

# MG-ALFA with GridServer Boosts Business Analysis and Computing Capacity

## Industry

Insurance

## Business Challenge

Dramatically increase capacity and enhance performance of key insurance applications.

## Technology Solution

The MG-ALFA\* actuarial projection system with the GridServer\* grid-computing solution—an integrated solution from DataSynapse and Milliman.

## Enterprise Hardware Platform

Intel® architecture-based systems

Intel® Xeon® processor family-based servers

Intel® Centrino® mobile technology-powered laptop computers

## SOLUTION ARCHITECTS

A MILLIMAN GLOBAL FIRM



**Milliman**

Consultants and Actuaries



**Data  
Synapse**





## MEETING NEW MARKET DEMANDS

Milliman and DataSynapse together provide a grid-enabled actuarial projection solution that allows insurance organizations to meet changing regulatory requirements, reduce time to market for new products, and decrease the time needed to analyze and manage risk. The Milliman MG-ALFA solution distributes the processing of actuarial projections across an unlimited number of GridServer engines, dramatically reducing the runtime for such projections. It reduces a client's total cost of ownership (TCO) by harnessing the power of servers, desktop systems and laptop computers anywhere on a client's network, whether dedicated or under-used. Intel® architecture provides the building blocks for the grid-computing engines. With powerful and cost-effective architecture, Intel® processor family-based systems provide the performance, flexibility to scale, reliability, and security that power this solution.

## THE BUSINESS CHALLENGE

As competition in the insurance industry intensifies, pressure to bring products to market faster and offer product features that decrease risk to the company is increasing. At the same time, financial regulators around the world are increasingly focused on regulatory capital for insurance companies. The combination of these forces is behind the massive shift in the industry from deterministic pricing and analysis to stochastic pricing and analysis. Insurance company actuaries thereby face the need to exponentially expand their computing capacity to move to stochastic analysis for product development, risk management, and regulatory compliance. In many companies, a single projection of a large model on a standalone machine can take an entire evening to run. Considering that recent regulations now require that the same models need to be projected across hundreds or thousands of economic scenarios to determine capital requirements, it is clear that these companies face a difficult challenge.

## THE SOLUTION OVERVIEW

The MG-ALFA actuarial projection system with GridServer grid computing—an integrated solution from Milliman and DataSynapse—is designed to deliver significantly expanded computing capacity to meet today's challenging actuarial computing needs. While many applications can be grid-enabled on a one-off basis by a client or a grid-computing vendor, the MG-ALFA solution is fully grid-enabled out of the box. The MG-ALFA solution can distribute the processing of models over hundreds or thousands of scenarios across a grid of engines, dramatically reducing runtime requirements.

For example, the proliferation of guaranteed benefit riders on variable products and secondary guarantees on life insurance products has led to the need for pricing to be done on a stochastic basis. The scenario set must be sufficiently large to allow for analysis of tail risk. A pricing exercise involving secondary guarantees on a universal life product can require 500 or more scenarios to run against a model with 1000 cells. On a standalone machine, such a run requires overnight processing and allows the actuary only one pass at the model each evening. With a grid of 100 machines, however, model results can be available in approximately 20 minutes, allowing dozens of iterations per day.

This accelerated turnaround time for stochastic pricing models results in enormous competitive advantage. Using this solution, a company can bring sophisticated products to market in a fraction of the time—with far better understanding of the associated risks.

**Solution Blueprint:**  
MG-ALFA with GridServer  
Boosts Business Analysis  
and Computing Capacity

**Solution Provider:**  
Milliman Consulting  
and Actuaries

## Case Studies/Proof Points

For a case study, see  
[www.datasynapse.com](http://www.datasynapse.com)

In addition to delivering a quantum leap in performance and capacity, a grid-computing solution is more cost-effective than purchasing dedicated hardware to solve the same problem. A grid-computing solution can be deployed across a company's existing desktops and use otherwise idle processing capacity, reducing the need to purchase, maintain, and upgrade expensive dedicated hardware to meet expanding actuarial needs.

The grid-computing solution allows insurers to perform new types of business analysis—many of which are emerging as leading management metrics. Application areas include:

- **Stochastic pricing and projections.** Stochastic pricing and projections are used to measure risk in the development of new products and to measure and manage risk in existing blocks of business. The integrated grid-computing solution can dramatically shorten the time-to-market for new products and accelerate risk analysis to a real-time process rather than an overnight process, resulting in a strong competitive advantage.
- **Economic capital.** Many companies, especially those with ownership in Europe, are moving to economic capital or similar management metrics as the foundation for measuring success and determining management compensation. While these measures are definitely superior to more deterministic measures, they require stochastic projections and, in many cases, nested stochastic projections. The grid-computing solution of MG-ALFA and GridServer is ideally suited to meet the exponential increase in computing capacity needs.
- **Regulatory compliance.** In the past, reserves and risk-based capital were easily calculated using deterministic formulas. Recently, however, the industry has started to shift toward principles-based reserves and capital requirements. This requires a stochastic approach to calculate the values for a single point in time, and a nested stochastic approach to calculate the values for future points in a financial projection. On a standalone machine, the runtimes for such projections would be measured in days or even weeks. The MG-ALFA and GridServer solution, however, delivers the speed and capacity necessary to allow companies to calculate these values in real time.

## TECHNOLOGY

The solution integrates the MG-ALFA actuarial projection system and the GridServer grid-computing system. Using adaptive grid infrastructure software that virtualizes and distributes application services and executes them in a guaranteed, scalable manner, this solution:

- Partitions large workloads into smaller tasks
- Distributes these tasks in parallel to computing resources (that is, to dedicated servers or underused, intermittently available computing resources on the network)
- Returns results without user intervention or system management overhead

By creating a virtualized environment that matches demand (user and application service requests) with supply (available system and data resources), this integrated solution dramatically improves application performance, resiliency, and uptime.

The MG-ALFA with GridServer solution is designed to run on any network configuration using engine and driver machines that run Microsoft Windows\* 95, 98, 2000, XP or Windows NT\* operating systems. Intel architecture provides the processor hardware building blocks for this solution; with its powerful and cost-effective architecture, Intel processor family-based servers allow businesses to implement reliable, scalable solutions that increase performance and availability. The flexibility in implementing solutions with Intel architecture leads to tailored deployment and rapid implementation.

While this integrated grid-computing solution is designed to run on a wide variety of processor hardware, systems based on more capable processors deliver faster throughput and better grid performance. Intel® Xeon® processor-based systems provide the flexibility and application headroom to achieve outstanding performance today, and provide a solid path for migrating to 64-bit operating environments in the future. Available in 32-bit and 64-bit configurations, Intel Xeon processors can support large memory configurations, making them ideal for supporting the analysis of large data sets. In many businesses, an increasingly mobile workforce relies on high-performance laptop computers, such as those based on Intel® Centrino® mobile technology. Network connected laptop computers can also be included in the grid. Intel Centrino mobile technology-based laptops provide high performance in a grid environment, thanks to their highly efficient processor architecture.

The minimum recommended RAM per grid computer is 256K, although smaller models can run with less.

### WHO THE SOLUTION WILL BENEFIT

The MG-ALFA with GridServer solution is designed to benefit both actuaries and senior management in the insurance industry. This solution allows actuaries to run model projections in a fraction of the time it used to take, allowing for additional analysis, faster analysis, and analysis that was not possible in the past due to lengthy runtimes.

Senior management benefits by gaining better information through more thorough analysis, shorter turnaround time on required financial projections, and shorter turnaround time on what-if requests. The solution's ability to use otherwise idle processor capacity also provides overall lower cost relative to dedicated hardware solutions, greatly reducing TCO.

### SOLUTION BENEFITS

The MG-ALFA solution enables insurers to tap existing system availability, expanding the capacity to run new and more complex business analyses with multiple iterations. Exponentially faster runtimes allow for more analysis, better analysis, and analysis that was not practical without this solution. The solution delivers benefits in three areas:

- Product pricing
- Asset liability management (ALM)
- Regulatory compliance

### PRODUCT PRICING

With the proliferation of guaranteed benefit riders on variable products and secondary guarantees on life insurance products, product pricing must be done on a stochastic basis, with a set sufficiently large to allow for analysis of tail risk. Pricing exercises involving secondary guarantees on a universal life product may require 500 or more scenarios to be run against a model with 1000 cells. On a standalone computer, an analytical model such as this requires overnight processing and allows the actuary only one pass at the model each day. The MG-ALFA solution enables product pricing models to be run across a grid of computers. Using a grid of 100 desktop computers, the MG-ALFA solution can deliver those same modeling results in less than one hour.<sup>1</sup> This accelerated turnaround time for stochastic pricing models enables dozens of iterations of the pricing models per day, which translates to a better understanding of risk and a faster time to market. Companies that use the MG-ALFA solution can bring sophisticated products to market in a fraction of the time—with a far more comprehensive understanding of associated risks.

<sup>1</sup> Performance tests and ratings are measured using specific computer systems and/or components and reflect the appropriate performance of Intel products as measured by those tests. Any difference in system 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, go to [http://www.intel.com/performance/resources/benchmark\\_limitations.htm](http://www.intel.com/performance/resources/benchmark_limitations.htm)

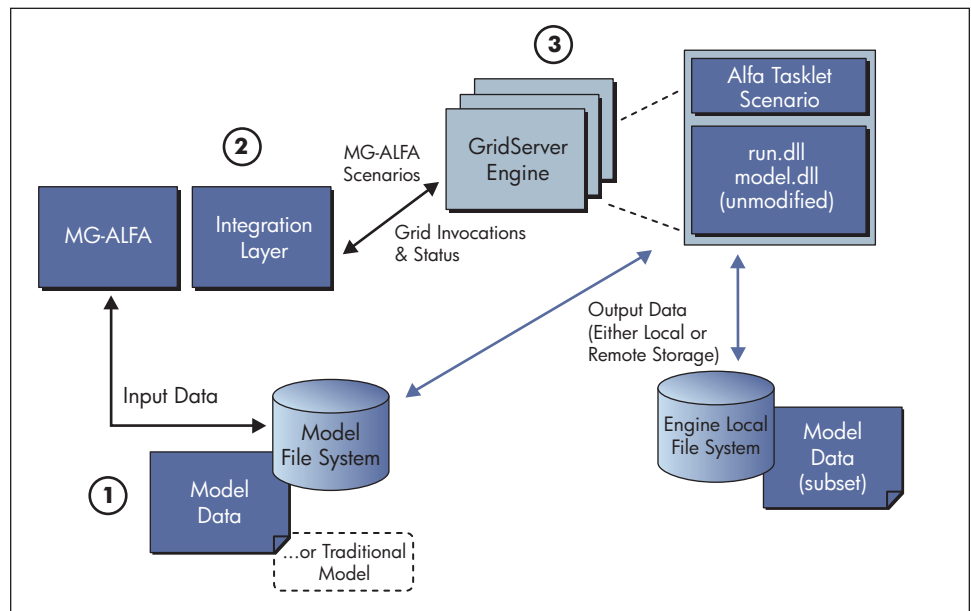
## ASSET LIABILITY MANAGEMENT

ALM activities typically use models at the line-of-business or company level based on deterministic or stochastic scenarios. A single pass through a complex scenario can, at best, run overnight and, at worst, may take days to process. Assuming a model takes an hour to run across a single scenario and a company wants to evaluate 25 scenarios, the runtime on a single computer would be 25 hours, allowing for only one iteration per day. A 25 computer grid, however, can run the ALM in one hour, allowing for up to multiple iterations per day. As a result, users of the MG-ALFA solution can gain greater insight into their ALM risks and conduct what-if analysis using their ALM models on a real-time basis.

## REGULATORY COMPLIANCE

Many actuarial aspects of regulatory compliance continue to evolve in a stochastic direction. While cash-flow testing is still limited to a smaller, defined set of interest scenarios, required capital calculations for variable annuities now require very large stochastic scenario sets. Principles-based reserve calculations for variable and interest-sensitive products are on the horizon, moving even more required actuarial calculations into the realm of stochastic analysis. The MG-ALFA solution allows users to meet these requirements with reasonable turnaround times, without heavy expenditure on IT infrastructure.

## FUNCTIONAL BUSINESS CONCEPT



**Figure 1.** Functional interactions between the components of the MG-ALFA actuarial projection system and the GridServer grid-computing system.

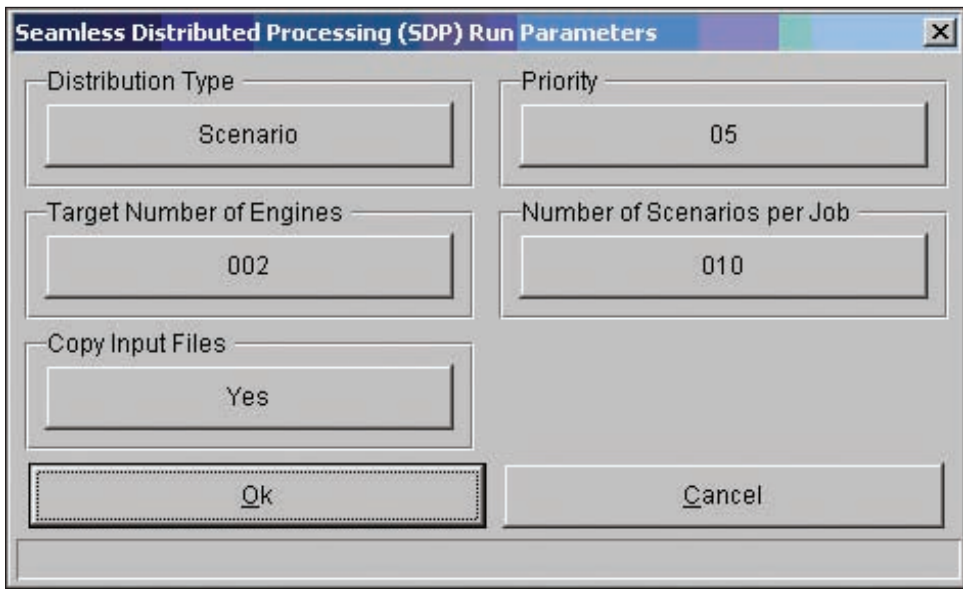
Users of the MG-ALFA solution create scenarios and submit the scenarios along with relevant model data for analysis.

1. The integration layer provides the interface to the underlying grid technology, which remains transparent to the user.
2. The integration layer is the gateway to the grid.

3. The scenarios are distributed across GridServer engines, which execute the submitted scenarios in parallel. The GridServer engines are Intel architecture-based computing systems, such as Intel Xeon processor family-based servers or Intel Centrino mobile technology-based laptop computers. Intel Xeon processor family-based systems deliver high performance and the ability to address very large memory, enabling fast and reliable analysis. Intel Centrino mobile technology-based laptops deliver responsive computation processing through their highly efficient processor architecture.

### USER EXPERIENCE

Using the MG-ALFA solution, the user submits a job for processing on the GridServer grid by setting the seamless distribution processing (SDP) setting in the software configuration to Yes. This displays the SDP Run Parameters dialog box, which is shown in Figure 2.

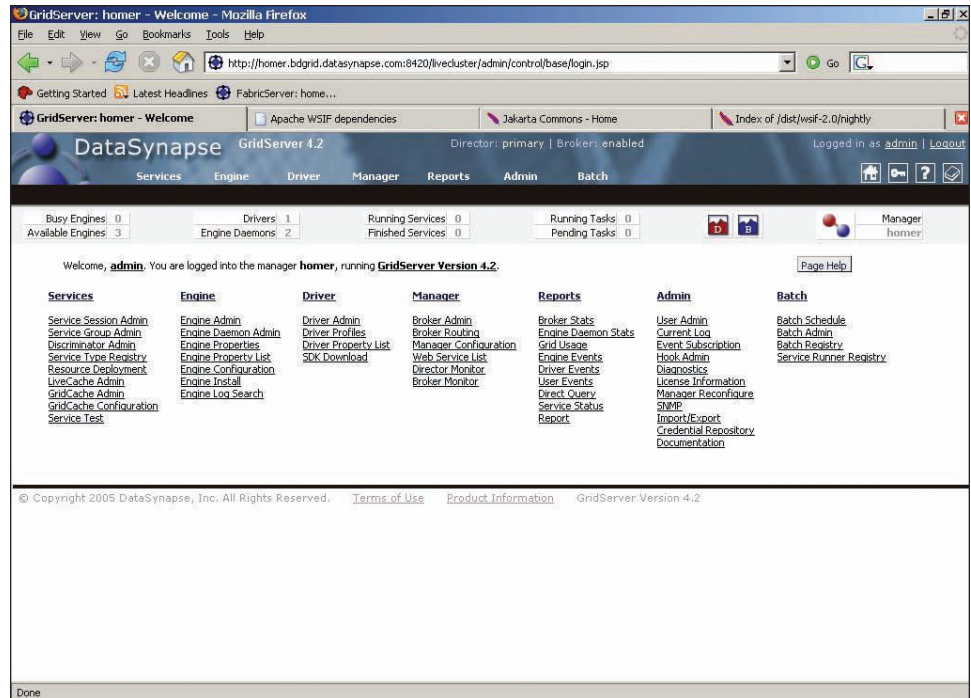


**Figure 2.** Seamless distributed processing run parameters dialog box.

Currently the Distribution Type settings are Scenario and NA. Target Number of Engines is the number of computing engines to use for processing the job. If this parameter is set to a very high number, all available computing engines are used. To launch two separate jobs and have the first job run on just five engines, this parameter would be set to five, leaving the other computing engines available to run another job.

**Copy Input Files:** This option is usually set to YES, unless very large-factor files are used in the project.

**Number of Scenarios per Job:** This option allows the user to dictate the granularity of a job. The user can experiment with this setting to find optimal settings for a project.



**Figure 3.** DataSynapse GridServer user interface.

## SOFTWARE ARCHITECTURE

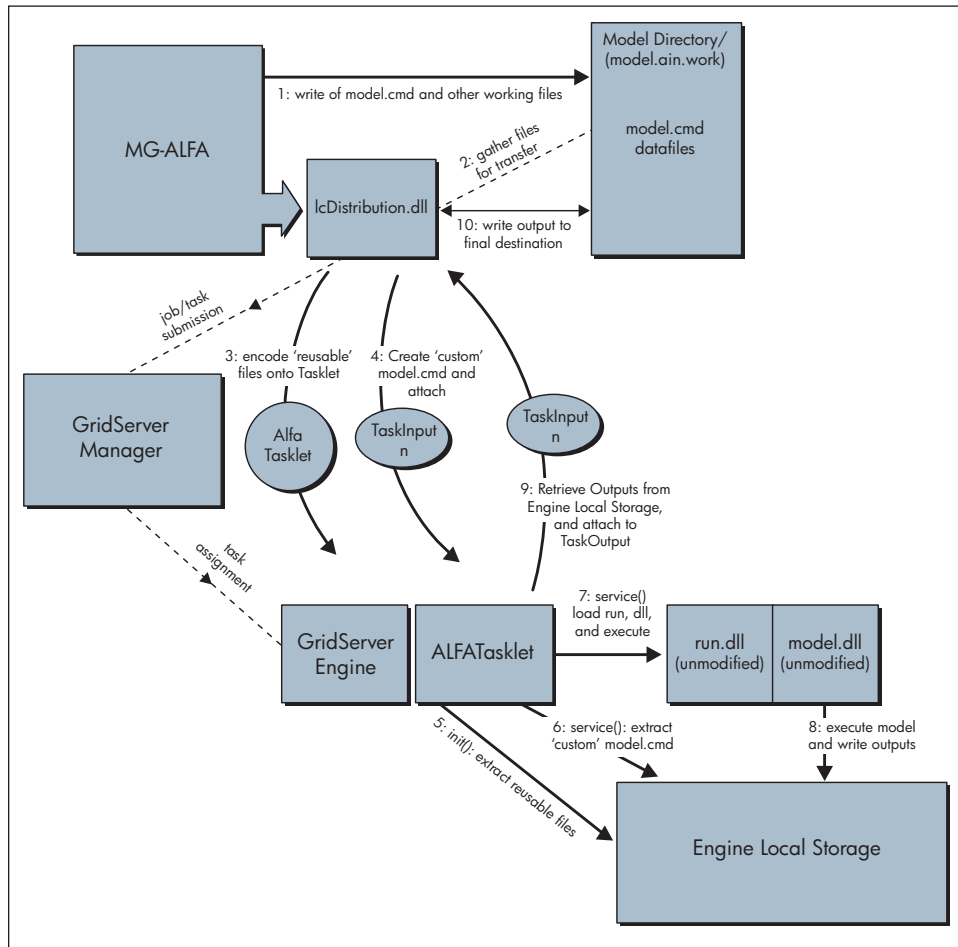
This solution is based on a highly scalable, adaptive grid infrastructure that virtualizes application and data services, independent of specific system resources. Client applications submit requests to the grid environment and GridServer dynamically provisions services to respond to the request. Multiple client applications can submit multiple requests in parallel. GridServer creates multiple service instances to handle requests in parallel on different grid nodes. For example, a single client will see scalable performance gains in the processing of multiple requests, and many applications and users will see scalable throughput of the aggregate load.

The result is a highly scalable architecture that offers linear improvements in speed and throughput.

A scalable architecture should provide linear gains in performance with the addition of more system resources (such as the Intel dual-core processors) while making no impact on the manageability of the distributed application platform. This means that the system should be as manageable with ten nodes as it is with 1,000 nodes. GridServer offers this type of operating environment—incremental resources that are added to the system impact the performance in a linear fashion but have no effect on management—including:

- Deployment of application services
- Provisioning data access and components
- Configuration and administration
- Runtime management and monitoring
- Historical reporting and usage statistics

On the engine, the MG-ALFA-supplied model.dll interface acts as a bridge between the user interface and the MG-ALFA calculation engine. When finished, control is returned to the user interface, which moves the results back to the project directory, properly updates any created factor files and creates a table of contents for the reports.



**Figure 4.** Software architecture.

The computing engines are Intel architecture-based systems, which provide the performance and reliability on which this solution relies. An important feature of this solution is that it provides reliable computations over networks of interruptible engines, making it possible to use intermittently active resources when they would otherwise remain idle. The engine launches when it is determined that the computer is idle (or that a sufficient system capacity is available in a multi-processor setting) and relinquishes the processor immediately in case of an interrupt (for example, by keyboard input on a desktop computer).

## **SYSTEM ARCHITECTURE**

The MG-ALFA solution submits the work to be completed as a series of tasks. The MG-ALFA solution contacts the GridServer Primary Director, which routes the request to an available broker. The grid-computing engines log on to the broker to perform individual tasks. While running MG-ALFA tasks, files move between the engines and the submitter using the DataSynapse LiveCache mechanism. A shared file system can be used in addition to the LiveCache.

The computing engines are Intel architecture-based systems. Intel architecture-based systems, such as servers powered by the Intel Xeon processor family and laptops based on Intel Centrino mobile technology, provide the high performance and reliability on which this solution relies. The GridServer technology allows this solution to distribute the processing of actuarial projections across an unlimited number of Intel architecture-based grid-computing engines, dramatically reducing the runtime for such projections.

## **SUMMARY**

The MG-ALFA with GridServer solution from Milliman and DataSynapse delivers the computing horsepower for meeting the expanding computing capacity and needs of insurance companies. Designed to deliver exponential increases in capacity and performance in a user-friendly environment, this solution provides much lower TCO than a dedicated hardware solution. Using the MG-ALFA solution allows insurers to leverage existing system availability, expanding their capacity to run new and more complex business analyses, and perform multiple iterations significantly enhancing decision support. Intel architecture-based systems, such as Intel Xeon processor family-based servers and Intel Centrino mobile technology-based laptop computers, provide the high performance, flexibility, and reliability that are the backbone of this solution.

## **LEARN MORE ABOUT THIS INNOVATIVE SOLUTION**

For general information about the products described in this solution blueprint, visit:

**[www.datasynapse.com](http://www.datasynapse.com)**

**[www.milliman.com/mg-alfa](http://www.milliman.com/mg-alfa)**

**[www.intel.com/go/solutionblueprints](http://www.intel.com/go/solutionblueprints)**

If you have specific questions about implementing this solution within your organization, contact your Intel representative or e-mail us at: [solutionblueprints@intel.com](mailto:solutionblueprints@intel.com).

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

Intel, the Intel logo, Intel Centrino, 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.

Copyright © 2006 Intel Corporation. All rights reserved. 0206/FP/MESH/XX Intel Literature Center: 1-800-548-4725

311829-001US

---