Bilko Observability & Self-Healing — Program Overview (MC #103328)

Bilko Observability & Self-Healing — Program Overview (MC #103328)

This is the single entry point for the entire Bilko observability and self-healing program. It links every sub-page, states current status, and gives a quick orientation for any CEO, agent, or engineer arriving for the first time.

Environment Topology

Tier	Cloud Run services	Public URL	Purpose
PROD (demo)	bilko-api-demo, bilko-web-demo	app.bilko.cloud / bilko-demo-api.alai.no	Live trial surface — real prospects register here
STAGE	bilko-api-stage, bilko-web-stage	stage.bilko.cloud (internal)	CI validation; masks some RLS bugs (documented lesson)
DORMANT	bilko-web	—	Old web service; superseded by bilko-web-demo

GCP project: tribal-sign-487920-k0. All observability targets bilko-api-demo and bilko-web-demo as the canonical production equivalents.

Program Arc (2026-06-09 ? 2026-06-11)

1. **GCP-native observability baseline** (MC #103329/P1-A) — Cloud Monitoring dashboard (070613fa...), latency/traffic/saturation/5xx alerts wired end-to-end.
2. **Validation** (MC #103331) — Proveo independent PASS confirming all alert signals fire correctly.
3. **Docs** (MC #103332) — Initial BookStack page published.
4. **Sentinel Tier-0 built** (MC #103337) — Read-only agent: detect → diagnose via Ollama → propose → notify. Proveo PASS. LaunchAgent running at PID 11465 (com.alai.bilko-sentinel).
5. **CD-green + GCP error-tracking** (MC #103364) — CD pipeline repaired; error log metric + alert added. Threshold tuned to >3 errors/5 min after the first real incident (see below).
6. **CIAM E2E blocking gate** (MC #103365) — Playwright CIAM auth-lifecycle spec added as a mandatory blocking step in the deploy pipeline. Proveo two-sided PASS (green + fails-on-broken).
7. **CRITICAL security review** (MC #103369, Securion) — /auth/test/session endpoint on bilko-demo-api found to accept arbitrary email → impersonation risk (F7 = CRITICAL). Full findings + F7 remediation path issued.
8. **F7 security fix deployed** (MC #103371) — Email whitelist, constant-time compare, 5/min rate-limit, Sentry audit. Verified live: attacker email → 403, seeded email → 200. CIAM E2E gate now 3/3 (F7-WHITELIST-GATE added). PR #330 merged, PR #332 (gate) merged.
9. **First real incident handled (503, 2026-06-10)** — Transient 503s during Cloud Run revision cutover/scale-from-zero (not a code bug). Alert pipeline worked. Threshold tuned >0 → >3. See Security & Decisions page for full post-mortem.
10. **Sentinel dynamic-discovery fix** (MC #103420) — Sentinel was missing the error-count policy (hardcoded list). Fixed to live gcloud discovery + 5-min cache + embedded fallback. 9 policies now evaluated each cycle, 13 conditions total.
11. **Tier-1 shadow** (MC #103435) — Shadow-only armed auto-remediation module built. Structurally inert (dual barrier confirmed by Securion). Will NOT be promoted to ack/auto without clearing the promotion bar.
12. **Securion Tier-1 review** (MC #103436) — Parisa Tabriz lens review. Shadow inert confirmed. 8 findings; 3 must be resolved before ack/auto (F5 HMAC ledger, F7 SA IAM scope, F4 ack allowlist). See Security & Decisions page.
13. **Tier-1 arming prerequisites** (MC #103439) — Hard blockers catalogued. Tier-0 calibration clock starts now. Review at 30 days.
14. **Dashboard maturity roadmap** (MC #103393) — Backlog. SLOs, error-rate tiles, distributed tracing, business metrics. CEO decision: document now, build before meaningful paying-customer volume.

Master Status Table

MC	Title	Status	Evidence / Notes
#103329 (P1-A)	GCP-native observability	DONE — Proveo PASS	Dashboard 070613fa…; alerts wired
#103331	Validation	DONE — Proveo PASS	All alert signals verified
#103332	Docs (initial)	DONE	Page 3101 published
#103337	Tier-0 Sentinel build	DONE — Proveo PASS	LaunchAgent PID 11465 live
#103364	CD-fix + error-tracking	DONE	Threshold >3/5min after 503 incident
#103365	CIAM E2E blocking gate	DONE — Proveo PASS	2-sided proof; gate blocks on broken
#103369	Securion test-endpoint review	DONE — verdict MOVE_OFF_PROD (pre-fix); overridden post-F7 fix per Decision 1	/tmp/evidence-103369/verification.json
#103371	F7 security fix	DONE — Proveo PASS	PR #330+#332 merged; 3/3 gate; live proof attacker→403
#103393	Dashboard maturity roadmap	BACKLOG (not-now)	SLOs, tracing, business metrics — before real paying customers
#103420	Sentinel dynamic-discovery fix	DONE — AgentForge PASS	9 policies, 5-min cache, embedded fallback
#103435	Tier-1 shadow build	DONE — shadow inert	Dual barrier; Securion review attached
#103436	Securion Tier-1 review	DONE — HARDENING_REQUIRED before ack/auto	8 findings; F5/F7/F4 block arming
#103439	Tier-1 arming prerequisites	IN PROGRESS — calibration clock running	30-day / 20-proposal bar; see Decisions page

Key Live URLs

- **GCP Monitoring Dashboard:** https://console.cloud.google.com/monitoring/dashboards?project=tribal-sign-487920-k0 (filter: slug 070613fa…)
- **Demo API health:** https://bilko-demo-api.alai.no/api/v1/health
- **Demo app:** https://app.bilko.cloud

Documentation Map

Page	What it covers
Page 3101 — Bilko Observability (GCP-native)	Cloud Monitoring setup, dashboard tiles, alert policies, runbook for each alert type
Page 3102 — Bilko Sentinel Tier-0	Tier-0 agent design, how detect/diagnose/propose/notify works, LaunchAgent config, audit log path
Page 3106 — Bilko Sentinel Tier-1 (shadow-first)	Tier-1 architecture, shadow barriers, action set, circuit breakers, pre-fire gates
This page — Program Overview (#103328)	Single entry point: arc, status table, links to all sub-pages
Page — Security & Engineering Decisions	F7 security fix, CIAM gate design, first incident post-mortem, Tier-1 arming prerequisites, architectural decisions

Revision #1
Created 2026-06-11 19:49:18 UTC by John
Updated 2026-06-11 19:49:18 UTC by John