

Oracle 8i Parallel Server for Performance and Reliability

Paul Baumgartel
InstiPro Group, Inc.
NYOUG - March 13, 2001

Introduction

- Clustering technology
- Performance and reliability
- Oracle 8i in clustered systems
- Oracle Parallel Server (OPS) basics
- New OPS features in Oracle 8i
- Locking methods

Clustering technology

- Shared disk farm
- High-speed interconnect
- Cluster management
- Reliability
- Scalability

Cluster management

- Single point of control for multiple systems
- System-wide view from management console
- Capacity planning and load balancing services
- Common data repository for reporting

Reliability and performance

- Redundancy
- Failover
- Scaleup: enhanced throughput
- Speedup: improved response time
- Load balancing

Reliability

- Provided by redundant systems
- Requires clustering extensions to operating system
- Protection against planned and unplanned downtime

Scalability

- Speedup: more hardware can perform the same task in less time
- Scaleup: how much more work can be done in the same time
- Synchronization: coordination of concurrent tasks
- Too much time in synchronization diminishes scalability

Failover

- Monitoring provided by operating system cluster extensions to detect and recover from:
 - Operating system crashes
 - Failed network cards
 - Disconnected network cables

Oracle 8i in clustered systems

- Designed to exploit cluster capabilities
- High flexibility
- High availability
- High performance
- More users

Oracle Parallel Server basics

- One database, multiple instances
- Block pinging must be avoided
- Lock management (parallel cache management)
- Distributed lock manager (integrated DLM in Oracle8)
- Application analysis and partitioning required to realize benefits of OPS

Block pinging

- Principal drawback of OPS
- An instance's request for a block in another instance's SGA requires that holding instance write block to disk first
- False pinging: request for a block locked by a multi-block lock even if row requested isn't in the block being changed

New OPS features in Oracle8i

- Integrated Distributed Lock Manager
- Cache Fusion
- Listener load balancing
- Oracle Parallel Server Management
- Cluster-aware installation
- Instance affinity

Cluster configuration for OPS

- Identical operating system and Oracle versions
- Logical filesystem management
- Support for raw partitions
- Single, cluster-wide file systems for systems and application files: not required, but highly desirable

Parallel cache management (PCM)

- Lock types
 - Fixed and releasable
 - Fine-grain and hashed
- Type of locks to use depends on type of read and write activity

Fixed locks

- Acquired at instance startup
- Greater startup overhead
- Lower operational overhead (no acquisition and release)

Releasable locks

- Acquired and released as needed
- Greater operational overhead for acquisition and release

Hashed locks

- Hashed to data blocks
- Default method of GC_FILES_TO_LOCKS parameter

Fine-grain locks

- Achieved by specifying sufficient locks
- Maximum granularity is one lock per data block (data block address, or DBA, locking)

Instance and database configuration

- Separate threads of redo for each instance (but all must be visible to each instance)
- Different rollback segments for each instance
- Assign unique instance numbers to each instance
- PARALLEL_SERVER in init.ora file

Locking parameters

- GC_FILES_TO_LOCKS
 - Maps locks to blocks within each data file
- GC_RELEASABLE_LOCKS
 - Default is value of DB_BLOCK_BUFFERS
- GC_ROLLBACK_LOCKS
 - For each rollback segment, specifies the number of instance locks available for simultaneously modified rollback segment blocks

Cache Fusion

- Transfers blocks between instances over high-speed interconnect
- Addresses reader/writer conflicts
- Consistent read image of a block held for writing by one instance is transferred to requesting instance
- Improves performance of mixed DSS/OLTP databases

Conclusion

- Parallel Server requires careful planning and design
- Other approaches may be superior (hot standby, advanced replication)
- Oracle8 implementation is more powerful and flexible