



AI-Assisted Power BI for Business Analytics

Figure 1: AI-Assisted Power BI for Business Analytics

Lab 2: Prompt-Driven Report Creation

Module 2: Creating Reports with Copilot

Duration: 35 minutes **Dataset:** CloudRevenue (Microsoft Cloud Subscription Analytics) **Deliverable:** Published report + screenshot of best prompts

Learning Objectives

By the end of this lab, you will be able to:

- Generate multi-page reports using natural language prompts
- Apply iterative prompting to refine AI outputs
- Recognize what makes data “AI-ready”
- Evaluate and improve Copilot-generated reports

Before You Begin

Required Reading

1. **CloudRevenue BRD** - View BRD
 - Focus on Section 3: Report Requirements
 - Note the specific KPIs and visualizations requested
2. **CloudRevenue Data Dictionary** - View Data Dictionary
 - Know which columns contain the metrics you need
 - Understand the relationships for meaningful analysis

What You’ll Need

- Power BI Service account with edit permissions
- Access to CloudRevenue semantic model in course workspace
- Notepad or document to record your prompts

Lab Instructions

Part 1: Create Your First AI-Generated Report (10 minutes)

1. **Navigate to the Course Workspace**
 - Go to app.powerbi.com
 - Open the VT-MBA-PowerBI-[Section] workspace
2. **Start a New Report**
 - Find the semantic model named **CloudRevenue**
 - Click the **three dots (...)** menu
 - Select **Create report** → **Start from scratch**
 - You should see a blank canvas with the Copilot pane available
3. **Generate Your Initial Report**
 - Open the **Copilot pane** (sparkle icon)

- Try this initial prompt:
Create a report about cloud subscription revenue showing total revenue, revenue by product category, and revenue trends over time

- Press Enter and wait for Copilot to generate the report
- This may take 15-30 seconds

4. Review the Output

- How many pages did Copilot create?
- What visuals did it choose?
- Does it address the BRD requirements?

Screenshot #1: Capture your initial Copilot-generated report

Part 2: Iterative Refinement (15 minutes)

The first output is rarely perfect. Now you'll refine it through additional prompts.

Iteration 1: Add Missing Elements Review your report against the BRD. What's missing? Try prompts like:

"Add a page showing revenue by region with a map visualization"

"Include a table showing top 10 products by revenue"

"Add customer segment breakdown"

Record your prompt and the result: _____

Iteration 2: Improve Existing Visuals

"Make the revenue trend chart show monthly instead of daily"

"Add data labels to the bar chart"

"Change the color theme to be more professional"

Record your prompt and the result: _____

Iteration 3: Add Context and Filters

"Add slicers for year, region, and product category"

"Include a card showing year-over-year growth percentage"

"Add a comparison of this year vs last year"

Record your prompt and the result: _____

Iteration 4: Executive Polish

"Make this more executive-friendly"

"Simplify the layout - fewer visuals, bigger fonts"

"Add a title and summary section"

Record your prompt and the result: _____

Part 3: Prompt Engineering Experiments (5 minutes)

Try different ways of asking for the same thing. Compare results:

Experiment A: Specificity

Vague: "Show me sales data"

Specific: "Create a bar chart showing total revenue by product category for 2024, sorted descending, with data labels"

Which gave better results? Why?

Experiment B: Business Context

Technical: "Create a line chart of MRR over time"

Business: "Show me how monthly recurring revenue has trended this year so I can identify growth patterns"

Which gave better results? Why?

Experiment C: Iterative vs. Comprehensive

All at once: "Create a 5-page executive dashboard with revenue trends, regional breakdown, product analysis, customer segments, and KPI cards"

Step by step: "Create a page about revenue trends"

→ "Now add a page about regional performance"

→ "Add a product analysis page"

Which approach worked better? Why?

Part 4: Finalize and Publish (5 minutes)

1. Final Review

- Check that your report addresses the BRD requirements
- Ensure visuals are clear and readable
- Verify data accuracy on at least one metric

2. Save Your Report

- Click **File** → **Save**
- Name it: `CloudRevenue_[YourName]`
- Save to the course workspace

3. Capture Your Best Prompts

- **Screenshot #2:** Your final report
- **Screenshot #3:** The Copilot pane showing your most effective prompts

Lab Deliverable

Submit: Two items via Canvas

Item 1: Published Report

- Your report should be visible in the course workspace
- Instructor will verify: `CloudRevenue_[YourName]`

Item 2: Prompt Documentation

A document (1 page) containing:

1. **Screenshot of your final report** (full page view)
2. **Your 3 most effective prompts**
 - The exact text you used
 - Why it worked well
 - What Copilot generated
3. **Your 1 least effective prompt**
 - What you asked
 - Why it didn't work
 - How you recovered

Submission Format

- **File Name:** Lab2_Prompts_[YourName].docx or .pdf
- **Submit via:** Canvas → Assignments → Lab 2

Completion Checklist

Before submitting, confirm you have:

- Generated initial report with Copilot
- Refined through at least 4 iterative prompts
- Experimented with prompt variations
- Saved report to course workspace
- Captured screenshots of prompts and final report
- Documented best and worst prompts with explanations
- Submitted via Canvas

Prompting Competition (Optional)

Your instructor may run a class competition:

Challenge: Create the best CloudRevenue dashboard using only Copilot prompts

Judging Criteria:

Criterion	Weight
Addresses BRD requirements	30%
Visual clarity and design	25%
Insight quality	25%
Prompt efficiency (fewer = better)	20%

Prize: Bragging rights and prompt library contribution

Key Takeaways

1. **Start broad, then refine:** Initial prompts set direction; iterations add polish
2. **Specificity wins:** The more specific your prompt, the better the output
3. **Business context helps:** Telling Copilot *why* you need something improves results
4. **Iterate don't restart:** Build on what Copilot creates rather than starting over
5. **Document what works:** Your prompt library is a valuable personal asset

Prompting Pattern Library

Add these to your personal collection:

Report Creation Patterns

"Create a report about [topic] showing [metric1], [metric2], and [metric3]"

"Create a [audience]-friendly dashboard for [purpose]"

"Build a report that answers: [business question]"

Refinement Patterns

"Add a page focused on [specific area]"

"Change [visual] to show [different aggregation/dimension]"

"Make this more [executive/detailed/visual/data-dense]"

Comparison Patterns

"Compare [period A] to [period B]"

"Show [metric] by [dimension] vs [another dimension]"

"Highlight the difference between [segment A] and [segment B]"

Preparation for Next Lab

- Complete Prompting Journal entry #2
- Read: Microsoft Learn - AI Visuals in Power BI
- Think about: What factors might influence customer satisfaction?

"The best prompt is one that gets you 80% there in one shot, leaving you time to perfect the remaining 20%."