sageza

Instant Insight
March 5, 2010

IBM Announces Latest Offerings in eX5 Portfolio

By Clay Ryder

IBM has announced three new servers that are based on the fifth-generation IBM Enterprise X-Architecture chip (eX5). These are the four-processor IBM System x3850 X5, the BladeCenter HX5, and the System x3690 X5, which the company stated will be the most powerful two-processor server on the market. Each of these eX5 servers is equipped with a new independent memory-scaling technology, known as MAX 5, which allows processors on eX5 systems to access extended memory very quickly and enables these servers to offers six times more memory than comparable x86-based servers.

In addition to MAX5, IBM's new eX5 systems offer additional features that seek to improve the performance, cost, and flexibility for x86 workloads:

- eXFlash, a next-generation flash-storage technology, replaces previous generation storage and can slash storage costs substantially as each eXFlash can replace 80 JBODs and associated hardware and cabling.
- FlexNode, a physical partitioning capability, allows organizations to change their systems configuration from one system to two distinct systems and back again as desired. This enables organizations to run infrastructure applications by day and larger batch jobs by night on the same system.

The eX5 systems take advantage of integration with IBM middleware to create a virtualized environment providing a flexible, highly scalable system that can reduce the number of physical servers needed to support a given workload. IBM's Systems Director management suite has been upgraded to support eX5 technology and will allow users to pre-configure servers, remotely re-purpose systems, and set up automatic updates and recoveries. The company stated that its new eX5 servers are the result of a three-year engineering effort to improve the economics of operating enterprise-sized, x86-based systems.

Pricing/Availability

The IBM System x3850 X5, BladeCenter HX5, and System x3690 X5 are being previewed this week at the CeBIT trade show in Germany and will be officially available later this month and throughout the year. No pricing details were released.

Net/Net

At times it is difficult to get excited about Industry Standard x86-based servers. The market is saturated with solutions that seem to universally benefit from ever-increasing CPU performance combined with a price/performance ratio that continues to rise towards the stratosphere. All told, more applications are finding their way onto this platform every year as legacy UNIX and other systems are retired in favor of this widely deployed low acquisition cost option. The platform offers choice of operating system and components from thousands of suppliers, and now with virtualization, utilization is going up, and acquisition price points are going down. Overall, it sounds like a commodity market. So why get all worked up about a vendor's latest entry in this well populated market segment? The reasons are simple: this is NOT a commodity market, virtualization is here to stay, and disk drives are too darn slow.

The Sageza Group
931 Campbell Avenue
Los Altos CA 94024
510·675·0700 fax 501·640·8105

While x86 processors as a component could easily be considered a commodity, a processor alone does not make a server. All of the components on the motherboard and rack/chassis along with various interconnections have a substantial impact on the performance and capability of the server. While one industry-standard server would be expected to execute the same software as another, the performance and scalability of the systems is not automatically the same. Commodity CPU yes, commodity and hence undifferentiated server, no.

The growth in x86 CPU performance is undeniable. Quad-core and higher systems with lightning-fast clock speeds have achieved an astonishing level of performance at very compelling price points. However, to effectively harness this performance, the system performance as a whole must be balanced across all of its component parts in order to cost-effectively support the application workload. As computational performance has risen, there has been a corresponding gap in overall system utilization as newer and faster systems are limited by the mundanity of relatively slow I/O access, memory swaps, disk reading and writing, network access, etc. To address this utilization gap, many have embraced virtualization to get more virtual servers out of the physical server. But this success in modestly raising CPU utilization has caused a new scaling constraint, one of limited memory.

Virtualization allows many more logical servers and workloads to be discretely serviced by a physical server. However, each of these workloads requires RAM, and the total amount of RAM is limited by the number of DIMMs and memory interconnects to the CPU that can be installed on a motherboard. Historically this has limited most x86 servers to 256GB or less, a shortcoming that MAX5 seeks to address. By increasing the amount of RAM available to the CPU up to a maximum of 1536GB, the server can support a larger virtual machine footprint and hence more workloads. This can help make use of computational resources that are otherwise idle or saturated with workloads that are more memory- than computation-intensive. Increasing the yield of virtual servers supported per physical server can reduce software licensing fees for organizations with per-socket licensing regimes. The number of users supported on a server can increase, but without an increase in socket count and hence corresponding license expense. To our way of thinking, this gives eX5-based systems a compelling advantage over traditional solutions, one that not only is technically crafty, but yields financial dividends as well.

Lastly, with increased computational prowess and sufficient memory in place, the last leg of the balancing trifecta becomes evident: access speed to out-of-memory storage. In the grand scheme of things, disk drives are just much slower than memory or CPU cache. Although many steps have been taken to improve spindle-based storage access speed, it still remains the turtle to the computational hare. Advances in flash technology have sought to address this reality, and this is where eXFlash plays its hand. Besides its higher IOP rating, each eXFlash can replace 80JBODs, which translates into a substantial reduction in energy, floor space, and wiring requirements. Again, the performance improvement is notable, but the underlying acquisition and operational economics are the most compelling aspects of this technology.

So all told, we believe that eX5 illustrates once again that just because a server is based upon x86 processors it is not automatically a commodity solution. The unique scaling and performance technologies in these offerings serve notice that a holistic and balanced systems approach can yield very different results than a solution based simply upon assembling industry-standard technologies. We believe organizations that are turning to workload optimization and systems efficiency as part of their strategic IT course will likely find much to rejoice in with the latest eX5 servers. Likewise, for those who are merely seeking to achieve a higher degree of vertical scaling from x86 solutions, eX5 redefines the current limitations on single system scaling in a notable fashion, one which we believe will be well received.

The Sageza Group
931 Campbell Avenue
Los Altos CA 94024
510·675·0700 fax 501·640·8105