

## **NetBuild: Transparent Cross-Platform Access to Computational Software Libraries**

Keith Moore

Jack Dongarra

*Innovative Computing Laboratory*

University of Tennessee

Email: {moore,dongarra}@cs.utk.edu

Telephone: +1 865 974 8295

Fax: +1 865 974 8296

*Summary:* NetBuild is a suite of tools which automate the process of selecting, locating, downloading, configuring, and installing computational software libraries from over the Internet, and which aid in the construction and cataloging of such libraries. Unlike many other tools, NetBuild is designed to work across a wide variety of computing platforms, and perform fine-grained matching to find the most suitable version of a library for a given target platform. We describe the architecture of NetBuild and its initial implementation.

## 1. Overview

### A. Description and Goals

As computationally intensive modeling and simulation become increasingly important staples of scientific life across nearly every domain and discipline, the challenge of deploying the software that encodes the underlying computational science has become increasingly complex and pressing. The ubiquity of the Internet and the establishment of large software repositories, such as Netlib [1], have made it relatively easy to provide basic access to such software. However, a rapidly growing and more diverse user community still has to cope with the process of collecting, compiling, installing, configuring, and managing the software on a constantly changing and increasingly elaborate computing infrastructure.

For all of its promise, the development of Grid-enabled computing [2] may actually exacerbate this problem. A program may require certain software libraries be installed on large numbers of diverse computing platforms. These platforms are administered by a variety of concerns, and therefore have wide variation in their library support. It is difficult to keep track of which libraries are installed on each platform. In general, the more libraries on which a program depends, the fewer environments in which it can run.

In addition, an increased acceptance of software reuse has enabled creation of more and more sophisticated software packages, built out of numerous components obtained from diverse sources. This in turn has created a configuration management problem for users and system administrators, who must devote valuable time to obtaining and installing those components, tracking changes (including enhancements, bug fixes, and sometimes security fixes) to those components. Version conflicts, where one software package requires version X of a library while another package on the same platform requires version Y, are common.

Packages such as ATLAS [3,4] can increase the effective performance of many platforms by optimizing the parameters of the computation for that specific platform. However, for the same reason, these packages can be sensitive to specific platform characteristics (CPU features, cache sizes, etc.). These packages must therefore be configured on a per-platform basis, using fairly fine-grained attributes of each platform to identify the proper version of the library.

To enable scientists and engineers to spend more of their time thinking about their research and less time managing the quirks and intricacies of their computing tools, new approaches must be found to automate these tasks. The NetBuild tools attempt to address the problem of automatic library identification, configuration, and installation. Though these tools can be used independently of the Grid, we believe that they are well-positioned to address problems of large-scale complex distributed software systems of which the Grid is an example.

### B. Services Provided

The NetBuild suite of tools provides the following services:

- Automatic tagging (at compilation time) of computational software libraries with catalog information, including an indication of the target platform characteristics for which the library is suitable,
- Replication of libraries to multiple servers from which the library can be downloaded,
- Indexing of libraries according to targeted platform characteristics, and indexing of locations of each instance of a library,
- Automated searching of libraries for candidates meeting target platform criteria,
- Matching of candidate libraries against fine-grained characteristics of a target platform,
- Downloading of a selected candidate library from an appropriate location,
- Cryptographic verification of the authenticity and integrity of a library,

- Linking of the downloaded library into an executable image, and
- Run-time dynamic loading of a downloaded library into a running program.

From the perspective of the author of a scientific computing program, NetBuild provides the ability to incorporate standardized computational software libraries, at link time or run time, without imposing the requirement that they be pre-installed on a target platform.

### **C. Systems/Sites/Users Served**

We initially intend to make the most popular of the mathematical software libraries on Netlib available for use from NetBuild, with libraries pre-built for the most popular computing platforms. These libraries will be available on all Netlib mirrors. There are efforts underway to equip Netlib with facilities to allow computational software developers to compile and test their libraries on a wide range of target platforms.

Once NetBuild is more stable, we intend to make the entire suite of tools available to others, as free software, via the Internet. This will allow other developers to build NetBuild-compatible libraries, and other repositories to establish NetBuild-compatible servers.

### **D. Status**

NetBuild currently exists as a prototype, with limited distribution to alpha testers.

## **2. Architecture**

The following tools comprise the principal components of NetBuild:

- The NetCompile tool assists the library developer in constructing libraries that can be used by NetBuild. It performs several functions, including:
  - Invoking the compiling system's compilers, linkers, and library managers with appropriate options to produce either code which is portable across a variety of target environments, or code which is optimal for a particular target environment, according to the developer's preference,
  - Concurrently with compilation of an object module, producing meta-data which precisely describes both the compilation environment and the target platform(s) for which the object module is presumed to be appropriate,
  - Cryptographically signing the library and metadata, to allow users to verify these components for authenticity and integrity,
  - Incorporating the results into a package from which individual components can be extracted, and which is usable by NetBuild.
- The NetBuild tool assists the computational scientist in incorporating NetBuild libraries into his or her programs. Its functions include
  - Intercepting calls to the build platform's compilers and linkers,
  - Identifying libraries that are needed by the software and which are not installed on the build platform,
  - Searching for instances of the missing libraries which are compatible with the build platform and appropriate for the target platform,
  - Downloading those libraries to the build platform,
  - Verifying their signatures for authenticity and integrity,

- Invoking the build platform's compilers and linkers with appropriate options to allow the downloaded libraries to be linked into the compiled program.
- NetLoader is a set of library routines which are analogous to the NetBuild tool, but which operate at run-time. Specifically, when asked to load a particular library, NetLoader:
  - Searches for instances of that particular library which are appropriate for the target platform,
  - Downloads the library to the target platform,
  - Verifies the signature on that library for authenticity and integrity, and
  - Incorporates that library into the running program

In addition the NetBuild system includes a server which indexes object files according to their characteristics, and which answers queries from NetBuild tools and NetLoader library calls for locations of libraries matching particular characteristics. The libraries are provided using ordinary HTTP servers.

### 3. Implementation

#### NetBuild tool

The NetBuild tool may be invoked in either of two ways - explicitly, by use of the `netbuild` command, or implicitly. In the latter case the user or program invokes a compiler or linker as it normally would, but the `PATH` environment variable has been modified to contain a directory of *shims* that intercept calls to the compiler and linker. Users do not normally include the shim directory in their `PATHs`; this is done only when the compiler is some descendant subprocess of an explicit invocation of `netbuild`. This allows compilers and linkers called by `make` and other tools to be invoked in the NetBuild environment simply by typing `netbuild make`, without changes to the Makefiles or other compilation scripts.

#### General

`netbuild` is a C program that runs on several UNIX-derived and UNIX-like platforms. It works as follows:

- The program checks the name by which it was invoked. If that name ends in `netbuild`, it adds the directory containing NetBuild's shims to the `PATH` environment variable, and treats the remainder of the command-line as a command to be invoked with the modified `PATH`.
- Otherwise, NetBuild parses the command-line arguments as if it were the compiler or linker, identifying options that specify libraries to be linked.
- For each of these libraries, NetBuild determines whether those libraries are already installed on the local system.
- For each of the libraries that are not installed, NetBuild consults one or more network servers in an attempt to find libraries which match the characteristics of the target platform. When it finds such a library it will download it to the local system. Previously downloaded libraries are cached so they are not downloaded again if they have not changed.
- The authenticity and integrity of the libraries is verified, and if valid, the libraries are installed in local directories which are private to NetBuild.
- The system compiler or linker is then invoked with extra options to cause the newly-downloaded libraries to be linked in along with the resident ones.

#### Option parsing

Since `netbuild` is invoked as if it were the normal system compiler or linker, it needs to be able to understand

options that vary from one compiler or linker to another. `netbuild` therefore has a configurable parser for command-line options. The parser can be configured on a per-host, per-platform, and per-compiler basis.

NetBuild need not understand the full syntax and semantics of each option, but it does need to know which options require additional arguments (so that subsequent arguments beginning with a hyphen are not treated as separate options), which options specify libraries to be linked, and which options specify local directories which should be searched. In the future, other extensions may be necessary. For instance, it may also need to be aware of options which specify variants of the compiler's target platform, so that it can use the correct libraries if the specified target is different than the default one.

### Searching for local libraries

`netbuild` must search local directories to determine whether some of the requested libraries are already resident. Since these directories vary from one target platform to another and from one compiler to another on the same platform, the list of directories which `netbuild` consults is configurable. In addition, any directories specified on the command-line are also consulted. Finally, since naming conventions vary from one platform to another, `netbuild` can be configured to understand the file naming conventions for libraries on the local platform. For instance, library `xyz` might be matched by any of `libxyz.a`, `libxyz.so`, or `libxyz.so.1.2`.

### Identifying target platform characteristics

NetBuild first uses the GNU `config.guess` program to determine a canonical name for the target platform. This is a string which is generally of the form `CPUtype-vendor-OSname` where `OSname` also contains a version number. This is nowhere nearly precise enough for NetBuild to be able to find the "best" match for a particular target platform, but this naming scheme is more-or-less consistently used by a large number of software packages. In many cases (especially where performance is not critical) this is sufficient to find an appropriate match. In other cases, it is useful to obtain a list of candidate libraries against which more precise matching can be performed.

If multiple candidate libraries matching the canonical platform name exist, NetBuild compares the actual characteristics of the target platform against each library's *constraint expression* to determine the "best" library. This constraint expression is of a mathematical expression in terms of real-valued attributes of the library and target platform. This expression is evaluated to produce a numerical "match quality". A match quality of zero indicates that the candidate library is not suitable for the target platform. The library with the highest nonzero match quality is assumed to be the best match. The set of attributes which can be used in computing match quality is platform-specific and extensible, but may include:

- CPU architecture family (e.g. Intel IA32),
- instruction set version (where later versions are supersets of earlier ones)
- instruction set extensions (e.g. MMX, 3DNow!),
- sizes of various caches,
- CPU chip vendor,
- CPU chip version,
- number of processing elements,
- primary memory required,
- working set size required,
- disk space required,
- operating system,

- operating system version,
- operating system features/extensions,
- compiler,
- compiler version,
- compiler ABI or calling sequence,
- source language,
- object file format (e.g. ELF, COFF, a.out)

It can also be necessary to distinguish between attributes of the platform on which the library was compiled, actual attributes of the build platform and build tools, actual attributes of the target platform, and assumed attributes of the target platform for which the library author configured the code.

### **Identifying available network-resident libraries**

NetBuild currently consults an external web server to identify which libraries might be available for a particular platform. In order to find which versions of library "xyz" are available, NetBuild accesses a URL which is constructed from the library name, the target platform canonical name, and a list of features of the target platform. Accessing that URL causes a "common gateway interface" (CGI) script to be invoked, which returns a list of URLs of suitable candidate libraries in order of preference. In other words, the library matching algorithm is currently implemented on the NetBuild server rather than on the client. This is a temporary measure designed to allow the NetBuild developers to continually fine-tune the matching algorithms and the attributes of the libraries without having to constantly update code in the field that is being tested.

### **Library container file format**

NetBuild currently expects downloaded libraries to be in gzip-compressed tar format. The actual library archive is one component of the tar file, the metadata is another, and the signature on the library and metadata are yet another component. The archive may also contain other components. After downloading the tar file the archive is extracted, renamed as necessary, verified against its signature, and copied to the cache directory.

We used tar because it allowed for easy prototyping. However, the tar format has a number of shortcomings for this purpose, and we expect to replace it with a different format before a public release of NetBuild.

### **Caching**

NetBuild caches files that are downloaded from the network so that they are not downloaded again unless necessary. The cache is currently maintained on a per-user basis, in each user's private filestore, due to security concerns associated with maintaining a shared cache. Libraries downloaded from servers are stored in a directory whose name is derived from a hash of the (canonicalized) URL from which the library was obtained; a separate metadata file contains the last-modified date of that URL. Subsequent attempts to download that file use the HTTP "if-modified-since" directive which causes the file to be downloaded only if it has been changed. Note that the last change date of the container file which is downloaded may be different than the last change date of the actual library; thus it is possible for the metadata (and signature) to be updated even though the object library is unchanged. We can use this feature, for example, to change the constraints on use of a library to reflect experience with the library after it's publication.

### **Authenticity and Integrity verification**

We currently use GNU Privacy Guard (GPG) [5] to verify digital signatures on NetBuild libraries. Compatible signatures can be created with GPG or any of several PGP variants. Because the trust model for netbuild libraries is different from that of normal PGP signatures (just because you trust a signature on a library to be authentic does not mean you trust that it's okay to execute that library on your computer), the signatures used by NetBuild are kept on a

separate key ring in a separate directory.

GPG is used in NetBuild prototypes because it is easy to interface to, portable, readily available, and presumably free of patent issues. However the current implementation requires that GPG be installed in addition to NetBuild. To make it easier for the user to install NetBuild, it would be preferable for the signature verification code to be incorporated directly. This would allow NetBuild to support additional signature formats, and additional certificate formats such as X.509v3.

### **NetCompile**

The operation of NetCompile is similar to NetBuild, but the function is different. Instead of supplying missing libraries, NetCompile's job is to invoke the native compilers and library managers in such a way that:

- The compiler is optimally tuned for the assumed target environment, or to produce portable code;
- The characteristics of the compiler and the resulting object files are recorded,
- The characteristics of the assumed target are recorded,
- The resulting library and metadata are signed, and
- The result is packaged for easy downloading by NetBuild.

Whereas NetBuild tries to transparently preserve the native interface to each compiler or linker, NetCompile tries to provide a platform- and compiler-independent interface to compilers, linkers, and library managers.

### **NetLoader**

NetLoader consists of an application programmer's interface to the portions of NetBuild which download and verify code, along with a dynamic library loader. However, the matching algorithms are subtly different than those used by NetBuild. For instance, NetLoader can only use dynamically-loadable libraries, while NetBuild can also use static libraries. Another difference is due to target platforms that support emulation of the operating environment for non-native images - for instance, a NetBSD platform can often run an executable image compiled for Linux. In this case it is necessary to load a library that is compatible with the emulated environment, rather than one which is compatible with the target's native environment. On the other hand, if NetBuild is asked to install a particular executable image, it might substitute one intended for a different platform if the target platform supports emulation of the other platform and no native executable is available.

### **Distribution and Location of Libraries**

The distribution system takes advantage of the fact that Netlib is already extensively mirrored. Each distinct version of a library is assigned a unique (and cryptic) filename which will never be reused, so if a file with the same name appears on a mirror site, it's very likely to be an identical copy of that library. The NetBuild servers will maintain a list of mirror sites and the directories on which they store Netlib files, to use in determining the probable location of the library on the mirror. The servers initially will use heuristics (such as the last-changed-date of the library) to guess whether a particular mirror site is likely to have a copy of a particular library. (If the client doesn't find a copy of the library at that mirror, or if the library is truncated or altered, the client will try to obtain the library from another mirror.)

We plan to periodically check Netlib mirror sites for integrity and currency of mirroring, and update our location database accordingly.

### **Related Work**

NetBuild bears a resemblance to several other works:

The NetLink project [6] has similar goals to those of NetBuild, in that it is also focused on easing the burden of maintaining software libraries used in scientific computing. Their stated objective is to "identify a data distribution architecture [...] that can help to centralize the library maintenance and tuning". Their work appears to be focused on

identifying an appropriate data distribution architecture, whereas our work to date has focused on making NetBuild as transparent as possible, and effective matching of object files with target platform characteristics. We intend to actively cooperate with the NetLink project as part of our continued work on NetBuild.

Libtool [7] is a package which eases the burden of development of software for multiple computing environments, by providing a uniform interface to a variety of compilers, linkers, and library archivers. Libtool is similar to NetBuild and NetCompile in that both sets of tools run the existing compilers, linkers, etc. within a wrapper. Libtool allows developers to write scripts and makefiles which invoke compilers and linkers via the `libtool` command; these scripts and makefiles will be portable to every platform supported by libtool. Libtool also includes the ability to dynamically load libraries created by Libtool. Libtool has been ported to a wide variety of platforms. We have liberally borrowed ideas from libtool for use in NetBuild's library-building function, extending it to support the additional features and cataloging needed by NetBuild. We intend to keep NetBuild's libraries compatible with those generated by Libtool in order to maximize potential for code reuse.

The NetBSD operating system [8] contains a "package" facility [9] which automates the process of downloading source code for a particular program or library, verifying its integrity, configuring it to run on NetBSD, compiling it, installing it, and cataloging it for configuration management purposes. The package facility also handles dependencies, so that any components needed by the component to be installed are automatically compiled and installed also. The facility supports both source-code packages and pre-compiled binary packages, and has successfully been used to ease installation of over two thousand different tools and libraries. The package facility has been ported to each of the several hardware platforms supported by NetBSD, as well as to the Solaris operating system. The NetBSD package system is itself descended from a similar "ports" system incorporated into FreeBSD; similar but independently-derived features exist in Debian and RedHat Linux. We plan to borrow technology from the NetBSD package system for use in NetBuild - specifically to support automatic compilation of libraries from source code when pre-compiled objects are not available.

#### 4. Status and Future Plans

The NetBuild tools currently exist as prototypes. We are currently experimenting with various data models for platform and compiler characteristics, and working on more automatic derivation of these characteristics at library build time. Our next step will be to construct a "compile and test zoo" for the purpose of automatically building and testing large numbers of libraries on various platforms.

The current implementation is intended as a proof-of-concept and a testbed for new features rather than a code base for use by ordinary users. As such, it is designed to be flexible and configurable and easily implemented, rather than (say) secure and robust. The publicly-distributed version of NetBuild will need to pay much more attention to security and robustness issues..

#### 5. References

1. Browne S, Dongarra J, Grosse E, Rowan T. 1995. "The Netlib Mathematical Software Repository," *D-Lib Magazine*, September 1995. <http://www.dlib.org/dlib/september95/netlib/09browne.html>
2. Foster I, Kesselman C, eds. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufman: San Francisco, 1999.
3. Whaley, R C, Dongarra J. 1998. "Automatically Tuned Linear Algebra Software". Proceedings Supercomputing 1998 conference. [http://www.supercomp.org/sc98/TechPapers/sc98\\_FullAbstracts/Whaley814/INDEX.HTM](http://www.supercomp.org/sc98/TechPapers/sc98_FullAbstracts/Whaley814/INDEX.HTM)
4. Whaley R, Petitet A, Dongarra J. 2001. "Automated Empirical Optimization of Software and the ATLAS Project" *Parallel Computing* 2001; 27(1-2):3-25.
5. Gnu Privacy Guard. <http://www.gnupg.org/>
6. Holmqvist I, Lindström E. 1998. "NetLink: A Modern Data Distribution Approach Applied to Transparent Access of High Performance Software Libraries, in Kågström, et. al. (eds), *Applied Parallel Computing: Large Scale Scientific and Industrial Problems*, Lecture Notes in Computer Science, Springer-Verlag, 1998;1541:248-254.



7. GNU libtool website and documentation. <http://www.gnu.org/software/libtool/libtool.html> Free Software Foundation, 2001.
8. NetBSD - <http://www.netbsd.org/>
9. Feyrer H, Crooks A. ``Documentation on the NetBSD Package System"  
<ftp://ftp.netbsd.org/pub/NetBSD/packages/pkgsrc/Packages.txt>