

# INFORMIX-OnLine XPS: A Dynamically Scalable RDBMS for Open Parallel Platforms

Hannes Spintzik  
921 SW Washington Street, Suite 670, Portland, OR 97205, USA  
hannes@informix.com

## Introduction

The INFORMIX-OnLine Extended Parallel Server (OnLine XPS) is the latest in the series of OnLine database servers that are based upon the INFORMIX Dynamic Scalable Architecture (DSA).

Previous OnLine DSA server versions have demonstrated the effectiveness of the high performance Parallel Data Query (PDQ) technology on SMP systems. The key components of PDQ are multi-threaded process groups, table partitioning, pipelined hash-partitioned operators, light access methods, and parallel resource management.

OnLine XPS extends that core PDQ technology from SMP systems to massively parallel clusters of SMP or uniprocessor systems. Further, OnLine XPS leverages the inherent fault tolerant characteristics of those systems to provide database high availability for open systems.

## Scalability

OnLine XPS provides linear performance scaleups and speedups for both decision support and transaction workloads for very large database enterprises.

Scalable database parallelism is accomplished by avoiding bottlenecks through the partitioning of data, control and execution. Shared disk subsystems are not a requirement.

Data partitioning supports round robin, expression based, hash, range, and hybrid schemes. Index partitions are stored along with their corresponding data partitions. Control partitioning distributes the scheduling workload while execution partitioning establishes wide-scale independent dataflows.

These enabling mechanisms are integrated into the OnLine XPS server subsystems, rather than adding layers of functionality on top of local OnLine servers.

All this builds the foundation for the algorithms that thrive in parallel, partitioned environments: the hash-partitioned algorithms.

## Dynamic Resource Management

Large enterprise systems need to support different classes of users with different privileges. The characteristics defined for such classes will qualify, enhance or constrain what an individual user can do as well as how many resources will be assigned to that user's tasks.

For example, priority will be defined as a resource directive in one class and a scheduling imperative in another. Also, assigned resources may be dynamically changed while a task is executing to adjust to changing workloads.

## High Availability

An OnLine XPS server consists of coservers. A coserver is a process group (an OnLine instance) that resides on a given hardware node.

Coservers are the unit of high availability. That is, OnLine XPS masks failures at the granularity of an individual coserver. Data access is automatically restored when coservers fail. The re-establishment of such data access is accomplished by assigning the affected database portions to other active coservers.

The most critical aspect of commercial grade high availability is the platform support of cluster membership protocols.

The most critical aspect of commercial grade high availability is the platform support of I/O clean transitions of device ownership between coservers.

## Ease of Use

OnLine XPS presents a single server image to general users and also to database administrators.

OnLine XPS simplifies administrative tasks by introducing the concept of a database slice. An OnLine XPS administrator can manage physical database objects at an abstract level, focusing on a slice of a physical database rather than at the detailed level of the individual storage objects and devices.

INFORMIX-OnLine XPS delivers powerful database capabilities to the open systems VLDB market: virtually unlimited performance scalability, high availability, dynamic resource control 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 VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.  
Proceedings of the 21st VLDB Conference  
Zurich, Switzerland, 1995