



Gelato Gazette

Advancing Linux Itanium



FREE

JANUARY 31, 2005

VOLUME 1, ISSUE 1

OSDL ENTERPRISE LINUX SUMMIT 2005 SPECIAL EDITION

From the Pipeline

Itanium Build System for Eclipse

Gelato has submitted patches to the Eclipse build system and written detailed instructions to facilitate compilation of the IA64 Eclipse native binaries. The Eclipse Foundation is in the process of integrating our changes and will be using Gelato's regular builds for their IA64 releases.

Kevin Cernekee cernekee@gelato.org

Ubuntu: IA64 Beta Version

The Ubuntu IA64 port team recently announced ISOs are available for testing. Ubuntu is a newly-created Linux distribution, which has received a great amount of enthusiasm from the Linux community.

Thierry Simonnet tsimonnet@yahoo.com

See PIPELINE page 2

Gelato Gazette

Published periodically by the Gelato Federation for the Linux Itanium user community. The content of this publication is taken from information available on the Gelato website.

Gelato Website: www.gelato.org
General Inquiries: info@gelato.org

The Scheduler

May 2-6, Grid Asia 2005, Biopolis, Singapore

On the afternoon of May 4, Wilson Rivera (UPRM) and Jon Lau (NGO) will co-host a Gelato grid Birds of a Feather (BOF) presentation.

May 23-25, Gelato Federation Meeting, Cupertino, California, USA

See article on this page.

June 21-24, ISC2005, Heidelberg, Germany

The International Supercomputer Conference is the leading supercomputing event in Europe.

October 3-5, Gelato Federation Meeting, Porto Alegre, Brazil

See article on this page.

International Linux Itanium Meetings

The Gelato Federation has set its second quarter meeting for May 23-25, 2005 in Cupertino, California, USA. Hewlett-Packard, the Federation's founding sponsor, will be hosting this event. Presentations will focus on:

- Details of the IA64 architecture and the upcoming release of Montecito
- Demonstrations of compiler optimizations and performance analysis tools leading to minimized latencies
- Scalability solutions

Twice a year, Gelato holds technically rich, educational, international meetings focused on the Linux Itanium platform. Locations in the past have ranged from Sydney, Australia to Beijing, China to Champaign, Illinois. In October, the Catholic University of Rio Grande do Sul will host a meeting in Porto Alegre, Brazil. To view content from previous meetings or to learn more about the Gelato Federation, visit www.gelato.org.

Gelato Federation meetings are open to current members and invited guests. If you are interested in attending and are not a current member, please contact Gelato's Managing Director, Mark K. Smith, at mksmith@gelato.org.

Compiler-Driven Performance on Itanium

The continuing desire for greater computational power and the current complexities of the Pentium and Xeon processors beg the question: How far can x86 be pushed? Point of fact, the current Pentium processor uses a RISC core with an x86 front-end. The Itanium processor is a long-term solution to this problem developed by Intel and Hewlett-Packard. Designed for high performance and continued scalability, the Itanium is revolutionary. Its architecture depends on the compiler to create highly-parallelized code while hardware support in the form of predication, speculation, and rotating registers make the compiler's job easier.

Outstanding Itanium performance is possible but only through an informed understanding of the Itanium architecture and careful selection of a compiler. To accomplish the goal of highly-parallelized code, the compiler must be completely aware of the Itanium architecture and effectively utilize all available hardware support. Given that GCC does not currently exploit any Itanium

hardware support, it is no surprise that experimental evidence clearly indicates that the performance compiler of choice for the Linux/Itanium platform is not GCC, but rather Intel's proprietary compiler (icc).

Since icc is proprietary, what options are available for someone interested in producing freely distributable binaries in the open-source community? First, use an open-source, high-performance compiler designed for the Linux Itanium platform, such as OpenIMPACT available at <http://www.gelato.uiuc.edu>. However, open-source compilers are generally not as well tested as icc and offer less support. Second, encourage changes in GCC to bring performance on the Itanium to a more acceptable level. Given that the majority of Itaniums used today are running Linux and that GCC is the standard choice for a compiler, it makes sense to consider this alternative. To this end, Gelato and Hewlett-Packard recently held a workshop with leading international compiler experts for the purpose of considering improvements to GCC critical to the Itanium.

Optimized Binary Depot for Linux on Itanium

The Gelato team is serving up Vanilla. Not a flavor of ice cream but a project that will create highly-optimized binaries for critical Linux utilities on the Linux Itanium platform and will give developers compiler tips to create efficient binaries. The Vanilla Project is composed of three stages. First, determine compiler flags and recorded speedup. Second, isolate critical areas in the code and hand tweak the code for even greater performance. Third, produce LSB-compliant binaries able to run any LSB-compliant or certified Linux distribution. Gelato is currently in the first stage, which includes careful documentation of the compiler options used. The team anticipates adding four "compiler recipes" a month.

What Is Gelato?

Gelato is the global research community dedicated to advancing the Linux Itanium platform through collaborative relationships targeting real-world problems and solutions. Through an open exchange of ideas and results, Gelato brings together those developing and using Linux on Itanium for the purpose of fostering Itanium-specific research, focus groups, and collaborative discussions. Participation is invited from all interested organizations and individuals.

CHECK US OUT ONLINE @ WWW.GELATO.ORG

Gelato—Three Years and Growing Strong

In January 2002, Hewlett-Packard together with a group of research organizations from around the world agreed to create a federation, called Gelato, dedicated to enabling scalable, open-source commodity computing solutions for Linux on Itanium to address real-world problems in academic, government, and industrial research.

Within these three years, the Gelato Federation has attracted two new sponsors—Intel Corporation and SGI—and has grown to thirty-nine members. The Federation delivers lasting value through the careful balance of short-term goals (meeting current real-world problems) and long-term goals (cutting edge research). Essential to Gelato is the notion of an open exchange of ideas and solutions, which is expressed in part by Gelato's commitment to facilitating the development of open-source software for the high performance computing community.

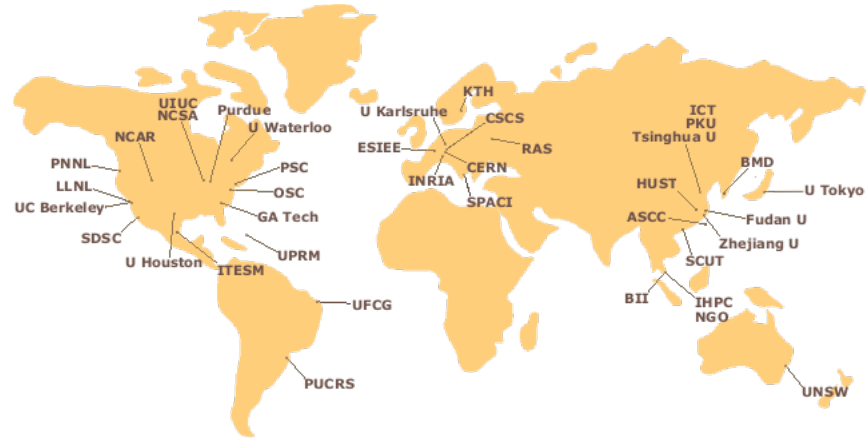
With members from almost every continent on the globe, the work of the Federation centers around seven main focus areas: scalability, compilers, parallel filesystems, performance, cluster and grid computing, Linux Itanium standards, and Linux Itanium tool integration. Within these areas, Gelato Members are applying new technology and solutions to specific applications in bioinformatics, high-energy physics, and atmospheric sciences. Here is a list of current projects and their member institution.

Compilers

- OpenIMPACT Compiler (UIUC)
- ORC-OpenMP—ORC as an OpenMP Compiler (Tsinghua U)
- SWP—Software Pipeline (Tsinghua U)
- C++ Compiler Optimization (RAS)

Parallel Filesystems

- IA64 Lustre (PNNL)
- Clusterfile Parallel Filesystem (U Karlsruhe)
- Parallel I/O Benchmarking (OSC)



Scalability and Performance in a Cluster

- OSCAR Linux Clustering Solution (NCSA)
- Buster—Parallel Debugger for Cluster (Tsinghua U)
- ChaRM—Checkpoint and Rollback Recovery (Tsinghua U)
- Linux Clustering Solution (INRIA)
- Coordinated Process Scheduling (U Karlsruhe)
- JavaParty (U Karlsruhe)
- NPACI Rocks Cluster Distribution (SDSC)
- Development of Cluster Web Interface (RAS)
- Clconfig—Cluster Description Language (OSC)
- Performance Tools—PAPI and I2prof (KTH)
- CPAD—Resource Management in Clusters (PUCRS)

Scalability and Performance in a Box

- μ C++—High-Performance User-Level Concurrency (U Waterloo)
- Understanding, Improving, and Implementing High-Performance Operating System Interfaces (U Waterloo)
- New Virtual Memory Infrastructure (UNSW)
- User Level Device Drivers (UNSW)
- Microstate Accounting (UNSW)
- Preemption Support (UNSW)
- Large Block Devices (UNSW)

- Spectral Toolkit (NCAR)
- NCL—NCAR Command Language (NCAR)
- ESMF—Earth System Modeling Framework (NCAR)
- Scalable Servers (U Karlsruhe)
- PeSO—OS Scalability (PUCRS)
- Crystallographic FFTs Library (UPRM)

Grid Computing

- CERNLIB for CERN's FORTRAN Applications (CERN)
- CLHEP—CLASS Library for High-Energy Physics (CERN)
- European DataGRID Software (CERN)
- GEANT4—Physics Simulation Toolkit (CERN)
- ROOT—OO Framework (CERN)
- Silver Integration with Globus Security Infrastructure (GSI) (OSC)
- Grid Computing—GraDSOft and an FFT Library (KTH)
- Grid Performance Modeling (U Houston)
- OurGrid (UFCG)

Gelato offers participants the opportunity to take a leading role in the worldwide development and promulgation of Linux on Itanium, the next generation of high-performance computing. For more information, visit www.gelato.org.

From the Pipeline

PIPELINE from Front

uC++ now a part of Debian GNU/Linux

The University of Waterloo's concurrent programming language uC++ is now available pre-packaged to users of Debian GNU/Linux's "testing" distribution (code-named "sarge"). The package includes the u++ command to compile uC++ programs, and can be installed on both the ia64 and i386 architectures by running `apt-get install u++`. Also, the uC++ Annotated Reference Manual is available in the u++-doc package.

uC++ (pronounced: micro-see-plus-plus) is an extension of the C++ programming language that supports concurrent and real-time programming, as well as advanced control flow techniques. More information on uC++ is available from Waterloo's Gelato portal, gelato.uwaterloo.ca.

MySQL Performance on Itanium 2

DIKU, the Computer Science Department of the University of Copenhagen, recently completed a tech report describing the impact of the Intel compiler on the performance of MySQL on Itanium 2. This report can be found on the Gelato website at http://www.gelato.org/pdf/mysql_itanium2_perf.pdf.

This work is part of the ongoing Badger Project. Current activities include main memory performance of MySQL on Itanium 2, a profiling framework for MySQL on Itanium 2, and a study of MySQL cluster database. More information can be found at www.distlab.dk/badger/.

Updated libm Now in CVS Tree of glibc

Thanks to the help of HJ Lu and Jakub Jelinek, the updated libm has made it into the CVS tree of glibc. This libm is based on Intel's libm v2.1 and greatly improves both speed and accuracy of many routines. If you'd like

to see this libm in your favorite distro, you may want to let your distributor know, since it's normally takes some time for distros to update their libc.

Free Software for Linux Open Source Developers

Intel has expanded its offerings of free Linux tools for non-commercial software development. This offering is provided to developers who are developing software on their own time without compensation. Visit <http://developer.intel.com/software/products/noncom/> for details.

Gelato Grid Focus Group

Gelato members and others interested in grid computing are suggested to subscribe the Gelato Grid mailing list. General information about the mailing list is at: <https://www.gelato.unsw.edu.au/mailman/listinfo/gelato-grid>. To post to this list, send your comments via email to: gelato-grid@gelato.unsw.edu.au.