ARJUN TRIVEDI

SENIOR SECURITY RESEARCH LEADER | AGENTIC AI AND AI SECURITY | DETECTION ENGINEERING
People Leadership | Microsoft Security | Multi-Agent Systems | Threat Research | Cloud and XDR
Hyderabad, India | +91 9999691778 | arjuntrivedi42@yahoo.com | LinkedIn | GitHub | PyPI | Newsletter

EXECUTIVE PROFILE

Security research and agentic AI leader with 12+ years of progressive cybersecurity experience and 7+ years at Microsoft. Directly manages and technically leads a 15-person global engineering and research team. Brings a distinctive combination of malware reverse engineering, nation-state threat hunting, cloud and XDR detection, applied machine learning, AI red teaming, MCP tool development, and autonomous security-agent architecture. Known for converting ambiguous research problems into production capabilities, measurable defender outcomes, and scalable teams.

12+ Years	7+ Years	15	200+
CYBERSECURITY	MICROSOFT	ENGINEERS LED	KQL DETECTIONS

LEADERSHIP AND TECHNICAL CAPABILITIES

AI Security and Agentic Systems
Multi-agent architecture, autonomous investigation, MCP and FastMCP, RAG and context grounding, NL2KQL, agent evaluation, prompt injection, jailbreak, tool-abuse, authorization, and excessive-agency testing

Detection and Threat Research
Detection engineering, Microsoft Sentinel, Defender XDR, KQL, UEBA, identity and AiTM detection, threat hunting, MITRE ATT&CK, nation-state tracking, incident response, reverse engineering, and malware analysis

Leadership and Product Delivery
Direct people management, hiring, mentoring, technical strategy, roadmap ownership, cross-functional execution, customer engagement, research-to-product transfer, and executive communication

Cloud and Engineering
Azure, AWS, GCP, Okta, Splunk, Elastic Security, Azure Data Explorer, Python, PowerShell, TypeScript, FastAPI, Docker, Azure Container Apps, Bicep, GitHub Actions, and managed identity

PROFESSIONAL EXPERIENCE

Microsoft | Senior Security Researcher | Detection Engineering and AI Security — Aug 2020 - Present

Global team leadership | Multi-cloud security | AI for Security | India

- **People and technical leadership:** Directly manage and technically lead a 15-person global detection engineering and security research team; own strategy, hiring, performance, mentoring, prioritization, and delivery across product engineering, data science, operations, and customer stakeholders.
- **Microsoft Project Perception:** Contributed to the creation of research and hunting agent capabilities for **Project Perception**, applying organizational context, detection knowledge, and evidence-grade reasoning to agentic defense workflows while preserving human strategy and oversight.
- **AI security and red teaming:** Lead hands-on assessments for prompt injection, jailbreaks, unsafe tool invocation, authorization boundary failures, indirect instruction attacks, and excessive agency; translate findings into guardrails, evaluation scenarios, approval gates, and safer agent behavior.
- **Security Copilot and grounded AI:** Integrated retrieval-augmented threat summarization and refined context, token, and vector handling for Microsoft Security Copilot, improving reported summarization accuracy by 35% for SOC analysts.
- **Applied ML and data science:** Developed and validated Random Forest, Logistic Regression, Gradient Boosting, Isolation Forest, and One-Class SVM models on security telemetry; partnered with data scientists to improve model quality and reduce false positives by 30%.
- **Detection engineering at scale:** Designed 200+ KQL detection and correlation rules across Microsoft Sentinel and XDR, reducing incident response time by up to 50% and strengthening coverage for identity threats, AiTM phishing, UEBA, cloud abuse, and advanced persistence.
- **Nation-state and multi-cloud research:** Tracked actor tactics, techniques, procedures, and infrastructure across Azure, AWS, GCP, Okta, IBM, and Oracle; converted campaign findings into hunts, detections, remediation guidance, and product improvements protecting enterprise environments.
- **Productized cloud research:** Co-created and shipped **Multi Cloud Attack Coverage Essentials** for Microsoft Sentinel, translating cross-cloud threat research into deployable analytics and investigation content adopted by 1,000+ enterprises.
- **Security automation and MSTICPy:** Built 20+ Python and Jupyter automation assets, including MSTICPy contributions, for hunting, enrichment, YARA generation, repetitive analysis, and blue-team workflows, reducing research effort by up to 60%.
- **Research to production:** Directed R&D and proof-of-concept development for detection and threat-research teams, defining production requirements and enabling reliable handoff to engineering for integration into Sentinel, Defender, XDR, and Security Copilot experiences.
- **Enterprise customer leadership:** Partnered with 50+ enterprises to resolve complex SIEM and EDR escalations, identify recurring product gaps, and convert customer evidence into reusable analytics and engineering priorities, improving reported CSAT by 20%.
- **Community and capability building:** Mentor researchers and engineers, deliver technical and executive briefings, and represent AI-driven security innovation through conferences, internal forums, podcasts, hackathons, and customer engagements.

Microsoft | Cyber Threat Hunter | Microsoft Threat Experts — Mar 2019 - Aug 2020

Windows, mobile, and cross-platform threat hunting | India

- Pioneer member of Microsoft Threat Experts, helping design and operationalize the on-demand expert response service for active human-operated attacks against enterprise customers.
- Led the Tier 1 hunting team as subject matter expert, establishing investigation standards, coaching analysts, and improving threat-advisory quality and service-level performance.
- Reduced analyst effort and advisory turnaround through Jupyter hunting notebooks, cloud automation, scripting, and reusable investigation workflows.

Microsoft | Cyber Threat Hunter | Microsoft Threat Experts | ***Continued***

- Built automated detection systems for novel malicious behaviors and indicators of compromise using telemetry analytics, machine learning, and structured human-review pipelines.
- Analyzed platform telemetry for anomalies associated with unwanted software and emerging attacker behavior, improving detection coverage across Windows, mobile, and cross-platform environments.
- Designed reverse engineering review processes, quality controls, and educational material for large contractor teams, enabling scalable and consistent malware triage.

Wipro | Malware Analyst **Jun 2017 - Feb 2019**

India

- Performed malware analysis, reverse engineering, and threat hunting for advanced threats; translated technical findings into indicators, remediation guidance, and customer security response.
- Reverse engineered Android applications and native libraries using IDA Pro and Ghidra, contributing to detection and mitigation of harmful applications in the Android ecosystem.
- Conducted PE and non-PE analysis, threat-log analysis, behavioral investigation, and malware-lab architecture for repeatable static and dynamic research.
- Provided functional leadership by mentoring analysts, setting technical direction, prioritizing investigations, and collaborating with engineering, policy, and legal teams on harmful behavior and product vulnerabilities.
- Authored technical white papers and investigation reports for technical and nontechnical stakeholders, improving understanding of emerging threats and recommended controls.

Sophos | Threat Researcher **Sep 2014 - May 2017**

Malware research and behavioral detection | India

- Analyzed bots, worms, trojans, packed executables, and malicious documents through static and dynamic reverse engineering, disassembly, debugging, unpacking, sandboxing, and network analysis.
- Built static and dynamic behavioral patterns based on file, registry, process, and network activity to improve high-confidence malware detection.
- Evaluated existing analysis and detection systems, recommended improvements, and supported information-security architecture decisions for global endpoint protection products.
- Provided technical investigation and detection expertise for products protecting more than 100 million users across 150 countries.

ADDITIONAL EARLY TECHNICAL FOUNDATIONS

- **Government defense malware-detection project:** Contributed Python automation, exploit proof-of-concept tooling, malicious-file reverse engineering, signature development, malware-lab design, incident analysis, and technical recommendations for a confidential defense engagement.
- **Cybersecurity and programming instructor:** Taught cybersecurity, ethical hacking, C, C++, Java, PHP, and Android fundamentals, building an early foundation in mentoring and technical communication.

FLAGSHIP AGENTIC AI AND SECURITY ENGINEERING

ASRS: Autonomous Security Reasoning System | **Public demo** | **TypeScript** | **Multi-agent security operations**. Designed a multi-SIEM autonomous investigation platform spanning Microsoft Sentinel, Splunk, Elastic Security, and CrowdStrike Falcon. Specialized Triage, Deep Analysis, and Verdict agents correlate alerts, entities, evidence timelines, confidence scores, and MITRE ATT&CK mappings through MCP-enabled workflows

MCP KQL Server | **Python** | **PyPI 2.1.5** | **Official MCP Registry** | **MseeP verified** | **23 stars**. Created a production-oriented MCP server for schema-grounded NL2KQL, strict validation and repair, semantic schema memory, secure Kusto execution, visual analysis, and Azure Container Apps deployment using managed identity and infrastructure as code

FastMCP ThreatIntel | **Python** | **PyPI 0.2.7** | **MseeP verified** | **39 stars** | **12 forks**. Created an AI-ready threat-intelligence MCP server integrating VirusTotal, OTX, AbuseIPDB, and IPinfo with confidence-scored attribution, MITRE ATT&CK mapping, IOC relationships, D3.js visualization, and STIX output

PromptBook and Assured Arc | **AI productivity and agent engineering**. Built a privacy-first prompt-engineering workspace and a portable GitHub Copilot workbench for agent skills, MCP configuration, repeatable workflows, security review, and multi-agent change evaluation

RECOGNITION, SPEAKING, AND COMMUNITY LEADERSHIP

- **Microsoft Global Hackathon 2023 winner** for the MSTICPy AI Assistant; served as a hackathon judge and participated in the Microsoft Security Hackathon 2025.
- **Podcast guest:** **AZ Security, Episode 27**, discussing MCP KQL Server, advanced threat hunting, and AI-enabled SOC automation; speaker at Microsoft and external security events.
- **Author and practitioner:** Created the **MCP KQL Server newsletter** with 1,900+ subscribers and published 15+ articles across Microsoft Tech Community and a personal security blog, including analysis of practical risks in cybersecurity agents.
- Mentored 50+ security professionals and built an active professional community around detection engineering, AI for Security, MCP tooling, and cloud threat research.

PROFESSIONAL DEVELOPMENT

- **SANS FOR578: Cyber Threat Intelligence** | On-demand training
- **Microsoft SkillUp AI** | Applied AI and responsible AI development
- **Cyber Security and Ethical Hacking** | Appin Technologies

EDUCATION

M.Sc., Information Technology | Amity University | 2015 - 2017 | Distance learning

B.Sc., Computer Science (IT) | K.R. Doshi Group of Computer Science College, M.K. Bhavnagar University | 2013 - 2015

SELECTED PUBLIC EVIDENCE

Project Perception | ASRS demo | Official MCP Registry | MCP KQL on PyPI | ThreatIntel on PyPI | GitHub portfolio | Technical blog