

John DeSanto

Principal Engineer | Tech Lead

📞(415) 420-6742 | ✉️jdesanto@gmail.com | 🌐johndesanto | 📍Seattle, WA

Summary

As a software engineer with over 20 years in the field, I have built high-performance systems across search, AI, and data science domains. My expertise includes developing scalable applications in Python, Go, and JavaScript, with a particular focus on machine learning integration and cloud infrastructure. I have driven technical innovation as a tech lead and engineering manager, leveraging my specialized knowledge in ontologies and semantic modeling to enhance application performance in complex domains.

Skills

- **Programming Languages:** Python, Go, Javascript
- **Cloud & DevOps:** Oracle Cloud (OCI), AWS, ETL, Apache Airflow, Docker, CI/CD
- **Data Engineering:** ETL Pipelines, MySQL, PostgreSQL, MongoDB, Redis, Neo4j
- **Machine Learning & AI:** PyTorch, Matplotlib, LLM Integration, Responsible AI, LangChain
- **Specialized Knowledge:** Ontology engineering, Semantic Web, RDF, OWL, SPARQL, TopBraid

Experience

Technical Ontology Developer

1/2024 - 3/2025

Hyperthesis - Los Angeles, CA

Developed software for a semantic search engine focused on health outcomes.

- Engineered semantic relationships from academic publications to enhance disease reasoning capabilities, improving classification accuracy by 20%. Integrated multiple foundational ontologies with subject specific.
- Developed custom Python toolchain for bidirectional RDFS-Excel conversion, enabling efficient collaboration with domain experts.
- Created novel algorithms for PDF section boundary detection using document outline features, enhancing document parsing precision by 27%.

Principal Engineer / Manager

3/2020 - 11/2023

Oracle - Seattle, WA

Developed and maintained Oracle's Accelerated Data Science Python SDK and OCI's hosted notebook service. Managed 4-6 engineers.

- Ensured computing environments met and exceeded customer needs with current versions of machine learning and data science libraries like PyTorch, Matplotlib, and Scikit-learn. Improvements increased downloads by 20% while reducing core environment update times by 33%.
- Improved key components of the data science SDK to simplify important tasks, including enhancing hyperparameter tuning methods and authentication.
- Led Oracle's Guardian AI open-source effort, a unified responsible AI library in Python for detecting bias and privacy risks in machine learning models and LLM output.

Staff Engineer

11/2018 - 2/2020

Medium - San Francisco, CA

Managed and optimized Medium's ETL infrastructure as part of the data engineering team.

- Maintained and improved 200+ Scala/Spark ETL pipelines, including a full database migration (~20TB) from Amazon Redshift to Snowflake on AWS.
- Enhanced the in-house job orchestration framework to support PySpark ETL pipelines and single-instance Python ML jobs, significantly improving job flexibility for ML engineers.
- Led the transition to Apache Airflow, improving infrastructure efficiency and reliability.

Principal Developer, Foundation

5/2016 - 9/2018

EvengX (formerly TrustPipe) - San Francisco, CA

Designed and built the core of EvengX's network intelligence platform for packet analysis and threat detection.

- Developed APIs in Python and Go to identify malware in network packet data using enhanced K-means clustering, significantly improving detection speed and accuracy.
- Built and maintained processing/testing pipelines for supervised model training and continuous integration, enabling rapid model retraining as new data sets were added.

Staff Software Engineer, Architecture

3/2015 - 5/2016

Vungle - San Francisco, CA

Developed and improved Vungle's advertising platform and SDK in Scala, Go, and Python.

- Improved ad customization and reconciliation via optimizations to Vungle's Kafka infrastructure.
- Team lead for designing and implementing Vungle's first advertising exchange. We built an advertising exchange from the ground up in two quarters in Go allowing Vungle to improve valuations on ad inventory.

Software Architect

5/2011 - 2/2015

AdPredictive - San Francisco, CA

Developed and optimized AdPredictive's RTB platform in Python and Javascript on AWS.

- Implemented VAST/VPAID compliance to support improved metrics for video completion and in-video events.
- Created analytics dashboards to monitor VAST-specific metrics, including completion rates, error rates, and interactive engagement.
- Improved profit margins by implementing adjustable bid pricing for high-value users.

Ontologist

(see below for dates)

Full lifecycle development for custom ontologies.

- Created and improved application specific ontologies using industry standard tools Protégé and TopBraid Composer, and semantic web standards such as RDF, RDFS, OWL, and SPARQL.
- Set assessment standards with stakeholders and subject matter experts.
- Built extensive tooling for Excel import/export, visualization, A/B comparison, and querying.

Roles (both contract and full-time)

Electric Capital (Remote)

6/2024 – 9/2024

- Crypto project taxonomy for industry trend analysis.
- Added upper level classifications to a previously flat taxonomy improving segmentation between projects.

Gro-Technology (Remote)

7/2019 - 11/2020

- Food ontology for climate risk and food sourcing analytics.
- Created toolchains to normalize data using the ontology, allowing reasoning across facets

Search Strategy Solutions (San Francisco, CA)

3/2011 - 9/2013

- Taxonomies of retail items to improve search engine performance. Topics: tech retail (printers, headphones, etc.), home retail (home furnishings).

NASA (Mountain View, CA)

8/2007 - 8/2008

- Ontologies for NASA's Constellation program and space shuttle.

Unicru (San Francisco, CA)

10/2002 - 7/2004

- Job role and skills ontology supporting a CRM system for resumes. Designed role and skill ontologies for technology, hospitality, and nursing.

Kanisa (San Jose, CA)

6/2024 – 9/2024

- Product and issue taxonomy supporting queries and recommendations for Apple Computer support KB.

Cycorp (Austin, TX)

9/1998 - 3/2001

- Knowledge engineering for the world's largest and most comprehensive knowledge base
- Created bioterrorism ontology components and modeled complex threat vectors and response protocols
- Led ontology development for CycSecure, implementing cyber threat reasoning capabilities

Education

- **The Johns Hopkins University** - M.A. Philosophy
- **The University of Colorado, Boulder** - M.A. Mathematics
- **The University of Michigan** - B.S. Mathematics and Philosophy