Masstor Systems Corporation

Origins
Masstor Systems Corporation was founded in 1976 with a charter to provide trillion-bit (terabit) Mass Storage Systems (MSS) for large mainframe computer systems. Competitive MSS products at the time were the Ampex Terabit Memory (TBM), IBM 3850 MSS ("honeycomb"), CDC 38500, Cal Comp Automated Tape Library (ATL), and Precision Instruments (???).

Traditional Offline Tape Libraries
Large mainframe computer centers in the 1960s-1970s were accumulating huge offline tape libraries containing hundreds of thousands of 2400-foot reels of magnetic tape recorded at 6250 bpi (or less). Associated with these offline tape libraries were rows upon rows of tape drives. Human operators would "mount" a single reel of tape on a tape drive as requested in the Job Control Language (JCL) for a batch computer job. This manual process caused delays and errors in processing while the job waited for operators ("tape librarians") to find and fetch the appropriate reels of tape from the tape library, mount and "ready" these tapes on the appropriate tape drives, place physical "sticky" labels on the tapes and return them to the tape library when the job completed. Lost or incorrectly labeled tapes where a frequent cause of job failure. Tape media errors were another cause of job failure.

The cost of maintaining huge tape libraries was growing rapidly due to labor, space, equipment, and lost processing costs. To overcome these problems and to meet a market need, IBM, CDC, Ampex, and others began developing alternatives to offline tape libraries – with so-called Mass Storage Systems (MSS).

Ampex TBM
One of the first commercially-successful products in this area was the Ampex TBM (Terabit Memory System) – first delivered to NSA in 1970. TBM stored data on 2-inch wide 3600-foot reels of tape using the same transverse scan technology that Ampex used for its commercial TV recording products. TBM tapes were organized into about 45,000 fixed-size numbered data blocks of about 1 million bits of data. A TBM tape reel could store about 44 billion bits (about 5.5 GB) of data. Each data block occupied about 1-inch of tape, and data blocks could be searched (scanned) at a rate of 1,000 inches per second.

There were two general architectural techniques for delivering Ampex TBM data to application programs:

- Direct Attach – in which application programs (or their agents) were "TBM aware". These applications controlled the TBM directly, spun TBM tape reels at 1,000 inches per second to find the requested data, and delivered the data directly from TBM tape to the application program. The problem with this approach is that a TBM tape can only be in one position at a time. This architecture is therefore not practical in situations with more than a few (3 or 4) application programs running at the same time.

- Indirect Attach – via a so-called File Processor (or File Server) which reduced some of the burden of cataloging data location and managing TBM operations. Even with a File Processor, application programs had to be somewhat "File Processor aware", and this was not a practical solution for the general mainframe computing environment.

IBM 3850 MSS

IBM's entry into the MSS market was the IBM 3850 MSS – developed in the late 1960s as the so-called "Comanche" project. The IBM 3850 was announced for commercial availability in October 1974 and withdrawn in August 1986. IBM 3850 data cartridges were circular cylinders, 2-inches in diameter and 4-inches long, each holding a spool of 770-inches of tape using helical scan technology. Cartridges were stored in a 2-dimensional array of hexagonal (rather than square) bins to save space – and thus giving the appearance of a "honey comb" by which the product became known. Rather than delivering data directly from cartridges to application programs, data was "staged" from cartridge to 3330 disk drives. Application programs (batch jobs) requested MSS data by a simple modification to the JCL (UNIT=MSS). A mechanical robot within the 3850 would find and fetch the appropriate cartridge and mount it in an available read/write device. The operating system (MVS) would then copy ("stage") data from the cartridge to available 3330 disk space and release the cartridge (and read/write device) for use by other jobs. Upon completion of the job, any modified data on the 3330 staging disks would be "de-staged" back to the appropriate 3850 cartridge. User (application program) transparency was a major feature of the IBM 3850.

CIA Oracle MSS with Ampex TBM

In 1971, the CIA funded a project named "Oracle" to automate the needs of its huge offline tape library in a computer center consisting mostly of high-end IBM mainframes. Ampex was the winning bidder of an RFP in 1972 with its TBM system. IBM was not able to bid its 3850 MSS because it was still under development. The proposed Ampex TBM for Oracle at the CIA was unlike any of its "direct attached" predecessors because it employed "data staging". The Oracle architecture employed techniques that were similar to what IBM would announce 2-years later with the 3850:

- Simple extension to JCL (UNIT=MSS) to request data from the Oracle.
- Cataloged Datasets which pointed to the Oracle.
- Data pre-staging from TBM tape to 3330 disk storage prior to job initiation.
- Data post-staging from 3330 disk storage to TBM tape upon job termination.

The File Processor (FP) in Oracle was more sophisticated than any of its TBM predecessors. The FP (consisting of several PDP-11 mini computers) was connected to the Ampex TBM hardware on one side and to a pool of shared 3330 staging disk hardware on the other. The FP maintained a relational database which translated abstract dataset requests (by application programs) into physical locations on TBM tapes.

Oracle included extensions to the IBM operating system (MVS) such that application program requests for MSS data were recognized before the job went into execution.

The FP received these MSS dataset requests, located the data on TBM tape, allocated space on the shared 3330 staging disk pool, copied ("staged") the data from TBM tape to 3330 disk, and then released the job for execution. User (application program) transparency was a major feature of the CIA Oracle MSS with Ampex TBM.

Masstor Founders and Key Employees
Masstor Systems Corporation was a spin-off from the Ampex TBM department – with permission and cooperation of Ampex.
Co-founders of Masstor:

- Erik Salbu, President and Chairman. Mr. Salbu was previously the TBM Department Manager at Ampex.
  Mr. Salbu managed the TBM department at Ampex since its origin in the mid 1960s.
  Prior to Ampex, Mr. Salbu was a key development engineer at IBM – working on 3330 disk storage systems.
- Lynn Shirley, VP Business Development. Mr. Shirley was previously TBM Business Manager at Ampex.
  Prior to Ampex, Mr. Shirley was a consultant for Informatics and worked on the Oracle project at the CIA.
- Bob Howie, VP Software Development. Mr. Howie was previously TBM Software Manager at Ampex.
  Prior to Ampex, Mr. Howie was the architect and technical manager of the Oracle project at the CIA.

Key Persons and Early Employees at Masstor:

- Bill Piper, VP Finance.
- Jerry Miller
- Terje Iversen
- Rick Tavan
- Bob Stevens
- Who else?

Masstor Products

- M850 – Jerry?
- SVSS – Rick?
- NETIO – Stephens?

Masstor Customers

- Shell Oil – SVSS
- NCAR
- NOAA/NESS
-

Masstor IPO

Food for Thought – Notes to Myself - Topics and Terms:

File Server
File Sharing
Data Computer
Staging
Tape Usage – 80% less than 10% used – etc.
Networking
Customers – NSA, CCA, SSA, CIA, NOAA, NCAR,
About Oracle
Tape Emulation
Disk Emulation
Channel Simulator
Applied Materials and Versatime?
X
Y
Z

Need:

Architecture Diagrams
Photographs
Customer Lists
Financial Stuff – Yearly Statements
War Stories and Fun Stuff
X
Y
Z