

Noelle Milton Vega
Transforming Legacy Business Processes to
Strategic Cognitive Systems with
Agentic Generative AI
CV and Portfolio: <https://jupyter.ai>
LinkedIn & Contact: <https://linkedin.com/in/nmvega> (not up to date)

COVER LETTER

Hello:

Thank you in advance for your time. As you review the enclosed resume, I invite you to explore the portfolio it references at <https://jupyter.ai>, where you will find an up-to-date copy of my resume along with select examples of my career's body-of-work. The portfolio includes real-world projects, architectural diagrams, letters of recommendation, whitepapers, certifications, tutorials, and more. I believe these materials will offer valuable insight into my experience and capabilities.

Early on I observed numerous examples of great leadership, and what distinguished the truly successful from the rest was a higher level of contribution towards an organization's most important goals. I fashioned my career after those observations, and my body-of-work has grown to include consultancy, advisory, entrepreneurship and key driver-roles within the heart of many businesses ranging from global Fortune-500 firms to startups.

You may recognize this through my full-cycle collaboration with professionals along the organizational ladder – advising executives on business-value at one end, guiding technology personnel at the other, and interfacing with key stakeholders in between.

I have strategically helped shape business directions towards more sustainable and certain futures • Identified solutions that streamline day-to-day operations • Mentored and guided teams towards successful finish lines • Driven consensus across functional organizations • And hands-on conceived, architected and implemented technology platforms that made enterprise-critical goals lasting realities.

By faithfully applying time-tested best-practices and playbooks, I've been able to replicate those outcomes across disparate enterprises and verticals, including Financial Services, Pharmaceuticals, Telecommunications and Cable-TV Media & Entertainment. Though I remain technical, these days I'm equally sought because I understand the impact that strategic-technologies have on businesses ability to stand tall during these times.

I'd be delighted to discuss your initiatives and review ways in which my technical and business-technical experiences can be of value. I have included an abbreviated variation of my résumé to provide you with additional background that can form the basis for our discussion.

Sincerely,
Noelle Milton Vega

Noelle Milton Vega
Transforming Legacy Business Processes to
Strategic Cognitive Systems with
Agentic Generative AI
CV and Portfolio: <https://jupyter.ai>
LinkedIn & Contact: <https://linkedin.com/in/nmvega> (not up to date)

SUMMARY COVER

Masters in Electrical & Computer Engineering from Boston University and a Bachelor's in Electrical Engineering from Rensselaer Polytechnic Institute (RPI). My career is defined by cross-industry success as a principal architect, engineer, and strategic advisor to Fortune 500 biggest brands, leading high-value, enterprise-critical projects. Explore select highlights of my work - including real-world architectures, whitepapers, commendations, certifications, and more - at my portfolio / CV site: <https://jupyter.ai>

OBJECTIVE

Engaging organizations in full-cycle capacities to implement or transform legacy business-data architectures, strategies, and governance into **modern reasoning systems** with **Agentic Generative AI**, driving sustained competitive advantage.

EDUCATION

- **Boston University (BU)** -- M.S. Electrical & Computer Engineering (Boston, Massachusetts)
- **Rensselaer Polytechnic Institute (RPI)** -- B.S. Electrical Engineering (Troy, New York)
- **George Westinghouse High School** -- High School Valedictorian (Brooklyn, N.Y.)

INDUSTRY VERTICALS

• Financial Services	• Cable, Media & Entertainment	• Pharmaceuticals
• Telecommunications	• Stock & Commodities Exchanges	• Defense and R&D Labs

BODY-OF-WORK PORTFOLIO SAMPLES

- PivotMath Math Platform – Designed & Coded a guided math practice for Middle Schoolers. [Video](#)
- <https://jupyter.ai> :: Portfolio Website Landing Page
- <https://jupyter.ai/RndLab> :: R&D Sandbox with servers and full-stack technologies to prototype customer solutions
- <https://verilabs.jupyter.ai> :: Cloud-in-a-Box SaaS platform to Build & Prototype against full-stack IT environments
- <https://jupyter.ai/deloitte> :: Full-Cycle Letter-of-Commendation from **Deloitte** for a Mission Critical Engagement
- <https://jupyter.ai/bofa> :: Architecture for a Realtime Trade-Anomaly-Analytics CEP Platform for **Bank of America**
- <https://jupyter.ai/nymex> :: Multi-Datcenter SAN for Realtime BC/DR for **The New York Mercantile Exchange**
- <https://jupyter.ai/netapp> :: Monitoring & Alerting Product Developed for **Network Appliance (NetApp)**
- <https://jupyter.ai/ubs> :: In-Memory Reference Data MDM Lookup System for **United Bank of Switzerland (UBS)**
- <https://jupyter.ai/data-pipeline> :: Full-Stack Data Pipeline prototype for a client with Hive, Spark & Kafka

CERTIFICATION LINKS

- [Coursera: Neural Networks & Deep Learning | Andrew Ng](#)
- [Coursera: Improving Deep Neural Networks: Hyperparameters, Regularization & Optimization | Andrew Ng](#)
- [Coursera: Structuring Machine Learning Projects | Andrew Ng](#)
- [Amazon Web Services \(AWS\): Certified Solutions Architect – Associate \(CSA-A\)](#)
- [Amazon Web Services \(AWS\): Certified Developer -- Associate \(CD-A\)](#)
- [Minority Owned Business Enterprise Certifications - \(City of NY; State of NY NMSDC\)](#)

DATA ARCHITECTURE TECHNOLOGIES: TOOLS AND PLATFORMS DEVELOPED OR USED OVER TIME:

AGENTIC GENERATIVE AI ECOSYSTEM: (N8N • MULTI-MODAL MODELS • OLLAMA • GPT4ALL • HUGGINGFACE • QDRANT • PINECONE • SUPABASE • OPENAI API • RAG • ETC.) • AI ML OPS • DOCKER • JUPYTERLAB • VSCODE • PYTHON • PYTORCH • SCIKIT-LEARN • JAVA • AWS • SPARK • LINUX (FEDORA) • BUSINESS CONTINUITY & DISASTER RECOVERY • ELT • NOSQL • FIBRE-CHANNEL WAN / SANs

Noelle Milton Vega
Transforming Legacy Business Processes to
Strategic Cognitive Systems with
Agentic Generative AI
CV and Portfolio: <https://jupyter.ai>
LinkedIn & Contact: <https://linkedin.com/in/nmvega> (not up to date)

EXPERIENCES

JUPYTER, LLC CONSULTANCY PRESENT

Consultancy focused enabling business-critical technology platforms and architectures with Agentic Generative AI. Except where noted, engagements are / were performed via my JUPYTER, LLC consultancy.

FLUR.EE – AGENTIC GENERATIVE AI ARCHITECT JULY 2024 – PRESENT (3RD PARTY CONSULTANT)

For a Fluree financial client, I developed a LangGraph-powered document compliance analysis agent that analyzed thousands of in-force vendor and client communications - including flagging deprecated terms, signatures from departed executives, and even outdated logos. The solution leveraged LangChain for supplemental capabilities (e.g., chat templates, ChatOllama, etc.), Postgres with pgvector to store relational data as well as to semantically index reference documents against which target documents were compared. Pydantic models mirroring the Postgres schema were programmatically generated (via a db introspection script) and injected into the LLM system prompt, enabling it to generate SQL queries from natural language user queries. This solution replaced manual offshore third-party review teams, cutting project costs by 63% while eliminating human error and significantly decreasing time-to-completion. It also established a reliable audit trail and a reusable system for future operations.

TIAA BANK | TIAA FINANCIAL SERVICES | TIAA CORPORATE – DIGITAL TRANSFORMATION ARCHITECT JANUARY 2023 – JULY 2024 (3RD PARTY CONSULTANT)

Led a dual-phase project that transitioned an enterprise-critical function from a spreadsheet-driven process to an online platform enabling collaborative projects, workflows, tasks, approvals, messaging, custom forms, lineage tracking, and other transformative benefits. **Phase 1 (Assessment):** Gathered and analyzed business & functional requirements, and dataflows to identify potential solutions; authored user stories for POCs; and selected a solution based on multi-dimensional criteria—technical, functional, cost, time-to-benefit, usability, and low-code/no-code capabilities. **Phase 2 (Implementation):** Collaborated closely with the vendor (Adobe) to master and implement Adobe Workfront; deployed business user stories; documented end-to-end processes; and trained business users, successfully transitioning the critical function from spreadsheets to Workfront's collaborative platform.

ZETTALABS, INC (now FLUR.EE) – DIRECTOR, PRODUCT DATA ARCHITECTURE & PLATFORMS NOVEMBER 2019 – SEPTEMBER 2022 (3RD PARTY CONSULTANT)

At ZettaLabs, I contributed to the flagship product, ZettaSense, a suite of tools including **Z/S Ingest** (data logistics), **Z/S Classify** (discovery and classification of raw data into business data domains), and **Z/S Resolve** (mastering, merging, and deduplication of data). In this small, niche company serving Financial Services and Pharma, I pivoted between three distinct roles: (1) Designing cloud data architectures and platforms for Z/S deployment, (2) Providing professional integration for Z/S customers, and (3) Developing PySpark and Kafka solutions for various Z/S stack components.

SCOTIABANK (BANK OF NOVA SCOTIA) – DIRECTOR DATA ANALYTICS, STANDARDS & PRACTICES APRIL 2019 – OCTOBER 2019 (3RD PARTY CONSULTANT)

Reporting directly to the global Chief Data Officer within the Data Governance Office, I lead a team focused on a Data Lake platform dedicated to data governance transformation initiatives. This platform serves as a single source of truth for Customer and Transaction data across various bank lines of business (LOBs). Through user stories, LOB data is coalesced, mastered, merged, and quality-monitored. Technologies used include: Cloudera Hadoop Data Lake, Database CDC with Attunity/Qlik Replicate & Compose, Sqoop, Apache Kafka for data logistics (including database CDC) and Data Lake ingest, PySpark/Python (and libraries) atop Spark 2.x for batch and streaming analytics, HiveQL, SparkSQL, Jupyter Notebook, and more. Programmatic interaction with the final platform is

Noelle Milton Vega
Transforming Legacy Business Processes to
Strategic Cognitive Systems with
Agentic Generative AI

CV and Portfolio: <https://jupyter.ai>

LinkedIn & Contact: <https://linkedin.com/in/nmvega> (not up to date)

similar to the tutorial I published here -- <https://jupyter.ai/data-pipeline> -- which is built atop my personal R&D Lab here: <https://jupyter.ai/rndlab/> (also: <https://verilabs.jupyter.ai>)

BRITISH TELECOM – DATA STREAMING AND ANALYTICS
JUNE 2018 – MARCH 2019 (3RD PARTY CONSULTANT)

A restrictive NDA and project sensitivity prevent divulging business-level detail; however, I lead a project as well as wrote Python 3.x code to drive a Kafka 2.x, Spark 2.x, Spark MLlib and MongoDB 4.x real-time analytics stack to identify correlations between transactions

MODERNIZING MEDICINE – RDBMS STREAMING CHANGE DATA CAPTURE (CDC) POC ON AWS
NOVEMBER 2017 – MARCH 2018 (3RD PARTY CONSULTANT)

Designed and implemented a Change Data Capture (CDC) proof-of-concept to generate real-time change-data event streams from updates to multiple MySQL database shards. The architecture included: (1) Kafka Connect connectors extracting data from MySQL bin logs and producing change-data events to Kafka topics, deployed as Docker containers in AWS ECS; (2) A Kafka platform powered by Confluent's managed Kafka service atop AWS; (3) A custom Python program to consume Kafka data, serialize it as Apache Parquet objects, and store them in AWS S3 for efficient querying by business intelligence teams via Apache Impala. The objective of this new platform was to: (1) Eliminate an 8-hour AWS EMR batch job with a continuous data stream for near real-time data updates, and to (2) Implement a low-touch, operator-free solution using a managed-services and quasi-serverless architecture. The architecture was eventually adopted and operationalized.

SYNERGY – AWS ARCHITECTURE FOR REPORTS JUNE 2016 – (CONSULTANT. ONGOING VIA RETAINER)

Designed an AWS-based capture and reporting platform to track member engagement and ROI. Key attributes, including member visit dates, check-in/check-out times, membership levels, and facility feature utilization (e.g., trainers, classes), are added or updated in DynamoDB through AWS API Gateway calls, triggering IAM-role enabled AWS Lambda functions. Local and Global Secondary Indices are utilized to optimize query performance. Analytical results are fetched programmatically using boto3 Python scripts or by scheduling CloudWatch Events (cron jobs) to trigger Lambda-driven DynamoDB queries, which are saved as CSV files in AWS S3. Completion of each query triggers an AWS SNS notification, emailing a business-owner with a restricted S3 URL to the generated CSV.

UNITED BANK OF SWITZERLAND (UBS) - BUSINESS PROCESS ACCELERATION AND S/W DEVELOPMENT
JANUARY 2015 – MARCH 2016 (3RD PARTY CONSULTANT)

Engaged by United Bank of Switzerland's Evidence Lab for a two-phased project. Phase 1 (audit): Involved an in-depth audit of business processes and IT capabilities in New York and London to identify tools that could accelerate financial research analyst activities. The 152-page audit report guided the development of a solution comprising a data transformation pipeline, a mini-Master Data Management (MDM) system, and a taxonomy-centric data catalog. Phase 2 (delivery): Involved the development & delivery of this solution using pure Python-3 and Oracle. The platform enabled the business to ingest arbitrary web-derived data, cleanse, normalize, and enrich it, and store the final result in a vetted data store for BI tools. The solution is modeled after this deliverable: <https://jupyter.ai/ubs>.

BANK OF AMERICA – TRADE STREAM ANOMALY & FRAUD DETECTION
FEBRUARY 2014 – NOVEMBER 2014 (3RD PARTY CONSULTANT)

Engaged by Bank of America's trading division to assess requirements for replacing an aging FIX trade analytics system. The new system needed to identify wash sales, order-to-execution latencies, account limit breaches, and other trade anomalies in real-time. Worked with LOBs to secure buy-in and funding for the Complex Event Processing solution I proposed, which leveraged open-source big data technologies including Apache Storm, Apache Cassandra, and Redis. After gaining consensus, I successfully designed, built, and delivered the system, capable of processing 12 million events per second. This high-visibility open-source project also demonstrated the viability of a "BUILD" over "BUY" approach bank-wide. A detailed PDF of the R&D conducted during the final design phase is available here: <https://jupyter.ai/bofa>

Noelle Milton Vega
Transforming Legacy Business Processes to
Strategic Cognitive Systems with
Agentic Generative AI

CV and Portfolio: <https://jupyter.ai>

LinkedIn & Contact: <https://linkedin.com/in/nmvega> (not up to date)

Noelle Milton Vega
Transforming Legacy Business Processes to
Strategic Cognitive Systems with
Agentic Generative AI

CV and Portfolio: <https://jupyter.ai>

LinkedIn & Contact: <https://linkedin.com/in/nmvega> (not up to date)

MERRILL LYNCH – ARCHITECT (RECOMMENDER SYSTEM)
MARCH 2013 – FEBRUARY 2014 (3RD PARTY CONSULTANT)

Engaged to design a recommendation system for wealth clients browsing Merrill Lynch's financial research portal, similar to Amazon's "You may also like" feature. The system leveraged a combination of Item and User-Similarity collaborative filtering algorithms, along with K-Means clustering, to recommend investment research publications from a database of hundreds of thousands of items. Inputs included client demographics, investment objectives, and publication attributes such as title, sector, industry, author, and age of publication. The solution was implemented using Apache Mahout atop Hadoop, alongside Python-3 and Scikit-Learn. A Talend pipeline facilitated the shuttling of raw input data and computed results between HDFS and the Enterprise Data Warehouse.

SELECT LEGACY EXPERIENCES (PREVIOUS ONES OMITTED FOR RESUME BREVITY)

REUTERS --> NASDAQ – ARCHITECT: ECN EXCHANGE | SANs | DATACENTER & BC/DR MIGRATION
JUNE 2004 – 2006 (EMPLOYEE – REUTERS, THEN NASDAQ, POST-ACQUISITION)

INET was an electronic trading platform (ECN) developed and autonomously operated within REUTERS, executing orders at sub-millisecond speeds and handling billions of dollars in daily trade volume. As the then world's fastest exchange, it was acquired by NASDAQ in 2005 to replace its aging exchange infrastructure, and today it is not only implemented at NASDAQ but also at other exchanges globally. Invited by dot-com era peers, including the then-CTO of INET and NASDAQ, I held multiple pre- and post-acquisition roles, including: **(1)** Writing C-based code to control a headless compute grid of embedded Linux systems, including the orders and execution matching engine; **(2)** Performance engineering to implement end-to-end tests and optimizations to enhance the ECN platform's speed; **(3)** Designing SAN Storage architecture to capture real-time quotes, orders, and execution transactions; **(4)** Leading the post-acquisition datacenter buildout and migration of INET's infrastructure from Exchange Place, NJ, to Ashburn, VA, including managing the critical stock exchange cutover during a 3-day weekend.

NYMEX / CME GROUP – BUSINESS CONTINUITY / DISASTER RECOVERY DATACENTER ARCHITECT
SEPTEMBER 2002 – APRIL 2004 (3RD PARTY CONSULTANT)

The New York Mercantile Exchange (NYMEX) engaged me to design and implement a post-9/11 Business Continuity and Disaster Recovery (BC/DR) datacenter platform, along with developing disaster response policies and procedures. The final platform featured a Storage Area Network (WAN / SAN) connecting primary and backup datacenters, enabling real-time synchronous replication of commodities trades and mission-critical database transactions between them. This solution, along with its processes and procedures, mitigated operational, reputational, financial, and legal risks. Success was measured using key performance indicators (KPIs), including Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO), incorporated into bi-annual failover tests, which I authored. Full-cycle responsibilities included: Presenting solutions to CEO, CIO, and CTO to gain consensus • Justifying CAPEX and OPEX (ROI/TCO) to plan and secure a 3-year budget from the CFO • Conducting infrastructure capacity planning and right-sizing, including a pay-as-you-grow procurement model • Sourcing vendors and evaluating product fit • Collaborating with facilities personnel and trade unions to prepare and install incoming infrastructure • Architecting, deploying, and documenting the inter-datacenter SAN platform (as seen in the linked design) • Documenting policies and procedures for failover scenarios. These efforts were pivotal in helping NYMEX recover from disasters. A draft of the final platform design can be viewed here: <https://jupyter.ai/nymex>

Noelle Milton Vega
Transforming Legacy Business Processes to
Strategic Cognitive Systems with
Agentic Generative AI
CV and Portfolio: <https://jupyter.ai>
LinkedIn & Contact: <https://linkedin.com/in/nmvega> (not up to date)

CABLEVISION SYSTEMS – DATACENTER PLATFORM ARCHITECT (CABLETV MASTER HEADEND)
MARCH 2001 – JUNE 2002 (3RD PARTY CONSULTANT)

Engaged in all aspects of the design, development, and deployment of Cablevision's inaugural Optimum.TV systems, the nation's first Video-On-Demand Cable TV product. Activities took place at the Master Headend datacenter and included: Designing storage and streaming systems for on-demand ordering and consumption of digital content • Deploying systems to enable at-home self-provisioning of new cable boxes • Designing highly available database clusters handling customer accounts, digital orders, and subscriptions • Integration with the Conditional Access System (CAS) to control subscriber access to programming via encryption • Building a Fibre Channel WAN / SAN for database and digital content replication to a backup datacenter • And collaborating with Madison Square Garden (owned by Cablevision) to offer live and on-demand MSG event programming. These efforts were pivotal in launching the nation's first VoD cable TV product.

DELOITTE & TOUCHE – APPLICATION PLATFORM ARCHITECT; TEAM BUILDER & MANAGER
JUNE 1997 – FEBRUARY 1998 (INDEPENDENT SUB-CONTRACTOR)

As primary consultant, Deloitte & Touche was responsible for building a PeopleSoft / Oracle application that would service Adecco Group's (formerly Olsten) 1,100+ staffing offices globally. As subcontractor to Deloitte, I was retained to architect and deploy the enterprise backend compute, storage, SAN and database farms that would be the foundation of that application. Responsibilities included: Working with stakeholders to obtain capacity growth projections for right-sized, pay-as-you-grow procurement; Building the SUN, EMC, SAN, Veritas, Oracle-based compute and storage farms and documenting them; Building and training the permanent systems team that would operate the farms once I completed the project. A summary of this full-cycle work was captured by the Deloitte Managing Director in his letter-of-commendation: <https://jupyter.ai/deloitte>

Noelle Milton Vega
Transforming Legacy Business Processes to
Strategic Cognitive Systems with
Agentic Generative AI
CV and Portfolio: <https://jupyter.ai>
LinkedIn & Contact: <https://linkedin.com/in/nmvega> (not up to date)

GENERAL CAREER PROGRESSION

- My career began as a Sonar Communications Engineer for G.E. Aerospace, where I contributed to a secret clearance defense project. This was followed by a role as a Digital Signal Processing Engineer at Sarnoff/RCA Labs in Princeton, NJ. There, I participated in R&D efforts for DirecTV as part of my Applied Master's Degree thesis.
 - Globalization and macroeconomic shifts prompted a pivot towards enterprise technology architecture, where I expanded my expertise across Information Technology, Strategic Business-Technology, and Business Optimization. This journey began at Sun Microsystems, where I supported Tier-1 enterprise clients in planning, prioritizing, procuring, and delivering high-stakes multi-million-dollar initiatives. I continued these full-cycle engagements through my consultancy, JUPYTER, LLC.
 - Post my Sun Microsystems departure, I founded JUPYTER, LLC (originally Rensselaer Technology Group, LTD), where I have served as a principal consultant to Fortune-500 companies on mission-critical projects. Similar to my time at SUN, my work spans a variety of industries, including Financial Services, Pharmaceuticals, Cable TV, Media & Entertainment, Securities Exchanges, Commodities Exchanges, and Telecommunications.
 - To date, I've led many diverse projects, including: Investment Bank Trading Floor Systems Design, Data Center Acquisition, Architecture & Migrations, Video-On-Demand Cable TV Systems, AI-based Recommendation Systems, Business Continuity & Disaster Recovery (BC/DR) SAN Architectures, and Real-Time Trade Stream Analytics Platforms. For additional examples of my work, please refer to: <https://jupyter.ai>
 - This enduring success across time, tools and technologies can be attributed to my unwavering application of time-tested best-practices and playbooks honed through those high-impact, marquee engagements.
-



CAPABILITIES STATEMENT

JUPYTER, LLC

Mission-Critical Technology Consulting • Online I.T. Instruction and Assessments

ABOUT JUPYTER (1997 – PRESENT)

Latest version of this statement: <https://jupyter.ai/nmvega>

Originally formed as **RENSSELAER TECHNOLOGY GROUP, LTD** in 1997 and reformed as **JUPYTER, LLC** in 2013; **JUPYTER** has been principal architect, engineer and strategic advisor to Fortune-500's biggest brands, on their high-value mission critical technology initiatives.

Successful outcomes span major verticals, including: CableTV, Media & Entertainment • Pharmaceuticals • Financial Services • Securities Exchanges • Commodities Exchanges • Telecommunications.

Projects types are equally diverse and include: Financial Trading Systems • Data Center Architectures • Backend Platforms for VoD CableTV Services • Artificial Intelligence-based Recommender Systems • Multi-Datacenter Business Continuity & Disaster Recovery Platforms • Fibre Channel SAN Data Architectures • Real-Time Complex Event Processing (CEP) Platforms for Trade Stream Analytics • Backend Software Development • Cloud Big-Data Platforms for Streaming & Batch Analytics • and more.

JUPYTER is also the company behind the **VERILABS** platform and product (jupyter.ai/verilabs); a Cloud-Native platform for Online Technology Education • Proficiency Exams • and Rapid Solutions Prototyping, on virtually any Information Technology Ecosystem.

Over a near 25-year span – which has witnessed unrelenting change in technologies, tools, techniques and trends – JUPYTER credits it's sustained success to adopting, adapting and faithfully applying the time-tested best practices and playbooks learned during those many marquee engagements. Select real-world artifacts resulting from this body-of-work is found here: <https://jupyter.ai>

CORE CONSULTANCY COMPETENCIES • PRODUCTS

CONSULTANCY

- ❖ C-Suite Advisory for Strategic Business-Technology Adoption
- ❖ Batch & Streaming Data Analytics (Data Engineering & Science)
- ❖ Cloud Data-Platform Design: Analytics • Data Stores • Logistics
- ❖ Event and Data-Driven Backend Software Development
- ❖ Cloud-Native, Containerized Designs via Infrastructure as Code
- ❖ Business Continuity / Disaster Recovery Platforms & Processes
- ❖ Multi-Datacenter Storage Area Network (SAN) Design

PRODUCTS

- ❖ **VERILABS** platform: Rapid Solutions Prototyping & Simulation
- ❖ **VERILABS** platform: I.T. Instruction and Skills Assessment

WHAT SETS JUPYTER APART

- **Time-Tested / Past-Performance Proven:** Founded by the former lead architect to Sun Microsystems' Northeast Tier-1 clients, JUPYTER' successful body-of-work spans 23-years, 6-verticals, and many enterprises-critical successes. As evidence of past performance, we offer a commendation letter from **Deloitte** for a delivered mission-critical project: <https://jupyter.ai/deloitte>
- **Mitigated Risk:** We use our **VERILABS** cloud platform to emulate client I.T. environments, as well as to prototype various solutions prior to implementing one; leading to optimized final solutions, rehearsed deployments and overall risk-mitigated adoption.
- **Reduced Time-To-Benefit:** Projects with procurement delays can continue prototyping atop **VERILABS' R&D Lab**, applying results that emerge from it onto *eventually-procured* assets. This strategy was adopted on a **Bank of America** greenfield project, resulting in an on-time delivery of these lab results: <https://jupyter.ai/bofa>
- **Transparency:** Where client permission was granted, we share [select artifacts from engagements](#) that document our deliberative approach to solutioning. Assets include diagrams, whitepapers, code fragments, technologies & techniques, explainers, Jupyter Notebooks, and more. For instance, the following is the Multi-Datacenter SAN that we architected for **The New York Mercantile Exchange / CME Group** in order to implement its Business Continuity / Disaster Recovery goals: <https://jupyter.ai/nymex>
- **Better ROI:** We over-deliver results that reliably operate for years. Accompanying every delivered platform, we provide operator training, guides, well-planned future OpEx budgets, and more. Because details matter, we always outperform even large rivals.

GENERAL INFORMATION

Registered Company Name: **JUPYTER, LLC**
Year Incorporated (as **LLC**): 2013 (Originally formed in 1997 as a **LTD**)
State of Incorporation: NEW YORK STATE
Corporation Type: LIMITED LIABILITY COMPANY (LLC)
D-U-N-S Number: 111811606
CAGE: 8SVR8
In Business Since 1997: Originally formed as **RENSSELAER TECHNOLOGY GROUP, LTD** in 1997; then reformed as **JUPYTER, LLC** in 2013.

CONTACT INFORMATION

Voice: (212)-882-1276
Fax: (646)-787-0760
Email: hello@jupyter.ai
Web-1: <https://jupyter.ai>
Web-2: <https://aicopilot.me>



CAPABILITIES STATEMENT

JUPYTER, LLC

Mission-Critical Technology Consulting • Online I.T. Instruction and Assessments

SELECT PAST-PERFORMANCE ENGAGEMENTS

FORTUNE-500 END-CLIENT	ENGAGEMENT ROLE(S)
❖ Bank of America	❖ Data platform design for real-time trade analytics
❖ Cablevision Systems (now Altice)	❖ Backend systems design for On-Demand CableTV
❖ NY Mercantile Exchange (NYMEX / CME Group)	❖ Multi-Datcenter Disaster Recovery Platform Design
❖ REUTERS INET -> NASDAQ Stock Exchange	❖ Engineering for Stock Exchange Network (ECN)
❖ Bank of Nova Scotia – Chief Data Officer (CDO)	❖ Director Data Analytics, Standards & Processes
❖ Network Appliance (NetApp)	❖ Product Creator for OnTAP Mgmt Software
❖ Deloitte & Touche	❖ Datacenter Compute, SAN and Storage Architectures
❖ Merrill Lynch -- Investment Research	❖ Machine Learned Article Recommender Engine

NORTH AMERICAN INDUSTRY CLASSIFICATIONS (NAICS)

- ❖ 541511 – Custom Computer Programming Services
- ❖ 541690 – Other Scientific and Technical Consulting Services
- ❖ 541618 – Other Management Consulting Services
- ❖ 611710 – Educational Support Services (Information Technology)
- ❖ 541512 – Computer Systems Design Services
- ❖ 611420 – Computer Training (I.T. Skill Training & Assessment)
- ❖ 611519 – Technical & Trade Schools
- ❖ 541519 – Other Computer Related Services
- ❖ 518210 – Data Processing, Hosting, and Related Services

SELECT PAST-PERFORMANCE END-CLIENTS



SELECT TECHNOLOGY CERTIFICATIONS

<https://jupyter.ai/aws-certifications>

- ❖ AWS: Certified Developer Associate
- ❖ AWS: Certified Solutions Architect Associate

<https://jupyter.ai/deeplearning-certification>

- ❖ Coursera: Certified Neural Network & Deep Learning (Andrew Ng)
- ❖ Coursera: Certified Structured Deep Learning Projects (Andrew Ng)
- ❖ Coursera: Certified Neural Network Optimization (Andrew Ng)

M / WBE CERTIFICATIONS (LOCAL)

- [M/WBE: City of New York, Small Business Services \(SBS\)](#)

M / WBE CERTIFICATIONS (REGIONAL & STATE)

- [M/WBE: NMSDC of New York & New Jersey](#)
- M/WBE: State of New York (In-Progress; Pending)

M / WBE CERTIFICATIONS (FEDERAL)

- M/WBE: SBA 8a (In-Progress; Pending)