

Win in the flat world

Nomura derives superior business performance and scalability using Grid Computing

Abstract

Nomura Securities partnered with Infosys to address business performance and scalability issues in its Interest Rate Risk Analysis application. Using Infosys' High Performance Computing solutions, Nomura was able improve the performance of the application by more than 150 times by porting the application to a Grid environment, thus demonstrating the Proof of Concept (PoC). The solution improved the scalability of the application by supporting several hundreds of products as opposed to under a hundred earlier, enabling Nomura to easily introduce new products, more simulations and research for products involving complex calculations, thus removing a constraint to business growth.



The client

Nomura Securities Co., Ltd. (NSC) plays a central role in the Nomura Group's core securities business. As the leading securities and investment banking company in Japan, they provide individual investors and corporate clients with a broad range of services, including investment advisory services and securities underwriting.

Background

Nomura was facing performance and scalability issues with its current Risk Analysis application for Mortgage-backed Securities. The C/C++ based application, running on Solaris faces time constraints, thereby impacting the processing of additional products and constraining their business growth. The performance of the application was also impacting the time taken by traders and research analysts to do risk analysis relating to market sensitivities for the securities, as it took several hours to generate risk reports for under hundred products. As this was one of the roadblocks to business growth, it was imperative for Nomura to improve the performance of the application.

In the quest to improve performance and scalability, Nomura was faced with a number of challenges. The key challenges were:

- The application involved numerous mathematical functions. So, the opportunities for performance improvement using traditional approaches, like code tuning, were extremely limited.
- The application did not cater to running multiple concurrent threads. Hence, multithreading was not a viable option to improve performance and scalability.
- Floating point accuracy of the output was crucial to the calculations which limited the options for tuning the compiler.
- Parallelizing the application was an option. Choosing the right degree of parallelism was a challenge. It was constrained by the availability of the CPU and by the capability of the middleware to handle tasks fired

Infosys' Solution

Nomura engaged with Infosys to address these performance and scalability issues. As part of the engagement, Infosys used a pro-active and holistic performance engineering approach to amplify the performance and scalability of the application. The approach involved:

- Analyzing the application with the objective of introducing parallelization to address performance issues.
- Evaluating different degrees of parallelization, using tools such as Platform's Symphony and Gemstone's Gemfire, to determine optimal performance
- Addressing compiler issues by changing from GNU GCC to Intel ICC and tuning the compiler to maximize performance without sacrificing floating point accuracy

- Using Infosys' patent-pending Grid Application Migration Framework to determine the best approach for porting the application to run in a Grid environment
- Identifying critical sections of the application that could be distributed across a Grid.
- Evaluating and using commercially available Grid middleware platforms to maintain grid clusters and facilitate load balancing
- Introducing a data caching mechanism to reduce performance overhead by limiting data movement between nodes of the Grid

Benefits

Nomura derived significant improvement in application performance (PoC level) among other benefits:

- The time taken to generate reports dropped from several hours to about 1 minute for under hundred products; reducing the processing time to about 1/200 of the original processing time. This improvement allows the application to process 500 products in less than 8 minutes
- Time reduction may allow more simulations and research for those kinds of products involving complex mathematical calculations.
- The solution offers non-intrusive scalability to address future application needs as it can scale simply by adding nodes to the Grid
- Migrations costs will be reduced as the solution is upward compatible (with middleware versions)
- The solution can support both 32 and 64 bit processor architectures, broadening the range of hardware that can be adopted in the future

Technology Highlights

- Grid with dozens of CPU cores running on Xeon processors
- SUSE Linux
- C / C++
- Platform Symphony
- Global Parallel File System
- Gemstone Gemfire
- Support for 32 bit & 64 bit processor architectures
- Intel ICC
- Infosys Grid Application Migration Framework as an analysis tool

Client Testimonial

“This PoC has demonstrated the feasibility of using grids to enhance the performance of applications involving complex calculations. We believe this research by enhancing the performance of the application significantly, will enable us to make decisions quicker and improve our ability to introduce more products and understand and research those interest rate risks within a short span of time. Infosys’ thought leadership in the High Performance Computing space and its performance engineering methodology were critical in achieving these results. Another key factor was the extensive domain knowledge demonstrated by Infosys’ in IT, financial transactions and the finance domain.”

- Toshiaki Nomura, Quantitative Research, Nomura Securities

About Infosys’ Grid Application Migration Framework (GAMF)

Infosys’ GAMF, used for code analysis and design of parallel code, analyzes the application portfolio and profiles code for Grid enablement. It takes legacy programs in C, C++, or Java as input and generates a Directed Acyclic Graph (DAG) which depicts both task and data dependencies among the components of the program. Grid tasks of proper granularity are generated using a set of DAG reducing and clustering algorithms

The framework simulates the actual Grid execution by taking the reduced task graph as input and schedules the tasks on different processors in the Grid. The performance data is then analyzed to study the benefits of porting the application to Grid.

© 2008 Infosys Technologies Limited.

ALL RIGHTS RESERVED

Copyright in whole and in part of this document “**Nomura derives superior business performance and scalability using Grid Computing**” belongs to Infosys Technologies Limited. This work may not be used, sold, transferred, adapted, abridged, copied or reproduced in whole or in part in any manner or form or in any media without the prior written consent of Infosys Technologies Limited.