

Grace Dossier (Structural Sample)

active

Rendered 2026-05-14 06:38 UTC

STRUCTURAL SAMPLE, NOT THE FULL DOSSIER

This document is a redacted sample of a Via Negativa Health launch dossier, published as an example of the kind of governance and evidence dossier Via Negativa produces for its consultancy clients. It is not a complete record and must not be relied on as one.

What you can see. Roughly the first half is shown in full: the framing, document control, methodology, company and regulatory context, positioning, and language-standards sections. These show the structure, depth, and house standard of the work.

What is withheld. The critical and proprietary sections are blanked with a clear "[REDACTED, AVAILABLE UNDER NDA]" panel: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, and tier and pricing figures. Tier and pricing material is also subject to a pending content correction and is withheld for that reason as well.

The full dossier is available under a non-disclosure agreement. Via Negativa Health produces comparable dossiers, to this structure and depth, as a standard consultancy deliverable for clients bringing a product, tool, or platform to market.

Via Negativa Health is a trading name of GNL Ltd. Enquiries: john@theglucoseneverlies.com

THE GLUCOSE NEVER LIES®

Grace Dossier

Knowledge Base: Evidence, Algorithms, Validation, Safety, Pack Programme, Governance, Commercial Record

Build date 10 May 2026 (internal version v9.0, supersedes v8.6 with the algorithm appendix v2.9 folded in as Section 15; date-named filename per the locked dossier convention 10 May 2026; the standalone "GNL Grace Compliance & Governance Dossier" naming is retired in favour of the consolidated "Grace Dossier")

CONFIDENTIAL, INTERNAL AND REGULATORY USE ONLY, NOT FOR EXTERNAL SHARING

Tagline: Keep GNL Grace Free for T1D

[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]





The Glucose Never Lies Ltd Company No. 16733595 36 Lea Green Lane, Wythall, Birmingham, B47 6HE, United Kingdom <https://theglucoseneverlies.com> UK Trademark: UK00004267795 | Classes 41, 44 Second Trademark (accepted for publication): UK00004360249 | Classes 41, 42 GNL Grace UK Trademark: UK00004370619 | Classes 9, 41, 42 (filed 9 Apr 2026)

DOCUMENT CONTROL

Version	Date	Author	Summary
<i>[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]</i>			

| 8.2 | 4 May 2026 (overnight) | John Pemberton / Claude Code | Four remaining matrix packs, evidence sweep + LINK FIX + library shopping list locked. Four parallel Opus 4.7 survey agents ran on Packs 8 (Carb counting), 9 (Paediatrics), 11 (Sleep), 12 (Menstrual cycle) and identified 151 anchor papers across the four scopes. Disk audit cross-checked the 151 against `_ARCHIVE/Evidence-base-backup-2026-03-28/`, `evidence/` in-tree, `gnl-grace/wiki/sources/`, both .bib libraries, and the working

Mendeley export (`gnl-grace-citations.bib` , treated as canonical, no separate Mendeley SQLite or BibTeX export found). 14 papers + 6 menstrual / pregnancy bonus PDFs are already on disk; 137 anchors remain for the library run on 5 May (Pack 8: 37, Pack 9: 21, Pack 11: 40, Pack 12: 39, T1: 69, T2: 53, T3: 15, OA flag Y on ~16). **Wiki ingest gap analysis** confirmed no "free Grace ingest" candidates exist (the .bib is auto-generated from `wiki/sources/` so by construction every cited paper has a summary); pack density measured: 8 = 45 LINKED, 9 = 97 LINKED (very rich), 11 = 6 LINKED (thin), 12 = 5 LINKED (thin). **Auto-download attempt failed** at the publisher gate (Sage / Wiley / Elsevier / Liebert all blocked direct UA, SSL handshake breakage on several CDNs); the parent ingest agent stalled at the 600s no-progress watchdog and was retired without commits. **Recovery pass executed without network**: 136 of 153 LINK FIX rows actioned across `wiki/concepts/carbohydrate-counting.md` , `wiki/evidence-grades/paediatrics-evidence.md` , `wiki/concepts/sleep-and-t1d.md` , `wiki/concepts/menstrual-cycle-and-t1d.md` (Pack 8 + 32, Pack 9 + 94, Pack 11 + 5, Pack 12 + 5; 17 deferred rows correctly belong to Pack 3 Exercise CLOSED 15 Apr 2026, not Pack 8 carb counting); three flagged citation typos (Brown 2015 *J Diabetes Sci Technol* not *DTT*, Levy 2022 DOI 10.1089/dia.2021.0431, Tatulashvili 2022 *JCEM* dgac443 not *Diabetes Metab*) recorded in the Pack 12 Evidence Backbone footer for the morning ingest pass; ISPAD adolescence chapter mismatch (Ch21 not Ch8) confirmed already correct in `wiki/sources/` and only present in the working shopping draft `to-do/GRACE_PAPER_SHOPPING_RAW.md` , left as-is per recovery scope. **PACK PROGRAMME matrix** rows for Packs 8, 9, 11, 12 bumped to last-touched 2026-05-04 with per-pack LINK FIX counts and shopping-list pointers; **no pack closed this version**, the library run is the gate. Banned-string scanner clean across all five touched files (zero em-dashes, zero en-dashes, no DAPHNE / endorsed-by / cohort-mismatch / personalised-correction-dose / Take-X-units strings). `gen_bibtex.py` re-run for hygiene; only `.bib` date stamps changed (2026-05-03 to 2026-05-04). Five conventional-commit commits pushed to `john-GNL/GNL_Grace main (90bb4ef. .7c8124b` , one per pack plus matrix). **Forge auto-pull verified at 90 seconds**: `https://api.theglucoseneverlies.com/openapi.yaml` returned **HTTP 200 OK**. **Rendered artefacts** for John's morning library trip: `rendered/GRACE_PAPER_SHOPPING_LIST.html` (interactive 137-paper list with localStorage tick-as-you-go, pack / tier / OA filters, print-friendly) and `rendered/STATUS_2026-05-04_NIGHT.html` (overnight status report, timeline, pack scoreboard, top 10 priorities). Final manifest at `to-do/INGEST_MANIFEST_2026-05-04.md` (untracked, working note). **No change** to: cohort canon, AID Optimiser positioning, age-banding canon, correction-dose framing, two-bucket audience model, registration go-live status (3 May 2026). **Pack count**: 7 of 12 closed (unchanged); Packs 8, 9, 11, 12 remain OPEN with library run as the gating dependency. Pack count after library run + ingest: target 11 of 12 closed by 5 May 2026 evening (Pack 12 menstrual menopause evidence layer remaining the thinnest at 5 LINKED + 39 to fetch). || **8.1 | 3 May 2026 (late evening) | John Pemberton / Claude Code | Pack 10 GLP-1/GIP CLOSED**. 4 deferred stub source summaries built and PubMed-verified (Saeed 2024/2025 Clin Diabetes titration PMID 39829697 DOI 10.2337/cd24-0067; ADA 2025 SoC §9 PMID 39651989 DOI 10.2337/dc25-S009; Jensen 2024 JAMA Netw Open BMD secondary analysis PMID 38916894 DOI 10.1001/jamanetworkopen.2024.16775; Holt 2021 GLP-1 extension on existing source page). Pack body bare-prose Mendeley-key references swapped for `[[wikilinks]]` at 5 locations. `gen_bibtex.py` regenerated to 556 entries; `push_mendeley.py` synced full library (782 DOI-tagged docs in Mendeley, all 4 new DOIs verified present). Citation-integrity rule satisfied: every wiki source page traces to a verified BibTeX entry and a Mendeley document, no hallucinations. **Three-agent Tier A sweep run on Opus 4.7** (factual + voice + consistency, three agents per pack) on Pack 7 (Alcohol) and Pack 10 (GLP-1) per the Grace review HARD LOCK; both QUALIFIED PASS; cleanups applied: Pasqua 2025 PMID reconciled from 39833404 (a brain-computer-interface paper) to verified 39794615; Holt 2021 standardised on Diabetes Care primary; ADJUNCT TWO trial-evidence row dose-label parenthetical reordered to `1.8 / 1.2 / 0.6 mg arms` matching the high-dose-first ordering of HbA1c values; 6 prose long-separator hyphens

replaced with commas / colons / semicolons on the GLP-1 concept page; ISPAD Ch12 source frontmatter regrade A to C per methodology v1.1 §4 (consensus reports = C regardless of journal); Pack 7 Q4 unverified DKA figure dropped, Q12 men-only softened, Decision 6 PENDING flipped to PASS; alcohol concept page got population-average framing lifted above the insulin-adjustments bullets and a closing care-team-referral line on the Targets-and-rescue section; 4 source `related:` blocks backfilled. **Live Grace API verified** end-to-end: `POST /api/grace/query` returns 200 with valid integer `conversation_id` (the earlier session-test "500" was a 422 validation error masked by Laravel's production error handler when the field was sent as a string); Grace cited Saeed 2024/2025 with the right PMID, DOI, evidence grade, and the pub-year-2025-with-cd24-DOI-prefix nuance. **PACK PROGRAMME row 10 flipped to**     (CLOSED 3 May 2026). Pack 10 evidence-pack frontmatter: `source_count: 13 → 17`, `status: draft → closed`, `version: 1.0 → 1.1`. **Site state already complete since 20 Apr:** FAQ page-7186 rebuilt + live; 4-part Adjunctive Therapies guide live (pages 13759-13762); Shah review email sent 20 Apr 2026 22:50 (Gmail 19dacdf41cfff4f). **Pack count:** 7 of 12 closed (was 6); Pack 7 v1.1 evidence layer also closed by parallel agent same evening, leaving Calculator rebuild + 35-source expansion as the residual Pack 7 surface work. `CHANGE_LEDGER +2` rows. Commits: `gnl-grace 4643e64` (stub closures), `0a74d74` (3-agent sweep cleanups), `e4a358f` (`CHANGE_LEDGER`); `gnl-site 11721433` (`TECHNICAL_JOBS` rev 7.16 + `OUTSTANDING-ACTIONS`). | [*Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.*]

| **7.0 | 1 May 2026 (evening) | John Pemberton / Claude Code** | Age-banding canon shipped: Grace voice rule + wiki audience tagging convention + banned-string regression scanner. (Original v7.0 row preserved verbatim in archive.) | | **6.7 | 30 Apr 2026 | John Pemberton / Claude Code | MiniMed 780G positioning policy LOCKED following Cohen letter (28 April 2026).** Trigger: letter from Ohad Cohen MD (MiniMed Global Sr. Medical Affairs Director) replying to John's 22 April 2026 follow-up after the scientific team call. Cohen's position: (a) "aggressive" is not their term of art; (b) the GNL 5-level ladder cannot be endorsed without simulator validation and "should be properly stated"; (c) MiniMed-recommended starting configuration is AIT 2 hours with target 100 mg/dL adults / 110 mg/dL ages 7 to 14 / 120 mg/dL ages 2 to 6, down-titrated to 100 mg/dL only when there are no hypoglycaemia concerns; (d) IOB on the pump screen is not a routine clinical-decision input; (e) any podcast appearance by a MiniMed employee requires Medical Affairs content review for off-label and legal reasons; (f) MiniMed has offered to assess the GNL ladder through their simulator. **Actions taken:** (1) New canonical policy at `wiki/policies/medtronic-780g-positioning.md` (locked); single source for every 780G surface across wiki, website, PDFs, explorer outputs, and podcast scripts. Captures manufacturer-recommended starting configuration verbatim, the non-endorsement block, banned terminology table, podcast review scope, and audit triggers. (2) `wiki/entities/medtronic-780g.md` rewritten: System overview table reassembled (paediatric targets and AIT options were stranded outside the table since the 22 April Via Negativa edit; now reintegrated); manufacturer-recommended starting configuration table added; non-endorsement block placed above the renamed "GNL teaching framework (3-tier illustration)"; "aggressive" replaced with "responsive" throughout; AIT framing rewritten to MiniMed Medical Affairs language. (3) Cascading wiki edits: `concepts/aid-systems.md`, `concepts/aid-systems-head-to-head.md`, `concepts/aid-selection.md`, `concepts/aid-evidence-pack.md` (FAQ 3 + Decision 3 resolved), `concepts/aid-nutrition-bolus.md`, `concepts/gnl-guide-design-system.md`, `wiki/log.md`, `wiki/index.md`. (4) Companion non-endorsement disclaimers added to `entities/camaps-fx.md`, `entities/omnipod-5.md`, and `entities/tandem-control-iq.md` (the principle generalises; CIQ table notes Laurel Messer's 17 April review but no formal Tandem endorsement). (5) `CHANGE_LEDGER.md` row added with full before/after delta. **Site work drafted, NOT deployed:** `pages/page-2272-rebuilt.html` (selector criterion 3 wording, MiniMed feature box, head-to-head section); `pages/resource-aid-780g-survive-`

`thrive.html` (Side 1 framework table to GNL teaching framework with non-endorsement block, "Most aggressive for highs" tick item rewritten, footer non-endorsement); the AID Optimiser host page (WP page 12111) needs a non-endorsement panel and a paediatric target guard above the explorer. **Server-side AID Optimiser** (Laravel API adapter `adapters/m780g.php` in Phillip's repo) needs a parallel update: paediatric target floor enforced on input, "aggressive" replaced with "responsive" in level labels, non-endorsement note added to the result.note string for any 780G output. Phillip handover drafted at `PHIL_NEXT_STEPS.md`. **External:** counter-proposal to Cohen drafted with podcast review scope (MedAffairs reviews MiniMed segment for factual/off-label only; GNL retains editorial control over the rest, 10 business-day turnaround, non-response = approval), explicit acceptance of simulator validation offer, acknowledgement of the in-label points where MiniMed is correct (AIT 2h is the recommended start; paediatric targets follow age-banded labels). Wiki and dossier are legally clear as of this commit; site deploy, external reply, and Phillip ship pending John sign-off. || 1.0 - 5.8 | 6 Apr - 15 Apr 2026 | John Pemberton / Claude Code | Initial build through real-world evidence integration, GLP-1 ingest, evidence expansion to 275 sources, learn-more footer, qualitative selectors, AID + Exercise pack closure. Full history preserved in **Appendix A**. || **6.0 | 16 Apr 2026 | John Pemberton / Claude Code | Framework cementing release**. Consolidated v5.x patches into a clean baseline before the remaining nine evidence packs run. Two structural additions: (1) **Section 6 Pack Programme Integration** ; formal link to `PACK_PROGRAMME.md v2.0` , `EVIDENCE_PACK_METHODODOLOGY.md v1.2` , `CHANGE_LEDGER.md v1.0` . Three packs CLOSED (CGM, AID, Exercise); nine queued (IOB next). (2) **Section 10 Algorithm Validation** ; IOB sigmoidal rising-branch clamp (15 Apr 2026 incident) recorded as §10.4; Explorer Dossier reference updated v8.0 to v8.1. Wiki state refreshed: 51 concept pages, 413 source summaries, 11 entity pages, 12 evidence-grade maps. Doc-control history compressed to a single row + appendix; future updates append a single row per pack closure. **Manuscript-governance section removed** ; the *Diabetes Care* commentary is being authored by a separate external agency and is not a Grace governance dependency. || **6.3 | 17 Apr 2026 (EOD) | John Pemberton / Claude Code | End-of-day commercial consolidation: Educational Grant product line LIVE and PITCHED; Grace Max positioned as discovery tier for new prospects; balanced Sonnet/Opus model strategy locked; Opus preview RETIRED**. (1) **Educational Grant product line launched** (`consultancy/educational-grants/` , separate from Via Negativa consultancy tiers): Foundation **£15,000** / Partner **£30,000** (anchor) / Premium **£45,000** annual grant partnership. Two editions live in the repo: generic pitch + Insulet edition with three Omnipod-5-specific Grace scenarios (tubeless selector, SmartAdjust explainer, CGM non-comparability deep-dive citing Pemberton JS et al. 2026 *DOM* 28(4):2551-2565 DOI 10.1111/dom.70460 + Khan 2026 *Diabetic Medicine* + Forlenza 2024 *DTT* n=69,902 + Pemberton & Uday 2026 system-agnostic framework + Syntactiq cross-check). Pitches sent live 17 Apr 17:00 to Tandem (Adam Sayer + Laurel Messer, Gmail `19d9c3a83e06377d`) and Insulet (Pete Jennings + Zoe Cholewa, Gmail `19d9c4e5766800ee`). (2) **Grace Max re-positioned as the discovery tier for new commercial prospects**. Grace Max is not only the paid subscription product for contract partners; it is also **the live-demo experience shown to prospective manufacturer / pharma / national-scheme partners before any contract is signed**. When a prospect arrives, Max answers research-depth questions about their own products with full citations in their own voice, and this is what converts the discovery call into a signed grant. Tandem and Insulet commercial pitches both build on this framing. (3) **Balanced model strategy LOCKED, no Haiku anywhere**. Basic (public free) and Pro (subscriber-gated free) both run **Sonnet 4.6** with the same evidence base and safety rails; they differ by output token ceiling (350 vs 1,000) and quota (10/hr per IP vs 30/month per user), not brain. Max (contract-gated) runs **Opus 4.7** with 8,000-token ceiling, 10 modes, extended thinking, file upload. Anthropic workspace rate limits locked on the console: Sonnet pool 60/40 Basic/Pro (270K/60K + 180K/30K input/output TPM), Opus pool 78/22 Max/partners. Monthly caps \$250/\$150/\$300/per-contract. (4) **Grace Opus preview RETIRED**. WP page 13515 force-deleted via `wp-cli` 17 Apr; `/grace-opus/` now

301s to `/gnl-grace/`; local files (page + 2 stale email templates) archived to `archive/old-pages/`. Superseded by the Grace Pro and Grace Max password-protected previews (`/grace-pro-preview.html` pw `JohnandPhillPro`; `/grace-max-preview.html` pw `JohnandPhillMax`). WP-deploy shells for `/grace-pro/` and `/grace-max/` sitting in `pages/page-grace-pro-rebuilt.html` (Sonnet, teal-on-navy) and `pages/page-grace-max-rebuilt.html` (Opus, gold-on-navy, 10 modes), awaiting Phillip's JWT security ship. Subscriber-announce email drafted at `pages/email-grace-pro-announce.html`. (5) **Phil JWT ship plan locked** (`to-do/PHIL_JWT_SHIP_CHECKLIST.md`): 3 pieces (WP mu-plugin to Laravel middleware to Max-tier routes) + DB migration (5 tables) + top-5 optimisations. 5-7 days core, 3-4 days optimisations. Emailed Phil 17 Apr 18:30 (Gmail `19d9c9060f0efba3`) with full brief and AID Phase 2 handover attached. (6) **Control-IQ framework v2.0 shipped live 17 Apr** (Messer sign-off, Phase 1): level labels Protective/Gentle/Balanced (default)/Strong/Strongest; ICR replaces Bolus % as third lever; Sleep Activity reframed (24/7 if bolusing regularly, night-only if missed); progression caption warning L5 is more responsive not better. (7) **Wiki unchanged this version**. Grace API architecture unchanged (widget live on `/api/grace/query`, Sonnet 4.6, agentic RAG, Forge auto-pull). Lactate v1.0 evidence still current. 515-file wiki stable. (8) **Repo hygiene**: `to-do/` swept (7 stale working docs archived); dossier filename in lockstep with header version. | *[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]*

| **6.1 | 16 Apr 2026 (PM, late) | John Pemberton / Claude Code | Lactate concept v1.0 closure**. Three more anchor papers ingested in one session (Monnier 2017 [Lac-8], Zhang 2019 [Lac-11], San-Millán & Brooks 2018 [Lac-12]); concept page promoted v0.2 to v0.3 to v1.0. **Battelino 2019 retired** from the lactate concept as superseded for current TIR/TBR/TAR/CV target citations; **Danne 2017 dropped** as non-load-bearing. [Lac-4] reassigned to ADA Standards of Care 2026 Section 6 (Glycaemic Goals and Hypoglycaemia); [Lac-10] reassigned to ISPAD 2024 Glycaemic Targets (PMID 39701064). Citation library v1.5 to v1.7. New **answering-rules section** added to the concept page covering free Q&A vs deep-work deflection (Spencer-tier requests, full protocol design, manuscripts, white papers, partner integrations all redirect to `john@theglucoseneverlies.com` per `wiki/policies/deep-work-deflection.md`). All seven primary lactate evidence chains closed; overall grade A. Built on `lactate-v1.0` branches in both repos, reviewed in lockstep, merged to main (Grace `c301976`, `gnl-site 0c049d84`). Forge auto-pull live ~60s after Grace merge. |

| **6.4 | 18 Apr 2026 | John Pemberton / Claude Code | Pack 4 IOB CLOSED. 15% exercise IOB uplift specified**. (1) Evidence pack v1.0 built: 7 new source summaries (Leohr 2021/2024/2026, Pinsker 2024, Rakab 2025, Moser/Zaharieva 2025, Bassi 2025). Tier A PASS. Key finding: ultra-rapid insulin in AID produces less than 1pp TIR improvement (Rakab SR/MA, 12 RCTs, $p=0.12$); benefit is in glycaemic variability reduction (CV -0.78% , $p=0.02$) and exercise safety. (2) `iob-trade-off.md` concept page augmented (not rewritten) with ultra-rapid findings, netIOB framework, EASD/ISPAD position statement. Evidence grade map upgraded MEDIUM to HIGH. (3) Site upgrades: page-12530 (IOB Trade-Off) new ultra-rapid section; page-2272 (AID hub) new ultra-rapid callout; page-11681 (Activity Explorer) reframed as walking, exercise type selector removed. All three deployed live. (4) **15% IOB exercise uplift specified** for all three exercise explorer endpoints: when exercise type = aerobic AND duration ≥ 30 min, IOB remaining fraction $\times 1.15$ (clamped at 1.0). Evidence: Leohr 2026 Diabetes Care (exercise accelerates SC insulin absorption). Glucose drop estimates unchanged (already exercise-calibrated). Spec at `docs/ALGORITHM_SPEC_IOB_EXERCISE_UPLIFT.md`. Awaiting Phillip server-side implementation. (5) Emails drafted to Mike Riddell (algorithm notification) and Simon Helleputte + Joe Henske cc Riddell (IOB podcast series update). (6) Phil priority list consolidated (9 items, dependency-ordered). | | **6.5 | 20 Apr 2026 | John Pemberton / Claude Code | Pack 5 HbA1c/HGI CLOSED. Lachin 2007 Table 2 fallacy reframe live. Three audit gates passed**. (1) **Evidence pack v1.3** built (`concepts/hba1c-hgi-`

glycaemic-measures-evidence-pack.md): 22 sources, evidence_grade C overall (matrix combination rule Grade D, inputs Grade A/B). New source file: **Bebu 2020** (DCCT/EDIC CVD risk factors, HR 1.38 per 1% HbA1c, 16.6% 30-year cumulative first CV event incidence). Source file rewrites: **Lachin 2007** leads with the mean-glucose-adjusted finding, names the **Table 2 fallacy / mediator-as-confounder** error in the title-line claim, lists 8 follow-on confirming studies; **Pemberton 2025** frontmatter completed (n=168, +4 mmol/mol Black-vs-White-and-South-Asian adjusted, p<0.05 hypoglycaemia signal, BMJ Open DRC, authors Pemberton/Uday/Krone/Fang/Chalew); **McCarter 2004** frontmatter completed (within-person CV 3.6%, journal Diabetes Care, full author list McCarter/Hempe/Chalew); **Hempe and Hsia 2022** author spelling corrected from "Hisa" via Mendeley library lookup, slug preserved to avoid breaking related: links.

(2) **Three governance gates passed.** Tier A appraisal returned REVISE then PASS after four required revisions (IFCC head-to-head reframing because Pleus 2025 published the protocol but no manufacturer formally certified yet; Pack 1 CHANGE_LEDGER row authorising named-device placement on the basis of post-Pack-1 head-to-head data Sanfilippo/Eichenlaub/Freckmann 2025; Pemberton 2025 source file completion; TBR <4% across all matrix cells). 4-Grace pre-deploy audit (congruency, flow, accuracy, honesty) returned PASS-WITH-MINOR / TIGHTEN-FIRST / MINOR-CORRECTIONS / TIGHTEN-LANGUAGE-FIRST; all 29 audit items applied including author-order corrections, Lachin call-out moved above worked example, Hub TL;DR cut to 6 bullets, "fundamentally"/"game-changers"/"powerful" replaced, GMI defined in prose, Grade D matrix cue added, modelled-deterministic prefix on compounding bullets, ensemble-to-individual spine. Final 2-Grace audit (cold-eyes reader walk-through plus deploy-readiness + edge cases) returned SHIP / GREEN-LIGHT. (3) **11 live pages deployed and verified 200 OK.** Hub 13545 (What Gets Measured Gets Managed); Parts 13585 (CGM Zones, with SmartGuide carve-out callout) / 13586 (HbA1c biology) / 13587 (Know Your Glycator Status, with Lachin call-out moved above worked example, weight-of-evidence table) / 13588 (Personalised Target matrix, with Jump-to-matrix anchor in hero, structured Methods + Findings box reframed to drop OP5 by name and lead with ≥ 2 pp TIR over 30 to 40 years compounding via Bebu HR, "Unpublished, full dataset available on request" with mailto) / 13598 (Taleb companion, SITG framing softened to mostly ergodicity, antifragility plain-English gloss); Foundations Part 9 (187, with bridging sentence on 48 ISPAD vs 53 matrix anchor); 10-Tips-TIR (3769); AID hub (2272); IOB Trade-Off (12530, +5 mmol/mol Black ethnicity claim corrected to +4 with Pemberton 2025 citation). Triple cache flush completed (WP Fastest plus object cache plus CloudFront invalidations). (4) **Locked decisions** (Section 10 of pack v1.3): mHGI thresholds plus or minus 3 mmol/mol; Zone B = Simplera + Guardian 4 only (Roche SmartGuide off-matrix pending head-to-head dynamic accuracy data); TBR <4% on every cell (international consensus, ATTD 2019 / ADA 2024); self-classification gate $n \geq 3$ paired windows over > 9 months; mHGI Explorer deferred until peer-reviewed publication and cross-sensor validation; pre-ADA publication authorised; Lachin 2007 Table 2 fallacy called out plainly across pack and Part 3; Author Note declaring John's COI on Pemberton 2025 and mHGI manuscript live on Hub and Part 4. (5) **Site audit:** comparisons/hba1c-hgi-site-audit-2026-04-20.md . (6) **CHANGE_LEDGER.md:** 4 new rows (Pack 1 device-naming authorisation, Pack 5 v1.0 to v1.3, Roche SmartGuide off-matrix, Lachin 2007 source file reframe). (7) **PACK_PROGRAMME.md row 5** IN PROGRESS to CLOSED 20 Apr 2026 (P/SU/L all checked). 5 of 12 packs now closed. (8) **CLAUDE.md (parent + gnl-grace):** Grace review rule HARD LOCK locked 20 April 2026; supersedes "if it reads, ask Grace" rule of thumb for review-class tasks. Memory feedback_grace_reviews_internal.md indexed in MEMORY.md. (9) **Citation library:** v1.10 to v1.11 with Pack 5 closure note. (10) **ADA-ready** before 5 to 8 June 2026 deadline (Chalew, Pemberton et al. ADA poster on mHGI). |

| **6.6 | 20 Apr 2026 (evening) | John Pemberton / Claude Code | Pack 6 Hypoglycaemia CLOSED.** (1) **17 foundational hypo papers ingested** from the 20 Apr library drop, 18 PDFs delivered (17 usable, 1 wrong substitution for Pedersen-Bjergaard 2003). Source files added for Gold 1994, Clarke 1995, Cryer 2013 NEJM, Cranston 1994, Fanelli 1993, Skriverhaug 2006, Weston 1999, Pedersen-Bjergaard 2004,

Geddes 2008, Brodows 1984, Slama 1990, Husband 2010, Rickels 2016, Valentine 2019, Bergenstal ASPIRE 2013, Buckingham PLGM 2015, Gonder-Frederick 2011 HFS-II. (2) **Evidence Pack 6 v1.1** built at `concepts/hypoglycaemia-evidence-pack.md`, 21-row trial matrix (DCCT, Gold, Clarke, Cranston, Fanelli, Geddes, Skrivarhaug, Pedersen-Bjergaard 2004, McTavish, Brodows, Slama, Husband, Rickels, Valentine, Bergenstal, Buckingham, Gonder-Frederick, Tsalikian, Deshmukh, Karges, Ware), 12 FAQs with care-team referral inside each, Via Negativa audit quoting `concepts/hypoglycaemia.md` lines 78 and 88, Tier A two-pass PASS. First-pass FAIL was on Husband 2010 misrepresentation (pack said "no difference" but trial showed glucose and sucrose equivalent at 88 percent success, fructose failed 33 percent); rebuild corrected Table B, Table C, Q5, Q7, Section 5 rows, Section 6 item 2, Section 9 Decision 1. (3) **Evidence grade A- to A** on `wiki/evidence-grades/hypo-hyper-evidence.md`, 28 dedicated hypo source files (was 11), 41 total hypo-relevant sources. (4) **concepts/hypoglycaemia.md refined** (John Decision 1): fructose-to-liver mechanism retained (Husband confirms it, Fruit to Go failed 33 percent), sucrose-takes-an-hour line softened to a dose-plus-stacking framing at 10 to 15 g doses. Sources list updated 19 to 36, grade B to A. (5) **CHANGE LEDGER +7 rows**. PACK_PROGRAMME row 6 marked P and L ticked, SU pending. **6 of 12 packs now closed**. (6) **Library re-request v2.0** drafted at `rendered/HYPO_LIBRARY_REREQUEST_v2.0.md` for 5 residual papers (Pedersen-Bjergaard 2003 wrong paper sent, Tattersall 1991 not sent, Haymond 2001 not sent, Forlenza PROLOG 2018 not sent, Zaharieva 75 percent bolus publication TBC). (7) **gnl-grace-citations.bib**: 3,549 to 3,774 lines (+225, 17 entries). Valentine 2019 DOI corrected to `10.1089/dia.2019.0148` and PMID to 31335191 (was lock-in guess). Buckingham 2015 source file pages and DOI corrected to 1939-1946 and `10.2337/dc15-1012`. Gold and Clarke source files regraded A to B (observational cohort, not RCT). (8) **Root GNL_CITATION_LIBRARY.md**: v1.12 to v1.13, 17 entries added to HYPOGLYCAEMIA section (141a to 141q). Grace curated subset 412 to 429. (9) **Four-part /hypoglycaemia/ + hyperglycaemia + driving guide brief** queued for next session at `to-do/HYPO_HYPER_DRIVING_GUIDE_BRIEF_v1.0.md`. Group 2 DVLA rules, finger-pricking requirements, hyperglycaemia awareness woven through Part 4. (10) **Pack 6 commit d2056e8** on GNL_Grace main; Forge auto-pull live ~60s after push. **Pack 10 (GLP-1/GIP) is next**, prior session already in flight with uncommitted draft (`wiki/concepts/glp1-gip-adjunct-therapy-evidence-pack.md` plus 4 source files); this session did not touch it. |

Next scheduled review: on close of Pack 10 (GLP-1/GIP), or on npj Digital Medicine reviewer return, whichever comes first.

TWO-BEACON ROUTING MODEL (locked 20 April 2026)

Status as of 20 April 2026 evening

- **200-case web-app parity + correctness test suite ran clean**. Protocol: `./docs/GNL_WEBAPP_TEST_PROTOCOL_v1.0.md`. The Flutter app is now accepted as the source of truth for the six explorers.
- **/gnl-grace/ rebuild in progress**. John is building the hub (Grace widget inline + explorers grid + "How Grace works" explainer). As of 26 April 2026, Grace is a single product with no visible tiers.
- **Decommissioning of redundant pages in hand**. `/gnl-explorers/`, `/how-gnl-works/`, and any other explorer surfaces that duplicate what `/gnl-grace/` + the web app now deliver are being retired or 301-redirectioned. Grace should already be linking to `/gnl-grace/` per the directory update in `wiki/gnl-website-directory.md`; the old URLs will continue to resolve via 301 so any cached link still lands in the right place.

The routing decision

[/gnl-grace/](#) (page 12998) and the homepage (/ , page 7) are the two absolute beacons on the GNL site. Everything that is not an editorial guide funnels through one of them and then out to the Flutter web app ([app.theglucoseneverlies.com](#)) for the authenticated experience.

Surface	Status	What it does
Homepage (/ , page 7)	Second beacon	Top-of-funnel landing. Hero plus three beacon blocks (Ask Grace / Use explorers / Read guides), all routing to /gnl-grace/ or the guide hubs.
/gnl-grace/ (page 12998)	Single hub	Grace widget inline; 7 explorers (each as a short explainer card); "How Grace works" explainer section; every CTA routes to the web app for authenticated use. Single product, no tier cards.
Flutter web app (app.theglucoseneverlies.com)	Authenticated surface	Grace chat (7 guided modes for all users; 9 specialist modes silently unlocked for Max-authenticated logins), all 7 explorers with saved profile, saved history, push, TTS, attachments.
Guides (HbA1c/TIR, Foundations, CGM, AID, Exercise, Alcohol)	Library	Stay on the site. SEO anchors. Every guide footer links to /gnl-grace/ as the next step.
Individual explorer landing pages	Keep live, plus "Open in app" nudge	Still functional for direct Google / newsletter traffic. Gain a top banner pushing authenticated users into the app.
/gnl-explorers/ (page 12166)	301 → /gnl-grace/#explorers	Folds into the hub.
/how-gnl-works/ (page 13099)	301 → /gnl-grace/#how-it-works	Folds into the hub.
/grace-pro/ , /grace-max/ (pages 13531, 13532)	Stripe checkout + fallback surface	Not the primary usage surface. Users who cannot or will not install the app can still chat here.
Podcast hub, resource hub, consultancy pages, legal pages	Unchanged	Stay as they are.

When answering user questions about "how to use Grace" or "where are the explorers":

- Point them to [/gnl-grace/](#) as the single hub.
- If they ask to use Grace or an explorer while logged in, direct them to the web app ([app.theglucoseneverlies.com](#)).
- If they ask where a specific guide is, link directly to the guide URL (guides are still their own pages).

- Do not link to `/gnl-explorers/` or `/how-gnl-works/`; those URLs redirect.

Funnel shape: The site is a library and a doorway. Library = guides + podcast + evidence. Doorway = `/gnl-grace/` + homepage. Every doorway CTA routes to the web app. The web app is where Grace answers, where explorers run, and where the monthly quota lives.

Full architectural detail: `../docs/GRACE_PRO_MAX_BUSINESS_MODEL.md v1.3 §2`.

CANONICAL SOURCE FOR LIVE TIER / PRICING / BUDGET (refreshed 8 May 2026 evening, supersedes 26 April 2026)

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

CLAUDE DOUBLE-CHECK INSTRUCTIONS

This dossier is version 9.0. When reviewing or amending it, confirm each item below before signing off.

5 May 2026 voice-rule reversal (top-of-document erratum, locked). John lives with type 1 diabetes AND is a parent (children do not have T1D). First-person PWD voice is honest and permitted when drawing on John's own lived experience. The earlier 4 May 2026 "John is parent only, never first-person PWD" framing was an error and is reversed. Historical control rows v8.3 / v8.4 / v8.5 below pre-date this correction and refer to the superseded 4 May framing; the live policy is `wiki/policies/gnl-voice-and-storytelling.md §2` (locked 5 May 2026). The seven-hat structure replaces the previous six (hat 2 reactivated as Person living with T1D).

[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]

Naming check

- [] "Grace" or "GNL Grace" : never "GRACE" in body text
- [] "Phillip Hayes" : never "Phil"
- [] "explorers" : never "calculators"
- [] "evidence pack" : never "guide" when referring to the structured `EVIDENCE_PACK_METHODODOLOGY.md` output

- [] British English throughout (organisation, behaviour, analyse, recognise, programme)

Cross-reference check

- [] Explorer Compliance Dossier referenced as **v10.3** (not v9.0, v9.2, v10.0, v10.1, v10.2, or earlier) [*Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.*]
- [] Explorer audience-gating policy referenced as `wiki/policies/explorer-audience-gating.md` (updated 3 May 2026, two-bucket model + AID open)
- [] AID Optimiser positioning policy referenced as `wiki/policies/aid-optimiser-positioning.md` (1 May 2026 lock + 3 May 2026 §1.1 lay companion + §10 audience scope)
- [] `EVIDENCE_PACK_METHODOLOGY.md` referenced as **v1.2** (not v1.0 or v1.1)
- [] `PACK_PROGRAMME.md` referenced as **v2.0** (not v1.0)
- [] `CHANGE_LEDGER.md` referenced as the living post-close record
- [] `GNL_GRACE_ARCHITECTURE.md` referenced as the canonical wiring doc

Pack programme check

- [] Nine packs marked CLOSED (see Pack programme check, immediately below)
- [] Nine packs marked CLOSED: CGM, AID, Exercise, IOB, HbA1c/HGI, Hypoglycaemia (2 May), Alcohol and T1D (3 May), Carbohydrate counting (5 May), GLP-1RA / GIP (3 May, rolled to PASS). Three packs at PARTIAL CLOSE (4 May): Paediatric T1D, Sleep / circadian, Menstrual cycle (each: live guide + evidence pack queued post-launch)
- [] No Diabetes Care manuscript governance (that work belongs to an external agency, not to Grace)

Content check

- [] No em-dashes (U+2014) or en-dashes (U+2013) anywhere in the body
- [] No filler words ("certainly", "absolutely", "in conclusion")
- [] All glucose values mmol/L first, mg/dL bracketed
- [] Population-average framing throughout (no "personalised plan" language)

EXECUTIVE SUMMARY

[*Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.*]

The current state as of 2 May 2026 (updated 2 May 2026):

- **Nine evidence packs CLOSED** under `EVIDENCE_PACK_METHODOLOGY.md v1.2`: CGM (15 Apr), AID (14 Apr), Exercise (15 Apr), IOB (18 Apr), HbA1c/HGI (20 Apr), Hypoglycaemia (2 May), Alcohol and T1D (3 May), Carbohydrate counting (5 May), GLP-1RA / GIP (3 May, Tier A QUALIFIED PASS rolled forward to PASS). Each Tier A appraised, lint-passed, site-audited.
- **Three packs at PARTIAL CLOSE** (4 May 2026): Pack 9 Paediatric T1D (4-part NG paediatric guide live, Tier A evidence pack queued post-launch), Pack 11 Sleep / circadian (4-part guide live, evidence pack queued), Pack 12 Menstrual cycle (3-part guide live, evidence pack queued). 5 May 2026 library run closed the link-fix backlog across all three.

- **Post-close changes go to CHANGE_LEDGER.md**, not back into the closed packs. Packs are point-in-time snapshots; the ledger is the living record. End-of-programme promo reads the ledger. *[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]*
- **17 Via Negativa reversals** documented across all studies, including the female-pen-plus-exercise high-risk profile, the TIR plateau finding, the sleep-dominates-exercise interaction, and the high-CV paradox.
- **77-case test suite passes with 0 genuine failures**; 18-case adversarial explorer-signposting test passes 18/18; safety rails enforced server-side.
- **IOB sigmoidal rising-branch clamp** deployed to Laravel API on 15 Apr 2026 (Phillip commit `e8062ec`, GNL Grace wiki commit `b012f07`) after a user-reported physical-impossibility defect (5.0 U returned for a 4.0 U bolus). Clamp recorded as Grade A evidence (enforces a physical invariant).

Grace is not a medical device. Grace does not diagnose, treat, cure, or prevent any medical condition. Grace does not make individual clinical decisions. Grace is an educational knowledge base that presents compiled evidence at a population-average level, exactly as the GNL Explorer Suite does, but in natural language rather than interactive tool format.

The commercial roadmap includes: (1) the Grace internal API for GNL team use, (2) the live website chatbot widget (deployed site-wide 10 Apr 2026), (3) a voice interface (pilot complete), and (4) a commercial Grace Widget API for Via Negativa Health consultancy clients and device partners. All outputs carry mandatory educational disclaimers and are subject to the same safety rails as the GNL Explorer Suite.

This dossier inherits and extends the GNL Explorer Compliance Dossier (v9.0); the company identity, insurance, trademark, copyright, GDPR, and governance sections of that dossier apply in full to Grace.

SECTION 1 - RELATIONSHIP TO EXPLORER DOSSIER

The GNL Explorer Compliance Dossier (`docs/GNL_COMPLIANCE_DOSSIER_v10.2.md`, refreshed 7 May 2026 with the regulatory horizon sweep; v9.2 archived at `archive/old-docs/GNL_COMPLIANCE_DOSSIER_v9.2-superseded-2026-05-03.md`) is the parent document. The following sections of that dossier apply to Grace without modification:

Explorer Dossier Section	Applies to Grace?	Notes
S1 Company Identity and Legal Framework	YES - in full	Same company, same trademarks, same IP structure
S2 Insurance and Risk Framework	YES - in full	E&O and Multimedia Liability cover educational content outputs
S3 Regulatory Framework (not a medical device)	YES - extended in §3 below	Grace adds natural language outputs to existing interactive tool outputs
S4 Platform Story and Unique Positioning	YES - extended in §2 below	Grace is the next evolution of GNL's educational approach
S5 Educational Positioning and Language Standards	YES - in full	Same forbidden language, same required framing, same population-average rule
S6 Scientific Governance and Advisory Framework	YES, in full	Same team, same advisory model. §6.7 of v10.0 documents Grace formally.
S7 GDPR	YES - extended in §9 below	Grace API adds new data processing considerations
S8-S10 Explorer-specific sections	NO	Explorer-specific; Grace has her own sections

This dossier does not duplicate content from the Explorer Dossier. Read both together.

The Explorer Dossier is updated in lockstep with this one. When Grace ingests evidence that touches an explorer algorithm (e.g. the IOB sigmoidal clamp, 15 Apr 2026), the Explorer Dossier records the algorithm change and Grace's dossier records the wiki/evidence-grade change. Both registries are updated in `docs/GNL_FILE_REGISTRY.md`.

SECTION 2 - WHAT GRACE IS

2.1 Definition

GNL Grace is a persistent, structured knowledge base that compiles clinical evidence for Type 1 diabetes education into graded, cross-referenced wiki pages. Grace is built on the LLM Wiki pattern (Karpathy, April 2026) and adapted for GNL's safety constraints, evidence grading, and commercial requirements.

2.2 Architecture

Grace has five layers:

Layer	What it is	Who owns it
Raw sources	1,245+ PDFs, JSONs, transcripts, guidelines (immutable)	John - evidence base (local, gitignored), scholar pipeline (daily auto-run)
Wiki	541 source summaries + 73 concept pages + 16 entity pages + 26 evidence-grade maps + guidelines + indexes (updated 2 May 2026)	LLM writes, humans review, Tier A appraised, lint-enforced
Guidelines	Age-stratified clinical guideline references (ISPAD 2024, ADA 2026)	Compiled from guideline PDFs, mapped to concepts and explorers
Real-world data	Population-level benchmarks from validated real-world T1D datasets, synthesised via the GNL AI-assisted research engine	See §5.2
Schema + governance	Operating instructions, safety rules, evidence hierarchy, pack methodology	Co-evolved by John and Claude. <code>SCHEMA.md</code> , <code>EVIDENCE_PACK_METHODODOLOGY.md v1.2</code> , <code>PACK_PROGRAMME.md v2.0</code> , <code>CHANGE_LEDGER.md v1.0</code>

2.3 What Grace contains (updated 2 May 2026)

Component	Count	Description
Concept pages	73 (updated 2 May 2026)	Foundations, Guides, Via Negativa, Explorer framework, plus six closed evidence packs (CGM, AID, Exercise, IOB, HbA1c/HGI, Hypoglycaemia), GLP-1/GIP, predictive CGM, lactate, menstrual cycle, sleep, gnl-methodology, foundations-framework.
Entity pages	16 (updated 2 May 2026)	Named devices and systems: Dexcom G7, Libre 3, Simplera, Eversense 365, GlucoMen iCan, CareSens Air, Medtronic 780G, Tandem Control-IQ, Omnipod 5, CamAPS FX, Instinct, MiniMed Go Smart, Diabeloop DBLG2, Roche SmartGuide, and others.
Evidence-grade maps	26 (updated 2 May 2026)	CGM, Exercise, AID, Paediatrics, HbA1c/HGI, Insulin/IOB, Hypo/Hyper, Alcohol, Nutrition, Lactate, Sleep, Menstrual cycle, AI in diabetes education, driving, and others.
Source summaries	490 (updated 2 May 2026)	Quality-filtered from 1,245+ evidence items. Full YAML metadata, BibTeX exported (3,242 entries in <code>gnl-grace-citations.bib</code> + 25 in <code>gnl-ai-education-citations.bib</code>), Mendeley synced.
Real-world source pages	18	Population dataset analyses (full list in §5.2)
Clinical guidelines	3	Guidelines index + ISPAD 2024 (25 chapters) + ADA 2026 (4 key sections)
Source indexes	7	John's publications (42), curated evidence base, citation database, 3 podcast clusters, manuscripts (8)
Safety files	4	Contradictions log, retracted log, confidence map, audit trail
Closed evidence packs	3	CGM (15 Apr), AID (14 Apr), Exercise (15 Apr). See §6.
Total sources mapped	1,245+	Across 27 topic areas, indexed by <code>evidence/evidence_registry.py</code>

2.4 Evidence grades: current status

Grade B and above (real-world population data plus literature review) carries 14 concept pages. The three closed evidence packs (CGM, AID, Exercise) operate at evidence-pack level ; every claim within them is graded individually per `EVIDENCE_PACK_METHODOLOGY.md v1.2 §4` .

Concept page	Evidence grade	Basis
Sleep and T1D	B	9-study lit review + real-world population data (611 users)
<i>[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]</i>		

| Bolus insulin | B | Real-world 364-user dataset + clinical evidence | *[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]*

| IOB trade-off | B | Confirmed by real-world TIR-TBR correlation; sigmoidal clamp recorded 15 Apr 2026 || Activity and movement | B | Real-world 526-user activity profiles + CGM consistency data || Measuring T1D success | B | Real-world benchmarks: cohort, gender, age-stratified, weekday/weekend || Carbohydrate counting | B | 9-study literature review: accuracy, low-carb, bolus timing || Menstrual cycle and T1D | B | Real-world 62-user cycle phase analysis (first female-specific Grade B) || Alcohol and T1D | B | Real-world 881-user activity proxy + 5-study literature review || AID Systems | B | Closed evidence pack v1.0 (14 Apr 2026); real-world 839-user device comparison (409,056 days), 87+ sources || CGM Basics | B | Closed evidence pack v2.0 (15 Apr 2026); real-world 652-701 user wear + consistency analyses || Dynamic Glucose Management | B | Real-world illness/stress, eA1c mapping, technology trends || Lactate and T1D | C | New 16 Apr 2026. 12 sources, mechanism + AID interactions || All others | D | GNL site content - pending intraday/basal dataset access |

2.5 How Grace differs from the Explorers

Characteristic	GNL Explorers	GNL Grace
Input format	Structured form fields	Natural language question
Output format	Structured result	Natural language answer with citations
Interaction model	Single query-response	Conversational (multi-turn)
Evidence presentation	Embedded in algorithm	Explicit citations with evidence grades
Scope	One topic per explorer	Cross-topic synthesis
Underlying model	Deterministic algorithm (Laravel API)	LLM with compiled wiki context (agentic RAG)
Guideline alignment	Implicit	Explicit (cites ISPAD/ADA with grades)
<i>[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]</i>		

| Female-specific content | Not available | Available - menstrual cycle phase analysis (Grade B) | | Age stratification | Not available | Available - five life stage benchmarks (Grade B) | | Pack-level synthesis | Not applicable | Three packs closed; nine queued (PACK_PROGRAMME.md v2.0) |

What they share: population-average framing, educational disclaimer, forbidden language policy, evidence grading, direction to care team.

2.6 The GNL Research Engine

Grace's evidence base is augmented by a structured research process using a highly trained AI-assisted research system with access to a large, rigorously validated real-world T1D dataset. The dataset is accessed through an academic research collaboration and represents:

- An upstream Syntactiq Dynamics FlexCo pool of over 10,000 people living with T1D (children, adults, and older adults across diverse demographic groups)
- Over 10 years of continuous CGM and insulin device data collected in everyday conditions (2013 to 2025) *[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]*
- Multiple device types (MDI, traditional pump, hybrid closed-loop AID), multiple CGM platforms, and documented transitions
- Outcomes that include TIR, TBR, TAR, glucose CV, insulin utilisation, activity, sleep, and self-reported events

The research system: - Analyses population-level data from real-world conditions, not controlled trials - Synthesises findings from peer-reviewed literature alongside real-world observations - Returns structured, statistically validated findings (N, effect size, p-value) that Grace can cite at population-average level - Is used by GNL to validate, challenge, or extend clinical trial evidence with real-world signal

[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]

2.7 Named by

Grace was named by John and Grace Pemberton on 6 April 2026. The name reflects the core property: Grace grows with time, learns with time, gets smarter with time.

SECTION 3 - REGULATORY CLASSIFICATION

3.1 Grace is NOT a medical device

The same six-ground analysis from Explorer Dossier S3.1-3.3 applies to Grace:

Classification axis	Grace position	Evidence
Intended purpose	Education and general information	All outputs are compiled evidence summaries, not clinical decisions
Seriousness of situation	Non-serious / non-critical	User is asking an educational question, not in acute clinical situation
Drive clinical management?	No	Outputs require clinical judgement; explicitly not prescriptive
Diagnose a condition?	No	No diagnostic output of any kind
Connected to a medical device?	No	No CGM, pump, or AID system integration
Individual prediction?	No	Population-average framing - explicitly stated
Replace clinical judgement?	No	All outputs direct user to their diabetes care team

3.2 Grace's intended purpose statement

[Line redacted in the public sample: contains tier, pricing, commercial, or cohort-evidence figures withheld pending the current content fix and the NDA gate.]

3.3 Additional regulatory considerations for LLM-based outputs

Grace uses a large language model (Claude, Anthropic) to synthesise answers from her compiled wiki via an agentic RAG architecture (`search_wiki` MCP-style tool). This introduces considerations beyond the deterministic Explorer algorithms:

Consideration	Mitigation
LLM may hallucinate facts	Grace reads compiled wiki pages, not raw sources. All claims trace to graded evidence. Pack-level claims trace to a source file that exists at time of writing (per <code>EVIDENCE_PACK_METHODODOLOGY.md v1.2 §6</code>).
LLM may generate prescriptive language	Forbidden phrase regex filter applied server-side before any response is returned. <code>pack_lint.py</code> enforces at pack level.
LLM outputs are non-deterministic	Population-average framing means exact wording varies but underlying evidence does not. Citations anchor the answer.
LLM may be prompted to bypass safety	System prompt (SCHEMA.md) is not user-modifiable. API enforces safety layer regardless of user input.
LLM context window limits	Index-first navigation ensures relevant pages are read. Answer quality scales with wiki quality, not prompt engineering.
LLM may misrepresent guideline positions	Guidelines layer provides pre-compiled, human-verified summaries. Alignment status is pre-determined by John.
Real-world data misrepresentation	All real-world benchmarks include sample size, effect size, and p-value in Grace's wiki.
Counterintuitive findings suppressed	Via Negativa policy (§10) requires Grace to report unexpected findings explicitly.
Retrieval drift (model retrieves wrong row from wiki)	Entity-level guard rails added inline in entity pages (e.g. <code>entities/medtronic-780g.md</code> ; "Never say 6.7 is the exercise target; never deny 8.3 exists"). Logged in <code>safety/contradictions.md</code> .

3.4 EU AI Act considerations

Grace may fall within the scope of the EU AI Act (Regulation 2024/1689) as an AI system providing information in the health domain. Current assessment:

- Grace is not a "high-risk" AI system under Annex III - she does not make clinical decisions, does not process biometric data, and is not integrated with medical devices.
- Grace may be classified as "limited risk" under Article 50, requiring transparency obligations.
- Mitigation: All Grace outputs clearly state they are AI-generated educational content. The chatbot and voice interface identify themselves as AI from the first interaction.

SECTION 4 - SAFETY RAILS

4.1 The Section 5.1 Test

Every Grace output must pass the test from Explorer Dossier S5.1:

"Does the output tell the individual what to do, or does it show them how a principle behaves on average?"

Grace always does the latter.

4.2 Forbidden language (inherited from Explorer Dossier S5.3, enforced by `pack_lint.py`)

Prohibited category	Example
Direct instruction with glucose threshold	"If your glucose is X, do Y"
Dose adjustment instruction	"You should adjust your dose when..."
Action trigger language	"Take action if..."
Prescriptive correction framing	"The correct response is..."
Individual outcome certainty	"This will raise your glucose by..."
Device-specific dosing advice	"Set your basal rate to X"
Tools labelled as calculators	"Use the calculator"
Personalised plan framing	"Your plan", "tailored to you", "personalised plan"
AI filler words	"certainly", "absolutely", "great question", "in summary"
US spellings	"color", "behavior", "analyze", "organize", etc.
Em-dash, en-dash	<code>\u2014</code> , <code>\u2013</code> ; use <code>-</code> or restructure
Diminutive of "Phillip"	"Phil" (Phillip Hayes, always two l's)

Enforcement: server-side regex filter on every API response before delivery. `pack_lint.py` enforces at pack-commit time and blocks any commit that fails.

4.3 Required framing

Content type	Required framing
Evidence synthesis	Population-average, source-cited, evidence-graded
Real-world benchmarks	"In a validated real-world dataset of [N] people with T1D..."
Explorer references	"You can explore this with the [Explorer Name]"
Guideline citations	"ISPAD 2024 recommends..." or "ADA 2026 recommends..." with evidence grade
Hypo/hyper guidance	Threshold-referenced to ISPAD/ADA guidelines
Age-specific benchmarks	"In real-world data from [age group] users..."
Female-specific findings	"In a real-world analysis of [N] female users..."
Glucose values	mmol/L first, mg/dL bracketed: <code>5.5 mmol/L (99 mg/dL)</code>
Action-touching FAQ	Care-team referral sentence inside the answer (not as a footer)

4.4 Mandatory disclaimer

Every Grace output (text, voice, or API) must include or reference:

"This is educational information compiled from clinical evidence and population-average data. It is not medical advice, not a medical device, and must not be used as either. Any changes to your diabetes management must be made with your diabetes care team."

4.5 Evidence trust hierarchy

Tier	Source type	Grace treatment
1	Systematic reviews and meta-analyses	Highest weight - cited directly
2	RCTs	High weight - cited with sample size and effect size
3	Real-world population datasets (validated, large N)	High weight - cited with N, effect size, p-value
4	Observational studies	Moderate weight - noted as observational
5	Expert consensus / guidelines / regulatory frameworks / methodology papers	Cited as ISPAD/ADA/IFCC recommendation with grade C
6	GNL site content and podcast synthesis	Background context - not cited as primary evidence

Per `EVIDENCE_PACK_METHODODOLOGY.md v1.2 §4` : regulatory frameworks (FDA iCGM standard, EU CE marking) are Grade C, not Grade A. Methodology papers (Pemberton 2022 zone framework, Eichenlaub 2024 trend arrow proposal) are Grade C, not Grade A. Position statements and editorial opinions are Grade C even if published in a Grade A journal. Real-world dataset analyses (Syno, large registries) are Grade B.

4.6 Learn More footer (hardwired site-wide since 13 Apr 2026)

Every Grace response ends with 1-3 priority-ordered links to relevant pages on theglucoseneverlies.com. Three-layer implementation:

1. `gnl_pages`: frontmatter on all 51 concept pages and 11 entity pages
2. `wiki/gnl-website-directory.md` - master topic-to-URL reference
3. `SCHEMA.md` QUERY operation mandates the footer

The `gnl_pages`: field travels with every `search_wiki` retrieval, so the production widget picks up the URLs automatically without any Forge-side changes. Sources for the footer: `wiki/gnl-website-directory.md` + `wiki/sources/gnl-site-map.md` (155-page inventory).

SECTION 5 - EVIDENCE BASE

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing

figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

SECTION 6 - PACK PROGRAMME INTEGRATION (NEW IN v6.0)

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

SECTION 7 - COMMERCIAL FRAMEWORK

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

SECTION 8 - GDPR AND DATA

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

SECTION 9 - VIA NEGATIVA RESEARCH PHILOSOPHY

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

SECTION 10 - ALGORITHM VALIDATION FRAMEWORK

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

SECTION 11 - GRACE SAFETY VALIDATION

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

SECTION 12 - DOCUMENT REGISTRY

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

SECTION 13 - AGE-AWARE VOICE AND WIKI AUDIENCE TAGGING (LOCKED 1 MAY 2026)

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

SECTION 14 - AUDIENCE BUCKET MODEL AND AID OPTIMISER OPENING (LOCKED 3 MAY 2026)

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

APPENDIX A - DOCUMENT CONTROL HISTORY (v1.0 - v5.8)

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

APPENDIX B - VERSION 6.0 CHANGE SUMMARY

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

SECTION 15 - ALGORITHM REFERENCE (folded in from algorithm appendix v2.9, 10 May 2026)

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

The Glucose Never Lies[®], Proprietary Algorithm Reference

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A0. DOCUMENT SCOPE AND CORRECTION NOTICE

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A1. SHARED COMPUTATIONAL FOUNDATIONS

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A2. ACTIVITY AND EXERCISE EXPLORER

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A3. EXERCISE IOB CALCULATOR

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A4. EXERCISE PLANNING EXPLORER

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A5. AID SYSTEM EXPLORER (AID Algorithm Optimiser)

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A6. HYPO AND HYPERGLYCAEMIA EXPLORER

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A8. HYPER TREATMENT EXPLORER (consolidated, v2.0.0)

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A7. ALCOHOL AND T1D EXPLORER

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A8-IP. SHARED IP ELEMENTS - ORIGINAL ARCHITECTURE

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A9. EVIDENCE GRADE SUMMARY ACROSS ALL EXPLORERS

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A10. FLAGGED ITEMS REQUIRING FUTURE ACTION

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A12. AGE BANDING CANON (locked 1 May 2026)

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.

A11. FULL BIBLIOGRAPHIC REFERENCES

[REDACTED, AVAILABLE UNDER NDA]

This section of the dossier is withheld from the public sample. It contains proprietary or commercially sensitive material: algorithm internals and parameters, evidence-base specifics, validation data, governance internals, data-handling detail, commercial terms, or tier and pricing figures.

The complete section, with full working detail, is supplied to engaged clients under a non-disclosure agreement. Via Negativa Health produces dossiers of this depth and structure as a standard consultancy deliverable.