



Db2 z/OS Data Sharing Implementation, Performance and Availability

This seminar provides a review of the basics of data sharing, its implementation and various aspects of performance. Data sharing was introduced to allow multiple Db2 subsystems running on multiple z/OS systems to concurrently read or write shared Db2 databases. Data sharing in a parallel sysplex environment allows for extended processing, improved availability, and design flexibility. Db2 data sharing brings many new issues to the table such as coupling facilities, sysplex timers, group buffer pools, and global locking. We will cover all of the components necessary for operating in a data sharing environment, how Db2 uses these components and how they implemented. There are also many considerations with performance and recovery that must be considered when operating in a data sharing environment. Over the past few releases there have been several new enhancements especially in the areas regarding performance and availability. This seminar will cover enhancements up to and current with Db2 13.

Data Sharing Overview

- The Parallel Sysplex
- DB2 in the Parallel Sysplex
- SYSPLEX Implementation
- Parallel Sysplex Hardware
- Db2 Data Sharing In a Parallel Sysplex
- Sysplex Coupling Facility
- Coupling Facility Contents
- Global Locking
- Group Buffer Pools
- System Communication Area
- Db2's Use of Coupling Facility
- N-Way Db2 Data Sharing
- Data Sharing Group
- Shared Data Architecture
- Critical Success Factors

Data Sharing Design

- Planning for data sharing
- Naming conventions
- Coupling facility architecture
- Inter-Db2 locking
- Inter-Db2 Coherency
- GBP Caching
- Inter-Db2 communication
- EDM pool coherency
- On-line Utilities
- Recovery Log Environment
- Using the data sharing group
- Distributed connections to Db2 data sharing

Data Sharing Performance

- Lock Avoidance
- Locking design
- Lock architecture
- Lock contention and resolution
- Local buffer pools
- Group buffer pools
- Monitoring and Tuning GBP
- Physical lock (P Lock) issues

Application Enablement for Performance

- Commit strategies
- BIND considerations
- SQL level isolation
- Uncommitted read
- Access path issues
- Lock sizing
- Lock avoidance
- Thread reuse

Data Sharing Availability

- Availability planning
- Handling outages
- Connectivity failures
- Structure failures
- LPL issues
- Group restart
- Forcing Db2
- Log archiving and recovery
- Active log issues
- Inter-system partitioned table issues
- Affinities vs dynamic workload balancing

Recovery

- Concepts
- LRSNs
- Log Merging
- Disaster Recovery

Coupling Facility Tuning

- Structure Placement
- Structure Sizing
- Duplexing



Db2 z/OS Data Sharing Implementation, Performance and Availability

Db2 10 features related to data sharing

- Subgroup attach
- LRSN assignment improvement
- Member alias support
- Modify DDF
- Restart light enhancements
- GBP dependency optimization
- More threads support per DB2

Member consolidation considerations

- Availability
- Memory
- Capacity
- Deactivating and deleting members
- Delete and destroy
- Goals and strategies for migration and feature exploitation

Db2 11 - Data Sharing Performance and Availability Features

- Restart light with CASTOUT option
- Conditional propagation of child update locks
- Improved performance in handling lock waiters
- Increase in maximum number of CF lock table entries
- Throttle batched unlock requests
- Improved IRLM resource hash table algorithm
- Group buffer pool write-around protocol
- Improved castout processing
- Improved DELETE_NAME performance
- Automatic LPL recovery at end of restart
- Log record sequence number spin avoidance

Db2 12 - Enhancements

- Improved lock avoidance checking
- Reduced get pages and CF requests (in-memory indexes)
- Improved insert space search – reduced p-lock contention
- RUNSTATS and UNLOAD ISOLATION(UR) - avoid CF page registration
- DDF shared session data across group
- Data sharing peer recovery option
- Retry of automatic LPL GRECP recovery