

Offered by **Chris Adamson** through **TDWI Onsite Education**
The companion course to the book *Star Schema: The Complete Reference*.

Course Agenda, Day 1 of 2

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

Course Agenda, Day 2 of 2

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 (Types 4 & 5)
 - 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 (Type 3)
 - 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

- The “flattening” option
- The bridge table
- Use of the bridge
- Avoiding double counting with an allocation factor
- Avoiding double counting by hiding the bridge (sandbox)
- The attribute bridge
 - A dimension attribute repeats for a single dimension row
 - The attribute bridge
 - Use of the bridge and avoidance of double-counting
 - Bridge table vs. factless fact table
- The hierarchy bridge
 - Recursive relationships (ragged hierarchies, unbalanced hierarchies, variable depth hierarchies and instance hierarchies)
 - Flattening and backfilling
 - The hierarchy bridge and its structure
 - Using the bridge to “roll up” a recursive hierarchy
 - Using the bridge to “roll down” the hierarchy
 - ETL implications of hierarchy change and slow changes

Module Eight

Scaling Dimensional Designs

- Scaling and scope
 - Conformance across subject areas
 - Conformed dimensions that are not identical
 - Conformed rollups and overlapping dimensions
 - An advanced conformance matrix
- Derived schemas and performance
 - Merged fact tables
 - Pivoted fact tables
 - Set operations on fact tables
 - Sliced or partitioned fact tables
- Aggregate schemas and performance
 - N-way aggregate design
 - Aggregate portfolio and impact on throughput
 - Aggregates and aggregate navigation

Module Nine

Design Tasks and Deliverables

- Design and data architecture
 - Dimensional design as architecture task
 - Architecture and incremental implementation
 - Implementation without up-front architecture activity
- Design activities
 - Key design activities
 - Resources required for design success

- Estimating duration of design activities
- Requirements, top level design, detailed design and roadmap
- Documenting requirements (Conceptual design)
 - Subject area definition
 - Metric groups and metric definition
 - Business hierarchies
 - Conformance
- Top level design
 - Key elements of dimension table design
 - Key elements of fact table design
 - Documenting what is not visible
 - The conformance matrix
 - The design review
- Detailed design
 - Attribute level design of fact and dimension tables
 - Technical design elements (DBA perspective)
 - Business definitions (business perspective)
 - Source to target mappings (integration perspective)
- Roadmap
 - Project definitions
 - Managing scope of implementations
- Implementation approaches
 - Time-boxed projects with iteration
 - Agile projects (with and without up-front discovery)
 - Typical team member roles for implementation projects

Appendix A

Exercises

- Ten modeling exercises designed to reinforce concepts of each module
 - Two are completed during the class (one each day)
 - Students can work through remaining exercises on their own
- Exercise sessions
 - Students are divided into groups of 4-5 to complete the exercises
 - Flip chart or white board must be provided for each group
 - Exercise is discussed by class as a whole upon completion

Appendix B

Exercise Solutions

- Solution pages for each exercise
- Example models and discussion points