


ACTUARIAL HIGH-PERFORMANCE COMPUTING  
**INCREASINGLY ACCESSIBLE AND INDISPENSABLE**  
BY PAT RENZI AND JIM BRACKETT



What does running an insurance risk model have in common with ray tracing to create photorealistic computer graphics? Millions of data points and the technology it takes to analyze them.

Today, competitive pressures and regulatory changes in the life insurance industry are demanding increasingly complex analysis and, as a result, more sophisticated actuarial tools. But knowing what needs to be understood and having the technology to analyze it are often two different things.

Milliman has created programs that generate financial projections to support highly complex risk analysis. Dedicated cluster computing has become an essential tool for this level of modeling, but the costs have been prohibitive to many in the insurance industry. Recently, Milliman teamed with Microsoft® to provide a scalable cluster computing solution that opens the door to highly sophisticated analyses to nearly everyone in the insurance business, both large and small companies.

\* \* \*

**A BRIEF HISTORY OF CLUSTERING** It was only 20 years ago that insurers began running seven scenarios to model risk and return, up from a long-time standard of evaluating a single scenario. Over the years, technology has provided the industry with tools

to increase the accuracy of projections, and eventually insurers began running 50 scenarios. In today's highly demanding risk-analysis environment, it can now take 1,000 scenarios and millions of data points to effectively manage risk and return. In fact, a *1,000-scenario model with reserves and capital based on 1,000 paths at each valuation point for a 30-year monthly projection requires the cash flows for each policy to be projected 360 million times.*

For the insurance industry, modeling continues to grow more sophisticated. Some of today's complex models exceed the capabilities of desktop computers and even enterprise computing resources. To meet these challenges, actuaries harness computing power through a multitude of machines, or a high-performance computer cluster. There are companies that are running stochastic and nested stochastic projections on clusters with as many as 1,500 PCs.

**A 1,000-scenario model with reserves and capital based on 1,000 paths at each valuation point for a 30-year monthly projection requires the cash flows for each policy to be projected 360 million times.**

That level of sophisticated analysis is familiar ground to a range of industries. Oil and gas exploration is the oldest and largest user of high-performance computing, with some estimating that as many as 30,000 servers have been employed by a single company for seismic analysis and reservoir modeling. High-performance computing is used for molecular modeling and protein folding in drug design, and it is playing a critical role in the world's largest particle physics microscope, the Large Hadron Collider.

\* \* \*

**THE NEW AFFORDABILITY** For the widespread insurance industry, high-performance computer clusters come at a cost that most have not been able to afford. Even the deployment and maintenance of smaller-scale, 50-PC, high-performance computers demand an investment in hardware and expertise outside the reach of many insurance firms. In fact, some large companies struggle to dedicate the required resources.

Most small and mid-sized firms are working to keep up with this evolution in the industry, but the magnitude of the investment of resources it requires will pose challenges. To help provide a more cost-effective solution and, as a result, much greater access

#### **HIGH-PERFORMANCE, OVERNIGHT**

The high-capacity analysis provided by cluster computing also is essential to run scenarios for market exposure hedging. Milliman's MG-Hedge® provides a market scenario generator that produces hundreds of millions of time series paths based either upon a risk-neutral methodology or upon a "realistic scenario" methodology. With this system, clients are able to look at risks based on their actual portfolios and determine the optimal course of action. *The sophisticated modeling provided by MG-Hedge – which for any given application might include stochastic volatility and/or stochastic interest rates in a stochastic-on-stochastic framework – requires analysis of sometimes billions of data points.*

MG-Hedge was designed from the beginning to model actual data over many "what-if" scenarios, rather than use aggregate data or make broad assumptions. With this type of analysis, as with insurance policy modeling, increasing the capacity for data processing provides better information and decisions. Running on numerous systems also provides the opportunity to maximize resiliency. If there is a problem with one computer, there are many others available to ensure that the integrity of the process is preserved.

In addition to processing enormous amounts of data, MG-Hedge must generate information that is highly time-sensitive. Our clients' hedging strategies need to be determined overnight. As a result, MG-Hedge was built to always run on high-performance cluster computing.

to important grid technology, Milliman worked with Microsoft to integrate its financial modeling tool, MG-ALFA® (Asset Liability Financial Analysis), with Windows Compute Cluster Server 2003. MG-ALFA supports complex stochastic and nested stochastic projections, which can require hundreds of hours of computing time without cluster computing. The integration of MG-ALFA with Microsoft's Windows Compute Cluster Server 2003 provides distribution of sophisticated analysis to clusters ranging from a single node to several thousand simultaneously.

Through the solution's integrated Job Scheduler, both job creation and submission can be performed directly from a desk-side application, which helps make complex models accessible to a wider swath of the industry.

"Microsoft and Milliman are working closely together to address the growing technological challenges presented by increasingly complex actuarial analysis," said Jeff Wierer, senior product manager of high-performance computing at Microsoft Corp. "Together, we're able to provide our customers with a solution using a familiar user interface that easily integrates with their existing system."

In order to make its solution more widely affordable, Microsoft released the Windows Compute Cluster Edition (CCE) of Windows Server 2003, which is fully compatible with existing 64-bit versions of Windows Server 2003 and runs a full range of actuarial modeling programs. This version of Windows Server significantly reduces the software cost for implementing high-performance computer clusters and can support processor counts in the hundreds or thousands. Microsoft combines CCE with the Microsoft Compute Cluster Pack (CCP) as the components of Windows Compute Cluster Server 2003. The system serves to control and mediate all access to cluster resources as well as provide a single point of management, deployment, and job scheduling for the computing cluster.

The ability to run highly sophisticated and specialized analyses quickly and accurately helps level the playing field. The digital divide of high-performance cluster computing has threatened to leave many in the insurance industry behind. Increasingly, fast and precise risk modeling is becoming mission-critical. Eventually, the information gained through high-performance cluster computing will likely be as central to the insurance industry as it is to theoretical physics and computer gaming. **M**

**PAT RENZI** is a principal with the Seattle office of Milliman and is responsible for the oversight of the marketing, development, planning, and client service for Milliman's MG-ALFA product.

**JIM BRACKETT** manages the Milliman Financial Technology practice in Chicago. He and his team specialize in the implementation of high-performance, enterprise-scale distributed computing systems.