IDC EMEA, 389 Chiswick High Road, London, W4 4AE, U.K. Tel.+44.20.8987.7100 www.idc.com

WHITE PAPER

The Role of x86 Processors in the Virtualized, Energy-Constrained Datacenter

Sponsored by: AMD

Giorgio Nebuloni Thomas Meyer
November 2008

IDC OPINION

The x86 market has gained momentum in Europe in recent years, and thanks to the rapid adoption of 64-bit processing and the always-improving multicore systems, x86 has come to represent the majority of server activity, in unit terms as well as from a business perspective. Standardization, a more price-aggressive approach, and the greater flexibility of this type of system have driven this success, and x86 servers are now omnipresent and able to handle a variety of demanding tasks, from network to HPC and from ERP to databases.

At the same time, the proliferation of x86 servers has caused one of the major challenges. Due to the high volumes of machines shipped in recent years, end users are now turning to virtualization solutions to optimize their computing infrastructure and power management tools to keep electricity bills under control.

As a result, energy efficiency and virtualization now take center stage in the x86 server development model, turning from accessory issues to factors that modify and influence the way systems and microprocessors are designed and marketed. In particular, x86 CPU makers are taking charge of the dual responsibility of building computing engines that can contain their energy needs while at the same time supporting resource-consuming virtualized environments.

While multicore technology and the gradual shrinking that is taking place in the size of the microprocessors are prerequisite to address these problems, IDC maintains that CPU manufacturers must follow a specific path in research and development, both in their own product roadmaps and in their relationships with OEMs and channel partners.

IN THIS WHITE PAPER

This IDC White Paper analyzes the latest technology and market developments in the x86 server space, which has been significantly impacted by the spread of virtualization solutions and the rise of energy issues. The paper puts AMD's latest 45nm quad-core AMD Opteron Processor, codenamed "Shanghai," in context in this environment and analyzes its opportunities in the European market. It also shows the benefits that customers can realize from the new processor.

SITUATION OVERVIEW

In recent years, the European server market has seen many significant changes, the most relevant of which is the increasingly widespread adoption of x86-based systems throughout the region. In terms of units, IDC has seen an ever-increasing dominance of x86 units, which by the second quarter of 2008 accounted for around 95% of overall shipments in Western Europe.

This situation has been driven by the impressive growth in x86 unit shipments, at an average of 8.5% year over year (CAGR) from 2004 to 2007, and has resulted in a profoundly modified market landscape. According to IDC server installed base data, 6.8 million (91%) of the 7.5 million servers installed and active in Western Europe were equipped with x86 CPUs. More importantly, IDC estimates that the volume of business linked to x86 machines outstrips that of any other server type, as the asset value of x86-based servers topped €10 billion for the first time in 2006, overtaking RISC-based systems (see Figure 1).

FIGURE 1

Western Europe — All Servers — Asset Value Growth by CPU Type, 2003–2007

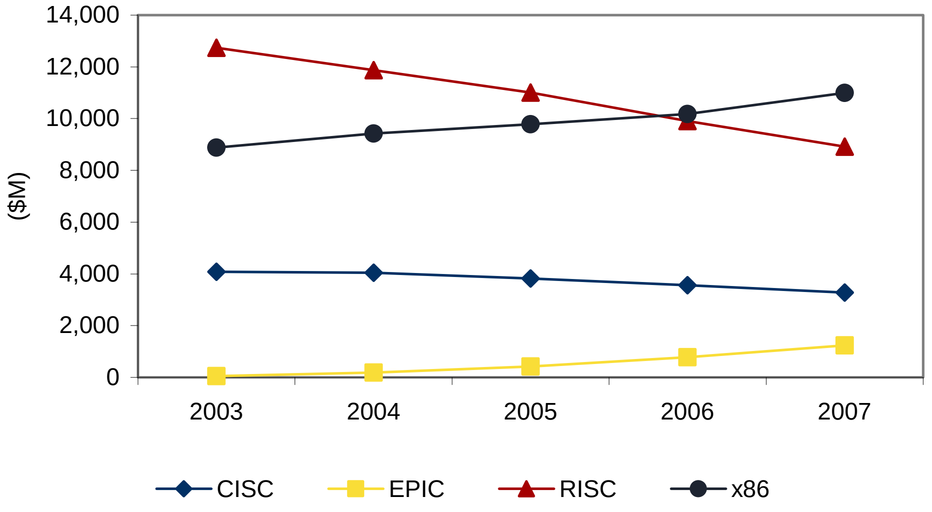

Source: IDC, 2008

Drivers of x86 Proliferation

The main reasons for the major success of x86 platforms are the boost in performance shown by x86 processors in the last five years and the fact that these systems have been marketed at more and more accessible prices, thus becoming viable solutions both for small budget customers and mid to large enterprises with increasingly extensive needs. Along with lower prices, users can benefit from the more open and dynamic architecture that comes with an equally dynamic supplier context, where competition happens at all levels (OS, CPU, server system, applications, etc.).

The adoption of now-standard client-server architecture in midsize enterprises, and subsequently in SMBs as well, produced the first uptake of x86 servers in the European market. Systems based on x86 architecture were originally deployed to carry out low-end operations and acted mainly as a connection node between the clients distributed in the enterprise.

Around three years ago, x86 server performance was significantly improved by the combined introduction of 64-bit and multicore technologies, which have now become standard in the European server market. x86 64-bit processing laid the basis for systems able to cope with much higher needs of processing power. For example, starting with the first AMD 64-bit processors, they quickly found resonance in the server market, benefiting from the easy upgradeability and the protection of the investment they granted.

The increase in physical and virtual memory address space ensured by x86-64 CPUs was paralleled by the boom in multicore processors. The percentage of dual-core machines sold in Western Europe was only 14% in 2006 but grew to more than 60% in 2007. In the first half of 2008, the vast majority of x86 machines shipped were equipped with either dual-core or quad-core processors, with just 20% still shipping with single-core CPUs.

The scalability of systems is an important factor in the development of the x86 landscape and in particular the extension to all areas of the datacenter. Recently, vendor focus has been on facilitating both the scale out and scale up of x86 systems, and the market has seen a variety of systems targeted at these architectures. Advances in the relationship between cores and memory bandwidth and I/O availability are examples of this.

As the computing power of x86 machines grew, both server manufacturers and customers began to employ x86-based systems for operations that went beyond the usually associated tasks (file sharing, messaging, printing). x86 stacks today are able to handle critical applications such as ERP, databases, or BI systems, and with the introduction of blade servers, x86-based machines are now also successfully employed to perform high-performance computing (HPC) tasks in financial and R&D environments.

This means that x86 platforms now represent a viable alternative to midrange or high-end systems in a number of contexts, as confirmed by recent IDC research in Europe. According to the *IDC European Datacenter Power and Cooling Survey* conducted in 2008 focusing on 102 large and very large companies in the U.K. and Germany, x86 rack servers are the preferred server platform in 73% of the smaller datacenters (up to 499 square meters), 60% of midsize datacenters (500–999 square meters), and 47% of large datacenters (1,000+ square meters).

Disruptive Factors in the x86 Space

As mentioned above, in the general frame of the x86 story, two aspects have been gaining increasing importance in the last two years, changing from secondary concerns into critical factors that are reshaping the server and semiconductor industries: virtualization and energy issues.

Virtualization

Until two years ago, the expansion of x86 machines in the European installed base had coincided with an approach based on a one-application-per-server ratio, often accompanied by server over-provisioning in order to guarantee resource availability in case of peak usage. This approach, combined with growth in computing power that has not always been matched by the rise in demand of the software stack, has resulted in a mass of under-utilized machines, draining energy and management resources from the IT department. IDC calculates that the number of x86 shipments in Western Europe grew by more than 25% from 2004 to 2007.

x86 server virtualization is based on tackling this situation, allowing IT managers to consolidate the workload of several physical machines on one system, for example, leveraging the abstraction layer provided by virtualization software. On the wave of server consolidation, virtualization solutions from VMware, Citrix, and Microsoft have been conquering their own space in the market, generating hundreds of millions in revenue in the European markets alone.

According to recent research from IDC, server virtualization is being adopted at an ever-faster pace by European organizations and is continuing on the hyper-growth path. While the percentage of installed servers virtualized in Europe until 2007 was lower than 10%, IDC's 2008 European Virtualization Survey indicates a very strong increase among large and medium enterprises: European organizations that reported having adopted server virtualization technology revealed that around 35% of the machines purchased in 2007 were virtualized, with the number expected to grow to 52% when considering those bought in 2008.

The latest IDC studies all highlight that virtualization usage patterns are changing quickly, as users become acquainted with the technology and the software platforms become more powerful and stable. In addition to consolidation and test and development, virtualization solutions are increasingly used to tackle disaster recovery and high availability (HA) issues. In 2007, less than 10% of European companies regarded availability of applications and systems as one of the three main reasons to deploy virtualization, but in 2008 the percentage has grown to almost 30% (according to the 2007 and 2008 surveys). The possibility of replicating or migrating entire virtual environments with little or no service disruption is redefining the way HA infrastructures are designed and built. Of course, the underlying hardware platforms have to be able to support this, in terms of both performance and compatibility.

In particular, x86 hardware is subject to two main strains coming from an extensive deployment of virtualized environments:

- **Higher standard compliance** — x86 servers have to be more homogenous and standard-adherent than before, especially in their CPU components, because they have to stack up as a uniformed pool of resources on which virtual machines can be moved smoothly, without service disruption.

- **Greater reliability** — Though quite self-explanatory, a physical server operating five virtual machines is five times more critical than a traditional one-OS server; this is often an afterthought. As its criticality grows along with the increased number and importance of workloads it has to support, an x86 server needs to be increasingly reliable. This last point is particularly important, as, according to IDC, almost half of the Western European enterprises adopting virtualization are doing so on legacy servers (see Figure 2).

FIGURE 2

Western Europe — Hardware Reuse in Virtualized Environments, 2008

When virtualizing a server do you usually:

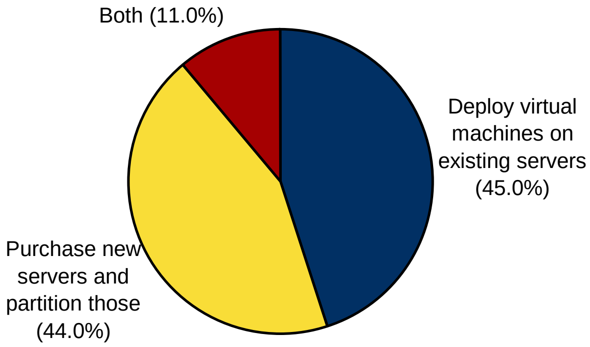

Source: IDC Virtualization Survey, 2008

In fact, what few still recognize is that virtualization is bound to change not only the utilization rate of x86 servers, but also and more fundamentally the *way* these servers are utilized. This in turn results in considerable changes in the way machines and CPUs have to be designed and built. While server manufacturers have only recently begun to propose regular portfolio products aimed at easing the handling of virtualized environments (e.g., increasing the I/O bandwidth and the RAM capabilities), CPU makers have started to face the issue much earlier. With its native multithreading approach, the multicore technology presents itself as a natural hardware platform to run several virtualized environments, splitting the workloads between the various cores.

This is not enough, however, and for the last two years x86 processor manufacturers have been working to optimize their platforms to cope with heavy virtualization loads, "assisting" the hypervisors and reducing their natural performance overheads. The importance that virtualization has acquired for the server industry is such that specialized benchmarks have been or are about to be issued to evaluate platform performance in virtualized environments, such as VMware's VMark or the upcoming SPEC benchmark.

Energy Issues in Datacenters

Recent IDC studies estimate that approximately 16.3TWh of electricity was consumed by servers in 16 European countries in 2007, 13% more than in 2006 (see Figure 3). This number has been corroborated by end-user surveys in the region. IT managers of large and very large companies in Germany and the U.K. maintain that their datacenter power consumption has been growing by 23% to 34% in 2007 compared to 2006.

FIGURE 3

Western Europe Server Electricity Consumption and Expenses, 2006–2007

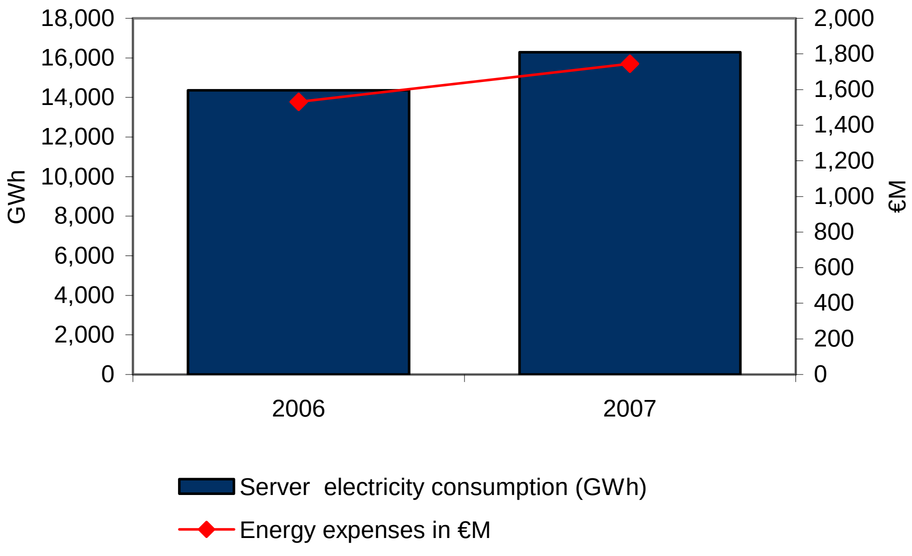

Source: IDC Energy footprint of WE Server IB, 2008

Such a rise is determined mainly by the expanding number of installed servers in the region, increased cooling needs, and the growth of average power needed by every single server due to increased hardware features. Poorly designed facilities multiply the effect of energy consumption at server level, making it a primary goal for end users to rearrange their datacenters.

With energy consumption figures growing at that pace, the energy woes of datacenters are rapidly turning into financial liabilities for organizations. In 2007, around €4.4 billion was needed to power datacenters in Western Europe, of which €1.6 billion was exclusively for servers. If the trend continues without major improvements, IDC believes the cost of powering datacenters in the region could top €7.5 billion by 2012, which would be equal to 80% of the amount of money spent to buy new servers.

IDC maintains that most of the burden of improving IT energy efficiency lies on organizations such as the Green Grid (thegreengrid.com). IDC advises customers to undertake in-depth TCO calculations that comprise energy consumption levels before acquiring new hardware, but also to reassess their organizational functions (e.g., improving communications between IT and facility managers).

That said, the energy issue is also a technology problem, and x86 machines, covering a large part of the server installed base, play an important role in this picture. As denser, more powerful, and more numerous x86 machines hit the market, it becomes increasingly crucial for manufacturers to keep the energy footprint of each of those systems under control.

In this context, consumption at a CPU level impacts significantly on the values at system and rack level in a chain effect. The higher the CPU power envelope is the higher the cooling requirements in the servers, which in turn increases the energy consumption at server and rack level. This is the reason for server and CPU manufacturers building products that can contain energy requirements while at the same time aiming to deliver performance increases. This is why new metrics to measure hardware products, such as performance per watt, have been taking root in the marketplace, which is directing increasing attention towards energy-related issues.

AMD Multicore Path

AMD has been active in the multicore space since 2005, when it started shipping its first dual-core 64-bit Opteron CPUs for servers, the first product of this type to market. Over time, the AMD Opteron Processor family has been enhanced in a process that had focused on performance while maintaining a stable power envelope (usually by adding cores, increasing clock-speed, or improving the micro-architecture and memory bandwidth), shrinking the size of the microprocessors, and maintaining a stable socket infrastructure. This minimized the impact of training, certification, and new software images required by new platforms.

However, stiff competition in the processor environment, including very aggressive schedules, clearly has had some impact along the way. For example, AMD postponed the launch of its first 65nm quad-core Opteron CPU, codenamed "Barcelona," in 2007 due to a problem that was discovered very late. Wisely, AMD decided not to ship the initial processors and pulled them off production to resolve the issue. Shipments of the new processors reached full volume some months later than planned, in 2Q08, finding the top tier OEMs (Dell, Fujitsu Siemens, HP, IBM, and Sun Microsystems) ready to integrate the new processor and hit the market with Barcelona-based products of each form factor: blades, racks, and tower servers.

The delays in shipping Barcelona coincided with the work to integrate the $5.4 billion acquisition of graphics maker ATI. As a result, the company significantly reviewed both its internal engineering structure and the business processes that accompany any CPU launch. With this approach, AMD focused on the tangibles with the priority of keeping its processor roadmap on track, well aware that in the CPU marketplace, the timing in shipping and promoting a new generation of processors is essential.

From an organizational point of view, AMD applied the lessons learned for the development of the subsequent CPU, codenamed "Shanghai." The company determined to give the ownership and responsibility to one single person for both the evaluation and testing legs of the new processor. Raghuram Tupuri, lead engineer on the Shanghai project, established a cross-functional team composed of members of all the different engineering groups, managing to keep tight to the deadline and optimizing the evaluation phases.

The modified approach to the OEM validation process proved to be just as pivotal in keeping true to the roadmap. For the first time, AMD undertook the OEM validation process right from the beginning, whereas in previous projects samples were sent to partners only after an internal probe cycle. According to the company, this also helped to reestablish the trust of hardware vendors and ISVs.

Quad-Core AMD Opteron Processor Codenamed "Shanghai"

As planned, AMD started shipping its new Shanghai processor in November 2008. Shanghai represents the 45nm shrink of Barcelona and exhibits significant advantages compared to the predecessor. According to AMD, Shanghai's improvements will enable up to 35% computing performance gains and up to 20% overall power savings.

The first processors launched in 4Q08 will be in the 75W range, with versions running in higher and lower power envelopes starting to ship in 1Q09. Compared to Barcelona, at 65nm, one of Shanghai's main differences is the smaller 45nm build that allows the significant power savings. In addition, Shanghai will offer an extended L3 cache memory of 6MB, twice as large as Barcelona, and the new HyperTransport technology version 3.0 that will directly connect cores, caches, and I/O subsystems at a maximum speed of 17.6GB/sec. Like its predecessor, Shanghai is built on AMD's Direct Connect Architecture, linking the CPU directly with the memory DIMMs and allowing a faster I/O throughput.

As with Barcelona, Shanghai will be compatible with the AMD Socket F platform, supporting in-place upgrades in a wide range of Opteron-based servers. AMD expects the largest portion of the server product shifts towards Shanghai to happen in the first months of 2009, as more versions of the CPU are validated and released, including a one-socket version compatible with smaller server systems. AMD plans to gradually dedicate its Barcelona processor to low-end servers, while covering more sophisticated and demanding systems with the new Shanghai CPU. As of November 2008, all major OEM vendors have completed the validation process for the new CPU platform and have started shipping Shanghai-based server products.

FIGURE 4

Shanghai Architecture

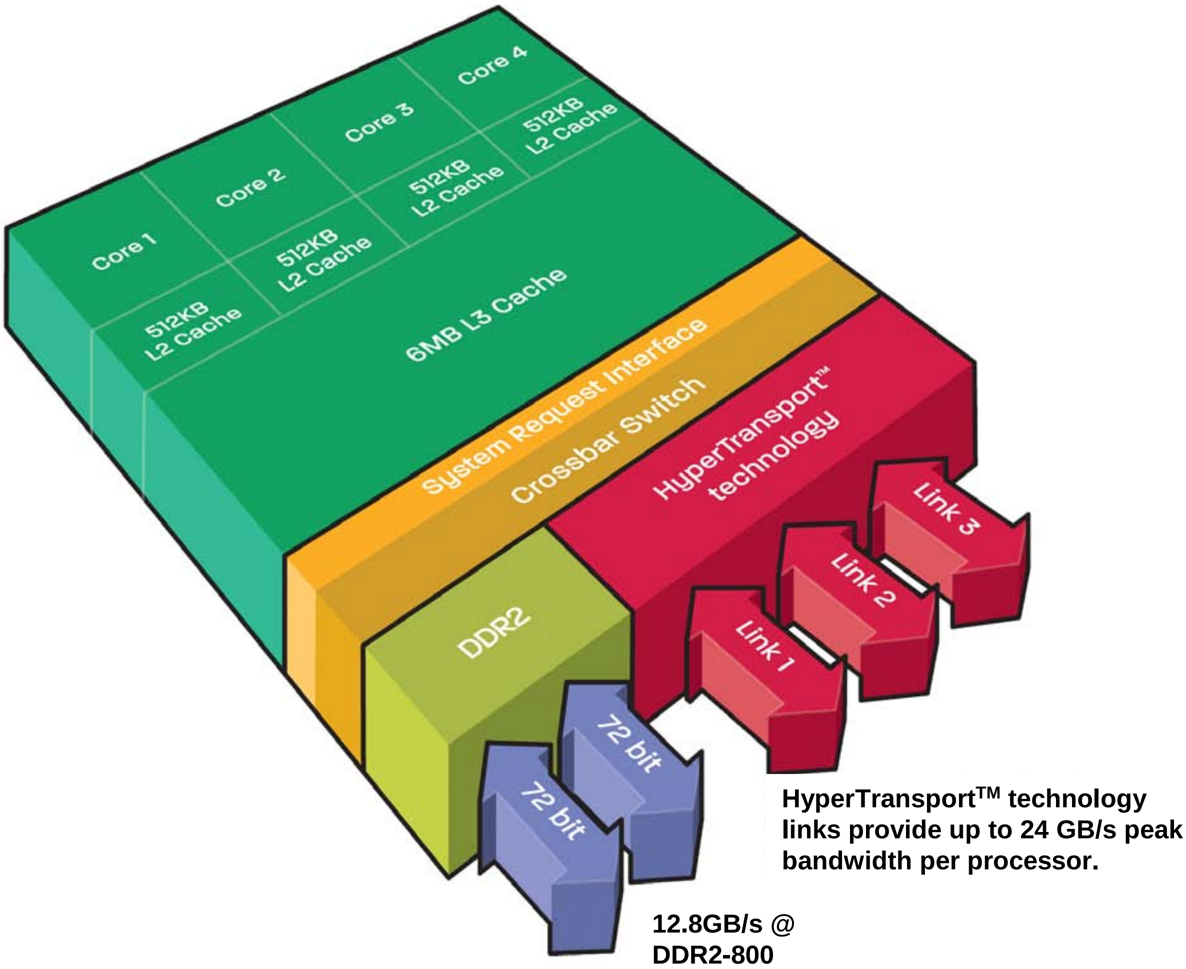

Source: AMD, 2008

As Shanghai hits the streets, AMD engineering teams are already working on future server processor products aimed at tackling future issues of the computing infrastructure.

- The company plans to launch "Fiorano," the next generation of its Socket F platform designed for Shanghai and the future CPU generations, in 2009. Istanbul, the 6-core 45nm evolution of Shanghai, is also planned to ship in 2009, and will include advanced management tools and an enhanced architecture to improve data transmission.

- AMD also has new server processor products on the roadmap with increased core counts, to be released in 2010 with a brand-new socket platform codenamed Maranello. This platform will offer enhanced virtualization capabilities and support DDR3 memory.

Responding to the Challenge of Energy Efficiency

Following a now established approach, AMD continued to intensively focus on the energy efficiency aspect of its 45nm quad-core processor. The cornerstone of the improved performance-per-watt ratio is the reduced architecture, which alone provides significant savings compared to the larger 65nm build. On top of that, AMD continues to use low-power DDR2 memory technology and, more importantly, it associates the CPU hardware with an extended internal power management system branded CoolCore, whose main aim is to switch off all the unused parts of the CPU (e.g., entire cores, floating point units [FPUs], individual gates) saving further on energy. The advantage of this technology, which is embedded in all of AMD's new server CPUs, is that it does not require any action from the user, does not need any software, and has no impact on performance but reduces energy consumption.

Along with CoolCore technology, AMD includes further power management tools for both the cores and the memory, such as Dual Dynamic Power Management, which is used by AMD PowerNow technology. For some time, AMD has been working with OEM partners to integrate the functionality of such energy-efficiency tools with common management suites (e.g., HP System Insight Manager, Dell OpenManage, IBM Tivoli) that therefore gain greater visibility and operational space at CPU level. IDC believes this to be a necessary step, as end users enter the phase of energy awareness and decide to clamp down on power waste in the datacenter.

Shanghai in Virtualized Environments

AMD has customarily been able to make the most of its Direct Connect Architecture in virtual environments where I/O performance is pivotal. The company is working on different areas to extend the impact of its processors on virtualized servers. One of the latest features of the AMD V virtualization set is the Rapid Virtualization Indexing technology. This allows the processor to enhance its performance when handling applications that require more direct access to the VM, reducing memory overheads originally associated with the translation of instruction codes from physical to virtual. Between the sample workloads that can benefit from RVI are database workloads, environments running application virtualization engines, and Web servers.

AMD is developing I/O virtualization features aimed at bypassing some of the performance and scalability limits of the current technology via Direct Device Assignment, a technology that is able to slash I/O overheads by mapping the physical Network Interface Cards directly to a virtual machine. Direct Device Assignment will enable VMDirectPath and will be available in the Fiorano Socket platform in 2009.

As discussed earlier in this paper, virtualization increasingly becomes an instrument to turn servers into a common, rapidly accessible pool of resources, making compatibility between different systems a primary issue. AMD's attempt to tackle this problem is labeled Extended Migration. AMD Extended Migration allows a VMware-based virtualized environment to enable Vmotion (the feature that makes it possible to execute live migration of virtual machines across the hardware pool) across different generations of Opteron processors, whether single, dual, or quad-core. In this way, customers can deploy large virtualized environments on their AMD-enabled x86 infrastructure without worrying about compatibility issues for VM live migration. The processor families supported go back to the three-year-old series such as Rev E and Rev F, and move on to Barcelona, Shanghai, and all the upcoming Opteron processors.

Building out the solutions portfolio to address efficient and optimized virtualization, AMD is also partnering with Microsoft, which it has supported for the launch of the Hyper-V virtualization technology for Windows Server 2008, and works closely with Citrix for application and server virtualization. IDC expects further announcements around optimized solutions in the near future.

CHALLENGES AND OPPORTUNITIES

With the introduction of Shanghai in the portfolio, AMD is back on track in CPU technology development. IDC believes that Shanghai could help AMD defend its stronghold in the high end of the x86 market and support its presence in the blade server landscape. At the same time, the broader portfolio allows AMD to undertake the further crucial challenge of shifting Barcelona products towards the large portion of the x86 market represented by two-socket 2U machines, where the company has traditionally presented a less effective front.

AMD's quad-core products are of course not without competition, and the pressure for tight schedules, performance increases, and broader functions for the CPU remain. AMD also faces a very significant installed base challenge. That said, IDC believes AMD has major opportunities in specific areas that play to its key strengths. This includes areas such as high-performance computing and virtualized infrastructure environments, while in the long run stepping up its presence in the large volume two-socket market will become key. Leveraging Shanghai's capabilities, AMD also intends to better address Java environments, a marketplace where it has not had a significant presence.

In an ecosystem where refreshment cycles and timing turn out to be as crucial as the quality of the product itself, the on-schedule launch of Shanghai is at least as important as the addition of new technology features in terms of image and credibility. In 2008, the company finds itself at an inflection point, mirrored by organization leadership changes, by the recently revealed spin-off of the manufacturing branch, and by the improvement of financial results announced for the last quarter.

IDC believes that AMD's future success and competitive prowess will rely heavily on its ability to keep delivering products on schedule and to maintain tight relationships with hardware and software partners, focusing on execution and process optimization. At the same time, AMD needs to stay close to the user side and deliver value for the pain points customers are trying to address. Products need to remain relevant and integrated, as Shanghai is promising with its focus on energy efficiency, virtualization, and workload optimization.