

Informix Online XPS

Bob Gerber

921 SW Washington Street, Suite 670, Portland, OR 97205
gerber@informix.com

Introduction

Parallel database systems have delivered on the promise of providing excellent, linear performance speedups for both decision support and transaction processing workloads across a variety of parallel system platforms. The challenge is now to extend those results across the entire open systems spectrum of loosely coupled environments while realizing the inherent high availability of such systems. The Informix eXtended Parallel Server (XPS) will accomplish that goal.

XPS is the third and latest in the series of Online database servers that are based upon the Informix Dynamic Scalable Architecture (DSA). Earlier Online.DSA servers have proven the effectiveness of the SMP-based high performance parallel data query (PDQ) technology that is the foundation of XPS: multi-threaded process groups, table partitioning, pipelined hash-partitioned operators, light access methods, and parallel resource management. XPS extends that PDQ technology across the continuum from small clusters of large SMP systems to massively parallel clusters of SMP or uniprocessor systems.

Partitioning is the key

The key, enabling idea behind scalable database parallelism is to avoid bottlenecks by partitioning data, control and execution. The more independent the elements of the system, the better:

- Data partitioning minimizes I/O bottlenecks, establishes the base for parallel executions and improves table manageability.
- Control partitioning minimizes locking and scheduling bottlenecks.
- Execution partitioning establishes the wide-scale independent dataflows that characterize successful parallel database systems.

And, of course, there are the algorithms that thrive in parallel, partitioned environments: the hash-partitioned algorithms.

XPS partitioning enables users to effectively invest in incremental, additional hardware capabilities. The ability to predictably grow the processing power of a VLDB application is a tangible benefit of great value.

More than the sum of the parts

Individual components may fail, but the XPS system will persist. For the open systems market, that's unique. Previously, there's been a choice: a highly available database system or an open systems database system. XPS delivers the advantages of both.

The individual components of XPS are called coservers. A coserver is a process group that resides on a given hardware node and presents to the user a single system image of the larger, composite XPS super-server.

The XPS processing fabric is coserver-reinforced as opposed to being merely coserver-aggregated. That is, a failed coserver's neighbors recover the affected portions of an XPS database and then load balance the additional data access and processing load until that coserver is returned to an active state. The most critical aspect of that process is the guarantee of commercial grade, I/O clean transitions of device ownership between coservers.

Taming the physical elements

XPS introduces new capabilities that enable users to conquer the challenges that result from the sheer multiplicity of objects at the physical level of a VLDB system. The goal is to reduce the complexity of creating and managing a very large physical database to a close approximation of the level of complexity of a small, single node database system.

The elegance and power of the logical levels of a relational database are evidenced by the commercial success of these systems. The physical sediments of relational databases are another matter. Database and system administrators must often delve into the individual particles of such sediments to effectively manage a physical database. XPS simplifies such tasks by introducing the concept of a database slice (DBslice for short). An XPS administrator can create, alter and manage physical database objects at an abstract level, focusing on a slice of a physical database rather than at the detailed level of the much more numerous individual storage objects and devices.

Simplicity is a virtue

Online servers are simple to use. The Online XPS server follows in that tradition. Simplicity of use, however, should not be confused with simplicity of function. XPS delivers powerful database capabilities to the open systems VLDB market: virtually unlimited performance scalability, high availability and ease of use.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association of Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

SIGMOD '95, San Jose, CA USA

© 1995 ACM 0-89791-731-6/95/0005..\$3.50