

Chris Turner

Profile

Hands-on AI/ML and Generative AI engineering leader with 15+ years spanning AI/ML, data engineering, cloud platforms, and cybersecurity, including leadership roles at PwC and AWS. Led architecture and production deployment of agentic AI systems spanning managed AWS AI services (Amazon Bedrock, Nova, AgentCore) to custom open source model deployments on EKS, ECS, and on-prem air-gapped environments, for clients across government, healthcare, oil and gas, and insurance.

Amazon OpenSearch and Splunk SME at AWS, specializing in agentic AI, multi-agent orchestration, RAG, LLM development, and MLOps automation, with deep experience integrating AI into cybersecurity and threat intelligence workloads.

Social Media

 LinkedIn

 Medium Tech Blog (Pending)

 Public Github

 Portfolio Site (Card)

Work Experience

10/2025 – 08/2026
Remote

Senior Manager - Generative AI Engineering - Automation and Cognitive Solutions
PwC

AI Systems Architecture: Designed and deployed LLM-driven ecosystems spanning Amazon Bedrock, OpenSearch, and open source model serving for Retrieval-Augmented Generation (RAG), multimodal understanding, and contextual search; engineered custom agentic retrieval strategies tailored to each agent variation; started each initiative by partnering with business owners and end users to map current-state processes, understand where teams struggled to apply AI, quantify inefficiencies, and prioritize high-impact use cases.

Advanced Multi-Agent Orchestration: Led development of advanced multi-agent systems using CrewAI, LangGraph, LangChain, and Strands Agents; built custom MCP servers connecting agents to custom machine learning models and third-party AI tools, federated behind unified gateway layers for agent-accessible tooling; designed agent workflows around how users actually worked and built a reusable library of prompt patterns, agent templates, and playbooks teams could pick up directly.

AWS Managed Agent Infrastructure (Bedrock AgentCore): Architected agent workloads against Bedrock AgentCore’s service contract, including session-isolated Runtime hosting, AgentCore Memory for durable multi-turn state via LangGraph checkpointing, and Identity for OAuth-based inbound and outbound auth so agents never hold credentials; designed AgentCore Gateway federation to expose private, cluster-hosted MCP servers as governed agent tools; evaluated managed Runtime against Kubernetes self-hosting and selected EKS for production control, keeping workloads portable across both.

Custom & Open Source Models: Architected customized deployments of open source models from Hugging Face, including fine-tuning, model optimization, and offline air-gapped serving through JFrog Artifactory and Nexus proxy repositories for restricted environments.

Centralized Routing, Observability & Governance: Centralized all LLM calls through a single routing layer using LiteLLM and OpenRouter, with hybrid routing between frontier and open source models to reduce inference cost; instrumented agent observability with Portkey and Langfuse for token-level telemetry, tracing, and cost metrics, using usage and evaluation data to track adoption and see where outputs fell short for users; enforced Amazon Bedrock Guardrails and policy controls in production, and helped teams understand AI capabilities and limits within privacy, security, and responsible-use constraints so they could adopt AI with confidence.

Agent Deployment & MLOps: Deployed custom agents in containerized environments with session isolation, memory, and identity management; implemented custom OAuth authentication for secure agent access; automated delivery with Terraform, AWS CDK, CloudFormation, CodeBuild, and CI/CD pipelines running GenAI microservices with FastAPI

Work Experience

on Kubernetes/EKS.

Document AI & Workflow Automation: Built custom Airflow workflows for OCR extraction and intelligent document processing, unifying structured and unstructured data for large-scale embedding generation feeding agentic pipelines, replacing manual document handling steps in existing business workflows.

Full-Stack AI Engineering: Delivered end-to-end AI applications with React front ends and Open WebUI custom chat interfaces over streaming FastAPI agent endpoints, giving business users a familiar interface to use AI in their daily work.

Leadership, Adoption & Delivery: Hands-on technical leader delivering advanced agentic AI use cases for enterprise clients across government, healthcare, oil and gas, and insurance; hired and led AI-forward delivery teams of data scientists, AI engineers, deployment engineers, cloud engineers, and full-stack engineers; architected systems, wrote production code, and championed engineering standards for RAG, embeddings, and agentic design across PwC's AI practice; owned the relationship between AI platforms and the business teams using them from first training session through ongoing use; ran onboarding sessions, office hours, hands-on demos, and team-specific training tailored to how each group worked; served as the primary channel bringing user feedback back to product, engineering, and design; translated technical capabilities into business value for leadership and measured adoption after go-live.

Previous Experience

02/2022 – 10/2025

Global

AI/ML & Data Analytics Solution Architect

Amazon Web Services (AWS)

Customer Discovery & AI Use-Case Strategy: Trusted technical advisor to Fortune 500 business, security, and analytics teams; ran discovery sessions to understand how teams worked, where they struggled to apply AI/ML, and which use cases would deliver the most value; channeled customer and partner feedback to the OpenSearch and Bedrock service teams to shape the product roadmap.

GenAI Delivery & Migrations: Delivered search and GenAI solutions for Zscaler, American Express, McKinsey, Cisco, and Redpanda, and led Splunk-to-OpenSearch migrations for Western Union, Citibank, Adobe, Yahoo, and Capital One; started each engagement by mapping existing analyst and business processes with the teams who would use the system, staying engaged from first session through production rollout.

Enterprise RAG & LLM Systems: Built enterprise RAG and LLM systems with LangChain, Amazon OpenSearch, GPT-4, and Claude for semantic search, summarization, and document intelligence; Redis-backed vector memory and conversational cache; AWS CDK stacks for one-click Bedrock + OpenSearch RAG environments; fine-tuning and inference with Hugging Face Transformers and SageMaker; paired with customer teams on prompt design and evaluation so they could extend and operate the systems themselves.

Search & Retrieval Engineering: Recognized Amazon OpenSearch SME advancing neural, multimodal, and hybrid (BM25 + neural reranking) search and contextual ranking; directed global OpenSearch initiatives delivering AWS-native analytics and AI/ML.

GenAI Reliability & Evaluation: Defined SLOs for latency, throughput, and cost; bias, drift, and regression detection via SageMaker with Grafana/Prometheus; A/B harnesses for retrieval quality, policy adherence, and hallucination reduction; cut manual review cycles 60% after studying how reviewers used and corrected model outputs.

MLOps & Data Platforms: MLflow, Apache Spark, Airflow, Kubeflow on EKS, and SageMaker Pipelines for batch training, real-time inference, automated retraining, and lineage; petabyte-scale event pipelines on AWS Glue and Apache Iceberg.

Enablement, Training & Adoption: Presented at AWS re:Invent (2023–2024), Splunk .conf, and PartnerCast; built hands-on workshops teaching enterprise teams to apply LLMs to their own workflows; led proofs of concept, immersion days, and follow-up sessions tailored to each team; built reusable prompt libraries and agent templates teams could pick up directly, and tracked adoption after rollout.

01/2019 – 02/2022

Mgr. Cyber Analytics & Solutions Development @ Yahoo | Toyota | Dept. Homeland Security

Booz Allen Hamilton

- **Managed and matured cybersecurity architecture and data engineering** for Booz Allen's Commercial Cyber Fusion Center (CFC), leading enterprise-scale Splunk deployments (AWS/on-prem), integrating CrowdStrike, and developing advanced analytics, dashboards, security essentials, and automated workflows for threat detection, rapid

 Previous Experience

- response, incident management, and vulnerability assessments
- **Engineered AWS-based AI/ML-driven cyber intelligence capabilities**, creating data-driven analytics and machine-learning models to proactively detect sophisticated cyber threats, anomalies, and TTPs from diverse threat actors (nation-state, criminal, hacktivist); collaborated with cross-functional teams (cybercrime, e-crime, government, intelligence)
 - **Provided senior leadership, strategic oversight, and technical mentorship**, leading cybersecurity engagements, proposal development, client presentations, risk identification and mitigation, attack surface management, and utilization of collaboration platforms (Jira, ServiceNow, Confluence), enabling cybersecurity program maturity, threat intelligence, and enhanced decision-making for commercial and government clients

03/2018 – 01/2019
Tampa, United States

Global Cyber Investigative Intelligence/Cyber Analytics Sr. Program Lead
Citibank

- Led data science and machine learning initiatives for a Global Cyber Investigative Intelligence program; developed cyber intelligence, OSINT, investigative analytics, and threat intelligence using data lake architectures for efficient analysis.
- Tracked nation-state actors/APTs by analyzing SOC reports to identify tactics, techniques, and procedures (TTPs); conducted cyber threat research, including dark web analysis, forensic data evaluation, and penetration test reviews to detect security gaps.
- Enhanced cybersecurity analytics by quantifying cyber intelligence into mathematical models; applied GDPR knowledge to implement security controls for EU data compliance, and recommended actionable controls for security improvements.

11/2015 – 03/2018
Tampa, United States

Senior Security Engineer/Analyst
Centene

Identity and Access Management (IAM) with SailPoint

- Delivered comprehensive IAM solutions integrating critical infrastructure with SailPoint IdentityIQ, automating role assignments and policy enforcement.
- Drove best practices in Role-Based Access Control (RBAC), creating custom IAM attributes and security policies within SailPoint to effectively manage access risk and compliance.

01/2013 – 11/2015

Infrastructure IT Support
Raymond James Financial

06/2007 – 06/2009

Minor League Professional baseball player
Houston Astros Minor League

 Strength & Skills

– AI/ML, GENERATIVE AI & DATA ENGINEERING

Generative AI, Agentic AI, AI Adoption & Enablement, RAG, LangGraph, LangChain, CrewAI, Strands Agents, Multi-Agent Orchestration, Custom MCP Servers (FastMCP)	Amazon Bedrock, AgentCore (Gateway, Runtime, Memory, Identity), SageMaker, Hugging Face Fine-Tuning, vLLM, Ollama, TensorFlow, PyTorch, Air-Gapped Serving, Amazon Comprehend Medical	LiteLLM, OpenRouter, Hybrid Model Routing, Langfuse, Portkey, Token Cost Telemetry, Bedrock Guardrails, PHI De-identification, AWS HealthLake, FHIR
---	---	---

– DATA, RETRIEVAL & PLATFORM ENGINEERING

OpenSearch (Neural, Hybrid, k-NN), Semantic Search, S3 Vectors, Embedding Strategy, Reranking, Semantic Deduplication, OCR / Intelligent Document Processing	Glue, Kinesis, MSK, Airflow, Apache Iceberg, S3 + Parquet (Hive-Partitioned), Data Lakes, Python, SQL, Pandas, NumPy, FastAPI, React	Kubernetes / EKS, ECS, GitOps, Open WebUI, OAuth Agent Identity, Terraform, AWS CDK, CloudFormation, CodeBuild, CodePipeline, CI/CD Automation
--	--	--

🏆 Achievements

AI/ML, Analytics & Enablement Leadership — Led 150+ strategic partner and client engagements across AI/ML, data analytics, and cybersecurity, including IBM's EMEA cybersecurity restructuring and PwC's Cyberfusion Center; mentored 50+ mentees, onboarded hundreds of AWS partners, and trained ~1,000 AWS employees. Built an internal Generative AI Salesforce analytics tool, transforming opportunity management and strategic forecasting with LLM-powered insights.

AWS Technical Content Creation & Public Enablement — Featured presenter and technical content builder at AWS re:Invent (2023–2024), Splunk Conference 2024, and AWS Summit, focused on generative AI, cybersecurity, and data platforms. Developed hands-on labs and technical enablement content across AWS Workshop Studio, Solutions Library, and GitHub Samples.

AWS Skill Builder – PartnerCast Series (2023–2024) — Delivered global technical PartnerCast sessions on advanced fraud detection, real-time threat analytics, and LLM-enhanced search with OpenSearch, SageMaker, and Bedrock. Sessions include:

- **Advanced Fraud Detection with Amazon Kinesis, MSK, OpenSearch & Fraud Detector**
- **Enabling Semantic Search with OpenSearch & Amazon SageMaker**
- **Advanced Search with LLMs in OpenSearch & Bedrock**
- **Generative AI for Splunk: Cyber Insights with OpenSearch**
- **Real-Time Threat Detection with OpenSearch Security Analytics**

🎓 Education

University of South Florida, Tampa, FL

Master of Science (M.S.), Cybersecurity (Cyber Intelligence)

Bachelor of Arts (B.A.), Criminology

Hillsborough Community College, Tampa, FL

Associate of Science (A.S.), Network Security & Digital Forensics

🏆 Certificates

AWS Certified: AWS Data Analytics, AWS Solutions Architect, AWS Associate, Developer Associate, Cloud Practitioner • Splunk Power User