Minimum Viable AIMS: ISO 42001 Required Documentation Artifacts for Beginners

Aleksandra Drobnjak

Security Engineer

About me


women in CYBERSECURITY
WiCyS
PROFESSIONAL MENTORSHIP 25-26
Proud to be a MENTOR
meetup
HACKTHEBOX
AWS community builders
MEETUPS
CISA
LEAD AUDITOR
CERTIFIED
ISO/IEC 42001:2023
AUDITOR
CERTIFIED
/IEC 27001:2022

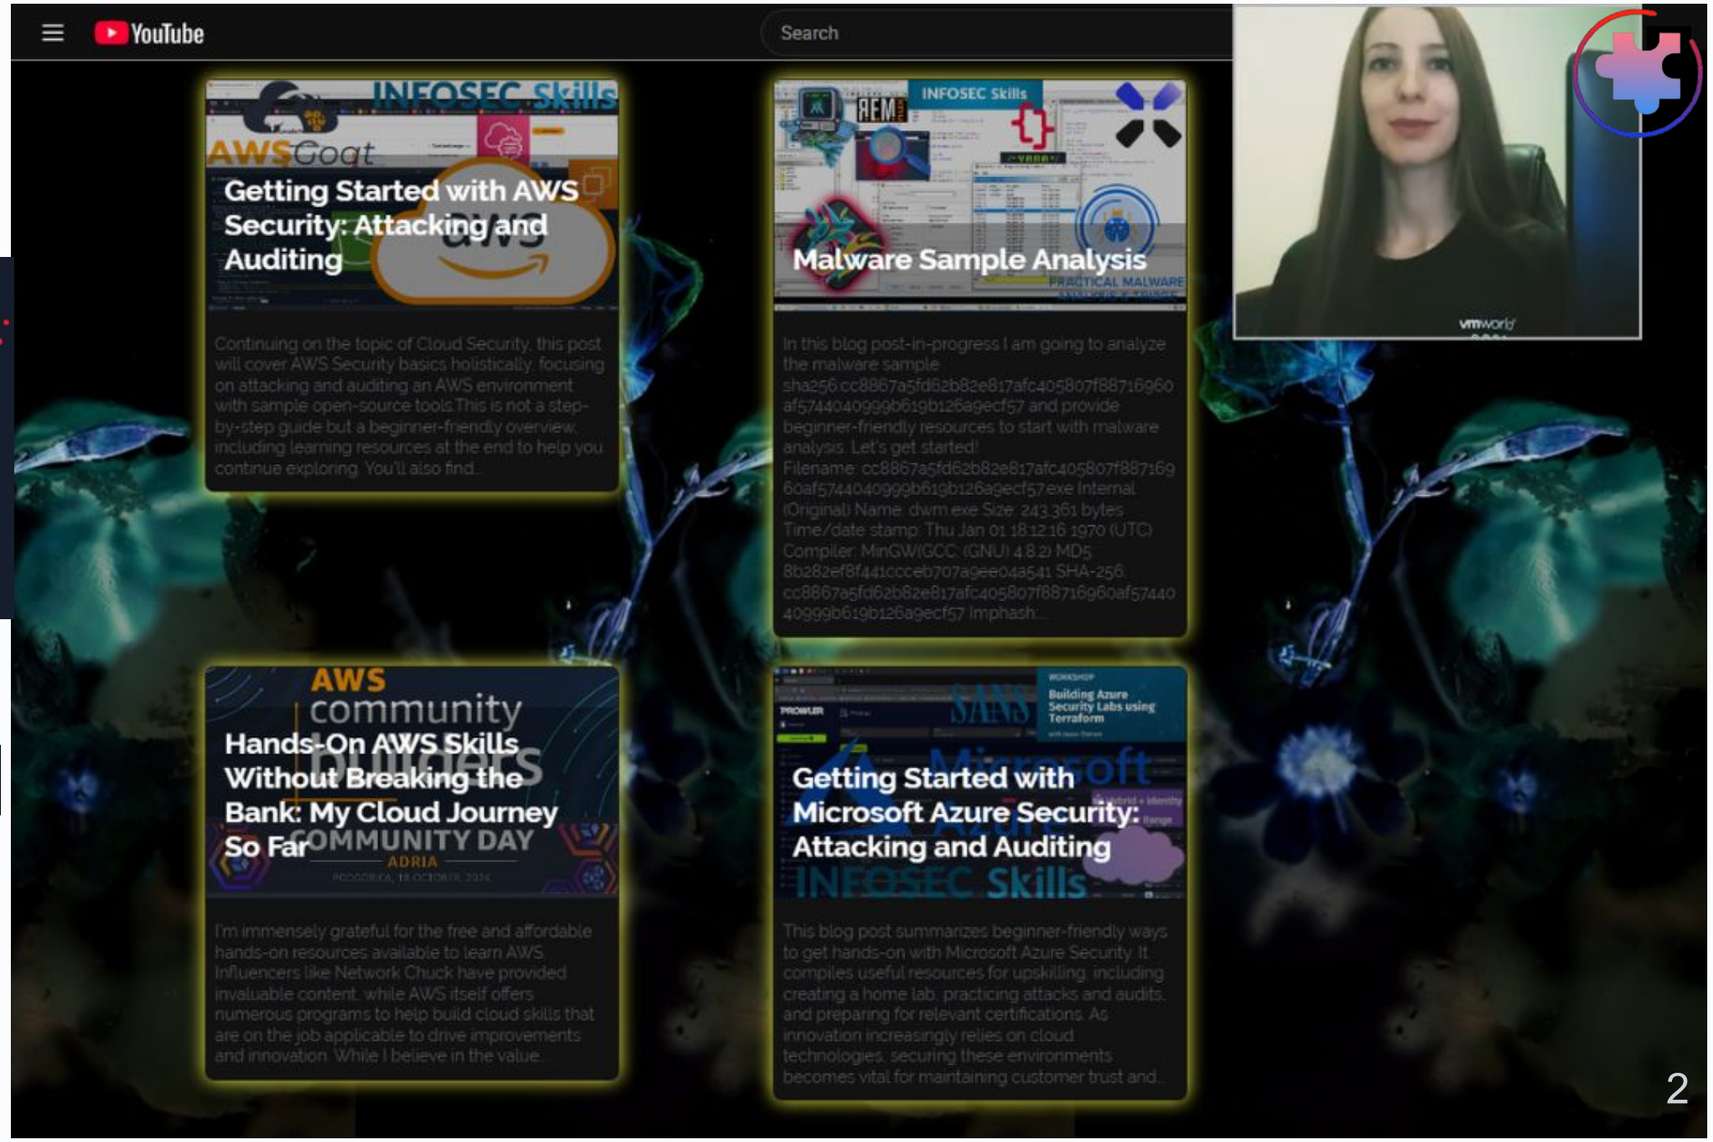
YouTube
Search
INFOSEC Skills
AWS Goat
Getting Started with AWS Security: Attacking and Auditing
Continuing on the topic of Cloud Security, this post will cover AWS Security basics holistically, focusing on attacking and auditing an AWS environment with sample open-source tools.This is not a step-by-step guide but a beginner-friendly overview, including learning resources at the end to help you continue exploring. You'll also find...
INFOSEC Skills
Malware Sample Analysis
PRACTICAL MALWARE
In this blog post-in-progress I am going to analyze the malware sample sha256:cc8867a5fd62b82e817afc405807f887169 60af5744040999b619b126a9ecf57 and provide beginner-friendly resources to start with malware analysis. Let's get started! Filename: cc8867a5fd62b82e817afc405807f887169 60af5744040999b619b126a9ecf57.exe Internal (Original) Name: dwm.exe Size: 243,361 bytes Time/date stamp: Thu Jan 01 18:12:16 1970 (UTC) Compiler: MinGW(GCC (GNU) 4.8.2) MD5: 8b282ef8f441ccceb707a9ee04a541 SHA-256: cc8867a5fd62b82e817afc405807f887169 60af5744040999b619b126a9ecf57 Imphash:...
vmworld
AWS community builders
Hands-On AWS Skills Without Breaking the Bank: My Cloud Journey So Far
COMMUNITY DAY
ADRIA
I'm immensely grateful for the free and affordable hands-on resources available to learn AWS. Influencers like Network Chuck have provided invaluable content, while AWS itself offers numerous programs to help build cloud skills that are on the job applicable to drive improvements and innovation. While I believe in the value
Building Azure Security Labs using Terraform
Microsoft
Getting Started with Microsoft Azure Security: Attacking and Auditing
INFOSEC Skills
This blog post summarizes beginner-friendly ways to get hands-on with Microsoft Azure Security. It compiles useful resources for upskilling, including creating a home lab, practicing attacks and audits, and preparing for relevant certifications. As innovation increasingly relies on cloud technologies, securing these environments becomes vital for maintaining customer trust and...

What is AIMS & ISO/IEC 42001

- AIMS = how an organization manages AI systems over time
- ISO/IEC 42001 — voluntary, certifiable, published 2023
- Integrates with ISO/IEC 27001 (ISMS)
- Auditable clauses: 4–10 → source of required documentation artifacts
- Annex A: 38 controls, not all mandatory
- Adaptable to different organizational maturity levels
- IAF-accredited certification iafcertsearch.org : Stage 1 → Stage 2 → surveillance

Why this matters

- Many organizations still lack AI policies or audits for shadow AI
- Shadow AI breaches expose more sensitive data and drive higher remediation costs
- Public scrutiny is rising with bias and discrimination cases revealing lifecycle governance gaps
- Governance oversight for incorporating technical frameworks like OWASP ML Top 10 and MITRE ATLAS
- Structured AI governance improves operational efficiency, accelerates procurement, and builds customer trust

Amazon Textract

Amazon Transcribe

Amazon Bedrock

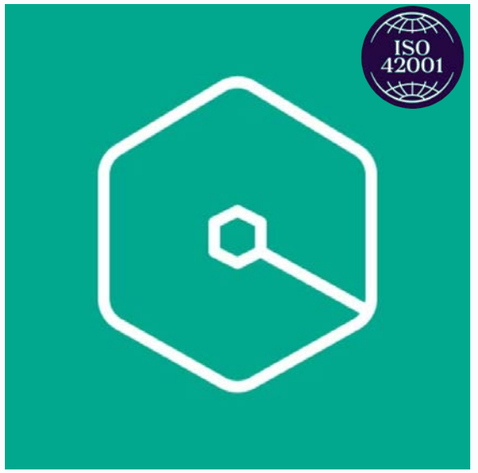

Amazon Q Business

The Minimum Viable AIMS Documentation Artifact Map

- AIMS scope (Clause 4.3)
- AI policy (Clause 5.2)
- Actions to address AI risks & opportunities (Clause 6.1.1)
- AI risk assessment process documentation (Clause 6.1.2)
- AI risk treatment process including Statement of Applicability (Clause 6.1.3)
- AI objectives and related plans (Clause 6.2)
- Personnel competency (Clause 7.2)
- Documentation deemed necessary by the organization (Clause 7.5)
- Documentation ensuring processes are carried out as planned (Clause 8.1)
- AI risk register (Clause 8.2)
- AI risk treatment plans (Clause 8.3)
- AI system impact assessment(s) (Clause 8.4)
- Results of monitoring & measurement activities (Clause 9.1)
- Internal audit program & results (Clause 9.2)
- Management review results (Clause 9.3)
- Nonconformities, corrections & corrective actions (Clause 10.2)

Artifact 1: AIMS Scope (Clause 4.3)

Problem:

> "Models moved from experimentation into production without a formal handoff or ownership"

AIMS Scope Statement — Connected42 Corp.:

PROOF SNIPPET

In Scope:

- Enterprise-managed AI platforms with centralized governance
- ML pipelines deployed via approved CI/CD environments

Out of Scope:

- Experimental AI projects to be decommissioned by a specific date

Scope Change Approval:

- Product Owner (business impact)
- ML Platform Lead (technical readiness)
- GRC Lead (policy & compliance alignment)

Scope Owner:

- Head of Platform Engineering

Artifact 2: AI Policy (Clause 5.2)

Problem:

> "ML teams follow different rules for data usage, retraining, approvals, and deployment"

PROOF SNIPPET

AI Policy:

- Defined stakeholder roles and responsibilities
- Core principles covering transparency, fairness, human oversight
- Escalation and consequences in case of policy violations

Artifact 3: Planning Actions to address AI Risks & Opportunities (Clause 6.1.1)

Problem:

> "Risks were discussed, but never became pipeline controls or work in progress"

PROOF SNIPPET

Risks and Opportunities Register for AIMS document:

Risk: Sensitive data entered into unmanaged external AI tools

Decision: Treat

Action: Restrict AI usage to enterprise-managed platforms

Opportunity: Standardized AI governance accelerates enterprise sales

Decision: Pursue

Action: Publish AI governance overview for customer due diligence

Artifact 4: AI Risk Assessment Process (Clause 6.1.2)

Problem:

> "Model changes and launches are approved inconsistently, risk gates are vibe-subjective"

PROOF SNIPPET

Risk Assessment Criteria:

- Impact: Customer / Legal / Operational
- Likelihood: Rare → Frequent
- Risk Level: Low / Medium / High

Required Before Promotion:

- Reviewer assigned
- Risk rating documented
- Approval recorded

Artifact 5: AI Risk Treatment Process including Statement of Applicability (Clause 6.1.3)

Problem:

> "Teams can't justify why some MLOps controls exist and others were skipped"

PROOF SNIPPET

Statement of Applicability (SoA):

Control A.5.5 Assessing societal impacts of AI systems

Applicable: Yes

Status: Implemented

Justification: Includes relevant environmental, ethical, and societal impact review

Artifact 6: AI Objectives & Plans to Achieve Them (Clause 6.2)

Problem:

> "We're unsure whether leadership is aligned with the metrics our MLOps teams track"

PROOF SNIPPET

AI Objective:

- Reduce model-related incidents by 30%
- Detect model drift within 24 hours
- Patch critical ML infrastructure vulnerabilities within 48 hours

Artifact 7: Personnel Competency (Clause 7.2)

Problem:

> "People influencing AI risk decisions have unclear competency so governance depends on individual judgement"

PROOF SNIPPET

Competency Record:

Role: AI Risk & Governance Lead

Required Competency Areas:

- AI Program Management
- Business Continuity And Incident Response
- Privacy, Ethical, Trust And Safety Controls for AI system

Evidence:

- ISACA Advanced in AI Security Management (AAISM) exam passed

Artifact 8: Documented Information Deemed Necessary by the Organization (Clause 7.5)

Problem:

> "Critical ML decisions live in Slack DMs, across emails, or tribal knowledge — not in version centralized documentation"

Additional Documented Information:

PROOF SNIPPET

- Licensed copy of ISO/IEC 42001:2023 standard
- AI System Inventory record
- AI System Customer Feedback Form

Artifact 9: Operational Documentation (Clause 8.1)

Problem:

> "Incident escalated and took longer than expected to resolve because rollback plans weren't tested"

Operational Control:
Change Management Requirement

- Impact assessment required before deployment
- Rollback procedure defined and tested quarterly
- Approval required from ML Platform Lead
- Change Type: Model Update
- Rollback Verified: Yes

Artifact 10: AI Risk Register (Clause 8.2)

Problem:

> "ML risks are scattered across different platform tickets so some may wither away — no central ownership"

Risk Register:

Risk: Toxic language output bypassing moderation controls

Impact: High (Reputational and regulatory exposure)

Treatment: Input/output validation controls + retraining trigger

Status: In Progress

Artifact 11: AI Risk Treatment Plans (Clause 8.3)

Problem:

> "Risks are identified, but mitigation rarely moves beyond initial registration"

Risk Treatment Plan:

Risk: Third-party attempts to manipulate model outputs via poisoned datasets

Treatment Strategy: Mitigate

Controls Applied: robust access controls + insider risk training + anomaly detection + alerting

Review Cycle: Quarterly

Target Completion: Q3 2026

Status: In Progress

Owner: ML Security Lead

Artifact 12: AI System Impact Assessment(s) (Clause 8.4)

Problem:

> "Models perform as designed — but their real-world impact isn't evaluated holistically"

PROOF SNIPPET

AI System Impact Assessment:

System: Content Recommendation AI

Risk Level: Medium

Stakeholders: Customers, Trial Users

Impact Areas: Social, Ethical

Identified Risk: Promotion of harmful or biased content

Mitigation: Content filtering + bias mitigation strategy

Status: Closed

Artifact 13: Results of Monitoring & Measurement (Clause 9.1)

Problem:

> "Metrics exist — but there is no established review cadence, and decisions aren't really driven by monitoring results"

PROOF SNIPPET

AI KPI sheet:

Monitoring & Measurement Record:

Parameter: Response Accuracy Rate (%)

KPI: >=95%

Red Flag: >65%

Action Plan: Model retraining & validation updates

Monitoring: Monthly

Target: Q4 2026

Artifact 14: Internal Audit Program & Audit Results (Clause 9.2)

Problem:

> "The pipeline changed, docs didn't, and different teams do it their own way now"

PROOF SNIPPET

Internal Audit Record:

- Audit area: Change Management & Governance
- Finding: Approval evidence missing for recent update
- Status: Open — corrective action assigned
- Audit area: Governance Documentation
- Finding: AI Policy lacks documented executive approval
- Status: In progress

Artifact 15: Management Review Results (Clause 9.3)

Problem:

> "Leadership only hears about ML risk after a major incident"

PROOF SNIPPET

Management Review Record:

Reviewed Inputs:

- Monitoring trends (model drift alerts)
- Internal audit findings
- KPI performance

Decision: Increase monitoring cadence for production models

Target Review: Next Quarterly Cycle

Action Owner: Head of AI Operations

Artifact 16: Nonconformities & Corrective Actions (Clause 10.2)

Problem:

> "Model drift kept reappearing because the root cause was never formally addressed"

Register of Nonconformities and Improvements:

Finding: Drift alerts ignored due to unclear ownership

Type: Minor Nonconformity

Root Cause: Monitoring responsibility not formally assigned

Corrective Action: Assign Drift Owner + enforce monthly review cadence

Status: In Progress

Effectiveness Review Date: Q3 Review Cycle

Owner: Head of MLOps Governance

PROOF SNIPPET

Takeaways

Minimum viable ≠ minimal value

Most failures stem from ownership, traceability, and feedback gaps

Integrate AIMS into existing processes, don't reinvent them

Executive sponsorship and training are critical success factors

Prioritize controls proportionate to business impact

These artifacts are not bureaucracy — they're how ML survives reality

Thank you!