

# Christopher Nielson

Portfolio: [cjnielson.com](http://cjnielson.com)

LinkedIn: [linkedin.com/in/christopherjnielson](https://www.linkedin.com/in/christopherjnielson)

GitHub: [github.com/chrisjnielson44](https://github.com/chrisjnielson44)

Pittsburgh, PA

Mobile: +1 239-307-8146

Email: [cjnielson44@gmail.com](mailto:cjnielson44@gmail.com)

---

## SUMMARY

Software Engineer building infrastructure for production AI agents and agentic workflows. Lead development of BNY Mellon's Enterprise Risk MCP Server, a secure orchestration platform that enables LLM agents to retrieve enterprise data, invoke analytical tools, execute governed workflows, and return structured outputs across critical financial-risk systems. Experienced designing Python/FastAPI services, TypeScript platforms, containerized AI services, evaluation workflows, observability patterns, RBAC, audit logging, and human-in-the-loop controls. Strong track record debugging complex agent failures across retrieval, tool selection, prompts, permissions, service integrations, logs, and execution traces in regulated production environments.

---

## EXPERIENCE

- Bank of New York Mellon**, Software Engineer — Risk Engineering Pittsburgh, PA — **June 2024–Present**
- **Agent Infrastructure & Tool Execution Platform:** Lead development of BNY Mellon's Enterprise Risk MCP Server, a centralized agent orchestration layer that enables LLM agents to securely retrieve enterprise financial-risk data, invoke analytical tools, and return structured outputs. Built authentication, entitlements, RBAC, audit logging, schema validation, and governed tool interfaces for traceable production workflows.
  - **Production Backend APIs for Agentic Systems:** Design and build Python FastAPI services and TypeScript/Next.js platform interfaces that expose risk data, tool execution, and agent workflows to internal users and downstream systems. Define API contracts, structured payloads, error handling, and service integration patterns across PostgreSQL, Snowflake, and enterprise data services.
  - **Agent Reliability, Observability & Evaluation:** Debug production AI workflows using application logs, execution traces, agent histories, and service telemetry to identify failures in retrieval, prompt behavior, tool selection, permissions, and backend integrations. Build evaluation scenarios and monitoring patterns to improve reliability, explainability, and safe human-reviewed execution.
  - **Containerized AI Platform Engineering:** Deploy and operate containerized AI services using Docker, Kubernetes, CI/CD pipelines, Linux environments, and Nvidia GPU infrastructure. Contribute to deployment patterns, anomaly detection, event-driven reporting, and operational controls for high-stakes financial-risk workflows.
  - **Multi-Agent Orchestration & Context Systems:** Develop LangGraph-based agents that coordinate retrieval, tool calls, structured generation, and human-in-the-loop review for financial reporting and risk analysis. Design GraphRAG and knowledge graph patterns that give agents precise, permission-aware context across Credit, Market, Counterparty Credit, Treasury, Model, Operational, and Technology Risk.
  - **0-to-1 Applied AI Systems:** Built LLM-powered backend tools for counterparty risk, including a system comparing daily Potential Future Exposure values across FX risk engines and a natural-language-to-SQL agent that inspected schemas, generated queries, and returned real-time insights. Applied multi-LLM evaluation, market data, trade data, and structured validation to improve consistency and explainability.

---

## SKILLS SUMMARY

- **AI / ML:** OpenAI SDK, Claude Agent SDK, Langfuse, Scikit-learn, NLP, Multi-Agent Workflows, Model/Agent Debugging
- **Agent Infrastructure:** MCP Server Development, Tool Execution, Tool Calling, Agent Orchestration, LangGraph, Structured Outputs, Schema Validation, Agent Evaluation, Failure Analysis, Human-in-the-Loop Controls
- **Backend / APIs:** FastAPI, REST APIs, Microservices, Node.js, Next.js, React, API Contracts, Service Boundaries, Error Handling
- **Infrastructure / Observability:** Docker, Kubernetes, Linux, GitHub/GitLab CI/CD, Nvidia CUDA/GPU Infrastructure, Cloud Platforms, Monitoring, Production Logs, Service Telemetry
- **Data / Retrieval:** PostgreSQL, Snowflake, MongoDB, Oracle SQL, MSSQL, Large Financial Datasets, RAG, GraphRAG, Knowledge Graphs, Retrieval Quality
- **Languages:** Python, TypeScript, SQL, JavaScript, Java, C++, C

---

## EDUCATION

**Carnegie Mellon University, School of Computer Science**

*Graduate Certificate, Machine Learning & Data Science*

Pittsburgh, Pennsylvania

*May 2026*

**Florida State University, College of Arts and Sciences**

*Bachelor of Science, Computer Science — Minor in Business & Mathematics*

Tallahassee, Florida

*December 2024*

---

## LANGUAGES

- **Spanish:** Fluent — Read, Write, Speak

---

## HONORS AWARDS

- **Carnegie Mellon University Graduate Certificate Scholarship** 2025 – 2026
- **Florida Bright Futures Academic Scholar — Full Tuition** August 2021 – December 2024