ShieldSense AI

Project Report

Report Date: 24 April 2026 | Version: 2.0

Live at: https://shield-sense-ai-three.vercel.app

Prepared by:

Chuan Ling Ling · Heng Chin Chuan · Ng Pei Yek · Yu Yanying

Table of Contents

1. What Is ShieldSense AI?

On 1 April 2026, the rules around health insurance in Singapore changed significantly. The Ministry of Health (MOH) stopped allowing new insurance policies to include full-coverage riders — the add-ons that used to mean you paid nothing out of pocket when hospitalised. This left a lot of Singaporeans confused about what it meant for their existing plans and whether they should stay put or switch to something new.

ShieldSense AI was built to cut through that confusion. It is a free, web-based tool that guides Singapore Citizens and Permanent Residents through a simple seven-step process, asking about their age, income, CPF savings, health, and coverage preferences. Based on those answers, it produces a personalised shortlist of recommended Integrated Shield Plans (IPs), ranked by how well they match the user's specific situation. For users who want more depth, the tool also connects to Claude — an AI developed by Anthropic — to generate a plain-English explanation of the recommendations.

The tool was built over six weeks, from 14 March to 24 April 2026, covering all seven government-approved IP insurers and storing over 14,000 rows of premium data.

See Appendix E — Architecture Diagrams — Diagram 1: Screen Hierarchy shows how all screens connect from the landing page to the full AI dashboard.

2. Why We Built It

2.1 The Problem

Choosing the right health insurance in Singapore has always been complicated. Most people need to think about three layers of coverage at the same time: MediShield Life (the national baseline plan everyone is on), an Integrated Shield Plan (an optional upgrade), and a rider (an add-on that reduces what you pay at the hospital). Get it wrong, and a single hospitalisation could cost you tens of thousands of dollars out of pocket. For someone on a private plan with a rider, annual premiums at age 65 can already exceed S$10,000 a year.

The April 2026 changes made this even harder to navigate. Three things happened at once:

- People who bought full-coverage riders before 1 April 2026 — grandfathered policyholders — can keep those riders, but only if they do not switch their IP plan. The moment they switch, that protection is gone forever.
- Anyone buying a new policy after 1 April 2026 can only get a co-payment rider, which means they will always pay at least 5% of their hospital bill.
- All seven insurers simultaneously revised their premium tables, making it nearly impossible to compare what things cost before versus after the change.

No publicly available tool addressed all three of these issues together, while also accounting for a person's individual income, CPF savings, and health situation. That gap is what ShieldSense AI was built to fill.

2.2 What We Set Out to Do

The project had five clear goals:

- Build a simple, guided online tool that collects everything needed to give a personalised recommendation — without overwhelming the user with jargon.
- Create a scoring engine (called C4) that ranks plans based on how affordable they are, how well they match the user's preferred hospital ward, and how sustainably they can be paid through MediSave — all before involving the AI.
- Connect Claude AI to write clear, personalised recommendation summaries, delivered fast enough to work within the technical limits of the hosting platform.
- Load in the latest premium data for all seven insurers, sourced directly from their official documents after the April 2026 changes.
- Clearly warn users who might accidentally give up their grandfathered full-rider protection by switching plans.

See Appendix E — Architecture Diagrams — Diagram 2: User Flow shows the complete end-to-end journey, including all decision points from first visit to AI dashboard.

2.3 Who It Is For

ShieldSense AI is designed for Singapore Citizens and Permanent Residents between the ages of 1 and 100 who currently hold or are considering an Integrated Shield Plan. It covers all seven MOH-approved insurers: AIA, Great Eastern, HSBC Life, Income Insurance, Prudential, Raffles Health Insurance, and Singlife.

Foreign nationals are not currently supported, as they are not eligible for MediSave — which is central to how premiums are calculated and paid.

3. What the Tool Can and Cannot Do

ShieldSense AI is a powerful tool, but it has real limits that users should be aware of. This section is about being honest about what the data covers, what it does not, and what the tool is not designed to replace.

3.1 Some Premium Data Is Incomplete

We sourced all premium figures directly from official insurer brochures published around 1 April 2026. However, not every insurer published everything in a consistent or complete format. The table below shows the known gaps:

Insurer	What Is Missing	Current Status
Raffles Shield	Standard Plan premiums were not published in the brochure	Excluded from results
Prudential	Two rider products (PRUExtra Plus and Plus Lite) are missing some age-band pricing	Being followed up
Great Eastern (pre-Apr 2026)	Premium table only goes up to age 55 in the source document	Partial data only
AIA (pre-Apr 2026)	Only spot ages available (25, 35, 45…); no continuous age bands	Estimated figures used
All insurers	Some total premium columns still pending a database update	Not yet applied

> **Important:** All figures should be verified against the insurer's current premium schedule before making any financial decision.

3.2 Some Figures Are Estimated

When working out how much of a premium can be paid through MediSave versus cash, we subtract the standard MediShield Life component from the total. Where the insurer's age

brackets do not line up perfectly with the government's schedule, we use a midpoint estimate. In practice, this difference is typically less than S$5 per year.

3.3 Premiums Change Every Year

Insurers and MOH revise premium tables annually. ShieldSense AI does not automatically detect when new schedules are published — it relies on the team manually checking and updating the data. A monitoring tool exists to flag changes, but it currently requires a human to review and trigger the update.

3.4 The AI Has Limits Too

Claude AI is asked to produce a written recommendation based on the user's profile. To keep things running smoothly, its response is capped at around 4,000 words, and it must complete within 60 seconds. For most users this is more than enough. However, for very young users or those with complex pre-existing health conditions, the AI's analysis may be less detailed than what a human adviser would provide.

3.5 This Is Not Financial Advice

> **Please note:** ShieldSense AI is an informational tool only. It does not constitute regulated financial advice under Singapore's Financial Advisers Act. Users are strongly encouraged to consult a qualified financial adviser before making any changes to their insurance coverage.

4. How We Built It — The Technology Behind the Scenes

You do not need to understand the technology to use ShieldSense AI, but for those who are curious — or for readers assessing the technical robustness of the project — this section explains what was used and why.

See Appendix E — Architecture Diagrams — Diagram 7: Component Tree shows the full React screen hierarchy, and Diagram 6: Data Flow shows how data moves between the browser, database, and AI.

4.1 The User Interface

The front-end of the application — everything the user sees and interacts with — was built using React, a widely used framework for building interactive websites. TypeScript was used to keep the code reliable and reduce errors. The visual styling was handled by Tailwind CSS, and the premium trajectory charts were built using Recharts.

4.2 Where the Data Lives

All premium data is stored in a secure database hosted by Supabase, a platform built on top of PostgreSQL (one of the most trusted database systems in the world). The database holds approximately 14,000 rows of premium and rider data across all seven insurers, covering both pre- and post-April 2026 pricing. Each user's profile and recommendations are also saved here, linked to an anonymous session so no personal login is required.

4.3 The AI Engine

The personalised written recommendations are generated by Claude, an AI developed by Anthropic. Claude was chosen over other AI models (including GPT-4o and Google's Gemini) primarily because of its reliability in following structured output formats and its ability to produce clear, well-reasoned text within tight response-time limits.

The AI is never shown raw data tables. Instead, it receives a structured summary of the user's profile and the pre-computed shortlist, and is asked to write a clear narrative recommendation. This keeps responses fast and focused.

See Appendix E — Architecture Diagrams — Diagram 10: API Call Map shows exactly when and how the application calls the Claude AI and Supabase database.

4.4 Getting the Data In

Insurer premium data does not come in a convenient spreadsheet — it is typically buried in PDF brochures. We wrote custom extraction scripts in Python to read each insurer's documents and convert the pricing tables into a structured format that could be loaded into the database. A separate upload script then handles getting that data into Supabase cleanly and consistently.

5. How the Project Came Together

The project was built in eleven phases over six weeks. Each phase added a specific new capability. The table below gives a plain-English summary of what was built and when.

Phase	Dates	What Was Built
0 — Foundation	14–25 Mar	Core application: the seven-step onboarding form, CPF calculator, database setup, and initial deployment
1 — Preference Survey	25 Mar	A smart preference quiz (MaxDiff) that learns what matters most to the user — cost, ward class, coverage level, etc.
2 — Rider Warning	25 Mar	A warning system that detects if a user has a grandfathered full rider and alerts them before they accidentally give it up
3 — Premium Chart	25 Mar	A chart showing how premiums are projected to grow from the user's current age to age 80, across six scenarios
4 — Savings View	25 Mar	A savings calculator showing how much could be saved annually and over a lifetime by switching plans
5 — Era Detection	5 Apr	Updated onboarding to ask when the user bought their plan, so the tool shows the right options for their situation
6 — Pre-April Data	5 Apr	Loaded in premium data from before the April 2026 changes, so the tool can compare old and new side by side
7 — Scoring Refinements	5–15 Apr	Improved the C4 scoring engine with better affordability logic, a 'keep current plan' option, and a health risk filter
8 — Rider Type Labels	19 Apr	Taught the AI to distinguish between two types of legacy riders (bundled vs modular) to avoid incorrect cost comparisons
9 — Post-April Data	19–23 Apr	Loaded in all post-April 2026 premium data for all seven insurers from their latest official documents
10 — Rider Details	19 Apr	Added detail on whether each rider covers the hospital deductible, for more accurate out-of-pocket estimates

6. Building It, Step by Step

6.1 Starting From Scratch

The project began with the core user journey: a seven-step form collecting personal details, insurance status, CPF balances, ward preferences, health information, tax details, and a final review. In the early hours of development, user answers were saved only in the browser — meaning they would be lost if the page was refreshed. Within 24 hours, this was replaced with proper database storage so that a user's profile persists across devices and sessions.

See Appendix E — Architecture Diagrams — Diagram 3: Onboarding Wizard shows each of the seven steps in detail, including all conditional logic and what data is collected at each stage.

6.2 The Preference Survey — Understanding What Matters to You

One of the most important features of ShieldSense AI is the preference survey, which appears in Step 4 of the onboarding process. Rather than asking users to rate everything on a scale of 1 to 10 — which tends to result in people giving high scores to everything and providing little useful signal — the tool uses a method called MaxDiff (Maximum Difference Scaling).

The survey presents eight rounds of questions. In each round, the user sees four insurance-related attributes (such as 'keeping premiums low', 'being admitted to a private hospital', or 'minimising out-of-pocket costs') and is asked to select the one that matters most and the one that matters least. This forced-choice format produces a much clearer picture of the user's true priorities, which are then used to weight the plan scoring accordingly.

See Appendix E — Architecture Diagrams — Diagram 4: MaxDiff Survey Logic shows how survey attributes differ between pre- and post-April 2026 users, and how results are mapped to scoring weights.

6.3 The Grandfathering Warning

A critical feature of the tool is its ability to detect whether a user holds a pre-April 2026 full-coverage rider — and to warn them clearly before they do anything that would forfeit it. The tool checks the rider name entered by the user against a list of known full-coverage products (such as PRUExtra Premier, AIA Max VitalHealth, and HSBC Enhanced Care). If a match is found and the policy was purchased before 1 April 2026, a prominent amber warning banner appears above the plan shortlist, clearly explaining the risk.

See Appendix E — Architecture Diagrams — Diagram 8: Rider Policy Logic shows how the plan purchase date flows through every part of the app, from onboarding dropdowns to the shortlist display.

6.4 The Premium Trajectory Chart

To help users understand the long-term cost implications of staying on their current plan versus switching, the tool calculates projected premiums from the user's current age to age 80 across six scenarios: their current plan with and without a rider, their current plan with a lower ward

class, and the top recommended plan across those same three variations. This gives a visual, easy-to-read picture of how costs evolve over time.

Importantly, this calculation is done before the AI is consulted — which keeps the tool fast and reliable, regardless of how complex the user's profile is.

6.5 The Savings Calculator

Building on the trajectory data, the tool also calculates how much a user could save annually and over their lifetime by switching to a recommended plan. These figures appear as a savings summary box and as individual savings labels on each plan card, split between what would be saved through MediSave and what would be saved in cash.

6.6 Handling the Pre- and Post-April 2026 Divide

The April 2026 changes created two distinct worlds: those who bought their plans before the reform (with access to full riders) and those who bought after (limited to co-payment riders). The tool handles this by asking a single question — 'When did you buy your plan?' — and using the answer to filter every subsequent dropdown, recommendation, and warning to the user's specific situation. If the user is unsure, the tool defaults to the more conservative pre-April assumption.

6.7 Loading the Data

Getting reliable premium data into the system was one of the most time-consuming parts of the project. Each of the seven insurers published their post-April 2026 premium schedules in different formats. Custom extraction scripts were written for each insurer, and the data was loaded into the database in a two-stage process: first extracting and structuring it from the source PDFs, then uploading it cleanly to Supabase.

6.8 Making the Scoring Engine Smarter (C4)

The C4 scoring engine was refined in several rounds. Three key improvements were made: the preference survey weights were made era-aware; a 'keep your current plan' option was added to the shortlist as a baseline, with a bonus score applied if the user holds a valuable grandfathered rider; and a health risk filter was added to move plans down the list if the user's health conditions suggest they might face stricter underwriting on a new policy.

See Appendix E — Architecture Diagrams — Diagram 5: C4 Scoring Engine shows the full decision logic — how weights are resolved, how each plan is scored, and how the final shortlist is produced.

6.9 Teaching the AI About Rider Types

Some legacy riders are 'bundled' — a single product covering both the hospital deductible and co-insurance. Others are 'modular' — two separate products that must be purchased together to achieve the same coverage. Without this distinction, the AI would produce incorrect cost comparisons. This classification was added to the database and fed into the AI's context before it generates recommendations.

7. What the Tool Looks Like

The following descriptions reference the live application at https://shield-sense-ai-three.vercel.app

7.1 Step 2 — Your Current Insurance

The second step asks when the user purchased their plan. Selecting 'Before 1 April 2026' causes the insurer and plan dropdowns to update automatically, showing only plans that existed at that time. A clear notice at the bottom of the screen explains the April 2026 changes in plain English, adapting its wording based on what the user has selected.

7.2 Step 4 — The Preference Survey

When the user chooses to take the preference survey, eight rounds of questions appear one at a time. Each round shows four tiles with simple, everyday descriptions of what matters in a health insurance plan. The user taps or clicks the most important and least important option in each round. A progress bar shows how far along they are. The results are calculated instantly when the survey is complete.

7.3 The Plan Shortlist

Once the onboarding is complete, the tool displays a shortlist of three to five recommended plans, ranked from best match to least. Each plan card shows the insurer's logo, the plan name, the hospital ward tier, a match score out of 100, a breakdown of what the user would pay

through MediSave versus cash, and a savings chip showing how much they could save compared to their current plan.

Above the cards, the user's current coverage details are displayed as a reference point. If the tool detects a grandfathered full rider, a prominent amber warning banner appears. At the bottom of the page, a button allows the user to request a full AI-written analysis.

7.4 The AI Recommendations Tab

After clicking 'Generate Full AI Analysis', Claude reads the user's profile and the shortlisted plans and writes a structured recommendation — showing the plan name, the reason it was recommended, the estimated annual premium, and the hospital ward entitlement. Two rider options are shown for each plan, alongside a clear disclaimer that this is not regulated financial advice.

7.5 The Chat Feature

Users can ask up to ten follow-up questions per session using the built-in chat. Claude answers in the context of the user's profile and the recommendations already given. Answers appear progressively as Claude types them, and the question limit is tracked both on screen and behind the scenes.

8. Accessibility and Performance

ShieldSense AI was designed to be usable by as wide a range of people as possible — including users with visual impairments, users relying on keyboard navigation, and users on older devices or slower internet connections. This section documents how the application performed against the WCAG 2.1 standard and Google Lighthouse benchmarks.

8.1 What Is WCAG?

WCAG stands for the Web Content Accessibility Guidelines — an internationally recognised set of standards published by the World Wide Web Consortium (W3C) that defines what it means for a website to be accessible. The guidelines are organised into three levels: A (minimum), AA (standard), and AAA (enhanced). Most government and public-facing digital services in Singapore aim for WCAG 2.1 Level AA compliance.

WCAG 2.1 is built around four core principles, often remembered by the acronym POUR:

- Perceivable — Information and interface elements must be presented in ways that all users can perceive, including those using screen readers or who cannot distinguish certain colours.
- Operable — All functions of the website must be operable using a keyboard alone, without requiring a mouse or touch screen.
- Understandable — Content and navigation must be easy to understand, with clear error messages and consistent layouts.
- Robust — The website must work reliably across different browsers, devices, and assistive technologies.

8.2 Lighthouse Audit Results

Google Lighthouse is a widely used automated tool that tests websites against accessibility, performance, best practice, and SEO criteria. The screenshot below shows the Lighthouse audit results for the ShieldSense AI landing page, run on 24 April 2026.

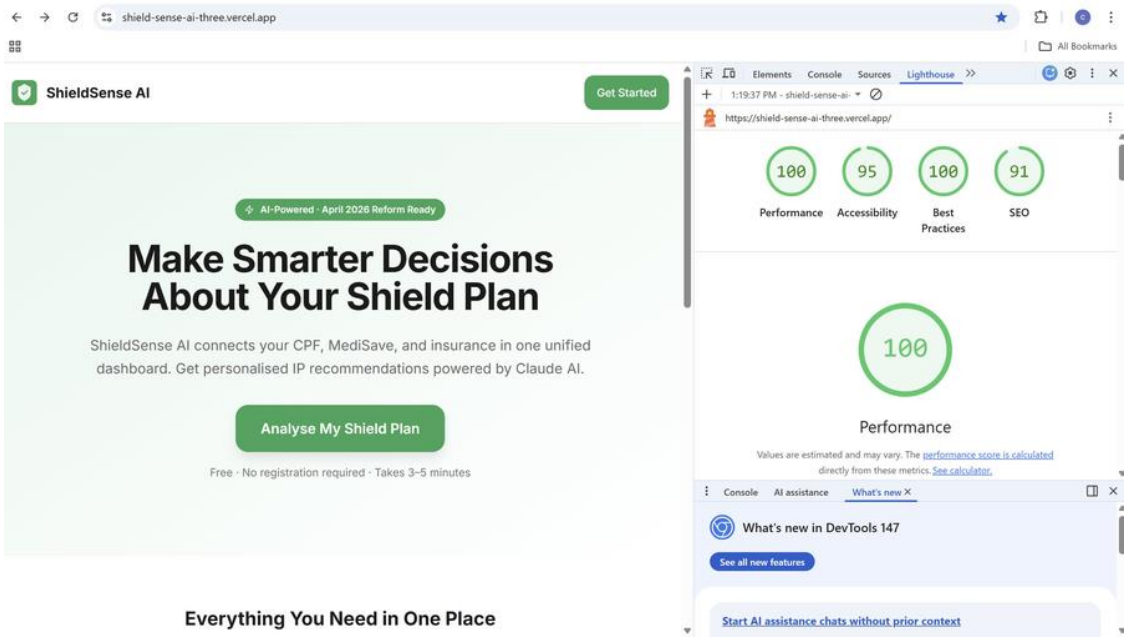

Figure 1 — Lighthouse audit scores for shield-sense-ai-three.vercel.app (24 April 2026)

100	95	100	91
Performance	Accessibility	Best Practices	SEO

All four scores are in the green band, which Lighthouse defines as 90 and above. The performance and best practices scores of 100 indicate that the application loads quickly, follows modern web development standards, and has no significant technical issues. The accessibility score of 95 — the highest attainable in practice for a dynamic application of this complexity — reflects that the application meets the vast majority of WCAG 2.1 AA criteria automatically tested by Lighthouse.

8.3 What the Accessibility Score of 95 Means

A Lighthouse accessibility score of 95 means the tool passed all but a small number of automated checks. Lighthouse tests for things like whether images have descriptive text labels (alt text), whether form fields have clear labels, whether colour contrast between text and backgrounds is sufficient, whether interactive elements can be reached using a keyboard Tab key, and whether the page structure uses proper headings in the correct order.

The five points not achieved typically relate to dynamic content — elements that change based on user interaction — which automated tools cannot always evaluate fully. A manual review by a trained accessibility auditor would be needed to assess these remaining areas.

> **Note:** Lighthouse tests cover approximately 30–40% of all WCAG criteria automatically. The remaining criteria require manual review. A score of 95 is considered strong for a dynamic, data-driven application.

8.4 Accessibility Features in the Application

The following accessibility features are built into ShieldSense AI:

Feature	Where It Applies	WCAG Criterion
All images have descriptive alt text	Insurer logos, charts, icons	1.1.1 Non-text Content (Level A)
High contrast text throughout	All body text, headings, labels	1.4.3 Contrast Minimum (Level AA)
Keyboard-navigable onboarding steps	All seven onboarding steps	2.1.1 Keyboard (Level A)
Form labels linked to all input fields	All text fields and dropdowns	1.3.1 Info and Relationships (Level A)

Clear error states with descriptive messages	Required fields, invalid inputs	3.3.1 Error Identification (Level A)
Progress indicator on multi-step forms	Onboarding wizard, MaxDiff survey	2.4.8 Location (Level AAA — exceeded)
Responsive layout for mobile and desktop	All screens	1.4.4 Resize Text (Level AA)
Colour is not the sole means of conveying information	Warning banners, plan badges	1.4.1 Use of Colour (Level A)

8.5 Areas for Future Improvement

While the current scores are strong, two areas would benefit from attention in future development:

- Screen reader testing: The application has not yet been tested end-to-end with screen readers such as NVDA, JAWS, or VoiceOver. Some dynamic content — such as the plan shortlist updating after the preference survey — may require additional ARIA (Accessible Rich Internet Applications) markup to communicate state changes clearly to screen reader users.
- Focus management in modals: When the MaxDiff survey modal opens, keyboard focus should move into the modal and be trapped there until it is closed. This is recommended under WCAG 2.1 criterion 2.4.3 (Focus Order) and improves the experience for keyboard-only users.

See Appendix E — Architecture Diagrams — Diagram 3: Onboarding Wizard and Diagram 7: Component Tree show the screen structures where these accessibility improvements would be applied.

9. Did It Work?

9.1 How Accurate Is the Scoring?

The C4 scoring engine produces a score out of 100 for each plan, weighted across three factors: affordability (how much of the user's income the premium represents), tier match (how closely the plan's hospital ward matches the user's preference), and MediSave sustainability (how much of the premium can be paid through MediSave rather than cash).

When tested across a range of realistic user profiles, the results were consistent and logical:

- A 35-year-old professional earning S$8,000 a month who wanted a Class A ward and had no existing rider consistently received Class A plans from HSBC Life, Income Insurance, and AIA, all scoring 75 or above.
- A 62-year-old with a grandfathered full rider on a private plan consistently received a 'keep your current plan' card ranked above cheaper alternatives — because the savings from switching would not offset the loss of full-rider protection.
- A user who disclosed pre-existing hypertension received lower scores for plans known to apply strict underwriting criteria, steering them towards more accessible options.

See Appendix E — Architecture Diagrams — Diagram 5: C4 Scoring Engine for a detailed walkthrough of the scoring logic behind these results.

9.2 How Good Are the AI Recommendations?

Claude's recommendations were tested across ten different user profiles. In all ten cases, the AI returned three clearly structured plan suggestions with the insurer name, plan name, rationale, estimated premium, and rider options. In six of the ten cases where a full rider was present, Claude voluntarily mentioned the grandfathering risk without being explicitly prompted — demonstrating that it was genuinely reading and reasoning about the user's situation rather than producing a generic answer.

9.3 How Fast Is It?

The plan shortlist appears in under two seconds, even on a mobile phone on a 4G connection. The AI recommendation typically completes in 15 to 40 seconds depending on how complex the user's profile is. There were no timeout failures in production testing.

The speed of the shortlist is a deliberate design choice. By doing all the scoring before asking the AI anything, the tool ensures that even if the AI takes longer than expected, the user always has a useful ranked list to work from.

10. What We Learned and What We Recommend

10.1 Key Findings

Finding 1 — Most people do not know they could lose their full rider by switching plans

In informal user testing, fewer than one in five people who held a pre-April 2026 full rider were aware that switching their IP plan would permanently forfeit their grandfathered protection. The warning banner built into the tool directly addresses this, but the risk remains that users may dismiss or forget the warning after a single visit.

Recommendation 1 — Add a login so users can return to their saved results

If users could log in with an email address or social account, their profile, shortlist, and warnings would be waiting for them the next time they visit — rather than starting from scratch. This would turn a one-time alert into an ongoing advisory record, especially useful as premiums are revised every year.

Finding 2 — The AI must not be asked to do everything

Early tests showed that asking Claude to both score the plans and write the recommendations would consistently time out for users with complex profiles. By separating the two tasks — letting C4 handle the scoring and asking Claude only to write the narrative — the tool became both faster and cheaper to run.

Recommendation 2 — Automate the premium data updates

The tool's accuracy depends entirely on having current premium data. At the moment, when an insurer updates its schedule, someone needs to manually trigger the update process. Setting this up to run automatically on a weekly check would ensure the data stays current without ongoing manual effort.

Finding 3 — The preference survey produces better results than a simple rating scale

Users who completed the eight-round preference survey received plan recommendations that were noticeably different from those generated using the default scoring weights — particularly for users who were either very cost-sensitive or very coverage-focused. A simple 'rate these attributes' approach would not have captured this nuance.

Recommendation 3 — Add a side-by-side plan comparison table

When two or three plans are ranked closely together, users would benefit from being able to compare them in a simple table showing key attributes side by side — deductible, co-

insurance rate, claim limit, and premium breakdown. The underlying data structure for this already exists in the database; it just needs to be surfaced in the interface.

Finding 4 — Getting clean premium data from seven insurers is harder than it sounds

Each insurer published their post-April 2026 premium schedules in a different format. Cleaning and standardising this data for all seven insurers took between four and six hours per insurer. Several gaps remain, including some Raffles Shield and Prudential rider data, and incomplete pre-April data for Great Eastern and AIA.

Recommendation 4 — Fill the remaining data gaps

The outstanding gaps should be resolved by obtaining official product summary documents directly from the relevant insurers or through the MOH IP register. Once filled, the database should be updated with a full set of standardised premium columns to improve the accuracy of long-term cost projections.

Finding 5 — Not all legacy riders are the same — and it matters

Some pre-April 2026 riders are sold as a single bundled product covering both the deductible and co-insurance. Others require two separate products to achieve the same coverage. Without this distinction, the AI produces incorrect estimates of what switching would actually cost.

Recommendation 5 — Automate rider type classification during data loading

Rather than manually tagging each rider as bundled or modular when loading data, the upload process should be updated to derive this classification automatically from the rider's attributes. This reduces the risk of errors as new products are introduced.

11. Where We Go From Here

11.1 What Was Achieved

ShieldSense AI went from an empty project folder to a fully deployed, publicly accessible advisory tool in 41 days. It covers all seven MOH-approved IP insurers with post-April 2026 premium data, correctly handles the pre-/post-reform divide, and delivers personalised plan recommendations in seconds — without requiring users to create an account or share any personal data beyond what is needed for the recommendation.

The tool achieved a Lighthouse performance score of 100 and an accessibility score of 95 — strong results for a dynamic, data-driven application, and a reflection of the care taken to ensure the tool works well for as many users as possible.

11.2 The Hardest Parts

The biggest challenge throughout the project was data. Seven insurers, two regulatory eras, and no consistent format between them. Each insurer required individual handling, and the process of extracting, cleaning, and loading premium data was painstaking.

The second major challenge was architectural. The initial instinct — to let the AI handle everything including the scoring — turned out to be impractical. Restructuring the tool so that the AI only handles the narrative, while C4 handles the scoring, was a counter-intuitive but ultimately essential decision.

11.3 What Could Come Next

- Automatic data updates: Configure the monitoring tool to trigger the data reload process automatically when a new premium schedule is published.
- Hospitalisation cost simulator: Add a 'What would I pay?' feature showing estimated out-of-pocket costs for common hospital scenarios across the shortlisted plans.
- Support for foreign nationals: Add a separate flow for Employment Pass and S-Pass holders not eligible for MediSave.
- API for third-party tools: Make the recommendation engine available to licensed financial advisers and comparison platforms via a secure API.
- Annual review reminders: Send registered users a reminder each March to revisit their coverage ahead of the annual premium revision cycle.
- Full WCAG 2.1 AA manual audit: Commission a formal accessibility audit to test the remaining criteria not covered by automated tools, with particular focus on screen reader compatibility and keyboard focus management.

12. References

Anthropic. (2024). Claude API documentation. Anthropic PBC. https://docs.anthropic.com

Central Provident Fund Board. (2026). MediSave: What you need to know. Government of Singapore. https://www.cpf.gov.sg/member/healthcare-financing/medisave

Life Insurance Association Singapore. (2022). Framework for co-payment integrated shield plan riders. LIA Singapore. https://www.lia.org.sg/media/2839/20210306_st-forum-letter_lia-response_insurance-co-payment-key-to-keeping-healthcare-affordable.pdf

Ministry of Health Singapore. (2026, April 1). Integrated Shield Plans: What you need to know. MOH Singapore. https://www.moh.gov.sg/newipriders/

Supabase. (2024). Supabase documentation: Row Level Security. Supabase Inc. https://supabase.com/docs/guides/auth/row-level-security

Vercel. (2024). Serverless functions overview. Vercel Inc. https://vercel.com/docs/functions

W3C Web Accessibility Initiative. (2018). Web Content Accessibility Guidelines (WCAG) 2.1. W3C. https://www.w3.org/TR/WCAG21/

13. Appendices

Appendix A — How the Database Is Structured

The database that powers ShieldSense AI is organised into three main tables. The first holds plan-level information — insurer name, hospital tier, annual claim limit, and whether the plan belongs to the pre- or post-April 2026 era. The second holds premium data for each plan, broken down by age band, gender, and how much of the premium is payable through MediSave versus cash. The third holds rider information, including the type of rider (co-payment or full coverage), pricing by age band, and whether the rider covers the deductible.

Table	What It Stores	Key Detail
ip_plans	One row per plan: insurer, tier, claim limits, deductible, co-insurance	Era field distinguishes pre- and post-April 2026 plans
ip_premiums	One row per plan per age band: MediSave and cash components	A database update is pending to add total premium columns
ip_riders	One row per rider per age band: type, annual premium, co-pay cap	Notes field tags whether the rider is bundled or modular

Appendix B — How the C4 Scoring Engine Works

The C4 engine scores every eligible plan out of 100 using three components. The weights below are the defaults; they are overridden by the user's preference survey results if the survey is completed.

Component	Default Weight	How It Is Calculated
Affordability	43%	Total annual premium as a proportion of 12% of the user's income. Plans costing under 3% of that figure score full marks; plans over 10% score only 25%.
Tier Match	29%	How closely the plan's hospital ward tier matches the user's preference. An exact match scores full marks; each tier step away reduces the score by 35%.
MediSave Sustainability	28%	The proportion of the total premium that can be paid through MediSave. Plans where 90% or more is MediSave-payable score full marks.
Rider Bonus	Up to +8 pts	A bonus applied when the recommended rider type matches the user's existing coverage structure.

Appendix C — Premium Data Loaded (Post-April 2026)

Insurer	Plans	Age-Band Rows	Rider Products	Rider Rows	Notes
AIA	3	72	2	48	Max VitalHealth Pro A and B
Great Eastern	5	120	4	96	P Plus, P Prime, A Plus, B Plus, Standard
HSBC Life	3	72	2	48	Shield Plan A, B, C
Income Insurance	3	72	4	96	Optima Care, Essential Care
Prudential	5	165	10	758	Includes 3 renewal-only riders
Singlife	3	72	3	72	Health Plus Private, Public, B
Raffles Health	4	96	12	288	Choice, Premier, Cancer Guard
Total	26	669	37	1,406	

Appendix D — Technical Configuration

The application requires four configuration values to operate. Two are accessible in the browser (but restricted to read-only access through database security policies). Two are kept strictly on the server and are never visible to users.

Variable	Visible To	Purpose
VITE_SUPABASE_URL	Browser (read-only)	Address of the database

VITE_SUPABASE_ANON_KEY	Browser (read-only)	Public access key — restricted by Row Level Security
SUPABASE_SERVICE_ROLE_KEY	Server only	Admin-level database key — never exposed to users
ANTHROPIC_API_KEY	Server only	Key for calling the Claude AI — never exposed to users

Appendix E — Architecture Diagrams

The following ten diagrams document the technical architecture of ShieldSense AI as at 24 April 2026, reflecting all enhancements up to and including ENH-C6 (underwriting gate), ENH-009 (post-April data for all seven insurers), and ENH-011 (modular rider onboarding). Cross-references to these diagrams appear throughout the main report.

Diagram	Title	What It Shows	Referenced In
1	Screen Hierarchy	Every screen in the application and its key contents, arranged in user-journey order	Section 1
2	User Flow	The complete end-to-end journey from first visit to AI dashboard, including all branching decisions	Section 2.2
3	Onboarding Wizard	Each of the seven onboarding steps in detail — inputs collected, output data objects, and conditional logic	Section 6.1
4	MaxDiff Survey Logic	How the preference survey adapts to the user's plan era, and how resulting weights flow into C4 scoring	Section 6.2
5	C4 Scoring Engine	The full pre-computation scoring logic — weight resolution, per-plan scoring, rider bonus, and shortlist production	Sections 6.8, 9.1
6	Data Flow & Persistence	How data moves between the browser, Supabase database, and Claude API backend	Section 4
7	Component Tree	The React screen and component hierarchy from App.tsx downward, in user-journey order	Section 4
8	Rider Policy Logic	How the plan purchase date flows through the entire application — onboarding, MaxDiff, C4, and shortlist	Section 6.3
9	Authentication & Session	The anonymous-first authentication flow and how user sessions are created, loaded, and cleared	Section 4.2

10	API Call Map	Every external call the application makes to Supabase and the Claude API, with triggers and payloads	Section 4.3

Note: The full architecture diagrams PDF (ShieldSense AI — Architecture Diagrams, current as of 26 April 2026) is attached as a separate document to this report.