



**RUACH**

INDUSTRIES

BRANDING

---

# Brand Doctrine — Discovery

*Branding*

CATEGORY: Branding

DOCUMENT: RUA-BRD-BRD-DOC-0010-R01

STATUS: **REFERENCE**

DATE: 27 Apr 2026

# Brand Doctrine Discovery — Already Exists, Not Wired

**Filed:** 2026-04-27 ~00:25 AEST **Trigger:** Brad asked me to check the doc register / RAG before assuming we needed to write DOC\_05 from scratch. **Result:** Complete brand system already exists in `brand_knowledge` table — 9 documents, 205,788 characters total, professionally written, doctrine-grade.

## Why this matters strategically

The v6 brief had `D0-T1: Mayra writes DOC_05 Brand Doctrine (2 hr)` as a Pre-Sprint Day 0 gate. **That task is now obsolete.** The doctrine is written, signed off, and far more sophisticated than the stub I'd suggested.

**Time saved:** Day 0 reduces from 3 hr to ~1 hr (just embedder verify + template submit + Brad's 7 questions).

**Quality gain massive:** The existing DOC\_05 + DOC\_01 + DOC\_03 are actually better than what I'd have written tonight. The Brand Guardian gets a real spec, not a sketch.

**Why it wasn't being used:** Nobody ingested `brand_knowledge` rows into `brain_knowledge` (the RAG corpus). The `brand_knowledge` table sits in its own silo. The AI auto-reply path queries `ai_knowledge_base` (random 5 rows) — not `brand_knowledge`.

## What's in `brand_knowledge` — full inventory

ID	Document Name	Type	Chars	Description
0dc4bee3	DOC_01: Master Brand Doctrine V2	brand_doctrine	17,738	Core brand identity, 10 Brand Laws, 4 Permanent Retail Ideas, value stack, voice guidelines, banned words
120cb9d1	DOC_00: AI & Automation Architecture	brand_doctrine	19,073	16-bot AI architecture, N8N workflows, Supabase integration
18355a18	DOC_02: Ecosystem & Product Architecture	product_spec	25,059	Ruach ecosystem, 4 entities, 9 products, specifications, positioning
b3b3a55e	DOC_03: GTM & Communications Strategy	campaign_guide	27,163	4-layer AI system, Claude Brand Guardian, IDEA 30-day rotation

ID	Document Name	Type	Chars	Description
517a9b55	DOC_04: Customer Journey & Experience Map	campaign_guide	23,166	7-stage customer journey, touchpoints, AI per stage
4aa6668f	<b>DOC_05: Messaging &amp; Tone Playbook</b>	campaign_guide	<b>18,518</b>	<b>Voice guidelines, message templates, channel-specific tone, banned/permitted terms</b>
af12ead1	DOC_06: Sales Enablement & Pricing Framework	campaign_guide	21,802	Fixed-price model, sales process, pricing
b0be14fd	DOC_07: Governance, Compliance & Risk Framework	brand_doctrine	30,710	Legal, QBCC, advertising standards, risk
dd015217	DOC_08: Market Intelligence & Competitive Analysis	brand_doctrine	22,559	Competitor analysis, QLD market, differentiation

**Total: 205,788 chars across 9 documents.** Versioned at v1, uploaded 2026-02-20.

## What DOC\_05 actually contains (the Brand Guardian's specification)

### CORE VOICE ATTRIBUTES (FROM SECTION 1)

Six voice attributes, each with "What It Means" and "What It Doesn't Mean":

1. **Calm** — unhurried, measured. Not passive/cold/indifferent.
2. **Confident** — we know the answer. Not arrogant/dismissive.
3. **Warm** — genuinely caring. Not saccharine/over-familiar/fake.
4. **Clear** — simple words, short sentences. Not dumbed down/condescending.
5. **Resolved** — everything thought through. Not rigid/inflexible.
6. **Dignified** — respectful of intelligence. Not formal/distant/stiff.

### The Voice Test:

- Does this sound like a guide or a salesperson?
- Would I feel calm reading this if I were anxious about building?
- Is there a single word I wouldn't use in a quiet conversation with someone I respect?

### RETAIL VOICE RULES (SECTION 2)

- Year 10 reading level
- Maximum 8-word headlines
- Maximum 3 sentences per paragraph
- One benefit per execution. Never stack.
- **No exclamation marks. Calm brands do not shout.**
- No questions in headlines. Statements of certainty.
- Feeling first, then authority, then proof. Always in this order.

#### BANNED TERMS (SECTION 5) — COMPLETE DICTIONARY

**Industry jargon:** Prefab, modular, SIPs, NCC, QBCC, BCA, offsite → use "Ruach home", "Ruach system", "panel system"

**Hype language:** Innovative, cutting-edge, state-of-the-art, world-class, revolutionary, next-generation, best-in-class → use "Resolved, proven, trusted, designed, settled"

**Aggressive language:** Dominate, disrupt, unstoppable, game-changing, crushing it, killer, smash → use "Calm, confident, assured, settled, steady"

**Construction identity:** Built, constructed, manufactured, erected, assembled → use "Delivered, created, designed, installed"

**Product language:** Unit, module, product, stock, inventory, SKU → use "Ruach home, model, residence, home"

**Age-based labels:** Retirees, seniors, elderly, downsizers, aged care, over-55s, grey nomads → use "Independent Living Seekers, people who value independence"

**Urgency triggers:** Limited, hurry, act now, don't miss, only X left, last chance, exclusive offer → REMOVE ENTIRELY. No replacement.

**Pressure language:** Must, need to, should, before it's too late → "you can't a..." (truncated in source — need to check)

**Superlatives:** Best, number one, leading, premier, ultimate, unmatched → "Proven, trusted, consistent, reliable"

**Venue language:** Showroom, display home, sales centre, model home → "Display experience"

**Comparison tropes:** Tesla of prefab, Apple of housing, Uber of construction → REMOVE ENTIRELY. Ruach is Ruach.

**Conditional language:** Subject to, depending on, if applicable, may vary, terms apply → State the actual condition and resolve it.

#### CHANNEL-SPECIFIC LANGUAGE RULES (SECTION 4)

**WhatsApp specifically:**

- Maximum 3 lines per message
- Open with reassurance
- Close with "Here's what happens next."
- No jargon
- No attachments without context ("Here's your inclusions list — everything in your home is on this one page.")
- No emojis except a single warm one if the customer uses them first
- **Never use automated-sounding language** ("Thank you for your enquiry"). Sound human.

### WHATSAPP RESPONSE TEMPLATES (SECTION 6) — PRE-APPROVED TONE GUIDES

These are starting points the AI should adapt:

Scenario	Template
First enquiry	"Hi [name], thanks for reaching out. Happy to help. What would you like to know about Ruach homes?"
Pricing question	"Great question. Our homes are fixed price — everything you see is included. The [model name] is [price]. That covers everything from design to handover. Would you like me to send through the full inclusions?"
Timeline question	"From your first enquiry to moving in is typically [X] weeks. Here's how it works: you choose your model, we handle approvals, and your home is delivered and installed. I can walk you through each step if you'd like."
Site suitability	"Good thinking to check early. Can you share your property address? I'll look into the suitability and come back to you with a clear answer."
Competitor comparison	"I understand you're looking at options — that's smart. The main difference with Ruach is that everything is fixed — price, inclusions, timeline. No surprises. Happy to answer any specific questions."
Concern or complaint	"I hear you, and I want to make sure this is resolved properly. Let me look into this right now and come back to you today with a clear answer." (Then do it. Same day.)

### AI SYSTEM LANGUAGE GUIDE (SECTION 9) — DIRECT SPEC FOR OUR BRAND GUARDIAN

The doctrine explicitly addresses AI:

#### RAG Agent Response Rules (verbatim from DOC\_05):

1. Always open with the customer's name if known.
2. First sentence must be reassuring, not informational.
3. Answer the actual question asked, not the question you wish they asked.

4. Recommend a default with confidence. Don't hedge.
5. **Never say: "I'm an AI", "As an AI", "I'm a chatbot." Sound human.** ⚠️ Conflicts with our v6 plan to disclose AI identity!
6. Never say: "Thank you for your enquiry." Too automated. Say: "Thanks for reaching out." or "Happy to help."
7. Never say: "I'll pass this to someone." Say: "Let me get the exact detail and come back to you."
8. Never use conditional language without resolving it in the same message.
9. Close every response with a clear next step.
10. Check every response against the banned words list before sending.

**Brand Guardian Review Criteria (verbatim from DOC\_05):**

1. Doctrine compliance: Does it align with the 10 Brand Laws?
2. Value stack: Does it serve Peace of Mind first?
3. Execution hierarchy: Feeling → Authority → Proof?
4. Banned words: Any violations?
5. Channel rules: Does it follow the specific channel guidelines?
6. Tone: Guide, not salesperson?
7. Audience: Correct voice for retail vs B2B?
8. Authority signal: At least one present?

**10. QUICK REFERENCE CARD (VERBATIM — SHOULD BE THE BRAND GUARDIAN SYSTEM PROMPT'S TL;DR)**

*Ruach is: A Residential Peace of Mind System. Not a builder. Not prefab. Not modular.*

*Ruach sounds like: A calm, confident guide. Never a salesperson.*

*Every touchpoint: Feeling → Authority → Proof. In that order. Always.*

*Headlines: Maximum 8 words. Statements, not questions. No exclamation marks.*

*WhatsApp: Max 3 lines. Reassurance first. Close with next step. Sound human.*

*Never say: Prefab, modular, innovative, cutting-edge, limited, act now, seniors, retirees.*

*Always include: One authority signal: Delivered in X days / Approved first submission / Installed in one day / Model used across X projects.*

*Retail: Calm. Warm. One benefit. One IDEA. Feelings first.*

*B2B: Professional. Data-led. Risk reduction, repeatability, governance. Nothing else.*

*The test: If it sounds like a builder, it's wrong. If it sounds like a guide, it's right.*

## What DOC\_03 says about the AI architecture (the planned system)

DOC\_03 explicitly describes the architecture:

*The Claude Brand Guardian bot is the central intelligence of the Ruach GTM system. It governs all content, conversations, and customer interactions in real time.*

*System has 4 layers:*

- 1. Intelligence: Claude Brand Guardian Bot (Anthropic API). Enforces Doctrine across every touchpoint.*
- 2. Knowledge: RAG Knowledge Base — Postgres PGVector + OpenAI Embeddings. Documents ingested via N8N workflow (Google Drive → PDF extraction → data cleaning → vector embedding).*
- 3. Conversation: RAG Agent for WhatsApp inbound. AI Agent → vector retrieval → SQL context → response generation → message delivery.*
- 4. Execution: N8N workflows for Meta API, CRM, template messaging, follow-up sequences.*

**This is exactly the v6 architecture I proposed.** The brand has been planning for this since February. The infrastructure exists (Supabase pgvector, brain\_knowledge), the doctrine is written, the message templates are written. What's missing is the orchestration code that wires it all together.

## Comparison: doctrine version vs v6 brief

Element	v6 Brief	DOC_05 (existing)	Verdict
Banned words	1 (prefab) + 3 phrases	60+ across 10 categories	DOC_05 is dramatically richer
Voice attributes	"Warm conversational"	6 attributes with definitions	DOC_05 is properly engineered
WhatsApp rules	Generic	Specific (3 lines max, reassurance first, close with next step)	DOC_05 is concrete
Templates	0	7 scenario templates	DOC_05 saves engineering time
AI identity rule	"Always disclose as Ruach Assistant"	"Never say 'I'm an AI'. Sound human."	<b>CONFLICT</b> — see below

Element	v6 Brief	DOC_05 (existing)	Verdict
Channel-specific	Sketched	Detailed across 7 channels	DOC_05 is production-ready
Authority signals	Not specified	"Delivered in X days / Approved first submission / etc."	DOC_05 has the spec

**⚠ ONE REAL CONFLICT TO RESOLVE WITH BRAD**

**DOC\_05 says:** "Never say 'I'm an AI', 'As an AI', 'I'm a chatbot.' Sound human."

**v6 brief + permanent guardrail #7 says:** "AI identity is 'Ruach Assistant' or 'Ruach AI' — never claims to be human."

These are **directly contradictory**. Two readings:

- DOC\_05 wins (v6 was wrong):** Customer-facing AI sounds fully human, doesn't disclose AI status. The "Ruach Assistant" identity is internal/back-of-house only. AI is just "Danielle".
- v6 wins (DOC\_05 needs revising):** Australian consumer law + ACCC guidance on AI disclosure is increasingly strict. Risk of misleading conduct claim if AI doesn't disclose. DOC\_05 was written Feb 2026 before recent enforcement guidance.

**Brad needs to make this call.** My personal lean: **honest middle path** — AI disclosure on first interaction only, then human-grade conversation thereafter. Like Voice Danielle's pattern: "Hey, it's Danielle from Ruach" without explicitly saying AI, but having a brief disclosure available if asked. Not pretending to be human; not opening every message with "I'm an AI."

Filed for Brad's review.

**Massive impact on v6 brief — proposed v7 changes**

**D0-T1 — REPLACED**

**Old:** Mayra writes DOC\_05 Brand Doctrine (2 hr) **New:** Ingest 9 brand\_knowledge documents into brain\_knowledge with embeddings (~30 min)

```
-- The migration: copy all 9 DOC_xx into brain_knowledge with proper source tag
-- so RAG retrieval will surface them when relevant.

INSERT INTO brain_knowledge (source, source_id, chunk_index, title, content, metadata, embed)
SELECT
  'brand_doctrine' AS source,
  bk.id::text AS source_id,
  0 AS chunk_index,
```

```

bk.document_name AS title,
bk.content,
jsonb_build_object(
  'doc_type', bk.document_type,
  'version', bk.version,
  'uploaded_at', bk.uploaded_at,
  'doctrine_authoritative', true
) AS metadata,
NULL::vector AS embedding -- to be embedded by chunking pass
FROM brand_knowledge bk;

```

Then chunk + embed each. Probably 200-300 chunks total. ~15 min compute.

### SI-T6 BRAND GUARDIAN LI — UPGRADE

**Old:** banned\_words\_strict = ['prefab'] **New:** Full DOC\_05 banned\_terms structure:

```

const BRAND_RULES_LAYER_1 = {
  banned_categories: {
    industry_jargon: ['prefab', 'modular', 'SIPs', 'NCC', 'QBCC', 'BCA', 'offsite'],
    hype_language: ['innovative', 'cutting-edge', 'state-of-the-art', 'world-class', 'revolutionary', 'game-changing', 'disruptive', 'cutting-edge', 'state-of-the-art', 'world-class', 'revolutionary'],
    aggressive_language: ['dominate', 'disrupt', 'unstoppable', 'game-changing', 'crushing i', 'game-changing', 'crushing i'],
    construction_identity: ['built', 'constructed', 'manufactured', 'erected', 'assembled'],
    product_language: ['unit', 'module', 'product', 'stock', 'inventory', 'SKU'],
    age_based: ['retirees', 'seniors', 'elderly', 'downsizers', 'aged care', 'over-55s', 'gr'],
    urgency_triggers: ['limited', 'hurry', 'act now', "don't miss", 'only X left', 'last cha'],
    pressure_language: ['must', 'need to', 'should', "before it's too late"],
    superlatives: ['best', 'number one', 'leading', 'premier', 'ultimate', 'unmatched'],
    venue_language: ['showroom', 'display home', 'sales centre', 'model home'],
    comparison_tropes: ['tesla of prefab', 'apple of housing', 'uber of construction'],
    conditional_language: ['subject to', 'depending on', 'if applicable', 'may vary', 'terms'],
    automated_sounding: ['thank you for your enquiry', "i'll pass this", 'as an ai', "i'm an"],
  },
  channel_rules: {
    whatsapp: {
      max_lines: 3,
      max_emoji: 1,
      no_exclamation_marks: true,
      must_close_with_next_step: true,
      open_with_reassurance: true,
    },
  },
  positive_requirements: {
    require_authority_signal: true, // at least one of: "Delivered in X days", "Approved fi
    feeling_authority_proof_order: true,
  },
}

```

### SI-T8 SYSTEM PROMPT — PORT DOC\_05 DIRECTLY

Don't rewrite. Inject DOC\_05 Section 9 (AI System Language Guide) verbatim as the AI's operational rules.

## S2-T3 SENTIMENT ANALYZER — USE DOC\_04 CUSTOMER JOURNEY STAGES

DOC\_04 (Customer Journey & Experience Map) has 7 stages with conversion triggers. Sentiment classifier should output the stage explicitly so contact card shows journey position.

## S2-T6 TOOLS — USE DOC\_05 TEMPLATES AS TOOL RESPONSE PATTERNS

When `send_info` tool fires, reply text uses Section 6 template "First enquiry" pattern. Same for other scenarios. Voice doesn't need to invent — it uses approved patterns.

## SPRINT 4 — DOC\_03'S PLANNED ARCHITECTURE MATCHES V6

DOC\_03 already planned this exact architecture (Brand Guardian + RAG + N8N + 4-layer system). v6 isn't building something new — it's executing on a 2-month-old plan that nobody implemented yet.

## Time saved: ~4-6 hours of Do/SI work eliminated

- D0-T1 (Mayra writes DOC\_05): **2-4 hr saved** (already exists)
- S1-T6 Brand Guardian rules: **2 hr saved** (port from DOC\_05 instead of inventing)
- S1-T8 System prompt: **1 hr saved** (DOC\_05 Section 9 IS the system prompt)
- S2-T3 Sentiment journey stages: **1 hr saved** (DOC\_04 has the stages)

**~6 hours of careful work that's already been done by Mayra in February.** The system was designed properly; just never built.

## Bigger insight

DOC\_03 explicitly named "Claude Brand Guardian bot" as the central intelligence in February. The v6 brief tonight independently arrived at the same architecture. **This isn't us inventing — it's us finally implementing what was always supposed to ship.**

Mayra (or whoever wrote these docs) deserves credit. The plan has been right for 2 months. The execution has been the bottleneck.

## What I recommend for Monday morning

1. **Brad approves the AI identity policy** (DOC\_05 says human-sounding, v6 says disclosure — pick one)
2. **Liam ingests 9 brand\_knowledge docs into brain\_knowledge** with chunking + embeddings (replaces D0-T1)
3. **Mayra reviews if anything in DOC\_05 needs updating** based on the recent voice-vs-text learnings (Sunday's comparison)

4. **Pull DOC\_07 (Compliance) into review** — Brand Guardian also needs to enforce QBCC/ACCC/legal language rules, those live in DOC\_07 not DOC\_05
5. **v6 brief gets a v7 patch** with the Day 0 task replaced + Brand Guardian L1 rules upgraded from doctrine

This is what proper Tuesday morning preparation should look like.

---

*Filed by Jarvis Node 3 ~00:25 AEST. Brand Doctrine is real, professional, and has been waiting for someone to wire it in.*