



- COLUMNS
- COMMUNITY POSTS
- ISSUES
- BOOKS
- FORUM
- FS NEWS

Ads by Google

Source Code Analysis Tool

Covers 100% of the execution path to resolve defects. Get free trial. Coverity.com

Free programming magazine

Claim a free subscription to VSJ The UK's leading developer title www.vsj.co.uk

NI Graphical Programming

Powerful, Intuitive Programming w/ LabVIEW- Online Tutorials & Evals. www.ni.com

Debug Trace Monitor

Collect diagnostics output from C#, VB.NET, VB6, VBA, Script, COM, ... www.devtracer.com

Jobs at Sophos

Career opportunities at Sophos in IT, Software Development & Testing www.sophos.com/careers

- Tracking Worms, Spam and Malware ...
- How to Restore Lost Bookmarks in ...
- Google AdSense for RSS Feeds Coming ...

Sponsored by Novell.

Best voted contents

- Open letter to standards professionals, developers, and activists
  - Pieter Hintjens, 2008-05-13
- The 2008 Google Summer of Code: 21 Projects I'm Excited About
  - Andrew Min, 2008-05-13
- The Bizarre Cathedral - 6

Home

# Autotools: a practitioner's guide to Autoconf, Automake and Libtool



by John Calcote

There are few people who would deny that Autoconf, Automake and Libtool have revolutionized the free software world. While there are many thousands of Autotools advocates, some developers absolutely *hate* the Autotools, with a passion. Why? Let me try to explain with an analogy.

In the early 1990's I was working on the final stages of my bachelor's degree in computer science at Brigham Young University. I took a 400-level computer graphics class, wherein I was introduced to C++, and the object-oriented programming paradigm. For the next 5 years, I had a love-hate relationship with C++. I was a pretty good C coder by that time, and I thought I could easily pick up C++, as close in syntax as it was to C. How wrong I was. I fought late into the night, more often than I'd care to recall, with the C++ compiler over performance issues.

The problem was that the most fundamental differences between C++ and C are not obvious to the casual observer. Most of these differences are buried deep within the C++ language specification, rather than on the surface, in the language syntax. The C++ compiler generates code beneath the covers at a level never even conceived of by C compiler writers. This level of code generation provides functionality in a few lines of C++ code that requires dozens of lines of C code. Oh, yes—you can write object-oriented software in C. But you are required to manage all of the details yourself. In C++, these details are taken care of for you by the compiler. The advantages should be clear.

But this high-level functionality comes at a price—you have to learn to understand what the compiler is doing for you, so you can write your code in a way that complements it. Not surprisingly, often the most intuitive thing to do in this situation for the new C++ programmer is to inadvertently write code that works against the underlying infrastructure generated by the compiler.

And therein lies the problem. Just as there were many programmers then (I won't call them software engineers—that title comes with experience, not from a college degree) complaining of the nightmare that was C++, so likewise there are many programmers today complaining of the nightmare that is the Autotools. The differences between make and

Static Code Analysis Tool

Deliver bug-free source code. Download whitepaper from Coverity. Coverity.com

UK Software Developers

Experienced software developers ASP.NET, SQL, C#, Expertise www.pbcsystems.co.uk

Process Mapping Software

MS Visio Process Library solutions with web-enabled keyword search www.triaster.co.uk

UK Scrum Training

Certified Scrum Master training throughout the UK and Europe www.agilebear.com

Ads by Google



From the FSM staff...

The Top 10 Everything (Dave). The good, the bad and the ugly.

Free Software news (Dave & Bridget). All about free software -- free as in freedom!

Book Reviews: Illiterarty (Bridget). Book reviews, blogs, and short stories.

Hot topics - last 60 days

Installing an all-in-one printer device in Debian Ryan Cartwright, 2008-05-05

Things you miss with GNU/Linux Ryan Cartwright, 2008-05-01

Why Microsoft should not



Ryan Cartwright, 2008-05-25

o Digital Rights

Management (DRM): is it in its death throes?

Gary Richmond, 2008-05-07

Automake are very similar to the differences between C and C++. The most basic single-line Makefile.am generates a Makefile.in file (an Autoconf template) containing nearly 350 lines of make script.

### Who should read this book

This book is written for the open source software package maintainer. I'm purposely not using the terms "free software" or "proprietary software that's free". The use of the term "open source" is critical in this context. You see, open source defines a type of software distribution channel. One in which the primary method of obtaining software functionality is downloading a source archive, unpacking, building and installing the built products on your system. Free software may be published in binary form. Proprietary software may be given away. But open source software implies source-level distribution.

Source-level distribution relegates a particular portion of the responsibility of software development to the end-user that has traditionally been assumed by the software developer. But end-users are not developers, so most of them won't know how to properly build your package. What to do, what to do... The most widely adopted approach from the earliest days of the open source movement was to make the package build process as simple as possible for the end user, such that she could perform a few well-known steps and have your package cleanly installed on her system.

Most packages are built using makefiles, and the make utility is as pervasive a tool as anything else that's available. It's very easy to type make—but that's not the problem. The problem crops up when the package doesn't build successfully, because of some unanticipated difference between the user's system and the developer's system.

Thus was born the ubiquitous configure script—initially a simple shell script that configured the end-user's environment so that the make utility could successfully build a source package on the end-user's system. Hand-coded configure scripts helped, but they weren't the final answer. They fixed about 65 percent of the problems resulting from system configuration differences—and they were a pain in the neck to write properly. Dozens of changes were made incrementally over a period of years, until the script would work properly on most systems anyone cared about. But the entire process was clearly in need of an upgrade.

Do you have any idea of the number of build-breaking differences there are between existing systems today? Neither do I, but there is a handful of developers in the world who know a large percentage of these differences. Between them and the free software community, the Autotools were born. The Autotools were designed to create configure scripts and makefiles that work correctly and provide significant chunks of valuable end-user functionality under most circumstances, and on most systems—even systems not initially considered (or even known about) by the package maintainer.

### Buzz authors

Marco Marongiu

Pieter Hintjens

drascus321

dogboi

Richard Rothwell

### All news

மஸ்ரீ ராமதாஸ்: Party time with Ubuntu...

Sridhar Dhanapalan: Open CeBIT

Martin Meredith: Welcome to the Family, Synergy

Stephen Stalcup: Next Membership meeting for the Americas

Andres Rodriguez: Hello Planet Ubuntu!!

more



- 14 20 Essential KDE Applications - Review
- 14 10 Reasons to Love Debian
- 16 Burn Your CDs and DVDs! K3b Review
- 17 Novell's Moonlight: Crippled and Defective by Design™
- 18 India Appeals Against OOXML, Joining Brazil, South Africa, Maybe More
- 17 Wiping your disk drive clean
- 20 Dear Google: Is AGPL Evil?
- 19 Slashdot: World's Third-Largest Population Appeals ISO's Decision

Get this widget »

### Why Microsoft should not lose (and free software will still win)

Ryan Cartwright, 2008-04-21

### Drigg (the pligg alternative) vs. Pligg: why should people switch?

Tony Mobily, 2008-04-13

### Beyond Synaptic - using apt for better package management

Ryan Cartwright, 2008-04-03

### Hot topics - last 21 days

#### Dubious ads in Free Software Magazine

Tony Mobily, 2008-05-25

#### The Bizarre Cathedral - 6

Ryan Cartwright, 2008-05-25

#### Open letter to standards professionals, developers, and activists

Pieter Hintjens, 2008-05-13

#### The Bizarre Cathedral - 4

Ryan Cartwright, 2008-05-11



Odiogo



So, returning to that passionate hate felt by some developers toward the Autotools: If you get your head screwed on straight about the primary purpose of the Autotools, then hate quickly turns into respect—and even appreciation. Often the root of such hate is a simple misunderstanding of the rationale behind the Autotools. The purpose of the Autotools is not to make life simpler for the package maintainer (although it really does in the long run). *The purpose of the Autotools is to make life simpler for the end-user.*

To drive my point home, I'll wager that you'll never see a Linux distribution packager spouting hateful sentiment on the Autotools mailing lists. These people are in a class of engineers by themselves. They're generally quiet on mailing lists—asking an occasional well-considered question when they really need to—but lurking and learning, for the most part. Packagers grasp the advantages of the Autotools immediately. They embrace them by studying them until they know them like an expert C++ programmer knows his compiler. They don't write many Autoconf input scripts, but they do patch a lot of them.

How do you become such an expert? I recommend you start with this book. I've organized it in the best way I know how to help you get your head around the functionality provided by the Autotools. But don't stop there. Pick up the manuals. They're free, and links are provided in the References section of this book, but they're easy to find with a simple internet query. I've left a LOT of details out of this book, because my purpose is to quickly get you up to speed on understanding and using the Autotools. The Autotools manuals are well-written and concise, but more importantly, they're complete. After reading this book, they should be a cake walk.

Then study open source and free software packages that use the Autotools. See what other experts have done. Learning by example is an excellent way to begin to retain the information you've read. Finally, instrument some of your own projects with the Autotools. Doing is by far the best way to learn. The initial reading will reduce the frustration of this exercise to something bearable.

Above all, remember why you're doing this—because you want your end-user's experience with your package to be as delightful as possible. No open source project was ever successful until it had a well-established user base, and you don't get there by alienating your users. You do it by creating a user build, installation and operation experience that shines. You'll still need to handle the operation experience, of course, but Autotools can provide a *great* multi-platform build and installation experience—with far less effort on your part.

### **The book that was never to be**

I've wondered often during the last eight years how strange it is that the *only* third-party book on Autotools that I've been able to discover is the New Rider's 2000 publication of GNU AUTOCONF, AUTOMAKE and LIBTOOL, affectionately known in the community as "The Goat Book".

I've been in this industry for 25 years, and I've worked with free software for quite some time now. I've learned a lot about free software maintenance and development—most of it, unfortunately, by trial and error. Had there *been* other books on the topic, I would have snatched them all up immediately, rather than spend hours—even days sometimes—trying to get the Autotools to do something I could have done in a makefile in a few minutes.

I've been told by publishers that there is simply no market for such a book. In fact, one editor told me that he himself had tried unsuccessfully to entice authors to write this book a few years ago. His authors wouldn't finish the project, and the publisher's market analysis indicated that there was very little interest in the book.

No interest?! Let's analyze this picture: There are nearly 200,000 free software projects on sourceforge.net alone. If only 10 percent of those are still active, that's still 20,000 live projects. If 80 percent of those are Linux or Unix based packages, that's 16,000 free software packages that might use the Autotools. And that's only sourceforge.net. Each of those packages has at least one maintainer—often two or three. Each of those maintainers *probably* uses (or has tried to use) the Autotools. Many of them have a fairly solid understanding of the Autotools by now, but at what expense in time and effort did they gain this understanding?

Publishers believe that free software developers tend to disdain written documentation—perhaps they're right. Interestingly, books on Perl sell like Perl's going out of style—which is actually somewhat true these days—and yet people are still buying enough Perl books to keep their publishers happy. All of this explains why there are ten books on the shelf with animal pictures on the cover for perl, but literally nothing for free software developers.

The authors of the Goat Book, Gary Vaughan, Ben Elliston, Tom Tromey and Ian Lance Taylor, are well known in the industry, to say the least—indeed, they're probably the best people I know of to write a book on the Autotools. But, as fast as free software moves these days, a book published in 2000 might as well have been published in 1980. Nevertheless, because of the need for *any* book on this subject, [the Goat Book](#) is still being sold new in bookstores. In fairness to the authors, they have maintained an [online version](#) through February of 2006.

The biggest gripe I have with the Goat Book is the same gripe I have with the GNU manuals themselves. I'm talking about the sheer volume of information that is assumed to be understood by the reader. The Goat Book is written in a very non-linear fashion, so it's difficult to learn anything from it. It's a great reference, but a terrible tutorial. Perhaps the authors were targeting an audience that had already graduated to more advanced topics. In either case, the Goat Book, while being very complete from a content perspective, is definitely not a great learning resource for the beginner.

And yet a large percentage of their readership today are

young people just starting out with Unix and Linux, and most of their issues center around Unix utilities not generally associated with the Autotools. Take `sed`, for example: What a dream of a tool to work with—I love it! More to the point however, a solid understanding of the *basic* functionality of `sed`, `m4`, shell script and other utilities is critical to understanding the proper use of the Autotools. The Goat Book does cover the `m4` macro processor in great detail, but it's not clear to the uninitiated that one might do well to start with Chapter 21. Understanding how something works under the covers is often a good way to master a topic, but a general introduction at an appropriate point in higher-level discussions can make all the difference to a beginner.

Existing GNU documentation is more often reference material than solution-oriented instruction. What we need is a cookbook-style approach, covering real problems found in real projects. As each recipe is mastered, the reader makes small intuitive leaps—I call them minor epiphanies. Put enough of these under your belt and overall mastery of the Autotools is ultimately inevitable.

Let me give you another analogy: I'd been away from math classes for about three years when I took my first college calculus course. I struggled the entire semester with little progress. I understood the theory, but I had trouble with the homework. I just didn't have the background I needed. So the next semester, I took college algebra and trigonometry as half-semester classes each ("on the block", to use the vernacular). At the end of that semester I tried calculus again. This time I did very well—finishing the class with a solid A grade. What was missing the first time? Just basic math skills. You'd think it wouldn't have made that much difference, but it really does.

The same concept applies to understanding the Autotools. You need a solid understanding of the tools upon which the Autotools are built in order to become proficient with the Autotools themselves. For example, here's a message I came across a few days ago while I was perusing the Autoconf mailing list:

```
>>> If I do this:
>>>
>>> AC_CHECK_FUNC(
>>>   [chokeme],
>>>   [],
>>>   []
>>> )
>>>
>>> It will yield shell code that ends in:
>>>
>>> if
>>> :
>>> else
>>>
>>> fi
>>>
>>> Which produces a configure script that dies
>>> with:
>>> "syntax error near unexpected token `fi'"
```

```
>>>
>>> Is this an autoconf bug, or user error on
>>> my part?

>> The else part is not empty, it consists of
>> explicit whitespace. When collecting arguments
>> only unquoted leading whitespace is skipped by
>> m4, trailing whitespace (quoted or not) is
>> preserved. You need to put the closing paren
>> immediately after the closing quote of the
>> argument.

> Is that something I should always do? I've been
> consistently putting the closing paren on its
> own line. Is that a "never"?
```

```
You can instead use dnl to ignore the trailing
whitespace, provided the closing paren is in
column 1.
```

Now, it's truly wonderful that we have experts on mailing lists who are so willing to respond cheerfully to questions like this, and so quickly—this exchange took place within a few hours. However, without looking, I submit that similar questions have probably been asked dozens of times in the last 5 years. Not because mailing list posters don't read the archives (although I'll admit that they probably don't often do so), but rather because this problem can rear its ugly head in many different ways, none of which look remotely related to each other in the eyes of the uninitiated.

Here are some of the problems with the response to this request: Does the original poster (OP) even know what m4 is? If so, does he realize he's running it when he executes "autoconf" to generate his configure script? Alright, suppose he does; either way, he's clearly not an m4 expert or he wouldn't have needed help with this issue to begin with.

Does the OP understand the concept of quoting as it relates to m4 or to Autoconf? Perhaps he's always simply copied one configure.ac script to another, modifying as little as possible to get it to work with a new project. Given the high-level nature of configure.ac, this is entirely possible (I've done it myself). If so, he may just assume that the square brackets are necessary around each parameter in an Autoconf macro. Given the nature of the question, I'd say the OP believes that the entirety of each parameter is contained within the brackets, so this assumption is not at all improbable.

Another problem is seen in the final response where the OP is told, "...instead use dnl to ignore the trailing whitespace..." If the OP didn't understand m4 whitespace rules, he probably doesn't know about the m4 built-in macro, `dnl`. If that's the case, then this response made no sense to him whatsoever. Even if he did understand what he was to do—perhaps based on having seen `dnl` being used in other configure.ac scripts, apparently as a secondary form of comment delimiter—he probably doesn't understand the full impact or use of this macro. Regardless, you can bet there are other mailing list

readers who experienced far more confusion over this exchange.

This book attempts to alleviate some of the confusion and reduce the existing learning curve by presenting the Autotools in a manner conducive to an open source beginner learning how to use them.

## How this book is organized

Chapter 1 presents a general overview of the packages that are considered part of the GNU Autotools. This chapter describes the interaction between these packages, and the files consumed by and generated by each one. In each case, I've provided a graphic depiction of the flow of data from hand-coded input files, to final output files. Don't worry if you feel overwhelmed after reading Chapter 1. The details will become clear later. I recommend that you give this chapter a quick read to start with, and then come back to it later, after you've read the rest of this book.

Chapter 2 covers free software project structure and organization. This chapter also goes into detail on the GNU coding standards and the Filesystem Hierarchy Standard documents, both of which have played vital roles in the design of the Autotools. It presents some fundamental tenets upon which the design of each of the Autotools is based. With these concepts, you'll be prepared to understand some of the most fundamental rationale behind architectural decisions made by the Autotools developers. This chapter designs a simple project (jupiter) from start to finish using a hand-coded configure script and makefiles. It builds on jupiter in a step-wise fashion, as we begin to discover useful functionality to make our's and our end-users' tasks simpler, relative to the jupiter project. The project is built on principles taken from these two documents. As a side benefit, the GNU manuals for the Autotools should begin to make a lot more sense to you.

Chapters 3, 4 and 5 cover the basic purpose and use of the GNU Autoconf, Automake and Libtool packages, respectively. If you already have a basic familiarity with these packages, you can probably skip these chapters, but please feel free to revisit them if you find yourself in over your head with the remaining chapters.

Chapter 6 takes an existing complex open source project (FLAIM) through the process of converting from a hand-coded build system to an Autotools build system, from start to finish. The example provided by this chapter will use the concepts presented in previous chapters to take it from the original hand-coded makefiles to a complete Autotools project, implementing all of the features provided by the original build system. This process should help you to understand how you might "autoconfiscate" one of your own existing complex projects.

Chapter 7 is where we really begin to break ground on Autoconf. This chapter goes into detail on the m4 macro language and how it's used by Autoconf to generate your configure script from your configure.ac file. A proper grounding

in m4 will leave you very comfortable with a discussion on the internal workings of Autoconf.

Chapter 8 describes the sed utility and how it's used by your Autoconf-generated configure script to generate a Makefile from a Makefile.in template. In addition, I'll show you how you can use the same technique to extend your own Makefiles (and your Automake Makefile.am files).

Chapter 9 covers the process of writing your own Autoconf macros, thereby allowing you to provide other projects with canned versions of your cool solutions.

Chapter 10 discusses the [AC Autoconf Macro Archive](#)—an online repository of reusable Autoconf components (m4 macros) that package maintainers can drop in for various bits of useful functionality. A base library of such components is provided by the archive, wherein each macro is discussed in detail, along with the problem it's designed to solve, and the mechanisms used to solve it.

Finally, the References section includes relevant links to the best material on Autotools available on the internet, including manuals and tutorials.

---

01. [Chapter 1: A brief introduction to the GNU Autotools](#)

---

02. [Chapter 2: Project management and the GNU coding standards](#)

---

03. [Chapter 3: Configuring your project with Autoconf](#)

---

04. [Chapter 4: Automatically writing makefiles with Automake](#)

---

05. [Chapter 5: Building shared libraries with Libtool](#)

---

06. [Chapter 6: An autotools example](#)

---

07. [Chapter 7: The m4 macro processor](#)

---

08. [Chapter 8: You sed what?! Awk com'on!](#)

---

09. [Chapter 9: Reusing your solutions](#)

---

10. [Chapter 10: A catalog of reusable solutions](#)

---

11. [References](#)

---

[Chapter 1: A brief introduction to the GNU Autotools >](#)

Login or register to post comments    12029 reads



#### Chapter 5+

Submitted by [slowfranklin](#) on Tue, 2008-05-20 10:15.

Vote!

This is great stuff!  
But chapters 5+ are missing?

Login or register to post comments



#### They are coming...

Submitted by [admin](#) on Wed, 2008-05-21 01:04.



Vote!

Hi,

The book is not complete yet. the other chapters will come soon :-D

Merc.

[Login](#) or [register](#) to post comments

---



- COLUMNS
- COMMUNITY POSTS
- ISSUES
- BOOKS
- FORUM
- FS NEWS

Ads by Google

**IT Support London & SE**  
 Local, cost effective IT support for your 2003 or SBS network  
[www.1st-solution.co.uk](http://www.1st-solution.co.uk)

**Pump Software**  
 Streamline your pump sales. Select, configure, quote, and order online.  
[www.bigmachines.com](http://www.bigmachines.com)

**Configuration Software**  
 Product & Service Configurator. Configure manufacturing & more.  
[www.Config8or.com/manufacturir](http://www.Config8or.com/manufacturir)

**Why Is My Computer Slow?**  
 Perform Registry Compression, Defrag, Manage & Cleanup Your PC.  
[www.RegClean.com](http://www.RegClean.com)

**Get NetOp Software**  
 Free Trials | Gold Partner for full range of NetOp Products  
[www.bluebeamsecurity.com](http://www.bluebeamsecurity.com)

- [Microsoft May Seek New Yahoo ...](#)
- [Building Semantics is Different ...](#)
- [Lux: multi-touch for OS X](#)

Sponsored by **Novell**.

### Best voted contents

- [The Bizarre Cathedral - 3](#)  
 Ryan Cartwright, 2008-05-05
- [The Bizarre Cathedral - 2](#)  
 Ryan Cartwright, 2008-04-27
- [Indexing offline CD-ROM archives](#)  
 Terry Hancock, 2008-05-03
- [Microsoft and free software? I don't think so...](#)  
 Terry Hancock, 2008-04-26

Home » Autotools: a practitioner's guide to Autoconf, Automake and Libtool

## Chapter 1: A brief introduction to the GNU Autotools



by John Calcote

I'm going to make a rather broad and sweeping statement here: If you're writing free or open source software targeting Unix or Linux systems, then you should be using the GNU Autotools. I'm sure I sound a bit biased, but I'm not. And I shouldn't be, given the number of long nights I've spent working around what appeared to be shortcomings in the Autotools system. Normally, I would have been angry enough to toss the entire project out the window and write a good hand-coded makefile and configure script. But the one fact that I always came back to was that there are literally thousands of projects out there that appear to be very successfully using the Autotools. This was too much for me. My pride would never let me give up.

### Who should use the Autotools?

The Autotools are supposed to make projects simpler for the maintainer, right? And the answer to that question is a definitive "No". Don't misunderstand me here—the Autotools do make your life easier in the long run, but for different reasons than you may first realize. The primary goal of the Autotools is not to make project maintenance simpler, although I honestly believe the system is as simple as it can be, given the functionality it provides. It took me a while to figure this out, and really, it was one of my most significant Autotools epiphanies. Ultimately, I came to understand that the purpose of the Autotools is two-fold: First, to make life easier for your users, and second, to make your project more portable—even to systems on which you've never tested, installed or even built your code.

Well then, what if you don't work on free or open source software? Do you still care about these goals? What if you're writing proprietary software for Unix or Linux systems? Then, I say, you would probably still benefit to some degree from using the Autotools. Even if you only ever intend to target a single distribution of Linux, the Autotools will provide you with a build environment that is flexible enough to allow your project to build successfully on future versions or distributions with virtually no changes to the build scripts. And, let's be honest here—you really *can't* know in advance whether or not your management will want your software to run on other platforms in the future. This fact alone is enough to warrant my statement.

From the FSM staff...

[The Top 10 Everything](#) (Dave). The good, the bad and the ugly.

[Free Software news](#) (Dave & Bridget). A site about short stories and writing.

[Book Reviews: Illiterary](#) (Bridget). Book reviews, blogs, and short stories.

Hot topics - last 60 days

[Installing an all-in-one printer device in Debian](#) Ryan Cartwright, 2008-05-05

[What is the free software community?](#) Tony Mobily, 2008-03-29

[Things you miss with GNU/Linux](#)

## Who should NOT use the Autotools?

About the only scenario where it makes sense NOT to use the Autotools is the one in which you are writing software for non-Unix platforms only—Microsoft Windows comes to mind. Some people will tell you that the Autotools can be used successfully on Windows as well, but my opinion is that the POSIX/FHS approach to software build management is just too alien for Windows development. While it *can* be done, the tradeoffs are way too significant to justify shoe-horning a Windows project into the Autotools build paradigm.

I've watched some project managers develop custom versions of the Autotools which allow the use of all native Windows tools. These projects were maintained by people who spent much of their time tweaking the tools and the build environment to do things it was never intended to do, in a hostile and foreign environment. Quite frankly, Microsoft has some of the best tools on the planet for Windows software development. If I were developing a Windows software package, I'd use Microsoft's tools exclusively. In fact, I often write portable software that targets both Linux and Windows. In these cases, I maintain two separate build environments—one for Windows, and one based on the Autotools for everything else.

The original reasons for using GNU tools to build Windows software were that GNU tools were free, and Microsoft tools were expensive. This reason is no longer valid, as Microsoft makes the better part of their tools available for free download today. This was a smart move on their part—but it took them long enough to see the value in it.

## Your choice of language

One other important factor in the decision to use or not use the Autotools with your project is your language of choice. Let's face it, the Autotools were designed by GNU people to manage GNU projects. There are two factors that determine the importance of a computer language within the GNU community:

1. Are there any GNU packages written in the language?
2. Does the GNU compiler tool set support the language?

Autoconf provides native support for the following languages based on these two criteria:

- C
- C++
- Objective C
- Fortran
- Fortran 77
- Erlang

By “native support”, I mean that Autoconf will compile, link and run source-level feature checks in these languages.

If you want to build a Java package, you can configure

### Buzz authors

Marco Marongiu

Pieter Hintjens

drascus321

dogboi

Richard Rothwell

### All news

John Crawford: Arizona  
LoCo Hardy Release Party

John Crawford: Ubuntu  
Weekly Newsletter #91

Matthew Helmke: Morocco  
blocks Google Maps

Andy Price: Synchronicity  
and Sharing

Jono Bacon: Beware False  
Prophets

[more](#)



Ryan Cartwright, 2008-05-01

**How do you replace Microsoft Outlook? Groupware applications**

Ryan Cartwright, 2008-03-20

**Drigg (the pligg alternative) vs. Pligg: why should people switch?**

Tony Mobily, 2008-04-13

Hot topics - last 21 days

**Installing an all-in-one printer device in Debian**

Ryan Cartwright, 2008-05-05

**Things you miss with GNU/Linux**

Ryan Cartwright, 2008-05-01

**Digital Rights Management (DRM): is it in its death throes?**

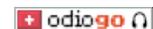
Gary Richmond, 2008-05-07

**Open letter to standards professionals, developers, and activists**

Pieter Hintjens, 2008-05-13



Odiogo



15 How Microsoft Uses  
Novell to Fight  
GNU/Linux, Xen to  
Fight VMWare and  
GNU/Linux

15 Why You Should Reject  
Novell's Moonlight

17 University of Havana  
Finally Switches to  
Free Software

15 Please Welcome  
Digistan

17 Content Protection  
madness on Vista

25 Asus to embed Linux  
into all motherboards

23 Make Your Distro Free  
of Miguel de Icaza's  
junk code

16 Firefox 3 Release  
Candidate now  
available for download

[Get this widget »](#)

Automake to do so, but you can't ask Autoconf to compile, link or run Java-based checks. Java simply isn't supported natively at this time by Autoconf. I believe it's important to point out here that the very nature of the Java language and virtual machine specifications make it far less likely that you'll need to perform a Java-based Autoconf check in the first place.

There is work being actively done on the gcj compiler and tool set, so it's not unreasonable to think that some native Java support will be added to Autoconf at some future date, but gcj is a bit immature yet, and currently very few (if any) GNU packages are written in Java, so the issue is not critical to the GNU community.

That said, there is currently rudimentary support in Automake for both GNU (gcj) and non-GNU Java compilers and VM's. I've used it myself on a project, and it works well, as long as you don't try to push it too far. Given the history of the GNU project, I think it's safe to say that this functionality will definitely improve with age.

If you're into Smalltalk, ADA, Modula, LISP, Forth, or some other non-mainstream language, well then you're probably not too concerned about porting your code to dozens of platforms and CPUs.

As an aside, if you *are* using a non-mainstream language, and you are in fact concerned about the portability of your build systems, then please consider adding support for your language to the Autotools. This is not as daunting a task as you may think, and I guarantee that you'll be an Autotools expert when you're finished. If you think this statement is funny, then consider how Erlang support made it into the Autotools. I'm betting most developers have never heard of Erlang, but members of the Erlang community thought it was important enough to add Erlang support themselves.

## Generating your package build system

The GNU Autotools framework is comprised of three main packages, each of which provides and relies on several smaller components. The three main packages are Autoconf, Automake and Libtool. These packages were invented in that order, and evolved over time. Additionally, the tools in the Autotools packages can depend on or use utilities and functionality from the gettext, m4, sed, make and perl packages, as well as others.

It's very important at this point to distinguish between a maintainer's system and an end-user's system. The design goals of the Autotools specify that an Autotools-generated build system rely only on readily available, preinstalled tools on the host machine. Perl is only required on machines that maintainers use to create distributions, not on end-user machines that build packages from resulting release distributions packages. A corollary to this is that end-users' machines need not have the Autotools installed.

If you've ever downloaded, built and installed software from a "tarball"—a compressed archive with a .tar.gz, .tgz or .tar.bz2

extension—then you’re probably aware of the fact that there is a common theme to this process. It usually looks something like this:

```
$ gzip -cd hackers-delight-1.0.tar.gz | tar -xvf -  
...  
$ cd hackers-delight-1.0  
$ ./configure  
$ make all  
$ sudo make install
```

*NOTE: I have to assume some level of knowledge on your part, and I’m stating right now that this is it. If you’ve performed this sequence of commands before and you know what it means, and if you have a basic understanding of the software development process, then you’ll have no trouble following the content of this book.*

Most developers know and understand the purpose of the `make` utility. But what’s the point of the `configure` script? The use of configuration scripts (generally named `configure`) started a long time ago on Unix systems due to the variety imposed by the fast growing and divergent set of Unix and Unix-like platforms. It’s interesting to note that while Unix systems have generally followed the defacto-standard Unix kernel interface for decades, most software that does anything significant generally has to stretch outside of these more or less standardized boundaries. Configuration scripts are hand-coded shell scripts designed to determine platform-specific characteristics, and to allow users to choose package options before running `make`.

This approach worked well for decades. With the advent of dozens of Linux distributions, the explosion of feature permutations has made writing a decent portable configuration script very difficult—much more so than writing the `makefiles` for a new project. Most people have come up with configuration scripts for their projects using a well-understood and pervasive technique—copy and modify a similar project’s script. By the early 90’s it was becoming apparent to many developers that project configuration was going to become painful if something weren’t done to ease the burden of writing massive shell scripts to manage configuration options—both those related to platform differences, and those related to package options.

## Autoconf

[Autoconf](#) changed this paradigm almost overnight. A quick glance at the `AUTHORS` file in the [Savannah Autoconf project](#) repository will give you an idea of the number of people that have had a hand in the making of Autoconf. The original author was David MacKenzie, who started the Autoconf project in 1991. While configuration scripts were becoming longer and more complex, there were really only a few variables that needed to be specified by the user. Most of these were simply choices to be made regarding components, features and options: Where do I find libraries and header files? Where do I want to install my finished product? Which

optional components do I want to build into my products? With Autoconf, instead of modifying, debugging and losing sleep over literally thousands of lines of supposedly portable shell script, developers can write a short meta-script file, using a concise macro-based language, and let Autoconf generate a perfect configuration script.

A generated configuration script is more portable, more correct, and more maintainable than a hand-code version of the same script. In addition, Autoconf often catches semantic or logic errors that the author would have spent days debugging. Another benefit of Autoconf is that the shell code it generates is as portable as possible between systems that supply any form of the Bourne shell. Mistakes made in portability between shells are by far the most common, and unfortunately the most difficult to find, because no one programmer has access to all versions or brands of Bourne-like shells in existence.

Autoconf generated configure scripts provide a common set of options that are important to all portable, free, open source, and proprietary software projects running on LSB-compliant systems. These include options to modify “standard locations”, a concept I’ll cover in more detail in Chapter 2. Autoconf generated configure scripts also provide project-specific options. These are defined in the `configure.ac` file for each project. I’ll detail this process in Chapter 3.

The Autoconf package provides several programs. Autoconf itself is written in Bourne shell script, while the others are perl scripts.

- `autoconf`
- `autoheader`
- `autom4te`
- `autoreconf`
- `autoscan`
- `autoupdate`
- `ifnames`

### Autoheader

The `autoheader` utility generates a C language header file *template* from `configure.ac`. This template file is usually called `config.h.in`. We’ll cover `autoheader` in greater detail in Chapter 3.

### Autom4te

The `autom4te` utility is a cache manager used by most of the other Autotools. In the early days of Autoconf there was really no need for such a cache, but because most of the Autotools use constructs found in `configure.ac`, the cache speeds up access by successive programs to `configure.ac` by about 40 percent or more. I won’t spend a lot of time on `autom4te` (which is pronounced “automate”, by the way), because it’s mainly used internally by the Autotools, and the only sign you’re given that it’s working is the existence of an `autom4te.cache` directory in your top-level project directory.

## Autoreconf

The autoreconf program can be used to execute the configuration tools in the Autoconf, Automake and Libtool packages as required by the project. The purpose of autoreconf is to minimize the amount of regeneration that needs to be done, based on timestamps, features, and project state. Think of autoreconf as an Autotools bootstrap utility. If all you have is a configure.ac file, running autoreconf will run the tools you need in order to run configure and then make.

## Autoscan

The autoscan program is used to generate a reasonable configure.ac file for a new project. We'll spend some time on autoscan later in Chapter's 3 and 6, as we go through the process of setting up the Autotools on a basic project.

## Autoupdate

The autoupdate utility is used to update your configure.ac or template (\*.in) files to the syntax of the current version of the Autotools. I'll cover autoupdate in more detail in Chapter 2.

## ifnames

The ifnames program is a small, and generally under-utilized program that accepts a list of source files names on the command line, and displays a list of C preprocessor definitions and their containing files on the stdout device. This utility is designed to help you determine what to put into your configure.ac and Makefile.am files for the sake of portability. If your project has already been written with some level of portability in mind, ifnames can help you find out where those attempts are located in your source tree, and what the names of potential portability definitions might be.

Of the tools in this list, only autoconf and autoheader are used directly by the project maintainer while generating a configure script, and actually, as we'll see later, only autoreconf really needs to be called directly. The following diagram shows the interaction between input files and the Autoconf and autoheader programs to generate product files:

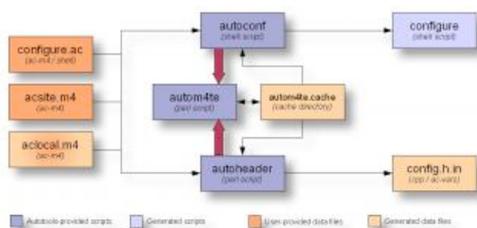


Figure 1: Autoconf and autoheader data flow diagram

*NOTE: I'll follow this data flow diagram format through the rest of this book. Darker colored boxes represent objects that are provided either by the user or by an Autotools package. Lighter*

*shades of the same colors represent generated objects of the same type.*

These tools' primary task is to generate a configure script that can be used by you or others to configure a project build directory. The configure script generated does not rely in any way on the Autotools themselves. In fact, Autoconf is *specifically designed* to generate configure scripts that will run on all Unix or Unix-like platforms that support a Bourne shell. You should be able to generate a configure script from Autoconf, and then successfully execute that script on a machine which does not have the Autotools installed. Not surprisingly, this is actually a common use-case in the free software world, so it's also a well-tested use case.

As you can see in this diagram, Autoconf and autoheader are called by the user. These tools take their input from your project's `configure.ac` file, and various Autoconf-flavored m4 macro definition files. They use `autom4te` to maintain cache information. Autoconf generates your configure script, a very portable Bourne shell script that provides your project with configuration capabilities. Autoheader generates the `config.h.in` template based on macro definitions in `configure.ac`.

You may have noticed the apparent identity crisis being suffered by the `aclocal.m4` input file. Is that a bit of a blush on that box—is it a generated file, or a user-provided file? Well, the answer is that it's both, and I'll explain this in more detail in the next section.

## Automake

So, what's so difficult about writing a makefile? Well, actually, once you've done it a few times, writing a *basic* makefile for a new project is really rather trivial. The problems occur when you try to do more than just the basics. And let's face it—what project maintainer has ever been satisfied with just a basic makefile?

The single most significant difference between a successful free software project and one that rarely gets a second glance can be found deep in the heart of project maintenance details. These details include providing the so-called “standard make targets”. Potential users become disgusted with a project fairly easily—especially when certain bits of expected functionality are missing or improperly written. Users have come to expect certain more or less standard make targets. A make target is a goal specified on the make command line:

```
$ make install
```

In this example, `install` is the goal or target. Common make targets include `all`, `clean` and `install`, among others. You'll note that none of these are *real* targets. A real target is a file produced by the build system. If you're building an executable called `doofabble`, then you'd expect to be able to type:

```
$ make doofabble
```

This would generate an actual executable file called `doofabble`. But specifying real targets on the make command line is more work than necessary. Each project must be built differently—make `doofabble`, make `foodabble`, make `abfooble`, etc. Why not just type `make` or `make all`, if there is more than one binary to be made? So `all` has become an expected pseudo-target, but “expected” doesn’t mean “automatic”.

Supporting the expected set of standard targets can be fairly challenging. As with configuration scripts, the most widely used implementation is one written in the late 80’s and copied from project to project throughout the internet. Why? Because writing it yourself is error prone. In fact, copying it is just as error-prone. It’s like getting a linked-list implementation right the first time. The process is well-understood by any veteran software engineer, but it still rarely happens. Object-oriented programming languages like C++ and Java provide libraries and templates for these constructs now—not because they’re hard to implement by hand, but because doing so is error-prone, and there’s no point in re-inventing the wheel—yet again.

Automake’s job is to convert a much simplified specification of your project’s build process into standard boilerplate makefile syntax that always works correctly the first time, and provides all the standard functionality expected of a free software project. In actuality, Automake creates projects that support guidelines defined in the GNU Coding Standards, which I’ll cover in greater detail in Chapter 2.

The Automake package provides the following tools in the form of perl scripts:

- `automake`
- `aclocal`

The primary task of the Automake program is to generate standard makefile templates (named `Makefile.in`) from high-level build specification files (named `Makefile.am`). One of the most interesting and useful aspects of the way Automake works is that the `Makefile.am` input files are mostly just regular makefiles. If you put only the few required Automake definitions in a `Makefile.am`, you’ll get a `Makefile.in` file containing several hundred lines of makefile code. But if you add additional makefile syntax to your `Makefile.am` files, this code will be transferred to the most functionally correct location in the resulting `Makefile.in`. In fact, you can (if you wish) write pure make syntax in your `Makefile.am` files, and they’ll work just fine (as long as you actually write them correctly, that is). This pass-through feature gives you the power and flexibility to extend Automake’s functionality with your project’s own special requirements.

### **Aclocal**

The `aclocal` utility is actually documented by the GNU manuals as a temporary work-around for a certain lack of flexibility in `Autoconf`. `Autoconf` was designed and written first, and then a few years later, the idea for Automake was conceived as an

add-on for Autoconf. But Autoconf was really not designed to be extensible on the scale required by Automake.

Automake adds an extensive set of macros to those provided by Autoconf. The originally documented method for adding user-defined macros to an Autoconf project was to create a file called `aclocal.m4` in the same directory as `configure.ac`. Any user-provided extension macros were to be placed in this file, and Autoconf would automatically read it while processing `configure.ac`. From the perspective of the Automake designers, this existing extension mechanism was too good to pass up. But requiring the user to add an `m4_include` line to `aclocal.m4` seemed a bit brittle. Instead, the `aclocal` utility was designed to create a project's `aclocal.m4` file, containing all the required Automake macros. Since Automake's `aclocal` utility basically took over `aclocal.m4` for its own purposes, it was also designed to read a new user-provided macro file called `acinclude.m4`.

Essentially, `aclocal`'s job is to create an `aclocal.m4` file by consolidating various macro files from installed Autotool packages and user-specified locations, such that Autoconf can find them all in one place.

For the sake of modularity, the Autoconf manual is still unaware of the `aclocal` utility—for the most part. The current revision of the manual rants a bit on the subject of where `aclocal` functionality should actually be. Automake's manual originally suggested that you should rename `aclocal.m4` to `acinclude.m4` when adding Automake to an existing Autoconf project. This method is still followed rigorously in new projects.

However, the latest documentation from both sets of tools suggests that the entire `aclocal/acinclude` paradigm is now obsolete, in favor of a newer method of specifying a *directory* containing `m4` macro files. The current recommendation is that you create a directory in your project directory called simply `m4` (`acinclude` seems more appropriate to this author), and add macros in the form of individual `.m4` files to this directory. All files in this directory will be gathered into `aclocal.m4` before Autoconf processes your `configure.ac` file. Ultimately, `aclocal` will be replaced by functionality in Autoconf itself. (Given the fairly complex nature of `aclocal` functionality, and given that most of the other tools are already written in perl, I'm guessing that Autoconf will be rewritten in perl, at this point.)

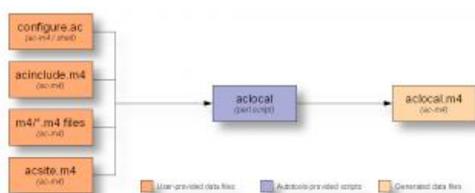


Figure 2: Aclocal data flow diagram

With `aclocal` behind us, it should be more apparent now why the `aclocal.m4` box in the Autoconf data flow diagram of Figure 1 above couldn't decide which color it should be. When used

without Automake and Libtool, the `aclocal.m4` file is written by hand, but when used in conjunction with Automake and Libtool, the file is generated by the `aclocal` utility, and `acinclude.m4` is used to provide project-specific macros.

## Libtool

How do you build shared libraries on different Unix platforms without adding a lot of very platform-specific conditional code to your build system and source code? This is the question that the Libtool package tries to address.

There's a significant amount of visible functionality in Unix and Unix-like platforms that is the same from one platform to another. However, one very significant difference is how shared libraries are built, named and managed. Some platforms don't even provide native shared libraries (although it's rare these days). Some platforms name their libraries `libsomething.so`, while others use `something.o`. Some use `libsomething.a`, while others use `libsomething.sa`. Some platforms provide `libdl` (`dlopen/dlsym/dlclose`) to allow software to dynamically load and access library functionality at runtime. Others provide other mechanisms—or none at all.

All of these differences have been carefully considered by the authors of the Libtool project. Dozens of platforms are currently supported by Libtool, and adding support for new platforms is done via the open source way—someone who cares (and knows how) supplies a patch to the Libtool mailing list, and the maintainers look it over and apply it to the source code for the next release.

Libtool not only provides a set of Autoconf macros that hide library naming differences in makefiles, but it also provides an optional library of dynamic loader functionality that can be added to your programs, allowing you to write more portable *runtime* dynamic shared object management code.

The `libtool` package provides the following programs, libraries and header files:

- `libtool` (program)
- `libtoolize` (program)
- `ltdl` (static and shared libraries)
- `ltdl.h` (header)

The `libtool` shell script is a generic version of Libtool designed to be used by programs on your platform. There's nothing specific to a project in this particular copy of `libtool`.

## Libtoolize

The `libtoolize` shell script is used to prepare your project to use Libtool. In reality, `libtoolize` generates a custom version of the `libtool` script in your project directory. This script is then executed at the appropriate time by Automake-generated makefiles.

## The Libtool C API—`ltdl`

The Libtool package also provides the ltdl library and header files, which provide a consistent run-time shared object manager across platforms. The ltdl library may be linked statically or dynamically into your programs, giving them a consistent runtime shared library access interface from one platform to another.

The following data flow diagram illustrates the interaction between Automake and Libtool scripts and input files to create products used by users to configure and build your project:

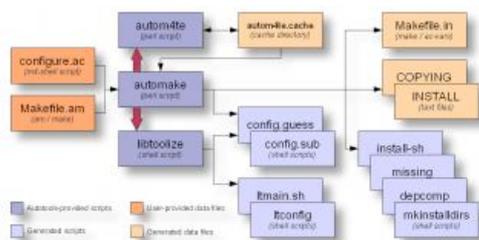


Figure 3: Automake and Libtool data flow diagram

Automake and Libtool are both standard pluggable options that can be added to `configure.ac` with a few simple macro calls.

## Building your package

While, as maintainer, you probably build your software packages a lot more often than do your users, you also have the advantage of being intimately familiar with your project's components, architecture and build system. That's why you ought to be concerned that your users' build experience is much simpler than yours. (And it wouldn't hurt a bit if you got some benefit from this concern, as well.)

## Running configure

Once the Autotools have finished their work, you're left with a shell script called `configure`, and one or more `Makefiles.in` files. These product files are intended to be packages with project release distribution packages. Your users download these packages, unpack them, and run `configure` and `make`. The `configure` script generates `Makefiles` from the `Makefile.in` files. It also generates a `config.h` header file from the `config.h.in` file built by `autoheader`.

So why didn't the Autotools just generate the makefiles directly to be shipped with your release? One reason is that without makefiles, you can't run `make`. This means that you're forced to run `configure` first, after you download and unpack a project distribution package. `Makefile.in` files are nearly identical to the makefiles you might write by hand, except that you didn't have to. And they do a lot more than most people are willing to hand code into a set of makefiles. Another reason is that the `configure` script may then insert platform-characteristics and user-specified optional features directly into your makefiles,

making them more specifically tailored to the platforms on which they are being used.

The following diagram illustrates the interaction between configure and the scripts that it executes during the build process to create your Makefiles and your config.h header file:

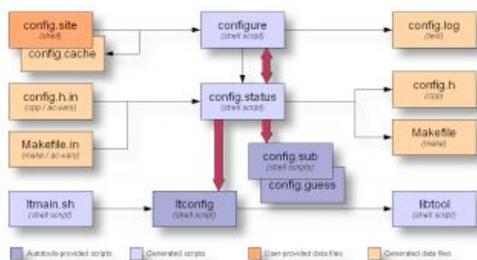


Figure 4: Configure script data flow diagram

The configure script appears to have this weird sort of incestuous relationship with another script called config.status. I'll bet you've always thought that your configure script generated your makefiles. As it turns out, the only file (besides a log file) that configure generates is config.status. The configure script's function is to determine platform characteristics and features available, as specified in configure.ac. Once it has this information, it generates config.status such that it contains all of the check results, and then calls it. The newly generated config.status file uses the check information (now embedded within it) to generate platform-specific config.h and makefiles, as well as any other files specified for instantiation in configure.ac. As the double ended red arrow shows, config.status can also call configure. When used with the `--recheck` option, config.status will call configure with the same command line options with which it was originally generated.

The configure script also generates a log file called config.log, which contains very useful information about why a particular execution of configure failed on your user's platform. As maintainer, you can use this information to help you debug user problems. Just ask them to send you their config.log file. The problem is often in plain sight. Another nice feature of config.log is that it logs how configure was executed—which command line options were used.

From a user perspective, this could be really handy, as he comes back from a long vacation, and can't remember what options he used to generate the project build directory. But Autoconf-generated configure scripts make it even simpler than this. If you need to re-generate makefiles and config.h header files for some reason, just type `./config.status` in the project build directory. The output files will be generated using the same options originally used to generate the config.status file.

### Remote build directories

A little-known feature of Autotools build environments is that

they need not be generated within a project source directory tree. That is, a user may execute configure remotely, and generate a full build environment within a remote build directory.

In the following example, Joe User downloads doofabble 3.0 and unpacks it. Then he creates two sibling directories called doofabble-3.0.debug and doofabble-3.0.release. He cd's into doofabble-3.0.debug, executes doofabble-3.0's configure script remotely with a doofabble-specific debug option, and then runs make. Finally, he switches over to the doofabble-3.0.release directory and does the same thing, this time running configure without the debug option enabled:

```
$ tar -zxvf doofabble-3.0.tar.gz
$ mkdir doofabble-3.0.debug
$ cd doofabble-3.0.debug
$ ../doofabble-3.0/configure --enable-debug
$ make
...
$ cd ..
$ mkdir doofabble-3.0.release
$ cd doofabble-3.0.release
$ ../doofabble-3.0/configure
$ make
...
```

Users don't often care about remote build functionality because all they generally want to do is configure, make and install your code on their own platforms. Maintainers, on the other hand should find remote build functionality very useful, as it allows them to, 1) maintain a reasonably pristine source tree, and 2) maintain multiple build environments for their project, each with potentially complex configuration options. Rather than reconfigure a single build environment, they may simply switch between build directories configured in multiple different ways.

## Running make

Finally, you run make. Just plain old make. In fact, the Autotools designers went to a LOT of trouble to ensure that you didn't need any special version or brand of make. You don't need GNU make—you can use Solaris make, or BSD Unix make if you wish (read, "if you must").

The following diagram depicts the interaction between the make utility and the generated makefiles during the build process to create your project products:



Figure 5: Make data flow diagram

This diagram shows make running several generated scripts, but these are all really ancillary to the make process.

## Summary

In this chapter I've presented a high-level overview of the Autotools to give you a feel for how everything ties together.

In the next chapter, we'll begin creating a hand-coded build system for a toy project. The idea is that you'll become familiar with the requirements of a reasonable build system, and how much can be done for you by the Autotools.

Too many developers these days start out with the Autotools, not having aquired through the "school of hard knocks" the experience to know what it's really doing for them. This can lead to frustration, and a negative attitude. In the next chapter, you'll become familiar with the rationale for a lot of the original design of the Autotools. In understanding this background information, my hope is that any potential negative bias you may already have for the Autotools will be tempered a bit.

◀ Autotools: a practitioner's guide to Autoconf, Automake and Libtool

u  
p

Chapter 2: Project management and the GNU coding standards

▶

[Login or register to post comments](#)

8858 reads



- COLUMNS
- COMMUNITY POSTS
- ISSUES
- BOOKS
- FORUM
- FS NEWS

Ads by Google

**Digital Command Control**

Digitrains supply & install systems and decoders; help & advice is free www.digitrains.co.uk

**Why Is My Computer Slow?**

Perform Registry Compression, Defrag, Manage & Cleanup Your PC. www.RegClean.com

**PostgreSQL Courses**

DBA and Developer courses from the experts www.2ndQuadrant.com

**Zip & Unzip Free Download**

Download free zip & unzip software. Free Full Version, no need to buy UberDownloads.com/Unzip

**Registry Scanners Exposed**

Don't Download Any Registry Scanner Until You See This Shocking Review! RegistryCleanerGuide.org

- Microsoft May Seek New Yahoo ...
- Building Semantics is Different ...
- Lux: multi-touch for OS X

Sponsored by **Novell.**

**Best voted contents**

- [The Bizarre Cathedral - 3](#)  
Ryan Cartwright, 2008-05-05
- [The Bizarre Cathedral - 2](#)  
Ryan Cartwright, 2008-04-27
- [Indexing offline CD-ROM archives](#)  
Terry Hancock, 2008-05-03
- [Microsoft and free software? I don't think so...](#)  
Terry Hancock, 2008-04-26

Home » Autotools: a practitioner's guide to Autoconf, Automake and Libtool

# Chapter 2: Project management and the GNU coding standards



by John Calcote

In Chapter 1, I gave a brief overview of the Autotools and some of the resources that are currently available to help reduce the learning curve. In this chapter, we're going to step back a little and examine project organization techniques that are applicable to all projects, not just those whose build system is managed by the Autotools.

When you're done reading this chapter, you should be familiar with the common make targets, and why they exist. You should also have a solid understanding of why projects are organized the way they are. Trust me—by the time you finish this chapter, you'll already be well on your way to a solid understanding of the GNU Autotools.

The information provided by this chapter comes primarily from two sources:

- [The GNU Coding Standards Document](#)
- [The Filesystem Hierarchy Standard](#)

In addition, you may find the GNU make manual very useful, if you'd like to brush up on your make syntax:

- [The GNU Make Utility Manual](#)

## Creating a new project directory structure

There are two questions to ask yourself when setting up a new open source software (OSS) project build system:

- What platforms will I target?
- What do my users expect?

The first is an easy question to answer - you get to decide, but don't be too restrictive. Free software projects become great due to the number of people who've adopted them. Limiting the number of platforms arbitrarily is the direct equivalent of limiting the number of users. Now, why would you want to do that?!

The second question is more difficult, but not unsolvable. First, let's narrow the scope to something manageable. We really mean to say, "What do my users expect of my build system?" A common approach for many OSS developers of determining these expectations is to download, unpack, build and install about a thousand different packages. You think I'm kidding? If

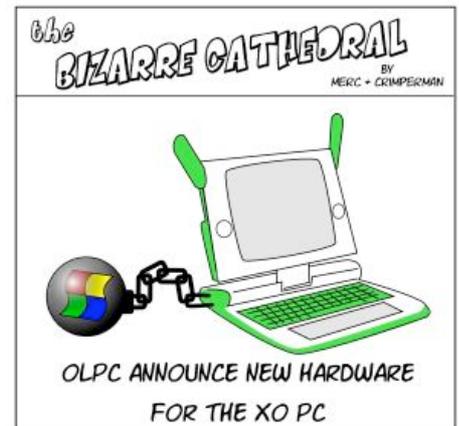
**Need EDM Now?**

Enterprise Decision Management Accelerate your decision process! www.zementis.com

**Registry Scanners Exposed**

Don't Download Any Registry Scanner Until You See This Shocking Review! RegistryCleanerGuide.org

Ads by Google



http://www.freesoftwaremagazine.com (c) Copyright 2008 Ryan Cartwright CC: BY-NC-SA

From the FSM staff...

[The Top 10 Everything](#) (Dave). The good, the bad and the ugly.

[Free Software news](#) (Dave & Bridget). A site about short stories and writing.

[Book Reviews: Illiterary](#) (Bridget). Book reviews, blogs, and short stories.

Hot topics - last 60 days

[Installing an all-in-one printer device in Debian](#) Ryan Cartwright, 2008-05-05

[What is the free software community?](#) Tony Mobily, 2008-03-29

[Things you miss with GNU/Linux](#)



you do this, eventually, you will come to know intuitively what your users expect of your build system. Unfortunately, package configuration, build and install processes vary so far from the “norm” that it’s difficult to come to a solid conclusion about what the norm really is when using this technique.

A better way is to go directly to the source of the information. Like many developers new to the OSS world, I didn’t even know there was a source of such information when I first started working on OSS projects. As it turns out, the source is quite obvious, after a little thought: The Free Software Foundation (FSF), better known as the GNU project. The FSF has published a document called [The GNU Coding Standards](#), which covers a wide variety of topics related to writing, publishing and distributing free software—specifically for the FSF. Most non-GNU free software projects align themselves to one degree or another with the GNU Coding Standards. Why? Well...just because they were there first. And because their ideas make sense, for the most part.

## Project structure

We’ll start with a simple example project, and build on it as we continue our exploration of source-level software distribution. OSS projects generally have some sort of catchy name—often they’re named after some past hero or ancient god, or even some made-up word—perhaps an acronym that can be pronounced like a real word. I’ll call this the jupiter project, mainly because that way I don’t have to come up with functionality that matches my project name! For jupiter, I’ll create a project directory structure something like this:

```
$ cd projects
$ mkdir -p jupiter/src
$ touch jupiter/Makefile
$ touch jupiter/src/Makefile
$ touch jupiter/src/main.c
$ cd jupiter
$
```

Woot! One directory called `src`, one C source file called `main.c`, and a makefile for each of the two directories. Minimal yes, but hey, this is a new project, and everyone knows that the key to a successful OSS project is evolution, right? Start small and grow as needed (and, as you have time and inclination).

We’ll start with support for the most basic of targets in any software project: `all` and `clean`. As we progress, it’ll become clear that we need to add a few more important targets to this list, but for now, these will get us going. The top-level Makefile does very little at this point, merely passing requests for `all` and `clean` down to `src/Makefile` recursively. In fact, this is a fairly common type of build system, known as a *recursive build system*. Here are the contents of each of the three files in our project:

### Makefile

## Buzz authors

Marco Marongiu

Pieter Hintjens

drascus321

dogboi

Richard Rothwell

## All news

John Crawford: Arizona  
LoCo Hardy Release Party

John Crawford: Ubuntu  
Weekly Newsletter #91

Matthew Helmke: Morocco  
blocks Google Maps

Andy Price: Synchronicity  
and Sharing

Jono Bacon: Beware False  
Prophets

more



## GNU/Linux

Ryan Cartwright, 2008-05-01

**How do you replace Microsoft Outlook? Groupware applications**

Ryan Cartwright, 2008-03-20

**Drigg (the pligg alternative) vs. Pligg: why should people switch?**

Tony Mobily, 2008-04-13

Hot topics - last 21 days

**Installing an all-in-one printer device in Debian**

Ryan Cartwright, 2008-05-05

**Things you miss with GNU/Linux**

Ryan Cartwright, 2008-05-01

**Digital Rights Management (DRM): is it in its death throes?**

Gary Richmond, 2008-05-07

**Open letter to standards professionals, developers, and activists**

Pieter Hintjens, 2008-05-13



Odiogo



- 15 How Microsoft Uses Novell to Fight GNU/Linux, Xen to Fight VMWare and GNU/Linux
- 15 Why You Should Reject Novell's Moonlight
- 17 University of Havana Finally Switches to Free Software
- 15 Please Welcome Digistan
- 17 Content Protection madness on Vista
- 25 Asus to embed Linux into all motherboards
- 23 Make Your Distro Free of Miguel de Icaza's junk code
- 16 Firefox 3 Release Candidate now available for download

Get this widget »

```
all clean jupiter:
    $(MAKE) -C src $@
```

### src/Makefile

```
all: jupiter

jupiter: main.c
    gcc -g -O0 -o $@ $+

clean:
    -rm jupiter
```

### src/main.c

```
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char * argv[])
{
    printf("Hello from %s!\n", argv[0]);
    return 0;
}
```

At this point, you may need to stop and take a refresher course in `make` syntax. If you're already pretty well versed on `make`, then you can skip the sidebar entitled, "Some makefile basics". Otherwise, give it a quick read, and then we'll continue building on this project.

### Some makefile basics

For those like myself who use `make` only when they have to, it's often difficult to remember exactly what goes where in a makefile. Well, here are a few things to keep in mind. Besides comments, which begin with a HASH mark, there are only three types of entities in a makefile:

- variable assignments
- rules
- commands

*NOTE: There are also conditional statements, but for the purposes of this chapter, we need not go into conditionals. See the GNU make manual for more information.*

Commands always start with a TAB character. Any line in a makefile beginning with a TAB character is ALWAYS considered by `make` to be a command. Commands should always be associated with a preceding rule. In fact, if you insert a TAB character before a line not preceded by a rule, you'll get a very strange error message from `make`. GNU `make` is a little better these days about figuring out what you did wrong and telling you.

The general layout of a makefile is:

```
var1=val1
var2=val2
```

```

...

rule1
    cmd1a
    cmd1b
    ...

rule2
    cmd2a
    cmd2b
    ...

```

Variable assignments may take place at any point in the makefile, however you should be aware that `make` reads each makefile twice. The first pass gathers variables, and the second pass resolves dependencies defined by the rules. So regardless of where you put your variable definitions, `make` will act as though they'd all been declared at the top. Furthermore, `make` binds variable references to values at the very last minute—at the point they are actually used in a command or rule. So, in general, variables may be assigned values by reference to other variables that haven't even been assigned yet. Thus, the order of variable assignment isn't really that important.

`make` is a rule-based command engine. The rules indicate when and which commands should be executed. When you prefix a line with a TAB character, you're telling `make` that you want it to execute these statements from a shell according to the rules specified on the line above.

*NOTE: The fact that commands are required to be prefixed with an essentially invisible character is one of the most frustrating aspects of makefile syntax to both neophytes and experts alike. The error messages generated by the `make` utility when a required TAB is missing or when an unintentional TAB is inserted are obscure at best. As mentioned earlier, GNU `make` does a better job with such error messages these days. Nonetheless, be careful to use TAB characters properly in your makefiles—only before commands, which in turn immediately follow rules.*

Of the remaining lines, those containing an EQUAL sign are variable definitions. Variables in makefiles are nearly identical to shell or environment variables. In Bourne shell syntax, you'd reference a variable in this manner: `${my_var}`. In a makefile, the same syntax applies, except you would use parentheses instead of french braces: `$(my_var)`. As in shell syntax, the delimiters are optional, but should be used to avoid ambiguous syntax, thus `$my_var` is functionally equivalent to `$(my_var)`.

One caveat: If you ever want to use a shell variable inside a `make` command, you need to escape the DOLLAR sign by doubling it. For instance, `$$shell_var`. This need arises occasionally, and it nearly always catches me off-guard the first time I use it in a new project.

Variables may be defined anywhere and used anywhere within a makefile. By default, `make` will read the entire

process environment into the `make` variable table before processing the makefile, so you can access any environment variables as if they were defined in the makefile itself. In addition, `make` defines several useful variables of its own, such as the `MAKE` variable, whose value is the file system path used to invoke the current `make` process.

All other lines in a makefile are rules of one type or another. The rules used in my examples are known as “common” `make` rules, containing a single COLON character. The COLON character separates targets on the left from dependencies on the right. Targets are products—generally file system entities that can be produced by running one or more commands, such as a C compiler. Dependencies are source objects, or objects from which targets may be created. These may be computer language source files, or anything really that can be used by a command to generate a target object.

For example, a C compiler takes dependency `main.c` as input, and generates target `main.o`. A linker takes dependency `main.o` as input, and generates a named executable target, `jupiter` in these examples:

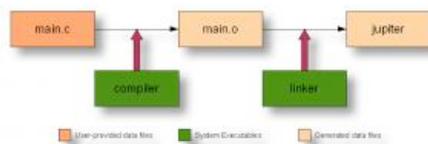


Figure 6: Compile and link process

The `make` utility implements some fairly complex logic to determine when a rule should be run based on whether the target exists or is older than its dependencies, but the syntax is trivial enough:

```
jupiter: main.o
    ld main.o ... -o jupiter

main.o: main.c
    gcc -c -g -O2 -o main.o main.c
```

This sample makefile contains two rules. The first says that `jupiter` depends on `main.o`, and the second says that `main.o` depends on `main.c`. Ultimately, of course, `jupiter` depends on `main.c`, but `main.o` is a necessary intermediate dependency, because there are two steps to the process—compile and link—with an intermediate result in between. For each rule, there is an associated list of commands that `make` uses to build the target from the list of dependencies.

Of course, there is an easier way—`gcc` (as with most compilers) will call the linker for you—which, as you can probably tell from the elipsis in my example above, is *very* desirable. This alleviates the need for one of the rules, and

provides a convenient way of adding more dependent files to the single remaining rule:

```
sources = main.c print.c display.c

jupiter: $(sources)
    gcc -g -O2 -o jupiter $(sources)
```

In this example, I've added a `make` variable to reduce redundancy. We now have a list of source files that is referenced in two places. But, it seems a shame to be required to reference this list twice in this manner, when the `make` utility knows which rule and which command it's dealing with at any moment during the process. Additionally, there may be other objects in the dependency list that are not in the `sources` variable. It would be nice to be able to reference the entire dependencies list without duplicating the list.

As it happens, there are various “automatic” variables that can be used to reference portions of the controlling rule during the execution of a command. For example `$(@)` (or the more common syntax `$@`) references the current target, while `$(+)` references the current list of dependencies:

```
sources = main.c print.c display.c

jupiter: $(sources)
    gcc -g -O2 -o $@ $+
```

If you enter “`make`” on the command line, the `make` utility will look for the first target in a file named “`Makefile`” in the current directory, and try to build it using the rules defined in that file. If you specify a different target on the command line, `make` will attempt to build that target instead.

Targets need not be files only. They can also be so-called “phony targets”, defined for convenience, as in the case of `all` and `clean`. These targets don't refer to true products in the file system, but rather to particular outcomes—the directory is “cleaned”, or “all” desirable targets are built, etc.

In the same way that dependencies may be listed on the right side of the COLON, rules for multiple targets with the same dependencies may be combined by listing targets on the left side of the COLON, in this manner:

```
all clean jupiter:
    $(MAKE) -C src $@
```

The `-C` command-line option tells `make` to change to the specified directory before looking for a makefile to run.

GNU Make is significantly more powerful than the original Unix `make`, although completely backward compatible, as long as GNU extensions are avoided. The [GNU Make manual](#) is available online. O'Reilly has an excellent book on [the original Unix make utility](#) and all of its many nuances. They also have a more recent [book written](#)

specifically for GNU make that covers GNU Make extensions.

## Creating a source distribution archive

It's great to be able to type "make all" or "make clean" from the command line to build and clean up this project. But in order to get the jupiter project source code to our users, we're going to have to create and distribute a source archive.

What better place to do this than from our build system. We could create a separate script to perform this task, and many people have done this in the past, but since we have the ability, through phony targets, to create arbitrary sets of functionality in make, and since we already have this general purpose build system anyway, we'll just let make do the work for us.

Building a source distribution archive is usually relegated to the `dist` target, so we'll add one. Normally, the rule of thumb is to take advantage of the recursive nature of the build system, by allowing each directory to manage its own portions of a global process. An example of this is how we passed control of building jupiter down to the `src` directory, where the jupiter source code is located. However, the process of building a compressed archive from a directory structure isn't really a recursive process—well, okay, yes it is, but the recursive portions of the process are tucked away inside the `tar` utility. This being the case, we'll just add the `dist` target to our top-level makefile:

### Makefile

```
package = jupiter
version = 1.0
tarname = $(package)
distdir = $(tarname)-$(version)

all clean jupiter:
    $(MAKE) -C src $@

dist: $(distdir).tar.gz

$(distdir).tar.gz: $(distdir)
    tar chof - $(distdir) | \
    gzip -9 -c >$(distdir).tar.gz
    rm -rf $(distdir)

$(distdir):
    mkdir -p $(distdir)/src
    cp Makefile $(distdir)
    cp src/Makefile $(distdir)/src
    cp src/main.c $(distdir)/src

.PHONY: all clean dist
```

Now, there are a couple of noteworthy items in this version of the makefile. The first is that we've added a new construct, the

`.PHONY` rule. At least it seems like a rule—it contains a COLON character, anyway. The `.PHONY` rule is a special kind of rule called a dot-rule, which is built in to `make`. The `make` utility understands several different dot-rules. The purpose of the `.PHONY` rule is simply to tell `make` that certain targets don't generate file system objects, so `make` won't go looking for product files in the file system that are named after these targets.

The second point of interest is the use of a leading DASH character in some of our commands. A leading DASH character tells `make` to not care about the status code of the associated command. Normally `make` will stop execution with an error message on the first command that returns a non-zero status code to the shell. We use a leading DASH character on the initial `rm` commands because we want to delete previously created product files that may or may not exist, and `rm` will return an error if we attempt to delete a non-existent file.

Another such character that you may encounter is the leading ATSIGN (`@`) character. A command prefixed with an ATSIGN character tells `make` not to print the command as it executes it. Normally `make` will print each command (unless you've given `make` a command line parameter (`-s`), or used another dot-rule, the `.SILENT` rule, to keep it quiet). A leading ATSIGN tells `make` that you *never* want to see this command. This is a common thing to do on `echo` commands—you don't want `make` to print `echo` commands because then your message will be printed twice, and that's just ugly.

We've added the new `dist` target in the form of three rules for the sake of readability, modularity and maintenance. This is a great rule of thumb to following in any software engineering process: Build large processes from smaller ones, and reuse the smaller processes where it makes sense to do so.

The `dist` target depends on the existence of the ultimate goal, a source-level compressed archive package, `jupiter-1.0.tar.gz`—also known as a “tarball”. We've added a `make` variable for the version number to ease the process of updating the project version later, and We've used another variable for the package name for the sake of possibly porting this makefile to another project. We've also logically split the functions of package name and tar name, in case we want them to be different later—the default tar name is the package name. Finally, We've combined references to these variables into a `distdir` variable to reduce duplication and complexity in the makefile.

The rule that builds the tarball indicates how this should be done with a command that uses the `gzip` and `tar` utilities to create the file. But, notice also that the rule has a dependency—the directory to be archived. We don't want everything in our project to go into our tarball—only exactly those files that are necessary for the distribution. Basically, this means any file required to build and install our project. We certainly don't want object files and executables from our last build attempt to end up in the archive, so we have to build a directory

containing exactly what we want to ship. This pretty much means using individual `cp` commands, unfortunately.

Since there's a rule in the makefile that tells how this directory should be created, `make` runs the commands for this rule *before* running the commands for the current rule. The `make` utility runs rules to build dependencies recursively until the actual requested target's commands can be run.

## Forcing a rule to run

There's a subtle flaw in the `$(distdir)` target that may not be obvious, but it will rear its ugly head at the worst times. If the archive directory already exists when you type `make dist`, then `make` won't try to create it. Try this:

```
$ mkdir jupiter-1.0
$ make dist
tar chof - jupiter-1.0 | gzip -9 -c >jupiter-1.0...
rm -rf jupiter-1.0 &> /dev/null
$
```

Notice that the `dist` target didn't copy any files—it just built an archive out of the existing `jupiter-1.0` directory, which was empty. Our end-users would have gotten a real surprise when they unpacked this tarball!

The problem is that the `$(distdir)` target is a real target with no dependencies, which means that `make` will consider it up-to-date as long as it exists in the file system. We could add `$(distdir)` to the `.PHONY` rule, but this would be a lie—it's not a phony target, it's just that we want to force it to be rebuilt every time.

The proper way to ensure it gets rebuilt is to have it not exist before `make` attempts to build it. A common method for accomplishing this task is to create a true phony target that will run every time, and add it to the dependency chain at or above the `$(distdir)` target. For obvious reasons, a commonly used name for this sort of target is "FORCE":

### Makefile

```
...
$(distdir).tar.gz: FORCE $(distdir)
    tar chof - $(distdir) |\
        gzip -9 -c >$(distdir).tar.gz
    rm -rf $(distdir)

$(distdir):
    mkdir -p $(distdir)/src
    cp Makefile $(distdir)
    cp src/Makefile $(distdir)/src
    cp src/main.c $(distdir)/src

FORCE:
    -rm $(distdir).tar.gz &> /dev/null
    -rm -rf $(distdir) &> /dev/null
```

```
.PHONY: FORCE all clean dist
```

The FORCE commands are executed every time because FORCE is a phony target. By making FORCE a dependency of the tarball, we're given the opportunity to delete any previously created files and directories before make begins to evaluate whether or not these targets' commands should be executed. This is really much cleaner, because we can now remove the "pre-cleanup" commands from all of the rules, except for FORCE, where they really belong.

There are actually more accurate ways of doing this—we could make the `$(distdir)` target dependent on all of the files in the archive directory. If any of these files are newer than the directory, the target would be executed. This scheme would require an elaborate shell script containing sed commands or non-portable GNU Make functions to replace file paths in the dependency list for the copy commands. For our purposes, this implementation is adequate. Perhaps it would be worth the effort if our project were huge, and creating an archive directory required copying and/or generating thousands of files.

The rule for building the archive directory is the most frustrating of any in this makefile—it contains commands to copy files *individually* into the distribution directory. What a sad shame! Everytime we change the file structure in our project, we have to update this rule in our top-level makefile, or we'll break our `dist` target.

But, there's nothing to be done for it. We've made the rule as simple as possible. Now, we just have to remember to manage this process properly. But unfortunately, breaking the `dist` target is not the worst thing that could happen if we forget to update the `distdir` rule's commands. The `dist` target may continue to *appear* to work, but not actually copy all of the required files into the tarball. This will cause us some embarrassment when our users begin to send us emails asking why our tarball doesn't build on their systems.

In fact, this is a far more common possibility than that of breaking the `dist` target, because the more common activity while working on a project is to add files to the project, not move them around or delete them. New files will not be copied, but the `dist` rule won't notice the difference.

If only there were some way of unit-testing this process. As it turns out, there is a way of performing a sort of self-check on the `dist` target. We can create yet another phony target called "`distcheck`" that does exactly what our users will do—unpack the tarball, and build the project. If the build process fails, then the `distcheck` target will break, telling us that we forgot something crucial in our distribution.

### Makefile

```
...
distcheck: $(distdir).tar.gz
    gzip -cd $+ | tar xvf -
    $(MAKE) -C $(distdir) all clean
```

```

rm -rf $(distdir)
@echo "*** Package $(distdir).tar.gz\
ready for distribution."
...
.PHONY: FORCE all clean dist distcheck

```

Here, we've added the `distcheck` target to the top-level makefile. Since the `distcheck` target depends on the `tarball` itself, it will first build a tarball using the same targets used by the `dist` target. It will then execute the `distcheck` commands, which are to unpack the tarball it just built and run "make all clean" on the resulting directory. This will build both the `all` and `clean` targets, successively. If that process succeeds, it will print out a message, telling us that we can sleep well knowing that our users will probably not have a problem with this tarball.

Now all we have to do is remember to run "make `distcheck`" *before* we post our tarballs for public distribution!

## Unit testing anyone?

Some people think unit testing is evil, but really—the only rationale they can come up with for not doing it is laziness. Let's face it—proper unit testing is hard work, but it pays off in the end. Those who do it have learned a lesson (usually as children) about the value of delayed gratification.

A good build system is no exception. It should incorporate proper unit testing. The commonly used target for testing a build is the `check` target, so we'll go ahead and add the `check` target in the usual manner. The test should probably go in `src/Makefile` because `jupiter` is built in `src/Makefile`, so we'll have to pass the `check` target down from the top-level makefile.

But what commands do we put in the `check` rule? Well, `jupiter` is a pretty simple program—it prints out a message, "Hello from <path>jupiter!", where <path> is variable, depending on the location from which `jupiter` was executed. We could check to see that `jupiter` actually does output such a string. We'll use the `grep` utility to test our assertion.

### Makefile

```

...
all clean check jupiter:
    $(MAKE) -C src $@
...
.PHONY: FORCE all clean dist distcheck

```

### src/Makefile

```

...
check: all
    ./jupiter | grep "Hello from .*jupiter!"
    @echo "*** ALL TESTS PASSED ***"
...

```

```
.PHONY: all clean check dist distcheck
```

Note that `check` is dependent on `all`. We can't really test our products unless they've been built. We can ensure they're up to date by creating such a dependency. Now `make` will run commands for `all` if it needs to before running the commands for `check`.

There's one more thing we could do to enhance our system a bit. We can add the `check` target to the `make` command in our `distcheck` target. Between the `all` and `clean` targets seems appropriate:

### Makefile

```
...
distcheck: $(distdir).tar.gz
    gzip -cd $+ | tar xvf -
    $(MAKE) -C $(distdir) all check clean
    rm -rf $(distdir)
    @echo "*** Package $(distdir).tar.gz\
    ready for distribution."
...
```

Now, when we run “`make distcheck`”, our entire build system will be tested before packaging is considered successful. What more could you ask for?!

## Installing products

Well, we've now reached the point where our users' experiences with our project should be fairly painless—even pleasant, as far as building the project is concerned. Our users will simply unpack the distribution tarball, change into the distribution directory, and type “`make`”. It can't really get any simpler than that.

But still we lack one important feature—installation. In the case of the `jupiter` project, this is fairly trivial - there's only one executable, and most users could probably guess that this file should be copied into either the `/usr/bin` or `/usr/local/bin` directory. More complex projects, however could cause our users some real consternation when it comes to where to put user and system binaries, libraries, header files, and documentation, including man pages, info pages, pdf files, and README, INSTALL and COPYRIGHT files. Do we really want our users to have to figure all that out?

I don't think so. So we'll just create an `install` target that manages putting things where they go, once they're built properly. Why not just make installation part of the `all` target? A few reasons, really. First, build and installation are separate logical concepts. Remember the rule: Break up large processes into smaller ones and reuse the smaller ones where you can. The second reason is a matter of rights. Users have rights to build in their own home directories, but installation often requires root-level rights to copy files into system directories. Finally, there are several reasons why a user may wish to build, but not install.

While creating a distribution package may not be an inherently recursive process, installation certainly is, so we'll allow each subdirectory in our project to manage installation of its own components. To do this, we need to modify both makefiles. The top-level makefile is easy. Since there are no products to be installed in the top-level directory, we'll just pass on the responsibility to `src/Makefile` in the usual way:

### Makefile

```
...
all clean install jupiter:
    $(MAKE) -C src $@
...
.PHONY: FORCE all clean dist distcheck install
```

### src/Makefile

```
all: jupiter

jupiter: main.c
    gcc -g -O0 -o $@ $+

clean:
    -rm jupiter

install:
    cp jupiter /usr/bin
    chown root:root /usr/bin/jupiter
    chmod +x /usr/bin/jupiter

.PHONY: all clean install
```

In the top-level makefile, we've added `install` to the list of targets passed down to `src/Makefile`. In both files we've added `install` to the phony target list.

As it turns out, installation was a bit more complex than simply copying files. If a file is placed in the `/usr/bin` directory, then the root user should own it so that only the root user can delete or modify it. Additionally, we should ensure that the `jupiter` binary is executable, so we use the `chmod` command to set the mode of the file to executable. This is probably redundant, as the linker ensures that `jupiter` gets created as an executable file, but it never hurts to be safe.

Now our users can just type the following sequence of commands, and have our project built and installed with the correct system attributes and ownership on their platforms:

```
$ tar -zxvf jupiter-1.0.tar.gz
$ cd jupiter-1.0
$ make all
$ sudo make install
$
```

All of this is well and good, but it could be a bit more flexible with regard to *where* things get installed. Some of our users may be okay with having `jupiter` installed into the `/usr/bin`

directory. Others are going to ask us why we didn't put it into the `/usr/local/bin` directory—after all, it was built locally. Well, we could change the target directory to `/usr/local/bin`, but then others will ask us why we didn't just put it into the `/usr/bin` directory. This is the perfect situation for a little command-line flexibility.

Another problem we have with these makefiles is the amount of stuff we have to do to install files. Most Unix systems provide a system-level program called “`install`”, which allows a user to specify, in an intelligent manner, various attributes of the files being installed. The proper use of this utility could simplify things a bit. We'll just make a few changes to include to use of the `install` utility:

### Makefile

```
...
prefix=/usr/local
export prefix

all clean install jupiter:
    $(MAKE) -C src $@
...
```

### src/Makefile

```
...
install:
    mkdir -p $(prefix)/bin
    install -m 0755 jupiter $(prefix)/bin
...
```

If you're astute, you may have noticed that we declared and assigned the `prefix` variable in the top-level makefile, but we *referenced* it in `src/Makefile`. This is possible because we used the `export` keyword in the top-level makefile to export this make variable to the shell that `make` spawns when it executes itself in the `src` directory. This is a nice feature of `make` because it allows us to define all of our user variables in one obvious location—at the top of the top-level makefile.

We've now declared our `prefix` variable to be `/usr/local`, which is very nice for those who want `jupiter` to be installed in `/usr/local/bin`, but not so nice for those who just want it installed in `/usr/bin`. Fortunately, `make` allows the definition of make variables on the command line, in this manner:

```
$ sudo make prefix=/usr install
...
```

Variables defined on the command line *override* those defined in the makefile. Thus, users who want to install `jupiter` into their `/usr/bin` directory now have the option of specifying this on the `make` command line when they install `jupiter`.

Actually, with this system in place, our users may install `jupiter`

into *any* directory they choose, including a location in their home directory, for which they do not need additional rights granted. This is the reason for the addition of the `mkdir -p` command. We don't actually know *where* the user is going to install jupiter now, so we have to be prepared for the possibility that the location may not yet exist.

A bit of trivia about the `install` utility—it has the interesting property of changing the ownership of any file it copies to the owner and group of the containing directory. So it automatically sets the owner and group of our installed files to `root:root` if the user tries to use the default `/usr/local` prefix, or to the user's id and group if she tries to install into a location within her home directory. Nice, huh?

## Uninstalling a package

What if a user doesn't like our package after it's been installed, and she just wants to get it off her system? This is fairly likely with the jupiter package, as it's rather useless and takes up valuable space in her `bin` directory. In the case of your projects however, it's more likely that she wants to install a newer version of your project cleanly, or she wants to change from the test build she downloaded from your website to a professionally packaged version of your project provided by her Linux distribution. We really should have an `uninstall` target, for these and other reasons:

### Makefile

```
...
all clean install uninstall jupiter:
    $(MAKE) -C src $@
...
.PHONY: FORCE all clean dist distcheck
.PHONY: install uninstall
```

### src/Makefile

```
...
uninstall:
    -rm $(prefix)/bin/jupiter

.PHONY: all clean install uninstall
```

And, again, this particular target will require root-level rights if the user is using a system prefix, such as `/usr` or `/usr/local`. The list of things to maintain is getting a out of hand, if you ask me. We now have two places to update when changing our installation processes—the `install` and `uninstall` targets. Unfortunately, this is really about the best we can hope for when writing our own makefiles, without resorting to fairly complex shell script commands. Hang in there—in Chapter 6, I'll show you how this example can be rewritten in a much simpler way using Automake.

## The Filesystem Hierarchy Standard

By the way, where am I getting these directory names from? What if some Unix system out there doesn't use `/usr` or `/usr/local`? Well, in the first place, this is another reason for providing the `prefix` variable—to handle those sorts of situations. However, most Unix and Unix-like systems nowadays follow the Filesystem Hierarchy Standard (FHS), as closely as possible. The FHS defines a number of “standard places”, including the following root-level directories:

- `/bin`
- `/etc`
- `/home`
- `/opt`
- `/sbin`
- `/srv`
- `/tmp`
- `/usr`
- `/var`

This list is not exhaustive. I've only mentioned the ones most relevant to our purposes. In addition, the FHS defines several standard locations beneath these root-level directories. For instance, the `/usr` directory should contain the following sub-directories:

- `/usr/bin`
- `/usr/include`
- `/usr/lib`
- `/usr/local`
- `/usr/sbin`
- `/usr/share`
- `/usr/src`

The `/usr/local` directory should contain a structure very similar to the `/usr` directory structure, so that if the `/usr/bin` directory (for instance) is an NFS mount, then `/usr/local/bin` (which should always be local) may contain local copies of some programs. This way, if the network is down, the system may still be used to a degree.

Not only does the FHS define these standard locations, but it also explains in fair detail what they are for, and what types of files should be kept there. All in all, the FHS leaves just enough flexibility and choice to you as a project maintainer to keep your life interesting, but not enough to make you lose sleep at night, wondering if you're installing your files in the right places.

Before I found out about the FHS, I relied on my personal experience to decide where files should be installed in my projects. Mostly I was right, because I'm a careful guy, but I have gone back to some of my past projects with a bit of chagrin and changes things once I read the FHS document. I heartily recommend you become thoroughly familiar with this document if you seriously intend to develop open source software.

## Supporting standard targets and variables

In addition to those I've already mentioned, the GNU Coding Standards document lists some important targets and variables that you should support in your projects, mainly because everyone else does and your users will expect them.

Some of the chapters in the GNU Coding Standards should be taken with a grain of salt (unless you're actually working on a GNU sponsored project, in which case, you're probably not reading this book because you need to). For example, you probably won't care much about the C source code formatting suggestions in Chapter 5. Your users certainly won't care, so you can use whatever source code formatting style you wish.

That's not to say that all of Chapter 5 is worthless. Sections 5.5 and 5.6, for instance, provide excellent information on C source code portability between POSIX-oriented platforms and CPU types. Section 5.8 gives some tips on using GNU software to internationalize your program. This is excellent material.

While Chapter 6 discusses documentation the GNU way, some sections of Chapter 6 describe various top-level text files found commonly in projects, such as the AUTHORS, NEWS, INSTALL, README and ChangeLog files. These are all bits that the well-read OSS user expects to see in any decent OSS project.

But, the *really* useful information in the GNU Coding Standards document begins in Chapter 7, "The Release Process". The reason why this chapter is so critical to you as an OSS project maintainer, is that it pretty much defines what your users will expect of your project's build system. Chapter 7 *is* the defacto-standard for user options provided by packages using source-level distribution.

Section 7.1 defines the configuration process, about which we haven't spent much time so far in this chapter, but we'll get to it. Section 7.2 covers makefile conventions, including all of the "standard targets" and "standard variables" that users have come to expect in OSS packages. Standard targets defined by the GNU Coding Standards document include:

- all
- install
- install-html
- install-dvi
- install-pdf
- install-ps
- uninstall
- install-strip
- clean
- distclean
- mostlyclean
- maintainer-clean
- TAGS
- info
- dvi
- html
- pdf
- ps

- `dist`
- `check`
- `installcheck`
- `installdirs`

Note that you don't need to support *all* of these targets, but you should consider supporting those which make sense for your project. For example, if you build and install HTML pages in your project, then you should probably consider supporting the `html` and `install-html` targets. Autotools projects support these, and more. Some of these are useful to users, while others are only useful to maintainers.

Variables that your project should support (as you see fit) include the following. I've added the default values for these variables on the right. You'll note that most of these variables are defined in terms of a few of them, and ultimately only one of them, `prefix`. The reason for this is (again) flexibility to the end user. I call these "prefix variables", for lack of a better name:

```

prefix           = /usr/local
exec-prefix      = $(prefix)
bindir           = $(exec_prefix)/bin
sbindir          = $(exec_prefix)/sbin
libexecdir       = $(exec_prefix)/libexec
datarootdir      = $(prefix)/share
datadir          = $(datarootdir)
sysconfdir       = $(prefix)/etc
sharedstatedir   = $(prefix)/com
localstatedir    = $(prefix)/var
includedir       = $(prefix)/include
oldincludedir    = /usr/include
docdir           = $(datarootdir)/doc/$(package)
infodir          = $(datarootdir)/info
htmldir          = $(docdir)
dvidir           = $(docdir)
pdfdir           = $(docdir)
psdir            = $(docdir)
libdir           = $(exec_prefix)/lib
lisplibdir       = $(datarootdir)/emacs/site-lisp
localedir        = $(datarootdir)/locale
mandir           = $(datarootdir)/man
manNdir          = $(mandir)/manN   (N = 1..9)
manext           = .1
manNext          = .N             (N = 1..9)
srcdir           = (compiled project root)

```

Autotools projects support these and other useful variables automatically. Projects that use Automake get these variables for free. Autoconf provides a mid-level form of support for these variables. If you write your own makefiles and build system, you should support as many of these as you use in your build and install processes.

To support the variables and targets that we've used so far in the jupiter project, we need to add the `bindir` variable, in this manner:

**Makefile**

```
...
prefix = /usr/local
exec_prefix = $(prefix)
bindir = $(exec_prefix)/bin
export prefix exec_prefix bindir
...
```

**src/Makefile**

```
...
install:
    mkdir -p $(bindir)
    install -m 0755 jupiter $(bindir)

uninstall:
    -rm $(bindir)/jupiter
...
```

Note that we have to export `prefix`, `exec_prefix` and `bindir`, even though we only use `bindir` explicitly. The reason for this is that `bindir` is defined in terms of `exec_prefix`, which is itself defined in terms of `prefix`. So when `make` runs the `install` command, it will first resolve `bindir` to `$(exec_prefix)/bin`, and then to `$(prefix)/bin`, and finally to `/usr/local/bin`—`src/Makefile` obviously needs access to all three variables during this process.

How do such recursive variable definitions make life better for the end-user? The user can change the root install location from `/usr/local` to `/usr` by simply typing:

```
$ make prefix=/usr install
...
```

The ability to change these variables like this is particularly useful to a Linux distribution packager, who needs to install packages into very specific system locations:

```
$ make prefix=/usr sysconfdir=/etc install
...
```

**Getting your project into a Linux distro**

The dream of every OSS maintainer is that his or her project will be picked up by a Linux distribution. When a Linux “distro” picks up your package for distribution on their CD’s and DVD’s, your project will be moved magically from the realm of tens of users to that of tens of thousands of users—almost overnight.

By following the GNU Coding Standards with your build system, you remove many barriers to including your project in a Linux distro, because distro packagers (employees of the company, whose job it is to professionally package your project as RPM or APT packages) will immediately know what to do with your tarball if it follows all the usual conventions.

And, in general, packagers get to decide based on needed functionality, and their feelings about your package, whether or not it should be included in their flavor of Linux.

Section 7.2.4 of the GNU Coding Standards talks about the concept of supporting “staged installations”. This is a concept easily supported by a build system, but which if neglected, will almost always cause major problems for Linux distro packagers.

Packaging systems such as the Redhat Package Manager (RPM) system accept one or more tarballs, a set of patches and a specification file (in the case of RPM, called an “rpm spec file”). The spec file describes the process of building and installing your package. In addition, it defines all of the products installed into the targeted installation directory hierarchy. The package manager uses this information to install your package into a temporary directory, from which it pulls the specified binaries, storing them in a special binary archive that the package installer (eg., rpm) understands.

To support staged installation, all you really need to do is provide a variable named “DESTDIR” in your build system that is a sort of super-prefix to all of your installed products. To show you how this is done, I’ll add staged installation support to the jupiter project. This is so trivial, it only requires two changes to `src/Makefile`:

#### **src/Makefile**

```
...
install:
    mkdir -p $(DESTDIR)$(bindir)
    install -m 0755 jupiter $(DESTDIR)$(bindir)
...
```

As you can see, I’ve added the `$(DESTDIR)` prefix to the commands in our `install` target that reference any installation paths. You’ll perhaps have noticed that I didn’t need to change the `uninstall` target. The reason for this is that package managers don’t care how your package is uninstalled, as they only install it so they can copy the products from a temporary install directory, which they then delete entirely after the package is created. Package managers like RPM use their own rules for removing products from a system.

At this point, an RPM spec file (for example) could provide the following text as the installation commands for the jupiter package:

```
%install
make prefix=/usr DESTDIR=%BUILDDIR install
```

But don’t worry about package manager file formats. Just focus on providing staged installation functionality through the `DESTDIR` variable.

You may be wondering why this functionality could not be provided by the `prefix` variable. Well, for one thing, not every path in a system-level installation is defined relative to the `prefix` variable. The system configuration directory

(`sysconfdir`), for instance, is often defined simply as `/etc` by packagers. Defining `prefix` to anything other than `/` will have little effect on `sysconfdir` during staged installation, unless a build system uses `$(DESTDIR)$(sysconfdir)` to reference the system configuration directory. Other reasons for this will become more clear as we talk about project configuration in the next section.

## Build versus installation prefix overrides

At this point, I'd like to digress slightly for just a moment to explain an illusive (or at least non-obvious) concept regarding the prefix and other path variables defined by the GNU Coding Standards document.

In the preceeding examples, I've always used prefix overrides on the `make install` command line, like this:

```
$ make prefix=/usr install
...
```

The question I wish to address is: What's the difference between using a prefix override for `make all` and `make install`? In our small sample makefiles, we've managed to avoid using prefixes in any targets not related to installation, so it may not be clear at this point that a prefix is *ever* useful during the build stages.

One key use of prefix variables during the build stage is to substitute paths into source code at compile time, in this manner:

```
main.o : main.c
        gcc -DCFGDIR=\"$(sysconfdir)\" -o $@ $+
```

In this example, I'm defining a C preprocessor variable called `CFGDIR` on the compiler command line for use by `main.c`. Presumably, there's some code in `main.c` that looks like this:

```
#ifndef CFGDIR
# define CFGDIR "/etc"
#endif

char cfgdir[FILENAME_MAX] = CFGDIR;
```

Later in the code, the C global variable `cfgdir` might be used to access the application's configuration file.

Okay, with that background then, would you ever want to use *different* prefix variable overrides on the build and installation command lines? Sure—Linux distro packagers do this all the time in RPM spec files. During the build stage, the actual runtime directories are hard-coded into the executable by using a command like this:

```
%build
%setup
./configure prefix=/usr sysconfdir=/etc
make
```

The RPM build process installs these executables into a stage directory, so it can copy them out. The corresponding installation command looks like this:

```
%install
rm -rf %BUILDDROOT%
make DESTDIR=%BUILDDROOT% install
```

I mentioned the DESTDIR variable previously as a tool used by packagers for staged installation. This has the same effect as using:

```
%install
rm -rf %BUILDDROOT%
make prefix=%BUILDDROOT%/usr \
  sysconfdir=%BUILDDROOT%/etc install
```

The key take-away point here is this: Never recompile from an `install` target in your makefiles. Otherwise your users won't be able to access your staged installation features when using prefix overrides.

Another reason for is to allow the user to install into a grouped location, and then create links to the actual files in the proper locations. Some people like to do this, especially when they are testing out a package, and want to keep track of all of its components. For example, some Linux distributions provide a way of installing multiple versions of some common packages. Java is a great example here. To support using multiple versions or brands (perhaps Sun Java vs IBM Java), the Linux distribution provides a script set called the "alternatives" scripts, which allows a user (running as root) to swap all of the links in the various system directories from one grouped installation to another. Thus, both sets of files may be installed in different auxiliary locations, but links in the true installation locations can be changed to refer to each group at different times.

One final point about this issue. If you're installing into a system directory hierarchy, you'll need root permissions. Often people run `make install` like this:

```
$ sudo make install
...
```

If your `install` target depends on your build targets, and you've neglected to build beforehand, then `make` will happily build your program before installing it, but the local copies will all be owned by `root`. Just an inconvenience, but easily avoided by having `'make install'` fail for lack of things to install, rather than simply jump right into a build while running as root.

## Standard user variables

There's one more topic I'd like to cover before we move on to configuration. The GNU Coding Standards document defines a set of variables that are sort of sacred to the user. That is, these variables should be used by a GNU build system, but never modified by a GNU build system. These are called "user variables", and they include the following for C and C++

programs:

```
CC          - the C compiler
CFLAGS     - C compiler flags
CXX        - the C++ compiler
CXXFLAGS   - C++ compiler flags
LDFLAGS    - linker flags
CPPFLAGS   - C preprocessor flags
...
```

This list is by no means comprehensive, and ironically, there isn't a comprehensive list to be found in the GCS document. Interestingly, most of these user variables come from the documentation for the make utility. You can find a fairly complete list of program name and flag variables in section 10.3 of the GNU make manual. The reason for this is that these variables are used in the built-in rules of the make utility.

For our purposes, these few are sufficient, but for a more complex makefile, you should become familiar with the larger list so that you can use them as the occasion arises. To use these in our makefiles, we'll just replace "gcc" with `$(CC)`, and then set `CC` to the gcc compiler at the top of the makefile. We'll do the same for `CFLAGS` and `CPPFLAGS`, although this last one will contain nothing by default:

#### **src/Makefile**

```
...
CC      = gcc
CFLAGS = -g -O2
...
jupiter: main.c
        $(CC) $(CFLAGS) $(CPPFLAGS) -o $@ $+
...
```

The reason this works is that the make utility allows such variables to be overridden by the environment. Environment and make command-line variable assignments always override values set in the makefiles themselves. Thus, to change the compiler and set some compiler flags, a user need simply type:

```
$ make CC=gcc3 CFLAGS=-g -O0 CPPFLAGS=-dtest
```

In this case, our user has decided to use gcc version 3 instead of 4, and to disable optimization and leave the debugging symbols in place. She's also decided to enable the "test" option through the use of a preprocessor definition. These variables are set on the make command line, but may also be exported in the environment before running make.

## **Configuring your package**

The GNU Coding Standards document describes the configuration process in section 7.1, "How Configuration Should Work". Up to this point, we've been able to do about everything we've wanted to do with the jupiter project using only makefiles. You might be wondering at this point what

configuration is actually for! The opening paragraphs of Section 7.1 state:

Each GNU distribution should come with a shell script named `configure`. This script is given arguments which describe the kind of machine and system you want to compile the program for.

The `configure` script must record the configuration options so that they affect compilation.

One way to do this is to make a link from a standard name such as `config.h` to the proper configuration file for the chosen system. If you use this technique, the distribution should not contain a file named `config.h`. This is so that people won't be able to build the program without configuring it first.

Another thing that `configure` can do is to edit the makefiles. If you do this, the distribution should not contain a file named `Makefile`. Instead, it should include a file `Makefile.in` which contains the input used for editing. Once again, this is so that people won't be able to build the program without configuring it first.

So then, the primary tasks of a typical `configure` script are to:

- generate files from templates containing replacement variables,
- generate a C language header file (often called `config.h`) for inclusion by project source code,
- set user options for a particular `make` environment—such as debug flags, etc.,
- set various package options as environment variables,
- and test for the existence of tools, libraries, and header files.

For complex projects, `configure` scripts often generate the project makefile(s) from one or more templates maintained by the project manager. A makefile template contains configuration variables in an easily recognized (and substituted) format. The `configure` script replaces these variables with values determined during configuration—either from command line options specified by the user, or from a thorough analysis of the platform environment. Often this analysis entails such things as checking for the existence of certain include files and libraries, searching various file system paths for required utilities and tools, and even running small programs designed to indicate the feature set of the shell, C compiler, or desired libraries.

The tool of choice here for variable replacement has, in the past, been the `sed` stream editor. A simple `sed` command can replace all of the configuration variables in a makefile template in a single pass through the file. In the latest version of Autoconf (2.62, as of this writing) prefers `awk` to `sed` for this process. The `awk` utility is *almost* as pervasive as `sed`

these days, and it much more powerful with respect to the operations it can perform on a stream of data. For the purposes of the jupiter project, either one of these tools would suffice.

## Summary

At this point, we've created a complete project build system by hand—with one important exception. We haven't designed a `configure` script according to the design criteria specified in the GNU Coding Standards document that works with this build system. We could do this, but it would take a dozen more pages of text to build one that even comes close to conforming to these specifications.

There are yet a few key build system features related specifically to the makefiles that are indicated as being desirable by the GNU Coding Standards. Among these is the concept of VPATH building. This is an important feature that can only be properly illustrated by actually writing a `configure` script that works as specified by the GNU Coding Standards.

Rather than spend this time and effort, I'd like to simply move on to a discussion of Autoconf in Chapter 3, which will allow us to build one of these `configure` scripts in as little as two or three lines of code, as you'll see in the opening paragraphs of that chapter. With that step behind us, it will be trivial to add VPATH building, and other features, to the jupiter project.

[◀ Chapter 1: A brief introduction to the GNU Autotools](#)

u  
p

[Chapter 3: Configuring your project with autoconf ▶](#)

[Login or register to post comments](#)

3158 reads



- COLUMNS
- COMMUNITY POSTS
- ISSUES
- BOOKS
- FORUM
- FS NEWS

Ads by Google

**Linux**  
 Make The Web Tier A Cooler Place. Try Sun Hardware For Free! Get Started Now  
[uk.sun.com](http://uk.sun.com)

**Ubuntu linux**  
 Multi-platform Remote Access. No installation necessary. Free Trial!  
[www.Bomgar.com/linux](http://www.Bomgar.com/linux)

**Server Monitoring Tools**  
 Download Our Server Monitoring Software Risk-Free. Get Started Now  
[www.uptimesoftware.com](http://www.uptimesoftware.com)

**Full SUSE 11.0 DVD £6**  
 Plus 2 additional full distro DVDs Don't waste time downloading Linux!  
[linux-magazine.co.uk](http://linux-magazine.co.uk)

**Core 2 Duo Servers £49**  
 Unlimited bandwidth.Windows/Linux. Full remote control, Private LAN.  
[www.Fasthosts.co.uk/DedicatedSer](http://www.Fasthosts.co.uk/DedicatedSer)

- [Microsoft May Seek New Yahoo ...](#)
- [Building Semantics is Different ...](#)
- [Lux: multi-touch for OS X](#)

Sponsored by **Novell**.

Best voted contents

- [The Bizarre Cathedral - 3](#)  
 Ryan Cartwright, 2008-05-05
- [The Bizarre Cathedral - 2](#)  
 Ryan Cartwright, 2008-04-27
- [Indexing offline CD-ROM archives](#)  
 Terry Hancock, 2008-05-03
- [Microsoft and free software? I don't think so...](#)  
 Terry Hancock, 2008-04-26

Home » [Autotools: a practitioner's guide to Autoconf, Automake and Libtool](#)

# Chapter 3: Configuring your project with autoconf



by [John Calcote](#)

We should all be very grateful to David MacKenzie for having the foresight to—metaphorically speaking—stop and sharpen the ax. Otherwise we'd still be writing (copying) and maintaining long, complex hand-coded `configure` scripts today.

Before Automake, Autoconf was used alone, and many legacy open source projects have never really made the transition to the full Autotools suite. As a result, it would not be uncommon to find an open source project containing a file called `configure.in` (the older naming convention used by Autoconf) and hand-written `Makefile.in` templates.

## Configure scripts, the Autoconf way

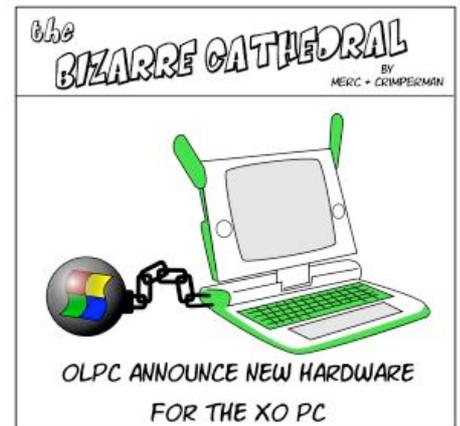
It's instructive for this and other reasons that will become clear shortly, to spend some time just focusing on the use of Autoconf alone. Exploring in this manner can provide a fair amount of insight into the operation of Autoconf by exposing aspects of this tool that are often hidden by Automake and other add-on tools.

The input to Autoconf is ... (drum roll please) ... shell script. Man, what an anti-climax! Okay, so it's not pure shell script. That is, it's shell script with macros, plus a bunch of macro definition files—both those that ship with an Autoconf distribution, as well as those that you or I write. The macro language used is called M4. ("M-what?!", you ask?) The M4 utility is a general purpose macro language processor that was originally written by none other than Brian Kernighan and Dennis Ritchie in 1977. (The name M4 means "m plus 4 more letters" or the word "Macro" - cute, huh? As a point of interest, this naming convention is a fairly common practice in some software engineering domains. For example, the term *internationalization* is often abbreviated *i18n*, and the term *localization* is sometimes replaced with *l10n*, for the sake of brevity. The use of the term *m4* here is no-doubt a play on this concept.)

Some form of the M4 macro language processor is found on every Unix and Linux variant (as well as other systems) in use today. In fact, this proliferation is the primary reason for its use in Autoconf. The design goals of Autoconf included primarily that it should run on all systems without the addition of

**PostgreSQL Replication**  
 Slony, Warm Standby & Multi-master Training from the Developers  
[www.2ndQuadrant.com](http://www.2ndQuadrant.com)

**VMware Certified Training**  
 VMware VCP courses. We train the the trainers, and VMware themselves  
[www.magirus-training.co.uk](http://www.magirus-training.co.uk)



<http://www.freesoftwaremagazine.com> (c) Copyright 2008 Ryan Cartwright CC: BY-NC-SA

From the FSM staff...

[The Top 10 Everything](#) (Dave). The good, the bad and the ugly.

[Free Software news](#) (Dave & Bridget). A site about short stories and writing.

[Book Reviews: Illiteraty](#) (Bridget). Book reviews, blogs, and short stories.

Hot topics - last 60 days

[Installing an all-in-one printer device in Debian](#)  
 Ryan Cartwright, 2008-05-05

[What is the free software community?](#)  
 Tony Mobily, 2008-03-29

[Things you miss with](#)

Free Open Source Book Samples...  
 Exclusive discount offer to FSM Readers  
**PACKT**

Download it.  
 It's open source.  
 It's yours.

**VYATTA**  
 FREE Open Source ROUTER FIREWALL VPN

complex tool chains and utility sets. Autoconf depends on the existence of relatively few tools, including M4, `sed` and now in version 2.62, the `awk` utility. Most of the Autotools (Autoconf being the exception) rely on the existence of a perl processor, as well.

*NOTE: Do not confuse the requirements of the Autotools with the requirements of the scripts and makefiles generated by them. The Autotools are maintainer tools, while the resulting scripts and makefiles are end-user tools. We can reasonably expect a higher level of installed functionality on development systems than we can on end-user systems. Nevertheless, the Autotools design goals still include a reliance only on a minimal set of pre-installed functionality, much of which is part of a default installation.*

While it's true that `configure.ac` is written in shell script sprinkled with M4 syntax, the proper use of the M4 macro processor is the subject of Chapter 7. Because I want to stick to Autoconf in this chapter, I'll gloss over some key concepts related to M4, which I'll cover in more detail in Chapter 7. This chapter is designed to help you understand Autoconf concepts, however, so I will cover minor aspects of M4 as it makes sense to do so.

## The smallest `configure.ac` file

The simplest possible `configure.ac` file has just two lines:

```
$ cat configure.ac
AC_INIT([jupiter], [1.0])
AC_OUTPUT
$
```

*NOTE: This chapter builds on the Jupiter project begun in Chapter 2.*

To those new to Autoconf, these two lines appear to be a couple of function calls, perhaps in the syntax of some obscure computer language. Don't let this appearance throw you—these are M4 macro expansions. The macros are defined in files distributed with Autoconf. The definition of `AC_INIT` is found in `$PREFIX/share/autoconf/autoconf/general.m4`, while `AC_OUTPUT` is defined in `status.m4`, in the same directory.

M4 macros are similar in many ways to macros defined in C language source files for the C preprocessor, which is also a text replacement tool. This isn't surprising, given that both M4 and `cpp` were originally designed by Kernighan and Ritchie.

The square brackets around the parameters are used by Autoconf as a quoting mechanism. Such quotes are only really necessary in cases where the context of the macro call could cause an ambiguity that the macro processor may resolve incorrectly (usually without telling you). We'll discuss M4 quoting in much more detail in Chapter 7. For now, just use Autoconf quotes (`[` and `]`) around every argument to ensure that the expected macro expansions are generated.

## Buzz authors

[Marco Marongiu](#)

[Pieter Hintjens](#)

[drascus321](#)

[dogboi](#)

[Richard Rothwell](#)

## All news

[John Crawford: Arizona LoCo Hardy Release Party](#)

[John Crawford: Ubuntu Weekly Newsletter #91](#)

[Matthew Helmke: Morocco blocks Google Maps](#)

[Andy Price: Synchronicity and Sharing](#)

[Jono Bacon: Beware False Prophets](#)

[more](#)



## GNU/Linux

[Ryan Cartwright, 2008-05-01](#)

[How do you replace Microsoft Outlook? Groupware applications](#)

[Ryan Cartwright, 2008-03-20](#)

[Drigg \(the pligg alternative\) vs. Pligg: why should people switch?](#)

[Tony Mobily, 2008-04-13](#)

## Hot topics - last 21 days

[Installing an all-in-one printer device in Debian](#)

[Ryan Cartwright, 2008-05-05](#)

[Things you miss with GNU/Linux](#)

[Ryan Cartwright, 2008-05-01](#)

[Digital Rights Management \(DRM\): is it in its death throes?](#)

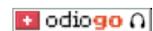
[Gary Richmond, 2008-05-07](#)

[Open letter to standards professionals, developers, and activists](#)

[Pieter Hintjens, 2008-05-13](#)



[Odiogo](#)



- 15 [How Microsoft Uses Novell to Fight GNU/Linux, Xen to Fight VMWare and GNU/Linux](#)
- 15 [Why You Should Reject Novell's Moonlight](#)
- 17 [University of Havana Finally Switches to Free Software](#)
- 15 [Please Welcome Digistan](#)
- 17 [Content Protection madness on Vista](#)
- 25 [Asus to embed Linux into all motherboards](#)
- 23 [Make Your Distro Free of Miguel de Icaza's junk code](#)
- 16 [Firefox 3 Release Candidate now available for download](#)

[Get this widget »](#)

As with `cpp` macros, M4 macros may or may not take parameters. And (also as with `cpp`) when they do, then a set of parentheses must be used when passing the arguments. In both M4 and `cpp`, the opening parenthesis must immediately follow the macro name, with no intervening white space. When they don't accept parameters, the parenthesis are simply omitted. Unlike `cpp`, M4 has the ability to specify *optional* parameters, in which case, you may omit the parenthesis if you choose not to pass a parameter.

The result of passing this `configure.ac` file through Autoconf is essentially the same file (now called `configure`), only with these two macros fully expanded.

Now, if you've been programming in C for many years, as I have, then you've no doubt run across a few C preprocessor macros from the dark regions of the lower realm. I'm talking about those truly evil `cpp` macros that expand into one or two pages of C code! You know the ones I'm talking about—they should really have been written as C functions, but the author was overly worried about performance!

Well baby, you ain't seen *nothin'* yet! These two M4 macros expand into a file containing over 2200 lines of Bourne shell script that's over 60K bytes in size! Interestingly, you wouldn't really know this by looking at their definitions. They're both fairly short—only a dozen or two lines each. The reason for this apparent disparity is simple—they're written in a modular fashion, each macro expanding several others, which in turn expand several others, and so on.

## Executing Autoconf

Running Autoconf couldn't be simpler. Just execute `autoconf` in the same directory as your `configure.ac` file. While we *could* do this for each example in this chapter, we're going to use the `autoreconf` (capitalization added for emphasis) command instead of the `autoconf` command. The reason for this is that running `autoreconf` has exactly the same effect as running `autoconf`, except that `autoreconf` will also do "the right thing" when you start adding Automake and Libtool functionality to your build system. `autoreconf` is the recommended method for executing the Autotools tool chain, and it's smart enough to only execute the tools that you need, in the order that you need them, and with the options that you need (with one exception that I'll mention shortly here).

```
$ autoreconf
$ ls -lp
autom4te.cache/
configure
configure.ac
$
```

First, notice that `autoreconf` operates at exactly the same level of verbosity as the tools it runs. By default, zero. If you want to see something happening, use the `-v` or `--`

verbose option. If you want `autoreconf` to run the other Autotools in verbose mode, add `-vv` to the command line. (You may also pass `--verbose --verbose`, but this syntax seems a bit verbose to me—sorry, I couldn't resist!)

First, notice that Autoconf creates a directory called `autom4te.cache`. This is the `autom4te` (pronounced “automate”) cache directory. This cache is used to speed up access to `configure.ac` by successive executions of utilities in the Autotools tool chain. I'll cover `autom4te` in greater detail in Chapter 9, where I'll show you how to write your own Autoconf macros that are “environmentally friendly”.

## Executing `configure`

If you recall from the last section of Chapter 2, the GNU Coding Standards document indicates that `configure` should generate a script called `config.status`, whose job it is to generate files from templates. Well, this is exactly the sort of functionality found in an Autoconf-generated `configure` script. An `autoconf configure` script has two primary tasks:

- perform requested checks
- generate, and then call `config.status`

The results of all of the checks performed by the `configure` script are written, as environment variable settings to the top of `config.status`, which uses the values in these environment variables as replacement text for Autoconf substitution variables used in template files (`Makefile.in`, `config.h.in`, etc).

When you execute `configure`, it tells you that it's creating the `config.status` file. In fact, it also creates a log file called `config.log` that has several important attributes:

```
$ ./configure
configure: creating ./config.status
$
$ ls -lp
autom4te.cache/
config.log
config.status
configure
configure.ac
$
```

The `config.log` file contains the following information:

- the command line used to invoke `configure` (Very handy!)
- information about the platform on which `configure` was executed
- information about the core tests executed by `configure`
- the line number in `configure` at which `config.status` is generated

At this point in the log file, `config.status` takes over

generating log information—it adds the command line used to invoke `config.status`. After `config.status` generates all of the files from their templates, control then returns to `configure`, which adds the following information to the log:

- the cache variables used by `config.status` to perform its tasks
- the list of output variables that may be replaced in templates
- the exit code returned by `configure` to the shell

This information is invaluable when debugging a `configure` script and its associated `configure.ac` file.

## Executing `config.status`

Now that you know how `configure` works, you can probably see that there might be times when you'd be tempted to simply execute `config.status` yourself, rather than going to all the trouble of having `configure` perform all those time-consuming checks first. And right you'd be. This was exactly the intent of the Autoconf designers—and the authors of the GNU Coding Standards, by whom these design goals were originally conceived.

There are in fact, times when you'd just like to manually regenerate all of your output files from their corresponding templates. But, far more importantly, `config.status` can be used by your makefiles to regenerate themselves individually from their templates, when `make` determines that something in a template file has changed.

Rather than call `configure` to perform needless checks (your environment hasn't changed, has it? Just your template files), your makefiles should be written in a way that ensures that output files are dependent on their templates. If a template file changes (because, for example, you modified one of your `Makefile.in` templates), then `make` calls `config.status` to regenerate this file. Once the `Makefile` is regenerated, then `make` re-executes the original `make` command line—basically, it restarts itself.

Let's take a look at the relevant portion of just such a `Makefile.in` template:

```
...
Makefile: Makefile.in config.status
    ./config.status Makefile
...
```

An interesting bit of `make` functionality is that it always looks for a rule with a target named "`Makefile`". Such a rule allows `make` to regenerate the source makefile from its template, in the event that the template changes. It does this *before* executing either the user's specified targets, or the default target, if none was given.

This example indicates that its makefile is dependent on a `Makefile.in` template. Note that `Makefile` is also

dependent on `config.status`. After all, if `config.status` is regenerated by the `configure` script, then it may generate a makefile differently—perhaps something in the compilation environment changed, such as when a new package is added to the system, so that `configure` can now find libraries and headers not previously found. In this case, Autoconf substitution variables may have different values. Thus, `Makefile` should be regenerated if either `Makefile.in` or `config.status` changes.

Since `config.status` is itself a generated file, it stands to reason that this line of thinking can be carried to the `configure` script as well. Expanding on the previous example:

```
...
Makefile: Makefile.in config.status
    ./config.status $@

config.status: configure
    ./config.status --recheck
...
```

Since `config.status` is a dependency of the `Makefile` rule, then `make` will check for a rule whose target is `config.status` and run its commands if the dependencies of `config.status` (`configure`) are newer than `config.status`.

## Adding some real functionality

Well, it's about time we move forward and put some true functionality into this `configure.ac` file. I've danced around the topic of having `config.status` generate a makefile up to this point. Here's the code to actually make this happen in `configure.ac`. It constitutes a single additional macro expansion between the original two lines:

```
$ cat configure.ac
AC_INIT([jupiter], [1.0])
AC_CONFIG_FILES([Makefile
                 src/Makefile])
AC_OUTPUT
$
```

This code assumes we have templates for `Makefile` and `src/Makefile`, called `Makefile.in` and `src/Makefile.in`, respectively. These files look exactly like their `Makefile` counterparts, with one exception: Any text we want Autoconf to replace should be marked as Autoconf substitution variables, using the `@xxxxx@` syntax.

To create these files, I've merely renamed the existing makefiles to `Makefile.in` in the top-level and `src` directories. By the way, this is a common practice when "autoconfiscating" a project. Next, I added a few Autoconf substitution variables to replace our original default values. I've also added the makefile regeneration rules from above to each

of these templates, with slight file path differences to account for their different positions relative to `config.status` and `configure`:

### Makefile.in

```
# Makefile: generated from Makefile.in by autoconf

# Package-related substitution variables
package      = @PACKAGE_NAME@
version      = @PACKAGE_VERSION@
tarname      = @PACKAGE_TARNAME@
distdir      = $(tarname)-$(version)

# Prefix-related substitution variables
prefix       = @prefix@
exec_prefix  = @exec_prefix@
bindir       = @bindir@

all clean check install uninstall jupiter:
    $(MAKE) -C src $@

dist: $(distdir).tar.gz

$(distdir).tar.gz: FORCE $(distdir)
    tar chof - $(distdir) | \
    gzip -9 -c >$(distdir).tar.gz
    rm -rf $(distdir)

$(distdir):
    mkdir -p $(distdir)/src
    cp configure $(distdir)
    cp Makefile.in $(distdir)
    cp src/Makefile.in $(distdir)/src
    cp src/main.c $(distdir)/src

distcheck: $(distdir).tar.gz
    gzip -cd $+ | tar xvf -
    cd $(distdir); ./configure
    $(MAKE) -C $(distdir) all check clean
    rm -rf $(distdir)
    @echo "*** Package $(distdir).tar.gz\
    ready for distribution."

Makefile: Makefile.in config.status
    ./config.status $@

config.status: configure
    ./config.status --recheck

FORCE:
    -rm $(distdir).tar.gz &> /dev/null
    -rm -rf $(distdir) &> /dev/null

.PHONY: FORCE all clean check dist distcheck
.PHONY: install uninstall
```

**src/Makefile.in**

```

# Makefile: generated from Makefile.in by autoconf

# Package-related substitution variables
package      = @PACKAGE_NAME@
version      = @PACKAGE_VERSION@
tarname      = @PACKAGE_TARNAME@
distdir      = $(tarname)-$(version)

# Prefix-related substitution variables
prefix       = @prefix@
exec_prefix  = @exec_prefix@
bindir       = @bindir@

all: jupiter

jupiter: main.c
        gcc -g -O0 -o $@ $+

clean:
        -rm jupiter

check: all
        ./jupiter | grep "Hello from .*jupiter!"
        @echo "*** All TESTS PASSED"

install:
        mkdir -p $(bindir)
        install -m 0755 jupiter $(bindir)

uninstall:
        -rm $(bindir)/jupiter

Makefile: Makefile.in ../config.status
        cd .. && ../config.status $@

../config.status: ../configure
        cd .. && ../config.status --recheck

.PHONY: all clean check install uninstall

```

I've removed the `export` statement in the top-level `Makefile.in`, and added a copy of all of the substitution variables into `src/Makefile.in`. Since `config.status` is generating both of these files, we can reap excellent benefits by substituting everything into both files. The primary advantage of doing this is that we can now run `make` in any sub-directory, and not be concerned about environment variables that would have been passed down by a higher-level makefile.

I've also changed the distribution targets a bit. Rather than distribute makefiles, we now want to distribute `Makefile.in` templates, as well as the `configure` script. In addition, the `distcheck` target needs to be updated so that it runs the `configure` script before running `make`.

## Generating files from templates

I'm now generating makefiles from `Makefile.in` templates. The fact is, however, that *any* (white space delimited) file you list in `AC_CONFIG_FILES` will be generated from a file of the same name with a `.in` extension, and found in the same directory. The `.in` extension is the default template naming pattern for `AC_CONFIG_FILES`, but this default behavior may be overridden, if you wish. I'll get into the details shortly.

Autoconf generates `sed` or `awk` expressions into the resulting `configure` script, which then copies them into the `config.status` script. The `config.status` script uses these tools to perform this simple string replacement.

Both `sed` and `awk` are text processing tools that operate on file streams. The advantage of a stream editor (`sed` is actually named after the concept of Stream EDitor) is that it replaces text patterns in a byte stream. Thus, both `sed` and `awk` can operate on huge files, because they don't need to load an entire file into memory in order to process it. The expression list passed to `sed` or `awk` by `config.status` is built by Autoconf from a list of variables defined by various macros, many of which we'll cover in greater detail in this chapter.

The important thing to notice here is that the Autoconf variables are the *only* items replaced in `Makefile.in` while generating the makefile. The reason this is important is that it helps you to realize the flexibility you have when allowing Autoconf to generate a file from a template. This flexibility will become more apparent as we get into various use cases for the pre-defined Autoconf macros, and later in Chapter 9 when we delve into the topic of writing your own Autoconf macros.

At this point, we've created a basic `configure.ac` file, and we can indeed run `autoreconf`, the generated `configure` script, and then `make` to build the jupiter project.

The idea that I want to promote at this point is that *this simple three-line `configure.ac` file generates a `configure` script that is fully functional, according to the definition of a `configure` script given in Chapter 7 of the the GNU Coding Standards document*. The resulting `configure` script runs various system checks and generates a `config.status` file, which can replace a fair number of substitution variables in a set of specified template files in a build system. That's a lot of stuff for three lines of code. (You'll recall my comments in the introduction to this book about C++ doing a lot for you with just a few lines of code?)

## Adding VPATH build functionality

Okay, you may recall at the end of Chapter 2, I mentioned that we hadn't yet covered a key concept—that of VPATH builds. A VPATH build is a way of using a construct supported by `make` (VPATH) to configure and build your project in a directory other than the source directory. Why is this important? Well, for several reasons. You may need to:

1. maintain a separate debug configuration,
2. test different configurations, side by side,
3. keep a clean source directory for patch diffs after local modifications,
4. or build from a read-only source directory.

These are all great reasons, but won't we have to change our entire build system to support this type of remote build? As it turns out, it's quite simple using the `VPATH` statement. `VPATH` is short for virtual path, meaning virtual search path. A `VPATH` statement contains a colon-separated list of places to look for dependencies, when they can't be found relative to the current directory:

```
VPATH = some/path:some/other/path:yet/another/path
```

```
jupiter : main.c
        gcc ...
```

In this (contrived) example, if `make` can't find `main.c` in the current directory while processing the rule, it will look for `some/path/main.c`, and then for `some/other/path/main.c`, and finally for `yet/another/path/main.c`, before finally giving up in despair—okay, perhaps only with an error message about not knowing how to make `main.c`.

"Nice feature!", you say? Nicer than you think, because with just a couple of lines of additional code and a few simple modifications, I can now completely support remote builds in my jupiter project build system:

### **Makefile.in**

```
...
# VPATH-related substitution variables
srcdir      = @srcdir@
VPATH       = @srcdir@
...
$(distdir):
    -rm -rf $(distdir) &&> /dev/null
    mkdir -p $(distdir)/src
    cp $(srcdir)/configure $(distdir)
    cp $(srcdir)/Makefile.in $(distdir)
    cp $(srcdir)/src/Makefile.in $(distdir)/src
    cp $(srcdir)/src/main.c $(distdir)/src
...
Makefile: Makefile.in config.status
    ./config.status $@

config.status: configure
    ./config.status --recheck
...
```

### **src/Makefile.in**

```
...
# VPATH-related substitution variables
srcdir      = @srcdir@
```

```

VPATH          = @srcdir@
...
jupiter: main.c
               gcc -g -O0 -o $@ $(srcdir)/main.c
...
Makefile: Makefile.in ../config.status
               cd .. && ../config.status src/$@

../config.status: ../configure
               cd .. && ../config.status --recheck
...

```

That's it. Really. When `config.status` generates a file, it replaces an Autoconf substitution variable called `@srcdir@` with the relative path to the template's source directory. Each makefile will get a different value for `@srcdir@`, depending on the relative location of its template.

The rules then for supporting VPATH builds in your make system are as follows:

1. Set a make variable, `srcdir` to the `@srcdir@` substitution variable.
2. Set `VPATH` to `$(srcdir)`.
3. Prefix all file dependencies used *in commands* with `$(srcdir)/`.

If the source directory is the same as the build directory, then the `@srcdir@` substitution variable degenerates to `.`, so all of these `$(srcdir)/` prefixes will degenerate to `./`, which is just so much harmless baggage.

A quick example is the easiest way to show you how this works. Now that `jupiter` is fully functional with respect to VPATH builds, let's just give it a try. Start in the `jupiter` project directory, create a subdirectory called `build`, and then change into that directory. Now run `configure` using a relative path, and then list the current directory contents:

```

$ pwd
../prj/jupiter
$ mkdir build
$ cd build
$ ../configure
configure: creating ./config.status
config.status: creating Makefile
config.status: creating src/Makefile
$ ls -lp
config.log
config.status
Makefile
src/
$

```

Our entire build system seems to have been constructed by `configure` and `config.status` within the `build` subdirectory, just as it should be. What's more, it actually works:

```

$ pwd
../prj/jupiter/build

```

```

$ make
make -C src all
make[1]: Entering directory `../prj/jupiter/bui...
gcc -g -O2 -o jupiter ../../src/main.c
make[1]: Leaving directory `../prj/jupiter/bui...
$
$ ls -lp src
jupiter
Makefile
$

```

VPATH builds work, not just from sub-directories of the project directory, but from anywhere you can access the project directory, using either a relative or an absolute path. This is just one more thing that Autoconf does for you in Autoconf-generated `configure` scripts. Just imagine managing proper relative paths to source directories in your own hand-coded `configure` scripts!

### Let's take a breather

At this point, I'd like you to stop and consider what you've seen so far: I've shown you a mostly complete build system that includes most of the features outlined in the GNU Coding Standards document. The features of the jupiter project's make system are all fairly self-contained, and reasonably simple to grasp. The most difficult feature to implement by hand is the `configure` script. In fact, writing a `configure` script by hand is so labor intensive relative to the simplicity of the Autoconf version that I just skipped over the hand-coded version entirely.

If you've been one to complain about Autoconf in the past, I'd like you to consider what you have to complain about now. You now know how to get very feature-rich configuration functionality *in just three lines of code*. Given what you know now about how `configure` scripts are meant to work, can you see the value in Autoconf?

Most people never have trouble with that portion of Autoconf that I've covered up to this point. The trouble is that most people don't create their build systems in the manner I've just shown you. They try to copy the build system of another project, and then tweak it to make it work in their own project. Later when they start a new project, they do the same thing again. Are they going to run into problems? Sure—the stuff they're copying was often never meant to be used the way they're trying to use it.

I've seen projects in my experience whose `configure.ac` file contained junk that had nothing to do with the project to which it belonged. These left-over bits came from the previous project, from which `configure.ac` was copied. But the maintainer didn't know enough about Autoconf to remove the cruft. It's better to start small, and add what you need, than to start with a full-featured build system, and try to pare it down to size.

Well, I'm sure you're feeling like there's a lot more learn about

Autoconf. And you're right, but what additional Autoconf macros are appropriate for the jupiter project?

## An even quicker start with autoscan

The simplest way to create a (mostly) complete `configure.ac` file is to run the `autoscan` utility, which, if you remember from Chapter 1, is part of the Autoconf package.

First, we'll clean up the droppings from our earlier experiments, and then run the `autoscan` utility in the `jupiter` directory. Note here that I'm NOT deleting my original `configure.ac` file - I'll just let `autoscan` tell me what's wrong with it. In less than a second I'm left with a couple of new files in the top-level directory:

```
$ rm config.* Makefile src/Makefile ...
$
$ ls -lp
configure.ac
Makefile.in
src/
$
$ autoscan
configure.ac: warning: missing AC_CHECK_HEADERS
    ([stdlib.h]) wanted by: src/main.c:2
configure.ac: warning: missing AC_HEADER_STDC
    wanted by: src/main.c:2
configure.ac: warning: missing AC_PROG_CC
    wanted by: src/main.c
configure.ac: warning: missing AC_PROG_INSTALL
    wanted by: Makefile.in:11
$
$ ls -lp
autom4te.cache/
autoscan.log
configure.ac
configure.scan
Makefile.in
src/
$
```

*NOTE: I've wrapped some of the output lines for the sake of column width during publication.*

`autoscan` creates two files called `configure.scan`, and `autoscan.log` from a project directory hierarchy. The project may already be instrumented for Autotools, or not. It doesn't really matter because `autoscan` is decidedly non-destructive. It will never alter any existing files in a project.

`autoscan` generates a warning message for each issue discovered in an existing `configure.ac` file. In this example, `autoscan` noticed that `configure.ac` really should be using the `AC_CHECK_HEADERS`, `AC_HEADER_STDC`, `AC_PROG_CC` and `AC_PROG_INSTALL` macros. It made these assumptions based on scanning our existing `Makefile.in` templates

and C source files, as you can see by the comments after each warning statement. You can always see these messages (in even greater detail) by examining the `autoscan.log` file.

Now let's examine the generated `configure.scan` file. `autoscan` has added more text to `configure.scan` than was originally in our `configure.ac` file, so it's probably easier for us to just overwrite `configure.ac` with `configure.scan` and then change the few bits of information that are specific to jupiter:

```
$ mv configure.scan configure.ac
$
$ cat configure.ac
#          -*- Autoconf -*-
# Process this file with autoconf to produce ...

AC_PREREQ(2.61)
AC_INIT(FULL-PACKAGE-NAME, VERSION,
        BUG-REPORT-ADDRESS)
AC_CONFIG_SRCDIR([src/main.c])
AC_CONFIG_HEADERS([config.h])

# Checks for programs.
AC_PROG_CC
AC_PROG_INSTALL

# Checks for libraries.

# Checks for header files.
AC_HEADER_STDC
AC_CHECK_HEADERS([stdlib.h])

# Checks for typedefs, structures, and compiler ...

# Checks for library functions.

AC_CONFIG_FILES([Makefile
                 src/Makefile])
AC_OUTPUT
$
```

*NOTE: The contents of your `configure.ac` file may differ slightly from mine, depending on the version of Autoconf you have installed. I have version 2.62 of GNU Autoconf installed on my system (the latest, as of this writing), but if your version of `autoscan` is older (or newer), you may see some minor differences.*

I'll then edit the file and change the `AC_INIT` macro to reflect the jupiter project parameters:

```
$ head configure.ac
#          -*- Autoconf -*-
# Process this file with autoconf to produce ...

AC_PREREQ([2.61])
AC_INIT([jupiter], [1.0], [bugs@jupiter.org])
```

```
AC_CONFIG_SRCDIR([src/main.c])
AC_CONFIG_HEADERS([config.h])
$
```

The `autoscan` utility really does a lot of the work for you. The GNU Autoconf manual states that you should hand-tailor this file to your project before using it. This is true, but there are only a few key issues to worry about (besides those related to the `AC_INIT` macro). We'll take each of these as we come to them, starting at the top of the file.

### Trying out `configure`

I like to experiment, so the first thing I'd do at this point would be to try to run `autoreconf` on this new `configure.ac`, and then try to run the generated `configure` script to see what happens. If `autoscan` is all it's cracked up to be, then the resulting `configure` script should generate some makefiles for me:

```
$ autoreconf
$ ./configure
checking for gcc... gcc
checking for C compiler default output file name...
checking whether the C compiler works... yes
checking whether we are cross compiling... no
checking for suffix of executables...
checking for suffix of object files... o
checking whether we are using the GNU C compiler...
checking whether gcc accepts -g... yes
checking for gcc option to accept ISO C89...
configure: error: cannot find install-sh or
  install.sh in "." "../.." "../../.."
$
```

Well, we didn't get too far. I mentioned the `install` utility in Chapter 1, and you may have already been aware of it. It appears here that Autoconf is looking for a shell script called `install-sh` or `install.sh`.

Autoconf is all about portability, and unfortunately, the `install` utility is not as portable as we'd like it to be. From one platform to another, critical bits of installation functionality are just different enough to cause problems, so the Autotools provide a shell script called `install-sh` (deprecated name: `install.sh`) that acts as a wrapper around the platform `install` utility. This wrapper script masks important differences among various versions of `install`.

`autoscan` noticed that we used the `install` program in our `src/Makefile.in` template, and generated an expansion of the `AC_PROG_INSTALL` macro into the `configure.scan` file based on this observation. The problem is that the generated `configure` script couldn't find the `install-sh` wrapper script.

This seems to be a minor defect in Autoconf—if Autoconf expects `install-sh` to be in our project directory, then it should just put it there, right? Well, `autoreconf` has a

command line option, `--install`, which is supposed to install missing files like this for me. I'll give it a try. Here's a before-and-after picture of my directory structure:

```
$ ls -lp
autoscan.log
configure.ac
Makefile.in
src/
$
$ autoreconf --install
$ ls -lp
autom4te.cache/
autoscan.log
config.h.in
configure
configure.ac
Makefile.in
src/
$
```

Hmmm. It didn't seem to work, as there's no `install-sh` file in the directory after running `autoreconf --install`. This is, in my opinion, a defect in both `autoreconf` and `autoconf`. You see, when `autoreconf` is used with the `--install` command-line option, it *should* install all auxilliary files required by all Autoconf macros used in `configure.ac`. The trouble is, this auxilliary-file-installation functionality is actually a part of Automake, not Autoconf. So when you use `--install` on the `autoreconf` command-line, it passes tool-specific `install-missing-files` options down to each of the tools that it calls. This technique would have worked just fine, except that Autoconf doesn't *provide* an option to install any missing files.

Worse still, the GNU Autoconf manual tells you in Section 5.2.1 under `AC_PROG_INSTALL`, "Autoconf comes with a copy of `install-sh` that you can use." But this is a lie. In fact, it's Automake and Libtool that come with copies of `install-sh`, not Autoconf.

We could just copy `install-sh` from the Automake installation directory (`PREFIX/share/automake...`), but let's try running `automake` instead:

```
$ automake --add-missing --copy
configure.ac: no proper invocation of AM_INIT_...
configure.ac: You should verify that configure...
configure.ac: that aclocal.m4 is present in th...
configure.ac: and that aclocal.m4 was recently...
configure.ac:11: installing `./install-sh'
automake: no `Makefile.am' found for any confi...
$
```

Ignoring the warnings indicating that I've not yet configured my project properly for Automake, I can now see that `install-sh` was copied into my project root directory:

```
$ ls -lp
autom4te.cache/
```

```

autoscan.log
configure.ac
configure.scan
install-sh
Makefile.in
src/
$

```

So why didn't `autoreconf --install` do this for me? Isn't it supposed to run all the programs that it needed to, based on my `configure.ac`? As it happens, it was exactly *because* my project was not configured for Automake, that `autoreconf` failed to run `automake --add-missing --copy`. `Autoreconf` saw no reason to run `automake` because `configure.ac` doesn't contain the requisite macros for initializing Automake.

And therein lies the defect. First, `Autoconf` should ship with `install-sh`, since it provides a macro that requires it, and because `autoscan` adds that macro based on the contents of a `Makefile.in` template. In addition, `Autoconf` should provide an "add-missing" command-line option, and `autoreconf` should use it when called with the `--install` option. This is most likely an example of the "work-in-progress" nature of the Autotools.

Before `autoreconf` came along, maintainers used a shell script, often called `autogen.sh` to run all of the Autotools required for their projects in the proper order. To solve this problem temporarily, I'll just add a simple temporary `autogen.sh` script to the project root directory:

```

$ echo "automake --add-missing --copy
autoreconf --install" > autogen.sh
chmod 0755 autogen.sh
$

```

If you don't want to see all the error messages from `automake`, just redirect the `stderr` and `stdout` output to `/dev/null`. By the way, the `--copy` option tells Automake to copy the missing auxiliary files, rather than simply creating links to them from where they're installed in `/usr/share/automake`.

Eventually, we'll be able to get rid of `autogen.sh` file, and just run `autoreconf --install`, but for now, this will solve our missing files problems. Hopefully, you read this section before scratching your head too much over the missing `install-sh` script. We can now run our newly generated `configure` script without errors. I'll cover the details of properly using the `AC_PROG_INSTALL` macro shortly. I'll cover Automake in much greater detail in Chapter 4.

## Updating `Makefile.in`

Okay, so how do the additional macros added by `autoscan` affect our `make` system? Well, we have some new files to consider. For one, the `config.h.in` file is generated for us

now by `autoheader`. We can assume that `autoreconf` now executes `autoheader` for us when we run it. Additionally, we have a new file in our project called `install-sh`.

Anything provide by, or generated by the Autotools should be copied into the archive directory so that it can be shipped with our release tarballs. So, we should add these two files. Note that we don't need to install `autogen.sh`, as it's purely a maintainer tool—our users shouldn't ever need to execute it:

```
...
srcdir      = @srcdir@
VPATH      = $(srcdir)
...
$(distdir):
    mkdir -p $(distdir)/src
    cp $(srcdir)/configure $(distdir)
    cp $(srcdir)/config.h.in $(distdir)
    cp $(srcdir)/install-sh $(distdir)
    cp $(srcdir)/Makefile.in $(distdir)
    cp $(srcdir)/src/Makefile.in $(distdir)/src
    cp $(srcdir)/src/main.c $(distdir)/src
...
```

If you're beginning to think that this could become a maintenance nightmare, then you're right. I warned you that the `$(distdir)` target was painful to maintain in Chapter 2. Luckily the `distcheck` target still exists, and still works as designed. It would have caught this problem, because the distribution build will not work without these additional files, and certainly the `check` target wouldn't work, if the build didn't work. When we discuss Automake in Chapter 4, much of this mess will be cleared up.

## Initialization and package information

The first section in our new `configure.ac` file (copied from `configure.scan`) contains Autoconf initialization macros. These are required for all projects. Let's consider each of these macros individually.

### AC\_PREREQ

The `AC_PREREQ` macro simply defines the lowest version of Autoconf that may be used to successfully process the `configure.ac` script. The manual indicates that `AC_PREREQ` is the only macro that may be used before `AC_INIT`. The reason for this should be obvious—you'd like to be able to ensure you're using a late enough version of Autoconf before you begin processing any other macros, which may be version dependent. As it turns out, `AC_INIT` is not version dependent anyway, so you may place it first, if you're so inclined. I happen to prefer the way `autoscan` generates the file, so I'll leave it alone.

### AC\_INIT

The `AC_INIT` macro (as its name implies) initializes the Autoconf system. It accepts up to four arguments (`autoconf` only generates a call with the first three), `PACKAGE`, `VERSION`, and optional `BUG-REPORT` and `TARNAME` arguments. The `PACKAGE` argument is intended to be the name of your package. It will end up (in a canonicalized form) as the first part of the name of an Automake-generated release distribution tarball when you run `make dist`.

In fact, by default, Automake-generated tarballs will be named `TARNAME-VERSION.tar.gz`, but `TARNAME` is set to a canonicalized form of the `PACKAGE` string (lower-cased with all punctuation converted to underscores), unless you specify `TARNAME` manually, so bear this in mind when you choose your package name and version string. Incidentally, M4 macro arguments, including `PACKAGE` and `VERSION`, are just strings. M4 doesn't attempt to interpret any of the text that it processes.

The optional `BUG-REPORT` argument is usually set to an email address, but it can be any text really. An Autoconf substitution variable called `PACKAGE_BUGREPORT` will be created for it, and that variable will be added to a `config.h.in` template as a C preprocessor string, as well. The intent is that you use the variable in your code (or in template text files anywhere in your project) to present an email address for bug reports at appropriate places—possibly when the user requests help or version information from your application.

While the `VERSION` argument can be anything you like, there are a few free software conventions that will make life a little easier for you if you follow them. The widely used convention is to pass in `MAJOR.MINOR` (eg., `1.2`). However, there's nothing that says you can't use `MAJOR.MINOR.REVISION` if you want, and there's nothing wrong with this approach. None of the resulting `VERSION` macros (Autoconf, shell or `make`) are parsed or analysed anywhere—only used in various places as replacement text, so if you want, you can even add non-numeric text into this macro, such as `0.15.alpha1`, which is useful occasionally.

*Note that the RPM package manager does indeed care what you put in the version string. For the sake of RPM, you may wish to limit the version string text to only alpha-numeric and periods—no dashes or underscores, unfortunately.*

Autoconf will generate the substitution variables `PACKAGE_NAME`, `PACKAGE_VERSION`, `PACKAGE_TARNAME`, `PACKAGE_STRING` (a stylized concatenation of the package name and version information) and `PACKAGE_BUGREPORT` from arguments to `AC_INIT`.

### **AC\_CONFIG\_SRCDIR**

The `AC_CONFIG_SRCDIR` macro is just a sanity check. Its purpose is to ensure that the generated `configure` script knows that the directory on which it is being executed is in fact the correct project directory. The argument can be a relative path to any source file you like - I try to pick one that sort of

defines the project. That way, in case I ever decide to reorganize source code, I'm not likely to lose it in a file rename. But it doesn't really matter, because if you do rename the file or move it to some other location some time down the road, you can always change the argument passed to `AC_CONFIG_SRCDIR`. Autoconf will tell you immediately if it can't find this file—after all, that's the purpose of this macro in the first place!

## The instantiating macros

Before we dive into the details of `AC_CONFIG_HEADERS`, I'd like to spend a little time on the framework provided by Autoconf. From a high-level perspective, there are four major things happening in `configure.ac`:

1. Initialization
2. File instantiation
3. Check requests
4. Generation of the `configure` script

We've pretty much covered initialization—there's not much to it, although there are a few more macros you should be aware of. (Check out the GNU Autoconf manual to see what these are—look up `AC_COPYRIGHT`, for an example.) Now, let's move on to file instantiation.

There are actually four so-called instantiating macros, which include `AC_CONFIG_FILES`, `AC_CONFIG_HEADERS`, `AC_CONFIG_COMMANDS` and `AC_CONFIG_LINKS`. An instantiating macro is one which defines one or more tags, usually referring to files that are to be translated, by the generated `configure` scripts, from a template containing Autoconf substitution variables.

*NOTE: You might need to change the name of `AC_CONFIG_HEADER` (singular) to `AC_CONFIG_HEADERS` (plural) in your version of `configure.scan`. This was a defect in `autoscan` that had not been fixed yet in Autoconf version 2.61. I reported the defect and a patch was committed. Version 2.62 works correctly. If your `configure.scan` is generated with a call to `AC_CONFIG_HEADER`, just change it manually. Both macros will work, as the singular version was the older name of this macro, but the older macro is less functional than the newer one.*

These four instantiating macros have an interesting signature in common:

```
AC_CONFIG_XXXS([tag ...], [commands], [init-cmds])
```

For each of these four macros, the `tag` argument has the form, `OUT[:INLIST]` where `INLIST` has the form, `IN0[:IN1:...:INn]`. Often, you'll see a call to one of these macros with only a single simple argument, like this:

```
AC_CONFIG_HEADERS([config.h])
```

In this case, `config.h` is the `OUT` portion of the above specification. The default `INLIST` is the `OUT` portion with `“.in”`

appended to it. So the above call is equivalent to:

```
AC_CONFIG_HEADERS([config.h:config.h.in])
```

What this means is that `config.status` will contain shell code that will generate `config.h` from `config.h.in`, substituting all Autoconf variables in the process. You may also provide a list of input files to be concatenated, like this:

```
AC_CONFIG_HEADERS([config.h:cfg0:cfg1:cfg2])
```

In this example, `config.status` will generate `config.h` by concatenating `cfg0`, `cfg1` and `cfg2`, after substituting all Autoconf variables. The GNU Autoconf manual calls this entire “OUT:INLIST” thing a “tag”.

What’s that all about? Why not call it a file? Well, the fact is, this parameter’s primary purpose is to provide a sort of command-line target name—much like `Makefile` targets. It also *happens* to be used as a file system name, *if the associated macro happens* to generate file system names, as is the case when calling `AC_CONFIG_HEADERS`, `AC_CONFIG_FILES` and `AC_CONFIG_LINKS`.

But `AC_CONFIG_COMMANDS` doesn’t actually generate any files. Rather, it runs arbitrary shell code, as specified by the user in the macro. Thus, rather than name this first parameter after a secondary function (the generation of files), the manual refers to it by its primary purpose - as a command line tag-name that may be specified on the `config.status` command line. Here’s an example:

```
./config.status config.h
```

This `config.status` command line will regenerate the `config.h` file based on the macro call to `AC_CONFIG_HEADERS` in `configure.ac`. It will *only* regenerate `config.h`. Now, if you’re curious like me, you’ve already been playing around a little, and have tried typing `./config.status --help` to see what options are available when executing `config.status`. You may have noticed that `config.status` has a help signature like this:

```
$ ./config.status --help
`config.status' instantiates files from templates
according to the current configuration.
```

```
Usage: ./config.status [OPTIONS] [FILE]...
```

```
  -h, --help          print this help, then exit
```

```
...
```

```
  --file=FILE[:TEMPLATE]
```

```
...
```

```
Configuration files:
```

```
  Makefile src/Makefile
```

```
Configuration headers:
```

```
  config.h
```

*NOTE: I left out portions of the help display irrelevant to this*

*discussion.*

I'd like you to notice a couple of things about this help display. First, `config.status` is designed to give you custom help about this particular project's `config.status` file. It lists "Configuration files" and "Configuration headers" that you may use as tags. Oddly, given the "tag" nomenclature used in the manual so rigorously, the help line still refers to such tags as `[FILE]`s in the "Usage:" line. Regardless, where the usage specifies `[FILE]`s you may use one or more of the listed configuration files, headers, links, or commands displayed below. In this case, `config.status` will only instantiate those objects. In the case of commands, it will execute the commands specified by the tag passed in the associated expansion of the `AC_CONFIG_COMMANDS` macro.

Each of these macros may be used multiple times in a `configure.ac` script. The results are cumulative. This means that we can use `AC_CONFIG_FILES` as many times as we need to in our `configure.ac` file. Reasons why we may want to use it more than once are not obvious right now, but I'll get to them eventually.

Another noteworthy item here is that there is a `--file` option. Now why would `config.status` allow us to specify files either with or without the `--file=` in front of them? Well, these are actually different usages of the `[FILE]` option, which is why it would make more sense for the usage text to read:

```
$ ./config.status --help
...
Usage: ./config.status [OPTIONS] [TAG]...
```

When `config.status` is called with tag names on the command line, only those tags listed in the help text as available configuration files, headers, links and commands may be used as tags. When you execute `config.status` with the `--file=` option, you're really telling `config.status` to generate a new file *not already associated with any of the calls to instantiating macros in your `configure.ac`*. The file is generated from a template using configuration options and check results determined by the last execution of the `configure` script. For example, I could execute `config.status` like this:

```
./config.status --file=extra:extra.in
```

*NOTE: The default template name is the file name with a ".in" suffix, so this call could have been made without using the ":extra.in" portion of the option.*

Let's get back to the instantiating macro signature. The `tag` argument has a complex format, but it also represents multiple tags. Take another look:

```
AC_CONFIG_xxxS([tag ...], [commands], [init-cmds])
```

The elipsis after `tag` indicates there may be more than one, and in fact, this is true. The `tag` argument accepts multiple

tag specifications, separated by white space or new-line characters. Often you'll see a call like this:

```
AC_CONFIG_FILES([Makefile
                 src/Makefile
                 lib/Makefile
                 etc/project.cfg])
```

Each entry here is one tag specification, which if fully specified would look like this:

```
AC_CONFIG_FILES([Makefile:Makefile.in
                 src/Makefile:src/Makefile.in
                 lib/Makefile:lib/Makefile.in
                 etc/proj.cfg:etc/proj.cfg.in])
```

One more point to cover. There are two optional arguments that you'll not often see used in the instantiating macros, `commands` and `init-cmds`. The `commands` argument may be used to specify some arbitrary shell code that should be executed by `config.status` just before the files associated with the tags are generated. You'll not often see this used with the file generating instantiating macros, but in the case of `AC_CONFIG_COMMANDS`, which generates no files by default, you almost always see this argument used, because a call to this macro is basically useless without it! In this case, the tag argument becomes a way of telling `config.status` to execute a set of shell commands.

The `init-cmds` argument is used to initialize shell variables at the top of `config.status` with values available in `configure.ac` and `configure`. It's important to remember that all calls to instantiating macros share a common namespace along with `config.status`, so choose shell variable names carefully.

The old adage about the relative value of a picture vs. an explanation holds true here, so let's try a little experiment. Create a test version of your `configure.ac` file containing only the following lines:

```
AC_INIT(test, 1.0)
AC_CONFIG_COMMANDS([abc],
                   [echo "Testing $mypkgname"],
                   [mypkgname=$PACKAGE_NAME])
AC_OUTPUT
```

Then execute `autoreconf`, `configure`, and `config.status` in various ways to see what happens:

```
$ autoreconf
$ ./configure
configure: creating ./config.status
config.status: executing abc commands
Testing test
$ ./config.status
config.status: executing abc commands
Testing test
$ ./config.status --help
`config.status' instantiates files from templates
```

```

according to the current configuration.

Usage: ./config.status [OPTIONS] [FILE]...
...
Configuration commands:
  abc

Report bugs to <bug-autoconf@gnu.org>.
$ ./config.status abc
config.status: executing abc commands
Testing test
$

```

As you can see here, executing `configure` caused `config.status` to be executed with no command line options. There are no checks specified in `configure.ac`, so executing `config.status` has nearly the same effect. Querying `config.status` for help indicates that “abc” is a valid tag, and executing `config.status` with that tag simply runs the associated commands.

Okay, enough fooling around. The important points to remember here are:

1. Both `configure` and `config.status` may be called individually to perform their individual tasks.
2. The `config.status` script generates all files from templates.
3. The `configure` script performs all checks and then executes `config.status`.
4. `config.status` generates files based on the last set of check results.
5. `config.status` may be called to execute file generation or command sets specified by any of the tag names specified in any of the instantiating macro calls.
6. `config.status` may generate files not associated with any tags specified in `configure.ac`.
7. `config.status` can be used to call `configure` with the same set of command line options used in the last execution of `configure`.

## AC\_CONFIG\_HEADERS

As you’ve no doubt concluded by now, the `AC_CONFIG_HEADERS` macro allows you to specify one or more header files to be generated from template files. You may write multiple template header files yourself, if you wish. The format of a configuration header template is very specific:

```

/* Define as 1 if you have unistd.h. */
#undef HAVE_UNISTD_H

```

Multiple such statements may be placed in your header template. The comments are optional, of course. Let’s try another experiment. Create a new `configure.ac` file with the following contents:

```

AC_INIT([test], [1.0])
AC_CONFIG_HEADERS([config.h])

```

```
AC_CHECK_HEADERS([unistd.h foobar.h])
AC_OUTPUT
```

Now create a configuration header template file called `config.h.in`, which contains the following two lines:

```
#undef HAVE_UNISTD_H
#undef HAVE_FOOBAR_H
```

Finally, execute the following commands:

```
$ autoconf
$ ./configure
checking for gcc... gcc
...
checking for unistd.h... yes
checking for unistd.h... (cached) yes
checking foobar.h usability... no
checking foobar.h presence... no
checking for foobar.h... no
configure: creating ./config.status
config.status: creating config.h
$ cat config.h
/* config.h. Generated from ... */
#define HAVE_UNISTD_H 1
/* #undef HAVE_FOOBAR_H */
```

You can see that `config.status` generated a `config.h` file from your `config.h.in` template file. The contents of this header file are based on the checks executed by the `configure` script. Since the shell code generated by `AC_CHECK_HEADERS([unistd.h foobar.h])` was able to locate a `unistd.h` header file in the standard include directory, the corresponding `#undef` statement was converted into a `#define` statement. Of course, no `foobar.h` header was found in the system include directory, as you can also see by the output of `configure`, so its definition was left commented out in the template.

Thus, you may add this sort of code to appropriate C source files in your project:

```
#if HAVE_CONFIG_H
# include <config.h>
#endif

#if HAVE_UNISTD_H
# include <unistd.h>
#endif
```

### Using autoheader to generate an include file template

Maintaining your `config.h.in` template is more pain than necessary. After all, most of the information you need is already encapsulated in your `configure.ac` script, and the format of `config.h.in` is very strict. For example, you may not have any leading or trailing white space on the `#undef` lines.

Fortunately, the autoheader utility will generate an include

header template for you based on your `configure.ac` file contents. Back to the command prompt for another quick experiment. This one is easy—just *delete* your `config.h.in` template before you run `autoheader` and `autoconf`, like this:

```
$ rm config.h.in
$ autoheader
$ autoconf
$ ./configure
checking for gcc... gcc
...
checking for unistd.h... yes
checking for unistd.h... (cached) yes
checking foobar.h usability... no
checking foobar.h presence... no
checking for foobar.h... no
configure: creating ./config.status
config.status: creating config.h
$ cat config.h
/* config.h. Generated from config.h.in... */
/* config.h.in. Generated from configure.ac... */
...
/* Define to 1 if you have... */
/* #undef HAVE_FOOBAR_H */

/* Define to 1 if you have... */
#define HAVE_UNISTD_H 1

/* Define to the address where bug... */
#define PACKAGE_BUGREPORT ""

/* Define to the full name of this package. */
#define PACKAGE_NAME "test"

/* Define to the full name and version... */
#define PACKAGE_STRING "test 1.0"

/* Define to the one symbol short name... */
#define PACKAGE_TARNAME "test"

/* Define to the version... */
#define PACKAGE_VERSION "1.0"

/* Define to 1 if you have the ANSI C... */
#define STDC_HEADERS 1
```

*NOTE: Here again, I encourage you to use `autoreconf`, which will automatically run `autoheader` for you if it notices an expansion of the `AC_CONFIG_HEADERS` macro in your `configure.ac` script.*

You may also want to take a peek at the `config.h.in` template file generated by `autoheader`. In the meantime, here's a much more realistic example of using a generated `config.h` file for the sake of portability of project source code.

```
AC_INIT([test], [1.0])
```

```
AC_CONFIG_HEADERS([config.h])
AC_CHECK_HEADERS([dlfcn.h])
AC_OUTPUT
```

The `config.h` file is obviously intended to be included in your source code in locations where you might wish to test a configured option in the code itself using the C preprocessor. Using this `configure.ac` script, Autoconf will generate a `config.h` header file with appropriate definitions for determining at compile time if the current system provides the `dlfcn` interface. To complete the portability check, you can add the following code to a source file in your project that uses dynamic loader functionality:

```
#if HAVE_CONFIG_H
# include <config.h>
#endif

#if HAVE_DLFCN_H
# include <dlfcn.h>
#else
# error Sorry, this code requires dlfcn.h.
#endif
...
#if HAVE_DLFCN_H
    handle = dlopen(
        "/usr/lib/libwhatever.so", RTLD_NOW);
#endif
...
```

If you already had code that included `dlfcn.h` then `autoscan` will have generated a `configure.ac` call to `AC_CHECK_HEADERS`, which contains `dlfcn.h` as one of the header files to be checked. Your job as the maintainer is to add the conditional to your source code around the existing use of the `dlfcn` header inclusion and the `libdl` API calls. This is the crux of Autoconf provided portability.

Your project may be able to get along at compile time without the dynamic loader functionality if it must, but it would be nice to have it. Perhaps, your project will function in a limited manner without it. Sometimes you just have to bail out with a compiler error (as this code does) if the key functionality is missing. Often this is an acceptable first-attempt solution, until someone comes along and adds support to the code base for some other dynamic loader service that is perhaps available on non-`dlfcn`-oriented systems.

*NOTE: If you have to bail out with an error, it's best to do so at configuration time, rather than at compile time. I'll cover examples of this sort of activity shortly.*

One obvious flaw in this source code is that `config.h` is only included if `HAVE_CONFIG_H` is defined in your compilation environment. But wait...doesn't that definition happen in `config.h`? Well, no, not in the case of this particular definition. `HAVE_CONFIG_H` must be either defined by you manually, if you're writing your own makefiles, or by Automake-generated makefiles automatically on the compiler

command line. (Are you beginning to get the feeling that Autoconf really shines when used in conjunction with Automake?)

HAVE\_CONFIG\_H is part of a string of definitions passed on the compiler command line in the Autoconf substitution variable @DEFS@. Before Autoheader and AC\_CONFIG\_HEADERS, all of the compiler configuration macros were added to the @DEFS@ variable. You can still use this method if you don't use AC\_CONFIG\_HEADERS in `configure.ac`. but it's not the recommended method nowadays.

## Back to VPATH builds for a moment

Regarding VPATH builds, we haven't yet covered how to get the preprocessor to properly locate our generated `config.h` file. This file, being a generated file, will be found in the same relative position in the build directory structure, as its counterpart template file, `config.h.in`. The template is located in the top-level source directory (unless we choose to put it somewhere else), so the generated file will be in the top-level build directory. Well, that's easy enough—it's always one level up from `src/Makefile`.

Let's consider where we might have include files in our project. We might add an internal header file to the current source directory. We obviously now have a `config.h` file in our top-level build directory. We might also create a top-level source include directory for library interface header files. Which order should we care about these files?

The order we place include directives (`-I<path>`) options on the compiler command line is the order in which they will be searched. The proper preprocessor include paths should include the current build directory (`.`), the source directory (`$(srcdir)`), and the top-level build directory (`..`), in that order:

```
...
jupiter: main.c
        gcc -g -O0 -I. -I$(srcdir) -I..\
        -o $@ $(srcdir)/main.c
...
```

It appears we need an additional rule for VPATH builds:

1. Add preprocessor commands for the current build, associated source and top-level build directories, in that order.

## Checks for compilers

The `AC_PROG_CC` macro ensures that we have a working C language compiler. This call was added to `configure.scan` when `autoscan` noticed that we had C source files in our project directory. If we'd had files suffixed with `.cxx` or `.C` (an upper-case `.C` extension indicates a C++ source file), it would have inserted a call to the

AC\_PROG\_CXX macro, as well as a call to AC\_LANG([C++]).

This macro looks for gcc and then cc in the system search path. If neither of these are found, it looks for other C compilers. When a compatible compiler is found, it sets a well-known variable, \$CC to the full path of the program, with options for portability, if necessary.

AC\_PROG\_CC accepts an optional parameter containing an ordered list of compiler names. For example, if you used AC\_PROG\_CC([cc cl gcc]), then the macro would expand into shell code that searched for cc, cl and gcc, in that order.

The AC\_PROG\_CC macro also defines the following Autoconf substitution variables:

- @CC@ (full path of compiler)
- @CFLAGS (eg., -g -O2 for gcc)
- @CPPFLAGS@ (empty by default)
- @EXEEXT@ (eg., .exe)
- @OBJEXT@ (eg., .o)

AC\_PROG\_CC configures these substitution variables, but unless I used them in my Makefile.in templates, I'm just wasting time running configure. I'll add a few of these as make variables to my src/Makefile.in template, and then consume them, like this:

```
# Tool-related substitution variables
CC          = @CC@
CFLAGS      = @CFLAGS@
CPPFLAGS    = @CPPFLAGS@
...
jupiter: main.c
    $(CC) $(CFLAGS) $(CPPFLAGS) \
    -I. -I$(srcdir) -I..\
    -o $@ $(srcdir)/main.c
```

## Checking for other programs

Now, let's return to the AC\_PROG\_INSTALL macro. As with the AC\_PROG\_CC macro, the other AC\_PROG\_\* macros set and AC\_SUBST various environment variables that point to the located utility. To make use of this check, you need to use these Autoconf substitution variables in your Makefile.in templates, just as we did with CC, CFLAGS, and CPPFLAGS above:

```
...
# Tool-related substitution variables
CC          = @CC@
CFLAGS      = @CFLAGS@
CPPFLAGS    = @CPPFLAGS@
INSTALL     = @INSTALL@
INSTALL_DATA = @INSTALL_DATA@
INSTALL_PROG = @INSTALL_PROG@
INSTALL_SCRIPT = @INSTALL_SCRIPT@
...
```

```
install:
    mkdir -p $(bindir)/jupyter
    $(INSTALL_PROG) -m 0755 $(bindir)/jupyter
```

The value of `@INSTALL@` is obviously the path of the located `install` script. The value of `@INSTALL_DATA@` is `${INSTALL} -m 0644`. Now, you'd think that the values of `@INSTALL_PROG@` and `@INSTALL_SCRIPT@` would be `${INSTALL} -m 0755`, but they're not. These are just set to `${INSTALL}`. Oversight? I don't know.

Other important utility programs you might need to check for are `lex`, `yacc`, `sed`, `awk`, etc. If so, you can add calls to `AC_PROG_LEX`, `AC_PROG_YACC`, `AC_PROG_SED`, or `AC_PROG_AWK` yourself. There are about a dozen different programs you can check for using these more specialized macros. If such a program check fails, then the resulting `configure` script will fail with a message indicating that the required utility could not be found, and that the build may not continue until it's been properly installed.

As with the other program and compiler checks, in `Makefile.in` templates, you should use the `make` variables `$(LEX)` and `$(YACC)` to invoke these tools (note that Automake does this for you), as these Autoconf macros will set the values of these variables according to the tools it finds installed on your system *if they are not already set in your environment*.

Now, this is a key aspect of `configure` scripts generated by Autoconf—you may *always* override anything `configure` will do to your environment by exporting or setting an appropriate output variable *before* you execute `configure`.

For example, perhaps you would like to build with a very specific version of `bison` that you've installed in your own home directory:

```
$ cd jupyter
$ YACC="$HOME/bin/bison -y" ./configure
$ ...
```

This will ensure that `YACC` is set the way you want for your makefiles, and that `AC_PROG_YACC` does essentially nothing in your `configure` script.

If you need to check for the existence of a program not covered by these more specialized macros, you can call the generic `AC_CHECK_PROG` macro, or you can write your own special purpose macro (I'll cover writing macros in Chapter 9).

Key points to take away:

1. `AC_PROG_*` macros check for the existence of programs.
2. If a program is found, a substitution variable is created.
3. Use these variables in your `Makefile.in` templates to execute the program.

### A common problem with Autoconf

Here's a common problem that developers new to the Autotools consistently encounter. Take a look at the formal definition of `AC_CHECK_PROG` found in the GNU Autoconf manual. *NOTE: In this case, the square brackets represent optional parameters, not Autoconf quotes.:*

```
AC_CHECK_PROG(variable, prog-to-
check-for, value-if-found,
[value-if-not-found], [path],
[reject])
```

Check whether program `prog-to-check-for` exists in `PATH`. If it is found, set `variable` to `value-if-found`, otherwise to `value-if-not-found`, if given. Always pass over `reject` (an absolute file name) even if it is the first found in the search path; in that case, set `variable` using the absolute file name of the `prog-to-check-for` found that is not `reject`. If `variable` was already set, do nothing. Calls `AC_SUBST` for `variable`.

I can extract the following clearly defined functionality from this description:

1. If `prog-to-check-for` is found in the system search path, then `variable` is set to `value-if-found`, otherwise it's set to `value-if-not-found`.
2. If `reject` is specified (as a full path), then skip it if it's found first, and continue to the next matching program in the system search path.
3. If `reject` is found first in the path, and then another match is found besides `reject`, set `variable` to the absolute path name of the second (non-reject) match.
4. If `variable` is already set by the user in the environment, then `variable` is left untouched (thereby allowing the user to override the check by setting `variable` before running `autoconf`).
5. `AC_SUBST` is called on `variable` to make it an Autoconf substitution variable.

At first read, there appears to be a terrible conflict of interest here: We can see in point 1 that `variable` will be set to one or the other of two specified values, based on whether or not `prog-to-check-for` is found in the system search path. But then in point 3 it states that `variable` will be set to the full path of some program, but only if `reject` is found first and skipped. Clearly the documentation needs a little work.

Discovering the real functionality of `AC_CHECK_PROG` is as easy as reading a little shell script. While you could spend your time looking at the definition of `AC_CHECK_PROG` in `/usr/share/autoconf/autoconf/programs.m4`

, the problem with this approach is that you're one level removed from the actual shell code performing the check. Wouldn't it be better to just look at the resulting shell script generated by AC\_CHECK\_PROG? Okay, then modify your new `configure.ac` file in this manner:

```
...
AC_PREREQ(2.59)
AC_INIT([jupiter], [1.0],
        [jupiter-devel@lists.example.com])
AC_CONFIG_SRCDIR([src/main.c])
AC_CONFIG_HEADER([config.h])

# Checks for programs.
AC_PROG_CC
AC_CHECK_PROG([bash_var], [bash], [yes],
              [no],, [/usr/sbin/bash])
...
```

Now just execute `autoconf` and then open the resulting `configure` script and search for something specific to the definition of `AC_CHECK_PROG`. I used the string `"ac_cv_prog_bash_var"`, a shell variable generated by the macro call. You may have to glance at the definition of a macro to find reasonable search text:

```
$ autoconf
$ vi -c /ac_cv_prog_bash_var configure
...
# Extract the first word of "bash", so it can be
# a program name with args.
set dummy bash; ac_word=$2
echo "$as_me:$LINENO: checking for $ac_word" >&6
echo $ECHO_N "checking for $ac_word... $ECHO_C" \
>&6
if test "${ac_cv_prog_bash_var+set}" = set; then
  echo $ECHO_N "(cached) $ECHO_C" >&6
else
  if test -n "$bash_var"; then
    # Let the user override the test.
    ac_cv_prog_bash_var="$bash_var"
  else
    ac_prog_rejected=no
    as_save_IFS=$IFS; IFS=$PATH_SEPARATOR
    for as_dir in $PATH
    do
      IFS=$as_save_IFS
      test -z "$as_dir" && as_dir=.
      for ac_exec_ext in ''\
$ac_executable_extensions;
      do
        if $as_executable_p \
"$as_dir/$ac_word$ac_exec_ext"; then
          if test "$as_dir/$ac_word$ac_exec_ext" = \
"/usr/sbin/bash"; then
            ac_prog_rejected=yes
            continue
          fi
        ac_cv_prog_bash_var="yes"
```

```

        echo "$as_me:$LINENO: found\
$as_dir/$ac_word$ac_exec_ext" >&5
        break 2
    fi
done
done

if test $ac_prog_rejected = yes; then
    # We found a bogon in the path, so make sure
    # we never use it.
    set dummy $ac_cv_prog_bash_var
    shift
    if test $# != 0; then
        # We chose a different compiler from the
        # bogus one. However, it has the same
        # basename, so the bogon will be chosen
        # first if we set bash_var to just the
        # basename; use the full file name.
        shift
        ac_cv_prog_bash_var=\
"$as_dir/$ac_word${1+ ' '}$@"
    fi
fi
test -z "$ac_cv_prog_bash_var" &&\
ac_cv_prog_bash_var="no"
fi
fi
bash_var=$ac_cv_prog_bash_var
if test -n "$bash_var"; then
    echo "$as_me:$LINENO: result: $bash_var" >&5
    echo "${ECHO_T}$bash_var" >&6
else
    echo "$as_me:$LINENO: result: no" >&5
    echo "${ECHO_T}no" >&6
fi
...

```

Wow! You can immediately see by the opening comment that `AC_CHECK_PROG` has some undocumented functionality: You can pass in arguments with the program name if you wish. But why would you want to? Well, look farther. You can probably fairly accurately deduce that the `reject` parameter was added into the mix in order to allow your `configure` script to search for a particular version of a tool. (Could it possibly be that someone might really rather use the GNU C compiler instead of the Solaris C compiler?)

In fact, it appears that `variable` really is set based on a *tri-state condition*. If `reject` is not used, then `variable` can only be either `value-if-found` or `value-if-not-found`. But if `reject` is used, then `variable` can also be the full path of the first program found that is not `reject`! Well, that is exactly what the documentation stated, but examining the generated code yields insight into the authors' intended use of this macro. We probably should have called `AC_CHECK_PROG` this way, instead:

```
AC_CHECK_PROG([bash_shell],[bash -x],[bash -x],,,
```

```
[ /usr/sbin/bash ]
```

Now it makes more sense, and you can see by this example that the manual is in fact accurate, if not clear. If `reject` is not specified, and `bash` is found in the system path, then `bash_shell` will be set to `bash -x`. If it's *not* found in the system path, then `bash_shell` will be set to the empty string. If, on the other hand, *reject is* specified, and the undesired version of `bash` is found *first* in the path, then `bash_shell` will be set to the full path of the *next* version found in the path, along with the originally specified arguments (`-x`). The `bash_shell` variable may now be used by the rest of our script to run the desired `bash` shell, if it doesn't test out as empty. Wow! No wonder it was hard to document in a way that's easy to understand! But quite frankly, a good example of the intended use of this macro, along with a couple of sentences of explanation would have made all the difference.

## Checks for libraries and header files

Does your project rely on external libraries? Most non-trivial projects do. If you're lucky, your project relies only on libraries that are already widely available and ported to most platforms.

The choice to use an external library or not is a tough one. On the one hand, you'll want to reuse code that provides functionality—perhaps significant functionality that you need and don't really have the time or expertise to write yourself. Reuse is one of the hallmarks of the free software world.

On the other hand, you don't want to depend on functionality that may not exist on all of the platforms you wish to target, or that requires significant porting effort on your part to make these libraries available on all of your target platforms.

Occasionally, library-based functionality can exist in slightly different forms on different platforms. These different forms may be functionally compatible, but have different API signatures. For example, POSIX threads (`pthread`) versus a native threading library. For basic multi-threading functionality, many threading libraries are similar enough to be almost drop-in replacements of each other.

To illustrate this concept, We'll add some trival multi-threading capabilities to the `jupiter` project. We want to have `jupiter` print its message using a background thread. To do this, We're going to need to add the `pthread` library to our project build system. If we weren't using the Autotools, we'd just add it to our linker command line in the makefile:

```
jupiter: main.c
    $(CC) ... -lpthreads ...
```

But what if a system doesn't support `pthread`? We might want to support native threads on a non-`pthread` system—say Solaris native threads, using the `libthreads` library.

To do this, we'll first modify our `main.c` file such that the

printing happens in a secondary thread, like this:

```
#include <stdio.h>
#include <stdlib.h>
#include <pthread.h>

static void * print_it(void * data)
{
    printf("Hello from %s!\n", (char *)data);
    return 0;
}

int main(int argc, char * argv[])
{
    pthread_t tid;
    pthread_create(&tid, 0, print_it, argv[0]);
    pthread_join(tid, 0);
    return 0;
}
```

Now, this is clearly a ridiculous use of a thread. Nonetheless, it *is* the prototypical form of thread usage. Consider the case where `print_it` did some long calculation, and `main` had other things to do while `print_it` performed this calculation. On a multi-processor machine, this could literally double the throughput of such a program.

What we now need is a way of determining which libraries should be added to the compiler command line. Enter Autoconf and the `AC_CHECK_*` macros. The `AC_SEARCH_LIBS` macro allows us to check for key functionality within a list of libraries. If the function exists within one of the specified libraries, then an appropriate command line option is added the `@LIBS@` substitution variable. The `@LIBS@` variable should be used in a `Makefile.in` template on the compiler (linker) command line. Here is the formal definition of `AC_SEARCH_LIBS`, again from the manual:

```
AC_SEARCH_LIBS(function, search-
libs, [action-if-found], [action-
if-not-found], [other-libraries])
Search for a library defining function if it's not
already available. This equates to calling
AC_LINK_IFELSE([AC_LANG_CALL([ ],
[function])]) first with no libraries, then for
each library listed in search-libs. Add -
llibrary to LIBS for the first library found to
contain function, and run action-if-found. If
function is not found, run action-if-not-
found. If linking with the library results in
unresolved symbols that would be resolved by
linking with additional libraries, give those libraries
as the other-libraries argument, separated
by spaces: e.g., -lXt -lX11. Otherwise, this
macro fails to detect that function is present,
because linking the test program always fails with
unresolved symbols.
```

Wow, that's a lot of stuff for one macro. Are you beginning to

see why the generated `configure` script is so large? Essentially, what you get by calling `AC_SEARCH_LIBS` for a particular function is that the proper linker command line arguments (eg., `-lpthread`), for linking with a library containing the desired function, are added to a substitution variable called `@LIBS@`. Here's how we'll use `AC_SEARCH_LIBS` in our `configure.ac` file:

```
...
# Checks for libraries.
AC_SEARCH_LIBS([pthread_create], [pthread])
...
```

Of course, we'll have to modify `src/Makefile.in` again to make proper use of the now populated `LIBS` make variable:

```
...
# Tool-related substitution variables
CC          = @CC@
LIBS        = @LIBS@
CFLAGS      = @CFLAGS@
CPPFLAGS    = @CPPFLAGS@
...
jupiter: main.c
    $(CC) $(CFLAGS) $(CPPFLAGS) \
    -I. -I$(srcdir) -I..\ \
    -o $@ $(srcdir)/main.c $(LIBS)
...
```

Note that we added `$(LIBS)` *after* the source file on the compiler command line. Generally, the linker cares about object file order, and searches them for required functions in the order they are specified on the command line. Since we want `main.c` to be the primary source of object code for `jupiter`, we'll continue to add additional objects, including libraries, after this file on the command line.

### Right or just good enough?

We could just stop at this point. We've done enough to make this build system properly use `pthread` on most systems. If a library is needed, it'll be added to the `@LIBS@` variable, and subsequently used on our compiler command line. In fact, this is the point at which many maintainers *would* stop. The problem is that stopping here is just about the build-system equivalent of not checking the return value of `malloc` in a C program (and there are many developers out there who don't give this process the credit it deserves either). It *usually* works fine. It's just during those few cases where it fails that you have a real problem.

Well, we want to provide a good user experience, so we'll take `jupiter`'s build system to the "next level". However, in order to do this, we need to make a design decision: In case `configure` fails to locate a `pthread` library on a user's system, should we fail the build process, or build a `jupiter` program without multi-threading? If we fail the build, it will generally be obvious to the user, because the build has stopped with an error message—although, perhaps not a very

user-friendly one. At this point, either the compile process or the link process will fail with a cryptic error message about a missing header file or an undefined symbol. If we choose to build a single-threaded version of jupiter, we should probably display some clear message that we're moving ahead without threads, and why.

There's another potential problem also. Some users' systems may have a `pthread` library installed, but not have the `pthread.h` header file installed properly. This can happen for a variety of reasons, but the most common is that the executable package was installed, but not the developer package. Executable binaries are often packaged independently of static libraries and header files. Executables are installed as part of a dependency chain for a higher level consuming application, while developer packages are often only installed directly by a user. For this reason, Autoconf provides checks for both libraries and header files. The `AC_CHECK_HEADERS` macro is used to ensure the existence of a particular header file.

Autoconf checks are very thorough. They generally not only ensure the existence of a file, but also that the file is in fact the one we're looking for. They do this by allowing us to make some assertions about the file, which are then verified by the macro. Additionally, the `AC_CHECK_HEADERS` macro doesn't just scan the file system for the requested header. Rather, it builds a short test program in the appropriate language, and then compiles it to ensure that the compiler can both find the file, and use it. Similarly, `AC_SEARCH_LIBS` is built around an attempt to link to the specified libraries, and import the requested symbols.

Here is the formal definition of `AC_CHECK_HEADERS`, as found in the GNU Autoconf manual:

```
AC_CHECK_HEADERS(header-file...,
  [action-if-found], [action-if-not-
  found], [includes = default-
  includes]) For each given system header file
  header-file in the blank-separated argument
  list that exists, define HAVE_header-file (in
  all capitals). If action-if-found is given, it is
  additional shell code to execute when one of the
  header files is found. You can give it a value of
  break to break out of the loop on the first match.
  If action-if-not-found is given, it is
  executed when one of the header files is not
  found.
```

Normally, this macro is called only with a list of desired header files in the first argument. Remaining arguments are optional, and not often used. The reason for this is that the macro is very functional when used in this manner. We'll add a check for the `pthread` library using `AC_CHECK_HEADERS` to our `configure.ac` file.

If you're the jump-right-in type, then you've noticed by now that `configure.ac` already calls `AC_CHECK_HEADERS` for `stdlib.h`. No problem—we'll just add `pthread.h` to

the list, using a space to separate the file names, like this:

```
...
# Checks for header files.
AC_HEADER_STDC
AC_CHECK_HEADERS([stdlib.h pthread.h])
...
```

I like to make my packages available to as many people as possible, so let's go ahead and use the dual-mode build approach, where we can at least provide *some* form of jupiter to users without pthreads. To accomplish this, we'll need to add some conditional compilation to our `main.c` file:

```
#include <stdio.h>
#include <stdlib.h>

#if HAVE_PTHREAD_H
# include <pthread.h>
#endif

static void * print_it(void * data)
{
    printf("Hello from %s!\n", (char *)data);
    return 0;
}

int main(int argc, char * argv[])
{
#if HAVE_PTHREAD_H
    pthread_t tid;
    pthread_create(&tid, 0, print_it, argv[0]);
    pthread_join(tid, 0);
#else
    print_it(argv[0]);
#endif
    return 0;
}
```

In this version of `main.c`, we've added a couple of conditional checks for the existence of the header file. The `HAVE_PTHREAD_H` macro will be defined to the value 1 in the `config.h.in` template, if the `AC_CHECK_HEADERS` macro locates the `pthread.h` header file, otherwise the definition will be added as a comment in the template. Thus, we'll need to include the `config.h` file at the top of our `main.c` file:

```
#if HAVE_CONFIG_H
# include <config.h>
#endif
...
```

Recall that `HAVE_CONFIG_H` must be defined on the compiler command line. Autoconf populates the `@DEFS@`, substitution variable with this definition, if `config.h` is available. If you choose not to use the `AC_CONFIG_HEADERS` macro in your `configure.ac`, then `@DEFS@` will contain all of the definitions generated by

all of the various check macros you do use. In our example here, we've used `AC_CONFIG_HEADERS`, so our `config.h.in` template will contain most of these definitions, and `@DEFS@` will only contain `HAVE_CONFIG_H`. This is actually a nice way to go because it significantly shortens the compiler command line. An additional benefit is that it becomes very simple to take a snapshot of the template, and modify it by hand for non-Autotools platforms, such as Microsoft Windows, which doesn't require as dynamic a configuration process as does Unix/Linux. We'll make the necessary change to our `src/Makefile.in` template, like this:

```
...
# Tool-related substitution variables
CC          = @CC@
DEFS        = @DEFS@
LIBS        = @LIBS@
CFLAGS      = @CFLAGS@
CPPFLAGS    = @CPPFLAGS@
...
jupiter: main.c
    $(CC) $(CFLAGS) $(DEFS) $(CPPFLAGS) \
    -I. -I$(srcdir) -I..\
    -o $@ $(srcdir)/main.c $(LIBS)
...
```

Now, we have everything we need to conditionally build `jupiter`. If the end-user's system has `pthread` functionality, she'll get a version of `jupiter` that uses multiple threads of execution, otherwise, she'll have to settle for serialized execution. The only thing left is to add some code to the `configure.ac` script that displays a message during configuration, indicating that we're defaulting to serialized execution if the library is not found.

Another point to consider here is what it means to have the header file installed, but no library. This is very unlikely, but it can happen. However, this is easily remedied by simply skipping the header file check entirely if the library isn't found. We'll reorganize things a bit to handle this case also:

```
...
# Checks for libraries.
have_pthreads=no
AC_SEARCH_LIBS([pthread_create], [pthread],
    [have_pthreads=yes])

# Checks for header files.
AC_HEADER_STDC
AC_CHECK_HEADERS([stdlib.h])

if test "x${have_pthreads}" = xyes; then
    AC_CHECK_HEADERS([pthread.h], [],
        [have_pthreads=no])
fi

if test "x${have_pthreads}" = xno; then
    echo "-----"
```

```

    echo " Unable to find pthreads on this system. "
    echo " Building a single-threaded version. "
    echo "-----"
fi
...

```

Let's run `autoreconf` and `configure` and see what additional output we get now:

```

$ autoreconf
$ ./configure
checking for gcc... gcc
...
checking for library... pthread_create... -lpthread
...
checking pthread.h usability... yes
checking pthread.h presence... yes
checking for pthread.h... yes
configure: creating ./config.status
config.status: creating Makefile
...

```

Of course, if your system doesn't have pthreads, you'll get something a little different. To emulate this, I'll rename my pthreads libraries (both shared and static), and then rerun `configure`:

```

$ su
Password:
# mv /usr/lib/libpthread.so ...
# mv /usr/lib/libpthread.a ...
# exit
exit
$ ./configure
checking for gcc... gcc
...
checking for library... pthread_create... no
...
checking for stdint.h... yes
checking for unistd.h... yes
checking for stdlib.h... (cached) yes
-----
Unable to find pthreads on this system.
Building a single-threaded version.
-----
configure: creating ./config.status
config.status: creating Makefile
config.status: creating src/Makefile
config.status: creating config.h

```

Of course, if we had chosen to fail the build if we couldn't find the pthread header file or libraries, then our source code would have been simpler—no need for conditional compilation. We could change our `configure.ac` file to look like this, instead:

```

...
# Checks for libraries.
have_pthreads=no

```

```

AC_SEARCH_LIBS([pthread_create], [pthread],
               [have_pthreads=yes])

# Checks for header files.
AC_HEADER_STDC
AC_CHECK_HEADERS([stdlib.h])

if test "x${have_pthreads}" = xyes; then
    AC_CHECK_HEADERS([pthread.h], [],
                    [have_pthreads=no])
fi

if test "x${have_pthreads}" = xno; then
    echo "-----"
    echo " The pthread library and header file is "
    echo " required to build jupiter. Stopping... "
    echo " Check 'config.log' for more information. "
    echo "-----"
    (exit 1); exit 1;
fi
...

```

I could have used a couple of macros provided by Autoconf for the purpose of printing messages to the console:

`AC_MSG_WARNING` and `AC_MSG_ERROR`, but I don't really care for these macros, because they tend to be single-line-oriented. This is especially a problem in the case of the warning message, which merely indicates that we're continuing, but we're building a single-threaded version of jupiter. Such a single-line message could zip right by in a large configuration process, without even being noticed by the user.

In the case where we decide to terminate with an error, this is less of a problem, because—well, we terminated. But, for the sake of consistency, I like all of my messages to look the same. There is a note in the GNU Autoconf manual that indicates that some shells are not able to properly pass the value of the `exit` parameter to the shell, and that `AC_MSG_ERROR` has a work-around for this problem. Well, the funny code after the `echo` statements in this last example *is* this very work-around, copied right out of a test `configure` script that I created that uses `AC_MSG_ERROR`.

This last topic brings to light a general lesson regarding Autoconf checks. Checks do just that—they check. It's up to the maintainer to add code to do something based on the results of the check. This isn't strictly true, as `AC_SEARCH_LIBS` adds a library to the `@LIBS@` variable, and `AC_CHECK_HEADERS` adds a preprocessor definition to the `config.h.in` template. However, regarding the flow of control within the `configure` process, all such decisions are left to the developer. Keep this in mind while you're designing your `configure.ac` script, and life will be simpler for you.

## Supporting optional features and packages

Alright, we've covered the cases in jupiter where a pthread

library exists, and where it doesn't exist. I'm satisfied, at this point, that we've done just about all we can to manage both of these cases very well. But what about the case where the user wants to deliberately build a single-threaded version of jupiter, even in the face of an existing pthread library? Do we add a note to jupiter's README file, indicating that the user should rename her pthread libraries in this case? I don't think so.

Autoconf provides for both optional features, and optional sub-packages with two new macros: `AC_ARG_ENABLE` and `AC_ARG_WITH`. These macros are designed to do two things: First, to add help text to the output generated when you enter `configure --help`, and second, to check for the specified options, `--enable-feature [=yes|no]`, and `--with-package [=arg]` on the configure script's command line, and set appropriate environment variables within the script. The values of these variables may be used later in the script to set or clear various preprocessor definitions or substitution variables.

`AC_ARG_WITH` is used to control the use of optional sub-packages which may be consumed by your package. `AC_ARG_ENABLE` is used to control the inclusion or exclusion of optional features in your package. The choice to use one or the other is often a matter of perspective and sometimes simply a matter of preference, as they provide somewhat overlapping sets of functionality. For instance, in the jupiter package, it could be justifiably argued that jupiter's use of pthreads constitutes the use of an external package. However, it could just as well be said that asynchronous processing is a feature that might be enabled.

In fact, both of these statements are true, and which type of option you use should be dictated by a high-level architectural perspective on the software in question. For example, the pthread library supplies more than just thread creation functions. It also provides mutexes and condition variables, both of which may be used by a library package that doesn't create threads. If a project provides a library that needs to act in a thread-safe manner within a multi-threaded process, then it will probably use one or more mutex objects. But it may never create a thread. Thus, a user may choose to disable asynchronous execution within this library package at configuration time, but the package may still need to link the pthread library in order to access the mutex functionality from an unrelated portion of the code.

From this perspective, it makes more sense to specify `--enable-async-exec` than `--with-pthreads`. Indeed, from a purist's perspective, this rationale is always sound, even in cases where a project only uses pthreads to create threads. When writing software, you won't often go wrong by siding with the purist. While some of their choices may seem arbitrary—even ridiculous, they're almost always vindicated at some point in the future.

So, when do we use `AC_ARG_WITH`? Generally, when a choice should be made between implementing functionality one way or another. That is, when there is a choice to use one package or another, or to use an external package, or an

internal implementation. For instance, if jupiter had some reason to encrypt a file, it might be written to use either an internal encryption algorithm, or an external package, such as openssl. When it comes to encryption, the use of a widely understood package can be a great boon toward gaining community adoption of your package. However, it can also be a hindrance to those who don't have access to a required external package. Giving your users a choice can make all the difference between them having a good or bad experience with your package.

These two macros have very similar signatures, so I'll just list them here together:

```
AC_ARG_WITH(package, help-string,
  [action-if-given], [action-if-not-given])

AC_ARG_ENABLE(feature, help-string,
  [action-if-given], [action-if-not-given])
```

As with many Autoconf macros, these may be used in a very simple form, where the check merely sets environment variables:

- `${withval}` and `${with_package}`
- `${enableval}` and `${enable_feature}`

They can also be used in a more complex form, where these environment variables are used by shell script in the optional arguments. In either case, as usual, the resulting variable must be used to act on the results of the check, or performing the check is pointless.

### Coding up the feature

Okay, we've decided that we should use `AC_ARG_ENABLE`. Do we enable or disable the "async-exec" feature by default? The difference in how these two cases are encoded is limited to the help text and to the shell script that we put into the `action-if-not-given` argument. The help text describes the available options and the default value, and the shell script indicates what we want to do if the option is NOT specified. Of course, if it is specified, we don't need to assume anything.

Say we decide that asynchronous execution is a risky feature. In this case, we want to disable it by default, so we might add code like this to our `configure.ac` script:

```
AC_ARG_ENABLE([async-exec],
  [ --enable-async-exec      enable async exec],
  [async_exec=${enable_val}],
  [async_exec=yes])
```

On the other hand, if we decide that asynchronous execution is a fairly fundamental part of jupiter, then we'd like it to be enabled by default. In this case we'd use code like this:

```
AC_ARG_ENABLE([async-exec],
```

```

[ --disable-async-exec    disable async exec],
[ async_exec=${enable_val}],
[ async_exec=no])

```

There are a couple of really neat features of this macro that I'd like to point out:

- Regardless of the help text, the user may always use the syntactical standard formats, `--enable-option[=yes|no]` or `--disable-option[=yes|no]`. In either case, the `[=yes|no]` portion is optional.
- Inverse logic is handled transparently—that is, the value of `${enableval}` always represents the user's answer to the question, "Should it be enabled?". For instance, even if the user enters something like `--disable-option=no`, the value of `${enableval}` will still be set to `yes`.

These features of `AC_ARG_ENABLE` and `AC_ARG_WITH` make a maintainer's life a *lot* simpler.

Now, the only remaining question is, do we check for the library and header file, regardless of the user's desire for this feature, or do we only check for them if the user indicates that the "async-exec" feature should be enabled. Well, in this case, it's purely a matter of preference, as we're using the pthread library only for this feature. Again, if we were also using the pthread library for non-feature-specific reasons, then this question would be answered for us.

In cases where we need the library even if the feature is disabled, we add the `AC_ARG_ENABLE` macro, as in the example above, and then an additional `AC_DEFINE` macro to define a `config.h` definition specifically for this feature. Since we don't really want to enable the feature if the library or header file is missing—even if the user specifically requested it—we also need to add some shell code to turn the feature off if either of these are missing:

```

...
# Checks for headers.
AC_HEADER_STDC

# Checks for command line options
AC_ARG_ENABLE([async-exec],
  [ --disable-async-exec    disable async exec],
  [ async_exec=${enableval}],
  [ async_exec=yes])

have_pthreads=no
AC_SEARCH_LIBS([pthread_create], [pthread],
  [have_pthreads=yes])

if test "x${have_pthreads}" = xyes; then
  AC_CHECK_HEADERS([pthread.h], [],
    [have_pthreads=no])
fi

```

```

if test "x${have_pthreads}" = xno; then
  if test "x${async_exec}" = xyes; then
    echo "-----"
    echo "Unable to find pthreads on this system."
    echo "Building a single-threaded version.  "
    echo "-----"
  fi
  async_exec=no
fi

if test "x${async_exec}" = xyes; then
  AC_DEFINE([ASYNC_EXEC], 1, [async exec enabled])
fi

# Checks for headers.
AC_CHECK_HEADERS([stdlib.h])
...

```

I've also added an additional test for a `yes` value in `async_exec` around the `echo` statements within the last test for `have_pthreads`. The reason for this is that this text really belongs to the feature, not the `pthreads` library test. Remember, we're trying to create a logical separation between testing for `pthreads`, and testing for the requirements of the feature.

Of course, now we also have to modify `main.c` such that it uses this new definition, as follows:

```

...
#if HAVE_PTHREAD_H
# include <pthread.h>
#endif

static void * print_it(void * data)
{
    printf("Hello from %s!\n", (char *)data);
    return 0;
}

int main(int argc, char * argv[])
{
    #if ASYNC_EXEC
        pthread_t tid;
        pthread_create(&tid, 0, print_it, argv[0]);
        pthread_join(tid, 0);
    #else
        print_it(argv[0]);
    #endif
    return 0;
}

```

Notice that I left the `HAVE_PTHREAD_H` check around the inclusion of the header file. This is so as to facilitate the use of `pthread.h` in other ways besides for this feature.

In order to check for the library and header file only if the feature is enabled, we merely have to wrap the original check code in a test of `async_exec`, like this:

```

...
if test "x${async_exec}" = xyes; then
  have_pthreads=no
  AC_SEARCH_LIBS([pthread_create], [pthread],
    [have_pthreads=yes])

  if test "x${have_pthreads}" = xyes; then
    AC_CHECK_HEADERS([pthread.h], [],
      [have_pthreads=no])
  fi

  if test "x${have_pthreads}" = xno; then
    echo "-----"
    echo "Unable to find pthreads on this system."
    echo "Building a single-threaded version.    "
    echo "-----"
    async_exec=no
  fi
fi
...

```

This time, I've removed the test for `async_exec` from the `echo` statements, or more appropriately, I've moved the original check from around the `echo` statements, to around the entire set of checks.

## Checks for typedefs and structures

I've spent a fair amount of time during my career writing cross-platform networking software. One key aspect of networking software is that the data sent in network packets from one machine to another needs to be formatted in an architecture-independent manner. If you're trying to use C-language structures to format network messages, one of the first road blocks you generally come to is the complete lack of basic C-language types that have the same size from one platform to another. The C language was purposely designed such that the sizes of its basic integer types are implementation-defined. The designers did this to allow an implementation to use sizes for `char`, `short`, `int` and `long` that are optimal for the platform. Well, this is great for optimizing software for one platform, but it entirely discounts the need for sized types when moving data *between* platforms.

In an attempt to remedy this shortcoming in the language, the C99 standard provides just such sized types, in the form of the `intX_t` and `uintX_t` types, where `x` may be one of `8`, `16`, `32` or `64`. While many compilers provide these types today, some are still lagging behind. GNU C, of course, has been at the fore front for some time now, providing the C99 sized types along with the `stdint.h` header file in which these types are supposed to be defined. As time goes by, more and more compilers will support C99 types completely. But for now, it's still rather painful to write portable code that uses these and other more recently defined integer-based types.

To alleviate the pain somewhat, Autoconf provides macros for determining whether such integer-based types exist on a

user's platform, defining them appropriately if they don't exist. To ensure, for example, that `uint16_t` exists on your target platforms, you may use the following macro expansion in your `configure.ac` file:

```
AC_TYPE_UINT16_T
```

This macro will ensure that either `uint16_t` is defined in the appropriate header files (`stdint.h`, or `inttypes.h`), or that `uint16_t` is defined in `config.h` to an appropriate basic integer type that actually *is* 16 bits in size and unsigned in nature.

The compiler tests for such integer-based types is done almost universally by a generated configure script using a bit of C code that looks like this:

```
...
int main()
{
    static int test_array
        [1 - 2 * !((uint16_t) -1 >> (16 - 1) == 1)];
    test_array[0] = 1;
    return 0;
}
```

Now, if you study this code carefully, you'll notice that the important line is the one on which `test_array` is declared (*Note that I've wrapped this line for publication purposes*). Autoconf is relying on the fact that all C compilers will generate an error if you attempt to define an array with a negative size. An even more thorough examination of the bracketed expression will prove to you that this expression really is a compile-time expression. I don't know if this could have been done with simpler syntax or not, but it's a fact proven over the last several years, that this code does the trick on all compilers currently supported by Autoconf—which is most of them. The array is defined with a non-negative size if (and only if) the following two conditions are met:

1. `uint16_t` is in fact defined in one of the included header files.
2. the actual size of `uint16_t` really is 16 bits; no more, no less.

Code that relies on the use of this macro might contain the following construct:

```
#if HAVE_CONFIG_H
# include <config.h>
#endif
#if HAVE_STDINT_H
# include <stdint.h>
#endif
...
#if defined UINT16_MAX || defined uint16_t
// code using uint16_t
#else
// complicated alternative using >16-bit unsigned
#endif
```

There are a few dozen such type checks available in Autoconf. You should familiarize yourself with Section 5.9 of the GNU Autoconf manual, so that you have a working knowledge of what's available. I recommend you don't commit such checks to memory, but rather just know about them, so that they'll come to mind when you need to use them. Then go look them up for the exact syntax, when you do need them.

In addition to these type-specific checks, there is also a generic type check macro, `AC_CHECK_TYPES`, which allows you to specify a comma-separated list of questionable types that your project needs. Note that this list is comma-separated, not space separated, as in the case of most of these sorts of check lists. This is because type definitions (like `struct fooble`) may have embedded spaces. Since they are comma-delimited, you will need to *always* use the square bracket quotes around this parameter (if, that is, you list more than one type in the parameter).

```
AC_CHECK_TYPES(types, [action-if-
found], [action-if-not-found],
[includes = 'default-includes'])
```

If you don't specify a list of include files in the last parameter, then the default includes are used in the compiler test. The default includes are used via the macro `AC_INCLUDES_DEFAULT`, which is defined as follows (in version 2.62 of Autoconf):

```
#include <stdio.h>
#ifdef HAVE_SYS_TYPES_H
# include <sys/types.h>
#endif
#ifdef HAVE_SYS_STAT_H
# include <sys/stat.h>
#endif
#ifdef STDC_HEADERS
# include <stdlib.h>
# include <stddef.h>
#else
# ifdef HAVE_STDLIB_H
# include <stdlib.h>
# endif
#endif
#ifdef HAVE_STRING_H
# if !defined STDC_HEADERS && defined HAVE_MEMORY_H
# include <memory.h>
# endif
# include <string.h>
#endif
#ifdef HAVE_STRINGS_H
# include <strings.h>
#endif
#ifdef HAVE_INTTYPES_H
# include <inttypes.h>
#endif
#ifdef HAVE_STDINT_H
# include <stdint.h>
#endif
#ifdef HAVE_UNISTD_H
```

```
# include <unistd.h>
#endif
```

If you know that your type is not defined in one of these header files, then you should specify one or more include files to be included in the test, like this:

```
AC_CHECK_TYPES([struct doodah], [], [], [
#include<doodah.h>
#include<doodahday.h>])
```

The interesting thing to note here is the way I wrapped the last parameter of the macro over three lines in `configure.ac`, with no indentation. This time I didn't do it for publication reasons. This text is included verbatim in the test source file. Since some compilers have a problem with placing the POUND sign anywhere but the first column it's a good idea to tell Autoconf to start each include line in column one, in this manner.

Admittedly, these are the sorts of things that developers complain about regarding Autoconf. When you do have problems with such syntax, your best friend is the `config.log` file, which contains the exact source code for all failed tests. You can simply look at this log file to see how Autoconf formatted the test, possibly incorrectly, and then fix your check in `configure.ac` accordingly.

## The AC\_OUTPUT macro

The `AC_OUTPUT` macro expands into the shell code that generates the `configure` script, based on all the data specified in all previous macro expansions. The important thing to note here is that all other macros must be used before `AC_OUTPUT` is expanded, or they will be of little value to your `configure` script.

Additional shell script may be placed in `configure.ac` after `AC_OUTPUT` is expanded, but this additional code will not affect the configuration or the file generation performed by `config.status`.

I like to add some echo statements after `AC_OUTPUT` to indicate to the user how the system is configured, based on their specified command line options, and perhaps additional useful targets for `make`. For example, one of my projects has the following text after `AC_OUTPUT` in `configure.ac`:

```
echo \
"-----

${PACKAGE_NAME} Version ${PACKAGE_VERSION}

Prefix: '${prefix}'.
Compiler: '${CC} ${CFLAGS} ${CPPFLAGS}'

Package features:
  Async Execution: ${async_exec}

Now type 'make @<:@<target>@:>@'
  where the optional <target> is:
```

```

all             - build all binaries
install        - install everything

```

-----"

This is a really handy `configure` script feature, as it tells the user at a glance just what happened during configuration. Since variables such as `debug` are set on or off based on configuration, the user can see if the configuration he asked for actually took place.

By the way, in case you're wondering what those funny character sequences are around the word `<target>`, they're called quadrigraph sequences or simply quadrigraphs, and serve the same purpose as escape sequences. Quadrigraphs are a little more reliable than escaped characters, or escape sequences because they're never subject to ambiguity. They're converted to proper characters at a very late stage by M4, and so are not subject to mis-interpretation.

The sequence, `@<:@` is the quadrigraph sequence for the open square bracket (`[`) character, while `@:>@` is the quadrigraph for the close square bracket (`]`) character. These quadrigraphs will *always* be output by Autoconf (M4) as literal bracket characters.

There are a few other quadrigraphs. We'll see some of them in Chapter 9 when we begin to examine the process of writing our own Autoconf macros. If you're interested, check out section 8.1.5 of the GNU Autoconf manual.

*NOTE: Version 2.62 of Autoconf does a much better job of deciphering the user's intent with respect to the use of square brackets than previous versions of Autoconf. Where you might have needed to use a quadrigraph in the past to force Autoconf to display a square bracket, you may now use the character itself. Most of the problems that occur are a result of not properly quoting arguments.*

## Does (project) size matter?

An issue that might have occurred to you by now is the size of my toy project. I mean, c'mon! One source file?! But, I've used `autoscan` to autoconfiscate projects with several hundred C++ source files, and some pretty complex build steps. It takes a few seconds longer to run `autoscan` on a project of this size, but it works just as well. For a basic build, the generated `configure` script only needed to be touched up a bit—project name, version, etc.

To add in compiler optimization options for multiple target tool sets, it took a bit more work. I'll cover these sorts of issues in Chapter 6 where I'll show you how to autoconfiscate a real project.

## Summary

In this chapter, I've covered about a tenth of the information in the GNU Autoconf manual, but in much greater analytical

detail than the manual. For the developer hoping to quickly bootstrap into Autoconf, I believe I've covered one of the more important tenths. But this statement in no way alleviates a responsible software engineer from studying the other nine tenths, as time permits, of course.

For example, I didn't go into detail about the differences between searching for a function and searching for a library. In general, `AC_SEARCH_LIBS` should be used to check for a function you need, but expect in one or more libraries. The `AC_FUNC_*` macros are available to check for very specific portability-related functionality, such as `AC_FUNC_ALLOCA`, which exists on some platforms, but not others. The `AC_CHECK_FUNC` macro is used for determining if a function not supported by one of the more specific `AC_FUNC_*` macros is available in the C (or C++) standard libraries. I recommend reading through Section 5.5 of the GNU Autoconf manual to familiarize yourself with what's available within these special function checks.

Another topic on which I didn't spend much time was that of checking for compiler characteristics. Section 5.10 covers these issues completely. Given what you've learned after reading this chapter, reading these sections of the manual should be pretty straight-forward.

In fact, once you're comfortable with the material in this and the preceding chapters of this book, I'd highly recommend spending a fair amount of time in Chapter 5 of the GNU Autoconf manual. Doing so will make you the Autoconf expert you never thought you could be, by filling in all of the missing details.

The next chapter takes us away from Autoconf for a while, as we get into Automake, an Autotools tool chain add-on enhancement for the make utility.

[◀ Chapter 2: Project management and the GNU coding standards](#)

u  
p

[Chapter 4: Automatically writing makefiles with Automake ▶](#)

[Login or register to post comments](#) 1647 reads



- COLUMNS
- COMMUNITY POSTS
- ISSUES
- BOOKS
- FORUM
- FS NEWS

Ads by Google

**IT Support London & SE**  
Local, cost effective IT support for your 2003 or SBS network  
www.1st-solution.co.uk

**Get NetOp Software**  
Free Trials | Gold Partner for full range of NetOp Products  
www.bluebeamsecurity.com

**Free Windows Server 2003**  
89 Page Reference Guide. Download Now.  
www.PreLogic.com

**Pump Software**  
Streamline your pump sales. Select, configure, quote, and order online.  
www.bigmachines.com

**Product Configurator**  
Configure, Calculate, Visualize Generate Sales Documents  
www.sofon.com

- [Microsoft May Seek New Yahoo ...](#)
- [Building Semantics is Different ...](#)
- [Lux: multi-touch for OS X](#)

Sponsored by **Novell**

Best voted contents

- [The Bizarre Cathedral - 3](#)  
Ryan Cartwright, 2008-05-05
- [The Bizarre Cathedral - 2](#)  
Ryan Cartwright, 2008-04-27
- [Indexing offline CD-ROM archives](#)  
Terry Hancock, 2008-05-03
- [Microsoft and free software? I don't think so...](#)  
Terry Hancock, 2008-04-26

Home » Autotools: a practitioner's guide to Autoconf, Automake and Libtool

# Chapter 4: Automatically writing makefiles with Automake



by John Calcote

Most of the general complaints I've ever seen aimed at the Autotools are ultimately associated with Automake, in the final analysis. The reason for this is simple: Automake provides the highest level of abstraction over the build system. This high level of abstraction is both apparent, and actual. And yet a solid understanding of the inner workings of Automake can provide us with the one of the most satisfying auto-generated build system experiences, because we can feel comfortable using the features of Automake to their maximum potential, and extending it where our projects require.

Shortly after Autoconf was well on its way to success in the GNU world, David MacKenzie began work on a new tool—a tool for automatically generating makefiles for a GNU project. MacKenzie's work on Automake lasted about a year during 1994, ending around November of that year. A year later, during November of 1995, Tom Tromej (of RedHat and Cygnus fame) took over development of the Automake project. Tromej really had very much a defining role in Automake. In fact, although MacKenzie wrote the initial version of Automake in Bourne shell script, Tromej completely rewrote the tool in Perl over the following year. Tromej continued to maintain and enhance Automake during the next 5 years.

*NOTE: Do not confuse the requirements of Automake on the project maintainer with the requirements of a generated build system on the end user. Perl is required by Automake, not by the generated build system.*

Around February of 2000, Alexandre Duret-Lutz began to take a more active role in the development of the Automake project, and by the end of that year, had pretty much taken over project maintenance. Duret-Lutz's role as project lead lasted until about mid-2007. Since then, the project has been maintained by Eric Blake of the Free Software Foundation (FSF), with strong leadership (and most of the repository check-in's, for that matter) from automake mailing list contemporaries such as Ralf Wildenhues and Akim Demaille. *(Many heartfelt thanks to Ralf for kindly answering so many seemingly trivial questions while I worked on this book.)*

Sometime early during development of the GNU Coding Standards (GCS), it became clear to MacKenzie that much of a GNU project makefile was fairly boilerplate in nature. This is because the GCS guidelines are fairly specific about how and where a project's products should be built, tested, and installed. These conditions have allowed Automake syntax to be concise—in fact, it's terse, almost to a fault. One Automake statement represents a *lot* of functionality. The nice thing, however, is that once you understand it, you can get a fairly complete, complex and functionally correct build system up and

**Free Windows Server 2003**  
89 Page Reference Guide. Download Now.  
www.PreLogic.com

**Cisco NetFlow Challenge**  
Immediate results with NetFlow. Take our analysis challenge.  
Plixer.com/NetFlowChallenge

**Cisco**  
Expert design, installation, and support for your Cisco network  
www.L3n.co.uk

**Compiler consulting**  
Custom-built compilers and tools, support and maintenance services.  
www.excelsior-usa.com



http://www.freesoftwaremagazine.com (c) Copyright 2008 Ryan Cartwright CC: BY-NC-SA

From the FSM staff...

[The Top 10 Everything](#) (Dave). The good, the bad and the ugly.

[Free Software news](#) (Dave & Bridget). A site about short stories and writing.

[Book Reviews: Illiterarty](#) (Bridget). Book reviews, blogs, and short stories.

Hot topics - last 60 days

[Installing an all-in-one printer device in Debian](#)  
Ryan Cartwright, 2008-05-05

[What is the free software community?](#)  
Tony Mobily, 2008-03-29

[Things you miss with GNU/Linux](#)  
Ryan Cartwright, 2008-05-01

[How do you replace Microsoft Outlook? Groupware applications](#)  
Ryan Cartwright, 2008-03-20

[Drigg \(the pligg alternative\) on Debian](#)



Buzz authors

[Marco Marongiu](#)

[Pieter Hintjens](#)

[drascus321](#)

running in short order—I mean on the order of minutes, not hours or days.

## Getting down to business

Let's face it, writing a makefile is hard. Oh, the initial writing is fairly simple, but getting it right is often very difficult—the devil, as they say, is in the details. Like any high-level programming language, make syntax is often conducive to formulaic composition. That's just a fancy way of saying that once you've solved a “make problem”, you're inclined to memorize the solution and apply the same formula the next time that problem crops up—which happens quite often when writing build systems.

So what advantages does Automake give us over our hand-coded Makefile.in templates, anyway? Well, that's pretty easy to answer with a short example. Consider the following changes to the files in our project directory structure (these commands are executed from jupiter's top-level directory):

```
$ rm autogen.sh Makefile.in src/Makefile.in
$ echo "SUBDIRS = src" > Makefile.am
$ echo "bin_PROGRAMS = jupiter
> jupiter_SOURCES = main.c" > src/Makefile.am
$ touch NEWS README AUTHORS ChangeLog
$ vi configure.ac
...
AC_INIT([Jupiter], 1.0, [bugs@jupiter.org])
AM_INIT_AUTOMAKE
AC_CONFIG_SRCDIR([src/main.c])
...
$ autoreconf -i
$
```

The “rm” command deletes our hand-coded Makefile.in templates and the autogen.sh script we wrote to ensure that all the support scripts and files were copied into the root of our project directory. We won't be needing this script anymore because we're upgrading jupiter to Automake proper.

For the sake of brevity in the text, I used “echo” statements to write the new Makefile.am files, but you may, of course, use an editor if you wish. *NOTE: There is a hard carriage-return after “bin\_PROGRAMS = jupiter” in the third line. The shell will continue to accept input after the carriage return until the quotation is closed on the following line.*

The “touch” command is used to create new empty versions of the NEWS, README, AUTHORS and ChangeLog files in the project root directory. These files are required by the GCS for all GNU projects. While they're NOT required for non-GNU programs, they've become something of an institution in the FOSS world—you'd do well to have these files, properly formatted, in your project, as users have come to expect them. The GCS document covers the format and contents of these files. Section 6 covers the NEWS and ChangeLog files, and Section 7 covers the README and INSTALL files. The AUTHORS file is a list of people (names and optional email addresses) to whom attribution should be given.

## Enabling Automake in configure.ac

Finally, I've added a single line to the configure.ac file, “AM\_INIT\_AUTOMAKE” between the AC\_INIT and

dogboi

Richard Rothwell

## All news

John Crawford: Arizona LoCo Hardy Release Party

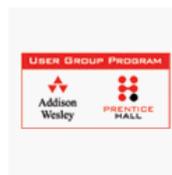
John Crawford: Ubuntu Weekly Newsletter #91

Matthew Helmke: Morocco blocks Google Maps

Andy Price: Synchronicity and Sharing

Jono Bacon: Beware False Prophets

more



**LINUXUSER**  
*Download edition available now!*

**alternative) vs. Piigg: why should people switch?**

Tony Mobily, 2008-04-13

Hot topics - last 21 days

**Installing an all-in-one printer device in Debian**

Ryan Cartwright, 2008-05-05

**Things you miss with GNU/Linux**

Ryan Cartwright, 2008-05-01

**Digital Rights Management (DRM): is it in its death throes?**

Gary Richmond, 2008-05-07

**Open letter to standards professionals, developers, and activists**

Pieter Hintjens, 2008-05-13



DEDICATED SERVER

Odiogo



- 15 How Microsoft Uses Novell to Fight GNU/Linux, Xen to Fight VMWare and GNU/Linux
- 15 Why You Should Reject Novell's Moonlight
- 17 University of Havana Finally Switches to Free Software
- 15 Please Welcome Digistan
- 17 Content Protection madness on Vista
- 25 Asus to embed Linux into all motherboards
- 23 Make Your Distro Free of Miguel de Icaza's junk code
- 16 Firefox 3 Release Candidate now available for download

Get this widget »

AC\_CONFIG\_SRCDIR statements. Besides the normal requirements of an Autoconf input file, this is the *only* line that's required to enable Automake in a project that's already configured with Autoconf. The AM\_INIT\_AUTOMAKE macro accepts an optional argument—a white-space separated list of option tags, which can be passed into this macro to modify the general behavior of Automake. The following is a comprehensive list of options for Automake version 1.10:

- gnits
- gnu
- foreign
- cygnus
- ansi2knr
- path/ansi2knr
- check-news
- dejagnu
- dist-bzip2
- dist-lzma
- dist-shar
- dist-zip
- dist-tarZ
- filename-length-max=99
- no-define
- no-dependencies
- no-dist
- no-dist-gzip
- no-exeext
- no-installinfo
- no-installman
- nostdinc
- no-texinfo.tex
- readme-alpha
- std-options
- subdir-objects
- tar-v7
- tar-ustar
- tar-pax
- <version>
- -W<category>
- --warnings=<category>

I won't spend a lot of time on the option tag list at this point. For a detailed description of each option, check out Chapter 17 of the GNU Automake manual. I will, however, point out a few of the most useful options.

The `check-news` option will cause "make dist" to fail if the current version doesn't show up in the first few lines of the `NEWS` file. The `dist-*` tags can be used to change the default distribution package type. Now, these are handy because often developers want to distribute tar.bz2 files, rather than tar.gz files. By default, "make dist" builds a tar.gz file. You can override this by using "make dist-bzip2", but this is more painful than it needs to be for projects that like to use bzip2 by default. The `readme-alpha` option can be used to temporarily alter the behavior of the build and distribution process during alpha releases of a project. First, a file named `README-alpha`, found in the project root directory, will be distributed automatically while using this option. This option will also alter the expected versioning scheme of the project.

The `<version>` option is actually a placeholder for a numeric version number. This value represents the lowest

version number of Automake that is acceptable for this project. For instance, if 1.10 is passed as a tag, then Automake will fail if it's version is less than 1.10. The `-w<category>` and `--warnings=<category>` options indicate that the project would like to use Automake with various warning categories enabled.

### What we get from Automake

The last line of the example executes the `"autoreconf -i"` command, which, as we've already discussed in prior chapters, regenerates all Autotools-generated files according to the `configure.ac` file. This time, with the inclusion of the `AM_INIT_AUTOMAKE` statement, the `"-i"` option properly tells Automake to add any missing files. The `"-i"` option need only be used once in a newly checked out work area. Once the missing utility files have been added, the `"-i"` option may be dropped.

These few commands create for us an Automake-based build system containing everything that we wrote into our original `Makefile.in` templates, *except that this one is more correct and functionally complete*. A quick glance at the resulting generated `Makefile.in` template shows us that, from just a couple of input lines, Automake has done a significant amount of work for us. The resulting top-level `Makefile.in` template (remember, the `configure` script turns these templates into `Makefiles`), is nearly 18K in size. Our original files were only a few hundred bytes long.

A generated Automake build system supports the following important `make` targets—and this list is *not* comprehensive:

- `all`
- `distdir`
- `install`
- `install-strip`
- `install-data`
- `install-exec`
- `uninstall`
- `install-dvi`
- `install-html`
- `install-info`
- `install-ps`
- `install-pdf`
- `installdirs`
- `check`
- `installcheck`
- `mostlyclean`
- `clean`
- `distclean`
- `maintainer-clean`
- `dvi`
- `pdf`
- `ps`
- `info`
- `html`
- `tags`
- `ctags`
- `dist`
- `dist-bzip2`
- `dist-gzip`
- `dist-lzma`

- `dist-shar`
- `dist-zip`
- `dist-tarZ`
- `uninstall`

As you can see, this goes a bit beyond what we provided in our hand-coded `Makefile.in` templates. And Automake writes all of this functionality automatically, correctly and quickly for each project that you instrument in the manner outlined above.

### So, what's in a `Makefile.am` file?

You'll no doubt recall from Chapter 3 that Autoconf accepts shell script, sprinkled with M4 macros, and generates the same shell script with those macros fully expanded into additional shell script. Likewise, Automake accepts as input a makefile, sprinkled with Automake commands. As with Autoconf, the significance of this statement is that Automake input files are nothing more or less than makefiles with additional syntax.

One very significant difference between Autoconf and Automake is that Autoconf generates NO output text except for the existing shell script in the input file, plus any additional shell script resulting from the expansion of embedded M4 macros. Automake, on the other hand, assumes that all makefiles should contain a minimal infrastructure designed to support the GCS, in addition to any targets and variables that you specify.

To illustrate this point, I'll create a `temp` directory in the root of the `jupiter` project, and add an *empty* `Makefile.am` file to that directory. Then I'll add this new `Makefile.am` to my project, like this:

```
$ mkdir temp
$ touch temp/Makefile.am
$ echo "SUBDIRS = src temp" > Makefile.am
$ vi configure.ac
...
AC_CONFIG_FILES([Makefile
                  src/Makefile
                  temp/Makefile])
...
$ autoreconf
$ ./configure
...
$ ls -lsh temp
total 20K
 12K Makefile
   0 Makefile.am
 8.0K Makefile.in
$
```

Thus we can see that Automake considers a certain amount of support code to be indispensable in every makefile. Even with an empty `Makefile.am` file, we end up with about 12K of code in the resulting `Makefile`, which is generated by `configure` (`config.status`) from an 8K `Makefile.in` template. Incidentally, it's fairly instructive to examine the contents of this `Makefile.in` template to see the Autoconf substitution variables that are passed in, as well as the framework code that Automake generates.

Since the `make` utility uses a fairly rigid set of rules for processing makefiles, Automake takes some minor "literary

license” with your additional make code. Specifically, two basic rules are followed by Automake when generating Makefile.in templates from Makefile.am files that contain additional non-Automake-specific syntax (rules, variables, etc):

1. Make *variables* that you define in your Makefile.am files are placed at the top of the resulting Makefile.in template, immediately following any Automake-generated variable definitions.
2. Make *rules* that you specify in your Makefile.am files are placed at the end of the resulting Makefile.in template, immediately following any Automake-generated rules.

Make doesn't care where rules are located relative to one another, because it reads all of the rules and stores them in an internal database before processing any of them. Variables are treated in a similar manner. To prove this to yourself, try referencing a variable in a makefile before its definition. Make binds values to variable references at the last possible moment, right before command lines containing these references are passed to the shell for execution.

Often, you won't need to specify anything besides a few Automake commands within a given Makefile.am, but there are frequent occasions when you will want to add your own make targets. This is because, while Automake does a lot for you, it can't anticipate *everything* you might wish to do in your build system. It's in this “grey” area where most developers begin to complain about Automake.

We'll spend the rest of this chapter examining the functionality provided by Automake. Later, we'll get into some tricks you can use to significantly enhance existing Automake functionality.

## Analyzing our new build system

Now let's spend some time looking at what we put into those two simple Makefile.am files. We'll start with the top-level file, with its single line of Automake code:

### **Makefile.am**

```
SUBDIRS = src
```

It's pretty easy to divine the primary purpose of this line of text just by looking at the text itself. It appears to be indicating that we have a sub-directory in our project called `src`. In fact, this line tells Automake several things about our project:

1. There are one or more immediate sub-directories containing Makefile.am files to be processed, in addition to this file.
2. Directories in this space-delimited list are to be processed in the order specified.
3. Directories in this list are to be recursively processed for all primary make targets.
4. Directories in this list are to be treated as part of the project distribution.

`SUBDIRS` is clearly a make variable, but it's more than just a make variable. The `SUBDIRS` variable is recognized by Automake to have special meaning, besides the intrinsic meaning associated with make variables. As we continue to study Automake constructs, this theme will come up over and over again. Most Automake statements are, in fact, just make

variables with special meaning to Automake.

Another point about the `SUBDIRS` variable is that it may be used in an arbitrarily complex directory structure, to process `Makefile.am` files within a project. You might say that `SUBDIRS` is the “glue” that links `Makefile.am` files together in a project’s directory hierarchy.

One final point about `SUBDIRS` is that the current directory is *implicitly listed last* in the `SUBDIRS` list, meaning that the current directory will be built *after* all of the directories listed in the `SUBDIRS` variable. You may change this implied ordering if you wish, by using “.” (meaning the current directory) anywhere in the list. This is important because it’s sometimes necessary to build the current directory before one or more subdirectories.

Let’s move down a level now into the `src` directory. The `src/Makefile.am` file contains slightly more code for us to examine; two lines rather than one:

#### **src/Makefile.am**

```
bin_PROGRAMS = jupiter
jupiter_SOURCES = main.c
```

### **Primaries**

The first line, “`bin_PROGRAMS = jupiter`” lists the products generated by this `Makefile.am` file. Multiple files may be listed in this variable definition, separated by white space. The variable name itself is made up of two parts, the installation location, `bin`, and the product type, `PROGRAMS`. GNU Automake documentation calls the product type portion of these variables a “primary”. The following is a list of valid primaries for version 1.10 of Automake:

- `PROGRAMS`
- `LIBRARIES`
- `LISP`
- `PYTHON`
- `JAVA`
- `SCRIPTS`
- `DATA`
- `HEADERS`
- `MANS`
- `TEXINFOS`

*NOTE: Libtool adds `LTLIBRARIES` to the primaries list supported by Automake. We’ll examine this and other Automake extensions provided by Libtool in Chapter 5.*

You could consider primaries to be “product classes”, or types of products that might be generated by a build system. This being the case, it’s pretty clear that not all product classes are handled by Automake. What differentiates one class of product from another? Basically differences in handling semantics during build and installation. `PROGRAMS`, for example are built using different compiler and linker commands than are `LIBRARIES`. Certainly `LISP`, `JAVA` and `PYTHON` products are handled differently—the build system uses entirely different tool chains to build these types of products. And `SCRIPTS`, `DATA` and `HEADERS` aren’t generally even built (although they might be), but rather simply copied into appropriate installation directories.

PROGRAMS also have different execution, and thus installation, semantics from LISP, PYTHON and JAVA programs. Products that fit into the PROGRAMS category are generally executable by themselves, while LISP, JAVA and PYTHON programs require virtual machines and interpreters.

What makes this set of primaries important? The fact that they cover 99 percent of the products created in official GNU projects. If your project generates a set of products that define their own product class, or use a product class not listed in this set of primaries, then you might do well to simply stick with Autoconf until support is added to Automake for your product class. Another option is to add support yourself to Autoconf for your product class, but doing so requires a deep knowledge of both the product class and the Automake perl script. I believe it's fair to say, however, that this set of primaries covers a wide range of currently popular product classes.

## Prefixes

Supported installation locations are provided by the GCS document. This is the same list that I provided to you in Chapter 2. I'll relist them here for convenience:

- bindir
- sbindir
- libexecdir
- datarootdir
- datadir
- sysconfdir
- sharedstatedir
- localstatedir
- includedir
- oldincludedir
- docdir
- infodir
- htmldir
- dvidir
- pdfdir
- psdir
- libdir
- lispdir
- localedir
- mandir
- manNdir

You may have noticed that I left a few entries out of this version of the list. Essentially, all entries ending in `dir` are viable prefixes for Automake primaries. Besides these standard GCS installation locations, three other installation locations are defined by Automake to have enhanced meaning:

- pkglibdir
- pkgincludedir
- pkgdatadir

The `pkg` versions of the `libdir`, `includedir` and `datadir` prefixes are designed to install products into subdirectories of these installation locations that are named after the package. For example, for the `jupiter` project, the `pkglibdir` installation location would be found in `$(exec-prefix)/lib/jupiter`, rather than the usual `$(exec-prefix)/lib` directory.

If this list of installation locations isn't comprehensive enough, don't worry—Automake provides a mechanism for you to define your own installation directory prefixes. Any make variable you define in your Makefile.am file that ends in “dir” can be used as a valid primary prefix. To reuse the example found in the GNU Automake manual, let's say you wish to install a set of XML files into an “xml” directory within the system data directory. You might use this code to do so:

```
xmldir = $(datadir)/xml
xml_DATA = file1.xml file2.xml file3.xml ...
```

Note that the same naming conventions are used with custom installation locations as with the standard locations. Namely, that the variable ends with `dir`, but the `dir` portion of the variable name is left off when using it as a primary prefix.

There are also several prefixes with special meanings not related to installation locations:

- `check`
- `noinst`
- `EXTRA`

The `check` prefix indicates products that are built only for testing purposes, and thus will not be installed at all. Products listed in primary variables that are prefixed with “`check`” aren't even built if the user never types “`make check`”.

The “`noinst`” prefix indicates that the listed products should be built, but not installed. For example, a static so-called “convenience” library might be built as an intermediate product, and used in other stages of the build process to build final products. Such libraries are not designed to be installed, so the prefix shouldn't be an installation location. The “`noinst`” prefix serves this purpose.

The “`EXTRA`” prefix is used to list programs that are conditionally built. This is a difficult concept to explain in a tutorial paragraph, but I'll give it a try: All product files must be listed statically (as opposed to being calculated at build-time) in order for Automake to generate a Makefile.in template that will work for any set of input commands. However, a project maintainer may elect to allow some products to be built conditionally, based on configuration options given to the configure script. If some products are listed in variables generated by the configure script, then these products should also be listed in a primary prefixed with “`EXTRA`”, like this:

```
EXTRA_PROGRAMS = myoptionalprog
bin_PROGRAMS myprog $(optional_programs)
```

Here, it is assumed that the “`optional_programs`” variable is defined in the configure script, and listed in an `AC_SUBST` macro. This way, Automake can know in advance that “`myoptionalprog`” *may* be built, and so generate rules to build it. Any program that may or may not be built, based on configuration options should be specified in “`EXTRA_PROGRAMS`”, so that Automake can generate a makefile that *could* build it if requested to do so.

### “Super” prefixes

Some primaries allow a sort of “super” prefix to be prepended to a prefix/PRIMARY variable specification. Such modifiers may be used together on the same variable where it makes

sense. Thus, these “super” prefixes modify the normal behaviour of a prefix/PRIMARY specification. The existing modifiers include:

- `dist`
- `nodist`
- `nobase`

The `dist` modifier indicates a set of files that should be distributed (that is, included in the distribution package when “make dist” is executed). The `dist` modifier is used with files that are normally not distributed, but may be used explicitly anywhere for clarity. For instance, assuming that some source files for a product should be distributed, and some should not (perhaps they’re generated), the following rules might be used:

```
dist_jupiter_SOURCES = file1.c file2.c
nodist_jupiter_SOURCES = file3.c file4.c
```

While the `dist` prefix is redundant in this example, it is nonetheless useful to the casual reader.

The `nobase` modifier is used to suppress the removal of path information from installed header files that are obtained from subdirectories by a `Makefile.am` file. For instance, assume that installable jupiter project header files exist in a subdirectory of the `src` directory “jupiter”:

```
nobase_dist_include_HEADERS = \
    jupiter/jupiter_interface.h
```

Normally, such a header file would be installed into the `/usr(/local)/include` directory as simply `jupiter_interface.h`. However, if the `nobase` modifier is used, then the extra path information would not be removed, so the final resting place of the installed header would instead be, `/usr(/local)/include/jupiter/jupiter_interface.h`.

Notice also in this example that I combined the use of the `nobase` modifier with that of the `dist` modifier—just to show the concept.

## Product sources

The second line in `src/Makefile.am` is “`jupiter_SOURCES = main.c`”. This variable lists the source files used to build the `jupiter` program. Like product variables made from prefixes and primaries, this type of variable is derived from two parts, the product name, `jupiter` in this case, and the dependent type. I call it the “dependent type” because this variable lists source files on which the product depends. Ultimately, Automake adds these files to make rule dependency lists.

The `EXTRA` prefix may also be used sometimes as a super prefix modifier. When used with a product `SOURCES` variable (eg., `jupiter_SOURCES`), `EXTRA` can be used to specify extra source files that may or may not be used, which are directly associated with the `jupiter` product:

```
EXTRA_jupiter_SOURCES = possibly.c
```

In this case, `possibly.c` may or may not be compiled—perhaps based on an `AC_SUBST` variable.

## Unit tests - supporting “make check”

I mentioned earlier that this Automake-generated build system provided the same functionality as our hand-coded build system. Well, I wasn't completely truthful when I said that. For the most part, that was an accurate statement, but what's still missing is our simple-minded “make check” functionality. The “check” target is indeed supported by our new Automake build system, but it's just not hooked up to any real functionality. Let's do that now.

You'll recall in Chapter 2 that we added code to our `src/Makefile` to run the `jupiter` program and check for the proper output string when the user entered “make check”. We did this with a fairly simple addition to our `src/Makefile`:

```
...
check: all
    ./jupiter | grep "Hello from .*jupiter!"
    @echo "*** ALL TESTS PASSED ***"
...
```

As it turns out, Automake has some solid support for unit tests. Unfortunately, the documentation consists of Chapter 15 of the GNU Automake manual—a single page of text—half of which is focused on the obscure DejaGNU test suite syntax. Nevertheless, adding unit tests to a `Makefile.am` file is fairly trivial. To add our simple “grep test” back into our new Automake-generated build system, I've added a few more lines to the bottom of the `src/Makefile.am` file:

### `src/Makefile.am`

```
bin_PROGRAMS = jupiter
jupiter_SOURCES = main.c
jupiter_CPPFLAGS = -I$(top_srcdir)/common
jupiter_LDADD = ../common/libjupcommon.a

check_SCRIPTS = greptest.sh
TESTS = $(check_SCRIPTS)

greptest.sh:
    echo './jupiter | grep \
    "Hello from .*jupiter!"' > greptest.sh
    chmod +x greptest.sh

CLEANFILES = greptest.sh
```

The “`check_SCRIPTS`” line is clearly a prefixed primary. The “`SCRIPT`” primary indicates a “built” script, or a script that is somehow generated at build time. Since the prefix is “check”, we know that scripts listed in this line will only be built when the user enters “make check” (or “make distcheck”). However, this is as far as Automake goes in supporting such built scripts with Automake-specific syntax. You must supply a make rule for building the script yourself.

Furthermore, since you supplied the rule to generate the script, you must also supply a rule for cleaning the file. Automake provides an extension to the generated “`clean`” rule, wherein all files listed in a special “`CLEANFILES`” variable are added to the list of automatically cleaned files.

The “`TESTS`” line is the important line here, in that it indicates which targets are built and executed when a user enters

“make check”. Since the “check\_SCRIPTS” variable contains a complete list of these targets, I simply reused its value here.

Generating scripts or data files in this manner is a very useful technique. I’ll present some more interesting ways of doing this sort of thing in Chapter 8.

## Adding complexity with convenience libraries

Well, jupiter is fairly trivial, as free software projects go. In order to highlight some more of the key features of Automake, we’re going to have to expand jupiter into something a little bit more complex (if not functional).

We’ll start by adding a convenience library, and having jupiter consume this library. Essentially, we’ll move the code in main.c to a library source file, and then call the function in the library from jupiter’s main routine. Start with the following commands, executed from the top-level project directory:

```
$ mkdir common
$ touch common/jupcommon.h
$ touch common/print.c
$ touch common/Makefile.am
```

Add the following text to the .h and .c files:

### common/jupcommon.h

```
int print_routine(char * name);
```

### common/print.c

```
#include <jupcommon.h>

#if HAVE_CONFIG_H
# include <config.h>
#endif

#include <stdio.h>
#include <stdlib.h>

#if HAVE_PTHREAD_H
# include <pthread.h>
#endif

static void * print_it(void * data)
{
    printf("Hello from %s!\n", (char *)data);
    return 0;
}

int print_routine(char * name)
{
    #if ASYNC_EXEC
        pthread_t tid;
        pthread_create(&tid, 0, print_it, name);
        pthread_join(tid, 0);
    #else
        print_it(name);
    #endif
    return 0;
}
```

As promised, `print.c` is merely a copy of `main.c`, with a couple of small modifications. First, I renamed `main` to `print_routine`, and second, I added the inclusion of the `jupcommon.h` header file at the top. This header file (as you can see) merely provides `print_routine`'s prototype to the new `src/main.c`, where it's called from `main`. Modify `src/main.c` to look like this:

#### **src/main.c**

```
#include <jupcommon.h>

int main(int argc, char * argv[])
{
    print_routine(argv[0]);
    return 0;
}
```

And now for the new `common/Makefile.am` file; add the following text to this file:

#### **common/Makefile.am**

```
noinst_LIBRARIES = libjupcommon.a
libjupcommon_a_SOURCES = jupcommon.h print.c
```

Let's take a look at this file for a minute. You'll recall from our discussion of Automake primaries and prefixes that the first line indicates the products to be built and installed by this `Makefile.am` file. In this case, the `noinst` prefix indicates that this library should not be installed at all. This is because we're creating a "convenience" library, or a library designed solely to make using the source code in the common directory more convenient for two or more consumers. (Granted, we only have one consumer at this point—the jupiter program—but later on we'll add another consumer of this library, and then it will make more sense.)

The library we're creating will be called `libjupcommon.a`—this is a static library, also known as an "object archive". Object archives are merely packages containing one or more object (`.o`) files. They can't be executed, or loaded into a process address space, as can shared libraries. They can only be added to a linker command line. The linker is smart enough to realize that such archives are merely groups of object files. The linker extracts the object files it needs to complete the linkage process when building a program or shared library.

The second line represents the list of source files associated with this library. I chose to place both the header and the C source file in this list. I could have chosen to use a `noinst_HEADERS` line for the header file, but it was unnecessary because the `libjupcommon_a_SOURCES` list works just as well. The appropriate time to use `noinst_HEADERS` is when you have a directory that contains no source (`.c`) files—such as an internal include directory. Personally, I don't care for this style of project directory structure organization. I prefer to place private header files right along side of the source code they represent. As a result, I never seem to need `noinst_HEADERS` in my projects.

Notice the format of the `libjupcommon_a_SOURCES` variable. Automake transforms library and program names in the product list into derived variable names by converting all

characters except for letters, numbers and at-signs (@) into underscore characters. Thus, a library named `libc++.a` generates a `SOURCES` variable called `“libc___a_SOURCES”` (there are three consecutive underscores in that variable name).

Clean up your top-level project directory, removing all files and directories except those that we’ve written by hand so far. Also remove all `Makefile.in` files in the top-level directory and in sub-directories. The top-level directory should look like this when you’re done:

```
$ ls -lF
AUTHORS
ChangeLog
common/
configure.ac
COPYING
INSTALL
src/
Makefile.am
NEWS
README
```

Edit the `“SUBDIRS”` variable in the top-level `Makefile.am` file to include the new common directory that we just added:

#### **Makefile.am**

```
SUBDIRS = common src
```

Now we have to add some additional information to the `src/Makefile.am` file so that the generated Makefile can find the new library and header file we created in the common directory. Add two more lines to the end of the existing file, in this manner:

#### **src/Makefile.am**

```
bin_PROGRAMS = jupiter
jupiter_SOURCES = main.c
jupiter_CPPFLAGS = -I$(top_srcdir)/common
jupiter_LDADD = ../common/libjupcommon.a
```

Like the `“jupiter_SOURCES”` variable, these two new variables are obviously derived from the program name. The `“jupiter_CPPFLAGS”` variable is used to add product-specific C preprocessor flags to the compiler command line for all source files that are built for the jupiter program. The `“jupiter_LDADD”` variable is used to add libraries to the linker command line for the jupiter program.

These product-specific option variables are used to pass options to the compiler and linker command lines. The option variables currently supported by Automake for programs include:

- `program_CCASFLAGS`
- `program_CFLAGS`
- `program_CPPFLAGS`
- `program_CXXFLAGS`
- `program_FFLAGS`
- `program_GCJFLAGS`
- `program_LFLAGS`
- `program_OBJCFLAGS`
- `program_RFLAGS`

- `program_UPCFLAGS`
- `program_YFLAGS`

For static library products use `library_LIBADD`, instead of `program_LDADD`. The `_LIBADD` variable for libraries allows us to specify additional object files and static libraries that should be added to the static archive we're currently building. This can be handy for combining multiple convenience libraries. Consider the difference between these cases: The `library_LIBADD` variable is merely allowing us to specify already built objects—either libraries or actual object modules—to the library we're currently building. This can't be accomplished with the `library_SOURCES` variable, because `library_SOURCES` members are compiled, whereas `library_LIBADD` members are already built.

Additionally, the `program_LDADD` variable generally expects linker command line options such as `-lz` (to add the `libz` library to the linker's library specification for this program), while the `library_LIBADD` variable is formatted as a list of fully specified objects (eg., `libabc.a file1.o`). This rule isn't particularly strict however as we'll see shortly here. Quite frankly, it doesn't really matter, as long as the final command line composed by Automake from all of these variables makes sense to the linker.

### File-level option variables

Often you'll see unprefixed variables like `AM_CPPFLAGS` or `AM_LDFLAGS` used in a `Makefile.am`. This is the per-file form of these flags, rather than the per-product form. The per-file forms are used when the developer wants the same set of flags to be used for all products within a given `Makefile.am` file.

Sometimes you need to set a group of preprocessor flags for all products in a `Makefile.am` file, but add additional flags for one particular target. When you use a per-product flag variable, you need to include the per-file variable explicitly, like this:

```
AM_CFLAGS = ... some flags ...
program_CFLAGS = ... more flags ... $(AM_CFLAGS)
```

User variables, such as "CFLAGS" should never be modified by configuration scripts or makefiles. These are reserved for the end-user, and will be always be appended to the per-file or per-product versions of these variables.

Regarding the "jupiter\_LDADD" variable, `../common/libjupcommon.a` merely adds an object to the linker command line, so that code in this library may become part of the final program. Note that this sort of syntax is really only necessary for libraries built as part of your own package. If you're linking your program with a library that's installed on the user's system, then the configure script should have found it, and automatically added an appropriate reference to the linker's command line.

In the "jupiter\_CPPFLAGS" variable, the `-I$(top_srcdir)/common` directive tells the C preprocessor to add a search directory to its list of locations in which to look for header file references. Specifically, it indicates that header files referenced in C source files with angle brackets (< and >) should be searched for in this include search path. Header files referenced with double-quotes are

not searched for, but merely expected to exist in the specified directory, relative to the directory containing the referencing source file.

Getting back to our example—edit the `configure.ac` file; add a reference to the `AC_CONFIG_FILES` macro for the new generated `common/Makefile`, in this manner:

#### **configure.ac**

```
...
AC_CONFIG_FILES([Makefile
                 common/Makefile
                 src/Makefile])
...
```

Okay, now let's give our updated build system a try. I'm adding the `-i` option to the `autoreconf` command so that it will install any additional missing files that might be required after our enhancements:

```
$ autoreconf -i
configure.ac:6: installing `./missing'
configure.ac:6: installing `./install-sh'
common/Makefile.am:1: library used but `RANLIB'
      is undefined. The usual way to define
      `RANLIB' is to add `AC_PROG_RANLIB' to
      `configure.ac' and run `autoconf' again.
common/Makefile.am: installing `./depcomp'
src/Makefile.am:3: compiling `main.c' with
      per-target flags requires `AM_PROG_CC_C_O' in
      `configure.ac'
autoreconf: automake failed with exit status: 1
```

Well, it appears that we're still not done yet. Since we've added a new type of entity to our build system—static libraries—Automake (via `autoreconf`) tells us that we need to add a new macro to the `configure.ac` file. The `AC_PROG_RANLIB` macro is a standard program check macro, just like `AC_PROG_YACC` or `AC_PROG_LEX`. There's a lot of history behind the use of the `ranlib` utility on archive libraries. I won't get into whether it's still useful with respect to modern development tools. It seems however, that wherever you see it used in modern Makefiles, there's always a comment about running `ranlib` in order to “add karma” to the archive. You be the judge...

Additionally, we need to add the Automake macro, `AM_PROG_CC_C_O`, because this macro defines constructs in the resulting configure script that support the use of per-product flags, such as `jupiter_CPPFLAGS`. Let's add these two macros to our `configure.ac` script:

#### **configure.ac**

```
...
# Checks for programs.
AC_PROG_CC
AC_PROG_INSTALL
AC_PROG_RANLIB
AM_PROG_CC_C_O
...
```

Alright, once more then, but this time I'm adding the `--force` option, as well as the `-i` option to the `autoreconf` command line to keep it quiet about adding files that already

exist. (This seems like a pointless option to me, because the entire purpose of the “-i” option is to *add-missing* files, not to add *\_all* files that are required, regardless of whether they already exist, or not, and then complain if they do exist.):

```
$ autoreconf -i --force
configure.ac:15: installing `./compile'
```

Blessed day! It works. And it really wasn't too bad, was it? Automake told us exactly what we needed to do. (I always find it ironic when a piece of software tells you how to fix your input file—why didn't it just do what it knew you wanted it to do, if it understood your intent without the correct syntax?! Okay, I understand the “purist” point of view, but why not just do “the right thing”, with a side-line comment about your ill-formatted input text? Eventually, you'd be annoyed enough to fix the problem anyway, wouldn't you? Of course you would!)

### A word about the utility scripts

It seems that Automake has added yet another missing file—the “compile” script is a wrapper around some older compilers that do not understand the use of both “-c” and “-o” on the command line to name the object file differently than the source file. When you use product-specific flags, Automake has to generate code that may compile source files multiple times with different flags for each file. Thus it has to name the files differently for each set of flags it uses. The requirement for the compile script actually comes from the inclusion of the `AM_PROG_CC_C_O` macro.

At this point, we have the following Autotools-added files in the root of our project directory structure:

- `compile`
- `depcomp`
- `install-sh`
- `missing`

These are all scripts that are executed by the configure script, and by the generated Makefiles at various points during the end-user build process. Thus, the end-user will need these files. We can only get these files from Autotools. Since the user shouldn't be required to have Autotools installed on the final target host, we need to make these files available to the user somehow.

These scripts are automatically added (by “make dist”) to the distribution tarball. So, do we check them in to the repository, or not? The answer to this question is debatable, but generally I recommend against doing this. Anyone who will be creating a distribution tarball should also have the Autotools installed, and should be working from a repository work area. As a result, this maintainer will also be running “`autoreconf -i (--force)`” to ensure that she has the latest updated Autotools-provided utility scripts. Checking them in will only make it more probable that they become out of date as time goes by.

As mentioned in Chapter 2, this sentiment goes for the configure script as well. Some people argue that checking the utility and configure scripts into the project repository is beneficial, because it ensures that someone checking out a work area can build the project from the work area without having the Autotools installed. But is this really important? Shouldn't developers and maintainers be expected to have

more advanced tools? My personal philosophy is that they should. Yours may differ. Occasionally, an end user will need to build a project from a work area, but this should be the exceptional case, not the typical case. If it is the typical case, then there are bigger problems with the project than can be solved in this discussion.

### What goes in a distribution?

In general, Automake determines automatically what should go into a distribution created with “make dist”. This is because Automake is very aware of every single file in the build process, and what it’s used for. Thus, it need not be told explicitly which files should be in the package, and which should be left behind.

An important concept to remember is that Automake wants to know statically about every source file used to build a product, and about every file that’s installed. This means, of course, that all of these files must somehow be specified at some point in a Makefile.am primary variable. This bothers some developers—and with good reason. There are cases where dozens of installable files are generated by tools using long, apparently random and generally unimportant naming conventions. Listing such generated files statically in a primary variable is problematic, to say the least.

We’ll cover techniques that can be used to work around such problem cases later in this book. At this point, however, I’d like to introduce the `EXTRA_DIST` variable for those cases where file system entities are not part of the Automake build process, but should be distributed with a distribution tarball. The `EXTRA_DIST` variable contains a space-delimited list of files and directories which should be added to the distribution package when “make dist” is executed.

```
EXTRA_DIST = windows
```

This might be used to add, for example, a `windows` build directory to the distribution package. Such a directory would be otherwise ignored by Automake, and then your windows users would be upset when they unpacked your latest tarball. Note in this example that “`windows`” is a directory, not a file. Automake will automatically and recursively add every file in this directory to the distribution package.

### Summary

In this chapter, we’ve covered a fair number of details about how to instrument a project for Automake. The project we chose to instrument happened to already be instrumented for Autoconf, which is the most likely scenario, as you’ll probably be adding Autoconf functionality to your bare projects first in most cases.

What I’ve explicitly *not* covered are situations where you need to extend Automake to handle your special cases, although I’ve hinted at this sort of thing from time to time.

In the next chapter, we’ll examine adding Libtool to the jupiter project, and then in Chapter 6, we’ll Autotool-ize a real-world project, consisting of several hundred source files and a custom build system that takes the form of a GNU makefile designed to use native compilers on multiple platforms including Solaris, AIX, Linux, Mac OS and Windows, among

others. I'll warn you up front that we'll be remaining true to the original mission statement of this book in that we'll not be trying to get Autotools to build Windows products.

[◀ Chapter 3: Configuring your project with autoconf](#)

u  
p

[Chapter 5: Building shared libraries once using autotools ▶](#)

[Login or register to post comments](#)

1162 reads

---



- COLUMNS
- COMMUNITY POSTS
- ISSUES
- BOOKS
- FORUM
- FS NEWS

Ads by Google

**Configuration Software**  
 Product & Service Configurator. Configure manufacturing & more.  
[www.Configuration.com/manufacturir](http://www.Configuration.com/manufacturir)

**Pump Software**  
 Streamline your pump sales. Select, configure, quote, and order online.  
[www.bigmachines.com](http://www.bigmachines.com)

**Free ALM Starter**  
 Integrated bug tracking, VCS, code to rev links, easy future mobility  
[www.polarion.com](http://www.polarion.com)

**Network Config tool**  
 Free Configuration Management tool for Routers Switches Firewalls etc  
[deviceexpert.com/network-config](http://deviceexpert.com/network-config)

**Download Remote Access**  
 Network Management Software & IT Tools for Free. Download Now!  
[Spiceworks.com](http://Spiceworks.com)

- [Tracking Worms, Spam and Malware ...](#)
- [Google AdSense for RSS Feeds Coming ...](#)
- [Hotmail, The E-Mail You Never ...](#)

Sponsored by **Novell**

### Best voted contents

- [Open letter to standards professionals, developers, and activists](#)  
Pieter Hintjens, 2008-05-13
- [The 2008 Google Summer of Code: 21 Projects I'm Excited About](#)  
Andrew Min, 2008-05-13
- [The Bizarre Cathedral -](#)

Home » Autotools: a practitioner's guide to Autoconf, Automake and Libtool

## Chapter 5: Building shared libraries with Libtool



by [John Calcote](#)

The person who invented the concept of shared libraries should be given a raise...and a bonus. The person who decided that shared library management and naming conventions should be left to the implementation should be flogged. Sorry—that opinion is the result of too much negative experience on my part with building shared libraries for multiple platforms *without* the aid of Libtool. The very existence of Libtool stands as a witness to the truth of this sentiment. Libtool exists for one purpose only—to provide a standardized, abstract interface for developers desiring to create portable shared libraries. It abstracts both the shared library build process, and the programming interfaces used to dynamically load and access shared libraries at run time.

Before we get into a discussion of the proper use of Libtool, we should probably spend a few minutes on the features and functionality provided by shared libraries, so that you will understand the scope of the material I'm covering here.

### The benefits of shared libraries

Shared libraries provide a way to ship reusable chunks of functionality in a convenient package that can be loaded into a process address space, either automatically at program load time by the operating system loader, or by code in the application itself, when it decides to load and access the library's functionality. The point at which an application binds functionality from a shared library is very flexible, and determined by the developer, based on the design of the program and the needs of the end-user.

The interfaces between the program executable and modules defined as shared libraries must be well-designed by virtue of the fact that shared library interfaces must be well-specified. This rigorous specification promotes good design practices. When you use shared libraries, you're essentially forced to be a better programmer.

Shared libraries may be (as the name implies) shared among processes. This sharing is very literal. The code segments for a shared library can be loaded once into physical memory pages. That same memory pages can then be mapped into the process address spaces for multiple programs. The data pages must, of course, be unique per process, but global data



<http://www.freesoftwaremagazine.com> © Copyright 2008 Tony Mobley & Ryan Cartwright CC: BY-NC-SA

### From the FSM staff...

[The Top 10 Everything](#) (Dave). The good, the bad and the ugly.

[Free Software news](#) (Dave & Bridget). All about free software -- free as in freedom!

[Book Reviews: Illiterary](#) (Bridget). Book reviews, blogs, and short stories.

[Hot topics - last 60 days](#)

[Installing an all-in-one printer device in Debian](#)  
Ryan Cartwright, 2008-05-05

[Things you miss with GNU/Linux](#)  
Ryan Cartwright, 2008-05-01

Free Open Source Book Samples...  
 Exclusive discount offer to FSM Readers  
**PACKT**

Download it.  
 It's open source.  
 It's yours.

**VYATTA**  
 FREE Open Source ROUTER FIREWALL VPN

- o Ryan Cartwright, 2008-05-25
- o **Book Review: Practical Guide to Ubuntu Linux by <Mark G. Sobell</>**  
Alan Berg, 2008-05-10

segments are often small compared to the code segments of a shared library. This is true efficiency.

Shared libraries are easily updated during program upgrades. The base program may not have changed at all between two revisions of a software package. A new version of a shared library may be laid down on top of the old version, as long as its interfaces have not been changed. When interfaces are changed, two shared libraries may co-exist side-by-side, because the versioning scheme used by shared libraries (and supported by Libtool) allows the shared library files to be named differently, but treated as the same library. Older programs may continue to use older versions of the library, while newer programs may use the newer versions.

If a software package specifies a well-defined "plug-in" interface, then shared libraries can be used to implement user-configurable loadable functionality. This means that additional functionality can become available to a program *after* it's been released, and third-parties can even add functionality to your program, if you publish a document describing your plug-in interface specification.

### How shared libraries work

As I mentioned above, the way a POSIX-based operating system implements shared libraries varies from platform to platform, but the general idea is the same for all platforms. The following discussion applies to shared library references that are resolved by the linker while the program is being built, and by the operating system loader at program load time.

### Dynamic linking at load time

As a program executable image is being built, the linker maintains a table of unresolved function entry points and global data references. If the linker can't find a function that is called by, or a global data item that is referenced by code within the program, it adds the missing symbol name to this table of undefined references. The linker will exit with an error message if there are any outstanding undefined references in this table after all of the object code has been analyzed. The linker then combines all of the object code containing all located references into an executable program binary image. This image may then be loaded and executed by a user. It is entirely self-contained and depends only upon itself.

Assuming that all undefined references *are* resolved during the linking process, if the list of objects to be linked contains one or more shared libraries, the linker will build the executable image from all *non-shared* objects specified on the linker command line. This includes all individual .o files and all static library archives. However it will add two tables to the binary image header; the first is the table of outstanding external references—those found in shared libraries, and the second is a table of shared library names and versions in which the outstanding undefined references were found during the linking process.



### Buzz authors

Marco Marongiu

Pieter Hintjens

drascus321

dogboi

Richard Rothwell

### All news

ம ஸ்ரீ ராமதாஸ்: Party time with Ubuntu...

Sridhar Dhanapalan: Open CeBIT

Martin Meredith: Welcome to the Family, Synergy

Stephen Stalcup: Next Membership meeting for the Americas

Andres Rodriguez: Hello Planet Ubuntu!!

more



### Why Microsoft should not lose (and free software will still win)

Ryan Cartwright, 2008-04-21

### Drigg (the pligg alternative) vs. Pligg: why should people switch?

Tony Mobily, 2008-04-13

### Beyond Synaptic - using apt for better package management

Ryan Cartwright, 2008-04-03

### Hot topics - last 21 days

### Dubious ads in Free Software Magazine

Tony Mobily, 2008-05-25

### The Bizarre Cathedral - 6

Ryan Cartwright, 2008-05-25

### Open letter to standards professionals, developers, and activists

Pieter Hintjens, 2008-05-13

### The Bizarre Cathedral - 4

Ryan Cartwright, 2008-05-11



### Odiogo



14 20 Essential KDE Applications - Review

14 10 Reasons to Love Debian

16 Burn Your CDs and DVDs! K3b Review

17 Novell's Moonlight: Crippled and Defective by Design™

18 India Appeals Against OOXML, Joining Brazil, South Africa, Maybe More

17 Wiping your disk drive clean

20 Dear Google: Is AGPL Evil?

19 Slashdot: World's Third-Largest Population Appeals ISO's Decision

Get this widget »

When the operating system loader attempts to load this executable image, it must resolve the remaining outstanding references to symbols imported from the shared libraries named in the executable header. If the loader can't resolve all of the references, then a load error occurs, and the entire process is terminated with an operating system error.

Note here that these external symbols are not tied to a specific shared library. The operating system will stop loading shared libraries as soon as it is able to resolve all of the outstanding symbol references. Usually, this happens after the last indicated shared library is loaded into the process address space, but there are exceptions.

*NOTE: This process differs a bit from the way a Windows operating system resolves symbols in Dynamic Link Libraries (DLLs). On Windows, a particular symbol is tied by the linker at program build time to a specifically named DLL.*

This system has both pros and cons. For example, on some systems, such unbound symbols can be satisfied by a library specified by the user. That is, a user can entirely *replace* a library at runtime by simply preloading one that contains the same symbols. Unfortunately, free-floating symbols can also lead to problems. For instance, two libraries can provide the same symbol name, and the dynamic loader can inadvertently bind an executable to the symbol from the wrong library. At best, this will cause a program crash when the wrong arguments are passed to the mis-matched function. At worst, this can lead to security risks, because the mis-matched function might be used to capture passwords and security credentials passed by the unsuspecting program.

C-language symbols do not include parameter information, so it's rather likely that symbols will clash in this manner. C++ symbols are a bit safer, in that the entire function signature (minus the return type) is encoded into the symbol name. However, even C++ is not immune to hackers replacing a security function with their own version of that function.

### **Automatic dynamic linking at run time**

The operating system loader can also use a very late form of binding, often referred to as "lazy binding". In this situation, the entries in the jump table in the program header for external references are initialized to code in the dynamic loader itself.

At the point the program first calls such a lazy entry, the call will be routed to the loader, which will then determine the actual address of the function, reset the entry in the jump table, and redirect to the proper function entry point. The next time this happens, the jump table entry will be correctly initialized, and the program will jump directly to the called function.

This lazy binding mechanism makes for very fast program startup, because shared libraries whose symbols are not bound until they're needed aren't even loaded until they're first referenced by the application program. Now, consider this— they may *never* be referenced. Which means they may never

be loaded, saving both time and space.

The problems with this method should be obvious, at this point. While using automatic run-time dynamic linking can give you faster load times, and better performance and space efficiency, it can also cause abrupt terminations of your application without warning. If the loader can't find the requested symbol—perhaps the required library is missing—then it has no recourse except to abort the process. Why not ensure that all symbols exist when the program is loaded? If the loader resolved all symbols at load time, then it might as well populate the jump table entries at that point, as well. After all, it had to load all the libraries any way to ensure that the symbols actually exist. This then defeats the purpose of this binding method.

The moral of this story is that you get what you pay for. If you don't want to pay the insurance premium for longer up-front load times, and more space consumed (even if you may never really need it), then you may have to take the hit of a missing symbol at run time—a program crash, essentially.

### Manual dynamic linking at run time

Manual run-time dynamic linking is done a bit differently. In this case, the linker finds no external symbols in shared libraries, because the program doesn't call any shared library functions directly. Rather, shared library functions are referenced through a function pointer that is populated by the application program itself at run time.

The way this works is that a program calls an operating system function to manually load a shared library into its own process address space. This system function returns a "handle", or an opaque value representing the loaded library, to the calling program. The program then calls another function to import a symbol from the library referred to by that handle. If all goes well, the operating system returns the address of the requested function or data item in the desired library. The program may then call the function, or access the global data item through this pointer.

If something goes wrong in one of these two steps—say the library could not be found, or the symbol was not found within the library, then it becomes the responsibility of the program to define the results—perhaps display an error message, indicating that the program was not configured correctly.

This is a little nicer than the way automatic dynamic run-time linking works because, while the loader has no option but to abort, the application, having a higher-level viewpoint, can handle the problem much more gracefully. The drawback, of course, is that you as the programmer have to manage the process of loading libraries and importing symbols within your application code.

### Using Libtool

An entire book could be written on the details of shared libraries and their implementations on various systems. This

short primer will suffice for our needs, so let's move on to how Libtool can be used to make a package maintainer's life easier.

The Libtool project was started in 1996 by Gordon Matzigkeit. Libtool was designed to extend Automake, but can be used independently as well in hand-coded makefiles. The Libtool projects is currently maintained by Bob Friesenhahn, Peter O'Gorman, Gary V. Vaughan and Ralf Wildenhues.

### Abstracting the build process

First, let's look at how Libtool helps during the build process. Libtool provides a script (`ltmain.sh`) that `config.status` executes in a Libtool-enabled project. The `ltmain.sh` script builds a custom version of the `libtool` script, specifically for your package. This `libtool` script is then used by your project's makefiles to build shared libraries specified with the `LTLIBRARIES` primary. The `libtool` script is really just a fancy wrapper around the compiler, linker and other tools. The `ltmain.sh` script should be shipped in a distribution tarball, as part of your end-user build system. Automake-generated rules ensure that this happens as it should.

The `libtool` script insulates the build system author from the nuances of building shared libraries on multiple platforms. This script accepts a well-defined set of options, converting them to appropriate platform- and linker-specific options on the host platform and tool set. Thus, the maintainer need not worry about the specifics of building shared libraries on each platform. She need only understand the available `libtool` script options. These are well specified in the GNU Libtool manual.

On systems that don't support shared libraries at all, the `libtool` script uses appropriate commands and options to build and link static libraries. This is all done in such a way that the maintainer is isolated from the difference between building shared libraries and static libraries.

You can emulate building your package on a static-only system by using the `--disable-shared` option on the `configure` command line for your project. This will cause Libtool to assume that shared libraries cannot be built on the target system.

### Abstraction at run-time

Libtool can also be used to abstract the programming interfaces supplied by the operating system for loading libraries and importing symbols. Programmers who've ever dynamically loaded a library on a Linux system are familiar with the standard Linux shared library API, including the functions, `dlopen`, `dlsym` and `dlclose`. These functions are provided by a shared library usually named `"d1"`. Unfortunately, not all POSIX systems that support shared libraries provide the `d1` library, or functions using these names.

To address these differences, Libtool provides a shared library called “`ltdl`”, which provides a clean, portable library management interface, very similar to the `dlopen` interface provided by the Linux loader. The use of this library is optional, of course, but highly recommended, because it provides more than just a common API across shared library platforms. It also provides an abstraction for manual run-time dynamic linking between shared library and non-shared library platforms.

What!? How can that work? On systems that don’t provide shared libraries, Libtool actually creates internal symbol tables within the executable containing all of the code that would otherwise be found in shared libraries, on systems that support shared libraries. By using these symbol tables on these platforms, the `ltdlopen` and `ltdlsym` functions can make your code appear to be loading and importing symbols, when in fact, the “load” function does nothing, and the “import” function hands you back the address of some code that’s been linked right into your program.

The `ltdl` library is, of course, not really necessary for packages that don’t use manual run-time dynamic linking. But if your package does—perhaps by providing a plug-in interface of some sort, then you’d be well-advised to use the API provided by `ltdl` to managed loading and linking to your plug-in modules—even if you only target systems that provide good shared library services. Otherwise, your source code will have to consider the differences in shared library management between your many target platforms. At the very least, some of your users will have to put on their “developer” hats, and attempt to modify your code so that it works on their odd-ball platforms.

## A word about the latest Libtool

The most current version of Libtool is version 2.2. However, many popular Linux distributions are still shipping packages containing Libtool version 1.5, so many developers don’t know about the changes between these two versions. The reason for this is that certain backward-compatibility issues were introduced after version 1.5 that make it difficult for Linux distros to support the latest version of Libtool. This probably won’t happen until all (or almost all) of the packages they provide have updated their `configure.ac` files to properly use the latest version of Libtool.

This is somewhat of a “chicken-and-egg” scenario—if distros don’t ship it, how will developers ever start using it on their own packages? So it’s not likely to happen any time soon. If you want to make use of the latest Libtool version while developing your packages (and I highly recommend that you do so), you’ll probably have to download, build and install it manually, or look for an updated Libtool package from your distribution provider.

Downloading, building and installing Libtool manually is trival:

```
$ wget ftp.gnu.org/gnu/libtool/libtool-2.2.tar.gz
```

```

...
$ tar xzf libtool-2.2.tar.gz
$ cd libtool-2.2
$ ./configure && make
...
$ sudo make install
...

```

Be aware that the default installation location (as with most of the GNU packages) is `/usr/local`. If you wish to install it into the `/usr` hierarchy, then you'll need to use the `--prefix=/usr` option on the `configure` command line.

You might also wish to use the `--enable-ltdl-install` option on the `configure` command line to install the `ltdl` library and header files into your `lib` and `include` directories.

## Adding shared libraries to Jupiter

Okay, now that we have that background behind us, let's take a look at how we can add a Libtool shared library to the Jupiter project. First, let's consider what we might do with a shared library. As mentioned above, we might wish to provide our users with some library functionality that their own applications could use. We might also have several applications in our package that need to share the same functionality. A shared library is a great tool for both of these scenarios, because we get the benefits of code reuse and memory savings, as shared code is amortized across multiple applications—both internal and external to our project.

We'll add a shared library to Jupiter that provides the print functionality we use in the `jupiter` application. We'll do this by having the new shared library call into the `libjupcommon.a` static library. Remember that calling a routine in a static library has the same effect as linking the object code for the called routine right into the calling application (or shared library, as the case may be). The called routine ultimately becomes an integral part of the calling binary image (program or shared library).

Additionally, we'll provide a public header file from the Jupiter project that will allow external applications to call this same functionality. By doing this, we can allow other applications to “display stuff” in the same way that the `jupiter` program “displays stuff”. (This would be significantly cooler if we were actually doing anything useful in `jupiter`!).

## Using the `LTLIBRARIES` primary

Automake has built-in support for Libtool. The `LTLIBRARIES` primary is provided by code in the Automake package, not the Libtool package. This really doesn't qualify as a pure extension, but rather more of an add-on package for Automake, where Automake provides the necessary infrastructure for that specific add-on package. You can't access the `LTLIBRARIES` primary functionality provided by Automake without Libtool, because the use of this primary

obviously generates make rules that call the `libtool` build script.

I state all of this here because it bothers me that you can't *really* extend the list of primaries supported by Automake without modifying the actual Automake source code. The fact that Automake is written in perl is somewhat of a boon, because it means that it's *possible* to do it. But you've *really* got to understand Automake source code in order to do it properly. I envision a future version of Automake whereby code may be added to an Automake extension file that will allow the dynamic definition of new primaries.

It's a bit like the old FOSS adage, generally offered to someone complaining about a particular package: "It's open source. Change it yourself!" This is very often easier said than done. Furthermore, what these people are really telling you is to change *your copy* of the source code for your own purposes, not to change the *master copy* of the source code. Getting your changes accepted into the master source base often depends more on the quality of your relationship with the current project maintainers than it does on the quality of your coding skills. I'm not complaining, mind you. I'm merely stating a fact that should not be overlooked when one is considering making changes to an existing free open source software project.

So why not ship Libtool as part of Automake, rather than as a separate package? Because Libtool can quite effectively be used independently of Automake. If you wish to try Libtool by itself, then please refer to the GNU Libtool manual for more information. The opening chapters in that manual describe the use of the `libtool` script as a stand-alone product.

### Public include directories

Earlier in this book, I made the statement that a directory in a project named "include" should only contain public header files—those that expose a public interface in your project. Well, now we're going to add just such a header file to the Jupiter project, so, we'll create an include directory. I'll add this directory at the top-level of the project directory structure.

If we had multiple shared libraries, we'd have a choice to make: Do we create separate include directories for each library in the library source directory, or do we add a single top-level include directory? I usually use the following rule of thumb to determine the answer to this question: If the libraries are designed to work together as a group, and if consuming applications generally use the libraries as a group, then I use a single top-level include directory. If, on the other hand, the libraries can be used independently, and if they offer fairly autonomous sets of functionality, then I provide individual include directories in my project's library subdirectories.

In the end, it really doesn't matter much because the header files for these libraries will be installed in entirely different directory structures than those in which they exist within your project. In fact, make sure you don't inadvertently use the

same file name for headers in two different libraries in your project, or you'll probably have problems installing these files. They generally end up all together in the `"$(prefix)/include"` directory.

I'll also add a directory for the new Jupiter shared library, called `"libjup"`. These changes require adding references to these new directories to the top-level `Makefile.am` file's `SUBDIRS` variable, and then adding corresponding makefile references to the `AC_CONFIG_FILES` macro in the `configure.ac` file. Let's do that now:

```
$ mkdir include
$ mkdir libjup
$ echo "SUBDIRS = common include libjup src" \
  > Makefile.am
$ echo "include_HEADERS = libjupiter.h" \
  > include/Makefile.am
$ vi configure.ac
...
AC_CONFIG_FILES([Makefile
                  common/Makefile
                  include/Makefile
                  libjup/Makefile
                  src/Makefile])
...
```

The `include` directory's `Makefile.am` file is trivial, containing only a single line wherein the public header file, `"libjupiter.h"` is referred to in an Automake `HEADERS` primary. Note that we're using the `include` prefix on this primary. You'll recall that the `include` prefix indicates that files specified in this primary are destined to be installed in the `"$(includedir)"` directory (eg., `/usr/local/include`). The `HEADERS` primary is much like the `DATA` primary, in that it specifies a set of files that are to be treated simply as data to be installed without modification or pre-processing.

The `libjup/Makefile.am` file is a bit more complex, containing four lines, as opposed to the usual one or two lines:

```
lib_LTLIBRARIES = libjupiter.la
libjupiter_la_SOURCES = jup_print.c
libjupiter_la_LDFLAGS = ../common/libjupcommon.a
libjupiter_la_CPPFLAGS = -I$(top_srcdir)/include \
  -I$(top_srcdir)/common
```

Let's analyze this file, line by line. The first line is the primary line, and contains the usual prefix for libraries. The `"lib"` prefix indicates that the referenced products are to be installed in the `"$(libdir)"` directory. We might also have used the `pkglib` prefix to indicate that we wanted our libraries installed into the `"$(prefix)/lib/jupiter"` directory. Here, we're using the `LTLIBRARIES` primary, rather than the older `LIBRARIES` primary. The use of this primary tells Automake to generate rules that use the `libtool` script, rather than calling the compiler, linker and `ar` utilities to generate the products.

The second line lists the sources that are to be used for the first (and only) product. The third line indicates a set of linker options for this product. In this case, we're specifying that the `libjupcommon.a` static library should be linked into (become part of) the `libjupiter.so` shared library.

There's an important concept regarding the `*_LDFLAGS` variable that you should strive to understand completely: Libraries that are consumed within, and yet built as part of the same project, should be referenced internally, using relative paths within the *build* directory hierarchy. Libraries that are external to a project generally need not be referenced explicitly at all, as the `$(LIBS)` variable should already contain the appropriate `"-L"` and `"-l"` options for those libraries. These options come from an attempt to find these libraries in the `configure` script, using the appropriate `AC_CHECK_LIBS`, or `AC_SEARCH_LIBS` macros.

The fourth line indicates a set of C preprocessor flags that are to be used on the compiler command line for locating our shared library header files. These options indicate, of course, that the top-level `include` and `common` directories should be searched by the pre-processor for header file references in the source code.

Here's the new source file, `jup_print.c`:

#### **libjup/jup\_print.c**

```
#include <libjupiter.h>
#include <jupcommon.h>

int jupiter_print(char * name)
{
    print_routine(name);
}
```

We need to include the new shared library header file for access to the `jupiter_print` function's public prototype. This leads us to another general software engineering principle. I've heard it called many names, but the one I tend to use the most is "The DRY Principle", which is an acronym that stands for Don't Repeat Yourself. C function prototypes are very useful, because when used correctly, they enforce the fact that the public's view of a function is the same as the package maintainer's view. So often, I've seen source code for a function where the source file doesn't include the public prototype for the function. It's easy to make a small change in the function or prototype, and then not duplicate it in the other location—unless you've included the public header file in the source file containing the function.

We need the static library header file, because we call its function from within our public library function. Note also that I placed the public header file first—there's a good reason for this—another general principle: By placing the public header file first in the source file, I can allow the compiler to check that the use of this header file doesn't depend on any other files in the project.

If the public header file has a hidden dependency on some construct (a typedef, structure or pre-processor definition) defined in `jupcommon.h`, and if I include it after `jupcommon.h`, then the dependency would be hidden by the fact that the required construct is already available in the translation unit when the compiler begins to process the public header file, by virtue of the order in which they were included.

Next, I'll modify the `jupiter` application's `main` function so that it calls into the shared library instead of calling into the common static library:

#### **src/main.c**

```
#include <libjupiter.h>

int main(int argc, char * argv[])
{
    jupiter_print(argv[0]);
    return 0;
}
```

Here, I've changed the print function from `print_routine`, found in the static library, to `jupiter_print`, as provided by our new shared library. I've also changed the header file included at the top from `libjupcommon.h` to `libjupiter.h`.

My choices of names for the public function and header file were arbitrary, but based on a desire to provide a clean, rational and informational public interface. The name `libjupiter.h` very clearly indicates that this header file provides the public interface for the `libjupiter.so` shared library. I try to name library interface functions in such a way that they are clearly part of an interface. How you choose to name your public interface members—files, functions, structures, typedefs, pre-processor definitions, global data, etc—is up to you, but you should consider using a similar philosophy. Remember, the goal is to provide a great end-user experience.

Finally, the `src/Makefile.am` file must also be modified to use our new shared library, rather than the `libjupcommon.a` static library.

#### **src/Makefile.am**

```
bin_PROGRAMS = jupiter
jupiter_SOURCES = main.c
jupiter_CPPFLAGS = -I$(top_srcdir)/include
jupiter_LDADD = ../libjup/libjupiter.la
...
```

In this file, I've changed the `jupiter_CPPFLAGS` variable so that it now refers to the new `include` directory, rather than the `common` directory. I've also changed the `jupiter_LDADD` variable so that it refers to the new Libtool shared library object, rather than the `libjupcommon.a` static library. All else remains the same. Note that these

changes are both obvious and simple. The syntax for referring to a Libtool library is identical to that referring to an older static library. Only the library extension is different. The Libtool library extension, “.la” refers to “Libtool Archive”.

Taking a step back for a moment: Do we actually need to make this change? No, of course not. The `jupiter` application will continue to work just fine the way it was originally set up—linking the code for the static library’s `print_routine` directly into the application works equally well to calling the new shared library routine (which ultimately contains the same code). In a real project, we might actually leave it the way it was. Why? Because both public entry points, `main` and `jupiter_print` call the exactly same function (`print_routine`) in `libjupcommon.a`, so the functionality is identical. Why add the overhead of a call through the public interface? Here, the reason is purely educational.

In this situation, you might now consider simply moving the code from the static library into the shared library, thereby removing the need for the static library entirely. Again, I’m going to beg your indulgence with my contrived example. In a more complex project, we could very well have a need for this sort of configuration, so I’m going to leave it the way it is for the sake of its educational value to us.

## Reconfigure and build

Well, let’s give it a try and see where we stand at this point. Since we added a major new component to our project build system (Libtool), I’ll add the “-i” option to the `autoreconf` command, just in case new files need to be installed:

```
$ autoreconf -i
$ ./configure
...
checking for ld used by gcc...
checking if the linker ... is GNU ld... yes
checking for BSD- or MS-compatible name lister...
checking the name lister ... interface...
checking whether ln -s works... yes
checking the maximum length of command line...
checking whether the shell understands some XSI...
checking whether the shell understands "+="...
checking for ...ld option to reload object files...
checking how to recognize dependent libraries...
checking for ar... ar
checking for strip... strip
checking for ranlib... ranlib
checking command to parse ...nm -B output...
...
checking for dlfcn.h... yes
checking for objdir... .libs
checking if gcc supports -fno-rtti...
checking for gcc option to produce PIC... -fPIC
checking if gcc PIC flag -fPIC -DPIC works...
checking if gcc static flag -static works...
checking if gcc supports -c -o file.o... yes
```

```

checking if gcc supports -c -o file.o... yes
checking whether ... linker ... supports shared...
checking whether -lc should be explicitly linked...
checking dynamic linker characteristics...
checking how to hardcode library paths...
checking whether stripping libraries is possible...
checking if libtool supports shared libraries...
checking whether to build shared libraries...
checking whether to build static libraries...
...
$

```

The first noteworthy item here is that Libtool adds *significant* overhead to the configuration process. I've only shown the output lines here that are new since we added Libtool. All I've added to the `configure.ac` script is the reference to the `LT_INIT` macro. I've nearly doubled my `configure` script output. This should give you some idea of the number of system characteristics that must be examined to create portable shared libraries. Libtool does a lot of the work for us. Let's continue.

*NOTE: In the following output examples, I've wrapped long output lines to fit publication formatting, and I've added blank lines between output lines for readability. I've also removed some unnecessary text, such as long directory names—again to increase readability.*

```

$ make
...
Making all in libjup
make[2]: Entering directory `../libjup'

/bin/sh ../libtool --tag=CC --mode=compile gcc
-DHAVE_CONFIG_H -I. -I../libjup -I..
-I../include -I../common -g -O2
-MT libjupiter_la-jup_print.lo -MD -MP -MF
.deps/libjupiter_la-jup_print.Tpo -c
-o libjupiter_la-jup_print.lo
`test -f 'jup_print.c'
  || echo '../libjup/'`jup_print.c

libtool: compile: gcc -DHAVE_CONFIG_H -I.
-I../libjup -I.. -I../include
-I../common -g -O2 -MT
libjupiter_la-jup_print.lo -MD -MP -MF
.deps/libjupiter_la-jup_print.Tpo -c
../libjup/jup_print.c -fPIC -DPIC
-o .libs/libjupiter_la-jup_print.o

libtool: compile: gcc -DHAVE_CONFIG_H -I.
-I../libjup -I.. -I../include
-I../common -g -O2 -MT
libjupiter_la-jup_print.lo -MD -MP -MF
.deps/libjupiter_la-jup_print.Tpo -c
../libjup/jup_print.c
-o libjupiter_la-jup_print.o >/dev/null 2>&1

mv -f .deps/libjupiter_la-jup_print.Tpo

```

```

    .deps/libjupiter_la-jup_print.Plo

/bin/sh ../libtool --tag=CC --mode=link gcc -g
-O2 ../common/libjupcommon.a -o libjupiter.la
-rpath /usr/local/lib libjupiter_la-jup_print.lo
-lpthread

*** Warning: Linking ... libjupiter.la against the
*** static library libjupcommon.a is not portable!

libtool: link: gcc -shared
    .libs/libjupiter_la-jup_print.o
    ../common/libjupcommon.a -lpthread
-Wl,-soname -Wl,libjupiter.so.0
-o .libs/libjupiter.so.0.0.0

../ld: ../common/libjupcommon.a(print.o):
relocation R_X86_64_32 against `a local symbol'
can not be used when making a shared object;
recompile with -fPIC

../common/libjupcommon.a: could not read symbols:
Bad value

collect2: ld returned 1 exit status
make[2]: *** [libjupiter.la] Error 1
...

```

Well, that wasn't a very pleasant experience! It appears that we have some errors to fix. Let's take them one at a time, from top to bottom.

The first point of interest is that the `libtool` script is being called with a `--mode=compile` option, which causes `libtool` to act as a wrapper script around a somewhat modified version of our standard `gcc` command line. You can see the effects of this statement in the next two compiler command lines. *Two lines?* That's right. It appears that `libtool` is causing the compile operation to occur twice.

A careful examination of the differences between these two command lines shows us that the first compiler command is using two additional flags: `-fPIC` and `-DPIC`. The first line also appears to be directing the output file to a `.libs` subdirectory, whereas, the second line is saving it in the current directory. Finally, both the `STDOUT` and `STDERR` output is redirected to `/dev/null` in the second line.

This double-compile "feature" has caused a fair amount of anxiety on the Libtool mailing list over the years. Mostly, this is due to a lack of understanding of what it is that Libtool is trying to do, and why it's necessary. Using various `configure` script command line options provided by Libtool, you can force a single compilation, but doing so brings with it a certain loss of functionality, which I'll explain here shortly.

The next line renames the dependency file from `*.Tpo` to `*.Plo`. Dependency files contain `make` rules that declare dependencies between source files and header files. These

are generated by the C preprocessor when the “-MT” compiler option is used. They are then included in makefiles so that the `make` utility can properly recompile a source file if one or more of its include dependencies have been modified since the last build. This is not really germane to our Libtool discussion, so I'll not go into any more detail here, but check the GNU Make manual for more information. The point is that one Libtool command may (and often does) execute a string of shell commands.

The next line is another call to the `libtool` script, this time using the “-mode=link” option. This option generates a call to execute the compiler in “link” mode, passing all of the libraries and linker options specified in the `Makefile.am` file.

And finally, we come to our first problem—a portability warning about linking a shared library against a static library. Specifically, this warning is about linking against a *non-Libtool static library*. Shortly, we'll begin to see why this might be a problem. Notice that this is not an error. Were it not for additional errors we'll encounter later, the library would be built for us in spite of this warning.

After the portability warning, `libtool` attempts to link the requested objects together into a shared library named “`libjupiter.so.0.0.0`”. But here we run into the real problem—a linker error indicating that somewhere from within `libjupcommon.a`—and more specifically within `print.o`—an Intel object relocation cannot be performed because the original source file (`print.c`) was apparently not compiled correctly. Libtool (actually, the linker) is kind enough to tell us what we need to do to fix the problem. It indicates that we need to compile the source code using a “-fPIC” compiler option.

Now, if you don't know anything about the “-fPIC” option, then you'd be wise at this point to open the man page for `gcc` and study it, before willy-nilly inserting compiler or linker options until the warning or error disappears, as many new programmers are wont to do. Software engineers should understand the meaning and nuances of *every* command line option used by the tools in their projects' build systems. Why? Because, otherwise, they don't really know what they have when their build completes. It may work the way it should—but if it does, it's simply by luck, rather than by design. Good engineers know their tools, and the best way to learn is to study error messages and their fixes until the problem is well-understood, before moving on.

### So what is “PIC” code?

When operating systems create new process address spaces, they always load the executable images at the same memory address. This magic address is system-specific. Compilers and linkers know this, and they know what that address is on a given system. Therefore, when they generate internal references to, say function calls, for example, they can generate those references as *absolute* addresses. If you were

able somehow to load the executable at a different location in memory it would simply not work properly, because the absolute addresses within the code would be incorrect. At the very least, the program will crash when it jumps to the wrong location during a function call.

Consider Figure 1 below for a moment. Given a system whose magic executable load address is `0x10000000`, this diagram depicts two process address spaces on such a system. In the process on the left, an executable image is loaded correctly at address `0x10000000`. At some point in the code a “`jmp`” instruction tells the processor to transfer control to the absolute address `0x10001000`, where it continues executing instructions in another area of the program. In the process on the right, the program is (somehow) loaded incorrectly at address `0x20000000`. When that same absolute branch instruction is encountered, the processor jumps to address `0x10001000`, because the address is hard-coded into the program. This, of course, fails—often spectacularly.

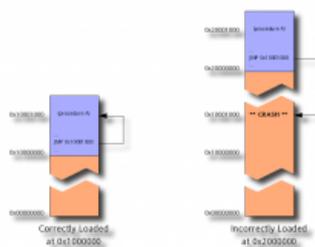


Figure 1: Absolute addressing in executable images

When a shared library is built for certain types of hardware (Intel x86 and x86\_64 included), the address at which the library will be loaded within a process address space cannot be known by the compiler or the linker beforehand. This is because many libraries may be loaded into any given process, and the order in which they are loaded depends on how the *executable* is built, not the library. Furthermore, who’s to say which library owns location “A”, and which one owns location “B”? The fact is, libraries may be loaded *anywhere* into a process where there is space for it at the time it’s loaded. Only the operating system loader knows where it will finally reside—and then only just before it’s actually loaded.

As a result, shared libraries can only be built from a special class of object file called “PIC” objects. PIC stands for “Position-Independent Code”, and implies that absolute references in the object code are not *really* absolute. When the “`-fPIC`” option is used on the compiler command line, the compiler will use somewhat less efficient relative addresses in code branches, rather than the usual absolute addresses. Such position-independent code may be loaded anywhere, because no where within the code will you find a reference to an absolute address.

There are various nuances to generating and using position-independent code, and you should become familiar with them all before using them, so that you can choose the option that

is most appropriate for your situation. For example, the GNU C compiler also supports a “-fpic” option, which uses a slightly quicker system supported, but more limited mechanism to accomplish relocatable object code. Wikipedia has a very informative page on [position-independent code](#) (although I find its treatment of Windows DLLs to be somewhat less than accurate).

### Fixing the jupiter “PIC” problem

From what we now understand, one way to fix our linker error is to add the “-fPIC” option to the compiler command line for the source files that comprise the `libjupcommon.a` static library. Let's try that:

#### `common/Makefile.am`

```
noinst_LIBRARIES = libjupcommon.a
libjupcommon_a_SOURCES = jupcommon.h print.c
libjupcommon_a_CFLAGS = -fPIC
```

And now we'll try the build again:

```
$ autoreconf
$ make
...
gcc -DHAVE_CONFIG_H -I. -I../common -I.. -fPIC
-g -O2 -MT libjupcommon_a-print.o -MD -MP -MF
.deps/libjupcommon_a-print.Tpo -c
-o libjupcommon_a-print.o `test -f 'print.c' ||
echo '../common/'`print.c
...
/bin/sh ../libtool --tag=CC --mode=link gcc -g
-O2 ../common/libjupcommon.a -o libjupiter.la
-rpath /usr/local/lib libjupiter_la-jup_print.lo
-lpthread

*** Warning: Linking ... libjupiter.la against the
*** static library libjupcommon.a is not portable!

libtool: link: gcc -shared
.libs/libjupiter_la-jup_print.o
../common/libjupcommon.a -lpthread -Wl,-soname
-Wl,libjupiter.so.0 -o .libs/libjupiter.so.0.0.0

libtool: link: (cd .libs && rm -f libjupiter.so.0
&& ln -s libjupiter.so.0.0.0 libjupiter.so.0)

libtool: link: (cd .libs && rm -f libjupiter.so
&& ln -s libjupiter.so.0.0.0 libjupiter.so)

libtool: link: ar cru .libs/libjupiter.a
../common/libjupcommon.a
libjupiter_la-jup_print.o

libtool: link: ranlib .libs/libjupiter.a

libtool: link: (cd .libs && rm -f libjupiter.la
&& ln -s ../libjupiter.la libjupiter.la)
```

...

Well, now we have a shared library, built properly with position-independent code, as per system requirements. However, we still have that strange warning about the portability of linking a Libtool library against a static library. The problem here is not in what we're doing, but rather in the way we're doing it. You see, the PIC concept does not apply to all hardware architectures. Some CPUs don't support any form of absolute addressing in their instruction sets. As a result, native compilers for those platforms don't support a "`-fPIC`" option—it has no meaning for them.

If we tried (for example) to compile our code on an IBM RS/6000 system using the native IBM compiler, it would "hiccup" when it came to the "`-fPIC`" option because it doesn't make sense to support such an option on a system where *all* code is automatically generated as position-independent code. One way we could get around this problem would be to make the "`-fPIC`" option conditional in our `Makefile.am` file, based on the type of the target system, and the tools we're using. But that's exactly the sort of problem that Libtool was designed to address! We'd have to account for all of the different Libtool target system types and tool sets in order to handle the entire set of conditions that Libtool handles.

The way around this portability problem then is to let Libtool generate our static library also. Libtool makes a distinction between static libraries that are installed as part of a developer's kit, and static libraries used only internally in a project. It calls such internal static libraries "convenience" libraries, and whether or not such a convenience library is generated depends on the prefix used with the `LT_LIBRARIES` primary. If the `noinst` prefix is used, then Libtool assumes that we want a convenience library because there's no point in generating a shared library that will never be installed. Thus, convenience libraries are always generated as static archives.

The reason for distinguishing between internal convenience static libraries and other forms of static library is that convenience libraries are always built, whereas non-convenience static libraries are only built if the "`--enable-static`" option is specified on the `configure` command line (or conversely, if the "`--disable-static`" option is NOT specified).

### Customizing Libtool with `LT_INIT` options

Default values for enabling or disabling static and shared libraries can be specified in the argument list passed into the `LT_INIT` macro in the `configure.ac` file. Let's take a quick look at the `LT_INIT` macro. This macro may be used with or without arguments. When used with arguments, it accept a single argument, which is a white-space separated list of key words. The following key words are valid:

- `dlopen` — Enable checking for "`dlopen`" support.

This option should be used if the package makes use of the “-dlopen” and “-dlpreopen” libtool flags, otherwise libtool will assume that the system does not support dlopening.

- `disable-fast-install` — Change the default behavior for `LT_INIT` to disable optimization for fast installation. The user may still override this default, depending on platform support, by specifying “--enable-fast-install” to configure.
- `shared` — Change the default behavior for `LT_INIT` to enable shared libraries. This is the default on all systems where Libtool knows how to create shared libraries. The user may still override this default by specifying “--disable-shared” to configure.
- `disable-shared` — Change the default behavior for `LT_INIT` to disable shared libraries. The user may still override this default by specifying “--enable-shared” to configure.
- `static` — Change the default behavior for `LT_INIT` to enable static libraries. This is the default on all systems where shared libraries have been disabled for some reason, and on most systems where shared libraries have been enabled. If shared libraries are enabled, the user may still override this default by specifying “--disable-static” to configure.
- `disable-static` — Change the default behavior for `LT_INIT` to disable static libraries. The user may still override this default by specifying “--enable-static” to configure.
- `pic-only` — Change the default behavior for libtool to try to use only PIC objects. The user may still override this default by specifying “--without-pic” to configure.
- `no-pic` — Change the default behavior of libtool to try to use only non-PIC objects. The user may still override this default by specifying “--with-pic” to configure.

*NOTE: I've omitted the description for the `win32-dll` option, because it doesn't apply to this book.*

Now, let's return to our project. The conversion from an older static library to a new Libtool convenience library is simple enough—all we have to do is add “`LT`” to the primary name and remove the “`-fPIC`” option, and the associated variable, as there were no other options being used in that variable:

#### **common/Makefile.am**

```
noinst_LTLIBRARIES = libjupcommon.la
libjupcommon_la_SOURCES = jupcommon.h print.c
```

Now when we try to build, here's what we get:

```

$ autoreconf
$ ./configure
...
$ make
...
/bin/sh ../libtool --tag=CC --mode=compile gcc
  -DHAVE_CONFIG_H -I. -I../common -I..
  -g -O2 -MT print.lo -MD -MP -MF .deps/print.Tpo
  -c -o print.lo ../common/print.c

libtool: compile: gcc -DHAVE_CONFIG_H -I.
  -I../common -I.. -g -O2 -MT print.lo -MD -MP
  -MF .deps/print.Tpo -c ../common/print.c
  -fPIC -DPIC -o .libs/print.o
...
/bin/sh ../libtool --tag=CC --mode=link gcc -g -O2
  -o libjupcommon.la print.lo -lpthread

libtool: link: ar cru .libs/libjupcommon.a
  .libs/print.o
...
/bin/sh ../libtool --tag=CC --mode=link gcc -g -O2
  ../common/libjupcommon.la -o libjupiter.la
  -rpath /usr/local/lib libjupiter_la-jup_print.lo
  -lpthread

libtool: link: gcc -shared
  .libs/libjupiter_la-jup_print.o
  -Wl,--whole-archive
  ../common/.libs/libjupcommon.a
  -Wl,--no-whole-archive -lpthread -Wl,-soname
  -Wl,libjupiter.so.0 -o .libs/libjupiter.so.0.0.0
...

```

We can see that the common library is now built as a static convenience library because the `ar` utility is used to build `libjupcommon.a`. Libtool also seems to be building files with new and different extensions. A closer look will discover extensions such as “.lo” and “.la”. If you take a closer look at these files, you’ll find that they’re actually descriptive text files containing object and library meta data. Let’s look at the `common/libjupcommon.la` file:

#### **common/libjupcommon.la**

```

# libjupcommon.la - a libtool library file
# Generated by ltmain.sh (GNU libtool) 2.2
#
# Please DO NOT delete this file!
# It is necessary for linking the library.

# The name that we can dlopen(3).
dlname=''

# Names of this library.
library_names=''

# The name of the static archive.

```

```

old_library='libjupcommon.a'

# Linker flags that can not go in dependency_libs.
inherited_linker_flags=''

# Libraries that this one depends upon.
dependency_libs=' -lpthread'
...

```

The various fields in these files help the linker—or rather the `libtool` wrapper script—to determine certain options that would otherwise have to be remembered by the developer and passed on the command line to the linker. For instance, the library’s shared name and static name are remembered here, as well as any required libraries. The list of library dependencies is remembered here as well. In this library, for example, we can see that `libjupcommon.a` depends on the `pthread` library. But, using Libtool, we don’t have to pass a “`-lpthread`” option on the `libtool` command line because `libtool` can see in this meta data file that the linker will need this, so it passes the option for us.

Making these files human-readable was a minor stroke of genius, as they can tell us a lot about our Libtool libraries, at a glance. These files are designed to be installed with their associated binaries, and in fact, the `make install` rules generated by Automake for Libtool libraries do just this.

## The Libtool library versioning scheme

If you’ve spent any time at all working at the Linux command prompt, then you’ll certainly recognize this series of executable and link names. *NOTE: There’s nothing special about `libz`—I am merely using this library as a common example:*

```

$ ls -dal /lib/libz*
... /lib/libz.so.1 -> libz.so.1.2.3
... /lib/libz.so.1.2 -> libz.so.1.2.3
... /lib/libz.so.1.2.3

```

If you’ve ever wondered what this means, then read on. Libtool provides a versioning scheme for shared libraries that has become prevalent in the Linux world. Other operating systems use different versioning schemes for shared libraries, but the one defined by Libtool has become so popular that people often associate it with Linux, rather than with Libtool. This is not entirely an unfair assessment because the Linux loader honors this scheme to a certain degree. But to be completely fair, it’s Libtool that should be given the credit for this versioning scheme.

One interesting aspect of this scheme is that, if not understood properly, people can easily mis-use or abuse the system without intending to. People who don’t understand this system tend to think of the numeric values as *major*, *minor* and *revision*, when in fact, these values have very specific meaning to the operating system loader, and must be updated appropriately for each new library version in order to keep from

confusing the loader.

I remember a meeting I had at work one day several years ago with my company's corporate versioning committee. This committee's job was to come up with software versioning policy for the company as a whole (a grandiose vision, in and of itself). They wanted us to ensure that the version numbers incorporated into our shared library names were in alignment with the corporate software versioning standard. It took me the better part of a day to convince them that a shared library version was not related to a product version in any way, nor should such a relationship be established, or enforced by them or anyone else.

Here's why. The version number on a shared library is not really a library version, but rather an interface version. The interface I'm referring to here is the application programming interface (API) presented by a library to the potential user—a programmer wishing to call functions in the interface. As the GNU Libtool manual points out, a program has a single well-defined entry point (usually called "main", in the C language). But a shared library has multiple entry points that are generally not standardized in a widely understood manner. This makes it much more difficult to determine if a particular version of a library is "interface-compatible" with another version of the same library.

### Microsoft DLL versioning

Consider Microsoft Windows Dynamic Link Libraries (DLLs). These are shared libraries in every sense of the word. They provide a proper application programming interface. But unfortunately, Microsoft has in the past provided no integrated DLL interface versioning scheme. As a result, Windows developers often refer to DLL versioning issues (tongue-in-cheek, I'm sure) as "DLL hell".

As a partial fix to this problem, on Windows systems, DLLs can be installed into the same directory as the program that uses them, and the Windows operating system loader will always attempt to use the local copy first before searching for a copy in the system path. This alleviates part of the problem because a specific version of the library can be installed with the package that requires it.

This is not a good long-term solution, however, because one of the major benefits of shared libraries is that they can be shared—both on disk, and in memory. If every application has its own copy of a different version of the library, then this benefit of shared libraries is lost—both on disk and in memory.

Recently, Microsoft invented the concept of the "Side-by-Side Cache" (sometimes referred to as "SxS"), which allows developers to associate a unique identification value (a GUID, in fact) with a particular version of a DLL installed in a system location. This location is named by the dll name and version identifier. Applications built against SxS-versioned libraries have meta data stored in their

executable headers that indicate the particularly versioned DLLs that they require. If the right version is found (by newer OS loaders) in the SxS cache, then they load and use it. Based on policy in the executable header meta data, they can then revert to the older scheme of looking for a local and then a global copy of the DLL. This is a vast improvement over earlier solutions—providing a very flexible versioning system.

Linux and other Unix-like systems that support shared libraries manage interface versions using the Libtool versioning scheme. In this scheme, shared libraries are said to support a range of interface versions, each identified by a unique integer value. If any aspect of an interface changes in any way between public releases, then it can no longer be considered the same interface. It becomes a new interface, identified by a new integer interface value. To make the interface versioning process comprehensible to the human mind, each public release of a library wherein the interface has changed simply acquires the next consecutive interface version number. Thus, a given shared library may support versions 2-5 of an interface.

Libtool shared libraries follow a naming convention that encodes the interface range supported by a particular shared library. A shared library named `“libname.so.1.2.3”` contains the library interface version number, `“1.2.3”`. These three values are respectively called the library interface `“current”`, `“revision”` and `“age”` values.

The `“current”` value represents the current interface version number. This is the value that changes each time a new interface version must be declared because the interface has changed in any way since the last interface. Consider a shared library wherein the developer has added a new function to the set of functions exposed by this library since the last release. The interface can't be considered the same in this new version as it was in the previous version because there's one additional function. Thus, its `“current”` number must be increased from `“0”` to `“1”`.

The `“age”` value represents the number of back-versions supported by the shared library. In mathematical terms, the library is said to support the interface range, `“current” - “age”` through `“current”`. In the example I just gave, a new function was added to the library, so the interface presented in this version of the library is not the same as that presented in the previous version. However, the previous version of the interface is still fully supported because the previous interface is a proper subset of the current interface. Thus, this library could conceivably be named `“libname.so.1.0.1”`, where the range of supported interfaces is 1 - 1 (or 0) through 1, inclusive.

The `“revision”` value merely represents a serial revision of the current interface. That is, if no changes are made to a shared library's interface between releases, then the library name should change in some manner, but both the `“current”` and `“age”` values would be the same, as the

interface has not changed. The “`revision`” value is incremented to reflect the fact that this is a new release of the same interface. If two libraries exist on a system with the same name, “`current`” and “`age`” values, then the operating system loader will always select the library with the higher “`revision`” value.

To simplify the release process for shared libraries, the GNU Libtool manual provides an algorithm that should be followed step-by-step for each new version of a library that is about to be publically released. I'll reproduce the algorithm verbatim here for your information:

1. Start with version information of “`0 : 0 : 0`” for each libtool library. [This is done automatically by simply omitting the “`-version`” option from the list of linker flags passed to the `libtool` script.]
2. Update the version information only immediately before a public release of your software. More frequent updates are unnecessary, and only guarantee that the “`current`” interface number gets larger faster.
3. If the library source code has changed at all since the last update, then increment “`revision`” (“`c:r:a`” becomes “`c:r+1:a`”).
4. If any interfaces [exported functions or data] have been added, removed, or changed since the last update, increment “`current`”, and set “`revision`” to 0.
5. If any interfaces have been added since the last public release, then increment “`age`”.
6. If any interfaces have been removed since the last public release, then set “`age`” to 0.

Let's analyze this process a bit. This is an algorithm, and as such it is designed to be followed step by step, as opposed to jumping directly to the steps that appear to apply to your case. For example, if you removed any API functions from your library since the last release, you would not simply jump to the last step and set “`age`” to zero. Rather, you would follow all of the steps properly until you reached the last step, and *then* set “`age`” to zero.

In greater detail: Let's assume that this is the second release of a library, and that the first release was named “`libexample.so.0.0.0`”, and that one new function was added to the API during this development cycle, and one old function was deleted. The effect on this release of the library would be as follows:

1. (n/a)
2. (n/a)
3. `libexample.so.0.0.0` ->  
`libexample.so.0.1.0` (library source was changed)
4. `libexample.so.0.1.0` ->  
`libexample.so.1.0.0` (library interface was modified)
5. `libexample.so.1.0.0` ->  
`libexample.so.1.0.1` (one new function was added)

```
6. libexample.so.1.0.1 ->
   libexample.so.1.0.0 (one old function was
   removed)
```

Why all the “hoop jumping”? Well, because, as I alluded to earlier, the versioning scheme is honored by the linker and the operating system loader. When the linker creates the library name table in an executable image header, it writes the list of supported library versions along side of each entry in this table. When the loader searches for a matching library, it looks for the latest version of the library required by the executable. Thus, `libname.so.0.0.0` can coexist in the same directory as `libname.so.1.0.0` without any problem. Programs that need the earlier version (which supports *only* the later interface because of the deleted function) will properly and automatically have it loaded into their process address space, just as will programs that require the later version properly have the “1.0.0” version loaded.

One more point regarding interface versioning. Once you fully understand Libtool versioning, you’ll find that even the above algorithm does not cover all possible interface modification scenarios. Consider, for example, version “0.0.0” of a shared library that you maintain. Now, assume you add a new function to the interface for the next public release. This second release is properly named version “1.0.1”, because the library supports both interfaces 0 and 1. Just before the third release of the library, you realize that you didn’t really need that new function after all, and so you remove it. Assume also that this is the only change made to the library interface in this release. The above algorithm would have this release named version “2.0.0”. But in fact, you’ve merely removed the second interface, and are now presenting the original interface once again. Technically, this library should be properly named version “0.1.0”, as it presents a second release of version 0 of the shared library interface.

## Using `libltdl` to “`dlopen`” a shared library

### Summary

### Chapter Notes

As with the last three chapters, the notes in this chapter are comprised of the complete source files for the Jupiter project. These files include all of the additional code and changes made while adding Libtool shared libraries to the project, as well as the conversion of the common static library over to a Libtool convenience library.

Again, as usual, no formatting has been performed on these files, so you may cut and paste to your heart’s delight, but do beware of spaces in front of makefile command lines.

#### `configure.ac`

```
#                                     -*- P
# Process this file with autoconf to produce a config
```

```

AC_PREREQ([2.61])
AC_INIT([Jupiter], [1.0], [bugs@jupiter.org])
AM_INIT_AUTOMAKE
LT_PREREQ([2.2])
LT_INIT

AC_CONFIG_SRCDIR([src/main.c])
AC_CONFIG_HEADERS([config.h])

# Checks for programs.
AC_PROG_CC
AC_PROG_RANLIB
AC_PROG_INSTALL
AM_PROG_CC_C_O

# Checks for header files (1).
AC_HEADER_STDC

# Checks for command line options
AC_ARG_ENABLE([async-exec],
  [AS_HELP_STRING([--disable-async-exec],
    [disable asynchronous execution @<:@default: no@:
    [async_exec=${enableval}]],
  [async_exec=yes])

if test "x${async_exec}" = xyes; then
  have_pthreads=no
  AC_SEARCH_LIBS([pthread_create], [pthread],
    [have_pthreads=yes])

  if test "x${have_pthreads}" = xyes; then
    AC_CHECK_HEADERS([pthread.h], [],
      [have_pthreads=no])
  fi

  if test "x${have_pthreads}" = xno; then
    echo "-----"
    echo "Unable to find pthreads on this system."
    echo "Building a single-threaded version.    "
    echo "-----"
    async_exec=no
  fi
fi

if test "x${async_exec}" = xyes; then
  AC_DEFINE([ASYNC_EXEC], 1, [async exec enabled])
fi

# Checks for header files (2).
AC_CHECK_HEADERS([stdlib.h])

# Checks for libraries.
# Checks for typedefs, structures, and compiler chara
# Checks for library functions.

AC_CONFIG_FILES([Makefile
                 common/Makefile

```

```

        include/Makefile
        libjup/Makefile
        src/Makefile])

AC_OUTPUT

echo \
"-----

${PACKAGE_NAME} Version ${PACKAGE_VERSION}

Prefix: '${prefix}'.
Compiler: '${CC} ${CFLAGS} ${CPPFLAGS}'

Package features:
  Async Execution: ${async_exec}

Now type 'make @<:@<target>@:>@'
  where the optional <target> is:
    all             - build all binaries
    install         - install everything

-----"

```



### Makefile.am

```
SUBDIRS = common include libjup src
```

### common/Makefile.am

```
noinst_LTLIBRARIES = libjupcommon.la
libjupcommon_la_SOURCES = jupcommon.h print.c
```

### common/print.c

```

#include <jupcommon.h>

#if HAVE_CONFIG_H
# include <config.h>
#endif

#include <stdio.h>
#include <stdlib.h>

#if HAVE_PTHREAD_H
# include <pthread.h>
#endif

static void * print_it(void * data)
{
    printf("Hello from %s!\n", (char *)data);
    return 0;
}

int print_routine(char * name)
{
    #if ASYNC_EXEC
        pthread_t tid;

```

```

        pthread_create(&tid, 0, print_it, name);
        pthread_join(tid, 0);
#else
        print_it(name);
#endif
        return 0;
}

```

#### **common/jupcommon.h**

```

#ifndef JUPCOMMON_H_INCLUDED
#define JUPCOMMON_H_INCLUDED

int print_routine(char * name);

#endif /* JUPCOMMON_H_INCLUDED */

```

#### **include/Makefile.am**

```
include_HEADERS = libjupiter.h
```

#### **include/libjupiter.h**

```

#ifndef LIBJUPITER_H_INCLUDED
#define LIBJUPITER_H_INCLUDED

int jupiter_print(char * name);

#endif /* LIBJUPITER_H_INCLUDED */

```

#### **libjup/Makefile.am**

```

lib_LTLIBRARIES = libjupiter.la
libjupiter_la_SOURCES = jup_print.c
libjupiter_la_LDFLAGS = ../common/libjupcommon.la
libjupiter_la_CPPFLAGS = -I$(top_srcdir)/include -I$(

```



#### **libjup/jup\_print.c**

```

#include <libjupiter.h>
#include <jupcommon.h>

int jupiter_print(char * name)
{
    print_routine(name);
}

```

#### **src/Makefile.am**

```

bin_PROGRAMS = jupiter
jupiter_SOURCES = main.c
jupiter_CPPFLAGS = -I$(top_srcdir)/include
jupiter_LDADD = ../libjup/libjupiter.la

check_SCRIPTS = greptest.sh
TESTS = $(check_SCRIPTS)

```

```
greptest.sh:
    echo './jupiter | grep "Hello from .*jupiter!'
    chmod +x greptest.sh
```

```
CLEANFILES = greptest.sh
```



### src/main.c

```
#include <libjupiter.h>

int main(int argc, char * argv[])
{
    jupiter_print(argv[0]);
    return 0;
}
```

[◀ Chapter 4:  
Automatically writing  
makefiles with  
Automake](#)

u  
p

[Chapter 6: An  
autotools example ▶](#)

[Login or register to post comments](#) 1138 reads