

GNU Status Reports: January 2011

maintainers@gnu.org

GNU Status Reports: January 2011

Copyright © 2011 Free Software Foundation, Inc.

Copying and distribution of this file, with or without modification, are permitted in any medium without royalty provided the copyright notice and this notice are preserved.

Table of Contents

GNU Status Reports: January 2011	1
Binutils (http://www.gnu.org/software/binutils)	1
Coreutils (http://www.gnu.org/software/coreutils)	2
Electric (http://www.gnu.org/software/electric)	3
Freefont (http://www.gnu.org/software/freefont)	3
Gawk (http://www.gnu.org/software/gawk).....	3
GCC (http://www.gnu.org/software/gcc).....	4
GNUCOMM (http://www.gnu.org/software/gnucomm)	5
GNUtrition (http://www.gnu.org/software/gnutrition)...	5
Grep (http://www.gnu.org/software/grep).....	5
GSASL (http://www.gnu.org/software/gsas1).....	6
GSEGrafix (http://www.gnu.org/software/gsegrafix)	6
GSL (http://www.gnu.org/software/gsl)	7
GSRC (http://www.gnu.org/software/gsrc).....	7
Guile (http://www.gnu.org/software/guile).....	8

Hello (http://www.gnu.org/software/hello)	8
Hurd (http://www.gnu.org/software/hurd)	9
Texinfo (http://www.gnu.org/software/texinfo)	10
XBoard (http://www.gnu.org/software/xboard)	11

GNU Status Reports: January 2011

This document collects status reports for GNU packages.

This is a revival of the GNU Status Reports from the historical *GNU's Bulletins*, <http://www.gnu.org/bulletins>. The goal is to provide GNU-wide news and information from time to time, from as many packages as possible.

<http://www.gnu.org/manual> lists all GNU packages, with links to online manuals and home pages. All GNU packages can be accessed on the web via <http://www.gnu.org/software/pkgname>, as shown in the headings here. This report includes items for only a few of the hundreds of GNU packages; we hope more will be represented in future installments.

See also <http://planet.gnu.org>, which aggregates the individual GNU news items posted on Savannah (<http://savannah.gnu.org>).

See also the info-gnu@gnu.org mailing list, where announcements of new GNU releases are posted. Subscribe or view the archives at <http://lists.gnu.org/mailman/listinfo/info-gnu>.

The aim of the present report is to be somewhat higher-level and more general than the others, although there is inevitably some overlap.

Questions, comments, and suggestions about this document in general are welcome; please email maintainers@gnu.org. Issues for specific packages should of course be addressed via their usual routes.

Binutils (<http://www.gnu.org/software/binutils>)

From Nick Clifton: The GNU Binary Utilities project (**binutils** to its friends) is continuing active development. Version 2.21 has just been released containing many bug fixes and some new features:

- Support has been added for the TMS320C6000 (TI C6X) processor family and Renesas RX processor families.
- Support has been added for the Alpha VMS target.
- The linker option ‘`--add-needed`’ has been renamed to ‘`--copy-dt-needed-entries`’ in order to avoid confusion with the ‘`--as-needed`’ option.
- The linker’s ‘`.def`’ file syntax has been extended by allowing the syntax `== id` for imports and exports. This allows aliases of the import/export table names to be written in PE image.
- The `objcopy` program now supports a ‘`--interleave-width`’ option to allow copying a range of bytes from the input to the output with the ‘`--interleave`’ option.
- The `readelf` tool can now display ARM unwind tables using the ‘`--unwind`’ option. It can also display dynamic symbol tables with the ‘`--dyn-syms`’ option.
- A new tool `elfedit` has been added to directly manipulate ELF format binaries.
- The `addr2line` tool now accepts an ‘`--addresses`’ option to display the address before function name or source filename. It also accepts ‘`--pretty-print`’ to have a more human readable output.
- `gas` now understands an extended syntax in the `.section` directive flags for COFF targets that allows the section’s alignment to be specified.

- `gas` also has a new command line option ‘`--compress-debug-sections`’, which requests compression of DWARF debug information sections in the relocatable output file. Compressed debug sections are supported by `readelf`, `objdump`, and `gold`, but not currently by GNU `ld`.
- The `gold` linker now supports `MEMORY` regions in linker scripts.

Meanwhile work continues on `ld` to add support for plugins and compressed debug sections. `gold` is improving and soon will be able to replace `ld` for linking important projects like the Linux kernel.

The next release of the binutils is expected in the middle of 2011.

Coreutils (<http://www.gnu.org/software/coreutils>)

From Jim Meyering and Pádraig Brady: We’ve just released ‘`coreutils-8.9`’, which fixes a number of relatively important bugs (http://savannah.gnu.org/forum/forum.php?forum_id=6662 for 8.8, http://savannah.gnu.org/forum/forum.php?forum_id=6679 for 8.9). Most were in `sort` and were introduced via the parallelization support that was added in ‘`coreutils-8.6`’.

The only significant non-bug-fix change was to add a useful set of features to `split` that lets you split input into n roughly-equal pieces, with options to split on line boundaries or not, and, when honoring line boundaries, to distribute lines in a round-robin fashion or not. See the manual via ‘`info split`’ for details and examples, or online at http://www.gnu.org/software/coreutils/manual/html_node/split-invocation.html.

Some of the hardest work on `coreutils` is knowing what to reject and providing appropriate justification to the contributors. The contributions below all came in recently, and while good ideas, they were not included for various reasons detailed on the mailing list.

`test -ed` Using `stat` in a shell function was deemed sufficient.

`touch -R` ‘`find . -exec touch -am {} +`’ is more general.

`df autoscale`
`df -h` was thought good enough.

`df/ls --blocksize={decimal,binary}`
 Though more correct, it was deemed overkill.

`sha1sum --raw | base64`.
 ‘`openssl dgst -sha1 -binary $file | openssl enc -base64`’ was deemed available enough.

`command --examples`
 This would need to be accepted into the GNU Coding Standards first.

On the upcoming front, we should soon resume work on getting the FIEMAP (efficient sparse copying) code into the master development branch <http://git.savannah.gnu.org/cgi/coreutils.git/log/?h=fiemap-copy>.

Electric (<http://www.gnu.org/software/electric>)

From Stuart Rubin: The Electric VLSI Design System is a complete CAD package for integrated-circuit design. It lets you draw schematics, draw IC layout, and even use textual hardware-description-languages. It has many tools for analysis and synthesis, and can interface to many more.

Electric has been free software from the Free Software Foundation for 12 years now and continues to be developed and distributed. Currently, the Electric development team is employed by Oracle, and they have just released version 9.00.

Freefont (<http://www.gnu.org/software/freefont>)

From Steve White: In GNU FreeFont release 20100919, several language scripts were added:

- Cherokee, Osmanya, Unified Canadian Aboriginal Syllabics, Tifinagh, Armenian (serif), Vai, Kayah Li.
- Coptic, Glagolitic.
- Buginese and Persian were patched.
- Many characters were added to old Cyrillic ranges.

In addition, Sinhala was replaced by a version taken from the T_EX font `sinh` by Yannis Haralambous.

Several of the symbol ranges were expanded, and a few glyphs were corrected.

Several technical improvements were made to font tables, some of which should improve display.

For more information, see https://savannah.gnu.org/forum/forum.php?forum_id=6518.

Gawk (<http://www.gnu.org/software/gawk>)

From Arnold Robbins: Gawk has been seeing lots of progress recently. A major new release is in the works and should be ready within a few more months. New stuff includes:

- New internals, providing along with them a debugger!
- Built-in file inclusion
- Indirect function calls
- BEGINFILE and ENDFILE patterns
- FPAT variable for content-based field splitting
- IPv6 support
- Arrays of arrays
- *Lots* of code cleanup and doc clean up too
- Other features and changes, see NEWS in the git repo
- Moved from CVS to GIT

GCC (<http://www.gnu.org/software/gcc>)

From Gerald Pfeifer: The last couple of months have seen several GCC (GNU Compiler Collection) releases, first GCC 4.4.5 in October, then GCC 4.5.2 in December and we are working full steam to brush up the first release in the 4.6 series (<http://gcc.gnu.org/gcc-4.6/changes.html>).

GCC 4.6 will bring support for the Xilinx MicroBlaze softcore processor on the backend and a new frontend (and run-time library) for the Go language, plus many new Objective-C features, such as declared and synthesized properties, dot syntax, fast enumeration, optional protocol methods, method/protocol/class attributes, class extensions, and a new GNU Objective-C runtime API.

GCC 4.6 is also adding a new optimization level `-Ofast` that is `-O3` with options that can affect standards compliance but result in better optimized code. The link-time optimizations (LTO) we introduced with GCC 4.5 have been improved, stabilized and reduced in terms of resource consumption which allows GCC to build itself or Mozilla Firefox with LTO. As a user you can control how many compilations to execute via `-flto=N`. Inlining heuristics and interprocedural optimizations have been tuned heavily, too.

Back to frontends, stack checking for Ada has been tightened and now detects all overflows on several platforms; experimental support for the upcoming C1X standard has been added, new function attributes `leaf` and `callee_pop_aggregate`, and a new warning `-Wdouble-promote`. Support for the upcoming C++0x standard has seen a lot of love and care and C++ code generation can be tuned more finely (cf. new options `-fstrict-enums` and `-fnothrow-opt`).

On systems supporting the new libquadmath library, GNU Fortran now also supports a quad-precision, `kind=16` floating-point data type (`REAL(16)`, `COMPLEX(16)`). The `-fwhole-file` command-line option is now enabled by default which improves code generation and diagnostics, and both Fortran 2003 and Fortran 2008 support have seen their fair dose of enhancements.

In terms of platform enhancements, `-march=core2` (and `-mtune=core2`) provides support for Intel Core 2, `-march=corei7` for Intel i3/i5/i7 and `-march=btver1` gives you support for AMD Bobcat (family 14). Darwin, FreeBSD, MinGW and Cygwin now all support `__float128` on 32-bit x86 targets. MIPS now supports the Loongson 3A processor and s390x added full support for the zEnterprise z196 processor with a number of new instructions facilities.

As far as operating systems go, GCC provides a convenient way of building native libraries and applications for the Android platform and on Windows `#pragma push_macro` & `#pragma pop_macro`, decimal floating point and more have been added.

Should you find any issue testing this forthcoming release, or any other, please make use of our updated Bugzilla at <http://gcc.gnu.org/bugzilla>.

GNUCOMM

(<http://www.gnu.org/software/gnucomm>)

From David Sugar: GNUCOMM is the umbrella name for the GNU Telephony system. It includes several packages: Bayonne, Sipwitch, and supporting libraries ccAudio, ccRTP, ccScript, and CommonCPP.

The ZRTP communication protocol is spreading as a free (as in freedom) means to communicate privately with GNU GPL licensed software. We are close to releasing ZRTP for Android (`libcppzrtp` already builds on the Android NDK), along with a complete Android client built off of CSipSimple for Android.

We of course continue to maintain full compatibility and interoperability with the emerging IETF ZRTP draft standards, and with ZRTP for Android will now cover even more users. Previously we saw GNU ZRTP4J, which is used in SIP Communicator (a cross-platform Java client for GNU/Linux, OS/X, Windows, and others), as well as Twinkle (a QT ZRTP agent), so ZRTP for Android will extend this reach further.

In GNU uCommon, there is also greater emphasis on secure communication and cryptographic practices in application development, using a new secure library which was introduced earlier this year. It also acts as a generic wrapper for either GNUTLS or OpenSSL (with a configure preference for the former). This is still evolving. It fits into our overall effort in developing secure communications and promoting privacy in application development by default by making it easier to do so, an overall goal for GNU Telephony.

GNUtrition

(<http://www.gnu.org/software/gnutrition>)

From Jason Self: I'm working to rewrite GNUtrition in C with SQLite instead of depending on an external SQL server, a modern UI, and some extra features. I plan to unveil the new version at LibrePlanet in March.

The food database is being upgraded. The current version of GNUtrition uses a very old version (from 1999) so it's time to upgrade. This means more information on more foods—2,000 or so new food items.

Grep (<http://www.gnu.org/software/grep>)

From Jim Meyering: Grep saw several maintenance releases in 2010. There have been a few minor bug fixes since the last release, and there's at least one pending, unfixed bug that is triggered when performing a case-insensitive search that finds a match on a line containing a Turkish upper-case I (<http://thread.gmane.org/gmane.comp.gnu.grep.bugs/3413/focus=3417>). Once that is fixed, I expect to make another release.

GSASL (<http://www.gnu.org/software/gsas1>)

From Simon Josefsson: GSASL is the GNU Simple Authentication and Security Layer library. During December 2010 a new stable release (1.6.0) was published, which includes support for the two most recent SASL mechanisms: SCRAM and GS2. These are the new generation of SASL mechanisms, intended to replace (over time) the older CRAM-MD5, DIGEST-MD5 and GSS-API mechanisms.

The development series for the next stable branch has been started already, and plans include improved GS2 support for non-Kerberos mechanisms, support for the two SAML mechanisms that are discussed in the IETF, and more. Now is a good time to join the mailing list and contribute to the project!

GSEGrafix

(<http://www.gnu.org/software/gsegrafix>)

From Spencer Buckner: GSEGrafix is a GNOME application which uses an anti-aliased GNOME canvas for creating scientific and engineering plots. The program is written in C and reads ASCII parameter files and data files. The parameter files contain keywords and corresponding arguments for specifying plot parameters (such as data file names, data file formats, plot type, plot style, axis type, axis labels, and more). Eleven example plots, corresponding examples of Octave code or C code for creating the data files, and corresponding parameter files are included. The program can be run from a terminal window or from a graphical user interface.

The current version is `gsegrafix-1.0.5`. Version 1.0.5 added data files for Examples 1 and 7, for use in becoming familiar with the program. There are links to each of these files in the corresponding example descriptions. The program includes an HTML help file, and an XML help file was added in this version.

Version 1.0.6 is under development and includes two additional keywords: `background_color` and `background_image`. The former enables the background color of the plot window to be specified as either white, the default, or black. If black is chosen, the plot box, tick marks, axis labels, title, and text are white. The keyword `background_image` enables a background image, such as a map, to be displayed in the plot box. The image can be scaled in four different ways by specifying one of the parameter values: `center`, `fill`, `scale`, or `zoom`.

GSL (<http://www.gnu.org/software/gsl>)

From Brian Gough: The GNU Scientific Library is a numerical library written in C. It provides a large collection of mathematical routines such as random number generators, special functions, fourier transforms and least-squares fitting. There are over 1000 routines in total with an extensive test suite.

The role of the project is to provide an alternative to the well-known proprietary numerical libraries, which hinder scientific cooperation by preventing the sharing of scientific software.

The project has been in stable release for 10 years, and has maintained API stability throughout that time. Ongoing work continues on bug-fixing, particularly improving the stability of special functions, and a 1.14 maintenance release was made earlier this year.

Currently the project does not have sufficient contributors to tackle major developments, and is limited to bug fixing. A good knowledge of numerical methods is needed to work on the project—anyone interested in helping should contact bug-gsl@gnu.org for more information.

GSRC (<http://www.gnu.org/software/gsrc>)

From Brian Gough: The GNU Source Release Collection (GSRC) is a build system for GNU packages. The goal is to make it easy to install the latest GNU releases from source in your home directory on an existing distribution.

GSRC simplifies this process by automatically downloading and building packages and their dependencies. You only need a single `./configure && make` to get everything, as usual.

For those familiar with BSD systems, this is similar to BSD Ports. GSRC uses GAR, a build system written in GNU Make by Nick Moffit, and GARstow by Adam Sampson. GSRC is also similar to the GARNOME system for GNOME.

There are currently 136 GNU packages in GSRC which build and install successfully, including GCC, Guile, GnuPG, Emacs, coreutils, GnuTLS and many others. Packages are updated several times a week, as new releases come up.

The whole collection is also built regularly on the GCC compile farm and a weekly report posted on <http://planet.gnu.org> with the current build status of all GNU packages (<http://chapters.gnu.org/~bjg/gsrc/summary/>), you can also check the complete build logs there.

Goals for the project are to support the remaining GNU packages—particularly those with complex dependencies, such as Gnash and GnuCash, and to also make the latest alpha releases available as well (this has already been done for a few projects). Additional help from people familiar with those packages is welcome.

Guile (<http://www.gnu.org/software/guile>)

From Ludovic Courtès: GNU Guile is an implementation of the Scheme programming language that can be embedded in applications to allow users to extend them. Extensibility empowers users to leverage software freedom by making it easier to start making changes. Just as Emacs Lisp allowed unanticipated applications to be written within the Emacs environment, Guile should do the same for other GNU applications. It should spread the Emacs nature to the rest of GNU.

Guile has undergone major changes over the last couple of years, resulting in 15 pre-releases—labeled 1.9.x. This will lead to the 2.0 stable series, scheduled for release in the first quarter of 2011. The 1.8 stable series was released in 2006. It offered many improvements such as support for exact rational numbers, multi-threading, and a better C programming interface.

Guile 2.0 is a breakthrough in Guile’s history. First and foremost, it is based on a compiler and a virtual machine, and comes with a powerful read-eval-print loop (REPL) and debugger. The compiler compiles Scheme code to bytecode, applying well-known optimizations. As a result, Scheme code runs noticeably faster with Guile 2.0. Compilation can occur transparently: when the compiled form of a module is not found in cache, it is automatically compiled before being run.

In addition, Guile 2.0 has a wealth of exciting new features. Hygienic macros are supported natively and well integrated with the module system—they are what makes Scheme itself extensible. It supports Unicode, as well as most of R6RS, the latest revision of the Scheme programming language.

A dynamic foreign function interface (FFI) is provided, allowing bindings to C libraries to be written entirely in Scheme. Also new in 2.0 is support for composable continuations (a.k.a. “prompts”), a novel control flow mechanism. New modules include web tools for HTTP and URIs, XML/SXML, pattern matching, LALR parsing, purely functional hash lists, and code coverage, to name a few.

While Guile 2.0 remains primarily an implementation of Scheme, its compiler comes with a front-end for ECMAScript—one more choice for users seeking to extend their applications. An Emacs Lisp front-end is also available. It is the result of work towards the unification of Guile and Emacs, the flagships of extensibility in GNU.

2011 promises to be another exciting year for Guile. We believe Guile 2.0 is a great tool to push extensibility further and enhance software freedom!

Hello (<http://www.gnu.org/software/hello>)

From Karl Berry: GNU Hello remains a complete example package, attempting to exemplify some best practices for GNU packages. For example, it uses Automake for configuration, Gettext for internationalization, Help2man to avoid maintaining separate man pages, and Gnulib to share common sources.

For examples in other programming languages, please see Gettext. The Automake manual also contains several examples worked from scratch.

Hurd (<http://www.gnu.org/software/hurd>)

From Olaf Buddenhagen, Arne Babenhausner, Thomas Schwinge: Yeah, that's quite right: this project is still alive!

According to our mission statement, the goal is creating a *general-purpose kernel suitable for the GNU operating system, which is viable for everyday use, and gives users and programs as much control over their computing environment as possible*. It has a unique multi-server microkernel-based architecture—bringing advanced operating system research to the mainstream. More concretely, it's a collection of user-space server processes that run on the GNU Mach microkernel.

The Hurd doesn't fully deliver on the *everyday usability* goal yet, but it is seeing continuous improvement—and 2010 has been no exception. Let's take a look at the progress throughout the year.

- Apart from having done a lot of other work, Samuel Thibault, our Jack of all trades, merged his development branch that added Xen domU support to GNU Mach, which makes it possible to run a GNU/Hurd system as a Xen guest. Development of this started in 2007, and since then it has been heavily tested by using it for the Debian GNU/Hurd build servers, most of our public GNU/Hurd systems, http://www.gnu.org/software/hurd/public_hurd_boxen.html, and the Hurd project's wiki web server.
- We had Zheng Da work on a new hardware device driver framework, which is based on the Dresden L4 (Fiasco) group's DDE project, and allows running modern Linux kernel drivers as user-space server processes. Many network cards already work perfectly with this new framework. (It has not yet been integrated into the mainstream Hurd code base, so it needs to be compiled and set up by hand.) Other driver classes, such as hard disk controllers, will require further work.
- As in the previous years, we again participated in the Google Summer of Code 2010. Olaf Buddenhagen is our main guy for organizing this.

Jérémie Koenig ported the modern Debian Installer to Debian GNU/Hurd. Installation images using the new installer are replacing the previous CD images, which were using an installer based on the old Debian boot floppies (and running under the Linux kernel)—Philip Charles has been maintaining these single-handedly for almost ten years! The new installer images are available from <http://people.debian.org/~sthibault/hurd-i386/installer/cdimage/>.

Emilio Pozuelo Monfort was investigating specific compatibility problems exposed by the extensive test suites coming with some software packages. Emilio's analysis uncovered a number of programming errors in the Hurd code, and he fixed several of them. As these typically affected other programs too, this improved stability and compatibility in general.

- Jérémie Koenig created a new implementation of a `procfs` translator, which is considerably more robust and efficient than the previous one. Tools such as `top` can now be used without problems.

Some other translators (`gopherfs`, `netio`, `tarfs`) which have been created by external contributors in the past have been fixed up by Manuel Menal, and packaged in Debian. Thus, some of the results of Hurd's extensible architecture are now easier to access, and these updated translators can serve as examples for other developers to implement their own ideas.

- In addition to various general stability, compatibility, and portability fixes, several people (Samuel Thibault, Pino Toscano, Emilio Pozuelo Monfort, and others) have been working on fixing issues with specific Debian packages. So far, about 68% of all Debian packages are also available for Debian GNU/Hurd.
- Michael Walker started the Arch Hurd distribution, and together with other enthusiastic Arch developers (Allan McRae, Matthias Lanzinger, Alexander Preisinger, Stephen Gilles, Diego Nieto Cid) they got it working in an amazingly short amount of time, both as an installable system, and a live CD. So now there is a choice between two well-featured distributions for the Hurd. These new people of course also help forwarding Hurd development in general—Diego in particular contributed various patches to the Hurd console and other components.
- Carl Fredrik Hammar finished and presented his thesis, “Generalizing mobility for the Hurd”, <http://lists.gnu.org/archive/html/bug-hurd/2010-01/msg00078.html>, and passed with distinction.

This is a very short digest of what happened in the last year. You can read our regular “Month of the Hurd” at <http://www.gnu.org/software/hurd/news.html>, or by subscribing to our RSS feed at <http://www.gnu.org/software/hurd/index.rss>.

If you are interested, for example, in doing a university project on a multi-server microkernel-based operating system, or if you are interested in contributing to Hurd development in general, please see <http://www.gnu.org/software/hurd/contributing.html>. Or just talk to us at bug-hurd@gnu.org or the #hurd IRC channel on freenode.

Texinfo (<http://www.gnu.org/software/texinfo>)

From Karl Berry: We hope the next major release of Texinfo will be made in 2011, after several years of development. Thanks to my colleague Patrice Dumas, the Makeinfo program has been completely rewritten in Perl while retaining essentially complete compatibility. The new program is based on and is as compatible as possible with `texi2html`. It will provide vastly better support for HTML customization, Unicode, additional back-ends, and so on. If you want to try it now, it’s in the development sources.

‘`texinfo.tex`’, as a separable file, has been updated separately on an ongoing basis. The current version is available from <http://ftpmirror.gnu.org/texinfo/texinfo.tex> (and from Gnulib), and we encourage you to always get the current version when making a release.

XBoard (<http://www.gnu.org/software/xboard>)

From Arun Persaud: XBoard is a graphical user interface to chess in all its major forms (and many others). Over the last year XBoard development has seen three new releases fixing lots of bugs and including many new features with another major release just around the corner. We now support many more variations of chess (xiangqi, shogi, makruk, as well as better support for variations like crazyhouse, bughouse and many more). Also, the support for computer chess has been improved.

An ongoing effort is to merge the code of XBoard and its so-called “Winboard” companion back into one code base and adapt it to use GTK for the GUI front end. If you are interested in this and are willing to help, feel free to contact us at xboard-devel@gnu.org. We can always use more help! You can of course also contact us to discuss other issues/ideas too.