

It's Magic: SourceMage GNU/Linux as HPC Cluster OS

Jörg Cassens and Zoran Constantinescu*
Norwegian University of Science and Technology (NTNU)
7491 Trondheim, Norway
{cassens|zoran}@idi.ntnu.no

July 25, 2003

Abstract

The goal of the presentation is to give an overview about how to build a commodity PC based GNU/Linux cluster for High Performance Computing (HPC) in a research environment. Due to the extreme flexibility of the GNU/Linux operating system and the large variety of hardware components, building a cluster for High Performance Computing (HPC) is still a challenge in many cases. At the Division of Intelligent Systems at the Norwegian University of Science and Technology (NTNU), we have build a 40 node HPC cluster for research purposes using the source-based GNU/Linux distribution Source Mage.

We describe a methodology for designing and installing a highly customized GNU/Linux cluster.

Different types of Linux distributions will be mentioned, binary-based and source-based, with their advantages and disadvantages.

The presentation will focus on using SourceMage for HPC, specifying the 'magical' ideas behind it: the ease of upgrading to the latest available version of the source code, a packaging system for keeping track of dependencies, optimized compiles for the hardware architecture used, easy integration of new packages, amongst others.

1 History

- **Version 1.0:** Original version published at LinuxTag 2003.
- **Version 1.1:** Nicer PDF output, removed some typos, included the GNU FDL.

*Copyright © 2003 Jörg Cassens and Zoran Constantinescu. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled 'GNU Free Documentation License'.

2 Introduction

High performance computing (HPC) has evolved into a small, stable, and high-priced market. When looking at HPC systems, we see two developments:

- A reduced number of vendors supporting the traditional type parallel computers and an
- increased performance and capacity of clusters based on off-the-shelf components in terms of CPU power, storage capacity, OS abilities, network, and availability of parallel software.

The Beowulf class of parallel computing machines started as a small research project at NASA in 1993, with the goal of building a one GFlops parallel workstation for less than \$50,000. The idea was to use commodity off-the-shelf (COTS) hardware and software configured as a cluster of machines, and which exhibits excellent price-performance, provides a robust environment and has broad applicability.

In June 2002, there were about 80 Beowulfs in the top500 (the list of most powerful computers, see [\[8\]](#)), and the Beowulf population is estimated at several thousands.

Two things are the main driving technologies for the success of Beowulf type clusters: cheap hardware and open source software. Most Beowulfs use mass market commodity off the shelf PC technology as hardware platform.

The main advantage is that by using no customized components it allows to buy the hardware from multiple vendors. This approach exploits components that are widely accepted by industry standards and benefit from prices resulting from heavy competition and mass production. Recent advances in performance of such components make them even more appealing for such clusters. One big advantage is their rapid response to technology trends. The latest advances in microprocessors, storage or other mass market technologies tend to find their way into PCs much faster than to other platforms. Technologies as old as only a few months can easily be integrated in Beowulf class systems.

Beowulf exploits readily available, usually free software systems. The success comes from the unification of public domain parallel tools and applications for the scientific software community. Such tools and applications are the result of decades of parallel processing research and on many attempts to use loosely coupled computers for a variety of applications. Some of these components include: parallel communication libraries, parallel file systems, tools for configuring, scheduling, managing and tuning parallel applications, higher level scientific libraries, etc. The quality of many such publicly available software is comparable with that offered by many commercial software vendors.

Most Beowulf class systems employ UNIX-like operating systems, with source code availability and low or no cost. The cost issue is very important because paying an OS license for each node of a large cluster (hundreds of nodes) could be prohibitively expensive. Source code availability is important because it allows custom modification to facilitate optimized performance. The most widely used operating systems in Beowulf clusters are Linux and BSD. There are also a few Beowulf clusters using Solaris or Windows as operating system.

One other important feature of Beowulf type clusters is that they enable do-it-yourself cluster computing, which fits perfectly for many academic and research environments.

The focus of most Beowulf class systems is science and engineering applications. These applications are mostly floating point intensive and require substantial amount of memory. There are many new applications in the area of data processing, especially in the new biocomputing field, where such systems are beginning to be used. Beowulfs are also finding significant use in computer science research, for the development of new tools and environments, such as metacomputing. One of fastest growing area of Beowulf use is in the computer science education. These systems offer flexibility and price advantages over other commercial offerings, which makes them ideal for pedagogical purposes.

3 Cluster Architectures

Not every Cluster is built as a High Performance Computing (HPC) resource. The two most common cluster forms are:

- **Load Balancer:** Distribute uniform workload, e.g. web server farms.
- **Computational Cluster:** Conquer a computationally difficult problem by using several machines.

In this paper, we take a look at the second type: A computational cluster as an alternative to a conventional supercomputer.

The underlying architectural idea in computers like the common desktop PC is the use of one instruction on one piece of data at any given time. For HPC, different architectures were traditionally discussed and used.

Take for example one piece of data and apply different processing directives at any given time, and you get the multiflow machine.

The most common form of traditional supercomputers used the same instruction on different data at the same time. This can most easily be achieved by having several CPUs working together. This type can additionally be divided into systems where all CPUs share access to the same memory and systems where each system sees a different memory. Applications can be of a type where memory access is shared or of a type where the teamworking is done by sending messages between them.

The theoretic step from a traditional supercomputer to cluster architectures seems now not to be very far fetched: what if the different memories and the different CPUs were not part of one single computer, but if one would take a collection of computers on a network that can function as a single computing resource with a distributed memory architecture.

3.1 ClustIS Architecture

For the original ClustIS project (see [1]), we have chosen simplified Desktop-PC's from one of our university's regular hardware suppliers. The computers have AMD Athlon

CPU's, 100Mbit/Sec ethernet cards, Hard disks, a simple graphics card and up to 2GB of RAM. The cases are of standard Desktop type with a standard Power Supply and no CD-ROM or Soundcards. They are placed side-by-side on workshop shelves. This reduced the costs considerably compared to rack servers.

One of the biggest advantages of specialized server hardware is the extended mean time between failures: the power supply is of bigger dimension, casing and cabling are of better quality, and so on.

Downtime of computing nodes is no issue for COTS clusters: just replace it with a spare one or wait for the repair. This is different for the master node, which is in charge of tasks like job scheduling. Therefore, our master node is a server-type rack mounted computer. We have a second, smaller rack mount PC integrated into the cluster. Under normal operating conditions, it works as a file server (NFS-mounts) for both the cluster and the Division of Intelligent Systems, offering an easy exchange of and access to data for the users of the cluster. In case of a failure of the master node, it can take over the jobs of the master node as well.

During the use of the cluster, a collaboration with another organizational unit of our Department was established. Therefore, some rack mounted PC's originally delivered for another cluster were integrated into ClustIS. Our concept proved to be flexible enough to integrate these new systems despite of some hardware differences.

Per may 2003, ClustIS consists of the following systems:

- Master: AMD Dual Athlon MP 2100+ (1.66 GHz), 2GB RAM, 80+120GB IDE HD, 1*Gigabit Ethernet, 1*100MBit Ethernet
- Node type 1: (16 nodes: 1...16) AMD Athlon XP 1700+ (1.46 GHz), 2GB RAM, 1*40GB IDE HD, 1*100MBit Ethernet
- Node type 2: (12 nodes: 17...28) AMD Athlon XP 1700+ (1.46 GHz), 1GB RAM, 1*40GB IDE HD, 1*100MBit Ethernet
- Node type 3: (8 nodes: 29...36) AMD Athlon MP 1600+ (1.4 GHz), 1GB RAM, 1*18GB SCSI HD, 2*100MBit Ethernet
- Node type 4: (1 node: 37) AMD Dual Athlon MP 1600+ (1.4 GHz), 1GB RAM, 3*18GB SCSI HD, 2*100MBit Ethernet
- Storage: AMD Athlon XP 1700+ (1.46 GHz), 0.5GB RAM, 8+2*80GB IDE HD, 1*Gigabit Ethernet, 1*100MBit Ethernet

Important for clusters is the network interconnect used. For ClustIS, we went for a low cost alternative by using a private switched 100Mbit/sec Ethernet between the nodes and giving both the Master and the Storage Node Gigabit Ethernet access to the switch.

This network type seems suitable for the size of our cluster and the need of the applications we run. For larger clusters or applications which depend on a high bandwidth or low latency interconnect, other network types might be more useful. These solutions exist (see e.g. Myrinet), but are more expensive.

4 GNU/Linux

Most of Beowulf clusters are using GNU/Linux as an operating system because of its performance, availability of source code, better device support, and wider user acceptance. The extremely low cost is also very important. The key features for GNU/Linux in research are: sophistication, accessibility and low cost.

Currently there are many Linux distributions available. The choice of which distribution to use in a cluster environment is influenced by many factors. For example, which software packages come with a distribution and/or the familiarity with a certain distribution.

4.1 GNU/Linux Distributions

There are basically two types of Linux distributions: binary based and source based. A binary distribution consists of a certain number of precompiled software packages, while a source based distribution is compiled at install time using the latest source code of the software.

Most binary distributions offer support for cluster use by including the necessary tools and libraries. System suppliers offers solutions based on existing commercial binary distributions like Red Hat or SuSE. Or Linux distributors offer complete systems with hard- and software, like in the case of the PPC-based Yellowdog distributions.

These system suppliers often focus on stability. The setup might therefore include older versions of software which is known to harmonize with each other and the hardware. This is a perfectly valid goal and often suits the needs in e.g. engineering domains or whenever the cluster is the tool, and not the research object.

Research systems, especially systems focusing on research in distributed or parallel systems itself, might be in need for more "bleeding edge" versions. It is not always easy to combine a binary distribution in a stable way with newer software versions, given the complexity of dependencies.

A source based distribution consists of the source code of the packages which are compiled and configured at install time. The source codes for the packages are always downloaded directly from the software authors' homepages and mirrors. This means that it will always be getting the latest and greatest version of software packages, unlike other distributions that ship with outdated packages. Then, they are compiled with the architecture and optimizations that the system administrator specifies. Finally, they are installed, tracked, and archived for easy removal and upgrades. See table 1 for an overview about the differences between source based and binary distributions.

A source based distribution offers the system administrator more control over the GNU/Linux system. At the same time, it offers enhanced performance and customization. The cost for this is the time required to compile the software packages. However, time for compilation can be allocated from the computer when this is not used.

Having always the latest version of the software, built from the current sources, makes the operating system free of known vulnerabilities and exploits.

There are a few source based GNU/Linux distributions: Source Mage, Gentoo, Lunar and Source Mage. They follow the same procedure of installation, by downloading the source code and compiling it. The major differences between these distributions are

	Binary	Source
Bytes to download	less	more
Time to compile	short	long
Install time	short	long
Latest software versions	no	yes
Compilation logs	no	yes
Optimized binaries	maybe	yes
Architecture specific binaries	maybe	yes

Table 1: *Small comparison table: Binary and source based distributions*

the way software packages are described (source code location, compile options and procedure, dependencies) and the number of available packages, i.e. which software packages can be installed using one distribution.

4.2 SourceMage

We decided to use the Source Mage GNU/Linux (SMGL) distribution, a source based distribution. It has many powerful features that set it apart from other distros.

Up-to-date packages can be auto-built using the optimization settings and build-time functionality the administrator wanted, rather than what other distro creator thought would be best for him. For example, if you don't want GNOME on the system, the applications won't have optional GNOME support enabled.

It has an advanced package management system, supporting a number of advanced features including dependencies, fine-grained package management, path sandboxing, safe unmerging, system profiles, virtual packages, config file management.

Source Mage is a simple yet powerful source based distribution for advanced systems administration. The package management system is called sorcery. Sorcery is comprised of modular, easily modified, command line and menu driven BASH scripts. A system administrator wielding sorcery can keep FHS 2.2 complaint Source Mage boxes current with the latest stable software releases.

Source Mage offers most of the features of modern GNU/Linux distributions. There is a menu driven installer on the Installation/Rescue CD image which simplifies the creation of a new box. After installation Source Mage has both command line and menu driven source management programs, making it easy to use both for beginners or more advanced administrators. It can also be used in automatic scripts.

Software installed by sorcery, with few exceptions, are the authors' latest stable release. Compiling software with the optimizations, options, and architecture that the sys admin specifies is how a Source Mage box achieves excellent run-time performance. Sorcery can discover and automatically fix broken library dependencies caused by an upgrade or removal of a library.

The Grimoire contains all of the "spells" that Source Mage uses. Each of these spells contains the instructions for downloading the source code, for compiling, and for installing a certain software package. By "cast"-ing the name of a program, the system

will download, compile and install the program. For the first step, "wget" is used to download the source code from its web or ftp site. Next, following the instructions contained in the spell, the program is configured and compiled. Some of the spells, notably the kernel, require user input for further configuration options. Most spells are configured automatically, without needing human input beyond the initial command. The last step is the installation. The spell also contains instructions about where to put the compiled files. Logs documenting the filenames and locations of all these files are written to /var/log/sorcery/ so they can later be reversed, either manually or using the "dispel" command.

There are several cases in which casting of individual packages can become more complicated:

1. **Dependencies:** One spell may depend on another. Fortunately, cast takes care of dependencies automatically. You will be prompted to cast dependencies, which come in two flavors: required and optional. The dependencies you select will be cast before the main spell, so that everything will work right. You can say "no" to a required dependency, but it's not recommended for the faint of heart. Say "yes" to required dependencies unless you know what you're doing!
2. **Post-Install setup:** Sometimes spells require post-install configuration and testing that can't be taken care of by cast. An example is Free Type2 – after casting this spell, you will be reminded that you need to specify the font path of your X server manually, and run ttmkfdir in your True Type font directory. Because there is no agreed-upon standard location for True Type fonts in X-windows, cast cannot do these things for you. It doesn't know where you've decided to install your fonts.
3. **Supported Spells:** Sometimes you will need to recompile a spell that is a dependency for other programs. After the spell is recast, you may be prompted to recast the spells that depend on it in order to make certain they function with the new version. This is important for programs that are statically linked against common libraries, and particularly necessary when taking care of security vulnerabilities. For example, following the discovery of a severe vulnerability in the commonly used library zlib, a sorcery update recompiled zlib, and then prompted to recompile affected programs, thereby ensuring that the vulnerability was eliminated on that box.

5 Cluster Software

We use a job scheduling system in order to share the resources fair between the users. The OpenPBS (see [5]) is such a general job scheduler: let's say you need 5 nodes for 15 hours, you describe that in your job script, then submit it to the scheduler. When the requested number of nodes are available (i.e. not used by other jobs), your job (i.e. script) will be started. You will get also at that point a list of nodes 'dedicated' to your job (no other jobs will be scheduled on those nodes by PBS). If your job doesn't finish

after the ‘time’ you specified when submitting the job, it will be stopped by the PBS. The same applies to memory limits and other resources.

On the cluster now, the PBS is scheduling jobs on nodes 9 and up. The first 8 nodes can be used freely for testing purposes, running small programs, debugging etc., and are not touched by PBS. User are discouraged to run ‘long’ jobs on those nodes (i.e. for more than a few hours).

Clusters are NUMA (Non-Uniform Memory Access) architectures. When we want to run some kind of parallelization, the first question that arises is: How does Information flow in Parallel Programs?

5.1 Parallelization techniques

When fully implemented, ClustIS will offer three different answers to this problem:

1. Message Passing
2. Distributed Shared Memory
3. Distribution instead of Parallelization

5.1.1 Message Passing

The application consists of different process running on the different nodes. In order to exchange information with the other parts, they exchange messages. Two standards are important in this respect:

- **MPI (Message Passing Interface):** proposed as a standard for writing message passing programs.
- **PVM (Parallel Virtual Machine):** library enabling collection of heterogeneous computers to be used as one concurrent computational resource, see [6]

Today, the most common used form for parallel computing on ClustIS is the use of the MPICH implementation of MPI, see [3]. MPICH is heavily used in student assignments in courses on parallel computing.

5.1.2 Distributed Shared Memory

The underlying idea here is to implement a kernel extension distributing applications transparent to the user. This technique is not implemented on ClustIS yet, but we plan to use it least on a part of the cluster because of its ease of use for programmers not used to parallel programming.

We cite the general idea from the openMosix website (see [4]): ‘openMosix is a Linux kernel extension. [...]

Once you have installed openMosix, the nodes in the cluster start talking to one another and the cluster adapts itself to the workload. Processes originating from any one node, if that node is too busy compared to others, can migrate to any other node. openMosix continuously attempts to optimize the resource allocation.

There is no need to program applications specifically for openMosix. Since all openMosix extensions are inside the kernel, every Linux application automatically and transparently benefits from the distributed computing concept of openMosix. The cluster behaves much as does a Symmetric Multi-Processor, but this solution scales to well over a thousand nodes which can themselves be SMPs.'

5.1.3 Distributed Computing

The third answer is not to parallelize a program at all, but to distribute it. Many applications have the need for high computing power, but it is sometimes sufficient to let the same algorithm run on different data sets instead of parallelizing an algorithm. We have an solution developed in-house for this kind of applications.

Q²ADPZ ['kwɔd 'pi: 'si:] is a modular C++ implementation of a free, open source, multi-user, multi-platform system for distributing computing requests in a TCP/IP network. The users of the system can submit, monitor, and control computing tasks (grouped into jobs) to be executed by computers participating in the Q²ADPZ system in form of dynamic shared libraries, executables, or interpreted programs (including Java applications).

Users can provide software, hardware, and platform requirements for each task and the proper computer is automatically selected. The system automatically delivers the input and output data files. Computers executing tasks detect users logging in, and the tasks are terminated or moved to other computers to minimize the disturbance of regular computer users. Q²ADPZ can operate both in conditions of an open Internet environment or of a closed local TCP/IP network.

The internal communication protocol is based on optionally encrypted XML messages. The system provides basic statistics information on usage accounting. Several user modes are supported: from novice users submitting simple binary executable programs to advanced users who can alter the internal communication interfaces for their special needs. We are currently using the system for research tasks in the areas of large scale scientific visualization, evolutionary computation, and simulation of complex neural network models.

5.2 Application Areas

ClustIS is used both in education and research. The Division for Complex Datasystems (KDS) is using it in its undergraduate courses on parallel and distributed computing.

The Division of Intelligent Systems (DIS) developed the setup and is maintaining the cluster primarily for research purposes. In Intelligent Systems research, there are several areas in need for computational power:

- Data Mining
- Bioinformatics (Protein Folding)
- Large Scale Visualization
- Machine Learning (Cross Validation)

- Genetic Algorithms and Genetic Programming
- Artificial Life

An application domain on ClustIS with a huge need for computational power is data mining in real medical data. The task is to identify the role of different proteins in gene expression. A rough set approach to data mining is used with the in-house developed ROSETTA C++ library, a collection of C++ classes and routines that enable discernibility-based empirical modeling and data mining (see [7]).

Another domain is empirical work on machine learning algorithms. We use Weka, [9], a framework for Machine Learning written in Java Implements which implements all common ML algorithm. Its distributed version enables the cross validation step (multiple testing of the performance of algorithms on different learning and test data sets) to be distributed to the cluster nodes. We included our own framework CREEK (a knowledge intensive Case-Based Reasoner, [2]) into the range of algorithms to compare its performance with other approaches.

ClustIS is, amongst others, on a daily basis also used for computation intensive tasks in Genetic Algorithms, Visualization, and simulation of distributed systems.

6 GNU Free Documentation License

GNU Free Documentation License Version 1.2, November 2002

Copyright (C) 2000,2001,2002 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or

(for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

C. State on the Title page the name of the publisher of the Modified Version, as the publisher.

D. Preserve all the copyright notices of the Document.

E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.

G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.

H. Include an unaltered copy of this License.

I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.

L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.

M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.

N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.

O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and

disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (c) YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

References

- [1] <http://clustis.idi.ntnu.no/>. Website, May 2003.
- [2] <http://www.grupper.ai/cbr/>. Website, May 2003.
- [3] <http://www-unix.mcs.anl.gov/mpi/mpich/>. Website, May 2003.
- [4] <http://www.openmosix.org/>. Website, May 2003.
- [5] <http://www.openpbs.org/>. Website, May 2003.
- [6] <http://www.epm.ornl.gov/pvm/>. Website, May 2003.
- [7] <http://rosetta.sourceforge.net/>. Website, May 2003.
- [8] <http://www.top500.org/>. Website, May 2003.
- [9] <http://www.cs.waikato.ac.nz/ml/weka/>. Website, May 2003.