

# Server-based Flash Poised to Change the Data Center; Interview with Fusion-io CMO Rick White Part I

This past Monday EMC created a fair amount of buzz in the storage industry with its VFCache [announcement](#) that in essence validates the emergence of server-based flash technology in the enterprise. But does EMC VFCache go far enough? Fusion-io, who arguably invented this space, argues, “**Definitely not!**” In this first of a multi-part interview series with Fusion-io’s Chief Marketing Officer, [Rick White](#), we talk about server-based flash technology and why it is poised to change enterprise data centers.

**Jerome:** *Rick, thanks for joining me today. It has been awhile since we last spoke and, when we did, it was only briefly when we were at [VMworld](#) shortly after [Fusion-io](#) had gone public. So tell me what it has been like at Fusion-io since you have gone public?*

**Rick:** Jerome, great to speak with you again as well. Interesting thing about going public is that you raise your visibility. The frustrating thing is fame is an ugly thing. It is a double-edged sword. It is like being the [Jonas Brothers](#). I can fill concerts but everyone hates me. We seem to be the company everyone loves to hate at this point.

**Jerome:** *So why are you the company “everyone loves to hate” as you put it? What have you done that is so disruptive?*

**Rick:** When we came up with this concept about five years we were looking at things like Amdahl’s law and how to get I/O into the processor. It is the whole reason CPUs have [L1 cache](#).

It is really expensive but it is as close to the processor as you get. [L2 cache](#) is not as expensive but it is not as close either. Then computers go to [L3 cache](#) and then to RAM.

But then computers go directly from RAM straight to disk. No one has used disk as a scratch base for memory for two decades. The last time I did that was 1982 or 1984. It just gets too slow.

**This is what Fusion-io offers: a new memory tier, one that is thousands of times faster than disk.** This puts us back to where disk was twenty (20) years ago and is persistent.

**Flash is a perfect memory tier** as it is lower cost than RAM and it will always be ahead of RAM in terms of density. So we now have this great density and power profile that is better than RAM and offers performance that is way better than disk. This gives us a great memory tier.

Ironically, when Fusion-io got started, it went to fabricators like [Micron](#) to convince them to work with us. They said, “*We are nuts.*” Unfortunately without the support of fabricators like Micron and without motherboards that have sockets for our concept around a [DIMM](#), we had to take our dual-inline memory module, build it ourselves and then put it on a [PCI Express](#) carrier card. Then we came out with one that could hold two of our DIMMs. Now we have one that holds eight of these DIMMs.

We do believe at some point that we will see DIMMs down on the motherboard. But what is interesting is what was once this weird niche is suddenly the hottest thing around. We basically invented the sector and I remember all of the grief we got for it. And now others are saying, “*Alright, this is the place to go.*” and they use Fusion-io as their measuring stick and they love to hate us. It is an interesting place to be – post-public and having the visibility.

**Jerome:** *So what exactly are you changing that has everyone so*

*up in arms and why is EMC having to respond?*

**Rick:** We are not going to go out and change our lives. ***This is about changing the data center.*** This is about a chance to be a part of the history of technology. We believe this new technology, this new storage medium, ***this new memory tier is going to be very, very important going forward.***

Having this new high-speed memory tier is not going to be enough. This is a lot like the x86 processor. Suddenly you have this cheap, commodity processor and you can build computers and eventually servers out of this architecture.

**Jerome:** *You say this is a lot like the introduction of the x86 processor. Can you elaborate?*

**Rick:** Mainframe guys said back in the day, “*Ha, ha, ha, isn’t this cute! Ha, ha, ha, what a toy! The X86 is a toy.*” Now we look back at when megaflops meant something on the mainframe and who is laughing now?

The graphics workstation industry said the same thing. Silicon Graphics’ position was, “*This is our [Indigo 2](#) \$48,000 workstation. Here is how many polygons per second we can do.*” They were replaced by a workstation with an [Nvidia](#) graphics card running Windows NT.

This reason why this happened to both the mainframe and graphic workstations are almost exactly the same. The mainframe was displaced because processing was decentralized from the mainframe. Suddenly I could put it out wherever. I could run processing at someone’s desk and run an application at their desk.

Using a mainframe everyone had to log into a terminal and, because it was so expensive, everyone had to share it by doing time slicing and run batch jobs. They shared a centralized processing unit. x86s decentralized processing allowing

processes to move throughout the business.

In the case of graphics, graphics performance was decoupled from a proprietary box and put onto a card that could go into any machine.

So as Fusion-io looked at what it was doing, we said, *“We are fundamentally decoupling I/O performance from the SAN.”*

*In [Part II](#) of this interview series with Fusion-io’s CMO Rick White, we will discuss why this decoupling of I/O performance from storage is necessary and why this creates a new tier of memory as opposed to a new tier of storage.*

In [Part III](#) of this series, Rick explains the new Fusion-io Octal drive, what makes it different from Fusion-io’s earlier ioDrives and how Fusion-io is going to market with it.

*In [part IV](#) of this interview series, Rick and I discuss why Fusion-io is opening up its virtual storage library (VSL) APIs to developers.*

*In the final [Part V](#) of our interview series, Rick provides Fusion-io’s takes EMC’s Project Lightning (now known as VFCache) and the gap that persists between SSD providers and Fusion-io’s ioMemory.*