IDC
Analyze the Future

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

TECHNOLOGY ASSESSMENT

Multicore Processing Scenarios, 2005–2009: Disrupting the IT Market in Three Generations?

Vernon Turner Matthew Eastwood

IDC OPINION

Multicore processing, and the technology enhancements that multicore processing will enable across the IT industry, could be one of the most significant developments in the past 40 years. Consequently, IDC believes the following:

- Multicore processing could create highly disruptive environments across the enterprise landscape and accelerate architectural convergence between enterprise and consumer technologies.

- Incumbents in microprocessors, operating platforms, systems management, security, application development tools, network management, and services delivery will be challenged by start-ups, new industry partnerships, and customer requirements that will push their customer delivery mechanisms to the limit and customer experiences to new expectations.

- Multicore processing will be a cornerstone of transforming the arthritic IT structures of today into the dynamic IT of tomorrow. To that end, as each generation of multicore processors is deployed, customers will demand more agility, better utilization, and higher efficiency from their IT infrastructure. In fact, end users are telling IDC that they expect to realize significant price/performance improvements from these systems and they fully expect that multicore systems will further accelerate the adoption of server virtualization technologies. (Please note that this document builds on IDC's opinion published in *Multicore Processing: Disruption or Distraction for the IT Infrastructure?,* IDC #32278, November 2004, which describes fundamental concepts, reasoning, and challenges of multicore processing.)

Filing Information: July 2005, IDC #33789, Volume: 1, Tab: Markets

IN THIS STUDY

This study covers the impact that multicore processors and their processing abilities will have on a broad IT ecosystem as new generations of server technology are rolled out. IDC feels that the impact on the IT infrastructure will accelerate with each generation of multicores that comes to market. For that reason, we expect that the IT industry will experience significant disruption by the third generation of multicore technology in the 2009–2010 time frame. The approach to understanding the market impact will be based on adoption scenarios because the technology has no prior track record of success or failure. This document describes scenarios created by the first three generations of multicore processing; however, it does not attempt to generate new server market forecasts as a result of these scenarios.

SITUATION OVERVIEW

What Is Multicore Processing, and What Is a "Thread"?

In its first iteration of widespread multicore processing, the industry will see most products that are dual core in design. A dual-core processor integrates two separate logical processing units on a single chip. A dual-core processor can often outperform a single-core processor on most multithreaded applications while running at lower clock speeds and consuming less power.

A thread is a single stream of instructions that the operating system sends to the processor for execution. Most current processors execute a single thread or, with some high-end models, two threads at once. However, in a multicore environment, the capacity to process much more work becomes significantly greater as more and more cores are addressable to the system. Because of the architecture of the dual- and multicore chips, discrete threads can be allocated to the distinct cores within the processor, thereby increasing speed and reducing latency in data and transaction processing. Some systems vendors and silicon suppliers have made announcements that show there will be immediate market segmentation based on the number of threads to cores. This change will lead to immediate product and market differentiation and positioning.

Multicore processing systems will create new, challenging issues for IT suppliers and users alike. IDC has identified many key areas that will be impacted by the rollout and adoption of this technology. Although some observers feel that multicore processing spells an end to Moore's law, whereby processor performance was based on improved circuit density and speed, IDC believes that multicore processors and their compute capability will provide a potentially potent one-two punch to commodity-based systems over established proprietary systems often found in midrange and high-end systems. Consequently, we believe that there will be significant market dynamics and changes to customer behaviors that will alter the IT landscape by the end of the decade.

However, to gain the full benefit of multicore processing, IDC believes that the product offerings need to have the following:

- Support for multiple operating systems
- Substantial bus bandwidth to and from main memory, as well as companion chips
- A flexible, on-chip input/output (I/O) interface
- A real-time resource management system for real-time applications
- Embedded on-chip hardware support for security systems, which offer intellectual property protection

This research document acknowledges that there will be a parallel market impact through consumer markets, and discusses these briefly. However, the focus of this study is enterprise infrastructure.

Current Announcements

2005 has been the watershed year for the multicore processing industry. Table 1 shows the current market offerings and highlights key features (number of cores, threads, coherent scalability). While IBM has been shipping multicore-based RISC-based platforms for four years, it is only now that there is widespread choice for the IT industry. The offerings summarized in the table are described in additional detail below:

- Azul Systems of Mountain View, California, officially launched the Azul Compute Appliance 1920, which has 192 processor cores and 64GB of memory, and the Azul Compute Appliance 3840, which has 384 processor cores and 128GB or 256GB of memory. The Vega Multicore processor at the heart of the Azul Compute Appliance is based on 64-bit multicore technology. Vega is designed with 24 cores, each equipped with a single thread, and it is optimized for Java virtual machine execution. Systems include multiple chips, each with 24 cores (providing 4x, 8x, and 16x multiples of the 24-core design).
- AMD got a jump on the volume x86 server market and has been first to market with processors for x86 volume systems. One of AMD's initial dual-core features is that discrete threads can be allocated to the distinct cores within the processor, thereby increasing speed and reducing latency in data and transaction processing.
- Intel's dual-core Montecito processor, which was announced to the market in September 2004 at the Intel Developer's Forum, is expected to ship in volume later this year. As to be expected, Intel has an opportunity to leverage its installed base of more than 90% of the total x86 market as a target for this new processor. IDC expects that this transition to dual-core systems will occur rapidly, including a significant majority of all Intel x86 systems shipped by the end of 2006.

- IBM, Sony, and Toshiba Corp. unveiled a multicore offering, code-named Cell, which they are jointly developing for next-generation computing applications and digital consumer electronics. Cell is a multicore chip comprising a 64-bit Power processor core and multiple processor cores capable of floating point processing. Cell is optimized for compute-intensive workloads and broadband rich media applications, including computer entertainment, movies, and other forms of digital content.

- Sun Microsystems announced that Niagara's multiple cores and multithreading capabilities are expected to perform at lower clock speeds in a smaller power envelope than is typically found in the server platforms where Sun traditionally competes. Niagara will be able to execute 32 threads at once. Each chip will have eight cores, each capable of executing four threads simultaneously.

TABLE 1

Multicore Industry Map: 2005

	Workload Type	No. of Coherent Cores/Chip	Multichip Coherent Scalability	Max. No. of Coherent Cores/System	System Vendor
Azul Vega	Application tier	24	Yes	384	Azul
Intel x86 Dempsey	General purpose	2	Yes	16	All major vendors
Intel IA-64 Montecito	General purpose	2	Yes	128	HP, Fujitsu
AMD x86 Opteron	General purpose	2	Yes	16	Most major vendors (not Dell)
Sun US-IV	General purpose	2	Yes	144	Sun, Fujitsu
Sun Niagara	Web tier	8*	No	8	Sun
IBM Power	General purpose	2**	Yes	64	IBM
IBM Cell	Gaming/media	1***	No	1	IBM, Sony
ClearSpeed	Embedded	1****	No	NA	Boards from ClearSpeed
Broadcom SiByte	Networking	4	Yes	16	OEMs

* 4 threads/core

** 2 threads/core

*** 9 cores not coherent

**** 96 ALUs

Source: IDC, Azul Systems, and IDG Ventures, 2005

FUTURE OUTLOOK

The future outlook is described by the impact of each generation of multicore processors as seen through various adoption scenarios. Consequently, no market forecast for unit shipments or revenue appears in this document, but the market impacts are clearly highlighted.

First Generation of Multicore Processing: 2005–2006

The first generation of multicore processors has been launched or has been given market availability. Although IBM (the first to market with a RISC-based platform in 2001), HP, and Sun have been shipping dual-core RISC processors for some time, AMD, Azul Systems, and Intel are also laying important groundwork for the rest of the industry. AMD and Intel are tackling the market by addressing a dual-core platform aimed at the 2-way and 4-way market space, while Azul is approaching the J2EE Virtual Machine world with its Network Attached Processing technology. Although the go-to-market strategies for all of these vendors vary significantly, it is clear that success will depend on a simple, nondisruptive implementation model. IDC expects some or all of the following to play out in this phase:

The first generation of multicore processors has been launched or has been given market availability. AMD, Azul Systems, and Intel are laying the groundwork for the rest of the industry.

- There will be little overall market reaction or impact to the first-generation offerings as customers, independent software vendors (ISVs), OEMs, and channel partners evaluate the offerings and look for early end-user adopters.

- Marketing position for these offerings will be around nonintrusive deployment, significant power-savings messages, and other server consolidation and simplification initiatives.

- Systems management capabilities won't be stretched by these new systems — in other words, there wont be a rush of new offerings. Some workload migration from mid- to high-end systems will occur, but not at blinding speed. Consequently, vendors such as BMC, CA, IBM, and Mercury will be evaluating their software licensing strategies.

- Dual-core systems will provide added compute capacity, and it will be important for systems vendors to ensure that ongoing systems utilization levels don't fall. As part of the justification of dual-cores systems, server and workload consolidation initiatives will remain important. To that end, partitioned servers supported by products such as VMware will remain strong, driven in part by VMware's socket-based product licensing strategy. Additionally, systems vendors and channel partners will expand added-value services to implement and support the buildout of a virtual server consolidation environment.

- The blades market will increasingly be incorporated into virtual server consolidation environments and serve as a launching pad for multicore processing in the volume server market. The power and heat advantages of dual-core systems will be particularly attractive for blade suppliers. Customer education programs that explain the benefits of virtualization and dual-core processing in blades takes will become critical.

- Customers will be looking to see how systems and software vendors plan to implement dual-core platforms in order to understand if programming models are maintained.

- Software licensing will not be an issue for the early adopters because the software vendors have largely confirmed ongoing licensing policies. Microsoft, however, will receive kudos from the end-user community for stating that it will continue to license its products based on the number of sockets. IDC expects that database and enterprise application vendors such as Oracle and SAP will continue to maintain a wait-and-see approach in order to fully understand market impact. IBM has taken the middle ground by announcing a hybrid licensing model that prices software for x86 based on sockets but maintains a cores-based approach for RISC systems. IDC also believes that ISVs will increasingly review their software license prices based on the number of sockets being physically or logically partitioned.

- This phase of dual-core offerings will remain an IT operations initiative, although vendors will quickly need to raise the overall positioning to the lines of business and to the CIO/CFO to educate them on the long-term infrastructure impact.

- Dual-core systems and subsequent multicore systems will require faster/larger memory configurations with larger I/O bandwidth capabilities. In this initial phase, vendors will be concerned about increasing pricing for these configuration changes without being able to demonstrate or articulate improved business value for these systems. Consequently, many vendors will meet this requirement by using older I/O and memory offerings rather than increasing the price of the platforms. To the well-informed IT buyer, this change will mark the beginning of being able to procure tremendous value–priced systems offering a lot of compute cycles without major price increases because vendors have yet to establish the overall value proposition.

Second Generation Of Multicore Processing: 2006–2007

The second generation of multicore processing won't be witnessed by the arrival of a specific product, but rather a series of events that will have much more of an impact on IT infrastructure. From a product standpoint, IDC expects the following market conditions to appear:

The second generation of multicore processing won't be witnessed by the arrival of a specific product, but rather a series of events that will have much more of an impact on IT infrastructure.

- Four-way core or quad-core platforms for the x86 processors will be delivered (probably by AMD first, and quickly followed by Intel). Dell and HP will challenge each other with platform offerings based on their market appeal and not necessarily technology differentiation, whereas IBM will attempt to create a gap based on its X3 Summit chipset as it goes into multicore mode.

- Sun Microsystems and IBM will ship RISC-based platforms that will raise multicore deployment into the middle tier and high-end enterprises. Although IBM has been shipping multicore platforms for some time (in both the pSeries and zSeries platforms), this phase in the technology adoption will force IBM to further explore the need for SMP-based platform strategy as designed and built today.

- In general, smaller profit pools will be available to suppliers in the market as workloads typically deployed on large-scale SMPs begin to move to other, generally lower-margin, systems. As a result, server vendors will struggle to maintain the appropriate R&D mix as good-enough technologies permeate a wider spectrum of the market. Suppliers must be careful to balance the needs of their most leading-edge customers with the needs of the market as a whole. IDC believes the gap between these two types of customers will only grow.

IDC expects products from start-up companies such as Azul Systems will now be familiar to the ISV community and OEMs with a visible channel distribution phase as they attempt to expand their installed base beyond early adopters and gain widespread market acceptance. IDC also expects that there will be several new entrants (venture capital funded) that will attempt to further segment the markets using their performance-based or application-based benefits. However, the good news about all of the above product offerings is that CIOs and IT channels will now be well versed in multicore processing technology and the benefits afforded to IT infrastructures. IDC believes that this will create an emerging services implementation opportunity for channel partners and integrators, and therefore bring more attention to the market space.

In the second phase of multicore processing, several market dynamics will begin to emerge. Up to this point, software licensing issues by the major systems and ISV providers have been largely ignored. (Early positioning by major players such as Microsoft smoothed the initial offering by simply pricing the platform licensing based on sockets.) However, if the software industry waits until the second phase to address licensing, it runs the risk of trying to figure out whether it is missing revenue opportunities as 1-socket systems begin to look and behave like today's 4-way servers and 2-socket systems begin to look and behave like a traditional 8-way server. At this point in the adoption curve, it is clear that workload shifts will occur and force the software licensing issue as vendors try to evolve their licensing models. However, IDC believes that the industry will remain fragmented for the duration of this phase, with no consistent pricing model.

At this point in the adoption curve, it is clear that workload shifts will occur and force the software licensing issue as vendors try to evolve their licensing models. However, IDC believes that the industry will remain fragmented for the duration of this phase, with no consistent pricing model.

In this phase, IDC expects end users to be familiar with the more complex price/performance curves associated with multicore technology. In one of the scenarios that IDC expects to see, the following will occur:

- Early IT adopters will experiment with moving workloads off 2-way and 4-way systems onto 1-socket servers with four processing cores. Although we don't expect the results to be initially attractive, we do think that the learning experience will accelerate the 2-way server market to run 4-way and 8-way workloads. Here, the software industry will begin to take multicore processing seriously and wonder if indeed its licensing models are appropriate for the marketplace. Debates between software vendors and IT procurement/finance departments will be in full swing as both jockey for the best position that suits their arguments. The overall discussion will be centered around placing business value on the capability that multicore processing can bring.

- However, despite workload migration evaluations that the second phase will bring, it is important to realize that by the end of this phase IDC expects about 95% of all servers shipped will be dual-core or multicore enabled. Although much of the software in the installed base will be largely unable to exploit the underlying technology, hardware vendors will have to begin to ship richer configurations to support the expected increase in virtualization workloads and the associated increases in systems utilizations. Consequently, the challenge for vendors will be to know how much (if any) to increase their pricing and how to best pass on these increases to their customers. Will this slow the decline in hardware pricing after years of double-digit annual percentage declines? Will customers buy into the marketing reasons for the pricing increases? Will the hardware prices be matched by software vendors' cashing in? Will this combination create a procurement headache for end users and stall the whole market adoption process? Is this market driven by a very strong price elasticity model, and if so, is there a potential for margins to be squeezed as many vendors look for alternatives to make up for revenue losses?

- Systems management capabilities for volume servers will become even more critical because there will be an expectation that servers will run more varied, more mission-critical workloads and applications. There will be an expectation that the server will be able to run at a higher overall utilization within acceptable performance levels. To that end, there will be increased opportunity for products offered by VMware, XEN, and Microsoft and other virtualization and partitioning capabilities available in the market. Additionally, good provisioning management of these additional compute resources will begin to be another mandatory or minimal offering.

- The adoption of bladed computing will receive a shot in the arm from the second generation of multicore processing. However, vendors that have amassed their own provisioning, monitoring, and virtualization capabilities will be able to sustain wider profit margins through these systems management capabilities rather than through the hardware alone. IDC expects server vendors to struggle more as time passes in a multicore environment if they cannot offer up their own systems management stack. Although plenty of industry-standard offerings will be available, it will simply mean a transfer of profit to someone else in a technology ecosystem as the market experiences a limited profit pool potential overall. Consequently, bladed computing offerings in a multicore environment will be a significant differentiator between the "haves" and the "have nots."

- In the second phase, the storage industry will start to use dual-core and multicore processors in disk storage subsystems to allow the building of both larger-capacity systems and more highly managed systems that are rich in virtualization and systems management capabilities.

The Dawn of the Disruptive Technology Begins to Be Felt by the Industry

The second phase of multicore product offerings will be marked by the disruptive challenge that the products bring to today's scale-out and scale-up offerings. Here lies the true business value for the industry as a whole. It is here that multicore

The second phase of multicore product offerings will be marked by the disruptive challenge that the products bring to today's scale-out and scale-up offerings.

processing can be sustaining as scale-out clustering moves into new areas, and at the same time it can be disruptive to the market as improved virtualization drives better utilization and, consequently, IT consolidation. Equally important is the disruptive beginning of multicore processing to the traditional scale-up (or SMP) compute model. Here, IDC believes that the midrange server platforms of today will begin to seriously challenge the viability of comparable high-end systems. Initially, the competition will appear in similar markets. For example, we could expect to see today's 16-way servers begin to challenge the 32-way server market for their workloads. However, IDC expects to see technology developed for high-performance computing markets transcend into the enterprise market as very powerful and dense compute systems are built on industry-standard platform offerings.

The challenge, of course, will be to ensure that the "abilities" (availability, scalability, and so on) are maintained, but the future SMP-based platforms should be considerably less expensive and capable of doing nearly as much work. This future compute model will be highly disruptive to many players in the solution stack — the hardware platform, systems management, operating systems, database, and application providers. Another key angle to market disruption will be the window of opportunity for a supplier such as Dell. IDC can see a scenario whereby Dell could move much higher up in the enterprise workload markets using multicore processing and, consequently, be very disruptive to major vendors in the mid- to high-end server market tier (such as IBM, Sun, and HP). This scenario is based on Dell's ability to assemble either a strong scale-out or scale-up value proposition and offer confederation systems management technologies for a relatively low product margin. The response from the competition will be based on a complete solution offering and integrated systems platforms through ease of management and linkages to storage and network offerings. However, the door will be opened and the barriers to entry to offer SMP-like systems will be lower than ever before. Consequently, we believe a major market shift could be on the horizon.

Third Generation of Multicore Processing: 2007–2009

The third generation of multicore processing will have major implications for the IT infrastructure. By then there will be dual-core, quad-core, and octi-core (8-way) systems from traditional platform suppliers as well as the specialists offering even more very dense multicore platforms. However, by this time in the market development and adoption cycle we are assuming that I/O interconnects and memory latency issues will be overcome and supported to create well-balanced systems. As you can imagine, though, this market will create a wide range of challenges and opportunities. IDC believes that the effects of this technology will be driven by innovation and disruption. For example:

- "Systems on a chip" will be widely available. As a result, new markets will open up and serve customer needs that are not being met today. These offerings will begin to manifest themselves in systems for the broader small and medium-sized business market. This scenario calls for an increased level of interest by channel partners, integrators, and service providers.

- Functions on a core will begin to be available. For example, some of the cores may be dedicated to security, and we could see a shift in the software functions offered today as these functions become "embedded."

- Systems management, provisioning, and virtualization engines could be supported on cores, allowing new products to be created. However, we don't believe that different cores will support dedicated and different operating systems (i.e., cores optimized for specific operating systems and dynamically configured to meet new requirements). We believe that multioperating system support will be based on traditional partitioning techniques. Overall, major R&D effort on systems management capabilities will continue because this is where much of the market differentiation will appear and, consequently, the revenue opportunities will exist.

- Dense core offerings will create new form factors that currently are not possible. As a result, the industry will be able to serve new customers, markets, and applications. This change will drive a new innovation cycle in the systems industry and potentially allow a clear link between consumer-based multicore systems and enterprise-class systems.

- Specific-function multicore processing environments will be well established. For example, platforms to optimize Java, J2EE, or highly threaded applications will become part of the infrastructure as long as they support the existing programming models and can be implemented in a nondisruptive manner.

However, the biggest shift in the marketplace may be the following scenario. By the time the third iteration of multicore processing has arrived, today's approach to building large-scale SMP systems may have reached a crossroads. Although the scenario does not call for an end to SMP platforms as we know them, the multicore processing capabilities could be strong enough to run a majority of the workloads that currently require these platforms. These capabilities, combining multithreading, multicore, scalability, security, systems management, and so on, at prices that make total cost of operations and ROIs very attractive while maintaining the application programming models, could be highly disruptive to the current SMP market model. Additionally, the following market conditions could accelerate this scenario:

By the time the third iteration of multicore processing has arrived, today's approach to building large-scale SMP systems may have reached a crossroads.

- Investment by ISVs and systems management software companies will have been focused on the emergence of the new technologies and platforms, and the current SMP market model may no longer be attractive for sustainable growth opportunities. Consequently, vendors will face issues regarding investment protection and competitive offerings that make it harder to justify the current investment model.

- New application programming models and their requirements would confirm that highly threaded and multicore platforms could be the optimal platforms to provide network-based computing or processing. Using a combination of multicore product offerings could significantly improve the IT business value equation.

- Large configurations of multicore processors managed by grid-computing-like infrastructures (this assumes we continue to see the transformation of grid computing functions from technical computing to enterprise computing) may offer up another IT infrastructure model that will aggressively compete with today's SMP infrastructure model.

- Bladed architectures based on multicores will become the normal IT infrastructure offering, combining the benefits of virtualization, power, management, and density into a multioperating system environment supporting large storage and network needs through integrated switch technologies.

Consequently, the disruption of the IT infrastructure market could be widespread and beneficial to many suppliers. However, it is clear the suppliers that have the most intellectual property (IP) on the dynamic IT stacks will be the natural leaders and able to control entry to and exit from a market that was thought to have reduced the barriers to entry. In fact, we believe that the race to own IP will create significant barriers to entry but shore up profit pools that currently are open to many entrants.

ESSENTIAL GUIDANCE

Every IT supplier should be aware of the potential impact that multicore processing could bring to its markets. As each iteration of the technology is released, its capability extends further and further into the enterprise, eventually impacting the high-end systems market. As the impact is felt with each tier of the datacenter, opportunity and challenges grow, and it is important to understand now what the available market size and opportunities may become.

As each iteration of the technology is released into the market, its capability extends further and further into the enterprise, eventually impacting the high-end systems market.

IT customers will rapidly gravitate toward multicore processing because it has the ability to be implemented seamlessly and give the majority of workloads a performance boost. Although there will be some workloads that won't see immediate improvements, the technology will soon address those needs and, consequently, the technology will be pervasive.

However, IDC believes that ISVs and systems management suppliers should avoid attempting to confuse end users through complicated pricing models. The consequence of not heeding this guidance will be a backlash from CIOs and major openings by competitors in the same market space.

LEARN MORE

Related Research

- *Multicore Processing: Disruption or Distraction for the IT Infrastructure?* (IDC #32278, November 2004)

Published Under Services: Enterprise Server Fundamentals; Modular Server Solutions