

DATABASES

INTERMEDIATE

4 DAYS

IBM Db2 Administration

Four days of Db2 for Linux, UNIX and Windows administration covering architecture and instances, storage and buffer pools, object management, data movement, security, locking, backup and rollforward recovery, monitoring, optimiser tuning, partitioning and BLU Acceleration, and high availability with HADR and pureScale. Built for teams running Db2 in banking, insurance and telecommunications.

Why Db2 Administration Training?

Recovery rehearsed under realistic conditions

Participants perform rollforward recovery to a point in time and recover a database from a lost tablespace.

HADR taught as an operational skill

Teams configure high availability disaster recovery, take over under failure and reintegrate the failed node.

Built for regulated estates

Authorisation, audit and separation of duties are treated as first-class topics, which matters for the sectors that still run Db2.

Scarce skills, ageing teams

Db2 expertise is concentrated in a shrinking group of engineers. This course is aimed squarely at succession and knowledge transfer.

What you'll be able to do

- ✓ Explain Db2 architecture including instances, databases and process model
- ✓ Install, configure and manage instances and database manager configuration
- ✓ Design storage groups, tablespaces and container layouts
- ✓ Size and tune buffer pools and database memory
- ✓ Create and manage tables, indexes, views and other database objects
- ✓ Move data using LOAD, IMPORT, EXPORT and ingest utilities
- ✓ Implement authentication, authorisation and audit controls
- ✓ Explain locking, isolation levels and concurrency behaviour
- ✓ Design a backup strategy and configure logging appropriately
- ✓ Perform crash, version and rollforward recovery to a point in time
- ✓ Monitor a database using monitoring table functions and diagnostic tooling
- ✓ Interpret access plans and tune queries using the Db2 optimizer
- ✓ Apply table partitioning, MDC and BLU Acceleration where appropriate
- ✓ Configure and operate HADR and understand pureScale architecture

Who this is for

- Database administrators supporting Db2 for Linux, UNIX and Windows
- Administrators cross-training onto Db2 from another platform

- Infrastructure and platform engineers in banking, insurance and telecommunications
- Teams planning succession or knowledge transfer on a legacy Db2 estate
- Application DBAs taking on production Db2 responsibility

Curriculum

01 Db2 Architecture and Editions

- Instances, databases and the relationship between them
- Process and threading model and agent management
- Memory areas: instance, database and application shared memory
- Db2 editions, deployment options and licensing considerations
- Db2 on distributed platforms compared with Db2 for z/OS
- **Hands-on:** inspect a running instance's configuration and memory layout

02 Installation, Instances and Configuration

- Installation methods and response file based deployment
- Creating, starting, stopping and dropping instances
- Database manager configuration parameters and their scope
- Database configuration parameters and automatic tuning options
- Registry variables and environment configuration
- **Hands-on:** create an instance and database with a scripted standard configuration

03 Storage Groups, Tablespaces and Containers

- Storage groups and multi-temperature data placement
- Automatic storage, database managed and system managed tablespaces
- Container definition, extent size and prefetch size
- Page size selection and its downstream constraints
- Monitoring tablespace usage and growth
- **Hands-on:** design and implement a tablespace layout for a defined workload

04 Buffer Pools and Memory Management

- Buffer pool creation, sizing and page size alignment
- Self tuning memory manager and where to override it
- Sort heap, lock list and package cache configuration
- Prefetching, cleaners and I/O parallelism
- Buffer pool hit ratio interpretation and its limits
- **Hands-on:** size buffer pools against a working set and measure the effect

05 Managing Database Objects

- Schemas, tables, data types and constraints
- Index types, design considerations and index-only access
- Views, materialized query tables and sequences
- Table reorganisation, RUNSTATS and object maintenance
- Online schema change options and their locking behaviour
- **Hands-on:** build a schema and implement its maintenance routine

06 Data Movement Utilities

- EXPORT, IMPORT and their transactional differences
- LOAD utility phases, performance and table states
- Resolving load pending and set integrity pending conditions
- Ingest utility for continuous low latency loading
- db2move, db2look and cloning a database structure
- **Hands-on:** load a large dataset and resolve a deliberate set integrity condition

07 Security, Authentication and Authorisation

- Authentication types and where authentication actually occurs
- Instance and database level authorities and separation of duties
- Privileges, roles and group based access
- Label based access control and row and column access control
- Db2 audit facility configuration and evidence collection
- **Hands-on:** implement a separated duty model and produce an audit trail

08 Locking, Concurrency and Isolation

- Isolation levels and the anomalies each permits
- Lock modes, compatibility and lock escalation
- Currently committed semantics and reader writer behaviour
- Deadlock detection, lock timeouts and diagnostic logging
- Application patterns that create avoidable contention
- **Hands-on:** reproduce lock escalation and eliminate it

09 Logging, Backup and Recovery Strategy

- Circular versus archive logging and the recovery each permits
- Log file sizing, retention and archive destinations
- Offline, online, incremental and delta backups
- Backup compression, encryption and to Tivoli or object storage
- Designing a strategy against stated recovery objectives
- **Hands-on:** configure archive logging and take a verified online backup

10 Recovery Scenarios

- Crash recovery and the restart process
- Version recovery from a full backup
- Rollforward recovery to end of logs or a point in time
- Tablespace level restore and rollforward
- Redirected restore for relocating containers
- db2dart, INSPECT and handling corruption
- **Hands-on:** recover a database to a specified timestamp after a simulated failure

11 Monitoring and Diagnostics

- Monitoring table functions and administrative views
- Event monitors and workload level capture
- db2diag log, db2pd and first occurrence data capture
- db2top and command line diagnostics
- Db2 Data Management Console for graphical monitoring
- **Hands-on:** triage an unknown performance incident from diagnostics

12 Optimizer and Query Tuning

- Optimization classes and their trade-offs
- Explain facility, db2exfmt and reading an access plan
- Statistics, distribution statistics and column group statistics
- Access methods, join methods and sort avoidance
- Statement concentrator, package cache and parameter markers
- **Hands-on:** tune a slow query set from access plan evidence

13 Partitioning, MDC and BLU Acceleration

- Range partitioned tables, partition elimination and roll-in and roll-out
- Multidimensional clustering and block indexes
- Database partitioning feature for shared nothing scale out
- BLU Acceleration, column organised tables and their workload fit
- Choosing between row and column organisation
- **Hands-on:** partition a large table and verify partition elimination

14 High Availability and Disaster Recovery

- HADR architecture, roles and synchronisation modes
- Configuring a primary and standby pair
- Takeover, forced takeover and reintegrating a failed primary
- Multiple standbys, reads on standby and delayed replay
- TSA and Pacemaker for automated failover
- pureScale architecture and where it applies
- **Hands-on:** configure HADR, fail over and reintegrate the original primary

15 Maintenance, Upgrades and Day-Two Operations

- Fix pack strategy and applying updates safely
- Version upgrade planning, pre-upgrade checks and fallback
- Automatic maintenance windows and policy configuration
- Workload management for resource isolation
- Capacity planning, health monitoring and alerting
- Runbook and documentation standards for knowledge transfer
- **Hands-on:** capstone covering a maintenance plan and a documented upgrade path