

Data Warehousing Designs for Performance

By David Beulke

Abstract:

Does your DW and BI system need a refresh or are you wondering what are the latest DW and BI techniques for the different DB2 z/OS, Linux UNIX and Window (LUW) or DPF environments?

The **Data Warehousing Designs for Performance** presents the design architectures, performance implications and advanced Temporal topics experienced building a state of the art data warehouse and business intelligence system. Learn how to design your DW and BI architecture on any of the DB2 platforms, z/OS, LUW or DPF and understand the infrastructure considerations for maximum effectiveness, performance and overall return on investment.

This class highlights the critical platform design considerations for analyzing large amounts of data by many concurrent users. It also helps you understand the different performance impacts of database designs, ETL strategies index options, MQTs, DB2 10 Temporal tables and the many new technology options.

Big Data, predictive analytics, Cloud and NoSQL considerations have forced DW and BI strategies to evolve and you need to learn how to design your system for these latest advanced technologies, techniques and capabilities. Through this class, you will learn about the best performing architectures, the advanced design options, the latest techniques for maximizing performance and getting the fastest return-on-investment from your data warehouse and business intelligence solution.

Outline

Chapter 1: DW & BI Performance Overview

- System Considerations
- Capacity Requirements
- Physical Characteristics
- End-User Considerations
- Platforms and Connectivity
- Maintenance Aspects
- Summary Points

Chapter 2: Performance Architectures

- Architectures Trends
- Hardware Architectures
- Architecture Considerations
- DPF Node Considerations
- Storage Architectures
- Processing Architectures

Chapter 3: Extract Performance Characteristics

- Extract Source Performance
- Extract Validation for Quality
- Extract Data Manipulation
- Extract Scheduling Performance

Chapter 4: Designs Options for Performance

- Scope Design Considerations
- DW Fact Table Types
- Fact Table Designs and Grain
- Dimension Types, Flat & SCD Considerations
- Code Table Designs
- Clustering and Sequence Design Options
- Star versus Snowflake Schemas
- Answer Optimization – History Tables

Chapter 5: Key Designs

- General Key Designs
- Different Key Types
- Index Keys Considerations
- Key Designs for Partitioning
- DB2 10 Index Advantages

Chapter 6: Aggregates Performance

- Aggregates Design
- MDC-Multi-Dimensional Clustering Tables
- MQT-Materialized Query Tables Designs
- MQTs- Performance Processing Strategies
- OLAP functions and Aggregates & MQTs
- Operational DW MQT Designs

Chapter 7: Temporal Data Designs

- Temporal Table Design Overview
- History tables consideration
- Temporal Considerations
- BUSINESS_TIME and SYSTEM_TIME
- Bi-Temporal Considerations
- Temporal SQL Considerations

Appendix 1: Multi-Thread Prototype Process

Appendix 2: Partition Keys and Aggregates

Appendix 3: DW Capacity Planning

Appendix 4: OLAP Tool Evaluation

Additional or custom material substituted per request