Can IBM Revitalize 8P with x86?

The 8 socket market was predominantly proprietary RISC/UNIX, but do today's scale up workloads signal a rebirth of x86 8P?

Executive Summary

The need for greater agility is driving IT strategies. The importance of scale up applications is growing, but a different innovation cadence is needed. Applications and data streams are becoming more complex, requiring additional processing power and larger memory footprints for these scale up applications. With cloud-based applications, large virtualization pools, and ERP/database all running on x86 platforms, there is a need for 8 CPU servers that can run standards-based backend workloads or create a large consolidation point for virtualized workloads. The scalability of 8 CPU x86 servers combined with their robust platform availability make them an excellent choice for deploying critical applications. In an era of rapid innovation and changing IT patterns, these products can be the right choice for demanding applications.

Application Needs Continue to Scale Up

There was much focus on scale out over the last decade, as x86 servers have helped to drive that trend. But not all workloads scale out. Many workloads, like business intelligence, analytics, ERP, and CRM currently require a larger single instance footprint and do not lend themselves to a scale out environment. While workloads like web serving, HPC and some scale out databases can be spread out and handled by a larger number of parallel nodes, classic database-driven applications require the consistency and low latency that forces them onto a single instance for best results. To drive the best performance and scalability for these single instance applications, a larger number of CPUs (and cores) needs to be employed— along with a very large memory footprint. This combination of scalability and performance is perfect for today's increasingly data-driven applications.

With business cycles accelerating alongside the increasing need to stay ahead of the competition, businesses seek to be more agile. For businesses to take advantage of these fast moving opportunities, managing and responding to data is critical. Decision support systems demand the scalability of 4 and 8 CPU systems, as the applications that aid in this process all favor single instances with large amounts of resources. The memory footprint for many of these applications is a key performance driver: loading more of the application or dataset into memory greatly accelerates performance versus trying to retrieve the data from the slower I/O subsystem.

In virtualization and cloud environments, the desire to reduce cost-per-VM is driving businesses toward pooling more instances on fewer large and robust physical hosts. By scaling up consolidation, businesses are able to drive better ROI.

Key applications for supporting today's agile businesses are favoring industry-standard platforms based on 4 and 8 CPU scalability.

- **Relational database management system (RDBMS)** vendors have seen the steady shift to x86 over the years, driven mainly by the maturation of Linux (and Windows Server for Microsoft SQL Server). We are at a point today where almost all of the growth in these applications is driven on standardized platforms with almost all being scale up deployments.
- **Enterprise resource planning (ERP) and customer relationship management (CRM)** applications are tied tightly to backend databases. One major ERP vendor stated that 80-90% of their customers run on Windows Server platforms for both the database and the application.
- **In-memory databases, applications, and appliances** are finding the large memory footprint of x86 servers to be on par (or better than) offerings of RISC platforms. Some, like SAP HANA, are offered only on x86 platforms.
- **Highly consolidated datacenters** with high VM density on larger physical hosts (versus larger numbers of smaller hosts in a cluster) are finding less complex management by using highly scalable systems.

Figure 1: Application Growth and Scalability

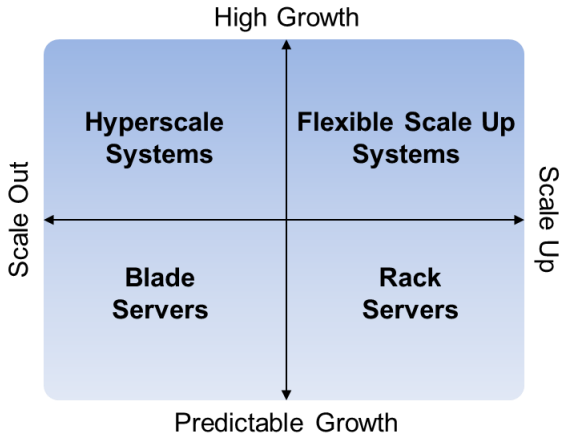

In stable growth environments, traditional rack and blade servers are fine depending on whether applications will need to scale up with rack systems (larger memory footprints, more/faster CPUs, or more internal storage) or scale out with blades (more dense physical systems to spread the workload).

High growth environments need different systems to match rapidly changing needs.

Scale out workloads typically feature purpose-built, stripped down dense server solutions built around dual socket architectures (think Google or Facebook). But these scale out workloads typically have lower system requirements, because the application load can be spread across multiple system images. The hole in the market was typically in the high growth/scale up quadrant, where businesses with unclear growth trajectories

were forced either to overbuy up front or to scrap deployed systems and replace them with larger, more expensive systems as workloads grew. Neither of these choices helped optimize the budget. Choices for flexible scale up architectures really did not exist.

Standardized Systems Are the Clear Choice

For years, proprietary platforms were the only choice for highly scalable enterprise applications. But as x86 systems took over the lower end of the market, their pace of innovation (12-18 months for x86 vs. 3-4 years for RISC) allowed them to catch up with proprietary systems' capabilities. The open nature of x86 systems brought an extra element of total cost of ownership (TCO) savings. They lowered acquisition costs for both the platform and software, and they also eased operational costs through better manageability and power consumption.

While pockets of RISC remain in areas where applications are very specialized and difficult to port (like telecom), these proprietary systems have seen a slow, gradual decline as x86 has taken over more of the market. Likewise, Itanium has slowly lost support: its largest proponent (HP) indicated it would move away from Itanium in the future, possibly accelerated by Oracle's decision to stop development on Itanium. Momentum for scalable applications is moving to x86.

- **Resources** The CPU/core density and memory footprints are ideal for scale up applications, because of their ability to handle large numbers of simultaneous requests and cache more data into system memory.
- **Applications/APIs** Consistency with frontend applications and tools of x86 platforms makes development, integration, and support easier.
- **Virtualized Environments** The cost factors that are driving consolidation, virtualization, and cloud lend themselves best to cost-effective x86 systems.
- **Innovation** With an innovation cycle of 12-18 months, x86 platforms are in a better position to keep pace with the demands of quickly changing markets.
- **Economics** The volume and velocity of innovation bring better economics for standardized systems, including development, deployment, management, and support.

Scalable Applications Require Resilient, Scalable Platforms

x86 platforms have become faster, more scalable, and more reliable—which is why most applications in the datacenter are running on these platforms today. However, for the most scalable applications, only the latest Intel Xeon E7-8800/4800 v2 series processors deliver both the horsepower and the highest level of reliability required. Highly scalable applications use multiple software threads to handle large numbers of concurrent users or highly parallel instruction processing. And in heavily-virtualized environments, the ability to scale up to a high number of cores enables better efficiency for densely-packed virtual machines. With more cores, VMs and cloud instances can take advantage of thread affinity and spend less time "core jumping".

Memory density and speed are critical for these enterprise applications. With the ability to load terabytes of data and applications into memory, platforms can scale up more easily as user and application loads increase. When more data are held in-memory, fewer swaps are required out to system drives, which will ultimately boost performance (even when compared to the access times of the fastest enterprise-class SSD drives).

While performance and scalability are important considerations for highly scalable applications, in reality one of the most critical facets of the platform is **availability**. While x86 servers have made tremendous strides in the past decade toward closing the gap with more expensive proprietary systems, not all x86 platforms can deliver the same level of availability. Today's high end 4 and 8 CPU servers built on Xeon E7 processors have RAS features that exceed standard 2 CPU x86 servers—putting them on par, if not above, RISC/Itanium offerings. In addition to the resiliency that Intel Xeon E7 platform technologies provide, manufacturers like IBM add additional RAS features to push system uptime and fault tolerance even higher.

The IBM x3950 X6 Server

Clearly not every workload in the datacenter requires the strength of an 8 CPU server, which is why most enterprise workloads are running 2 CPU servers. But these commodity servers do not meet the requirements for highly scalable applications which require larger CPU, memory, and I/O configurations along with additional RAS and serviceability features that boost uptime. This is where systems like the IBM x3950 X6 help deliver the capabilities of a more expensive proprietary system—with a TCO savings and innovation cycle that more closely resembles what businesses see in their commodity 2 CPU systems.

Figure 2: IBM x3950 X6 Chassis

Performance and Scalability

The product of years of innovation, IBM recently introduced the sixth generation of Enterprise X-Architecture technology. IBM has a distinct advantage, as this is not only the **first** 8 CPU server built on Intel's most recent E7 CPU (the Xeon E7-8800 v2), but it is also the **only** one, as Dell, HP, and Cisco have yet to introduce 8 CPU products.

The IBM x3950 X6 is designed as a full 8 CPU system with a modular capability that allows a business to purchase the right amount of performance and scalability up front, and then easily add more capacity in a modular fashion when applications require additional scalability.

The system is built around a flexible chassis that has slots for 8 Compute Books. Each Compute Book is a self-contained module holding an Intel Xeon E7 CPU and the associated memory, allowing the system to scale up evenly as CPU demands increase. Managing the performance expectations of scalable applications is typically handled through the CPU-to-memory ratio. So tying the memory to the CPU complex is helpful not only from an architectural standpoint (common memory bus) but also from a system planning standpoint. Each Compute Book supports up to 24 DIMMs per CPU, giving the module the ability to address 1.5TB of memory per CPU, for a total system footprint of 12TB of memory.

Figure 2: Individual Compute Book

In addition to the Compute Books, the system also features two Storage Books to handle up to 16 hot-swap hard drives or 32 SSDs. Equally important are the 6 I/O Books (2 standard), allowing for the expansion of up to 18 PCIe I/O cards plus 2 additional Mezzanine cards.

Together, these modular options allow IT organizations to configure a base system to meet existing needs and then easily scale up the system as those needs change. Contrary to most existing x86 systems that are deployed and then never touched again, the IBM x3950 X6 is designed to be updated. It is more like the experience that businesses have with scalable RISC/UNIX systems—without the price premium.

To help drive greater performance, IBM has integrated a feature called eXFlash Memory Channel that allows up to 32 eXFlash DIMMs to be installed directly on the system memory bus. With the ability to write block data directly to the eXFlash DIMMs, the system allows for up to 12.8TB of low-latency storage to be tied directly to the system bus, greatly accelerating data reads and writes. In many cases, this memory technology may allow businesses to bypass SAN storage completely, thus creating a more self-contained application solution within a single chassis.

The flexibility of the x3950 X6 is demonstrated through the IBM FlexNode capability which allows an 8 CPU system to be partitioned into two 4 socket servers. This partitioning gives businesses a more flexible way to deploy servers. With FlexNode, a single chassis can be split, allowing for 2 clustered 4 socket servers to be deployed. This flexibility reduces capital cost and physical management costs, because there are fewer physical assets to manage.

Businesses with a need for future scalability (but unclear growth patterns at the point of purchase) can deploy the 4 processor IBM x3850 X6 with Intel Xeon E7-8800 v2 CPUs, then upgrade the server later to a x3950 X6 (8 processor) while maintaining the original serial number of the server. This upgrade path makes servicing, management, and asset tracking easier, so businesses can leverage their initial investments better.

In the future, as new technology is available such as new CPU models, IBM x3950 X6 systems can be easily upgraded as well—simply by replacing the pluggable compute books, again preserving the original system serial number and the initial investment in the platform.

High Availability/RAS

Scale out applications achieve availability through system redundancy. But that requires software, network connectivity, switches, and other components—all of which add latency and cost. Scale up applications achieve this level of redundancy through system-level functionality that is integrated directly into the server to bring lower latency along with redundancy. These RAS features are designed not only to keep the server online but also to minimize downtime from scheduled service events.

The IBM x3950 X6 begins with a strong set of RAS features integrated into the system as part of the Intel Run Sure technology. IBM then builds onto those capabilities with additional RAS features designed to enhance uptime even more.

Figure 4: Uptime Features and Capabilities

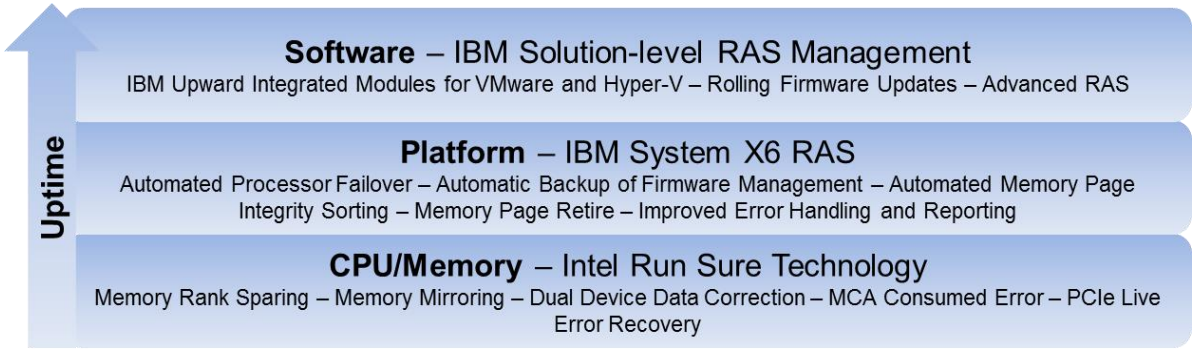

The Run Sure Technology features bring the Machine Check Architecture that RISC and Itanium systems have had for years into an x86 platform. These features help the

Xeon E7 systems identify and sustain operation despite problems that may have resulted in blue/purple screen errors on single or dual socket x86 systems.

Then IBM adds to the reliability capabilities with features designed to track issues in the memory and then report back to the system about memory page locations for the highest reliability, as well as those locations that should be avoided. Although this feature is native to some of the more modern operating systems, IBM firmware is optimized to allow it also to work with older legacy operating systems that do not feature page retire support.

Then, on top of all of these hardware features, IBM adds additional integration features for virtualization environments that provide higher levels of integration and manageability. These features tie in the management functionality of the x3950 X6 server with the management consoles of Microsoft and VMware management tools.

IBM Support

Support for IBM and Lenovo systems is provided worldwide by IBM Global Services, the world's largest business and technology services provider, employing over 190,000 people across more than 170 countries. As a worldwide organization, IBM Global Services has deep experience in both software and enterprise architecture, making them a critical partner in keeping scale up platforms running smoothly for years.

True Customer Benefits for Scale Up Applications

To help organizations adapt to rapidly changing business environments, compressed cycles, and accelerated decision making, companies need to be agile. They need to move quickly, with purpose. Having the right systems in place to support these actions is critical.

Cost-optimized, fit-for-purpose IT solutions like the IBM x3950 X6—with eight CPUs and massive memory scalability—help consolidate large streams of data and run scale up applications more effectively. The system performance and scalability help close the gap between changing market demands and having actionable data in hand to make the right decisions.

Although businesses are increasingly moving to cloud architectures (be it public, private, or hybrid), the enterprise applications and databases that will feed many of these cloud delivery systems will remain in the corporate datacenters, closer to the business. But because they need to interface with cloud services, a common underlying architecture, operating system, and set of APIs help ease the communications between these two parts of the business. This is why x86 platforms running the backend applications makes so much sense when connecting to web or cloud applications.

Through the transition to x86 from proprietary platforms, businesses took advantage of a faster cycle of innovation where new technologies are introduced every 12-18 months instead of every 3-4 years. With the IBM x3950 X6, the modular chassis enables

multiple generations of technology to be supported over the years. Thus, a business can take advantage of future technologies through the modular swap of Processor, Storage, or I/O Books instead of having to recycle the whole system in a forklift replacement. Because many businesses segment their capital expenditure and upgrade budgets, an upgrade down the road may not face the same administrative hurdles—especially if the asset tag is not changing on the chassis. This provides an opportunity for a business to better plan for a future capacity upgrade at the time of acquisition and better budget for the solution over its useful life.

System modularity also helps reduce total cost of ownership by reducing downtime for servicing. Modules can be swapped quickly, bringing the system back online with minimal interruption to operations.

Together, all of these business benefits clearly make the IBM x3950 X6 a compelling choice for hosting scalable applications.

Summary

Businesses have come to rely on standardized technology for their enterprise applications and virtualized environments, with only highly scalable systems not being addressed adequately by x86 platforms. But with today's enterprise applications changing at the same time that proprietary systems are waning, it is harder than ever to recommend a proprietary system for scalable applications.

With so many new applications running in virtualized or cloud environments, there are benefits to using a common x86 architecture between the scale out frontends and scale up back-ends of the datacenter.

This server is a viable choice that has the economics and innovation of commodity x86 servers combined with the scalability, performance, and RAS of proprietary systems. Businesses who are investigating platforms for their scalable applications need to investigate the IBM x3950 X6.

For more information: http://www.ibm.com/systems/x/hardware/enterprise/x3950x6/

Important Information About This Paper

Author
John Fruehe, Senior Analyst at Moor Insights & Strategy

Review
Patrick Moorhead, President & Principal Analyst at Moor Insights & Strategy

Editor
Scott McCutcheon, Director of Research at Moor Insights & Strategy

Inquiries
Please contact us here if you would like to discuss this report and Moor Insights & Strategy will promptly respond.

Citations
This note or paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

Disclosures
Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. This paper was commissioned by IBM. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER
The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness or adequacy of such information and shall have no liability for errors, omissions or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and

forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.