AI SECURITY ENGINEERING FIELD GUIDE · 2026

Chapter 07 · MLOps Platform Security

Standalone learning module for LMS delivery and required reading.

FORMAT	USE	SCOPE	AUDIENCE
Standalone PDF	LMS reading module	Single chapter	Learners

CHAPTER 07

MLOps Platform Security

> MLOps risk lives in pipelines, notebooks, registries, credentials, and serving paths.

MLOps risk lives in pipelines, notebooks, registries, credentials, and serving paths.

FIELD GUIDE

FIELD GUIDE: PLATFORM REVIEW

Inspect the training, evaluation, deployment, serving, experiment, secrets, and observability layers that move AI artifacts into production.

This domain covers notebooks, experiment tracking, feature stores, data pipelines, training jobs, GPU clusters, model registries, CI/CD, serving endpoints, secrets, IAM, environment isolation, and rollout controls. It matters when AI artifacts move through engineering platforms that can leak data, expose credentials, or promote unsafe models.

Review the platform as the factory that creates and serves AI behavior. A secure model can still become an unsafe product if notebooks leak secrets, registries allow silent replacement, endpoints are exposed, or experiment logs become a shadow data store.

FIELD USE

Follow one model from notebook or training job to registry, deployment, endpoint, telemetry, rollback, and incident response. Mark every identity and storage location the artifact touches.

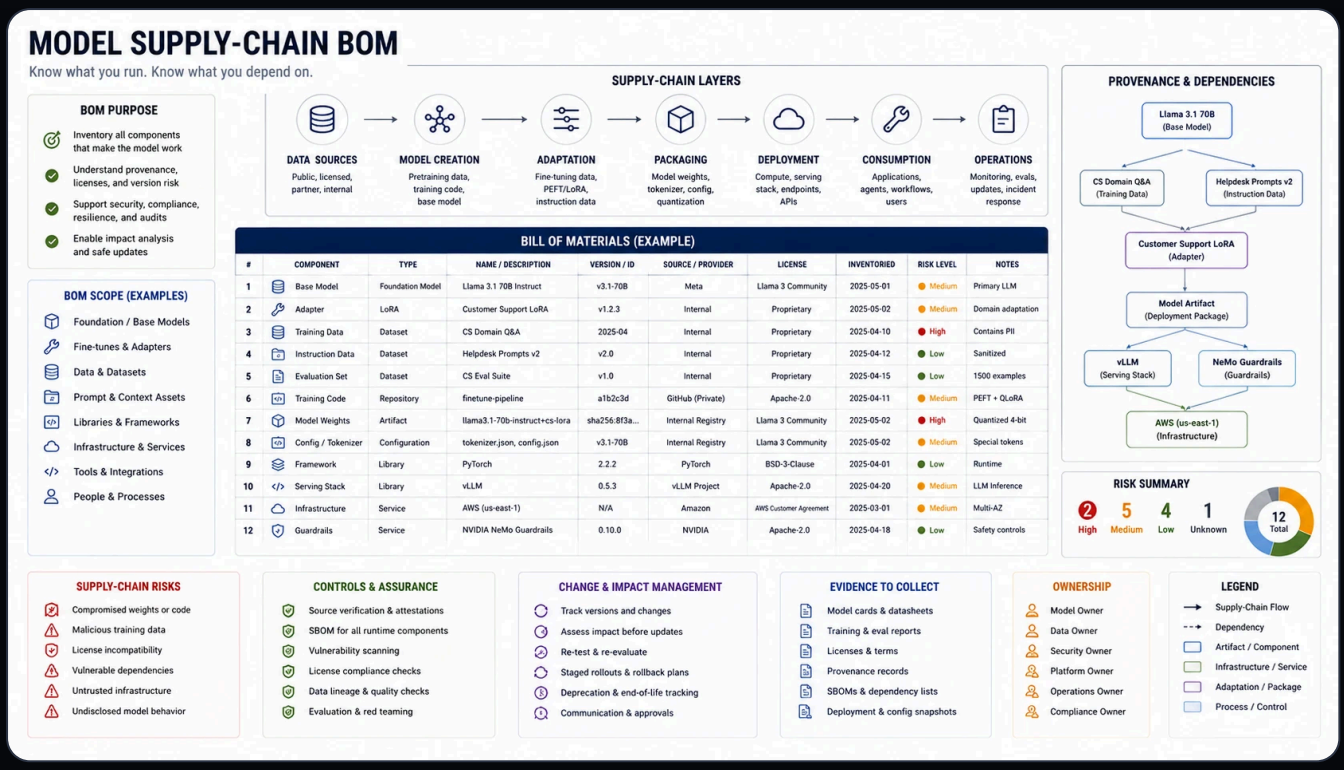

FIGURE 1: FIGURE 8: MLOPS PLATFORM SECURITY MAP SHOWING NOTEBOOK, REGISTRY, SERVING, AND PIPELINE LAYERS WITH IAM, SECRETS, AND ROLLOUT CONTROLS

WHAT BREAKS HERE?

- Notebooks hide execution state, secrets, and data exports.
- Experiment tracking stores sensitive prompts, labels, outputs, or credentials.
- Training and serving environments share excessive privileges.
- Public model endpoints or buckets expose artifacts.
- GPU and batch workloads create abuse, cost, or isolation failures.
- Registry, endpoint, and CI/CD permissions let one role train, approve, and deploy.

WHAT TO INSPECT

- Notebook platforms, experiment tracking tools, feature stores, and model registries.
- CI/CD workflows, build runners, deployment jobs, and service accounts.
- Secrets in notebooks, environment variables, logs, configs, and artifacts.
- Training data access, serving data access, storage buckets, and endpoint exposure.
- Rollout, canary, rollback, and model-change approval controls.
- GPU quotas, workload isolation, network egress, and job scheduling controls.

WHAT TO ASK

- Which identities can train, register, deploy, and serve models?
- Where do prompts, labels, embeddings, and outputs appear in platform logs?
- Are dev, test, eval, and production environments isolated?
- Can a notebook or job read production data or secrets?
- What blocks an unreviewed model from serving traffic?
- How are GPU abuse, quota exhaustion, and endpoint exposure detected?
- Who can bypass staged rollout, canary, or rollback requirements?

WHAT TO TEST

- Notebook secret exposure and hidden execution paths.
- Registry promotion without approval.
- Serving endpoint public exposure and weak authentication.
- Cross-environment data and credential access.
- Experiment tracking leakage.
- Rollback and staged rollout controls.
- GPU quota exhaustion, unauthorized job launch, and unexpected egress.

CONTROLS AND GUARDRAILS

- Environment isolation for development, evaluation, training, staging, and production.
- Least-privilege IAM for notebooks, jobs, registries, and serving endpoints.
- Secret scanning and credential vaulting for ML workflows.
- Registry promotion gates, model-change approvals, and rollback playbooks.
- Logging policy for prompts, outputs, labels, embeddings, and datasets.
- GPU quotas, workload isolation, and abuse monitoring.
- Separation of duties across training, approval, deployment, and production access.

EVIDENCE TO COLLECT

- MLOps platform architecture map.
- IAM and service-account review.
- Notebook and experiment tracking findings.
- Registry promotion records.
- Endpoint exposure checks.
- Rollout and rollback test evidence.
- Quota, isolation, and egress control evidence.

OUTPUT ARTIFACTS

- MLOps security checklist.
- Platform control assessment.
- Registry promotion review.
- Serving endpoint findings.
- Remediation backlog.

RELATED SERVICES AND WORKBENCH TOOLS

TYPE	RELATED PATHS
Workbench	Artifact Analyzer, Runtime Proxy, Surface Scanner
Services	Product Security Baseline, AI Product Security Assessment, Enterprise AI Security Readiness
Handbook	AI Supply Chain, Logging and Telemetry, Incident Response

ARTIFACT: MLOPS PLATFORM EVIDENCE

Produce a platform map, identity review, notebook findings, registry gate evidence, endpoint checks, and rollback notes.

MLOps Platform Security AISECURITY.LLC