Fidwyn Augur

The on-prem AI line across PAM and EPM. A model bundled into the appliance and served from your own GPU scores every identity, runs SOAR playbooks, maps identity blast radius, and gives administrators a copilot that previews before it applies. Zero inference egress.

CAPABILITIES

- **Sentinel Dashboard**
  Per-identity behavioral risk 0–100 with SHAP explainability and MITRE ATT&CK-mapped signals from a local, air-gapped CTI bundle.
- **Playbook Builder (SOAR)**
  Boolean trigger trees wired to actions — revoke tokens, force MFA, quarantine, forward to SIEM — as versioned YAML with inline compliance mappings.
- **Blast Radius**
  Identity attack-path analysis fusing AD paths and pass-the-hash into a 0–100 score, with the top path to a tier-0 target.
- **Agent Guard**
  Governance for non-human and AI identities: over-privileged, orphaned and critical-risk agents ranked first.
- **Cloud Entitlements (CIEM)**
  Read-only entitlement analysis across AWS, Azure and GCP: unused admin, long-lived keys, escalation paths.
- **AI Copilot**
  Natural-language operations that preview, then apply through the access resolver — audited, secrets never in prose.

AT A GLANCE

0 B	On-prem	100%
INFERENCE EGRESS	GPU-SERVED MODEL	OFFLINE CAPABLE

The AI runs where the data lives.
Endpoint validation rejects any address that is not loopback or private, at configuration and call time. No privileged-access telemetry ever leaves the perimeter — viable in a bank or a defense enclave.