THE BUILDER'S STACK · THE FRONTIER · AGENTS AS USERS

The Agent Access Policy

One policy that decides how agents may use the product: what they may touch, how they prove who sent them, where they stop and ask, and what the traffic costs. We sign it before launch and reopen it when the protocols move.

PRODUCT	OWNER	DATE

1 What agents may do

Each action the product exposes gets a ruling: open to agents or not, and how much rope it gets. Autonomy is autonomous, confirm-first, or human-only.

ACTION	OPEN TO AGENTS	AUTONOMY
	☐	
	☐	
	☐	
	☐	
	☐	
	☐	

2 Authentication tiers

Access scales with proof: the less an agent can show about who sent it, the less it may do.

TIER	ALLOWED TO
Anonymous read	
Declared bot	
Signed agent	
Delegated customer agent	

3 Delegation proof

A delegated agent carries proof before it acts on a customer's behalf: the scopes it holds, the money it may spend, and the moment its mandate ends.

REQUIRED SCOPES	SPEND CEILING	EXPIRY

THE CHECKPOINT RULE (WHICH CONSEQUENTIAL ACTIONS STOP AND WAIT FOR THE HUMAN)

4 Limits and refusals

Every tier gets a rate limit, and every refusal comes back structured, so a legitimate agent knows what to fix.

TIER	RATE LIMIT
Anonymous read	
Declared bot	
Signed agent	
Delegated customer agent	

THE STRUCTURED REFUSAL FORMAT (MACHINE-READABLE CODE, REASON, AND THE PATH TO MORE ACCESS)

5 Injection and abuse defenses

An agent surface is an attack surface. Check a defense only when it is live, and give it one owner.

	DEFENSE	OWNER
☐	**Server-side authorization** every call is authorized at the server, never by the prompt	
☐	**Content sanitization** everything the model reads is treated as untrusted input	
☐	**Probe monitoring** we watch for agents testing the fences, and log what they try	

6 The pricing stance

We charge agent traffic on purpose, not by accident.

HOW AGENT TRAFFIC IS CHARGED	WHAT STAYS FREE

7 Metrics to watch

The numbers that say the policy is working, each with the target it must reach.

METRIC	TARGET
Agent share of traffic	
Agent task completion	
Legitimate-agent block rate	

REVIEW

This policy reopens on a cadence and on a trigger: a protocol shift, a new agent tier, or a metric off its target.

REVIEW CADENCE

Protocol facts verified against the Interop Ledger dated

SIGNATURE	DATE