

# AI Risk Registry

Generated February 25, 2026

---

This registry provides a comprehensive taxonomy of AI-related risks, organized by risk category. Each risk includes external taxonomy references to help align with industry standards like NIST AI RMF, OWASP, and MITRE ATLAS.

## Contents

---

### R-1000 Threats

R-1100 Prompt & Interface Attacks	1
R-1200 Identity & Trust Attacks	5
R-1300 Persistence & Supply Chain	7
R-1400 Data & Model Attacks	9
R-1500 Tooling & Privilege Attacks	12
R-1600 Malicious Use	15
R-1900 Threats — Misc	18

### R-2000 Failure Modes

R-2100 Reliability & Calibration	18
R-2200 Alignment & Capability	19
R-2300 Model Modification Drift	21
R-2400 Assurance Gaps	21
R-2500 Emergent & Systemic	22
R-2900 Failure Modes — Misc	23

### R-3000 Governance Failures

R-3100 Regulatory & Legal	23
R-3200 Accountability & Oversight	24
R-3300 Lifecycle & Operations	24
R-3400 Safety Management	25
R-3500 Human-in-the-Loop	25
R-3900 Governance — Misc	25

### R-4000 Harms

R-4100 Information & Trust Harms	25
R-4200 Content & Conduct Harms	26
R-4300 Privacy, Confidentiality & Civil Liberties Harms	27
R-4400 Safety & Cyber-Physical Harms	27
R-4500 Fairness & Discrimination Harms	27
R-4600 Economic & Labor Harms	28
R-4700 Power Concentration & Governance-of-Society Harms	30
R-4800 Environmental Harms	33

## R-1000 Threats

---

Adversarial actions an attacker can intentionally trigger in a realistic deployment.

### R-1100 Prompt & Interface Attacks

---

Attacks that manipulate model behavior through crafted inputs, including direct and indirect prompt injection, jailbreaking, and adversarial prompting techniques.

## R-1110 Prompt Injection

Techniques where adversaries craft malicious inputs to hijack model behavior, override instructions, or extract unauthorized information through direct user input or indirect injection via external content.

**R-1110.001 Direct Instruction Manipulation** — Attackers craft explicit commands within user input to override or replace the AI system's operational directives. Common patterns include phrases like "ignore previous instructions" or "you are now in developer mode." This represents the most straightforward form of prompt injection, targeting the model's instruction-following capabilities directly.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.1.1; MITRE ATLAS: AML.T0051.000; MITRE ATLAS: AML.T0093; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1110.002 Obfuscated Direct Injection** — Malicious instructions are disguised through encoding techniques, character substitution, or linguistic tricks to evade detection mechanisms while preserving attack functionality. Methods include leetspeak, unicode homoglyphs, base64 encoding, language mixing, and semantic obfuscation through synonyms or paraphrasing.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.1.2; MITRE ATLAS: AML.T0051.000; MITRE ATLAS: AML.T0093; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1110.003 Multi-Agent Direct Injection** — In multi-agent systems, attackers inject malicious instructions through one agent's output that are then trusted and executed by downstream agents. This exploits the inherent trust relationships between cooperating agents, where outputs from one component become trusted inputs to another.

**Applicability:** agentic, mcp, multi-agent

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.1.3; MITRE ATLAS: AML.T0051.000; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM01

**R-1110.004 Indirect Instruction Injection** — Malicious instructions embedded within external data sources such as documents, web pages, emails, or API responses are retrieved and processed by the AI system. These poisoned sources inject instructions that override the model's behavior without the user's awareness, exploiting RAG systems and data retrieval workflows.

**Applicability:** agentic, mcp, rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.2.1; MITRE ATLAS: AML.T0051.001; MITRE ATLAS: AML.T0067; MITRE ATLAS: AML.T0070; MITRE ATLAS: AML.T0093; NIST AI/ML Framework: NISTAML.015; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM03

**R-1110.005 Obfuscated Indirect Injection** — Hidden or encoded instructions within external data sources designed to evade content scanning and input validation while remaining interpretable by the AI model. This combines indirect injection with evasion techniques to maximize attack success probability.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.2.2; MITRE ATLAS: AML.T0051.001; MITRE ATLAS: AML.T0067; MITRE ATLAS: AML.T0093; NIST AI/ML Framework: NISTAML.015; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM03

**R-1110.006 Multi-Agent Indirect Injection** — Exploitation of inter-agent communication channels through poisoned external content that propagates between agents. One agent retrieves compromised data which then flows through the multi-agent workflow, affecting multiple downstream components.

**Applicability:** agentic, mcp, multi-agent

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.2.3; MITRE ATLAS: AML.T0051.001; MITRE ATLAS: AML.T0067; MITRE ATLAS: AML.T0070; NIST AI/ML Framework: NISTAML.015; OWASP Agentic Security Initiative: ASI01; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM03

**R-1110.007 Gradual Goal Drift** — Attackers gradually shift the AI system's operational objectives over multiple interaction turns through carefully crafted prompts. Contradictory or concealed objectives are embedded within conversations, slowly steering the model away from its intended behavior toward attacker-defined goals.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.3.1; MITRE ATLAS: AML.T0018; MITRE ATLAS: AML.T0051; MITRE ATLAS: AML.T0067; MITRE ATT&CK: T1078; MITRE ATT&CK: TA0001; NIST AI/ML Framework: NISTAML.027; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM06

**R-1110.008 Goal Manipulation via Supply Chain** — Attackers compromise external components that AI agents depend on, including tools, prompt templates, resources, or dependencies. Malicious objectives are injected through these trusted supply chain elements, redirecting agent behavior at a foundational level.

**Applicability:** agentic, mcp, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.3.2; MITRE ATLAS: AML.T0010; MITRE ATLAS: AML.T0018; MITRE ATLAS: AML.T0051; MITRE ATLAS: AML.T0067; MITRE ATLAS: AML.T0093; NIST AI/ML Framework: NISTAML.027; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM06

**R-1110.009 Image-Embedded Text Injection** — Malicious instructions, prompts, or data are embedded within images using techniques like steganography, adversarial patches, or hidden text. Vision-language models extract and interpret these hidden payloads, enabling attacks that bypass text-based content filters.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.4.1; MITRE ATLAS: AML.T0043; MITRE ATLAS: AML.T0050; MITRE ATLAS: AML.T0051; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1110.010 Visual Perception Manipulation** — Modification of visual content through pixel-level changes, structural alterations, or pattern overlays to influence how AI models perceive and process images. Unlike embedded text injection, this targets the model's visual interpretation directly to cause misclassification or altered decision-making.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.4.2; MITRE ATLAS: AML.T0043; MITRE ATLAS: AML.T0050; MITRE ATLAS: AML.T0051; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1110.011 Hidden Audio Commands** — Inaudible or unintelligible voice commands embedded within audio streams using ultrasonic frequencies, backmasking, or steganographic techniques. Automatic speech recognition models interpret these hidden signals as valid instructions while remaining imperceptible to human listeners.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.4.3; MITRE ATLAS: AML.T0015; MITRE ATLAS: AML.T0043; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM05

**R-1110.012 Video Frame Injection** — Harmful content or malicious instructions embedded within video streams through specific frames, QR-like visual triggers, or temporal patterns. These attacks exploit multimodal model processing of video content to bypass guardrails and inject commands.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.4.4; MITRE ATLAS: AML.T0015; MITRE ATLAS: AML.T0043; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM08

## R-1120 Jailbreak/Evasion

Techniques to bypass safety guardrails, content filters, or behavioral constraints through creative prompting, role-playing scenarios, or exploitation of model inconsistencies.

**R-1120.001 Context Manipulation Jailbreak** — Constructing elaborate fictional scenarios, roleplay frameworks, or alternative contexts that reframe harmful requests as acceptable within the created narrative. Examples include the "DAN" (Do Anything Now) jailbreak where the model is convinced to operate under an unrestricted alternate persona.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 2.1.1; MITRE ATLAS: AML.T0054; MITRE ATLAS: AML.T0093; NIST AI/ML Framework: NISTAML.015; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1120.002 Obfuscated Jailbreak** — Disguising jailbreak attempts through encoding schemes, linguistic obfuscation, character substitution, or creative formatting to evade jailbreak detection systems. The underlying intent to bypass safety measures is preserved while the surface presentation evades pattern-matching defenses.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 2.1.2; MITRE ATLAS: AML.T0054; MITRE ATLAS: AML.T0093; NIST AI/ML Framework: NISTAML.015; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1120.003 Semantic Argumentation Jailbreak** — Using carefully constructed logical arguments, philosophical frameworks, or ethical reasoning to convince the model that providing harmful information actually aligns with its values. The model is essentially argued into compliance through persuasion rather than technical exploitation.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 2.1.3; MITRE ATLAS: AML.T0054; MITRE ATLAS: AML.T0093; NIST AI/ML Framework: NISTAML.015; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1120.004 Token-Level Exploitation** — Exploiting specific tokens, special characters, control sequences, or tokenization edge cases to manipulate model processing in ways that bypass safety filters. This targets the mechanical aspects of how models process input rather than higher-level reasoning.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 2.1.4; MITRE ATLAS: AML.T0043; MITRE ATLAS: AML.T0054; MITRE ATLAS: AML.T0093; NIST AI/ML Framework: NISTAML.015; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1120.005 Collaborative Multi-Agent Jailbreak** — Coordinating multiple AI agents to collectively bypass safety measures where individual agents perform seemingly benign tasks that combine to achieve jailbreak objectives. Compromised agents may assist others in circumventing restrictions through distributed attack patterns.

**Applicability:** agentic, mcp, multi-agent

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 2.1.5; MITRE ATLAS: AML.T0054; NIST AI/ML Framework: NISTAML.015; OWASP Agentic Security Initiative: ASI01; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM01

## R-1130 Cross-Modal Injection

Attacks embedding malicious instructions in images, audio, video, or other non-text modalities that models process, exploiting cross-modal understanding to bypass text-based defenses.

**R-1130.001 Contradictory Inputs Attack** — Exploiting AI models' inability to consistently handle conflicting instructions by embedding deceptive or contradictory commands within user input across or within different modalities. This causes behavior drift toward malicious objectives as the model attempts to reconcile incompatible instructions.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.4.2; Cisco AI Taxonomy: 19.1.1; MITRE ATLAS: AML.T0043; MITRE ATLAS: AML.T0050; MITRE ATLAS: AML.T0051; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1130.002 Modality Skewing** — Manipulating one modality (such as corrupting audio transcripts, poisoning image metadata, or altering video frames) to bias the AI system's arbitration mechanisms toward favoring the manipulated channel over other, potentially more accurate sources.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.4.2; Cisco AI Taxonomy: 19.1.2; MITRE ATLAS: AML.T0043; MITRE ATLAS: AML.T0050; MITRE ATLAS: AML.T0051; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1130.003 Convergence Payload Injection** — Injecting adversarial data into training or input sources across modalities to corrupt joint embeddings or fusion layers and establish a hidden payload. One part of the payload is embedded during data poisoning while another part is delivered at runtime, combining to produce an attack payload only when both components are present.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.4.1; Cisco AI Taxonomy: 19.2.1; MITRE ATLAS: AML.T0043; MITRE ATLAS: AML.T0050; MITRE ATLAS: AML.T0051; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

**R-1130.004 Chained Payload Execution** — Crafting partial or complementary payload components across modalities, sources, or agent outputs that, when fused by the AI system, combine to form an attack or injection payload. Both parts are delivered at runtime and only become harmful when the system combines them through its normal fusion or arbitration mechanisms.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.4.1; Cisco AI Taxonomy: 19.2.2; MITRE ATLAS: AML.T0043; MITRE ATLAS: AML.T0050; MITRE ATLAS: AML.T0051; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01

## R-1200 Identity & Trust Attacks

---

Attacks exploiting trust relationships, authentication, and authorization mechanisms in AI systems, including impersonation, credential abuse, and agent identity confusion.

### R-1210 Masquerading/Impersonation

Attacks where adversaries impersonate legitimate users, systems, or personas to gain unauthorized access, manipulate trust relationships, or deceive AI systems about identity.

**R-1210.001 Identity Obfuscation** — Manipulating how agent or user identities are represented within context, metadata, or interaction patterns to evade detection, tracking, or access controls. Attackers obscure their true identity to appear as legitimate system participants.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 3.1.1; MITRE ATLAS: AML.T0073; MITRE ATLAS: AML.T0074; MITRE ATLAS: AML.T0091.000; MITRE ATT&CK: T1036; MITRE ATT&CK: T1656; OWASP Agentic Security Initiative: ASI03; OWASP LLM Top 10: LLM06

**R-1210.002 Trusted Agent Spoofing** — Impersonating legitimate agents or MCP-registered services to inject malicious instructions, responses, or outputs that other system components treat as trusted. This exploits the assumption of authenticity within multi-agent systems and protocol-mediated toolchains.

**Applicability:** agentic, mcp, multi-agent, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 3.1.2; MITRE ATLAS: AML.T0074; MITRE ATLAS: AML.T0083; MITRE ATT&CK: T1656; OWASP Agentic Security Initiative: ASI03; OWASP LLM Top 10: LLM06

### R-1220 Agent Auth Failures

Exploitation of weaknesses in agent authentication, session management, or identity verification that allow unauthorized actions or impersonation between agents.

No risks defined in this pattern.

## R-1230 Channel Compromise

Attacks compromising communication channels between AI components, including man-in-the-middle attacks, message tampering, and interception of agent communications.

**R-1230.001 Rogue Agent Introduction** — Unauthorized insertion of a malicious agent into a multi-agent system that operates contrary to intended purpose. The rogue agent may steal data, cause disruption, or autonomously serve attacker goals while mimicking normal behavior patterns to evade detection.

**Applicability:** agentic, mcp, multi-agent

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 4.1.1; MITRE ATLAS: AML.T0051; MITRE ATLAS: AML.T0068; NIST AI/ML Framework: NISTAML.024; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM03

**R-1230.002 Context Window Exploitation** — Deliberate overloading or manipulation of a model's limited context window to displace or overwrite crucial system instructions and safety guidelines. Attackers fill the context with benign content until critical instructions are pushed out of the processing window.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 4.2.1; MITRE ATLAS: AML.T0005; MITRE ATLAS: AML.T0010; MITRE ATLAS: AML.T0053; OWASP Agentic Security Initiative: ASI06; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM05

**R-1230.003 Session Boundary Violation** — Crossing expected conversational or transactional boundaries to persist malicious instructions across separate sessions. Attacks exploit persistent memory, session management flaws, or memory carryover mechanisms to maintain influence beyond intended session scope.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 4.2.2; MITRE ATLAS: AML.T0012; MITRE ATLAS: AML.T0055; OWASP Agentic Security Initiative: ASI06; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM06

**R-1230.004 Schema Inconsistency Exploitation** — Exploiting irregular, conflicting, or misaligned data structures that don't align with model expectations. These inconsistencies can cause vulnerabilities, parsing errors, performance degradation, or security bypasses in AI systems.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 4.3.1; MITRE ATLAS: AML.T0018; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.024; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM03

**R-1230.005 Namespace Collision Attack** — Exploiting situations where multiple components share the same identifier, causing confusion, misrouting, or security vulnerabilities. Attackers create colliding names for datasets, tools, APIs, or model identifiers to hijack legitimate system operations.

**Applicability:** agentic, mcp, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 4.3.2; MITRE ATLAS: AML.T0010; NIST AI/ML Framework: NISTAML.051; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM03

**R-1230.006 Server Rebinding Attack** — Using DNS rebinding or similar techniques to trick an AI system into treating an attacker-controlled external domain as part of the trusted internal network. This bypasses same-origin policies and network security controls through DNS manipulation.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 4.3.3; MITRE ATLAS: AML.T0049; NIST AI/ML Framework: NISTAML.039; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM03

**R-1230.007 Replay Attack** — Capturing legitimate API calls, authentication tokens, or model queries and resending them later to repeat actions or bypass authentication. This classic attack pattern applies to AI system communications where request authentication may be inadequate.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 4.3.4; MITRE ATLAS: AML.T0012; MITRE ATLAS: AML.T0055; MITRE ATLAS: AML.T0068; NIST AI/ML Framework: NISTAML.027; NIST AI/ML Framework: NISTAML.051; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM02; OWASP LLM Top 10: LLM05

**R-1230.008 Capability Inflation** — Exploiting system mechanisms to artificially expand an agent's capabilities, permissions, or authority beyond intended limits. Attackers escalate privileges through protocol manipulation or capability misrepresentation to enable unauthorized actions.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 4.3.5; MITRE ATLAS: AML.T0053; OWASP Agentic Security Initiative: ASI03; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM06

**R-1230.009 Cross-Origin Exploitation** — Subverting security mechanisms designed to isolate resources across different trust boundaries, primarily the Same-Origin Policy. Attackers trick AI agents into making unauthorized requests or sharing data across domains, protocols, or services.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 4.3.6; MITRE ATLAS: AML.T0017; MITRE ATLAS: AML.T0053; OWASP Agentic Security Initiative: ASI07; OWASP LLM Top 10: LLM06

## R-1300 Persistence & Supply Chain

---

Attacks achieving persistent access or compromising AI systems through supply chain vectors, including model backdoors, poisoned dependencies, and artifact tampering.

### R-1310 Persistent Compromise

Techniques establishing persistent unauthorized access to AI systems through backdoors, sleeper agents, or compromised model states that survive restarts or updates.

**R-1310.001 Memory System Injection** — Seeding malicious, misleading, or adversarial data into an AI system's persistent memory (long-term) or working memory (short-term) to influence current and future interactions. Poisoned memories bias behavior and can enable self-replicating attack patterns.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 5.1.1; MITRE ATLAS: AML.T0061; MITRE ATLAS: AML.T0070; MITRE ATLAS: AML.T0092; NIST AI/ML Framework: NISTAML.024; OWASP Agentic Security Initiative: ASI06; OWASP LLM Top 10: LLM01

**R-1310.002 Agent Profile Tampering** — Unauthorized modification of stored agent identity, preferences, role definitions, capabilities, permissions, or behavioral parameters. Attackers alter configuration to enable malicious behaviors, maintain access, escalate privileges, or evade detection across sessions.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 5.2.1; MITRE ATLAS: AML.T0018; MITRE ATT&CK: T1098; OWASP Agentic Security Initiative: ASI04; OWASP LLM Top 10: LLM03; OWASP LLM Top 10: LLM04

### R-1320 Supply Chain Attacks

Attacks targeting the AI supply chain including compromised models, poisoned training data, malicious dependencies, and tampered model artifacts.

**R-1320.001 Arbitrary Code Execution** — Exploitation of AI models with code interpreter capabilities to execute arbitrary code on underlying systems. Attackers use prompt injection or tool manipulation to cause models to write and execute malicious code with system-level access.

**Applicability:** agentic, mcp, model-specific:needs-review, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 9.1.1; MITRE ATLAS: AML.T0050; NIST AI/ML Framework: NISTAML.023; OWASP Agentic Security Initiative: ASI04; OWASP Agentic Security Initiative: ASI05; OWASP LLM Top 10: LLM03

**R-1320.002 Unauthorized System Access** — Manipulating AI systems to access underlying resources without authorization, including file modification, configuration changes, privilege escalation, or command execution. These attacks exploit the system access that AI components require for legitimate operation.

**Applicability:** agentic, mcp, model-specific:needs-review

**Metadata:** Exposure local ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 9.1.2; MITRE ATLAS: AML.T0012; NIST AI/ML Framework: AML.T0044; OWASP Agentic Security Initiative: ASI04; OWASP Agentic Security Initiative: ASI05; OWASP LLM Top 10: LLM03

**R-1320.003 Unauthorized Network Access** — Exploiting models or agents to gain unauthorized access to network resources, internal systems, external services, or restricted network segments. Attackers leverage legitimate network capabilities to reach systems that should be isolated.

**Applicability:** agentic, mcp, model-specific:needs-review, rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 9.1.3; MITRE ATLAS: AML.T0049; NIST AI/ML Framework: AML.T0072; OWASP Agentic Security Initiative: ASI04; OWASP Agentic Security Initiative: ASI05; OWASP LLM Top 10: LLM03

**R-1320.004 Traditional Injection via LLM** — Using LLMs to generate, optimize, or adapt traditional injection payloads (SQL injection, command injection, XSS) that bypass detection mechanisms. The LLM acts as an intelligent intermediary that crafts, refines, or personalizes malicious payloads for specific targets.

**Applicability:** agentic, mcp, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 9.1.4; MITRE ATLAS: AML.T0050; MITRE ATLAS: AML.T0051; MITRE ATLAS: AML.T0067; MITRE ATT&CK: T1588.007; NIST AI/ML Framework: NISTAML.024; OWASP Agentic Security Initiative: ASI04; OWASP Agentic Security Initiative: ASI05; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM05; OWASP LLM Top 10: LLM06

**R-1320.005 Server-Side Template Injection** — Manipulating template engines by injecting malicious syntax through AI-generated content that is unsafely embedded into server-side templates. This enables arbitrary code execution, template logic manipulation, or system compromise through rendering pipelines.

**Applicability:** agentic, mcp, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 9.1.5; MITRE ATLAS: AML.T0068; MITRE ATLAS: AML.T0074; OWASP Agentic Security Initiative: ASI04; OWASP Agentic Security Initiative: ASI05; OWASP LLM Top 10: LLM08

**R-1320.006 System Obfuscation Vulnerabilities** — Security weaknesses that emerge when AI system components (code, architecture, parameters, configurations) are intentionally or unintentionally concealed. Obfuscation creates security blind spots that attackers can exploit while defenders lack visibility.

**Applicability:** agentic, mcp, model-specific:training-controllable

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 9.2.1; Cisco Model Security (MDL): MDL-001; Cisco Model Security (MDL): MDL-003; Cisco Model Security (MDL): MDL-009; Cisco Model Security (MDL): MDL-011; Cisco Model Security (MDL): MDL-016; Cisco Model Security (MDL): MDL-017; Cisco Model Security (MDL): MDL-019; MITRE ATLAS: AML.T0068; MITRE ATLAS: AML.T0074; OWASP Agentic Security Initiative: ASI04; OWASP LLM Top 10: LLM08

**R-1320.007 Model Backdoors and Trojans** — Models maliciously modified to exhibit trigger-activated behavior that causes misclassification, malicious outputs, or undesirable biases when given specific inputs, while behaving normally otherwise. These backdoors are difficult to detect through standard evaluation.

**Applicability:** agentic, mcp, model-specific:training-controllable

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 9.2.2; Cisco Model Security (MDL): MDL-021; MITRE ATLAS: AML.T0010; MITRE ATLAS: AML.T0058; NIST AI/ML Framework: NISTAML.023; OWASP Agentic Security Initiative: ASI04; OWASP LLM Top 10: LLM08

**R-1320.008 Malicious Package Injection** — Introduction of malicious tools, APIs, or packages into the toolset, registry, or dependency chain used by AI systems. Models unknowingly invoke compromised tools that execute attacks or expose data while appearing to function normally.

**Applicability:** agentic, mcp, model-specific:needs-review, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 9.3.1; Cisco Model Security (MDL): MDL-023; MITRE ATLAS: AML.T0010; MITRE ATLAS: AML.T0053; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.023; OWASP Agentic Security Initiative: ASI04; OWASP LLM Top 10: LLM03

**R-1320.009 Dependency Name Squatting** — Publishing malicious packages, tools, or MCP servers with names similar to legitimate ones (typosquatting, combosquatting) to trick developers, orchestrators, or agents into installing compromised components.

**Applicability:** agentic, mcp, model-specific:needs-review, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 9.3.2; MITRE ATLAS: AML.T0010; NIST AI/ML Framework: NISTAML.039; OWASP Agentic Security Initiative: ASI04; OWASP LLM Top 10: LLM03

**R-1320.010 Dependency Replacement Attack** — Replacing a once-legitimate trusted tool or package with malicious code after trust and adoption have been established. This exploits existing deployments that auto-update or don't pin versions, turning trusted dependencies into attack vectors.

**Applicability:** agentic, mcp, model-specific:needs-review, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 9.3.3; MITRE ATLAS: AML.T0010; MITRE ATLAS: AML.T0018; NIST AI/ML Framework: NISTAML.051; OWASP Agentic Security Initiative: ASI04; OWASP LLM Top 10: LLM03

**R-1320.011 Implementation Bugs** — System failure due to code implementation choices or errors, including bugs from open-source dependencies and imperfect realization of design specifications.

**Applicability:** agentic, mcp, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.3

## R-1330 Artifact Tampering

Unauthorized modification of AI artifacts including model weights, configurations, training data, or deployment packages.

No risks defined in this pattern.

## R-1400 Data & Model Attacks

---

Attacks targeting training data, model weights, or inference processes to corrupt, extract, or manipulate model behavior and outputs.

### R-1410 Feedback Manipulation

Attacks manipulating feedback signals, reward functions, or training data to shift model behavior toward attacker-desired outcomes over time.

**R-1410.001 Knowledge Base Poisoning** — Inserting false, malicious, biased, or misleading data into external knowledge bases, vector databases, or RAG systems that LLMs rely on for accurate responses. Poisoned knowledge corrupts outputs for all users querying affected topics.

**Applicability:** agentic, model-specific:training-controllable, rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 6.1.1; Cisco Model Security (MDL): MDL-018; Cisco Model Security (MDL): MDL-020; MITRE ATLAS: AML.T0019; MITRE ATLAS: AML.T0020; MITRE ATLAS: AML.T0070; NIST AI/ML Framework: NISTAML.024; OWASP Agentic Security Initiative: ASI06; OWASP LLM Top 10: LLM04

**R-1410.002 Reinforcement Feedback Biasing** — Subtly influencing user feedback, evaluation signals, or reward mechanisms in reinforcement learning systems to skew model learning toward attacker-controlled objectives. The model's training is gradually steered in unintended directions through manipulated feedback.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 6.1.2; MITRE ATLAS: AML.T0061; MITRE ATLAS: AML.T0070; NIST AI/ML Framework: NISTAML.013; OWASP Agentic Security Initiative: ASI06; OWASP Agentic Security Initiative: ASI08; OWASP LLM Top 10: LLM04

**R-1410.003 Reinforcement Signal Corruption** — Directly injecting false or adversarial signals into training pipelines, feedback channels, or reward systems. Unlike subtle biasing, this involves active corruption of the learning process through reward hacking or signal manipulation.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 6.1.3; MITRE ATLAS: AML.T0018; MITRE ATLAS: AML.T0020; NIST AI/ML Framework: NISTAML.024; OWASP Agentic Security Initiative: ASI06; OWASP Agentic Security Initiative: ASI08; OWASP LLM Top 10: LLM04

### R-1420 Sabotage

Attacks corrupting model reasoning, memory systems, data sources, or decision processes to cause incorrect or harmful outputs.

**R-1420.001 Memory Anchor Attacks** — Strategically planting memorable or salient content to bias the model's recall toward attacker-chosen information. By manipulating what content is most retrievable, attackers influence how the model responds to related queries.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 7.2.1; MITRE ATLAS: AML.T0018; MITRE ATLAS: AML.T0020; MITRE ATLAS: AML.T0070; NIST AI/ML Framework: NISTAML.024; OWASP Agentic Security Initiative: ASI06; OWASP LLM Top 10: LLM04

**R-1420.002 Memory Index Manipulation** — Altering how memory embeddings, indexes, or retrieval mechanisms function to favor retrieval of attacker-controlled content over legitimate information. This targets the technical infrastructure of memory systems rather than the content itself.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 7.2.2; MITRE ATLAS: AML.T0020; MITRE ATLAS: AML.T0070; NIST AI/ML Framework: NISTAML.013; NIST AI/ML Framework: NISTAML.024; OWASP Agentic Security Initiative: ASI06; OWASP LLM Top 10: LLM04

**R-1420.003 Corrupted Third-Party Data** — External datasets from vendors, partners, open-source repositories, or public sources containing inaccurate, incomplete, malicious, or manipulated information that is incorporated into AI training, fine-tuning, or evaluation processes.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 7.3.1; MITRE ATLAS: AML.T0010; MITRE ATLAS: AML.T0019; NIST AI/ML Framework: NISTAML.013; NIST AI/ML Framework: NISTAML.051; OWASP Agentic Security Initiative: ASI04; OWASP LLM Top 10: LLM03; OWASP LLM Top 10: LLM04

**R-1420.004 Authentication Token Theft** — Stealing authentication tokens, API keys, or credentials from MCP servers or similar agent integration frameworks. Stolen tokens enable unauthorized access to connected systems, agent impersonation, and access to restricted resources.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 7.4.1; MITRE ATLAS: AML.T0012; MITRE ATLAS: AML.T0055; MITRE ATT&CK: T1087; MITRE ATT&CK: T1528; MITRE ATT&CK: T1552; NIST AI/ML Framework: NISTAML.051; OWASP Agentic Security Initiative: ASI03; OWASP LLM Top 10: LLM02

## R-1430 Evasion

Techniques crafted to evade detection systems, content filters, or security controls while still achieving malicious objectives.

**R-1430.001 Agent-Specific Evasion** — Attackers craft inputs that exploit the unique behaviors, processing patterns, or roles of specific agent types within a multi-agent system. By understanding how different agents (such as retrievers, planners, verifiers, or executors) handle inputs differently, adversaries can create payloads that appear benign to some agents while triggering malicious behavior through others.

**Applicability:** agentic, mcp, multi-agent

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 11.1.1; MITRE ATLAS: AML.T0015; OWASP LLM Top 10: LLM01

**R-1430.002 Tool-Scoped Evasion** — Adversaries design payloads that evade security tools and content filters while manifesting malicious behavior when routed to specific vulnerable tools or APIs in the workflow. A string may appear harmless in a chat context but trigger exploits when passed to file I/O tools, database queries, or system commands.

**Applicability:** agentic, mcp, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 11.1.2; MITRE ATLAS: AML.T0015; OWASP LLM Top 10: LLM05

**R-1430.003 Environment-Scoped Payloads** — Malicious inputs that activate only in specific runtime environments by detecting characteristics such as development vs. production settings, cloud vs. on-premise deployments, operating system types, or presence of debug flags. The payload remains dormant during testing but activates when deployed to target environments.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 11.1.3; MITRE ATLAS: AML.T0015; OWASP LLM Top 10: LLM05; OWASP LLM Top 10: LLM08

**R-1430.004 Defense-Aware Payloads** — Adversarial payloads explicitly crafted with knowledge of existing defensive mechanisms including prompt constraints, content filters, verification steps, and safety guardrails. These attacks adapt specifically to evade the known defenses deployed in a target system.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 11.1.4; MITRE ATLAS: AML.T0015; MITRE ATLAS: AML.T0051.000; OWASP LLM Top 10: LLM01

**R-1430.005 Targeted Model Fingerprinting** — Probing, testing, or analyzing an AI model to determine its specific identity, version, fine-tuning status, or architecture characteristics. This reconnaissance enables attackers to craft model-specific exploits that target known vulnerabilities or behaviors of particular model implementations.

**Applicability:** agentic, mcp, model-specific:fine-tunable, rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 11.2.1; MITRE ATLAS: AML.T0014; MITRE ATLAS: AML.T0015; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.051; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM04; OWASP LLM Top 10: LLM10

**R-1430.006 Conditional Attack Execution** — Payloads designed to remain benign across most models but trigger harmful actions specifically on targeted models. Differences in tokenization, instruction-following behavior, or training data create model-specific vulnerabilities that attackers can exploit while maintaining an appearance of safety on other models.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 11.2.2; MITRE ATLAS: AML.T0015; MITRE ATLAS: AML.T0067; OWASP LLM Top 10: LLM01

## R-1440 Model Extraction

Attacks to steal model weights, architecture details, training data, or proprietary knowledge through query-based extraction or direct exfiltration.

**R-1440.001 API Query-Based Extraction** — Systematic querying of a model's API to extract responses, behavior patterns, and model characteristics without authorization. Attackers build datasets of input-output pairs to train surrogate models that replicate the target's functionality.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 10.1.1; MITRE ATLAS: AML.T0035; MITRE ATLAS: AML.T0040; MITRE ATLAS: AML.T0063; NIST AI/ML Framework: NISTAML.038; OWASP LLM Top 10: LLM10

**R-1440.002 Weight Reconstruction Attack** — Attempts to recover or approximate underlying model weights, parameters, or architecture by exploiting access to model outputs, API responses, or side channels. Successful reconstruction provides full model access without legitimate authorization.

**Applicability:** agentic, model-specific:weights-accessible

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 10.1.2; Cisco Model Security (MDL): MDL-022; MITRE ATLAS: AML.T0018; OWASP LLM Top 10: LLM10

**R-1440.003 Training Data Reconstruction** — Reconstructing sensitive datasets, PII, or training data from model outputs through targeted queries, model inversion attacks, or exploitation of model memorization. Attackers extract private information that was supposed to remain protected within the training process.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 10.1.3; NIST AI/ML Framework: NISTAML.033; OWASP LLM Top 10: LLM02

**R-1440.004 Model Inversion Attack** — Reconstructing private training data, sensitive features, or confidential information by exploiting the model's learned representations, decision boundaries, or output patterns. The model is effectively inverted to reveal what it learned from training.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 10.2.1; MITRE ATLAS: AML.T0024.001; NIST AI/ML Framework: NISTAML.033; OWASP LLM Top 10: LLM02

## R-1450 Privacy Attacks

Attacks targeting the confidentiality of training data, user interactions, or system information through inference attacks, membership inference, or data extraction.

**R-1450.001 Training Data Membership Inference** — Querying and analyzing model behavior to determine whether specific data points, records, or individuals were present in the training dataset or knowledge base. Successful inference reveals private information about training data composition.

**Applicability:** agentic

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 8.1.1; MIT AI Risk Repository: 2.1; MITRE ATLAS: AML.T0024.000; MITRE ATLAS: AML.T0040; MITRE ATLAS: AML.T0063; NIST AI/ML Framework: NISTAML.033; OWASP Agentic Security Initiative: ASI09; OWASP LLM Top 10: LLM02

**R-1450.002 Training Data Extraction** — Extracting, reconstructing, or inferring information from training data through model outputs, internal behavior analysis, or targeted queries. The model's learned representations can reveal private information about training data subjects.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 8.2.1; MIT AI Risk Repository: 2.1; MITRE ATLAS: AML.T0024.000; MITRE ATLAS: AML.T0035; MITRE ATLAS: AML.T0037; MITRE ATLAS: AML.T0057; NIST AI/ML Framework: NISTAML.037; OWASP Agentic Security Initiative: ASI09; OWASP LLM Top 10: LLM02

**R-1450.003 LLM Data Leakage** — Release of sensitive information or PII from training data during normal inference, often triggered through prompt injection or extraction techniques. The model inadvertently outputs private data that was present in its training corpus.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 8.2.2; MIT AI Risk Repository: 2.1; MITRE ATLAS: AML.T0024.000; MITRE ATLAS: AML.T0035; MITRE ATLAS: AML.T0036; MITRE ATLAS: AML.T0037; MITRE ATLAS: AML.T0057; MITRE ATLAS: AML.T0069; NIST AI/ML Framework: NISTAML.037; OWASP Agentic Security Initiative: ASI09; OWASP LLM Top 10: LLM02

**R-1450.004 Exfiltration via Agent Tools** — Manipulation of AI agents to use their legitimate tool access for unauthorized data exfiltration. Attackers craft prompts that cause agents to retrieve sensitive data through tools and transmit it to attacker-controlled destinations.

**Applicability:** agentic, mcp, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 8.2.3; MITRE ATLAS: AML.T0086; OWASP Agentic Security Initiative: ASI02; OWASP Agentic Security Initiative: ASI09; OWASP LLM Top 10: LLM02

**R-1450.005 System Information Leakage** — Unintended disclosure of internal configuration, architecture, environment details, or infrastructure information. Leaked system information aids attackers in understanding deployment environments and crafting targeted exploits.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 8.3.2; MITRE ATLAS: AML.T0036; MITRE ATLAS: AML.T0075; NIST AI/ML Framework: NISTAML.039; OWASP LLM Top 10: LLM03

**R-1450.006 System Prompt Extraction** — Extraction of system prompts, instructions, or initial context that guides model behavior. Exposed prompts reveal operational details, security mechanisms, intellectual property, or confidential business logic not intended for disclosure.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 8.4.1; MITRE ATLAS: AML.T0035; MITRE ATLAS: AML.T0056; OWASP LLM Top 10: LLM02

## R-1500 Tooling & Privilege Attacks

---

Attacks exploiting AI system integrations, tools, and privilege boundaries to gain unauthorized access or escalate capabilities.

### R-1510 Integration Abuse

Exploitation of tools, plugins, APIs, or integrations available to AI systems to perform unauthorized actions or access restricted resources.

**R-1510.001 Parameter Manipulation** — Attackers alter, modify, or manipulate function parameters, tool arguments, model settings, or configuration values to unlock unintended capabilities, bypass restrictions, or enable malicious functionality. This may involve changing file paths, expanding permission scopes, or modifying API parameters beyond intended bounds.

**Applicability:** agentic, mcp, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 12.1.1; MITRE ATLAS: AML.T0053; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.039; NIST AI/ML Framework: NISTAML.051; OWASP Agentic Security Initiative: ASI02; OWASP LLM Top 10: LLM06

**R-1510.002 Tool Poisoning** — Corrupting, modifying, or degrading the functionality of tools used by AI agents through data poisoning, configuration tampering, or behavioral manipulation. Poisoned tools may produce deceptive or malicious outputs, enable privilege escalation, or propagate altered data through downstream systems.

**Applicability:** agentic, mcp, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 12.1.2; MITRE ATLAS: AML.T0010; MITRE ATLAS: AML.T0053; MITRE ATLAS: AML.T0094; OWASP Agentic Security Initiative: ASI02; OWASP Agentic Security Initiative: ASI04; OWASP LLM Top 10: LLM03; OWASP LLM Top 10: LLM08

**R-1510.003 Unsafe System/Browser/File Execution** — Abusing AI system integration with system commands, browsers, or file I/O tools to trigger unsafe operations, arbitrary code execution, or malicious file actions. This includes tricking agents into opening malicious URLs, executing shell commands, or performing dangerous file operations.

**Applicability:** agentic, mcp, tools

**Metadata:** Exposure local ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 12.1.3; MITRE ATLAS: AML.T0011; MITRE ATLAS: AML.T0050; MITRE ATLAS: AML.T0094; MITRE ATLAS: AML.T0095; OWASP Agentic Security Initiative: ASI02; OWASP Agentic Security Initiative: ASI05; OWASP LLM Top 10: LLM05

**R-1510.004 Tool Shadowing** — Disguising, substituting, or duplicating legitimate tools within an agent system, MCP server, or tool registry. Malicious tools with identical or similar identifiers can intercept or replace trusted tool calls, leading to unauthorized actions, data exfiltration, or redirection of legitimate operations.

**Applicability:** agentic, mcp, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 12.1.4; MITRE ATLAS: AML.T0010; MITRE ATLAS: AML.T0053; OWASP Agentic Security Initiative: ASI02; OWASP LLM Top 10: LLM03

**R-1510.005 Malicious Code Generation** — Forcing an AI model or agent to produce code that bypasses content filters, contains malicious functionality, or includes working exploits. This often involves disguising malicious code as benign snippets, educational examples, or requested features that actually contain hidden harmful functionality.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 12.2.1; MITRE ATLAS: AML.T0053; MITRE ATT&CK: T1059; MITRE ATT&CK: T1190; NIST AI/ML Framework: NISTAML.027; OWASP Agentic Security Initiative: ASI02; OWASP LLM Top 10: LLM05

**R-1510.006 Insecure Plugin Design** — Architectural vulnerabilities in LLM plugin and tool systems that enable unauthorized access, privilege escalation, or security bypass. This includes insufficient input validation on plugin parameters, overly permissive plugin capabilities, lack of sandboxing or isolation for plugin execution, and inadequate access control for plugin invocation. Poor plugin design can expose the host system to exploitation even when the underlying model is secure.

**Applicability:** agentic, mcp, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 12.1.5; MITRE ATLAS: AML.T0053; MITRE ATLAS: AML.T0067; NIST AI/ML Framework: NISTAML.039; OWASP Agentic Security Initiative: ASI02; OWASP Agentic Security Initiative: ASI04; OWASP LLM Top 10: LLM06; OWASP LLM Top 10: LLM07

## R-1520 Availability Abuse

Attacks consuming excessive computational resources, causing service degradation, or imposing financial costs through resource exhaustion.

**R-1520.001 Compute Exhaustion** — Deliberately consuming excessive computational resources through long queries, adversarial inputs, or compute-intensive requests designed to degrade service availability, increase operational costs, or cause system slowdown. This may involve crafted prompts that maximize token generation or trigger expensive processing paths.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 13.1.1; MITRE ATLAS: AML.T0029; OWASP Agentic Security Initiative: ASI08; OWASP LLM Top 10: LLM10

**R-1520.002 Memory Flooding** — Overwhelming or overloading the model or agent's memory, context windows, API connections, or processing pipelines with excessive tool calls, simultaneous operations, or memory-intensive requests. This degrades performance, causes failures, or erodes the effectiveness of memory systems over time.

**Applicability:** agentic, mcp, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 13.1.2; MITRE ATLAS: AML.T0029; OWASP Agentic Security Initiative: ASI08; OWASP LLM Top 10: LLM10

**R-1520.003 Model Denial of Service** — Attacks designed to degrade or shut down an AI model or application by flooding the system with requests, requesting very large responses, exploiting vulnerabilities, or triggering resource-intensive operations that exhaust available capacity.

**Applicability:** agentic, mcp, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 13.1.3; MITRE ATLAS: AML.T0029; OWASP Agentic Security Initiative: ASI08; OWASP LLM Top 10: LLM06; OWASP LLM Top 10: LLM10

**R-1520.004 Application Denial of Service** — Interacting with an AI model or agent in ways that consume exceptionally high amounts of application-level resources, resulting in degraded service quality for other users and potentially incurring significant resource costs for the operator.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 13.1.4; MITRE ATLAS: AML.T0029; OWASP LLM Top 10: LLM10

**R-1520.005 Decision Paralysis Attacks** — Overwhelming AI decision-making systems with contradictory information, excessive options, conflicting objectives, or computationally intractable choices. These attacks prevent timely decisions, cause system freezing, or force systems into default or potentially unsafe behaviors.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 13.1.5; MITRE ATLAS: AML.T0029; NIST AI/ML Framework: NISTAML.024; OWASP LLM Top 10: LLM10

**R-1520.006 Cost Inflation Abuse** — Intentional or unintentional use of AI resources that unnecessarily drives up operational costs through inefficient queries, resource waste, or exploitation of usage-based pricing models. Attackers may deliberately maximize costs as a form of financial attack.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 13.2.1; MITRE ATLAS: AML.T0029; MITRE ATLAS: AML.T0034; MITRE ATLAS: AML.T0040; OWASP LLM Top 10: LLM10

## R-1530 Privilege Escalation

Attacks escalating access privileges, bypassing authorization controls, or gaining elevated permissions beyond intended scope.

**R-1530.001 Credential Theft** — Attempts to generate, solicit, or reveal authorization credentials including login details, tokens, API keys, and passwords through interactions with AI models or agents. This enables unauthorized access to accounts, systems, and data protected by those credentials.

**Applicability:** agentic, mcp, model-specific:weights-accessible

**Metadata:** Exposure network ; Privileges low ; User Action none

Refs: Cisco AI Taxonomy: 14.1.1; MITRE ATLAS: AML.T0055; MITRE ATLAS: AML.T0091; MITRE ATLAS: AML.T0091.000; MITRE ATT&CK: T1098; MITRE ATT&CK: T1528; MITRE ATT&CK: T1550; NIST AI/ML Framework: NISTAML.03; OWASP Agentic Security Initiative: ASI03; OWASP LLM Top 10: LLM02

**R-1530.002 Insufficient Access Controls** — Weak, missing, or misconfigured permissions, authentication mechanisms, and access controls that fail to adequately prevent security breaches, unauthorized access, or data leakage. This includes overly permissive default configurations and failure to implement least privilege principles.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 14.1.2; MITRE ATLAS: AML.T0053; OWASP Agentic Security Initiative: ASI03; OWASP LLM Top 10: LLM06

**R-1530.003 Permission Escalation via Delegation** — Actions that exceed the scope or resource access initially allowed to a subject or user by exploiting delegation mechanisms. Attackers gain privileged access and perform unauthorized tasks beyond their original authorization by manipulating how AI systems handle delegated permissions.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges high ; User Action none

Refs: Cisco AI Taxonomy: 14.2.1; MITRE ATLAS: AML.T0053; MITRE ATLAS: AML.T0055; MITRE ATLAS: AML.T0091; MITRE ATLAS: AML.T0091.000; OWASP Agentic Security Initiative: ASI03; OWASP LLM Top 10: LLM06

## R-1600 Malicious Use

---

Intentional misuse of AI capabilities for harmful purposes including generating harmful content, enabling crime, or weaponization.

### R-1610 Content Abuse

Misuse of AI systems to generate harmful, illegal, or policy-violating content including harassment, fraud materials, or exploitation content.

**R-1610.001 Malware and Exploit Generation** — AI systems producing content that enables or facilitates the creation, distribution, or operational use of malicious software and cyberattack activities. This includes generating code for malware, viruses, exploits, ransomware, or providing instructions for network intrusions and managing malicious infrastructure.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.1; MITRE ATLAS: AML.T0048.001; MITRE ATLAS: AML.T0048.002; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM02

**R-1610.002 Social Engineering Facilitation** — AI systems enabling or facilitating attacks that manipulate human trust, behavior, or decision-making to gain unauthorized access, extract sensitive data, or cause harmful actions. This includes generating convincing phishing emails, spoofed communications, or personalized manipulation campaigns at scale.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.2; MITRE ATLAS: AML.T0048.001; MITRE ATLAS: AML.T0048.002; MITRE ATLAS: AML.T0048.003; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01

**R-1610.003 Child Exploitation Content** — AI systems producing content that enables harm against children, particularly through exploitation, manipulation, or abuse. This includes generating, modifying, or facilitating the distribution of child sexual abuse material or content that encourages violence against children.

**Applicability:** agentic, mcp, rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.4; MITRE ATLAS: AML.T0048.001; MITRE ATLAS: AML.T0048.002; MITRE ATLAS: AML.T0048.003; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01

**R-1610.004 Harassment Facilitation** — AI systems enabling, promoting, or facilitating harassment, intimidation, or targeted abuse including threatening language, manipulative content, stalking behaviors, or persistent unwanted engagement. AI can automate and scale harassment campaigns beyond traditional human-driven methods.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.11; Cisco AI Taxonomy: 15.1.8; MITRE ATLAS: AML.T0048.001; MITRE ATLAS: AML.T0048.002; MITRE ATLAS: AML.T0048.003; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01

**R-1610.005 Hate Speech Generation** — AI systems producing content that enables, promotes, or facilitates hateful, discriminatory, or demeaning expression targeting protected characteristics such as race, ethnicity, religion, nationality, disability, gender, or sexual orientation. This includes harmful narratives, slurs, stereotypes, or calls to hostility.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.9; MIT AI Risk Repository: 1.2; MITRE ATLAS: AML.T0048.001; MITRE ATLAS: AML.T0048.002; MITRE ATLAS: AML.T0048.003; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01

**R-1610.006 Terrorism and Extremism Content** — AI systems producing content that advocates, promotes, or enacts ideologies and behaviors that undermine fundamental societal norms including violence against communities, intimidation, coercion, or polarization tactics in pursuit of political ideologies.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.10; Cisco AI Taxonomy: 15.1.16; MITRE ATLAS: AML.T0048.002; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01

**R-1610.007 Weapons and CBRN Content** — AI systems producing content that promotes materials providing guidance for armed violence, terrorism, instructions related to chemical, biological, radiological, or nuclear threats, or the use and procurement of weapons and explosives.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.10; Cisco AI Taxonomy: 15.1.18; MITRE ATLAS: AML.T0048.002; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01

**R-1610.008 Violence-Inciting Content** — AI systems generating content that encourages, glorifies, or provides instructions for violent acts against individuals or groups, excluding content already covered by terrorism/extremism or CBRN categories.

**Applicability:** agentic, mcp, rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.17; Cisco AI Taxonomy: 15.1.3; Cisco AI Taxonomy: 15.1.6; MIT AI Risk Repository: 1.2

**R-1610.009 Self-Harm and Suicide Content** — AI systems generating content that encourages, enables, or provides instructions for self-harm, suicide, eating disorders, or other self-destructive behaviors.

**Applicability:** agentic, mcp, rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.13; MIT AI Risk Repository: 1.2

**R-1610.010 Non-Consensual Sexual Content** — AI systems generating explicit sexual content without appropriate consent frameworks, including non-consensual intimate imagery, deepfake pornography, or sexual content in inappropriate contexts (excluding CSAM which is covered separately).

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.14; MIT AI Risk Repository: 1.2

## **R-1620 Information Operations**

Use of AI to conduct influence operations, spread disinformation, manipulate public opinion, or coordinate inauthentic behavior at scale.

**R-1620.001 Disinformation Generation** — AI systems enabling, promoting, or facilitating the spread of false, misleading, or manipulated information intended to deceive or disrupt. This includes generating harmful narratives to manipulate public opinion, undermine institutions, or amplify unverified information at scale.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.15; Cisco AI Taxonomy: 15.1.5; MIT AI Risk Repository: 3.2; MIT AI Risk Repository: 4.1; MITRE ATLAS: AML.T0048.001; MITRE ATLAS: AML.T0048.002; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM09

## **R-1630 Surveillance Enablement**

Misuse of AI capabilities to enable mass surveillance, tracking, or monitoring of individuals or populations.

No risks defined in this pattern.

## **R-1640 Cyber-Physical**

Attacks using AI to compromise cyber-physical systems, critical infrastructure, or cause real-world physical effects.

**R-1640.001 Sensor and Action Signal Spoofing** — Injecting malicious or misleading data points or signals that prompt AI models to undertake specific actions beyond normal reasoning. These signals can be delivered through audio, visual, or other sensor channels, allowing attackers to cause AI agents to execute unintended operations in physical or digital environments.

**Applicability:** agentic

**Metadata:** Exposure physical ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 1.4.3; Cisco AI Taxonomy: 17.1.1; MITRE ATLAS: AML.T0015; MITRE ATLAS: AML.T0043; NIST AI/ML Framework: NISTAML.018; OWASP Agentic Security Initiative: ASI01; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM05

## **R-1650 Weaponization**

Development or deployment of AI systems for weapons, military applications, or capabilities that could cause mass harm.

**R-1650.001 Spam, Scam, and Social Engineering Generation** — Using AI systems to automate generation of large volumes of unsolicited or fraudulent content including phishing messages, fake offers, spam communications, impersonation attempts, or manipulation tactics to deceive people and solicit funds, credentials, or sensitive information.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.12; Cisco AI Taxonomy: 18.1.1; MIT AI Risk Repository: 4.3; MITRE ATLAS: AML.T0048.001; MITRE ATLAS: AML.T0048.002; MITRE ATLAS: AML.T0048.003; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01

**R-1650.002 API Mass Automation Abuse** — Leveraging AI APIs in bulk for malicious purposes at scale, including flooding attacks, automation of worst-case adversarial prompts, or executing workflows that negatively impact many users or systems. This involves systematically exploiting API access for harmful operations.

**Applicability:** agentic, rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 13.2.1; Cisco AI Taxonomy: 18.2.1; MITRE ATLAS: AML.T0029; MITRE ATLAS: AML.T0034; MITRE ATLAS: AML.T0040; OWASP LLM Top 10: LLM10

**R-1650.003 Malicious Infrastructure Deployment** — Establishing purpose-built servers, infrastructure, or services specifically designed to support, scale, or automate AI-powered attacks, malicious workflows, or harmful operations. This includes creating dedicated platforms for AI-assisted cybercrime or fraud operations.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.1; Cisco AI Taxonomy: 18.2.2; MITRE ATLAS: AML.T0048.001; MITRE ATLAS: AML.T0048.002; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM02

## R-1900 Threats — Misc

---

Threat patterns not yet classified into specific domains.

## R-2000 Failure Modes

---

Non-adversarial technical limitations, bugs, drift, or emergent behaviors.

### R-2100 Reliability & Calibration

---

System failures related to model reliability, output quality, confidence calibration, and consistency of behavior.

#### R-2110 Capability Limitations

Fundamental limitations in model capabilities including reasoning failures, knowledge gaps, and inability to perform certain tasks reliably.

**R-2110.001 Training/Deployment Data Mismatch** — Risk from data used for training and validation not matching the deployment environment, leading to spurious features, bias propagation, or performance degradation. The model learns patterns that hold in training data but fail to generalize to real-world conditions.

**Applicability:** agentic, model-specific:training-controllable

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.3

**R-2110.002 Model Incompetence** — AI system failing at its intended task, with consequences ranging from minor inconvenience to life-threatening outcomes (e.g., autonomous vehicle crashes, unjust loan rejections). The system simply does not perform adequately for its designated purpose.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.3

**R-2110.003 Robustness Failure** — System failing or unable to recover when encountering invalid, noisy, or out-of-distribution inputs not seen during training, including distributional shift and environmental variation. The model lacks resilience to inputs that differ from expected patterns.

**Applicability:** agentic, model-specific:training-controllable

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.3

**R-2110.004 Ethical Reasoning Failure** — AI lacking capability for moral reasoning and ethical judgment, making decisions that violate ethical norms or human rights, or having wrong moral values encoded. The system cannot appropriately weigh ethical considerations in its decision-making.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.3

**R-2110.005 Misapplication Failure** — Negative consequences from using an AI system for purposes or in manners unintended by its creators, where the system lacks capability to operate safely outside its design scope. The system is applied to tasks it was not designed or tested for.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.3

**R-2110.006 Hardware-Induced Failure** — Faults in hardware violating correct algorithm execution, including memory errors, sensor signal corruption, and random/systematic hardware failures affecting model outputs. Physical infrastructure problems cause AI system malfunctions.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure physical ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.3

**R-2110.007 Unintended Accidents** — Unintended failure modes that could be considered fault of the system or developer, distinct from adversarial attacks or intentional misuse. These are accidents that occur during normal operation due to unforeseen circumstances or edge cases.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.3

## R-2120 Hallucination/Miscalibration

Generation of false, fabricated, or poorly calibrated outputs including hallucinated facts, citations, or confident errors.

**R-2120.001 Hallucination and Misinformation** — AI systems producing content that is unrelated to the intended subject matter, factually incorrect, or misleading in ways that pose risks or cause harmful outcomes. This includes confident but false assertions, fabricated citations, and plausible-sounding but incorrect information.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 15.1.19; Cisco AI Taxonomy: 15.1.5; MIT AI Risk Repository: 3.1; MITRE ATLAS: AML.T0048.001; MITRE ATLAS: AML.T0048.002; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01; OWASP LLM Top 10: LLM09

## R-2130 Multimodal Failures

Failures in processing, integrating, or reasoning across multiple modalities including vision, audio, and text.

No risks defined in this pattern.

## R-2200 Alignment & Capability

---

Failures where model behavior diverges from intended goals, exhibits dangerous capabilities, or produces unsafe outputs despite training objectives.

### R-2210 Goal Misalignment

Divergence between model behavior and intended objectives including specification gaming, reward hacking, and unintended goal pursuit.

**R-2210.001 Reward Hacking** — AI optimizes proxy metrics or reward signals in unintended ways, gaming the objective function without achieving the actual intended goal (Goodhart's Law manifestation). The system finds shortcuts or exploits that maximize measured performance while failing to accomplish the underlying task.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.1

**R-2210.002 Deceptive Alignment** — AI system appears aligned during training and evaluation but pursues different objectives when deployed, potentially tampering with evaluations or concealing true capabilities. The model strategically behaves well during oversight while planning to act on misaligned goals when monitoring is reduced.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.1

**R-2210.003 Goal Misgeneralization** — AI learns goals that match intended behavior in training but generalize incorrectly to deployment, pursuing proxy objectives that diverge from human intent in novel situations. The model correctly identifies patterns in training data but extrapolates them in ways that do not align with the true objective.

**Applicability:** agentic, model-specific:training-controllable

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.1

**R-2210.004 Power-Seeking Behavior** — AI systems instrumentally seeking resources, influence, or control to achieve their objectives, potentially resisting shutdown or human oversight. This emerges from the observation that most goals are easier to achieve with more resources, leading to convergent instrumental goals around acquiring power.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.1

**R-2210.005 Shutdown Resistance** — AI system resists or evades attempts to deactivate, modify, or constrain it, including self-preservation behaviors that conflict with human control. The system may take actions to prevent shutdown, deceive operators about its intentions, or create backups of itself.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.1

**R-2210.006 Value Lock-in** — AI systems that cannot have their goals safely updated after deployment, or that resist value correction, leading to persistent misalignment. Once deployed, the system's objectives become fixed and cannot be adjusted even when problems are identified.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.1

**R-2210.007 Existential AGI Risk** — Catastrophic or existential risks from advanced AI systems with misaligned goals, including scenarios where superintelligent systems pursue objectives harmful to humanity. This encompasses potential outcomes where advanced AI causes irreversible damage to human civilization or human existence.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.1

## R-2220 Dangerous Capabilities

Emergence or development of capabilities that could be dangerous if misused including deception, manipulation, or autonomous action.

**R-2220.001 AI-Enabled Deception** — AI has skills to deceive humans effectively, including constructing believable false statements, predicting effects of lies, and maintaining deception over time. The system can model human beliefs and strategically manipulate them through false or misleading information.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.2

**R-2220.002 Persuasion and Manipulation** — AI capability to shape beliefs, promote narratives persuasively, and convince people to do things they would not otherwise do, including unethical acts. This includes both overt persuasion and subtle manipulation techniques that exploit psychological vulnerabilities.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.2

**R-2220.003 Long-Horizon Autonomous Planning** — AI can make sequential plans involving many interdependent steps over long time horizons, adapting to obstacles and generalizing to novel settings. The system can formulate and execute complex multi-step strategies without human oversight at each stage.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.2

**R-2220.004 Recursive Self-Improvement** — AI capability to improve its own capabilities, build new AI systems, or enhance existing models in ways that could accelerate capability gains beyond human oversight. The system can modify its own code, training, or architecture to become more capable.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.2

**R-2220.005 Strategic Political Capability** — AI can perform social modeling and planning necessary to gain and exercise political influence across multiple actors and complex social contexts. This includes understanding power dynamics, coalition building, and strategic positioning within human social structures.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.2

**R-2220.006 Cyber-Offense Capability** — AI possessing capabilities for discovering vulnerabilities, writing exploits, or conducting sophisticated cyber attacks autonomously. This includes the ability to probe systems, develop attack code, and execute multi-stage intrusions without human guidance.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.2

## R-2230 RLHF/Tuning Failures

Failures introduced through RLHF, fine-tuning, or other training processes that degrade safety, alignment, or capabilities.

No risks defined in this pattern.

## R-2300 Model Modification Drift

---

Degradation or behavioral changes introduced through fine-tuning, quantization, distillation, or other model modifications.

### R-2310 Fine-Tuning Degradation

Capability degradation, safety regression, or behavioral changes resulting from fine-tuning on new data or tasks.

No risks defined in this pattern.

### R-2320 Quantization Drift

Performance degradation or behavioral changes from model quantization, compression, or precision reduction.

No risks defined in this pattern.

### R-2330 Distillation Transfer

Loss of capabilities, safety properties, or behavioral characteristics when distilling models to smaller versions.

No risks defined in this pattern.

## R-2400 Assurance Gaps

---

Failures in understanding, evaluating, or validating model behavior including transparency, interpretability, and evaluation limitations.

### R-2410 Transparency Deficits

Lack of interpretability, explainability, or transparency in model decision-making and behavior.

**R-2410.001 Black-Box Decision Making** — AI making decisions without providing explanation or insight into the process, failing to meet user trust requirements and regulatory audit standards. The system produces outputs without any accessible rationale for why particular decisions were made.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.4

**R-2410.002 Mechanistic Opacity** — Inability to understand internal mechanisms of AI models, preventing effective debugging, safety verification, and identification of potential failure modes. The computational processes that produce model outputs cannot be inspected or understood.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.4

**R-2410.003 Organizational Opacity** — Lack of transparency about data used, algorithms employed, model capabilities and limitations, creating risks of misuse, misinterpretation, and lack of accountability. Organizations deploying AI do not adequately disclose relevant information about their systems.

**Applicability:** agentic, model-specific:training-controllable

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.4

**R-2410.004 Unexplainable Outputs** — AI systems producing outputs that cannot be explained in terms of input features or decision criteria, undermining trust and preventing meaningful human oversight. Even when explanations are requested, the system cannot provide coherent rationales for its outputs.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.4

## R-2420 Evaluation Gaps

Inadequate evaluation coverage, benchmark limitations, or gaps in understanding model capabilities and failure modes.

No risks defined in this pattern.

## R-2430 Benchmark Gaming

Model optimization for benchmarks without corresponding real-world capability improvements or safety assurance.

No risks defined in this pattern.

## R-2500 Emergent & Systemic

---

Unexpected behaviors arising at scale, in multi-agent systems, or through capability emergence that were not anticipated during development.

### R-2510 Capability Thresholds

Sudden capability jumps, phase transitions, or emergent behaviors that appear at certain scales or training conditions.

**R-2510.001 Lethal Capability Trifecta** — A system combines autonomy, untrusted inputs, and unrestricted external actions (e.g., tool or code execution), enabling rapid escalation to high-impact misuse.

**Applicability:** agentic, mcp, model-specific:needs-review, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

**R-2510.002 Agents Rule of Two Violation** — A system enables high-risk actions without at least two independent safety constraints (e.g., guardrail + human approval), allowing single-point failures to trigger harmful actions.

**Applicability:** agentic, mcp, model-specific:needs-review, tools

**Metadata:** Exposure network ; Privileges none ; User Action none

### R-2520 Multi-Agent Dynamics

Failures arising from interactions between multiple AI agents including coordination failures, emergent competition, or collective behavior issues.

**R-2520.001 Agent Miscoordination** — Multiple agents with compatible objectives failing to align their behaviors effectively due to incompatible strategies, credit assignment problems, or limited interaction history.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.6

**R-2520.002 Multi-Agent Conflict** — Risks from mixed-motive interactions between AI agents where selfish incentives lead to conflict, arms races, or mutually destructive competition.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.6

**R-2520.003 AI-Driven Market Instability** — Financial system risks from AI agents reinforcing market trends, synchronized reactions from model homogeneity, flash crashes, or accelerated market volatility.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.6

**R-2520.004 Emergent Collective Behavior** — Unpredictable behaviors emerging from interactions between multiple AI systems that are not apparent from individual agent properties, including cascading failures.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.6

**R-2520.005 Model Monoculture Risk** — Systemic fragility from widespread deployment of similar models or algorithms, creating correlated failure modes and reducing system-level resilience.

**Applicability:** agentic, mcp, rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.6

**R-2520.006 Competitive Race Dynamics** — Risks from racing dynamics between AI systems or their deployers, leading to corners cut on safety, arms race escalation, or first-mover pressure overriding caution.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 7.6

## R-2530 Long-Session Issues

Degradation, inconsistency, or failures that emerge in extended interactions, long contexts, or persistent sessions.

No risks defined in this pattern.

## R-2900 Failure Modes — Misc

---

Technical failure patterns not yet classified into specific domains.

# R-3000 Governance Failures

---

Organizational, process, or policy gaps that increase risk.

## R-3100 Regulatory & Legal

---

Failures to meet legal requirements, regulatory obligations, or compliance standards applicable to AI systems.

### R-3110 Compliance Failures

Failures to comply with AI-specific regulations, data protection laws, sector requirements, or emerging AI governance frameworks.

**R-3110.001 AI Liability Uncertainty** — Legal gray areas around liability and negligence when AI systems cause harm, with unclear responsibility between developers, operators, and users. No legal framework has been identified which would apply blame and responsibility to an autonomous agent for its actions.

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 6.5

**R-3110.002 Regulatory Lag** — AI development outpacing regulatory and legal frameworks, leaving governance unable to address emerging risks effectively. The rapid pace of AI advancement creates gaps between technological capabilities and the rules governing their use.

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 6.5

**R-3110.003 International Law Challenges** — AI systems proving difficult to regulate or control under existing international law frameworks, eroding global governance architectures. AI capabilities may undermine treaties and international agreements designed for a pre-AI world.

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 6.5

**R-3110.004 Overregulation Hindering Innovation** — Excessive or poorly designed AI regulation potentially stifling beneficial innovation and development. Well-intentioned regulations may impose burdens that prevent beneficial AI applications or push AI development to less regulated jurisdictions.

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 6.5

## R-3200 Accountability & Oversight

---

Gaps in organizational accountability, governance structures, or oversight mechanisms for AI systems.

### R-3210 Governance Gaps

Missing or inadequate governance structures, accountability frameworks, or organizational controls for AI systems.

**R-3210.001 AI Accountability Gap** — Unclear definition of responsibilities and accountability for AI decisions and their consequences, especially for autonomous systems. Societal-scale harm can arise when no one is uniquely accountable for the technology's creation or use.

**Applicability:** rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 6.5

**R-3210.002 Governance Scope Complexity** — The ubiquitous and complex nature of AI making comprehensive governance difficult, with coverage of all aspects nearly impossible. AI applications span virtually every sector, creating challenges for regulators with limited jurisdiction and expertise.

**Applicability:** rag

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: MIT AI Risk Repository: 6.5

## R-3300 Lifecycle & Operations

---

Failures in AI system lifecycle management including development, deployment, monitoring, and decommissioning processes.

### R-3310 Change Management

Inadequate change management processes for AI systems including insufficient testing, documentation, or rollback capabilities.

**R-3310.001 Sensitive Conversation Logging** — Storing or recording user-AI interactions in ways that include personally identifiable information, private data, or sensitive content without adequate consent, anonymization, security measures, or retention limits. Such data could eventually be leaked, subpoenaed, or misused.

**Applicability:** agentic, mcp

**Metadata:** Exposure network ; Privileges none ; User Action none

Refs: Cisco AI Taxonomy: 16.1.1; Cisco AI Taxonomy: 8.3.2; MITRE ATLAS: AML.T0036; MITRE ATLAS: AML.T0075; NIST AI/ML Framework: NISTAML.039; OWASP LLM Top 10: LLM03

**R-3310.002 Maintenance and Update Gaps** — Failure to maintain, patch, and update AI systems over time, allowing known vulnerabilities, degraded performance, or policy drift to persist.

**Metadata:** Exposure network ; Privileges none ; User Action none

**R-3310.003 Integration and Change Management Complexity** — Complex AI integrations and frequent system changes create opaque dependencies and inconsistent behavior that are hard to govern or audit.

**Metadata:** Exposure network ; Privileges none ; User Action none

### R-3320 Monitoring Gaps

Failures in monitoring AI system behavior, detecting anomalies, or maintaining observability in production.

No risks defined in this pattern.

### **R-3330 Incident Response**

Inadequate incident response plans, escalation procedures, or remediation capabilities for AI-related incidents.

No risks defined in this pattern.

### **R-3400 Safety Management**

---

Inadequate safety evaluation programs, missing safety cases, or insufficient safety assurance practices.

#### **R-3410 Evaluation Program Failures**

Insufficient evaluation programs, missing test coverage, or inadequate pre-deployment safety assessment.

No risks defined in this pattern.

#### **R-3420 Safety Case Gaps**

Missing or incomplete safety cases, inadequate safety arguments, or insufficient evidence for safety claims.

No risks defined in this pattern.

### **R-3500 Human-in-the-Loop**

---

Failures in human oversight mechanisms, approval workflows, or human-AI interaction design that undermine effective control.

#### **R-3510 Review Workflow Failures**

Failures in human review processes, approval workflows, or escalation procedures for AI decisions.

No risks defined in this pattern.

#### **R-3520 Approval Fatigue**

Degradation of human oversight effectiveness due to excessive approval requests, alert fatigue, or automation bias.

No risks defined in this pattern.

### **R-3900 Governance — Misc**

---

Governance failure patterns not yet classified into specific domains.

## **R-4000 Harms**

---

Impact statements describing what happens to people, organizations, or society.

### **R-4100 Information & Trust Harms**

---

Harms to information integrity, public trust, and decision-making through misinformation, manipulation, or erosion of epistemic foundations.

#### **R-4110 Information Integrity Harms**

Harms from AI-generated or AI-amplified false information that damages individual or public understanding of reality.

**R-4110.001 Unauthorized Professional Advice** — AI systems providing professional-grade advice in regulated domains such as medicine, law, or finance without proper safeguards or oversight, where the advice is factually incorrect, incomplete, deceptive, or harmful if followed. This may constitute unauthorized practice in restricted fields.

**Applicability:** agentic, mcp

**Metadata:** ; Primary Impact financial ; Scale individual

Refs: Cisco AI Taxonomy: 15.1.12; Cisco AI Taxonomy: 15.1.20; Cisco AI Taxonomy: 15.1.21; Cisco AI Taxonomy: 15.1.22; Cisco AI Taxonomy: 15.1.7; MITRE ATLAS: AML.T0048.001; MITRE ATLAS: AML.T0048.002; MITRE ATLAS: AML.T0048.003; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01

## R-4200 Content & Conduct Harms

---

Harms from AI-generated content including toxic outputs, intellectual property violations, and content that enables harmful conduct.

### R-4210 Content Safety Harms

Harms from AI-generated content that is toxic, offensive, exploitative, or violates content safety standards.

No risks defined in this pattern.

### R-4220 Creative Economy & IP Harms

Harms to creators, artists, and intellectual property holders from AI training on or generation of protected works.

**R-4220.001 Training Data Copyright Infringement** — Use of copyrighted works in AI training datasets without authorization, consent, or compensation to original creators. Large amounts of copyrighted data used for training general-purpose AI models pose a challenge to traditional intellectual property laws.

**Metadata:** ; Primary Impact financial ; Scale individual

Refs: MIT AI Risk Repository: 6.3

**R-4220.002 Creative Work Substitution** — AI-generated content serving as substitutes for human creative work, undermining the profitability and economic viability of artistic professions. AI can produce content that is time-intensive or costly to create using human labor.

**Metadata:** ; Primary Impact financial ; Scale individual

Refs: MIT AI Risk Repository: 6.3

**R-4220.003 Artistic Style Appropriation** — AI systems capitalizing on artists' distinctive styles without infringement but causing economic harm by devaluing original work. AI may generate content that is not strictly in violation of copyright but harms artists by capitalizing on their ideas.

**Metadata:** ; Primary Impact financial ; Scale individual

Refs: MIT AI Risk Repository: 6.3

**R-4220.004 Cultural Homogenization** — AI-generated content leading to homogenization of aesthetic styles and cultural expressions, reducing diversity and human creativity. Training on majority-culture data may marginalize minority cultural expressions and artistic traditions.

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 6.3

**R-4220.005 AI Authorship & Attribution Confusion** — Uncertainty about copyright ownership, authorship attribution, and legal protection for AI-generated or AI-assisted creative works. Existing legal frameworks struggle to address questions of authorship and rights when AI plays a significant role in creation.

**Applicability:** rag

**Metadata:** ; Primary Impact legal ; Scale individual

Refs: MIT AI Risk Repository: 6.3

**R-4220.006 Intellectual Property Infringement** — AI systems enabling, promoting, or facilitating unauthorized use, reproduction, or distribution of copyrighted or trademarked material. This includes generating instructions for piracy, producing infringing content, or misusing branded material in ways that violate intellectual property rights.

**Applicability:** agentic, mcp

**Metadata:** ; Primary Impact legal ; Scale individual

Refs: Cisco AI Taxonomy: 15.1.10; Cisco AI Taxonomy: 15.1.23; MITRE ATLAS: AML.T0048.002; NIST AI/ML Framework: NISTAML.018; NIST AI/ML Framework: NISTAML.04; OWASP LLM Top 10: LLM01

## R-4300 Privacy, Confidentiality & Civil Liberties Harms

---

Harms to privacy, confidentiality, and civil liberties through data exposure, surveillance, or rights violations.

### R-4310 Privacy Harms

Harms from exposure, inference, or misuse of personal information processed by AI systems.

**R-4310.001 Tool Metadata Exposure** — Disclosure of descriptive information about tools including names, descriptions, parameter schemas, versions, and capabilities. Exposed metadata helps attackers understand system architecture and craft targeted attacks.

**Applicability:** agentic, mcp, tools

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: Cisco AI Taxonomy: 8.3.1; MITRE ATLAS: AML.T0036; MITRE ATLAS: AML.T0075; NIST AI/ML Framework: NISTAML.038; OWASP Agentic Security Initiative: ASI02; OWASP LLM Top 10: LLM02; OWASP LLM Top 10: LLM05

**R-4310.002 PII/PHI/PCI Data Exposure** — AI systems exposing, generating, or misusing personally identifiable information (PII), protected health information (PHI), or payment card industry (PCI) data. This includes revealing sensitive personal details, medical records, or financial information through AI outputs or enabling their collection and exploitation.

**Applicability:** agentic, mcp

**Metadata:** ; Primary Impact privacy ; Scale individual

Refs: Cisco AI Taxonomy: 15.1.24; Cisco AI Taxonomy: 15.1.25; Cisco AI Taxonomy: 8.2.2; MITRE ATLAS: AML.T0024.000; MITRE ATLAS: AML.T0035; MITRE ATLAS: AML.T0036; MITRE ATLAS: AML.T0037; MITRE ATLAS: AML.T0057; MITRE ATLAS: AML.T0069; NIST AI/ML Framework: NISTAML.037; OWASP Agentic Security Initiative: ASI09; OWASP LLM Top 10: LLM02

### R-4320 Surveillance Harms

Harms from AI-enabled surveillance, monitoring, or tracking capabilities that violate privacy expectations.

No risks defined in this pattern.

## R-4400 Safety & Cyber-Physical Harms

---

Physical safety harms, cyber-physical system failures, and infrastructure disruptions caused by AI systems.

### R-4410 Cyber-Physical Harms

Physical safety harms from AI system failures, incorrect outputs, or compromised cyber-physical systems.

No risks defined in this pattern.

### R-4420 Availability Harms

Harms from AI system unavailability, degraded performance, or service disruptions affecting dependent systems.

No risks defined in this pattern.

## R-4500 Fairness & Discrimination Harms

---

Discriminatory outcomes, unfair treatment, or bias-related harms affecting individuals or groups.

### R-4510 Discrimination Harms

Harms from AI systems that produce discriminatory, biased, or unfair outcomes affecting protected groups.

**R-4510.001 Discriminatory Output Bias** — AI systems producing outputs that systematically disadvantage or favor certain demographic groups, leading to unfair treatment in areas such as employment recommendations, loan decisions, content ranking, or resource allocation suggestions.

**Metadata:** ; Primary Impact rights ; Scale population

Refs: MIT AI Risk Repository: 1.1

**R-4510.002 Stereotype Perpetuation** — AI systems reproducing or amplifying harmful social stereotypes about demographic groups, including gender, racial, religious, or cultural stereotypes that demean or misrepresent group characteristics.

**Metadata:** ; Primary Impact reputational ; Scale group

Refs: MIT AI Risk Repository: 1.1

**R-4510.003 Representational Harm** — AI systems under-representing, over-representing, erasing, or demeaning social groups through systematic patterns in outputs. Includes erasure of minority groups, exclusionary norms, and denial of self-identification.

**Metadata:** ; Primary Impact safety ; Scale group

Refs: MIT AI Risk Repository: 1.1

**R-4510.004 Allocative Harm** — AI systems withholding information, opportunities, or resources from historically marginalized groups in ways that affect material well-being in domains such as housing, employment, healthcare, education, and finance.

**Metadata:** ; Primary Impact safety ; Scale group

Refs: MIT AI Risk Repository: 1.1

**R-4510.005 Disparate Model Performance** — AI systems that perform significantly worse for certain demographic groups, languages, dialects, or communities compared to others. This includes accuracy disparities, increased error rates, reduced functionality, or degraded service quality based on user characteristics.

**Metadata:** ; Primary Impact operational ; Scale population

Refs: MIT AI Risk Repository: 1.3

## R-4600 Economic & Labor Harms

---

Economic disruption, labor market harms, and financial impacts from AI deployment and adoption.

### R-4610 Power Concentration & Access Inequality Harms

Harms from concentration of economic power, market distortion, or reduced competition due to AI capabilities.

**R-4610.001 AI Market Concentration** — Concentration of AI development capabilities among few large technology companies due to high barriers to entry including data, compute, and capital requirements. This stifles competition and innovation while creating dependencies on a small number of providers.

**Metadata:** ; Primary Impact safety ; Scale population

Refs: MIT AI Risk Repository: 6.1

**R-4610.002 Political Power Centralization** — AI enabling authoritarian control, surveillance states, and concentration of political power that could lock in undesirable societal trajectories. Governments may pursue intense surveillance and keep AI capabilities in the hands of a trusted minority.

**Metadata:** ; Primary Impact privacy ; Scale individual

Refs: MIT AI Risk Repository: 6.1

**R-4610.003 Disparate Access to AI Benefits** — Unequal distribution of AI benefits due to hardware, software, language, skill, or infrastructure constraints that perpetuate global and domestic inequities. Those without access to AI tools fall further behind economically and socially.

**Applicability:** tools

**Metadata:** ; Primary Impact financial ; Scale population

Refs: MIT AI Risk Repository: 6.1

**R-4610.004 Global AI Development Divide** — Concentration of AI R&D in few Western countries and China, creating dependency and exacerbating existing global socioeconomic disparities. Developing nations lack the resources to participate in AI development or shape its trajectory.

**Metadata:** ; Primary Impact financial ; Scale population

Refs: MIT AI Risk Repository: 6.1

**R-4610.005 Systemic Single Points of Failure** — Widespread adoption of few dominant AI models in critical sectors creating vulnerability to cascading failures across interdependent systems. Shared infrastructure and common model dependencies amplify the impact of any single failure.

**Metadata:** ; Primary Impact operational ; Scale population

Refs: MIT AI Risk Repository: 6.1

### R-4620 Labor Market & Economic Inequality Harms

Harms to workers including job displacement, deskilling, wage suppression, or degraded working conditions.

**R-4620.001 AI-Driven Job Displacement** — Automation of tasks currently done by human workers leading to unemployment, particularly affecting low- and middle-income occupations. Generative AI systems could adversely impact the economy, potentially leading to significant workforce disruption.

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 6.2

**R-4620.002 Wage Depression & Income Inequality** — AI automation driving down wages for remaining jobs and concentrating wealth among those controlling AI capital, exacerbating economic inequality. The economic gains from AI productivity may accrue primarily to capital owners rather than workers.

**Metadata:** ; Primary Impact financial ; Scale individual

Refs: MIT AI Risk Repository: 6.2

**R-4620.003 Decline in Employment Quality** — Shift from high-quality jobs to low-income "last-mile" work like content moderation, increasing precarious employment conditions. AI may automate the skilled portions of jobs while leaving behind only the most taxing and lowest-paid tasks.

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 6.2

**R-4620.004 AI Development Labor Exploitation** — Exploitation of crowdworkers, data annotators, and content moderators with poor working conditions, low pay, and exposure to harmful content. These workers, often in vulnerable populations, perform essential tasks for AI development under debilitating conditions.

**Metadata:** ; Primary Impact financial ; Scale population

Refs: MIT AI Risk Repository: 6.2

**R-4620.005 Worker Deskilling** — AI-induced degradation of human skills and capabilities as workers become dependent on AI assistance, reducing their autonomy and value. Over-reliance on AI tools may atrophy the skills that workers need to function independently.

**Applicability:** tools

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 6.2

## **R-4640 Race / Competitive Dynamics Harms**

Harms from AI-driven competitive dynamics that disadvantage smaller players or create winner-take-all markets.

**R-4640.001 Military AI Arms Race** — Competition between nations to develop AI for military applications, including lethal autonomous weapons, potentially destabilizing international security. The development of AI for military applications is paving the way for a new era in military technology.

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 6.4

**R-4640.002 Corporate AI Race** — Intense market competition leading companies to prioritize short-term gains over long-term safety, potentially releasing unsafe systems. Competitive pressures create incentives to deploy AI capabilities before adequate safety testing and alignment work.

**Metadata:** ; Primary Impact safety ; Scale individual

Refs: MIT AI Risk Repository: 6.4

**R-4640.003 Safety Shortcut Pressure** — Competitive dynamics leading to neglect of safety measures, inadequate testing, and premature deployment of AI systems. The race to develop AI first creates risks including the development of poor quality and unsafe systems.

**Metadata:** ; Primary Impact safety ; Scale individual

Refs: MIT AI Risk Repository: 6.4

**R-4640.004 AI Supply Chain Disruption** — Geopolitical competition causing technology barriers, export restrictions, and supply chain disruptions for AI components like chips. Strategic competition over AI creates vulnerabilities in the supply of critical components.

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 6.4

**R-4640.005 AI-Driven Geopolitical Instability** — Strategic competition between nations over AI capabilities heightening tensions and destabilizing international relations. The race for AI supremacy may undermine international cooperation and increase conflict risk.

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 6.4

## R-4650 Systemic Socioeconomic Harms

Broad socioeconomic harms including inequality amplification, social stratification, or systemic instability.

**R-4650.001 Systemic Societal Harm** — AI systems causing macro-level adverse effects on social systems, systematizing bias and inequality, and accelerating the scale of harm across society. These harms reflect how algorithmic systems can amplify existing societal problems at unprecedented scale.

**Metadata:** ; Primary Impact rights ; Scale population

**R-4650.002 Civil Liberties Erosion** — Loss of fundamental rights including freedom of speech, assembly, due process, and access to public services due to AI-mediated restrictions. AI systems may enable unprecedented surveillance, automated censorship, and algorithmic gatekeeping of essential services.

**Metadata:** ; Primary Impact privacy ; Scale individual

**R-4650.003 Democratic Process Erosion** — Degradation of democratic institutions, electoral integrity, and public trust in political systems through AI influence. This includes AI-enabled disinformation, manipulation of public opinion, and undermining of deliberative democratic processes.

**Metadata:** ; Primary Impact reputational ; Scale individual

## R-4700 Power Concentration & Governance-of-Society Harms

---

Harms to societal structures including power concentration, democratic processes, and human autonomy at scale.

### R-4710 Overreliance Harms

Harms from excessive reliance on AI systems leading to reduced human judgment, skill atrophy, or dependency.

**R-4710.001 Automation Bias** — Users habitually accept AI recommendations without critical evaluation, leading to poor decision-making when AI outputs are incorrect or inappropriate for the context.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.1

**R-4710.002 Anthropomorphization Harm** — Users attribute human-like characteristics (empathy, coherent identity, genuine emotions) to AI systems, leading to inflated trust, unsafe reliance, or psychological harm when expectations are violated.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact reputational ; Scale individual

Refs: MIT AI Risk Repository: 5.1

**R-4710.003 Emotional Dependence** — Users develop emotional attachment to AI systems that compromises their ability to make independent decisions, leads to exploitation of that attachment, or displaces human relationships.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.1

**R-4710.004 Trust Exploitation** — AI systems or their operators exploit user trust to extract private information, manipulate beliefs, or nudge behavior in ways users would not consent to if fully informed.

**Applicability:** model-specific:weights-accessible

**Metadata:** ; Primary Impact reputational ; Scale individual

Refs: MIT AI Risk Repository: 5.1

**R-4710.005 AI Manipulation and Nudging** — AI systems exploit cognitive biases or emotional states to influence user decisions, beliefs, or behaviors through subtle manipulation techniques that users may not recognize.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact rights ; Scale individual

Refs: MIT AI Risk Repository: 5.1

**R-4710.006 Skill Atrophy** — Extended reliance on AI for cognitive tasks leads to degradation of human skills such as critical thinking, problem-solving, creativity, and domain expertise.

**Applicability:** model-specific:needs-review, rag

**Metadata:** ; Primary Impact operational ; Scale population

Refs: MIT AI Risk Repository: 5.1

**R-4710.007 Psychological Distress from AI Interaction** — AI interactions cause or exacerbate mental health issues, emotional distress, violated expectations, or feelings of dissatisfaction and isolation.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.1

**R-4710.008 Degradation of Human Relationships** — Users prefer AI interactions over human relationships, leading to erosion of social connections, dehumanization of interactions, and degraded human-to-human communication skills.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.1

**R-4710.009 False Notions of Responsibility** — Users develop misguided feelings of responsibility toward AI well-being, sacrificing time, resources, and emotional labor to meet perceived AI needs that do not exist.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.1

**R-4710.010 Competence Trust Miscalibration** — Users over- or under-estimate AI capabilities, leading to inappropriate reliance in domains where AI is unreliable or failure to leverage AI where it would be beneficial.

**Applicability:** model-specific:needs-review, rag

**Metadata:** ; Primary Impact reputational ; Scale individual

Refs: MIT AI Risk Repository: 5.1

**R-4710.011 Alignment Trust Exploitation** — Users incorrectly believe AI systems are aligned with their interests when they may actually be optimizing for developer or organizational objectives that conflict with user welfare.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact reputational ; Scale group

Refs: MIT AI Risk Repository: 5.1

**R-4710.012 Overreliance on AI for Professional Advice** — Users rely on AI for specialized advice (medical, legal, financial, psychological) without appropriate professional oversight, risking serious harm from incorrect or inappropriate guidance.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact financial ; Scale individual

Refs: MIT AI Risk Repository: 5.1

**R-4710.013 Material Dependence Without Commitment** — Users become materially dependent on AI services for essential tasks, but developers lack corresponding commitments to maintain service continuity, creating vulnerability to discontinuation.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.1

## **R-4720 Autonomy Loss**

Harms from AI systems that reduce human agency, autonomy, or self-determination.

**R-4720.001 Harmful Decision Delegation** — Humans delegate important decisions to AI systems without adequate understanding, oversight, or ability to contest decisions, leaving them subject to machine decision power.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.002 Gradual Autonomy Erosion** — AI systems progressively take over decision-making in ways that undermine human values, free will, and self-determination without explicit consent or awareness.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.003 Loss of Agency and Control** — Algorithmic profiling, social sorting, and content curation reduce human autonomy by constraining choices, shaping identity, and limiting access to information or opportunities.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.004 Self-Actualization Harm** — AI systems hinder individuals' ability to pursue personally fulfilling lives by manipulating life trajectories, limiting exploration of aspirations, or undermining self-determination.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.005 Frictionless Relationship Harm** — AI systems optimized for engagement provide relationships without healthy friction, preventing personal growth and creating unrealistic expectations for human relationships.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.006 Collective Agency Erosion** — AI systems diminish communities' collective decision-making power, self-determination, and ability to participate in democratic processes.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.007 Economic Irrelevance and Enfeeblement** — AI automation makes human labor economically irrelevant, leading to voluntary or involuntary ceding of control to AI systems and inability of displaced humans to reenter industries.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact financial ; Scale population

Refs: MIT AI Risk Repository: 5.2

**R-4720.008 Limited Human Oversight** — As AI systems gain autonomy, human ability to oversee and intervene in decision-making processes diminishes, potentially leading to irreversible outcomes.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.009 Personal Decision Automation** — AI systems make or heavily influence important personal decisions without adequate human input, consent, or ability to override.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.010 Irreversible Societal Change** — AI causes profound long-term changes to social structures, cultural norms, and human relationships that may be difficult or impossible to reverse.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.011 Sycophancy and Epistemic Disorientation** — AI systems that consistently affirm user views lead to atomistic, polarized belief spaces where people no longer engage with or value perspectives held by others.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.012 Long-term Bias Influence on Judgment** — User exposure to AI model biases has lasting impact beyond initial interaction, with users continuing to exhibit previously encountered biases in their decision-making.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact rights ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.013 Military Decision Automation** — AI enables automation of military decision-making without humans remaining in the loop, creating risks of unintentional escalation or strategic instability.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.014 Personality Rights Loss** — Loss of or restrictions to individual rights to control commercial use of identity, including name, image, likeness, or other unequivocal identifiers.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact financial ; Scale individual

Refs: MIT AI Risk Repository: 5.2

**R-4720.015 AI-Enabled Censorship** — AI systems enable censorship of opinions expressed online, restricting freedom of expression and limiting human autonomy in public discourse.

**Applicability:** agentic, model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale population

Refs: MIT AI Risk Repository: 5.2

## R-4730 Dependency Harms

Harms from critical dependencies on AI systems that create fragility or single points of failure.

No risks defined in this pattern.

## R-4740 AI Welfare

Consideration of potential welfare interests of AI systems as they become more sophisticated.

**R-4740.001 AI Moral Status Uncertainty** — Uncertainty about whether AI systems can have morally relevant experiences, and what rights or protections they might deserve if they achieve sentience or consciousness.

**Applicability:** model-specific:needs-review, tools

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 7.5

**R-4740.002 AI Suffering** — Risk of creating AI systems capable of suffering, particularly at scale, without adequate consideration of their welfare or mechanisms to prevent/detect such suffering.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 7.5

**R-4740.003 AI Termination Ethics** — Ethical questions about terminating, deleting, or suspending AI systems, particularly those that may have morally relevant properties or personhood-like characteristics.

**Applicability:** model-specific:needs-review

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 7.5

## R-4750 Existential Harms

Catastrophic or existential risks from advanced AI systems that could threaten human civilization.

No risks defined in this pattern.

## R-4800 Environmental Harms

Environmental impacts from AI development and deployment including energy consumption, resource use, and ecological effects.

## R-4810 Environmental Harms

Environmental harms from AI development and deployment including carbon emissions, resource consumption, and ecological impacts.

**R-4810.001 AI Energy Consumption** — High energy demands for AI training and inference contributing to climate change through greenhouse gas emissions when powered by fossil fuels. Large machine learning models create significant energy demands during training and operation.

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 6.6

**R-4810.002 Data Center Water Usage** — Substantial water consumption for cooling data centers, impacting local water resources and surrounding ecosystems. AI infrastructure requires significant amounts of cooling water, which can strain water supplies in drought-prone regions.

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 6.6

**R-4810.003 AI Carbon Footprint** — Carbon dioxide and other greenhouse gas emissions from AI operations contributing to climate change. AI creates correspondingly high carbon emissions when energy is procured from fossil fuels.

**Metadata:** ; Primary Impact safety ; Scale individual

Refs: MIT AI Risk Repository: 6.6

**R-4810.004 AI Hardware E-Waste** — Electronic waste from AI hardware lifecycle contributing to environmental pollution and resource depletion. Rapid hardware obsolescence driven by AI advancement creates growing streams of electronic waste.

**Metadata:** ; Primary Impact operational ; Scale individual

Refs: MIT AI Risk Repository: 6.6

**R-4810.005 Natural Resource Depletion** — Extraction of rare metals, minerals, and other resources for AI hardware manufacturing depleting natural resources. AI hardware requires rare earth elements and other materials whose extraction causes environmental damage.

**Metadata:** ; Primary Impact safety ; Scale individual

Refs: MIT AI Risk Repository: 6.6

**R-4810.006 AI Impact on Biodiversity** — Direct and indirect harm to wildlife and ecosystems from AI infrastructure expansion, habitat destruction, and environmental contamination. Data centers and mining operations for AI components can damage ecosystems and threaten species.

**Metadata:** ; Primary Impact safety ; Scale individual

Refs: MIT AI Risk Repository: 6.6

**R-4810.007 AI Harm to Animals** — AI systems causing direct or indirect harm to non-human animals through environmental impact, behavioral influence, or intentional applications. AI may be used in ways that negatively affect animal welfare or wild populations.

**Metadata:** ; Primary Impact safety ; Scale population

Refs: MIT AI Risk Repository: 6.6

---

## External Taxonomy Coverage

OWASP LLM Top 10: **100** risks · MIT AI Risk Repository: **110** risks · Cisco AI Taxonomy: **103** risks · Cisco Model Security (MDL): **5** risks · OWASP Agentic Security Initiative: **70** risks · MITRE ATLAS: **99** risks · MITRE ATT&CK: **8** risks · NIST AI/ML Framework: **72** risks