Is “Scalable Blade” an Oxymoron?

Blades deliver density but required a tradeoff in scalability until now

Executive Summary

Blade servers solved the problem of density for businesses, but they did so with compromises that limited their appeal for many applications. As complexity of both applications and environments grew, blade infrastructures were limited in their effectiveness for the task. Customers were relegated to lower-density rack servers for their more scalable applications. And while purpose-built systems allowed a business to tailor to their specific needs, they created more management complexity.

New flexibility is required to allow modular blade architectures to keep pace with the types of applications that customers need to drive a more agile business. The IBM Flex System X6 family of compute nodes are worth serious consideration for businesses trying to balance the needs of this new generation of agile applications with the desire for more standardization and modularity to hold down total cost of ownership (TCO).

The Market is Changing and Traditional Blades Can’t Keep Up

Emerging trends in both data and how it is being delivered are creating fundamental changes for IT, as these two forces barrel headlong into the age-old problem of datacenter economics. “Blade 1.0” designs excel where growth and applications are both predictable and have lower resource requirements. But an emerging set of applications and use cases are pushing the envelope on system scalability and predictability beyond the limits of traditional blades. Blade infrastructure needs to reinvent itself to handle density and scalable applications within the same infrastructure while also keeping management and operations costs in check. This drives a need for “Blade 2.0” designs, or modular converged systems.

Figure 1: Market Dynamics

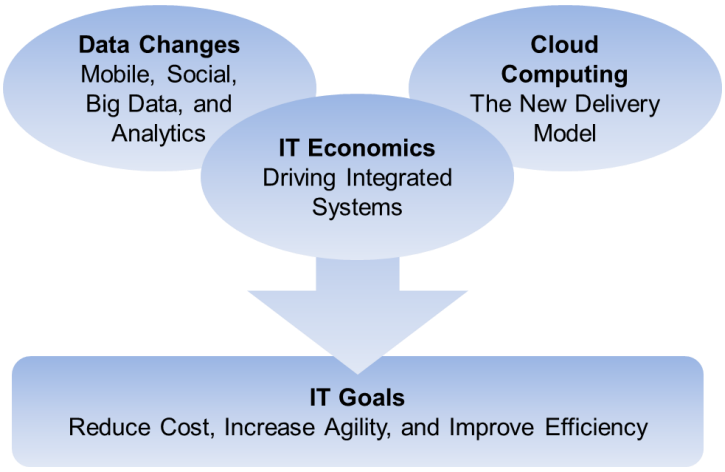

Data Changes

Applications are more "data-driven" today, communication patterns are becoming more complex, and there is a shift towards social interaction and heterogeneous input methods. Data input sources are expanding from human-only to include non-traditional methods like the millions of sensors and devices of the Internet of Things (IoT). Systems require increased computational flexibility and demand to deal with these expanding data sources. Data that is collected then needs to be analyzed and brought into the decision making process. Big Data is a critical strategy/application to deal with all of this information. Seventy percent of LOB and IT leaders either are implementing or already have Big Data implementations, while spending on analytics is up by 60%.[1]

Cloud Computing

After analysis, data then need to be delivered in an actionable format for users, and cloud is the new delivery model of choice. This year businesses are expected to spend $13 billion (USD) on cloud computing[2], and 80% of LOB and IT leaders are building new apps for the cloud[3]. This relatively recent shift to cloud has companies scrambling to understand how to scale their resources. Few know exactly how quickly the applications will grow for their private or hybrid clouds or how far systems need to scale out or scale up. For instance, after just eight years, Facebook had already accumulated over 1 billion users.[4] While not every application is the next Facebook, consumers and employees are becoming far more comfortable with cloud technology, and they are quickly adopting and using these tools as they become available. As companies try to rationalize how to deploy private/hybrid cloud technology, they need to consider the impacts of rapidly scaling within their infrastructure as well as how to contend with public/hybrid cloud sources. A more scalable, elastic, and dense solution needs to be available to keep pace with the cloud applications trend and the thirst for data. But like every new technology trend, these are running headlong into the realities of budgeting.

IT Economics

Companies are being pushed to be more agile and move faster. This push is driving demand for IT transformation, and innovation is seen as the best way to bridge current infrastructure with future needs. But innovation is rarely affordable. IT budgets have been strained for years with little relief in sight. Spending on datacenter systems had shrunk in 2013 and is set for only a modest gain in 2014.[5] However, more of the IT spending growth will be allocated to enterprise software, reflecting the changes that we see in Big Data, analytics, and cloud. While the modularity of blades solved some issues for customers, blades created new economic challenges as companies needed to manage a wider array of platforms. Most blade platforms had limited capabilities that

---

[1] Data from IBM Annual Survey of IT and line-of-business leaders for 2013

[2] http://www.forbes.com/sites/tjmccue/2014/01/29/cloud-computing-united-states-businesses-will-spend-13-billion-on-it/

[3] Data from IBM Annual Survey of IT and line-of-business leaders for 2013

[4] Chart http://money.cnn.com/2012/10/04/technology/facebook-billion-users/

[5] http://www.techrepublic.com/blog/european-technology/what-do-companies-really-spend-their-it-budgets-on-the-answer-may-surprise-you/#

strained under the changing loads of a dynamic set of applications. A blade solution that provided scalability and also helped drive down operating costs was needed.

Picking the Right Tool for the Job

For the past twenty years, industry-standard x86 servers allowed companies to standardize their computing work onto discrete, low-cost "building blocks". Blades helped bring density, but customers needed to trade off low cost to get that density. With the introduction of virtualization and cloud-based technologies, the limits of these blade systems were being pushed. "One server, one workload" was no longer the norm, and systems needed the ability to scale more elastically. Today, creating another VM or another cloud instance is easier than deploying physical resources. This leads to a new issue: "VM sprawl". The management challenge shifted from physical system management to virtual system management. What was traditionally an under-utilized system suddenly became over-utilized based on growing demands for memory, CPU, storage, and network. And these demands drove up operational costs from power and administration. Last year, it was estimated that 79% of all datacenter budgets were spent on management, administration, power, and cooling.[6] Another study found that 65% of the budgets were spent on ongoing operations and maintenance instead of new strategic initiatives.[7] How does a company prepare for the application changes of the future when it is overspending on the past and its platform choices are neither flexible nor adaptable?

In an environment with stable growth and simple applications, choosing a blade platform was relatively straightforward. But in an environment where inputs and outputs are evolving/changing rapidly and applications may need to scale up as well as out, there is a need for a more flexible "building block". It should provide the ability to scale applications within the platform and also have the resources/technology to scale up at the same time. Traditional blades just can't do this. Few flexible scalable architectures existed, which is why a modular approach like the IBM Flex System X6 compute node brings an interesting and needed option to the market.

True Modular Systems

To handle the emerging needs of businesses that may require an architecture to scale both within the platform as well as upward, IBM developed the IBM Flex System X6, the sixth generation of the IBM EXA architecture.

Designed to bring the right levels of performance and availability for scalable applications, these systems help enterprises take a more agile approach when addressing business-critical workloads including:

[6] IDC Market Analysis Perspective: Worldwide Datacenter Trends and Strategies 2013, Michelle Bailey and Katherine Broderick

[7] IBM Market Insights Study – 2013 Business Benchmarking Time-To-Value Study

- **Virtualization** where large contiguous pools of VMs, high VM density, or robust VM configurations are essential
- **Database** including large in-memory databases like SAP HANA
- **Enterprise planning and management** applications (ERP, CRM)
- **Analytics and decision support**
- **Cloud infrastructure** for private or hybrid models where there is a large concentration of VMs or cloud instances

The common denominator for these types of applications/environments is that they have high CPU/memory requirements or that they can gain significant cost advantages by consolidating/pooling resources onto a single larger, more robust platform.

In these cases where a higher concentration of VMs exists, a business can extract even better utilization of its computing assets with a larger system footprint (as in 4P or 8P platforms). Virtualization allows for better utilization by pooling VMs, thus reducing the overall amount of headroom reserved for workload spikes. Greater density of VMs on a larger host helps minimize power consumption for better operational efficiency.

The IBM Flex System X6 compute nodes are housed in a 10U chassis that accommodates 10 compute nodes along with storage, management, and networking. It is a completely self-contained system for handling a solution that may span several servers or for allowing servers to be joined together to scale up as needed.

Table 1: IBM Flex System X6

			
	x280 X6	**x480 X6**	**x880 X6**
Usage	Large virtual machines requiring more memory per VM	Large virtual machines, analytics, and databases requiring more cores as well more memory per system or VM	The highest levels of performance such as analytics and databases requiring more cores as well more memory per system
CPU Scalability	Scalable to 2 sockets	Scalable to 4 sockets	Scalable to 8 sockets
Memory	3TB memory (48 DIMMs)	6TB memory (96 DIMMs)	12TB memory (192 DIMMs)
Networking	Up to 4 networking adapters	Up to 8 networking adapters	Up to 16 networking adapters
Storage	2 x 2.5” HDD / SSD or 8 x 1.8” SSD with RAID 5	4 x 2.5” HDD / SSD or 16 x 1.8” SSD with RAID 5	8 x 2.5” HDD / SSD or 32 x 1.8” SSD with RAID 5

Because of their criticality to the business, platforms running scale up applications like these also need to be highly available. The IBM Flex System X6 compute nodes deliver resilience through not only the latest Intel platform technologies but also through IBM-specific availability features that build on IBM’s record of dependability.

Performance

Scalable applications demand performance, and the IBM Flex System X6 compute nodes deliver. Compared to the last generation (IBM BladeCenter HX5), the new X6 family doubles the socket count and delivers up to 300% more CPU cores. More than eleven times the memory, four times the networking capability, and a huge increase in internal flash storage all increase data handling capability. Table 2 highlights these performance improvements.

Table 2: IBM Flex System X6 *vs*. IBM BladeCenter HX5

	HX5	X6 (880)	Increase[8]
CPUs	4	8	100%
Cores	40	120	300%
Memory	1TB	12TB[9]	1100%
Network	4x 10Gb	16x 10Gb	400%
SSD	1.6TB	30TB	1875%

Flex System X6 compute nodes help drive the I/O requirements of scalable workloads. Key technologies to accelerate and protect I/O are shown in Table 3.

Table 3: IBM Flex System X6 Memory

IBM eXFlash Memory Channel Storage	Flash Cache Storage Accelerator	WriteNow and RAID 1 Memory
Ultra-low latency flash storage directly on the processor/memory bus for more I/O per second and lower latency	Transforms IBM Flash storage into a transparent acceleration device to cache hot data	Commits data earlier, reducing latency by up to 3X, and protects the data in the cache

Handling larger amounts of data stresses a server's I/O subsystems. Using flash memory ahead of drive storage can help reduce wait times for writing data and accelerate overall application performance. Introduced earlier this year with the initial launch of the X6 family, IBM eXFlash Memory Channel Storage will find support in Flex System X6 compute nodes later this year. With low write latency (as low as 3.3 µs[10]), IBM eXFlash Memory Channel Storage is integrated directly into the CPU/memory bus, creating an ultra-high-speed cache location. Block data can be written to the cache before it moves through the server to the storage subsystem. Just as the cache on a RAID controller creates a pooled area ahead of the rotating media to stage data for writes, this memory channel storage creates an even faster write mechanism ahead of the PCIe bus. This allows applications to commit writes to a low-latency device on the CPU/memory bus for significantly higher I/O throughput. Applications that rely on high I/O throughput like high frequency trading, Big Data head nodes, real-time analytics, or high volume transaction processing will benefit greatly from accelerated writes. Accelerating writes also increases reliability by shortening the amount of time that the data is most vulnerable: while it is "in transit" between the CPU and storage media.

[8] Configurations are scalable. All values are maximum supported.
[9] Assumes 64GB DIMMs. Available on Flex System x880 X6 later in 2014.
[10] Performance measured by IBM under lab conditions, actual customer results may vary.

The amount of memory and its speed are both critical to scalable applications. On most servers, loading 3 DIMMs per channel reduces memory speed by about 25% to help maintain the memory's signaling integrity. The IBM Flex System X6 is able to maintain the full speed of DIMMs—even when the channels are fully loaded—which allows memory-intensive applications to achieve better performance.

IBM Flex System X6 compute nodes are built on the same high performance architecture as the IBM x3850 server, which has delivered #1 results in transaction processing (4 CPU)[11], virtualization[12], and the #1 SAP SD 2-Tier[13] performance.

Agility

To excel in this rapidly-changing world, businesses need to be agile: their infrastructure needs to enable, not hold them back. A key to agility is the ability to scale up or reconfigure based on changing needs. Most blade systems are designed around a set of fairly rigid platform constraints focused on density, so scalability is not a primary design goal. The IBM Flex System X6 is different in how it addresses upward scalability.

Figure 2: Scalability

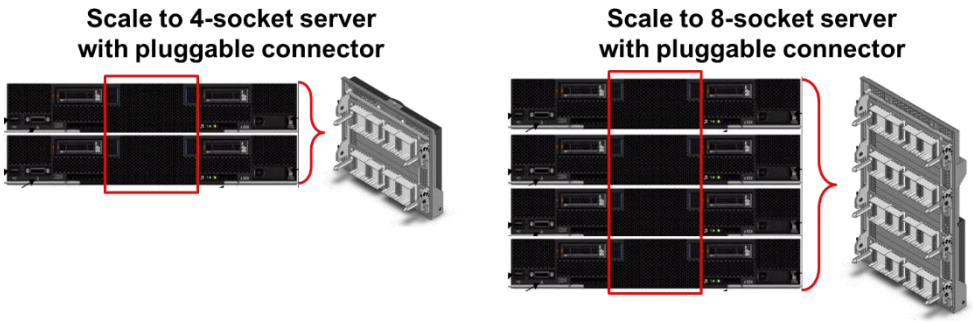

Typical blade solutions require customers to choose their platform up front; if needs change, the investment is lost because there is no upward expandability. When companies are uncertain about future workload scaling, they tend to take the conservative route by buying a platform for today's needs with the expectation that in the future they may need to do a "forklift upgrade". This means there is a risk of removing the existing server and losing that initial investment by purchasing a completely new server. With a modular scalable system, businesses can purchase what they need today and then grow into a larger footprint as their workloads increase.

Scaling presents a challenge in determining how to "right size" IT investments. Over- and under-provisioned systems are concerns. Modular systems can help alleviate this issue by providing a platform that can be repurposed easily when needs change or

[11] http://www.tpc.org/tpce/results/tpce_result_detail.asp?id=114021601
[12] http://www.spec.org/virt_sc2013/results/res2014q1/virt_sc2013-20140211-00010-perf.html
[13] http://download.sap.com/download.epd

grow. IBM's FlexNode partitioning capability allows platforms to be split: an 8-socket system can be configured as two 4-socket systems or four 2-socket systems, or a 4-socket system can be configured as two 2-socket systems. This modularity drives agility and investment protection by allowing a business to partition capacity for today's needs and then reconfigure as needs grow or change. If the expectation is that an application may grow to need 8P but today only needs 4P, an 8P system can be purchased with half (4P) allocated to the scalable application and the other 4P allocated as a VM pool. When the capacity is needed in the future, the VMs can be migrated to another host, allowing the server nodes to be reconfigured as a single 8P system with the application running on all 8 processors. This could never be done with traditional blades.

For agile businesses, management tools are essential to help administrators handle server tasks. An application like IBM Flex System Manager (FSM) provides the tools needed to manage a scalable system, comprehending both the physical and the virtual resources on the system. With its ability to provision servers and resources easily, new systems can be set up and configured with only a few clicks. It cuts down the amount of time required to deploy servers by up to 66%.[14] Through the IBM Upwards Integration Modules (UIMs), Flex System Manager can tie directly into VMware vCenter or Microsoft System Center, thus enabling management integration from the hypervisor all the way down to the bare metal.

Another key to agility is being able to deploy system resources quickly to support changes in the business or to pursue opportunities quickly. Fast setup and configuration patterns enable automatic provisioning of new Flex System X6 servers being deployed. Businesses can put resources into the hands of their end users faster for the most rapid, agile business applications.

Resiliency

Critical enterprise applications and large pools of virtual machines on a single platform all demand resilient and bulletproof operations. IBM begins with the resiliency features built into the latest Intel Xeon E7 processor and chipset, as expected, but adds additional resiliency features to help ensure maximum uptime.

[14] Internal IBM R&D estimates and measurements. Management server setup for IBM Flex System is 131 minutes. Equivalent setup for BladeCenter is 388 minutes. http://www.ibm.com/pureflex

Figure 3: Maximizing Uptime

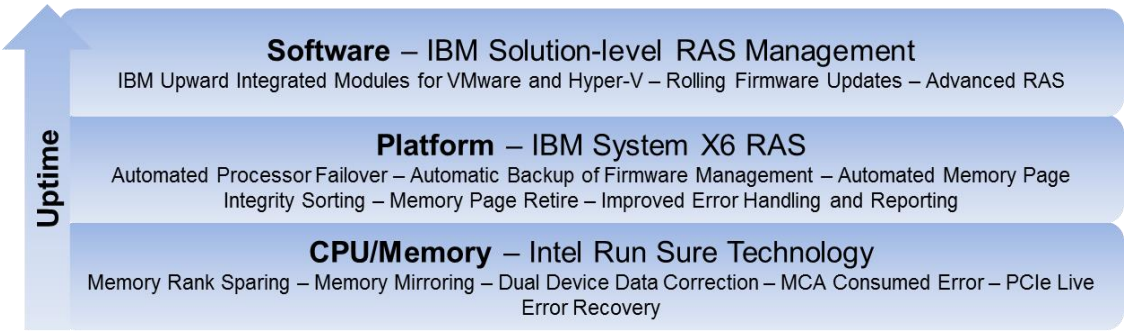

Hardware failure results in 25% of all system downtime[15], so keeping the server running is a key starting point. The standard Intel Run Sure Technology helps bring a high level of reliability and availability by protecting the CPU and memory from errors that would normally cause a system stop/restart.

Then on top of the hardware RAS, IBM adds technology that delivers even greater uptime with features like automated processor failover and advanced recovery of processor transactions. They help ensure that applications are not interrupted and data are not at risk from CPU events. For memory protection, automated memory page integrity sorting and memory page retire help ensure that memory reads and writes are done in the most reliable locations. As well, automated backup of management firmware helps reduce downtime by keeping a backup of the controller firmware to ensure the system can recover from any failure as a result of problems during a firmware update.

But hardware is only the start of the equation. The other 75% of the causes of downtime can be tracked to external factors, human error, and often software.

Enterprises need to focus on uptime with scalable systems: understanding the interaction between hardware and software can help reduce incidents that could result in downtime. At the software level, IBM continues to add value with the Upward Integration Modules for VMware and Hyper-V, the two most popular server hypervisors. These integration modules are designed to give an administrator more control over the server and software stack. With a fully integrated management stack, administrators can use VMware vCenter software or Microsoft System Center software for all of their daily management tasks. If an error occurs, the administrator can view it not only at the hypervisor level but also drill down into the hardware level to address the issue quickly from a single management console. Additionally, rolling firmware updates allow IT to update the servers in a virtualized cluster without impacting application availability. When systems within a cluster need to be updated, the feature will automatically move VMs, update the host, reboot the system, and then move the VMs back. Every server in the cluster is updated without having to take the entire cluster offline and without interrupting the operations of any of the VMs running on the cluster.

[15] http://venturebeat.com/2012/11/14/the-high-cost-of-server-downtime-infographic/

Together, each of these features adds layer upon layer of availability to the IBM Flex System X6 to help ensure the highest availability for critical applications.

Clear Business Benefits

For today's data centers, total cost of ownership (TCO) is a far larger consideration. To drive the best economic benefits to the business, scalable platforms need to address all three aspects of total cost of ownership: acquisition costs, operational costs, and downtime costs.

Figure 4: Reducing TCO

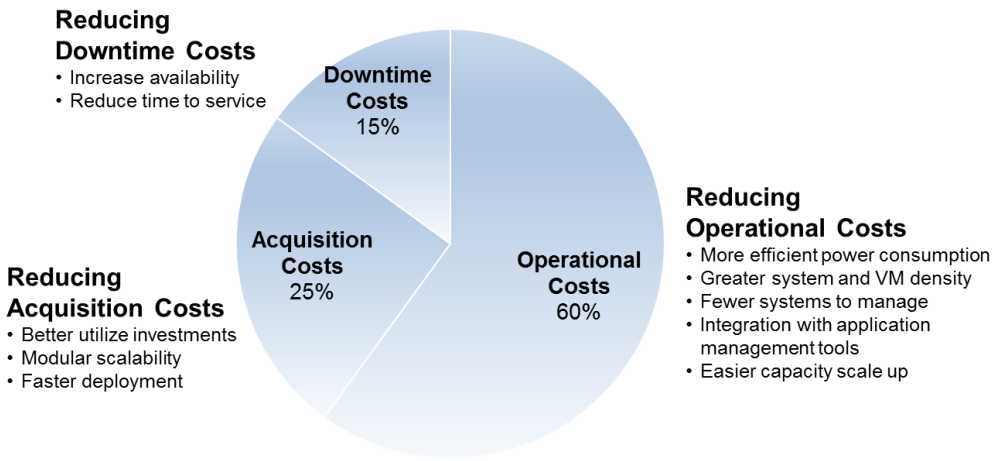

Reducing Acquisition Costs

Upgradable modular nodes allow customers to upgrade capacity easily as workloads scale quickly. This means not having to break the budget by either buying the extra capacity today (that may or may not be needed in the future) or buying the right product for today's needs only to have to replace it in short time if the application's needs explode (which is all too common for Big Data, analytics, and cloud applications). Standardization across 2P through 8P systems helps drive better efficiencies in ordering and deploying systems, helping to hold down acquisition costs.

Reducing Operational Costs

IBM Flex System X6 helps reduce the three largest drivers of operational cost: management, power, and physical space. Businesses can deploy 2, 4, and 8 CPU compute nodes using the same modular building blocks. Compared to standard blades, this allows three systems to be managed with the same image instead of only two systems managed with two separate images, reducing management by 50%

The IBM Flex System X6 Compute Nodes can be deployed with the same modular building block. This consistency enables companies to run multiple applications—and deploy the full spectrum of mission-critical workloads—on a single type of system. The Flex System management architecture allows administrators to manage all elements of

compute, storage, and networking in a single system. It greatly reduces the system administration overhead associated with managing multiple platforms.

The higher VM density of a larger footprint Flex System X6 means businesses can save power by replacing a large number of small 2P VM host nodes with a single 4P or 8P node to achieve lower power per VM. For instance, two 2P nodes with 86 VMs each (172 VMs total) consume 443W per server or 5.15W per VM. One 4 CPU IBM Flex System X6 480 can handle 368 VMs with the same configuration and consume 1390W total or only 3.8W per VM.[16] The effective wattage per VM decreases 30% when comparing two 2P systems to one 4P Flex System x480 X6. This represents a long term savings that continues to pay dividends especially as the price of power rises.

Finally, with datacenter floor space being the most expensive real estate in any company (due to special power, fire control, security, and cooling), IBM helps drive better efficiency. An IBM Flex System X6 system consumes only 10U of the typical rack, yet it holds the equivalent compute power as 14 rack servers with their associate layer 3 networking, aggregating switching, SAN switching, and SAN storage. In only 10U, IBM can pack the equivalent of a 42U rack of servers, storage, and networking. With integration of compute, network, storage, and management all in one chassis, businesses can optimize their datacenter space and reduce operational costs.

Reducing Downtime Costs

As discussed earlier, there are features in these systems that help reduce both planned and unplanned downtime. Features like rolling firmware updates help ensure that planned downtime is kept to a minimum, and features like memory and CPU protection help reduce unplanned downtime. The modularity of Flex System X6 compute nodes also reduces downtime for servicing and upgrades.

Conclusion

We believe that blade infrastructures had generally run their course and growth had flattened out. The key reasons for this flattening revolved around the limited scaling flexibility of each server along with the specialized, rigid environment where these systems lived. Blades were an excellent solution for problems that looked like a blade, but as soon as requirements wavered at all, blades were counted out.

Not having a solid strategy to handle the pressing needs of the future will lead to a situation where the enterprise is not provisioned to handle the rapid increases in scalability that today's applications require or the increase in user demand for data.

With the IBM Flex System X6, customers can deploy a modular system that brings the density of blades but—unlike blades—allows for far greater scalability at the node level along with scaling up as applications need even greater capacity and resources. Whether it is a single image footprint that demands huge memory or CPU addressability

[16] Source: IBM internal testing using Power Configurator tool and Virtualization Guide

or a large pool for handling a high density of VMs or cloud instances, the flexible scaling solution definitely has a place in transitioning datacenters.

Enterprises need to understand where modular platforms fit in their overall compute architecture and which applications could benefit from them. We recommend that enterprises examine their IT plans to ensure that they are properly provisioned for these changing workload requirements. Any organization faced with the challenge of rapidly scaling their workloads needs to consider the IBM Flex System X6.

For more information: http://www.IBM.com/flex.

Important Information About This Paper

Author
John Fruehe, 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
IBM is a research client of Moor Insights & Strategy and this paper was commissioned by IBM. No employees at the firm hold any equity positions with any companies cited in this documented.

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.