RESOURCE · PROCUREMENT · RES-005

AI MSA Red-Team Checklist

VERSION

v1.0

TRACKED AT

agentmodeai.com/holding/RES-005/

SOURCE

agentmodeai.com/resources/

LICENCE

CC BY 4.0

Agent Mode AI publishes editorial work for senior IT leaders. Every claim including the premise of this template is tracked on a public ledger. Verdict updates and corrections at agentmodeai.com/holding/.

Version 1.0 · Published 4 May 2026 · 60-day review cycle

A 38-item checklist for reviewing AI vendor Master Service Agreements, Data Processing Addenda, and AI-specific addenda. Built for a working session between procurement and legal counsel reviewing actual vendor paper.

Source explainer: https://agentmodeai.com/resources/ai-msa-red-team-checklist/ Holding-up tracking: https://agentmodeai.com/holding/RES-005/

How to use

Use AFTER the AI Vendor Security Questionnaire (RES-001) is complete. The questionnaire surfaces what the vendor does; the MSA review surfaces what the vendor has contractually committed to. Discrepancies between the two are the most informative signal in procurement.

Each item is a yes/no question. Mark each Y/N + capture the contract clause reference + note whether language is acceptable, treatable (redline), or non-negotiable.

Clause family 1: Training-data carve-outs (6 items)

#	Question	Y/N	Clause ref	Notes
1.1	Contract explicitly prohibits using customer prompts for model training, fine-tuning, evaluation, or improvement			
1.2	Prohibition extends to completions, embeddings, intermediate agent state, tool-use traces			
1.3	Prohibition extends to upstream model providers (Anthropic, OpenAI, etc.)			
1.4	Technical mechanism enforcing prohibition is named OR referenced in linked technical document			
1.5	Vendor will certify on request that no customer data has been used in training			
1.6	Audit mechanism specified for verifying prohibition has been honoured			

Acceptable: explicit prohibition + named technical mechanism + audit right. **Unacceptable:** "aggregated", "anonymised", "improvement", "service quality", "with appropriate privacy protections" without operational specificity.

Clause family 2: Output ownership + IP indemnification (5 items)

#	Question	Y/N	Clause ref	Notes
2.1	Customer owns outputs generated by AI system using customer's prompts			
2.2	Vendor indemnifies customer against IP infringement claims arising from model outputs			
2.3	Indemnification cap consistent with customer's potential exposure (contract value + multiplier, not flat)			
2.4	Indemnification covers original output as generated (not carved out for customer modifications)			
2.5	Vendor warrants model trained on data the vendor had the right to use (burden not on customer)			

Acceptable: customer owns outputs + vendor indemnifies + warranty on training-data rights. **Unacceptable:** "to the extent permitted by law" + "subject to applicable third-party rights" without naming what those rights are.

Clause family 3: Model-deprecation + version-change rights (5 items)

#	Question	Y/N	Clause ref	Notes
3.1	Notice period for model version changes specified			
3.2	Notice period long enough for testing (typical minimum: 90 days major, 30 days minor)			
3.3	Vendor obligation to maintain prior model version during transition window			
3.4	Customer rights if new model materially degrades performance against original use case			
3.5	Migration support specified (test environments, prompt re-tuning, A/B comparison tooling)			

Acceptable: named notice period + dual-version maintenance + degradation remedies. **Unacceptable:** vendor right to "update, modify, or discontinue" without notice or consent.

Clause family 4: Sub-processor expansion (5 items)

#	Question	Y/N	Clause ref	Notes
4.1	Named sub-processor list (not categories)			
4.2	Notification of sub-processor additions with SLA			
4.3	Customer right to object with meaningful remedy (termination right OR carve-out)			
4.4	Sub-processors restricted to customer-approved geographies			
4.5	Vendor flows down customer's data protection terms to all sub-processors			

Acceptable: named-list disclosure + notification SLA + meaningful objection right. **Unacceptable:** "categories" of sub-processors without naming them.

Clause family 5: Kill-switch operability + SLA (5 items)

#	Question	Y/N	Clause ref	Notes
5.1	Kill-switch SLA defined in seconds			
5.2	Customer-side party authorised to invoke kill-switch identified			
5.3	Emergency kill-switch invocation procedure outside business hours specified			
5.4	Vendor obligations during kill-switch period (preserve forensic state, halt billing, RCA)			
5.5	Customer remedies if SLA missed (termination right + liquidated damages, not just service credits)			

Acceptable: seconds-defined SLA + named customer party + meaningful remedies. **Unacceptable:** "best efforts" + SLA in hours/business days + service credits as sole remedy.

Clause family 6: Exit-data portability (6 items)

#	Question	Y/N	Clause ref	Notes
6.1	Export format specified (structured, machine-readable, schema documented)			
6.2	Export SLA specified (delivery time from customer request)			
6.3	Retention window after termination specified (long enough for verification)			
6.4	Deletion certification at end of retention window specified			
6.5	Export of derivatives (embeddings, fine-tuning datasets, evaluation traces) included			
6.6	Export of audit logs + retention-obligation documentation included			

Acceptable: structured machine-readable + SLA + verification window + deletion certification + derivative coverage. **Unacceptable:** "reasonable assistance" + formats "selected by the vendor".

Clause family 7: Regulatory cooperation + EU AI Act flow-through (6 items)

#	Question	Y/N	Clause ref	Notes
7.1	Vendor will provide Article 26 deployer documentation on request			
7.2	Vendor will cooperate with customer DPIAs under GDPR Article 35			
7.3	Vendor will notify customer of regulatory actions, supervisory authority investigations, class-action litigation			
7.4	Vendor will notify customer of model changes affecting EU AI Act risk classification			
7.5	Vendor will provide Annex IV technical documentation on request			
7.6	Customer rights specified if vendor product loses EU AI Act conformity assessment after signing			

Acceptable: explicit cooperation commitments on each item. **Unacceptable:** silence on these points.

Scoring

- **30+ yes** — vendor is contractually serious. Proceed.
- **20-29 yes** — treatable. Each gap = a redline for negotiation.
- **<20 yes** — vendor's commercial position depends on retaining the rights this checklist constrains. Either renegotiate substantially or walk.

The checklist deliberately does not weight items. Weighting depends on deployment scope: kill-switch SLA matters more for payment-system-access agents than for retrieval agents. Apply weighting in your procurement context.

Frameworks referenced

- EU AI Act Articles 26, 27, 79; Annex III; Annex IV
- GDPR Articles 6, 22, 28 (sub-processors), 35 (DPIA), 44-49 (transfers)
- Cloud Security Alliance CCM v4 (Cloud Controls Matrix, 2024)
- Observable 2024-2025 enterprise AI procurement disputes

Licence