

Noelle Milton Vega
Business-Strategic Infrastructure
Architect, Developer, Data Engineer



E: nmvega@jupyter.ai
P: +1 (212)-882-1276
W: <https://jupyter.ai> (CV Portfolio)
W: <https://jupyter.ai/legacy> (Legacy)
W: <https://rndbox.io> (Cloud-Native)
W: [Personal Cloud-Native R&D Lab](#)
W: [LinkedIn Profile](#) (not updated)

URLs for career body-of-work samples:

DELOITTE: [Letter-Of-Commendation](#)
NYMEX: [Multi-Datacenter BC/DR SAN](#)
BofA: [Realtime Trade Anomaly CEP](#)
UBS: [MDM / Lookup system in Python](#)
NETAPP: [Product Devel & Marketing](#)
R&D: [M/L for Human Learning Curve](#)

COVER LETTER

Dear Reader:

As you read this cover letter, please also peruse the sample-works portfolio I have provided for you here: <https://jupyter.ai>. There you will find the [latest copy of this resume](#); A personal [R&D Lab](#) I invested in to assist clients and employers with their initiatives; and a small sample of writings, examples and letters-of-commendations for works I have performed across companies and industries. Thank you.

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 a career after that observation, and my body-of-work has now grown to include consultancy, advisory, entrepreneurship and key driver-roles within the heart of many businesses ranging from global Fortune-500 firms to startups.

The reader will recognize this through full-cycle collaborations with professionals all along the organizational ladder – advising executives on business-vision at one end, guiding technology personnel at the other, and interfacing with key stakeholders everywhere 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 personally conceived, architected and implemented greenfield business-technology platforms that made enterprise mission-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 deeply technical, these days I'm sought for my understanding of the impact that strategic-technology has on businesses ability to stand tall during our emergent times: my greatest values emerge in settings where decision-makers lean on my thought-leadership, knowing that my "in-the-trenches" architectural past will help guide them to optimal business-technology outcomes.

I'd be delighted to discuss your challenges and review business-facing ways in which my breadth & depth 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 further discussion.

Sincerely,

Noelle Milton Vega

Noelle Milton Vega
Business-Strategic Infrastructure
Architect, Developer, Data Engineer



E: nmvega@jupyter.ai

P: +1 (212)-882-1276

W: <https://jupyter.ai> (CV Portfolio)

W: <https://jupyter.ai/legacy> (Legacy)

W: <https://rndbox.io> (Cloud-Native)

W: [Personal Cloud-Native R&D Lab](#)

W: [LinkedIn Profile](#) (not updated)

URLs for career body-of-work samples:

DELOITTE: [Letter-Of-Commendation](#)

NYMEX: [Multi-Datacenter BC/DR SAN](#)

BofA: [Realtime Trade Anomaly CEP](#)

UBS: [MDM / Lookup system in Python](#)

NETAPP: [Product Devel & Marketing](#)

R&D: [M/L for Human Learning Curve](#)

GENERAL CAREER PROGRESSION

Career began as Sonar Communications Engineer for **G.E. Aerospace** (a secret clearance defense project); followed by work as Digital Signal Processing engineer for **Sarnoff/RCA Labs** (in Princeton, NJ) during its R&D work on **DirecTV** as part of my Masters Degree thesis.

Globalization and other macro-economic changes forced a career pivot towards enterprise technology architecture, where with increasing breadth and depth, I navigated my work along Information Technology; Strategic Business-Technology; and Business Optimization roles. These roles began while working for **SUN MICROSYSTEMS**, where I helped their Tier-1 enterprise clients plan, prioritize, procure and deliver high-stakes multi-million initiatives. I would later continue these full-cycle roles under my personal consultancy, PRISMALYTICS.

Departed SUN MICROSYSTEMS to found **PRISMALYTICS, LLC** (originally named **RENSSELAER TECHNOLOGY GROUP, LTD**), where I've been principal consultant to Fortune-500's biggest brands on their mission-critical projects. As with SUN, successful outcomes span every major vertical, including: Healthcare; Financial Services; Pharmaceuticals, CableTV, Media & Entertainment; Securities Exchanges; Commodities Exchanges and Telecommunications.

Projects -- now numbering in the dozens -- were diverse and include: Investment Bank Trading Floor Systems Design • Data Center Acquisition, Architecture & Migrations • Video On Demand Cable TV Systems • Artificial Intelligence-based Recommendation Systems • Business Continuity & Disaster Recovery SAN Architectures • Real-Time Trade Stream Analytics Platform; and more. URL links on the **left** and in **blue** above provide a few examples of work performed.

I credit sustained success to rigorously adopting, adapting, and faithfully applying the time-tested best-practices and playbooks learned during those many marquee engagements.

Noelle Milton Vega
Business-Strategic Infrastructure
Architect, Developer, Data Engineer



E: nmvega@jupyter.ai
P: +1 (212)-882-1276
W: <https://jupyter.ai> (CV Portfolio)
W: <https://jupyter.ai/legacy> (Legacy)
W: <https://rndbox.io> (Cloud-Native)
W: [Personal Cloud-Native R&D Lab](#)
W: [LinkedIn Profile](#) (not updated)

URLs for career body-of-work samples:

DELOITTE: [Letter-Of-Commendation](#)
NYMEX: [Multi-Datacenter BC/DR SAN](#)
BofA: [Realtme Trade Anomaly CEP](#)
UBS: [MDM / Lookup system in Python](#)
NETAPP: [Product Devel & Marketing](#)
R&D: [M/L for Human Learning Curve](#)

EDUCATION AND INDUSTRY VERTICALS — (PLEASE REVIEW BODY-OF-WORK COVER LETTER ON PAGES 1 & 2)

RENSELAER POLYTECHNIC INSTITUTE (RPI) -- B.S. ELECTRICAL ENGINEERING (TROY, NEW YORK)
BOSTON UNIVERSITY (BU) -- M.S. ELECTRICAL & COMPUTER ENGINEERING (BOSTON, MASSACHUSETTS)
HIGH SCHOOL VALEDICTORIAN - GEORGE WESTINGHOUSE HIGH SCHOOL (BROOKLYN, N.Y.)
NEW YORK UNIVERSITY (NYU) -- SCHOOL OF CONTINUING EDUCATION COURSES
U.S. GOV'T SECRET CLEARANCE (COMMUNICATIONS ENGINEER @ G.E. AEROSPACE DEFENSE). *INACTIVE*
CITIZENSHIP(S): U.S.A. CITIZEN (BY BIRTH)
FLUENT LANGUAGES SPOKEN: ENGLISH, SPANISH (PUERTO RICAN DESCENT)

FINANCIAL SERVICES • PHARMACEUTICALS • CABLE TV MEDIA & ENTERTAINMENT
SECURITIES EXCHANGES • COMMODITIES EXCHANGES • TELECOMMUNICATIONS
ENTERPRISE COMPUTE & STORAGE VENDORS • GOVERNMENT SECTOR
RESEARCH & DEVELOPMENT LABORATORIES • DEFENSE INDUSTRY

TECHNOLOGIES: TOOLS, PLATFORMS, CERTIFICATIONS (EMPLOYED OVER YEARS ACROSS DIFFERENT ENGAGEMENTS)

AMAZON WEB SERVICES (AWS) • LINUX / LXC CONTAINERS • DOCKER • PYTHON 3 • UNIX BASH & CLI
AWS-BASED HYBRID DATACENTER ARCHITECTURES • AWS-BASED DATACENTER MITRATIONS
AWS-BASED BUSINESS CONTINUITY & DISASTER RECOVERY (BC/DR) • CLOUDERA CDH AND HADOOP ECOSYSTEM
AWS-BASED REALTIME STREAM COMPLEX EVENT PROCESSORS (C.E.P.) • E.T.L. • NOSQL DATABASES
LONG DISCANCE MULTI-SITE FIBRE-CHANNEL SAN ARCHITECTURES FOR BC/DR
AWS TOOLS (ABBREV): SQS | ECS | SNS | DYNAMODB | ELASTICACHE | EC2 | VPC | KINESIS | DIRECT CONNECT | ETC.
OPEN SOURCE TOOLS (ABBREV): MONGODB | C* | REDIS | STORM | KAFKA | SPARK | SCIKIT-LEARN | TENSORFLOW | ETC.

CLICKABLE URLS to CERTIFICATIONS:

- >> [MINORITY OWNED BUSINESS ENTERPRISE CERTIFICATIONS - \(FEDERAL, STATE, LOCAL\)](#)
- >> [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 \(CSA-A\) \(2018 VERSION\)](#)
- >> [AMAZON WEB SERVICES \(AWS\): CERTIFIED DEVELOPER -- ASSOCIATE \(CD-A\) \(2018 VERSION\)](#)
- >> [NEW YORK UNIVERSITY \(NYU\): CERTIFICATE PYTHON-3 PROGRAMMING \(2008\)](#)
- >> [SUN MICROSYSTEMS \(SUN\): CERTIFICATE ENTERPRISE-STRATEGIC TECHNOLOGY ADOPTION & ARCHITECTURE](#)

EXPERIENCE (INCLUDES URLS TO DOCUMENT / ILLUSTRATION ARTIFACTS WHERE AVAILABLE):

[PRISMALYTICS, LLC - CONSULTANCY](#) — 1998 - PRESENT

Following 4-years as lead architect for **SUN MICROSYSTEMS'** northeast enterprise clients, I departed to found the **PRISMALYTICS, LLC** consultancy. Engagements below marked as “(CONSULTANT)” were performed under C2C contracts with this consultancy; usually as a 3rd-party to prime contractors. Some engagements were performed remotely. During periods when not consulting, efforts focused on entrepreneurial ventures.

[SCOTIABANK \(BANK OF NOVA SCOTIA\) – DIRECTOR DATA ANALYTICS, STANDARDS & PRACTICES](#) APRIL 2019 – PRESENT (CONSULTANT)

Reporting directly to the global Chief Data Officer (within the Data Governance Office), direct a team that works on a Hadoop Datalake platform dedicated to **data governance** initiatives. The platform provides *single-source-of-truth* consumption repositories for Customer and Transaction data originating across bank LOBs. Through ‘user-stories’, LOB data is coalesced, mastered, merged and quality-monitored. Technologies include: Cloudera Hadoop Data Lake; Database CDC with Attunity/Qlik Replicate & Compose; Sqoop; Apache Kafka for data-logistics (including database CDC) and Datalake ingest; PySpark / Python-3 (and Python libraries) atop Spark 2.x for Batch and Streaming analytics; HiveQL; SparkSQL; Jupyter Notebook and more. The end-to-end work bears similarity to the R&D Lab tutorial / example I published here: <https://jupyter.ai/data-engineering-on-movielens-dataset/>

[BRITISH TELECOM – DATA STREAMING AND ANALYTICS](#) JUNE 2018 – MARCH 2019 (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 MLib** and **MongoDB 4.x** real-time analytics stack to identify correlations between transactions.

Noelle Milton Vega
Business-Strategic Infrastructure
Architect, Developer, Data Engineer



E: nmvega@jupyter.ai

P: +1 (212)-882-1276

W: <https://jupyter.ai> (CV Portfolio)

W: <https://jupyter.ai/legacy> (Legacy)

W: <https://rndbox.io> (Cloud-Native)

W: [Personal Cloud-Native R&D Lab](#)

W: [LinkedIn Profile](#) (not updated)

URLs for career body-of-work samples:

DELOITTE: [Letter-Of-Commendation](#)

NYMEX: [Multi-Datacenter BC/DR SAN](#)

BofA: [Realtime Trade Anomaly CEP](#)

UBS: [MDM / Lookup system in Python](#)

NETAPP: [Product Devel & Marketing](#)

R&D: [M/L for Human Learning Curve](#)

[MODERNIZING MEDICINE – RDBMS STREAMING CHANGE DATA CAPTURE \(CDC\) POC ON AWS](#) NOVEMBER 2017 – MARCH 2018 (CONSULTANT)

Remotely designed a Change Data Capture POC to generate real-time change-data-event-streams from updates made to multiple MySQL database shards. The architected POC consisted of (1) Kafka Connect connectors tapping into MySQL bin logs and producing change-data-events onto Kafka topics; these connectors are Docker containers running inside **AWS ECS**, (2) A **Kafka** platform implemented via Confluent's managed Kafka service on AWS, (3) A custom **Python-3** program to consume data from Kafka topics and serialize them as **Apache Parquet** objects in **AWS s3 buckets** (enabling business intelligence teams to efficiently query datasets via **Apache Impala**). Success was to be measured across two KPIs: (1) Replacing an 8-hours long **AWS EMR** batch job with a continuous data stream, in turn replacing hours-stale datasets with near real-time equivalents and (2) A virtually operator-free (or low-touch) solution by way of a managed-services / quasi-serverless centric design.

[SYNERGY FITNESS – AWS Serverless Architecture for Analytic Reports](#)

JUNE 2016 – AS NEEDED (CONSULTANT—RECURRING / RE-ENGAGED AS NEEDED)

Designed AWS-based capture and report platform to track member engagement and R.O.I. Attributes including member visit dates, check-in/check-out times, membership level, facility feature utilization (e.g. trainers, classes, etc) are added or updated to **DynamoDB** via calls to **AWS API Gateway** that trigger **IAM-Role** enabled **AWS Lambda functions** to update the database. Local and Global Secondary Indices are used to accommodate various queries efficiently. Analytic results are programmatically fetched using boto3 Python programs and/or using a **CloudWatch Events** (cron jobs) to periodically trigger Lambda-driven DynamoDB queries that are saved to AWS s3 in CSV format; the completion of which triggers a **AWS SNS** notification to email a business-owner email address with the CSV's s3 restricted URL.

[UNITED BANK OF SWITZERLAND - Business Process Acceleration and Software Development](#)

JANUARY 2015 – MARCH 2016 (CONSULTANT)

Retained by **United Bank of Switzerland** 's **Evidence Lab** for a two-phased project. Phase-1 consisted of an intensive audit of business processes and I.T. capability in New York and London in order to determine what tools could be developed to accelerate financial research analyst activity. The 152-page audit report steered development of a solution consisting of a data transformation pipeline, a mini-Master Data Management (MDM) system, and a taxonomy-centric data catalog. This solution was developed during phase-2 of the project using pure Python-3 and Oracle. The platform enabled the business to receive arbitrary web-derived data, to cleanse, normalize and enrich it; and ingest the final result into a vetted data store which B.I. tools would read. The deliverable was modeled after this solution: <https://jupyter.ai/mdm-in-python>

[BANK OF AMERICA – Trade Stream Anomaly & Fraud Detection](#)

FEBRUARY 2014 – NOVEMBER 2014 (CONSULTANT)

Engaged with Bank of America's trading division to understand requirements for replacing an aging FIX trade analytics system. The replacement needed to identify wash-sales, orders-to-execution latencies, account limit breaches and other **trade anomalies in real-time**. Next, worked with LOBs to secure buy-in and funding for the Complex Event Processing solution I proposed, which was based on open-source big data technologies including Apache Storm, Apache Cassandra and Redis. After reaching consensus, successfully designed, built and delivered the system which was able to detect and process 12-million events/second. Success of this high-visibility open-source project had a bank-wide side-effect of proving the viability "BUILD" over "BUY" option. The following PDF shows some of the R&D I performed during the final design phase: <https://jupyter.ai/apache-storm/>

[MERRILL LYNCH – Architect \(Recommender System\)](#)

MARCH 2013 – FEBRUARY 2014 (CONSULTANT)

Engaged to design a system to recommend investment research publications to wealth clients browsing Merrill Lynch's financial research portal. (Similar to Amazon's "You make also like" feature). The recommender used a combination Item and User-Similarity collaborative filtering algorithms, as well as K-Means clustering to select from 100's-of-thousands of research publications to recommend for purchase. Inputs to the algorithms included client demographics, client investment objectives and publication attributes such as title text, sector, industry, author, age-of-publication, etc. The system was implemented using a combination of Apache Mahout on top of Hadoop as well as Python-3 and Scikit-Learn. A Talend pipeline was used to shuttle raw input data and computed results between HDFS and a data-warehouse.

Noelle Milton Vega
Business-Strategic Infrastructure
Architect, Developer, Data Engineer



E: nmvega@jupyter.ai
P: +1 (212)-882-1276
W: <https://jupyter.ai> (CV Portfolio)
W: <https://jupyter.ai/legacy> (Legacy)
W: <https://rndbox.io> (Cloud-Native)
W: [Personal Cloud-Native R&D Lab](#)
W: [LinkedIn Profile](#) (not updated)

URLs for career body-of-work samples:

DELOITTE: [Letter-Of-Commendation](#)
NYMEX: [Multi-Datacenter BC/DR SAN](#)
BofA: [Realtime Trade Anomaly CEP](#)
UBS: [MDM / Lookup system in Python](#)
NETAPP: [Product Devel & Marketing](#)
R&D: [M/L for Human Learning Curve](#)

SELECT LEGACY EXPERIENCES (PREVIOUS EXPERIENCES OMITTED FOR BREVITY):

[REUTERS --acquisition--> NASDAQ – Stock Exchange Architect | Performance | SAN | Datacenter Migration](#)
JUNE 2004 – 2006 (EMPLOYEE – REUTERS, then NASDAQ after acquisition)

INET was an electronic trading platform / Stock Exchange developed and autonomously operated within REUTERS. It executed orders at sub-millisecond speeds and transacted billions of dollars in daily trade volume. Being the fastest exchange in the world, NASDAQ acquired it in 2005 to replace its own aging exchange (and today runs NASDAQ and other exchanges globally). I was invited to work there by my peers of the dot-com era, (including the -then CTO of INET and subsequently of NASDAQ), and my pre and post-acquisition roles included: **Writing software in C** to manage these black-box (zero human-interface) grid of embedded Linux systems; **Performance engineering** to instrument end-to-end tests & measurements aimed at making the platform still faster; **SAN Storage Architecture** to store quotes, orders and execution transactions in real-time; and **Datacenter Migration** to help with post-acquisition planning and execution of the INET datacenter from New Jersey to Ashburn, Virginia, including the critical stock exchange cutover during a 3-day weekend.

[NYMEX / CME GROUP – Business Continuity / Disaster Recovery Architect \(BC/DR\)](#)
SEPTEMBER 2002 – APRIL 2004 (CONSULTANT)

The New York Mercantile Exchange (NYMEX) engaged me to design and implement a post 9/11 BC/DR datacenter platform, as well as design policies and procedures to follow in case of disaster. The final platform consisted of a Storage Area Network that connected primary and backup datacenters, and real-time synchronously replicated commodities trades and mission critical database transactions between them. A draft of the final platform design is seen here (with permission): <https://jupyter.ai/nymex>

This platform along with accompanying processes and procedures mitigated operational, reputational, financial and legal risk. KPI metrics against which the success of this solution was measured included Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO); both of which were incorporated into bi-annual failover scenarios and tests.

FULL-CYCLE RESPONSIBILITIES INCLUDED:

- Solution presentations to CEO/CIO/CTO to gain consensus
- CAPEX and OPEX justification (ROI / TCO) to plan and procure 3-year budget from the CFO
- Infrastructure capacity planning and right-sizing (including a pay-as-you-grow approach)
- Sourcing vendors and vetting respective product lines for fit & fitness.
- Working with facilities personnel and trade-unions to prepare and place incoming infrastructure
- Architecting, Deploying and Documenting the inter-datacenter SAN platform (see above URL)

[CABLEVISION SYSTEMS – Senior Product Architecture and Market Deployment Strategist](#)
MARCH 2001 – JUNE 2002 (CONSULTANT)

Engaged in all facets of the design, development and deployment of Cablevision's inaugural Optimum.TV systems, the nation's first Video-On-Demand CableTV product. Activities occurred at the Master Headend datacenter and included, designing storage and streaming systems for on-demand ordering of digital content; systems and software to enable at-home self-provisioning of new cable boxes; highly-available databases for customer accounts, digital orders and subscriptions; the Conditional Access System (CAS) used to control subscriber access to programming via encryption; replication of databases and digital content to a backup datacenter; integration with Madison Square Garden (then owned by Cablevision) to offer subscriptions to live and on-demand programming of MSG events; and more.

[DELOITTE & TOUCHE – Application Platform Architect; Team Builder & Manager](#)
JUNE 1997 – FEBRUARY 1998 (SUB-CONTRACTOR via EMPLOYEE FOR SUN MICROSYSTEMS)

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. Duties included: Working with stakeholders to get capacity growth projections for right-sized, pay-as-you-grow budget and purchases; Building the SUN, EMC, SAN, Veritas, Oracle-based farms and documenting them; Building and training the permanent systems team that would operate the farm once I completed the project. A summary of this work was captured by the Deloitte Managing Director here: <https://jupyter.ai/deloitte-commendation-letter>