

# Akshat Dharmesh Patel

AI Engineer | Secure Agentic AI & Automation

☎ 240-351-7070

✉ [patelaksht24@gmail.com](mailto:patelaksht24@gmail.com)

🌐 [linkedin.com/in/akshat-patel64](https://www.linkedin.com/in/akshat-patel64)

🌐 [Portfolio Website](#)

## Summary

---

**AI Engineer with a security-first engineering background** building staff-facing agents, governed AI workflows, and secure automation systems across high-paced environments. Hands-on experience with **Microsoft Foundry, Microsoft Copilot Studio, Microsoft Teams agent deployment, Power Platform, prompt engineering, workflow automation, vulnerability management, cloud security, and AI governance**. Built and deployed AI-enabled workflows that return staff time by automating repetitive marketing, reporting, and operational tasks while applying **least privilege, approved data handling, access-control review, and OWASP/NIST-aligned risk thinking**.

## Technical Skills

---

### AI Engineering & Agents

Microsoft Foundry, Foundry Agent Service, Microsoft Copilot Studio, Microsoft Teams agents, Power Platform, AI workflow design, prompt engineering, custom agents/copilots, tool/function calling, knowledge grounding, Retrieval-Augmented Generation (RAG), LLM response evaluation, guardrails, workflow automation

### Programming & Data

Python, SQL, PowerShell, Bash, JavaScript, Pandas, Excel automation, JSON, REST APIs, data validation, reporting automation

### Cloud & AI Infrastructure

AWS, Amazon Bedrock, Lambda, S3, RDS, VPC, IAM, WAF, KMS, CloudTrail, CloudWatch, Secrets Manager, Route 53, Docker, Terraform, Amazon EKS

### MLOps / Observability

OpenTelemetry, Prometheus, Grafana, GitHub Actions, response monitoring, tracing, alerting, reliability validation, deployment rollback testing

### AI Security & Governance

NIST AI RMF, OWASP Top 10 for LLM Applications, prompt-injection risk, data-classification review, RBAC, least privilege, SSO/MFA, access policies, secrets management, Secure SDLC

### Security Operations

Rapid7 InsightVM, Splunk, Security Onion, ELK Stack, CrowdStrike, Snort, Wire-shark, vulnerability management, SIEM-aligned triage, endpoint security

### Security Testing

Burp Suite, OWASP ZAP, Nmap, Metasploit, Nikto, Autopsy, FTK Imager

### Enterprise Platforms

Microsoft 365, Google Workspace, Intune MDM, JAMF, HelpScout, 1Password

### Certifications

AWS Certified Security – Specialty, CompTIA Security+, AWS Certified Cloud Practitioner, ISC2 CC, Microsoft Azure AI Engineer Associate (in progress)

## Experience

---

### Sr. IT Security & AI Automation Specialist

July 2025 – Present

*Maryland Athletics*

- Built and deployed staff-facing **Microsoft Foundry / Microsoft Copilot Studio agents** in **Microsoft Teams**, translating departmental workflows into governed AI automations with reusable instructions, controlled data handling, and user-focused interaction design.
- Delivered AI-enabled workflow automation for Athletics staff, helping return time to Marketing, Development, HR, Business, and Leadership teams by reducing repetitive manual work across link generation, campaign review, reporting, and operational information retrieval.
- Built a **tracking-link generator agent** for Marketing that standardized UTM campaign URL naming and reduced link creation time from **5–8 minutes to 1–2 minutes** per request.
- Designing an **AI-assisted Email Performance and Revenue Tracking workflow** that ingests **2 weekly Excel exports** (KPI + revenue), validates campaign data, rolls up staggered sends, evaluates **CTR, unsubscribe, and bounce-rate thresholds**, and produces AP-style insights for faster campaign review.
- Run **AI office hours** and discovery sessions across Marketing, Development, HR, Business, and Leadership, documenting **8+ AI use cases** and converting ambiguous staff pain points into secure automation requirements, agent designs, and implementation plans.
- Evaluate Microsoft-approved and enterprise GenAI tools against institutional policy, data classification, access-control requirements, and approved-use guidance, producing practical recommendations for secure AI adoption.
- Apply a security-first agent design mindset across AI workflows, including least-privilege access, approved data sources, user permission boundaries, prompt/output review, and risk-aware deployment planning.
- Lead weekly **vulnerability management** and remediation reviews using **Rapid7 InsightVM**, partnering with DIT and system owners to prioritize, validate, and track critical findings across **600+ endpoints, servers, and SaaS-connected systems**.
- Perform exploit-based validation and false-positive triage of reported vulnerabilities, investigating endpoint and suspicious-activity issues involving disk encryption, Intune enrollment, access policies, and SIEM telemetry.

*Maryland Athletics*

- Supported security alert review and case documentation by analyzing endpoint, authentication, and network activity for escalation to central IT and follow-up remediation.
- Provided **Tier 1 / Tier 2 IT support** in a fast-paced higher-ed environment, resolving most Tier 1 tickets within **1 hour** and troubleshooting device compliance, enrollment, and access issues across **Intune and JAMF**-managed systems.
- Built and maintained department-specific **application allow lists** and supported software review requests for internal and third-party tools, strengthening endpoint control and reducing unauthorized software usage.
- Coordinated a **data inventory audit** across **500+ assets** and **50+ staff and vendors**, improving visibility into system ownership, data handling, retention, and operational usage.
- Documented incidents and support cases using **NIST** and **MITRE ATT&CK** terminology and delivered **4 security awareness newsletters** on social engineering and typosquatting to improve suspicious-email reporting.

**Cybersecurity Co-op | Detection Engineering & Secure Application Review August 2022 – May 2023***Dotsquares*

- Engineered and deployed a **Snort-based IDS/IPS** on a live university network segment, writing and tuning custom rules across **7+ attack categories**, including DoS/DDoS, backdoor, SMTP, application-detection, and local rule sets.
- Validated detection and prevention effectiveness through controlled attack-scenario testing, packet blocking, alert-trigger validation, and rule tuning, achieving **95%+ detection and blocking coverage** for defined unauthorized-traffic scenarios.
- Built a **real-time monitoring dashboard** with a mobile-friendly interface to surface alerts, logs, and packet activity from **1M+ daily captured packets**, improving investigation speed for non-command-line users.
- Identified and helped remediate **30+ critical vulnerabilities** in a university **e-governance platform**, including misconfigurations, authentication/session weaknesses, SQL injection, XSS, and input-validation flaws.

**Security Researcher | Privacy-Preserving Cloud Search****July 2021 – August 2022***Charotar University of Science and Technology*

- Led a **13.5-month** cloud-security research effort on privacy-preserving search over encrypted cloud data, comparing **SSE, PEKS**, proxy re-encryption, and multi-user searchable-encryption models across security, privacy, and cloud-usage trade-offs.
- Analyzed key leakage risks including **index privacy, search-pattern leakage, access-pattern leakage, keyword-guessing attacks, and file-injection attacks**, connecting cryptographic research to practical cloud data-protection requirements.
- Synthesized mitigation approaches including **ORAM, PIR**, secure indexing, proxy re-encryption, forward/backward privacy, and multi-user access control for privacy-preserving cloud-storage designs.
- First-authored a **Springer Nature** conference paper, *Privacy Challenges and Solutions in Implementing Searchable Encryption for Cloud Storage*, published in **ICTIS 2023 (LNNS Vol. 719)**.

**Projects****MomentumEngine – Secure Agentic AI Research Platform** | AWS, Amazon Bedrock, Agentic AI, Guardrails

- Architected a secure agentic AI platform that combines an **8-factor deterministic ranking engine** with LLM-assisted agents for candidate comparison, explanation generation, announcement synthesis, journaling, and weekly review.
- Designed a **4-layer trust architecture** across raw inputs, normalized claims, canonical facts, and traceable agent outputs, using freshness scoring, anomaly checks, read-only actions, and evidence-backed responses to reduce hallucination risk and preserve human decision control.
- Evaluated historical setups and decision outputs, with the platform identifying early momentum candidates targeting **6%+ moves** and achieving **83.4% setup accuracy** in historical evaluation.

**SecureScalr – Secure AWS Architecture** | AWS, WAF, IAM, Cloud Security

- Designed and validated a secure multi-tier AWS e-commerce architecture with private EC2 application tiers, ALB + Auto Scaling, Multi-AZ RDS, CloudFront-backed S3, **WAF, ACM, IAM roles**, and **Secrets Manager**; sustained **55K+ JMeter requests** while blocking **1,500+** simulated malicious requests.

**TeleTrackr – EKS Observability Pipeline** | EKS, OpenTelemetry, Prometheus, Grafana

- Built a cloud-native observability pipeline on **Amazon EKS** that unified **metrics, traces, and alerts** using **OpenTelemetry**, Prometheus, Grafana, Helm, GitHub Actions, custom PrometheusRule alerts, and Alertmanager notifications for production-style reliability monitoring.

**SnortEduGuard – AI-Aware Academic IDS** | Snort 3, Flask, spaCy, Detection Engineering

- Built a real-time academic intrusion detection system using **Snort 3**, Python, Flask, and **spaCy-powered Smart Search**, with **40+ custom rules** to detect AI tools, VPNs, study-help platforms, C2 behavior, and unauthorized exam-environment activity.

**Education & Publication****University of Maryland, College Park****August 2023 – May 2025***Master of Engineering in Cybersecurity***Charotar University of Science and Technology****July 2019 – May 2023***Bachelor of Technology in Electronics and Communication*