HIBILTER · HUMAEL — WHITE PAPER

Use cases

Sovereign AI for defence: agentic decision-superiority for air defence and national security

Why defence AI is won on sovereignty, trust and speed — and how agentic, human-commanded systems deliver decision-superiority across ISR, IFF, counter-UAS and command.

Ankesh Tiwari · 28 Jun 2026 · 16 min read · hibilter.com

EXECUTIVE SUMMARY

In defence, the decisive advantage is no longer the platform — it is the speed and trustworthiness of the decision. But defence AI cannot be bought off the commercial shelf: it must be sovereign, because the data cannot leave the perimeter; trustworthy, because a commander must be able to explain and command it; and resilient, because it has to work air-gapped and under electronic contest. This paper maps where agentic AI changes the modern fight — ISR sensor fusion, IFF friend-or-foe, counter-UAS against drone swarms, multi-domain threat assessment, and mission command — and sets out the design principles that make autonomy acceptable in defence. It introduces Hibilter's Vajra line, Sentinel and Aarambh, as agentic systems built sovereign and human-in-command from the ground up.

Decision-superiority is the modern advantage

Generations of advantage came from better platforms — faster aircraft, longer-range missiles, stealthier ships. That edge is narrowing. What now decides engagements is the speed and quality of the decision: who assembles the picture first, identifies friend from foe first, and acts inside the adversary's loop. Artificial intelligence matters in defence precisely because it can compress that loop — but only if it does so without surrendering control.

The defence problem AI must actually solve

Defence AI is a different problem from enterprise AI, and the differences are disqualifying, not cosmetic. A system can be the most accurate in the world and still be unusable if it fails any one of three tests.

- **Sovereign** — the data, the model and the inference stay inside the national or operational perimeter. A capability that depends on someone else's cloud is a capability an adversary can deny or compromise.
- **Trustworthy** — the system must explain its reasoning in doctrinal terms and operate under human command, because no commander will stake lives on a black box.
- **Resilient** — it must run air-gapped, at the edge, and keep working while the very links and sensors it depends on are jammed, spoofed or degraded.

Where agentic AI changes the fight

Agentic AI — coordinated teams of specialised agents that observe, reason and act under guardrails — maps directly onto the hardest, most time-critical tasks in air defence and national security.

- **ISR sensor fusion** — combining radar, transponder, satellite and SIGINT into one trusted air picture.

- **IFF friend-or-foe** — resolving identity in real time, the call on which every engagement depends.
- **Counter-UAS** — detecting, classifying and defeating drone swarms at machine speed.
- **Multi-domain threat assessment** — ranking risk across air, surface, cyber and electronic warfare.
- **Mission command** — turning intelligence into assignments, timelines and decision packages.
- **After-action and doctrine** — reconstructing every engagement for review and learning.

The drone swarm forced the change

Nothing exposed the limits of human-speed, platform-centric defence like the autonomous drone swarm: many cheap, small, coordinated objects that saturate sensors, invert the cost-exchange, and overwhelm a single operator's ability to triage and decide. Counter-swarm is the clearest case where machine-speed, human-commanded autonomy is not a convenience but a necessity.

Vajra: the strike, built sovereign

Hibilter's **Vajra** line applies agentic AI to exactly these problems, and is sovereign and human-in-command by design rather than by configuration.

- **Vajra Sentinel** — autonomous counter-UAS: multi-sensor fusion, an agentic kill-chain, and layered, non-kinetic-first, ROE-checked defeat of hostile drone swarms.
- **Vajra Aarambh** — sovereign defence intelligence: ISR fusion, real-time IFF, multi-domain threat assessment and mission command-and-control in one trusted air picture.

> *In defence, autonomy is only acceptable when it is sovereign, explainable, and under command.*

Sovereign, air-gapped, human-in-command — the non-negotiables

These are not premium features; they are the price of entry. Vajra runs on sovereign models inside the perimeter, air-gapped, with a complete audit trail. Its reasoning is doctrine-grounded, so recommendations arrive in the language the commander already uses. And it informs rather than overrides: the human owns every consequential decision, with the system surfacing the decision and the evidence rather than a stream of confirmations.

Deployment posture

Defence AI has to meet the force where it operates. Vajra is edge-capable and deploys on-premise or air-gapped, integrates with existing sensors and command systems as a software overlay rather than a

rip-and-replace, and is built for classified environments where nothing — no telemetry, no source, no model weights — leaves the building.

Time-to-decision

the loop compressed — see first, decide first

IFF accuracy

friend / foe / unknown resolved under contest

Leakers

threats reaching the asset — the only score that counts

Sovereign & audited

air-gapped operation, complete reconstructable trail

Conclusion

Defence AI is won on sovereignty, trust and speed — in that order. Agentic, human-commanded systems deliver decision-superiority without surrendering control: they fuse what the force can sense, resolve friend from foe, defeat the swarm at machine speed, and keep a commander in command throughout. See first, decide first — sovereign, and under command.