

## Storage of data from TERA

In one day, the TERA Computing Center produces a volume of data equivalent to what is stored in the French National Library in Paris, i.e. several terabytes<sup>1</sup>. Over months and years this production rate will generate a gigantic volume of data, 6 petabytes, which must remain online and accessible to users at all times. Hence, an architecture for data management at several levels, including dedicated hardware and software components, has been associated with the Tera-10 computer.



Disc array (sub-assembly)

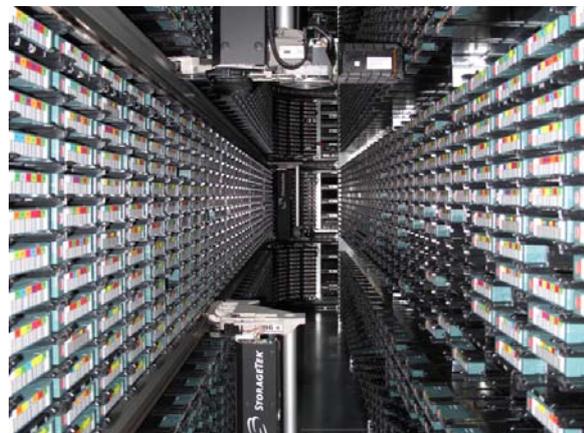
The first component is an array of 10,000 discs intended to accommodate the most recently created or used files. These discs constitute the first level of the storage hierarchy through which the data flows.

The files cannot reside long on these discs - that would require there to be several million of them. This is why they are transferred on to high capacity magnetic tapes, of order 800 GB per tape, which are stored in a storage silo. This second component constitutes the second storage level.

Each silo has several thousand slots containing the media. Eight robotic hands, genuine "automated librarians", search for the desired tapes and insert them in the readers assigned to them by the data management software.



Silos Sun/StorageTek



Interior of a Sun/StorageTek silo: Cartridge storage slots and hand (centre)

The third component is the software. This is the HPSS - High Performance Storage System -, a hierarchical storage manager. It controls all of the storage discs, tapes and robots used. As a function of the frequency of use of files, their age and size, the files are deposited onto media which have increasing capacity but which are correspondingly slower to access. This system constitutes a very high capacity data storage hierarchy.

Hence, a file spends all its life in this system, being first born on a disc and then ageing on tapes, thus releasing disc space for new files. The HPSS software is operated on a small cluster of servers which is dedicated to it.

The storage network interconnects the computers and the storage systems; it is the last component of this assembly. Based on the InfiniBand technology, it offers both very high bit rates and quick response times, and constitutes the high speed artery which is indispensable to the traffic of the data flows generated.

---

<sup>1</sup> 1 Gigabyte = 10<sup>9</sup> bytes; 1 Terabyte = 10<sup>12</sup> bytes; 1 Petabyte = 10<sup>15</sup> bytes