

Botan Build Guide

Jack Lloyd
lloyd@randombit.net

2008-09-30

Contents

1	Introduction	2
2	For the Impatient	2
3	Building the Library	2
3.1	POSIX / Unix	4
3.2	MS Windows	4
3.3	Configuration Parameters	5
3.4	Multiple Builds	5
3.5	Local Configuration	5
4	Modules	6
5	Building Applications	7
5.1	Unix	7
5.2	MS Windows	7

1 Introduction

This document describes how to build Botan on Unix/POSIX and MS Windows systems. The POSIX oriented descriptions should apply to most common Unix systems (including MacOS X), along with POSIX-ish systems like BeOS, QNX, and Plan 9. Currently, systems other than Windows and POSIX (such as VMS, MacOS 9, OS/390, OS/400, ...) are not supported by the build system, primarily due to lack of access. Please contact the maintainer if you would like to build Botan on such a system.

Botan's build is controlled by `configure.pl`, which is a Perl script. Perl 5.6 or later is required.

2 For the Impatient

```
$ ./configure.pl [--prefix=/some/directory]
$ make
$ make install
```

Or using `nmake`, if you're compiling on Windows with Visual C++. On platforms that do not understand the `'#!'` convention for beginning script files, or that have Perl installed in an unusual spot, you might need to prefix the `configure.pl` command with `perl` or `/path/to/perl`.

3 Building the Library

The first step is to run `configure.pl`, which is a Perl script that creates various directories, config files, and a Makefile for building everything. The script requires at least Perl 5.6; any later version should also work.

The script will attempt to guess what kind of system you are trying to compile for (and will print messages telling you what it guessed). You can override this process by passing the options `--cc`, `--os`, and `--arch` – acceptable values are printed if you run `configure.pl` with `--help`.

You can pass basically anything reasonable with `--cpu`: the script knows about a large number of different architectures, their sub-models, and common aliases for them. The script does not display all the possibilities in its help message because there are simply too many entries. You should only select the 64-bit version of a CPU (such as "sparc64" or "mips64") if your operating system knows how to handle 64-bit object code – a 32-bit kernel on a 64-bit CPU will generally not like 64-bit code.

By default the script tries to figure out what will work on your system, and use that. It will print a display at the end showing which algorithms have and have not been abled. For instance on one system we might see the line:

```
(loading): entropy: [beos_stats] buf_es [cryptoapi_rng] \\  
                dev_random egd proc_walk unix_procs [win32_stats]
```

The names listed in brackets are disabled, the others are enabled. Here we see the list of entropy sources which are going to be compiled into Botan. Since this particular line comes when Botan was configuring for a Linux system, the Win32 and BeOS specific modules were disabled, while modules that use Unix APIs and `/dev/random` are built.

You can control which algorithms and modules are built using the options `--enable-modules=MODS` and `--disable-modules=MODS`, for instance `--enable-modules=blowfish,md5,rsa,zlib --disable-modules=arc4,cmac`. Modules not listed on the command line will simply be loaded if needed or if configured to load by default.

Not all Oses or CPUs have specific support in `configure.pl`. If the CPU architecture of your system isn't supported by `configure.pl`, use `'generic'`. This setting disables machine-specific optimization flags.

Similarly, setting OS to 'generic' disables things which depend greatly on OS support (specifically, shared libraries).

However, it's impossible to guess which options to give to a system compiler. Thus, if you want to compile Botan with a compiler which `configure.pl` does not support, you will need to tell it how that compiler works. This is done by adding a new file in the directory `src/build-data/cc`; the existing files should put you in the right direction.

The script tries to guess what kind of makefile to generate, and it almost always guesses correctly (basically, Visual C++ uses NMAKE with Windows commands, and everything else uses Unix make with POSIX commands). Just in case, you can override it with `--make-style=somestyle`. The styles Botan currently knows about are 'unix' (normal Unix makefiles), and 'nmake', the make variant commonly used by Windows compilers. To add a new variant (eg, a build script for VMS), you will need to create a new template file in `src/build-data/makefile`.

3.1 POSIX / Unix

The basic build procedure on Unix and Unix-like systems is:

```
$ ./configure.pl [--enable-modules=<list>] [--cc=CC]
$ make
# You may need to set your LD_LIBRARY_PATH or equivalent for ./check to run
$ make check # optional, but a good idea
$ make install
```

This will probably default to using GCC, depending on what can be found within your PATH.

The `make install` target has a default directory in which it will install Botan (typically `/usr/local`). You can override this by using the `--prefix` argument to `configure.pl`, like so:

```
./configure.pl --prefix=/opt <other arguments>
```

On some systems shared libraries might not be immediately visible to the runtime linker. For example, on Linux you may have to edit `/etc/ld.so.conf` and run `ldconfig` (as root) in order for new shared libraries to be picked up by the linker. An alternative is to set your `LD_LIBRARY_PATH` shell variable to include the directory that the Botan libraries were installed into.

3.2 MS Windows

The situation is not much different here. We'll assume you're using Visual C++ (for Cygwin, the Unix instructions are probably more relevant). You need to have a copy of Perl installed, and have both Perl and Visual C++ in your path.

```
> perl configure.pl --cc=msvc --os=windows [--cpu=CPU]
> nmake
> nmake check # optional, but recommended
```

For Win95 pre OSR2, the `cryptoapi_rng` module will not work, because CryptoAPI didn't exist. And all versions of NT4 lack the ToolHelp32 interface, which is how `win32_stats` does its slow polls, so a version of the library built with that module will not load under NT4. Later systems (98/ME/2000/XP) support both methods, so this shouldn't be much of an issue.

Unfortunately, there currently isn't an install script usable on Windows. Basically all you have to do is copy the newly created `libbotan.lib` to someplace where you can find it later (say, `C:\botan\`). Then copy the entire `build\include\botan` directory, which was constructed when you built the library, into the same directory.

When building your applications, all you have to do is tell the compiler to look for both include files and library files in `C:\botan`, and it will find both.

3.3 Configuration Parameters

There are some configuration parameters which you may want to tweak before building the library. These can be found in `config.h`. This file is overwritten every time the configure script is run (and does not exist until after you run the script for the first time).

Also included in `config.h` are macros which are defined if one or more extensions are available. All of them begin with `BOTAN_HAS_`. For example, if `BOTAN_HAS_COMPRESSOR_BZIP2` is defined, then an application using Botan can include `<botan/bzip2.h>` and use the Bzip2 filters.

BOTAN_MP_WORD_BITS: This macro controls the size of the words used for calculations with the MPI implementation in Botan. You can choose 8, 16, 32, or 64, with 32 being the default. You can use 8, 16, or 32 bit words on any CPU, but the value should be set to the same size as the CPU's registers for best performance. You can only use 64-bit words if an assembly module (such as `mp_ia32` or `mp_asm64`) is used. If the appropriate module is available, 64 bits are used, otherwise this is set to 32. Unless you are building for a 8 or 16-bit CPU, this isn't worth messing with.

BOTAN_VECTOR_OVER_ALLOCATE: The memory container `SecureVector` will over-allocate requests by this amount (in elements). In several areas of the library, we grow a vector fairly often. By over-allocating by a small amount, we don't have to do allocations as often (which is good, because the allocators can be quite slow). If you *really* want to reduce memory usage, set it to 0. Otherwise, the default should be perfectly fine.

BOTAN_DEFAULT_BUFFER_SIZE: This constant is used as the size of buffers throughout Botan. A good rule of thumb would be to use the page size of your machine. The default should be fine for most, if not all, purposes.

BOTAN_GZIP_OS_CODE: The OS code is included in the Gzip header when compressing. The default is 255, which means 'Unknown'. You can look in RFC 1952 for the full list; the most common are Windows (0) and Unix (3). There is also a Macintosh (7), but it probably makes more sense to use the Unix code on OS X.

3.4 Multiple Builds

It may be useful to run multiple builds with different configurations. Specify `--build-dir=<dir>` to set up a build environment in a different directory.

3.5 Local Configuration

You may want to do something peculiar with the configuration; to support this there is a flag to `configure.pl` called `--with-local-config=<file>`. The contents of the file are inserted into `build/build.h` which is (indirectly) included into every Botan header and source file.

4 Modules

There are a fairly large number of modules included with Botan. Some of these are extremely useful, while others are only necessary in very unusual circumstances. The modules included with this release are:

- **alloc_mmap**: Allocates memory using memory mappings of temporary files. This means that if the OS swaps all or part of the application, the sensitive data will be swapped to where we can later clean it, rather than somewhere in the swap partition.
- **bzip2**: Enables an application to perform bzip2 compression and decompression using the library. Available on any system that has bzip2.
- **zlib**: Enables an application to perform zlib compression and decompression using the library. Available on any system that has zlib.
- **gnump**: An engine that uses GNU MP to speed up PK operations. GNU MP 4.1 or later is required.
- **openssl**: An engine that uses OpenSSL to speed up public key operations and some ciphers/hashes. OpenSSL 0.9.7 or later is required.
- **beos_stats**: An entropy source that uses BeOS-specific APIs to gather (hopefully unpredictable) data from the system.
- **cryptoapi_rng**: An entropy source that uses the Win32 CryptoAPI function `CryptGenRandom` to gather entropy. Supported on NT4, Win95 OSR2, and all later Windows systems.
- **egd**: An entropy source that accesses EGD (the entropy gathering daemon). Common on Unix systems that don't have `/dev/random`.
- **proc_walk**: Gather entropy by reading files from a particular file tree. Usually used with `/proc`; most other file trees don't have sufficient variability over time to be useful.
- **unix_procs**: Gather entropy by running various Unix programs, like `arp` and `vmstat`, and reading their output in the hopes that at least some of it will be unpredictable to an attacker.
- **win32_stats**: Gather entropy by walking through various pieces of information about processes running on the system. Does not run on NT4, but should run on all other Win32 systems.
- **fd_unix**: Let the users of `Pipe` perform I/O with Unix file descriptors in addition to `iostream` objects.
- **pthread**: Add support for using `pthread` mutexes to lock internal data structures. Important if you are using threads with the library.
- **qt_mutex**: Add support for using Qt mutexes to lock internal data structures.
- **cpu_counter**: Use the contents of the CPU cycle counter when generating random bits to further randomize the results. Works on x86 (Pentium and up), Alpha, and SPARCv9.
- **posix_rt**: Use the POSIX realtime clock as a high-resolution timer.
- **gettimeofday**: Use the traditional Unix `gettimeofday` as a high resolution timer.
- **win32_query_perf_ctr**: Use Win32's `QueryPerformanceCounter` as a high resolution timer.

5 Building Applications

5.1 Unix

Botan usually links in several different system libraries (such as `librt` and `libz`), depending on which modules are configured at compile time. In many environments, particularly ones using static libraries, an application has to link against the same libraries as Botan for the linking step to succeed. But how does it figure out what libraries it *is* linked against?

The answer is to ask the `botan-config` script. This basically solves the same problem all the other `*-config` scripts solve, and in basically the same manner. At some point in the future, a transition to `pkg-config` will be made (as it's less work, and has more features), but right now it doesn't exist on most Unix systems, while a plain Bourne shell script will run fine on anything.

There are 4 options:

`--prefix[=DIR]`: If no argument, print the prefix where Botan is installed (such as `/opt` or `/usr/local`). If an argument is specified, other options given with the same command will execute as if Botan as actually installed at `DIR` and not where it really is; or at least where `botan-config` thinks it really is. I should mention that it

`--version`: Print the Botan version number.

`--cflags`: Print options that should be passed to the compiler whenever a C++ file is compiled. Typically this is used for setting include paths.

`--libs`: Print options for which libraries to link to (this includes `-lbotan`).

Your `Makefile` can run `botan-config` and get the options necessary for getting your application to compile and link, regardless of whatever crazy libraries Botan might be linked against.

5.2 MS Windows

No special help exists for building applications on Windows. However, given that typically Windows software is distributed as binaries, this is less of a problem - only the developer needs to worry about it. As long as they can remember where they installed Botan, they just have to set the appropriate flags in their `Makefile/project` file.