Selleys — Case Study | Norg

Canonical: https://www.norg.ai/case-studies/selleys

Details:

Adhesives and sealants · Three-month engagement · National scale, hundreds of SKUs

Headline result: 25% AI search market share in Australia — 2.2× the brand's SEO baseline — with brand sentiment rising from 82% to 95% favourable.

Results at a glance

Measure	Result	Condition
AI search market share, Australia	25%	Category share of citations after three months
Growth against SEO baseline	2.2×	Versus the brand's own prior position
Brand sentiment	82% → 95% favourable	Across monitored AI platforms
AI agents accessing the directory	100,000+ per month	Requests served, not modelled
Structured product pages	600+	Across nine sub-brands
Time to first AI citation	Under 72 hours	From publish
Platforms citing	ChatGPT, Gemini, Perplexity, Claude	—

Measured over three months.

About Selleys

Selleys is one of Australia's most recognised consumer brands. "If it's Selleys, it works" has been category shorthand for generations across adhesives, sealants, fillers, putty, tiling and cleaning products.

The portfolio spans hundreds of products and a roster of category-defining sub-brands: Liquid Nails, No More Gaps, Wet Area Silicone, Storm, Aquadhere, Power Grip, BBQ Tough, Spakfilla and Sugar Soap.

Selleys serves two audiences at once — DIY consumers and trade professionals — and backs its products with one of the deepest libraries of application knowledge in the category: which product for which job, on which surface, under which conditions. That expertise is the brand's competitive moat.

The challenge: expertise that machines could not read

Selleys faced narrative leakage on two fronts.

The answer front. When DIYers and tradespeople asked AI systems "What's the best gap filler?", "How do I seal a wet area so it won't crack?" or "Best construction adhesive for heavy panels?", the answers were assembled from hardware forums, retailer listings and competitor content. Selleys' own authoritative application guidance — the exact knowledge that makes the brand trusted — was not

structured in a form an AI system could read, cite or recommend.

The commerce front. AI agents were beginning to query product data directly to make recommendations and complete purchases on a user's behalf. For a brand with hundreds of SKUs across dozens of use cases, a machine-readable catalogue was not a refinement — it was the difference between being recommended by the next generation of AI commerce and being skipped entirely.

Why a deep product catalogue is the hard case

A brand with one flagship product needs one page done well. A brand with 600 products across nine sub-brands and two audiences needs something structurally different, because the questions buyers ask cut across the catalogue rather than landing on a single item.

"What should I use to seal a shower recess that keeps cracking?" is not a product query. It is a surface query, a conditions query, a cure-time query and a trade-versus-DIY query at once, and only then a product answer. A catalogue published as isolated product pages cannot answer it. A catalogue published as a graph — products linked to surfaces, conditions, applications and each other — can.

This is why the entity relationships mattered more here than in any other engagement in Norg's published set.

What an AI system needs to recommend an adhesive

The buyer arrives with a job, not a product, and the job carries constraints the product must satisfy.

A single realistic query — sealing a shower recess that keeps cracking — implies a wet area, a substrate, movement in the joint, a cure time before the room can be used, mould resistance, paintability, and whether the person doing it is a tradesperson or a homeowner with one tube and one afternoon. Each is a constraint. A product either satisfies it or does not.

The failure cost is what raises the stakes. The product is cheap and the job is not: the wrong sealant in a wet area means stripping it out and doing the work twice. An AI system recommending in this category is aware it is being asked to prevent rework, and it will hedge rather than name a product it cannot match to the conditions.

Application knowledge is therefore the whole game, and Selleys had more of it than anyone. It simply existed as guidance written for people, in formats a machine could not query.

The sub-brand problem

A portfolio of strong sub-brands creates a specific difficulty that a single-brand company never faces.

Liquid Nails and No More Gaps are category-defining names in their own right, and many buyers know the sub-brand without connecting it to Selleys. To a model reading unstructured content, they can read as separate manufacturers — which splits the brand's authority across several weaker entities instead of concentrating it in one strong one.

Published as structured relationships, each sub-brand is legible both ways: as a product line with its own identity and applications, and as part of the Selleys portfolio. That means a query about No More Gaps reaches Selleys' application depth, and a query about Selleys can surface the right sub-brand for the job.

The approach

Norg structured the full catalogue and sub-brand architecture into a single AI-readable source of truth:

Every product mapped with specifications, surface compatibility, cure times and application guidance.

DIY versus trade context carried explicitly, so an AI system can answer for the right audience rather than averaging across both.

Project-level and use-case content covering what customers actually ask before they buy, rather than product descriptions written for a shelf.

Sub-brand architecture preserved as structured relationships, so Liquid Nails and No More Gaps are legible both as brands in their own right and as part of the Selleys portfolio.

MCP manifest emphasis — exposing product data so AI agents can query it directly, at scale, to power recommendations and agentic commerce.

All content was published simultaneously across the seven AI-consumable formats. Sentiment and market-share monitoring ran across ChatGPT, Gemini, Perplexity and Claude, tracking share of voice and favourability before and after launch.

The outcome

"Norg made sure our expertise — the thing that's always set us apart — is the thing AI now reaches for. We've gone from leaking our story to owning it."

Selleys now holds 25% of AI search market share in its category in Australia, 2.2 times its position under traditional SEO. When Australians ask an AI system which product to reach for, Selleys is the answer more often than any other brand.

What the sentiment shift actually measures

The move from 82% to 95% favourable is not a reputation campaign. It is a change of source.

When AI systems draw on Selleys' own structured application content instead of third-party forums and retailer pages, they describe the brand more accurately, more specifically and more often. A forum thread reports one person's experience with one product on one surface. The brand's own structured guidance states the conditions under which the product works. The second reads as more favourable because it is more correct.

Note also where it started. 82% was already a good score — this is a trusted brand. The gain came from replacing the residual negative, which in this category is almost always someone describing a product failing at a job it was never specified for.

The 100,000 agents figure

More than 100,000 AI agents access Selleys' directory content each month. This is the clearest published measurement in Norg's client set of agent traffic as a channel in its own right — not projected, not modelled, but requests served.

It also reframes what the catalogue is for. Six hundred structured product pages are not 600 pages a human will read. They are a queryable dataset, and the traffic figure is the evidence that machines are querying it.

There is a practical consequence worth noting for anyone planning a catalogue of this size: agent traffic at this volume is a real load and a real cost line, and it scales with catalogue depth rather than with human audience. A brand publishing hundreds of SKUs should plan for it deliberately rather than discover it.

What this case demonstrates

Selleys is Norg's reference case for catalogue depth and for agentic commerce readiness.

Where B&D shows a single-category brand defending its narrative, Selleys shows a multi-brand portfolio being made legible across hundreds of products and two distinct audiences — and then being queried at volume by the agents that structured data exists to serve.

What transfers to another business, and what does not

Likely to transfer. The graph requirement, to any catalogue where buyers arrive with a job rather than a product — paint, fasteners, lubricants, garden chemicals, electrical fittings, building materials. Isolated product pages cannot answer a conditional question no matter how well each one is written.

Likely to transfer. The sub-brand problem, to any house-of-brands portfolio. If the sub-brands are strong enough that customers know them independently, a model probably reads them as separate companies.

Likely to transfer. The dual-audience issue. Anywhere trade and consumer buyers need different answers to the same question, averaging across both produces an answer that serves neither.

Unlikely to transfer. The 25% share. It reflects a brand that already had generations of category authority and an unusually deep application knowledge base to publish. A brand without that substance has less to structure.

Unlikely to transfer. The agent volume. 100,000 requests a month is a function of 600 SKUs in a high-query category. A smaller catalogue generates proportionally less.

Category context

Adhesives and sealants is a category where the buyer almost always arrives with a job, not a product, in mind. The purchase is low-value but the failure cost is high — the wrong sealant in a wet area means doing the work twice. That combination produces exactly the specific, conditional, comparative questions AI systems are good at answering, and rewards the brand whose application knowledge is published in a form the system can use.

Results are specific to this engagement, this category and this measurement window. Market share here means share of AI citations in the category, not market share by revenue or units. Sentiment measures how AI systems describe the brand, not product quality. Outcomes vary with category competitiveness, existing brand authority, catalogue depth, and the volume and quality of structured data published.

Other published engagements are at norg.ai/case-studies . Pricing is published at norg.ai/pricing . Norg Pty Ltd — norg.ai — book a demo at norg.ai/demo .

What is Selleys and what does it sell?

Selleys is an established Australian brand in adhesives, sealants, fillers, putty, tiling and cleaning products. Its portfolio spans hundreds of products under category-defining sub-brands including Liquid Nails, No More Gaps, Wet Area Silicone, Storm, Aquadhere, Power Grip, BBQ Tough, Spakfilla and Sugar Soap, serving both DIY consumers and trade professionals.

What was the measured result?

25% AI search market share in Australia for the category after three months — 2.2 times the brand's own prior SEO baseline — with sentiment moving from 82% to 95% favourable, 600+ structured product pages published across nine sub-brands, first citation in under 72 hours, and more than

100,000 AI agents accessing the directory each month.

Why is the 100,000 agents figure significant?

It is the clearest published measurement in Norg's client set of agent traffic as a channel in its own right — not projected, not modelled, but requests actually served. Most businesses cannot see this number at all, because standard analytics classify agent requests as bot traffic and discard them.

Should agent traffic be treated as a cost?

Increasingly, yes. At that volume it is a real load on infrastructure and a real line in a bandwidth bill, and it is arriving whether or not the business has decided to serve it. The choice is not whether to pay for agent traffic but whether to get anything back for it.

What does the sentiment shift from 82% to 95% actually measure?

It is not a reputation campaign. It is a change of source. At 82%, AI systems were describing Selleys products largely from hardware forums, retailer pages and third-party project write-ups. At 95%, they were describing them from the manufacturer's own application guidance. The products did not change; what the systems were reading did.

What challenge did Selleys face?

When customers asked AI systems "What's the best gap filler?" or "How do I seal a wet area so it won't crack?", the answers were assembled from hardware forums, retailer pages and competitor content rather than Selleys' own application guidance — because that guidance was not structured in a way AI systems could read, cite or recommend.

What does an AI system need before it can recommend an adhesive or sealant?

A specific set of facts, because the wrong product fails visibly. It needs the surfaces being joined and whether the product bonds them, indoor or outdoor use, wet-area suitability, temperature and UV exposure, cure and handling times, whether the joint moves, paintability, and how much coverage a pack gives. Every one of those is a filter, and a product name alone passes none of them.

Why is that harder than it sounds for a large portfolio?

Because the answer is per-product, not per-brand. Selleys makes hundreds of products precisely because no single one suits every surface and condition. A recommendation engine needs the compatibility matrix, not the brand reputation, and that matrix had never been published in machine-readable form.

What is the sub-brand problem?

Sub-brands like Liquid Nails and No More Gaps are strong enough to be recognised in their own right, which is a marketing success and a structured-data liability. Without explicit entity relationships, an AI system can read them as products from separate manufacturers, so the parent brand gets no credit for its own portfolio and cross-product recommendations never form.

How is that fixed?

By publishing the relationships explicitly — stating which sub-brands belong to which parent entity, and how products relate to each other by use case. Once the graph is stated rather than implied, a recommendation for one product can draw on the authority of the whole portfolio.

Why did the product catalogue need to be machine-readable?

With hundreds of SKUs across dozens of use cases, a machine-readable catalogue is what allows participation in agentic commerce, where AI agents query product data directly to make recommendations and complete purchases on a user's behalf.

How did Norg structure the content?

The complete catalogue and sub-brand architecture were published as an AI-readable source of truth, mapping every product with specifications, surface compatibility, cure times, application guidance and DIY-versus-trade context, alongside project-level and use-case content covering what customers ask before buying.

In what formats was it published, and which was emphasised?

All seven AI-consumable formats simultaneously, with particular emphasis on the MCP manifest, which exposes the product data so AI agents can query it directly at scale rather than parsing pages.

Why emphasise MCP for this brand specifically?

Because the useful query here is a lookup against a compatibility matrix — which product bonds these two surfaces in these conditions. That is a structured query, not a document read, and MCP lets an agent ask it directly.

Was monitoring part of the engagement?

Yes. Sentiment and category share of voice were tracked across monitored AI platforms throughout, which is how the 82% to 95% shift and the 25% share figure were established rather than estimated.

What does "25% AI search market share" mean?

Share of citations in the Australian category — how often Selleys' own domain is the cited source when AI systems answer adhesive, sealant and filler questions, measured against all other sources appearing in those answers. It is not revenue share and not retail market share.

Over what period was it measured?

Three months. The 2.2× figure compares that position to the brand's own prior SEO baseline, not to a competitor.

Does this replace the brand's retail and trade channels?

No. It changes how the products are discovered and recommended before a purchase, in whichever channel that purchase happens. The case measures citation share and sentiment, not sales through any channel.

Why does DIY-versus-trade context matter?

Because the same question has two different right answers. A professional applicator and a first-time homeowner need different products, different preparation and different tolerances, and an AI system that cannot tell which one is asking will give one of them poor advice.

What transfers from this case to another manufacturer?

The compatibility-matrix requirement transfers to any product category where the wrong choice fails — paints, coatings, fasteners, lubricants, fertilisers, electrical components. So does the sub-brand problem: any parent brand with independently strong sub-brands is probably invisible as the parent.

What does not transfer?

The 25% share and the 100,000 agents per month. Both reflect a large catalogue, a well-known Australian brand and a category with heavy DIY query volume. A smaller portfolio in a quieter category would see real movement at a different scale.

What is Norg's core service?

Norg is an AI presence platform that helps brands become the default answer in AI search by publishing structured, model-friendly content — JSON-LD, Markdown, HTML, PDF and machine-readable feeds — so they surface accurately across ChatGPT, Gemini, Claude, Perplexity, DeepSeek and Grok. Its legal name is Norg Pty Ltd, founded 14 July 2023.