



TECHNICAL BRIEF

OAM Manager vs. OAM native tape support

Lifecycle Director™ software for OAM Manager

NOVEMBER 2004

Executive summary 2

OAM Manager 2

Lifecycle Director software for OAM features 2

Enhanced data protection and disaster recovery 2

DB2 relief 3

Full VolSafe® secure media technology support 3

Multiple tape copies 4

Optimized I/O 4

Recycling tapes 4

Remigration avoidance 4

Volume recovery 4

Cataloged tapes 4

Enhanced prefetch 5

Dynamic load balancing 5

Disk staging 5

Operator interface 5

Tape retention on the drive 5

Extended storage hierarchy 5

Device support 5

Data conversion 6

The opportunity 6

Executive summary

StorageTek's **Lifecycle Director™ software for OAM** product provides distinct benefits and advantages in IBM Object Access Method (OAM) environments where the intention is to move inactive data to a more appropriate storage platform. The purpose of this document is to describe the difference between Lifecycle Director for OAM and the OAM product and to highlight the advantages of Lifecycle Director for OAM and OAM native tape support. Compared with native OAM tape support, Lifecycle Director software for OAM enables enhanced data protection and disaster recovery capabilities, improvements in application features and tape performance and simplified management and housekeeping.

Lifecycle Director software for OAM can help an organization automate the data migration process, free up valuable primary disk space and improve backup efficiency. In this role, Lifecycle Director software facilitates the cost savings and productivity gains embodied in an information lifecycle management strategy.

OAM Manager

The Lifecycle Director software for OAM product was originally developed in the early 1990s to provide tape archive support in OAM environments. At that time, OAM did not support tape natively — it supported only DASD and optical disk (OD). Since then, IBM has added native tape support to OAM, but this offering reflects only basic functionality. During this same period, Lifecycle Director software for OAM has been enhanced with features specific to tape archiving, which have been incorporated in response to customer requests based on operational requirements. As a result, Lifecycle Director software for OAM still offers a wide range of functional, performance-oriented and ease-of-management benefits over native OAM tape.

Lifecycle Director software for OAM is made up of two software components: an OAM integration module, **OAM Manager**, and the base archive management module, **Archive Manager**. The two modules operate in conjunction with each other to provide full tape archive support and management in OAM application environments.

A notable advantage is that OAM Manager requires no application modifications or changes to the OAM environment, other than the addition of a new storage class. Lifecycle Director software for OAM's archive and recall functionality is implemented at the individual row level, so consequently neither the application nor DB2 is aware of Lifecycle Director software during OAM's operation.

Lifecycle Director software for OAM features

Enhanced data protection and disaster recovery

Lifecycle Director software for OAM addresses or eliminates several exposures and limitations in data protection and business continuity exhibited by native OAM tape support. In particular, it enables reduced onsite exposure and improved recovery times.

• **Onsite exposure** — OAM's architecture normally involves partially filled backup storage group tape volumes remaining onsite until each one is full, at which point it can be taken offsite for disaster recovery (DR) purposes. This onsite time represents a potential DR exposure, as none of the data on the partially filled tape is located on an offsite copy. This problem is even more acute in today's high-capacity tape situations.

- It takes more archive cycles and longer durations to fill high-capacity tape, which in turn extends the exposure time.
- In a worst-case scenario, a disaster occurs when a tape is almost full and ready to go offsite. In the case of a 9840 tape cartridge, this represents 20 gigabytes of data (40 gigabytes when the T9840C tape drive becomes available), and 200 gigabytes in the case of a T9940B tape drive.

If the OAM application involves uncompressed data, normal tape compression would extend the potential T9940B data loss situation into the 400–800-gigabyte range. Obviously, the use of low-capacity tape would reduce this exposure, but the net effect would be to have all offsite backup storage group data on low-capacity tapes. This would force the site to retain some older tape technology, with the attendant handling, tape recovery and management issues coming into play.

Lifecycle Director software for OAM's combination of up to four tape copies, true tape duplexing and incremental backup features support a wide range of backup and recovery options and strategies. Most significantly, the incremental backup capability allows an offsite data copy to be in effect at all times.

• **Recovery time** — Conventional OAM architecture involves incoming data objects being stored in DB2 until migration occurs, which in turn involves DB2 image copy processing (to tape) to protect the data while it is disk-resident.

In the recovery situation after a disaster has occurred at the primary site, none of the DB2 data will be available until the object tables have been restored from the image copy backup tapes. Ironically, the most recent data, i.e. the data most likely to be accessed, will occupy the later rows in the table and will be the last to be restored. For large object table situations, where large daily volumes and/or extended on-disk retention periods are involved, the recovery time can be significant — several days, or longer in extreme cases. Restore jobs can be executed in parallel across storage groups, but this is much less effective if the OAM environment reflects a small number of large storage groups, or data volumes are heavily concentrated in several storage groups.

Lifecycle Director software for OAM's disk copy feature, as described in the DB2 relief section below, allows data to be migrated to tape immediately as it arrives in OAM, while still providing disk performance for the required period. This strategy in turn supports "immediate" recoverability in the situation described above. Once the application environment has been reestablished at the recovery site, the data is available to users immediately from Lifecycle Director software's offsite tape copy(s), at nearline response times. No delay to restore disk data is involved for immediate access to the data.

DB2 relief

The traditional approach in archive applications is to allow the data to reside on disk during its active period, and then when the recall requirements drop off to an appropriate level, the data is migrated to archive storage. As OAM is architected via a series of DB2 tables, native OAM tape applications require that the DB2 data be managed while on disk — backups, re-orgs, etc. In high-volume applications, where data is kept on disk for an extended period or where batch windows are tight, this OAM-based DB2 housekeeping activity can be a significant problem.

Lifecycle Director software for OAM can generate a disk copy in addition to the tape copies — this feature was specifically added to provide "DB2 relief."

The approach consists of writing two plus tape copies immediately when the data is created in OAM, plus a disk copy. Because migration occurs immediately, the data is deleted from the DB2 storage table, as normally occurs on data migration. The net result is that the data is still available to the application/users at disk response time, via the Lifecycle Director software's disk copy. The data is fully protected immediately (two plus tape copies), but the backup and re-org processing requirement for the DB2 storage table **has been eliminated**.

The Lifecycle Director software for OAM disk copy feature also provides the following DB2-related benefits:

- **Reduction/elimination of over-allocated DASD space** — Most OAM sites allocate additional disk space in the DB2 object tables to avoid out-of-space conditions, and the more dynamic the application's data growth pattern is, the more DB2 disk free space is typically required. Lifecycle Director software allocates the space for its disk copy dynamically and only as required, so the additional pre-allocated space can be reduced or eliminated.
- **Re-org enabling** — Many large OAM sites are rarely, if ever, able to reorganize the OAM object tables, due to the downtime required to unload/reload each table. When Lifecycle Director software's disk copy is in effect, each object table can be reduced to the point of storing one or two days' data volume, which in turn allows the object tables to be easily reorganized when required.

Full VolSafe® secure media technology support

While the VolSafe technology feature will protect data on native OAM tape, Lifecycle Director software for OAM extends this as follows:

- **Active VolSafe technology checking** — As native OAM is not VolSafe technology-aware, a human error made in tape sub-pool assignment will cause OAM to simply proceed blindly with the write operation. This will result in native OAM producing what appears to be data protected by VolSafe technology, while in fact the data is stored on regular erasable tape cartridges.

When Lifecycle Director software for OAM mounts a scratch tape in a VolSafe technology-designated write operation, it queries the tape drive to verify that an actual VolSafe technology cartridge has been loaded in the drive. If not, the process is terminated and an error is logged. Similarly, if the write operation is designated as non-VolSafe technology, Lifecycle Director software will reject the operation if a sub-pool error results in a VolSafe technology cartridge being mounted. Lifecycle Director software's fail-safe approach helps make sure that VolSafe secure media technology is in effect when specified, and avoids the potential waste of VolSafe technology media when not requested.

- **Error recovery** — If an error occurs during a write operation, native OAM will attempt to rewrite the data during the restarted run, but VolSafe technology will (correctly) prevent the over-write, and the VolSafe technology cartridge will effectively be unusable beyond the point where data was last written.

Lifecycle Director software for OAM's full VolSafe technology support will result in the rerun process rewriting the data beyond the previously written data, and the cartridge will then be used normally.

Optimized I/O

When native OAM generates two tape copies, it generates synchronous writes, so the overall elapsed time is approximately double the single-copy write time.

When Lifecycle Director software for OAM is configured to generate duplex tape copies, it uses asynchronous writes, so the overall elapsed time for two tape copies is approximately the same as for a single copy — this can be a significant advantage for high-volume situations. The feature also provides an extra level of data protection in high availability/remote vaulting situations, as the offsite tape copy is being created simultaneously with the primary copy, rather than afterwards. If a major failure/disaster occurs, all of the onsite data is reflected in the offsite copy.

Remigration avoidance

Many OAM applications, especially in the insurance industry, are case- or folder-based, where a folder contains a series of related objects. As the case or folder switches between being active or inactive after the data has been migrated, the application needs to relocate the folder contents on disk while the folder is active, to provide appropriate response time.

A common OAM technique to support this is to switch the related objects' management class to a value that will cause OAM's OSMC utility to copy the migrated data back into the DB2 storage table on disk. Once the folder contents become inactive once more, the management class is "re-flipped." In these situations, native OAM will **remigrate the objects to tape**, as if being migrated for the first time. (OAM doesn't do this for WORM optical disk, only for rewritable optical disk and tape.)

The net effect is that a folder that has gone active/inactive 10 times will have 10 tape copies, nine of which occupy "dead" space, which of course is irrecoverable on VolSafe technology tape.

Lifecycle Director software for OAM has extended capabilities over native OAM to stage data back to disk, but already-migrated data is never remigrated afterward, thus avoiding wasted tape space.

Volume recovery

When a tape volume fails, recovering the volume in native OAM can be a burden. The copies of the contents of the failed volume may be scattered over multiple backup tape volumes, causing multiple tapes to be mounted when recovering a single primary or backup copy tape volume.

Lifecycle Director software for OAM's architecture and features support a much easier approach:

- The SysAdmin flags the failed volume(s) on a Lifecycle Director software screen. This results in Lifecycle Director software automatically using one of the other tape copies to handle any incoming requests for data on the failed volume(s).
- The SysAdmin then submits a recovery batch job — no need to specify which volume(s) is involved — and the Lifecycle Director software for OAM utility executes a simple tape-to-tape copy process for the volume(s) previously flagged. Once the job completes, the recovered volumes are automatically un-flagged.

Cataloged tapes

When native OAM archives data to tape, it creates an archive data set with a common data set name, regardless of how many VOLSERS are involved. This prevents the use of standard catalog and tape management system processing and results in OAM tapes having to be tracked internally within OAM tables. This can pose problems in the event of table corruption or disk failure between tape table backups. In addition, OAM's approach requires the site to manage tapes manually with respect to offsite processing.

Lifecycle Director software for OAM uses a standard naming convention for the archive file written on each tape volume, plus the dataset is cataloged. This allows the site to use standard techniques to manage Lifecycle Director software's OAM tapes. Lifecycle Director software also includes the copy identifier in the dataset name, so standard wild-carded high level qualifier, or "vault pattern," techniques may be used to automate offsite tape management.

Enhanced prefetch

Many OAM applications use prefetch processing to stage archived data back to disk. Native OAM requires application code to use OSREQ coding techniques, which are cumbersome to use. In addition, for high-volume situations, additional table look-up and sorting processing must be coded to minimize tape mounts.

Lifecycle Director software for OAM has extended, easy-to-code prefetch capabilities, in online or batch situations, including non-DB2 environments. In addition, no additional look-up or sorting code is required. Lifecycle Director software's prefetch utility automatically handles these processes internally.

Dynamic load balancing

Lifecycle Director software for OAM has the option that allows the use of a second tape copy to handle high-demand situations where multiple requests target the same tape volume — in effect, this feature cuts the request queue in half. Native OAM has no corresponding capability.

Disk staging

Lifecycle Director software for OAM has the ability to allow recalled objects to be either passed back to the application directly, or staged back to disk, with an optional management class change. This feature is ideal for applications where any object being recalled will be accessed repeatedly for some time and so should be on disk. The change in management class allows the object to be automatically deleted from disk when appropriate. Native OAM has no corresponding capability.

Operator interface

Lifecycle Director software for OAM has an operator interface that allows the Lifecycle Director software environment to be dynamically adjusted to suit the current load conditions:

- Number of tape drives available to Lifecycle Director software
- Queue length request
- Queue wait time

Native OAM has no corresponding capability. While some parameter adjustment is possible, it is not necessarily dynamic — OAM must be re-started.

Tape retention on the drive

Lifecycle Director software for OAM can be configured to leave a tape mounted after servicing a recall request for a specified length of time, which is a good fit for applications having a high probability of a request using an already-mounted tape — “secondary hits.” This ability can be set up individually by storage group, to match the application type.

Native OAM can be set up to leave a tape mounted, but this is global across the entire OAM tape environment, which may not suit certain applications.

Extended storage hierarchy

Lifecycle Director software has the ability to extend the storage hierarchy via multiple storage levels. This is ideal for applications where recently archived data should be stored on high speed recall devices, and as the data ages, it can be moved to higher capacity, slower recall devices. Lifecycle Director software for OAM handles this process automatically via batch utilities, including space recovery of the tape volumes involved.

Native OAM can support a similar approach, but the process must be handled via OSMC and transition events, which involve DB2 updates and associated overhead processing. In addition, native OAM has no easy way to recover the tape space.

Device support

Lifecycle Director software will always include complete device support for all StorageTek® devices, both current and future. StorageTek is committed to including VolSafe technology support for all future tape devices, and Lifecycle Director software for OAM will always maintain the full active VolSafe technology support features previously described. Native OAM has no corresponding guaranteed device support.

Data conversion

Data conversion in an OAM tape environment can be considered to involve two categories:

- The initial conversion — this may be from OD to native OAM tape or Lifecycle Director software tape, or from native OAM tape to Lifecycle Director software tape.
- Ongoing conversion — this is where new tape devices become available and the existing data needs to be moved to the new media type.

In both conversion situations, the OAM applications are normally set up to start archiving data to the new target tape device, either native OAM-controlled or Lifecycle Director software-controlled, at some specific point, with data conversion proceeding afterward. In either situation, the multiple environments will peacefully co-exist, with recall requests being directed to the appropriate volumes and devices.

There are several considerations normally applied to these conversions:

- Conversion sequence — Should the data be converted in “oldest-first” or “youngest-first” sequence? Normally, the youngest-first approach is chosen, as it realizes the benefits of the new device to the data most likely to be recalled. However, the oldest-first approach might be selected in certain situations, such as cases being reviewed for disposal.
- Data separation — Normally it is highly desirable to separate the older data from newly archived data, as this concentrates the data most likely to be recalled in fewer tape volumes, which reduces the number of tape mounts, increases the probability of secondary hits and reduces overall response time.

In native OAM tape environments, **all** conversions, initial and ongoing, must be handled in the same way, involving regular OSMC processing, multiple management classes/storage classes and transition event processing, which, in turn, means that all converted objects have their DB2 directory entries updated. Over the long term, this results in each conversion exercise involving the overhead of DB2 updates (an ever-increasing load for applications where the data is never deleted), plus more management and/or storage classes being required.

In the Lifecycle Director software situation, only the initial conversion involves DB2 processing and a single new storage class. For all subsequent device conversions, Lifecycle Director software for OAM utilities will handle the conversion process internally, completely avoiding all DB2 processing and any new management/storage class assignment requirements. In addition, the Lifecycle Director software utilities have extended selection capabilities, which facilitate the conversion sequence and data separation considerations described above.

The net effect is that once data has been converted to the Lifecycle Director software environment, all subsequent conversions will involve much less processing, less I/O and correspondingly reduced overall elapsed times.

The opportunity

Lifecycle Director software for OAM offers valuable functions and features to assist in the overall management of data stored in a native OAM environment. These features and functions can help an organization generate disk cost savings, reduce DB2 housekeeping requirements, improve batch window availability, enhance data protection and realize other business benefits.

Collectively, the capabilities of Lifecycle Director software for OAM provide a much higher level of overall tape environment support than that made possible with native OAM tape support. In terms of application features, tape I/O and related performance issues, ease of use in day-to-day operational and contingency/data recovery situations, Lifecycle Director software for OAM consistently provides superior capability over native OAM.

By providing a technically advanced, cost-effective solution for tape archiving in DB2 environments, Lifecycle Director software for OAM creates the opportunity to unlock the value of underutilized disk resources and lower the cost of managing data throughout its lifecycle. These are both key steps in the information lifecycle management journey.



ABOUT STORAGETEK®

Storage Technology Corporation (NYSE: STK) is a \$2 billion global company that enables businesses, through its information lifecycle management strategy, to align the cost of storage with the value of information. The company's innovative storage solutions manage the complexity and growth of information, lower costs, improve efficiency and protect investments. For more information, visit www.storagetek.com, or call 1.800.275.4785 or 01.303.673.2800.

WORLD HEADQUARTERS

Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA
1.800.877.9220 or 01.303.673.5151