



ABOUT THE AUTHOR

George Teixeira is the CEO, President and Co-founder of DataCore Software Corporation. Mr. Teixeira is an expert in the field of storage virtualization, parallel I/O and real time systems. He is responsible for the creation and execution of the overall strategic direction and vision for DataCore Software.

To learn more visit:

www.datacore.com

Why Parallel I/O Software and Moore's Law Enable Virtualization and Software-Defined Data Centers to Achieve their Potential

By DataCore CEO & President George Teixeira

In the software-defined data center (SDDC), all elements of the infrastructure, including networking, compute, servers and storage, are virtualized and delivered as a service. Virtualization both at the server and storage level are therefore critical components on the journey to a SDDC since they shield users from underlying hardware complexity and enable greater productivity through software automation and agility.

Today's demanding applications, especially within virtualized environments, require high performance from storage to keep up with the rate of data acquisition and unpredictable demands of enterprise workloads. In a world that requires near instant response times and increasingly faster access to data, the needs of business-critical tier 1 enterprise applications, such as databases including SQL, Oracle and SAP, have been largely unmet.

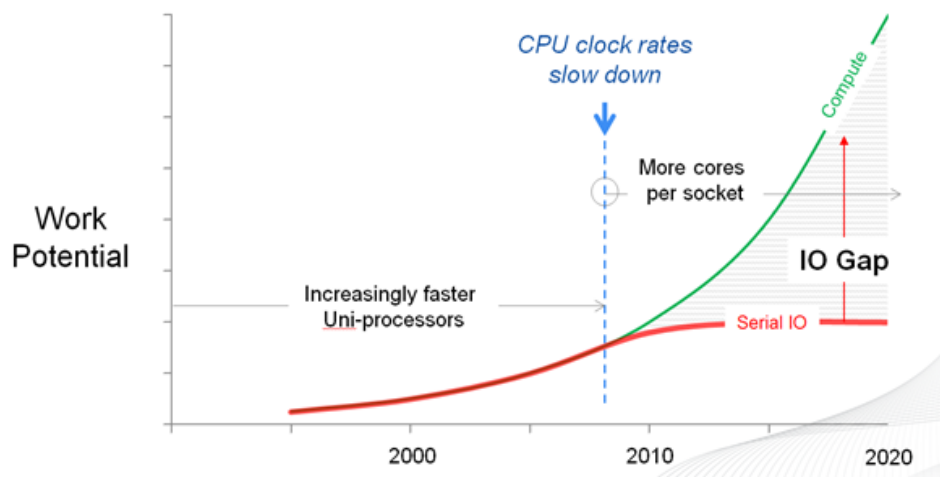
Overcoming the I/O Bottleneck in a Virtual World

The major bottleneck holding back the industry is I/O performance. This is because current systems still rely on device -level optimizations tied to specific disk and flash technologies since they don't have software optimizations that can fully harness the latest advances in more powerful server system technologies such as multicore architectures. Therefore, they have not been able to keep up with the pace of Moore's Law.

"By taking full advantage of the processing power offered by multicore servers, parallel I/O software is the key enabler for a true software-defined data center. This is due to the fact that it avoids any special hardwiring that impedes achieving the benefits of virtualization while it unlocks the underlying hardware power to achieve a dramatic acceleration in I/O and storage performance"

- George Teixeira,
President & CEO
DataCore Software

IO Starved Virtualization Servers



Many have tried to address the performance problem at the device level by adding solid-state storage (flash) to meet the increasing demands of enterprise applications or by hard-wiring these fast devices to virtual machines (VMs) in hyper-converged systems. However, improving the performance of the storage media—which replacing spinning disks with flash attempts to do—only addresses one aspect of the I/O stack: read performance. Hard-wiring flash to VMs, a la Tintri and VMware EVO:RAIL seems to be a contradiction that goes against the movement of virtualization, in which technology is elevated to a software-defined level above the hard-wired and physical aware level, and it also adds complexity and vendor specific lock-ins between the hypervisor and device levels. That's no way to move forward on virtualization and the transformation to a true software-defined data center!

Another approach to improving I/O performance is through I/O parallelization that builds on virtualization's ability to decouple software advances from hardware innovations. This method takes into consideration the industry shift from single to multicore processors as the norm in modern server technologies by using software to drive parallel I/O across all of the cores. Today's commodity servers already cost-effectively provide a wealth of powerful multiprocessor capabilities, but most of it sits parked and in idle mode, unexploited. By taking full advantage of the processing power offered by multicore servers, parallel I/O software is the key enabler for a true software-defined data center. This is due to the fact that it avoids any special hardwiring that impedes achieving the benefits of virtualization while it unlocks the underlying hardware power to achieve a dramatic acceleration in I/O and storage performance – allowing the force of Moore's Law to solve the I/O bottleneck problem holding back the IT industry and the realization of Software-Defined Data Centers.

Why Parallel I/O?

Let's begin with the introduction of parallel processing technology. More than three decades ago, computer scientists were working diligently to increase performance. During this time, which pre-dates the PC revolution, the goal of developers was to create superfast computing platforms by operating multiple processors in parallel. This primarily involved optimizing system architectures by using the power of multiple microprocessors to achieve faster computations and I/O performance in the form of parallel computing. Essentially, workloads could be spread across multiple processors in order to be executed faster – making them ideal platforms for “parallelization.”

This was eventually supplanted by PCs running on a network, and parallel architectures were replaced by frequency scaling (driving faster and faster clock frequency to achieve performance gains). The thing that ultimately frustrated parallel computing was the pace of clock speed increases for uniprocessors which was a byproduct of Moore's Law, the observation made by Intel co-founder Gordon Moore that the number of transistors per square inch on integrated circuits would double approximately every 2 years, and the prediction that this would continue for the foreseeable future.

As a result, instead of using more expensive and complex multiprocessors, uniprocessors advanced quickly and systems were able to accomplish most processing tasks by sharing a common fast processor, the CPU. Workloads were executed on these chips sequentially rather than in parallel, and the industry shifted its focus to amplifying clock speeds instead of extracting increased performance from multiprocessor designs.

However, with the advancement of Moore's Law over the years came an unseen side effect. The “tick tock” cycle of innovation hit a wall when it encountered overwhelming power and heat issues and increasing clock speeds fundamentally came to a halt – and, in fact, started going backwards a few years ago. Nevertheless, transistor densities still continued to hold to Moore's law and doubled every 18 to 24 months with the advances coming from more cores being available to do the work versus faster clock speed on single CPUs, and this became the new norm.

It became commercially more profitable to take two processors and put them on the same die (rather than use that budget to advance the silicon production methodology to increase the performance of a single processor). For the first time, without it being widely observed, Moore's Law (or our perception of it) had fundamentally changed. Computers have now been completely revolutionized -- evolving from one to many core systems with 10s to 100s of cores available to do work – and without a lot of hype, multicore computing has become the dominant form of computing today.

“Parallel I/O software will be the next ‘killer app’ that will allow the industry to build on virtualization and fully utilize the multiple cores that are readily available, and allow for greater productivity and further consolidation of server technology by allowing more applications, VMs and workloads to run on these platforms without being bottlenecked by I/O. Harnessing today's powerful multicore servers to do parallel I/O is the way to unleash the full power of Moore's Law and make it possible through virtualization to increase business productivity --and therefore drive the promise of Software Defined Data Centers .”

- George Teixeira,
President & CEO
DataCore Software

ABOUT DATACORE

DataCore is a leader in software-defined storage. The company's storage virtualization and Hyperconverged Virtual SAN solutions empower organizations to seamlessly manage and scale their data storage architectures, delivering massive performance gains at a fraction of the cost of solutions offered by legacy storage hardware vendors. Backed by 10,000 customer sites around the world, DataCore's adaptive and self-learning and healing technology takes the pain out of manual processes and helps deliver on the promise of the new software defined data center through its hardware agnostic architecture.

www.datacore.com

Virtualization + Parallel I/O = the 'Killer Apps' Driving Data Center Business Productivity

The virtual server revolution largely driven by VMware and Microsoft hypervisors became the "killer app" that could exploit a greater degree of the new multicore capabilities. As timesharing went from providing terminals using time slices of large uniprocessors to everyone having a PC, the logical next step was to go parallel and put multiple PCs (VMs or virtual machines) to work and use multiple processors on these new powerful servers to drive even more VMs. This drove down multicore costs via volume manufacturing, and became the driver of the virtualization revolution where one could consolidate servers and administration for cost savings, while increasing productivity through the utilization of more powerful multicore platforms. In a marketing phrase – it became "do more with less."

The downside is that virtualization and server consolidation creates a workload blender effect in which more and more of the application I/O workloads are concentrated on the same system and must be dealt with. All of those VMs and their applications become easily bottlenecked going through a common "I/O straw." As processors and memory have dramatically increased in speed, this I/O straw continues to bottleneck performance -- especially when it comes to the critical business applications driving databases and on-line transaction workloads.

The primary element that is missing in this equation is software that can take advantage of the powerful multicore/parallel processing infrastructure. What is needed is a Software-Defined Storage stack with parallel I/O capabilities to remove the roadblocks holding back application performance. This capability will be a game-changer for the industry, similar to the way in which VMware was able to better utilize computing and consolidate servers by running many VMs on a single system.

Parallel I/O software will be the next 'killer app' that will allow the industry to build on virtualization and fully utilize the multiple cores that are readily available, and allow for greater productivity and further consolidation of server technology by allowing more applications, VMs and workloads to run on these platforms without being bottlenecked by I/O. Harnessing today's powerful multicore servers to do parallel I/O is the way to unleash the full power of Moore's Law and make it possible through virtualization to increase business productivity --and therefore drive the promise of Software Defined Data Centers .