It's Not the Cores – It's the Systems

Recent research conducted by Gabriel Consulting has identified confusion within the IT management discipline regarding how to absorb and incorporate a fundamental change in the way Moore's law is playing out in the x86 microprocessor market. As Intel and AMD have succeeded in shrinking the circuits on their silicon wafers using sub-100 nanometer (a nanometer is a billionth of a meter) manufacturing process technologies, performance gains delivered by the perpetual rise in processor clock frequencies is being replaced by a multiplication of the number of CPUs or 'cores' plugged into a single processor socket. The net effect here—of chip-level transistors doubling every two years—with multi-core designs, now breaks a traditional and pervasive x86-based application programming style using low levels of concurrency or parallelism.

Adding to the confusion, Intel has fueled the fire with recent announcements concerning their 80-core 'concept car' processor, and its introduction of a quad-core processor. AMD will have a quad-core in the second half of '07, with IBM, Sun, and HP (with Itanium) pushing out their quad-core (or greater) offerings as quickly as possible. All of this multi-core hoopla has led to an industry mindset that "more cores are more better". Bad grammar aside, we're a little concerned that some customers may see these nifty quad-core processors as the answer to every problem in enterprise computing – "4x the power at the same price – it's the miracle we've been praying for!" Just like every new technology in our industry, the initial hype is becoming a little bit overblown.

Quad-core processors are certainly a worthy technology, and a significant advance, but, in their current instantiation, they are certainly not the be-all, end-all of system technology. Many are touting these systems as no-brainer replacements for larger systems, the typical comparison being a dual-socket quad-core (8 total execution units) server vs. four-socket dual-core (also 8 total execution units) systems. However, we aren't so sure this is the right move, mainly because the current crop of quad-core boxes, from a *system perspective*, don't quite measure up to their four-socket brethren. This may sound mildly heretical, but we have our reasons, mostly based in what we see as the challenges data centers are facing.

Bottlenecks & Benchmarks

The quad-core is an advance of processor technology – obvious, right? But, the reality is that the majority of x86 systems in data centers are not CPU bound. Current processors have more than enough power to handle virtually any workload thrown at them from simple file serving to heavy transaction processing. Most users report that they reach the limits of their systems' memory and I/O capacity long before they overrun the processor. This is our major quarrel with the rush towards quad-core servers today – the quad-core systems in general just aren't as robust as quad-socket systems in terms of system size and configurability, we'll talk about this in detail later on in this paper.

Gabriel
CONSULTING GROUP

The early quad-core systems look pretty good from a benchmark perspective. These servers even look pretty good on the new consolidation benchmarks. However, like all benchmarks, these workloads are fairly simple when compared to the complex workloads in the real world. Consider that many of today's applications were written years ago for 16- or 32-bit processor designs, and adapted to take advantage of new instructions and memory addressing schemes incorporated into the new '64-bit' x86 designs. A significant percentage of the enterprise applications (ERP, CRM, BI, etc.) were also ported from older UNIX or other legacy platforms and may still contain artifacts of those code bases.

Don't get us wrong, benchmarks are a good starting place to evaluate how well these new systems perform, but we aren't sure that these limited tests are comprehensive enough, or that there are enough data points to help customers sort out current purchasing decisions. In short, benchmarks are no substitute for several months of grinding through enterprise applications in a pre-sales, proof-of-concept engagement.

Changing Workloads & Usage Models

As x86-based servers continue to move up the enterprise, they are increasingly being utilized for mission critical enterprise workloads. Our recent survey of x86 data centers clearly illustrates this trend…

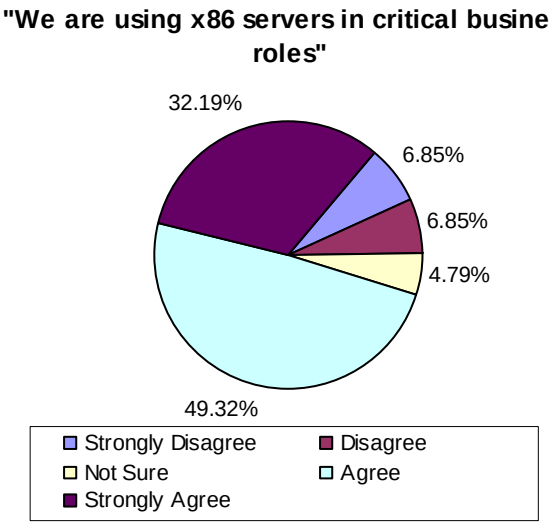

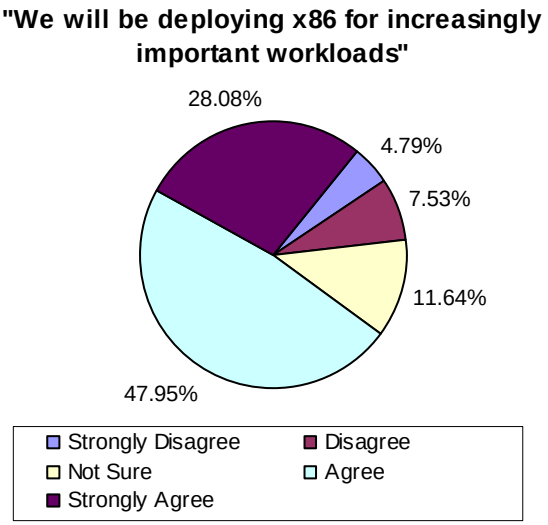

Over 80% of our survey respondents said they were using x86 servers to handle critical business workloads (above left), moreover, 76% of these same respondents said they were moving increasingly important workloads to x86 platforms. In the same survey, 73% of participants agreed that "properly managed, x86 servers are just as available/reliable as UNIX servers", with another 72% saying that x86 servers, with correct configuration, can meet or beat UNIX system performance. All of this data points to x86 platforms becoming trusted and important data center players – much more than just file and print servers.

Workloads on x86 systems have not only become more important, they are also multiplying at a frightening pace. The still prevalent single-application-to-single-server usage model has spawned huge numbers of systems in most organizations. This 'server sprawl' results in extra labor for

system management, spiraling power/cooling demands, and rapidly increasing floor space requirements. To add insult to injury, the average utilization rate of x86 systems hovers somewhere around 6% - which is an astonishingly low number. It isn't that these systems were intentionally over-configured, rather they were sized to handle the highest possible peak load – plus about 25% extra capacity 'to grow on' for safety's sake. Depressing the average utilization rates even further are large numbers of test and development systems that seem to hide in the nooks and crannies of every organization. The problem is that all systems, production, development, and testing systems alike, regardless of utilization, have to be managed, maintained, secured, powered, cooled, and given a home somewhere on the data center floor. The expanding numbers of boxes and increasing complexity of the interrelationships are pushing data centers to the breaking point in terms of both labor and facilities overhead.

Guidelines for Implementing Virtualization Technology

Solutions to x86 server sprawl are finally emerging and being implemented. Virtualization technology makes it possible for multiple workloads to be safely and efficiently hosted on single x86 systems. Currently the most popular x86 virtualization product is EMC's VMware software, which allows users to run multiple instances of Windows and/or Linux operating systems in partitions on a single system. Properly implemented, the benefits from x86 virtualization can be profound. We have spoken to customers who routinely pull off 16-to-1 server consolidation ratios, decreasing system footprints, overall power/cooling requirements, while increasing utilization rates five-fold. The impact on overall costs is also significant, with customers citing savings in the 20-30% neighborhood.

We believe x86 virtualization is inevitable and will gradually become the dominant new application deployment usage model based on data from our primary research efforts.

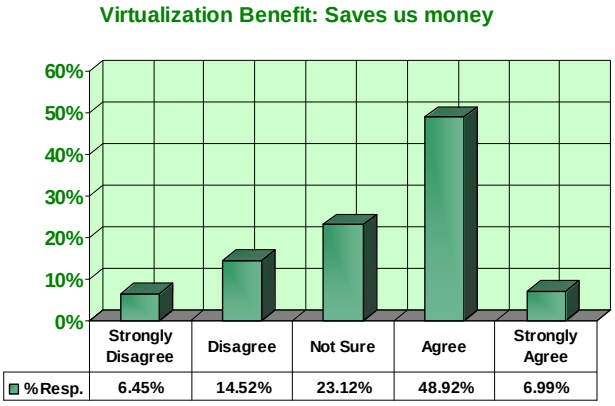

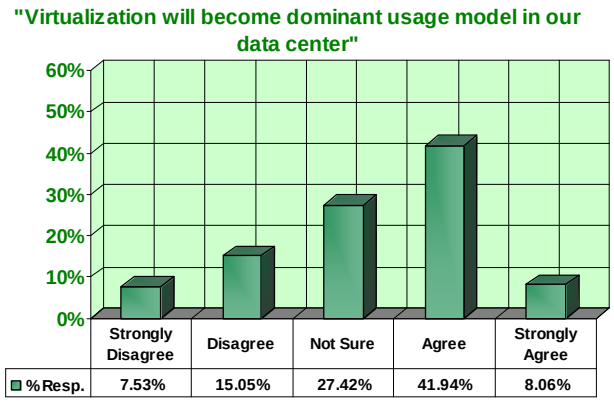

In our most recent poll of enterprise x86 data center personnel, we found that a large majority have already begun to implement x86 server virtualization/consolidation. Almost 60% of our survey population say that this move to virtualization is saving them money overall. More than half of the respondents believe that virtualized x86 will become the dominant usage model in their organization – and we expect this number to rise significantly over time.

Given that consolidation is going to be a fact of life in most data centers, the fundamental question facing IT managers is the definition of an optimal consolidation platform. We offer two pieces of broad pieces of advice:

1) Buy Scalability Larger systems (greater than 2-socket) allow for much greater consolidation ratios and provide headroom to handle potential simultaneous peak loads on multiple applications. One of the hallmarks of today's Internet-centric workloads is that they're inherently unpredictable, exhibiting both high peaks and low valleys of server usage. Most organizations lack the time and resources to conduct exhaustive pre-consolidation studies to characterize workload behavior. In many cases, x86 workloads change so much over time that historical usage stats offer little or no help in predicting future needs. Because of this, we find many organizations implement x86 consolidation by way of installing scalable x86 gear and then layering on multiple workloads. Virtualization software, coupled with solid system/workload management software, has the ability to closely monitor workloads and automatically change resource allocations to ensure that important applications always have enough resource to hit quality of service requirements. However, customers are finding that characterizing performance in virtualized systems prior to actual virtualization, is difficult – memory and I/O requirements are generally much higher than anticipated. This underscores the need for scalable hardware that not only offers the ability to run multiple processors, but can be configured with large amounts of memory and I/O capacity.

2) Buy Availability This is intuitive for most customers. If you're putting many eggs in a single basket, the basket needs to have a very solid handle, reinforced sides, and a double bottom. Virtualization provides many benefits, but it also raises the stakes in terms of reliability and availability. On a virtualized system, a hardware outage can take down scores of workloads that may be supporting business critical operations. It is lunacy to put these applications on systems that do not offer state-of-the-art RAS (Reliability, Availability, and Serviceability) features.

This brings us back to why we think the current quad-core systems might not be the best choice for enterprise x86 workloads. The systems can not be scaled to the same level as current four-socket systems, and they do not offer the same level of RAS features. In short, our quarrel isn't with quad-core processors; it's all about the systems in which they are being placed.

To illustrate our point, let's consider a range of IBM systems. IBM offers both scalable four-socket, dual-core processor based systems along with dual-socket, quad-core based servers.

	IBM x3650	IBM x3850	IBM x3950
Processor Sockets	2	4	4-32 total (with expansion chassis)
Processor options	Intel **quad-core** Xeon, 2.66GHz	Intel dual-core Xeon 3.5GHz	Intel dual-core Xeon 3.5GHz
Memory slots (max) **Memory capacity (max)**	12/48GB	16/64GB	16 slots per chassis 64GB/512GB
Memory type	Fully buffered DIMM 667 MHz	PC2-3200 DDR II	PC2-3200 DDR II
Network interface	Dual Gigabit Ethernet	Dual Gigabit Ethernet	Dual Gigabit Ethernet (per chassis)
Expansion slots/type	4-PCI-Express **or** 2 PCI-X and 2-PCI-Express	6 total, 4 PCI-X, 2 PCI-Express	6 PCI-X
Hot Swap Components	Power supply, fans, hard drives	Power supplies, fans, memory, hard drives, PCI-Express/PCI-X adapters	Power supplies, fans, memory, hard drives, PCI-X adapters

The servers in the table above represent three different ways to build an 8-core, Intel-based system. There are a fair number of similarities between the products; all have dual power supplies, built-in RAID support, and dual gigabit Ethernet network adapters. But it's the differences that we're concerned about.

The first area of concern is total memory capacity of the quad-core systems in general vs. the more scalable enterprise servers. For example, the IBM x3650 offers a total of twelve memory slots, while the x3850 and x3950 servers have 16 memory slots – 30% more. While we agree that 48GB of total memory is quite a bit (even though it takes ultra-expensive 4GB memory sticks to get to this level), we'd be more comfortable with the system if it offered more. The x3850 and x3950, with 16 memory slots, is better – but we'd still want greater capacity. Why? It's because virtualization is extremely memory intensive. The first thing we hear from clients beginning to implement x86 virtualization is an exclamation (sometimes obscene) about how much memory virtualized systems need. Memory is the first (and biggest) bottleneck in a virtualized system. Each consolidated workload will require at least 1GB of memory. Scalable applications and other memory intensive workloads may require double or triple that figure. When you pile more apps on a system, the total memory requirement is significantly more than just the sum of memory the apps used in separate boxes. Thus total memory capacity is a real concern and this is one case where size really matters, and those four extra memory slots can provide the critical performance margin needed to handle simultaneous peaks in important workloads.

Another important consideration is I/O capacity. As consolidation becomes more widespread, we expect to see customers who do not use scalable platforms to have problems with I/O bottlenecks. For some, it might be hard to imagine, but we're already seeing some evidence of this in the field. The enterprise-class systems, like the x3850 and x3950 in the table above, have 50% more I/O expansion slots than the smaller quad-core x3650. Again, as a system hosting a single workload, four expansion slots might well be overkill. But consider that systems in our brave new virtualized world will probably be hosting 20, 40, 60 or more applications – and that many apps will definitely use a whole bunch of I/O. Add in the fact that many of these workloads will have unpredictable spikes, and it just makes sense to have a system that can handle a big I/O load.

The enterprise servers also add a lot of business value in terms of RAS features that generally aren't found in the current quad-core servers. As we discussed above, virtualized servers demand a level of RAS that is much higher than the garden variety x86 server. Enterprise servers, such as the IBM x3850 and x3950, offer Predictive Failure Analysis (PFA) on almost all of the components on the system. PFA constantly probes components for hints of impending failure, giving administrators notice when a failure is imminent. Armed with this knowledge, administrators can then hot-swap a wide range of components (power supplies, memory, drives, fans, and even I/O adapters) without having to power the system down. This radically increases uptime, reducing both planned and unplanned outages. While some of the quad-core systems offer a measure of PFA and hot-swap, they have generally been engineered to hit a price point rather than provide the highest level of RAS.

One thing we haven't yet discussed is the cost difference between the new quad-core servers and what we've been referring to as enterprise systems. As we just stated above, the quad-core boxes come in at a pretty attractive price point – fully configured, from $12k-$15k. The enterprise systems, at the same configuration level will probably cost $5k-$10k more. The price delta will be higher or lower, depending on specific configurations, enterprise purchasing agreements, and negotiating skills. We're not sure that price differences are all that relevant a topic in this discussion. For a reasonably sized enterprise, the cost of server hardware has become a much smaller proportion of most projects. In discussions with several clients, they confirmed that the

server share of costs in a typical deal usually ranges anywhere from 10-25% of the total. Software and services costs are usually **much** higher than hardware costs. With this in mind, the savings provided by a quad-core system just doesn't seem to make much difference in the big picture.

Summary/Recommendations

While quad-core systems might well become the norm in the future, the technology can't simply be dropped into the same server design that worked for their single-core predecessors. Gabriel Consulting believes tomorrow's x86-based server offerings will look very different than they do today. The traditional and obsessive focus on low cost in the single and two-processor space will abate resulting in larger and more reliable rack-based products. The need for larger than four-processor SMP servers will likely be rediscovered by existing vendors and may offer opportunities for new players servicing the needs of large enterprise organizations.

The cost to customers for larger x86 servers will continue to decrease from a price/performance perspective as new virtualization technology allows customers to achieve much higher system utilization rates. We expect to see fierce competition between virtualization vendors that will result in better products in terms of feature/functions and business value. In turn, this will force application vendors to both embrace the technology and support their code running in virtualized partitions and adapt their licensing models – where necessary – to equitably charge for the value derived by end users.

The lack of experience and perspective monitoring the performance of x86 software workloads and the immature nature of today's hardware technology will depress the potential performance gains to be had with the use of current quad-core processor technology in more traditional x86 network server roles. We'd feel better about recommending them as enterprise systems if they offered more robust configurations and thus the ability to handle more work while providing headroom to guard against unanticipated workload spikes. We would also like to see a fuller range of RAS features, given that these systems may well be hosting some of the most important applications in the enterprise. In this soon-to-be-consolidated world, selecting the right server is becoming more, not less, important and mistakes can be costly. With this in mind, we recommend that customers purchase the highest level of scale and RAS available.

Gabriel
CONSULTING GROUP
phone / 503.372.9389
gcginfo@gabrielconsultinggroup.com
www.gabrielconsultinggroup.com