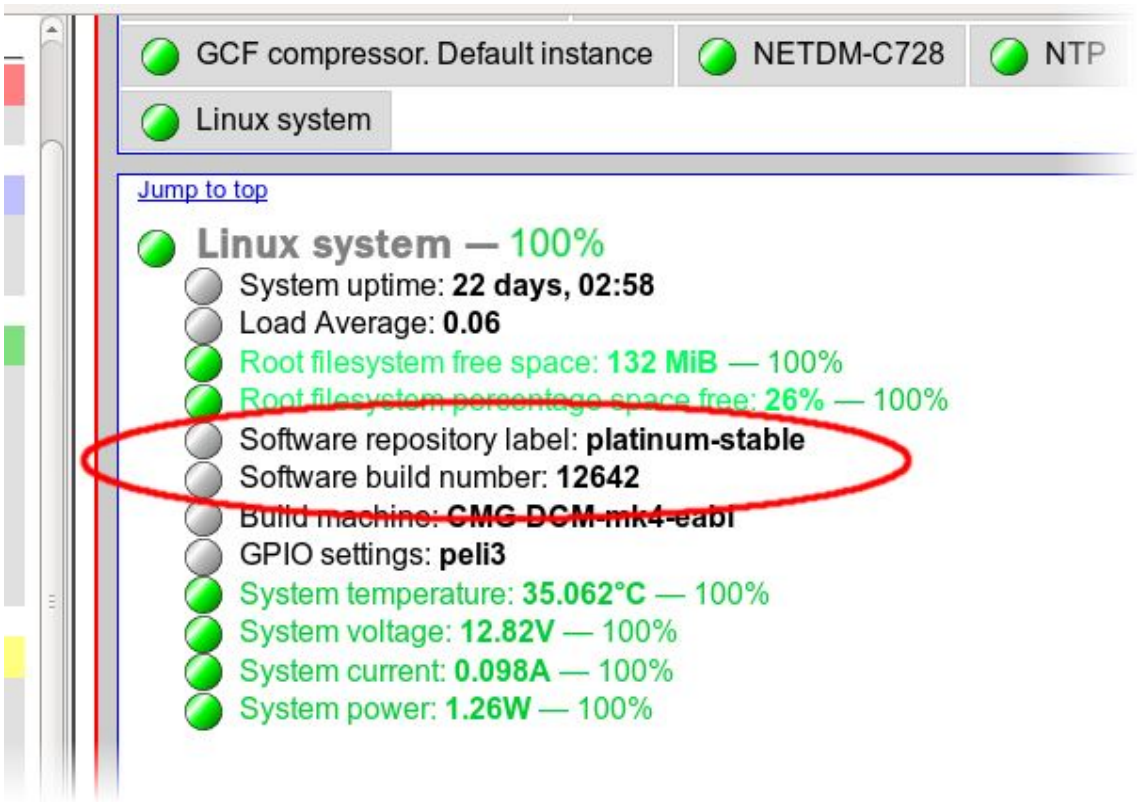

In the example above, the version is 12642.

The same information can be obtained at the command line by using `gconfig` and selecting "Version and serial number information". If you just want quick access to the software build number, this is contained in the file `/etc/build.version`, which can be read with the command:

```
eam2243 ~ # cat /etc/build.version
# Overall build version
BUILD_LABEL="platinum-stable"
BUILD_VERSION="12642"
eam2243 ~ #
```

The version, in this case, can be seen to be 12642.

5.3 Upgrade Methods

This section describes three different methods of upgrading the firmware.

In order to be upgraded, the unit needs access to the latest version of the firmware. If an internet connection is available, Güralp Systems Ltd's software repository can be used. This method is described in section 5.3.1 on page 54.

If a number of units share a common network but that network is not connected to the internet, you can make you own copy of the software repository on a PC or laptop, which can be connected the network either permanently or temporarily, and use that as the upgrade source. This method is described in section 5.3.2 on page 55.

If one or more units are to be upgraded but Internet access is not available, the new firmware can be copied to a USB storage device, such as a memory stick, and the upgrade performed from that. This method is described in section 5.3.3 on page 59.

5.3.1 Upgrading via the internet

In order to upgrade over the Internet from Güralp Systems Ltd's software repository, the unit must have its networking properly configured. In particular, a DNS (Domain Name Service) server and a default gateway (or default route) must both be configured. It is advisable to check these before proceeding or in the event of problems preventing a successful upgrade.

To check for correct configuration of both of these items, issue the command:

```
ping -c3 rsync.guralp.com
```

This will send three “echo request” packets to the GSL upgrade server and listen for responses. If both the DNS server and the correct gateway (default router) are configured, the output will look like this:

```
eam2010 ~ # ping -c3 rsync.guralp.com
PING rsync.guralp.com (80.68.92.160): 56 data bytes
64 bytes from 80.68.92.160: seq=0 ttl=55 time=58.280 ms
64 bytes from 80.68.92.160: seq=1 ttl=55 time=66.845 ms
64 bytes from 80.68.92.160: seq=2 ttl=55 time=56.413 ms
--- rsync.guralp.com ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 56.413/60.512/66.845 ms
eam2010 ~ #
```

If the DNS server is correctly configured but the gateway is not correctly configured, the output will look like this:

```
eam2010 ~ # ping -c3 rsync.guralp.com
PING rsync.guralp.com (80.68.92.160): 56 data bytes
ping: sendto: Network is unreachable
eam2010 ~ #
```

If you are using DHCP, it is advisable to correct this problem by reconfiguring the DHCP server to supply the correct route. If you are using static addressing, enter the address of the Internet gateway router in the “Default route (gateway)” field of the network interface configuration form. See section 7.1.2 on page 72 for more details.

If the DNS server is not configured correctly (or at all), the output will look like this:

```
eam2010 ~ # ping -c3 rsync.guralp.com
ping: bad address 'rsync.guralp.com'
eam2010 ~ #
```

If you are using DHCP, it is advisable to correct this problem by reconfiguring the DHCP server to supply the correct name-server details. If you are using static addressing, enter the address of a suitable DNS server in the "Nameserver" field of the network interface configuration form (only available in expert mode). See section 7.1.2 on page 72 for more details.

Note: If your internet connection is via a firewall, you must ensure that port 873 (rsync) is open for TCP connections initiated by the acquisition module. The following paragraphs explain how to check this.

To check that your internet connection is correctly configured, enter the command:

```
rsync rsync.guralp.com::
```

The output should look like this:

```
eam2010 ~ # rsync rsync.guralp.com::
platinum-crosslib    Libraries required for cross-compiling
on the Platinum platform
platinum-stable Stable filesystems for the Platinum platform
platinum-builder     Open source parts of Platinum firmware
build tree (large)
eam2010 ~ #
```

Any other output indicates a problem with your firewall – consult your network administrator or see the trouble-shooting advice in section 16.4 on page 270.

Once the network has been checked, you can proceed to upgrade the unit by following the instructions in section 5.4 on page 61.

5.3.2 Upgrading from a local mirror

Setting up a mirror involves three steps:

1. Downloading the mirror content.
2. Setting up a local rsync server.
3. Configuring the client acquisition devices to use the new server.

Note: Due to Windows file-system restrictions, it is only possible to build a mirror server on a PC or laptop running Linux.

5.3.2.1 Downloading the mirror content

The mirror can occupy a significant amount of disk-space, depending on which architectures you need to support. See the sections for each architecture (below) for the current space requirements. You should pick a file-system with ample space in which to store your own copy. In order to simplify the download, we recommend that you start with an empty directory each time. If you wish to make a fresh copy after a new firmware release, it is much easier to create this in an empty directory than to "update" an existing mirror. You can keep multiple, simultaneous versions of the firmware if you wish and tell each acquisition device which version to use when upgrading.

The server on which you create the mirror should have access to the Internet during the download step but does not need Internet access while it is acting as an upgrade server. It does, of course, need to be accessible by your networked acquisition devices. It is possible to create the mirror content on a removable mass storage device attached to an Internet-connected computer and then move the device to a different computer when performing the upgrades. The removable mass storage device must use a Linux-compatible file-system, such as ext3. FAT- and NTFS-formatted devices will not work properly for this purpose.

Create the mirror directory and use the cd command to make it your current directory. As root, enter one or more of the following command sequences to download the mirror content. Each sequence downloads the files for a particular architecture. If you know, for example, that you will never want to upgrade a CMG-NAM64, you can omit the commands for this architecture.

Note: Be careful not to omit the final '.' or the space before it in the rsync commands below.

CMG-DCMs

This architecture currently requires around 50MB of mass storage device space for the mirror. The commands to download the content are:

```
GSLSRC=rsync.guralp.com/platinum-stable/CMG-DCM-mk2x
rsync -EgHloprtv --exclude resolv.conf rsync://$GSLSRC .
```

CMG-EAMs

There are currently two sets of firmware available for the CMG-EAM: a frozen image of build 3801 (for users who do not yet wish to upgrade to post 10,000 builds) and the latest build.

If you do not wish to upgrade beyond build 3801 yet, the files for build 3801 require around 53MB of storage space for the mirror. Use the following commands:

```
GSLSRC=rsync.guralp.com/platinum-stable/CMG-DCM-mk4
rsync -EgHloprtv --exclude resolv.conf rsync://$GSLSRC .
```

If the systems being upgraded are already at build 10,000 or later, you only require the latest build, which occupies around 70MB of space. Use the following commands:

```
GSLSRC=rsync.guralp.com/platinum-stable/CMG-DCM-mk4-eabi
rsync -EgHloprtv --exclude resolv.conf rsync://$GSLSRC .
```

If you are upgrading systems from earlier than build 3801 to the current build, you will need both sets of firmware, requiring around 123MB of storage. Use the following commands (note the additional punctuation):

```
GSLSRC='rsync.guralp.com/platinum-stable/CMG-DCM-mk4*'
rsync -EgHloprtv --exclude resolv.conf rsync://"$GSLSRC" .
```

CMG-NAMs

This architecture currently requires around 94MB of mass storage device space for the mirror.

```
GSLSRC=rsync.guralp.com/platinum-stable/CMG-NAM
rsync -EgHloprtv --exclude resolv.conf rsync://$GSLSRC .
```

CMG-NAM64s

This architecture currently requires around 125MB of mass storage device space for the mirror.

```
GSLSRC=rsync.guralp.com/platinum-stable/CMG-NAM64
rsync -EgHloprtv --exclude resolv.conf rsync://$GSLSRC .
```

5.3.2.2 Setting up a local rsync server

Your local rsync server is configured by creating the file `/etc/rsyncd.conf`. If the serving host already runs an rsync server, you should modify this file (by adding an extra module) in order to allow access from the acquisition devices to the mirror directory: we assume that you have the knowledge to do this without further assistance. This section covers setting up a new, dedicated rsync server.

You will need to choose a TCP port number which will not conflict with another service on your network. The port number should be greater than 1024 in order to avoid additional complexity. Consult your network administrator for an available port or simply try 61616 and, if you get an error saying that the port is in use when you attempt to start the server, choose a different random number in the range 49152 - 65535. 61616 will be used in the following example and should be replaced with the port number you have chosen or been allocated. If there are firewalls between your server and the acquisition devices, you will need to open channels through them for this port.

You will also need to choose a module name for the server. This can be any descriptive string but, for simplicity, it is best to stick to numbers, lower-case letters and hyphens (-). The name `platinum-local-mirror` has been used in the following example and should be replaced with the module name you have chosen.

Create the file `/etc/rsyncd.conf` with the following contents:

```
port = 61616
[platinum-local-mirror]
path = /path/to/your/local/mirror/directory
comment = GSL-EAM firmware
numeric ids = yes
log file = /path/to/writeable/log/file
timeout = 600
hosts allow = *
```

Consult the manual page for `rsyncd.conf(5)` for details of further options that you can use in this file, including optional security improvements. This is available on-line at http://man-wiki.net/index.php/5:rsyncd.conf.

Once the `/etc/rsyncd.conf` file is ready, you can start the rsync server with the command

```
sudo rsync --daemon
```

If you want to run the rsync server permanently, it is possible to start it via inetd, xinetd or an rc script. Consult the manual page for `rsyncd.conf(5)` for further details.

5.3.2.3 Configuring the upgrade system to use the new server

The standard upgrade source must be over-ridden: on each system to be upgraded. Create the file `/etc/conf.d/upgrade.local` with the following contents:

```
RSYNC_HOST="address.of.my.server”
RSYNC_PORT="61616"
RSYNC_MODULE="platinum-local-mirror"
```

replacing:

- *`address.of.my.server`* with the DNS name or IP address of the mirror server;
- *`61616`* with the port number you chose earlier; and
- *`platinum-local-mirror`* with the module name you chose earlier;

Note the quotation marks around the variables.

The acquisition devices can now be updated from the mirror by following the instructions in section 5.4 on page 61. Note that the configuration files

`/etc/conf.d/upgrade.local` on each acquisition devices will not be disturbed by the upgrade process and, so, only need to be created once.

5.3.3 Upgrading from a USB storage device

For situations where it is either impossible or undesirable to upgrade over a network, Güralp Systems Ltd can supply the latest Platinum firmware on a USB memory stick, along with an appropriate adaptor cable, part number CAS-DCM-0038.

The adapter cable is required when upgrading the firmware of most acquisition module units but not when upgrading a CMG-NAM.

You will need both physical access and command-line access to the device being upgraded. Command-line access may be via ssh or a serial connection.

To upgrade the firmware from a USB storage device:

1. Power up the USB ports (CMG-EAM only). The hardware for other acquisition modules do not have control over the 5V USB supply and power is always provided to the USB ports.

Depending on the revision of CMG-EAM firmware installed before the upgrade, there are three possible commands for powering up the USB ports. Ignoring any error messages, enter these commands at the terminal:

```
echo 1 > /sys/class/gpio/USBpowerB/level
ioline -L USBA_power -o 1
ioline -L USBB_power -o 1
```

Note: Using the wrong command is completely harmless and merely results in an error message, so it is easiest - and perfectly safe - to simply enter all three commands at the terminal.

2. Connect the firmware pod.

If you are upgrading a CMG-EAM, CMG-DAS or CMG-DCM, the supplied cable (CAS-DCM-0038) should be attached to the 6-pin bayonet USB connector of the module. The firmware pod is then connected to the end of this cable.

A CMG-NAM unit has standard USB sockets and these can be used to directly connect the firmware pod. You may use the USB socket on the front (or any on the rear) of the module.

3. Wait for the device to be scanned and registered by the operating system. You can confirm this by looking for USB mass storage registration entries in the system log file, `/var/log/messages`.

4. Mount the file system by entering the following command:

```
mount -t ext2 -o ro /dev/disk/by-label/Pt-firmware /mnt
```

Note: Some early Platinum releases only had ext3 file-system support. If this is the case, the above command will result in an "Invalid argument" error message. If this occurs, simply change **`ext2`** in the above command to **`ext3`**.

5. Run the upgrade script.

This is not the same script that is used for network upgrades but it takes the same optional arguments. These arguments are documented in section 5.4 on page 61, which should be read before proceeding.

Once you have decided which argument, if any, you wish to use, run the script with the command:

```
/mnt/upgrade optional_argument
```

6. Ensure there are no error messages and then reboot the device. Once the unit has rebooted, the upgrade process is complete.

Note: Certain upgrades (depending on the initial version number) make it difficult or impossible to reboot the system via the web interface. In these cases, the system can be rebooted from the command line (with the `reboot` command) or by power-cycling.

5.3.4 U3 USB mounting problems

U3 was a method of launching Windows applications from special USB flash drives. A U3 flash drive presents itself to the host system as a USB hub with a CD drive and standard USB mass storage device attached.

Reformatting the drive will remove some of the software (the hidden "SYSTEM" folder), but not all of it. The virtual CD-ROM drive cannot be removed by reformatting because it is presented to the host system as a physical device attached to a USB hub.

U3 Tool is an open-source management tool (for Windows and Linux) that allows the locked U3 partition to be removed.

The U3 tool is available from http://resourcelessness/projects/u3-tool/files/.

5.3.4.1 MS Windows

Open a command prompt and change the directory to that where the U3 tool was installed. To see all options, enter the command:

```
u3-toolmaker
```

To repartition the device, enter the command:

```
u3-toolmaker -p 0 E
```

replacing ***E*** with the device's drive letter

5.3.4.2 Linux 2.6.20+

To see all options, enter the command:

```
u3-tool
```

To repartition the device enter the command:

```
./u3_tool -U /dev/sg3
```

Replace **/dev/sg3** with the SCSI generic device associated with your device. The correct device can be deduced from the output from the command `dmesg`.

5.3.4.3 Linux 2.4

To unlock the “secured data partition” using Linux's USB subsystem:

```
# modprobe -r usb-storage
# ./u3_tool -u scan
# modprobe usb-storage
```

After unlocking, the device can be used normally.

5.4 Upgrade Types

There are three different types of upgrade, each of which is described below. When upgrading via the web interface, the desired type is selected by pressing the appropriate button. When upgrading directly from the command line or from a USB storage device, the required type is selected by the use or omission of command-line arguments.

5.4.1 Standard upgrade

The standard upgrade brings the firmware to the latest revision while respecting and preserving all configuration settings.

Note: All files on any connected hard drive or USB drive are left untouched, as are any files in the directories `/home`, `/root`, `/var` and `/usr/local`. In addition, any file with an extension of `.local` will be preserved: this is the mechanism by which most configuration settings are safeguarded.

To perform a standard upgrade from the web interface, select:

Tools → Firmware

The following screen will be displayed:

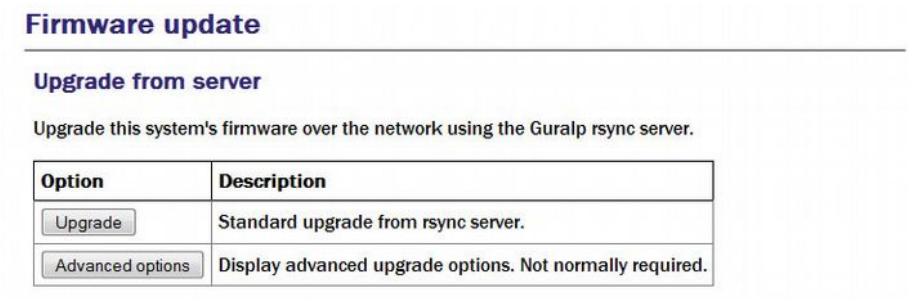

Press the Upgrade button and watch the screen for any error messages.

To perform the same upgrade from the command line, simply enter the command:

```
eam999 ~ # upgrade
```

with no arguments. Watch the screen for any error messages and then reboot to complete the process.

Note: Certain upgrades (depending on the initial version number) make it difficult or impossible to reboot the system via the web interface. In these cases, the system can be rebooted from the command line (with the `reboot` command) or by power-cycling.

5.4.2 Upgrade and restore defaults

The standard upgrade respects and preserves user configuration settings. In some circumstances it may be necessary to overwrite these settings and return all configuration settings to their factory defaults. The unit is not completely restored to factory condition: this allows for the possibility of implementing customisations and user-developed scripts which persist across upgrades.

Note: All files on any connected hard drive or USB drive are left untouched, as are any files in the directories `/home`, `/root`, `/var` and `/usr/local`. Files with the extension `.local` are deleted.

To restore defaults while upgrading from the web interface, select:

Tools → Firmware

and then click Advanced options

The following screen will be displayed:

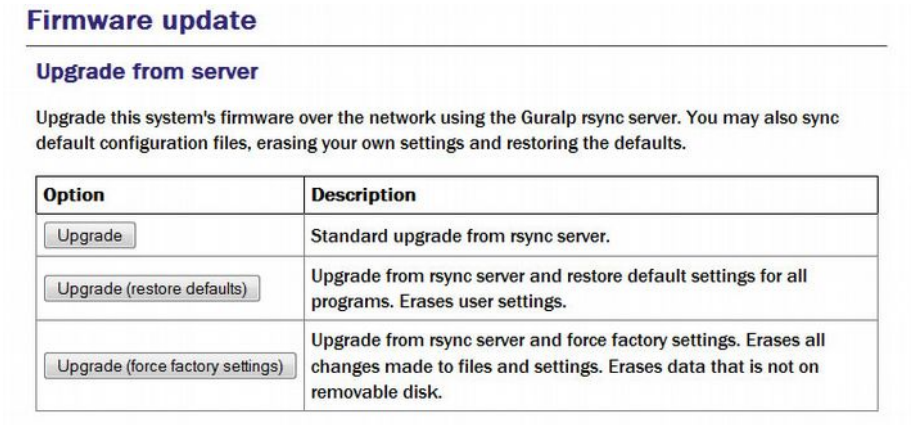

The first button, Upgrade, does exactly the same as the similar button on the previous screen. The Upgrade (restore defaults) button performs the action described in this section.

To perform this action from the command line, invoke the upgrade script with the argument `--restore-defaults`:

```
eam999 ~ # upgrade --restore-defaults
```

Watch the screen for any error messages and then reboot the unit to complete the process.

Note: Certain upgrades (depending on the initial version number) make it difficult or impossible to reboot the system via the web interface. In these cases, the system can be rebooted from the command line (with the `reboot` command) or by power-cycling.

5.4.3 Upgrade and force factory defaults

The third upgrade option effectively wipes everything (other than files on any connected hard drive or USB drive) while installing the new firmware revision, leaving the unit as it would be delivered. Avoid using this option if you have made any customisations to your unit or installed any scripts. If in doubt, please consult Güralp Systems Ltd technical support for advice before proceeding. Conversely, if you have made changes which you believe have adversely affected the unit but are having trouble undoing them, this option lets you start with a clean slate.

To invoke this option from the web interface, select "Firmware" from the "Tools" menu and then click the Advanced options button. From the resulting screen, press the Upgrade (force factory settings) button. Watch the screen for any error messages and then reboot the unit to complete the process.

Note: Certain upgrades (depending on the initial version number) make it difficult or impossible to reboot the system via the web interface. In these cases, the system can be rebooted from the command line (with the `reboot` command) or by power-cycling.

To perform this action from the command line, invoke the upgrade script with the argument `--force-factory-settings`:

```
eam999 ~ # upgrade --force-factory-settings
```

Watch the screen for any error messages and then reboot the unit to complete the process.

5.5 Upgrade logs

The upgrade process stores all progress and error messages in the file **`/var/log/upgrade.log`**. If you suspect that there has been a problem with an upgrade or you wish to have full details of what has changed, you can inspect this file by issuing the command

```
eam999 ~ # less /var/log/upgrade.log
```

You can scroll forward through the file simply by pressing the space key. For more control, you can move forward and backwards, line by line, with the ↑ and ↓ keys or, page by page, with the Page Up and Page Down keys.

The Q key should be used to return to the command line.

If you wish to obtain a copy of this log file, it can be copied from the system to an external computer either via the serial port (see section 11.3.1.4 on page 157) or over the network (see section 11.3.1.2 on page 152).

6 Data Handling

6.1 Introduction

The data handling system of the acquisition module is very flexible, due to its modular software architecture. In the simplest configuration, all data flowing through the acquisition module is routed through a single multiplexor module called `gdi-base`. This communicates directly with all input modules, which handle the various incoming data streams, and all output modules, which convert the data into the required formats. All incoming data are stored and accessed internally in an intermediate format, regardless of the format in which it was originally received.

Note: The sole exception to this is incoming CD1.1 data which, for reasons related to frame signing (an authentication technology), do not pass through `gdi-base`. CD1.1 data-handling and the associated system configuration are covered in a separate manual, MAN-EAM-1100, which is available for download from www.guralp.com.

The diagram below shows the basic internal organisation of an acquisition system, ignoring the CD1.1 subsystem:

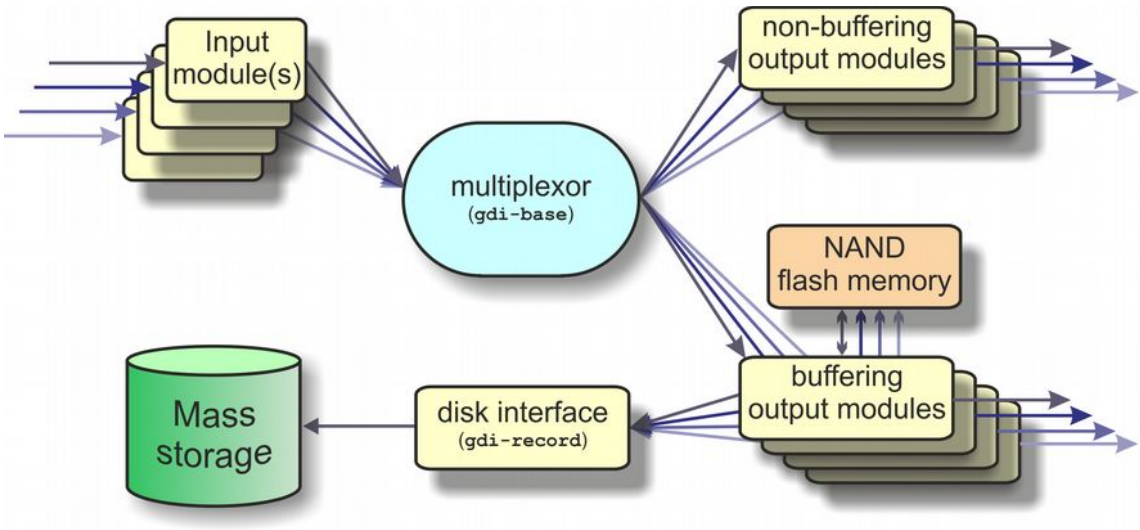

The multiplexor makes incoming data available to the output modules. These come in two flavours: simple modules (such as those for WIN, GSMS and QSCD) simply convert the data streams and output them in the required format; other modules maintain a ring-buffer which is used to, for example, satisfy BRP back-fill requests. The ring-buffers use the NAND flash memory. These output modules also send data to `gdi-record`, which handles all mass storage device writing requests, regardless of format.

The `gdi-base` and `gdi-record` programs are designed to be stateless, so that the data on the mass storage device are always consistent. This means the system is tolerant of power interruptions and of the mass storage device being removed without notice.

Any number of input modules can be configured to acquire data in any supported format from any source, simultaneously. These modules convert their data and pass it to the multiplexor. Data can be acquired in any of the following formats, from multiple sources:

- BRP via serial lines
- Scream, via Ethernet or ppp
- GDI-link, via Ethernet or ppp
- CD1.1

The architecture has been designed to support the addition of extra formats simply by adding input modules. Please contact Güralp Systems if you have requirements which are currently unsupported.

Any number of output modules can be configured to send data in any supported format to any destination. The following data formats are supported:

- GCF output via serial port or TCP stream
- GCF output via Scream (TCP or UDP)
- SEEDlink
- CD1.1;
- GDI (Güralp Data Interconnect)
- GDI-link
- WIN format output
- QSCD - Quick Seismic Characteristic Data (designed by KIGAM) output
- GSMS (Güralp Seismic Monitoring System) output

In the default, factory configuration, the acquisition module is configured to receive serial GCF input on all of its serial ports except Data Out. There is a single chain of data through the multiplexor to a Scream server. Data are also recorded to the configured mass storage device in GCF format. This is shown in the following diagram:

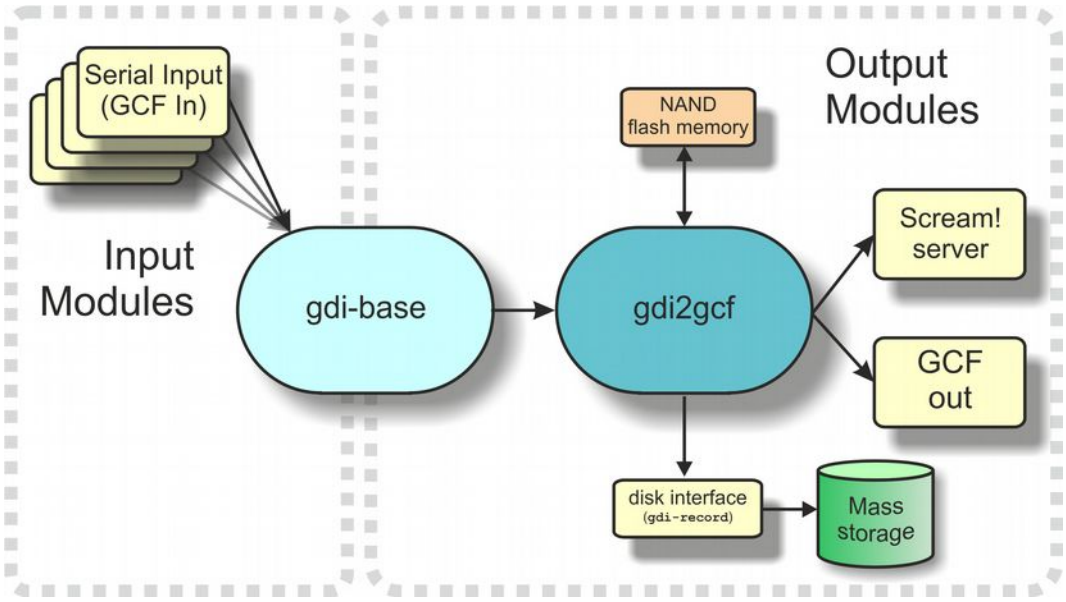

The `gdi2gcf` module, know as the GCF compressor, is responsible for re-blocking GDI samples into GCF blocks. It provides data to all GCF output modules as well as to the `gdi-record` module, which writes GCF files to the mass storage device. It has a number of configuration options, which are described in section 11.2 on page 138.

A similar arrangement applies to miniSEED data: the `gdi2miniseed` module provides data to the SEEDlink server and to `gdi-record`.

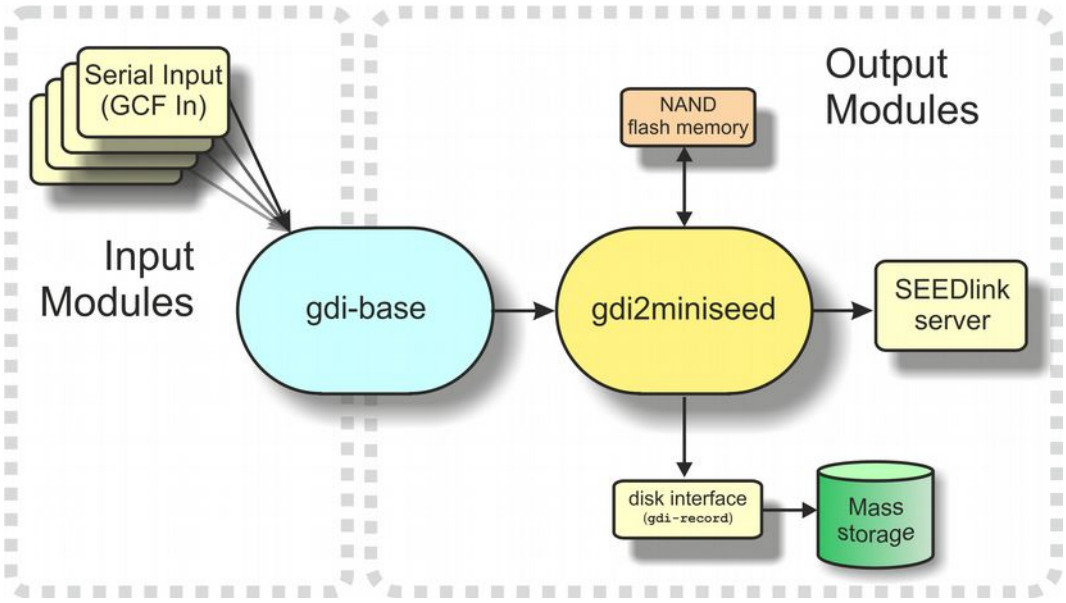

6.2 Configuring gdi-base

The `gdi-base` module requires no configuration in most applications. For very complex situations, however, it may be necessary to create additional

instances. The information in this section is provided the interests of completeness.

To configure a gdi-base instance on the acquisition module from the web interface, select:

Configuration → Service → Expert → Advanced

or

Configuration → All options → System services → Expert → Advanced

To configure an instance from the command line, start gconfig and select "System services" from the top level menu.

Now select "gdi-base -- Data transport and channel manager".

The screen shows a list of all gdi-base instances that have been configured. Select the appropriate link to edit an existing instance or to create a new one.

6.2.1 Configurable parameters

The configurable parameters for gdi-base are contained in a single form:

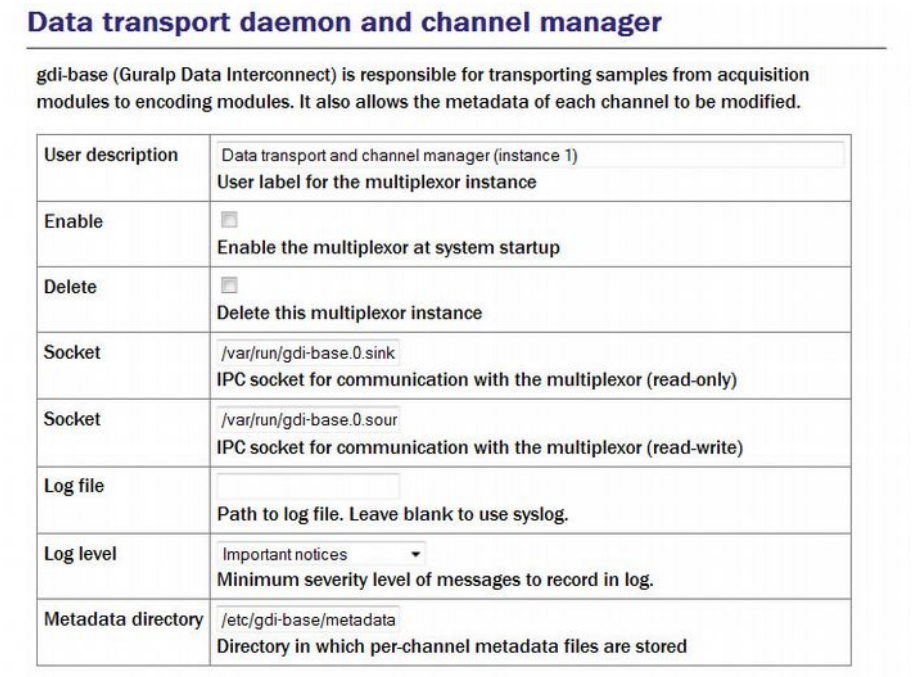

Data transport daemon and channel manager

gdi-base (Guralp Data Interconnect) is responsible for transporting samples from acquisition modules to encoding modules. It also allows the metadata of each channel to be modified.

User description	Data transport and channel manager (instance 1) User label for the multiplexor instance
Enable	☐ Enable the multiplexor at system startup
Delete	☐ Delete this multiplexor instance
Socket	/var/run/gdi-base.0.sink IPC socket for communication with the multiplexor (read-only)
Socket	/var/run/gdi-base.0.sour IPC socket for communication with the multiplexor (read-write)
Log file	 Path to log file. Leave blank to use syslog.
Log level	Important notices Minimum severity level of messages to record in log.
Metadata directory	/etc/gdi-base/metadata Directory in which per-channel metadata files are stored

User description: Sets the name of the instance; this should be set to a meaningful name for the function it will be performing.

Socket: Allows specification of the inter-process communication socket used by this instance. This should not normally be changed.

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field

can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to **`/var/log/messages`**.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

Metadata directory: Allows specification of the directory where this instance stores intermediate information. This should not normally be changed

6.3 Using compressors

The two compressors mentioned at the beginning of section 6 on page 65 convert from GDI to GCF and from GDI to miniSEED format. They provide channel filtering, channel name mapping and data buffering for the `gdi-record` module, which writes GCF and miniSEED files to the mass storage device.

One instance each of `gdi2gcf` and `gdi2miniseed` are present in the default configuration. Additional instances may be created as required. This will be necessary if you have different channel filtering requirements for, say, recording and transmitting or if you need different transmitters to send different sets of channels to different destinations. The configuration page for every transmitter has, in the “Expert mode” options, a drop-down menu which allows the operator to select which compressor instance to use for its input. The gdi-record configuration page has a similar facility.

Instances of either compressor are “dependant services”, meaning that they do not need to be (and should not be) configured to start automatically when the system boots. They will be started whenever a client service, such as a connected transmitter (or gdi-record), starts.

Configuration of the GCF compressor, `gdi2gcf`, is described in section 12.1.1 on page 162.

Configuration of the miniSEED compressor, `gdi2miniseed`, is described in section 12.2.1 on page 180.

7 Networking Configuration

Platinum firmware includes comprehensive support for Ethernet networking. Features include VLAN (virtual network) support, a PPP implementation (IP over serial lines), an `iptables` firewall and IPV6 support.

Minimal network configuration is described in section 3.2 on page 20. Those steps will allow you to communicate with your device over a network. The configuration changes made in that way will not, however, survive a reboot. To make the configuration permanent, follow the procedures in this section.

7.1 Configuring physical network interfaces

Most acquisition modules have a single physical network interface. CMG-NAMs are typically equipped with multiple physical network interfaces. Platinum firmware follows the standard Linux convention of naming the first physical network interface present on a system `eth0` and subsequent interfaces `eth1`, `eth2`, etc.

To configure a physical network interface from the web interface, select:

Configuration → Networking → Interfaces

or

Configuration → All options → Networking

The following screen will be displayed:

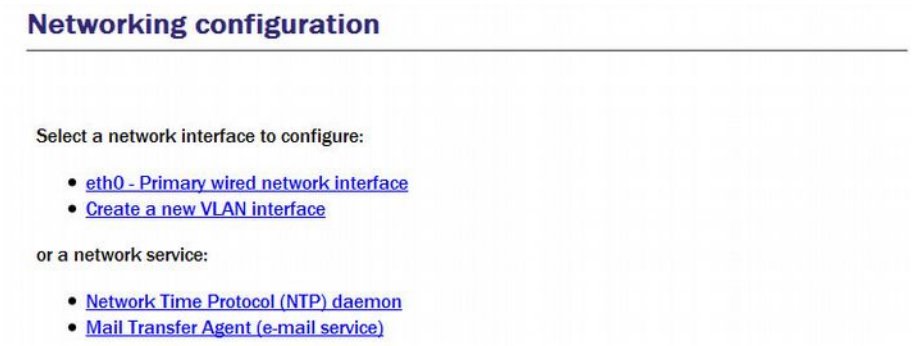

To configure a physical network interface from the command line, start gconfig and select “Networking” from the top level menu.

The first link on the screen takes you to the configuration page for the first physical network interface. If your hardware has multiple physical interfaces, you may need to create configurations for them using the “Create a new interface” button. Once created, the can be configured in an identical manner to eth0, as described below.

7.1.1 Configurable parameters in simple mode

The configurable parameters for physical network interface have two tabs in simple mode: Interface and Static IP.

7.1.1.1 Interface

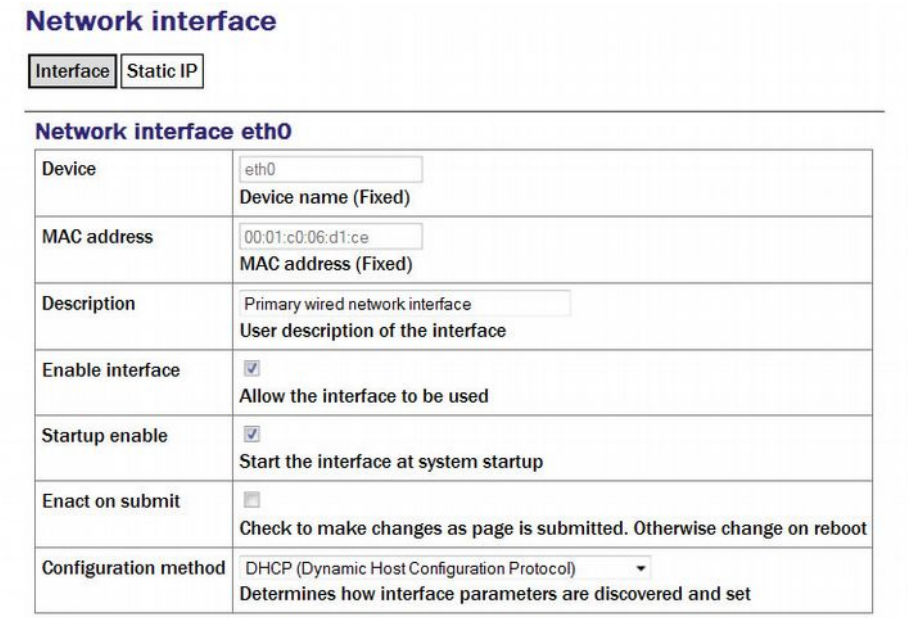

Device: Not editable. It displays the name of the network interface being configured.

MAC address: Not editable. It shows the *Media Access Control* address of the adapter's hardware. It is often useful to know this when configuring DHCP servers: by binding an IP address to a particular MAC address, the DHCP administrator can ensure that the device retains the same IP address across reboots.

Description: Allows the operator to modify the description of this interface in configuration dialogues and error messages. This is of limited value when there is only a single interface but, for example, when a CMG-NAM has multiple interfaces, it may be useful to rename them in order to reflect their logical function rather than their physical position.

Enable interface: Enables the interface when checked. No other configuration settings are changed when the interface is disabled, allowing use of the interface to be suspended without deleting the configuration.

Startup enable: Controls whether the interface is enabled automatically when the unit boots.

Enact on submit: Controls whether changes made using the rest of this form take effect immediately or are only written to the configuration files. When

this box is cleared, changes will only take effect the next time the unit is booted or the interface is reconfigured with this box ticked.

Configuration method: The drop-down menu offers the following choices:

- **Static:** The interface will take its address and routing parameters from values entered by the operator.
- **DHCP (Dynamic Host Configuration Protocol):** The interface will attempt to obtain its address and routing parameters from a DHCP server.
- **Unconfigured but powered up (possible VLAN trunk):** The interface will not be used directly but is available for carrying virtual network (VLAN) traffic.
- **Powered off:** The interface will not be used and the interface chip is disabled, reducing the total power consumption by around 200mW.

7.1.1.2 Static IP

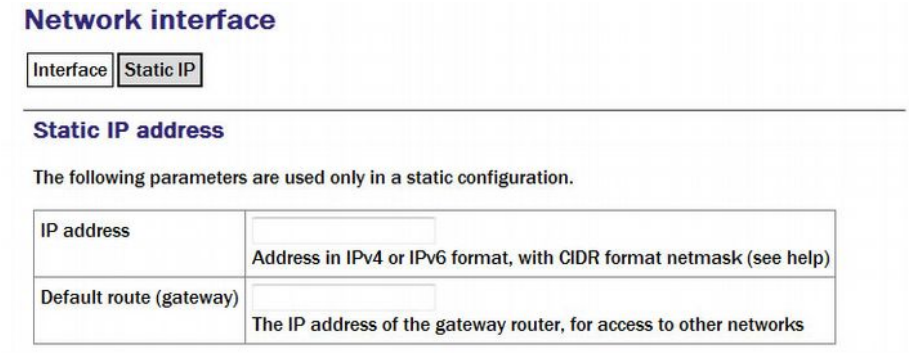

IP address: Only used if the **Configuration method** drop-down menu is set to "Static". The address should be entered in CIDR format, where the address is followed by a slash and then the number of bits defining the netmask, e.g. `192.168.0.1/24` for IPV4 or `2001:db8:/32` for IPV6.

For more information about CIDR, please refer to:
http://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing.

Default route (gateway): Should be populated with the IP address of the default router. If more complicated routing configurations are required, these can be entered in expert mode.

7.1.2 Configurable parameters in expert mode

When in expert mode new fields are available on the "Interface" and "Static IP" tabs and there are three new tabs: DCHP options, IP aliasing and Routes.

7.1.2.1 Interface (Expert)

Media type/speed: This drop-down menu offers the following options for controlling the communication speed and duplex mode of the network link:

- Automatically detected and set
- Restrict speed to 10Mbps. Recommended to save power
- Fixed 100baseTx, full duplex
- Fixed 100baseTx, half duplex
- Fixed 10baseTx, full duplex
- Fixed 10baseTx, half duplex

MTU: Allows the Maximum Transfer Unit to be set for the network link. This parameter controls the maximum packet size used for outgoing network packets. If any segment of a link between two systems has a restriction on packet size, larger packets flowing across the link must be fragmented - broken into smaller parts - and then re-assembled on arrival. This is inefficient and can badly affect the throughput of a link. In such situations, it makes sense to restrict the maximum packet size at the sender (to match the limitation) so that all packets can pass unimpeded.

There is no method to empirically determine the optimum MTU for a given link from the acquisition module itself but, if the link is to a PC (or, in the case of, say, a link between two acquisition modules, one end can temporarily be replaced by a PC) the PC can be used to investigate the link properties and the correct value can be obtained.

For more information, please see http://www.dslreports.com/faq/695. Note that, if testing from a PC running Windows, the MTU of the Windows PC should be set to 1500 before starting the test.

7.1.2.2 DCHP options

Extra dhcpcd arguments: Used to change the operation of the DHCP client. Please see http://man-wiki.net/index.php/8:dhcpcd for information about what options can be added here.

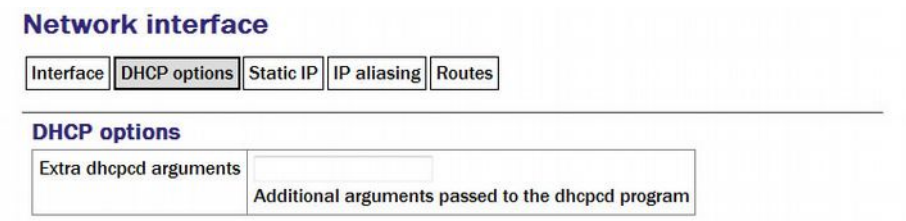

7.1.2.3 Static IP (Expert)

Extra ip addr arguments: Used to tune the operation of the network interface. A non-standard broadcast address can be specified by entering `broadcast` *`broadcast_address`*. For other settings that can be used here, please see http://linux.die.net/man/8/ip.

Nameserver: Used to specify the IP address of the DNS server for your network. This field must be set correctly before internet firmware upgrades can be used. A secondary DNS server's address can be added in the **Backup nameserver** field.

Default route (gateway): Populate with the IP address of the gateway router, for access to other networks or to the Internet. This field must be set correctly before internet firmware upgrades can be used.

IP route arguments: Used to modify the invocation of the `ip route add` command in order to, e.g., set the *route metric*. The options that can be set here are mostly highly technical and should rarely be required. Please see http://linux.die.net/man/8/ip for more information.

7.1.2.4 IP aliasing

IP aliasing is used to add extra addresses to this interface, a technique known as multi-homing. By default, the table displays three blank rows but, should you need more, complete the first three and submit the form: it will be re-drawn with extra blank rows. Alternatively, clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

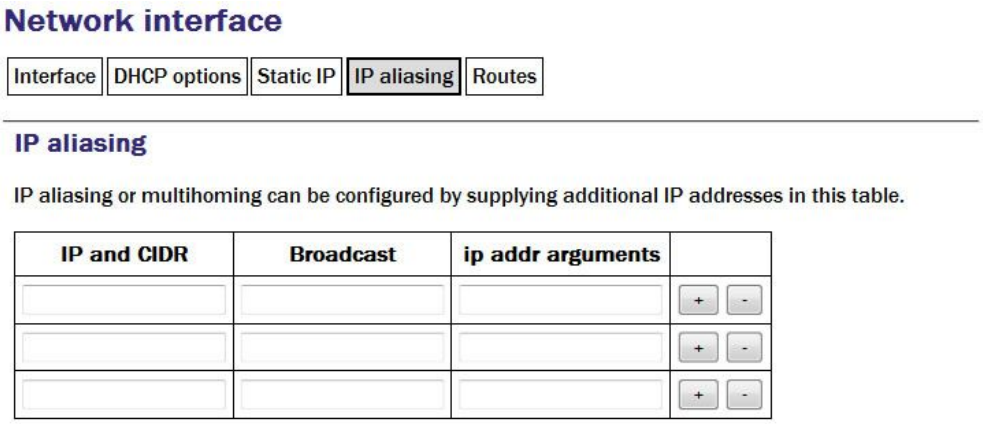

IP and CIDR: The address should be entered in CIDR format, where the address is followed by a slash and then the number of bits defining the netmask, e.g. `192.168.0.1/24` for IPV4 or `2001:db8::/32` for IPV6.

Broadcast: Enter the broadcast address to be associated with this address on the interface.

IP addr arguments: This field can be used to tune the operation of the network interface. For settings that can be used here, please see http://linux.die.net/man/8/ip.

7.1.2.5 Routes

Routes are used to add extra network and host routes to allow access to networks other than those connected via the default router specified earlier,

or to force packets to traverse a particular route despite the default router setting. By default, the table displays three blank rows but, should you need more, complete the first three and submit the form: it will be redrawn with extra blank rows. Alternatively, clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

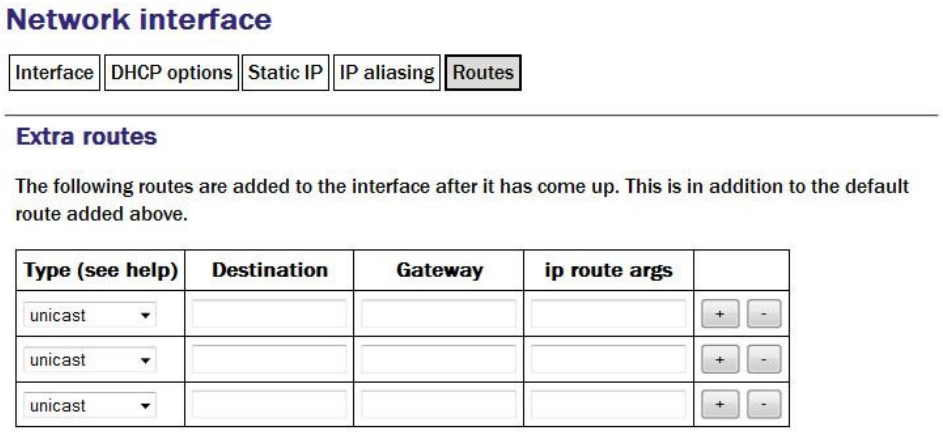

Type: This drop-down menu offers the following choices:

- **unicast** - This is the normal setting for a host or network route. The route entry describes real paths to the destinations specified in the **Destination** column.
- **unreachable** - these destinations are unreachable. Packets are discarded and the ICMP message *host unreachable* is generated. An *EHOSTUNREACH* error may appear in `/var/log/messages`.
- **blackhole** - these destinations are unreachable. Packets are discarded silently. An *EINVAL* error may appear in `/var/log/messages`.
- **prohibit** - these destinations are unreachable. Packets are discarded and the ICMP message *communication administratively prohibited* is generated. An *EACCES* error may appear in `/var/log/messages`.
- **local** - the destinations are assigned to this host. The packets are looped back and delivered locally.
- **broadcast** - the destinations are broadcast addresses. The packets are sent as link broadcasts.

Destination: The host or network to which this route offers access should be entered here in CIDR format, where the address is followed by a slash and then the number of bits defining the netmask, e.g. **`192.168.0.1/24`** for IPV4 or **`2001:db8:/32`** for IPV6.

Gateway: The IP address of the host which serves as the gateway to the specified destination.

IP route args: Used to modify the invocation of the associated `ip route add` command in order to, e.g., set the *route metric*. The options that can be set here are mostly highly technical and should rarely be required. Please see http://linux.die.net/man/8/ip for more information.

7.2 Wireless Networking

CMG-EAM and CMG-DAS units equipped with wireless networking hardware can function as clients for Wireless Access Points or participate in ad-hoc wireless networks.

To configure a wireless network interface from the web interface, select:

Configuration → Networking → Interfaces

or

Configuration → All options → Networking

The following screen will be displayed:

Now select 'wlan0 – Wireless interface wlan0':

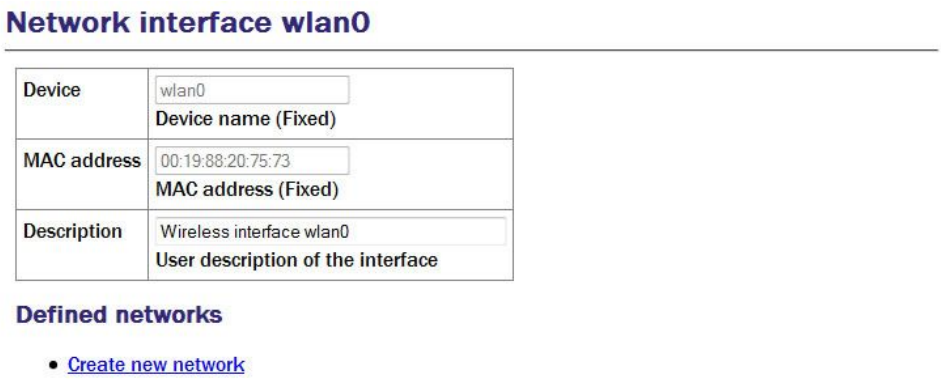

You can change the description of the wireless network if you wish.

To configure a wireless network interface from the command line, start `gconfig` and select “Networking” from the top level menu:

Existing defined networks can be reconfigured by clicking on the link and new networks can be defined by clicking on the “Create new network” link.

7.2.1 Configurable parameters in simple mode

The configurable parameters for wireless networking have three tabs in simple mode: Network, Access Point and IP Address.

7.2.1.1 Network

Network description: Enter a descriptive name of the network.

7.2.1.2 Access Point

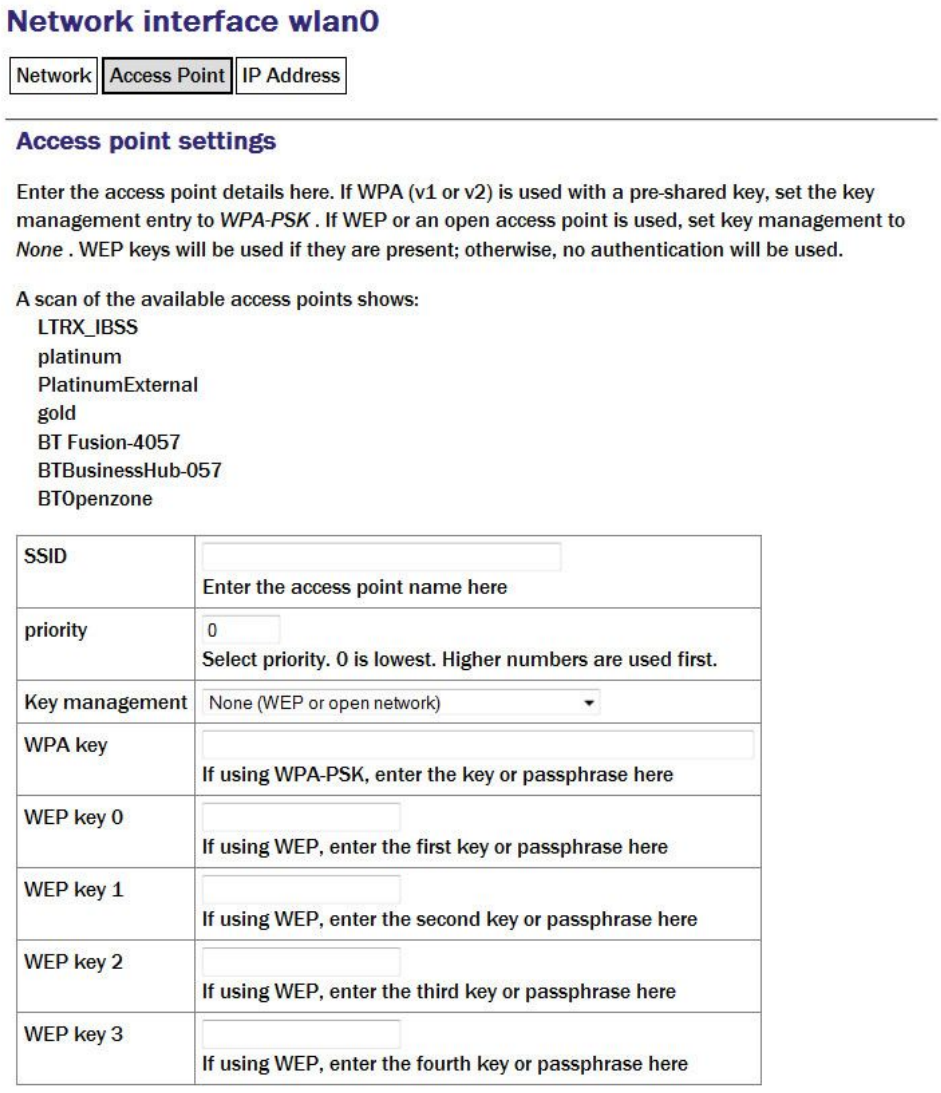

SSID: Specifies the Service Set Identifier (network name) for the desired wireless network.

Priority: Used to decide which network is chosen for connection when several configured networks are detected: connections to networks with higher priority numbers are attempted first.

Key management: Drop-down menu used to select which authentication strategy is used for this network. If WPA-PSK (WiFi Protected Access, pre-shared keys) is chosen, the following fields allow the operator to specify one or more access keys (passphrases).

7.2.1.3 IP address

The IP address screen allows the selection of DHCP (Dynamic Host Configuration Protocol) or static addressing. If DHCP is chosen, a DHCP server will provide the unit with an IP address as well as the IP address of a default router (gateway) and, optionally, other parameters such as the address of a DNS server (name-server).

If static addressing is selected, the desired IP address and default router must be specified in the following two fields.

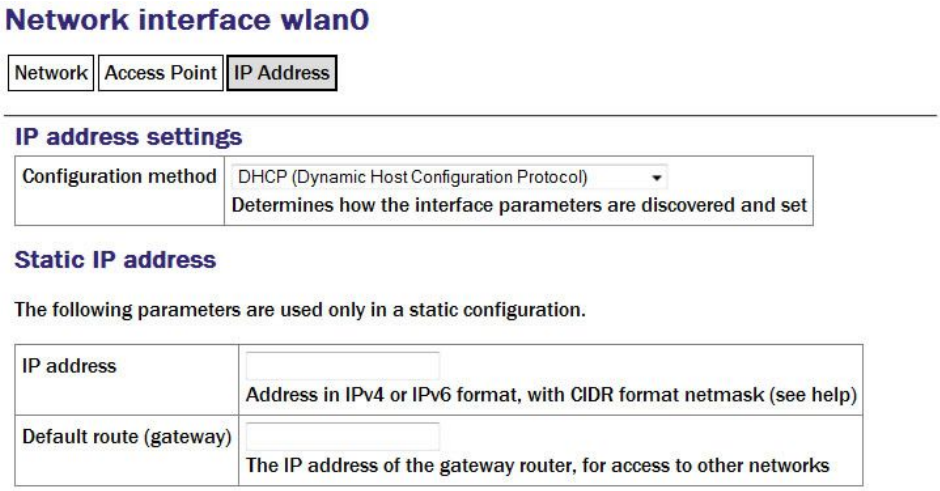

7.2.2 Configurable parameters in expert mode

When in expert mode new fields are available on the network and IP address forms and there are two new tabs: IP aliasing and Routes.

7.2.2.1 Network

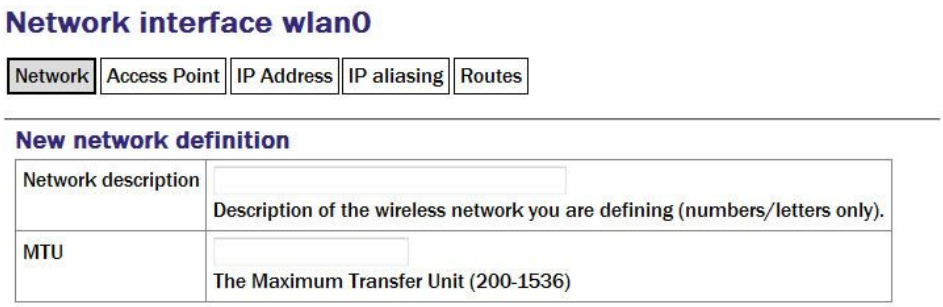

MTU: This field allows you to define a Maximum Transmission Unit (the size of the largest packet that a network protocol can transmit) for the associated network. A value which is too small will cause inefficient packetisation whereas one which is too large will cause packet fragmentation and possible disruption to communication. See the discussion in section 7.1.2 on page 72 for more information.

7.2.2.2 IP Address

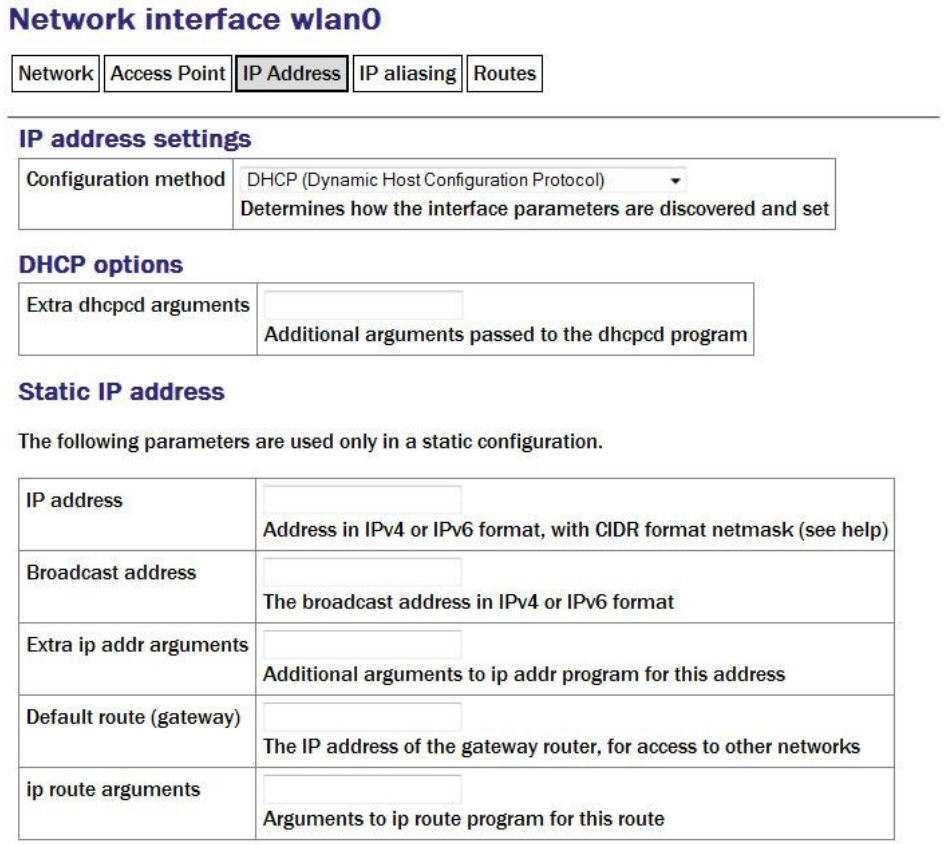

Extra dhcpcd options: Used to change the operation of the DHCP client. Please see http://man-wiki.net/index.php/8:dhcpcd for information about what options can be added here.

Extra IP addr arguments: Used to tune the operation of the network interface. A non-standard broadcast address can be specified by entering `broadcast broadcast_address`. For other settings that can be used here, please see http://linux.die.net/man/8/ip.

IP route arguments: Used to modify the invocation of the `ip route add` command in order to, e.g., set the *route metric*. The options that can be set here are mostly highly technical and should rarely be required. Please see http://linux.die.net/man/8/ip for more information.

7.2.2.3 IP aliasing

The IP aliasing table is used to add extra addresses to this interface, a practice known as multi-homing. By default, the table displays three blank rows but, should you need more, complete the first three and submit the form: it will be redrawn with extra blank rows. Alternatively, clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

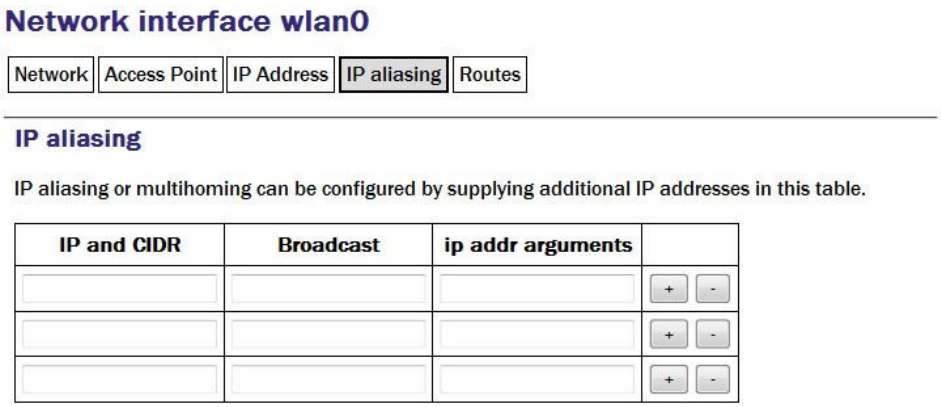

IP and CIDR: The address should be entered in CIDR format, where the address is followed by a slash and then the number of bits defining the netmask, e.g. `192.168.0.1/24` for IPV4 or `2001:db8::/32` for IPV6.

Broadcast: Enter the broadcast address to be associated with this address on the interface.

IP addr arguments: This field can be used to tune the operation of the network interface. For settings that can be used here, please see http://linux.die.net/man/8/ip.

7.2.2.4 Extra routes

Routes are used to add extra network and host routes to allow access to networks other than those connected via the default router specified earlier, or to force packets to traverse a particular route despite the default router setting. By default, the table displays three blank rows but, should you need more, complete the first three and submit the form: it will be redrawn with extra blank rows. Alternatively, clicking the + button on any row will open

a new row. In the same way, rows can be deleted by clicking the corresponding - button.

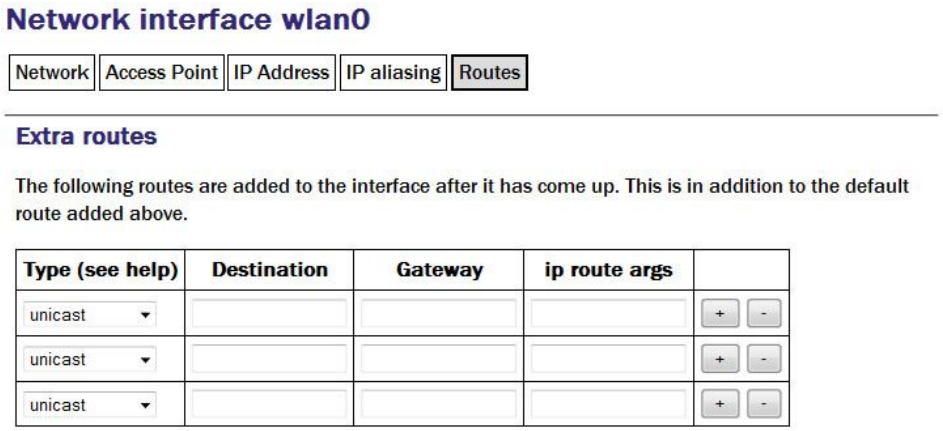

Type: The drop-down menu offers the following choices:

- **unicast:** Normal setting for a host or network route. The route entry describes real paths to the destinations specified in the **Destination** column.
- **Unreachable:** Destinations are unreachable. Packets are discarded and the ICMP message *host unreachable* is generated. An *EHOSTUNREACH* error may appear in `/var/log/messages`.
- **Blackhole:** Destinations are unreachable. Packets are discarded silently. An *EINVAL* error may appear in `/var/log/messages`.
- **Prohibit:** Destinations are unreachable. Packets are discarded and the ICMP message *communication administratively prohibited* is generated. An *EACCES* error may appear in `/var/log/messages`.
- **Local:** Destinations are assigned to this host. The packets are looped back and delivered locally.
- **Broadcast:** The broadcast addresses. The packets are sent as link broadcasts.

Destination: The host or network to which this route offers access should be entered here in CIDR format, where the address is followed by a slash and then the number of bits defining the netmask, e.g. **`192.168.0.1/24`** for IPV4 or **`2001:db8:/32`** for IPV6.

Gateway: Enter the IP address of the host which serves as the gateway to the specified destination.

ip route args: This field field can be used to modify the invocation of the associated `ip route add` command in order to, e.g., set the *route metric*. The options that can be set here are mostly highly technical and should rarely be required. Please see http://linux.die.net/man/8/ip for more information.

7.3 Virtual network (VLAN) interfaces

Platinum firmware supports the use of Virtual Local Area Networks (VLANs) to partition network traffic on the same physical subnet. Virtual interfaces can be created and assigned to a particular VLAN tag (ID) and a particular physical interface. A full discussion of VLANs is beyond the scope of this document.

To configure a virtual network interface from the web interface, select:

Configuration → Networking → Interfaces

or

Configuration → All options → Networking

The following screen is displayed:

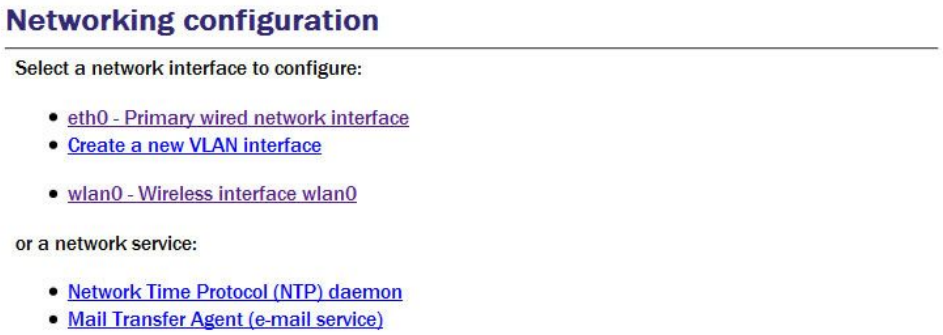

To configure a VLAN interface from the command line, start gconfig and select “Networking” from the top level menu.

Now select 'Create a new VLAN interface'

7.3.1 Configurable parameters in simple mode

The configurable parameters for virtual network interfaces have two tabbed pages in simple mode: Interface and Static IP.

7.3.1.1 VLAN Interface

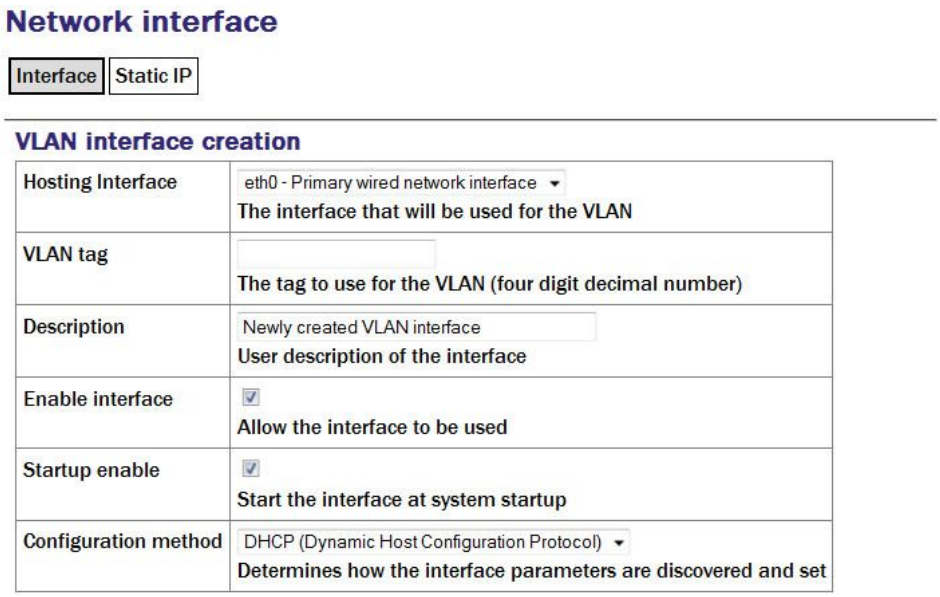

Hosting Interface drop-down menu is populated with a list of the physical interfaces present on the system. Select the physical interface over which you wish this virtual interface's traffic to flow.

VLAN tag text field should be populated with the required VLAN tag. This identifies packets sent over this interface as belonging to the specified VLAN.

Description text field can be edited to provide a more useful name for the interface.

Enable interface check-box can be ticked in order to enable the interface or cleared in order to disable it. No other configuration settings are changed when the interface is disabled, allowing use of the interface to be suspended without deleting the configuration.

Startup enable check-box controls whether the interface is enabled automatically when the unit boots.

Configuration method drop-down menu offers the following choices:

- **Static** - The interface will take its address and routing parameters from values entered by the operator.
- **DHCP (Dynamic Host Configuration Protocol)** - The interface will attempt to obtain its address and routing parameters from a DHCP server with a matching VLAN tag.
- **Powered off** - The interface will not be used and the interface chip is disabled, reducing the total power consumption by around 200mW.

7.3.1.2 VLAN Static IP

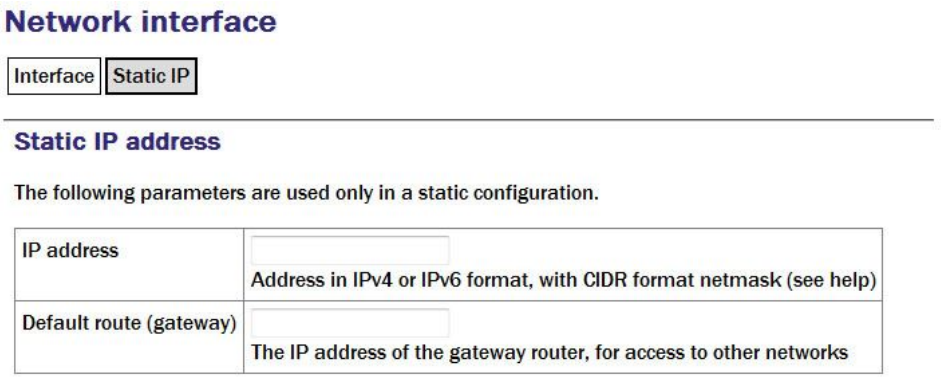

If DHCP is not being used, the **IP address** text field should be populated with the required address in CIDR format, where the address is followed by a slash and then the number of bits defining the netmask, e.g. **`192.168.0.1/24`** for IPV4 or **`2001:db8:/32`** for IPV6.

For more information about CIDR, please refer to http://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing.

The **Default route (gateway)** field should be populated with the IP address of the default router for this VLAN. If more complicated routing configurations are required, these can be entered in expert mode.

7.3.2 Configurable parameters in expert mode

A set of additional parameters are available when in expert mode. These are identical to the additional parameters on the physical interface configuration screen, as described in section 7.1.2 on page 72 and are not discussed further here.

7.4 Network Time Protocol (NTP)

The Network Time Protocol (NTP) is a method of synchronising the clocks of computer systems over networks, including those with variable latency, such as packet-switched networks. Platinum firmware include a fully-featured NTP implementation, which can be used to keep the system clock synchronised to external time sources, such as Internet NTP servers, connected digitisers and connected GPS receivers.

Most acquisition modules include a battery-backed real-time clock (RTC) module which can retain system time with tolerable accuracy during periods of power loss. CMG-DCMs will revert to January the 1st, 1970 after each power-cycle.

The system clock is used to provide time-stamps for log-file messages and can also be used to generate NMEA and PPS signals, emulating a GPS receiver, in

order to synchronise external equipment, such as a CMG-DM24 digitiser module. This technique is described in section 9.6 on page 124.

The NTP subsystem displays its status in a panel of the “System status” display of the web interface. This panel includes the current system date and time, the lock status, estimated error and current clock source.

Platinum systems support two distinct timing modes, NTP and Manual. NTP is recommended for nearly all installations. Manual mode requires that the system clock be set from the command-line at least once per day and is, by its nature, considerably less accurate than NTP mode.

To configure NTP from the web interface, select:

Configuration → Data handling → Timing

or

Configuration → All options → Networking → Network Time Protocol daemon

To configure NTP from the command line, start gconfig, select “Networking” from the top level menu and then select “ Network Time Protocol (NTP) daemon”.

7.4.1 Configurable parameters

The configurable parameters for timing are presented in three tabs: General, NTP and Manual.

7.4.1.1 General tab

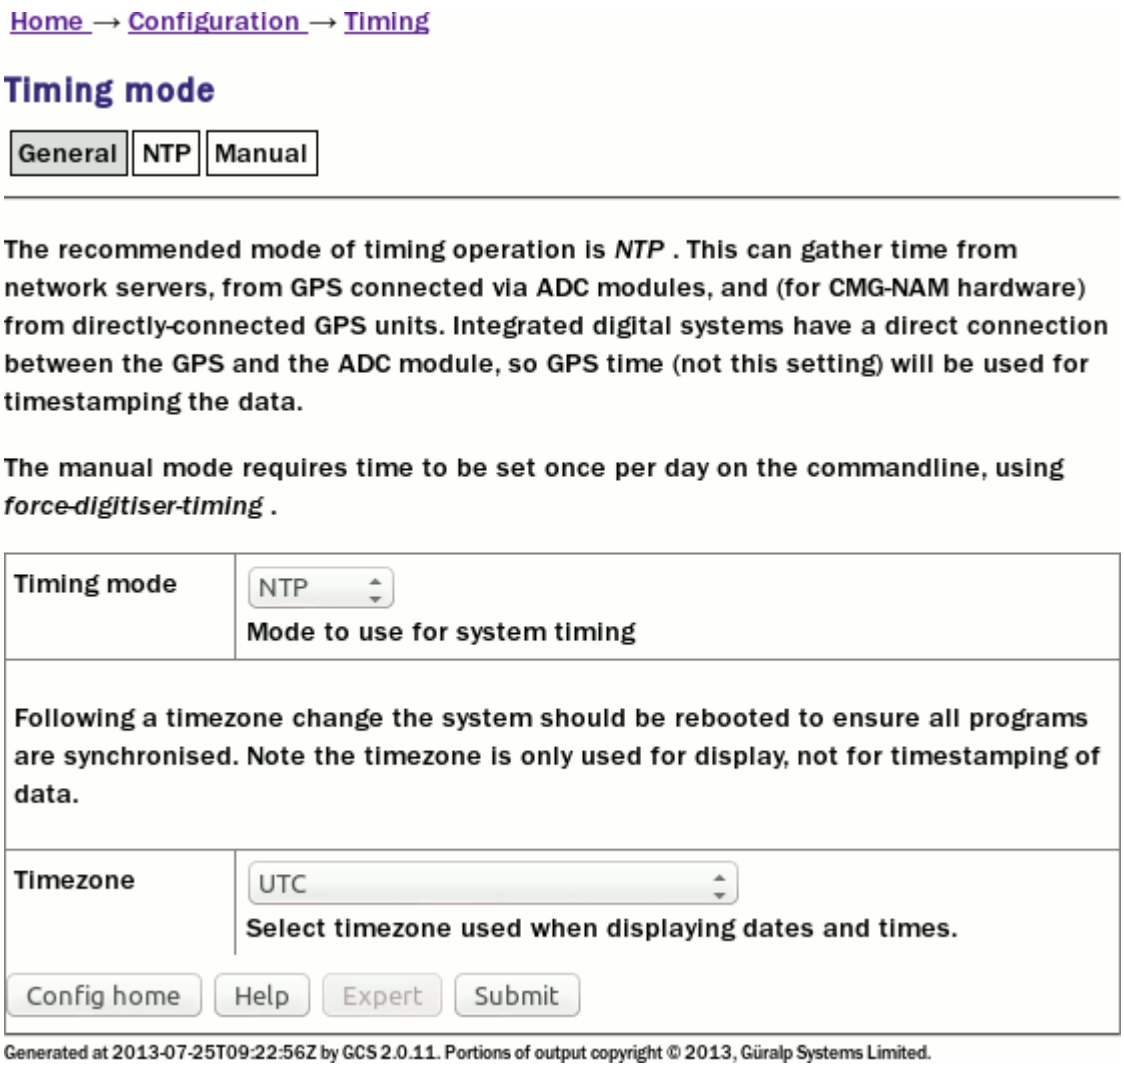

The first drop-down menu on the "General" tab, "Timing mode", allows selection between NTP and Manual. NTP is recommended for nearly all applications.

If NTP is selected, further configuration should be carried out on the NTP tab. Likewise, if Manual is selected, further configuration should be carried out on the Manual tab.

The second drop-down menu, "Timezone", allows selection of the time zone used when dates and times are displayed in web pages, in log-files and on the command-line. This is a system-wide setting that affects all users.

Note: Changing the time zone of the acquisition unit does not affect the timestamps of recorded data in any way. The ADC module has its own clock which always runs in UTC.

7.4.1.2 NTP tab

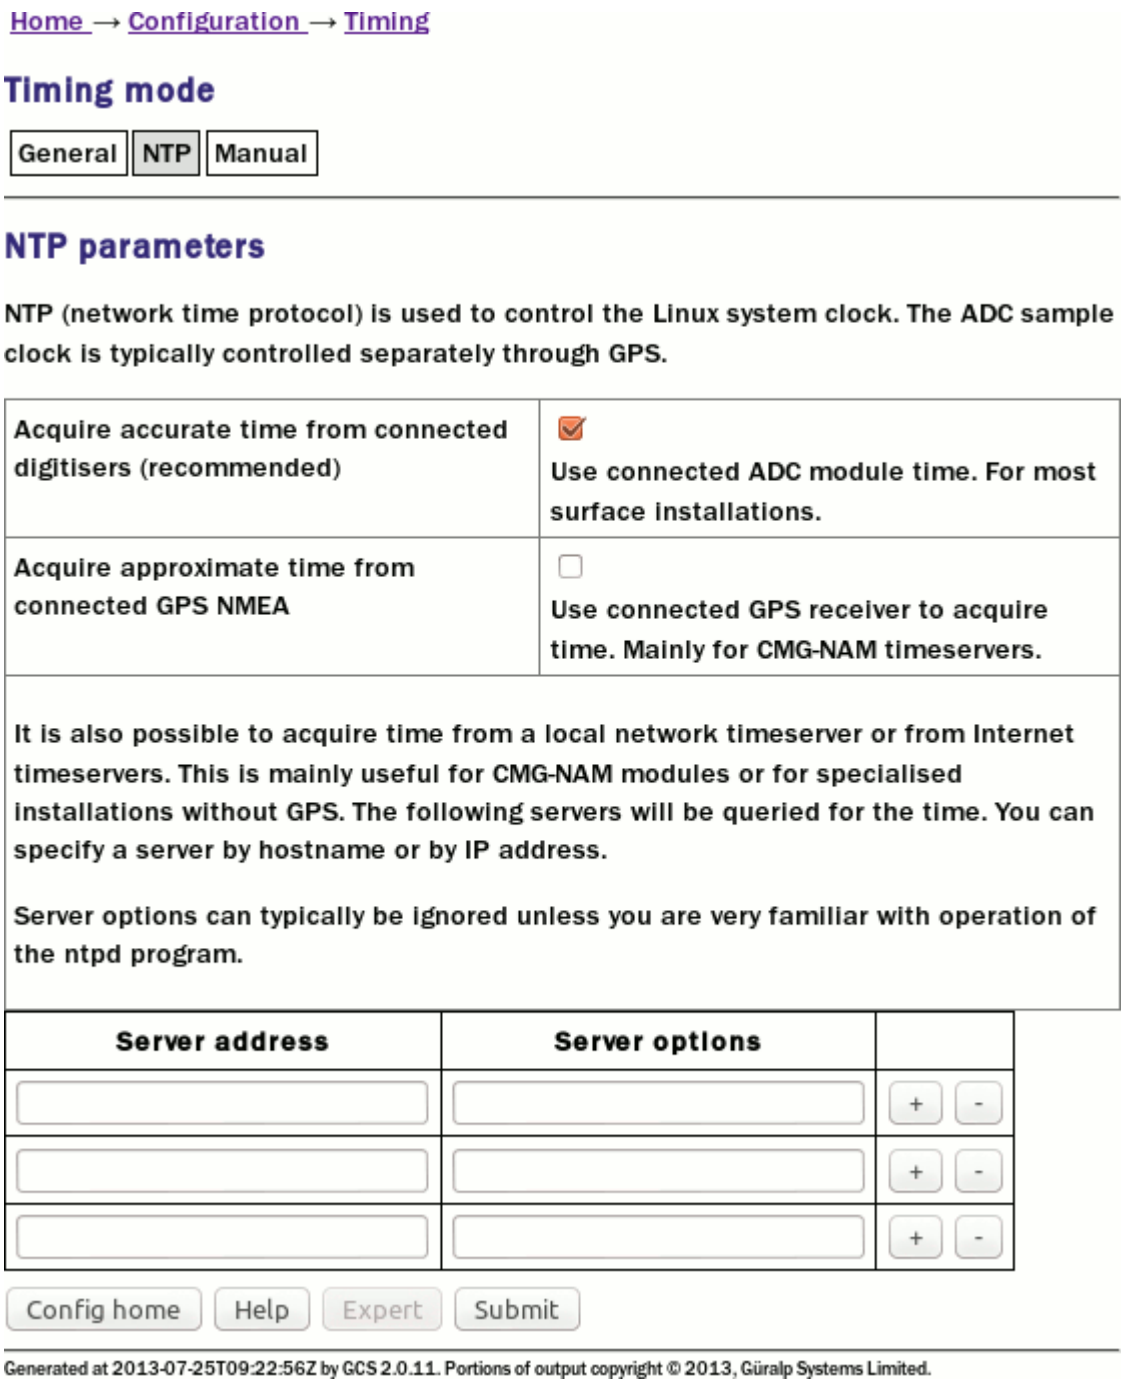

The NTP tab contains the following configurable options:

Acquire time from connected digitisers: This check-box tells the system that one or more attached digitisers are to be considered as accurate clock sources. For this to work, the digitiser must produce "RTSTATUS" packets. CMG-CD24 digitisers and digital sensors incorporating them, such as the CMG-6TD, will do this unconditionally when running firmware version 279 or later. CMG_DM24 digitisers can have these packets enabled or disabled via

software (using the command `RTSTATUS ENABLE`). They are automatically enabled if the digitiser is ever configured via the interface in Platinum.

Acquire approximate time from connected GPS: This check-box tells the system that a GPS receiver has been attached to one of the serial ports and is to be used as a clock source. The serial port used must be configured with a "Port function" of "NMEA in. Receive GPS data for NTP" and, when used with Güralp supplied GPS receivers, must be set to 4,800 baud operation, as described in 9.6 on page 124. No further configuration is required.

Note: This setting should only be used on CMG-NAMs and CMG-NAM64s.

Server address: Allows a number of Internet or network-accessible NTP servers to be listed for use as clock sources. You can specify these servers by either IP address or hostname. If names are used, they must either be listed in the local hosts file (`/etc/hosts`) or resolvable via the Domain Name Service (DNS). Clicking the + button on any row will open a new, blank row. In the same way, rows can be deleted by clicking the corresponding - button.

7.4.2 Configurable parameters in expert mode

The only difference between standard and expert mode on this screen is the addition of a **Server options** column to the NTP servers table.

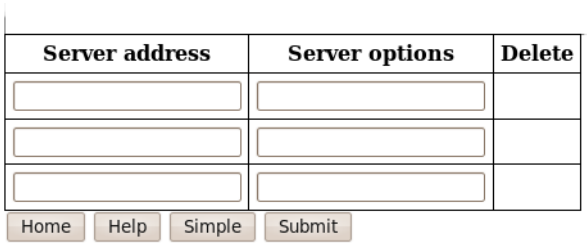

This text field can be used to provide additional control over how the NTP daemon uses the server. The options are described in the standard `ntp.conf` manual page, available on-line at http://linux.die.net/man/5/ntp.conf.

7.5 Email configuration

Platinum firmware is capable of sending system alerts via email over the Internet or over a local area network.

This feature is currently unused but is provided for future expansion.

To configure email from the web interface, select:

Configuration → Networking → Mail

or

Configuration → All options → Networking → Mail Transfer Agent (e-mail service)

To configure email from the command line, start gconfig, select Networking from the top level menu and then select Mail Transfer Agent (e-mail service).

7.5.1 Configurable parameters

The configurable parameters for email are contained on a single form.

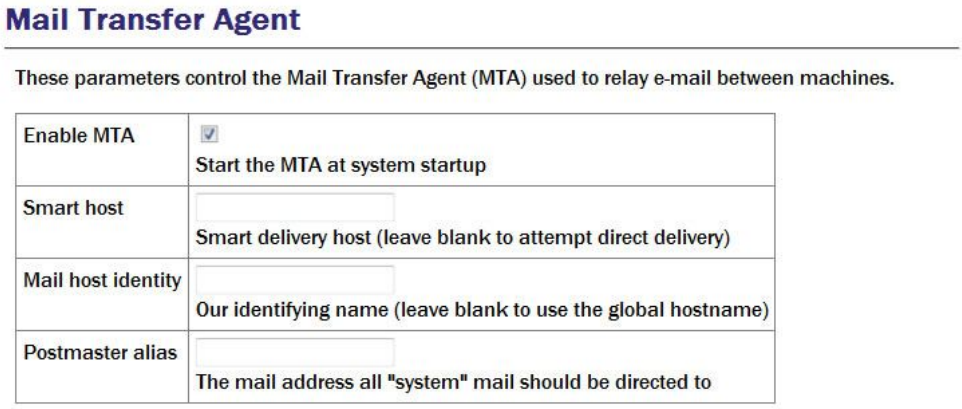

Enable MTA: Used to control whether the mail transfer agent is started automatically at boot time. If this check-box is left clear, the MTA can still be started manually from the services menu (see section 14.3.5 on page 234).

Smart host: Most email configurations use a “smart host” to route mail. This can greatly simplify the administration: only one node on a network, the smart host, needs to be configured to know about any intricacies of the system and all other machines need only know the location of the smart host. If the Smart host text field is populated with the name or address of of such a host, all mail is sent directly to that host for further routing. If this field is left blank, the MTA will attempt to use DNS to discover the mail host(s) for any given address and then deliver mail directly.

Mail host identity: Specifies the hostname from which outgoing emails should appear to originate. If this field is left blank, the real hostname is used.

Postmaster alias: Specifies the address to which all internally generated mail should be sent: This should be set to the email address of the acquisition module's administrator.

7.6 Configuring the SSH Server

The acquisition module has an ssh server running on its Ethernet port which allows remote terminal access.

The ssh server, sshd, can not currently be configured using `gconfig` although it can be configured via the web interface. If web access is unavailable, it is possible to configure sshd from the command line by directly editing the configuration files.

7.6.1 Configuring sshd via the web interface

To configure the SSH server from the web interface, select:

Configuration → Networking → SSH server

The screen is not reproduced in this document as it is particularly large, due to the amount of explanatory text. Each option is, however, discussed below.

The version of `sshd` installed (openSSH) supports both version 1 and version 2 of the ssh protocol. Version 1 has some well-known weaknesses and should be avoided if at all possible, but some commercially available systems still do not support v2, so v1 is supported here for compatibility. The **Enable SSH Protocol v1** check-box should be cleared unless your ssh client cannot support v2 or cannot be upgraded to support it. Click the **Change server options** button to commit this change.

If you want to download the ssh server's public key to allow the connecting host to check and verify the CMG-EAM's identity, use the relevant **Download server public key** button – there is one each for protocol versions 1 and 2. There is also the capability to command the CMG-EAM to create a new private/public key pair from this screen.

To configure password-less login to the CMG-EAM, you can upload the public key of the connecting machine to the CMG-EAM using the **New client key** section. Browse the connecting host's file system for the key file (usually named `id_dsa.pub`) and upload it here. This will allow password-less root access to the system from that machine.

Uploaded client keys are displayed in the **Authorised client keys** section. Any existing authorised keys can be removed: Select the check-box next to the key you wish to remove and click **Remove selected keys**.

Note: Password-less login via ssh v2 is, perhaps counter-intuitively, the most secure way to access your acquisition module. There is a useful discussion of the ssh protocol and full details of its usage at the site `http://tinyurl.com/whyssh`

Note: Systems are shipped with a pre-authorised client key to which Güralp Systems' support staff have the matching key. This allows us to access your unit and reset the root password should you forget it. You are free to delete this key if you wish.

There is a second (and significantly more complicated) way of resetting the root password if you have physical access to the system. Please contact support if you find yourself in this situation.

7.6.2 Configuring sshd from the command line

This is a complex issue and use of the web interface is strongly encouraged unless you are familiar with Linux text editors, configuration files and sshd itself. The configuration file is located at `/etc/ssh/sshd_config` and its syntax and semantics are described at http://man-wiki.net/index.php/5:sshd_config. More detailed discussion is beyond the scope of this document.

7.7 Working with PPP

PPP, or Point-to-Point Protocol, is a data link protocol that can carry IP packets over a serial link between two networking nodes. It can provide connection authentication, transmission encryption privacy, and compression. Platinum firmware includes an implementation of PPP which can be used to provide network links between sites or to connect to an Internet Service Provider (ISP). A number of "chat scripts" are provided, allowing connection negotiation and establishment over PSTN, GPRS and satellite modems.

7.7.1 Setting up a PPP Connection

To configure a PPP connection, you will need a user ID and authentication code (the PAP secret) as required by the remote server. In addition, a dial up (chat) script specific to the service you are using must be created. If one does not already exist for your service, please contact Güralp support.

To configure a PPP connection from the web interface, select:

Configuration → Data handling → Serial ports → Port...

or

Configuration → All options → Serial ports → Port...

To configure a PPP connection from the command line, start gconfig, select Serial ports from the top level menu and and select the port to which the modem is connected. Change the function of the port to PPP. PPP network connection, with the correct baud rate for the modem. Click Submit to save

these changes. Go back to the configuration of the serial port and click on PPP network configuration.

7.7.2 Configurable parameters

The configurable parameters for each port on the PPP connection screen have three tabs: General, IP and Authentication.

7.7.2.1 General

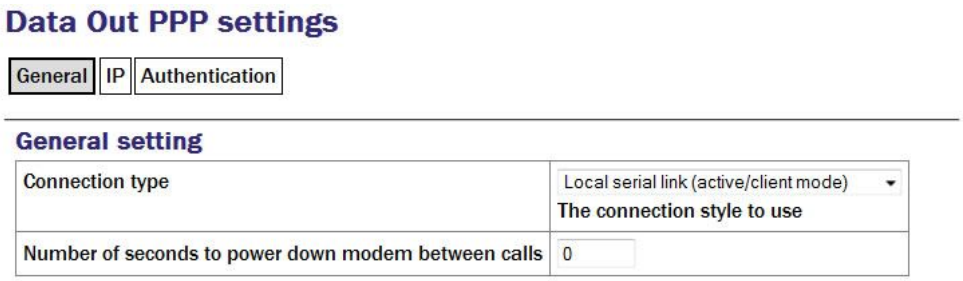

Connection type has the following options:

- Local serial link (active/client mode)
- Local serial link (passive/server mode)
- GPRS connection via Vodafone
- GPRS connection via T-Mobile

The choices available from this menu reflect the chat scripts installed on your system. If you wish to use a satellite modem or GPRS with an ISP other than those listed, please contact Güralp Systems Ltd technical support.

Number of seconds to power down modem between calls: Set the desired power-down time for the modem, if required. This can be useful with modems that are prone to entering undefined states and ensures a clean reset before each new connection.

7.7.2.2 IP

The IP addresses and routing section handles the network configuration. In active/client mode or when connecting to an ISP, the remote PPP daemon will set these parameters, in which case this section can be left blank.

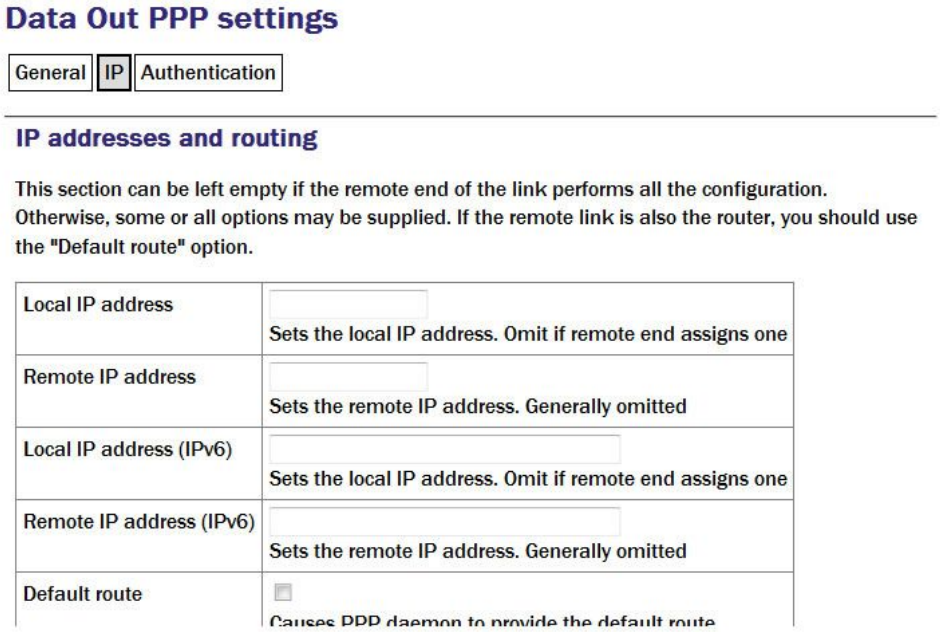

If you are using PPP between, say, two acquisition modules, one should be designated the client and the other the server. The "IP addresses and routing" parameters for the client should be left blank and those for the server completed as follows:

IPv4 networking: the **Local IP address** field should be populated with an IP address from an otherwise unused reserved class C network address range, such as 192.168.123.1 (with no CIDR postfix) and the **Remote IP address** field should be populated with an IP address from the same network, such as 192.168.123.2 - this address is provided to the client at connection initiation.

IPv6 networking: the **Local IP address IPv6** and **Remote IP address IPv6** fields should be populated.

Default route: If selected the PPP daemon will modify the routing table on successful connection, setting the remote end of the PPP link as the default gateway.

7.7.2.3 Authentication

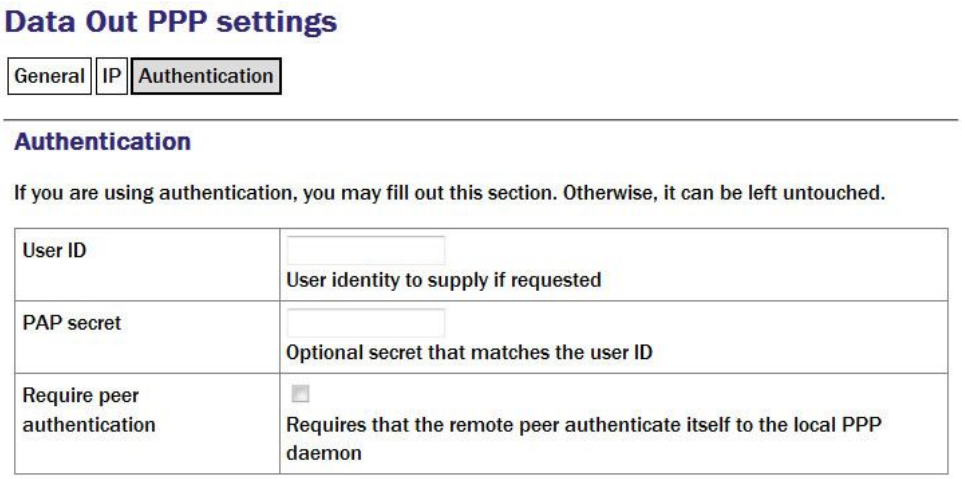

User ID / PAP secret: Use those given to you by your service provider in the appropriate fields.

The standard Linux commands `ppp-on`, `ppp-off`, `ip`, `ping`, and `traceroute` are available from the command line for use in controlling and testing PPP connections but it is also possible to configure a "watchdog" service to monitor a PPP connection and automatically restart it should it fail. This is described in the next section.

7.7.3 Monitoring a PPP connection

PPP connections can be monitored and, should they fail for any reason, automatically restarted.

To configure this PPP connection monitor from the web interface, select:

Configuration → Services → Network → pppd-watchdog – PPP link watchdog

or

Configuration → All options → System services → Network → pppd-watchdog → PPP link watchdog

To configure a PPP connection monitor from the command line, start gconfig, select Services from the top level menu.

You must create a separate watchdogs for each PPP connection if you are running multiple PPP instances. This screen allows you to select any of the existing watchdog services for re-configuration or to create a new watchdog service.

7.7.4 Configurable parameters in simple mode

The configurable parameters for the PPP daemon watchdog are contained in a single form:

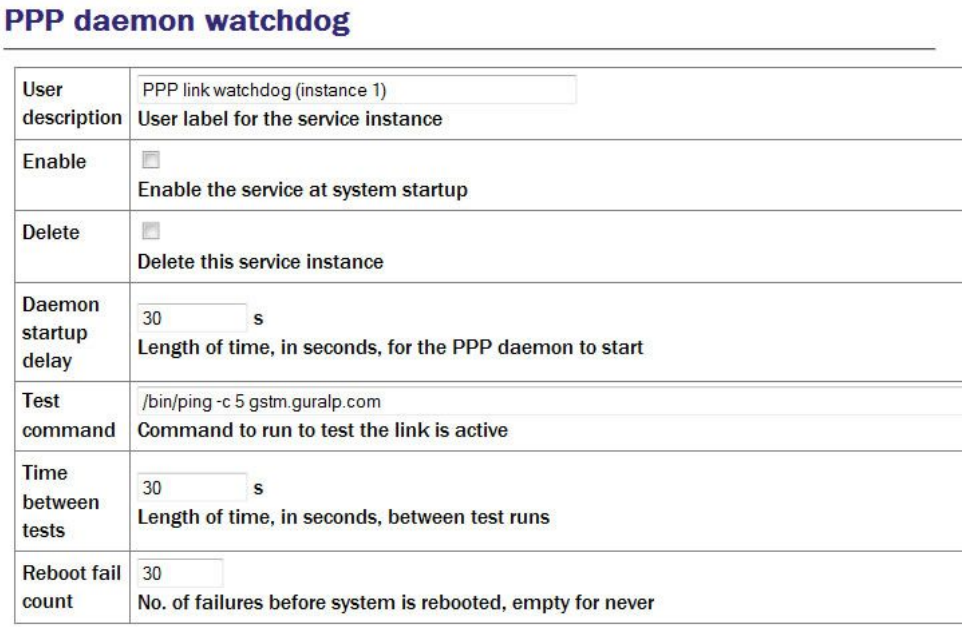

User description: If you are configuring a number of watchdogs give each of them memorable names.

Enable: Should normally be ticked but can be left clear if you only want to use the associated PPP connection occasionally.

Delete: Deletes the instance when the form is submitted.

Daemon startup delay: If the PPP connection relies on a modem link for its transport, there may be a significant delay between instructing the PPP link to start and the connection being established. So that the watchdog does not falsely detect a failed link during this period, it can be instructed to sit idle for a number of seconds before it begins to test the link. The length of the required delay should be entered into the field.

Test command: Once the start-up delay time has elapsed, the watchdog periodically tests the connection. To ensure that there is a valid end-to-end connection where, for example, a multi-hop link is in use, the exact method of testing is user-configurable: any valid command can be entered into the field and its exit status is taken to represent the link status (zero for link up, non-zero for down). The most common method used is to use the `ping` command to verify ICMP connectivity to the ultimate remote host, but you are free to use the command or script of your choice here, so long as it returns a non-zero exit status on link failure.

Time between tests: Determines how often the configured test is applied. It can be set high to conserve bandwidth or set low to improve failure response times. It can also be used to keep a sparsely-used link alive where a "disconnect-on-inactivity" feature would otherwise interrupt it.

Reboot fail count: If the link test fails repeatedly, the acquisition module is rebooted. Enter the number of failed tests before reboot.

The ppp watchdog service can be started, stopped and restarted using the **Services** page under the **Control** menu. See section 14.3.5 on page 234.

7.7.5 Configurable parameters in expert mode

A number of additional parameters are displayed when in expert mode:

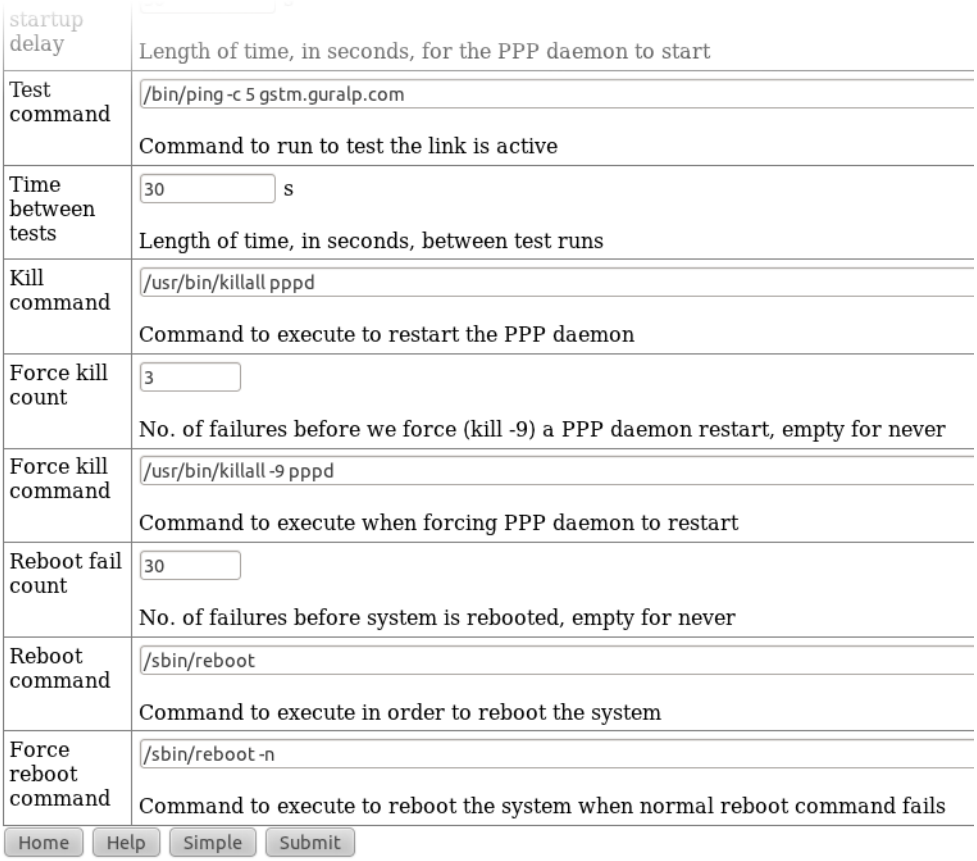

Kill command: The command which the watchdog should use to attempt to restart the PPP daemon if it is determined that the link has failed.

Force kill count: If the watchdog determines that the Kill command is not working after this number of attempts, it will try the Force kill command.

Force kill command: The command which the watchdog should use to attempt to restart the PPP daemon if it is determined that the link has failed and several attempts to restart it with the Kill command have failed to have a positive effect.

Reboot command: The command to use to reboot the system if repeated Force kills have been ineffective.

Force reboot command: The command to use to reboot the system if the normal reboot command fails to shut down the system.

7.8 Configuring TCP to serial converters

The acquisition module can act as a TCP to serial converter, effectively transporting data between one (or more) of its serial ports and a TCP connection. There are two different modes of operation, as detailed below. Any number of serial ports can be configured as TCP converters, as long as the TCP port numbers do not clash.

In “Simple server” mode, the acquisition module listens for incoming TCP connections and, should it receive one that matches its configured rules, accepts the connection and begins copying data between the serial port and the TCP connection.

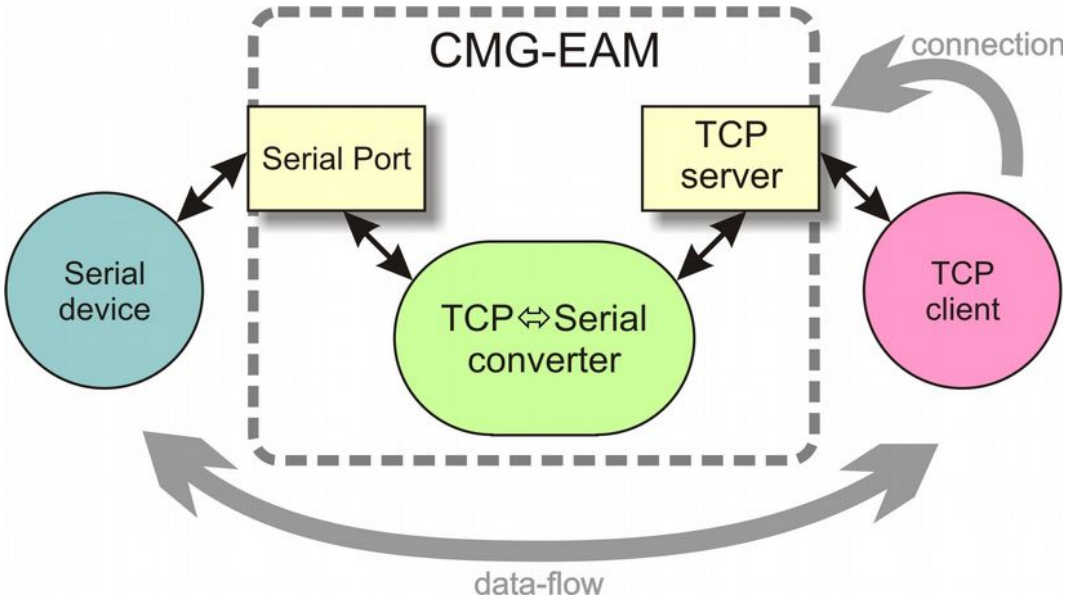

The acquisition module can be configured to only listen on particular addresses and ports, to only accept connections from certain addresses or blocks of addresses and to reject connections from certain addresses or blocks of addresses.

In “Simple client mode”, the acquisition module will connect to an external TCP server on a particular address and port and then copy data bidirectionally between the serial port and the network port.

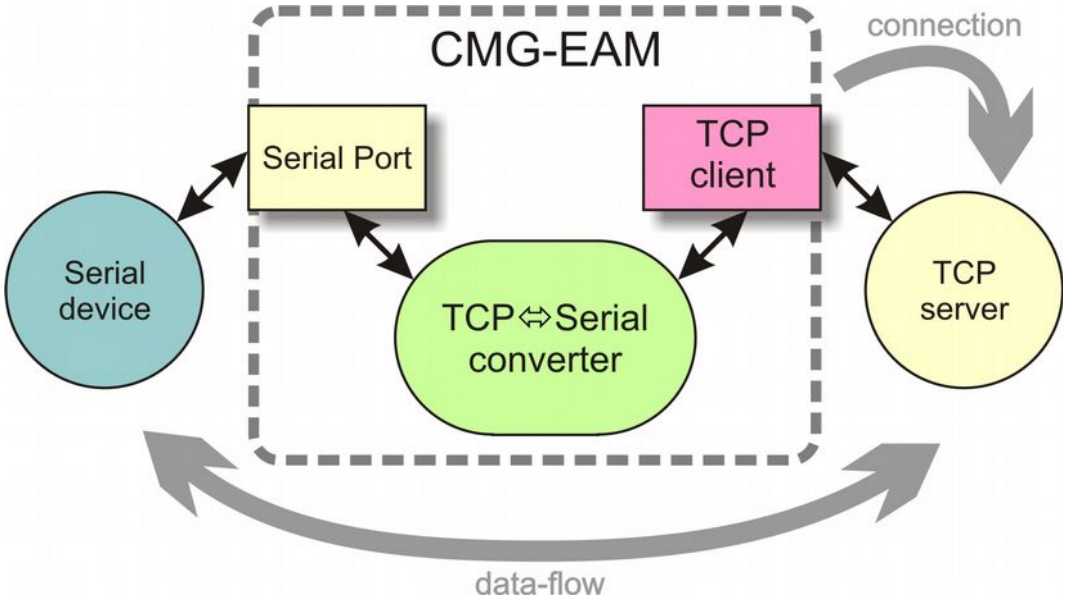

To configure the converter from the web interface, select:

Configuration → Serial Ports

then choose the required port. Set the function to “tcp serial converter”, select the baud rate, and save the settings. You can then use the TCP serial converter settings button at the bottom of the page to configure the converter.

The converter's configuration page allows you to choose the mode at the top (“Operation mode”). The other options on the page are only required in certain modes; see below for which modes require which options.

7.8.1 Simple server mode

In Simple server mode, the converter opens the serial port and creates a TCP server socket. Whenever a client connects to the socket, the converter reads raw data from the serial port and writes it to the client, and reads raw data from the client and writes it to the serial port. The serial port hardware control lines cannot be read or altered in this mode.

Simple server mode has two relevant options: the list of addresses to listen to, and an optional list of addresses to filter. The server can listen on multiple simultaneous local ports and addresses (although only one client can be active at a time).

The “Bind host” option is usually left blank. If specified, it is the name or IP address on which this server socket will listen. For example, if you specify “localhost” here, then this socket will only listen for incoming connections on the loopback address, and not on the external Ethernet port. Leave it blank to listen to all addresses.

The “Bind service” option must be specified. It is the TCP port number (1-65535) or service name (such as “tcpserial”) for the socket. This can be

anything you choose, although we recommend that you use the names `tcpserial`, `tcpserial1`, `tcpserial2` and so on through to `tcpserial15`, which are pre-defined to correspond to port numbers 10002, 10003, through to 10017.

The mapping from port names to port numbers is configured by the conventional Linux file `/etc/services` which can be edited from the command line if required.

If desired, you can configure a list of addresses from which to accept connections. If no addresses are configured, then all incoming addresses will be accepted. Otherwise, connections will only be accepted if they match an entry in the table with its Reject box unticked. Entries are matched in order; as soon as a match is made, the connection is accepted or rejected, and no further processing is done.

The "IP addresses" fields can each specify a host name, an IP address or an IP address range (given in CIDR format). For example, to accept connections from LAN addresses, you can add the addresses:

- **`10.0.0.0/8`**
  (anything from 10.0.0.0 to 10.255.255.255);
- **`172.16.0.0/12`**
  (anything from 172.16.0.0 to 172.31.255.255);
- **`192.168.0.0/16`**
  (anything from 192.168.0.0 to 192.168.255.255); and
- **`127.0.0.1`**
  (loopback address).

7.8.2 Simple client mode

This mode of operation is similar to simple server, except that the acquisition module establishes an outgoing TCP client connection rather than listening on a socket. It writes raw data from the serial port to the remote server, and writes raw data from the remote server to the serial port. It does not support the querying or setting of the serial port hardware control lines.

In this mode, only a single option needs to be provided: the contact details for the remote server (IP address and port). The format of this option is "host,service". The host may be a hostname or an IP address. The service may be a TCP port number or a service name from `/etc/services`.

8 Digitiser Configuration

8.1 Configuring digitisers using the web interface

The configuration interface can be used to configure the digitiser module in a DAS or any serially attached Güralp digitiser, such as CMG-DM24 or CMG-CD24. The internal digitiser module in a CMG-DAS is, effectively, serially connected so both internal and external digitisers are handled identically.

To configure a digitiser using the web interface select:

Configuration → Instruments → Port A instrument....

The list alters dynamically to reflect the system's embedded and attached devices. For every digitiser detected, an entry appears which allows you to configure the digitiser.

Note: To control (as opposed to configure) the digitiser and its attached instrument (sensor locking, mass centring, etc.) see section 14.3.2 on page 224.

The information shown on this screen is retrieved from the digitiser using a sequence of background commands over a serial communications line and may, therefore, take a few seconds to display. A progress indicator is displayed during this process. It is possible to display this sequence of commands (together with the responses received from the digitiser) and this may be useful both for learning the command-line interface of the digitiser and for debugging any unexpected behaviour. To do this, select “Show full digitiser dialogue in future form submissions” from the miscellaneous section near the bottom of the configuration screen.

Note: This is a web form and, therefore, no changes will be made to the actual configuration of the digitiser until the form is submitted – i.e. until the operator scrolls to the bottom of the page and clicks Submit changes & Reboot digitiser.

8.1.1 Configurable parameters

The digitiser configuration screen is too large to be reproduced in one illustration in this manual. It is, therefore, shown here in sections.

8.1.1.1 Identity

The first section displays the digitiser's identification string and serial number, which can be edited. It also displays the digitiser's software version:

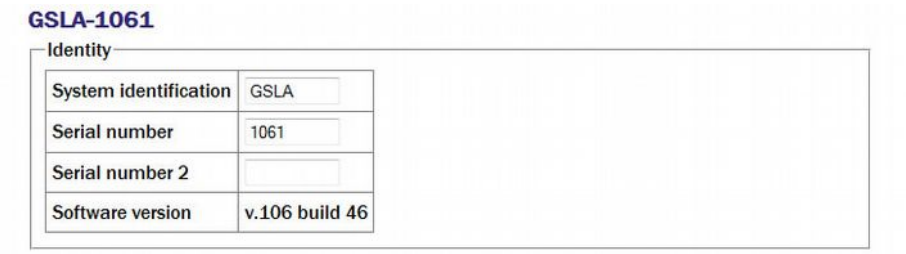

System identification / Serial number: Make any desired changes and then click on the Submit changes & Reboot digitiser button at the bottom of the screen. If the digitiser is running in dual serial mode, both serial numbers are displayed on this screen in separate rows.

8.1.1.2 Connected devices

Sensor type: Has no effect on the acquisition module's operation and acts as a memo field.

Timing source: Set to "NMEA protocol GPS" (which should be used for all GPS devices) or "None", for situations where there is no timing source.

GPS power-cycling: GPS units can be turned off to save power in battery-powered environments. In order to keep the internal clock synchronised, the GPS unit is regularly turned on for long enough to obtain an accurate time and then turned off. The GPS power-cycling drop-down allows you to select the intervals at which this happens (1, 2, 3, 4, 6, 8, 21 or 24 hours) or whether to leave the GPS constantly powered up.

Mass auto-centring: The drop-down menu can be used to configure (or disable) a digitiser function which automatically initiates a round of centring when the mass position (averaged over several minutes) drifts further than the specified distance from zero. The excursion required to trigger centring is expressed as a percentage of full scale.

8.1.1.3 InfoBlocks

InfoBlocks (information blocks) are one-kilobyte areas of storage within the digitiser which can hold arbitrary data. In some applications, such as when a digitiser is generating strong motion packets, they should hold structured information about the attached sensors. Newer digital sensors have their calibration information pre-loaded into the InfoBlock. The format of the data is given at http://www.guralp.com/information-blocks-from-guralp-digitizers/ but you can add any additional information that you find useful, such as asset tracking information, provided you don't exceed the one-kilobyte limit.

Digitisers automatically transmit the InfoBlock, if it is set, at every reboot. The block is transmitted with a special stream ID ending "`IB`". Retransmission of the InfoBlock can be requested at any time by issuing the digitiser command `SENDINFO`. If a copy of Scream receives an InfoBlock, it will scan it for calibration values and incorporate any that it finds into its calvals.txt file. This information is used for displaying streams in ground units (rather than counts) amongst other things.

There are one or two info blocks per digitiser and the display will recognise this fact. The Display device info blocks button(s) show the contents of the info block(s) and allows you to upload new data to them, should you wish. When the button is clicked, the following is displayed:

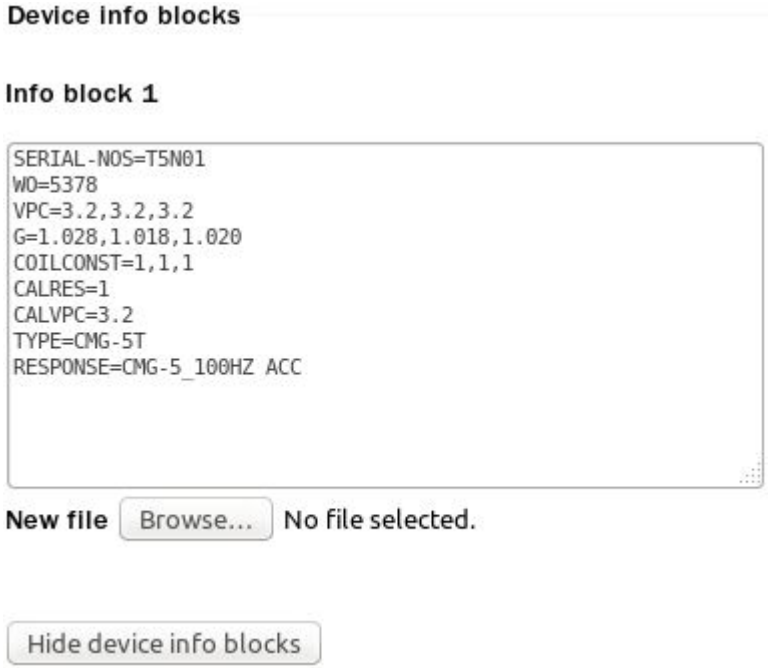

A file containing the desired InfoBlock text can be uploaded by clicking the Browse... button: the web-browser's normal upload facility is used.

Clicking the Hide device info blocks button returns the display to the previous state.

8.1.1.4 Causal filtering (low latency) and Strong motion mode

The next section of the screen controls the filtering mode of the digitiser and optionally enables strong motion calculations. Causal filtering mode (previously called "low latency" mode) is intended for use with strong motion calculations. In this mode, the last stage of the digital filtering is changed from finite impulse response (FIR, acausal) to infinite impulse response (IIR, causal) and packets are output each second at twenty samples per second in order to achieve very near real time data.

These modes cannot be selected unless the first decimator output is set to two hundred samples per second. The Enable causal filtering button is enabled only when this condition is satisfied.

CMG-DM24 digitisers can operate in "normal" mode or "strong motion" mode. In strong motion mode, additional "strong motion packets" are generated; these carry derived and resultant data.

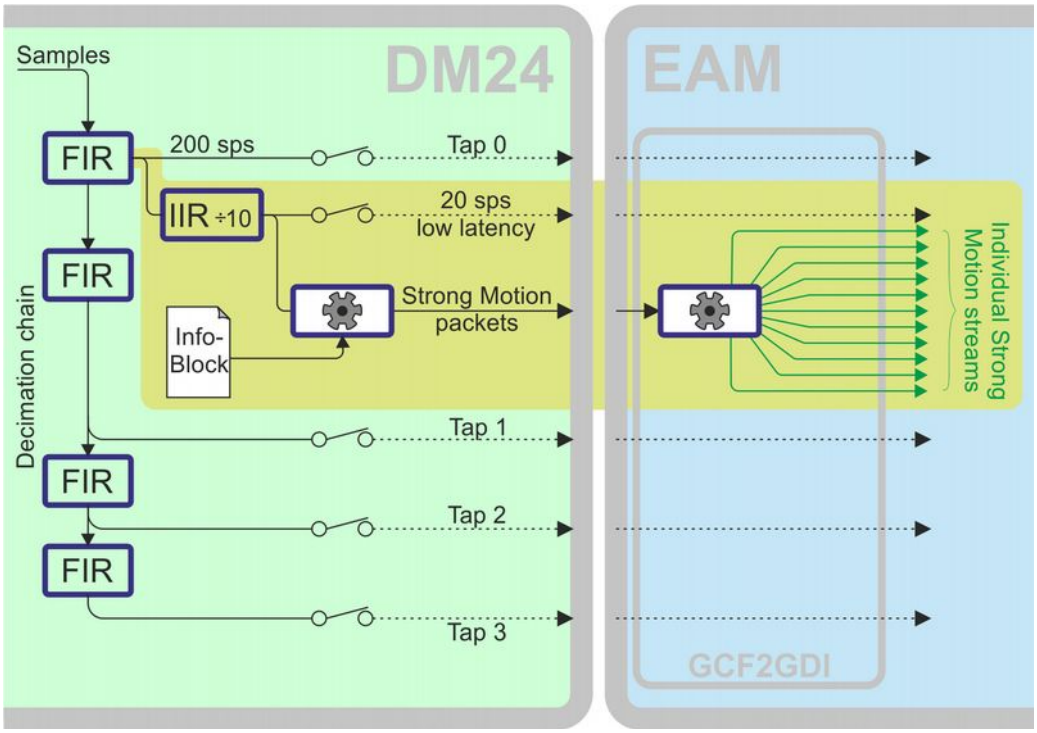

The DM24 uses the calibration information from the InfoBlock (see section 8.1.1.3 on page 102) to compute values such as minima, maxima, averages and two- and three-dimensional resultants in ground units, all of which are contained in the strong motion packets. The InfoBlock must be correctly populated for this to work. The acquisition system demultiplexes the strong motion packets and generates individual GCF streams for each type of data. Because GCF streams can only carry integer values, the floating point values from the strong motion packets are each multiplied by 32767 before being packed as GCF data.

Strong motion mode requires that the highest tap (the one with the fastest sample rate) is set to 200 samples per second. The 200 sps data are passed through a fast-response, divide-by-ten IIR filter in order to produce the low-latency 20 sps data which are used for the calculations. This 20 sps low-latency stream can be enabled for direct output: it's stream ID will have the suffix 'E'.

Strong motion mode can be configured using buttons on the web page. If the highest tap is configured for 200 samples per second operation, the following is displayed:

If the highest tap is not set to 200 samples per second, an additional button is displayed:

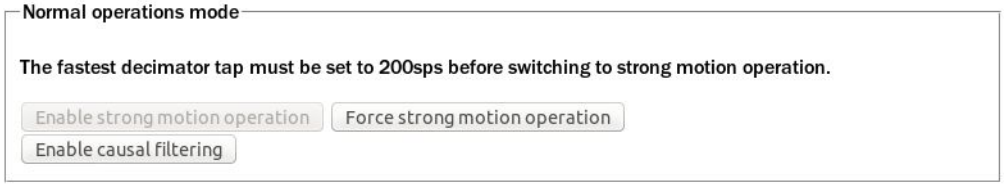

The Force strong motion operation first reconfigures the taps so that the highest tap is set to 200 samples per second and then enables strong motion mode. If the highest tap is already set to 200 samples per second, the Enable strong motion operation button is active and, when clicked, enables strong motion mode.

When in strong motion mode, this section of the screen displays as:

The “Transmit 20sps low-latency waveforms” check-box enables or disables the 20 samples per second stream, which is used internally to provide the data for the strong-motion calculations but can also be transmitted, recorded, and used like any other output stream.

Clicking the Revert to normal operation button cancels strong motion mode and returns the CMG-DM24 to normal operation.

The thirty additional streams produced when in strong motion mode are shown in the following table, where 𝕊 represents the System ID and ℂ identifies the component ('Z' for vertical, 'N' for North/South or 'E' for East/West).

Stream ID	Data content
𝕊ℂO	Windowed minimum for component
𝕊2O	Windowed minimum for horizontal resultant
𝕊3O	Windowed minimum for 3-dimensional resultant
𝕊ℂP	Windowed PGA for component
𝕊2P	Windowed PGA for horizontal resultant
𝕊3P	Windowed PGA for 3-dimensional resultant
𝕊ℂQ	Windowed maximum for component
𝕊2Q	Windowed maximum for horizontal resultant
𝕊3Q	Windowed maximum for 3-dimensional resultant
𝕊ℂR	Windowed RMS for component
𝕊2R	Windowed RMS for horizontal resultant
𝕊3R	Windowed RMS for 3-dimensional resultant
𝕊ℂS	Windowed SI for component
𝕊2S	Windowed SI for horizontal resultant
𝕊3S	Windowed SI for 3-dimensional resultant
𝕊ℂT	Windowed average for component
𝕊2T	Windowed average for horizontal resultant
𝕊3T	Windowed average for 3-dimensional resultant

8.1.1.5 Causal filtering

The normal decimation chain on a CMG-DM24 uses acausal finite impulse response (FIR) filters. Acausal FIR filters have significantly higher latency than causal FIR filters. The last filter in a decimation chain has the biggest effect on the overall latency of the chain. The CMG-DM24 offers a single causal FIR filter which can be connected at any point in the decimation chain to provide a lower-latency output.

To enable the causal FIR filter, ensure that the CMG-DM24 is in normal operation mode (not strong motion mode) and then click Enable causal filtering. Scroll to the bottom of the page and click Submit changes & Reboot digitiser.

Once the digitiser has rebooted, a new section appears on the screen, as shown below:

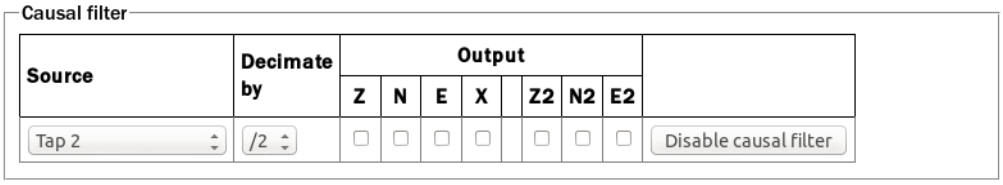

The “Source” drop-down menu allows you to select the tap that will be used as the input to the causal filter (which need not, therefore, be the final filter in the decimation chain). The “Decimate by” drop-down menu allows you to select the divisor implemented by the filter. The “Output” check-boxes allow you to select which components to output via the causal filter.

The causal filter may be disabled by clicking the Disable causal filter button, then scrolling to the bottom of the page and clicking Submit changes & Reboot digitiser.

8.1.1.6 Compression mode

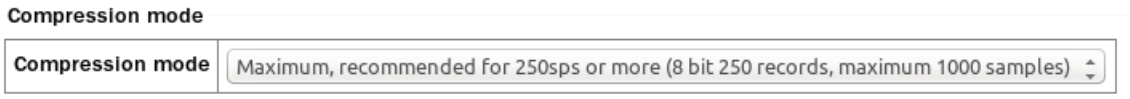

Compression mode: Controls how samples are packed, affecting both data latency and line utilisation. GCF packets contain a 32-bit starting value and then a series of differences between consecutive samples. When the input signal is relatively quiet, these differences can often be expressed as 8-bit quantities. When the input signal includes large transients, the differences are transmitted as 32-bit quantities. For intermediate level signals, 16-bit values can be used. This is known as compression as it can “compress” four samples into the space otherwise occupied by a single value. The digitiser can be configured to limit the amount of compression used, decreasing latency.

The difference values are stored in records, which are four bytes long and, so, may contain four 8-bit differences, two 16-bit differences or one 32-bit difference. A GCF packet can contain up to 250 records so the maximum number of samples in a packet is between 250 (when 32-bit differences are used) and 1000 (for eight-bit differences).

Packets must start on whole-second boundaries, so they are not always filled. In addition, it is possible to configure the digitiser to further restrict the number of records in a packet in order to decrease latency.

The drop-down menu controls both of these settings and normally offers the following choices:

- **Recommended for 200sps or less (8 bit 20 records, maximum 80 samples)**: - Represents the best compromise for throughput: allowing 8-bit compression potentially increases the number of samples per

packet and limiting packets to twenty records guarantees reasonably low latency.

- **Maximum, recommended for 250sps or more (8 bit 250 records, maximum 1000 samples)**: - Optimises line utilisation, allowing maximum compression and minimising the number of packet headers transmitted.
- **Minimum for lowest latency (32 bit 5 samples)**: - Disables compression, forcing samples to be transmitted as 32-bit differences. Latency is reduced by limiting the number of records to one per packet.

If the system is connected to a digitiser that is using a different combination of compression control and sample limits, it will appear as an extra item in the drop-down menu, labelled **Custom**. For example: **Custom (16-bit 40 samples max)** would appear if these settings had been manually configured from the digitiser's command line.

8.1.1.7 Decimator outputs

The decimator outputs control which digitiser taps have been configured to output data, both in continuous and triggered states.

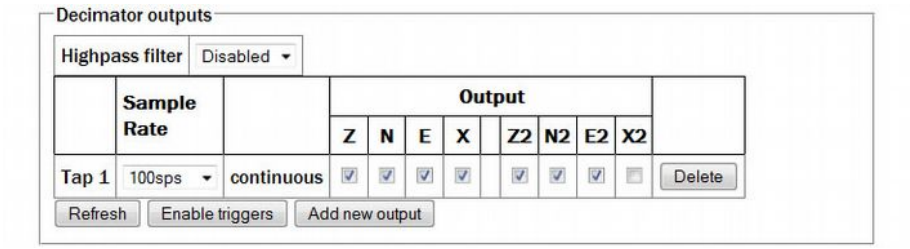

On four-channel digitisers, only the columns Z, N, E and X are displayed. On seven-channel digitisers, Z, N, E and X represent the outputs from SENSOR A and Z2, N2, E2 and X2 represent the outputs from SENSOR B.

Note: Do not confuse SENSOR A and SENSOR B, which are analogue inputs to the digitiser, with Port A and Port B, which are digital inputs to the EAM.

An optional **Highpass filter** can be applied to eliminate the effect of any DC offset in the sensor's output. The available corner frequencies, 100, 300 and 1000 seconds, can be selected from the drop-down menu.

The table below the filter configuration shows the currently configured continuous and triggered outputs with components in columns and taps in rows. The number of output columns increases when using seven-channel digitisers.

Extra taps can be added with the Add new output button. The rates available at each tap are dependant on the rate selected at the previous tap: the base sampling rate is 2000 samples per second and each tap can be configured to

divide this by either 2, 4 or 5. The available rates are shown in the table below, along with a way to configure each, although there are sometimes very many different ways to configure any given rate.

Desired output rate	Intermediate steps
4	400, 100, 20
5	400, 100, 20
8	400, 200, 40
10	400, 100, 50
16	400, 80
20	400, 100
25	400, 100
40	400, 200
50	400, 250
80	400
100	400
125	500
200	400
250	500
400	*tap 1*
500	*tap 1*
1000	*tap 1*

The triggering settings are normally hidden but can be revealed by clicking the Enable triggers or View trigger settings button. The following extra dialogues are displayed:

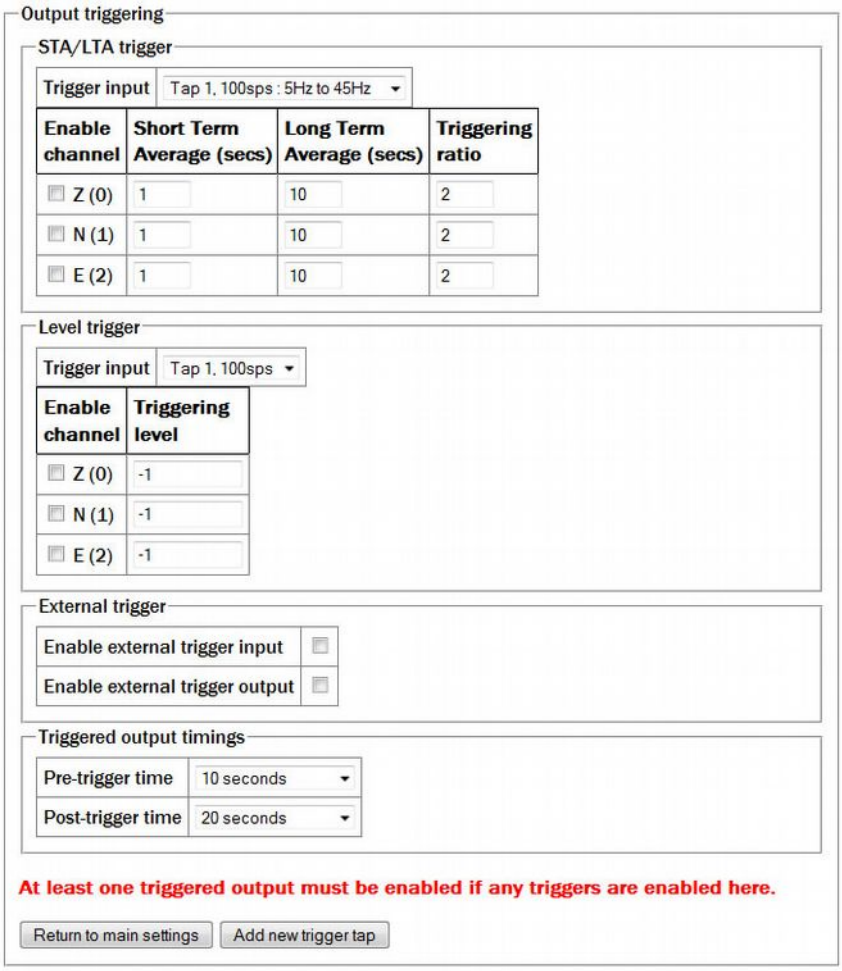

STA/LTA triggers: Activate when the ratio of the short-term average to the long-term average (of the input signal) exceeds a configured value.

Trigger input: The drop-down menu allows the selection of the input signal to be used for these calculations, along with one of three filter settings. The available filters have lower corner frequencies based on fixed fractions of the sample frequency: specifically, 5%, 15% and 25% of the sample rate. The upper corner is 0.9 times the Nyquist frequency (45% of the sample rate).

For example, if a tap configured for 25 samples per second, the three filters offered will have pass-bands of:

- 1.25 Hz to 11.25 Hz (5% of 25 Hz is 1.25 Hz and, as the Nyquist frequency for 25 samples per second is 12.5 Hz, $0.9 \times 12.5 = 11.25$);
- 2.5 Hz to 11.25 Hz (15% to 45% of the sample rate); and
- 6.125 Hz to 11.25 Hz (25% to 45% of the sample rate).

Note that the filtering specified here is applied to the specified tap output before being used for trigger input but does not affect the tap output if it is also used to generate continuous streams.

The periods over which the short term and long term averages are computed, along with the triggering ratio itself, can be altered by changing values in the table. The check-boxes next to each component are used to enable or disable the use of that component's output in the STA/LTA triggering algorithm.

The tap to be used as input should be selected from the **Trigger input** drop-down menu. All configured taps are available, regardless of whether they have been selected for continuous output or not.

The check-boxes labelled with component designators are used to include or exclude the associated component from the triggering algorithm. If one is ticked, a trigger will be activated if that component's instantaneous output exceeds the value entered into the **Triggering level** field in the same row.

Full coverage of external triggering is beyond the scope of this manual and the interested reader is referred to the relevant digitiser manual.

The **pre-trigger time** and **post-trigger time** drop-down menus control the amount of data transmitted around each trigger period. The options offered range from 5 seconds to 4 minutes. The "0 seconds" setting disables the feature.

The Add new trigger tap button inserts and extra row in the decimator output table and returns the display to the main digitiser configuration settings screen.

8.1.1.8 Multiplexor channels

The next section of the main display shows and controls the transmission of data from the auxiliary and state-of-health channels of the digitiser:

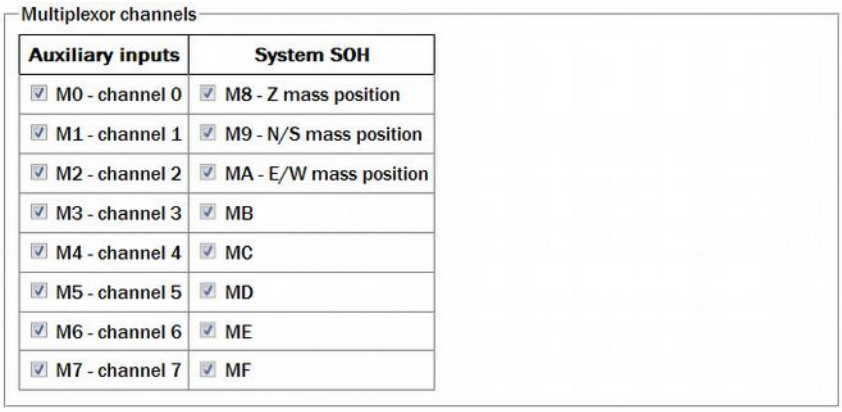

Inputs M3 through M7 and MB through MF are typically derived by digitising (at 16-bit resolution) the analogue inputs on the "Auxiliary" connector of the digitiser although, in some configurations, they may be connected to internal sensors. For example, ME is often used for an internal temperature sensor and M8, M9 and MA provide mass position data from the first sensor. Output from each displayed channel can be enabled by ticking the associated check-box or disabled by clearing it.

8.1.1.9 Transmission mode

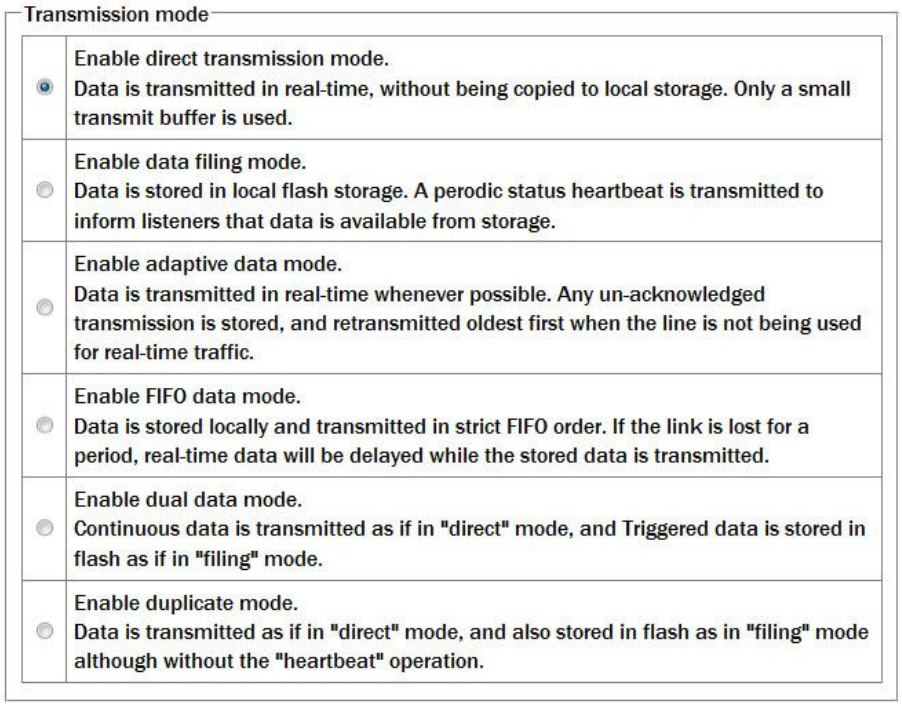

The Transmission mode selections are exactly as defined on the form.

Note that **Direct** is the default and preferred method of transmission

8.1.1.10 Storage Mode

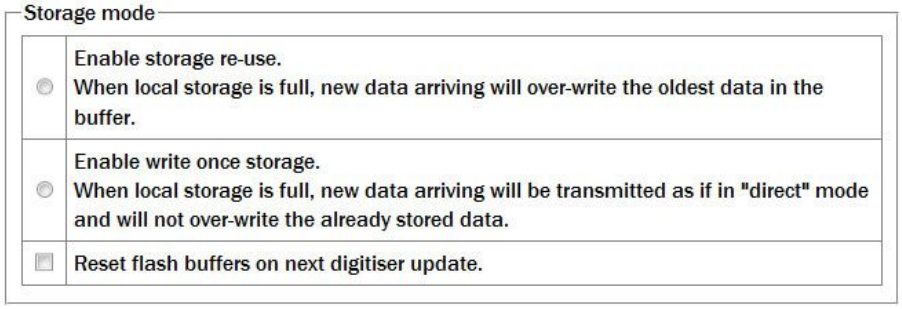

In the storage mode section, selecting the **Reset flash buffers on next digitiser update** check-box will cause all data in the flash storage to be erased and the read pointers to be reset.

Note: Use this facility with caution. Data will be erased so, if they are important, ensure that they have been flushed to storage and are readable before clearing the buffer.

8.1.1.11 Transmission Parameters

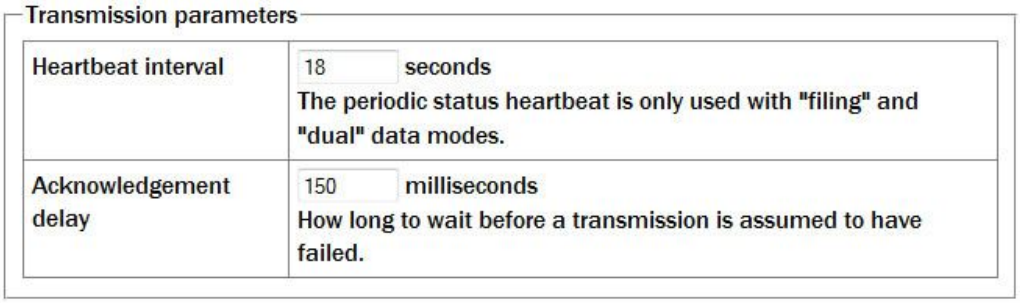

Heartbeat interval: When the digitiser is in the “filing” or “dual” transmission modes regular heart-beat messages are sent. This allows software such as Scream to be aware of the devices even though they are not sending sampled waveform data. The frequency of these messages can be set to an integer number of seconds.

Acknowledgement delay: When the digitiser is in the “adaptive” or “FIFO” transmission modes, special action is taken if data cannot be transmitted. The acknowledgement delay field controls how long the digitiser waits for an acknowledgement packet before assuming that the link has failed. This should be set to an integer number of milliseconds.

8.1.1.12 Ports

The Ports section of the web page allows control of the baud rates of the digitiser's serial ports:

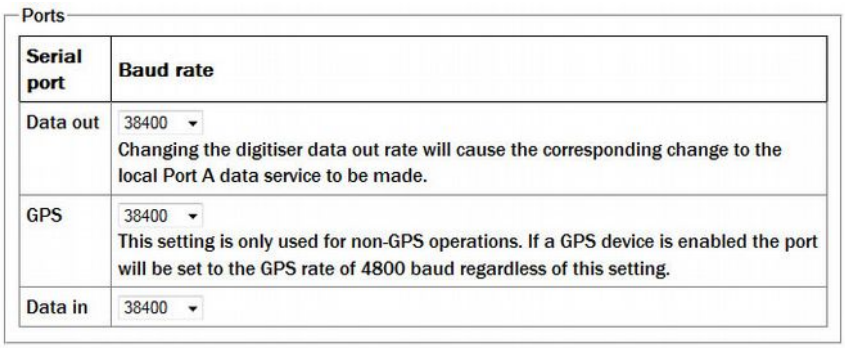

For a stand-alone digitiser or digital instrument, the “GPS” and “Data in” ports are exposed on external connectors. The “Data out” port is assumed to be connected to the acquisition module and the rate set here must match that set for the appropriate serial port (see section 10.1 on page 128 for details of reconfiguring serial ports).

Note: For CMG-DAS systems, the digitiser's “Data out” port is connected internally to the CMG-EAM module's “Port A” and both ports must use the same Baud rate. The digitiser's “Data in” port is used to provide a console for the digitiser (without interrupting seismic data transmission) and is connected internally to the CMG-EAM module's “Port B”. Again, both ports must use the same Baud rate. The digitiser's “GPS” port is connected to the CMG-EAM module's “Port C”. This connection can be used in two ways: the acquisition module module can share with the digitiser the data from the physical GPS receiver and use it as an NTP clock source (see section 7.4 on page 84); or, alternatively, the acquisition module module can be synchronised to another time source (such as Internet NTP) and provide NMEA signals to the digitiser module. In either case, the digitiser module's “GPS” port and the acquisition module module's “Port C” must use the same Baud rate.

If a stand-alone digitiser or digital instrument is fitted with a Lantronix Ethernet or WiFi option, it uses the “Data out” port settings for its internal communications with the digitiser. Changing the associated Baud rate requires making a network connection to the Lantronix unit's web interface and selecting the matching baud rate from its control page.

8.1.1.13 Miscellaneous features and Submission

The final section of the digitiser control web page is entitled “Miscellaneous features”. This section displays a warning in red if a discrepancy is detected between the EAM's time and the digitiser's own clock. If the two clocks have reasonable synchronisation, this message is suppressed. A typical warning looks like this:

> Digitiser clock is displaced by more than 5 minutes from the system clock. (Plus 7 minutes).

This section of the page is shown here without the warning:

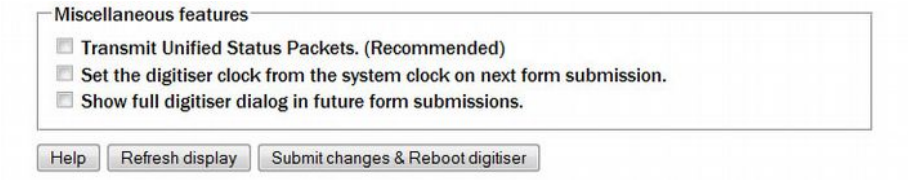

The first check-box enables the transmission of Unified Status Packets. Unified Status Packets are a machine-readable representation of the data carried in the normal, human-readable status streams and allow programs such as Scream to access complete and consistent state-of-health information regardless of any status stream customisations.

The second check-box allows the one-time re-synchronisation of the digitiser to the EAM's system clock. The third toggles display of the underlying dialogue with the digitiser, as described at the beginning of this section.

Note that, as with all web interfaces, options selected on this screen will not take effect until the page is submitted.

An extra button at the bottom of this page, Refresh display, causes the system to re-query the digitiser for its current configuration settings, allowing the refreshing of the web page display with up-to-date information.

Note: Submitting this page, whether or not changes have been made, will reboot the digitiser.

8.2 Configuring digitisers from the command line.

Platinum provides two command-line tools to allow configuration and control of digitisers and the instruments connected to them, `adc-command` (see section 8.2.1 on page 114) and `data-terminal` (see section 8.2.2 on page 114). The `adc-command` feature allows specific instrument control commands to be sent to digitisers while `data-terminal` exposes the digitiser's console, allowing arbitrary commands to be sent.

In addition, the **`dm24-upgrade`** tool (see section 8.2.3 on page 115) allows the firmware of attached digitisers to be upgraded from the command line of the acquisition device.

8.2.1 adc-command

The **`adc-command`** tool allows a number of instrument mass control commands to be issued to attached digitisers. For further details, see section 14.3.2.2 on page 227.

8.2.2 data-terminal

Platinum provides a tool, **`data-terminal`**, which allows direct access to the command-line of any serially attached digitiser. This gives the greatest level of control but also involves the most complexity.

Interactions with the digitiser's command-line are beyond the scope of this document: please consult the relevant digitiser manual for information on this topic. This section discusses use of the **`data-terminal`** tool only.

To invoke the tool, enter the command

```
data-terminal
```

You will be presented with a menu listing all digitisers to which a connection can be made:

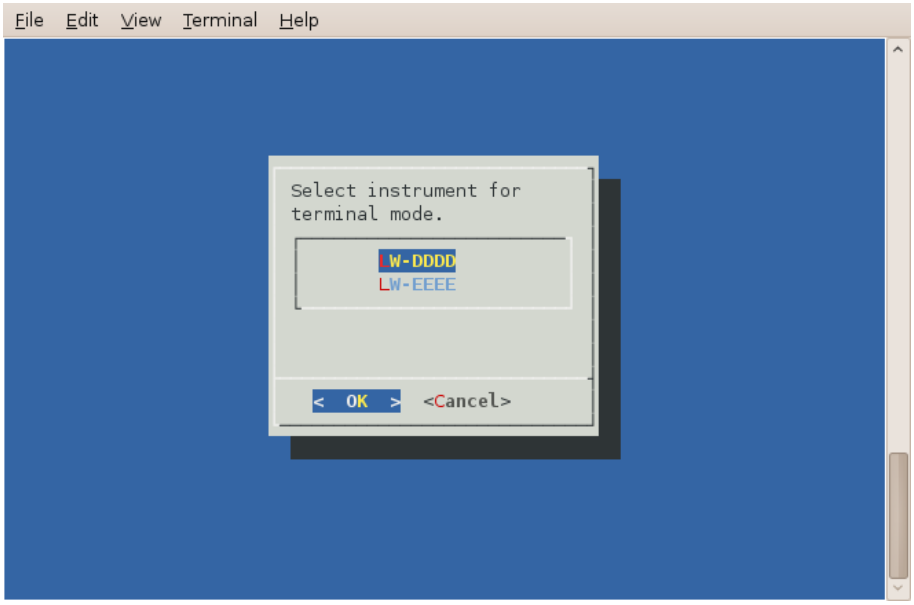

Select the required digitiser from the menu. The data-terminal program will suspend any service running on the associated port and start a minicom session with the correct communications parameters already set.

The use of minicom is described in section 16.3 on page 267. When you have finished configuring the digitiser, key Ctrl + A then Q to exit. Any previously running service will be restarted.

8.2.3 dm24-upgrade

The dm24-upgrade tool provides a command-line facility to upgrade the firmware on attached CMG-DM24 and CMG-CD24 digitisers. Each release of Platinum firmware contains the latest firmware images for both digitiser types. You may wish to upgrade your Platinum firmware (see section 5 on page 51 for details) before using this command.

To perform a simple upgrade, enter the command

```
dm24-upgrade id
```

replacing ***id*** with the port identifier (typically PortA or PortB) or the digitiser's ID (as reported by, for example, the `data-terminal` command described in the previous section).

More complex operations are possible; these are invoked by placing command-line arguments between the command and the port identifier, as in:

```
dm24-upgrade arguments id
```

The available arguments are described in the following paragraphs.

`--trashfram` — Perform a hard factory reset, all parameters will be lost.

--ids *sys ser* — Specify new system ID and serial number to set following a hard factory reset. This option will attempt to preserve the previous values if new values are not specified)

--samp # # # # — Change the sample speeds for each of the four decimation taps. This option requires four numeric values as supplied for the `SAMPLES/SEC` command described in the DM24 manual.

--cont # # # # — Change which streams are produced as continuous output from each of the decimation taps. This option requires four numeric values as supplied for the `CONTINUOUS` command described in the DM24 manual.

--trig # # # # — Change which streams are produced as triggered output from each of the decimation taps. This option requires four numeric values as supplied for the `TRIGGERED` command described in the DM24 manual.

--gps-baud # — Set the line speed for the GPS port. This option requires a single numeric value: the new line speed to use.

--in-baud # — Set the line speed for the DATA IN port. This option requires a single numeric value: the new line speed to use.

--upgrade — Upgrade firmware files only (this is the default).

--downgrade — Allow firmware files to be downgraded as well as upgraded.

--force — Firmware files are to be loaded even if they appear identical, or the installed version number cannot be decoded.

--boot *file* — Specify a source file for DM24mk3 bootstrap. This option requires a single value, ***file***, which is the path to the firmware image file to be loaded.

--firm *file* — Specify a source file for DM24mk3 or CD24 firmware. This option requires a single value, ***file***, which is the path to the firmware image file to be loaded.

--dsp *file* — Specify a source file for DM24mk3 DSP code. This option requires a single value, ***file***, which is the path to the firmware image file to be loaded.

--auto-baud — Scan for digitiser baud rate and adjust configuration. If possible.

--verbose — Show most of the digitiser dialog.

--debug — Output some additional debug information from the underlying expect script.

8.3 Configuration for a second instrument

When the acquisition module detects that an attached (or built-in) digitiser supports a second instrument (SENSOR B), additional items appear on the web interface menu for the digitiser.

To use a second instrument, click the link for the associated digitiser in the "System Setup" sub-menu of the "Configuration" menu. The resulting screen will display a text-field, **Serial number 2**. Populate this field with the serial number of the second sensor (or any desired value) and submit the page, rebooting the digitiser.

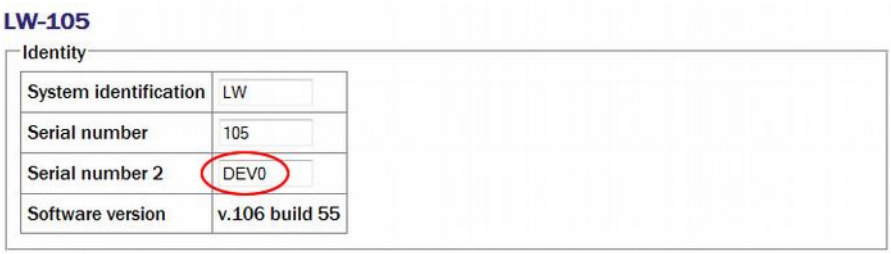

Wait a short while and then use the refresh facility in your web browser to reload the main menu. An extra device will appear in the "Control" section of the main menu and an extra link for the digitiser will appear in the "System Setup" sub-menu of the "Configuration" menu:

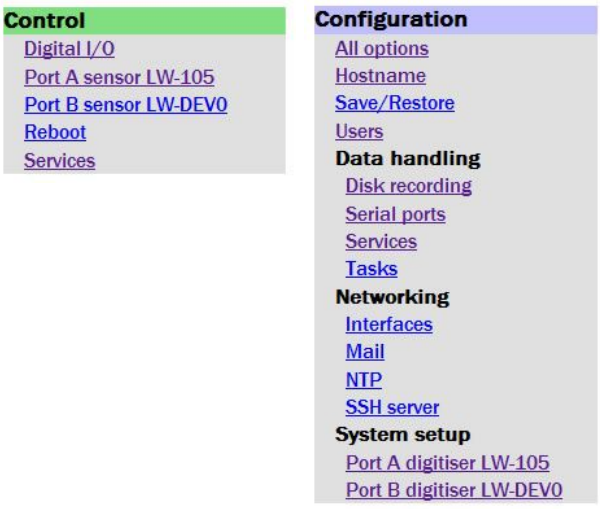

Although there are two links on the 'Control' section, only the first sensor can be controlled.

Both links for the digitiser take you to the same page and either may be used (This allows recovery from the situation where two physical digitisers with the same serial numbers have been connected).

If either of these links are selected, the resulting page differs slightly from that described in Section 8.1 on page 100.

The “Device info blocks” section now displays the status of two Info blocks - one for each connected sensor. The “Display device info blocks” button takes you to a screen from where you can edit both.

Extra columns appear in the “Decimator outputs” configuration section:

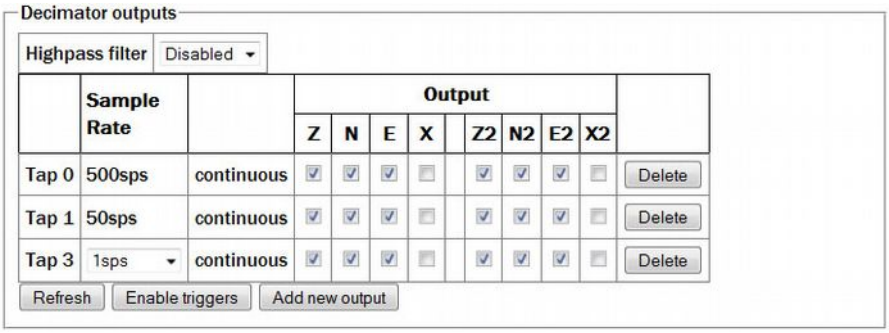

Streams from the second sensor are enabled or disabled by ticking or clearing the check-boxes in the columns labelled Z2, N2 and E2. An X2 stream will appear in this table when a seven- or eight-channel digitiser is detected, but it is not available on a seven-channel digitiser.

Notes:

1. The second sensor is assumed to be an accelerometer. No provision is made for mass control (locking, unlocking and centring) of SENSOR 2.

2. The Sensor type drop-down menu in the digitiser configuration page refers to the first sensor only. SENSOR 2 is assumed to be an accelerometer.

3. InfoBlock 1 refers to SENSOR 1 and InfoBlock 2 refers to SENSOR 2. Values entered in these InfoBlocks are passed to Scream, which will apply them to the correct data streams.

4. Multiplexor outputs M8 (Z mass position), M9 (N/S mass position) and MA (E/W mass position) refer to SENSOR 1 only.

5. When configuring triggering, input streams are available from both sensors (at the selected tap). When a trigger condition is detected, configured outputs from both sensors are enabled, regardless of which sensor generated the trigger.

9 Digitiser Synchronisation

Accurate time-stamping of samples is essential to seismology. Güralp Systems Ltd recommend the use of GPS receivers for generating clock signals wherever possible: they are the most accurate time sources available for all practical purposes (but see below: GPS is best used indirectly). Where GPS receivers cannot be used but an internet connection is available, Network Time Protocol (NTP) can produce acceptable results. Platinum firmware can produce NTP-synchronised NMEA output for use with GPS-capable devices.

All acquisition modules have an internal clock which is used to time-stamp log-file entries (but not data samples). This clock is managed by the NTP subsystem but need not necessarily use Internet NTP servers (and normally doesn't).

DAS units, Cylindrical Digitisers and Integrated Instruments have two clocks: the digitiser clock and the acquisition module clock: the former is used to time-stamp data samples and the latter to time-stamp log-file entries. The acquisition module clock is also managed by NTP. The two clocks can be synchronised in a number of ways.

9.1 Overview and important notes

Acquisition modules can synchronise their internal clocks to three different time sources:

- special timing signals, known as RTSTATUS packets, produced by a digitiser
- internet NTP
- signals from a GPS receiver

GPS receivers produce two signals:

- NMEA, which is a serial, ASCII data stream consisting of a sequence of “sentences” which convey information such as the time (to the nearest second), the position, the number of satellites visible and much more
- PPS, which is a one-pulse-per-second signal which accurately signals the start of each second

GSL digitisers, such as the DMG-DM24, use both signals to accurately synchronise their internal clocks.

Because of a limitation in the Linux kernel on ARM processors, the CMG-DCM and CMG-EAM cannot make use of the PPS signal. This means that a DCM or EAM synchronised directly to GPS can be several hundreds of milliseconds adrift. While this is not crucial - the clock is only used to

time-stamp log-file entries - it does make direct synchronisation to GPS the least accurate of the three options available.

We recommend the following synchronisation strategy:

- If GPS is available, synchronise the digitiser(s) directly to GPS and synchronise the acquisition module to the RTSTATUS packets from the digitiser.
- If GPS is not available but an internet connection can be used, synchronise the acquisition module to internet NTP and enable the NMEA output, to which the digitiser can then be synchronised.
- If neither GPS nor internet NTP are available, contact technical support for advice.

9.2 RTSTATUS packets

Where a CMG-EAM or CMG-DCM is used with an external GPS-synchronised GSL digitiser, the digitiser can emit special synchronisation packets called RTSTATUS packets. These are transmitted along the same link as the data packets. Platinum units can use these as a time source for NTP: see section 7.4 on page 84 for more details.

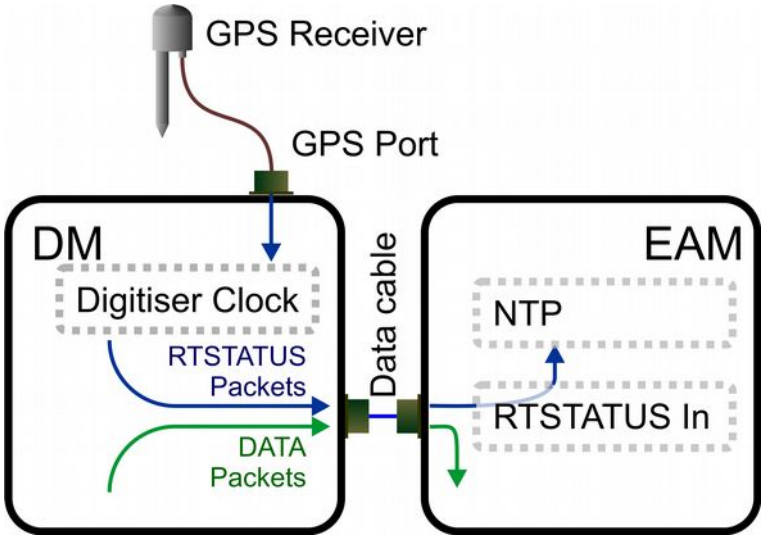

Note: RTSTATUS packets are available with MkIII DM24 units. Earlier units, such as MkIIs, are not capable of generating these packets. With MkIIs, however, the system clock can be set from GCF status block timestamps by enabling this feature on the expert settings of the appropriate serial port GCF input page.

Note: RTSTATUS packets are available with CD24 units running firmware revision 279 and above.

This is also the recommended configuration for CMG-DAS units, where GPS reception is available.

9.3 Using NTP with CMG-NAM units

Where GPS reception is not practical but an internet connection is available, NTP can be used to synchronise the Platinum clock, which can then generate NMEA output. This NMEA data-stream can be fed to the GPS input of the digitiser module using an external cable from a serial port.

Note: This technique is only applicable to CMG-NAM and CMG-NAM64s, where the Linux kernel can correctly handle the incoming PPS from the GPS receiver.

The data flow when Internet NTP is used is illustrated below:

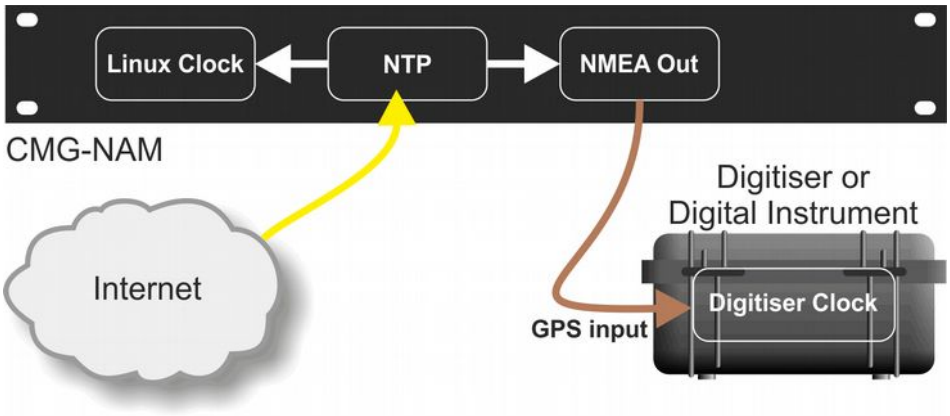

To configure the NTP subsystem, see section 7.4 on page 84.

To configure NMEA output, see section 9.7 on page 125.

9.4 Using GPS with Cylindrical Digitisers

Güralp Systems Ltd's cylindrical digitisers provide a CMG-DM24 and a CMG-EAM in a single package. An internal, bi-directional connection is available between Port C of the EAM and the digitiser module. This connection can operate in one of two modes:

- The EAM's NTP subsystem can provide NMEA to the digitiser module's clock circuitry. In this case, the external GPS socket is automatically disconnected. This is the recommended configuration when GPS synchronisation is impossible but internet NTP is available; or
- An external GPS receiver can provide input to both the digitiser module's clock circuitry and the EAM's NTP subsystem. This configuration is possible but not recommended.

Note: Where GPS is available, you should always synchronise the EAM to the digitiser's RTSTATUS packets rather than directly to GPS, as explained in section 9.1 on page 120.

The data flow when a GPS receiver is used is illustrated below:

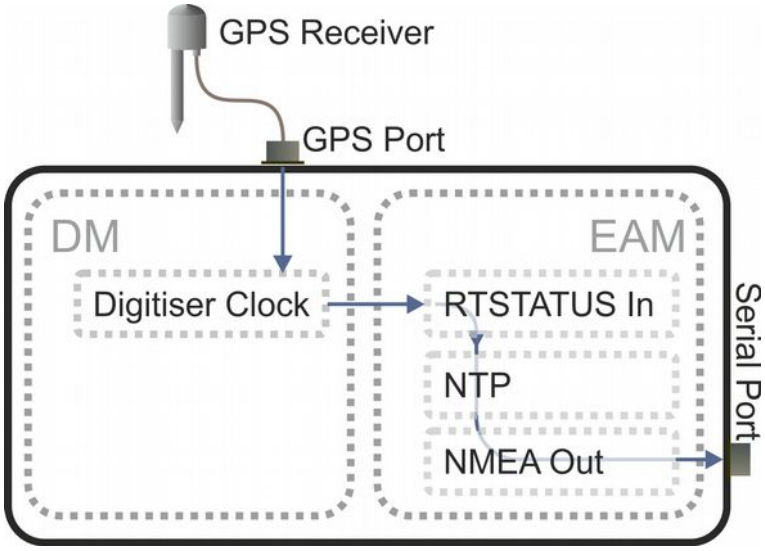

If required, the NTP subsystem can provide NMEA output via a serial port which can then be used to synchronise an additional digitiser. This, however, is optional and no serial port is dedicated to this use.

To configure the NTP subsystem, see section 7.4 on page 84.

To configure NMEA as an NTP clock source, see section 9.6 on page 124.

To configure NMEA output, see section 9.7 on page 125.

9.5 Using NTP with Cylindrical Digitisers

Note: Where GPS is available, you should always synchronise the EAM to GPS, via the digitiser's RTSTATUS packets, rather than to NTP.

Please see the discussion of synchronisation options available with Cylindrical Digitisers in the previous section.

The data flow when NTP is used as the primary clock source is illustrated below:

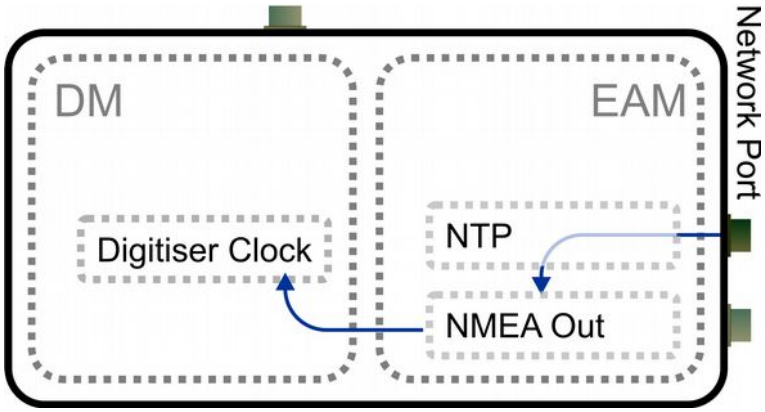

To configure the NTP subsystem, see section 7.4 on page 84.

To configure NMEA output, see section 9.7 on page 125.

Note: The external GPS connector is disconnected when Port C of the EAM is set to “NMEA Out” and connected to the digitiser's GPS input in all other cases.

9.6 Configuring NMEA as an NTP clock source

To configure NTP to use NMEA as a clock source, two steps are required. First, tick the **Acquire time from connected GPS** check-box in the NTP configuration page as described in section 7.4 on page 84.

Secondly, configure the relevant serial port as an NMEA input.

To do this using the web interface select:

Configuration → Serial ports

or

Configuration → All options → Serial ports

To configure an NMEA input from the command line, start gconfig and select “Serial ports” from the top level menu.

Select a serial port from then click NMEA output settings.

Note: For cylindrical digitisers, this will be Port C. For more details of Cylindrical Digitisers, please see section 15.1 on page 250.

9.6.1 Configurable parameters

Port function: Set to "NMEA in. Receive GPS data for NTP"

Port speed: Set to 4800.

Click Submit to save the changes.

9.7 Configuring NMEA output

Platinum can generate simulated GPS data (NMEA-0183) to synchronise a connected digitiser's clock. In this case, the internal clock of the acquisition module is used as a reference for the digitiser. In order to provide a sufficiently accurate time-stamp, the clock must be controlled using NTP (See section 7.4 on page 84).

To configure NMEA Output using the web interface select:

Configuration → Serial ports

or

Configuration → All options → Serial ports

To configure NMEA Output from the command line, start gconfig and select "Serial ports" from the top level menu.

Next, select the serial port from which you want to output NMEA. Only one port can be used for NMEA output at any time: the timing constraints are such that a single processor cannot produce the pulse-per-second (PPS) signal on two ports simultaneously with sufficient accuracy.

Note: For Cylindrical Digitisers, Port C should be used in order to provide NMEA output to the internal digitiser module. If a GPS receiver is used, the "Data Out" Port (exposed as the DATA connector) can be used to provide timing to additional, external digitisers.

9.7.1 Configurable parameters in simple mode

Port function: Set to "NMEA out. NMEA (time + fixed position) output"

Port speed: Set to 4800.

Click Submit to save these changes.

Go back to the configuration of the serial port and click on "NMEA output settings". You will see this screen:

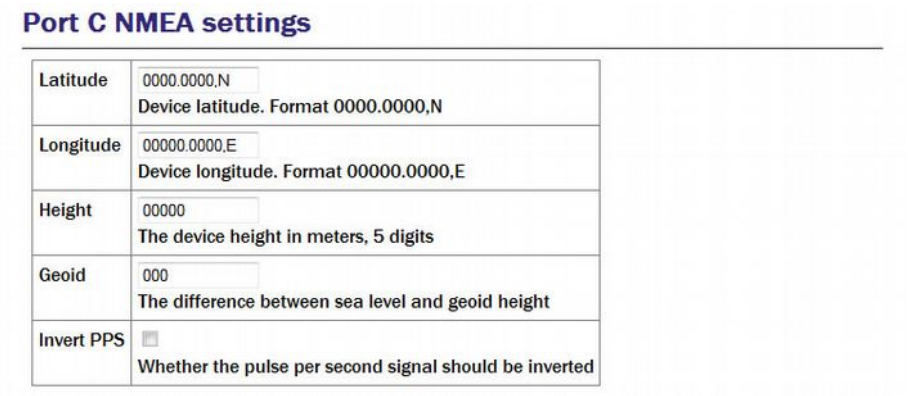

Here, you can configure the NMEA sentences that will be sent to the digitiser. You can specify the location (latitude, longitude, elevation), the geoid (the offset of the location from the theoretical earth surface) and whether to invert the Pulse-Per-Second signal (if unchecked, the PPS line will be briefly asserted each second, on the second, and held to ground at other times). It is not essential that the position string sent matches the physical location of the digitiser, as only the GPS time signal is used by the digitiser. Click Submit to save the changes.

9.7.2 Configurable parameters in expert mode

The following additional parameters are available in expert mode:

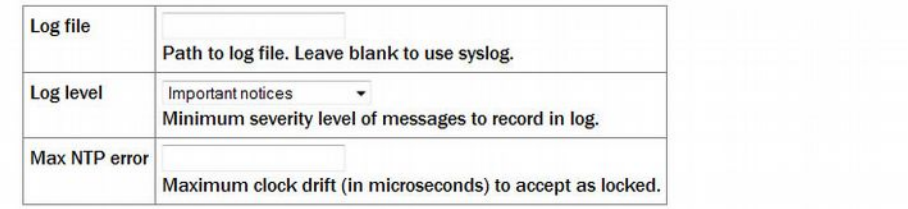

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

Max NTP error: Controls the accuracy of synchronisation which must be achieved by NTP before the resulting NMEA sentences will will indicate that the “GPS” is locked.

10 Receiving Data

The modular architecture of Platinum software allows seismic data to be received simultaneously from a number of sources and using a number of protocols. Extra protocols can be implemented by request: please contact Güralp Systems Ltd for more details.

At the time of writing, Platinum firmware is shipped with support for CD1.1, Güralp Compressed Format (GCF) data received over serial ports, GCF data received over a network using the Block Recovery Protocol (BRP) and GCF data forwarded from a copy of Scream or a Scream server, such as a CMG-EAM.

The use of CD1.1 is covered in a separate manual, MAN-EAM-1100. The use of the other receivers is described in this section.

10.1 GCF from serial devices

Any or all of the serial ports may be configured to receive GCF data from a serially attached digitiser or digital instrument.

To configure a port for this purpose from the web interface, select:

Configuration → Serial ports

or

Configuration → All options → Serial ports

The following screen is displayed:

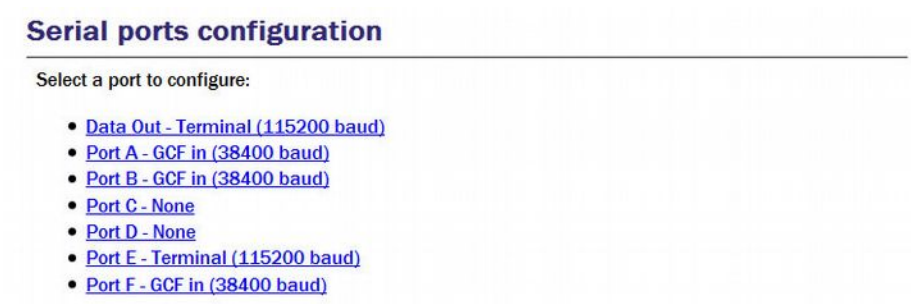

To configure a GCF input port from the command line, start `gconfig` and select “Serial ports” from the top level menu.

Each port on the system is listed along with its function and line speed. Any port can be used for any function with the exception of the console port, which is dedicated to the terminal function, and the internal ports used for inter-module communications in CMG-DAS units.

Select the link for the serial port you wish to configure for GCF input.

Note: When configuring units without a dedicated console port, such as the CMG-DCM, take care not to "lock yourself out" of the system by, eg, configuring all serial ports for non-terminal functions before completing network access configuration.

10.1.1 Configurable parameters in simple mode

The configurable parameters for GCF input ports are contained in a single form in simple mode:

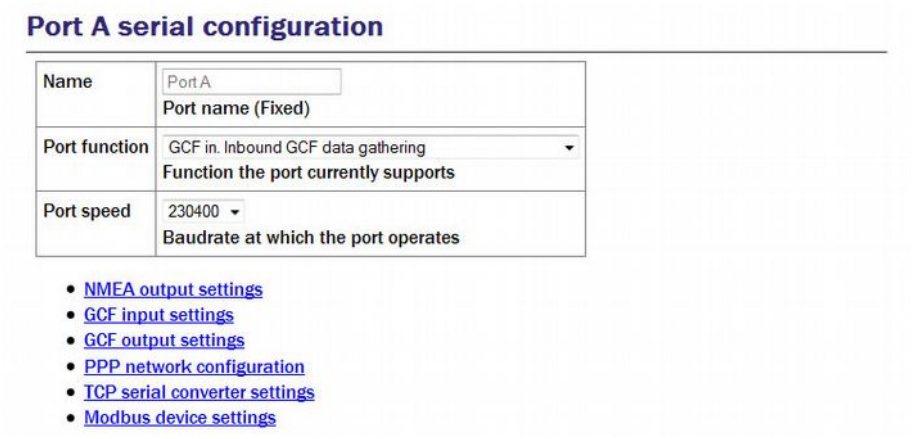

Port Function: Select "GCF in. Inbound GCF data gathering".

Port speed: Set the appropriate Baud rate from the drop-down menu

Click on the GCF input settings link.

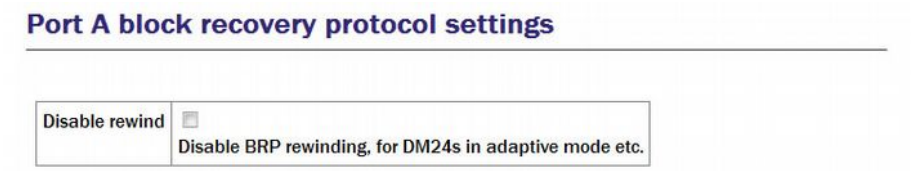

In simple mode, the only option available is to disable BRP rewinding. In some modes, some digitisers will not allow BRP to rewind to earlier blocks. In these modes, missed packets will, instead, be sent at a later time. However, the log-file will accumulate many entries about sending NAKs and giving up. These may be avoided by telling the receiver that its digitiser is using one of these modes and that rewinding will not work. The log messages are harmless, so leave this check-box clear if unsure.

10.1.2 Configurable parameters in expert mode

The following additional options appear in expert mode:

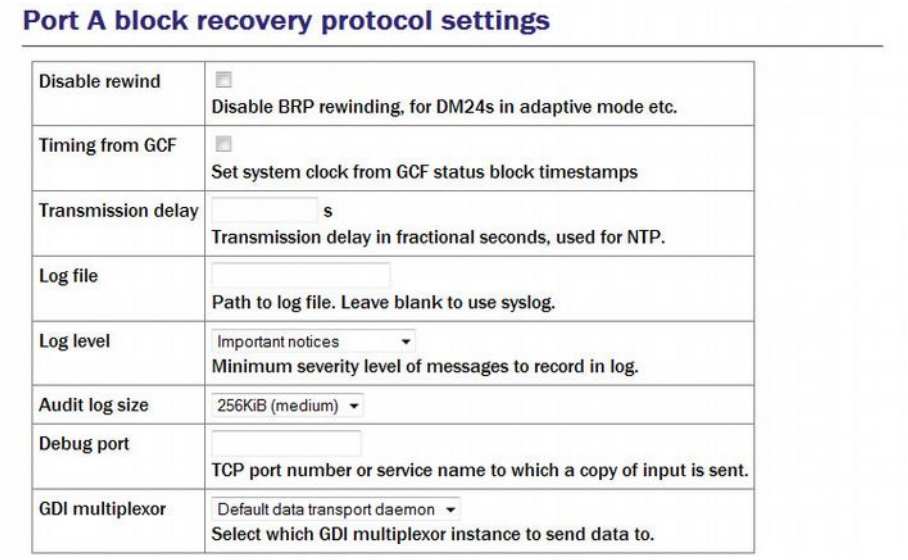

Port A block recovery protocol settings

Disable rewind	☐ Disable BRP rewinding, for DM24s in adaptive mode etc.
Timing from GCF	☐ Set system clock from GCF status block timestamps
Transmission delay	s Transmission delay in fractional seconds, used for NTP.
Log file	Path to log file. Leave blank to use syslog.
Log level	Important notices Minimum severity level of messages to record in log.
Audit log size	256KiB (medium)
Debug port	TCP port number or service name to which a copy of input is sent.
GDI multiplexor	Default data transport daemon Select which GDI multiplexor instance to send data to.

Transmission delay: Allows the operator to specify the total delay incurred during packet transmission from the attached digitiser. Digitisers can produce special "RTSTATUS" packets which can be used to synchronise the NTP subsystem and, hence, the system clock, with the digitiser's own GPS-synchronised clock (see section 7.4 on page 84). Unlike normal NTP peer dialogues, there is no transmission delay discovery mechanism so, for optimal accuracy, it is important to specify the value here. The ordinary delay associated with packet transmission down a "short" serial cable is already calculated and used, so this field only needs populating if additional delays generated by, say, modems or radio links are encountered.

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

Audit log size: The GCF input subsystem keeps its own audit log, independent from the system log. The contents of this log are available using the "GCF Audit Log viewer" facility as described in section 14.4.5 on page 241. The amount of data retained is controlled by the drop-down menu where the choices are:

- 64Kib (small)
- 256Kib (medium)
- 2MiB (large)
- 16MiB (huge)

Debug port: It is possible to copy all incoming data, verbatim, to a network port, which can be specified in this field. This is an advanced debugging technique which is beyond the scope of this manual.

GDI multiplexer: In most configurations, all data from all inputs is sent to a single multiplexor which then feeds all outputs, as described in section 6 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this receiver. The menu offers a list of currently configured multiplexers.

10.2 BRP - GCF From Network Devices

The acquisition module can receive data from network enabled instruments such as the CMG-6TD and networked digitisers such as the CMG-DM24. Data can be received from any number of sources, by creating multiple GCF BRP receiver instances.

To set up a GCF BRP receiver on the acquisition module, select

Configuration → Services → GCF

or

Configuration → All options → System services → GCF

To configure a receiver from the command line, start `gconfig` and select "System services" from the top level menu.

Select "gcf-in-brp GCF BRP network client". The screen shows a list of all GCF BRP receiver instances that have been configured

To configure a new GCF BRP receiver instance, select "Create new service instance". The screen allows you to configure the parameters of the service

10.2.1 Configurable parameters in simple mode

The configurable parameters for GCF BRP receiver are contained in a single form in simple mode:

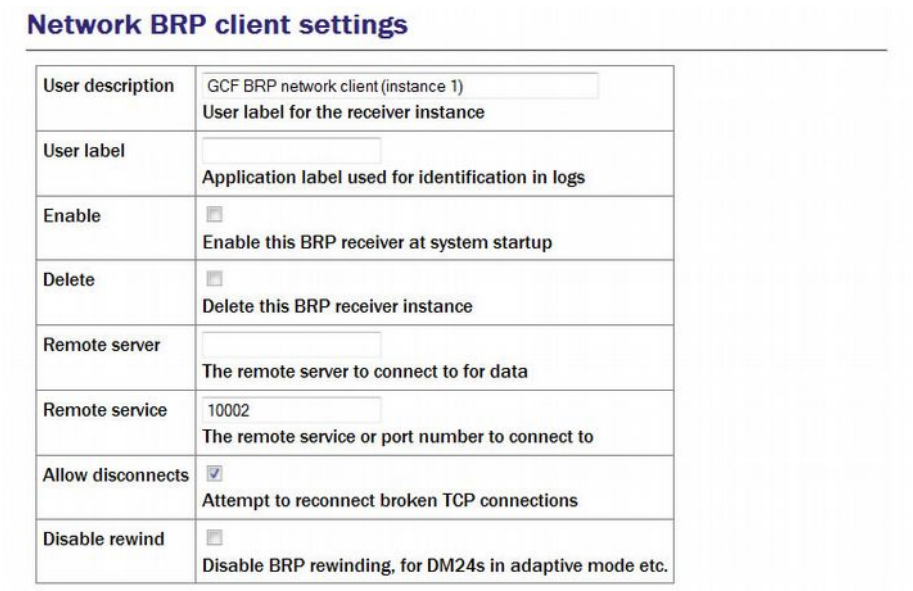

User description: Sets the name of the service; this should be set to a meaningful name for the data that it will be receiving, such as the IP or hostname of the network digitiser.

User label: Identify the particular client instance in log-files (optional).

Enable: Causes this service to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

Remote Server: Specify the hostname or IP address of the network digitiser

Remote service: Specify the port (name or number) that the digitiser is transmitting on.

Allow disconnects: If checked, the instance will attempt to automatically recover from lost connections by trying to reconnect to the server.

Disable rewind: If checked no attempts will be made to request missing data blocks. This should only be selected if the server is unable to fulfil such requests.

10.2.2 Configurable parameters in expert mode

A number of additional configuration parameters are available by clicking the "Expert" button at the bottom of the form.

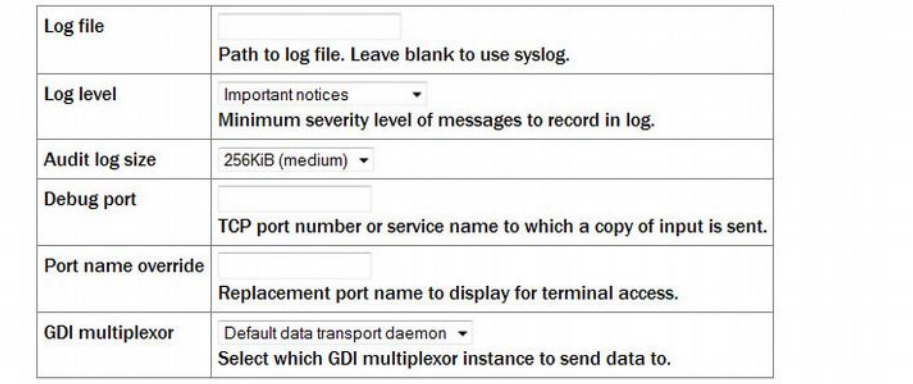

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

Audit log size: The GCF input subsystem keeps its own audit log, independent from the system log. The contents of this log are available using the "GCF Audit Log viewer" facility as described in section 14.4.5 on page 241. The amount of data retained is controlled by the drop-down menu, where the choices are:

- 64Kib (small)
- 256Kib (medium)
- 2MiB (large)
- 16MiB (huge)

Debug port: It is possible to copy all incoming data, verbatim, to a network port, which can be specified in the text field. This is an advanced debugging technique beyond the scope of this manual.

Port name override: Allows the operator to specify a descriptive name for this data source. If left blank, it will be labelled with the IP address and service number of the source device, and this label will appear in, for example, the GDI channels display and the network tree in Scream.

GDI multiplexer: In most configurations, all data from all inputs is sent to a single multiplexor which then feeds all outputs, as described in section 6 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this receiver. The menu offers a list of currently configured multiplexers.

10.3 Data from Scream servers

The acquisition module has the ability to receive data over the network from Scream servers. Data can be received from a number of Scream servers using a single Scream client.

To set up a Scream client on the acquisition module, select

Configuration → Services → GCF

or

Configuration → All options → System services → GCF

To configure a Scream client from the command line, start `gconfig` and select “System services” from the top level menu.

Select “gcf-in-scream -- GCF Scream network client”. The screen shows a list of all Scream network client instances that have been configured.

To configure a Scream receiver, select “Create new service instance”.

10.3.1 Configurable parameters

The configurable parameters for Scream network clients have three tabbed pages: General, Network and Servers:

10.3.1.1 General

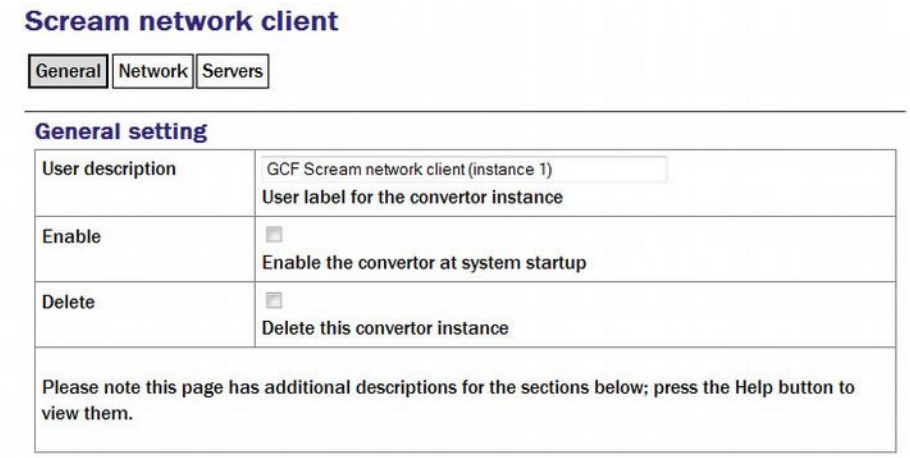

User description: Sets the name of the service. This should be set to a meaningful name for the data that it will be receiving, such as the IP or hostname of the Scream server.

Enable: Causes this service to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

10.3.1.2 Network

Local address: If the acquisition module has multiple IP addresses, you can optionally restrict the client so that incoming connections are only listened for on one address. Leave blank to listen on all available interfaces.

Local service: Enter the UDP/TCP port number on which the server is to listen for data requests. Port numbers can be mapped to names using the standard Linux /etc/services file, which can be edited from the command line. Leave blank to use the default scream port, 1567.

These two fields can normally be left blank.

10.3.1.3 Servers

On the “Servers” page, you specify the details of any Scream servers from which you want to pull data.

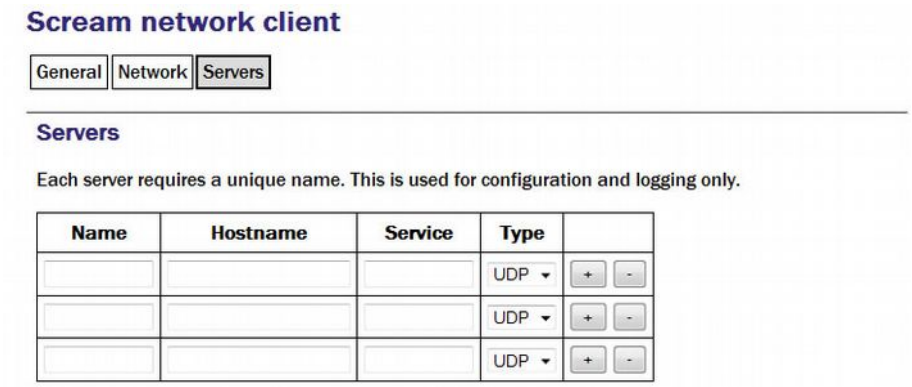

Name: A descriptive name for identification purposes.

Hostname: Enter the DNS name or IP address of the desired server.

Service: Enter the UDP/TCP port number on which the server is listening for data requests. Port numbers can be mapped to names using the standard Linux `/etc/services` file, which can be edited from the command line.

Type: Select whether you wish to use UDP packets or TCP connections. With UDP packets, the GCF protocol keeps track of which packets have been received and automatically requests retransmission of any missing data. TCP, on the other hand, is a connection-orientated protocol which handles packet sequencing and retransmission itself (at the cost of a little extra network overhead).

11 Recording and Retrieving Data

Data can be recorded to internal and external storage, in raw GCF format or in miniSEED format. Data can be browsed via the web interface or copied to external computers for further processing.

11.1 Preparing removable mass storage devices

When a new removable mass storage device is to be used with a acquisition module, it must first be formatted for use. The mass storage device can be formatted by any computer, but the acquisition module also has the capability of formatting the mass storage device itself. The acquisition module accepts mass storage devices formatted in either ext3 format (which is faster and more reliable, but can only be read under Linux systems) or VFAT format (slower and arguably less reliable, but can be read under all operating systems). To prepare the mass storage device on a PC, simply format it with a single partition containing either of the above file-systems; it can then be inserted directly into an acquisition module.

Note: When using removable mass storage devices from acquisition modules and CMG-DCMs with PCs, you may need to provide a power supply for the mass storage device.

When using a six-circuit (powered) IEEE1394 FireWire interface, the mass storage device can draw its power from the host PC. When connected to a four-circuit (un-powered) FireWire interface, such as Sony i.Link, external power needs to be applied as described below. Power also needs to be supplied when using the USB interface.

A power supply of between 4.5V and 30V DC should be connected to the 2.1 mm barrel connector (uppermost in the picture). The central pin of the connector should be connected to the positive supply line.

***IMPORTANT**: Do not connect anything to the larger barrel connector, if fitted. This was used for the heater and temperature sensor on units manufactured before 2013. On later units, it is replaced with a clear plastic window which allows easy viewing of the internal diagnostic LEDs.*

To prepare a mass storage device using the web interface select:

Tools → Removable disk → Format disk

After a delay, while the mass storage device powers up, the following screen will appear:

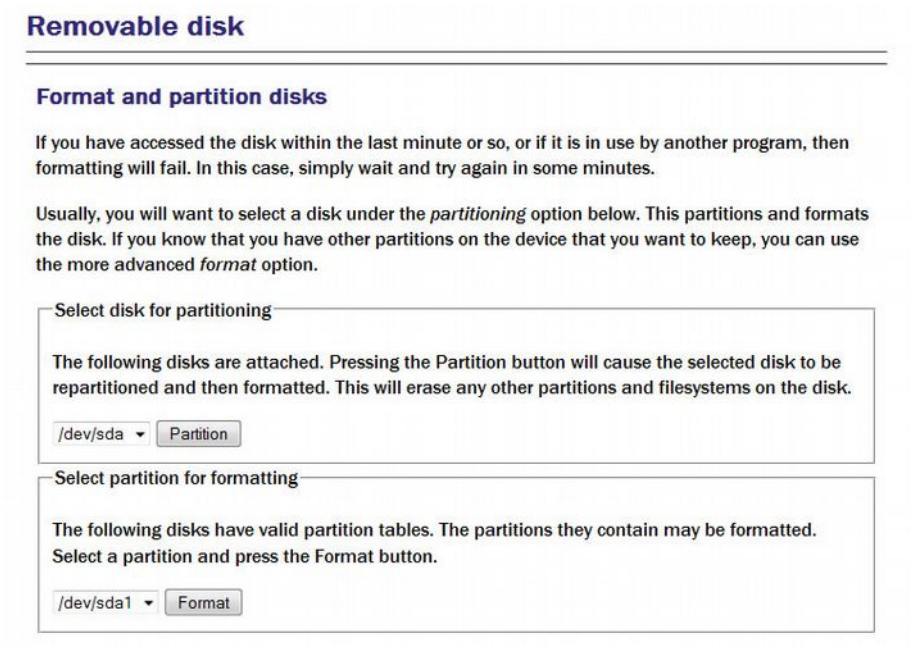

The drop-down menus for device selection use the Linux device naming convention, where `/dev/sda` is the first device, `/dev/sdb` is the second and so on. Individual partitions on devices are identified by an appended number, so `/dev/sda1` is the first partition on the first device and `/dev/sdb3` is the third partition on the second device.

The Partition button causes the selected device to be repartitioned with a single partition which is then reformatted. If you are using a device with existing partitions that you wish to preserve, you should use the Format button instead.

Successful completion of the format is signalled by a short, on-screen message. The mass storage device is then ready for use.

Mass storage devices can also be formatted from the command line with the command

```
eam999 ~ # Pt-storage --format
```

11.2 Recording data

All data recording is performed by the `gdi-record` service. Data are recorded first to a buffer held in flash memory. When the buffer utilisation exceeds a configurable percentage, data are flushed to the hard drive. In low power applications, the hard drive will normally be powered down and, so,

must be powered up and mounted before use. This is handled by the `Pt-storage` service. The disk is dismounted and powered down once the flush is complete. An option exists to provide permanent power to the disk.

Older removable hard drives were equipped with a temperature sensor and heating element. In low temperature applications, the drive would be automatically warmed to a safe operating temperature before power was applied. Modern GSL removable hard drives do not require pre-heating.

The options that control this process are all on one page but, given its size, it is shown here in sections.

To configure data storage using the web interface, select

Configuration → Storage and recording

or

Configuration → All options → Storage and recording

To configure recording from the command line, start `gconfig` and select "Storage and recording" from the top level menu.

11.2.1 Configurable parameters

The configuration screen for recording data has five tabbed pages: Storage device, MiniSEED, GCF, Non-seismic data, NFS and Advanced. No expert mode screens are required.

Selecting

Configuration → Storage and recording

displays the following screen:

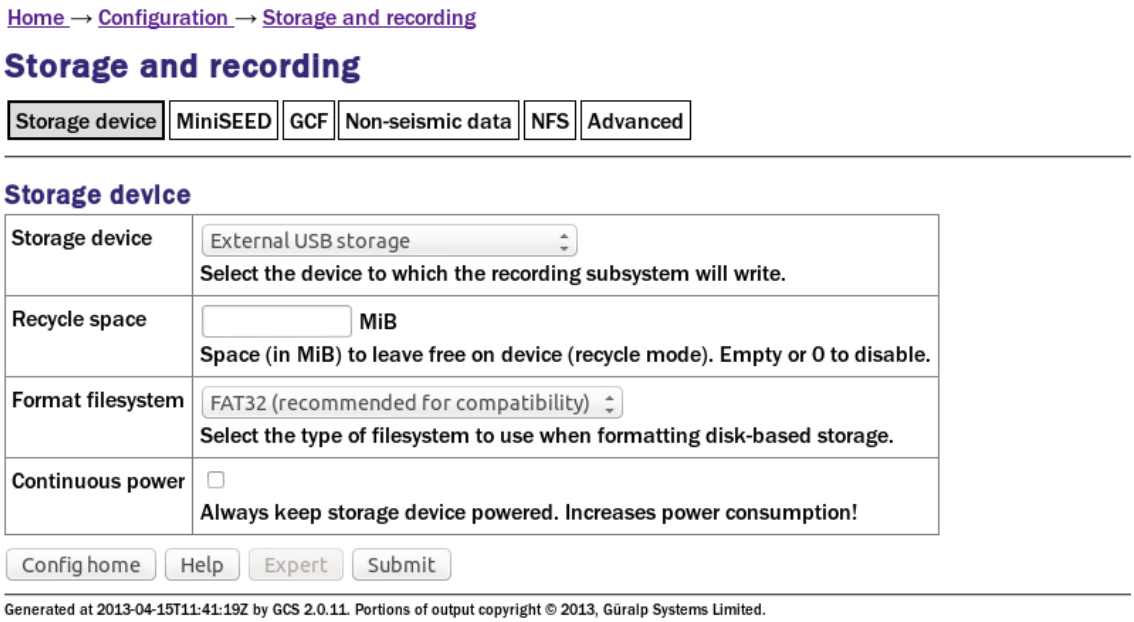

11.2.1.1 Storage device

The “Storage device” tab has the following fields:

Storage device: This drop-down menu lets you choose from the storage options available on your system. The options may include one or more of:

- “Removable USB disk in internal tray” - this is the default setting for CMG_EAMs and CMG-DCMs. Data are written to the removable USB/FireWire disk and can be read via the web interface, the command line or by removing the mass storage device and attaching it to an external USB or FireWire host, such as a PC or laptop.;
- “External USB drive on mil-spec connector” - data will be stored on external media, which should be attached to the USB connector. The Pt-storage daemon handles the required operations for mounting and un-mounting the file-system;
- “Internal USB storage” - for cylindrical digitisers only, data can be stored on internal USB-accessible flash memory. It can be retrieved by a USB host (such as a laptop) connected via the GPIO connector. See section 15.1.3.1 on page 256 for more details;
- “Ring buffer on flash module” - some systems are equipped with an external flash module to extend the storage capacity beyond that available on the normal file-system. Data written here are accessible via the web interface or, from the command line, under the path `/media/flash_module`;
- “Record files under /var/spool/recdata” - this is the only option on CMG-NAMs and CMG-NAM64s and it is not available on other platforms. Data are written directly to the internal hard drive and are accessible via the web interface or, from the command line, under the path `/var/spool/recdata`.
- “Remote NFS mount” - this option causes data to be written to a remote, network-accessible storage device that uses NFS as the underlying protocol, such as NAS units. See section 11.2.1.5 on page 144 for configuration details.

Recycle space: If this check-box is set to zero, `gdi-record` will write to the storage device until it is full and then stop writing. The module will check periodically for free space and start writing again when it can. This prioritises the earliest data at the expense of the most recent data.

If the “recycle space” field is set to any other numerical value, it is interpreted as the amount of free space, in MebiBytes, to leave on the disk. When this value is approached, older data are deleted to make space for the current data. This prioritises the most recent data at the expense of the earliest data.

Format filesystem: This drop-down menu allows you to select the format used when the storage device is reformatted. The current choices are FAT32 and ext4. The ext4 filesystem is more robust and should be used where possible but it is not currently supported by Microsoft windows. If you need to read the storage device under windows (for example, by using the GPIO/USB cable), then you should choose FAT32.

Continuous power: In order to reduce power consumption, `gdi-record` does not write continuously to the hard drive. Data are buffered in flash memory and, at a configurable interval, these data are checked to see which complete files can be written to a mass storage device. The storage device is powered down when not in use. If power is not a consideration, certain options (such as retrieving data to remote systems using `scp`) are easier if the power to the storage device is left on continuously. Tick this check-box to enable continuous power to the storage device.

11.2.1.2 MiniSEED

For some applications, it is more convenient to store the data directly in mini-SEED format. The second tab of the "Storage and recording" page controls options related to recording in this format. If recording in mini-SEED format is enabled, a GDI to Mini-SEED compressor (converter) is started automatically. See section 12.2 on page 179 for more details.

Selecting the "MiniSEED" tab displays the following screen:

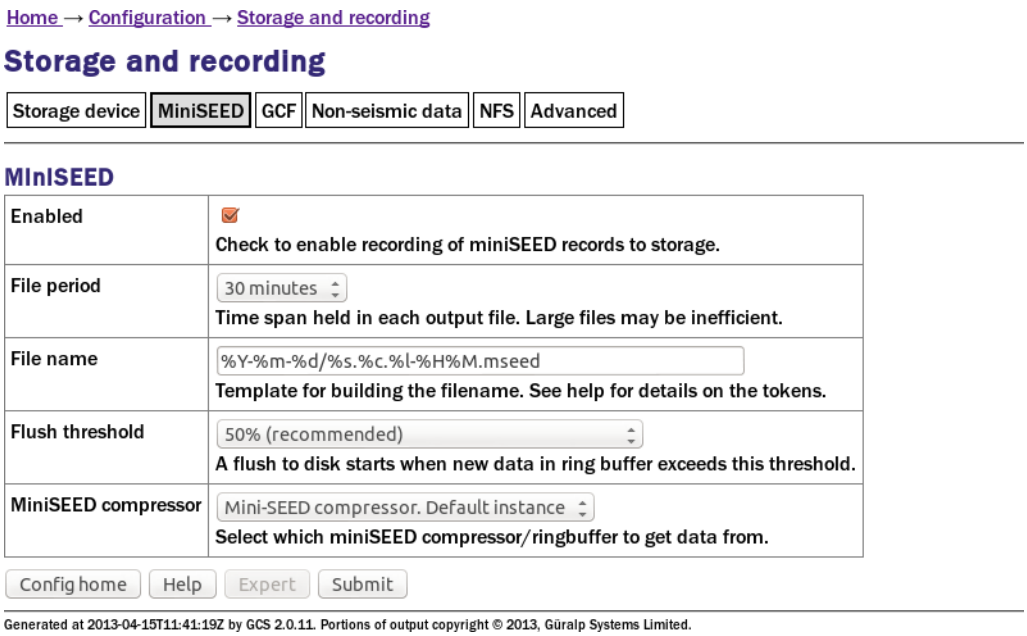

Enabled: MiniSEED recording will be enabled if this check-box is ticked and disabled otherwise.

File period: This drop-down menu offers a choice of file periods: the options are 15 or 30 minutes and 1, 2, 3, 4, 6, 12 or 24 hours. If set, for example, to 15 minutes, four files will be generated each hour whereas, if set to 24 hours,

only one file will be generated per day. The files may contain a single component or multiple components: see the next field description for details. The default file period is 30 minutes.

File name: The contents of this field determine how the data files are named. Tokens within the field are replaced by parameters derived from the data while all other characters are coped verbatim into the resulting file-names. It is possible to organise files into hierarchical directories by using forward-slash characters ('/') in this field. Please see section 11.2.2 on page 146 for details of the tokens that can be used.

Note that if this field contains any token that identifies a stream, individual files will be created for each stream. If no such token is present, multiple streams will be recorded into each file.

Flush threshold: This drop-down menu allows control over the percentage utilisation of the miniSEED ring-buffer which triggers a flush to the mass storage device. The default is 50% but you can also choose 10%, 25% or 75%.

MiniSEED compressor: This drop-down menu allows the operator to select which instance of `gdi2miniseed` is used as the data source. The menu will offer all configured instances.

11.2.1.3 GCF

GCF is the native recording format of Platinum firmware. It can be read directly by Scream and many software packages are available (from the GSL web-site) for converting it into other formats.

Selecting the "GCF" tab displays the following screen:

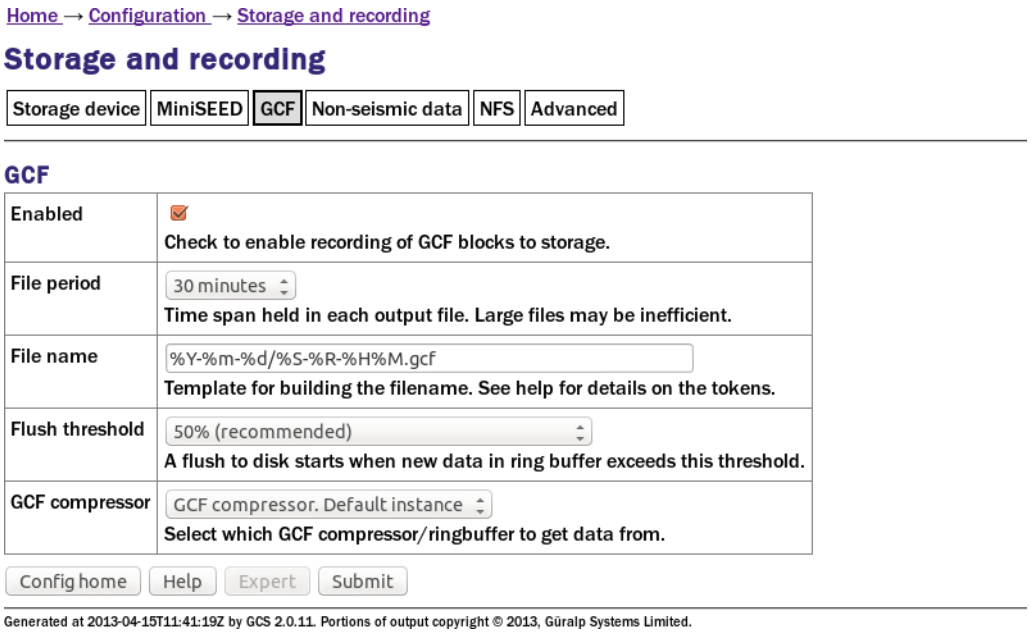

Enabled: GCF recording will be enabled if this check-box is ticked and disabled otherwise.

File period: This drop-down menu offers a choice of file periods: the options are 15 or 30 minutes and 1, 2, 3, 4, 6, 12 or 24 hours. If set, for example, to 15 minutes, four files will be generated each hour whereas, if set to 24 hours, only one file will be generated per day. The files may contain a single component or multiple components: see the next field description for details. The default file period is 30 minutes.

File name: The contents of this field determine how the data files are named. Tokens within the field are replaced by parameters derived from the data while all other characters are coped verbatim into the resulting file-names. It is possible to organise files into hierarchical directories by using forward-slash characters ('/') in this field. Please see section 11.2.2 on page 146 for details of the tokens that can be used.

Note that if this field contains any token that identifies a stream, individual files will be created for each stream. If no such token is present, multiple streams will be recorded into each file.

Flush threshold: This drop-down menu allows control over the percentage utilisation of the GCF ring-buffer which triggers a flush to the mass storage device. The default is 50% but you can also choose 10%, 25% or 75%.

GCF compressor: This drop-down menu allows the operator to select which instance of `gdi2gcf` is used as the data source. The menu will offer all configured instances.

11.2.1.4 Non-seismic data

Selecting the “Non-seismic data” tab displays the following screen:

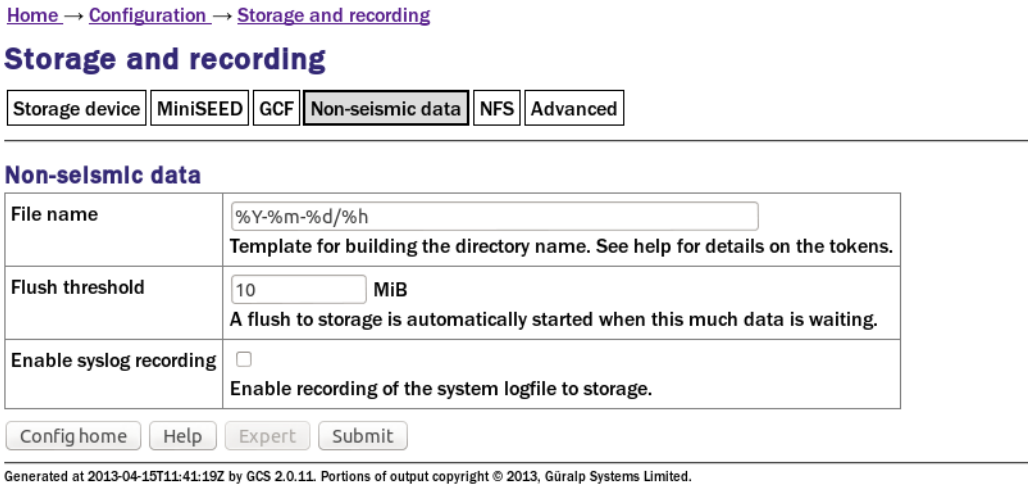

This facility can be used to copy data files from your own applications and, optionally, the system log files to the mass storage device, simplifying transfer of these data from the unit to external systems.

Files to be copied should be created in the directory `/var/spool/to-rdisk`.

The following options are available:

File name: When building the directory in to which files are copied, any tokens in the template string are replaced with strings based on the current date/time. The following tokens are supported:

%%	a literal percentage sign
%h	system hostname
%Y	4-digit year number
%m	2-digit month number (01-12)
%d	2-digit day of month (01-31)
%j	3-digit ordinal day/day of year (001-366)
%H	2-digit hour (00-23)
%M	2-digit minute (00-59)
%I	4-digit ISO year number (for use with ISO week)
%W	2-digit ISO week number (01-52)
%w	1-digit ISO day of week (1/Monday-7/Sunday)

Forward-slashes (*'/'*) will cause subdirectories to be created. If using recycle mode, be sure to start the file-name template with the date in either year-month-day, year-ordinal or year-week-weekday format so that the sorting algorithm can identify the oldest files correctly.

Flush threshold: Files from `/var/spool/to-rdisk` are copied to the mass storage device when their total size exceeds the threshold size defined (in Mebibytes) in this field.

Enable syslog recording: If this checkbox is ticked, the system log (`/var/log/messages`) is also copied to the mass storage device. This option consumes very little space and can be a significant help when debugging any problems.

11.2.1.5 NFS

Platinum systems can write to external, network-accessible NFS storage systems, such as NAS systems. This option is primarily intended for use with CMG-NAMs and CMG-NAM64s in data-centre applications. GSL can supply suitable NAS systems for this purpose.

To use external NFS storage, select “Remote NFS mount” in the “Storage device” drop-down menu of the “Storage device” tab (see section 11.2.1.1 on page 140) and complete the fields in the “NFS” tab as described below.

Selecting the NFS tab causes the following screen to be displayed:

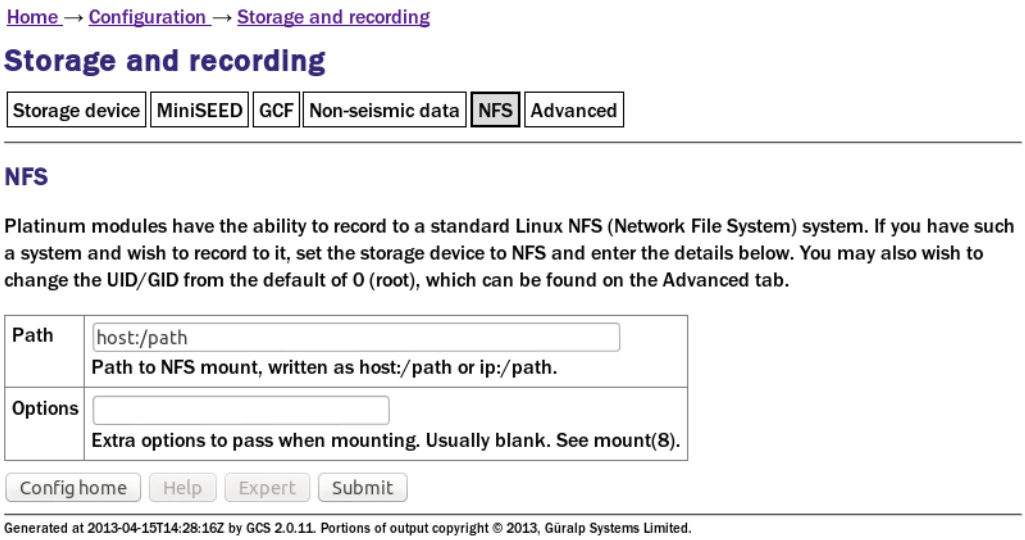

Path: Use this field to specify the NFS host and the path on that host to use for storing data. For example, if the data are to be stored in a directory called `GCF-data` on a NAS host called `nas0.datacentre.example.com`, enter

```
nas0.datacentre.example.com:/GCF-data
```

in this field.

Options: The external storage provided by the NFS server will be prepared for access using the standard Linux mount command (see http://linux.die.net/man/8/mount for details) invoked without non-mandatory options. If additional options are required (for example, to specify a hard mount rather than a soft mount), they can be entered here.

11.2.1.6 Advanced

The "Advanced" tab of the "Storage and recording" screen contains a number of specialised options. Selecting the tab displays the following screen:

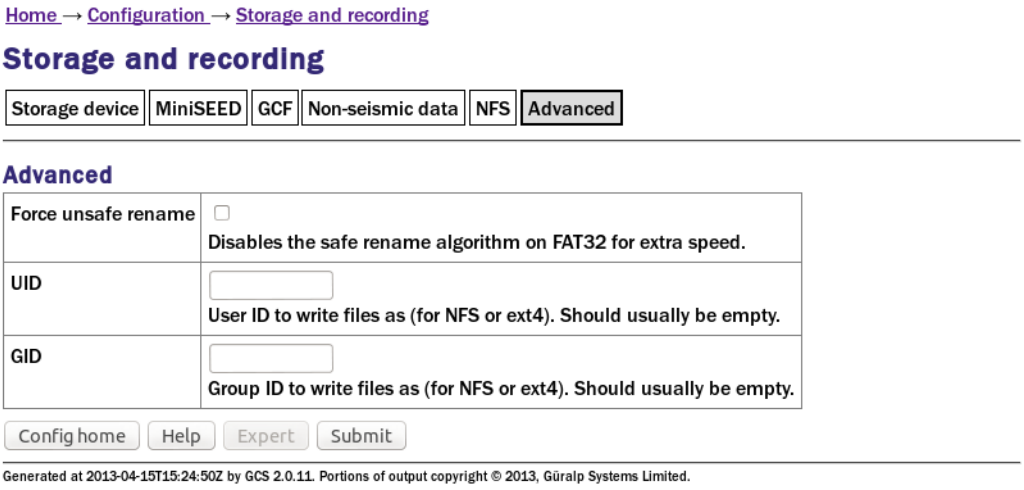

Force unsafe rename: The recording system takes considerable care to avoid data loss in the event of power disconnection or storage device removal. When the storage device is formatted as a FAT32 volume, part of this care involves a "safe rename" routine, which takes a little time to complete. When the power budget is particularly tight, this extra time costs extra power as the storage device needs to be powered up for longer. Ticking this check-box disables the safe rename routine, reducing the amount of time for which the storage device needs to be powered at the cost of increased vulnerability to power outages or unexpected storage device disconnection.

UID and **GID**: When recording to an NFS or ext4 volume, files are normally written as user "root" and group "root". If this is undesirable, another user ID and/or group ID can be specified using these two fields.

11.2.2 File name escape sequences

Escape sequences (tokens) are used in the "File name" fields in the MiniSEED and GCF recording configuration pages. The escape sequences all begin with a percent character ('%') and are used to insert variable data such as the date or stream name into the file or directory name; each escape sequence is replaced with the relevant value. Any non-escape sequence characters are copied verbatim into the name, as are unrecognised escape sequences. All numbers are decimal and will have leading zeroes added to fill the required number of digits, where appropriate.

The available escape sequences are listed below. Items marked † only have values in GCF context and will be replaced by the relevant number of spaces if used with miniSEED. Items marked ‡ only have values in miniSEED context:

%d	2 digit day of month (01-31)
%H	2 digit hour in 24 hour clock (00-23)
%j	3 digit Julian day
%m	2 digit month (01-12)
%M	2 digit minute (00-59)
%y	2 digit year i.e. without century digits (00-99)
%Y	4 digit year
%S	GCF System ID †
%C	GCF Stream ID †
%s	5 char SEED station identifier (spaces are removed from all SEED Ids) ‡
%c	3 char SEED channel identifier ‡
%n	2 char SEED network identifier ‡

%l	2 char SEED location identifier ‡
%b	block length
%f	bits/sample or compression format
%p	samples/second
%%	% (a literal percent sign)

If the format string ends in a `.extension` (without any escape sequences in the extension) then this extension will be noted and used in some other locations – e.g. for the top level date directory.

The default format strings are:

- GCF Directory format: `%Y%j-gcfraw`
- GCF File format: `%Y%jT%H%MZ.gcf`
- miniSEED Directory format: `%Y%j-mseed`
- miniSEED File format: `%Y%j-%H%M-%s-%c-%n-%l.mseed`

Slashes "/" will cause subdirectories to be created. Using them as date separators will have unintended and, usually, undesirable consequences.

11.2.2.1 Some examples

The GCF default, `%Y%jT%H%MZ.gcf`, includes the data and time but no stream identifiers, resulting in a single file containing all streams per recording period. This will produce file-names like:

```
2008315T1445Z.gcf
```

If you are recording at two different sample rates, including a `%p` will cause all streams at the same rate to be grouped into one file per rate per recording period. Using `%Y%jT%H%MZ_%p.gcf` might produce files like:

```
2008315T1445Z_200.gcf and
2008315T1445Z_50.gcf
```

If you need one file per stream, include the `%C` sequence. Using the format `%Y%jT%H%MZ_%C.gcf` would produce files like:

```
2008315T1445Z_406N2.gcf
2008315T1445Z_406E2.gcf and
2008315T1445Z_406Z2.gcf
```

Adding a %S would add the system ID, so `%Y%jT%H%MZ_%S_%C.gcf` would produce files like:

```
2008315T1445Z_EKA_406N2.gcf
```

The miniSEED default `%Y%jT%H%MZ-%s-%c-%n-%l.mseed` includes the date, the time and the complete SCNL identification. The date format matches that used by the GCF recorder. This will produce file-names like:

```
2008315T1442Z-TEST1-BHE-NN-LL.mseed
```

To combine all the channels from a given station simply omit the channel marker from the file name format string:

```
%Y%jT%H%MZ-%s--%n-%l.mseed
```

It is recommended that the "--" is left in place to highlight the omitted channel id. This will produce file-names like:

```
2008315T1442Z-TEST1--NN-LL.mseed
```

If you specifically want to include a marker to identify that it contains all channels, the use of a lower case string will differentiate it from a regular channel name, which is always presented in upper case.

```
%Y%jT%H%MZ-%s-all-%n-%l.mseed
```

yields file-names like:

```
2008315T1442Z-TEST1-all-NN-LL.mseed
```

If you prefer human readable dates, rather than using the Julian date

```
%Y_%m_%d-%H:%M-%s-%c-%n-%l.mseed
```

yields file-names like:

```
2008_08_14-14:42-TEST1-BHE-NN-LL.mseed
```

Note: Using / as a date separator will split the data into sub-directories, which may not be the desired result.

It is often required to separate the data into sub directories by network and station prefix. In this case, it is recommended that the network and station id are still included in the file-name so that the contents of the file are still recognisable even if it is moved to a different location.

```
%n_%s/%Y%jT%H%MZ-%s-%c-%n-%l.mseed
```

will store the data like this:

```
NN_TEST1
          2008315T1442Z-TEST1-BHE-NN-LL.mseed
          2008315T1452Z-TEST1-BHE-NN-LL.mseed
          2008315T1502Z-TEST1-BHE-NN-LL.mseed
          ...
NN_TEST2
          2008315T1442Z-TEST2-BHE-NN-LL.mseed
```

```
2008315T1452Z-TEST2-BHE-NN-LL.mseed
2008315T1502Z-TEST2-BHE-NN-LL.mseed
...
```

11.3 Retrieving data

Data are recorded first to buffers held in flash memory. There are separate buffers for GCF and miniSEED data and their sizes are defined in the configuration pages for their respective compressors (See section 12.1.1 on page 162 for the GCF compressor and section 12.2.1 on page 180 for the miniSEED compressor). When the buffer utilisation exceeds a configurable percentage (as specified in the relevant tabs of the **Configuration → Storage and recording page**), a process is triggered to flush the data to the hard drive. In low power applications, the hard drive will normally be powered down and, so, must be powered up and mounted before use. It is then dismounted and powered down once the flush is complete.

If you wish to work with data on the removable drive, it must first be powered up. This is done automatically when using the web interface but must be done manually when working from the command line. If you wish to work with recent data, a manual flush should first be performed in order to move the data from the buffer memory to the drive.

Facilities exist to aid automation of data downloads. See section 11.3.1.3 on page 156 for details.

11.3.1 Retrieving data from the removable drive

Note: Older removable hard-drives had internal heaters and temperature sensors. It can take several seconds to pre-heat and power up one of these drives. Be prepared for short delays when using some of the following commands with these drives.

Data from the removable drive can be retrieved using the web interface, using network file transfer tools or over a serial port.

- The web interface is most convenient if you only want to download one or two files. It is not suitable for large numbers of files or for automated downloads.
- Using the serial interface is slow and also not suitable for large numbers of files or for automated downloads.
- The use of network transfer tools, such as scp, sftp or rsync, is recommended in most cases. If large numbers of files are to be downloaded or if the process is to be automated, a special user can be created to simplify the process. When this user logs in over ssh (as used by rsync), scp or sftp, the storage is mounted automatically and

dismounted when the user logs out again. See section 11.3.1.3 on page 156 for details.

11.3.1.1 Downloading over a network, using the web interface

To retrieve data from the removable drive using the web interface, select:

Tools → Storage and recording

The following screen is displayed:

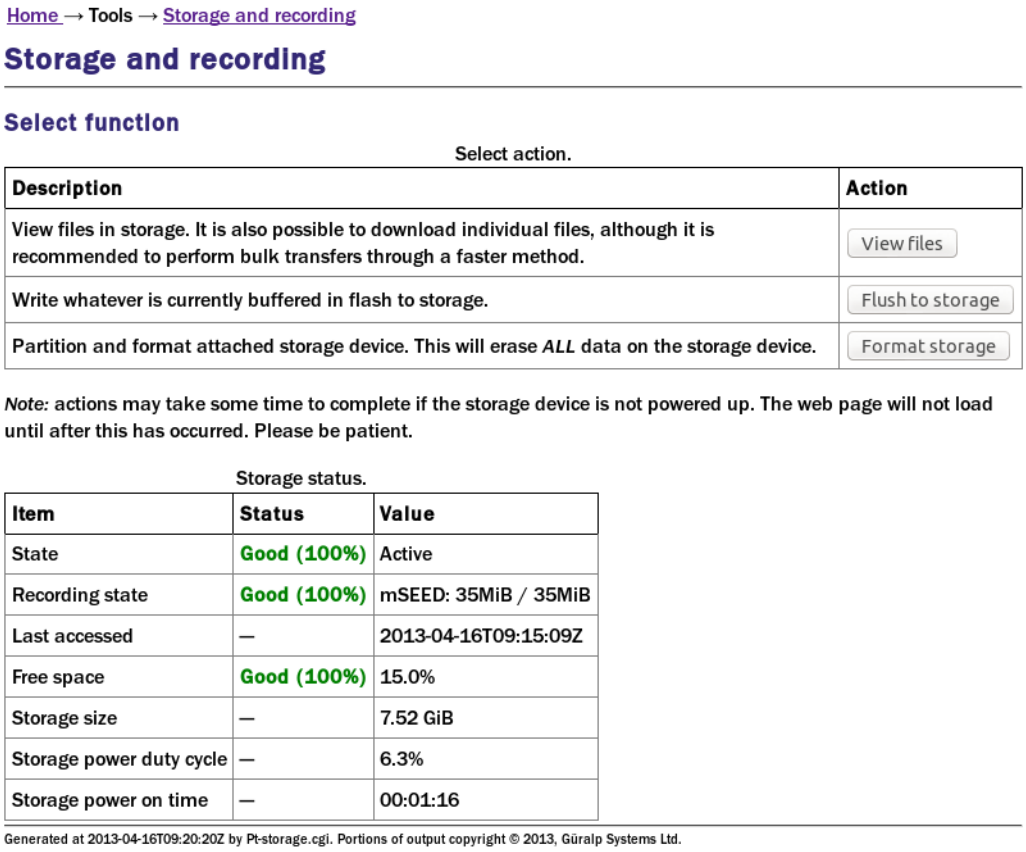

Home → Tools → Storage and recording

Storage and recording

Select function

Select action.

Description	Action
View files in storage. It is also possible to download individual files, although it is recommended to perform bulk transfers through a faster method.	View files
Write whatever is currently buffered in flash to storage.	Flush to storage
Partition and format attached storage device. This will erase *ALL* data on the storage device.	Format storage

Note: actions may take some time to complete if the storage device is not powered up. The web page will not load until after this has occurred. Please be patient.

Storage status.

Item	Status	Value
State	Good (100%)	Active
Recording state	Good (100%)	mSEED: 35MiB / 35MiB
Last accessed	—	2013-04-16T09:15:09Z
Free space	Good (100%)	15.0%
Storage size	—	7.52 GiB
Storage power duty cycle	—	6.3%
Storage power on time	—	00:01:16

Generated at 2013-04-16T09:20:20Z by Pt-storage.cgi. Portions of output copyright © 2013, Güralp Systems Ltd.

The top of the screen offers various action buttons while the bottom half displays some status information.

In the example given above, one can see that the storage is currently mounted and in use (“Active”) and that a flush of MiniSEED data is almost complete (“35MiB/35MiB”). The size and utilisation of the device is also given.

The “storage power duty cycle” shown is the ratio of the total time during which the storage has been powered up to the system uptime (the time since the last boot) expressed as a percentage. In power-critical applications, this provides useful feedback when tuning the ring-buffer sizes and fill thresholds (see section 11.2.1 on page 139 for details of these parameters).

If you wish to retrieve the most recent data, first click the Flush to storage button to copy all pending data from the buffer memory to the hard drive. A

progress screen will display messages as the various stages of the process complete.

Once the flush process has completed, as shown above, return to the main disk menu by clicking on “Removable disk” on the “Tools” menu.

You can now click the View files button. This will power up any connected mass storage devices and, after a short delay, present a list of attached devices and their details (filesystem, free space, etc):

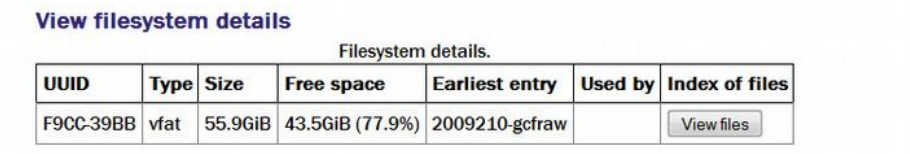

Clicking on any of the View files buttons takes you to the “Removable disk file index” screen, which displays folders and files within the selected filesystem. Subdirectories (folders) have a Follow button next to them and files have a Download button.

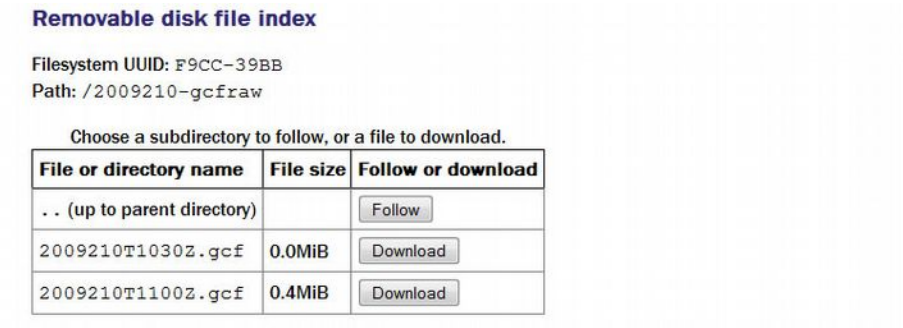

If the displayed directory contained subdirectories, you could continue to navigate down them using the Follow buttons. When files are present, as in the above screen-shot, they each have an associated Download button which invokes your web browser's standard download facility to copy the data to your computer.

11.3.1.2 Downloading over a network, using the command line

If very recent data are required, start by flushing the buffers with the command:

```
eam999 ~ # gdi-record --flush
```

Before data can be retrieved from the removable drive, it must first be brought to operating temperature (for heated drives), powered up and the relevant file-systems mounted. A utility, `Pt-storage`, is provided to accomplish this.

When invoked without arguments, `Pt-storage` displays a menu:

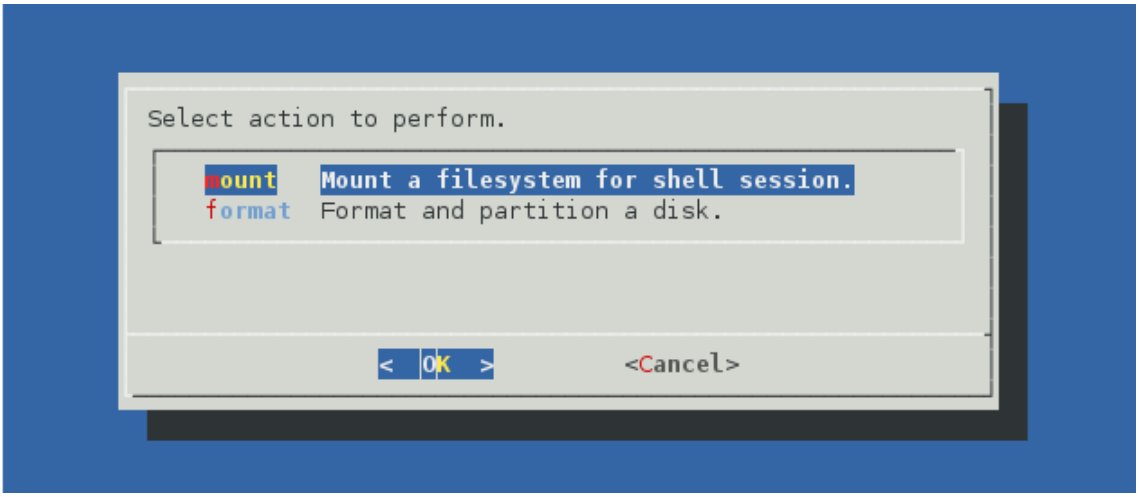

The “mount” option, which can be selected with the M key, displays a menu of available file-systems on removable media. When you choose your desired file-system, it is mounted under `/media` and you will be told the exact mount point and then presented with a command prompt. This is a sub-shell: the file-system will remain mounted and the mass storage device will remain

powered up until you exit the shell (with the `exit` command or by keying Ctrl + D).

The “format” option, which can be selected with the F key, prepares a drive for use, as described in section 11.1 on page 137.

When using shell scripts, all of these menu functions can be accessed by passing the function name as an argument to `Pt-storage`. For example:

```
eam999 ~ # Pt-storage --mount
```

performs the same function as selecting ”mount” from the Pt-storage menu.

Note: Older removable hard-drives had internal heaters and temperature sensors. It can take several seconds to pre-heat and power up one of these drives. Be prepared for short delays when using some of the following commands with these drives.

Once the mass storage device is powered up and the relevant file-system mounted, the recorded files can be browsed with standard Linux shell commands such as cd and `ls`. They can be copied to a remote PC using the network or over the serial port (as described in section 11.3.1.4 on page 157).

To copy files over the network, the use of scp or rsync is recommended. The scp program is most convenient to use and can copy single files or recursively copy directories. The rsync program is more complicated but is ideal when a remote copy of the data is to be updated regularly, since it minimises the traffic over the network by only copying new or changed files. The same rsync syntax can be reused regardless of changing filenames.

For Linux users, scp is installed by default or available as an optional package in most distributions. Consult your operating system documentation for more details. For Windows users, the WinSCP package is recommended: this has the additional benefit of providing a graphical, explorer-like interface for browsing files on the acquisition module. WinSCP can be downloaded for free from http://winscp.net.

The screen-grab below shows a complete session recorded from a Linux PC. The operator connects to a acquisition module, powers up the mass storage devices, downloads all recorded files, powers down the mass storage devices and then disconnects.

```
fish@fish-desktop:~/tmp$ ssh root@51.187.130.165
dcm105 ~ # rdisk mount
Connected to server. Powering up disks and requesting details...
```

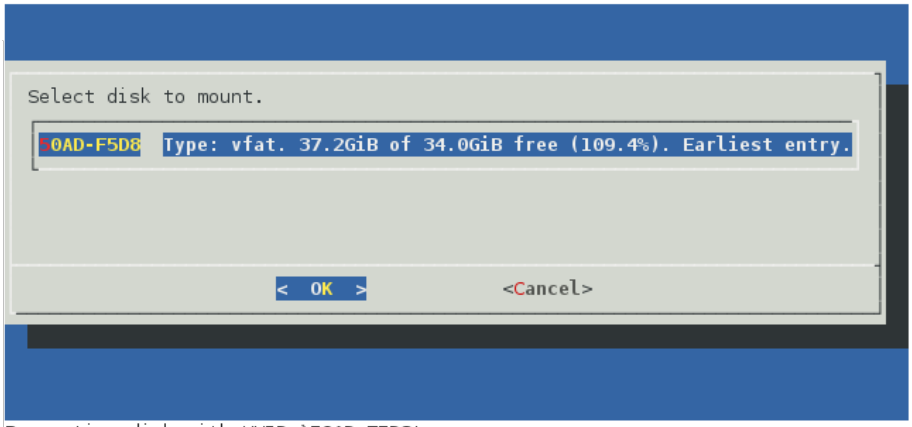


```
Requesting disk with UUID `50AD-F5D8'...
The disk is mounted at: /media/50AD-F5D8
A new bash session has been created.

When you exit from this bash session (which you can do with `exit' or Ctrl-D),
the disk will be unmounted. It will remain mounted (and thus powered up) until
you exit. Then you will be returned to your original shell.
dcm105 ~ # ~^Z [suspend ssh]

[1]+  Stopped                 ssh root@51.187.130.165
fish@fish-desktop:~/tmp$ scp -r root@51.187.130.165:/media/50AD-F5D8 .
2009336T1030Z.gcf                             100% 1886KB 314.3KB/s   00:06
2009336T1100Z.gcf                             100% 1781KB 356.2KB/s   00:05
2009336T1130Z.gcf                             100% 1910KB 382.0KB/s   00:05
2009336T1200Z.gcf                             100% 1928KB 385.6KB/s   00:05
2009336T1230Z.gcf                             100% 1957KB 326.2KB/s   00:06
2009336T1300Z.gcf                             100% 1889KB 377.8KB/s   00:05
fish@fish-desktop:~/tmp$ fg
ssh root@51.187.130.165

dcm105 ~ # exit

Disk no longer required by this shell session.

dcm105 ~ # logout
Connection to 51.187.130.165 closed.
fish@fish-desktop:~/tmp$ fg
```

Note the use of ssh's tilde ('~') escape followed by Ctrl + Z to suspend the ssh session and return to the calling PC, in order to run the scp command with the mass storage device still mounted. The `fg` command (foreground) returns control to the ssh session.

Windows users should follow the same procedure to log in, power up and mount the file system. At this point, rather than suspending the ssh shell, they can connect with WinSCP and navigate to the `/media` directory. The recorded files will then be displayed and can be dragged and dropped to appropriate locations on the PC:

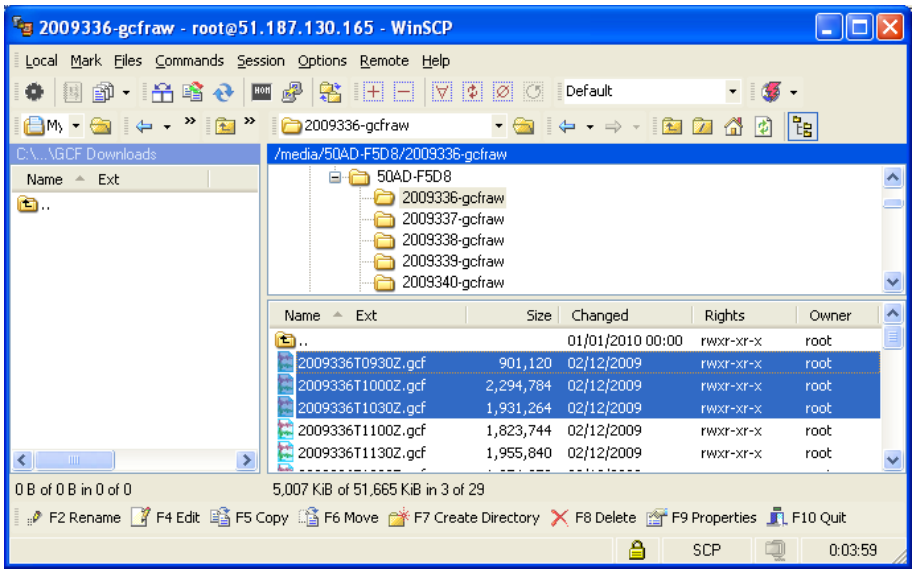

When the transfer is complete, return to the ssh session and power down the mass storage devices (with the **`exit`** command or by keying Ctrl + D).

Using rsync is very similar: simply replace the invocation of scp in the above instructions with an appropriate rsync command.

For Linux users, the simplest usage is

```
me@mypc:~/dl $ rsync -avz root@51.187.130.165:/media/*/
```

This will copy all files from all removable drives to the current directory on the invoking computer. Only the differences are transferred, making this particularly efficient when used regularly. For more advanced usage, please see the rsync manual, available on-line at http://man-wiki.net/index.php/1:rsync-2006.11.06

Windows users can download a free port of rsync using the cygwin library (see http://www.cygwin.com/ for more information) or use one of several free, graphical interfaces, such as grsync, available from http://sourceforge.net/projects/grsync-win/files

The following screen-shot shows grsync about to download all data from a removable drive. Note the password prompt appearing in the separate console window.

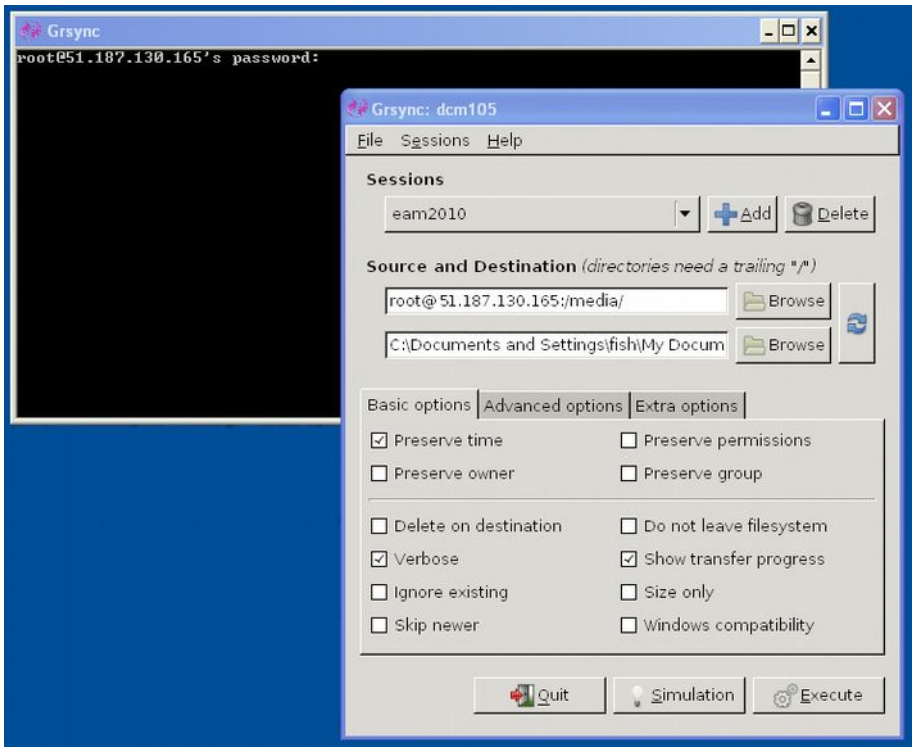

11.3.1.3 Automating network downloads

If data are to be repeatedly downloaded from a Platinum system, it is possible to automate the mounting and dismounting of the storage device, which considerably simplifies the operation. This is accomplished by means of a special user identity: the special user is configured so that the storage device is mounted automatically when the user logs in and dismounted automatically when the user logs out. Network file transfer protocols such as scp, sftp and rsync all involve a hidden login so, as long as they authenticate as the special user, they can be invoked at any time without the need for a command-line session to handle device operations.

To create the special user, run the command `Pt-storage-adduser` followed by the desired username – this command will prompt interactively for a password for the new user. The password will not be echoed to the screen.

Note: If the system is to be connected permanently to the Internet, it is necessary to choose a strong password. Since the password will almost certainly be programmed into the client file transfer software, there is no requirement for it to be memorable. GSL recommend the use of password-generation software for this purpose. There are a number of good, free password-generating applications and web-sites available on the Internet.

A typical invocation looks like this:

```
eam999# Pt-storage-adduser dskusr
New password:
Repeat password:
Password changed
New user 'dskusr' added and ready
eam999#
```

The chosen username, “dskusr” in this example, can be anything you wish but must consist of lower-case letters and digits only.

The new user can now be used with file transfer commands. For example, from a Linux box, you could enter commands like

```
me@mypc:~/dl $ ssh dskusr@eam999 gdi-record --flush
me@mypc:~/dl $ rsync -avz dskusr@eam999:/media/*/
```

This will flush outstanding data from the ring-buffers and then synchronise all files from the mass storage device to the current directory on the invoking computer.

It is also possible to eliminate the password prompt. To do this:

- Run the following commands on the Platinum system:
  `mkdir -m 700 ~dskusr/.ssh > ~dskusr/.ssh/authorized_keys`
  `chmod 600 ~dskusr/.ssh/authorized_keys`

  replacing *`dskusr`* with the name of your own special user.
- create an ssh key pair on the client PC, using the `ssh-keygen` command on Linux or the PuTTYgen application on windows.
- Copy the newly created public key into the authorized_keys file.

The new user should now be able to invoke ssh, scp, sftp or rsync without being prompted for a password.

Once you are happy that password-less logins are working, you can remove the password for the new user by editing `/etc/shadow` on the Platinum system and replacing the encrypted password – the part of the line between the first and second colons – with an exclamation mark ('**!**').

11.3.1.4 Downloading over a serial port, using the command line

In situations where it is not convenient to use the network interface, files can be downloaded from the removable mass storage device using one of three file-transfer protocols, X-modem, Y-modem or Z-modem. None of these protocols were ever rigidly standardised so, if you are not using one of the terminal emulators discussed in this section, you may need to experiment a little: this is reflected in the huge range of optional arguments that the transfer program accepts.

The X-modem transmitter is invoked as **sx**, the Y-modem transmitter as **sy** and the Z-modem transmitter as **sz**. They are, in fact, all implemented by the same program so the detailed help message (displayed with the **--help** option) describes options relevant to all three protocols.

Linux users are advised to use the minicom terminal emulator. This includes X-modem support and its use is described below. For Windows users, we describe the use of HyperTerminal, which is supplied with many Windows systems.

For minicom users, the X-modem protocol should first be configured: start minicom and type Ctrl + A (minicom's escape sequence) and then O to display the options menu:

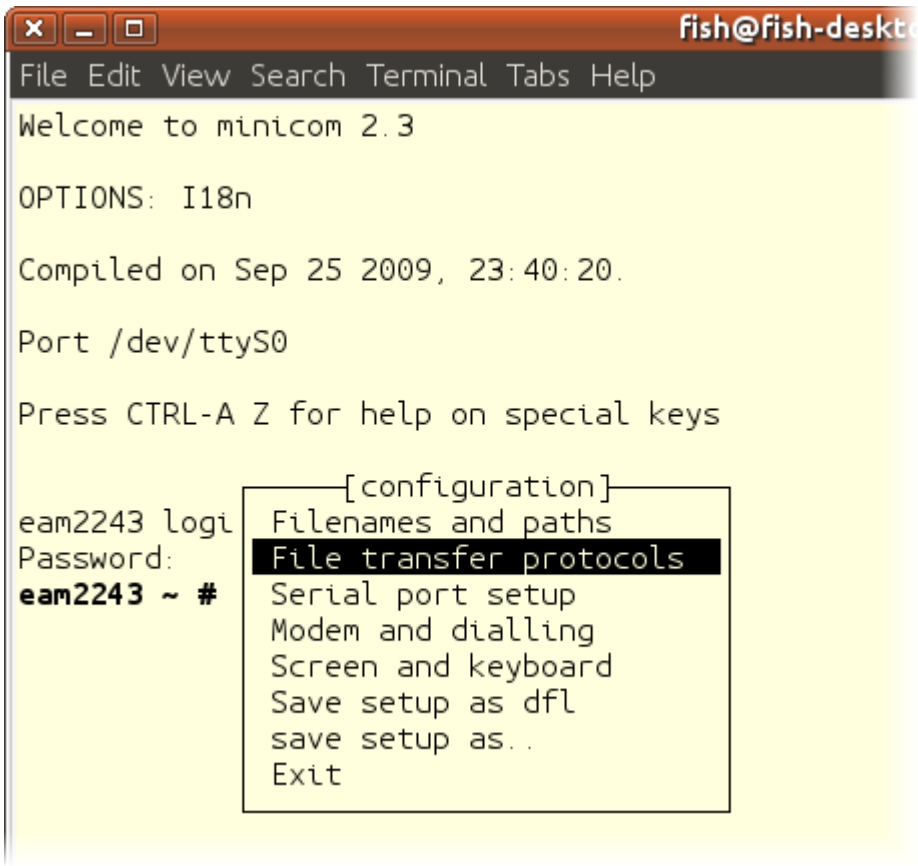

Select “File transfer protocols” and ensure that the command used for X-modem transfers is set to **/usr/bin/sx -vv**

Before attempting to download files, you should flush the recording buffers and mount the mass storage device. The buffers can be flushed with the command:

```
eam2010 ~ # gdi-record --flush
```

and the mass storage device can be mounted and held in a powered-up state by use of the **Pt-storage** facility, as described in section 11.3.1.2 on page 152.

Once the mass storage device is mounted, you can proceed to download files by entering commands like

```
eam2010 ~ # sx path-to-file
```

and then immediately typing Ctrl + A followed by R to activate minicom's receive file function. Select “xmodem” from the resulting menu and then enter a name for the downloaded file. The transfer should start immediately with a progress indicator displayed:

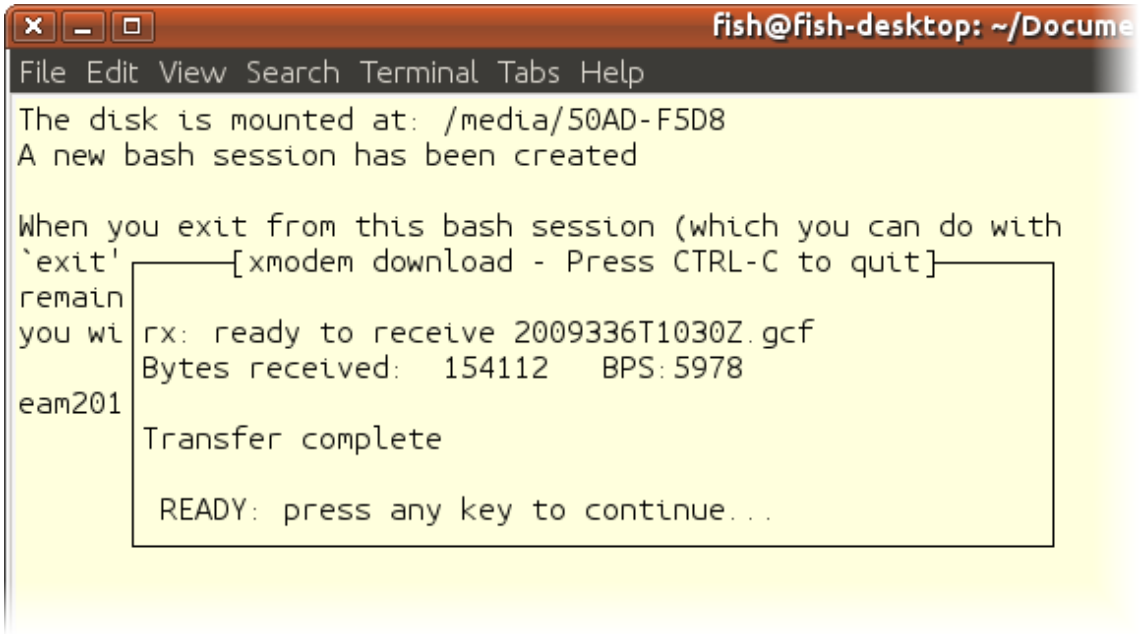

HyperTerminal users should flush the buffers, mount the mass storage device, type the command

```
eam2010 ~ # sx path-to-file
```

and then immediately select “Receive File...” from the “Transfer” menu. A dialogue asking for the destination directory name is followed by another asking for the destination file name:

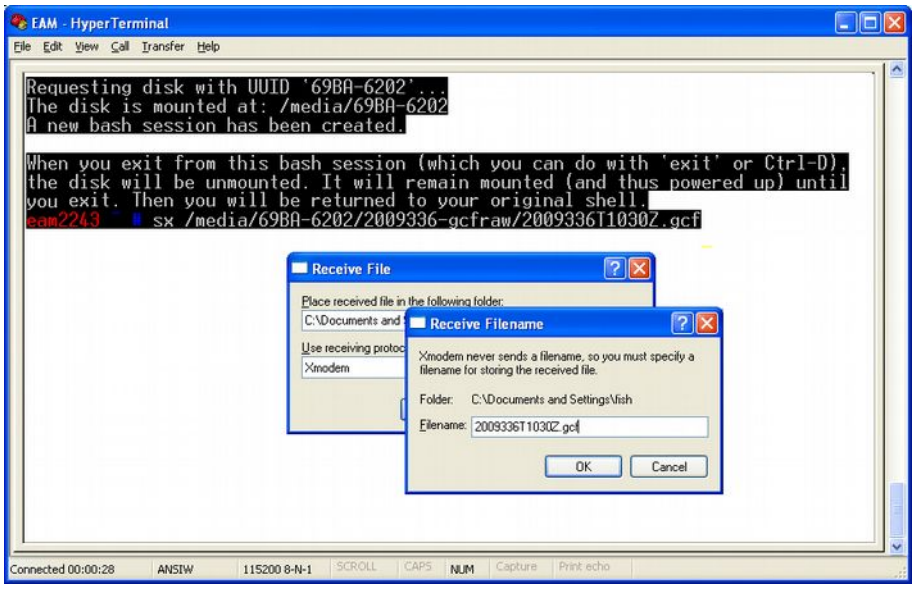

When you click the “OK” button, a progress indicator appears:

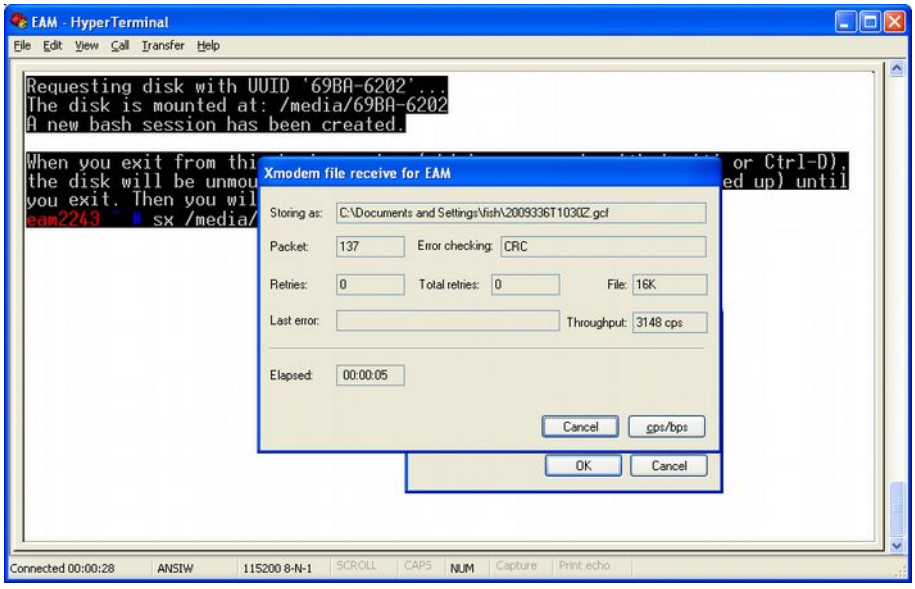

You will be returned to the prompt when the transfer is complete.

11.3.2 Reading the removable drive on other computers

The mass storage device may be disconnected from the system at any time without risk of data loss. In practice, users will normally manually flush the memory contents to mass storage device (using either the `gdi-record --flush` command or the Flush to storage button on the “Removable disk” page of the web interface) and allow that process to complete before removing the drive.

The drive can then be connected to any computer that supports external USB or FireWire storage devices. In some circumstances, you may need to provide a power connection to the drive: see section 11.1 on page 137 for more details.

The mass storage device may have been formatted as either a VFAT or ext3 volume. Windows computers cannot read ext3 volumes without additional software such as "Explore2fs" or "DiskInternals Linux Reader". See http://www.howtoforge.com/access-linux-partitions-from-windows for more details about these two packages.

11.3.3 Accessing internal storage directly

Certain systems, such as integrated instruments and cylindrical digitisers are fitted with an internal Flash memory device which is accessible via USB. It can be written to by selecting "Internal USB storage" from the "Recording destination" drop-down menu on the "Disk recording" page (see section 11.2 on page 138).

When a USB host, such as a laptop or PC, is connected to the GPIO port (the pin-out of which is given in section 16.5.6 on page 278) internal circuitry detects the USB power and automatically connects the Flash memory directly to the GPIO socket, causing it to behave identically to a standard USB memory stick. This method can be used even if no power is supplied to the EAM.

When no power is detected at the GPIO port, the Flash memory is available to the system as if it were a standard removable mass storage device. All of the mass storage device recording options described above (in section 11.2 on page 138) will apply to this device, as will the other data access options described in section 11.3.1 on page 149.

12 Transmitting Data

Received data may be re-transmitted in near real time in one or more of a number of different formats. By default, a GCF Scream server and GDI Link transmitter are instantiated, both configured to forward all received data. Any other desired transmitters must be configured and enabled before use.

The following transmission services are currently available:

- GCF BRP serial server - see section 12.1.2 on page 166
- GCF BRP network server - see section 12.1.3 on page 169
- GCF Scream network server - see section 12.1.4 on page 174
- SEEDlink - see section 12.2 on page 179
- EarthWorm – see section 12.3 on page 186
- CD1.1 - this is covered in a separate manual, MAN-EAM-1100.
- GSMS: Güralp Seismic Monitoring System - see section 12.4 on page 192
- QSCD: Quick Seismic Characteristic Data - see section 12.5 on page 196
- WIN sender - see section 12.6 on page 198

12.1 GCF

Three GCF servers are available, a GCF BRP serial server, which can output data over any available serial port, a GCF BRP network server and a GCF Scream network server. All of these take their input from one or more GCF compressors. These servers are described in the next section.

12.1.1 The GCF compressor

The GCF compressor, `gdi2gcf`, converts from GDI to GCF format. It provides channel filtering, channel name mapping and data buffering for the gdi-record service, which writes GCF files to mass storage device.

One instance of `gdi2gcf` is present in the default configuration. Additional instances may be created as required. This will be necessary if you have different channel filtering requirements for, say, recording and transmitting or if you need different transmitters to send different sets of channels. The configuration page for every transmitter has, in the “Expert mode” options, a drop-down menu which allows the operator to select which compressor instance to use for its input. The `gdi-record` configuration page has a similar facility.

Instances of the GCF compressor are “dependant services”, meaning that they do not need to be (and should not be) configured to start automatically when

the system boots. They will be started whenever a client service, such as a connected transmitter, starts.

To configure a GCF compressor on the acquisition module, select

Configuration → Services → GCF

or

Configuration → All Options → System services → GCF

To configure an instance from the command line, start gconfig and select "System services" from the top level menu.

Select the "gdi2gcf - GCF compressor" link. The screen shows a list of all GCF compressor instances that have been configured:

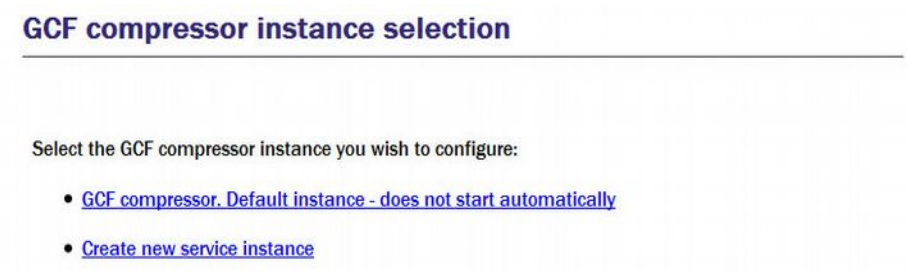

To configure any existing instance, click on its link. To configure an additional `gdi2gcf` instance, select "Create new service instance". The resulting screen allows you to configure the parameters of the selected instance.

12.1.1.1 Configurable parameters in simple mode

The configurable parameters for the GCF compressor in simple mode have two tabbed pages: General and Channels.

General

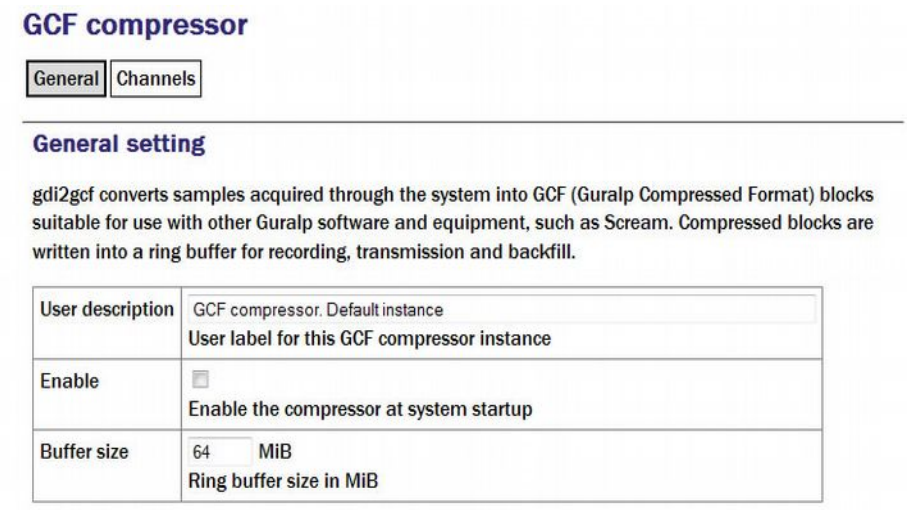

User description: Sets the name of the service. If multiple instances are created, this can be set to meaningful names for the data that each will be handling.

Enable: Causes this service to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

Buffer size: GCF data from the compressor are held in a ring buffer from which all client services, such as GCF transmitters or the mass storage device recorder, read. The number in the field determines both how long a communication link can be down before data are lost and the time interval between "flush to disk" operations (which can affect power consumption).

The next section of this page, shown overleaf, contains a drop-down menu and a table which allow the operator to control which channels are transmitted and, optionally, renaming them.

Channels

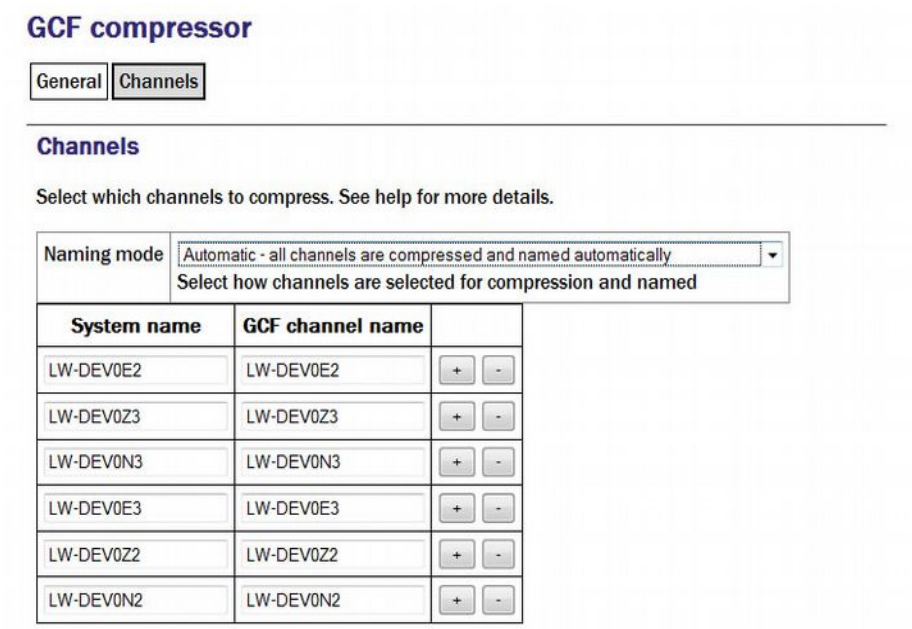

Naming mode: The drop-down menu offers three choices:

- **Automatic - all channels are compressed and named automatically:** Offers no filtering and uses the system-generated names for each channel as forwarded by `gdi-base`.
- **Semi-automatic - all channels are compressed, names may be mapped below:** One or more of the channels may be renamed by adding entries to the mapping table. If you wish to use this mode, it may be useful first

to run the system in automatic mode for a short while: this will populate the mapping table with an entry for each currently known channel, which can serve as the basis for your own mapping table.

- **Manual - only channels named below are compressed:** Offers both channel filtering and name mapping. If you wish to use this mode, it may be useful first to run the system in automatic mode for a short while to populate the mapping table with an entry for each currently known channel; this can serve as the basis for your own mapping table.

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

Deleted channels will be transmitted unmapped in "Semi-automatic" mode and not transmitted in "Manual" mode.

12.1.1.2 Configurable parameters in expert mode

The following additional parameters are available in expert mode:

Advanced

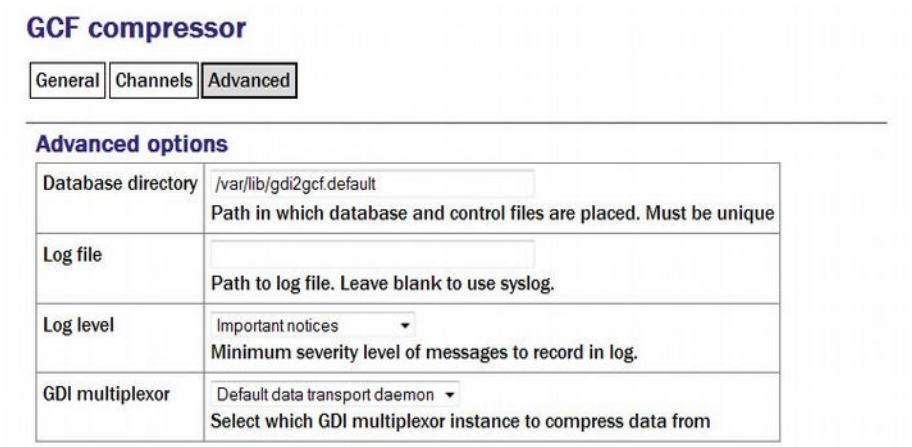

Database directory: Can be used to control the location of the ring-buffer and associated files. In most configurations, the default location is adequate but if, for example, a very large ring-buffer is desired and the optional extra flash memory module is fitted, it may be desirable to use the extra memory for this purpose. To do this, enter into this field the path to a unique directory under **`/media/flash_module`**.

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to **`/var/log/messages`**.

Log level: The drop-down menu controls the level of detail present in log messages. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

GDI multiplexer: In most configurations, all data for all compressors are taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this compressor. The menu offers a list of currently configured multiplexers.

12.1.2 GCF BRP Serial Server

The GCF BRP serial server transmits Güralp Compressed Format (GCF) data using the Block Recovery Protocol (BRP) over any available serial port.

To configure a GCF BRP serial server from the web interface, select:

Configuration → Serial ports

or

Configuration → All options → Serial ports

To configure a GCF BRP serial server from the command line, start gconfig and select **Serial ports** from the top-level menu.

Now select the serial port over which you wish to transmit GCF.

Port function: Select **GCF out. Outbound GCF data transmission** from the drop-down menu.

Port speed: Choose the required Baud rate from the drop-down menu.

Now select the **GCF output settings** link from the list at the bottom of the page. The next screen allows you to configure the GCF BRP serial server instance which will run on the previously selected port. The screen illustrated here is for an instance running on the DATA OUT port: the screens for other serial ports are practically identical.

12.1.2.1 Configurable parameters in simple mode

The configurable parameters for the GCF output settings in simple mode have two tabbed pages: General and Filtering.

General

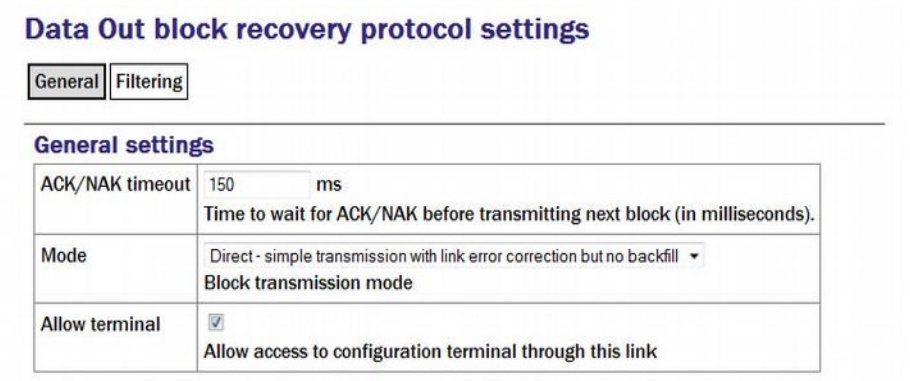

ACK/NAK timeout: Populated with an integer value which specifies the number of milliseconds the server should wait for an acknowledgement packet before transmitting the next block.

Mode: The drop-down menu controls the BRP transmission mode of the server. At present, the only available choice is “Direct - simple transmission with link error correction but no backfill”. Future implementations will offer additional options.

Allow terminal: If checked, the server will allow remote clients access to the source digitiser's command line for configuration purposes.

Filtering

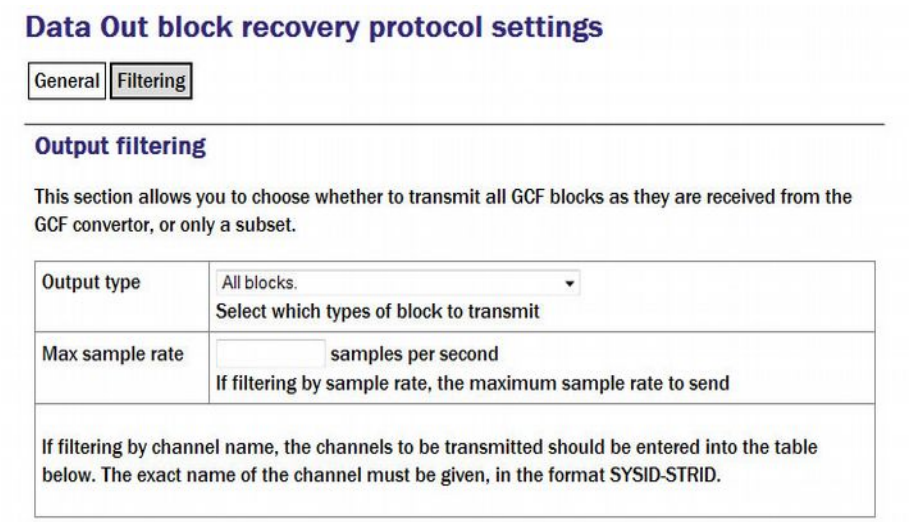

Output type: The drop-down menu offers a choice of:

- **All blocks** - filtering by block type is disabled
- **Only status blocks** - no data blocks are transmitted
- **Only blocks below a certain sample rate** - the threshold (inclusive) rate is specified in the following text field.

- **Only blocks matching a list of channel names** - offering the highest granularity of control.

Max sample rate: If the output type is set to **Only blocks below a certain sample rate**, the text field is used to specify the inclusive threshold, above which data are not transmitted.

If the Output type field is set to **Only blocks matching a list of channel names**, the channel names must be specified in the channel name table:

12.1.2.2 Configurable parameters in expert mode

The following additional configuration parameters in expert mode

Advanced

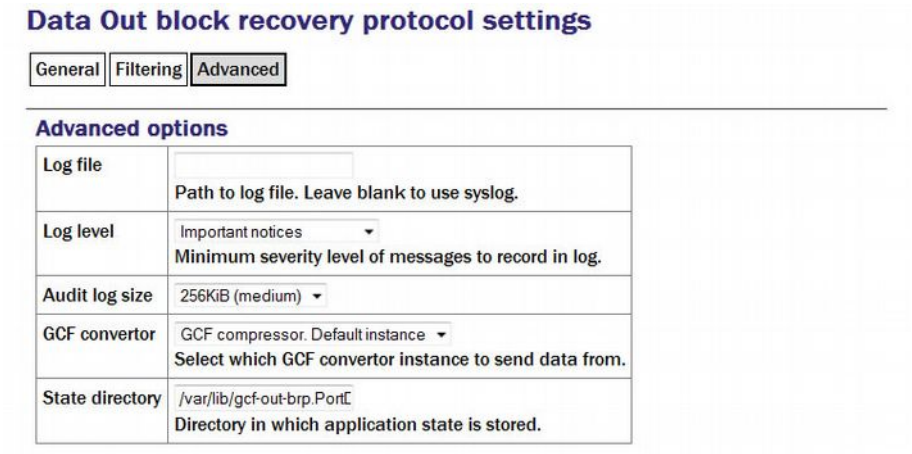

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information;
- Informational messages;
- Important notices; or
- Warnings

Audit log size: The GCF BRP serial sender keeps its own audit log, independent from the system log. The contents of this log are available using the "GCF Audit Log viewer" facility as described in section 14.4.5 on page 241. The amount of data retained is controlled by the drop-down menu, where the choices are:

- 64Kib (small)
- 256Kib (medium)
- 2MiB (large)
- 16MiB (huge)

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

State directory: The GCF BRP protocol requires the transmitter to store some state information. By default, this is held in the directory **`/var/lib/gcf-out-brp.`*`PPP`*** where *PPP* is the port name. The text field can be changed to cause this information to be stored elsewhere (typically on another device). This may be useful for managing storage utilisation in complex configurations.

12.1.3 GCF BRP Network Server

The GCF BRP network server transmits Güralp Compressed Format (GCF) data using the Block Recovery Protocol (BRP) over an Ethernet network.

To configure a GCF BRP network server from the web interface, select:

Configuration → Services → GCF

or

Configuration → All options → System services → GCF

To configure a GCF BRP network server from the command line, start gconfig and select "System services" from the menu.

Now select “gcf-out-brp”. The screen shows a list of all GCF BRP configured server instances.

You can reconfigure any existing service by clicking on its menu entry. To configure a new GCF BRP server instance, select “Create new service instance.”

12.1.3.1 Configurable parameters in simple mode

The configurable parameters for the GCF BRP network server have four tabbed pages in simple mode: General, Network, Protocol and Filter.

General

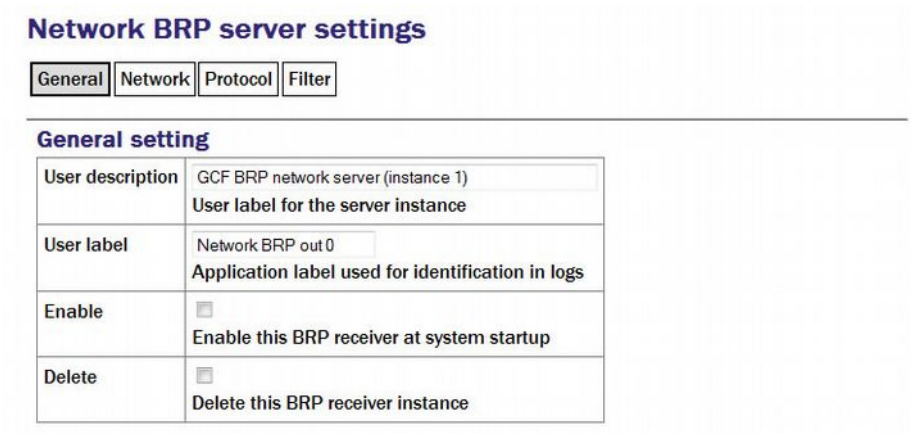

User description: Used to rename the service in configuration menus to something more indicative of its function.

User label: Used to provide a shorter but still potentially more useful name for use in log files.

Enable: Causes this service to start automatically when the system is re-booted if checked. If this check-box is cleared, the service will need to be started manually (from the “Control” → “Services” menu)

Delete: Causes the configuration for this instance to be removed from the system if ticked when the form is submitted.

Network

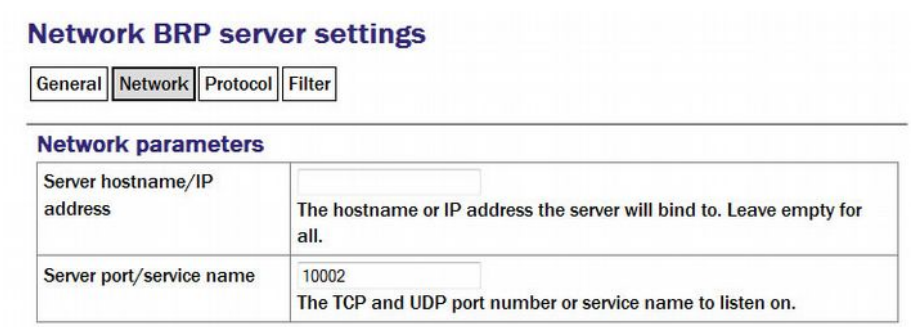

Server hostname/IP address: Used to restrict the server to listen for incoming connection requests only via particular network interfaces. If multiple interfaces or addresses are configured for this system, entering the IP address or associated hostname of one of them prevents connection attempts made to all other addresses. If left blank, connection requests will be considered from all interfaces.

Server port/service name: Must be populated with the service name or port number on which it will listen for incoming connections. This must not be used by any other service on this system.

Protocol

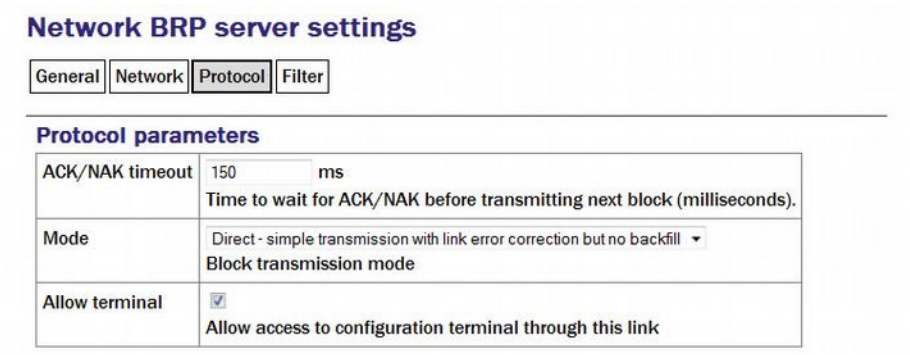

ACK/NAK timeout: Should be populated with an integer value which specifies the number of milliseconds the server should wait for an acknowledgement packet before transmitting the next block.

Mode: The drop-down menu controls the BRP transmission mode of the server. At present, the only available choice is “Direct - simple transmission with link error correction but no backfill”. Future implementations will offer additional options.

Filter

The “Output filtering” section allows the operator to control which data are transmitted, selecting by block type, sample rate or channel name.

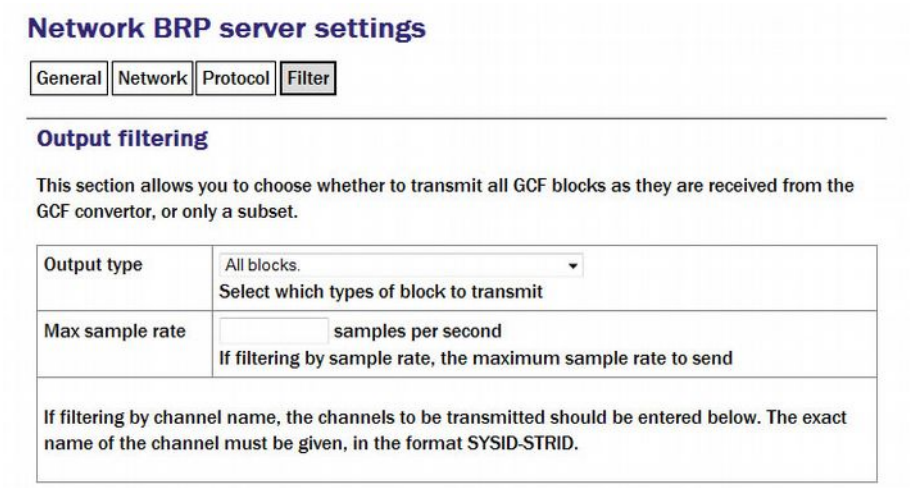

Output type: The drop-down menu offers a choice of:

- **All blocks** - filtering by block type is disabled
- **Only status blocks** - no data blocks are transmitted
- **Only blocks below a certain sample rate** - the threshold (inclusive) rate is specified in the following text field.
- **Only blocks matching a list of channel names** - offering the highest granularity of control.

Max sample rate: If the output type is set to **Only blocks below a certain sample rate**, the text field is used to specify the inclusive threshold, above which data are not transmitted.

If the Output type field is set to **Only blocks matching a list of channel names**, the channel names must be specified in the channel name table:

Channels should be specified by giving their system ID and their stream ID, separated by a hyphen ('-').

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

If the form is submitted when the table is full, extra blank lines are appended.

12.1.3.2 Configurable parameters in expert mode

The following additional configuration parameters are available by clicking the Expert button at the bottom of the form and selecting the 'Advanced' tab.

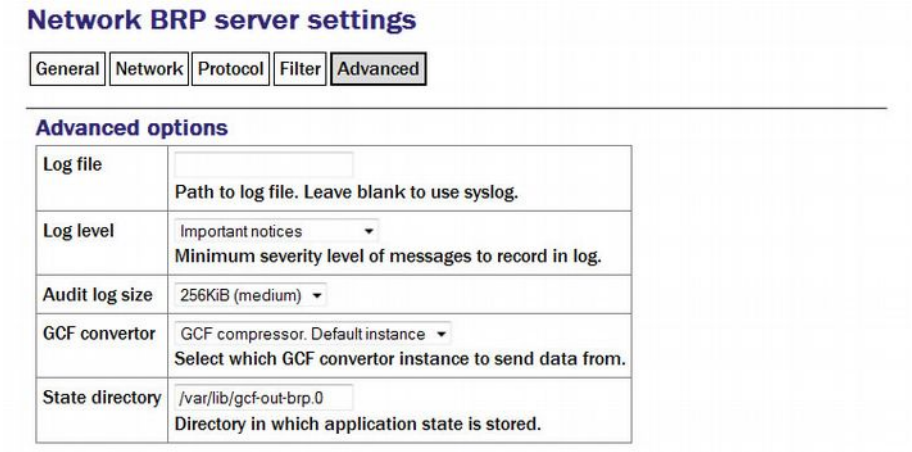

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

Audit log size: The GCF BRP sender keeps its own audit log, independent from the system log. The contents of this log are available using the "GCF Audit Log viewer" facility as described in section 14.4.5 on page 241. The amount of data retained is controlled by the drop-down menu, where the choices are:

- 64Kib (small)
- 256Kib (medium)
- 2MiB (large)
- 16MiB (huge)

Debug port: It is possible to copy all incoming data, verbatim, to a network port, which can be specified in the text field. This is an advanced debugging technique which is beyond the scope of this manual.

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

State directory: The GCF BRP protocol requires the transmitter to store some state information. By default, this is held in the directory `/var/lib/gcf-out-brp.`*n* where *n* is the instance number (counting from zero for the first instance). The text field can be used to cause this information to be stored elsewhere: typically on another device. This may be useful for managing storage utilisation in complex configurations.

12.1.4 GCF Scream Server

The GCF Scream network server transmits Güralp Compressed Format (GCF) data in the native Scream protocol over an Ethernet network.

To configure a GCF Scream network server from the web interface, select:

Data transfer/recording → Services

or

Configuration → All options → System services

To configure a GCF Scream network server from the command line, start `gconfig` and select "System services" from the top level menu.

Now select "gcf-out-scream" from the System Services menu. The next screen shows a list of all GCF Scream server instances that have been configured:

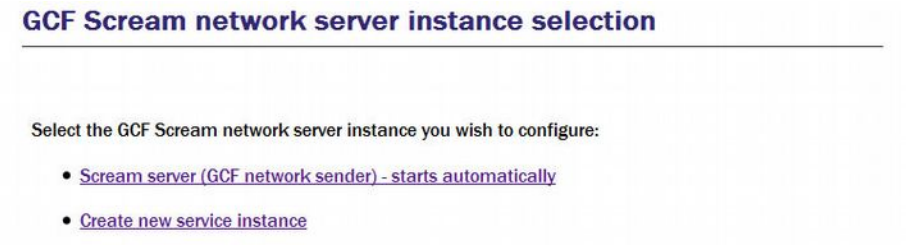

You can reconfigure any existing service by clicking on its menu entry. To configure a new GCF Scream server instance, select "Create new service instance". The following screen allows you to configure the parameters of the service. It is a large form and is shown here in parts.

12.1.4.1 Configurable parameters in simple mode

The configurable parameters for GCF Scream network server have four tabbed pages in simple mode: General, Terminal, Push and Filter

General

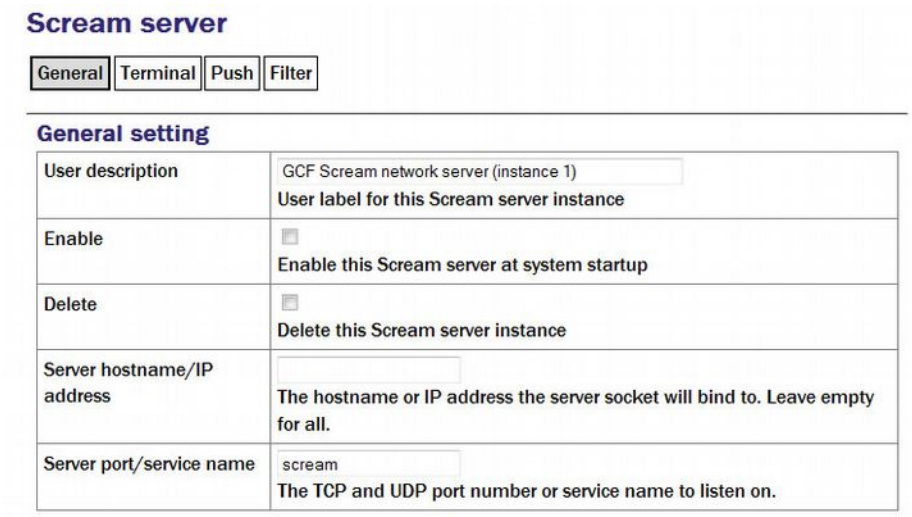

User description: Used to rename the service in configuration menus and log files to something more indicative of its function.

The first instance can neither be disabled nor deleted but, if subsequent instances are created, two additional check-boxes appear on their associated configuration menu:

Enable: Causes this service to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

Server hostname/IP address: Used to restrict the server to listen for incoming connection requests only via particular network interfaces. If multiple interfaces or addresses are configured for this system, entering the IP address or associated hostname of one of them prevents connection attempts made to all other addresses. If left blank, connection requests will be considered from all interfaces.

Server port/service name: Must be populated with the service name or port number on which it will listen for incoming connections. This must not be used by any other service on this system.

Terminal

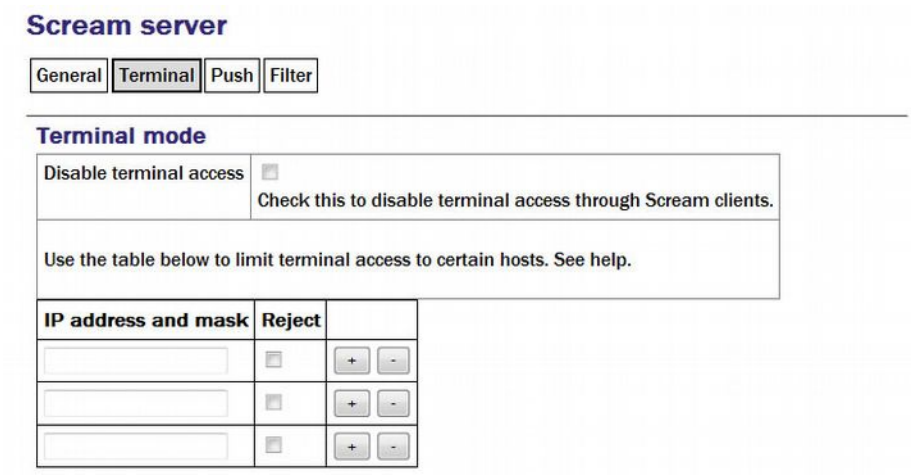

Disable terminal access: Used to prevent clients of this server from accessing the command line of the originating digitiser via the connection, allowing support for data consumers who should not be allowed to reconfigure the data sources. If this check-box is ticked, the table below can be used to either prohibit terminal access to all hosts except those listed (by ticking the relevant **Reject** check-boxes) or to allow access to all hosts except those lists (by leaving the relevant **Reject** check-boxes clear).

Push

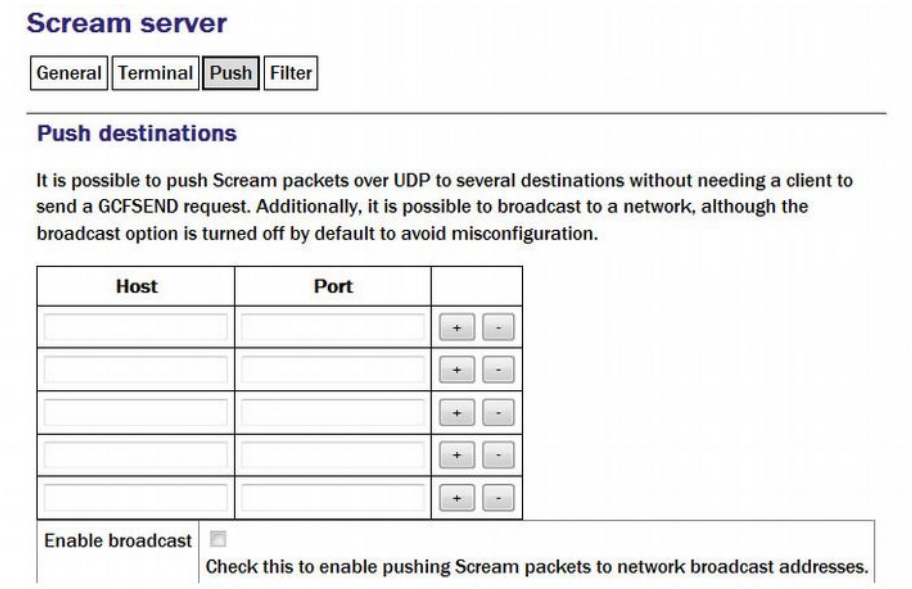

The scream server is capable of both responding to data requests from clients (PULL mode) and of sending data uninvited to remote destinations (PUSH mode). The table below is used to list any PUSH mode clients. For each, a **Host** must be specified as either an IP address or hostname and a **Port** must be given as either a service number or name.

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

Enable broadcast: By default, the server will not send to network broadcast addresses. This behaviour can be enabled by ticking the check-box.

Filter

The "Output filtering" section allows the operator to control which data are transmitted, selecting by block type, sample rate or channel name.

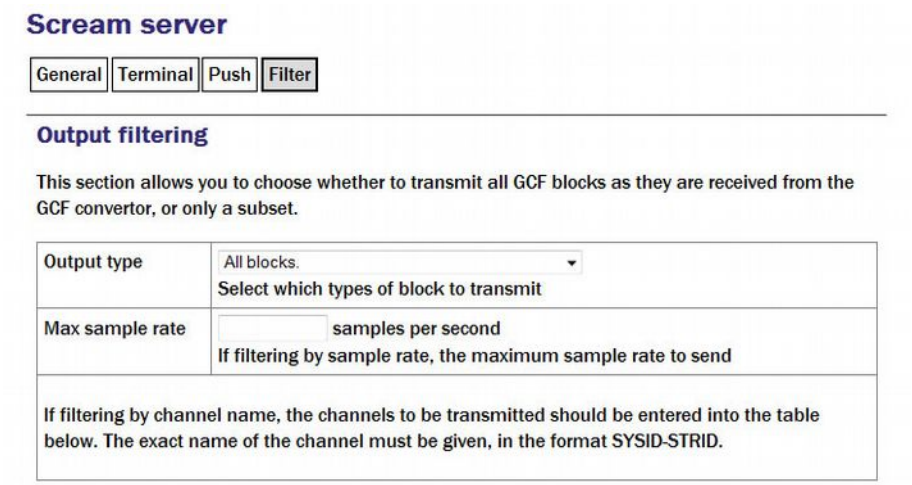

Output type: The drop-down menu offers a choice of:

- **All blocks** - filtering by block type is disabled
- **Only status blocks** - no data blocks are transmitted
- **Only blocks below a certain sample rate** - the threshold (inclusive) rate is specified in the following text field.
- **Only blocks matching a list of channel names** - offering the highest granularity of control.

Max sample rate: If the output type is set to **Only blocks below a certain sample rate**, the text field is used to specify the inclusive threshold, above which data are not transmitted.

If the Output type field is set to **Only blocks matching a list of channel names**, the channel names must be specified in the channel name table:

Channels should be specified by giving their system ID and their stream ID, separated by a hyphen ('-').

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

12.1.4.2 Configurable parameters in expert mode

The following additional configuration parameters are in expert mode:

Advanced

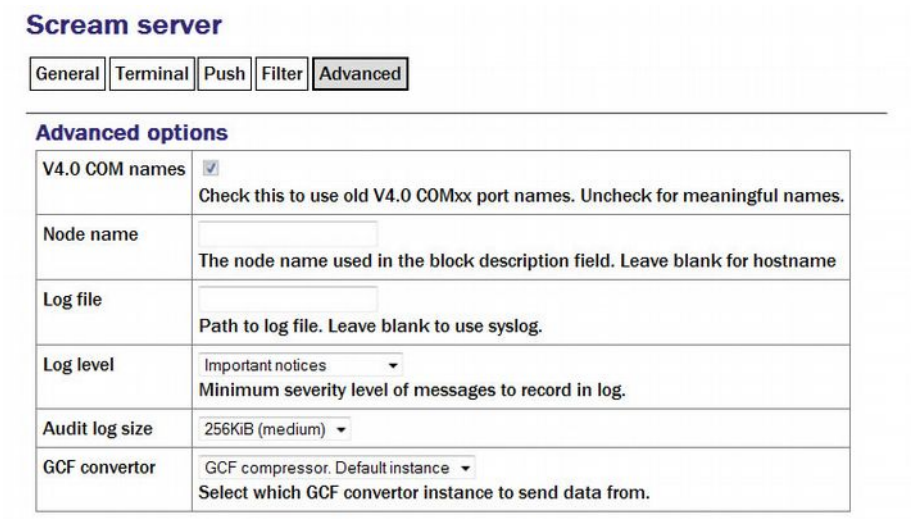

V4.0 COM names: Early versions of the Scream protocol expected all data to originate from COM ports and the port number was used to identify data sources. Version 4.5 and above of the protocol allow for a much more flexible naming scheme. The check-box can be cleared to enable advanced naming or ticked to retain compatibility with earlier versions of the protocol.

Node name: The protocol includes a description field identifying the source of each block. By default, this is set to the host name of the originating machine but it can be over-ridden by entering a value in the field.

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field

can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

Audit log size: The GCF Scream sender keeps its own audit log, independent from the system log. The contents of this log are available using the "GCF Audit Log viewer" facility as described in section 14.4.5 on page 241. The amount of data retained is controlled by the drop-down menu, where the choices are:

- 64Kib (small)
- 256Kib (medium)
- 2MiB (large)
- 16MiB (huge)

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

12.2 SEEDlink

The Standard for the Exchange of Earthquake Data (SEED) is an international standard format for the exchange of digital seismological data developed by the USGS and adopted as a standard by the Federation of Digital Broad-Band Seismograph Networks (FDSN). MiniSEED is a stripped-down version of SEED, which only contains waveform data, without the station and channel metadata that are included in full SEED.

Incoming data (in any format other than CD1.1) is converted first into GDI format. In order to transmit SEEDlink data or record it to mass storage device,

it must be converted into miniSEED format; this is done by the `gdi2miniseed` module, known as the GDI Mini-SEED compressor.

12.2.1 The GDI Mini-SEED compressor

A default instance of the GDI Mini-SEED compressor is provided. Further instances can be created if required for complex implementations. Although the default instance is not set to start automatically, it is a necessary prerequisite for both SEEDlink transmission and recording, so it will be started as a dependant service when required.

To configure a GDI Mini-SEED compressor from the web interface, select:

Configuration → Services → Mini-SEED

or

Configuration → All options → System services → Mini-SEED

To configure a GDI Mini-SEED compressor from the command line, start gconfig and select "System services" from the top level menu.

Now select **gdi2miniseed - Mini-SEED compressor**.

The screen shows a list of all Mini-SEED compressor instances that have been configured:

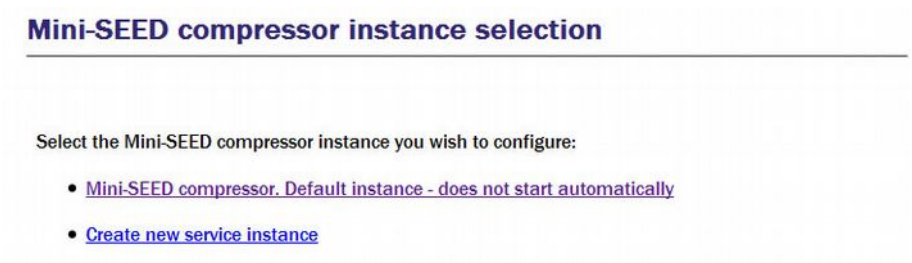

Although the default instance is marked as "does not start automatically", it will be started if a dependant service is started.

You can reconfigure any existing compressor by clicking on its menu entry. To configure a new Mini-SEED compressor instance, select **Create new service instance**. The following screen allows you to configure the parameters of the compressor. The first part of the screen is shown below.

12.2.1.1 Configurable parameters in simple mode

The configurable parameters for the Mini-SEED compressor have two tabbed pages in simple mode: General and Channels.

General

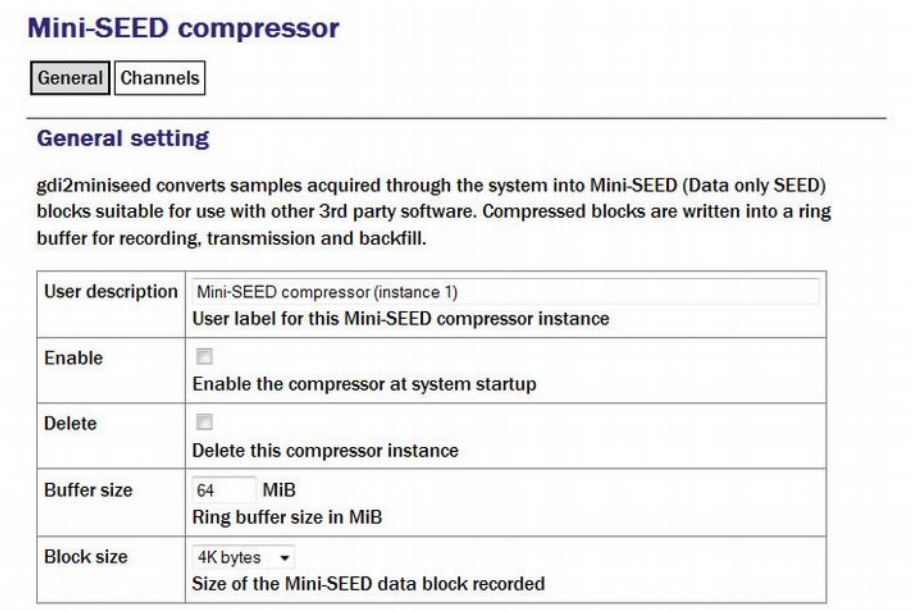

User description: Set to a meaningful name for the data the compressor will handle.

Enable/Disable: The server can be enabled or disabled at boot-up.

Buffer size: Data converted by the compressor are written to a ring-buffer which is read by both the miniSeed recorder and the SEEDlink transmitter. The size of this buffer can be set using the text entry field, which accepts an integer number of mebibytes. Records can be extracted from this buffer to a file: see section 14.4.4 on page 238 for more details.

Block size: The SEED block size is set in the compressor and can not be changed by subsequent software modules. This has the important implication that, if data are to be transmitted using the SEEDlink server, this parameter ***must*** be set to 512 bytes. The size is controlled by the drop-down menu and the possible choices range from 256 bytes to 8K bytes, doubling at each step. The default value is 4K bytes: this is chosen as the optimal for mass storage device recording.

Channels

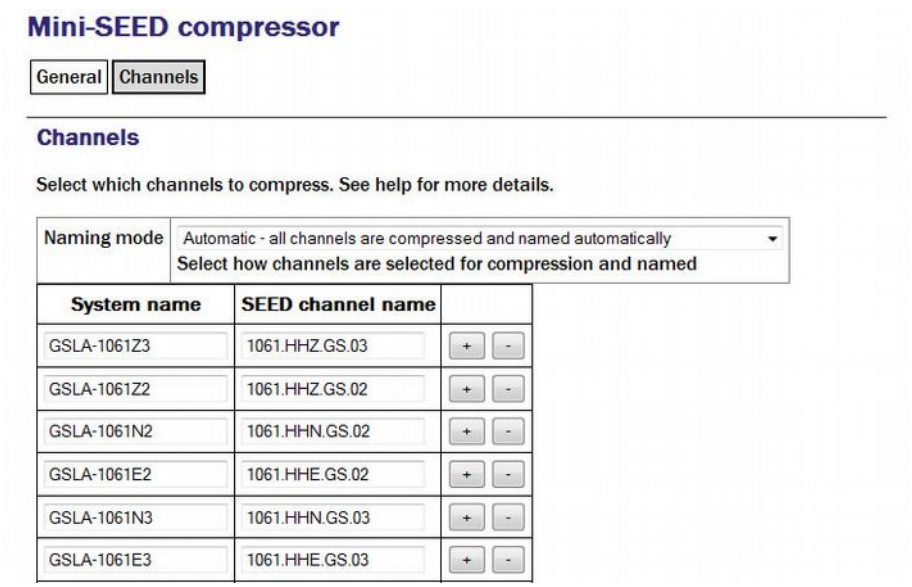

The **Naming mode** drop-down menu offers three choices:

- "Automatic - all channels are compressed and named automatically". This mode offers no filtering and provides system-generated names for each channel forwarded by gdi-base.
- "Semi-automatic - all channels are compressed, names may be mapped below". In this mode, one or more of the channels may be renamed by adding entries to the mapping table. If you wish to use this mode, it may be useful first to run the system in automatic mode for a short while: this will populate the mapping table with an entry for each currently known channel, which can serve as the basis for your own mapping table.
- "Manual - only channels named below are compressed". This mode offers both channel filtering and name mapping. If you wish to use this mode, it may be useful first to run the system in automatic mode for a short while: this will populate the mapping table with an entry for each currently known channel, which can serve as the basis for your own mapping table.

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

If the form is submitted when the table is full, extra blank lines are appended.

Deleted channels will be transmitted unmapped in "Semi-automatic" mode and not transmitted in "Manual" mode.

12.2.1.2 Configurable parameters in expert mode

The following additional configuration parameters are available by clicking the "Expert" button at the bottom of the form and selecting the 'Advanced' tab.

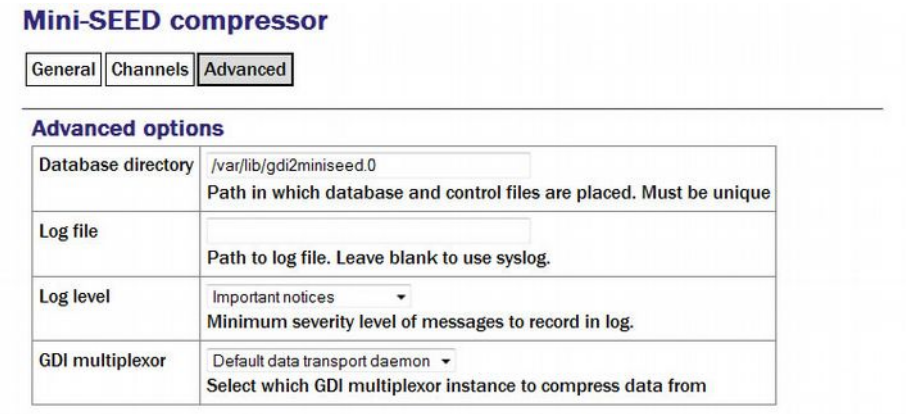

The **Database directory** field can be used to control the location of the ring-buffer and associated files. In most configurations, the default location is adequate but if, for example, a very large ring-buffer is desired and the optional extra flash memory module is fitted, it may be desirable to use the extra memory for this purpose. To do this, enter into this field the path to a unique directory under `/media/flash_module`.

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

12.2.2 The SEEDlink server

The SEEDlink server transmits data in miniSEED format (i.e. no station/channel metadata) over the network to remote data consumers. The data are generated by a GDI Mini-SEED compressor instance.

Note: The SEEDlink server requires data in 512 byte blocks - the compressor must be reconfigured from its default setting (4 Kbytes) if the SEEDlink server is to be used: see the previous section for details.

The system will prevent you from configuring a SEEDlink server unless the chosen compressor is set to prepare 512-byte blocks. The system does not stop you from subsequently reconfiguring the compressor but, if you change the block size, the SEEDlink server will fail.

If you want to write miniSEED data with a block-size other than 512 bytes *and* run a SEEDlink server, you should instantiate different compressors for each.

A single SEEDlink server instance takes data from a single compressor instance and can serve multiple, simultaneous clients. If it is required to serve different channels to different clients, multiple server instances should be configured, each receiving data from a different compressor instance (the channel selection is controlled by the compressor, not the server). A server has a configured "Organization" name: if data are to appear to come from multiple organizations, multiple server instances should be configured: they can share a compressor instance if they will be serving the same channels or a number of compressor instances can be used.

To configure a SEEDlink server from the web interface, select:

Data transfer/recording → Services

or

Configuration → All options → System services

To configure a SEEDlink server from the command line, start `gconfig` and select "System services" from the top level menu.

Now select "seedlink-out -- SEEDlink network server" from the System Services menu. The next screen shows a list of all SEEDlink server instances that have been configured:

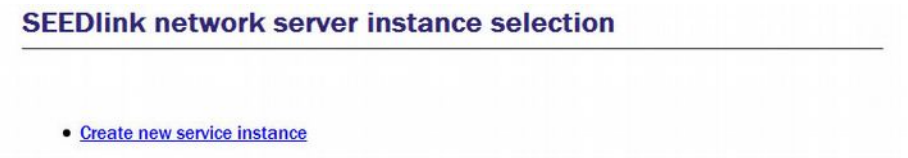

You can reconfigure any existing service by clicking on its menu entry. To configure a new SEEDlink server, select "Create service instance". The following screen, shown overleaf, allows you to configure the parameters of the server.

12.2.2.1 Configurable parameters in simple mode

The configurable parameters for the SEEDlink network server are contained in a single form in simple mode:

General

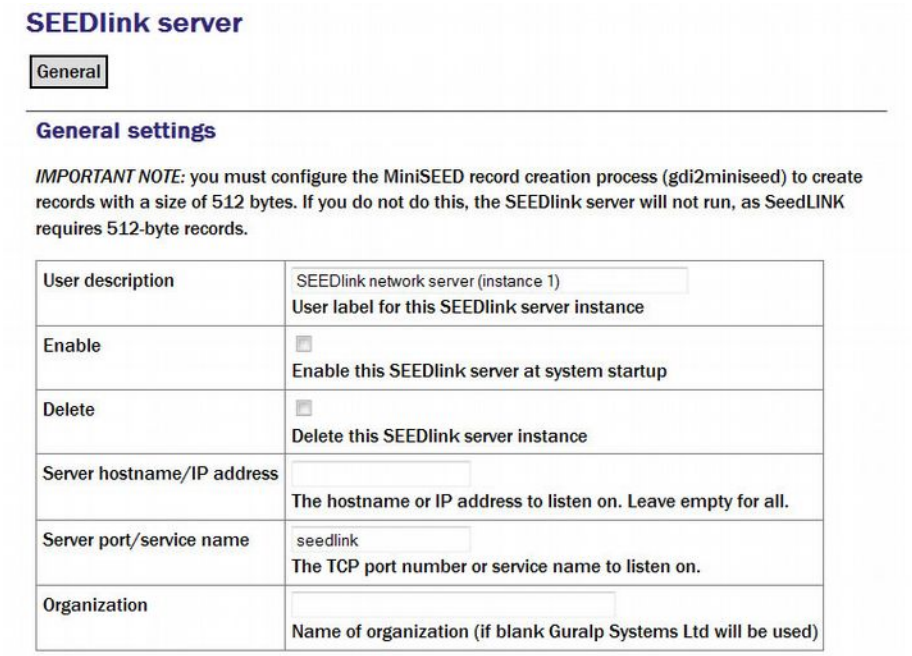

User description: Set to a meaningful name for the server data.

Enable: Causes this service to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

Server hostname/IP address: To configure the server to listen for incoming data requests only on a specific IP address, set this (or the associated host name) in the text field. By default it will listen on all configured interfaces.

Server Port/service name: Set the port (port number or service name) that you want the server to listen on in the text field.

Organization: The server identifies itself to clients with an organization name, this should be entered into the text-field. If left blank, the value will default to "Güralp Systems Ltd".

12.2.2.2 Configurable parameters in expert mode

The following additional configuration parameters are available by clicking the “Expert” button at the bottom of the form and selecting the 'Advanced' tab.

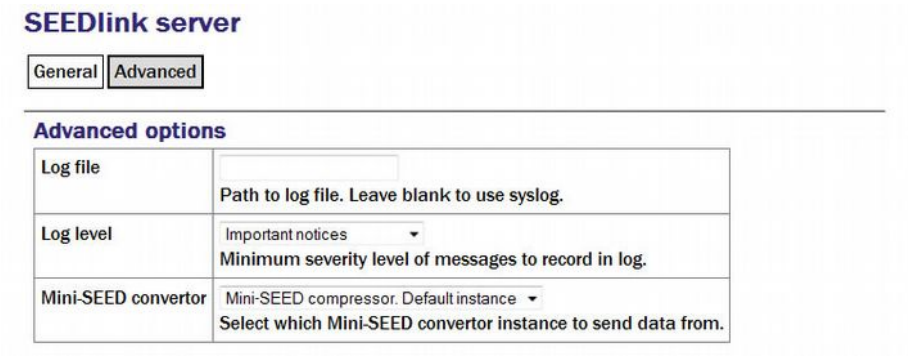

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

12.3 EarthWorm

EarthWorm is a suite of automated earthquake processing software developed by Instrument Software Technologies, Inc. For more information, please see www.isti.com/products/earthworm.

The EarthWorm sender allows Platinum systems to send data directly to an EarthWorm installation.

To configure an EarthWorm sender from the web interface, select:

Configuration → Services → Miscellaneous

or

Configuration → All options → System services → Miscellaneous

To configure an EarthWorm sender from the command line, start gconfig and select “System services” from the top level menu.

Now select “gdi2ew -- Earthworm sender”. The screen shows a list of all EarthWorm sender instances that have been configured:

You can reconfigure any existing service by clicking on its menu entry. To configure a new EarthWorm sender, select “Create service instance”. The following screen allows you to configure the parameters of the sender.

12.3.1 Configurable parameters in simple mode

The configurable parameters for the EarthWorm sender have four tabbed pages in simple mode: General, Settings, Channels and Heartbeats.

12.3.1.1 General tab

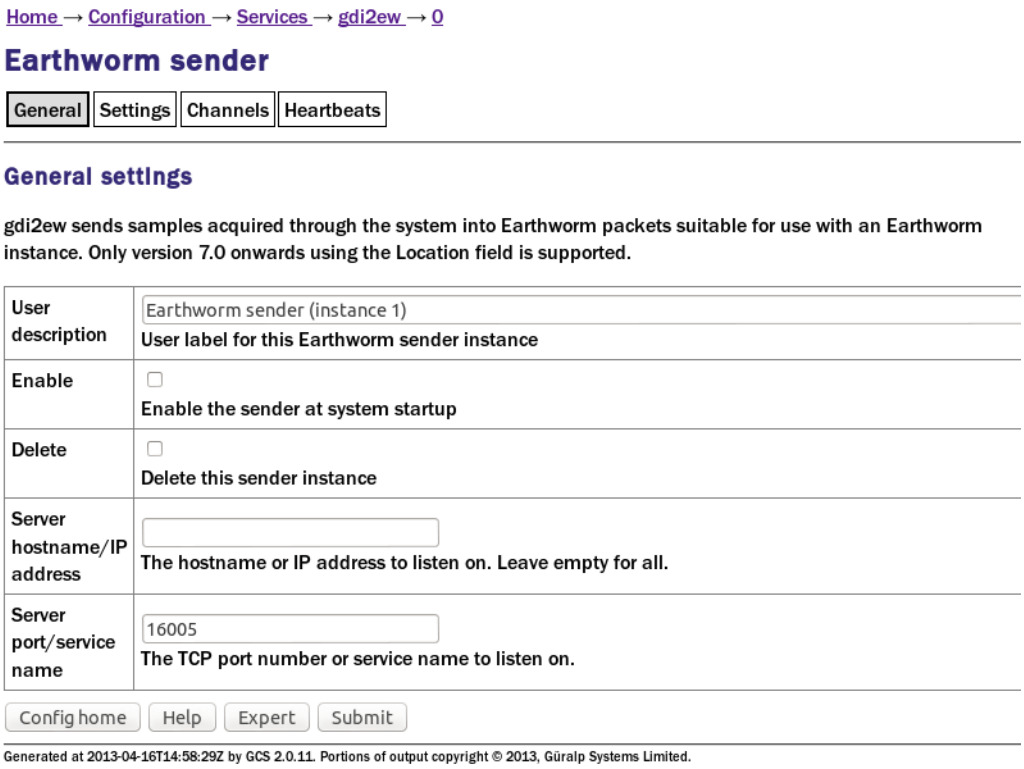

User description: This field allows the user to configure a meaningful name for this EarthWorm sender instance.

Enable: When ticked, causes this service instance to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

Server hostname / IP address: By default, the service instance will listen for incoming connections on all available interfaces. If this field is populated with an IP address, the instance will only listen for connections addressed to the specified IP address. This option is only useful on systems with multiple IP addresses.

Server port/service name: This field specifies the port on which the service instance listens for incoming connections. The port can be specified as either a name or a number. The mapping from port names to port numbers is configured by the conventional Linux file `/etc/services` which can be edited from the command line if required.

12.3.1.2 Settings tab

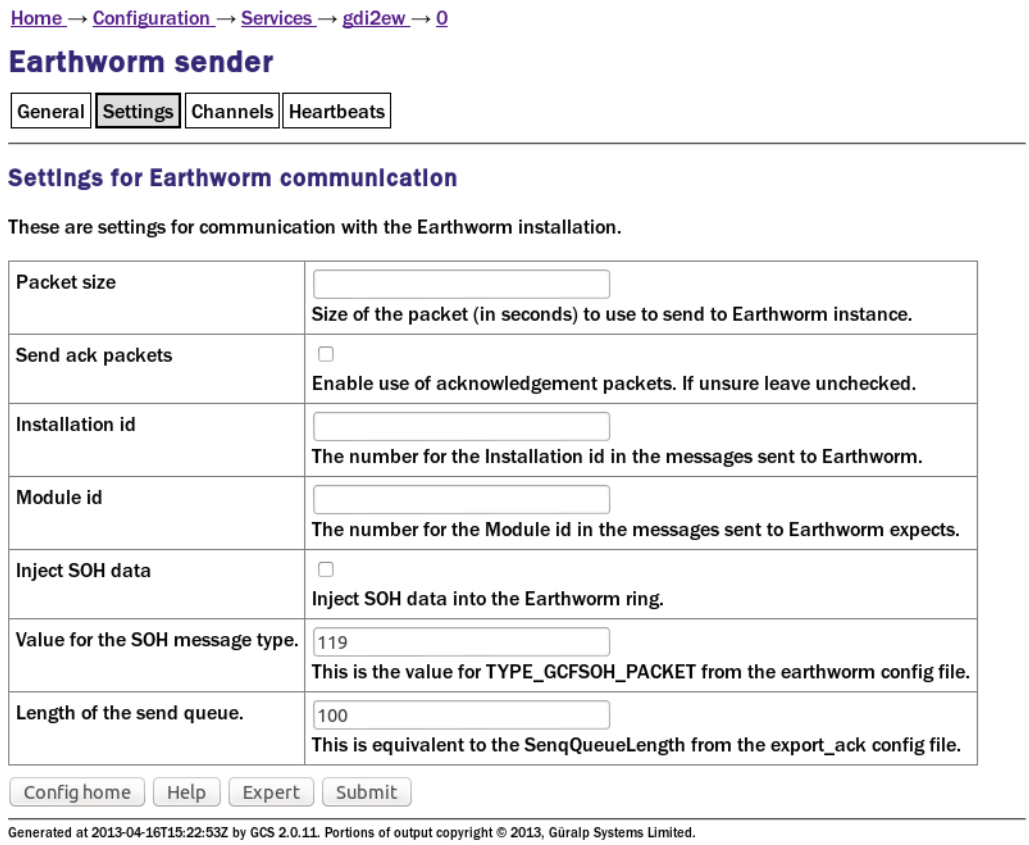

Packet size: This field controls the size of the data packets that are transmitted. The size should be specified as a duration, in seconds.

Send ack packets: This check-box, when ticked, enables the use of acknowledgement packets. These are not normally required when the connection between the client and server is reliable.

Installation id: This field allows the user to specify the installation ID that will be supplied to the EarthWorm installation.

Module id: This field allows the user to specify the module ID that will be supplied to the EarthWorm installation.

Inject SOH data: This check-box, when ticked, causes the service instance to generate special state-of-health (SoH) packets and send them to the EarthWorm installation.

Value for the SOH message type: State-of-health packets should have a distinct message type to distinguish them from seismic data packets. If the previous check-box is ticked, use this field to specify the message type. This value must match the value of `TYPE-GCFSOH_PACKET` in the EarthWorm configuration file.

Length of the send queue: This field controls the size of the buffer used for sending data and should match the parameter `SendQueueLength` in the EarthWorm `export_ack` configuration file.

12.3.1.3 Channels tab

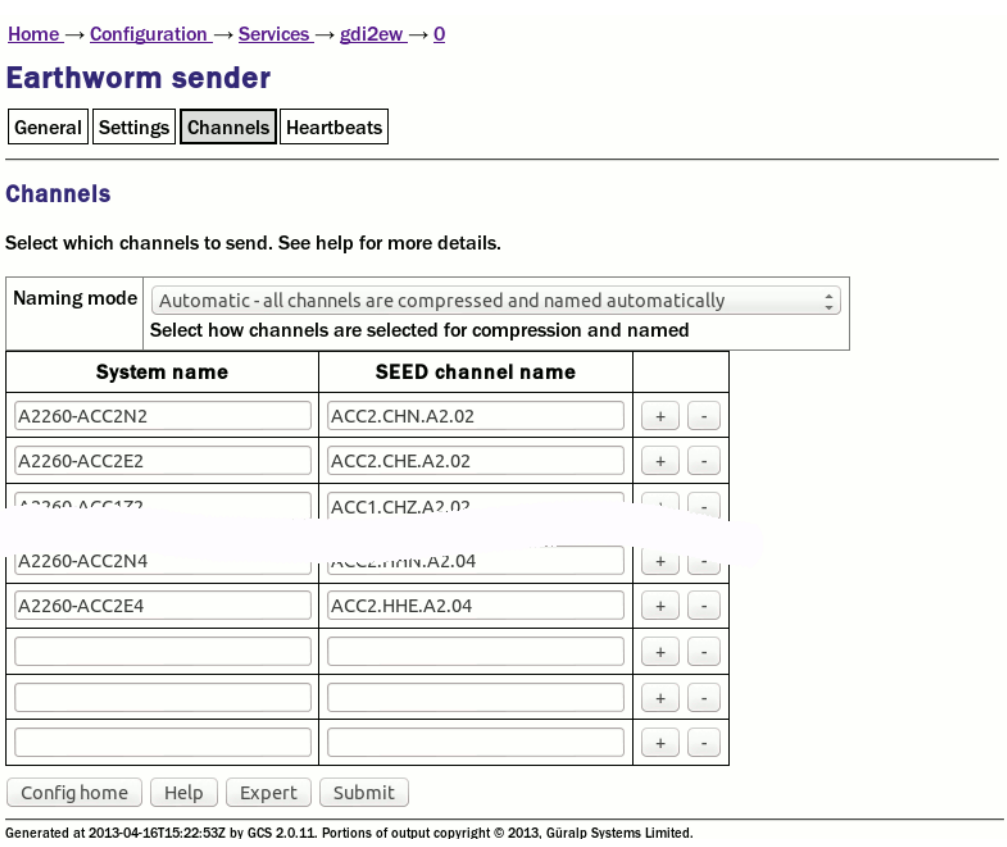

The channels tab contains a drop-down menu and a table which allow the operator to control which channels are transmitted and, optionally, rename them.

Naming mode: The drop-down menu offers three choices:

- **Automatic - all channels are compressed and named automatically:** Offers no filtering and uses the system-generated names for each channel as forwarded by `gdi-base`.
- **Semi-automatic - all channels are compressed, names may be mapped below:** One or more of the channels may be renamed by adding entries to the mapping table. If you wish to use this mode, it may be useful first to run the system in automatic mode for a short while: this will populate the mapping table with an entry for each currently known channel, which can serve as the basis for your own mapping table.
- **Manual - only channels named below are compressed:** Offers both channel filtering and name mapping. If you wish to use this mode, it may be useful first to run the system in automatic mode for a short while to populate the mapping table with an entry for each currently known channel; this can serve as the basis for your own mapping table.

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

Channels which do not appear in the mapping table will be transmitted unmapped in “Semi-automatic” mode and not transmitted in “Manual” mode.

12.3.1.4 Heartbeats tab

EarthWorm has a “heart-beats” facility which can be used for monitoring connections. Heartbeats are generated both by the client and by the server. A heartbeat is a small packet sent at regular intervals so that an alert can be raised if a heartbeat packet does not arrive within a configurable time. Heartbeat packets contain configurable messages.

The following fields can be configured:

Export heartbeat timeout: This parameter controls how often this sender instance will broadcast heartbeat packets to the EarthWorm installation.

Export heartbeat message: This field allows the user to set the message contained within outgoing heartbeat packets.

Import heartbeat timeout: This parameter controls how often this sender instance should expect to receive heartbeat packets from the EarthWorm installation.

Import heartbeat message: This field allows the user to set the message that this sender instance should expect to see within incoming heartbeat packets.

Heartbeat debug messages: If this check-box is ticked, debugging messages are generated for every heartbeat packet sent, received or expected but missed.

Debugging messages will normally be written to /var/log/messages but a different file can be specified in Expert mode.

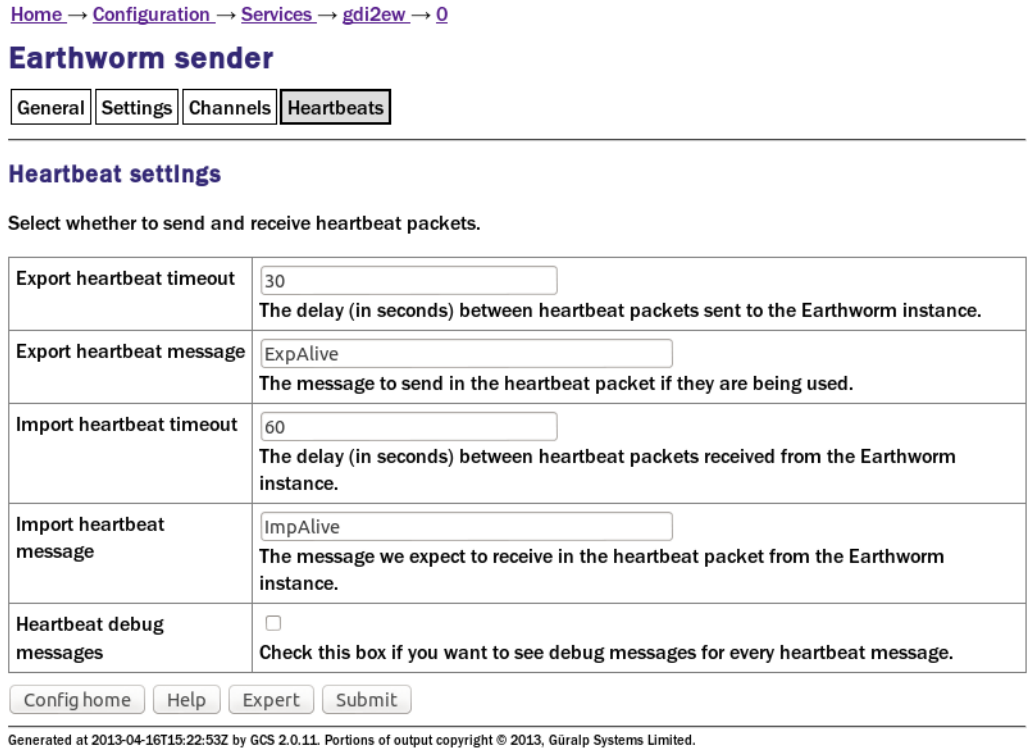

12.3.2 Configurable parameters in Expert mode

Clicking the Expert button will display an additional tab, “Advanced”, with three extra fieds.

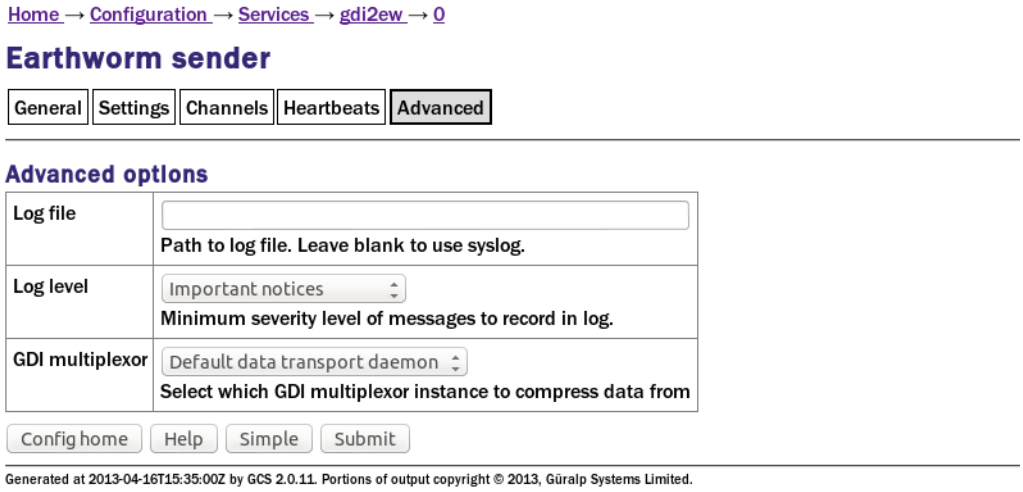

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this sender from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging.

If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

12.4 Güralp Seismic Monitoring System

GSMS is a protocol designed by Güralp Systems to send real time, low latency strong motion data.

To configure a GSMS server from the web interface, select:

Configuration → Services → Miscellaneous

or

Configuration → All options → System services → Miscellaneous

To configure a GSMS server from the command line, start `gconfig` and select "System services" from the top level menu.

Now select "gsms-out -- GSMS sender". The screen shows a list of all GCF Scream server instances that have been configured:

You can reconfigure any existing service by clicking on its menu entry. To configure a new GSMS server, select "Create service instance". The following screen allows you to configure the parameters of the server.

12.4.1 Configurable parameters in simple mode

The configurable parameters for the GSMS server have four tabbed pages in simple mode: General, Network, Push and Channels.

12.4.1.1 General

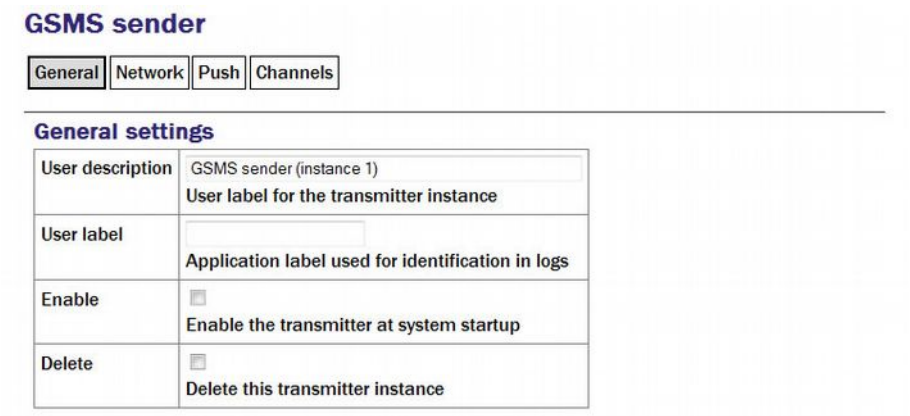

User description: Set to a meaningful name for the data that it will serve by populating the text field.

User Label: Can be filled in with a name which will then be used to identify this instance in log files.

Enable: Causes this service to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

12.4.1.2 Network

Bind host: configure the server to listen for incoming data requests only on a specific IP address. By default it will listen on all configured interfaces.

Service Port: Set the port (port number or service name) that you want the server to listen on in the text field.

12.4.1.3 Push

Push allows the server to pro-actively send data to remote GSMS receivers.

Protocol: Select TCP or UDP - this must match the receiver's setting.

Push host: Enter IP addresses (or host names).

Service: Enter port numbers (or service names)

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

If the form is submitted when the table is full, extra blank lines are appended.

12.4.1.4 Channels

The GSMS server need not send all data from all channels to its clients. It is possible to select which channels are transmitted.

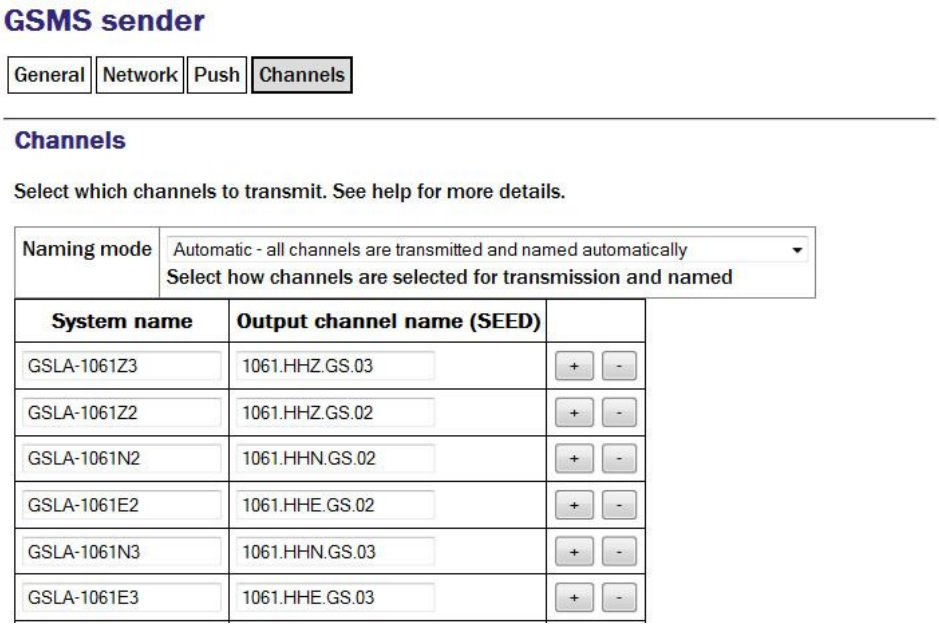

Select one of the three different naming modes:

- Automatic: all channels are transmitted and named automatically
- Semi-automatic: all channels are transmitted and names can be mapped using a configuration table
- Manual: only channels named in the configuration table are transmitted.

The software will attempt to populate the table based on incoming data streams so it is a good idea to configure all input sources and run the system for a few minutes before completing this table.

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

12.4.2 Configurable parameters in expert mode

The following additional configuration parameters are available by clicking the “Expert” button at the bottom of the form and selecting the 'Advanced' tab.

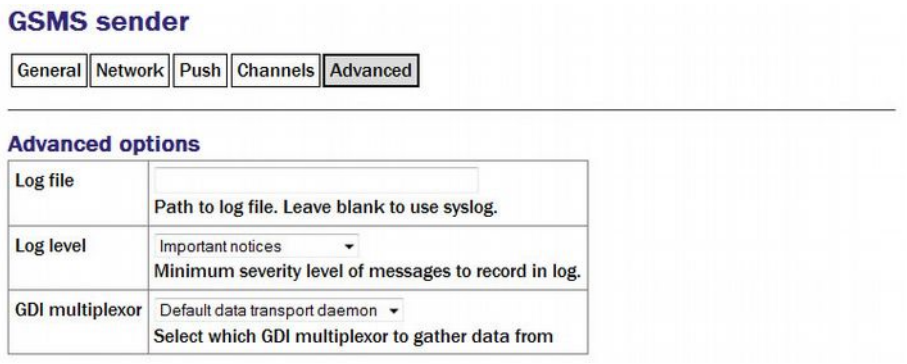

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

12.5 Quick Seismic Characteristic Data

QSCD is a protocol developed by KIGAM (http://www.kigam.re.kr/eng) to send strong motion results, which are computed every second.

To set up a QSCD server on the acquisition module, first configure the relevant strong motion data sources as described in section 8.1.1.4 on page 103, then, from the web interface, select:

Configuration → Services → Miscellaneous

or

Configuration → All options → System services → Miscellaneous

To configure a QSCD server from the command line, start `gconfig` and select "System services" from the top level menu.

Now select "qscd-out -- KIGAM QSCD (Quick Seismic Characteristic Data) sender " from the System Services menu. The next screen shows a list of all QSCD server instances that have been configured:

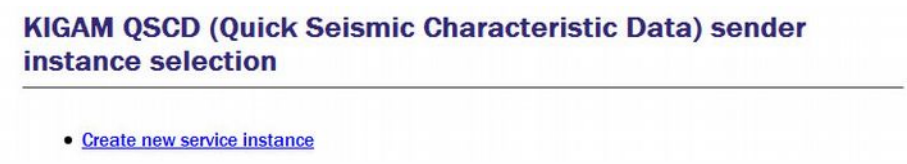

You can reconfigure any existing service by clicking on its menu entry. To configure a new QSCD server, select "Create service instance". The following screen allows you to configure the parameters of the server. As it is a large screen, it is shown here in pieces.

12.5.1 Configurable parameters in simple mode

The configurable parameters for the QSCD server have three tabbed pages in simple mode: General, Network and Channels.

General

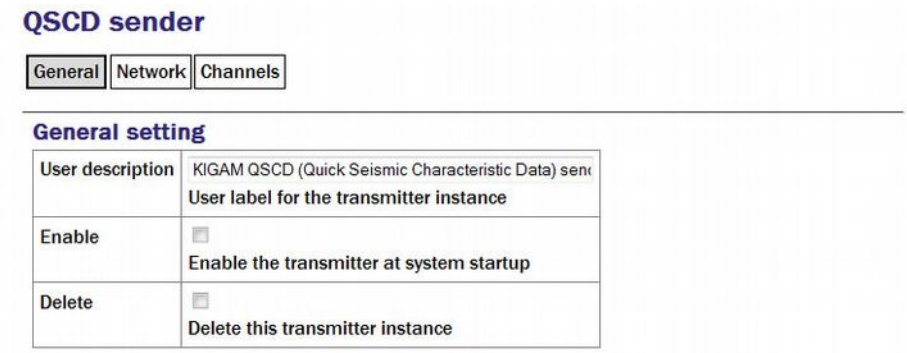

User description: Set to a meaningful name for the data that it will serve.

Enable: Causes this service to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

Network

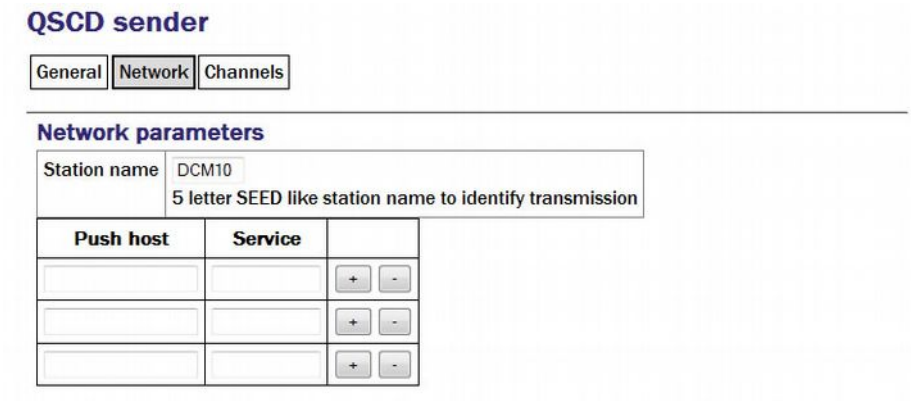

Station name: Like SEED, QSCD links require a unique name to identify the source of the data.

To send QSCD data to remote hosts, enter their DNS names or IP addresses in the table, with the associated service name or port number for each. Port names and numbers are associated with each other in the standard Linux `/etc/services` file.

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

Channels

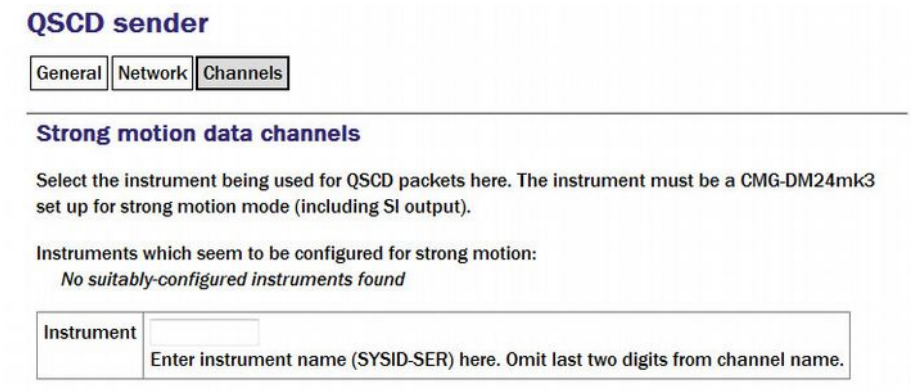

Instrument: The acquisition module scans all incoming data and prepares a list, in the correct format, of the names of instruments which are sending strong motion results. Enter one of these names in the field.

Note: The QSCD protocol only supports a single instrument. If you need to transmit results from multiple instruments, you should configure multiple QSCD sender instances, one for each instrument.

12.5.2 Configurable parameters in expert mode

The following additional configuration parameters are available by clicking the “Expert” button at the bottom of the form and selecting the 'Advanced' tab.

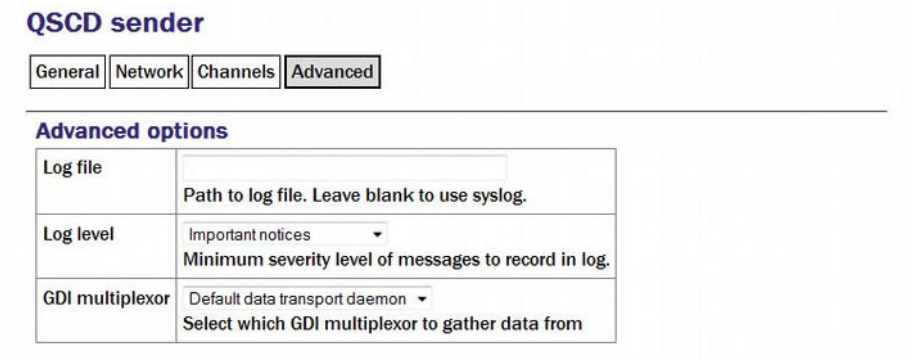

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

12.6 WIN Sender

WIN is a Japanese seismic data format.

To set up a WIN server on the acquisition module using the web interface select:

Configuration → Services → Miscellaneous

or

Configuration → All options → System services → Miscellaneous

To configure a WIN server from the command line, start `gconfig` and select "System services" from the top level menu.

Now select "win-out - WIN sender ". The screen shows a list of all WIN server instances that have been configured:

You can reconfigure any existing service by clicking on its menu entry. To configure a new WIN sender, select "Create service instance". The following screen allows you to configure the parameters of the sender. It is shown here in parts.

12.6.1 Configurable parameters in simple mode

The configurable parameters for the WIN sender have three tabbed pages in simple mode: General, Network and Channels.

General

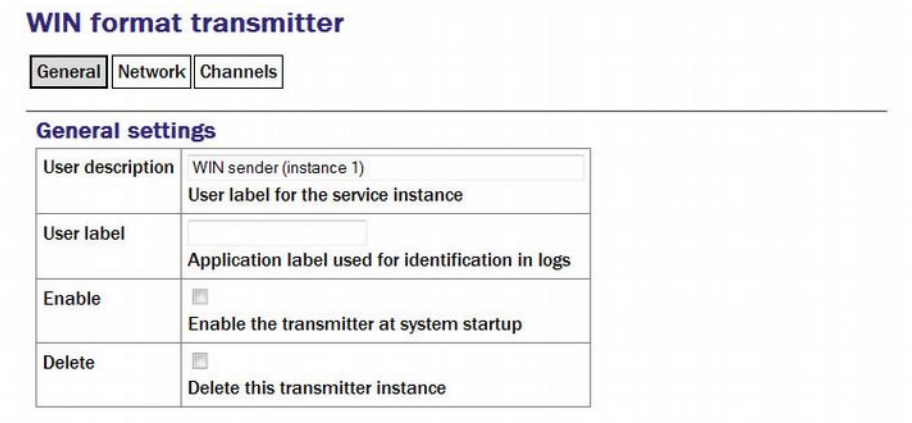

User description: Set to a meaningful name for the data that it will send.

User label: Can be set to distinguish this instance from others in the log files.

Enable: Causes this service to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

Network

The WIN transmitter can be configured to be either a TCP server to multiple clients, or a UDP sender to a single address. If you want to sent the data to multiple clients, set up the acquisition module as a TCP server and the remote machines as clients that connect to it.

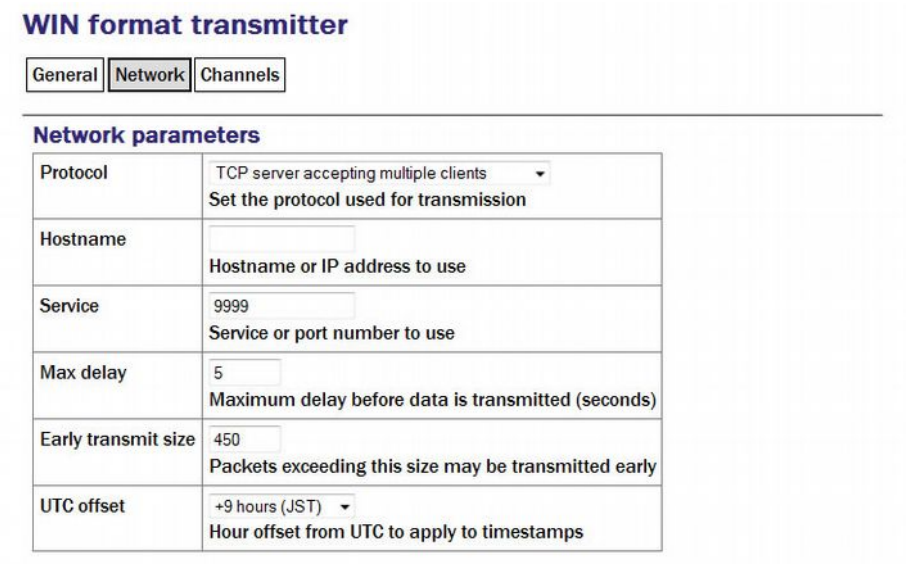

Protocol: To configure the sender as a TCP server, select "TCP server accepting multiple clients" from the drop-down list.

Hostname: To use a specific IP address to listen for requests from clients, set this in the box. By default it will listen on all interfaces.

Service: Set to the port that you want the server to listen on in the box.

If you only want to send the data to a single UDP server, select "UDP datagrams sent to specified address" from the **Protocol** drop-down list. Configure the remote machine's hostname or IP address in the **Hostname** box and set the port number that the remote machine will listen on in the **Service** box.

Max delay: The WIN sender will buffer up data before it is sent so that outgoing packets have a second's worth of data from all channels. If no data are received from some channels within a certain time limit, the data from other channels will be transmitted anyway. This limit is specified by the value in the field and defaults to five seconds.

Early transmit size: If a packet in construction exceeds the size specified the packet will also be sent early.

UTC Offset: The WIN Format uses the local time in order to time-stamp packets. The offset of the local time-zone from UTC used in the GCF data is specified in the box.

Channels

Note: Previous versions of the firmware required this mapping to be entered in SEED notation but this is no longer the case.

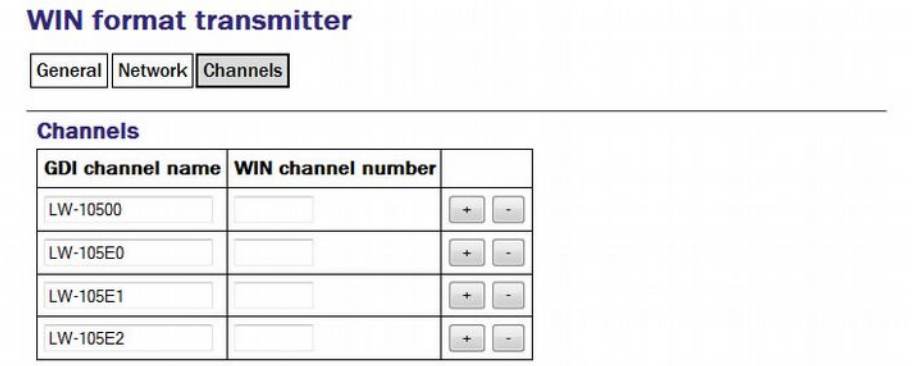

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

12.6.2 Configurable parameters in expert mode

The following additional configuration parameters are available by clicking the “Expert” button at the bottom of the form and selecting the 'Advanced' tab.

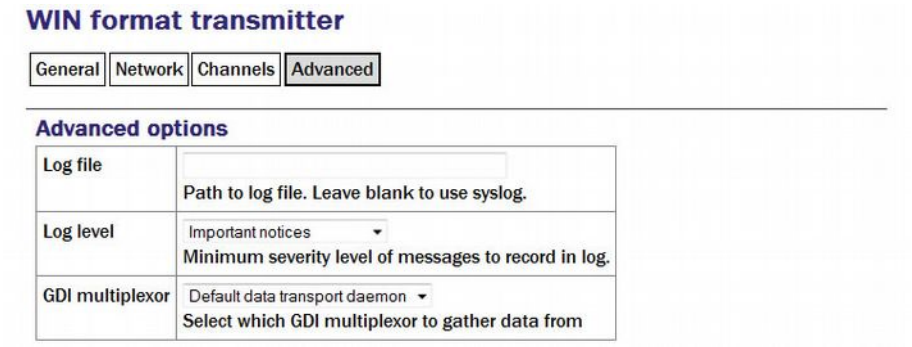

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages

- Important notices
- Warnings

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

13 Building Networks

13.1 GDI-link

The GDI-link protocol provides the most efficient means of exchanging data between two systems running Platinum firmware. GDI is the native data format of the central data multiplexer, the `gdi-base` module, and GDI-link allows highly efficient, low latency data exchange between two such multiplexers without the overhead of any additional protocol conversion. State of health information is attached to samples before transmission.

GDI links have transmitters, which send data, and receivers which receive it. These terms do not refer to the direction of initiation of the network connection: a receiver can initiate a connection to a transmitter and vice versa.

A single GDI-link receiver can accept data from multiple transmitters and a single transmitter can send data to multiple receivers, allowing maximum flexibility in configuring seismic networks.

Caution: The GDI link transmitter generates back-fill files which it does not delete. If ignored, these files will accumulate and eventually fill the file system. It is necessary to configure a directory cleaner to remove these files. See section 13.1.1.3 on page 207 for details.

13.1.1 The GDI-link transmitter

To configure a GDI-link transmitter via the web interface select:

Configuration → Services → GDI

or

Configuration → All options → System services → GDI

Now select "gdi-link-tx - GDI link transmitter". The screen shows a list of all GDI link transmitter instances that have been configured on the acquisition module.

In most circumstances you will only need a single GDI link transmitter but this screen allows you to create more if desired.

To configure the transmitter, click on the link corresponding to the required instance.

13.1.1.1 Configurable parameters in simple mode

The configurable parameters for the GDI link transmitter have four tabbed pages in simple mode: General, Network, Backfill and Push.

General

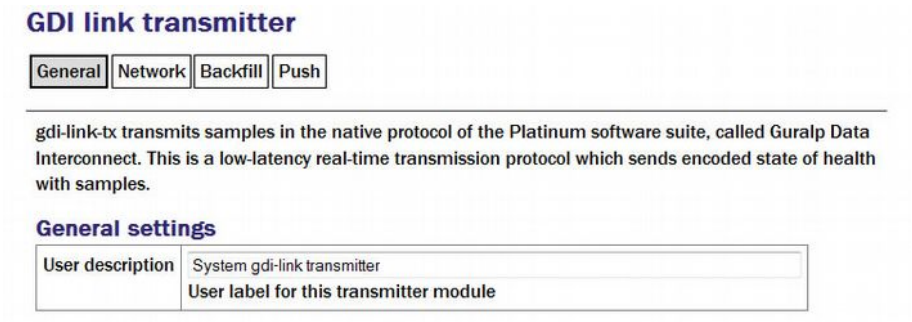

User description: Can be changed if desired. This may be useful if you have multiple instances. This description is seen when viewing running services or configuring instances. It is not seen by the clients.

Subsequent instances can be enabled or disabled with a check-box but this is absent from the page for the default instance because the default instance is always enabled.

Network

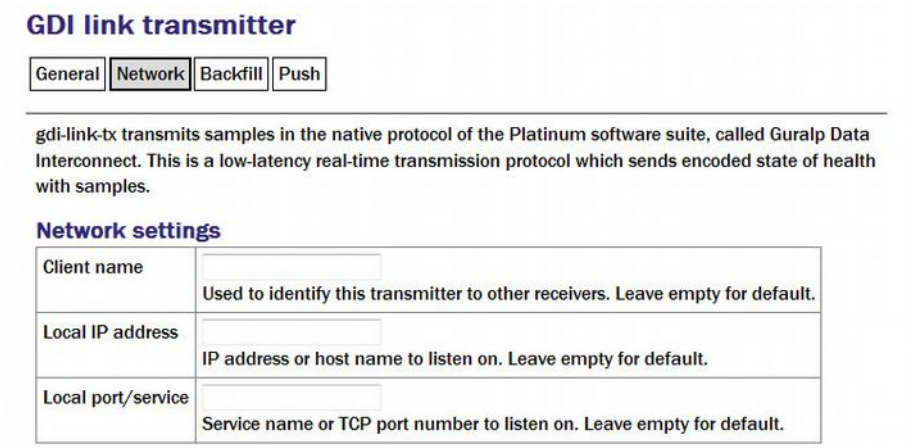

Client name: Set to the instance name as seen by the client. A suitable default is used if this field is left blank.

Local IP address: If the acquisition module has multiple network addresses, it can be restricted to listen for incoming connections on only one of them by entering its address here. If left blank, the transmitter will listen on all available instances.

Local port/service: The default service (port) for the transmitter is 1565 but an alternative port can be entered here if required.

Back-fill

Back-fill is the process whereby missing data are recovered.

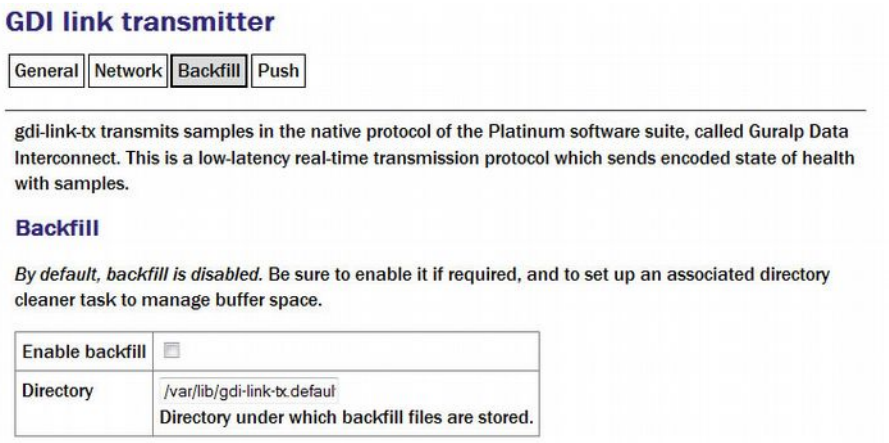

Enable backfill: If selected enables the back-fill transmission.

Directory: The default directory for back-fill data.

Caution: Unless, managed, back-fill files will accumulate ultil they fill the disk, which will disable the system. It is essential to configure a directory-cleaner instance to manage these files: see section 13.1.1.3 on page 207 for details.

Push

The push tab contains a table within which you can configure the GDI link clients to which this transmitter should send data.

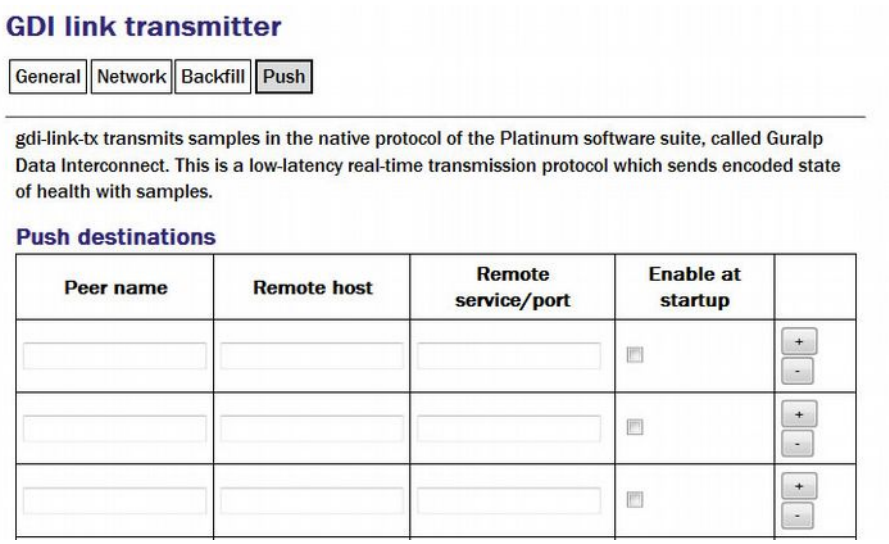

For each client, you should set:

- **Peer name**: Should match the server name configured in the client (receiver) at the remote end of the link
- **Remote host**: The DNS name or IP address of the GDI link client
- **Remote service/port**: The default is 1566 but, if you have configured a different port on the GDI link client, you should enter the same port here.
- **Enable at startup** Controls whether the service enables at startup.

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

13.1.1.2 Additional options available in expert mode

The following additional configuration parameters are available by clicking the “Expert” button at the bottom of the form and selecting the 'Advanced' tab.

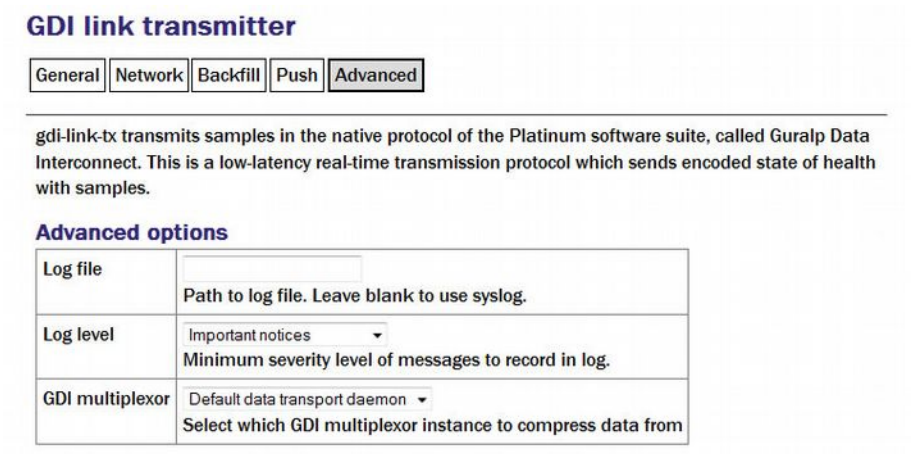

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information;
- Informational messages;
- Important notices; or
- Warnings

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

13.1.1.3 Managing back-fill files

Unless, managed, back-fill files will accumulate ultil they fill the disk, which will disable the system. It is essential to configure a directory-cleaner instance to manage these files. One directory cleaner instance is required for each `gdi-link-tx` instance.

Instructions for creating a new directory cleaner instance are in section 14.5.1 on page 246. Configure the parameters are follows:

- **Directory**: this should be the location of the back-fill directory for the `gdi-link-tx` instance. The first instance uses the directory `/var/lib/gdi-link-tx.default` by default but, if you have changed this, you should specify the new location here.
- **File sorting**: set to "Lexical".
- **Maximum used space** or **Maximum number of files**: use either or both of these to limit how much back-fill is retained. The `gdi-link-tx` back-fill process generates one file per hour so setting Maximum number of files to 24, for example, will guarantee 24 hours of back-fill.

Note: The size of the back-fill files varies with data compressibility. If you specify a large number of files to be retained, you should monitor operation for several hours and then check that the configured number of files multiplied by their average size does not exceed the free space on the disk.

13.1.2 The GDI link receiver

To configure a GDI link receiver via the web interface select:

Configuration → Services → GDI

or

Configuration → All options → System services → GSI

Now select "gdi-link-rx - GDI link receiver".

The screen shows a list of all GDI link receiver instances that have been configured on the acquisition module.

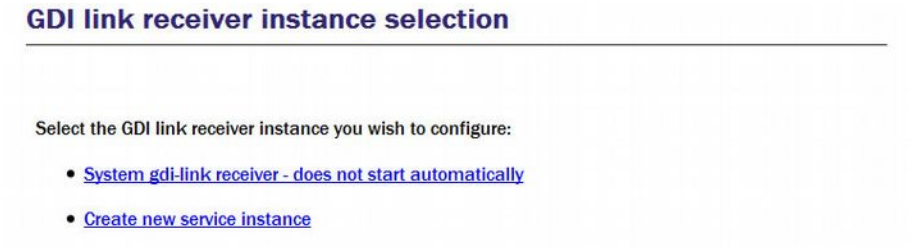

In most cases, you will only need a single instance and you can enable and reconfigure the Default Instance for your requirements.

13.1.2.1 Configurable parameters in simple mode

The configurable parameters for the GDI link receiver have three tabbed pages in simple mode: General, Network and Servers.

General

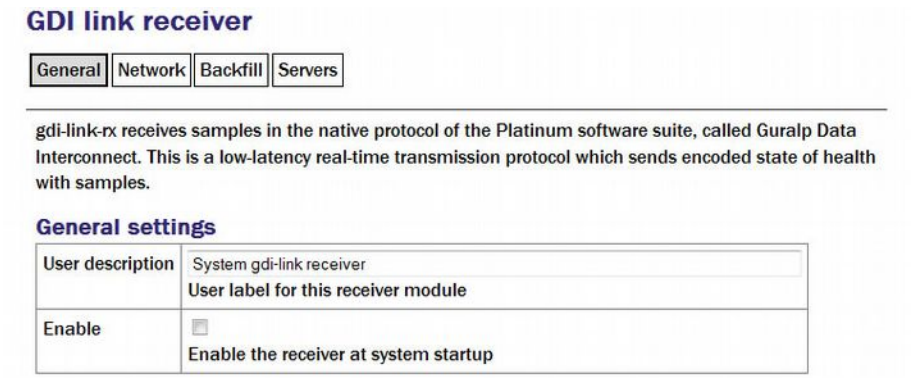

User description: Enter a descriptive name for the instance: this is useful if you are configuring multiple instances but, in most cases, this can be set to the hostname of the acquisition module.

Enable: Enables the receiver at system startup if selected.

Network

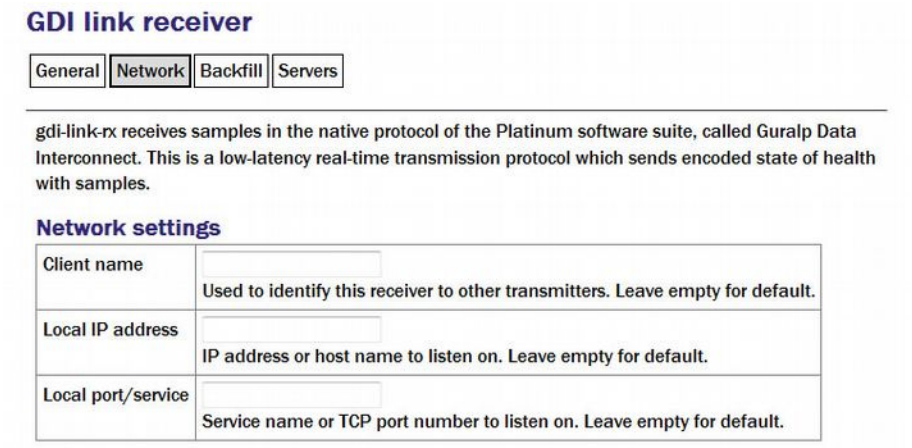

Client name: Optional name visible from the GDI link server.

Local IP address: If the GSL-EAM has multiple network addresses, you can limit the GDI link receiver to use only one of them by entering it in this field. If left blank, the receiver will listen on all configured addresses.

Local port/service: The default GDI link port is 1566 but this can be over-ridden if desired - you would want to do this if you had multiple instances running on the same address by entering a port name or number.

Back-fill

Back-fill is the process whereby missing data are recovered. It can be disabled if desired but, in most cases, you should leave this enabled.

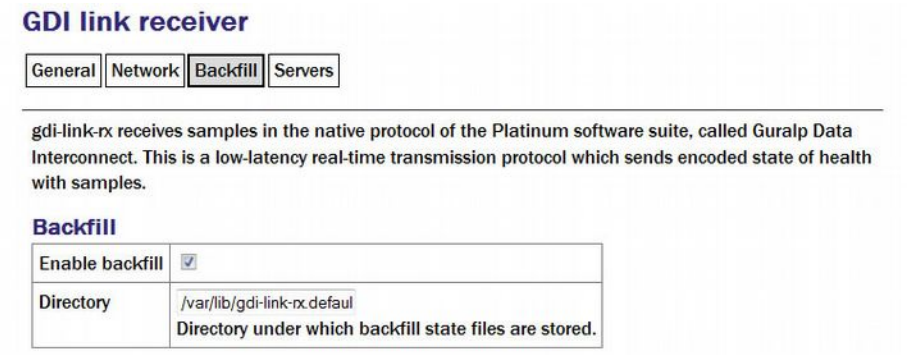

Enable backfill: If selected enables the back-fill transmission.

Directory: The default directory for back-fill data.

Servers

The servers tab contains a table within which you can configure the GDI link servers to which this receiver should listen.

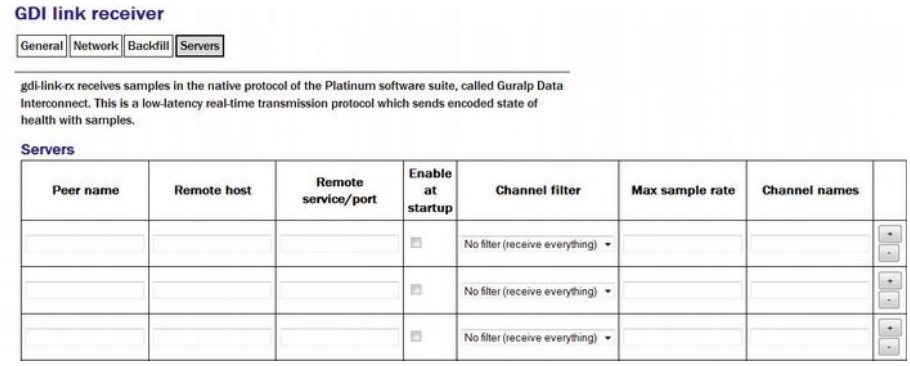

For each client, you should set:

- **Peer name**: Should match the server name configured in the client (receiver) at the remote end of the link
- **Remote host**: The DNS name or IP address of the GDI link server

- **Remote service/port**: The default is 1566 but, if you have configured a different port on the GDI link client, you should enter the same port here.
- **Enable at startup:** Controls whether the service is enabled at start-up.
- **Channel filtering**: The drop down has three options: No filtering, By sample rate or By channel names.
- **Max sample rate:** Enter the maximum sample rate you with to receive from the client
- **Channel name:** Enter the channels names for the data streams you wish to receive.

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

13.1.2.2 Configurable parameters in expert mode

The following additional configuration parameters are available by clicking the “Expert” button at the bottom of the form and selecting the 'Advanced' tab.

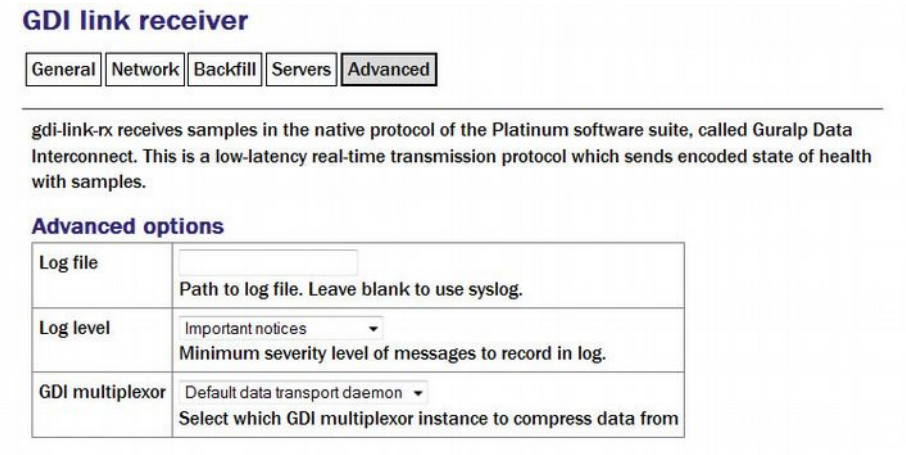

Log file: It may sometimes be desirable, for debugging purposes, to separate log messages for this transmitter from the standard system log. The text field can be populated with a path name which will then be used for dedicated logging. If left blank, logging occurs (via the standard Linux syslog facility) to `/var/log/messages`.

Log level: The drop-down menu controls the level of detail present in log messages, whether to syslog or to a dedicated log file. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational message

- Important notices
- Warnings

GDI multiplexer: In most configurations, all data for all transmitters is taken from a single multiplexor, as described in section 6.1 on page 65. For more complex configurations, it is possible to configure multiple multiplexers, each with their own set of input and output services. In these situations, the drop-down menu can be used to select a multiplexer instance with which to associate this transmitter. The menu offers a list of currently configured multiplexers.

13.2 Güralp Secure TCP Multiplexer

The Güralp Secure TCP Multiplexer (GSTM) is a method by which TCP connections can be tunnelled in both directions over a single TCP connection. It is an essential tool in situations where local network service providers cannot provide fixed (static) IP addresses.

For example, in an installation involving a single, central data collection point and multiple, remote sensor sites it is sometimes impractical for the sensor sites to be allocated static IP addresses. The use of GSTM allows the remote sites each to initiate a single GSTM TCP connection to the central site. Once established, further TCP connections can be initiated in either direction: their packets are tunnelled over the GSTM link.

If no sites in an array can be assigned fixed IPs, including the central data collection point, a GSL-EAM or GSL-NAM can be installed anywhere that has a fixed IP address and used as a communications hub. All sites initiate GSTM connections to the hub, which can then act as a communications router, forwarding individual connections as required.

The initial link is established from a GSTM client to a GSTM server.

13.2.1 The GSTM Client

To configure a GSTM client from the web interface, select:

Configuration → Services → Network

or

Configuration → All options → System services → Network

Now select “gstm-client - Güralp secure TCP multiplexor client”.

The screen shows a list of all GSTM client instances that have been configured on the acquisition module. To create or edit a GSTM client click the appropriate link.

13.2.1.1 Configurable parameters

The configurable parameters for the GSTM client have three tabbed pages: General, Server and Link.

General tab

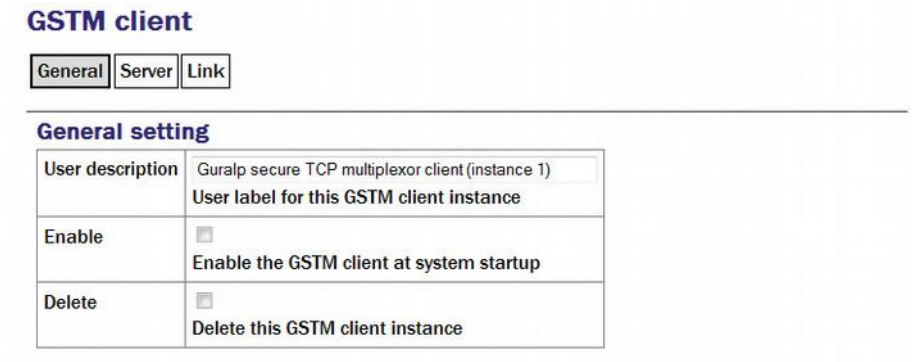

User description: Allows you to enter a mnemonic description of this instance, which may be useful if you intend to run multiple instances.

Enable: Causes this service to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

Server tab

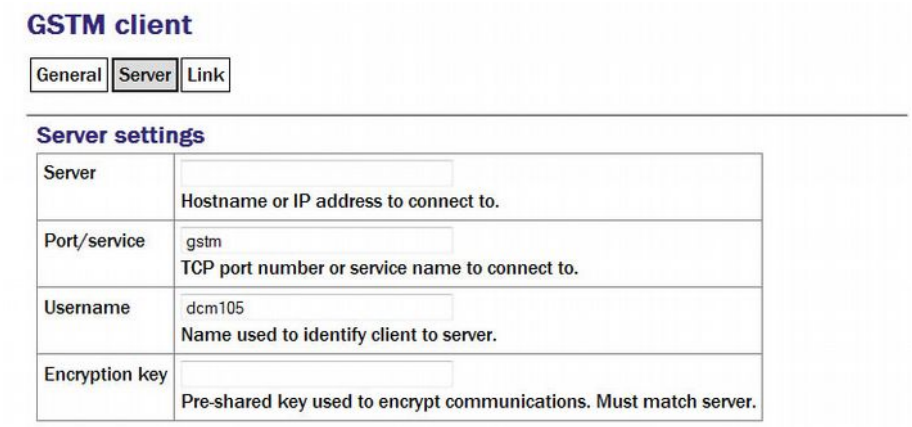

Server: The client will automatically connect to the GSTM server specified by a DNS name or IP address entered into the "Server" field.

Port/service: The connection to the server will be made using the TCP port specified by a service name or number in the "Port/service" field.

Username: The client identifies itself to the server using a username: this can usefully be set to the hostname of the acquisition module.

Note: The username is the means by which the server refers to this client.

Encryption key: GSTM communication is encrypted using TLS. Each end of any GSTM link needs to be configured with the same pre-shared key. If the server has already been configured, the server administrator will give you a value for the "Encryption Key" field; Otherwise, enter a random string into this field and let the person administering the server know what you have used. A command-line tool, gstm-genpsk, is provided to generate suitable random strings.

Link tab

Watchdog interval: If a configured link carries no traffic for an extended period, the client will send "watchdog" packets to the server. This serves two functions: it reassures the client that the link is still usable and it defeats any "automatic disconnect on idle" mechanisms which may be active on some links. The time, in seconds, between such watchdog probes can be configured by entering a value in the "Watchdog interval" field.

Restart interval: If the GSTM link fails for any reason, it is automatically restarted. There may be situations where the link cannot be restarted so, to prevent almost continuous restart attempts and consequent processor thrashing, a time delay is implemented between a link failure and a restart attempt. This defaults to thirty seconds but a different value can be configured if desired by entering it in the "Restart interval" field.

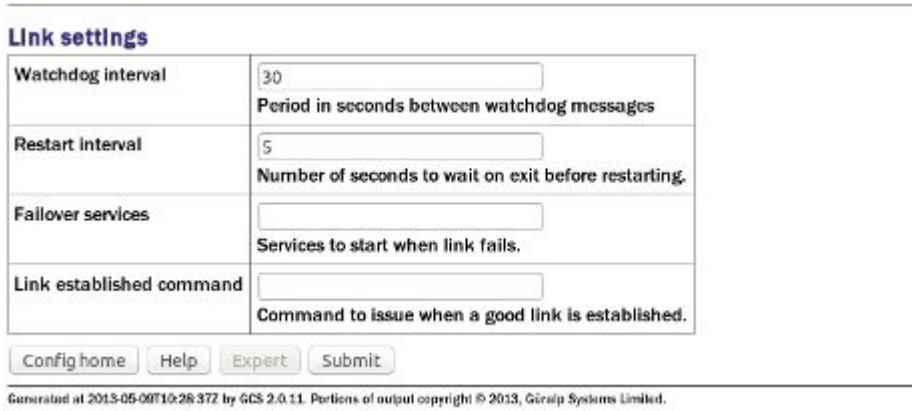

Failover services: If the watchdog packets do not elicit a response from the server, the link is assumed to have failed and, optionally, an additional service can be started in response. This will typically be another GSTM client in order to establish a back-up link. The GSTM client to be started should be identified here by its service descriptor, which takes the form `gstm-client.`*n* where *n* is an integer: 0 denotes the first configured client instance, 1 the second and so on.

Link established command: When the GSTM link is established or re-established, it is possible to run an arbitrary command. Any text entered in the "Link established command" field is passed to the Linux shell for execution, so this can be a single command or the path to a shell script to execute multiple commands. Please contact Güralp support if you need assistance with this feature.

13.2.2 The GSTM Server

GSTM clients initiate connections to GSTM servers.

To configure a GSTM server from the web interface, select:

Configuration → All options → System services → Network

or

Configuration → Services → Network

Now select "Güralp secure TCP multiplexor server".

The resulting screen shows a list of all GSTM server instances that have been configured on the acquisition module. To create or edit a GSTM server click the appropriate link.

13.2.2.1 Configurable parameters in simple mode

The configurable parameters for the GSTM server have four tabbed pages in simple mode: General, Server, Clients and Forwards.

General tab

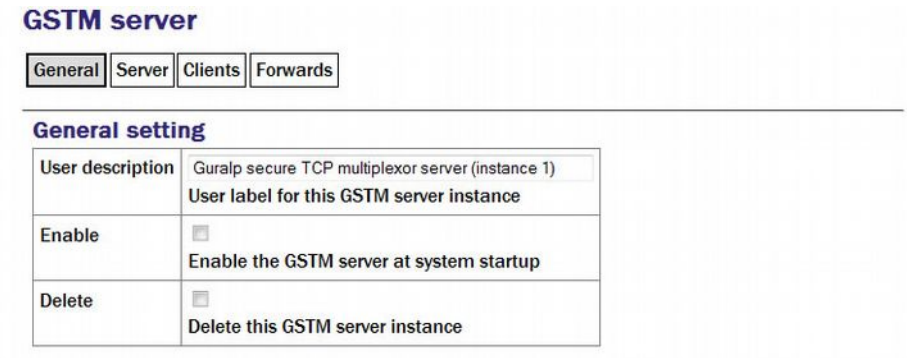

User description: Enter meaningful names here to help distinguish if several instances are to be created.

Enable: Causes this service to start automatically when the system is re-booted. If this check-box is cleared, the service will need to be started manually (from the "Control" → "Services" menu)

Delete: Causes the configuration for this instance to be removed from the system when the form is submitted.

Server tab

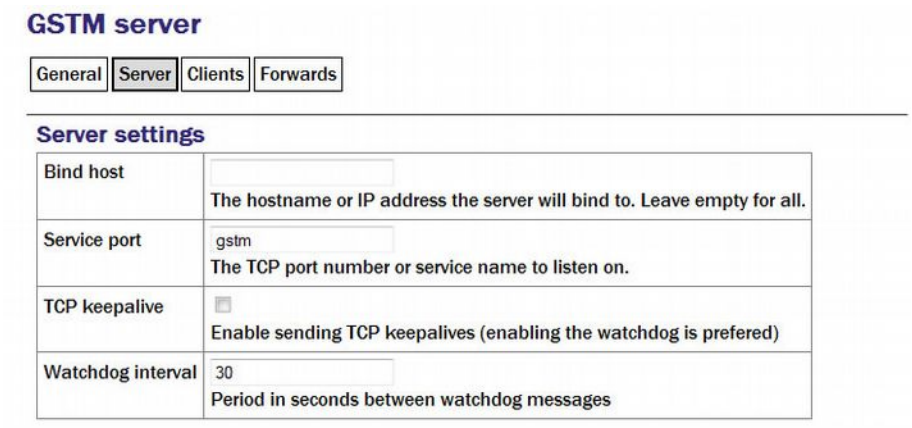

Bind host: If the acquisition module has multiple IP addresses, the GSTM server can be constrained to listen on only one of them by entering its address in the field. If this field is left empty, the server will listen on all available IP addresses.

"Bind host" Enter the service name or port number on which you want the server instance to listen. If you are configuring multiple server instances, each needs a unique service/port. The service name to port number mapping is stored in the standard Linux file `/etc/services`, which can be edited from the command line.

TCP keepalive: The server is capable of generating TCP keep-alive packets in order to defeat any automatic "disconnect on idle" mechanisms which may be present on the link. Tick the check-box to enable this feature.

Watchdog interval: Like the GSTM client, the GSTM server also generates watchdog packets to monitor for link failure. These may also be more effective at maintaining a link than TCP keep-alives because some network devices automatically block and/or spoof keep-alive packets. Watchdog packets are sent after a certain amount of time when the link appears to be idle and at regular intervals thereafter until traffic is detected. This time interval can be configured by entering an integer value, in seconds, in the field.

Clients tab

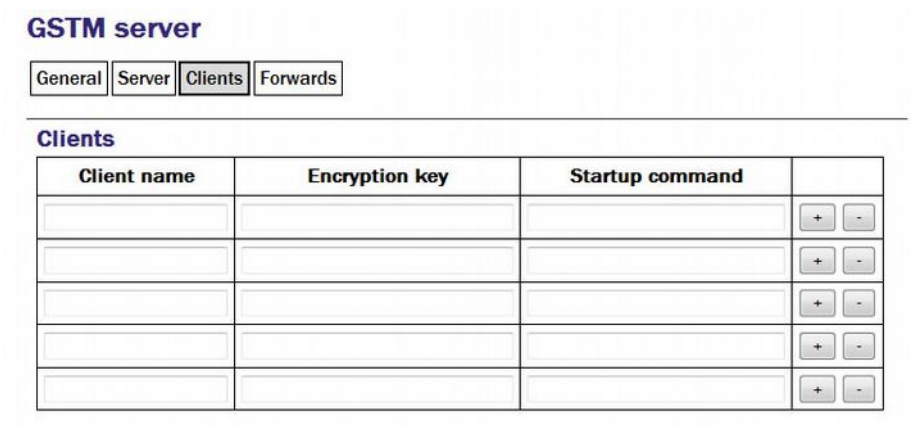

A single GSTM server instance can accept simultaneous connections from multiple clients. For each client enter:

Client name: Should contain the username, as configured in the GSTM client.

Encryption key: Should match that configured in the client (see the notes about the client configuration in section 13.2.1 on page 211 for more information).

Startup command: The GSTM server can run an arbitrary command when a client successfully initiates communication. Any text entered into this column is passed to the Linux shell for execution. The path to a shell script can be entered here if it is required to run multiple commands. Contact Güralp support if you need assistance with this feature.

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

Forwards tab

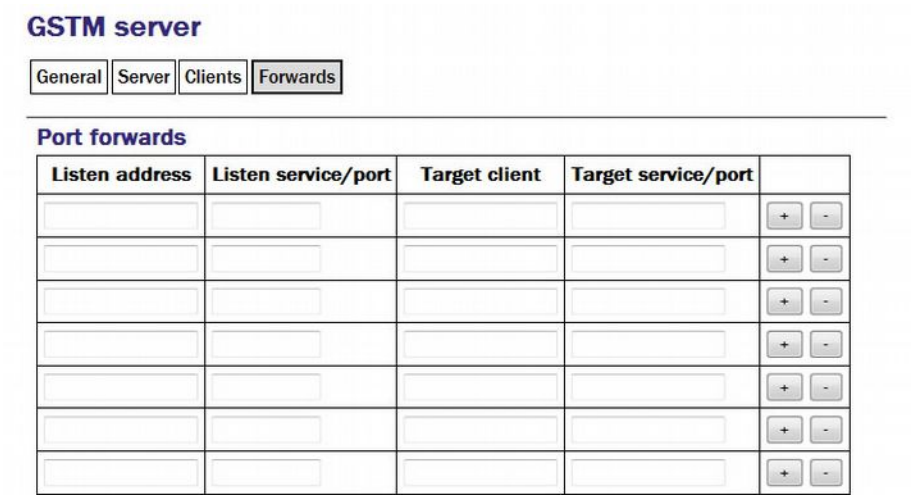

An active GSTM link can forward arbitrary TCP connections to the clients from any host that can access the server. For each desired port forward, enter the following information:

Listen address: By default, the port-forward will listen for incoming connections on all available interfaces. If this field is populated with an IP address, the port-forward will only listen for connections addressed to that address. This option is only useful on systems with multiple IP addresses.

Listen service/port: The port-forward will listen for incoming connections addressed to this port and forward those connections across the network to the client and port specified in the subsequent fields. The port can be specified as either a name or a number. The mapping from port names to port numbers is configured by the conventional Linux file `/etc/services` which can be edited from the command line if required.

Target client: Connections accepted by the port-forward will be forwarded to the client identified by the contents of this field. The field should be populated with the **Username** that was configured on the "Server" tab of the GSTM client (as described in section 13.2.1.1 on page 212).

Note: The client is specified by **Username** rather than I.P. address, as might be expected. This is to support clients that might be connecting from dynamic I.P. addresses.

Target service/port: Connections accepted by the port-forward will be forwarded to this port on the remote client. The port can be specified as either a name or a number. The mapping from port names to port numbers is configured by the conventional Linux file `/etc/services` which can be edited from the command line if required.

Clicking the + button on any row will open a new row. In the same way, rows can be deleted by clicking the corresponding - button.

14 Monitoring Operations

This chapter details how to monitor and control the acquisition module. Some functionality is only available from the command line (when connected via a serial cable or via SSH over the network - see section 3.2.5 on page 27 for details on how to do this) and other features are only available via the web interface. This chapter will describe both.

14.1 Diagnostics and the Summary screen

14.1.1 System Status

To view the overall system status click on the “Home” link in the breadcrumb trail or the “System status” link in the main menu.

The top part of the status screen is a tabbed list of devices connected to the acquisition module. Clicking on a tab will display a status report for that device.

The system has predefined warning and error levels which are displayed as coloured dots and are defined as:

- **Green:** Level is 70% or above – system OK
- **Amber:** Level is below 70% - system warning
- **Red:** Level is below 40% - system error/malfunction
- **Grey:** There is no status information.

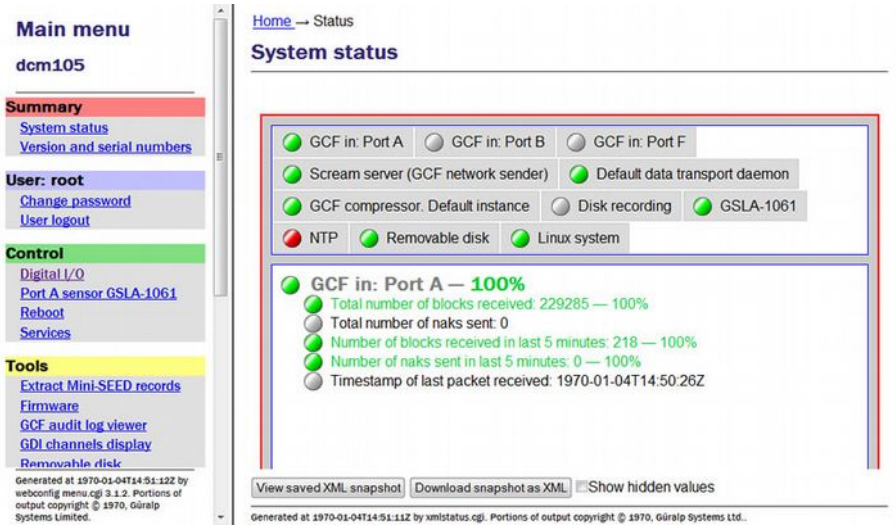

14.1.2 System Log

The most important source of diagnostic and debugging information is the system log facility ("syslog"). This logs all messages from programs and from the Linux kernel. At present, this can only be viewed from the command line.

To view the system logs, you can use the **`tail`**, **`less`**, **`grep`** or **`vi`** commands to inspect the file `/var/log/messages` - older files are available as `/var/log/messages.1`, `/var/log/messages.2` etc.

Use the following syntax to view the files:

- **`tail /var/log/messages`**
  Views the last few entries.
- **`tail -f /var/log/messages`**
  As above but, after printing the last ten lines of the file, continues to run, printing each new line as it is added to the log. Type Ctrl + C to stop the output and return to the command line.
- **`less /var/log/messages`**
  Views the whole log file; use the ↑, ↓, Home and End keys to navigate and the Q key to exit.
- **`vi /var/log/messages`**
  Those users familiar with the `vi` text editor may wish to use it as the most powerful way to view log entries.
- **`grep -i '`*`string`*`' /var/log/messages`**
  Searches for a string or pattern in the log file. This search is case insensitive (-i flag).

`grep` is a very powerful tool for searching for patterns. For more information, see the section on Regular Expressions in the **`grep`** manual page at http://man-wiki.net/index.php/1:grep

14.1.3 Incoming Data

The status web-page has one box for each GCF acquisition process. This box will be updated every minute to reflect the number of packets that have been acquired.

To view details of incoming GCF format data using the command line, enter the command **`gdi-top`**. This will displays real-time information about each packet arriving, until interrupted by the operator typing Q.

14.1.4 Software build number

The “Software build number”, as displayed on the “Linux system” tab of the status screen, may be useful for your own records or when requesting technical support. See section 5.2 on page 52 for more details.

14.2 Warning and error monitoring

In addition to the status indicators listed in section 14.1.1 on page 218, the status of all components can be monitored and, if the level falls below a set percentage, a selected output line can be switched from low to high for use by an external indication system.

To configure the status iolines using the web interface select:

Configuration → All options → Status iolines assertion configuration

To configure the status iolines from the command line, start gconfig and select “Status iolines assertion configuration” from the top level menu.

14.2.1 Configurable parameters in simple mode

The configurable parameters for recording data in simple mode are shown below:

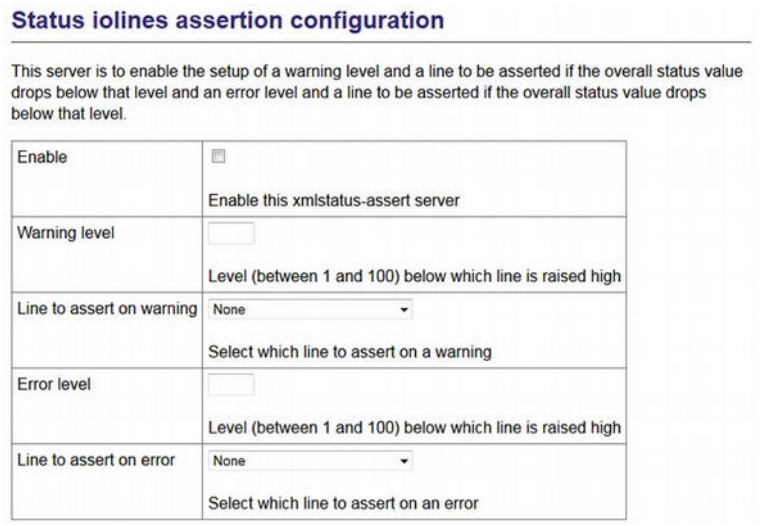

Enable: Activates the assertion server.

Warning level: Set the level at which the warning signal is to be triggered.

Line to assert on warning: Select the output line from the drop-down list. All output lines are displayed so care should be taken to choose a suitable output line.

Error level: Set the level at which the warning signal is to be triggered. The error level must be less than the warning level.

Line to assert on error: Select the output line from the drop-down list. All output lines are displayed so care should be taken to choose a suitable output line.

14.2.2 Configurable parameters in expert mode

An additional field is displayed in expert mode:

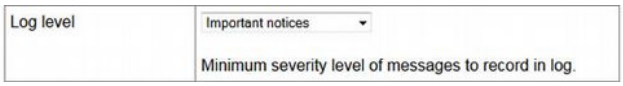

Log level: The drop-down menu controls the level of detail present in log messages. Not all of the standard syslog logging levels are available. The menu offers a choice (in order of decreasing detail) of:

- Debugging information
- Informational messages
- Important notices
- Warnings

14.3 The Control Menu

The "Control" menu of the web interface is a dynamic menu with content that changes depending on the which devices are attached. Two items on this menu are always present: "Reboot" and "Services". CMG-NAM units fitted with RAID arrays will also have a "RAID array services" menu item. Other items will appear as required, depending on both the underlying hardware and attached devices.

14.3.1 Digital I/O (power control and anti-tamper monitoring)

The acquisition module hardware can be fitted with optional sensors to monitor and switch the voltages and currents being supplied to the CMG-EAM and also to devices connected to the CMG-EAM ports, such as digitisers. A program on the CMG-EAM runs constantly in the background and monitors the sensors if they are fitted. The same program can monitor the anti-tampering lines, where fitted.

To control the Digital I/O from the web interface, select:

Control → Digital I/O

The screen displays the current digital I/O status for each monitored line.

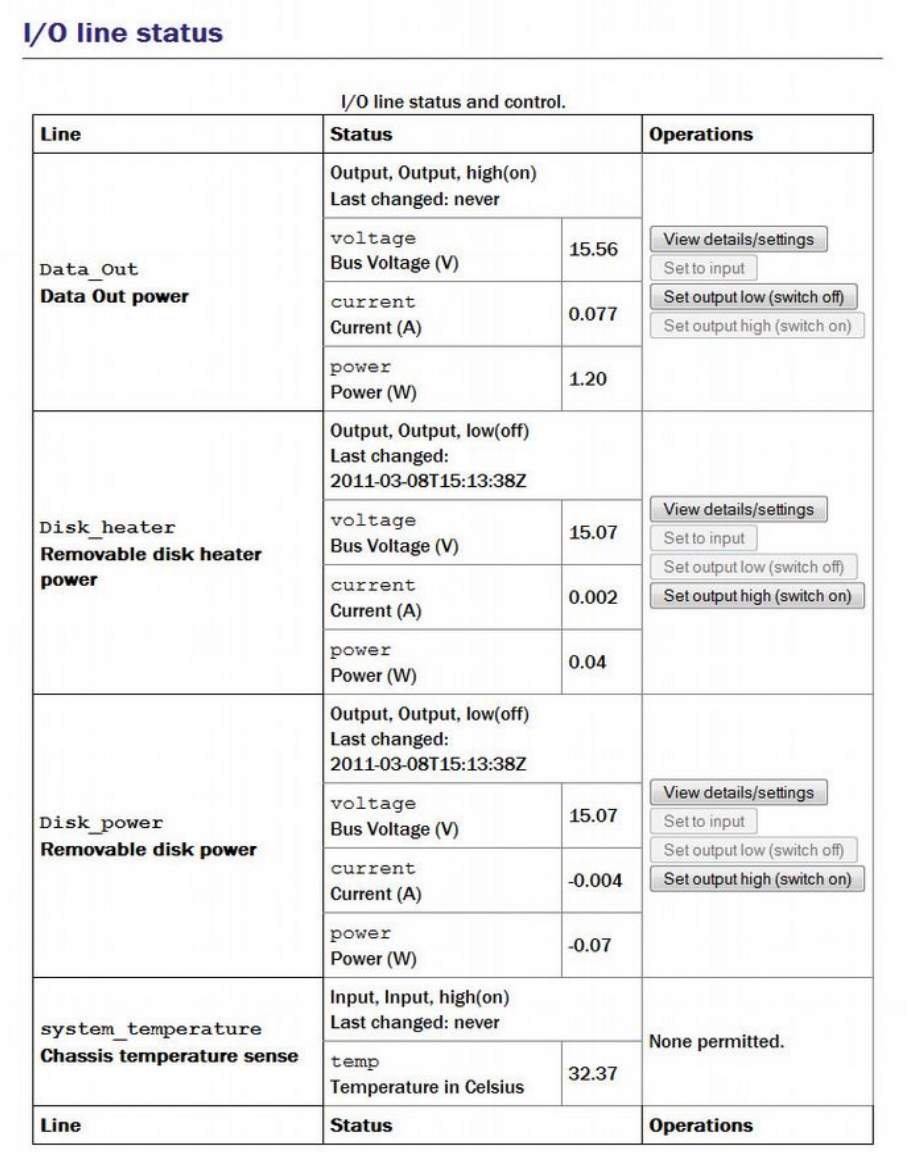

I/O line status

I/O line status and control.

Line	Status		Operations
Data_Out **Data Out power**	Output, Output, high(on) Last changed: never		
	voltage Bus Voltage (V)	15.56	View details/settings Set to input
	current Current (A)	0.077	Set output low (switch off) Set output high (switch on)
	power Power (W)	1.20	
Disk_heater **Removable disk heater power**	Output, Output, low(off) Last changed: 2011-03-08T15:13:38Z		
	voltage Bus Voltage (V)	15.07	View details/settings Set to input
	current Current (A)	0.002	Set output low (switch off) Set output high (switch on)
	power Power (W)	0.04	
Disk_power **Removable disk power**	Output, Output, low(off) Last changed: 2011-03-08T15:13:38Z		
	voltage Bus Voltage (V)	15.07	View details/settings Set to input
	current Current (A)	-0.004	Set output low (switch off) Set output high (switch on)
	power Power (W)	-0.07	
system_temperature **Chassis temperature sense**	Input, Input, high(on) Last changed: never		None permitted.
	temp Temperature in Celsius	32.37	
Line	**Status**		**Operations**

Clicking the “View details/settings” button produces the following screen:

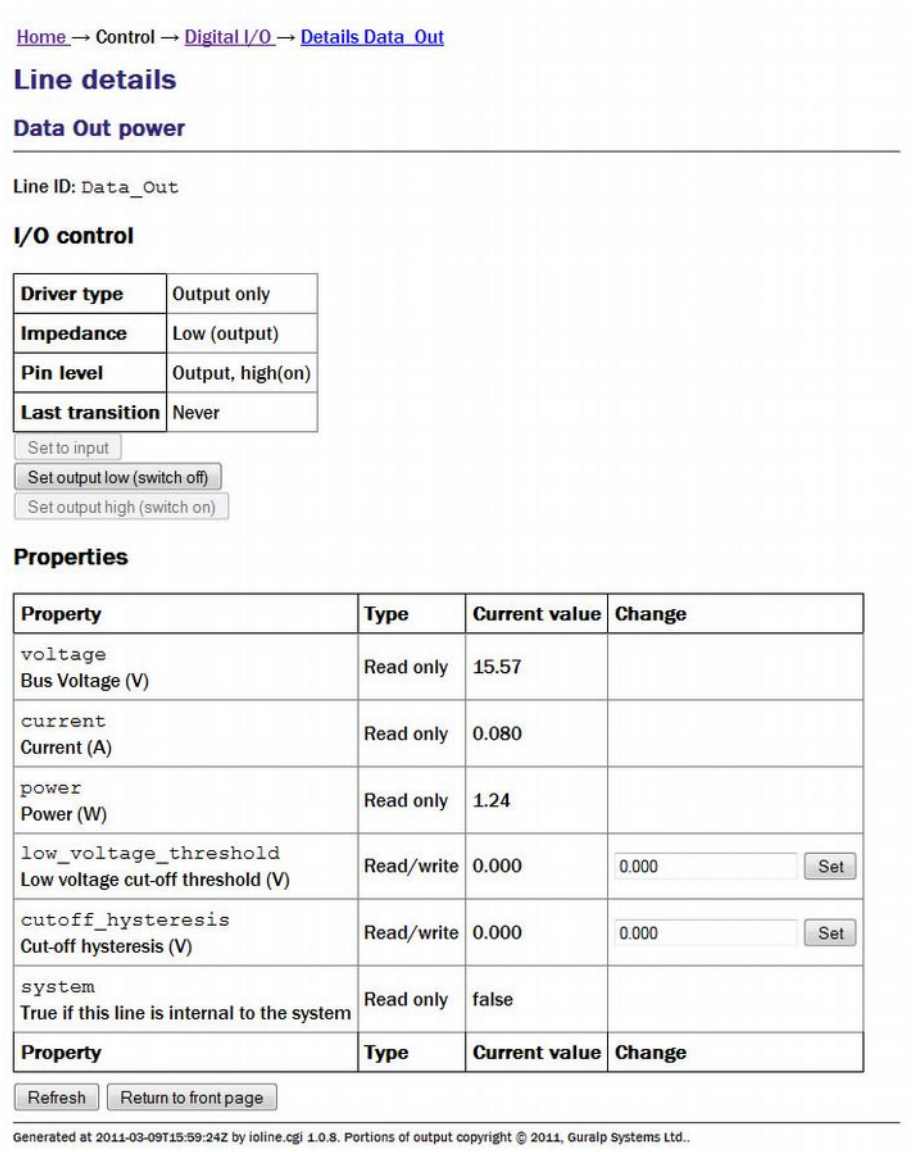

Home → Control → Digital I/O → Details Data_Out

Line details

Data Out power

Line ID: Data_Out

I/O control

Driver type	Output only
Impedance	Low (output)
Pin level	Output, high(on)
Last transition	Never

Set to input
Set output low (switch off)
Set output high (switch on)

Properties

Property	Type	Current value	Change
voltage Bus Voltage (V)	Read only	15.57	
current Current (A)	Read only	0.080	
power Power (W)	Read only	1.24	
low_voltage_threshold Low voltage cut-off threshold (V)	Read/write	0.000	0.000 Set
cutoff_hysteresis Cut-off hysteresis (V)	Read/write	0.000	0.000 Set
system True if this line is internal to the system	Read only	false	
Property	**Type**	**Current value**	**Change**

Refresh Return to front page

Generated at 2011-03-09T15:59:24Z by ioline.cgi 1.0.8. Portions of output copyright © 2011, Guralp Systems Ltd..

The first part of the page duplicates the information shown on the main screen. The remainder of the page provides a more detailed report.

Note: The user must be in the 'Peripheral control (gpio)' group to make any changes. All other users will only be able to view the settings.

Low voltage cut-off threshold: It is possible to protect attached devices from under-voltages by setting a value. If the monitored voltage falls below this value, it is automatically turned off.

Cut-off hysteresis: To prevent rapid, repeated power-cycling of attached devices, a hysteresis value should be set. The monitored voltage must rise

above the sum of the threshold voltage and the hysteresis voltage before the supply will be re-enabled.

Click the "Return to front page" button to display the main Digital I/O summary again.

14.3.2 Digitiser/Sensor Control

The acquisition module allows *control* of attached digitisers and sensors. To *configure* digitisers, see section 8on page 100.

Note: ***Support for "active high" sensors***: Control of sensors is implemented using logic lines in the cable between the sensor and the digitiser, one for each function (lock, unlock, centre, etc). Standard Güralp sensors use "active low" logic, which means that the logic lines are normally floating at +5V with respect to the logic ground (the sensors have a pull-up resistor and the digitiser presents a high impedance). When a particular function is to be activated, the digitiser grounds the relevant logic line, triggering the appropriate action at the sensor.

Some special-order Güralp sensors use "active high" logic, where the lines are normally grounded and allowed (*not* driven) to float high in order to trigger the associated control function. CMG-DM24S*x*EAM units can be configured to support active high instruments by connecting pins V and Y of the SENSOR-A connector using a special cable. These cables can be ordered from GSL (as "active high sensor cables for DM24S*x*EAM") or manufactured by modifying standard GSL sensor cables.

Sensors can be controlled via the web interface or the command-line, The web interface is simpler and requires no detailed knowledge of the attached devices. The command line interface is more powerful but requires detailed knowledge of the digitiser's command line interface and the manual for the digitiser in question should be referred to for further details.

14.3.2.1 Instrument Control - Web interface

The main menu of the web interface adapts to include additional options when the system detects attached digitisers and/or digital sensors.

Extra items in the "Control" menu are associated with digitiser serial numbers. If a digitiser has two sensors attached to it, it is recommended that an extra serial number be added to the digitiser. See section 8.3 on page 117 (or the relevant digitiser manual) for information on how to configure this.

The sensors are listed in the following format:

```
Port A instrument C914-3K55
```

and indicate the physical connection (Port A) to the attached device, the device type (instrument) and its serial number (C194-3K55). If the instrument is connected via TCP, the EAM connection is shown as host:port.

Note: If using a DM24 MkIII seven-channel digitiser, there is no control over the second sensor (Sensor B). Any attempt to control the second sensor will act on Sensor A.

Selecting a device takes you to the "Instrument Control" page. From here, one can query mass positions, lock, unlock and centre the sensor masses and perform calibration functions.

The form is long and is displayed here in sections:

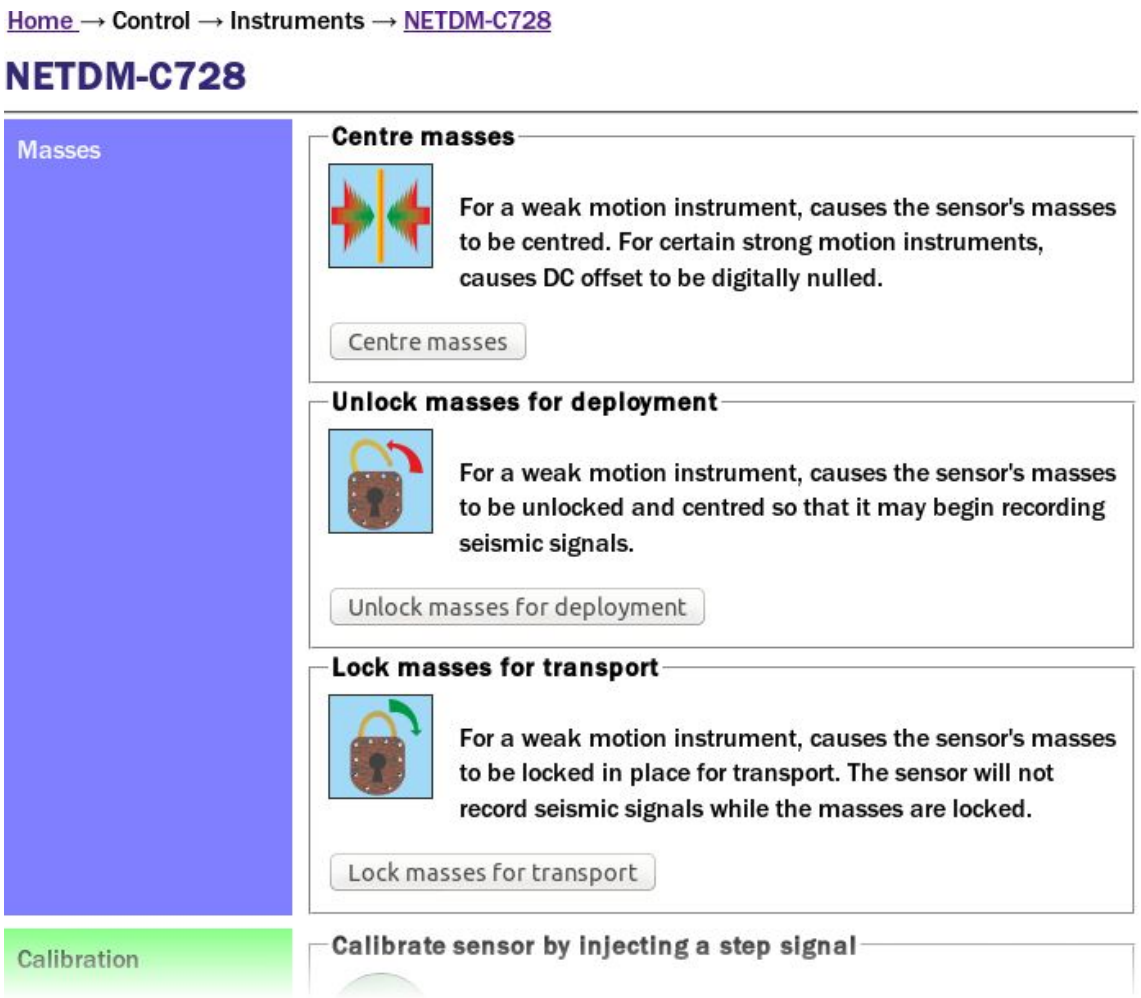

The first part of the screen offers controls to centre, unlock and lock the instrument's masses.

Note: These controls invoke the appropriate digitiser commands. An indication of success means that the digitiser has accepted the command but does not imply that it has executed correctly.

The next part of the screen provides controls to initiate calibration. The step-calibration section provides controls for selection of components, duration and amplitude. Not all instruments support calibration of individual components.

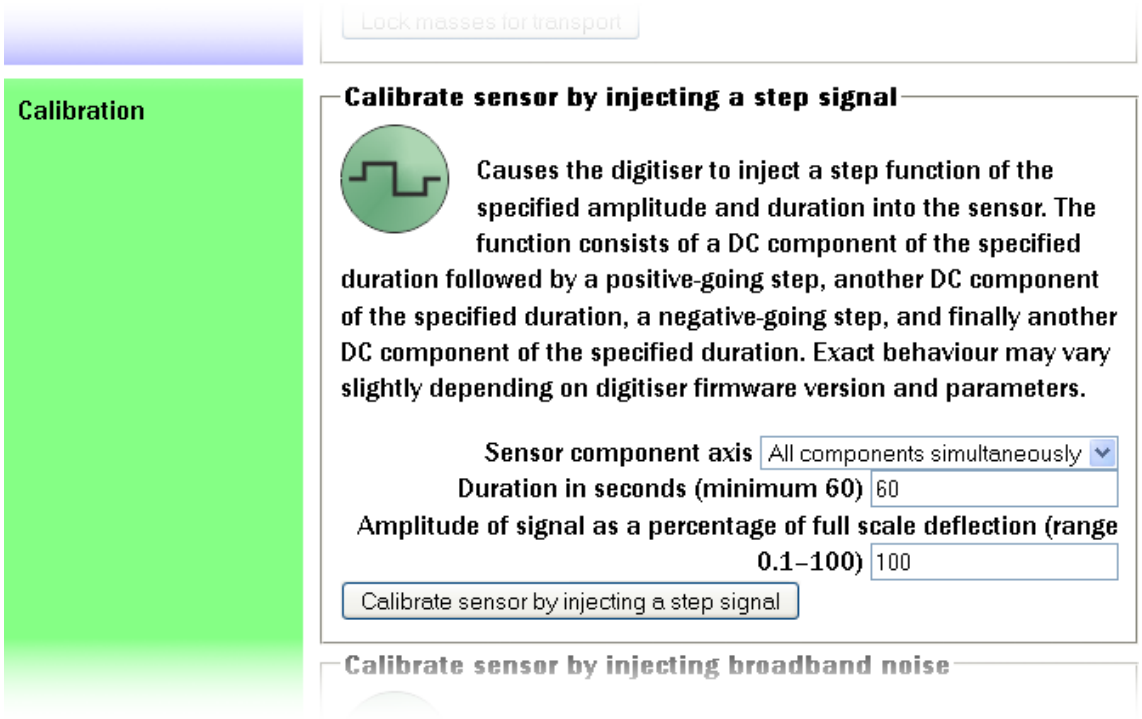

The broadband noise calibration section provides controls for selection of components, duration and amplitude. Not all instruments support calibration of individual components.

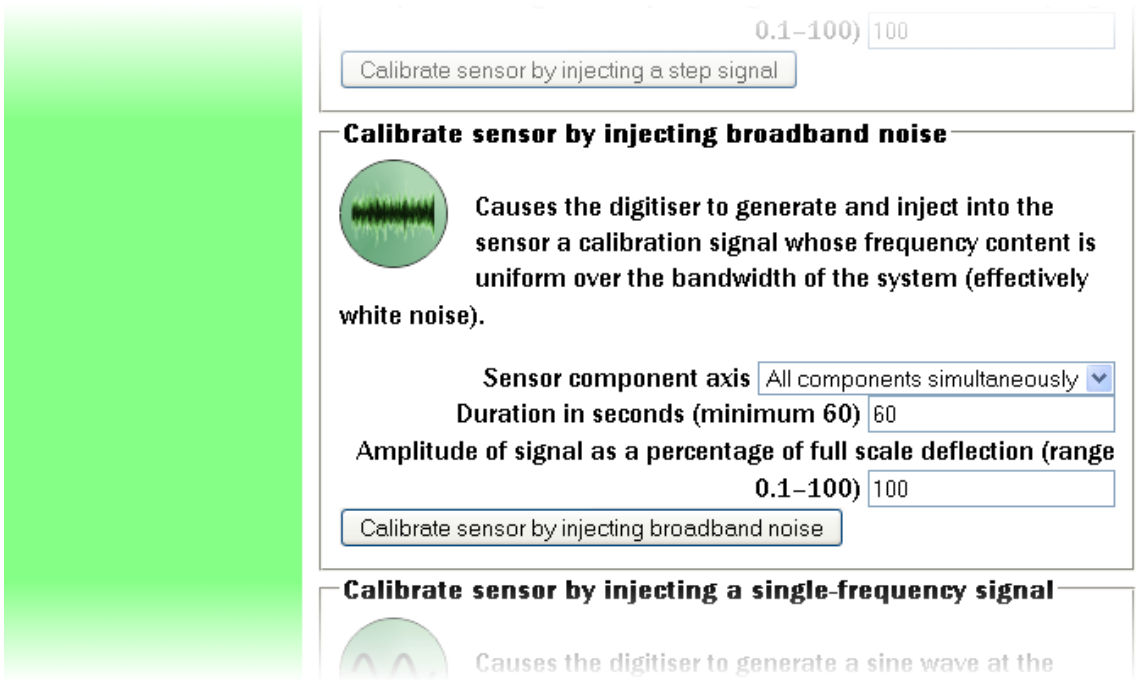

The sine-wave calibration section provides controls for selection of components, duration, amplitude and frequency. Not all instruments support calibration of individual components.

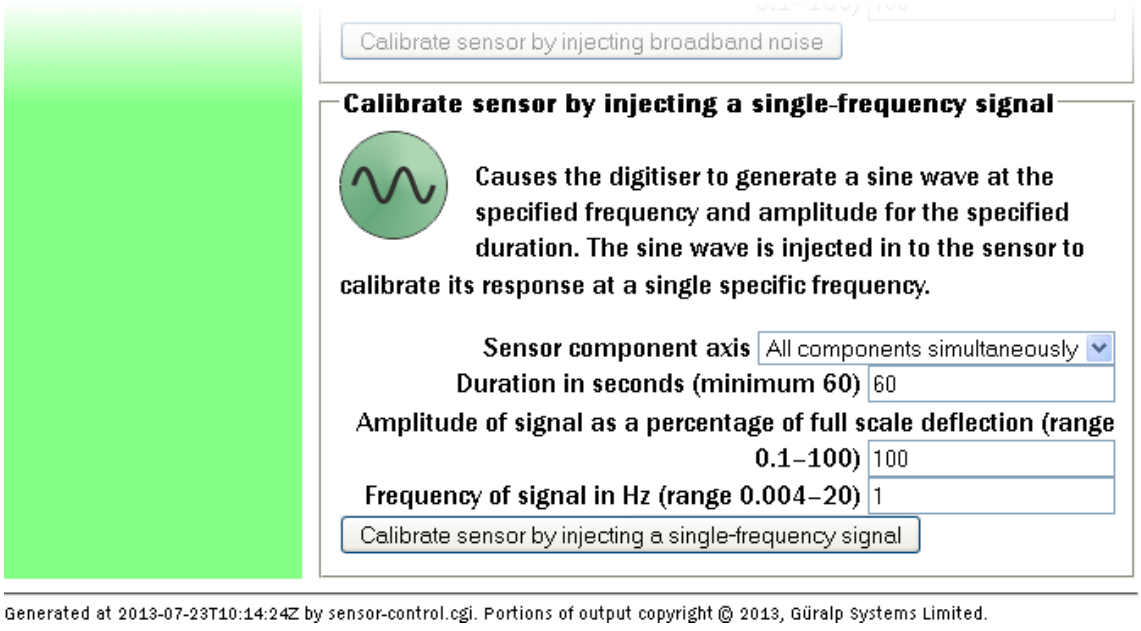

Note: The **Amplitude** value in all calibration dialogues refers to an arbitrary full-scale output from the calibration circuitry. It can be left at 100% unless clipping is observed, in which case it should be reduced until an undistorted output is observed.

14.3.2.2 Instrument Control - Command line

Platinum provides both a high-level and a low-level interface to connected digitisers. The low-level interface involves interacting directly with the command-line of the digitiser and is described at the end of this section.

The high-level interface is provided by the `adc-command` command, which takes a number of sub-commands as described below. Each command must be directed to a specific sensor and this must be specified as the first argument. When invoked with no arguments, or with the `--help` argument, a list of available targets (referred to as "modules") is displayed:

```
eam2010 ~ # adc-command --help
usage: adc-command <module> <command> [options...]

======== Available ADC modules ========
C914-3K55
C914-3K56
```

This is followed by a list of available sub-commands. In the descriptions below, ***module*** should be replaced by a name from the list provided by invoking `adc-command --help`.

- **adc-command *module* mass-centre**
  Perform a centring operation on the sensor's masses. This sub-command takes no options.
- **adc-command *module* mass-lock**
  Lock the masses for transportation. This sub-command takes no options.
- **adc-command *module* mass-unlock**
  Unlock the masses for deployment. This sub-command takes no options.
- **adc-command *module* calib-sine *arguments***
  Perform a sine-wave calibration according to the specified ***arguments***, which are given as a space-separated list of key=value pairs:
  - `component=[ ALL | Z | N/S | E/W ]` - specify the component(s) to be calibrated
  - `duration=`*m* - specify the duration, in minutes, where *m* is an integer value. For accurate results, the duration should be significantly longer than the response time of the instrument.
  - `amplitude=`*a* - specify the amplitude, where *a* is the integer percentage of the full output of the calibration signal generator. This can normally be given as 100 but should be reduced if clipping is noticed.
  - `freq=`*n* - specify the frequency (in integer Hertz) or period (in integer seconds) of the calibration signal, according to the setting of...
  - `freq_or_period=[ Hz | sec ]` - specifies whether the value associated with freq is interpreted as Hertz or seconds.
- **adc-command *module* calib-step *arguments***
  Perform a step calibration according to the specified ***arguments***, which are given as a space-separated list of key=value pairs:
  - `component=[ ALL | Z | N/S | E/W ]` - specify the component(s) to be calibrated
  - `duration=`*m* - specify the duration, in minutes, where *m* is an integer value. The calibration signal will first go negative for this duration, then back to zero for the same duration, then go positive for this duration, then return to zero again for the same duration. For accurate results, the duration should be significantly longer than the response time of the instrument.
  - `amplitude=`*a* - specify the amplitude, where *a* is the integer percentage of the full output of the calibration signal generator. This can normally be given as 100 but should be reduced if clipping is noticed.

- **`adc-command` *`module`* `calib-noise` *`arguments`***
  Perform a broad-band noise calibration according to the specified ***`arguments`***, which are given as a space-separated list of key=value pairs:
  - `component=[ ALL | Z | N/S | E/W ]` - specify the component(s) to be calibrated
  - `duration=`*`m`* - specify the duration, in minutes, where *`m`* is an integer value. For accurate results, the duration should be significantly longer than the response time of the instrument.
  - `amplitude=`*`a`* - specify the amplitude, where *`a`* is the integer percentage of the full output of the calibration signal generator. This can normally be given as 100 but should be reduced if clipping is noticed.

The acquisition module also provides the ability to connect to the terminal of any connected Güralp digitisers in order to configure their operation and control the attached sensors. To do this, connect to the acquisition module terminal as in section 3.3 on page 30 and run the "data-terminal" command.

```
eam999 ~ # data-terminal
```

Select the desired digitiser (using the  and  keys and Enter to select) from the list that is presented:

This will launch a minicom session (see section 16.3 on page 33), allowing you to communicate with the digitiser terminal. For example:

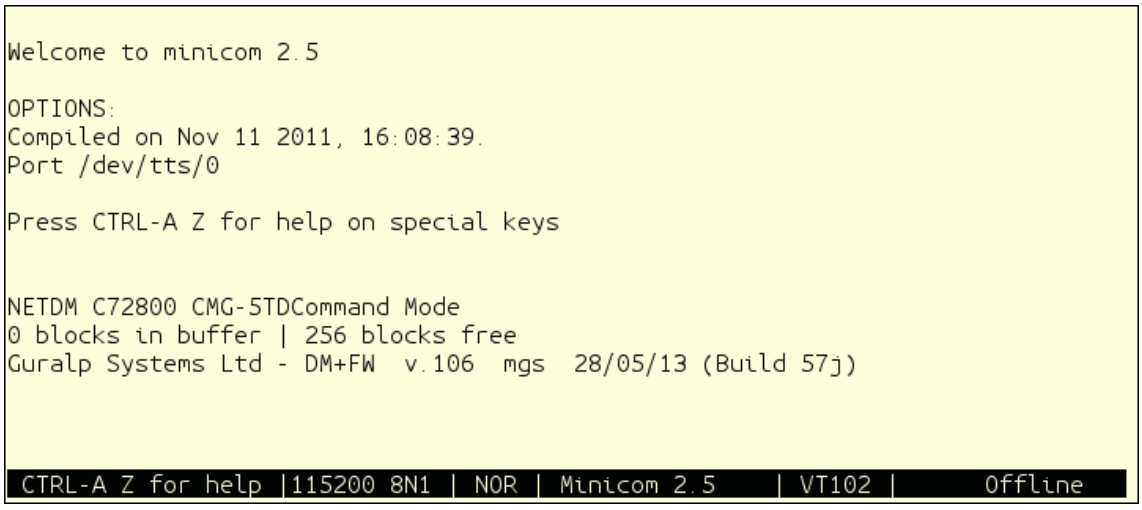

If the session closes due to a time-out (or you close it manually by issuing the `GO` command) then you will see the message `Killed by signal 15` and minicom will exit shortly thereafter.

14.3.3 Upgrading digitiser firmware

The latest digitiser firmware for each type of digitiser is included in the Platinum distribution so no Internet connection is necessary for the following tasks. You may wish first to upgrade the Platinum firmware (see section 5 on page 51), so as to be sure that you have the latest digitiser firmware available.

14.3.3.1 Upgrading via the web interface

To upgrade the firmware of a digitiser using the web interface select:

Configuration → Instruments → Port A instrument....

(or any other page under Configuration→Instruments). The system will check the version of the firmware currently loaded. If a newer version is available, an alert is displayed at the top of the page:

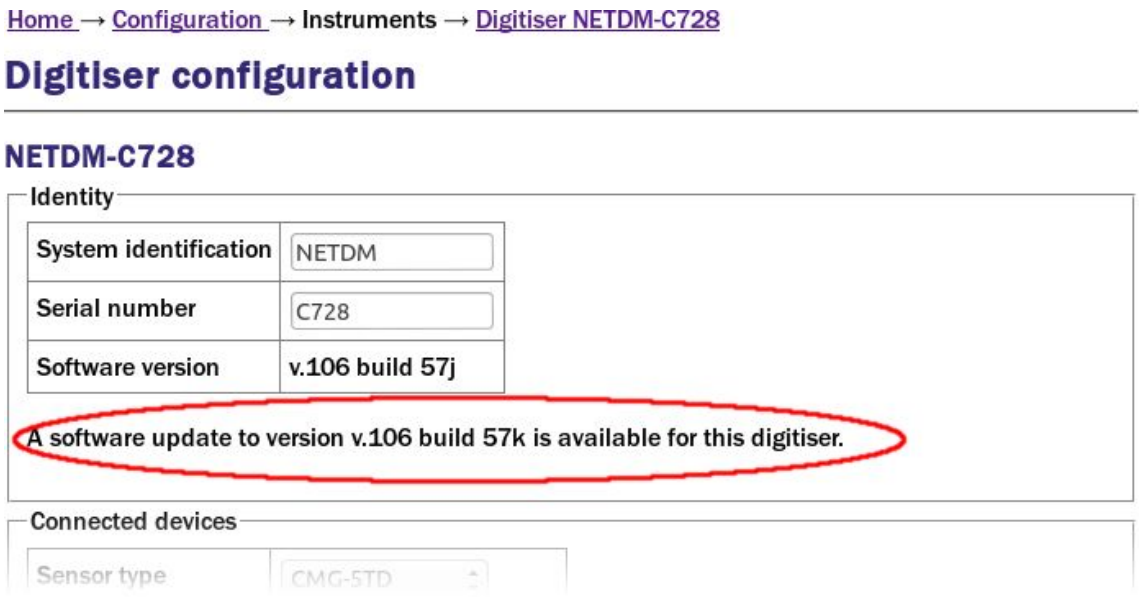

and a check-box is displayed at the bottom of the page:

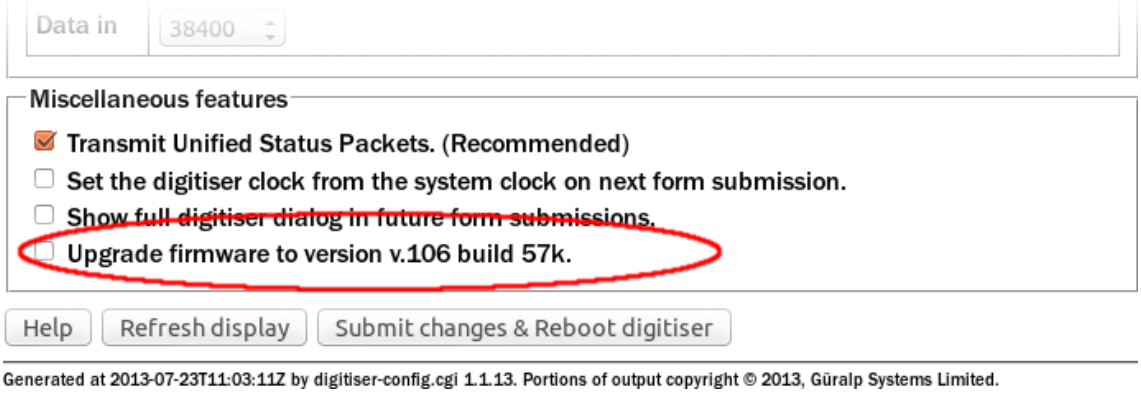

To upgrade, tick the check-box and click Submit changes & Reboot digitiser.

14.3.3.2 Upgrading from the command line

The dm24-upgrade command can be used to update the firmware of attached DM24 and CD24 digitisers. When invoked without arguments, it displays a detailed usage message.

When invoked with the name of a port (e.g. `Port A` or `PortA`) as an argument, the command will identify the type of attached digitiser and the current firmware revision and will upgrade it if necessary. For example, the command

```
dm24-upgrade PortA
```

will upgrade the firmware of the digitiser attached to Port A of the acquisition module, if necessary. The latest digitiser firmware for each type of digitiser is included in the Platinum distribution so no Internet connection is necessary.

Note: The currently-loaded DSP firmware version is identified from the label that was entered when the firmware was last upgraded. Because this label is not checked, any text could have been entered, which may confuse the upgrade command. If this is the case, please use the `–force` option, described below, to over-ride the version checking.

For more advanced uses, the following options are available. Multiple options can be separated by spaces and the command should always end with a port specifier. This can be in the form `Port A` or `PortA` or `/dev/ttyS0` (the Linux port name).

- `--trashfram`

  Performs a hard factory reset on a CMG-DM24 digitiser. Note that the values of all configuration parameters will be lost. Some of the following options allow basic reconfiguration following the reset.

- `--ids` *`sys ser`*

  A hard reset will normally erase the system ID and serial number of the digitiser. The `dm24-upgrade` command, when invoked with `--trashfram`, attempts to discover these values and reinstate them after the reset. This option allows you to specify new values for the system ID and serial number to be used instead of the discovered values. The argument *`sys`* specifies the new System ID to use and the argument *`ser`* specifies the serial number.

- `--samp # # # #`

  This argument allows you to specify new sample speeds to set for the decimation filter chain. This argument requires requires four numerical values to be supplied using the same format as for the

`SAMPLES/SEC` command (described in the DM24 manual). For example, the command

```
dm24-upgrade --samp 1000 500 100 50 PortA
```

will set tap 0 to 1,000 samples per second and taps 1, 2 and 3 to 500, 100 and 50 samples per second respectively.

- `--cont # # # #`

  This option specifies which streams will be enabled for continuous output at each tap. This argument requires requires four numerical values to be supplied using the same format as for the `CONTINUOUS` command (described in the DM24 manual). For example, the command

  ```
  dm24-upgrade --cont 1 0 0 7
  ```

  will turn on continuous output of the vertical component at tap 0 and of the three triaxial components at tap 3

- `--trig # # # #`

  This option specifies which streams will be enabled for triggered output at each tap. This argument requires requires four numerical values to be supplied using the same format as for the `TRIGGERED` command (described in the DM24 manual).

- `--gps-baud #`

  This option takes a single numerical argument which specifies the Baud rate of the GPS port.

- `--in-baud #`

  This option takes a single numerical argument which specifies the Baud rate of the Data in port.

- `--upgrade`

  Upgrade firmware files only (this is the default).

- `--downgrade`

  This option over-rides internal version number checks, allowing firmware to be downgraded as well as upgraded.

- `--force`

  This option causes firmware files are to be loaded even if they appear identical to the versions already loaded, or in cases where the installed version number cannot be decoded.

- `--boot` *`file`*

  This option specifies an alternative source file for DM24 Mk3 bootstrap firmware. If not specified, the upgrade command will look in `/usr/share/firmware/CMG-DM24mk3/` for files with names of the form `dm24mk3-boot*.img` and use the most recent version that it finds. It uses a lexical sort to determine which is the most recent. The argument *`file`* should be the path to a valid DM24 bootstrap firmware image.

- `--firm` *`file`*

  This option specifies an alternative source file for DM24 Mk3 system firmware or for CD24 firmware. If not specified, the upgrade command will identify the digitiser type and, for DM24s, look in `/usr/share/firmware/CMG-DM24mk3/` for files with names of the form `dm24mk3-[0-9]*.img` and use the most recent version that it finds. For CD24s, it will look for files with names of the form `dm32_`*`xx`*`_*.hex` in `/usr/share/firmware/CMG-CD24/` where represents the hardware type (e.g. 4A, 8mR, etc) as determined by interrogating the digitiser. In either case, the program uses a lexical sort to determine which file contains the most recent firmware buid. The argument *`file`* should be the path to a valid firmware image.

- `--dsp` *`file`*

  This option specifies an alternative source file for DM24 Mk3 DSP firmware. If not specified, the upgrade command will look in `/usr/share/firmware/CMG-DM24mk3/` for files with names of the form `dm24mk3-dsp*.img` and use the most recent version that it finds. It uses a lexical sort to determine which is the most recent. The argument *`file`* should be the path to a valid DM24 DSP firmware image.

- `--auto-baud`

  Scan the digitiser output to attempt to determine the Baud rate and adjust the configuration of the relevant serial port to match.

- `--verbose`

  Show most of the digitiser dialogue while the command is running.

- `--debug`

  This option turns on the production of some additional debugging information.

The `dm24-upgrade` command is the fastest and simplest way to upgrade the firmware on GSL-manufactured digitisers.

14.3.4 Rebooting

The “Reboot” item on the “Control” menu allows acquisition modules to be rebooted. CMG-NAMs can be both rebooted and powered off. To reboot from the command line prompt, use the reboot command.

```
eam999 ~ # reboot
```

14.3.5 Services

To control the service from the web interface, select:

Control → Services

This screen gives a list of all configured services: services are the background programs that read, convert and and write data and carry out the individual functions of the acquisition module.

The services are presented in three columns. In the first is given the name of the service and, in italics, its description. The second column shows the word “Stopped” in red for any services which are not running and, for those which are running, the PID (process ID, a unique number which the operating system uses to keep track of running programs) and the date and time that this instance of the service was started. The third column has buttons allowing you to stop, start or re-start each service.

It is possible to monitor and control services from the command line using the `ps` command and various scripts in `/etc/init.local` and `/etc/init.d`. This should be familiar to Linux users but full details are beyond the scope of this manual.

14.3.6 RAID Array Services

RAID arrays provide increased data security at the cost of extra storage devices. They can prevent the loss of data in the event of a single drive

failure. The "RAID array" item on the "Control" menu will only be displayed on CMG-NAMs with RAID fitted. It displays a page which reports the status of and allows simple control of the fitted RAID array. The status of swap partitions are also reported on this page.

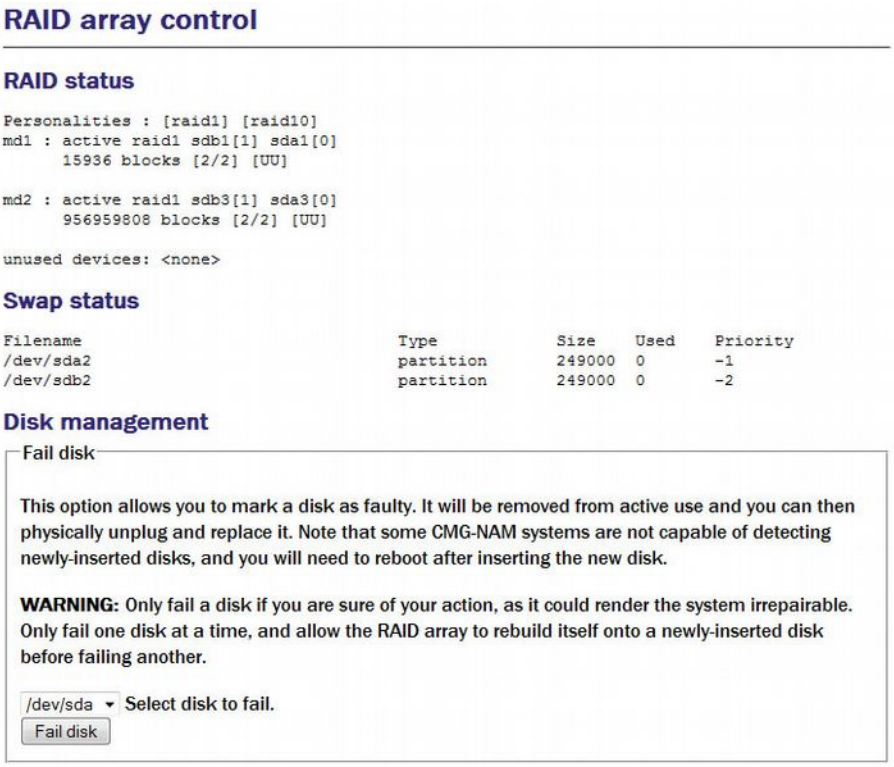

14.4 Tools Menu

14.4.1 CD1.1 log analyser

The CD1.1 log analyser is covered in a separate manual, the CD1.1 operations guide, MAN-EAM-1100, which is available for download from http://www.guralp.com/documents/MAN-EAM-1100.pdf

14.4.2 Environment logs

Platinum firmware has the facility to record environmental information such as temperature and supply voltage. Samples are recorded every ten seconds. The data can be subsequently downloaded or graphed. The parameters available for monitoring vary between models.

To configure the environment logs from the web interface, select:

Configuration → All options → GPIO labels and power switch settings → Logging

The screen displays the environment logging options:

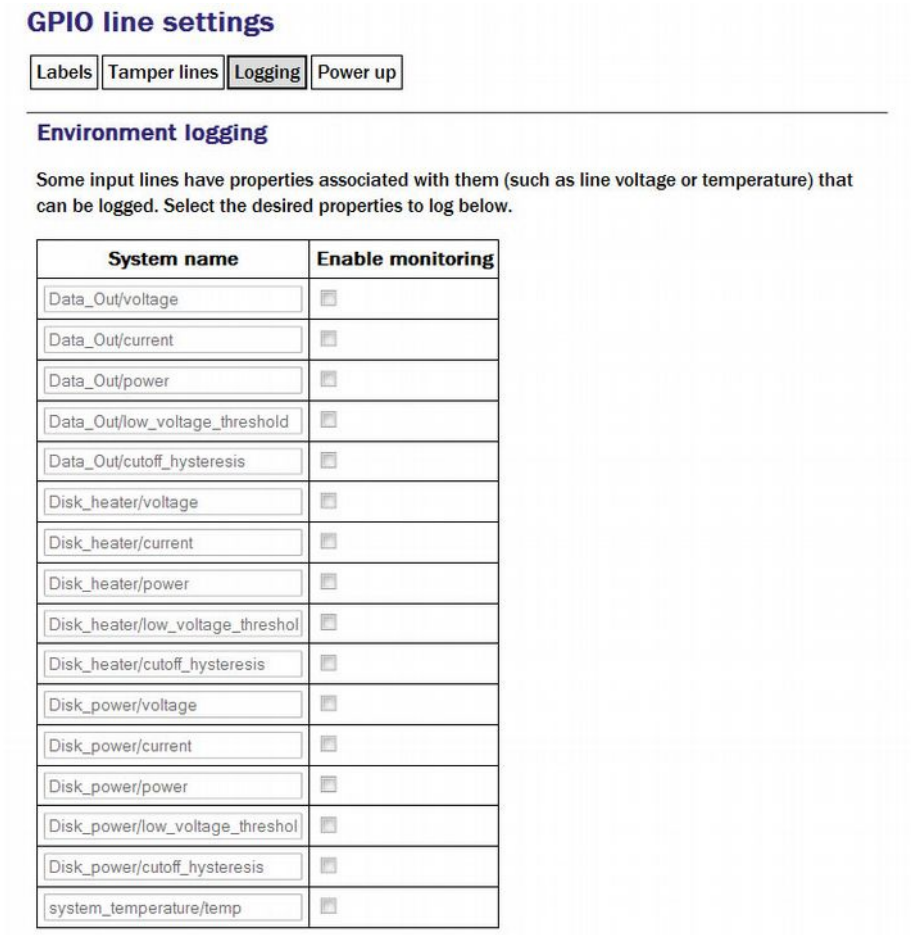

Tick the check-box next to each parameter you wish to monitor, then scroll to the bottom of the page and click Submit to record your choices.

14.4.3 Retrieving environment log data

Before data can be retrieved from the environment logs, the parameters that are to be logged must be selected, as described in the previous section.

To view or retrieve the logged data from the web interface, select:

Tools → Environment logs

The screen displays the query options:

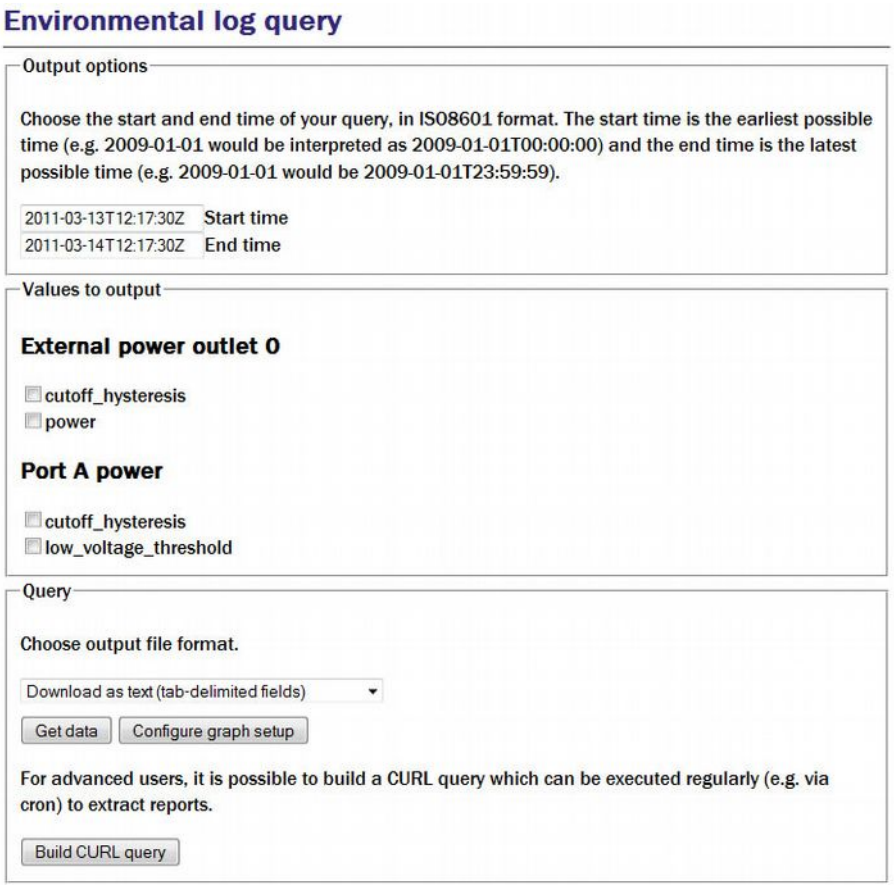

Output options: Enter the start and end times of the query. If you enter the same date without times (e.g. 2011-03-23) into both boxes, the query will select times covering 24 hours from midnight to midnight on the selected date.

Values to output: Select the value required by ticking the boxes. To change the available values, the environmental logging options will need updating as described in section 14.4.2 on page 235.

Query: Select the output file format from the drop down menu. The options are:

- Download as text (tab-delimited fields)
- Download as CSV (comma-delimited text file)
- Download as XML
- View as table (HTML)
- View as graph plot (PNG)
- Download PNG graph plot

Clicking on 'Get data' performs out the selected query. If the selected format is PNG, extra configurations options are available by clicking 'Configure graph setup'.

14.4.3.1 Build CURL query

It is possible to schedule environmental data extraction using cron, the system job-scheduler. When the job is scheduled, the data are extracted from the server using a set of cURL instructions.

cURL is an open source project for transferring data from a server using a command line. At a very basic level, all you need to do is type `curl` at the command line followed by the output to retrieve. For example:

```
curl -output test.html www.example.com
```

would take the contents of the URL www.example.com and output it to the filename `test.html`. There are a whole range of arguments you can insert between the command and the URL. These are all explained in the cURL manual, available at: http://curl.haxx.se/docs/manual.html

When you click on 'Build CURL query', Platinum generates a cURL query based on the options you have selected on the environmental log page. For example:

```
curl --output eam3057-2011-03-13T12:40:30Z.txt \
--form-string output=txt \
--form-string start=2011-03-13T12:40:30Z \
--form-string end=2011-03-14T12:40:30Z \
--form-string value='Data Out power/power' \
--form-string value='Data Out power/voltage' \
--form-string action='Get data' \
--digest --user root:rootme \
http://eam3057/cgi-bin/envirolog.cgi
```

In this example, the device serial number is eam3057. The data are output as text. The start and end times are 24 hours apart. There are two environmental values selected. The username and password are set to the system defaults.

Copying the generated code into ta erminal emulator connected to the EAM will output the file in the same way that the 'Get data' button does on the environmental log page.

If the cURL command is to be used for automated scheduled data extraction, some editing is required. Contact Güralp Systems for assistance on how to tailor the command for your specific requirements.

14.4.4 Extract MiniSEED records

The MiniSEED compressor, `gdi2miniseed`, uses a ring buffer to store miniSEED records after conversion from GDI. The default size of the buffer is

64 M but can be changed from the compressor's configuration screen, as described in section 12.2.1 on page 180. Records are never erased from the ring-buffer; the oldest records are over-written when the buffer becomes full. In a steady-state configuration, therefore, a fixed number of records, covering a fixed time period, are always present in the buffer.

If 512-byte records have been configured (as required for SEEDlink) a tool is available to extract some or all of these records and either download them (using the operator's browser) or write them to a file as a MiniSEED volume. Such files can then be manually copied from the system to removable storage, using the `cp` or `mv` commands, or to another system over the network, using `scp` or rsync, for example.

To extract MiniSEED records using the web interface select:

Tools → Extract MiniSEED records

The following screen is displayed:

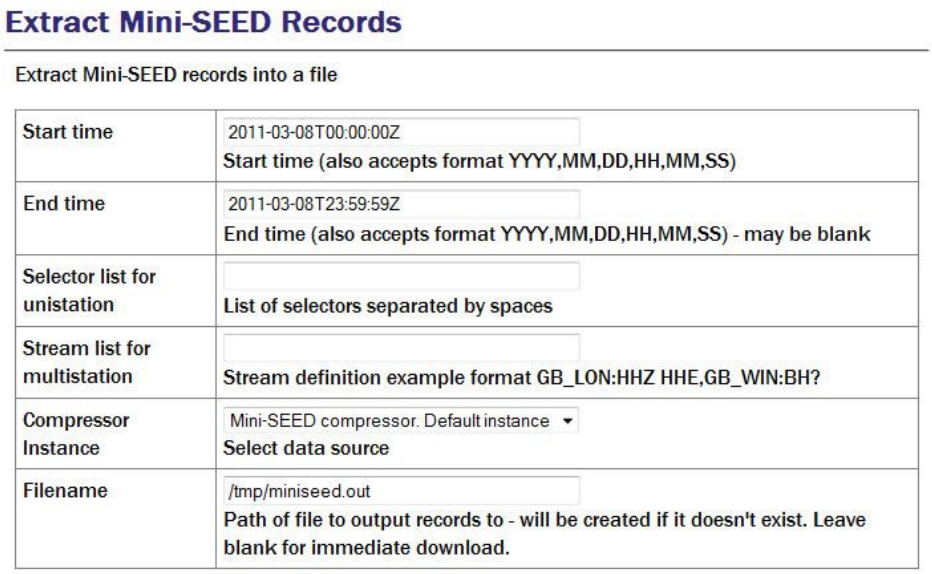

Start time and **End time:** Used to prevent the extraction of records whose timestamps fall outside the specified interval. The **End time** text-box may be left blank, in which case data up to the most recent are included in the extraction. Times may be specified in one of two formats: *YYYY*-*MM*-*DD*T*hh*:*mm*:*ss*Z or *YYYY*,*MM*,*DD*,*hh*,*mm*,*ss* (where *YYYY* represents the four-digit year, *MM* the two-digit month (January is month 1), *DD* is the two-digit day-of month (counting from 1) and *hh*, *mm* and *ss* are the two-digit hour, minute and second, respectively. The data and time are UTC and all numbers should be padded with leading zeroes to make up the required length.

The data can also be filtered to extract one or more specific streams. The selection mechanism supports both the ubiquitous multistation mode and the

less common unistation mode. Separate text fields are provided for specifying selections for each.

Selector list for unistation: Accepts a space-separated list of channel specifications. The channel name can be given in full or wild-cards and negations can be used: A question-mark character (?) is a single wild-card and will match any character in the same position; an exclamation-mark character (!) can be used to deselect a channel, so

`BHZ`	selects just channel **`BHZ`**
`BH?`	selects channels **`BHZ`**, **`BHN`** and **`BHE`**
`HH? !HHZ`	selects **`HHE`** and **`HHN`** but not **`HHZ`**
`!HHE`	selects all channels except HHE

Leaving this text field blank will select all channels.

Stream list for multistation: Can be left blank (to select all streams) or populated with a comma-separated list of stream specifications. Each specification is in the form ***`NET_STA`*** (where ***`NET`*** is the network identifier and ***`STA`*** is the station identifier, as used in S.C.N.L notation ***`)`***, optionally followed immediately by a colon (:) and then a space-separated list of channel selectors, as described in the previous paragraph.

For example, say a system is processing the following channels (in S.C.N.L notation):

```
B675.HHZ.KA.04
B675.CHZ.KA.0G
B675.SOH.KA.00
B675.HHE.KA.04
B675.HHN.KA.04
B675.BHE.KA.0M
B675.MME.KA.0A
B675.MMZ.KA.08
2200.BHN.GS.06
2200.HHE.GS.04
2200.HHN.GS.04
2200.BHE.GS.06
2200.BHZ.GS.06
2200.BHX.GS.06
2200.MMZ.GS.08
2200.MMN.GS.09
```

- Leaving both the unistation and multistation lists blank will select all data.
- A specification of `GS_2200` would select all data from station 2200.

- To select all data for HHE channels (from both stations) you could either use a unistation selector of `HHE` or a multistation selector of `GS_2200:HHE,KA_B675:HHE` - the second method would allow other streams to be selected in addition.
  - The specification `GS_2200:HHE BH?,KA_B675:HHE` will select data from the following channels:

    ```
    2200.HHE.GS.04
    2200.BHN.GS.06
    2200.BHE.GS.06
    2200.BHZ.GS.06
    2200.BHX.GS.06
    B675.HHE.KA.04
    ```

  - The specification `GS_2200:BH? !BHZ` will select data from the following channels:

    ```
    2200.BHN.GS.06
    2200.BHE.GS.06
    2200.BHX.GS.06
    ```

Database directory: The drop-down menu is used to select which `gdi2miniseed` compressor instance serves as the data source (each has its own ring-buffer).

Filename: If left blank, the extracted data are made available through the web-browser's file download facility and will typically be saved to the operator's PC's hard drive - details vary between browsers and the exact location can typically be configured by the user. If this field is populated with a file name, the data will be written to that file on the EAM. Any existing file with the same path name will be overwritten.

This functionality is also available from the command line with the `extract_miniseed` command. A usage message is displayed if the command is invoked with a `-h` argument.

14.4.5 GCF Audit Log Viewer

Detailed information about every GCF packet sent or received are stored on the GSL-EAM and can be viewed with the GCF Audit Log Viewer.

To access the GCF Audit Log Viewer from the web interface, select:

Tools → GCF audit log viewer

To access the same information from the command line, enter the command

```
gcflogview
```

The initial screen displays all GCF data sources and sinks in a table, together with some summary information. In the example given, it can be

seen that Ports B and F are inactive, Port A was receiving GCF data until 16:36 and the default instance of the Scream network server was sending GCF data until the same time.

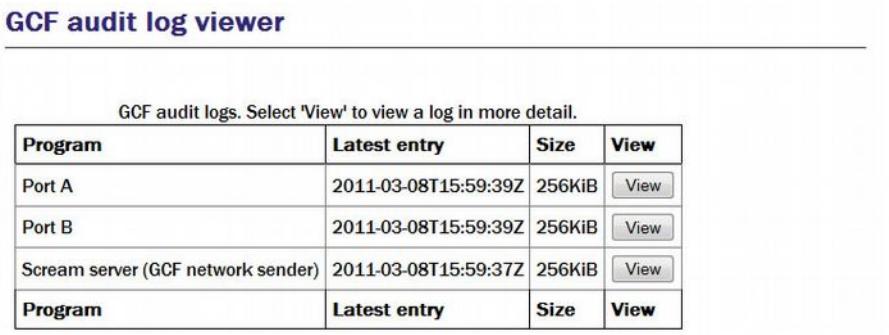

The "Size" column shows the size of the log buffer allocated to each data source or sink. The log buffer size can be changed from the relevant service or port configuration screens in expert mode.

For example, to allocate a larger log buffer to the GCF receiver running on Port A, click on "Serial ports" from the main menu, then on "Port A - GCF in", "GCF input settings" and then click the "Expert" button. You will see a drop-down selection list labelled "Audit log size" from which you can select 64Kib, 256Kib, 2MiB or 16MiB.

To change the GCF audit log buffer size for the Scream network server from the web interface, select:

Configuration → Data transfer/recording → Services

Now then click on "GCF Scream network server". Click on the entry for the instance you wish to change and then click the "Expert" button at the bottom of the page. You will see a drop-down selection list labelled "Audit log size" from which you can select 64Kib, 256Kib, 2MiB or 16MiB.

Each entry in the table has a "View" button, which shows detail from the relevant log at block (packet) level). The view for Port A is shown here:

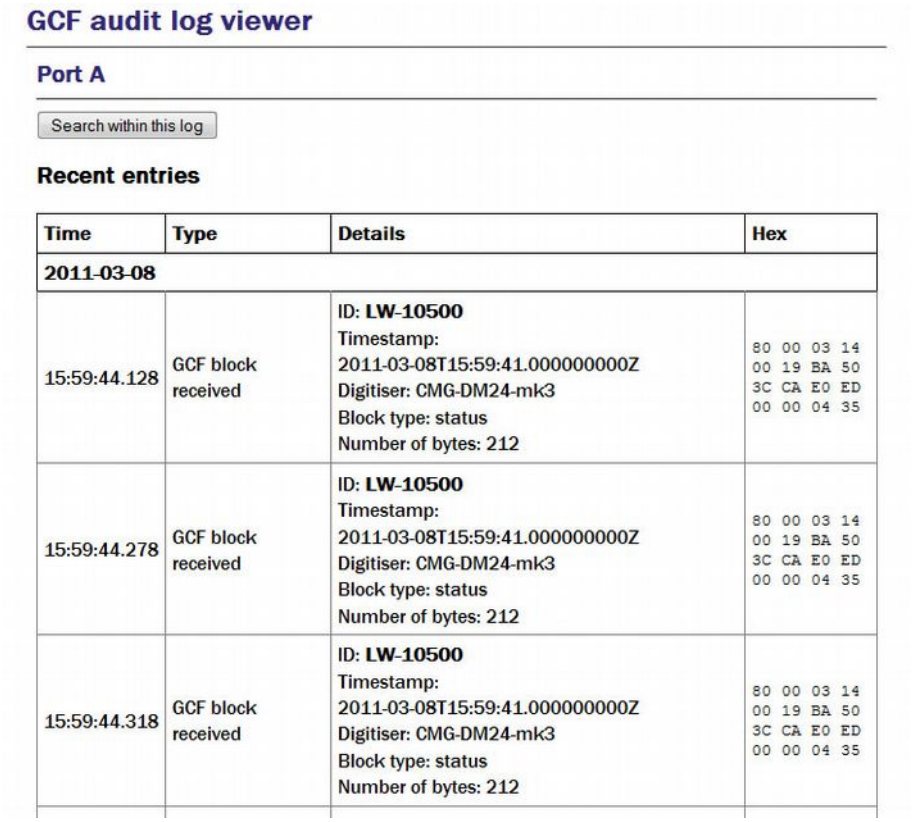

GCF audit log viewer

Port A

Search within this log

Recent entries

Time	Type	Details	Hex
2011-03-08			
15:59:44.128	GCF block received	ID: **LW-10500** Timestamp: 2011-03-08T15:59:41.000000000Z Digitiser: CMG-DM24-mk3 Block type: status Number of bytes: 212	80 00 03 14 00 19 BA 50 3C CA E0 ED 00 00 04 35
15:59:44.278	GCF block received	ID: **LW-10500** Timestamp: 2011-03-08T15:59:41.000000000Z Digitiser: CMG-DM24-mk3 Block type: status Number of bytes: 212	80 00 03 14 00 19 BA 50 3C CA E0 ED 00 00 04 35
15:59:44.318	GCF block received	ID: **LW-10500** Timestamp: 2011-03-08T15:59:41.000000000Z Digitiser: CMG-DM24-mk3 Block type: status Number of bytes: 212	80 00 03 14 00 19 BA 50 3C CA E0 ED 00 00 04 35

The first column shows the time-that the block was received (not the time-stamp on the block itself) and the second column shows the event type - “GCF block received” in most cases.

The “Details” column shows the stream ID from the received data block and the block time-stamp. The digitiser ID is shown if it is encoded in the block; otherwise, a best guess is displayed. The rest of the entry shows the block type, sample rate, compression level and the number of samples in the data block.

A hexadecimal display of the block header is shown in the final column.

14.4.6 GDI Channels Display

It is often useful, particularly when configuring a acquisition module for a complex array, to see a list of the Stream IDs, or channel names, which the CMG-EAM is receiving. The GDI Channels Display feature allows you to view a list of all active channels, together with some additional detail about each.

To access the GDI Channels Display from the web interface, select:

Tools → GDI Channels Display

Similar information is available from the command line via the command **`gdi-dump -l`** but the format is optimised for automated processing rather than human consumption. Giving the **`--help`** option provides usage details. Use of the web interface, however, is recommended.

The following summary screen is displayed:

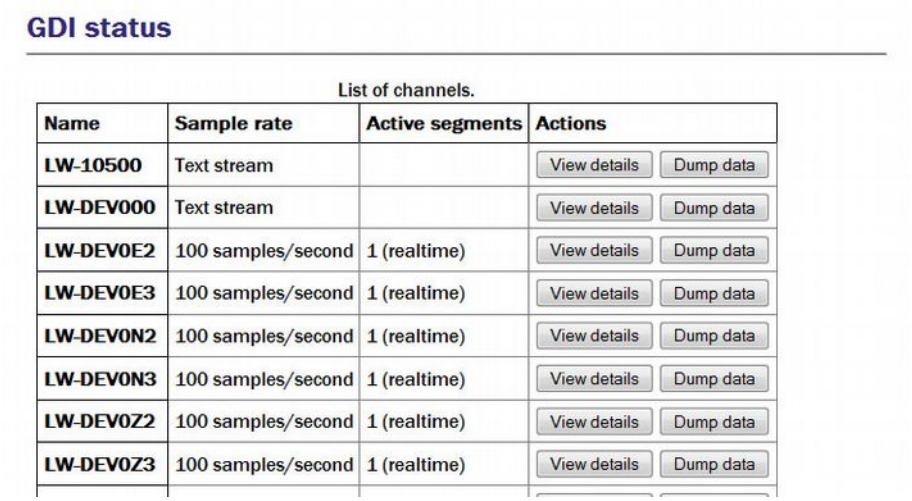

The first two columns show the names of the channels, together with information about the data format. The Active segments column shows details of data currently being received. A segment is a contiguous sequence of blocks so any data being back-filled always requires separate segments.

For each channel, you have the option of viewing detailed information about the data or the data itself, by use of the “View details” and “Dump data” buttons.

14.4.6.1 View details

The “View details” button displays the following screen:

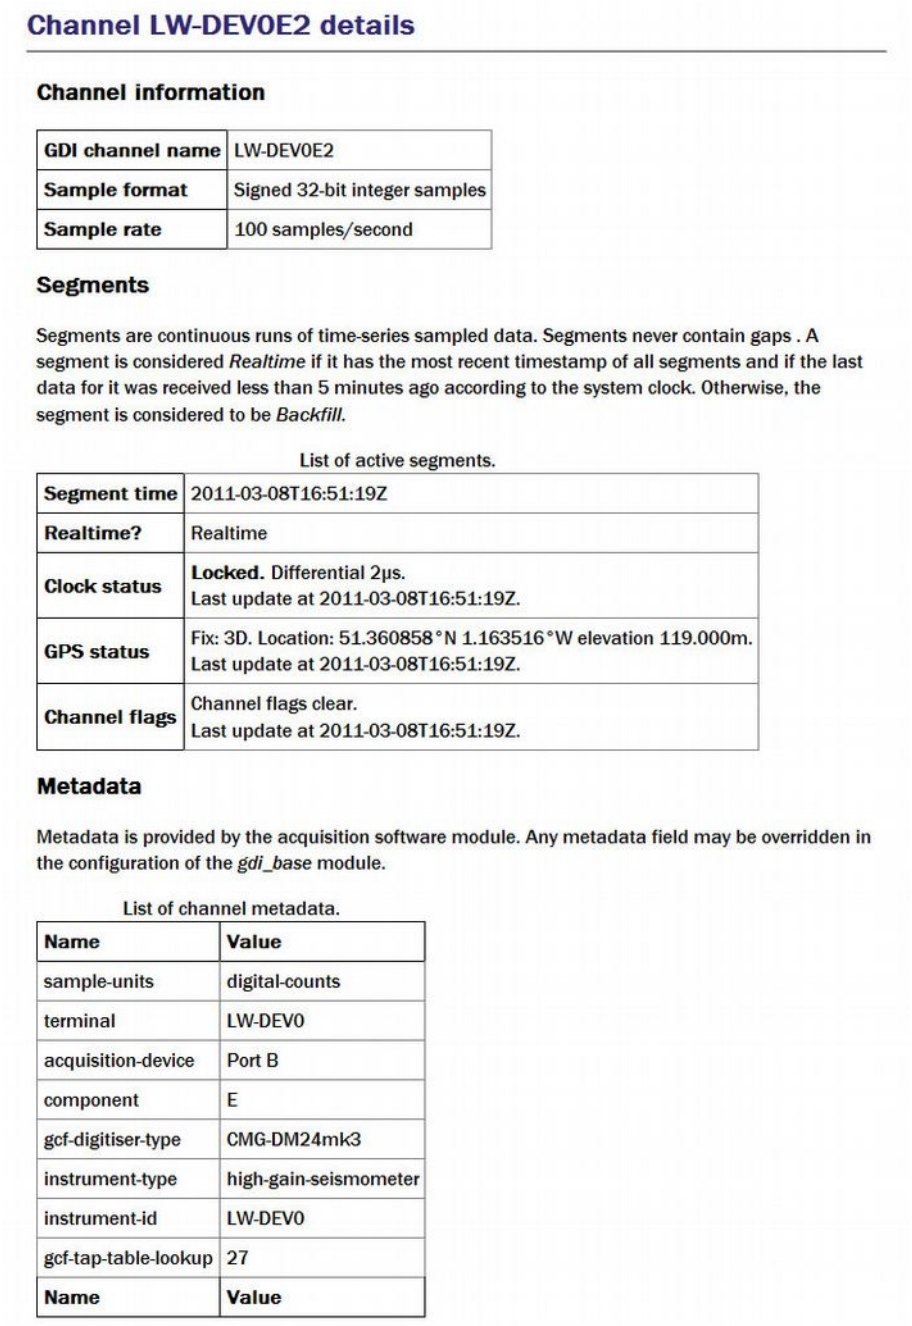

Channel LW-DEV0E2 details

Channel information

GDI channel name	LW-DEV0E2
Sample format	Signed 32-bit integer samples
Sample rate	100 samples/second

Segments

Segments are continuous runs of time-series sampled data. Segments never contain gaps . A segment is considered *Realtime* if it has the most recent timestamp of all segments and if the last data for it was received less than 5 minutes ago according to the system clock. Otherwise, the segment is considered to be *Backfill*.

List of active segments.

Segment time	2011-03-08T16:51:19Z
Realtime?	Realtime
Clock status	**Locked.** Differential 2µs. Last update at 2011-03-08T16:51:19Z.
GPS status	Fix: 3D. Location: 51.360858°N 1.163516°W elevation 119.000m. Last update at 2011-03-08T16:51:19Z.
Channel flags	Channel flags clear. Last update at 2011-03-08T16:51:19Z.

Metadata

Metadata is provided by the acquisition software module. Any metadata field may be overridden in the configuration of the *gdi_base* module.

List of channel metadata.

Name	**Value**
sample-units	digital-counts
terminal	LW-DEV0
acquisition-device	Port B
component	E
gcf-digitiser-type	CMG-DM24mk3
instrument-type	high-gain-seismometer
instrument-id	LW-DEV0
gcf-tap-table-lookup	27
Name	**Value**

The top section shows the channel name, sample format and sample rate, as seen on the previous screen.

The centre section shows detailed information decoded from the packet header for each active segment.

The final section shows the metadata associated with the stream derived from the configuration parameters of the relevant input module.

14.4.6.2 Data dump

Check-boxes are available to toggle the display of both metadata and sample data. These can be changed at any time and the “Restart dump” button used to refresh the display.

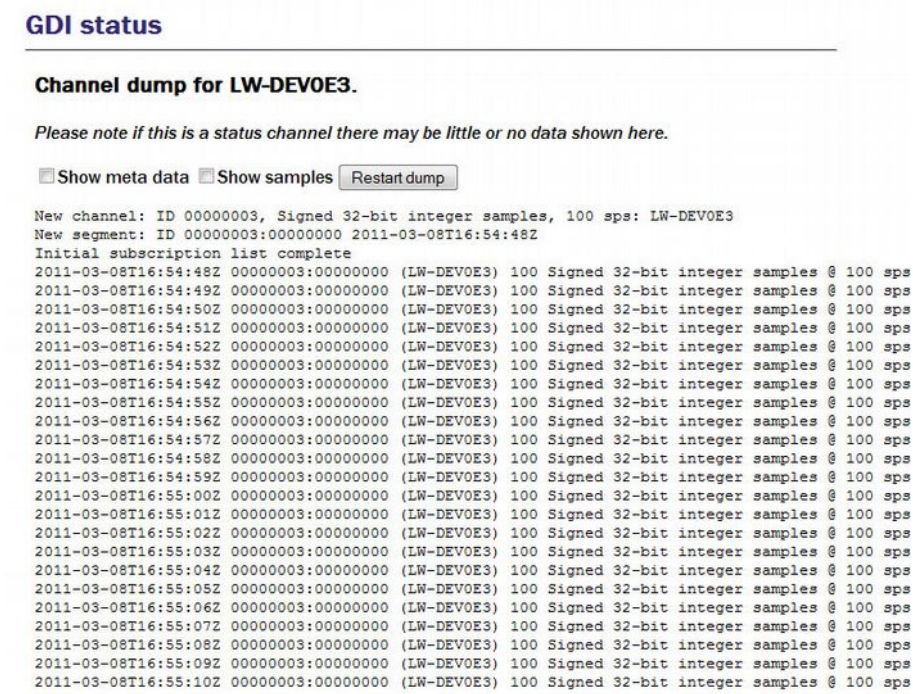

For sample data, each line displays the sample's time-stamp, the segment ID, the channel name in parentheses, the sample type and the actual sample value.

A button at the bottom of the screen allows the display to be refreshed with current data. There is also a button which, when clicked, returns the user to the main GDI Channels Display index page, so that another channel can be inspected.

14.4.7 Removable disk

Use of the mass storage device is described in chapter 11 on page 137.

14.5 Routine tasks

14.5.1 The directory cleaner

The directory cleaner is used to delete files which would otherwise accumulate and, eventually, fill the file system. Non-configurable, system-provided instances monitor three directories, `/var/lib/envirolog` (which holds environmental logging data), `/var/log/libterminal` (which records terminal interactions with digitisers) and `/var/lib/tamper` (which

records tamper-detection events). Additional instances can be created if required.

It is necessary to create a directory cleaner instance if you use the GDI link transmitter (`gdi-link-tx`). One directory cleaner instance is required for each `gdi-link-tx` instance. You may also use directory cleaner instances in association with your own scripts, to simplify programming.

To create a directory cleaner instance using the web interface, select:

Configuration → All Options → Routine tasks → Directory Cleaner

To configure an instance from the command line, start `gconfig` and select "Routine tasks" and then "Directory cleaner".

The screen shows a list of all directory cleaner instances that have been configured:

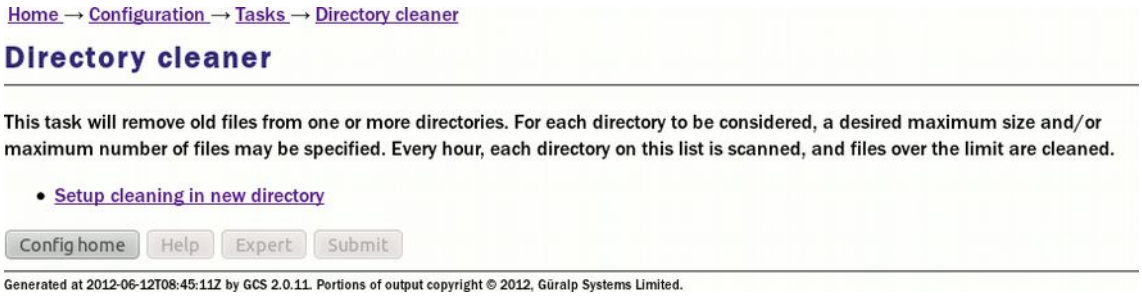

To edit an existing instance, click the appropriate link. To create a new instance, click "Setup cleaning in new directory". The resulting screen allows you to configure the parameters of the selected instance.

14.5.1.1 Configurable parameters

All parameters for a directory cleaner instance are displayed on a single screen:

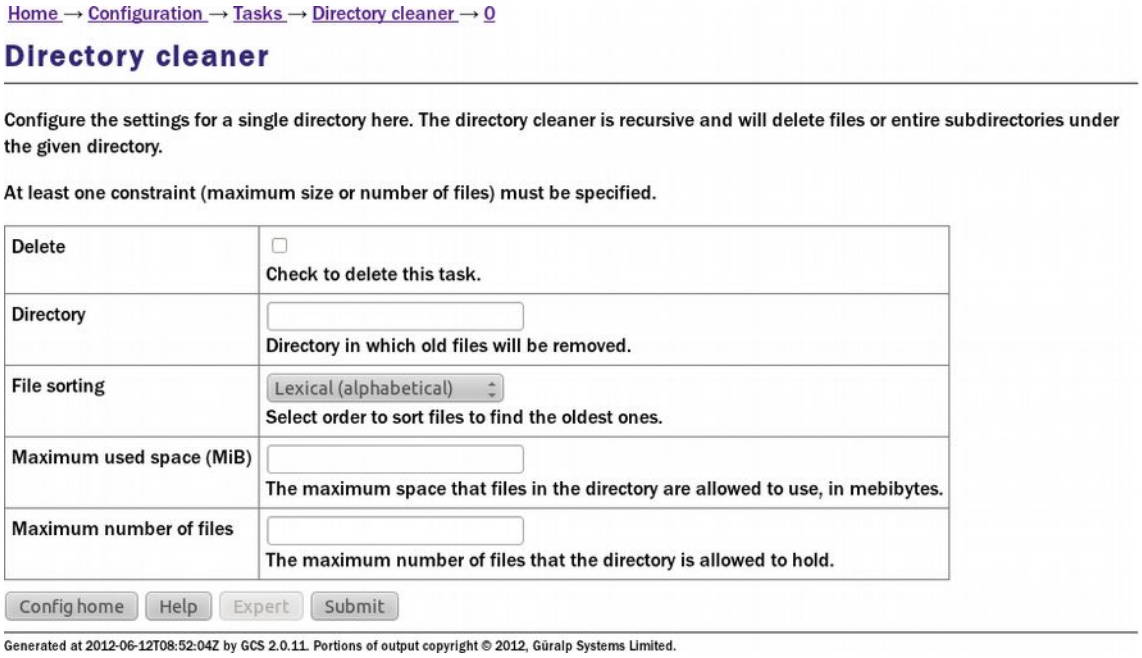

Delete: Causes this instance to be deleted from the system when the form is submitted.

Directory: Specifies the directory to be monitored and cleaned.

File sorting: Specifies the method used to determine which files should be deleted when the directory becomes too large. Several methods are available:

- **Lexical (alphabetical)**: this method sorts the files by their names in strict ASCII order. If the files are named after dates (in YYYYMMDD order) then this method will remove the oldest files first.
- **Reverse lexical**: this method also uses strict ASCII order. If the files are named after dates (in YYYYMMDD order) then this method will remove the newest files first.
- **By version number**: this method attempts to use a more natural sort order. It can cope with file-names like a0, a1, a2..... a9, a10, a11 and recognise the sequence. Files with lower numbers will be removed first. The exact algorithm comes from the standard GNU/Linux `strverscmp` function, for which documentation can easily be found on the web.
- **Reverse version number**: this method uses the same comparison function as the one above but deletes files in the reverse order.

- **File modification time**: this method ignores the file-names and deletes the files which were modified least recently.

Maximum used space (MiB): Specifies a threshold, in terms of storage space, that will trigger deletion in the directory. If the total size of the files in the directory is greater than the value specified here (in mebibytes), files are removed in the selected sort order until the threshold is no longer exceeded.

Maximum number of files: Specifies a threshold, in terms of number of files, that will trigger deletion in the directory. If the directory holds a larger number, excess files are deleted in the selected order until the specified number remain.

Note: At least one of **Maximum used space (MiB)** and **Maximum number of files** must be specified. If both are specified, they will both be respected.

Note: Each directory cleaner instance checks its configured directory once per hour so, if an application generates files rapidly, the total size or number of files in the directory may grow to considerably more than the configured threshold values.

15 Technical operation

15.1 Cylindrical Digitisers

Güralp Systems Ltd's cylindrical digitisers provide a CMG-DM24 and a CMG-EAM in a single package: a stainless steel or aluminium cylinder with an optional carrying/mounting bracket. An internal Spyrus PC card provides authentication and digital signing of CD1.1 frames and subframes.

The system is fitted with variable gain analogue inputs, tamper-detection lines on all key connectors and an internal USB storage device, which is available to external USB host devices such as laptop computers.

The connectors are significantly different from other packages. The analogue connectors are grouped at one end of the cylinder and correspond to similarly labelled connectors on a standard CMG-DM24:

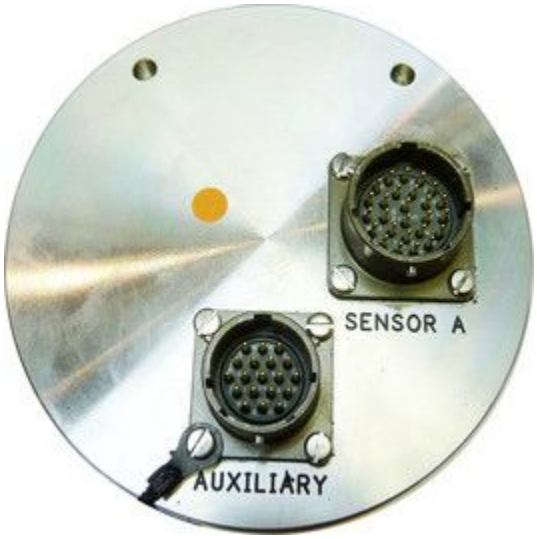

The illustration shows a four-channel unit; the seven-channel unit has an additional connector for sensor B:

The digital connectors are arranged at the other end of the cylinder:

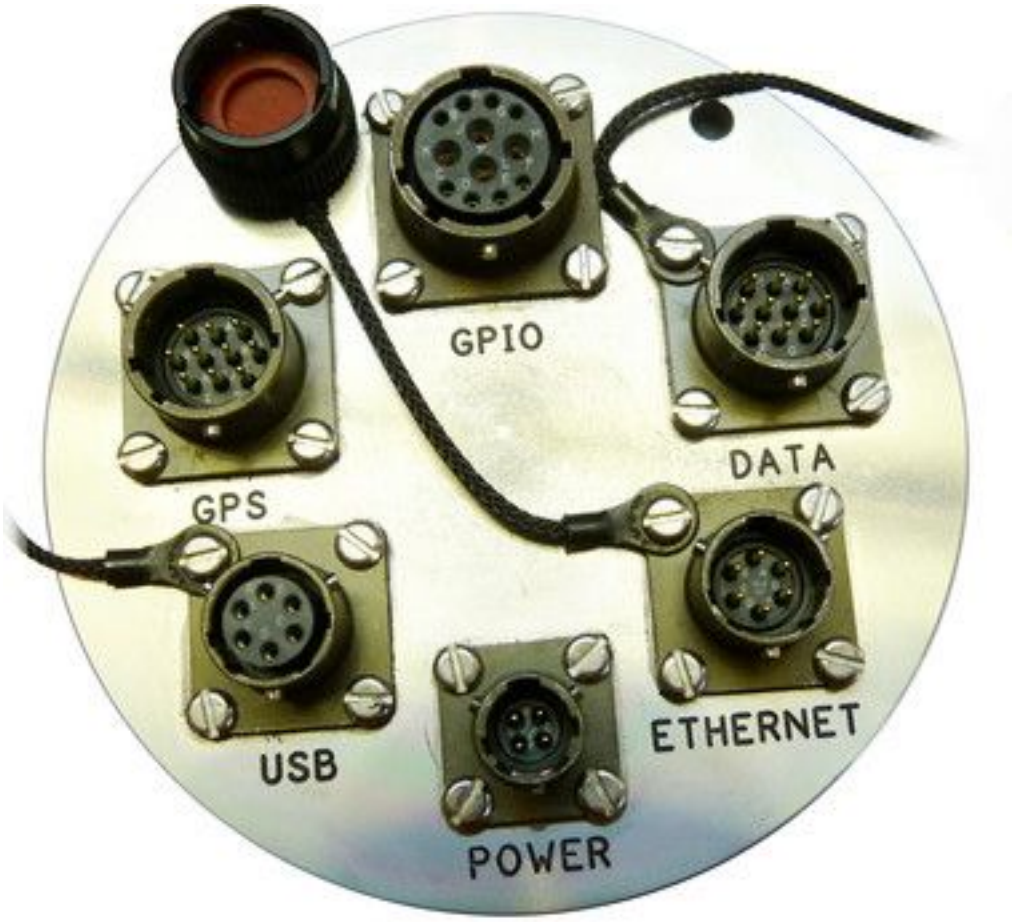

For more information on the ports see section 2.5on page 14. The connector pin-outs are given in section 14.1.1 on page 218 and listed in the table below:

Port	Section
Sensor	Section 16.5.12 on page 284
Auxiliary	Section 16.5.13 on page 285
Power	Section 16.5.9 on page 281
Data	Section 16.5.11 on page 283
Net	Section 16.5.4 on page 276
GPIO	Section 16.5.6 on page 278
USB	Section 16.5.8 on page 280
GPS	Section 16.5.10 on page 282

15.1.1 Internal Connections

Internally, the digitiser module and acquisition module are connected using three serial lines. The exact connections depend on the synchronisation mode, determined by the service running on Port C of the acquisition module.

In all cases, the GCF data output from the digitiser module is connected to Port A of the acquisition module, which should be set to “GCF in” at 38,400 Baud. It is currently possible to set the Baud rate of the digitiser's data output port and of the EAM's Port A independently, leading to a loss of data communication between the two modules if the two do not match. If you wish to have a higher transfer rate between the two modules, both ends must have their Baud rates increased separately.

The digitiser module also exposes a dedicated console connection, which is internally attached to Port B of the acquisition module. This can be accessed from the command line of the acquisition module or, if desired, made accessible over the network. If you wish to disallow network access, set the service on the EAM's Port B to “none”. To enable access over the network, set the service on the EAM's Port B to “TCP serial converter. Serial data to TCP link converter” and configure the converter according to the instructions in section 7.8 on page 97.

It is currently possible to set the Baud rate of the digitiser's console port and of the EAM's Port B independently, leading to a loss of data communication between the two modules if the two do not match. If you wish to have a higher transfer rate between the two modules, both ends must have their Baud rates increased separately. The digitiser module's Baud rate must be changed first.

Port C of the acquisition module is used for synchronisation and can either provide NMEA to the digitiser module or share incoming GPS data from an external receiver. Different internal connections are used in each case: an analogue switch is controlled by the service selected for the EAM's Port C.

When the Port C service is set to “NMEA In”, the connections are as follows:

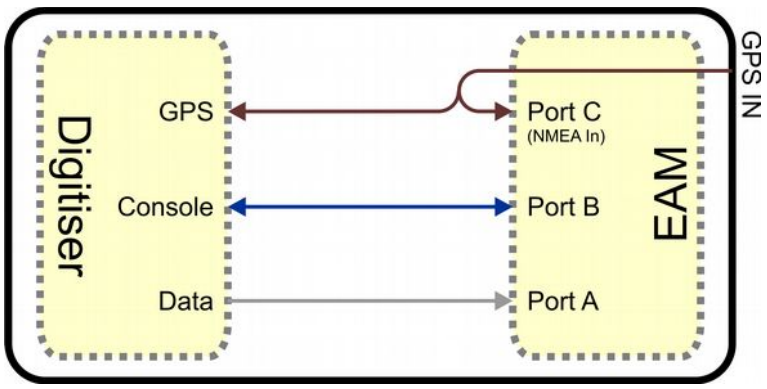

Note: this configuration is not recommended as a method of synchronising a acquisition module (see section 9 on page 120 for a detailed explanation).

Incoming data from an external GPS receiver is available to both the digitiser and acquisition modules. Both the digitiser's GPS input and Port C of the acquisition module must run at 4,800 Baud. This should never be changed.

If an external GPS receiver is available, it should be used to synchronise the digitiser: the EAM should then be synchronised the digitiser's RTSTATUS packets. In this case, the service on Port C can be set to “none”.

If use of a GPS receiver is impractical but internet-derived NTP synchronisation is available, this can be used as the clock source for both the digitiser module and the acquisition module. By setting the service on Port C to “NMEA out”, the following connections are enabled:

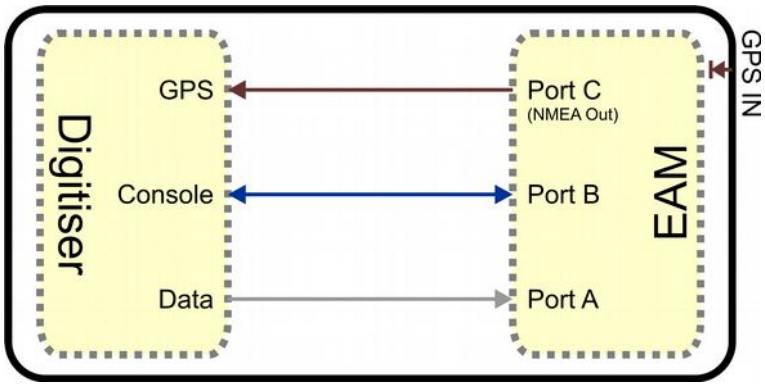

Both the digitiser's GPS input and Port C of the acquisition module must run at 4,800 Baud. This should never be changed. Note that, in this mode, the external GPS socket is disconnected: it cannot be used as an output for additional digitisers.

15.1.2 Variable Gain Inputs

The cylindrical digitiser is fitted with a programmable gain differential input amplifier which can be set to ×1, ×2, ×4, ×8, ×16, ×32 or ×64 gain operation.

The gain can be set individually for each input channel, either using the Platinum web configuration interface or directly from the digitiser's command line. In either case, the digitiser module must be re-booted before the new value will take effect.

The gain settings are reported in the status stream at boot time:

```
ADC #1 Version 760303
ADC o/s nulls 0 0 0 0
4 channel system
```

```
Gain Control : E8
Gain settings :  Ch#0 *1  Ch#1 *1  Ch#2 *1  Ch#3 *1
```

In the example above, all channels are set to unity gain (channel 0 is the vertical component and 1, 2 and 3 are the North/South, East/West and auxiliary/calibration channels, respectively. On a seven-channel digitiser, channels 4, 5 and 6 are the vertical, North/South and East/West components for the second instrument). If variable-gain-aware firmware (v106b42 and above) is loaded on a digitiser without variable-gain hardware, the text "No gain stage" will appear in this position in the boot status stream.

The selected gain setting is encoded into the GCF headers by appropriating bits from the System ID which must, therefore, be chosen to be five (or fewer) characters long. See the note at the end of this section for more information. The InfoBlock should be changed to reflect the amended System ID but the gain figure taken from the calibration document should be used unchanged, regardless of the variable gain setting chosen. Similarly the "calvals" file in Scream should not be changed, other than to reflect the System ID; Scream can deduce the variable gain settings in use from the GCF block headers and automatically take account of these during calibration operations.

To change the gain using the web interface select:

Configuration → System setup

Now select the digitiser from the list. Once connected, scroll down to the Connected devices. The following section appears:

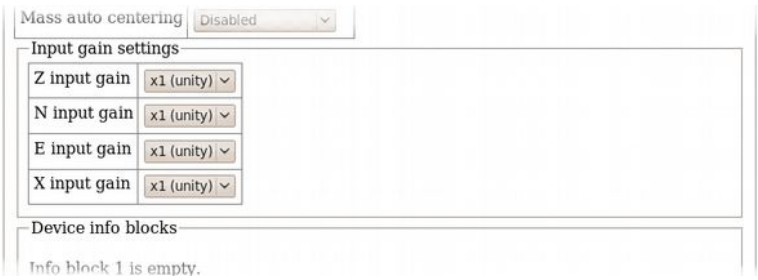

(This table is extended to show three additional components when a seven-channel digitiser is detected.)

From here, the gain can be set individually for each component. If the Submit button is clicked, the changes will be stored in the digitiser module's configuration but will not take effect until the module is rebooted. If an immediate change is required, the Submit changes & Reboot digitiser button should be used instead.

To change the gain using the command line, use the `data-terminal` command to connect to the digitiser, as described in section 8.2 on page 114 and issue one of the following two commands.

To simultaneously set all channels to the same gain, enter the command:

***gain* *gains**

where ***gain*** is one of `1`, `2`, `4`, `8`, `16`, `32` or `64`. For example, to select ×8 gain on all channels, enter the command

***8* *gains**

The digitiser must be rebooted before the change will take effect.

To set the gain for an individual channel, enter the command:

***channel gain* *gain**

where ***channel*** is one of 0 (vertical), 1 (North/South), 2 (East/West) or 3 (auxiliary/calibration). On seven channel digitisers, this parameter can also be one of 4 (vertical), 5 (North/South) or 6 (East/West), referring to the components from the second instrument. ***gain*** is one of `1`, `2`, `4`, `8`, `16`, `32` or `64`. For example, to select ×16 gain on just the vertical channel, enter the command

***0 16* *gains**

The digitiser must be rebooted before the change will take effect.

Software developers working with GCF packets can decode the selected gain setting from the GCF header as follows:

If the most significant bit of the System ID is zero, variable gain is not used. If the most significant two bits of the System ID are 10 or 11, the next three bits encode the gain, using this code:

Bits 2, 3 & 4	Gain
000	not fitted
001	×1
010	×2
011	×4
100	×8
101	×16
110	×32
111	×64

15.1.3 USB operations

The Cylindrical Digitiser can behave as a USB storage device (via the GPIO connector) or as a USB host (via the USB connector).

15.1.3.1 USB device mode

The Cylindrical Digitiser is fitted with an internal Flash memory device which is accessible via USB. It can be written to by selecting “Internal USB storage” from the “Recording destination” drop-down menu on the “Disk recording” page (see section 11.2 on page 138).

When a USB host, such as a laptop or PC, is connected to the GPIO port (who's pin-out is given in section 16.5.6 on page 278) internal circuitry detects the USB power and automatically connects the Flash memory to the GPIO socket, causing it to behave identically to a standard USB memory stick.

When no power is detected at the GPIO port, the Flash memory is available to the system as if it were a standard removable mass storage device. All of the mass storage device recording options described in section 11.2 (on page 138) will apply to this device.

15.1.3.2 USB host mode

If a USB storage device is connected to the USB port (see section 16.5.8 on page 280 for the pin-out), it will be mounted under `/media`. It can be used to store seismic data by selecting “External USB drive on mil-spec connector” from the “Recording destination” drop-down menu on the “Disk recording” menu (see section 11.2 on page 138).

15.2 DCM

The DCM is a fully-functional, Linux-based computer system especially designed for handling seismic data. It can collect and store data from several sources and, if required, output it in your preferred format to other locations on your network or on the Internet. The CMG-DCM receives data from one or more of the following sources:

- Digitizer connected through a serial link
- Computer running a Scream server
- CD1.0 or CD1.1 transmitter (optionally)
- Another DCM or AM

All the received data are stored in files in the on-board Flash memory. There are two banks of Flash memory available, which are accessible as `/nand0` and `/nand1` in the Linux file tree. Data are normally stored as GCF (Güralp Compressed Format) files.

As an option, you may be able to configure the DCM to use the *miniSEED* or *sac* formats instead.

In *automatic* mode, when the Flash memory becomes more than 75% full, the oldest data files are moved to the DCM's primary hard disk until it is less than

50% full. If you prefer, you can configure the DCM to write to the hard disk at set intervals.

Writing to the hard disk is performed robustly, so that no data will be lost if a write is aborted. This means that you can safely swap hardware in and out at any time. Stand-alone DCM modules use off-the-shelf Lacie hard disks, which can be easily removed and installed in most conditions. You can specify other models of IDE / USB or IEEE 1394 2.5" disk at manufacture. If an internal disk is not present, and the module has a USB *host* interface, it will look for hard disks connected to its external USB port.

Once the data are stored on the DCM, whether in Flash memory or on the hard disk, it can be retrieved by

- A remote computer running Güralp Systems' Scream, or other GCF-compatible software;
- Another DCM or AM, also using GCF;
- Setting up a CD1.0 or CD1.1 transmitter on the DCM;
- Direct file transfer (using SSH, HTTP, HTTPS, etc.,)
- Optionally requesting the data using SeedLink or AutoDRM.

A PC running Güralp Systems' Scream software can not only collect data from the DCM, but also configure the module and any instruments attached to it.

You may need to enable and configure some of these methods before you can use them.

Most installations of the DCM will not require any more complex setting up than the Web configuration system can offer. However, in some cases you may need to take advantage of the flexibility offered by the underlying Linux operating system.

For more information on using and configuring the CMG-DCM see the manual MAN-DCM-0001.

15.3 24 Channel DAS

The CMG-DM24S24EAM comprises an EAM, interface board and 4 DM24 digitisers.

Originally built to process geophone signals, the DAS uses a system of breakout boxes to transmit the analogue signals to the module. The Pelicase and breakout boxes are shown in the images below

Each of the four breakout boxes has terminal connectors for up to 6 channels and connects to the DAS using 32 pin mil spec connectors.

Other ports on the DAS are: Data out, Ethernet, USB and GPS and triggering.

For more information on the ports see section 2.5 on page 14. The connector pinouts are the same as those for the Cylindrical Digitiser and are are given in 16.5 on page 273 and listed in the table below:

Port	Section
Sensor	Section 16.5.14 on page 286
Data	Section 16.5.2 on page 274
Ethernet	Section 16.5.4 on page 276
USB	Section 16.5.3 on page 275
GPS	Section 16.5.7 on page 279

The system architecture is detailed in the image below:

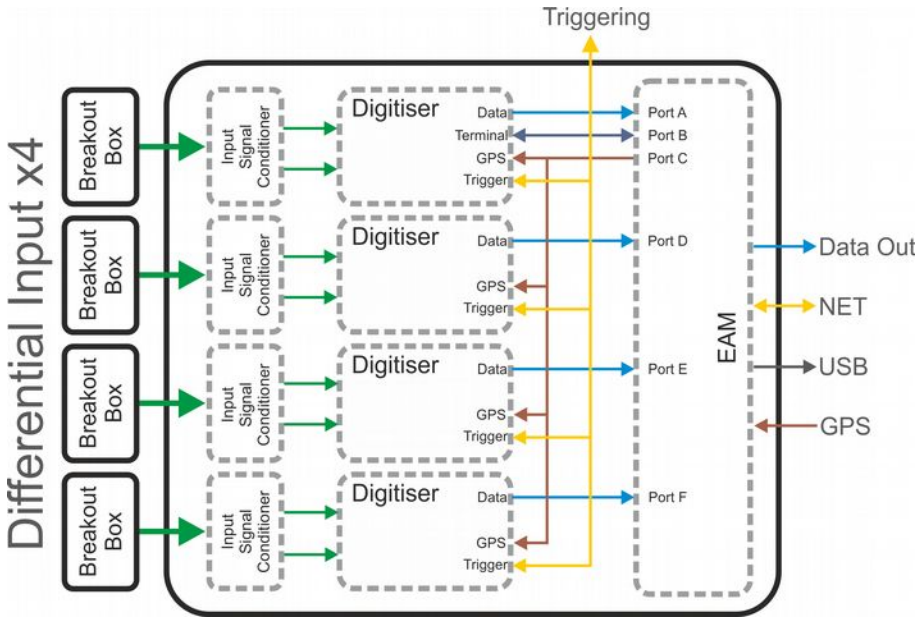

15.4 Instruments with integrated CMG-EAMs

Güralp Systems digital instruments, such as the CMG-5TDE (shown below), have an integrated digitiser and EAM.

The internal connections between the CMG-DM24 modules and the CMG-EAM module are identical to those in an Cylindrical Digitiser, as described in section 15.1 on page 250.

For more information on the ports see section 2.5 on page 14. The connector pinouts are the same as those for the Cylindrical Digitiser and are are given in 16.5 on page 273 and listed in the table below:

Port	Section
Data	Section 16.5.11 on page 283
Net	Section 16.5.4 on page 276
GPIO	Section 16.5.6 on page 278
USB	Section 16.5.8 on page 280
GPS	Section 16.5.10 on page 282

16 Appendices

16.1 Appendix A - Setting the System Identity (Hostname)

The system identity is pre-set at the factory to contain a device-type identifier and the system's serial number, such as “EAM2010”. It should not be necessary to change this but, should you desire to, the following procedure can be used.

To set the system identity using the web interface select:

Configuration → Hostname

or

Configuration → All options → System identity (hostname)

To set the system identity from the command line, start `gconfig` and select “System identity (hostname)” from the top level menu.

The following screen is displayed:

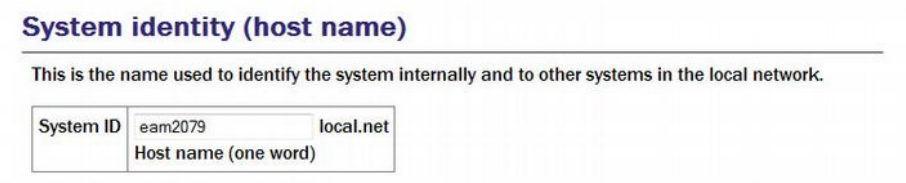

Use this screen to set the hostname of the system.

16.2 Appendix B - Using third-party terminal emulators

There are a number of terminal emulator programs that you can use to access the serial ports of the digitiser and the optional networking interface. The terminal emulator built into Scream is recommended but, if this is not available, there are a variety of alternatives. Three of these are detailed below.

16.2.1 Hyperterminal, as provided with Windows XP.

1. Click on **Start** and then **Run**.
2. Enter 'hypertrm' and click on OK.

3. The program will ask you for a name for the connection:

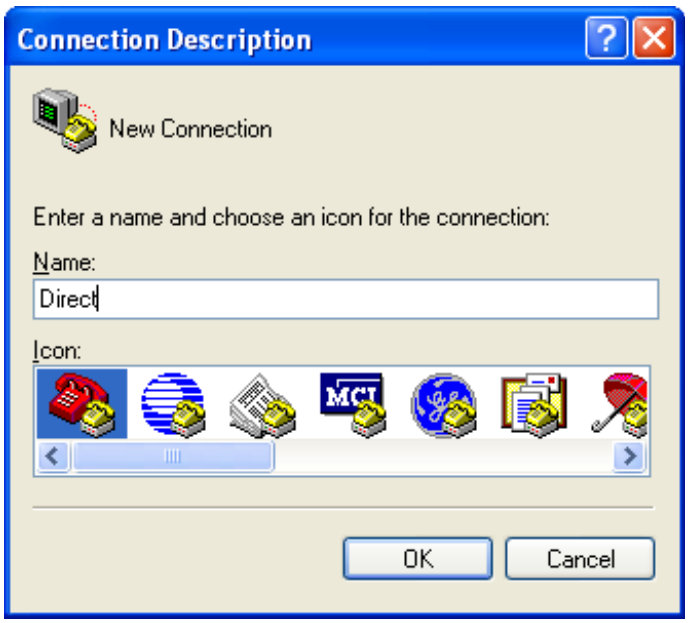

Enter any suitable name, then click OK.

4. You will then be prompted to enter COM port and modem details:

The **Country/region**, **Area code** and **Phone number** fields can be ignored: they are only used when working with modem connections. Select the name of the correct COM port from the **Connect using** drop-down menu, then click OK.

5. Enter the port configuration settings:

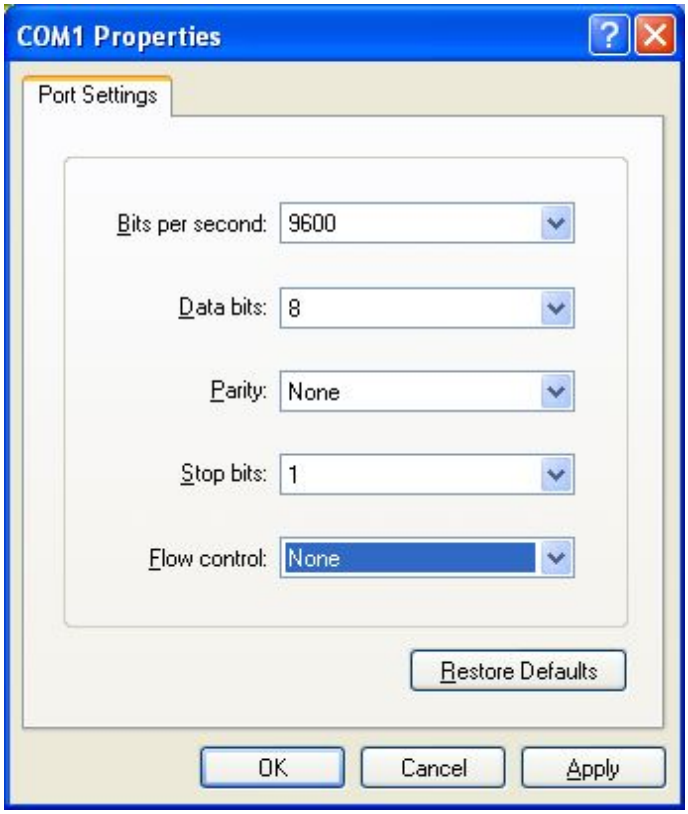

Ensure the following parameters are set:

Bits per second: 38400 for DM24, 115200 for DCM, 19200 for CD24
Data bits: 8
Parity: None
Stop bits: 1
Flow control: None

Caution: If the port is re-configured or the settings changed, access may no longer be possible.

6. Click on OK and the program will then connect, providing you with a terminal emulator screen, from which you can access the command line of your system.

16.2.2 Using Hyperterminal with Windows Vista or Windows 7.

HyperTerminal is not provided with the Windows Vista or Windows 7 operating systems but the necessary files can copied from the i386 directory of the Windows XP CD, if you have one available. The two files you will need are:

hypertrm.dll
hypertrm.exe.

To use Hyperterminal with Windows Vista or Windows 7:

1. Copy the two files into your windows/system32 directory.

2. To access HyperTerminal, use Windows+R on your keyboard. Enter 'hypertrm' and click on **OK**.

3. If Windows open a security warning window click on **Run**. You may also be asked if you wish to use HyperTerminal as a the default terminal window.

4. Now follow the instructions given in section 16.2.1, above.

16.2.3 Using PuTTY for Windows

PuTTY is a free terminal package for windows which is useful if HyperTerminal is not available. It can be downloaded from www.chiark.greenend.org.uk. The easiest package to use is the 'windows installer'. Install PuTTY by following the on-screen instructions.

1. Start PuTTY by clicking on the desktop icon or Start-menu entry

2. Click on **Serial** at the bottom of the category menu on the left:

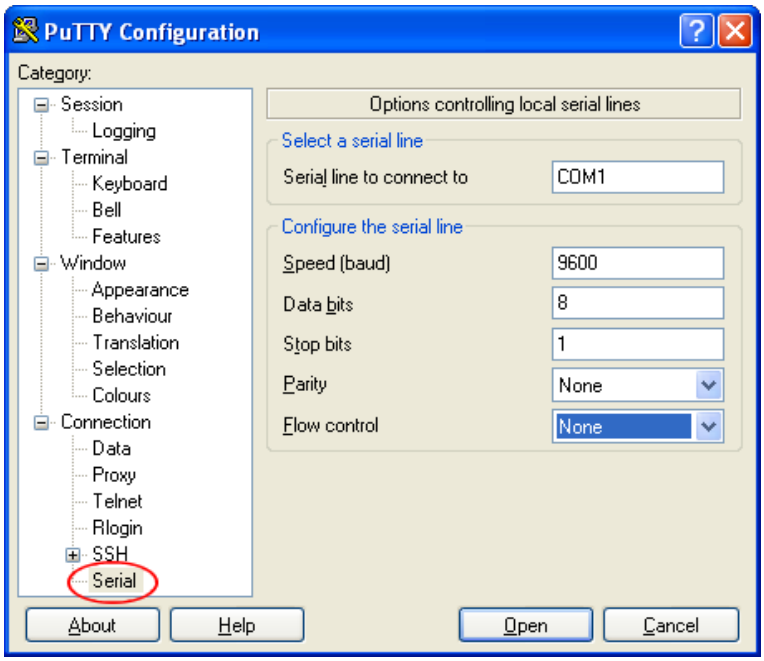

3. Select a serial line to connect to. This is usually COM1.

4. Configure the serial line with the following settings:

Speed (baud): 38400 for DM24, 115200 for DCM, 19200 for CD24

Data bits: 8

Stop bits: 1

Parity: None

Flow control: None

> **Caution:** If the port is re-configured or the settings changed, access may no longer be possible.

5. Click on **Session** in the main menu.

6. In the right hand part of the window use the radio buttons to select the Serial Connection type. Check the **Serial line** and **Speed** fields are correct.

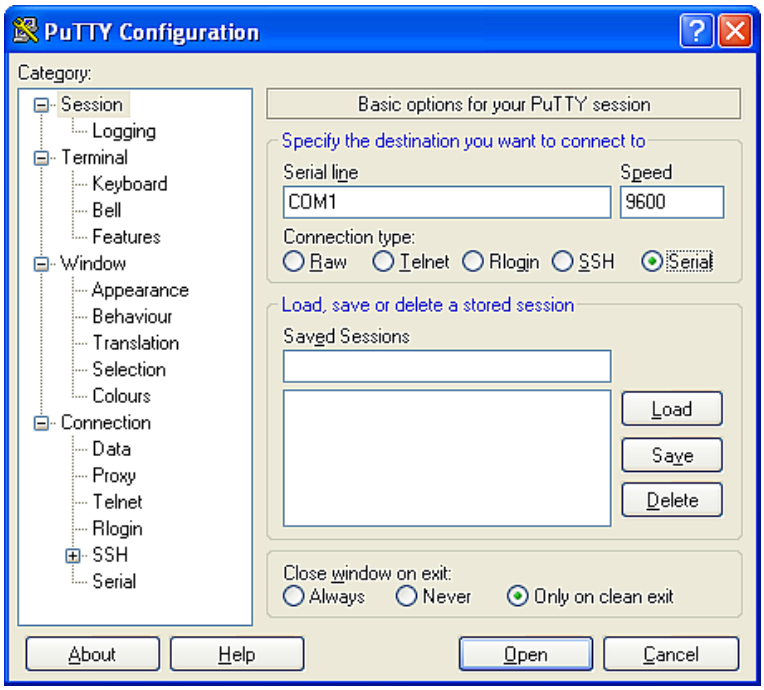

7. To save the settings enter a suitable session name in the **Saved Sessions** field then click **Save**.

The next time you start PuTTY, your saved session will appear in the list and you can simply double-click it to open a new session with the same settings.

8. Click the **Open** button to start the terminal emulator.

16.2.4 Mincom for Linux

Under Unix or Linux, Miquel van Smoorenburg's *minicom* terminal emulator (more details from http://alioth.debian.org/projects/minicom) is recommended, although most terminal emulators can be used. An extract from Minicom's user manual is reproduced in section 16.3 on page 267.

16.3 Appendix C - Using Minicom

The acquisition module includes the Linux program minicom as a terminal emulator for use with serial devices, including Güralp digitisers. The following is part of the minicom man page.

Minicom is window based. To pop up a window with the function you want, hold the Ctrl key while typing A (from now on, we will use Ctrl + A to denote this), and then the function key (a-z or A-Z). By pressing Ctrl + A first and then Z, a help screen comes up with a short summary of all commands.

For every menu the following keys can be used:

Key	Action
↑ or K	UP
↓ or J	DOWN
← or H	LEFT
→ or L	RIGHT
↵	CHOOSE
Esc	CANCEL

The screen is divided into two portions: the upper 24 lines are the terminal-emulator screen. In this window, ANSI or VT100 escape sequences are interpreted. If there is a line left at the bottom, a status line is placed there. If this is not possible the status line will be showed every time you press Ctrl + A. On terminals that have a special status line, it will be used if the termcap information is complete and the -k flag has been given. Possible commands are listed next, in alphabetical order.

Key	Action
Ctrl + A	Pressing Ctrl + A a second time will just send a Ctrl + A to the remote system. If you have changed your "escape character" to something other than Ctrl + A, this works analogously for that character.
A	Toggle 'Add Linefeed' on/off. If it is on, a linefeed is

	added before every carriage return displayed on the screen.
B	Gives you a scroll back buffer. You can scroll up with U, down with D, a page up with B, a page down with F and, if you have them, the ← → ↑ ↓ and Page Up Page Down keys can also be used. You can search for text in the buffer with S (case-sensitive) or ⇧ + S (case-insensitive). N will find the next occurrence of the string. C will enter citation mode. A text cursor appears and you specify the start line by hitting ↵ key. Then scroll back mode will finish and the contents with prefix '>' will be sent.
C	Clears the screen.
E	Toggle local echo on and off.
F	A break signal is sent.
I	Toggle the type of escape sequence that the cursor keys send between normal and applications mode. (See also the comment about the status line below).
J	Jump to a shell. On return, the whole screen will be redrawn.
K	Clears the screen, runs kermit and redraws the screen upon return.
L	Turn Capture file on off. If turned on, all output sent to the screen will be captured in the file too.
O	Configure minicom. Puts you in the configuration menu.
P	Communication Parameters. Allows you to change the bps rate, parity and number of bits.
Q	Exit minicom without resetting the modem. If macros changed and were not saved, you will have a chance to do so.
R	Receive files. Choose from various protocols (external). If you have the filename selection window and the prompt for download directory enabled, you'll get a selection window for choosing the directory for downloading. Otherwise the download directory defined in the Filenames and paths menu will be used.

S	Send files. Choose the protocol like you do with the receive command. If you don't have the file-name selection window enabled (in the File transfer protocols menu), you'll just have to write the file-name(s) in a dialog window. If you have the selection window enabled, a window will pop up showing the file-names in your upload directory. You can tag and un-tag file-names by pressing space, and move the cursor up and down with ↑/↓ or J/K. The selected file-names are shown highlighted. Directory names are shown [within brackets] and you can move up or down in the directory tree by pressing space twice. Finally, send the files by pressing ↵ or quit by pressing Esc.
T	Choose Terminal emulation: Ansi(color) or vt100. You can also change the backspace key here, turn the status line on or off, and define delay (in milliseconds) after each newline if you need that.
W	Toggle line-wrap on/off.
X	Exit minicom, reset modem. If macros changed and were not saved, you will have a chance to do so.
Y	Paste a file. Reads a file and sends its contents just as if it would be typed in.
Z	Pop up the help screen.

16.4 Appendix D - Troubleshooting

This section of the manual comprises answers to the questions most frequently asked of our support team.

16.4.1 Upgrades report “Temporary failure in name resolution”

If an attempt to run the upgrade process produces error messages like

```
rsync: getaddrinfo: rsync.guralp.com 873: Temporary failure
in name resolution
rsync error: error in socket IO (code 10) at
clientserver.c(122) [receiver=3.0.2]
```

the Domain Name Service (DNS) client is misconfigured. If you are running DHCP, this may be a problem with your DHCP server not providing a nameserver (or providing an incorrect one). If you are using static addressing, check and correct the

16.4.2 Upgrades report “Network is unreachable”

If an attempt to run the upgrade process produces error messages like

```
rsync: failed to connect to rsync.guralp.com: Network is
unreachable (101)
rsync error: error in socket IO (code 10) at
clientserver.c(122) [receiver=3.0.2]
```

the network routing is misconfigured. If you are running DHCP, this may be a problem with your DHCP server not providing a default route (or providing an incorrect one). If you are using static addressing, check and correct the **Default route (gateway)** field in the “Network interface” configuration page, as described in section 7.1.1 on page 71.

16.4.3 Upgrades report “rsync error”

If an attempt to run the upgrade process (using the GSL rsync server over the internet) produces error messages like

```
rsync error: received SIGINT, SIGTERM or SIGHUP (code 20) at
rsync.c(541)
```

the likelihood is that a firewall is blocking traffic on the rsync port. Ask your network administrator to permit the EAM to open TCP connections to host `rsync.guralp.com` on port 873.

16.4.4 Errors during upgrade: “directory not empty”

This message can occur under two separate sets of circumstances.

On CMG-DCMs running firmware earlier than Build 3696, this message could appear in relation to files with the `.so` extension and/or files in the

`/usr/share/terminfo` directory. Such units should have their firmware upgraded twice to the latest version. The first upgrade installs a new upgrade script and the second upgrade runs the new script, which fixes this problem.

In all other cases, this message signifies nothing of concern and can be safely ignored.

16.4.5 Upgrade completes but build version remains at 3801

Build 3801 was the last Platinum build that used the OABI software architecture. Subsequent versions use EABI. The two software architectures are not compatible so the standard upgrade method cannot be used to move from build 3801 (or earlier) to the current build. Performing a standard upgrade on a system running build 3801 will check and repair any corrupt files but will not install the latest firmware.

An extra step is required to install the latest firmware. This is fully described in section 5.1 on page 51. Once you have performed this step, subsequent upgrades will behave as expected.

16.4.6 Regaining access when “locked out”

It is, on some units, possible to configure the network and serial ports in such a way as to make reconfiguration apparently impossible. This section provides directions for regaining control of such units.

On cylindrical digitisers and instruments with integrated DAS, such as the CMG-5TDE, the GPIO connector provides a dedicated serial console (running at 38,400 Baud) which cannot be reconfigured. The GPIO connector on these units can, therefore, always be used to achieve command-line access.

On peli-cased acquisition modules, the 9-pin 'D'-connector under the lid is intended as a dedicated console connector and is preconfigured to run at 38,400 Baud. It is recommended that the port is not reconfigured for other purposes but there is nothing in the firmware to prevent this, so it is still possible to inadvertently mis-configure the system.

There are two common reasons why one might not be able to use a serial port: incorrect port function and incorrect baud rate. This section may help in recovering from these situations.

Connect to the EAM via a serial port, open a terminal emulator (not Scream) and observe the traffic appearing in the window.

Carry out the appropriate procedures depending on the traffic:

16.4.6.1 Short repeated bursts of traffic

If you see a constant stream of unreadable traffic it is likely that the output is set to GCF.

1. Connect the serial port to a PC running Scream and navigate to File > Setup. Select the 'COM Ports' tab and use 'Auto Detect' from the baud rate drop down menu. Allow Scream to determine the baud rate and note the value that it uses.

2. Close Scream and start your terminal emulator. Configure it to the baud rate noted in the previous step using the appropriate commands.

3. Interrupt the GCF data stream by quickly typing:

```
forcegetty
```

It may take several attempts for this to succeed. It can be quicker to copy the command from a text editor and repeatedly paste it into the terminal window.

4. When the interrupt is accepted, the the login prompt will appear.

5. Login to the EAM as normal and then enter the command:

```
gconfig
```

6. You can now configure the EAM as required.

16.4.6.2 A single burst of traffic

If you see a short burst of unreadable traffic and no password prompt it is likely that the baud rate is wrong.

1. Change the baud rate using the appropriate commands for your terminal emulator. Try each plausible baud rate in turn.

2. At each speed, press Enter a few times in the terminal window until the login prompt appears.

3. Login to the EAM and enter the command:

```
gconfig
```

4. You can now configure the EAM as required.

16.4.6.3 No output and connection not possible

If there is no output at all and you are unable to connect to the EAM, please contact Güralp Systems for assistance.

16.5 Appendix E - Connector pinouts

16.5.1 Peli-case: PORTs A, B, C....

These are standard 10-pin “mil-spec” sockets, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-12-10S although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-12-10P and are available from Amphenol, ITT Cannon and other manufacturers.

The pin-out is such that the port can be connected to the serial output of a DM24 digitizer using a straight-through cable.

Pin	Function
A	Power 0 V
B	Power input +10 to +35VDC
C	RS232 RTS
D	RS232 CTS
E	RS232 DTR
F	RS232 DSR
G	RS232 ground
H	RS232 CD
J	RS232 transmit
K	RS232 receive

Wiring details for the compatible plug, ***-12-10P, as seen from the cable end (i.e. during assembly).

16.5.2 Peli-case: Data Out port

This is a standard 10-pin “mil-spec” plug, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-12-10P although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-12-10S and are available from Amphenol, ITT Cannon and other manufacturers.

The pin-out is the same as the serial output of a DM24 digitizer, allowing you to insert a acquisition module into a pre-existing installation and maintain connectivity.

Pin	Function
A	Power input, 0 V
B	Power input +10 to +35VDC
C	RS232 RTS
D	RS232 CTS
E	RS232 DTR
F	RS232 DSR
G	RS232 ground
H	RS232 CD
J	RS232 receive
K	RS232 transmit

Wiring details for the compatible socket, ***-12-10S, as seen from the cable end (i.e. during assembly).

16.5.3 Peli-case: USB

This is a standar d 6-pin “mil-spec” socket, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-10-06S although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-10-06P and are available from Amphenol, ITT Cannon and other manufacturers.

Pin	Function
A	+5 V DC (USB Type A pin 1)
B	Data –ve (USB Type A pin 2)
C	Data +ve (USB Type A pin 3)
D	0 V (USB Type A pin 4)
E	Shielding
F	*not connected*

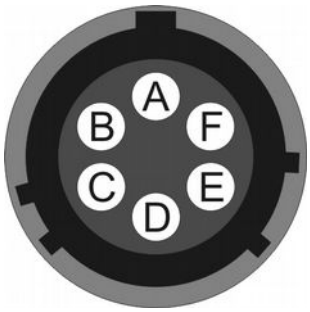

Wiring details for the compatible plug, ***-10-06P, as seen from the cable end (i.e. during assembly).

16.5.4 Peli-case: Network

This is a standard 6-pin “mil-spec” plug, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-10-06P although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-10-06S and are available from Amphenol, ITT Cannon and other manufacturers.

Pin	Function
A	*not connected*
B	Data transmit +ve (RJ45 pin 1)
C	Data receive +ve (RJ45 pin 3)
D	*not connected*
E	Data receive –ve (RJ45 pin 6)
F	Data transmit –ve (RJ45 pin 2)

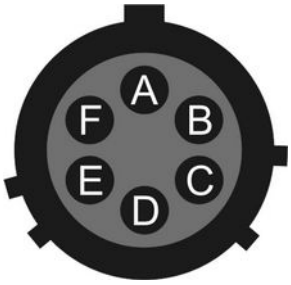

Wiring details for the compatible socket, ***-10-06S, as seen from the cable end (i.e. during assembly).

16.5.5 Peli-case: Console

This is a standard DE9F (TIA-574) sub-miniature (D-sub) line socket, conforming to DIN 41652 and MIL-DTL-24308. They are very widely available, as are suitable mating connectors.

The console port is pre-configured to run at 38,400 Baud with 8 data bits, no parity bit and one stop bit. It is strongly recommended that these settings are never changed, so that access to the configuration system can be gained via this port if all other routes are unavailable.

Pin	Function
1	*not connected*
2	RS232 transmitted data
3	RS232 received data
4	*not connected*
5	Ground
6	*not connected*
7	*not connected*
8	*not connected*
9	*not connected*

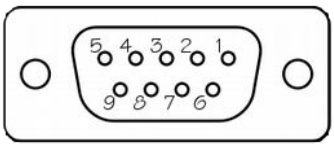

Wiring details for the compatible plug, DE9M, as seen from the cable end (i.e. during assembly).

16.5.6 Cylinder: GPIO

These are standard 12-pin “mil-spec” sockets, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-14-12S although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-14-12P and are available from Amphenol, ITT Cannon and other manufacturers.

The USB lines provide external host access to the internal USB memory device. When power is sensed on pin J, an internal switch disconnects the memory device from the internal circuitry and connects it to this socket.

Pin	Function
A	USB Data -ve (USB Type A pin 2) - see text above.
B	USB Data +ve (USB Type A pin 3) - see text above.
C	Anti-tamper line 4
D	Anti-tamper line 3
E	Anti-tamper line 2
F	Anti-tamper line 1
G	Console transmit (RS232 TXD)
H	Console receive (RS232 RXD)
J	USB Power input (USB Type A pin 1) - see text above.
K	USB Ground (USB Type A pin 4)
L	Anti-tamper line 0
M	Ground

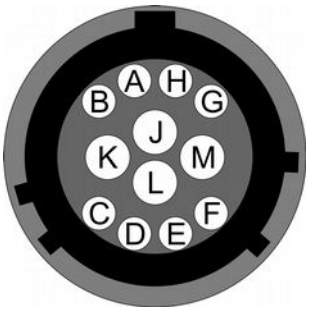

Wiring details for the compatible plug, ***-14-12P, as seen from the cable end (i.e. during assembly).

16.5.7 Cylinder: GPS

This is a standard 10-pin “mil-spec” plug, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-12-10P although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-12-10S and are available from Amphenol, ITT Cannon and other manufacturers.

The pin-out is the same as the GPS input of a DM24 digitizer.

Pin	Function
A	Power 0 V
B	Power +12VDC
C	1pps signal
D	*not connected*
E	Digitizer console transmit
F	Digitizer console receive
G	RS232 ground
H	Digitizer console ground
J	RS232 transmit to GPS
K	RS232 receive from GPS

Wiring details for the compatible socket, ***-12-10S, as seen from the cable end (i.e. during assembly).

16.5.8 Cylinder: USB

This is a standard 6-pin “mil-spec” socket, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-10-06S although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-10-06P and are available from Amphenol, ITT Cannon and other manufacturers.

Pin	Function
A	+5 V DC (USB Type A pin 1)
B	Data –ve (USB Type A pin 2)
C	Data +ve (USB Type A pin 3)
D	0 V (USB Type A pin 4)
E	Shielding
F	*not connected*

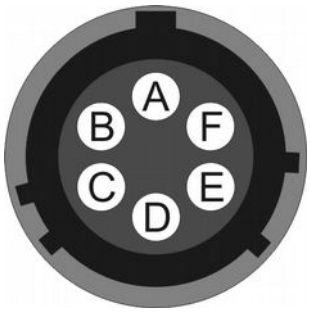

Wiring details for the compatible plug, ***-10-06P, as seen from the cable end (i.e. during assembly).

16.5.9 Cylinder: Power

This is a standard 4-pin “mil-spec” plug, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-08-04P although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-08-04S and are available from Amphenol, ITT Cannon and other manufacturers.

Pin	Function
A	Ground
B	Power input +10 to +36VDC
C	Anti-tamper line 5
D	External switched power output #1

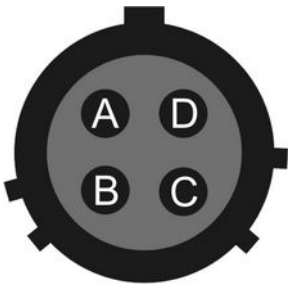

Wiring details for the compatible socket, ***-08-04S, as seen from the cable end (i.e. during assembly).

16.5.10 Cylinder: Ethernet

This is a standard 6-pin “mil-spec” plug, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-10-06P although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-10-06S and are available from Amphenol, ITT Cannon and other manufacturers.

Pin	Function
A	Ground
B	Data transmit +ve (RJ45 pin 1)
C	Data receive +ve (RJ45 pin 3)
D	External switched power output #0
E	Data receive –ve (RJ45 pin 6)
F	Data transmit –ve (RJ45 pin 2)

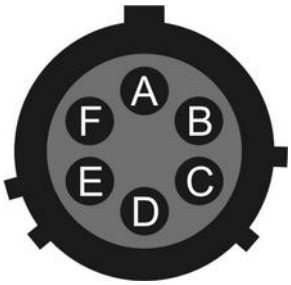

Wiring details for the compatible socket, ***-10-06S, as seen from the cable end (i.e. during assembly).

16.5.11 Cylinder: Data

This is a standard 10-pin “mil-spec” plug, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-12-10P although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-12-10S and are available from Amphenol, ITT Cannon and other manufacturers.

Pin	Function
A	Power input, 0 V
B	Power input +10 to +36VDC
C	RS232 CTS
D	RS232 RTS
E	External trigger output
F	External trigger output
G	RS232 ground
H	External trigger input
J	RS232 receive
K	RS232 transmit

Wiring details for the compatible socket, ***-12-10S, as seen from the cable end (i.e. during assembly).

16.5.12 Sensor Port

This is a standard 26-pin “mil-spec” plug, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-16-26P although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-16-26S and are available from Amphenol, ITT Cannon and other manufacturers.

Pin	Function	Pin	Function
A	Vertical velocity +ve	P	Calibration signal
B	Vertical velocity –ve	R	Vertical calibration enable
C	N/S velocity +ve	S	N/S calibration enable
D	N/S velocity –ve	T	E/W calibration enable
E	E/W velocity +ve	U	Centre
F	E/W velocity –ve	V	Active-high – see section 14.3.2 on page 224
G	Vertical mass position	W	Unlock
H	*not connected*	X	Lock
J	N/S mass position	Y	Logic signal ground
K	Busy indicator LED	Z	*not connected*
L	E/W mass position	a	*not connected*
M	*not connected*	b	Power 0 V
N	Signal ground	c	Power output +10 to +24VDC

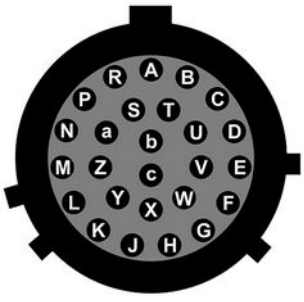

Wiring details for the compatible socket, ***-16-26S, as seen from the cable end (i.e. during assembly).

16.5.13 Cylinder: Auxiliary Input

This is a standard 19-pin “mil-spec” plug, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-14-19P although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-14-19S and are available from Amphenol, ITT Cannon and other manufacturers.

Pin	Function	Pin	Function
A	Optional *SENSOR B* auxiliary / calibration channel +ve	L	*SENSOR A* signal ground
B	Optional SENSOR B auxiliary / calibration channel –ve	M	Optional *SENSOR B* Mux channel `M6`
C	Optional *SENSOR B* signal ground	N	Digital ground
D	Optional *SENSOR B* Mux channel `M3`	P	*not connected*
E	Optional *SENSOR B* Mux channel `M4`	R	*SENSOR A* Mux channel `MB`
F	Optional *SENSOR B* Mux channel `M7`	S	*SENSOR A* Mux channel `MC`
G	*not connected*	T	*SENSOR A* Mux channel `MF`
H	Optional *SENSOR B* Mux channel `M5`	U	*SENSOR A* Mux channel `MD`
J	*SENSOR A* auxiliary / calibration channel +ve	V	*SENSOR A* Mux channel `ME`
K	*SENSOR A* auxiliary / calibration channel –ve		

Wiring details for the compatible socket, ***-14-19S, as seen from the cable end (i.e. during assembly).

16.5.14 DM24S24EAM: Sensor Inputs

This is a standard 32-pin “mil-spec” plug, conforming to MIL-DTL-26482 (formerly MIL-C-26482). A typical part-number is 02E-16-32P although the initial “02E” varies with manufacturer.
Suitable mating connectors have part-numbers like ***-16-32S and are available from Amphenol, ITT Cannon and other manufacturers.

Pin	Function	Pin	Function
A	Channel 0 +ve	T	E/W calibration enable
B	Channel 0 –ve	U	Centre
C	Channel 1 +ve	V	*not connected*
D	Channel 1 –ve	W	Unlock
E	Channel 2 +ve	X	Lock
F	Channel 2 –ve	Y	Logic signal ground
G	Vertical mass position	Z	*not connected*
H	*not connected*	a	*not connected*
J	N/S mass position	b	Power 0 V
K	Busy indicator LED	c	Power output +10 to +24VDC
L	E/W mass position	d	Channel 3 +ve
M	*not connected*	e	Channel 3 –ve
N	Signal ground	f	Channel 4 +ve
P	Calibration signal	g	Channel 4 –ve
R	Vertical calibration enable	h	Channel 5 +ve
S	N/S calibration enable	j	Channel 5 –ve

Wiring details for the compatible socket, ***-16-32S, as seen from the cable end (i.e. during assembly).

16.6 Appendix F – Open source software and the GPL

16.6.1 Introduction

Platinum firmware contains open-source programs covered by the GNU Public License (GPL). The exact details vary from release to release but full, current details are always available from our web site at http://www.guralp.com/platinum/opensource/.

16.6.2 Physical copies of source code

Under the terms of the GPL, GSL has an obligation to offer to ship the source code for such programs (and associated build scripts) on a physical storage medium customarily used for information interchange, such as a CD. If you would like the source code in physical format, please contact support@guralp.com. A small charge will be made for administration, the medium and postage.

16.6.3 The GNU General Public License

The license itself is available on the web at www.gnu.org/licenses.

17 Revision history

February 2014	E	Added GPL information and warnings about back-fill with gdi-link-tx.
November 2013	D	Significant rewrite to cover many new features and revised user interface.
August 2011	C	Revised firmware mirror server instructions. Added PPP watchdog expert mode details.
May 2011	B	Hardware information updated. Update of Platinum overview. System status details updated. New section on U3 USB. Formatting updates.
Dec 2010	A	New document