



REWRITING VENDOR CONTRACTS FOR THE AI AGE

A five-part guide to allocating risk in contracts for AI-based products



INTRODUCTION

A decade ago, signing an agreement with a SaaS vendor meant contracting for a defined product, a stable architecture and a largely predictable risk profile. However complex the product, its behavior was broadly understood when the contract was signed.

AI is changing that.

AI systems are probabilistic, not deterministic. They evolve through retraining, depend on changing data, and may produce outputs that were neither programmed nor foreseeable at signing. Standard SLA credits, fee-based liability caps and output disclaimers in vendor contracts are typically built for a different kind of product. They offer little protection when a system learns, adapts, and generates results with legal or financial consequences that nobody anticipates.

This five-part guide examines where traditional technology contracts break down when used for AI-based products. It uses a single illustrative fact pattern set in Indian e-commerce and marketplace retail. Each section covers a practical scenario, the clause that failed, and a drafting fix.

The aim is to propose principles for re-engineering vendor contracts for the AI age.

Sections in this guide

- **AI output risk:** Who should bear the cost of a harmful automated decision?
- **Regulatory exposure and audit rights:** Can the customer obtain the evidence a regulator asks for?
- **IP ownership and fine-tuning data:** Who owns the value created from customer data?
- **M&A and assignment risk:** What happens when the AI vendor is acquired by a competitor?
- **Operational risk and remedies:** Do SLA credits and liability caps reflect the real harm?

A note on the illustration

Orenda Commerce Private Limited (Orenda) is a fictional Indian e-commerce marketplace. It lists roughly 90 lakh SKUs from about 78,000 sellers and runs its own fulfilment network. All companies, products, figures and incidents in this guide are invented for illustration.

Section I - AI Output Risk

Overview

- Orenda signed a three-year MSA with Halcyon Labs. Halcyon supplied an AI catalogue integrity and seller-risk platform called Sentra.
- Sentra screens listings, scores sellers for counterfeit and quality risk, and applies enforcement action automatically. It takes about 1.1 lakh enforcement decisions a day. None are reviewed by a person before they take effect.
- Four months into the arrangement, Sentra classified Lumen Living, a premium home furnishing seller, as a repeat counterfeit offender. The classification came from a mismatched brand authorization record. It was wrong.
- Sentra suppressed the seller's listings and published an 'authenticity under review' banner on 1,900 live product pages. The banner stayed up for 11 days, across a festive sale window.
- Lumen Living sent Orenda a legal notice. It claimed Rs. 4.5 crore for defamation and lost sales.

Halcyon's response: It relied on the following clause of the MSA to shift liability to Orenda: 'Orenda assumes sole responsibility for reviewing, validating and acting upon any output generated by the Platform.'

Failure point: The clause assumed that Orenda could review every output before it caused harm. At 1.1 lakh automated enforcement decisions a day, that assumption was incorrect. It was also the wrong assumption commercially: Orenda bought the platform precisely to remove human review from the loop. A liability clause must work in the real setting in which the AI system is used. This one did not.

What the MSA should say

- The vendor must implement and maintain accuracy controls, evidentiary thresholds and output guardrails. These controls must be designed to prevent baseless, discriminatory, defamatory or otherwise unlawful classifications and public-facing statements.
- The vendor must indemnify Orenda against third-party claims arising from outputs that reasonable guardrails should have caught.
- Orenda is responsible for outputs generated within, and in compliance with, the vendor's documented guardrail specifications.
- Where an output triggers an irreversible or publicly visible action, the contract should require a confidence threshold, a human review gate, or both. The parties should agree which action classes fall on which side of that line.
- The vendor must give at least 30 days' prior written notice of any change to its guardrails, confidence thresholds, decision thresholds or model version.
- The vendor must support rapid reversal: a defined rollback mechanism, a maximum time to reverse an erroneous action, and preservation of the record of the reversed decision.

Section II - Regulatory Exposure & Audit Rights

Overview

- A seller association complained to the consumer protection regulator about automated delisting, suppression and ranking on large marketplaces. Orenda was named.
- The regulator opened a formal inquiry. It asked Orenda for three sets of records:
 - (a) All automated seller enforcement decisions for the previous 18 months.
 - (b) The decision logic, model specifications, accuracy benchmarks and measured error rates of every AI system used in enforcement and ranking.
 - (c) Evidence of the human oversight applied to automated enforcement actions.
- Orenda could produce the decision records. It could not produce the rest.

Halcyon's response: The MSA gave Orenda access to Halcyon's systems and processes as they relate to the processing of Orenda's data. It did not give Orenda access to the AI model itself. Orenda had never secured rights to inspect the model specifications or decision logic.

Failure point: The inquiry sought three things: decision records, decision logic and proof of human oversight. Orenda could provide only the first.

The problem lay in the audit clause. It had been drafted for a data center audit. It focused on what data the vendor held and how that data was processed. It did not account for an AI system, where processing involves a probabilistic inference engine that may not be fully explainable, even by the vendor.

An audit right that covers the data store but not the decision engine is incomplete. The absence of a contractual right to access this information is not a defence before a regulator.

There was a second gap. Halcyon had no obligation to help Orenda respond to an inquiry arising from Sentra. Orenda had 21 days to respond, but no contractual right to compel the vendor's cooperation.

What the MSA should say

- Audit rights must expressly cover the model's decision logic, configuration, system prompts, thresholds and output guardrail specifications.
- They must cover accuracy benchmarks, testing records, measured error rates and filtering documentation.
- The vendor must maintain a version history and change log for the model, its configuration and its guardrails, and must retain that record for a defined period after termination.
- The vendor must maintain enough model documentation to meet applicable Indian regulatory requirements.
- The vendor must help Orenda respond to regulatory inquiries arising from the platform. The clause should set clear response times, calibrated to the shortest statutory deadline Orenda realistically faces.
- Orenda must be allowed to conduct an independent technical audit at least once a year, assessing accuracy, bias and guardrail effectiveness, with reasonable cooperation from the vendor. The parties should agree who bears the cost, and that a failed audit shifts that cost to the vendor.

Section III - IP Ownership & Fine-Tuning Data Rights

Overview

- Orenda separately deployed Atlas, a returns and claims adjudication model supplied by Quillon AI.
- To make Atlas work for its categories, Orenda's risk team built a proprietary fine-tuning dataset. It contained 9.6 lakh adjudicated return and claim files. Each file was annotated with the fraud signal detected, the category reason code, the adjudication outcome, the seller liability finding and the appeal result.
- The project cost Orenda an estimated Rs. 6.2 crore. Its purpose was to make Atlas more accurate for Orenda's business.
- Six weeks after completion, Quillon sent a formal letter claiming a right to use the fine-tuning dataset to train models across its client base.

Quillon's response: The MSA allowed it to use 'customer data' and transaction logs to improve the platform and its services generally. The fine-tuning dataset was, on its reading, a transaction log over which Quillon could exercise that right.

Failure point: A single AI deployment creates at least four distinct layers of value:

- Orenda input data.
- Model architecture and what the model learns.
- Outputs generated by the model.
- Feedback and telemetry, meaning technical and usage data collected automatically while the system operates.

The MSA did not clearly allocate the IP created at each layer. Its broad language allowed Quillon to treat the fine-tuning dataset as customer data or a transaction log. It could then use Orenda's investment to improve a broader model, including one deployed for Orenda's competitors.

What the MSA should say

- Define four separate categories: Customer Input Data, Model Architecture and Weights, Outputs, and Feedback and Telemetry.
- Orenda Input Data must remain exclusively owned by Orenda. Model Architecture and Weights may remain with the vendor unless the parties agree otherwise.
- Rights in outputs should be assigned to Orenda, subject to applicable law.
- Feedback and Telemetry must be narrowly defined so that they cannot become a route for Orenda IP to leak into the vendor's wider models.
- A fine-tuning dataset compiled by Orenda must be treated as Orenda IP, regardless of the infrastructure on which it was processed. Annotation and labelling work should be expressly captured, since that is where most of the effort is spent.
- The contract must expressly prohibit the use of customer data, including transaction and decision logs, to train or improve foundation models or generalised models made available to third parties.
- After termination, the vendor must delete all customer input data, embeddings, fine-tuning outputs and derivative model weights, and certify deletion in writing within 30 days.
- Any IP created from Orenda data in breach of the clause must automatically be assigned to Orenda.

Section IV - M&A Risk & Assignment

Overview

- Halcyon announced that it was being acquired by Northwind Retail Technologies.
- Northwind is a Chennai-based retail technology company. It supplies catalogue and seller-risk tooling to Orenda's direct competitors.
- The MSA allowed Halcyon to assign the agreement without Orenda's consent where an entity acquired all or substantially all of Halcyon's business or assets.
- The agreement contained no other restrictions linked to an acquisition. It did not require data segregation or deletion at the time of the deal.
- Northwind would now have access to Orenda's seller risk scores, counterfeit patterns, category-level enforcement thresholds, high-return postcodes and festive demand curves that could be shared with Orenda's competitors.
- Orenda could switch vendors. But the move would cost between Rs. 22 crore and Rs. 31 crore and take five to six months.

Where the contract failed

- The assignment was contractually valid.
- Nothing stopped Northwind, which now controlled Orenda's complete enforcement and seller-risk dataset, from using that intelligence to improve products sold to Orenda's competitors.
- The standard M&A assignment carve-out assumes the acquirer is neutral: a private equity fund, or a strategic buyer in an unrelated sector. That assumption is dangerous. An AI vendor's value often lies in the data and model intelligence accumulated from customers like Orenda. That same value is what attracts the buyer, and the likely buyer is therefore someone in or adjacent to the customer's own market.
- The contract also missed two practical complications:
 1. Switching would cost Rs. 22 crores to Rs. 31 crores and take five to six months. Orenda could not exit quickly, and Northwind knew that.
 2. The AI supply chain kept running. Every enforcement decision taken after the announcement generated more data flowing into Northwind's ecosystem. The clock was running.

What the MSA should say

- A change of control must require prior written consent where the acquirer is a competitor or provides services to the customer's competitors. Competitors may be defined by named sectors, or by a list reviewed each year.
- Orenda has an unconditional right to terminate within 30 days after a qualifying change of control is announced.
- Data segregation duties must apply immediately on announcement, whether or not Orenda terminates: logical separation, access restriction to the incumbent team, and a bar on cross-use for the acquirer's other products.
- Exit assistance must be a standing obligation with agreed scope, duration and pricing, so that the cost of leaving does not become the vendor's leverage.

Section V - Operational Risk & Remedy Design

Overview

- Orenda's fulfilment arm signed a five-year MSA with Beacon Systems for an AI fulfilment orchestration platform called Torent. The annual contract value was Rs. 14 crores.
- Torent handled slotting, dispatch sequencing, carrier allocation and settlement across 12 fulfilment centres and Orenda's 60 largest brand partners.
- At 02:40 on the second Saturday of a festive sale, Torent went completely dark. The outage lasted 41 hours.
- Ridgeway Appliances, Orenda's largest brand partner, missed a contractually committed dispatch window and a full peak-sale weekend. It sent a formal breach notice claiming Rs. 12 crores in damages.
- **The MSA's remedy:** SLA credits were Orenda's 'sole and exclusive remedy' for availability failures. The maximum credit was 25% of the monthly fee, or Rs. 35 lakhs, whichever was lower.

Orenda estimated that a full platform failure in peak season could expose it to losses of Rs. 80 crores to Rs. 120 crores.

Where the contract failed

- The sole-and-exclusive-remedy clause limited Orenda's options to claim compensation from Beacon.
- Fee-based liability caps were built for a world of predictable software risk. They often fail when applied to AI-dependent systems.
- An operational failure may trigger downstream claims, regulatory disclosure duties and reputational harm that far exceed the subscription fee.
- Here, a 41-hour outage led to a Rs. 12 crore downstream claim. The maximum SLA credit was Rs. 35 lakhs. The gap between the contractual remedy and the real harm was enormous.
- The SLA was also uniform across the year. It priced an ordinary Tuesday and a festive weekend identically, though the loss profile differs by an order of magnitude.

What the clauses should say

- The MSA must include an AI performance warranty. The system must maintain the agreed baseline accuracy and operational parameters.
- The system must not materially degrade after a model update or infrastructure change. Degradation should be a defined, measurable breach, not a matter of impression.
- Liability caps should be layered. For example: 100% of annual fees for general claims, 150% for data breaches and 200% for IP infringement.
- SLA credits must not be the sole and exclusive remedy for a mission-critical system.
- The cap and sole-remedy language should not apply to wilful abandonment, gross negligence, or outages that exceed a defined duration.
- Availability commitments and remedies should be tiered by business period, with heightened obligations for identified peak windows.

FIVE PRINCIPLES

Taken together, these scenarios make one point clear: an AI contract cannot treat the product as something that will remain fixed after signing. The system may be updated, retrained or used in new ways, and its behaviour may change over time. The contract needs to account for that.

1

Draft for actual deployment. If the contract assumes that a person will review the AI's decisions, that safeguard must be workable at the scale and level of autonomy at which the system will operate.

2

Access to information. Regulators may also ask how the system reaches decisions, how it has been tested and what oversight is in place. The contract should provide access to the information needed to answer those questions.

3

Deal with each layer of IP separately. Input data, model weights, outputs and usage data are different assets. The contract should clearly state who owns each of them and how they may be used.

4

Plan for the vendor to be acquired by a competitor. Customer data, usage patterns and the insights drawn from them may be commercially valuable. The contract should address what happens to these assets if ownership of the vendor changes.

5

Set liability by reference to the possible harm, not only the contract fee. If a failure could cause losses far greater than the fees paid, a fee-based liability cap may leave the customer carrying most of the risk.

This guide provides general information only and does not constitute legal advice. The drafting positions should be tailored to the particular transaction and checked against the law in force at the time.