

WHITE PAPER

How Multicore Processing and 64-Bit Operating Systems Enhance a Utility Computing Environment

Sponsored by: Advanced Micro Devices, Inc. and Egenera

Kelly Quinn

Jed Scaramella

John Humphreys

October 2005

IDC OPINION

In today's market efficiency, effectiveness and the integration of the IT service with the business process are all top of mind with technology leaders. As organizations wrestle with how to better align IT with their business and drive economies of scale into their IT environments, increasingly utility computing is an option to consider. As the concept of a pool of resources develops, technologies like multicore and 64-bit computing in the x86 space are helping to mature the utility environment, expand the pool of potential applications and dissolve even more barriers in today's compartmentalized IT environments. While there are clearly challenges to creating an IT utility, some of the early customer concerns around the capabilities of the technology are being addressed, and users are encouraged to consider how and when their organization would benefit from adoption.

IN THIS WHITE PAPER

In the IDC white paper, *The Next Evolution in Enterprise Computing: The Convergence of Multicore x86 Processing and 64-Bit Operating Systems*, sponsored by Advanced Micro Devices, Inc., the impact of these two converging technologies on enterprise computing was discussed at length (see http://enterprise.amd.com/downloadables/IDC_Multi-Core_64-bit_White_Paper.pdf). The paper *The Impact of Multicore Processing and 64-bit Operating Systems Deployed in a Utility Computing Environment* examines the impact these combined technologies bring to utility computing environments.

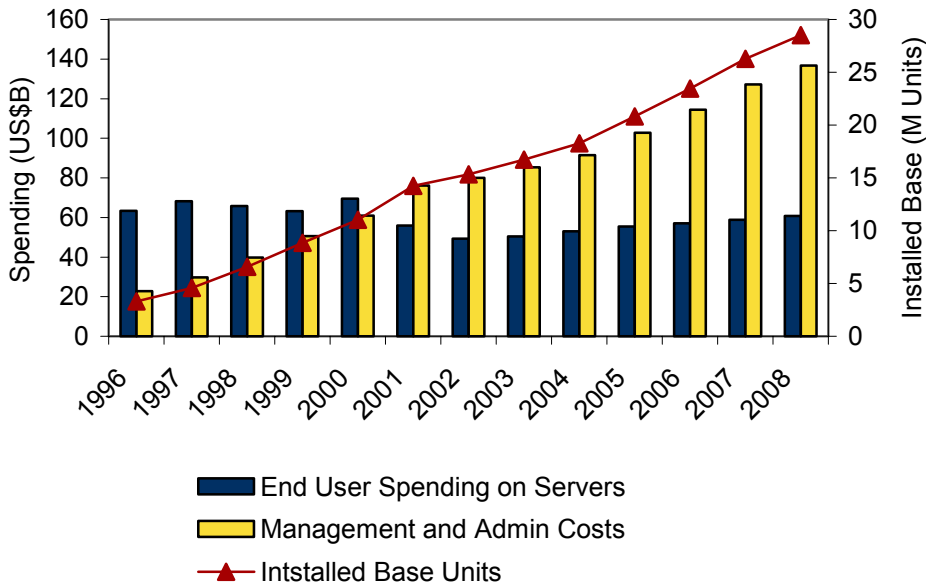
SITUATION OVERVIEW

Today's IT managers face opposing priorities from multiple stakeholders within the enterprise. Business unit managers expect IT to respond quickly to business priorities, while financial managers are continuously seeking ways to lower costs and improve efficiency in the datacenter. As a means to reconcile these concerns, there is a move underway to manage IT more comprehensively. As part of this transition, CIOs and IT managers are increasingly looking at how IT can better support their businesses. This includes investigations into:

- ☒ **Increasing availability of applications.** Businesses are increasingly operating around the clock, with little to no downtime. IT resources and services are required to be available to end users on a 24 x 7 basis.
- ☒ **Improving performance of applications.** Business applications are increasingly linked to partners, suppliers, and customers. The performance of these applications is critical to deriving business benefits from the IT infrastructure.
- ☒ **Designing scalability in infrastructure.** As new applications and users are added, the ability to easily scale the IT infrastructure at minimal cost becomes significant.
- ☒ **Aligning IT resources with business requirements.** IT infrastructure is required to be increasingly flexible to better address business priorities. To realize full value, managers must be able to utilize IT resources in multiple ways.
- ☒ **Simplifying infrastructure management.** Demand is increasing for tools that automate day-to-day tasks in order to assist system administrators in streamlining processes and minimizing errors.
- ☒ **Reducing total cost of ownership.** Operational and administrative costs represent a significant and increasing portion of the total cost of IT expenditures (see Figure 1). To reduce expenses, IT managers must find solutions that control management costs.

FIGURE 1

The Growing Cost of Managing a Distributed Environment



Source: IDC, 2005

Beginning the Process

Traditionally, companies overbuilt infrastructure to ensure that capacity was available to meet peak demand. Consequently, IT environments have been managed inefficiently and yielded low utilization rates. As a first step to addressing this and the aforementioned challenges, IT organizations have been consolidating servers into centralized locations so that computing resources can then be dedicated to specific applications and users. By reducing the number of server systems, companies are able to minimize traditional compute silos as well as lower systems administration costs. While this is a good first step, computing resources remain dedicated to specific applications and users — even with consolidation. Only by having an environment where applications and users are not dedicated to specific compute resources can these silos be eliminated.

To combat this situation, IT managers are turning to utility computing technologies, which pool hardware resources and share compute cycles. With utility computing, many companies have been able to realize cost reductions in servers, achieve better use of currently deployed assets and improve the service levels of the business services they host.

However, even as IT managers are seeking solutions to decrease costs, increase utilization rates, and increase agility, they are faced with escalating demand from end users. In addition to expecting faster and more powerful computing to address large data sets, end users are expecting greater reliability and continuity of operations from their computing infrastructure.

Overview of Utility Computing

Utility computing is the pooling of infrastructure to dynamically provision applications. Utility computing allows IT managers to meet changes in business priorities and fluctuations in demand. For example, a commercial bank provisions a large portion of their IT resources to processing customer transactions during business hours. Leveraging utility computing, these same assets will be reallocated to account for reconciliation calculation and statement processing after business hours. Today, companies are enabling utility computing (a few applications at a time) through the use of software and hardware that automates the provisioning, virtualization, and management capabilities of IT. Largely, enterprises are installing internal systems, software, and processes to configure the IT utility. In this way, computing resources can deliver higher value as they are linked directly to business objectives.

Currently, utility computing is being deployed across a wide spectrum of applications and enterprises. Several early adopters have experienced the greatest progress by establishing utility computing architectures for mission-critical applications, such as business processing and decision support workloads. However, there is increasing activity from mainstream companies utilizing "sand box deployments" in testing environments.

Broadly, users can expect to derive two major benefits from utility computing: the ability to address cost, resource utilization, and management efficiency and the ability to enhance service levels and responsiveness. First, an increased number of assets can be managed by a given set of IT personnel, reducing operational and administrative expenses to help lower operational costs. Second, through provisioning, virtualization and service-level automation, computing resources can yield improved utilization rates and impact capital costs. The increased flexibility and ease of management can produce significant time-to-response advantages as well as better uptime on key business services.

At the highest level, utility computing represents an infrastructure that is a virtualized pool of resources that can be automatically provisioned based on predefined business priorities and provides accompanying services (support, metering, and billing). As a result, an IT environment is created that can ensure service levels and control costs, leading to high availability of computing resources and an increased level of business agility.

Utility Computing in Practice

One company that is helping to lead the cultivation of this market is Egenera. Founded in 2000, Egenera introduced its first product in 2001. The BladeFrame system was devised as a platform for utility computing. The crown jewel of the platform is the PAN Manager — the management, provisioning and scale-out virtualization software that turns the BladeFrame system into a pool of compute power. PAN Manager's automatic management capabilities facilitate efforts to decouple applications from servers and network components.

Egenera is able to leverage the combination of 64-bit computing and multicore processors to power the utility computing environment. BladeFrame with PAN

Manager has the ability to scale applications quickly and precisely depending on demand. In addition, this highly scalable solution provides the flexibility and sophistication to support mission-critical applications.

Overview of Benefits of Multicore and 64-bit Systems

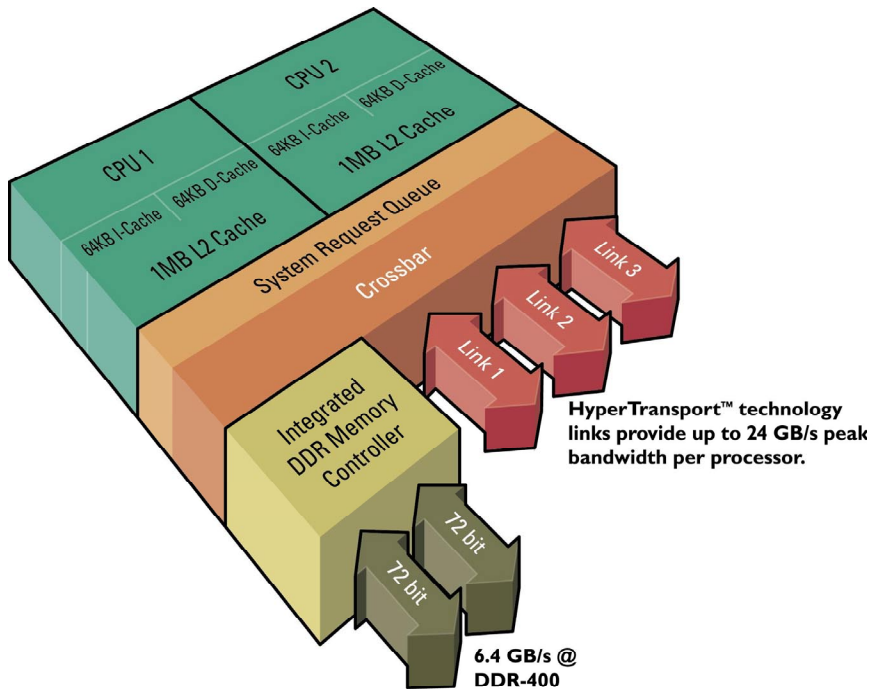
The Dual-Core AMD Opteron™ processor with Direct Connect Architecture integrates two separate logical processing units on a single chip that share memory, I/O and caching (see Figure 2). The result is a processor that consumes the equivalent power of a single core, yet provides significantly more processing power than a single-core processor. Since the power consumption per core is reduced, performance per watt is increased. Users can dedicate individual cores to different applications, providing a natural platform for consolidation — the first step to utility computing. In addition, the AMD processors produce less heat, thereby allowing vendors to manufacture servers that are more dense and require less physical space in the datacenter.

Along with multicore processors, the innovation of 64-bit computing provides a leap forward for x86 server technology. 64-bit architecture provides considerably higher performance than 32-bit architecture. As a result, processing power that was previously only available on expensive mainframe servers is now possible for x86 servers.

In addition, 64-bit architecture significantly increases addressable physical and virtual memory, which in turn greatly improves overall available memory. The AMD Opteron™ processor enables simultaneous high-performance 32- and 64-bit computing, providing investment protection for customers. Not only is the technology backwards compatible with older applications and 32-bit operating systems, it is also designed to allow future expansion of I/O, memory, storage, and processing — which can alleviate customer concerns about their ability to seamlessly increase computing capabilities down the road.

FIGURE 2

Dual-Core AMD Opteron™ Processor Architecture



Source: AMD, 2005

"Carry the One" with AMD Direct Connect Architecture

AMD has made Direct Connect Architecture a mainstay in its microprocessors, and its inclusion in AMD's 32- and 64-bit, multicore design will further enhance its benefits. Direct Connect Architecture features an integrated memory controller that directly links together the microprocessor, its memory, and the CPU's I/O. This design reduces latency between the memory controller and the microprocessor and can improve overall system performance. By avoiding bottlenecks in the system, Direct Connect Architecture will be able to fully realize both the 64-bit memory advantage and the improved computational power inherent in a multicore system.

The Multiplicative Effect

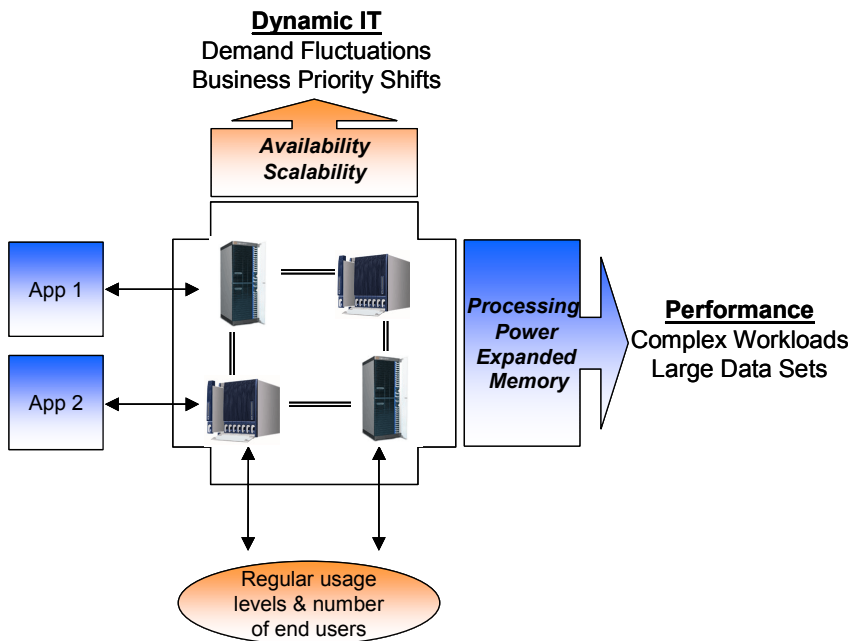
As stated earlier, the goal of utility computing is to directly link IT infrastructure with business priorities to deliver the benefits of increased operational efficiency and technical agility. To that end, the technical advancements of 64-bit and multicore processors increase the impact of x86 servers in utility environments and provide even greater flexibility to the enterprise employing utility computing (see Figure 3).

Specifically, the business and operational benefits multicore and 64-bit computing amplify in utility computing environments include:

- ☒ **Improved performance.** Increased computing and memory capacity enable process-intensive workloads and database applications to migrate to utility computing environments.
- ☒ **Flexibility and availability.** Utility computing provides agility in IT infrastructure to meet peaks in demand across the broadest swath of 32- and 64-bit applications, enabling IT managers the ability to re-provision resources according to shifting business needs.
- ☒ **Lower total cost of ownership.** Utility computing facilitates consolidation efforts and simplifies infrastructure, which in turn can improve operational efficiencies and reduce management expenses.
- ☒ **Increased control over IT.** Through access to granular levels of compute resources, IT managers are able to align IT supply to business need, as well as quickly test, deploy, and "right size" production capacity during a new application roll out.

FIGURE 3

Multicore and 64-Bit in a Utility Computing Environment



Source: IDC, 2005

The Customer Perspective

The expanded memory footprint of 64-bit computing allows large data sets to be moved entirely into memory for faster computer processing, which enables utility

computing to tackle a broader set of workloads. This offers users the flexibility to run smaller infrastructure applications as well as complex numerical calculations involved with business-critical applications, such as business intelligence and database workloads, on a single infrastructure.

When leveraged in a utility computing environment, the addition of dual-core technology also affords users even more flexibility. Customers can allocate CPU cores to specific applications with the confidence that their needs for reliability, availability, capacity, and performance can be met.

IDC believes that incorporating dual-core and 64-bit capabilities into planning for a utility environment will help users further expand the benefits they receive. Users can expect to see benefits from the increased levels of uniformity of infrastructure, the commoditization of the components that make up the infrastructure, and the ability to decouple the application from the underlying infrastructure. The benefits will include both operational and capital cost reductions as well as improved service levels and greater responsiveness from the business.

The combination of dual-core and utility computing provides step-function enhancements in the power of computing and the ability to control and direct that power. Once applications are unencumbered by the tethers of physical servers, they can be deployed in testing environments on different CPU configurations quickly and efficiently. Additionally, multicore defines utility computing at the granular level, enabling the precise matching of computing supply with the exact needs of end users. Companies can identify the precise computing support required to roll out the application as well as the level of support required for future application scaling while maintaining the benefits of reliability, availability, and fail-over protection.

CHALLENGES/OPPORTUNITIES

Utility computing adoption within an enterprise may be inhibited by internal issues as some business units may want to own and host their own computing resources. A "shared" infrastructure may not be a concept that all managers support. Resistance may also be felt from financial managers unwilling to invest in new tools with an ROI that may be difficult to quantify. Additionally, with some systems, security becomes a significant concern. A high level of security will be required in resource sharing, especially in applications that touch external parties. Platforms should be evaluated with this in mind.

Software licensing is another obstacle to overcome. Microsoft, a large player in the x86 market with Windows, announced in 2004 it would charge on a "per-processor," rather than "per-core," basis. To ensure successful mainstream adoption, other software vendors need to follow Microsoft's lead.

The incorporation of business processes from outside the enterprise will be a key driver in the adoption of utility computing. For example, the flexibility, availability, and scalability of a utility computing environment are ideal to support supply-chain integration. It is vital that the services have consistent interoperability and dependable security.

Non-x86 multicore processors have been available in the market for some time and provide a proven track record in terms of reliability and performance. Some customers may consider this fact in decisions concerning mission-critical applications. However, IDC believes the benefits from the industry-standard architecture and utility management capabilities can deliver the price/performance improvements necessary to drive adoption.

CONCLUSION

IT managers will continue to face pressure to deliver increasing levels of services and at the same time lower expenses. Computing resources must increase the availability and performance of applications, while maintaining enough flexibility to align with changing business needs. As enterprises initiated consolidation efforts, utility computing enabled IT assets to be deployed more efficiently by improving the levels of service to business units.

The enhancements of 64-bit and multicore processors bring utility computing to the next level. x86 servers will play an increased role in the datacenter. A wider range of workloads from network to process intensive can be handled by the expanded memory footprint provided by 64-bit computing. Database applications, once the domain of mainframes, will be able to migrate to utility computing environments.

Augmented by 64-bit and multicore processing, utility computing offers IT managers the flexibility in infrastructure to handle demand peaks. In addition, with the increased ability to control infrastructure, IT managers will now be able to align resources to business needs while lowering total cost of ownership.

Customers considering utility concepts in their own IT infrastructures are encouraged to monitor the technical advancements that are not only broadening the applicability of utility computing, but also maturing the software stack that is enabling its deployment in support of ever-larger and more mission-critical workloads.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2005 IDC. Reproduction without written permission is completely forbidden.