

Whitepapers - Free to Download

Value-driven quality management for complex systems

SC10 Student Cluster Competition

Original URL: http://www.theregister.co.uk/2009/09/25/six_flash_locations/

Balancing act for servers that flash the cache Lazy storage arrays making data-hungry servers wait

By [Chris Mellor](#)

Posted in [Storage](#), [25th September 2009 09:38 GMT](#)

[Sign up for The Reg enterprise storage newsletter](#)

Comment Multiple flash locations in the server-storage stack are upsetting balanced I/O conventions and making overall system design much more difficult.

Designing and implementing server-to-storage systems is going to become very much harder, because existing assumptions about server and storage I/Os per second (IOPS) handling are being swept aside. Virtualised, multi-core, cache-enhanced servers can digest and generate data faster than many hard drive storage arrays can deliver and receive it. How do you build balanced server-storage systems when this is happening?

Whitepapers - Free to Download

Does virtualization pay off for small business?

2010 Nemertes PilotHouse Awards

Is Your Network Ready For IP Telephony?

Is there a (Blue) Cheetah in your future?

Power and Cooling Capacity Management for Data Centers

Let's look at the basic server-to-external-storage stack. We start with a server motherboard containing four, six, even eight-core servers which talk via memory and a PCIe bus to an I/O adapter. This links across a network to a storage array controller, which is connected to hard disk drive shelves. Imagine these are virtualised servers and its easy to see, put crudely, a server engine room capable of handling thousands of I/Os per second (IOPS) talking to a

storage array only capable of handling hundreds of IOPS.

To get the storage array IOPS capacity up to the server's IOPS level will require a large increase in array spindles or solid state drive (SSD) use, and increased controller I/O capability. This can be done by using flash memory caching but, unfortunately, flash storage can also be used as a server cache, in various places on the server side of the stack, to increase the server's IOPS rate, thus unbalancing things again.

We can envisage server motherboards coming with flash modules on them, acting as a cache between the server engine and its data store. Intel's Braidwood technology is heading this way, it appears, and will increase a server's hunger for data. Then there could be a flash store connected to the server's PCIe bus: another cache. Fusion-io and Violin Memory both have such flash-based I/O acceleration cards, as does Oracle's Sun-powered Exadata server. This caching also increases the I/O burden on the connected storage.

Next there could be a flash cache on the I/O adapter used to connect to the storage. Adaptec is showing the way with its data conditioning ideas and an Intel-supplied flash cache addition to its 5000 and 2000 RAID controllers. We have three potential tiers of flash cache here all increasing the burden on the storage array.

Skipping across the network link to the storage array we could have a flash I/O accelerator card in the array controller. This is what NetApp's PAM (Performance Acceleration Module) is, admittedly in DRAM form initially, but now with a flash version announced, too. That makes four flash locations with the fifth being SSDs actually replacing hard disk drives and providing a so-called tier zero of storage. Virtually every storage array supplier is doing this, with the majority using STEC SSDs and a few settling on Intel.

Finally the sixth flash location is a replacement for the hard drive array itself. This is what Texas Memory Systems (RamSan), Sun (FlashFire), and also Violin Memory (1010 Memory Appliance with network head) are doing, and also what Fusion-io might do with its ioSAN. These can be viewed as flash data stores, not caches, coming in either direct-attached flash (flash DAS) or network-attached (a flash SAN) form.

It seems a nonsense to have flash in all six locations. We might say we could sensibly add a flash cache to servers and/or have either a flash store or a flash-enhanced storage array. The flash-enhanced storage array would have a cache for hot data or else incorporate a flash store to augment a high-capacity and cheap bulk data SATA hard drive store.

A system vendor - say Cisco (soon), Dell, HP, IBM or Oracle/Sun - would design systems to provide server and storage IOPS-handling and data capacity in the most balanced, scalable and cost-effective way. Oracle's Exadata 2 is an example of this. A server and storage system integrator and VAR will also need to do it, but not have the same resources as tier one systems providers, and probably do a clumsier job. A data centre operator buying servers and storage separately will also have a complicated balancing act to accomplish balanced IOPS handling across its servers and storage. Tools are going to be needed.

Its not clear where these tools will come from and it's likely that the need to accomplish a technically different balancing act may help drive a focus on integrated server and storage systems from suppliers and away from the idea of treating servers and storage as relatively separate purchases. This aspect of flash could prove to be quite disruptive. ®

Related stories

[Oracle promises fresh approach to storage](#) (20 September 2010)

http://www.theregister.co.uk/2010/09/20/exalogic_storage/

[Toshiba takes Violin Memory stake](#) (1 April 2010)

<http://www.theregister.co.uk/2010/04/01/toshiba-violin/>

[Adaptec's caching card uses MS code](#) (26 March 2010)

http://www.theregister.co.uk/2010/03/26/adaptec_maxiq_microsoft/

[LSI and Seagate take on Fusion-io with flash](#) (16 March 2010)

http://www.theregister.co.uk/2010/03/16/lsi_seagate_pcie_ssd/

[STEC becalmed as Fusion-io streaks ahead](#) (5 March 2010)

http://www.theregister.co.uk/2010/03/05/stec_fusion_io/

[Seagate expanding into PCIe flash](#) (27 January 2010)

http://www.theregister.co.uk/2010/01/27/seagate_lsi_pcie/

[Bendy flash memory raises prospect of flexible displays](#) (16 December 2009)

http://www.theregister.co.uk/2009/12/16/flexi_flash/

[Fusion-io whips out fast gov-grade ioDrive](#) (23 November 2009)

http://www.theregister.co.uk/2009/11/23/fusion_iodrive_octal/

[MySpace replaces disks with flash](#) (14 October 2009)

http://www.theregister.co.uk/2009/10/14/myspace_fusionio/

[Sun debuts FlashFire, calls record books](#) (12 October 2009)

http://www.theregister.co.uk/2009/10/12/sun_flashfire/

[IBM readies Exadata killer](#) (5 October 2009)

http://www.theregister.co.uk/2009/10/05/ibm_db2_pure_scale/

[STEC notches up LSI win](#) (28 September 2009)

http://www.theregister.co.uk/2009/09/28/stec_lsi/

[Dataram makes SAN flash cache sandwich](#) (28 September 2009)

<http://www.theregister.co.uk/2009/09/28/dataram/>

[IBM punts fixed-cost BlackBerry support plan](#) (25 September 2009)

http://www.theregister.co.uk/2009/09/25/ibm_fixedcost_blackberry_support_plan/

[SandForce gets SMART](#) (24 September 2009)

http://www.theregister.co.uk/2009/09/24/smart_sandforce/

[EMC puts future on FAST forward](#) (22 September 2009)

http://www.theregister.co.uk/2009/09/22/emc_fast/

[HP's full flash house](#) (18 September 2009)

http://www.channelregister.co.uk/2009/09/18/hp_samsung_flash/

[Violin Memory grabs Fusion-io's former chief](#) (15 September 2009)

http://www.theregister.co.uk/2009/09/15/violin_gets_basile/

[Adaptec bolts SSDs onto lightning RAIDs](#) (9 September 2009)

http://www.theregister.co.uk/2009/09/09/adaptec_maxiq/

[NetApp heading cloudwards](#) (25 August 2009)

http://www.theregister.co.uk/2009/08/25/netapp_ontap_8/

[STEC goes on range extension rampage](#) (12 August 2009)

http://www.theregister.co.uk/2009/08/12/stec_range_extension/

[Fusion-io flush with flash success after demo](#) (28 July 2009)

http://www.theregister.co.uk/2009/07/28/fusionio_tpch_result/

[WhipTail promises no bull 6TB SSD](#) (10 July 2009)

http://www.theregister.co.uk/2009/07/10/whiptail_ssd/

[Texas goes large with supersize SSD](#) (22 April 2009)

http://www.channelregister.co.uk/2009/04/22/tms_5tb_ssd/

[Super Talent flashes the big whopper](#) (3 April 2009)

http://www.theregister.co.uk/2009/04/03/super_talent RAIDdrive/

[The 4 terabyte Violin](#) (10 November 2008)

http://www.theregister.co.uk/2008/11/10/violin_4tb_flash_memory_appliance/

© Copyright 1998–2010