



## ac3: Triumph in Teraflops—ac3 Breaks Australia's Teraflops Barrier

### SOLUTION SUMMARY

Challenge	The Australian Centre for Advanced Computing and Communications (ac3) built an Intel® architecture-based supercomputing cluster with the goal of obtaining the best possible performance from the hardware. When the center's staff realized how powerful the cluster was, they became determined to develop a high-performing supercomputer that would place in the TOP500 list of the world's most powerful computers.
Solution	Intel® Solution Services helped ac3 design the large cluster as well as tune the software to yield the maximum performance. The result was a performance of 1.095 teraflops on the LINPACK benchmark and the rank of 108 among the TOP500 (November 2003) supercomputers.
Cluster configuration	A cluster of 154 Dell PowerEdge* 1750 servers with dual Intel® Xeon™ processors, each with up to 1MB L2 cache: 152 compute nodes Two head nodes One file server connected to a Dell/EMC CX200 Fibre Channel storage system
Cluster interconnect	Gigabit Ethernet over a Foundry Networks Fastron* 1500 switch
Operating system	Red Hat Linux* 9 with the 2.4.20 kernel, Red Hat Enterprise Linux 2.1 for the file server
Programming environment	OSCAR 2.3 cluster software for Linux, Intel® C++ Compiler 7.1 for Linux, Intel® Math Kernel Library (MKL 6.1), Intel® Fortran Compiler 7.1 for Linux
Intel Solution Services provided	High-Performance Computing - Application platform

### Business Challenge

The Australian Centre for Advanced Computing and Communications (ac3) has big ambitions. Owned by the government of New South Wales (NSW), Australia, and eight NSW universities, it is creating a world-class, high-performance computing (HPC) center that will help promote economic and employment growth in Australia by providing computing resources and services to industry, commerce, government and academia.

In early 2003, a consortium of five NSW universities received funding from the Australian Research Council to purchase a supercomputer that would allow researchers to pursue cutting-edge research in areas such as photonics technology and drug design. The extensive HPC experience at ac3 made it a natural partner for the consortium, because ac3 has the expertise to help with requirements definition, system design, tender offer and acceptance and performance testing. Today, ac3 operates the supercomputer for the consortium.

"The consortium of five universities led by the University of Technology Sydney and ac3 wanted to build a very large and cost-effective supercomputer," says Dr. Vladas Leonas, chief technology officer at ac3. The initial design goal was to provide about 500

**"Once you own hardware of this complexity, there is a lot you can do to tweak its performance. Intel® Solution Services has tremendous expertise in high-performance computing, and we knew they could extract the maximum value from the Dell cluster by fine-tuning software and optimizing code using Intel® compilers designed for the latest systems."**

**Dr. Vladas Leonas**  
Chief Technology Officer  
Australian Centre for  
Advanced Computing and  
Communications

gigaflops of performance. Other requirements included 24x7 availability, the ability to upgrade the supercomputer as technology evolves and the scalability to handle growing workloads. Dr. Leonas adds, "Based on the consortium's applications, which include a lot of Gaussian code, we determined that a fast, 32-bit processor-based system would provide the best price/performance. Once that decision was made, there was no question that we wanted to build the supercomputer on a cluster of Intel® architecture-based servers to get the most power for the money."

To make the best possible design decisions and to tune the performance of the system once installed, ac3 looked for a consulting partner to supplement its own staff.

## Business Solution

ac3 found the perfect match in Intel® Solution Services. Says Dr. Leonas, "The Intel Solution Services team is not tied to any one server or software vendor. That independence proved invaluable in helping with the tender. It was also clear to the ac3 staff that Intel® compilers would be used because of their reputation for high performance. And who better to help tune software that used those compilers than consultants from Intel Solution Services?"

The initial design comprised six racks that included 128 compute servers plus two cluster head nodes and a file server. To increase scalability and to handle the large number of compute nodes, all servers were connected via a Foundry Networks FastIron\* 1500 Gigabit Ethernet switch. The Red Hat Linux\* 9 operating system was selected for its flexibility and its tight integration with OSCAR cluster software, which enables multiple computers to perform large-scale applications in parallel as one computer.

In June 2003, after the competitive bids were evaluated, ac3 chose Dell as the server vendor because of its extensive experience with large clusters. Adds Dr. Leonas, "Dell offered more cluster nodes than any other vendor." Dell proposed to provide Dell PowerEdge\* 1750 servers with dual Intel® Xeon™ processors, which are designed for cost-effective, high-performance computing environments where performance and scalability are key. The PowerEdge 1750 servers have the following configuration:

- Two Intel Xeon processors running at 3.06 GHz with Hyper-Threading Technology
- A 533 MHz front-side bus
- A mix of 512KB L2 cache and 1MB L2 cache
- Dual Gigabit Ethernet network interfaces
- 2GB to 4GB Dual Channel DDR RAM



*The ac3 cluster is comprised of 152 compute nodes and two head nodes, each with dual Intel® Xeon™ processors running at 3.06 GHz.*

All the equipment was installed and integrated in August and September 2003, so the acceptance testing and tuning phases could begin. Intel Solution Services, Dell, Foundry and ac3 engineers worked together to ensure that the system was well integrated and optimally configured. Intel Solution Services used the SPEC\* CPU2000 benchmark suite, pre-compiled with Intel® C++ Compilers 7.1 for Linux, to execute the acceptance tests on all machines of the cluster. All nodes successfully met the performance level specified in the tender document for integer and floating-point calculations.

## FINE-TUNING HELPS MAXIMIZE PERFORMANCE

With the acceptance testing completed, the team of Intel Solution Services consultants and ac3 engineers began their efforts to optimize the system throughput and determine the maximum performance. As the center's staff witnessed the exceptional performance of the cluster, they became determined to be named to the TOP500 list of the world's most powerful supercomputers.

The key to getting the most from an existing HPC server is software tuning. "There is a lot you can do to tweak the performance of hardware once you have it," Dr. Leonas says. "Intel Solution Services has tremendous expertise

in high-performance computing, and we knew its technical team could wring the maximum performance from our cluster by fine-tuning the application code and the Intel compiler options.”

The team used the LINPACK benchmark suite to measure cluster performance. When the tender was written, ac3 estimated that, after optimization, the 128-node Dell cluster would perform at approximately 600–700 gigaflops with Intel Xeon processors running at 3.06 GHz. With another university providing an additional 16 nodes, ac3 was able to start the testing with 144 compute nodes. The very first LINPACK tests showed performance of approximately 550–600 gigaflops.

Intel Solution Services consultants not only have a great deal of experience with the LINPACK benchmark, but they also can access Intel resources around the world. “We could change about 50 parameters in LINPACK,” Dr. Leonas says. “The Intel Solution Services team showed us the parameters we could ignore and those which were important. For example, we learned that if the amount of memory on a node increases too much, performance actually goes down. Intel Solution Services helped us understand how best to allocate jobs across nodes. The highly experienced consultants from Intel helped in many ways and passed the knowledge to our engineers and researchers so we could carry on after they were gone.”

**“The value of our ongoing relationship with Intel® Solution Services is absolutely huge. We will always strive to get the most out of our systems. It is very important to be able to turn to Intel for first-class information that we trust.”**

**Dr. Vladas Leonas**  
Chief Technology Officer  
Australian Centre for  
Advanced Computing and  
Communications

Throughout the LINPACK optimization process, Intel Solution Services consultants and the ac3 engineering team tuned and optimized the network, the operating system, the cluster software, the file server and the LINPACK benchmark to achieve maximum cluster efficiency. The first round of optimizations, which used the Intel® Math Kernel Library (MKL) 6.0 and “libgoto” BLAS performance library, yielded a performance of 960 gigaflops—tantalizingly close to the teraflops mark.

Ranking in the TOP500 (November 2003)	108
Peak LINPACK performance	1.095 teraflops
Number of compute nodes	152
Number of I/O nodes	2
Number of file servers	1

Figure 1. Facts about the ac3 supercomputer cluster

### **EXCITEMENT BUILDS AS TERAFLUPS PERFORMANCE NEARS**

“When the teraflops mark was in sight, we became really excited and decided to try to surpass it,” says Dr. Leonas. “We thought we needed a few more machines to achieve the one-teraflops level, so we ordered eight more compute nodes. But it turned out we did not need them. Upgrading to the newly released Intel MKL 6.1, along with additional tuning of the LINPACK benchmark suite by Intel Solution Services, put us over the top. Those steps alone enabled us to achieve 1.036 teraflops.”

The eight new nodes arrived a week after the cluster broke the teraflops barrier. Once they were integrated into the cluster—now totaling 152 compute nodes—Intel Solution Services performed a final round of LINPACK tests that resulted in the highest score of all—1.095 teraflops.

With the help of Intel Solution Services, ac3 achieved one of its main goals—a rank in the TOP500 list of supercomputers. When the TOP500 list was published in November 2003, ac3 ranked 108th among the fastest supercomputers in the world, higher than any other Australian computer on the list.

### **CONSORTIUM IS THRILLED WITH ITS NEW SUPERCOMPUTER**

LINPACK testing, originally used to confirm that the Dell computer cluster worked as planned, became a significant marketing benefit for both ac3 and the university consortium that funded the project.

Says Dr. Leonas, “The university consortium was thrilled to find that it now owns (according to the TOP500 list) the fastest supercomputer in Australia—a far more powerful machine than it originally envisioned. The cluster provides a staging area to prepare and produce research results, and its ability to facilitate first-class science research will allow the consortium to apply for more research funds.”

Breaking the teraflops barrier also is beneficial to ac3, in part because the feat was widely reported in the Australian press. "Our teraflops capability has given ac3 much greater visibility and prestige in Australia, which will help us attract new commercial customers for our co-location and hosting services," Dr. Leonas says. "We also are attracting new types of customers with larger problems who want to take advantage of the power of the Dell cluster."

#### **INTEL SOLUTION SERVICES WILL CONTINUE TO SUPPORT AC3**

Dr. Leonas anticipates an ongoing relationship with Intel Solution Services. "Our engineering team, already very strong, learned a great deal about optimizing a super-computer cluster from working with Intel Solution Services. The Intel team members are excellent teachers, so they helped us increase our capabilities. But the cluster will run a diverse set of applications, each with different requirements." Concludes Dr. Leonas, "We know that, going into the future, we can trust Intel Solution Services to provide first-class assistance as we evolve."

#### **LESSONS LEARNED**

- **Consultants can be valuable in the design phase.** Intel Solution Services helped ac3 and Dell design the optimal cluster configuration. They tackled issues such as server configuration, how to connect a large number of nodes for the best performance, which operating system to use and how to set up the file and I/O servers.
- **Once you have your hardware, most performance gains come from software tuning.** Intel Solution Services—with its extensive experience with the LINPACK benchmark suite and Intel compilers, libraries and other tools—can tune operating systems and software codes to achieve the maximum efficiency for performance and availability.
- **Keep up-to-date with the latest releases of operating systems, drivers, compilers and libraries.** ac3 gained a significant performance boost by upgrading to the latest compilers and libraries from Intel and other vendors as soon as they became available.

Intel<sup>®</sup>  
**solution**  
services

*Solution provided by*

Intel<sup>®</sup> Solution Services is Intel Corporation's worldwide professional services organization, helping enterprise companies capitalize on the full value of Intel<sup>®</sup> architecture through consulting focused on architecture transitions. For more information, visit

[www.intel.com/internetservices/intelsolutionservices](http://www.intel.com/internetservices/intelsolutionservices)

Intel works with the world's largest community of technology leaders and solution providers, from software and hardware to systems integration and services companies that all are working with Intel<sup>®</sup> products, technologies and services with a common goal of providing better, more agile and cost-effective business solutions.

**Find out more about a business solution that is right for your company by contacting your Intel representative, or visit the Intel<sup>®</sup> Business Computing Web site at [intel.com/ebusiness](http://intel.com/ebusiness) or Intel's industry-specific sites: [intel.com/modelschool](http://intel.com/modelschool); [intel.com/go/retail](http://intel.com/go/retail); [intel.com/go/manufacturing](http://intel.com/go/manufacturing); [intel.com/go/digitalmedia](http://intel.com/go/digitalmedia); [intel.com/go/finance](http://intel.com/go/finance); [intel.com/go/telco](http://intel.com/go/telco); [intel.com/go/hpc](http://intel.com/go/hpc)**



This document and the information given are for the convenience of Intel's customer base and are provided "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed or contained herein. Intel products are not intended for use in medical, life-saving, life-sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications, product descriptions and plans at any time, without notice.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference [www.intel.com/procs/perf/limits.htm](http://www.intel.com/procs/perf/limits.htm), or call (U.S.) 1-800-628-8686 or 1-916-356-3104.

Intel, the Intel logo and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

\*Other names and brands may be claimed as the property of others.