Workshop Expectations

Module 1 – AI-Ready Semantic Models

- What is Generative BI?
- Power BI Semantic Models are Foundational
- Why Technically Correct Models Can Fail
- Business Semantics vs. Technical Metadata
- What Does "Prep Data for AI" Do?
- Understanding AI Limitations

Module 2 – Our Realistic Case Study

- The Business Scenario
- Understanding the Model
- The Important Business Measures
- Establishing a Baseline for AI Readiness
- Hands-on Lab #1

Module 3 – Enriching Semantic Models

- Why Business-Friendly Names Matter
- Adding Descriptions
- Improving Measures for Business Questions
- Understanding Date Roles
- Organizing & Hiding Model Objects
- Hands-on Lab #2

Module 4 – Simplifying the AI Data Schema

- Introducing the AI Data Schema
- Choosing Relevant Tables & Columns
- Choosing the Measures
- Excluding Irrelevant Objects
- Balancing Simplicity & Analytical Flexibility
- Hands-on Lab #3

© David Langer

Module 5 – Communicating Business Rules to AI

- Introducing AI Instructions
- Providing Business Context
- Defining Metrics & Their Meaning
- Communicating Date-Role Rules
- Adding Analysis Rules & Guardrails
- Writing Clear, Testable Instructions
- Hands-On Lab #4

Module 6 – Verified Answers

- What Are Verified Answers?
- When Should You Use Them?
- Creating a Verified Answer
- Verified Answers vs. AI Instructions
- Limitations & Practical Considerations

Module 7 – Testing

- Why the AI Configuration Must Be Tested
- Realistic Stakeholder Questions
- Developing Expected Answers
- Validating Copilot's Responses
- Diagnosing Model vs. Configuration Issues
- Refining & Retesting
- Hands-On Lab #5

Module 8 – Workshop Wrap-Up

- Applying What You've Learned
- Understanding Current Limitations
- Introducing Fabric IQ & Fabric Data Agents
- Continue Your Learning

© David Langer