

## Mainframe Data Library Centera and OAM

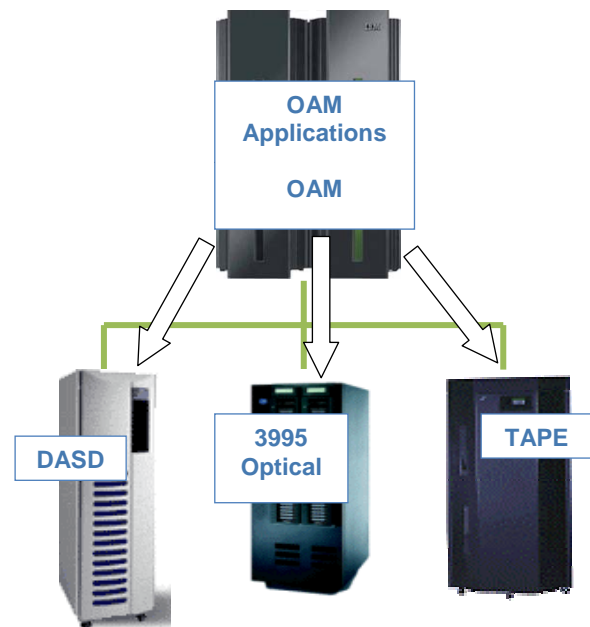
This paper describes the use of Mainframe Data Library with EMC Centera Support with IBM's Object Access Method (OAM)

### IBM'S OAM

IBM's Object Access Method (OAM) is an access method designed to support a class of data referred to as objects. An object is any stream of data bytes. The content and format of an object is unknown to the access method. There are no restrictions on the data that makes up an object. Objects differ from traditional mainframe data sets in that there is generally no concept of a record within an object. Objects can vary in size from less than a kilobyte to many megabytes and, while individual objects are generally smaller than data sets, an application often needs to store a very large number of objects. An example of a typical object may be a compressed image such as a driver's license image or an X-ray.

OAM requires objects be assigned to a **collection**. A collection is a group of objects which have similar performance, availability, backup, retention, and class transition characteristics. A collection is used to catalog a large number of objects. Each object stored within a collection must be uniquely named. Each collection defined to OAM belongs to one, and only one, **Object Storage Group**. But an Object Storage Group may contain one or many collections.

OAM is a component of IBM's System Managed Storage (SMS) and uses the SMS managed storage hierarchy. As shown in the picture above, OAM supports a storage hierarchy which can include Direct Access Storage (DASD), 3995 Optical storage, and/or Tape including 3480, 3490, and 3590. OAM allows any of

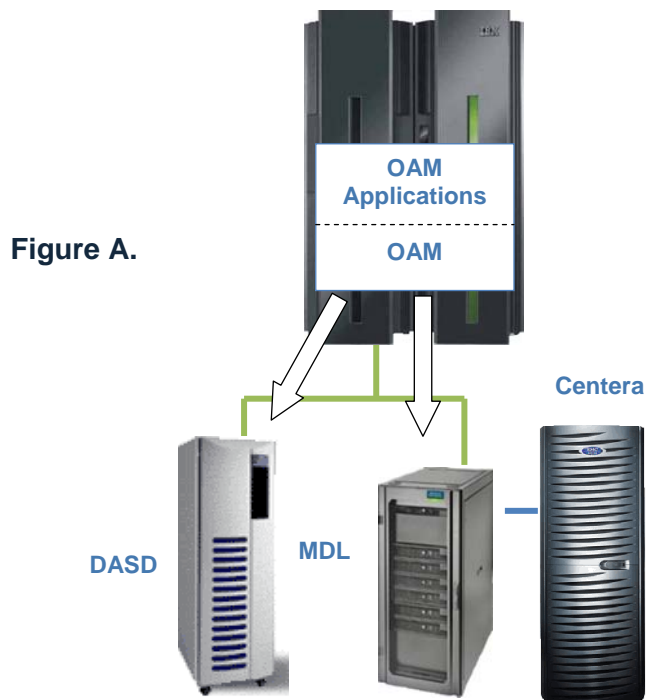


these three storage media types to be used as either the primary storage for an object or as backup storage. More important, the OAM application program interface (API) used by applications to write or retrieve objects is completely unaware of where in the storage hierarchy the object is actually stored—the application does not know whether the object is physically stored on DASD, optical (3995), or tape. OAM uses a DB2 database to keep track of the physical location of an object and it is the responsibility of the OAM administrator to ensure the collection is stored on the appropriate media to satisfy the performance and availability requirements of the application.

In a traditional OAM environment with DASD and 3995, an object may be initially stored on DASD with a backup copy stored on 3995. The object may remain on the DASD for some time in order to satisfy retrieval requirements. Eventually, when the object retrieval requirements lessen, the object may be removed from the DASD with permanent storage remaining on the 3995. From that point, however, retrieval will take longer, since the application is unaware that the object is no longer stored on DASD.

## Mainframe Data Library

The Bus-Tech Mainframe Data Library (MDL) with EMC Centera support appears to the mainframe as a collection of tape drives. The MDL is capable of emulating IBM 3480, 3490, or 3590 device types. As shown in Figure A, an MDL can be configured within SMS and OAM such that it can replace both IBM 3995 optical and real IBM tape drives in the storage hierarchy. Since OAM manages the placement of objects within the hierarchy, OAM applications are unaware that 3995 optical systems have been removed from the hardware configuration.



And, since MDL is emulating tapes and storing tape volumes as Centera objects, availability and retrieval of objects stored on Centera performs much more closely to objects stored on DASD. As a result, as objects migrate from DASD to MDL, there is minimal drop-off in the

service level when an object needs to be retrieved. Customers replacing an IBM 3995 with a Mainframe Data Library and EMC Centera have found that, while retrieval of an object from a 3995 may take up to 45 seconds, retrieval of an object from Centera occurs in 2 seconds or less 95% of the time.

## How it Works

The MDL appears to the mainframe as a collection of tape drives. When working with IBM's Object Access Method (OAM), these devices are configured as one or more Manual Tape Libraries (MTL). An MTL allows a group of one or more standalone tape drives to be managed by SMS. SMS treats the group (MTL) as if it were an automated tape library.

Once the drives are established as an MTL and defined to SMS, the MTL can be defined for use within an OAM storage group construct. The OAM storage group construct allows you to define a storage hierarchy and then manage the hierarchy as one single, high-capacity storage area. The OAM storage group is a set of volumes. The storage hierarchy allows the administrator to classify storage according to location as well as retrieval response times. A single storage group may include:

- DB2 Tables on DASD
- 3995 Optical Storage
- Tape Volumes

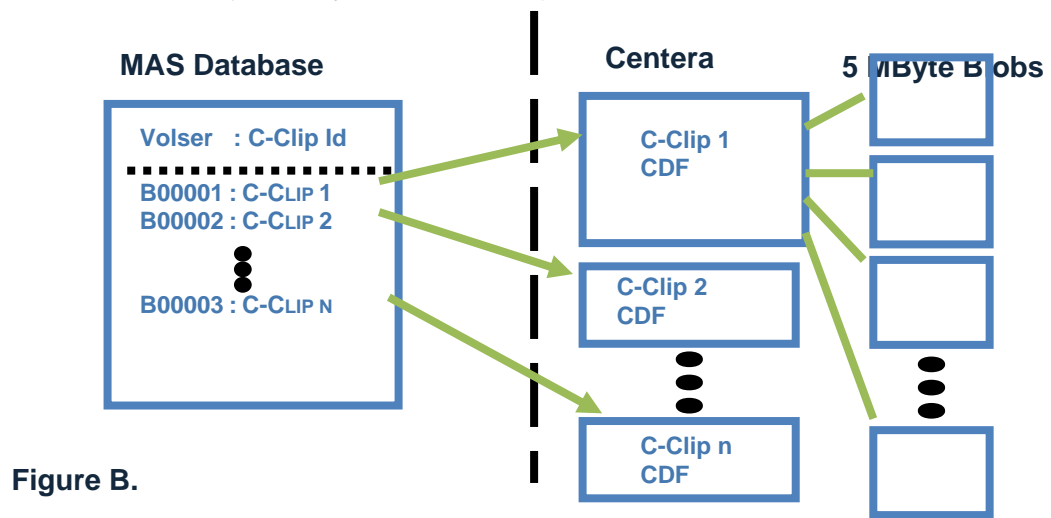
The OAM Storage Group keeps track of the physical location of each object within the group. During its lifetime, an object may move from one storage location to another or, more precisely, an object may move from DASD to tape.

Once an OAM Storage Group is defined using tape, and objects are assigned to the storage group, OAM will start writing objects to tape. When OAM first attempts to write an object to tape, it will request an SMS-managed scratch tape be mounted and it will then write the object to that tape. OAM will keep the tape mounted for a set period of time in anticipation that additional objects will be written. The tape will remain mounted and receiving objects until either the volume fills or the tape is idle for the set period of time without any object being written. If a tape has been unmounted, it will be remounted and extended as additional objects are written until OAM determines the volume is full. OAM will then mark the volume as "full" in its database and request a new scratch for the next object written.

OAM keeps track of the location of each object written to the storage group, storing the tape volume serial number and the tape block count for each object in its database. When retrieval of an object from tape is required, OAM requests the volume be mounted and then issues a block locate command to the tape drive in order to quickly locate the required data block containing the object.

The Mainframe Data Library stores each tape volume written by OAM as a separate object on Centera. The MDL maintains a cross-reference between mainframe volume serial number (volser) and Centera object ID. When the MDL receives a load mount request on the channel for a specific volser, a look-up is done to retrieve the Centera object ID and the object is opened. As shown in Figure B, data from a mainframe tape volume (volser) is written to Centera in 5 MByte BLOBs (Binary Large Objects). A single C-clip Descriptor File (CDF) is created for each volume. The CDF includes reference information to all BLOBs making up the volume. Included in the reference information is the first and last data set block numbers contained in the BLOB.

When the MDL receives a mount request from OAM to mount an existing volume, the MDL issues an “open” to Centera and reads the C-clip Descriptor File (CDF) into its memory. When OAM issues a “block locate” to locate the desired object, the MDL uses the reference information in the CDF to identify which 5 MByte BLOB contains the requested block, and then retrieves that specific BLOB from Centera. This means that the MDL never needs to read more than a single 5 MByte BLOB from Centera to begin retrieving an object requested by OAM. Regardless of the size of the tape volumes written by OAM, the MDL is able to retrieve objects anywhere in the tape in as little as 1 to 5 seconds.



**Figure B.**

Because of the improved performance when retrieving objects from tape volumes, customers find that they can maintain or improve their service level commitments while replacing aging optical and/or tape subsystems used by OAM with an MDL with EMC's Centera support. In addition, with retrieval times nearing that of DASD, customers may be able to minimize the time OAM objects need to be maintained in DB2 tables—with little or no impact to overall application performance.

## Summary and Conclusions

IBM's Object Access Method (OAM) differs from traditional mainframe file systems and access methods in that it is object-orientated rather than dataset-orientated. The size and content of individual objects is irrelevant to the access method.

OAM has always supported three different types of storage media within the OAM Storage Hierarchy: DASD (in the form of DB2 tables); optical (IBM 3995); and tape. More importantly, OAM applications are completely unaware of the storage media an object is actually stored on. Traditionally, because of the nature of the different media types, objects were stored on one type versus another based on their service requirements. Objects that needed to be retrieved in seconds were stored to DASD, while objects that could tolerate retrieval in minutes were stored to either optical or automated tape libraries.

To a mainframe, the Bus-Tech Mainframe Data Library with EMC Centera support looks in every way like a group of 3490/3590 tape drives. However, because of the near-line nature of the MDL with EMC Centera support, tape volumes stored on Centera can be quickly mounted and objects can be located with response times nearing that of online DASD. This lets customers maintain or improve their overall service level commitments while replacing optical or tape storage with newer, more robust technology.