

# MANUEL MUÑOZ JR

## Senior Business Intelligence Solutions Architect

*Power BI · Microsoft Fabric · Power Platform · AI-Augmented Analytics*

### ABOUT THIS DOCUMENT

This brief accompanies the full interactive portfolio at [munozdataworks.com](https://munozdataworks.com). It is designed to be read in 10–15 minutes and shared with hiring panels who may not have time to explore the site directly. It is not a resume — a separate, formal CV is available for download from the portfolio. This document instead provides a structured, comprehensive view of the elements that a resume cannot capture: technical breadth, architectural philosophy, verified credentials, demonstrated reasoning, and professional endorsements.

*Each section below corresponds to a section of the live portfolio and includes a link back to the relevant page when deeper exploration is warranted.*

---

### CONTACT

**Email** [manuel@munozdataworks.com](mailto:manuel@munozdataworks.com)

**Web** [www.munozdataworks.com](http://www.munozdataworks.com)

**LinkedIn** [in/manuel-munoz-jr](https://in.linkedin.com/in/manuel-munoz-jr)

**Location** Coastal North Carolina · Remote-first

## DOCUMENT MAP

# Contents

---

<b>01.</b> Professional Identity	3
<b>02.</b> Industry & Enterprise Footprint	4
<b>03.</b> Technical Architecture Stack	5
<b>04.</b> Specialties & Domain Strengths	8
<b>05.</b> Verified Certifications	9
<b>06.</b> Continuous Learning Snapshot	10
<b>07.</b> Architectural Methodology — Case Studies	11
<b>08.</b> Reasoning Samples — Selected Q&A	12
<b>09.</b> Professional Endorsements	14
<b>10.</b> Selected Writing & Thought Leadership	14
<b>11.</b> Engagement Modalities & Contact	16

## 01 · PROFESSIONAL IDENTITY

## How I Practice

---

I am a Senior Business Intelligence Solutions Architect specializing in the Microsoft Fabric ecosystem, with 15+ years architecting end-to-end analytics platforms for Fortune 100 organizations across Financial Services, Healthcare, Defense, and Food Manufacturing. **My core practice is the design of clinically honest data architectures** — semantic models, governance frameworks, and consumption layers that transform fragmented, manual, unreliable data processes into structured, governed, insight-ready systems.

My core expertise lies in building semantic models and designing clear, compelling visualizations that support operational monitoring and actionable executive decision-making — models that are human-readable, performance-optimized, and aligned to precise metric definitions that leaders can trust. I also create organizational efficiency through process automation, ensuring analytical workflows are consistent, scalable, and repeatable.

**Technical depth with strategic vision.** My work spans the full lifecycle of modern analytics architecture: data ingestion, storage, transformation, semantic modeling, metric development, dashboard design, analysis, and user experience. While my primary focus is BI architecture, I work closely with data engineering teams to guide upstream strategy — ensuring that Lakehouse design, Delta table structures, and transformation logic align with analytical and governance requirements. In addition to SQL-based shaping, I work extensively with Power Query Online and Dataflows Gen2, enabling scalable, maintainable, and highly automated transformation pipelines within Fabric.

**What sets me apart.** Beyond technical expertise, I bring a strategic mindset honed through years of working across diverse business domains and organizational challenges. I understand how to balance immediate business needs with long-term architectural vision — *the tension that separates solutions that work from solutions that last*. I'm equally committed to the human element: educating stakeholders, simplifying complexity through intuitive design, and ensuring that the solutions I deliver are actually used and valued by the people they're meant to serve.

What drives me is architecting concrete, scalable solutions from ambiguous business needs — designing the data, semantic, and consumption layers that produce reports, dashboards, and automated workflows as *downstream outputs of a coherent architectural foundation*. The result is analytical environments where data flows cleanly, models remain consistent across teams and time, and insights reach decision-makers with both technical integrity and narrative clarity.

*I am currently seeking roles where thoughtful architecture, semantic modeling excellence, and clarity-driven analytics are recognized as strategic drivers of organizational performance — full-time, contract, or advisory.*



02 · INDUSTRY & ENTERPRISE FOOTPRINT

# Where the Work Has Lived

Direct engagement experience across eight Fortune 100 and multinational organizations spanning Financial Services, Defense, Healthcare, and Food Manufacturing — sectors where auditability, lineage, and governance are not afterthoughts but design constraints. The summary below pairs each engagement with a representative architectural contribution.

Organization	Key Contributions
<b>FINANCIAL SERVICES</b>	
<b>BBVA Compass</b>	Designed and delivered enterprise BI and reporting solutions supporting risk and operational leadership; built standardized analytical frameworks that replaced ad-hoc Excel-based reporting across multiple lines of business.
<b>Santander Bank</b>	Architected operational risk and internal controls reporting platforms; designed semantic models and dashboards that translated regulatory requirements into actionable performance metrics for senior leadership.
<b>Wells Fargo</b>	Delivered enterprise analytics solutions integrating KPIs and Key Risk Indicators (KRIs) into governance-grade dashboards; designed reporting frameworks supporting compliance, audit, and senior-leadership decision cadences.
<b>USAA</b>	Designed risk and compliance visualization frameworks supporting enterprise risk and compliance initiatives; built executive and operational dashboards that translated regulatory and control requirements into clear performance signals, enabling proactive issue identification and accelerating leadership decision cycles.
<b>PayPal</b>	Architected a Sarbanes-Oxley (SOX) compliance dashboard for the Internal Audit department, consolidating control testing status, deficiency tracking, and remediation progress into a unified governance-grade reporting layer; replaced fragmented audit tracking with a single semantic model supporting both auditor workflows and senior leadership oversight.
<b>DEFENSE</b>	
<b>Northrop Grumman</b>	Built business intelligence and process automation solutions in a defense-sector environment requiring strict governance, data lineage, and auditability; contributed to analytical frameworks supporting program-level operational oversight.
<b>HEALTHCARE</b>	
<b>CVS Health</b>	Architected cybersecurity risk metrics and reporting platforms; designed unified semantic models with executive and operational reporting canvases connected through drillthrough patterns, replacing fragmented reporting and accelerating audit-readiness.
<b>FOOD MANUFACTURING</b>	
<b>McCormick</b>	Automated the end-to-end data collection and reporting of product defects across the company's regional and global manufacturing footprint; designed custom visualizations surfacing defect patterns by plant, product line, and root cause, replacing manual plant-by-plant reporting with a unified analytical platform supporting quality leadership and operational decision-making at scale.

## 03 · TECHNICAL ARCHITECTURE STACK

## Capabilities by Architectural Layer

The portfolio site documents 45 tools across the analytics lifecycle. The grouping below reflects how I think about that stack — by **architectural layer**, not as a flat skills list. Within each layer, the tools listed are those I reach for most often. [Full hover-detail matrix on the site](#) →

### DATA SOURCES & STORAGE

*The foundation of every analytics solution — relational, lakehouse, and cloud sources connected through governed, predictable patterns.*

**Microsoft Fabric / OneLake** — Primary cloud foundation; Lakehouses, Dataflows Gen2, and semantic models live here.

**ADLS Gen2** — Connected via Fabric for hierarchically organized large-scale datasets.

**SQL Server** — Primary relational source; tight integration with the Microsoft stack.

**Dataverse** — Reach for it when business units need governed, role-secured, relational tables.

**PostgreSQL / MySQL** — Open-source environments encountered in leaner consulting engagements.

### INGESTION & TRANSFORMATION

*Where reliability is engineered. Transformation logic lives upstream so semantic models stay clean, performant, and maintainable.*

**Power Query (M)** — The transformation layer that determines whether a model is reliable, performant, and maintainable.

**Dataflows Gen2** — Reusable Power Query logic shared across multiple Power BI semantic models — maintain once, not per-model.

**KNIME** — Preferred open-source ETL platform; visual, node-based pipelines for transparent, repeatable transformation.

**Alteryx Designer** — High-value tool for complex prep and blending; reserved for engagements that license it.

**SQL / T-SQL** — Direct source interaction for reconciliation, schema validation, and ad-hoc verification.

## MODELING & GOVERNANCE

*The semantic layer is where business logic becomes trustworthy. Star schema by default; deviation requires justification.*

**Power BI Semantic Models** — The reusable analytical layer; relationships, measures, and structure consumed by many reports.

**DAX** — The calculation language; measures, time intelligence, filter context, AI-assisted authoring.

**Star Schema / Dimensional Modeling** — Default modeling pattern; clean fact–dimension separation, predictable cardinality.

**Row-Level Security (RLS)** — Role-based filter expressions for multi-audience reports.

**Governed Data Marts** — Curated subject-area layer between raw sources and consumption.

## ANALYSIS & VISUALIZATION

*The visible layer — but only as the last act of a coherent architectural sequence.*

**Power BI Desktop & Service** — Primary platform for semantic design, DAX, governance, and refresh orchestration.

**Tableau Desktop & Server** — Secondary platform; deployed when client ecosystem standardizes on Tableau.

**Power BI Report Builder / SSRS** — Pixel-perfect paginated reports for operational and compliance distributions.

**Advanced Excel** — Now primarily a source feeding Power BI; analytical work has moved to purpose-built tools.

## AUTOMATION & DELIVERY

*Closing the loop between insight and action. Where analytics becomes something a business user can actually use.*

**Power Automate** — Workflow layer connecting data, people, and processes; closes the loop from insight to action.

**Power Apps** — Low-code applications putting data and automation in the hands of business users.

**Microsoft Forms / SharePoint** — Lightweight intake and content layers that feed downstream automation.

**Visio (BPMN 2.0 certified)** — Process and architecture documentation for governance-conscious environments.

## AI & INTELLIGENT TOOLS

*An embedded layer across the analytics workflow — accelerating authoring, reasoning, and conversational data access.*

**Microsoft Copilot / Copilot Studio** — Embedded intelligence across M365 and Power Platform; agents for documentation and policy-aligned lookup.

**Microsoft Fabric Data Agent** — Conversational intelligence over semantic models; the future-state of analytics consumption.

**Anthropic Claude** — Primary AI assistant for architecture reasoning, documentation, and metric interpretation.

**OpenAI ChatGPT · Google Gemini / NotebookLM** — Versatile, ecosystem-specific assistants used per task fit.

**Generative Prompting** — Treated as a design discipline — the difference between using AI and leveraging it.

## 04 · SPECIALTIES &amp; DOMAIN STRENGTHS

## Six Areas of Concentrated Practice

Distilled from the keyword index on the live portfolio. These are the areas where my architectural depth is greatest and where I deliver the strongest strategic and operational impact.

Specialty	Scope
<b>Semantic Modeling &amp; DAX</b>	Star-schema-first model design, DAX optimization, filter context, time intelligence, performance tuning.
<b>Microsoft Fabric Architecture</b>	Lakehouse, Dataflows Gen2, OneLake, semantic model design, deployment pipelines, capacity planning.
<b>Governance &amp; Security (RLS / OLS)</b>	Role-based access, row- and object-level security, lineage documentation, governed self-service BI.
<b>Enterprise Risk &amp; KRI Frameworks</b>	KRI/KPI separation, compliance-grade dashboards, audit-ready reporting, control deficiency tracking.
<b>Process Automation via Power Platform</b>	Power Apps, Power Automate, Dataverse, SharePoint — end-to-end low-code solutions on existing licenses.
<b>AI-Augmented &amp; Conversational BI</b>	Fabric Data Agent, Copilot Studio agents, prompt engineering, AI-assisted DAX and documentation.

05 · VERIFIED CREDENTIALS

## Nine Certifications, Independently Auditable

My approach to certifications is deliberate. For most of my career, the work itself was the credential — the field had few formal programs aligned to senior practitioner depth. As the Microsoft data platform matured, I invested in formal credentialing as *a formalization of expertise built across 15+ years of practice*, not an introduction to it.

Credential	Issuer	Year	Verify
Microsoft Certified AI Transformation Leader (AB-731)	Microsoft	2026	Verify →
Microsoft Certified Power BI Data Analyst (PL-300)	Microsoft	2023	Verify →
Microsoft Certified Power Platform Fundamentals (PL-900)	Microsoft	2025	Verify →
Microsoft Office Specialist — Access Expert	Microsoft	2021	Verify →
Tableau Desktop Specialist	Tableau	2022	Verify →
Alteryx Designer Core	Alteryx	2023	Verify →
KNIME Analytics Platform — Basic Proficiency	KNIME	2023	Verify →
CompTIA Data+	CompTIA	2023	Verify →
OCBE — Certified Expert in Business Process Management	Object Management Group	2008	Verify →

**Upcoming milestones (target July 2026): DP-600 Fabric Analytics Engineer Associate · DP-700 Fabric Data Engineer Associate**

Full credential catalog with verification links →

## 06 · CONTINUOUS LEARNING SNAPSHOT

## Sustained Mastery

*In a rapidly shifting data landscape, static expertise is technical debt.*

The metrics below are not aspirational — they are tracked on a live training dashboard linked from the portfolio site. Continuous investment in specialized workshops and advanced training is how every architectural decision stays anchored to current enterprise standards rather than yesterday's.

<b>293</b> MODULES	<b>12</b> COURSES	<b>1</b> ACTIVE NOW	<b>104+</b> HOURS
-----------------------	----------------------	------------------------	----------------------

### Recent & Active Coursework

- Microsoft Fabric: Implement and Manage Analytics Solutions — Coursera / Whizlabs · 12h · Completed APR-2026
- Fabric Analyst In A Day — Microsoft / Pragmatic Works · 8h · Completed APR-2026
- Modern Data Architecture Fundamentals with Microsoft Fabric — Microsoft / Coursera · 15h · Est. completion MAY-2026

[Live training dashboard with full course history →](#)

## 07 · ARCHITECTURAL METHODOLOGY

## Three Case Studies in Applied Practice

These are sample case studies built on synthetic data and fictional organizations to protect professional confidentiality. The methodologies, architectural decisions, and stack choices reflect actual professional practice. Full case studies on the portfolio site →

### Enterprise Risk Management (ERM) Framework — Meridian Health Network

SQL Server · Alteryx · Power Query · Power BI · Data Governance

**PROBLEM** Six business units tracking compliance risk through fragmented Excel workbooks. No standardized KRI definitions, no governed semantic model, audit cycles consuming weeks of manual reconciliation.

**DECISION** Replaced the fragmented model with a single governed semantic layer, standardized KRI definitions, and real-time executive dashboards integrated to the underlying control deficiency data.

**OUTCOME** **Audit preparation time reduced 92%. Reporting latency eliminated by 45 days. Single source of truth across all six business units.**

[View full case study \(PDF\) →](#)

### Multi-Channel Marketing Performance Platform — Prism Retail Group

Microsoft Fabric · Lakehouse · RFM Modeling · Power BI · CLV Modeling

**PROBLEM** Multi-brand retail group operating with siloed channel reporting. No unified customer intelligence layer; CLV and segmentation logic inconsistent across brands.

**DECISION** Designed a unified Fabric-based customer intelligence platform with governed CLV and RFM segmentation models feeding executive Power BI dashboards.

**OUTCOME** **Marketing leadership reallocated \$14.2M in spend by customer lifetime value. 34% ROAS improvement. 2.1× CLV uplift in the Champions segment.**

[View full case study \(PDF\) →](#)

## Capital Allocation & Expenditure Oversight — Vantage Industrial Partners

Power Apps · Dataverse · Power Automate · SharePoint · Power BI

**PROBLEM** \$38M in annual capital spend tracked through a single-analyst Excel process. A two-year-old Internal Audit material weakness finding remained open. No budget for a proper back-end system.

**DECISION** Engineered a governance-grade oversight platform entirely on existing Microsoft 365 licenses — Power Apps front end, Dataverse back end, Power Automate for approvals, SharePoint integration, Power BI executive reporting.

**OUTCOME** **Delivered in five months at zero incremental technology cost. Closed the two-year audit material weakness. Replaced a fragile, single-point-of-failure process with a fully governed enterprise platform.**

[View full case study \(PDF\) →](#)

**Why these three:** *Each illustrates a distinct mode of architectural decision-making — governance under regulatory pressure (Meridian), modernization of a fragmented analytical estate (Prism), and architectural judgment under hard budget constraints (Vantage).*

### 08 · REASONING SAMPLES

## How I Think Through the Work

These summaries distill responses given in actual interview sessions. They are intentionally compressed — the goal is to communicate *how I reason, not what I conclude*. Full responses, including specific examples and architectural detail, are available on the portfolio site. [Read the complete Q&A →](#)

### DATA ACCURACY & INTEGRITY

#### **Describe a time your data or analysis was wrong. How did you catch it?**

Caught a date-normalization defect in a healthcare risk dashboard before formal presentation — a silent join behavior was creating an artificial improvement trend. Pulled the deliverable, corrected the model, added a data quality check to the pipeline so the same edge case would surface as a warning on future refreshes. The architectural lesson is that errors caught are an opportunity to harden the system permanently, not just to fix the immediate output.

### ARCHITECTURE VS. QUICK WINS

#### **How do you decide between a 'quick win' and a long-term architectural solution?**

The decision starts with one question: have we been here before? Most one-time urgent requests are recurring requests in disguise. The right move is a five-minute conversation with the requester explaining that the urgent need can be met today and the recurring need can be eliminated permanently with a small architectural investment. Most managers respond well because the framing

protects them from a problem they haven't seen yet — a fragile, manual, single-analyst process they will eventually own.

#### STAKEHOLDER CONFLICT RESOLUTION

##### **How do you navigate conflicting requirements from different business units?**

When operational analysts and senior leadership pulled in opposite directions on the same reporting surface, the path of least resistance was building two separate dashboards — and quietly creating two versions of the truth. Instead, I built one governed semantic model with two presentation canvases, connected by drillthroughs, drilldown hierarchies, and direct hyperlinks back to systems of record. One source of truth, two audiences served, no duplicated maintenance burden.

#### TECHNICAL IMPLEMENTATION

##### **What is your strategy for optimizing DAX performance in complex semantic models?**

Push transformation upstream wherever possible — DAX runs at query time, so anything left for it to figure out on the fly costs the user. Keep the model lean: clean star schema, simple relationships, no bidirectional unless necessary. Avoid calculated columns and heavy iterators. Use variables for readability and recalculation avoidance. When something feels slow, instrument it with Performance Analyzer or DAX Studio rather than guessing.

#### ADOPTION & CHANGE MANAGEMENT

##### **How do you ensure user adoption when rolling out a new BI platform?**

Adoption is a design problem, not a training problem. Start with how end users actually make decisions and build the dashboard backward from that. Bring data to users through alerts and subscriptions rather than expecting them to seek it out. Keep visuals clean and intentional. Identify power users early as internal champions. Short, practical training that connects directly to daily work — not long generic sessions.

## 09 · PROFESSIONAL ENDORSEMENTS

## Verified Peer Recognition

These endorsements are published verbatim on the portfolio site with attribution. They represent peers and managers who have observed the work directly.

*"Manuel is one of the sharpest technical minds I've worked with. Beyond his deep expertise, he has a rare ability to step back, connect the dots, and account for the bigger picture including the intangible factors that often get overlooked. If you need clarity on a complex problem, he's the person you go to."*

**Nisar Ahmed** · Senior Manager, Cybersecurity Metrics & Reporting · *CVS Health*

*"Manuel did amazing work and was consistently recognized as a top performer. I was especially impressed with his end-to-end solutions, his attention to detail, and his ability to present information in a way that clearly told the story to key stakeholders. If I ever have the opportunity to work with or hire him again, I would do so without hesitation."*

**Jad Mahnken** · Manager, TA Production & Onboarding · *Trinity Health*

## 10 · SELECTED WRITING &amp; THOUGHT LEADERSHIP

## Technical Guides & Best Practice Briefs

Selected pieces drawn from training workshops and reference libraries. The full catalog of articles is published on LinkedIn and indexed on the portfolio site.

### Technical / How-To

#### Creating a Dynamic RANKX DAX Measure in Power BI

A walkthrough of dynamic ranking patterns in DAX — context-sensitive RANKX behavior, filter context preservation, and the design choices that determine whether ranking measures behave predictably across slicer combinations.

[Read on LinkedIn](#) →

### Best Practices

**Best Practices for Requirement Gathering in Data Analysis and Visualization**

Framework-level guidance on requirement gathering — the institutional knowledge that separates mature analytics teams from those perpetually firefighting technical debt. Drawn from years of stakeholder engagements where the cost of skipped discovery compounds across every downstream decision.

[Read on LinkedIn →](#)

## 11 · ENGAGEMENT &amp; CONTACT

## How to Reach Me

Open to opportunities where strategic thinking, technical rigor, and operational impact converge. Specifically:

- **Full-time:** Senior BI Solutions Architect roles — ideally Microsoft Fabric / Power Platform-centric, with semantic modeling and governance scope.
- **Contract:** Architecture, semantic modeling, governance design, and BI platform modernization engagements.
- **Advisory:** Strategic BI direction-setting, AI transformation roadmaps, and senior-leadership-facing analytics oversight.

**DIRECT CONTACT**

Email [manuel@munozdataworks.com](mailto:manuel@munozdataworks.com)

Portfolio [www.munozdataworks.com](http://www.munozdataworks.com)

LinkedIn [in/manuel-munoz-jr](https://www.linkedin.com/in/manuel-munoz-jr)

Resume (PDF) [Download from the portfolio site](#)

**A note on this brief:** *This document is a curated companion to the live portfolio. It is not a substitute for the site or the formal CV — each artifact serves a different audience and a different decision moment. If anything in this brief raises a question or invites a deeper look, the portfolio is built precisely for that exploration.*