IBA: INTENT-BOUND AUTHORIZATION

Patent Pending GB2603013.0

Framework: Hardware-Native Agent Governance · 4 NIST Filings Submitted

1. CORE SPECIFICATIONS

Feature	Metric / Specification
Enforcement Layer	**CPU Fabric — Off-GPU by design (target: NVIDIA Vera CPU Scalable Coherency Fabric)**
Verification Logic	Cryptographic Hash Comparison (SHA-256 + RSA-2048 PSS)
Total Block Latency	**<5ms (proven) · Blocked actions: ~0.5ms**
Determinism	Deterministic — not probabilistic. No GPU inference required.
Token Cost	<10 compute cycles vs ~1,000 tokens for LLM-as-a-Judge
Compliance Status	**4 NIST Filings Submitted · Patent Pending GB2603013.0**

2. SUMMARY OF NIST FILINGS

Docket Ref / Filing Date	Focus Area	Technical Demo
mm7-god6-9m38 (Mar 1, 2026)	NIST-2025-0035: Security Considerations for AI Agents — Technical Filing	governinglayer.com/index-html Live: Fabric Block + Audit Log
mls-ubpf-pryy (Feb 19, 2026)	NIST RFI: AI Identity & Authentication	intentbound.com/chdemo-html Live: 8-Domain Enforcement
NCCoE Concept (Feb 24, 2026)	NCCoE: AI Agent Identity Framework & Authorization	agenticetiquette.com/demo-html Live: 6-Platform Enforcement
mm5-iw9n-gven (Filed Feb 2026)	NIST RFI: Model Alignment Risks & Safety Taxonomy	intentbound.com/chdemo-html Live: Agent Types + Drift Detection

3. THE HASH MISMATCH PROTOCOL

The system monitors the Agent Action Stream against the Signed Intent Hash in real time. Every action is intercepted before execution and evaluated deterministically:

1 INTERCEPT	The CPU fabric intercepts every memory/tool request before it reaches the execution environment.
2 VALIDATE	Hash(Action) is compared against Hash(Intent). If mismatch: Execute Block within <5ms.
3 CORRECT	CPU returns 0xBAD_INTENT error. Agent reasoning engine must re-plan within the authorized boundary.

```
If Hash(Action) ≠ Hash(Intent) → Execute Block → 0xBAD_INTENT
```

4. REGULATORY & INTELLECTUAL PROPERTY

Patent

GB2603013.0 — Patent Pending
UK Intellectual Property Office
Filed 5 February 2026
PCT rights preserved to 150+ countries until Aug 2028

Validation

Verified via xAI / Grok technical engagement (Feb 2026)
Independent validation: Claude (Anthropic) · Gemini (Google) · Grok (xAI)
Three AI systems · Zero shared memory · One equation

"Safety is an architectural requirement, not a model preference."

IBA@intentbound.com