

Contact

jmsanguineti@hotmail.com

www.linkedin.com/in/jmsanguineti
(LinkedIn)

Top Skills

MLOps / LLMOps

RAG & Retrieval Systems

AWS Bedrock & AgentCore

Certifications

Data Science Professional
Certificate

IELTS Academics (7.0 overall)

Amazon Web Services Cloud
Practitioner

TOEIC (855)

Honors-Awards

Scholarship to study abroad under
the program PROIMES

Strathclyde SBS master scholarship

Juan Sanguineti

AI Platform Manager | Building enterprise AI agents at scale | AWS
Agentcore, LLMs, Multi-channel Orchestration
Argentina

Summary

Building AI agent platforms that take conversational AI from prototype to production. I lead the AI Platform team at Liberty Latin America, where we design the frameworks that let Data Science and Business teams ship production-grade AI agents in days, not months. Our platform handles the hard parts: agent orchestration with tool-calling and memory, multi-channel integration (WhatsApp, Telegram, Teams, web), and automated evaluation pipelines that catch regressions before users do. Background: I came to AI engineering through an unconventional path — philosophy (logic, epistemology) + business strategy (Strathclyde MBA) + self-taught Python and cloud architecture. That combination shapes how I approach AI systems: I care about what the agent actually accomplishes for the user, not just whether it technically works. What I build: Agentic workflows with retrieval, tool-calling, and multi-step reasoning. Production inference services with evaluation, tracing, and cost controls. Reusable infrastructure patterns on AWS (Bedrock, Lambda, DynamoDB, CDK). Stack: Python, AWS Bedrock, Strands Agents SDK, LangChain, RAG pipelines, CDK/IaC, Docker.

Experience

Liberty Latin America

5 years 1 month

AI Platform Manager

March 2026 - Present (8 months)

Panama

Building the internal AI agent platform that enables teams to ship productiongrade conversational agents at scale. Key deliverables:

- Agent orchestration framework with tool-calling, memory management, and multi-turn conversation handling
- Multi-channel integration layer — single agent codebase deployed to

WhatsApp, Telegram, Teams, and web

- Automated simulation and evaluation pipeline (LLM-as-judge) for agent quality monitoring.

Stack: Python, AWS Bedrock, Strands Agents SDK, AgentCore, Lambda, DynamoDB, S3, CDK, Docker.

Data Engineering Manager

January 2025 - March 2026 (1 year 3 months)

Panama

Leading the Data Engineering team responsible for ML pipelines, data architecture, and analytics infrastructure across the organization.

Key responsibilities:

- Design and optimization of automated ML pipelines using AWS SageMaker and modern data stack (Glue, Spark, Athena)
- Scalable, cost-efficient data architecture aligned with business objectives
- Data governance, security, and operational best practices

This role built the data foundation that later enabled the AI Platform initiative — the same infrastructure now powers agent memory, retrieval, and evaluation systems.

Stack: AWS (SageMaker, Glue, Athena, S3, Redshift), Spark, Python, Power BI, QuickSight.

Senior Data Engineer

October 2021 - January 2025 (3 years 4 months)

Panama

Designed and optimized automated ML pipelines integrated within a modern AWS data stack.

Architecture highlights:

- Data processing: AWS Glue, Spark SQL, Athena, S3, Data Catalog, Crawlers
- ML orchestration: SageMaker pipelines with containerized components (ECR)
- Event-driven workflows: SNS-triggered jobs, Secrets Manager for credentials
- Full Infrastructure as Code: CI/CD pipelines powered by AWS CDK
- Visualization: Power BI and Amazon QuickSight dashboards

This infrastructure became the backbone for the organization's subsequent AI and ML initiatives.

Accenture

Data Engineer

March 2021 - October 2021 (8 months)

Argentina

Led migration of on-premise data pipelines to AWS for a USA-based enterprise client.

Pipeline architecture:

- Ingestion: Raw data into S3, exposed via external tables in Apache Hive on EMR
- Processing: HiveQL transformations orchestrated by Control-M, Spark jobs on Parquet files
- Loading: Final transformations in Amazon Redshift stored procedures
- Result: Cloud-native data platform supporting downstream analytics at scale

ITE Consult

Data Engineer

September 2020 - March 2021 (7 months)

Argentina

Designed and implemented an ELT program comparing internal demand planning simulations against client ERP data.

Technical scope:

- REST API integration with client ERP systems
- Data transformation and loading into SAP HANA Cloud
- OLAP cube development in SAP Business Application Studio
- Interactive dashboards in SAP Analytics Cloud for stakeholder insights

Grupo Patagonia

Data Engineer

September 2019 - September 2020 (1 year 1 month)

Argentina

Built ETL processes and analytical reports using SQL and Tableau.

Business impact: Identified pricing optimization opportunities on inelastic-demand products, leading to 5% increase in gross profit through targeted margin adjustments.

Education

UC San Diego

Specialization, Big Data · (December 2021 - October 2022)

University of Michigan

Specialization, Python 3 Programming · (December 2020 - January 2021)

Duke University

Specialization, Excel to MySQL: Analytic Techniques for Business · (December 2019 - December 2020)

Strathclyde Business School

Msc, Business and Management · (2018 - 2019)

Universidad Nacional del Litoral

Licenciatura, Philosophy · (2008 - 2016)