



Porting Intel Applications to 32 bit Linux on POWER

Matthew Davis

IBM eServer Solutions Enablement

Special Notices

This publication/presentation was produced in the United States. IBM may not offer the products, programs, services or features discussed herein in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products, programs, services, and features available in your area. Any reference to an IBM product, program, service, or feature is not intended to state or imply that only IBM's product, program, service, or feature may be used. Any functionally equivalent product, program, service, or feature that does not infringe on IBM's intellectual property rights may be used instead.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. Sources for non-IBM list prices and performance numbers are taken from publicly available information including D.H. Brown, vendor announcements, vendor WWW Home Pages, SPEC Home Page, GPC (Graphics Processing Council) Home Page and TPC (Transaction Processing Performance Council) Home Page. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

Questions on the capabilities of non-IBM products should be addressed to suppliers of those products. IBM may have patents or pending patent applications covering subject matter in this presentation. Furnishing this presentation does not give you any license to these patents. Send license inquiries, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA. All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of a specific Statement of General Direction.

The information contained in this presentation has not been submitted to any formal IBM test and is distributed "AS IS." While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. The use of this information or the implementation of any techniques described herein is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The information contained in this document represents the current views of IBM on the issues discussed as of the date of publication. IBM cannot guarantee the accuracy of any information presented after the date of publication. The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: AIX, AIX 5L, AIX/6000, IBM, RS/6000, VisualAge, e-business (logo), POWER2 Architecture, PowerPC (logo), PowerPC 604, pSeries, SP, iSeries, OS/400, AS/400, POWER3, POWER4, RS64IV, POWER. A full list of U.S. trademarks owned by IBM may be found at <http://ibm.com/legal/copy/trade.html>. UNIX is a registered trademark of The Open Group. Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States and other countries. Lotus, Lotus Domino and Lotus Notes are trademarks or registered trademarks of Lotus Development Corporation. Tivoli, TME, TME 10 and TME 10 Global Enterprise Manager are trademarks or registered trademarks of Tivoli Systems, Inc. Other company, product and service names, which may be denoted by a double asterisk (**), may be trademarks or service marks of others. Linux is a registered trademark of Linus Torvalds. HP-UX and Tru64 are trademarks of HPQ in the United States and other countries. Solaris is a registered trademark of Sun Microsystems in the United States and other countries.

© Copyright IBM Corporation 2003. All rights reserved

Table of Contents

Helpful Terms and Definitions.....	3
POWER TM	3
Linux for POWER.....	3
GNU Compiler Collection (GCC).....	3
SUSE Linux Enterprise Server 8 (SLES 8).....	3
Red Hat Enterprise Linux 3 (RHEL 3).....	4
VisualAge Compiler Set (VA Compiler).....	4
IA32.....	4
GPL.....	4
Introduction.....	5
About POWER Architecture.....	5
About SLES 8 for POWER.....	5
About RHEL 3 for POWER.....	5
Compiler Sets and Architecture Awareness.....	6
Porting with the GCC Compiler Set.....	6
Porting with the VisualAge Compiler Set.....	6
Java.....	7
Summary.....	8
Additional Information.....	8
Where to Find Compilers.....	8
Community Resources.....	8
§About the author.....	8

Helpful Terms and Definitions

POWER TM

The 64-bit architecture RISC featured in many enterprise IBM servers.

Linux for POWER

The Linux operating system running on POWER hardware. It can be assumed in the context of this document that this refers explicitly to IBM iSeries and pSeries servers.

GNU Compiler Collection (GCC)

The open source GNU Compiler Set, including gcc, the GNU C Compiler, g++, the GNU C++ compiler, and g77, the GNU Fortran77 compiler.

SUSE Linux Enterprise Server 8 (SLES 8)

SUSE Linux Enterprise Server, version 8. This is the most recent enterprise Linux offering from SUSE Linux. In the context of this document, SLES 8 refers to SLES 8 for POWER architecture.

Red Hat Enterprise Linux 3 (RHEL 3)

Red Hat Enterprise Linux, version 3. This is the most recent enterprise Linux offering from Red Hat Linux. In the context of this document, RHEL 3 refers to RHEL 3 for POWER architecture.

VisualAge Compiler Set (VA Compiler)

This compiler set, functionally identical to IBM's high-performance compiler set for AIX, is now available for Linux on POWER. The VA compiler set includes xLC, the VA C compiler, xLC++, the VA C++ compiler, and xLF, the VA Fortran compiler.

IA32

The 32 bit Intel architecture.

GPL

The GNU General Public License, under which the Linux kernel and much of the software found in the open source community is licensed. More about the GPL can be found at <http://www.gnu.org>.

Introduction

The path to enablement of IA32 code on Linux on POWER travels through one of two compiler sets. The GNU Compiler Collection (GCC) is the Free Software Foundation's open source compiler set; this is the de facto standard for software compilation on Linux. GCC's strength rests in its standardization across several types of hardware, and therefore the robust portability of code written to be compiled with GCC compilers. Additionally, IBM has made available the VisualAge Compilers used traditionally on AIX to develop and deploy applications for Linux on POWER. This document reviews some tips for porting 32-bit software from Linux IA32 to Linux on POWER with both of these compiler sets.

About POWER Architecture

Performance Optimized with Enhanced RISC (POWER) architecture is the result of more than a decade of hardware engineering dating to the POWER1 chips released in 1990. The POWER1 was IBM's first commercial RISC processor and debuted at 25 Mhz clock speed. POWER4, the current IBM RISC processor offering operates at speeds between 1.1 and 1.45 Ghz. With the extensive parallelism and optimization in the hardware of the current implementations, POWER can deliver more computation per cycle than other chip architectures. Announced in 2001, POWER4 chips feature two 64 bit microprocessors per CPU board and support SMP configurations up to 32 way. In addition to the POWER4 hardware, Linux on POWER also runs on IBM eServer hardware featuring POWER3 and RS64IV hardware. These architectures both feature 64 bit SMP capability and copper fabrication technology, but whereas the RS64IV processors are still actively developed, POWER3 is considered the legacy predecessor to POWER4. Each of the PowerPC derived chips offers industry leading performance and provides a world class hardware platform for the deployment of Linux PowerPC.

About SLES 8 for POWER

SuSE Linux, available for more than 10 years, offers two levels of Linux distribution. SuSE Linux 8.0 is intended for workstation environments, and SuSE Linux Enterprise Server version 8 for enterprise systems. SLES 8 for POWER is the product of more than three years of collaboration between SuSE, IBM, and the open source community. Though SLES 8 boasts a number of accomplishments regarding 64 bit enablement of the Linux operating system, the ability of Linux on POWER to run both 32 and 64 bit applications simultaneously does not prohibit 32 bit applications from taking advantage of this mature Linux distribution. Combined with IBM's VA Compilers for performance critical applications, Linux on POWER environment offered by SLES 8 on IBM iSeries and pSeries hardware is an unparalleled platform for developing and deploying Linux solutions. For more information on SuSE Linux, see <http://www.suse.com/us/company/index.html>

About RHEL 3 for POWER

Red Hat Linux is widely recognized as an industry leader in both desktop and enterprise class Linux markets. Compiler and threading implementations in RHEL 3 differ from those in SLES 8. Native Posix Threads for Linux (NPTL), featured in RHEL 3, provide better vertical scaling and general performance improvements over the traditional Linux threading model.

Compiler Sets and Architecture Awareness

Unlike other Linux platforms, Linux PowerPC offers not only GCC, but also the IBM VA Compiler set. GCC offers robust portability of code intended for compilation in Linux, while the VA Compilers offer a substantial performance increase over GCC.

Both AIX and Linux on the POWER architecture are run as Big Endian (left to right assignment of addresses to bytes in a data word). As most developers familiar with cross-platform develop will attest, it is of utmost importance to consider endian friendly coding habits when deploying across Little Endian and Big Endian platforms. Much has been written on this topic, and for an excellent review of addressing endian concerns, please see chapter 2 of the [AIX 5L Porting Guide](#). Another practical and concise reference, [Migrating Tru64 Applications to AIX](#). These references may prove particularly useful when migrating code with the VA Compilers.

Porting with the GCC Compiler Set

In general, porting code with GCC compilers should be straightforward. Architecture may vary, and occasionally library version discrepancies may arise, but for the most part, GCC is GCC, no matter which architecture it runs on. Architecture specific flags are discouraged for compilation on Linux on POWER (e.g. `-m486`, `-powerpc`), as GCC does not have extensive processor maps for optimization routines on these architectures. Also, without architecture specific flags, binary compatibility will be greater across different models of POWER hardware.

On all architectures, libraries must be compiled with `-fPIC`; IA32 has this flag applied by default with GCC. On POWER, this flag specifies that the generated code will be used in a shared object.

Developers not familiar with the GPL should reference it carefully and make sure the terms and conditions of the license are understood.

Porting with the VisualAge Compiler Set

Compiling code with the VA compiler set has shown a significant performance increase on Linux for IBM eServers featuring POWER hardware. These compilers, migrated from the high performance compiler set used on AIX, have extensive optimization routines for POWER hardware, and therefore should be used for any application that wants to take full advantage of the POWER architecture's potential. However, software originally intended to be compiled with GCC compilers may need more attention when recompiled

with VA compilers. Makefiles need to be edited to reflect the correct path to the VA compilers, by default `/opt/ibmcmp/`. The correct optimization flags should also be set for the architecture (e.g. `-O3 -qarch=pwr4 -qtune=pwr4` for IBM POWER4 hardware; a complete list of these flags can be found by executing the compiler binary with no arguments). In addition to optimization modes for various POWER architecture derivatives, the VisualAge compilers can be instructed to compile software in common mode. This guarantees compatibility across all types of POWER architecture at the expense of performance otherwise gained through architecture-specific optimization.

The VisualAge Compilers do not create object code directly. First, `.w` code is generated, which is fed back to the compiler's backend optimizer. Once the most efficient compilation method is determined, the compiler generates `.o` files, and then the final binaries are generated.

A few tips to compiling GCC oriented code with the VisualAge compiler set are listed below:

C++ comments: GCC allows C++ style commenting to be used in C files by default, but this is not the case with XLC, the VisualAge C compiler. Since it is not economical to change all the comments in a group of source files to comply with C style commenting, XLC provides a `-q` argument to allow these comments, `-q cplusplusmt`. When C code is compiled with this flag, both C and C++ style commenting is interpreted.

Environment variables are often the easiest way to configure build scripts. As opposed to hand editing Configure scripts and Makefiles, it is recommend that relevant environment variables such as `$CC` and `$CFLAGS` be set and the Configure script be allowed to generate the Makefile without the need for hand editing.

Configure scripts also need to be aware of the proper platform type. This will be either `linux-powerpc-unknown-gnu` or `linux-powerpc64-unknown-gnu`. This should be set for compilation with either GCC or VA by appending the `-target=` flag to the Configure script.

Although extensive documentation for the VisualAge compilers exists, an argument list is provided at the console by running the compiler with no arguments, e.g.
`$COMPILER_PATH/bin/cc.`

Java

In previous distributions of Linux on POWER, the Blackdown JDK was the only available development kit. Blackdown, while renowned for offering a highly compatible JDK, has not boasted the performance the [IBM JDK](#) has delivered to Java applications on Linux Intel and AIX platforms. Yet, thanks to the joint effort of IBM, SuSE, and the open source community, Java migration to Linux on POWER is now powered by the IBM JDK version 1.4.1 running on either SLES 8 or RHEL 3. The IBM JDK offers the formal compliance to Java Standards and enterprise class performance expected from

IBM development kits. Additional performance gains may be found by recompiling requisite C libraries with the VA Compilers.

Developers not accustomed to developing with the IBM JDK should reference <http://ibm.com/developerworks/java>. This may be helpful for understanding optimal levels of garbage collection, JIT compiling, and other Java features.

Summary

Linux on POWER offers an enterprise class Linux environment, complete with 64 bit kernel, application environment, and toolchains. Unlike other Linux environments, Linux on POWER offers a twin compiler set representing the best of both worlds: ease of migration with open source code and high performance exploitation of the award winning high performance POWER architecture. As a result of more than three years of cooperation between IBM, SuSE, and the open source community, SLES 8 has emerged as a genuine enterprise class distribution and a leader in its field. RHEL 3 is now available as a Linux on POWER distribution from long-time Linux leader, Red Hat. Application performance available on the POWER architecture is now augmented by development tools never before offered on the Linux operating system. In summary, Linux on POWER is the leading platform for the deployment of high performance Linux applications.

Additional Information

Where to Find Compilers

The best source for GCC compilers for Linux on POWER is from a certified IBM Linux vendor, such as SuSE Linux, as these vendors have taken efforts to build GCC and its libraries to minimize any version discrepancies.

VisualAge compilers can be found at <http://ibm.com/software/awdtools/vacpp/>.

Community Resources

<http://gcc.gnu.org>

<http://linuxppc64.org>

<http://ibm.com/developerworks/linux/>

§About the author

Matt Davis is a Linux technical consultant in the IBM eServer Solutions Enablement team. As a member of the Linux on POWER project since its inception, he explored and tested emerging technology for Linux on POWER and wrote several reports summarizing his findings. These include Journaling File Systems for Linux PowerPC, Parallel Grid Computing with Linux PowerPC, Open Source Alternatives to Commercial Software for Linux PowerPC, and the Linux Solutions Catalog, to name a few. He came to IBM as an

intern during his tenure as a student at the University of Texas at Austin, from which he earned two degrees with academic honors.