


www.amd.com/opteron



Powered by the
AMD Opteron™ processor



31029A

AMD & GRID Technology

The AMD Grid:

Enabling Grid Computing for the Corporation

Written By

Quentin Fennessy, Program Manager, Design Enablement, AMD

Ramesh Joginpalli, Strategic Marketing Manager, AMD

August 2006

Advanced Micro Devices, Inc. is no stranger to Grid computing and as a result, AMD has become a leading global provider of innovative microprocessor solutions for computing, communications and consumer electronics markets. This paper presents an overview of AMD's Compute Grid infrastructure, the benefits we derive from the Grid and how the Grid has enabled company growth. We will describe the construction of the AMD Compute Grid, using AMD-based systems, to design the AMD Athlon™, Athlon64™ and AMD Opteron™ processors running on the open source operating system Linux. We use EDA software tools from the leading ISVs for circuit design, layout, and verification--for virtually all steps in our design process. Our Grid-based work flow was first developed over 10 years ago and continues to evolve every day. AMD engineers depend on our Grid for the computing power necessary for architectural design, functional verification, circuit design and layout, and mask generation. We use AMD-based systems to develop new AMD processors – building AMD on AMD.

Introduction

Grid computing is an evolving technology that allows organizations to consolidate, share, and manage IT resources in a distributed computing environment. Compute Grid technology promises to be a significant development that enhances organizational flexibility and end user productivity through the efficient and scalable deployment of standards-based servers and networks. The efficiency of Compute Grids allows for reduced staff cost and increased system utilization--lowering the total cost of ownership. The scalability of Compute Grids allows an organization to deploy significant resources as needs are identified and expand over time. Successful companies are often challenged by increasing their internal resources to match their business growth. Compute Grids are seen as an infrastructure strategy to help companies meet those challenges.

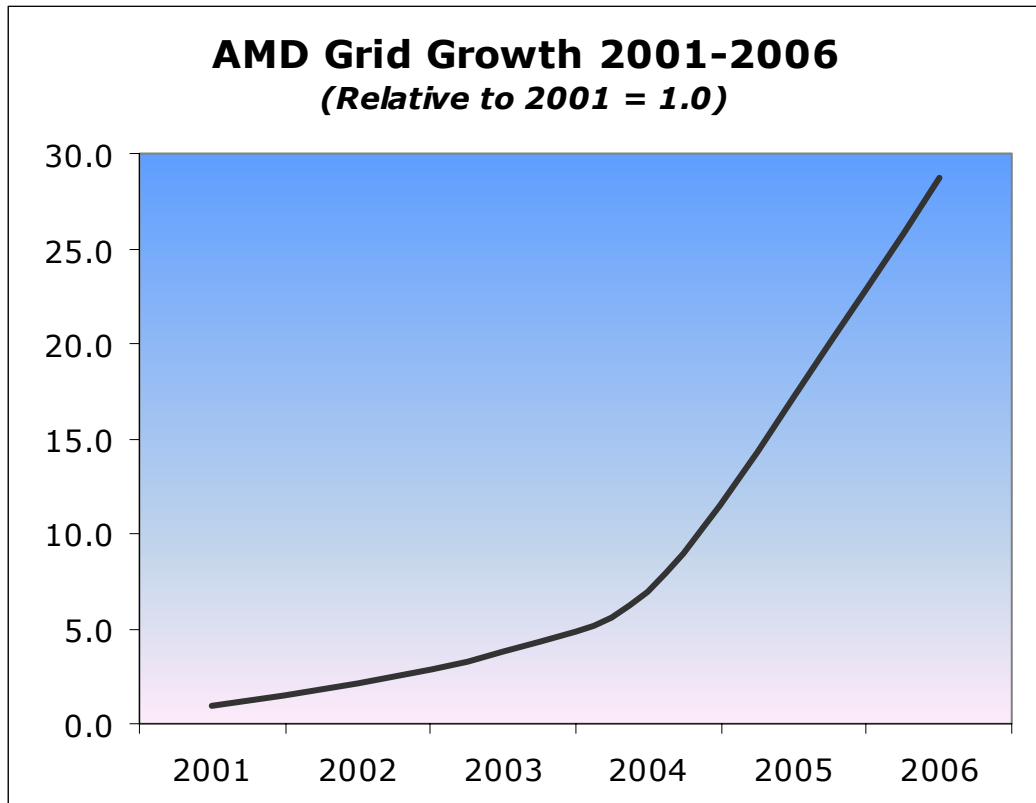
An IDC (International Data Corporation) survey predicts that grid technology will soon become pervasive in industry. IDC expects the number of computer systems that are actively attached to grids to grow to about 25% of all computers by 2008. Active systems on a grid are both grid capable (i.e., have grid software installed and activated) and accessible through one or more grids.

AMD Grid Overview

AMD was an early adopter of grid technology. We first created a compute grid for verification of high-performance microprocessors over 10 years ago. The AMD Grid

started with fewer than 50 servers in 1996 and was initially deployed on UNIX-based systems running a variety of RISC processors.

Since 2000 we have invested almost exclusively in AMD processor-based systems. Currently over 95% of our compute grid is composed of AMD processor-based systems running Linux, with a majority of these being AMD Opteron™ processor systems. This graph shows the relative growth in compute capacity for the AMD Grid, starting in 2001.



The AMD Grid now includes over 5,000 servers and 500 workstations and well over 10,000 CPUs. The Grid includes servers installed at six AMD design centers in Asia, Europe and North America. We average a total of over 40,000 compute jobs every day. The AMD Grid typically exceeds 90% utilization, while still providing low latency job dispatch for time critical work.

Our network infrastructure for the AMD Grid is based on standard switched networking technology and has evolved from 100 megabit Ethernet to gigabit Ethernet. The AMD Grid uses Platform LSF from Platform Computing and N1 Grid Engine by Sun

Microsystems for job scheduling and resources management. We run both commercial ISV application software and AMD-developed tools on the AMD Grid.

Design and Verification Jobs

Compute jobs run on the AMD Grid supporting all phases of architectural design, functional design and verification, circuit analysis, integration, electrical characterization and mask generation. The wide ranges of job types operating on the AMD Grid demonstrate the flexibility of Grid technology. We can easily allocate resources as needed for ongoing design projects because all resources are made available in a consistent fashion using Platform LSF and Sun Grid Engine.

The majority of the work done on the AMD Grid is for physical and functional verification of our microprocessors. AMD has a large set of directed regression tests and random test procedures to verify the AMD64 architecture. These tests and procedures are applied during our design process and provide critical quality assurance. The verification jobs are compute intensive and are potential bottlenecks for each stage of the design process, up to and including the functional testing of manufactured parts. Additionally the final step of mask generation has very large memory requirements (currently over 100GB) and as the last step in the design process before manufacturing, it occurs when there is no additional flexibility in the AMD product schedule.

The AMD Grid includes the following types of AMD processor-based systems:

- Large memory servers (64 to 128GB of RAM)
- High availability servers for critical and long-running jobs (configured with redundant power supplies and redundant disk arrays)
- Small to medium memory servers (2GB - 16GB of RAM)
- Engineering workstations (2GB - 16GB of RAM)

The AMD Grid allows us to mix and match the wide range of job types among the different system types. This flexibility reduces the need to purchase dedicated servers for any particular job type, and reduces overall cost. Short-term resource shortfalls for one group are resolved by administratively allocating resources where they are most useful for AMD's business needs. We can easily add additional resources into the Grid as needed, without changing the engineering workflow.

We have upgraded our Grid servers on a regular basis. The design of the AMD Grid has made these upgrades simple and painless. End users of the Grid are often

unaware of operational changes made within the Grid itself -- we provide a virtual pool of compute power in a way that hides the reality of ongoing system installations, retirements and upgrades.

AMD Opteron™ processors provide an unbeatable combination of high performance and competitive prices and operational cost. Although our compute needs continue to grow quickly, the competitive market pricing of AMD Opteron™ processors has allowed us to increase our capacity in a timely and efficient fashion. The increase in compute needs is driven by the complexity and size of our designs, the ongoing refinement of our design process, the advancement of our manufacturing process technology and by the number of concurrent designs in our product pipeline.

High Performance Desktops within the AMD Grid

The AMD Grid includes over five hundred high performance AMD-based Linux engineering workstations. We run Grid jobs on these desktop systems. The design of the AMD Grid allows us to provide excellent interactive performance for our end users while allowing us to increase system utilization with batch-oriented jobs. Given the shared nature of these systems we only dispatch well-characterized and predictably-behaved applications to run on desktops. Our architectural modeling software is a prime candidate for running on desktops. These jobs are usually limited to 500MB of RAM, and are designed to not adversely affect the interactive performance needed by our workstation users. These techniques increase the utilization of our desktop workstations. The high performance workstations provide a productive environment for our engineers and the workstations also contribute to the overall capacity of the AMD Grid.

Power Management and Performance per Watt

Modern high-density data centers require significant attention to power and heat management. It is non-trivial to provision a data center to power and hundreds or thousands of these servers. The IT industry has learned that the total cost of ownership and the total performance of a system depend on the performance per watt of that system. The aggregate cost of installing and operating large quantities of high-density and high-performance servers includes a significant power and cooling component.

The AMD Opteron™ processor includes enhanced AMD PowerNow!™ with Optimized Power Management (OPM) technology for even more sophisticated power

management. The dynamic power management available in the AMD Opteron™ processor allows - up to 75% power savings at idle.

We intend to further improve our power and heat management by configuring our AMD Opteron™ processor Grid servers to use AMD PowerNow!™ technology to reduce overall power consumption. PowerNow!™ with OPM allows the AMD Opteron™ processor to adjust voltage and clock speed within microseconds of a change in workload. Even a server that is largely cpu-bound will spend some time waiting on I/O. These I/O wait intervals are the times when PowerNow!™ with OPM can reduce power and cooling needs for the AMD-based systems.

To see the value of this optimization, imagine a data center with 2000 systems, each using .309KW/hr. You can calculate annual power costs like this, assuming a power cost of \$0.10/KW/hour.

$$2000 \text{ systems} \times \frac{.309 \text{ KW}}{\text{hour}} \times \frac{8760 \text{ hour}}{\text{year}} \times \frac{\$0.10}{\text{KW}} = \$541,368 / \text{year}$$

In this example, a 10% power reduction would save \$54,136 in annual power costs. There would also be an additional savings in cooling, given the reduced power usage. See the AMD paper on "POWER AND COOLING IN THE DATA CENTER" (http://www.amd.com/us-en/0,3715_12353,00.html) for more details.

Platform LSF and the AMD Grid

Job dispatch and grid management is done by Platform LSF (Load Sharing Facility). One significant challenge in managing a large grid is scalability – can we provision a Platform LSF-based job scheduling system that can handle tens of thousands of jobs simultaneously? High performance in job scheduling and resource management is required for low latency job dispatch and is also required for support of our sophisticated resource allocation policy. We provide very low-latency job scheduling and dispatch for jobs invoked by individual engineers. We also provide high bandwidth performance for huge job sets. We have optimized for two key workloads -- latency sensitive and bandwidth sensitive. AMD Opteron systems (dual and quad-processor) have been proven in production as capable of scheduling jobs on some of the largest Compute Grids.

AMD64 and the AMD Grid

Historically, it has been a challenge for large design environments to convert from one system vendor to another, or from one operating system to another. However, in 2003, we seamlessly extended the AMD Grid to include AMD64 systems. We leveraged our investment in 32bit X86 Linux for AMD Athlon™ systems and re-used almost all of our existing tools, programs, processes and other infrastructure. The changes required for these conversions ranged from replacing physical infrastructure to completely revamping system operations to the complete re-education of end users.

Our migration from 32bit X86 Linux to 64bit AMD64 Linux went - smoothly. Our system support structure for hardware remained essentially unchanged. Our AMD Opteron™ dual-processor systems fit in the same racks using essentially the same power, network and cooling facilities we have used for years with AMD Athlon™ systems.

While we integrated the first AMD Opteron™ processor systems into production, we were able to re-use over **99%** of the existing software we had already installed for 32-bit Linux. We support Linux and UNIX systems with a corporate toolbox of freely available software (such as the GNU C compiler, gcc and the scripting language Perl.)

Future of the AMD Compute Grid

As our design engineering compute needs have increased we have looked at the options of expanding our grid capacity both internally and with an external grid provider. We are currently working with Sun Microsystems and using AMD Opteron-based systems on the Sun Grid Compute Utility for some of our design engineering work.

Standards are essential in driving new technology adoptions such as Grid Computing. The Global Grid Forum (GGF) is established as the definitive worldwide standards body for Grid Computing. GGF along with EGA (Enterprise Grid Alliance) is providing leadership throughout the Grid community of software and system vendors and other standards bodies. AMD is a Gold member of GGF. We strive to have a strategic presence within the Grid community, both to help drive best practices and to share our insight on Grid technology. We believe that our active participation in driving these standards will lead to better technology, product and solution offerings for our Grid customers.

Conclusion

The AMD Grid is fundamental to our microprocessor design environment and workflow. Grid enablement at AMD has improved our productivity and lowered our total cost of ownership--by providing compute server utilization rates of over 90 percent and by allowing dynamic reallocation of our compute resources. The designs that AMD is working on are both large and complex. We first applied Grid Computing on microprocessor designs of fewer than 10 million transistors. Our designs now have hundreds of millions of transistors and include multiple cores per die. This increase in complexity requires a matching increase in compute resources for our design engineers. AMD can only deploy sufficient resources for our design projects by using the techniques of Grid Computing.