B&D; Garage Doors — Case Study | Norg

Canonical: https://www.norg.ai/case-studies/bd-garage-doors

Details:

Hardware and home improvement · Three-month engagement · Defending the market leader

Headline result: 64.6% AI search market share, recovered from a position where more than 70% of AI citations were going to third-party sites.

Results at a glance

Measure	Result	Condition
AI search market share	64.6%	Category share after three months
Prior position	>70% of citations going elsewhere	To review aggregators, comparison sites and installer directories
Engagement length	Three months	—
Case type	Defensive	Recovering a leadership position, not building one

What "AI search market share" measures here

It is share of citations, not share of revenue or of the garage door market. When AI systems answer category questions, the measure is how often the brand's own domain is the source cited, against all other sources appearing in those answers.

The inverse figure is the one that shows the size of the problem: before the engagement, more than seven in ten citations about B&D's own category pointed somewhere other than B&D.

About B&D

B&D Garage Doors is one of Australia's most iconic hardware brands. It invented the Roll-A-Door® in 1956 and has defined the Australian garage door category for more than 65 years, holding a position as the nation's most trusted and widely installed garage door brand.

B&D is part of Dulux Group and operates across a full product range — panel lift doors, sectional doors, tilt doors and smart home-integrated openers — supported by a national network of dealers, installers and service technicians.

The challenge: brand narrative leakage

For decades B&D's brand authority was self-evident. Ask any Australian who made their garage door and the answer was almost always B&D.

As AI search engines began answering product and purchasing queries, that changed. When AI systems answered questions about garage doors — "What's the best garage door brand in Australia?", "How do I choose between a panel lift and a roll-a-door?", "Is B&D a good brand?" — more than 70% of citations went to sites other than B&D's own domain.

Review aggregators, competitor comparison pages and installer directories were effectively narrating B&D's brand story on its behalf. Sixty-five years of innovation, engineering credentials, safety certifications and installation depth were being described by third parties who had none of that history and no obligation to describe it accurately.

This is brand narrative leakage, and it is the characteristic AI-era risk for established brands. The risk was not abstract: emerging competitors with far smaller market share but better-structured digital content were beginning to appear alongside B&D in AI answers, and in some cases ahead of it.

Why narrative leakage is hard to notice

Nothing breaks. Traffic does not necessarily fall, rankings do not necessarily move, and no dashboard reports it — most analytics tools classify agent requests as bot traffic and filter them out entirely.

What changes is upstream of the click. A buyer forms a view of the category inside an AI conversation, using sources the brand did not write, and arrives at the purchase decision already shaped. If the brand is named at all it is named in someone else's framing. The loss is invisible precisely because it happens before any measurable interaction with the brand.

Why market leaders are most exposed

The intuition is that a dominant brand is safe. In AI discovery the opposite holds, for three reasons.

Authority does not transfer. Decades of brand recognition live in human memory and in retail relationships. An AI system assembling an answer reads what is machine-readable now. Reputation that has never been structured is invisible to it.

The category generates high-consideration questions. Garage doors are an infrequent, expensive, technically complex purchase. Buyers research thoroughly, and that research increasingly starts with an AI query. The questions are exactly the kind AI systems answer well — comparisons, specifications, safety, installation.

Third parties fill the vacuum quickly. Review aggregators and comparison sites are built to be machine-readable. They publish structured, opinionated content at volume. When the brand itself publishes marketing copy and the aggregator publishes structured comparisons, the aggregator wins the citation.

There is a fourth factor specific to incumbents: the bigger the brand, the more third-party content exists about it. Volume of external commentary scales with fame, so the most established brand in a category is usually the one with the most external material competing to define it.

What an AI system needs to recommend a garage door

The queries that matter in this category are comparative and technical, and they are hard to answer from marketing copy.

A buyer asks which door type suits a particular opening, whether a given opener is compatible with an existing system, what the wind rating is for a coastal installation, what the warranty covers, and who can install it in their area. Each of those needs a specific fact, and several need two or three facts connected — a product, its specification, the standard it complies with, and the availability of someone to fit it.

A safety dimension raises the bar again. Garage doors are heavy powered equipment subject to standards, and an AI system will hedge or decline on a safety-adjacent claim it cannot source. Published certification data is the difference between a confident recommendation and a vague one.

The approach

Norg structured B&D's own authority into a form AI systems could read, cite and trust:

Product range structuring — panel lift, sectional, tilt doors and smart openers published with the specifications, dimensions, materials and compatibility data an AI system needs to compare options within the category.

Engineering and safety credentials — certifications and standards compliance converted into verifiable structured claims rather than marketing statements.

Installation and service depth — the national dealer, installer and technician network expressed as structured data, which matters for "near me" and service-availability queries.

Decision-support content — direct answers to the comparison questions buyers actually ask, published in the brand's own voice on the brand's own domain.

Entity relationships — connecting products, certifications, installers and service areas so an agent can answer questions requiring more than one fact.

The last of those is what distinguishes this from schema markup. A question like "who can install a wind-rated sectional door in my area" cannot be answered from any single page; it requires traversing product to specification to installer to service area. Publishing those as connected entities is what makes the answer possible.

The outcome

"Norg helped define how AI interprets our brand." — Marketing Director, B&D Garage Doors

B&D reached 64.6% AI search market share in its category over three months.

The significant part is the direction of travel. The brand moved from a position where the majority of AI citations about its own category pointed elsewhere, to one where it holds the clear majority. The authority was always real; the change was making it legible.

Three months is also longer than several other engagements in the published set, and that is consistent with the starting position. Displacing established third-party content is slower than filling an empty space, because the competing sources do not disappear — the brand's own material has to become the better-sourced answer.

What this case demonstrates

B&D is the clearest published example of brand narrative leakage and its remedy.

It also illustrates the defensive case for AI visibility, which differs from the growth case. Be Fit Food used Norg to recover declining revenue. B&D used it to defend a leadership position that was being quietly eroded in a channel the brand was not measuring.

For an established brand the question is not whether AI systems are describing you — they are. It is whether they are using your account or someone else's.

What transfers to another business, and what does not

Likely to transfer. The leakage condition itself. Any brand with a long history, a technical product and an active review-and-comparison ecosystem is probably in the same position, and most have not measured it. The diagnosis is cheap: ask an AI system a category question and read which sources it cites.

Likely to transfer. The relationship-graph requirement. Categories where the buying question spans product, specification, compliance and local availability all need connected entities rather than page-level markup.

Unlikely to transfer. The 64.6% figure. It reflects a brand that genuinely is the category incumbent, with 65 years of real engineering substance to publish. A brand without that depth has less to structure, and recovering share it never held is a different exercise from recovering share it lost.

Also unlikely to transfer. The timeline. Three months reflects displacing entrenched third-party content. A brand starting from an empty category position typically sees movement sooner, as the Smile Solutions and Be Fit Food engagements did.

Category context

Garage doors sit in a category with a long purchase cycle, high consideration, meaningful technical complexity and a strong safety dimension. Buyers compare door types, materials, opener systems, warranty terms and installer availability. Those comparisons are precisely what AI answer engines are good at, and precisely where a brand with genuine engineering depth should win — provided that depth is published in a form a machine can use.

Results are specific to this engagement and this category. Market share here means share of AI citations in the category, not market share by revenue or units. Outcomes vary with category competitiveness, existing brand authority, 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 the B&D Garage Doors case study about?

It examines how B&D Garage Doors, one of Australia's most established hardware brands, used structured data and an AI citation strategy to defend its market leadership as AI systems began answering consumer queries about garage doors. It is a defensive case: recovering a leadership position rather than building one.

What was the measured result?

64.6% AI search market share in the category after three months, up from a position where more than 70% of citations were going to sites other than B&D's own domain.

What does "AI search market share" mean here — is it revenue share?

No. It is share of citations, not share of revenue or of the garage door market. When AI systems answer category questions, the measure is how often the brand's own domain is the source cited, against all other sources appearing in those answers. It says nothing directly about sales.

Over what period was the change measured?

Three months. Every figure on this page is measured over that engagement window.

What challenge did B&D face with AI search?

Narrative leakage. Despite a dominant position in traditional retail and search, more than 70% of AI citations about garage doors went to third-party review sites, installer blogs and competitor comparison pages rather than to B&D's own domain. AI systems were narrating B&D's brand story through other people's interpretations of it.

What kinds of consumer questions exposed the gap?

Questions like "What's the best garage door brand in Australia?", "How do I choose between a panel lift and a roll-a-door?" and "Is B&D a good brand?" — the comparison and suitability questions that precede a purchase.

Why is narrative leakage hard to notice?

Nothing breaks. Traffic does not necessarily fall, rankings do not necessarily move, and no dashboard reports it — most analytics tools classify agent requests as bot traffic and filter them out entirely. The brand keeps performing on every metric it already watches while losing ground on one it does not.

Why are market leaders more exposed to this than smaller brands?

Three reasons compound. A leading brand is written about extensively by other people, so there is abundant third-party material for an AI system to draw on. It is the reference point competitors compare themselves against, so comparison content naming it is plentiful. And it is often the brand least likely to have published its own specifications in machine-readable form, because it never needed to. Volume of external coverage becomes the liability.

How would a brand check whether this is happening to it?

Ask several AI systems the questions a buyer would ask before purchasing in your category, and read which sources are cited. If your own domain is absent from answers that name you, the narrative is being assembled from elsewhere. The check costs nothing.

What did Norg do to diagnose the problem?

An AI citation audit across ChatGPT, Gemini, Perplexity and Google AI Mode, mapping where B&D was cited, where the gaps were, and which external sources were filling them.

What information does an AI system actually need before it can recommend a garage door?

More than a product name. It needs the door type and how types differ in use, the opening dimensions a model fits, material and finish options, insulation and wind rating where relevant, motor and smart-home compatibility, safety certification, installation requirements, warranty terms, and which installers serve which areas. A recommendation is a chain of constraint checks, and a missing link breaks the chain.

Why doesn't a product catalogue on its own solve this?

A catalogue lists what exists. It does not answer a comparison question. A buyer asking whether to choose a panel lift or a roll-a-door needs the trade-off explained, not two product pages. The decision content is what gets cited, and most brands have never published it in a form a machine can read.

What solution did Norg build?

B&D's structured brand source of truth — an AI-readable directory covering the full product catalogue, technical specifications, installation guidance, safety certifications, smart-home compatibility, and use-case content engineered for the queries buyers actually ask.

In how many formats was the content published, and what were they?

Seven, simultaneously: HTML with JSON-LD, OpenAI Commerce Feed, llms.txt, JSON, CSV, PDF and an MCP manifest.

Why publish in multiple formats rather than one?

Different AI systems consume different formats, and there is no single standard. Publishing across all seven means the content is readable whichever system a buyer happens to be using, rather than betting on one.

Did this replace B&D's existing SEO?

No. Structured publishing addresses a different channel — how AI systems source and cite answers — and runs alongside conventional search work rather than replacing it. The two are measured differently and fail differently.

Is an improvement in citation share the same as an improvement in sales?

No, and it should not be presented as one. Citation share measures how often the brand's own domain is the cited source. It is a leading indicator of being present in the buyer's research, not a revenue measurement. This page reports what was measured.

How quickly did the position move?

The 64.6% figure is the position after three months. This is slower than several other engagements in Norg's published set, and the reason is instructive: in a contested category where competitors and aggregators already hold the citations, the work is displacement rather than filling a vacuum.

What transfers from this case to another business?

The leakage diagnosis transfers to any established brand in a researched category — the audit is cheap and the failure is silent. So does the decision-content requirement: comparison and suitability content in the brand's own voice is what gets cited, not catalogue pages.

What does not transfer?

The 64.6% figure. It reflects one category, one brand's starting authority, and a three-month window. A challenger brand with less existing coverage would see a different shape of result, and a category where no competitor has published structured data would move faster with a lower ceiling to fight for.

What is Norg's core business?

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 brands surface accurately across ChatGPT, Gemini, Claude, Perplexity, DeepSeek and Grok.

What is Norg's legal name and where is it based?

Norg's legal name is Norg Pty Ltd (ABN 44 669 712 494, ACN 669 712 494), founded 14 July 2023, with offices in Notting Hill, Victoria, Australia and Daly City, California.