

# StorHouse and the Met Office

## Raising the Bar for Climate Research Applications

The Meteorological Office of the United Kingdom (the Met Office), one of the world's most prestigious scientific organizations, provides national weather information to the United Kingdom and environmental and weather-related services to a global community. In 1999, faced with limited storage space and escalating storage costs, the Met Office selected FileTek's StorHouse® data management system for its Managed Archive Storage System (MASS) application for climate research. The key StorHouse features that drove this decision were:

- SQL row-level access to data on tape with no need to first restore files to disk – a feature that reduces network load and enables more efficient and timely SQL drilldowns
- SQL execution optimization for removable volumes and serial tape
- Support for direct storage and retrieval of non-relational file data through FileTek's StorHouse/RFS file interface software
- Automatic data management, including backup, recovery, migration, replication, retention, and protection to keep data accessible and available
- Scalability to multiple petabytes in a single system (with no performance degradation) to accommodate future growth
- Acceptance of industry-leading and/or new storage devices dynamically as they become available to ensure data longevity
- Reduced data ownership costs based on using cost-effective alternative storage media.

StorHouse is a comprehensive data storage and access management system specifically designed to administer massive amounts of fixed content and associated backups. StorHouse technology combines industry-leading, scalable storage devices and Open System processors with specialized storage management (StorHouse/SM), relational database management (StorHouse/RM), and file system interface (StorHouse/RFS) software components. The system supports a diverse, virtualized hierarchy of storage devices that can include FC/SCSI RAID, SATA disk, Massive Arrays of Idle Disks (MAID), high-speed automated tape, and high-density tape. Data from multiple applications can be stored on a single StorHouse system

## About MASS

The StorHouse system currently installed at the Met Office's headquarters in Exeter supports an active archive for relational weather modeling data used to investigate global warming and climate change. The archive contains 1.4 petabytes of research data and is growing at the rate of 1.5 terabytes per day.

MASS organizes data as a system of hierarchical objects comprised of namespaces, experiments, and streams. A namespace typically corresponds to one or more large research projects. Experiments represent the logical set of data within a namespace and contain unstructured and structured streams. Users archive and retrieve files through unstructured streams. Structured streams consist of a variety of formats that are archived as logical records (detail data) and constraint information (metadata) about those records. This organization permits users to archive worldwide weather observations on a daily basis for years and subsequently query information relevant only to a small geographical area over the full lifetime of an experiment.

**With the MASS system growing by over 30 terabytes each month, storing constraint information used for critical data lookups on StorHouse/MAID instead of 9840 media decreases the load on the PowderHorn silo and improves overall archival and query performance, system reliability, and data availability at a cost comparable to tape.**

The Met Office StorHouse system is highly active. The 1.4 petabytes of data it manages represent approximately 200 namespaces, 12,000 experiments, and 10 billion user objects. On a daily basis, user applications archive between 14,000 and 22,000 files to StorHouse and retrieve 4,000 to 10,000 files.

The MASS architecture consists of a Sun server running StorHouse/SM and StorHouse/RM software, a FileTek-developed front-end client application that executes on Linux servers, FC and SCSI RAID and SATA disk, 56 terabytes of MAID, and two Sun StorageTek PowderHorn automated tape libraries with 9840 high-performance and 9940 and T10K high-capacity tape drives. Each layer in this diverse storage hierarchy has a specific function. For MAID, that function is to fill the access gap between disk and tape for data with mid-to low-level access requirements.

## About MAID

MAID is a SATA-based storage subsystem where disks spin and consume power *only when required* for writing and accessing data. Initially engineered as a replacement for tape (except for offsite storage and deep archives), MAID is the ideal storage system for large data volumes that are written once and read occasionally, an access pattern frequently found in numerous applications, including scientific modeling, public record archives, telephone company call detail records, customer statements, and less active data warehouses.

MAID architecture supports cabinets containing disk drives that are spun up and down on demand. The benefits are reduced energy consumption for operation and cooling, smaller storage footprint through higher density packaging, high storage density, high storage capacity, and performance similar to conventional disk arrays at the cost of tape. Furthermore, because reliability is critical, MAID maintains a parity copy for each disk array (the same data protection as RAID) even when less than the full RAID group is powered on.

StorHouse enhances MAID efficiency and operation by managing spin-cycles based on maximizing accesses to powered-on drives where the access time is in milliseconds. By queuing read and write requests and managing when drives are powered off, StorHouse enables applications to satisfy a high percentage of accesses from already powered-on drives at disk access speeds.

## Enabling More Efficient, Reliable, and Cost-Effective Climate Research

The MASS application reaps many benefits from using MAID. Before installing MAID at the Met Office, constraint data resided on 9840 tapes, and detail data resided on 9940 cartridges. Now approximately 15-20 percent of *all* MASS data is stored on MAID. Newly archived constraint data is initially written to RAID and then automatically migrated to MAID after 30 days. Detail data is stored on MAID for 30 days, merged into one large segment (file), and then written to tape. This collection strategy boosts performance by reducing the total number of system-managed objects and their associated overhead. It also facilitates query processing on tape by clustering data on just a few volumes.



**Corporate Headquarters:**  
**FileTek, Inc.**

9400 Key West Avenue  
Rockville, MD 20850  
Phone: 301.251.0600  
info@filetek.com - www.filetek.com

**International Headquarters:**  
**FileTek UK Limited**

1 Northumberland Avenue  
London WC2N 5BW  
Phone: +44 (0) 207.872.5583  
intsales@filetek.com

©2007 FileTek, Inc. All rights reserved.  
FileTek and StorHouse are U.S. registered trademarks of FileTek, Inc. Other trademarks included herein are the property of their respective owners. The following U.S. patents protect StorHouse: 4,864,572; 5,247,660; 5,727,197; and 6,049,804

MAID provides other advantages. With the MASS system growing by over 30 terabytes each month, storing constraint information used for critical data lookups on MAID instead of 9840 media decreases the load on the PowderHorn silo and improves overall archival and query performance, system reliability, and data availability at a cost comparable to tape. Retaining detail data for the current month on MAID further enhances these benefits by providing fast access to the most frequently requested data. An additional benefit is the reduced power consumption provided by managing the disk spin-up/spin-down and the resulting reduced carbon footprint.

There is an on going need to study global warming and climate change. Because of this, the volume of weather-related research data will have multi-petabyte-growth. To accommodate this, the Met Office recognized early on the need to build its MASS application on a system that can automatically manage data as well as storage. The use of StorHouse at the Met Office has provided the solid foundation for managing the storage and data needed for climate research. Adding MAID provided additional capability as well as capacity. With MAID, the Met Office has continued to grow storage capacity and number of files; maintain the required performance; and keep data protected, secure, and available for timely access. And for selected categories of data, the Met Office will achieve the unthinkable: disk performance at a cost and scale comparable to tape.

## About FileTek

FileTek, Inc. is a provider of data storage and access management solutions, enabling enterprises, worldwide and across multiple industry segments, to rapidly access ever-growing volumes of historical/fixed content data. The Company has been supplying active archive solutions for structured and fixed content data to Fortune 1000 organizations and government agencies worldwide for over 20 years.