

## "Dark" storage: wastefulness or just good engineering?

Written by Max Kalashnikov

Wednesday, June 24, 2009 - Last Updated Sunday, November 22, 2009

---

Having recently read more and more discussion about so-called dark storage, I've been reminded of something I routinely try to impress upon managers, especially clients: unless your use case is archiving, total bytes is a poor metric for storage.

In fact, the term "storage" itself may be partly to blame for the continued misconception. One need only glance at the prices of commodity disks to recognize that there isn't anything near a linear relationship between cost and bytes stored.

A quarter century ago was the golden age of the mini-computer, and the reign of the micro- was dawning. The Fujitsu Eagle was, at least in the semiconductor industry here in Silicon Valley, very popular, so it will be my yardstick. At a third of a gigabyte in usable space and just under 1.9MB/s, one could read or write the whole thing in just under 3 minutes. Today, a 1.5TB Barracuda is 4500 times the size but only 66 times the throughput, so it takes over 3 hours to go through the whole thing. A 6th-generation 450GB Cheetah is better, at under an hour.

I like the Eagle's 3 minutes as a rule of thumb. That's 21GB on larger, modern, 7200 RPM disks, and I suggest that everything beyond that may as well be considered superfluous or archive storage. Accepting this measure end-to-end means that one would only want 72GB accessible to a host off each 4Gb/s FC or 216GB per 4x SAS. Ouch.

A whitepaper from Xiotech criticizes storage vendors' performance numbers as being misleading, since they are based on short-stroking benchmarks, rather than representing the performance of the whole disk.

I suggest that short-stroking disks as a matter of course and leaving the rest purposefully "dark" is smart engineering. Suddenly, those 160GB drives look much more appealing than the 1.5TB ones, at least for performance-sensitive uses, such as databases.

Certainly, there are use cases where data beyond the 3 minute limit is still useful: anything that rarely, if ever, gets read. That tends to include backups, archives, audit trails, and even database intent logs. One may be able to have all these coexist on the same spindles as the "high performance" uses, but it would require careful forethought and testing.

## **"Dark" storage: wastefulness or just good engineering?**

Written by Max Kalashnikov

Wednesday, June 24, 2009 - Last Updated Sunday, November 22, 2009

---

My 21GB example with a 160GB disk means 87% "dark," to simulate an Eagle. It's a high percentage but nothing to be alarmed about, as long as it's done with full awareness.