



Eli Castillo

VP of Engineering - Python

📞 (312) 555-0184 ✉️ eli.castillo.engineering@example.com

🌐 linkedin.com/example/elicastillo 📍 Chicago, IL 60606

STRENGTHS

- 🔇 **Quiet incident leadership**
During difficult valuation outages, separated signal from noise, gave teams a clear recovery path, and became a trusted voice in tense reviews.
- ✅ **Model delivery discipline**
Introduced versioned artifacts and test gates for pricing libraries; one small release check later saved a long production debate.
- 🗣️ **Cross-team translation**
Turned quantitative, trading, and risk questions into practical system choices, helping each group feel heard before a design moved forward.
- ♥️ **Capacity stewardship**
Used scheduling, data locality, and cost reviews across CPU and GPU pools; peers sought advice when demand spiked unexpectedly.
- 👥 **People-first coaching**
Guided engineers through architecture reviews and post-incident learning, giving quieter contributors room to build confidence and influence.

SKILLS

Python AWS GCP Kubernetes
Docker AWS Batch
Google Cloud Batch Terraform
Distributed systems
High-performance computing
CPU/GPU orchestration
Monte Carlo simulation VaR XVA
Performance tuning

SUMMARY

Engineering executive with 13+ years building and operating **Python** platforms for quantitative finance, **distributed systems**, and high-performance computation. Leads AWS and GCP architecture, Kubernetes operations, Docker delivery, **batch processing**, serverless services, infrastructure as code, and production reliability across massive compute grids. Recent work supports millions of trade revaluations across historical market scenarios, including Monte Carlo, VaR, and XVA workloads. Partners closely with **quantitative research**, trading, risk management, and technology leaders, turning valuation needs into resilient products with stronger capacity, cost control, recovery targets, and regulatory traceability. Brings hands-on technical depth plus calm coaching for engineers facing difficult production incidents. Ready to help build a pricing engine where speed, accuracy, ethics, and client protection matter every day.

EXPERIENCE

VP, Quantitative Platform Engineering

Argent Peak Markets 📅 July 2023 - Present 📍 Chicago, IL

Directs architecture, delivery, and operations for a multi-region AWS and GCP pricing platform. Owns Python services, Kubernetes clusters, Docker images, batch queues, serverless control-plane components, model releases, and infrastructure as code. Guides senior engineers while partnering with quantitative research, trading, **risk management**, and technology leadership on capacity, controls, reliability, and product priorities.

- Directed a multi-region **AWS** and **GCP** pricing platform supporting millions of trade revaluations across historical market scenarios, giving risk teams dependable valuation capacity.
- Set direction for Python services, **Kubernetes** clusters, Docker images, batch queues, and **serverless** control-plane components, creating consistent paths for quantitative model delivery.
- Established versioned model artifacts, automated test gates, release controls, and promotion standards, giving regulatory reviewers clearer lineage and stronger production confidence.
- Redesigned orchestration across CPU and GPU pools through scheduling, data locality, and workload isolation, improving Monte Carlo and XVA processing windows.
- Partnered with quantitative research, **trading**, and risk leaders on platform enhancements and valuation incidents, balancing urgent delivery with control requirements.
- Coached senior and mid-level engineers through architecture reviews, incident analysis, capacity planning, Terraform adoption, and Git-based infrastructure workflows.

Senior Director, Distributed Compute Engineering

Bluewater Quantitative Systems 📅 January 2020 - June 2023 📍 New York, NY

Owned engineering delivery for a cloud compute grid running Python **pricing libraries**, scenario analytics, and VaR calculations across containerized AWS workloads. Led architecture, scheduling, model packaging, data logistics, performance analysis, recovery planning, and cross-functional delivery for high-volume valuation operations.

- Owned a cloud compute grid for Python pricing libraries, **scenario analytics**, and **VaR calculations**, giving portfolio teams predictable execution across containerized AWS workloads.
- Designed Kubernetes scheduling and queue-management services that separated urgent intraday requests from overnight portfolio batches, preserving capacity during peak cycles.
- Introduced **Docker** build pipelines, model compatibility checks, and canary execution, reducing release uncertainty for quantitative libraries entering production.
- Reworked market-data and trade-data staging with object-storage manifests, runtime configuration validation, and lineage records, ensuring each calculation received correct inputs.
- Led investigations into serialization, network transfer, worker startup, and memory pressure, restoring predictable completion during high-volume valuation periods.
- Coordinated recovery exercises with operations, risk technology, and **infrastructure** teams, documenting **recovery time objectives** and remediation actions for critical services.

LANGUAGES

English Native

Spanish Proficient

MY CAREER



- VP, Quantitative Platform Engineering at Argent Peak Markets (3.1 Years)
- Senior Director, Distributed Compute Engineering at Bluewater Quantitative Systems (3.4 Years)
- Engineering Manager, Risk Computation at Redwood Harbor Bank (2.9 Years)
- Lead Software Engineer, Pricing Infrastructure at Summit Ridge Financial (2.4 Years)
- Software Engineer, Financial Systems at Crescent Point Analytics (1.8 Years)

Engineering Manager, Risk Computation

Redwood Harbor Bank 📅 January 2017 - December 2019 📍 Boston, MA

Managed a team delivering Python applications and distributed services for derivatives valuation, stress testing, and counterparty exposure analysis. Set engineering standards, guided AWS batch migration, supported quantitative model releases, and connected service design with trading, risk, compliance, and operational needs.

- Managed delivery of Python applications and distributed services for derivatives valuation, stress testing, and counterparty exposure analysis, linking **engineering** work with risk decisions.
- Migrated Monte Carlo workloads from managed servers into **AWS Batch** and **container** services, using reusable Terraform modules for repeatable deployments.
- Worked with quantitative **developers** on dependency packaging, numerical-output validation, and controlled release workflows, protecting model integrity during production changes.
- Built dashboards and alerts for queue depth, worker health, processing latency, cloud spend, and failed-task recovery, giving operators earlier signals during incidents.
- Facilitated design reviews with trading and risk stakeholders, translating business scenarios into service interfaces, data contracts, and testable acceptance criteria.
- Developed **coding**, **debugging**, and implementation standards while mentoring engineers on distributed-systems diagnosis and production change management.

Lead Software Engineer, Pricing Infrastructure

Summit Ridge Financial 📅 July 2014 - December 2016 📍 Charlotte, NC

Built pricing infrastructure for trade enrichment, market-scenario distribution, result collection, and quantitative model execution. Combined Python development with Docker, Kubernetes, **performance tuning**, validation, testing, documentation, and production support for derivatives operations.

- Developed **Python** services coordinating trade enrichment, market-scenario distribution, and result collection, giving derivatives teams a dependable path from inputs through valuation results.
- Implemented **Docker** execution environments and **Kubernetes** deployment templates, making quantitative model testing repeatable across development and controlled release stages.
- Tuned parallel task execution, serialization formats, and worker allocation, shortening overnight valuation cycles without changing model outputs.
- Added checks for market-data timestamps, trade attributes, model versions, and **runtime** parameters before jobs entered a calculation grid, preventing avoidable failures.
- Partnered with quantitative **analysts** on discrepancies between research notebooks and production runs through reproducible test cases and controlled replays.
- Authored runbooks, test standards, and deployment documentation, helping operations engineers support releases and investigate high-impact incidents.

Software Engineer, Financial Systems

Crescent Point Analytics 📅 September 2012 - June 2014 📍 Atlanta, GA

Built Python components for portfolio analytics and scenario processing under senior engineering and quantitative guidance. Contributed tested modules, batch workflows, diagnostics, Linux support, source-control practices, release checks, and post-deployment validation for financial research and scheduled valuation jobs.

- Built **Python** components for portfolio analytics and scenario processing under senior engineer and quantitative analyst guidance, supporting fixed-income and derivatives research.
- Translated calculation requirements into tested modules, service utilities, and batch workflows, giving research teams clearer paths from specifications to usable results.
- Created unit and integration tests for pricing inputs, numerical results, failure handling, and data-file compatibility, catching defects before scheduled runs.
- Improved job diagnostics through structured logs, correlation identifiers, and **runtime** summaries, helping analysts and operators trace issues with less interruption.
- Supported **Linux** compute environments, source-control workflows, release checklists, and post-deployment validation for scheduled valuation jobs.
- Presented implementation findings to **engineering** and quantitative teams, applying review feedback toward clearer interfaces and maintainable Python code.

PROJECTS

Global Pricing Grid Reliability Program 📅 2025

Built a cross-cloud reliability program for massive pricing workloads, connecting recovery objectives, model versioning, runtime data checks, capacity planning, and incident learning. Work reflected a simple belief: financial technology earns trust through dependable answers when pressure rises.

Scenario Data Lineage Exchange 📅 2022

Created a lineage approach for market data, trade data, model configurations, and calculation outputs across batch runs. Stronger provenance helped quantitative teams reproduce results, gave risk partners clearer evidence, and made difficult questions easier to answer.

GPU Monte Carlo Execution Framework 📅 2018

Developed a GPU-aware execution framework for Monte Carlo scenario processing, combining task partitioning, worker allocation, numerical validation, and controlled replay. Faster runs mattered, yet consistent results mattered more for people making risk decisions.

LEADERSHIP & AWARDS

- Engineering Leadership Recognition, Bluewater Quantitative Systems, 2022
- Technology Risk Innovation Award, Redwood Harbor Bank, 2019
- Outstanding Capstone Project, Georgia State University, 2012

EDUCATION

B.S. in Computer Science

Georgia State University 🎓 GPA: 3.7 📅 2012 📍 Atlanta, GA

Coursework: *Distributed Systems, Algorithms, Database Systems, Software Engineering*

CERTIFICATIONS

- AWS Certified Solutions Architect - Professional 📅 2022
- Certified Kubernetes Administrator (CKA) 📅 2020
- HashiCorp Certified: Terraform Associate 📅 2019

TECHNICAL SKILLS

- **Programming Languages:** Python, SQL, Bash
- **Cloud Platforms:** AWS, GCP, Amazon EC2
- **Batch Computing:** AWS Batch, Google Cloud Batch, job queues
- **Containers:** Docker, Kubernetes, container images
- **Serverless Services:** AWS Lambda, Google Cloud Functions, event-driven services
- **Infrastructure as Code:** Terraform, Git, reusable modules
- **Distributed Computing:** distributed systems, high-performance computing, parallel processing
- **Financial Modeling:** Monte Carlo simulation, VaR, XVA
- **Pricing Systems:** financial pricing models, derivatives valuation, scenario analytics
- **Data Logistics:** market data, trade data, object storage
- **Performance Engineering:** performance tuning, serialization, memory profiling
- **Observability:** structured logging, dashboards, alerting
- **Reliability Engineering:** high availability, recovery time objectives, failed-task recovery
- **Software Delivery:** CI/CD, automated testing, canary releases
- **Risk Controls:** model versioning, compliance controls, production governance

PROFESSIONAL AFFILIATIONS

- Mentor, Chicago Cloud Engineering Roundtable, 2024 - Present
- Volunteer Coach, Quantitative Computing Career Workshop, 2022 - Present
- Technical Advisor, Georgia State University Computer Science Alumni Network, 2020 - Present

ADDITIONAL INFORMATION

Work Status : Authorized to work in United States. No sponsorship required.

REFERENCES

AVAILABLE ON REQUEST