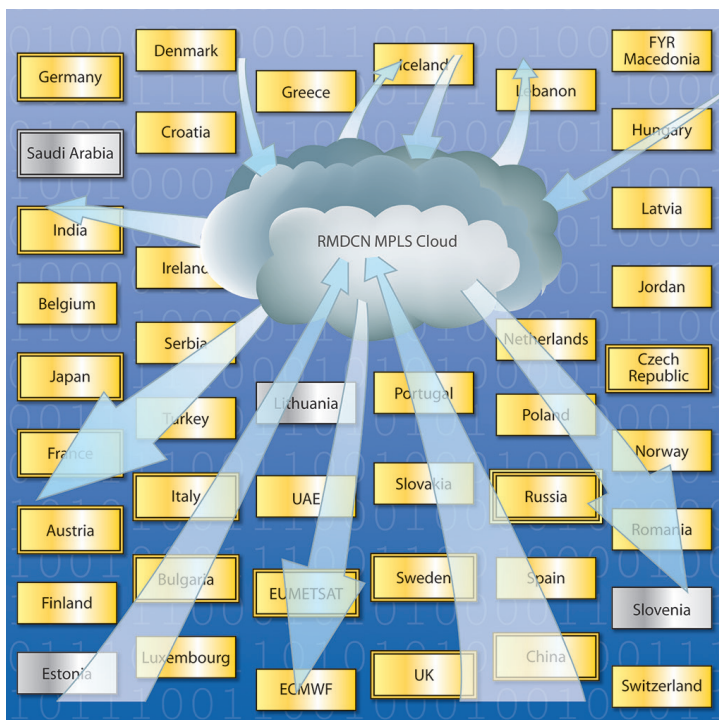




## COMPUTING

### New Automated Tape Library for the Disaster Recovery System



This article appeared in the Computing section of ECMWF Newsletter No. 113 – Autumn 2007, pp. 34-35.

## New Automated Tape Library for the Disaster Recovery System

Neil Storer

ECMWF recently replaced the Automated Tape Library (ATL) in the Disaster Recovery System (DRS) Building. In doing so it maintained the requirement that the specific tape media used by the DRS to hold secondary copies of the data should be different from that used by the Data Handling System (DHS) to hold the primary data. Here the history of the DRS will now be described followed by a comparison of the various tape drives and media technologies used at ECMWF.

In 1998 ECMWF procured its first ATL for the DRS with the specific purpose of providing storage space and magnetic tape drives to take secondary copies of important data in case the primary copy of the data was ever destroyed. Prior to then secondary copies were made using the same tape drives and media (IBM Magstar) as the primary copies.

Earlier experiences had shown the importance of having a secondary copy.

- Once some shelving holding magnetic tapes collapsed and although the data was not lost, all of the tapes had to be re-tensioned before they could be used again.
- In the early days of the DHS we had several instances of tapes holding primary data being damaged, or becoming unreadable.
- Many years ago ECMWF experienced several instances of stiction (static friction) on reel-to-reel magnetic tapes. Within certain batches of tapes the bonding between the various layers of the tape would weaken causing the magnetic oxide to flake off and stick to the tape drive heads as the tape passed over them. There was concern about the possibility of the secondary copy being written to a tape from the same batch as the primary copy, and both tapes exhibiting the same problem.

The need for more secure arrangements for making secondary copies led to an Invitation to Tender in 1998 for a new DRS (ECMWF/1998/168). The type of medium chosen was Sony's AIT (Advanced Intelligent Tape) 8-mm tape cartridges (see [www.aittape.com](http://www.aittape.com) for more information). Data is written onto AIT media using helical scan technology, where the tape is wrapped at an angle around a cylindrical drum with drive heads that are mounted flush with the drum's surface so that the tracks are striped at an angle across the width of the tape. The drum rotates quite quickly while the tape passes across it at a relatively slow pace. This is the way a VHS recorder works and it enables four to five more data to be packed onto a tape than using the linear recording method, such as was employed by the IBM Magstar tapes that were used for writing the primary data at the time. The AIT-1 cassettes could hold 25 GB of data in a very small 8-mm cassette, while the IBM Magstars could only hold 10 GB in a much larger 1/2-inch cassette. However, the AIT-1 drive could only transfer data at 3 MB/s, quite a bit less than the 10 MB/s of the Magstar. While this was fine for streaming data to/from tape, the drive was not designed for handling the type of data access patterns at which the Magstar drives excelled. The different types of tape cartridge can be seen in Figure 1. For comparison purposes Table 1 shows the characteristics of some of the latest generation of tape drives and media that are generally available.



**Figure 1** The different cartridges are shown: Magstar on the left, AIT in the middle and LTO-3 on the right.

Technology	Capacity (GB)	Performance (MB/s)
AIT-5	400	24
STK T10000	500	120
IBM TS1120	700	100
LTO-4	800	120

**Table 1** Characteristics of modern tape drives and media.

The Centre purchased an ADIC (Grau EMASS at the time) AML/J automated tape library, with twelve AIT-1 tape drives. Over the years these were replaced with AIT-2 drives and were later augmented by LTO-1 and later LTO-2 tape drives. “LTO” stands for Linear Tape Open and is an open standard set up by several tape manufacturers (see [www.lto.org](http://www.lto.org) for more information). On an LTO tape data is written in a series of tracks linearly along the tape. The tape itself passes very quickly over the stationary tape heads. LTO has the advantage of transferring data faster than AIT and storing more data per cartridge. In the end LTO was used for secondary copies of data, while AIT was used for other types of backup data. Both the AML/J and its replacement are linear robotic libraries (i.e. a set of cabinets connected in a straight line with internal racking for cartridges and tape drives and a “railway track” along which a robot runs). The robot travels in the x-direction, while the robot’s “hand” moves in the y-direction; in this way it can access any cassette and insert it into any tape drive that is installed in the library. The robotic hand mechanism can also rotate 180° in the z-direction to enable it to access cartridges and drives on both sides of the library. The robot and internal racking can be seen in Figure 2.

In 2005 ADIC informed its customers that the AML/J was coming towards end of life and would not be maintained from 2008 onwards, so in 2006 ECMWF issued an Invitation to Tender (ECMWF/2006/189) for a replacement ATL and tape drives. Tenderers were free to propose any tape technology but it was deemed to be highly desirable that they also equip the ATL with some LTO drives to facilitate the migration of the LTO cartridges in the AML/J. In the event all of the tenderers bid only LTO-3 tape technology, which contrasts with the IBM TS1120 enterprise level tape drives used by the DHS with the IBM 3592 tape cassettes to store the primary data. However none of the robots that were bid could read the barcode labels on our LTO-2 tape cassettes and as a consequence it was necessary to remove all of these labels and replace them with new ones. This manual work extended over a period of several weeks and the methodical and painstaking way this was done is a credit to all of the staff involved.



**Figure 2** The IBM TS3500 ATL located in the DRS Building. The internal racking is visible and the robotic “hand” mechanism can be seen at the bottom of the cabinet. LTO tapes are also stacked on the external racking to the right of the ATL.

The new ATL consists of an IBM TS3500 linear robotic library containing twenty LTO-3 tape drives. The LTO-3 drives can read legacy 100 GB native capacity LTO-1 cassettes, read/write the 200 GB LTO-2 tape cassettes from the AML/J, and write new LTO-3 tape cassettes at 400 GB native capacity. At 80 MB/s the performance of the drives is twice that of the old LTO-2 drives. The doubling in capacity comes from three factors.

- The tape is thinner (8  $\mu\text{m}$  versus 8.9  $\mu\text{m}$ ) and hence can be made longer, 680 m as opposed to 609 m.
- The number of tracks written across the tape has increased from 512 to 704.
- The linear density has improved from 7,398 bits/mm to 9,638 bits/mm.

LTO-4 tape drives are now making their way onto the market. The LTO-4 tape is even thinner, longer and has a much higher recording density and so is able store 800 GB of data (uncompressed) in a cartridge. These drives will provide an upgrade path for ECMWF in the future.

The library and its drives are controlled by HPSS, the underlying data management software used by ECFS and MARS to store data. A secondary copy of the “operational” MARS and ECFS data is always taken. For the remainder of the data it is left to the data’s owner to decide whether or not a secondary copy is made, the ECFS command “**ecp**” for example has a “**-b**” option to indicate that a secondary (or backup) copy should be created.

The new ATL was installed in the DRS Building on-site in the spring of this year and was run for a short period of time in parallel with the old AML/J. Once ECMWF was satisfied with its performance and reliability the AML/J was switched off and decommissioned. The old ATL provided an excellent service over its eight-year life at ECMWF and so far all indications are that the new one will provide just as good a service, if not better.

© Copyright 2016

European Centre for Medium-Range Weather Forecasts, Shinfield Park, Reading, RG2 9AX, England

The content of this Newsletter article is available for use under a Creative Commons Attribution-Non-Commercial-No-Derivatives-4.0-Unported Licence. See the terms at <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

The information within this publication is given in good faith and considered to be true, but ECMWF accepts no liability for error or omission or for loss or damage arising from its use.