

IBM smashes data storage drive record: 120 petabytes, the biggest hard drive ever?

✘ According to a [white paper study from IDC](#) there will be 1800 exabytes (EB) of digital data in 2011, up from 220 EB in 2007. Such growth places the data user under tremendous pressure to turn data into actionable insights quickly, while straining the world's IT infrastructure to its limits. This forces the data user to manage the explosive growth of data and storage using tools designed for easier times.

Smashing all known records by a multiple of 10, IBM Research Almaden, California, has developed hardware and software technologies that will allow it to strap together 200,000 hard drives to create a single storage cluster of 120 petabytes — or 120 million gigabytes.

The giant data container is expected to store around one trillion files and should provide the space needed to allow more powerful simulations of complex systems, like those used to model weather and climate.

A 120 petabyte drive could hold 24 billion typical five-megabyte MP3 files or comfortably swallow 60 copies of the biggest backup of the Web, the 150 billion pages that make up the Internet Archive's WayBack Machine.

120PB is the kind of capacity that you need to store global weather models or infinitely detailed weapon system simulations, both of which are usually carried out by government agencies or federally-funded research institutions. Alternatively, it could be used to store a large portion of the internet (or data about its users) for Google or Facebook, or another client with very deep pockets.

The largest systems currently in existence are generally around 15 petabytes — though, as of 2010, **Facebook** had a 21PB [Hadoop cluster](#), and by now it's probably significantly larger.

Exact details about the software and hardware isn't given by IBM, but we do know that it features a new-and-updated version of **IBM's General Parallel File System (GPFS)**.

GPFS is a volume-spanning file system which stores individual files across multiple disks — in other words, instead of reading a multi-terabyte high-resolution model at 100MB/sec from a single drive, the same file can be read in a massively parallel fashion from multiple disks.

The end result is read/write speeds in the region of several terabytes per second — and, as a corollary, the ability to create more than 30,000 files per second.

GPFS also supports redundancy and fault tolerance: when a drive dies, its contents are rebuilt on a replacement drive automatically by the governing computer.

GPFS has been already breaking world records by itself scoring a 37x record in file scanning (see [GPFS breaks file scanning record by 37x](#)) which means that brutal hard disk workforce is nothing if not paired with innovative software.

On the hard drive side of things, if we divide 120PB by 200,000 you get 630GB — and once you factor in redundancy, it's fairly safe to assume that the drives are all 1TB in size. We

IBM smashes data storage drive record: 120 petabytes, the biggest
hard drive ever?

also know that every single one of the 200,000 drives will be watercooled with presumably the largest and most complicated bit of plumbing ever attempted — but considering IBM's penchant for watercooling its top-end servers, that's hardly surprising (though we still hope to post a photo of the system once it's complete).

[Via: IBM Research [Almaden](#) and [GPFS](#)].