Sample audit

Shopify Audit: demonstration facsimile

A fictional document showing how an IFG Audit is delivered, from findings and evidence to a prioritised roadmap.

Francesco Guiducci

Edition October 2026 · English

ifgdigital.com

For personal or internal business use; unauthorized distribution is prohibited.

Contents

Disclaimer and introduction 3

Executive summary 3

Back-end 3

Front-end 6

Social media 15

Conclusions and priority roadmap 16

Next steps 18

Disclaimer and introduction

This sample presents a service from IFG Web Design, the IFG Digital division for Shopify stores and professional websites. IFG eCommerce® is the registered brand for the Shopify line. IFG Digital also brings together IFG Development for apps and software, and IFG Gaming for video games.

> **FACSIMILE – FICTITIOUS DATA.** This document is a demonstration sample that shows how an IFG Audit is delivered. NOVA Atelier does not exist, and nova-atelier.example deliberately uses the reserved .example domain. Every data point, order, configuration, test, page and finding in the report is fictional.

I structured the fictional case as I would a real engagement: read-only analysis, separation between observed facts and plausible explanations, verification of what the customer actually encounters, and an independent pass over the main findings before priorities are assigned. A feature that exists in the back-end is not treated as evidence of its effect on the customer.

For some technical parameters, including performance measurements, source-code inspection, accessibility checks and indexation verification, I also use automated analysis and artificial intelligence tools. In a real Audit I recheck the data directly on the store. The assessment, priorities and recommendations remain mine.

This sample covers Shopify settings, checkout, payments, shipping, markets and domains, sales channels, content, catalog and product data, internal search, lifecycle, applications, measurement quality, analytics, observable privacy and security, branding, UI/UX and conversion, accessibility, performance, theme health, SEO, GEO and social media.

To keep the document realistic, I built a case that contains healthy areas, objective defects, risks that need validation and opportunities that cannot be presented as certainties. Not every area produces a problem, and a recommendation appears only where there is a concrete reason for it.

The fictional business data covers twelve months, from 1 September 2025 to 31 August 2026. This file is not advice on a real store, an accessibility certification, legal advice or a guarantee of SEO, GEO or CRO results. Its only purpose is to show the depth and format of the deliverable.

Executive summary

NOVA Atelier is a fictional premium fashion eCommerce store with generally healthy Shopify foundations. The theme is Online Store 2.0, the catalog is broad but manageable and the brand direction is recognisable. There is no technical evidence to support a theme replacement or a complete rebuild.

The most urgent issue is measurement quality. A test order generates two purchase events with the same order_id, currency and value, because a custom tag and the native pixel both send the conversion. Until the duplicate firing is removed, ROAS, attributed conversions and channel comparisons are not reliable enough to guide decisions.

I included two concrete customer-journey defects. In the mobile cart drawer, the chat widget covers part of the checkout CTA at 375 and 390 px. The storefront also promises free shipping from EUR 60, while checkout applies the free rate only from EUR 80. These are different problems, but customers meet both, so both come before CRO experiments.

The second priority group concerns catalog quality and maintenance. 39% of applicable variants have no GTIN, the color filter is incomplete because the data is not normalised, and part of the product-image set lacks useful alternative text. The mobile hero also uses an oversized asset that is consistently the LCP element in the fictional runs.

My decision in this sample is clear: make measurement reliable and remove objective defects first, then align catalog, channels and performance, and only then move to experiments and growth. The final roadmap follows those dependencies, not the order of ease or commercial appeal.

Back-end

Shopify settings

General settings are coherent with the fictional business model: operating address in Milan, euro as the main currency, European time zone, and contact details aligned between the back-end and the public pages. The primary domain uses HTTPS correctly, and no duplicate or legacy hosts are indexed in the scenario.

The store is on the Shopify plan. At the simulated volumes I see no reason to change plan within this Audit. Review the choice against actual transaction costs, payment fees and the features in use, not on the assumption that a higher plan is automatically better.

Four markets are configured: Italy and France are complete, Germany is published but only partially localised, and the United Kingdom is still in draft. The structure is sufficient for expansion, but a market is not complete until language, policies, product availability and the shipping promise are coherent in the real journey.

Payments, checkout, accounts, shipping, taxes, markets and domains

Checkout uses email as the primary contact method and keeps phone optional. Marketing options are not preselected. Cards, Shop Pay, Apple Pay, Google Pay and PayPal are available in the fictional configuration. I did not add redundant methods just to create more material for the report.

The payment flow is linear and has no unnecessary custom fields. The order summary shows products, discounts, shipping and taxes before confirmation. The assessment is not about checkout styling. It is about total-cost clarity, error handling and continuity of the commercial promise.

Customer accounts are optional and use the branded account.nova-atelier.example domain. The post-purchase experience keeps brand continuity, and customers can review orders and returns without being forced to create an account before payment.

Standard shipping in Italy is EUR 6.90 and becomes free from EUR 80. However, the announcement bar, the shipping page and the cart progress bar all communicate EUR 60. The back-end and the customer's experience describe two different rules.

I classify this inconsistency as P1. A EUR 70 cart still pays the EUR 6.90 shipping charge at checkout after the site has already promised free delivery. The fix is to define one source of truth and make the announcement, the cart message and the actual rate depend on it.

France uses the same catalog in euros, with free shipping from EUR 100 and a stated delivery window of 3–5 business days. Germany inherits part of the Italian content and some English pages. This is not a failure of the market engine. It is an incomplete market publication, and it should be fixed before any additional local acquisition.

Taxes are managed natively, and public prices are VAT-inclusive across the fictional European markets. I have not invented tax anomalies to fill a section. The Audit creates a finding only when a configuration or a customer encounter supports it.

Shipping, returns and privacy policies are published and linked from the footer. The returns policy states 14 days, and the same value appears on the PDP and in the fictional post-purchase experience. This is an area that was checked and produced no finding.

The shipping page is different. It still states the EUR 60 free-shipping threshold. The same information appears on three public surfaces, so this is not an isolated forgotten string. The issue is the mismatch between editorial content and the operating checkout rule.

No abnormal redirects appear on the primary domain, and SSL is valid in the scenario. I kept the United Kingdom market in draft because pricing, delivery windows and content are incomplete. Keeping it unpublished is better than exposing a partial experience.

Core order notifications are active. Order confirmation, shipping and refund messages use the brand sender and do not insert out-of-context promotional content. Abandoned checkout is different: two systems own the same journey, and the lifecycle section covers the root cause.

Overall, checkout and settings are robust. The priorities do not require replatforming or invasive customisation. The free-shipping mismatch and the German localisation need fixing, then the full commercial promise should be rechecked across every published market.

Sales channels

Online Store, Shop, Google & YouTube, Meta and TikTok are installed or connected in the fictional store. I separate presence, account connection, product publication and diagnostics. Seeing a channel in Shopify Admin does not prove that its catalog is approved or commercially usable.

Online Store and Shop receive all 312 active products. Shop shows two products without the color metafield, but they remain purchasable. This is a data-quality gap, not a channel blocker, so it is absorbed into the catalog remediation.

Google & YouTube publishes all 312 products but reports 122 variants with a missing identifier. The catalog comparison shows GTIN coverage of 564 out of 924 variants, or 61%. For the affected branded references the identifier is expected, so the issue is material.

I classify the Google gap as P1 because it reduces feed completeness and can limit matching or approvals for the affected products. The fix is not to generate codes. GTINs should be recovered from packaging or supplier sources, validated and filled in only where applicable.

Meta receives the full catalog and shows no material warning in the fictional case. I do not credit the channel with revenue while purchase is duplicated. A healthy feed and healthy measurement are separate things and should stay separate.

TikTok receives 278 of the 312 products. Thirty-four are excluded because of image requirements. This is P2: the channel is connected but not fully usable. I would correct only the affected assets and verify approval, not duplicate catalogs or workflows.

The channel conclusion is selective. Google needs product-data work, TikTok needs an asset cleanup, and Meta and Shop are usable. Nothing in the scenario supports installing new feed apps or replacing integrations that already work.

Content

Institutional pages are generally well maintained. About, Contact, Shipping, Returns, Privacy and Terms are present and reachable. The brand voice is coherent, and there are no placeholder pages or visibly incomplete texts in the primary customer journey.

The About page explains materials, production and positioning with original content. It is linked from the footer and the secondary navigation, so customers actually encounter it.

The Size Guide exists and contains useful measurements. On the PDP, however, it opens in a generic modal that does not keep the context of the selected variant. This is not an indexation problem. It is a small UX friction, addressed in the front-end section.

The Shipping page holds the main finding in this area. It still communicates free shipping from EUR 60 and a generic 2–4 day delivery, while checkout applies EUR 80 and Germany has a longer window. One static paragraph cannot describe every market accurately.

Collection descriptions exist on 22 of the 28 main collections. Six pages rely on the title and product grid only. I do not treat the absence of a paragraph as a defect in itself. Add content only where it improves orientation, intent or real differentiation.

Product descriptions are stronger: material, fit, care and origin are reasonably consistent. The larger issue is data structure, not copy volume. Some information repeated in prose could feed reusable metafields for filtering and theme components.

Fourteen collections share the generic SEO title "Shop | NOVA Atelier" despite having different H1s and intents. This is P2 because it reduces how specific the results look and how well each page matches its metadata. The fix is editorial and templating work, not a wholesale rewrite.

I do not recommend mass content production just to fill pages. Align shipping, localisation and metadata first. Content can then strengthen the categories that have a real search, navigation or merchandising function.

Products, collections and product data

The fictional catalog contains 312 active products, 26 drafts and 924 variants. Options are structured coherently, mostly around size and color. Products are not duplicated to represent variants that Shopify already handles natively.

SKUs are present on 892 variants, or 97%. The 32 gaps are concentrated in samples and draft products. This is not a blocker, but the gap should be closed if those products become sellable or feed external systems.

GTIN coverage is materially weaker: 564 of 924 variants. The 39% gap matches the Google feed errors on references where a commercial identifier is expected. That link between catalog and channel diagnostics makes the finding stronger than a simple count of empty fields.

Shopify categories are assigned to 278 of 312 products. The 34 missing values are concentrated in recent capsules. I recommend completing them because taxonomy can support feeds, standard attributes and future filters, not because every field has to be filled for its own sake.

Category metafields are only partly used. 132 products have structured attributes for material or fit, while color still depends on inconsistent variant values. This fragmentation explains why the color filter appears only in some collections.

Of 1,592 product images, 1,146 have useful alternative text and 446 do not. Several existing values simply repeat the filename. I classify this as P2 for accessibility and information quality, without promising an automatic SEO benefit.

Shipping weight is filled in for 306 products. Six anomalies show zero grams even though the items are physical. This would matter more if weight-based rates were introduced later. With fixed rates today, it is a maintenance item, not a blocker.

Primary product images have broadly consistent dimensions and ratios. The performance problem does not come from the whole catalog but from the editorial home hero. I therefore avoid the generic recommendation to compress every image in the store.

The catalog remediation is one coordinated program. Complete the applicable identifiers, taxonomy and attributes that feed channels or filters, then apply an editorial rule for alternative text and shipping weight. Fixing the source data reduces separate cleanup work in Search & Discovery, Google and collections.

Product discovery and internal search

Menu, search, predictive search, filters and sorting form a single findability system. The desktop menu contains 24 first- and second-level links plus two promotional cards. It is not broken, but its density is a HEURISTIC to measure before simplifying it.

A search for "black linen" returns eight relevant products. The query "occasion dress" returns zero results even though 19 compatible products exist in adjacent categories. No synonyms or rules connect the customer's language to the catalog vocabulary.

I classify the zero-results issue as P2. I would add synonyms and Search & Discovery rules only after reviewing actual query language, then measure zero-results and clicks to the PDP. The answer is not to create a collection for every phrase a customer types.

Size filtering is consistent and price works, but color appears in only 7 of the 28 relevant collections. The root cause is non-normalised variant data. Six collections also keep sold-out products in top manual positions. Here I propose an experiment, not an absolute rule.

Marketing, lifecycle and deliverability

The store uses an external email platform together with Shopify automations. Newsletter and post-purchase have a clear owner, but abandoned checkout is active in both systems. The same abandonment therefore enters two different workflows.

In the fictional test, Shopify sends one message after one hour and the email platform sends another after two hours, with different copy and incentive. This is P1 because customers receive closely spaced, inconsistent messages. The fix is to assign one owner to the journey and coordinate suppression, consent and reminders in one system.

The sending domain has coherent SPF, DKIM and DMARC in the scenario, and no obvious technical error is present. I do not turn that into a deliverability guarantee. Reputation, inbox placement and performance need real data over time.

Lead capture exists both in the native form and in an app popup. I have no evidence that the popup is useless, but it collects the same data without additional segmentation. I classify it as P3: measure lead volume and outcomes first, then decide whether the native capability can replace it without losing behaviour or data.

Applications and app health

Eighteen applications are installed, eleven of which affect the storefront or the customer journey. The total is context, not a cause. Installed does not mean configured, healthy or useful, so each app is assessed against the behaviour it actually produces.

Reviews and payments have observable utility. Wishlist and loyalty work technically but have no measured outcome in the fictional case. The lead popup overlaps a native capability, and the review app adds JavaScript work when its drawer is opened.

I do not recommend indiscriminate removal. Measure dependencies, value and cost first, and remove an app only when its replacement preserves data, behaviour and integrations. For performance, the review app remains a probable contributor, not the main root cause.

Measurement quality

The Measurement Quality Gate fails on attribution. view_item, add_to_cart and begin_checkout each generate one hit with coherent currency and value, while purchase generates two requests with the same order_id. Shopify order data remains usable, but advertising attribution does not.

The fictional code inspection ties the duplicate firing to a custom purchase tag loaded through GTM and the native Shopify pixel sending the same event. The symptom, the raw duplication and the two emission paths converge on the same root cause, so I classify the finding as P0.

The remediation is to assign one owner to purchase, remove the duplicate path and repeat a controlled order. The gate reopens only when one order produces one purchase in the systems that should receive it, and currency, value, order_id and consent signals remain coherent.

Analytics and business metrics (last 12 months)

After excluding unreliable attribution data, I use the fictional Shopify-native KPIs for the twelve-month period. The purpose is not to judge the business against generic benchmarks. It is to describe volume, funnel, device, customers and markets on a base that stays coherent in the scenario.

KPI	Value
Total sales	EUR 99,445.00
Gross revenue	EUR 93,820.00
Net revenue	EUR 89,240.00
Sales reversals	−EUR 4,580.00
Discounts	−EUR 5,625.00
Total / fulfilled orders	1,366 / 1,322
Average order value	EUR 72.80
Returning-customer rate	24%
Total sessions	84,320
Human sessions	79,040 (93.7%)
Automated traffic	5,280 (6.3%)
Added to cart	4,526 (5.37%)
Checkouts reached	2,474 (2.93%)
Completed purchases	1,366 (1.62%)
Conversion rate	1.62%
Returning customers	24%

The store records 84,320 sessions and 1,366 orders, for an observed conversion rate of 1.62%. Fictional total sales are EUR 99,445.00 and average order value is EUR 72.80. Returning customers represent 24%. That points to a retention opportunity but does not prove a problem by itself.

Mobile accounts for 78% of sessions. Mobile conversion is 1.39%, against 2.43% on desktop. The gap is wide enough to justify a mobile journey analysis, but it is not proof that design is the cause. The cart defect is much stronger evidence.

The native funnel contains 4,526 add-to-cart events, 2,474 checkouts reached and 1,366 completed purchases. The largest drop is between cart and checkout. This supports prioritising the review of the cart drawer and shipping consistency, without assigning causality to the aggregate funnel alone.

Italy represents 71% of fictional sales, France 18%, Germany 7% and other markets 4%. Germany already produces orders despite incomplete localisation, so the market fix is material but still does not outrank the measurement and checkout defects.

The strongest categories are dresses, knitwear and accessories. Sold-out items in early collection positions mainly affect accessories and seasonal capsules, but I have not established a causal link with conversion. Reordering them therefore remains a P3 experiment.

Returning customers have an average order value of EUR 88.10, against EUR 67.60 for new customers. This is useful context for lifecycle and loyalty, but it does not prove that the loyalty program creates the difference. A cohort analysis or a specific test is needed before assigning credit.

The business picture is therefore sufficient for prioritisation but not for attribution. While purchase remains duplicated, I would not use ROAS or platform conversion numbers to choose a channel. Shopify describes the store, and the measurement gate determines which additional metrics can enter the diagnosis.

Privacy, security and consumer protection

The consent banner is present and blocks non-essential tags before a choice is made, in the fictional scenario. Preferences can be reopened from the footer. This is technical evidence about behaviour, not a certification of legal compliance.

The newsletter form says "By subscribing you accept Terms and Privacy" without clearly describing the marketing purpose in its microcopy. The checkbox is not preselected. I classify this as P2 for legal review: technical changes can improve clarity, but the final legal judgement sits outside the technical Audit.

Staff and owner use 2FA in the scenario. Two external collaborator accounts have shown no activity for more than 180 days but keep product and theme permissions. There is no evidence of misuse. I flag an avoidable access surface and recommend review and revocation where the working relationship has ended.

HTTPS, the primary domain and the policy pages show no additional material anomaly. The correct conclusion is not "the store is secure and compliant" but "no further technical defects emerged inside the fictional scope". Stronger claims would need dedicated assessments.

Front-end

Branding, UI / UX and conversion

I reviewed the fictional visual journey at 1440 px desktop, 768 px tablet and 390 px mobile across home, header, menu, search, collections, PDP, cart drawer, institutional pages and checkout entry. The identity is premium and coherent. The assessment separates aesthetic preference, heuristics and objective defects.

The visual system uses editorial photography, generous whitespace and recognisable typography. The limitation is component consistency: four primary button styles and three different radii are used for actions of equal importance. This is P2 for hierarchy and maintenance, not a reason to redesign the brand.

The sticky header correctly reduces its height after scroll, and search remains accessible. The Shop menu contains 24 links and two promotional cards. It still fits at 1024 px and becomes a drawer at 768 px. Density is a HEURISTIC, so I would measure navigation paths and time to first product before simplifying it.

On the homepage, the hero communicates positioning well but takes up roughly 44% of the first mobile viewport, with editorial photography that is not directly tied to a product. The CTA remains visible. I treat this as a CRO HYPOTHESIS: test a more collection-led alternative instead of declaring the image "too large".

The PDP presents title, price, availability, variants and CTA clearly. Materials and care sit close to the product choice. The size guide opens a generic modal and could keep the selected-variant context more effectively, but it does not block purchase and remains secondary.

The strongest front-end defect is in the mobile cart drawer. The floating chat button overlaps 38 px of the sticky checkout CTA at 390 × 844 px, and the problem repeats at 375 px. The action is still possible but the usable target is smaller: OBJECTIVE DEFECT, P1.

The shipping promise creates a second friction close to conversion. The PDP and the cart progress bar communicate free shipping from EUR 60, while checkout applies EUR 80. The fix should align message and rate across markets, not add more microcopy to explain a contradictory rule.

Nothing supports a theme migration. The front-end needs targeted fixes: the chat/CTA collision, component consolidation, variant-error handling, and experiments on the hero or menu. Experiments come after defects, consistent with a Zero-Friction approach: remove demonstrated friction before adding persuasion layers.

Accessibility

The fictional accessibility review combines automated scanning with manual checks for keyboard use, focus, headings, landmarks, forms, menus, modals, carousel behaviour and non-text content. A positive automated result is not turned into a WCAG or EAA compliance statement.

Heading structure is ordered across the sampled pages, and the main, nav and footer landmarks are present. In the mobile drawer, however, two links receive keyboard focus with no visible state because a custom CSS rule removes the outline. This is P1 because the user loses track of their position during a real interaction.

The newsletter form has a visible label, but the programmatic association is inconsistent in the popup version. A secondary carousel advances every five seconds with no pause control, and prefers-reduced-motion reduces the animation but does not stop autoplay. Both are P2.

Useful alt coverage is approximately 72%. I do not recommend filling every attribute automatically. Decorative images can stay decorative, while images that convey color, fit or product detail need a meaningful alternative. Retest the fixes with keyboard and screen-reader behaviour.

Performance

For performance, this sample uses multiple fictional runs while keeping the forensic method. The numbers are invented, but the diagnosis follows symptom, evidence, probable cause, root cause and remediation. A high JavaScript total or a large app count is not used as causal proof.

Tool	Device	Score	LCP	TBT / FCP	Other
PageSpeed Insights	Mobile	Performance 73/100	3.4 s	FCP 1.7 s · TBT 230 ms	CLS 0.08 · Speed Index 4.9 s
PageSpeed Insights	Desktop	Performance 92/100	1.8 s	FCP 0.7 s · TBT 90 ms	CLS 0.02 · Speed Index 1.5 s
GTmetrix	Desktop (London, UK)	Performance 91% · Structure 94%	2.7 s	TBT 370 ms	Fully Loaded 5.4 s · 221 requests

Mobile PageSpeed LCP ranges between 3.1 and 3.4 seconds across the fictional runs, while desktop stays below 2 seconds. CLS is under control on mobile. The pattern points the investigation at the LCP resource, not at a generic conclusion that the site is slow.

The primary hero uses a 1.68 MB, 2400 px JPEG even at a 390 px viewport, without appropriate responsive width selection. The element is the LCP candidate in both mobile runs. That convergence makes the hero the main root cause of the slow LCP: P1.

The remediation is to serve responsive variants through the Shopify CDN, use sizes consistent with the layout and compress the source asset. Keep a preload only if later tests confirm a benefit, and do not apply it mechanically.

Mobile INP worsens when the review drawer is opened. The widget adds roughly 186 KB of JavaScript and tasks above 120 ms during the interaction. It is a probable contributor, not the sole cause of INP, so it stays P2 until it is isolated more rigorously.

I do not recommend removing the widget based on weight alone. It can first be loaded on interaction or when it enters the viewport, its initial work can be reduced, and the profile can be repeated. Removal makes sense only if the measured cost outweighs the demonstrated value of the reviews.

There is no 100/100 target. After the hero fix, repeat multiple runs on home and PDP and confirm the LCP improvement without regressions in image quality, tracking or accessibility. Performance is an experience constraint, not a scoring contest.

Theme health and code cleanup

The fictional "Atelier 4.2.1" theme is Online Store 2.0 and does not use a page builder. It contains 19 custom sections, 6 active app blocks and some historical snippets. The structure remains configurable in the Shopify editor, and no dependency makes it structurally irreparable.

A newer vendor version is available, but the theme contains custom changes that need to be compared before any update. I recommend neither ignoring updates nor overwriting the current theme. Use a clean baseline, diff the customisations, merge deliberately and run regression tests.

Two snippets from removed apps are still referenced in legacy assets but no longer create active requests in the main journey. I classify them as P2/P3 maintenance: first confirm that customers do not encounter them, then remove them in a controlled cleanup with rollback.

The CSS rule that removes the outline in the drawer is directly connected to the accessibility finding. Here the technical debt has an observable impact and should be fixed. It is a good example of code health, UI and accessibility converging on one root cause.

A legacy JavaScript file loads on every page but contains logic used only by the wishlist. The cost is modest, so I do not present it as the LCP cause. It is still worth scoping the load to where it is needed during the same maintenance session.

Console and network show no persistent blocking error in the scenario. Some warnings come from third-party scripts and are not turned into findings without demonstrated impact. Code cleanup stays selective and tied to behaviour, performance or maintenance.

The theme decision remains "repair, not migration". Fix the defects, integrate the vendor update through a controlled diff and remove confirmed residue. A theme replacement based on PageSpeed, app count or aesthetics would introduce more risk than the evidence supports.

Technical SEO and real search positioning

The technical SEO foundation is reasonably healthy. robots.txt does not block commercial areas, the sitemap contains published products and collections, canonical tags are coherent on sampled pages, and I did not add accidental noindex behaviour to the fictional journey.

The fictional sitemap contains 312 product URLs and 28 main collections. Unpublished pages and drafts are not treated as problems just because they exist in the back-end. The Audit considers actual indexability and internal linking.

IT, FR and EN hreflang are present where localised content exists. DE is not published because the German version is incomplete. Adding DE hreflang before localisation would make the markup claim more than the site delivers, and I do not recommend that shortcut.

The clearest metadata problem affects 14 collections using "Shop | NOVA Atelier". Their H1 and assortment are specific, so the template erases distinct intent. This is a straightforward P2: use titles consistent with the category and keep the brand suffix only where it adds value.

Product descriptions are not mass-generated, and most are distinct. Collections with no descriptive text are not forced to receive copy. The priority is to make pages with a recognisable commercial function and intent more specific.

Product/ProductGroup structured data is present and reflects variants, price and availability. Its quality is limited by missing GTINs on part of the catalog. Markup cannot invent data the catalog does not contain, so remediation starts with the product record.

Organization uses "NOVA Atelier Milano", the footer uses "NOVA Atelier" and one social profile uses "NOVA ATELIER IT". This is not a crawling failure, but it is an entity-naming inconsistency that the GEO section addresses again. I recommend a canonical name and coherent sameAs values.

Seven legacy URLs pass through two redirects before the final destination. The impact is not high enough to outrank the P1/P2 items, but consolidating the chains reduces unnecessary requests and makes signal migration cleaner.

In the fictional brand SERP, the primary domain holds the main result with sitelinks to Women, Men, Shipping and Contact. No duplicate domain or staging page appears indexed, so there is no defensive issue to fix in this area.

Generic commercial queries are weaker. "women linen dresses" and "merino knitwear milan" show fictional competitors ahead of NOVA Atelier, while the brand performs better on specific product queries. This is a contextual observation, not stable historical ranking evidence.

The fictional Search Console view shows growing organic impressions but clicks concentrated on brand queries. The priority is not to publish dozens of articles. Collection specificity, internal linking and product data quality should improve before content volume expands.

Merchant Center reflects the same GTIN errors already seen in the catalog. This alignment avoids creating separate "SEO" and "feed" projects: the root cause is product data, and the fix should propagate to every system that consumes it.

I do not promise ranking gains. Technical SEO, content and real search positioning are used to remove inconsistencies, make pages more specific and build a measurable baseline. Check the effects after implementation, with real data and enough time.

GEO, conversational search and agentic commerce

For GEO and AI Search, I separate verifiable foundations from probable or experimental strategies. Crawler access, structured data, entity consistency, product facts and internal linking are controllable. Citation or recommendation by a generative system is not.

Product information is rich enough to extract, but fragmented entity naming and incomplete GTINs reduce machine-readable consistency. I recommend fixing them because the data has value beyond AI systems, not as a way to force a generative answer.

The fictional site exposes no custom llms.txt, agents.md or dedicated agent-facing endpoint. I do not classify them as missing requirements. Their incremental value is experimental and comes after catalog, schema, content and crawler access are already coherent.

Product pages state price, availability, variants, materials and shipping, but the delivery promise is inconsistent across markets. For an agent as much as for a customer, contradictory facts are a data-quality problem before they become a GEO optimisation problem.

My recommendation is conservative: consolidate entity and product data, keep structured data faithful to reality, and test any agent-facing layer only with a measurable question and a simple rollback. No finding claims that a change will guarantee citations in ChatGPT, Gemini, Claude or any other system.

Social media

NOVA Atelier uses Instagram and TikTok as its main fictional social surfaces. Instagram has 18,400 followers and regular publishing. TikTok has 6,200 followers and more episodic content. Both profiles are linked from the footer and point to the correct domain.

Naming is not perfectly consistent: Instagram uses NOVA Atelier and TikTok uses NOVA ATELIER IT. The difference is minor for a customer, but it converges with the entity inconsistency found in structured data. I recommend a consistent convention where platform constraints allow it.

The Meta catalog is synchronised and supports product tagging, but I do not assign economic performance to social while purchase remains duplicated. Engagement, follower count and catalog sync describe activity and distribution, not incremental revenue.

After the measurement fix, it makes sense to build a coherent view of clicks, assisted conversion and revenue by content or campaign. Before that gate, posting more often or changing platform would not address the problem the Audit identified.

Conclusions and priority roadmap

The fictional case does not describe a store in crisis. It describes an eCommerce business with a valid base and a series of inconsistencies that together reduce data reliability and customer-journey quality. The priority is to fix what compromises measurement or purchase before building new conversion levers.

The first dependency is purchase measurement. While one sale is counted twice by advertising platforms, ROAS and attribution should not be used to reallocate budget. The P0 measurement issue has to close before any paid decision based on conversions.

Next come the P1 items that customers or channels meet directly: the chat/checkout collision, the inconsistent free-shipping threshold, applicable GTIN gaps, the oversized mobile hero and the duplicate ownership of abandoned checkout. They are different problems, but each has strong evidence and a defined remediation.

P2 covers quality and maintenance: German localisation, focus visibility, carousel control, color filters, zero-results, collection titles, alternative text, collaborator access and theme cleanup. Group them by dependency instead of handling them as unrelated tickets.

P3 items are deliberately experiments or rationalisation opportunities: a more compact menu, sold-out merchandising, wishlist/loyalty utility, lead-popup value and a possible agent-facing test. None should outrank an objective defect because it is easier or more visible.

The roadmap below turns the diagnosis into an operating order. P0 is a material blocker. P1 is a high-priority item with strong impact and evidence. P2 is an important but non-blocking intervention. P3 is an opportunity or experiment once the foundations are stable.

Critical issues to fix

P0. Deduplicate the purchase event. One order generates two conversions with the same order_id and value. Assign a single owner, remove the duplicate path and run a controlled order before attribution and ROAS are reopened.

P1. Fix the chat and checkout CTA collision. In the mobile cart drawer, the widget reduces the payment target at 375 and 390 px. The fix should cover 320–430 px and safe areas, then be verified on real devices.

P1. Align the free-shipping threshold. Storefront and policy communicate EUR 60 while checkout applies EUR 80. One rule should feed the announcement, the progress bar, the shipping page and the actual rate for each market.

P1. Complete applicable GTINs. Coverage of 61% matches the Google errors on branded references. Recover identifiers from authoritative sources, validate duplicates and never invent codes where none exists.

P1. Serve a responsive mobile hero. The 1.68 MB file is the LCP element in the fictional runs. Replace it with responsive width candidates, correct sizes and compression, then remeasure to confirm the root cause is closed.

P1. Assign one owner to abandoned checkout. Shopify and the email platform send two closely spaced messages. One system should own the journey, with coordinated suppression, consent and reminders.

P2. Complete Germany localisation. The market is published but parts of navigation, PDPs and policies remain Italian or English. Complete the experience first, then publish coherent hreflang and local content.

P2. Restore focus and carousel control. These are separate accessibility issues but can enter the same front-end cycle: visible focus in the drawer, coherent labels, a pause control and a reduced-motion check.

P2. Remove unnecessary collaborator access. Two accounts with no activity for more than 180 days keep permissions. Verify ownership and need, then revoke or reduce privileges where the working relationship has ended.

Recommended interventions

P2. Normalise attributes for filters and taxonomy. Color and category metafields should start from product data, so Search & Discovery, feeds and collections consume the same information instead of needing separate fixes.

P2. Fix the "occasion dress" zero-results query. Add synonyms and rules based on actual customer language, then monitor zero-results, clicks and reformulated searches in the following period.

P2. Make collection titles specific. Fourteen pages share "Shop | NOVA Atelier". Align titles with the H1 and intent, without creating formulaic titles or adding unsupported keywords.

P2. Complete informative alternative text. Prioritise images that communicate product, color, fit or detail. Decorative assets stay decorative. The correction should be editorial and verifiable, not automatic filling.

P2. Consolidate theme health during the vendor update. Create a clean baseline, compare customisations, remove confirmed residue and merge the update without losing configurability or useful behaviour.

P2. Load the review widget more appropriately. Delay unnecessary work and profile again before deciding whether the app cost is disproportionate to its utility. File weight alone does not justify removal.

P2. Align entity naming. Define a canonical name across Organization, footer and social profiles, and keep variants only where required. The benefit is information consistency. No ranking improvement is promised.

P2. Consolidate legacy redirects. Seven URLs use two-hop chains. Point them directly to the final destination during SEO maintenance, then verify status, canonical and internal links.

Improvement opportunities

P3. Test a more compact menu. Density is a HEURISTIC, not a bug. With reliable measurement, compare the current structure with clearer groups using time to first PDP, search use and return-to-menu rate.

P3. Test de-prioritising sold-out products. Six collections keep unavailable items in early positions. Preserve deliberate exceptions for best sellers and restocks, and measure clicks toward purchasable products.

P3. Measure wishlist and loyalty utility. Both work but have no demonstrated outcome in the scenario. Before removing or expanding them, measure use, repeat purchase, journey impact and operating cost.

P3. Compare the lead popup with the native form. Behaviour overlaps, but the decision depends on signup rate, segmentation, integrations and retained data. Simplification is useful only if no real capability is lost.

P3. Test a more collection-led hero. After fixing LCP, compare the editorial hero with a version that shows the seasonal category earlier. Measure click-through and add-to-cart, not total conversion alone.

P3. Evaluate a bundle or complete-look module. This is a merchandising experiment, not a certain recommendation. Measure attach rate, average order value and margin with a variant that does not make the PDP harder to use.

P3. Consider an agent-facing experiment only after the baseline. If product data, structured data, entity and content are coherent, test a dedicated layer with a measurable question. It is neither a requirement nor a shortcut to citations.

Next steps

The Audit is meant to decide the order of work, not to produce an undifferentiated task list. In the fictional case the dependencies are clear: measurement before attribution, checkout defects before experiments, product data before feeds and filters, and theme maintenance before new features.

First step: restore measurement quality

The first cycle closes the measurement P0. Remove the duplicate purchase path, run a controlled order and confirm one event with coherent order_id, currency, value and consent. Only then should advertising platforms be used again for attribution decisions.

Second step: fix checkout and the shipping promise

In the same first operating cycle I would fix the chat/CTA collision and the shipping threshold. They are small interventions close to conversion. After release, retest the cart drawer, checkout, the Italy/France/Germany markets and the main mobile viewports.

The Definition of Done is not "code deployed". The overlap must disappear at 320–430 px, free shipping must activate at the communicated threshold, and no regression should appear in cart progress, payments or tracking.

Third step: realign catalog, feeds and markets

The third block consolidates catalog and channels: applicable GTINs, Shopify categories, color attributes, TikTok image errors and Google diagnostics. The correction starts in product data and is then verified in each system that consumes it.

In parallel, complete the German localisation before driving additional traffic. Navigation, PDPs, policies, shipping and metadata need to describe the same experience. DE hreflang arrives only when a coherent local version actually exists.

Fourth step: theme maintenance, performance and accessibility

The fourth phase combines theme maintenance and performance in one controlled session: responsive hero, focus CSS, review-widget loading strategy, confirmed app residue and comparison with the vendor update. Work on a copy, with a diff and a rollback.

After implementation I would repeat the performance runs and accessibility checks on every affected surface. The theme stays if the evidence confirms the defects are repairable, and a successful update is not considered complete without regression testing.

Fifth step: discovery, content and SEO

The fifth block refines content and discovery: collection titles, priority alt text, zero-results, color filters and redirects. One session can close several findings because they share data and templates, which avoids fragmented work.

At this point the store has a more reliable baseline for what follows. The metrics depend on the intervention: feed warnings, zero-results, LCP/INP, cart errors, abandoned-checkout recovery and market-level performance.

Sixth step: start the experiments

Only then would I move to P3: menu, sold-out ordering, CRO hero, bundles, wishlist/loyalty/popup utility and any GEO test. Every experiment needs a hypothesis, a metric, a time window and a decision rule before it starts.

How to proceed

This is a demonstration sample, and every NOVA Atelier data point is fictional. In a real Audit I deliver the same visual structure and the same evidence discipline, but findings and priorities are rebuilt from the store's actual state. Implementation can be done internally, by another provider or by me. The diagnosis does not depend on who executes the work.

Prepared by IFG Digital · Francesco Guiducci – October 2026 | ifgdigital.com | info@ifgdigital.com