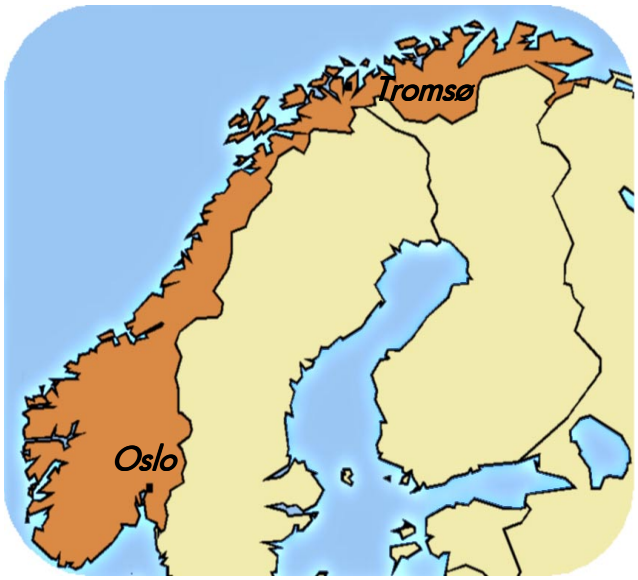

Itanium-powered performance computing in Norway with the universities of Tromsø and Oslo

success story

Overview

Intel® Itanium™ processor-powered Hewlett-Packard systems are being groomed to become a critical part of a powerful high performance computing (HPC) grid being built in Norway. A number of universities in Norway are pooling their resources to offer an HPC network, which will be used for both academic, and commercial research.

An early focus on the grid will be in bioinformatics – the collection, analysis and referencing of genetic data. As it matures it will also be employed in research across a wide number of other disciplines including chemistry, astrophysics, geophysics, financial simulations, and mathematics.

A grid is a distributed computing network that pulls together a wide variety of different and geographically dispersed computing resources to tackle numerous and large-scale compute intensive tasks.

The universities involved are:

- University of Oslo
- University of Tromsø
- University of Bergen

Norwegian University of Science and Technology (NTNU) in Trondheim

Challenge:

The Universities of Oslo and Tromsø are playing a critical part in an advanced High Performance Computing grid. They must ensure:

- **Outstanding price performance with the fastest throughput** – running compute intensive tasks for the academic and commercial community.
- **Freedom of choice** – Users can simply pick the platform that offers the best resources for the job while the four universities computing resources appear as homogenous as possible

Early adopters have big Itanium plans

Out of the four, the University of Oslo and the University of Tromsø have big plans for Itanium. Supercomputing specialists from these universities believe that Itanium may be the ideal match for academic supercomputing, and so they have teamed up with HP to be the first to incorporate Itanium processor-powered systems into the network.

For the good of the grid
Intel Itanium architecture is ideal for scientific computing in heterogeneous environments. Itanium offers:

- The flexibility of clustering that can handle diverse operating environments
- Outstanding performance and scalability
- The power of 64-bit computing in operating environments that traditionally have been 32-bit, like Linux and Windows
    - HP supports Windows, Linux, and HP-UX on the IPF architecture
- The ability to consolidate on a single architecture:
    - The Superdomes can be upgraded to a future generation of the IPF, the workstation clusters run on Itanium today, and future systems will run next generation McKinley processors

The University of Oslo, which is Norway's largest and oldest university, has generated no less than four Nobel Prize winners. It has a Biotechnology Center that does genetic and proteomic research that is currently playing a key role in the sequencing of a salmon genome, one of the higher animals to be sequenced. Its Faculty of Medicine runs its own neurobiology and cancer research programs. The University's Faculty of Math and Sciences is the home of the Institute of Theoretical Astrophysics where world-leading research into stellar and plasma physics takes place.

The University of Tromsø is the world's most Northern university, and many of its scientific research programs reflect its unique geographic position. Whether under the summer midnight sun or through long and dark winter days, the Faculty of Science is hard at work gathering huge amounts of data on Arctic research, biotechnology, and environmental issues, as well as auroral, solar, and high-energy physics.

Today, both Tromsø and Oslo rely on powerful consolidated PA-RISC server solutions for their performance computing on the grid, including two Superdomes running the HP-UX operating system.

However, while consolidation is important for the largest jobs, both universities are also interested in learning more about the power of clustered workstations They each have a cluster of HP workstation i2000 Itanium systems running Linux.

As early adopters of leading-edge technology, both Oslo and Tromsø's learning process with Itanium today will help ensure readiness for the leading performance and economical prices of Intel's next generation of Itanium processor, codenamed McKinley.

Clustering configurations for Itanium

Currently, HP and the University of Tromsø are testing a number of different clustering management applications. Tromsø is particularly interested in the National Partnership for Advanced Computational Infrastructure's ROCKs cluster management tool, which allows easy development of commodity off-the-shelf clusters.

Oslo is also testing Linux clustering software on their Itanium processor systems. In this case the SCore clustering software, which is being ported to the Itanium Processor Family by a Japanese university. Both Oslo and Tromsø also have HP-UX HP workstation i2000 systems, as they find that the HP-UX compilers provide exceptional performance.

The universities are discovering that Linux has the strongest clustering capacities and HP-UX has the strongest performance. As the grid develops, they will be able to consolidate on a single architecture while still using both operating systems by using Interoperable Message Passing Interface protocols implemented by the NOTUR consortium. The diverse operating environments on the grid such as HP-UX and Linux will be united.

Tromsø's HP Superdome (left) and HP workstation i2000 cluster

For example, Oslo plans to use the workload management tool PBS Pro, which establishes a common queuing environment across diverse HPC systems and architectures, on its Superdome and N-Class servers. The PBS Pro interface running on the servers will submit jobs to the SCore cluster.

Tromsø plan to use a meta-scheduler called Silver on top of PBS Pro to manage jobs on both the Superdome and the Itanium processor cluster from a single batch system as the cluster matures. They plan to upgrade their cluster to one based on the next-generation Itanium processor codenamed McKinley next year.

Tor Johansen, HPC Project Manager for the University of Tromsø, says: "We expect the Itanium processor family to dominate the HPC community for years to come and

we expect that our future systems will contain Itanium processors, whether they are clusters or consolidated servers. This is why we want to learn about the IPF as early as possible.

"The combination of having early access to the latest members of the Itanium processor family, and the direct access we have to key personnel at HP with the experience and knowledge in developing tools and products based on the Itanium processor specification is crucial to a group such as ours."

Flexible operating system choice

Lars Oftedal, Assistant Director at the Center for Information Technology Services, University of Oslo, says: "We find the flexibility of HP's operating system strategy very appealing. We have HP-UX at the university already, and we are very pleased with it. But we also have a large and very strong Linux community at the University of Oslo, and it would be very interesting to see what opportunities are created with 64-bit Linux and large memory systems.

"We are also looking at running 64-bit applications on Windows XP; Itanium gives us the opportunity to test their performance. Our scientists are also running some very large statistical problems on the statistical and graphing software SPSS, and we are very interested in seeing how these run on a parallel architecture… We will use the OS that is best suited for the application or where the most appropriate tools are available for our scientists who are developing their own applications."

"I believe that McKinley will be a very important technology to us – early indications are that it will have even greater performance at good prices. The HP roadmap is very impressive"

Arne Laukholm
Director
Center for Information Technology Services
University of Oslo
http://www.uio.no

"We have a lot of scientific applications that run on Windows 2000 and NT today that need a lot of CPU and I/O performance as well as lots of memory, such as SAS, SPSSx, S-plus, MatLab, Maple, and Origin. We believe that some of these applications will be a lot faster when running large data sets."

Arne Laukholm, Director for the University of Oslo's Center for Information Technology Services, says: "As we are funded with public money we must be very careful – value for money is absolutely critical to a university, and we cannot risk making the wrong choice on an architecture. This makes the Itanium processor family very interesting to us for two reasons.

"First, these processors will be produced in much greater volumes than other 64-bit processors, so price performance should be excellent because of the economies of scale.

"The second reason is that Itanium is very widely supported by the HPC community. We believe that the Itanium processor family will be very successful, so we are starting early to deploy it. We have been investing in HPC clustering for quite a while and Itanium processor technology will fit quite well into the clustering model.

Laukholm continues: "We are starting with Itanium now to ensure that we have the applications and clustering management software needed before the next generation of Itanium processor becomes available. I believe that McKinley will be a very important technology for us – early indications are that it will have even greater performance at good prices. The HP roadmap is very impressive."

The grid network is being developed by an alliance known as NOTUR – NO is an abbreviation of NOrwegian, TUR, is a contraction of TUngRegning, which means "supercomputing".

NOTUR, headed by NTNU, is made up of members of Norway's High Performance Computing Consortium, including all four universities, and is partially funded by a government body, the Norwegian Research Council. University students and research departments, private enterprise, and hospitals will all have access to NOTUR's grid. The developing grid already serves the operational weather forecast for Norway, HPC users at the four universities, and the Norwegian company STATOIL.

The expectation is to introduce the Itanium processor clusters as official member of the developing grid in early 2002 and will be available to the entire Norwegian research community for experimental usage.

Professor Bjørnar Pettersen, the leader of the NOTUR consortium, said: "The collaboration between HP and the Universities of Oslo and Tromsø has been successful right from the start. From NOTUR's perspective, this should serve as an example to the other sites and vendors. This collaboration is helping NOTUR stay up-to-date with the latest technology developments."

Applying research in life sciences industry

One enterprise user of the growing grid will be Sencel Bioinformatics AS. Sencel work closely with the University of Oslo. In fact, several Sencel employees also work at the University and much of their R&D is done though the University.

"I expect a performance of the 800MHz Itanium in the range of 600 to 800 million symbol comparisons per second with the Smith-Waterman algorithm. Thus, even the first generation processor in this family will probably be faster for this application than any other processor."

Torbjørn Rognes,
Chief Technology Officer
Sencel Bioinformatics AS
www.sencel.com

Sencel's bioinformatics software makes short work of searching genomic databases to find matched proteins, despite the computationally challenging nature of the task.

Torbjørn Rognes, Chief Technology Officer at Sencel, says: "The Itanium processor really seems to be ideal for this application and I think it will show extreme performance."

"The Itanium's two integer issue ports and two memory/integer Arithmetic Logic Units issue ports enable it to perform up to four operations (in addition to branches and floating-point operations) in each cycle," explains Rognes.

"My preliminary code sketches indicate that it is possible to keep these four issue ports busy nearly all the time. In addition, the instruction set has all the multimedia instructions that are needed, and more. Together, this makes it possible to write very efficient code."

Sencel Bioinformatics AS

The engineers at Sencel are experts at hand-coding optimized assembly language for specific processor architectures, which have been developed to exploit them at a fundamental level.

The Sencel code uses the processors' multimedia instructions to execute their proprietary ParAlign sequence comparison algorithms, 4-8 times faster than the Smith Waterman algorithm.

These compute intensive algorithms are the most effective and sensitive way to locate related proteins in a database.

Sencel scientists have studied the Itanium processor in detail and sketched some preliminary code for the core, and are very anxious to get their hands on a system.

"I expect a performance of the 800MHz Itanium in the range of 600 to 800 million symbol comparisons per second

with the Smith-Waterman algorithm. Thus, even the first generation processor in this family will probably be faster for this application than any other processor. With the 1GHz McKinley processor I expect at least an additional 50% boost due to the additional issue ports, increased clock frequency and other improvements."

Sencel plan to do their software development for Itanium on the Itanium workstations at the University of Oslo in collaboration with NOTUR and HP. They also hope to use the Itanium processor cluster at Tromsø as part of the development process.

Itanium takes HPC to the next level

The Itanium processor family has a lot to offer the HPC market:

- It is founded on an entirely new 64-bit architecture that emphasizes parallel performance, critical for today's compute-intensive scientific applications and data sets. It combines 64-bit power with the ability to deploy terabytes of memory and address space.
- It has the flexibility to be able to run a choice of operating systems, including Windows XP, Linux, and HP-UX.
- It offers leading floating point performance makes it a perfect match for scientific computing.
- It has backwards compatibility with 32 and 64-bit PA-RISC applications, and 32-bit Windows applications.
- It is equally at home in a large consolidated server as it is in a cluster, bringing new levels of 64-bit compute power to shared memory systems.

As the NOTUR grid is built, the Itanium systems will be put through their paces - Oslo and Tromsø alone easily have the potential to stretch the capabilities of the Itanium-powered grid.

For more information:

- University of Tromsø HPC: http://www.hpc.uit.no/
- University of Oslo HPC: http://www.hpc.uio.no/
- Norwegian HPC Consortium (NOTUR) – http://www.notur.org
- Sencel Bioinformatics AS: http://www.sencel.com
- HP and the Itanium processor family: http://www.hp.com/go/itanium