

TRISTAN CHIAPPISI

Principal Data Architect · Columbus, Ohio

614-657-0161 · tristan@tristanchiappisi.com · [linkedin.com/in/tristanchiappisi](https://www.linkedin.com/in/tristanchiappisi) · tristanchiappisi.com

Twenty-five years in IT, fifteen of them in data. Five employers. One patent. The work, in order, with the wins folded into the prose.

25 YEARS IN IT **15** YEARS IN DATA **5** EMPLOYERS **1** PATENT

CURRENT · 2021 → PRESENT

Where I work now and what the job actually is. Data architecture, engineering effectiveness, and AI at every layer of the work.

Principal Data Architect

at **Cetera Financial Group**

2021 → CURRENT · COLUMBUS, OHIO · FINANCIAL SERVICES

I lead data architecture for the customer-facing systems. Client portals, AWS managed services, Postgres. The work is partly design (what the data layer looks like, how it scales, how AI plugs in) and partly getting engineering teams faster without losing the audit trail.

AI is part of most of the work now. The engineering loop: code review, system design, test generation, documentation. The data layer: profiling, validation, enrichment. The strategy layer: transformation, roadmap, governance. The kind that compounds.

The platform I am proudest of is the data ingestion, governance, and processing platform on AWS managed services. Company-wide. Above that, the AI-enhanced data-conversion platform that runs on every M&A event. The previous baseline was months of manual rework per acquisition. The new baseline is days. Below both: the work of standardizing engineering practices around AI-assisted tooling. Documented patterns, CI checks that compound across teams, and a senior-advisor role on responsible AI adoption in a regulated environment. Less visible. More durable.

PRIOR · 2001 → 2021

Nine years at JP Morgan Chase, one year embedded at Microsoft, eleven years before that across a regional ISP and a computer-services consultancy where the data passion started. In reverse chronology.

Lead Software Developer · Architect · Data Scientist

at **JP Morgan Chase**

2015 → 2021 · COLUMBUS, OHIO · 6 YEARS

Six years at Chase in the data-science and analytics space. The title kept expanding. I started as Lead Software Developer and ended as Lead Architect and Data Scientist on the same team. The underlying problem stayed the same. Large-scale analytics, data and aggregations, and the question of how an organization keeps its data from sprawling sideways into duplicates.

The pieces that mattered most: a real-time machine-learning classification system (Python, scikit-learn, Flask, Kubernetes) that automatically routed raw text messages from an application supporting more than a hundred thousand devices. A complex event processor for logs that dispatched to vendors without a human in the loop. A state-based KPI system on branch entities that gave the business real-time analytics where it used to wait for nightly reports. A PowerShell ETL framework that ended up running more than a hundred packages in production. The first prediction engines I built end-to-end were for application-failure prevention.

The standardization work was at least as load-bearing as the products. Moved the org from a half-dozen ETL approaches onto SSIS, PETL, and a documented home-grown stack. Led the data-visualization standardization onto Power BI. Trained the team (and their managers) on Tableau before that.

Premier Field Engineer

at **Microsoft**

2014 → 2015 · SQL SERVER PRACTICE · 4 ENTERPRISE CUSTOMERS

Embedded with four enterprise SQL Server customers as Microsoft's on-site senior engineer. The job was triage and architecture, in that order. The immediate priority issues. The migration plans. The high-availability design across data centers. The training that meant the next priority issue went to one of theirs instead of me.

The largest piece was a multi-database, multi-terabyte SQL Server migration coordinated across the customer's staff, mine, and the calendar. Shipped with minimal downtime. The clustered VLDB high-availability and disaster-recovery work spanned multiple data centers and pulled the operations teams in on the design rather than after.

SQL Server Database Administrator

at **JP Morgan Chase**

2012 → 2014 · COLUMBUS, OHIO · CHASE.COM ALERTS

First job at Chase. Primary database developer for the Chase.com Alerts system, built on Service Broker (the queue layer) and SQLCLR (C# functions running inside SQL Server). The largest schema crossed a billion rows.

Distribution was handled with database partitioning, Service Broker, and clustered views across multiple servers.

In parallel, I supported more than a hundred SQL Server instances and a couple hundred databases as a working DBA. Built the cross-technology failover coordination system that handled DB2, MSSQL, and WebSphere as one orchestration unit during data-center cutovers. Designed the ETL approach across SSIS, PowerShell, and SQLCMD.

Before Chase

Computer-services consultancy, then a regional ISP

2001 → 2012 · 11 YEARS · WHERE THE DATA PASSION STARTED

An internship at a computer-services consultancy out of school. Then a regional ISP for most of the rest of those years. General-purpose IT was the day job. The data problems were what I gravitated to, and that is where the data career started.

STACK · WHAT I REACH FOR

Languages I still write, platforms I have shipped on, integration tools that have earned their place. Honest about which ones I live in and which ones I rent.

Skills

ARTIFICIAL INTELLIGENCE Applied work across engineering (Claude Code, MCP, tool-use, RAG, embeddings), data (AI-assisted profiling, validation, enrichment), and strategy (roadmap, governance, responsible-adoption framing for regulated environments).

LANGUAGES Python for AI/ML pipelines and automation. T-SQL for everything that hits a relational store. PowerShell when the shell is Windows. C# when SQLCLR is the right tool. Java when I have had to.

PLATFORMS MSSQL has been the longest relationship. Postgres and AWS are where most current work lives. Oracle, MariaDB, Snowflake, and Cassandra each shipped in production at various points. CloudFoundry was the Chase years. Splunk for the logs you cannot lose.

INTEGRATION AWS Glue for managed-services work. SSIS where the org already lives in it. Informatica when there is a legacy to respect. Event-driven where the pattern fits.

VISUALIZATION Power BI in the enterprise. Tableau where it is in place. Matplotlib and Plotly when the chart is the argument.

ALSO · OUTSIDE THE ROLE

A patent, a co-authored career-track workshop series, and the work that does not fit on a job description.

Patent · Distributed analytics storage engine

FILED DURING THE CHASE YEARS · PATENT RECORD TO BE LINKED

The invention is distribution at the storage layer with the consistency model analytics work actually needs. The full record is on file and will be linked here when I get the citation cleaned up.

Career-track workshops · Co-authored with Natalie Hilton

2023 → 2026 · CODEMASH · 5 SESSIONS · 240-MIN INTERACTIVE

Four years of co-authored career-track sessions at CodeMash. Communication, mindset, career conversations, working with others. Each session is a four-hour interactive workshop with **Natalie Hilton** as lead author and me as co-author. The full archive lives on the [Speaking](#) page.

The reason this is on a resume page rather than a side note: running these sessions is load-bearing for how I lead engineering teams day-to-day. The EQ vocabulary and the structured career conversations are the same practice in two settings.

Technical talks · Solo

2024 → 2026 · CODEMASH · STIR TREK · SQL TUNING, SCALING, QUANTUM

Solo technical talks on SQL performance tuning, scaling beyond relational stores, and an introduction to quantum computing for developers. *Caffeinate Your Queries* has been the recurring one. Three deliveries so far, more on the calendar. Full archive at [Speaking](#).

Consulting · AI transformation and data strategy

AVAILABLE · BY REFERRAL · REGULATED-INDUSTRY DATA ARCHITECTURE

Advisory work on AI transformation, data strategy, and the architectural decisions that connect the two.

Engagements are short and scoped: a quarter or two to deliver a 12-month AI roadmap, an AI playbook for each department and role so the team has a standard way of working with the tools, and a build/buy/defer architecture radar. Then hand back to the in-house team. By referral.