

Sybase® IQ 15 Very Large Database Management Option

Maximize the value of growing data assets



PRODUCT DATASHEET

AT A GLANCE

- Range Partitioning for greater precision and granularity of data
- Logical partitioning for enhanced manageability
- Shortened backup and restore cycles
- Easily configured and managed Information Lifecycle Management

The Sybase IQ Very Large Database Management (VLDB) Option enhances the manageability of very large data sets. This option enables you to partition data logically in a range of values and to partition data in logical groups of objects called table spaces. Range partitioning enables you to work at a higher level of granularity with the data in your analytics environment by defining manageable subsets that are relevant for a particular set of queries. Table spaces also enable you to define more manageable sets of data, and to treat these as separate databases when ease of administration and/or other factors require that some data be isolated from the entire data set.

By defining table spaces as read-only, read-write, offline, or online, you can shorten backup and restore cycles by backing up or restoring only the relevant and needed data. You can also perform data validation just within a particular DBSpace or table partition target.

These partitioning capabilities enable you to group data together to meet defined business and organizational needs such as information lifecycle management policies. You can manage these partitions of data independently for purposes of backup and recovery, allowing you to relocate less critical data to less expensive storage.

USING THE VLDB OPTION FOR INFORMATION LIFECYCLE MANAGEMENT

The Sybase IQ VLDB option enables you to implement Information Lifecycle Management policies and practices that provide tremendous benefits to your organization. Among these are:

- Reduced storage costs
- Reduced time and resource dedicated to backups
- Reduced downtime
- Reduced risk of Loss

Figure 1 shows how an information lifecycle management process would work, using data partitioning to separate current data from read-only data. Figure 2 shows how such a process enables you to maintain high load and query performance as data grows, while controlling storage costs as data ages.

POWERDESIGNER® FACILITATES ILM

Version 15.1 of Sybase PowerDesigner supports Information Lifecycle management in its data model. PowerDesigner enables you to specify various phases marking the age of the data and assigning it to appropriate storage. Once the data model is in place, PowerDesigner will generate the data partitioning, placement and movement scripts for Sybase IQ.

Contact us today at infobi@sybase.com.

Visit our web site at www.sybase.com/IQ to access analyst reports, white papers, and other information.

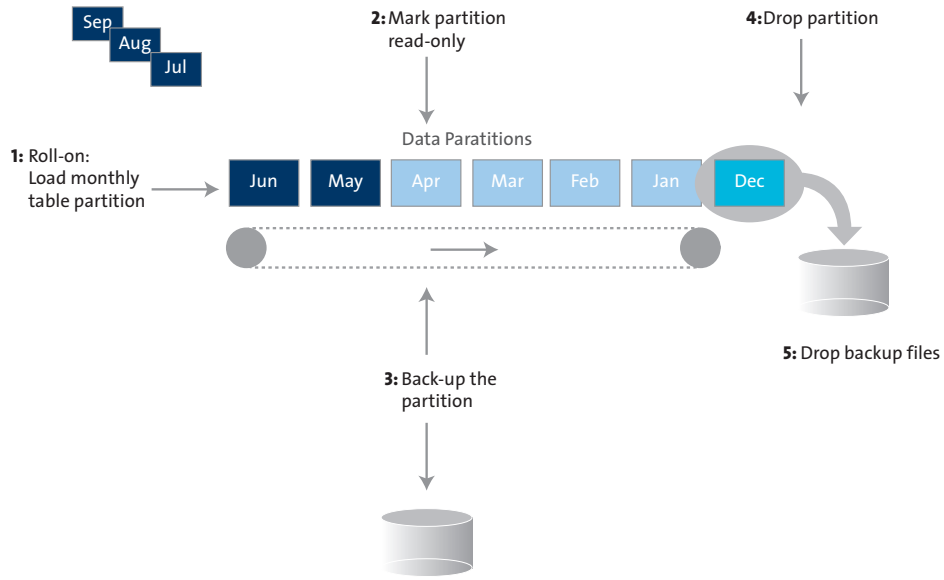


Figure 1: Partitioning current and read-only data as part of an overall information lifecycle management process

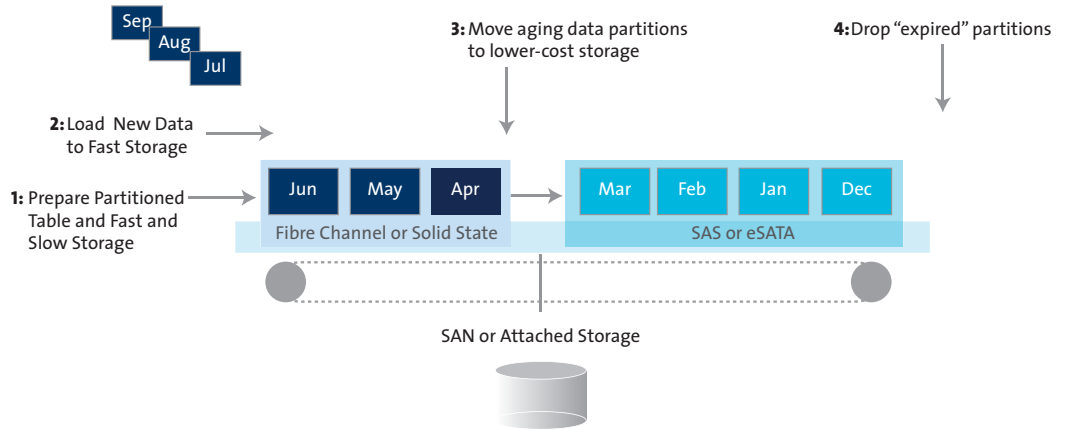


Figure 2: Data partitioning allows you to put older data on less expensive storage, maintaining high performance for current data at a relatively low cost

ABOUT SYBASE IQ

Sybase IQ is the world's leading column-based analytics server, designed specifically to deliver faster results for mission critical analytics and reporting solutions on standard hardware and operating systems. It works with diverse data—including unstructured data—and diverse data sources to deliver unsurpassed query performance at the lowest price/performance available.