

Marko Tomic

Head of Digital & Product · four-person team · six market configurations · four languages

markotomic.org · markoxtomic@gmail.com · linkedin.com/in/markoxtomic

I design systems that turn repetitive digital operations into measurable leverage, and I own the commercial outcome. What follows is the evidence in compressed form: the decision behind each system, and what changed as a result. Estimates are labelled. Unfinished work says so. The full case studies live at markotomic.org.

+52.8%

net revenue - five months on the new platform vs. the legacy shop's full 2025

220+ h

manual production removed by a metadata agent that earned its autonomy

9.5 h / wk

analytical capacity returned to an external specialist, read-only by design

2.1%

full-year forecast variance against actual net revenue

01 AI Operations

Three production systems, deliberately different in how much autonomy each is granted. The differences are the point.

- **Google Performance Intelligence.** Three specialist agents (Google Ads, GA4, Search Console) plus an orchestrator, integrated via MCP, OAuth and API, read-only throughout, running autonomously on cron. The orchestrator condenses and ranks what changed, then delivers to an external specialist who retains execution authority. The AI does not touch the ad accounts - that boundary is permanent by design. **~9.5 analytical hours returned per week** (estimate).
- **Multilingual Metadata Automation.** 463 products and 29 collections across four languages - 1,968 localised entities, 3,936 fields. Built around *product families* and reusable rules rather than per-product prompts: one abstraction applied 1,968 times. Structured accept/reject/correct review; after three to four calibration runs per language, human approval was removed and the workflow ran unsupervised. **220+ hours** of concentrated manual production avoided.
- **Semantic Commerce Architecture** (in progress). Parent and child topic hubs connecting products, materials, applications and editorial content across four languages, structured so retrieval systems can follow relationships rather than crawl isolated pages. Four prototypes per page; agents execute, the human judges. Estimated **≥4 working days saved per completed four-language hub**. No proven visibility result yet - the evidence does not exist.
- **Model evaluation.** Prototypes are not chosen on instinct. A fixed rubric - semantic quality, visual identity, hub consistency - applied across a panel of **ten models over six runs**. Kimi K3 and Opus 5 lead on semantics; Luna 5.6 on visual identity, usually first attempt; Luna and Fable 5 on consistency. Cost-normalised aggregate: Luna comfortably first. General benchmarks measure capability on tasks that are not this task.
- **Local inference.** LM Studio environment for commercially sensitive workloads. 8B in production; 12B evaluated and rejected - even at Q4 too slow to be usable, and a secure system people avoid secures nothing. Architecture follows the risk model.

02 International Commerce Transformation

Rebuilt the entire commerce stack on Shopify against where the business intended to sell, not what existed in the legacy database. 467 SKUs at ~5 variants each; German, English, French and Spanish live at launch; five markets including a Switzerland configuration with CHF, VAT treatment and market-specific catalogue rules. The legacy storefront stayed live throughout, cutover only at completion. Four calendar months, ~3.5 active.

- **+52.8% net revenue** - five months on the new platform versus the legacy shop's entire 2025. ~+73% traffic. ~+36% page speed. **Zero revenue downtime**.
- The UK was launched, tested and **discontinued** after weak adoption. Four markets retained - the architecture made market experimentation cheap enough to be worth running and closing.
- **Prior chapter (2022-23).** Four Magento storefronts (~1,570 products) consolidated into a single backend with Store Views; staged migrations, no revenue downtime. Business-side product owner, coordinating the external development agency, UAT and delegation.

03 Commercial Intelligence

The business had revenue data and no usable unit economics, so pricing rested on gross revenue. Source data was hostile: supplier invoices as irregular PDFs, tax treatment varying by supplier and market.

- **Margin & pricing intelligence.** Locally deployed application over **1,153 SKUs**. Parses PDF invoices (~1,000 in 74 seconds), derives net cost, VAT, shipping and price, applies mathematical optimisation, returns current versus recommended pricing and margin scenarios. Has informed real pricing decisions.
- **Revenue & cash-flow intelligence.** Local predictive planning application, built 2024 and validated against actual subsequent net revenue: ~6% variance quarterly, **~2.1% across a full year**. Natural-language query layer as an MVP.

04 Performance & Growth

Paid acquisition was underperforming and the measurement beneath it could not explain why. Budget was fixed and the internal capability did not exist, so the right answer was a specialist - and then the measurement environment and operating model that would let one perform. I owned

partner selection, onboarding, tracking architecture, measurement, commercial oversight and prioritisation; the specialist contributed the majority of campaign execution.

- **+371%** year-on-year ROAS improvement. **-29%** CPC. **1,127%** annual ROAS on a German-market Performance Max campaign covering bestseller categories, measured across a full year.
- Repair the operating model, then augment it. The performance agent above only produces value because this measurement layer was rebuilt first - an agent placed on top of a broken process automates the breakage.

05 Function, product and team

- **Building the function.** Rebuilt a solo operation into four specialists plus an apprentice: recruiting, hiring, onboarding, direct management, role definition, delegation and budget responsibility across media, technology and partners. ~47 SOPs documented and responsibilities cross-trained, reducing key-person dependency.
- **Commercial product.** Strategic supplier sourcing, development of new collections and products, upsell architecture across the range, and go-to-market planning that gets a new product earning sooner. In a catalogue business this is product management - deciding what should exist, at what price, against which demand.
- **Supporting systems.** EU market expansion research across France and Spain (~21-day sprint, leads and orders generated). A reusable 21-component web architecture across three properties, built to surface a restoration service the main catalogue could not - that category grew substantially afterwards.

06 Stated gaps

Named rather than found: no Product Manager title, five years of work containing product management; SME rather than enterprise, so no procurement cycle or formal security review; nothing built horizontally across many customers; no formal CS background - self-taught, building in AI-assisted environments and specifying precisely rather than competing as a standalone engineer; no marketplace or DTC subscription depth.

Compact extract of markotomic.org, where each system has its own case study. Prepared September 2026. Estimates are labelled; systems are in production unless marked in progress.