TECHNICAL BRIEF
March 2005

Lifecycle Director™ software for IBM DB2 supports Sysplex technology

1 Executive summary

Database applications contain the heart of today's businesses — information. To protect these applications, most companies employ redundancy in both the hardware and software arenas. While this works for many network and application server setups, redundancy in databases presents some unique challenges.

You can address these challenges by using very large databases. And you can mitigate the normally high costs of implementing these databases by employing Sysplex technology to link many smaller databases. Version 1.1 of Lifecycle Director™ software for IBM DB2 supports this cost-effective, scalable and high-performance solution.

This document discusses the advantages and implementation issues of using Lifecycle Director software in such a configuration.

2 The problem

Redundancy costs for mission-critical applications:

- *Increased latency (decreased performance)*
- *Additional hardware*
- *Additional software*

Performance cost of redundancy can become unacceptable as databases grow and transaction requirements increase. This is especially true with online transaction processing (OLTP) applications in which database updates need to be coordinated across multiple systems. And, of course, the cost of hardware and software redundancy can be unappealing.

For those applications that replicate to create redundant copies, data latency may increase, leaving the replicated data in an earlier state than the primary data — sometimes by hours. As the amount of information grows, you can provide performance and availability by designing larger databases (multiple tens of terabytes) so that information and processes to access the database can be added as needed.

3 Sysplex solution

With DB2 Universal Database (UDB) for OS/390 and z/OS using Parallel Sysplex technology, businesses have a cost-effective alternative that gives them the availability, scalability and high performance they require.

Sysplex support in Multiple Virtual Storage (MVS) systems allows individual system images to communicate with one another via a mechanism called the coupling facility (CF). This can be a physical device — such as the S/390 9674 coupling facility — or a logical partition (LPAR) of a processor enabled for this feature. (Note that MVS is used in this document as a generic term, covering all releases of the OS/390 and z/OS operating systems.)

MVS software allows applications running on any attached system to communicate with applications running on other attached systems via a number of structures in the coupling facility. Synchronization, automated error recovery and other processes are also supplied as part of the operating system.

4 Customer benefits

Sysplex facilities provide flexibility and scalability to support changing application requirements.

Before we get into the specific implementation features, here is a quick summary of the benefits provided by a Sysplex configuration:

- Continuous availability and processor redundancy — Any single system can be taken out of the Sysplex (such as for maintenance) without impacting the service provided by the overall Sysplex configuration.
- Improved performance and throughput — Load balancing allows the workload to be spread across all systems in the Sysplex, reducing average response times and improving transaction throughput.
- Flexibility of deployment — A Sysplex allows you to both mix different levels of hardware and software, and to dynamically and incrementally expand or reduce systems, as required.

For all these reasons, Sysplex is a good option for large database requirements and Lifecycle Director software provides all essential support.

5 Deploying Lifecycle Director software for IBM DB2 in a DB2 data-sharing environment

Each MVS system in the Sysplex has its own Lifecycle Director software control region and DB2 intercept, so that access to archived data for all archive-enabled tables on that system is scheduled via the local Lifecycle Director software control region, as with non-data-sharing implementations.

There is a key advantage to this arrangement: all local tables remain available even if a plex is in use. The only DB2 tables that cannot be accessed are those on the specific plex in use.

Lifecycle Director software for IBM DB2 minimizes contention to shared archived data using shared DASD in the Sysplex.

DB2 data sharing requires shared direct access storage device (DASD) between all systems in the Sysplex, so there are no issues involving access to Lifecycle Director software Archive Manager disk files — such as primary indexes and space management files — during object retrieval processing. This could also be implemented through a disk-mirroring mechanism in cases where shared DASD is not implemented. Because Archive Manager software requires only read access to these files for this type of processing, there is no data contention or integrity issue.

Shared DASD also gives all control regions in the Sysplex ready access to all Archive Manager disk *(K)* copy files, those containing disk copies of archived objects. Again, Archive Manager software requires only read access to these files, so all such access is satisfied from the local control region.

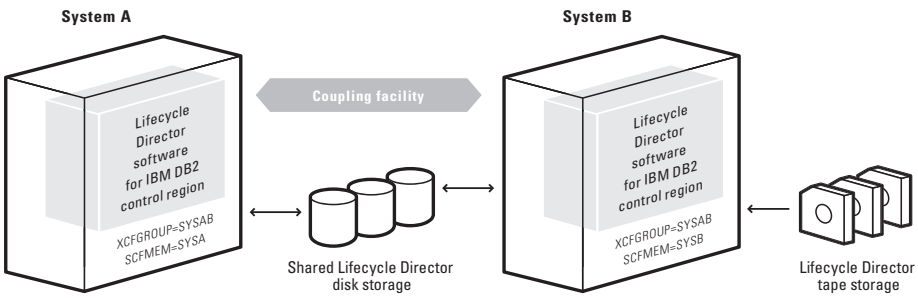

Figure 1. Lifecycle Director software for IBM DB2 configuration.

5.1 Handling concurrent tape requests

Access to records on tape devices are managed through a single control region.

The primary restriction when operating multiple control regions in a data-sharing configuration *without* Sysplex support is that shared access to a single tape volume is impossible. If concurrent requests arrive in more than one control region for retrieval of migrated rows from the same tape volume, only one of those control regions will be able to allocate the tape volume. The Lifecycle Director software system database (a DB2 database containing several tables) is implemented as a shared DB2 database, which allows control regions in each system in the Sysplex to gain shared concurrent access to the tables in the database. Sysplex support in Lifecycle Director software schedules all such requests via a single control region.

And lastly, execution of product utilities (for row archival, row restore, pre-fetch or database housekeeping) is not affected by implementation in a DB2 data-sharing environment. The product does not permit concurrent execution of utilities against a single DB2 table — this restriction will continue to be the case for data sharing implementations. A utility will not execute on one system against a shared DB2 table if the same table is being processed by a utility on another system.

6 Configuring Lifecycle Director software for IBM DB2 Sysplex support

Tape access requests are passed to the appropriate control region for processing. The results are returned to the requesting control region.

The Lifecycle Director software parameter library supplies parameters to control implementation of Sysplex support. The XCFGROUP and XCFMEMBER parameters in the ENVCNTL parameter library member identify the Sysplex group to which the control region is to belong and the member name the region is to be given in that group. When present, a Lifecycle Director software control region will be initialized as a member of a CF (coupling facility) Sysplex group. Absence of these parameters disables Sysplex support in a control region.

While there will still be one control region per MVS system, each can communicate with the others via structures in the coupling facility.

When Lifecycle Director software intercepts a retrieval request for a tape-resident migrated row and passes it to the local Lifecycle Director software control region for processing, the control region scheduler reviews tape activity across all members of the Lifecycle Director software data sharing group and assigns the retrieval request to the most appropriate tape reader task. If necessary, tape retrieval requests are passed from one control region to another for processing. Retrieved rows are then passed back to the invoking region for onward transmission to the calling application.

Lifecycle Director software selects the control region used to perform the tape retrieval in the following order:

Selecting the control region for tape data retrieval is performed using prioritized criteria.

- The control region currently processing the tape volume containing the required object
- The local control region, if it contains a waiting tape reader task with no tape volume mounted
- A remote control region that contains a waiting tape reader task with no tape volume mounted

- The control region containing the least recently used waiting tape reader task

There are two possible outcomes when a control region cannot be selected using the above rules: the retrieval task may be rejected with a “no resources available” condition (if TAPEWAIT=0), or the task may be placed in the held request queue for subsequent retry (if the TAPEWAIT value is non-zero).

The control region that initially receives the request always performs object storage or disk copy object retrieval processing.

7 Conclusion

Sysplex support in Lifecycle Director software mimimizes the possibility of multiple concurrent allocation requests for a single tape volume being issued by more than one Lifecycle Director software control region in the MVS Sysplex configuration. This scheduling, routing and balancing of requests prevents “not available” errors that would occur if more than one region requested the same tape at the same time.

This is simply another way the Sysplex configuration improves performance and throughput of retrieval requests. And that enhanced performance, along with superior redundancy and flexibility of deployment, sums up the advantages of a Lifecycle Director software-controlled Sysplex configuration.

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

FT 0011 A 04/05