

ARCHITECTURE

The Protected Region as a Unit of Deployment

Inari Watch is bought, deployed, and operated one protected region at a time. This primer covers what that unit is, why it's the right one, and what it actually takes to stand the system up in your region — coverage, depots, timeline, and your people.

What a protected region is

A protected region is a bounded geography with a single risk owner and an existing operational structure — a county, a utility service territory, or a federal installation. Three properties make it the right unit to buy and deploy against. It has a **defined boundary** you can draw on a map. It has an **identifiable risk owner** — an organization whose mandate is to protect it and which can actually procure the capability. And it has an **existing operational structure** — dispatch, an incident command process, control-room staff — that the system is meant to strengthen rather than replace.

The alternative to a regional unit is a sensor here and a drone there: orphaned tools that no one owns end to end and that never add up to a system. The region is the smallest unit at which detection, coordination, and response can function as one capability, under one authority, accountable to one owner. That is why we sell and deploy nothing smaller.

What coverage looks like

The sensing mesh is sized to the region's risk, not spread evenly across it. The highest-risk corridors — transmission lines, the wildland-urban interface, canyon outlets — get the densest coverage, on the order of one thermal camera per 5–10 km of corridor, sited for overlapping viewsheds. Lower-risk interior is covered more sparsely and backed by always-on federated satellite feeds (NASA FIRMS, NOAA GOES-R) and on-demand UAV patrol. Each sensor cluster runs on local edge compute and publishes normalized events rather than streaming raw video.

In practice a representative region is scoped at something like 10, 25, or 50 sensor sites depending on its size and terrain — a number we set with you against a coverage analysis of your specific risk map, not pulled from a catalog. Depots, which house the recon aircraft along with launch, recovery, and a local edge cluster, are positioned near the highest-risk zones so that a recon UAV can reach a priority point and return within a target response window. The number and spacing of depots follow from the region's size and the aircraft's endurance. In current deployments the aircraft detect and verify; the suppression-capable aircraft are a later, separately authorized addition, not part of an initial buy.

How it connects

Standing the system up is mostly integration, not greenfield installation. Inari Watch reads and writes the tools a region already runs — dispatch (CAD), the incident command common operating picture (ICS), airspace traffic (ADS-B), and active flight restrictions (NOTAMs and TFRs, treated as hard constraints) — and operates as a sanctioned resource inside them rather than as a parallel system. Because decisions are made on local edge compute, the system keeps working through the degraded connectivity a wildfire produces; the cloud supports it but is never a single point of failure.

What standing it up looks like

A deployment moves through a predictable sequence:

1. **Site survey and integration plan** — map the region's risk and viewsheds, inventory the existing feeds and systems, and define the authority chain.
2. **Connect and stand up** — integrate the existing systems, deploy the initial sensor clusters, and bring up the regional console and the coordination fabric.
3. **Shadow mode** — the system runs alongside your operation, correlating and recommending but authorizing nothing, while you measure it against the bar below.
4. **Supervised live** — once it is earning its recommendations, operators begin authorizing actions through it, still detection and verification only.
5. **Scale out** — extend coverage across the full region, and later, under separate authorization, add the suppression capability.

Shadow mode is the part that matters most to a buyer. You should never have to take performance on faith: the system proves itself by watching real operations, on your own data, before it touches a single decision. A first deployment typically reaches an initial operating capability within a fire season, with full-region scale-out following on the validated foundation.

What it does to your staffing

The system augments the people you have. It does not add a parallel staff and it does not replace experienced operators — pilot operators are drawn from your existing control room. Its job is to take the cognitive grunt work off them: the correlation and triage that otherwise buries a control room under a wall of raw alerts, so that human judgment is spent where it actually matters.

The bar we hold ourselves to in a pilot is specific and measured on the console's own interaction logs over a fire-season window: a reduction of at least 60% in operator-facing alerts per hour relative to the unfused baseline, a median time-to-triage under 30 seconds, and operators accepting the system's recommendations more than 70% of the time. In return, the system asks three things of your organization — operator training, a defined authorization chain that says who may approve what, and an agreed escalation path for when an operator is unavailable.

What buying it looks like

Cost has three parts: capital (sensor clusters, edge compute, recon aircraft, and depots), operating cost (maintenance, communications, software, and personnel), and a marginal cost per mission. Pricing is scoped to the region — its size, its risk map, and the coverage you need — rather than sold as a fixed package. The value is measured as avoided expected loss: structures and infrastructure preserved, evacuation and smoke exposure avoided, and the long-tail downstream costs that dwarf direct suppression spending. The first step is a discovery call to scope your region — its boundary, its risk map, the systems you already run, and the authority boundaries that matter to you.

About Inari Watch

Inari Watch is a veteran-founded company building an integrated system of systems for regional wildfire defense. We unify the sensors, aerial assets, and data feeds a region already depends on into one coordinated capability — connected by AI at the edge and bounded agentic coordination, with humans in command of every consequential decision.