# Database development: Requirements and interfaces

Dr. John Hodkinson

Database Services Group, UK Met. Office

**Summary:** *The data management system used at the UK Met. Office is described together with our use of the Internet.*

## 1 Introduction

The Meteorological Office is the United Kingdom's national meteorological service. As well as preparing weather forecasts for distribution to the Office's customers, both commercial and other parts of the Government, the Meteorological Office provides a climate research and prediction service for the UK Department of the Environment, Transport and Regions (DETR).

## 2 Computing Environment

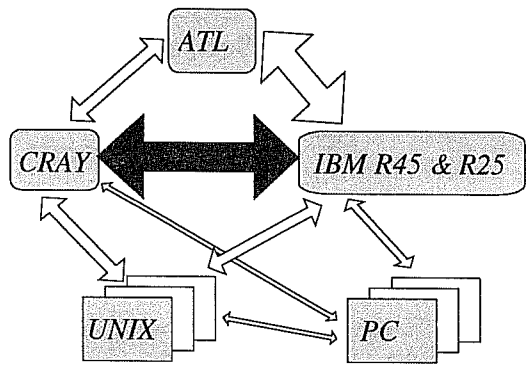


**Figure 2.1 Simplified diagram of the computing environment at the UK Met.Office.**

### 2.1 SGI/Cray T3E Supercomputer

A SGI/Cray T3E supercomputer running Unicos/mk was installed during 1997. It is configured with 880 processing elements. Attached is a disk pool of about 1400 Gigabytes, managed by SGI/Cray. A second SGI/Cray T3E supercomputer was installed in October 1999. This is configured with 640 programming elements running at higher clock speed than the original T3E. The total system performance is equivalent to the original T3E. The T3E's are connected with two HiPPI 100Mbyte/s interfaces. The supercomputers run the operational Weather Forecasting Model and the Climate Simulation Model. Output from the supercomputers is staged on the Cray disk pool. Most of the data on the Cray disk pool is then automatically passed to the GPCS disk pool and from there is archived in the ATL.

## 2.2 IBM Mainframe (GPCS)

The GPCS is Sysplex of two machines - an IBM 9672-R45 and IBM 9672-R25, running OS/390 with Open Edition. Combined the system has three Gbytes of memory. Attached is a disk pool of about 720 Gigabytes of disk storage. The GPCS is the main engine for centralised data handling and uses the GPCS disk pool to effect data transactions to and from the: ATL; Cray T3E and UNIX workstations. The GPCS is also used to pre and post process data for input to distributed applications, mainly on PCs or UNIX systems. The GPCS also hosts the general meteorological databases, the Numerical Weather Prediction (NWP) databases (known as field files), the real-time meteorological observation database (MetDB) and the climatological database (MIDAS).

Some of the GPCS disk pool is used for storage of permanent or relatively permanent data: flat files, databases, system catalogues etc. The rest is used for staging data on its way to or from the ATL.

## 2.3 Automated Tape Library (ATL)

The ATL is an ABBA/2 robotic system comprising: a GRAU Automatic Tape Library with 2 independent robots and 5 rotating towers, each with a tape cartridge capacity of 5,760 (28,800 in total). Each tape can hold up to 10Gbytes of data giving a potential storage capacity of 280 Tbytes. The ATL is connected to both the supercomputer and the GPCS by ESCON channels. Most data handling to and from the ATL is managed by the GPCS through its disk pool, although some data reaches the ATL directly from the supercomputer via the ESCON channels.

The Met.Office is purchasing a FileTek mass storage and management facility (known as MASS) to replace our current system. The main components of the system are:

- A StorageTek PowderHorn 9310-F Tape storage silo
- 32 9840 Tape drives and their 20Gb cartridges
- A Sun E6500 8-processor Unix server computer
- The FileTek StorHouse Storage Manager Software system
- The FileTek StorHouse Relational Manager Software, adapted for Met.Office use
- Goodwill from both the suppliers and the end users.

Using data compression, the 20Gb cartridges will hold 47.5Gb of Unified Model data. Only 5% on average of 'full' cartridges will be wasted. It is expected that 40Gb cartridges will come into use in 2001, and StorageTek is expected to introduce a 120Gb cartridge around 2003. Scientists in NWP will get a substantially improved facility because of the ability of accessing data through 'metadata' – in a similar fashion to ECMWF's MARS system. One aspect of MASS that may well receive interest is the ability of Relational Manager to retrieve fields of records that satisfy SQL statements in client programs. The only other site in the UK to select the FileTek software to date has already demonstrated substantially improved data analysis times when using this software.

## 2.4 Central Data Network (CDN)

Distributed computing platforms are connected to the Central Data Network (CDN). It is an Ethernet LAN operating at a nominal data rate of 10 Mbits per second (Bps) using Thick- and Thin-wire Ethernet, Fibre-optic cabling and multiple-64kbps circuits. It comprises approximately 1700 nodes.

## 2.5 UNIX Workstation and Server Systems

There are around 416 UNIX system users. The platforms are mainly HP-UNIX workstations or servers running HP-UX 10 and X-terminals. Most of these platforms do transactions to and from the GPCS and Cray disk pools. There are approximately four Tbytes of local disk storage.

## 2.6 PC Systems

There are over 1100 Windows NT PC users connected to servers on the CDN at Bracknell. Over 1.5 Tbytes of storage is available from over 20 servers. Six Tbytes of local disk storage is also distributed on the PCs.

## 3 Data storage

Observational data is stored separately for operational (MetDB) and for climatological need (MIDAS). Figure 3.1 summarises the flow of observational data through the two databases. In figure 3.1, 'TROPICS' is the main message-switch controlling the distribution of meteorological messages; it has many more recipients than indicated.

### 3.1 The MetDB

The primary purpose of the MetDB is to provide data to the operational NWP model. It ingests and stores all types of observational data, with the exception of satellite and radar imagery. It retains them for up to 5 years, with typically the most recent 10-15 days on-line. In addition, the MetDB also stores output from the analysis component of the operational NWP process, for verification purposes.

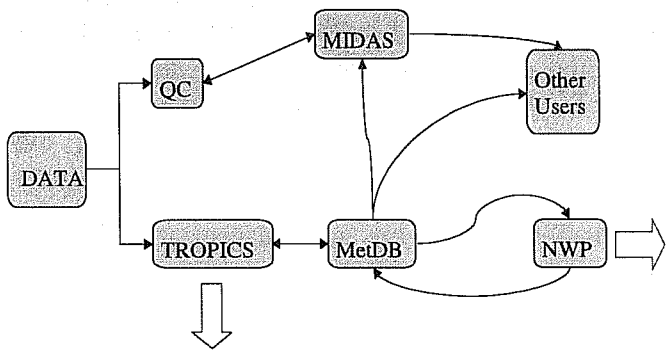


**Figure 3.1** Simplified diagram of data storage process at the UK Met. Office

Data are stored internally as BUFR BLOBs. Where the data arrive in an alphanumeric code (e.g. FM 12-XI SYNOP code) they are decoded and reformatted into BUFR before storage. Where the data arrive already coded into BUFR (e.g. many satellite data-types), they are usually stored *verbatim*. The ingestion and storage process is continuously running on the GPCS, and there is a daily transfer of older data to off-line storage. A general-purpose batch-retrieval interface is used, that makes the location of the data transparent to the user. The system is optimised to minimise the time taken to retrieve large volumes of essentially synoptic data.

### 3.2 MIDAS

The **M**eteorological **O**ffice **I**ntegrated **D**ata **A**rchiving **S**ystem (MIDAS) provides the long-term archive of observational data for climatological purposes. The data are stored in the most flexible available way using a commercial Relational Database Management System (RDBMS). Users have access to the data via SQL embedded in batch-software, or in a client-server environment using SQL (or any one of a number of GUI tools that generate SQL) via ODBC.

MIDAS holds indefinitely all data from UK sites and from UK-administered sites overseas, all marine observations and a selection of overseas surface and upper-air observations. It also holds monthly CLIMAT summaries. Data are ingested into MIDAS from the MetDB twice daily, and the automatic Q/C processing occurs daily. The original 'as-reported' data-values are retained, together with the latest quality-controlled version. In addition, data that arrive in slower time (from automatic logging equipment or in manuscript form) are transferred into MIDAS 'in arrears'.

### 3.3 Other Meteorological Databases

At the time of writing, there are currently separate databases at the UKMO for the storage, both short- and long-term, of:-

- Satellite imagery
- Radar data (raw reflectivities and derived rainfall-rates)
- NWP model-output.

Planning is underway for a rationalisation of these separate systems.

## 4 Our use of the Internet

There are a number of sophisticated visualisation packages for the displaying forecast fields, data, satellite and radar imagery and for producing products, i.e. the Met.Office products HORACE, PREVIN, MIST and NIMBUS. These, however, are relatively expensive and not readily available to employees or the public.

As such, we have made use of the Internet as a means of displaying data and products at low cost. The Internet is a relatively robust system. Despite talks of it collapsing a few years ago, it is still going strong. It also allows access to a large market at low cost. In the UK, there have been over 3 million new Internet users this year. This rapid uptake is down to the advent of free Internet service providers in the UK and lowering costs of telephone access. The Internet also allows access to a variety of systems using packages common to most PC's.

There is of course a downside to using the Internet. Security is an issue, though the careful use of firewalls should protect against attack from hackers. The Internet is also not an operational product and does not benefit from the same support as other systems. It also has variable throughput and uneven availability.

There are numerous uses we put the internet to including; simple and interactive documentation; access to data; monitoring of systems; problem reporting; phone directory; exchange of software; technical notes; user guides; codes information; minutes of meetings; displaying source code showing interactions between routines; station details; meteorological code information; CBT courses using internet; Euromet for NWP training; personnel manual; staff magazines; job vacancies; forecasts and animations; monthly briefs; service reports; helpdesk; tools and utilities; notice boards; newsletters; Web based (JAVA) server on demand to interface IBM; Web based access to e-mail; brochures; calendars of events; information on research programs and research; budgets; management reports; software download; software documentation; problem reporting; personal information - job descriptions; publications; search utilities; inventories.

## 5 Conclusion

The impact of Internet of growing importance as the number of users increases rapidly. It offers significant opportunities with marginal security risks. Uptake in some areas has been a little slow; and we need to dedicate more resources. The Internet gives us access to a much bigger market at lower cost. It gives us a highly visible public image using an interface with which people are familiar. There are potential security risks but there are solutions to these.

The UK government has a "Modernizing Government" initiative, which states that within 5 years 25% of government transactions with public must be taken electronically. As such our IT strategy is being reviewed to encourage the use of the Internet.

From a database point of view we are looking at rationalising to a single database with an intuitive interface, whilst maintaining portability.