



Flash Solid State Disks Now Cost Less Than Enterprise Hard Disks ... and Are 7x to 20x Faster

Solid State Disks (SSDs) are ready for mainstream business servers. By combining new MLC-based SSDs with EasyCo's MFT Flash management software, you can now improve the performance of your servers by 15x and save money over enterprise hard disks

Managed Flash Technology™ (MFT)

MFT™ software changes the topology of random writes. Traditional random writes put data back where it was found. MFT writes data in clusters as large linear writes, remembering the new location of the data. MFT writes randomly to Flash SSDs up to a thousand times faster than un-managed Flash SSDs. MFT offers two advantages for Flash SSDs. First, that random data can be written very quickly. MFT can perform 4K random writes at 10,000 to 25,000 IOPS to a single drive and basically uses the entire linear write throughput of the drive. Equally important, MFT reduces the number of write/erase cycles consumed with random writes extending the drive life by 75x for most applications.

Cost and Performance Comparison

The following table compares a small RAID array with SSDs and HDDs. Each array is designed for best performance.

	15K 2.5" SAS HDD	MFT-100 SSD w/ MFT
# of drives	8	6
Drive Size	36 GB	32 GB
Raid Level	10	5
Usable Capacity	144 GB	144 GB
Cost incl. raid controller	\$3,600	\$2,624
Average read/write IOPS	1,666	24,408
Cost/GB	\$25.00	\$18.22
Cost/IOPS	\$2.16	\$0.11
Drive Power Usage	48 watts	9 watts

15k rpm 2.5" SAS Drives

The HDD is Seagate's latest 2.5" Savvio 15K HDD. This 2.5" SAS drive delivers about 250 random read/write IOPS in a 15 mm thick 2.5" form factor drawing about 6 watts of power. This is the fastest "random access" hard drive currently manufactured. This 36GB drive costs about \$350.

Mtron 1000 series Flash SSDs

This comparison uses the MFT-1025-32 MLC SSD. This 2.5" SATA drive requires no cooling, draws 1.5 watts of power, and costs \$399 for a 32GB drive. MFT-1000 drives random read at about 10,000 4K IOPS, linear read at 100mb/sec, and linear write at 40mb/sec. Like all Flash SSDs, the 1000 series has a

slow native random write performance: 24 IOPS. By itself, an 1000 series drive performs at about 1/5th the performance of a 15K SAS drive in most server settings. With MFT, 1000 series drive is about 15x faster, for random read/write operations, than a 2.5" Savvio 15K HDD.

1000 series SSDs are manufactured using 10,000 endurance MLC (Multi Level Cell) Flash memory. While this is a high quality chip, calculations show that by itself, MLC SSDs have an unsatisfactory life in a server environment, and can in small sizes (e.g. 16gb) fail in as little as 77 days of saturated random writing.

Wear Life Comparison with MFT

As a simple rule of thumb, an MFT enabled 1000 series drive will last at least five years provided that it is over-written on average less than one and a half times every day with random data. Thus this example 144GB MFT RAID-5 array can have almost a quarter terabyte of 4K random writes daily and still last five years. For most real-world servers, this represents a life span of dozens or even hundreds of years.

When comparing random writes with 15K SAS drives, the SAS drives are simply incapable of randomly writing at these speeds. If your heavy random write application "can run", even saturated, on 15K SAS drives, then MLC SSDs will have adequate endurance for at least a 5+ year lifespan, if not much longer.

Conclusion

MLC Flash SSDs with MFT can power small business servers at 15x the performance and for less money than a same-sized 15K 2.5" SAS array. This is without even considering the smaller physical footprint of Flash SSDs and lower power consumption.

What is the value of happier customers, more productive employees, and lower power consumption to you?

Easy Computing Company

220 Stanford Drive
Wallingford PA 19086

Toll Free: 888-473-7866
Phone: (+1) 610-237-2000
Fax: (+1) 610-672-9549
E-mail: sales@easyco.com
web: http://EasyCo.com

An Expanded Comparison Table and Notes Regarding Evaluation Methods

	7200 2.5" SATA HDD	15K 2.5" SAS HDD	15K 2.5" SAS HDD	MFT-500	MFT-1000	MFT-3000
Raid Level	RAID-5	RAID-5	RAID-10	RAID-5	RAID-5	RAID-5
Drive Size	80 GB	36 GB	36 GB	32 GB	32 GB	32 GB
# of drives	5	5	8	6	6	6
Usable Capacity	320 GB	144 GB	144 GB	144 GB	144 GB	144 GB
Drive Cost	\$425	\$1,750	\$2,800	\$1,734	\$2,394	\$4,308
RAID Controller Cost	\$230	\$800	\$800	\$230	\$230	\$230
Cost incl. raid controller	\$655	\$2,550	\$3,600	\$1,964	\$2,624	\$4,538
Average IOPS; 70/30 rw mix	180	601	1,666	11,160	24,406	32,573
Cost/GB	\$2.05	\$17.71	\$25.00	\$13.63	\$18.22	\$31.51
Cost per rw IOPS	\$3.64	\$4.24	\$2.16	\$0.18	\$0.11	\$0.14
Drive Overwrites/Day; 100%	0.15	1.15	3.16	0.75	1.50	15.00
Flash Memory Type				MLC	MLC	SLC
Drive Power Usage	20 watts	30 watts	48 watts	9 watts	9 watts	9 watts

Size Normalization. The examples have all been normalized to a storage size of 144gb (with the exception of 7200 rpm drives). In servers, MFT technology recommends a 10% free space set aside to assure minimum performance. This has been encapsulated in the normalization, in the 5 32gb drives would have a storage space of 160gb. All gigabyte sizes are based upon a billion bytes used by drive manufacturers rather than the 10243 standard.

RAID specification and performance. Because the performance of RAID-10 is so dramatically different than RAID-5 in Hard Disks, we have shown both for SAS drives. We have not shown RAID-10 numbers for Flash Media because tests show only a marginal performance gain for RAID-10 which would result in a significantly higher IOPS and cost per usable gigabyte cost.

Component Prices used are based upon published and regularly available Internet offers for the applicable product. In the case of hard drives and controllers, these are prices from significant distributors. SAS controllers are significantly more expensive than dumb SATA controllers. In the case of flash prices, these are EasyCo's published retail prices.

Chip Grades and Usable Lives. The MFT-500, and MFT-1000 models all use Multi-Level Cell technology. Both are 10,000 erase life products, and should be suitable for 1.5x random over-writing of each drive every day for five years. The MFT-3000 uses Single Level Cell technology with a presumed life of 100,000 erase cycles and thus should be suitable for 15x overwrites in a day.

The first two Flash items reflect the requirements of about 95% of all servers. For instance, our IMAP/POP/SMTP email server receives about 11gb of updates daily and has 50gb of used space in a 64gb footprint. Our calculations forecast it will last 22 years with the currently installed MLC solution. The last Flash item is what should be used for Enterprise class 7x24 heavy duty machines.

Average IOPS uses our standard based upon the average data size and read/write mix used to test Microsoft Exchange Servers. This calculation method is fully described at <http://ManagedFlash.com/technical/driveperf/index.htm>. In the case of Flash drives, all IOPS performance results are based upon the per drive test numbers published in the same section as well as latencies arising in RAID arrays. Hard Disk results are manually computed from manufacturer's published data as well as noted behaviors of RAID arrays.

Computation of Gigabyte Costs is done by dividing the total price by the number of usable gigabytes.

Computation of IOPS (Input/Outputs per second Costs is determined by dividing the total price of the item by the number of IOPS.

Drive Overwrites per Day. This number estimates the number of permissible drive overwrites in a day using the average IOPS data mix described above. For Hard Disks, this number represents the amount of writing achievable in a day if the drive operates at a 100% duty cycle performing the Average IOPS data mix described above. For Flash Media, it represents the number of times the disk can be overwritten daily and still last five years.

Given the cost of hard disks, do Flash SSDs and MFT ever not make sense?