

STEC stamps its mark on PCIe flash

Meet Kronos and Kronos Turbo

By Chris Mellor • The Register

Posted in Storage, 5th August 2011 09:05 GMT

STEC has introduced Kronos and Kronos Turbo, their dynamic PCIe server flash duo. Actually STEC is launching triplets as there is a Kronos Bi-Turbo – two Turbos ganged together – and there is also flash caching software.

STEC has been absent from the fast-growing PCIe flash card space, one pioneered by Fusion-io in which servers use flash as an extension of memory to cache active data and get to it much faster than if they were to read it from hard drives. With disks, the server processor has to wait while the disk's read/write head goes to the right disk track and then wait again until the track's sector spins under the head. Flash can complete dozens of read I/Os while the disk drive is waiting for head and target track to meet.

STEC is introducing three Kronos PCIe flash cards, solid state accelerators called the Kronos, Kronos Turbo and Kronos Bi-Turbo PCIe SSA products. The Kronos PCIe SSA, a half-height X8 PCIe 2.0 card, holds up to 980GB of flash and this can be fast single level cell (SLC) or slower but cheaper multi-level cell (MLC) flash. The Kronos Turbo holds up to 1.96TB and this is a full height X8 PCIe 2.0 product featuring two clustered ASIC-based controllers. The Bi-Turbo holds up to 3.9TB and will be two full height cards.



STEC's Kronos PCIe flash product

The products have branded controller features. STEC says these are: "CellCare, Secure Array of Flash Elements (S.A.F.E.) and PowerSafe, [which] provide ... performance, endurance and data protection for ... enterprise and cloud computing applications."

- CellCare uses "signal processing and adaptive flash management algorithms" to extend NAND's endurance.
- SAFE provides the "ability to recover from NAND flash page, block, die and chip failures and maximises the Mean Time Between Failure (MTBF) and Mean Time To Data Loss (MTTDL)."

- PowerSafe is protection against power loss and preserves data in the device that hasn't yet been written to flash.

STEC is of the view that CellCare enables it to use 25 per cent less flash than it otherwise would to achieve the endurance rates it needs.

Performance

The base Kronos PCIe with SLC flash does up to 110,000 random read and 100,000 random write IOPS, and has a 1GB/sec sequential read and write bandwidth rating. The Kronos Turbo does 220,000 random read and 200,000 random write IOPS, and 2GB/sec sequential reads and writes. By extension the Kronos Bi-Turbo does 440,000 random read and 400,000 random write IOPS, and 4GB/sec sequential reads and writes.

The obvious competitive comparison is with Fusion-io's ioDrive which in 640GB MLC ioDrive Duo form does 196,000 random read IOPS using 512 byte blocks and has a bandwidth of 1.5GB/sec with 64KB blocks and 1GB/sec with sequential writes at that block size. As we can see direct random I/O comparisons are simply not possible; STEC is not furnishing us the block sizes involved in its numbers. But sequential I/O bandwidth figures should be comparable.

Here Fusion-io's 320GB SLC ioDrive Duo does 1.5GB/sec reads and writes and so is slower than the Kronos Turbo. Its 160GB SLC ioDrive does 770MB/sec sequential reads and writes, and is slower than the Kronos PCIe. We expect Fusion to announce a product refresh fairly soon and so cannot say how long-lived STEC's Kronos performance advantage over Fusion-io will be.

OCZ's recently announced [Z-Drive R4](#) claims up to 410,000 random read and write IOPS, and has a 2.8GB/sec bandwidth, positioned between the Kronos Turbo and Kronos Bi-Turbo. It can use SLC and MLC flash.

Server CPU offload

STEC says the Kronos has an ASIC on board and it does management processing, unlike the Fusion-io cards which, STEC says, rely on the host server to do this work and so bleed off CPU cycles that can't be used to run application software. However, this simple and neat story is weakened by STEC announcing caching software that runs in a host server and is optimised for working with STEC flash products, PCIe ones or storage array SSDs.

We note this also makes performance comparisons between Kronos and ioDrive products difficult, as STEC will say it delivers its IOPS and GB/sec without sucking up host CPU cycles whereas Fusion-io products suck up host CPU cycles. According to Stifel Nicolaus analyst Aaron Rakers: "STEC claims its solution will utilise about 37 per cent of the host CPU resources, while the Fusion-io solution would utilise up to 50 per cent of the CPU resources."

We need formal benchmarks here, such as SPC-1s, to work out the effects of such component differences.

STEC's marketing pitch is that customers with large data centres need fewer servers if using Kronos than if using ioDrive-based server acceleration.

Caching software

The EnhanceIO caching software resides in a host server underneath, as it were, the server's O/S and application software. It enables server apps to make better use of any flash in the server system, be it using the PCIe interface or being Fibre Channel, SATA or SAS interface SSDs.

The STEC claims it accelerates data response times and speeds read I/O and is supposed to be transparent to the server's O/S, with Windows, Linux and VMware supported, and transparent to the applications and underlying storage. The software gives the flash read, write back, and write-through cache functionality, and works with any vendor's flash products, although it is optimised for STEC NAND ones. Presumably the combination of EnhanceIO caching software and Kronos PCIe flash hardware is more beneficial to host servers and their applications than either alone.

STEC claims: "In a benchmark test of Oracle servers, STEC's EnhanceIO SSD Cache Software enabled 2.5 times faster throughput while serving 3.5 times more concurrent users. Through additional testing in a 10 Terabyte direct attached storage environment, EnhanceIO software provided between twice and ten times the performance (depending on the application) when compared to an 'all-HDD' solution, at about one-third the system cost."

It doesn't tell us the effect of adding EnhanceIO software caching to an existing STEC Kronos PCIe flash set-up though.

STEC's overall message here is: "EnhanceIO SSD Cache Software enables the best of both worlds – the full performance benefits of SSD technology coupled with the storage capacity benefits of HDDs."

The EnhanceIO product includes:

- a Performance Profiler to identify opportunities for server performance optimisation. It can enable customised solutions for specific server workloads.
- An SSD Cache Engine to drive the use of SSDs and manage all I/O activity according to policies established by the end-user enterprise.
- A GUI Manager: a browser-based tool used to set-up, manage and monitor the SSD cache.

STEC recently announced [poor results](#) with a pretty dire outlook and its share price was trashed by investors. If these PCIe flash drives are as good as claimed and can sustain their current competitive advantage – and if the EnhanceIO caching software does what STEC claims it does – then the share price might start its long march back to health.

STEC is trying to halt the share price fall by authorising the repurchase of up to \$15m of its common shares. It expects to complete these repurchases by the end of August 2011.

The Kronos products are warranted for five years. Kronos PCIe and Kronos Turbo PCIe are sampling now and Kronos Bi-Turbo should arrive in 2012, when STEC expects to start using 29-20nm process flash; it is currently using 39-30nm product. The EnhanceIO caching software is also sampling to customers now. ®