



TDWI Advanced Dimensional Modeling: Techniques for Practitioners

Training Calendar

Date	Training Time	Location
19 August 2019	2 Days	Bilginç IT Academy
04 November 2019	2 Days	Bilginç IT Academy

Training Details

Training Time	:	2 Days
Capacity	:	10
Prerequisites	:	This course assumes a basic understanding of dimensional modeling concepts, techniques, and terminology.

About Training

About Training

Students learn the complete set of best practices—from multiple fact table designs, to bridge tables, to advanced slow change processing. Students learn to match these techniques to real-world business complexity, and explore their impacts on BI and data integration tools. Students also learn how to fit dimensional modeling into agile development frameworks, and receive templates for capturing requirements and designs.

What You'll Learn

- Why most subject areas require multiple fact tables, and how to identify them
- When to use alternatives to the basic transaction fact table, including periodic snapshots,

accumulating snapshots, and type-specific stars

- How to cope with dimensional intricacy using techniques such as bridge tables, mini-dimensions, time-stamped dimensions, hybrid slow changes, and other slow change options
- Techniques to ensure your data warehouse will scale as new subject areas are added
- How design fits into development methods, who should be involved in design activities, and what tasks and outputs should be incorporated

Who Should Attend

- BI program managers
- Business analysts
- Data architects and modelers
- BI architects and developers
- Data integration architects and developers
- Project managers
- Database administrators
- “Power users” and business subject matter experts

Outline

Module One

- Fundamentals
- Dimensional Modeling
- Process measurement
- Uses beyond database design
- Information Architectures and Dimensional Data
- Inmon’s CIF Architecture
- Kimball’s Dimensional Bus
- Stand-alone Data Marts
- Fundamental Terms and Best Practices
- Facts and dimensions
- Surrogate keys, natural keys and slow change processing
- Fundamental concepts including grain, sparsity and additivity
- Best practices for fact tables and dimension tables
- Implementations
- Relational (star and snowflake)
- Multidimensional (cube)
- Cubes as primary store
- Cubes as derived data stores

Module Two

- Multiple Stars
- Multiple star solutions
- Designing multiple fact tables

- Identifying multiple processes
- Differences in dimensionality
- The pitfalls of single fact table design
- Using multiple stars
- How not to query multiple fact tables
- The concept of drilling across
- What you need to know about your query and reporting tools
- Conformance and business value
- High impact business questions span processes
- The concept of conformance
- Ensuring subject areas work together
- Enabling incremental implementation

Module Three

- Advanced Fact Table Design
 - Transaction schemas
 - Transaction grain
 - Shortcomings of transaction designs
 - Periodic Snapshots
 - Snapshot grain and period
 - Semi-additivity, density, and impact on BI
 - Building both transaction and snapshot schemas
 - Snapshots and averages
 - Accumulating Snapshots
 - Studying process efficiency
 - Accumulating metrics in a single row
 - Lag analysis
 - Impacts on slow change processing and data integration
 - Building both transaction and accumulating snapshots
 - Factless Fact Tables
 - Processes that seem to lack metrics
 - Factless fact tables that track events
 - Pros and cons of adding constant-value fact
 - Factless fact tables that track conditions
 - Comparing conditions to actual events
 - Heterogeneous Attributes
 - Attributes that vary based on category
 - The impact of modeling a single set of attributes
 - Core and custom dimensions
 - Core and custom fact tables
 - Alternatives to core and custom solutions

Module Four

- Design and Data Integration

- The data integration process
- Loading a dimension table
- Loading fact tables
- Key lookup processing
- Designing to aid key management
- The impact of slow changes on key lookups
- Adding attributes to aid lookup process
- Designing to reduce slow change bottlenecks
- The bottlenecks of slow change processing
- How dimension checksums aid slow change processing
- Specifying data transformation rules
- How to capture standard translations and data quality rules
- Why this must be done at design-time
- Invalid or late reference data
- Receiving valid facts with invalid or missing reference data
- Adding rows to dimension tables so that facts can be loaded
- Adjusting facts when reference data arrives
- Adding columns to support QA
- Housekeeping columns in dimension tables
- Housekeeping dimensions for fact tables

Module Five

- Advanced Dimension Design
- Understanding hierarchies
- Drilling with and without hierarchies
- Multiple hierarchies in one dimension
- Impact of hierarchies on BI, ETL and DBA perspectives
- Why to model hierarchies
- Snowflake schemas
- Dimension Reuse
- Modeling roles
- Querying with roles
- Nulls
- Problems introduced by NULL values
- Avoiding NULL dimension attributes
- Avoiding NULLs with optional relationships
- NULL facts

Module Six

- Advanced Slow Change Processing
- Time-stamped Dimensions
- Why type 2 is not good enough
- Time-stamped dimensions (transaction dimensions)
- Using with a fact table

- Hybrid attributes that behave like facts and dimensions
- Mini-dimensions
- Large and expanding dimension tables
- The use of a mini-dimension to stem growth
- Impact on schema capability
- Type 4 and Type 5 slow changes
- Loading the mini-dimension
- Current and Previous
- Limited access to unchanged value
- Hybrid Responses
- Tracking both Type 1 and Type 2 responses
- For dimension Attributes (Type 6)
- For entire dimension tables (Type 7)

Module Seven

- Bridge Table Design
- The dimension bridge
- When a single fact needs to reference more than one dimension row