

Trend Watch: Mission Critical x86

Over the past several years, our research has shown a definite trend of customers moving mission-critical workloads to x86 systems. We believe that the emergence of highly sophisticated x86 virtualization mechanisms, coupled with advances in x86 system and solution technology, is making customers see the platform in a new light. While some customers are consolidating important workloads onto fewer, but larger, x86 systems, others are choosing to deploy new mission-critical applications directly onto x86-based hosts.

In this report, we examine what's driving this trend and take a look at two of the market leaders, IBM and VMware, and discuss what they are bringing to the table...

It's no secret that customers in enterprises large and small utilize x86-based systems to run a plethora of mission-critical workloads. It wasn't always this way; as little as eight or nine years ago, customers shied away from using x86 for their most important applications.


There were solid technical reasons for this uncertainty. Linux was still in its infancy and, while it was developing fast, it was still new and unknown to a lot of enterprises. Microsoft was doing better in data centers with their Windows Server product line, but that too was still in the throes of security and scalability growing pains. Over time, both operating systems have improved markedly in terms of availability, scalability, and security. At the same time, the underlying x86 hardware has radically improved in terms of scalability, availability, and overall sophistication. And in the last few years, x86 virtualization has given customers the ability to run these servers at much higher utilization rates, while improving flexibility and application availability at the same time. The x86 platform has grown up and matured; no longer are they considered suitable only for edge servers handling low-level chores. They've reached the point where they are making inroads into data centers in a new role: as the host of choice for mission-critical applications.

We've been tracking this trend through our annual x86 server vendor preference surveys, where we ask the experts – IT managers, architects, and administrators – what's happening in the data center. Over 90% of our survey participants work with systems from multiple vendors, so they've 'been there, done that' and can give us the inside track on how they're managing their workloads. They can also tell us how their IT shops are evolving as they respond to the constantly changing demands of enterprise computing.



What we've found is an accelerating trend of customers telling us that they're putting more apps – and more *important* apps – onto their new and existing x86 servers.

"How many of your x86 workloads are business critical?"




In our most recent survey, we asked data center personnel to tell us how many of the applications hosted on their x86 servers are 'mission-critical' to their organization. On the chart at left, you can see that 99% of our survey respondents said that at least some x86 workloads are vital to their organization. But more telling is the fact that a third of the customers we surveyed said that more than 50% of the apps hosted on their x86 platforms are mission-critical. When we began asking this question three years ago, the number of customers who said that half or more of their x86 workloads are critical to their organization was 20%; we've watched this number grow

steadily in each survey. Most of these customers are SMBs with fewer than 1,000 employees; their size alone makes it more likely that they have an all-x86 infrastructure. However, we are seeing the trend toward mission-critical x86 strengthening in large and mid-sized enterprises as well.

We also asked these same customers about their future plans as they relate to x86 servers and the applications they plan to host on them. The trend towards mission-critical workloads moving onto x86

"Do you expect to use x86 to host more or fewer business critical apps in the future?"



really couldn't be clearer; almost 40% of our respondents say that more important apps will be deployed on x86 in the future.

It's obvious that the x86 platform has gained the confidence of customers over the past several years and is now fulfilling a vital role in the vast majority of IT infrastructures. There are several factors fueling this trend. The first, of course, is that the systems and software have gotten much better. Both Linux and Microsoft Server operating systems are much more scalable, secure, and reliable. There have also been great strides forward on the system front. While there are considerable differences between the offerings

from the various x86 vendors, the best systems offer true enterprise-class size, features, and availability that put them in the same league as other data center servers. Even more important is that some of these vendors aren't simply churning out lowest common denominator vanilla systems; they're adding features and capabilities that make them a much better fit for business-critical computing.

Another trend that is driving the move towards using x86 systems in mission-critical roles is the rise of x86 virtualization. In years past, customers mostly experimented with x86 virtualization. They began these efforts by combining multiple test and development workloads onto single systems; then they expanded to using virtualized systems to host multiple edge-of-data-center applications like file servers and web server instances.

But virtualization technology in x86 has advanced at a very fast clip – much faster than its evolution in mainframe or Unix systems. And it’s no longer just a mechanism for increasing system utilization; it also enhances server and application availability, workload management, and overall system manageability. In some ways, x86 virtualization as offered by VMware has surpassed the major Unix flavors in terms of features and capabilities. In our opinion, this is a major reason why customers are hosting more of their most important workloads on x86 systems.

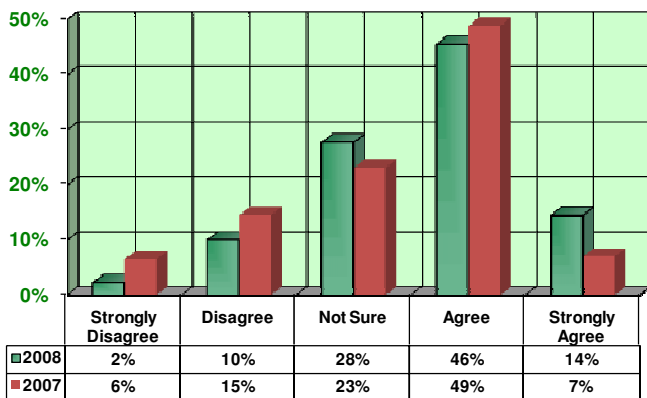
The payoff for data centers is significant. In our surveys, we’ve been asking customers whether they have actually seen the much-touted benefits from x86 virtualization. In short, we’ve found that...

- More than 60% of our survey respondents cite higher hardware utilization as a significant benefit. For most organizations, the average utilization rate of an x86 server hovers around 6%. With virtualization, this figure can be improved to 50% or more, meaning that customers get a bigger bang from their hardware buck.
- More than half of the respondents either agree or strongly agree that x86 virtualization makes their x86 server infrastructure easier to manage. With virtualization, many routine management chores are automated, allowing fewer administrators to manage more systems and workloads. We see roughly the same numbers of respondents saying that virtualization has made it easier for them to meet their SLA requirements.

The majority of x86 customers clearly also recognize that virtualization is saving them money, and this number has increased over just the past year. In 2007, 56% of our respondents said that x86

virtualization was resulting in significant savings. In 2008, that number climbed to 60%. At the same time, the number of respondents who believed that virtualization did *not* convey a cost benefit dropped almost in half; we can safely assume that many of them became believers in the cost efficiency of virtualization.

Virtualization Benefit: Saves Us Money

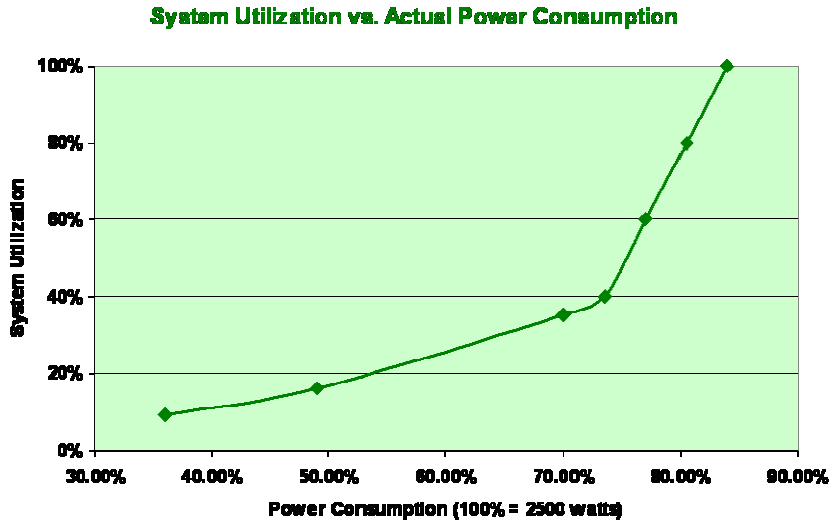


Cost is one of the big reasons why customers are using x86 systems as a host for mission-critical applications. Estimates concerning just how much money can be saved by going to a virtualized x86 infrastructure vary widely. Most customers we’ve spoken with who have performed internal cost studies estimate that their savings range from 20 – 35% of overall capital and operating expenses.

Overall hardware spending drops significantly, and oftentimes software costs drop as well when licenses are combined and redundancies are eliminated. Internal administrative and support costs are also reduced, although most companies do not accurately measure these effects – they tend to underestimate the impact of raising the application/administrator ratio on total costs. They also tend to overestimate future personnel needs and the costs associated with those.

Virtualization: A Pragmatic Shade of Green

There is yet another category of cost that is impacted by virtualization. One of the most highly publicized problems facing data centers today is the high cost of the energy used to power and cool systems. The chart below is instructive, as it depicts the actual power consumption of an x86 system, as measured by a GCG client. This picture illustrates several points: First, a system at idle is using significant power – about 35% of its total potential power draw. When CPU utilization hits 40%, the system is drawing roughly 75% of its power max. When system utilization moves above 40% or so, there is very little additional cost in terms of power.




When system utilization moves above 40% or so, there is very little additional cost in terms of power.

To make the point more clearly, the additional 60% of usable capacity of the system only costs an additional 25% in terms of electricity. Given that the average CPU utilization rate of an x86 server is somewhere around 6%, a consolidation ratio of 8:1 is easy to achieve and will provide considerable savings in terms of energy use. Simply virtualizing x86 workloads will provide much more overall

energy savings than just adding more energy-efficient servers. Of course, bringing in new and more efficient hardware coupled with virtualization is the most energy-efficient strategy of all.

Virtualized x86 will become the dominant x86 usage model in our organization



Virtualization also pays dividends on the data center floor. Consolidation allows customers to decommission older equipment and remove it entirely. Many facilities that find themselves space-constrained are turning to virtualization as a safe and easy alternative to building or buying more data center floor space. As we tell our clients, virtualization is the best and easiest route to fixing energy and space problems.

Perhaps the strongest indicator of the rate at which customers are adopting x86 server virtualization is in this chart where we asked customers whether virtualized x86 systems will become the standard in their

organization. In 2007, 50% of customers agreed or strongly agreed with this statement. In 2008, that number rose to 56%. Note that the number of customers who *disagreed* with this statement dropped as well – a signal that these respondents are becoming virtualization converts.

With sophisticated virtualization now in place, customers can take full advantage of the x86 platform for mission-critical workloads. Every major ISV now has Windows or Linux versions of their enterprise

applications, including entire ERP suites from industry heavyweights Oracle and SAP. Even the most sophisticated software packages are available in various Linux and Microsoft server versions.

Two companies in particular seem to be going after this burgeoning ‘mission-critical x86’ market segment. IBM and VMware, both leaders in their respective markets, have also established a close partner relationship. One of their joint goals is to make it easier for enterprise customers to move mission-critical applications to x86 systems. According to them, this will help customers reduce costs, increase IT flexibility, and even increase application availability. In the remainder of this research report, we take a look at what these companies bring to the table both separately and together.

Enterprise x86: Not a Commodity Play

Some vendors and industry pundits seem to believe that all x86 servers are created equally. We’ve always disputed this view, as our research shows that customers perceive significant differences between hardware platforms, even if they run the same operating systems and use the same basic parts. It’s much like automobiles...they all have tires, steering wheels, engines, and radios, but there are huge differences in how well each vehicle is suited to a particular task and, even with similar cars, how well each owner is served by that particular make of car. There are also considerable differences in how vendors approach the market and how they craft their product strategy to solve customer problems.

IBM is synonymous with data center and mission-critical computing. They just unveiled a new initiative, Dynamic Infrastructure, that puts forward IBM’s take on the changing world and how they will help customers deal with it. IBM sees a world that is constantly accelerating, with more interconnections being established between people, organizations, governments, and a myriad of newly-smart objects. With Dynamic Infrastructure, these connections and tools can be harnessed to improve life for everyone. According to IBM, their mantra in this initiative comes down to three easy-to-understand tasks: Reduce Cost, Improve Service, and Manage Risk. The purpose of Dynamic Infrastructure is to devote IBM’s considerable technology assets towards helping customers accomplish those goals.

Part of this, of course, is systems and solutions. In the x86 market IBM’s real differentiation is their focus on custom technology and their ability to provide a range of products that address different enterprise needs. IBM is the only major company that offers large-scale SMP x86 systems. Their model x3950, for example, can scale to a massive 16-socket (96 total cores) 1TB memory system; this can handle large numbers of even the most demanding enterprise applications. IBM, of course, also has a wide selection of smaller systems and a full line of blade products, all of which feature advanced technology aimed at supporting mission-critical workloads. In fact, IBM has been one of the leaders in the blade market with their BladeCenter family of products. These systems provide high compute density and the widest variety of blades in the industry.

While almost all of their competitors use standard motherboards and components, IBM invested heavily in designing their own system technology, which they refer to as their Enterprise X-Architecture. IBM is now shipping the fourth generation (eX4) of this chipset. The payoff for enterprise customers is found in higher performance, higher availability, and better system and environment management. On the performance side, for example, IBM’s eX4 chipset increases processor-to-processor bandwidth by providing an additional local cache directory within the chipset – helping to reduce latency and increase throughput. Cache misses are redirected to the chipset to locate the missing data, as opposed to first requesting the data from each CPU in succession before looking to main memory, as happens in most other x86-based systems. Thus CPU-to-CPU communication is reduced, according to IBM, by as much as 75%; more cycles are available to process real work. This extra

processing power gives the systems the headroom necessary to support multiple important applications – even under peak demand conditions.

IBM and VMware put this performance advantage to the test in a 2008 Microsoft Exchange benchmark. A four-socket IBM x3850 with Intel quad-core processors, running 8 virtual machines with VMware's ESXi 3.5, was able to host **16,000** Exchange mailboxes. This result more than doubled the next highest score for a single physical server. According to IBM and VMware engineers, there was still some headroom left in the system, meaning that in the real world, this solution may be able to scale even higher. Of course, with Intel's upcoming six-core Nehalem chips, we can expect to see higher numbers down the road.

IBM's custom technology also pays dividends in the area of system availability which, to our thinking, is in many ways more important than raw performance – particularly when discussing critical enterprise applications. IBM x86 servers have a number of availability-related features and capabilities that are simply not found in non-IBM commodity-based servers.

For example, IBM systems offer operating system independent memory mirroring, which helps reduce failures due to soft memory errors. Their Predictive Failure Analysis (PFA) warns of component failures before they happen. IBM also offers a large number of hot-swap components (more than most

competitors) which, when coupled with PFA, allow customers to change out failing components while the systems continue to operate.

In another recent GCG x86 survey ("Enterprise x86: What's Important"), 90% of respondents cited "High Quality Sophisticated Systems" as their most important decision factor. "System Availability Features" was the second most important factor. System purchase price came in seventh...

Innovative server technology is an important component in IBM's bid to become the mission-critical x86 platform of choice. IBM is taking this same approach in crafting the enterprise storage side of their hardware product line. Taking advantage of VMware's full suite of enterprise virtualization technology requires a shared storage infrastructure that is capable of supporting high I/O rates, providing redundant data paths to aid availability, and seamlessly scaling total storage according to enterprise growth at a reasonable cost.

IBM's DS4000 and DS5000 are 4Gbps Fibre Channel disk storage systems that are built with scalability in mind. With the DS5000, for example, customers can begin with a single enclosure containing up to 16 SATA or Fibre Channel drives and, over time, expand their environment to include a total of 16 enclosure with 256 drives (up to 168TB of raw capacity and 512 storage partitions). As these upgrades are performed, the existing data remains intact. The ability to add additional controllers as storage needs grow means that the controllers never become performance-choking bottlenecks.

As we noted in our discussion of servers above, for the mission-critical application, availability is perhaps even more important than performance. The DS4000/5000 family of products are chock full of features designed to ensure that the systems stay operational under almost any conditions. Like other storage subsystems, these systems support the full range of RAID protocols, and most components are hot swap or redundant. However, IBM goes an extra step with their Dynamic Capacity Expansion feature, which allows additional disk drive units to be added without interrupting existing workloads and business processes. Performance can be optimized using their Dynamic Segment Size feature; this

allows users to change the data stripe size on the fly, again without disrupting workloads or losing access to data. Customers can even change RAID levels dynamically.

IBM also offers a full line of SAN (storage area network) products that give customers centralized management and provision individual storage arrays, even those from other vendors. A solid and scalable SAN is a crucial factor in enterprise virtualization; it allows all systems to access storage on any or all attached storage repositories. The IBM SAN Volume Controller (SVC) combines storage from different arrays into pools of capacity and enables storage virtualization and consolidation, with centralized management. This makes it easier for administrators to set up and manage tiered storage shares for various systems and applications based on performance needs. Combining all of your storage into centralized pools eliminates the storage 'silos' that lead to over provisioning. With a SAN, customers can provision only the storage that is needed today, and dynamically add more capacity as necessary to handle future demand.

IBM has also invested in making system management and monitoring easier for customers – a must when hosting mission-critical applications in complex infrastructures. Their IBM Director management suite lets customers manage any IBM system (including RISC systems and even mainframes) from a single console. Director allows customers to easily manage and configure systems using a common and consistent set of tools.

A relatively new entry into IBM's system management product suite is their Active Energy Manager product. Active Energy Manager (AEM) monitors the actual power consumption, in real time, of IBM x86 systems. Energy usage is becoming a higher priority in most data centers these days, and AEM will help managers understand usage patterns and needs over time. It will also alert them to energy capacity issues before they become critical. AEM also gives customers the ability to cap power draw by servers or racks, and to optimize data center power/thermal requirements by suggesting workload changes. For example, using AEM, a manager may discover that systems in a particular rack are barely idling, but still powered up. He may then decide to move those workloads to another, more highly utilized set of servers, and completely power down the vacant rack.

“System Management & Management Features” are also very important to x86 purchasers. Sophisticated systems management automates routine tasks and allows more workloads to be managed by existing personnel – thus reducing overall TCO.

All of the above attributes serve to separate IBM's x86 mission-critical solutions from the crowd. A common theme in IBM's system technology, especially in x86 servers and storage, is harnessing advanced technology from their other platforms and moving it into less exotic and far less expensive systems. Much of the technology discussed above has roots in IBM mainframe and supercomputer systems; it has been adapted to improve the performance, availability, and manageability of their industry-standard x86 systems. This should give them an inside track with customers, both large and small, who are looking for systems to host mission-critical workloads.

VMware: Virtualization Leader

In x86 virtualization, there are really only three major vendors right now: VMware, Microsoft, and Citrix. In terms of market share, VMware is far and away the leader, with around 75% of the total market. They were the first vendor to offer a truly sophisticated x86 virtualization product, allowing

customers to run multiple Linux or Windows workloads on single servers – without compromising availability or security. VMware and virtualization changed the economics of x86 computing, making it much more efficient and economical. While this trend is nowhere near completion, it has already had a huge impact on the industry.

VMware is holding their lead over competitors due to their constant pushing of the technology envelope. VMware has added more capabilities and made their products more technically sophisticated, yet easier to implement and manage.

At its core, VMware's virtualization products give customers the ability to consolidate 5, 50, or more separate workloads onto a single server, depending, of course, on the specific applications and usage characteristics. They can easily push server utilization rates from mid-single digits to 50% or more, resulting in a huge savings on hardware purchases alone. Rolling out new applications is also a much quicker process using virtualization. A new virtual machine can be fully provisioned in minutes, while the alternative – purchasing, installing, and provisioning a new physical server – can take days or weeks.

VMware makes it easy to assign system resources to a particular workload to ensure it has enough capacity to handle business needs. It can reassign resources to handle workload spikes, giving a particular application more processing power or I/O on a temporary basis to cover higher usage – without operator intervention or interrupting other workloads on the physical system. VMware ESX with VMotion can even move entire running applications from one physical server to another – again, without human intervention or inconvenient reboots. This feature is most commonly used to take workloads off of a system so it can be taken down for maintenance purposes, or to move applications that need more hardware resource to larger systems.

Virtualization can also help increase application and system availability – a vital concern to any enterprise that is hosting mission-critical applications on x86 servers. VMware has enhanced HA to give customers true fault tolerance. In past years, this capability was very expensive to achieve and used for only the most important workloads (like ATM networks, for example). But VMware's new Fault Tolerance (FT) uses vLockstep technology to keep two VMs exactly in synch; if the primary fails, then the secondary VM running on a separate host automatically picks up the load. This is much like traditional clustering, but significantly less expensive to implement and far easier to manage and configure.

As we discussed above, one of the main problems customers are facing is excessive data center energy use. Some IT shops have simply run out of capacity to add more juice, while others are aghast at skyrocketing electricity bills. Single-application servers are energy hogs, using lots of power even when they are just idling. VMware is addressing this problem with their Distributed Power Manager product, which can automatically and dynamically relocate workloads onto fewer systems during off-peak hours. The process takes only seconds to execute and is fully scriptable; the unneeded hosts are then powered down completely. When user demand rises again, the process can be reversed to ensure that Service Level Agreements are met.

IBM + VMware

IBM and VMware have a much deeper and closer relationship than it may appear to the casual industry observer. Both companies have partnerships with competitive players and products of their own that could be seen to conflict with each other. IBM has partnerships with Microsoft (Hyper-V virtualization) and Citrix (Xen), along with their own virtualization and virtualization management products. VMware

is allied with every x86 system vendor, including IBM's archrival Hewlett Packard, and also sports their own competitive virtualization and workload management offerings. Even with all of these potential conflicts, there seems to be a special relationship between the two companies.

IBM was the first system vendor to endorse VMware and the first to provide comprehensive support for VMware products on their systems. IBM was also the first vendor to offer and fully support VMware on their blade systems and the first vendor to demonstrate VMware's lightweight embedded hypervisor. Both companies seem to gain many benefits from their close cooperation. IBM was VMware's first joint development partner, and they have worked together on R & D since the beginning of their partnership. Their goal is to develop end-to-end comprehensive solutions that take advantage of their respective strengths and the synergy that has arisen from their combined development and research efforts.

The partnership is paying off, particularly for large enterprise computing customers who need a seamless solution that reduces implementation time and makes it easy for them to realize the benefits of moving to a fully virtualized infrastructure.

GCG Summary & Recommendations

We believe that mission-critical applications will be hosted on x86 systems in greater numbers in the future because of the performance, availability, and ability to meet stringent service levels. We see more and more customers (particularly large enterprises) proving to themselves that these solutions are up to the task. They have found that modern x86-based solutions provide the availability, scalability, and performance necessary to handle their most important workloads. The advent of VMware's sophisticated and comprehensive x86 virtualization solutions coupled with IBM's enterprise-class hardware and software gives mid-sized and large data centers confidence that these solutions can meet the needs of even the most demanding computing environments.

The real driver behind this trend is simple economics. An IBM and VMware virtualized x86 solution offers the required size, performance, availability, and manageability to run mission-critical workloads – but, in many cases, at a much lower price point than alternatives. The troubled economic climate will accelerate this trend as more companies look to cut costs and increase efficiency.

Entire contents © 2001- 2009 Gabriel Consulting Group, Inc. All rights reserved. This document may not be reproduced or transmitted in any form by any means without prior written permission from the publisher. All trademarks and registered trademarks of the products and corporations mentioned are the property of the respective holders. The information contained in this publication has been obtained from sources believed to be reliable. Gabriel Consulting Group does not warrant the completeness, accuracy, or adequacy of this report and bears no liability for errors, omissions, inadequacies, or interpretations of the information contained herein. Opinions reflect the judgment of Gabriel Consulting Group at the time of publication and are subject to change without notice.

Gabriel
CONSULTING GROUP

phone / 503.372.9389
gcginfo@gabrielconsultinggroup.com
www.gabrielconsultinggroup.com