

Baler44 Application Note for 20091229 Release

**Quanterra, Inc.
325 Ayer Rd.
Harvard, MA 01451**

- 1. INTRODUCTION 3**
 - 1.1 SCOPE OF DOCUMENT 3
 - 1.2 FIRMWARE REVISION LEVELS AND SUPPORT TOOLS 3
- 2. RECORDING MODES OF THE BALER44 3**
 - 2.1 MEDIA TYPES AND CAPACITIES 4
- 3. BALER CONFIGURATION OPTION STRING: “CIRCULAR” MODE AND OTHER FEATURES 5**
- 4. DATA RETRIEVAL 7**
 - Directory Listing of PacketBaler44 Tag 11552* 8
 - INDEX OF /WDIR/ 9
 - 4.1 WGET 9
 - 4.2 WGET EXAMPLE 9
 - 4.3 CURL EXAMPLES 11
- 5. REMOTE ADMINISTRATION 11**
 - 5.1 EXAMPLES USING WINXX & UNIX PLATFORMS 12
 - 5.2 UPDATE OF OLDER BALER44 FIRMWARE VERSIONS: WHAT IF I DON’T HAVE LOGIN CREDENTIALS? 14
- 6. SECURITY CHECKLIST 14**

Revisions History

2010/01/13 Rev A Initial Release

2010/01/13 Rev B added discussion pertaining to update of Baler44 firmware prior to 20090920

2010/01/20 Rev C typographic changes

1. Introduction

1.1 Scope of Document

This document provides reference information in an FAQ format for updating information in the *Baler44-Operation-Guide-RevB.pdf*, available on the www.q330.com website. In addition you should consult the release notes for this package.

1.2 Firmware Revision Levels and Support Tools

Support of features discussed in this appnote requires minimum versions of Baler44 and Q330 firmware and support tools:

- Willard Version 1.148 build 661 (20091229)
- Q330 Version 1.137
- Baler44 Version 20091229

2. Recording Modes of the Baler44

As of the most recent B44 firmware (20091229) and Q330 firmware (1.137), the USB media can be used in two different modes:

1. The default "linear" mode is to fill one drive, then move to the next one. Recording stops when both are full. Files may be deleted manually (script-driven) over-the-air to make room for new ones. The total available storage is the sum of the two media.
2. A new mode, where one drive is treated as "circular" storage, always managed to maintain a few % free. The second drive is kept as a hot standby in case of failure of the first one, at which point the second drive becomes the active "circular" one. The total active storage is normally the size of one drive in this mode.

A fixed-term deployment would most likely use mode 1, where the total media capacity is sized to contain all the data from the deployment, e.g. a PASSCAL-type experiment, or a 2-year TA-type deployment. For example, the data for a TA deployment fits on one 16G drive.

A permanent installation, however, may treat the local storage as circular. In that case the media capacity controls the residence time. Particularly in the case of comparatively high data generation rate, say, 6 channels at 100sps or possibly 200sps for strong motion (equivalent to 20-40Gb/year), smaller-capacity, e.g. 4G, drive media may be overwritten many times.

2.1 Media types and capacities

The 4G & 8G USB media that we ship with systems are normally so-called "SLC" (Single Layer Cell) flash media, that are a high-integrity type of device; these cost about 3-4x per Gb vs. the type that you can get at Best Buy that are called "MLC" (Multi-Level Cell). We have successfully tested, and many deployments use, including the TA, two 16G MLC drives, for a total capacity of 32G in the mode in which they are being used.

The SLC media are available only in the 4G & 8G size (2x gives 16G total). The MLC are available in capacities up to 64Gb. The B44 can at present use a maximum of 2x16G=32G total drive size. They will actually operate and record on 32G media (total 2x32 = 64G), however, retrieval from these large media is not possible over the air, only by physical removal, due to limitations in the http & ftp servers in dealing with large file systems.

What do the recording modes have to do with media type?

The SLC media might be preferred for the permanent-station/circular mode, because of the many re-write (mode 2) operation. An 8G SLC drive would probably give more than 2 months minimum data residence time operated in circular mode with high sample-rate network rates. At a slightly higher risk to long-term drive integrity, the 16G MLC media could be used to double that residence time. Of course MLC media could be replaced every 2 years or so to mitigate wear-out. Because of the lower capital cost, and the fact that all of the stations will normally be operated with telemetry, and therefore it is possible to keep a close eye on the drive media condition using automated procedures, there is unlikely to be a serious risk in the use of the MLC media.

What about various drive media manufacturers?

Some work, some don't. KMI & Quanterra have screening facilities for various commercial MLC media. This involves running, over a wide temperature range, an accelerated whole-media random read/write pattern while powering the drive up and down thousands of times. This procedure has resulted in several recommended brands, and several "don't buy" brands. In particular, the drives that come off the shelf with various "accelerator" software packages, such as "U3" on Sandisk products, are generally not compatible. Ron Martinolich at KMI can provide some further information on available MLC media.

Is special configuration required to utilize multiple drive media?

There isn't any configuration required. Each time the media are powered up (whenever data are read are written) the drive population is checked. The B44 automatically adjusts to use the two media (if available) in accordance with the selected mode, "linear" or "circular". You can format two drives at the same time automatically. It is also possible to mix capacities.

3. Baler Configuration Option String: “circular” mode and other features

New Baler44 features are configured in a "baler configuration option string". This is a string containing multiple "name=value" pairs to control a range of new functions that may be added from time to time.

Where is the "baler configuration option string" set?

In the Willard Interfaces|Serial 2|Advanced area, as shown below. The “baler configuration option string” appears at the bottom of the form in the text box, and is greyed and cannot be edited, to prevent inadvertent changes, unless the "unlock" box is checked.

How does circular mode work?

In circular mode, an attempt is made to maintain the total utilization of the media at 5% free by deleting the oldest files. Because there is a primary (*data*) and secondary (*sdata*) set of data, there is more than one approach to deletion of the oldest files. When deleting, either files in all directories (*data,wfdisc,sdata,md5*) will be deleted, or all directories except those in the the *sdata* directory will be deleted. This latter mode specifies a "soft" circular operation. In this case, eventually, the entire media will be filled by files in the *sdata* directory. The default option string setting to enable circular recording is "circ=sdata". This will cause files in all directories, including the *sdata* directory, to be

deleted when space is needed. The parameter in the "circ=" specification should be thought of as the directory in which files will be searched to find the oldest. If starting at the *data* directory, the files in the *sdata* directory are omitted from the search, and hence will remain on the media.

Is it possible to see the names of the last file deleted?

Yes, in the B44 web page, the names of the last file written (most recent) and the last file deleted (oldest) are presented. Note in the example excerpt below, the media capacity is being maintained at 5%. This media is in "soft" circular mode at present, with files building in the *sdata* directory

```
Copyright Quanterra, Inc. BALER44-20091229-16K tag 29000 at 2010-01-13 16:54:41
last baler reboot: 2009-12-30 20:11:44   reboots: 128   runtime: 13.863d
MEDIA GOOD
MEDIA site 1 crc=0x5e9f IN USE state: ACTIVE   capacity=7963.078Mb   free=5.003%
upsvolts=13.123 primaryvolts=13.385 degc=18.539
baler cfg options:   circ=data
...
media last write time: 2010-01-13 14:03:22
last media file written: /mnt/data/QT-HRV_4-20100113140311
last media file deleted: /mnt/data/QT-HRV_4-20090408022248
```

Why would "circ=data" be used, not the default mode "circ=sdata"?

The residence time of data on the media in circular mode is determined by the data generation rate and the capacity. In the event that some important information in the *sdata* directory are desired to be maintained on the media and are in danger of being deleted because space is needed, soft circular operation may be specified to retain older data in the *sdata* directory. This should be considered as a temporary measure to allow time for the desired data to be recovered from the media, however, since eventually the media will fill in soft circular mode. In addition, it is possible for many files (>4000) to be created on the media in soft mode, which is beyond the capacity of the web server to display. The files may be present, although remote access is limited in this mode. Use soft circular only in an emergency situation to preserve *sdata* contents until recovery is possible.

Additional options and notes on circular mode.

- files in the *cont* directory are not included in the circular deletion. these will remain and take up negligible space.
- files in the *recover* directory resulting from media fsck recoveries are not included in the circular operation.
- options may be specified to target a specific level of free space in circular mode, and the "aggressivity" with which the target is approached.

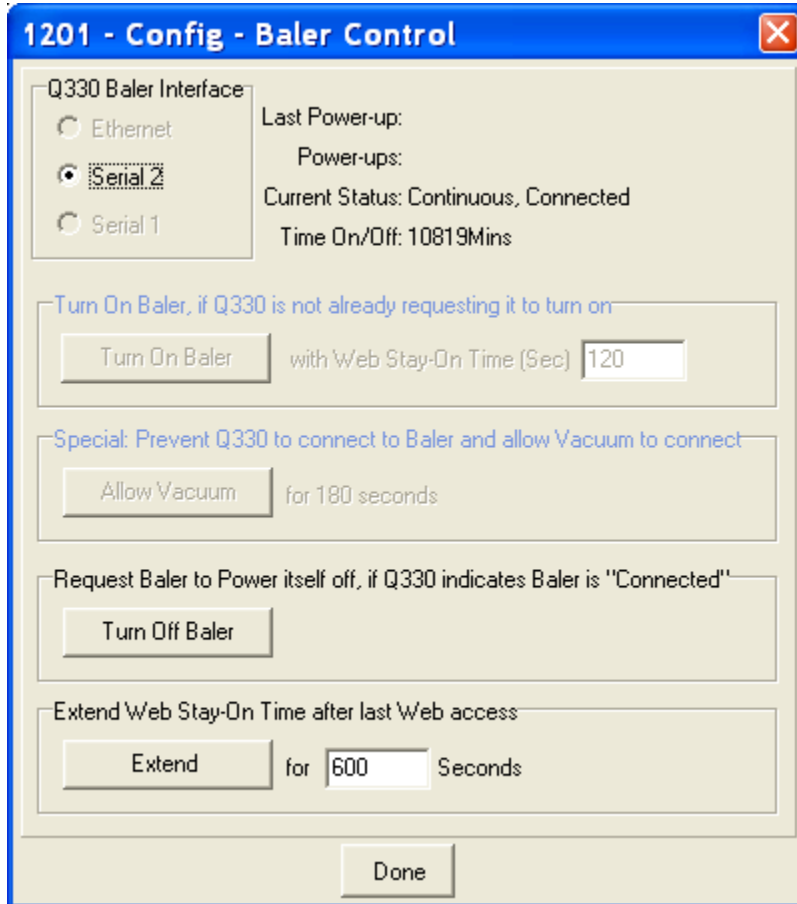
```
circ=sdata[,xx,rr]
```

xx is optional target free threshold in % free. default 5.

rr is optional max number of files removed per operation default 3. max 19.

Will the Baler recognize the Option string immediately when it is set?

The Baler44 recognizes the option string (and other settings) when it connects with the Q330. Therefore the Baler44 must be rebooted to read new option string settings. This does not result in loss of data. After saving Q330 settings to EEPROM, you may use the baler control command to turn off the baler. When the baler reboots, the option string settings will be read. Note, as in the example above, the option string settings are echoed in the Baler44 web page as a confirmation.



4. Data Retrieval

Does data retrieval from the Baler44 using http work the same way as from a Baler14?

No. Remote access on the B44 is a view directly in to the media file system using either http or ftp. There are just directories with MSEED files in them. The main difference is that on a B14, although you can specify down to the record (based on time) which data you want, you can never resume an interrupted transmission once begun because the B14 assembles a "virtual download" on the fly.

The B44 sends the files you ask for, and you can ask for them as many times as you like, and resume interrupted transmissions (using standard http and ftp resume methods). scripting ftp engines, or "wget --mirror" therefore allow you to move an arbitrarily large number of requested files without intervention, automatically taking care of retry.

The Q330 and Baler44 can be connected serially. In our legacy stations, we use the Baler14 where the baler is not on the ethernet at all and gets its http services via the Q330 ethernet and NAT to the serial interface. Can the Baler44 run this old method?

No. There's a better way with the Baler44. The approach is to turn around the connection. The B44 provides routing services from its Ethernet interface to the Q330's serial port. The advantages of this approach are that all of the Q330's services are available and in addition all of the baler services (http,ftp, ssh for maintenance) are available not throttled by the serial link or constrained by the limited Q330 MTU. Downloads from the B44 can therefore take place at maximum rates the ethernet links can sustain. For this to work the continuously powered Baler option must be selected (Willard|Configuration|Interfaces|Serial 2 or Ethernet tab|Advanced button).

My network restricts UDP traffic, so I can't talk to the Q330 for remote administration. What to do?

The B44 provides TCP access to the basic Q330 ports allowing access to the Q330 over networking infrastructure that may restrict UDP traffic. In Willard, check the "Tunnel330" box in the registration form. Similarly Pecos2 and lib330 provide this capability.

Is there some specific documentation of the URL's to access the B44 web interface?

The file/directory structure is described in the *Baler44 Operations Guide*. The web and ftp views into the file structure are verbatim, and use industry-standard ftp and http servers (unlike the B14). So, if you point at a B44 using a URL for the web server, e.g. <http://66.189.86.236:6381/>,

you get:

Directory Listing of PacketBaler44 Tag 11552

[activemedia](#)
[reservemedia](#)

[current status](#)

If you go to the "activemedia" you get the media root directory:

Index of /WDIR/

mode	links	bytes	last-changed	name
dr-x	8	32768	Jan 1 1970	./
dr-x	1	512	Jun 27 2008	../
dr-x	10	32768	Dec 31 00:42	admin/
dr-x	2	32768	Jan 13 17:02	cont/
dr-x	2	98304	Jan 13 17:02	data/
-r-x	1	1	Mar 7 2008	format_prevent*
dr-x	2	98304	Jan 13 16:58	recover/
dr-x	2	262144	Jan 13 17:02	sdata/
dr-x	2	196608	Jan 13 16:58	wfdisc/

...and so on. These look just like the directory structure described in the manual.

Since the file names are ever-time-increasing, requests for specific time ranges can be generated by sorting and selecting file names from the directory listings. Below is an example of moving a group of files matching an arbitrary name template, and hence timestamp range, using "wget".

4.1 wget

Off-the-shelf gnu wget can recover all (or selected) files from a B44. wget has a "mirror" mode where it looks at time stamps, caches them on the local receiving system (if an ftp url is used), and gets whatever is changed remotely. It will resume an interrupted transfer. A single cmd line can mirror the entire baler media image, or if you drill down the baler directory tree, only the desired directory or file names, such as a restricted date range, can be selected.

wget operates in a sense similarly to rsync, except that rsync does a 100% byte-for-byte check on contents. wget `--mirror` uses only length and timestamps to sync the local copy. In the latest B44 firmware, the *recover* directory contains md5 checksum files written for every file in the *data* directory. These can be used to support integrity checking, regardless of the protocol.

4.2 wget example

The example below shows resuming an aborted wget session. Note that the last file that was partially transferred is recovered again after wget is restarted and it consults the cached directory listing to figure out where to resume. Note that any matching url with wild cards can be specified. e.g. "...activemedia/data/*200906*" recovers data files for the month of June 2009.

```
steim@linux-elx4:~/mir> wget --mirror -w 2
"ftp://66.189.86.236:6382/activemedia/sdata/*"
--2009-07-02 09:43:11-- ftp://66.189.86.236:6382/activemedia/sdata/*
=> `66.189.86.236:6382/activemedia/sdata/.listing'
Connecting to 66.189.86.236:6382... connected.
Logging in as anonymous ... Logged in!
==> SYST ... done.      ==> PWD ... done.
```

```
==> TYPE I ... done. ==> CWD /activemedia/sdata ... done.
==> PASV ... done. ==> LIST ... done.
```

```
[          <=>
] 115,703    41.1K/s
in 2.7s
```

```
2009-07-02 09:43:15 (41.1 KB/s) - `66.189.86.236:6382/activemedia/sdata/.listing'
saved [115703]
```

```
Remote file no newer than local file `66.189.86.236:6382/activemedia/sdata/QT-
0505_4-S-20081114180731' -- not
retrieving.
```

```
Remote file no newer than local file `66.189.86.236:6382/activemedia/sdata/QT-
0505_4-S-20081114200942' -- not
retrieving.
```

```
Remote file no newer than local file `66.189.86.236:6382/activemedia/sdata/QT-
0505_4-S-20081114221148' -- not
retrieving.
```

```
Remote file no newer than local file `66.189.86.236:6382/activemedia/sdata/QT-
0505_4-S-20081115001009' -- not
retrieving.
```

```
The sizes do not match (local 50368) -- retrieving.
```

```
--2009-07-02 09:43:17-- ftp://66.189.86.236:6382/activemedia/sdata/QT-0505_4-S-
20081115021116
```

```
=> `66.189.86.236:6382/activemedia/sdata/QT-0505_4-S-20081115021116'
```

```
==> CWD not required.
```

```
==> PASV ... done. ==> RETR QT-0505_4-S-20081115021116 ... done.
```

```
Length: 196608 (192K)
```

```
100%[=====
======>] 196,608    57.9K/s
in 3.3s
```

Is it possible to effect retrieval of requested MSEED data recorded on the USB media down to the individual record level?

Yes. The *wfdisc* directory contains the *.wfdisc* index files corresponding to each data file. These are maps of each record in the files in the *data* directory. You can use these to locate a specific channel/time and then issue an http or ftp request with a range specifier to retrieve exactly the records you want. If you want to retrieve selective MSEED records, the process would then be to locate the nearest *wfdisc* files (these are stored as *.gz*'s, and are small), find the time range you want in the *wfdisc*s, and then issue the http or ftp request you need.

A tool that allows specifying range values for http and ftp transfers is "curl" .

4.3 curl examples

SINGLE RECORD, BEGINNING AT RECORD #10, HTTP:

```
$ curl -r40960-45055 -O http://66.189.86.236:6381/WDIR/data/QT-1013_4-20091001004439
% Total    % Received % Xferd Average Speed   Time    Time     Time  Current
           Dload  Upload  Total   Spent    Left  Speed
102 4096 102 4096  0    0   210    0 0:00:19 0:00:19 --:--:-- 1120
```

ENTIRE FILE, FTP:

```
$ curl -O ftp://
66.189.86.236:6382/activemedia/data/QT-1013_4-20090930000318
% Total    % Received % Xferd Average Speed   Time    Time     Time  Current
           Dload  Upload  Total   Spent    Left  Speed
100 4000k 100 4000k  0    0 55258    0 0:01:14 0:01:14 --:--:-- 63731
```

SINGLE RECORD, BEGINNING AT RECORD #10, FTP:

```
$ curl -r40960-45055 -O ftp://66.189.86.236:6382/activemedia/data/QT-1013_4-20090930000318
** Resuming transfer from byte position 40960
% Total    % Received % Xferd Average Speed   Time    Time     Time  Current
           Dload  Upload  Total   Spent    Left  Speed
102 4096 102 4096  0    0 1447    0 0:00:02 0:00:02 --:--:-- 6553
```

The Baler14 stored data on disk in files demultiplexed by channel. Can the Baler44 store data in this mode?

No. The Baler44 stores data in constant-length files containing MSEED records multiplexed by channel. These data are stored in exactly the order in which the data are generated. The data files once written to the USB media are static. The "view" of the file system, whether through remote access, such as http or ftp, or by direct access to the media when plugged into a host computer, is identical.

5. Remote Administration

*ftp appears to support only anonymous ftp (with "passive" mode). > Is this correct?
I can't delete files with anonymous ftp.*

Correct. This is a security feature to prevent arbitrary users from deleting data.

What about ssh/scp access?

ssh/scp access to the B44 is possible. The port numbers and procedures are not discussed in the general manual as a security precaution. The first step in managing ssh/scp access is to define a unique user name and password. This is specified in the baler configuration option string using the "auth=" keyword. When a single value is supplied, e.g. auth=pw, the value will be interpreted as a password to be used along with a default user name. If two values are specified, the values are interpreted as your specific username and password. Examples:

	resulting	
string	uname	password
auth=pw	default	pw
auth=un,pw	un	pw
auth=un,	un	un
auth=	default	default

Consult Quanterra for the default username

If you want to specify a password not in clear text, you can supply a linux-gnu standard encrypted password, that looks something like "\$1\$M8sOIuIvkt26N.olXGk2H0"

How can tasks such as remote firmware update or management of specific data files be accomplished using administrative (ssh/scp) access?

There are two ssh ports available for remote access:

baler baseport + 55 (default 6385) if USB media access is not required or desired
 baler baseport + 56 (default 6386) if USB media access is required

The main task for which ssh/scp will be used is firmware update. This comprises copying a Quanterra-supplied update file to the active media, and restarting the baler, which will initiate automatic installation. This uses the +56 port. Of course any media operation may be performed, including copying or deleting certain files.

Note that manipulation of the root file system and passwd and ssh database is not permitted by hardware write protection. Consult Quanterra.

5.1 Examples using WinXX & Unix platforms

A comprehensive ssh package, "putty" is available for windows machines. In this package is a set of command line tools. In particular the "plink" command allows remote command execution (and

specification of the login password on the command line), while "pscp" can be used to move files. That package is available at:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

The example below shows an initial contact (and key cache) to a remote station (just listing the media) followed by creation of the *admin* directory, sending a firmware update, and requesting a reboot.

```
C:\>plink -ssh -P 5386 -l username -pw password 166.241.252.234 "ls -l /mnt"
The server's host key is not cached in the registry. You
have no guarantee that the server is the computer you
think it is.
The server's rsa2 key fingerprint is:
ssh-rsa 1040 3f:9d:48:a7:fb:9f:34:2f:ac:fc:dd:42:76:3f:ac:cf
If you trust this host, enter "y" to add the key to
PuTTY's cache and carry on connecting.
If you want to carry on connecting just once, without
adding the key to the cache, enter "n".
If you do not trust this host, press Return to abandon the
connection.
Store key in cache? (y/n) y
drwxr-xr-x  2 root    root          16384 Jun 26 22:11 cont
drwxr-xr-x  2 root    root          16384 Jul  7 20:59 data
-rwxr-xr-x  1 root    root              1 Feb  5 21:43 format_prevent
drwxr-xr-x  2 root    root          16384 Feb  5 21:43 recover
drwxr-xr-x  2 root    root          16384 Jul  7 20:59 sdata
drwxr-xr-x  2 root    root          16384 Jul  7 20:59 wfdisc

C:\>plink -ssh -P 5386 -l username -pw password 166.241.252.234 "mkdir /mnt/admin"

C:\>pscp -scp -P 5386 -l username -pw password b44update-20090707A-16K-
sg.tar.jz.asc 166.241.252.234:/mnt/admin/b44update-20090707A-16K-sg.tar.jz.asc
b44update-20090707A-16K-s | 806 kB |  17.9 kB/s | ETA: 00:00:00 | 100%

C:\>plink -ssh -P 5386 -l username -pw password 166.241.252.234 "echo
>/tmp/turnoff"
```

Similarly from a unix host...

```
<sun> joe 15% /usr/local/bin/ssh -p 5386 -l admin 166.241.252.236 "ls /mnt"
The authenticity of host '166.241.252.236 (166.241.252.236)' can't be established.
RSA key fingerprint is 3f:9d:48:a7:fb:9f:34:2f:ac:fc:dd:42:76:3f:ac:cf.
Are you sure you want to continue connecting (yes/no)? y
Please type 'yes' or 'no': yes
Warning: Permanently added '166.241.252.236' (RSA) to the list of known hosts.
admin@166.241.252.236's password:
admin
cont
data
format_prevent
recover
sdata
wfdisc
<sun> joe 16%
```

5.2 Update of older Baler44 firmware versions: what if I don't have login credentials?

Above we've discussed remote update of Baler44 firmware using pscp or scp to copy Quanterra-supplied updates over an ssh connection. You may note that this is possible only if the ssh login credentials can be supplied to the Baler44 for ssh access. The method to control the ssh login credentials, using specification in the baler configuration option string, is also discussed above. Note, however, that the ability to set the Baler44 login credentials using the option string not only requires the most recent Q330 firmware modules and Willard support tool, but also requires recent Baler44 firmware (20090920 or later). The question arises, then, how can the Baler44 firmware be updated if you don't know the Baler44 login credentials, and the firmware presently installed on the Baler44 is older than 20090920, and therefore does not support specification of your own private login credentials? There are two approaches:

- The Baler44 firmware may be updated on-site using the update file copied onto the USB active media. This requires no login credentials.
- Connect the Baler44 with Q330 to the internet and assign publicly-addressable IP. Quanterra can perform the necessary update remotely to install current firmware.

Once the Baler44 firmware has been updated, the login credentials may be set to user preferences using the new capabilities of current firmware. Future updates can then be performed remotely using your private credentials.

If neither of these options are possible, consult Quanterra.

6. SECURITY CHECKLIST

- ssh/scp password. set this in the baler config option string.
- set a non-trivial Q330 authorization code
- use a private network
- if you use a public network and want to restrict access, the Baler44 has a simple source-IP based firewall capability that can be set to adjust access to the Q330 and Baler from up to 4 known IP's. By default this capability is turned off. This needs to be configured on-site using a terminal emulator, prior to field deployment. Consult Quanterra.

END OF DOCUMENT